

M-Flash: Fast Billion-Scale Graph Computation Using a Bimodal Block Processing Model

Hugo Gualdrón¹(✉), Robson Cordeiro¹(✉), Jose Rodrigues Jr.¹(✉),
Duen Horng (Polo) Chau²(✉), Minsuk Kahng²(✉), and U. Kang³(✉)

¹ University of Sao Paulo, Sao Carlos, SP, Brazil
{gualdron,robson,junio}@icmc.usp.br

² Georgia Institute of Technology, Atlanta, USA
{polo,kahng}@gatech.edu

³ Seoul National University, Seoul, Republic of Korea
ukang@snu.ac.kr

Abstract. Recent graph computation approaches have demonstrated that a single PC can perform efficiently on billion-scale graphs. While these approaches achieve scalability by optimizing I/O operations, they do not fully exploit the capabilities of modern hard drives and processors. To overcome their performance, in this work, we introduce the Bimodal Block Processing (*BBP*), an innovation that is able to boost the graph computation by minimizing the I/O cost even further. With this strategy, we achieved the following contributions: (1) M-Flash, the fastest graph computation framework to date; (2) a flexible and simple programming model to easily implement popular and essential graph algorithms, including the *first* single-machine billion-scale eigensolver; and (3) extensive experiments on real graphs with up to 6.6 billion edges, demonstrating M-Flash’s consistent and significant speedup. The software related to this paper is available at <https://github.com/M-Flash>.

Keywords: Graph algorithms · Graph processing · Graph mining · Complex networks

1 Introduction

Large graphs with *billions* of nodes and edges are increasingly common in many domains and applications, such as in studies of social networks, transportation route networks, citation networks, and many others. Distributed frameworks (find a thorough review in the work of Lu *et al.* [13]) have become popular choices for analyzing these large graphs. However, distributed approaches may not always be the best option, because they can be expensive to build [11], and hard to maintain and optimize.

These potential challenges prompted researchers to create single-machine, billion-scale graph computation frameworks that are well-suited to essential graph algorithms, such as eigensolver, PageRank, connected components and many others. Examples are GraphChi [11] and TurboGraph [5]. Frameworks

in this category define sophisticated processing schemes to overcome challenges induced by limited main memory and poor locality of memory access observed in many graph algorithms. However, when studying previous works, we noticed that despite their sophisticated schemes and novel programming models, they do not optimize for disk operations and data locality, which are the core of performance in graph processing frameworks.

In the context of *single-node*, *billion-scale*, graph processing frameworks, we present **M-Flash**, a novel scalable framework that overcomes critical issues found in existing works. The innovation of M-Flash confers it a performance many times faster than the state of the art. More specifically, our contributions include:

1. **M-Flash Framework & Methodology:** we propose a novel framework named M-Flash that achieves fast and scalable graph computation. M-Flash (<https://github.com/M-Flash>) introduces the Bimodal Block Processing, which significantly boosts computation speed and reduces disk accesses by dividing a graph and its node data into blocks (dense and sparse) to minimize the cost of I/O.
2. **Programming Model:** M-Flash provides a flexible and simple programming model, which supports popular and essential graph algorithms, e.g., PageRank, connected components, and the *first* single-machine eigensolver over billion-node graphs, to name a few.
3. **Extensive Experimental Evaluation:** we compared M-Flash with state-of-the-art frameworks using large graphs, the largest one having 6.6 billion edges (YahooWeb <https://webscope.sandbox.yahoo.com>). M-Flash was consistently and significantly faster than GraphChi [11], X-Stream [15], TurboGraph [5], MMap [12], and GridGraph [19] across all graph sizes. Furthermore, it sustained high speed even when memory was severely constrained, e.g., 6.4X faster than X-Stream, when using 4 GB of Random Access Memory (RAM).

2 Related Works

A typical approach to scalable graph processing is to develop a distributed framework. This is the case of Gbase [7], Powergraph, Pregel, and others [13]. Among these approaches, Gbase is the most similar to M-Flash. Despite the fact that Gbase and M-Flash use a block model, Gbase is distributed and lacks an adaptive edge processing scheme to optimize its performance. Such scheme is the greatest innovation of M-Flash, conferring to it the highest performance among existing approaches, as demonstrated in Sect. 4.

Among the existing works designed for single-node processing, some of them are restricted to SSDs. These works rely on the remarkable low-latency and improved I/O of SSDs compared to magnetic disks. This is the case of TurboGraph [5], which relies on random accesses to the edges — not well supported over magnetic disks. Our proposal, M-Flash, avoids this drawback by avoiding random accesses.

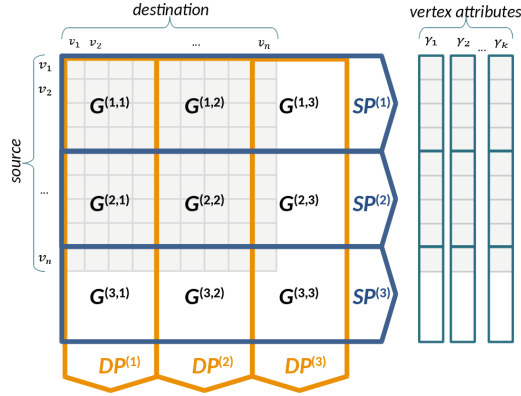


Fig. 1. Organization of edges and vertices in M-Flash. **Edges (left):** example of a graph's adjacency matrix (in light blue color) using 3 logical intervals ($\beta = 3$); $G^{(p,q)}$ is an edge block with source vertices in interval $I^{(p)}$ and destination vertices in interval $I^{(q)}$; $SP^{(p)}$ is a *source-partition* containing all blocks with source vertices in interval $I^{(p)}$; $DP^{(q)}$ is a *destination-partition* containing all blocks with destination vertices in interval $I^{(q)}$. **Vertices (right):** the data of the vertices as k vectors ($\gamma_1 \dots \gamma_k$), each one divided into β logical segments. (Color figure online)

GraphChi [11] was one of the first single-node approaches to avoid random disk/edge accesses, improving the performance over mechanical disks. GraphChi partitions the graph on disk into units called *shards*, requiring a preprocessing step to sort the data by source vertex. GraphChi uses a vertex-centric approach that requires a shard to fit entirely in memory, including both the vertices in the shard and all their edges (in and out). As we demonstrate, this fact makes GraphChi less efficient when compared to our work. M-Flash requires only a subset of the vertex data to be stored in memory.

