

Achieving Long-Term Surveillance in VigilNet

PASCAL VICAIRE, TIAN HE, QING CAO, TING YAN, GANG ZHOU, LIN GU, LIQIAN LUO, RADU STOLERU, JOHN A. STANKOVIC, TAREK F. ABDELZAHER
University of Virginia

Energy efficiency is a fundamental issue for outdoor sensor network systems. This article presents the design and implementation of multi-dimensional power management strategies in VigilNet, a major recent effort to support long-term surveillance using power-constrained sensor devices. A novel tripwire service is integrated with an effective sentry and duty cycle scheduling in order to increase the system lifetime, collaboratively. The tripwire service partitions a network into distinct, non-overlapping sections and allows each section to be scheduled independently. Sentry scheduling selects a subset of nodes, the sentries, which are turned on while the remaining nodes save energy. Duty cycle scheduling allows the active sentries themselves to be turned on and off, further lowering the average power draw. The multi-dimensional power management strategies proposed in this paper were fully implemented within a real sensor network system using the XSM platform. We evaluate key system parameters using a network of 200 XSM nodes in an outdoor environment, and an analytical probabilistic model. We evaluate network lifetime using a simulation of a 10,000 nodes network that uses measured XSM power values. These evaluations demonstrate the effectiveness of our integrated approach and identify a set of lessons and guidelines, useful for the future development of energy-efficient sensor systems. One of the key results indicates that the combination of the three presented power management techniques is able to increase the lifetime of a realistic network from 4 days to 200 days.

Categories and Subject Descriptors: C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design; C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed Applications*; C.4 [**Performance of Systems**]: Design Studies

General Terms: Performance, Measurement, Design, Experimentation, Algorithm

Additional Key Words and Phrases: Applications of sensor and actuator networks, energy and resource management, network protocols, coverage, connectivity, and longevity, sensor networks, energy conservation, tracking

1. INTRODUCTION

VigilNet is a recent major effort to support long-term military surveillance, using large-scale networks composed of tiny resource constrained sensors. Besides re-

Author's address: Pascal A. Vicaire, Computer Science Department University of Virginia 151 Engineer's way -P.O. Box 400740 Charlottesville VA 22904-4740 USA.

Part of this work was published in IEEE INFOCOM 2006.

This work was supported in part by NSF grants CNS-0435060, CNS-0614870, CNS-0626616, and DARPA grants IXO under the NEST project F336615-01-C-1905.

T. He is now at the University of Minnesota.

Q. Cao, L. Luo, and T. F. Abdelzaher are now at the University of Illinois, Urbana-Champaign.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0000-0000/20YY/0000-0111 \$5.00

quirements of accurate target tracking and classification [Liu et al. 2003], one of the key design goals of VigilNet is to achieve long-term surveillance in a realistic mission deployment. Due to the small form factor and low-cost requirements, sensor devices such as the XSM motes [Dutta et al. 2005] are normally equipped with limited power sources (e.g., two AA batteries). Moreover, because of the hostile environment and a large number of nodes deployed, currently it is not operationally and economically feasible to replace the power source without introducing enormous effort and elements of risk to the military personnel. In addition, the static nature of the nodes in the field prevents the scavenging of the power from ambient motion or vibration [Paradiso and Starner 2005][Roundy et al. 2006]. The small form factor and possible lack of the line of sight (e.g., deployment in the forest) make it difficult to harvest solar power. On the other hand, a 3~6 month system life span is essential to guarantee the effectiveness of normal military operations, which necessitates a 12~24 fold extension of the normal lifetime of active sensor nodes. Consequently, it is critical to investigate practical approaches of spending the power budget effectively.

Many solutions have been proposed for energy efficiency at various levels of the system architecture, ranging from the hardware design [CrossBow b][Dutta et al. 2005], coverage [Wang et al. 2003][Yan et al. 2003][Sichitiu 2004][Cardei et al. 2005] MAC [Polastre and Culler 2004][van Dam and Langendoen 2003][Ye et al. 2002], routing [Seada et al. 2004][Xu et al. 2001], data dissemination [Agarwal et al. 2004], data gathering [Yu et al. 2004][Choi and Das 2005], data aggregation [Madden et al. 2002][Shrivastava et al. 2004], data caching [Bhattacharya et al. 2003], topology management [Chen et al. 2001], clustering [Heinzelman et al. 2000], placement [Ganesan et al. 2004] [Bogdanov et al. 2004] to energy-aware applications [Szewczyk et al. 2004][Xu et al. 2004]. Instead of focusing on a single protocol, our answer to energy efficiency is an integrated multi-dimensional power management system. Our contributions are identified in the following aspects: 1) Our power management techniques have been fully implemented using the XSM platform. The techniques are used as part of VigilNet – a large-scale target detection and classification sensor network system that has been delivered to military agencies. 2) VigilNet takes a systematic approach, and the energy efficiency is not narrowly accounted for within a single protocol. We propose a novel tripwire service, integrated with an effective sentry and duty cycle scheduling to increase the system lifetime, collaboratively. The tripwire service partitions a network into distinct, non-overlapping sections and allows each section to be scheduled independently. Sentry scheduling selects a subset of nodes, the sentries, which are turned on while the remaining nodes save energy. Duty cycle scheduling allows the active sentries themselves to be turned on and off, further lowering the average power draw. 3) Tradeoffs are investigated to meet requirements of both surveillance performance and the network lifetime. We present a complete system with 40,000 lines of code, running on motes, which achieves performance and energy efficiency simultaneously. 4) We study key system parameters of VigilNet using a theoretical model and 200 XSM motes in an outdoor environment. We evaluate system lifetime using experimental power measurements from the XSM platform as inputs to a discrete event simulations of 10,000 nodes. Our results indicate that the proposed

combination of power management techniques can extend the lifetime of a realistic network from 4 days to 200 days¹.

The remainder of the paper is organized as follows: Section 2 categorizes power management features for different application scenarios. Section 3 describes the power management requirements in VigilNet. Section 4 introduces three power management strategies utilized in VigilNet, namely, the sentry service, the tripwire service and the duty cycle scheduling service. Section 5 describes the integrated power management architecture in VigilNet. Section 6 briefly discusses some additional energy efficient techniques applied in VigilNet. In Section 7, we analyze the target detection performance of Vigilnet through simplified deployment models. Section 8 addresses the tradeoff between energy efficiency and network performance. Section 9 details the VigilNet implementation. Section 10 provides the evaluation of a network of 200 XSM nodes as well as the results of the hybrid simulations of networks containing 10,000 nodes. Section 11 concludes the paper.

2. BACKGROUND

Power management is by no means a stand-alone research issue. It can be dramatically affected by the underlying system configuration and by the application requirements. These include the form factor [Kahn et al. 1999], hardware capability [CrossBow b], possibility of energy scavenging [Roundy et al. 2006][Kar et al. 2005], network/sensing topology and density [Wang et al. 2003], link quality [Keshavarzian et al. 2004], event patterns, node mobility, availability and accuracy of time synchronization [Maroti et al. 2004], real-time requirements and application domain [Szewczyk et al. 2004]. At the hardware level, multi-level sleep modes in the low power microcontroller [CrossBow b] enable software to control the rate of power dissipation. Fine-grained power control [Dutta et al. 2005] allows applications to activate hardware modules incrementally. Radio wakeup circuits [Gu and Stankovic 2004] achieve passive vigilance with a minimal power draw. Energy scavenging [Paradiso and Starner 2005] is also possible for some application scenarios, where ambient energy can be harvested. Sensing coverage schemes [Wang et al. 2003][Yan et al. 2003] exploit redundancy in the node deployment to activate only a subset of the sensor nodes. The coordinated scheduling of the sensor duty cycle [Cao et al. 2005] increases the probability of detection and reduces the detection delay with minimal power consumption. Communication protocols turn off the radio when a node is not the intended receiver [Ye et al. 2002]. Though many individual solutions are proposed, few real systems actually achieve power efficiency comprehensively, which makes the integrated approach in VigilNet novel and practically useful. Considering the diversity of the different approaches, we categorize power management strategies in the context of two types of systems: sampling systems and surveillance systems.

2.1 Power Management in Sampling Systems

Great Duck Island [Szewczyk et al. 2004] and Structural Monitoring [Xu et al. 2004] are typical sampling systems, which are deployed as distributed large-scale data

¹The lifetime is defined as the duration for which the network detects targets with a probability of 90%.

acquisition instruments. Power management strategies in these systems normally make use of the following techniques:

- Predefined sampling schedules:** Most environmental phenomena, such as temperature, exist ubiquitously over space and continuously over time. The static nature of these phenomena makes it sufficient to construct the data profile by sampling the environment within discrete time and space. Nodes can conserve energy by turning themselves off and on, according to a predefined schedule.
- Synchronized and coordinated operations:** Once the sampling interval is defined *a priori*, nodes can communicate in a synchronized fashion. With a precise time synchronization [Maroti et al. 2004], a receiver can turn on the radio module right before the message payload arrives. Consequently, we can avoid low-power listening over radio [Polastre and Culler 2004] during a non-active period. In addition, with the knowledge about the sending rate of individual nodes, we are able to estimate the radio link quality without control messages [Woo et al. 2003].
- Data aggregation and compression:** Since channel media access is costly, especially when the receiver is in a deep-sleep state [Polastre and Culler 2004], it is beneficial to send out one aggregate containing multiple sensor readings [Madden et al. 2002] [Shrivastava et al. 2004]. In addition, due to the locality of the sensed data, we can compress the total number of bits to be sent over the air. Since both aggregation and compression need to buffer a relatively large number of readings, which introduces a certain delay, they are not quite suitable for time-critical surveillance systems. However, they match most sampling systems very well.

2.2 Power Management in Surveillance System

On the other hand, operations in surveillance systems [Liu et al. 2003][Arora et al. 2003][Simon and et. al. 2004] [He et al. 2004], such as VigilNet, are event-driven in nature. In surveillance systems, we are more interested in the data profile between inception and conclusion of the transient events. These systems should remain dormant in the absence of the events of interest, and switch to an active state to obtain high fidelity in detection. Normally, the surveillance systems improve the system lifetime through the following approaches:

- Coverage control:** Surveillance systems are normally deployed with a high density (For instance, the default configuration of VigilNet [He et al. 2004] has 28 nodes per nominal radio range (30 m)) for the sake of robustness in detection and fine-grained sensing during tracking. We can increase the system lifetime by activating only a subset of nodes at a given point of time, waiting for potential targets.
- Duty cycle scheduling:** The duration of transient events within the area of surveillance is normally non-negligible. By coordinating nodes' sleep schedules, we can conserve energy without noticeably reducing the chance of detection. Duty cycle scheduling in surveillance systems is significantly different from sample scheduling in sampling systems. Indeed, the goal of sample scheduling is to get environmental information at a known rate. A sampling system can acquire a

sample at the required time and save power until the next sampling. By contrast, the goal of the surveillance system is typically to detect an object. The system can sense continuously to avoid missing targets. Alternatively, for long lasting events, the system can use duty cycle scheduling, which alternates periods of energy preservation with periods of continuous sensing. Note that the problem of coordinating nodes sleep schedules and maintaining appropriate coverage are tightly related as exposed in [Kumar et al. 2004].

- **Incremental activation:** The sampling systems are normally designed for data logging. At each sample instance, all sensors should be activated to obtain a complete data profile. In contrast, surveillance systems are designed to detect transient events of interest. It is sufficient to activate only a subset of sensors for the initial detection. After the initial detection, we can activate additional sensors to achieve a higher sensing fidelity and to perform target classification.

3. POWER MANAGEMENT REQUIREMENTS IN VIGILNET

Our power management strategies are motivated by a typical military surveillance application. The mission objective of such a system is to conduct remote, persistent, clandestine surveillance to a certain geographic region to acquire and verify enemy capabilities and transmit summarized intelligence worldwide in a near-real time manner. Several system requirements affect the power management design within VigilNet:

- **Continuous surveillance:** Due to the dynamic/transient nature of the event of interest, VigilNet is required to provide continuous surveillance. This requirement significantly affects the overall architecture of power management strategies and the degree of energy conservation that VigilNet can achieve.
- **Real-time:** As a real-time online system for target tracking, VigilNet is required to cope with rapidly moving targets in a responsive manner. The delays introduced by the power management directly affect the maximum target speed our VigilNet can track. It is an essential design tradeoff to balance between network longevity and responsiveness.
- **Rare and critical event detection:** Due to the nature of military surveillance, VigilNet deals with the rare event model. In this model, the total duration of events is small, compared to the overall system lifetime. On the other hand, events are so critical that the power management becomes a secondary consideration in the presence of events.
- **Stealthiness:** Deployed in hostile environments, it is vital for VigilNet to have a very low profile. Miniaturization makes nodes hard to detect, physically. However, radio messages can be easily intercepted if nodes frequently communicate. Power management protocols designed for VigilNet should maintain silence during surveillance in the absence of significant events.
- **Flexibility:** We envision the deployment of VigilNet using different densities, topologies, sensing capabilities and communication capabilities. Therefore, it is essential to design a power management architecture that is flexible enough to accommodate various system scenarios.

