

# Massive Streaming Data Analytics: A Case Study with Clustering Coefficients

David Ediger Karl Jiang Jason Riedy David A. Bader  
Georgia Institute of Technology  
Atlanta, GA, USA

**Abstract**—We present a new approach for parallel massive graph analysis of streaming, temporal data with a dynamic and extensible representation. Handling the constant stream of new data from health care, security, business, and social network applications requires new algorithms and data structures. We examine data structure and algorithm trade-offs that extract the parallelism necessary for high-performance updating analysis of massive graphs. Static analysis kernels often rely on storing input data in a specific structure. Maintaining these structures for each possible kernel with high data rates incurs a significant performance cost. A case study computing clustering coefficients on a general-purpose data structure demonstrates incremental updates can be more efficient than global recomputation. Within this kernel, we compare three methods for dynamically updating local clustering coefficients: a brute-force local recalculation, a sorting algorithm, and our new approximation method using a Bloom filter. On 32 processors of a Cray XMT with a synthetic scale-free graph of  $2^{24} \approx 16$  million vertices and  $2^{29} \approx 537$  million edges, the brute-force method processes a mean of over 50 000 updates per second and our Bloom filter approaches 200 000 updates per second.

## I. INTRODUCTION

The data deluge from a wide range of application domains from business and finance to computational biology and computer security requires development of new analysis tools and algorithms. To keep pace with the data, these tools must analyze the resulting interaction networks and graphs as data arrives in high-volume streams rather than in static snapshots. With the Facebook user base containing over 350 million people [12], Twitter boasting more than four billion tweets [21], and an estimated hundreds of millions of blogs on the Internet [23], the massive graph data sets must be analyzed faster than ever before. These massive networks store a wealth of community and relationship information. Data sets from the literature are often constructed statically from web crawls of a particular domain [3], email correspondence between colleagues [19], patent and literature citations [15], and biological networks [16]. Analysis of these data sets is carried out in a single-shot fashion. Changes to the input data induce global recomputation. Here we investigate another approach, computing incremental updates, along with the data structures necessary and assumptions useful for achieving high-performance analysis.

Current large graph analysis tools like Pajek [5] contain three phases. First, the tools preprocess data into appropriate data structures. Second, a graph kernel analyzes the data. Finally, the tools post-process and store analysis results for later presentation. These tools calculate static graph properties; any dynamic use assumes the properties change slowly relative to execution time. However, there is increasing interest in the temporal properties of the dynamic data set. With millions of users and billions of messages in a social network, even the preprocessing time is much larger than the time between potentially large graph changes. Repeatedly processing graphs as snapshots of a process cannot keep up with the data rates.

In this paper, we tackle these problems with a new computational approach for the analysis of complex graphs and networks with billions of vertices based on the streaming input of spatio-temporal data. Our approach accumulates as much of the recent graph data as possible in main memory. Once the reserved memory is full, older or uninteresting edges are aged off and removed. We update analytical kernels after each new edge or block of edges and attempt to detect significant changes in these metrics. We refer to this new approach as *massive streaming data analytics*.

To accommodate a stream of edge data, we present a new, extensible data structure for massive graphs: STINGER (Spatio-Temporal Interaction Networks and Graphs (STING) Extensible Representation) [1]. This data structure provides a compromise between list- and array-based graph representations to support both efficient updates and efficient analysis.

As a case study, we demonstrate the effectiveness of the streaming approach to computation and the STINGER dynamic graph representation using an important social network metric called clustering coefficients. Global and local (per-vertex) clustering coefficients give the analyst an idea of the “small world-ness” of the graph [25]. The metric is derived from a count of triangles in the graph. We present an efficient multithreaded algorithm to calculate and update the clustering coefficients in an undirected, unweighted graph of 16 million vertices and 537 million edges.

The terms *streaming* [2] and *semi-streaming* [14] appear in related literature to describe a model of computation with very restrictive properties on data accesses. In streaming

graph algorithms, the graph edges are read one-by-one in an arbitrary, unknown order. Streaming algorithms typically are limited to storing  $O(n)$  or  $O(n \text{ polylog } n)$  data, where  $n$  is the number of vertices, and taking at most logarithmically many passes over the data. The metric of interest must be maintained or approximated in this fashion without significant access to data other than the edge being observed at any given point in time. Streaming models have been applied to the approximation of local clustering coefficients [6], [9]. Current high-performance computer platforms like the Cray XMT and IBM Power 595 support enough main memory to store a significant amount of graph data at once. On these platforms, the streaming model’s restrictions are overly conservative. Our massive streaming data approach leverages the continued growth in available memory.

