TOPICS IN WIRELESS SENSOR NETWORKING

# EXPLOITING EDGE CAPABILITY FOR WIRELESS SENSOR NETWORKING

**SHIWEN MAO, AUBURN UNIVERSITY**
**Y. THOMAS HOU, VIRGINIA POLYTECHNIC INSTITUTE AND STATE UNIVERSITY**
**MIN-YOU WU, SHANGHAI JIAO TONG UNIVERSITY**

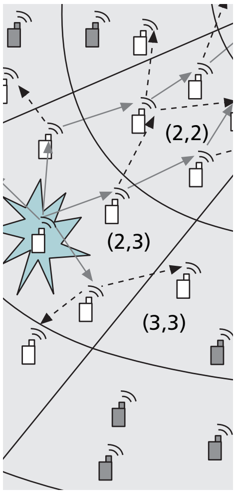

The authors propose to exploit capabilities at the network edge (i.e., an edge-based approach). They overview existing approaches to this end and present a novel edge-based routing protocol, called BeamStar, as a case study.

## ABSTRACT

As considerable progress is being made in wireless sensor networking, it is envisioned that sensor nodes will be on the cubic millimeter scale, posing stringent constraints on the processing, communication, and storage capabilities of sensor nodes. While it is important to continue pursuing novel algorithms and protocols to squeeze the most out of the existing design space (sensor nodes), it is equally important to explore a new design paradigm for future sensor networks to reduce the complexity burden on sensor nodes. We propose to exploit capabilities at the network edge (i.e., an edge-based approach). We overview existing approaches to this end and present a novel edge-based routing protocol, called BeamStar, as a case study. We show that exploiting edge capability provides a new dimension of freedom for wireless sensor networking, and is effective in relieving the processing, communication, and storage requirements of sensor nodes.

## INTRODUCTION

As considerable progress is being made in wireless sensor networking, it is envisioned that sensor nodes, with all the capabilities demonstrated today and new capabilities promised for tomorrow, will be on the scale of tiny dust, or cubic millimeter scale [1]. However, despite continued advances in micro-electromechanical systems (MEMS), low-power very large-scale integration (VLSI), and computing, it remains a formidable task to implement many of the capabilities on a sensor node on such a tiny scale. We observe that the main stream of sensor network research has focused their design space on sensor nodes. Under such an approach, the burden of achieving complex networking functions (e.g., topology control, medium access control [MAC], routing, localization, synchronization) rests solely on sensor nodes in the network core (i.e., a core-based paradigm). Such an approach relies heavily on future advances in silicon technology to reduce the size and cost of sensor nodes.

While it is important to continue pursuing novel algorithms and protocols to squeeze the most out of the existing design space (sensor nodes), we believe that it is equally important to explore a new design paradigm for future sensor networks to reduce the complexity burden on sensor nodes. We propose to explore capabilities at the network edge (i.e., an *edge-based approach*) for wireless sensor networking. This expanded design space has the potential of simplifying various algorithms and protocols on a sensor node, thus offering new possibilities to drive down its size and cost. The proposed approach represents a new paradigm for sensor network design.

In this article we present a brief overview of existing approaches to exploiting edge capability (or infrastructure support) in wireline and wireless networks. This interesting work demonstrates considerable benefits of the edge-based approach and motivates us to further explore this direction for efficient wireless sensor networking. As a case study, we present a novel edge-based routing protocol called *BeamStar* for wireless sensor networks. In BeamStar directional antenna and power control capabilities are employed at a base station. By having the base station scan the network using a power-controlled rotating directional antenna, each sensor node can derive its location information from received control messages. With such location information, sensor nodes can route reports to the base station via controlled broadcasting.

We show that such an edge-based approach enables extremely simple hardware and software design on sensor nodes, since many network control functions are effectively shifted to the base station. There is no need for sensor nodes to exchange control information for localization, synchronization, topology control, and routing. Furthermore, the proposed scheme is error-resilient, since each data packet may be forward-

---

*This research was supported in part by the U.S. National Science Foundation (NSF) under grants CNS-0347390, IIP-0738088, and ECCS-0802113.*

1536-1284/08/$25.00 © 2008 IEEE

Routing in wireless sensor networks is much more difficult than in traditional wireline/wireless networks, largely due to the stringent resource/hardware constraints, harsh deployment environments, unattended operations, as well as the unreliable nature of sensor nodes and wireless links.

ed along multiple paths constrained in a sector, yielding a better chance of successful delivery. The span of the mesh provides an elegant tradeoff between energy cost and reliability. Our performance studies show that BeamStar achieves a higher successful delivery ratio than a representative core-based approach at comparable energy cost, while greatly reducing the hardware and software requirements for a sensor node.