4. KEY POWER MANAGEMENT STRATEGIES IN VIGILNET

In order to achieve long-term surveillance that meets the military requirement (e.g., 3~6 months), an aggressive 12~24 fold life-time extension is essential. Our initial investigation [He et al. 2004] indicates that a single power management strategy is neither sufficient nor flexible. Therefore we restructure our prototype system described in [He et al. 2004] by adding a new combination of tripwire service and duty cycle scheduling. We believe this is the right direction to pursue. In this section, we detail three main strategies, namely the tripwire service, sentry service and duty cycle scheduling, before presenting the VigilNet architecture in the next section. In order to support these strategies, all nodes within VigilNet find their positions with an accuracy of 1~2 meters and they synchronize with each other within 1~10 milliseconds using the techniques described in [Stoleru et al. 2004] and [Maroti et al. 2004], respectively. Long-range communication devices are deployed as bases to relay sensor reports outside of the sensor field.

4.1 Tripwire Service

This section proposes a novel network-wide power management strategy called *Tripwire Service*. This service divides the sensor field into multiple sections, called *tripwire sections*, and applies different working schedules to each tripwire section. A tripwire section can be either in an active or a dormant state, at a given point of time. When a tripwire section is dormant, all nodes within this section are put into a deep-sleep state to save energy. Surveillance in active tripwire sections can be done by either turning all nodes on or applying coverage algorithms such as the sentry service discussed later in Section 4.2. The rationale behind the tripwire service is as follows. First, the network is divided into several tripwire sections for scalability purposes: each base station communicates with a limited number of nodes, which reduces congestion problems. Second, the network is divided into several tripwire sections for reliability purposes: the network is still functional even if some of the base stations fail. Third, as far as energy conservation is concerned, we observe that we can turn on and off various tripwire sections to save energy following a pattern depending on the characteristics of the targets. Consider the following example. Imagine that a target vehicle passes through a sensor field by following a road. We can divide the road into several sections. We can then turn on the sensors in one section and turn off the sensors in the other sections. As the target follows the road, it will necessarily pass through the section of the road where the sensors are on and be detected. Periodically, we can then turn on another section and turn off all the others. We are then able to balance energy consumption between the various sections and, at the same time, to still detect most targets.

4.1.1 Tripwire partition. The tripwire partition policy of VigilNet is based on Voronoi diagrams. A network with n bases is partitioned into n tripwire sections such that each tripwire section contains exactly one base i and every node in that tripwire section is closer to its base i than to any other base inside the sensor field. Every node in the network uniquely belongs to one and only one tripwire section. The rationale behind Voronoi partitioning is to reduce the energy consumption and the end-to-end delay in data delivery.

The positions of bases directly determine the layout of tripwire sections and affect

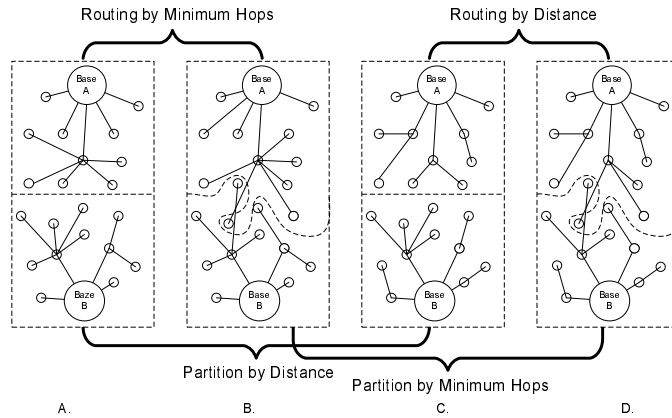


Fig. 1. Four different ways of implementing the tripwire service. For routing purposes, the tripwire service can either use the route with the minimum number of hops, or the route with smallest physical distance. The tripwire partitions can be determined using either the distance to the base or the minimum number of hops to the base.

the number of communication hops necessary for node-base communications. The optimal base placement method to minimize the average path length to the nearest base can be found in [Okabe et al. 2000]. In practice, the base placement strategy is normally determined by the mission plan and topology.

4.1.2 Tripwire partition mechanism. This section describes the mechanism to enforce the tripwire partition policy. At the beginning of the tripwire partition operation, each base broadcasts one initialization beacon, to its neighbors with a hop-count parameter initialized to one. Link Symmetry Detection [Zhou et al. 2004] is used to ensure beacons can only be received through high quality symmetric links. Each receiving node maintains the minimum hop-count value of all beacons it received from the nearest base, in terms of the physical distance, and ignores beacons with higher hop-count values and beacons from other bases. Beacons are flooded outward with hop-count values incremented at every intermediate hop. Through this mechanism, all nodes in the network get the shortest high quality path, in *hops*, to the nearest base, in *physical distance*. While the above mechanism is intuitive, the design deserves some further clarification. First, the boundaries between partitions are well delimited if we partition the network according to the physical distance between sensor nodes and bases (Figure 1A and 1C). If the communication hop is used instead, the radio irregularity and the interference cause partitions to interleave with each other (Figure 1B and 1D). This brings complexity and uncertainty to the design of optimal tripwire placement strategies. Note that, because we use the physical distance to partition the network, physical occlusions do not affect the partitioning of the network except if it prevents a node to communicate with its closest base in terms of physical distance (through a multi-hop communication path).

Second, it is beneficial to use hop counts to build diffusion trees within each

partition, because 1) the normal geographic-based routing does not guarantee high-quality shortest path to the root. 2) Due to the existence of high-quality long links, a smaller number of nodes become active backbone nodes in the hop-based routing than in the geographic-based routing. Finally, this design provides certain robustness in case of base failure. If a base fails, the sensor field can be easily repartitioned without this base.

4.1.3 Tripwire Duty Cycle. A tripwire section can be either in an active or a dormant state. The state of a tripwire section can change during what we call a system rotation. A tripwire schedule determines which tripwire sections are awake during which rotations. The tripwire schedule can be specified manually before deployment or it can be determined randomly. The tripwire schedule is stored on the base stations. During a system rotation, the tripwire bases configure the state of the tripwire nodes according to the tripwire schedule. The Tripwire Duty Cycle (TDC) is the percentage of time for which a given tripwire section is active. For instance, if a tripwire is active during 5 of 10 rotations, we say that the TDC equals 50%.

4.2 Sentry Services

In order to exploit the high node density within the sections, we design and implement a section-wide power management strategy, called *sentry service*. The main purpose of the sentry service is to select the subset of nodes, which we name the *sentries*, that is in charge of the surveillance. Sentry selection contains two phases. Nodes first exchange neighboring information through hello messages. In each hello message, a sender attaches its node ID, position, number of neighbors and its own energy readings. After the first phase, each node builds up a one-hop neighbor table. In the second phase, each node sets a delay timer. The duration of the timer is calculated based on the weighted energy rank R_{energy} and the weighted cover rank R_{cover} , as shown in Equation 1. The energy rank R_{energy} is assigned according to energy readings among neighboring nodes (e.g., the node with the highest energy reading within a neighborhood has a rank of 1). Similarly, the cover rank R_{cover} is assigned according to the number of neighbors within a the node's sensing range. As for the current implementation, we assign equal weights to both ranks.

$$T_{timer} = \frac{W_e \times R_{energy} + W_c \times R_{cover}}{(W_e + W_c) \times \#Neighbors} MaxDelay + Jitter \quad (1)$$

After the delay timer fires in one node, this node announces itself as *sentry* by sending out a declaration message. While other nodes, in the vicinity of the declaring node, cancel their timers and become dormant *non-sentry* nodes. The effective range, in physical distance, of a sentry's declaration message is named the range of vicinity (ROV). While the sentry selection can be straightforwardly implemented, the challenging part is to *choose* and to *enforce* the appropriate range of vicinity (ROV). This parameter directly affects the sentry density, and hence affects the lifetime of the network. In Section 7, we develop a theoretical model that we use in Section 7.5.1 to select an appropriate value for the ROV parameter.

4.2.1 How to enforce ROV. After we choose a ROV value, we need to enforce it during the sentry selection phase. Since the sensing range is normally smaller than

the radio range, directly using the radio range as the ROV cannot guarantee an effective coverage of the area. For example, the HMC1002 dual-axis magnetometer used by MICA2 has only 30-foot effective range for a moving car. If we use the Chipcon radio (>100 feet) to define the range of vicinity, less than 10% of area is covered by the sensors. There are two approaches to address these issues. The first approach is to reduce the radio emission power to emulate the ROV range. The power setting can be chosen in such a way that there is about one sentry within each sensing range. The second approach is to discard declaration messages from any sentry beyond the distance of ROV. The first approach achieves sensing coverage, without the location information of the nodes [He et al. 2003], while the second approach provides a more predictable sentry distribution, because the emulated ROV would be affected by the radio irregularity in the environment. Consequently, we adopt the second solution in our system, given the fact that localization [Stoleru et al. 2004] is supported in VigilNet. Note that in real environment, if some field areas do not contain any sensor nodes, it is then impossible to achieve full sensing coverage.

4.3 Sentry Duty Cycle Scheduling

The requirement for continuous sensing coverage in the sentry service imposes a theoretical upper bound on the system lifetime. This upper bound is decided by the total number of nodes deployed. Since a target normally stays in the sensing area of a sentry node for a non-negligible period of time, it is not necessary to turn sentry nodes on all the time. By using duty cycle scheduling, we are able to break the theoretical upperbound imposed by the full coverage algorithms [Yan et al. 2003]. Let T_{on} be the active duration and T_{off} be the inactive duration, then the Sentry Toggle Period (STP) is defined as $(T_{on} + T_{off})$, and the Sentry Duty Cycle (SDC) is defined as $\frac{T_{on}}{STP}$. Theoretically, the duty cycle scheduling can achieve unbounded energy conservation by lowering the SDC value. The paramount concern of this technique is that lowering the SDC value increases the detection delay and reduces the detection probability. We can either effectively implement random duty cycle scheduling or more sophisticated scheduling algorithms to coordinate node activities to maximize performance. In [Cao et al. 2005], we demonstrate a local optimal scheduling coordination algorithm to reduce the detection delay and increase the detection probability. We prove that, at relatively large SDC (e.g. $5\% < SDC$), the difference between random scheduling and optimal scheduling can be practically ignored. Since the random scheduling does not need control messages for coordination (more stealthy), and is not affected by time drift, we choose random scheduling over the coordinated one for the VigilNet implementation.

5. INTEGRATED SOLUTION: TRIPWIRE-BASED POWER MANAGEMENT WITH SENTRY SCHEDULING

To achieve an aggressive network lifetime extension, the VigilNet power management subsystem integrates the three strategies mentioned in previous sections into a multi-level architecture, as shown in Figure 2. At the top level, the tripwire service controls the network-wide distribution of power consumption among sections. The uniform discharge of energy across sections is achieved through the scheduling mechanism we discussed in Section 4.1.3. We use a Tripwire Duty Cycle (TDC),

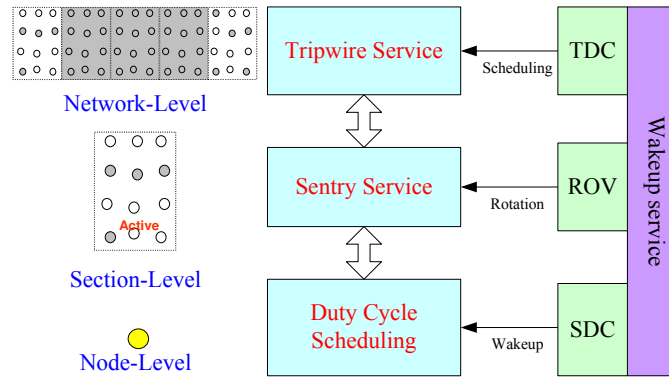


Fig. 2. Integrated power management architecture. The architecture integrates three power management strategies: tripwire service, sentry service, and duty cycle scheduling.

which is the percentage of active time for each tripwire section, to control the network-wide energy consumption rate. There are two special cases: when TDC equals 100%, the whole network becomes active and the tripwire service is merely a network partition service. When TDC equals 0%, the whole network is in dormant status and it can only be awoken by external sources. At the second level, the sentry service controls the power distribution within each section. The uniform discharge of energy in a section is achieved through automatic rotation strategies according to the remaining power within individual nodes. We use the Range of Vicinity (ROV) parameter to control the energy-burning rate of active sections. When ROV equals 0 meter, the sentry service is actually disabled and all nodes within the section are awake, providing the highest degree of coverage. At the third-level, duty cycle scheduling controls the energy consumption rate of individual sentry nodes by manipulating their wakeup/sleep schedule. The Sentry Duty Cycle (SDC) parameter, which quantifies the percentage of active time, is used to control the awareness of sentry nodes. The duty cycle scheduling can be disabled by setting SDC to 100%. By adopting different values for TDC, ROV and SDC, we can flexibly adjust our power management to accommodate different system scenarios.

