Research Statement

Erran L. Li (Li Li)
Room 4F535
Bell-Labs, Lucent
Holmdel, NJ 07733
Tel: +1 732-949-1282
erranlli@dnrc.bell-labs.com
http://www.bell-labs.com/user/erranlli/

My main research interests are networking, systems and distributed computing. My work involves both designing practical algorithms with sound theoretical foundations and building prototype systems. My work has spanned the application, routing, and transport layers. The main thrust of my work is developing efficient and enabling solutions for real networking problems. I briefly describe my current and previous work, and then discuss my plans for the future. Most of my work has been performed jointly with colleagues and all of it has benefited from their support and insight.

1. DISTRIBUTED NETWORK MONITORING WITH BOUNDED LINK UTILIZATION

The increasing demands for network Quality of service (QoS) support is placing a critical need on network management systems to monitor network performance. I address the problem of optimizing a scalable distributed polling system. The goal of the optimization is to reduce the cost of deployment of the measurement infrastructure by identifying a minimum poller set subject to bandwidth constraints on the individual links. I show that this problem is NP-hard and propose three different heuristics for solving it. I evaluate heuristics on both hierarchical and flat topologies with different network sizes under different polling bandwidth constraints. I have found that the heuristic of choosing the poller that can poll the maximum number of un-polled nodes was the best approach. My simulation studies show that the results obtained by the best heuristic is close to the lower bound obtained using LP relaxation [1].

2. PATH RESTORATION IN LABEL SWITCHED NETWORKS

The emerging Multi-Protocol Label Switching (MPLS) networks enable network service providers to route bandwidth guaranteed paths between customer sites. This basic Label Switched Path (LSP) routing is often enhanced using restoration routing, which sets up alternate LSPs to guarantee uninterrupted connectivity in case network links or nodes along the primary path fail. I address the problem of distributed routing of restoration paths, which can be defined as follows: given a request for a bandwidth guaranteed LSP between two nodes, find a primary LSP and a set of backup LSPs that protect the links along the primary LSP. A routing algorithm that computes these paths must optimize the restoration latency and the amount of bandwidth used. We introduce the concept of "backtracking" to bound the restoration latency. I consider three different cases characterized by a parameter called backtracking distance D: (1) no backtracking ($D = 0$), (2) limited backtracking ($D = k$), and (3) unlimited backtracking ($D = \infty$). I first show that joint optimization of primary and backup paths is NP-hard in all cases. I then consider algorithms that compute primary and backup paths in two separate steps. Using link cost metrics that capture bandwidth sharing, I devise heuristics for each case. My simulation study shows that these algorithms offer a way to tradeoff bandwidth to meet a range of restoration latency requirements [2].

3. IMPROVING THROUGHPUT IN 3G 1X-EVDO NETWORKS

In 3G data networks, the channel quality of a user is location dependent and time varying. Due to proportional fairness scheduling, users with poor channel quality receive low data rates and therefore bring down the total cell throughput. Existing network architectures for improving cell throughput require new hardware elements and therefore are prohibitively expensive. I propose a highly deployable software-only solution that improves the cell throughput with the simultaneous use of ad hoc networking technology in the 802.11 unlicensed spectrum. The basic idea is to select a proxy with good channel quality for a user with a poor channel. The packet is first delivered to the proxy and then forwarded to the actual destination through multiple wireless hops in the 802.11 unlicensed spectrum. I demonstrate the effectiveness of the architecture by enhancements specific to the HDR technology. The average network capacity of HDR downlink channel (about 600 Kbps) is far from the maximum (2.4 Mbps) that it can achieve. Simulations show that the architecture can improve the throughput of a HDR cell by upto 150%.

4. TOPOLOGY CONTROL AND ROUTING IN AD HOC NETWORKS

Networks that do not have a fixed infrastructure are called *ad hoc networks*. Ad hoc networks are expected to be widely deployed in applications such as community networks, disaster relief, tetherless classrooms, environment monitoring and forecasting, and battlefield situations. In ad hoc networks, not only is the power supply of individual nodes limited, but wireless bandwidth is also limited. In addition, channel conditions can vary greatly. Moreover, since nodes can be mobile, routes may constantly change. Thus, efficient protocols must take into account limited resources like network bandwidth and the energy available to each node. My thesis work focuses on the design and evaluation of algorithms and protocols to improve the performance of ad hoc networks. Specifically, I have developed novel topology-control and routing algorithms to reduce energy consumption of network nodes and increase network throughput.

