A Distributed, Utility-based Architecture for Task Assignment in Tactical WSNs*

Kangyuan Zhu and Stephen D. Patek
Systems and Information Engineering
University of Virginia
Charlottesville, VA
{kz3y, patek}@virginia.edu

Abstract

This paper develops an architecture for task assignment in wireless sensor networks in support of tactical operations. Our approach is based on a mathematical framework, which, based on multi-attribute utility theory, allows mobile end-users to easily express the value (utility) that they would attach to sensor information based on features associated with the data. We have applied the utility-based architecture framework in the design of a distributed task assignment algorithm for a simple pursuit-evasion (patrol) scenario. Simulation studies based on this scenario show that in terms of mission-level objectives utility-based task assignment can significantly improve upon a baseline solution, in which each sensor indiscriminately reports all data to all available users.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design; C.2.4 [**Distributed Systems**]: Distributed Applications

General Terms

Algorithms, Design

Keywords

Sensor networks, Task assignment, Utility

1 Introduction

The last decade has seen a tremendous amount of research and development directed toward sensing technologies, sensor development, and the application of wireless sensor networks (WSNs). These efforts have targeted applications in many domains, such as military applications, environmental applications, health applications, home applications, and other commercial applications [2, 9].

Tactical operations are an important military application of WSNs. The features of tactical operations are different from other applications of WSNs in that (a) distinct end-users are working amidst sensors in the field and may directly receive message from individual sensors in a decentralized fashion, (b) end-users are mobile and have dynamic features, (c) sensors make decisions about data dissemination based on both the information they collect and end-users' dynamic features, and (d) standard metrics of network performance, such as data delay and data throughput, may not appropriately reflect the overall system objectives. These features for tactical operations thrown into question the applicability of existing architectures for task assignment.

There have been a number of recent proposals for network architectures that explicitly take into consideration the applications and characteristics of WSNs [16, 2, 1, 9]. Lim [11] proposed an architecture for information dissemination in self-organizing sensor networks, which involves application systems, configurable systems, sensor networking, and physical device layers. Estrin et al. [6] proposed a localized algorithm of directed diffusion to establish flexible and efficient data delivery paths in WSNs. Subramanian and Katz [20] proposed a generic architecture for self-configurable systems where a large number of sensors coordinate amongst themselves to achieve a large sensing task.

Despite the diversity of available network architectures, so far WSNs are mainly used as fixed infrastructures for data collection for specific applications. Hence, data processing and/or other high-level application functions supports have to be integrated with the sensor networks [16, 9]. Recent research tends to integrate data processing requirements with network-level considerations. Ganesan et al. [7] presented a data handling architecture, DIMENSIONS, a system that provides a unified view of data handling in sensor networks, which has the ability to incorporate the resource constraints and spatio-temporal interpretation of the physical world. Madden and Franklin [12] developed the Fjords architecture for managing multiple queries over many sensors and showed that it can limit the sensor resources used while maintaining high query throughput. Shen et al. [19] introduced a sensor information networking architecture named SINA, with which users can access information using declarative queries, or perform tasks using programming scripts.

Although various data processing methods are integrated with network architectures, most current research on architectures only focuses on the collection of data, with the flow of information going mainly from sensors to the users through a fixed gateway. Thus, while there might be a large set of diverse applications supported by sensor data, the responsibility of the WSN is mainly to collect and transmit

*This work was supported in part by the Wireless Internet Center for Advanced Technology (WICAT), an NSF-sponsored IUCRC.

data. Future sensor networks are likely to be deployed in a manner that data collection is only part of the overall responsibility of the WSN, such as in tactical operations. In this case, current architectures may not reflect overall system objectives leaving significant opportunities for improvement.

In this paper we attempt to establish an architecture that can reflect overall system objectives for specific applications of WSNs. In particular, we develop a sensor-centric (decentralized), utility-based architecture for task assignment in wireless sensor networks in which sensors individually make judicious decisions about what information should be forwarded to a collection of receivers based on the "value" of the information encoded within a utility function.

