2005

Protecting the Communication Structure in Sensor Networks

S. Olariu
Old Dominion University

Q. Xu
Old Dominion University

M. Eltoweissy

A. Wadaa

Original Publication Citation

Olariu, S., Xu, Q., Eltoweissy, M., Wadaa, A., & Zomaya, A. Y. (2005). Protecting the communication structure in sensor networks. *International Journal of Distributed Sensor Networks, 1*(2), 187-203. doi: 10.1080/15501320590966440

International Journal of Distributed Sensor Networks, 1: 187–203, 2005

ISSN: 1550-1329 print/1550-1477 online
DOI: 10.1080/15501320590966440

Taylor & Francis
Taylor & Francis Group

Protecting the Communication Structure in Sensor Networks

S. OLARIU and Q. XU

Department of Computer Science, Old Dominion University, Norfolk, VA

M. ELTOWEISSY

Department of Computer Science, Virginia Tech, Falls Church, VA

A. WADAA

Intel Corporation, Hillsboro, Oregon

A.Y. ZOMAYA

University of Sydney, Sydney, Australia

In the near future wireless sensor networks will be employed in a wide variety of applications establishing ubiquitous networks that will pervade society. The inherent vulnerability of these massively deployed networks to a multitude of threats, including physical tampering with nodes exacerbates concerns about privacy and security. For example, denial of service attacks (DoS) that compromise or disrupt communications or target nodes serving key roles in the network, e.g. sink nodes, can easily undermine the functionality as well as the performance delivered by the network. Particularly vulnerable are the components of the communications or operation infrastructure. Although, by construction, most sensor network systems do not possess a built-in infrastructure, a virtual infrastructure, that may include a coordinate system, a cluster structure, and designated communication paths, may be established post-deployment in support of network management and operation. Since knowledge of this virtual infrastructure can be instrumental for successfully compromising network security, maintaining the anonymity of the virtual infrastructure is a primary security concern.

Somewhat surprisingly, in spite of its importance, the anonymity problem has not been addressed in wireless sensor networks. The main contribution of this work is to propose an energy-efficient protocol for maintaining the anonymity of the virtual infrastructure in a class of sensor network systems. Our solution defines schemes for randomizing communications such that the cluster structure, and coordinate system used remain undetectable and invisible to an observer of network traffic during both the setup and operation phases of the network.

Keywords wireless sensor networks; energy-efficient protocols; scalable protocols; traffic anonymity

This work was supported in part by a grant from the Commonwealth of Virginia Technology Research Fund (SE 2001-01) through the Commonwealth Information Security Center.

Address correspondence to S. Olariu, Department of Computer Science, Old Dominion University, Norfolk, VA, 23529.

1. Introduction

Recent advances in nano-technology made it technologically feasible and economically viable to develop low-power, battery-operated devices that integrate special-purpose computing with low-power sensing and wireless communications capabilities [16, 23, 24, 40, 49, 55, 59]. It is expected that these small devices, referred to as *sensor nodes*, will be mass-produced making production costs negligible [20, 44, 60, 76]. Individual sensor nodes have a non-renewable power supply and, once deployed, must work unattended. We envision a massive random deployment of sensor nodes, numbering in the thousands or tens of thousands. Aggregating sensor nodes into sophisticated computation and communication infrastructures, called sensor networks, will have a significant impact on a wide array of applications including military, scientific, industrial, health, and domestic. The fundamental goal of a wireless sensor network is to produce, over an extended period of time, meaningful global information from local data obtained by individual sensor nodes [1, 11, 15, 28, 49, 55, 63, 64, 70].

However, a wireless sensor network is only as good as the information it produces [2-6, 8, 9, 12-14, 17-19, 45]. In this respect, perhaps the most important concern is *information* security. Indeed, in most application domains sensor networks will constitute a mission critical component requiring commensurate security protection [41, 45, 48, 71-75]. Sensor network communications must prevent disclosure and undetected modification of exchanged messages. Due to the fact that individual sensor nodes are anonymous and that communication among sensors is via wireless links, sensor networks are highly vulnerable to security attacks. If an adversary can thwart the work of the network by perturbing the information produced, stopping production, or pilfering information, then the perceived usefulness of sensor networks will be drastically curtailed. Thus, security is a major issue that must be resolved in order for the potential of wireless sensor networks to be fully exploited [21-24, 29-39]. The task of securing wireless sensor networks is complicated by the fact that the sensors are mass-produced anonymous devices with a severely limited energy budget and initially unaware of their location [1, 41–43, 51–53].

