

Generalizing k -Betweenness Centrality Using Short Paths and a Parallel Multithreaded Implementation

Karl Jiang David Ediger David A. Bader
College of Computing
Georgia Institute of Technology
Atlanta, GA, USA

Abstract

We present a new parallel algorithm that extends and generalizes the traditional graph analysis metric of betweenness centrality to include additional non-shortest paths according to an input parameter k . Betweenness centrality is a useful kernel for analyzing the importance of vertices or edges in a graph and has found uses in social networks, biological networks, and power grids, among others. k -betweenness centrality captures the additional information provided by paths whose length is within k units of the shortest path length. These additional paths provide robustness that is not captured in traditional betweenness centrality computations, and they may become important shortest paths if key edges are missing in the data. We implement our parallel algorithm using lock-free methods on a massively multithreaded Cray XMT. We apply this implementation to a real-world data set of pages on the World Wide Web and show the importance of the additional data incorporated by our algorithm.

1. Introduction

The development of algorithms for the analysis of large graph data sets has driven much research in high performance computing today. Real-world networks from application domains, such as computational biology, economics, sociology, and computer networks, are sparse in nature and often exhibit “small world” properties. The algorithms often exhibit low amounts of temporal and spatial locality, while revealing little computation to hide the latency of memory accesses. Thus, their performance is often limited by the speed of main memory. Additionally, it can be difficult to obtain a balanced workload across processors in a parallel cluster computer when working with massive real-world graphs.

A useful analytic for large graphs has been the computation of betweenness centrality. As defined by Freeman in [7], betweenness centrality is a measure of the number of shortest paths in a graph passing through a given vertex. For a graph $G(V, E)$, let σ_{st} denote the number of shortest paths between vertices s and t , and $\sigma_{st}(v)$ the count of shortest

paths that pass through a specified vertex v . The betweenness centrality of v is defined as:

$$BC(v) = \sum_{s \neq v \neq t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}} \quad (1)$$

Betweenness centrality can be used to identify critical vertices in a network. High centrality scores indicate that a vertex lies on a considerable fraction of shortest paths connecting pairs of vertices. This metric has been applied extensively to the study of various networks including biological networks [9], sexual networks and the transmission of the AIDS virus [11], identifying key actors in terrorist networks [6], organizational behavior, and transportation networks [8].

In our earlier work, we developed the first parallel algorithm for betweenness centrality [1, 13]. In the remainder of this paper, we will motivate and present an extension of Freeman’s betweenness centrality and our previous algorithm. We generalize the definition to include paths in the graph whose length is within a specified value k of the length of the shortest path. We extend our recent parallel, lock-free algorithm for computing betweenness centrality to compute generalized k -betweenness centrality scores. Next, we will give details of an implementation of our new algorithm on the massively multithreaded Cray XMT and describe the performance effects of this extension in terms of execution time and memory usage. Last, we will compare the results of this algorithm on synthetic and real-world data sets.

2. Extending Betweenness Centrality

The traditional definition of betweenness centrality [7] enumerates all shortest paths in a graph and defines betweenness centrality in terms of the percentage of shortest paths passing through a vertex v . This metric has proved valuable for a number of graph analysis applications, but fails to capture the robustness of a graph. A vertex that lies on a number of paths whose length is just one greater than the shortest path receives no additional value compared to a vertex with an equally large number of shortest paths, but few paths of length one greater.

$$\delta_{s_k}(v) = \sum_{w \in Succ[v]} \left(\frac{\tau_{sv(k-\Delta D)}}{\sigma_{sw_k}} + \sum_{i=0}^k \sum_{j=0}^{k-i} [W(k-i-j, k-i) \cdot \tau_{sv(j-\Delta D)}] \frac{\delta_{s_i}(w)}{\tau_{sw_0}^{k-i+1}} \right) \quad (2)$$

$$W(n, d) = \begin{cases} \tau_{sw_0}^d, & n = 0 \\ -\sum_{i=1}^n \tau_{sw_i} \cdot W(n-i, d-1), & n > 0 \end{cases} \quad (3)$$

Figure 1. Recursive relation for the *dependency* of a vertex s on a vertex v in the graph.

We will define k -betweenness centrality in the following manner. For an arbitrary graph $G(V, E)$, let $d(s, t)$ denote the length of the shortest path between vertices s and t . We define σ_{st_k} to be the number of paths between s and t whose length is less than or equal to $d(s, t) + k$. Likewise, $\sigma_{st_k}(v)$ is the count of the subset of these paths that pass through vertex v . Therefore, k -betweenness centrality is given by:

$$BC_k(v) = \sum_{s \neq v \neq t \in V} \frac{\sigma_{st_k}(v)}{\sigma_{st_k}} \quad (4)$$

This definition of k -betweenness centrality subsumes Freeman's definition of betweenness centrality for $k = 0$. From this definition, it is clear that there are multiple ways to count paths whose length is within k of $d(s, t)$.