The remainder of this article is organized as follows. We first review related work. We then examine existing routing protocols and the expanded design space when edge capabilities are considered. BeamStar is presented as a case study. We then conclude this article.

## EXPLOITING INFRASTRUCTURE SUPPORT

The idea of taking advantage of intelligence at the network edge has been explored in prior networking research, such as in the Internet, mobile ad hoc networks, wireless sensor networks, and WiFi networks. We discuss these existing approaches in the following.

### AN EDGE-BASED APPROACH IN THE INTERNET

In the Internet, the design principle of implementing complex control algorithms at end hosts has long been adopted (e.g., in TCP congestion control) due to the high volume of traffic in the core. This approach has been adopted for quality of service (QoS) provisioning on both the data and controls plane [2, references therein]. The purpose is mainly for scalability in QoS provisioning and traffic management based on the concept of a stateless core.

Although the edge node within such an approach is tasked with more intensive packet processing and other control plane functions, there is hardly any direct interaction between an edge node and core nodes. In wireless sensor networks the network core (i.e., the sensor nodes) have very limited capabilities (in contrast to high-speed routers in the Internet core), while base stations are relatively more capable in terms of processing and communication power. It would be desirable to allow the edge system to directly convey control information to sensor nodes, with the potential of more efficiently accomplishing the various important tasks that are unique in such networks.

### INFRASTRUCTURE SUPPORT FOR IMPROVED NETWORK CAPACITY

In mobile ad hoc networks the benefits of using infrastructure support have also been observed [3, references therein]. In particular, these efforts show that infrastructure such as base stations can substantially increase network capacity beyond the limit found by Gupta and Kumar.

Under the so-called hybrid network architecture, base stations are deployed in a large-scale ad hoc network and connected with wired links. A mobile node therefore has two choices for routing its data:

- Through the infrastructure
- Via multihop relays

Liu, Liu, and Towsley [3] show that for an ad hoc network of $n$ nodes with $m$ base stations, if $m$ grows asymptotically faster than $\sqrt{n}$, the network transport capacity increases linearly with the number of base stations. Therefore, exploiting infrastructure support provides an effective solution to the fundamental capacity/scalability problem of large-scale ad hoc networks.

### BASE STATION SUPPORT FOR LOCALIZATION

From the perspective of practical protocol design, base station support has been exploited for localization [4–6]. In [4], Romer proposes using rotating lasers at a base station to scan a smart dust sensor network so that each node can estimate its location by measuring the timing information of the received laser beam. For wireless sensor networks Nasipuri and Li [5] present an interesting scheme where three or more directional antennas are used to scan the network in a synchronized manner. Each sensor node computes its location using angle of arrival and triangulation. Niculescu and Nath [6] present a VHF omnidirectional ranging (VOR) scheme for indoor positioning adopting multiple WiFi base stations with rotating directional antennas. Angles and ranges from the base stations are used to determine the location of a mobile node, while localization error can be as small as a couple of meters.

The objective of these efforts is twofold:

- Improving localization precision
- Reducing the hardware/software complexity of wireless nodes

It is worth noting that they rely on sensor nodes (or IEEE 802.11 users) to measure base station transmissions (e.g., time, angle, or range) and to estimate their own location based on the measurements.

These interesting schemes clearly demonstrate the benefits of exploiting edge capabilities. In the following we consider routing as a canonical application and discuss how to exploit edge capabilities from the routing perspective, using BeamStar [7] as a case study.

## ROUTING IN WIRELESS SENSOR NETWORKS

### CLASSIFICATION OF EXISTING APPROACHES

Routing has been the center of gravity for wireless sensor networking research. Routing in wireless sensor networks is much more difficult than in traditional wireline/wireless networks, largely due to the stringent resource/hardware constraints, harsh deployment environments, unattended operations, as well as the unreliable nature of sensor nodes and wireless links.

For sensor routing protocols, a common design objective is energy efficiency [8–11]. Sensor nodes have very limited battery power and are disposable. It is critical to extend the operation duration of sensor nodes (i.e., network lifetime) via energy-efficient routing [8]. Existing routing protocols proposed for wireless sensor networks can be roughly classified into three categories [8]

- Data-centric routing protocols (e.g., directed diffusion [DD] [9])
- Location-based routing protocols (e.g., Geographical Adaptive Fidelity [GAF] [10])

- Hierarchical routing protocols (e.g., Low-Energy Adaptive Clustering Hierarchy [LEACH] [11])

Under *data-centric routing*, a base station sends out queries to certain regions (via broadcast) and waits for data responses from the sensors in that region; or a sensor node broadcasts an advertisement for its newly acquired data and waits for a request from an interested data sink. The fundamental underlying technique under data-centric routing is *broadcast* or flooding, which is nevertheless undesirable for a sensor network operating under severe energy constraints. Under *location-based protocols*, geographical location information is used for routing. Such protocols share the common problem of obtaining location information in the first place. Generally, it is not cost effective and practical to employ GPS devices on sensor nodes. But the communication complexity associated with distributed localization algorithms may not be trivial. *Hierarchical protocols* are designed to alleviate the potential scalability problem, within which neighboring sensor nodes first self-organize into clusters, and a cluster head is then elected to relay information for all the nodes in that cluster. A considerable amount of sensor node activity is needed for cluster formation and maintenance.

