Scalable Coordination for Wireless Sensor Networks: Self-Configuring Localization Systems*

Nirupama Bulusu, Deborah Estrin, Lewis Girod
University of California at Los Angeles
email: {*bulusu,destrin,girod*}*@cs.ucla.edu*

John Heidemann
USC/ISI
email: *johnh@isi.edu*

Abstract

Pervasive networks of micro-sensors and actuators offer to revolutionize the ways in which we understand and construct complex physical systems. Sensor networks must be scalable, long-lived and robust systems, overcoming energy limitations and a lack of pre-installed infrastructure. We explore three themes in the design of self-configuring sensor networks: tuning *density* to trade operational quality against lifetime; using *multiple sensor modalities* to obtain robust measurements; and exploiting *fixed environmental characteristics*. We illustrate these themes through the problem of localization, which is a key building block for sensor systems that itself requires coordination.

Keywords - low-power wireless, sensor networks, localization, self-configuration, localized algorithms.

1 Introduction

Recent technological advances have fostered the emergence of small, low-power devices that integrate microsensing and actuation with on-board processing and wireless communications capabilities. When deployed in large numbers and embedded deeply within large-scale physical systems, these devices gain the ability to measure aspects of the physical environment in unprecedented detail. Through distributed coordination, pervasive networks of micro-sensors and actuators will revolutionize the ways in which we understand and construct complex physical systems.[EGHK99]

There are many potential applications of sensor networks: physiological monitoring; environmental monitoring (air, water, soil, chemistry); condition based maintenance; smart spaces; military surveillance; precision agriculture; transportation; factory instrumentation and inventory tracking. This paper will address requirements and design themes for these densely distributed, physically coupled and wireless sensor networks.

Due to the sheer numbers of nodes involved and the particular needs of applications (eg. emergency services), these systems must be *ad hoc deployable*. In extreme cases, nodes may be dropped from an aircraft in a remote terrain; however, even assuming individual placement, the scale of the system and variations in the environment require that they self-configure and adapt to their environment without user intervention. Because wiring is often impractical, nodes must be *untethered*. This requirement stems from many factors, including remoteness (wildlife monitoring), mobility, and the need for ad hoc deployment. Once deployed, these systems must operate despite being largely *unattended*, since nodes may be inaccessible, whether due to their tight physical coupling (large industrial plants, aircraft interiors) or inhospitable terrain (toxic or urban locations).

The above requirements impose substantial physical constraints at both the node and system levels. Nodes must be small for unobtrusive monitoring. Since they are untethered, their energy sources must be on-board, and is often relatively small. The system as a whole must tolerate ad hoc deployment and unattended operation without infrastructure support. Given such constraints, the network designers' goals shift towards extending system lifetime and robustness in the face of unpredictable dynamics, rather than focusing on optimizing channel throughput or minimizing node deployment.

Although in most systems centralized solutions are preferred for simplicity, several constraints of wireless sensor networks make centralization expensive and often infeasible. Node energy limitations place numerous constraints on communication [PK00]. In addition, radios used in sensor networks are often quite low bandwidth (10-20Kb/s). Finally, system dynamics (node movement or failure and changes in radio propagation) with large numbers of nodes make a global picture expensive to get and impossible to maintain.

Localization is an important building block for sensor networks and is itself a sensor network. We use it as our example to motivate the need for automatic *self-configuration* through adaptive localized algorithms. A *localized algorithm* is a distributed computation in which sensor nodes achieve a desired global objec-

*This research was supported by DARPA under grant DABT63-99-1-0011 as part of the SCADDS project, NSF grant ANI-9979457 as part of the SCOWR project, and was also made possible in part due to support from Cisco systems.

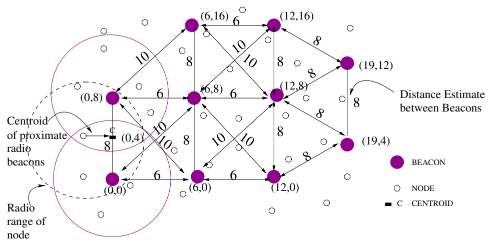

Figure 1: Localization - an example sensor network. Beacons self-organize into a coordinate system using pairwise distance estimates obtained by acoustic ranging. Other nodes may determine position to be the centroid of proximate radio beacons.

tive while constraining their communication to sensors within some neighborhood [EGHK99].

