

Designing Multithreaded Algorithms for Breadth-First Search and *st*-connectivity on the Cray MTA-2

David A. Bader Kamesh Madduri
College of Computing
Georgia Institute of Technology, Atlanta GA 30332
{bader, kamesh}@cc.gatech.edu

Abstract

Graph abstractions are extensively used to understand and solve challenging computational problems in various scientific and engineering domains. They have particularly gained prominence in recent years for applications involving large-scale networks. In this paper, we present fast parallel implementations of three fundamental graph theory problems, Breadth-First Search, st-connectivity and shortest paths for unweighted graphs, on multithreaded architectures such as the Cray MTA-2. The architectural features of the MTA-2 aid the design of simple, scalable and high-performance graph algorithms. We test our implementations on large scale-free and sparse random graph instances, and report impressive results, both for algorithm execution time and parallel performance. For instance, Breadth-First Search on a scale-free graph of 400 million vertices and 2 billion edges takes less than 5 seconds on a 40-processor MTA-2 system with an absolute speedup of close to 30. This is a significant result in parallel computing, as prior implementations of parallel graph algorithms report very limited or no speedup on irregular and sparse graphs, when compared to the best sequential implementation.

1 Introduction

Graph theoretic and combinatorial problems arise in several traditional and emerging scientific disciplines such as VLSI Design, Optimization, Databases and Computational Biology. Some examples include phylogeny reconstruction [36, 37], protein-protein interaction networks [43], placement and layout in VLSI chips [31], data mining [25, 27], and clustering in semantic webs. Graph abstractions are also finding increasing relevance in the domain of large-scale network analysis [14, 30]. Empirical studies show that many social and economic interactions tend to organize

themselves in complex network structures. These networks may contain billions of vertices with degrees ranging from small constants to thousands [8, 21]. The Internet and other communication networks, transportation and power distribution networks also share this property. The two key characteristics studied in these networks are *centrality* (which nodes in the graph are best connected to others, or have the most influence) and *connectivity* (how nodes are connected to one another). Popular metrics for analyzing these networks, like betweenness centrality [22, 10], are computed using fundamental graph algorithms like Breadth-First Search (BFS) and shortest paths.

In recognition of the importance of graph abstractions for solving large-scale problems on High Performance Computing (HPC) systems, several communities have proposed graph theoretic computational challenges. For instance, the recently announced 9th DIMACS Implementation Challenge [19] is targeted at finding shortest paths in graphs. The DARPA High Productivity Computer Systems (HPCS) [17] program has developed a synthetic graph theory benchmark called SSCA#2 [28, 5] which is composed of four kernels operating on a large-scale, directed multigraph. (We describe our implementation of SSCA#2 on symmetric multiprocessors in [6])

Graph theoretic problems are typically memory intensive, and the memory accesses are fine-grained and highly irregular. This leads to poor performance on cache-based systems. On distributed memory clusters, few parallel graph algorithms outperform their best sequential implementations due to long memory latencies and high synchronization costs. Parallel shared memory systems are a more supportive platform. They offer higher memory bandwidth and lower latency than clusters, as the global shared memory avoids the overhead of message passing. However, parallelism is dependent on the cache performance of the algorithm and scalability is limited in most cases. While it may be possible to improve the cache performance to a certain degree for some classes of graphs, there are no known general techniques for cache optimization because the memory

access pattern is largely dependent on the structure of the graph.

1.1 Preliminaries

The Cray 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 zero-overhead synchronization. The MTA-2 has no data cache; rather than using a memory hierarchy to hide latency, the MTA-2 processors use hardware multithreading to tolerate the latency. The low-overhead synchronization support complements multithreading and makes performance primarily a function of parallelism. Since graph algorithms often have an abundance of parallelism, these architectural features lead to superior performance and scalability.

The computational model for the MTA-2 is thread-centric, not processor-centric. A thread is a logical entity comprised of a sequence of instructions that are issued in order. An MTA processor has 128 hardware streams and one instruction pipeline. A stream is a physical resource (a set of 32 registers, a status word, and space in the instruction cache) that holds the state of one thread. Each stream can have up to 8 outstanding memory operations. Threads are mapped onto streams by the runtime system, and no thread is bound to a particular processor. System memory size and the inherent degree of parallelism within the program are the only limits on the number of threads used by a program.

