International Journal of Advances in Engineering and Management (IJAEM)
Volume 2, Issue 11, pp: 177-184 www.ijaem.net ISSN: 2395-5252

Authentic Data Dissemination Algorithm

Abba Hassan[1], Bulama Mohammed[2], Yakubu Bala, Mohammed[3]
[1,2,3] *Computer Science, School of Science & Technology, Abubakar Tatari Ali Polytechnic, Bauchi- Nigeria.*

Date of Submission: 05-12-2020 Date of Acceptance: 20-12-2020

ABSTRACT: Gossip and broadcast protocols are two information transmission problems defined for a group of individuals connected by a communication network. When gossiping, all users of the network know one thing of information and need to communicate it to everyone else. When broadcasting, a source node wants to send a message to all other nodes. Both are some of the most fundamental problems that arise in communication networks. This work examined problems that produce gossip and broadcasting in general. For example, the source node can have multiple messages. Many of the related works examined in the literature focused on homogenous networks. The developed algorithms are more applicable to data management on local networks. However, large storage systems generally consist of storage devices assembled over a wide network. Finding an appropriate model and developing algorithms for transmission that recognize the multidimensional nature of the communication network is an important part of this study. Problems with data collection over an extensive, largely overlooked network were also addressed. It is likely to become more important as the Internet becomes more embedded in everyday life. We consider a situation where large volumes of data must be moved from several different locations to a destination. In this work, we focus on two main properties: the available bandwidth may vary and the network may not select the best route for transferring data between two hosts. We focus on optimizing task completion time by redirecting data through intermediate hosts and showing that, under certain network conditions, we can reduce the total completion time by a factor of two. This is done by developing an approach to coordinate schedules of data collection using network streams. communicates serially with a microprocessor. The microprocessor monitors and reports the engine's performance and control the opening/closing of the engine valves. The ultimate goal is improved efficiency, decrease pollutants, and produce maximum power throughout the RPM range with a calmness engine.

KEYWORDS:AGossip; Broadcast; Algorithm; Data; Transmission.

I. INTRODUCTION

Broadcasting and gossip are some of the most basic problems that arise in communication networks. In this thesis, we mainly deal with issues related to broadcasting and gossip that arise when managing large amounts of data. We study various issues that relate to the dissemination and collection of data both in local networks and in large areas.

The issues of broadcasting and gossip have been studied for decades [66, 44, 48, 10, 11, 51]. The diffusion problem is defined as follows: There are n nodes, and a source node must pass an element to any other node. In the gossip problem, each node has an element that they want to communicate to all the others. We can treat the problem of gossip as simultaneously making programs. Communication is usually done in rounds, where one node in each round can send (or receive) one item at most to another node. A typical objective function is to reduce the number of rounds of communication. In this thesis we use this objective function as a performance metric. Another typical objective function considered in the literature is to reduce the number of calls made.

In order to correct the need for high data, in addition to high faults, we may want to keep multiple copies of the same file on some discrepancy. The disk usually has a storage limit and the number of users who can access the data from it at the same time. A data structure sets up how many copies of each file there should be, and what subset of disks it will display. Given the demand for data objects, it is difficult to compute an NP data model that maximizes the number of appeals [86, 39]. Golubchik et al. [39] developed a polynomial approximation equation for this problem. Let us first consider a simple and appealing example: suppose we have a collection of files stored on a disk at the beginning (this will be a repository) how we can do this model first?

Each item must be sent for a wrong order of pencil. In addition, we want to create output files as soon as possible, given the large volume of data involved. We call this problem only a multicast problem. (With multicast, we mean only one of the categories of the requested object.) This is a generalization to the single-source signal optimization problem by Cockayne, Thomason, and Farley [23, 29]. In some cases, product information is initially stored in multiple locations, and we also want to create a startup. We call this multicast problem. The problems of data migration are generalized to the three problems above, and a factor for estimation has been established. [59].

Another reason to investigate the issue of news media is because it is a key function of multimedia communication, such as MPI (Message Delivery) [77, 40, 52]. Broadcasting takes place when you need to quickly send data for processing across the network. This system makes it easy to see any multicast and multi-site multicast applications. We will present the problem as well as other general information 1.2.1. In the section. Remember that in some situations' latency is important when sending small files. In other situations, bandwidth is the most important thing when making big changes. By solving the number of combinations, we solve two types of problems.