### A UNIQUE CHALLENGE

Among many challenges for efficient routing protocol design, the stringent hardware and space constraints have received relatively less attention from the research community. This is a unique challenge for wireless sensor networks, especially when we consider drastically reducing the size of sensor nodes to very small scales.

For example, with the 0.13 μm complementary metal oxide semiconductor (CMOS) technology, a memory block of 512 × 128 bits alone will occupy an area of $0.22 \times 1.4$ mm$^2$, a considerable portion for smart dust sensor nodes on the cubic millimeter scale. For an operating sensor node, many other components are also indispensable, such as processor, sensing units, radio frequency (RF) circuits, mixed signal circuits, antenna, and battery. It would be a great challenge to integrate all of these components on a sensor node while keeping the size small. For a sensor node packaged with all of these components on a tiny scale, the processing, communication, and storage capabilities will be seriously limited. It is therefore crucial to reduce the complexity of routing protocols (as well as other algorithms for, e.g., synchronization and MAC), in order to achieve such reduction in size and cost.

### AN EXTENDED DESIGN SPACE

Toward this end, exploiting edge capabilities provides a new dimension of freedom for wireless sensor networking. For a sensor application that was previously solely implemented within the network core with localized algorithms, it is now possible to explore new algorithms that exploit capabilities at both the network edge and network core. This approach has the potential of yielding greatly simplified algorithms and protocols for many important functions, thereby relieving sensor nodes of processing, communication, and storage requirements.

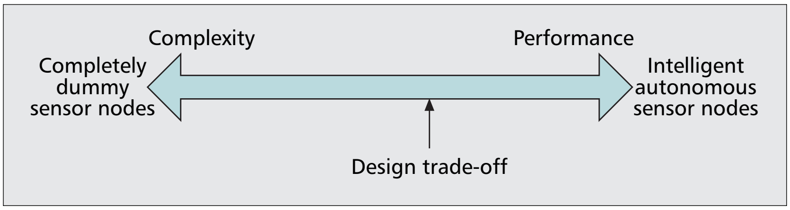

**Figure 1.** *The design space and trade-off between network edge and core.*

The extended design space is illustrated in Fig. 1. On one extreme, fully distributed protocols can be implemented in advanced sensor nodes without involvement of edge systems (as in many existing approaches). On the other end, network control is shifted to the base stations, thus allowing very simple sensor node design on tiny scales. The edge-system-based approach also allows centralized algorithms to be executed at the base station, which usually yields improved performance over distributed ones. For further improvements, multiple base stations can be deployed to form an upper tier over the underlying sensor network. A distributed algorithm can be adopted to coordinate the operation of the base stations and for improved reliability.

For such an extended design space, some interesting problems need to be addressed:

- How to effectively involve base stations in network control
- How to achieve a proper partition of basic functions between the edge and core
- How much simplification can be achieved in sensor node design

Clearly, for each application and network control task, the optimal division of functionalities between the network edge and network core should be carefully examined and handled differently.

## BEAMSTAR — A CASE STUDY

### A SIMPLE EXAMPLE

To have an edge-based system (e.g., a base station) participate in network control and routing, the first and foremost capability it must have is to locate a specific region in the network area and communicate with the sensor nodes located in that region. For this purpose, we could use directional antennas at a base station. Recent advances in low-cost computing technology and signal processing for arrays of antenna elements have made such smart directional antennas available for wireless communications.

We use Fig. 2 to illustrate the proposed approach. We assume line-of-sight conditions, and address the more challenging case of cluttered environments and channel dynamics later. The idea is to use power control and beam-forming at a base station to deliver location information to sensor nodes. Rather than finding accurate location for each sensor node as in [4–6], in BeamStar a base station logically partitions the network area into cells (or regions) via manipulation of beam radius, beamwidth, and beam orientation.