3. A Parallel Algorithm for k -Betweenness Centrality

Fast sequential and parallel algorithms are available for computing several variations of betweenness centrality. Brandes described a fast sequential algorithm for computing betweenness in $O(mn)$ time for an unweighted graph [4]. Madduri and Bader developed the first parallel betweenness centrality algorithm motivated by Brandes' approach that exploits both coarse- and fine-grained parallelism in low-diameter graphs in [1] and improved the performance of this algorithm using lock-free methods in [13]. In this paper, we extend the latter work to incorporate our new analytic of k -betweenness centrality.

We define $d(s, t)$ to be the length of the shortest path between s and t . Paths must be acyclic and directed outwards from the source vertex. Let τ_{st_k} be the number of paths between s and t with length equal to exactly $d(s, t) + k$, and let $\tau_{st_k}(v)$ be the number of these paths that pass through vertex v . Then, $\tau_{st_k}(v)$ is given by:

$$\tau_{st_k}(v) = \sum_{i=0}^k \tau_{sv_i} \cdot \tau_{vt_{k-i}} \quad (5)$$

Clearly, for $k = 0$ (where $\tau_{st_0} = \sigma_{st_0}$ by definition), we have reproduced the original value of $\sigma_{st}(v)$ from Brandes where $d(s, t) \geq d(s, v) + d(v, t)$ (the Bellman criterion).

The k -betweenness centrality of vertex v may be obtained by summing the pair-dependencies for that vertex:

$$BC_k(v) = \sum_{i=0}^k \sum_{s \neq v \neq t \in V} \delta_{st_k}(v) \quad (6)$$

$\delta_{st_k}(v)$ is given by a ratio of the number of paths whose length is equal to $d(s, t) + k$ passing through vertex v over the total count of the paths of length less than or equal to $d(s, t) + k$ between s and t .

$$\delta_{st_k}(v) = \frac{\tau_{st_k}(v)}{\sigma_{st_k}} \quad (7)$$

In his work, Brandes cleverly derives a recursive relation for the *dependency* of s on any other vertex v in the graph. Likewise, we have derived the general expression for any path length k greater than the shortest path. We define $\Delta D(w, v)$ as $d(s, v) - d(s, w) + 1$, where s is the source vertex and w is a neighbor of v . ΔD is bounded by k for neighbors lying on a $+k$ -path in which we are interested. We define $\delta_{s_k}(v)$, the dependence of s on v through paths of length $d(s, t) + k$, to be:

$$\delta_{s_k}(v) = \sum_{t \in V, t \neq s} \delta_{st_k}(v) \quad (8)$$

$$BC_k(v) = \sum_{i=0}^k \sum_{s \in V} \delta_{s_k}(v) \quad (9)$$

It follows that $BC_k(v)$ can be directly calculated from a sum of these dependence values. A formula for this relation is given in Fig. 1. Note that for negative i , σ_{sv_i} is defined to be zero. The reasoning for this equation is as follows: In [4], Brandes gives a definition for $\delta_{st}(v, \{v, w\})$ being the ratio of shortest paths from s to t that travel through both the vertex v and the edge $\{v, w\}$. Similarly we can define $\delta_{st_k}(v, \{v, w\})$ as the ratio of paths of length $d(s, t) + k$ from s to t that travel through v and $\{v, w\}$. If we do this, it is easy to see that we can express δ_{st_k} as a sum of $\delta_{st_k}(v, \{v, w\})$ across all neighbors w of v .

Now we see that if $t = w$ then $\delta_{st_k}(v, \{v, w\})$ is simply $\tau_{sv_{k-\Delta D(w, v)}} / \sigma_{sw_k}$. The interesting part is when $t \neq w$. In the shortest paths case, $\delta_{st}(v, \{v, w\})$ can be derived as the ratio of shortest paths to w that go through v times the fraction of shortest paths from s to t that go through w , via the Bellman criterion. However we do not have the luxury of doing this as $\delta_{st_k}(v, \{v, w\})$ depends not only

on $\tau_{st_k}(w)$ but on $\tau_{st_i}(w)$, $\forall i \leq k$. To construct this dependency we take our general formula for $\tau_{st_k}(w)$ and extend it to take into account the constraint that we need to pass through a certain edge $\{v, w\}$. To do this we can consider $\tau_{sw_k}(v, \{v, w\})$, which is exactly $\tau_{sv_{k-\Delta D(w,v)}}(v)$, and τ_{wt_k} , which is the number of shortest paths from w to t with length exactly $d(w, t) + k$, and which is unknown in our algorithm. It can be seen that for $w \neq t$:

$$\delta_{st_k}(v, \{v, w\}) = \frac{1}{\sigma_{st_k}} \sum_{i=0}^{k-\Delta D(w,v)} \tau_{sv_{k-\Delta D(w,v)-i}} \cdot \tau_{wt_i} \quad (10)$$

But here the problem is that we do not compute the value τ_{wt_k} at all. Therefore, we must derive it from other data. We want to get some expression in terms of $\tau_{st_k}(w)$, since this will reduce in summation to a recurrence on $\delta_{s_k}(w)$. From Eq. (5) we can derive a recurrence relation for τ_{wt_k} :

$$\tau_{wt_k} = \frac{\tau_{st_k}(w) - \sum_{i=1}^k \tau_{sw_i} \cdot \tau_{wt_{k-i}}}{\tau_{sw_0}} \quad (11)$$

Then we are prepared to calculate a formula for $\delta_{s_k}(v)$. Combining (11) with (10), and summing over all t 's, if we notice that $\delta_{s_k}(v) = \sum_t \frac{\tau_{st_k}(w)}{\sigma_{st_k}}$, we obtain equations (2) and (3) in Fig. 1.

