



# CUSPARSE LIBRARY

DU-06709-001\_v10.2 | October 2019

# TABLE OF CONTENTS

<b>Chapter 1. Introduction.....</b>	<b>1</b>
1.1. Naming Conventions.....	1
1.2. Asynchronous Execution.....	2
1.3. Static Library support.....	2
<b>Chapter 2. Using the cuSPARSE API.....</b>	<b>4</b>
2.1. Thread Safety.....	4
2.2. Scalar Parameters.....	4
2.3. Parallelism with Streams.....	4
2.4. Compatibility and Versioning.....	5
<b>Chapter 3. cuSPARSE Indexing and Data Formats.....</b>	<b>6</b>
3.1. Index Base Format.....	6
3.1.1. Vector Formats.....	6
3.1.1.1. Dense Format.....	6
3.1.1.2. Sparse Format.....	6
3.2. Matrix Formats.....	7
3.2.1. Dense Format.....	7
3.2.2. Coordinate Format (COO).....	7
3.2.3. Compressed Sparse Row Format (CSR).....	8
3.2.4. Compressed Sparse Column Format (CSC).....	9
3.2.5. Ellpack-Itpack Format (ELL).....	10
3.2.6. Hybrid Format (HYB) [DEPRECATED].....	11
3.2.7. Block Compressed Sparse Row Format (BSR).....	11
3.2.8. Extended BSR Format (BSRX).....	13
<b>Chapter 4. cuSPARSE Types Reference.....</b>	<b>15</b>
4.1. Data types.....	15
4.2. cusparseAction_t.....	15
4.3. cusparseDirection_t.....	15
4.4. cusparseHandle_t.....	16
4.5. cusparseHybMat_t [DEPRECATED].....	16
4.5.1. cusparseHybPartition_t [DEPRECATED].....	16
4.6. cusparseMatDescr_t.....	17
4.6.1. cusparseDiagType_t.....	17
4.6.2. cusparseFillMode_t.....	17
4.6.3. cusparseIndexBase_t.....	17
4.6.4. cusparseMatrixType_t.....	17
4.7. cusparseOperation_t.....	18
4.8. cusparsePointerMode_t.....	18
4.9. cusparseAlgMode_t.....	19
4.10. cusparseSolvePolicy_t.....	19
4.11. cusparseSolveAnalysisInfo_t [DEPRECATED].....	19

4.12. <code>cusparseColorInfo_t</code> .....	19
4.13. <code>csrv2Info_t</code> .....	19
4.14. <code>csrsm2Info_t</code> .....	20
4.15. <code>csric02Info_t</code> .....	20
4.16. <code>csrilu02Info_t</code> .....	20
4.17. <code>bsrv2Info_t</code> .....	20
4.18. <code>bsrsm2Info_t</code> .....	20
4.19. <code>bsric02Info_t</code> .....	20
4.20. <code>bsrilu02Info_t</code> .....	20
4.21. <code>csrgemm2Info_t</code> .....	20
4.22. <code>cusparseStatus_t</code> .....	21
<b>Chapter 5. cuSPARSE Management Function Reference.....</b>	<b>23</b>
5.1. <code>cusparseCreate()</code> .....	23
5.2. <code>cusparseDestroy()</code> .....	23
5.3. <code>cusparseGetErrorName()</code> .....	24
5.4. <code>cusparseGetErrorString()</code> .....	24
5.5. <code>cusparseGetProperty()</code> .....	24
5.6. <code>cusparseGetVersion()</code> .....	25
5.7. <code>cusparseGetPointerMode()</code> .....	25
5.8. <code>cusparseSetPointerMode()</code> .....	25
5.9. <code>cusparseGetStream()</code> .....	26
5.10. <code>cusparseSetStream()</code> .....	26
<b>Chapter 6. cuSPARSE Helper Function Reference.....</b>	<b>27</b>
6.1. <code>cusparseCreateColorInfo()</code> .....	27
6.2. <code>cusparseCreateHybMat()</code> [DEPRECATED].....	27
6.3. <code>cusparseCreateMatDescr()</code> .....	27
6.4. <code>cusparseCreateSolveAnalysisInfo()</code> [DEPRECATED].....	28
6.5. <code>cusparseDestroyColorInfo()</code> .....	28
6.6. <code>cusparseDestroyHybMat()</code> [DEPRECATED].....	28
6.7. <code>cusparseDestroyMatDescr()</code> .....	29
6.8. <code>cusparseDestroySolveAnalysisInfo()</code> [DEPRECATED].....	29
6.9. <code>cusparseGetLevelInfo()</code> .....	29
6.10. <code>cusparseGetMatDiagType()</code> .....	30
6.11. <code>cusparseGetMatFillMode()</code> .....	30
6.12. <code>cusparseGetMatIndexBase()</code> .....	30
6.13. <code>cusparseGetMatType()</code> .....	31
6.14. <code>cusparseSetMatDiagType()</code> .....	31
6.15. <code>cusparseSetMatFillMode()</code> .....	31
6.16. <code>cusparseSetMatIndexBase()</code> .....	32
6.17. <code>cusparseSetMatType()</code> .....	32
6.18. <code>cusparseCreateCsrsv2Info()</code> .....	32
6.19. <code>cusparseDestroyCsrsv2Info()</code> .....	33
6.20. <code>cusparseCreateCsrsm2Info()</code> .....	33

6.21. <code>cusparseDestroyCsrsm2Info()</code> .....	33
6.22. <code>cusparseCreateCsric02Info()</code> .....	33
6.23. <code>cusparseDestroyCsric02Info()</code> .....	34
6.24. <code>cusparseCreateCsrlu02Info()</code> .....	34
6.25. <code>cusparseDestroyCsrlu02Info()</code> .....	34
6.26. <code>cusparseCreateBsrsv2Info()</code> .....	35
6.27. <code>cusparseDestroyBsrsv2Info()</code> .....	35
6.28. <code>cusparseCreateBsrm2Info()</code> .....	35
6.29. <code>cusparseDestroyBsrm2Info()</code> .....	36
6.30. <code>cusparseCreateBsric02Info()</code> .....	36
6.31. <code>cusparseDestroyBsric02Info()</code> .....	36
6.32. <code>cusparseCreateBsrlu02Info()</code> .....	36
6.33. <code>cusparseDestroyBsrlu02Info()</code> .....	37
6.34. <code>cusparseCreateCsrgemm2Info()</code> .....	37
6.35. <code>cusparseDestroyCsrgemm2Info()</code> .....	37
6.36. <code>cusparseCreatePruneInfo()</code> .....	38
6.37. <code>cusparseDestroyPruneInfo()</code> .....	38
<b>Chapter 7. cuSPARSE Level 1 Function Reference.....</b>	<b>39</b>
7.1. <code>cusparse&lt;t&gt;axpyi()</code> .....	39
7.2. <code>cusparse&lt;t&gt;doti() [DEPRECATED]</code> .....	41
7.3. <code>cusparse&lt;t&gt;dotci() [DEPRECATED]</code> .....	42
7.4. <code>cusparse&lt;t&gt;gthr()</code> .....	44
7.5. <code>cusparse&lt;t&gt;gthrz()</code> .....	45
7.6. <code>cusparse&lt;t&gt;roti()</code> .....	46
7.7. <code>cusparse&lt;t&gt;sctr()</code> .....	48
<b>Chapter 8. cuSPARSE Level 2 Function Reference.....</b>	<b>50</b>
8.1. <code>cusparse&lt;t&gt;bsrmv()</code> .....	51
8.2. <code>cusparse&lt;t&gt;bsrxmv()</code> .....	54
8.3. <code>cusparse&lt;t&gt;bsrsv2_bufferSize()</code> .....	58
8.4. <code>cusparse&lt;t&gt;bsrsv2_analysis()</code> .....	61
8.5. <code>cusparse&lt;t&gt;bsrsv2_solve()</code> .....	64
8.6. <code>cusparseXbsrsv2_zeroPivot()</code> .....	68
8.7. <code>cusparse&lt;t&gt;csrmv() [DEPRECATED]</code> .....	70
8.8. <code>cusparse&lt;t&gt;csrmv_mp() [DEPRECATED]</code> .....	73
8.9. <code>cusparseCsrmvEx()</code> .....	76
8.10. <code>cusparse&lt;t&gt;csrv_analysis() [DEPRECATED]</code> .....	78
8.11. <code>cusparseCsrsv_analysisEx() [DEPRECATED]</code> .....	80
8.12. <code>cusparse&lt;t&gt;csrsv_solve() [DEPRECATED]</code> .....	81
8.13. <code>cusparseCsrsv_solveEx() [DEPRECATED]</code> .....	83
8.14. <code>cusparse&lt;t&gt;csrv2_bufferSize()</code> .....	84
8.15. <code>cusparse&lt;t&gt;csrv2_analysis()</code> .....	86
8.16. <code>cusparse&lt;t&gt;csrv2_solve()</code> .....	89
8.17. <code>cusparseXcsrv2_zeroPivot()</code> .....	93

8.18. <code>cusparse&lt;t&gt;gemvi()</code> .....	94
8.19. <code>cusparse&lt;t&gt;hybmv()</code> [DEPRECATED].....	96
8.20. <code>cusparse&lt;t&gt;hybsv_analysis()</code> [DEPRECATED].....	98
8.21. <code>cusparse&lt;t&gt;hybsv_solve()</code> [DEPRECATED].....	100
<b>Chapter 9. cuSPARSE Level 3 Function Reference.....</b>	<b>102</b>
9.1. <code>cusparse&lt;t&gt;bsrmm()</code> .....	103
9.2. <code>cusparse&lt;t&gt;bsrsm2_bufferSize()</code> .....	107
9.3. <code>cusparse&lt;t&gt;bsrsm2_analysis()</code> .....	110
9.4. <code>cusparse&lt;t&gt;bsrsm2_solve()</code> .....	113
9.5. <code>cusparseXbsrsm2_zeroPivot()</code> .....	116
9.6. <code>cusparse&lt;t&gt;csrmm()</code> [DEPRECATED].....	117
9.7. <code>cusparse&lt;t&gt;csrmm2()</code> [DEPRECATED].....	120
9.8. <code>cusparse&lt;t&gt;csrsm_analysis()</code> [DEPRECATED].....	124
9.9. <code>cusparse&lt;t&gt;csrsm_solve()</code> [DEPRECATED].....	126
9.10. <code>cusparse&lt;t&gt;csrsm2_bufferSizeExt()</code> .....	129
9.11. <code>cusparse&lt;t&gt;csrsm2_analysis()</code> .....	132
9.12. <code>cusparse&lt;t&gt;csrsm2_solve()</code> .....	135
9.13. <code>cusparseXcsrsm2_zeroPivot()</code> .....	138
9.14. <code>cusparse&lt;t&gt;gemmi()</code> .....	139
<b>Chapter 10. cuSPARSE Extra Function Reference.....</b>	<b>141</b>
10.1. <code>cusparse&lt;t&gt;csrgeam()</code> [DEPRECATED].....	142
10.2. <code>cusparse&lt;t&gt;csrgeam2()</code> .....	146
10.3. <code>cusparse&lt;t&gt;csrgemm()</code> [DEPRECATED].....	150
10.4. <code>cusparse&lt;t&gt;csrgemm2()</code> .....	154
<b>Chapter 11. cuSPARSE Preconditioners Reference.....</b>	<b>159</b>
11.1. Incomplete Cholesky Factorization: level 0.....	159
11.1.1. <code>cusparse&lt;t&gt;csric0()</code> [DEPRECATED].....	160
11.1.2. <code>cusparse&lt;t&gt;csric02_bufferSize()</code> .....	162
11.1.3. <code>cusparse&lt;t&gt;csric02_analysis()</code> .....	164
11.1.4. <code>cusparse&lt;t&gt;csric02()</code> .....	166
11.1.5. <code>cusparseXcsric02_zeroPivot()</code> .....	169
11.1.6. <code>cusparse&lt;t&gt;bsric02_bufferSize()</code> .....	171
11.1.7. <code>cusparse&lt;t&gt;bsric02_analysis()</code> .....	173
11.1.8. <code>cusparse&lt;t&gt;bsric02()</code> .....	176
11.1.9. <code>cusparseXbsric02_zeroPivot()</code> .....	180
11.2. Incomplete LU Factorization: level 0.....	180
11.2.1. <code>cusparse&lt;t&gt;csrilu0()</code> [DEPRECATED].....	181
11.2.2. <code>cusparseCsrilu0Ex()</code> [DEPRECATED].....	183
11.2.3. <code>cusparse&lt;t&gt;csrilu02_numericBoost()</code> .....	184
11.2.4. <code>cusparse&lt;t&gt;csrilu02_bufferSize()</code> .....	185
11.2.5. <code>cusparse&lt;t&gt;csrilu02_analysis()</code> .....	187
11.2.6. <code>cusparse&lt;t&gt;csrilu02()</code> .....	189
11.2.7. <code>cusparseXcsrilu02_zeroPivot()</code> .....	192

11.2.8. <code>cusparse&lt;t&gt;bsrilu02_numericBoost()</code> .....	193
11.2.9. <code>cusparse&lt;t&gt;bsrilu02_bufferSize()</code> .....	195
11.2.10. <code>cusparse&lt;t&gt;bsrilu02_analysis()</code> .....	198
11.2.11. <code>cusparse&lt;t&gt;bsrilu02()</code> .....	201
11.2.12. <code>cusparseXbsrilu02_zeroPivot()</code> .....	205
11.3. Tridiagonal Solve.....	205
11.3.1. <code>cusparse&lt;t&gt;gtsv()</code> [DEPRECATED].....	206
11.3.2. <code>cusparse&lt;t&gt;gtsv_nopivot()</code> [DEPRECATED].....	208
11.3.3. <code>cusparse&lt;t&gt;gtsv2_buffSizeExt()</code> .....	210
11.3.4. <code>cusparse&lt;t&gt;gtsv2()</code> .....	212
11.3.5. <code>cusparse&lt;t&gt;gtsv2_nopivot_bufferSizeExt()</code> .....	214
11.3.6. <code>cusparse&lt;t&gt;gtsv2_nopivot()</code> .....	216
11.4. Batched Tridiagonal Solve.....	217
11.4.1. <code>cusparse&lt;t&gt;gtsvStridedBatch()</code> [DEPRECATED].....	218
11.4.2. <code>cusparse&lt;t&gt;gtsv2StridedBatch_bufferSizeExt()</code> .....	220
11.4.3. <code>cusparse&lt;t&gt;gtsv2StridedBatch()</code> .....	222
11.4.4. <code>cusparse&lt;t&gt;gtsvInterleavedBatch()</code> .....	224
11.5. Batched Pentadiagonal Solve.....	226
11.5.1. <code>cusparse&lt;t&gt;gpsvInterleavedBatch()</code> .....	227
<b>Chapter 12. cuSPARSE Reorderings Reference.....</b>	<b>230</b>
12.1. <code>cusparse&lt;t&gt;csrcolor()</code> .....	231
<b>Chapter 13. cuSPARSE Format Conversion Reference.....</b>	<b>234</b>
13.1. <code>cusparse&lt;t&gt;bsr2csr()</code> .....	235
13.2. <code>cusparse&lt;t&gt;gebsr2gebsc()</code> .....	238
13.3. <code>cusparse&lt;t&gt;gebsr2gebsr()</code> .....	241
13.4. <code>cusparse&lt;t&gt;gebsr2csr()</code> .....	245
13.5. <code>cusparse&lt;t&gt;csr2gebsr()</code> .....	248
13.6. <code>cusparse&lt;t&gt;coo2csr()</code> .....	251
13.7. <code>cusparse&lt;t&gt;csc2dense()</code> .....	253
13.8. <code>cusparse&lt;t&gt;csc2hyb()</code> [DEPRECATED].....	255
13.9. <code>cusparse&lt;t&gt;csr2bsr()</code> .....	257
13.10. <code>cusparse&lt;t&gt;csr2coo()</code> .....	259
13.11. <code>cusparse&lt;t&gt;csr2csc()</code> [DEPRECATED].....	261
13.12. <code>cusparseCsr2cscEx()</code> [DEPRECATED].....	263
13.13. <code>cusparseCsr2cscEx2()</code> .....	264
13.14. <code>cusparse&lt;t&gt;csr2dense()</code> .....	266
13.15. <code>cusparse&lt;t&gt;csr2csr_compress()</code> .....	268
13.16. <code>cusparse&lt;t&gt;csr2hyb()</code> [DEPRECATED].....	272
13.17. <code>cusparse&lt;t&gt;dense2csc()</code> .....	274
13.18. <code>cusparse&lt;t&gt;dense2csr()</code> .....	276
13.19. <code>cusparse&lt;t&gt;dense2hyb()</code> [DEPRECATED].....	278
13.20. <code>cusparse&lt;t&gt;hyb2csc()</code> [DEPRECATED].....	280
13.21. <code>cusparse&lt;t&gt;hyb2csr()</code> [DEPRECATED].....	281

13.22. <code>cusparse&lt;t&gt;hyb2dense()</code> [DEPRECATED].....	282
13.23. <code>cusparse&lt;t&gt;nnz()</code> .....	284
13.24. <code>cusparseCreateIdentityPermutation()</code> .....	285
13.25. <code>cusparseXcoosort()</code> .....	286
13.26. <code>cusparseXcsrsort()</code> .....	288
13.27. <code>cusparseXcscsort()</code> .....	290
13.28. <code>cusparseXcsru2csr()</code> .....	293
13.29. <code>cusparseXpruneDense2csr()</code> .....	297
13.30. <code>cusparseXpruneCsr2csr()</code> .....	300
13.31. <code>cusparseXpruneDense2csrPercentage()</code> .....	303
13.32. <code>cusparseXpruneCsr2csrByPercentage()</code> .....	306
13.33. <code>cusparse&lt;t&gt;nnz_compress()</code> .....	309
<b>Chapter 14. cuSPARSE Generic API Reference.....</b>	<b>311</b>
14.1. Generic Types Reference.....	311
14.1.1. <code>cudaDataType_t</code> .....	311
14.1.2. <code>cusparseFormat_t</code> .....	312
14.1.3. <code>cusparseOrder_t</code> .....	312
14.1.4. <code>cusparseIndexType_t</code> .....	312
14.2. Sparse Vector APIs.....	312
14.2.1. <code>cusparseCreateSpVec()</code> .....	312
14.2.2. <code>cusparseDestroySpVec()</code> .....	313
14.2.3. <code>cusparseSpVecGet()</code> .....	313
14.2.4. <code>cusparseSpVecGetIndexBase()</code> .....	314
14.2.5. <code>cusparseSpVecGetValues()</code> .....	314
14.2.6. <code>cusparseSpVecSetValues()</code> .....	314
14.3. Sparse Matrix APIs.....	315
14.3.1. <code>cusparseCreateCoo()</code> .....	315
14.3.2. <code>cusparseCreateCooAoS()</code> .....	316
14.3.3. <code>cusparseCreateCsr()</code> .....	316
14.3.4. <code>cusparseDestroySpMat()</code> .....	317
14.3.5. <code>cusparseCooGet()</code> .....	317
14.3.6. <code>cusparseCooAoSGet()</code> .....	318
14.3.7. <code>cusparseCsrGet()</code> .....	319
14.3.8. <code>cusparseSpMatGetFormat()</code> .....	319
14.3.9. <code>cusparseSpMatGetIndexBase()</code> .....	320
14.3.10. <code>cusparseSpMatGetValues()</code> .....	320
14.3.11. <code>cusparseSpMatSetValues()</code> .....	320
14.3.12. <code>cusparseSpMatGetStridedBatch()</code> .....	320
14.3.13. <code>cusparseSpMatSetStridedBatch()</code> .....	321
14.4. Dense Vector APIs.....	321
14.4.1. <code>cusparseCreateDnVec()</code> .....	321
14.4.2. <code>cusparseDestroyDnVec()</code> .....	322
14.4.3. <code>cusparseDnVecGet()</code> .....	322

14.4.4. <code>cusparseDnVecGetValues()</code> .....	322
14.4.5. <code>cusparseDnVecSetValues()</code> .....	323
14.5. Dense Matrix APIs.....	323
14.5.1. <code>cusparseCreateDnMat()</code> .....	323
14.5.2. <code>cusparseDestroyDnMat()</code> .....	324
14.5.3. <code>cusparseDnMatGet()</code> .....	324
14.5.4. <code>cusparseDnMatGetValues()</code> .....	324
14.5.5. <code>cusparseDnSetValues()</code> .....	325
14.5.6. <code>cusparseDnMatGetStridedBatch()</code> .....	325
14.5.7. <code>cusparseDnMatSetStridedBatch()</code> .....	325
14.6. Generic API Functions.....	326
14.6.1. <code>cusparseSpVV()</code> .....	326
14.6.2. <code>cusparseSpMV()</code> .....	327
14.6.3. <code>cusparseSpMM()</code> .....	330
14.6.4. <code>cusparseConstrainedGeMM()</code> .....	332
14.7. Example of Generic API.....	334
14.7.1. CSR sparse matrix - vector multiplication.....	334
<b>Chapter 15. Appendix A: cuSPARSE Library C++ Example.....</b>	<b>336</b>
<b>Chapter 16. Appendix B: cuSPARSE Fortran Bindings.....</b>	<b>338</b>
16.1. Example B, Fortran Application.....	340
<b>Chapter 17. Appendix C: Examples of sorting.....</b>	<b>341</b>
17.1. COO Sort.....	341
<b>Chapter 18. Appendix D: Examples of prune.....</b>	<b>343</b>
18.1. Prune Dense to Sparse.....	343
18.2. Prune Sparse to Sparse.....	345
18.3. Prune Dense to Sparse by Percentage.....	347
18.4. Prune Sparse to Sparse by Percentage.....	349
<b>Chapter 19. Appendix E: Examples of gtsv.....</b>	<b>351</b>
19.1. Batched Tridiagonal Solver.....	351
<b>Chapter 20. Appendix F: Examples of gpsv.....</b>	<b>353</b>
20.1. Batched Penta-diagonal Solver.....	353
<b>Chapter 21. Appendix G: Examples of csrsm2.....</b>	<b>355</b>
21.1. Forward Triangular Solver.....	355
<b>Chapter 22. Appendix H: Acknowledgements.....</b>	<b>357</b>
<b>Chapter 23. Bibliography.....</b>	<b>358</b>

# Chapter 1. INTRODUCTION

The cuSPARSE library contains a set of basic linear algebra subroutines used for handling sparse matrices. It is implemented on top of the NVIDIA® CUDA™ runtime (which is part of the CUDA Toolkit) and is designed to be called from C and C++. The library routines can be classified into four categories:

- ▶ Level 1: operations between a vector in sparse format and a vector in dense format
- ▶ Level 2: operations between a matrix in sparse format and a vector in dense format
- ▶ Level 3: operations between a matrix in sparse format and a set of vectors in dense format (which can also usually be viewed as a dense tall matrix)
- ▶ Conversion: operations that allow conversion between different matrix formats, and compression of csr matrices.

The cuSPARSE library allows developers to access the computational resources of the NVIDIA graphics processing unit (GPU), although it does not auto-parallelize across multiple GPUs. The cuSPARSE API assumes that input and output data reside in GPU (device) memory, unless it is explicitly indicated otherwise by the string **DevHostPtr** in a function parameter's name (for example, the parameter **\*resultDevHostPtr** in the function **cusparse<t>doti()**).

It is the responsibility of the developer to allocate memory and to copy data between GPU memory and CPU memory using standard CUDA runtime API routines, such as **cudaMalloc()**, **cudaFree()**, **cudaMemcpy()**, and **cudaMemcpyAsync()**.

## 1.1. Naming Conventions

The cuSPARSE library functions are available for data types **float**, **double**, **cuComplex**, and **cuDoubleComplex**. The sparse Level 1, Level 2, and Level 3 functions follow this naming convention:

```
cusparse<t>[<matrix data format>]<operation>[<output matrix data format>]
```

where <t> can be **S**, **D**, **C**, **Z**, or **X**, corresponding to the data types **float**, **double**, **cuComplex**, **cuDoubleComplex**, and the generic type, respectively.

The `<matrix data format>` can be `dense`, `coo`, `csr`, `csc`, or `hyb`, corresponding to the dense, coordinate, compressed sparse row, compressed sparse column, and hybrid storage formats, respectively.

Finally, the `<operation>` can be `axpyi`, `doti`, `dotci`, `gthr`, `gthrz`, `roti`, or `sctr`, corresponding to the Level 1 functions; it also can be `mv` or `sv`, corresponding to the Level 2 functions, as well as `mm` or `sm`, corresponding to the Level 3 functions.

All of the functions have the return type `cusparseStatus_t` and are explained in more detail in the chapters that follow.

## 1.2. Asynchronous Execution

The cuSPARSE library functions are executed asynchronously with respect to the host and may return control to the application on the host before the result is ready. Developers can use the `cudaDeviceSynchronize()` function to ensure that the execution of a particular cuSPARSE library routine has completed.

A developer can also use the `cudaMemcpy()` routine to copy data from the device to the host and vice versa, using the `cudaMemcpyDeviceToHost` and `cudaMemcpyHostToDevice` parameters, respectively. In this case there is no need to add a call to `cudaDeviceSynchronize()` because the call to `cudaMemcpy()` with the above parameters is blocking and completes only when the results are ready on the host.

## 1.3. Static Library support

Starting with release 6.5, the cuSPARSE Library is also delivered in a static form as `libcusparse_static.a` on Linux and Mac OSes. The static cuSPARSE library and all others static maths libraries depend on a common thread abstraction layer library called `libculibos.a` on Linux and Mac and `culibos.lib` on Windows.

For example, on linux, to compile a small application using cuSPARSE against the dynamic library, the following command can be used:

```
nvcc myCusparseApp.c -lcusparse -o myCusparseApp
```

Whereas to compile against the static cuSPARSE library, the following command has to be used:

```
nvcc myCusparseApp.c -lcusparse_static -lculibos -o myCusparseApp
```

It is also possible to use the native Host C++ compiler. Depending on the Host Operating system, some additional libraries like `pthread` or `d1` might be needed on the linking line. The following command on Linux is suggested :

```
g++ myCusparseApp.c -lcusparse_static -lculibos -lcudart_static -lpthread -ldl -I <cuda-toolkit-path>/include -L <cuda-toolkit-path>/lib64 -o myCusparseApp
```

Note that in the latter case, the library `cuda` is not needed. The CUDA Runtime will try to open explicitly the `cuda` library if needed. In the case of a system which does not have

the CUDA driver installed, this allows the application to gracefully manage this issue and potentially run if a CPU-only path is available.

# Chapter 2.

# USING THE CUSPARSE API

This chapter describes how to use the cuSPARSE library API. It is not a reference for the cuSPARSE API data types and functions; that is provided in subsequent chapters.

## 2.1. Thread Safety

The library is thread safe and its functions can be called from multiple host threads. However, simultaneous read/writes of the same objects (or of the same handle) are not safe. Hence the handle must be private per thread, i.e., only one handle per thread is safe.

## 2.2. Scalar Parameters

In the cuSPARSE API, the scalar parameters  $\alpha$  and  $\beta$  can be passed by reference on the host or the device.

The few functions that return a scalar result, such as `doti()` and `nnz()`, return the resulting value by reference on the host or the device. Even though these functions return immediately, similarly to those that return matrix and vector results, the scalar result is not ready until execution of the routine on the GPU completes. This requires proper synchronization be used when reading the result from the host.

This feature allows the cuSPARSE library functions to execute completely asynchronously using streams, even when  $\alpha$  and  $\beta$  are generated by a previous kernel. This situation arises, for example, when the library is used to implement iterative methods for the solution of linear systems and eigenvalue problems [3].

## 2.3. Parallelism with Streams

If the application performs several small independent computations, or if it makes data transfers in parallel with the computation, CUDA streams can be used to overlap these tasks.

The application can conceptually associate a stream with each task. To achieve the overlap of computation between the tasks, the developer should create CUDA streams using the function `cudaStreamCreate()` and set the stream to be used by each individual cuSPARSE library routine by calling `cusparseSetStream()` just before calling the actual cuSPARSE routine. Then, computations performed in separate streams would be overlapped automatically on the GPU, when possible. This approach is especially useful when the computation performed by a single task is relatively small and is not enough to fill the GPU with work, or when there is a data transfer that can be performed in parallel with the computation.

When streams are used, we recommend using the new cuSPARSE API with scalar parameters and results passed by reference in the device memory to achieve maximum computational overlap.

Although a developer can create many streams, in practice it is not possible to have more than 16 concurrent kernels executing at the same time.

## 2.4. Compatibility and Versioning

The cuSPARSE APIs are intended to be backward compatible at the source level with future releases (unless stated otherwise in the release notes of a specific future release). In other words, if a program uses cuSPARSE, it should continue to compile and work correctly with newer versions of cuSPARSE without source code changes. cuSPARSE is not guaranteed to be backward compatible at the binary level. Using different versions of the `cusparse.h` header file and the shared library is not supported. Using different versions of cuSPARSE and the CUDA runtime is not supported. The APIs should be backward compatible at the source level for public functions in most cases

# Chapter 3.

# CUSPARSE INDEXING AND DATA FORMATS

The cuSPARSE library supports dense and sparse vector, and dense and sparse matrix formats.

## 3.1. Index Base Format

The library supports zero- and one-based indexing. The index base is selected through the `cusparseIndexBase_t` type, which is passed as a standalone parameter or as a field in the matrix descriptor `cusparseMatDescr_t` type.

### 3.1.1. Vector Formats

This section describes dense and sparse vector formats.

#### 3.1.1.1. Dense Format

Dense vectors are represented with a single data array that is stored linearly in memory, such as the following  $7 \times 1$  dense vector.

$$[1.0 \ 0.0 \ 0.0 \ 2.0 \ 3.0 \ 0.0 \ 4.0]$$

(This vector is referenced again in the next section.)

#### 3.1.1.2. Sparse Format

Sparse vectors are represented with two arrays.

- ▶ The *data array* has the nonzero values from the equivalent array in dense format.
- ▶ The *integer index array* has the positions of the corresponding nonzero values in the equivalent array in dense format.

For example, the dense vector in section 3.2.1 can be stored as a sparse vector with one-based indexing.

$$\begin{bmatrix} 1.0 & 2.0 & 3.0 & 4.0 \\ 1 & 4 & 5 & 7 \end{bmatrix}$$

It can also be stored as a sparse vector with zero-based indexing.

$$\begin{bmatrix} 1.0 & 2.0 & 3.0 & 4.0 \\ 0 & 3 & 4 & 6 \end{bmatrix}$$

In each example, the top row is the data array and the bottom row is the index array, and it is assumed that the indices are provided in increasing order and that each index appears only once.

## 3.2. Matrix Formats

Dense and several sparse formats for matrices are discussed in this section.

### 3.2.1. Dense Format

The dense matrix **x** is assumed to be stored in column-major format in memory and is represented by the following parameters.

<b>m</b>	(integer)	The number of rows in the matrix.
<b>n</b>	(integer)	The number of columns in the matrix.
<b>lidx</b>	(integer)	The leading dimension of <b>x</b> , which must be greater than or equal to <b>m</b> . If <b>lidx</b> is greater than <b>m</b> , then <b>x</b> represents a sub-matrix of a larger matrix stored in memory
<b>x</b>	(pointer)	Points to the data array containing the matrix elements. It is assumed that enough storage is allocated for <b>x</b> to hold all of the matrix elements and that cuSPARSE library functions may access values outside of the sub-matrix, but will never overwrite them.

For example,  $m \times n$  dense matrix **x** with leading dimension **lidx** can be stored with one-based indexing as shown.

$$\begin{bmatrix} X_{1,1} & X_{1,2} & \cdots & X_{1,n} \\ X_{2,1} & X_{2,2} & \cdots & X_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{m,1} & X_{m,2} & \cdots & X_{m,n} \\ \vdots & \vdots & \ddots & \vdots \\ X_{lidx,1} & X_{lidx,2} & \cdots & X_{lidx,n} \end{bmatrix}$$

Its elements are arranged linearly in memory in the order below.

$$[X_{1,1} \ X_{2,1} \ \cdots \ X_{m,1} \ \cdots \ X_{lidx,1} \ \cdots \ X_{1,n} \ X_{2,n} \ \cdots \ X_{m,n} \ \cdots \ X_{lidx,n}]$$



This format and notation are similar to those used in the NVIDIA CUDA cuBLAS library.

### 3.2.2. Coordinate Format (COO)

The  $m \times n$  sparse matrix **A** is represented in COO format by the following parameters.

<b>nnz</b>	(integer)	The number of nonzero elements in the matrix.
<b>cooValA</b>	(pointer)	Points to the data array of length <b>nnz</b> that holds all nonzero values of <b>A</b> in row-major format.
<b>cooRowIndA</b>	(pointer)	Points to the integer array of length <b>nnz</b> that contains the row indices of the corresponding elements in array <b>cooValA</b> .
<b>cooColIndA</b>	(pointer)	Points to the integer array of length <b>nnz</b> that contains the column indices of the corresponding elements in array <b>cooValA</b> .

A sparse matrix in COO format is assumed to be stored in row-major format: the index arrays are first sorted by row indices and then within the same row by compressed column indices. It is assumed that each pair of row and column indices appears only once.

For example, consider the following  $4 \times 5$  matrix **A**.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 \end{bmatrix}$$

It is stored in COO format with zero-based indexing this way.

$$\begin{aligned} \text{cooValA} &= [1.0 \ 4.0 \ 2.0 \ 3.0 \ 5.0 \ 7.0 \ 8.0 \ 9.0 \ 6.0] \\ \text{cooRowIndA} &= [0 \ 0 \ 1 \ 1 \ 2 \ 2 \ 2 \ 3 \ 3] \\ \text{cooColIndA} &= [0 \ 1 \ 1 \ 2 \ 0 \ 3 \ 4 \ 2 \ 4] \end{aligned}$$

In the COO format with one-based indexing, it is stored as shown.

$$\begin{aligned} \text{cooValA} &= [1.0 \ 4.0 \ 2.0 \ 3.0 \ 5.0 \ 7.0 \ 8.0 \ 9.0 \ 6.0] \\ \text{cooRowIndA} &= [1 \ 1 \ 2 \ 2 \ 3 \ 3 \ 3 \ 4 \ 4] \\ \text{cooColIndA} &= [1 \ 2 \ 2 \ 3 \ 1 \ 4 \ 5 \ 3 \ 5] \end{aligned}$$

### 3.2.3. Compressed Sparse Row Format (CSR)

The only way the CSR differs from the COO format is that the array containing the row indices is compressed in CSR format. The  $m \times n$  sparse matrix **A** is represented in CSR format by the following parameters.

<b>nnz</b>	(integer)	The number of nonzero elements in the matrix.
<b>csrValA</b>	(pointer)	Points to the data array of length <b>nnz</b> that holds all nonzero values of <b>A</b> in row-major format.
<b>csrRowPtrA</b>	(pointer)	Points to the integer array of length <b>m+1</b> that holds indices into the arrays <b>csrColIndA</b> and <b>csrValA</b> . The first <b>m</b> entries of this array contain the indices of the first nonzero element in the <b>i</b> th row for <b>i=i, ..., m</b> , while the last entry contains <b>nnz+csrRowPtrA(0)</b> . In general, <b>csrRowPtrA(0)</b> is 0 or 1 for zero- and one-based indexing, respectively.
<b>csrColIndA</b>	(pointer)	Points to the integer array of length <b>nnz</b> that contains the column indices of the corresponding elements in array <b>csrValA</b> .

Sparse matrices in CSR format are assumed to be stored in row-major CSR format, in other words, the index arrays are first sorted by row indices and then within the same row by column indices. It is assumed that each pair of row and column indices appears only once.

Consider again the  $4 \times 5$  matrix **A**.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 \end{bmatrix}$$

It is stored in CSR format with zero-based indexing as shown.

$$\begin{aligned} \text{csrValA} &= [1.0 \ 4.0 \ 2.0 \ 3.0 \ 5.0 \ 7.0 \ 8.0 \ 9.0 \ 6.0] \\ \text{csrRowPtrA} &= [0 \ 2 \ 4 \ 7 \ 9] \\ \text{csrColIndA} &= [0 \ 1 \ 1 \ 2 \ 0 \ 3 \ 4 \ 2 \ 4] \end{aligned}$$

This is how it is stored in CSR format with one-based indexing.

$$\begin{aligned} \text{csrValA} &= [1.0 \ 4.0 \ 2.0 \ 3.0 \ 5.0 \ 7.0 \ 8.0 \ 9.0 \ 6.0] \\ \text{csrRowPtrA} &= [1 \ 3 \ 5 \ 8 \ 10] \\ \text{csrColIndA} &= [1 \ 2 \ 2 \ 3 \ 1 \ 4 \ 5 \ 3 \ 5] \end{aligned}$$

### 3.2.4. Compressed Sparse Column Format (CSC)

The CSC format is different from the COO format in two ways: the matrix is stored in column-major format, and the array containing the column indices is compressed in CSC format. The  $m \times n$  matrix **A** is represented in CSC format by the following parameters.

<b>nnz</b>	(integer)	The number of nonzero elements in the matrix.
<b>cscValA</b>	(pointer)	Points to the data array of length <b>nnz</b> that holds all nonzero values of <b>A</b> in column-major format.
<b>cscRowIndA</b>	(pointer)	Points to the integer array of length <b>nnz</b> that contains the row indices of the corresponding elements in array <b>cscValA</b> .
<b>cscColPtrA</b>	(pointer)	Points to the integer array of length <b>n+1</b> that holds indices into the arrays <b>cscRowIndA</b> and <b>cscValA</b> . The first <b>n</b> entries of this array contain the indices of the first nonzero element in the <b>i</b> th row for <b>i=i,...,n</b> , while the last entry contains <b>nnz+cscColPtrA(0)</b> . In general, <b>cscColPtrA(0)</b> is 0 or 1 for zero- and one-based indexing, respectively.



The matrix **A** in CSR format has exactly the same memory layout as its transpose in CSC format (and vice versa).

For example, consider once again the  $4 \times 5$  matrix **A**.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 \end{bmatrix}$$

It is stored in CSC format with zero-based indexing this way.

```
cscValA = [1.0 5.0 4.0 2.0 3.0 9.0 7.0 8.0 6.0]
cscRowIndA = [0 2 0 1 1 3 2 2 3 ]
cscColPtrA = [0 2 4 6 7 9 ]
```

In CSC format with one-based indexing, this is how it is stored.

```
cscValA = [1.0 5.0 4.0 2.0 3.0 9.0 7.0 8.0 6.0]
cscRowIndA = [1 3 1 2 2 4 3 3 4 ]
cscColPtrA = [1 3 5 7 8 10 ]
```

Each pair of row and column indices appears only once.

### 3.2.5. Ellpack-Itpack Format (ELL)

*The storage format will be removed in the next major release*

An  $m \times n$  sparse matrix **A** with at most **k** nonzero elements per row is stored in the Ellpack-Itpack (ELL) format [2] using two dense arrays of dimension  $m \times k$ . The first data array contains the values of the nonzero elements in the matrix, while the second integer array contains the corresponding column indices.

For example, consider the  $4 \times 5$  matrix **A**.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 \end{bmatrix}$$

This is how it is stored in ELL format with zero-based indexing.

$$\begin{aligned} \text{data} &= \begin{bmatrix} 1.0 & 4.0 & 0.0 \\ 2.0 & 3.0 & 0.0 \\ 5.0 & 7.0 & 8.0 \\ 9.0 & 6.0 & 0.0 \end{bmatrix} \\ \text{indices} &= \begin{bmatrix} 0 & 1 & -1 \\ 1 & 2 & -1 \\ 0 & 3 & 4 \\ 2 & 4 & -1 \end{bmatrix} \end{aligned}$$

It is stored this way in ELL format with one-based indexing.

$$\begin{aligned} \text{data} &= \begin{bmatrix} 1.0 & 4.0 & 0.0 \\ 2.0 & 3.0 & 0.0 \\ 5.0 & 7.0 & 8.0 \\ 9.0 & 6.0 & 0.0 \end{bmatrix} \\ \text{indices} &= \begin{bmatrix} 1 & 2 & -1 \\ 2 & 3 & -1 \\ 1 & 4 & 5 \\ 3 & 5 & -1 \end{bmatrix} \end{aligned}$$

Sparse matrices in ELL format are assumed to be stored in column-major format in memory. Also, rows with less than **k** nonzero elements are padded in the **data** and **indices** arrays with zero and  $-1$ , respectively.

The ELL format is not supported directly, but it is used to store the regular part of the matrix in the HYB format that is described in the next section.

### 3.2.6. Hybrid Format (HYB) [DEPRECATED]

**[[DEPRECATED]]** *The storage format will be removed in the next major release*

The HYB sparse storage format is composed of a regular part, usually stored in ELL format, and an irregular part, usually stored in COO format [1]. The ELL and COO parts are always stored using zero-based indexing. HYB is implemented as an opaque data format that requires the use of a conversion operation to store a matrix in it. The conversion operation partitions the general matrix into the regular and irregular parts automatically or according to developer-specified criteria.

For more information, please refer to the description of **`cusparseHybPartition_t`** type, as well as the description of the conversion routines **`dense2hyb`**, **`csc2hyb`** and **`csr2hyb`**.

### 3.2.7. Block Compressed Sparse Row Format (BSR)

The only difference between the CSR and BSR formats is the format of the storage element. The former stores primitive data types (**`single`**, **`double`**, **`cuComplex`**, and **`cuDoubleComplex`**) whereas the latter stores a two-dimensional square block of primitive data types. The dimension of the square block is *blockDim*. The **`m`** × **`n`** sparse matrix **`A`** is equivalent to a block sparse matrix  $A_b$  with  $mb = \frac{m + blockDim - 1}{blockDim}$  block rows and  $nb = \frac{n + blockDim - 1}{blockDim}$  block columns. If  $m$  or  $n$  is not multiple of *blockDim*, then zeros are filled into  $A_b$ .

**`A`** is represented in BSR format by the following parameters.

<b><code>blockDim</code></b>	(integer)	Block dimension of matrix <b><code>A</code></b> .
<b><code>mb</code></b>	(integer)	The number of block rows of <b><code>A</code></b> .
<b><code>nb</code></b>	(integer)	The number of block columns of <b><code>A</code></b> .
<b><code>nnzb</code></b>	(integer)	The number of nonzero blocks in the matrix.
<b><code>bsrValA</code></b>	(pointer)	Points to the data array of length $nnzb * blockDim^2$ that holds all elements of nonzero blocks of <b><code>A</code></b> . The block elements are stored in either column-major order or row-major order.
<b><code>bsrRowPtrA</code></b>	(pointer)	Points to the integer array of length $mb+1$ that holds indices into the arrays <b><code>bsrColIndA</code></b> and <b><code>bsrValA</code></b> . The first $mb$ entries of this array contain the indices of the first nonzero block in the $i$ th block row for $i=1, \dots, mb$ , while the last entry contains $nnzb + bsrRowPtrA(0)$ . In general, <b><code>bsrRowPtrA(0)</code></b> is 0 or 1 for zero- and one-based indexing, respectively.
<b><code>bsrColIndA</code></b>	(pointer)	Points to the integer array of length <b><code>nnzb</code></b> that contains the column indices of the corresponding blocks in array <b><code>bsrValA</code></b> .

As with CSR format, (row, column) indices of BSR are stored in row-major order. The index arrays are first sorted by row indices and then within the same row by column indices.

For example, consider again the **4×5** matrix **A**.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 \end{bmatrix}$$

If *blockDim* is equal to 2, then *mb* is 2, *nb* is 3, and matrix **A** is split into **2×3** block matrix **A<sub>b</sub>**. The dimension of **A<sub>b</sub>** is **4×6**, slightly bigger than matrix **A**, so zeros are filled in the last column of **A<sub>b</sub>**. The element-wise view of **A<sub>b</sub>** is this.

$$\begin{bmatrix} 1.0 & 4.0 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 2.0 & 3.0 & 0.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 0.0 & 7.0 & 8.0 & 0.0 \\ 0.0 & 0.0 & 9.0 & 0.0 & 6.0 & 0.0 \end{bmatrix}$$

Based on zero-based indexing, the block-wise view of **A<sub>b</sub>** can be represented as follows.

$$A_b = \begin{bmatrix} A_{00} & A_{01} & A_{02} \\ A_{10} & A_{11} & A_{12} \end{bmatrix}$$

The basic element of BSR is a nonzero  $A_{ij}$  block, one that contains at least one nonzero element of **A**. Five of six blocks are nonzero in **A<sub>b</sub>**,

$$A_{00} = \begin{bmatrix} 1 & 4 \\ 0 & 2 \end{bmatrix}, A_{01} = \begin{bmatrix} 0 & 0 \\ 3 & 0 \end{bmatrix}, A_{10} = \begin{bmatrix} 5 & 0 \\ 0 & 0 \end{bmatrix}, A_{11} = \begin{bmatrix} 0 & 7 \\ 9 & 0 \end{bmatrix}, A_{12} = \begin{bmatrix} 8 & 0 \\ 6 & 0 \end{bmatrix}$$

BSR format only stores the information of nonzero blocks, including block indices (*i*, *j*) and values  $A_{ij}$ . Also row indices are compressed in CSR format.

$$\begin{aligned} \text{bsrValA} &= [A_{00} \ A_{01} \ A_{10} \ A_{11} \ A_{12}] \\ \text{bsrRowPtrA} &= [0 \ 2 \ 5] \\ \text{bsrColIndA} &= [0 \ 1 \ 0 \ 1 \ 2] \end{aligned}$$

There are two ways to arrange the data element of block  $A_{ij}$ : row-major order and column-major order. Under column-major order, the physical storage of **bsrValA** is this.

$$\text{bsrValA} = [1 \ 0 \ 4 \ 2 \ | 0 \ 3 \ 0 \ 0 \ | 5 \ 0 \ 0 \ 0 \ | 0 \ 9 \ 7 \ 0 \ | 8 \ 6 \ 0 \ 0]$$

Under row-major order, the physical storage of **bsrValA** is this.

$$\text{bsrValA} = [1 \ 4 \ 0 \ 2 \ | 0 \ 0 \ 3 \ 0 \ | 5 \ 0 \ 0 \ 0 \ | 0 \ 7 \ 9 \ 0 \ | 8 \ 0 \ 6 \ 0]$$

Similarly, in BSR format with one-based indexing and column-major order, **A** can be represented by the following.

$$A_b = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \end{bmatrix}$$

$$\text{bsrValA} = [1 \ 0 \ 4 \ 2 \ | 0 \ 3 \ 0 \ 0 \ | 5 \ 0 \ 0 \ 0 \ | 0 \ 9 \ 7 \ 0 \ | 8 \ 6 \ 0 \ 0]$$

```
bsrRowPtrA = [1   3   6]
bsrColIndA = [1   2   1   2   3]
```



The general BSR format has two parameters, `rowBlockDim` and `colBlockDim`. `rowBlockDim` is number of rows within a block and `colBlockDim` is number of columns within a block. If `rowBlockDim=colBlockDim`, general BSR format is the same as BSR format. If `rowBlockDim=colBlockDim=1`, general BSR format is the same as CSR format. The conversion routine `gebsr2gebsr` is used to do conversion among CSR, BSR and general BSR.



In the cuSPARSE Library, the storage format of blocks in BSR format can be column-major or row-major, independently of the base index. However, if the developer uses BSR format from the Math Kernel Library (MKL) and wants to directly interface with the cuSPARSE Library, then `cusparseDirection_t CUSPARSE_DIRECTION_COLUMN` should be used if the base index is one; otherwise, `cusparseDirection_t CUSPARSE_DIRECTION_ROW` should be used.

### 3.2.8. Extended BSR Format (BSRX)

BSRX is the same as the BSR format, but the array `bsrRowPtrA` is separated into two parts. The first nonzero block of each row is still specified by the array `bsrRowPtrA`, which is the same as in BSR, but the position next to the last nonzero block of each row is specified by the array `bsrEndPtrA`. Briefly, BSRX format is simply like a 4-vector variant of BSR format.

Matrix **A** is represented in BSRX format by the following parameters.

<code>blockDim</code>	(integer)	Block dimension of matrix <b>A</b> .
<code>mb</code>	(integer)	The number of block rows of <b>A</b> .
<code>nb</code>	(integer)	The number of block columns of <b>A</b> .
<code>nnzb</code>	(integer)	number of nonzero blocks in the matrix <b>A</b> .
<code>bsrValA</code>	(pointer)	Points to the data array of length $nnzb * blockDim^2$ that holds all the elements of the nonzero blocks of <b>A</b> . The block elements are stored in either column-major order or row-major order.
<code>bsrRowPtrA</code>	(pointer)	Points to the integer array of length <code>mb</code> that holds indices into the arrays <code>bsrColIndA</code> and <code>bsrValA</code> ; <code>bsrRowPtrA(i)</code> is the position of the first nonzero block of the <i>i</i> th block row in <code>bsrColIndA</code> and <code>bsrValA</code> .
<code>bsrEndPtrA</code>	(pointer)	Points to the integer array of length <code>mb</code> that holds indices into the arrays <code>bsrColIndA</code> and <code>bsrValA</code> ; <code>bsrRowPtrA(i)</code> is the position next to the last nonzero block of the <i>i</i> th block row in <code>bsrColIndA</code> and <code>bsrValA</code> .
<code>bsrColIndA</code>	(pointer)	Points to the integer array of length <code>nnzb</code> that contains the column indices of the corresponding blocks in array <code>bsrValA</code> .

A simple conversion between BSR and BSRX can be done as follows. Suppose the developer has a  $2 \times 3$  block sparse matrix  $A_b$  represented as shown.

$$A_b = \begin{bmatrix} A_{00} & A_{01} & A_{02} \\ A_{10} & A_{11} & A_{12} \end{bmatrix}$$

Assume it has this BSR format.

$$\begin{aligned}\text{bsrValA of BSR} &= [A_{00} \ A_{01} \ A_{10} \ A_{11} \ A_{12}] \\ \text{bsrRowPtrA of BSR} &= [0 \ 2 \ 5] \\ \text{bsrColIndA of BSR} &= [0 \ 1 \ 0 \ 1 \ 2]\end{aligned}$$

The **bsrRowPtrA** of the BSRX format is simply the first two elements of the **bsrRowPtrA** BSR format. The **bsrEndPtrA** of BSRX format is the last two elements of the **bsrRowPtrA** of BSR format.

$$\begin{aligned}\text{bsrRowPtrA of BSRX} &= [0 \ 2] \\ \text{bsrEndPtrA of BSRX} &= [2 \ 5]\end{aligned}$$

The advantage of the BSRX format is that the developer can specify a submatrix in the original BSR format by modifying **bsrRowPtrA** and **bsrEndPtrA** while keeping **bsrColIndA** and **bsrValA** unchanged.

For example, to create another block matrix  $\tilde{A} = \begin{bmatrix} O & O & O \\ O & A_{11} & O \end{bmatrix}$  that is slightly different from  $A$ , the developer can keep **bsrColIndA** and **bsrValA**, but reconstruct  $\tilde{A}$  by properly setting of **bsrRowPtrA** and **bsrEndPtrA**. The following 4-vector characterizes  $\tilde{A}$ .

$$\begin{aligned}\text{bsrValA of } \tilde{A} &= [A_{00} \ A_{01} \ A_{10} \ A_{11} \ A_{12}] \\ \text{bsrColIndA of } \tilde{A} &= [0 \ 1 \ 0 \ 1 \ 2] \\ \text{bsrRowPtrA of } \tilde{A} &= [0 \ 3] \\ \text{bsrEndPtrA of } \tilde{A} &= [0 \ 4]\end{aligned}$$

# Chapter 4. CUSPARSE TYPES REFERENCE

## 4.1. Data types

The `float`, `double`, `cuComplex`, and `cuDoubleComplex` data types are supported. The first two are standard C data types, while the last two are exported from `cuComplex.h`.

## 4.2. cusparseAction\_t

This type indicates whether the operation is performed only on indices or on data and indices.

Value	Meaning
<code>CUSPARSE_ACTION_SYMBOLIC</code>	the operation is performed only on indices.
<code>CUSPARSE_ACTION_NUMERIC</code>	the operation is performed on data and indices.

## 4.3. cusparseDirection\_t

This type indicates whether the elements of a dense matrix should be parsed by rows or by columns (assuming column-major storage in memory of the dense matrix) in function `cusparse[S|D|C|Z]nnz`. Besides storage format of blocks in BSR format is also controlled by this type.

Value	Meaning
<code>CUSPARSE_DIRECTION_ROW</code>	the matrix should be parsed by rows.
<code>CUSPARSE_DIRECTION_COLUMN</code>	the matrix should be parsed by columns.

## 4.4. cusparseHandle\_t

This is a pointer type to an opaque cuSPARSE context, which the user must initialize by calling prior to calling `cusparseCreate()` any other library function. The handle created and returned by `cusparseCreate()` must be passed to every cuSPARSE function.

## 4.5. cusparseHybMat\_t [DEPRECATED]

**[[DEPRECATED]]** *The data type will be removed in the next major release*

This is a pointer type to an opaque structure holding the matrix in HYB format, which is created by `cusparseCreateHybMat` and destroyed by `cusparseDestroyHybMat`.

### 4.5.1. cusparseHybPartition\_t [DEPRECATED]

**[[DEPRECATED]]** *The data type will be removed in the next major release*

This type indicates how to perform the partitioning of the matrix into regular (ELL) and irregular (COO) parts of the HYB format.

The partitioning is performed during the conversion of the matrix from a dense or sparse format into the HYB format and is governed by the following rules. When `CUSPARSE_HYB_PARTITION_AUTO` is selected, the cuSPARSE library automatically decides how much data to put into the regular and irregular parts of the HYB format. When `CUSPARSE_HYB_PARTITION_USER` is selected, the width of the regular part of the HYB format should be specified by the caller. When `CUSPARSE_HYB_PARTITION_MAX` is selected, the width of the regular part of the HYB format equals to the maximum number of non-zero elements per row, in other words, the entire matrix is stored in the regular part of the HYB format.

The *default* is to let the library automatically decide how to split the data.

Value	Meaning
<code>CUSPARSE_HYB_PARTITION_AUTO</code>	the automatic partitioning is selected ( <i>default</i> ).
<code>CUSPARSE_HYB_PARTITION_USER</code>	the user specified threshold is used.
<code>CUSPARSE_HYB_PARTITION_MAX</code>	the data is stored in ELL format.

## 4.6. cusparseMatDescr\_t

This structure is used to describe the shape and properties of a matrix.

```
typedef struct {
    cusparseMatrixType_t MatrixType;
    cusparseFillMode_t FillMode;
    cusparseDiagType_t DiagType;
    cusparseIndexBase_t IndexBase;
} cusparseMatDescr_t;
```

### 4.6.1. cusparseDiagType\_t

This type indicates if the matrix diagonal entries are unity. The diagonal elements are always assumed to be present, but if **CUSPARSE\_DIAG\_TYPE\_UNIT** is passed to an API routine, then the routine assumes that all diagonal entries are unity and will not read or modify those entries. Note that in this case the routine assumes the diagonal entries are equal to one, regardless of what those entries are actually set to in memory.

Value	Meaning
<b>CUSPARSE_DIAG_TYPE_NON_UNIT</b>	the matrix diagonal has non-unit elements.
<b>CUSPARSE_DIAG_TYPE_UNIT</b>	the matrix diagonal has unit elements.

### 4.6.2. cusparseFillMode\_t

This type indicates if the lower or upper part of a matrix is stored in sparse storage.

Value	Meaning
<b>CUSPARSE_FILL_MODE_LOWER</b>	the lower triangular part is stored.
<b>CUSPARSE_FILL_MODE_UPPER</b>	the upper triangular part is stored.

### 4.6.3. cusparseIndexBase\_t

This type indicates if the base of the matrix indices is zero or one.

Value	Meaning
<b>CUSPARSE_INDEX_BASE_ZERO</b>	the base index is zero.
<b>CUSPARSE_INDEX_BASE_ONE</b>	the base index is one.

### 4.6.4. cusparseMatrixType\_t

This type indicates the type of matrix stored in sparse storage. Notice that for symmetric, Hermitian and triangular matrices only their lower or upper part is assumed to be stored.

The whole idea of matrix type and fill mode is to keep minimum storage for symmetric/Hermitian matrix, and also to take advantage of symmetric property on SpMV

(Sparse Matrix Vector multiplication). To compute  $\mathbf{y} = \mathbf{A} * \mathbf{x}$  when  $\mathbf{A}$  is symmetric and only lower triangular part is stored, two steps are needed. First step is to compute  $\mathbf{y} = (\mathbf{L} + \mathbf{D}) * \mathbf{x}$  and second step is to compute  $\mathbf{y} = \mathbf{L}^T * \mathbf{x} + \mathbf{y}$ . Given the fact that the transpose operation  $\mathbf{y} = \mathbf{L}^T * \mathbf{x}$  is 10x slower than non-transpose version  $\mathbf{y} = \mathbf{L} * \mathbf{x}$ , the symmetric property does not show up any performance gain. It is better for the user to extend the symmetric matrix to a general matrix and apply  $\mathbf{y} = \mathbf{A} * \mathbf{x}$  with matrix type `CUSPARSE_MATRIX_TYPE_GENERAL`.

In general, SpMV, preconditioners (incomplete Cholesky or incomplete LU) and triangular solver are combined together in iterative solvers, for example PCG and GMRES. If the user always uses general matrix (instead of symmetric matrix), there is no need to support other than general matrix in preconditioners. Therefore the new routines, `[bsr|csr]sv2` (triangular solver), `[bsr|csr]ilu02` (incomplete LU) and `[bsr|csr]ic02` (incomplete Cholesky), only support matrix type `CUSPARSE_MATRIX_TYPE_GENERAL`.

Value	Meaning
<code>CUSPARSE_MATRIX_TYPE_GENERAL</code>	the matrix is general.
<code>CUSPARSE_MATRIX_TYPE_SYMMETRIC</code>	the matrix is symmetric.
<code>CUSPARSE_MATRIX_TYPE_HERMITIAN</code>	the matrix is Hermitian.
<code>CUSPARSE_MATRIX_TYPE_TRIANGULAR</code>	the matrix is triangular.

## 4.7. `cusparseOperation_t`

This type indicates which operations need to be performed with the sparse matrix.

Value	Meaning
<code>CUSPARSE_OPERATION_NON_TRANSPOSE</code>	the non-transpose operation is selected.
<code>CUSPARSE_OPERATION_TRANSPOSE</code>	the transpose operation is selected.
<code>CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE</code>	the conjugate transpose operation is selected.

## 4.8. `cusparsePointerMode_t`

This type indicates whether the scalar values are passed by reference on the host or device. It is important to point out that if several scalar values are passed by reference in the function call, all of them will conform to the same single pointer mode. The pointer mode can be set and retrieved using `cusparseSetPointerMode()` and `cusparseGetPointerMode()` routines, respectively.

Value	Meaning
<code>CUSPARSE_POINTER_MODE_HOST</code>	the scalars are passed by reference on the host.
<code>CUSPARSE_POINTER_MODE_DEVICE</code>	the scalars are passed by reference on the device.

## 4.9. `cusparseAlgMode_t`

This is type for algorithm parameter to `cusparseCsrsvEx()` and `cusparseCsrsvEx_bufferSize()` functions. Note that previously defined values `CUSPARSE_ALG0` (for naive algorithm) and `CUSPARSE_ALG1` (for merge path algorithm) are deprecated and replaced by named aliases specified below.

Value	Meaning
<code>CUSPARSE_ALG_NAIVE</code>	Use default naive algorithm.
<code>CUSPARSE_ALG.Merge_Path</code>	Use load-balancing algorithm that suits better for irregular nonzero-patterns.

## 4.10. `cusparseSolvePolicy_t`

This type indicates whether level information is generated and used in `csrsv2`, `csric02`, `csrilu02`, `bsrsv2`, `bsric02` and `bsrilu02`.

Value	Meaning
<code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code>	no level information is generated and used.
<code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code>	generate and use level information.

## 4.11. `cusparseSolveAnalysisInfo_t` [DEPRECATED]

This is a pointer type to an opaque structure holding the information collected in the analysis phase of the solution of the sparse triangular linear system. It is expected to be passed unchanged to the solution phase of the sparse triangular linear system.

## 4.12. `cusparseColorInfo_t`

This is a pointer type to an opaque structure holding the information used in `csrcolor()`.

## 4.13. `csrsv2Info_t`

This is a pointer type to an opaque structure holding the information used in `csrsv2_bufferSize()`, `csrsv2_analysis()`, and `csrsv2_solve()`.

## 4.14. `csrsm2Info_t`

This is a pointer type to an opaque structure holding the information used in `csrsm2_bufferSize()`, `csrsm2_analysis()`, and `csrsm2_solve()`.

## 4.15. `csric02Info_t`

This is a pointer type to an opaque structure holding the information used in `csric02_bufferSize()`, `csric02_analysis()`, and `csric02()`.

## 4.16. `csrilu02Info_t`

This is a pointer type to an opaque structure holding the information used in `csrilu02_bufferSize()`, `csrilu02_analysis()`, and `csrilu02()`.

## 4.17. `bsrsv2Info_t`

This is a pointer type to an opaque structure holding the information used in `bsrsv2_bufferSize()`, `bsrsv2_analysis()`, and `bsrsv2_solve()`.

## 4.18. `bsrsm2Info_t`

This is a pointer type to an opaque structure holding the information used in `bsrsm2_bufferSize()`, `bsrsm2_analysis()`, and `bsrsm2_solve()`.

## 4.19. `bsric02Info_t`

This is a pointer type to an opaque structure holding the information used in `bsric02_bufferSize()`, `bsric02_analysis()`, and `bsric02()`.

## 4.20. `bsrilu02Info_t`

This is a pointer type to an opaque structure holding the information used in `bsrilu02_bufferSize()`, `bsrilu02_analysis()`, and `bsrilu02()`.

## 4.21. `csrgemm2Info_t`

This is a pointer type to an opaque structure holding the information used in `csrgemm2_bufferSizeExt()`, and `csrgemm2()`.

## 4.22. `cusparseStatus_t`

This data type represents the status returned by the library functions and it can have the following values

Value	Description
<code>CUSPARSE_STATUS_SUCCESS</code>	The operation completed successfully
<code>CUSPARSE_STATUS_NOT_INITIALIZED</code>	The cuSPARSE library was not initialized. This is usually caused by the lack of a prior call, an error in the CUDA Runtime API called by the cuSPARSE routine, or an error in the hardware setup
<code>CUSPARSE_STATUS_NOT_INITIALIZED</code>	<p><b>To correct:</b> call <code>cusparseCreate()</code> prior to the function call; and check that the hardware, an appropriate version of the driver, and the cuSPARSE library are correctly installed</p> <p>The error also applies to generic APIs (<a href="#">Generic APIs reference</a>) for indicating a matrix/vector descriptor not initialized</p>
<code>CUSPARSE_STATUS_ALLOC_FAILED</code>	<p>Resource allocation failed inside the cuSPARSE library. This is usually caused by a device memory allocation (<code>cudaMalloc()</code>) or by a host memory allocation failure</p> <p><b>To correct:</b> prior to the function call, deallocate previously allocated memory as much as possible</p>
<code>CUSPARSE_STATUS_INVALID_VALUE</code>	<p>An unsupported value or parameter was passed to the function (a negative vector size, for example)</p> <p><b>To correct:</b> ensure that all the parameters being passed have valid values</p>
<code>CUSPARSE_STATUS_ARCH_MISMATCH</code>	<p>The function requires a feature absent from the device architecture</p> <p><b>To correct:</b> compile and run the application on a device with appropriate compute capability</p>
<code>CUSPARSE_STATUS_EXECUTION_FAILED</code>	<p>The GPU program failed to execute. This is often caused by a launch failure of the kernel on the GPU, which can be caused by multiple reasons</p> <p><b>To correct:</b> check that the hardware, an appropriate version of the driver, and the cuSPARSE library are correctly installed</p>
<code>CUSPARSE_STATUS_INTERNAL_ERROR</code>	<p>An internal cuSPARSE operation failed</p> <p><b>To correct:</b> check that the hardware, an appropriate version of the driver, and the cuSPARSE library are correctly installed. Also, check that the memory passed as a parameter to the routine is not being deallocated prior to the routine completion</p>
<code>CUSPARSE_STATUS_MATRIX_TYPE_NOT_SUPPORTED</code>	

Value	Description
	<p>The matrix type is not supported by this function. This is usually caused by passing an invalid matrix descriptor to the function</p> <p><b>To correct:</b> check that the fields in <code>cusparseMatDescr_t descrA</code> were set correctly</p>
<code>CUSPARSE_STATUS_NOT_SUPPORT</code>	The operation or data type combination is currently not supported by the function

# Chapter 5. CUSPARSE MANAGEMENT FUNCTION REFERENCE

The cuSPARSE functions for managing the library are described in this section.

## 5.1. `cusparseCreate()`

```
cusparseStatus_t  
cusparseCreate(cusparseHandle_t *handle)
```

This function initializes the cuSPARSE library and creates a handle on the cuSPARSE context. It must be called before any other cuSPARSE API function is invoked. It allocates hardware resources necessary for accessing the GPU.

Param.	In/out	Meaning
<code>handle</code>	IN	The pointer to the handle to the cuSPARSE context

See `cusparseStatus_t` for the description of the return status

## 5.2. `cusparseDestroy()`

```
cusparseStatus_t  
cusparseDestroy(cusparseHandle_t handle)
```

This function releases CPU-side resources used by the cuSPARSE library. The release of GPU-side resources may be deferred until the application shuts down.