6. OTHER ENERGY CONSERVATION TECHNIQUES

Besides the three main power management strategies, several other techniques have been integrated into various aspects of the VigilNet system. Similar techniques [Shrivastava et al. 2004][He et al. 2004][Xu et al. 2004][Polastre and Culler 2004] have been proposed in the literature and we provide this section for the completeness of the description of the VigilNet power management design and implementation.

- **Minimum connected dominating tree:** To ensure a swift delivery of messages, VigilNet requires an active diffusion tree over any active tripwire section. Since the communication range is normally much larger than the sensing range [CrossBow b][Dutta et al. 2005], it is possible to build a diffusion tree on

top of the sentry nodes. To reduce the energy spent during idle listening, VigilNet needs a tree with the minimum connected dominating set (a tree with minimum non-leaf nodes). Since it is a NP-Complete problem to find the minimum connected dominating set of a graph, we adopt a localized approximation as follows: during the building process, each node rebroadcasts the hop-count beacon after a certain time delay. The delay in one node is inversely proportional to the number of neighbors and the energy remaining. By doing so, the node with more neighbors and more energy left has a higher chance to become the parent node within the diffusion tree.

- Data aggregation:** The channel media access in wireless sensor network is relatively expensive. For example, in the Chipcon radio implementation for MICA2, to deliver a default payload size of 29 bytes, the total overhead is 17 bytes (37%!), including 8 bytes preamble, 2 bytes synchronization, 5 bytes header and 2 bytes CRC. This motivates us to utilize various kinds of aggregation techniques. The first technique we use is called Application-Independent Aggregation, which concatenate data from different modules into one aggregate, regardless of their semantics. For example, system-wide parameters can be sent with time synchronization messages. The second technique we use is called Application-Dependent Aggregation. The tracking subsystem in VigilNet performs in-network aggregation by organizing the nodes into groups. Instead of each node reporting its position separately, a leader node calculates the weighted center of gravity from multiple inputs and reports only one aggregate back to the base.
- Implicit acknowledgement:** Given that the sensor payload is very small, it might not be energy efficient to acknowledge every packet explicitly. Implicit acknowledgement can be achieved through several approaches. They differ in functionality and overhead. B-MAC [Polastre and Culler 2004] provides an efficient implementation of the CSMA protocol with radio-layer acknowledgement support. Observing that most of the packets need to be forwarded for routing, we alternatively implemented the acknowledgement as a special field in outgoing packets. When there are no outgoing packets for a period of time, a special acknowledgement packet is sent.
- Incremental detection:** Multi-sensing modalities are desired for achieving target classification. However, it is not necessary to activate all sensors only for detection. Among the three types of sensors in XSM motes, the optic TR230 PIR sensor has the longest detection range and a relatively low power consumption of 0.88mW. We use this sensor to support the initial detection and to incrementally wakeup other sensors for classification purposes.
- Passive wakeup circuitry :** Several efforts [Dutta et al. 2005][Goldberg et al. 2004][Gu and Stankovic 2004] have been made to support low-power passive wakeup by using an acoustic detector [Goldberg et al. 2004], an infrared sensor [Dutta et al. 2005] or a radio [Gu and Stankovic 2004]. Currently, the design [Dutta et al. 2005] of XSM motes is not mature enough for VigilNet to exploit this technology, but it is a very promising direction.

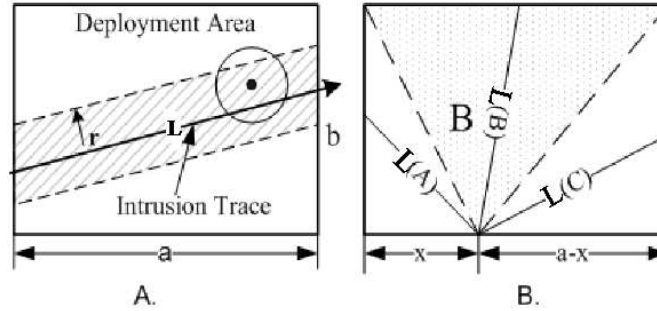


Fig. 3. Intrusion Model

7. ANALYSIS OF DETECTION PROBABILITY

7.1 Purpose

This section analyzes target detection performance of Vigilnet through simplified deployment models. This analysis serves two purposes. First, it can be used as a “feasibility check” before real systems, i.e., Vigilnet, are deployed. In most cases, real deployment can be both costly and time-consuming. While simulations can usually be an effective alternative, they cannot provide enough insight on the impact of different parameters, and tend to provide obscure results if such parameters are poorly chosen. Second, theoretical results make it possible to quantitatively analyze system behavior, and quickly provide insight on the effect of different parameter combinations. For these reasons, we consider it beneficial to include a section that is dedicated to the analysis of detection probability.

7.2 Preliminaries

To derive our analytical model, we first outline our assumptions. We consider a rectangular deployment area with side lengths of a and b . We assume that N sensor nodes are uniformly deployed in the area, the node density d is N/ab . Assuming that the area is considerably large, therefore, the number of nodes in an area of A ($A \ll ab$) can be approximated by a Poisson distribution with parameter $\lambda = dA$. We assume that the entry point of the intruding target (intruder) is uniformly distributed along all sides, and to make the problem tractable, we assume that the intruder moves along a straight line. The angle between the target direction and the side where the entry point is located is θ , which is also considered uniformly distributed, i.e., $\theta \in [0, \pi]$. The whole intruding scenario is shown in Figure 3A.

Observe that the area of nodes that can detect the intruder contains all points whose distances to the intruder’s locus are no larger than the sensing range r . If the length of the intruder’s locus in the deployment area is L , the detection area can be approximated by $2Lr$, without considering the edge effect. We now use a general result from theory of probability. If the probability of an event A occurring in a single experiment is p , and if the number of experiments conforms to a Poisson Distribution with parameter λ , the probability of event A occurring at least once in the series of experiments is:

$$P = 1 - e^{-p\lambda} \quad (2)$$

Based on this result, we know that for the detection area $2Lr$, the probability that at least one node is located in this area is $1 - e^{-2Lrd}$, where d is the density of awake sentries. We can now calculate the probability that an intruder is detected by at least one node by integrating over all entry points on the two adjacent sides of the area. We next consider two cases of the problem.

7.3 Probability of Intruder Detection using Constantly Awake Sentry Nodes

In this section, we consider the first case where all nodes are sentries that are constantly awake. As shown in Figure 3B, the deployment area is divided into three regions. The length of the intrusion trace is:

$$L(\theta, x) = \begin{cases} x/\cos\theta & \text{Locus} \in A \\ b/\sin\theta & \text{Locus} \in B \\ (x-a)/\cos\theta & \text{Locus} \in C \end{cases} \quad (3)$$

We can then calculate the expected detection probability of a randomized intruding target by integrating over all entry points and all incoming directions. Therefore, we have that:

$$Expected(P_{detection}) = 1 - \frac{F(a, b, r, d) + F(b, a, r, d)}{\pi(a + b)} \quad (4)$$

where $F(m, n, r, d) =$

$$\int_0^m \left[\int_0^{\arctan(\frac{n}{x})} e^{-\frac{2rx}{\cos\theta}} d\theta + \int_{\arctan(\frac{n}{x})}^{\pi - \arctan(\frac{n}{m-x})} e^{-\frac{2rn}{\sin\theta}} d\theta + \int_{\pi - \arctan(\frac{n}{m-x})}^{\pi} e^{-\frac{2r(x-m)}{\cos\theta}} d\theta \right] dx \quad (5)$$

One comment on this result is that this integral does not lead to a closed form. Therefore, it can be solved only numerically.

7.4 Probability of Intruder Detection under Sentry Duty Cycle Scheduling

In more general cases, it is not necessarily true that all nodes are awake at all times. To save energy, it is usually the case that some sentries employ certain duty cycle scheduling policies. We consider this case in this section.

We first present the duty scheduling model. The duty cycle of node is shown in Figure 4, which, at a random time point, has a probability of β of being awake. We also assume that this node follows a cycle T in the scheduling.

Our following analysis is based on our previous work [Cao et al. 2005], which analyzed detection delay distributions in several cases. We first briefly outline the results from [Cao et al. 2005], and then apply them in our next step analysis.

In our work of [Cao et al. 2005], we presented analysis for four types of target detection scenarios. These scenarios are shown in Figure 5. In the model assumed by Vigilnet deployment, we are considering Type I and Type III: on one hand, we are interested in the detection probability of a target when it passes the area covered by Vigilnet, i.e., Type I detection; on the other hand, we are interested in the expected detection delay *if the area is large enough*, i.e., Type III detection.

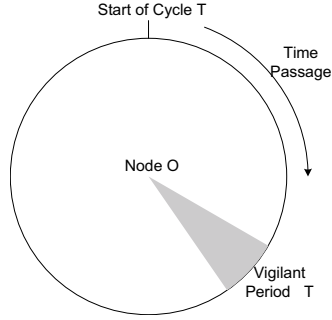


Fig. 4. Duty Cycle of a Node

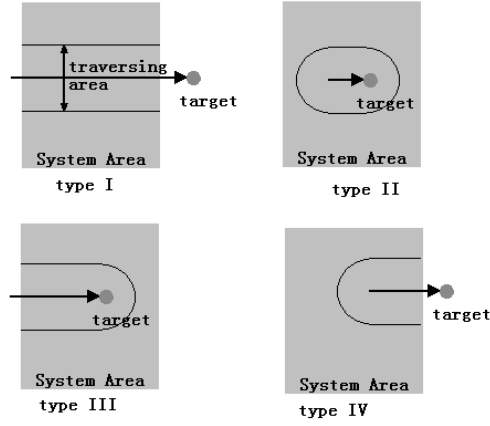


Fig. 5. Target Detection Scenarios

We now briefly outline the analysis results from [Cao et al. 2005]. The analysis classifies moving targets into two categories: *fast* and *slow*, based on its speed. Formally, [Cao et al. 2005] pointed out that the boundary speed v_0 between fast targets and slow targets as:

$$v_0 = \frac{2r}{(1 - \beta)T} \quad (6)$$

Those targets that are moving faster than v_0 are considered *fast*. Otherwise, they are considered *slow*. Fast and slow targets lead to different analysis results.

For type I fast target, with a deployment width of L , it has a detection probability as follows:

$$P_{detection}(v) = 1 - e^{-2rLdP} = 1 - e^{-2rLd(\beta + \frac{\pi r}{2vT})} \quad (7)$$

while for type I slow targets, the result is:

$$P_{detection}(v) = 1 - e^{-2rLdP} = 1 - e^{-2rLd(\beta + \frac{\pi r^2 + k(r, \alpha)}{2rvT})} \quad (8)$$

where

$$k(r, a) = 2a\sqrt{r^2 - a^2} - 2r^2 \cos^{-1}\left(\frac{a}{r}\right) \quad (9)$$

For type III fast targets, the expected detection delay is:

$$E(T_d) = \frac{e^{-\beta\pi r^2 d/2}}{(2r\beta v + \frac{\pi r^2}{T})d} \quad (10)$$

For slow type III targets, the expected detection delay is:

$$E(T_d) = \frac{e^{-\beta\pi r^2 d/2}}{(2r\beta v + \frac{\pi r^2}{T})d} \left[1 - \frac{m(r, \beta) e^{-(2r\beta v T + \pi r^2)(1-\beta)d/2}}{2r\beta v T + \pi r^2 + m(r, \beta)} \right] \quad (11)$$

We now calculate the expected detection probability based on the integration of L over all potential entry points and incoming directions. More specifically, we have: $Expected(P_{detection}(v)) =$

$$\frac{\int_0^\pi d\theta \int_0^a dx P(\theta, x) + \int_0^\pi d\theta \int_0^b dx P(\theta, x)}{(a+b)\pi} \quad (12)$$

where for fast targets,

$$P(\theta, x) = 1 - e^{-2rd(\beta + \frac{\pi r}{2vT})L(\theta, x)} \quad (13)$$

and for slow targets,

$$P(\theta, x) = 1 - e^{-2rL(\theta, x)d(\beta + \frac{\pi r^2 + k(r, a)}{2rvT})} \quad (14)$$

Since the symbolic integral result is not available for Equation 12, we use numeric integration to obtain the detection probabilities in the following section.

