Research Statement

Yong Liu
Department of Electrical & Computer Engineering
Polytechnic Institute of NYU
yongliu@poly.edu

My primary area of research is networking. The rapid growth and evolution of the Internet and other forms of networking, such as Peer-to-Peer (P2P), wireless and sensors networks, has continually created new challenges for networking researchers. My research goal is to analyze and improve existing and proposed protocols, and design new networks, protocols and networked applications. My research spans all network layers, from the application layer down to the physical layer. I have worked on problems in traditional IP networks, as well as newly emerging networks, such as P2P networks and online social networks. My research methodology is to combine analytical approaches, such as modelling, control, and optimization, with experimental evaluation, such as simulation and experimentation. Analytical models often provide fundamental understanding and subtle insights beyond intuition. Experimental evaluation is critical to validate analysis and answer *"what if"* questions for analytically untractable problems.

1 Research Accomplishments

In the following, I summarize my research accomplishments in three sub-areas.

Measurement, Analysis, and Design of P2P Video Streaming Systems

Key Contributions: measurement study of large-scale systems; development of analytical models and performance bounds for P2P streaming; design and prototyping of new P2P streaming solutions

Video is driving much of the innovations on the Internet. In recent years, P2P video streaming has successfully enabled large-scale video multicast on top of the "best-effort" Internet. Together with colleagues at Polytechnic Institute of NYU, I made significant research contributions on the measurement, analysis and design of P2P video systems.

We conducted the first thorough measurement study of large-scale P2P video systems. We developed passive and active measurement platforms and collected extensive traces of PPLive, one of the largest commercial P2P video streaming systems which offers hundreds of channels and attracts hundreds of thousands of users every day. Through trace analysis, we unveiled key factors leading to the successes of P2P video streaming. We also identified the performance problems of existing P2P streaming solutions. Our measurement study provide valuable inputs for research on P2P streaming and IPTV in general. Our papers are highly cited. One paper was awarded *the Best Paper in Multimedia Communications of IEEE Communications Society* in 2008.

On the theory side, we analytically studied the performance of P2P video streaming in dynamic and heterogeneous network environment. For single P2P streaming swarm, we built stochastic models to characterize the universal streaming probability under peer churn. We analytically evaluated the impact of video buffering and system size on the system-wide streaming quality. We derived the fundamental delay bounds of P2P live streaming, and showed that peer heterogeneity can be exploited to improve delay performance. The robustness of P2P streaming against variations in network bandwidth and delay are also analyzed. For P2P streaming systems offering multiple channels, we developed Jackson infinite server queuing network models. We analytically showed that existing P2P streaming designs are vulnerable to channel churn and heterogeneous channel popularity. We further proposed a new streaming design VUD, which radically decouples peer video upload from download. We showed that VUD can significantly improve the streaming performance of multi-channel systems. Our paper on the modeling of multi-channel P2P streaming systems is the sole winner of *the Best Paper Award in IEEE INFOCOM 2009*.

Insights obtained from our measurement and analysis allow us to design new P2P streaming solutions. We worked closely with industry on design and prototyping. In collaboration with Thomson Corporate Research, we designed a Hierarchically Clustered P2P Streaming System (HCPS) to achieve close to 100% peer uploading bandwidth utilization with dynamic peer arrivals and departures. We developed the HCPS prototype and conducted extensive experiments on the PlanetLab to investigate the scheduling sensitivity of P2P streaming. In collaboration with IBM Research, we developed a new incentive mechanism for P2P video-on-demand system. In collaboration with Verizon labs, we investigated the benefit of adopting P2P technology to transport and deliver voice and video services in the Verizon FiOS networks. In collaboration with Huawei, we studied the optimal caching strategy for multitree based P2P video streaming. Several novel ideas have been filed as patents.

Robust Network Design: the X-ities

Key Contributions: the interaction between the routing optimization of overlay networks and underlay networks; network traffic engineering under traffic variations and failures

From 2005 - 2010, I was supported by NSF grant CNS-0519998: "Network X-ities: – Foundations and Applications". The objective of this project is to design the future generation Internet to achieve good performance in the face of a complex, uncertain, error-prone, and ever-changing environment. The need for such "robust" network operation leads to a set of design considerations that we refer to as the network X-ities (since they all end in "ity"): non-fragility, manageability, diagnosability, optimizability, scalability, and evolvability.

While overlay networks improve the performance perceived by overlay users, a fundamentally important question is to understand how overlay networks might affect the operation of underlay networks. We systematically studied the interaction between the routing optimization of overlay networks and underlay networks within a game theoretic framework, and showed analytically that the selfish behavior of overlay networks can cause both huge cost increases and oscillations in underlay networks. More importantly, we have also identified cases where the interaction between the overlay and underlay networks is *inefficient*, i.e., both overlay and underlay users suffer performance degradation. The analysis has been confirmed by simulation studies. This work provides a starting point in the search for a rigorous and well-founded understanding of the interaction between overlay routing and traffic engineering. Various insights gained from this study can be used to guide the design of overlay/peer-peer networks.

Traffic Engineering (TE) determines routes for network users to minimize network cost, such as network-wide delay. Given the matrix of traffic demand between all node pairs, TE can calculate the optimal set of routes by solving a constrained optimization problem. However, in practice, the traffic demand continually changes, make it difficult to accurately estimate a traffic matrix. We studied the optimal routing with multiple traffic matrices. We extended existing methods based on a single demand matrix to find an optimal set of routes which minimize expected network cost when there are multiple traffic matrices. To address the inherent complexity of optimal routing with multiple traffic matrices, we proposed and evaluated a heuristic algorithm which gives a near optimal solution and requires much less computation.