Param.	In/out	Meaning
<code>handle</code>	IN	The handle to the cuSPARSE context

See `cusparseStatus_t` for the description of the return status

## 5.3. cusparseGetErrorName()

```
const char*
cusparseGetErrorString(cusparseStatus_t status)
```

The function returns the string representation of an error code enum name. If the error code is not recognized, "unrecognized error code" is returned.

Param.	In/out	Meaning
<b>status</b>	IN	Error code to convert to string
<b>const char*</b>	OUT	Pointer to a NULL-terminated string

## 5.4. cusparseGetErrorString()

```
const char*
cusparseGetErrorString(cusparseStatus_t status)
```

Returns the description string for an error code. If the error code is not recognized, "unrecognized error code" is returned.

Param.	In/out	Meaning
<b>status</b>	IN	Error code to convert to string
<b>const char*</b>	OUT	Pointer to a NULL-terminated string

## 5.5. cusparseGetProperty()

```
cusparseStatus_t
cusparseGetProperty(libraryPropertyType type,
                   int* value)
```

The function returns the value of the requested property. Refer to **libraryPropertyType** for supported types.

Param.	In/out	Meaning
<b>type</b>	IN	Requested property
<b>value</b>	OUT	Value of the requested property

**libraryPropertyType** (defined in **library\_types.h**):

Value	Meaning
<b>MAJOR_VERSION</b>	Enumerator to query the major version
<b>MINOR_VERSION</b>	Enumerator to query the minor version
<b>PATCH_LEVEL</b>	Number to identify the patch level

See [cusparseStatus\\_t](#) for the description of the return status

## 5.6. `cusparseGetVersion()`

```
cusparseStatus_t
cusparseGetVersion(cusparseHandle_t handle,
                  int*           version)
```

This function returns the version number of the cuSPARSE library.

Param.	In/out	Meaning
<code>handle</code>	IN	cuSPARSE handle
<code>version</code>	OUT	The version number of the library

See [cusparseStatus\\_t](#) for the description of the return status

## 5.7. `cusparseGetPointerMode()`

```
cusparseStatus_t
cusparseGetPointerMode(cusparseHandle_t handle,
                      cusparsePointerMode_t *mode)
```

This function obtains the pointer mode used by the cuSPARSE library. Please see the section on the [cusparsePointerMode\\_t](#) type for more details.

Param.	In/out	Meaning
<code>handle</code>	IN	The handle to the cuSPARSE context
<code>mode</code>	OUT	One of the enumerated pointer mode types

See [cusparseStatus\\_t](#) for the description of the return status

## 5.8. `cusparseSetPointerMode()`

```
cusparseStatus_t
cusparseSetPointerMode(cusparseHandle_t handle,
                      cusparsePointerMode_t mode)
```

This function sets the pointer mode used by the cuSPARSE library. The *default* is for the values to be passed by reference on the host. Please see the section on the [cublasPointerMode\\_t](#) type for more details.

Param.	In/out	Meaning
<code>handle</code>	IN	The handle to the cuSPARSE context
<code>mode</code>	IN	One of the enumerated pointer mode types

See [cusparseStatus\\_t](#) for the description of the return status

## 5.9. cusparseGetStream()

```
cusparseStatus_t
cusparseGetStream(cusparseHandle_t handle, cudaStream_t *streamId)
```

This function gets the cuSPARSE library stream, which is being used to execute all calls to the cuSPARSE library functions. If the cuSPARSE library stream is not set, all kernels use the default NULL stream.

Param.	In/out	Meaning
<b>handle</b>	IN	The handle to the cuSPARSE context
<b>streamId</b>	OUT	The stream used by the library

See [cusparseStatus\\_t](#) for the description of the return status

## 5.10. cusparseSetStream()

```
cusparseStatus_t
cusparseSetStream(cusparseHandle_t handle, cudaStream_t streamId)
```

This function sets the stream to be used by the cuSPARSE library to execute its routines.

Param.	In/out	Meaning
<b>handle</b>	IN	The handle to the cuSPARSE context
<b>streamId</b>	IN	The stream to be used by the library

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 6. CUSPARSE HELPER FUNCTION REFERENCE

The cuSPARSE helper functions are described in this section.

## 6.1. `cusparseCreateColorInfo()`

```
cusparseStatus_t  
cusparseCreateColorInfo(cusparseColorInfo_t* info)
```

This function creates and initializes the `cusparseColorInfo_t` structure to *default* values.

### Input

<code>info</code>	the pointer to the <code>cusparseColorInfo_t</code> structure
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.2. `cusparseCreateHybMat()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```
cusparseStatus_t  
cusparseCreateHybMat(cusparseHybMat_t *hybA)
```

This function creates and initializes the `hybA` opaque data structure.

### Input

<code>hybA</code>	the pointer to the hybrid format storage structure.
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.3. `cusparseCreateMatDescr()`

```
cusparseStatus_t  
cusparseCreateMatDescr(cusparseMatDescr_t *descrA)
```

This function initializes the matrix descriptor. It sets the fields **MatrixType** and **IndexBase** to the *default* values **CUSPARSE\_MATRIX\_TYPE\_GENERAL** and **CUSPARSE\_INDEX\_BASE\_ZERO**, respectively, while leaving other fields uninitialized.

#### Input

<code>descrA</code>	the pointer to the matrix descriptor.
---------------------	---------------------------------------

See `cusparseStatus_t` for the description of the return status

## 6.4. `cusparseCreateSolveAnalysisInfo()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```
cusparseStatus_t  
cusparseCreateSolveAnalysisInfo(cusparseSolveAnalysisInfo_t *info)
```

This function creates and initializes the solve and analysis structure to *default* values.

#### Input

<code>info</code>	the pointer to the solve and analysis structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.5. `cusparseDestroyColorInfo()`

```
cusparseStatus_t  
cusparseDestroyColorInfo(cusparseColorInfo_t info)
```

This function destroys and releases any memory required by the structure.

#### Input

<code>info</code>	the pointer to the structure of <code>csrcolor()</code>
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.6. `cusparseDestroyHybMat()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```
cusparseStatus_t  
cusparseDestroyHybMat(cusparseHybMat_t hybA)
```

This function destroys and releases any memory required by the **hybA** structure.

#### Input

<code>hybA</code>	the hybrid format storage structure.
-------------------	--------------------------------------

See `cusparseStatus_t` for the description of the return status

## 6.7. `cusparseDestroyMatDescr()`

```
cusparseStatus_t
cusparseDestroyMatDescr(cusparseMatDescr_t descrA)
```

This function releases the memory allocated for the matrix descriptor.

### Input

<code>descrA</code>	the matrix descriptor.
---------------------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 6.8. `cusparseDestroySolveAnalysisInfo()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseDestroySolveAnalysisInfo(cusparseSolveAnalysisInfo_t info)
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve and analysis structure.
-------------------	-----------------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 6.9. `cusparseGetLevelInfo()`

```
cusparseStatus_t
cusparseGetLevelInfo(cusparseHandle_t handle,
                     cusparseSolveAnalysisInfo_t info,
                     int *nlevels,
                     int **levelPtr,
                     int **levelInd)
```

This function returns the number of levels and the assignment of rows into the levels computed by either the `csrv_analysis`, `csrsm_analysis` or `hybsv_analysis` routines.

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	the pointer to the solve and analysis structure.

### Output

<code>nlevels</code>	number of levels.
----------------------	-------------------

<b>levelPtr</b>	integer array of <b>nlevels+1</b> elements that contains the start of every level and the end of the last level plus one.
<b>levelInd</b>	integer array of <b>m</b> (number of rows in the matrix) elements that contains the row indices belonging to every level.

See `cusparseStatus_t` for the description of the return status

## 6.10. `cusparseGetMatDiagType()`

```
cusparseDiagType_t
cusparseGetMatDiagType(const cusparseMatDescr_t descrA)
```

This function returns the **DiagType** field of the matrix descriptor **descrA**.

### Input

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

### Returned

	One of the enumerated diagType types.
--	---------------------------------------

## 6.11. `cusparseGetMatFillMode()`

```
cusparseFillMode_t
cusparseGetMatFillMode(const cusparseMatDescr_t descrA)
```

This function returns the **FillMode** field of the matrix descriptor **descrA**.

### Input

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

### Returned

	One of the enumerated fillMode types.
--	---------------------------------------

## 6.12. `cusparseGetMatIndexBase()`

```
cusparseIndexBase_t
cusparseGetMatIndexBase(const cusparseMatDescr_t descrA)
```

This function returns the **IndexBase** field of the matrix descriptor **descrA**.

### Input

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

### Returned

	One of the enumerated indexBase types.
--	--

## 6.13. `cusparseGetMatType()`

```
cusparseMatrixType_t  
cusparseGetMatType(const cusparseMatDescr_t descrA)
```

This function returns the **MatrixType** field of the matrix descriptor **descrA**.

### Input

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

### Returned

	One of the enumerated matrix types.
--	-------------------------------------

## 6.14. `cusparseSetMatDiagType()`

```
cusparseStatus_t  
cusparseSetMatDiagType(cusparseMatDescr_t descrA,  
                      cusparseDiagType_t diagType)
```

This function sets the **DiagType** field of the matrix descriptor **descrA**.

### Input

<b>diagType</b>	One of the enumerated diagType types.
-----------------	---------------------------------------

### Output

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

See `cusparseStatus_t` for the description of the return status

## 6.15. `cusparseSetMatFillMode()`

```
cusparseStatus_t  
cusparseSetMatFillMode(cusparseMatDescr_t descrA,  
                      cusparseFillMode_t fillMode)
```

This function sets the **FillMode** field of the matrix descriptor **descrA**.

### Input

<b>fillMode</b>	One of the enumerated fillMode types.
-----------------	---------------------------------------

### Output

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

See `cusparseStatus_t` for the description of the return status

## 6.16. `cusparseSetMatIndexBase()`

```
cusparseStatus_t
cusparseSetMatIndexBase(cusparseMatDescr_t descrA,
                        cusparseIndexBase_t base)
```

This function sets the **IndexBase** field of the matrix descriptor **descrA**.

### Input

<b>base</b>	One of the enumerated indexBase types.
-------------	--

### Output

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 6.17. `cusparseSetMatType()`

```
cusparseStatus_t
cusparseSetMatType(cusparseMatDescr_t descrA, cusparseMatrixType_t type)
```

This function sets the **MatrixType** field of the matrix descriptor **descrA**.

### Input

<b>type</b>	One of the enumerated matrix types.
-------------	-------------------------------------

### Output

<b>descrA</b>	the matrix descriptor.
---------------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 6.18. `cusparseCreateCsrsv2Info()`

```
cusparseStatus_t
cusparseCreateCsrsv2Info(csrsv2Info_t *info);
```

This function creates and initializes the solve and analysis structure of csrv2 to *default* values.

### Input

<b>info</b>	the pointer to the solve and analysis structure of csrv2.
-------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 6.19. `cusparseDestroyCsrsv2Info()`

```
cusparseStatus_t
cusparseDestroyCsrsv2Info(csrsv2Info_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve ( <code>csrsv2_solve</code> ) and analysis ( <code>csrsv2_analysis</code> ) structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.20. `cusparseCreateCsrsm2Info()`

```
cusparseStatus_t
cusparseCreateCsrsm2Info(csrsm2Info_t *info);
```

This function creates and initializes the solve and analysis structure of csrsm2 to *default* values.

### Input

<code>info</code>	the pointer to the solve and analysis structure of csrsm2.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.21. `cusparseDestroyCsrsm2Info()`

```
cusparseStatus_t
cusparseDestroyCsrsm2Info(csrsm2Info_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve ( <code>csrsm2_solve</code> ) and analysis ( <code>csrsm2_analysis</code> ) structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.22. `cusparseCreateCsric02Info()`

```
cusparseStatus_t
cusparseCreateCsric02Info(csric02Info_t *info);
```

This function creates and initializes the solve and analysis structure of incomplete Cholesky to *default* values.

#### Input

<code>info</code>	the pointer to the solve and analysis structure of incomplete Cholesky.
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.23. `cusparseDestroyCsric02Info()`

```
cusparseStatus_t  
cusparseDestroyCsric02Info(csric02Info_t info);
```

This function destroys and releases any memory required by the structure.

#### Input

<code>info</code>	the solve ( <code>csric02_solve</code> ) and analysis ( <code>csric02_analysis</code> ) structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.24. `cusparseCreateCsrilu02Info()`

```
cusparseStatus_t  
cusparseCreateCsrilu02Info(csrilu02Info_t *info);
```

This function creates and initializes the solve and analysis structure of incomplete LU to *default* values.

#### Input

<code>info</code>	the pointer to the solve and analysis structure of incomplete LU.
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.25. `cusparseDestroyCsrilu02Info()`

```
cusparseStatus_t  
cusparseDestroyCsrilu02Info(csrilu02Info_t info);
```

This function destroys and releases any memory required by the structure.

#### Input

<code>info</code>	the solve ( <code>csrillu02_solve</code> ) and analysis ( <code>csrillu02_analysis</code> ) structure.
-------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 6.26. `cusparseCreateBsrsv2Info()`

```
cusparseStatus_t
cusparseCreateBsrsv2Info(bsrsv2Info_t *info);
```

This function creates and initializes the solve and analysis structure of bsrsv2 to *default* values.

### Input

<code>info</code>	the pointer to the solve and analysis structure of bsrsv2.
-------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 6.27. `cusparseDestroyBsrsv2Info()`

```
cusparseStatus_t
cusparseDestroyBsrsv2Info(bsrsv2Info_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve ( <code>bsrsv2_solve</code> ) and analysis ( <code>bsrsv2_analysis</code> ) structure.
-------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 6.28. `cusparseCreateBsrsm2Info()`

```
cusparseStatus_t
cusparseCreateBsrsm2Info(bsrsm2Info_t *info);
```

This function creates and initializes the solve and analysis structure of bsrsm2 to *default* values.

### Input

<code>info</code>	the pointer to the solve and analysis structure of bsrsm2.
-------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 6.29. `cusparseDestroyBsrsm2Info()`

```
cusparseStatus_t
cusparseDestroyBsrsm2Info(bsrsm2Info_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve ( <code>bsrsm2_solve</code> ) and analysis ( <code>bsrsm2_analysis</code> ) structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.30. `cusparseCreateBsric02Info()`

```
cusparseStatus_t
cusparseCreateBsric02Info(bsric02Info_t *info);
```

This function creates and initializes the solve and analysis structure of block incomplete Cholesky to *default* values.

### Input

<code>info</code>	the pointer to the solve and analysis structure of block incomplete Cholesky.
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 6.31. `cusparseDestroyBsric02Info()`

```
cusparseStatus_t
cusparseDestroyBsric02Info(bsric02Info_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the solve ( <code>bsric02_solve</code> ) and analysis ( <code>bsric02_analysis</code> ) structure.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.32. `cusparseCreateBsrilu02Info()`

```
cusparseStatus_t
cusparseCreateBsrilu02Info(bsrilu02Info_t *info);
```

This function creates and initializes the solve and analysis structure of block incomplete LU to *default* values.

#### Input

<b>info</b>	the pointer to the solve and analysis structure of block incomplete LU.
-------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 6.33. cusparseDestroyBsrilu02Info()

```
cusparseStatus_t  
cusparseDestroyBsrilu02Info(bsrilu02Info_t info);
```

This function destroys and releases any memory required by the structure.

#### Input

<b>info</b>	the solve ( <b>bsrilu02_solve</b> ) and analysis ( <b>bsrilu02_analysis</b> ) structure.
-------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 6.34. cusparseCreateCsrgemm2Info()

```
cusparseStatus_t  
cusparseCreateCsrgemm2Info(csrgemm2Info_t *info);
```

This function creates and initializes analysis structure of general sparse matrix-matrix multiplication.

#### Input

<b>info</b>	the pointer to the analysis structure of general sparse matrix-matrix multiplication.
-------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 6.35. cusparseDestroyCsrgemm2Info()

```
cusparseStatus_t  
cusparseDestroyCsrgemm2Info(csrgemm2Info_t info);
```

This function destroys and releases any memory required by the structure.

#### Input

<b>info</b>	opaque structure of csrgemm2.
-------------	-------------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 6.36. `cusparseCreatePruneInfo()`

```
cusparseStatus_t  
cusparseCreatePruneInfo(pruneInfo_t *info);
```

This function creates and initializes structure of `prune` to *default* values.

### Input

<code>info</code>	the pointer to the structure of <code>prune</code> .
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 6.37. `cusparseDestroyPruneInfo()`

```
cusparseStatus_t  
cusparseDestroyPruneInfo(pruneInfo_t info);
```

This function destroys and releases any memory required by the structure.

### Input

<code>info</code>	the structure of <code>prune</code> .
-------------------	---------------------------------------

See `cusparseStatus_t` for the description of the return status

# Chapter 7.

# CUSPARSE LEVEL 1 FUNCTION REFERENCE

This chapter describes sparse linear algebra functions that perform operations between dense and sparse vectors.

## 7.1. `cusparse<t>axpyi()`

```
cusparseStatus_t
cusparseSaxpyi(cusparseHandle_t    handle,
              int             nnz,
              const float*    alpha,
              const float*    xVal,
              const int*      xInd,
              float*          y,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseDaxpyi(cusparseHandle_t    handle,
               int             nnz,
               const double*   alpha,
               const double*   xVal,
               const int*     xInd,
               double*         y,
               cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseCaxpyi(cusparseHandle_t    handle,
               int             nnz,
               const cuComplex* alpha,
               const cuComplex* xVal,
               const int*      xInd,
               cuComplex*      y,
               cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZaxpyi(cusparseHandle_t    handle,
               int             nnz,
               const cuDoubleComplex* alpha,
               const cuDoubleComplex* xVal,
               const int*      xInd,
               cuDoubleComplex* y,
               cusparseIndexBase_t idxBase)
```

This function multiplies the vector **x** in sparse format by the constant  $\alpha$  and adds the result to the vector **y** in dense format. This operation can be written as

$$y = y + \alpha * x$$

In other words,

```
for i=0 to nnz-1
    y[xInd[i]-idxBase] = y[xInd[i]-idxBase] + alpha*xVal[i]
```

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values of vector <b>x</b> .
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>y</b>	<type> vector in dense format.
<b>idxBase</b>	CUSPARSE_INDEX_BASE_ZERO or CUSPARSE_INDEX_BASE_ONE.

### Output

<b>y</b>	<type> updated vector in dense format (that is unchanged if <b>nnz == 0</b> ).
----------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 7.2. `cusparse<t>doti()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseSpVV()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseSdoti(cusparseHandle_t      handle,
              int                 nnz,
              const float*        xVal,
              const int*          xInd,
              const float*        y,
              float*              resultDevHostPtr,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseDdoti(cusparseHandle_t      handle,
              int                 nnz,
              const double*       xVal,
              const int*          xInd,
              const double*       y,
              double*             resultDevHostPtr,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseCdoti(cusparseHandle_t      handle,
              int                 nnz,
              const cuComplex*   xVal,
              const int*          xInd,
              const cuComplex*   y,
              cuComplex*          resultDevHostPtr,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZdoti(cusparseHandle_t      handle,
              int                 nnz,
              const cuDoubleComplex* xVal,
              const int*          xInd,
              const cuDoubleComplex* y,
              cuDoubleComplex*    resultDevHostPtr,
              cusparseIndexBase_t idxBase)

```

This function returns the dot product of a vector **x** in sparse format and vector **y** in dense format. This operation can be written as

$$\text{result} = \mathbf{y}^T \mathbf{x}$$

In other words,

```
for i=0 to nnz-1
    resultDevHostPtr += xVal[i]*y[xInd[i-idxBase]]
```

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values of vector <b>x</b> .
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>y</b>	<type> vector in dense format.
<b>resultDevHostPtr</b>	pointer to the location of the result in the device or host memory.
<b>idxBase</b>	CUSPARSE_INDEX_BASE_ZERO or CUSPARSE_INDEX_BASE_ONE.

## Output

<b>resultDevHostPtr</b>	scalar result in the device or host memory (that is zero if <b>nnz == 0</b> ).
-------------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 7.3. `cusparse<t>dotci()` [DEPRECATED]

[[DEPRECATED]] use [`cusparseSpVV\(\)`](#) instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseCdotci(cusparseHandle_t      handle,
              int                  nnz,
              const cuComplex*    xVal,
              const int*          xInd,
              const cuComplex*    y,
              cuComplex*          resultDevHostPtr,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZdotci(cusparseHandle_t      handle,
               int                  nnz,
               const cuDoubleComplex* xVal,
               const int*            xInd,
               const cuDoubleComplex* y,
               cuDoubleComplex*     resultDevHostPtr,
               cusparseIndexBase_t  idxBase)

```

This function returns the dot product of a complex conjugate of vector **x** in sparse format and vector **y** in dense format. This operation can be written as

$$\text{result} = \mathbf{x}^H \mathbf{y}$$

In other words,

```

for i=0 to nnz-1
  resultDevHostPtr += \xVal[i]*y[xInd[i-idxBase]]

```

- ▶ This function requires temporary extra storage that is allocated internally

- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values of vector <b>x</b> .
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>y</b>	<type> vector in dense format.
<b>resultDevHostPtr</b>	pointer to the location of the result in the device or host memory.
<b>idxBase</b>	CUSPARSE_INDEX_BASE_ZERO or CUSPARSE_INDEX_BASE_ONE.

### Output

<b>resultDevHostPtr</b>	scalar result in the device or host memory (that is zero if <b>nnz == 0</b> ).
-------------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 7.4. `cusparse<t>gthr()`

```

cusparseStatus_t
cusparseSgthr(cusparseHandle_t      handle,
              int                  nnz,
              const float*          y,
              float*                xVal,
              const int*             xInd,
              cusparseIndexBase_t   idxBase)

cusparseStatus_t
cusparseDgthr(cusparseHandle_t      handle,
              int                  nnz,
              const double*         y,
              double*               xVal,
              const int*             xInd,
              cusparseIndexBase_t   idxBase)

cusparseStatus_t
cusparseCgthr(cusparseHandle_t      handle,
              int                  nnz,
              const cuComplex*      y,
              cuComplex*             xVal,
              const int*             xInd,
              cusparseIndexBase_t   idxBase)

cusparseStatus_t
cusparseZgthr(cusparseHandle_t      handle,
              int                  nnz,
              const cuDoubleComplex* y,
              cuDoubleComplex*       xVal,
              const int*             xInd,
              cusparseIndexBase_t   idxBase)

```

This function gathers the elements of the vector **y** listed in the index array **xInd** into the data array **xVal**.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>y</b>	<type> vector in dense format (of <code>size≥max(xInd)-idxBase+1</code> ).
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>idxBase</b>	<code>CUSPARSE_INDEX_BASE_ZERO</code> or <code>CUSPARSE_INDEX_BASE_ONE</code> .

### Output

<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values that were gathered from vector <b>y</b> (that is unchanged if <b>nnz == 0</b> ).
-------------	---

See `cusparseStatus_t` for the description of the return status

## 7.5. `cusparse<t>gthrz()`

```

cusparseStatus_t
cusparseSgthrz(cusparseHandle_t handle,
               int nnz,
               float* y,
               float* xVal,
               const int* xInd,
               cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseDgthrz(cusparseHandle_t handle,
               int nnz,
               double* y,
               double* xVal,
               const int* xInd,
               cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseCgthrz(cusparseHandle_t handle,
               int nnz,
               cuComplex* y,
               cuComplex* xVal,
               const int* xInd,
               cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZgthrz(cusparseHandle_t handle,
               int nnz,
               cuDoubleComplex* y,
               cuDoubleComplex* xVal,
               const int* xInd,
               cusparseIndexBase_t idxBase)

```

This function gathers the elements of the vector **y** listed in the index array **xInd** into the data array **xVal**. Also, it zeros out the gathered elements in the vector **y**.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>y</b>	<type> vector in dense format (of <b>size≥max(xInd)-idxBase+1</b> ).

<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>idxBase</b>	<b>CUSPARSE_INDEX_BASE_ZERO</b> or <b>CUSPARSE_INDEX_BASE_ONE</b> .

**Output**

<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values that were gathered from vector <b>y</b> (that is unchanged if <b>nnz == 0</b> ).
<b>y</b>	<type> vector in dense format with elements indexed by <b>xInd</b> set to zero (it is unchanged if <b>nnz == 0</b> ).

See `cusparseStatus_t` for the description of the return status

## 7.6. `cusparse<t>roti()`

```

cusparseStatus_t
cusparseSroti(cusparseHandle_t      handle,
              int                  nnz,
              float*               xVal,
              const int*           xInd,
              float*               y,
              const float*         c,
              const float*         s,
              cusparseIndexBase_t  idxBase)

cusparseStatus_t
cusparseDroti(cusparseHandle_t      handle,
              int                  nnz,
              double*              xVal,
              const int*           xInd,
              double*              y,
              const double*        c,
              const double*        s,
              cusparseIndexBase_t  idxBase)

```

This function applies the Givens rotation matrix

$$G = \begin{pmatrix} c & s \\ -s & c \end{pmatrix}$$

to sparse **x** and dense **y** vectors. In other words,

```

for i=0 to nnz-1
    y[xInd[i]-idxBase] = c * y[xInd[i]-idxBase] - s*xVal[i]
    x[i]                = c * xVal[i]           + s * y[xInd[i]-idxBase]

```

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

**Input**

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values of vector <b>x</b> .
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>y</b>	<type> vector in dense format.
<b>c</b>	cosine element of the rotation matrix.
<b>s</b>	sine element of the rotation matrix.
<b>idxBase</b>	CUSPARSE_INDEX_BASE_ZERO or CUSPARSE_INDEX_BASE_ONE.

## Output

<b>xVal</b>	<type> updated vector in sparse format (that is unchanged if <b>nnz == 0</b> ).
<b>y</b>	<type> updated vector in dense format (that is unchanged if <b>nnz == 0</b> ).

See [cusparseStatus\\_t](#) for the description of the return status

## 7.7. `cusparse<t>sctr()`

```

cusparseStatus_t
cusparseSsctr(cusparseHandle_t      handle,
              int                 nnz,
              const float*        xVal,
              const int*          xInd,
              float*              y,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseDsctr(cusparseHandle_t      handle,
              int                 nnz,
              const double*       xVal,
              const int*          xInd,
              double*             y,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseCsctr(cusparseHandle_t      handle,
              int                 nnz,
              const cuComplex*   xVal,
              const int*          xInd,
              cuComplex*          y,
              cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZsctr(cusparseHandle_t      handle,
              int                 nnz,
              const cuDoubleComplex* xVal,
              const int*          xInd,
              cuDoubleComplex*    y,
              cusparseIndexBase_t idxBase)

```

This function scatters the elements of the vector **x** in sparse format into the vector **y** in dense format. It modifies only the elements of **y** whose indices are listed in the array **xInd**.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>nnz</b>	number of elements in vector <b>x</b> .
<b>xVal</b>	<type> vector with <b>nnz</b> nonzero values of vector <b>x</b> .
<b>xInd</b>	integer vector with <b>nnz</b> indices of the nonzero values of vector <b>x</b> .
<b>y</b>	<type> dense vector (of <b>size≥max(xInd) - idxBase+1</b> ).

<b>idxBase</b>	CUSPARSE_INDEX_BASE_ZERO or CUSPARSE_INDEX_BASE_ONE.
----------------	---

## Output

<b>y</b>	<type> vector with <b>nnz</b> nonzero values that were scattered from vector <b>x</b> (that is unchanged if <b>nnz == 0</b> ).
----------	--

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 8.

# CUSPARSE LEVEL 2 FUNCTION REFERENCE

This chapter describes the sparse linear algebra functions that perform operations between sparse matrices and dense vectors.

In particular, the solution of sparse triangular linear systems is implemented in two phases. First, during the analysis phase, the sparse triangular matrix is analyzed to determine the dependencies between its elements by calling the appropriate `csrsv_analysis()` function. The analysis is specific to the sparsity pattern of the given matrix and to the selected `cusparseOperation_t` type. The information from the analysis phase is stored in the parameter of type `cusparseSolveAnalysisInfo_t` that has been initialized previously with a call to `cusparseCreateSolveAnalysisInfo()`.

Second, during the solve phase, the given sparse triangular linear system is solved using the information stored in the `cusparseSolveAnalysisInfo_t` parameter by calling the appropriate `csrsv_solve()` function. The solve phase may be performed multiple times with different right-hand sides, while the analysis phase needs to be performed only once. This is especially useful when a sparse triangular linear system must be solved for a set of different right-hand sides one at a time, while its coefficient matrix remains the same.

Finally, once all the solves have completed, the opaque data structure pointed to by the `cusparseSolveAnalysisInfo_t` parameter can be released by calling `cusparseDestroySolveAnalysisInfo()`. For more information please refer to [3].

## 8.1. `cusparse<t>bsrmv()`

```

cusparseStatus_t
cusparseSbsrmv(cusparseHandle_t          handle,
               cusparseDirection_t    dir,
               cusparseOperation_t   trans,
               int                  mb,
               int                  nb,
               int                  nnzb,
               const float*          alpha,
               const cusparseMatDescr_t descr,
               const float*          bsrVal,
               const int*            bsrRowPtr,
               const int*            bsrColInd,
               int                  blockDim,
               const float*          x,
               const float*          beta,
               float*                y)

cusparseStatus_t
cusparseDbsrmv(cusparseHandle_t          handle,
               cusparseDirection_t    dir,
               cusparseOperation_t   trans,
               int                  mb,
               int                  nb,
               int                  nnzb,
               const double*         alpha,
               const cusparseMatDescr_t descr,
               const double*         bsrVal,
               const int*            bsrRowPtr,
               const int*            bsrColInd,
               int                  blockDim,
               const double*         x,
               const double*         beta,
               double*               y)

cusparseStatus_t
cusparseCbsrmv(cusparseHandle_t          handle,
               cusparseDirection_t    dir,
               cusparseOperation_t   trans,
               int                  mb,
               int                  nb,
               int                  nnzb,
               const cuComplex*      alpha,
               const cusparseMatDescr_t descr,
               const cuComplex*      bsrVal,
               const int*            bsrRowPtr,
               const int*            bsrColInd,
               int                  blockDim,
               const cuComplex*      x,
               const cuComplex*      beta,
               cuComplex*             y)

cusparseStatus_t
cusparseZbsrmv(cusparseHandle_t          handle,
               cusparseDirection_t    dir,
               cusparseOperation_t   trans,
               int                  mb,
               int                  nb,
               int                  nnzb,
               const cuDoubleComplex* alpha,
               const cusparseMatDescr_t descr,
               const cuDoubleComplex* bsrVal,
               const int*            bsrRowPtr,
               const int*            bsrColInd.

```

This function performs the matrix-vector operation

$$y = \alpha * \text{op}(A) * x + \beta * y$$

where  $A$  is an  $(mb * blockDim) \times (nb * blockDim)$  sparse matrix that is defined in BSR storage format by the three arrays **bsrVal**, **bsrRowPtr**, and **bsrColInd**; **x** and **y** are vectors;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

Several comments on **bsrmv()**:

- ▶ Only **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** is supported, that is

$$y = \alpha * A * x + \beta * y$$

- ▶ Only **CUSPARSE\_MATRIX\_TYPE\_GENERAL** is supported.
- ▶ The size of vector **x** should be  $(nb * blockDim)$  at least, and the size of vector **y** should be  $(mb * blockDim)$  at least; otherwise, the kernel may return **CUSPARSE\_STATUS\_EXECUTION\_FAILED** because of an out-of-bounds array.

For example, suppose the user has a CSR format and wants to try **bsrmv()**, the following code demonstrates how to use **csr2bsr()** conversion and **bsrmv()** multiplication in single precision.

```
// Suppose that A is m x n sparse matrix represented by CSR format,
// hx is a host vector of size n, and hy is also a host vector of size m.
// m and n are not multiple of blockDim.
// step 1: transform CSR to BSR with column-major order
int base, nnz;
int nnzb;
cuspDirection_t dirA = CUSPARSE_DIRECTION_COLUMN;
int mb = (m + blockDim-1)/blockDim;
int nb = (n + blockDim-1)/blockDim;
cudaMalloc((void**)&bsrRowPtrC, sizeof(int) * (mb+1));
cuspseXcsr2bsrNnz(handle, dirA, m, n,
    descrA, csrRowPtrA, csrColIndA, blockDim,
    descrC, bsrRowPtrC, &nnzb);
cudaMalloc((void**)&bsrColIndC, sizeof(int)*nnzb);
cudaMalloc((void**)&bsrValC, sizeof(float)*(blockDim*blockDim)*nnzb);
cuspseScsr2bsr(handle, dirA, m, n,
    descrA, csrValA, csrRowPtrA, csrColIndA, blockDim,
    descrC, bsrValC, bsrRowPtrC, bsrColIndC);
// step 2: allocate vector x and vector y large enough for bsrmv
cudaMalloc((void**)&x, sizeof(float)*(nb*blockDim));
cudaMalloc((void**)&y, sizeof(float)*(mb*blockDim));
cudaMemcpy(x, hx, sizeof(float)*n, cudaMemcpyHostToDevice);
cudaMemcpy(y, hy, sizeof(float)*m, cudaMemcpyHostToDevice);
// step 3: perform bsrmv
cuspseSbsrmv(handle, dirA, transA, mb, nb, nnzb, &alpha,
    descrC, bsrValC, bsrRowPtrC, bsrColIndC, blockDim, x, &beta, y);
```

## Input

<b>handle</b>	handle to the cuSPARSE library context.
---------------	---

<b>dir</b>	storage format of blocks, either <code>CUSPARSE_DIRECTION_ROW</code> or <code>CUSPARSE_DIRECTION_COLUMN</code> .
<b>trans</b>	the operation $\text{op}(A)$ . Only <code>CUSPARSE_OPERATION_NON_TRANSPOSE</code> is supported.
<b>mb</b>	number of block rows of matrix $A$ .
<b>nb</b>	number of block columns of matrix $A$ .
<b>nnzb</b>	number of nonzero blocks of matrix $A$ .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descr</b>	the descriptor of matrix $A$ . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<b>bsrVal</b>	<type> array of $\text{nnz}$ (= <code>csrRowPtrA(mb) - csrRowPtrA(0)</code> ) nonzero blocks of matrix $A$ .
<b>bsrRowPtr</b>	integer array of $\text{mb} + 1$ elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColInd</b>	integer array of $\text{nnz}$ (= <code>csrRowPtrA(mb) - csrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix $A$ .
<b>blockDim</b>	block dimension of sparse matrix $A$ , larger than zero.
<b>x</b>	<type> vector of $nb * \text{blockDim}$ elements.
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>y</b>	<type> vector of $mb * \text{blockDim}$ elements.

## Output

<b>y</b>	<type> updated vector.
----------	------------------------

See `cusparseStatus_t` for the description of the return status

## 8.2. `cusparse<t>bsrxmlv()`

```

cusparseStatus_t
cusparseSbsrxmlv(cusparseHandle_t
                  cusparseDirection_t
                  cusparseOperation_t
                  int
                  int
                  int
                  int
                  const float*
                  const cusparseMatDescr_t
                  const float*
                  const int*
                  const int*
                  const int*
                  const int*
                  int
                  const float*
                  const float*
                  float*)
                  handle,
                  dir,
                  trans,
                  sizeOfMask,
                  mb,
                  nb,
                  nnzb,
                  alpha,
                  descr,
                  bsrVal,
                  bsrMaskPtr,
                  bsrRowPtr,
                  bsrEndPtr,
                  bsrColInd,
                  blockDim,
                  x,
                  beta,
                  y)

cusparseStatus_t
cusparseDbsrxmlv(cusparseHandle_t
                  cusparseDirection_t
                  cusparseOperation_t
                  int
                  int
                  int
                  int
                  const double*
                  const cusparseMatDescr_t
                  const double*
                  const int*
                  const int*
                  const int*
                  const int*
                  int
                  const double*
                  const double*
                  double*)
                  handle,
                  dir,
                  trans,
                  sizeOfMask,
                  mb,
                  nb,
                  nnzb,
                  alpha,
                  descr,
                  bsrVal,
                  bsrMaskPtr,
                  bsrRowPtr,
                  bsrEndPtr,
                  bsrColInd,
                  blockDim,
                  x,
                  beta,
                  y)

cusparseStatus_t
cusparseCbsrxmlv(cusparseHandle_t
                  cusparseDirection_t
                  cusparseOperation_t
                  int
                  int
                  int
                  int
                  const cuComplex*
                  const cusparseMatDescr_t
                  const cuComplex*
                  const int*
                  const int*
                  const int*
                  const int*
                  int
                  const cuComplex*
                  const cuComplex*
                  cuComplex*)
                  handle,
                  dir,
                  trans,
                  sizeOfMask,
                  mb,
                  nb,
                  nnzb,
                  alpha,
                  descr,
                  bsrVal,
                  bsrMaskPtr,
                  bsrRowPtr,
                  bsrEndPtr,
                  bsrColInd,
                  blockDim,
                  x,
                  beta,
                  y)

cusparseStatus_t
cusparseZbsrxmlv(cusparseHandle_t
                  cusparseDirection_t
                  handle,
                  dir,

```

This function performs a **bsrmv** and a mask operation

$$y(\text{mask}) = (\alpha * \text{op}(A) * x + \beta * y)(\text{mask})$$

where  $A$  is an  $(mb * blockDim) \times (nb * blockDim)$  sparse matrix that is defined in BSRX storage format by the four arrays **bsrVal**, **bsrRowPtr**, **bsrEndPtr**, and **bsrColInd**; **x** and **y** are vectors;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

The mask operation is defined by array **bsrMaskPtr** which contains updated block row indices of  $y$ . If row  $i$  is not specified in **bsrMaskPtr**, then **bsrxmv()** does not touch row block  $i$  of  $A$  and  $y$ .

For example, consider the  $2 \times 3$  block matrix  $A$ :

$$A = \begin{bmatrix} A_{11} & A_{12} & O \\ A_{21} & A_{22} & A_{23} \end{bmatrix}$$

and its one-based BSR format (three vector form) is

$$\begin{aligned} \text{bsrVal} &= [A_{11} \ A_{12} \ A_{21} \ A_{22} \ A_{23}] \\ \text{bsrRowPtr} &= [1 \ 3 \ 6] \\ \text{bsrColInd} &= [1 \ 2 \ 1 \ 2 \ 3] \end{aligned}$$

Suppose we want to do the following **bsrmv** operation on a matrix  $\bar{A}$  which is slightly different from  $A$ .

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} := \text{alpha} * \left( \tilde{A} = \begin{bmatrix} O & O & O \\ O & A_{22} & O \end{bmatrix} \right) * \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} y_1 \\ \text{beta} * y_2 \end{bmatrix}$$

We don't need to create another BSR format for the new matrix  $\bar{A}$ , all that we should do is to keep **bsrVal** and **bsrColInd** unchanged, but modify **bsrRowPtr** and add an additional array **bsrEndPtr** which points to the last nonzero elements per row of  $\bar{A}$  plus 1.

For example, the following **bsrRowPtr** and **bsrEndPtr** can represent matrix  $\bar{A}$ :

$$\begin{aligned} \text{bsrRowPtr} &= [1 \ 4] \\ \text{bsrEndPtr} &= [1 \ 5] \end{aligned}$$

Further we can use a mask operator (specified by array **bsrMaskPtr**) to update particular block row indices of  $y$  only because  $y_1$  is never changed. In this case, **bsrMaskPtr** = [2] and **sizeOfMask**=1.

The mask operator is equivalent to the following operation:

$$\begin{bmatrix} ? \\ y_2 \end{bmatrix} := \text{alpha} * \begin{bmatrix} ? & ? & ? \\ O & A_{22} & O \end{bmatrix} * \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \text{beta} * \begin{bmatrix} ? \\ y_2 \end{bmatrix}$$

If a block row is not present in the **bsrMaskPtr**, then no calculation is performed on that row, and the corresponding value in **y** is unmodified. The question mark "?" is used to indicate row blocks not in **bsrMaskPtr**.

In this case, first row block is not present in **bsrMaskPtr**, so **bsrRowPtr[0]** and **bsrEndPtr[0]** are not touched also.

$$\begin{aligned} \text{bsrRowPtr} &= [? \quad 4] \\ \text{bsrEndPtr} &= [? \quad 5] \end{aligned}$$

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

A couple of comments on **bsrxmv()**:

- ▶ Only **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and **CUSPARSE\_MATRIX\_TYPE\_GENERAL** are supported.
- ▶ Parameters **bsrMaskPtr**, **bsrRowPtr**, **bsrEndPtr** and **bsrColInd** are consistent with base index, either one-based or zero-based. The above example is one-based.

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dir</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>trans</b>	the operation $\text{op}(A)$ . Only <b>CUSPARSE_OPERATION_NON_TRANSPOSE</b> is supported.
<b>sizeOfMask</b>	number of updated block rows of <b>y</b> .
<b>mb</b>	number of block rows of matrix <b>A</b> .
<b>nb</b>	number of block columns of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descr</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrVal</b>	<type> array of <b>nnz</b> nonzero blocks of matrix <b>A</b> .
<b>bsrMaskPtr</b>	integer array of <b>sizeOfMask</b> elements that contains the indices corresponding to updated block rows.
<b>bsrRowPtr</b>	integer array of <b>mb</b> elements that contains the start of every block row.
<b>bsrEndPtr</b>	integer array of <b>mb</b> elements that contains the end of every block row plus one.

<b>bsrColInd</b>	integer array of <code>nnzb</code> column indices of the nonzero blocks of matrix $A$ .
<b>blockDim</b>	block dimension of sparse matrix $A$ , larger than zero.
<b>x</b>	<type> vector of $nb * blockDim$ elements.
<b>beta</b>	<type> scalar used for multiplication. If <code>beta</code> is zero, <code>y</code> does not have to be a valid input.
<b>y</b>	<type> vector of $mb * blockDim$ elements.

See `cusparseStatus_t` for the description of the return status

## 8.3. `cusparse<t>bsrsv2_bufferSize()`

```

cusparseStatus_t
cusparseSbsrsv2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           cusparseOperation_t   transA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           float*                bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                  blockDim,
                           bsrv2Info_t            info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseDbsrsv2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           cusparseOperation_t   transA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           double*               bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                  blockDim,
                           bsrv2Info_t            info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseCbsrsv2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           cusparseOperation_t   transA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuComplex*             bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                  blockDim,
                           bsrv2Info_t            info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseZbsrsv2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           cusparseOperation_t   transA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*        bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                  blockDim,
                           bsrv2Info_t            info,
                           int*                 pBufferSizeInBytes)

```

This function returns size of the buffer used in **bsrsv2**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) * \mathbf{y} = \alpha \mathbf{x}$ .

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

Although there are six combinations in terms of parameter **trans** and the upper (lower) triangular part of **A**, **bsrsv2\_bufferSize()** returns the maximum size buffer among these combinations. The buffer size depends on the dimensions **mb**, **blockDim**, and the number of nonzero blocks of the matrix **nnzb**. If the user changes the matrix, it is necessary to call **bsrsv2\_bufferSize()** again to have the correct buffer size; otherwise a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation $\text{op}(A)$ .
<b>mb</b>	number of block rows of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) nonzero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb</b> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> ; must be larger than zero.

## Output

<b>info</b>	record of internal states based on different algorithms.
-------------	--

<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in the <code>bsrsv2_analysis()</code> and <code>bsrsv2_solve()</code> .
---------------------------	---

See `cusparseStatus_t` for the description of the return status

## 8.4. `cusparse<t>bsrsv2_analysis()`

```

cusparseStatus_t
cusparseSbsrsv2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        int                  mb,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const float*          bsrValA,
                        const int*            bsrRowPtrA,
                        const int*            bsrColIndA,
                        int                  blockDim,
                        bsrv2Info_t           info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseDbsrsv2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        int                  mb,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const double*         bsrValA,
                        const int*            bsrRowPtrA,
                        const int*            bsrColIndA,
                        int                  blockDim,
                        bsrv2Info_t           info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseDbsrsv2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        int                  mb,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const cuComplex*      bsrValA,
                        const int*            bsrRowPtrA,
                        const int*            bsrColIndA,
                        int                  blockDim,
                        bsrv2Info_t           info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseZbsrsv2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        int                  mb,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const cuDoubleComplex* bsrValA,
                        const int*            bsrRowPtrA,
                        const int*            bsrColIndA,
                        int                  blockDim,
                        bsrv2Info_t           info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

```

This function performs the analysis phase of **bsrsv2**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) \mathbf{y} = \alpha \mathbf{x}$ .

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **x** and **y** are the right-hand side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

The block of BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored.

It is expected that this function will be executed only once for a given matrix and a particular operation type.

This function requires a buffer size returned by **bsrsv2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsrsv2\_analysis()** reports a structural zero and computes level information, which is stored in the opaque structure **info**. The level information can extract more parallelism for a triangular solver. However **bsrsv2\_solve()** can be done without level information. To disable level information, the user needs to specify the policy of the triangular solver as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **bsrsv2\_analysis()** always reports the first structural zero, even when parameter **policy** is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. No structural zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if block **A(j, j)** is missing for some **j**. The user needs to call **cusparseXbsrsv2\_zeroPivot()** to know where the structural zero is.

It is the user's choice whether to call **bsrsv2\_solve()** if **bsrsv2\_analysis()** reports a structural zero. In this case, the user can still call **bsrsv2\_solve()**, which will return a numerical zero at the same position as a structural zero. However the result **x** is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation $\text{op}(A)$ .
<b>mb</b>	number of block rows of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .

<b>descrA</b>	the descriptor of matrix A. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix A.
<b>bsrRowPtrA</b>	integer array of <b>mb + 1</b> elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) column indices of the nonzero blocks of matrix A.
<b>blockDim</b>	block dimension of sparse matrix A, larger than zero.
<b>info</b>	structure initialized using <b>cusparseCreateBsrsv2Info()</b> .
<b>policy</b>	the supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is return by <b>bsrsv2_bufferSize()</b> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See **cusparseStatus\_t** for the description of the return status

## 8.5. `cusparse<t>bsrsv2_solve()`

This function performs the solve phase of **bsrsv2**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) * \mathbf{y} = \alpha * \mathbf{x}$ .

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

The block in BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored. Function **bsrsv02\_solve()** can support an arbitrary **blockDim**.

This function may be executed multiple times for a given matrix and a particular operation type.

This function requires a buffer size returned by **bsrsv2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **bsrsv2\_solve()** can be done without level information, the user still needs to be aware of consistency. If **bsrsv2\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **bsrsv2\_solve()** can be run with or without levels. On the other hand, if **bsrsv2\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **bsrsv2\_solve()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

The level information may not improve the performance, but may spend extra time doing analysis. For example, a tridiagonal matrix has no parallelism.

In this case, **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** performs better than **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**. If the user has an iterative solver, the best approach is to do **bsrsv2\_analysis()** with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** once. Then do **bsrsv2\_solve()** with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** in the first run, and with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** in the second run, and pick the fastest one to perform the remaining iterations.

Function **bsrsv02\_solve()** has the same behavior as **csrsv02\_solve()**. That is, **bsr2csr(bsrsv02(A)) = csrsv02(bsr2csr(A))**. The numerical zero of **csrsv02\_solve()** means there exists some zero **A(j, j)**. The numerical zero of **bsrsv02\_solve()** means there exists some block **A(j, j)** that is not invertible.

Function **bsrsv2\_solve()** reports the first numerical zero, including a structural zero. No numerical zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if **A(j, j)** is not invertible for some **j**. The user needs to call **cusparseXbsrsv2\_zeroPivot()** to know where the numerical zero is.

The function supports the following properties if **pBuffer != NULL**

- ▶ The routine requires no extra storage

- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

For example, suppose L is a lower triangular matrix with unit diagonal, then the following code solves  $L \cdot y = x$  by level information.

```

// Suppose that L is m x m sparse matrix represented by BSR format,
// The number of block rows/columns is mb, and
// the number of nonzero blocks is nnzb.
// L is lower triangular with unit diagonal.
// Assumption:
// - dimension of matrix L is m(=mb*blockDim),
// - matrix L has nnz(=nnzb*blockDim*blockDim) nonzero elements,
// - handle is already created by cusparseCreate(),
// - (d_bsrRowPtr, d_bsrColInd, d_bsrVal) is BSR of L on device memory,
// - d_x is right hand side vector on device memory.
// - d_y is solution vector on device memory.
// - d_x and d_y are of size m.
cusparseMatDescr_t descr = 0;
bsrsv2Info_t info = 0;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.;
const cusparseSolvePolicy_t policy = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;

// step 1: create a descriptor which contains
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has unit diagonal, specified by parameter CUSPARSE_DIAG_TYPE_UNIT
//   (L may not have all diagonal elements.)
cusparseCreateMatDescr(&descr);
cusparseSetMatIndexBase(descr, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatFillMode(descr, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr, CUSPARSE_DIAG_TYPE_UNIT);

// step 2: create a empty info structure
cusparseCreateBsrsv2Info(&info);

// step 3: query how much memory used in bsrsv2, and allocate the buffer
cusparseDbsrsv2_bufferSize(handle, dir, trans, mb, nnzb, descr,
    d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, &pBufferSize);

// pBuffer returned by cudaMalloc is automatically aligned to 128 bytes.
cudaMalloc((void**)&pBuffer, pBufferSize);

// step 4: perform analysis
cusparseDbsrsv2_analysis(handle, dir, trans, mb, nnzb, descr,
    d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim,
    info, policy, pBuffer);
// L has unit diagonal, so no structural zero is reported.
status = cusparseXbsrsv2_zeroPivot(handle, info, &structural_zero);
if (CUSPARSE_STATUS_ZERO_PIVOT == status){
    printf("L(%d,%d) is missing\n", structural_zero, structural_zero);
}

// step 5: solve L*y = x
cusparseDbsrsv2_solve(handle, dir, trans, mb, nnzb, &alpha, descr,
    d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info,
    d_x, d_y, policy, pBuffer);
// L has unit diagonal, so no numerical zero is reported.
status = cusparseXbsrsv2_zeroPivot(handle, info, &numerical_zero);
if (CUSPARSE_STATUS_ZERO_PIVOT == status){
    printf("L(%d,%d) is zero\n", numerical_zero, numerical_zero);
}

// step 6: free resources
cudaFree(pBuffer);
cusparseDestroyBsrsv2Info(info);
cusparseDestroyMatDescr(descr);
cusparseDestroy(handle);

```

**Input**

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation $\text{op}(A)$ .
<b>mb</b>	number of block rows and block columns of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb + 1</b> elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> , larger than zero.
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<b>x</b>	<type> right-hand-side vector of size <b>m</b> .
<b>policy</b>	the supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is returned by <b>bsrsv2_bufferSize()</b> .

**Output**

<b>y</b>	<type> solution vector of size <b>m</b> .
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 8.6. `cusparseXbsrsv2_zeroPivot()`

```
cusparseStatus_t
cusparseXbsrsv2_zeroPivot(cusparseHandle_t handle,
                           bsrsv2Info_t      info,
                           int*              position)
```

If the returned error code is **CUSPARSE\_STATUS\_ZERO\_PIVOT**, **position=j** means  $A(j,j)$  is either structural zero or numerical zero (singular block). Otherwise **position=-1**.

The **position** can be 0-based or 1-based, the same as the matrix.

Function **`cusparseXbsrv2_zeroPivot()`** is a blocking call. It calls **`cudaDeviceSynchronize()`** to make sure all previous kernels are done.

The **position** can be in the host memory or device memory. The user can set the proper mode with **`cusparseSetPointerMode()`**.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>info</b>	<b>info</b> contains a structural zero or numerical zero if the user already called <b><code>bsrv2_analysis()</code></b> or <b><code>bsrv2_solve()</code></b> .

### Output

<b>position</b>	if no structural or numerical zero, <b>position</b> is -1; otherwise if $A(j,j)$ is missing or $U(j,j)$ is zero, <b>position=j</b> .
-----------------	--

See **`cusparseStatus_t`** for the description of the return status

## 8.7. `cusparse<t>csrmv()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseSpMV\(\)` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrmv(cusparseHandle_t          handle,
               cusparseOperation_t    transA,
               int                   m,
               int                   n,
               int                   nnz,
               const float*          alpha,
               const cusparseMatDescr_t descrA,
               const float*          csrValA,
               const int*            csrRowPtrA,
               const int*            csrColIndA,
               const float*          x,
               const float*          beta,
               float*                y)

cusparseStatus_t
cusparseDcsrsv(cusparseHandle_t          handle,
               cusparseOperation_t    transA,
               int                   m,
               int                   n,
               int                   nnz,
               const double*         alpha,
               const cusparseMatDescr_t descrA,
               const double*         csrValA,
               const int*            csrRowPtrA,
               const int*            csrColIndA,
               const double*         x,
               const double*         beta,
               double*               y)

cusparseStatus_t
cusparseCcsrsv(cusparseHandle_t          handle,
               cusparseOperation_t    transA,
               int                   m,
               int                   n,
               int                   nnz,
               const cuComplex*      alpha,
               const cusparseMatDescr_t descrA,
               const cuComplex*      csrValA,
               const int*            csrRowPtrA,
               const int*            csrColIndA,
               const cuComplex*      x,
               const cuComplex*      beta,
               cuComplex*            y)

cusparseStatus_t
cusparseZcsrsv(cusparseHandle_t          handle,
               cusparseOperation_t    transA,
               int                   m,
               int                   n,
               int                   nnz,
               const cuDoubleComplex* alpha,
               const cusparseMatDescr_t descrA,
               const cuDoubleComplex* csrValA,
               const int*            csrRowPtrA,
               const int*            csrColIndA,
               const cuDoubleComplex* x,
               const cuDoubleComplex* beta,
               cuDoubleComplex*      y)

```

This function performs the matrix-vector operation

$$y = \alpha * \text{op}(A) * x + \beta * y$$

**A** is an  $m \times n$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **x** and **y** are vectors;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

When using the (conjugate) transpose of a general matrix or a Hermitian/symmetric matrix, this routine may produce slightly different results during different runs with the same input parameters. For these matrix types it uses atomic operations to compute the final result, consequently many threads may be adding floating point numbers to the same memory location without any specific ordering, which may produce slightly different results for each run.

If exactly the same output is required for any input when multiplying by the transpose of a general matrix, the following procedure can be used:

1. Convert the matrix from CSR to CSC format using one of the **csr2csc()** functions. Notice that by interchanging the rows and columns of the result you are implicitly transposing the matrix.
2. Call the **csrmv()** function with the **cusparseOperation\_t** parameter set to **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and with the interchanged rows and columns of the matrix stored in CSC format. This (implicitly) multiplies the vector by the transpose of the matrix in the original CSR format.

The function has the following properties when operation is different from **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** or matrix type is not **CUSPARSE\_MATRIX\_TYPE\_GENERAL**:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does **not** support asynchronous execution
- ▶ The routine does **not** support CUDA graph capture

otherwise:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation $\text{op}(A)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.

<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , <b>CUSPARSE_MATRIX_TYPE_SYMMETRIC</b> , and <b>CUSPARSE_MATRIX_TYPE_HERMITIAN</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>x</b>	<type> vector of <b>n</b> elements if $\text{op}(A)=A$ , and <b>m</b> elements if $\text{op}(A)=A^T$ or $\text{op}(A)=A^H$
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>y</b>	<type> vector of <b>m</b> elements if $\text{op}(A)=A$ , and <b>n</b> elements if $\text{op}(A)=A^T$ or $\text{op}(A)=A^H$

## Output

<b>y</b>	<type> updated vector.
----------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 8.8. `cusparse<t>csrmv_mp()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseCsrmvEx()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrmv_mp(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  int                   m,
                  int                   n,
                  int                   nnz,
                  const float*          alpha,
                  const cusparseMatDescr_t descrA,
                  const float*          csrValA,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  const float*          x,
                  const float*          beta,
                  float*                y)

cusparseStatus_t
cusparseDcsrmv_mp(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  int                   m,
                  int                   n,
                  int                   nnz,
                  const double*         alpha,
                  const cusparseMatDescr_t descrA,
                  const double*         csrValA,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  const double*         x,
                  const double*         beta,
                  double*               y)

cusparseStatus_t
cusparseCcsrmv_mp(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  int                   m,
                  int                   n,
                  int                   nnz,
                  const cuComplex*      alpha,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*      csrValA,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  const cuComplex*      x,
                  const cuComplex*      beta,
                  cuComplex*             y)

cusparseStatus_t
cusparseCcsrmv_mp(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  int                   m,
                  int                   n,
                  int                   nnz,
                  const cuDoubleComplex* alpha,
                  const cusparseMatDescr_t descrA,
                  const cuDoubleComplex* csrValA,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  const cuDoubleComplex* x,
                  const cuDoubleComplex* beta,
                  cuDoubleComplex*       y)

```

This function performs a load-balanced matrix-vector operation

$$y = \alpha * \text{op}(A) * x + \beta * y$$

**A** is an  $m \times n$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **x** and **y** are vectors;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

*Note:* This function is deprecated in favor of **cusparseCsrsvEx()** with **CUSPARSE\_ALG\_MERGE\_PATH** parameter which provides same functionality with better performance.

This routine was introduced specifically to address some of the loss of performance in the regular **csrmv()** code due to irregular sparsity patterns. The core kernels are based on the "MergePath" approach created by Duanne Merrill. By using this approach, we are able to provide performance independent of a sparsity pattern across data types.

Remark: **csrmv\_mp** only supports matrix type **CUSPARSE\_MATRIX\_TYPE\_GENERAL** and **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** operation.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation $\text{op}(A)$ . only support <b>CUSPARSE_OPERATION_NON_TRANSPOSE</b> .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>x</b>	<type> vector of <b>n</b> elements if $\text{op}(A) = A$ , and <b>m</b> elements if $\text{op}(A) = A^T$ or $\text{op}(A) = A^H$

<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>y</b>	<type> vector of <b>m</b> elements if $\text{op}(A) = A$ , and <b>n</b> elements if $\text{op}(A) = A^T$ or $\text{op}(A) = A^H$

## Output

<b>y</b>	<type> updated vector.
----------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 8.9. cusparseCsrmvEx()

```

cusparseStatus_t
cusparseCsrmvEx_bufferSize(cusparseHandle_t           handle,
                           cusparseAlgMode_t      alg,
                           cusparseOperation_t    transA,
                           int                   m,
                           int                   n,
                           int                   nnz,
                           const void*            alpha,
                           cudaDataType           alphatype,
                           const cusparseMatDescr_t descrA,
                           const void*            csrValA,
                           cudaDataType           csrValAtype,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           const void*            x,
                           cudaDataType           xtype,
                           const void*            beta,
                           cudaDataType           betatype,
                           void*                 Y,
                           cudaDataType           ytype,
                           cudaDataType           executiontype,
                           size_t*                bufferSizeInBytes)

cusparseStatus_t
cusparseCsrmvEx(cusparseHandle_t           handle,
                 cusparseAlgMode_t      alg,
                 cusparseOperation_t    transA,
                 int                   m,
                 int                   n,
                 int                   nnz,
                 const void*            alpha,
                 cudaDataType           alphatype,
                 const cusparseMatDescr_t descrA,
                 const void*            csrValA,
                 cudaDataType           csrValAtype,
                 const int*             csrRowPtrA,
                 const int*             csrColIndA,
                 const void*            x,
                 cudaDataType           xtype,
                 const void*            beta,
                 cudaDataType           betatype,
                 void*                 y,
                 cudaDataType           ytype,
                 cudaDataType           executiontype,
                 void*                 buffer)

```

This function is an extended version of **`cusparse<t>csrmv()`** which performs the matrix-vector multiply operation. For detailed description of the functionality, see **`cusparse<t>csrmv()`**. Also see **`cusparseAlgMode_t`** for **`alg`** parameter description

For alg **CUSPARSE\_ALG\_NAIVE**: for half-precision execution type, the minimum GPU architecture is SM\_53. Also, for both half-precision IO and execution, only CUSPARSE\_MATRIX\_TYPE\_GENERAL and CUSPARSE\_OPERATION\_NON\_TRANSPOSE are supported.

For alg **CUSPARSE\_ALG\_MERGE\_PATH**: half-precision is not supported, only **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and **CUSPARSE\_INDEX\_BASE\_ZERO** are supported. Input, output and execution types should be the same.

The function **`cusparseCsrsvEx_bufferSize`** returns the size of the workspace needed by **`cusparseCsrsvEx`**.

All pointers should be aligned with 128 bytes.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

#### Input specifically required by **`cusparseCsrsvEx`**

<b>alg</b>	Algorithm implementation for csrsv, see <b><code>cusparseAlgMode_t</code></b> for possible values.
<b>alphatype</b>	Data type of <b>alpha</b> .
<b>csrValAtype</b>	Data type of <b>csrValA</b> .
<b>xtype</b>	Data type of <b>x</b> .
<b>betatype</b>	Data type of <b>beta</b> .
<b>ytype</b>	Data type of <b>y</b> .
<b>executiontype</b>	Data type used for computation.
<b>bufferSizeInBytes</b>	Pointer to a <b>size_t</b> variable, which will be assigned with the size of workspace needed by <b><code>cusparseCsrsvEx</code></b> .
<b>buffer</b>	Pointer to workspace buffer

See **`cusparseStatus_t`** for the description of the return status

## 8.10. `cusparse<t>csrvs_analysis()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrvs2_analysis()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const float*           csrValA,
                        const int*              csrRowPtrA,
                        const int*              csrColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseDcsrsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const double*          csrValA,
                        const int*              csrRowPtrA,
                        const int*              csrColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseCcsrsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const cuComplex*       csrValA,
                        const int*              csrRowPtrA,
                        const int*              csrColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseZcsrsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const cuDoubleComplex* csrValA,
                        const int*              csrRowPtrA,
                        const int*              csrColIndA,
                        cusparseSolveAnalysisInfo_t info)

```

This function performs the analysis phase of the solution of a sparse triangular linear system

$$\text{op}(A) * \mathbf{y} = \alpha * \mathbf{x}$$

where **A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

The routine **csrv\_analysis** supports analysis phase of **csrv\_solve**, **csric0** and **csrilu0**. The user has to be careful of which routine is called after **csrv\_analysis**. The matrix descriptor must be the same for **csrv\_analysis** and its subsequent call to **csrv\_solve**, **csric0** and **csrilu0**.

For **csrv\_solve**, the matrix type must be **CUSPARSE\_MATRIX\_TYPE\_TRIANGULAR** or **CUSPARSE\_MATRIX\_TYPE\_GENERAL**.

For **csrilu0**, the matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**.

For **csric0**, the matrix type must be **CUSPARSE\_MATRIX\_TYPE\_SYMMETRIC** or **CUSPARSE\_MATRIX\_TYPE\_HERMITIAN**.

It is expected that this function will be executed only once for a given matrix and a particular operation type.

- ▶ This function requires a significant temporary amount of extra storage that is proportional to the matrix size
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation $\text{op}(A)$
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> for <b>csrv_solve</b> , <b>CUSPARSE_MATRIX_TYPE_SYMMETRIC</b> and <b>CUSPARSE_MATRIX_TYPE_HERMITIAN</b> for <b>csric0</b> , <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> for <b>csrilu0</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .

<b>info</b>	structure initialized using <code>cusparseCreateSolveAnalysisInfo</code> .
-------------	---

**Output**

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See `cusparseStatus_t` for the description of the return status

## 8.11. `cusparseCsrsv_analysisEx()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrsv2_analysis()` instead. *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseCsrsv_analysisEx(cusparseHandle_t handle,
                          cusparseOperation_t transA,
                          int m,
                          int nnz,
                          const cusparseMatDescr_t descrA,
                          const void* csrSortedValA,
                          cudaDataType csrSortedValAtype,
                          const int* csrSortedRowPtrA,
                          const int* csrSortedColIndA,
                          cusparseSolveAnalysisInfo_t info,
                          cudaDataType executionType)
```

This function is an extended version of `cusparse<t>csrsv_analysis()`. For detailed description of the functionality, see `cusparse<t>csrsv_analysis()`.

This function does not support half-precision execution type, but it supports half-precision IO with single precision execution.

- ▶ This function requires a significant temporary amount of extra storage that is proportional to the matrix size
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input specifically required by `cusparseCsrsv_analysisEx`

<code>csrSortedValAtype</code>	Data type of <code>csrSortedValA</code> .
<code>executionType</code>	Data type used for computation.