Much of the work in the media and negatives in literature focuses on homogeneous participation. Algorithms are better developed for data management in regional centers. However, for large-scale systems, storage capabilities are typically spread across the network, where data transfer across the network is much slower than data transfer. from a regional perspective. Therefore, we plan to develop algorithms to create data files faster than traditional networks.

II. RELATED WORK

Broadcasting efficiently is an essential operation and many works are devoted to this under a number of communication models (see [81, 44, 55, 9, 12] and references therein). For example, Elkin-Kortsarz [28] consider minimizing the broadcast time in arbitrarily connected graphs, with the property that only adjacent nodes in the graph may communicate. However, the approximation guarantee is). The postal model [9] captures the communication latency when passing a message, and optimal broadcast algorithm was developed. The LogP model [24] suggests an alternative framework when dealing with nodes in a single cluster, and it captures both the communication latency and throughput in the network. Broadcasting algorithms [55] for the LogP model have been developed and shown to be optimal. However, all algorithms under the above models only work under homogeneous environment.

Various models for heterogeneous environments have been proposed in the literature. One general model is the one proposed by Bar-Noy et al. [8] where the communication costs between links are not uniform. In addition, the sender may engage in another communication before the current one is complete. An approximation algorithm with a guarantee of O(logk) is given for the operation of performing a multicast of size k. Another simple model for heterogeneous networks of workstations was proposed by Banikazemi et al. [7]. In this model, heterogeneity among processors is modeled by a non-uniform speed of the sending processor. A heterogeneous cluster is defined as a collection of processors p1,p2,...,pn in which each processor is capable of communicating with any other processor. Each processor has a transmission time which is the time required to send a message to any other processor in the cluster. Thus, the time required for the communication is a function of only the sender. Each processor may send messages to other processors in order, and each processor may be receiving only one message at a time. They proposed a simple heuristic called the Fastest Node First (FNF) heuristic, which was studied further by Liu [73] and by Khuller and Kim [57, 62]. However, in this model it is assumed that the time taken by a processor to send a message to any other processor is the same. This is the main limitation of the model.

Lowekamp and Beguelin [74] considered the same two-tier communication network model. Several works have been done to provide ways to deliver data by not following the default network route. Although some works exist on multipoint-to-point aggregation mechanisms at the IP layer [5, 17], such solutions have focused on reduction of overheads due to small packets (e.g., ACKs) and usually require the use of an active networks framework which is not currently widely deployed over the public Internet. Another approach is application-level re-routing, which is used to improve end-to-end performance, or provide efficient fault detection and recovery for wide-area applications. For instance, in [84] the authors perform a measurement-based study of comparing end-to-end round-trip time, loss rate, and bandwidth of default routing vs alternate path routing. Their results show that in 30% to 80% of the cases, there is an alternate path with significantly superior quality. Their work provides evidence for existence of alternate paths which can outperform default.

III. METHODOLOGY

Suppose we have N disks and Δ data items. The single-source multicast problem is defined asfollows:

Single-source multicast. There are Δ data items stored on a single disk (the source). We need to send data item ito a specified subset D_i of disks. Figure 1.1 shows the initial and target layouts, and their corresponding D_i's for a single-source multicast instance when Δ is 4. Our goal is to find a schedule using the minimum number of rounds, that is, minimizing the makespan, subject to the following communication model.Our algorithm consists of two phases. In the first phase, we make exactly $b|D_i|/2c$ copies for all items i. Once each item ihas $b|D_i|/2c$ copies, in second phase we can finish migrating one item at each round by copying from the current copies to the remaining $b|D_i|/2c$ disks in D_i which have not received item ias yet and using the source disk to make another copy if $|D_i|$ is odd.

Phase I. At the t-th round, we do the following.

1. The source disk s creates a new copy for item t if $t \leq \Delta$.
2. For items j ($j < t$), double the number of copies until the number of copies becomes $b|D_j|/2c$.