4.1 Topology Control in Ad Hoc Networks

Most work on ad hoc networks has assumed fixed transmission power for all nodes. However, such fixed-power schemes have several drawbacks. In areas where the network is dense, throughput and energy-efficiency will be low due to severe interference; in areas where the network is sparse, the network may be disconnected. I have developed topology-control algorithms that aim to maintain global connectivity, conserve energy consumption, and balance throughput. I have worked on two types of topology control algorithms:

- Minimum-energy topology control algorithm using location information

  To reduce energy consumption, a good heuristic for topology control is to preserve the minimum-energy path between any source-destination pair. In [3], I identify the necessary and sufficient conditions to achieve a minimum-energy communication network. I use this characterization to construct a protocol that preserves the minimum energy property using only local information. Through detailed simulation, I have shown that the protocol significantly outperforms the one proposed earlier in the literature.

- Minimum energy competitive topology control algorithms using only directional information

  Location information is usually provided by GPS. A GPS unit can be expensive, and may not be available to every node. In [4, 5], based solely on directional information, I propose a simple distributed algorithm in which each node makes local decisions about its transmission power; these local decisions collectively guarantee global connectivity. Roughly speaking, the basic idea of the algorithm is that a node u transmits with the minimum power $p_{u,\alpha}$ required to ensure that in every cone of degree α around u, there is some node that u can reach with power $p_{u,\alpha}$. I show that taking $\alpha = 5\pi/6$ is a necessary and sufficient condition to guarantee that network connectivity is preserved. More precisely, if there is a path from s to t when every node communicates at maximum power then, if $\alpha \leq 5\pi/6$, there is still a path in the smallest symmetric graph G_α containing all edges (u, v) such that u can communicate with v using power $p_{u,\alpha}$. On the other hand, if $\alpha > 5\pi/6$, connectivity is not necessarily preserved. I also propose a set of optimizations that further reduce power consumption and prove that they retain network connectivity. simulation results demonstrate the effectiveness of the algorithm and the optimizations.

Recently I have designed an analytical framework for evaluating the performance of topology control algorithms using overall network throughput, and total energy consumption per packet delivered, as the metrics. The goal is to identify scenarios under which topology control can be used to improve the network performance. Based on studies with various traffic patterns and network loads, I have clearly identified scenarios under which network performance of ad hoc networks can be improved by using topology control [6].

4.2 Gossip-Based Ad Hoc Routing

Due to the changing and high variability nature of ad hoc networks, routing traffic can be a significant overhead and impediment to application performance. Many *ad hoc* routing protocols are based on (some variant of) flooding. Despite various optimizations, many routing messages are propagated unnecessarily. In [7], I propose a gossip-based approach to reduce the overhead of the routing protocols. The major findings are as follows. In large networks, there is a threshold effect: in almost every execution of the gossiping protocol, almost every node gets the message or almost no node does. The fraction of executions in which almost every node gets the message depends on the gossiping probability and the topology of the network. In the networks I have considered, using gossiping probability between 0.6 and 0.8 suffices to ensure that almost every node gets the message in almost every execution. For large networks, the simple gossiping protocol uses up to 35% fewer messages than flooding, with improved performance. Gossiping can also be combined with various optimizations of flooding to yield further benefits. My simulations show that adding gossiping to AODV results in significant performance improvement, even in networks as small as 150 nodes. The improvement should be even more significant in larger networks.

5. IP PAGING SERVICE FOR MOBILE HOSTS

In wireless networks with fixed infrastructure, mobile hosts must update the network with their current location in order to get packets delivered to them. Paging facilitates efficient power management at the mobile host by allowing the host to update the network less frequently at the cost of providing the network with only approximate location information. The network determines the exact location of a mobile host through paging before delivering packets destined to the mobile host. In current circuit-switched wireless networks, paging is implemented as a special purpose functionality in a centralized component inside the network. Given the emergence of different packet-switched wireless networks, in [8] I propose a novel router service called IP paging. This enables one common IP-based infrastructure to support different wireless interfaces such as CDMA, GPRS, Wireless LAN, etc. The major contributions are the design, implementation, and detailed performance evaluation, using measurements and simulation, of three IP-based paging protocols for mobile hosts.