Synchronization among threads within an executing program is easy and efficient because of special hardware support. Each 64-bit word of memory also has an associated *full/empty bit* which can be used to synchronize load and store operations. A synchronous load or store operation retries until it succeeds or traps. The thread that issued the load or store remains blocked until the operation completes, but the processor that issued the operation continues to issue instructions from non-blocked streams.

BFS [15] is one of the basic paradigms for the design of efficient graph algorithms. Given a graph $G = (V, E)$ (m edges and n vertices) and a distinguished source vertex s , BFS systematically explores the edges of G to *discover* every vertex that is reachable from s . It computes the *distance* (smallest number of edges) from s to each reachable vertex. It also produces a *breadth-first tree* with root s that contains all the reachable vertices. All vertices at a distance k (or *level* k) are first visited, before discovering any vertices at distance $k + 1$. The *BFS frontier* is defined as the set of vertices in the current level. Breadth-First Search works on both undirected and directed graphs. A queue-based sequential algorithm runs in optimal $O(m + n)$ time.

st-connectivity is a related problem, also applicable to both directed and undirected graphs. Given two vertices s and t , the problem is to decide whether or not they are con-

nected, and determine the shortest path between them, if one exists. It is a basic building block for more complex graph algorithms, has linear time complexity, and is complete for the class SL of problems solvable by symmetric, non-deterministic, log-space computations [32].

In this paper, we present fast parallel algorithms for Breadth-First Search and *st*-connectivity, for directed and undirected graphs, on the MTA-2. We extend these algorithms to compute single-source shortest paths, assuming unit-weight edges. The implementations are tested on four different classes of graphs – random graphs generated based on the Erdős-Rényi model, scale-free graphs, synthetic sparse random graphs that are hard cases for parallelization, and SSCA#2 benchmark graphs. We also outline a parallel implementation of BFS for handling high-diameter graphs.

1.2 Related Work

Distributed BFS [2, 38, 44] and *st*-connectivity [9, 23] are both well-studied problems, with related work on graph partitioning and load balancing schemes [3, 41] to facilitate efficient implementations. Other problems and algorithms of interest include shortest paths variants [20, 13, 40, 39, 34, 16] and external memory algorithms and data structures [1, 11, 33] for BFS. Several PRAM and BSP [18] algorithms have been proposed to solve this problem. However, there are very few parallel implementations that achieve significant parallel speedup on sparse, irregular graphs when compared against the best sequential implementations. In [4], we demonstrated superior performance for list ranking and connected components on the MTA-2 when compared with symmetric multiprocessor implementations, and attained considerable absolute speedups over the best sequential implementations. This work serves as the primary motivation for our current experimentation on the MTA-2.

2 A Multithreaded Approach to Breadth-First Search

Unlike prior parallel approaches to BFS, on the MTA-2 we do not consider load balancing or the use of distributed queues for parallelizing BFS. We employ a simple level-synchronized parallel algorithm (Alg. 1) that exploits concurrency at two key steps in BFS:

1. All vertices at a given *level* in the graph can be processed simultaneously, instead of just picking the vertex at the head of the queue (step 7 in Alg. 1)
2. The adjacencies of each vertex can be inspected in parallel (step 9 in Alg. 1).

We maintain an array d to indicate the level (or distance) of each visited vertex, and process the global queue Q ac-

Input: $G(V, E)$, source vertex s

Output: Array $d[1..n]$ with $d[v]$ holding the length of the shortest path from s to $v \in V$, assuming unit-weight edges

```

1 for all  $v \in V$  in parallel do
2    $d[v] \leftarrow -1$ ;
3  $d[s] \leftarrow 0$ ;
4  $Q \leftarrow \phi$ ;
5 Enqueue  $s \leftarrow Q$ ;
6 while  $Q \neq \phi$  do
7   for all  $u \in Q$  in parallel do
8     Delete  $u \leftarrow Q$ ;
9     for each  $v$  adjacent to  $u$  in parallel do
10      if  $d[v] = -1$  then
11         $d[v] \leftarrow d[u] + 1$ ;
12      Enqueue  $v \leftarrow Q$ ;

