

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 15-05-2007		2. REPORT TYPE Final		3. DATES COVERED (From - To) April 2006 - March 2007	
4. TITLE AND SUBTITLE Fundamentals of Combinatorial Optimization and Algorithm Design: Report for April 1, 2006 - March 31, 2007				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-06-1-0396	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Shepherd, Bruce Andrews, Matthew Chekuri, Chandra Wilfong, Gordon Zhang, Lisa				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
				8. PERFORMING ORGANIZATION REPORT NUMBER N00014-06-1-0396 - F001	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Alcatel-Lucent / Bell Laboratories 600 Mountain Ave Murray Hill, NJ 07974				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research One Liberty Center 875 North Randolph Street, Suite 1425 Arlington, VA 22203-1995				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
				12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release: distribution is Unlimited	
13. SUPPLEMENTARY NOTES None					
14. ABSTRACT The main research output this year included work on degree-constrained network flows which was accepted by STOC 2007. Shepherd will give a plenary talk on this topic at CANADAM. Lucent has also applied for a patent based on this work. Work was also done on protected buy-at-bulk network design problems by Antonakopoulos, Chekuri, Shepherd, and Zhang resulting in a manuscript, April 2007. Work was done by Andrews and Zhang on scheduling in wireless networks resulting in a manuscript. Work was done on flow-cut gaps in series-parallel graphs by Chekuri and Shepherd.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON Brian Moore
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) (908) 582 3970

MATHEMATICAL SCIENCES
BELL LABORATORIES
LUCENT TECHNOLOGIES
600 Mountain Ave.
Murray Hill, NJ 07974

**Fundamentals of Combinatorial Optimization and Algorithm
Design: Report for April 1 2006 -March 31 2007**

Amount Received: 50K

Duration: 1 Year

Key Personnel:

- Matthew Andrews, andrews@research.bell-labs.com
- Chandra Chekuri, 908-582 1204, chekuri@research.bell-labs.com
- Bruce Shepherd (PI), 908-582 4181, bruce.shepherd@mcgill.ca
- Gordon Wilfong, 908-582 3561, gtw@research.bell-labs.com
- Lisa Zhang, ylz@research.bell-labs.com

Funded through March 2007.

Authorization: Debasis Mitra, Vice-President

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

1 2005-06 Accomplishments

The main activities supported under this grant were research support (primarily for travel) and salary support (primarily for Shepherd and Wilfong). We summarize below the work completed and in progress during the year. Please contact Bruce Shepherd (bruce.shepherd@mcgill.ca) to obtain any of the work in progress as manuscripts.

We describe the main work accomplished this year.

The main research output this year included: (1) work on degree-constrained network flows (by Donovan, Shepherd, Vetta, Wilfong) which was accepted by STOC 2007, [25]. Shepherd will give a plenary talk on this topic at the upcoming CANADAM, Banff May 28-31, 2007. Lucent (Shepherd-Wilfong) also applied for a patent based on this work. (2) work on protected buy-at-bulk problems (by Antonakopoulos, Chekuri, Shepherd, Zhang) resulting in a manuscript, April 2007 [6]. (3) work of Andrews and Zhang on scheduling in wireless networks - see manuscript [2]. (4) work on flow-cut gaps in series-parallel graphs (by Chekuri, Shepherd), in progress.

One paper supported by the previous grant also appeared in journal form for the first time: Hardness of Robust Network Design, by C. Chekuri, G. Oriolo, M.G. Scutella, B. Shepherd, *Networks*, vol. 50, no. 1, (2007), 50-54

