

---

# LAGraph

**Tim Davis, Tim Mattson, Scott McMillan, Jim Kitchen, Erik Welch**

**Sep 20, 2022**



## CONTENTS:

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>LAGraph context and error handling</b>	<b>5</b>
2.1	LAGraph Context Functions . . . . .	5
2.2	Error handling . . . . .	7
<b>3</b>	<b>The Graph Object</b>	<b>13</b>
3.1	Basic Graph Functions . . . . .	15
<b>4</b>	<b>Algorithms</b>	<b>21</b>
4.1	Basic . . . . .	21
4.2	Advanced . . . . .	22
<b>5</b>	<b>Utility Functions</b>	<b>29</b>
5.1	Input/Output Functions . . . . .	29
5.2	Matrix Structure Functions . . . . .	32
5.3	Matrix Comparison Functions . . . . .	33
5.4	Introspecting Types . . . . .	35
5.5	Printing . . . . .	36
5.6	Pre-defined semirings . . . . .	38
<b>6</b>	<b>Experimental Algorithms</b>	<b>41</b>
<b>7</b>	<b>Installation</b>	<b>43</b>
<b>8</b>	<b>Acknowledgements</b>	<b>45</b>
<b>9</b>	<b>References</b>	<b>47</b>
<b>10</b>	<b>Example Usage</b>	<b>49</b>
	<b>Index</b>	<b>51</b>



The LAGraph library is a collection of high level graph algorithms based on the GraphBLAS C API. These algorithms construct graph algorithms expressed *in the language of linear algebra*. Graphs are expressed as matrices, and the operations over these matrices are generalized through the use of a semiring algebraic structure.

LAGraph is available at <https://github.com/GraphBLAS/LAGraph>. LAGraph requires SuiteSparse:GraphBLAS, available at <https://github.com/DrTimothyAldenDavis/GraphBLAS>.



## INTRODUCTION

A graph is a set of vertices and a set of edges between them. This pair of sets leads directly to the familiar picture of a graph as a set of dots connected by arcs (an undirected graphs) or arrows (a directed graph). You can also represent a graph in terms of matrices. Usually, this is done with an adjacency matrix where the rows and columns correspond to the vertices and the non-empty elements represent the edges between vertices. Since fully connected graphs (i.e., every vertex is connected to every other vertex) are rare, matrices used for Graphs are typically sparse (most elements are empty so it makes sense to only store the non-empty elements).

Representing a graph as a sparse matrix results in graph algorithms expressed in terms of linear algebra. For example, if a vector represents a set of vertices, multiplication of that vector by an adjacency matrix returns a vector of the neighbors of those vertices. A sequence of multiplications traverses the graph in a breadth first manner.

To cover the full range of graph algorithms, one additional ingredient is needed. We are used to thinking of matrix operations over real numbers: multiply pairs of matrix elements and then combine the resulting products through addition. There are times, however, when those operations do not provide the functionality needed by an algorithm. For example, it may be better to combine elements by only keeping the minimum value. The elements of the matrices may be Boolean values or integers or even a user-defined type. If the goal is to cover the full range of graph algorithms, therefore, we need a way to generalize the type and the operators to use instead of the usual addition and multiplication.

We do this through an algebraic semiring. This algebraic structure consists of (1) an operator corresponding to addition, (2) the identity of that operator, (3) an operator corresponding to multiplication, and (4) the identity of that operator. We are all familiar with the semiring used with real numbers consisting of  $(+,0,*,1)$ . A common semiring in graph algorithms is the so-called tropical semiring consisting of  $(\min, \text{infinity}, +, 0)$ . This is used in shortest path algorithms. These semirings give us a mathematically rigorous way to modify the operators used in our graph algorithms.

If you work with linear algebra, you most likely know about the Basic Linear Algebra subprograms or BLAS. Introduced in the 70's and 80's, the BLAS had a huge impact on the practice of linear algebra. By designing linear algebra in terms of the BLAS, an algorithm can be expressed at a high level leaving specialization to the low level details of a particular hardware platform to the BLAS. So if you want to use Linear Algebra for Graph Algorithms, it stands to reason that you need the Basic Linear Algebra Subprograms for Graph Algorithms. We call these the GraphBLAS ([www.graphblas.org](http://www.graphblas.org)).

The GraphBLAS define opaque types for a matrix and a vector objects. Since these objects are opaque, an implementation has the freedom to specialize the data structures as needed to optimize the software for a particular platform. The GraphBLAS are great for people interested in sparse linear algebra and designing their own graph algorithms. The GraphBLAS library, however, does not include any graph algorithms. The GraphBLAS provide a software framework for constructing graph algorithms, but it doesn't provide any actual Graph Algorithms. Since most people working with graphs use algorithms but don't develop them "from scratch", the GraphBLAS are not really useful to most people.

Hence, there is a need for a library of Graph Algorithms implemented on top of the GraphBLAS. We have created this library. It is called LAGraph. The LAGraph library is a library of functions that implement the most common high level graph algorithms used in graph analytics. It includes types, utility functions and everything needed to incorporate graph algorithms into your analytics work flows. The library uses the GraphBLAS objects (e.g., `GrB_matrix` and `GrB_vector`) inside the objects defined by LAGraph. Consequently, GraphBLAS and LAGraph functions can be freely mixed inside a single program.

A graph in LAGraph uses the LAGraph\_Graph data type. Unlike the GrB\_matrix object, an LAGraph\_Graph object is not opaque. The elements of the data structure are available to the user of the LAGraph library. The data associated with an LAGraph\_Graph is represented by an GrB\_matrix. The data structure includes information about the graph and key properties of the graph. For example, many algorithms require not only the matrix representing a graph, but also its transpose. These (and other) properties can be stored within the LAGraph\_Graph. Storage of properties such as the transpose of a matrix requires additional storage, but the performance impact can more than compensate for the cost associated with that extra memory.

The algorithms within LAGraph roughly break down into two categories: Basic (LAGraph\_\*) and advanced (LAGr\_\*). The idea is that users who are not familiar with the ways graph algorithms are implemented and just want to apply an algorithm to their graphs, would use the Basic interface. For advanced users who are comfortable working with key aspects of the algorithms they are working with might see a significant performance benefit from working with the advanced algorithm.

For example, the basic and advanced algorithms deal with the properties of an LAGraph graph differently. The basic algorithm assumes the user will not set-up the LAGraph\_Graph with the properties needed by an algorithm. Such properties will be computed as needed. An advanced user, however, may know that the string of operations in a workflow all requires a subset of key properties. By computing them in advanced and storing them with the LAGraph graph, the workflow can run much faster since it won't need to, for example, rearrange a matrix into its transpose for each algorithm in a workflow.



## LAGRAPH CONTEXT AND ERROR HANDLING

The sections below describe a set of functions that manage the LAGraph context within a user application, and discuss how errors are handled.

### 2.1 LAGraph Context Functions

int **LAGraph\_Init**(char \*msg)

**LAGraph\_Init**: initializes GraphBLAS and LAGraph. This method must be called before calling any other GrB\* or LAGraph\* method. It initializes GraphBLAS with GrB\_init and then performs LAGraph-specific initializations. In particular, the LAGraph semirings listed below are created. GrB\_init can also safely be called before calling *LAGr\_Init* or **LAGraph\_Init**.

**Parameters** *msg* – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_INVALID\_VALUE** – if **LAGraph\_Init** or **LAGr\_Init** has already been called by the user application.

int **LAGr\_Init**(GrB\_Mode mode, void \*(\*user\_malloc\_function)(size\_t), void \*(\*user\_calloc\_function)(size\_t, size\_t), void \*(\*user\_realloc\_function)(void\*, size\_t), void (\*user\_free\_function)(void\*), char \*msg)  
**LAGr\_Init**: initializes GraphBLAS and LAGraph. **LAGr\_Init** is identical to *LAGraph\_Init*, except that it allows the user application to specify the GraphBLAS mode. It also provides four memory management functions, replacing the standard malloc, calloc, realloc, and free. The functions *user\_malloc\_function*, *user\_calloc\_function*, *user\_realloc\_function*, and *user\_free\_function* have the same signature as the ANSI C malloc, calloc, realloc, and free functions, respectively. Only *user\_malloc\_function* and *user\_free\_function* are required. *user\_calloc\_function* may be NULL, in which case **LAGraph\_Calloc** uses **LAGraph\_Malloc** and **memset**. Likewise, *user\_realloc\_function* may be NULL, in which case **LAGraph\_Realloc** uses **LAGraph\_Malloc**, **memcpy**, and **LAGraph\_Free**.

**Parameters**

- **mode** – [in] the mode for GrB\_Init
- **user\_malloc\_function** – [in] pointer to a malloc function
- **user\_calloc\_function** – [in] pointer to a calloc function, or NULL
- **user\_realloc\_function** – [in] pointer to a realloc function, or NULL
- **user\_free\_function** – [in] pointer to a free function
- *msg* – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_INVALID\_VALUE** – if LAGraph\_Init or LAGr\_Init has already been called by the user application.
- **GrB\_NULL\_POINTER** – if user\_malloc\_function or user\_free\_function are NULL.

int **LAGraph\_Finalize**(char \*msg)

LAGraph\_Finalize: finish LAGraph and GraphBLAS. Must be called as the last LAGraph method. It calls GrB\_finalize and frees any LAGraph objects created by *LAGraph\_Init* or *LAGr\_Init*. After calling this method, no LAGraph or GraphBLAS methods may be used.

**Parameters** **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns** **GrB\_SUCCESS** – if successful.

int **LAGraph\_Version**(int version\_number[3], char \*version\_date, char \*msg)

LAGraph\_Version: determines the version of LAGraph. The version number and date can also be obtained via compile-time constants from LAGraph.h. However, it is possible to compile a user application that #includes one version of LAGraph.h and then links with another version of the LAGraph library later on, so the version number and date may differ from the compile-time constants.