Literature Review

The benefits of decentralization in WSNs are summarized in several papers [15, 17]. Wen and Sethares [21] proposed a decentralized algorithm for WSN clustering, with which each sensor adopts a random waiting timer and local criteria to determine whether to form a new cluster or to join a current cluster. Makerenko et al. [15] developed an approach adhering to scalable decentralized algorithms for both data fusion and the decision-making layers of the system for an indoor Active Sensor Network (ASN). Makarenko and Durrant-Whyte [14] presented an algorithm for Bayesian decentralized data fusion (BDDF) and its extension to information theoretic control. Sadagopan et al. [18] advocated a systematic decentralized approach to designing networks based on utility functions to achieve the global optimal load for data gathering tree. Mainland et al. [13] presented a self-organizing resource allocation for achieving efficient resource allocation in sensor networks based on the decentralized, utility-defined action selections of individual sensors. Ridley et al. [17] described the theoretical and practical development of a decentralized air and ground sensing network for target tracking and identification with the information-filter formulation of the Kalman filter algorithm and information-theoretic methods from Bayes theorem.

Utility theories also have been used in various fields of WSN to represent the value of information. Chen and Sha [5] formulated data transport problem in WSN as an optimization problem to achieve the maximal amount of utility collected at sinks subject to flow, energy, and channel bandwidth constraints. Byers and Nasser [4] presented a model for numbers of sensors participating sensing to define appropriate global objectives based on utility functions and specify the cost for energy consumption. Zhao et al. [22] introduced and developed the definition of information utility and several approximate measures of the information utility, with which the paper determined the participants in a sensor collaboration by dynamically optimizing the information utility of data. Kang and Li [8] used the information utility measurement for decision making in sensor selection of clustering with three key factors: sensing quality, communication cost and power level. Bian et al. [3] proposed a framework to select a sequence sensor sets which has the maximal utility for sensor selection techniques.

2 Architectural Issues

There are two main questions to resolve in designing a task assignment architecture for WSNs: (1) Where should the authority for making decisions about consumption of network resources lie? (2) If control authority is decentralized, what information needs to be exchanged in order to ensure effective and efficient operation of the network? Here, we have focused on understanding the performance that can be achieved with no explicit coordination between sensors. In Section 2.1 we briefly outline a baseline solution for data dissemination derived from existing WSN's network architectures. In Section 2.2 we develop the Sensor-Centric, Utility-Based (SCUB) architecture, which is completely decentralized and no explicit information exchanges between sensors.

2.1 Baseline Solution: Send-All-to-All

Given that there are multiple mobile end-users operating within the sensor network, a straightforward approach to data dissemination is to require all sensors to send all detections to all users, up to constraints on network resources including power, energy, bandwidth, and link availability. This "Send-All-to-All" architecture for task assignment is simple, and its main benefit is that sensors do not have to waste resources by communicating with one another (or with a base station) to determine whether it is appropriate to send sensor data to an individual user. On the other hand, all sensor information within this architecture is treated as equally important, when in reality information from neighboring sensors may be highly correlated and/or irrelevant to accomplishing mission objectives. In addition, the "Send-All-to-All" architecture assumes that all information is equally relevant to all end-users, where in reality some users may find the data more actionable than others. Implementation of this approach requires a routing protocol for sensor to end-user communications, and, in our evaluation of this scheme, we assume (as described below) that sensors are periodically informed about the positions of end-users.

2.2 SCUB Architecture

Instead of indiscriminately reporting all information to all end-users, the Sensor-Centric, Utility-Based architecture requires that decisions about data dissemination are based on a common model for the utility that end-users would attach to data given that it is received. The utility model serves to quantify the value that users attach to data based on a vector of features associated with the data (e.g. type of observation, location, etc.) and on a vector of user-adjustable parameter values that express the relative importance of data's features to users. By introducing a utility model into the framework, we hope to achieve a common value system for all sensors, so that, even though they are operating independently, they still have a means of discriminating between transmission opportunities and hopefully consuming network resources by transmitting only high-value observations.