MMap [12] introduced an interesting approach based on OS-supported mapping of disk data into memory (virtual memory). It allows graph data to be accessed as if they were stored in unlimited memory, avoiding the need to manage data buffering. Our framework uses memory mapping when processing edge blocks but, with an improved engineering, M-Flash consistently outperforms MMap, as we demonstrate.

GridGraph [19] divides the graphs into blocks and processes the edges reusing the vertices' values loaded in main memory (in-vertices and out-vertices). Furthermore, it uses a two-level hierarchical partitioning to increase the performance, dividing the blocks into small regions that fit in cache. When comparing GridGraph with M-Flash, both divide the graph using a similar approach with a two-level hierarchical optimization to boost computation. However, M-Flash adds a bimodal partition model over the block scheme to optimize even more the computation for sparse blocks in the graph.

GraphTwist [18] introduces a 3D cube representation of the graph to add support for multigraphs. The cube representation divides the edges using three

partitioning levels: slice, strip, and dice. These representations are equivalent to the block representation (2D) of GridGraph and M-Flash, with the difference that it adds one more dimension (slice) to organize the edge metadata for multi-graphs. The slice dimension filters the edges' metadata optimizing performance when not all the metadata is required for computation. Additionally, Graph-Twist introduces pruning techniques to remove some slices and vertices that they do not consider relevant in the computation.

M-Flash also draws inspiration from the edge streaming approach introduced by X-Stream's processing model [4, 15], improving it with fewer I/O operations for dense regions of the graph. Edge streaming is a sort of stream processing referring to unrestricted data flows over a bounded amount of buffering. As we demonstrate, this leads to optimized data transfer by means of less I/O and more processing per data transfer.

3 M-Flash

The design of M-Flash considers the fact that real graphs have a varying density of edges; that is, a given graph contains dense regions with many more edges than other regions that are sparse. In the development of M-Flash, and through experimentation with existing works, we noticed that these dense and sparse regions could not be processed in the same way. We also noticed that this was the reason why existing works failed to achieve superior performance. To cope with this issue, we designed M-Flash to work according to two distinct processing schemes: Dense Block Processing (DBP) and Streaming Partition Processing (SPP). For full performance, M-Flash uses a theoretical I/O cost-based scheme to decide the kind of processing to use in face of a given block, which can be dense or sparse. The final approach, which combines DBP and SPP, was named Bimodal Block Processing (BBP).

3.1 Graph Representation in M-Flash

A graph in M-Flash is a directed graph $G = (V, E)$ with vertices $v \in V$ labeled with integers from 1 to $|V|$, and edges $e = (\text{source}, \text{destination})$, $e \in E$. Each vertex has a set of attributes $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_K\}$; edges also might have attributes for specific processings.

Blocks in M-Flash: Given a graph G , we divide its vertices V into β intervals denoted by $I^{(p)}$, where $1 \leq p \leq \beta$. Note that $I^{(p)} \cap I^{(p')} = \emptyset$ for $p \neq p'$, and $\bigcup_p I^{(p)} = V$. Consequently, as shown in Fig. 1, the edges are divided into β^2 blocks. Each block $G^{(p,q)}$ has a *source node interval* p and a *destination node interval* q , where $1 \leq p, q \leq \beta$. In Fig. 1, for example, $G^{(2,1)}$ is the block that contains edges with source vertices in the interval $I^{(2)}$ and destination vertices in the interval $I^{(1)}$. We call this on-disk organization as *partitioning*. Since M-Flash works by alternating one entire block in memory for each running thread, the value of β is automatically determined by the following equation:

$$\beta = \left\lceil \frac{\phi(T + 1)|V|}{M} \right\rceil \tag{1}$$

where the constant 1 refers to the need of one buffer to store the input vertex values that are shared between threads (read-only), ϕ is the amount of data to represent each vertex, T is the number of threads, $|V|$ is the number of vertices, and M is the available RAM. For example, 4 bytes of data per node, 2 threads, a graph with 2 billion nodes, and for 1 GB RAM, $\beta = \lceil (4 \times (2 + 1) \times (2 \times 10^9)) / (2^{30}) \rceil = 23$, thus requiring $23^2 = 529$ blocks. The number of threads enters the equation because all the threads access the same block to avoid multiple seeks on disk, and they use an exclusive memory buffer to store the vertex data processed (one buffer per thread), so to prevent “race” conditions.

3.2 The M-Flash Processing Model

This section presents our proposed processing model. We first describe the two strategies targeted at processing dense or sparse blocks. Next, we present the novel cost-based optimization used to determine the best processing strategy.