7.5 Applications of Detection Probability Model

7.5.1 Choosing ROV. The appropriate ROV value can be chosen using our analytical intrusion detection model. This model describes the relationship between the detection probability, the sensing range and the sentry density. Since, theoretically, there is at most one sentry within each ROV range, according to the circle covering theorem [Williams 1979], the sentry density is upper bounded by $\frac{2\pi}{\sqrt{27}ROV^2}$. Given the area size, sensing range and sentry density, we can get the detection probability, as shown in Figures 8. For a typical deployment with 1000 nodes in $100 \times 1000 \text{ m}^2$ area, Figure 8 indicates how to choose the right combination of system parameters. For example, in order to achieve a 99% detection probability, we can choose either a sentry density of 0.008 nodes/ m^2 (ROV= 6 meters) with 8 meter sensing range or a lower density of 0.004 nodes/ m^2 (ROV=8.5 meters) with 14 meters sensing range. Note that sentries need to be able to establish a routing tree to the base in order to report target detection events. This may be a problem in networks with a very low density of nodes.

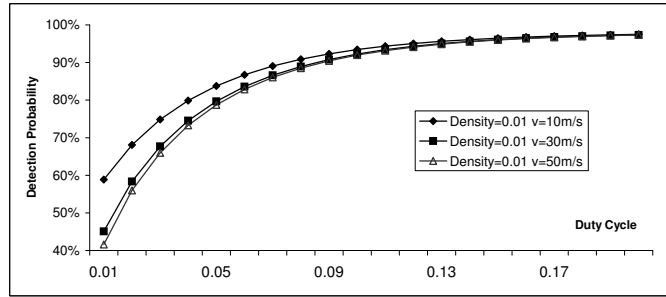


Fig. 6. Effect of sentry duty cycle (SDC, in % / 100) on detection probability according to node density (in number of nodes per m^2) and target velocity. This graph is generated using a theoretical model. A high detection probability can be achieved even with a relatively low duty cycle.

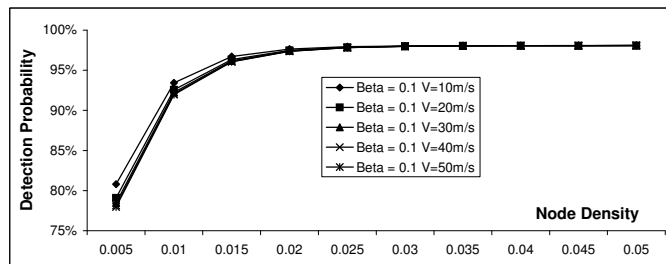


Fig. 7. Effect of sentry density (in number of nodes per m^2) on detection probability, according to wakening period ratio β and target velocity. This graph is generated using a theoretical model. The target velocity does not noticeably impact the detection probability when node density is reasonably high.

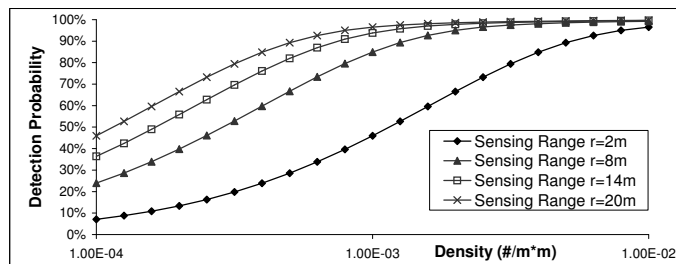


Fig. 8. Detection probability versus sentry density. This graph is generated using a theoretical model and indicates how to choose the density of the sensor network according to the sensing range to achieve a desired detection probability.

7.5.2 Impact of Duty Cycle Scheduling. We use the theoretical model developed in the previous sections to investigate the impact of the duty cycle scheduling. For a typical deployment in $100 \times 1000 \text{ m}^2$ area with a 10m sensing range, we analyze the effect of different factors on the detection probability. We focus on three parameters: the sentry density, target velocity and duty cycle percentage. Two interesting results are described in the following paragraph.

First, Figure 6 shows that a high detection probability (99%) can be achieved with a relatively low duty cycle (19%). In other words, 81% of the duty cycle can be saved without significantly impacting the probability of detection! Second, Figure 7 shows that the velocity of targets does not noticeably impact the detection probability when node density is reasonably high.

We emphasize here that the analytical results only give us a lower bound of the duty cycle. Due to sensor warm-up and calibration issues, the required SDC should be higher to achieve the same performance in reality.

8. TRADEOFF: PERFORMANCE VS. ENERGY EFFICIENCY

One key research challenge for VigilNet is to reconcile the need for network longevity with the need for fast and accurate target detection and classification. The former requires most sensor nodes to remain inactive, while the later requires many active sensor nodes. As we mentioned before, the event model directly affects the design of the power management. Energy efficiency can be comparatively easy to achieve if events of interests are ubiquitously present. The data quality of some events, such as events related to temperature sampling, is not directly correlated with the responsiveness of the system. However, in a surveillance system, *responsiveness* and *awareness* directly affect system performance, which includes the tracking performance and the target classification performance. The former can be measured in terms of detection probability and delay, and the later can be measured in terms of the number of nodes detecting external events simultaneously. We have investigated responsiveness in previous sections 4.2 and 4.3. This section focuses on how to improve the system awareness. In VigilNet, awareness is supported by the *on-demand wakeup service*. The on-demand control is stealthier compared to the periodic control [He et al. 2004], because wakeup beacons are sent only when events occur. To support the on-demand control, we need to guarantee the delivery of wakeup beacons. Because of the particular stealthiness requirement, the non-sentries cannot synchronize their clocks with their sentries by exchanging messages. Therefore, neighboring non-sentry sensor nodes can no longer have a sleep-wakeup cycle synchronized with each other due to their clock drift, and a sentry cannot keep track of which of its neighbors are awake. To guarantee delivery, a non-sentry periodically wakes up and checks radio activity (detects preamble bytes) once per checking period (e.g., every second). If no radio activity is detected, this node goes back to sleep, otherwise it remains active for a period of time, preparing for incoming targets. If a sentry node wants to wake up all neighboring nodes, it only needs to send out a message with a long preamble with a length equal to or longer than the checking period of non-sentry nodes. Since in the rare event model, the wakeup operations are done very infrequently, the long preamble doesn't introduce much energy consumption in sentry nodes. On the other hand, since the amount of time

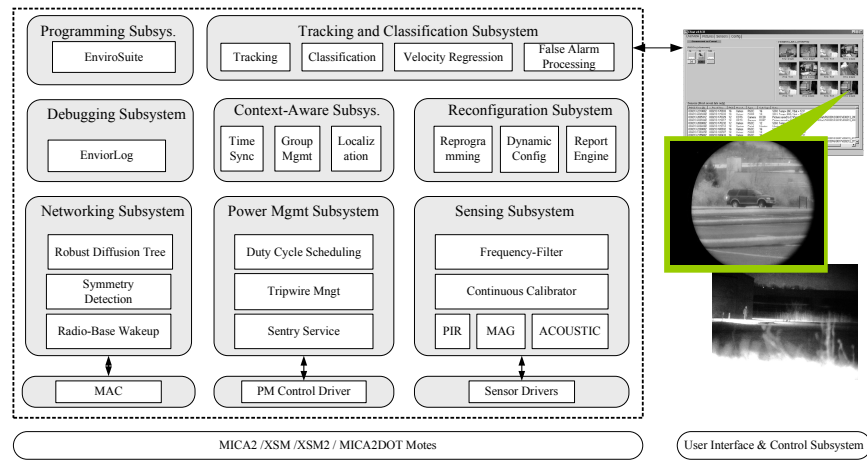


Fig. 9. The VigilNet system architecture. The three power management services described in this paper can run on a XSM node along with additional power management techniques (e.g., robust diffusion tree), context-aware services (time synchronization, group management, localization), reconfiguration subsystems (reprogramming, reconfiguration, reporting), and tracking and classification subsystems.

taken to check the radio activity is constant for a specific radio hardware, the length of the checking period determines the energy consumption in non-sentry nodes. In general, a long checking period leads to a lower energy consumption. However, to ensure that a sentry node wakes up neighboring non-sentry nodes before a target moves out of their sensing range, the checking period can not be arbitrarily long. Theoretically, the upperbound of checking period is $\frac{\sqrt{R^2-r^2}}{S}$, where R is radio range, r is sensing range of sentries and S is the speed of target. Due to the other delays, such as sensor warm-up time, the checking period should be smaller than this theoretical bound. In our implementation, non-sentry nodes have 1% duty cycle with 1 second checking period.

9. IMPLEMENTATION

The power management architecture described in Section 5 has been integrated into the VigilNet system. We have successfully transferred VigilNet to a military agency for deployment by the end of 2004. The overarching architecture of VigilNet is shown in Figure 9. The three power management techniques presented in this paper form the power management subsystem. Note that every component of the VigilNet system has been designed with power management in mind. Techniques employed in VigilNet components include those presented in Section 6.

VigilNet is built on top of the TinyOS operating system. TinyOS supports a lightweight event driven computation model with two-level scheduling. VigilNet is mostly written in NesC, a language derived from C that is especially designed for embedded programming. The VigilNet software is composed of about 40,000 lines of code, and supports multiple existing mote platforms including the MICA2 sensor node and the XSM sensor node. The compiled image occupies 83,963 bytes

of code memory and 3,586 bytes of data memory. In particular, the power management techniques (sentry service, duty cycle service, and tripwire service) occupy 6,472 bytes of code memory and 202 bytes of data memory. The small memory footprint of our power management techniques is one of their main advantages. Indeed, as VigilNet is a real application, it is essential that the power management techniques must be efficient enough to extend the lifetime of the network and have a small enough memory footprint to fit on XSM sensor nodes, which only have 4 kilobytes of RAM. Because of VigilNet requirements, not only power management techniques must use no more than the RAM and ROM available on the XSM sensor nodes, but also the power management techniques must leave a large part of the memory available to other VigilNet modules such as the networking subsystem, the target classification subsystem, the debugging subsystem, and the reconfiguration subsystem.

In Figure 9, the components labeled MAC, Power Management Control Driver, and Sensor Drivers are standard components found in the TinyOS operating system. The VigilNet software is organized into the following subsystems:

- The networking subsystem consists of the following components: a robust diffusion tree for routing data to and from the system base station, a symmetry detection protocol to limit message loss, and a radio-based wakeup system that is key to the power management strategies presented in this paper.
- The power management subsystem implements the VigilNet power management strategies: duty cycle scheduling, tripwire management, and sentry service (see Section 4).
- The sensing subsystem provides calibration and filtering services for the motion sensor (PIR), the magnetic sensor (MAG) and the acoustic sensor. We note that minimizing false alarms is critical in an outdoor environment. Indeed, false alarms generate unnecessary wakeup operations that reduce network lifetime.
- The debugging subsystem includes a tool called EnviroLog that logs sensor data into the Flash memory of the sensor nodes. The logged data is then replayed by the nodes to achieve repeatability while experimentally testing VigilNet.
- The context-aware subsystem includes time synchronization, group management, and localization services. Some of these services are essential to the functioning of the VigilNet power management strategies. For instance, the localization service provides information necessary for the tripwire partition and the sentry selection.
- The reconfiguration subsystem includes reprogramming, dynamic configuration, and reporting services. It basically establishes the communication between tripwire base stations and sensor nodes.
- The programming subsystem includes a suite of tools for facilitating the programming of sensor networks.
- The tracking and classification subsystems correlate data received from the sensor nodes to infer the trajectory of the tracked targets, their speed, and their type (human without weapon, human with weapon, car-sized vehicle, truck-sized vehicle).



Fig. 10. Location of deployment and a deployed XSM mote. Nodes are randomly placed roughly 10 meters apart, covering a 300 meter road that intersects a 200 meter road.

10. SYSTEM EVALUATION

This section presents experimental results that evaluate the performance of the power management subsystem. The experimental results in Section 10.1 are obtained through an actual deployment of 200 XSM motes, focusing on the sentry selection, tripwire partition and tracking delays. Other experiments in Section 10.2, especially those related to the system lifetime, require a significant amount of time. Unfortunately, we currently can not afford to deploy such a large system unattended for a long time. We have to conduct those evaluations through a hybrid approach, which uses basic measurements from a smaller number of motes as input to a simulator. By doing so, we can investigate the impact of different system configurations on the performance of power management.