The LAGraph\_Version method allows the library itself to be queried, after it is linked in with the user application.

The version\_number array is set to LAGRAPH\_VERSION\_MAJOR, LAGRAPH\_VERSION\_MINOR, and LAGRAPH\_VERSION\_UPDATE, in that order. The LAGRAPH\_DATE string is copied into the user-provided version\_date string, and is null-terminated.

**Parameters**

- **version\_number** – [out] an array of size 3; with the major, minor, and update versions of LAGraph, in that order.
- **version\_date** – [out] an array of size >= LAGraph\_MSG\_LEN, returned with the date of this version of LAGraph.
- **msg** – [inout] any error messages.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if version\_number or version\_date are NULL.

int **LAGraph\_GetNumThreads**(int \*nthreads\_outer, int \*nthreads\_inner, char \*msg)

LAGraph\_GetNumThreads determines the current number of OpenMP threads that can be used. See LAGraph\_SetNumThreads for a description of nthreads\_outer and nthreads\_inner.

**Parameters**

- **nthreads\_outer** – [out] number of threads for outer region.
- **nthreads\_inner** – [out] number of threads for inner region, or for the underlying GraphBLAS library.
- **msg** – [inout] any error messages.

**Returns**

- **GrB\_SUCCESS** – if successful.

- **GrB\_NULL\_POINTER** – if `nthreads_outer` or `nthreads_inner` are NULL.

int **LAGraph\_SetNumThreads**(int `nthreads_outer`, int `nthreads_inner`, char \*`msg`)

LAGraph\_SetNumThreads sets the current number of OpenMP threads that can be used by LAGraph and GraphBLAS. Two thread counts can be controlled:

#### Parameters

- **nthreads\_outer** – [in] number of threads to be used in outer regions of a nested parallel construct assuming that `nthreads_inner` is used in the inner region. The total number of threads used for an entire nested region in LAGraph is given by `nthreads_outer*nthreads_inner`. This product is also the # of threads that a flat parallel region in LAGraph may use.
- **nthreads\_inner** – [in] number of threads to be used in an inner region of a nested parallel construct, or for the # of threads to be used in each call to the underlying GraphBLAS library.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns** **GrB\_SUCCESS** – if successful.

## 2.2 Error handling

### LAGRAPH\_RETURN\_VALUES

Nearly all LAGraph methods return an int to denote their status, and have a final string (`msg`) that captures any error messages.

LAGraph has a single function that does not follow this rule. *LAGraph\_WallClockTime* has no error handling mechanism (it returns a value of type double, and does not have an final `msg` string parameter).

All other methods return an int to denote their status: zero if they are successful (which is the value of **GrB\_SUCCESS**), negative on error, or positive for an informational value (such as **GrB\_NO\_VALUE**). Integers in the range -999 to 999 are reserved for GraphBLAS **GrB\_Info** return values:

#### successful results:

- **GrB\_SUCCESS** = 0 // all is well
- **GrB\_NO\_VALUE** = 1 // A(i,j) requested but not there

#### errors:

- **GrB\_UNINITIALIZED\_OBJECT** = -1 // object has not been initialized
- **GrB\_NULL\_POINTER** = -2 // input pointer is NULL
- **GrB\_INVALID\_VALUE** = -3 // generic error; some value is bad
- **GrB\_INVALID\_INDEX** = -4 // row or column index is out of bounds
- **GrB\_DOMAIN\_MISMATCH** = -5 // object domains are not compatible
- **GrB\_DIMENSION\_MISMATCH** = -6 // matrix dimensions do not match
- **GrB\_OUTPUT\_NOT\_EMPTY** = -7 // output matrix already has values
- **GrB\_NOT\_IMPLEMENTED** = -8 // method not implemented

- GrB\_PANIC = -101 // unknown error
- GrB\_OUT\_OF\_MEMORY = -102 // out of memory
- GrB\_INSUFFICIENT\_SPACE = -103, // output array not large enough
- GrB\_INVALID\_OBJECT = -104 // object is corrupted
- GrB\_INDEX\_OUT\_OF\_BOUNDS = -105 // row or col index out of bounds
- GrB\_EMPTY\_OBJECT = -106 // an object does not contain a value

LAGraph returns any errors it receives from GraphBLAS, and also uses the GrB\_\* error codes in these cases:

- GrB\_INVALID\_INDEX: if a source node id is out of range
- GrB\_INVALID\_VALUE: if an enum to select an option is out of range
- GrB\_NOT\_IMPLEMENTED: if a type is not supported, or when SuiteSparse GraphBLAS is required.

Summary of return values for all LAGraph functions that return int:

- GrB\_SUCCESS if successful
- a negative GrB\_Info value on error (in range -999 to -1)
- a positive GrB\_Info value if successful but with extra information (in range 1 to 999)
- -1999 to -1000: a common LAGraph-specific error, see list above
- 1000 to 1999: if successful, with extra LAGraph-specific information
- <= -2000: an LAGraph error specific to a particular LAGraph method
- >= 2000: an LAGraph warning specific to a particular LAGraph method

Many LAGraph methods share common error cases, described below. These return values are in the range -1000 to -1999. Return values of -2000 or greater may be used by specific LAGraph methods, which denote errors not in the following list:

- **LAGRAPH\_INVALID\_GRAPH (-1000):** The input graph is invalid; the details are given in the error msg string returned by the method.
- **LAGRAPH\_SYMMETRIC\_STRUCTURE\_REQUIRED (-1001):** The method requires an undirected graph, or a directed graph with an adjacency matrix that is known to have a symmetric structure. LAGraph\_Cached\_IsSymmetricStructure can be used to determine this cached property.
- **LAGRAPH\_IO\_ERROR (-1002):** A file input or output method failed, or an input file has an incorrect format that cannot be parsed.
- **LAGRAPH\_NOT\_CACHED (-1003):** Some methods require one or more cached properties to be computed before calling them (AT, out\_degree, or in\_degree. Use LAGraph\_Cached\_AT, LAGraph\_Cached\_OutDegree, and/or LAGraph\_Cached\_InDegree to compute them.
- **LAGRAPH\_NO\_SELF\_EDGES\_ALLOWED (-1004):** Some methods requires that the graph have no self edges, which correspond to the entries on the diagonal of the adjacency matrix. If the G->nself\_edges cached property is nonzero or unknown, this error condition is returned. Use LAGraph\_Cached\_NSelfEdges to compute G->nself\_edges, or LAGraph\_DeleteSelfEdges to remove all diagonal entries from G->A.
- **LAGRAPH\_CONVERGENCE\_FAILURE (-1005):** An iterative process failed to converge to a good solution.
- **LAGRAPH\_CACHE\_NOT\_NEEDED (1000):** This is a warning, not an error. It is returned by LAGraph\_Cached\_\* methods when asked to compute cached properties that are not needed. These include G->AT and G->in\_degree for an undirected graph.

**LAGRAPH\_MSG\_LEN**

All LAGraph functions (except for *LAGraph\_WallClockTime* ) have a final msg parameter that is a pointer to a user-allocated string in which an algorithm-specific error message can be returned. If msg is NULL, no error message is returned. This is not itself an error condition, it just indicates that the caller does not need the message returned. If the message string is provided but no error occurs, an empty string is returned.

LAGRAPH\_MSG\_LEN is the minimum required length of a message string.

For example, the following call computes the breadth-first-search of an LAGraph\_Graph G, starting at a given source node. It returns a status of zero if it succeeds and a negative value on failure.

```
GrB_Vector level, parent ;
char msg [LAGRAPH_MSG_LEN] ;
int status = LAGr_BreadthFirstSearch (&level, &parent, G, src, msg) ;
if (status < 0)
{
    printf ("status %d, error: %s\n", status, msg) ;
    ... take corrective action here ...
}
```

Error handling is simplified by the *LAGRAPH\_TRY* / *LAGRAPH\_CATCH* mechanism described below. For example, assuming the user application #defines a single LAGRAPH\_CATCH mechanism for all error handling, the above example would become:

```
GrB_Vector level, parent ;
char msg [LAGRAPH_MSG_LEN] ;
#define LAGRAPH_CATCH(status) \
{ \
    printf ("status %d, error: %s\n", status, msg) ; \
    ... take corrective action here ... \
}
...
LAGRAPH_TRY (LAGr_BreadthFirstSearch (&level, &parent, G, src, msg)) ;
```

The advantage of the second use case is that the error-handling block of code needs to be written only once.

**LAGRAPH\_TRY(LAGraph\_method)**

LAGRAPH\_TRY: try an LAGraph method and check for errors.

In a robust application, the return values from each call to LAGraph and GraphBLAS should be checked, and corrective action should be taken if an error occurs. The LAGRAPH\_TRY and *GRB\_TRY* macros assist in this effort.

LAGraph and GraphBLAS are written in C, and so they cannot rely on the try/catch mechanism of C++. To accomplish a similar goal, each LAGraph file must #define its own file-specific macro called LAGRAPH\_CATCH. The typical usage of macro is to free any temporary matrices/vectors or workspace when an error occurs, and then “throw” the error by returning to the caller. A user application may also #define LAGRAPH\_CATCH and use these macros. The LAGRAPH\_CATCH macro takes a single argument, which is the return value from an LAGraph method.

