

# An Experimental Study of a Parallel Shortest Path Algorithm for Solving Large-Scale Graph Instances

Kamesh Madduri<sup>†</sup> and David A. Bader

Georgia Institute of Technology

and

Jonathan W. Berry

Sandia National Laboratories

and

Joseph R. Crobak

Rutgers University

---

We present an experimental study of the single source shortest path problem with non-negative edge weights (NSSP) on large-scale graphs using the  $\Delta$ -stepping parallel algorithm. We report performance results on the Cray MTA-2, a multithreaded parallel computer. The MTA-2 is a high-end shared memory system offering two unique features that aid the efficient parallel implementation of irregular algorithms: the ability to exploit fine-grained parallelism, and low-overhead synchronization primitives. Our implementation exhibits remarkable parallel speedup when compared with competitive sequential algorithms, for low-diameter sparse graphs. For instance,  $\Delta$ -stepping on a directed scale-free graph of 100 million vertices and 1 billion edges takes less than ten seconds on 40 processors of the MTA-2, with a relative speedup of close to 30. To our knowledge, these are the first performance results of a shortest path problem on realistic graph instances in the order of billions of vertices and edges.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Experimentation

Additional Key Words and Phrases: graph algorithms, shortest paths, parallel computing

---

## 1. INTRODUCTION

We present an experimental study of the  $\Delta$ -stepping parallel algorithm [Meyer and Sanders 2003] for solving the single source shortest path problem on large-scale graph instances. In addition to applications in combinatorial optimization problems, shortest path algorithms are finding increasing relevance in the domain of complex network analysis. Popular graph theoretic analysis metrics such as betweenness centrality [Freeman 1977; Brandes 2001] are based on shortest path algorithms. Our parallel implementation targets graph families that are representative of real-world, large-scale networks (for example, [Barabási and Albert 1999; Newman 2003]). Real-world graphs are typically characterized by a low diameter, heavy-tailed degree distributions modeled by power laws, and self-similarity. They are often very large, with the number of vertices and edges ranging from several hundreds of thousands to billions. On current workstations, it is not possible to do exact in-core computations on these graphs due to the limited physical memory. In such cases, parallel computing techniques can be applied to obtain exact solutions for memory and compute-intensive graph problems quickly. For instance, recent experimental studies on Breadth-First Search for large-scale graphs show that a parallel in-core implementation is two orders of magnitude faster than an optimized external memory implementation [Bader and Madduri 2006a; Ajwani et al. 2006]. In this paper, we present an efficient parallel implementation for the single source shortest paths problem that can handle scale-free instances in the order of billions of edges. In addition, we conduct an experimental study of performance on several other graph families, also used in the recent DIMACS Implementation Challenge [Demetrescu et al. 2006] on Shortest Paths. Please refer to our technical report [Madduri et al. 2006] for additional performance details.

Sequential algorithms for the single source shortest path problem with non-negative edge weights (NSSP) are studied extensively, both theoretically (for example, Dijkstra [1959], Fredman and Tarjan [1987], Raman [1997]) and experimentally (for example, Dial et al. [1979], Cherkassky et al. [1996], Goldberg [2001a]). Let  $m$  and  $n$  denote the number of edges and vertices in the graph respectively. Nearly all NSSP algorithms are based on the classical Dijkstra's algorithm. Using Fibonacci heaps [Fredman and Tarjan 1987], Dijkstra's

---

<sup>†</sup> Contact author: kamesh@cc.gatech.edu

algorithm can be implemented in  $O(m + n \log n)$  time. Thorup [1999] presents a linear time RAM algorithm for undirected graphs that differs significantly from Dijkstra’s approach. Instead of visiting vertices in the order of increasing distance, it traverses a *component tree*.

Parallel algorithms for solving NSSP are reviewed in detail by Meyer and Sanders [2003]. There are no known PRAM algorithms that run in sub-linear time and  $O(m + n \log n)$  work. Parallel priority queues [Driscoll et al. 1988; Brodal et al. 1998] for implementing Dijkstra’s algorithm have been developed, but these linear work algorithms have a worst-case time bound of  $\Omega(n)$ , as the only parallel operations they perform are edge relaxations. Several matrix-multiplication based algorithms [Han et al. 1997] proposed for the parallel All-Pairs Shortest Paths (APSP) involve running time and efficiency trade-offs. Parallel approximate NSSP algorithms [Klein and Subramanian 1997; Shi and Spencer 1999] based on the randomized Breadth-First search algorithm of Ullman and Yannakakis [1990] run in sub-linear time. However, it is not known how to use the Ullman-Yannakakis randomized approach for exact NSSP computations in sub-linear time.

Meyer and Sanders [2003] give the  $\Delta$ -stepping NSSP algorithm that divides Dijkstra’s algorithm into a number of *phases*, each of which can be executed in parallel. For random graphs with uniformly distributed edge weights, this algorithm runs in sub-linear time with linear average case work. Several theoretical improvements [Meyer and Sanders 2000; Meyer 2000; 2002] are given for  $\Delta$ -stepping (for instance, finding shortcut edges, adaptive bucket-splitting), but it is unlikely that they would be faster than the simple  $\Delta$ -stepping algorithm in practice, as the improvements involve sophisticated data structures that are hard to implement efficiently. On a random  $d$ -regular graph instance ( $2^{19}$  vertices and  $d = 3$ ), Meyer and Sanders report a speedup of 9.2 on 16 processors of an Intel Paragon machine, for a distributed memory implementation of the simple  $\Delta$ -stepping algorithm. For the same graph family, we are able to solve problems three orders of magnitude larger with near-linear speedup on the Cray MTA-2. For instance, we achieve a speedup of 14.82 on 16 processors and 29.75 on 40 processors for a random  $d$ -regular graph of size  $2^{29}$  vertices and  $d$  set to 3.

The literature contains few experimental studies on parallel NSSP algorithms (for example, Papaefthymiou and Rodrigue [1997], Träff [1995]). Prior implementations on distributed memory machines resorted to graph partitioning [Chandy and Misra 1982; Adamson and Tick 1991] and running a sequential NSSP algorithm on the sub-graph. The implementations perform well for certain graph families and problem sizes, but in the worst case there is no speedup.