In other words, if the current number of copies of item j is less than or equal to 2^{dj-2}every disk having item j makes another copy of it so that the number of copies is doubled. Otherwise if the current number of copies of item j is 2^{dj-1}, then only $b|D_j|/2c-2^{dj-1}$ disks need to make copies (thus the number of copies of item j becomes exactly $b|D_j|/2c$).,

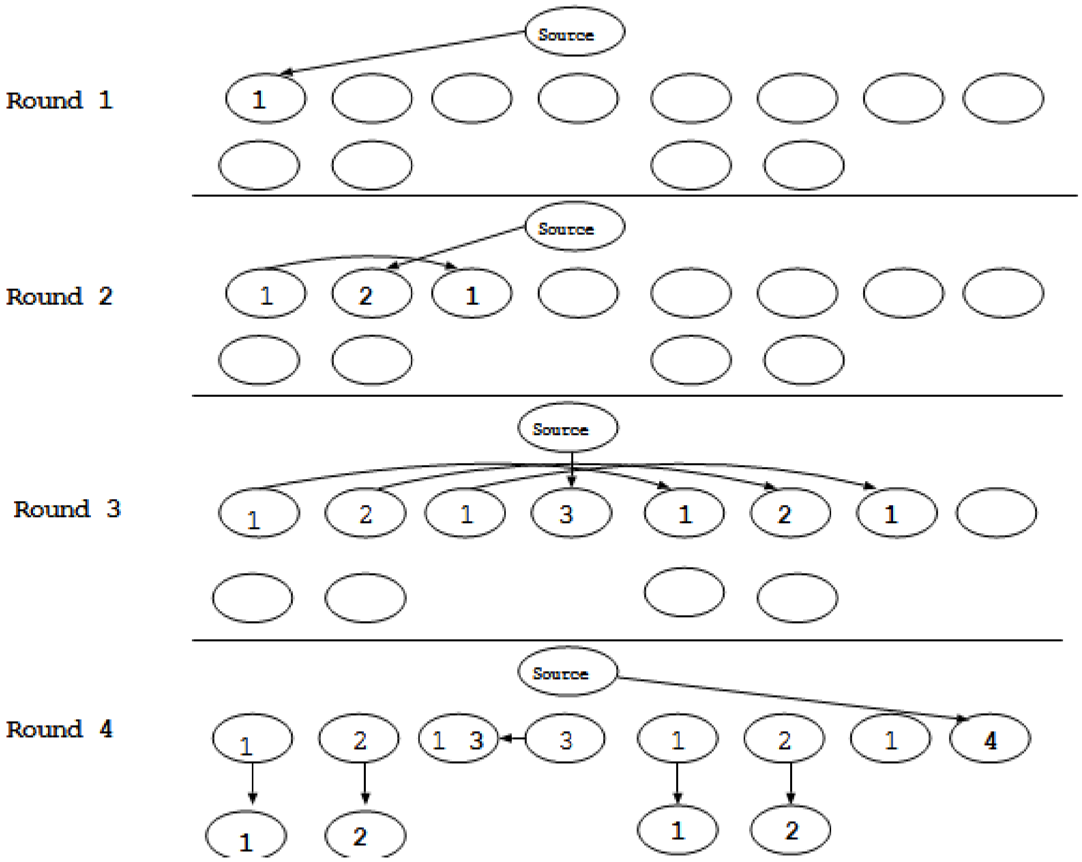

Figure 3.1: An example of a single-source broadcast instance.

It is easy to see that the second phase can be scheduled without conflicts as we deal with only one item in each round. For the first phase. For instance, if all D_i are identical and include all disks (thus the problem is the same as single-source broadcast [23, 29]) and $\Delta = 4$, $|D_i| = 12$ for each item i. At each round, the source disk makes a new copy. For other items, the numbers of copies are doubled if possible. Consider Round 4. Since there are four copies of item 1, only two copies need to be created to make $|D_i|/2 = 6$ copies. For items 2 and 3, we can double the number of copies, and a new copy for item 4 is created by the source disk. Without loss of generality, we assume that $|D_1| \geq |D_2| \geq \cdots \geq |D_\Delta|$ (otherwise renumber the items). Let d_i be the largest index such that $2^{di} \leq |D_i|$. For example, if $|D_i| = 12$, then $d_i = 3$.

