Pei Zhang
Department of Electrical Engineering, Princeton University
Olden Street, Princeton, NJ 08544
(609) 356 - 2525
peizhang@princeton.edu

Statement of Research Interest

Overview:

Mobile embedded devices have become an essential aspect of everyday life, evidenced by their increasing importance as drivers of many new applications, ranging from application-specific mobile sensor networks to ubiquitous mobile phone networks. I have every confidence that these networks will generate even more exciting new application opportunities in the near future; but at the same time they are challenged by multi-issues, which must be overcome before they can command widespread usage. One of the most prominent hurdles facing mobile network designers is the system functionality when node density, due to both logistical and technical reasons, becomes sparse. These sparse mobile systems face vastly different challenges in contrast to their more common dense/well-connected counterparts. Throughout my time at Princeton University under the guidance of Professor Margaret Martonosi in the areas of embedded systems and mobile sensor networks, I have been convinced that this problem is best resolved with a holistic, i.e. combining both hardware and software, approach.

My doctoral thesis, "Dynamic Management of Sparse Mobile Systems with Intermittent Connectivity", tackles many of the above-mentioned challenges in sparse mobile networks, enabling them to function under a full spectrum of densities, from the very dense to the very sparse.

ZebraNet Deployment: Impact of Sparseness and Mobility on Hardware and Software

When I first joined Princeton in 2002, I was immediately drawn to the ZebraNet Project with Professor Margaret Martonosi. The goal of the project is to create a comprehensive system that has:

- Spatial coverage of 1,000s of sq kilometers.
- Lightweight, infrastructure-less mobile hardware customized for wildlife tracking.
- Custom software/protocol for maximizing data reception.
- Long-term autonomous deployment.
- 2 deployments in Central Kenya in 2 years (2004-2005).

I couldn't help but be enticed with the comprehensive and real-world application nature of this project. Almost instantly, I committed myself to the design of the ZebraNet hardware along with some of its firmware. After months of 20-hour days with 4 major revisions, the first ever mobile sensor system designed and deployed for a sparse, hostile environment came into reality. In recognition of this design's originality and innovativeness, I was awarded the 3rd International Low Power Design Contest at ISLPED 2003. In addition, because of the novelty of my design's self-tuning solar energy scavenging unit, I was the winner of the Global Photonics Energy Corporation's (GPEC) Edith and Martin B. Stein Solar Energy Innovation Award.

The ZebraNet System was first deployed on wild zebras at the Sweetwater Ranch in central Kenya during January 2004 [1][2], where the deployment used only 7 tracking nodes inside the roughly 100km by 100km region. It was the first time ever that detailed nighttime animal movement data were recorded, and a first glimpse offered to biologists of detailed zebra social and movement behavior. With the valuable lessons learned, I later redesigned the system, moving much energy management responsibilities to hardware. Combined with other improvements, the new version became more tolerant of various failures caused by software, hardware and the inhospitable environment. This version was deployed with even greater success during the summer of 2005 [3]. It was through these experiences that I gained a firsthand knowledge of the real-world tradeoffs between hardware, firmware and middleware energy management techniques in delay-tolerant networks (DTN). Furthermore, I was later able to incorporate many of these lessons into the design of our second-generation middleware for DTN [4].

The immense challenge of designing a lightweight, long-operating device, coupled with that of carrying out novel research in an entirely unknown and hostile environment, led me to recognize the importance of grounding theoretical research in real-life implementations to uncover latent issues. For example, I realized that the lack of system-wide collaborative policies in DTNs seriously limited their full potential. As a result, my subsequent research has intelligent algorithms to bridge the gap between sparse DTNs and dense networks.

Adaptive Targeted System Lifetime: Energy Usage Adaptation

Through the data collected from ZebraNet deployments, I noticed that energy consumption was extremely volatile across nodes. Any slight difference in node position, animal habits, or solar panel exposure translated into wide variation on energy consumption and scavenging. Yet, due to the sparseness of the network, nodes were blind to the systemic energy level, let alone adjusting capabilities accordingly. To solve this problem, I developed Adaptive Targeted System Lifetime (A-TSL) [5], which

- automatically adapts individual nodes' energy usage to global system expectations, and
- defines system-wide expectations only in accordance with design *expectations*.