Dense Block Processing (DBP): Figure 2 illustrates the DBP; notice that vertex intervals are represented by vertical (Source I) and horizontal (Destination I) vectors. After partitioning the graph into *blocks*, we process them in a vertical zigzag order, as illustrated. There are three reasons for this order: (1) we store the computation results in the destination vertices; so, we can “pin” a

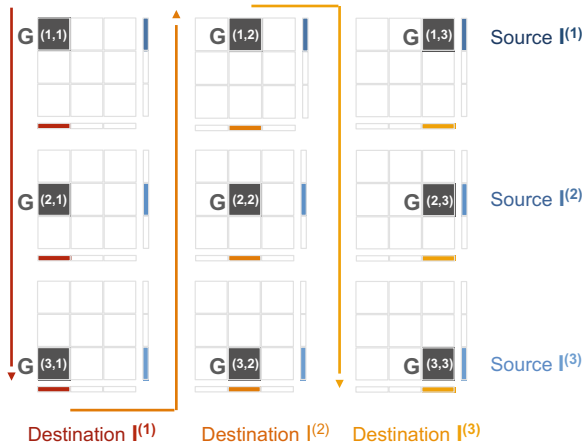


Fig. 2. M-Flash’s computation schedule for a graph with 3 intervals. Vertex intervals are represented by vertical (Source I) and horizontal (Destination I) vectors. Blocks are loaded into memory, and processed in a vertical zigzag manner, indicated by the sequence of red, orange and yellow arrows. This enables the reuse of input, as when going from $G^{(3,1)}$ to $G^{(3,2)}$, M-Flash reuses source node interval $I^{(3)}$, which reduces data transfer from disk to memory. (Color figure online)

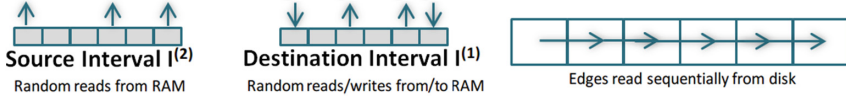


Fig. 3. Example of DBP I/O operations to process the *dense* block $G^{(2,1)}$.

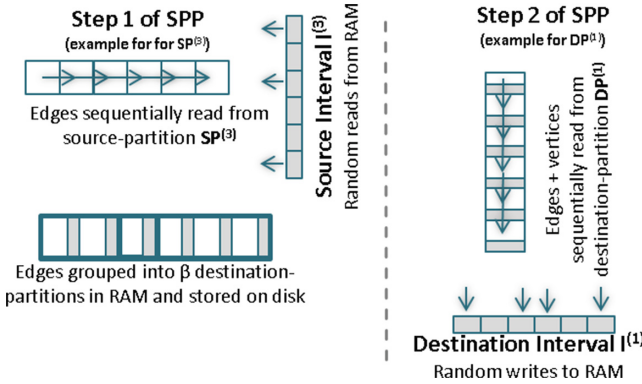


Fig. 4. I/O operations for SPP taking $SP^{(3)}$ and $DP^{(1)}$ as illustrative examples. Step 1: the edges of *source-partition* $SP^{(3)}$ are sequentially read and combined with the values of their source vertices from $I_{(3)}$. Next, edges are grouped by destination, and written to β files, one for each *destination partition*. Step 2: the files corresponding to *destination-partition* $DP^{(1)}$ are sequentially processed according to a given desired computation, with results written to destination vertices in $I_{(1)}$.

destination interval (e.g., $I^{(1)}$) and process all the vertices that are sources to this destination interval (see the red vertical arrow); (2) using this order leads to fewer reads because the attributes of the destination vertices (horizontal vectors in the illustration) only need to be read once, regardless of the number of source intervals. (3) after reading all the blocks in a column, we take a “U turn” (see the orange arrow) to benefit from the fact that the data associated with the previously-read source interval is already in memory.

Within a block, besides loading the attributes of the source and destination intervals of vertices into RAM, the corresponding edges $e = \langle source, destination, edge\ properties \rangle$ are sequentially read from disk, as explained in Fig. 3. These edges, then, are processed using a user-defined function so to achieve the desired computation. After all blocks in a column are processed, the updated attributes of the destination vertices are written to disk.

Streaming Partition Processing (SPP): The performance of DBP decreases for graphs with sparse blocks; this is because, for a given block, we have to read more data from the source intervals of vertices than from the very blocks of edges. In such cases, SPP processes the graph using partitions instead of blocks. A graph *partition* is a set of *blocks* sharing the same *source node interval* – a line in the logical partitioning, or, similarly, a set of *blocks* sharing the same

destination node interval – a column in the logical partitioning. Formally, a *source-partition* $SP^{(p)} = \bigcup_q G^{(p,q)}$ contains all the blocks with edges having source vertices in the interval $I^{(p)}$; a *destination-partition* $DP^{(q)} = \bigcup_p G^{(p,q)}$ contains all the blocks with edges having destination vertices in the interval $I^{(q)}$. For example, in Fig. 1, $SP^{(1)}$ is the union of blocks $G^{(1,1)}$, $G^{(1,2)}$, and $G^{(1,3)}$; meanwhile, $DP^{(3)}$ is the union of blocks $G^{(1,3)}$, $G^{(2,3)}$, and $G^{(3,3)}$. In a graph, hence, there are β *source-partitions* and β *destination-partitions*.

Considering the graph organized into partitions, SPP takes two steps (see Fig. 4). In the first step, for a given *source-partition* $SP^{(p)}$, it loads the values of the vertices of the corresponding interval $I^{(p)}$; next, it reads the edges of the partition $SP^{(p)}$ sequentially from disk, storing them in a buffer together with their source-vertex values. At this point, it sorts the buffer in memory, grouping the edges by destination. Finally, it stores the edges on disk into β files, one for each of the β *destination-partitions*. This processing is performed for each *source-partition* $SP^{(p)}$, $1 \leq p \leq \beta$, so to iteratively build the β *destination-partitions*.

In the second step, after processing the β *source-partitions* (each with β blocks), it is possible to read the β files according to their destinations, so to logically “build” β *destination-partitions* $DP^{(q)}$, $1 \leq q \leq \beta$, each one containing edges together with their source-vertex values. For each *destination-partition* $DP^{(q)}$, we read the vertices of interval $I^{(q)}$; next, we sequentially read the edges, processing their values through a user-defined function. This function uses the properties of the vertices and of the edges to perform specific computations whose results will update the vertices. Finally, SPP stores the updated vertices of interval $I^{(q)}$ back on disk.