## 8.12. `cusparse<t>csrsv_solve()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrsv2_analysis()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrsv_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     const float*             alpha,
                     const cusparseMatDescr_t descrA,
                     const float*             csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const float*             f,
                     float*                  x)

cusparseStatus_t
cusparseDcsrsv_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     const double*            alpha,
                     const cusparseMatDescr_t descrA,
                     const double*            csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const double*            f,
                     double*                 x)

cusparseStatus_t
cusparseCcsrsv_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     const cuComplex*          alpha,
                     const cusparseMatDescr_t descrA,
                     const cuComplex*          csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuComplex*          f,
                     cuComplex*                x)

cusparseStatus_t
cusparseZcsrsv_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     const cuDoubleComplex*    alpha,
                     const cusparseMatDescr_t descrA,
                     const cuDoubleComplex*    csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuDoubleComplex*    f,
                     cuDoubleComplex*          x)

```

This function performs the solve phase of the solution of a sparse triangular linear system

$$\text{op}(A) * \mathbf{x} = \alpha * \mathbf{f}$$

where **A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrSortedValA**, **csrSortedRowPtrA**, and **csrSortedColIndA**; **f** and **x** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

This function may be executed multiple times for a given matrix and a particular operation type.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation $\text{op}(A)$
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrSortedValA</b>	<type> array of <b>nnz</b> (= <b>csrSortedRowPtrA(m) - csrSortedRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrSortedRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrSortedColIndA</b>	integer array of <b>nnz</b> (= <b>csrSortedRowPtrA(m) - csrSortedRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<b>f</b>	<type> right-hand-side vector of size <b>m</b> .

### Output

<b>x</b>	<type> solution vector of size <b>m</b> .
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 8.13. `cusparseCsrsv_solveEx()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrsv2_analysis()` instead. *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseCsrsv_solveEx(cusparseHandle_t handle,
                      cusparseOperation_t transA,
                      int m,
                      const void* alpha,
                      const void* alphatype,
                      const void* descrA,
                      const void* csrSortedValA,
                      const void* csrSortedValAtype,
                      const int* csrSortedRowPtrA,
                      const int* csrSortedColIndA,
                      cusparseSolveAnalysisInfo_t info,
                      const void* f,
                      const void* ftype,
                      void* x,
                      const void* xtype,
                      const void* executiontype)
```

This function is an extended version of `cusparse<t>csrsv_solve()`. For detailed description of the functionality, see `cusparse<t>csrsv_solve()`.

This function does not support half-precision execution type, but it supports half-precision IO with single precision execution.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input specifically required by `cusparseCsrsv_solveEx`

<code>alphatype</code>	Data type of <code>alpha</code> .
<code>csrSortedValAtype</code>	Data type of <code>csrSortedValA</code> .
<code>ftype</code>	Data type of <code>f</code> .
<code>xtype</code>	Data type of <code>x</code> .
<code>executiontype</code>	Data type used for computation.

## 8.14. `cusparse<t>csrv2_bufferSize()`

```

cusparseStatus_t
cusparseScsrsv2_bufferSize(cusparseHandle_t          handle,
                           cusparseOperation_t    transA,
                           int                   m,
                           int                   nnz,
                           const cusparseMatDescr_t descrA,
                           float*                csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrv2Info_t            info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseDcsrv2_bufferSize(cusparseHandle_t          handle,
                           cusparseOperation_t    transA,
                           int                   m,
                           int                   nnz,
                           const cusparseMatDescr_t descrA,
                           double*               csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrv2Info_t            info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseCcsrv2_bufferSize(cusparseHandle_t          handle,
                           cusparseOperation_t    transA,
                           int                   m,
                           int                   nnz,
                           const cusparseMatDescr_t descrA,
                           cuComplex*             csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrv2Info_t            info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseZcsrv2_bufferSize(cusparseHandle_t          handle,
                           cusparseOperation_t    transA,
                           int                   m,
                           int                   nnz,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*       csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrv2Info_t            info,
                           int*                  pBufferSizeInBytes)

```

This function returns the size of the buffer used in `csrv2`, a new sparse triangular linear system  $\text{op}(\mathbf{A}) \cdot \mathbf{y} = \alpha \cdot \mathbf{x}$ .

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays `csrValA`, `csrRowPtrA`, and `csrColIndA`; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

Although there are six combinations in terms of the parameter `trans` and the upper (lower) triangular part of `A`, `csrv2_bufferSize()` returns the maximum size buffer of these combinations. The buffer size depends on the dimension and the number of nonzero elements of the matrix. If the user changes the matrix, it is necessary to call `csrv2_bufferSize()` again to have the correct buffer size; otherwise, a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>transA</code>	the operation $\text{op}(A)$ .
<code>m</code>	number of rows of matrix <code>A</code> .
<code>nnz</code>	number of nonzero elements of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> , while the supported diagonal types are <code>CUSPARSE_DIAG_TYPE_UNIT</code> and <code>CUSPARSE_DIAG_TYPE_NON_UNIT</code> .
<code>csrValA</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .

### Output

<code>info</code>	record of internal states based on different algorithms.
<code>pBufferSizeInBytes</code>	number of bytes of the buffer used in the <code>csrv2_analysis</code> and <code>csrv2_solve</code> .

See `cusparseStatus_t` for the description of the return status

## 8.15. `cusparse<t>csrv2_analysis()`

```

cusparseStatus_t
cusparseScsrsv2_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t      transA,
                        int                      m,
                        int                      nnz,
                        const cusparseMatDescr_t descrA,
                        const float*             csrValA,
                        const int*               csrRowPtrA,
                        const int*               csrColIndA,
                        csrv2Info_t              info,
                        cusparseSolvePolicy_t    policy,
                        void*                   pBuffer)

cusparseStatus_t
cusparseDcsrv2_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t      transA,
                        int                      m,
                        int                      nnz,
                        const cusparseMatDescr_t descrA,
                        const double*            csrValA,
                        const int*               csrRowPtrA,
                        const int*               csrColIndA,
                        csrv2Info_t              info,
                        cusparseSolvePolicy_t    policy,
                        void*                   pBuffer)

cusparseStatus_t
cusparseCcsrv2_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t      transA,
                        int                      m,
                        int                      nnz,
                        const cusparseMatDescr_t descrA,
                        const cuComplex*         csrValA,
                        const int*               csrRowPtrA,
                        const int*               csrColIndA,
                        csrv2Info_t              info,
                        cusparseSolvePolicy_t    policy,
                        void*                   pBuffer)

cusparseStatus_t
cusparseZcsrv2_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t      transA,
                        int                      m,
                        int                      nnz,
                        const cusparseMatDescr_t descrA,
                        const cuDoubleComplex*   csrValA,
                        const int*               csrRowPtrA,
                        const int*               csrColIndA,
                        csrv2Info_t              info,
                        cusparseSolvePolicy_t    policy,
                        void*                   pBuffer)

```

This function performs the analysis phase of **csrv2**, a new sparse triangular linear system **op(A) \*y = alpha\*x**.

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

It is expected that this function will be executed only once for a given matrix and a particular operation type.

This function requires a buffer size returned by **csrv2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **csrv2\_analysis()** reports a structural zero and computes level information that is stored in opaque structure **info**. The level information can extract more parallelism for a triangular solver. However **csrv2\_solve()** can be done without level information. To disable level information, the user needs to specify the policy of the triangular solver as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **csrv2\_analysis()** always reports the first structural zero, even if the policy is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. No structural zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if  $A(j, j)$  is missing for some  $j$ . The user needs to call **cusparseXcsrv2\_zeroPivot()** to know where the structural zero is.

It is the user's choice whether to call **csrv2\_solve()** if **csrv2\_analysis()** reports a structural zero. In this case, the user can still call **csrv2\_solve()** which will return a numerical zero in the same position as the structural zero. However the result **x** is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of $m + 1$ elements that contains the start of every row and the end of the last row plus one.

<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure initialized using <b>cusparseCreateCsrsv2Info()</b> .
<b>policy</b>	The supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is returned by <b>csrv2_bufferSize()</b> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See **cusparseStatus\_t** for the description of the return status

## 8.16. `cusparse<t>csrv2_solve()`

```

cusparseStatus_t
cusparseScsrsv2_solve(cusparseHandle_t          handle,
                      cusparseOperation_t    transA,
                      int                   m,
                      int                   nnz,
                      const float*          alpha,
                      const cusparseMatDescr_t descra,
                      const float*          csrValA,
                      const int*            csrRowPtrA,
                      const int*            csrColIndA,
                      csrv2Info_t           info,
                      const float*          x,
                      float*                y,
                      cusparseSolvePolicy_t policy,
                      void*                 pBuffer)

```

```

cusparseStatus_t
cusparseDcsrsv2_solve(cusparseHandle_t          handle,
                      cusparseOperation_t    transA,
                      int                   m,
                      int                   nnz,
                      const double*         alpha,
                      const cusparseMatDescr_t descra,
                      const double*         csrValA,
                      const int*            csrRowPtrA,
                      const int*            csrColIndA,
                      csrv2Info_t           info,
                      const double*         x,
                      double*               y,
                      cusparseSolvePolicy_t policy,
                      void*                 pBuffer)

```

```

cusparseStatus_t
cusparseCcsrsv2_solve(cusparseHandle_t          handle,
                      cusparseOperation_t    transA,
                      int                   m,
                      int                   nnz,
                      const cuComplex*      alpha,
                      const cusparseMatDescr_t descra,
                      const cuComplex*      csrValA,
                      const int*            csrRowPtrA,
                      const int*            csrColIndA,
                      csrv2Info_t           info,
                      const cuComplex*      x,
                      cuComplex*            y,
                      cusparseSolvePolicy_t policy,
                      void*                 pBuffer)

```

```

cusparseStatus_t
cusparseZcsrsv2_solve(cusparseHandle_t          handle,
                      cusparseOperation_t    transA,
                      int                   m,
                      int                   nnz,
                      const cuDoubleComplex* alpha,
                      const cusparseMatDescr_t descra,
                      const cuDoubleComplex* csrValA,
                      const int*            csrRowPtrA,
                      const int*            csrColIndA,
                      csrv2Info_t           info,
                      const cuDoubleComplex* x,
                      cuDoubleComplex*      y,
                      cusparseSolvePolicy_t policy,
                      void*                 pBuffer)

```

This function performs the solve phase of **csrv2**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) \mathbf{x} = \alpha \mathbf{x}$ .

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **x** and **y** are the right-hand-side and the solution vectors;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

This function may be executed multiple times for a given matrix and a particular operation type.

This function requires the buffer size returned by **csrv2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **csrv2\_solve()** can be done without level information, the user still needs to be aware of consistency. If **csrv2\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **csrv2\_solve()** can be run with or without levels. On the contrary, if **csrv2\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **csrv2\_solve()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

The level information may not improve the performance but spend extra time doing analysis. For example, a tridiagonal matrix has no parallelism. In this case, **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** performs better than **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**. If the user has an iterative solver, the best approach is to do **csrv2\_analysis()** with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** once. Then do **csrv2\_solve()** with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** in the first run and with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** in the second run, picking faster one to perform the remaining iterations.

Function **csrv2\_solve()** reports the first numerical zero, including a structural zero. If **status** is 0, no numerical zero was found. Furthermore, no numerical zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if  $\mathbf{A}(j, j)$  is zero for some  $j$ . The user needs to call **cusparseXcsrv2\_zeroPivot()** to know where the numerical zero is.

For example, suppose L is a lower triangular matrix with unit diagonal, the following code solves **L\*y=x** by level information.

```

// Suppose that L is m x m sparse matrix represented by CSR format,
// L is lower triangular with unit diagonal.
// Assumption:
// - dimension of matrix L is m,
// - matrix L has nnz number zero elements,
// - handle is already created by cusparseCreate(),
// - (d_csrRowPtr, d_csrColInd, d_csrVal) is CSR of L on device memory,
// - d_x is right hand side vector on device memory,
// - d_y is solution vector on device memory.

cusparseMatDescr_t descr = 0;
csrv2Info_t info = 0;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.;
const cusparseSolvePolicy_t policy = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans = CUSPARSE_OPERATION_NON_TRANSPOSE;

// step 1: create a descriptor which contains
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has unit diagonal, specified by parameter CUSPARSE_DIAG_TYPE_UNIT
//   (L may not have all diagonal elements.)
cusparseCreateMatDescr(&descr);
cusparseSetMatIndexBase(descr, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatFillMode(descr, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr, CUSPARSE_DIAG_TYPE_UNIT);

// step 2: create a empty info structure
cusparseCreateCsrsv2Info(&info);

// step 3: query how much memory used in csrv2, and allocate the buffer
cusparseDcsrsv2_bufferSize(handle, trans, m, nnz, descr,
    d_csrVal, d_csrRowPtr, d_csrColInd, &pBufferSize);
// pBuffer returned by cudaMalloc is automatically aligned to 128 bytes.
cudaMalloc((void**)&pBuffer, pBufferSize);

// step 4: perform analysis
cusparseDcsrsv2_analysis(handle, trans, m, nnz, descr,
    d_csrVal, d_csrRowPtr, d_csrColInd,
    info, policy, pBuffer);
// L has unit diagonal, so no structural zero is reported.
status = cusparseXcsrsv2_zeroPivot(handle, info, &structural_zero);
if (CUSPARSE_STATUS_ZERO_PIVOT == status){
    printf("L(%d,%d) is missing\n", structural_zero, structural_zero);
}

// step 5: solve L*y = x
cusparseDcsrsv2_solve(handle, trans, m, nnz, &alpha, descr,
    d_csrVal, d_csrRowPtr, d_csrColInd, info,
    d_x, d_y, policy, pBuffer);
// L has unit diagonal, so no numerical zero is reported.
status = cusparseXcsrsv2_zeroPivot(handle, info, &numerical_zero);
if (CUSPARSE_STATUS_ZERO_PIVOT == status){
    printf("L(%d,%d) is zero\n", numerical_zero, numerical_zero);
}

// step 6: free resources
cudaFree(pBuffer);
cusparseDestroyCsrsv2Info(info);
cusparseDestroyMatDescr(descr);
cusparseDestroy(handle);

```

Remark: **`csrv2_solve()`** needs more nonzeros per row to achieve good performance. It would perform better if more than 16 nonzeros per row in average.

The function supports the following properties if **`pBuffer != NULL`**

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> ( $= \text{csrRowPtrA}(m) - \text{csrRowPtrA}(0)$ ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> ( $= \text{csrRowPtrA}(m) - \text{csrRowPtrA}(0)$ ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<b>x</b>	<type> right-hand-side vector of size <b>m</b> .
<b>policy</b>	The supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is return by <b>csrv2_bufferSize</b> .

## Output

<b>y</b>	<type> solution vector of size <b>m</b> .
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 8.17. `cusparseXcsrv2_zeroPivot()`

```
cusparseStatus_t
cusparseXcsrv2_zeroPivot(cusparseHandle_t handle,
                           csrv2Info_t      info,
                           int*             position)
```

If the returned error code is **CUSPARSE\_STATUS\_ZERO\_PIVOT**, `position=j` means  $A(j,j)$  has either a structural zero or a numerical zero. Otherwise `position=-1`.

The `position` can be 0-based or 1-based, the same as the matrix.

Function `cusparseXcsrv2_zeroPivot()` is a blocking call. It calls `cudaDeviceSynchronize()` to make sure all previous kernels are done.

The `position` can be in the host memory or device memory. The user can set the proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	<code>info</code> contains structural zero or numerical zero if the user already called <code>csrv2_analysis()</code> or <code>csrv2_solve()</code> .

### Output

<code>position</code>	if no structural or numerical zero, <code>position</code> is -1; otherwise, if $A(j,j)$ is missing or $U(j,j)$ is zero, <code>position=j</code> .
-----------------------	---

See `cusparseStatus_t` for the description of the return status

## 8.18. `cusparse<t>gemvi()`

```

cusparseStatus_t
cusparseSgemvi_bufferSize(cusparseHandle_t      handle,
                           cusparseOperation_t transA,
                           int                  m,
                           int                  n,
                           int                  nnz,
                           int*                pBufferSize)

cusparseStatus_t
cusparseDgemvi_bufferSize(cusparseHandle_t      handle,
                           cusparseOperation_t transA,
                           int                  m,
                           int                  n,
                           int                  nnz,
                           int*                pBufferSize)

cusparseStatus_t
cusparseCgemvi_bufferSize(cusparseHandle_t      handle,
                           cusparseOperation_t transA,
                           int                  m,
                           int                  n,
                           int                  nnz,
                           int*                pBufferSize)

cusparseStatus_t
cusparseZgemvi_bufferSize(cusparseHandle_t      handle,
                           cusparseOperation_t transA,
                           int                  m,
                           int                  n,
                           int                  nnz,
                           int*                pBufferSize)

```

```

cusparseStatus_t
cusparseSgemvi(cusparseHandle_t      handle,
               cusparseOperation_t transA,
               int                  m,
               int                  n,
               const float*         alpha,
               const float*         A,
               int                  lda,
               int                  nnz,
               const float*         x,
               const int*           xInd,
               const float*         beta,
               float*               y,
               cusparseIndexBase_t  idxBase,
               void*                pBuffer)

```

```

cusparseStatus_t
cusparseDgemvi(cusparseHandle_t      handle,
               cusparseOperation_t transA,
               int                  m,
               int                  n,
               const double*        alpha,
               const double*        A,
               int                  lda,
               int                  nnz,
               const double*        x,
               const int*           xInd,
               const float*         beta,
               double*              y,
               cusparseIndexBase_t  idxBase,
               void*                pBuffer)

```

This function performs the matrix-vector operation

$$y = \alpha * \text{op}(A) * x + \beta * y$$

**A** is an  $m \times n$  dense matrix and a sparse vector **x** that is defined in a sparse storage format by the two arrays **xVal**, **xInd** of length **nnz**, and **y** is a dense vector;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

To simplify the implementation, we have not (yet) optimized the transpose multiple case. We recommend the following for users interested in this case.

1. Convert the matrix from CSR to CSC format using one of the **csr2csc()** functions. Notice that by interchanging the rows and columns of the result you are implicitly transposing the matrix.
2. Call the **gemvi()** function with the **cusparseOperation\_t** parameter set to **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and with the interchanged rows and columns of the matrix stored in CSC format. This (implicitly) multiplies the vector by the transpose of the matrix in the original CSR format.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

The function **cusparse<t>gemvi\_bufferSize()** returns size of buffer used in **cusparse<t>gemvi()**

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation $\text{op}(A)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>A</b>	the pointer to dense matrix <b>A</b> .
<b>lda</b>	size of the leading dimension of <b>A</b> .
<b>nnz</b>	number of nonzero elements of vector <b>x</b> .
<b>x</b>	<type> sparse vector of <b>nnz</b> elements of size <b>n</b> if $\text{op}(A) = A$ , and size <b>m</b> if $\text{op}(A) = A^T$ or $\text{op}(A) = A^H$
<b>xInd</b>	Indices of non-zero values in <b>x</b>
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>y</b>	<type> dense vector of <b>m</b> elements if $\text{op}(A) = A$ , and <b>n</b> elements if $\text{op}(A) = A^T$ or $\text{op}(A) = A^H$

<b>idxBase</b>	0 or 1, for 0 based or 1 based indexing, respectively
<b>pBufferSize</b>	number of elements needed the buffer used in <code>cusparse&lt;t&gt;gemvi()</code> .
<b>pBuffer</b>	working space buffer

**Output**

<b>y</b>	<type> updated dense vector.
----------	------------------------------

See `cusparseStatus_t` for the description of the return status

## 8.19. `cusparse<t>hybmv()` [DEPRECATED]

**[[DEPRECATED]]**. The routine will be removed in the next major release

```

cusparseStatus_t
cusparseShybmv(cusparseHandle_t          handle,
               cusparseOperation_t      transA,
               const float*             alpha,
               const cusparseMatDescr_t descrA,
               const cusparseHybMat_t   hybA,
               const float*             x,
               const float*             beta,
               float*                  y)

cusparseStatus_t
cusparseDhybmv(cusparseHandle_t          handle,
               cusparseOperation_t      transA,
               const double*            alpha,
               const cusparseMatDescr_t descrA,
               const cusparseHybMat_t   hybA,
               const double*            x,
               const double*            beta,
               double*                 y)

cusparseStatus_t
cusparseChybmv(cusparseHandle_t          handle,
               cusparseOperation_t      transA,
               const cuComplex*         alpha,
               const cusparseMatDescr_t descrA,
               const cusparseHybMat_t   hybA,
               const cuComplex*         x,
               const cuComplex*         beta,
               cuComplex*               y)

cusparseStatus_t
cusparseZhybmv(cusparseHandle_t          handle,
               cusparseOperation_t      transA,
               const cuDoubleComplex*   alpha,
               const cusparseMatDescr_t descrA,
               const cusparseHybMat_t   hybA,
               const cuDoubleComplex*   x,
               const cuDoubleComplex*   beta,
               cuDoubleComplex*         y)

```

This function performs the matrix-vector operation

$$y = \alpha * \text{op}(A) * x + \beta * y$$

**A** is an  $m \times n$  sparse matrix that is defined in the HYB storage format by an opaque data structure **hybA**, **x** and **y** are vectors,  $\alpha$  and  $\beta$  are scalars, and

$$\text{op}(A) = \begin{cases} A & \text{if transA == CUSPARSE_OPERATION_NON_TRANSPOSE} \end{cases}$$

Notice that currently only  $\text{op}(A) = A$  is supported.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$ (currently only $\text{op}(A) = A$ is supported).
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.
<b>x</b>	<type> vector of <b>n</b> elements.
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>y</b>	<type> vector of <b>m</b> elements.

### Output

<b>y</b>	<type> updated vector.
----------	------------------------

See [cusparseStatus\\_t](#) for the description of the return status

## 8.20. `cusparse<t>hybsv_analysis()` [DEPRECATED]

**[[DEPRECATED]]**. The routine will be removed in the next major release

```

cusparseStatus_t
cusparseShybsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        const cusparseMatDescr_t descrA,
                        cusparseHybMat_t       hybA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseDhybsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        const cusparseMatDescr_t descrA,
                        cusparseHybMat_t       hybA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseChybsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        const cusparseMatDescr_t descrA,
                        cusparseHybMat_t       hybA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseZhybsv_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        const cusparseMatDescr_t descrA,
                        cusparseHybMat_t       hybA,
                        cusparseSolveAnalysisInfo_t info)

```

This function performs the analysis phase of the solution of a sparse triangular linear system

$$\text{op}(A) * \mathbf{y} = \alpha * \mathbf{x}$$

**A** is an  $m \times m$  sparse matrix that is defined in HYB storage format by an opaque data structure **hybA**, **x** and **y** are the right-hand-side and the solution vectors,  $\alpha$  is a scalar, and

$$\text{op}(A) = \begin{cases} A & \text{if } \text{transA} == \text{CUSPARSE_OPERATION\_NON\_TRANSPOSE} \\ & \end{cases}$$

Notice that currently only  $\text{op}(A) = A$  is supported.

It is expected that this function will be executed only once for a given matrix and a particular operation type.

- ▶ This function requires a significant amount of extra storage that is proportional to the matrix size that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
---------------	---

<b>transA</b>	the operation $\text{op}(A)$ (currently only $\text{op}(A)=A$ is supported).
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and diagonal type <b>USPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.
<b>info</b>	structure initialized using <b>cusparseCreateSolveAnalysisInfo()</b> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See **cusparseStatus\_t** for the description of the return status

## 8.21. `cusparse<t>hybsv_solve()` [DEPRECATED]

**[[DEPRECATED]]**. The routine will be removed in the next major release

```

cusparseStatus_t
cusparseShybsv_solve(cusparseHandle_t          handle,
                     cusparseOperation_t   transA,
                     const float*          alpha,
                     const cusparseMatDescr_t descrA,
                     cusparseHybMat_t       hybA,
                     cusparseSolveAnalysisInfo_t info,
                     const float*          x,
                     float*                y)

cusparseStatus_t
cusparseDhybsv_solve(cusparseHandle_t          handle,
                     cusparseOperation_t   transA,
                     const double*         alpha,
                     const cusparseMatDescr_t descrA,
                     cusparseHybMat_t       hybA,
                     cusparseSolveAnalysisInfo_t info,
                     const double*         x,
                     double*               y)

cusparseStatus_t
cusparseChybsv_solve(cusparseHandle_t          handle,
                     cusparseOperation_t   transA,
                     const cuComplex*      alpha,
                     const cusparseMatDescr_t descrA,
                     cusparseHybMat_t       hybA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuComplex*      x,
                     cuComplex*             y)

cusparseStatus_t
cusparseZhybsv_solve(cusparseHandle_t          handle,
                     cusparseOperation_t   transA,
                     const cuDoubleComplex* alpha,
                     const cusparseMatDescr_t descrA,
                     cusparseHybMat_t       hybA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuDoubleComplex* x,
                     cuDoubleComplex*       y)

```

This function performs the solve phase of the solution of a sparse triangular linear system:

$$\text{op}(A) * \mathbf{y} = \alpha * \mathbf{x}$$

**A** is an  $m \times m$  sparse matrix that is defined in HYB storage format by an opaque data structure **hybA**, **x** and **y** are the right-hand-side and the solution vectors,  $\alpha$  is a scalar, and

$$\text{op}(A) = \begin{cases} A & \text{if } \text{transA} == \text{CUSPARSE\_OPERATION\_NON\_TRANPOSE} \\ A^T & \text{otherwise} \end{cases}$$

Notice that currently only  $\text{op}(A) = A$  is supported.

This function may be executed multiple times for a given matrix and a particular operation type.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$ (currently only $\text{op}(A)=A$ is supported).
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and the diagonal type is <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.
<b>info</b>	structure with information collected during the analysis phase (that should be passed to the solve phase unchanged).
<b>x</b>	<type> right-hand-side vector of size <b>m</b> .

### Output

<b>y</b>	<type> solution vector of size <b>m</b> .
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 9.

# CUSPARSE LEVEL 3 FUNCTION REFERENCE

This chapter describes sparse linear algebra functions that perform operations between sparse and (usually tall) dense matrices.

In particular, the solution of sparse triangular linear systems with multiple right-hand sides is implemented in two phases. First, during the analysis phase, the sparse triangular matrix is analyzed to determine the dependencies between its elements by calling the appropriate `csrsm_analysis()` function. The analysis is specific to the sparsity pattern of the given matrix and to the selected `cusparseOperation_t` type. The information from the analysis phase is stored in the parameter of type `cusparseSolveAnalysisInfo_t` that has been initialized previously with a call to `cusparseCreateSolveAnalysisInfo()`.

Second, during the solve phase, the given sparse triangular linear system is solved using the information stored in the `cusparseSolveAnalysisInfo_t` parameter by calling the appropriate `csrsm_solve()` function. The solve phase may be performed multiple times with different multiple right-hand sides, while the analysis phase needs to be performed only once. This is especially useful when a sparse triangular linear system must be solved for different sets of multiple right-hand sides one at a time, while its coefficient matrix remains the same.

Finally, once all the solves have completed, the opaque data structure pointed to by the `cusparseSolveAnalysisInfo_t` parameter can be released by calling `cusparseDestroySolveAnalysisInfo()`. For more information please refer to [3].

## 9.1. `cusparse<t>bsrmm()`

```

cusparseStatus_t
cusparseSbsrmm(cusparseHandle_t
                cusparseDirection_t
                cusparseOperation_t
                cusparseOperation_t
                int
                int
                int
                int
                const float*
                const cusparseMatDescr_t
                const float*
                const int*
                const int*
                int
                const float*
                int
                const float*
                float*
                int
                handle,
                dirA,
                transA,
                transB,
                mb,
                n,
                kb,
                nnzb,
                alpha,
                descrA,
                bsrValA,
                bsrRowPtrA,
                bsrColIndA,
                blockDim,
                B,
                ldb,
                beta,
                C,
                ldc)

cusparseStatus_t
cusparseDbsrmm(cusparseHandle_t
                cusparseDirection_t
                cusparseOperation_t
                cusparseOperation_t
                int
                int
                int
                int
                const double*
                const cusparseMatDescr_t
                const double*
                const int*
                const int*
                int
                const double*
                int
                const double*
                double*
                int
                handle,
                dirA,
                transA,
                transB,
                mb,
                n,
                kb,
                nnzb,
                alpha,
                descrA,
                bsrValA,
                bsrRowPtrA,
                bsrColIndA,
                blockDim,
                B,
                ldb,
                beta,
                C,
                ldc)

cusparseStatus_t
cusparseCbsrmm(cusparseHandle_t
                cusparseDirection_t
                cusparseOperation_t
                cusparseOperation_t
                int
                int
                int
                int
                const cuComplex*
                const cusparseMatDescr_t
                const cuComplex*
                const int*
                const int*
                int
                const cuComplex*
                int
                const cuComplex*
                cuComplex*
                int
                handle,
                dirA,
                transA,
                transB,
                mb,
                n,
                kb,
                nnzb,
                alpha,
                descrA,
                bsrValA,
                bsrRowPtrA,
                bsrColIndA,
                blockDim,
                B,
                ldb,
                beta,
                C,
                ldc)

```

This function performs one of the following matrix-matrix operations:

$$C = \alpha * \text{op}(A) * \text{op}(B) + \beta * C$$

**A** is an **mb**\***kb** sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **B** and **C** are dense matrices;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if transA == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if transA == CUSPARSE_OPERATION_TRANSPOSE (not supported)} \\ A^H & \text{if transA == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE (not supported)} \end{cases}$$

and

$$\text{op}(B) = \begin{cases} B & \text{if transB == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ B^T & \text{if transB == CUSPARSE_OPERATION_TRANSPOSE} \\ B^H & \text{if transB == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE (not supported)} \end{cases}$$

The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**.

The motivation of **transpose(B)** is to improve memory access of matrix **B**. The computational pattern of **A\*transpose(B)** with matrix **B** in column-major order is equivalent to **A\*B** with matrix **B** in row-major order.

In practice, no operation in an iterative solver or eigenvalue solver uses **A\*transpose(B)**. However, we can perform **A\*transpose(transpose(B))** which is the same as **A\*B**. For example, suppose **A** is **mb**\***kb**, **B** is **k**\***n** and **C** is **m**\***n**, the following code shows usage of **cusparseDbsrmm()**.

```
// A is mb*kb, B is k*n and C is m*n
const int m = mb*blockSize;
const int k = kb*blockSize;
const int ldb_B = k; // leading dimension of B
const int ldc_C = m; // leading dimension of C
// perform C:=alpha*A*B + beta*C
cusparseSetMatType(descrA, CUSPARSE_MATRIX_TYPE_GENERAL );
cusparseDbsrmm(cusparseHandle,
                CUSPARSE_DIRECTION_COLUMN,
                CUSPARSE_OPERATION_NON_TRANSPOSE,
                CUSPARSE_OPERATION_NON_TRANSPOSE,
                mb, n, kb, nnzb, alpha,
                descrA, bsrValA, bsrRowPtrA, bsrColIndA, blockSize,
                B, ldb_B,
                beta, C, ldc_C);
```

Instead of using **A\*B**, our proposal is to transpose **B** to **Bt** by first calling **cublas<t>gemm()**, and then to perform **A\*transpose(Bt)**.

```
// step 1: Bt := transpose(B)
    const int m = mb*blockSize;
    const int k = kb*blockSize;
    double *Bt;
    const int ldb_Bt = n; // leading dimension of Bt
    cudaMalloc((void**)&Bt, sizeof(double)*ldb_Bt*k);
    double one = 1.0;
    double zero = 0.0;
    cublasSetPointerMode(cublas_handle, CUBLAS_POINTER_MODE_HOST);
    cublasDgemm(cublas_handle, CUBLAS_OP_T, CUBLAS_OP_T,
                n, k, &one, B, int ldb_B, &zero, Bt, int ldb_Bt, Bt, ldb_Bt);

// step 2: perform C:=alpha*A*transpose(Bt) + beta*C
    cusparseDbsrmm(cusparse_handle,
                    CUSPARSE_DIRECTION_COLUMN,
                    CUSPARSE_OPERATION_NON_TRANSPOSE,
                    CUSPARSE_OPERATION_TRANSPOSE,
                    mb, n, kb, nnzb, alpha,
                    descrA, bsrValA, bsrRowPtrA, bsrColIndA, blockSize,
                    Bt, ldb_Bt,
                    beta, C, ldc);
```

The routine has the following properties if **blockDim > 1**:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dir</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation <b>op(A)</b> .
<b>transB</b>	the operation <b>op(B)</b> .
<b>mb</b>	number of block rows of sparse matrix <b>A</b> .
<b>n</b>	number of columns of dense matrix <b>op(B)</b> and <b>A</b> .
<b>kb</b>	number of block columns of sparse matrix <b>A</b> .
<b>nnzb</b>	number of non-zero blocks of sparse matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix <b>A</b> .

<b>bsrRowPtrA</b>	integer array of <code>mb</code> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> , larger than zero.
<b>B</b>	array of dimensions ( <code>ldb</code> , <code>n</code> ) if <code>op(B)=B</code> and ( <code>ldb</code> , <code>k</code> ) otherwise.
<b>ldb</b>	leading dimension of <b>B</b> . If <code>op(B)=B</code> , it must be at least <code>max(1, k)</code> If <code>op(B) != B</code> , it must be at least <code>max(1, n)</code> .
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>c</b> does not have to be a valid input.
<b>C</b>	array of dimensions ( <code>ldc</code> , <code>n</code> ).
<b>ldc</b>	leading dimension of <b>c</b> . It must be at least <code>max(1, m)</code> if <code>op(A)=A</code> and at least <code>max(1, k)</code> otherwise.

## Output

<b>c</b>	<type> updated array of dimensions ( <code>ldc</code> , <code>n</code> ).
----------	---

See `cusparseStatus_t` for the description of the return status

## 9.2. `cusparse<t>bsrsm2_bufferSize()`

```

cusparseStatus_t
cusparseSbsrsm2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t   dirA,
                           cusparseOperation_t  transA,
                           cusparseOperation_t  transX,
                           int                  mb,
                           int                  n,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           float*                bsrSortedValA,
                           const int*             bsrSortedRowPtrA,
                           const int*             bsrSortedColIndA,
                           int                   blockDim,
                           bsrsm2Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseDbsrsm2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t   dirA,
                           cusparseOperation_t  transA,
                           cusparseOperation_t  transX,
                           int                  mb,
                           int                  n,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           double*               bsrSortedValA,
                           const int*             bsrSortedRowPtrA,
                           const int*             bsrSortedColIndA,
                           int                   blockDim,
                           bsrsm2Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseCbsrsm2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t   dirA,
                           cusparseOperation_t  transA,
                           cusparseOperation_t  transX,
                           int                  mb,
                           int                  n,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuComplex*            bsrSortedValA,
                           const int*             bsrSortedRowPtrA,
                           const int*             bsrSortedColIndA,
                           int                   blockDim,
                           bsrsm2Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseZbsrsm2_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t   dirA,
                           cusparseOperation_t  transA,
                           cusparseOperation_t  transX,
                           int                  mb,
                           int                  n,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*      bsrSortedValA,
                           const int*             bsrSortedRowPtrA,
                           const int*             bsrSortedColIndA,
                           int                   blockDim,
                           bsrsm2Info_t          info,
                           int*                  pBufferSizeInBytes)

```

This function returns size of buffer used in **bsrsm2()**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) * \text{op}(\mathbf{X}) = \alpha \text{op}(\mathbf{B})$ .

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **B** and **X** are the right-hand-side and the solution matrices;  $\alpha$  is a scalar; and

$$\text{op}(\mathbf{A}) = \begin{cases} \mathbf{A} & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ \mathbf{A}^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ \mathbf{A}^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

Although there are six combinations in terms of parameter **trans** and the upper (and lower) triangular part of **A**, **bsrsm2\_bufferSize()** returns the maximum size of the buffer among these combinations. The buffer size depends on dimension **mb**, **blockDim** and the number of nonzeros of the matrix, **nnzb**. If the user changes the matrix, it is necessary to call **bsrsm2\_bufferSize()** again to get the correct buffer size, otherwise a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation <b>op(A)</b> .
<b>transX</b>	the operation <b>op(X)</b> .
<b>mb</b>	number of block rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>op(B)</b> and <b>op(X)</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) nonzero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb</b> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> ; larger than zero.

## Output

<b>info</b>	record internal states based on different algorithms.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>bsrsm2_analysis()</b> and <b>bsrsm2_solve()</b> .

See [cusparseStatus\\_t](#) for the description of the return status

## 9.3. `cusparse<t>bsrsm2_analysis()`

```

cusparseStatus_t
cusparseSbsrsm2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        cusparseOperation_t   transX,
                        int                  mb,
                        int                  n,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const float*          bsrSortedVal,
                        const int*            bsrSortedRowPtr,
                        const int*            bsrSortedColInd,
                        int                  blockDim,
                        bsrsm2Info_t          info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseDbsrsm2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        cusparseOperation_t   transX,
                        int                  mb,
                        int                  n,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const double*         bsrSortedVal,
                        const int*            bsrSortedRowPtr,
                        const int*            bsrSortedColInd,
                        int                  blockDim,
                        bsrsm2Info_t          info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseCbsrsm2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        cusparseOperation_t   transX,
                        int                  mb,
                        int                  n,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const cuComplex*      bsrSortedVal,
                        const int*            bsrSortedRowPtr,
                        const int*            bsrSortedColInd,
                        int                  blockDim,
                        bsrsm2Info_t          info,
                        cusparseSolvePolicy_t policy,
                        void*                pBuffer)

cusparseStatus_t
cusparseZbsrsm2_analysis(cusparseHandle_t          handle,
                        cusparseDirection_t    dirA,
                        cusparseOperation_t   transA,
                        cusparseOperation_t   transX,
                        int                  mb,
                        int                  n,
                        int                  nnzb,
                        const cusparseMatDescr_t descrA,
                        const cuDoubleComplex* bsrSortedVal,
                        const int*            bsrSortedRowPtr,
                        const int*            bsrSortedColInd.

```

This function performs the analysis phase of **bsrsm2()**, a new sparse triangular linear system  $\text{op}(\mathbf{A}) * \text{op}(\mathbf{X}) = \alpha \text{op}(\mathbf{B})$ .

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **B** and **X** are the right-hand-side and the solution matrices;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

and

$$\text{op}(X) = \begin{cases} X & \text{if transX == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ X^T & \text{if transX == CUSPARSE_OPERATION_TRANSPOSE} \\ X^H & \text{if transX == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE (not supported)} \end{cases}$$

and **op(B)** and **op(X)** are equal.

The block of BSR format is of size **blockDim\*blockDim**, stored in column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_ROW** or **CUSPARSE\_DIRECTION\_COLUMN**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored.

It is expected that this function will be executed only once for a given matrix and a particular operation type.

This function requires the buffer size returned by **bsrsm2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsrsm2\_analysis()** reports a structural zero and computes the level information stored in opaque structure **info**. The level information can extract more parallelism during a triangular solver. However **bsrsm2\_solve()** can be done without level information. To disable level information, the user needs to specify the policy of the triangular solver as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **bsrsm2\_analysis()** always reports the first structural zero, even if the parameter **policy** is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. Besides, no structural zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if block **A(j, j)** is missing for some **j**. The user must call **cusparseXbsrsm2\_query\_zero\_pivot()** to know where the structural zero is.

If **bsrsm2\_analysis()** reports a structural zero, the solve will return a numerical zero in the same position as the structural zero but this result **x** is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .

<b>transA</b>	the operation <code>op(A)</code> .
<b>transX</b>	the operation <code>op(B)</code> and <code>op(X)</code> .
<b>mb</b>	number of block rows of matrix <code>A</code> .
<b>n</b>	number of columns of matrix <code>op(B)</code> and <code>op(X)</code> .
<b>nnzb</b>	number of non-zero blocks of matrix <code>A</code> .
<b>descrA</b>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> , while the supported diagonal types are <code>CUSPARSE_DIAG_TYPE_UNIT</code> and <code>CUSPARSE_DIAG_TYPE_NON_UNIT</code> .
<b>bsrValA</b>	<type> array of <code>nnzb</code> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) nonzero blocks of matrix <code>A</code> .
<b>bsrRowPtrA</b>	integer array of <code>mb</code> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <code>nnzb</code> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) column indices of the nonzero blocks of matrix <code>A</code> .
<b>blockDim</b>	block dimension of sparse matrix <code>A</code> ; larger than zero.
<b>info</b>	structure initialized using <code>cusparseCreateBsrsm2Info</code> .
<b>policy</b>	The supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<b>pBuffer</b>	buffer allocated by the user; the size is return by <code>bsrsm2_bufferSize()</code> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See `cusparseStatus_t` for the description of the return status

## 9.4. `cusparse<t>bsrsm2_solve()`

```

cusparseStatus_t
cusparseSbsrsm2_solve(cusparseHandle_t           handle,
                      cusparseDirection_t    dirA,
                      cusparseOperation_t   transA,
                      cusparseOperation_t   transX,
                      int                  mb,
                      int                  n,
                      int                  nnzb,
                      const float*          alpha,
                      const cusparseMatDescr_t descrA,
                      const float*          bsrSortedVal,
                      const int*             bsrSortedRowPtr,
                      const int*             bsrSortedColInd,
                      int                  blockDim,
                      bsrsm2Info_t          info,
                      const float*          B,
                      int                  ldb,
                      float*                X,
                      int                  ldx,
                      cusparseSolvePolicy_t policy,
                      void*                pBuffer)

cusparseStatus_t
cusparseDbsrsm2_solve(cusparseHandle_t           handle,
                      cusparseDirection_t    dirA,
                      cusparseOperation_t   transA,
                      cusparseOperation_t   transX,
                      int                  mb,
                      int                  n,
                      int                  nnzb,
                      const double*         alpha,
                      const cusparseMatDescr_t descrA,
                      const double*         bsrSortedVal,
                      const int*            bsrSortedRowPtr,
                      const int*            bsrSortedColInd,
                      int                  blockDim,
                      bsrsm2Info_t          info,
                      const double*         B,
                      int                  ldb,
                      double*               X,
                      int                  ldx,
                      cusparseSolvePolicy_t policy,
                      void*                pBuffer)

cusparseStatus_t
cusparseCbsrsm2_solve(cusparseHandle_t           handle,
                      cusparseDirection_t    dirA,
                      cusparseOperation_t   transA,
                      cusparseOperation_t   transX,
                      int                  mb,
                      int                  n,
                      int                  nnzb,
                      const cuComplex*      alpha,
                      const cusparseMatDescr_t descrA,
                      const cuComplex*      bsrSortedVal,
                      const int*             bsrSortedRowPtr,
                      const int*             bsrSortedColInd,
                      int                  blockDim,
                      bsrsm2Info_t          info,
                      const cuComplex*      B,
                      int                  ldb,
                      cuComplex*             X,
                      int                  ldx.

```

This function performs the solve phase of the solution of a sparse triangular linear system:

$$\text{op}(A) * \text{op}(X) = \alpha * \text{op}(B)$$

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**; **B** and **X** are the right-hand-side and the solution matrices;  $\alpha$  is a scalar, and

$$\text{op}(A) = \begin{cases} A & \text{if transA == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if transA == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if transA == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

and

$$\text{op}(X) = \begin{cases} X & \text{if transX == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ X^T & \text{if transX == CUSPARSE_OPERATION_TRANSPOSE} \\ X^H & \text{not supported} \end{cases}$$

Only **op(A)=A** is supported.

**op(B)** and **op(X)** must be performed in the same way. In other words, if **op(B)=B**, **op(X)=X**.

The block of BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_ROW** or **CUSPARSE\_DIRECTION\_COLUMN**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored. Function **bsrsm02\_solve()** can support an arbitrary **blockDim**.

This function may be executed multiple times for a given matrix and a particular operation type.

This function requires the buffer size returned by **bsrsm2\_bufferSize()**.

The address of **pBuffer** must be multiple of 128 bytes. If it is not,

**CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **bsrsm2\_solve()** can be done without level information, the user still needs to be aware of consistency. If **bsrsm2\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **bsrsm2\_solve()** can be run with or without levels. On the other hand, if **bsrsm2\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **bsrsm2\_solve()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsrsm02\_solve()** has the same behavior as **bsrsv02\_solve()**, reporting the first numerical zero, including a structural zero. The user must call **cusparseXbsrsm2\_query\_zero\_pivot()** to know where the numerical zero is.

The motivation of **transpose(X)** is to improve the memory access of matrix **X**. The computational pattern of **transpose(X)** with matrix **X** in column-major order is equivalent to **X** with matrix **X** in row-major order.

In-place is supported and requires that **B** and **X** point to the same memory block, and **ldb=ldx**.

The function supports the following properties if **pBuffer != NULL**

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>transA</b>	the operation <b>op(A)</b> .
<b>transX</b>	the operation <b>op(B)</b> and <b>op(X)</b> .
<b>mb</b>	number of block rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>op(B)</b> and <b>op(X)</b> .
<b>nnzb</b>	number of non-zero blocks of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) non-zero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb + 1</b> elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> ; larger than zero.
<b>info</b>	structure initialized using <b>cusparseCreateBsrsm2Info()</b> .
<b>B</b>	<type> right-hand-side array.
<b>ldb</b>	leading dimension of <b>B</b> . If <b>op(B)=B</b> , <b>ldb &gt;= (mb*blockDim)</b> ; otherwise, <b>ldb &gt;= n</b> .
<b>idx</b>	leading dimension of <b>X</b> . If <b>op(X)=X</b> , then <b>idx &gt;= (mb*blockDim)</b> . otherwise <b>idx &gt;= n</b> .
<b>policy</b>	the supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user; the size is returned by <b>bsrsm2_bufferSize()</b> .

## Output

<b>X</b>	<type> solution array with leading dimensions <b>idx</b> .
----------	--

See `cusparseStatus_t` for the description of the return status

## 9.5. `cusparseXbsrsm2_zeroPivot()`

```
cusparseStatus_t
cusparseXbsrsm2_zeroPivot(cusparseHandle_t handle,
                           bsrsm2Info_t     info,
                           int*             position)
```

If the returned error code is `CUSPARSE_STATUS_ZERO_PIVOT`, `position=j` means  $A(j,j)$  is either a structural zero or a numerical zero (singular block). Otherwise `position=-1`.

The `position` can be 0-base or 1-base, the same as the matrix.

Function `cusparseXbsrsm2_zeroPivot()` is a blocking call. It calls `cudaDeviceSynchronize()` to make sure all previous kernels are done.

The `position` can be in the host memory or device memory. The user can set the proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	<code>info</code> contains a structural zero or a numerical zero if the user already called <code>bsrsm2_analysis()</code> or <code>bsrsm2_solve()</code> .

### Output

<code>position</code>	if no structural or numerical zero, <code>position</code> is -1; otherwise, if $A(j,j)$ is missing or $U(j,j)$ is zero, <code>position=j</code> .
-----------------------	---

See `cusparseStatus_t` for the description of the return status

## 9.6. `cusparse<t>csrmm()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseSpMM\(\)` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrmm(cusparseHandle_t           handle,
               cusparseOperation_t      transA,
               int                      m,
               int                      n,
               int                      k,
               int                      nnz,
               const float*             alpha,
               const cusparseMatDescr_t descrA,
               const float*             csrValA,
               const int*               csrRowPtrA,
               const int*               csrColIndA,
               const float*             B,
               int                      ldb,
               const float*             beta,
               float*                  C,
               int                      ldc)

cusparseStatus_t
cusparseDcsrmm(cusparseHandle_t           handle,
               cusparseOperation_t      transA,
               int                      m,
               int                      n,
               int                      k,
               int                      nnz,
               const double*            alpha,
               const cusparseMatDescr_t descrA,
               const double*            csrValA,
               const int*               csrRowPtrA,
               const int*               csrColIndA,
               const double*            B,
               int                      ldb,
               const double*            beta,
               double*                 C,
               int                      ldc)

cusparseStatus_t
cusparseCcsrmm(cusparseHandle_t           handle,
               cusparseOperation_t      transA,
               int                      m,
               int                      n,
               int                      k,
               int                      nnz,
               const cuComplex*         alpha,
               const cusparseMatDescr_t descrA,
               const cuComplex*         csrValA,
               const int*               csrRowPtrA,
               const int*               csrColIndA,
               const cuComplex*         B,
               int                      ldb,
               const cuComplex*         beta,
               cuComplex*               C,
               int                      ldc)

cusparseStatus_t
cusparseZcsrmm(cusparseHandle_t           handle,
               cusparseOperation_t      transA,
               int                      m,
               int                      n,
               int                      k,
               int                      nnz,
               const cuDoubleComplex*   alpha,
               const cusparseMatDescr_t descrA,
               const cuDoubleComplex*   csrValA,
               const int*               csrRowPtrA,
               const int*               csrColIndA,
               const cuDoubleComplex*   B,
               int                      ldb,
               const cuDoubleComplex*   beta,
               cuDoubleComplex*         C,
               int                      ldc)

```

This function performs one of the following matrix-matrix operations:

$$C = \alpha * \text{op}(A) * B + \beta * C$$

**A** is an  $m \times k$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **C** are dense matrices;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

When using the (conjugate) transpose of a general matrix or a Hermitian/symmetric matrix, this routine may produce slightly different results with the same input parameters during different runs of this function. For these matrix types it uses atomic operations to compute the final result; consequently, many threads may be adding floating point numbers to the same memory location without any specific ordering, which may produce slightly different results for each run.

If exactly the same output is required for any input when multiplying by the transpose of a general matrix, the following procedure can be used:

1. Convert the matrix from CSR to CSC format using one of the **csr2csc()** functions. Notice that by interchanging the rows and columns of the result you are implicitly transposing the matrix.
2. Call the **csrmmtm()** function with the **cusparseOperation\_t** parameter set to **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and with the interchanged rows and columns of the matrix stored in CSC format. This (implicitly) multiplies the vector by the transpose of the matrix in the original CSR format.

The function has the following properties for operation **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and the matrix type is different from **CUSPARSE\_MATRIX\_TYPE\_GENERAL**

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

otherwise:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$
<b>m</b>	number of rows of sparse matrix <b>A</b> .
<b>n</b>	number of columns of dense matrices <b>B</b> and <b>C</b> .
<b>k</b>	number of columns of sparse matrix <b>A</b> .

<b>nnz</b>	number of nonzero elements of sparse matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , <b>CUSPARSE_MATRIX_TYPE_SYMMETRIC</b> , and <b>CUSPARSE_MATRIX_TYPE_HERMITIAN</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>B</b>	array of dimensions ( <b>ldb</b> , <b>n</b> ).
<b>ldb</b>	leading dimension of <b>B</b> . It must be at least $\max(1, k)$ if $\text{op}(A)=A$ and at least $\max(1, m)$ otherwise.
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>c</b> does not have to be a valid input.
<b>C</b>	array of dimensions ( <b>ldc</b> , <b>n</b> ).
<b>ldc</b>	leading dimension of <b>c</b> . It must be at least $\max(1, m)$ if $\text{op}(A)=A$ and at least $\max(1, k)$ otherwise.

## Output

<b>c</b>	<type> updated array of dimensions ( <b>ldc</b> , <b>n</b> ).
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 9.7. `cusparse<t>csrmm2()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseSpMM\(\)` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrmm2(cusparseHandle_t          handle,
                 cusparseOperation_t    transA,
                 cusparseOperation_t    transB,
                 int                   m,
                 int                   n,
                 int                   k,
                 int                   nnz,
                 const float*          alpha,
                 const cusparseMatDescr_t descRA,
                 const float*          csrValA,
                 const int*             csrRowPtrA,
                 const int*             csrColIndA,
                 const float*          B,
                 int                   ldb,
                 const float*          beta,
                 float*                C,
                 int                   ldc)

cusparseStatus_t
cusparseDcsrmm2(cusparseHandle_t          handle,
                 cusparseOperation_t    transA,
                 cusparseOperation_t    transB,
                 int                   m,
                 int                   n,
                 int                   k,
                 int                   nnz,
                 const double*          alpha,
                 const cusparseMatDescr_t descRA,
                 const double*          csrValA,
                 const int*             csrRowPtrA,
                 const int*             csrColIndA,
                 const double*          B,
                 int                   ldb,
                 const double*          beta,
                 double*               C,
                 int                   ldc)

cusparseStatus_t
cusparseCcsrmm2(cusparseHandle_t          handle,
                 cusparseOperation_t    transA,
                 cusparseOperation_t    transB,
                 int                   m,
                 int                   n,
                 int                   k,
                 int                   nnz,
                 const cuComplex*       alpha,
                 const cusparseMatDescr_t descRA,
                 const cuComplex*       csrValA,
                 const int*              csrRowPtrA,
                 const int*              csrColIndA,
                 const cuComplex*       B,
                 int                   ldb,
                 const cuComplex*       beta,
                 cuComplex*             C,
                 int                   ldc)

cusparseStatus_t
cusparseZcsrmm2(cusparseHandle_t          handle,
                 cusparseOperation_t    transA,

```

This function performs one of the following matrix-matrix operations:

$$C = \alpha * \text{op}(A) * \text{op}(B) + \beta * C$$

**A** is an  $m \times k$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **C** are dense matrices;  $\alpha$  and  $\beta$  are scalars; and

$$\text{op}(A) = \begin{cases} A & \text{if } \text{transA} == \text{CUSPARSE_OPERATION_NON\_TRANSPOSE} \\ A^T & \text{if } \text{transA} == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if } \text{transA} == \text{CUSPARSE_OPERATION_CONJUGATE\_TRANSPOSE} \end{cases}$$

and

$$\text{op}(B) = \begin{cases} B & \text{if } \text{transB} == \text{CUSPARSE_OPERATION_NON\_TRANSPOSE} \\ B^T & \text{if } \text{transB} == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ B^H & \text{not supported} \end{cases}$$

If **op(B)=B**, **cusparse<t>csrm2()** is the same as **cusparse<t>csrm()**; otherwise, only **op(A)=A** is supported and the matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**.

The motivation of **transpose(B)** is to improve the memory access of matrix **B**. The computational pattern of **A\*transpose(B)** with matrix **B** in column-major order is equivalent to **A\*B** with matrix **B** in row-major order.

In practice, no operation in iterative solver or eigenvalue solver uses **A\*transpose(B)**. However we can perform **A\*transpose(transpose(B))** which is the same as **A\*B**. For example, suppose **A** is  $m \times k$ , **B** is  $k \times n$  and **C** is  $m \times n$ , the following code shows usage of **cusparseDcsrm()**.

```
// A is m*k, B is k*n and C is m*n
const int ldb_B = k; // leading dimension of B
const int ldc_ = m; // leading dimension of C
// perform C:=alpha*A*B + beta*C
cusparseSetMatType(descrA, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseDcsrm(cusparse_handle,
              CUSPARSE_OPERATION_NON_TRANSPOSE,
              m, n, k, nnz, alpha,
              descrA, csrValA, csrRowPtrA, csrColIndA,
              B, ldb_B,
              beta, C, ldc);
```

Instead of using **A\*B**, our proposal is to transpose **B** to **Bt** first by calling **cublas<t>geam()**, then to perform **A\*transpose(Bt)**.

```
// step 1: Bt := transpose(B)
    double *Bt;
    const int ldb_Bt = n; // leading dimension of Bt
    cudaMalloc((void**)&Bt, sizeof(double)*ldb_Bt*k);
    double one = 1.0;
    double zero = 0.0;
    cublasSetPointerMode(cublas_handle, CUBLAS_POINTER_MODE_HOST);
    cublasDgemm(cublas_handle, CUBLAS_OP_T, CUBLAS_OP_T,
        n, k, &one, B, int ldb_B, &zero, B, int ldb_Bt, Bt, ldb_Bt);

// step 2: perform C:=alpha*A*transpose(Bt) + beta*C
    cusparseDcsrmm2(cusparse_handle,
        CUSPARSE_OPERATION_NON_TRANSPOSE,
        CUSPARSE_OPERATION_TRANSPOSE
        m, n, k, nnz, alpha,
        descrA, csrValA, csrRowPtrA, csrColIndA,
        Bt, ldb_Bt,
        beta, C, ldc);
```

Remark: **cublas<t>geam()** and **cusparse<t>csrmm2()** are memory bound. The complexity of **cublas<t>geam()** is  $2*n*k$ , and the minimum complexity of **cusparse<t>csrmm2()** is about  $(nnz + nnz*n + 2*m*n)$ . If  $nnz$  per column ( $=nnz/k$ ) is large, it is worth paying the extra cost on transposition because **A\*transpose(B)** may be 2x faster than **A\*B** if the sparsity pattern of **A** is not good.

The function has the following properties for operation **CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and the matrix type is different from **CUSPARSE\_MATRIX\_TYPE\_GENERAL**

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

otherwise:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$
<b>transB</b>	the operation $\text{op}(B)$
<b>m</b>	number of rows of sparse matrix <b>A</b> .
<b>n</b>	number of columns of dense matrix $\text{op}(B)$ and <b>c</b> .
<b>k</b>	number of columns of sparse matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of sparse matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are

	CUSPARSE_INDEX_BASE_ZERO and CUSPARSE_INDEX_BASE_ONE.
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>B</b>	array of dimensions <b>(ldb, n)</b> if <b>op(B)=B</b> and <b>(ldb, k)</b> otherwise.
<b>ldb</b>	leading dimension of <b>B</b> . If <b>op(B)=B</b> , it must be at least <b>max(1, k)</b> if <b>op(A)=A</b> and at least <b>max(1, m)</b> otherwise. If <b>op(B) !=B</b> , it must be at least <b>max(1, n)</b> .
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>c</b> does not have to be a valid input.
<b>C</b>	array of dimensions <b>(ldc, n)</b> .
<b>ldc</b>	leading dimension of <b>c</b> . It must be at least <b>max(1, m)</b> if <b>op(A)=A</b> and at least <b>max(1, k)</b> otherwise.

## Output

<b>c</b>	<type> updated array of dimensions <b>(ldc, n)</b> .
----------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 9.8. `cusparse<t>csrsm_analysis()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrsm2_analysis()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrsm_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const float*           csrSortedValA,
                        const int*              csrSortedRowPtrA,
                        const int*              csrSortedColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseDcsrsm_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const double*          csrSortedValA,
                        const int*              csrSortedRowPtrA,
                        const int*              csrSortedColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseCcsrsm_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const cuComplex*        csrSortedValA,
                        const int*              csrSortedRowPtrA,
                        const int*              csrSortedColIndA,
                        cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseZcsrsm_analysis(cusparseHandle_t          handle,
                        cusparseOperation_t    transA,
                        int                   m,
                        int                   nnz,
                        const cusparseMatDescr_t descrA,
                        const cuDoubleComplex*  csrSortedValA,
                        const int*              csrSortedRowPtrA,
                        const int*              csrSortedColIndA,
                        cusparseSolveAnalysisInfo_t info)

```

This function performs the analysis phase of the solution of a sparse triangular linear system

$$\text{op}(A) * \mathbf{X} = \alpha * \mathbf{B}$$

with multiple right-hand sides, where **A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **x** are the right-hand-side and the solution dense matrices;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

It is expected that this function will be executed only once for a given matrix and a particular operation type.

- ▶ This function requires a significant amount of extra storage that is proportional to the matrix size
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure initialized using <b>cusparseCreateSolveAnalysisInfo()</b> .

### Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See **cusparseStatus\_t** for the description of the return status

## 9.9. `cusparse<t>csrsm_solve()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrsm2_solve()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrsm_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     int                      n,
                     const float*             alpha,
                     const cusparseMatDescr_t descrA,
                     const float*             csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const float*             B,
                     int                      ldb,
                     float*                  X,
                     int                      ldx)

cusparseStatus_t
cusparseDcsrsm_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     int                      n,
                     const double*            alpha,
                     const cusparseMatDescr_t descrA,
                     const double*            csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const double*            B,
                     int                      ldb,
                     double*                 X,
                     int                      ldx)

cusparseStatus_t
cusparseCcsrsm_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     int                      n,
                     const cuComplex*         alpha,
                     const cusparseMatDescr_t descrA,
                     const cuComplex*         csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuComplex*         B,
                     int                      ldb,
                     cuComplex*              X,
                     int                      ldx)

cusparseStatus_t
cusparseZcsrsm_solve(cusparseHandle_t           handle,
                     cusparseOperation_t      transA,
                     int                      m,
                     int                      n,
                     const cuDoubleComplex*   alpha,
                     const cusparseMatDescr_t descrA,
                     const cuDoubleComplex*   csrSortedValA,
                     const int*               csrSortedRowPtrA,
                     const int*               csrSortedColIndA,
                     cusparseSolveAnalysisInfo_t info,
                     const cuDoubleComplex*   B,
```

This function performs the solve phase of the solution of a sparse triangular linear system

$$\text{op}(A) * \mathbf{X} = \alpha * \mathbf{B}$$

with multiple right-hand sides, where **A** is an  $m \times n$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **x** are the right-hand-side and the solution dense matrices;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

This function may be executed multiple times for a given matrix and a particular operation type.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>transA</b>	the operation <b>op(A)</b> .
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>x</b> and <b>y</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_TRIANGULAR</b> and <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure with information collected during the analysis phase (that should be passed to the solve phase unchanged).
<b>B</b>	<type> right-hand-side array of dimensions ( <b>ldb</b> , <b>n</b> ).
<b>ldb</b>	leading dimension of <b>B</b> (that is $\geq \max(1, m)$ ).

## Output

<b>x</b>	<type> solution array of dimensions ( <b>idx</b> , <b>n</b> ).
----------	--

<b>lidx</b>	leading dimension of <b>x</b> (that is $\geq \max(1, m)$ ).
-------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 9.10. `cusparse<t>csrsm2_bufferSizeExt()`

This function returns the size of the buffer used in **csrsm2**, a sparse triangular linear system  $\text{op}(\mathbf{A}) * \text{op}(\mathbf{x}) = \alpha \text{op}(\mathbf{B})$ .

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **x** are the right-hand-side matrix and the solution matrix;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>algo</b>	algo = 0 is non-block version; algo = 1 is block version.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>transB</b>	the operation $\text{op}(B)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nrhs</b>	number of columns of right hand side matrix $\text{op}(\mathbf{B})$ .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>B</b>	<type> right-hand-side matrix. $\text{op}(\mathbf{B})$ is of size <b>m</b> -by- <b>nrhs</b> .
<b>ldb</b>	leading dimension of <b>B</b> and <b>x</b> .
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<b>policy</b>	The supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .

## Output

<b>info</b>	record of internal states based on different algorithms.
<b>pBufferSize</b>	number of bytes of the buffer used in the <code>csrsm2_analysis</code> and <code>csrsm2_solve</code> .

See `cusparseStatus_t` for the description of the return status

## 9.11. `cusparse<t>csrsm2_analysis()`

```

cusparseStatus_t
cusparseScsrsm2_analysis(cusparseHandle_t      handle,
                           int                algo,
                           int                transA,
                           cusparseOperation_t transB,
                           int                m,
                           int                nrhs,
                           int                nnz,
                           const float*       alpha,
                           const cusparseMatDescr_t descrA,
                           const float*       csrSortedValA,
                           const int*         csrSortedRowPtrA,
                           const int*         csrSortedColIndA,
                           const float*       B,
                           int                ldb,
                           csrsm2Info_t       info,
                           cusparseSolvePolicy_t policy,
                           void*              pBuffer)

cusparseStatus_t
cusparseDcsrsm2_analysis(cusparseHandle_t      handle,
                           int                algo,
                           int                transA,
                           cusparseOperation_t transB,
                           int                m,
                           int                nrhs,
                           int                nnz,
                           const double*      alpha,
                           const cusparseMatDescr_t descrA,
                           const double*      csrSortedValA,
                           const int*         csrSortedRowPtrA,
                           const int*         csrSortedColIndA,
                           const double*      B,
                           int                ldb,
                           csrsm2Info_t       info,
                           cusparseSolvePolicy_t policy,
                           void*              pBuffer)

cusparseStatus_t
cusparseCcsrsm2_analysis(cusparseHandle_t      handle,
                           int                algo,
                           int                transA,
                           cusparseOperation_t transB,
                           int                m,
                           int                nrhs,
                           int                nnz,
                           const cuComplex*   alpha,
                           const cusparseMatDescr_t descrA,
                           const cuComplex*   csrSortedValA,
                           const int*         csrSortedRowPtrA,
                           const int*         csrSortedColIndA,
                           const cuComplex*   B,
                           int                ldb,
                           csrsm2Info_t       info,
                           cusparseSolvePolicy_t policy,
                           void*              pBuffer)

cusparseStatus_t
cusparseZcsrsm2_analysis(cusparseHandle_t      handle,
                           int                algo,
                           int                transA,
                           cusparseOperation_t transB,
                           int                m,

```

This function performs the analysis phase of **csrsm2**, a sparse triangular linear system  $\text{op}(\mathbf{A}) * \text{op}(\mathbf{x}) = \alpha \text{op}(\mathbf{B})$ .

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **x** are the right-hand-side matrix and the solution matrix;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans == CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if trans == CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

It is expected that this function will be executed only once for a given matrix and a particular operation type.

This function requires a buffer size returned by **csrsm2\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **csrsm2\_analysis()** reports a structural zero and computes level information that is stored in opaque structure **info**. The level information can extract more parallelism for a triangular solver. However **csrsm2\_solve()** can be done without level information. To disable level information, the user needs to specify the policy of the triangular solver as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **csrsm2\_analysis()** always reports the first structural zero, even if the policy is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. No structural zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if  $\mathbf{A}(j, j)$  is missing for some  $j$ . The user needs to call **cusparseXcsrsm2\_zeroPivot()** to know where the structural zero is.

It is the user's choice whether to call **csrsm2\_solve()** if **csrsm2\_analysis()** reports a structural zero. In this case, the user can still call **csrsm2\_solve()** which will return a numerical zero in the same position as the structural zero. However the result **x** is meaningless.

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>algo</b>	algo = 0 is non-block version; algo = 1 is block version.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>transB</b>	the operation $\text{op}(B)$ .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nrhs</b>	number of columns of matrix $\text{op}(\mathbf{B})$ .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .

<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>B</b>	<type> right-hand-side matrix. <b>op(B)</b> is of size <b>m-by-nrhs</b> .
<b>ldb</b>	leading dimension of <b>B</b> and <b>x</b> .
<b>info</b>	structure initialized using <b>cusparseCreateCsrsv2Info()</b> .
<b>policy</b>	The supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is returned by <b>csrsm2_bufferSize()</b> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 9.12. `cusparse<t>csrsm2_solve()`

```

cusparseStatus_t
cusparseScsrsm2_solve(cusparseHandle_t          handle,
                      int                  algo,
                      int                  transA,
                      cusparseOperation_t   transB,
                      int                  m,
                      int                  nrhs,
                      int                  nnz,
                      const float*          alpha,
                      const cusparseMatDescr_t descrA,
                      const float*          csrSortedValA,
                      const int*            csrSortedRowPtrA,
                      const int*            csrSortedColIndA,
                      float*                B,
                      int                  ldb,
                      csrsm2Info_t          info,
                      cusparseSolvePolicy_t policy,
                      void*                pBuffer)

cusparseStatus_t
cusparseDcsrsm2_solve(cusparseHandle_t          handle,
                      int                  algo,
                      int                  transA,
                      transB,
                      int                  m,
                      int                  nrhs,
                      int                  nnz,
                      const double*         alpha,
                      const cusparseMatDescr_t descrA,
                      const double*         csrSortedValA,
                      const int*            csrSortedRowPtrA,
                      const int*            csrSortedColIndA,
                      double*               B,
                      int                  ldb,
                      csrsm2Info_t          info,
                      cusparseSolvePolicy_t policy,
                      void*                pBuffer)

cusparseStatus_t
cusparseCcsrsm2_solve(cusparseHandle_t          handle,
                      int                  algo,
                      transA,
                      transB,
                      int                  m,
                      int                  nrhs,
                      int                  nnz,
                      const cuComplex*      alpha,
                      const cusparseMatDescr_t descrA,
                      const cuComplex*      csrSortedValA,
                      const int*            csrSortedRowPtrA,
                      const int*            csrSortedColIndA,
                      cuComplex*            B,
                      int                  ldb,
                      csrsm2Info_t          info,
                      cusparseSolvePolicy_t policy,
                      void*                pBuffer)

cusparseStatus_t
cusparseZcsrsm2_solve(cusparseHandle_t          handle,
                      int                  algo,
                      transA,
                      transB,
                      int                  m,

```

This function performs the solve phase of **csrsm2**, a sparse triangular linear system  $\text{op}(\mathbf{A}) * \text{op}(\mathbf{x}) = \alpha \text{op}(\mathbf{B})$ .

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**; **B** and **x** are the right-hand-side matrix and the solution matrix;  $\alpha$  is a scalar; and

$$\text{op}(A) = \begin{cases} A & \text{if } \text{transA} == \text{CUSPARSE\_OPERATION\_NON\_TRANSPOSE} \\ A^T & \text{if } \text{transA} == \text{CUSPARSE\_OPERATION\_TRANSPOSE} \\ A^H & \text{if } \text{transA} == \text{CUSPARSE\_OPERATION\_CONJUGATE\_TRANSPOSE} \end{cases}$$

**transB** acts on both matrix **B** and matrix **x**, only

**CUSPARSE\_OPERATION\_NON\_TRANSPOSE** and **CUSPARSE\_OPERATION\_TRANSPOSE**. The operation is in-place, matrix **B** is overwritten by matrix **x**.

**ldb** must be not less than **m** if **transB = CUSPARSE\_OPERATION\_NON\_TRANSPOSE**. Otherwise, **ldb** must be not less than **nrhs**.

This function requires the buffer size returned by **csrsm2\_bufferSize()**.

The address of **pBuffer** must be multiple of 128 bytes. If it is not,

**CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **csrsm2\_solve()** can be done without level information, the user still needs to be aware of consistency. If **csrsm2\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **csrsm2\_solve()** can be run with or without levels. On the contrary, if **csrsm2\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **csrsm2\_solve()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

The level information may not improve the performance but spend extra time doing analysis. For example, a tridiagonal matrix has no parallelism.

In this case, **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** performs better than **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**. If the user has an iterative solver, the best approach is to do **csrsm2\_analysis()** with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** once. Then do **csrsm2\_solve()** with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL** in the first run and with **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL** in the second run, picking faster one to perform the remaining iterations.

Function **csrsm2\_solve()** reports the first numerical zero, including a structural zero. If **status** is 0, no numerical zero was found. Furthermore, no numerical zero is reported if **CUSPARSE\_DIAG\_TYPE\_UNIT** is specified, even if **A(j, j)** is zero for some **j**. The user needs to call **cusparseXcsrsm2\_zeroPivot()** to know where the numerical zero is.

**csrsm2** provides two algorithms specified by the parameter **algo**. **algo=0** is non-block version and **algo=1** is block version. non-block version is memory-bound, limited by bandwidth. block version partitions the matrix into small tiles and applies dense operations. Although it has more flops than non-block version, it may be faster if non-block version already reaches maximum bandwidth..

Appendix H shows an example of **csrsm2**.

The function supports the following properties if **pBuffer != NULL**

- ▶ The routine requires no extra storage

- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

**Input**

<b>handle</b>	handle to the cuSPARSE library context.
<b>algo</b>	algo = 0 is non-block version; algo = 1 is block version.
<b>transA</b>	the operation $\text{op}(A)$ .
<b>transB</b>	the operation $\text{op}(B)$ .
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>nrhs</b>	number of columns of matrix $\text{op}(B)$ .
<b>nnz</b>	number of nonzeros of matrix <b>A</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , while the supported diagonal types are <b>CUSPARSE_DIAG_TYPE_UNIT</b> and <b>CUSPARSE_DIAG_TYPE_NON_UNIT</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> ( $= \text{csrRowPtrA}(m) - \text{csrRowPtrA}(0)$ ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m</b> + 1 elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> ( $= \text{csrRowPtrA}(m) - \text{csrRowPtrA}(0)$ ) column indices of the nonzero elements of matrix <b>A</b> .
<b>B</b>	<type> right-hand-side matrix. $\text{op}(B)$ is of size <b>m</b> -by- <b>nrhs</b> .
<b>ldb</b>	leading dimension of <b>B</b> and <b>x</b> .
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<b>policy</b>	The supported policies are <b>CUSPARSE_SOLVE_POLICY_NO_LEVEL</b> and <b>CUSPARSE_SOLVE_POLICY_USE_LEVEL</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is returned by <b>csrsm2_bufferSize</b> .

**Output**

<b>x</b>	<type> solution matrix, $\text{op}(x)$ is of size <b>m</b> -by- <b>nrhs</b> .
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 9.13. `cusparseXcsrsm2_zeroPivot()`

```
cusparseStatus_t
cusparseXcsrsm2_zeroPivot(cusparseHandle_t handle,
                           csrsm2Info_t    info,
                           int*            position)
```

If the returned error code is **CUSPARSE\_STATUS\_ZERO\_PIVOT**, `position=j` means  $A(j,j)$  has either a structural zero or a numerical zero. Otherwise `position=-1`.

The `position` can be 0-based or 1-based, the same as the matrix.

Function `cusparseXcsrsm2_zeroPivot()` is a blocking call. It calls `cudaDeviceSynchronize()` to make sure all previous kernels are done.

The `position` can be in the host memory or device memory. The user can set the proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	<code>info</code> contains structural zero or numerical zero if the user already called <code>csrsm2_analysis()</code> or <code>csrsm2_solve()</code> .

### Output

<code>position</code>	if no structural or numerical zero, <code>position</code> is -1; otherwise, if $A(j,j)$ is missing or $U(j,j)$ is zero, <code>position=j</code> .
-----------------------	---

See `cusparseStatus_t` for the description of the return status

## 9.14. `cusparse<t>gemmi()`

```

cusparseStatus_t
cusparseSgemmi(cusparseHandle_t handle,
               int          m,
               int          n,
               int          k,
               int          nnz,
               const float* alpha,
               const float* A,
               int          lda,
               const float* cscValB,
               const int*   cscColPtrB,
               const int*   cscRowIndB,
               const float* beta,
               float*       C,
               int          ldc)

cusparseStatus_t
cusparseDgemmi(cusparseHandle_t handle,
               int          m,
               int          n,
               int          k,
               int          nnz,
               const double* alpha,
               const double* A,
               int          lda,
               const double* cscValB,
               const int*   cscColPtrB,
               const int*   cscRowIndB,
               const double* beta,
               double*      C,
               int          ldc)

cusparseStatus_t
cusparseCgemmi(cusparseHandle_t handle,
               int          m,
               int          n,
               int          k,
               int          nnz,
               const cuComplex* alpha,
               const cuComplex* A,
               int          lda,
               const cuComplex* cscValB,
               const int*   cscColPtrB,
               const int*   cscRowIndB,
               const cuComplex* beta,
               cuComplex*   C,
               int          ldc)

cusparseStatus_t
cusparseZgemmi(cusparseHandle_t handle,
               int          m,
               int          n,
               int          k,
               int          nnz,
               const cuDoubleComplex* alpha,
               const cuDoubleComplex* A,
               int          lda,
               const cuDoubleComplex* cscValB,
               const int*   cscColPtrB,
               const int*   cscRowIndB,
               const cuDoubleComplex* beta,
               cuDoubleComplex* C,
               int          ldc)

```

This function performs the following matrix-matrix operations:

$$C = \alpha * A * B + \beta * C$$

**A** and **C** are dense matrices; **B** is a  $k \times n$  sparse matrix that is defined in CSC storage format by the three arrays **cscValB**, **cscColPtrB**, and **cscRowIndB**);  $\alpha$  and  $\beta$  are scalars; and

Remark: **B** is base-0.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrices <b>B</b> and <b>C</b> .
<b>k</b>	number of columns of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of sparse matrix <b>B</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>A</b>	array of dimensions ( <b>lda</b> , <b>k</b> ).
<b>lda</b>	leading dimension of <b>A</b> . It must be at least <b>m</b>
<b>cscValB</b>	<type> array of <b>nnz</b> (= <b>cscColPtrB(k)</b> - <b>cscColPtrB(0)</b> ) nonzero elements of matrix <b>B</b> .
<b>cscColPtrB</b>	integer array of <b>k</b> + 1 elements that contains the start of every row and the end of the last row plus one.
<b>cscRowIndB</b>	integer array of <b>nnz</b> (= <b>cscColPtrB(k)</b> - <b>cscColPtrB(0)</b> ) column indices of the nonzero elements of matrix <b>B</b> .
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>C</b> does not have to be a valid input.
<b>C</b>	array of dimensions ( <b>ldc</b> , <b>n</b> ).
<b>ldc</b>	leading dimension of <b>C</b> . It must be at least <b>m</b>

### Output

<b>C</b>	<type> updated array of dimensions ( <b>ldc</b> , <b>n</b> ).
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 10.

# CUSPARSE EXTRA FUNCTION REFERENCE

This chapter describes the extra routines used to manipulate sparse matrices.

## 10.1. `cusparse<t>csrgeam()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrgeam2()` instead. *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseXcsrgeamNnz(cusparseHandle_t handle,
                     int m,
                     int n,
                     const cusparseMatDescr_t descrA,
                     int nnzA,
                     const int* csrRowPtrA,
                     const int* csrColIndA,
                     const cusparseMatDescr_t descrB,
                     int nnzB,
                     const int* csrRowPtrB,
                     const int* csrColIndB,
                     const cusparseMatDescr_t descrC,
                     int* csrRowPtrC,
                     int* nnzTotalDevHostPtr)
```

```
cusparseStatus_t
cusparseScsrgeam(cusparseHandle_t handle,
                  int m,
                  int n,
                  const float* alpha,
                  const cusparseMatDescr_t descrA,
                  int nnzA,
                  const float* csrValA,
                  const int* csrRowPtrA,
                  const int* csrColIndA,
                  const float* beta,
                  const cusparseMatDescr_t descrB,
                  int nnzB,
                  const float* csrValB,
                  const int* csrRowPtrB,
                  const int* csrColIndB,
                  const cusparseMatDescr_t descrC,
                  float* csrValC,
                  int* csrRowPtrC,
                  int* csrColIndC)
```

```
cusparseStatus_t
cusparseDcsrgeam(cusparseHandle_t handle,
                  int m,
                  int n,
                  const double* alpha,
                  const cusparseMatDescr_t descrA,
                  int nnzA,
                  const double* csrValA,
                  const int* csrRowPtrA,
                  const int* csrColIndA,
                  const double* beta,
                  const cusparseMatDescr_t descrB,
                  int nnzB,
                  const double* csrValB,
                  const int* csrRowPtrB,
                  const int* csrColIndB,
                  const cusparseMatDescr_t descrC,
                  double* csrValC,
                  int* csrRowPtrC,
                  int* csrColIndC)
```

```
cusparseStatus_t
cusparseCcsrgeam(cusparseHandle_t handle,
```

This function performs following matrix-matrix operation

$$C = \alpha * A + \beta * B$$

where **A**, **B**, and **C** are  $m \times n$  sparse matrices (defined in CSR storage format by the three arrays `csrValA`|`csrValB`|`csrValC`, `csrRowPtrA`|`csrRowPtrB`|`csrRowPtrC`, and `csrColIndA`|`csrColIndB`|`csrColIndC` respectively), and  $\alpha$  and  $\beta$  are scalars. Since **A** and **B** have different sparsity patterns, cuSPARSE adopts a two-step approach to complete sparse matrix **C**. In the first step, the user allocates `csrRowPtrC` of  $m + 1$  elements and uses function `cusparseXcsrgeamNnz()` to determine `csrRowPtrC` and the total number of nonzero elements. In the second step, the user gathers `nnzC` (number of nonzero elements of matrix **C**) from either (`nnzC = nnzTotalDevHostPtr`) or (`nnzC = csrRowPtrC(m) - csrRowPtrC(0)`) and allocates `csrValC`, `csrColIndC` of `nnzC` elements respectively, then finally calls function `cusparse[S|D|C|Z]csrgeam()` to complete matrix **C**.

The general procedure is as follows:

```
int baseC, nnzC;
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzC;
cusparseSetPointerMode(handle, CUSPARSE_POINTER_MODE_HOST);
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
cusparseXcsrgeamNnz(handle, m, n,
                    descrA, nnzA, csrRowPtrA, csrColIndA,
                    descrB, nnzB, csrRowPtrB, csrColIndB,
                    descrC, csrRowPtrC, nnzTotalDevHostPtr);
if (NULL != nnzTotalDevHostPtr){
    nnzC = *nnzTotalDevHostPtr;
} else {
    cudaMemcpy(&nnzC, csrRowPtrC+m, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&baseC, csrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzC -= baseC;
}
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnzC);
cudaMalloc((void**)&csrValC, sizeof(float)*nnzC);
cusparseScsrgeam(handle, m, n,
                 alpha,
                 descrA, nnzA,
                 csrValA, csrRowPtrA, csrColIndA,
                 beta,
                 descrB, nnzB,
                 csrValB, csrRowPtrB, csrColIndB,
                 descrC,
                 csrValC, csrRowPtrC, csrColIndC);
```

Several comments on `csrgeam()`:

- ▶ The other three combinations, NT, TN, and TT, are not supported by cuSPARSE. In order to do any one of the three, the user should use the routine `csr2csc()` to convert  $A|B$  to  $A^T|B^T$ .
- ▶ Only `CUSPARSE_MATRIX_TYPE_GENERAL` is supported. If either **A** or **B** is symmetric or Hermitian, then the user must extend the matrix to a full one and reconfigure the `MatrixType` field of the descriptor to `CUSPARSE_MATRIX_TYPE_GENERAL`.
- ▶ If the sparsity pattern of matrix **C** is known, the user can skip the call to function `cusparseXcsrgeamNnz()`. For example, suppose that the user has an iterative algorithm which would update **A** and **B** iteratively but keep the sparsity patterns.

The user can call function **`cusparseXcsrgeamNnz()`** once to set up the sparsity pattern of **C**, then call function **`cusparse[S|D|C|Z]geam()`** only for each iteration.

- ▶ The pointers **alpha** and **beta** must be valid.
- ▶ When **alpha** or **beta** is zero, it is not considered a special case by cuSPARSE. The sparsity pattern of **C** is independent of the value of **alpha** and **beta**. If the user wants  $C = 0 \times A + 1 \times B^T$ , then **`csr2csc()`** is better than **`csrgeam()`**.
- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of sparse matrix <b>A,B,C</b> .
<b>n</b>	number of columns of sparse matrix <b>A,B,C</b> .
<b>alpha</b>	<type> scalar used for multiplication.
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> only.
<b>nnzA</b>	number of nonzero elements of sparse matrix <b>A</b> .
<b>csrValA</b>	<type> array of <b>nnzA</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnzA</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>beta</b>	<type> scalar used for multiplication. If <b>beta</b> is zero, <b>y</b> does not have to be a valid input.
<b>descrB</b>	the descriptor of matrix <b>B</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> only.
<b>nnzB</b>	number of nonzero elements of sparse matrix <b>B</b> .
<b>csrValB</b>	<type> array of <b>nnzB</b> (= <b>csrRowPtrB(m) - csrRowPtrB(0)</b> ) nonzero elements of matrix <b>B</b> .
<b>csrRowPtrB</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndB</b>	integer array of <b>nnzB</b> (= <b>csrRowPtrB(m) - csrRowPtrB(0)</b> ) column indices of the nonzero elements of matrix <b>B</b> .
<b>descrC</b>	the descriptor of matrix <b>C</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> only.

### Output

<code>csrValC</code>	<type> array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) nonzero elements of matrix c.
<code>csrRowPtrC</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndC</code>	integer array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) column indices of the nonzero elements of matrixc.
<code>nnzTotalDevHostPtr</code>	total number of nonzero elements in device or host memory. It is equal to $(\text{csrRowPtrC}(m) - \text{csrRowPtrC}(0))$ .

See `cusparseStatus_t` for the description of the return status



This function performs following matrix-matrix operation

$$C = \alpha * A + \beta * B$$

where **A**, **B**, and **C** are  $m \times n$  sparse matrices (defined in CSR storage format by the three arrays **csrValA**|**csrValB**|**csrValC**, **csrRowPtrA**|**csrRowPtrB**|**csrRowPtrC**, and **csrColIndA**|**csrColIndB**|**csrColIndC** respectively), and  $\alpha$  and  $\beta$  are scalars. Since **A** and **B** have different sparsity patterns, cuSPARSE adopts a two-step approach to complete sparse matrix **C**. In the first step, the user allocates **csrRowPtrC** of  $m+1$  elements and uses function **cusparseXcsrgeam2Nnz()** to determine **csrRowPtrC** and the total number of nonzero elements. In the second step, the user gathers **nnzC** (number of nonzero elements of matrix **C**) from either (**nnzC=\*nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC(m)-csrRowPtrC(0)**) and allocates **csrValC**, **csrColIndC** of **nnzC** elements respectively, then finally calls function **cusparse[S|D|C|Z]csrgeam2()** to complete matrix **C**.

The general procedure is as follows:

```
int baseC, nnzC;
/* alpha, nnzTotalDevHostPtr points to host memory */
size_t bufferSizeInBytes;
char *buffer = NULL;
int *nnzTotalDevHostPtr = &nnzC;
cusparseSetPointerMode(handle, CUSPARSE_POINTER_MODE_HOST);
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
/* prepare buffer */
cusparseScsrgeam2_bufferSizeExt(handle, m, n,
    alpha,
    descrA, nnzA,
    csrValA, csrRowPtrA, csrColIndA,
    beta,
    descrB, nnzB,
    csrValB, csrRowPtrB, csrColIndB,
    descrC,
    csrValC, csrRowPtrC, csrColIndC
    &bufferSizeInBytes
);
cudaMalloc((void**)&buffer, sizeof(char)*bufferSizeInBytes);
cusparseXcsrgeam2Nnz(handle, m, n,
    descrA, nnzA, csrRowPtrA, csrColIndA,
    descrB, nnzB, csrRowPtrB, csrColIndB,
    descrC, csrRowPtrC, nnzTotalDevHostPtr,
    buffer);
if (NULL != nnzTotalDevHostPtr){
    nnzC = *nnzTotalDevHostPtr;
} else {
    cudaMemcpy(&nnzC, csrRowPtrC+m, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&baseC, csrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzC -= baseC;
}
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnzC);
cudaMalloc((void**)&csrValC, sizeof(float)*nnzC);
cusparseScsrgeam2(handle, m, n,
    alpha,
    descrA, nnzA,
    csrValA, csrRowPtrA, csrColIndA,
    beta,
    descrB, nnzB,
    csrValB, csrRowPtrB, csrColIndB,
    descrC,
    csrValC, csrRowPtrC, csrColIndC
    buffer);
```

Several comments on `csrgeam2()`:

- ▶ The other three combinations, NT, TN, and TT, are not supported by cuSPARSE. In order to do any one of the three, the user should use the routine `csr2csc()` to convert  $A \mid B$  to  $A^T \mid B^T$ .
- ▶ Only `CUSPARSE_MATRIX_TYPE_GENERAL` is supported. If either **A** or **B** is symmetric or Hermitian, then the user must extend the matrix to a full one and reconfigure the `MatrixType` field of the descriptor to `CUSPARSE_MATRIX_TYPE_GENERAL`.
- ▶ If the sparsity pattern of matrix **C** is known, the user can skip the call to function `cusparseXcsrgeam2Nnz()`. For example, suppose that the user has an iterative algorithm which would update **A** and **B** iteratively but keep the sparsity patterns. The user can call function `cusparseXcsrgeam2Nnz()` once to set up the sparsity pattern of **C**, then call function `cusparse[S|D|C|Z]geam()` only for each iteration.
- ▶ The pointers `alpha` and `beta` must be valid.
- ▶ When `alpha` or `beta` is zero, it is not considered a special case by cuSPARSE. The sparsity pattern of **C** is independent of the value of `alpha` and `beta`. If the user wants  $C = 0 \times A + 1 \times B^T$ , then `csr2csc()` is better than `csrgeam2()`.
- ▶ `csrgeam2()` is the same as `csrgeam()` except `csrgeam2()` needs explicit buffer where `csrgeam()` allocates the buffer internally.
- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of sparse matrix <b>A</b> , <b>B</b> , <b>C</b> .
<code>n</code>	number of columns of sparse matrix <b>A</b> , <b>B</b> , <b>C</b> .
<code>alpha</code>	<type> scalar used for multiplication.
<code>descrA</code>	the descriptor of matrix <b>A</b> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzA</code>	number of nonzero elements of sparse matrix <b>A</b> .
<code>csrValA</code>	<type> array of <code>nnzA</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <b>A</b> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnzA</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <b>A</b> .
<code>beta</code>	<type> scalar used for multiplication. If <code>beta</code> is zero, <code>y</code> does not have to be a valid input.
<code>descrB</code>	the descriptor of matrix <b>B</b> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzB</code>	number of nonzero elements of sparse matrix <b>B</b> .

<code>csrValB</code>	<type> array of <code>nnzB</code> ( $= \text{csrRowPtrB}(m) - \text{csrRowPtrB}(0)$ ) nonzero elements of matrix B.
<code>csrRowPtrB</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndB</code>	integer array of <code>nnzB</code> ( $= \text{csrRowPtrB}(m) - \text{csrRowPtrB}(0)$ ) column indices of the nonzero elements of matrix B.
<code>descrC</code>	the descriptor of matrix c. The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.

## Output

<code>csrValC</code>	<type> array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) nonzero elements of matrix c.
<code>csrRowPtrC</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndC</code>	integer array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) column indices of the nonzero elements of matrixc.
<code>nnzTotalDevHostPtr</code>	total number of nonzero elements in device or host memory. It is equal to <code>(csrRowPtrC(m) - csrRowPtrC(0))</code> .

See `cusparseStatus_t` for the description of the return status

## 10.3. `cusparse<t>csrgemm()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrgemm()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseXcsrgemmNnz(cusparseHandle_t          handle,
                     cusparseOperation_t    transA,
                     cusparseOperation_t    transB,
                     int                   m,
                     int                   n,
                     int                   k,
                     const cusparseMatDescr_t descrA,
                     int                   nnzA,
                     const int*            csrRowPtrA,
                     const int*            csrColIndA,
                     const cusparseMatDescr_t descrB,
                     int                   nnzB,
                     const int*            csrRowPtrB,
                     const int*            csrColIndB,
                     const cusparseMatDescr_t descrC,
                     int*                 csrRowPtrC,
                     int*                 nnzTotalDevHostPtr)

cusparseStatus_t
cusparseScsrgemm(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  cusparseOperation_t    transB,
                  int                   m,
                  int                   n,
                  int                   k,
                  const cusparseMatDescr_t descrA,
                  int                   nnzA,
                  const float*          csrValA,
                  const int*             csrRowPtrA,
                  const int*             csrColIndA,
                  const cusparseMatDescr_t descrB,
                  int                   nnzB,
                  const float*          csrValB,
                  const int*             csrRowPtrB,
                  const int*             csrColIndB,
                  const cusparseMatDescr_t descrC,
                  float*                csrValC,
                  const int*             csrRowPtrC,
                  int*                 csrColIndC)

cusparseStatus_t
cusparseDcsrgemm(cusparseHandle_t          handle,
                  cusparseOperation_t    transA,
                  cusparseOperation_t    transB,
                  int                   m,
                  int                   n,
                  int                   k,
                  const cusparseMatDescr_t descrA,
                  int                   nnzA,
                  const double*          csrValA,
                  const int*              csrRowPtrA,
                  const int*              csrColIndA,
                  const cusparseMatDescr_t descrB,
                  int                   nnzB,
                  const double*          csrValB,
                  const int*              csrRowPtrB,
                  const int*              csrColIndB,
                  const cusparseMatDescr_t descrC,
                  double*               csrValC.

```

This function performs following matrix-matrix operation:

$$C = \text{op}(A) * \text{op}(B)$$

where  $\text{op}(A)$ ,  $\text{op}(B)$  and  $C$  are  $m \times k$ ,  $k \times n$ , and  $m \times n$  sparse matrices (defined in CSR storage format by the three arrays **csrValA**|**csrValB**|**csrValC**, **csrRowPtrA**|**csrRowPtrB**|**csrRowPtrC**, and **csrColIndA**|**csrColIndB**|**csrColIndC** respectively. The operation is defined by

$$\text{op}(A) = \begin{cases} A & \text{if trans == CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if trans != CUSPARSE_OPERATION_NON_TRANSPOSE} \end{cases}$$

There are four versions, NN, NT, TN, and TT. NN stands for  $C = A^* B$ , NT stands for  $C = A^* B^T$ , TN stands for  $C = A^T * B$  and TT stands for  $C = A^T * B^T$ .

The cuSPARSE library adopts a two-step approach to complete sparse matrix. In the first step, the user allocates **csrRowPtrC** of  $m+1$  elements and uses the function **cusparseXcsrgemmNnz()** to determine **csrRowPtrC** and the total number of nonzero elements. In the second step, the user gathers **nnzC** (the number of nonzero elements of matrix **C**) from either (**nnzC==nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC(m)-csrRowPtrC(0)**) and allocates **csrValC** and **csrColIndC** of **nnzC** elements respectively, then finally calls function **cusparse[S|D|C|Z]csrgemm()** to complete matrix **C**.

The general procedure is as follows:

```
int baseC, nnzC;
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzC;
cusparseSetPointerMode(handle, CUSPARSE_POINTER_MODE_HOST);
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
cusparseXcsrgemmNnz(handle, transA, transB, m, n, k,
                     descrA, nnzA, csrRowPtrA, csrColIndA,
                     descrB, nnzB, csrRowPtrB, csrColIndB,
                     descrC, csrRowPtrC, nnzTotalDevHostPtr );
if (NULL != nnzTotalDevHostPtr){
    nnzC = *nnzTotalDevHostPtr;
} else{
    cudaMemcpy(&nnzC, csrRowPtrC+m, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&baseC, csrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzC -= baseC;
}
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnzC);
cudaMalloc((void**)&csrValC, sizeof(float)*nnzC);
cusparseScsrsgemm(handle, transA, transB, m, n, k,
                  descrA, nnzA,
                  csrValA, csrRowPtrA, csrColIndA,
                  descrB, nnzB,
                  csrValB, csrRowPtrB, csrColIndB,
                  descrC,
                  csrValC, csrRowPtrC, csrColIndC);
```

Several comments on **csrsgemm()**:

- ▶ Although NN, NT, TN and TT are supported, only the NN version is implemented. For the NT, TN and TT versions, **csr2csc()** is used to transpose the relevant matrices, followed by a call to the NN version of **csrsgemm()**.
- ▶ The NN version needs working space of size **nnzA** integers at least.

- Only `CUSPARSE_MATRIX_TYPE_GENERAL` is supported. If either **A** or **B** is symmetric or Hermitian, the user must extend the matrix to a full one and reconfigure the `MatrixType` field descriptor to `CUSPARSE_MATRIX_TYPE_GENERAL`.
- This function requires temporary extra storage that is allocated internally
- The routine does *not* support asynchronous execution
- The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>transA</code>	the operation $\text{op}(A)$
<code>transB</code>	the operation $\text{op}(B)$
<code>m</code>	number of rows of sparse matrix $\text{op}(A)$ and $c$ .
<code>n</code>	number of columns of sparse matrix $\text{op}(B)$ and $c$ .
<code>k</code>	number of columns/rows of sparse matrix $\text{op}(A)$ / $\text{op}(B)$ .
<code>descrA</code>	the descriptor of matrix <b>A</b> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzA</code>	number of nonzero elements of sparse matrix <b>A</b> .
<code>csrValA</code>	<type> array of <code>nnzA</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <b>A</b> .
<code>csrRowPtrA</code>	integer array of $\tilde{m}+1$ elements that contains the start of every row and the end of the last row plus one. $\tilde{m}=m$ if <code>transA == CUSPARSE_OPERATION_NON_TRANSPOSE</code> , otherwise $\tilde{m}=k$ .
<code>csrColIndA</code>	integer array of <code>nnzA</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <b>A</b> .
<code>descrB</code>	the descriptor of matrix <b>B</b> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzB</code>	number of nonzero elements of sparse matrix <b>B</b> .
<code>csrValB</code>	<type> array of <code>nnzB</code> nonzero elements of matrix <b>B</b> .
<code>csrRowPtrB</code>	integer array of $\tilde{k}+1$ elements that contains the start of every row and the end of the last row plus one. $\tilde{k}=k$ if <code>transB == CUSPARSE_OPERATION_NON_TRANSPOSE</code> , otherwise $\tilde{k}=n$
<code>csrColIndB</code>	integer array of <code>nnzB</code> column indices of the nonzero elements of matrix <b>B</b> .
<code>descrC</code>	the descriptor of matrix <b>c</b> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.

## Output

<code>csrValC</code>	<type> array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) nonzero elements of matrix c.
<code>csrRowPtrC</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndC</code>	integer array of <code>nnzC</code> ( $= \text{csrRowPtrC}(m) - \text{csrRowPtrC}(0)$ ) column indices of the nonzero elements of matrix c.
<code>nnzTotalDevHostPtr</code>	total number of nonzero elements in device or host memory. It is equal to <code>(csrRowPtrC(m) - csrRowPtrC(0))</code> .

See `cusparseStatus_t` for the description of the return status

## 10.4. `cusparse<t>csrgemm2()`

```

cusparseStatus_t
cusparseScsrgemm2_bufferSizeExt(cusparseHandle_t          handle,
                                 int                  m,
                                 int                  n,
                                 int                  k,
                                 const float*         alpha,
                                 const cusparseMatDescr_t descrA,
                                 int                  nnZA,
                                 const int*           csrRowPtrA,
                                 const int*           csrColIndA,
                                 const cusparseMatDescr_t descrB,
                                 int                  nnzB,
                                 const int*           csrRowPtrB,
                                 const int*           csrColIndB,
                                 const float*         beta,
                                 const cusparseMatDescr_t descrD,
                                 int                  nnzD,
                                 const int*           csrRowPtrD,
                                 const int*           csrColIndD,
                                 csrgemm2Info_t       info,
                                 size_t*              size_t*
pBufferSizeInBytes)

cusparseStatus_t
cusparseDcsrgemm2_bufferSizeExt(cusparseHandle_t          handle,
                                 int                  m,
                                 int                  n,
                                 int                  k,
                                 const double*        alpha,
                                 const cusparseMatDescr_t descrA,
                                 int                  nnZA,
                                 const int*           csrRowPtrA,
                                 const int*           csrColIndA,
                                 const cusparseMatDescr_t descrB,
                                 int                  nnzB,
                                 const int*           csrRowPtrB,
                                 const int*           csrColIndB,
                                 const double*        beta,
                                 const cusparseMatDescr_t descrD,
                                 int                  nnzD,
                                 const int*           csrRowPtrD,
                                 const int*           csrColIndD,
                                 csrgemm2Info_t       info,
                                 size_t*              size_t*
pBufferSizeInBytes)

cusparseStatus_t
cusparseCcsrgemm2_bufferSizeExt(cusparseHandle_t          handle,
                                 int                  m,
                                 int                  n,
                                 int                  k,
                                 const cuComplex*     alpha,
                                 const cusparseMatDescr_t descrA,
                                 int                  nnZA,
                                 const int*           csrRowPtrA,
                                 const int*           csrColIndA,
                                 const cusparseMatDescr_t descrB,
                                 int                  nnzB,
                                 const int*           csrRowPtrB,
                                 const int*           csrColIndB,
                                 const cuComplex*     beta,
                                 const cusparseMatDescr_t descrD,
                                 int                  nnzD,
                                 const int*           csrRowPtrD,
                                 const int*           csrColIndD,
                                 csrgemm2Info_t       info,
                                 size_t*              size_t*
pBufferSizeInBytes)

```

This function performs following matrix-matrix operation:

$$C = \alpha * A * B + \beta * D$$

where **A**, **B**, **D** and **C** are  $m \times k$ ,  $k \times n$ ,  $m \times n$  and  $m \times n$  sparse matrices (defined in CSR storage format by the three arrays **csrValA**|**csrValB**|**csrValD**|**csrValC**, **csrRowPtrA**|**csrRowPtrB**|**csrRowPtrD**|**csrRowPtrC**, and **csrColIndA**|**csrColIndB**|**csrColIndD**|**csrColIndC** respectively).

We provide **csrgemm2** as a generalization of **csrgemm**. It provides more operations in terms of **alpha** and **beta**. For example,  $C = -A * B + D$  can be done by **csrgemm2**.

The **csrgemm2** uses **alpha** and **beta** to support the following operations:

<b>alpha</b>	<b>beta</b>	<b>operation</b>
NULL	NULL	invalid
NULL	!NULL	$C = \beta * D$ , <b>A</b> and <b>B</b> are not used
!NULL	NULL	$C = \alpha * A * B$ , <b>D</b> is not used
!NULL	!NULL	$C = \alpha * A * B + \beta * D$

The numerical value of **alpha** and **beta** only affects the numerical values of **C**, not its sparsity pattern. For example, if **alpha** and **beta** are not zero, the sparsity pattern of **C** is union of **A\*B** and **D**, independent of numerical value of **alpha** and **beta**.

The following table shows different operations according to the value of **m**, **n** and **k**

<b>m, n, k</b>	<b>operation</b>
$m < 0$ or $n < 0$ or $k < 0$	invalid
$m$ is 0 or $n$ is 0	do nothing
$m > 0$ and $n > 0$ and $k$ is 0	invalid if <b>beta</b> is zero; $C = \beta * D$ if <b>beta</b> is not zero.
$m > 0$ and $n > 0$ and $k > 0$	$C = \beta * D$ if <b>alpha</b> is zero. $C = \alpha * A * B$ if <b>beta</b> is zero. $C = \alpha * A * B + \beta * D$ if <b>alpha</b> and <b>beta</b> are not zero.

This function requires the buffer size returned by **csrgemm2\_bufferSizeExt()**.

The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

The cuSPARSE library adopts a two-step approach to complete sparse matrix. In the first step, the user allocates **csrRowPtrC** of  $m+1$  elements and uses the function **cusparseXcsrgemm2Nnz()** to determine **csrRowPtrC** and the total number of nonzero elements. In the second step, the user gathers **nnzC** (the number of nonzero elements of matrix **C**) from either (**nnzC=\*nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC(m) - csrRowPtrC(0)**) and allocates **csrValC** and **csrColIndC** of **nnzC** elements respectively, then finally calls function **cusparse[S|D|C|Z]csrgemm2()** to evaluate matrix **C**.

The general procedure of **C=-A\*B+D** is as follows:

```
// assume matrices A, B and D are ready.
int baseC, nnzC;
csrgemm2Info_t info = NULL;
size_t bufferSize;
void *buffer = NULL;
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzC;
double alpha = -1.0;
double beta = 1.0;
cusparseSetPointerMode(handle, CUSPARSE_POINTER_MODE_HOST);

// step 1: create an opaque structure
cusparsCreateCsrgemm2Info(&info);

// step 2: allocate buffer for csrgemm2Nnz and csrgemm2
cusparsDcsrgemm2_bufferSizeExt(handle, m, n, k, &alpha,
    descrA, nnzA, csrRowPtrA, csrColIndA,
    descrB, nnzB, csrRowPtrB, csrColIndB,
    &beta,
    descrD, nnzD, csrRowPtrD, csrColIndD,
    info,
    &bufferSize);
cudaMalloc(&buffer, bufferSize);

// step 3: compute csrRowPtrC
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
cusparsXcsrgemm2Nnz(handle, m, n, k,
    descrA, nnzA, csrRowPtrA, csrColIndA,
    descrB, nnzB, csrRowPtrB, csrColIndB,
    descrD, nnzD, csrRowPtrD, csrColIndD,
    descrC, csrRowPtrC, nnzTotalDevHostPtr,
    info, buffer );
if (NULL != nnzTotalDevHostPtr){
    nnzC = *nnzTotalDevHostPtr;
} else{
    cudaMemcpy(&nnzC, csrRowPtrC+m, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&baseC, csrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzC -= baseC;
}

// step 4: finish sparsity pattern and value of C
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnzC);
cudaMalloc((void**)&csrValC, sizeof(double)*nnzC);
// Remark: set csrValC to null if only sparsity pattern is required.
cusparsDcsrgemm2(handle, m, n, k, &alpha,
    descrA, nnzA, csrValA, csrRowPtrA, csrColIndA,
    descrB, nnzB, csrValB, csrRowPtrB, csrColIndB,
    &beta,
    descrD, nnzD, csrValD, csrRowPtrD, csrColIndD,
    descrC, csrValC, csrRowPtrC, csrColIndC,
    info, buffer);

// step 5: destroy the opaque structure
cusparsDestroyCsrgemm2Info(info);
```

Several comments on **csrgemm2()**:

- ▶ Only the NN version is supported. For other modes, the user has to transpose **A** or **B** explicitly.
- ▶ Only **CUSPARSE\_MATRIX\_TYPE\_GENERAL** is supported. If either **A** or **B** is symmetric or Hermitian, the user must extend the matrix to a full one and reconfigure the **MatrixType** field descriptor to **CUSPARSE\_MATRIX\_TYPE\_GENERAL**.

- if `csrValC` is zero, only sparsity pattern of `C` is calculated.

The functions `cusparseXcsrgeam2Nnz()` and `cusparse<t>csrgeam2()` supports the following properties if `pBuffer != NULL`

- The routine requires no extra storage
- The routine supports asynchronous execution
- The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of sparse matrix <code>A</code> , <code>D</code> and <code>C</code> .
<code>n</code>	number of columns of sparse matrix <code>B</code> , <code>D</code> and <code>C</code> .
<code>k</code>	number of columns/rows of sparse matrix <code>A</code> / <code>B</code> .
<code>alpha</code>	<type> scalar used for multiplication.
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzA</code>	number of nonzero elements of sparse matrix <code>A</code> .
<code>csrValA</code>	<type> array of <code>nnzA</code> nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnzA</code> column indices of the nonzero elements of matrix <code>A</code> .
<code>descrB</code>	the descriptor of matrix <code>B</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzB</code>	number of nonzero elements of sparse matrix <code>B</code> .
<code>csrValB</code>	<type> array of <code>nnzB</code> nonzero elements of matrix <code>B</code> .
<code>csrRowPtrB</code>	integer array of <code>k+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndB</code>	integer array of <code>nnzB</code> column indices of the nonzero elements of matrix <code>B</code> .
<code>descrD</code>	the descriptor of matrix <code>D</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> only.
<code>nnzD</code>	number of nonzero elements of sparse matrix <code>D</code> .
<code>csrValD</code>	<type> array of <code>nnzD</code> nonzero elements of matrix <code>D</code> .
<code>csrRowPtrD</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndD</code>	integer array of <code>nnzD</code> column indices of the nonzero elements of matrix <code>D</code> .

<b>beta</b>	<type> scalar used for multiplication.
<b>descrC</b>	the descriptor of matrix c. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> only.
<b>info</b>	structure with information used in <b>csrgemm2Nnz</b> and <b>csrgemm2</b> .
<b>pBuffer</b>	buffer allocated by the user; the size is returned by <b>csrgemm2_bufferSizeExt</b> .

## Output

<b>csrValC</b>	<type> array of <b>nnzC</b> nonzero elements of matrix c.
<b>csrRowPtrC</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndC</b>	integer array of <b>nnzC</b> column indices of the nonzero elements of matrix c.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>csrgemm2Nnnz</b> and <b>csrgemm2</b> .
<b>nnzTotalDevHostPtr</b>	total number of nonzero elements in device or host memory. It is equal to <b>(csrRowPtrC(m) - csrRowPtrC(0))</b> .

See **cusparseStatus\_t** for the description of the return status

# Chapter 11. CUSPARSE PRECONDITIONERS REFERENCE

This chapter describes the routines that implement different preconditioners.

In particular, the incomplete factorizations are implemented in two phases. First, during the analysis phase, the sparse triangular matrix is analyzed to determine the dependencies between its elements by calling the appropriate `csrv_analysis()` function. The analysis is specific to the sparsity pattern of the given matrix and the selected `cusparseOperation_t` type. The information from the analysis phase is stored in the parameter of type `cusparseSolveAnalysisInfo_t` that has been initialized previously with a call to `cusparseCreateSolveAnalysisInfo()`.

Second, during the numerical factorization phase, the given coefficient matrix is factorized using the information stored in the `cusparseSolveAnalysisInfo_t` parameter by calling the appropriate `csrilo0()` or `csric0()` function.

The analysis phase is shared across the sparse triangular solve, and the incomplete factorization and must be performed only once. The resulting information can be passed to the numerical factorization and the sparse triangular solve multiple times.

Finally, once the incomplete factorization and all the sparse triangular solves have completed, the opaque data structure pointed to by the `cusparseSolveAnalysisInfo_t` parameter can be released by calling `cusparseDestroySolveAnalysisInfo()`.

## 11.1. Incomplete Cholesky Factorization: level 0

Different algorithms for ic0 are discussed in this section.

## 11.1.1. `cusparse<t>csric0()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csric02_solve()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsric0(cusparseHandle_t          handle,
                cusparseOperation_t   trans,
                int                  m,
                const cusparseMatDescr_t descrA,
                float*               csrValM,
                const int*            csrRowPtrA,
                const int*            csrColIndA,
                cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseDcsric0(cusparseHandle_t          handle,
                cusparseOperation_t   trans,
                int                  m,
                const cusparseMatDescr_t descrA,
                double*              csrValM,
                const int*            csrRowPtrA,
                const int*            csrColIndA,
                cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseCcsric0(cusparseHandle_t          handle,
                cusparseOperation_t   trans,
                int                  m,
                const cusparseMatDescr_t descrA,
                cuComplex*            csrValM,
                const int*            csrRowPtrA,
                const int*            csrColIndA,
                cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseZcsric0(cusparseHandle_t          handle,
                cusparseOperation_t   trans,
                int                  m,
                const cusparseMatDescr_t descrA,
                cuDoubleComplex*      csrValM,
                const int*            csrRowPtrA,
                const int*            csrColIndA,
                cusparseSolveAnalysisInfo_t info)

```

This function computes the incomplete-Cholesky factorization with 0 fill-in and no pivoting:

$$\text{op}(A) \approx R^T R$$

**A** is an  $m \times m$  Hermitian/symmetric positive definite sparse matrix that is defined in CSR storage format by the three arrays `csrValM`, `csrRowPtrA`, and `csrColIndA`; and

$$\text{op}(A) = \begin{cases} A & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_NON\_TRANSPOSE} \\ A^T & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_TRANSPPOSE} \\ A^H & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_CONJUGATE\_TRANSPPOSE} \end{cases}$$

Notice that only a lower or upper Hermitian/symmetric part of the matrix **A** is actually stored. It is overwritten by the lower or upper triangular factors  $R^T$  and  $R$ , respectively.

A call to this routine must be preceded by a call to the **csrv\_analysis()** routine.

The matrix descriptor for **csrv\_analysis()** and **csric0()** must be the same. Otherwise, runtime error would occur.

The function supports the following properties if **pBuffer != NULL**

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>trans</b>	the operation <b>op (A)</b> .
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix types are <b>CUSPARSE_MATRIX_TYPE_SYMMETRIC</b> and <b>CUSPARSE_MATRIX_TYPE_HERMITIAN</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValM</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>info</b>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).

### Output

<b>csrValM</b>	<type> matrix containing the incomplete-Cholesky lower or upper triangular factor.
----------------	--

See **cusparseStatus\_t** for the description of the return status

## 11.1.2. `cusparse<t>csric02_bufferSize()`

```

cusparseStatus_t
cusparseScsric02_bufferSize(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           float*               csrValA,
                           const int*            csrRowPtrA,
                           const int*            csrColIndA,
                           csric02Info_t         info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseDcsric02_bufferSize(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           double*              csrValA,
                           const int*            csrRowPtrA,
                           const int*            csrColIndA,
                           csric02Info_t         info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseCcsric02_bufferSize(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           cuComplex*           csrValA,
                           const int*            csrRowPtrA,
                           const int*            csrColIndA,
                           csric02Info_t         info,
                           int*                 pBufferSizeInBytes)

cusparseStatus_t
cusparseZcsric02_bufferSize(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*      csrValA,
                           const int*            csrRowPtrA,
                           const int*            csrColIndA,
                           csric02Info_t         info,
                           int*                 pBufferSizeInBytes)

```

This function returns size of buffer used in computing the incomplete-Cholesky factorization with 0 fill-in and no pivoting:

$$A \approx LL^H$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**.

The buffer size depends on dimension **m** and **nnz**, the number of nonzeros of the matrix. If the user changes the matrix, it is necessary to call **csric02\_bufferSize()** again to have the correct buffer size; otherwise, a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>nnz</b>	number of nonzeros of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .

### Output

<b>info</b>	record internal states based on different algorithms.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>csric02_analysis()</b> and <b>csric02()</b> .

See [cusparseStatus\\_t](#) for the description of the return status

### 11.1.3. `cusparse<t>csric02_analysis()`

```

cusparseStatus_t
cusparseScsric02_analysis(cusparseHandle_t handle,
                           int m,
                           int nnz,
                           const cusparseMatDescr_t descrA,
                           const float* csrValA,
                           const int* csrRowPtrA,
                           const int* csrColIndA,
                           csric02Info_t info,
                           cusparseSolvePolicy_t policy,
                           void* pBuffer)

cusparseStatus_t
cusparseDcsric02_analysis(cusparseHandle_t handle,
                           int m,
                           int nnz,
                           const cusparseMatDescr_t descrA,
                           const double* csrValA,
                           const int* csrRowPtrA,
                           const int* csrColIndA,
                           csric02Info_t info,
                           cusparseSolvePolicy_t policy,
                           void* pBuffer)

cusparseStatus_t
cusparseCcsric02_analysis(cusparseHandle_t handle,
                           int m,
                           int nnz,
                           const cusparseMatDescr_t descrA,
                           const cuComplex* csrValA,
                           const int* csrRowPtrA,
                           const int* csrColIndA,
                           csric02Info_t info,
                           cusparseSolvePolicy_t policy,
                           void* pBuffer)

cusparseStatus_t
cusparseZcsric02_analysis(cusparseHandle_t handle,
                           int m,
                           int nnz,
                           const cusparseMatDescr_t descrA,
                           const cuDoubleComplex* csrValA,
                           const int* csrRowPtrA,
                           const int* csrColIndA,
                           csric02Info_t info,
                           cusparseSolvePolicy_t policy,
                           void* pBuffer)

```

This function performs the analysis phase of the incomplete-Cholesky factorization with 0 fill-in and no pivoting:

$$A \approx LL^H$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**.

This function requires a buffer size returned by `csric02_bufferSize()`. The address of `pBuffer` must be multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

Function `csric02_analysis()` reports a structural zero and computes level information stored in the opaque structure `info`. The level information can extract more parallelism during incomplete Cholesky factorization. However `csric02()` can be done without level information. To disable level information, the user must specify the policy of `csric02_analysis()` and `csric02()` as `CUSPARSE_SOLVE_POLICY_NO_LEVEL`.

Function `csric02_analysis()` always reports the first structural zero, even if the policy is `CUSPARSE_SOLVE_POLICY_NO_LEVEL`. The user needs to call `cusparseXcsric02_zeroPivot()` to know where the structural zero is.

It is the user's choice whether to call `csric02()` if `csric02_analysis()` reports a structural zero. In this case, the user can still call `csric02()`, which will return a numerical zero at the same position as the structural zero. However the result is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows and columns of matrix <code>A</code> .
<code>nnz</code>	number of nonzeros of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>info</code>	structure initialized using <code>cusparseCreateCsric02Info()</code> .
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user; the size is returned by <code>csric02_bufferSize()</code> .

## Output

<b>info</b>	number of bytes of the buffer used in <b>csric02_analysis()</b> and <b>csric02()</b> .
-------------	--

See [cusparseStatus\\_t](#) for the description of the return status

#### 11.1.4. `cusparse<t>csric02()`

```

cusparseStatus_t
cusparseSscsic02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  float*              csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csric02Info_t        info,
                  cusparseSolvePolicy_t policy,
                  void*                pBuffer)

cusparseStatus_t
cusparseDscsic02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  double*              csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csric02Info_t        info,
                  cusparseSolvePolicy_t policy,
                  void*                pBuffer)

cusparseStatus_t
cusparseCscsic02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  cuComplex*           csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csric02Info_t        info,
                  cusparseSolvePolicy_t policy,
                  void*                pBuffer)

cusparseStatus_t
cusparseZscsic02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  cuDoubleComplex*     csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csric02Info_t        info,
                  cusparseSolvePolicy_t policy,
                  void*                pBuffer)

```

This function performs the solve phase of the computing the incomplete-Cholesky factorization with 0 fill-in and no pivoting:

$$A \approx LL^H$$

This function requires a buffer size returned by `csric02_bufferSize()`. The address of `pBuffer` must be a multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

Although `csric02()` can be done without level information, the user still needs to be aware of consistency. If `csric02_analysis()` is called with policy `CUSPARSE_SOLVE_POLICY_USE_LEVEL`, `csric02()` can be run with or without levels. On the other hand, if `csric02_analysis()` is called with `CUSPARSE_SOLVE_POLICY_NO_LEVEL`, `csric02()` can only accept `CUSPARSE_SOLVE_POLICY_NO_LEVEL`; otherwise, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

Function `csric02()` reports the first numerical zero, including a structural zero. The user must call `cusparseXcsric02_zeroPivot()` to know where the numerical zero is.

Function `csric02()` only takes the lower triangular part of matrix `A` to perform factorization. The matrix type must be `CUSPARSE_MATRIX_TYPE_GENERAL`, the fill mode and diagonal type are ignored, and the strictly upper triangular part is ignored and never touched. It does not matter if `A` is Hermitian or not. In other words, from the point of view of `csric02()` `A` is Hermitian and only the lower triangular part is provided.



In practice, a positive definite matrix may not have incomplete cholesky factorization. To the best of our knowledge, only matrix `M` can guarantee the existence of incomplete cholesky factorization. If `csric02()` failed cholesky factorization and reported a numerical zero, it is possible that incomplete cholesky factorization does not exist.

For example, suppose `A` is a real  $m \times m$  matrix, the following code solves the precondition system  $M \cdot y = x$  where `M` is the product of Cholesky factorization `L` and its transpose.

$$M = LL^H$$

```

// Suppose that A is m x m sparse matrix represented by CSR format,
// Assumption:
// - handle is already created by cusparseCreate(),
// - (d_csrRowPtr, d_csrColInd, d_csrVal) is CSR of A on device memory,
// - d_x is right hand side vector on device memory,
// - d_y is solution vector on device memory.
// - d_z is intermediate result on device memory.

cusparseMatDescr_t descr_M = 0;
cusparseMatDescr_t descr_L = 0;
csric0Info_t info_M = {0};
csrv2Info_t info_L = {0};
csrv2Info_t info_Lt = {0};
int pBufferSize_M;
int pBufferSize_L;
int pBufferSize廖;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.0;
const cusparseSolvePolicy_t policy_M = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_L = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy廖 = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans_L = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseOperation_t trans廖 = CUSPARSE_OPERATION_TRANSPOSE;

// step 1: create a descriptor which contains
// - matrix M is base-1
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has non-unit diagonal
cusparseCreateMatDescr(&descr_M);
cusparseSetMatIndexBase(descr_M, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_M, CUSPARSE_MATRIX_TYPE_GENERAL);

cusparseCreateMatDescr(&descr_L);
cusparseSetMatIndexBase(descr_L, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_L, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_L, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr_L, CUSPARSE_DIAG_TYPE_NON_UNIT);

// step 2: create an empty info structure
// we need one info for csric02 and two info's for csrv2
cusparseCreateCsric02Info(&info_M);
cusparseCreateCsrsv2Info(&info廖);
cusparseCreateCsrsv2Info(&info廖);

// step 3: query how much memory used in csric02 and csrv2, and allocate the
// buffer
cusparseDcsric02_bufferSize(handle, m, nnz,
    descr_M, d_csrVal, d_csrRowPtr, d_csrColInd, info_M, &bufferSize_M);
cusparseDcsrv2_bufferSize(handle, trans廖, m, nnz,
    descr廖, d_csrVal, d_csrRowPtr, d_csrColInd, info廖, &pBufferSize廖);
cusparseDcsrv2_bufferSize(handle, trans廖, m, nnz,
    descr廖, d_csrVal, d_csrRowPtr, d_csrColInd, info廖, &pBufferSize廖);

pBufferSize = max(bufferSize_M, max(pBufferSize廖, pBufferSize廖));

// pBuffer returned by cudaMalloc is automatically aligned to 128 bytes.
cudaMalloc((void**)&pBuffer, pBufferSize);

// step 4: perform analysis of incomplete Cholesky on M
//           perform analysis of triangular solve on L
//           perform analysis of triangular solve on L'
// The lower triangular part of M has the same sparsity pattern as L, so
// we can do analysis of csric02 and csrv2 simultaneously.

cusparseDcsric02_analysis(handle, m, nnz, descr_M,
    d_csrVal, d_csrRowPtr, d_csrColInd, info_M,
    policy_M, pBuffer);
status = cusparseXcsric02_zeroPivot(handle, info_M, &structural_zero);
if (CUSPARSE_STATUS_ZERO_PIVOT == status)

```

The function supports the following properties if `pBuffer != NULL`

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows and columns of matrix <code>A</code> .
<code>nnz</code>	number of nonzeros of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA_valM</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>info</code>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user; the size is returned by <code>csric02_bufferSize()</code> .

### Output

<code>csrValA_valM</code>	<type> matrix containing the incomplete-Cholesky lower triangular factor.
---------------------------	---

See `cusparseStatus_t` for the description of the return status

## 11.1.5. `cusparseXcsric02_zeroPivot()`

```
cusparseStatus_t
cusparseXcsric02_zeroPivot(cusparseHandle_t handle,
                           csric02Info_t    info,
                           int*             position)
```

If the returned error code is `CUSPARSE_STATUS_ZERO_PIVOT`, `position=j` means `A(j,j)` has either a structural zero or a numerical zero; otherwise, `position=-1`.

The `position` can be 0-based or 1-based, the same as the matrix.

Function `cusparseXcsric02_zeroPivot()` is a blocking call. It calls `cudaDeviceSynchronize()` to make sure all previous kernels are done.

The `position` can be in the host memory or device memory. The user can set proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	<code>info</code> contains structural zero or numerical zero if the user already called <code>csric02_analysis()</code> or <code>csric02()</code> .

### Output

<code>position</code>	if no structural or numerical zero, <code>position</code> is -1; otherwise, if $A(j, j)$ is missing or $L(j, j)$ is zero, <code>position=j</code> .
-----------------------	---

See `cusparseStatus_t` for the description of the return status

## 11.1.6. `cusparse<t>bsric02_bufferSize()`

```

cusparseStatus_t
cusparseSbsric02_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           float*                bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsric02Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseDbsric02_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           double*               bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsric02Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseCbsric02_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuComplex*            bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsric02Info_t          info,
                           int*                  pBufferSizeInBytes)

cusparseStatus_t
cusparseZbsric02_bufferSize(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*       bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsric02Info_t          info,
                           int*                  pBufferSizeInBytes)

```

This function returns the size of a buffer used in computing the incomplete-Cholesky factorization with 0 fill-in and no pivoting

$$A \approx LL^H$$

**A** is an  $(\text{mb} * \text{blockDim}) * (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**.

The buffer size depends on the dimensions of **mb**, **blockDim**, and the number of nonzero blocks of the matrix **nnzb**. If the user changes the matrix, it is necessary to call **bsric02\_bufferSize()** again to have the correct buffer size; otherwise, a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows and block columns of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb + 1</b> elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> , larger than zero.

### Output

<b>info</b>	record internal states based on different algorithms.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>bsric02_analysis()</b> and <b>bsric02()</b> .

See **cusparseStatus\_t** for the description of the return status

## 11.1.7. `cusparse<t>bsric02_analysis()`

```

cusparseStatus_t
cusparseSbsric02_analysis(cusparseHandle_t      handle,
                           cusparseDirection_t   dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           const float*          bsrValA,
                           const int*            bsrRowPtrA,
                           const int*            bsrColIndA,
                           int                  blockDim,
                           bsric02Info_t         info,
                           cusparseSolvePolicy_t policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseDbsric02_analysis(cusparseHandle_t      handle,
                           cusparseDirection_t   dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           const double*         bsrValA,
                           const int*            bsrRowPtrA,
                           const int*            bsrColIndA,
                           int                  blockDim,
                           bsric02Info_t         info,
                           cusparseSolvePolicy_t policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseCbsric02_analysis(cusparseHandle_t      handle,
                           cusparseDirection_t   dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           const cuComplex*      bsrValA,
                           const int*            bsrRowPtrA,
                           const int*            bsrColIndA,
                           int                  blockDim,
                           bsric02Info_t         info,
                           cusparseSolvePolicy_t policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseZbsric02_analysis(cusparseHandle_t      handle,
                           cusparseDirection_t   dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           const cuDoubleComplex* bsrValA,
                           const int*            bsrRowPtrA,
                           const int*            bsrColIndA,
                           int                  blockDim,
                           bsric02Info_t         info,
                           cusparseSolvePolicy_t policy,
                           void*                pBuffer)

```

This function performs the analysis phase of the incomplete-Cholesky factorization with 0 fill-in and no pivoting

$$A \approx LL^H$$

**A** is an **(mb\*blockDim) x (mb\*blockDim)** sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**. The block in BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored.

This function requires a buffer size returned by **bsric02\_bufferSize90**.

The address of **pBuffer** must be a multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsric02\_analysis()** reports structural zero and computes level information stored in the opaque structure **info**. The level information can extract more parallelism during incomplete Cholesky factorization. However **bsric02()** can be done without level information. To disable level information, the user needs to specify the parameter **policy** of **bsric02[\_analysis]** as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **bsric02\_analysis** always reports the first structural zero, even when parameter **policy** is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. The user must call **cusparseXbsric02\_zeroPivot()** to know where the structural zero is.

It is the user's choice whether to call **bsric02()** if **bsric02\_analysis()** reports a structural zero. In this case, the user can still call **bsric02()**, which returns a numerical zero in the same position as the structural zero. However the result is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows and block columns of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix <b>A</b> .

<b>bsrRowPtrA</b>	integer array of <code>mb</code> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix A; must be larger than zero.
<b>info</b>	structure initialized using <code>cusparseCreateBsic02Info()</code> .
<b>policy</b>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<b>pBuffer</b>	buffer allocated by the user; the size is returned by <code>bsic02_bufferSize()</code> .

## Output

<b>info</b>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------	---

See `cusparseStatus_t` for the description of the return status

## 11.1.8. `cusparse<t>bsric02()`

```

cusparseStatus_t
cusparseSbsric02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descrA,
                  float*                bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsric02Info_t          info,
                  cusparseSolvePolicy_t   policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseDbsric02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descrA,
                  double*               bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsric02Info_t          info,
                  cusparseSolvePolicy_t   policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseCbsric02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descrA,
                  cuComplex*            bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsric02Info_t          info,
                  cusparseSolvePolicy_t   policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseZbsric02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descrA,
                  cuDoubleComplex*       bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsric02Info_t          info,
                  cusparseSolvePolicy_t   policy,
                  void*                 pBuffer)

```

This function performs the solve phase of the incomplete-Cholesky factorization with 0 fill-in and no pivoting

$$A \approx LL^H$$

**A** is an  $(mb * blockDim) \times (mb * blockDim)$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**. The block in BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored.

This function requires a buffer size returned by **bsric02\_bufferSize()**.

The address of **pBuffer** must be a multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **bsric02()** can be done without level information, the user must be aware of consistency. If **bsric02\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **bsric02()** can be run with or without levels. On the other hand, if **bsric02\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **bsric02()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsric02()** has the same behavior as **csric02()**. That is, **bsr2csr(bsric02(A)) = csric02(bsr2csr(A))**. The numerical zero of **csric02()** means there exists some zero  $L(j, j)$ . The numerical zero of **bsric02()** means there exists some block  $L_j, j$  that is not invertible.

Function **bsric02** reports the first numerical zero, including a structural zero. The user must call **cusparseXbsric02\_zeroPivot()** to know where the numerical zero is.

The **bsric02()** function only takes the lower triangular part of matrix **A** to perform factorization. The strictly upper triangular part is ignored and never touched. It does not matter if **A** is Hermitian or not. In other words, from the point of view of **bsric02()**, **A** is Hermitian and only the lower triangular part is provided. Moreover, the imaginary part of diagonal elements of diagonal blocks is ignored.

For example, suppose **A** is a real m-by-m matrix, where **m=mb\*blockDim**. The following code solves precondition system **M\*y = x**, where **M** is the product of Cholesky factorization **L** and its transpose.

$$M = LL^H$$

```

// Suppose that A is m x m sparse matrix represented by BSR format,
// The number of block rows/columns is mb, and
// the number of nonzero blocks is nnzb.
// Assumption:
// - handle is already created by cusparseCreate(),
// - (d_bsrRowPtr, d_bsrColInd, d_bsrVal) is BSR of A on device memory,
// - d_x is right hand side vector on device memory,
// - d_y is solution vector on device memory.
// - d_z is intermediate result on device memory.
// - d_x, d_y and d_z are of size m.
cusparseMatDescr_t descr_M = 0;
cusparseMatDescr_t descr_L = 0;
bsric02Info_t info_M = 0;
bsrvsv2Info_t info_L = 0;
bsrvsv2Info_t info廖 = 0;
int pBufferSize_M;
int pBufferSize_L;
int pBufferSize廖;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.;
const cusparseSolvePolicy_t policy_M = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_L = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy廖 = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans_L = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseOperation_t trans廖 = CUSPARSE_OPERATION_TRANSPOSE;
const cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;

// step 1: create a descriptor which contains
// - matrix M is base-1
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has non-unit diagonal
cusparseCreateMatDescr(&descr_M);
cusparseSetMatIndexBase(descr_M, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_M, CUSPARSE_MATRIX_TYPE_GENERAL);

cusparseCreateMatDescr(&descr_L);
cusparseSetMatIndexBase(descr_L, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_L, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_L, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr_L, CUSPARSE_DIAG_TYPE_NON_UNIT);

// step 2: create a empty info structure
// we need one info for bsric02 and two info's for bsrsv2
cusparseCreateBsric02Info(&info_M);
cusparseCreateBsrsv2Info(&info廖);
cusparseCreateBsrsv2Info(&info廖);

// step 3: query how much memory used in bsric02 and bsrsv2, and allocate the
// buffer
cusparseDbsric02_bufferSize(handle, dir, mb, nnzb,
    descr_M, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info_M,
    &bufferSize_M);
cusparseDbsrvsv2_bufferSize(handle, dir, trans廖, mb, nnzb,
    descr廖, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info廖,
    &pBufferSize廖);
cusparseDbsrvsv2_bufferSize(handle, dir, trans廖, mb, nnzb,
    descr廖, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info廖,
    &pBufferSize廖);

pBufferSize = max(bufferSize_M, max(pBufferSize廖, pBufferSize廖));

// pBuffer returned by cudaMalloc is automatically aligned to 128 bytes.
cudaMalloc((void**)&pBuffer, pBufferSize);

// step 4: perform analysis of incomplete Cholesky on M
//           perform analysis of triangular solve on L
//           perform analysis of triangular solve on L'
// The lower triangular part of M has the same sparsity pattern as L' so

```

The function supports the following properties if `pBuffer != NULL`

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>dirA</code>	storage format of blocks, either <code>CUSPARSE_DIRECTION_ROW</code> or <code>CUSPARSE_DIRECTION_COLUMN</code> .
<code>mb</code>	number of block rows and block columns of matrix <code>A</code> .
<code>nnzb</code>	number of nonzero blocks of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>bsrValA</code>	<type> array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) nonzero blocks of matrix <code>A</code> .
<code>bsrRowPtrA</code>	integer array of <code>mb + 1</code> elements that contains the start of every block row and the end of the last block row plus one.
<code>bsrColIndA</code>	integer array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix <code>A</code> .
<code>blockDim</code>	block dimension of sparse matrix <code>A</code> , larger than zero.
<code>info</code>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user, the size is returned by <code>bsric02_bufferSize()</code> .

## Output

<code>bsrValA</code>	<type> matrix containing the incomplete-Cholesky lower triangular factor.
----------------------	---

See `cusparseStatus_t` for the description of the return status

## 11.1.9. `cusparseXbsric02_zeroPivot()`

```
cusparseStatus_t
cusparseXbsric02_zeroPivot(cusparseHandle_t handle,
                           bsric02Info_t    info,
                           int*             position)
```

If the returned error code is **CUSPARSE\_STATUS\_ZERO\_PIVOT**, **position=j** means  $\mathbf{A}(j,j)$  has either a structural zero or a numerical zero (the block is not positive definite). Otherwise **position=-1**.

The **position** can be 0-based or 1-based, the same as the matrix.

Function **`cusparseXbsric02_zeroPivot()`** is a blocking call. It calls **`cudaDeviceSynchronize()`** to make sure all previous kernels are done.

The **position** can be in the host memory or device memory. The user can set the proper mode with **`cusparseSetPointerMode()`**.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>info</b>	<b>info</b> contains a structural zero or a numerical zero if the user already called <b><code>bsric02_analysis()</code></b> or <b><code>bsric02()</code></b> .

### Output

<b>position</b>	if no structural or numerical zero, <b>position</b> is -1, otherwise if $\mathbf{A}(j,j)$ is missing or $\mathbf{L}(j,j)$ is not positive definite, <b>position=j</b> .
-----------------	---

See **`cusparseStatus_t`** for the description of the return status

## 11.2. Incomplete LU Factorization: level 0

Different algorithms for ilu0 are discussed in this section.

## 11.2.1. `cusparse<t>csrilu0()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrilu02_solve()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsrilu0(cusparseHandle_t          handle,
                  cusparseOperation_t   trans,
                  int                  m,
                  const cusparseMatDescr_t descrA,
                  float*               csrValM,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseDcsrilu0(cusparseHandle_t          handle,
                  cusparseOperation_t   trans,
                  int                  m,
                  const cusparseMatDescr_t descrA,
                  double*              csrValM,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseCcsrilu0(cusparseHandle_t          handle,
                  cusparseOperation_t   trans,
                  int                  m,
                  const cusparseMatDescr_t descrA,
                  cuComplex*            csrValM,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  cusparseSolveAnalysisInfo_t info)

cusparseStatus_t
cusparseZcsrilu0(cusparseHandle_t          handle,
                  cusparseOperation_t   trans,
                  int                  m,
                  const cusparseMatDescr_t descrA,
                  cuDoubleComplex*      csrValM,
                  const int*            csrRowPtrA,
                  const int*            csrColIndA,
                  cusparseSolveAnalysisInfo_t info)

```

This function computes the incomplete-LU factorization with 0 fill-in and no pivoting:

$$op(A) \approx LU$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays `csrValM`, `csrRowPtrA`, and `csrColIndA`; and

$$op(A) = \begin{cases} A & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_NON\_TRANSPOSE} \\ A^T & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_TRANSPPOSE} \\ A^H & \text{if } \text{trans} == \text{CUSPARSE\_OPERATION\_CONJUGATE\_TRANSPPOSE} \end{cases}$$

Notice that the diagonal of lower triangular factor  $L$  is unitary and need not be stored. Therefore, the input matrix is overwritten with the resulting lower and upper triangular factors  $L$  and  $U$ , respectively.