In this paper, we explore coordination in wireless sensor networks based on adaptive localized algorithms that exploit both the local processing available at each node as well as the redundancy available in densely distributed sensor networks. We introduce the design themes of density, multiple sensor modalities and adaptation to fixed environments, and show how they can be applied to build self-configuring localization systems.

2 Node Localization

Unlike the Internet, wireless sensor networks are organized around the naming of data, not nodes [EGHK99]. Nodes are neither unique nor reliable; applications express a need for a particular data element or type of data by naming it directly. By eliminating *indirection*, e.g. the mapping from a name to a node address to a route, a sensor network can eliminate the maintenance overhead associated with constructing and maintaining these mappings and directory services.

Because sensor data are intrinsically associated with the physical context of the phenomena being sensed, *spatial coordinates* are often a natural way to name data. Spatial coordinates are also employed by collaborative signal processing algorithms (e.g. beamforming) that combine data from multiple sensor nodes for such tasks as target tracking. Furthermore, geographic assistance in ad hoc routing promises significant reductions in energy consumption [KK00, XHE01].

The problem of estimating spatial coordinates is known as localization, and has generated much interest in recent years [BP00, BHE00, DPG01, Gir00, WJH97, NB00, PCB00, SHS01]. When sensor nodes are deployed in an unplanned topology, there is no *a priori* knowledge of location. Device constraints such as cost, form factor (including antenna size) and power consumption may preclude the use of GPS on all nodes. Moreover GPS does not work indoors, under water, or in the presence of overhead obstructions such as dense foliage. Thus, in many scenarios, sensor network nodes will need to determine their relative positions and self-organize into a spatial coordinate system without relying on remote infrastructures such as GPS.

Such a localization system is in itself an example of a wireless sensor network, as it involves a collection of networked nodes collaborating to achieve a higher level task: a coordinate system based on sensory measurements of the physical environment (such as signal strength, signal propagation characteristics, or packet delivery rates).

A sensor network may be organized as a tiered architecture of nodes, perhaps with a mix of small PC-class nodes (32-bit CPUs, 10^7 bytes RAM/Flash) and smaller nodes such as UCB Motes [HSW$^+$00] (8-bit CPUs, 10^3 bytes RAM, 10^5 bytes Flash). By mixing node sizes, very small-form-factor nodes can be densely deployed and physically co-located with targets, while larger but more capable nodes are still available when needed [CEE$^+$01]. Because individual node capabilities are quite varied, we require a federation of localization approaches (see Figure 1).

Multilateration algorithms: In one approach, sensor nodes measure a sufficient number of pair-wise distance estimates, and then use multilateration algorithms for position estimation. One promising ranging technology uses a combination of radio and acoustic signals, using the much faster radio signal to establish time reference and the time of flight of acoustic signals to estimate distance[Gir00]. Sensor nodes with these capabilities can independently form a relative coordinate system.

Proximity-based Localization: Such nodes can act as beacons for smaller devices such as the UCB motes that may not have the hardware capability for acoustic ranging. A beacon would periodically broadcast its position. By listening to broadcasts from a collection of nearby beacons and inferring proximity to those beacons with *low* message loss, each node could estimate its position to be the centroid of its proximate beacons [BHE00].

Iterative Multilateration: If the density of beacons is not sufficient in some areas of the sensor network, the proximity-based localization can be augmented by the more costly and perhaps less precise approach of iterative multilateration, in which beacon information is propagated through multiple hops to enable location estimation in areas of low beacon density[SHS01].

All the above approaches are sensitive to environmental vagaries. The problems that arise include (i) incorrect range measurements due to non line-of-sight conditions for acoustic ranging, (ii) error in RSSI-based ranging caused by variations in channel parameters across different environments ($1/r^n$ models), (iii) poor correlation between RSSI and distance owing to multipath interference and fading, and (iv) insufficient number of reachable beacons or interference amongst densely deployed beacons for proximity-based localization. We argue below that such problems call for the use of self-configuring localized algorithms.