```

Algorithm 1: Level-synchronized Parallel BFS

cordingly. Alg. 1 is however a very high-level representation, and hides the fact that thread-safe parallel insertions to the queue and atomic updates of the distance array d are needed to ensure correctness. Alg. 2 details the MTA-2 code required to achieve this (for the critical steps 7 to 12), which is simple and very concise. The loops will not be automatically parallelized as there are dependencies involved. The compiler can be forced to parallelize them using the *assert parallel* directive on both the loops. In this case, the compiler automatically collapses the inner for loop, and dynamically schedules the loop iterations. Thus, the implementation is independent of the vertex degree distribution. We do not need to bother about load balancing in case of graph families with skewed degree distributions, such as real-world scale-free graphs.

We use the low-overhead synchronization calls `int_fetch_add`, `readfe()`, and `writteef()` to atomically update the value of d , and insert elements to the queue in parallel. `int_fetch_add` offers synchronized updates to data representing shared counters, without using locks. The `readfe` operation atomically reads data from a memory location only after that location's full/empty bit is set full, and sets it back to empty. If the bit is not full to start with, the thread executing the read operation suspends in hardware and is later retried. Similarly, a `writteef` writes to a memory location when the full/empty bit is empty and then sets it to full. A `readfe` should be matched with a `writteef`, or else the program might deadlock.

We observe that the above parallelization scheme will not work well for high-diameter graphs (for instance, consider a chain of vertices with bounded degree). In case of high-diameter graph families, the number of vertices at each BFS level is typically a small number. We do not have suffi-

```

/* While the Queue is not empty */
#pragma mta assert parallel
#pragma mta block dynamic schedule
for ( $i = \text{startIndex}$ ;  $i < \text{endIndex}$ ;  $i++$ )
   $u = Q[i]$ ;
  /* Inspect all vertices adjacent to  $u$  */
  #pragma mta assert parallel
  for ( $j = 0$ ;  $j < \text{degree}[u]$ ;  $j++$ )
     $v = \text{neighbor}[u][j]$ ;
    /* Check if  $v$  has been visited yet? */
     $\text{dist} = \text{readfe}(\&d[v])$ ;
    if ( $\text{dist} == -1$ )
      writteef(&d[v],  $d[u] + 1$ );
    else
      writteef(&d[v],  $\text{dist}$ );
  /* Enqueue  $v$  */
   $Q[\text{int\_fetch\_add}(\&\text{count}, 1)] = v$ ;

```

Algorithm 2: MTA-2 parallel C code for steps 7-12 in Alg. 1

cient parallelism in the level-synchronized approach to saturate the MTA-2 system. For arbitrary sparse graphs, Ullman and Yannakakis offer high-probability PRAM algorithms for transitive closure and BFS [42] that take $\tilde{O}(n^\epsilon)$ time with $\tilde{O}(mn^{1-2\epsilon})$ processors, provided $m \geq n^{2-3\epsilon}$. The key idea here is as follows. Instead of starting the search from the source vertex s , we expand the frontier up to a distance d in parallel from a set of randomly chosen *distinguished* vertices (that includes the source vertex s also) in the graph. We then construct a new graph whose vertices are the distinguished vertices, and we have edges between these vertices if they were pair-wise reachable in the previous step. Now a set of *superdistinguished* vertices are selected among them and the graph is explored to a depth t^2 . After this step, the resulting graph would be dense and we can determine the shortest path of the source vertex s to each of the vertices. Using this information, we can determine the shortest paths from s to all vertices.