10.1 Field Evaluation

The field evaluation was done as part of a technical transition on December 2004, when we deployed 200 XSM motes on a dirt T-shape road (200 meters by 300 meters). The XSM mote is designed by the joint efforts of Ohio State University [Dutta et al. 2005] and CrossBow Inc, which features an Atmel ATmega128L microcontroller and a Chipcon 433MHz CC1000 radio. Its sensing suite includes magnetic, acoustic, photo, temperature and passive infrared sensors (PIR). Figure 10 displays the environment where our system was located and the picture of one of the XSM motes. Nodes are randomly placed roughly 10 meters apart, covering one 300-meter road and one 200-meter road.

10.1.1 *Effectiveness of the tripwire partition.* One snapshot of the network layout collected by our graphical user interface is shown in Figure 11. We placed 200 XSM sensor nodes and 3 mica2dot base nodes in the field. Accordingly, the network

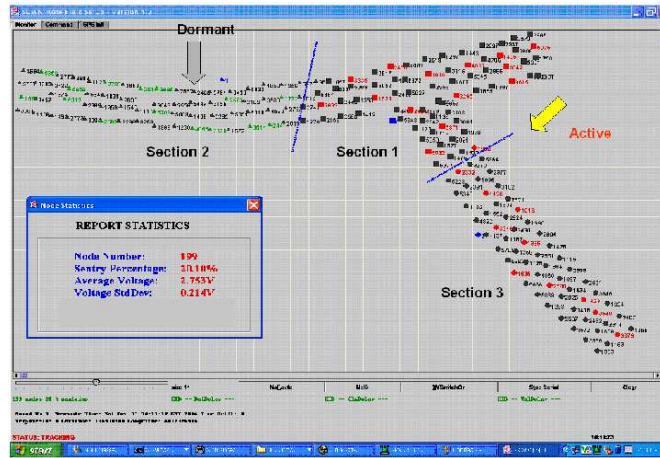


Fig. 11. Effectiveness of tripwire partition. All nodes attach to their nearest base node through their shortest path.

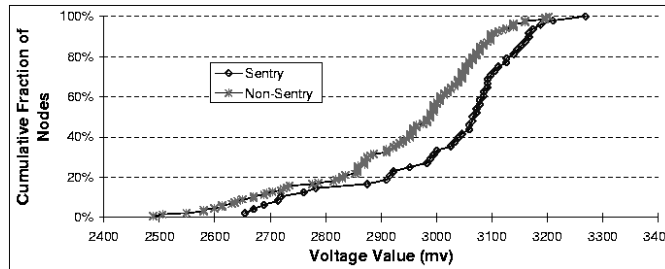


Fig. 12. Effectiveness of sentry selection. This graph is generated using a real deployment of 200 XSM motes. Nodes with a high level of remaining energy have a high probability of being chosen as sentries.

is divided into three sections. The layout indicates that the Voronoi-based tripwire partitioning is very effective and that all nodes attach to their nearest base node through the shortest path.

10.1.2 *Effectiveness of the sentry selection.* In this experiment, we evaluate the effectiveness of sentry selection. Figure 12 plots the cumulative distribution function of the node voltages within the network. The left curve is the voltage CDF of non-sentry nodes and the right curve is the voltage CDF for sentry nodes. It confirms that our sentry selection process is effective and that nodes with high remaining energy have a high probability of being chosen as sentries. For instance, none of nodes with a voltage below 2.65V is chosen as a sentry. Figure 12 further confirms that it is not the case that nodes with high voltages are always selected as sentries, due to the random jitter introduced in Equation 1 and to the localized selection process on a non-uniform distribution of XSM motes.

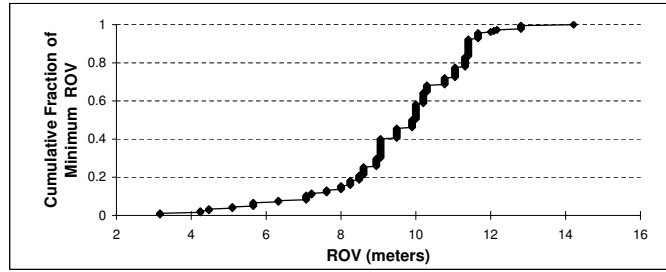


Fig. 13. ROV enforcement results. This graph is generated using a real deployment of 200 XSN nodes. For a ROV of 10 meters, the average minimum distance between sentry pairs is 9.57m with a standard deviation of 1.88m.

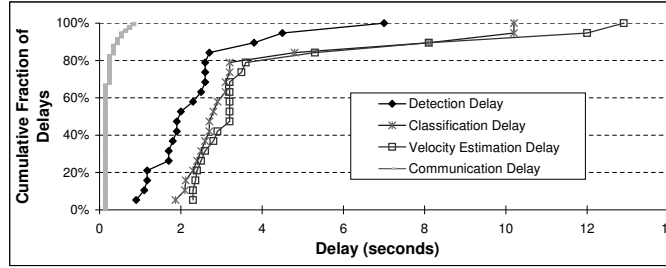


Fig. 14. Distribution of different delays. This graph is generated using a real deployment of 200 XSN nodes. About 80% of detections are done within 2 seconds. More than 80% of target classifications and velocity estimations are done within 4 seconds.

10.1.3 *Effectiveness of ROV enforcement.* We also investigate the effectiveness of enforcing the Range of Vicinity (ROV), when we set the system parameter ROV to 10 meters. Figure 13 shows the cumulative distribution function of the distance between a sentry and the sentry that is the closest to it. The average distance is 9.57 meters with 1.88 meters standard deviation. We note that due to the radio irregularity introduced by ground effects in outdoor environments, a small percentage of sentry nodes (4.4%) cannot reach all the nodes that are very close to it (distance < 5m).

10.1.4 *Delays under power management.* In this experiment, we investigate various delays under power management. When a target enters the surveillance area, a detection report is issued first, followed by classification reports. Finally, after sufficient information is gathered, velocity reports are issued. Figure 14 illustrates the cumulative distribution of different delays. The communication delay (leftmost curve) is much smaller compared with other delays. About 80% of detections are done within 2 seconds. Over 80% of the classification and velocity estimations are made within 4 seconds.

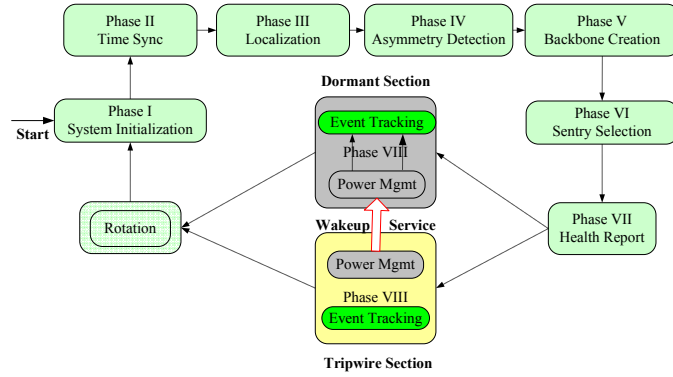


Fig. 15. Phase transition and rotation. The simulator emulates the multi-phase operations of VigilNet.

Table I. Key system parameters. Unless mentioned otherwise, the default values in this table are used in all experiments.

Parameter	Definition	Default Value
SDC	Sentry duty cycle (see 4.3)	25%
STP	Sentry toggle period (see 4.3)	1 second
SSA	Sentry service activation	True
TN	Number of tripwire partitions in the network	1
TDC	Tripwire duty cycle percentage (see 4.1.3)	100%
VS	Target Speed	4 m/s
RN	Number of system rotations per day	1
SR	Sensing Range	10 meters
ROV	Range of Vicinity (see 4.2)	10 meters
RR	Radio Range	30 meters

10.2 Hybrid Evaluation

In the hybrid evaluation, we use experimental measurements from the XSM platform (see Table II) as inputs to the discrete event simulator we built. This simulator emulates the multi-phase VigilNet operations as shown in Figure 15. We distribute 10,000 nodes randomly within a square of edge 1000 meters. The initialization consists of a sequence of phases (from Phase I to VII) in three minutes, before VigilNet enters into the surveillance phase (phase VIII). The system rotates periodically to introduce system-wide soft-states and balance the power. The number of rotations per day is referred to as RN in Table I.

A target enters the network area at a random point on one of the edges and exists the network area at a random point on the opposite edge. The trajectory of the target is a straight line with a constant speed. There is at most one target within the sensor field at any point in time. The entry of targets in the sensor field is evenly distributed throughout the day.

The simulated sensors have a startup time of 1ms. A target is detected when it is within the sensing range of an active node for at least 5ms and when that node can reach its tripwire base station to report the event. The sensory parameters (1ms startup time and 5ms detection time) are typical of the acoustic sensor of the XSM

Table II. Power consumption according to the mote state. This table describes the various sleep states and active states of the sensor nodes. We obtained the power consumption values by empirically measuring the power consumption of XSM nodes.

Node state	Radio State (Messages per second)	Processor State	Sensors State	Total Power
Init	receive (2)	active	off	49.449mW
SentrySleep	off (0)	sleep	off	42μW
NonSentrySleep	LPL (0)	sleep	off	450μW
AwakeComm	receive (2)	active	off	49.449mW
AwakeCommSensing	receive (2)	active	on	71.45mW
AwakeSensing	receive (0)	active	on	70.01mW

platform, which can sample the environment with a frequency of 8192Hz [Dutta et al. 2005]. For the simulation, we assume that environmental noise is low enough not to generate false alarms.

Our simulator correctly models message exchanges for discovering the shortest path to the base. It also models the wake-up of nodes on the path from a node to a base station: nodes use this communication path to notify the base of target detection events. The simulator takes into consideration the energy losses that occur while finding the shortest path to the base, when nodes wake-up nodes on their communication path to the base, during the startup phase of sensors, and during system rotations which, among other tasks, change the state of tripwire sections from turned on to turned off and vice-versa. At this point, our simulator does not model packet loss, node mobility, environmental noise, and physical obstacles. We assume accurate localization, circular radio ranges, and circular sensing ranges. In future work, we plan to increase the complexity of our simulator and to investigate how additional parameters, such as environmental noise and packet loss, affect system performance.

10.2.1 *Battery model and sleep states.* We obtained similar empirical power consumption measurements as reported in [Dutta et al. 2005], which provides very complete analysis of XSM motes. XSM motes use two standard AA (A91) batteries. Each battery has an energy capacity uniformly chosen between 2,848 mAh and 2,852 mAh [Energizer]. However, to better model reality [CrossBow a], we suppose that a mote dies when it has used 85% of the available energy.

The sensor nodes are in one of six power consumption states at any time. These states are the various sleep states and active states that a node can enter. We list and detail the power consumption of these six states in Table II. When a message is transmitted, the radio switches to the transmit state for 30ms (a typical time required by XSM nodes to send a message under the MAC contention). The indicated number of messages per second in Table II is an upper bound result from empirical observations.

10.2.2 *Performance metrics and system parameters.* We investigate three major performance metrics under different system configurations. 1) Detection Probabil-

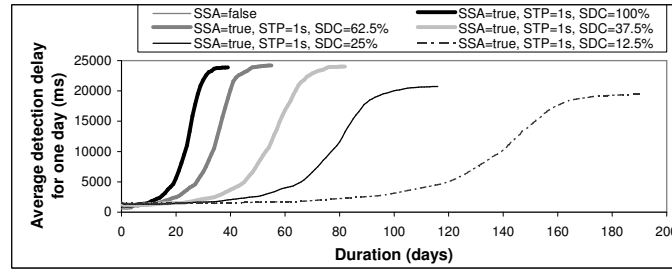


Fig. 16. Influence of sentry duty cycle (SDC) on average detection delay (ADD).

ity (**DP**), which is the percentage of successful detections among all targets that enter into the system during one day. 2) Average detection Delay (**ADD**), which is the average time elapsed between the entrance of a target into the area and its detection by one of the sensor nodes. 3) Network lifetime (**NL**), which is defined as the number of days for which the detection probability of a target remains greater than 90%. The key system parameters are listed in Table I. Unless mentioned otherwise, the default values in Table I are used in all experiments.

10.2.3 Impact of the sentry service and duty cycle scheduling. In this section, we evaluate the energy savings achieved by the sentry service and the duty cycle scheduling. In particular, we study the influence of the activation of the sentry service (SSA), of the sentry duty cycle (SDC), and of the sentry toggle period (STP) on energy consumption. System specifications require the system to support 10 targets per day ($VN=10$). However, to obtain an accurate estimation of the detection probability, we need a higher number of targets. As a consequence, we simulate one hundred targets per day to obtain a good estimation of the detection probability. Only the energy consumption of the 10 first targets is taken into consideration so as to satisfy system requirements. As previously mentioned, we use a network of 10,000 nodes randomly distributed within a square of 1km edge length. Each node has a radio range of 30m. With such a configuration, nodes have an average of 27.5 neighbors within their communication range, and an average of 3.1 neighbors within their sensing range. The density of the the simulated deployment is of a similar order of magnitude as the density of the real deployment, which was approximately of 20 nodes per communication range. When the sentry service is activated ($SSA = true$), 37% of the nodes are initially sentries.