For the case when $k = 0$, this formula quickly reduces to shortest path dependency of s on v from Brandes' betweenness centrality algorithm:

$$\delta_s(v) = \sum_w \frac{\sigma_{sv}}{\sigma_{sw}} \cdot (1 + \delta_s(w)) \quad (12)$$

The recursive relation involves a combinatorial number of terms and quickly becomes intractable for large values of k (see next subsection on complexity analysis). For our purposes, small values of k (less than 5) may be interesting in practice, meaning that we are concerned with how a few edge changes could heavily influence a vertex's importance. For these, it may be possible to directly derive the expression for $\delta_{s_k}(v)$.

Now that a recurrence relation for the δ values is established, we can summarize the algorithm outlined in Algorithms 1, 2, and 3. In the first stage we calculate τ_{sv_k} , $\forall v, k$ from a particular source vertex s . Then we use these τ values to recursively calculate the $\delta_{s_k}(v)$, $\forall v, k$. The first step is achieved by using breadth-first search to traverse the graph. In the second stage, we traverse the graph in backward order from which it was explored during the search stage. We repeat this for each source vertex s and sum the δ values for each vertex v to obtain the BC_k value.

We must modify the graph traversal phase from that of our previous work in [13] in order to correctly propagate values of τ_{st_k} . When $k = 0$, it suffices to do a single breadth first search and propagate the value of τ from each level to the next. For $k > 0$, we must do $k + 1$ breadth first searches.

Algorithm 1 A level-synchronous parallel algorithm for computing k -betweenness centrality of vertices in unweighted graphs.

Input: $G(V, E), k$

Output: $kBC[1..n]$, where $kBC[v]$ gives the k -centrality score ($BC_k(v)$) for vertex v

- 1: **for all** $v \in V$ **in parallel do**
 - 2: $kBC[v] \leftarrow 0$
 - 3: **for all** $s \in V$ **do**
 - 4: **Initialization**
 - 5: **for all** $t \in V$ **in parallel do** $d[t] \leftarrow -1$
 - 6: **for** $0 \leq i \leq k$ **in parallel do**
 - 7: $Succ[i][t] \leftarrow$ empty multiset, $\tau[i][t] \leftarrow 0$,
 - 8: $\tau[0][s] \leftarrow 1$, $d[s] \leftarrow 0$
 - 9: $phase \leftarrow 0$, $S[phase] \leftarrow$ empty stack
 - 10: push $s \rightarrow S[phase]$
-

Algorithm 2 Part II - Graph traversal for shortest path discovery and counting

- 1: $count \leftarrow 1$
 - 2: **while** $count > 0$ **do**
 - 3: $count \leftarrow 0$
 - 4: **for all** $v \in S[phase]$ **in parallel do**
 - 5: **for each neighbor** w **of** v **in parallel do**
 - 6: **if** $d[w] < 0$ **then**
 - 7: push $w \rightarrow S[phase + 1]$
 - 8: $count \leftarrow count + 1$
 - 9: $d[w] \leftarrow d[v] + 1$
 - 10: $\Delta D = d[v] - d[w] + 1$
 - 11: **if** $\Delta D \leq \min(k, 1)$ **then**
 - 12: $\tau[\Delta D][w] \leftarrow \tau[\Delta D][w] + \tau[0][v]$
 - 13: **if** $\Delta D \leq k$ **then**
 - 14: append $w \rightarrow Succ[\Delta D][v]$
 - 15: $phase \leftarrow phase + 1$
 - 16: **for** $1 \leq i \leq k$ **do**
 - 17: **for** $0 \leq p < phase$ **do**
 - 18: **for all** $v \in S[p]$ **in parallel do**
 - 19: **for all** $w \in Succ[0][v]$ **in parallel do**
 - 20: $\tau[i][w] \leftarrow \tau[i][w] + \tau[i][v]$
 - 21: **if** $i < k$ **then**
 - 22: **for** $0 < j \leq i + 1$ **in parallel do**
 - 23: **for all** $w \in Succ[j][v]$ **in parallel do**
 - 24: $\tau[i + 1][w] = \tau[i + 1][w] + \tau[i + 1 - j][v]$
-