2 Degree-constrained network flows

We consider the single-sink multicommodity network flow problem. We have a directed network (graph) $G = (V, A)$ with sink node t . Each node $v \in V$ wants to route r_v units of flow to the sink; this is termed the *demand* of node v . Furthermore, each node $v \in V$ has a fixed uniform capacity (by scaling we may take this capacity to be 1 as we allow fractional r_v values). Our interest lies in examining *bounded degree* or *degree-constrained flows*, that is, feasible flows whose support graphs have bounded outdegree¹ at every node. Since we only examine node congestion problems we assume our graphs have no parallel (that is, in the same direction) arcs or loops. Consequently, the maximum outdegree of any node is less than n , where n is the number of nodes in the network. We classify such flows as follows. The class of network flows with outdegree at most d is denoted by \mathcal{C}_d ; we call such flows *d-furcated*. The cases $d \in \{1, 2, \infty\}$ are of particular interest to us². These flow classes are:

- \mathcal{C}_∞ (*Fractional Flows*): flow from a node v to the sink t may be routed fractionally along *any* path; in particular, v may send flow on any number of outgoing arcs.
- \mathcal{C}_1 (*Confluent Flows*): flow from v to t must be routed on a *unique* path; in particular, v sends flow on at most one outgoing arc.
- \mathcal{C}_2 (*Bifurcated Flows*): flow may be sent from v on at most two outgoing arcs.

¹Note that we make no restriction on the indegree of a node.

²Note that by assumption, we have $\mathcal{C}_\infty = \mathcal{C}_{n-1}$.

Bounded degree flows are natural and elegant combinatorial objects in their own right. Interest in them, however, is primarily motivated by certain distributed routing protocols. For example, consider how confluent flows are produced by the *open shortest path first* (OSPF) protocol. This protocol is essentially a distributed implementation of Dijkstra’s algorithm. Consequently, for a specific network destination t , it populates each router v ’s next hop entry for t with some neighbour u of v for which there is a shortest path from v to t through u . In this context “shortest” is determined with respect to some costs on the links (arcs) and, in intra-domain networks, these costs may be altered by the network operator to achieve better traffic flow through the network (for example, in work by Thorup et al.). Hence, under these constraints the collective flow destined for t is routed along a directed arborescence (rooted at t); that is, we have a confluent flow.

In most intra-domain networks, however, flows with higher but bounded degrees are allowed. For example, if there is more than one “shortest path” from v to t , operators may place two or more next hops for t in the routing table. Traffic to t is then typically split using a round-robin approach. This motivates the present work. We wish to develop some of the network flow theory underlying the basic question: what happens if we allow multiple next hops per destination in our routing tables? In particular, how does network performance improve as the permissible outdegree increases? To answer this question we must adopt some performance measure to compare our various traffic flows. A variety of such measures could be used; we follow a common approach of where the performance measure applied is worst case congestion of a node. The *congestion* of a node is its load divided by its capacity, where the *load* of a node v is just the total flow value through v (this includes the demand of v itself). Hence, in a uniform capacity network we wish to minimise the maximum load of a node.³

2.1 Our Results and Previous Work

Our goal is to assess the cost, in terms of congestion, of restricting network flows to only route on a bounded number of arcs out of any node. Specifically, what is the cost of routing using the confluent or d -furcated flow constraint? As in [11], [10], we consider uniform node capacity (equivalently, *uncapacitated*) networks exclusively. Suppose that G contains a fractional flow satisfying all the demands in which no node has load more than 1. One may then ask, is there a d -furcated flow that routes all the demands and has low congestion at every node?

Therefore we are interested in the *congestion gap*, $\gamma(d)$, of a flow class \mathcal{C}_d , which is the worst ratio, over any network and any set of demands, of the congestion of an optimal flow in \mathcal{C}_d to the congestion of an optimal flow in \mathcal{C}_∞ . This question was first considered by Chen, Rajaraman and Sundaram who showed there always exists a confluent flow with congestion $O(\sqrt{n})$. This was subsequently improved by Chen et al. [10] who proved a congestion bound of $O(\log n)$ and gave an example to show that this result is tight. (More precisely, they gave a bound of $O(\log k)$ where k is the number

³Congestion can also be defined in terms of congestion along a link. Note that for confluent flows, the maximum load on an arc in the network must occur on some link into the destination t . Thus if all link and node capacities are 1, then the worst case node and link congestion problems are identical.