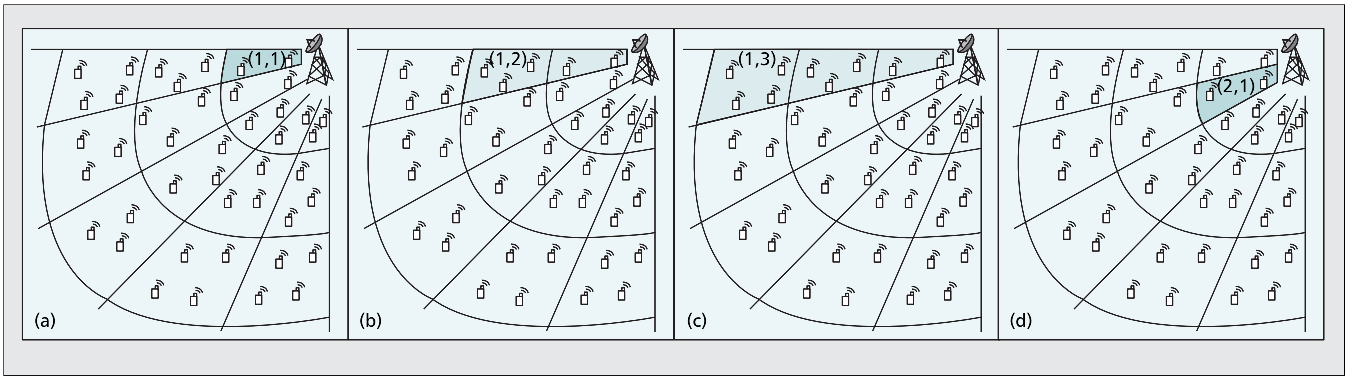

**Figure 2.** *An edge-assisted network organization and location service.*

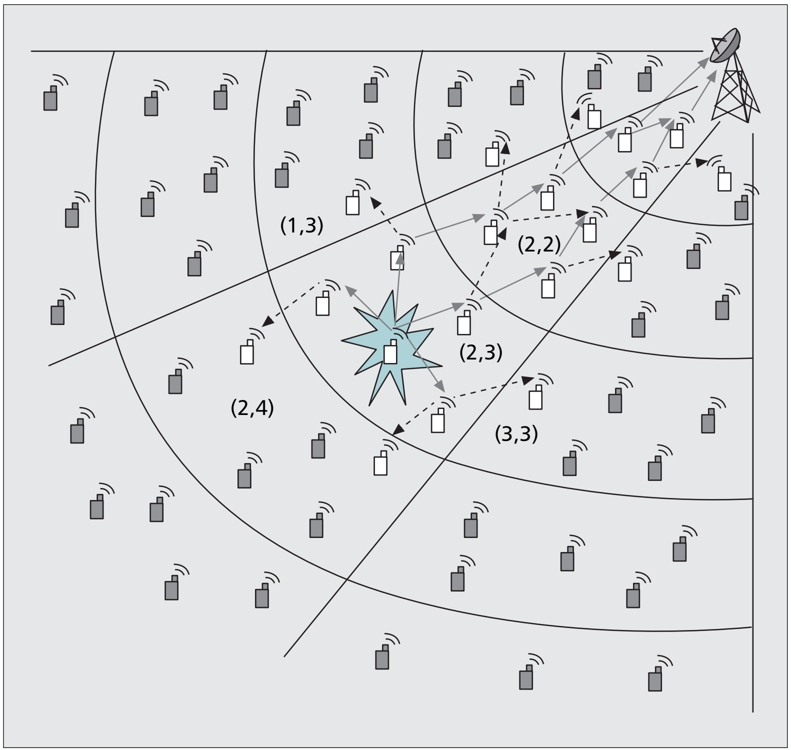

**Figure 3.** *Illustration of data forwarding after an event is detected.*

In Fig. 2a the base station makes a transmission (to the shaded area) by adjusting beam orientation and beamwidth with proper transmit power. The broadcast message carries phase and distance information, denoted (sector number, ring number) = (1,1) to all nodes within this region. The sector number corresponds to phase angle, and the ring number represents the radius from the base station. Sensors receiving this message will store (1,1) as their ID. Subsequently, in Fig. 2b the base station adjusts its transmit power to the next higher level and sends out a message with message (1,2). Note that those sensor nodes that have acquired ID (1,1) from the earlier transmission will ignore the second message. The remaining nodes without location information will store this message, (1,2), as their ID. Figures 2c and 2d illustrate subsequent base station transmissions for the other two cells with messages (1,3) and (2,1), respectively.

For packet forwarding, the base station can explicitly tell each region which sensor reports (i.e., reports from a specific neighboring region) should be forwarded (termed *forwarding rules*) during the configuration phase. Once location information and forwarding rules are distributed, the base station may change its antenna to omnidirectional mode for reception.

The routing process is illustrated in Fig. 3. Suppose a sensor node in a region with ID (2, 3) detects an event. Each sensor that detects this event will generate a report, as well as start a backoff timer with a value reversely proportional to the detected signal strength. The sensor node whose timer expires first will broadcast the report. Other sensors will cancel their transmission if they overhear this transmission. Thus, the amount of traffic generated by a single event will be reduced. Upon receiving this packet, the neighboring sensor nodes (with the same sector and ring numbers) will further broadcast it. Based on the packet signature (e.g., timestamp and other information), each sensor node will broadcast this packet only once.