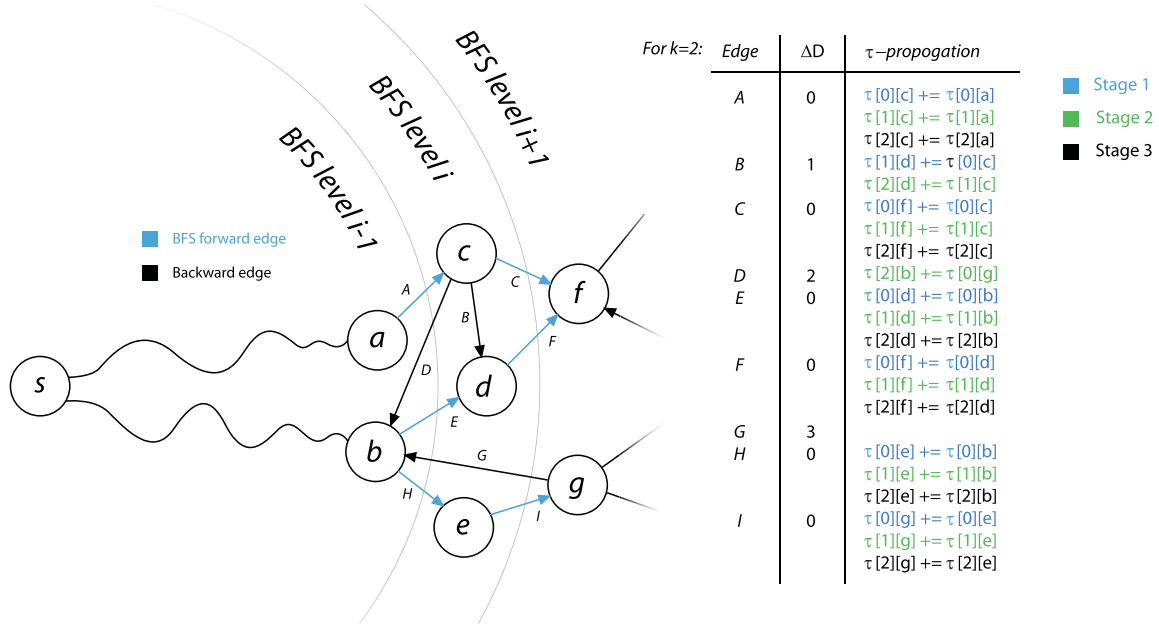


Figure 2. Illustration of τ propagation in the $k = 2$ case. Shown is a segment of the breadth-first search. The table represents the ΔD value as well as the propagation occurring as the result of each edge. The color of the τ propagation represents in which stage that addition will occur.

Algorithm 3 Part III - Dependency accumulation by back-propagation

```

1: phase  $\leftarrow$  phase - 1
2:  $\delta[i][t] \leftarrow 0 \forall t \in V, 0 \leq i \leq k$ 
3: for  $0 \leq k' \leq k$  do
4:   p  $\leftarrow$  phase
5:   while p > 0 do
6:     for all v  $\in S[p]$  in parallel do
7:       for  $0 \leq d \leq k'$  in parallel do
8:         for all w  $\in Succ[d][v]$  in parallel do
9:           for  $0 \leq i \leq k' - d$  do
10:            sum  $\leftarrow$  0
11:            e  $\leftarrow k' - d - i$ 
12:            for  $0 \leq j \leq e$  do
13:              sum  $\leftarrow$  sum + W(e - j, e, w,  $\tau$ ) *  $\tau[j][v]$ 
14:               $\delta[k'][v] \leftarrow \delta[k'][v] + sum * \frac{\delta[i][w]}{\tau[0][w]^{e+1}}$ 
15:               $\delta[k'][v] \leftarrow \delta[k'][v] + \frac{\tau[k'-d][v]}{\sum_i \tau[i][w]}$ 
16:            kBC[v]  $\leftarrow$  kBC[v] +  $\delta[k'][v]$ 
17:           p  $\leftarrow$  p - 1

```

Notice that when we are updating τ values, for neighbors on the next level in the breadth-first search, we have that $\tau_{sw_i} = \tau_{sv_i}, \forall i$, based on our generalized Bellman criterion. But we must also increment the τ values for neighbors in the current or previous levels of the search, which can be a problem since these neighbors then also need to pass on

Algorithm 4 Function $W(k, d, w, \tau)$: a recursive method producing a polynomial expansion with constant subscript sum.

Input: k, d, w , two-dimensional array τ
Output: A multivariate polynomial in $\tau[x_i][w]$ where each term has sum of exponents d and also in each term $\sum_i x_i = k$, evaluated with values from τ .

```

1: if k = 0 then return  $\tau[0][w]^d$ 
2: else
3:   sum  $\leftarrow$  0
4:   for  $0 < i \leq k$  in parallel do
5:     sum  $\leftarrow$  sum +  $\tau[i][w] * W(k - i, d - 1, w, \tau)$ 
   return sum

```

these values to their neighbors.

