SIS-External Research Statement

Rajesh Krishna Balan
rajesh@smu.edu.sg
http://apollo.smu.edu.sg

1 Research Philosophy

My research philosophy revolves around the belief that successful systems research has to demonstrate two things: 1) a novel and interesting approach to an important research problem, and 2) usability. I strongly believe that for a system to be successful, it has to be usable by the target audience. On the other hand, a highly usable system that doesn't advance the state of the art in any way is not particularly interesting from a research perspective either.

My research style is strongly systems-oriented with a large emphasis on hands-on system implementation. A real implementation is crucial, in my opinion, for testing complicated distributed systems as the complexity inherent in these systems is incredibly hard to capture by simulation studies alone. A real implementation is also vital in validating the usability of the system. In addition, the software artifacts from such implementations are important contributions that drive future research and even product development in the area. I am also a strong proponent of collaborative research. Modern systems are so complicated that it is impossible for one person to have all the answers to every problem encountered. As such, I tend to work with about 2-4 other people, each with slightly different skills and backgrounds, to tackle a large systems problem. This multi-disciplinary approach is very apparent in all my research undertakings. I describe some of these undertakings along with avenues of future research in the next few sections.

My research revolves around the theme of mobility. Around this theme, I have built clusters of research. In particular, I have research thrusts in the areas of *Applications for Mobility*, *Understanding Mobility in Organizations*, and *Infrastructure for Mobility*.

2 Applications for Mobility

2.1 Digital Wallet

Motivation and Background

Consider the following scenario: *"Jill is at the supermarket checkout line. She fumbles through her wallet to find credit card X, rejecting many other cards in the process, to pay for the transaction. Later in the day, she falls victim to a pickpocket who steals her wallet. Jill is now in a state of*

panic; she has to remember which cards she had in her wallet and then manually cancel those cards."

The above scenario highlights problems with a physical wallet; namely that finding particular items is time consuming, and revocating a lost wallet is extremely hard. In addition, managing multiple monetary and identification implements is not easy. Monetary implements include cash, debit and credit cards, and stored value cards while identification includes national and/or state identification cards and driver's licenses.

A solution would be to replace the physical wallet with a digital wallet integrated into an existing mobile device like a cell phone. This digital wallet would allow the owner to carry multiple monetary and identification implements. These implements could be quickly searched by name, type, or other keywords. In addition, with the right software, these implements could be managed far more effectively. Finally, security would be enhanced as all data on the digital wallet would be encrypted and back up options would make recovering from loss easier. In this research, I investigate the feasibility of moving your entire physical wallet into your cell phone.

Status

With the help of various student teams, we have developed a prototype digital wallet solution called *mFerio* [1, 2] that performs NFC-based peer-to-peer mobile payments. In addition, we have also developed another prototype, called *pFerio* [2] that provides consumers with point-of-sale information about appropriate payment options.

2.2 User-Centric Mobile Authentication Mechanisms

In this project, the idea is to create user-centric authentication mechanisms for mobile phones. The problem statement is as follows: everyone owns a phone – in fact for many people, this is their main computation device. The phone frequently contains information that is sensitive and potentially embarrassing if released to the general public. However, very few people lock their phones. The reasons for this are numerous, but one common reason is that the standard phone locking techniques, involving numeric key codes, are tedious, slow, and hard to remember. Hence, most people just ignore the authentication process and leave their phones unlocked.

We are testing alternative methods of authentication that are more user-centric; i.e., they employ methods that are easy for humans to understand and use while still preserving good security properties. In our current prototype, we plan to use history coupled with data mining to ask users questions that are specific to them. For example, questions such as "who did you call at around 8 p.m. last night?", or "When you send a SMS, how do you access the "send SMS" feature?".

3 Understanding Mobility in Organizations

3.1 Globally Distributed Software Development

Spurred by the desire to leverage global resources and the technological advancements to accomplish it, software development projects are increasingly being distributed across multiple development centers across the globe. In particular, this globalization has resulted in the emergence of large distributed software development efforts. These are efforts where the client is located in one country but the bulk of the software development for the client is performed in another country.

This separation of development effort from client locality is due to many factors. However, one of the prime reasons is cost. It is usually cheaper to develop software in third world or developing countries (such as India or China) instead of rich labour expensive countries (such as the US) where the clients are located. It is well documented that distributed teams bring new challenges to project management.