3 *st*-connectivity and Shortest Paths

We can easily extend the Breadth-First Search algorithm for solving the *st*-connectivity problem too. A naïve implementation would be to start a Breadth-First Search from s , and stop when t is visited. However, we note that we could run BFS concurrently both from s and to t , and if we keep track of the vertices visited and the expanded frontiers on both sides, we can correctly determine the shortest path between s and t . The key steps are outlined in Alg. 3 (termed STCONN-FB), which has both high-level details as well as MTA-specific synchronization constructs. Both s and t are

Input: $G(V, E)$, vertex pair (s, t)

Output: The smallest number of edges $dist$ between s and t , if they are connected

```

1 for all  $v \in V$  in parallel do
2    $color[v] \leftarrow WHITE$ ;
3    $d[v] \leftarrow 0$ ;
4  $color[s] \leftarrow RED$ ;  $color[t] \leftarrow GREEN$ ;  $Q \leftarrow \phi$ ;  $done \leftarrow FALSE$ ;  $dist \leftarrow \infty$ ;
5  $Enqueue\ s \leftarrow Q$ ;  $Enqueue\ t \leftarrow Q$ ;
6 while  $Q \neq \phi$  and  $done = FALSE$  do
7   for all  $u \in Q$  in parallel do
8      $Delete\ u \leftarrow Q$ ;
9     for each  $v$  adjacent to  $u$  in parallel do
10       $color \leftarrow readfe(\&color[v])$ ;
11      if  $color = WHITE$  then
12         $d[v] \leftarrow d[u] + 1$ ;
13         $Enqueue\ v \leftarrow Q$ ;
14         $writeef(\&color[v], color[u])$ ;
15      else
16        if  $color \neq color[u]$  then
17           $done \leftarrow TRUE$ ;
18           $tmp \leftarrow readfe(\&dist)$ ;
19          if  $tmp > d[u] + d[v] + 1$  then
20             $writeef(\&dist, d[u] + d[v] + 1)$ ;
21          else
22             $writeef(\&dist, tmp)$ ;
23         $writeef(\&color[v], color)$ ;

```

Algorithm 3: st -connectivity (STCONN-FB): concurrent BFSes from s and t

added to the queue initially, and newly discovered vertices are either colored RED (for vertices reachable from s) or GREEN (for vertices that can reach t). When a *back edge* is found in the graph, the algorithm terminates and the shortest path is evaluated. As in the previous case, we encounter nested parallelism here and apply the same optimizations. The pseudo-code is elegant and concise, but must be carefully written to avoid the introduction of race conditions and potential deadlocks (see [7] for an illustration).

We also implement an improved algorithm for st -connectivity (STCONN-MF, denoting *minimum frontier*, detailed in Alg. 4) that is suited for graphs with irregular degree distributions. In this case, we maintain two different queues Q_s and Q_t and expand the smaller frontier (Q in Alg. 4 is either Q_s or Q_t , depending on the values of $extentS$ and $extentT$) on each iteration. Thus, STCONN-MF visits fewer vertices and edges compared to STCONN-FB. Alg. 3 and 4 are discussed in more detail in an extended version of this paper [7].

Input: $G(V, E)$, vertex pair (s, t)

Output: The smallest number of edges $dist$ between s and t , if they are connected

```

1 for all  $v \in V$  in parallel do
2    $color[v] \leftarrow WHITE$ ;
3    $d[v] \leftarrow 0$ ;
4  $color[s] \leftarrow GRAY$ ;  $color[t] \leftarrow GRAY$ ;  $Q_s \leftarrow \phi$ ;  $Q_t \leftarrow \phi$ ;
5  $done \leftarrow FALSE$ ;  $dist \leftarrow -1$ ;
6  $Enqueue\ s \leftarrow Q_s$ ;  $Enqueue\ t \leftarrow Q_t$ ;  $extentS \leftarrow 1$ ;  $extentT \leftarrow 1$ ;
7 while ( $Q_s \neq \phi$  or  $Q_t \neq \phi$ ) and  $done = FALSE$  do
8   Set  $Q$  appropriately;
9   for all  $u \in Q$  in parallel do
10     $Delete\ u \leftarrow Q$ ;
11    for each  $v$  adjacent to  $u$  in parallel do
12       $color \leftarrow readfe(\&color[v])$ ;
13      if  $color = WHITE$  then
14         $d[v] \leftarrow d[u] + 1$ ;
15         $Enqueue\ v \leftarrow Q$ ;
16         $writeef(\&color[v], color[u])$ ;
17      else
18        if  $color \neq color[v]$  then
19           $dist \leftarrow d[u] + d[v] + 1$ ;
20           $done \leftarrow TRUE$ ;
21           $writeef(\&color[v], color)$ ;
22     $extentS \leftarrow |Q_s|$ ;  $extentT \leftarrow |Q_t|$ ;