of nodes with outgoing arcs to the sink.) Hence the congestion gap $\gamma(1)$ between fractional flows and confluent flows is $\Theta(\log n)$ in an uncapacitated (i.e., uniform capacity) network. Thus, the gap between flows in \mathcal{C}_1 and flows in \mathcal{C}_∞ is unbounded but, evidently, as the maximum outdegree of a flow is allowed to increase, the congestion gap tends to one. However, it was not known whether obtaining a bounded congestion gap required allowing an unbounded maximum degree. Perhaps surprisingly, we prove that a bounded congestion gap can be obtained with bounded outdegrees. In fact, the congestion gap is all but eliminated if we allow for bifurcated rather than confluent flows: Given a fractional flow of congestion one, there is a bifurcated flow with congestion at most two. Thus, our main result is that the congestion gap $\gamma(2)$ between flows in \mathcal{C}_2 and flows in \mathcal{C}_∞ is at most two in uncapacitated networks. We also show that this bound is tight. Moreover, our techniques show the rate at which the congestion gap is eliminated as d grows; the congestion gap $\gamma(d)$ between d -furcated flows and fractional flows is at most $1 + \frac{1}{d-1}$.

Our proof is algorithmic and so provides a factor 2 (respectively, factor $1 + \frac{1}{d-1}$) approximation algorithm for finding a minimum congestion bifurcated (respectively, d -furcated) flow in a single-sink multicommodity flow problem. Finally, we show that this problem is maxSNP-hard.

3 Protected Buy at Bulk

The telecommunications industry is the inspiration for numerous network optimization problems. In this paper, we consider buy-at-bulk network design problems that arise in the design and operation of modern optical core networks [15]. These networks are characterized by the following two salient features: (i) very high capacity achieved via DWDM (Dense Wavelength Division Multiplexing) based optical transmission technology and (ii) expensive equipment exhibiting economies of scale. In such networks, each link carries enormous amounts of traffic and hence the failure of a link or a node represents an unacceptable degradation of service. Therefore, fault tolerance is an integral part of the design. Although there are a variety of ways to ensure fault tolerance, one of the most commonly used solutions in optical core networks is to set up, for each commodity, so-called *dedicated* or *1+1* protection. This amounts to reserving a pair of disjoint paths between the source and destination nodes of each commodity. The popularity of the 1+1 model comes from its operational simplicity and high restoration speed. Disjointness may be defined in several ways, according to requirements of the commodity in question. For instance, the commonly used measures include “site-disjointness”, where the two paths do not share any common nodes; edge-disjointness, where the two paths do not share any common links; and cable or fiber-disjointness, where the two paths must use distinct fibers/cables if they go through the same link. In this context, a central problem faced by network operators and equipment vendors is to build a cost-effective and bandwidth-efficient network that supports a multitude of traffic at the desired level of protection. The network operators look to utilize their network resources as efficiently as possible, and the equipment vendors seek to find innovative cost advantages to obtain a competitive edge in bidding for contracts from the network providers.

We give a formal description of the optimization problem that abstracts the above problem. The input consists of an undirected edge-weighted graph $G = (V, E)$ and a set of h node pairs $s_1 t_1, s_2 t_2, \dots, s_h t_h$ representing different traffic demands. Each pair has a non-negative demand value $d(i)$ that needs to be routed between s_i and t_i and also specifies a protection requirement. In this paper we restrict our attention to the 1+1 model and, for simplicity, we assume each demand requires node-disjoint protection. A feasible solution consists of a collection of path pairs $(P_1, Q_1), \dots, (P_h, Q_h)$, where P_i and Q_i are internally node-disjoint paths between s_i and t_i and each carries a reserved bandwidth of $d(i)$. If e is an edge of the network used by any of these paths, and they induce a requirement of (say) b_e units of bandwidth on e , then equipment that can support this requirement has to be purchased. Now, let us discuss the cost model $f(b_e)$ for purchasing bandwidth. In this paper we focus on a simple cost model, namely the single-cable cost model: bandwidth can be purchased in *cables*, i.e. units of fixed capacity k , at a price that varies only by edge. Thus, if the cost of purchasing one cable on edge e is c_e , then $f(b_e) = \lceil \frac{b_e}{k} \rceil c_e$. The objective is to minimize the total cost $\sum_e f(b_e)$ over all possible choices of $(P_1, Q_1), \dots, (P_h, Q_h)$. The single-cable cost function closely models DWDM networks, where each optical fiber carries the same number of wavelengths k and each edge e has a fixed cost c_e for deploying an extra fiber (see [15]).