Implementations of PRAM graph algorithms for arbitrary sparse graphs are typically memory intensive, and the memory accesses are fine-grained and highly irregular. This often leads to poor performance on cache-based systems. On distributed memory clusters, few parallel graph algorithms outperform the best sequential implementations due to long memory latencies and high synchronization costs [Bader et al. 2005]. Parallel shared memory systems are a more supportive platform. They offer higher memory bandwidth and lower latency than clusters, and the global shared memory greatly improves developer productivity. However, parallelism is dependent on the cache performance of the algorithm and scalability is limited in most cases.

We present our shortest path implementation results on the Cray MTA-2, a massively multithreaded parallel machine. The MTA-2 is a high-end shared memory system offering two unique features that aid considerably in the design of irregular algorithms: fine-grained parallelism and low-overhead word-level synchronization. The MTA-2 has no data cache; rather than using a memory hierarchy to reduce latency, the MTA-2 processors use hardware multithreading to tolerate the latency. The word-level synchronization support complements multithreading and makes performance primarily a function of parallelism. Since graph algorithms have an abundance of parallelism, yet often are not amenable to partitioning, the MTA-2 architectural features lead to superior performance and scalability. Our recent results highlight the exceptional performance of the MTA-2 for implementations of key combinatorial optimization and graph theoretic problems such as list ranking [Bader et al. 2005], connected components [Bader et al. 2005; Berry et al. 2006], subgraph isomorphism [Berry et al. 2006], Breadth-First Search and *st*-connectivity [Bader and Madduri 2006a].

The main contributions of this paper are as follows:

- *An experimental study of solving the single-source shortest paths problem in parallel using the  $\Delta$ -stepping algorithm.* Prior studies have predominantly focused on running sequential NSSP algorithms on graph families that can be easily partitioned, whereas we also consider several arbitrary, sparse graph instances. We also analyze performance using machine-independent algorithmic operation counts.
- *Demonstration of the power of massive multithreading for graph algorithms on highly unstructured instances.* We achieve impressive performance on low-diameter random and scale-free graphs.
- *Solving NSSP for large-scale realistic graph instances in the order of billions of edges.*  $\Delta$ -stepping on a synthetic directed scale-free graph of 100 million vertices and 1 billion edges takes 9.73 seconds on 40 processors of the MTA-2, with a relative speedup of approximately 31. These are the first results that we

are aware of, for solving instances of this scale and also achieving near-linear speedup. Also, the sequential performance of our implementation is comparable to competitive NSSP implementations.

## 2. REVIEW OF THE $\Delta$ -STEPPING ALGORITHM

Let  $G = (V, E)$  be a graph with  $n$  vertices and  $m$  edges. Let  $s \in V$  denote the source vertex. Each edge  $e \in E$  is assigned a non-negative real weight by the length function  $l : E \rightarrow \mathbb{R}$ . Define the *weight of a path* as the sum of the weights of its edges. The single source shortest paths problem with non-negative edge weights (NSSP) computes  $\delta(v)$ , the weight of the *shortest* (minimum-weighted) path from  $s$  to  $v$ .  $\delta(v) = \infty$  if  $v$  is unreachable from  $s$ . We set  $\delta(s) = 0$ .

Most shortest path algorithms maintain a *tentative distance* value for each vertex, which are updated by *edge relaxations*. Let  $d(v)$  denote the tentative distance of a vertex  $v$ .  $d(v)$  is initially set to  $\infty$ , and is an upper bound on  $\delta(v)$ . *Relaxing* an edge  $\langle v, w \rangle \in E$  sets  $d(w)$  to the minimum of  $d(w)$  and  $d(v) + l(v, w)$ . Based on the manner in which the tentative distance values are updated, most shortest path algorithms can be classified into two types: *label-setting* or *label-correcting*. Label-setting algorithms (for instance, Dijkstra's algorithm) perform relaxations only from *settled* ( $d(v) = \delta(v)$ ) vertices, and compute the shortest path from  $s$  to all vertices in exactly  $m$  edge relaxations. Based on the values of  $d(v)$  and  $\delta(v)$ , at each iteration of a shortest path algorithm, vertices can be classified into *unreached* ( $d(v) = \infty$ ), *queued* ( $d(v)$  is finite, but  $v$  is not settled) or *settled*. Label-correcting algorithms (e.g., Bellman-Ford) relax edges from unsettled vertices also, and may perform more than  $m$  relaxations. Also, all vertices remain in a *queued* state until the final step of the algorithm.  $\Delta$ -stepping belongs to the label-correcting type of shortest path algorithms.

---

### Algorithm 1: $\Delta$ -stepping algorithm

---

**Input:**  $G(V, E)$ , source vertex  $s$ , length function  $l : E \rightarrow \mathbb{R}$

**Output:**  $\delta(v)$ ,  $v \in V$ , the weight of the shortest path from  $s$  to  $v$

```

1 foreach  $v \in V$  do
2    $heavy(v) \leftarrow \{\langle v, w \rangle \in E : l(v, w) > \Delta\};$ 
3    $light(v) \leftarrow \{\langle v, w \rangle \in E : l(v, w) \leq \Delta\};$ 
4    $d(v) \leftarrow \infty;$ 
5  $relax(s, 0);$ 
6  $i \leftarrow 0;$ 
7 while  $B$  is not empty do
8    $S \leftarrow \phi;$ 
9   while  $B[i] \neq \phi$  do
10     $Req \leftarrow \{\langle w, d(v) + l(v, w) \rangle : v \in B[i] \wedge \langle v, w \rangle \in light(v)\};$ 
11     $S \leftarrow S \cup B[i];$ 
12     $B[i] \leftarrow \phi;$ 
13    foreach  $\langle v, x \rangle \in Req$  do
14       $relax(v, x);$ 
15     $Req \leftarrow \{\langle w, d(v) + l(v, w) \rangle : v \in S \wedge \langle v, w \rangle \in heavy(v)\};$ 
16    foreach  $\langle v, x \rangle \in Req$  do
17       $relax(v, x);$ 
18     $i \leftarrow i + 1;$ 
19 foreach  $v \in V$  do
20    $\delta(v) \leftarrow d(v);$ 