```

Algorithm 4: st -connectivity (STCONN-MF): alternate BFSes from s and t

4 Experimental Results

This section summarizes the experimental results of our BFS and st -connectivity implementations on the Cray MTA-2. We report results on a 40-processor MTA-2, with each processor having a clock speed of 220 MHz and 4GB of RAM. From the programmer's viewpoint, the MTA-2 is however a global shared memory machine with 160GB memory.

We test our algorithms on four different classes of graphs (see Fig. 1):

- **Random graphs (RAND1):** We construct random graphs of n vertices and m directed edges. The m edges are added by randomly choosing a pair of vertices each time.
- **Scale-free graphs (SF-RMAT),** used to model real-world large-scale networks: These graphs are generated using the R-MAT graph model [12]. They have a significant number of vertices of very high degree, although the majority of vertices are low-degree ones. The degree distribution plot on a log-log scale is a

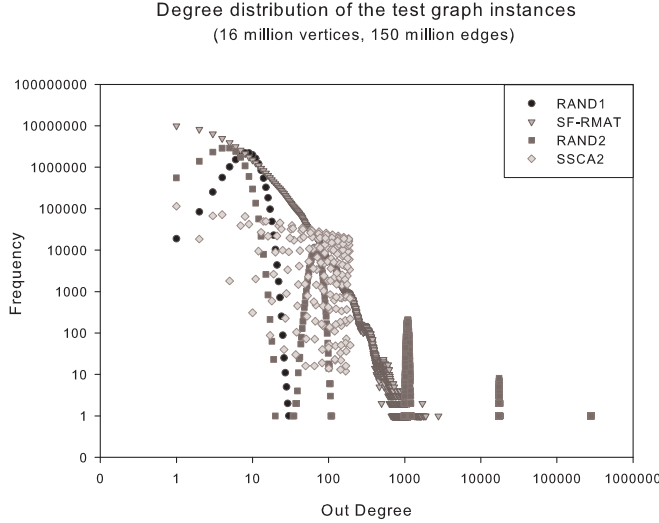


Figure 1. Degree distributions of the four test graph classes

straight line with a heavy tail, as seen in Fig. 1.

- Synthetic sparse random graphs that are hard cases for parallelization and load balancing (RAND2): As in scale-free graphs, a considerable percentage of vertices are high-degree ones, but the degree distribution is different.
- DARPA SSCA#2 benchmark (SSCA2) graphs: A typical SSCA#2 graph consists of a large number of highly interconnected clusters of vertices. The clusters are sparsely connected, and these inter-cluster edges are randomly generated. The cluster sizes are uniformly distributed and the maximum cluster size is a user-defined parameter. For the graph used in the performance studies in Fig. 1, we set the maximum cluster size to 10.

We generate directed graphs in all four cases. Our algorithms work for both directed and undirected graphs, as we store all the neighbors of each vertex in the internal representation of the graph. In this section, we report parallel performance results for directed graphs only. With minor changes to the code, we can apply it to undirected graphs also.


Fig. 2(a) plots the execution time and speedup attained by the Breadth-First Search algorithm on a random graph of 200 million vertices and 1 billion edges. The plot in the inset shows the scaling when the number of processors is varied from 1 to 10, and the main plot for 10 to 40 processors. We define the *Speedup* on p processors of the MTA-2 as the ratio of the execution time on p processors to that on one processor. Since the computation on the MTA is thread-centric, system utilization is also an important met-

ric to study. We observed utilization close to 97% for single processor runs. The system utilization was consistently high (around 80% for 40 processor runs) across all runs. We achieve a speedup of nearly 10 on 10 processors for random graphs, 17 on 20 processors, and 28 on 40 processors. This is a significant result, as random graphs have no locality and such instances would offer very limited on no speedup on cache-based SMPs and other shared memory systems. In case of the 40 processor run, there is a drop in efficiency as the parallelization overhead is comparable to the BFS execution time. However, even for the 40 processor runs, the hot spot due to the BFS queue index does not lead to any performance degradation.