Bimodal Block Processing (BBP): Schemes *DBP* and *SPP* improve the performance in complementary circumstances. But, *How can we decide which processing scheme to use when we are given a graph block to process?* To answer this question, we join *DBP* and *SPP* into a single scheme – the Bimodal Block Processing (BBP). The combined scheme uses the theoretical I/O cost model proposed by Aggarwal and Vitter [1] to decide for *SBP* or *SPP*. In this model, the I/O cost of an algorithm is equal to the number of *disk blocks* with size B transferred between disk and memory plus the number of non-sequential reads (seeks). Since we use this model to choose the scheme with the smaller cost, we need to algebraically determine the cost of each scheme, as follows.

For processing a graph $G = \{V, E\}$, *DBP* performs the following operations: it reads the $|V|$ vertices β times and it writes the $|V|$ vertices once; it also reads the $|E|$ edges once – *disk blocks* of size B , vertices and edges with constant sizes omitted from the equation for simplification. β^2 seeks are necessary because the edges are read sequentially. Hence, the I/O cost for *DBP* is given by:

$$\mathcal{O}(\text{DBP}(G)) = \mathcal{O}\left(\frac{(\beta + 1)|V| + |E|}{B} + \beta^2\right) \quad (2)$$

In turn, *SPP* initially reads the $|V|$ source vertices and the $|E|$ edges; then, still in its first step, it sorts (simple shuffling) the $|E|$ edges grouping them by

destination into a set of edges and vertices $|\hat{E}|$, writing them to disk. In its second step, it reads the \hat{E} edges/vertices to perform the update operation, writing $|V|$ destination vertices back to disk. The I/O cost for *SPP* comes to be:

$$\mathcal{O}(\text{SPP}(G)) = \mathcal{O}\left(\frac{2|V| + |E| + 2|\hat{E}|}{B} + \beta\right) \tag{3}$$

Equations 2 and 3 define the I/O cost for one processing iteration over the whole graph G . However, in order to decide in relation to the graph blocks, we are interested in the costs of Eqs. 2 and 3 divided by the number of graph blocks β^2 . The result, after the appropriate algebra, reduces to Eqs. 4 and 5.

$$\mathcal{O}(\text{DBP}(G^{(p,q)})) = \mathcal{O}\left(\frac{\vartheta\phi(1 + 1/\beta) + \xi\psi}{B}\right) \tag{4}$$

$$\mathcal{O}(\text{SPP}(G^{(p,q)})) = \mathcal{O}\left(\frac{2\vartheta\phi/\beta + 2\xi(\phi + \psi) + \xi\psi}{B}\right) \tag{5}$$

where ξ is the number of edges in $G^{(p,q)}$, ϑ is the number of vertices in the interval, and ϕ and ψ are, respectively, the number of bytes to represent a vertex and an edge e . Once we have the costs per graph block of DBP and SPP, we can decide between one and the other by simply analyzing the ratio SPP/DBP:

$$\mathcal{O}\left(\frac{\text{SPP}}{\text{DBP}}\right) = \mathcal{O}\left(\frac{1}{\beta} + \frac{2\xi}{\vartheta} \left[1 + \frac{\psi}{\phi}\right]\right) \tag{6}$$

This ratio leads to the final decision equation:

$$\text{BlockType}(G^{(p,q)}) = \begin{cases} \text{sparse,} & \text{if } \mathcal{O}\left(\frac{\text{SPP}}{\text{DBP}}\right) < 1 \\ \text{dense,} & \text{otherwise} \end{cases} \tag{7}$$

We apply Eq. 6 to select the best option according to Eq. 7. With this scheme, BBP is able to select the best processing scheme for each graph block. In Sect. 4, we demonstrate that this procedure yields a performance superior than the current state-of-the-art frameworks.

Algorithm 1. *MAlgorithm*: Algorithm Interface for coding in M-Flash

initialize (Vertex v)
process (Vertex u , Vertex v , EdgeData $data$)
gather (Accum v_1 , Accum v_2 , Accum v_{out})
apply (Vertex v)

Algorithm 2. PageRank in M-Flash

degree(v): = out degree for Vertex v
initialize(v): $v.value = 0$
process($u, v, data$): $v.value += u.value / degree(u)$
gather(v_1, v_2, v_{out}): $v_{out} = v_1 + v_2$
apply(v): $v.value = 0.15 + 0.85 * v.value$

Algorithm 3. Algorithm M-Flash

Input: Graph $G(V, E)$ and vertex attributes γ
Input: user-defined *MAlgorithm* program
Input: memory size M and number of iterations *iter*
Output: vector v with vertex results

- 1: set ϕ from γ attributes, and β using Eq. 1. $\vartheta = |V|/\beta$
- 2: execute graph preprocessing and *partitioning*
- 3: **for** $i = 1$ to *iter* **do**
- 4: execute the first step of *SPP* (Fig. 4) to process the sparse source-partitions
- 5: **for** $q = 1$ to β **do**
- 6: load vertex values of destination interval $I^{(q)}$
- 7: initialize $I^{(q)}$ of v using *MAlgorithm.initialize*
- 8: **if** there is a sparse destination-partition associated with $I^{(q)}$ **then**
- 9: **for each** edge
- 10: invoke *MAlgorithm.process* storing results in vector v
- 11: **if** q is odd **then**
- 12: partition-order = $\{1$ to $\beta\}$
- 13: **else**
- 14: partition-order = $\{\beta$ to $1\}$
- 15: **for** $p = \{\text{partition-order}\}$ **do**
- 16: **if** $G^{(p,q)}$ is dense **then**
- 17: load vertex values of source interval $I^{(p)}$
- 18: **for each** edge in $G^{(p,q)}$
- 19: invoke *MAlgorithm.process* storing results in vector v
- 20: invoke *MAlgorithm.gather* for $I^{(q)}$ of v
- 21: invoke *MAlgorithm.apply* for $I^{(q)}$ of v
- 22: store interval $I^{(q)}$ of vector v