At the boundaries of this region, suppose only sensor nodes in regions with ID (2, 2) are configured to relay this packet. Packets overheard by sensor nodes in other neighboring regions will be discarded. Such a decision can easily be made via a simple comparison of locally stored forwarding rules and the identifier of the last relay node. By following this process of broadcast, reception, drop, or rebroadcast among the sensor nodes, the report (or a small number of duplicate copies) will be delivered to the base station, from region to region.

The proposed approach has the intrinsic property of forwarding a number of duplicate packets without the need to flood the entire network. By relaying a small number of duplicate copies of the same packet toward the base station, potential loss of a sensed event can be greatly reduced. This property is highly desirable in harsh environments, in which a high probability of node failure and packet loss is expected.

## BEAMSTAR PERFORMANCE

***Reduced Complexity*** — The use of base station capabilities has greatly reduced the BeamStar routing complexity at a sensor node. Unlike core-based approaches, there is no need for peer sensor nodes to exchange information and calculate their location information. Instead, location information is directly delivered to the sensors in

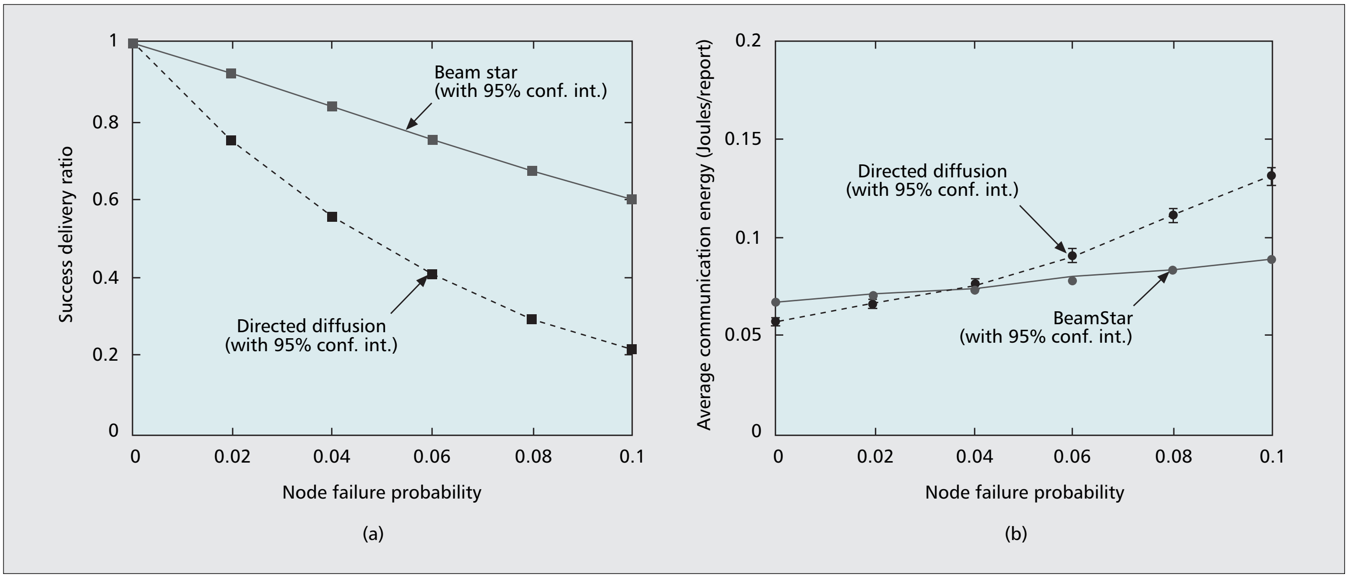

**Figure 4.** *Comparison with directed diffusion: a) successful delivery ratio; b) average communication energy consumption.*

control messages transmitted by the base station. It can be shown that the computational complexity of BeamStar is $O(1)$ and the storage requirement is also $O(1)$, while its control overhead is comparable to existing core-based approaches (e.g., DD [9]) [7].

Note that in addition to routing, some other important network control functions can also be greatly simplified. For example, *synchronization* can now be achieved with very low overhead by broadcasting control messages from the base station. Again, each sensor node only needs to receive several such control messages in order to synchronize with the base station clock. In addition, MAC protocol design can also be greatly simplified. The simple *broadcast-after-random-delay* scheme is sufficient for BeamStar. Those MAC protocols that implement sleep scheduling for sensor nodes can also be simplified with assistance from the base station.