Fig. 2(b) gives the BFS execution time for a Scale-free graph of 200 million vertices and 1 billion edges, as the number of processors is varied from 1 to 40. Our algorithm is independent of the degree distribution, and so the performance for scale-free graphs is identical to random graphs. Fig. 2(c) plots the multi-processor BFS speedup for SF-RMAT and RAND2 graph instances of the same size. The execution time and speedup are nearly identical. Fig. 2(d) summarizes the BFS performance for SSCA#2 graphs. The performance on SSCA#2 graphs are comparable to random graphs of similar sizes.

Fig. 2(e) and 2(f) show the performance of BFS as the edge density is varied for SSCA#2 and RAND2 graphs respectively. We vary the user-defined *MaxClusterSize* parameter in case of SSCA#2 graphs, and show that the implementation performs well for all m values. We observe similar performance for a RAND2 graph of 2.147 billion edges, when the number of vertices is varied from 16 million to 536 million.


Fig. 3 summarizes the performance of the *st*-connectivity algorithms. Note that both the *st*-connectivity algorithms are based on BFS, and if BFS is implemented efficiently, we would expect *st*-connectivity also to perform well. Fig. 3(a) shows the performance of STCONN-MF on scale-free graphs as the number of processors is varied from 1 to 10. Note that the execution times are highly dependent on (s, t) pair we choose. In this particular case, just 45,000 vertices were visited in a graph of 134 million vertices. The *st*-connectivity algorithm shows near-linear scaling with the number of processors. The actual execution time is bounded by the BFS time, and is dependent on the shortest path length and the degree distribution of the vertices in the graph. In Fig. 3(b), we compare the performance of the two algorithms, concurrent Breadth-First Searches from s and t (STCONN-FB), and expanding the smaller frontier in each iteration (STCONN-MF). Both of them scale linearly with the number of processors for a problem size of 134 million vertices and 805 million edges. STCONN-MF performs slightly better for this graph instance. They were also found to perform comparably for the other graph families.




(a) Execution time and speedup for Random (RAND1) graphs: 1-10 processors (inset), and 10-40 processors




(b) Execution time and speedup for Scale-free (SF-RMAT) graphs: 1-10 processors (inset), and 10-40 processors




(c) Parallel performance comparison: RAND2 and SF-RMAT graphs



(d) Execution time and speedup for SSCA2 graphs



(e) Execution time variation as a function of average degree for RAND2 graphs



(f) Execution time variation as a function of average degree for SSCA2 graphs

Figure 2. Breadth First Search Performance Results




Figure 3. *st*-connectivity Performance Results

5 Conclusions

We present fast multithreaded algorithms for fundamental graph theory problems. Our implementations show strong scaling for irregular and sparse graphs chosen from four different graph classes, and also achieve high system utilization. The absolute execution time values are significant; Problems involving large graphs of billions of vertices and edges can be solved in seconds to minutes. With its latency tolerant processors, high bandwidth network, global shared memory and fine-grained synchronization, the MTA-2 is the first parallel machine to perform extraordinarily well on sparse graph problems. It may now be possible to tackle several key PRAM algorithms [26, 29, 24, 35] that have eluded practical implementations so far. Another attractive feature of the MTA-2 is the ease of programming. It is possible to write concise and elegant code, focusing on exploiting the concurrency in the problem, rather than optimizing for cache locality (or minimizing communication in distributed memory systems).

Acknowledgements

This work was supported in part by NSF Grants 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 would like to thank Richard Russell for sponsoring our MTA-2 accounts. We are grateful to John Feo for providing the SSCA#2 graph generator source. We acknowledge the significant algorithmic inputs and MTA-2 programming help from Jonathan Berry and

Bruce Hendrickson. Finally, we would like to thank Simon Kahan, Petr Konecny, John Feo and other members of the Cray Eldorado team for their valuable advice and several suggestions on optimizing code for the MTA-2.

