Research Statement

Seungjoon Lee

Current Work and Interests

My research interests include protocols and algorithms for networks and distributed systems. In my recent work, I have focused on wireless, mobile networks and peer-to-peer systems, and dealt with various issues in different networking environments.

Theoretical analysis and experimental validation are essential to successful protocol and system design. After identifying the main issues in a research problem, I look for theoretical insights through formal analysis. Then, I perform detailed experiments using simulations or implementations to validate my proposed solution. In my work, I provide general solutions that work seamlessly with existing network components on various platforms, so that when needed, they can be readily deployed in current systems.

Wireless Networks

My thesis work concerns the area of wireless networks. Most current wireless networks are based on infrastructure. For example, cellular networks use base stations, and wireless local area networks (WLANs) use access points. However, I believe that in the future, it will be prevalent that different types of devices will communicate with each other without using infrastructure. While multihop wireless networks facilitate a fast and flexible way of communication in this scenario, they pose a number of interesting challenges due to node mobility and their decentralized nature.

For my thesis, I have worked on a general framework to deal with heterogeneity in wireless networks. Many early schemes in wireless networks assume homogeneous node capability or link quality. In practice, however, nodes in the same network can have different capabilities (e.g., differing battery capacity, processing power). Link qualities such as the average packet error rate can also vary widely among wireless links. As a result, these early schemes often suffer from significant performance degradation in practical wireless networks. I have designed and evaluated a set of protocols and systems that can cope with such heterogeneity. For example, I have proposed efficient networking schemes for routing, virtual routing backbone construction in multihop wireless networks, and multihop WLAN extension. This work also includes a logical abstraction framework that simplifies and unifies the access to link-level attributes in protocol implementations. I have performed various experiments using simulations and real wireless devices, and the results demonstrate that the proposed solutions achieve significant performance improvements over previous best schemes. My work in wireless networks has been published in **ACM MobiHoc 2005**, **IEEE Infocom 2004**, **ACM MobiHoc 2000**, **Elsevier Computer Networks Journal**. Some of my work is currently under review. Below I describe one specific scheme to handle heterogeneity in node capability and two protocols to handle heterogeneity in link quality.

Construction of a Virtual Backbone In most multihop wireless networks, there is no routing infrastructure, and wireless nodes act as routers. In such networks, however, sufficient redundancy often exists, such that only a subset of nodes can participate in routing and provide the service to the other nodes. This set of nodes is called a *virtual backbone*, and finding a “good” backbone has been a subject of intensive research. By using the service provided by the backbone, non-backbone nodes can save their resources (e.g., battery power, processing). In contrast, backbone nodes spend more resources, and thus it is beneficial to include only high-capability nodes in the backbone.

In this work, I have designed a distributed scheme (called TRUNC-K) for the virtual backbone construction. The TRUNC-K scheme has a number of desirable properties. From a *theoretical perspective*, TRUNC-K can construct a connected backbone that includes only high-capability nodes. We also can show that under reasonable assumptions, the resulting set is essentially the smallest possible in a graph-theoretic sense, which leads to maximum overall resource saving by non-backbone nodes. From a *systems perspective*, TRUNC-K enables a parameterized algorithm where we can trade off efficiency for resilience such that the resulting backbone can keep providing the service in a wide variety of network environments (e.g., dynamic networks with mobile nodes). For the evaluation, I have applied the TRUNC-K scheme to scenarios of saving energy in mobile nodes to increase overall network lifetime. Extensive simulation results show that TRUNC-K can increase the network lifetime by 20–220% over best prior schemes without adversely affecting delivery ratio or control overhead. The TRUNC-K backbone maintains connectivity better in the case of node mobility as well.

Efficient Geographic Routing Finding a good path from source to destination is a basic research topic in multihop wireless networks. Most earlier schemes are based on the network-wide flooding of control packets. Recently, geographic routing has emerged as an attractive paradigm, where nodes use the position information of their local neighbors, making the routing procedure more scalable in larger wireless networks. However, as in most flooding-based routing protocols, existing geographic routing protocols usually attempt to minimize the hop count. Without considering link quality, resulting paths are likely to include low-quality links (e.g., those with high error rate), which can disrupt successful end-to-end communication.

In this work, I have designed a simple link metric for geographic routing that considers both hop count and link quality. I have proven that in certain idealized settings, this simple metric provides the optimal trade-off between hop count and link cost, and finds minimum-cost paths. This general metric is applicable to various types of link cost such as link error rate, energy consumption, and link delay. I have evaluated this scheme using simulations in a number of different scenarios, and demonstrated significant performance improvement in various settings. For example, in high-noise environment with frequent packet losses, my proposed scheme can lead to 81% higher delivery ratio over existing schemes. I have also validated the feasibility of the proposed schemes through a simple measurement study, and in my continued work, I am implementing and evaluating this scheme on a real testbed.