We solve this by only passing on different levels of τ to forward vertices and to those that are behind the breadth first search frontier. Specifically, the forward propagation always trails the backward propagation by one level. For example, in the first step, we calculate and forward-propagate τ_0 and back-propagate τ_1 . See Fig. 2 for an illustration of this process. We back-propagate τ_k to neighbors of the current vertex with appropriate ΔD and which have already been explored. We forward-propagate τ_k to neighboring vertices on the shortest path that we discover during the breadth first search. During the first breadth first search we store $k + 1$ successor (or child) arrays. When we find a neighbor during breadth first search whose $\Delta D \leq k$, we append that

neighbor’s index to the ΔD^{th} successor array.

Therefore, we need not re-run the breadth first search after the first time: as long as we synchronize on the phases, we may directly scan the successor arrays to perform our propagation, avoiding contention on a common vertex queue, and also allowing us to ignore far away neighbors and jump exactly to the neighbors of interest. In each case we exploit parallelism in the traversal by exploring the neighbors of the current level of vertices concurrently. For a small-world graph, where graph diameter is small, the number of levels in the breadth-first search is correspondingly small, and parallelism is thus high. In the first traversal, all the vertices must add newly discovered vertices to an atomically accessed queue, which is the main bottleneck in the search. Since that work is done in the first phase, we can store some data so that subsequent searches do not rely on this queue and thus do not have a sequential bottleneck to their parallelism.

For the δ -accumulation step, we start by performing shortest-path accumulation as before, in Brandes’s original algorithm. However, having these δ_0 values allows us to repeat the backward traversal to recursively calculate δ_1 , and so on and so forth. Notice that in Fig. 1, the δ_k value of the current vertex sometimes relies on the δ_k values of vertices after, before, and on the same level of the breadth first search. This may seem to preclude recursion, however as it works out, for the same level of δ , the $-\Delta D$ term on the subscript of τ_{sv} cancels out any dependencies on vertices on the same level or lower, meaning we can perform the recursive step as before, so long as we only calculate one level of δ per recursive traversal.

Complexity analysis: In Brandes’ original algorithm the betweenness centrality of any vertex can be calculated in $O(mn)$ time and $O(m+n)$ space. It is easily seen that the memory requirements of our new algorithm are exactly $k+1$ times the original, as we make $k+1$ copies of the arrays for τ, δ , and the *successor* and *successor_count* arrays, one for each level between 0 and k . The sequential *time* analysis is a bit more difficult. The breadth-first search phase is roughly $k+1$ times $O(mn)$, as we essentially traverse the graph $k+1$ times (and in the worst case we traverse each edge during each iteration). However in the accumulation phase we are actually weighed down by the calculation of the W function which grows combinatorically: more specifically, the number of terms in $W(k, d)$ grows as the number of integer partitions of k , which is $O(e^{\sqrt{k}})$. So the amount of work in this algorithm is $O(e^{\sqrt{k}}mn)$. For large k we can see that this becomes intractable to compute. However at small k it is quite manageable. For example, if we are only interested in $k < 5$ then the multiplier is bounded by $e^2 \ll mn$.

4. Computing k -Betweenness on the Cray XMT

Current hardware utilize a hierarchy of caches to give the illusion of reducing the gap to memory by exploiting the principles of spatial and temporal locality. This approach works well when the memory access pattern is predictable or when application codes demonstrate significant levels of temporal or spatial locality. However, graph analysis kernels usually exhibit fairly low levels of spatial or temporal locality [2], and execution on these platforms is limited by the speed of the memory subsystem. Hardware multithreading has been shown to be effective for producing efficient implementations of parallel graph algorithms when a significant amount of parallelism can be revealed [12].

The Cray XMT [10] uses massive numbers of hardware threads to *tolerate* latency to main memory. The XMT uses a 500 MHz 64-bit Threadstorm processor that supports 128 hardware streams of execution. Context switching between threads is lightweight and requires a single clock cycle. Each processor can support up to 16 GB of main memory that is hashed and globally addressed. The memory has a 128 KB, 4-way set associative data buffer that caches local data only. The system is built around Cray’s XT infrastructure and can scale to 8,024 processors.

The Cray XT infrastructure provides the I/O facilities and the interconnection network for the XMT. The system utilizes the Seastar-2 interconnection network that connects nodes in a 3D-torus topology. As a result, per-processor bisection bandwidth decreases as the number of processors is increased. Service nodes provide access to a Lustre file system for the storage and retrieval of large data sets.

A novel feature of the Cray XMT is its support for lightweight word-level synchronization mechanisms. Each 64-bit word of memory has a full/empty bit associated with it. C-language primitives are provided to the programmer for managing this bit and mutual exclusion locks are common. The architecture also provides an atomic `int_fetch_add` instruction that allows for integer read-modify-write operations. This is commonly used for zero-overhead shared data structures like queues and stacks. See [13] for a more detailed explanation of the synchronization and lock-free methods used in this algorithm.