3.3 Programming Model in M-Flash

The M-Flash’s computational model, which we named *MAlgorithm* (short for *Matrix Algorithm Interface*) is shown in Algorithm 1. Since *MAlgorithm* is a vertex-centric model, it stores computation results in the destination vertices, allowing for a vast set of iterative computations, such as PageRank, Random Walk with Restart, Weakly Connected Components, Sparse Matrix Vector Multiplication, Eigensolver, and Diameter Estimation.

The *MAlgorithm* interface has four operations: **initialize**, **process**, **gather**, and **apply**. The *initialize* operation, optionally, loads the initial value of each destination vertex; the *process* operation receives and processes the data from incoming edges (neighbors) – this is where the desired processing occurs; the *gather* operation joins the results from the multiple threads so to consolidate a single result; finally, the *apply* operation is able to perform finalizing operations, such as normalization – apply is optional.

3.4 System Design and Implementation

M-Flash starts by preprocessing an input graph dividing the edges into β partitions and counting the number of edges per logical block (β^2 blocks), at the same time that the blocks are classified as sparse or dense using Eq. 7. Note that M-Flash does *not* sort the edges during preprocessing, it simply divides them into β^2 blocks, $\beta^2 \ll |V|$. In a second preprocessing, M-Flash processes the graph according to the organization given by the concept of *source-partition* as seen in Sect. 3.2. At this point, blocks are only a logical organization, while partitions are physical. The *source-partitions* are read and, whenever a dense block is found, the corresponding edges are extracted from the partition and a file is created for this block in preparation for DBP; the remaining edges in the *source-partition* will be ready for processing using SPP. Notice that, after the second preprocessing, the logical blocks classified as dense, are materialized into physical files. The total I/O cost for preprocessing is $\frac{4|E|}{B}$, where B is the size of each block transferred between disk and memory. Algorithm 3 shows the pseudo-code of M-Flash.

4 Evaluation

We compare M-Flash (<https://github.com/M-Flash>) with multiple state-of-the-art approaches: GraphChi, TurboGraph, X-Stream, MMap, and GridGraph. For a fair comparison, we used the experimental setups recommended by the respective authors. GridGraph did not publish nor share its code, so the comparison is based on the results reported in its publication. We omit the comparison with GraphTwist because it is not accessible and its published results are based on a hardware that is less powerful than ours. We use four graphs at different scales (See Table 1), and we compare the runtimes of all approaches for two well-known essential algorithms PageRank (Subsect. 4.2) and Weakly Connected Components (Subsect. 4.3). To demonstrate how M-Flash generalizes to more algorithms, we implemented the Lanczos algorithm (with *selective orthogonalization*), which is one of the most computationally efficient approaches to computing eigenvalues and eigenvectors [8] (Subsect. 4.4). To the best of our knowledge, M-Flash provides the **first design and implementation** of Lanczos that can handle graphs with more than one billion nodes. Next, in Subsect. 4.5, we show that M-Flash maintains its high speed even when the machine has little RAM (including extreme cases, like 4 GB), in contrast to the other methods. Finally, through a theoretical analysis of I/O, we show the reasons for the performance increase using the BBP strategy (Subsect. 4.6).

4.1 Experimental Setup

All experiments ran on a standard personal computer equipped with a four-core Intel i7-4500U CPU (3 GHz), 16 GB RAM, and 1 TB 540-MB/s (max) SSD disk. Note that M-Flash does *not* require an SSD to run, which is not the case for

Table 1. Graph datasets used in our experiments.

Graph	Nodes	Edges	Size
LiveJournal [2]	4,847,571	68,993,773	Small
Twitter [10]	41,652,230	1,468,365,182	Medium
YahooWeb	1,413,511,391	6,636,600,779	Large
R-Mat (Synthetic graph)	4,000,000,000	12,000,000,000	Large

all frameworks, like TurboGraph. We used an SSD, nevertheless, to make sure that all methods can perform at their best. Table 1 shows the datasets used in our experiments. GraphChi, X-Stream, MMap, and M-Flash ran on Linux Ubuntu 14.04 (x64). TurboGraph ran on Windows (x64). All the reported times correspond to the average time of three **cold** runs, that is, with all caches and buffers purged between runs to avoid any potential advantage due to caching or buffering.

Table 2. Runtime (in seconds) with 8 GB of RAM. The symbol “-” indicates that the corresponding system failed to process the graph or the information is not available in the respective papers.

	GraphChi	X-Stream	TurboGraph	MMap	GridGraph	M-Flash
PageRank						
LiveJournal (10 iter.)	33.1	10.5	7.9	18.2	6.4	<u>5.3</u>
Twitter (10 iter.)	1,199	962	241	186	269	<u>173</u>
YahooWeb (1 iter.)	642	668	628	1,245	235.95	<u>195</u>
R-Mat (1 iter.)	2,145	1,360	-	-	-	<u>745</u>
Connected Components						
LiveJournal (Union Find)	3.2	5.7	4.4	10.7	4.4	<u>1.3</u>
Twitter (Union Find)	70	1,038	128	45	287	<u>25</u>
YahooWeb (WCC - 1 iter.)	668	889	-	-	-	<u>125</u>
R-Mat (WCC - 1 iter.)	3,334	2,167.63	-	-	-	<u>663.17</u>

4.2 PageRank

Table 2 presents the PageRank runtime of all the methods, as discussed next.