In many application domains, safeguarding *output data assets*, data produced by the sensor network and consumed by the end user (application), against loss or corruption is a major security concern. In these application domains, a sensor network is deployed into a hostile target environment for a relatively extended amount of time. The network self-organizes and works to generate output data that is of import to the application. For example, a sensor network may be deployed across a vast expanse of enemy territory ahead of a planned attack; the network system monitors the environment and produces reconnaissance data that is key to a mission planning application. Periodically, during the network lifetime, a mobile gateway, mounted on a person, land or airborne vehicle, or a satellite, collects the output data assets from the network system, to maintain an up to date state. This means the network system must store the output data assets from the time it is produced until it is collected. Therefore, securing the output data assets in the network is an important problem in this class of applications [7, 10, 15, 26, 27, 34 ,45, 48, 50-52, 54, 56–58, 62, 63].

We view an attack on the output data assets in the sensor network as a type of denial of service attacks. This is based on the abstraction that output data is stored in a logical *repository*, and, that *access* to this output data repository constitutes, in effect, a service provided by the network system to the application; corruption or loss of output data denies the application access to that service. Many wireless sensor networks are mission-oriented, must work unattended, and espouse data-centric processing. Consequently, they are significantly different in their characteristics from ad-hoc networks; security solutions designed specifically for wireless sensor networks are therefore required.

1.1 What is Anonymity?

With the recent growth and wide acceptance of novel applications including e-banking, e-commerce, e-voting, and e-government, concerns about *privacy* and security have become extremely important. In all these applications *anonymity* protects the identity of the sender or receiver and guarantees that both parties involved in a communication transaction remain anonymous to each other. Recent years have seen a flurry of activity and many anonymous communication systems have been developed for the Internet [3, 7, 8, 9, 11, 32, 33, 34, 41, 43, 61, 62]. Most of the work on anonymity is concerned with *sender* anonymity, *receiver* anonymity, and *mutual* anonymity. Of the various forms of anonymity mentioned above, sender anonymity seems to be the most critical in current applications. In *e-voting*, for example, a vote should not be traceable back to the voter that cast it. Likewise, users may not wish to disclose their identities when visiting web sites. Quite recently, *traffic* anonymity has also received well-deserved attention in the literature. Indeed, if an adversary can identify traffic patterns and routers effecting them, the security risks to the traffic increase multifold. In a homeland security context, a group of terrorists may attempt a surgical operation taking out those routers whose loss will inflict the heaviest damage on the network.

Recently, the problem of securing ad-hoc networks has received a great deal of well-deserved attention in the literature [5, 6, 9, 29, 30, 33, 42, 48]. Somewhat surprisingly, however, in spite of its importance, the anonymity problem has not been addressed in wireless sensor networks.

We view this work as an initial contribution towards developing a lightweight solution for the anonymity problem in wireless sensor networks. In this paper, we focus primarily on *structure* anonymity in a wireless sensor network. Although the exact statement of the problem we address are presented in Sec. 5, suffice it to say that the basic infrastructure of a wireless sensor network that is constructed immediately after deployment is a coordinate system that affords natural clustering. Our main contribution is to protect the coordinate system from an external adversary by making the process of acquiring the coordinate system anonymous and, thus, invisible, to the adversary.

2. The Network Model

The network model used in this paper is a derivative of the model introduced in [46, 67]. Specifically, we assume a class of wireless sensor networks consisting of a large number of sensors nodes randomly deployed in the environment of interest. A training process, as explained below, establishes a coordinate system and defines a clustering of all nodes. In post training, the network undergoes multiple operation cycles during its lifetime. The training process also endows the role of *sink* upon one or more of the defined clusters. The sink role is transient, however, since new sink clusters are designated at the beginning of each operation cycle. Each sink cluster, henceforth called sink, acts as a repository for a portion of the sensory data, generated in the network during an operation cycle. At the end of an operation cycle, each sink transfers data stored in its repository to a *gateway*. In the following we describe the three primary entities in our network model in more detail.