The novelty of A-TSL lies in the fact that it enables each node to operate at its maximum capacity while maintaining energy consumption at the *same* system-wide goal, thereby reducing energy variations without any loss of functionality. When compared with nodes running tuned-energy-reduction policies, my approach improves sink data reception by more than 50%. What is more, it reduces system energy usage variation by more than 5.5X in many cases, all accomplished at only negligible overhead.

Delay-Tolerant Collaboration through Estimations and Predictions

While A-TSL works well for achieving static system goals, to broaden its scope for applications, there needs to be a method for *dynamic* information passage. While the very sparse nature of DTN tends to prohibit instant information sharing, storing stale information aggregates errors to an intolerable degree, even within a short period. Therefore, an effective method for collaborative localization, sensor calibration, routing, etc. would not only prove instrumental in obtaining accurate system power measurement, but also considerably improve performance in many other aspects that would otherwise be impossible for sparse DTNs. Through my research, I developed a novel concept of *Delay-Tolerant Collaboration*, which utilizes three steps to enable collaboration in DTNs [6].

- Prediction Phase: predicts changes in parameters as time passes, to maintain accuracy during periods of disconnection;
- Filter Phase: filters incoming information from encountered neighbors to prevent information from being repetitively merged;
- Merging Phase: updates prediction by incorporating information from neighbors.

One particularly interesting case of this idea is embodied in my research on Low-density Collaborative Ad-Hoc Localization Estimation (LOCALE) [7]. Each node under LOCALE not only collaborates with occasional neighbors, but also *actively predicts* its own position. To give even further credibility, location estimation is kept in the form of a multi-dimensional probability density function, providing information on both the location *and* the accuracy of such information. Overall, LOCALE yields more than 27X better accuracy in location estimation compared with existing GPS-less approaches, accomplishing, for the first time, a reliable GPS-less localization method for sparse DTNs. Moreover, by allowing nodes to refine location estimates collaboratively, LOCALE also reduces the need for location beacons by as much as 64X, while reducing power consumption of GPS per-node by as much as 150X.

I am extremely excited about the original and promising concepts proposed in my current research and am currently applying my experiences to the SARANA project, A Space Aware and Resource Aware Dynamic Network Architecture. It is a ubiquities system that enables simple deployment of spatially-aware and resource-aware applications on dynamic and heterogeneous devices. I am especially enthusiastic on continuing to expand my research coverage into new and unexplored areas.

Future work:

I believe that as wireless systems become more widely applied, the variety of their applications would necessarily proliferate, hence the importance for greater flexibility of such networks by making them not only diversity *tolerant* but also diversity *aware*. While previous work has focused on situation awareness in high-processing-power devices such as robotic networks, the few works on low-end sensor networks are mostly limited to low-diversity networks. Noticing the great potential in this largely unexplored area, I would like to apply my past experiences to the research on situation-aware in *low-capability*, *high-diversity* sensor systems. Similar to my previous work, I would approach the awareness problem in a vertical manner, by developing a comprehensive low-processing-power framework that incorporates both investigation into various awareness issues, and a hardware platform that would enable awareness and further new applications as a driver for my research.

Situation Awareness and Uncertainty:

In most mobile networks, it is highly probable for each node to encounter a vast variety of situations calling for different kinds of awareness, which can generally be classified into three categories.

- Self awareness (i.e. movement pattern)
- Global Situation awareness (i.e. system phases)
- Neighbor awareness (i.e. local information, security, etc.)

Self Aware: In mobile networks, unpredictable node movement patterns can vastly hamper system performance. Fortunately, while these patterns are sometimes uncontrollable in the real world, they generally can be learned, with a prominent example of such mobility-aware being routing algorithms. However, how multi-hop and destination-focused routes can be discovered and maintained in a sparse DTN is still an open question.