```

---

### Algorithm 2: The *relax* routine in the $\Delta$ -stepping algorithm

---

**Input:**  $v$ , weight request  $x$

**Output:** Assignment of  $v$  to appropriate bucket

```

1 if  $x < d(v)$  then
2    $B[\lfloor d(v)/\Delta \rfloor] \leftarrow B[\lfloor d(v)/\Delta \rfloor] \setminus \{v\};$ 
3    $B[\lfloor x/\Delta \rfloor] \leftarrow B[\lfloor x/\Delta \rfloor] \cup \{v\};$ 
4    $d(v) \leftarrow x;$ 

```

---

The  $\Delta$ -stepping algorithm (see Alg. 1) is an “approximate bucket implementation of Dijkstra’s algorithm” [Meyer and Sanders 2003]. It maintains an array of buckets  $B$  such that  $B[i]$  stores the set of vertices  $\{v \in V : v \text{ is queued and } d(v) \in [i\Delta, (i+1)\Delta]\}$ .  $\Delta$  is a positive real number that denotes the “bucket width”.

In each *phase* of the algorithm (the inner *while* loop in Alg. 1, lines 9–14, when bucket  $B[i]$  is not empty), all vertices are removed from the current bucket, added to the set  $S$ , and *light* edges ( $l(e) \leq \Delta$ ,  $e \in E$ ) adjacent to these vertices are relaxed (see Alg. 2). This may result in new vertices being added to the current bucket, which are deleted in the next phase. It is also possible that vertices previously deleted from the current bucket may be reinserted, if their tentative distance is improved. *Heavy* edges ( $l(e) > \Delta$ ,  $e \in E$ ) are not relaxed in a phase, as they result in tentative values outside the current bucket. Once the current bucket remains empty after relaxations, all heavy edges out of the vertices in  $S$  are relaxed at once (lines 15–17 in Alg. 1). The algorithm continues until all the buckets are empty.

Observe that edge relaxations in each phase can be done in parallel, as long as individual tentative distance values are updated atomically. The number of phases bounds the parallel running time, and the number of *reinsertions* (insertions of vertices previously deleted) and *rerelaxations* (relaxation of their out-going edges) costs an overhead over Dijkstra’s algorithm. The performance of the algorithm also depends on the value of the bucket-width  $\Delta$ . For  $\Delta = \infty$ , the algorithm is similar to the Bellman-Ford algorithm. It has a high degree of parallelism, but is inefficient compared to Dijkstra’s algorithm.  $\Delta$ -stepping tries to find a good compromise between the number of parallel phases and the number of re-insertions. Theoretical bounds on the number of phases and re-insertions, and the average case analysis of the parallel algorithm are presented in [Meyer and Sanders 2003]. We summarize the salient results.

Let  $d_c$  denote the maximum shortest path weight, and  $P_\Delta$  denote the set of paths with weight at most  $\Delta$ . Define a parameter  $l_{max}$ , an upper bound on the maximum number of edges in any path in  $P_\Delta$ . The following results hold true for any graph family.

- The number of buckets in  $B$  is  $\lceil d_c/\Delta \rceil$ .
- The total number of reinsertions is bounded by  $|P_\Delta|$ , and the total number of rerelaxations is bounded by  $|P_{2\Delta}|$ .
- The number of phases is bounded by  $\frac{d_c}{\Delta} l_{max}$ , i.e., no bucket is expanded more than  $l_{max}$  times.

For graph families with random edge weights and a maximum degree of  $d$ , Meyer and Sanders [2003] theoretically prove that  $\Delta = \theta(1/d)$  is a good compromise between work efficiency and parallelism. The sequential algorithm performs  $O(dn)$  expected work divided between  $O(\frac{d_c}{\Delta} \cdot \frac{\log n}{\log \log n})$  phases *with high probability*. In practice, in case of graph families for which  $d_c$  is  $O(\log n)$  or  $O(1)$ , the parallel implementation of  $\Delta$ -stepping yields sufficient parallelism for our parallel system.

### 3. PARALLEL IMPLEMENTATION OF $\Delta$ -STEPPING

The bucket array  $B$  is the primary data structure used by the parallel  $\Delta$ -stepping algorithm. We implement individual buckets as *dynamic arrays* that can be resized when needed and iterated over easily. To support constant time insertions and deletions, we maintain two auxiliary arrays of size  $n$ : a mapping of the vertex ID to its current bucket, and a mapping from the vertex ID to the position of the vertex in the current bucket. All new vertices are added to the end of the array, and deletions of vertices are done by setting the corresponding locations in the bucket and the mapping arrays to  $-1$ . Note that once bucket  $i$  is finally empty after a light edge relaxation phase, there will be no more insertions into the bucket in subsequent phases. Thus, the memory can be reused once we are done relaxing the light edges in the current bucket. Also observe that all the insertions are done in the relax routine, which is called once in each phase, and once for relaxing the heavy edges.

We implement a timed pre-processing step to *semi-sort* the edges based on the value of  $\Delta$ . All the light edges adjacent to a vertex are identified in parallel and stored in contiguous virtual locations, and so we visit only light edges in a phase. The  $O(n)$  work pre-processing step scales well in parallel on the MTA-2.

We also support fast parallel insertions into the request set  $R$ .  $R$  stores  $\langle v, x \rangle$  pairs, where  $v \in V$  and  $x$  is the requested tentative distance for  $v$ . We add a vertex  $v$  to  $R$  only if it satisfies the condition  $x < d(v)$ . We do not store duplicates in  $R$ . We use a sparse set representation similar to one used by Briggs and Torczon [1993] for storing vertices in  $R$ . This sparse data structure uses two arrays of size  $n$ : a *dense* array that contiguously stores the elements of the set, and a *sparse* array that indicates whether the vertex is a member of the set. Thus, it is easy to iterate over the request set, and membership queries and insertions are constant time. Unlike other Dijkstra-based algorithms, we do not relax edges in one step. Instead, we inspect adjacencies (light edges) in each phase, construct a request set of vertices, and then relax *vertices* in the relax step.