A call to this routine must be preceded by a call to the `csrv_analysis()` routine.

The matrix descriptor for `csrv_analysis()` and `csrilu0()` must be the same. Otherwise, runtime error would occur.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>trans</code>	the operation <code>op (A)</code> .
<code>m</code>	number of rows and columns of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValM</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>info</code>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).

### Output

<code>csrValM</code>	<type> matrix containing the incomplete-LU lower and upper triangular factors.
----------------------	--

See `cusparseStatus_t` for the description of the return status

## 11.2.2. `cusparseCsrilu0Ex()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>csrilu02_solve()` instead. *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseCsrilu0Ex(cusparseHandle_t handle,
                  cusparseOperation_t trans,
                  int m,
                  const cusparseMatDescr_t descrA,
                  void * csrSortedValA_ValM,
                  cudaDataType csrSortedValA_ValMtype,
                  const int * csrSortedRowPtrA,
                  const int * csrSortedColIndA,
                  cusparseSolveAnalysisInfo_t info,
                  cudaDataType executiontype)
```

This function is an extended version of `cusparse<t>csrilu0()`. For detailed description of the functionality, see `cusparse<t>csrilu0()`.

This function does not support half-precision execution type, but it supports half-precision IO with single precision execution.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input specifically required by `cusparseCsrilu0Ex`

<code>csrSortedValA_ValMtype</code>	Data type of <code>csrSortedValA_ValM</code> .
<code>executiontype</code>	Data type used for computation.

### 11.2.3. `cusparse<t>csrilonumericBoost()`

```

cusparseStatus_t
cusparseScsrilonumericBoost(cusparseHandle_t handle,
                             csrilon2Info_t info,
                             int enable_boost,
                             double* tol,
                             float* boost_val)

cusparseStatus_t
cusparseDcsrilonumericBoost(cusparseHandle_t handle,
                             csrilon2Info_t info,
                             int enable_boost,
                             double* tol,
                             double* boost_val)

cusparseStatus_t
cusparseCcsrilonumericBoost(cusparseHandle_t handle,
                             csrilon2Info_t info,
                             int enable_boost,
                             double* tol,
                             cuComplex* boost_val)

cusparseStatus_t
cusparseZcsrilonumericBoost(cusparseHandle_t handle,
                             csrilon2Info_t info,
                             int enable_boost,
                             double* tol,
                             cuDoubleComplex* boost_val)

```

The user can use a boost value to replace a numerical value in incomplete LU factorization. The `tol` is used to determine a numerical zero, and the `boost_val` is used to replace a numerical zero. The behavior is

if `tol >= fabs(A(j,j))`, then `A(j,j)=boost_val`.

To enable a boost value, the user has to set parameter `enable_boost` to 1 before calling `csrilon2()`. To disable a boost value, the user can call `csrilon2_numericBoost()` again with parameter `enable_boost=0`.

If `enable_boost=0`, `tol` and `boost_val` are ignored.

Both `tol` and `boost_val` can be in the host memory or device memory. The user can set the proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

#### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	structure initialized using <code>cusparseCreateCsrilon2Info()</code> .
<code>enable_boost</code>	disable boost by <code>enable_boost=0</code> ; otherwise, boost is enabled.

<b>tol</b>	tolerance to determine a numerical zero.
<b>boost_val</b>	boost value to replace a numerical zero.

See [cusparseStatus\\_t](#) for the description of the return status

## 11.2.4. `cusparse<t>csrilu02_bufferSize()`

```

cusparseStatus_t
cusparseScsrilu02_bufferSize(cusparseHandle_t          handle,
                             int                  m,
                             int                  nnz,
                             const cusparseMatDescr_t descrA,
                             float*               csrValA,
                             const int*            csrRowPtrA,
                             const int*            csrColIndA,
                             csrilu02Info_t        info,
                             int*                 int*
pBufferSizeInBytes)

cusparseStatus_t
cusparseDcsrilu02_bufferSize(cusparseHandle_t          handle,
                             int                  m,
                             int                  nnz,
                             const cusparseMatDescr_t descrA,
                             double*              csrValA,
                             const int*            csrRowPtrA,
                             const int*            csrColIndA,
                             csrilu02Info_t        info,
                             int*                 int*
pBufferSizeInBytes)

cusparseStatus_t
cusparseCcsrilu02_bufferSize(cusparseHandle_t          handle,
                             int                  m,
                             int                  nnz,
                             const cusparseMatDescr_t descrA,
                             cuComplex*           csrValA,
                             const int*            csrRowPtrA,
                             const int*            csrColIndA,
                             csrilu02Info_t        info,
                             int*                 int*
pBufferSizeInBytes)

cusparseStatus_t
cusparseZcsrilu02_bufferSize(cusparseHandle_t          handle,
                             int                  m,
                             int                  nnz,
                             const cusparseMatDescr_t descrA,
                             cuDoubleComplex*      csrValA,
                             const int*            csrRowPtrA,
                             const int*            csrColIndA,
                             csrilu02Info_t        info,
                             int*                 int*
pBufferSizeInBytes)

```

This function returns size of the buffer used in computing the incomplete-LU factorization with 0 fill-in and no pivoting:

$$A \approx LU$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**.

The buffer size depends on the dimension **m** and **nnz**, the number of nonzeros of the matrix. If the user changes the matrix, it is necessary to call **csrilu02\_bufferSize()** again to have the correct buffer size; otherwise, a segmentation fault may occur.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows and columns of matrix <b>A</b> .
<b>nnz</b>	number of nonzeros of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m + 1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .

### Output

<b>info</b>	record internal states based on different algorithms.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>csrilu02_analysis()</b> and <b>csrilu02()</b> .

See **cusparseStatus\_t** for the description of the return status

## 11.2.5. `cusparse<t>csrilu02_analysis()`

```

cusparseStatus_t
cusparseScsrilu02_analysis(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           const float*          csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrilu02Info_t         info,
                           cusparseSolvePolicy_t  policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseDcsrilu02_analysis(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           const double*          csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrilu02Info_t         info,
                           cusparseSolvePolicy_t  policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseCcsrilu02_analysis(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           const cuComplex*       csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrilu02Info_t         info,
                           cusparseSolvePolicy_t  policy,
                           void*                pBuffer)

cusparseStatus_t
cusparseZcsrilu02_analysis(cusparseHandle_t          handle,
                           int                  m,
                           int                  nnz,
                           const cusparseMatDescr_t descrA,
                           const cuDoubleComplex* csrValA,
                           const int*             csrRowPtrA,
                           const int*             csrColIndA,
                           csrilu02Info_t         info,
                           cusparseSolvePolicy_t  policy,
                           void*                pBuffer)

```

This function performs the analysis phase of the incomplete-LU factorization with 0 fill-in and no pivoting:

$$A \approx LU$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**.

This function requires the buffer size returned by `csrilon2_bufferSize()`. The address of `pBuffer` must be a multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

Function `csrilon2_analysis()` reports a structural zero and computes level information stored in the opaque structure `info`. The level information can extract more parallelism during incomplete LU factorization; however `csrilon2()` can be done without level information. To disable level information, the user must specify the policy of `csrilon2()` as `CUSPARSE_SOLVE_POLICY_NO_LEVEL`.

It is the user's choice whether to call `csrilon2()` if `csrilon2_analysis()` reports a structural zero. In this case, the user can still call `csrilon2()`, which will return a numerical zero at the same position as the structural zero. However the result is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows and columns of matrix <code>A</code> .
<code>nnz</code>	number of nonzeros of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>info</code>	structure initialized using <code>cusparseCreateCsrilon2Info()</code> .
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user, the size is returned by <code>csrilon2_bufferSize()</code> .

## Output

<code>info</code>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------------	---

See `cusparseStatus_t` for the description of the return status

### 11.2.6. `cusparse<t>csrilu02()`

```

cusparseStatus_t
cusparseScsrilu02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  float*              csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csrilu02Info_t       info,
                  cusparseSolvePolicy_t policy,
                  void*               pBuffer)

cusparseStatus_t
cusparseDcsrilu02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  double*             csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csrilu02Info_t       info,
                  cusparseSolvePolicy_t policy,
                  void*               pBuffer)

cusparseStatus_t
cusparseCcsrilu02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  cuComplex*          csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csrilu02Info_t       info,
                  cusparseSolvePolicy_t policy,
                  void*               pBuffer)

cusparseStatus_t
cusparseZcsrilu02(cusparseHandle_t          handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  cuDoubleComplex*    csrValA_valM,
                  const int*           csrRowPtrA,
                  const int*           csrColIndA,
                  csrilu02Info_t       info,
                  cusparseSolvePolicy_t policy,
                  void*               pBuffer)

```

This function performs the solve phase of the incomplete-LU factorization with 0 fill-in and no pivoting:

$$A \approx LU$$

**A** is an  $m \times m$  sparse matrix that is defined in CSR storage format by the three arrays `csrValA_valM`, `csrRowPtrA`, and `csrColIndA`.

This function requires a buffer size returned by `csrilon2_bufferSize()`. The address of `pBuffer` must be a multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

The matrix type must be `CUSPARSE_MATRIX_TYPE_GENERAL`. The fill mode and diagonal type are ignored.

Although `csrilon2()` can be done without level information, the user still needs to be aware of consistency. If `csrilon2_analysis()` is called with policy `CUSPARSE_SOLVE_POLICY_USE_LEVEL`, `csrilon2()` can be run with or without levels. On the other hand, if `csrilon2_analysis()` is called with `CUSPARSE_SOLVE_POLICY_NO_LEVEL`, `csrilon2()` can only accept `CUSPARSE_SOLVE_POLICY_NO_LEVEL`; otherwise, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

Function `csrilon2()` reports the first numerical zero, including a structural zero. The user must call `cusparseXcsrilon2_zeroPivot()` to know where the numerical zero is.

For example, suppose **A** is a real  $m \times m$  matrix, the following code solves precondition system  $\mathbf{M}^*\mathbf{y} = \mathbf{x}$  where **M** is the product of LU factors **L** and **U**.

```
// Suppose that A is m x m sparse matrix represented by CSR format,
// Assumption:
// - handle is already created by cusparseCreate(),
// - (d_csrRowPtr, d_csrColInd, d_csrVal) is CSR of A on device memory,
// - d_x is right hand side vector on device memory,
// - d_y is solution vector on device memory.
// - d_z is intermediate result on device memory.

cusparseMatDescr_t descr_M = 0;
cusparseMatDescr_t descr_L = 0;
cusparseMatDescr_t descr_U = 0;
csrilu02Info_t info_M = 0;
csrvs2Info_t info_L = 0;
csrvs2Info_t info_U = 0;
int pBufferSize_M;
int pBufferSize_L;
int pBufferSize_U;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.;
const cusparseSolvePolicy_t policy_M = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_L = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_U = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans_L = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseOperation_t trans_U = CUSPARSE_OPERATION_NON_TRANSPOSE;

// step 1: create a descriptor which contains
// - matrix M is base-1
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has unit diagonal
// - matrix U is base-1
// - matrix U is upper triangular
// - matrix U has non-unit diagonal
cusparseCreateMatDescr(&descr_M);
cusparseSetMatIndexBase(descr_M, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_M, CUSPARSE_MATRIX_TYPE_GENERAL);

cusparseCreateMatDescr(&descr_L);
cusparseSetMatIndexBase(descr_L, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_L, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_L, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr_L, CUSPARSE_DIAG_TYPE_UNIT);

cusparseCreateMatDescr(&descr_U);
cusparseSetMatIndexBase(descr_U, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_U, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_U, CUSPARSE_FILL_MODE_UPPER);
cusparseSetMatDiagType(descr_U, CUSPARSE_DIAG_TYPE_NON_UNIT);

// step 2: create a empty info structure
// we need one info for csrilu02 and two info's for csrvs2
cusparseCreateCsrilu02Info(&info_M);
cusparseCreateCsrsv2Info(&info_L);
cusparseCreateCsrsv2Info(&info_U);

// step 3: query how much memory used in csrilu02 and csrvs2, and allocate the
// buffer
cusparseDcsrilu02_bufferSize(handle, m, nnz,
    descr_M, d_csrVal, d_csrRowPtr, d_csrColInd, info_M, &pBufferSize_M);
cusparseDcsrsv2_bufferSize(handle, trans_L, m, nnz,
    descr_L, d_csrVal, d_csrRowPtr, d_csrColInd, info_L, &pBufferSize_L);
cusparseDcsrsv2_bufferSize(handle, trans_U, m, nnz,
    descr_U, d_csrVal, d_csrRowPtr, d_csrColInd, info_U, &pBufferSize_U);

pBufferSize = max(pBufferSize_M, max(pBufferSize_L, pBufferSize_U));

// pBuffer returned by cudaMalloc is automatically aligned to 128 bytes.
cudaMalloc((void**)&pBuffer, pBufferSize);

// step 4: perform analysis of incomplete Cholesky on M
```

The function supports the following properties if `pBuffer != NULL`

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows and columns of matrix <code>A</code> .
<code>nnz</code>	number of nonzeros of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA_valM</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m + 1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>info</code>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user; the size is returned by <code>csrilu02_bufferSize()</code> .

### Output

<code>csrValA_valM</code>	<type> matrix containing the incomplete-LU lower and upper triangular factors.
---------------------------	--

See `cusparseStatus_t` for the description of the return status

## 11.2.7. `cusparseXcsrilu02_zeroPivot()`

```
cusparseStatus_t
cusparseXcsrilu02_zeroPivot(cusparseHandle_t handle,
                           csrilu02Info_t   info,
                           int*            position)
```

If the returned error code is `CUSPARSE_STATUS_ZERO_PIVOT`, `position=j` means `A(j, j)` has either a structural zero or a numerical zero; otherwise, `position=-1`.

The `position` can be 0-based or 1-based, the same as the matrix.

Function **`cusparseXcsrilu02_zeroPivot()`** is a blocking call. It calls **`cudaDeviceSynchronize()`** to make sure all previous kernels are done.

The **`position`** can be in the host memory or device memory. The user can set proper mode with **`cusparseSetPointerMode()`**.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

#### Input

<b><code>handle</code></b>	handle to the cuSPARSE library context.
<b><code>info</code></b>	<code>info</code> contains structural zero or numerical zero if the user already called <code>csrilu02_analysis()</code> or <code>csrilu02()</code> .

#### Output

<b><code>position</code></b>	if no structural or numerical zero, <code>position</code> is -1; otherwise if $A(j,j)$ is missing or $U(j,j)$ is zero, <code>position=j</code> .
------------------------------	--

See **`cusparseStatus_t`** for the description of the return status

### 11.2.8. **`cusparse<t>bsrilu02_numericBoost()`**

```

cusparseStatus_t
cusparseSbsrilu02_numericBoost(cusparseHandle_t handle,
                                bsrilu02Info_t info,
                                int enable_boost,
                                double* tol,
                                float* boost_val)

cusparseStatus_t
cusparseDbsrilu02_numericBoost(cusparseHandle_t handle,
                                bsrilu02Info_t info,
                                int enable_boost,
                                double* tol,
                                double* boost_val)

cusparseStatus_t
cusparseCbsrilu02_numericBoost(cusparseHandle_t handle,
                                bsrilu02Info_t info,
                                int enable_boost,
                                double* tol,
                                cuComplex* boost_val)

cusparseStatus_t
cusparseZbsrilu02_numericBoost(cusparseHandle_t handle,
                                bsrilu02Info_t info,
                                int enable_boost,
                                double* tol,
                                cuDoubleComplex* boost_val)

```

The user can use a boost value to replace a numerical value in incomplete LU factorization. Parameter `tol` is used to determine a numerical zero, and `boost_val` is used to replace a numerical zero. The behavior is as follows:

if `tol >= fabs(A(j,j))`, then reset each diagonal element of block `A(j,j)` by `boost_val`.

To enable a boost value, the user sets parameter `enable_boost` to 1 before calling `bsriliu02()`. To disable the boost value, the user can call `bsriliu02_numericBoost()` with parameter `enable_boost=0`.

If `enable_boost=0`, `tol` and `boost_val` are ignored.

Both `tol` and `boost_val` can be in host memory or device memory. The user can set the proper mode with `cusparseSetPointerMode()`.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>info</code>	structure initialized using <code>cusparseCreateBsriliu02Info()</code> .
<code>enable_boost</code>	disable boost by setting <code>enable_boost=0</code> . Otherwise, boost is enabled.
<code>tol</code>	tolerance to determine a numerical zero.
<code>boost_val</code>	boost value to replace a numerical zero.

See `cusparseStatus_t` for the description of the return status

## 11.2.9. `cusparse<t>bsr_ilu02_bufferSize()`

```

cusparseStatus_t
cusparseSbsrilu02_bufferSize(cusparseHandle_t handle,
                             cusparseDirection_t dirA,
                             int mb,
                             int nnzb,
                             const cusparseMatDescr_t descrA,
                             float *bsrValA,
                             const int *bsrRowPtrA,
                             const int *bsrColIndA,
                             int blockDim,
                             bsr_ilu02Info_t info,
                             int *pBufferSizeInBytes);

cusparseStatus_t
cusparseDbsrilu02_bufferSize(cusparseHandle_t handle,
                             cusparseDirection_t dirA,
                             int mb,
                             int nnzb,
                             const cusparseMatDescr_t descrA,
                             double *bsrValA,
                             const int *bsrRowPtrA,
                             const int *bsrColIndA,
                             int blockDim,
                             bsr_ilu02Info_t info,
                             int *pBufferSizeInBytes);

cusparseStatus_t
cusparseCbsrilu02_bufferSize(cusparseHandle_t handle,
                             cusparseDirection_t dirA,
                             int mb,
                             int nnzb,
                             const cusparseMatDescr_t descrA,
                             cuComplex *bsrValA,
                             const int *bsrRowPtrA,
                             const int *bsrColIndA,
                             int blockDim,
                             bsr_ilu02Info_t info,
                             int *pBufferSizeInBytes);

cusparseStatus_t
cusparseZbsrilu02_bufferSize(cusparseHandle_t handle,
                             cusparseDirection_t dirA,
                             int mb,
                             int nnzb,
                             const cusparseMatDescr_t descrA,
                             cuDoubleComplex *bsrValA,
                             const int *bsrRowPtrA,
                             const int *bsrColIndA,
                             int blockDim,
                             bsr_ilu02Info_t info,
                             int *pBufferSizeInBytes);

```

This function returns the size of the buffer used in computing the incomplete-LU factorization with 0 fill-in and no pivoting

$$A \approx LU$$

**A** is an  $(\text{mb} * \text{blockDim}) * (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**.

The buffer size depends on the dimensions of **mb**, **blockDim**, and the number of nonzero blocks of the matrix **nnzb**. If the user changes the matrix, it is necessary to call **bsrilu02\_bufferSize()** again to have the correct buffer size; otherwise, a segmentation fault may occur.

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows and columns of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) nonzero blocks of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb</b> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<b>bsrColIndA</b>	integer array of <b>nnzb</b> ( $= \text{bsrRowPtrA}(\text{mb}) - \text{bsrRowPtrA}(0)$ ) column indices of the nonzero blocks of matrix <b>A</b> .
<b>blockDim</b>	block dimension of sparse matrix <b>A</b> , larger than zero.

### Output

<b>info</b>	record internal states based on different algorithms.
<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in <b>bsrilu02_analysis()</b> and <b>bsrilu02()</b> .

### Status Returned

<b>CUSPARSE_STATUS_SUCCESS</b>	the operation completed successfully.
<b>CUSPARSE_STATUS_NOT_INITIALIZED</b>	the library was not initialized.
<b>CUSPARSE_STATUS_ALLOC_FAILED</b>	the resources could not be allocated.
<b>CUSPARSE_STATUS_INVALID_VALUE</b>	invalid parameters were passed ( <b>mb</b> , <b>nnzb</b> $\leq 0$ ), base index is not 0 or 1.

<b>CUSPARSE_STATUS_ARCH_MISMATCH</b>	the device only supports compute capability 2.0 and above.
<b>CUSPARSE_STATUS_INTERNAL_ERROR</b>	an internal operation failed.
<b>CUSPARSE_STATUS_MATRIX_TYPE_NOT_SUPPORTED</b>	the matrix type is not supported.

## 11.2.10. `cusparse<t>bsrilu02_analysis()`

```

cusparseStatus_t
cusparseSbsrilu02_analysis(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           float*                bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsrilu02Info_t         info,
                           cusparseSolvePolicy_t   policy,
                           void*                 pBuffer)

cusparseStatus_t
cusparseDbsrilu02_analysis(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           double*               bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsrilu02Info_t         info,
                           cusparseSolvePolicy_t   policy,
                           void*                 pBuffer)

cusparseStatus_t
cusparseCbsrilu02_analysis(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuComplex*             bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsrilu02Info_t         info,
                           cusparseSolvePolicy_t   policy,
                           void*                 pBuffer)

cusparseStatus_t
cusparseZbsrilu02_analysis(cusparseHandle_t          handle,
                           cusparseDirection_t    dirA,
                           int                  mb,
                           int                  nnzb,
                           const cusparseMatDescr_t descrA,
                           cuDoubleComplex*        bsrValA,
                           const int*             bsrRowPtrA,
                           const int*             bsrColIndA,
                           int                   blockDim,
                           bsrilu02Info_t         info,
                           cusparseSolvePolicy_t   policy,
                           void*                 pBuffer)

```

This function performs the analysis phase of the incomplete-LU factorization with 0 fill-in and no pivoting

$$A \approx LU$$

**A** is an  $(\text{mb} * \text{blockDim}) \times (\text{mb} * \text{blockDim})$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**. The block in BSR format is of size **blockDim\*blockDim**, stored as column-major or row-major as determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored.

This function requires a buffer size returned by **bsrilu02\_bufferSize()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsrilu02\_analysis()** reports a structural zero and computes level information stored in the opaque structure **info**. The level information can extract more parallelism during incomplete LU factorization. However **bsrilu02()** can be done without level information. To disable level information, the user needs to specify the parameter **policy** of **bsrilu02[\_analysis]** as **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**.

Function **bsrilu02\_analysis()** always reports the first structural zero, even with parameter **policy** is **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**. The user must call **cusparseXbsrilu02\_zeroPivot()** to know where the structural zero is.

It is the user's choice whether to call **bsrilu02()** if **bsrilu02\_analysis()** reports a structural zero. In this case, the user can still call **bsrilu02()**, which will return a numerical zero at the same position as the structural zero. However the result is meaningless.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dirA</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows and block columns of matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of <b>nnzb</b> (= <b>bsrRowPtrA(mb) - bsrRowPtrA(0)</b> ) nonzero blocks of matrix <b>A</b> .

<code>bsrRowPtrA</code>	integer array of <code>mb</code> + 1 elements that contains the start of every block row and the end of the last block row plus one.
<code>bsrColIndA</code>	integer array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix <code>A</code> .
<code>blockDim</code>	block dimension of sparse matrix <code>A</code> , larger than zero.
<code>info</code>	structure initialized using <code>cusparseCreateBsrilu02Info()</code> .
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user, the size is returned by <code>bsrilu02_bufferSize()</code> .

## Output

<code>info</code>	structure filled with information collected during the analysis phase (that should be passed to the solve phase unchanged).
-------------------	---

See `cusparseStatus_t` for the description of the return status

## 11.2.11. `cusparse<t>bsr_ilu02()`

```

cusparseStatus_t
cusparseSbsrilu02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descry,
                  float*                bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsr_ilu02Info_t        info,
                  cusparseSolvePolicy_t  policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseDbsrilu02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descry,
                  double*               bsrValA,
                  const int*             bsrRowPtrA,
                  const int*             bsrColIndA,
                  int                   blockDim,
                  bsr_ilu02Info_t        info,
                  cusparseSolvePolicy_t  policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseCbsrilu02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descry,
                  cuComplex*             bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                   blockDim,
                  bsr_ilu02Info_t         info,
                  cusparseSolvePolicy_t  policy,
                  void*                 pBuffer)

cusparseStatus_t
cusparseZbsrilu02(cusparseHandle_t          handle,
                  cusparseDirection_t    dirA,
                  int                   mb,
                  int                   nnzb,
                  const cusparseMatDescr_t descry,
                  cuDoubleComplex*       bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                   blockDim,
                  bsr_ilu02Info_t         info,
                  cusparseSolvePolicy_t  policy,
                  void*                 pBuffer)

```

This function performs the solve phase of the incomplete-LU factorization with 0 fill-in and no pivoting

$$A \approx LU$$

**A** is an  $(mb * blockDim) \times (mb * blockDim)$  sparse matrix that is defined in BSR storage format by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**. The block in BSR format is of size **blockDim \* blockDim**, stored as column-major or row-major determined by parameter **dirA**, which is either **CUSPARSE\_DIRECTION\_COLUMN** or **CUSPARSE\_DIRECTION\_ROW**. The matrix type must be **CUSPARSE\_MATRIX\_TYPE\_GENERAL**, and the fill mode and diagonal type are ignored. Function **bsrilu02()** supports an arbitrary **blockDim**.

This function requires a buffer size returned by **bsrilu02\_bufferSize()**. The address of **pBuffer** must be a multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Although **bsrilu02()** can be used without level information, the user must be aware of consistency. If **bsrilu02\_analysis()** is called with policy **CUSPARSE\_SOLVE\_POLICY\_USE\_LEVEL**, **bsrilu02()** can be run with or without levels. On the other hand, if **bsrilu02\_analysis()** is called with **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**, **bsrilu02()** can only accept **CUSPARSE\_SOLVE\_POLICY\_NO\_LEVEL**; otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Function **bsrilu02()** has the same behavior as **csriliu02()**. That is, **bsr2csr(bsrilu02(A)) = csriliu02(bsr2csr(A))**. The numerical zero of **csriliu02()** means there exists some zero  $U(j, j)$ . The numerical zero of **bsrilu02()** means there exists some block  $U(j, j)$  that is not invertible.

Function **bsrilu02** reports the first numerical zero, including a structural zero. The user must call **cusparseXbsrilu02\_zeroPivot()** to know where the numerical zero is.

For example, suppose **A** is a real m-by-m matrix where  $m=mb*blockDim$ . The following code solves precondition system  $\mathbf{M} \cdot \mathbf{y} = \mathbf{x}$ , where **M** is the product of LU factors **L** and **U**.

```

// Suppose that A is m x m sparse matrix represented by BSR format,
// The number of block rows/columns is mb, and
// the number of nonzero blocks is nnzb.
// Assumption:
// - handle is already created by cusparseCreate(),
// - (d_bsrRowPtr, d_bsrColInd, d_bsrVal) is BSR of A on device memory,
// - d_x is right hand side vector on device memory.
// - d_y is solution vector on device memory.
// - d_z is intermediate result on device memory.
// - d_x, d_y and d_z are of size m.
cusparseMatDescr_t descr_M = 0;
cusparseMatDescr_t descr_L = 0;
cusparseMatDescr_t descr_U = 0;
bsrilu02Info_t info_M = 0;
bsrvs2Info_t info_L = 0;
bsrvs2Info_t info_U = 0;
int pBufferSize_M;
int pBufferSize_L;
int pBufferSize_U;
int pBufferSize;
void *pBuffer = 0;
int structural_zero;
int numerical_zero;
const double alpha = 1.;
const cusparseSolvePolicy_t policy_M = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_L = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const cusparseSolvePolicy_t policy_U = CUSPARSE_SOLVE_POLICY_USE_LEVEL;
const cusparseOperation_t trans_L = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseOperation_t trans_U = CUSPARSE_OPERATION_NON_TRANSPOSE;
const cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;

// step 1: create a descriptor which contains
// - matrix M is base-1
// - matrix L is base-1
// - matrix L is lower triangular
// - matrix L has unit diagonal
// - matrix U is base-1
// - matrix U is upper triangular
// - matrix U has non-unit diagonal
cusparseCreateMatDescr(&descr_M);
cusparseSetMatIndexBase(descr_M, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_M, CUSPARSE_MATRIX_TYPE_GENERAL);

cusparseCreateMatDescr(&descr_L);
cusparseSetMatIndexBase(descr_L, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_L, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_L, CUSPARSE_FILL_MODE_LOWER);
cusparseSetMatDiagType(descr_L, CUSPARSE_DIAG_TYPE_UNIT);

cusparseCreateMatDescr(&descr_U);
cusparseSetMatIndexBase(descr_U, CUSPARSE_INDEX_BASE_ONE);
cusparseSetMatType(descr_U, CUSPARSE_MATRIX_TYPE_GENERAL);
cusparseSetMatFillMode(descr_U, CUSPARSE_FILL_MODE_UPPER);
cusparseSetMatDiagType(descr_U, CUSPARSE_DIAG_TYPE_NON_UNIT);

// step 2: create a empty info structure
// we need one info for bsrilu02 and two info's for bsrvs2
cusparseCreateBsrilu02Info(&info_M);
cusparseCreateBsrsv2Info(&info_L);
cusparseCreateBsrsv2Info(&info_U);

// step 3: query how much memory used in bsrilu02 and bsrvs2, and allocate the
// buffer
cusparseDbsrilu02_bufferSize(handle, dir, mb, nnzb,
    descr_M, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info_M,
    &pBufferSize_M);
cusparseDbsrvs2_bufferSize(handle, dir, trans_L, mb, nnzb,
    descr_L, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info_L,
    &pBufferSize_L);
cusparseDbsrvs2_bufferSize(handle, dir, trans_U, mb, nnzb,
    descr_U, d_bsrVal, d_bsrRowPtr, d_bsrColInd, blockDim, info_U,
    &pBufferSize_U);

```

The function supports the following properties if `pBuffer != NULL`

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>dirA</code>	storage format of blocks: either <code>CUSPARSE_DIRECTION_ROW</code> or <code>CUSPARSE_DIRECTION_COLUMN</code> .
<code>mb</code>	number of block rows and block columns of matrix <code>A</code> .
<code>nnzb</code>	number of nonzero blocks of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>bsrValA</code>	<type> array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) nonzero blocks of matrix <code>A</code> .
<code>bsrRowPtrA</code>	integer array of <code>mb + 1</code> elements that contains the start of every block row and the end of the last block row plus one.
<code>bsrColIndA</code>	integer array of <code>nnzb</code> (= <code>bsrRowPtrA(mb) - bsrRowPtrA(0)</code> ) column indices of the nonzero blocks of matrix <code>A</code> .
<code>blockDim</code>	block dimension of sparse matrix <code>A</code> ; must be larger than zero.
<code>info</code>	structure with information collected during the analysis phase (that should have been passed to the solve phase unchanged).
<code>policy</code>	the supported policies are <code>CUSPARSE_SOLVE_POLICY_NO_LEVEL</code> and <code>CUSPARSE_SOLVE_POLICY_USE_LEVEL</code> .
<code>pBuffer</code>	buffer allocated by the user; the size is returned by <code>bsrilu02_bufferSize()</code> .

## Output

<code>bsrValA</code>	<type> matrix containing the incomplete-LU lower and upper triangular factors.
----------------------	--

See `cusparseStatus_t` for the description of the return status

## 11.2.12. `cusparseXbsrili02_zeroPivot()`

```
cusparseStatus_t
cusparseXbsrili02_zeroPivot(cusparseHandle_t handle,
                           bsrili02Info_t info,
                           int* position)
```

If the returned error code is **CUSPARSE\_STATUS\_ZERO\_PIVOT**, **position=j** means  $A(j,j)$  has either a structural zero or a numerical zero (the block is not invertible). Otherwise **position=-1**.

The **position** can be 0-based or 1-based, the same as the matrix.

Function **`cusparseXbsrili02_zeroPivot()`** is a blocking call. It calls **`cudaDeviceSynchronize()`** to make sure all previous kernels are done.

The **position** can be in the host memory or device memory. The user can set proper the mode with **`cusparseSetPointerMode()`**.

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>info</b>	<b>info</b> contains structural zero or numerical zero if the user already called <b><code>bsrili02_analysis()</code></b> or <b><code>bsrili02()</code></b> .

### Output

<b>position</b>	if no structural or numerical zero, <b>position</b> is -1; otherwise if $A(j,j)$ is missing or $U(j,j)$ is not invertible, <b>position=j</b> .
-----------------	--

See **`cusparseStatus_t`** for the description of the return status

## 11.3. Tridiagonal Solve

Different algorithms for tridiagonal solve are discussed in this section.

### 11.3.1. `cusparse<t>gtsv()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>gtsv2()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseSgtsv(cusparseHandle_t handle,
              int             m,
              int             n,
              const float*   dl,
              const float*   d,
              const float*   du,
              float*         B,
              int             ldb)

cusparseStatus_t
cusparseDgtsv(cusparseHandle_t handle,
              int             m,
              int             n,
              const double*  dl,
              const double*  d,
              const double*  du,
              double*        B,
              int             ldb)

cusparseStatus_t
cusparseCgtsv(cusparseHandle_t handle,
              int             m,
              int             n,
              const cuComplex* dl,
              const cuComplex* d,
              const cuComplex* du,
              cuComplex*      B,
              int             ldb)

cusparseStatus_t
cusparseZgtsv(cusparseHandle_t handle,
              int             m,
              int             n,
              const cuDoubleComplex* dl,
              const cuDoubleComplex* d,
              const cuDoubleComplex* du,
              cuDoubleComplex*    B,
              int             ldb)

```

This function computes the solution of a tridiagonal linear system with multiple right-hand sides:

$$A * Y = X$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **x**. Notice that solution **y** overwrites right-hand-side matrix **x** on exit.

Assuming **A** is of size **m** and base-1, **dl**, **d** and **du** are defined by the following formula:

$$dl(i) := A(i, i-1) \text{ for } i=1, 2, \dots, m$$

The first element of `dl` is out-of-bound (`dl(1) := A(1, 0)`), so `dl(1) = 0`.

`d(i) = A(i, i)` for `i=1, 2, ..., m`

`du(i) = A(i, i+1)` for `i=1, 2, ..., m`

The last element of `du` is out-of-bound (`du(m) := A(m, m+1)`), so `du(m) = 0`.

The routine does perform pivoting, which usually results in more accurate and more stable results than `cusparse<t>gtsv_nopivot()` at the expense of some execution time

- ▶ This function requires significant amount of temporary extra storage (`min(m, 8) * (3+n) * sizeof(<type>)`)
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>n</code>	number of right-hand sides, columns of matrix <code>B</code> .
<code>dl</code>	<code>&lt;type&gt;</code> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<code>d</code>	<code>&lt;type&gt;</code> dense array containing the main diagonal of the tri-diagonal linear system.
<code>du</code>	<code>&lt;type&gt;</code> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<code>B</code>	<code>&lt;type&gt;</code> dense right-hand-side array of dimensions <code>(lDb, n)</code> .
<code>lDb</code>	leading dimension of <code>B</code> (that is $\geq \max(1, m)$ ).

## Output

<code>B</code>	<code>&lt;type&gt;</code> dense solution array of dimensions <code>(lDb, n)</code> .
----------------	--

See `cusparseStatus_t` for the description of the return status

### 11.3.2. `cusparse<t>gtsv_nopivot()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>gtsv2_nopivot()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseSgtsv_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const float* dl,
                      const float* d,
                      const float* du,
                      float* B,
                      int ldb)

cusparseStatus_t
cusparseDgtsv_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const double* dl,
                      const double* d,
                      const double* du,
                      double* B,
                      int ldb)

cusparseStatus_t
cusparseCgtsv_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const cuComplex* dl,
                      const cuComplex* d,
                      const cuComplex* du,
                      cuComplex* B,
                      int ldb)

cusparseStatus_t
cusparseZgtsv_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const cuDoubleComplex* dl,
                      const cuDoubleComplex* d,
                      const cuDoubleComplex* du,
                      cuDoubleComplex* B,
                      int ldb)

```

This function computes the solution of a tridiagonal linear system with multiple right-hand sides:

$$A * Y = X$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **x**. Notice that solution **y** overwrites right-hand-side matrix **x** on exit.

The routine does not perform any pivoting and uses a combination of the Cyclic Reduction (CR) and the Parallel Cyclic Reduction (PCR) algorithms to find the solution. It achieves better performance when  $m$  is a power of 2.

This routine requires a significant amount of temporary extra storage ( $m \times (3+n) \times \text{sizeof}(<\text{type}>)$ ). The temporary storage is managed by `cudaMalloc` and `cudaFree` which stop concurrency. The user can call `cusparse<t>gtsv2_nopivot()` which has no explicit `cudaMalloc` and `cudaFree`.

- ▶ This function requires a significant amount of temporary extra storage ( $m \times (3+n) \times \text{sizeof}(<\text{type}>)$ ).
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>n</code>	number of right-hand sides, columns of matrix <code>B</code> .
<code>dL</code>	<code>&lt;type&gt;</code> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<code>d</code>	<code>&lt;type&gt;</code> dense array containing the main diagonal of the tri-diagonal linear system.
<code>dU</code>	<code>&lt;type&gt;</code> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<code>B</code>	<code>&lt;type&gt;</code> dense right-hand-side array of dimensions <code>(1dB, n)</code> .
<code>1dB</code>	leading dimension of <code>B</code> . (that is $\geq \max(1, m)$ ).

### Output

<code>B</code>	<code>&lt;type&gt;</code> dense solution array of dimensions <code>(1dB, n)</code> .
----------------	--

See `cusparseStatus_t` for the description of the return status

### 11.3.3. `cusparse<t>gtsv2_buffSizeExt()`

```

cusparseStatus_t
cusparseSgtsv2_bufferSizeExt(cusparseHandle_t handle,
    int m,
    int n,
    const float* dl,
    const float* d,
    const float* du,
    const float* B,
    int ldb,
    size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseDgtsv2_bufferSizeExt(cusparseHandle_t handle,
    int m,
    int n,
    const double* dl,
    const double* d,
    const double* du,
    const double* B,
    int ldb,
    size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseCgtsv2_bufferSizeExt(cusparseHandle_t handle,
    int m,
    int n,
    const cuComplex* dl,
    const cuComplex* d,
    const cuComplex* du,
    const cuComplex* B,
    int ldb,
    size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseZgtsv2_bufferSizeExt(cusparseHandle_t handle,
    int m,
    int n,
    const cuDoubleComplex* dl,
    const cuDoubleComplex* d,
    const cuDoubleComplex* du,
    const cuDoubleComplex* B,
    int ldb,
    size_t* bufferSizeInBytes)

```

This function returns the size of the buffer used in `gtsv2` which computes the solution of a tridiagonal linear system with multiple right-hand sides.

$$A * X = B$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **X** overwrites right-hand-side matrix **B** on exit.

- ▶ The routine requires no extra storage

- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>n</code>	number of right-hand sides, columns of matrix <code>B</code> .
<code>dl</code>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<code>d</code>	<type> dense array containing the main diagonal of the tri-diagonal linear system.
<code>du</code>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<code>B</code>	<type> dense right-hand-side array of dimensions ( <code>ldb</code> , <code>n</code> ).
<code>ldb</code>	leading dimension of <code>B</code> (that is $\geq \max(1, m)$ ).

### Output

<code>pBufferSizeInBytes</code>	number of bytes of the buffer used in the <code>gtsv2</code> .
---------------------------------	--

See `cusparseStatus_t` for the description of the return status

### 11.3.4. `cusparse<t>gtsv2()`

```

cusparseStatus_t
cusparseSgtsv2(cusparseHandle_t handle,
               int          m,
               int          n,
               const float* dl,
               const float* d,
               const float* du,
               float*       B,
               int          ldb,
               void         pBuffer)

cusparseStatus_t
cusparseDgtsv2(cusparseHandle_t handle,
               int          m,
               int          n,
               const double* dl,
               const double* d,
               const double* du,
               double*      B,
               int          ldb,
               void         pBuffer)

cusparseStatus_t
cusparseCgtsv2(cusparseHandle_t handle,
               int          m,
               int          n,
               const cuComplex* dl,
               const cuComplex* d,
               const cuComplex* du,
               cuComplex*    B,
               int          ldb,
               void         pBuffer)

cusparseStatus_t
cusparseZgtsv2(cusparseHandle_t handle,
               int          m,
               int          n,
               const cuDoubleComplex* dl,
               const cuDoubleComplex* d,
               const cuDoubleComplex* du,
               cuDoubleComplex*   B,
               int          ldb,
               void         pBuffer)

```

This function computes the solution of a tridiagonal linear system with multiple right-hand sides:

$$A * X = B$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **X** overwrites right-hand-side matrix **B** on exit.

Assuming **A** is of size **m** and base-1, **dl**, **d** and **du** are defined by the following formula:

```
dl(i) := A(i, i-1) for  $i=1, 2, \dots, m$ 
```

The first element of dl is out-of-bound (**dl(1) := A(1, 0)**), so **dl(1) = 0**.

```
d(i) = A(i, i) for  $i=1, 2, \dots, m$ 
```

```
du(i) = A(i, i+1) for  $i=1, 2, \dots, m$ 
```

The last element of du is out-of-bound (**du(m) := A(m, m+1)**), so **du(m) = 0**.

The routine does perform pivoting, which usually results in more accurate and more stable results than **cusparse<t>gtsv\_nopivot()** or **cusparse<t>gtsv2\_nopivot()** at the expense of some execution time.

This function requires a buffer size returned by **gtsv2\_bufferSizeExt()**.

The address of **pBuffer** must be multiple of 128 bytes. If it is not,

**CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	the size of the linear system (must be $\geq 3$ ).
<b>n</b>	number of right-hand sides, columns of matrix <b>B</b> .
<b>dl</b>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<b>d</b>	<type> dense array containing the main diagonal of the tri-diagonal linear system.
<b>du</b>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<b>B</b>	<type> dense right-hand-side array of dimensions ( <b>ldb</b> , <b>n</b> ).
<b>ldb</b>	leading dimension of <b>B</b> (that is $\geq \max(1, m)$ ).
<b>pBuffer</b>	buffer allocated by the user, the size is return by <b>gtsv2_bufferSizeExt</b> .

## Output

<b>B</b>	<type> dense solution array of dimensions ( <b>ldb</b> , <b>n</b> ).
----------	--

See **cusparseStatus\_t** for the description of the return status

### 11.3.5. `cusparse<t>gtsv2_nopivot_bufferSizeExt()`

```

cusparseStatus_t
cusparseSgtsv2_nopivot_bufferSizeExt(cusparseHandle_t handle,
                                      int m,
                                      int n,
                                      const float* dl,
                                      const float* d,
                                      const float* du,
                                      const float* B,
                                      int ldb,
                                      size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseDgtsv2_nopivot_bufferSizeExt(cusparseHandle_t handle,
                                      int m,
                                      int n,
                                      const double* dl,
                                      const double* d,
                                      const double* du,
                                      const double* B,
                                      int ldb,
                                      size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseCgtsv2_nopivot_bufferSizeExt(cusparseHandle_t handle,
                                      int m,
                                      int n,
                                      const cuComplex* dl,
                                      const cuComplex* d,
                                      const cuComplex* du,
                                      const cuComplex* B,
                                      int ldb,
                                      size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseZgtsv2_nopivot_bufferSizeExt(cusparseHandle_t handle,
                                      int m,
                                      int n,
                                      const cuDoubleComplex* dl,
                                      const cuDoubleComplex* d,
                                      const cuDoubleComplex* du,
                                      const cuDoubleComplex* B,
                                      int ldb,
                                      size_t* bufferSizeInBytes)

```

This function returns the size of the buffer used in `gtsv2_nopivot` which computes the solution of a tridiagonal linear system with multiple right-hand sides.

$$A * X = B$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **X** overwrites right-hand-side matrix **B** on exit.

- ▶ The routine requires no extra storage

- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>n</code>	number of right-hand sides, columns of matrix <code>B</code> .
<code>dl</code>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<code>d</code>	<type> dense array containing the main diagonal of the tri-diagonal linear system.
<code>du</code>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<code>B</code>	<type> dense right-hand-side array of dimensions <code>(lDb, n)</code> .
<code>lDb</code>	leading dimension of <code>B</code> . (that is $\geq \max(1, m)$ ).

### Output

<code>pBufferSizeInBytes</code>	number of bytes of the buffer used in the <code>gtsv2_nopivot</code> .
---------------------------------	--

See `cusparseStatus_t` for the description of the return status

### 11.3.6. `cusparse<t>gtsv2_nopivot()`

```

cusparseStatus_t
cusparseSgtsv2_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const float* dl,
                      const float* d,
                      const float* du,
                      float* B,
                      int ldb,
                      void* pBuffer)

cusparseStatus_t
cusparseDgtsv2_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const double* dl,
                      const double* d,
                      const double* du,
                      double* B,
                      int ldb,
                      void* pBuffer)

cusparseStatus_t
cusparseCgtsv2_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const cuComplex* dl,
                      const cuComplex* d,
                      const cuComplex* du,
                      cuComplex* B,
                      int ldb,
                      void* pBuffer)

cusparseStatus_t
cusparseZgtsv2_nopivot(cusparseHandle_t handle,
                      int m,
                      int n,
                      const cuDoubleComplex* dl,
                      const cuDoubleComplex* d,
                      const cuDoubleComplex* du,
                      cuDoubleComplex* B,
                      int ldb,
                      void* pBuffer)

```

This function computes the solution of a tridiagonal linear system with multiple right-hand sides:

$$A * X = B$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **x** overwrites right-hand-side matrix **B** on exit.

The routine does not perform any pivoting and uses a combination of the Cyclic Reduction (CR) and the Parallel Cyclic Reduction (PCR) algorithms to find the solution. It achieves better performance when  $m$  is a power of 2.

This function requires a buffer size returned by `gtsv2_nopivot_bufferSizeExt()`. The address of `pBuffer` must be multiple of 128 bytes. If it is not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>n</code>	number of right-hand sides, columns of matrix <code>B</code> .
<code>dl</code>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<code>d</code>	<type> dense array containing the main diagonal of the tri-diagonal linear system.
<code>du</code>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<code>B</code>	<type> dense right-hand-side array of dimensions <code>(ldb, n)</code> .
<code>ldb</code>	leading dimension of <code>B</code> . (that is $\geq \max(1, m)$ ).
<code>pBuffer</code>	buffer allocated by the user, the size is return by <code>gtsv2_nopivot_bufferSizeExt</code> .

### Output

<code>B</code>	<type> dense solution array of dimensions <code>(ldb, n)</code> .
----------------	---

See `cusparseStatus_t` for the description of the return status

## 11.4. Batched Tridiagonal Solve

Different algorithms for batched tridiagonal solve are discussed in this section.

## 11.4.1. `cusparse<t>gtsvStridedBatch()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparse<t>gtsv2StridedBatch()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseSgtsvStridedBatch(cusparseHandle_t handle,
                           int             m,
                           const float*   dl,
                           const float*   d,
                           const float*   du,
                           float*         x,
                           int            batchCount,
                           int            batchStride)

cusparseStatus_t
cusparseDgtsvStridedBatch(cusparseHandle_t handle,
                           int             m,
                           const double*  dl,
                           const double*  d,
                           const double*  du,
                           double*        x,
                           int            batchCount,
                           int            batchStride)

cusparseStatus_t
cusparseCgtsvStridedBatch(cusparseHandle_t handle,
                           int             m,
                           const cuComplex* dl,
                           const cuComplex* d,
                           const cuComplex* du,
                           cuComplex*      x,
                           int            batchCount,
                           int            batchStride)

cusparseStatus_t
cusparseZgtsvStridedBatch(cusparseHandle_t      handle,
                           int             m,
                           const cuDoubleComplex* dl,
                           const cuDoubleComplex* d,
                           const cuDoubleComplex* du,
                           cuDoubleComplex*      x,
                           int            batchCount,
                           int            batchStride)

```

This function computes the solution of multiple tridiagonal linear systems for  $i=0, \dots, \text{batchCount}$ :

$$A^{(i)} * y^{(i)} = x^{(i)}$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **x**. Notice that solution **y** overwrites right-hand-side matrix **x** on exit. The different matrices are assumed to be of the same size and are stored with a fixed **batchStride** in memory.

The routine does not perform any pivoting and uses a combination of the Cyclic Reduction (CR) and the Parallel Cyclic Reduction (PCR) algorithms to find the solution. It achieves better performance when  $m$  is a power of 2.

- ▶ This function requires a significant amount of temporary extra storage ( $((batchCount \times (4 \times m + 2048)) \times sizeof(<type>))$ ).
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	the size of the linear system (must be $\geq 3$ ).
<code>dl</code>	$<type>$ dense array containing the lower diagonal of the tri-diagonal linear system. The lower diagonal $dl^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location $dl + batchStride \times i$ in memory. Also, the first element of each lower diagonal must be zero.
<code>d</code>	$<type>$ dense array containing the main diagonal of the tri-diagonal linear system. The main diagonal $d^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location $d + batchStride \times i$ in memory.
<code>du</code>	$<type>$ dense array containing the upper diagonal of the tri-diagonal linear system. The upper diagonal $du^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location $du + batchStride \times i$ in memory. Also, the last element of each upper diagonal must be zero.
<code>x</code>	$<type>$ dense array that contains the right-hand-side of the tri-diagonal linear system. The right-hand-side $x^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location $x + batchStride \times i$ in memory.
<code>batchCount</code>	number of systems to solve.
<code>batchStride</code>	stride (number of elements) that separates the vectors of every system (must be at least $m$ ).

## Output

<code>x</code>	$<type>$ dense array that contains the solution of the tri-diagonal linear system. The solution $x^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location $x + batchStride \times i$ in memory.
----------------	--

See `cusparseStatus_t` for the description of the return status

## 11.4.2. `cusparse<t>gtsv2StridedBatch_bufferSizeExt()`

```

cusparseStatus_t
cusparseSgtsv2StridedBatch_bufferSizeExt(cusparseHandle_t handle,
                                         int m,
                                         const float* dl,
                                         const float* d,
                                         const float* du,
                                         const float* x,
                                         int batchCount,
                                         int batchStride,
                                         size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseDgtsv2StridedBatch_bufferSizeExt(cusparseHandle_t handle,
                                         int m,
                                         const double* dl,
                                         const double* d,
                                         const double* du,
                                         const double* x,
                                         int batchCount,
                                         int batchStride,
                                         size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseCgtsv2StridedBatch_bufferSizeExt(cusparseHandle_t handle,
                                         int m,
                                         const cuComplex* dl,
                                         const cuComplex* d,
                                         const cuComplex* du,
                                         const cuComplex* x,
                                         int batchCount,
                                         int batchStride,
                                         size_t* bufferSizeInBytes)

cusparseStatus_t
cusparseZgtsv2StridedBatch_bufferSizeExt(cusparseHandle_t handle,
                                         int m,
                                         const cuDoubleComplex* dl,
                                         const cuDoubleComplex* d,
                                         const cuDoubleComplex* du,
                                         const cuDoubleComplex* x,
                                         int batchCount,
                                         int batchStride,
                                         size_t* bufferSizeInBytes)

```

This function returns the size of the buffer used in `gtsv2StridedBatch` which computes the solution of multiple tridiagonal linear systems for  $i=0,\dots,\text{batchCount}$ :

$$A^{(i)} * y^{(i)} = x^{(i)}$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**d1**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **x**. Notice that solution **y** overwrites right-hand-side matrix **x** on exit. The different matrices are assumed to be of the same size and are stored with a fixed **batchStride** in memory.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>n</b>	the size of the linear system (must be $\geq 3$ ).
<b>d1</b>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The lower diagonal $d1^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <b>d1+batchStride*i</b> in memory. Also, the first element of each lower diagonal must be zero.
<b>d</b>	<type> dense array containing the main diagonal of the tri-diagonal linear system. The main diagonal $d^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <b>d+batchStride*i</b> in memory.
<b>du</b>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The upper diagonal $du^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <b>du+batchStride*i</b> in memory. Also, the last element of each upper diagonal must be zero.
<b>x</b>	<type> dense array that contains the right-hand-side of the tri-diagonal linear system. The right-hand-side $x^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <b>x+batchStride*i</b> in memory.
<b>batchCount</b>	number of systems to solve.
<b>batchStride</b>	stride (number of elements) that separates the vectors of every system (must be at least <b>m</b> ).

## Output

<b>pBufferSizeInBytes</b>	number of bytes of the buffer used in the <b>gtsv2StridedBatch</b> .
---------------------------	--

See [cusparseStatus\\_t](#) for the description of the return status

### 11.4.3. `cusparse<t>gtsv2StridedBatch()`

```

cusparseStatus_t
cusparseSgtsv2StridedBatch(cusparseHandle_t handle,
                           int m,
                           const float* dl,
                           const float* d,
                           const float* du,
                           float* x,
                           int batchCount,
                           int batchStride,
                           void* pBuffer)

cusparseStatus_t
cusparseDgtsv2StridedBatch(cusparseHandle_t handle,
                           int m,
                           const double* dl,
                           const double* d,
                           const double* du,
                           double* x,
                           int batchCount,
                           int batchStride,
                           void* pBuffer)

cusparseStatus_t
cusparseCgtsv2StridedBatch(cusparseHandle_t handle,
                           int m,
                           const cuComplex* dl,
                           const cuComplex* d,
                           const cuComplex* du,
                           cuComplex* x,
                           int batchCount,
                           int batchStride,
                           void* pBuffer)

cusparseStatus_t
cusparseZgtsv2StridedBatch(cusparseHandle_t handle,
                           int m,
                           const cuDoubleComplex* dl,
                           const cuDoubleComplex* d,
                           const cuDoubleComplex* du,
                           cuDoubleComplex* x,
                           int batchCount,
                           int batchStride,
                           void* pBuffer)

```

This function computes the solution of multiple tridiagonal linear systems for  $i=0, \dots, \text{batchCount}$ :

$$A^{(i)} * y^{(i)} = x^{(i)}$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**dl**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **x**. Notice that solution **y** overwrites right-hand-side matrix **x** on exit. The different matrices are assumed to be of the same size and are stored with a fixed **batchStride** in memory.

The routine does not perform any pivoting and uses a combination of the Cyclic Reduction (CR) and the Parallel Cyclic Reduction (PCR) algorithms to find the solution. It achieves better performance when  $m$  is a power of 2.

This function requires a buffer size returned by `gtsv2StridedBatch_bufferSizeExt()`. The address of `pBuffer` must be multiple of 128 bytes. If it is not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>n</code>	the size of the linear system (must be $\geq 3$ ).
<code>dl</code>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The lower diagonal $dl^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <code>dl+batchStride*i</code> in memory. Also, the first element of each lower diagonal must be zero.
<code>d</code>	<type> dense array containing the main diagonal of the tri-diagonal linear system. The main diagonal $d^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <code>d+batchStride*i</code> in memory.
<code>du</code>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The upper diagonal $du^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <code>du+batchStride*i</code> in memory. Also, the last element of each upper diagonal must be zero.
<code>x</code>	<type> dense array that contains the right-hand-side of the tri-diagonal linear system. The right-hand-side $x^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <code>x+batchStride*i</code> in memory.
<code>batchCount</code>	number of systems to solve.
<code>batchStride</code>	stride (number of elements) that separates the vectors of every system (must be at least <code>n</code> ).
<code>pBuffer</code>	buffer allocated by the user, the size is return by <code>gtsv2StridedBatch_bufferSizeExt</code> .

## Output

<code>x</code>	<type> dense array that contains the solution of the tri-diagonal linear system. The solution $x^{(i)}$ that corresponds to the $i^{\text{th}}$ linear system starts at location <code>x+batchStride*i</code> in memory.
----------------	--

See `cusparseStatus_t` for the description of the return status

## 11.4.4. `cusparse<t>gtsvInterleavedBatch()`

```

cusparseStatus_t
cusparseSgtsvInterleavedBatch_bufferSizeExt(cusparseHandle_t handle,
                                             int algo,
                                             int m,
                                             const float* dl,
                                             const float* d,
                                             const float* du,
                                             const float* x,
                                             int batchCount,
                                             size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseDgtsvInterleavedBatch_bufferSizeExt(cusparseHandle_t handle,
                                             int algo,
                                             int m,
                                             const double* dl,
                                             const double* d,
                                             const double* du,
                                             const double* x,
                                             int batchCount,
                                             size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseCgtsvInterleavedBatch_bufferSizeExt(cusparseHandle_t handle,
                                             int algo,
                                             int m,
                                             const cuComplex* dl,
                                             const cuComplex* d,
                                             const cuComplex* du,
                                             const cuComplex* x,
                                             int batchCount,
                                             size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseZgtsvInterleavedBatch_bufferSizeExt(cusparseHandle_t handle,
                                             int algo,
                                             int m,
                                             const cuDoubleComplex* dl,
                                             const cuDoubleComplex* d,
                                             const cuDoubleComplex* du,
                                             const cuDoubleComplex* x,
                                             int batchCount,
                                             size_t* pBufferSizeInBytes)

```

```

cusparseStatus_t
cusparseSgtsvInterleavedBatch(cusparseHandle_t handle,
                             int algo,
                             int m,
                             float* dl,
                             float* d,
                             float* du,
                             float* x,
                             int batchCount,
                             void* pBuffer)

cusparseStatus_t
cusparseDgtsvInterleavedBatch(cusparseHandle_t handle,
                             int algo,

```

This function computes the solution of multiple tridiagonal linear systems for  $i=0, \dots, \text{batchCount}$ :

$$A^{(i)} * x^{(i)} = b^{(i)}$$

The coefficient matrix **A** of each of these tri-diagonal linear system is defined with three vectors corresponding to its lower (**d1**), main (**d**), and upper (**du**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **x** overwrites right-hand-side matrix **B** on exit.

Assuming **A** is of size **m** and base-1, **d1**, **d** and **du** are defined by the following formula:

**d1(i) := A(i, i-1)** for  $i=1, 2, \dots, m$

The first element of **d1** is out-of-bound (**d1(1) := A(1, 0)**), so **d1(1) = 0**.

**d(i) = A(i, i)** for  $i=1, 2, \dots, m$

**du(i) = A(i, i+1)** for  $i=1, 2, \dots, m$

The last element of **du** is out-of-bound (**du(m) := A(m, m+1)**), so **du(m) = 0**.

The data layout is different from **gtsvStridedBatch** which aggregates all matrices one after another. Instead, **gtsvInterleavedBatch** gathers different matrices of the same element in a continuous manner. If **d1** is regarded as a 2-D array of size **m-by-batchCount**, **d1(:, j)** to store **j-th** matrix. **gtsvStridedBatch** uses column-major while **gtsvInterleavedBatch** uses row-major.

The routine provides three different algorithms, selected by parameter **algo**. The first algorithm is **cuThomas** provided by **Barcelona Supercomputing Center**. The second algorithm is LU with partial pivoting and last algorithm is QR. From stability perspective, cuThomas is not numerically stable because it does not have pivoting. LU with partial pivoting and QR are stable. From performance perspective, LU with partial pivoting and QR is about 10% to 20% slower than cuThomas.

This function requires a buffer size returned by

**gtsvInterleavedBatch\_bufferSizeExt()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Appendix F shows an example of **gtsvInterleavedBatch**. If the user prepares aggregate format, one can use **cublasXgemm** to get interleaved format. However such transformation takes time comparable to solver itself. To reach best performance, the user must prepare interleaved format explicitly.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>algo</b>	algo = 0: cuThomas (unstable algorithm); algo = 1: LU with pivoting (stable algorithm); algo = 2: QR (stable algorithm)
<b>m</b>	the size of the linear system.

<b>d1</b>	<type> dense array containing the lower diagonal of the tri-diagonal linear system. The first element of each lower diagonal must be zero.
<b>d</b>	<type> dense array containing the main diagonal of the tri-diagonal linear system.
<b>du</b>	<type> dense array containing the upper diagonal of the tri-diagonal linear system. The last element of each upper diagonal must be zero.
<b>x</b>	<type> dense right-hand-side array of dimensions <b>(batchCount, n)</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is return by <b>gtsvInterleavedBatch_bufferSizeExt</b> .

## Output

<b>x</b>	<type> dense solution array of dimensions <b>(batchCount, n)</b> .
----------	--

See [cusparseStatus\\_t](#) for the description of the return status

## 11.5. Batched Pentadiagonal Solve

Different algorithms for batched pentadiagonal solve are discussed in this section.

### 11.5.1. `cusparse<t>gpsvInterleavedBatch()`

```
cusparseStatus_t  
cusparseSgpsvInterleavedBatch(cusparseHandle_t handle,  
                                int          algo,  
                                int          m,  
                                float*       ds,  
                                float*       dl)
```

This function computes the solution of multiple penta-diagonal linear systems for  $i=0, \dots, \text{batchCount}$ :

$$A^{(i)} * x^{(i)} = b^{(i)}$$

The coefficient matrix **A** of each of these penta-diagonal linear system is defined with five vectors corresponding to its lower (**ds**, **d1**), main (**d**), and upper (**du**, **dw**) matrix diagonals; the right-hand sides are stored in the dense matrix **B**. Notice that solution **x** overwrites right-hand-side matrix **B** on exit.

Assuming **A** is of size **m** and base-1, **ds**, **d1**, **d**, **du** and **dw** are defined by the following formula:

**ds(i) := A(i, i-2)** for  $i=1, 2, \dots, m$

The first two elements of **ds** is out-of-bound (**ds(1) := A(1, -1)**, **ds(2) := A(2, 0)**), so **ds(1) = 0** and **ds(2) = 0**.

**d1(i) := A(i, i-1)** for  $i=1, 2, \dots, m$

The first element of **d1** is out-of-bound (**d1(1) := A(1, 0)**), so **d1(1) = 0**.

**d(i) = A(i, i)** for  $i=1, 2, \dots, m$

**du(i) = A(i, i+1)** for  $i=1, 2, \dots, m$

The last element of **du** is out-of-bound (**du(m) := A(m, m+1)**), so **du(m) = 0**.

**dw(i) = A(i, i+2)** for  $i=1, 2, \dots, m$

The last two elements of **dw** is out-of-bound (**dw(m-1) := A(m-1, m+1)**, **dw(m) := A(m, m+2)**), so **dw(m-1) = 0** and **dw(m) = 0**.

The data layout is the same as **gtsvStridedBatch**.

The routine is numerically stable because it uses QR to solve the linear system.