Observe that, even in the single-cable setting, the buy-at-bulk problem captures as special cases some well-known NP-hard connectivity problems, such as the minimum-cost Steiner tree and the minimum-cost Steiner forest problems. Moreover, Andrews [1] has shown that even the single-cable problem without protection constraints is hard to approximate to within an $\Omega(\log^{1/4-\epsilon} n)$ factor; this separates the approximability of the buy-at-bulk problem from those of connectivity problems. In the connectivity setting, survivability and protection constraints have long been studied and include classical problems. Jain [33] devised the important iterative rounding method that yields a 2-approximation algorithm for the survivable network design problem (SNDP), in which the goal is to find a minimum-cost subgraph that satisfies given edge connectivity requirements between each pair of nodes in a graph. In [28] this technique was extended to handle node connectivity, when the requirements are restricted to be in the set $\{0, 1, 2\}$.

Buy-at-bulk network design without protection has received substantial attention in the past decade, including some recent work on super-constant lower bounds in the simplest setting [1] and poly-logarithmic upper bounds in the most general non-uniform setting [13, 14]. On the contrary, the variant with protection has not been so far considered in the literature on approximation algorithms. One reason for this is the difficulty of the buy-at-bulk problem, even without protection constraints. Although the first approximation algorithm for the multiple-cable setting appeared in 1997 [7], the algorithm was based on a technique that was not sufficiently flexible. It is only recently that alternative algorithms [9, 13] were developed that not only handled the non-uniform cost functions, but also provided new algorithmic approaches and insights. Further, for SNDP, the iterative rounding method of Jain [33] and the earlier primal-dual approach [39] strongly rely on the structural properties on the underlying linear program, which do not hold for the buy-at-bulk problem.

Our primary motivation to study this problem arose while developing a sequence

of optical network design tools at Bell Labs. We realized the ubiquity of the 1+1 model in practice, the lack of theoretical understanding of protected buy-at-bulk network design and a dearth of useful heuristic methods for the problem. Most algorithms used in practice are based on simple ad hoc methods combining greedy algorithms, local improvement and some enumeration. We hope this paper serves as a starting point in addressing the challenges from the theoretical point of view, as well as in providing insights that lead to more sophisticated and effective heuristics.

Results. We give approximation algorithms for buy-at-bulk network design in the 1+1 protection model for the single-cable setting. Observe that the 1+1 edge-disjoint protection problem can be reduced in a straightforward fashion to the 1+1 node-disjoint protection problem. In fact, for the edge-disjoint case the arguments can be substantially simplified; however, our focus here is on the node version, as it is also of greater practical significance.

Our first result is for the single-sink problem. This is the special case of the problem where all the pairs have one terminal node in common. In other words, the pairs are st_1, st_2, \dots, st_h and s is the common sink. Note that this problem is APX-hard. We present an $O(1)$ approximation algorithm for it and also establish an $O(1)$ integrality gap for a natural linear programming relaxation.

Our second result is an $O(\log^3 h)$ approximation for the multi-commodity problem. In particular, we show that that an α approximation for the single-sink problem via a natural LP relaxation yields an $O(\alpha \log^3 h)$ approximation for the multi-commodity problem, and combine this with our result above.

At present, the table below summarizes the known results on buy-at-bulk network design.