Implementation and optimization: In terms of implementation on the XMT, the underlying structure is similar to that used in our previous work [13]. As before, utilization of the child arrays allows us to update our δ -values without locking. The further optimizations in our code for k -betweenness can be seen as utilizing the fact that we will mostly be using small k values. As one can see in Algorithm 3, there are several nested loops, however most of them are very simple for small k and unfurl quickly. By manually coding these loops for smaller values of k' , we significantly reduce the execution time since the time to set up and iterate

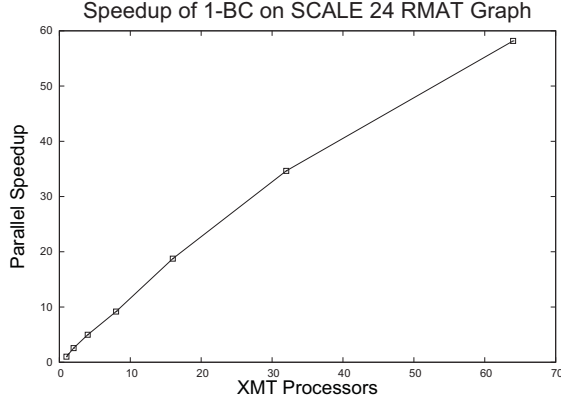


Figure 3. Parallel scaling on the Cray XMT, for an R-MAT generated graph of SCALE 24 (2^{24} vertices). Scaling is linear up to 32 processors and speedup is roughly 58 at all 64 processors (about 6400 hardware threads). $k = 1$, $K4approx = 8$ for balanced runtime and complexity (single node time 181 minutes).

over the small number of loop iterations quickly outstrips the actual useful work inside of them. For a SCALE 20 R-MAT (Recursive MATrix graph generator [5]) graph (having 2^{20} vertices and 2^{23} edges), the time to compute 1-betweenness drops by a factor of two with this optimization. We use this R-MAT generator to realize inputs that are similar to small-world graphs in degree distribution.

We begin unrolling at the loop over d values. Notice that for $d = k'$ the loops collapse into a very nice form, in that it becomes the following simplified algorithm:

Algorithm 5 Loop at line 8 from Algorithm 3, reduced form. Note that $W(0, 0) = 1$.

- 1: **for all** $w \in Succ[k'][v]$ **in parallel do**
 - 2: $\delta[k'][v] \leftarrow \delta[k'][v] + \tau[0][v] * \left(\frac{\delta[0][w]}{\tau[0][w]} + \frac{1}{\sum_i \tau[i][w]} \right)$
-

Just from this simple equation we have reduced three loops into one. Now since we have a formulation for $d = k'$ in the d loop, we may pull it out of the loop and run the rest of the loop from 0 to $k' - 1$. Similarly we may wish to pull the formulation when $d = k' - 1$, which we have done, or beyond, depending on the number of iterations we wish to unroll by hand. However, as i in $d = k' - i$ grows, the complexity of the resulting reduced formula grows at an exponential rate, meaning that for larger i this reduction is intractable.

In addition, just as we wish to take advantage of small k values in this outer loop, we may also utilize this property in writing our recursive W function that creates a τ polynomial

with an exponentially growing number of terms. For small values of k' this will already be built into the reduced loop iterations, however since larger inputs to the W function recurse downward to more manageable numbers, it will be convenient to quickly return the correct answer at lower input sizes, saving a couple recursion steps.

Apart from manually optimizing for lower loop iterations, other considerations were taken for this architecture. For example, the `malloc` system call is rather expensive in the case of these lightweight threads: if we have a `malloc` in parallel its cost compared to the number of operations in that thread's lifetime can actually be quite significant. Initially temporary arrays were kept to store the number of children accumulated for a particular vertex during graph traversal; however since temporary arrays require dynamic allocation, we modified the code to skip these temporary arrays and access the source arrays directly, at the cost of addressing the larger array repeatedly. Reorganizing memory accesses to avoid dynamic allocation within the loop reduced runtime by more than 75%. Furthermore, since the system has a plentiful amount of memory which can in essence be accessed with minimal latency, we are encouraged to utilize extra memory in lieu of performing extra calculations (as long as we have sufficient network bandwidth): thus, the expression $\sum_i \tau[i][w]$ in Algorithm 3 is precomputed and stored in an array for all values of w .

In Fig. 3 we show the parallel scaling of our optimized code on the 64-processor Cray XMT. We have reduced the execution time from nearly an hour down to a few minutes for this problem. To accommodate the more than 6,400 hardware thread contexts available on this machine, we run multiple breadth-first searches in parallel and instruct the compiler to schedule main loop iterations using *loop futures*. This is necessary since each breadth first search is dependent upon a single vertex queue which must be accessed atomically and which quickly becomes a hotspot. By doing this, however, the memory footprint is multiplied by the number of concurrent searches. On 64 processors, a graph of SCALE 24 takes about 190 seconds to run for $k = 1$ approximate betweenness. This approximation is based on selecting a random sample of source vertices s , in this case, when $K4approx = 8$, the number of starting vertices is $2^8 = 256$. The plot shows good scaling up to our machine size.