Since almost all ranging techniques rely on signal propagation characteristics, they are susceptible to external factors such as interference, multipath effects and changes in temperature and humidity. The signal propagation characteristics of both radio and acoustic signals may change with variations in the surrounding environment. These physical effects are difficult to predict and can lead to incorrect range measurements which would greatly affect the quality of localization in our *multilateration* and *iterative multilateration* approaches discussed above.

Beacon placement and density can significantly affect the quality of localization in our *proximity-based localization* approach[BHE01a]. We cannot rely on a uniform placement of beacons as propagation characteristics of low power radio communications can significantly affect the visibility of beacons even when uniformly placed. An intuitive way to solve this problem is to deploy a large number of beacons. The problem with this approach is that we cannot have all deployed beacons turned on simultaneously because of the possibility of interference among several beacons vying for the communication channel as well as excessive energy use. In simple and small structured environments (indoors, factory automation plants etc.), we could perhaps model and carefully place exactly the right number of beacons. However, this approach is not useful when we deploy large systems in dynamic, unpredictable environments.

Two characteristics are clearly desirable in practical solutions to the problems described above. First, the overall system must dynamically and autonomously adapt (*self-configure*) and reconfigure to the particulars of its environmental setting. Second, due to scalability and energy-efficiency considerations, and because environmental characteristics can vary widely even within a single region of interest, self-configuration must be achieved by individual beacons using *localized algorithms*.

3 Design Themes

Because sensing and actuation define a physical scope to a node's influence, localized algorithms provide a natural design paradigm for physically distributed sensor networks. More importantly, localized algorithms are attractive because of their scalability and robustness. Localized algorithms scale well with network size since scaling is influenced by density rather than physical extent, therefore algorithm complexity grows with the degree and not total size of the graphs. Since they are self-configuring, they can also be self-*re*-configuring and thus can be robust to network partitions and node failures.

In this section, we elaborate upon a few design themes that arise in the application of adaptive localized algorithms for scalable coordination and self-configuration in wireless sensor networks.

3.1 Density

Density is an important parameter in physically distributed systems, both density of the solution space and of nodes. We formalize these notions below.

3.1.1 Solution Space Density

Localized algorithms are most effective when the problem solution space is dense, that is, a given problem has a large number of satisfying solutions. Since localized algorithms work with limited local information, we cannot use them to solve problems wherein we need to determine a global optimal solution. Because we do not have accurate, general, models of the physical world, measurement is needed and is well suited to localized algorithms. An example is our self-configuring beacon placement algorithms wherein beacons locally measure their neighbors and message loss to determine their roles (be active or passive).

3.1.2 Network Density

Localized algorithms are more effective when the network itself is dense. We can express the network density $\mu(R)$ in terms of number of nodes per nominal coverage area. Thus, if N nodes are scattered in a region of area A, and the nominal range of each node is R,

$$\mu(R) = \frac{N \cdot \pi \cdot R^2}{A}. \tag{1}$$

Note that the range R can be either the range of a particular sensor or the radio transmission range (idealized with circular propagation). In each case, the associated network density will be different.

Various phenomena *saturate* at a certain critical network density particular to them. Beyond this critical node density, additional nodes do not necessarily provide additional sensing, communication or coverage *fidelity* and are essentially interchangeable. For instance, Kleinrock and Silvester show that in a wireless network with a uniform distribution of nodes, when $\mu(R)$ is 6 nodes, the probability that a node is connected reaches 1 [KS78] regardless of actual node placement.

For problems where such a critical saturation density exists, the solution space density S can be related to network density. Let λ be the critical density required to accomplish a certain task. Is $\mu(R) > \lambda$, only a subset of λ nodes in any local neighborhood of size $\mu(R)$ needs to participate in the task.

The size of the solution space S, is the number of distinct subsets of nodes that could be active in any neighborhood.

$$S = \binom{\mu(R)}{\lambda} \tag{2}$$

In other words, for a given λ, $S \propto \mu(R)^{\lambda}$, and grows rapidly with $\mu(R)$.