This paper presents data structures to support our approach and a case study with clustering coefficients. Section II describes STINGER, an extensible representation that accumulates the dynamic, streaming graph and supports efficient analysis kernels. We outline assumptions and methods for extracting parallelism in Section III. As an example of our approach, Section IV considers updating clustering coefficients with streaming data. Section V details the implementation of STINGER and streaming clustering coefficients on the Cray XMT, a massively multithreaded architecture for high performance graph analysis. Using our framework, Section VI compares our three methods for updating local clustering coefficients.

## II. STINGER: A GENERAL-PURPOSE DATA STRUCTURE FOR DYNAMIC GRAPHS

Traditional graph data structures choose between efficient traversal or efficient modification. For example, a full adjacency matrix permits  $O(1)$  edge insertion or removal but requires  $O(n^2)$  storage and  $O(n)$  time to traverse all edges from any vertex, where  $n$  is the number of vertices. Adjacency lists or arrays require only  $O(n + m)$  storage, where  $m$  is the number of edges, and permit  $O(d_v)$  traversal of edges out of vertex  $v$ , where  $d_v$  is the degree of vertex  $v$ . Modifying the graph, however, can require  $O(n + m)$  time. STINGER [1] is a general-purpose graph data structure that aims to support efficient, multithreaded traversal concurrently with efficient insertion and deletion of edges.

Data structures that focus on graph traversal typically store edges or destination vertices in a packed array. The most common high-performance structure for static graph analysis borrows from sparse matrices and uses a compressed sparse row (CSR) format. In CSR form, each edge’s end vertex is stored in a single, packed array within a contiguous section corresponding to the edge’s source vertex. Inserting or deleting edges requires changing the end vertex array’s length and shifting data throughout that array. This




Figure 1. An undirected graph and a representation in STINGER with edge blocks holding two edges. The blanks are holes storing negative numbers rather than end vertex indices. Our tests use an edge block size of 100.

not only requires a large amount of data motion but also complicates concurrent access by readers.

The primary traditional approach for representing dynamic graphs uses a linked list for storing end vertices. Insertion and deletion while supporting concurrent readers is well-understood [24]. List traversal is expensive, however, often incurring  $O(n)$  cache misses. Many graph analysis kernels spend most of their time traversing the edge lists.

We developed the STINGER data structure to support efficient edge insertion and deletion with concurrent readers. STINGER takes the efficient element of CSR, storing end vertices in arrays, and loosens other requirements. STINGER also borrows from the list structure and stores edge end vertices as a list of arrays. Each vertex points to a list of fixed-size end vertex arrays; see Figure 1. Arrays need not be fixed size, but we consider this case for simplicity. This is a common mechanism for representing dynamically sized lists or arrays while supporting rapid traversal.

The arrays are permitted to have holes or blanks represented by negative entries. To delete an edge, the end vertex is found and replaced atomically by a negative number. Inserting an edge requires replacing an empty slot or possibly adding a new edge block into the linked list. We assume that a single process manages all graph updates and ensures writing does not suffer from race conditions.

Insertion and deletion can occur concurrently with reader access. By default, low-level consistency is not enforced. In a massive sparse graph, graph updates very rarely will conflict with readers. For our applications, the graph is already assumed to be an approximate model of some real-world phenomenon, and analysis kernels must account for some inconsistencies.

Other information is associated with each edge: a weight and the most recent time stamp. We do not use this information here and do not discuss the relevant consistency issues. STINGER models multi-graphs, graphs with multiple,

distinct edges between the same vertices, by associating a numeric type value with each edge. We do not use these edge types in this paper; multiple edges are treated as a single connection. Extra information is stored with the per-vertex index, including current in- and out-degrees. The degrees are updated by atomic operations but are not necessarily consistent with respect to the edge list. Analysis kernels need to handle extra or missing edges when walking the edge list.

Alternative graph data structures include forms of binary trees [18]. Trees pay an extra cost in keeping some order on the edges. On our target platform, the Cray XMT, the maintenance cost is substantial and prohibitive. STINGER’s linked array structure permits simple multithreaded traversal. Similar work in cache-oblivious algorithms often uses trees where the leaves are ordered arrays with blank entries [7]. The blank entries limit data movement when inserting a new edge into the ordered array. We are investigating whether STINGER can take advantage of a similar technique for accelerating intersections of edge lists. More radical alternatives exist, including representations using sparse certificates specific to different analysis kernels [11]. Our target is to support massive graphs, so we must support a wide variety of algorithms with a single in-memory structure. STINGER is a compromise that permits dynamic updates while supporting a wide variety of analytical algorithms on a single copy.