IV. RESULT AND DISCUSSION

We prove that our algorithm uses at most Δ more rounds than the optimal solution for single-source multicasting. Let us denote the optimal makespan of a migration instance I as C(I).

Theorem 3.4.1 For any migration instance I, $C(I) \geq \max_1 \leq i \leq \Delta(i + blog|D_i|c)$.

Proof Consider the instance where there is no overlap among D_i's. After a disk in D_i receives ifrom s for the first time, we need at least $blog|D_i|c$ more rounds to make all disks in D_i receive ieven if s copies item iseveral times after the first copy. Therefore, $C(I) \geq \max_1 \leq i \leq \Delta(f(i)+blog|D_i|c)$ where f(i) is the round when D_i receives the first copy from s. Because s can be involved in copying only one item at a time, f(i) =6 f(j) if i=6 j. Also copying the same item from s more than once during the first Δ rounds will only increase f(i) of some sets. Therefore, C(I) can be minimized by choosing f(i) as a permutation of 1,...,Δ. Now we show that $\max_1 \leq i \leq \Delta(f(i) + blog|D_i|c) \geq \max_1 \leq i \leq \Delta(i + blog|D_i|c)$ for any permutation f(i). Suppose there is a set D_i that f(i) =6 iwhen $\max_1 \leq i \leq \Delta(f(i) + blog|D_i|c)$ is minimum. Let D_i be the set which have the smallest f(i) among such sets. Then f(i) <iand there should be a $D_j$such that j = f(i) and f(j) > j. Even if we exchange the order of two sets, the value does not increase because;

$\max(f(i) + blog|D_i|c, f(j) + blog|D_j|c) = f(j) + blog|D_j|c$

$\geq \max(j + blog|D_j|c, f(j) + blog|D_i|c)$. Thus when f(i) = ifor all i, $\max_1 \leq i \leq \Delta(f(i) + blog|D_i|c)$ is minimized. ut

Lemma 3.4.2 The total makespan of our algorithm is at most $\max_1 \leq i \leq \Delta(i + blog|D_i|c) + \Delta$.

Proof In phase I, D_i receives ifrom s at the i-th round for the first time. Because the number of copies doubles until it reaches $b|D_i|/2c$, the number of copies of item ireaches $b|D_i|/2c$ in i+ $blog|D_i|c$ rounds. Phase II takes at most Δ rounds because we finish one item in each round. Therefore, the lemma follows.ut**Corollary 3.4.3** The total makespan of our algorithm is at most the optimal makespan plus Δ.

Proof Follows from Lemma 3.4.1 and Lemma 3.4.2. ut

Theorem 3.4.4 We have a 2-approximation algorithm for the single-source multicasting problem.

Proof Because $\Delta \leq \max_1 \leq i \leq \Delta(i + blog|D_i|c)$, the algorithm is 2-approximation.tu

V. CONCLUSION

Broadcasting and gossiping problems resemble some of the data dissemination problems we considered. However, previous works have mainly concentrated on assuming two parties may exchange all the information they know in constant time, and assuming the underlying communication model is homogeneous. Our work addressed these assumptions, and provides a broadcasting algorithm that is more applicable on a wide-area network. Moreover, to collect data from a set of hosts in a large-scale public network like the Internet, we addressed two key problems: the available bandwidth can fluctuate, and the network may not choose the best route to transfer the data between two hosts. we considered the single-source multicast problem, where there is one source disk s that has all Δ items and others do not have any item in the beginning, and we would like to send item ito disks in set D_i. We developed an algorithm where D_i can be an arbitrary subset of disks.

The number of rounds required by our algorithm is at most Δ + OPT where OPT is the minimum number of rounds required for this problem. Our algorithm is obviously a 2-approximation for the problem, since Δ is a lower bound on the number of rounds required by the optimal solution.

REFERENCES

[1]. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. Network Flows: Theory, Algorithms, and Applications. Prentice Hall, Englewood Cliffs, NJ, 1993.