2.1 The Sensor Node

We assume a sensor to be a device that possesses three basic capabilities; sensing, computing, and communication. At any point in time, a sensor node is either performing one of a fixed set of *sensor primitive operations*, or is idle (asleep).

We assume that individual sensor nodes have four fundamental constraints:

- Sensors are anonymous; initially a sensor node has no unique identifier,
- Each sensor has a limited non-renewable energy budget,
- Each sensor attempts to maximize the time it is in sleep mode; a sensor wakes up at specific (possibly random) points in time for short intervals under the control of a timer, and
- Each sensor has a modest transmission range, perhaps a few meters with the ability to send and receive over a wide range of frequencies. In particular, communication among sensor nodes in the sensor network must be multi-hop.

2.2 The Sink

The sink granularity in our network model is a cluster in a trained network. For each operation cycle a number of clusters are designated to serve as sinks. This means that all nodes in a sink cluster serve as sink nodes. Coarser sink granularity supports higher sink storage capacity, and thus potentially longer operation cycles. However, coarser sink granularity, as envisioned here, comes at a potentially higher risk of anonymity attacks due to the space correlation among sink nodes. If it is discovered that a node, x, is a sink node then it immediately follows that there exists at least one other sink node (typically more) in the vicinity of x. Thus, the success of an anonymity attack in our model is sensitive to the probability of identifying the *first* node in a sink. One approach to decrease the latter probability is to have only a subset of the nodes in a sink cluster serve as sink nodes. There is a tradeoff between sink granularity, and hence its capacity and operation longevity, and the amount of security a sink has against anonymity attacks.

A notable advantage of our sink model is that sinks are dynamic configurations of regular nodes in the sensor network, as opposed to being, for example, special supernodes. This makes the system less complex to design and operate, and eliminates a would-be attractive focus for anonymity attacks. Another major advantage is that the sink, in our model, is not a single point of failure, in two distinct ways. First, the network uses multiple sink entities, as opposed to a single sink entity. Second, a sink entity is not a single node; rather it is a set of nodes.

The solution proposed in this paper for the anonymity problem uses whole cluster granularity for the sink. It encompasses techniques for randomly and securely choosing sinks for an operation cycle, and maintaining the anonymity of these sinks.

2.3 The Gateway

The gateway is an entity that connects the sensor network system to the outside world. The gateway is not constrained in mobility, energy, computation, or communication capabilities. There are two basic functions for the gateway in our network model:

(i)*Training*. The gateway performs the network training process, post deployment. For training purposes the gateway is assumed to be able to send long-range, possibly, directional broadcasts to all sensors. It should be noted that our training process does not involve any transmission from the sensor nodes deployed in the network. Thus, in principle, the gateway does not have to be geographically co-located with the sensor nodes during the training process.

(ii)*Harvesting* (data collection). At the end of an operation period, the gateway must collect sensory data stored in each sink. In a simple collection scenario the gateway traverses the deployment environment to collect data from all the sinks. If we assume

that nodes comprising a sink perform collaborative data fusion, then the harvesting process requires a relatively short period of time. Specific harvesting protocols are beyond the scope of this paper.

(iii)*Harvesting* (data collection). At the end of an operation period, the gateway must collect sensory data stored in each sink. In a simple collection scenario the gateway traverses the deployment environment to collect data from all the sinks. If we assume that nodes comprising a sink perform collaborative data fusion, then the harvesting process requires a relatively short period of time. Specific harvesting protocols are beyond the scope of this paper.

In our proposed solution for the anonymity problem discussed in detail in Sec. 4, the gateway is assigned the additional function of a group key management server for the sensor network.

3. Network Organization and Clustering

Figure 1(a) features an untrained sensor network immediately after deployment in an environment that is represented here as a 2-dimensional plane. For simplicity, we assume that the trainer is centrally located relative to all deployed nodes. Namely, considering the nodes to be points in the plane, we assume, in this paper, that the trainer is located at the center of the smallest circle in the plane that contains all deployed nodes. It should be noted that this is not a necessary condition. To be specific, our proposed training scheme is applicable if the trainer is located either at the center of, or at a point outside of the smallest circle that contains all deployed nodes.

The primary goal of training is to establish a *coordinate system*, to provide the nodes *location awareness* in that system, and to organize the nodes into *clusters*. The coordinate system, and clustering are briefly explained next. We refer the interested reader to [20] for an in-depth description of the training process.