All vertices in the request set  $R$  are relaxed in parallel in the relax routine. In this step, we first delete a vertex from the old bucket, and then insert it into the new bucket. Instead of performing individual insertions, we first determine the expansion factor of each bucket, expand the buckets, and add then all vertices into their new buckets in one step. Since there are no duplicates in the request set, no synchronization is involved for updating the tentative distance values.

To saturate the MTA-2 processors with work and to obtain high system utilization, we need to minimize the number of phases and non-empty buckets, and maximize the request set sizes. Entering and exiting a parallel phase involves a negligible running time overhead in practice. However, if the number of phases is  $O(n)$ , this overhead dominates the actual running time of the implementation. Also, we enter the relax routine once every phase. The number of implicit barrier synchronizations in the algorithm is proportional to the number of phases. Our implementation reduces the number of barriers. Our source code for the  $\Delta$ -stepping implementation, along with the MTA-2 graph generator ports, is freely available online [Madduri 2006].

## 4. EXPERIMENTAL SETUP

### 4.1 Platforms

We report parallel performance results on a 40-processor Cray MTA-2 system with 160 GB uniform shared memory. Each processor has a clock speed of 220 MHz and support for 128 hardware threads. The  $\Delta$ -stepping code is written in C with MTA-2 specific pragmas and directives for parallelization. We compile it using the MTA-2 C compiler (Cray Programming Environment (PE) 2.0.3) with `-O3` and `-par` flags.

The MTA-2 code also compiles and runs on sequential processors without any modifications. Our test platform for the sequential performance results is one processor of a dual-core 3.2 GHz 64-bit Intel Xeon machine with 6GB memory, 1MB cache and running RedHat Enterprise Linux 4 (linux kernel 2.6.9). We compare the sequential performance of our implementation with the DIMACS reference solver [Demetrescu et al. 2006]. Both the codes are compiled with the Intel C compiler (icc) Version 9.0, with the flags `-O3`.

### 4.2 Problem Instances

We evaluate sequential and parallel performance on several graph families. Some of the generators and graph instances are part of the DIMACS Shortest Path Implementation Challenge benchmark package [Demetrescu et al. 2006]:

- Random graphs*: Random graphs are generated by first constructing a Hamiltonian cycle, and then adding  $m - n$  edges to the graph at random. The generator may produce parallel edges as well as self-loops. We define the random graph family *Random $_4$ - $n$*  such that  $n$  is varied,  $\frac{m}{n} = 4$ , and the edge weights are chosen from a uniform random distribution.
- Grid graphs*: This synthetic generator produces two-dimensional meshes with grid dimensions  $x$  and  $y$ . *Long- $n$*  ( $x = \frac{n}{16}$ ,  $y = 16$ ) and *Square- $n$*  grid ( $x = y = \sqrt{n}$ ) families are defined, similar to random graphs.
- Road graphs*: Road graph families with transit time (*USA-road- $t$* ) and distance (*USA-road- $d$* ) as the length function.

In addition, we also study the following families:

- Scale-free graphs*: We use the R-MAT graph model [Chakrabarti et al. 2004] for real-world networks to generate scale-free graphs. We define the family *ScaleFree $_4$ - $n$*  similar to random graphs.
- Log-uniform weight distribution*: The above graph generators assume randomly distributed edge weights. We report results for an additional *log-uniform* distribution also. The generated integer edge weights are of the form  $2^i$ , where  $i$  is chosen from the uniform random distribution  $[1, \log C]$  ( $C$  denotes the maximum integer edge weight). We define *Random $_4$ logUnif- $n$*  and *ScaleFree $_4$ logUnif- $n$*  families for this weight distribution.

### 4.3 Methodology

For sequential runs, we report the execution time of the reference DIMACS NSSP solver (an efficient implementation of Goldberg’s algorithm [Goldberg 2001b], which has expected-case linear time for some inputs) and the baseline Breadth-First Search (BFS) on every graph family. The BFS running time is a natural lower bound for NSSP codes and is a good indicator of how optimized the shortest path implementations are. It is reasonable to directly compare the execution times of the reference code and our implementation: both use a similar adjacency array representation for the graph, are written in C, and compiled and run in identical experimental settings. Note that our implementation is optimized for the MTA-2 and we make no modifications to the code before running on a sequential machine. The time taken for semi-sorting and

mechanisms to reduce memory contention on the MTA-2 both constitute overhead on a sequential processor. Also, our implementation assumes real-weighted edges, and we cannot use fast bitwise operations. By default, we set the value of  $\Delta$  to  $\frac{m}{m}$  for all graph instances. We will show that this choice of  $\Delta$  may not be optimal for all graph classes and weight distributions.

On the MTA-2, we compare our implementation running time with the execution time of a multithreaded level-synchronized breadth-first search [Bader and Madduri 2006b], optimized for low-diameter graphs. The multithreaded BFS scales as well as  $\Delta$ -stepping for all the graph instances considered, and the execution time serves as a lower bound for the shortest path running time.

The first run on the MTA-2 is usually slower than subsequent ones (by about 10% for a typical  $\Delta$ -stepping run). So we report the average running time for 10 successive runs. We run the code from three randomly chosen source vertices and average the running time. We found that using three sources consistently gave us execution time results with little variation on both the MTA-2 and the reference sequential platform. We tabulate the sequential and parallel performance metrics in [Madduri et al. 2006], and report execution time in seconds.

## 5. RESULTS AND ANALYSIS

### 5.1 Sequential Performance

First we present the performance results of our implementation on the reference sequential platform, experimenting with various graph families. Fig. 1 compares the execution time across graph instances of the same size, but from different families. The DIMACS reference code is about 1.5 to 2 times faster than our implementation for large problem instances in each family. The running time on the *Random4-n* is slightly higher than the rest of the families. For additional details such as performance as we vary the problem size for BFS,  $\Delta$ -stepping, and the DIMACS implementation execution times, please refer to Section B.1 of [Madduri et al. 2006]. Our key observation is that the ratio of the  $\Delta$ -stepping execution time to the BFS time varies between 3 and 10 across different problem instances.