Multihop WLAN Architecture Wireless Local Area Networks (WLANs) are the most popular method for wireless Internet access. In the current WLAN architecture, each user device uses a *direct* wireless link to reach one of the access points (APs) that are connected to wired infrastructure. As discussed before, some nodes may suffer from low-quality links, and previous literature reports that such nodes can significantly deteriorate the overall network performance, extending to nodes with high quality links.

I have proposed a new multihop WLAN architecture in which nodes may selectively use intermediate nodes to reach an AP. Through detailed measurement study in an existing WLAN, I have demonstrated that the use of multihop paths can improve the throughput performance as well as coverage area. I have designed a protocol that seamlessly enables the multihop extension in existing WLANs. Simulation results show that nodes using the multihop extension achieve significant performance improvements. More interestingly, the use of multihop extension also positively affects the performance of other clients that are *not* even aware of the multihop extension. This is because the nodes using a multihop path can transmit at a higher rate among multiple transmission rates available in the wireless system (e.g., IEEE 802.11), which in turn allows unaware nodes to transmit more frequently and achieve higher throughput.

Multicast in Multihop Wireless Networks For my Master's thesis, I designed a multicast protocol for mobile multihop wireless networks. The new protocol utilizes spatial locality in mobile nodes to reduce the control overheads for the maintenance of the data delivery structure. I performed extensive simulation experiments to show that the significant reduction in control overheads eventually results in better delivery performances (e.g., delivery ratio and latency) over existing schemes in dynamic networks.

Peer-to-peer Systems

Peer-to-peer systems are designed to utilize resources available at end nodes. In contrast to traditional client-and-server systems, a node in a peer-to-peer system acts the role of both content server and client. In addition to my thesis work on wireless networks, I have also worked extensively on peer-to-peer systems.

Resilient Application-layer Multicast and Media Streaming For efficient data delivery to multiple users, multicast services at the network layer were proposed more than 15 years ago, but are not widely available yet. Recently, many protocols have been proposed to provide multicast services at the application layer. In most of these protocols, end nodes form a tree for data delivery. However, end nodes in a multicast tree may result in more frequent node failure or ungraceful departure, and the loss of a non-leaf node in the tree can cause all nodes in the subtree to lose messages for a long period of time.

To enable effective recovery without incurring large latency or control overheads, my colleagues and I have co-designed PRM (Probabilistic Resilient Multicast), a general recovery framework for any application-layer multicast scheme. A key component of PRM is randomized forwarding, in which each node selects a few random peers in a tree and sends the current message with low probability (e.g., three random peers with 1% probability). Through detailed analysis, we have proven that PRM scales well and can achieve a near-perfect delivery ratio with negligible overheads in a large group. We have also performed detailed simulations on top of the NICE application-layer multicast protocol, and the results show that PRM achieves significant performance improvement: high delivery ratios (above 97%) with low latency bounds (600 ms) while incurring minimal overhead (below 5%) in a large dynamic multicast group up to 8000 nodes. For real-network experiments, we have incorporated PRM into the Darwin Streaming Server from Apple. Our experiment results on a testbed up to 64 nodes show that PRM achieves high delivery ratios with minimal overheads in real-life networking scenarios. This work has been published in **IEEE/ACM Transactions on Networking**, **ACM Nossdav 2004**, and **ACM Sigmetrics 2003**.

Identification of Cooperative Peer Groups In decentralized peer-to-peer systems, some nodes may misbehave and not follow the protocol to enhance their gain. For a peer-to-peer system to be stable and useful, there needs to be a mechanism in which nodes can find trustworthy peers with which to perform successful transactions. This work proposes a decentralized and efficient scheme to store user reputation information and infer trustworthiness in order to identify malicious nodes in the system. The simulation results show that the proposed scheme scales well (e.g., up to thousands of nodes) and quickly identifies cooperative peer groups while the requirements at each node for storage, processing, and bandwidth consumption are minimal. This work has been published in **Elsevier Computer Networks Journal** and **IEEE Infocom 2003**.

Future Research

I intend to continue designing efficient network systems. Specifically, I shall build on my current research work: wireless networks and peer-to-peer systems.