4 Flow-Cut Gap for Series Parallel Graphs

Partly motivated by their recent work with Khanna on throughput maximization (see appendix) Chekuri and Shepherd have been studying whether planar graphs have a constant flow-cut gap. That is, does there exist a constant $C \geq 1$ such that the following holds. For any instance G, H of supply/demand graphs is it the case that if the cut condition for H holds in G , then one can fractionally route $\frac{1}{C}$ amount of flow simultaneously in G for each demand edge $f \in E(H)$. Even the case of series-parallel graphs, the result was nontrivial to establish [30]. This previous work relies on metric embeddings. In contrast, Chekuri and Shepherd are examining the primal formulation

	Single sink	Multi-Commodity
Uniform	$O(1)$ [29] $\Omega(1)$	$O(\log n)$ [7] $\Omega(\log^{1/4-\epsilon} n)$ [1]
NonUniform	$O(\log h)$ [35] $\Omega(\log \log n)$ [24]	$O(\log^4 h)$ [13] $\Omega(\log^{1/2-\epsilon} n)$ [1]

Table 1: Approximability of unprotected buy-at-bulk network design.

where one simultaneously seeks a routing in G with good integrality properties.
This work is in progress.

Workshop/Conference Travel

Matthew Andrews:

1. Workshop on: *Adversarial modeling and analysis of communication networks*, Bertinoro, Italy, Nov. 2006.

Chandra Chekuri:

1. 38th Symposium on the Theory of Computing (STOC 2006), Seattle, WA, May 21-23, 2006.

Lisa Zhang:

1. Workshop on: *Adversarial modeling and analysis of communication networks*, Bertinoro, Italy, Nov. 2006.
2. International Math. Programming Symposium, Rio de Janeiro, Brazil, July 30-August 4, 2006.

References

- [1] M. Andrews. Hardness of buy-at-bulk network design. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science*, pages 115–124, Rome, Italy, October 2004.
- [2] M. Andrews and Lisa Zhang. Scheduling algorithms for wireless data systems. Manuscript, 2006.
- [3] Matthew Andrews and Lisa Zhang, *Wavelength assignment in optical networks with fixed fiber capacity*, ICALP (Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, eds.), Lecture Notes in Computer Science, vol. 3142, Springer, 2004, pp. 134–145.
- [4] Matthew Andrews and Lisa Zhang, *Bounds on fiber minimization in optical networks*, Proceedings of 2005 IEEE INFOCOM. Miami, FL, March 2005., 2005.
- [5] Elliot Anshelevich and Lisa Zhang, *Path decomposition under a new cost measure with applications to optical network design*, ESA (Susanne Albers and Tomasz Radzik, eds.), Lecture Notes in Computer Science, vol. 3221, Springer, 2004, pp. 28–39.
- [6] S. Antonakopolous, C. Chekuri, B. Shepherd, L. Zhang. Buy-at-bulk network design with protection. Manuscript, April 2007.
- [7] B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 542–547, Miami Beach, FL, October 1997.
- [8] W. Ben-Ameur, H. Kerivin, Routing of uncertain demands, *Optimization and Engineering*.
- [9] M. Charikar and A. Karagiozova. On non-uniform multicommodity buy-at-bulk network design. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, Baltimore, MD, May 2005.
- [10] J. Chen, R. Kleinberg, L. Lovasz, R. Rajaraman, R. Sundaram and A. Vetta, “(Almost) tight bounds and existence theorems for confluent flows”, *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, pp529-538, 2004.
- [11] J. Chen, R. Rajaraman, and R. Sundaram, “Meet and merge: approximation algorithms for confluent flow”, *Proceedings of the 35th ACM Symposium on Theory of Computing (STOC)*, pp373-382, 2003.
- [12] S. Chawla, R. Krauthgamer, Ravi Kumar, Y. Rabani, and D. Sivakumar. On the Hardness of Approximating Multicut and Sparsest-Cut. Submitted, November 2004.
- [13] C. Chekuri, M. Hajiaghayi, G. Kortsarz, and M. Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design problems. In *Proceedings of the 47th Annual Symposium on Foundations of Computer Science*, 2006.