[2]. D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient overlay networks. In Proceedings of the 18th ACM Symposium on Operating Systems Principles, pages 131–145, Banff, Canada, Oct. 2001.

[3]. E. Anderson, J. Hall, J. Hartline, M. Hobbs, A. Karlin, J. Saia, R. Swaminathan, and J. Wilkes. An experimental study of data migration algorithms. In Proceedings of the 5thInternational Workshop on Algorithm Engineering, LNCS 2141, pages 145–158, New York, NY, 2001. Springer-Verlag.

[4]. T. Anderson, D. Culler, D. Patterson, and the NOW team. A case for NOW (networks of workstations). IEEE Micro, 15(1):54–64, Feb. 1995.

[5]. B. R. Badrinath and P. Sudame. Gathercast: An efficient mechanism for multi-point to

point aggregation in IP networks. Technical Report DCS-TR-362, Computer Science Department, Rutgers University, Piscataway, NJ, July 1998.

[6]. B. Baker and R. Shostak. Gossips and telephones. Discrete Mathematics, 2:191–193, 1972.

[7]. M. Banikazemi, V. Moorthy, and D. K. Panda. Efficient collective communication on heterogeneous networks of workstations. In Proceedings of the 1998 International Conference on Parallel Processing, pages 460–467. IEEE Computer Society, 1998.

[8]. A. Bar-Noy, S. Guha, J. Naor, and B. Schieber. Multicasting in heterogeneous networks. In Proceedings of the 30th Annual ACM Symposium on Theory of Computing, pages 448–453,1998.

[9]. A. Bar-Noy and S. Kipnis. Designing broadcast algorithms in the postal model for messagepassing systems. Mathematical Systems Theory, 27(5), 1994.

[10]. J. Bermond, L. Gargano, and S. Perennes. Optimal sequential gossiping by short messages.Discrete Applied Mathematics, 86(2–3):145–155, 1998.

[11]. J. Bermond, L. Gargano, A. A. Rescigno, and U. Vaccaro. Fast gossiping by short messages. SIAM Journal on Computing, 27(4):917–941, 1998.

[12]. P. Bhat, C. Raghavendra, and V. Prasanna. Efficient collective communication in distributed heterogeneous systems. Journal of Parallel and Distributed Computing, 63(3):251–263, 2003.

[13]. S. Bhattacharjee, W. C. Cheng, C.-F. Chou, L. Golubchik, and S. Khuller. Bistro: A platform for building scalable wide-area upload applications. ACM SIGMETRICS Performance Evaluation Review, 28(2):29–35, Sep. 2000.

[14]. J. A. Bondy and U. S. R. Murty. Graph Theory with applications. Elsevier, 1976.

[15]. J. Bruck, D. Dolev, C. Ho, M. Rosu, and R. Strong. Efficient message passing interface(MPI) for parallel computing on clusters of workstations. Journal of Parallel and Distributed Computing, 40(1):19–34, Jan. 1997.

[16]. R. T. Bumby. A problem with telephones. SIAM Journal on Algebraic and Discrete Methods,2(1):13–18, Mar. 1981.

[17]. K. Calvert, J. Griffioen, B. Mullins, A. Sehgal, and S. Wen. Concast: Design and implementation of an active network service. IEEE Trans. on Networking, 19:426–437, 2000.

[18]. R. L. Carter and M. E. Crovella. Measuring bottleneck link speed in packet-switched networks. Performance Evaluation, 27–28:297–318, 1996.

[19]. W. Cheng, C. Chou, L. Golubchik, S. Khuller, and Y.-C. Wan. A coordinated data collection approach: Design, evaluation, and comparison. IEEE Journal on Selected Areas in Communications—Special Issue on Design, Implementation and Analysis of CommunicationProtocols, 22(10):2004–2018, Dec. 2004.

[20]. W. C. Cheng, C.-F. Chou, L. Golubchik, S. Khuller, and Y.C. Wan. On a graph-theoretic approach to scheduling large-scale data transfers. Technical Report CS-TR-4322, University of Maryland, College Park, MD, Jan. 2002.