This function requires a buffer size returned by

**gpsvInterleavedBatch\_bufferSizeExt()**. The address of **pBuffer** must be multiple of 128 bytes. If it is not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Appendix G shows an example of **gpsvInterleavedBatch**. If the user prepares aggregate format, one can use **cublasXgemm** to get interleaved format. However such transformation takes time comparable to solver itself. To reach best performance, the user must prepare interleaved format explicitly.

The function supports the following properties if **pBuffer != NULL**

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>algo</b>	only support algo = 0 (QR)
<b>m</b>	the size of the linear system.

<b>ds</b>	<type> dense array containing the lower diagonal (distance 2 to the diagonal) of the penta-diagonal linear system. The first two elements must be zero.
<b>d1</b>	<type> dense array containing the lower diagonal (distance 1 to the diagonal) of the penta-diagonal linear system. The first element must be zero.
<b>d</b>	<type> dense array containing the main diagonal of the penta-diagonal linear system.
<b>du</b>	<type> dense array containing the upper diagonal (distance 1 to the diagonal) of the penta-diagonal linear system. The last element must be zero.
<b>dw</b>	<type> dense array containing the upper diagonal (distance 2 to the diagonal) of the penta-diagonal linear system. The last two elements must be zero.
<b>x</b>	<type> dense right-hand-side array of dimensions <b>(batchCount, n)</b> .
<b>pBuffer</b>	buffer allocated by the user, the size is return by <b>gpsvInterleavedBatch_bufferSizeExt</b> .

## Output

<b>x</b>	<type> dense solution array of dimensions <b>(batchCount, n)</b> .
----------	--

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 12.

# CUSPARSE REORDERINGS REFERENCE

This chapter describes the reordering routines used to manipulate sparse matrices.

## 12.1. `cusparse<t>csrcolor()`

```

cusparseStatus_t
cusparseScsrcolor(cusparseHandle_t           handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  const float*        csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  const float*        fractionToColor,
                  int*                ncolors,
                  int*                coloring,
                  int*                reordered,
                  cusparseColorInfo_t info)

cusparseStatus_t
cusparseDcsrcolor(cusparseHandle_t           handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  const double*       csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  const double*       fractionToColor,
                  int*                ncolors,
                  int*                coloring,
                  int*                reordered,
                  cusparseColorInfo_t info)

cusparseStatus_t
cusparseCcsrcolor(cusparseHandle_t           handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*    csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  const cuComplex*    fractionToColor,
                  int*                ncolors,
                  int*                coloring,
                  int*                reordered,
                  cusparseColorInfo_t info)

cusparseStatus_t
cusparseZcsrcolor(cusparseHandle_t           handle,
                  int                 m,
                  int                 nnz,
                  const cusparseMatDescr_t descrA,
                  const cuDoubleComplex* csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  const cuDoubleComplex* fractionToColor,
                  int*                ncolors,
                  int*                coloring,
                  int*                reordered,
                  cusparseColorInfo_t info)

```

This function performs the coloring of the adjacency graph associated with the matrix A stored in CSR format. The coloring is an assignment of colors (integer numbers) to nodes, such that neighboring nodes have distinct colors. An approximate coloring algorithm is used in this routine, and is stopped when a certain percentage of nodes has been colored. The rest of the nodes are assigned distinct colors (an increasing sequence of integers numbers, starting from the last integer used previously). The last two auxiliary routines can be used to extract the resulting number of colors, their assignment and the associated reordering. The reordering is such that nodes that have been assigned the same color are reordered to be next to each other.

The matrix A passed to this routine, must be stored as a general matrix and have a symmetric sparsity pattern. If the matrix is nonsymmetric the user should pass  $A+A^T$  as a parameter to this routine.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix <b>A</b> .
<b>nnz</b>	number of nonzero elements of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> - <b>csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .
<b>fractionToColor</b>	fraction of nodes to be colored, which should be in the interval [0.0,1.0], for example 0.8 implies that 80 percent of nodes will be colored.
<b>info</b>	structure with information to be passed to the coloring.

### Output

<b>ncolors</b>	The number of distinct colors used (at most the size of the matrix, but likely much smaller).
<b>coloring</b>	The resulting coloring permutation
<b>reordering</b>	The resulting reordering permutation (untouched if NULL)

See `cusparseStatus_t` for the description of the return status

# Chapter 13.

## CUSPARSE FORMAT CONVERSION

### REFERENCE

This chapter describes the conversion routines between different sparse and dense storage formats.

**coosort**, **csrssort**, **cscsort**, **csru2csr** and **csr2csc\_indexOnly** are sorting routines without malloc inside, the following table estimates the buffer size

<b>routine</b>	<b>buffer size</b>	<b>maximum problem size if buffer is limited by 2GB</b>
<b>coosort</b>	<b>&gt; 16*n bytes</b>	125M
<b>csrssort or cscsort</b>	<b>&gt; 20*n bytes</b>	100M
<b>csru2csr</b>	<b>'d' &gt; 28*n bytes ; 'z' &gt; 36*n bytes</b>	71M for 'd' and 55M for 'z'
<b>csr2csc_indexOnly</b>	<b>&gt; 16*n bytes</b>	125M

## 13.1. `cusparse<t>bsr2csr()`

```

cusparseStatus_t
cusparseSbsr2csr(cusparseHandle_t           handle,
                  cusparseDirection_t    dir,
                  int                   mb,
                  int                   nb,
                  const cusparseMatDescr_t descrA,
                  const float*           bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                    blockDim,
                  const cusparseMatDescr_t descrC,
                  float*                 csrValC,
                  int*                   csrRowPtrC,
                  int*                   csrColIndC)

cusparseStatus_t
cusparseDbsr2csr(cusparseHandle_t           handle,
                  cusparseDirection_t    dir,
                  int                   mb,
                  int                   nb,
                  const cusparseMatDescr_t descrA,
                  const double*          bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                    blockDim,
                  const cusparseMatDescr_t descrC,
                  double*                csrValC,
                  int*                   csrRowPtrC,
                  int*                   csrColIndC)

cusparseStatus_t
cusparseCbsr2csr(cusparseHandle_t           handle,
                  cusparseDirection_t    dir,
                  int                   mb,
                  int                   nb,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*       bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                    blockDim,
                  const cusparseMatDescr_t descrC,
                  cuComplex*             csrValC,
                  int*                   csrRowPtrC,
                  int*                   csrColIndC)

cusparseStatus_t
cusparseZbsr2csr(cusparseHandle_t           handle,
                  cusparseDirection_t    dir,
                  int                   mb,
                  int                   nb,
                  const cusparseMatDescr_t descrA,
                  const cuDoubleComplex* bsrValA,
                  const int*              bsrRowPtrA,
                  const int*              bsrColIndA,
                  int                    blockDim,
                  const cusparseMatDescr_t descrC,
                  cuDoubleComplex*        csrValC,
                  int*                   csrRowPtrC,
                  int*                   csrColIndC)

```

This function converts a sparse matrix in BSR format that is defined by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA**) into a sparse matrix in CSR format that is defined by arrays **csrValC**, **csrRowPtrC**, and **csrColIndC**.

Let **m** ( $=mb * \text{blockDim}$ ) be the number of rows of **A** and **n** ( $=nb * \text{blockDim}$ ) be number of columns of **A**, then **A** and **C** are  $m * n$  sparse matrices. The BSR format of **A** contains **nnzb** ( $=\text{bsrRowPtrA}[mb] - \text{bsrRowPtrA}[0]$ ) nonzero blocks, whereas the sparse matrix **A** contains **nnz** ( $=nnzb * \text{blockDim} * \text{blockDim}$ ) elements. The user must allocate enough space for arrays **csrRowPtrC**, **csrColIndC**, and **csrValC**. The requirements are as follows:

**csrRowPtrC** of **m+1** elements

**csrValC** of **nnz** elements

**csrColIndC** of **nnz** elements

The general procedure is as follows:

```
// Given BSR format (bsrRowPtrA, bsrColIndA, bsrValA) and
// blocks of BSR format are stored in column-major order.
cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;
int m = mb*blockDim;
int nnzb = bsrRowPtrA[mb] - bsrRowPtrA[0]; // number of blocks
int nnz = nnzb * blockDim * blockDim; // number of elements
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnz);
cudaMalloc((void**)&csrValC, sizeof(float)*nnz);
cusparseSbsr2csr(handle, dir, mb, nb,
                 descrA,
                 bsrValA, bsrRowPtrA, bsrColIndA,
                 blockDim,
                 descrC,
                 csrValC, csrRowPtrC, csrColIndC);
```

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution if **blockDim == 1**
- ▶ The routine does *not* support CUDA graph capture if **blockDim == 1**

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dir</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows of sparse matrix <b>A</b> .
<b>nb</b>	number of block columns of sparse matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> .
<b>bsrValA</b>	<type> array of $nnzb * \text{blockDim} * \text{blockDim}$ nonzero elements of matrix <b>A</b> .
<b>bsrRowPtrA</b>	integer array of <b>mb+1</b> elements that contains the start of every block row and the end of the last block row plus one of matrix <b>A</b> .
<b>bsrColIndA</b>	integer array of <b>nnzb</b> column indices of the nonzero blocks of matrix <b>A</b> .

<b>blockDim</b>	block dimension of sparse matrix <b>A</b> .
<b>descrC</b>	the descriptor of matrix <b>C</b> .

## Output

<b>csrValC</b>	<type> array of <b>nnz</b> ( $=\text{csrRowPtrC}[m] - \text{csrRowPtrC}[0]$ ) nonzero elements of matrix <b>C</b> .
<b>csrRowPtrC</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one of matrix <b>C</b> .
<b>csrColIndC</b>	integer array of <b>nnz</b> column indices of the nonzero elements of matrix <b>C</b> .

See [cusparseStatus\\_t](#) for the description of the return status

## 13.2. `cusparse<t>gebsr2gebsc()`

```

cusparseStatus_t
cusparseSgebsr2gebsc_bufferSize(cusparseHandle_t handle,
                                int mb,
                                int nb,
                                int nnzb,
                                const float* bsrVal,
                                const int* bsrRowPtr,
                                const int* bsrColInd,
                                int rowBlockDim,
                                int colBlockDim,
                                int* pBufferSize)

cusparseStatus_t
cusparseDgebsr2gebsc_bufferSize(cusparseHandle_t handle,
                                int mb,
                                int nb,
                                int nnzb,
                                const double* bsrVal,
                                const int* bsrRowPtr,
                                const int* bsrColInd,
                                int rowBlockDim,
                                int colBlockDim,
                                int* pBufferSize)

cusparseStatus_t
cusparseCgebsr2gebsc_bufferSize(cusparseHandle_t handle,
                                int mb,
                                int nb,
                                int nnzb,
                                const cuComplex* bsrVal,
                                const int* bsrRowPtr,
                                const int* bsrColInd,
                                int rowBlockDim,
                                int colBlockDim,
                                int* pBufferSize)

cusparseStatus_t
cusparseZgebsr2gebsc_bufferSize(cusparseHandle_t handle,
                                int mb,
                                int nb,
                                int nnzb,
                                const cuDoubleComplex* bsrVal,
                                const int* bsrRowPtr,
                                const int* bsrColInd,
                                int rowBlockDim,
                                int colBlockDim,
                                int* pBufferSize)

```

```

cusparseStatus_t
cusparseSgebsr2gebsc(cusparseHandle_t handle,
                      int mb,
                      int nb,
                      int nnzb,
                      const float* bsrVal,
                      const int* bsrRowPtr,
                      const int* bsrColInd,
                      int rowBlockDim,
                      int colBlockDim,
                      float* bscVal,
                      int* bscRowInd,
                      int* bscColPtr,
                      cusparseAction_t copyValues,

```

This function can be seen as the same as `csr2csc()` when each block of size `rowBlockDim*colBlockDim` is regarded as a scalar.

This sparsity pattern of the result matrix can also be seen as the transpose of the original sparse matrix, but the memory layout of a block does not change.

The user must call `gebsr2gebsc_bufferSize()` to determine the size of the buffer required by `gebsr2gebsc()`, allocate the buffer, and pass the buffer pointer to `gebsr2gebsc()`.

- ▶ The routine requires no extra storage if `pBuffer != NULL`
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>mb</code>	number of block rows of sparse matrix <code>A</code> .
<code>nb</code>	number of block columns of sparse matrix <code>A</code> .
<code>nnzb</code>	number of nonzero blocks of matrix <code>A</code> .
<code>bsrVal</code>	<type> array of <code>nnzb*rowBlockDim*colBlockDim</code> nonzero elements of matrix <code>A</code> .
<code>bsrRowPtr</code>	integer array of <code>mb+1</code> elements that contains the start of every block row and the end of the last block row plus one.
<code>bsrColInd</code>	integer array of <code>nnzb</code> column indices of the non-zero blocks of matrix <code>A</code> .
<code>rowBlockDim</code>	number of rows within a block of <code>A</code> .
<code>colBlockDim</code>	number of columns within a block of <code>A</code> .
<code>copyValues</code>	<code>CUSPARSE_ACTION_SYMBOLIC</code> or <code>CUSPARSE_ACTION_NUMERIC</code> .
<code>baseIdx</code>	<code>CUSPARSE_INDEX_BASE_ZERO</code> or <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>pBufferSize</code>	host pointer containing number of bytes of the buffer used in <code>gebsr2gebsc()</code> .
<code>pBuffer</code>	buffer allocated by the user; the size is return by <code>gebsr2gebsc_bufferSize()</code> .

### Output

<code>bscVal</code>	<type> array of <code>nnzb*rowBlockDim*colBlockDim</code> non-zero elements of matrix <code>A</code> . It is only filled-in if <code>copyValues</code> is set to <code>CUSPARSE_ACTION_NUMERIC</code> .
<code>bscRowInd</code>	integer array of <code>nnzb</code> row indices of the non-zero blocks of matrix <code>A</code> .

<b>bscColPtr</b>	integer array of <code>nb+1</code> elements that contains the start of every block column and the end of the last block column plus one.
------------------	--

See `cusparseStatus_t` for the description of the return status

### 13.3. `cusparse<t>gebsr2gebsr()`

```

cusparseStatus_t
cusparseSgebsr2gebsr_bufferSize(cusparseHandle_t
                                cusparseDirection_t
                                int
                                int
                                int
                                const cusparseMatDescr_t
                                const float*
                                const int*
                                const int*
                                int
                                int
                                int
                                int
                                int
                                int*)
                                handle,
                                dir,
                                mb,
                                nb,
                                nnzb,
                                descrA,
                                bsrValA,
                                bsrRowPtrA,
                                bsrColIndA,
                                rowBlockDimA,
                                colBlockDimA,
                                rowBlockDimC,
                                colBlockDimC,
                                pBufferSize)

cusparseStatus_t
cusparseDgebsr2gebsr_bufferSize(cusparseHandle_t
                                cusparseDirection_t
                                int
                                int
                                int
                                const cusparseMatDescr_t
                                const double*
                                const int*
                                const int*
                                int
                                int
                                int
                                int
                                int
                                int*)
                                handle,
                                dir,
                                mb,
                                nb,
                                nnzb,
                                descrA,
                                bsrValA,
                                bsrRowPtrA,
                                bsrColIndA,
                                rowBlockDimA,
                                colBlockDimA,
                                rowBlockDimC,
                                colBlockDimC,
                                pBufferSize)

cusparseStatus_t
cusparseCgebsr2gebsr_bufferSize(cusparseHandle_t
                                cusparseDirection_t
                                int
                                int
                                int
                                const cusparseMatDescr_t
                                const cuComplex*
                                const int*
                                const int*
                                int
                                int
                                int
                                int
                                int
                                int*)
                                handle,
                                dir,
                                mb,
                                nb,
                                nnzb,
                                descrA,
                                bsrValA,
                                bsrRowPtrA,
                                bsrColIndA,
                                rowBlockDimA,
                                colBlockDimA,
                                rowBlockDimC,
                                colBlockDimC,
                                pBufferSize)

cusparseStatus_t
cusparseZgebsr2gebsr_bufferSize(cusparseHandle_t
                                cusparseDirection_t
                                int
                                int
                                int
                                const cusparseMatDescr_t
                                const cuDoubleComplex*
                                const int*
                                const int*
                                int
                                int
                                int
                                int
                                int
                                int*)
                                handle,
                                dir,
                                mb,
                                nb,
                                nnzb,
                                descrA,
                                bsrValA,
                                bsrRowPtrA,
                                bsrColIndA,
                                rowBlockDimA,
                                colBlockDimA,
                                rowBlockDimC,
                                colBlockDimC,
                                pBufferSize)

```

This function converts a sparse matrix in general BSR format that is defined by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA** into a sparse matrix in another general BSR format that is defined by arrays **bsrValC**, **bsrRowPtrC**, and **bsrColIndC**.

If **rowBlockDimA=1** and **colBlockDimA=1**, **cusparse[S|D|C|Z]gebsr2gebsr()** is the same as **cusparse[S|D|C|Z]csr2gebsr()**.

If **rowBlockDimC=1** and **colBlockDimC=1**, **cusparse[S|D|C|Z]gebsr2gebsr()** is the same as **cusparse[S|D|C|Z]gebsr2csr()**.

**A** is an  $m \times n$  sparse matrix where  $m (=mb * \text{rowBlockDim})$  is the number of rows of **A**, and  $n (=nb * \text{colBlockDim})$  is the number of columns of **A**. The general BSR format of **A** contains  $nnzb (=bsrRowPtrA[mb] - bsrRowPtrA[0])$  nonzero blocks. The matrix **C** is also general BSR format with a different block size,  $\text{rowBlockDimC} * \text{colBlockDimC}$ . If  $m$  is not a multiple of **rowBlockDimC**, or  $n$  is not a multiple of **colBlockDimC**, zeros are filled in. The number of block rows of **C** is  $mc (= (m + \text{rowBlockDimC} - 1) / \text{rowBlockDimC})$ . The number of block columns of **C** is  $nc (= (n + \text{colBlockDimC} - 1) / \text{colBlockDimC})$ . The number of nonzero blocks of **C** is  $nnzc$ .

The implementation adopts a two-step approach to do the conversion.

First, the user allocates **bsrRowPtrC** of  $mc + 1$  elements and uses function **cusparseXgebsr2gebsrNnz()** to determine the number of nonzero block columns per block row of matrix **C**. Second, the user gathers  $nnzc$  (number of non-zero block columns of matrix **C**) from either ( $nnzc = *nnzTotalDevHostPtr$ ) or ( $nnzc = bsrRowPtrC[mc] - bsrRowPtrC[0]$ ) and allocates **bsrValC** of  $nnzc * \text{rowBlockDimC} * \text{colBlockDimC}$  elements and **bsrColIndC** of  $nnzc$  integers. Finally the function **cusparse[S|D|C|Z]gebsr2gebsr()** is called to complete the conversion.

The user must call **gebsr2gebsr\_bufferSize()** to know the size of the buffer required by **gebsr2gebsr()**, allocate the buffer, and pass the buffer pointer to **gebsr2gebsr()**.

The general procedure is as follows:

```
// Given general BSR format (bsrRowPtrA, bsrColIndA, bsrValA) and
// blocks of BSR format are stored in column-major order.
cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;
int base, nnzc;
int m = mb*rowBlockDimA;
int n = nb*colBlockDimA;
int mc = (m+rowBlockDimC-1)/rowBlockDimC;
int nc = (n+colBlockDimC-1)/colBlockDimC;
int bufferSize;
void *pBuffer;
cusparseSgebsr2gebsr_bufferSize(handle, dir, mb, nb, nnzb,
    descrA, bsrValA, bsrRowPtrA, bsrColIndA,
    rowBlockDimA, colBlockDimA,
    rowBlockDimC, colBlockDimC,
    &bufferSize);
cudaMalloc((void**)&pBuffer, bufferSize);
cudaMalloc((void**)&bsrRowPtrC, sizeof(int)*(mc+1));
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzc;
cusparseXgebsr2gebsrNnz(handle, dir, mb, nb, nnzb,
    descrA, bsrRowPtrA, bsrColIndA,
    rowBlockDimA, colBlockDimA,
    descrC, bsrRowPtrC,
    rowBlockDimC, colBlockDimC,
    nnzTotalDevHostPtr,
    pBuffer);
if (NULL != nnzTotalDevHostPtr){
    nnzc = *nnzTotalDevHostPtr;
} else{
    cudaMemcpy(&nnzc, bsrRowPtrC+mc, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&base, bsrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzc -= base;
}
cudaMalloc((void**)&bsrColIndC, sizeof(int)*nnzc);
cudaMalloc((void**)&bsrValC, sizeof(float)*(rowBlockDimC*colBlockDimC)*nnzc);
cusparseSgebsr2gebsr(handle, dir, mb, nb, nnzb,
    descrA, bsrValA, bsrRowPtrA, bsrColIndA,
    rowBlockDimA, colBlockDimA,
    descrC, bsrValC, bsrRowPtrC, bsrColIndC,
    rowBlockDimC, colBlockDimC,
    pBuffer);
```

- ▶ The routines require no extra storage if **pBuffer != NULL**
- ▶ The routines do *not* support asynchronous execution
- ▶ The routines do *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dir</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows of sparse matrix <b>A</b> .
<b>nb</b>	number of block columns of sparse matrix <b>A</b> .
<b>nnzb</b>	number of nonzero blocks of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are

	CUSPARSE_INDEX_BASE_ZERO and CUSPARSE_INDEX_BASE_ONE.
<b>bsrValA</b>	<type> array of $\text{nnzb} * \text{rowBlockDimA} * \text{colBlockDimA}$ non-zero elements of matrix A.
<b>bsrRowPtrA</b>	integer array of $\text{mb} + 1$ elements that contains the start of every block row and the end of the last block row plus one of matrix A.
<b>bsrColIndA</b>	integer array of $\text{nnzb}$ column indices of the non-zero blocks of matrix A.
<b>rowBlockDimA</b>	number of rows within a block of A.
<b>colBlockDimA</b>	number of columns within a block of A.
<b>descrC</b>	the descriptor of matrix c. The supported matrix type is CUSPARSE_MATRIX_TYPE_GENERAL. Also, the supported index bases are CUSPARSE_INDEX_BASE_ZERO and CUSPARSE_INDEX_BASE_ONE.
<b>rowBlockDimC</b>	number of rows within a block of c.
<b>colBlockDimC</b>	number of columns within a block of c.
<b>pBufferSize</b>	host pointer containing number of bytes of the buffer used in <code>gebsr2gebsr()</code> .
<b>pBuffer</b>	buffer allocated by the user; the size is return by <code>gebsr2gebsr_bufferSize()</code> .

## Output

<b>bsrValC</b>	<type> array of $\text{nnzc} * \text{rowBlockDimC} * \text{colBlockDimC}$ non-zero elements of matrix c.
<b>bsrRowPtrC</b>	integer array of $\text{mc} + 1$ elements that contains the start of every block row and the end of the last block row plus one of matrix c.
<b>bsrColIndC</b>	integer array of $\text{nnzc}$ block column indices of the nonzero blocks of matrix c.
<b>nnzTotalDevHostPtr</b>	total number of nonzero blocks of c. <code>*nnzTotalDevHostPtr</code> is the same as <code>bsrRowPtrC[mc] - bsrRowPtrC[0]</code> .

See `cusparseStatus_t` for the description of the return status

## 13.4. `cusparse<t>gebsr2csr()`

```

cusparseStatus_t
cusparseSgebsr2csr(cusparseHandle_t          handle,
                    cusparseDirection_t    dir,
                    int                  mb,
                    int                  nb,
                    const cusparseMatDescr_t descrA,
                    const float*           bsrValA,
                    const int*              bsrRowPtrA,
                    const int*              bsrColIndA,
                    int                   rowBlockDim,
                    int                   colBlockDim,
                    const cusparseMatDescr_t descrC,
                    float*                csrValC,
                    int*                 csrRowPtrC,
                    int*                 csrColIndC)

cusparseStatus_t
cusparseDgebsr2csr(cusparseHandle_t          handle,
                    cusparseDirection_t    dir,
                    int                  mb,
                    int                  nb,
                    const cusparseMatDescr_t descrA,
                    const double*          bsrValA,
                    const int*              bsrRowPtrA,
                    const int*              bsrColIndA,
                    int                   rowBlockDim,
                    int                   colBlockDim,
                    const cusparseMatDescr_t descrC,
                    double*               csrValC,
                    int*                 csrRowPtrC,
                    int*                 csrColIndC)

cusparseStatus_t
cusparseCgebsr2csr(cusparseHandle_t          handle,
                    cusparseDirection_t    dir,
                    int                  mb,
                    int                  nb,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*        bsrValA,
                    const int*              bsrRowPtrA,
                    const int*              bsrColIndA,
                    int                   rowBlockDim,
                    int                   colBlockDim,
                    const cusparseMatDescr_t descrC,
                    cuComplex*             csrValC,
                    int*                 csrRowPtrC,
                    int*                 csrColIndC)

cusparseStatus_t
cusparseZgebsr2csr(cusparseHandle_t          handle,
                    cusparseDirection_t    dir,
                    int                  mb,
                    int                  nb,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* bsrValA,
                    const int*              bsrRowPtrA,
                    const int*              bsrColIndA,
                    int                   rowBlockDim,
                    int                   colBlockDim,
                    const cusparseMatDescr_t descrC,
                    cuDoubleComplex*        csrValC,
                    int*                 csrRowPtrC,
                    int*                 csrColIndC)

```

This function converts a sparse matrix in general BSR format that is defined by the three arrays **bsrValA**, **bsrRowPtrA**, and **bsrColIndA** into a sparse matrix in CSR format that is defined by arrays **csrValC**, **csrRowPtrC**, and **csrColIndC**.

Let **m** ( $=mb * \text{rowBlockDim}$ ) be number of rows of **A** and **n** ( $=nb * \text{colBlockDim}$ ) be number of columns of **A**, then **A** and **C** are  $m * n$  sparse matrices. The general BSR format of **A** contains **nnzb** ( $=\text{bsrRowPtrA}[mb] - \text{bsrRowPtrA}[0]$ ) non-zero blocks, whereas sparse matrix **A** contains **nnz** ( $=nnzb * \text{rowBlockDim} * \text{colBlockDim}$ ) elements. The user must allocate enough space for arrays **csrRowPtrC**, **csrColIndC**, and **csrValC**. The requirements are as follows:

**csrRowPtrC** of **m+1** elements

**csrValC** of **nnz** elements

**csrColIndC** of **nnz** elements

The general procedure is as follows:

```
// Given general BSR format (bsrRowPtrA, bsrColIndA, bsrValA) and
// blocks of BSR format are stored in column-major order.
cuspDirection_t dir = CUSPARSE_DIRECTION_COLUMN;
int m = mb*rowBlockDim;
int n = nb*colBlockDim;
int nnzb = bsrRowPtrA[mb] - bsrRowPtrA[0]; // number of blocks
int nnz = nnzb * rowBlockDim * colBlockDim; // number of elements
cudaMalloc((void**)&csrRowPtrC, sizeof(int)*(m+1));
cudaMalloc((void**)&csrColIndC, sizeof(int)*nnz);
cudaMalloc((void**)&csrValC, sizeof(float)*nnz);
cuspSgebsr2csr(handle, dir, mb, nb,
    descrA,
    bsrValA, bsrRowPtrA, bsrColIndA,
    rowBlockDim, colBlockDim,
    descrC,
    csrValC, csrRowPtrC, csrColIndC);
```

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>dir</b>	storage format of blocks, either <b>CUSPARSE_DIRECTION_ROW</b> or <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>mb</b>	number of block rows of sparse matrix <b>A</b> .
<b>nb</b>	number of block columns of sparse matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>bsrValA</b>	<type> array of $nnzb * \text{rowBlockDim} * \text{colBlockDim}$ non-zero elements of matrix <b>A</b> .

<b>bsrRowPtrA</b>	integer array of $\text{mb}+1$ elements that contains the start of every block row and the end of the last block row plus one of matrix <b>A</b> .
<b>bsrColIndA</b>	integer array of $\text{nnzb}$ column indices of the non-zero blocks of matrix <b>A</b> .
<b>rowBlockDim</b>	number of rows within a block of <b>A</b> .
<b>colBlockDim</b>	number of columns within a block of <b>A</b> .
<b>descrC</b>	the descriptor of matrix c. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .

## Output

<b>csrValC</b>	<type> array of $\text{nnz}$ non-zero elements of matrix C.
<b>csrRowPtrC</b>	integer array of $\text{m}+1$ elements that contains the start of every row and the end of the last row plus one of matrix c.
<b>csrColIndC</b>	integer array of $\text{nnz}$ column indices of the non-zero elements of matrix c.

See [cusparseStatus\\_t](#) for the description of the return status

## 13.5. `cusparse<t>csr2gebsr()`

```
cusparseStatus_t
cusparseScsr2gebsr_bufferSize(cusparseHandle_t handle,
                               cusparseDirection_t dir,
                               int m,
                               int n,
                               const cusparseMatDescr_t descrA,
                               const float* csrValA,
                               const int* csrRowPtrA,
                               const int* csrColIndA,
                               int rowBlockDim,
                               int colBlockDim,
                               int* pBufferSize)

cusparseStatus_t
cusparseDcsr2gebsr_bufferSize(cusparseHandle_t handle,
                               cusparseDirection_t dir,
                               int m,
                               int n,
                               const cusparseMatDescr_t descrA,
                               const double* csrValA,
                               const int* csrRowPtrA,
                               const int* csrColIndA,
                               int rowBlockDim,
                               int colBlockDim,
                               int* pBufferSize)

cusparseStatus_t
cusparseCcsr2gebsr_bufferSize(cusparseHandle_t handle,
                               cusparseDirection_t dir,
                               int m,
                               int n,
                               const cusparseMatDescr_t descrA,
                               const cuComplex* csrValA,
                               const int* csrRowPtrA,
                               const int* csrColIndA,
                               int rowBlockDim,
                               int colBlockDim,
                               int* pBufferSize)

cusparseStatus_t
cusparseZcsr2gebsr_bufferSize(cusparseHandle_t handle,
                               cusparseDirection_t dir,
                               int m,
                               int n,
                               const cusparseMatDescr_t descrA,
                               const cuDoubleComplex* csrValA,
                               const int* csrRowPtrA,
                               const int* csrColIndA,
                               int rowBlockDim,
                               int colBlockDim,
                               int* pBufferSize)
```

```
cusparseStatus_t  
cusparseXcsr2gebsrNnz(cusparseHandle_t handle,  
                      cusparseDirection_t dir,  
                      int m,  
                      int n,  
                      const cusparseMatDescr_t descrA,  
                      const int* csrRowPtrA,  
                      const int* csrColIndA,  
                      const cusparseMatDescr_t descrC,  
                      int* bsrRowPtrC,
```

This function converts a sparse matrix **A** in CSR format (that is defined by arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**) into a sparse matrix **C** in general BSR format (that is defined by the three arrays **bsrValC**, **bsrRowPtrC**, and **bsrColIndC**).

The matrix **A** is a  $m \times n$  sparse matrix and matrix **C** is a  $(mb * \text{rowBlockDim}) \times (nb * \text{colBlockDim})$  sparse matrix, where **mb** ( $= (m + \text{rowBlockDim} - 1) / \text{rowBlockDim}$ ) is the number of block rows of **C**, and **nb** ( $= (n + \text{colBlockDim} - 1) / \text{colBlockDim}$ ) is the number of block columns of **C**.

The block of **C** is of size **rowBlockDim\*colBlockDim**. If **m** is not multiple of **rowBlockDim** or **n** is not multiple of **colBlockDim**, zeros are filled in.

The implementation adopts a two-step approach to do the conversion. First, the user allocates **bsrRowPtrC** of **mb+1** elements and uses function **cusparseXcsr2gebsrNnz()** to determine the number of nonzero block columns per block row. Second, the user gathers **nnzb** (number of nonzero block columns of matrix **C**) from either **(nnzb=\*nnzTotalDevHostPtr)** or **(nnzb=bsrRowPtrC[mb]-bsrRowPtrC[0])** and allocates **bsrValC** of **nnzb\*rowBlockDim\*colBlockDim** elements and **bsrColIndC** of **nnzb** integers. Finally function **cusparse[S|D|C|Z]csr2gebsr()** is called to complete the conversion.

The user must obtain the size of the buffer required by **csr2gebsr()** by calling **csr2gebsr\_bufferSize()**, allocate the buffer, and pass the buffer pointer to **csr2gebsr()**.

The general procedure is as follows:

```
// Given CSR format (csrRowPtrA, csrColIndA, csrValA) and
// blocks of BSR format are stored in column-major order.
cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;
int base, nnzb;
int mb = (m + rowBlockDim-1)/rowBlockDim;
int nb = (n + colBlockDim-1)/colBlockDim;
int bufferSize;
void *pBuffer;
cusparseScsr2gebsr_bufferSize(handle, dir, m, n,
    descrA, csrValA, csrRowPtrA, csrColIndA,
    rowBlockDim, colBlockDim,
    &bufferSize);
cudaMalloc((void**)&pBuffer, bufferSize);
cudaMalloc((void**)&bsrRowPtrC, sizeof(int) * (mb+1));
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzb;
cusparseXcsr2gebsrNnz(handle, dir, m, n,
    descrA, csrRowPtrA, csrColIndA,
    descrC, bsrRowPtrC, rowBlockDim, colBlockDim,
    nnzTotalDevHostPtr,
    pBuffer);
if (NULL != nnzTotalDevHostPtr){
    nnzb = *nnzTotalDevHostPtr;
} else{
    cudaMemcpy(&nnzb, bsrRowPtrC+mb, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&base, bsrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzb -= base;
}
cudaMalloc((void**)&bsrColIndC, sizeof(int)*nnzb);
cudaMalloc((void**)&bsrValC, sizeof(float)*(rowBlockDim*colBlockDim)*nnzb);
cusparseScsr2gebsr(handle, dir, m, n,
    descrA,
    csrValA, csrRowPtrA, csrColIndA,
    descrC,
    bsrValC, bsrRowPtrC, bsrColIndC,
    rowBlockDim, colBlockDim,
    pBuffer);
```

The routine **`cusparseXcsr2gebsrNnz()`** has the following properties:

- ▶ The routine requires no extra storage
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine **`cusparse<t>csr2gebsr()`** has the following properties:

- ▶ The routine requires no extra storage if `pBuffer != NULL`
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>dir</code>	storage format of blocks, either <code>CUSPARSE_DIRECTION_ROW</code> or <code>CUSPARSE_DIRECTION_COLUMN</code> .
<code>m</code>	number of rows of sparse matrix <code>A</code> .
<code>n</code>	number of columns of sparse matrix <code>A</code> .

<b>descrA</b>	the descriptor of matrix A. The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<b>csrValA</b>	<type> array of <code>nnz</code> nonzero elements of matrix A.
<b>csrRowPtrA</b>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one of matrix A.
<b>csrColIndA</b>	integer array of <code>nnz</code> column indices of the nonzero elements of matrix A.
<b>descrC</b>	the descriptor of matrix c. The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<b>rowBlockDim</b>	number of rows within a block of c.
<b>colBlockDim</b>	number of columns within a block of c.
<b>pBuffer</b>	buffer allocated by the user, the size is return by <code>csr2gebsr_bufferSize()</code> .

## Output

<b>bsrValC</b>	<type> array of <code>nnzb*rowBlockDim*colBlockDim</code> nonzero elements of matrix c.
<b>bsrRowPtrC</b>	integer array of <code>mb+1</code> elements that contains the start of every block row and the end of the last block row plus one of matrix c.
<b>bsrColIndC</b>	integer array of <code>nnzb</code> column indices of the nonzero blocks of matrix c.
<b>nnzTotalDevHostPtr</b>	total number of nonzero blocks of matrix c. Pointer <code>nnzTotalDevHostPtr</code> can point to a device memory or host memory.

See `cusparseStatus_t` for the description of the return status

## 13.6. `cusparse<t>coo2csr()`

```
cusparseStatus_t
cusparseXcoo2csr(cusparseHandle_t handle,
                  const int* cooRowInd,
                  int nnz,
                  int m,
                  int* csrRowPtr,
                  cusparseIndexBase_t idxBase)
```

This function converts the array containing the uncompressed row indices (corresponding to COO format) into an array of compressed row pointers (corresponding to CSR format).

It can also be used to convert the array containing the uncompressed column indices (corresponding to COO format) into an array of column pointers (corresponding to CSC format).

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>cooRowInd</b>	integer array of <b>nnz</b> uncompressed row indices.
<b>nnz</b>	number of non-zeros of the sparse matrix (that is also the length of array <b>cooRowInd</b> ).
<b>m</b>	number of rows of matrix <b>A</b> .
<b>idxBase</b>	<b>CUSPARSE_INDEX_BASE_ZERO</b> or <b>CUSPARSE_INDEX_BASE_ONE</b> .

### Output

<b>csrRowPtr</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
------------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 13.7. `cusparse<t>csc2dense()`

```

cusparseStatus_t
cusparseScsc2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const float*          cscValA,
                    const int*             cscRowIndA,
                    const int*             cscColPtrA,
                    float*                A,
                    int                  lda)

cusparseStatus_t
cusparseDcsc2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const double*          cscValA,
                    const int*              cscRowIndA,
                    const int*              cscColPtrA,
                    double*                A,
                    int                  lda)

cusparseStatus_t
cusparseCcsc2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*        cscValA,
                    const int*              cscRowIndA,
                    const int*              cscColPtrA,
                    cuComplex*              A,
                    int                  lda)

cusparseStatus_t
cusparseZcsc2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* cscValA,
                    const int*              cscRowIndA,
                    const int*              cscColPtrA,
                    cuDoubleComplex*        A,
                    int                  lda)

```

This function converts the sparse matrix in CSC format that is defined by the three arrays **`cscValA`**, **`cscColPtrA`**, and **`cscRowIndA`** into the matrix **`A`** in dense format. The dense matrix **`A`** is filled in with the values of the sparse matrix and with zeros elsewhere.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>cscValA</b>	<type> array of <b>nnz</b> (= <b>cscColPtrA(m) - cscColPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>cscRowIndA</b>	integer array of <b>nnz</b> (= <b>cscColPtrA(m) - cscColPtrA(0)</b> ) row indices of the nonzero elements of matrix <b>A</b> .
<b>cscColPtrA</b>	integer array of <b>n+1</b> elements that contains the start of every row and the end of the last column plus one.
<b>lda</b>	leading dimension of dense array <b>A</b> .

## Output

<b>A</b>	array of dimensions ( <b>lda, n</b> ) that is filled in with the values of the sparse matrix.
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 13.8. `cusparse<t>csc2hyb()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsc2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const float*        cscValA,
                  const int*          cscRowIndA,
                  const int*          cscColPtrA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseDcsc2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const double*       cscValA,
                  const int*          cscRowIndA,
                  const int*          cscColPtrA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseCcsc2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*    cscValA,
                  const int*          cscRowIndA,
                  const int*          cscColPtrA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseZcsc2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const cuDoubleComplex* cscValA,
                  const int*          cscRowIndA,
                  const int*          cscColPtrA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

```

This function converts a sparse matrix in CSC format into a sparse matrix in HYB format. It assumes that the **hybA** parameter has been initialized with the `cusparseCreateHybMat()` routine before calling this function.

- ▶ This function requires temporary extra storage that is allocated internally

- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix <code>A</code> .
<code>n</code>	number of columns of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>cscValA</code>	<type> array of <code>nnz</code> (= <code>cscColPtrA(m) - cscColPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>cscRowIndA</code>	integer array of <code>nnz</code> (= <code>cscColPtrA(m) - cscColPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>cscColPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>userEllWidth</code>	width of the regular (ELL) part of the matrix in HYB format, which should be less than the maximum number of nonzeros per row and is only required if <code>partitionType == CUSPARSE_HYB_PARTITION_USER</code> .
<code>partitionType</code>	partitioning method to be used in the conversion (please refer to <code>cusparseHybPartition_t</code> for details).

### Output

<code>hybA</code>	the matrix <code>A</code> in HYB storage format.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 13.9. `cusparse<t>csr2bsr()`

```

cusparseStatus_t
cusparseXcsr2bsrNnz(cusparseHandle_t          handle,
                     cusparseDirection_t    dir,
                     int                  m,
                     int                  n,
                     const cusparseMatDescr_t descrA,
                     const int*            csrRowPtrA,
                     const int*            csrColIndA,
                     int                  blockDim,
                     const cusparseMatDescr_t descrC,
                     int*                 bsrRowPtrC,
                     int*                 nnzTotalDevHostPtr)

cusparseStatus_t
cusparseScsr2bsr(cusparseHandle_t          handle,
                  cusparseDirection_t    dir,
                  int                  m,
                  int                  n,
                  const cusparseMatDescr_t descrA,
                  const float*           csrValA,
                  const int*             csrRowPtrA,
                  const int*             csrColIndA,
                  int                  blockDim,
                  const cusparseMatDescr_t descrC,
                  float*                bsrValC,
                  int*                 bsrRowPtrC,
                  int*                 bsrColIndC)

cusparseStatus_t
cusparseDcsr2bsr(cusparseHandle_t          handle,
                  cusparseDirection_t    dir,
                  int                  m,
                  int                  n,
                  const cusparseMatDescr_t descrA,
                  const double*          csrValA,
                  const int*             csrRowPtrA,
                  const int*             csrColIndA,
                  int                  blockDim,
                  const cusparseMatDescr_t descrC,
                  double*               bsrValC,
                  int*                 bsrRowPtrC,
                  int*                 bsrColIndC)

cusparseStatus_t
cusparseCcsr2bsr(cusparseHandle_t          handle,
                  cusparseDirection_t    dir,
                  int                  m,
                  int                  n,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*       csrValA,
                  const int*             csrRowPtrA,
                  const int*             csrColIndA,
                  int                  blockDim,
                  const cusparseMatDescr_t descrC,
                  cuComplex*            bsrValC,
                  int*                 bsrRowPtrC,
                  int*                 bsrColIndC)

cusparseStatus_t
cusparseZcsr2bsr(cusparseHandle_t          handle,
                  cusparseDirection_t    dir,
                  int                  m,
                  int                  n,

```

This function converts a sparse matrix in CSR format that is defined by the three arrays `csrValA`, `csrRowPtrA`, and `csrColIndA` into a sparse matrix in BSR format that is defined by arrays `bsrValC`, `bsrRowPtrC`, and `bsrColIndC`.

**A** is an  $m \times n$  sparse matrix. The BSR format of **A** has **mb** block rows, **nb** block columns, and **nnzb** nonzero blocks, where  $\text{mb} = (\lfloor (m+blockDim-1) / blockDim \rfloor)$  and  $\text{nb} = (\lfloor (n+blockDim-1) / blockDim \rfloor)$ .

If **m** or **n** is not multiple of **blockDim**, zeros are filled in.

The conversion in cuSPARSE entails a two-step approach. First, the user allocates `bsrRowPtrC` of **mb+1** elements and uses function `cusparseXcsr2bsrNnz()` to determine the number of nonzero block columns per block row. Second, the user gathers **nnzb** (number of non-zero block columns of matrix C) from either ( $\text{nnzb} = *nnzTotalDevHostPtr$ ) or ( $\text{nnzb} = bsrRowPtrC[\text{mb}] - bsrRowPtrC[0]$ ) and allocates `bsrValC` of  $\text{nnzb} \times \text{blockDim} \times \text{blockDim}$  elements and `bsrColIndC` of **nnzb** elements. Finally function `cusparse[S|D|C|Z]csr2bsr90` is called to complete the conversion.

The general procedure is as follows:

```
// Given CSR format (csrRowPtrA, csrColIndA, csrValA) and
// blocks of BSR format are stored in column-major order.
cusparseDirection_t dir = CUSPARSE_DIRECTION_COLUMN;
int base, nnzb;
int mb = (m + blockDim-1)/blockDim;
cudaMalloc((void**)&bsrRowPtrC, sizeof(int) * (mb+1));
// nnzTotalDevHostPtr points to host memory
int *nnzTotalDevHostPtr = &nnzb;
cusparseXcsr2bsrNnz(handle, dir, m, n,
                     descrA, csrRowPtrA, csrColIndA,
                     blockDim,
                     descrC, bsrRowPtrC,
                     nnzTotalDevHostPtr);
if (NULL != nnzTotalDevHostPtr){
    nnzb = *nnzTotalDevHostPtr;
} else{
    cudaMemcpy(&nnzb, bsrRowPtrC+mb, sizeof(int), cudaMemcpyDeviceToHost);
    cudaMemcpy(&base, bsrRowPtrC, sizeof(int), cudaMemcpyDeviceToHost);
    nnzb -= base;
}
cudaMalloc((void**)&bsrColIndC, sizeof(int)*nnzb);
cudaMalloc((void**)&bsrValC, sizeof(float)*(blockDim*blockDim)*nnzb);
cusparseScsr2bsr(handle, dir, m, n,
                  descrA,
                  csrValA, csrRowPtrA, csrColIndA,
                  blockDim,
                  descrC,
                  bsrValC, bsrRowPtrC, bsrColIndC);
```

The routine `cusparse<t>csr2bsr()` has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally if **blockDim > 16**
- ▶ The routine does *not* support asynchronous execution if **blockDim == 1**
- ▶ The routine does *not* support CUDA graph capture if **blockDim == 1**

The routine `cusparseXcsr2bsrNnz()` has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution

- The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>dir</code>	storage format of blocks, either <code>CUSPARSE_DIRECTION_ROW</code> or <code>CUSPARSE_DIRECTION_COLUMN</code> .
<code>m</code>	number of rows of sparse matrix <code>A</code> .
<code>n</code>	number of columns of sparse matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> .
<code>csrValA</code>	<type> array of <code>nnz</code> ( $=\text{csrRowPtrA}[m] - \text{csrRowPtr}[0]$ ) non-zero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> column indices of the non-zero elements of matrix <code>A</code> .
<code>blockDim</code>	block dimension of sparse matrix <code>A</code> . The range of <code>blockDim</code> is between 1 and <code>min(m, n)</code> .
<code>descrC</code>	the descriptor of matrix <code>c</code> .

### Output

<code>bsrValC</code>	<type> array of <code>nnzb*blockDim*blockDim</code> nonzero elements of matrix <code>c</code> .
<code>bsrRowPtrC</code>	integer array of <code>mb+1</code> elements that contains the start of every block row and the end of the last block row plus one of matrix <code>c</code> .
<code>bsrColIndC</code>	integer array of <code>nnzb</code> column indices of the non-zero blocks of matrix <code>c</code> .
<code>nnzTotalDevHostPtr</code>	total number of nonzero elements in device or host memory. It is equal to <code>(bsrRowPtrC[mb] - bsrRowPtrC[0])</code> .

See `cusparseStatus_t` for the description of the return status

## 13.10. `cusparse<t>csr2coo()`

```
cusparseStatus_t
cusparseXcsr2coo(cusparseHandle_t handle,
                  const int* csrRowPtr,
                  int nnz,
                  int m,
                  int* cooRowInd,
                  cusparseIndexBase_t idxBase)
```

This function converts the array containing the compressed row pointers (corresponding to CSR format) into an array of uncompressed row indices (corresponding to COO format).

It can also be used to convert the array containing the compressed column indices (corresponding to CSC format) into an array of uncompressed column indices (corresponding to COO format).

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>csrRowPtr</b>	integer array of $m+1$ elements that contains the start of every row and the end of the last row plus one.
<b>nnz</b>	number of nonzeros of the sparse matrix (that is also the length of array <b>cooRowInd</b> ).
<b>m</b>	number of rows of matrix <b>A</b> .
<b>idxBase</b>	<b>CUSPARSE_INDEX_BASE_ZERO</b> or <b>CUSPARSE_INDEX_BASE_ONE</b> .

### Output

<b>cooRowInd</b>	integer array of <b>nnz</b> uncompressed row indices.
------------------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 13.11. `cusparse<t>csr2csc()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseCsr2cscEx2()` instead. *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsr2csc(cusparseHandle_t      handle,
                  int                m,
                  int                n,
                  int                nnz,
                  const float*       csrVal,
                  const int*         csrRowPtr,
                  const int*         csrColInd,
                  float*             cscVal,
                  int*               cscRowInd,
                  int*               cscColPtr,
                  cusparseAction_t   copyValues,
                  cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseDcsr2csc(cusparseHandle_t      handle,
                  int                m,
                  int                n,
                  int                nnz,
                  const double*      csrVal,
                  const int*         csrRowPtr,
                  const int*         csrColInd,
                  double*            cscVal,
                  int*               cscRowInd,
                  int*               cscColPtr,
                  cusparseAction_t   copyValues,
                  cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseCcsr2csc(cusparseHandle_t      handle,
                  int                m,
                  int                n,
                  int                nnz,
                  const cuComplex*   csrVal,
                  const int*         csrRowPtr,
                  const int*         csrColInd,
                  cuComplex*         cscVal,
                  int*               cscRowInd,
                  int*               cscColPtr,
                  cusparseAction_t   copyValues,
                  cusparseIndexBase_t idxBase)

cusparseStatus_t
cusparseZcsr2csc(cusparseHandle_t      handle,
                  int                m,
                  int                n,
                  int                nnz,
                  const cuDoubleComplex* csrVal,
                  const int*         csrRowPtr,
                  const int*         csrColInd,
                  cuDoubleComplex*   cscVal,
                  int*               cscRowInd,
                  int*               cscColPtr,
                  cusparseAction_t   copyValues,
                  cusparseIndexBase_t idxBase)

```

This function converts a sparse matrix in CSR format (that is defined by the three arrays `csrVal`, `csrRowPtr`, and `csrColInd`) into a sparse matrix in CSC format (that is defined by arrays `cscVal`, `cscRowInd`, and `cscColPtr`). The resulting matrix can also be seen as the transpose of the original sparse matrix. Notice that this routine can also be used to convert a matrix in CSC format into a matrix in CSR format.

- ▶ This function requires a significant amount of extra storage that is proportional to the matrix size
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix $A$ .
<code>n</code>	number of columns of matrix $A$ .
<code>nnz</code>	number of nonzero elements of matrix $A$ .
<code>csrVal</code>	<type> array of <code>nnz</code> (= <code>csrRowPtr(m) - csrRowPtr(0)</code> ) nonzero elements of matrix $A$ .
<code>csrRowPtr</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColInd</code>	integer array of <code>nnz</code> (= <code>csrRowPtr(m) - csrRowPtr(0)</code> ) column indices of the nonzero elements of matrix $A$ .
<code>copyValues</code>	<code>CUSPARSE_ACTION_SYMBOLIC</code> or <code>CUSPARSE_ACTION_NUMERIC</code> .
<code>idxBase</code>	<code>CUSPARSE_INDEX_BASE_ZERO</code> or <code>CUSPARSE_INDEX_BASE_ONE</code> .

### Output

<code>cscVal</code>	<type> array of <code>nnz</code> (= <code>cscColPtr(n) - cscColPtr(0)</code> ) nonzero elements of matrix $A$ . It is only filled in if <code>copyValues</code> is set to <code>CUSPARSE_ACTION_NUMERIC</code> .
<code>cscRowInd</code>	integer array of <code>nnz</code> (= <code>cscColPtr(n) - cscColPtr(0)</code> ) column indices of the nonzero elements of matrix $A$ .
<code>cscColPtr</code>	integer array of <code>n+1</code> elements that contains the start of every column and the end of the last column plus one.

See `cusparseStatus_t` for the description of the return status

## 13.12. `cusparseCsr2cscEx()` [DEPRECATED]

**[[DEPRECATED]]** use `cusparseCsr2cscEx2()` instead. *The routine will be removed in the next major release*

```
cusparseStatus_t
cusparseCsr2cscEx(cusparseHandle_t handle,
                  int m,
                  int n,
                  int nnz,
                  const void* csrSortedVal,
                  cudaDataType csrSortedValtype,
                  const int* csrSortedRowPtr,
                  const int* csrSortedColInd,
                  void* cscSortedVal,
                  cudaDataType cscSortedValtype,
                  int* cscSortedRowInd,
                  int* cscSortedColPtr,
                  cusparseAction_t copyValues,
                  cusparseIndexBase_t idxBase,
                  cudaDataType executiontype)
```

This function is an extended version of `cusparse<t>csr2csc()`. For detailed description of the functionality, see `cusparse<t>csr2csc()`.

This function does not support half-precision execution type, but it supports half-precision IO with single precision execution.

- ▶ This function requires a significant amount of extra storage that is proportional to the matrix size
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

See `cusparseStatus_t` for the description of the return status

## 13.13. cusparseCsr2cscEx2()

```
cusparseStatus_t
cusparseCsr2cscEx2_bufferSize(cusparseHandle_t      handle,
                               int               m,
                               int               n,
                               int               nnz,
                               const void*       csrVal,
                               const int*        csrRowPtr,
                               const int*        csrColInd,
                               void*             cscVal,
                               int*              cscColPtr,
                               int*              cscRowInd,
                               cudaDataType      valType,
                               cusparseAction_t  copyValues,
                               cusparseIndexBase_t idxBase,
                               cusparseCsr2CscAlg_t alg,
                               size_t*           bufferSize)
```

```
cusparseStatus_t
cusparseCsr2cscEx2(cusparseHandle_t      handle,
                    int               m,
                    int               n,
                    int               nnz,
                    const void*       csrVal,
                    const int*        csrRowPtr,
                    const int*        csrColInd,
                    void*             cscVal,
                    int*              cscColPtr,
                    int*              cscRowInd,
                    cudaDataType      valType,
                    cusparseAction_t  copyValues,
                    cusparseIndexBase_t idxBase,
                    cusparseCsr2CscAlg_t alg,
                    void*             buffer)
```

This function converts a sparse matrix in CSR format (that is defined by the three arrays **csrVal**, **csrRowPtr**, and **csrColInd**) into a sparse matrix in CSC format (that is defined by arrays **cscVal**, **cscRowInd**, and **cscColPtr**). The resulting matrix can also be seen as the transpose of the original sparse matrix. Notice that this routine can also be used to convert a matrix in CSC format into a matrix in CSR format.

For alg **CUSPARSE\_CSR2CSC\_ALG1**: it requires extra storage proportional to the number of nonzero values **nnz**. It provides in output always the same matrix.

For alg **CUSPARSE\_CSR2CSC\_ALG2**: it requires extra storage proportional to the number of rows **m**. It does not ensure always the same ordering of CSC column indices and values. Also, it provides better performance than **CUSPARSE\_CSR2CSC\_ALG1** for regular matrices.

It is executed asynchronously with respect to the host, and it may return control to the application on the host before the result is ready.

The function `cusparseCsr2cscEx2_bufferSize()` returns the size of the workspace needed by `cusparseCsr2cscEx2()`. User needs to allocate a buffer of this size and give that buffer to `cusparseCsr2cscEx2()` as an argument.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context
<code>m</code>	number of rows of the CSR input matrix; number of columns of the CSC output matrix
<code>n</code>	number of columns of the CSR input matrix; number of rows of the CSC output matrix
<code>nnz</code>	number of nonzero elements of the CSR and CSC matrices
<code>csrVal</code>	value array of size <code>nnz</code> of the CSR matrix; of same type as <code>valType</code>
<code>csrRowPtr</code>	integer array of size <code>m + 1</code> that contains the CSR row offsets
<code>csrColInd</code>	integer array of size <code>nnz</code> that contains the CSR column indices
<code>valType</code>	value type for both CSR and CSC matrices
<code>copyValues</code>	<code>CUSPARSE_ACTION_SYMBOLIC</code> or <code>CUSPARSE_ACTION_NUMERIC</code>
<code>idxBase</code>	Index base <code>CUSPARSE_INDEX_BASE_ZERO</code> or <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>alg</code>	algorithm implementation. see <code>cusparseCsr2CscAlg_t</code> for possible values.
<code>bufferSize</code>	number of bytes of workspace needed by <code>cusparseCsr2cscEx2()</code>
<code>buffer</code>	pointer to workspace buffer

See `cusparseStatus_t` for the description of the return status

## 13.14. `cusparse<t>csr2dense()`

```

cusparseStatus_t
cusparseScsr2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const float*          csrValA,
                    const int*             csrRowPtrA,
                    const int*             csrColIndA,
                    float*                A,
                    int                  lda)

cusparseStatus_t
cusparseDcsr2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const double*          csrValA,
                    const int*              csrRowPtrA,
                    const int*              csrColIndA,
                    double*                A,
                    int                  lda)

cusparseStatus_t
cusparseCcsr2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*        csrValA,
                    const int*              csrRowPtrA,
                    const int*              csrColIndA,
                    cuComplex*              A,
                    int                  lda)

cusparseStatus_t
cusparseZcsr2dense(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* csrValA,
                    const int*              csrRowPtrA,
                    const int*              csrColIndA,
                    cuDoubleComplex*        A,
                    int                  lda)

```

This function converts the sparse matrix in CSR format (that is defined by the three arrays **csrValA**, **csrRowPtrA**, and **csrColIndA**) into the matrix **A** in dense format. The dense matrix **A** is filled in with the values of the sparse matrix and with zeros elsewhere.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix $A$ .
<b>n</b>	number of columns of matrix $A$ .
<b>descrA</b>	the descriptor of matrix $A$ . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<b>csrValA</b>	<type> array of $\text{nnz}$ (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix $A$ .
<b>csrRowPtrA</b>	integer array of $m+1$ elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	integer array of $\text{nnz}$ (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix $A$ .
<b>lda</b>	leading dimension of array <code>matrixA</code> .

## Output

<b>A</b>	array of dimensions <code>(lda, n)</code> that is filled in with the values of the sparse matrix.
----------	---

See `cusparseStatus_t` for the description of the return status

## 13.15. `cusparse<t>csr2csr_compress()`

```

cusparseStatus_t
cusparseScsr2csr_compress(cusparseHandle_t          handle,
                           int                  m,
                           int                  n,
                           const cusparseMatDescr_t descrA,
                           const float*          csrValA,
                           const int*             csrColIndA,
                           const int*             csrRowPtrA,
                           int                  nnzA,
                           const int*             nnzPerRow,
                           float*                csrValC,
                           int*                 csrColIndC,
                           int*                 csrRowPtrC,
                           float                tol)

cusparseStatus_t
cusparseDcsr2csr_compress(cusparseHandle_t          handle,
                           int                  m,
                           int                  n,
                           const cusparseMatDescr_t descrA,
                           const double*         csrValA,
                           const int*            csrColIndA,
                           const int*            csrRowPtrA,
                           int                  nnzA,
                           const int*            nnzPerRow,
                           double*               csrValC,
                           int*                 csrColIndC,
                           int*                 csrRowPtrC,
                           double                tol)

cusparseStatus_t
cusparseCcsr2csr_compress(cusparseHandle_t          handle,
                           int                  m,
                           int                  n,
                           const cusparseMatDescr_t descrA,
                           const cuComplex*      csrValA,
                           const int*            csrColIndA,
                           const int*            csrRowPtrA,
                           int                  nnzA,
                           const int*            nnzPerRow,
                           cuComplex*            csrValC,
                           int*                 csrColIndC,
                           int*                 csrRowPtrC,
                           cuComplex             tol)

cusparseStatus_t
cusparseZcsr2csr_compress(cusparseHandle_t          handle,
                           int                  m,
                           int                  n,
                           const cusparseMatDescr_t descrA,
                           const cuDoubleComplex* csrValA,
                           const int*            csrColIndA,
                           const int*            csrRowPtrA,
                           int                  nnzA,
                           const int*            nnzPerRow,
                           cuDoubleComplex*      csrValC,
                           int*                 csrColIndC,
                           int*                 csrRowPtrC,
                           cuDoubleComplex        tol)

```

This function compresses the sparse matrix in CSR format into compressed CSR format. Given a sparse matrix A and a non-negative value threshold(in the case of complex values, only the magnitude of the real part is used in the check), the function returns a sparse matrix C, defined by

$$C(i,j) = A(i,j) \text{ if } |A(i,j)| > \text{threshold}$$

The implementation adopts a two-step approach to do the conversion. First, the user allocates `csrRowPtrC` of  $m+1$  elements and uses function `cusparse<t>nzz_compress()` to determine `nzzPerRow`(the number of nonzeros columns per row) and `nzzC`(the total number of nonzeros). Second, the user allocates `csrValC` of `nzzC` elements and `csrColIndC` of `nzzC` integers. Finally function `cusparse<t>csr2csr_compress()` is called to complete the conversion.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix A.
<code>n</code>	number of columns of matrix A.
<code>descrA</code>	the descriptor of matrix A. The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA</code>	<type> array of <code>nzz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) elements of matrix A.
<code>csrColIndA</code>	integer array of <code>nzz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the elements of matrix A.
<code>csrRowPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>nzzA</code>	number of nonzero elements in matrix A.
<code>nzzPerRow</code>	this array contains the number of elements kept in the compressed matrix, by row.
<code>tol</code>	on input, this contains the non-negative tolerance value used for compression. Any values in matrix A less than or equal to this value will be dropped during compression.