As a more formal description of the utility framework, we assume that there is a set S of sensors that constitutes the sensor network. Each sensor $s \in S$ keeps in memory a history O_s of recent observations, and for each observation $o \in O_s$ there is a set of messages M_{so} that could be sent based on those observations. (It may be possible to send a number of different messages based on the same observation, e.g. imagery with different levels of compression. In the "Send-All-to-All" architecture, the strategy for sending imagery would have to be predefined.) The set $M_s = \cup_{o \in O_s} M_{so}$ represents the collection of all messages that are currently available for sensor s to send based on its recent history of observations. Let G_s represent the set of users that are potential recipients of messages $m \in M_s$ where each "user" may actually correspond to a group of end-users, applications, and/or software agents that could benefit by receiving the message m. In choosing which, if any, of the messages $m \in M_s$ to transmit to user $g \in G_s$, the sensor must compute the utility $u_g(m)$ that g would place on each message m if it were to be received. Mathematically, the function u_g takes the form $u_g(m) = U_g(f(m); \theta_g)$, where $f(m) = (f_1(m), f_2(m), \ldots f_F(m))$ is a vector of features associated with the message, and $\theta_g = (w_1, w_2, \ldots, w_W)$ is a vector of user-adjustable parameters that help to quantify the utility of the message m for user g. Features associated with messages are quantitative characteristics of the data associated with the message, such as the location of observation and age of the observation. The user adjustable parameters in θ_g include various scaling coefficients and weights, which ensure that the overall assessment of utility $u_g(m)$ is normalized to lie between zero (corresponding to no value) and one (corresponding to maximal value).

Having the ability to compute $u_g(m)$ for all possible messages, sensors have a quantitative means of discriminating between messages, and any decision rule could be implemented as a policy for determining which message to send to which user next.

2.2.1 *Optimal SCUB Task Assignment*

Since the SCUB architecture dictates that sensors act independently in forwarding observation data to users, what we mean by "optimal" task assignment is best characterized in game-theoretic terms. Our goal here is mainly to suggest a mathematical framework for understanding the general issues in SCUB task assignment.

We assume that each sensor s maintains a history H_s of messages that have already been transmitted to users in G_s. Thus, in considering whether to send a given message $m \in M_{so}$ to a particular user g the sensor can determine from H_s whether it has already sent a related message to that particular user, thereby having the ability to avoid self-generated redundancy in its transmissions.

Since the essence of the task assignment problem is the allocation of scarce WSN resources, let R denote the set of common resources associated with the sensor network. Let C_r denote the capacity of resource $r \in R$, and let x_r denote the current utilization of resource r. For example, if resource r corresponds to a communications link in the network, then C_R would correspond to the bandwidth associated with the link and x_r would correspond to the percentage of that bandwidth that is currently being utilized. In this case, x_r naturally characterizes the state of that communications channel, and the task assignment policy embedded within the sensor simply will not consider sending any message that would cause x_r to exceed the capacity C_r of the channel. (The message may be transmitted later.) For resources like bandwidth, the utilization level x_r may fluctuate up and down according to the need for sensors to transmit observation data. For other resources, like total battery energy, the consumption of the resource is monotonic.

Optimal SCUB resource management involves each sensor sending messages to users, subject to resource constraints, so as to maximize the expected aggregate accumulated utility associated with the information that is transmitted. To account for the independent operation of each sensor, we assume that each sensor $s \in \mathcal{S}$ implements a Strategy

$$\mu_s(M_s, G_s, H_s, (x_r)_{r \in \mathcal{R}})$$

that prescribes the message $m \in M_s$ to be transmitted to user $g \in G_s$ next (if any) depending on the current consumption of resources $(x_r)_{r \in \mathcal{R}}$ and on the history H_s of messages sent so far. Thus, we seek to compute a profile of strategies $(\mu_{s_1}, \mu_{s_2}, \ldots, \mu_{s_S})$ that maximizes:

$$F(\mu_{s_1}, \ldots, \mu_{s_{|S|}}) = \mathbb{E}\left\{ \int_0^T g_{\mu_{s_1}, \mu_{s_2}, \ldots, \mu_{s_{|S|}}}(t)dt \right\}, \qquad (1)$$

where T is the random time horizon of the problem and $g_{\mu_{s_1}, \mu_{s_2}, \ldots, \mu_{s_{|S|}}}(t)$ is the random impulse train associated with discrete chunks of utility associated the transmit-decisions that are made by sensors in the network under the strategies $\mu_{s_1}, \mu_{s_2}, \ldots, \mu_{s_{|S|}}$. Note that, in addition to accounting for the possibly randomized behavior of the sensors in the network, the expectation Equation (1) involves uncertainty over the random nature of observations made by individual sensors, which in turn depends on the randomness associated with the environment. $F(\mu_{s_1}, \ldots, \mu_{s_{|S|}})$ can be thought of as a mission-level utility function, representing the collective goals and objectives of all end-users for the duration of the systems lifetime. Optimal SCUB task assignment, can thus be thought of as an identical-interests game, in which each sensor $s \in S$ is a player implementing a strategy μ_s to achieve a high-value global equilibrium.