***Comparison with a Core-Based Approach*** — We compare the performance of BeamStar with that of DD [9], one of the most popular data-centric routing protocols for wireless sensor networks. Both protocols are implemented using the OPNET Modeler. The results in Fig. 4 are obtained from a 256-node network randomly deployed over a 500 × 500 $m^2$ field. In BeamStar, the base station partitions the network into seven rings and 12 sectors. As in [9], we let one base station stay at a corner and one source node stay at the diagonal corner of the field. The same energy consumption model from [9] is used in all the simulations. Each simulation is repeated 10 times, and 95 percent confidence intervals are computed and plotted.

In Fig. 4a we plot the successful delivery ratios vs. node failure probability $p$. For the loss-free case, both protocols achieve near 100 percent delivery. As node failure probability increases, however, the BeamStar curve becomes consistently higher than the DD curve, indicating that for bursty losses, forwarding packets via a broadcast mesh is much more effective than unicast routing.

In Fig. 4b we present the average communication energy consumption results. We find that the average communication energy is comparable for the two protocols for low $p$ values. As $p$ increases, the DD average communication energy becomes larger, while the BeamStar average communication energy increases at a slower rate. This is in fact due to the different routing mechanisms used in these protocols. BeamStar achieves reliability by constrained broadcast. Although each broadcast transmission is less reliable, forwarding reports using a mesh is more effective in combating fragile nodes. The communication energy consumption is amortized over a larger number of successfully delivered reports.

More simulation results of BeamStar can be found in [7]. We find that BeamStar consistently achieves a higher successful delivery ratio than the core-based approach at comparable energy costs. BeamStar also enables much simpler hardware and software requirements on a sensor node. As a result, the size and cost of sensor nodes can be reduced substantially.

## ADDITIONAL CAPABILITIES

Having described the basic operation of BeamStar, we discuss several additional functions that are also enabled by the edge-based approach, but are otherwise more difficult to accomplish under the traditional core-based approach:

- *Query for events*. To query a specific region, the base station can simply adjust its beam direction $\alpha$ and transmit power $R$ to transmit a query message carrying the ID of the target region. Upon receiving such a query, sensors with the target ID will respond with sensor reports. Location error can easily be handled by having the base station scan the target region with $\alpha \pm \Delta\alpha$ and $R \pm \Delta R$, where $\Delta\alpha$ and $\Delta R$ are estimated location errors.
- *Hole detection and mitigation*. Detection of connectivity holes can be achieved by querying the target regions from the base station. Absence of response when a timer expires indicates loss of connection in that sector.

We find that BeamStar consistently achieves higher successful delivery ratio than the core-based approach at comparable energy cost. BeamStar also enables much simpler hardware and software requirements on a sensor node. As a result, the size and cost of sensor node can be reduced substantially.

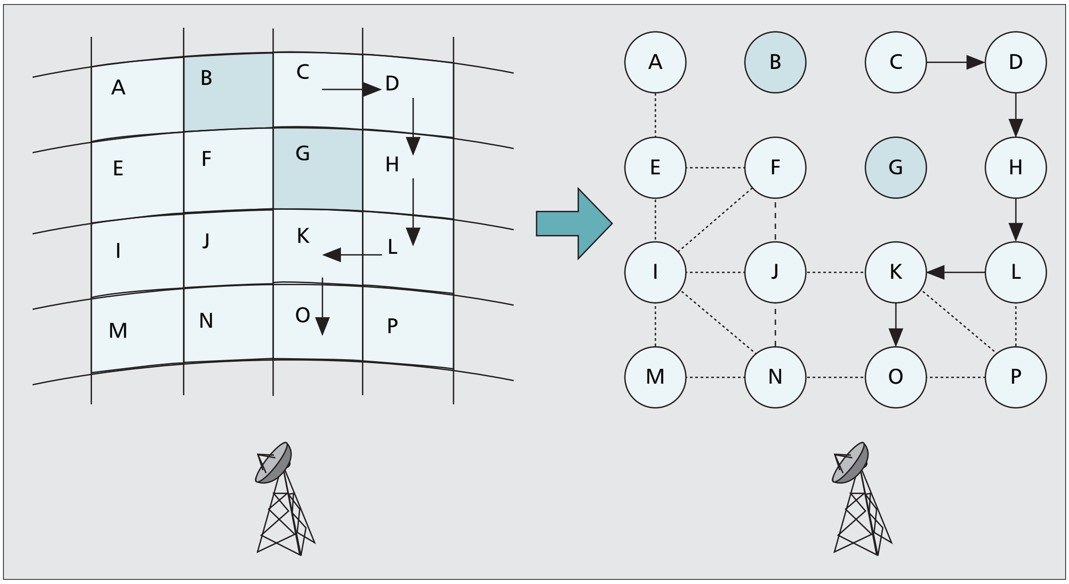

**Figure 5.** *Illustration of edge-assisted optimal routing.*

Once such a hole is detected, the base station can reconfigure the routing behavior of neighboring regions to bypass the hole or repartition the network to restore connectivity.

