

Cricket: A Location-Support System for Pervasive Mobile Computing

9807-04

Progress Report: July 1, 2001—December 31, 2001

Hari Balakrishnan and John Guttag

Project Overview

Cricket is an indoor location system that uses a combination of RF and ultrasound hardware. Cricket has a unique decentralized architecture using autonomous beacons transmitting location and spatial information to passive listeners (attached to static and mobile devices). The Cricket architecture provides a degree of scalability and autonomy (ease of deployment) well beyond other location systems, and facilitates the preservation of user privacy to a degree not seen in other systems. We believe that these features augur well for its widespread deployment, provided many technical challenges can be solved. Solving these problems and building a true production system are the goals of this project; we are fundamentally interested in real users and real developers using Cricket, rather than in implementing “only” a proof-of-concept prototype. We started working on Cricket (in its current form) in mid-1999 (very early stage designs) and came up with a high-level solution (untested) around September 1999 under the NTT-funded WIND project. Since then, we have been working on both the core Cricket system (design, implementation, evaluation, deployment) and on support for location-aware applications (INS for resource discovery, stream migration, connection migration). The overarching long-term goal of the research is to design context-aware systems and applications and understand how such systems should be engineered.

Progress Through December 2001

1. **Deployment.** We fabricated a few hundred units and have deployed them in various parts of our building. Cricket was used as a building-block technology in an undergraduate class on Project Oxygen. It has also been disseminated to about 15 other research groups outside MIT, who are using it for their own research in pervasive computing.
2. **Algorithms.** In collaboration with Profs. Teller and Demaine and students, we have been developing *self-configuration* algorithms to avoid manually configuring Cricket beacons with their locations. While we don't have an implementable algorithm yet, we have made significant progress on understanding why the problem is hard. We also have practical algorithms for important special cases.
3. **Cricket software compass.** We have made progress on designing a production-quality compass, which can enable a variety of new applications. One of our colleagues, Seth Teller, discussed this at the recent partnership meeting at NTT. This has turned out to be a very hard problem---doing ultrasound signal

processing (phase processing) in real-time isn't easy, but we now have methods to do this. We plan to engineer this over the next few months, according to our current plan.

4. **CricketNav.** We have designed and implemented CricketNav, a complete navigation system that builds on location and orientation capabilities of Cricket. This involves the development of good algorithms to estimate coordinate position in Cricket, and practical path planning methods that provide a good user interface to users; it's easy to swamp users with lots of useless information, and the challenge is in figuring out what's actually relevant at any stage.
5. **Design of INS/Twine (Scalable intentional resource discovery):** INS is a peer-to-peer system where the set of nodes in a community collaborate with each other to provide a system-wide service. The big challenge facing both INS and other existing discovery systems (e.g., Jini, IETF's Service Location Protocol, Universal Plug-and-Play, Salutation, and various other research proposals) is that they do not scale well, or if they claim to, they make unrealistic assumptions (such as, for example, that network proximity is equivalent to geographic proximity, which is clearly not true). The combination of INS and Cricket enables a variety of location-aware resource discovery applications. We have prototyped a version of INS that scales to a very large number of nodes and resources, using ideas from peer-to-peer location algorithms. Our algorithm builds on Chord, which was also developed under an NTT-funded project last year.

Research Plan for the Next Six Months

1. Auto-configuration in Cricket. We have preliminary algorithms for this.
2. Measurements/experiments with Cricket. We now have enough of a deployment that serious "non-lab" experiments can be done.
3. Continue development of Compass.
4. Monitoring and control applications. So far, our application efforts have been targeted towards human-centric "finding" applications. A large class of location-aware applications are in monitoring and control, e.g., large-scale asset tracking, theft prevention, sensor/actuator monitoring, etc. We plan to understand this domain better and investigate challenges therein. As a starting point, we have started working with the people in the LCS/AI reading room to design and automatic tracking system for all the books, journals, periodicals, etc. in the room.