3.1.3 Controlling Density by Changing Radio Power

How can we control network density? In many next generation sensor nodes, we can realistically expect the radio transmit power level of a node to be software-controllable [HSW$^+$00]. Assuming the receiver and transmitter gains remain the same, the nominal transmission range of a radio R is typically a function of its transmit power level P_t. For instance, according to the Friis Free Space radio propagation model [Rap96], the received power at distance d,

$$P_r(d) \quad \propto \quad \frac{P_t}{d^2} \tag{3}$$

If the threshold power for reception is P_{th}, then

$$P_r(R) \quad = \quad P_{th} \tag{4}$$

Thus

$$R \quad \propto \quad P_t^{\frac{1}{2}} \tag{5}$$

However, at very short ranges radio shadowing effects can attenuate specific frequencies, so the use of frequency hopping techniques is important. The correlation of range with transmit power in many cases may be non ideal, non radial, even non monotonic and concave. However, multiple power levels can still be used as a coarse adjustment of network density. For instance, if $R^2 \propto P_t$, then doubling the transmit power level can achieve twice the network density in Eqn. 1.

Additionally, multiple power levels can be used to get more information about the system. For instance, multihop communication breaks down at network partitions as there are no nodes *en route*. By transmitting at a higher power, information about such partitions can be communicated. Finally, multiple power levels can be used to construct a tiered architecture that takes advantage of heterogeneous capabilities and reach[CEE$^+$01].

Assuming beacons are distributed uniformly at random, proximity-based localization[BHE00] saturates at a certain beacon density[BHE01a] When there are redundant beacons, the system will be expending unnecessary energy and beacons may interfere with one another by congesting the communication channel. Our challenge is to find the right balance of beacons that provide basic beacon coverage and are conducive to good localization quality. Based on our principles of solution space and node density, we know localized algorithms may be applied here.

Our solution is termed STROBE, for Selectively TuRning Off BEacons[BHE01b]. Our basic approach is to extend system lifetime by exploiting the redundancy provided by dense sensor networks. Beacons in STROBE can be in either one of two active states (LISTEN-BEACON, BEACON-ONLY)[1] or in a passive state (SLEEP) and transition between these states depending on the number of their active neighbor beacons. While maintaining the desired threshold localization granularity across the terrain, STROBE both reduces the self-interference amongst several transmitting beacons and improves system lifetime by probabilistically turning off redundant beacons. By tuning it based on system and node energy consumption parameters, STROBE can be made energy efficient.

[1]Two active states are needed because listening incurs energy cost.

3.2 Multiple Sensor Modalities

Any individual mode of sensing can be blocked or confused by the environment. Leveraging multiple sensor modalities is one way to achieve robustness despite unpredictable environmental characteristics [GE01]. For every sensory system, there exists a set of environmental conditions that will confuse it, and a subset of those in which it fails to identify that it is confused. However, different sensory modalities are often *orthogonal* to each other, in the sense that their sets of failure conditions are largely disjoint. We contend that we can improve the quality of our sensor observations through coordination and communication, with significantly less effort relative to the effort required to incrementally improve the sensors on their own, simply by using these "orthogonal" modalities to identify each others' failure modes and reject bad data.

Based on these ideas, we are developing a prototype ad hoc deployable multimodal localization system [GE01] that is composed of many stand-alone acoustic ranging units and a few acoustic ranging units with cameras. In general, acoustic ranging performance suffers when the "line of sight" (LOS) path is obstructed. Acoustic range measurements in obstructed conditions often consistently detect longer reflected paths, leading to unbounded range error. Because they measure the long path consistently, it can be very difficult to identify these errors based exclusively on analysis of acoustic data.

However, suppose each camera's field of view contains several ranging units, which might be identified by a characteristic pattern strobed on an IR LED. Any ranging unit that the camera can see has a high probability of LOS to the camera, and thus in those cases, an accurate range can be determined with acoustics. Additionally, using angular displacement, a camera can estimate the range between any two ranging units in its field of view. By using the relatively coarse angular information from the camera, ranging units would be able to identify and ignore large errors resulting from obstructed conditions. Additionally, in a more complex scenario, two cameras might coordinate to formulate a 3D model of the terrain and thus determine the location of obstructing features, applying the techniques of Kanade[KON92].