Comparison with x86: We ported our code to x86 cores using OpenMP, relying solely on running multiple breadth-first searches for parallelism. We do not provide detailed scaling results here since running on an eight core 2.0GHz Intel Clovertown Xeon we see no parallel speedup. On a single Xeon core we were able to run 1-BC on a SCALE 21 RMAT graph in 2306 seconds, as compared to 1100 seconds on a single 500MHz XMT processor (here in both cases we only run one BFS at a time).

Percentile	$k = 1$	$k = 2$
90th	513	683
95th	96	142
99th	11	12

Figure 4. The number of vertices ranked in selected percentiles for $k = 1$ and $k = 2$ whose betweenness centrality score was 0 for $k = 0$ (traditional BC). There were 14,320 vertices whose traditional BC score was 0, but whose BC_k score for $k = 1$ was greater than 0. The ND-www graph contains 325,729 vertices and 1,497,135 edges.

5. Evaluating k -Betweenness

In order to explore the effect of various values of k on the calculation of k -betweenness centrality, we apply our Cray XMT implementation to the ND-www graph data set [3]. This graph represents the hyperlink connections of web pages on the Internet. It is a directed graph with 325,729 vertices and 1,497,135 edges. Its structure demonstrates a power-law distribution in the number of neighbors. The graph displays characteristics typical of scale-free graphs found in social networks, biological networks, and computer networks.

To examine the graph data, we ran k -betweenness centrality for k from 0 (traditional betweenness centrality) to 2. The betweenness scores are compared for each value of k . An analysis directly follows. Also, after computing betweenness centrality for $k = 0$, we remove one or more of the highest ranking vertices and re-examine the results.

Looking at the highest ranking vertices going from $k = 0$ to $k = 2$, the subset of vertices and the relative rankings change little. This would seem to indicate that the paths k longer than the shortest path lie along the same vertices as the shortest path in this graph. Moreover, as predicted, the traditional betweenness centrality metric fails to capture all of the information in the graph. When examining the BC_k score for $k > 0$ of vertices whose score for $k = 0$ was 0 (no shortest paths pass through these vertices), it becomes clear that a number of very important vertices in the graph are not counted in traditional betweenness centrality. For $k = 1$, 417 vertices are ranked in the top 10 percent, but received a score of 0 for $k = 0$. In the 99th percentile are 11 vertices. Likewise, 12 vertices received a traditional BC score of 0, but ranked in the top 1 percent for $k = 2$. Fig. 5 shows 14,320 vertices whose betweenness centrality score for $k = 0$ was 0, but had a k -betweenness centrality score of greater than 0 for $k = 1$.

Given that the execution time of this algorithm grows exponentially with k , it is desirable to understand the effect of choosing a given value for k . Fig. 6 shows that increasing k from 0 to 1 captures significantly more path data. However,

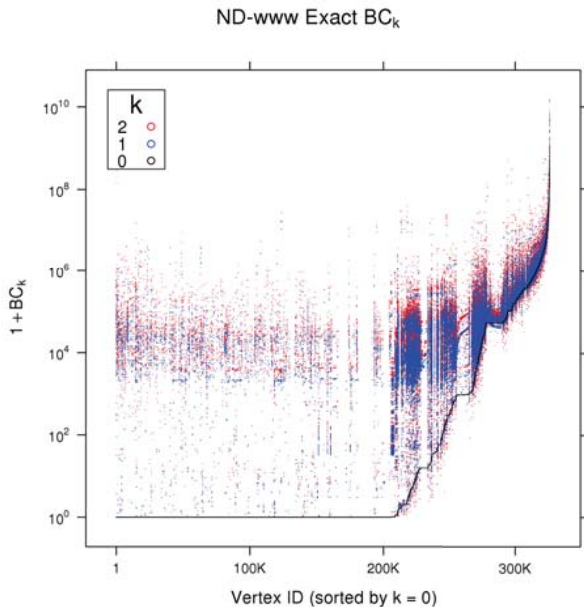


Figure 5. Per-vertex betweenness centrality scores for $k = 0, 1$, and 2 , sorted in ascending order for $k = 0$. Note the number of vertices whose score is several orders of magnitude larger for $k = 1$ or 2 than for traditional betweenness centrality.

increasing from $k = 1$ to $k = 2$ displays much less change. It is reasonable to believe that small values of k for some applications may capture an adequate amount of information while remaining computable.

The vertices that get overlooked by traditional betweenness centrality but are captured by k -betweenness centrality play an important role in the network. They do not lie along any shortest paths, but they lie along paths that are very close to the shortest path. If an edge is removed that breaks one or more shortest paths, these vertices could likely become very central to the graph. The traditional definition of betweenness centrality fails to capture this subtle importance, but k -betweenness centrality is more robust to noisy data and makes it possible to identify these vertices.