2.2.2 *Sub-Optimal SCUB Task Assignment*

Realistically, a complete mathematical specification of probabilistic model of Equation (1) can only be made in the context of specific applications built around prohibitively specific scenarios. Moreover, even if the were available, the computation of an optimal profile of task assignment strategies most likely be intractable. Consequently, in the remainder of the paper (as a first cut analysis), we only consider a heuristic strategy designed to improve the chances of achieving system-level objective, while falling short of constituting an optimal solution. We refer to the strategy as "Myopic Utility Maximization," in which, at every transmission opportunity, each sensor will send the message m^* that offers the highest utility to a corresponding user g^*. In other words, the message/user pair (m^*, g^*) for the next transmission is a maximizing solution to the optimization problem $\max_{g \in G_s, m \in M_s} u_g(m)$.

We emphasize that myopic utility maximization is not the only task assignment strategy admitted by the SCUB architecture. For example, in on-going work, we are evaluating

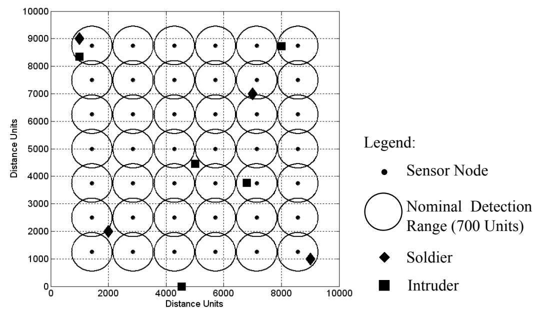

Figure 1. Sensor Network Configuration

other heuristics that incorporate a notion of cost or disutility associated with the consumption of WSN resources that offsets the raw utility that end-users perceive in receiving sensor data. The trick to such an approach is to adaptively adjust resource costs to reflect long-term objectives.

3 Numerical Evaluation

To illustrate the application of the SCUB architecture, we have built a simulation test bed, which allows us to compare SCUB with the "Send-All-to-All" architecture for task assignment. We briefly describe the test bed in Section 3.1 and utility models in Section 3.2. We present representative simulation results in Section 3.3, referring the reader to [23] for more details and additional simulation results.

3.1 Patrol Scenario Test Bed

We have developed a simulation test bed based on the "patrol scenario" of Figure 1, which involves (a) soldiers patrolling subregions within a larger Area of Interest (AOI), (b) a network of sensors reporting detections of intruders to soldiers, and (c) the pursuit and capture of intruders that traverse the AOI. The scenario is coded in JAVA as a discrete simulation event with a simulation interval of 1 second.

We assume that all sensors within the AOI are tripwire sensors with 360 degree field of view and given detection radius, evenly distributed within the AOI and powered by two AA batteries. Each sensor has the ability to make detections of intruders, to receive, transmit, and forward messages, and to engage in antenna idling. Such function consume different amounts of power and energy. The sensors have ability to detect the intruders with a fixed detection probability once the intruders are within their detection ranges. Intruders randomly appear at the boundary of the AOI and then randomly cross the AOI until captured or until from the system by leaving the AOI. Soldiers are assigned different sub-areas (rectangular) within the AOI and randomly patrol within their sub-areas until they are cued as to the presence of intruders by sensors, at which time they will commence to pursue the intruders, moving at a faster speed than the intruders. Soldiers in pursuit of intruders may require more than one detection report before the intruder is within visual range, in which case the soldier will pursue the intruders without needing (or wanting) additional messages. If a soldier receives multiple messages from different sensors, the soldier only moves in the direction of the nearest intruder. Upon capturing an intruder or upon concluding that the targeted intruder has escaped, soldiers become idle and proceed to patrol randomly until they cued by another detection report from one or more sensors. While in pursuit, soldiers are allowed to leave their assigned sub-areas.