### 5.2 $\Delta$ -stepping analysis

To better understand the algorithm performance across graph families, we use machine-independent algorithm operation counts. The parallel performance is dependent on the value of  $\Delta$ , the number of phases, the size of the request set in each phase. Fig. 2 plots the size of the light request set in each phase, for different graph families. By default,  $\Delta$  is set to 0.25 for all runs. If the request set size is less than 10, it is not plotted. Consider the random graph family (Fig. 2(a)). It executes in 84 phases, and the request set sizes vary from 0 to 27,000. Observe the recurring pattern of three bars stacked together in the plot. This indicates that all the light edges in a bucket are relaxed in roughly three phases, and the bucket then becomes empty. The size of the relax set is relatively high for several phases, which provides scope for exploiting multithreaded parallelism. The relax set sizes of a similar problem instance from the Long grid family (Fig. 2(b)) stands in stark contrast to that of the random graph. It takes about 200,000 phases to execute, and the maximum request size is only 15. Both of these values indicate that our implementation would fare poorly on long grid graphs (e.g. meshes with a very high aspect ratio). On square grids (Fig. 2(c)),  $\Delta$ -stepping takes fewer phases, and the request set sizes go up to 500. For a road network instance (NE USA-road-d, Fig. 2(d)), the algorithm takes 23,000 phases to execute, and only a few phases (about 30) have request sets greater than 1000 in size.

Fig. 3 plots several key  $\Delta$ -stepping operation counts for various graph classes. All synthetic graphs are roughly of the same size. Fig. 3(a) plots the average shortest path weight for various graph classes. Scale-free and Long grid graphs are on the two extremes. A log-uniform edge weight distribution also results in low average edge weight. The number of phases (see Fig. 3(b)) is highest for Long grid graphs. The number of buckets shows a similar trend as the average shortest path weight. Fig. 3(d) plots the total number of insertions for each graph family. The number of vertices is  $2^{20}$  for all graph families (slightly higher for the road network), and so  $\Delta$ -stepping results in roughly 20% overhead in insertions for all the graph families




Fig. 1. Sequential performance of our  $\Delta$ -stepping implementation on the core graph families. All the synthetic graphs are directed, with  $2^{20}$  vertices and  $\frac{m}{n} \approx 4$ . FLA(d) and FLA(t) are road networks corresponding to Florida, with 1070376 vertices and 2712768 edges

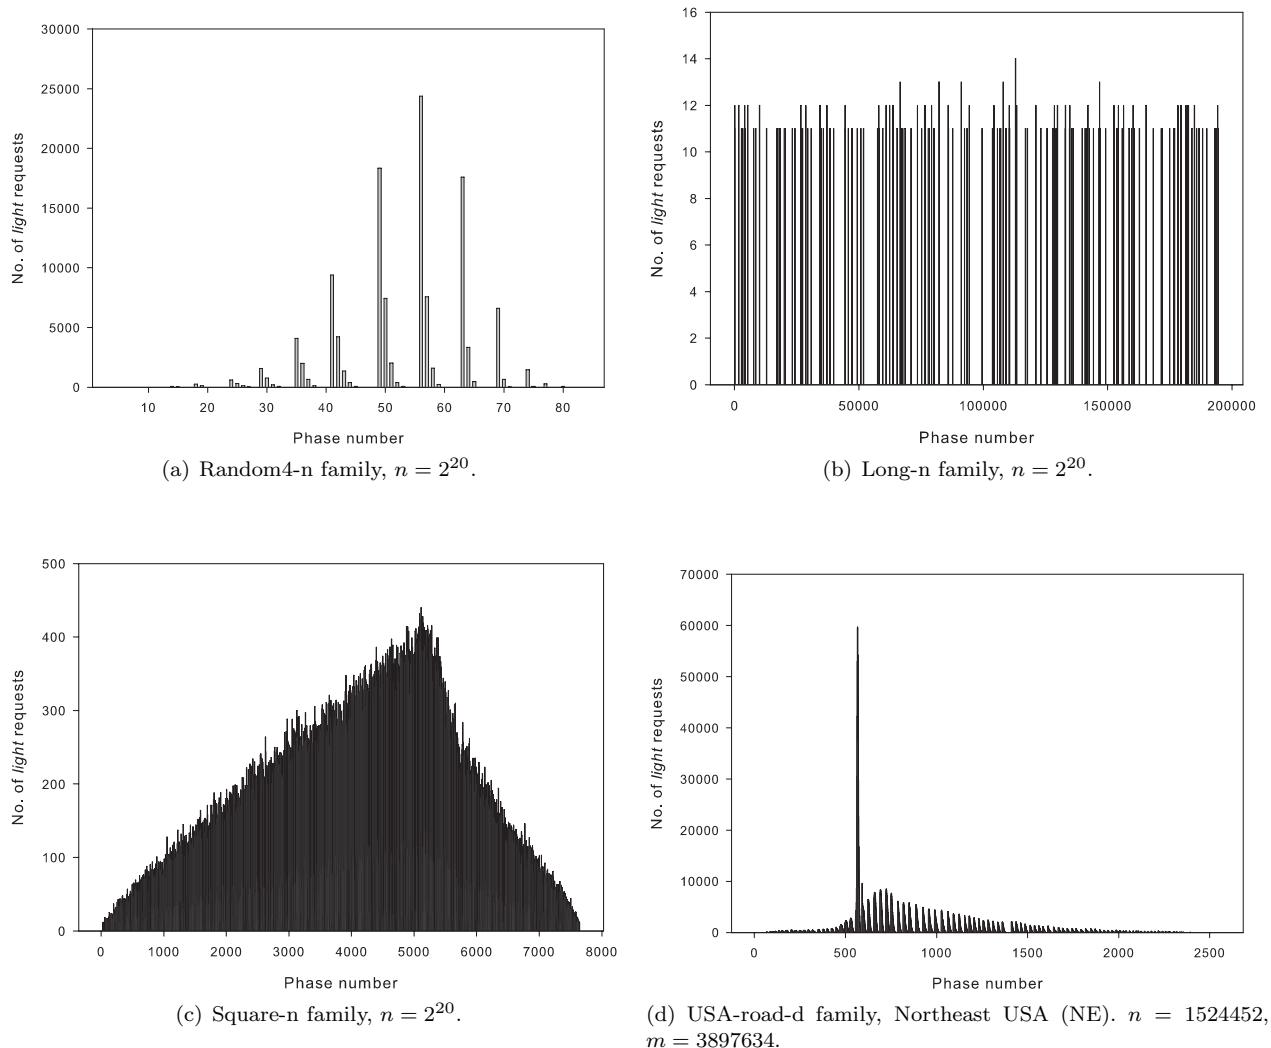


Fig. 2.  $\Delta$ -stepping algorithm: Size of the light request set at the end of each phase, for the core graph families. Request set sizes less than 10 are not plotted.