References

- [1] J. M. Abello and J. S. Vitter, editors. *External memory algorithms*. American Mathematical Society, Boston, MA, USA, 1999.
- [2] B. Awerbuch and R. G. Gallager. A new distributed algorithm to find breadth first search trees. *IEEE Trans. Inf. Theor.*, 33(3):315–322, 1987.
- [3] B. Awerbuch, A. V. Goldberg, M. Luby, and S. A. Plotkin. Network decomposition and locality in distributed computation. In *IEEE Symp. on Foundations of Computer Science*, pages 364–369, 1989.
- [4] D. Bader, G. Cong, and J. Feo. On the architectural requirements for efficient execution of graph algorithms. In *Proc. 34th Int'l Conf. on Parallel Processing (ICPP)*, Oslo, Norway, June 2005.
- [5] D. Bader, J. Feo, and *et al.* HPCS SSCA#2 Graph Analysis Benchmark Specifications v1.1, Jan. 2006.
- [6] D. Bader and K. Madduri. Design and implementation of the HPCS graph analysis benchmark on symmetric multiprocessors. In *Proc. 12th Int'l Conf. on High Performance Computing*, Goa, India, Dec. 2005. Springer-Verlag.
- [7] D. Bader and K. Madduri. Designing multithreaded algorithms for Breadth-First Search and *st*-connectivity on the Cray MTA-2. Technical report, Georgia Tech, Feb. 2006.
- [8] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, October 1999.
- [9] G. Barnes and W. L. Ruzzo. Deterministic algorithms for undirected *s-t* connectivity using polynomial time and sub-linear space. In *Proc. 23rd Annual ACM Symp. on Theory of*

Computing, pages 43–53, New York, NY, USA, 1991. ACM Press.