A typical example of a user function that calls LAGraph might #define LAGRAPH\_CATCH as follows. Suppose workvector is a GrB\_vector used for computations internal to the mybfs function, and W is a (double \*) space created by malloc.

```
// an example user-defined LAGRAPH_CATCH macro, which prints the error
// then frees any workspace or results, and returns to the caller:
```

(continues on next page)

(continued from previous page)

```

#define LAGRAPH_CATCH(status) \
{ \
    printf ("LAGraph error: (%d): file: %s, line: %d\n%s\n", \
        status, __FILE__, __LINE__, msg) ; \
    GrB_free (*parent) ; \
    GrB_free (workvector) ; \
    LAGraph_Free ((void **) &W, NULL) ; \
    return (status) ; \
}

// an example user function that uses LAGRAPH_TRY / LAGRAPH_CATCH
int mybfs (LAGraph_Graph G, GrB_Vector *parent, int64_t src)
{
    GrB_Vector workvector = NULL ;
    double *W = NULL ;
    char msg [LAGRAPH_MSG_LEN] ;
    (*parent) = NULL ;
    LAGRAPH_TRY (LAGr_BreadthFirstSearch (NULL, parent, G, src, true,
        msg)) ;
    // ...
    return (GrB_SUCCESS) ;
}

```

LAGRAPH\_TRY is defined as follows:

```

#define LAGRAPH_TRY(LAGraph_method) \
{ \
    int LG_status = LAGraph_method ; \
    if (LG_status < GrB_SUCCESS) \
    { \
        LAGRAPH_CATCH (LG_status) ; \
    } \
}

```

### GRB\_TRY(GrB\_method)

GRB\_TRY: LAGraph provides a similar functionality as *LAGRAPH\_TRY* for calling GraphBLAS methods, with the GRB\_TRY macro. GraphBLAS returns info = 0 (GrB\_SUCCESS) or 1 (GrB\_NO\_VALUE) on success, and a value < 0 on failure. The user application must #define GRB\_CATCH to use GRB\_TRY.

GraphBLAS and LAGraph both use the convention that negative values are errors, and the LAGraph\_status is a superset of the GrB\_Info enum. As a result, the user can define LAGRAPH\_CATCH and GRB\_CATCH as the same operation. The main difference between the two would be the error message string. For LAGraph, the string is the last parameter, and LAGRAPH\_CATCH can optionally print it out. For GraphBLAS, the GrB\_error mechanism can return a string.

GRB\_TRY is defined as follows:

```

#define GRB_TRY(GrB_method) \
{ \
    GrB_Info LG_GrB_Info = GrB_method ; \
    if (LG_GrB_Info < GrB_SUCCESS) \
    { \
        GRB_CATCH (LG_GrB_Info) ; \
    } \
}

```

(continues on next page)

(continued from previous page)

```

    }
    \
}

```





## THE GRAPH OBJECT

The fundamental object in LAGraph is the LAGraph\_Graph.

typedef struct *LAGraph\_Graph\_struct* \*LAGraph\_Graph

struct **LAGraph\_Graph\_struct**

LAGraph\_Graph: a representation of a graph.

The LAGraph\_Graph G contains a GrB\_Matrix G->A as its primary component. For graphs represented with adjacency matrices, A(i,j) denotes the edge (i,j). Unlike GrB\_\* objects in GraphBLAS, the LAGraph\_Graph data structure is not opaque. User applications have full access to its contents.

An LAGraph\_Graph G contains two kinds of components:

- a. Primary components of the graph, which fully define the graph.
- b. Cached properties of the graph, which can be recreated any time.

### Primary Components

GrB\_Matrix **A**

the adjacency matrix of the graph

*LAGraph\_Kind* **kind**

the kind of graph

### Cached Properties

All of these components may be deleted or set to ‘unknown’ at any time. For example, if AT is NULL, then the transpose of A has not been computed. A scalar cached property of type LAGraph\_Boolean would be set to LAGRAPH\_UNKNOWN to denote that its value is unknown.

If present, the cached properties must be valid and accurate. If the graph changes, these cached properties can either be recomputed or deleted to denote the fact that they are unknown. This choice is up to individual LAGraph methods and utilities.

LAGraph methods can set non-scalar cached properties only if they are constructing the graph. They cannot modify them or create them if the graph is declared as a read-only object in the parameter list of the method.

GrB\_Matrix **AT**

AT = A', the transpose of A, with the same type.

**GrB\_Vector out\_degree**

a GrB\_INT64 vector of length m, if A is m-by-n. where out\_degree(i) is the number of entries in A(i,:). If out\_degree is sparse and the entry out\_degree(i) is not present, then it is assumed to be zero.

**GrB\_Vector in\_degree**

a GrB\_INT64 vector of length n, if A is m-by-n. where in\_degree(j) is the number of entries in A(:,j). If in\_degree is sparse and the entry in\_degree(j) is not present, then it is assumed to be zero. If A is known to have a symmetric structure, the convention is that the degree is held in G->out\_degree, and in G->in\_degree is left as NULL.

**LAGraph\_Boolean is\_symmetric\_structure**

For an undirected graph, this cached property will always be implicitly true and can be ignored. The matrix A for a directed weighted graph will typically be unsymmetric, but might have a symmetric structure. In that case, this scalar cached property can be set to true. By default, this cached property is set to LAGRAPH\_UNKNOWN.

**int64\_t nself\_edges**

number of entries on the diagonal of A, or LAGRAPH\_UNKNOWN if unknown. For the adjacency matrix of a directed or undirected graph, this is the number of self-edges in the graph.

**GrB\_Scalar emin**

minimum edge weight: value, lower bound, or unknown

**LAGraph\_State emin\_state**

- VALUE: emin is equal to the smallest entry, min(G->A)
- BOUND: emin <= min(G->A)
- UNKNOWN: emin is unknown

**GrB\_Scalar emax**

maximum edge weight: value, upper bound, or unknown

**LAGraph\_State emax\_state**

- VALUE: emax is equal to the largest entry, max(G->A)
- BOUND: emax >= max(G->A)
- UNKNOWN: emax is unknown

**enum LAGraph\_Boolean**

LAGraph\_Boolean: a boolean value (true or false) or unknown (-1).

*Values:*

**enumerator LAGraph\_FALSE**

the Boolean value is known to be false.

**enumerator LAGraph\_TRUE**

the Boolean value is known to be true.

**enumerator LAGraph\_BOOLEAN\_UNKNOWN**

Boolean value is unknown.

**enum LAGraph\_State**

LAGraph\_State describes the status of a cached property of a graph. If the cached property is computed in floating-point arithmetic, it may have been computed with roundoff error, but it may still be declared as “value” if the roundoff error is expected to be small, or if the cached property was computed as carefully as possible (to within reasonable roundoff error). The “bound” state indicates that the cached property is an upper or lower bound, depending on the particular cached property. If computed in floating-point arithmetic, an “upper bound” cached property may be actually slightly lower than the actual upper bound, because of floating-point roundoff.

*Values:*

**enumerator LAGraph\_VALUE**

cached property is a known value.

**enumerator LAGraph\_BOUND**

cached property is a bound. The bound is upper or lower, depending on the particular cached property.

**enumerator LAGraph\_STATE\_UNKNOWN**

the property is unknown.

**enum LAGraph\_Kind**

LAGraph\_Kind: the kind of a graph. Currently, only two types of graphs are supported: undirected graphs and directed graphs. Edge weights are assumed to be present. Unweighted graphs can be represented by setting all entries present in the sparsity structure to the same value, typically 1. Additional types of graphs will be added in the future.

*Values:*

**enumerator LAGraph\_ADJACENCY\_UNDIRECTED**

undirected graph.  $G \rightarrow A$  is square and symmetric; both upper and lower triangular parts are present.  $A(i,j)$  is the edge  $(i,j)$ . Results are undefined if  $A$  is unsymmetric.

**enumerator LAGraph\_ADJACENCY\_DIRECTED**

$G \rightarrow A$  is square;  $A(i,j)$  is the edge  $(i,j)$ .

directed graph.

**enumerator LAGraph\_KIND\_UNKNOWN**

unknown kind of graph (-1).

## 3.1 Basic Graph Functions

**int LAGraph\_New**(LAGraph\_Graph \*G, GrB\_Matrix \*A, LAGraph\_Kind kind, char \*msg)

LAGraph\_New: creates a new graph G. The cached properties  $G \rightarrow AT$ ,  $G \rightarrow out\_degree$ , and  $G \rightarrow in\_degree$  are set to NULL, and scalar cached properties are set to LAGRAPH\_UNKNOWN.

**Parameters**

- **G** – [out] handle to the newly created graph, as &G.
- **A** – [inout] adjacency matrix. A is moved into G as  $G \rightarrow A$ , and A itself is set to NULL to denote that is now a part of G. That is,  $\{ G \rightarrow A = A ; A = NULL ; \}$  is performed. When G is deleted,  $G \rightarrow A$  is freed. If A is NULL, the graph is invalid until  $G \rightarrow A$  is set.

- **kind** – [in] the kind of graph to create. This may be LAGRAPH\_UNKNOWN, which must then be revised later before the graph can be used.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is null.

int **LAGraph\_Delete**(*LAGraph\_Graph* \*G, char \*msg)

LAGraph\_Delete: frees a graph G. The adjacency matrix G->A and the cached properties G->AT, G->out\_degree, and G->in\_degree are all freed.

**Parameters**

- **G** – [inout] handle to the graph to be free. \*G is NULL on output. To keep G->A while deleting the graph G, use: { A = G->A ; G->A = NULL ; LAGraph\_Delete (&G, msg) ; }
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns** **GrB\_SUCCESS** – if successful.

int **LAGraph\_DeleteCached**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_DeleteCached: frees all cached properties of a graph G. The graph is still valid. This method should be used if G->A changes, since such changes will normally invalidate G->AT, G->out\_degree, and/or G->in\_degree.

**Parameters**

- **G** – [inout] handle to the graph to be modified. The graph G remains valid on output, but with all cached properties freed. G may be NULL, in which case nothing is done.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns** **GrB\_SUCCESS** – if successful.

int **LAGraph\_Cached\_AT**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_Cached\_AT: constructs G->AT, the transpose of G->A. This matrix is required by some of the algorithms. Basic algorithms may construct G->AT if they require it. The matrix G->AT is then available for subsequent use. If G->A changes, G->AT should be freed and recomputed. If G->AT already exists, it is left unchanged (even if it is not equal to the transpose of G->A). As a result, if G->A changes, G->AT should be explicitly freed.