Figure 16, Figure 17, and Figure 18 show the variations of the average detection delay, detection probability, and network lifetime, according to the sentry duty cycle. Remind that we define the network lifetime as the number of days for which the detection probability remains greater than 90%. Figure 17 takes a closer look at a particular section of Figure 16. We first observe that without any sentry service ($SSA = false$), the lifetime of the network is short: all the nodes run out of energy after only four days.

The activation of the sentry service increases the lifetime of the network by approximately seven times. This may seem surprising: with a percentage of sentries

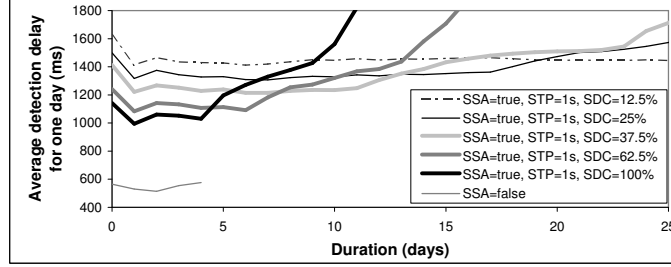


Fig. 17. Influence of sentry duty cycle (SDC) on average detection delay (ADD).

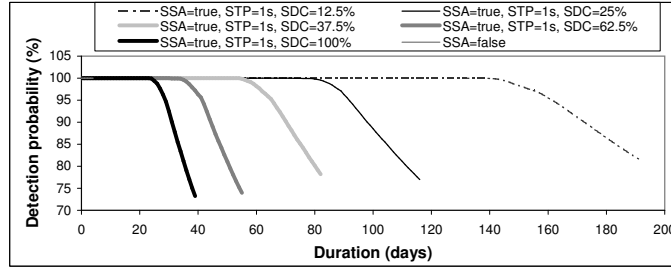


Fig. 18. Influence of sentry duty cycle (SDC) on detection probability (DP).

of 37%, we would expect the network to live only between two and three times longer. However, 37% is the initial percentage of sentry nodes. This percentage decreases as time passes. Indeed, the high initial percentage of sentries is due to network zones characterized by a low density of nodes. These zones rapidly run out of energy as only a small number of nodes can share the sensing task. Consequently, the percentage of sentry nodes decreases.

The activation of the sentry service also increases the average detection delay of a target. This could be expected as the first nodes that the target encounters may be dormant.

The use of duty cycle scheduling ($SDC \neq 100\%$) significantly improves the network lifetime. For instance, with a duty cycle of 12.5%, the lifetime of the network is multiplied by about five times. This may be surprising: we would expect the network lifetime when $SDC=12.5\%$ to be approximately eight times the network lifetime when $SDC=100\%$. The observed values are due to the energy consumed during the rotation phase and when target detection occurs. These tasks consume a non-negligible amount of energy and therefore impose a limit on network lifetime.

We remark that during the first four days of network operation, the average detection delay is shorter when the sentry duty cycle is higher. This could be expected as when a target enters the sensing range of a sentry node, this node may be in a dormant state. We note that the difference between the average detection delays for different values of SDC is no more than one second. This can be explained

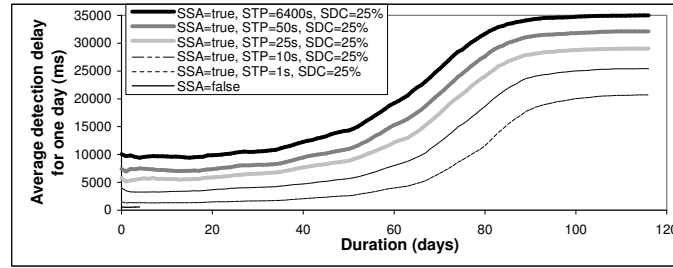


Fig. 19. Influence of sentry toggle period (STP) on average detection delay (ADD).

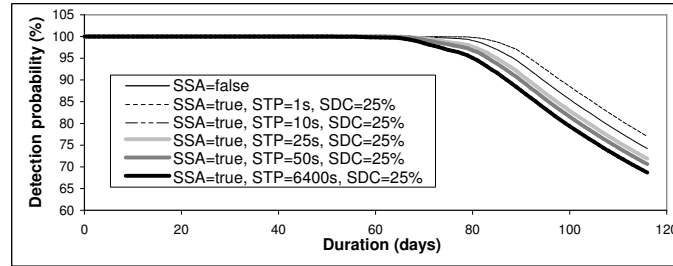


Fig. 20. Influence of sentry toggle period (STP) on detection delay (DP).

by the short sentry toggle period (1s).

Figure 18 shows the influence of the sentry duty cycle (SDC) on the detection probability. We observe that, for all configurations, the initial detection probability is 100%. As nodes start to run out of power, the detection probability decreases until all the nodes become dysfunctional. We note that the high detection probability is not only due to the energy conservation scheme but also due to the size of the network. On average, during the network lifetime, the successful detections reported in Figure 17 occur between 0.5s and 2s after the target entered the square area. After the network lifetime, VigilNet can still detect the targets, however the delay increase gradually as shown in Figure 16.

In Figure 19 and Figure 20, we study the effect of the sentry toggle period (STP) on the average detection delay and the detection probability. We fix the sentry duty cycle at 25%. We observe that a greater toggle period negatively impacts the average detection delay. Indeed, if the toggle period is small (e.g., 1s), a dormant sentry, having a target entering its sensing range, wakes up with a high probability before this target exits the sensing range. Conversely, if the toggle period is big (e.g., 6400s), a dormant sentry has a low probability of being waken up before the target leaves its sensing range.

Guidelines: From the analysis of this section, we can conclude the following. First, to reduce detection delay, we must choose a sentry toggle period as small as possible. Second, the detection probability increases with the size of the network.

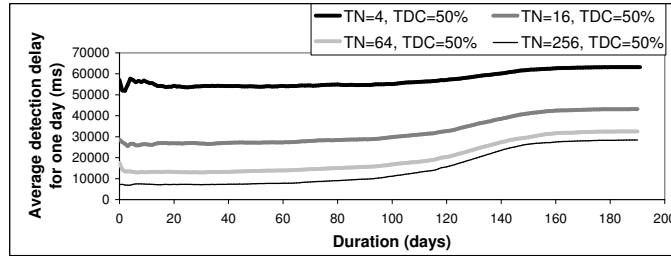


Fig. 21. Influence of number of tripwires (TN) on average detection delay (ADD).

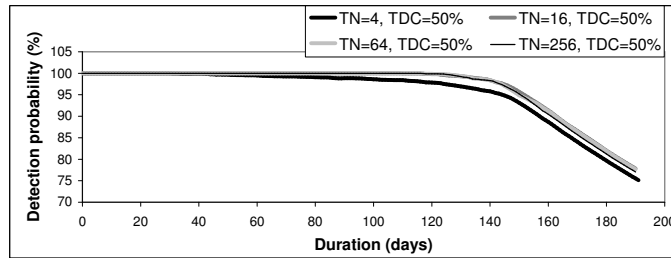


Fig. 22. Influence of number of tripwires (TN) on detection probability (DP).

Third, to increase the network lifetime, we advise to select a small sentry duty cycle. However, note that the time during which a sentry remains awake cannot be arbitrarily small, as it is limited by the time necessary to warm up the sensors and by the time necessary to gather enough sensor data to infer whether there is a target or not. Consequently, rapid sensor wake up and quick target detection algorithms are features that can significantly extend the lifetime of a sensor network. Effort in this direction is worthwhile.

10.2.4 Impact of the tripwire service. We investigate both grid and random placement of tripwire bases. In the case of $\text{Tripwire Number}(\text{TN}) \geq 16$, the two placement strategies generate similar results. For $\text{TN} < 16$, the grid topology performs better. This result could be expected. Indeed, when the network contains few tripwire bases and when these bases are not evenly distributed, the target may encounter only a small number of tripwire partitions while crossing the field. If those partitions are in a sleeping state, the target can cross the whole sensor field without being detected.

Due to the space constraints, we report here only the results concerning the grid tripwire topology.

We configure the wireless sensor network as in Table I. Figures 21 and 22 display the influence of the number of tripwires on the average detection delay and the detection probability. The tripwire duty cycle is 50%. We observe that a small number of tripwires negatively impact the average detection delay, because the

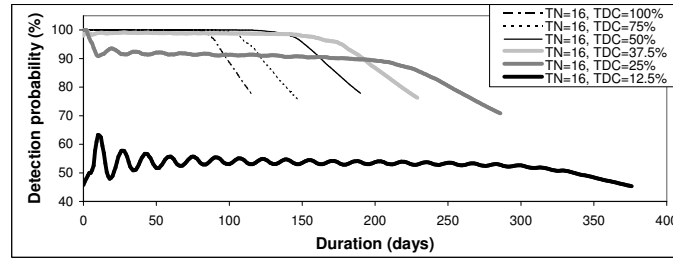


Fig. 23. Influence of tripwire duty cycle (TDC) on detection probability (DP).

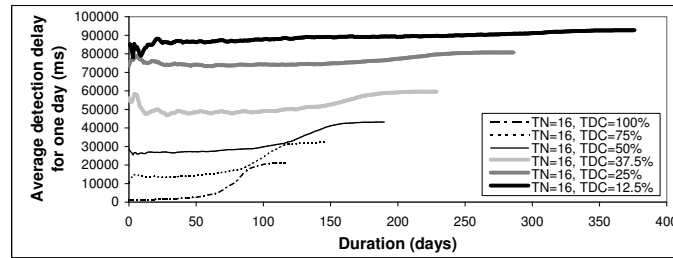


Fig. 24. Influence of tripwire duty cycle (TDC) on average detection delay (ADD).

target may enter the network through a large dormant tripwire section.

In Figures 23 and 24, we study the impact of the tripwire duty cycle on the performance of the network. We choose a tripwires number of sixteen. As we would expect, the smaller the tripwire duty cycle, the longer the lifetime of the network. For instance, when the tripwire duty cycle equals 25%, the network lifetime is about twice the lifetime obtained when the tripwire duty cycle equals 100%. One could expect a multiplication of the lifetime by four times, but this would not take into consideration the energy consumed during the rotation phase and when target detection occurs. Indeed, as an example, when nodes detect a target, they switch to the AwakeSensing state specified in Table II, and the AwakeSensing state (70.01mW) consumes about 1667 times as much energy as the SentrySleep state ($42\mu\text{W}$).

Additionally, we observe that the average detection delay is significantly longer when the tripwire duty cycle is small. Indeed, when the tripwire duty cycle is small, a relatively small portion of the network is awake at any given time, and the target may cover a bigger part of the network without being detected. Finally, we notice that a tripwire duty cycle of less than 25% seriously impacts the detection probability during the first weeks of network operation. This is due to the fact that, with such low levels of tripwire activity, large zones of the network may remain dormant for an extended period of time, producing the possibility that the target crosses the network exclusively through such zones.

Guidelines: From this section, we can conclude the following. First, for a fixed

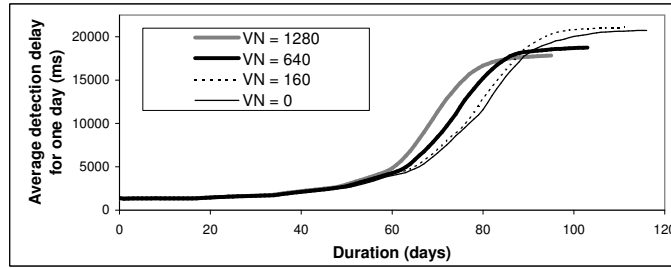


Fig. 25. Influence of number of targets per day (VN) on average detection delay (ADD).

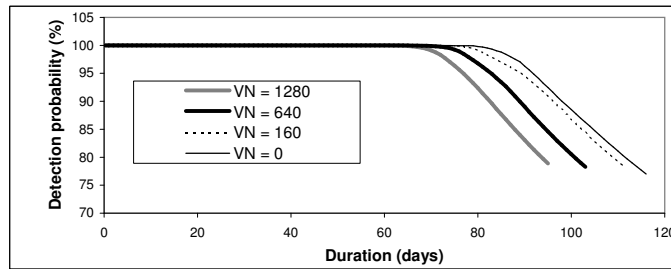


Fig. 26. Influence of number of targets per day (VN) on detection probability (DP).

tripwire duty cycle value, the presence of a large number of tripwire bases decreases detection delay, but does not increase the lifetime of the network significantly. Second, a low tripwire duty cycle increases the network lifetime, but also increases the detection delay and decreases the detection probability. Tripwire duty cycle below 25% is particularly detrimental to performance.