### III. FINDING PARALLELISM IN STREAMS AND ANALYTICS

We consider a single, unified input stream of edge insertions and deletions. This provides a synchronization point for analysis but also a bottleneck. For high performance, we need both to expose parallelism within the analytic kernels and to extract some parallelism from the sequential stream for updating the STINGER structure. We make two primary assumptions that help extract parallelism from streaming data: Changes in the stream are scattered widely enough in the massive graph that batches of them are sufficiently independent to expose parallelism. Analysis kernels have small support and small effect, and so a change to the graph only requires access to local portions and affects only a small portion of the output.

To extract parallelism from the stream, we assume the changes are somewhat scattered in the graph. In a low-diameter graph with high degree vertices, as is the case with many social networks, changes may not be completely independent, but there is potential for updating separate STINGER edge lists simultaneously. Considering the stream as batches of updates creates some loss of temporal resolution but exposes more parallelism in data structure and kernel updates. If the graph updates do not interact, then there is little temporal information lost by executing the updates together.

Analytical kernels with small support lend themselves to similar scattering across the graph. For example, per-vertex scores that depend on a fixed radius like Section IV’s local clustering coefficients naturally parallelize over batches of affected vertices. On graphs with millions or billions of vertices, the number of changes to the vertex scores will be only slightly more than the batch size for a particular set of edge insertions or deletions.

Large-support kernels like  $k$ -betweenness centrality [17] pose a more difficult challenge. They depend on paths potentially crossing the entire graph and require large-scale recalculation. A small change may update analysis results across the entire graph. Experience with  $k$ -betweenness centrality performance leads us to limit ourselves currently to kernels with small support.

We expect typical massive graph streaming analytics to fit into the following framework:

- 1: Take a section of the incoming stream as a batch.
- 2: Split the batch into per-vertex STINGER updates.
- 3: If necessary, save data (e.g. degrees) to permit incremental computation.
- 4: Process all the data structure updates.
- 5: Update analytics on the altered portion of the graph.
- 6: Transfer changed results to a monitoring process.

Sections V and VI investigate steps 2–5 for a simple analytic: local clustering coefficients.

### IV. ALGORITHM FOR UPDATING CLUSTERING COEFFICIENTS

Clustering coefficients measure the density of closed triangles in a network and are one method for determining if a graph is a small-world graph [25]. We adopt the terminology of [25] and limit our focus to *undirected* and *unweighted* graphs. A triplet is an ordered set of three vertices,  $(i, v, j)$ , where  $v$  is considered the focal point and there are undirected edges  $\langle i, v \rangle$  and  $\langle v, j \rangle$ . An open triplet is defined as three vertices in which only the required two are connected, for example the triplet  $(m, v, n)$  in Figure 2. A closed triplet is defined as three vertices in which there are three edges, or Figure 2’s triplet  $(i, v, j)$ . A triangle is made up of three closed triplets, one for each vertex of the triangle.




Figure 2. There are two triplets around  $v$  in this unweighted, undirected graph. The triplet  $(m, v, n)$  is open, there is no edge  $\langle m, n \rangle$ . The triplet  $(i, v, j)$  is closed.

The global clustering coefficient  $C$  is a single number describing the number of closed triplets over the total number of triplets,

$$C = \frac{\text{number of closed triplets}}{\text{number of triplets}} = \frac{3 \times \text{number of triangles}}{\text{number of triplets}}. \quad (1)$$

The local clustering coefficient  $C_v$  is defined similarly for each vertex  $v$ ,

$$C_v = \frac{\text{number of closed triplets centered around } v}{\text{number of triplets centered around } v}. \quad (2)$$

Let  $e_k$  be the set of neighbors of vertex  $k$ , and let  $|e|$  be the size of set  $e$ . Also let  $d_v$  be the degree of  $v$ , or  $d_v = |e_v|$ . We show how to compute  $C_v$  by expressing it as

$$C_v = \frac{\sum_{i \in e_v} |e_i \cap (e_v \setminus \{v\})|}{d_v(d_v - 1)} = \frac{T_v}{d_v(d_v - 1)}. \quad (3)$$

To update  $C_v$  as edges are inserted and deleted, we maintain the degree  $d_v$  and the triangle count  $T_v$  separately.

For the remainder of the paper, we concentrate on the calculation of local clustering coefficients. Computing the global clustering coefficient requires an additional sum reduction over the numerators and denominators.

An inserted edge increments the degree of each adjacent vertex, and a deleted edge decrements the degrees. Updating the triangle count  $T_v$  is more complicated. Algorithm 1 provides the general framework. Acting on edge  $\langle u, v \rangle$  affects the degrees only of  $u$  and  $v$  but may affect the triangle counts of all neighbors. With atomic increment operations available on most high-performance platforms, all of Algorithm 1's loops can be parallelized fully.

