# Report on DIMACS Workshop on Pervasive Networking (May 21, 2001)

**Workshop Organizers:**
Badri Nath, Rutgers University
Deborah Estrin, University of California at Los Angeles

**Report Assistant:**
Xiaodong Sun, Rutgers University

## 1 Workshop Focus

Progress in technology will make it possible to network and monitor every physical object that surrounds us. This large scale sensor network will become a reality in the near future due to the availability of cheap, low-power sensors with wireless access. This ubiquitous network will result in computing environments in places such as automobiles, homes, roads, infrastructures, buildings. These pervasive computing environments will contain millions of components that are distributed throughout physical space, allowing users to query the physical surroundings and get continuous information about the physical space at a level of detail that is not yet possible.

These future computing environments must be robust, deployable and maintainable. They must be self-configuring, self-monitoring and scalable. They must monitor performance, schedule resources, predict and recover and/or work around failures, and exhibit predictable behaviors; they must be long-lived and therefore low-power.

To realize the above vision, there are several research challenges in the area of low-cost networking, disposable networks, application level routing, information management and robust access protocols and design.

At this workshop, the talks addressed research challenges and problems in this new paradigm of a networked physical world, especially in the following areas:

- important unsolved problems (e.g., paging channels, network management)
- new directions
- low-maintenance networks (disposable networks)
- inter-disciplinary approaches
- promising models/modeling tools
- application drivers
- networking approaches

## 2 Talk Abstracts

The following were prepared by the speakers and mildly edited.

### Bluetooth Vs. 802.11b: State-Of-The-Art And Future Trends

*Pravin Bhagwat, ReefEdge Networks*

The promise of untethered computing in the workplace is becoming a reality. IEEE 802.11b, the 11Mbps wireless LAN standard, has finally arrived, and early market response has been positive. As the WLAN market takes off, Bluetooth, another emerging standard for short-range wireless networking, is also gathering force. Several vendors have demonstrated Bluetooth products, including cordless headsets, PCMCIA cards, and LAN access points. Both standards are competing for the same airwaves, but are they also chasing the same market? Will Bluetooth and 802.11b complement each other, or will one technology eventually displace the other?

### Distributed Query Planning In A Sensor Network

*Philippe Bonnet, Chang Choi, Cornell University Daniel Mosse, University of Pittsburgh*

Large collections of sensors with computation and multi-hop communications capabilities form the underlying computing infrastructure for a new generation of area monitoring applications. These applications rely on filtering, correlation and aggregation queries to access sensor data regardless of the physical organization of the underlying sensor network. Because each sensor disposes of a finite amount of energy provided by its batteries and because it has been shown that the energy cost of sending data is orders of magnitude higher than the cost of processing data locally, queries over a sensor network are best executed in a distributed manner, within the sensor network.

In this talk, we will briefly review the data model and the query language we have proposed for expressing queries over sensor networks. We will then show that classical distributed query processing techniques are challenged in this new environment and we will introduce an architecture for distributed query processing within a sensor network. We will detail the recent progresses we have made concerning two aspects of distributed query processing namely a sensor node index structure that can be used to route query fragments without flooding the network and hierarchical scheduling algorithms for planning distributed query execution given delay constraints.

### PARO: A Power-Aware Routing Optimization Scheme for Ad Hoc Networks

*Andrew T. Campbell, Columbia University, and Javier Gomez, IBM Research*

Due to the fact that mobile ad hoc nodes have a critical need to preserve battery power, MANET routing protocols need to consider power saving techniques during operations. In this talk we will discuss PARO, a Power-Aware Routing Optimization protocol that minimizes the transmission power necessary to forward packets between wireless devices. Using PARO, intermediate nodes can forward packets between source-destination pairs thus reducing the aggregate transmission power consumed by wireless devices. An important property of PARO is that it outperforms traditional broadcast-based routing protocols due to its power efficient point-to-point on-demand nature. The protocol is designed to operate as a stand-alone multihop routing protocol for local-area wireless networks (e.g., single-hop home networks, single-hop sensor networks, WLANS, etc.) and as a power-aware enhancement for routing in wide-area MANETs.

### Scalable And Adaptive Power Management For Multihop Ad Hoc Networks

*Benjie Chen, Kyle Jamieson, Hari Balakrishnan, Robert Morris, MIT*

We present Span, a distributed coordination technique for multi-hop ad hoc wireless networks that reduces energy consumption without significantly diminishing the capacity or connectivity of the network. Span builds on the observation that when a region of a shared-channel network has a sufficient density of nodes, only a small number of them need be on at any time to forward traffic for active connections. This provides a way to save energy and increase system lifetime by turning off "redundant" nodes, without substantially altering the overall transit capacity of the network.