**Parameters**

- **G** – [inout] graph for which G->AT is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_CACHE\_NOT\_NEEDED** – if G->kind is LA-Graph\_ADJACENCY\_UNDIRECTED.

- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_IsSymmetricStructure**(*LAGraph\_Graph* G, char \*msg)

*LAGraph\_Cached\_IsSymmetricStructure*: determines if the sparsity structure of G->A is symmetric (ignoring its values). If G->kind denotes that the graph is undirected, this cached property is implicitly true (and not checked). Otherwise, this method determines if the structure of G->A for a directed graph G has a symmetric sparsity structure. No work is performed if the cached property is already known.

#### Parameters

- **G** – [inout] graph for which G->is\_symmetric\_structure is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_OutDegree**(*LAGraph\_Graph* G, char \*msg)

*LAGraph\_Cached\_OutDegree*: computes G->out\_degree. No work is performed if it already exists in G.

#### Parameters

- **G** – [inout] graph for which G->out\_degree is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_InDegree**(*LAGraph\_Graph* G, char \*msg)

*LAGraph\_Cached\_InDegree* computes G->in\_degree. No work is performed if it already exists in G. If G is undirected, G->in\_degree is never computed and remains NULL (the method returns **LAGRAPH\_CACHE\_NOT\_NEEDED**). No work is performed if it is already exists in G.

Performance note: for SuiteSparse:GraphBLAS, if G->A is held by row (the default format), then computing G->in\_degree is fastest if G->AT is known. If G->AT will be needed anyway, compute it first with *LAGraph\_Cached\_AT*, and then call *LAGraph\_Cached\_Indegree*. This is optional; if G->AT is not known, then G->in\_degree is computed from G->A instead.

#### Parameters

- **G** – [inout] graph for which G->in\_degree is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.

- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_CACHE\_NOT\_NEEDED** – if G->kind is LA-Graph\_ADJACENCY\_UNDIRECTED.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_NSelfEdges**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_Cached\_NSelfEdges: computes G->nself\_edges, the number of diagonal entries that appear in the G->A matrix. For an undirected or directed graph with an adjacency matrix G->A, these are the number of self-edges in G. No work is performed if it is already computed.

**Parameters**

- **G** – [inout] graph for which G->nself\_edges is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_EMin**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_Cached\_EMin: computes the G->emin, the smallest entry in G->A. Not computed if G->emin already exists.

**Parameters**

- **G** – [inout] graph for which G->emin is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NOT\_IMPLEMENTED** – if G does not have a built-in real type: GrB\_(BOOL, INT8, INT16, INT32, INT64, UINT8, UINT16, UINT32, UINT64, FP32, OR FP64).
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_Cached\_EMax**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_Cached\_EMax: computes the G->emax, the largest entry in G->A. Not computed if G->emax already exists.

**Parameters**

- **G** – [inout] graph for which G->emax is computed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NOT\_IMPLEMENTED** – if G does not have a built-in real type: GrB\_(BOOL, INT8, INT16, INT32, INT64, UINT8, UINT16, UINT32, UINT64, FP32, OR FP64).
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_DeleteSelfEdges**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_DeleteSelfEdges: removes any diagonal entries from G->A. Most cached properties are cleared or set to LAGRAPH\_UNKNOWN. G->nself\_edges is set to zero, and G->is\_symmetric\_structure is left unchanged.

#### Parameters

- **G** – [inout] graph for which G->A is modified.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (G->A missing, or G->kind not a recognized kind).

int **LAGraph\_CheckGraph**(*LAGraph\_Graph* G, char \*msg)

LAGraph\_CheckGraph: determines if a graph is valid. Only basic checks are performed on the cached properties, taking O(1) time.

#### Parameters

- **G** – [in] graph to check.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid: G->A missing, G->kind not a recognized kind, G->AT present but has the wrong dimensions or its type does not match G->A, G->in\_degree/out\_degree present but with the wrong dimension or type (in/out\_degree must be GrB\_INT64).





## ALGORITHMS

Algorithms come in two flavors: *Basic* and *Advanced*.

### 4.1 Basic

Basic algorithms are meant to be easy to use. A single basic algorithm may encompass many underlying Advanced algorithms, each with various parameters that may be controlled. For the Basic API, these parameters are determined automatically. Cached graph properties may be determined, and as a result, the graph *G* is both an input and an output of these methods, since they may be modified.

LAGraph Basic algorithms are named with the `LAGraph_*` prefix.

int **LAGraph\_TriangleCount**(uint64\_t \*ntriangles, *LAGraph\_Graph* G, char \*msg)

LAGraph\_TriangleCount: count the triangles in a graph. This is a Basic algorithm (*G*->nself\_edges, *G*->out\_degree, *G*->is\_symmetric\_structure are computed, if not present).

#### Parameters

- **ntriangles** – [out] the number of triangles in *G*.
- **G** – [inout] the graph, which must be undirected, or directed but with a symmetric structure. No self loops can be present.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if *G* or *ntriangles* are NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if *G* is invalid (*LAGraph\_CheckGraph* failed).
- **LAGRAPH\_NO\_SELF\_EDGES\_ALLOWED** – if *G* has any self-edges.
- **LAGRAPH\_SYMMETRIC\_STRUCTURE\_REQUIRED** – if *G* is directed with an un-symmetric *G*->*A* matrix.

## 4.2 Advanced

The Advanced algorithms require the caller to select the algorithm and choose any parameter settings. G is not modified, and so it is an input-only parameter to these methods. If an Advanced algorithm requires a cached graph property to be computed, it must be computed prior to calling the Advanced method.

Advanced algorithms are named with the `LAGr_*` prefix, to distinguish them from Basic algorithms.

int **LAGr\_SortByDegree**(int64\_t \*\*P, const *LAGraph\_Graph* G, bool byout, bool ascending, char \*msg)

`LAGr_SortByDegree` sorts the nodes of a graph by their out or in degrees. The graph `G->A` itself is not changed. Refer to `LAGr_TriangleCount` for an example of how to permute `G->A` after calling this function. The output `&P` must be freed by `LAGraph_Free`. This method requires `G->out_degree` or `G->in_degree` to already be computed.

### Parameters

- **P** – [out] permutation of the integers 0..n-1.
- **G** – [in] graph of n nodes.
- **byout** – [in] if true, sort by out-degree, else sort by in-degree.
- **ascending** – [in] if true, sort in ascending order, else descending.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_NULL\_POINTER** – if P or G are NULL.
- **LAGRAPH\_NOT\_CACHED** – if `G->in_degree` or `G->out_degree` is not computed (whichever one is required).
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (*LAGraph\_CheckGraph* failed).

int **LAGr\_SampleDegree**(double \*sample\_mean, double \*sample\_median, const *LAGraph\_Graph* G, bool byout, int64\_t nsamples, uint64\_t seed, char \*msg)

`LAGr_SampleDegree` computes an estimate of the median and mean of the out or in degree, by randomly sampling the `G->out_degree` or `G->in_degree` vector. This method requires `G->out_degree` or `G->in_degree` to already be computed.

### Parameters

- **sample\_mean** – [out] sampled mean of the degree.
- **sample\_median** – [out] sampled median of the degree.
- **G** – [in] graph to sample.
- **byout** – [in] if true, sample out-degree, else sample in-degree.
- **nsamples** – [in] number of samples to take.
- **seed** – [in] random number seed.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_NULL\_POINTER** – if `sample_mean`, `sample_median`, or G are NULL.
- **LAGRAPH\_NOT\_CACHED** – if `G->in_degree` or `G->out_degree` is not computed (whichever one is required).

- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (*LAGraph\_CheckGraph* failed).

int **LAGr\_BreadthFirstSearch**(GrB\_Vector \*level, GrB\_Vector \*parent, const *LAGraph\_Graph* G, GrB\_Index src, char \*msg)

LAGr\_BreadthFirstSearch: breadth-first search of a graph, computing the breadth-first-search tree and/or the level of the nodes encountered. This is an Advanced algorithm. G->AT and G->out\_degree are required to use the fastest push/pull method when using SuiteSparse:GraphBLAS. If these cached properties are not present, or if a vanilla GraphBLAS library is being used, then a push-only method is used (which can be slower). G is not modified; that is, G->AT and G->out\_degree are not computed if not already cached.

#### Parameters

- **level** – [out] If non-NULL on input, on successful return, it contains the levels of each node reached. The src node is assigned level 0. If a node i is not reached, level(i) is not present. The level vector is not computed if NULL.
- **parent** – [out] If non-NULL on input, on successful return, it contains parent node IDs for each node reached, where parent(i) is the node ID of the parent of node i. The src node will have itself as its parent. If a node i is not reached, parent(i) is not present. The parent vector is not computed if NULL.
- **G** – [in] The graph, directed or undirected.
- **src** – [in] The index of the src node (0-based)
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_INVALID\_INDEX** – if src is invalid.
- **GrB\_NULL\_POINTER** – if both level and parent are NULL, or if G is NULL.
- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid (*LAGraph\_CheckGraph* failed).

int **LAGr\_ConnectedComponents**(GrB\_Vector \*component, const *LAGraph\_Graph* G, char \*msg)

LAGr\_ConnectedComponents: connected components of an undirected graph. This is an Advanced algorithm (G->is\_symmetric\_structure must be known).