3.1 The Coordinate System

The training process establishes a polar coordinate system as exemplified by Figure 1 (b). The coordinate system divides the sensor network area into equiangular *wedges*. In turn, these wedges are divided into sectors by means of concentric circles or *coronas* centered

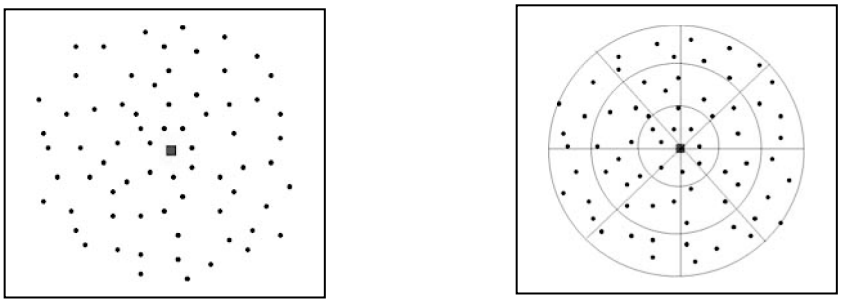

(a) **(b)**

FIGURE 1 (a) An untrained sensor network with a centrally located trainer; (b) A trained sensor network

at the trainer location. Corona radii can be determined based on several criteria, e.g. in [46, 67] they are designed to maximize the efficiency of *sensors-to-sink* multihop communication. The intersection of every wedge and corona defines a unique sector; each sector is uniquely identifiable by the combination of its unique wedge identifier, and unique corona identifier. The training process guarantees that each node belongs to one and only one sector in the coordinate system, and that each node knows the identity of its sector [46, 67].

Let c, and w be, respectively, the set of coronas, and the set of wedges defined by the training process. The resulting coordinate system can thus be formally represented by $\{(r_0, r_1, ..., r_{|c|-1}), \theta\}$, where r_i is the radius of corona i, $0 \leq i \leq |c| - 1$, θ is the wedge angle, and $|w| = 2\pi/\theta$. A fundamental assumption here is that any coordinate system is designed such that all nodes located in the same sector can communicate using direct (single hop) transmission.

3.2. Clustering

A major advantage of our coordinate system is that sectors implement the concept of clustering (at no additional cost). A sector effectively constitutes a cluster; clusters are disjointed, and are uniquely identifiable. All nodes located in the same sector are members of the same cluster, and have the same location coordinates, namely, the corona and wedge identifiers corresponding to that sector. This clustering scheme is ideally suited for sensor nodes that are intrinsically anonymous. Wadaa *et al.* proposed in [67] a scalable training protocol where each untrained node incurs a communication cost equal to $log|w| + log|c|$, and the nodes do not transmit any messages during the training process.

4. The Work Model

The work model defines how sensor nodes work collaboratively to generate and store sensory data during an operation cycle. We propose a work model that is based on two principles:

- *Intra-cluster* activity generates sensory data of interest to the application, and
- *Inter-cluster* activity transports each granule of sensory data to a sink, randomly chosen from among those available, for storage.

4.1 Intra-cluster Activity

In our model, the sensory data resulting from intra-cluster activity encodes states of a process of interest. Namely, we assume that the goal of intra-cluster activity is to monitor a process (or phenomenon), and report on its *local state* at any point in time. The state space of the phenomenon is given by $\{s_0, s_1, s_2, \ldots, s_z\}$. State s_0 denotes the *normal state*, and each of s_i, $1 \leq i \leq z$ denotes an *exception state.* The assumption here is each state s_i, $1 \leq i \leq z$, corresponds to an application-defined exception of a particular type. The normal state corresponds to the fact that no exception of any type is detected.

We propose a transaction-based model for managing the computation and reporting of target process states. The model is a specialization of a transaction-based management model for sensor networks introduced in [46, 67]**.** In this model, intra-cluster activity proceeds as follows. For a given cluster, subsets of nodes located in the cluster dynamically band together forming *workforces*.