[21]. W.C. Cheng, C.-F. Chou, L. Golubchik, S. Khuller, and Y.-C. Wan. Large-scale data collection: A coordinated approach. In Proceedings of IEEE INFOCOM, Mar.–Apr. 2003.

[22]. C.-F. Chou, Y.-C. Wan, W. C. Cheng, L. Golubchik, and S. Khuller. A performance study of a large-scale data collection problem. In Proceedings of the 7th International Workshop on Web Content Caching and Distribution, pages 259–272, Aug. 2002.

[23]. E. J. Cockayne and A. G. Thomason. Optimal multi-message broadcasting in complete graphs. Utilitas Mathematica, 18:181–199, 1980.

[24]. D. E. Culler, R. M. Karp, D. A. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian, and T. von Eicken. LogP: Towards a realistic model of parallel computation. In Proceedings of the 4th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pages 1–12, 1993.

[25]. G. Dantzig. Linear Programming and Extensions. Princeton University Press, Princeton, NJ,1963.

[26]. Y. Dinitz, N. Garg, and M. Goemans. On the single source unsplittable flow problem. In Proceedings of 39th Annual IEEE Symposium on Foundations of Computer Science, pages 290–299, 1998.

[27]. A. B. Downey. Using pathchar to estimate Internet link characteristics. In Proceedings of ACM SIGCOMM 1999, pages 241–250, 1999.

[28]. M. Elkin and G. Kortsarz. A sublogarithmic approximation algorithm for the undirected telephone broadcast problem: a path out of a jungle. In Proceedings of the 14th AnnualACM-SIAM Symposium on Discrete Algorithms, pages 76–85, 2003.

[29]. A. M. Farley. Broadcast time in communication networks. SIAM Journal on Applied Mathematics, 39(2):385–390, 1980.

[30]. A. M. Farley and S. Hedetniemi. Broadcasting in grid graphs. In Proceedings of the 9thSE Conference on Combinatorics, Graph Theory and Computing, pages 275–288. Utilitas Mathematica, 1978.

[31]. A. M. Farley and A. Proskurowski. Gossiping in grid graphs. Journal on CombinatorialInformation Systems Science, 5:161–172, 1980.

[32]. L.R. Ford and D.R. Fulkerson. Flows in Networks. Princeton University Press, Princeton,NJ, 1962.

[33]. I. Foster and C. Kesselman. The Grid: Blueprint for a New Computing Infrastructure. MorganKaufmann, 1998.

[34]. P. Fraigniaud and E. Lazard. Methods and problems of communication in usual networks. Discrete Applied Mathematic, 53(1-3):79–133, 1994.

[35]. P. Francis, S. Jamin, V. Paxson, L. Zhang, D. Gryniewicz, and Y. Jin. Internet Distance MapService. http://idmaps.eecs.umich.edu/, 1999.

[36]. M. R. Garey, R. L. Grahma, and J. D. Ullman. An analysis of some packing algorithms. Combinatorial Algorithms (Courant Computer Science Symposium, No. 9), pages 39–47, 1972.

[37]. A. V. Goldberg. An efficient implementation of a scaling minimum-cost flow algorithm. Journal of Algorithms, 22(1):1–29, 1997.

[38]. A.V.Goldberg.AndrewGoldberg'sNetworkOptimizationLibrary.http://www.avglab.com/andrew/soft.html, 2001.

[39]. L. Golubchik, S. Khanna, S. Khuller, R. Thurimella, and A. Zhu. Approximation algorithms for data placement on parallel disks. In Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 223–232, 2000.

[40]. W. Gropp, E. Lusk, N. Doss, and A. Skjellum. A high-performance, portable implementation of the MPI message passing interface standard. Parallel Computing, 22(6):789–828, 1996.

[41]. A. Hajnal, E. C. Milner, and E. Szemeredi. A cure for the telephone disease. Canadian Mathematical Bulletin, 15(3):447–450, 1972.

[42]. J. Hall, J. Hartline, A. R. Karlin, J. Saia, and J. Wilkes. On algorithms for efficient data migration. In Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 620–629, 2001.

[43]. F. Harary and A. J. Schwenk. The communication problem on graphs and digraphs. Journal of The Franklin Institute, 297:491–495, 1974.