#### Parameters

- **component** – [out] component(i)=s if node i is in the component whose representative node is s. If node i has no edges, it is placed in its own component, and thus the component vector is always dense.
- **G** – [in] input graph to find the components for. The graph must be undirected, or G->is\_symmetric\_structure must be true.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G or component are NULL.
- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid (*LAGraph\_CheckGraph* failed).
- **LAGRAPH\_SYMMETRIC\_STRUCTURE\_REQUIRED** – if G is directed with an un-symmetric G->A matrix.

int **LAGr\_SingleSourceShortestPath**(GrB\_Vector \*path\_length, const *LAGraph\_Graph* G, GrB\_Index src, GrB\_Scalar Delta, char \*msg)

LAGr\_SingleSourceShortestPath: single-source shortest paths. This is an Advanced algorithm (G->emin is required for best performance). The graph G must have an adjacency matrix of type GrB\_INT32, GrB\_INT64, GrB\_UINT32, GrB\_UINT64, GrB\_FP32, or GrB\_FP64. If G->A has any other type, GrB\_NOT\_IMPLEMENTED is returned.

#### Parameters

- **path\_length** – [out] path\_length(i) is the length of the shortest path from the source node to node i. The path\_length vector is dense. If node (i) is not reachable from the src node, then path\_length(i) is set to INFINITY for GrB\_FP32 and FP32, or the maximum integer for GrB\_INT32, INT64, UINT32, or UINT64.
- **G** – [in] input graph.
- **src** – [in] source node.
- **Delta** – [in] for delta stepping.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G or path\_length are NULL.
- **GrB\_INVALID\_INDEX** – if src is invalid.
- **GrB\_EMPTY\_OBJECT** – if Delta does not contain a value.
- **GrB\_NOT\_IMPLEMENTED** – if the type is not supported.
- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid (*LAGraph\_CheckGraph* failed).

int **LAGr\_Betweenness**(GrB\_Vector \*centrality, const *LAGraph\_Graph* G, const GrB\_Index \*sources, int32\_t ns, char \*msg)

LAGr\_Betweenness: betweenness centrality metric. This methods computes an approximation of the betweenness-centrality metric of all nodes in the graph. Only a few given source nodes are used for the approximation. This is an Advanced algorithm (G->AT is required).

#### Parameters

- **centrality** – [out] centrality(i) is the metric for node i.
- **G** – [in] input graph.
- **sources** – [in] source vertices to compute shortest paths, size ns
- **ns** – [in] number of source vertices.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G, centrality, and/or sources are NULL.
- **GrB\_INVALID\_INDEX** – if any source node is invalid.
- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid (*LAGraph\_CheckGraph* failed).

- **LAGRAPH\_NOT\_CACHED** – if G->AT is required but not present.

int **LAGr\_PageRank**(GrB\_Vector \*centrality, int \*iters, const *LAGraph\_Graph* G, float damping, float tol, int itermax, char \*msg)

LAGr\_PageRank: computes the standard PageRank of a directed graph G. Sinks (nodes with no out-going edges) are properly handled. This method should be used for production, not for the GAP benchmark. This is an Advanced algorithm (G->AT and G->out\_degree are required).

#### Parameters

- **centrality** – [out] centrality(i) is the PageRank of node i.
- **iters** – [out] number of iterations taken.
- **G** – [in] input graph.
- **damping** – [in] damping factor (typically 0.85).
- **tol** – [in] stopping tolerance (typically 1e-4).
- **itermax** – [in] maximum number of iterations (typically 100).
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G, centrality, and/or iters are NULL.
- **LAGRAPH\_NOT\_CACHED** – if G->AT is required but not present, or if G->out\_degree is not present.
- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid ( *LAGraph\_CheckGraph* failed).

int **LAGr\_TriangleCount**(uint64\_t \*ntriangles, const *LAGraph\_Graph* G, *LAGr\_TriangleCount\_Method* \*method, *LAGr\_TriangleCount\_Presort* \*presort, char \*msg)

LAGr\_TriangleCount: count the triangles in a graph (advanced API).

#### Parameters

- **ntriangles** – [out] the number of triangles in G.
- **G** – [in] The graph, which must be undirected or have G->is\_symmetric\_structure true, with no self loops. G->nself\_edges, G->out\_degree, and G->is\_symmetric\_structure are required.
- **method** – [inout] specifies which algorithm to use, and returns the method chosen. If NULL, the AutoMethod is used, and the method is not reported. Also see the LAGr\_TriangleCount\_Method enum description.
- **presort** – [inout] controls the presort of the graph, and returns the presort chosen. If NULL, the AutoSort is used, and the presort method is not reported. Also see the description of the LAGr\_TriangleCount\_Presort enum.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G or ntriangles are NULL.

- **LAGRAPH\_INVALID\_GRAPH** – Graph is invalid ( *LAGraph\_CheckGraph* failed).
- **LAGRAPH\_NO\_SELF\_EDGES\_ALLOWED** – if G has any self-edges, or if G->nself\_edges is not computed.
- **LAGRAPH\_SYMMETRIC\_STRUCTURE\_REQUIRED** – if G is directed with an un-symmetric G->A matrix.
- **LAGRAPH\_NOT\_CACHED** – if G->out\_degree is not present in G.
- **GrB\_INVALID\_VALUE** – method or presort are invalid.

enum **LAGr\_TriangleCount\_Method**

LAGr\_TriangleCount\_Method: an enum to select the method used to count the number of triangles.

*Values:*

enumerator **LAGr\_TriangleCount\_AutoMethod**  
auto selection of method

enumerator **LAGr\_TriangleCount\_Burkhardt**  
 $\text{sum}(\text{sum}((A^2) .* A)) / 6$

enumerator **LAGr\_TriangleCount\_Cohen**  
 $\text{sum}(\text{sum}((L * U) .* A)) / 2$

enumerator **LAGr\_TriangleCount\_Sandia\_LL**  
 $\text{sum}(\text{sum}((L * L) .* L))$

enumerator **LAGr\_TriangleCount\_Sandia\_UU**  
 $\text{sum}(\text{sum}((U * U) .* U))$

enumerator **LAGr\_TriangleCount\_Sandia\_LUT**  
 $\text{sum}(\text{sum}((L * U') .* L))$

enumerator **LAGr\_TriangleCount\_Sandia\_ULT**  
 $\text{sum}(\text{sum}((U * L') .* U))$

enum **LAGr\_TriangleCount\_Presort**

LAGr\_TriangleCount\_Presort: an enum to control if/how the matrix is sorted prior to counting triangles.

*Values:*

enumerator **LAGr\_TriangleCount\_NoSort**  
no sort

enumerator **LAGr\_TriangleCount\_Ascending**  
sort by degree, ascending.

enumerator **LAGr\_TriangleCount\_Descending**  
sort by degree, descending.

enumerator **LAGr\_TriangleCount\_AutoSort**  
auto selection of presort: No presort is done for the Burkhardt or Cohen methods, and no sort is done for the Sandia\_\* methods if the sampled mean out-degree is  $\leq 4 * \text{the sample median out-degree}$ . Otherwise: sort

in ascending order for Sandia\_LL and Sandia\_LUT, descending ordering for Sandia\_UU and Sandia\_ULT. On output, presort is modified to reflect the sorting method used (NoSort, Ascending, or Descending).





## UTILITY FUNCTIONS

## 5.1 Input/Output Functions

int **LAGraph\_MMRead**(GrB\_Matrix \*A, FILE \*f, char \*msg)

LAGraph\_MMRead: reads a matrix in MatrixMarket format. The file format used here is compatible with all variations of the Matrix Market “coordinate” and “array” format (<http://www.nist.gov/MatrixMarket>), for sparse and dense matrices respectively. The format is fully described in [LAGraph/Doc/MatrixMarket.pdf](#), and summarized here (with extensions for LAGraph).

## 5.1.1 First Line

The first line of the file starts with %%MatrixMarket, with the following format:

```
%%MatrixMarket matrix <fmt> <type> <storage>
```

**<fmt>** One of:

- coordinate : sparse matrix in triplet form
- array : dense matrix in column-major form

Both formats are returned as a GrB\_Matrix.

If not present, defaults to coordinate.

**<type>** One of:

- real : returns as GrB\_FP64
- integer : returns as GrB\_INT64
- pattern : returns as GrB\_BOOL
- complex : *currently not supported*

The return type can be modified by the %%GraphBLAS structured comment described below.

If not present, defaults to real.

**<storage>** One of:

- **general** the matrix has no symmetry properties (or at least none that were exploited when the file was created).

- **Hermitian** square complex matrix with  $A(i,j) = \text{conj}(A(j,i))$ . All entries on the diagonal are real. Each off-diagonal entry in the file creates two entries in the GrB\_Matrix that is returned.
- **symmetric**  $A(i,j) == A(j,i)$ . Only entries on or below the diagonal appear in the file. Each off-diagonal entry in the file creates two entries in the GrB\_Matrix that is returned.
- **skew-symmetric**  $A(i,j) == -A(j,i)$ . There are no entries on the diagonal. Only entries below the diagonal appear in the file. Each off-diagonal entry in the file creates two entries in the GrB\_Matrix that is returned.

The Matrix Market format is case-insensitive, so “hermitian” and “Hermitian” are treated the same.

If not present, defaults to general.

Not all combinations are permitted. Only the following are meaningful:

- (1) (coordinate or array) x (real, integer, or complex) x (general, symmetric, or skew-symmetric)
- (2) (coordinate or array) x (complex) x (Hermitian)
- (3) (coordinate) x (pattern) x (general or symmetric)

### 5.1.2 Second Line

The second line is an optional extension to the Matrix Market format:

```
%%GraphBLAS type <entrytype>
```

**<entrytype>** One of the 11 built-in types (bool, int8\_t, int16\_t, int32\_t, int64\_t, uint8\_t, uint16\_t, uint32\_t, uint64\_t, float, or double).

If this second line is included, it overrides the default GraphBLAS types for the Matrix Market <type> on line one of the file: real, pattern, and integer. The Matrix Market complex <type> is not yet supported.

### 5.1.3 Other Lines

Any other lines starting with “%” are treated as comments, and are ignored. Comments may be interspersed throughout the file. Blank lines are ignored. The Matrix Market header is optional in this routine (it is not optional in the Matrix Market format). The remaining lines are space delimited, and free format (one or more spaces can appear, and each field has arbitrary width).