3.2 SCUB Utility Model

We model the utility of a message m for soldier g as a multiplicative function of predefined marginal utilities that are associated with three features of interest as follows.

- *Feature 1: Age of Associated Observation* - We assume that the marginal utility that soldier g attaches to this feature is linear and saturates at zero, i.e.

  $$u_g^1(f_1(m)) = [1 - w_1 f_1(m)]^+,$$

  where $f_1(m)$ is the age of message m in seconds; w_1 is a g-specified coefficient that describes the maximum age a message can have and still be of any value. We set $w_1 = 1/5$ for all soldiers in the scenario.

- *Feature 2: Distance to sub-Patrol Area* - We assume that the marginal utility associated with the distance to the patrol area of soldier g, $u_g^2(f_2(m)) = 1$, when the observation m is made within the patrol area of soldier g. Otherwise, it is linear, saturating at zero, i.e.

  $$u_g^2(f_2(m)) = [1 - w_2 f_2(m)]^+,$$

  where $f_2(m)$ refers to the shortest Euclidean distance from the observation position to the upper left or lower right coordinates of the solider's sub-patrol area (rectangular); w_2 describes the largest Euclidean distance that a detection can have and still have value. We set $w_2 = 1/5000$ for all soldiers in the scenario.

- *Feature 3: Distance to Soldier Location* - We assume that the marginal utility for this feature is also linear, saturating at zero, i.e.

  $$u_g^3(f_3(m)) = [1 - w_3 f_3(m)]^+,$$

  where $f_3(m)$ is the Euclidean distance from the observation position to solider's position; w_3 describes the largest distance to soldier g that a detection can have and still have value. We set $w_3 = 1/14000$ for all soldiers in the scenario

Drawing upon insights from decision theory (see, for example, [10]), we model soldier-g's overall utility for message m as a multiplicative function of his marginal utilities, i.e.

$$u_g(m) = \frac{1}{k_0}\left[-1 + \prod_{i=1}^{3}\left(1 + k_0 k_i u_g^i(m)\right)\right],$$

where (a) the number 3 denotes the number of features, (b) the parameters k_1, k_2, k_3 are weighting parameters set by g to reflect the importance of each feature, and (c) the parameter k_0 is a constant set to be a solution to the equation

$$1 + k_0 = \prod_{i=1}^{3}(1 + k_0 k_i).$$

Each weighting parameter k_i $(i = 1, 2, 3)$ should be chosen to reflect the overall utility of a message where the i-th feature has its best possible value and all the other features have their worst possible values. In this scenario, we arbitrarily set $k_1 = 0.3, k_2 = 0.3, k_3 = 0.3, k_0 = 0.36$ for all the soldiers within the entire mission period.

3.3 Simulation Experiments

In this section, we present preliminary experimental results that validate the assertion that SCUB architecture can outperform the baseline "Send-All-to-All" approach, at least in terms of system-level performance metrics.

3.3.1 Experimental Metrics and Variables

In our preliminary experiments we have focused on two main performance metrics: (i) success rate and (ii) number of sensors alive, both plotted as a function of time for the duration of the network. Success rate is computed according to the following equation:

$$SR^{t_i,t_j} = \frac{N^{t_i,t_j}_{captured}}{N^{t_i}_{atLarge} + N^{t_i,t_j}_{generated} - N^{t_j}_{atLarge}},$$

where SR^{t_i,t_j} denotes the average success rate in time period $(t_i, t_j]$; $N^{t_i,t_j}_{captured}$ denotes the total number of intruders captured in time period $(t_i, t_j]$; $N^{t_i}_{atLarge}$ represents the number of intruders at large in time t_i; $N^{t_i,t_j}_{generated}$ denotes the total number of intruders generated in time period $(t_i, t_j]$. We computed the averages above within windows of $28,800$ seconds.

For the other metric, number of sensors alive, we sampled the number of sensors that have sufficient energy remaining to continue making detections throughout the lifetime of the network. We computed average number of sensors alive over time windows of 200 seconds.

With the two metrics in mind, we designed an experimental study to show the effect of two important scenario parameters: (i) number of soldiers within the system and (ii) intruder arrival probability (i.e. the likelihood that a new intruder appears on the boundary of the AOI in any discrete time step of the simulation- akin to intruder arrival rate). We conducted 20 independent trials, each time initializing each sensor with full battery power and then running the system until all energy from all sensors is depleted. Representative results from this experiment are shown in the next section.