[44]. S.M. Hedetniemi, S.T. Hedetniemi, and A.L. Liestman. A survey of gossiping and broadcasting in communication networks. Networks, 18:319–349, 1988.

[45]. D.S. Hochbaum, editor. Approximation Algorithms for NP-Hard Problems. PWS Publishing,1997.

[46]. B. Hoppe and E. Tardos.´ Polynomial time algorithms for some evacuation problems. In Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 512– 521, 1994.

[47]. B. Hoppe and E. Tardos. The quickest transshipment problem. In´Proceedings of the 5thAnnual ACM-SIAM Symposium on Discrete Algorithms, pages 443–441, 1995.

[48]. J. Hromkovic, R. Klasing, B. Monien, and R. Peine. Dissemination of information in interconnection networks (broadcasting and gossiping). In D.-Z. Du and D.F. Hsu, editors, Combinatorial Network Theory, pages 125–212. Kluwer Academic Publishers, Netherlands, 1996.

[49]. http://www.cc.gatech.edu/fac/Ellen.Zegura/graphs.html.Georgia Tech Internetwork topology generator.

[50]. http://www.isi.edu/nsnam/ns/. The Network Simulator—ns-2.

[51]. C. A. J. Hurkens. Spreading gossip efficiently. NieuwArchiefvoorWiskunde, 5(1):208–210,2000.

[52]. P. Husbands and J. Hoe. MPI-StarT: Delivering network performance to numerical applications. In Proceedings of the IEEE/ACM SC98 Conference, page 17, Nov. 1998.

[53]. V. Jacobson. pathchar. http://www.caida.org/tools/utilities/others/pathchar/, 1997.

[54]. D. S. Johnson. Approximation algorithms for combinatorial problems. Journal of Computer and System Sciences, 9:256–278, 1974.

[55]. R. Karp, A. Sahay, E. Santos, and K. Schauser. Optimal broadcast and summation

in the LogP model. In Proceedings of the 5th Annual ACM Symposium on Parallel Algorithms and Architectures, pages 142–153, 1993.

[56]. S. Kashyap and S. Khuller. Algorithms for non-uniform size data placement on parallel disks.In Proceedings of the 23rd Annual Conference on Foundations of Software Technology andTheoretical Computer Science, LNCS 2914, pages 265–276, New York, NY, 2003. Springer-Verlag.

[57]. S. Khuller and Y. A. Kim. On broadcasting in heterogeneous networks. In Proceedings of the15th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1011–1020. Society for Industrial and Applied Mathematics, 2004.

[58]. S. Khuller, Y. A. Kim, and Y.-C. Wan. On generalized gossiping and broadcasting. In Proceedings of the 11th Annual European Symposium on Algorithms (ESA'03), LNCS 2832, pages 373–384, 2003.

[59]. S. Khuller, Y. A. Kim, and Y.-C. Wan. Algorithms for data migration with cloning. SIAMJournal on Computing, 33(2):448–461, 2004.

[60]. S. Khuller, Y. A. Kim, and Y.-C. Wan. Broadcasting on networks of workstations. Manuscript,2005.

[61]. S. Khuller, Y. A. Kim, and Y.-C. Wan. On generalized gossiping and broadcasting. Accepted for publications in Journal of Algorithms, 2005.

[62]. S. Khuller, Y. A. Kim, and G. J. Woeginger. Approximation schemes for broadcasting in heterogeneous networks. In Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'04), LNCS 3122, pages 163–170, 2004.

[63]. T. Kielmann, H. Bal, and S. Gorlatch. Bandwidth-efficient collective communication for clustered wide area systems. In Proceedings of the 14th International Symposium on Parallel and Distributed Processing, page 492, 2000.

[64]. T. Kielmann, R. Hofman, H. Bal, A. Plaat, and R. Bhoedjang. MagPIe: MPI's collective communication operations for clustered wide area systems. In Proceedings of the 7th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pages 131–140, 1999.

[65]. J. M. Kleinberg. Single-source unsplittable flow. In Proceedings of the 37th Annual IEEESymposium on Foundations of Computer Science, pages 68–77, 1996.