Wireless and mobile networks are playing an important role in today's Internet and will become more crucial in the future as different types of wireless devices become ubiquitous. In future wireless networks, more devices will communicate directly with each other without using infrastructure. This diversity of wireless devices brings new challenges, including how to combine various devices and wireless technologies into a single network and induce cooperation among nodes if necessary. In addition, future user requirements will push wireless network services beyond the current simple best-effort access mechanism to the Internet. For example, future mobile, wireless networks will be able to meet evolving user demands for multimedia content based on certain Quality-of-Service (QoS) guarantees, at any time and any place. Peer-to-peer systems will grow beyond simple file-sharing applications and constitute an important service platform in the wired Internet. To achieve this goal, we need to address many interesting research issues, such as efficient usage of network resource, fault tolerance, and cooperation schemes. Interestingly, peer-to-peer systems share many core aspects with infrastructure-less wireless networks; both are decentralized and subject to dynamic topology change. As a result, I believe that a good understanding of one area can often lead to solutions in the other area. Below, I describe a few specific research directions that I wish to pursue in the short-term and long-term future.

Short-Term Plans

Design of Wireless Integration Sublayer As discussed before, there are a number of link-level characteristics that upper-level protocols can use to improve the performance of wireless networks. As more upper-level networking protocols start to use such information independently, popular functionality (e.g., link error rate estimation) will likely be duplicated over many protocol components. A uniform set of interfaces in an *integration sublayer* that provide access mechanisms to low-level information will eliminate this inefficiency. It will also provide high-level abstraction and help upper-level protocols focus on their own goal, thus dramatically aiding the overall design process. Moreover, in the future more wireless nodes will be equipped with multiple wireless cards using different standards (e.g., IEEE 802.11 and IEEE 802.16), and the integration of multiple access technologies will become an important issue in achieving high performance. As a preliminary work, I have applied this concept to the implementation of geographic routing and backbone construction work. I intend to continue on this work, so that this sublayer approach will become a widely accepted practice in the wireless networking community.

Cooperation Schemes Based on Incentives Most operations in peer-to-peer systems and mobile wireless networks involve cooperation of participating nodes. For example, data dissemination in peer-to-peer systems or basic routing operations in wireless networks require nodes in the system to follow respective protocols. Although most existing schemes in both areas implicitly assume the cooperation of participants, self-interested nodes may not follow the protocols without sufficient incentives to do so. Recently, several schemes have been proposed to address some aspects of this issue. However, there needs to be more research to better understand various networking scenarios from this perspective. Ideally, the solutions should require minimal support from infrastructure or central authority. Currently, I am designing an incentive mechanism that induces selfish nodes to cooperate and construct a virtual backbone in multihop wireless networks. I am using game theory to insure that it is in the nodes' best interest to follow the protocol for backbone construction as well as to participate in routing and forwarding on behalf of other nodes. From this perspective, I intend to explore other scenarios in peer-to-peer systems and mobile wireless networks.

Multimedia Streaming in Multihop Wireless Networks With future advances in technology, mobile devices will become more powerful, and users will become interested in more diverse computing needs, such as accessing multimedia contents even when the users are mobile. I wish to work on providing multimedia streaming services in multihop mobile wireless networks. When coupled with mobility and the shared wireless medium, providing QoS to each user in this setting becomes even more challenging than in wired networks. Also, when multiple users want the same information, well-designed multicasting schemes in mobile multihop wireless networks will be crucial to the efficient use of networking resources. Besides practical importance, this research will help the networking research community to expand its understanding of how related components integrate to form a whole system.

Long-Term Plans

Pervasive Large-scale Wireless Networks In the future, small embedded computing devices with wireless communication capability will become more prevalent. These devices will enable the automated and remote control of living environments and the transparent use of the Internet. As these devices multiply, some network design issues, such as scalability and fault resilience, will become even more challenging in large-scale networks. Although some recent research efforts have tried to address issues in large-scale sensor networks, the communication model in most cases is rather simplistic (e.g., simple message exchanges only to/from a gateway node). Future networks will need to support various user requirements and more general and dynamic communication patterns. I intend to develop pervasive networking systems that provide consistent and reliable network experiences to users at any time and any place. To achieve this goal, I believe that many challenging questions need to be addressed. For example:

- What is the necessary infrastructure? How do we place it in the deployment phase and locate it when needed, so that users can transparently enjoy their computing environments without interruption?
- What is the best way for the data and control message exchanges to occur between nodes in general traffic patterns? To what extent do user and system behaviors need to be disseminated?
- What is a good model that can jointly describe user behavior (e.g., mobility, communication patterns) and operations of embedded devices (e.g., duty cycles, data exchange)?

I believe that these research issues are essential to the future evolution of networking systems. Through my research, I hope to stimulate intellectual activities in the research community as well as make a practical impact on the future of networking systems.