- [10] U. Brandes. A faster algorithm for betweenness centrality. *J. Mathematical Sociology*, 25(2):163–177, 2001.
- [11] A. L. Buchsbaum, M. Goldwasser, S. Venkatasubramanian, and J. Westbrook. On external memory graph traversal. In *Proc. 11th Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 859–860, Philadelphia, PA, USA, 2000.
- [12] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-MAT: A recursive model for graph mining. In *Proc. 4th SIAM Intl. Conf. on Data Mining*, Florida, USA, Apr. 2004.
- [13] B. Cherkassky, A. Goldberg, and T. Radzik. Shortest paths algorithms: theory and experimental evaluation. *Mathematical Programming*, 73:129–174, 1996.
- [14] T. Coffman, S. Greenblatt, and S. Marcus. Graph-based technologies for intelligence analysis. *Communications of the ACM*, 47(3):45–47, 2004.
- [15] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Inc., Cambridge, MA, 1990.
- [16] A. Crauser, K. Mehlhorn, U. Meyer, and P. Sanders. A parallelization of Dijkstra’s shortest path algorithm. In *Proc. 23rd Int’l Symp. on Mathematical Foundations of Computer Science*, pages 722–731, London, UK, 1998. Springer-Verlag.
- [17] DARPA Information Processing Technology Office. High productivity computing systems project, 2004. <http://www.darpa.mil/ipto/programs/hpcs/>.
- [18] F. Dehne, A. Ferreira, E. Cáceres, S. W. Song, and A. Roncato. Efficient parallel graph algorithms for coarse-grained multicompilers and BSP. *Algorithmica*, 33:183–200, 2002.
- [19] C. Demetrescu, A. Goldberg, and D. Johnson. 9th DIMACS implementation challenge – Shortest Paths. <http://www.dis.uniroma1.it/~challenge9/>.
- [20] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 1:269–271, 1959.
- [21] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proc. ACM SIGCOMM*, pages 251–262, 1999.
- [22] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40(1):35–41, 1977.
- [23] H. Gazit and G. L. Miller. An improved parallel algorithm that computes the BFS numbering of a directed graph. *Inf. Process. Lett.*, 28(2):61–65, 1988.
- [24] J. Greiner. A comparison of data-parallel algorithms for connected components. In *Proc. 6th Ann. Symp. Parallel Algorithms and Architectures (SPAA-94)*, pages 16–25, Cape May, NJ, June 1994.
- [25] A. Inokuchi, T. Washio, and H. Motoda. An apriori-based algorithm for mining frequent substructures from graph data. In *Proc. 4th European Conf. on Principles of Data Mining and Knowledge Discovery*, pages 13–23, London, UK, 2000.
- [26] J. JáJá. *An Introduction to Parallel Algorithms*. Addison-Wesley Publishing Company, New York, 1992.
- [27] G. Karypis, E. Han, and V. Kumar. Chameleon: Hierarchical clustering using dynamic modeling. *Computer*, 32(8):68–75, 1999.
- [28] J. Kepner, D. P. Koester, and *et al.* HPCS Scalable Synthetic Compact Application (SSCA) Benchmarks, 2004. <http://www.highproductivity.org/SSCABmks.htm>.
- [29] P. Klein and J. Reif. An efficient parallel algorithm for planarity. *J. Comp. and System. Sci.*, 37(2):190–246, 1988.
- [30] V. Krebs. Mapping networks of terrorist cells. *Connections*, 24(3):43–52, 2002.
- [31] T. Lengauer. *Combinatorial algorithms for integrated circuit layout*. John Wiley & Sons, Inc., New York, NY, USA, 1990.
- [32] H. R. Lewis and C. H. Papadimitriou. Symmetric space-bounded computation (extended abstract). In *Proc. 7th Colloquium on Automata, Languages and Programming*, pages 374–384, London, UK, 1980.
- [33] U. Meyer. External memory BFS on undirected graphs with bounded degree. In *Proc. 12th Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 87–88, Philadelphia, PA, USA, 2001.
- [34] U. Meyer and P. Sanders. Δ -stepping: a parallelizable shortest path algorithm. *J. Algorithms*, 49(1):114–152, 2003.
- [35] G. L. Miller and J. H. Reif. Parallel tree contraction and its application. In *Proc. 26th Ann. IEEE Symp. Foundations of Computer Science*, pages 478–489, Portland, OR, Oct. 1985.
- [36] B. Moret, D. Bader, and T. Warnow. High-performance algorithm engineering for computational phylogenetics. In *Proc. Int’l Conf. on Computational Science*, volume 2073–2074 of *Lecture Notes in Computer Science*, San Francisco, CA, 2001. Springer-Verlag.
- [37] B. M. Moret, D. Bader, T. Warnow, S. Wyman, and M. Yan. GRAPPA: a high-performance computational tool for phylogeny reconstruction from gene-order data. In *Proc. Botany*, Albuquerque, NM, Aug. 2001.
- [38] P. M. Pardalos, M. G. Resende, and K. G. Ramakrishnan, editors. *Parallel Processing of Discrete Optimization Problems: DIMACS Workshop April 28-29, 1994*. American Mathematical Society, Boston, MA, USA, 1995.
- [39] R. Seidel. On the all-pairs-shortest-path problem in unweighted undirected graphs. *J. Comput. Syst. Sci.*, 51(3):400–403, 1995.
- [40] M. Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM*, 46(3):362–394, 1999.
- [41] J. L. Träff. An experimental comparison of two distributed single-source shortest path algorithms. *Parallel Comput.*, 21(9):1505–1532, 1995.
- [42] J. Ullman and M. Yannakakis. High-probability parallel transitive closure algorithms. In *Proc. 2nd Ann. Symp. Parallel Algorithms and Architectures (SPAA-90)*, pages 200–209, New York, NY, USA, 1990. ACM Press.
- [43] A. Vazquez, A. Flammini, A. Maritan, and A. Vespignani. Global protein function prediction in protein-protein interaction networks. *Nature Biotechnology*, 21(6):697–700, June 2003.
- [44] A. Yoo, E. Chow, K. Henderson, W. McLendon, B. Hendrickson, and Ü. V. Çatalyürek. A scalable distributed parallel breadth-first search algorithm on BlueGene/L. In *Proc. Supercomputing (SC 2005)*, Seattle, WA, Nov. 2005.