Span is a distributed, randomized algorithm where nodes make local decisions on whether to sleep, or to join the forwarding backbone as a coordinator. Each node adapts to its local topology, basing its decision on an estimate of how many of its neighbors will benefit from it being awake, and the amount of energy available to it. We give a randomized algorithm where coordinators rotate with time, demonstrating how localized node decisions lead to a connected, capacity-preserving global topology. For example, for a practical range of node densities, our simulations show that the system lifetime with Span is more than a factor of two better than without, while maintaining the forwarding capacity of the original network. Span integrates nicely with 802.11– when run in conjunction with the 802.11 power saving mode, Span improves both communication latency and capacity without much degradation in system lifetime.

### Autoconfiguration, Registration and Mobility Management For Pervasive Computing

*Subir Das and Anthony Mcauley, Telcordia Technologies, Inc., Archan Misra, IBM T.J. Watson, Sajal K. Das, University of Texas at Arlington*

In the vision of pervasive computing, users will exchange information and control their environments from anywhere using various wireline/wireless networks and computing devices. We believe that current protocols, such as DHCP, PPP and Mobile IP, must be enhanced to support pervasive network access.

In particular, this presentation identifies three fundamental functions: autoconfiguration, registration, and mobility management, that need such enhancements. Realizing that the IP autoconfiguration capabilities must be extended to configure routers and large dynamic networks, we first describe our autoconfiguration solution based on the Dynamic Configuration and Distribution (DCDP). Secondly, we discuss why providing user-specific services over a common infrastructure needs a uniform registration protocol, independent of the mobility and configuration mechanisms. We present an initial version of the Basic User Registration Protocol (BURP), which provides secure client-network registration and interfaces to AAA protocols such as Diamter. Finally, we discuss the Dynamic Mobility Agent (DMA) architecture, which provides a hierarchical and scalable mobility management framework. The DMA approach allows individual users to customize their own mobility-related features, such as paging, fast handoffs and QoS support, over a common access infrastructure and to select multiple globalbinding protocols as appropriate.

### Prediction-Based Monitoring In Sensor Networks

*Samir Goel, Tomasz Imielinski, Rutgers University*

In this paper we discuss the problem of monitoring data sensed in large sensor networks. A sensor typically runs a better having a limited lifetime. In order to increase the lifetime of a sensor it is important that the mechanisms used in monitoring them be energy-efficient. In this paper, we propose a new paradigm called Prediction-based monitoring for energy-efficient monitoring. We show that the paradigm can be visualized as a watching of a "sensor movie" and that concepts from MPEG may be applied to it. We have implemented the proposed algorithms in a testbed of Rene motes (http://tinyos.millennium.berkeley.edu). Experimental results show that the proposed solutions help to significantly cut-down the number of data transmissions, considerably increasing sensor lifetimes, and thereby, the lifetime of the networks formed from these sensors.

### Challenges In Geographic Routing: Sparse Networks, Radio Obstacles, and Traffic Provisioning

*Brad Karp, ICSI, Berkeley, California*

Greedy Perimeter Stateless Routing (GPSR) has been shown to scale well for networks of numerous rapidly moving routers, by virtue of its small state storage requirements at routers, low routing protocol overhead, and the robustness with which it delivers applications' packets. The use of geographic information in packet forwarding decisions is central to these properties. We identify three challenges in the advancement of such geographic routing systems: sparse topologies, where greedy forwarding is possible more rarely than on dense ones; radio-attenuating obstacles, that introduce non-uniformity in radio ranges; and traffic engineering for these non-hierarchical routing systems. Using simulation, we describe the role of network density in the performance of geographic routing schemes, and show that GPSR performs well on sparse networks. We describe the difficulties presented by non-uniform radio ranges. Finally, we harness the correlation between geography and capacity in wireless networks to propose a "geographic" approach to traffic engineering that increases the capacity of a geographically routed network by "spatially" distributing flows.

### Compiler Support For Power And Energy Management

*Uli Kremer, Rutgers University*

Effective power and energy management is important to prolong batter life and to reduce heat dissipation. Developing compile-time techniques for application specific power and energy management is an exciting challenge. In this talk, I will discuss several possible optimization techniques, with an emphasis on remote task execution. Experiments to support the benefit analysis were performed on SKIFF, a StrongARM based single board system developed at Compaq's Cambridge Research Lan (CRL). An addition, we report performance numbers on Compaq's iPAQ handheld PC.

### Hierarchical Power-Aware Routing In Sensor Networks

*Qun Li, Javed Aslam, Daniela Rus, Department of Computer Science, Dartmouth College*

This paper discusses online power-aware routing in large sensor networks. We seek to optimize the lifetime of the network. We develop an approximation algorithm called max-min $zP_{\min}$ that has a good empirical competitive ratio. To ensure scalability, we introduce a hierarchical algorithm, which is called zone-based routing.