### Output

<code>csrValC</code>	on output, this array contains the typed values of elements kept in the compressed matrix. Size = <code>nzzC</code> .
----------------------	---

<b>csrColIndC</b>	on output, this integer array contains the column indices of elements kept in the compressed matrix. Size = nnzC.
<b>csrRowPtrC</b>	on output, this integer array contains the row pointers for elements kept in the compressed matrix. Size = m+1

See [cusparseStatus\\_t](#) for the description of the return status

The following is a sample code to show how to use this API.

```
#include <stdio.h>
#include <sys/time.h>
#include <cusparse.h>

#define ERR_NE(X,Y) do { if ((X) != (Y)) { \
                           fprintf(stderr,"Error in %s at %s:%d \
\n", __func__, __FILE__, __LINE__); \
                           exit(-1); } } while(0)
#define CUDA_CALL(X) ERR_NE((X),cudaSuccess)
#define CUSPARSE_CALL(X) ERR_NE((X),CUSPARSE_STATUS_SUCCESS)
int main(){
    int m = 6, n = 5;
    cusparseHandle_t handle;
    CUSPARSE_CALL(cusparseCreate(&handle));
    cusparseMatDescr_t descrX;
    CUSPARSE_CALL(cusparseCreateMatDescr(&descrX));
    // Initialize sparse matrix
    float *X;
    CUDA_CALL(cudaMallocManaged( &X, sizeof(float) * m * n ));
    memset( X, 0, sizeof(float) * m * n );
    X[0 + 0*m] = 1.0; X[0 + 1*m] = 3.0;
    X[1 + 1*m] = -4.0; X[1 + 2*m] = 5.0;
    X[2 + 0*m] = 2.0; X[2 + 3*m] = 7.0; X[2 + 4*m] = 8.0;
    X[3 + 2*m] = 6.0; X[3 + 4*m] = 9.0;
    X[4 + 3*m] = 3.5; X[4 + 4*m] = 5.5;
    X[5 + 0*m] = 6.5; X[5 + 2*m] = -9.9;
    // Initialize total_nnz, and nnzPerRowX for cusparseSdense2csr()
    int total_nnz = 13;
    int *nnzPerRowX;
    CUDA_CALL( cudaMallocManaged( &nnzPerRowX, sizeof(int) * m ) );
    nnzPerRowX[0] = 2; nnzPerRowX[1] = 2; nnzPerRowX[2] = 3;
    nnzPerRowX[3] = 2; nnzPerRowX[4] = 2; nnzPerRowX[5] = 2;

    float *csrValX;
    int *csrRowPtrX;
    int *csrColIndX;
    CUDA_CALL( cudaMallocManaged( &csrValX, sizeof(float) * total_nnz ) );
    CUDA_CALL( cudaMallocManaged( &csrRowPtrX, sizeof(int) * (m+1) ) );
    CUDA_CALL( cudaMallocManaged( &csrColIndX, sizeof(int) * total_nnz ) );
```

Before calling this API, call two APIs to prepare the input.

```
/** Call cusparseSdense2csr to generate CSR format as the inputs for
cusparseScsr2csr_compress */
CUSPARSE_CALL( cusparseSdense2csr( handle, m, n, descrX, X,
                                    m, nnzPerRowX, csrValX,
                                    csrRowPtrX, csrColIndX ) );
float tol = 3.5;
int *nnzPerRowY;
int *testNNZTotal;
CUDA_CALL( cudaMallocManaged( &nnzPerRowY, sizeof(int) * m ) );
CUDA_CALL( cudaMallocManaged( &testNNZTotal, sizeof(int) ) );
memset( nnzPerRowY, 0, sizeof(int) * m );
// cusparseSnnz_compress generates nnzPerRowY and testNNZTotal
CUSPARSE_CALL( cusparseSnnz_compress( handle, m, descrX, csrValX,
                                      csrRowPtrX, nnzPerRowY,
                                      testNNZTotal, tol ) );

float *csrValY;
int *csrRowPtrY;
int *csrColIndY;
CUDA_CALL( cudaMallocManaged( &csrValY, sizeof(float) * (*testNNZTotal) ) );
CUDA_CALL( cudaMallocManaged( &csrRowPtrY, sizeof(int) * (m+1) ) );
CUDA_CALL( cudaMallocManaged( &csrColIndY, sizeof(int) * (*testNNZTotal) ) );

CUSPARSE_CALL( cusparseScsr2csr_compress( handle, m, n, descrX, csrValX,
                                         csrColIndX, csrRowPtrX,
                                         total_nnz, nnzPerRowY,
                                         csrValY, csrColIndY,
                                         csrRowPtrY, tol ) );

/* Expect results
nnzPerRowY: 0 2 2 2 1 2
csrValY: -4 5 7 8 6 9 5.5 6.5 -9.9
csrColIndY: 1 2 3 4 2 4 4 0 2
csrRowPtrY: 0 0 2 4 6 7 9
*/
cudaFree(X);
cusparseDestroy(handle);
cudaFree(nnzPerRowX);
cudaFree(csrValX);
cudaFree(csrRowPtrX);
cudaFree(csrColIndX);
cudaFree(csrValY);
cudaFree(nnzPerRowY);
cudaFree(testNNZTotal);
cudaFree(csrRowPtrY);
cudaFree(csrColIndY);
return 0;
}
```

## 13.16. `cusparse<t>csr2hyb()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseScsr2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const float*        csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseDcsr2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const double*       csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseCcsr2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const cuComplex*    csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseZcsr2hyb(cusparseHandle_t          handle,
                  int                 m,
                  int                 n,
                  const cusparseMatDescr_t descrA,
                  const cuDoubleComplex* csrValA,
                  const int*          csrRowPtrA,
                  const int*          csrColIndA,
                  cusparseHybMat_t    hybA,
                  int                 userEllWidth,
                  cusparseHybPartition_t partitionType)

```

This function converts a sparse matrix in CSR format into a sparse matrix in HYB format. It assumes that the **hybA** parameter has been initialized with `cusparseCreateHybMat()` routine before calling this function.

This function requires some amount of temporary storage and a significant amount of storage for the matrix in HYB format. It is executed asynchronously with respect to the host and may return control to the application on the host before the result is ready.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix <code>A</code> .
<code>n</code>	number of columns of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrValA</code>	<type> array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>csrRowPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndA</code>	integer array of <code>nnz</code> (= <code>csrRowPtrA(m) - csrRowPtrA(0)</code> ) column indices of the nonzero elements of matrix <code>A</code> .
<code>userEllWidth</code>	width of the regular (ELL) part of the matrix in HYB format, which should be less than maximum number of nonzeros per row and is only required if <code>partitionType == CUSPARSE_HYB_PARTITION_USER</code> .
<code>partitionType</code>	partitioning method to be used in the conversion (please refer to <code>cusparseHybPartition_t</code> for details).

## Output

<code>hybA</code>	the matrix <code>A</code> in HYB storage format.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 13.17. `cusparse<t>dense2csc()`

```

cusparseStatus_t
cusparseSdense2csc(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const float*          A,
                    int                  lda,
                    const int*            nnzPerCol,
                    float*               cscValA,
                    int*                 cscRowIndA,
                    int*                 cscColPtrA)

cusparseStatus_t
cusparseDdense2csc(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const double*         A,
                    int                  lda,
                    const int*            nnzPerCol,
                    double*              cscValA,
                    int*                 cscRowIndA,
                    int*                 cscColPtrA)

cusparseStatus_t
cusparseCdense2csc(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*      A,
                    int                  lda,
                    const int*            nnzPerCol,
                    cuComplex*            cscValA,
                    int*                 cscRowIndA,
                    int*                 cscColPtrA)

cusparseStatus_t
cusparseZdense2csc(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* A,
                    int                  lda,
                    const int*            nnzPerCol,
                    cuDoubleComplex*      cscValA,
                    int*                 cscRowIndA,
                    int*                 cscColPtrA)

```

This function converts the matrix **A** in dense format into a sparse matrix in CSC format. All the parameters are assumed to have been pre-allocated by the user, and the arrays are filled in based on **nnzPerCol**, which can be precomputed with `cusparse<t>nnz()`.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution

- The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix <code>A</code> .
<code>n</code>	number of columns of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>A</code>	array of dimensions <code>(lda, n)</code> .
<code>lda</code>	leading dimension of dense array <code>A</code> .
<code>nnzPerCol</code>	array of size <code>n</code> containing the number of nonzero elements per column.

### Output

<code>cscValA</code>	<type> array of <code>nnz</code> ( $= \text{cscRowPtrA}(m) - \text{cscRowPtrA}(0)$ ) nonzero elements of matrix <code>A</code> . It is only filled in if <code>copyValues</code> is set to <code>CUSPARSE_ACTION_NUMERIC</code> .
<code>cscRowIndA</code>	integer array of <code>nnz</code> ( $= \text{cscRowPtrA}(m) - \text{cscRowPtrA}(0)$ ) row indices of the nonzero elements of matrix <code>A</code> .
<code>cscColPtrA</code>	integer array of <code>n+1</code> elements that contains the start of every column and the end of the last column plus one.

See `cusparseStatus_t` for the description of the return status

## 13.18. `cusparse<t>dense2csr()`

```

cusparseStatus_t
cusparseSdense2csr(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const float*          A,
                    int                  lda,
                    const int*            nnzPerRow,
                    float*               csrValA,
                    int*                csrRowPtrA,
                    int*                csrColIndA)

cusparseStatus_t
cusparseDdense2csr(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const double*         A,
                    int                  lda,
                    const int*            nnzPerRow,
                    double*              csrValA,
                    int*                csrRowPtrA,
                    int*                csrColIndA)

cusparseStatus_t
cusparseCdense2csr(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*      A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cuComplex*            csrValA,
                    int*                csrRowPtrA,
                    int*                csrColIndA)

cusparseStatus_t
cusparseZdense2csr(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cuDoubleComplex*      csrValA,
                    int*                csrRowPtrA,
                    int*                csrColIndA)

```

This function converts the matrix **A** in dense format into a sparse matrix in CSR format. All the parameters are assumed to have been pre-allocated by the user and the arrays are filled in based on **nnzPerRow**, which can be pre-computed with `cusparse<t>nnz()`.

This function requires no extra storage. It is executed asynchronously with respect to the host and may return control to the application on the host before the result is ready.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>A</b>	array of dimensions ( <b>lda</b> , <b>n</b> ).
<b>lda</b>	leading dimension of dense array <b>A</b> .
<b>nnzPerRow</b>	array of size <b>n</b> containing the number of non-zero elements per row.

### Output

<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> – <b>csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m+1</b> elements that contains the start of every column and the end of the last column plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m)</b> – <b>csrRowPtrA(0)</b> ) column indices of the non-zero elements of matrix <b>A</b> .

See `cusparseStatus_t` for the description of the return status

## 13.19. `cusparse<t>dense2hyb()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseSdense2hyb(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const float*          A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cusparseHybMat_t      hybA,
                    int                  userEllWidth,
                    cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseDdense2hyb(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const double*         A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cusparseHybMat_t      hybA,
                    int                  userEllWidth,
                    cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseCdense2hyb(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuComplex*      A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cusparseHybMat_t      hybA,
                    int                  userEllWidth,
                    cusparseHybPartition_t partitionType)

cusparseStatus_t
cusparseZdense2hyb(cusparseHandle_t          handle,
                    int                  m,
                    int                  n,
                    const cusparseMatDescr_t descrA,
                    const cuDoubleComplex* A,
                    int                  lda,
                    const int*            nnzPerRow,
                    cusparseHybMat_t      hybA,
                    int                  userEllWidth,
                    cusparseHybPartition_t partitionType)

```

This function converts matrix **A** in dense format into a sparse matrix in HYB format. It assumes that the routine **`cusparseCreateHybMat()`** was used to initialize the opaque structure **hybA** and that the array **nnzPerRow** was pre-computed with **`cusparse<t>nnz()`**.

This function requires some amount of temporary storage and a significant amount of storage for the matrix in HYB format. It is executed asynchronously with respect to the host and may return control to the application on the host before the result is ready.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
<code>m</code>	number of rows of matrix <code>A</code> .
<code>n</code>	number of columns of matrix <code>A</code> .
<code>descrA</code>	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> .
<code>A</code>	array of dimensions <code>(lda, n)</code> .
<code>lda</code>	leading dimension of dense array <code>A</code> .
<code>nnzPerRow</code>	array of size <code>m</code> containing the number of nonzero elements per row.
<code>userEllWidth</code>	width of the regular (ELL) part of the matrix in HYB format, which should be less than maximum number of nonzeros per row and is only required if <code>partitionType == CUSPARSE_HYB_PARTITION_USER</code> .
<code>partitionType</code>	partitioning method to be used in the conversion (please refer to <code>cusparseHybPartition_t</code> for details).

### Output

<code>hybA</code>	the matrix <code>A</code> in HYB storage format.
-------------------	--

See `cusparseStatus_t` for the description of the return status

## 13.20. `cusparse<t>hyb2csc()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseShyb2csc(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t  hybA,
                  float*                 cscValA,
                  int*                   cscRowIndA,
                  int*                   cscColPtrA)

cusparseStatus_t
cusparseDhyb2csc(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t  hybA,
                  double*                cscValA,
                  int*                   cscRowIndA,
                  int*                   cscColPtrA)

cusparseStatus_t
cusparseChyb2csc(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t  hybA,
                  cuComplex*              cscValA,
                  int*                   cscRowIndA,
                  int*                   cscColPtrA)

cusparseStatus_t
cusparseZhyb2csc(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t  hybA,
                  cuDoubleComplex*        cscValA,
                  int*                   cscRowIndA,
                  int*                   cscColPtrA)

```

This function converts a sparse matrix in HYB format into a sparse matrix in CSC format.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>descrA</b>	the descriptor of matrix <b>A</b> in Hyb format. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.

### Output

<code>cscValA</code>	<type> array of <code>nnz</code> (= <code>cscColPtrA(m) - cscColPtrA(0)</code> ) nonzero elements of matrix <code>A</code> .
<code>cscRowIndA</code>	integer array of <code>nnz</code> (= <code>cscColPtrA(m) - cscColPtrA(0)</code> ) column indices of the non-zero elements of matrix <code>A</code> .
<code>cscColPtrA</code>	integer array of <code>m+1</code> elements that contains the start of every column and the end of the last row plus one.

See `cusparseStatus_t` for the description of the return status

## 13.21. `cusparse<t>hyb2csr()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseShyb2csr(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t   hybA,
                  float*                  csrValA,
                  int*                   csrRowPtrA,
                  int*                   csrColIndA)

cusparseStatus_t
cusparseDhyb2csr(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t   hybA,
                  double*                 csrValA,
                  int*                   csrRowPtrA,
                  int*                   csrColIndA)

cusparseStatus_t
cusparseChyb2csr(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t   hybA,
                  cuComplex*              csrValA,
                  int*                   csrRowPtrA,
                  int*                   csrColIndA)

cusparseStatus_t
cusparseZhyb2csr(cusparseHandle_t          handle,
                  const cusparseMatDescr_t descrA,
                  const cusparseHybMat_t   hybA,
                  cuDoubleComplex*         csrValA,
                  int*                   csrRowPtrA,
                  int*                   csrColIndA)

```

This function converts a sparse matrix in HYB format into a sparse matrix in CSR format.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>descrA</b>	the descriptor of matrix <b>A</b> in Hyb format. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.

## Output

<b>csrValA</b>	<type> array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) nonzero elements of matrix <b>A</b> .
<b>csrRowPtrA</b>	integer array of <b>m+1</b> elements that contains the start of every column and the end of the last row plus one.
<b>csrColIndA</b>	integer array of <b>nnz</b> (= <b>csrRowPtrA(m) - csrRowPtrA(0)</b> ) column indices of the nonzero elements of matrix <b>A</b> .

See [cusparseStatus\\_t](#) for the description of the return status

## 13.22. `cusparse<t>hyb2dense()` [DEPRECATED]

**[[DEPRECATED]]** *The routine will be removed in the next major release*

```

cusparseStatus_t
cusparseShyb2dense(cusparseHandle_t      handle,
                   const cusparseMatDescr_t descrA,
                   const cusparseHybMat_t   hybA,
                   float*                  A,
                   int                      lda)

cusparseStatus_t
cusparseDhyb2dense(cusparseHandle_t      handle,
                   const cusparseMatDescr_t descrA,
                   const cusparseHybMat_t   hybA,
                   double*                 A,
                   int                      lda)

cusparseStatus_t
cusparseChyb2dense(cusparseHandle_t      handle,
                   const cusparseMatDescr_t descrA,
                   const cusparseHybMat_t   hybA,
                   cuComplex*              A,
                   int                      lda)

cusparseStatus_t
cusparseZhyb2dense(cusparseHandle_t      handle,
                   const cusparseMatDescr_t descrA,
                   const cusparseHybMat_t   hybA,
                   cuDoubleComplex*         A,
                   int                      lda)

```

This function converts a sparse matrix in HYB format (contained in the opaque structure ) into matrix **A** in dense format. The dense matrix **A** is filled in with the values of the sparse matrix and with zeros elsewhere.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>descrA</b>	the descriptor of matrix <b>A</b> in Hyb format. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> .
<b>hybA</b>	the matrix <b>A</b> in HYB storage format.
<b>lda</b>	leading dimension of dense array <b>A</b> .

### Output

<b>A</b>	array of dimensions ( <b>lda</b> , <b>n</b> ) that is filled in with the values of the sparse matrix.
----------	---

See [cusparseStatus\\_t](#) for the description of the return status

## 13.23. `cusparse<t>nnz()`

```

cusparseStatus_t
cusparseSnnz(cusparseHandle_t      handle,
            cusparseDirection_t   dirA,
            int                  m,
            int                  n,
            const cusparseMatDescr_t descrA,
            const float*          A,
            int                  lda,
            int*                 nnzPerRowColumn,
            int*                 nnzTotalDevHostPtr)

cusparseStatus_t
cusparseDnnz(cusparseHandle_t      handle,
            cusparseDirection_t   dirA,
            int                  m,
            int                  n,
            const cusparseMatDescr_t descrA,
            const double*         A,
            int                  lda,
            int*                 nnzPerRowColumn,
            int*                 nnzTotalDevHostPtr)

cusparseStatus_t
cusparseCnnz(cusparseHandle_t      handle,
            cusparseDirection_t   dirA,
            int                  m,
            int                  n,
            const cusparseMatDescr_t descrA,
            const cuComplex*     A,
            int                  lda,
            int*                 nnzPerRowColumn,
            int*                 nnzTotalDevHostPtr)

cusparseStatus_t
cusparseZnnz(cusparseHandle_t      handle,
            cusparseDirection_t   dirA,
            int                  m,
            int                  n,
            const cusparseMatDescr_t descrA,
            const cuDoubleComplex* A,
            int                  lda,
            int*                 nnzPerRowColumn,
            int*                 nnzTotalDevHostPtr)

```

This function computes the number of nonzero elements per row or column and the total number of nonzero elements in a dense matrix.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

<code>handle</code>	handle to the cuSPARSE library context.
---------------------	---

<b>dirA</b>	direction that specifies whether to count nonzero elements by <b>CUSPARSE_DIRECTION_ROW</b> or by <b>CUSPARSE_DIRECTION_COLUMN</b> .
<b>m</b>	number of rows of matrix <b>A</b> .
<b>n</b>	number of columns of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>A</b>	array of dimensions <b>(lda, n)</b> .
<b>lda</b>	leading dimension of dense array <b>A</b> .

## Output

<b>nnzPerRowColumn</b>	array of size <b>m</b> or <b>n</b> containing the number of nonzero elements per row or column, respectively.
<b>nnzTotalDevHostPtr</b>	total number of nonzero elements in device or host memory.

See [cusparseStatus\\_t](#) for the description of the return status

## 13.24. `cusparseCreateIdentityPermutation()`

```
cusparseStatus_t
cusparseCreateIdentityPermutation(cusparseHandle_t handle,
                                  int                n,
                                  int*               p);
```

This function creates an identity map. The output parameter **p** represents such map by **p = 0:1:(n-1)**.

This function is typically used with **coosort**, **csrssort**, **cscsort**, **csr2csc\_indexOnly**.

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

parameter	device or host	description
<b>handle</b>	host	handle to the cuSPARSE library context.
<b>n</b>	host	size of the map.

## Output

parameter	device or host	description
<b>p</b>	device	integer array of dimensions <b>n</b> .

See `cusparseStatus_t` for the description of the return status

## 13.25. `cusparseXcoosort()`

```

cusparseStatus_t
cusparseXcoosort_bufferSizeExt(cusparseHandle_t handle,
                                int             m,
                                int             n,
                                int             nnz,
                                const int*     cooRows,
                                const int*     cooCols,
                                size_t*        pBufferSizeInBytes)

cusparseStatus_t
cusparseXcoosortByRow(cusparseHandle_t handle,
                      int             m,
                      int             n,
                      int             nnz,
                      int*            cooRows,
                      int*            cooCols,
                      int*            P,
                      void*           pBuffer)

cusparseStatus_t
cusparseXcoosortByColumn(cusparseHandle_t handle,
                        int             m,
                        int             n,
                        int             nnz,
                        int*            cooRows,
                        int*            cooCols,
                        int*            P,
                        void*           pBuffer);

```

This function sorts COO format. The sorting is in-place. Also the user can sort by row or sort by column.

**A** is an  $m \times n$  sparse matrix that is defined in COO storage format by the three arrays `cooVals`, `cooRows`, and `cooCols`.

There is no assumption for the base index of the matrix. `coosort` uses stable sort on signed integer, so the value of `cooRows` or `cooCols` can be negative.

This function `coosort()` requires buffer size returned by `coosort_bufferSizeExt()`. The address of `pBuffer` must be multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

The parameter `P` is both input and output. If the user wants to compute sorted `cooval`, `P` must be set as  $0:1:(nnz-1)$  before `coosort()`, and after `coosort()`, new sorted value array satisfies `cooval_sorted = cooval(P)`.

Remark: the dimension `m` and `n` are not used. If the user does not know the value of `m` or `n`, just passes a value positive. This usually happens if the user only reads a COO array first and needs to decide the dimension `m` or `n` later.

Appendix D provides a simple example of `coosort()`.

- ▶ The routine requires no extra storage if `pBuffer != NULL`
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

### Input

parameter	device or host	description
<code>handle</code>	host	handle to the cuSPARSE library context.
<code>m</code>	host	number of rows of matrix <code>A</code> .
<code>n</code>	host	number of columns of matrix <code>A</code> .
<code>nnz</code>	host	number of nonzero elements of matrix <code>A</code> .
<code>cooRows</code>	device	integer array of <code>nnz</code> unsorted row indices of <code>A</code> .
<code>cooCols</code>	device	integer array of <code>nnz</code> unsorted column indices of <code>A</code> .
<code>P</code>	device	integer array of <code>nnz</code> unsorted map indices. To construct <code>cooVal</code> , the user has to set <code>P=0:1:(nnz-1)</code> .
<code>pBuffer</code>	device	buffer allocated by the user; the size is returned by <code>coosort_bufferSizeExt()</code> .

### Output

parameter	device or host	description
<code>cooRows</code>	device	integer array of <code>nnz</code> sorted row indices of <code>A</code> .
<code>cooCols</code>	device	integer array of <code>nnz</code> sorted column indices of <code>A</code> .
<code>P</code>	device	integer array of <code>nnz</code> sorted map indices.
<code>pBufferSizeInBytes</code>	host	number of bytes of the buffer.

See `cusparseStatus_t` for the description of the return status

## 13.26. `cusparseXcsrsort()`

```
cusparseStatus_t
cusparseXcsrsort_bufferSizeExt(cusparseHandle_t handle,
                                int             m,
                                int             n,
                                int             nnz,
                                const int*      csrRowPtr,
                                const int*      csrColInd,
                                size_t*         pBufferSizeInBytes)

cusparseStatus_t
cusparseXcsrsort(cusparseHandle_t          handle,
                  int             m,
                  int             n,
                  int             nnz,
                  const cusparseMatDescr_t descrA,
                  const int*      csrRowPtr,
                  int*            csrColInd,
                  int*            P,
                  void*           pBuffer)
```

This function sorts CSR format. The stable sorting is in-place.

The matrix type is regarded as **CUSPARSE\_MATRIX\_TYPE\_GENERAL** implicitly. In other words, any symmetric property is ignored.

This function **`csrsort()`** requires buffer size returned by **`csrsort_bufferSizeExt()`**. The address of **pBuffer** must be multiple of 128 bytes. If not, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

The parameter **P** is both input and output. If the user wants to compute sorted **csrVal**, **P** must be set as 0:1:(nnz-1) before **`csrsort()`**, and after **`csrsort()`**, new sorted value array satisfies **csrVal\_sorted = csrVal(P)**.

The general procedure is as follows:

```
// A is a 3x3 sparse matrix, base-0
//   | 1 2 3 |
// A = | 4 5 6 |
//   | 7 8 9 |
const int m = 3;
const int n = 3;
const int nnz = 9;
csrRowPtr[m+1] = { 0, 3, 6, 9}; // on device
csrColInd[nnz] = { 2, 1, 0, 0, 2, 1, 1, 2, 0}; // on device
csrVal[nnz] = { 3, 2, 1, 4, 6, 5, 8, 9, 7}; // on device
size_t pBufferSizeInBytes = 0;
void *pBuffer = NULL;
int *P = NULL;

// step 1: allocate buffer
cusparseXcsort_bufferSizeExt(handle, m, n, nnz, csrRowPtr, csrColInd,
    &pBufferSizeInBytes);
cudaMalloc( &pBuffer, sizeof(char)* pBufferSizeInBytes);

// step 2: setup permutation vector P to identity
cudaMalloc( (void**)&P, sizeof(int)*nnz);
cusparseCreateIdentityPermutation(handle, nnz, P);

// step 3: sort CSR format
cusparseXcsort(handle, m, n, nnz, descrA, csrRowPtr, csrColInd, P, pBuffer);

// step 4: gather sorted csrVal
cusparseDgthr(handle, nnz, csrVal, csrVal_sorted, P, CUSPARSE_INDEX_BASE_ZERO);
```

- ▶ The routine requires no extra storage if **pBuffer != NULL**
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

parameter	device or host	description
<b>handle</b>	host	handle to the cuSPARSE library context.
<b>m</b>	host	number of rows of matrix <b>A</b> .
<b>n</b>	host	number of columns of matrix <b>A</b> .
<b>nnz</b>	host	number of nonzero elements of matrix <b>A</b> .
<b>csrRowsPtr</b>	device	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColInd</b>	device	integer array of <b>nnz</b> unsorted column indices of <b>A</b> .
<b>P</b>	device	integer array of <b>nnz</b> unsorted map indices. To construct <b>csrVal</b> , the user has to set <b>P=0:1:(nnz-1)</b> .
<b>pBuffer</b>	device	buffer allocated by the user; the size is returned by <b>csrsort_bufferSizeExt()</b> .

## Output

parameter	device or host	description
<b>csrColInd</b>	device	integer array of <b>nnz</b> sorted column indices of <b>A</b> .

<b>P</b>	device	integer array of <b>nnz</b> sorted map indices.
<b>pBufferSizeInBytes</b>	host	number of bytes of the buffer.

See `cusparseStatus_t` for the description of the return status

## 13.27. `cusparseXcscsort()`

```
cusparseStatus_t
cusparseXcscsort_bufferSizeExt(cusparseHandle_t handle,
                                int m,
                                int n,
                                int nnz,
                                const int* cscColPtr,
                                const int* cscRowInd,
                                size_t* pBufferSizeInBytes)
```

```
cusparseStatus_t
cusparseXcscsort(cusparseHandle_t handle,
                  int m,
                  int n,
                  int nnz,
                  const cusparseMatDescr_t descrA,
                  const int* cscColPtr,
                  int* cscRowInd,
                  int* P,
                  void* pBuffer)
```

This function sorts CSC format. The stable sorting is in-place.

The matrix type is regarded as `CUSPARSE_MATRIX_TYPE_GENERAL` implicitly. In other words, any symmetric property is ignored.

This function `cscsort()` requires buffer size returned by `cscsort_bufferSizeExt()`. The address of `pBuffer` must be multiple of 128 bytes. If not, `CUSPARSE_STATUS_INVALID_VALUE` is returned.

The parameter `P` is both input and output. If the user wants to compute sorted `cscVal`, `P` must be set as `0:1:(nnz-1)` before `cscsort()`, and after `cscsort()`, new sorted value array satisfies `cscVal_sorted = cscVal(P)`.

The general procedure is as follows:

```
// A is a 3x3 sparse matrix, base-0
//   | 1 2 |
// A = | 4 0 |
//   | 0 8 |
const int m = 3;
const int n = 2;
const int nnz = 4;
cscColPtr[n+1] = { 0, 2, 4}; // on device
cscRowInd[nnz] = { 1, 0, 2, 0}; // on device
cscVal[nnz] = { 4.0, 1.0, 8.0, 2.0 }; // on device
size_t pBufferSizeInBytes = 0;
void *pBuffer = NULL;
int *P = NULL;

// step 1: allocate buffer
cusparseXcscsort_bufferSizeExt(handle, m, n, nnz, cscColPtr, cscRowInd,
    &pBufferSizeInBytes);
cudaMalloc( &pBuffer, sizeof(char)* pBufferSizeInBytes);

// step 2: setup permutation vector P to identity
cudaMalloc( (void**)&P, sizeof(int)*nnz);
cusparseCreateIdentityPermutation(handle, nnz, P);

// step 3: sort CSC format
cusparseXcscsort(handle, m, n, nnz, descrA, cscColPtr, cscRowInd, P, pBuffer);

// step 4: gather sorted cscVal
cusparseDgthr(handle, nnz, cscVal, cscVal_sorted, P, CUSPARSE_INDEX_BASE_ZERO);
```

- ▶ The routine requires no extra storage if **pBuffer != NULL**
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

## Input

parameter	device or host	description
<b>handle</b>	host	handle to the cuSPARSE library context.
<b>m</b>	host	number of rows of matrix <b>A</b> .
<b>n</b>	host	number of columns of matrix <b>A</b> .
<b>nnz</b>	host	number of nonzero elements of matrix <b>A</b> .
<b>cscColPtr</b>	device	integer array of <b>n+1</b> elements that contains the start of every column and the end of the last column plus one.
<b>cscRowInd</b>	device	integer array of <b>nnz</b> unsorted row indices of <b>A</b> .
<b>P</b>	device	integer array of <b>nnz</b> unsorted map indices. To construct <b>cscVal</b> , the user has to set <b>P=0:1:(nnz-1)</b> .
<b>pBuffer</b>	device	buffer allocated by the user; the size is returned by <b>cscsort_bufferSizeExt()</b> .

## Output

parameter	device or host	description
<b>cscRowInd</b>	device	integer array of <b>nnz</b> sorted row indices of <b>A</b> .

<b>P</b>	<b>device</b>	integer array of <b>nnz</b> sorted map indices.
<b>pBufferSizeInBytes</b>	<b>host</b>	number of bytes of the buffer.

See [cusparseStatus\\_t](#) for the description of the return status

## 13.28. cusparseXcsru2csr()

```

cusparseStatus_t
cusparseCreateCsru2csrInfo(csru2csrInfo_t *info);

cusparseStatus_t
cusparseDestroyCsru2csrInfo(csru2csrInfo_t info);

cusparseStatus_t
cusparseScsru2csr_bufferSizeExt(cusparseHandle_t handle,
                                int m,
                                int n,
                                int nnz,
                                float* csrVal,
                                const int* csrRowPtr,
                                int* csrColInd,
                                csru2csrInfo_t info,
                                size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseDcsru2csr_bufferSizeExt(cusparseHandle_t handle,
                                int m,
                                int n,
                                int nnz,
                                double* csrVal,
                                const int* csrRowPtr,
                                int* csrColInd,
                                csru2csrInfo_t info,
                                size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseCcsru2csr_bufferSizeExt(cusparseHandle_t handle,
                                int m,
                                int n,
                                int nnz,
                                cuComplex* csrVal,
                                const int* csrRowPtr,
                                int* csrColInd,
                                csru2csrInfo_t info,
                                size_t* pBufferSizeInBytes)

cusparseStatus_t
cusparseZcsru2csr_bufferSizeExt(cusparseHandle_t handle,
                                int m,
                                int n,
                                int nnz,
                                cuDoubleComplex* csrVal,
                                const int* csrRowPtr,
                                int* csrColInd,
                                csru2csrInfo_t info,
                                size_t* pBufferSizeInBytes)

```

```

cusparseStatus_t
cusparseScsru2csr(cusparseHandle_t handle,
                  int m,
                  int n,
                  int nnz,
                  const cusparseMatDescr_t descrA,
                  float* csrVal,
                  const int* csrRowPtr,
                  int* csrColInd,
                  csru2csrInfo_t info,
                  void* pBuffer)

```

This function transfers unsorted CSR format to CSR format, and vice versa. The operation is in-place.

This function is a wrapper of **csrsort** and **gthr**. The usecase is the following scenario.

If the user has a matrix **A** of CSR format which is unsorted, and implements his own code (which can be CPU or GPU kernel) based on this special order (for example, diagonal first, then lower triangle, then upper triangle), and wants to convert it to CSR format when calling CUSPARSE library, and then convert it back when doing something else on his/her kernel. For example, suppose the user wants to solve a linear system **Ax=b** by the following iterative scheme

$$x^{(k+1)} = x^{(k)} + L^{(-1)*} (b - Ax^{(k)})$$

The code heavily uses SpMv and triangular solve. Assume that the user has an in-house design of SpMV (Sparse Matrix-Vector multiplication) based on special order of **A**. However the user wants to use CUSPARSE library for triangular solver. Then the following code can work.

```

do
    step 1: compute residual vector           r      =      b      -      A      x(k) by
    step 2: B := sort(A), and L is lower triangular part of B
            (only sort A once and keep the permutation vector)
            step 3: solve                  z      =      L(-1)      *      (      b
            step 4: add correction          x(k+1) =      x(k)      +      z
            step 5: A := unsort(B)
            (use permutation vector to get back the unsorted CSR)
until convergence

```

The requirements of step 2 and step 5 are

1. In-place operation.
2. The permutation vector **P** is hidden in an opaque structure.
3. No **cudaMalloc** inside the conversion routine. Instead, the user has to provide the buffer explicitly.
4. The conversion between unsorted CSR and sorted CSR may needs several times, but the function only generates the permutation vector **P** once.
5. The function is based on **csrsort**, **gather** and **scatter** operations.

The operation is called **csru2csr**, which means unsorted CSR to sorted CSR. Also we provide the inverse operation, called **csr2csru**.

In order to keep the permutation vector invisible, we need an opaque structure called **csru2csrInfo**. Then two functions (**cusparseCreateCsru2csrInfo**, **cusparseDestroyCsru2csrInfo**) are used to initialize and to destroy the opaque structure.

**cusparse[S|D|C|Z]csru2csr\_bufferSizeExt** returns the size of the buffer. The permutation vector **P** is also allcated inside **csru2csrInfo**. The lifetime of the permutation vector is the same as the lifetime of **csru2csrInfo**.

**`cusparse[S|D|C|Z]csru2csr`** performs forward transformation from unsorted CSR to sorted CSR. First call uses `csrssort` to generate the permutation vector `P`, and subsequent call uses `P` to do transformation.

**`cusparse[S|D|C|Z]csr2csru`** performs backward transformation from sorted CSR to unsorted CSR. `P` is used to get unsorted form back.

The routine `cusparse<t>csru2csr()` has the following properties:

- ▶ The routine requires no extra storage if `pBuffer != NULL`
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine `cusparse<t>csr2csru()` has the following properties if `pBuffer != NULL`:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

The following tables describe parameters of `csr2csru_bufferSizeExt` and `csr2csru`.

## Input

parameter	device or host	description
<code>handle</code>	host	handle to the cuSPARSE library context.
<code>m</code>	host	number of rows of matrix <code>A</code> .
<code>n</code>	host	number of columns of matrix <code>A</code> .
<code>nnz</code>	host	number of nonzero elements of matrix <code>A</code> .
<code>descrA</code>	host	the descriptor of matrix <code>A</code> . The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<code>csrVal</code>	device	<type> array of <code>nnz</code> unsorted nonzero elements of matrix <code>A</code> .
<code>csrRowsPtr</code>	device	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColInd</code>	device	integer array of <code>nnz</code> unsorted column indices of <code>A</code> .
<code>info</code>	host	opaque structure initialized using <code>cusparseCreateCsru2csrInfo()</code> .
<code>pBuffer</code>	device	buffer allocated by the user; the size is returned by <code>csru2csr_bufferSizeExt()</code> .

## Output

parameter	device or host	description
<code>csrVal</code>	device	<type> array of <code>nnz</code> sorted nonzero elements of matrix <code>A</code> .
<code>csrColInd</code>	device	integer array of <code>nnz</code> sorted column indices of <code>A</code> .

<code>pBufferSizeInBytes</code>	host	number of bytes of the buffer.
---------------------------------	------	--------------------------------

See `cusparseStatus_t` for the description of the return status

## 13.29. cusparseXpruneDense2csr()

```
cusparseStatus_t  
cusparseHpruneDense2csrNnz(cusparseHandle_t handle,  
                             int m,  
                             int n,  
                             const __half* A,  
                             int lda,  
                             const __half* threshold,  
                             const cusparseMatDescr_t descrC,  
                             int* csrRowPtrC,  
                             int* nnzTotalDevHostPtr,  
                             void* pBuffer)
```

This function prunes a dense matrix to a sparse matrix with CSR format.

Given a dense matrix **A** and a non-negative value **threshold**, the function returns a sparse matrix **C**, defined by

$$C(i,j) = A(i,j) \text{ if } |A(i,j)| > \text{threshold}$$

The implementation adopts a two-step approach to do the conversion. First, the user allocates **csrRowPtrC** of **m+1** elements and uses function **pruneDense2csrNnz()** to determine the number of nonzeros columns per row. Second, the user gathers **nnzC** (number of nonzeros of matrix **C**) from either (**nnzC=nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC[m] - csrRowPtrC[0]**) and allocates **csrValC** of **nnzC** elements and **csrColIndC** of **nnzC** integers. Finally function **pruneDense2csr()** is called to complete the conversion.

The user must obtain the size of the buffer required by **pruneDense2csr()** by calling **pruneDense2csr\_bufferSizeExt()**, allocate the buffer, and pass the buffer pointer to **pruneDense2csr()**.

Appendix E.1 provides a simple example of **pruneDense2csr()**.

The routine **cusparse<t>pruneDense2csrNnz()** has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine **cusparse<t>DpruneDense2csr()** has the following properties:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

parameter	device or host	description
<b>handle</b>	host	handle to the cuSPARSE library context.
<b>m</b>	host	number of rows of matrix <b>A</b> .
<b>n</b>	host	number of columns of matrix <b>A</b> .
<b>A</b>	device	array of dimension (lda, n).
<b>lda</b>	device	leading dimension of <b>A</b> . It must be at least max(1, m).
<b>threshold</b>	host or device	a value to drop the entries of <b>A</b> . <b>threshold</b> can point to a device memory or host memory.
<b>descrC</b>	host	the descriptor of matrix <b>C</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>pBuffer</b>	device	buffer allocated by the user; the size is returned by <b>pruneDense2csr_bufferSizeExt()</b> .

### Output

parameter	device or host	description
<code>nnzTotalDevHostPtr</code>	device or host	total number of nonzero of matrix c. <code>nnzTotalDevHostPtr</code> can point to a device memory or host memory.
<code>csrValC</code>	device	<type> array of <code>nnzc</code> nonzero elements of matrix c.
<code>csrRowsPtrC</code>	device	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<code>csrColIndC</code>	device	integer array of <code>nnzc</code> column indices of c.
<code>pBufferSizeInBytes</code>	host	number of bytes of the buffer.

See `cusparseStatus_t` for the description of the return status

### 13.30. cusparseXpruneCsr2csr()

This function prunes a sparse matrix to a sparse matrix with CSR format.

Given a sparse matrix **A** and a non-negative value **threshold**, the function returns a sparse matrix **C**, defined by

$$C(i,j) = A(i,j) \text{ if } |A(i,j)| > \text{threshold}$$

The implementation adopts a two-step approach to do the conversion. First, the user allocates **csrRowPtrC** of **m+1** elements and uses function **pruneCsr2csrNnz()** to determine the number of nonzeros columns per row. Second, the user gathers **nnzC** (number of nonzeros of matrix **C**) from either (**nnzC=nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC[m]-csrRowPtrC[0]**) and allocates **csrValC** of **nnzC** elements and **csrColIndC** of **nnzC** integers. Finally function **pruneCsr2csr()** is called to complete the conversion.

The user must obtain the size of the buffer required by **pruneCsr2csr()** by calling **pruneCsr2csr\_bufferSizeExt()**, allocate the buffer, and pass the buffer pointer to **pruneCsr2csr()**.

Appendix E.2 provides a simple example of **pruneCsr2csr()**.

The routine **cusparse<t>pruneCsr2csrNnz()** has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine **cusparse<t>pruneCsr2csr()** has the following properties:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

parameter	device or host	description
<b>handle</b>	<b>host</b>	handle to the cuSPARSE library context.
<b>m</b>	<b>host</b>	number of rows of matrix <b>A</b> .
<b>n</b>	<b>host</b>	number of columns of matrix <b>A</b> .
<b>nnzA</b>	<b>host</b>	number of nonzeros of matrix <b>A</b> .
<b>descrA</b>	<b>host</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<b>device</b>	<type> array of <b>nnzA</b> nonzero elements of matrix <b>A</b> .
<b>csrRowsPtrA</b>	<b>device</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	<b>device</b>	integer array of <b>nnzA</b> column indices of <b>A</b> .
<b>threshold</b>	<b>host or device</b>	a value to drop the entries of <b>A</b> . <b>threshold</b> can point to a device memory or host memory.

<b>descrC</b>	<b>host</b>	the descriptor of matrix c. The supported matrix type is <code>CUSPARSE_MATRIX_TYPE_GENERAL</code> . Also, the supported index bases are <code>CUSPARSE_INDEX_BASE_ZERO</code> and <code>CUSPARSE_INDEX_BASE_ONE</code> .
<b>pBuffer</b>	<b>device</b>	buffer allocated by the user; the size is returned by <code>pruneCsr2csr_bufferSizeExt()</code> .

## Output

parameter	device or host	description
<b>nnzTotalDevHostPtr</b>	<b>device or host</b>	total number of nonzero of matrix c. <code>nnzTotalDevHostPtr</code> can point to a device memory or host memory.
<b>csrValC</b>	<b>device</b>	<type> array of <code>nnzc</code> nonzero elements of matrix c.
<b>csrRowsPtrC</b>	<b>device</b>	integer array of <code>m+1</code> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndC</b>	<b>device</b>	integer array of <code>nnzc</code> column indices of c.
<b>pBufferSizeInBytes</b>	<b>host</b>	number of bytes of the buffer.

See `cusparseStatus_t` for the description of the return status

### 13.31. `cusparseXpruneDense2csrPercentage()`

```

cusparseStatus_t
cusparseHpruneDense2csrByPercentage_bufferSizeExt(cusparseHandle_t
    handle,
                                int
    m,
                                int
    n,
                                const __half*
    A,
                                int
    lda,
                                float
    percentage,
                                const
cusparseMatDescr_t descrC,
                                const __half*
    csrValC,
                                const int*
    csrRowPtrC,
                                const int*
    csrColIndC,
                                pruneInfo_t
    info,
                                size_t*
    pBufferSizeInBytes)

cusparseStatus_t
cusparseSpruneDense2csrByPercentage_bufferSizeExt(cusparseHandle_t
    handle,
                                int
    m,
                                int
    n,
                                const float*
    A,
                                int
    lda,
                                float
    percentage,
                                const
cusparseMatDescr_t descrC,
                                const float*
    csrValC,
                                const int*
    csrRowPtrC,
                                const int*
    csrColIndC,
                                pruneInfo_t
    info,
                                size_t*
    pBufferSizeInBytes)

cusparseStatus_t
cusparseDpruneDense2csrByPercentage_bufferSizeExt(cusparseHandle_t
    handle,
                                int
    m,
                                int
    n,
                                const double*
    A,
                                int
    lda.

```

This function prunes a dense matrix to a sparse matrix by percentage.

Given a dense matrix **A** and a non-negative value **percentage**, the function computes sparse matrix **C** by the following three steps:

Step 1: sort absolute value of **A** in ascending order.

```
key := sort( |A| )
```

Step 2: choose threshold by the parameter **percentage**

```
pos = ceil(m*n*(percentage/100)) - 1
pos = min(pos, m*n-1)
pos = max(pos, 0)
threshold = key[pos]
```

Step 3: call **pruneDense2csr()** by with the parameter **threshold**.

The implementation adopts a two-step approach to do the conversion.

First, the user allocates **csrRowPtrC** of **m+1** elements and uses function **pruneDense2csrNnzByPercentage()** to determine the number of nonzeros columns per row. Second, the user gathers **nnzC** (number of nonzeros of matrix **C**) from either (**nnzC=nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC[m]-csrRowPtrC[0]**) and allocates **csrValC** of **nnzC** elements and **csrColIndC** of **nnzC** integers. Finally function **pruneDense2csrByPercentage()** is called to complete the conversion.

The user must obtain the size of the buffer required

by **pruneDense2csrByPercentage()** by calling

**pruneDense2csrByPercentage\_bufferSizeExt()**, allocate the buffer, and pass the buffer pointer to **pruneDense2csrByPercentage()**.

Remark 1: the value of **percentage** must be not greater than 100. Otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Remark 2: the zeros of **A** are not ignored. All entries are sorted, including zeros. This is different from **pruneCsr2csrByPercentage()**

Appendix E.3 provides a simple example of **pruneDense2csrNnzByPercentage()**.

The routine **cusparse<t>pruneDense2csrNnzByPercentage()** has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine **cusparse<t>pruneDense2csrByPercentage()** has the following properties:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

### Input

parameter	device or host	description
-----------	----------------	-------------

<b>handle</b>	<b>host</b>	handle to the cuSPARSE library context.
<b>m</b>	<b>host</b>	number of rows of matrix <b>A</b> .
<b>n</b>	<b>host</b>	number of columns of matrix <b>A</b> .
<b>A</b>	<b>device</b>	array of dimension (lda, n).
<b>lda</b>	<b>device</b>	leading dimension of <b>A</b> . It must be at least max(1, m).
<b>percentage</b>	<b>host</b>	percentage <=100 and percentage >= 0
<b>descrC</b>	<b>host</b>	the descriptor of matrix c. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>pBuffer</b>	<b>device</b>	buffer allocated by the user; the size is returned by <b>pruneDense2csrByPercentage_bufferSizeExt()</b> .

## Output

<b>parameter</b>	<b>device or host</b>	<b>description</b>
<b>nnzTotalDevHostPtr</b>	<b>device or host</b>	total number of nonzero of matrix c. <b>nnzTotalDevHostPtr</b> can point to a device memory or host memory.
<b>csrValC</b>	<b>device</b>	<type> array of <b>nnzC</b> nonzero elements of matrix c.
<b>csrRowsPtrC</b>	<b>device</b>	integer array of <b>m+1</b> elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndC</b>	<b>device</b>	integer array of <b>nnzC</b> column indices of c.
<b>pBufferSizeInBytes</b>	<b>host</b>	number of bytes of the buffer.

See [cusparseStatus\\_t](#) for the description of the return status

### 13.32. cusparseXpruneCsr2csrByPercentage()

This function prunes a sparse matrix to a sparse matrix by percentage.

Given a sparse matrix **A** and a non-negative value **percentage**, the function computes sparse matrix **C** by the following three steps:

Step 1: sort absolute value of **A** in ascending order.

```
key := sort( |csrValA| )
```

Step 2: choose threshold by the parameter **percentage**

```
pos = ceil(nnzA*(percentage/100)) - 1
pos = min(pos, nnzA-1)
pos = max(pos, 0)
threshold = key[pos]
```

Step 3: call **pruneCsr2csr()** by with the parameter **threshold**.

The implementation adopts a two-step approach to do the conversion.

First, the user allocates **csrRowPtrC** of **m+1** elements and uses function **pruneCsr2csrNnzByPercentage()** to determine the number of nonzeros columns per row. Second, the user gathers **nnzC** (number of nonzeros of matrix **C**) from either (**nnzC=nnzTotalDevHostPtr**) or (**nnzC=csrRowPtrC[m]-csrRowPtrC[0]**) and allocates **csrValC** of **nnzC** elements and **csrColIndC** of **nnzC** integers. Finally function **pruneCsr2csrByPercentage()** is called to complete the conversion.

The user must obtain the size of the buffer required by **pruneCsr2csrByPercentage()** by calling **pruneCsr2csrByPercentage\_bufferSizeExt()**, allocate the buffer, and pass the buffer pointer to **pruneCsr2csrByPercentage()**.

Remark 1: the value of **percentage** must be not greater than 100. Otherwise, **CUSPARSE\_STATUS\_INVALID\_VALUE** is returned.

Appendix E.4 provides a simple example of **pruneCsr2csrByPercentage()**.

The routine **cusparse<t>pruneCsr2csrNnzByPercentage()** has the following properties:

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution
- ▶ The routine does *not* support CUDA graph capture

The routine **cusparse<t>pruneCsr2csrByPercentage()** has the following properties:

- ▶ The routine requires no extra storage
- ▶ The routine supports asynchronous execution
- ▶ The routine supports CUDA graph capture

## Input

parameter	device or host	description
<b>handle</b>	host	handle to the cuSPARSE library context.
<b>m</b>	host	number of rows of matrix <b>A</b> .
<b>n</b>	host	number of columns of matrix <b>A</b> .

<b>nnzA</b>	<b>host</b>	number of nonzeros of matrix A.
<b>descrA</b>	<b>host</b>	the descriptor of matrix A. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	<b>device</b>	<type> array of nnzA nonzero elements of matrix A.
<b>csrRowsPtrA</b>	<b>device</b>	integer array of m+1 elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndA</b>	<b>device</b>	integer array of nnzA column indices of A.
<b>percentage</b>	<b>host</b>	percentage <=100 and percentage >= 0
<b>descrC</b>	<b>host</b>	the descriptor of matrix c. The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> , Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>pBuffer</b>	<b>device</b>	buffer allocated by the user; the size is returned by <b>pruneCsr2csrByPercentage_bufferSizeExt()</b> .

## Output

<b>parameter</b>	<b>device or host</b>	<b>description</b>
<b>nnzTotalDevHostPtr</b>	<b>device or host</b>	total number of nonzero of matrix c. <b>nnzTotalDevHostPtr</b> can point to a device memory or host memory.
<b>csrValC</b>	<b>device</b>	<type> array of nnzC nonzero elements of matrix c.
<b>csrRowsPtrC</b>	<b>device</b>	integer array of m+1 elements that contains the start of every row and the end of the last row plus one.
<b>csrColIndC</b>	<b>device</b>	integer array of nnzC column indices of c.
<b>pBufferSizeInBytes</b>	<b>host</b>	number of bytes of the buffer.

See [cusparseStatus\\_t](#) for the description of the return status

## 13.33. `cusparse<t>nnz_compress()`

```

cusparseStatus_t
cusparseSnnz_compress(cusparseHandle_t          handle,
                     int                  m,
                     const cusparseMatDescr_t descr,
                     const float*          csrValA,
                     const int*             csrRowPtrA,
                     int*                  nnzPerRow,
                     int*                  nnzC,
                     float                 tol)

cusparseStatus_t
cusparseDnnz_compress(cusparseHandle_t          handle,
                     int                  m,
                     const cusparseMatDescr_t descr,
                     const double*          csrValA,
                     const int*             csrRowPtrA,
                     int*                  nnzPerRow,
                     int*                  nnzC,
                     double                tol)

cusparseStatus_t
cusparseCnnz_compress(cusparseHandle_t          handle,
                     int                  m,
                     const cusparseMatDescr_t descr,
                     const cuComplex*       csrValA,
                     const int*             csrRowPtrA,
                     int*                  nnzPerRow,
                     int*                  nnzC,
                     cuComplex             tol)

cusparseStatus_t
cusparseZnnz_compress(cusparseHandle_t          handle,
                     int                  m,
                     const cusparseMatDescr_t descr,
                     const cuDoubleComplex* csrValA,
                     const int*             csrRowPtrA,
                     int*                  nnzPerRow,
                     int*                  nnzC,
                     cuDoubleComplex        tol)

```

This function is the step one to convert from csr format to compressed csr format.

Given a sparse matrix A and a non-negative value threshold, the function returns nnzPerRow(the number of nonzeros columns per row) and nnzC(the total number of nonzeros) of a sparse matrix C, defined by

$$C(i,j) = A(i,j) \text{ if } |A(i,j)| > \text{threshold}$$

A key assumption for the cuComplex and cuDoubleComplex case is that this tolerance is given as the real part. For example tol = 1e-8 + 0\*i and we extract cureal, that is the x component of this struct.

- ▶ This function requires temporary extra storage that is allocated internally
- ▶ The routine does *not* support asynchronous execution

- The routine does *not* support CUDA graph capture

### Input

<b>handle</b>	handle to the cuSPARSE library context.
<b>m</b>	number of rows of matrix <b>A</b> .
<b>descrA</b>	the descriptor of matrix <b>A</b> . The supported matrix type is <b>CUSPARSE_MATRIX_TYPE_GENERAL</b> . Also, the supported index bases are <b>CUSPARSE_INDEX_BASE_ZERO</b> and <b>CUSPARSE_INDEX_BASE_ONE</b> .
<b>csrValA</b>	csr noncompressed values array
<b>csrRowPtrA</b>	the corresponding input noncompressed row pointer.
<b>tol</b>	non-negative tolerance to determine if a number less than or equal to it.

### Output

<b>nnzPerRow</b>	this array contains the number of elements whose absolute values are greater than tol per row.
<b>nnzC</b>	host/device pointer of the total number of elements whose absolute values are greater than tol.

See [cusparseStatus\\_t](#) for the description of the return status

# Chapter 14.

# CUSPARSE GENERIC API REFERENCE

Starting with CUDA 10.1, the cuSPARSE Library provides a new set of API to compute sparse vector-vector (SpVV), matrix-vector (SpMV), and matrix-matrix multiplication (SpMM), in addition to the existing legacy API. They are considered a *preview feature* i.e. the data structures and APIs for them are subject to change and may not be compatible with future releases.

**LIMITATION:** The generic APIs are currently available for all platforms except Windows. Using these APIs in any other systems will result in compile-time or run-time failures. Their support will be extended in the next releases.

## 14.1. Generic Types Reference

The cuSPARSE generic type references are described in this section.

### 14.1.1. `cudaDataType_t`

The section describes types shared by multiple CUDA Libraries and defined in the header file `library_types.h`. The `cudaDataType_t` type is an enumerator to specify the data precision. It is used when the data reference does not carry the type itself (e.g `void *`). For example, it is used in the routine `cusparseSpMM()`.

Value	Meaning	Data Type	Header
<code>CUDA_R_16F</code>	The data type is 16-bit floating-point	<code>__half</code>	<code>cuda_fp16.h</code>
<code>CUDA_C_16F</code>	The data type is 16-bit complex floating-point	<code>__half2</code>	<code>cuda_fp16.h</code>
<code>CUDA_R_32F</code>	The data type is 32-bit floating-point	<code>float</code>	
<code>CUDA_C_32F</code>	The data type is 32-bit complex floating-point	<code>cuComplex</code>	<code>cuComplex.h</code>
<code>CUDA_F_64F</code>	The data type is 64-bit floating-point	<code>double</code>	
<code>CUDA_C_64F</code>	The data type is 64-bit complex floating-point	<code>cuDoubleComplex</code>	<code>cuComplex.h</code>
<code>CUDA_R_8I</code>	The data type is 8-bit integer	<code>int8_t</code>	<code>stdint.h</code>
<code>CUDA_R_32I</code>	The data type is 32-bit integer	<code>int32_t</code>	<code>stdint.h</code>

## 14.1.2. `cusparseFormat_t`

This type indicates the format of the sparse matrix.

Value	Meaning
<code>CUSPARSE_FORMAT_COO</code>	The matrix is stored in Coordinate (COO) format organized in <i>Structure of Arrays (SoA)</i> layout
<code>CUSPARSE_FORMAT_COO_AOS</code>	The matrix is stored in Coordinate (COO) format organized in <i>Array of Structures (SoA)</i> layout
<code>CUSPARSE_FORMAT_CSR</code>	The matrix is stored in Compressed Sparse Row (CSR) format

## 14.1.3. `cusparseOrder_t`

This type indicates the memory layout of a dense matrix. Currently, only column-major layout is supported.

Value	Meaning
<code>CUSPARSE_ORDER_COL</code>	The matrix is stored in column-major

## 14.1.4. `cusparseIndexType_t`

This type indicates the index type for representing the sparse matrix indices.

Value	Meaning
<code>CUSPARSE_INDEX_16U</code>	16-bit unsigned integer [1, 65535]
<code>CUSPARSE_INDEX_32I</code>	32-bit signed integer [1, $2^{31} - 1$ ]
<code>CUSPARSE_INDEX_64I</code>	64-bit signed integer [1, $2^{63} - 1$ ]

## 14.2. Sparse Vector APIs

The cuSPARSE helper functions for sparse vector descriptor are described in this section.

### 14.2.1. `cusparseCreateSpVec()`

```
cusparseStatus_t
cusparseCreateSpVec(cusparseSpVecDescr_t* spVecDescr,
                    int64_t           size,
                    int64_t           nnz,
                    void*            indices,
                    void*            values,
                    cusparseIndexType_t idxType,
                    cusparseIndexBase_t idxBase,
                    cudaDataType     valueType)
```

This function initializes the sparse matrix descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	OUT	Sparse vector descriptor
<code>size</code>	HOST	IN	Size of the sparse vector
<code>nnz</code>	HOST	IN	Number of non-zero entries of the sparse vector
<code>indices</code>	DEVICE	IN	Indices of the sparse vector. Array of size <code>nnz</code>
<code>values</code>	DEVICE	IN	Values of the sparse vector. Array of size <code>nnz</code>
<code>idxType</code>	HOST	IN	Enumerator specifying the data type of <code>indices</code>
<code>idxBase</code>	HOST	IN	Enumerator specifying the the base index of <code>indices</code>
<code>valueType</code>	HOST	IN	Enumerator specifying the datatype of <code>values</code>

See `cusparseStatus_t` for the description of the return status

### 14.2.2. `cusparseDestroySpVec()`

```
cusparseStatus_t
cusparseDestroySpVec(cusparseSpVecDescr_t spVecDescr)
```

This function releases the host memory allocated for the sparse vector descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	IN	Sparse vector descriptor

See `cusparseStatus_t` for the description of the return status

### 14.2.3. `cusparseSpVecGet()`

```
cusparseStatus_t
cusparseSpVecGet(const cusparseSpVecDescr_t spVecDescr,
                 int64_t*                      size,
                 int64_t*                      nnz,
                 void**                        indices,
                 void**                        values,
                 cusparseIndexType_t*          idxType,
                 cusparseIndexBase_t*          idxBase,
                 cudaDataType*                valueType)
```

This function returns the fields of the sparse vector descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	IN	Sparse vector descriptor
<code>size</code>	HOST	OUT	Size of the sparse vector
<code>nnz</code>	HOST	OUT	Number of non-zero entries of the sparse vector
<code>indices</code>	DEVICE	OUT	Indices of the sparse vector. Array of size <code>nnz</code>
<code>values</code>	DEVICE	OUT	Values of the sparse vector. Array of size <code>nnz</code>
<code>idxType</code>	HOST	OUT	Enumerator specifying the data type of <code>indices</code>

Param.	Memory	In/out	Meaning
<code>idxBase</code>	HOST	OUT	Enumerator specifying the the base index of <code>indices</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>values</code>

See `cusparseStatus_t` for the description of the return status

#### 14.2.4. `cusparseSpVecGetIndexBase()`

```
cusparseStatus_t
cusparseSpVecGetIndexBase(const cusparseSpVecDescr_t spVecDescr,
                           cusparseIndexBase_t*           idxBase)
```

This function returns the `idxBase` field of the sparse vector descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	IN	Sparse vector descriptor
<code>idxBase</code>	HOST	OUT	Enumerator specifying the the base index of <code>indices</code>

See `cusparseStatus_t` for the description of the return status

#### 14.2.5. `cusparseSpVecGetValues()`

```
cusparseStatus_t
cusparseSpVecGetValues(const cusparseSpVecDescr_t spVecDescr,
                       void**                  values)
```

This function returns the `values` field of the sparse vector descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	IN	Sparse vector descriptor
<code>values</code>	DEVICE	OUT	Values of the sparse vector. Array of size <code>nnz</code>

See `cusparseStatus_t` for the description of the return status

#### 14.2.6. `cusparseSpVecSetValues()`

```
cusparseStatus_t
cusparseSpVecSetValues(cusparseSpVecDescr_t spVecDescr,
                       void*                  values)
```

This function set the `values` field of the sparse vector descriptor `spVecDescr`.

Param.	Memory	In/out	Meaning
<code>spVecDescr</code>	HOST	IN	Sparse vector descriptor
<code>values</code>	DEVICE	IN	Values of the sparse vector. Array of size <code>nnz</code>

See `cusparseStatus_t` for the description of the return status

## 14.3. Sparse Matrix APIs

The cuSPARSE helper functions for sparse matrix descriptor are described in this section.

### 14.3.1. `cusparseCreateCoo()`

```
cusparseStatus_t
cusparseCreateCoo(cusparseSpMatDescr_t* spMatDescr,
                  int64_t           rows,
                  int64_t           cols,
                  int64_t           nnz,
                  void*            cooRowInd,
                  void*            cooColInd,
                  void*            cooValues,
                  cusparseIndexType_t cooIdxType,
                  cusparseIndexBase_t idxBase,
                  cudaDataType     valueType)
```

This function initializes the sparse matrix descriptor `spMatDescr` in the COO format (Structure of Arrays layout).

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	OUT	Sparse matrix descriptor
<code>rows</code>	HOST	IN	Number of rows of the sparse matrix
<code>cols</code>	HOST	IN	Number of columns of the sparse matrix
<code>nnz</code>	HOST	IN	Number of non-zero entries of the sparse matrix
<code>cooRowInd</code>	DEVICE	IN	Row indices of the sparse matrix. Array of size <code>nnz</code>
<code>cooColInd</code>	DEVICE	IN	Column indices of the sparse matrix. Array of size <code>nnz</code>
<code>cooValues</code>	DEVICE	IN	Values of the sparse matrix. Array of size <code>nnz</code>
<code>cooIdxType</code>	HOST	IN	Enumerator specifying the data type of <code>cooRowInd</code> and <code>cooColInd</code>
<code>idxBase</code>	HOST	IN	Enumerator specifying the base index of <code>cooRowInd</code> and <code>cooColInd</code>
<code>valueType</code>	HOST	IN	Enumerator specifying the datatype of <code>cooValues</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.2. `cusparseCreateCooAoS()`

```
cusparseStatus_t
cusparseCreateCooAoS(cusparseSpMatDescr_t* spMatDescr,
                     int64_t           rows,
                     int64_t           cols,
                     int64_t           nnz,
                     void*            cooInd,
                     void*            cooValues,
                     cusparseIndexType_t   cooIdxType,
                     cusparseIndexBase_t    idxBase,
                     cudaDataType        valueType)
```

This function initializes the sparse matrix descriptor **spMatDescr** in the COO format (Array of Structures layout).

Param.	Memory	In/out	Meaning
<b>spMatDescr</b>	HOST	OUT	Sparse matrix descriptor
<b>rows</b>	HOST	IN	Number of rows of the sparse matrix
<b>cols</b>	HOST	IN	Number of columns of the sparse matrix
<b>nnz</b>	HOST	IN	Number of non-zero entries of the sparse matrix
<b>cooInd</b>	DEVICE	IN	<Row, Column> indices of the sparse matrix. Array of size <b>nnz</b>
<b>cooValues</b>	DEVICE	IN	Values of the sparse matrix. Array of size <b>nnz</b>
<b>cooIdxType</b>	HOST	IN	Enumerator specifying the data type of <b>cooInd</b>
<b>idxBase</b>	HOST	IN	Enumerator specifying the base index of <b>cooInd</b>
<b>valueType</b>	HOST	IN	Enumerator specifying the datatype of <b>cooValues</b>

See `cusparseStatus_t` for the description of the return status

### 14.3.3. `cusparseCreateCsr()`

```
cusparseStatus_t
cusparseCreateCsr(cusparseSpMatDescr_t* spMatDescr,
                  int64_t           rows,
                  int64_t           cols,
                  int64_t           nnz,
                  void*            csrRowOffsets,
                  void*            csrColInd,
                  void*            csrValues,
                  cusparseIndexType_t   csrRowOffsetsType,
                  cusparseIndexType_t   csrColIndType,
                  cusparseIndexBase_t    idxBase,
                  cudaDataType        valueType)
```

This function initializes the sparse matrix descriptor **spMatDescr** in the CSR format.

Param.	Memory	In/out	Meaning
<b>spMatDescr</b>	HOST	OUT	Sparse matrix descriptor

Param.	Memory	In/out	Meaning
<code>rows</code>	HOST	IN	Number of rows of the sparse matrix
<code>cols</code>	HOST	IN	Number of columns of the sparse matrix
<code>nnz</code>	HOST	IN	Number of non-zero entries of the sparse matrix
<code>csrRowOffs</code>	DEVICE	IN	Row offsets of the sparse matrix. Array of size <code>rows</code> + 1
<code>csrColInd</code>	DEVICE	IN	Column indices of the sparse matrix. Array of size <code>nnz</code>
<code>csrValues</code>	DEVICE	IN	Values of the sparse matrix. Array of size <code>nnz</code>
<code>csrRowOffsT</code>	HOST	IN	Enumerator specifying the data type of <code>csrRowOffsets</code>
<code>csrColIndT</code>	HOST	IN	Enumerator specifying the data type of <code>csrColInd</code>
<code>idxBase</code>	HOST	IN	Enumerator specifying the base index of <code>csrRowOffsets</code> and <code>csrColInd</code>
<code>valueType</code>	HOST	IN	Enumerator specifying the datatype of <code>csrValues</code>

See `cusparseStatus_t` for the description of the return status

#### 14.3.4. `cusparseDestroySpMat()`

```
cusparseStatus_t
cusparseDestroySpMat(cusparseSpMatDescr_t spMatDescr)
```

This function releases the host memory allocated for the sparse matrix descriptor `spMatDescr`.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor

See `cusparseStatus_t` for the description of the return status

#### 14.3.5. `cusparseCooGet()`

```
cusparseStatus_t
cusparseCooGet(const cusparseSpMatDescr_t spMatDescr,
               int64_t*           rows,
               int64_t*           cols,
               int64_t*           nnz,
               void**            cooRowInd,
               void**            cooColInd,
               void**            cooValues,
               cusparseIndexType_t* idxType,
               cusparseIndexBase_t* idxBase,
               cudaDataType*      valueType)
```

This function returns the fields of the sparse matrix descriptor `spMatDescr` stored in COO format (Array of Structures layout).

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor

Param.	Memory	In/out	Meaning
<code>rows</code>	HOST	OUT	Number of rows of the sparse matrix
<code>cols</code>	HOST	OUT	Number of columns of the sparse matrix
<code>nnz</code>	HOST	OUT	Number of non-zero entries of the sparse matrix
<code>cooRowInd</code>	DEVICE	OUT	Row indices of the sparse matrix. Array of size <code>nnz</code>
<code>cooColInd</code>	DEVICE	OUT	Column indices of the sparse matrix. Array of size <code>nnz</code>
<code>cooValues</code>	DEVICE	OUT	Values of the sparse matrix. Array of size <code>nnz</code>
<code>cooIdxType</code>	HOST	OUT	Enumerator specifying the data type of <code>cooRowInd</code> and <code>cooColInd</code>
<code>idxBase</code>	HOST	OUT	Enumerator specifying the base index of <code>cooRowInd</code> and <code>cooColInd</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>cooValues</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.6. `cusparseCooAoSGet()`

```
cusparseStatus_t
cusparseCooAoSGet(const cusparseSpMatDescr_t spMatDescr,
                  int64_t*                      rows,
                  int64_t*                      cols,
                  int64_t*                      nnz,
                  void**                        cooInd,
                  void**                        cooValues,
                  cusparseIndexType_t*          idxType,
                  cusparseIndexBase_t*          idxBase,
                  cudaDataType*                 valueType)
```

This function returns the fields of the sparse matrix descriptor `spMatDescr` stored in COO format (Structure of Arrays layout).