LiveJournal (small graph): Since the whole graph fits in RAM, all approaches finish in seconds. Still, M-Flash was the fastest, up to 6X faster than GraphChi, 3X than MMap, and 2X than X-Stream.

Twitter (medium graph): The edges of this graph do not fit in RAM (it requires 11.3 GB) but its node vectors do. M-Flash had a similar performance if compared to MMap, however, MMap is not a generic framework, rather it is based on dedicated implementations, one for each algorithm. Still, M-Flash was faster. In comparison to GraphChi and X-Stream, the related works that offer generic programming models, M-Flash was the fastest, 5.5X and 7X faster, respectively.

YahooWeb (large graph): For this billion-node graph, neither its edges nor its node vectors fit in RAM; this challenging situation is where M-Flash has notably outperformed the other methods. The results of Table 2 confirm this claim, showing that M-Flash provides a speed that is 3X to 6.3X faster than those of the other approaches.

R-Mat (Synthetic large graph): For our big graph, we compared only GraphChi, X-Stream, and M-Flash because TurboGraph and MMap require indexes or auxiliary files that exceed our current disk capacity. GridGraph was not considered in the comparison because its paper does not provide information about R-Mat graphs with a similar scale. Table 2 shows that M-Flash is 2X and 3X faster than X-Stream and GraphChi respectively.

4.3 Weakly Connected Components (WCC)

When there is enough memory to store all the vertex data, the *Union Find* algorithm [16] is the best option to find all the WCCs in one single iteration. Otherwise, with memory limitations, an iterative algorithm produces identical solutions. Hence, in this round of experiments, we use Algorithm *Union Find* to solve WCC for the small and medium graphs, whose vertices fit in memory; and we use an iterative algorithm for the YahooWeb graph.

Table 2 shows the runtimes for the LiveJournal and Twitter graphs with 8 GB RAM; all approaches use Union Find, except X-Stream. This is because of the way that X-Stream is implemented, which handles only iterative algorithms. In the WCC problem, M-Flash is again the fastest method with respect to the entire experiment: for the LiveJournal graph, M-Flash is 3x faster than GraphChi, 4.3X than X-Stream, 3.3X than TurboGraph, and 8.2X than MMap. For the Twitter graph, M-Flash's speed is 2.8X faster than GraphChi, 41X than X-Stream, 5X than TurboGraph, 2X than MMap, and 11.5X than GridGraph.

In the results of the YahooWeb graph, one can see that M-Flash was significantly faster than GraphChi, and X-Stream. Similarly to the PageRank results, M-Flash is pronouncedly faster: 5.3X faster than GraphChi, and 7.1X than X-Stream.

4.4 Spectral Analysis Using the Lanczos Algorithm

Eigenvalues and eigenvectors are at the heart of numerous algorithms, such as singular value decomposition (SVD) [3], spectral clustering, triangle counting [17], and tensor decomposition [9]. Hence, due to its importance, we demonstrate M-Flash over the *Lanczos algorithm*, a state-of-the-art method for eigen computation. We implemented it using method *Selective Orthogonalization (LSO)*. To the best of our knowledge, M-Flash provides the **first design and implementation** that can handle Lanczos for graphs with more than one billion nodes. Different from the competing works, M-Flash provides functions for basic vector operations using secondary memory. Therefore, for the YahooWeb graph, we are not able to compare it with the other frameworks using only 8 GB of memory.

To compute the top 20 eigenvectors and eigenvalues of the YahooWeb graph, one iteration of *LSO* over M-Flash takes 737s when using 8 GB of RAM. For a comparative panorama, to the best of our knowledge, the closest comparable result of this computation comes from the HEigen system [6], at 150s for one iteration; note however that, it was for a much smaller graph with 282 million edges (23X fewer edges), using a *70-machine* Hadoop cluster, while our experiment with M-Flash used a single personal computer and a much larger graph.

4.5 Effect of Memory Size

Since the amount of memory strongly affects the computation speed of single-node graph processing frameworks, here, we study the effect of memory size. Figure 5 summarizes how all approaches perform under 4 GB, 8 GB, and 16 GB of RAM when running one iteration of PageRank over the YahooWeb graph. M-Flash continues to run at the highest speed even when the machine has very little RAM, 4 GB in this case. Other methods tend to slow down. In special, MMap does not perform well due to *thrashing*, a situation when the machine spends a lot of time on mapping disk-resident data to RAM or unmapping data from RAM, slowing down the overall computation. For 8 GB and 16 GB, respectively, M-Flash outperforms all the competitors for the most challenging graph, the YahooWeb. Notice that all the methods, but for M-Flash and X-Stream, are strongly influenced by restrictions in memory size; according to our analyses, this is due to the higher number of data transfers needed by the other methods when not all the data fit in the memory. Despite that X-Stream worked efficiently for any memory setting, it still has worse performance if compared to M-Flash because it demands three full disk scans in every case – actually, the innovations of M-Flash, as presented in Sect. 3, were designed to overcome such problem.

4.6 Theoretical (I/O) Analysis

Following, we show the theoretical scalability of M-Flash when we reduce the available memory at the same time that we demonstrate why the performance of M-Flash improves when we combine DBP and SPP into BBP, instead of using

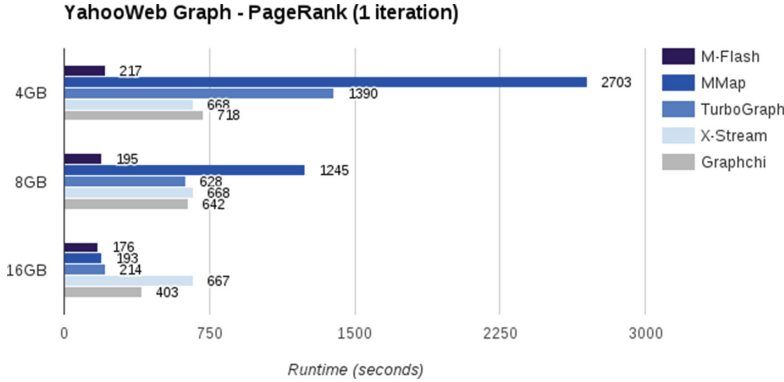


Fig. 5. Runtime comparison for PageRank over the YahooWeb graph. M-Flash is significantly faster than all the state-of-the-art competitors for three different memory settings, 4 GB, 8 GB, and 16 GB.