During network failures and/or maintenance, the amount of available network resources, such as routers and links, decreases. Congestion will accumulate in some portions of the network. While re-routing can alleviate the congestion to some extent, it is often the case that user perceived delay performance degrades due to the reduced network capacity. To meet the Service Level Agreement (SLA) with its customers, a network service provider may choose to block some traffic at network ingress points. The decision to be made is two-fold: a) how much traffic can be admitted for different customers? b) how should the admitted traffic be routed to meet the SLA? These two sub-problems are coupled together. In a joint work with AT&T research labs, we formulated it as a problem of utility maximization through joint rate control and routing. Several variations of original problem have been investigated to provide practical solutions for existing routing schemes, such as OSPF and MPLS.

In wireless sensor networks, the sensing and routing are strongly coupled through the power constraint on individual nodes. We extended previous studies on the interaction between rate control and routing in wired-line network to investigate the trade-off between sensing and routing in wireless sensor networks with directional antennas. We developed a distributed gradient-based algorithm that iteratively adjusts the per-node amount of energy allocated between sensing and communication to reach the system-wide optimum. We prove that our algorithm converges to the maximum system utility. We quantitatively demonstrate the energy balance achieved by this algorithm in a network of small, energy-constrained X-band radars, connected via point-to-point 802.11 links with directional antennas.

Modelling and Simulation of Large Scale Computer Networks

Key Contributions: development of scalable network fluid models; fluid and hybrid network simulators for large, high speed networks.

Networks, and the Internet in particular, have experienced exponential growth, becoming one of the largest engineering systems ever constructed. Understanding the behavior of this large system is of critical importance. The challenge we face is moving from a packet-based modelling abstraction to a modelling abstraction that is computationally tractable in today's networks with gigabit-rate links. Based on the previous fluid dynamic models of Transmission Control Protocols (TCP) and Active Queuing Management (AQM) schemes, we developed *topology-aware* fluid flow models to accurately abstract out the interaction between network components, such as between TCP at the network edge and AQM in network core. The proposed models are scalable with link speed and network size. It is proven to accurately capture the *average* behavior of large scale networks.

Based on the fluid flow models, we have also built a scalable network simulator that is capable of simulating large networks consisting of tens of thousands of routers and millions of TCP flows *in real time*. We also successfully integrated the fluid models with several major packet-level network simulators, such as *ns-2*, *QualNet* and *pdns*, to conduct hybrid network simulations with multi-level resolutions. Our simulators have been adopted by researchers from different areas to conduct simulations involving large scale networks.

Other than the previous areas, I have also made unique contributions to a set of challenging research problems:

- Application-level Relays for High Bandwidth Data Transport over Wide Area Networks;
- Information-theoretic Approaches for Network Measurement;
- Control-Theoretic Analysis of AQM Performance under Unresponsive Traffic;
- Statistical Analysis of the Origin of Long-range Dependency;
- Perturbation Analysis for Fluid Queuing Systems;
- Sample Path Analysis of Prioritized Fluid Queues

2 Plan for Future Research

My past work has demonstrated my ability to conduct interdisciplinary research in large-scale networked systems. Looking into the future, I am excited by the possibilities of working on a broad range of challenging issues in various networked systems. The following is a partial list of research projects in near future.

- **Next-Generation P2P Streaming.** P2P streaming is expected to be more prevalent in the future Internet. Other than large-scale IPTV services, P2P streaming is a natural solution for distributing User Generated Content (UGC). P2P streaming is also an efficient solution for new networked applications developed around user-user interactions, such as multi-party video conferencing and Massively Multiplayer Online Games (MMOG), etc. I have been funded by the National Science Foundation (under contract CNS-0953682) to develop new theory and design for the next generation P2P streaming systems in the next five years.
- **Online Social Neworks.** The phenomenally popular online social networks, such as Facebook and Twitter, present new opportunities for networking research. Recently, we proposed a new OSN-based P2P incentive paradigm: Networked Asynchronous Bilateral Trading (NABT). NABT exploits trust between friends in OSN. It allows a peer to trade with her friends asynchronously. It also enables a peer to trade with a remote peer through intermediaries. Our preliminary results demonstrated that NABT has much higher trading efficiency than the “tit-for-tat” type of synchronous bilateral trading mechanism. We will further investigate NABT and the integration of OSN with P2P in a NSF funded project (CNS-1018032) in the next three years. In a separate project, we proposed a Bayesian-inference based recommendation system for OSNs.
- **Content Distribution in Wireless Networks.** While wireless networks have brought us the convenience of mobility and new applications, they are still hamstrung by bandwidth bottlenecks. The capacities of wireless networks are challenged by an increasing array of date-intensive network applications, such as video, music, web, and games, on the new generation multi-mode wireless devices (such as the iPhone). Unlike in wireline networks, content distribution in wireless networks is tightly coupled with operation at the lower layers. We will address the challenge of content distribution in multi-channel wireless networks by integrating application layer solutions, such as overlay multicast, content caching, and P2P sharing, with optimal cross-layer link and channel scheduling.

3 Summary

In summary, I am confident of my ability to carry out independent research. I also believe many research problems from different fields are intrinsically closely coupled and should be solved drawing on expertise from multiple disciplines. Thus, I look forward to collaborating with researchers with a wide variety of research interests. My combined experience, in electrical & computer engineering and computer science, allows me to continue to make unique research contributions in my future career.