### Connection Establishment In The Bluetooth System

*Ivana Maric, WINLAB, Rutgers University*

Any communication between hosts in the Bluetooth system is preceded by a connection establishment procedure comprising two steps. The first step, referred to as Inquiry, enables hosts to learn about other hosts in their range. Inquiry can be omitted when not needed. The second step in which a host synchronizes and starts communication with the intended host is called Page. The complexity of the connection establishment procedure stems from the fact that hosts use different hopping sequences before a connection is established. The set-up delay is also affected by the requirement of low-power consumption and the possible connections that hosts may already have. The delay introduced by the procedure may impose constraints on applications such as ad hoc networking and LAN access.

In this paper, the connection establishment procedure in the Bluetooth system is modeled. The analytical results on mean page duration are derived for paging modes defined by the Bluetooth Specification. In a channel with packet errors, we derive analytical results for the average page duration as a function of the length of a page scan interval.

Analytical and simulation results show the range of values for the average page duration between two hosts.

### Communication Support For Location-Centric Collaborative Signal Processing In Sensor Networks

*P. Ramanathan, K.-C. Wang, K.K. Saluja, T. Clouqueur, University of Wisconsin-Madison*

Smart, low-cost devices, containing multiple sensors integrated with short-range wireless communication and significant embedded processing capabilities are expected to be a technological reality in the near future. These devices can be deployed in large numbers to form a wireless ad hoc network with significant compute capabilities to carry out mission-critical tasks such as monitoring, detection, and tracking of threats in the deployed area. In such networks, each device by itself may not be able to provide useful information without collaboration with other devices. At the same time, due to the large ad hoc nature of these networks, it is a formidable challenge for a programmer to develop efficient distributed algorithms and implementations without a simple, but flexible, programming model.

This paper describes an approach called location-centric computing for addressing this challenge. The proposed approach is based on the observation that distributed computations in sensor networks typically require tight collaboration among devices in a certain geographic area and not among an arbitrarily specified set of devices. The paper describes a set of communication primitives well-suited for such location-centric distributed computing. The paper also compares the number of messages required for an example collaborative signal processing application, namely target tracking, in the proposed approach with that required in other well-known approaches in literature. The comparison shows that the proposed approach is better in terms of number of messages required for this application.

### Cooperative Computing: A Computing Model For Large Networks Of Embedded Systems

*Phillip Stanley-Marbell, Cristian Borcea, Kiran Nagaraja, Liviu Iftode, Rutgers University*

To benefit from the aggregated computing resources deployed in large networks of embedded systems, new computing models and system architectures must be employed, which will necessarily be different from the traditional distributed computing models.

Presented is a model and architecture for distributed computing over large scale ad-hoc networks of embedded systems. In Cooperative Computing, applications execute on a set of nodes identified by their properties, and are responsible for the discovery of these nodes in the network. The system architecture we propose to implement this model is called Smart Messages. Smart messages are comprised of code and data and execute on each hop in their path through the network. Nodes support the execution of smart messages by providing a virtual machine and a name-based memory region called the Tag Space.

We believe the Cooperative Computing and Smart Messages can provide an adequate programmable infrastructure for pervasive networks of embedded systems such as sensor networks and computational fabrics. We are currently prototyping this architecture using uCsimm microcontrollers, uCLinux, Sun Microsystem's KVM virtual machine and Bluetooth wireless networking.

#### Automatic Configuration Management In Dynamic Short-Lived Networks

*Raj Rajagopalan, Telcordia Technologies*

Today's network management techniques are designed for long-term usage of networks and while this approach is acceptable for networks in general, it is quite inadequate in some common network scenarios where the network exists for a very small period of time. Such networks are quite common in the military in forward-deployed networks in combat or rescue situations as well as in the civilian sector, examples being disaster areas, trade shows, touring groups, and sporting events. We call networks with a defined lifetime "Rapid Response Networks" or "Disposable Networks" and argue that radically new management techniques are needed for success in mission-critical tasks in such arenas. We propose that research needs to be done on the question of what would we do for the specified span of time. This turns out to be fertile ground with many open questions, some that can be solved quickly and others that need concerted and collaborative effort. This is joint work with Badri Nath, Nick Maxemchuk, Muthu Muthukrishnan and Fred Roberts. We will also briefly review some of the lessons learned in creating new tools for configuration management in an ongoing project at Telcordia Technologies that focuses on automating security policy management.

## 3 Open Questions and New Directions

### 3.1 What Does the Future Wireless Network Look Like?

With the emergence of the wireless network standards like Bluetooth, IEEE 802.11b and the new wave of various products and applications on wireless networks, it is quite natural to ask what the wireless network will look like in the next 10 years. The talk by Pravin Bhagwat gave some predictions relevant to this question. While

IEEE 802.11b aims at high capacity and long range, Bluetooth enjoys the advantage of low cost and low power consumption due to the potential of building everything into a single chip. On the other hand, the single chip architecture puts limits on the bandwidth and communication range.