with random edge weights. Note the number of insertions for graphs with log-uniform weight distributions.  $\Delta$ -stepping performs a significant amount of excess work for these families, because the value of  $\Delta$  is quite high for this particular distribution.

### 5.3 Parallel Performance

We present the parallel scaling of the  $\Delta$ -stepping algorithm in detail. We ran  $\Delta$ -stepping and the level-synchronous parallel BFS on all graph instances described in Section 4.2 (see [Madduri et al. 2006] for complete tabular results from all experiments). We define the speedup on  $p$  processors of the MTA-2 as the ratio of the execution time on 1 processor to the execution time on  $p$  processors. In all graph classes except long grids, there is sufficient parallelism to saturate a single processor of the MTA-2 for reasonably large problem instances.

As expected,  $\Delta$ -stepping performs best for low-diameter random and scale-free graphs with randomly distributed edge weights (see Fig. 4(a) and 4(b)). We achieve a speedup of approximately 31 on 40 processors for a directed random graph of nearly a billion edges, and the ratio of the BFS and  $\Delta$ -stepping execution time is a constant factor (about 3-5) throughout. The implementation performs equally well for scale-free graphs, that are more difficult for partitioning-based parallel computing models to handle due to the irregular

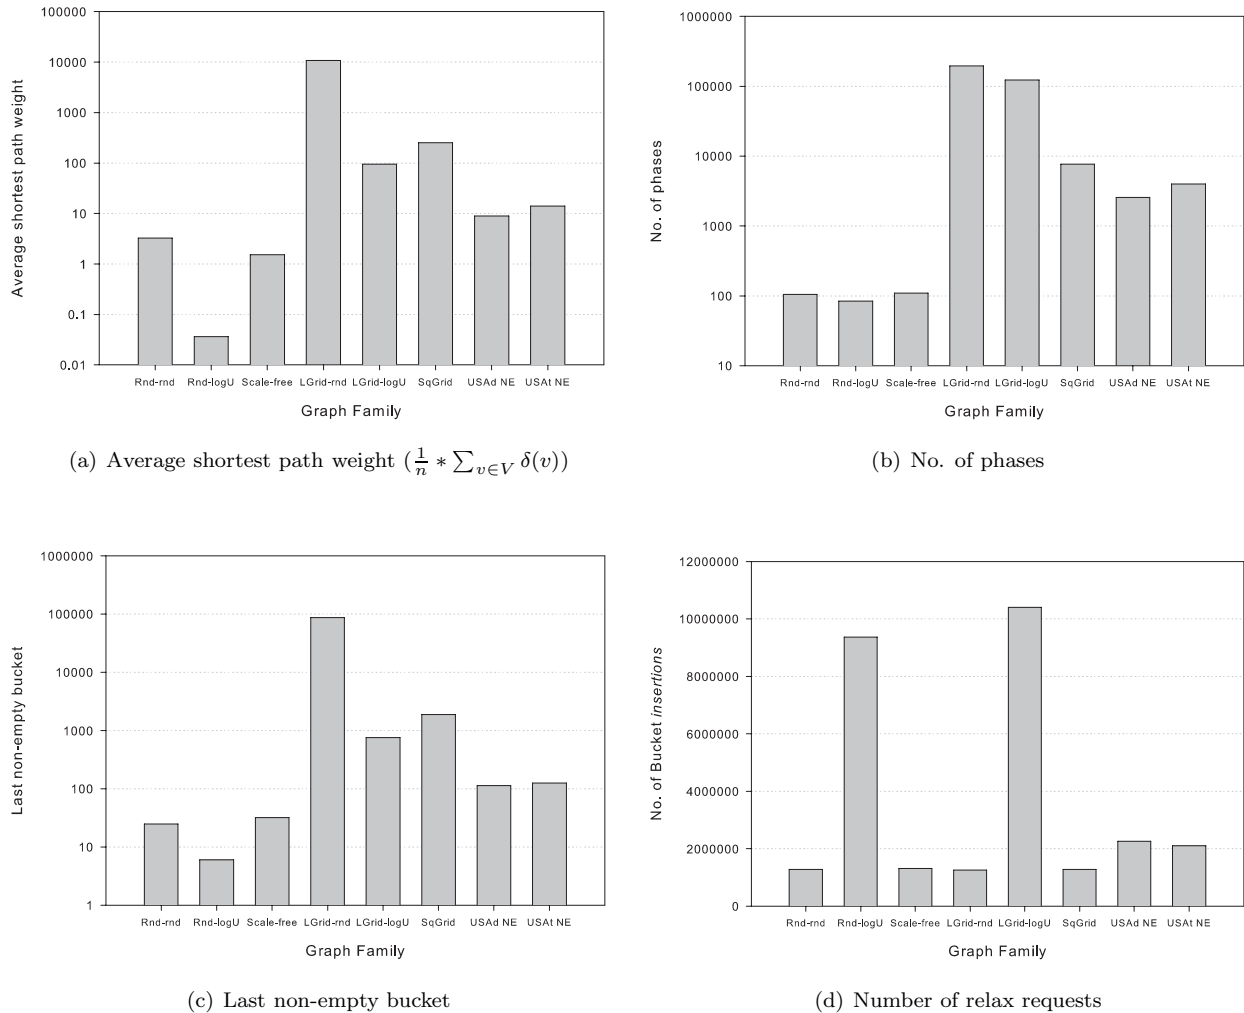


Fig. 3.  $\Delta$ -stepping algorithm performance statistics for various graph classes. All synthetic graph instances have  $n$  set to  $2^{20}$  and  $m \approx 4n$ . Rnd-rnd: Random graph with random edge weights, Rnd-logU: Random graphs with log-uniform edge weights, Scale-free: Scale-free graph with random edge weights, Lgrid: Long grid, SqGrid: Square grid, USA NE: 1524452 vertices, 3897634 edges. Plots (a), (b), (c) are on a log scale, while (d) uses a linear scale.