- [14] C. Chekuri, M. Hajiaghayi, G. Kortsarz, and M. Salavatipour. Approximation algorithms for node-weighted buy-at-bulk network design. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2007
- [15] C. Chekuri, P. Claisse, R. Essiambre, S. Fortune, D. Kilper, W. Lee, K. Nithi, I. Saniee, F.B. Shepherd, C. White, G. Wilfong, and L. Zhang. Design tools for transparent optical networks. *Bell Labs Technical Journal*, 11(2):129–143, 2006.
- [16] C. Chekuri, S. Khanna, and F.B. Shepherd. The All-or-nothing multicommodity flow problem, *Proceedings of the 36th ACM Symposium on Theory of Computing*, 2004.
- [17] C. Chekuri, S. Khanna, and F.B. Shepherd, Edge-disjoint paths in planar undirected graphs, *FOCS*, 2005.
- [18] C. Chekuri S. Khanna and F.B. Shepherd. Multicommodity Flow and Treewidth Manuscript Jan. 2006.
- [19] C. Chekuri, S. Khanna and F.B. Shepherd. Edge-disjoint paths in planar graphs with constant congestion *STOC 2006*.
- [20] C. Chekuri, S. Khanna and F.B. Shepherd. An $O(\sqrt{rtn})$ Approximation and Integrality Gap for Disjoint Paths and Unsplittable Flow. *Theory of Computing 2/7*, pp. 137-146 (2006)
- [21] C. Chekuri and F.B. Shepherd, Splitting-off for Approximate Integer Decomposition of Network Design Problems. Submitted.
- [22] C. Chekuri, G. Oriolo, M.G. Scutella, and F.B. Shepherd. Hardness of Robust Network Design, To appear *Networks*. Preliminary version in *INOC 2005*.
- [23] C. Chekuri, S. Khanna, and F.B. Shepherd, Multicommodity flow, well-linked terminals, and routing problems, *STOC 2005*.
- [24] J. Chuzhoy, A. Gupta, J. Naor, and A. Sinha. On the approximability of some network design problems. *Proc. of ACM-SIAM SODA*, 943–951, 2005.
- [25] P. Donovan, B. Shepherd, A. Vetta, G. Wilfong, Degree-constrained network flows, *STOC 2007*.
- [26] Cooperative Data-Optical InterNetworking: Distributed Multilayer Optimization. E. Elwalid, D. Mitra, Q. Wang. *Proceedings of INFOCOM 2006*.
- [27] Thomas Erlebach, Aris Pagourtzis, Katerina Potika, and Stamatis Stefanakos, *Resource allocation problems in multifiber wdm tree networks*, WG 2003, 2003, pp. 218–229.
- [28] L. Fleischer, K. Jain, and D.P. Williamson. An iterative rounding 2-approximation algorithm for the element connectivity problem. *JCSS*, 72(5):838–867, 2006.

- [29] S. Guha, A. Meyerson, and K. Munagala. A constant factor approximation for the single sink edge installation problem. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 383–388, 2001.
- [30] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, Trees and 11 Embeddings. *Combinatorica*, 24(2):233–269, 2004. Preliminary version in *Proc. of FOCS*, 1999.
- [31] Venkat Guruswami, Sanjeev Khanna, Raj Rajaraman, Bruce Shepherd, Mihalis Yannakakis. Near-Optimal Hardness Results and Approximation Algorithms for Edge-Disjoint Paths and Related Problems, *Journal of Computer and System Sciences*, Volume 67, Issue 3, Nov. 2003, Pages 473–652. (also in *Proc. 31st Annual ACM Symposium on Theory of Computing (STOC)* , pp. 19–28, 1999)
- [32] Mohammad Hajiaghayi, Robert Kleinberg, Tom Leighton, and Harald Racke. Oblivious Routing in Node-Capacitated and Directed Graphs. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 782–790.
- [33] K. Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21 (1), 39–60, 2001.
- [34] A. McGregor and B. Shepherd. Island hopping and path colouring, with applications to WDM network design Bell Labs Technical Report. November 2005.
- [35] A. Meyerson, K. Munagala, and S. Plotkin. Cost-distance: Two metric network design. *Proc. of IEEE FOCS*, 383–388, 2000.
- [36] H. Okamura and P.D. Seymour. Multicommodity Flows in Planar Graphs *Journal of Combinatorial Theory B*, 31:75–81, 1981.
- [37] Bruce Shepherd and Adrian Vetta., *Lighting fibers in a dark network.*, *Journal on Selected Areas in Communication* **22** (2004), no. 9, 1583–1588.
- [38] Peter Winkler and Lisa Zhang, *Wavelength assignment and generalized interval graph coloring.*, *SODA*, 2003, pp. 830–831.
- [39] D.P. Williamson, M.X. Goemans, M. Mihail and V.V. Vazirani. A primal-dual approximation algorithm for generalized Steiner network problems. *Combinatorica*, 15:534–454, December 1995.