10.2.5 Impact of the targets number and speed. In this section, we study the effect of the number of targets per day and of the target speed on performance. The configuration of the network is the same as in Table I. Figure 25 and Figure 26 report the results of an experiment varying the number of targets per day (VN) from 0 to 1280. Surprisingly, the number of targets per day influences only moderately the lifetime of the network. For instance, when the number of targets per day increases from 0 to 1280, the network lifetime reduces by only 21 days. The reason is that nodes need no more than 5 seconds to detect, classify and report one target.

Figures 27 and 28 show the influence of target speed on average detection delay and detection probability. We observe that a high target speed decreases the detection delay. This may be surprising as when the target speed increases, it spends less time within the sensing range of a given sensor, thereby decreasing the probability of being detected. However, as the target speed increases, it covers more nodes in a shorter amount of time. The effect of target speed on detection probability is insignificant for $VS \leq 16m/s$. We recall that the sensing range of a sensor is ten meters in this experiment. At a speed of 16m/s, the target spends a maximum

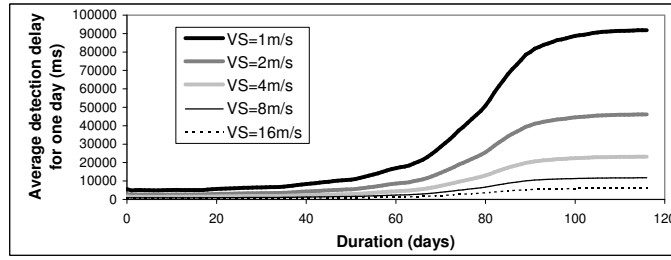


Fig. 27. Influence of target speed (VS) on average detection delay (ADD).

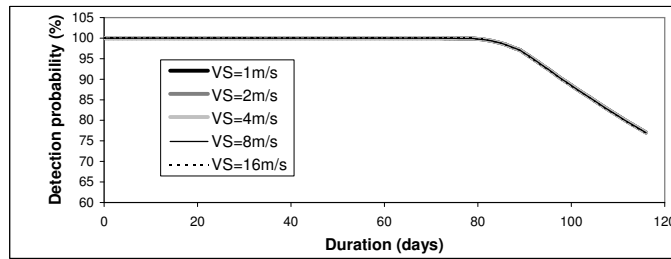


Fig. 28. Influence of target speed (VS) on detection probability (DP).

of 1250ms within the sensing range of a given sensor. This time should remain bigger than the 5 millisecond target detection time. Otherwise, the target cannot be detected. Note that we did not realize experiments using higher speeds because of the fundamental limitations that other components (such as group management) of our system imposes on the maximum trackable speed.

Guidelines: To summarize the results from this section, we can say the following. First, the number of targets per day impacts only moderately the lifetime and performance of the network. As a result, a network designed to handle a specific number of targets per day could cope with an unexpected increase in the frequency of targets. Second, it takes more time to detect slow targets than faster ones. Third, a network with characteristics similar to the one defined in this experiment can handle targets with speeds typical of moving terrestrial objects (up to at least 16 meters per second i.e. 35.8 miles per hour).

10.2.6 Impact of rotation number. Figures 29 and 30 characterize the influence of the number of rotations per day on the performance of the network. The network parameters are the same as in Table I. During a rotation, the network realizes essential operations such as spanning tree construction and/or healing, time synchronization, and management of sentries, tripwires, and groups. It appears that rotations are costly operations that significantly reduce the network lifetime.

Guidelines: In brief, the rate of system-wide rotations should be small and the duration of the rotation should be short. In our real network, one rotation per day proved to be sufficient to fulfill system requirements.

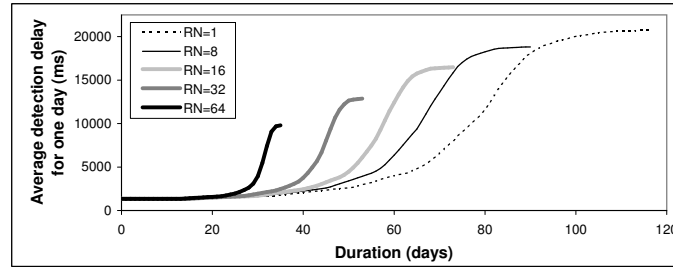


Fig. 29. Influence of number of rotations per day (RN) on average detection delay (ADD).

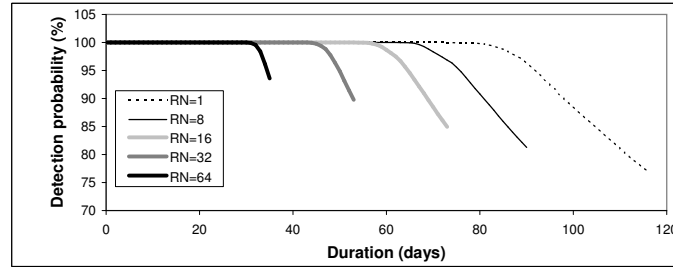


Fig. 30. Influence of number of rotations per day (RN) on detection probability (DP).

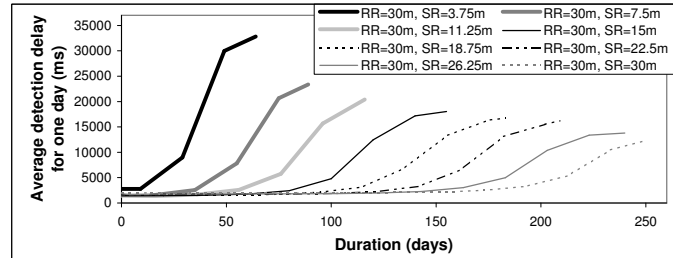


Fig. 31. Influence of sensing radius (SR) on average detection delay (ADD).

10.2.7 *Impact of the sensing range.* Figures 31 and 32 show the result of an experiment characterizing the influence of the sensing range on performance. The network parameters are the same as in Table I. The radio range remains equal to 30m and we vary the sensing range from 3.75m to 30m. We observe that a bigger sensing range dramatically improves network lifetime, average detection delay and detection probability. For instance the lifetime of the network when the sensing range equals 30m is about five times the lifetime of the network when the sensing range equals 3.75m. This is not surprising as an increased sensing range reduces the number of necessary sentries to maintain full coverage.

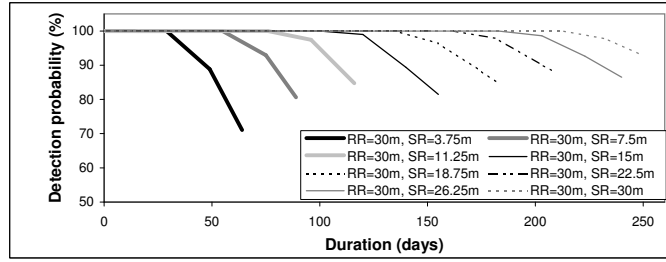


Fig. 32. Influence of sensing radius (SR) on detection probability (DP).

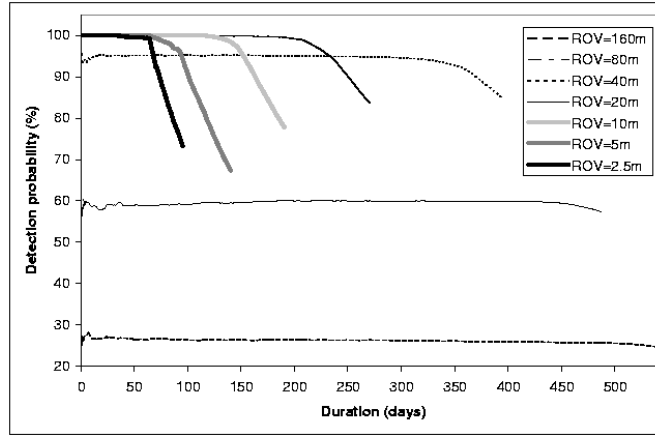


Fig. 33. Influence of range of vicinity (ROV) on detection probability (DP).

Note that, typically, we can modify the sensing range of an acoustic sensor by modifying its amplification factor. Similarly, we can modify the sensing range of magnetic and PIR sensors by modifying the value of their detection threshold. Increasing the sensing range of such sensors increases their sensibility to environmental noise and can in turn increase the rate of false alarms (false positives) depending on the amount of environmental noise. False alarms produce an increase in power draw as triggered sensors are turned on unnecessarily. For our simulations, we assume that environmental noise is low enough not to generate false alarms. The study of the effect of environmental noise on false alarm rate, and the study of the effect of false alarm rate on power consumption are left for future work.

Guidelines: In conclusion, investing in high range sensors is cost effective: even though such sensors are more expensive, the resulting network lasts longer.

10.2.8 *Impact of the Range of Vicinity.* Figures 33 and 34 show the influence of the range of vicinity (ROV) on the performance of the network. We configure the wireless sensor network as in Table I except for the tripwire duty cycle percentage

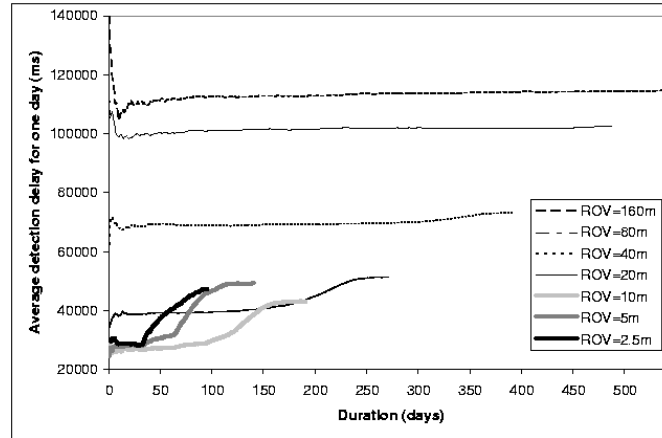


Fig. 34. Influence of range of vicinity (ROV) on average detection delay (ADD).

(TDC) that is set to be 50% instead of 100%. We recall that ROV is the radio range used during the execution of the sentry selection protocol (see Section 4.2). If ROV is greater than the sensing range, the selected sentries monitor an area that they cannot entirely sense: targets are more likely to escape detection. Note that the greater ROV is, the less the number of selected sentries is: the lifetime of the network is extended. If the sensing range is greater than ROV, selected sentries provide redundant coverage for some area of the network. Note that ROV is different from the communication range, which is used to transmit sensor data back to a base station. We can use a ROV different from the communication range by dynamically modifying the value of the radio transmission power.

Figure 33 confirms what we stated in the previous paragraph. Recalling that the lifetime of the network is defined as the number of days for which the detection probability of a target remains greater than 90%, we observe that the lifetime of the network is monotonically increasing from 73 days to 369 days as ROV increases from 2.5 meters to 40 meters. However, the lifetime falls sharply to 0 days for an ROV of 80 meters or 160 meters. This is not surprising: for such ROV values, a large area of the network is not monitored and the initial detection probability is too small.

Figure 34 shows the influence of ROV on the average detection delay. Varying ROV from 2.5 meters to 10 meters, we observe that, for a given day, the greater the ROV, the lesser the average detection delay. This is explained by the fact that, when ROV is lesser than the sensing range (10 meters), smaller ROV values mean more redundancy: nodes consume the totality of their energy resources more rapidly, and the detection delay increases accordingly. Varying ROV from 20 meters to 160 meters, we observe that, for a given day, the greater the ROV, the greater the average detection delay. This is explained by the fact that, when ROV is greater than the sensing range (10 meters), greater ROV values mean less coverage, which results into greater detection delays.

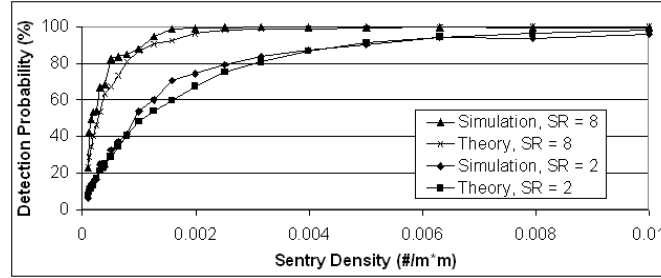


Fig. 35. Comparison Between Theory and Simulation: Detection Probability Versus Sentry Density When Using the Sentry Service. For a sensing range of 2m, the Pearson Correlation Coefficient (R) between theory and simulation equals 0.994. For a sensing range of 8m, R=0.984.