6. RESOURCE ALLOCATION IN WIRELESS MULTIMEDIA NETWORKS

The allocation of scarce spectral resources to support as many user applications as possible while maintaining reasonable quality of service is a fundamental problem in wireless communication. In [9], I argue that the problem is best formulated in terms of decision theory. I propose a scheme that takes decision-theoretic concerns (like preferences) into account and discuss the difficulties and subtleties involved in applying standard techniques from the theory of Markov Decision Processes (MDPs) in constructing an algorithm that is decision-theoretically optimal. As an example of the proposed framework, I construct such an algorithm under some simplifying assumptions. Additionally, I present analysis and simulation results that show that the algorithm meets its design goals. Finally, I investigate how far from optimal one well-known heuristic is. The main contribution is in providing insight and guidance for the design of near-optimal admission-control policies.

7. HIGH-PERFORMANCE DISTRIBUTED OBJECTS OVER SYSTEM AREA NETWORKS

Just as the advent of high-speed networks shifts the performance bottleneck to the protocol stacks, the availability of commercial user-level networking shifts the bottleneck to the distributed-object infrastructures. In [10], I describe an approach to building high-performance, commercial distributed object systems over system area networks (SANs) with user-level networking. The specific platforms that I use in the study are the Virtual Interface Architecture (VIA) and Microsoft's Distributed Component Object Model (DCOM). The target application environments are physically secure server clusters consisting of homogeneous machines connected by high-speed system area network. The goal is to implement a set of software modules that can be integrated into existing DCOM infrastructure. These modules will be loaded for inter-server communications within a cluster, while client machines outside the cluster still contact the server machines through traditional protocol stacks running over traditional networks.

I have conducted a detailed functional and performance analysis of DCOM and apply optimizations at several layers to take full advantage of modern high-speed networks, while preserving the full set of DCOM features including security and different threading models. After extensive runtime and transport optimization, the system achieves a round-trip latency of 72 microseconds for null DCOM calls, more than 5 times faster than current implementation on the same network. By eliminating buffer copying at the marshaling layer, the system achieves an application bandwidth of 86.1 megabytes per second, an improvement by more than a factor of 7 over the current implementation.

8. RESEARCH PLANS

In the future, I would like to expand in a number of directions. In wireless ad hoc networks, there are many avenues for further work on topology control, gossip-based routing, and other issues. I will next discuss some of the research problems.

1. *MAC protocols that interact well with topology control*: The standard IEEE 802.11 does not work correctly with topology control algorithms. The reason is that the mechanism in IEEE 802.11 that deals with the hidden-terminal problem (RTS/CTS/ACK) assumes that every node has the same transmission range. This is not the case for networks using topology control, as each node determines its own transmission range. I would like to develop MAC protocols that work efficiently with topology control.

2. *Topology Control Aware Routing Protocols*: Current routing protocols are independent of topology control. However, better routing decisions can be made if routing is made aware of the topology control algorithm. I would like to improve routing protocols to take advantage of topology control.

3. *Adaptive Gossip-based Routing Protocol*: Our gossip-based work assumes a fixed gossip probability. I am interested in designing mechanism to make the gossip probability adapt to the network conditions.

4. *Security and incentive issues*: As ad hoc networks gain momentum for deployment, security and incentive issues have become more and more pressing. I would like to explore effective mechanisms in these two areas.

In wireless networks with fixed infrastructure, I would like to continue my work on the hybrid 3G-1xEVDO and 802.11 networks. Some problems I would like to explore in the future are:

1. *Scheduling in hybrid wireless networks*: Previous scheduling mechanisms determine which active user receives data from the base station using the channel quality of the actual destination. With 802.11 ad hoc forwarding, the downlink scheduling of HDR base station needs to consider both the proxy channel quality and the requirements of the actual destination. I would like to investigate efficient scheduling mechanisms that work well in this context.