3.3.2 Preliminary Results

Figures 2-5 illustrate the relative performance of SCUB with myopic utility maximization and the baseline "Send-All-to-All" solution. Figure 2 shows that, for both architectures, higher success rates can be achieved by deploying larger numbers of soldiers, although the magnitude of improvement associated with adding one more soldier is significantly larger for SCUB than for "Send-All-to-All." Figure 3 shows that, for "Send-All-to-All" but *not* for SCUB, larger numbers of soldiers in the system result in more rapid depletion of sensor-energy. For the "Send-All-to-All" architecture, energy load increases dramatically with the number of receivers; however, for SCUB, the more soldiers there are, the shorter the routes are from individual sensors to the receivers with highest utility. Figures 4 and 5 show that SCUB continues to outperform "Send-All-to-All" for varying rates at which intruders arrive, both in terms of success rate and numbers of sensors alive. We note from Figure 4 that for increasing intruder arrival probabilities, the success rates for "Send-All-to-All" decrease more dramatically than for SCUB. From Figure 5, we see that larger intruder arrival probabilities cause both architectures to deplete sensor energy more rapidly, since higher intruder arrival rates provide sensors with more opportunities to detect intruders. We point out that success rates in both Figures 2 and 4 (especially for "Send-All-to-All") decrease rapidly toward the end of the experiment due to fact that fewer sensors are alive, and intruders are more likely to escape without detection.

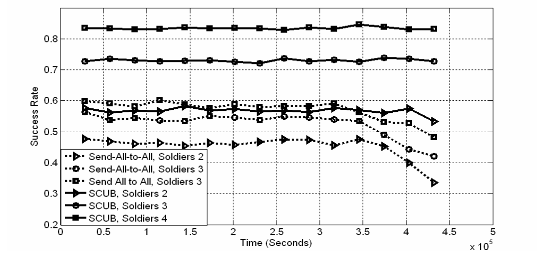

Figure 2. Success rate as a function of number of soldiers

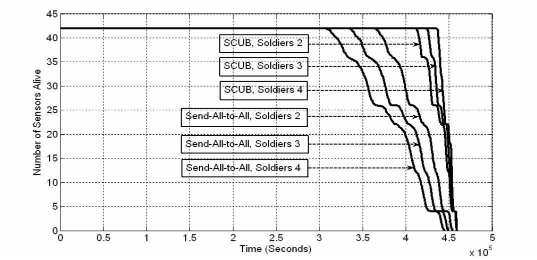

Figure 3. Number of sensors alive as a function of number of soldiers

4 Conclusions

Our work on architectures for task assignment in tactical WSNs, particularly on the SCUB architecture, is at a very preliminary stage. SCUB itself seems to represent a

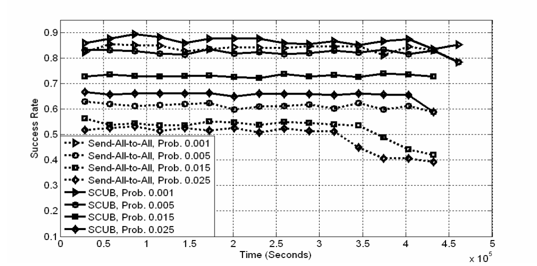

Figure 4. Success rate as a function of intruder arrival probability

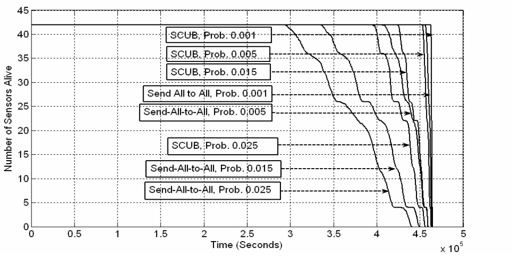

Figure 5. Number of sensors alive as a function of intruder arrival probability

reasonable approach to the problem of allocating the scarce resources of sensor networks for tactical operations, being highly decentralized and yet cognizant of the diverse information requirements of multiple users. Tactical operations bring into focus the notion that a wireless sensor network is far more than just a highly constrained data communications infrastructure, and an important point of departure in our work is our focus on application-layer (mission-level) performance metrics, rather than on lower-level performance metrics, such as the throughput and message delay.