### 5.1.4 Coordinate Format

For coordinate format, the first non-comment line must appear, and it must contain three integers:

```
nrows ncols nvals
```

For example, a 5-by-12 matrix with 42 entries would have:

```
5 12 42
```

Each of the remaining lines defines one entry. The order is arbitrary. If the Matrix Market <type> is real or integer, each line contains three numbers: row index, column index, and value. For example, if  $A(3,4)$  is equal to 5.77, a line:

```
3 4 5.77
```

would appear in the file. The indices in the Matrix Market are 1-based, so this entry becomes A(2,3) in the GrB\_Matrix returned to the caller. If the <type> is pattern, then only the row and column index appears. If <type> is complex, four values appear. If A(8,4) has a real part of 6.2 and an imaginary part of -9.3, then the line is:

```
8 4 6.2 -9.3
```

and since the file is 1-based but a GraphBLAS matrix is always 0-based, one is subtracted from the row and column indices in the file, so this entry becomes A(7,3). Note however that LAGraph does not yet support complex types.

### 5.1.5 Array Format

For array format, the first non-comment line must appear, and it must contain just two integers:

```
nrows ncols
```

A 5-by-12 matrix would have this as the first non-comment line after the header:

```
5 12
```

Each of the remaining lines defines one entry, in column major order. If the <type> is real or integer, this is the value of the entry. An entry if <type> of complex consists of two values, the real and imaginary part (not yet supported). The <type> cannot be pattern in this case.

### 5.1.6 Infinity & Not-A-Number

For both coordinate and array formats, real and complex values may use the terms INF, +INF, -INF, and NAN to represent floating-point infinity and NaN values, in either upper or lower case.

According to the Matrix Market format, entries are always listed in column-major order. This rule is followed by *LAGraph\_MMWrite*. However, *LAGraph\_MMRead* can read the entries in any order.

#### Parameters

- **A** – [out] handle of the matrix to create.
- **f** – [inout] handle to an open file to read from.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if A or f are NULL.
- **LAGRAPH\_IO\_ERROR** – if the file could not be read or contains a matrix with an invalid Matrix Market format.
- **GrB\_NOT\_IMPLEMENTED** – if the type is not supported. Complex types (GxB\_FC32 and GxB\_FC64 in SuiteSparse:GraphBLAS) are not yet supported.

int **LAGraph\_MMWrite**(GrB\_Matrix A, FILE \*f, FILE \*fcomments, char \*msg)

LAGraph\_MMWrite: writes a matrix in MatrixMarket format. Refer to [LAGraph\\_MMRead](#) for a description of the output file format. The MatrixMarket header line always appears, followed by the second line containing the GraphBLAS type:

```
%%GraphBLAS type <entrytype>
```

### Parameters

- **A** – [in] matrix to write.
- **f** – [inout] handle to an open file to write to.
- **fcomments** – [in] optional handle to an open file containing comments; may be NULL.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if A or f are NULL.
- **LAGRAPH\_IO\_ERROR** – if the file could not be written to.
- **GrB\_NOT\_IMPLEMENTED** – if the type is not supported. Complex types (GxB\_FC32 and GxB\_FC64 in SuiteSparse:GraphBLAS) are not yet supported.

double **LAGraph\_WallClockTime**(void)

LAGraph\_WallClockTime returns the current wall clock time. Normally, this is simply a wrapper for `omp_get_wtime`, if OpenMP is in use. Otherwise, an OS-specific timing function is called. Note that unlike all other LAGraph functions, this function does not return an error condition, nor does it have a msg string parameter. Instead, it returns the current wall clock time (in seconds) since some fixed point in the past.

**Returns** the current wall clock time.

## 5.2 Matrix Structure Functions

int **LAGraph\_Matrix\_Structure**(GrB\_Matrix \*C, GrB\_Matrix A, char \*msg)

LAGraph\_Matrix\_Structure: returns the sparsity structure of a matrix A as a boolean (GrB\_BOOL) matrix C. If A(i,j) appears in the sparsity structure of A, then C(i,j) is set to true. The sparsity structure of A and C are identical.

### Parameters

- **C** – [out] A boolean matrix with same structure of A, with C(i,j) true if A(i,j) appears in the sparsity structure of A.
- **A** – [in] matrix to compute the structure for.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if A or C are NULL.

int **LAGraph\_Vector\_Structure**(GrB\_Vector \*w, GrB\_Vector u, char \*msg)

LAGraph\_Vector\_Structure: returns the sparsity structure of a vector u as a boolean (GrB\_BOOL) vector w. If u(i) appears in the sparsity structure of u, then w(i) is set to true. The sparsity structure of u and w are identical.

#### Parameters

- **w** – [out] A boolean vector with same structure of u, with w(i) true if u(i,j) appears in the sparsity structure of u.
- **u** – [in] vector to compute the structure for.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if w or u are NULL.

## 5.3 Matrix Comparison Functions

int **LAGraph\_Matrix\_IsEqual**(bool \*result, const GrB\_Matrix A, const GrB\_Matrix B, char \*msg)

LAGraph\_Matrix\_IsEqual compares two matrices for exact equality. If the two matrices have different data types, the result is always false (no typecasting is performed). Only the 11 built-in GrB\* types are supported. If both A and B are NULL, the return value is true. If A and/or B are floating-point types and contain NaN's, result is false.

#### Parameters

- **result** – [out] true if A and B are exactly equal, false otherwise.
- **A** – [in] matrix to compare.
- **B** – [in] matrix to compare.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

#### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if result is NULL.
- **GrB\_NOT\_IMPLEMENTED** – if A or B has a user-defined type.

int **LAGraph\_Matrix\_IsEqualOp**(bool \*result, const GrB\_Matrix A, const GrB\_Matrix B, const GrB\_BinaryOp op, char \*msg)

LAGraph\_Matrix\_IsEqualOp compares two matrices using the given binary operator. The op may be built-in or user-defined. The two matrices may have different types and still be determined to be equal. To be equal, two matrices must have the same sparsity structure, and op(a<sub>ij</sub>,b<sub>ij</sub>) must return true for all pairs of entries a<sub>ij</sub> and b<sub>ij</sub> that appear in the structure of both A and B. The matrices A and/or B can have any type, as long as they are valid inputs to the op. If both A and B are NULL, the return value is true.

#### Parameters

- **result** – [out] true if A and B are equal (per the op), false otherwise.
- **A** – [in] matrix to compare.
- **B** – [in] matrix to compare.

- **op** – [in] operator for the comparison.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if result or op are NULL.

int **LAGraph\_Vector\_IsEqual**(bool \*result, const GrB\_Vector u, const GrB\_Vector v, char \*msg)

LAGraph\_Vector\_IsEqual compares two vectors for exact equality. If the two vectors have different data types, the result is always false (no typecasting is performed). Only the 11 built-in GrB\* types are supported. If both u and v are NULL, the return value is true. If u and/or v are floating-point types and contain NaN's, result is false.

**Parameters**

- **result** – [out] true if u and v are exactly equal, false otherwise.
- **u** – [in] vector to compare.
- **v** – [in] vector to compare.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if result is NULL.
- **GrB\_NOT\_IMPLEMENTED** – if u or v has a user-defined type.

int **LAGraph\_Vector\_IsEqualOp**(bool \*result, const GrB\_Vector u, const GrB\_Vector v, const GrB\_BinaryOp op, char \*msg)

LAGraph\_Vector\_IsEqualOp compares two vectors using the given binary operator. The op may be built-in or user-defined. The two vectors may have different types and still be determined to be equal. To be equal, two vectors must have the same sparsity structure, and op(ui,vi) must return true for all pairs of entries ui and vi that appear in the structure of both u and v. The vectors u and/or v can have any type, as long as they are valid inputs to the op. If both u and v are NULL, the return value is true.

**Parameters**

- **result** – [out] true if u and v are equal (per the op), false otherwise.
- **u** – [in] vector to compare.
- **v** – [in] vector to compare.
- **op** – [in] operator for the comparison.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if result or op are NULL.

## 5.4 Introspecting Types

int **LAGraph\_NameOfType**(char \*name, GrB\_Type type, char \*msg)

LAGraph\_NameOfType returns the name of a GraphBLAS type as a string. The names for the 11 built-in types (GrB\_BOOL, GrB\_INT8, etc) correspond to the names of the corresponding C types (bool, int8\_t, etc).

### Parameters

- **name** – [out] name of the type: user provided array of size at least LAGRAPH\_MAX\_NAME\_LEN.
- **type** – [in] GraphBLAS type to find the name of.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if name or type are NULL.

int **LAGraph\_TypeFromName**(GrB\_Type \*type, char \*name, char \*msg)

LAGraph\_TypeFromName: returns the GrB\_Type corresponding to its name. That is, given the string “bool”, this method returns GrB\_BOOL.

### Parameters

- **type** – [out] GraphBLAS type corresponding to the given name string.
- **name** – [in] name of the type: a null-terminated string.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if name or type are NULL.

int **LAGraph\_SizeOfType**(size\_t \*size, GrB\_Type type, char \*msg)

LAGraph\_SizeOfType: returns sizeof(...) of a GraphBLAS GrB\_Type. For example, if given the GrB\_Type of GrB\_FP64, the value sizeof(double) is returned.

### Parameters

- **size** – [out] size of the type
- **type** – [in] GraphBLAS type to find the size of.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

### Returns

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if size or type are NULL.

int **LAGraph\_Matrix\_TypeName**(char \*name, GrB\_Matrix A, char \*msg)

LAGraph\_Matrix\_TypeName: returns the name of the GrB\_Type of a GrB\_Matrix.