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>rows</code>	HOST	OUT	Number of rows of the sparse matrix
<code>cols</code>	HOST	OUT	Number of columns of the sparse matrix
<code>nnz</code>	HOST	OUT	Number of non-zero entries of the sparse matrix
<code>cooInd</code>	DEVICE	OUT	<Row, Column> indices of the sparse matrix. Array of size <code>nnz</code>
<code>cooValues</code>	DEVICE	OUT	Values of the sparse matrix. Array of size <code>nnz</code>
<code>cooIdxType</code>	HOST	OUT	Enumerator specifying the data type of <code>cooInd</code>
<code>idxBase</code>	HOST	OUT	Enumerator specifying the base index of <code>cooInd</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>cooValues</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.7. `cusparseCsrGet()`

```
cusparseStatus_t CUSPARSEAPI
cusparseCsrGet(const cusparseSpMatDescr_t spMatDescr,
               int64_t* rows,
               int64_t* cols,
               int64_t* nnz,
               void** csrRowOffsets,
               void** csrColInd,
               void** csrValues,
               cusparseIndexType_t* csrRowOffsetsType,
               cusparseIndexType_t* csrColIndType,
               cusparseIndexBase_t* idxBase,
               cudaDataType* valueType);
```

This function returns the fields of the sparse matrix descriptor `spMatDescr` stored in CSR format.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>rows</code>	HOST	OUT	Number of rows of the sparse matrix
<code>cols</code>	HOST	OUT	Number of columns of the sparse matrix
<code>nnz</code>	HOST	OUT	Number of non-zero entries of the sparse matrix
<code>csrRowOffs</code>	DEVICE	OUT	Row offsets of the sparse matrix. Array of size <code>rows + 1</code>
<code>csrColInd</code>	DEVICE	OUT	Column indices of the sparse matrix. Array of size <code>nnz</code>
<code>csrValues</code>	DEVICE	OUT	Values of the sparse matrix. Array of size <code>nnz</code>
<code>csrRowOffsT</code>	HOST	OUT	Enumerator specifying the data type of <code>csrRowOffsets</code>
<code>csrColIndT</code>	HOST	OUT	Enumerator specifying the data type of <code>csrColInd</code>
<code>idxBase</code>	HOST	OUT	Enumerator specifying the base index of <code>csrRowOffsets</code> and <code>csrColInd</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>csrValues</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.8. `cusparseSpMatGetFormat()`

```
cusparseStatus_t
cusparseSpMatGetFormat(const cusparseSpMatDescr_t spMatDescr,
                      cusparseFormat_t* format)
```

This function returns the `format` field of the sparse matrix descriptor `spMatDescr`.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>format</code>	HOST	OUT	Enumerator specifying the storage format of the sparse matrix

See `cusparseStatus_t` for the description of the return status

### 14.3.9. `cusparseSpMatGetIndexBase()`

```
cusparseStatus_t
cusparseSpMatGetIndexBase(const cusparseSpMatDescr_t spMatDescr,
                           cusparseIndexBase_t*           idxBase)
```

This function returns the `idxBase` field of the sparse matrix descriptor `spMatDescr`.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>idxBase</code>	HOST	OUT	Enumerator specifying the base index of the sparse matrix

See `cusparseStatus_t` for the description of the return status

### 14.3.10. `cusparseSpMatGetValues()`

```
cusparseStatus_t CUSPARSEAPI
cusparseSpMatGetValues(cusparseSpMatDescr_t spMatDescr,
                      void**                  values)
```

This function returns the `values` field of the sparse matrix descriptor `spMatDescr`.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>values</code>	DEVICE	OUT	Values of the sparse martix. Array of size <code>nnz</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.11. `cusparseSpMatSetValues()`

```
cusparseStatus_t CUSPARSEAPI
cusparseSpMatSetValues(cusparseSpMatDescr_t spMatDescr,
                      void*                  values)
```

This function sets the `values` field of the sparse matrix descriptor `spMatDescr`.

Param.	Memory	In/out	Meaning
<code>spMatDescr</code>	HOST	IN	Sparse matrix descriptor
<code>values</code>	DEVICE	IN	Values of the sparse martix. Array of size <code>nnz</code>

See `cusparseStatus_t` for the description of the return status

### 14.3.12. `cusparseSpMatGetStridedBatch()`

```
cusparseStatus_t
cusparseSpMatGetStridedBatch(const cusparseSpMatDescr_t spMatDescr,
                             int*                  batchCount)
```

This function returns the **batchCount** field of the sparse matrix descriptor **spMatDescr**.

Param.	Memory	In/out	Meaning
<b>spMatDescr</b>	HOST	IN	Sparse matrix descriptor
<b>batchCount</b>	HOST	OUT	Number of batches of the sparse matrix

See [cusparseStatus\\_t](#) for the description of the return status

### 14.3.13. `cusparseSpMatSetStridedBatch()`

```
cusparseStatus_t
cusparseSpMatSetStridedBatch(cusparseSpMatDescr_t spMatDescr,
                             int           batchCount)
```

This function sets the **batchCount** field of the sparse matrix descriptor **spMatDescr**.

Param.	Memory	In/out	Meaning
<b>spMatDescr</b>	HOST	IN	Sparse matrix descriptor
<b>batchCount</b>	HOST	IN	Number of batches of the sparse matrix

See [cusparseStatus\\_t](#) for the description of the return status

## 14.4. Dense Vector APIs

The cuSPARSE helper functions for dense vector descriptor are described in this section.

### 14.4.1. `cusparseCreateDnVec()`

```
cusparseStatus_t
cusparseCreateDnVec(cusparseDnVecDescr_t* dnVecDescr,
                    int64_t          size,
                    void*            values,
                    cudaDataType     valueType)
```

This function initializes the dense vector descriptor **dnVecDescr**.

Param.	Memory	In/out	Meaning
<b>dnVecDescr</b>	HOST	OUT	Dense vector descriptor
<b>size</b>	HOST	IN	Size of the dense vector
<b>values</b>	DEVICE	IN	Values of the dense vector. Array of size <b>size</b>
<b>valueType</b>	HOST	IN	Enumerator specifying the datatype of <b>values</b>

See [cusparseStatus\\_t](#) for the description of the return status

## 14.4.2. `cusparseDestroyDnVec()`

```
cusparseStatus_t
cusparseDestroyDnVec(cusparseDnVecDescr_t dnVecDescr)
```

This function releases the host memory allocated for the dense vector descriptor `dnVecDescr`.

Param.	Memory	In/out	Meaning
<code>dnVecDescr</code>	HOST	IN	Dense vector descriptor

See `cusparseStatus_t` for the description of the return status

## 14.4.3. `cusparseDnVecGet()`

```
cusparseStatus_t
cusparseDnVecGet(const cusparseDnVecDescr_t dnVecDescr,
                 int64_t* size,
                 void** values,
                 cudaDataType* valueType)
```

This function returns the fields of the dense vector descriptor `dnVecDescr`.

Param.	Memory	In/out	Meaning
<code>dnVecDescr</code>	HOST	IN	Dense vector descriptor
<code>size</code>	HOST	OUT	Size of the dense vector
<code>values</code>	DEVICE	OUT	Values of the dense vector. Array of size <code>nnz</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>values</code>

See `cusparseStatus_t` for the description of the return status

## 14.4.4. `cusparseDnVecGetValues()`

```
cusparseStatus_t
cusparseDnVecGetValues(const cusparseDnVecDescr_t dnVecDescr,
                      void** values)
```

This function returns the `values` field of the dense vector descriptor `dnVecDescr`.

Param.	Memory	In/out	Meaning
<code>dnVecDescr</code>	HOST	IN	Dense vector descriptor
<code>values</code>	DEVICE	OUT	Values of the dense vector

See `cusparseStatus_t` for the description of the return status

## 14.4.5. `cusparseDnVecSetValues()`

```
cusparseStatus_t
cusparseDnVecSetValues(cusparseDnVecDescr_t dnVecDescr,
                      void*           values)
```

This function set the **values** field of the dense vector descriptor **dnVecDescr**.

Param.	Memory	In/out	Meaning
<b>dnVecDescr</b>	HOST	IN	Dense vector descriptor
<b>values</b>	DEVICE	IN	Values of the dense vector. Array of size <b>size</b>

The possible error values returned by this function and their meanings are listed below :

See [cusparseStatus\\_t](#) for the description of the return status

## 14.5. Dense Matrix APIs

The cuSPARSE helper functions for dense matrix descriptor are described in this section.

### 14.5.1. `cusparseCreateDnMat()`

```
cusparseStatus_t
cusparseCreateDnMat(cusparseDnMatDescr_t* dnMatDescr,
                    int64_t          rows,
                    int64_t          cols,
                    int64_t          ld,
                    void*           values,
                    cudaDataType    valueType,
                    cusparseOrder_t order)
```

The function initializes the dense matrix descriptor **dnMatDescr**.

Param.	Memory	In/out	Meaning
<b>dnMatDescr</b>	HOST	OUT	Dense matrix descriptor
<b>rows</b>	HOST	IN	Number of rows of the dense matrix
<b>cols</b>	HOST	IN	Number of columns of the dense matrix
<b>ld</b>	HOST	IN	Leading dimension of the dense matrix ( <b>ld</b> $\geq$ <b>rows</b> )
<b>values</b>	DEVICE	IN	Values of the dense matrix. Array of size <b>size</b>
<b>valueType</b>	HOST	IN	Enumerator specifying the datatype of <b>values</b>
<b>order</b>	HOST	IN	Enumerator specifying the memory layout of the dense matrix

See [cusparseStatus\\_t](#) for the description of the return status

## 14.5.2. `cusparseDestroyDnMat()`

```
cusparseStatus_t
cusparseDestroyDnMat(cusparseDnMatDescr_t dnMatDescr)
```

This function releases the host memory allocated for the dense matrix descriptor `dnMatDescr`.

Param.	Memory	In/out	Meaning
<code>dnMatDescr</code>	HOST	IN	Dense matrix descriptor

See `cusparseStatus_t` for the description of the return status

## 14.5.3. `cusparseDnMatGet()`

```
cusparseStatus_t
cusparseDnMatGet(const cusparseDnMatDescr_t dnMatDescr,
                 int64_t* rows,
                 int64_t* cols,
                 int64_t* ld,
                 void** values,
                 cudaDataType* type,
                 cusparseOrder_t* order)
```

This function returns the fields of the dense matrix descriptor `dnMatDescr`.

Param.	Memory	In/out	Meaning
<code>dnMatDescr</code>	HOST	IN	Dense matrix descriptor
<code>rows</code>	HOST	OUT	Number of rows of the dense matrix
<code>cols</code>	HOST	OUT	Number of columns of the dense matrix
<code>ld</code>	HOST	OUT	Leading dimension of the dense matrix ( <code>ld</code> $\geq$ <code>rows</code> )
<code>values</code>	DEVICE	OUT	Values of the dense matrix. Array of size <code>ld</code> * <code>cols</code>
<code>valueType</code>	HOST	OUT	Enumerator specifying the datatype of <code>values</code>
<code>order</code>	HOST	OUT	Enumerator specifying the memory layout of the dense matrix

See `cusparseStatus_t` for the description of the return status

## 14.5.4. `cusparseDnMatGetValues()`

```
cusparseStatus_t CUSPARSEAPI
cusparseDnMatGetValues(const cusparseDnMatDescr_t dnMatDescr,
                      void** values)
```

This function returns the `values` field of the dense matrix descriptor `dnMatDescr`.

Param.	Memory	In/out	Meaning
<code>dnMatDescr</code>	HOST	IN	Dense matrix descriptor

Param.	Memory	In/out	Meaning
<b>values</b>	DEVICE	OUT	Values of the dense matrix. Array of size <code>ld * cols</code>

See `cusparseStatus_t` for the description of the return status

#### 14.5.5. `cusparseDnSetValues()`

```
cusparseStatus_t CUSPARSEAPI
cusparseDnMatSetValues(cusparseDnMatDescr_t dnMatDescr,
                      void* values)
```

This function sets the **values** field of the dense matrix descriptor **dnMatDescr**.

Param.	Memory	In/out	Meaning
<b>dnMatDescr</b>	HOST	IN	Dense matrix descriptor
<b>values</b>	DEVICE	IN	Values of the dense matrix. Array of size <code>ld * cols</code>

See `cusparseStatus_t` for the description of the return status

#### 14.5.6. `cusparseDnMatGetStridedBatch()`

```
cusparseStatus_t
cusparseDnMatGetStridedBatch(const cusparseDnMatDescr_t dnMatDescr,
                             int* batchCount,
                             int64_t* batchStride)
```

The function returns the number of batches and the batch stride of the dense matrix descriptor **dnMatDescr**.

Param.	Memory	In/out	Meaning
<b>dnMatDescr</b>	HOST	IN	Dense matrix descriptor
<b>batchCount</b>	HOST	OUT	Number of batches of the dense matrix
<b>batchStride</b>	HOST	OUT	Address offset between a matrix and the next one in the batch

See `cusparseStatus_t` for the description of the return status

#### 14.5.7. `cusparseDnMatSetStridedBatch()`

```
cusparseStatus_t
cusparseDnMatSetStridedBatch(cusparseDnMatDescr_t dnMatDescr,
                            int batchCount,
                            int64_t batchStride)
```

The function sets the number of batches and the batch stride of the dense matrix descriptor **dnMatDescr**.

Param.	Memory	In/out	Meaning
<b>dnMatDescr</b>	HOST	IN	Dense matrix descriptor

Param.	Memory	In/out	Meaning
<code>batchCount</code>	HOST	IN	Number of batches of the dense matrix
<code>batchStride</code>	HOST	IN	Address offset between a matrix and the next one in the batch $\text{batchStride} \geq \text{ld} * \text{cols}$

See `cusparseStatus_t` for the description of the return status

## 14.6. Generic API Functions

The sparse vector-vector, matrix-vector, matrix-matrix multiplication functions are described in this section.

### 14.6.1. `cusparseSpVV()`

```
cusparseStatus_t
cusparseSpVV_bufferSize(cusparseHandle_t          handle,
                      cusparseOperation_t    opX,
                      const cusparseSpVecDescr_t vecX,
                      const cusparseDnVecDescr_t vecY,
                      void*                  result,
                      cudaDataType           computeType,
                      size_t*                bufferSize)
```

```
cusparseStatus_t
cusparseSpVV(cusparseHandle_t          handle,
            cusparseOperation_t    opX,
            const cusparseSpVecDescr_t vecX,
            const cusparseDnVecDescr_t vecY,
            void*                  result,
            cudaDataType           computeType,
            void*                  externalBuffer)
```

The function performs the dot product of a sparse vector `vecX` and a dense vector `vecY`

$$\text{result} = \sum_{i=1}^{\text{size}} \mathbf{X}[i] \times \text{op}(\mathbf{Y}[i])$$

The function `cusparseSpVV_bufferSize()` returns the size of the workspace needed by `cusparseSpVV()`

The arrays representing the sparse vector `vecX` must be aligned to 16 bytes

Param.	Memory	In/out	Meaning
<code>handle</code>	HOST	IN	Handle to the cuSPARSE library context
<code>opX</code>	HOST	IN	Operation <code>op(x)</code> that is non-transpose or conjugate transpose
<code>vecX</code>	HOST	IN	Sparse vector <code>x</code>
<code>vecY</code>	HOST	IN	Dense vector <code>y</code>
<code>result</code>	HOST or DEVICE	OUT	The resulting dot product

Param.	Memory	In/out	Meaning
<code>computeType</code>	HOST	IN	Enumerator specifying the datatype in which the computation is executed
<code>bufferSize</code>	HOST	OUT	Number of bytes of workspace needed by <code>cusparseSpVV</code>
<code>externalBu</code>	DEVICE	IN	Pointer to workspace buffer

`cusparseSpVV` supports the sparse vector `vecX` with 32-bit indices (`CUSPARSE_INDEX_32I`) and 64-bit indices (`CUSPARSE_INDEX_64I`)

The datatypes combinations currently supported for `cusparseSpVV` are listed below :

x/Y	computeType/result
<code>CUDA_R_16F</code>	<code>CUDA_R_16F</code>
<code>CUDA_R_8I</code>	<code>CUDA_R_32I</code>
<code>CUDA_R_8I</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_16F</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_32F</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_64F</code>	<code>CUDA_R_64F</code>
<code>CUDA_C_16F</code>	<code>CUDA_C_16F</code>
<code>CUDA_C_32F</code>	<code>CUDA_C_32F</code>
<code>CUDA_C_64F</code>	<code>CUDA_C_64F</code>

See `cusparseStatus_t` for the description of the return status

## 14.6.2. `cusparseSpMV()`

```
cusparseStatus_t
cusparseSpMV_bufferSize(cusparseHandle_t          handle,
                      cusparseOperation_t    opA,
                      const void*            alpha,
                      const cusparseSpMatDescr_t matA,
                      const cusparseDnVecDescr_t vecX,
                      const void*            beta,
                      const cusparseDnVecDescr_t vecY,
                      cudaDataType           computeType,
                      cusparseSpMVAAlg_t     alg,
                      size_t*                bufferSize)
```

```
cusparseStatus_t
cusparseSpMV(cusparseHandle_t          handle,
            cusparseOperation_t    opA,
            const void*            alpha,
            const cusparseSpMatDescr_t matA,
            const cusparseDnVecDescr_t vecX,
            const void*            beta,
            const cusparseDnVecDescr_t vecY,
            cudaDataType           computeType,
            cusparseSpMVAAlg_t     alg,
            void*                  externalBuffer)
```

This function performs the multiplication of a sparse matrix **matA** and a dense vector **vecX**

$$\mathbf{Y} = \alpha \mathbf{op}(\mathbf{A}) \cdot \mathbf{X} + \beta \mathbf{Y}$$

where **op (A)** is a sparse matrix with dimensions  $m \times k$ , **x** is a dense vector of size  $k$ , **y** is a dense vector of size  $m$ , and  $\alpha$  and  $\beta$  are scalars. Also, for matrix **A**

$$\mathbf{op}(A) = \begin{cases} A & \text{if } \mathbf{op}(A) == \text{CUSPARSE\_OPERATION\_NON\_TRANSPOSE} \\ A^T & \text{if } \mathbf{op}(A) == \text{CUSPARSE\_OPERATION\_TRANSPOSE} \\ A^H & \text{if } \mathbf{op}(A) == \text{CUSPARSE\_OPERATION\_CONJUGATE\_TRANSPOSE} \end{cases}$$

When using the (conjugate) transpose of the sparse matrix **A**, this routine may produce slightly different results during different runs with the same input parameters.

The function **cusparseSpMV\_bufferSize()** returns the size of the workspace needed by **cusparseSpMV()**

Param.	Memory	In/out	Meaning
<b>handle</b>	HOST	IN	Handle to the cuSPARSE library context
<b>opA</b>	HOST	IN	Operation <b>op (A)</b>
<b>alpha</b>	HOST or DEVICE	IN	$\alpha$ scalar used for multiplication
<b>matA</b>	HOST	IN	Sparse matrix <b>A</b>
<b>vecX</b>	HOST	IN	Dense vector <b>x</b>
<b>beta</b>	HOST or DEVICE	IN	$\beta$ scalar used for multiplication
<b>vecY</b>	HOST	OUT	Dense vector <b>y</b>
<b>computeType</b>	HOST	IN	Enumerator specifying the datatype in which the computation is executed
<b>alg</b>	HOST	IN	Enumerator specifying the algorithm for the computation
<b>bufferSize</b>	HOST	OUT	Number of bytes of workspace needed by <b>cusparseSpMV</b>
<b>externalBu</b>	DEVICE	IN	Pointer to workspace buffer

**cusparseSpMV** currently supports the sparse matrix **matA** with 32-bit indices (**CUSPARSE\_INDEX\_32I**) and 64-bit indices (**CUSPARSE\_INDEX\_64I**).

The datatypes combinations currently supported for **cusparseSpMV** are listed below :

<b>computeType</b>	<b>A/X</b>	<b>Y</b>
<b>CUDA_R_16F</b>	<b>CUDA_R_16F</b>	<b>CUDA_R_16F</b>
<b>CUDA_R_32I</b>	<b>CUDA_R_8I</b>	<b>CUDA_R_32I</b>
<b>CUDA_R_32F</b>	<b>CUDA_R_16F</b>	<b>CUDA_R_16F</b>
<b>CUDA_R_32F</b>	<b>CUDA_R_8I</b>	<b>CUDA_R_32F</b>
<b>CUDA_R_32F</b>	<b>CUDA_R_16F</b>	<b>CUDA_R_32F</b>
<b>CUDA_R_32F</b>	<b>CUDA_R_32F</b>	<b>CUDA_R_32F</b>
<b>CUDA_R_64F</b>	<b>CUDA_R_64F</b>	<b>CUDA_R_64F</b>

<b>computeType</b>	<b>A/X</b>	<b>Y</b>
CUDA_C_16F	CUDA_C_16F	CUDA_C_16F
CUDA_C_32F	CUDA_C_32F	CUDA_C_32F
CUDA_C_64F	CUDA_C_64F	CUDA_C_64F

The sparse matrix formats currently supported are listed below :

<b>Format</b>	<b>Notes</b>
CUSPARSE_FORMAT_COO	May produce slightly different results during different runs with the same input parameters
CUSPARSE_FORMAT_COO_	May produce slightly different results during different runs with the same input parameters
CUSPARSE_FORMAT_CSR	

**cusparseSpMV** routine runs for the following algorithm :

<b>Format</b>	<b>Notes</b>
CUSPARSE_MV_ALG_DEF	Default algorithm for any sparse matrix format
CUSPARSE_COOMV_ALG	Default algorithm for COO sparse matrix format
CUSPARSE_CSRMV_ALG1	Default algorithm for CSR sparse matrix format
CUSPARSE_CSRMV_ALG2	Algorithm 2 for CSR sparse matrix format. May provide better performance for irregular matrices

See **cusparseStatus\_t** for the description of the return status

### 14.6.3. `cusparseSpMM()`

```

cusparseStatus_t
cusparseSpMM_bufferSize(cusparseHandle_t          handle,
                      cusparseOperation_t    opA,
                      cusparseOperation_t    opB,
                      const void*            alpha,
                      const cusparseSpMatDescr_t matA,
                      const cusparseDnMatDescr_t matB,
                      const void*            beta,
                      cusparseDnMatDescr_t   matC,
                      cudaDataType           computeType,
                      cusparseSpMMAlg_t      alg,
                      size_t*                bufferSize)

cusparseStatus_t
cusparseSpMM(cusparseHandle_t          handle,
             cusparseOperation_t    opA,
             cusparseOperation_t    opB,
             const void*            alpha,
             const cusparseSpMatDescr_t matA,
             const cusparseDnMatDescr_t matB,
             const void*            beta,
             cusparseDnMatDescr_t   matC,
             cudaDataType           computeType,
             cusparseSpMMAlg_t      alg,
             void*                  externalBuffer)

```

This function performs the multiplication of a sparse matrix **matA** and a dense matrix **matB**

$$\mathbf{C} = \alpha \text{op}(\mathbf{A}) \cdot \text{op}(\mathbf{B}) + \beta \mathbf{C}$$

where **op (A)** is a sparse matrix with dimensions  $m \times k$ , **op (B)** is a dense matrix of size  $k \times n$ , **C** is a dense matrix of size  $m \times n$ , and  $\alpha$  and  $\beta$  are scalars. Also, for matrix **A** and **B**

$$\text{op}(A) = \begin{cases} A & \text{if } \text{op}(A) == \text{CUSPARSE_OPERATION_NON\_TRANSPOSE} \\ A^T & \text{if } \text{op}(A) == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if } \text{op}(A) == \text{CUSPARSE_OPERATION_CONJUGATE\_TRANSPOSE} \end{cases}$$

$$\text{op}(B) = \begin{cases} B & \text{if } \text{op}(B) == \text{CUSPARSE_OPERATION_NON\_TRANSPOSE} \\ B^T & \text{if } \text{op}(B) == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ B^H & \text{if } \text{op}(B) == \text{CUSPARSE_OPERATION_CONJUGATE\_TRANSPOSE} \end{cases}$$

When using the (conjugate) transpose of the sparse matrix **A**, this routine may produce slightly different results during different runs with the same input parameters.

The function `cusparseSpMM_bufferSize()` returns the size of the workspace needed by `cusparseSpMM()`

Param.	Memory	In/out	Meaning
<b>handle</b>	HOST	IN	Handle to the cuSPARSE library context
<b>opA</b>	HOST	IN	Operation <b>op (A)</b>

Param.	Memory	In/out	Meaning
<code>alpha</code>	HOST or DEVICE	IN	$\alpha$ scalar used for multiplication
<code>matA</code>	HOST	IN	Sparse matrix <b>A</b>
<code>matB</code>	HOST	IN	Dense matrix <b>B</b>
<code>beta</code>	HOST or DEVICE	IN	$\beta$ scalar used for multiplication
<code>matC</code>	HOST	OUT	Dense matrix <b>C</b>
<code>computeType</code>	HOST	IN	Enumerator specifying the datatype in which the computation is executed
<code>alg</code>	HOST	IN	Enumerator specifying the algorithm for the computation
<code>bufferSize</code>	HOST	OUT	Number of bytes of workspace needed by <code>cusparseSpMM</code>
<code>externalBu</code>	DEVICE	IN	Pointer to workspace buffer

`cusparseSpMM` currently supports the sparse matrix `matA` with 32-bit indices (`CUSPARSE_INDEX_32I`).

computeType	A/B	C
<code>CUDA_R_16F</code>	<code>CUDA_R_16F</code>	<code>CUDA_R_16F</code>
<code>CUDA_R_32I</code>	<code>CUDA_R_8I</code>	<code>CUDA_R_32I</code>
<code>CUDA_R_32F</code>	<code>CUDA_R_16F</code>	<code>CUDA_R_16F</code>
<code>CUDA_R_32F</code>	<code>CUDA_R_8I</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_32F</code>	<code>CUDA_R_16F</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_32F</code>	<code>CUDA_R_32F</code>	<code>CUDA_R_32F</code>
<code>CUDA_R_64F</code>	<code>CUDA_R_64F</code>	<code>CUDA_R_64F</code>
<code>CUDA_C_16F</code>	<code>CUDA_C_16F</code>	<code>CUDA_C_16F</code>
<code>CUDA_C_32F</code>	<code>CUDA_C_32F</code>	<code>CUDA_C_32F</code>
<code>CUDA_C_64F</code>	<code>CUDA_C_64F</code>	<code>CUDA_C_64F</code>

The sparse matrix formats currently supported are listed below :

Format	Notes
<code>CUSPARSE_FORMAT_COO</code>	May produce slightly different results during different runs with the same input parameters
<code>CUSPARSE_FORMAT_CSR</code>	

The datatypes combinations currently supported for `cusparseSpMM` are listed below :

`cusparseSpMM` routine runs for the following algorithm :

Format	Notes
<code>CUSPARSE_MM_ALG_DEFAULT</code>	Default algorithm for any sparse matrix format

Format	Notes
CUSPARSE_COOMM_ALG1	Default algorithm for COO sparse matrix format. It supports batched computation. May produce slightly different results during different runs with the same input parameters
CUSPARSE_COOMM_ALG2	Algorithm 2 for COO sparse matrix format. It supports batched computation. It provides deterministic result, and requires additional memory
CUSPARSE_COOMM_ALG3	Algorithm 3 for COO sparse matrix format. May provide better performance for large matrices. May produce slightly different results during different runs with the same input parameters
CUSPARSE_CSRMM_ALG1	Default algorithm for CSR sparse matrix format

See `cusparseStatus_t` for the description of the return status

#### 14.6.4. `cusparseConstrainedGeMM()`

```

cusparseStatus_t
cusparseConstrainedGeMM(cusparseHandle_t           handle,
                        cusparseOperation_t      opA,
                        cusparseOperation_t      opB,
                        const void*              alpha,
                        const cusparseDnMatDescr_t matA,
                        const cusparseDnMatDescr_t matB,
                        const void*              beta,
                        cusparseSpMatDescr_t     matC,
                        cudaDataType             computeType,
                        void*                   externalBuffer)

cusparseStatus_t
cusparseConstrainedGeMM_bufferSize(cusparseHandle_t           handle,
                                    cusparseOperation_t      opA,
                                    cusparseOperation_t      opB,
                                    const void*              alpha,
                                    const cusparseDnMatDescr_t matA,
                                    const cusparseDnMatDescr_t matB,
                                    const void*              beta,
                                    cusparseSpMatDescr_t     matC,
                                    cudaDataType             computeType,
                                    size_t*                  bufferSize)

```

This function performs the multiplication of **matA** and **matB**, followed by an element-wise multiplication with the sparsity pattern of **matC**. Formally, it performs the following operation:

$$\mathbf{C} = \alpha(\text{op}(\mathbf{A}) \cdot \text{op}(\mathbf{B})) \circ \text{spy}(\mathbf{C}) + \beta \mathbf{C}$$

where  $\text{op}(\mathbf{A})$  is a dense matrix of size  $m \times k$ ,  $\text{op}(\mathbf{B})$  is a dense matrix of size  $k \times n$ ,  $\mathbf{C}$  is a sparse matrix of size  $m \times n$ ,  $\alpha$  and  $\beta$  are scalars,  $\circ$  denotes the Hadamard (entry-wise) matrix product, and  $\text{spy}(\mathbf{C})$  is the sparsity pattern matrix of  $\mathbf{C}$  defined as:

$$\text{spy}(\mathbf{C})_{ij} = \begin{cases} 0 & \text{if } \mathbf{C}_{ij} = 0 \\ 1 & \text{otherwise} \end{cases}$$

Matrices  $\text{op}(A)$  and  $\text{op}(B)$  are defined as

$$\text{op}(A) = \begin{cases} A & \text{if } \text{opA} == \text{CUSPARSE_OPERATION_NON_TRANSPOSE} \\ A^T & \text{if } \text{opA} == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ A^H & \text{if } \text{opA} == \text{CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

$$\text{op}(B) = \begin{cases} B & \text{if } \text{opB} == \text{CUSPARSE_OPERATION_NON_TRANSPOSE} \\ B^T & \text{if } \text{opB} == \text{CUSPARSE_OPERATION_TRANSPOSE} \\ B^H & \text{if } \text{opB} == \text{CUSPARSE_OPERATION_CONJUGATE_TRANSPOSE} \end{cases}$$

Param.	Memory	In/out	Meaning
<code>handle</code>	HOST	IN	Handle to the cuSPARSE library context
<code>opA</code>	HOST	IN	Enumerator specifying the operation $\text{op}(A)$ . Has to be <code>CUSPARSE_OPERATION_NON_TRANSPOSE</code>
<code>opB</code>	HOST	IN	Enumerator specifying the operation $\text{op}(B)$ . Has to be <code>CUSPARSE_OPERATION_NON_TRANSPOSE</code>
<code>alpha</code>	HOST or DEVICE	IN	Scalar $\alpha$ that scales the matrix product
<code>matA</code>	HOST	IN	Dense matrix A.
<code>matB</code>	HOST	IN	Dense matrix B.
<code>beta</code>	HOST or DEVICE	IN	Scalar $\beta$ that scales the accumulation matrix
<code>matC</code>	HOST	IN/OUT	Sparse matrix C.
<code>computeType</code>	HOST	IN	Enumerator specifying the datatype used to execute the computation
<code>bufferSize</code>	HOST	OUT	Size of <code>externalBuffer</code> in bytes
<code>externalBu</code>	DEVICE	IN	Pointer to a workspace buffer of at least <code>bufferSize</code> bytes

Currently, this function only supports `opA == CUSPARSE_OPERATION_NON_TRANSPOSE` and `opB == CUSPARSE_OPERATION_NON_TRANSPOSE`. Attempting to pass a different operator will cause a `CUSPARSE_STATUS_NOT_SUPPORTED` error.

Currently supported combinations of data and index types:

computeType	A/B	C (value)	C (index)
<code>CUDA_R_32F</code>	<code>CUDA_R_32F</code>	<code>CUDA_R_32F</code>	<code>CUSPARSE_INDEX_32I</code>
<code>CUDA_R_64F</code>	<code>CUDA_R_64F</code>	<code>CUDA_R_64F</code>	<code>CUSPARSE_INDEX_32I</code>
<code>CUDA_C_32F</code>	<code>CUDA_C_32F</code>	<code>CUDA_C_32F</code>	<code>CUSPARSE_INDEX_32I</code>
<code>CUDA_C_64F</code>	<code>CUDA_C_64F</code>	<code>CUDA_C_64F</code>	<code>CUSPARSE_INDEX_32I</code>

Currently supported sparse matrix formats:

Format	Notes
<code>CUSPARSE_FORMAT_CSR</code>	The column indices in each row must be sorted

See `cusparseStatus_t` for the description of the return status

## 14.7. Example of Generic API

### 14.7.1. CSR sparse matrix - vector multiplication

This section provides a simple example in the C programming language of `cusparseSpMV()`.

The example computes  $A \cdot X = Y$  where **A** is a sparse matrix in CSR format, **x** and **y** dense vectors.

$$\begin{pmatrix} 1 & 0 & 2 & 3 \\ 0 & 4 & 0 & 0 \\ 5 & 0 & 6 & 7 \\ 0 & 8 & 0 & 9 \end{pmatrix}_A \cdot \begin{pmatrix} 1 \\ 2 \\ 3 \\ 4 \end{pmatrix}_X = \begin{pmatrix} 19 \\ 8 \\ 51 \\ 52 \end{pmatrix}_Y$$

```
// *** spmv_example.c ***
// How to compile (assume CUDA is installed at /usr/local/cuda/)
// nvcc spmv_example.c -o spmv_example -L/usr/local/cuda/lib64 -lcusparse -lcudart
// or, for C compiler
// cc -I/usr/local/cuda/include -c spmv_example.c -o spmv_example.o -std=c99
// nvcc -lcusparse -lcudart spmv_example.o -o spmv_example
#include <cuda_runtime.h> // cudaMalloc, cudaMemcpy, etc.
#include <cusparse.h> // cusparseSpMV
#include <stdio.h> // printf
#include <stdlib.h> // EXIT_FAILURE

#define CHECK_CUDA(func)
{
    cudaError_t status = (func);
    if (status != cudaSuccess) {
        printf("CUDA API failed at line %d with error: %s (%d)\n",
               LINE, cudaGetErrorString(status), status);
        return EXIT_FAILURE;
    }
}

#define CHECK_CUSPARSE(func)
{
    cusparseStatus_t status = (func);
    if (status != CUSPARSE_STATUS_SUCCESS) {
        printf("CUSPARSE API failed at line %d with error: %s (%d)\n",
               LINE, cusparseGetErrorString(status), status);
        return EXIT_FAILURE;
    }
}

int main() {
    // Host problem definition
    const int A_num_rows = 4;
    const int A_num_cols = 4;
    const int A_num_nnz = 9;
    int hA_csrOffsets[] = { 0, 3, 4, 7, 9 };
    int hA_columns[] = { 0, 2, 3, 1, 0, 2, 3, 1, 3 };
    float hA_values[] = { 1.0f, 2.0f, 3.0f, 4.0f, 5.0f,
                          6.0f, 7.0f, 8.0f, 9.0f };
    float hX[] = { 1.0f, 2.0f, 3.0f, 4.0f };
    const float result[] = { 19.0f, 8.0f, 51.0f, 52.0f };
    float alpha = 1.0f;
    float beta = 0.0f;
    //-----
    // Device memory management
    int *dA_csrOffsets, *dA_columns;
    float *dA_values, *dX, *dY;
    CHECK_CUDA( cudaMalloc((void**) &dA_csrOffsets,
                           (A_num_rows + 1) * sizeof(int)) )
    CHECK_CUDA( cudaMalloc((void**) &dA_columns, A_num_nnz * sizeof(int)) )
    CHECK_CUDA( cudaMalloc((void**) &dA_values, A_num_nnz * sizeof(float)) )
    CHECK_CUDA( cudaMalloc((void**) &dX, A_num_cols * sizeof(float)) )
    CHECK_CUDA( cudaMalloc((void**) &dY, A_num_rows * sizeof(float)) )

    CHECK_CUDA( cudaMemcpy(dA_csrOffsets, hA_csrOffsets,
                           (A_num_rows + 1) * sizeof(int),
                           cudaMemcpyHostToDevice) )
    CHECK_CUDA( cudaMemcpy(dA_columns, hA_columns, A_num_nnz * sizeof(int),
                           cudaMemcpyHostToDevice) )
    CHECK_CUDA( cudaMemcpy(dA_values, hA_values,
                           A_num_nnz * sizeof(float), cudaMemcpyHostToDevice) )
    CHECK_CUDA( cudaMemcpy(dX, hX, A_num_rows * sizeof(float),
                           cudaMemcpyHostToDevice) )
    //-----
    // CUSPARSE APIs
    cusparseHandle_t handle = 0;
```

## Chapter 15.

# APPENDIX A: CUSPARSE LIBRARY C++ EXAMPLE

For sample code reference please see the example code below. It shows an application written in C++ using the cuSPARSE library API. The code performs the following actions:

1. Creates a sparse test matrix in COO format.
2. Creates a sparse and dense vector.
3. Allocates GPU memory and copies the matrix and vectors into it.
4. Initializes the cuSPARSE library.
5. Creates and sets up the matrix descriptor.
6. Converts the matrix from COO to CSR format.
7. Exercises Level 1 routines.
8. Exercises Level 2 routines.
9. Exercises Level 3 routines.
10. Destroys the matrix descriptor.

## 11. Releases resources allocated for the cuSPARSE library.

```
//Example: Application using C++ and the CUSPARSE library
//-----
#include <stdio.h>
#include <stdlib.h>
#include <cuda_runtime.h>
#include "cusparse.h"

#define CLEANUP(s)
do {
    printf ("%s\n", s);
    if (yHostPtr)          free(yHostPtr);
    if (zHostPtr)          free(zHostPtr);
    if (xIndHostPtr)       free(xIndHostPtr);
    if (xValHostPtr)       free(xValHostPtr);
    if (cooRowIndexHostPtr) free(cooRowIndexHostPtr);
    if (cooColIndexHostPtr) free(cooColIndexHostPtr);
    if (cooValHostPtr)     free(cooValHostPtr);
    if (y)                 cudaFree(y);
    if (z)                 cudaFree(z);
    if (xInd)              cudaFree(xInd);
    if (xVal)              cudaFree(xVal);
    if (csrRowPtr)         cudaFree(csrRowPtr);
    if (cooRowIndex)        cudaFree(cooRowIndex);
    if (cooColIndex)        cudaFree(cooColIndex);
    if (cooVal)             cudaFree(cooVal);
    if (descr)             cusparseDestroyMatDescr(descr);
    if (handle)            cusparseDestroy(handle);
    cudaDeviceReset();
    fflush (stdout);
} while (0)

int main() {
    cudaError_t cudaStat1, cudaStat2, cudaStat3, cudaStat4, cudaStat5, cudaStat6;
    cusparseStatus_t status;
    cusparseHandle_t handle=0;
    cusparseMatDescr_t descr=0;
    int * cooRowIndexHostPtr=0;
    int * cooColIndexHostPtr=0;
    double * cooValHostPtr=0;
    int * cooRowIndex=0;
    int * cooColIndex=0;
    double * cooVal=0;
    int * xIndHostPtr=0;
    double * xValHostPtr=0;
    double * yHostPtr=0;
    int * xInd=0;
    double * xVal=0;
    double * y=0;
    int * csrRowPtr=0;
    double * zHostPtr=0;
    double * z=0;
    int n, nnz, nnz_vector;
    double dzero =0.0;
    double dtwo =2.0;
    double dthree=3.0;
    double dfive =5.0;

    printf("testing example\n");
    /* create the following sparse test matrix in COO format */
    /* |1.0 2.0 3.0|
     |   4.0 |
     |5.0 6.0 7.0|
     |   8.0 9.0| */
    n=4; nnz=9;
    cooRowIndexHostPtr = (int *) malloc(nnz*sizeof(cooRowIndexHostPtr[0]));
    cooColIndexHostPtr = (int *) malloc(nnz*sizeof(cooColIndexHostPtr[0]));
    cooValHostPtr      = (double *)malloc(nnz*sizeof(cooValHostPtr[0]));
    if ((!cooRowIndexHostPtr) || (!cooColIndexHostPtr) || (!cooValHostPtr)) {
        CLEANUP("Host malloc failed (matrix)");
        return 1;
    }
    cooRowIndexHostPtr[0]=0; cooColIndexHostPtr[0]=0; cooValHostPtr[0]=1.0;
    cooRowIndexHostPtr[1]=0; cooColIndexHostPtr[1]=2; cooValHostPtr[1]=2.0;
    cooRowIndexHostPtr[2]=0; cooColIndexHostPtr[2]=3; cooValHostPtr[2]=3.0;
    cooRowIndexHostPtr[3]=0; cooColIndexHostPtr[3]=1; cooValHostPtr[3]=4.0;
    cooRowIndexHostPtr[4]=1; cooColIndexHostPtr[4]=0; cooValHostPtr[4]=5.0;
    cooRowIndexHostPtr[5]=1; cooColIndexHostPtr[5]=2; cooValHostPtr[5]=6.0;
    cooRowIndexHostPtr[6]=2; cooColIndexHostPtr[6]=0; cooValHostPtr[6]=7.0;
    cooRowIndexHostPtr[7]=2; cooColIndexHostPtr[7]=1; cooValHostPtr[7]=8.0;
    cooRowIndexHostPtr[8]=2; cooColIndexHostPtr[8]=2; cooValHostPtr[8]=9.0;
}
```

# Chapter 16.

## APPENDIX B: CUSPARSE FORTRAN BINDINGS

The cuSPARSE library is implemented using the C-based CUDA toolchain, and it thus provides a C-style API that makes interfacing to applications written in C or C++ trivial. There are also many applications implemented in Fortran that would benefit from using cuSPARSE, and therefore a cuSPARSE Fortran interface has been developed.

Unfortunately, Fortran-to-C calling conventions are not standardized and differ by platform and toolchain. In particular, differences may exist in the following areas:

Symbol names (capitalization, name decoration)

Argument passing (by value or reference)

Passing of pointer arguments (size of the pointer)

To provide maximum flexibility in addressing those differences, the cuSPARSE Fortran interface is provided in the form of wrapper functions, which are written in C and are located in the file `cusparse_fortran.c`. This file also contains a few additional wrapper functions (for `cudaMalloc()`, `cudaMemset`, and so on) that can be used to allocate memory on the GPU.

The cuSPARSE Fortran wrapper code is provided as an example only and needs to be compiled into an application for it to call the cuSPARSE API functions. Providing this source code allows users to make any changes necessary for a particular platform and toolchain.

The cuSPARSE Fortran wrapper code has been used to demonstrate interoperability with the compilers g95 0.91 (on 32-bit and 64-bit Linux) and g95 0.92 (on 32-bit and 64-bit Mac OS X). In order to use other compilers, users have to make any changes to the wrapper code that may be required.

The direct wrappers, intended for production code, substitute device pointers for vector and matrix arguments in all cuSPARSE functions. To use these interfaces, existing applications need to be modified slightly to allocate and deallocate data structures in GPU memory space (using `CUDA_MALLOC()` and `CUDA_FREE()`) and to copy data between GPU and CPU memory spaces (using the `CUDA_MEMCPY()` routines). The sample wrappers provided in `cusparse_fortran.c` map device pointers to the OS-

dependent type **size\_t**, which is 32 bits wide on 32-bit platforms and 64 bits wide on a 64-bit platforms.

One approach to dealing with index arithmetic on device pointers in Fortran code is to use C-style macros and to use the C preprocessor to expand them. On Linux and Mac OS X, preprocessing can be done by using the option '**-cpp**' with g95 or gfortran. The function **GET\_SHIFTED\_ADDRESS()**, provided with the cuSPARSE Fortran wrappers, can also be used, as shown in example B.

Example B shows the the C++ of example A implemented in Fortran 77 on the host. This example should be compiled with **ARCH\_64** defined as 1 on a 64-bit OS system and as undefined on a 32-bit OS system. For example, on g95 or gfortran, it can be done directly on the command line using the option **-cpp -DARCH\_64=1**.

## 16.1. Example B, Fortran Application

```

c      #define ARCH_64 0
c      #define ARCH_64 1

program cusparse_fortran_example
implicit none
integer cuda_malloc
external cuda_free
integer cuda_memcpy_c2fort_int
integer cuda_memcpy_c2fort_real
integer cuda_memcpy_fort2c_int
integer cuda_memcpy_fort2c_real
integer cuda_memset
integer cusparse_create
external cusparse_destroy
integer cusparse_get_version
integer cusparse_create_mat_descr
external cusparse_destroy_mat_descr
integer cusparse_set_mat_type
integer cusparse_get_mat_type
integer cusparse_get_mat_fill_mode
integer cusparse_get_mat_diag_type
integer cusparse_set_mat_index_base
integer cusparse_get_mat_index_base
integer cusparse_xcoo2csr
integer cusparse_dsctr
integer cusparse_dcsrmv
integer cusparse_dcsrmm
external get_shifted_address
#endif ARCH_64
integer*8 handle
integer*8 descrA
integer*8 cooRowIndex
integer*8 cooColIndex
integer*8 cooVal
integer*8 xInd
integer*8 xVal
integer*8 y
integer*8 z
integer*8 csrRowPtr
integer*8 ynp1
#else
integer*4 handle
integer*4 descrA
integer*4 cooRowIndex
integer*4 cooColIndex
integer*4 cooVal
integer*4 xInd
integer*4 xVal
integer*4 y
integer*4 z
integer*4 csrRowPtr
integer*4 ynp1
#endif
integer status
integer cudaStat1, cudaStat2, cudaStat3
integer cudaStat4, cudaStat5, cudaStat6
integer n, nnz, nnz_vector
parameter (n=4, nnz=9, nnz_vector=3)
integer cooRowIndexHostPtr(nnz)
integer cooColIndexHostPtr(nnz)
real*8 cooValHostPtr(nnz)
integer xIndHostPtr(nnz_vector)
real*8 xValHostPtr(nnz_vector)
real*8 yHostPtr(2*n)
real*8 zHostPtr(2*(n+1))
integer i, j
integer version, mtype, fmode, dtype, ibase
real*8 dzero, dtwo, dthree, dfive
real*8 epsilon

```

# Chapter 17.

## APPENDIX C: EXAMPLES OF SORTING

### 17.1. COO Sort

This chapter provides a simple example in the C programming language of sorting of COO format.

A is a 3x3 sparse matrix,

$$A = \begin{pmatrix} 1.0 & 2.0 & 0.0 \\ 0.0 & 5.0 & 0.0 \\ 0.0 & 8.0 & 0.0 \end{pmatrix}$$

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include coosort.cpp
 * g++ -o coosort.cpp coosort.o -L/usr/local/cuda/lib64 -lcusparse -lcudart
 *
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;

    cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
    cudaError_t cudaStat1 = cudaSuccess;
    cudaError_t cudaStat2 = cudaSuccess;
    cudaError_t cudaStat3 = cudaSuccess;
    cudaError_t cudaStat4 = cudaSuccess;
    cudaError_t cudaStat5 = cudaSuccess;
    cudaError_t cudaStat6 = cudaSuccess;

    /*
     * A is a 3x3 sparse matrix
     *   | 1 2 0 |
     * A = | 0 5 0 |
     *   | 0 8 0 |
     */
    const int m = 3;
    const int n = 3;
    const int nnz = 4;

#if 0
/* index starts at 0 */
    int h_cooRows[nnz] = {2, 1, 0, 0 };
    int h_cooCols[nnz] = {1, 1, 0, 1 };
#else
/* index starts at -2 */
    int h_cooRows[nnz] = {0, -1, -2, -2 };
    int h_cooCols[nnz] = {-1, -1, -2, -1 };
#endif
    double h_cooVals[nnz] = {8.0, 5.0, 1.0, 2.0 };
    int h_P[nnz];

    int *d_cooRows = NULL;
    int *d_cooCols = NULL;
    int *d_P = NULL;
    double *d_cooVals = NULL;
    double *d_cooVals_sorted = NULL;
    size_t pBufferSizeInBytes = 0;
    void *pBuffer = NULL;

    printf("m = %d, n = %d, nnz=%d \n", m, n, nnz );

```

```
/* step 1: create cusparse handle, bind a stream */
cudaStat1 = cudaStreamCreateWithFlags(&stream, cudaStreamNonBlocking);
assert(cudaSuccess == cudaStat1);

status = cusparseCreate(&handle);
assert(CUSPARSE_STATUS_SUCCESS == status);

status = cusparseSetStream(handle, stream);
assert(CUSPARSE_STATUS_SUCCESS == status);

/* step 2: allocate buffer */

```

# Chapter 18.

## APPENDIX D: EXAMPLES OF PRUNE

### 18.1. Prune Dense to Sparse

This section provides a simple example in the C programming language of pruning a dense matrix to a sparse matrix of CSR format.

A is a 4x4 dense matrix,

$$A = \begin{pmatrix} 1.0 & 0.0 & 2.0 & -3.0 \\ 0.0 & 4.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 6.0 & 7.0 \\ 0.0 & 8.0 & 0.0 & 9.0 \end{pmatrix}$$

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include prunedense_example.cpp
 * g++ -o prunedense_example.cpp prunedense_example.o -L/usr/local/cuda/lib64
 -lcusparse -lcudart
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

void printMatrix(int m, int n, const float*A, int lda, const char* name)
{
    for(int row = 0 ; row < m ; row++){
        for(int col = 0 ; col < n ; col++){
            float Areg = A[row + col*lda];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

void printCsr(
    int m,
    int n,
    int nnz,
    const cusparseMatDescr_t descrA,
    const float *csrValA,
    const int *csrRowPtrA,
    const int *csrColIndA,
    const char* name)
{
    const int base = (cusparseGetMatIndexBase(descrA) != CUSPARSE_INDEX_BASE_ONE)? 0:1;

    printf("matrix %s is %d-by-%d, nnz=%d, base=%d\n", name, m, n, nnz, base);
    for(int row = 0 ; row < m ; row++){
        const int start = csrRowPtrA[row] - base;
        const int end = csrRowPtrA[row+1] - base;
        for(int colidx = start ; colidx < end ; colidx++){
            const int col = csrColIndA[colidx] - base;
            const float Areg = csrValA[colidx];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;
    cusparseMatDescr_t descrC = NULL;
```

```
cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
cudaError_t cudaStat1 = cudaSuccess;
cudaError_t cudaStat2 = cudaSuccess;
cudaError_t cudaStat3 = cudaSuccess;
cudaError_t cudaStat4 = cudaSuccess;
cudaError_t cudaStat5 = cudaSuccess;
const int m = 4;
const int n = 4;
const int lda = m;
/*
 *      | 1 0 2 -3 |
 *      | 0 4 0 0 |
 *      | 5 0 6 7 |
 *      | 0 8 0 9 |
 */
```

## 18.2. Prune Sparse to Sparse

This section provides a simple example in the C programming language of pruning a sparse matrix to a sparse matrix of CSR format.

A is a 4x4 sparse matrix,

$$A = \begin{pmatrix} 1.0 & 0.0 & 2.0 & -3.0 \\ 0.0 & 4.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 6.0 & 7.0 \\ 0.0 & 8.0 & 0.0 & 9.0 \end{pmatrix}$$

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include prunecsr_example.cpp
 * g++ -o prunecsr_example.cpp prunecsr_example.o -L/usr/local/cuda/lib64 -
lcusparse -lcudart
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

void printCsr(
    int m,
    int n,
    int nnz,
    const cusparseMatDescr_t descrA,
    const float *csrValA,
    const int *csrRowPtrA,
    const int *csrColIndA,
    const char* name)
{
    const int base = (cusparseGetMatIndexBase(descrA) != CUSPARSE_INDEX_BASE_ONE) ? 0:1;

    printf("matrix %s is %d-by-%d, nnz=%d, base=%d, output base-1\n", name, m,
n, nnz, base);
    for(int row = 0 ; row < m ; row++){
        const int start = csrRowPtrA[row] - base;
        const int end   = csrRowPtrA[row+1] - base;
        for(int colidx = start ; colidx < end ; colidx++){
            const int col = csrColIndA[colidx] - base;
            const float Areg = csrValA[colidx];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;
    cusparseMatDescr_t descrA = NULL;
    cusparseMatDescr_t descrC = NULL;

    cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
    cudaError_t cudaStat1 = cudaSuccess;
    const int m = 4;
    const int n = 4;
    const int nnzA = 9;
/*
 *      | 1 0 2 -3 |
 *      | 0 4 0 0 |
 * A = | 5 0 6 7 |
 *      | 0 8 0 9 |
 */
}
```

```
const int csrRowPtrA[m+1] = { 1, 4, 5, 8, 10};
const int csrColIndA[nnzA] = { 1, 3, 4, 2, 1, 3, 4, 2, 4};
const float csrValA[nnzA] = {1, 2, -3, 4, 5, 6, 7, 8, 9};

int* csrRowPtrC = NULL;
int* csrColIndC = NULL;
float* csrValC = NULL;

int *d_csrRowPtrA = NULL;
```

### 18.3. Prune Dense to Sparse by Percentage

This section provides a simple example in the C programming language of pruning a dense matrix to a sparse matrix by percentage.

A is a 4x4 dense matrix,

$$A = \begin{pmatrix} 1.0 & 0.0 & 2.0 & -3.0 \\ 0.0 & 4.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 6.0 & 7.0 \\ 0.0 & 8.0 & 0.0 & 9.0 \end{pmatrix}$$

The percentage is 50, which means to prune 50 percent of the dense matrix. The matrix has 16 elements, so 8 out of 16 must be pruned out. Therefore 7 zeros are pruned out, and value 1.0 is also out because it is the smallest among 9 nonzero elements.

```

/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include prunedense2csrbyP.cpp
 * g++ -o prunedense2csrbyP.cpp prunedense2csrbyP.o -L/usr/local/cuda/lib64 -
lcusparse -lcudart
*/
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

void printMatrix(int m, int n, const float*A, int lda, const char* name)
{
    for(int row = 0 ; row < m ; row++){
        for(int col = 0 ; col < n ; col++){
            float Areg = A[row + col*lda];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

void printCsr(
    int m,
    int n,
    int nnz,
    const cusparseMatDescr_t descrA,
    const float *csrValA,
    const int *csrRowPtrA,
    const int *csrColIndA,
    const char* name)
{
    const int base = (cusparseGetMatIndexBase(descrA) != CUSPARSE_INDEX_BASE_ONE) ? 0:1 ;

    printf("matrix %s is %d-by-%d, nnz=%d, base=%d, output base-1\n", name, m,
n, nnz, base);
    for(int row = 0 ; row < m ; row++){
        const int start = csrRowPtrA[row] - base;
        const int end   = csrRowPtrA[row+1] - base;
        for(int colidx = start ; colidx < end ; colidx++){
            const int col = csrColIndA[colidx] - base;
            const float Areg = csrValA[colidx];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;
    cusparseMatDescr_t descrC = NULL;
    pruneInfo_t info = NULL;

    cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
    cudaError_t cudaStat1 = cudaSuccess;
    cudaError_t cudaStat2 = cudaSuccess;
    cudaError_t cudaStat3 = cudaSuccess;
    cudaError_t cudaStat4 = cudaSuccess;
    cudaError_t cudaStat5 = cudaSuccess;
    const int m = 4;
    const int n = 4;
    const int lda = m;

    /*          | 1 0 2 -3 |
 *          | 0 4 0 0 |

```

## 18.4. Prune Sparse to Sparse by Percentage

This section provides a simple example in the C programming language of pruning a sparse matrix to a sparse matrix by percentage.

A is a 4x4 sparse matrix,

$$A = \begin{pmatrix} 1.0 & 0.0 & 2.0 & -3.0 \\ 0.0 & 4.0 & 0.0 & 0.0 \\ 5.0 & 0.0 & 6.0 & 7.0 \\ 0.0 & 8.0 & 0.0 & 9.0 \end{pmatrix}$$

The percentage is 20, which means to prune 20 percent of the nonzeros. The sparse matrix has 9 nonzero elements, so 1.4 elements must be pruned out. The function removes 1.0 and 2.0 which are first two smallest numbers of nonzeros.

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include prunecsr2csrByP.cpp
 * g++ -o prunecsr2csrByP.cpp prunecsr2csrByP.o -L/usr/local/cuda/lib64 -
lcusparse -lcudart
*/
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

void printCsr(
    int m,
    int n,
    int nnz,
    const cusparseMatDescr_t descrA,
    const float *csrValA,
    const int *csrRowPtrA,
    const int *csrColIndA,
    const char* name)
{
    const int base = (cusparseGetMatIndexBase(descrA) != CUSPARSE_INDEX_BASE_ONE) ? 0:1 ;

    printf("matrix %s is %d-by-%d, nnz=%d, base=%d, output base-1\n", name, m,
n, nnz, base);
    for(int row = 0 ; row < m ; row++){
        const int start = csrRowPtrA[row] - base;
        const int end   = csrRowPtrA[row+1] - base;
        for(int colidx = start ; colidx < end ; colidx++){
            const int col = csrColIndA[colidx] - base;
            const float Areg = csrValA[colidx];
            printf("%s(%d,%d) = %f\n", name, row+1, col+1, Areg);
        }
    }
}

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;
    cusparseMatDescr_t descrA = NULL;
    cusparseMatDescr_t descrC = NULL;
    pruneInfo_t info = NULL;

    cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
    cudaError_t cudaStat1 = cudaSuccess;
    const int m = 4;
    const int n = 4;
    const int nnzA = 9;
/*
 *      | 1 0 2 -3 |
 *      | 0 4 0 0   |
 * A = | 5 0 6 7   |
 *      | 0 8 0 9   |
 */
}
```

```
const int csrRowPtrA[m+1] = { 1, 4, 5, 8, 10};
const int csrColIndA[nnzA] = { 1, 3, 4, 2, 1, 3, 4, 2, 4};
const float csrValA[nnzA] = {1, 2, -3, 4, 5, 6, 7, 8, 9};

int* csrRowPtrC = NULL;
int* csrColIndC = NULL;
float* csrValC = NULL;
```

# Chapter 19.

## APPENDIX E: EXAMPLES OF GTSV

### 19.1. Batched Tridiagonal Solver

This section provides a simple example in the C programming language of `gtsvInterleavedBatch`.

The example solves two linear systems and assumes data layout is NOT interleaved format. Before calling `gtsvInterleavedBatch`, `cublasXgemm` is used to transform the data

layout, from aggregate format to interleaved format. If the user can prepare interleaved format, no need to transpose the data.

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include gtsv.cpp
 * g++ -o gtsv gtsv.o -L/usr/local/cuda/lib64 -lcusparse -lcublas -lcudart
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>
#include <cublas_v2.h>

/*
 * compute | b - A*x|_inf
 */
void residual_eval(
    int n,
    const float *dl,
    const float *d,
    const float *du,
    const float *b,
    const float *x,
    float *r_norminf_ptr)
{
    float r_norminf = 0;
    for(int i = 0 ; i < n ; i++){
        float dot = 0;
        if (i > 0 ){
            dot += dl[i]*x[i-1];
        }
        dot += d[i]*x[i];
        if (i < (n-1) ){
            dot += du[i]*x[i+1];
        }
        float ri = b[i] - dot;
        r_norminf = (r_norminf > fabs(ri)) ? r_norminf : fabs(ri);
    }

    *r_norminf_ptr = r_norminf;
}

int main(int argc, char*argv[])
{
    cusparseHandle_t cusparseH = NULL;
    cublasHandle_t cublasH = NULL;
    cudaStream_t stream = NULL;
```

```
cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
cublasStatus_t cublasStat = CUBLAS_STATUS_SUCCESS;
cudaError_t cudaStat1 = cudaSuccess;

const int n = 3;
const int batchSize = 2;
/*
 *      |   1   6   0   |   |   1   |   | -0.603960 |
 * A1 = |   4   2   7   |, b1 = |   2   |, x1 = |  0.267327 |
 *      |   0   5   3   |   |   3   |   |  0.554455 |
 *
 *      |   8   13   0   |   |   4   |   | -0.063291 |
 * A2 = |  11   9   14   |, b2 = |   5   |, x2 = |  0.346641 |
 *      |   0   12   10  |   |   6   |   |  0.184031 |
 */

/*
 * A = (dl, d, du), B and X are in aggregate format
 */
    const float dl[n * batchSize] = {0, 4, 5, 0, 11, 12};
    const float d[batchSize] = {1, 2, 3, 2, 0, 10};
```

# Chapter 20.

## APPENDIX F: EXAMPLES OF GPSV

### 20.1. Batched Penta-diagonal Solver

This section provides a simple example in the C programming language of `gpsvInterleavedBatch`.

The example solves two penta-diagonal systems and assumes data layout is NOT interleaved format. Before calling `gpsvInterleavedBatch`, `cublasXgemm` is used to

transform the data layout, from aggregate format to interleaved format. If the user can prepare interleaved format, no need to transpose the data.

```

/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 * nvcc -c -I/usr/local/cuda/include gpsv.cpp
 * g++ -o gpsv gpsv.o -L/usr/local/cuda/lib64 -lcusparse -lcublas -lcudart
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>
#include <cublas_v2.h>

/*
 * compute | b - A*x|_inf
 */
void residual_eval(
    int n,
    const float *ds,
    const float *dl,
    const float *d,
    const float *du,
    const float *dw,
    const float *b,
    const float *x,
    float *r_nrminf_ptr)
{
    float r_nrminf = 0;
    for(int i = 0 ; i < n ; i++){
        float dot = 0;
        if (i > 1 ){
            dot += ds[i]*x[i-2];
        }
        if (i > 0 ){
            dot += dl[i]*x[i-1];
        }
        dot += d[i]*x[i];
        if (i < (n-1) ){
            dot += du[i]*x[i+1];
        }
        if (i < (n-2) ){
            dot += dw[i]*x[i+2];
        }
        float ri = b[i] - dot;
        r_nrminf = (r_nrminf > fabs(ri))? r_nrminf : fabs(ri);
    }
    *r_nrminf_ptr = r_nrminf;
}

int main(int argc, char*argv[])
{
    cusparseHandle_t cusparseH = NULL;
    cublasHandle_t cublasH = NULL;
    cudaStream_t stream = NULL;
}

```

```

cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
cublasStatus_t cublasStat = CUBLAS_STATUS_SUCCESS;
cudaError_t cudaStat1 = cudaSuccess;

const int n = 4;
const int batchSize = 2;

/*
 *      | 1   8   13   0   |   | 1   |   | -0.0592 |
 * A1 = | 5   2   9   14 |, b1 = | 2   |, x1 = | 0.3428 |
 *      | 11   6   3   10  |   | 3   |   | -0.1295 |
 *      | 0   12   7   4   |   | 4   |   | 0.1982 |
 */

```

# Chapter 21.

## APPENDIX G: EXAMPLES OF CSRSM2

### 21.1. Forward Triangular Solver

This section provides a simple example in the C programming language of csrsm2.

The example solves a lower triangular system with 2 right hand side vectors.

```
/*
 * How to compile (assume cuda is installed at /usr/local/cuda/)
 *   nvcc -c -I/usr/local/cuda/include csrsm2.cpp
 *   g++ -o csrsm2 csrsm2.o -L/usr/local/cuda/lib64 -lcusparse -lcudart
 */
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>
#include <cuda_runtime.h>
#include <cusparse.h>

/* compute | b - A*x|_inf */
void residual_eval(
    int n,
    const cusparseMatDescr_t descrA,
    const float *csrVal,
    const int *csrRowPtr,
    const int *csrColInd,
    const float *b,
    const float *x,
    float *r_nrminf_ptr)
{
    const int base = (cusparseGetMatIndexBase(descrA) != CUSPARSE_INDEX_BASE_ONE)? 0:1;
    const int lower = (CUSPARSE_FILL_MODE_LOWER == cusparseGetMatFillMode(descrA))? 1:0;
    const int unit = (CUSPARSE_DIAG_TYPE_UNIT == cusparseGetMatDiagType(descrA))? 1:0;

    float r_nrminf = 0;
    for(int row = 0 ; row < n ; row++){
        const int start = csrRowPtr[row] - base;
        const int end   = csrRowPtr[row+1] - base;
        float dot = 0;
        for(int colidx = start ; colidx < end; colidx++){
            const int col = csrColInd[colidx] - base;
            float Aij = csrVal[colidx];
            float xj  = x[col];
            if( (row == col) && unit ){
                Aij = 1.0;
            }
            int valid = (row >= col) && lower ||
                        (row <= col) && !lower ;
            if( valid ){
                dot += Aij*xj;
            }
        }
        float ri = b[row] - dot;
        r_nrminf = (r_nrminf > fabs(ri))? r_nrminf : fabs(ri);
    }
    *r_nrminf_ptr = r_nrminf;
}

int main(int argc, char*argv[])
{
    cusparseHandle_t handle = NULL;
    cudaStream_t stream = NULL;
    cusparseMatDescr_t descrA = NULL;
    csrsm2Info_t info = NULL;
```

```
cusparseStatus_t status = CUSPARSE_STATUS_SUCCESS;
cudaError_t cudaStat1 = cudaSuccess;
const int nrhs = 2;
const int n = 4;
const int nnzA = 9;
const cusparseSolvePolicy_t policy = CUSPARSE_SOLVE_POLICY_NO_LEVEL;
const float h_one = 1.0;

/*      | 1 0 2 -3 |
*      | 0 4 0 0 |
```

# Chapter 22.

## APPENDIX H: ACKNOWLEDGEMENTS

NVIDIA would like to thank the following individuals and institutions for their contributions:

- ▶ The cusparse<t>gtsv implementation is derived from a version developed by Li-Wen Chang from the University of Illinois.
- ▶ the cusparse<t>gtsvInterleavedBatch adopts cuThomasBatch developed by Pedro Valero-Lara and Ivan Martínez-Pérez from Barcelona Supercomputing Center and BSC/UPC NVIDIA GPU Center of Excellence.

# Chapter 23.

# BIBLIOGRAPHY

- [1] N. Bell and M. Garland, “[Implementing Sparse Matrix-Vector Multiplication on Throughput-Oriented Processors](#)”, Supercomputing, 2009.
- [2] R. Grimes, D. Kincaid, and D. Young, “[ITPACK 2.0 User’s Guide](#)”, Technical Report CNA-150, Center for Numerical Analysis, University of Texas, 1979.
- [3] M. Naumov, “[Incomplete-LU and Cholesky Preconditioned Iterative Methods Using cuSPARSE and cuBLAS](#)”, Technical Report and White Paper, 2011.
- [4] Pedro Valero-Lara, Ivan Martínez-Pérez, Raül Sirvent, Xavier Martorell, and Antonio J. Peña. NVIDIA GPUs Scalability to Solve Multiple (Batch) Tridiagonal Systems. Implementation of cuThomasBatch. In Parallel Processing and Applied Mathematics - 12th International Conference (PPAM), 2017.

### **Notice**

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

### **Trademarks**

NVIDIA and the NVIDIA logo are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

### **Copyright**

© 2007-2019 NVIDIA Corporation. All rights reserved.