The search in line 5 can be implemented in several ways. A brute-force method simply iterates over every element in  $e_v$  for each  $x$ , explicitly searching for all new closed triplets given a new edge  $\langle u, v \rangle$ . The running time of the algorithm is  $O(d_u d_v)$ , which may be problematic when two high-degree vertices are affected.

If the edge list is kept sorted as in a static computation, the intersection could be computed more efficiently in  $O(d_u + d_v)$  time. However, the cost of keeping our dynamic data structure sorted outweighs the update cost. We can, however, accelerate the method to  $O((d_u + d_v) \log d_u)$  by sorting the current edge list of  $d_v$  and searching for neighbors with bisection. The sorting routine can employ a parallel sort, and iterations of the search loop can be run in parallel given atomic addition / subtraction operations.

#### *Approximating Clustering Coefficients using a Bloom Filter*

We present a novel set intersection approximation algorithm with constant-time search and query properties and an extremely high degree of accuracy. In addition to our dynamic data structure described in Section II, we summarize neighbor lists with Bloom filters [8], a probabilistic data

---

**Algorithm 1** An algorithmic framework for updating local clustering coefficients. All loops can use atomic increment and decrement instructions to decouple iterations.

---

**Input:** Edge  $\langle u, v \rangle$  to be inserted (+) or deleted (-), local clustering coefficient numerators  $T$ , and degrees  $d$

**Output:** Updated local triangle counts  $T$  and degrees  $d$

```

1:  $d_u \leftarrow d_u \pm 1$ 
2:  $d_v \leftarrow d_v \pm 1$ 
3:  $count \leftarrow 0$ 
4: for all  $x \in e_v$  do
5:   if  $x \in e_u$  then
6:      $T_x \leftarrow T_x \pm 1$ 
7:      $count \leftarrow count \pm 1$ 
8:  $T_u \leftarrow T_u \pm count$ 
9:  $T_v \leftarrow T_v \pm count$ 

```

---

structure that gives false positives (but never false negatives) with some known probability. We then query against this Bloom filter to determine if the intersection exists.

Edge arrays could be represented as bit arrays. In one extreme, each neighbor list could be an array using one bit per vertex as well as an edge list. Then  $|e_u \cap e_v|$  can be computed in  $O(\min\{d_u, d_v\})$  time by iterating over the shorter edge list and checking the bit array. However, maintaining  $O(n)$  storage per source vertex is infeasible for massive graphs.

Instead, we approximate an edge list by inserting its vertices into a Bloom filter. A Bloom filter is also a bit array but uses an arbitrary, smaller number of bits. Each edge list  $e_v$  is summarized with a Bloom filter for  $v$ . A hash function maps a vertex  $w \in e_v$  to a specific bit in this much smaller array. With fewer bits, there may be hash collisions where multiple vertices are mapped to the same bit. These will result in an overestimate of the number of intersections.

A Bloom filter attempts to reduce the occurrence of collisions by using  $k$  independent hash functions for each entry. When an entry is inserted into the filter, the output of the  $k$  hash functions determines  $k$  bits to be set in the filter. When querying the filter to determine if an edge exists, the same  $k$  hash functions are used and each bit place is checked. If any bit is set to 0, the edge cannot exist. If all bits are set to 1, the edge exists with a high probability.

Bloom filters have several parameters useful to fix a given probability of failure. In-depth description of Bloom filters' theory is beyond the scope of this paper, but a few useful features include the following: Bloom filters never yield false negatives where an edge is ignored, only false positives where a non-existent edge is counted. The probability of falsely returning membership is approximately  $(1 - e^{-kd_u/m})^k$  where  $m$  is the length of the filter. This can be optimized by setting  $k$  to an integer near  $\ln 2 \times m/d$  [13], choosing  $d$  according to the expected degrees in the graph.

Our initial implementation uses two hash functions,  $k = 2$ , and a 1 MiB filter. The probability of a false-positive will vary depending on the degree of the vertex. In a scale-free graph with an average degree of 30 and a maximum degree of 200,000, the average false-positive rate will be  $5 \times 10^{-11}$  and the worst-case rate will be  $2 \times 10^{-3}$ .

When intersecting two high-degree vertices, the Bloom filter holds a slight asymptotic advantage over sorting one edge list, but a multithreaded implementation benefits from additional parallelism throughout the Bloom filter’s operations. Note that entries cannot be deleted from a Bloom filter. A new filter is constructed for each application of Algorithm 1, so we never need to delete entries.