5 Acknowledgements

The authors gratefully acknowledge the assistance of Xinyan Pan in helping to formulate the "patrol scenario."

6 References

[1] A. Ahmed and M. Eskicioglu. Current researches on sensor networks. Technical report, Telecommunication Research Labs,Winnipeg, Canada, 2004.

[2] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: a survey. *Computer Networks: the International Journal of Distributed Informatique*, 38(4):393–422, March 2002.

[3] F. Bian, D. Kempe, and R. Govindan. Utility-based sensor selection. In *The 5th International Symposium on IPSN*, Apr. 2006.

[4] J. Byers and G. Nasser. Utility-based decision-making in wireless sensor networks. *Mobile and Ad Hoc Networking and Computing, IEEE*, August 2000.

[5] W. Chen and L. Sha. An energy-aware data-centric generic utility based approach in wireless sensor networks. In *Proceedings of the 3rd international symposium on IPSN*, Apr. 2004.

[6] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: scalable coordination in sensor networks. In *ACM International Conference on Mobile Computing and Networks*, pages 263–270, 1999.

[7] D. Ganesan, D. Estrin, and J. Heidemann. Dimensions: why do we need a new data handling architecture for sensor networks? *Computer Communication Review*, 33:143–148, 2003.

[8] H. Kang and X. Li. Power-aware sensor selection in wireless sensor networks. In *The 5th International Conference on IPSN, Work-in-Progress*, 2006.

[9] H. Karl and A. Willig. A short survey of wireless sensor networks. Technical report, Technical University Berlin, Berlin, Germany, 2003.

[10] R. Keeney and H. Raiffa. *Decisions with multiple objectives: Preferences and value tradeoffs*. John Wiley and Sons, New York, 1976.

[11] A. Lim. Distributed services for information dissemination in self-organizing sensor networks. *Journal of the Franklin Institute*, 338(6), Sept. 2001.

[12] S. Madden and M. Franklin. Fjording the Stream: An Architecture for Queries Over Streaming Sensor Data. In *ICDE*, pages 555–566, 2002.

[13] G. Mainland, D. Parkes, and M. Welsh. Decentralized, adaptive resource allocation for sensor networks. In *Proceedings of the 2nd USENIX/ACM Symposium on Networked Systems Design and Implementation*, 2005.

[14] A. Makarenko and H. Durrant-Whyte. Decentralized data fusion and control in active sensor networks. In *Proceedings of the 7th Intl Conf on Information Fusion*, 2004.

[15] A. Makerenko, E. Nettleton, B. Grocholsky, S. Sukkarieh, and H. Durrant-Whyte. Building a decentralized active sensor networks. In *11th Intl Conf on Advanced Robotics*, Coimbra, Portugal, 2003.

[16] H. Qi, S. Iyengar, and K. Chakrabarty. Distributed sensor networks-a review of recent research. *Journal of the Franklin Institute*, 338(6):655–668, Sept. 2001.

[17] M. Ridley, E. Nettleton, S. Sukkarieh, and H. Durrant-Whyte. Tracking in decentralised air-ground sensing networks. In *Proceedings of the 5th Intl Conf on Information Fusion*, 2002.

[18] N. Sadagopan, M. Singh, and B. Krishnamachari. Decentralized utility-based sensor network design. *Mobile Networks and Applications*, 11(3), 2006.

[19] C. Shen, C. Srisathapornphat, and C. Jaikaeo. Sensor information networking architecture and applications. *IEEE Personal Communications*, 8(4), Aug. 2001.

[20] L. Subramanian and R. H. Katz. An architecture for building self-configurable systems. In *ACM/IEEE Workshop on Mobile Ad Hoc Networking and Computing*, August 2000.

[21] Y. Wen and W. Sethares. Automatic decentralized clustering for wireless sensor networks. *EURASIP Journal on Wireless Communications and Networking*, 5, 2005.

[22] F. Zhao, J. Shin, and J. Reich. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Processing Magazine*, 19(2), Mar. 2002.

[23] K. Zhu and S. D. Patek. SCUB: A Distributed, Utility-based Architecture for Closed Loop Control with Wireless Sensor Networks. Technical report, UVA, 2006.