degree distribution. The execution time on 40 processors of the MTA-2 for the scale-free graph instance is within 9% (a difference of less than one second) of the running time for a random graph and the speedup is approximately 30 on 40 processors. We have already shown that the execution time for smaller graph instances on a sequential machine is comparable to the DIMACS reference implementation, a competitive NSSP algorithm. Thus, achieving a speedup of 30 for a realistic scale-free graph instance of one billion edges (Fig. 4(b)) is a substantial result. To our knowledge, these are the first results to demonstrate near-linear speedup for such large-scale unstructured graph instances.

In case of all the graph families, the relative speedup increases as the problem size is increased (for e.g., on 40 processors, the speedup for a *Random* $_4$ - $n$  instance with  $n = 2^{21}$  is just 3.96, whereas it is 31.04 for  $2^{28}$  vertices). This is because there is insufficient parallelism in a problem instance of size  $2^{21}$  to saturate 40 processors of the MTA-2. As the problem size increases, the ratio of  $\Delta$ -stepping execution time to multithreaded BFS running time decreases. On an average,  $\Delta$ -stepping is 5 times slower than BFS for this graph family.

For random graphs with a log-uniform weight distribution,  $\Delta$  set to  $\frac{n}{m}$  results in a significant amount of additional work. The  $\Delta$ -stepping to BFS ratio is typically 40 in this case, about 8 times higher than the






Fig. 4.  $\Delta$ -stepping execution time and relative speedup on the MTA-2 for Random4-n (left) and ScaleFree4-n (right) graph instances (directed graph,  $n=2^{28}$  vertices and  $m = 4n$  edges, random edge weights).

corresponding ratio for random graphs with random edge weights. However, the execution time scales well with the number of processors for large problem sizes.

In case of *Long-n* graphs and  $\Delta$  set to  $\frac{n}{m}$ , there is insufficient parallelism to fully utilize even a single processor of the MTA-2. The execution time of the level-synchronous BFS also does not scale with the number of processors. In fact, the running time goes up in case of multiprocessor runs, as the parallelization overhead becomes significant. Note that the execution time on a single processor of the MTA-2 is two orders of magnitude slower than the reference sequential processor. In case of square grid graphs, there is sufficient parallelism to utilize up to 4 processors for a graph instance of  $2^{24}$  vertices. For all smaller instances, the running time does not scale for multiprocessor runs. The ratio of the running time to BFS is about 5 in this case, and the  $\Delta$ -stepping MTA-2 single processor time is comparable to the sequential reference platform running time for smaller instances. For the road networks, we note that the execution time does not scale well with the number of processors, as the problem instances are quite small. We observe better performance (lower execution time, better speedup) on USA-road-d graphs than on USA-road-t graphs.

## 6. CONCLUSIONS AND FUTURE WORK

In this paper, we present an efficient implementation of the parallel  $\Delta$ -stepping NSSP algorithm along with an experimental evaluation. We study the performance for several graph families on the Cray MTA-2, and observe that our implementation execution time scales very well with number of processors for low-diameter sparse graphs. Few prior implementations achieve parallel speedup for NSSP, whereas we attain near-linear speedup for several large-scale low-diameter graph families. We also analyze the performance using platform-independent  $\Delta$ -stepping algorithm operation counts such as the number of *phases* and the *request set sizes* to explain performance across graph instances.

We intend to further study the dependence of the bucket-width  $\Delta$  on the parallel performance of the algorithm. For high diameter graphs, there is a trade-off between the number of phases and the amount of work done (proportional to the number of bucket insertions). The execution time is dependent on the value of  $\Delta$  as well as the number of processors. We need to reduce the number of phases for parallel runs and increase the system utilization by choosing an appropriate value of  $\Delta$ . Our parallel performance studies have been restricted to the Cray MTA-2 in this paper. We have designed and have a preliminary implementation of  $\Delta$ -stepping for multi-core processors and symmetric multiprocessors (SMPs), and for future work we will analyze its performance.

## ACKNOWLEDGMENTS

This work was supported in part by NSF Grants CNS-0614915, CAREER CCF-0611589, ACI-00-93039, NSF DBI-0420513, ITR ACI-00-81404, ITR EIA-01-21377, Biocomplexity DEB-01-20709, ITR EF/BIO 03-31654, and DARPA Contract NBCH30390004. We acknowledge the algorithmic inputs from Bruce Hendrickson of Sandia National Laboratories. We would also like to thank John Feo of Cray for helping us optimize the

MTA-2 implementation. Sandia is a multipurpose laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