- *Granularity control of region size*. By manipulating the beam direction with overlapped transmissions, sectors with a span smaller than the beamwidth can be created. Combined with adjusting transmit power, uneven partition of the network regions can also be created [7].

We refer interested readers to [7] for details. In addition, we envision several interesting extensions of the basic BeamStar protocol, which are presented below.

***Computing and Distributing an Optimal Routing Path*** — The idea of reconfiguring routing behavior of certain regions can be further exploited, particularly if certain objectives or constraints should be met. For example, we could bypass certain spots in the network that are known to be unreliable due to a harsh or hostile environment. We could also achieve load balancing, maximize network lifetime, and optimize many other routing-related objectives.

When sensors relay reports toward a base station, they can piggyback useful information in the report packets. Such information may include:

- Connectivity with regard to a neighboring region, which can be inferred when overhearing transmissions from sensors in that region.
- Energy status of the region, such as minimum/maximum/total/average remaining energy among the sensor nodes in that region. For example, the sensor report could have a few "minimum remaining energy" fields, one for each region it traverses. When a sensor receives a report, it updates the field with its own residual energy if it is lower than the value in the report. The minimum remaining energy of a region will finally be updated in the report as it is forwarded by the sensors in that region.
- Other information such as congestion, delay, and loss rate if necessary, in a similar manner.

Upon receiving such piggybacked information, the base station can derive a virtual graph where each region is represented by a virtual node, as illustrated in Fig. 5. The base station will derive the metrics for the virtual nodes (e.g., remaining energy) and virtual links (e.g., how robust the connection is to a neighboring region) from feedback information. The base station then computes "optimal" routes for this virtual graph. Note that *centralized QoS routing algorithms* can be used at the base station, which usually achieve better performance than their distributed counterparts; and computational complexity is not a big issue since the base station is much more powerful than a bunch of sensor nodes in terms of computational power. Once the optimal paths are computed, the base station will use directional antennas to modify the forwarding rules of the sensors in certain cells, thus establishing the optimal paths.

***Deploying Multiple Base Stations*** — We have so far assumed a single base station as a reference model to illustrate how complexities on sensor nodes can be greatly reduced with the edge-based paradigm. In a highly cluttered environment, however, there may be no direct connection from the base station to certain areas in the field. Another potential problem is reliability: the base station becomes a single point of failure. Furthermore, a single base station has only limited coverage and is not scalable if the network size increases.

To address all these issues, multiple base stations can be deployed. A simple solution is to have each base station cover a specific area (possibly overlapping) and work independently. As a result, the impact of obstacles in the network area could be effectively mitigated. An interesting problem is how to determine the optimal location for each base station, given the network environment map and a total cost constraint, which can be formulated as an optimization and network planning problem and solved at the base stations without involving any sensor nodes.

With a more sophisticated approach, base stations can be synchronized and cooperatively control the operation of the wireless sensor network. Prior work [4–6] has provide strong evidence that such collaboration can greatly increase localization precision. The base stations can form an overlay upper tier that can improve the scalability of such networks, in addition to load balancing.

## PROPAGATION CONSIDERATIONS

An important practical consideration is the nonideal (i.e., non-free space) characteristics of wireless signal propagation. Generally, a directional channel is more reliable than an omnidirectional antenna channel due to more focused transmissions, as well as the general absence of mobility in sensor networks. However, if the environment is highly cluttered, sensor locations might be less accurate due to:

- Obstacles and various propagation conditions within different sectors
- Multiple receptions caused by multipath reflections

In a cluttered environment the contour of a ring will be indented behind an obstacle. Likewise, if the path losses at different directions are different, the contour of the rings will also be irregular. For multipath reflections, a reflected control message from a lower numbered sector (or in the exactly opposite line-of-sight direction) will be ignored [7]. If a sensor receives a control message reflected from an obstacle in a higher numbered sector, those sensors lying between this sensor and the reflecting obstacle will also get this control message if signal propagation is continuous. These sensors will set their *sector number* to the higher one received from this reflection. Thus, multipath reflections from a higher numbered sector will make the sector partitions irregular along the reflection trajectory.

Therefore, the joint effect of the nonidealities is that the partitioned regions will be irregular. However, BeamStar does not require regular partitioning of the network. This is quite different from the class of localization papers that aim to provide accurate location for sensor nodes [4–6]. Rather, BeamStar only requires roughly partitioning the network into regions, which allows sensor reports to be forwarded to the base station region by region. This is due to the benefits of the constrained broadcast adopted in BeamStar.