Background and Findings on Related Problems in Previous Year's Work.

Integrity Gaps and Approximation Algorithms for Maximum Throughput Problems

In a throughput problem we are given a graph G and terminals $s_i, t_i, i = 1, 2, \dots, k$, and we wish to find a maximum *routable subset* of $\{1, 2, \dots, k\}$. The notion of “routability” depends on the application at hand. For instance, in the edge-disjoint paths problem (EDP) a subset S is routable, if there is a collection of edge-disjoint paths $P_i : i \in S$, such that each P_i joins s_i, t_i . EDP is one of the most fundamental problems in combinatorial optimization. Apart from its applications in VLSI design and network design, it is also intrinsic to many approaches for solving other applied problems such as scheduling. For instance, Shepherd and Matthew Andrews created a scheduling system within Lucent, based on EDP.

In [31] it was shown that in directed graphs it is hard to approximate (in polytime, if P not equal to NP) the optimum to within a factor of $m^{5-\epsilon}$ for any $\epsilon > 0$, and hence EDP is terribly difficult with respect to polytime approximations. Their reduction however breaks down if small edge congestions are allowed, and also breaks down for undirected graphs. In fact, the gap for undirected graphs is only known to lie in the range 2 to $m^{5-\epsilon}$.

In [16], we show the first positive results in this direction by showing that the all-or-nothing multicommodity flow problem does admit a poly-logarithmic (i.e., polynomial in the variable $\log(n)$) approximation for general graphs – please refer to the paper: <http://cm.bell-labs.com/cm/ms/who/bshep/pub.html>. (The all-or-nothing flow problem is the throughput problem where a subset is routable, if the $s_i - t_i$ pairs admit a multicommodity flow. Note that unlike EDP, determining whether a subset is routable is easy for this problem, thus the hard part is the subset selection aspect.) The techniques in this paper used recent results of Räcke on oblivious routing as well as some interesting graph theoretic clustering results in 2-connected graphs.

Ideas from the work on all-or-nothing flow inspired the authors to take a look at EDP itself. Using schemes of Robertson, Seymour and Thomas, the authors have since proved the following theorem in planar graphs. We call a set X *well-linked* in G , if for any subset S with $|S \cap X| \leq |X - S|$, we have $|\delta(S)| \geq |S \cap X|$. We show that if S is well-linked in a planar graph G , there is a constant C (about 10,000 at present unfortunately) such that for any matching M of size $|S|/C$ on S , we can find paths connecting the endpoints of M , such that each edge lies in at most 2 of them. This result implies (again using Räcke's results) that EDP can be approximated to within a polylogarithmic factor where only \sqrt{n} was previously best known. Alternatively, this can be viewed as saying that the natural LP formulation for EDP has a polylogarithmic integrality gap for planar graphs with all capacities at least 2. In contrast, if some capacities are at most 1, then this gap may be $\Omega(\sqrt{n})$, exponentially larger. Thus, one goal for further work is to find cutting planes that tightens the integrality gap for arbitrary capacities. We mention that recently J. Kleinberg has used our framework to strengthen the planar result to the case of Eulerian planar graphs (not just all capacities

at least 2). The techniques used, also show that there is a constant approximation for product multicommodity flow in planar graphs, thus giving a new proof of an earlier result of Klein, Plotkin and Rao for uniform multicommodity flow in planar graphs. This work appears in [17].