As pointed out by Bhagwat, even though there are new developments that blur the distinctions between these two standards, the best strategy for Bluetooth would still be low cost and low energy consumption.

The challenges for the future wireless network include:

- Since multiple standards will co-exist in the next few years, there is a big challenge of interoperability between the standards.
- Since the throughput will degenerate when devices using different standards are very close, collaborative channel sharing methods are needed.
- How to deal with the expectations that are hyped up for the next generation wireless network?
- How to deal with the conflicting interests within the SIG (special interest group)?

## 3.2 Configuration Management

### 3.2.1 Understanding the configuration

An important issue for improving the existing wireless networks is to understand how they work. The talk of Ivana Maric presented models for the connection establishment procedure in the Bluetooth system and analytical and simulation results about the range of values for the average page duration between two hosts.

### 3.2.2 Automatic configuration management

In traditional networks, the administrators are allowed to turn off the network temporarily to do manual configurations. In the wireless networks, this is almost impossible due to their dynamic nature. Therefore, its very important to study how to enhance the current protocols, such as DHCP, PPP and Mobile IP, to support pervasive network access. Talks by Subir Das and Raj Rajagopalan investigated possible approaches to this problem.

## 3.3 Power and Energy Management

Due the observation that the growth of the computing power of the chips is much faster than growth of the life of the batteries, power and energy management is a very important issue for wireless networks.

### 3.3.1 Power-aware routing

In designing the algorithm for routing in a wireless network, the lifetime of the network is one of the main concerns. Talks by Andrew Campbell, Kyle Jamison and Qun Li proposed various ways to organize the wireless networks so that power is saved and hence the lifetime of the network is optimized.

### 3.3.2 Support from compiler

Compilers sometimes can give better prediction about the applications than real-time operating systems because they can scan through the whole code before compiling. This advantage is significant especially for hand-held devices where normally there is only one application running. Uli Kremer investigated the power saving potential of compiling the application into two pieces, one of which is to be run on a more powerful remote machine.

### 3.3.3 Support from query planning

The talk by Phillipe Bonnet showed that the existing algorithms in data-mining that filter out the useful information in the intermediate stage can help to reduce the amount of the data transmitted in the wireless networks and thus save power consumption.

### 3.3.4 Use prediction: idea from MPEG

The idea of compressing the data in MPEG format is to use the previous frame of a picture as a prediction for the current frame and just transmit the difference between the two frames. The talk by Samir Goel showed how this idea is used to save power in the network of sensors. The central server broadcasts a prediction on the next messages the sensors will send based on previous messages received from the sensors. The sensor sends a message to the server only if the message to be sent is different from the prediction by the server. The algorithms are implemented in a testbed of Rene motes (http://tinyos.millennium.berkeley.edu) and experimental results show significant savings.

## 3.4 Various routing algorithms

Because wireless networks sometimes can be so unstable that it is very hard to have a centralized control of routing, various routing algorithms for wireless networks were proposed.

### 3.4.1 Self-routing

The wireless networks have the following characters:

- Limited resources
- Large scale
- Heterogeneous
- Volatile

- Unattended

A typical example is a network of sensors and a typical application can be either data collection or ad-hoc routing. The talk by Cristian Borcea proposed the use of so-called smart messages. A typical smart message will contain a signature for authentication, a table of estimated resources needed, data blocks and code blocks to be executed at the nodes. Each node will have an embedded system for the code carried by the messages to be executed but does not provide any routing information. Message routing is solved by sending a much smaller probing message to find the routing path. The open issues are

- Trade-off between flexibility and migration overhead.
- Partial successful execution, e.g., averaging the temperature over most of the sensors.
- How to make it more secure?
- Simulation of various applications and routing algorithms.
- Integration of energy into the model.

### 3.4.2 Routing with geographic information

A wireless network normally does not providing information like the routing table due its dynamic nature. How to make use of the available location information is an interesting problem. Talks by Brad Karp and Parmesh Ramanathan gave very interesting insight into this problem.

Karp identified three challenges for geographic routing systems:

- Sparse networks where greedy forwarding fails.
- Radio-attenuating obstacles which introduce non-uniformity in radio ranges.
- Traffic engineering for these non-hierarchical routing systems

Karp also talked about the routing problem in so-called roof-top networks. The greedy algorithm usually causes traffic congestion in the central area of the roof-top networks. How to use geographic information to get around the crowded central part of the network is a challenge.

Ramanathan observed that distributed computations in sensor networks typically require tight collaboration among devices in a certain geographic area and not among an arbitrarily specified set of devices. This observation leads to an approach called location-centric computing.