[66]. W. Knodel. New gossips and telephones. Discrete Mathematics, 13:95, 1975.

[67]. D. W. Krumme, G. Cybenko, and K. N. Venkataraman. Gossiping in minimal time. SIAMJournal on Computing, 21(1):111–139, 1992.

[68]. K. Lai and M. Baker. Nettimer: A tool for measuring bottleneck link bandwidth. In Proceedings of the USENIX Symposium on Internet Technologies and Systems, pages 123–134, Mar., 2001.

[69]. K. Lebensold.Efficient communication by phone calls. Studies in Applied Mathematics,52:345–358, 1973.

[70]. H. M. Lee and G. J. Chang. Set to set broadcasting in communication networks. DiscreteApplied Mathematics, 40(4):411–421, 1992.

[71]. T. Leighton. The challenges of delivering content on the Internet. In Proceedings of ACM SIGMETRICS 2001, page 214, Cambridge, MA, June 2001.

[72]. D. Liben-Nowell. Gossip is synteny: incomplete gossip and an exact algorithm for syntenic distance. In Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 177–185, 2001.

[73]. Pangfeng Liu. Broadcast scheduling optimization for heterogeneous cluster systems. Journal of Algorithms, 42(1):135–152, 2002.

[74]. B. Lowekamp and A. Beguelin. ECO: Efficient collective operations for communication on heterogeneous networks. In Proceedings of the 10th International Parallel Processing Symposium, pages 399–405, Honolulu, HI, Apr. 1996.

[75]. B. A. Mah. pchar. http://www.kitchenlab.org/www/bmah/Software/pchar/, 2001.

[76]. G. De Marco, L. Gargano, and U. Vaccaro. Concurrent multicast in weighted networks. In Proceedings of the 6th Scandanavian Workshop on Algorithm Theory, LNCS 1432, pages 193–204. Springer-Verlag, 1998.

[77]. Message passing interface forum, Mar. 1994.

[78]. K. Moore, J. Cox, and S.Green. SONAR—A network proximity service. IETF Internet-Draft,1996.

[79]. C. H. Papadimitriou. Computational complexity. Addison-Wesley, 1994.

[80]. J. Pruyne and M. Livny. Interfacing condor and PVM to harness the cycles of workstation clusters. Journal on Future Generations of Computer Systems, 12(1):67–85, 1996.
[81]. D. Richards and A. L. Liestman. Generalization of broadcasting and gossiping. Networks,18:125–138, 1988.
[82]. D. Rubenstein, J. F. Kurose, and D. F. Towsley. Detecting shared congestion of flows via end-to-end measurement. In Proceedings of ACM SIGMETRICS 2000, pages 145–155, 2000.
[83]. S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, and J. Zahorjan. Detour: A case for informed Internet routing and transport. IEEE Micro Magazine, 19:50–59, Jan. 1999.
[84]. S. Savage, A. Collins, and E. Hoffman. The end-to-end effects of Internet path selection. In Proceedings of ACM SIGCOMM 1999 Conference on Applications, technologies, architectures, and protocols for computer communication, pages 289–299, 1999.
[85]. S. Seshan, M. Stemm, and R.H. Katz. SPAND: Shared passive network performance discovery. In Proceedings of the 1st USENIXSymp. on Internet Technologies and Systems, pages 285– 294, Dec. 1997.
[86]. H. Shachnai and T. Tamir. On two class-constrained versions of the multiple knapsack problem. Algorithmica, 29:442–467, 2001.
[87]. P. J. Slater, E. Cockayne, and S.T. Hedetniemi. Information dissemination in trees. SIAMJournal on Computing, 10:692–701, 1981.
[88]. R. Tijdeman. On a telephone problem. Nieuw Archief voor Wiskunde, 19(3):188–192, 1971.
[89]. V. V. Vazirani. Approximation Algorithms. Springer-Verlag, 2001.
[90]. V. G. Vizing. On an estimate of the chromatic class of a p-graph (Russian). Diskret. Analiz.,3:25–30, 1964.
[91]. R. Wolski and M. Swany. Network Weather Service. http://nws.cs.ucsb.edu/, 1999.