Continuing work has recently involved exploring decomposition methods for high-degree constant expansion graphs. Some progress has been made in particular on shortening of the proofs of Räcke's celebrated result (however this thread is still under development!) Most recently, the authors have found a direct decomposition into well-linked sets that avoids the use of Räcke's decomposition. One interesting side-effect is that it allows the authors to obtain similar polylog-approximation algorithms for the node-disjoint versions of the problem. This is not possible by Räcke's decomposition since in that case $\Omega(\sqrt{n})$ gaps are known in the node versions of oblivious routing [32]. This work is appearing in STOC 2005 in the paper *Multicommodity flow, well-linked terminals, and routing problems* [23].

Robust Network Optimization

During Gianpaolo Oriolo's visit to Bell Labs, the authors explored the problem of robust network design.

Network designers have traditionally adopted the view that an accurate estimate for point-to-point traffic is given a priori. With the increasing importance of flexible services (such as VPNs or remote storage/computing), there has been increasing interest in the design of networks for situations where traffic patterns are either not well known a priori or changing rapidly. In these settings the network should be "dimensioned" (i.e., assigned capacity) to support not just one traffic matrix, but a larger class of matrices determined by the application. This results in a *robust optimization* problem, where we are given a *universe* \mathcal{U} of demand matrices (normally specified as a convex region), and the goal is to design a minimum cost network so that every demand matrix in \mathcal{U} can be *supported*. The simplest form of this problem where fractional capacities are allowed was introduced by Ben-Ameur and Kerivin [8], but only recently was it shown to be NP-hard [22] by Chekuri, Oriolo, Scutella and Shepherd. The problem considered is how to allocate fractional link capacities that are sufficient to support every demand, i.e., so that there is a multicommodity flow for every demand matrix in the universe \mathcal{U} . It is perhaps the fact that capacities are allowed to be fractional that makes this result somewhat surprising.

The main complication in proving hardness is the fact that there is no max-flow min-cut theorem for multicommodity flow. To overcome this, the authors introduce a "trick" of considering a more simply analyzed class \mathcal{U} of matrices they call *single-source hose demand matrices*. This single source problem is defined as follows. There is a given root node r with a specified marginal traffic value b_r , and each other node also has a marginal value b_v . The authors normally consider that b_v 's are 0, 1 and that b_r is some integer. A matrix d is then a *single-source hose demand* (i.e., it will be in our universe \mathcal{U}) if $d_{ij} = 0$ unless $i = r$ and $\sum_j d_{rj} \leq b_r$ and each $d_{ri} \leq b_i$. We

then ask for the minimum fractional capacity so that every such single-source demand matrix can be routed.

There are some special cases of this single-source robust design problem of interest. In the case where all b_v 's are 1 and b_r is also, then the optimal fractional capacity allocation is obviously just a minimum spanning tree. If only some of the b_v 's are 1, then it is the fractional relaxation for the undirected Steiner tree problem. In the case, where all b_v 's are 1, and $b_r = n$, then the optimal solution is a shortest path tree rooted at r . So at the extreme values of $b_r = 1$ or n , the problem is well-known to be polytime computable.

The authors show that the problem becomes hard for values of $b_r = cn$ where $c \in [1/2, 1)$, and suspect this continues to hold for every $c \in (0, 1)$. In this case, [22] gives a reduction from checking whether a graph has the expander property. In fact, based on results in [12] the authors show that (assuming a certain conjecture in complexity theory), if $b_r = n/2$, then for any $\epsilon > 0$ it is hard to approximate the network to within a factor of $2 - \epsilon$. Moreover, they exhibit a matching 2-approximation algorithm in this special case.