Modifications to Algorithm 1 for supporting a Bloom filter are straight-forward. After line 3, initialize the Bloom filter using vertices in  $e_u$ :

- 1: **for all**  $y \in e_u$  **do**
- 2:   **for**  $i = 1 \rightarrow k$  **do**
- 3:     Set bit  $H_i(y)$  in  $B_x$  to 1

Then implement the search in line 5 as follows:

- 1: **for**  $i = 1 \rightarrow k$  **do**
- 2:   **if** bit  $H_i(x) = 0$  **then**
- 3:     Skip to next  $x$

## V. MULTITHREADED PLATFORMS AND IMPLEMENTATIONS

Our implementation is based on multithreaded, shared-memory parallelism. The single code base uses different compiler directives, or pragmas, to expose the threaded parallelism. We use Cray’s compiler (version 6.3.1) and its pragmas for the massively multithreaded Cray XMT, and we use OpenMP [20] via the GNU C compiler (version 4.4.1) for a comparison on an Intel Nehalem E5530-based commodity platform.

The Cray XMT provides an ideal platform for massively multithreaded massive graph analysis. Each Threadstorm processor contains 128 *hardware streams* that maintain a thread context. Context switches between threads occur every cycle, with a new thread selected from the pool of streams ready to execute.

In this architecture, multithreading is used to hide some or all of the long latency of memory accesses. There is no cache in the processors; all latency is handled by threading. The XMT features a large, globally shared memory that is hashed to break up locality and alleviate hot-spotting. Synchronization takes place at the level of 64-bit words, and lightweight primitives such as atomic fetch-and-add are provided to the programmer. The cost of synchronization is amortized over the cost of memory access. Combined, these features enable the algorithm designer to implement highly scalable parallel algorithms for analyzing massive graphs.

The Cray XMT used for these experiments contains 64 Threadstorm processors running at 500 MHz. The globally

addressable shared memory totals 512 GiB and can hold graph data structures containing more than 2 billion vertices and 17 billion edges.

The Intel Nehalem E5530 is a 2.4GHz quad-core processor with *hyperthreading* [4]. Each physical core holds two thread contexts and switches on memory stalls. The context switches are not as frequent as on the Cray XMT, and only two contexts are available to hide memory latencies. Each core has 256 KiB of level two cache, and each processor module shares 8 MiB of level three cache. The platform tested has two E5530s, a total of eight cores and 16 threads, with 12 GiB of main memory.

Our threaded implementation is straight-forward. Each undirected edge  $\langle u, v \rangle$  is added to or removed from the data structure using two threads, one to work from each end vertex. No explicit locking is involved. The STINGER structure requires only ordered, atomic read/write of 64-bit integers (*e.g.* end vertices, timestamps) and atomic increment/decrement of counters (*e.g.* degrees, offsets). Both the Cray intrinsics and the OpenMP pragmas can express the specific operations we need. For brevity, we use the GCC/Intel intrinsic functions similar to the Cray intrinsics.

The algorithm to update the triangle counts  $T$  above uses appropriate pragmas to parallelize the outer loops. Inner loops are not parallelized under OpenMP; the target platform has insufficient threading resources to benefit from that level of parallelism. However, the inner loops are parallelized on the XMT by a loop collapse [22]. An atomic increment updates the count of a shared neighbor  $T_w$ .

Local clustering coefficients’ properties help us batch the input data. Recomputing changed coefficients only at the end of the batch’s edge actions frees us to reorder the insertions and deletions. Reordering repeated insertions and removals of the same edge may alter the edge’s auxiliary data, however, so we must take some care to resolve those in sequence order. After resolving actions on the same edge, we process all removals before all insertions to open edge slots and delay memory allocation.

The batch algorithm is as follows:

- 1: Transform undirected edges  $\langle i, j \rangle$  into pairs of directed edges  $i \rightarrow j$  and  $j \rightarrow i$  because STINGER stores directed edges.
- 2: Group edges within the batch by their source vertex.
- 3: Resolve operations on the same edge in sequence order.
- 4: Apply all removals, then all insertions to the STINGER structure.
- 5: Recompute the triangle counts and record which vertices are affected.
- 6: Recompute the local clustering coefficients of the affected vertices.

In step 5, we use slight variations of the previous algorithms. The source vertex’s neighboring vertices are gathered only once, and the array is re-used across the inner loop.

The sorted update and Bloom filter update compute their summary data using the source vertex rather than choosing the larger list.

## VI. PERFORMANCE

Our test data is generated by sampling from a Kronecker product using the RMAT recursive matrix generator [10] with probabilities  $A = 0.55$ ,  $B = 0.1$ ,  $C = 0.1$ , and  $D = 0.25$ . Each generated graph has a few vertices of high degree and many vertices of low degree. Given the RMAT scale  $k$ , the number of vertices  $n = 2^k$ , and an edge factor  $f$ , we generate  $f \cdot n$  unique edges for our initial graph. We then select a fraction  $\rho e$  of those edges and add them to a deletion queue. For these experiments,  $\rho = 0.0625$ .