3.3 Adapting to a Fixed Environment

Tolerance of random placement or high node mobility are not the only reasons to design sensor networks to be self-configuring. Even in cases where they are placed uniformly and do not move, nodes must independently self-organize to coordinate for collaborative sensing functions.

The environments in which these systems are expected to operate will be time-varying due to RF vagaries and other environmental dynamics. In addition to time-varying components, many characteristics of the environment will be a function of fixed elements, such as trees or hills on a terrain. Although time-varying effects can be analysed statistically, errors and distortions resulting from fixed elements must be compensated by detecting and adapting to these conditions. An approach aimed at characterizing the environment has the potential to improve sensing fidelity as well as energy efficiency. For example, in the multimodal localization system [GE01] previously described, nodes could retain long-term information about non line of sight pairs detected when obstructions change slowly.

3.3.1 Adapting to System Characteristics

It is difficult to design localized algorithms that both empirically adapt to a wide range of environments and converge to a desired global behavior over that entire range. Some information about the system can significantly help the convergence of localized algorithms. External system information may be provided in several ways. Some examples are: (i) Instead of treating all nodes uniformly, perform edge detection to distinguish boundary nodes. (ii) Use information about partitions or other nodes. (iii) Use long range radios or tiered architectures to balance energy efficiency with convergence.

We will illustrate the first example with the context of self-configuring beacon placement. Our simulations show that to improve the system lifetime with STROBE, it is important to distinguish boundary beacons (through edge detection) from other beacons. Consider the following example. All beacons form a linear chain of D hops. Each beacon has a nominal transmission range R and their regions of coverage overlap. Suppose only 1 in every 2 beacons needs to be active to achieve our threshold localization quality. Let ϕ be the total energy of each beacon and $\dot{\phi}_a$ be the rate of energy dissipation when a beacon is active. To keep things simple, we assume the rate of energy dissipation when a beacon is in the SLEEP state $\dot{\phi}_s$ is negligible.

Case 1: Treating nodes uniformly.

Boundary beacons estimate a lower neighborhood size and are always active. In that case, expected lifetime of a boundary beacon B,

$$L_b = \frac{\phi}{\dot{\phi}_a}. \tag{6}$$

After time L_b the boundary beacon will die. In this period, other beacons were active only half the time.

Expected lifetime of a beacon that is k hops from the boundary can be derived as,

$$L_k = 2 \cdot L_b \cdot (1 - (1/2)^{k+1}). \tag{7}$$

In this case, beacons die successively at times $L_b, 1.5L_b, \cdots, \approx 2L_b$. These cascading failures lead to a non-uniform behavior across the network.

Case 2: Edge detection.

If a boundary beacon can detect it is at the boundary, it can adjust its duty cycle to be active only half the time. Lifetime of boundary beacons

$$\hat{L_b} = 2 \cdot \frac{\phi}{\dot{\phi}_a} \tag{8}$$

$$= 2 \cdot L_b \tag{9}$$

Lifetime of a beacon (k hops from the boundary)

$$\hat{L_c} = 2 \cdot L_b. \tag{10}$$

Case 1 leads to a cascading failure that does not occur in Case 2. Thus we achieve uniform behavior across the system by distinguishing boundary nodes, and improve system lifetime.

3.3.2 Adapting to the Wireless Channel

Savvides et al. [SHS01] propose an approach by which nodes in a wireless network can improve the accuracy of their RSSI based location estimates (discussed in Section 2) by dynamically deriving (learning) the surrounding wireless channel properties. The algorithm starts with an initial guess of channel properties[2] and tries to obtain node position estimates through a sequence of successive multilateration. The initial set of position estimates can now be used to obtain an initial estimate of the channel properties by providing two crucial components: (i) A large set of inputs for the estimation of the channel parameters. (ii) A corresponding error variance that is used as a weight for each input in the channel model estimator.

