Statement of Research Interest
Shah Asaduzzaman

Research Area and Approach

My primary research interest is in the area of networks and distributed computing systems. New generations of networked computing systems, aided by the global presence of the Internet infrastructure, are playing an increasingly important role in the society. Such increasing reach is making the demand for robust operation of such systems more stringent than ever. Failures or anomalies in these systems can now have enormous cost and serious consequences on our society. On the other hand, the dramatic increase in heterogeneity of the networked devices such as handhelds and embedded computers is making the building of integrated information systems ever more challenging.

Investigating the principles underlying the design and implementation of such robust and pervasive networked computing systems is the general goal of my research. How to build a reliable system using unreliable components and how local algorithms based on local knowledge can derive globally emergent system characteristics such as reliability, availability, efficient resource utilization and quality assurance, are the central questions that drive most of my research activities. Because real-world systems are complex, answering these questions requires careful design and implementation that attends the details of the real scenario. I also use theoretical analysis and simulation that allows exploration of design parameters in wider ranges and in isolation, and helps to understand the impact of each parameter on the observed behavior of the system.

Current and Past Research

Postdoctoral Research

My general research agenda as a postdoctoral fellow is to investigate ways to construction of efficient and robust peer-to-peer overlay networks for various applications including video streaming, distributed services and location based search. Part of my postdoctoral research is *funded by a postdoctoral fellowship from Fonds Quebecoise de la Recherche sur la Nature et les Technologies (FQRNT).*

In one of the projects, I have investigated the ways to seamlessly integrate the server and peer-to-peer resources in internet based video streaming platforms. As part of this project I have devised a clustered peer-to-peer overlay that efficiently carries the bulk content among server nodes in wide area deployment and at the same time creates a robust delivery network among locally available peer-to-peer resources. The proposed architecture and its evaluation have been presented in the papers **C2** and **J1**.

In another project, I have extended the idea of clustered overlays to design an adaptive and efficient overlay for a fully decentralized platform for location based search. This extends the previously designed distributed indexing schemes for geographic databases. The new design allows autonomous management of local information and self-organized growth and retraction of the overlay network, without presence of any special coordinating entities. The details are presented in the technical report **C8**. Currently, I am working on generalizing the search infrastructure for metric space data.

Based on my recent understanding on clustered overlays and its application in location based routing, I have investigated how off-the-shelf geographic coordinates of peers can be utilized to create efficient routing overlays. The results and arguments in favor of this concept are presented in the articles **C1** and **C9**.

In another direction, I am investigating the mechanisms to provide access-control in the cloud-based content sharing platforms, like Amazon S3, Google docs or YouTube, without trusting the cloud service provider as the gatekeeper. Recent works suggest that use of cryptography may provide an elegant way of access control in such untrusted environments. The research goal here is to devise a complete set of protocols that serve as an efficient mechanism for enforcing flexible access control policies in a decentralized manner in absence of a trusted centralized entity.

PhD Thesis

In my PhD thesis, I introduced a new bi-modal architecture for a geographically distributed and cost-effective service hosting platform. Hosting platforms based on dedicated resources, although able to provide controlled performance, lack scalability in case of highly variable user demands. Peer-to-peer platforms on the other hand self scales with user demand by utilizing public resources, but unable to provide any performance guarantee for the applications. The proposed architecture utilizes a combination of statically provisioned dedicated resources and widely available opportunistic public resources to provide quality assured services. The core idea is that through dynamic management of a combination of these two classes of resources, one can gain from the scalability of the public resources and achieve assured quality services by masking their unreliable behavior with the controlled performance of the dedicated resources. The project was partially *funded by JW McConnell Doctoral Fellowship from McGill University.*

I explored the combination in two different platforms and applications. In the first case a combination of a dedicated cluster of computers and idle capacities of user computers have been exploited to build a platform to serve compute intensive applications with response time guarantees. The proposed resource management policies for this platform have been presented in conference articles **C6**, **C7**, **C10** and journal articles **J2**, **J3**. In the other case, a platform with a geographically distributed collection of compute servers inter-connected with a combination of private dedicated links and best-effort links over Internet have been utilized to serve distributed stream processing applications that requires simultaneous allocation of computing and communication resources. Evaluation of the resource allocation policies of this stream processing platform have been presented in conference papers **C2**, **C4** and in the journal paper **J5**. In both cases, I have observed that by designing appropriate resource management policies, the combination can be utilized to increase the overall resource utilization and throughput of the system as well as to increase the client satisfaction in terms of fulfillment of the service agreements.

Future Research Directions

I envision my future research to span across some interrelated sub-areas of networked and distributed computing systems. The unifying theme of the research will be construction of efficient and robust distributed computing platforms with self-organized autonomous entities.

In continuation of my current focus in peer-to-peer organization of distributed systems, my near term plan is to work on the research issues that include the following –

- Peer-to-peer networking platforms for service hosting, with emphasis on routing locality
- Real-time message routing in peer-to-peer networks for games and virtual environments
- Exploiting social acquaintances for content routing in peer-to-peer content sharing networks
- Privacy and access control in social networks

In another track, extending my PhD research in resource management in distributed systems, I have plan to work on –

- Energy efficient task scheduling in cluster/cloud computing platforms
- Discovery and composition tools for web services
- Distributed stream processing for sensor networks

Additionally, I have interest to work in the area that studies the structure and dynamics of large scale networks in general, ranging from transport network to social interaction network and biological signal network. The main motivation here is to reveal the underlying properties of these self-organizing networks and exploit them for improved performance, reliability and security of the evolving and self-organizing networked computing systems.