After generating the initial graph, we generate 1024 actions (edge insertions or deletions) for edge-by-edge runs and 1 million actions for batched runs. With probability  $\rho$ , a deletion is taken from the deletion queue. Otherwise an insertion is generated with the same RMAT generator and parameters. The edge to be inserted may already exist in the graph. Inserted edges are appended to the deletion queue with probability  $\rho$ . There are no self-loops in our generated edges, but the implementation does handle self-loop cases by ignoring edges  $\langle v, v \rangle$ .

Clearly the Cray XMT with a large main memory can tackle far larger problems than the Intel-based platform. The latter is limited to scale  $k = 21$  and edge factor  $f = 16$  with our current test harness. On the Cray XMT, our largest experiments are sixteen times larger, with  $k = 24$  and  $f = 32$ , and are limited more by our testing code’s structure than the XMT’s architecture.

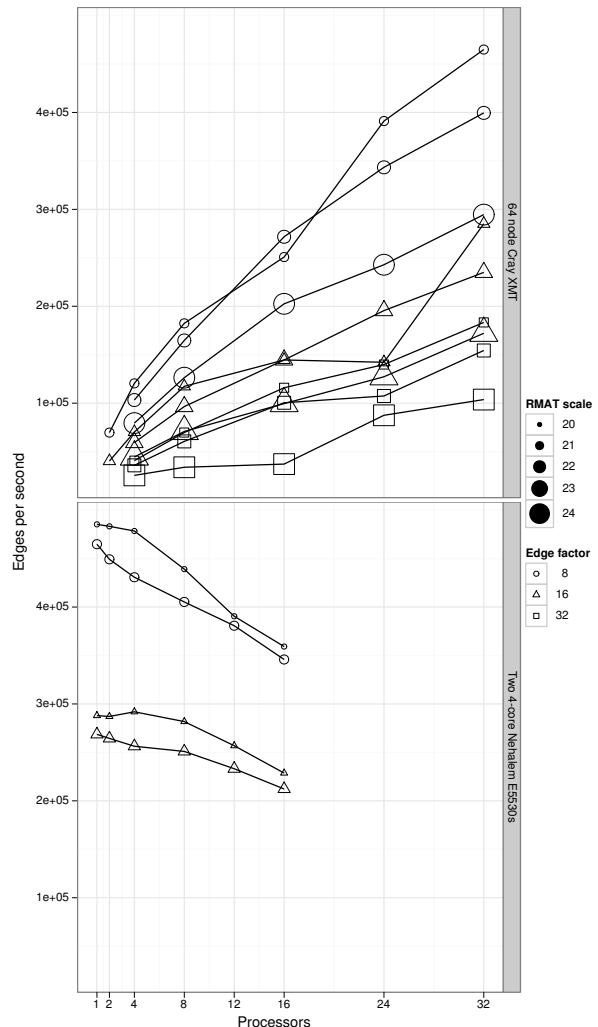
### A. Scalability of the Initial Computation

We begin by computing the correct clustering coefficients for our initial graph. While not the focus of this paper, performance on the initial computation shows interesting behavior on the two test platforms.

The initial clustering coefficients algorithm simply counts all triangles. For each edge  $\langle u, v \rangle$ , we count the size of the intersection  $|e_u \cap e_v|$ . This is a static computation, so we use a packed representation with sorted edge arrays for efficiency. The algorithm as a whole runs in  $O(\sum_v d_v^2)$  time where  $v$  ranges across the vertices and the structure is pre-sorted. The multithreaded implementation also is straightforward; we parallelize over the vertices. The  $2^{21} \approx 2$  million vertices in the smallest case provides sufficient parallelism for both platforms.

The initial computation has not been seriously tuned for performance on either platform. The algorithm itself is somewhat coarse-grained with sufficiently sized chunks of work to amortize run-time overhead. In Figure 3, the Cray XMT’s performance improves with increasing processors. Thread creation and management overhead appears responsible for the Intel Nehalem’s decreasing performance.

Figure 3. Performance of the initial clustering coefficient computations, normalized for problem size by presenting the number of edges in the graph divided by the total computation time. The Cray XMT (top) scales well as additional processors are added, while the Intel Nehalem (bottom) platform’s threading overheads decrease performance.



### B. Number of Individual Updates per Second

Unlike calculating the triangle counts  $T$  for the entire graph, updating  $T$  for an individual edge insertion or deletion exposes a variable amount of fine-grained parallelism. We present results showing how aggregate performance of a single edge insertion or deletion stays relatively constant.