## REFERENCES

- ADAMSON, P. AND TICK, E. 1991. Greedy partitioned algorithms for the shortest path problem. *Internat. J. Parallel Program* 20, 4, 271–278.
- AJWANI, D., DEMENTIEV, R., AND MEYER, U. 2006. A computational study of external-memory BFS algorithms. In *Proc. 17th Ann. Symp. Discrete Algorithms (SODA-06)*. ACM Press, Miami, FL, 601–610.
- BADER, D., CONG, G., AND FEO, J. 2005. On the architectural requirements for efficient execution of graph algorithms. In *Proc. 34th Int'l Conf. on Parallel Processing (ICPP)*. IEEE Computer Society, Oslo, Norway.
- BADER, D. AND MADDURI, K. 2006a. Designing multithreaded algorithms for breadth-first search and st-connectivity on the Cray MTA-2. In *Proc. 35th Int'l Conf. on Parallel Processing (ICPP)*. IEEE Computer Society, Columbus, OH.
- BADER, D. AND MADDURI, K. 2006b. Parallel algorithms for evaluating centrality indices in real-world networks. In *Proc. 35th Int'l Conf. on Parallel Processing (ICPP)*. IEEE Computer Society, Columbus, OH.
- BARABÁSI, A.-L. AND ALBERT, R. 1999. Emergence of scaling in random networks. *Science* 286, 5439, 509–512.
- BERRY, J., HENDRICKSON, B., KAHAN, S., AND KONECNY, P. 2006. Graph software development and performance on the MTA-2 and Eldorado. In *Proc. Cray User Group meeting (CUG 2006)*. CUG Proceedings, Lugano, Switzerland.
- BRANDES, U. 2001. A faster algorithm for betweenness centrality. *J. Mathematical Sociology* 25, 2, 163–177.
- BRIGGS, P. AND TORCZON, L. 1993. An efficient representation for sparse sets. *ACM Lett. Program. Lang. Syst.* 2, 1-4, 59–69.
- BRODAL, G., TRÄFF, J., AND ZAROLIAGIS, C. 1998. A parallel priority queue with constant time operations. *Journal of Parallel and Distributed Computing* 49, 1, 4–21.
- CHAKRABARTI, D., ZHAN, Y., AND FALOUTSOS, C. 2004. R-MAT: A recursive model for graph mining. In *Proc. 4th SIAM Intl. Conf. on Data Mining (SDM)*. SIAM, Orlando, FL.
- CHANDY, K. AND MISRA, J. 1982. Distributed computation on graphs: Shortest path algorithms. *Communications of the ACM* 25, 11, 833–837.
- CHERKASSKY, B., GOLDBERG, A., AND RADZIK, T. 1996. Shortest paths algorithms: theory and experimental evaluation. *Mathematical Programming* 73, 129–174.
- DEMETRESCU, C., GOLDBERG, A., AND JOHNSON, D. 2006. 9th DIMACS implementation challenge – Shortest Paths. <http://www.dis.uniroma1.it/~challenge9/>.
- DIAL, R., GLOVER, F., KARNEY, D., AND KLINGMAN, D. 1979. A computational analysis of alternative algorithms and labeling techniques for finding shortest path trees. *Networks* 9, 215–248.
- DIJKSTRA, E. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271.
- DRISCOLL, J., GABOW, H., SHRAIRMAN, R., AND TARJAN, R. 1988. Relaxed heaps: An alternative to fibonacci heaps with applications to parallel computation. *Communications of the ACM* 31, 11, 1343–1354.
- FREDMAN, M. AND TARJAN, R. 1987. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM* 34, 596–615.
- FREEMAN, L. 1977. A set of measures of centrality based on betweenness. *Sociometry* 40, 1, 35–41.
- GOLDBERG, A. 2001a. Shortest path algorithms: Engineering aspects. In *ISAAC 2001: Proc. 12th Int'l Symp. on Algorithms and Computation*. Springer-Verlag, London, UK, 502–513.
- GOLDBERG, A. 2001b. A simple shortest path algorithm with linear average time. In *9th Ann. European Symp. on Algorithms (ESA 2001)*. Lecture Notes in Computer Science, vol. 2161. Springer, Aachen, Germany, 230–241.
- HAN, Y., PAN, V., AND REIF, J. 1997. Efficient parallel algorithms for computing the all pair shortest paths in directed graphs. *Algorithmica* 17, 4, 399–415.
- KLEIN, P. AND SUBRAMANIAN, S. 1997. A randomized parallel algorithm for single-source shortest paths. *J. Algs.* 25, 2, 205–220.
- MADDURI, K. 2006. 9th DIMACS implementation challenge: Shortest Paths.  $\Delta$ -stepping C/MTA-2 code. <http://www.cc.gatech.edu/~kamesh/research/DIMACS-ch9>.
- MADDURI, K., BADER, D., BERRY, J., AND CROBAK, J. 2006. Parallel shortest path algorithms for solving large-scale instances. Technical report, Georgia Institute of Technology. Sept.
- MEYER, U. 2000. Heaps are better than buckets: parallel shortest paths on unbalanced graphs. In *Proc. 7th International Euro-Par Conference (Euro-Par 2001)*. Springer-Verlag, Manchester, United Kingdom, 343–351.
- MEYER, U. 2002. Buckets strike back: Improved parallel shortest-paths. In *Proc. 16th Int'l Parallel and Distributed Processing Symp. (IPDPS)*. IEEE Computer Society, Fort Lauderdale, FL, 1–8.
- MEYER, U. AND SANDERS, P. 2000. Parallel shortest path for arbitrary graphs. In *Proc. 6th International Euro-Par Conference (Euro-Par 2000)*. Lecture Notes in Computer Science, vol. 1900. Springer-Verlag, Munich, Germany, 461–470.
- MEYER, U. AND SANDERS, P. 2003.  $\Delta$ -stepping: a parallelizable shortest path algorithm. *J. Algs.* 49, 1, 114–152.
- NEWMAN, M. 2003. The structure and function of complex networks. *SIAM Review* 45, 2, 167–256.
- PAPAEFTHYMIU, M. AND RODRIGUE, J. 1997. Implementing parallel shortest-paths algorithms. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* 30, 59–68.
- RAMAN, R. 1997. Recent results on single-source shortest paths problem. *SIGACT News* 28, 61–72.
- SHI, H. AND SPENCER, T. 1999. Time-work tradeoffs of the single-source shortest paths problem. *J. Algorithms* 30, 1, 19–32.
- THORUP, M. 1999. Undirected single-source shortest paths with positive integer weights in linear time. *Journal of the ACM* 46, 3, 362–394.
- TRÄFF, J. L. 1995. An experimental comparison of two distributed single-source shortest path algorithms. *Parallel Computing* 21, 9, 1505–1532.
- ULLMAN, J. AND YANNAKAKIS, M. 1990. High-probability parallel transitive closure algorithms. In *Proc. 2nd Ann. Symp. Parallel Algorithms and Architectures (SPAA-90)*. ACM, Crete, Greece, 200–209.