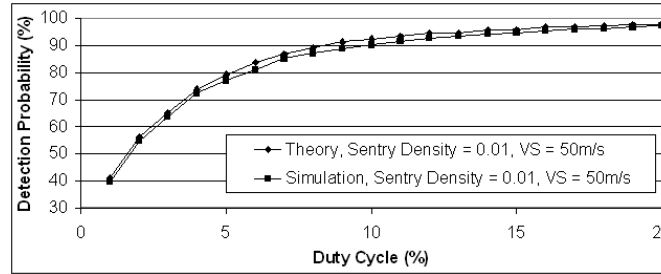


Fig. 36. Comparison Between Theory and Simulation: Detection Probability Versus Duty Cycle When Using the Sentry Service and Duty Cycle Scheduling. In this graph, the Pearson Correlation Coefficient between theory and simulation equals 0.999.

Guidelines: To conclude, we can say that increasing ROV up to the sensing range results in an extension of the lifetime and in an amelioration of the average detection delay. However, special care must be taken when increasing ROV beyond the sensing range. Then, even though sensor nodes consume energy at a slower pace, the average detection delay and the detection probability drop significantly.

10.2.9 *Comparison Between Theory and Simulation.* In Section 7, we derived a theoretical model allowing us to determine the impact of the sentry service and of the duty cycle scheduling service on the detection probability. We now quantify how results obtained through our theoretical model correlate with results obtained with our hybrid simulator.

In Figure 35, we use both theory and simulation to determine the influence of the sentry density on the detection probability when we use only the sentry service, for sensing ranges of 2m and 8m. We consider a sensor network of 1000 nodes randomly distributed within an area of $100m \times 1000m$. We deactivate the duty cycle and tripwire services.

In Figure 36, we use both theory and simulation to determine the influence of the

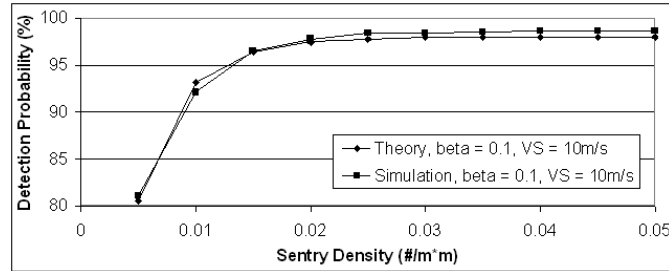


Fig. 37. Comparison Between Theory and Simulation: Detection Probability Versus Sentry Density When Using the Sentry Service and Duty Cycle Scheduling. In this graph, the Pearson Correlation Coefficient between theory and simulation equals 0.996.

duty cycle service on the detection probability with a sentry density of 0.01 sentries per m^2 , a target speed of 50m/s, and a sensing range of 10m. We consider a sensor network of 5000 nodes randomly distributed within an area of $100m \times 1000m$. Both the sentry service and the duty cycle service are activated, but not the tripwire service. We simulate 1000 targets crossing the sensor field. The duty cycle is expressed in percentage of the toggle period.

In Figure 37, we use both theory and simulation to determine the influence of the sentry density on the detection probability when we use both the sentry service and the duty cycle service. The system parameters are the same as in the previous experiment except that the vehicle speed is 10m/s, and $\beta = 0.1$ (β is defined in Section 7.4). In this experiment and the previous one, we set the ROV parameter of the simulator appropriately so as to obtain the desired sentry density.

From the described experiments, we obtain a total of 4 pairs of curves comparing theoretical and simulated results. The Pearson correlation coefficient for these 4 pairs of curves varies from 0.984 to 0.999, indicating a very high correlation between theoretical and simulation predictions, and providing evidence for the correctness of both our theoretical derivations and our simulator implementation.

11. CONCLUSION

This paper presents a recent major effort to address the energy efficiency for outdoor long-term surveillance. We investigated the power management at the network, section and node level by using a novel tripwire service, sentry service and duty cycle scheduling, respectively. We implemented our system using the XSM platform, and deployed a network of 200 nodes in an outdoor environment. We used the real network of 200 nodes and an analytical probabilistic model to evaluate key system parameters, and we used a hybrid simulation of 10,000 nodes to estimate network lifetime under various settings and conditions. Our results demonstrate the effectiveness of our approach and identify several useful guidelines and lessons for the future development of energy-efficient sensor systems.

REFERENCES

- AGARWAL, M., CHO, J. H., GAO, L., AND WU, J. 2004. Energy-Efficient Broadcast in Wireless Ad hoc Networks with Hitch-hiking. In *IEEE Infocom*.
- ARORA, A., DUTTA, P., AND ET.AL, S. B. 2003. A Wireless Sensor Network for Target Detection, Classification, and Tracking. *Computer Networks (Elsevier) Systems*.
- BHATTACHARYA, S., KIM, H., PRABH, S., AND ABDELZAHER, T. 2003. Energy-Conserving Data Placement and Asynchronous Multicast in Wireless Sensor Networks. In *MobiSys*.
- BOGDANOV, A., MANEVA, E., AND RIESENFELD, S. 2004. Power-Aware Base Station Positioning for Sensor Networks. In *IEEE Infocom*.
- CAO, Q., ABDELZAHER, T., HE, T., AND STANKOVIC, J. 2005. Towards Optimal Sleep Scheduling in Sensor Networks for Rare Event Detection . In *IPSN*.
- CAO, Q., YAN, T., ABDELZAHER, T., AND STANKOVIC, J. 2005. Analysis of Target Detection Performance for Wireless Sensor Networks. In *DCOSS*.
- CARDEI, M., THAI, M., AND WU, W. 2005. Energy-Efficient Target Coverage in Wireless Sensor Networks. In *IEEE Infocom*.
- CHEN, B., JAMIESON, K., BALAKRISHNAN, H., AND MORRIS, R. 2001. Span: An Energy-Efficient Coordination Algorithm for Topology Maintenance in Ad Hoc Wireless Networks. In *MOBICOM*.
- CHOI, W. AND DAS, S. 2005. A Novel Framework for Energy-Conserving Data Gathering in Wireless Sensor Networks. In *IEEE Infocom*.
- CrossBow. *Mica2 aa battery pack service life test*. CrossBow. Available at <http://www.xbow.com/Support/>.
- CrossBow. *Mica2 data sheet*. CrossBow. Available at <http://www.xbow.com>.
- DUTTA, P., GRIMMER, M., ARORA, A., BIBY, S., AND CULLER, D. 2005. Design of a Wireless Sensor Network Platform for Detecting Rare, Random, and Ephemeral Events. In *IPSN*.
- Energizer. *Energizer current batteries datasheets*. Energizer. Available at <http://www.energizer.com>.
- GANESAN, D., CRISTESCU, R., AND BEREFULL-LOZANO, B. 2004. Power-Efficient Sensor Placement and Transmission Structure for Data Gathering under Distortion Constraints. In *IPSN*.
- GOLDBERG, D. H., ANDREAS G. ANDREOU, P. J., POULIQUEN, P. O., RIDDLE, L., AND ROSASCO, R. 2004. A Wake-up Detector for an Acoustic Surveillance Sensor Network: Algorithm and VLSI Implementation. In *IPSN*.
- GU, L. AND STANKOVIC, J. A. 2004. Radio-Triggered Wake-Up Capability for Sensor Networks. In *Proceedings of RTAS*.
- HE, T., BLUM, B. M., STANKOVIC, J. A., AND ABDELZAHER, T. F. 2004. AIDA: Adaptive Application Independent Data Aggregation in Wireless Sensor Networks. *ACM Transactions on Embedded Computing System, Special issue on Dynamically Adaptable Embedded Systems*.
- HE, T., HUANG, C., BLUM, B. M., STANKOVIC, J. A., AND ABDELZAHER, T. 2003. Range-Free Localization Schemes in Large-Scale Sensor Networks. In *Proc. of the Intl. Conference on Mobile Computing and Networking (MOBICOM)*.
- HE, T., KRISHNAMURTHY, S., STANKOVIC, J. A., AND ABDELZAHER, T. 2004. An Energy-Efficient Surveillance System Using Wireless Sensor Networks. In *MobiSys*.
- HEINZELMAN, W., CHANDRAKASAN, A., AND BALAKRISHNAN, H. 2000. Energy-Efficient Communication Protocol for Wireless Microsensor Networks. In *HICSS*.
- KAHN, J. M., KATZ, R. H., AND PISTER, K. S. J. 1999. Next Century Challenges: Mobile Networking for Smart Dust. In *MOBICOM*.
- KAR, K., KRISHNAMURTHY, A., AND JAGGI, N. 2005. Dynamic Node Activation in Networks of Rechargeable Sensors. In *IEEE Infocom*.
- KESHAVARZIAN, A., UYSAL-BIYIKOGLU, E., HERRMANN, F., AND MANJESHWAR, A. 2004. Energy-Efficient Link Assessment in Wireless Sensor Networks. In *IEEE Infocom*.
- KUMAR, S., LAI, T. H., AND BALOGH, J. 2004. On k-Coverage in a Mostly Sleeping Sensor Network. In *10th Annual International Conference on Mobile Computing and Networking*.

- LIU, J., REICH, J., AND ZHAO, F. 2003. Collaborative In-Network Processing for Target Tracking. *J. on Applied Signal Processing*.
- MADDEN, S., FRANKLIN, M., HELLERSTEIN, J., AND HONG, W. 2002. TAG: A Tiny Aggregation Service for Ad-Hoc Sensor Networks. In *Operating Systems Design and Implementation*.
- MAROTI, M., KUSY, B., SIMON, G., AND LEDECZI, A. 2004. The Flooding Time Synchronization Protocol. In *SenSys*.
- OKABE, A., BOOTS, B., SUGIHARA, K., AND CHIU, S. N. 2000. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley.
- PARADISO, J. A. AND STARNER, T. 2005. Energy Scavenging for Mobile and Wireless Electronics. *IEEE Pervasive Computing* 4, 1.
- POLASTRE, J. AND CULLER, D. 2004. Versatile Low Power Media Access for Wireless Sensor Networks. In *SenSys*.
- ROUNDY, S., WRIGHT, P. K., AND RABAEEY, J. 2006. A Study of Low Level Vibrations as a Power Source for Wireless Sensor Nodes. *Computer Communications* 26, 11.
- SEADA, K., ZUNIGA, M., HELMY, A., AND KRISHNAMACHARI, B. 2004. Energy Efficient Forwarding Strategies for Geographic Routing. In *SenSys*.
- SHRIVASTAVA, N., BURAGOHAJ, C., SURI, S., AND AGRAWAL, D. 2004. Medians and Beyond: New Aggregation Techniques for Sensor Networks. In *SenSys*.
- SICHTIU, M. L. 2004. Cross-Layer Scheduling for Power Efficiency in Wireless Sensor Networks. In *IEEE Infocom*.
- SIMON, G. AND ET. AL. 2004. Sensor Network-Based Countersniper System. In *SenSys*.
- STOLERU, R., HE, T., AND STANKOVIC, J. A. 2004. Walking GPS: A Practical Solution for Localization in Manually Deployed Wireless Sensor Networks. In *1st IEEE Workshop on Embedded Networked Sensors EmNetS-I*.
- SZEWczyk, R., MAINWARING, A., ANDERSON, J., AND CULLER, D. 2004. An Analysis of a Large Scale Habit Monitoring Application. In *SenSys*.
- VAN DAM, T. AND LANGENDOEN, K. 2003. An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *SenSys*.
- WANG, X., XING, G., ZHANG, Y., LU, C., PLESS, R., AND GILL, C. 2003. Integrated Coverage and Connectivity Configuration in Wireless Sensor Networks. In *SenSys*.
- WILLIAMS, R. 1979. *Circle Coverings*. New York: Dover.
- WOO, A., TONG, T., AND CULLER, D. 2003. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *SenSys*.
- XU, N., RANGWALA, S., CHINTALAPUDI, K. K., GANESAN, D., BROAD, A., GOVINDAN, R., AND ESTRIN, D. 2004. A Wireless Sensor Network for Structural Monitoring. In *SenSys*.
- XU, Y., HEIDEMANN, J., AND ESTRIN, D. 2001. Geography-informed Energy Conservation for Ad Hoc Routing. In *MobiCom*.
- YAN, T., HE, T., AND STANKOVIC, J. 2003. Differentiated Surveillance Service for Sensor Networks. In *SenSys*.
- YE, W., HEIDEMANN, J., AND ESTRIN, D. 2002. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *INFOCOM*.
- YU, Y., KRISHNAMACHARI, B., AND PRASANNA, V. K. 2004. Energy-Latency Tradeoffs for Data Gathering in Wireless Sensor Networks. In *IEEE INFOCOM*.
- ZHOU, G., HE, T., AND STANKOVIC, J. A. 2004. Impact of Radio Irregularity on Wireless Sensor Networks. In *MobiSys*.