DBP or SPP alone. Here, we use a measure that we named *t-cost*; 1 unit of *t-cost* corresponds to three operations, one reading of the vertices, one writing of the vertices, and one reading of the edges. In terms of computational complexity, *t-cost* is defined as follows:

$$t\text{-cost}(G(E, V)) = 2|V| + |E| \tag{8}$$

Notice that this cost considers that reading and writing the vertices have the same cost; this is because the evaluation is given in terms of computational complexity. For more details, please refer to the work of McSherry *et al.* [14], who draws the basis of this kind of analysis.

We measure the *t-cost* metric to analyze the theoretical scalability for processing schemes *DBP* only, *SPP* only, and *BBP*. We perform these analyses using MatLab simulations that were validated empirically. We considered the characteristics of the three datasets used so far, LiveJournal, Twitter, and YahooWeb. For each case, we calculated the *t-cost* (*y*-axis) as a function of the available memory (*x*-axis), which, as we have seen, is the main constraint for graph processing frameworks.

Figure 6 shows that, for all the graphs, *DBP*-only processing is the least efficient when memory is reduced; however, when we combine *DBP* (for dense region processing) and *SPP* (for sparse region processing) into *BBP*, we benefit from the best of both worlds. The result corresponds to the best performance, as seen in the charts. Figure 7 shows the same simulated analysis – *t-cost* (*y*-axis) in function of the available memory (*x*-axis), but now with an extra variable: the density of hypothetical graphs, which is assumed to be uniform in each analysis. Each plot, from (a) to (d) considers a different density in terms of average vertex degree, respectively, 3, 5, 10, and 30. In each plot, there are two curves, one corresponding to *DBP*-only, and one for *SPP*-only; and, in dark blue, we depict the behavior of M-Flash according to combination *BBP*. Notice that, as

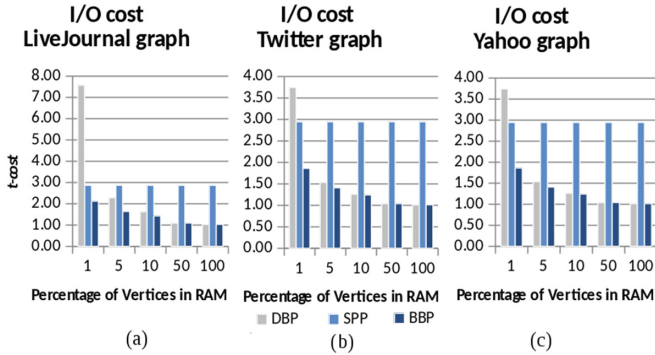


Fig. 6. I/O cost using *DBP*, *SPP*, and *BBP* for LiveJournal, Twitter and YahooWeb Graphs using different memory sizes. *BBP* model always performs fewer I/O operations on disk for all memory configurations.

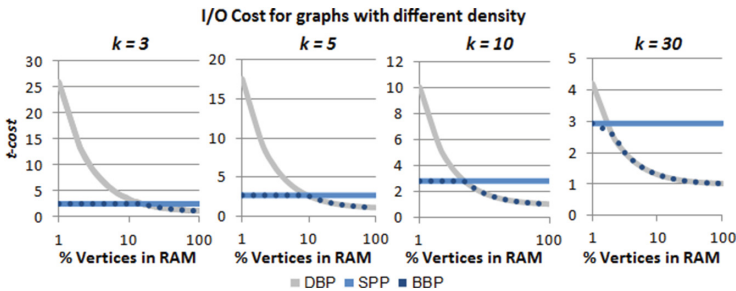


Fig. 7. I/O cost using *DBP*, *SPP*, and *BBP* for a graph with average degree (density) $k = \{3, 5, 10, 30\}$, where $|E| \approx k|V|$, and varying amount of memory (Color figure online)

the amount of memory increases, so does the performance of *DBP*, which takes less and less time to process the whole graph (decreasing curve). *SPP*, in turn, has a steady performance, as it is not affected by the amount of memory (light blue line). In dark blue, one can see the performance of *BBP*; that is, which kind of processing will be chosen by Eq. 7 at each circumstance. For sparse graphs, Figs. 7(a) and (b), *SPP* answers for the greater amount of processing; while the opposite is observed in denser graphs, Figs. 7(c) and (d), when *DBP* defines almost the entire dark blue line of the plot.

These results show that the graph processing must take into account the density of the graph at each moment (block) so to choose the best strategy. It also explains why M-Flash improves the state of the art. It is *important* to note that no former algorithm considered the fact that most graphs present varying density of edges (dense regions with many more edges than other regions that are sparse). Ignoring this fact leads to a decreased performance in the form of

a higher number of data transfers between memory and disk, as we empirically verified in the former sections.