Using these inputs, the channel model estimator can produce a new estimate of the channel properties which can be used in subsequent multilaterations. The process is repeated until the values of the channel model, and consequently position estimates converge to a specified tolerance.

This makes it a versatile solution that even without prior calibration can work in many different settings where the propagation channel properties are different. Furthermore, if the sensors are deployed over a wide area, the signal propagation characteristics may vary

[2]For instance, parameters such as the additive Gaussian channel noise in the log-normal shadowing model[Rap96].

widely even across the region of interest. Calculating the propagation characteristics locally yields better accuracy in the node location estimates.

4 Conclusions

Localization is a key building block for sensor network applications and is a sensor network in and of itself. We exemplified three design themes that will be important in wireless sensor networks generally - density, multiple sensor modalities for robust measurements and adapting to fixed environmental features.

References

[BHE00] Nirupama Bulusu, John Heidemann, and Deborah Estrin. Gps-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, October 2000.

[BHE01a] Nirupama Bulusu, John Heidemann, and Deborah Estrin. Adaptive beacon placement. In *Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS-21)*, pages 489–498, Phoenix, Arizona, USA, April 2001.

[BHE01b] Nirupama Bulusu, John Heidemann, and Deborah Estrin. Density adaptive algorithms for beacon placement in wireless sensor networks. Technical Report UCLA-CS-010013, University of California Los Angeles, May 2001.

[BP00] Paramvir Bahl and Venkata N. Padmanabhan. Radar: An in-building user location and tracking system. In *Proceedings of the IEEE INFOCOM 2000*, volume 2, pages 775–84, March 2000.

[CEE+01] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *ACM SIGCOMM Workshop on Data Communications in Latin America and the Carribean*, April " 3–5" 2001.

[DPG01] Lance Doherty, Kristofer S.J. Pister, and Laurent El Ghaoui. Convex position estimation in wireless sensor networks. In *Proceedings of IEEE INFOCOM 2001*, volume 3, pages 1655–1663, Anchorage, Alaska, April 22-26 2001.

[EGHK99] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99)*, pages 263–270, Seattle, WA, USA, August " 15–20" 1999. ACM.

[GE01] Lewis Girod and Deborah Estrin. Robust range estimation using acoustic and multimodal sensing. In *Submission to IEEE/RSJ International Conference on Intelligent Robots and Systems(IROS 2001)*, March 2001.

[Gir00] Lewis Girod. Development and characterization of an acoustic rangefinder. Technical Report USC-CS-TR-00-728, University of Southern California, April 2000.

[HSW+00] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors. In *Proceedings of the Ninth International Conference on Arhitectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.

[KK00] Brad Karp and H.T. Kung. Gpsr: Greedy perimeter stateless routing for wireless networks. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, pages 243–254, N.Y., August 2000. ACM.

[KON92] Takeo Kanade, Masatoshi Okutomi, and T. Nakahara. A multiple-baseline stereo method. In *Proceedings of the 1992 DARPA Image Understanding Workshop*, pages 409–426, San Diego, CA, USA, 1992.

[KS78] Leonard Kleinrock and John Silvester. Optimum transmission radii for packet radio networks or why six is a magic number. In *Proceedings of National Telecommunications Conference*, pages 4.3.1–4.3.5, 1978.

[NB00] Dragos Niculescu and B. R. Badrinath. Ad hoc positioning system (aps). Technical Report LCSR-TR-435, Laboratory of Computer Science Research, Rutgers University, 2000.

[PCB00] Nissanka Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location support system. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*, pages 32–43, Boston, MA, August 2000. ACM.

[PK00] Gregory J. Pottie and William J. Kaiser. Wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.

[Rap96] Theodore S. Rappaport. *Wireless Communications: Principles and Practice*. Prentice Hall PTR, 1996.

[SHS01] Andreas Savvides, Chih-Chieh Han, and Mani B. Srivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2001)*, Rome, Italy, July 2001. ACM. To appear.

[WJH97] Andy Ward, Alan Jones, and Andy Hopper. A new location technique for the active office. *IEEE Personal Communications Magazine*, 4(5):42–47, October 1997.

[XHE01] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking*, Rome, Italy, July 2001. ACM. To appear.