When a vertex of high betweenness is removed from the graph, it causes a number of changes in betweenness scores for all values of k that we are studying. Many vertices gain a small number of shortest paths and their ranking is fairly unchanged. In general, those vertices ranked highest on the list remain at the top. This would seem to indicate that there is a network of short paths between vertices of extremely high betweenness. Interestingly, however, other vertices jump wildly within the rankings. Often, several of these are neighbors of the removed vertex. This underscores the previous conclusion that a vertex of relatively little

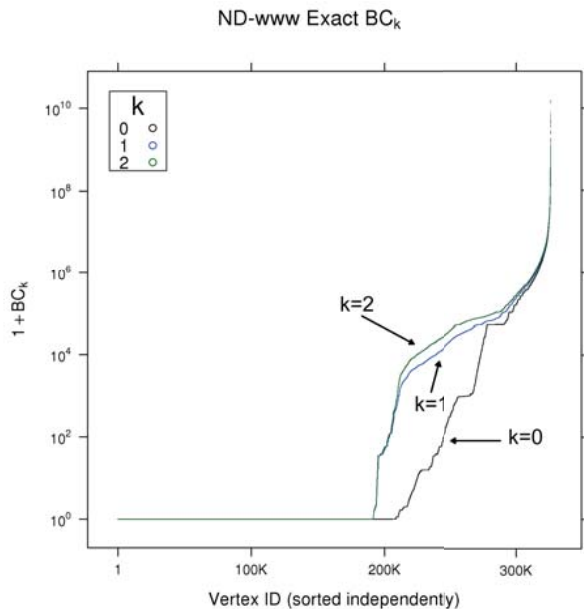


Figure 6. Per-vertex betweenness centrality scores for $k = 0, 1,$ and 2 , sorted independently in ascending order for each value of k .

importance in the graph can become extremely important if the right vertex or combination of vertices are removed. Future work will study the effect of removing edges of high betweenness, rather than vertices, on the rankings of k -betweenness centrality in these real-world networks.

6. Conclusions and Future Work

Betweenness centrality is known to be a useful metric for graph analysis. We have extended traditional betweenness centrality to the general case taking into account additional paths longer than the shortest paths. We analyze the effects of this metric on real world graphs such as the World Wide Web graph using the massively multithreaded Cray XMT. We believe that this new tool will have important consequences in the study of contingency analysis and planning, failover routing in computer networks, and extended relationships in social networks. In future work, we will apply this new tool to a wider range of graphs stemming from real world data in an effort to understand the impact that various types of paths have on the structure, robustness, and resiliency of the graph.

Acknowledgments

This work was supported in part by the PNNL CASS-MT Center and NSF Grant CNS-0614915. We would like to thank PNNL for providing access to the Cray XMT. 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 for massively multithreaded systems.

References

- [1] D.A. Bader and K. Madduri. Parallel algorithms for evaluating centrality indices in real-world networks. In *Proc. 35th Int'l Conf. on Parallel Processing (ICPP)*, Columbus, OH, August 2006. IEEE Computer Society.
- [2] D.A. Bader and K. Madduri. SNAP, Small-world Network Analysis and Partitioning: an open-source parallel graph framework for the exploration of large-scale networks. In *Proc. Int'l Parallel and Distributed Processing Symp. (IPDPS 2008)*, Miami, FL, April 2008.
- [3] A.-L. Barabási. Network databases. <http://www.nd.edu/~networks/resources.htm>, 2007.
- [4] U. Brandes. A faster algorithm for betweenness centrality. *J. Mathematical Sociology*, 25(2):163–177, 2001.
- [5] 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, April 2004. SIAM.
- [6] T. Coffman, S. Greenblatt, and S. Marcus. Graph-based technologies for intelligence analysis. *Communications of the ACM*, 47(3):45–47, 2004.
- [7] L.C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.
- [8] R. Guimerà, S. Mossa, A. Turtschi, and L.A.N. Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences USA*, 102(22):7794–7799, 2005.
- [9] H. Jeong, S.P. Mason, A.-L. Barabási, and Z.N. Oltvai. Lethality and centrality in protein networks. *Nature*, 411:41–42, 2001.
- [10] P. Konecny. Introducing the Cray XMT. In *Proc. Cray User Group meeting (CUG 2007)*, Seattle, WA, May 2007. CUG Proceedings.
- [11] F. Liljeros, C.R. Edling, L.A.N. Amaral, H.E. Stanley, and Y. Åberg. The web of human sexual contacts. *Nature*, 411:907–908, 2001.
- [12] K. Madduri, D.A. Bader, J.W. Berry, J.R. Crobak, and B.A. Hendrickson. Multithreaded algorithms for processing massive graphs. In D.A. Bader, editor, *Petascale Computing: Algorithms and Applications*, chapter 12, pages 237–262. Chapman and Hall/CRC, 2007.
- [13] K. Madduri, D. Ediger, K. Jiang, D.A. Bader, and D.G. Chavarría-Miranda. A faster parallel algorithm and efficient multithreaded implementations for evaluating betweenness centrality on massive datasets. In *Proc. Workshop on Multi-threaded Architectures and Applications (MTAAP'09)*, Rome, Italy, May 2009.