Global Situation Aware: Global situation awareness has always been a critical issue for both controlled and uncontrolled mobile networks. As devices move around, the network varies in density over time, therefore the system needs to be able to predict and adapt to diverse situations. Furthermore, the same node can be mobile or fixed at different phases of the system. How global situation awareness can best be exploited in both hardware and software remains to be explored.

Neighbor Aware*:* As mobile networks become more popular, attackers will naturally find them more visible targets, which would render the motives and trustworthiness of collaborators dubious, thereby arousing additional uncertainty. Whereas eBay-style rating systems might be used to maintain mutual trust, how to best accomplish this in sparse networks, where the problem is again compounded by the lack of direct verification methods, is still debatable.

While these different categories can appear unrelated at first sight, they all share the same fundamental characteristic as involvement of uncertain predictions. For instance, knowledge of movement pattern only provides an uncertain prediction of nodes' locations in future periods; the security issue is simply another way of stating the uncertainty around a collaborator's motives. How these predictions and uncertainties are best calculated and represented in low-power sensor networks is largely unexplored. Furthermore, while these areas are interesting research subjects in their own right, the interplay between different kinds of system awareness also raises important questions: how does a device decide and allocate its resources? How can nodes recover when a prediction turns out to be wrong? More importantly, they provide many opportunities for awareness-enabled services in mobile systems.

Hardware and Applications:
These situation aware services open up an entirely new class of low-processing-power sensor applications, with one possibility in sight being guidance of controlled node movement. This will enable mobile networks to self-deploy and reconfigure to reduce coverage holes, which is especially useful in hazardous situations and dynamic monitoring applications. Such networks have had limited success in the past due to the confined roaming capability of wheeled vehicles and the difficulties in coordinating multiple mobile devices. However, with recent advancements in battery technology, this has been made possible by small, inexpensive electric helicopters. This, coupled with the latest development of light-weight sensor nodes, not only provides a general platform for various controlled mobile networks, but also creates an extremely dynamic platform that necessitates a range of diversity aware algorithms.

My passion in research is to develop not just theoretically sound but also practical solutions that have *real* impact on *real* systems. As devices become ever more powerful, it is natural for them to be used in ever more diverse environments, highlighting the need for these applications to develop greater awareness of, as well as reacting to, their operating conditions. I believe that these areas point out some potentially fruitful, if not essential, steps toward future mobile systems.

References:

1. T. Liu, C. Sadler, P. Zhang, and M. Martonosi. "Implementing Software on Resource-Constrained Mobile Sensors: Experiences with Impala and ZebraNet", Mobisys 2004. The Second International Conference on Mobile Systems, Applications, and Services. June, 2004.
2. Pei Zhang, Christopher M. Sadler, Steve A. Lyon, and Margaret Martonosi. "Hardware Design Experiences in ZebraNet", SenSys 2004. The Second ACM Conference on Embedded Networked Sensor Systems. Nov, 2004.
3. Yong Wang and Pei Zhang and Ting Liu and Chris Sadler and Margaret Martonosi. "Movement Data Traces from Princeton ZebraNet Deployments". CRAWDAD Database. http://crawdad.cs.dartmouth.edu/. 2007.
4. Pei Zhang, Chris Sadler and Margaret Martonosi. "Middleware for Long-term Deployment of Delay-tolerant Sensor Networks", The first International Workshop on Middleware for Sensor Networks (MidSens'06). Nov, 2006.
5. Pei Zhang and Margaret Martonosi. "Energy Adaptation Techniques to Optimize Data Delivery in Store-and-Forward Sensor Networks", SenSys 2007. The Fourth ACM Conference on Embedded Networked Sensor Systems. Nov, 2006.
6. Pei Zhang and Margaret Martonosi. "Collaborative Parameter Estimation for Sparse Mobile Networks", under submission.
7. Pei Zhang and Margaret Martonosi. "LOCALE: Collaborative Localization Estimation for Sparse Mobile Sensor Networks", under submission.*