Table I summarizes the sequential complexity of our update algorithms. Figure 4 presents boxplots summarizing the updates per second achieved on our test platforms. Figure 5 shows local recomputation’s speed-up of locally relative to globally recomputing the graph’s clustering coefficients. Boxes in Figures 4 and 5 span the 25% – 75% quartiles

Figure 4. Updates per second by algorithm.





Figure 5. Speed up of incremental, local updates relative to recomputing over the entire graph.



of the update times. The bar through the box shows the median. The lines stretch to the farthest non-outlier, those within  $1.5 \times$  the distance between the median and the closest box side. Points are outliers.

In Figure 4, we see the Cray XMT keeps a steady update rate on this relatively small problem regardless of the number of processors. The outliers with 16 processors are a result of sharing resources with other users. The Bloom filter shows the least variability in performance. Figure 5 shows that local recomputation accelerates the update rate typically by a factor of over a thousand.

The Intel Nehalem results degrade with additional threads. The noise at 12 and 16 threads results from over-allocation and scheduling from hyperthreading. The Intel Nehalem outperforms the Cray XMT by several orders of magnitude, but can only hold a graph of approximately 2 million

Algorithm	Update complexity
Brute force	$O(d_u d_v)$
Sorted list	$O((d_u + d_v) \log d_u)$ , $d_u < d_v$
Bloom filter	$O(d_u + d_v)$

Table I  
SUMMARY OF UPDATE ALGORITHMS

Batch Size	Brute force	Bloom filter
Edge by edge	90	60
Batch of 1000	25,100	83,700
Batch of 4000	50,100	193,300

Table II  
COMPARISON OF SINGLE EDGE VERSUS BATCHED EDGE OPERATIONS ON 32 XMT PROCESSORS, RMAT 24 INPUT, IN UPDATES PER SECOND

vertices. This Cray XMT is capable of holding a graph in memory up to 135 million vertices.

Table II shows performance obtained from batching operations and extracting parallelism on the Cray XMT. The sorting algorithm was not considered for batching. Notice that increasing the batch size greatly improves performance. For the Bloom filter, this comes at the cost of a proportional increase in memory footprint. A batch size of 4000 required choosing a filter size of 1 MiB to fit within the system's available memory. Even so, we encountered no false positives over 1 million edge actions. Increasing the batch size intuitively improves scalability since data parallelism is increased in the update phase.

## VII. CONCLUSIONS AND FUTURE WORK

We have presented a new *massive streaming data analytics* framework for graph analysis and demonstrated its ability to speed-up a fundamental analysis metric over the traditional static recomputation. To store the graph data, we have offered STINGER as a compromise between efficient graph traversal and flexible edge insertions and deletions for a wide variety of graph algorithms. We have presented clustering coefficients as a case study for STINGER and our new analytical framework.

We are able to handle individual updates rapidly enough for simple analysis (Figure 4). Updating the clustering coefficients after each edge insertion or deletion duplicates setup time, so the sorted update and Bloom filter algorithms perform relatively poorly.

The serial stream contains enough parallelism when batched to exploit the Cray XMT's massively multithreaded architecture. We achieve a speed-up of  $550 \times$  over edge-by-edge updates. The update rates of nearly 200,000 updates per second almost match gigabit Ethernet packet rates.

False-positives from Bloom filters may introduce an approximation. A modestly sized filter produces an exact result with no false-positives from our sampled scale-free networks. The Bloom filter approach achieved a  $4 \times$  speed-up over the brute-force method on the Cray XMT.

In future work, we will use this framework to explore techniques for determining when an update or a series of updates in a temporal data stream represents a significant change in the structure of the graph or a relevant subgraph.

#### ACKNOWLEDGMENTS

This work was supported in part by the PNNL CASS-MT Center and NSF Grant CNS-0614915. We thank PNNL and Cray for providing access to Cray XMT platforms. We are grateful to Kamesh Madduri, Daniel Chavarría, Jonathan Berry, Bruce Hendrickson, John Feo, Jeremy Kepner, and John Gilbert, for discussions on large-scale graph analysis and algorithm design.