4.7 Preprocessing Time

Table 3 shows the preprocessing times for each graph using 8 GB of RAM. As one can see, M-Flash has a competitive preprocessing runtime. It reads and writes two times the entire graph on disk, which is the third best performance, after MMap and X-Stream. GridGraph and GraphTwist, in turn, demand a preprocessing that divides the graph using blocks in a way similar to M-Flash. We did not compare preprocessing with these frameworks because, as already discussed, we do not have their source code. Despite the extra preprocessing time required by M-Flash – if compared to MMap and X-Stream, the total processing time (preprocessing + *processing with only one iteration*) for algorithms PageRank and WCC over the YahooWeb graph, is of 1,460s and 1,390s, still, 29% and 4% better than the total time of MMap and X-Stream respectively. Note that the algorithms are iterative and M-Flash needs only one iteration to overcome its competitors.

Table 3. Preprocessing time (seconds).

	LiveJournal	Twitter	YahooWeb	R-Mat
GraphChi	23	511	2,781	7,440
X-Stream	5	131	865	2,553
TurboGraph	18	582	4,694	-
MMap	17	372	636	-
M-Flash	10	206	1,265	4,837

5 Conclusions

We proposed M-Flash, a *single-machine, billion-scale* graph computation framework that uses a block partition model to optimize the disk I/O. M-Flash uses an innovative design that takes into account the variable density of edges observed in the different blocks of a graph. Its design uses Dense Block Processing (DBP) when the block is dense, and Streaming Partition Processing (SPP) when the block is sparse. In order to take advantage of both worlds, it uses the combination of DBP and SPP according to the Bimodal Block Processing (BBP) scheme, which is able to analytically determine whether a block is dense or sparse, so to trigger the appropriate processing. To date, our proposal is the first framework that considers a bimodal approach for I/O minimization, a fact that, as we demonstrated, granted M-Flash the best performance compared to the state of the art (GraphChi, X-Stream, TurboGraph, MMap, and GridGraph); notably, even when memory is severely limited.

The findings observed in the design of M-Flash are a step further in determining an ultimate graph processing paradigm. We expect the research in this field to consider the criterion of block density as a mandatory feature in any such framework, consistently advancing the research on high-performance processing.

Acknowledgments. This work received support from Brazilian agencies CNPq (grant 444985/2014-0), Fapesp (grants 2016/02557-0, 2014/21483-2), and Capes; from USA agencies NSF (grants IIS-1563816, TWC-1526254, IIS-1217559), and GRFP (grant DGE-1148903); and Korean (MSIP) agency IITP (grant R0190-15-2012).

References

1. Aggarwal, A., Vitter, J.: The input/output complexity of sorting and related problems. *Commun. ACM* **31**, 1116–1127 (1988)
2. Backstrom, L., Huttenlocher, D., Kleinberg, J., Lan, X.: Group formation in large social networks: membership, growth, and evolution. In: *KDD*, pp. 44–54 (2006)
3. Berry, M.: Large-scale sparse singular value computations. *Int. J. High Perform. Comput. Appl.* **6**(1), 13–49 (1992)
4. Cheng, J., Liu, Q., Li, Z., Fan, W., Lui, J., He, C.: Venus: vertex-centric streamlined graph computation on a single PC. In: *IEEE International Conference on Data Engineering*, pp. 1131–1142 (2015)
5. Han, W.S., Lee, S., Park, K., Lee, J.H., Kim, M.S., Kim, J., Yu, H.: Turbograph: a fast parallel graph engine handling billion-scale graphs in a single PC. In: *KDD*, pp. 77–85 (2013)
6. Kang, U., Meeder, B., Papalexakis, E., Faloutsos, C.: Heigen: spectral analysis for billion-scale graphs. *IEEE TKDE* **26**(2), 350–362 (2014)
7. Kang, U., Tong, H., Sun, J., Lin, C.Y., Faloutsos, C.: Gbase: an efficient analysis platform for large graphs. *VLDB J.* **21**(5), 637–650 (2012)
8. Kang, U., Tsourakakis, C., Faloutsos, C.: Pegasus: a peta-scale graph mining system implementation and observations. In: *ICDM*, pp. 229–238. *IEEE* (2009)
9. Kolda, T., Bader, B.: Tensor decompositions and applications. *SIAM Rev.* **51**(3), 455–500 (2009)
10. Kwak, H., Lee, C., Park, H., Moon, S.: What is twitter, a social network or a news media? In: *WWW*, pp. 591–600. *ACM* (2010)
11. Kyrola, A., Blelloch, G., Guestrin, C.: Graphchi: large-scale graph computation on just a PC. In: *OSDI*, pp. 31–46. *USENIX Association* (2012)
12. Lin, Z., Kahng, M., Sabrin, K., Chau, D.H., Lee, H., Kang, U.: Mmap: fast billion-scale graph computation on a PC via memory mapping. In: *BigData* (2014)
13. Lu, Y., Cheng, J., Yan, D., Wu, H.: Large-scale distributed graph computing systems: an experimental evaluation. *VLDB* **8**(3), 281–292 (2014)
14. McSherry, F., Isard, M., Murray, D.G.: Scalability! but at what cost. In: *HotOS* (2015)
15. Roy, A., Mihailovic, I., Zwaenepoel, W.: X-stream: edge-centric graph processing using streaming partitions. In: *SOSP*, pp. 472–488. *ACM* (2013)
16. Tarjan, R.E., van Leeuwen, J.: Worst-case analysis of set union algorithms. *J. ACM* **31**(2), 245–281 (1984)
17. Tsourakakis, C.: Fast counting of triangles in large real networks without counting: algorithms and laws. In: *ICDM*, pp. 608–617. *IEEE* (2008)

18. Zhou, Y., Liu, L., Lee, K., Zhang, Q.: Graphtwist: fast iterative graph computation with two-tier optimizations. *Proc. VLDB Endowment* **8**(11), 1262–1273 (2015)
19. Zhu, X., Han, W., Chen, W.: Gridgraph: large-scale graph processing on a single machine using 2-level hierarchical partitioning. In: *USENIX ATC 2015*, pp. 375–386. USENIX Association (2015)