### Parameters

- **name** – [out] name of the type of the matrix: user provided array of size at least LAGRAPH\_MAX\_NAME\_LEN.
- **A** – [in] GraphBLAS matrix to find the type name of.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if name or A are NULL.

int **LAGraph\_Vector\_TypeName**(char \*name, GrB\_Vector v, char \*msg)

LAGraph\_Vector\_TypeName: returns the name of the GrB\_Type of a GrB\_Vector.

**Parameters**

- **name** – [out] name of the type of the vector: user provided array of size at least LAGRAPH\_MAX\_NAME\_LEN.
- **v** – [in] GraphBLAS vector to find the type name of.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if name or v are NULL.

int **LAGraph\_Scalar\_TypeName**(char \*name, GrB\_Scalar s, char \*msg)

LAGraph\_Scalar\_TypeName: returns the name of the GrB\_Type of a GrB\_Scalar.

**Parameters**

- **name** – [out] name of the type of the scalar: user provided array of size at least LAGRAPH\_MAX\_NAME\_LEN.
- **s** – [in] GraphBLAS scalar to find the type name of.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if name or s are NULL.

## 5.5 Printing

int **LAGraph\_Graph\_Print**(const *LAGraph\_Graph* G, *LAGraph\_PrintLevel* pr, FILE \*f, char \*msg)

LAGraph\_Graph\_Print: prints the contents of a graph to a file in a human- readable form. This method is not meant for saving a graph to a file; see *LAGraph\_MMWrite* for that method.

**Parameters**

- **G** – [in] graph to display.
- **pr** – [in] print level.



- **f** – [inout] handle to an open file to write to.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if G or f are NULL.
- **LAGRAPH\_INVALID\_GRAPH** – if G is invalid (*LAGraph\_CheckGraph* failed).
- **GrB\_NOT\_IMPLEMENTED** – if G->A has a user-defined type.
- **LAGRAPH\_IO\_ERROR** – if the file could not be written to.

int **LAGraph\_Matrix\_Print**(const GrB\_Matrix A, *LAGraph\_PrintLevel* pr, FILE \*f, char \*msg)

LAGraph\_Matrix\_Print displays a matrix in a human-readable form. This method is not meant for saving a GrB\_Matrix to a file; see *LAGraph\_MMWrite* for that method.

**Parameters**

- **A** – [in] matrix to display.
- **pr** – [in] print level.
- **f** – [inout] handle to an open file to write to.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if A or f are NULL.
- **GrB\_NOT\_IMPLEMENTED** – if A has a user-defined type.
- **LAGRAPH\_IO\_ERROR** – if the file could not be written to.

int **LAGraph\_Vector\_Print**(const GrB\_Vector v, *LAGraph\_PrintLevel* pr, FILE \*f, char \*msg)

LAGraph\_Vector\_Print displays a vector in a human-readable form. This method is not meant for saving a GrB\_Vector to a file. To perform that operation, copy the GrB\_Vector into an n-by-1 GrB\_Matrix and use *LAGraph\_MMWrite*.

**Parameters**

- **v** – [in] vector to display.
- **pr** – [in] print level.
- **f** – [inout] handle to an open file to write to.
- **msg** – [inout] any error messages.

**Returns** any GraphBLAS errors that may have been encountered.

**Returns**

- **GrB\_SUCCESS** – if successful.
- **GrB\_NULL\_POINTER** – if v or f are NULL.
- **GrB\_NOT\_IMPLEMENTED** – if v has a user-defined type.
- **LAGRAPH\_IO\_ERROR** – if the file could not be written to.

enum **LAGraph\_PrintLevel**

LAGraph\_PrintLevel: an enum to control how much to print in LAGraph\_\*\_Print methods.

*Values:*

enumerator **LAGraph\_SILENT**

nothing is printed.

enumerator **LAGraph\_SUMMARY**

print a terse summary.

enumerator **LAGraph\_SHORT**

short description, about 30 entries.

enumerator **LAGraph\_COMPLETE**

print the entire contents of the object.

enumerator **LAGraph\_SHORT\_VERBOSE**

short, but with “%.15g” for doubles.

enumerator **LAGraph\_COMPLETE\_VERBOSE**

complete, but “%.15g” for doubles.

## 5.6 Pre-defined semirings

LAGraph adds the following pre-defined semirings. They are created by *LAGr\_Init* or *LAGraph\_Init*, and freed by *LAGraph\_Finalize*.

- **LAGraph\_plus\_first\_T**: Uses the *GrB\_PLUS\_MONOID\_T* monoid and the corresponding *GrB\_FIRST\_T* multiplicative operator:

```
LAGraph_plus_first_int8
LAGraph_plus_first_int16
LAGraph_plus_first_int32
LAGraph_plus_first_int64
LAGraph_plus_first_uint8
LAGraph_plus_first_uint16
LAGraph_plus_first_uint32
LAGraph_plus_first_uint64
LAGraph_plus_first_fp32
LAGraph_plus_first_fp64
```

- **LAGraph\_plus\_second\_T** Uses the *GrB\_PLUS\_MONOID\_T* monoid and the corresponding *GrB\_SECOND\_T* multiplicative operator:

```
LAGraph_plus_second_int8
LAGraph_plus_second_int16
LAGraph_plus_second_int32
LAGraph_plus_second_int64
LAGraph_plus_second_uint8
LAGraph_plus_second_uint16
```

(continues on next page)

(continued from previous page)

```
LAGraph_plus_second_uint32
LAGraph_plus_second_uint64
LAGraph_plus_second_fp32
LAGraph_plus_second_fp64
```

- **LAGraph\_plus\_one\_T**: Uses the *GrB\_PLUS\_MONOID\_T* monoid and the corresponding *GrB\_ONEB\_T* multiplicative operator:

```
LAGraph_plus_one_int8
LAGraph_plus_one_int16
LAGraph_plus_one_int32
LAGraph_plus_one_int64
LAGraph_plus_one_uint8
LAGraph_plus_one_uint16
LAGraph_plus_one_uint32
LAGraph_plus_one_uint64
LAGraph_plus_one_fp32
LAGraph_plus_one_fp64
```

- **LAGraph\_any\_one\_T**: Uses the *GrB\_MIN\_MONOID\_T* for non-boolean types or *GrB\_LOR\_MONOID\_BOOL* for boolean, and the *GrB\_ONEB\_T* multiplicative op.

These semirings are very useful for unweighted graphs, or for algorithms that operate only on the sparsity structure of unweighted graphs:

```
LAGraph_any_one_bool
LAGraph_any_one_int8
LAGraph_any_one_int16
LAGraph_any_one_int32
LAGraph_any_one_int64
LAGraph_any_one_uint8
LAGraph_any_one_uint16
LAGraph_any_one_uint32
LAGraph_any_one_uint64
LAGraph_any_one_fp32
LAGraph_any_one_fp64
```



## EXPERIMENTAL ALGORITHMS

LAGraph includes a set of experimental algorithms and utilities, in the LAGraph/experimental folder. The include file appears in LAGraph/include/LAGraphX.h. These methods are in various states of development, and their C APIs are not guaranteed to be stable. They are not guaranteed to have all of their performance issues resolved. However, they have been tested, debugged, and mostly benchmarked.

New algorithms and utilities can be contributed by placing them in the experimental/algorithm or experimental/utility folder. Tests for new algorithms or utilities should be placed in the experimental/test folder, and benchmark programs that exercise their performance (and typically check results) should be placed in the experimental/benchmark folder.

An simple example algorithm and its test and benchmark is provided, which serves as a template for creating new algorithms:

- **algorithm/LAGraph\_HelloWorld.c** a simple “algorithm” that merely creates a copy of the G->A adjacency matrix.
- **benchmark/helloworld2\_demo.c** a benchmark program illustrates how to write a main program that loads in a graph, calls an algorithm, and checks and prints the result. If any file appears in the benchmark folder with a name ending in \_demo.c, then the CMake script will find it and compile it.
- **benchmark/helloworld\_demo.c** another benchmark program. This one relies on internal utilities. See the file for details.
- **test/test\_HelloWorld.c** a test program, using the acutest test suite. If any file appears in experimental/test with the prefix test\_\*, the CMake script will compile it and include it in the “make test” target.



## INSTALLATION

LAGraph is available at <https://github.com/GraphBLAS/LAGraph>. Be sure to check out the default `stable` branch, or use one of the stable releases. LAGraph requires SuiteSparse:GraphBLAS, available at <https://github.com/DrTimothyAldenDavis/GraphBLAS>.

To compile and install LAGraph, you must first compile and install a recent version of SuiteSparse:GraphBLAS. Place LAGraph and GraphBLAS in the same folder, side-by-side. Compile and (optionally) install SuiteSparse:GraphBLAS (see the documentation in SuiteSparse:GraphBLAS for details). At least on Linux or Mac, if GraphBLAS is not installed system-wide, LAGraph can find it if GraphBLAS appears in the same folder as LAGraph, so you do not need system privileges to use GraphBLAS.

LAGraph includes a *CMakeLists.txt* file that does the bulk of the work to build the package. Also included is a very simple *Makefile* that simplifies the use of *make* for Linux and MacOS. In Linux or Mac, you can use it to run these commands:

```
cd LAGraph
make
make test
```

If you have system admin privileges, you can then install LAGraph:

```
sudo make install
```

On Windows, the *CMakeLists.txt* file can be imported into MS Visual Studio, and LAGraph can be built directly from there.





## ACKNOWLEDGEMENTS

This work is funded in part by:

1. The United States Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. [DM22-0790]
2. The United States Department of Defense under Contract No. FA8650-18-2-7835 and HR0011-18-3-0007
3. NVIDIA, Intel, MIT Lincoln Laboratory, MathWorks, IBM, and Julia Computing.
4. The National Science Foundation (1514406, OAC-1740333, CCF-1629657).
5. The United States Department of Energy, Office of Science, ASCR Contract No. DE-AC02-05CH11231.