Deploying multiple base stations can mitigate the effects of nonideal propagation, as shown in [6]. In addition, there are several software tools for radio propagation prediction and network planning (e.g., SitePlanner from Wireless Valley and WinProp from AWE Communications). Such tools take the environment map as input, and produce an accurate estimate of signal strengths in the field. Equipped with such tools, the base station can make accurate calibration for the distortion caused by multipath reflections and obstacles, and adjust its beamwidth, transmit power levels, or location accordingly to mitigate these effects.

## CONCLUSION

In this article we propose exploring new design space for efficient wireless sensor networking. We review existing approaches in wireline and wireless networks from this perspective. We then consider routing as a canonical example and present BeamStar, a novel edge-based routing protocol. We show that BeamStar can greatly simplify localization, synchronization, and routing, thus enabling considerable size/cost reduction at sensor nodes. While our focus in this article is on routing issues, it is conceivable that such an edge-based paradigm can be exploited to simplify other complex tasks, such as code distribution, programmability, tasking, and query.

While our focus is on routing issues, it is conceivable that such edge-based paradigm can be exploited to simplify other complex tasks, such as code distribution, programmability, tasking and query.

## REFERENCES

[1] B. Warneke *et al.*, "Smart Dust, Communicating with a Cubic-Millimeter Computer," *IEEE Computer*, vol. 34, no. 1, Jan. 2001, pp. 44–51.

[2] V. Firoiu *et al.*, "Theories and Models for Internet Quality of Service," *Proc. IEEE*, vol. 90, no. 9, Sept. 2002, pp. 1565–91.

[3] B. Liu, Z. Liu, and D. Towsley, "On the Capacity of HYBrid Wireless Networks," *Proc. IEEE INFOCOM '03*, San Francisco, CA, Mar. 2003, pp. 1543–52.

[4] K. Romer, "The Lighthouse Location System for Smart Dust," *Proc. ACM MobiSys '03*, San Francisco, CA, May 2003, pp. 15–30.

[5] A. Nasipuri and K. Li, "A Directionality based Location Discovery Scheme for Wireless Sensor Networks," *Proc. ACM WSNA '02*, Sept. 2002, pp. 105–11.

[6] D. Niculescu and B. Nath, "VOR Base Stations for Indoor 802.11 Positioning," *Proc. ACM MobiCom '04*, Sept./Oct. 2004, pp. 58–69.

[7] S. Mao and Y. T. Hou, "Beamstar: An Edge-Based Approach to Routing in Wireless Sensor Networks," *IEEE Trans. Mobile Computing*, vol. 6, no. 11, Nov. 2007, pp. 1284–96.

[8] K. Akkaya and M. Younis, "A Survey on Routing Protocols for Wireless Sensor Networks," *Elsevier Ad Hoc Networks Journal*, vol. 3, no. 3, May 2005, pp. 325–49.

[9] C. Intanagonwiwat *et al.*, "Directed Diffusion for Wireless Sensor Networking," *IEEE/ACM Trans. Net.*, vol. 11, no. 1, Feb. 2003, pp. 2–16.

[10] Y. Xu, J. Heidemann, and D. Estrin, "Geography-Informed Energy Conservation for Ad Hoc Routing," *Proc. ACM MobiCom'01*, Rome, Italy, July 2001, pp. 70–84.

[11] W. B. Heinzelman, A. P. Chandrakasan, and H. Balakrishnan, "An Application-Specific Protocol Architecture for Wireless Microsensor Networks," *IEEE Trans. Wireless Commun.*, vol. 1, no. 4, Oct. 2002, pp. 660–70.

## BIOGRAPHIES

SHIWEN MAO [S'99, M'04] received a Ph.D. degree in electrical and computer engineering from Polytechnic University, Brooklyn, New York, in 2004. He was a research scientist at Virginia Tech, Blacksburg, and currently is an assistant professor in the Department of Electrical and Computer Engineering, Auburn University, Alabama. His research interests include cross-layer optimization in multihop wireless networks and multimedia communications. He was a co-recipient of the 2004 IEEE Communications Society Leonard G. Abraham Prize in the field of communications systems.

Y. THOMAS HOU [S'91, M'98, SM'04] (thou@vt.edu) is an associate professor of Electrical and Computer Engineering at Virginia Tech. His current research interests are radio resource management and networking for cognitive radio networks, optimization and algorithm design for wireless ad hoc and sensor networks, and video communications over dynamic ad hoc networks. From 1997 to 2002 he was a researcher at Fujitsu Laboratories of America, Sunnyvale, California. He is a recipient of an Office of Naval Research (ONR) Young Investigator Award (2003) and a National Science Foundation (NSF) CAREER Award (2004).

MIN-YOU WU [M'86, SM'96] is an IBM Chair Professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. He serves as chief scientist at the Grid Center of Shanghai Jiao Tong University. His research interests include grid computing, wireless networks, sensor networks, overlay networks, multimedia networking, parallel and distributed systems, and compilers for parallel computers.