#### REFERENCES

- [1] D. A. Bader, J. Berry, A. Amos-Binks, D. Chavarría-Miranda, C. Hastings, K. Madduri, and S. C. Poulos, “STINGER: Spatio-Temporal Interaction Networks and Graphs (STING) Extensible Representation,” Georgia Institute of Technology, Tech. Rep., 2009.
- [2] Z. Bar-Yossef, R. Kumar, and D. Sivakumar, “Reductions in streaming algorithms, with an application to counting triangles in graphs,” in *Proc. 13th Ann. Symp. Discrete Algorithms (SODA-02)*. San Francisco, CA: Society for Industrial and Applied Mathematics, Jan. 2002, pp. 623–632.
- [3] A.-L. Barabási, “Network databases,” 2007, <http://www.nd.edu/~networks/resources.htm>.
- [4] K. J. Barker, K. Davis, A. Hoisie, D. J. Kerbyson, M. Lang, S. Pakin, and J. C. Sancho, “A performance evaluation of the Nehalem quad-core processor for scientific computing,” *Parallel Processing Letters*, vol. 18, no. 4, pp. 453–469, 2008.
- [5] V. Batagelj and A. Mrvar, “Pajek – program for large network analysis,” *Connections*, vol. 21, no. 2, pp. 47–57, 1998.
- [6] L. Becchetti, P. Boldi, C. Castillo, and A. Gionis, “Efficient semi-streaming algorithms for local triangle counting in massive graphs,” in *KDD '08: Proceeding of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY, USA: ACM, 2008, pp. 16–24.
- [7] M. A. Bender, J. T. Fineman, S. Gilbert, and B. C. Kuszmaul, “Concurrent cache-oblivious b-trees,” in *SPAA '05: Proceedings of the seventeenth annual ACM symposium on Parallelism in algorithms and architectures*. New York, NY, USA: ACM, 2005, pp. 228–237.
- [8] B. H. Bloom, “Space/time trade-offs in hash coding with allowable errors,” *Commun. ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [9] L. S. Buriol, G. Frahling, S. Leonardi, A. Marchetti-Spaccamela, and C. Sohler, “Counting triangles in data streams,” in *PODS '06: Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. New York, NY, USA: ACM, 2006, pp. 253–262.
- [10] D. Chakrabarti, Y. Zhan, and C. Faloutsos, “R-MAT: A recursive model for graph mining,” in *Proc. 4th SIAM Intl. Conf. on Data Mining (SDM)*. Orlando, FL: SIAM, Apr. 2004.
- [11] D. Eppstein, Z. Galil, G. Italiano, and A. Nissenzweig, “Sparsification: a technique for speeding up dynamic graph algorithms,” *J. ACM*, vol. 44, no. 5, pp. 669–696, 1997.
- [12] Facebook, “User statistics,” February 2010, <http://www.facebook.com/press/info.php?statistics>.
- [13] L. Fan, P. Cao, J. Almeida, and A. Z. Broder, “Summary cache: A scalable wide-area web cache sharing protocol,” in *IEEE/ACM Transactions on Networking*, 1998, pp. 254–265.
- [14] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang, “On graph problems in a semi-streaming model,” *Theor. Comput. Sci.*, vol. 348, no. 2, pp. 207–216, 2005.
- [15] “Infovis databases,” 2005, <http://iv.slis.indiana.edu/db/index.html>.
- [16] H. Jeong, S. Mason, A.-L. Barabási, and Z. Oltvai, “Lethality and centrality in protein networks,” *Nature*, vol. 411, pp. 41–42, 2001.
- [17] K. Madduri, D. Ediger, K. Jiang, D. Bader, and D. Chavarría-Miranda, “A faster parallel algorithm and efficient multithreaded implementations for evaluating betweenness centrality on massive datasets,” in *Proc. Workshop on Multithreaded Architectures and Applications (MTAAP'09)*, Rome, Italy, May 2009.
- [18] K. Madduri and D. A. Bader, “Compact graph representations and parallel connectivity algorithms for massive dynamic network analysis,” in *23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Rome, Italy, May 2009. [Online]. Available: <http://www.cc.gatech.edu/~bader/papers/DynamicSNAP.html>
- [19] M. Newman, “Scientific collaboration networks: II. shortest paths, weighted networks and centrality,” *Phys. Rev. E*, vol. 64, p. 016132, 2001.
- [20] *OpenMP Application Program Interface; Version 3.0*, OpenMP Architecture Review Board, May 2008.
- [21] N. Reed, “Gigatweet,” 2008, <http://popacular.com/gigatweet/>.
- [22] M. Ringenburt and S.-E. Choi, “Optimizing loop-level parallelism in Cray XMT™ applications,” in *Cray User's Group*, May 2009.
- [23] Technorati, “State of the blogosphere,” 2008, <http://technorati.com/blogging/state-of-the-blogosphere/>.
- [24] J. Valois, “Lock-free linked lists using compare-and-swap,” in *Proc. 14th Ann. ACM Symp. on Principles of Distributed Computing*, Ottawa, Canada, Aug. 1995, pp. 214–222.
- [25] D. Watts and S. Strogatz, “Collective dynamics of small world networks,” *Nature*, vol. 393, pp. 440–442, 1998.