CarNet: A Scalable Wireless Network Infrastructure
MIT2000-06

Robert Morris and Frans Kaashoek

The Carnet project is intended to explore the construction of large self-configuring networks that require little or no pre-deployed infrastructure. Part of the exploration involves the design, analysis, and implementation of network protocols; this stage of the project is well under way. This work has raised a number of key questions that cannot be answered by simulation, since the answers involve individual and aggregate user behavior. The only way to learn about these issues is to deploy a real network, encourage its use, and learn the resulting lessons. We know of no published experiences with large ad-hoc nets.

The questions we'll be able to answer with a deployed network go to the heart of the feasability of ad-hoc networking:

1. How much forwarding load will users impose on each other? This depends on the locality of user communication patterns; lack of realistic models for this has prevented any useful simulation of ad-hoc protocols to date.

2. How well do our power-saving protocols avoid draining batteries? Again, the answer (and choice of protocol) depend on user communication and movement patterns.

3. What is the best way to estimate the positions of units that don't have GPS? Such estimation is needed both for geographic routing and for position-aware applications. We've simulated a number of estimation algorithms, but can't seriously evaluate them without more experience of real spatial distribution of users.

4. Are our designs for routing around holes in the geographic distribution of users likely to be practical? Again, simulation-based evaluations are as much influenced by user movement models as by algorithms, so an empirical basis for such models is important.

5. How much privacy will users demand? Our system tracks user locations, which may be objectionable. We can provide various kinds of location anonymity, but how much to focus on this problem depends on user reactions.

6. Finally, building and maintaining a real system will expose any misconceptions we have about the system design, and help direct our efforts along truly important avenues.

Thus our plan for the 2nd year of the project is to deploy and analyze a large ad-hoc network running our protocols. We'll distribute a few hundred hand-held units running our networking software to students. We'll implement a basic set of useful applications for these units: voice and text chat, moving map navigation displays, and a simple geographic resource locator. In addition, the hand-helds will have traditional

applications such as an e-mail reader, a web browser, and a calendar manager. We may also be able to include our network software in hand-helds deployed by the Oxygen Project.

In summary, our proposal for the second year of funding is to deploy a large network with our protocols, watch people use it, and use the results to validate and correct the system design.

Progress:

Since NTT started funding CarNet in July of 2000 we've made progress on a number fronts. First, we've implemented most of a prototype based on Compaq iPaq hand-held computers running Linux and 802.11 wireless radios. Our ad-hoc networking software implements geographic routing and self-configuration. At the moment the main application is Internet connectivity by way of a few wired gateways. Second, we've worked out a technique that allows geographic forwarding even when some nodes do not know their own locations. Third, we've developed techniques for saving battery power in ad-hoc networks; existing techniques (e.g. those used by 802.11) don't work in multi-hop networks. Fourth, we've studied the amount of network capacity likely to be available in ad-hoc networks using analysis and simulation. We've submitted papers on the last three areas to conferences.

Links to our published work can be found at this URL:

http://www.pdos.lcs.mit.edu/grid/

Collaboration:

Robert Morris visited NTT this January to give a talk about CarNet.