However, much of the research in this area has focused on theoretical models of software development. In particular, there has been precious little empirical work in quantifying the true impact of distributed software development on metrics such as quality, productivity, and cost estimation. We aim to rectify this gap in this research.

Status

Our initial results have been very promising. We analyzed the data from 42 distributed projects from the same Indian company. From this analysis, we were able to quantify the impact of distributing software development on quality and productivity. We were also able to provide solutions for mitigating the loss in quality and productivity. This analysis was very well received by the software engineering community (published at FSE 2007 [3]).

We then did a preliminary analysis of the effects of software process choice on various metrics associated with software development. This initial work has been well received by the software engineering community (appearing at ICSE 2009 [4]) and we are proceeding with future analysis.

In the short term moving forward, we plan to collect data from other Indian companies and use that data to create models that will help these companies accurately predict the costs associated with distributed software development. The companies are particularly interested in this aspect of the work.

3.2 Enterprise System Migration

In this research area, we tackle the problem of migrating large enterprise systems. Examples of these systems are the ERP and CRM systems used by many companies and institutions. These systems tend to be large, expensive, and complicated to maintain. Many companies already have

large existing investments in their current IT infrastructure. This infrastructure tends to be big, complex, mission critical, and hard to understand fully.

Unfortunately, these systems still need to be upgraded or changed completely as and when necessary. When these moments arise, companies need to be able to accurately assess the costs involved in choosing various system migration options. These options could include a) leaving the system as it is, b) upgrading the system to the latest version of the existing software, or c) migrating the current setup to a completely new system.

Currently, there are very few ways for companies to make informed estimation decisions regarding costs when facing this situation. They usually have to call in external consultants who may or may not provide correct answers. Optimally, the company should be able to reason about these decisions internally.

This domain is particularly challenging as large enterprise systems tend to have a large amount of business logic embedded in them. The exact nature of this logic has a very large impact on the cost of the migration. For example, migrating a system that has very little embedded business logic is much easier and less costly than migrating a system that has been heavily customized by the organization. However, this business logic is hard to discover as different stakeholders in the company (accounting, finance, etc.) understand different portions of the system well. In addition, these enterprise systems tend to be mission critical. Hence, it is not possible to shut them down to understand what they do carefully or run tests that could have side effects.

4 Infrastructure for Mobility

4.1 Massively Multiplayer Mobile Game Infrastructures

Almost everyone in Singapore carries a cellphone and most of those phones are 3G enabled. A potential "killer" application for these mobile phones might be massively multiplayer games. These games, when deployed on desktop computer, have proved to be lucrative revenue streams. However, how do you support thousands and even millions of mobile players in the same game? In this project, we look at the infrastructure-related questions related to this problem.

Status

We have obtained MOE Tier-2 funding for this project. Currently, we are looking at techniques to conserve energy on mobile devices when playing games.

4.2 Analysis of Transportation Systems

Transportation networks are large complicated entities. Identifying and eliminating inefficiencies in these networks requires a truly multi-disciplinary approach – ISM, data mining, and optimiza-

tion approaches to identify the inefficiencies, and systems approaches to re-work the existing approaches to correct those inefficiencies.

These transportation problems are also very interesting research avenues as they require new, novel, and interesting approaches to solve some of the problems inherent in these networks. In this work, we focus particularly on taxi networks.

Status

Our initial work analyzed the movement patterns, over a single day, of a taxi fleet. We found that the taxi system is mostly efficient. However, inefficiencies do exist and can be corrected. The results of this analysis were presented as a poster at MobiSys 2008.

References

[1] R. K. Balan and N. Ramasubbu. The digital wallet: Opportunities and prototypes. *IEEE Computer*, 42(2):100–102, Apr. 2009.

[2] R. K. Balan, N. Ramasubbu, K. Prakobphol, N. Christin, and J. Hong. mFerio: The design and evaluation of a peer-to-peer mobile payment system. In *Proceedings of the 7th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, Krakow, Poland, June 2009.

[3] N. Ramasubbu and R. K. Balan. Globally distributed software development project performance: An empirical analysis. In *Proceedings of the 6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE)*, Dubrovnik, Croatia, Sept. 2007. [Awarded ACM SIGSOFT Distinguished Paper Award].

[4] N. Ramasubbu and R. K. Balan. The effect of process choice in high maturity environments: An empirical analysis. In *Proceedings of the 31st International Conference on Software Engineering (ICSE)*, Vancouver, Canada, May 2009.