2. *Security and incentive engineering in hybrid wireless networks*: In hybrid wireless networks, in order to foster a node to forward other node's packet, incentive mechanism must be in place. I would like to investigate monetary mechanisms as well as non-monetary schemes. For monetary schemes, I am interested in applying auction theory to the problem. For non-monetary schemes, I am investigating schemes that give user scheduling priorities based on the amount of packets it forwarded for other nodes. Security mechanisms must also be in place for correct accounting.

In the long term, I would like to investigate how the 4th generation wireless networks should be built. I would look at the physical layer technologies (such as OFDM, BLAST, Smart Antenna and Software Radio) and their implication on higher-layer protocols. In addition, I am interested in exploring the right architecture that accommodates such physical layer technologies.

The Internet has been evolving from a set of wires and switches that carry packets to a sophisticated infrastructure that delivers a set of complex value-added services to customers. I am currently developing scalable solutions to enable such services. Some of the problems that I am investigating are:

1. *Fundamental tradeoff in provisioning multicast in virtual private networks*: To support multicast routing in virtual private networks, traditional multicasting protocols must keep information in the core routers for each customer and each multicast group. This approach does not scale. There is a tradeoff between how much state we can put in the core router and how efficiently we can reduce redundant packets. I am investigating algorithms that make a good tradeoff between network state and bandwidth usage.

2. *Scalable virtual-router based VPN Provisioning*: In the context of virtual router based VPN, the virtual router is a scarce resource that must be used effectively. The problem is formulated as a variant of the facility location problem. I am interested in developing approximation algorithms for it.

3. *Scalable mechanisms for IP traceback*: IP traceback mechanisms enable the network to trace the denial-of-service attack to its origins. Current solutions can not scale to large-scale distributed attacks. I am investigating scalable mechanisms for IP traceback.

In the long term, I am very interested in solving other scalability problems in the Internet.

9. REFERENCES

[1] L. Li, M. Thottan, B. Yao, and S. Paul. Distributed network monitoring with bounded link utilization in ip networks. In *Proceedings of the 22nd IEEE INFOCOM*, 2003.

[2] L. Li, M. Buddhikot, C. Chekuri, and K. Guo. Routing bandwidth guaranteed paths with local restoration in label switched networks. In *Proceedings of the 10th IEEE International Conference on Network Protocols (ICNP)*, pages 110–120, 2002.

[3] L. Li and J. Y. Halpern. Minimum energy mobile wireless networks revisited. In *Proceedings of the 37th IEEE International Conference on Communications (ICC)*, pages 278–283, June 2001.

[4] R. Wattenhofer, L. Li, P. Bahl, and Y. M. Wang. Distributed topology control for power efficient operation in multihop wireless ad hoc networks. In *Proceedings of the 20th IEEE INFOCOM*, pages 1388–1397, April 2001.

[5] L. Li, J. Y. Halpern, P. Bahl, Y. M. Wang, and R. Wattenhofer. Analysis of distributed topology control algorithms for wireless multi-hop networks. In *Proceedings of the 20th ACM Symposium on Principle of Distributed Computing (PODC)*, pages 264–273, August 2001.

[6] L. Li and P. Sinha. Throughput and energy efficiency in topology-controlled multihop wireless networks. submitted for publication, 2000.

[7] Z. Haas, J. Y. Halpern, and L. Li. Gossiping-based ad hoc routing. In *Proceedings of the 21th IEEE INFOCOM*, pages 1707–1716, June 2002.

[8] R. Ramjee, L. Li, T. LaPorta, and S. Kasera. IP paging service for mobile hosts. In *Proceedings of the 7th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 332–344, July 2001.

[9] Z. Haas, J. Y. Halpern, L. Li, and S. B. Wicker. A decision theoretic-approach to resource allocation in wireless multimedia networks. In *Proceedings of the 4th international workshop on Discrete algorithms and methods for mobile computing and communications (dialM)*, pages 86–95, August 2000.

[10] A. Forin, G. Hunt, L. Li, and Y. M. Wang. High-performance distributed objects over system area networks. In *Proceedings of the 3rd USENIX Windows NT Symposium*, pages 21–30, July 1999.