Research and Teaching Statement

Lewis D. Girod

December, 2005

Research

My research interests are focused on *embedded sensor and actuator networks*. These systems have much in common with previous work in networking and distributed systems, but their applications motivate different design choices. Whereas traditional distributed systems applications tend to be *virtual* and tend to emphasize network transparency, embedded sensor systems are tied to the *physical* world, both by their embedded nature and by their application.

Coupling the system and the application to the physical world opens up a rich new application space in which results from many fields may be applied, including distributed systems, signal processing, control systems, and robotics. In my initial exploration of this space, I have found that physical coupling has a broad impact on the way these results are applied and adapted: for example, a physically coupled system is concerned about spatial neighbors in addition to network neighbors. Through the development and study of new applications, I will characterize and abstract the new principles and primitives that enable robust embedded sensor systems to scale and flourish. In my vision these systems engender a future world in which active man–made systems form a ubiquitous part of the environment, tantamount to a new robotic ecology[1].

In my research I propose a specific approach and a guiding principle: to approach this goal from a *systems* perspective, synthesizing results from signal processing, robotics, networks and distributed systems, with the guiding principle never to stray far from experiments grounded in reality. Retaining a grounding in experimentation ensures a focus on solving real problems. A systems approach to this vision explicitly acknowledges that the complexity of our target system implies that environmental changes and hardware or software component failures will be the common case. I hope to realize this vision by applying robotics and signal processing results within a resilient system design informed by networking and distributed systems design principles. Ultimately I believe that while simple solutions make good demos, properties of resilience are pre–requisite to our vision of ubiquity, in order to survive exposure to a broad array of environmental conditions and states of system health.

In the culmination of my Ph.D. research, I developed an initial instantiation of this vision in *Acoustic ENSBox*[1], a self–calibrating distributed acoustic array system. In this system,

[1] *ENSBox* stands for Embedded Networked Sensing Box, a generic platform for ENS applications.

the sensor nodes autonomously coordinate to calibrate their relative positions and orientations by emitting calibration signals into the environment and detecting those signals at neighboring nodes. These nodes implement a distributed feedback system based on a distributed systems approach, adapting to broadly varying environmental conditions, changing system membership, and failures of nodes, components, communications links, and software.

Acoustic ENSBox has many of the properties of the systems I wish to study: each node implements considerable signal processing locally and shares only summary information with neighbors, and the system collectively computes a consistent system map, via a multi–hop ad–hoc wireless network. However, with Acoustic ENSBox as a platform I can continue to push farther toward this vision, developing applications that perform complex tasks grounded in the physical world. For example, an autonomous, rapidly deployable perimeter security application might combine coordinated acoustic sensors, cameras, PIR sensors, limited–motion actuators and fully autonomous robots, all integrated into a distributed system performing a specific task. Such a system has immediate commercial application, for example as an inexpensive way to provide 24–hour security to construction sites or event sites, as well as military and police applications.

I plan to investigate several key problems in this area. First, I plan to continue to seek out the common communications abstractions and models that best support these applications: to do for embedded sensing applications what sockets, TCP and the client–server model did for networking applications. My Ph.D. work explored the application of reliable multicast mechanisms to a specific application. After a few more applications exist, factoring out the right interface will be an interesting intellectual challenge.

Second, I plan to consider the impact of disconnected operation and duty cycling on these communications primitives and on the systems themselves. As a rule, duty cycling is difficult to integrate into systems that are general–purpose. Most of the successful low–energy systems in use today are specialized, with a very simple scheme for duty cycling. By considering several distinct applications, I intend to learn more about how to implement duty cycling while retaining sufficient generality for a broad range of applications.

Third, I plan to consider some problems of security, reliability, and assurance that arise in these types of system. Without some answer to these issues, it will be difficult for such systems to gain commercial acceptance. For example, users of these systems might want the system developer to offer quantified assurances about the system's performance that define success and specify a probability of failure. However, because of the impact of a complex and dynamic environment, it is not yet clear how those assurances might be assessed.

In summary, I have pursued this vision since beginning my Ph.D. in 1998, focusing my primary efforts on developing collaborative localization systems with minimal deployment requirements, while keeping in mind this larger vision. Three years in industry from 2000 to 2003 allowed me to better understand the commercial implications of this research, and have inspired me enhance the impact of my work by choosing interesting problems with clear applications. Along the way, I have invested a great deal of thought and effort in the underpinnings of embedded sensor systems, in order to build systems that could be tested in realistic conditions. Despite the effort involved in building a complete vertical application, the reward has been a deeper insight into the true requirements, and an invaluable understanding of what problems most need to be solved. I plan to continue this strategy in my ongoing research.

Teaching

One of the most rewarding aspects of an academic position is the opportunity to teach and interact with students. Whether or not they realize it, students have the freedom to explore and to think about problems in new ways. As teachers, we have the opportunity to guide students' discoveries, and learn a great deal in the process. I would be most interested in teaching classes in the area of systems such as networking, operating systems, and embedded systems, especially those classes with a strong lab component.

Unquestionably my favorite aspect of teaching is in working with students individually as an advisor. In my tenure at UCLA I have been very active in giving advice and help to more junior graduate students, offering technical suggestions and ideas as well as helping them to refine and explain their projects. In many cases, I have been able to encourage students to fit their projects together to avoid overlap and increase the utility and power of their work. As a faculty member, I expect working with students to be one of the most rewarding aspects, both because it represents an opportunity to push forward research that I may not have time to explore, and because of the fresh ideas that the students will bring.

A second aspect of teaching that I enjoy is the development of a well–designed course and associated materials. In my experience, I have learned a great deal about a subject in the process of organizing it for presentation. I also enjoy planning out homeworks, quizzes, and projects with care to eliminate busywork, errors and unnecessary confusion, allowing students to get right to the heart of the problems. This can be a time–consuming process and may require additional resources, but I believe that it vastly improves the student experience.

A third aspect of teaching that interests me is an idea that I would like to explore on a much longer–term basis. In my experience in school and in industry, I have found that computer science curricula do not always teach the skills required to be a good programmer or development engineer. In my case, I found myself learning these skills not from coursework but by watching others and following their example, and from a hodgepodge of online opinion. I believe that CS curricula as currently designed do a good job of presenting the theory, models, and abstractions that underlie the field, but fall short at teaching the practice of programming and development.

One possible solution would be a fifth–year practicum in which students work as an intern or apprentice, while taking coursework that is very focused on practical techniques. The development of the coursework itself would be a valuable contribution to the software and IT industries, because it would provide a focus for standardization of this practical knowledge. Today, this information is scattered among innumerable in–house and for–pay training seminars and certification programs, and often suffers from inconsistencies and religious disputes.

Such a plan is not achievable quickly, but I believe that it would have a beneficial effect, both in the effectiveness and reliability of industrial work, and in the overall satisfaction of computer science students.

References

[1] Gregory Pottie and Rodney Brooks. Towards a robotic ecology. DARPA ISAT Study, 1999.