## REFERENCES

- Graph algorithms in the language of linear algebra, Jeremy Kepner and John Gilbert, SIAM, 2011.
- Mathematical Foundations of the GraphBLAS, Jeremy Kepner, Peter Aaltonen, David Bader, Aydin Buluc, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, Scott McMillan, Jose Moreira, John D. Owens, Carl Yang, Marcin Zalewski, Timothy Mattson, IEEE High Performance Extreme Computing, 2016
- Design of the GraphBLAS API for C, Aydin Buluc, Tim Mattson, Scott McMillan, Jose Moreira, and Carl Yang Graph Algorithms Building Blocks workshop at IPDPS, 2017
- LAGraph: A Community Effort to Collect Graph Algorithms Built on Top of the GraphBLAS T Mattson, TA Davis, M Kumar, A Buluç, S McMillan, J Moreira, C Yang GrAPL workshop at IPDPS 2019. <https://people.eecs.berkeley.edu/~aydin/LAGraph19.pdf>
- LAGraph: Linear Algebra, Network Analysis Libraries, and the Study of Graph Algorithms, Gabor Szarnyas, David A. Bader, Timothy A. Davis, James Kitchen, Timothy G. Mattson, Scott McMillan, Erik Welch, GrAPL workshop at IPDPS 2021
- Algorithm 1000: SuiteSparse:GraphBLAS: Graph Algorithms in the Language of Sparse Linear Algebra Timothy A. Davis, ACM Trans. Math. Softw., vol 45, no 4, Dec. 2019. <https://doi.org/10.1145/3322125>
- Hoke, Tanner (2022). An Implementation of Fast Graphlet Transform in GraphBLAS. Undergraduate Research Scholars Program, Texas A&M University. <https://hdl.handle.net/1969.1/196609>
- Konduri, Pranav S (2022). An Implementation of the Parallel K-core Decomposition Algorithm in GraphBLAS. Undergraduate Research Scholars Program, Texas A&M University. <https://hdl.handle.net/1969.1/196516>
- Waheed, Abeer (2022). TriPoll in GraphBLAS. Undergraduate Research Scholars Program, Texas A&M University. <https://hdl.handle.net/1969.1/196576>



## EXAMPLE USAGE

Note that this simple example does not check any error conditions.

```
#include "LAGraph.h"

int main (void)
{
    // initialize LAGraph
    char msg [LAGRAPH_MSG_LEN] ;
    LAGraph_Init (msg) ;
    GrB_Matrix A = NULL ;
    GrB_Vector centrality = NULL ;
    LAGraph_Graph G = NULL ;

    // read a Matrix Market file from stdin and create a graph
    LAGraph_MMRead (&A, stdin, msg) ;
    LAGraph_New (&G, &A, LAGraph_ADJACENCY_UNDIRECTED, msg) ;

    // compute the out-degree of every node
    LAGraph_Cached_OutDegree (G, msg) ;

    // compute the pagerank
    int niters = 0 ;
    LAGr_PageRank (&centrality, &niters, G, 0.85, 1e-4, 100, msg) ;

    // print the result
    LAGraph_Vector_Print (centrality, LAGraph_COMPLETE, stdout, msg) ;

    // free the graph, the pagerank, and finish LAGraph
    LAGraph_Delete (&G, msg) ;
    GrB_free (&centrality) ;
    LAGraph_Finalize (msg) ;
}
```

genindex



## G

GRB\_TRY (*C macro*), 10

## L

LAGr\_Betweenness (*C++ function*), 24

LAGr\_BreadthFirstSearch (*C++ function*), 23

LAGr\_ConnectedComponents (*C++ function*), 23

LAGr\_Init (*C++ function*), 5

LAGr\_PageRank (*C++ function*), 25

LAGr\_SampleDegree (*C++ function*), 22

LAGr\_SingleSourceShortestPath (*C++ function*), 23

LAGr\_SortByDegree (*C++ function*), 22

LAGr\_TriangleCount (*C++ function*), 25

LAGr\_TriangleCount\_Method (*C++ enum*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_AutoMethod  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Burkhardt  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Cohen  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Sandia\_LL  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Sandia\_LUI  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Sandia\_ULI  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Method::LAGr\_TriangleCount\_Sandia\_UU  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Presort (*C++ enum*), 26

LAGr\_TriangleCount\_Presort::LAGr\_TriangleCount\_Ascending  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Presort::LAGr\_TriangleCount\_AutoSort  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Presort::LAGr\_TriangleCount\_Descending  
(*C++ enumerator*), 26

LAGr\_TriangleCount\_Presort::LAGr\_TriangleCount\_NoSort  
(*C++ enumerator*), 26

LAGraph\_Boolean (*C++ enum*), 14

LAGraph\_Boolean::LAGraph\_BOOLEAN\_UNKNOWN  
(*C++ enumerator*), 14

LAGraph\_Boolean::LAGraph\_FALSE (*C++ enumera-  
tor*), 14

LAGraph\_Boolean::LAGraph\_TRUE (*C++ enumera-  
tor*), 14

LAGraph\_Cached\_AT (*C++ function*), 16

LAGraph\_Cached\_EMax (*C++ function*), 18

LAGraph\_Cached\_EMin (*C++ function*), 18

LAGraph\_Cached\_InDegree (*C++ function*), 17

LAGraph\_Cached\_IsSymmetricStructure (*C++  
function*), 17

LAGraph\_Cached\_NSelfEdges (*C++ function*), 18

LAGraph\_Cached\_OutDegree (*C++ function*), 17

LAGraph\_CheckGraph (*C++ function*), 19

LAGraph\_Delete (*C++ function*), 16

LAGraph\_DeleteCached (*C++ function*), 16

LAGraph\_DeleteSelfEdges (*C++ function*), 19

LAGraph\_Finalize (*C++ function*), 6

LAGraph\_GetNumThreads (*C++ function*), 6

LAGraph\_Graph (*C++ type*), 13

LAGraph\_Graph\_Print (*C++ function*), 36

LAGraph\_Graph\_struct (*C++ struct*), 13

LAGraph\_Graph\_struct::A (*C++ member*), 13

LAGraph\_Graph\_struct::AT (*C++ member*), 13

LAGraph\_Graph\_struct::emax (*C++ member*), 14

LAGraph\_Graph\_struct::emax\_state (*C++ mem-  
ber*), 14

LAGraph\_Graph\_struct::emin (*C++ member*), 14

LAGraph\_Graph\_struct::emin\_state (*C++ mem-  
ber*), 14

LAGraph\_Graph\_struct::in\_degree (*C++ member*),  
14

LAGraph\_Graph\_struct::is\_symmetric\_structure  
(*C++ member*), 14

LAGraph\_Graph\_struct::kind (*C++ member*), 13

LAGraph\_Graph\_struct::nself\_edges (*C++ mem-  
ber*), 14

LAGraph\_Graph\_struct::out\_degree (*C++ mem-  
ber*), 13

LAGraph\_Init (*C++ function*), 5

LAGraph\_Kind (*C++ enum*), 15

LAGraph\_Kind::LAGraph\_ADJACENCY\_DIRECTED  
(*C++ enumerator*), 15

LAGraph\_Kind::LAGraph\_ADJACENCY\_UNDIRECTED  
(*C++ enumerator*), 15

LAGraph\_Kind::LAGraph\_KIND\_UNKNOWN (C++ *enumerator*), 15

LAGraph\_Matrix\_IsEqual (C++ *function*), 33

LAGraph\_Matrix\_IsEqualOp (C++ *function*), 33

LAGraph\_Matrix\_Print (C++ *function*), 37

LAGraph\_Matrix\_Structure (C++ *function*), 32

LAGraph\_Matrix\_TypeName (C++ *function*), 35

LAGraph\_MMRead (C++ *function*), 29

LAGraph\_MMWrite (C++ *function*), 31

LAGRAPH\_MSG\_LEN (C *macro*), 8

LAGraph\_NameOfType (C++ *function*), 35

LAGraph\_New (C++ *function*), 15

LAGraph\_PrintLevel (C++ *enum*), 37

LAGraph\_PrintLevel::LAGraph\_COMPLETE (C++ *enumerator*), 38

LAGraph\_PrintLevel::LAGraph\_COMPLETE\_VERBOSE (C++ *enumerator*), 38

LAGraph\_PrintLevel::LAGraph\_SHORT (C++ *enumerator*), 38

LAGraph\_PrintLevel::LAGraph\_SHORT\_VERBOSE (C++ *enumerator*), 38

LAGraph\_PrintLevel::LAGraph\_SILENT (C++ *enumerator*), 38

LAGraph\_PrintLevel::LAGraph\_SUMMARY (C++ *enumerator*), 38

LAGRAPH\_RETURN\_VALUES (C *macro*), 7

LAGraph\_Scalar\_TypeName (C++ *function*), 36

LAGraph\_SetNumThreads (C++ *function*), 7

LAGraph\_SizeOfType (C++ *function*), 35

LAGraph\_State (C++ *enum*), 14

LAGraph\_State::LAGraph\_BOUND (C++ *enumerator*), 15

LAGraph\_State::LAGraph\_STATE\_UNKNOWN (C++ *enumerator*), 15

LAGraph\_State::LAGraph\_VALUE (C++ *enumerator*), 15

LAGraph\_TriangleCount (C++ *function*), 21

LAGRAPH\_TRY (C *macro*), 9

LAGraph\_TypeFromName (C++ *function*), 35

LAGraph\_Vector\_IsEqual (C++ *function*), 34

LAGraph\_Vector\_IsEqualOp (C++ *function*), 34

LAGraph\_Vector\_Print (C++ *function*), 37

LAGraph\_Vector\_Structure (C++ *function*), 32

LAGraph\_Vector\_TypeName (C++ *function*), 36

LAGraph\_Version (C++ *function*), 6

LAGraph\_WallClockTime (C++ *function*), 32