Periodically, members of each workforce collaborate to perform an instance of a state computation transaction preloaded into each node. The transaction computes and reports

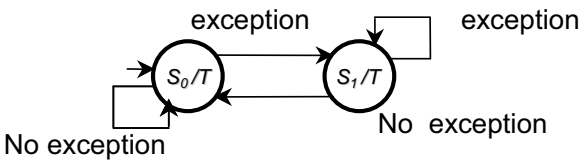

FIGURE 2 A model for workforce behaviour assuming a two-state target process and a canonical transaction

the local process state. Note that the system allows for a fresh transaction to be downloaded to the nodes at the beginning of each operation cycle. In the simplest case, performing an instance of the state computation transaction entails that each member in the corresponding workforce perform a sensing operation and formulate a node report. A specific member of the workforce, designated as a transaction instance manager, then receives all node reports, and formulates a *Transaction Instance Report (TIR)*. The TIR is the encoding of the local process state of interest at the time. This TIR is subsequently transported to a sink for storage. In principle, after transmitting the TIR the corresponding workforce disbands. For simplicity, we assume that at most one transaction instance is in progress in a given cluster at any given point in time.

The state diagram in Figure 2 represents the behavior of an arbitrary workforce in a cluster assuming the special case of a target process that has one normal state, $s0$, and one exception state, $s1$. T in the figure denotes the state computation transaction.

Two important design parameters in the above transaction-based model are *workforce size*, and *workforce setup*. These are discussed in the following.

Workforce size. This is application dependent. Applications negotiate QoS requirements, for example, the 'confidence level' associated with TIRs. These QoS parameters are, in turn, mapped to constraints at the transaction level. A constraint we assume in the anonymity solution proposed in this paper is that the average workforce size must be larger than or equal to λ, where λ is derived from application-level QoS parameters.

Workforce setup. This is governed by criteria such as load balancing or energy conservation, and thus represents a system concern. We distinguish two basic approaches to workforce setup, *static* and *dynamic*. The dynamic approach trades dynamic load balancing at the node level, and robustness to node failures for added delay time and energy overhead of the setup protocol. The static approach trades savings in energy and time of static setup for dynamic load balancing at the node level and robustness to node failures. In our proposed anonymity solution, we organize the node population of a cluster, prior to each operation cycle, into a fixed number of static workforces that have the same size on the average. Workforces are tasked to work (do an instance of the state computation transaction) in a round-robin fashion. The goal of our approach is to eliminate the energy and time overhead of dynamic setup while achieving load balancing at the coarser granularity of a workforce.

4.2 Inter-cluster Activity

As indicated earlier, the goal of inter-cluster activity in our work model is to route TIRs from their clusters of origin to the sinks, by means of multi-hop communication. We define a hop in a route as a direct transmission from one cluster to a *neighbor* cluster. A cluster u_j is a neighbor of a cluster u_i if and only if both u_j and u_i are located in the same corona, or the same wedge, for all i and j, $i \neq j$. It follows that in any instance of the

coordinate system defined in Sec. 3, each cluster has either three or four neighbors. The set of all neighbors of a cluster u_i is called the *neighborhood* of u_i. In the our proposed anonymity solution we define a distributed inter cluster routing protocol that yields optimal routes in terms of the number of hops from source to sink, and is highly scalable in the number of clusters in the network. Scalability can be attributed to two characteristics of the protocol. First, the protocol uses *no* dynamic global or regional state information. This eliminates the need for control messages to support routing. Second, the protocol uses distributed (incremental) route computation; the destination of hop i, computes, in turn, the destination of hop $i+1, i \geq 1$. In the remainder of this paper we assume the availability of a MAC layer that supports inter-cluster communication.

5. The Anonymity Problem

In this section we formulate a definition for the anonymity problem addressed in this paper. The problem is defined in the context of the network system described earlier. First, we introduce an anonymity threat model.

5.1 The Anonymity Threat Model

The threat model assumed in this paper emanates from a data-centric view of the sensor network. The model is predicated on the assumption that *the end-goal of anonymity attacks on the sensor network is to identify and eliminate the minimum number of nodes to inflict maximum loss of data assets*; eliminating a node means disabling that node so that it is permanently non operational. In our sensor network model, TIRs are the data assets of import to the end user (application). For any operation cycle, if a sink suffers a permanent failure before transferring the contents of its TIR repository to the gateway, then a portion of the data assets corresponding to the cycle is irrevocably lost. Therefore, sinks, or nodes comprising them to be precise, are the assumed targets of anonymity attacks in our model. Specifically, the goal of the adversary system is to eliminate all sink nodes. There are two main approaches to eliminating sink nodes:

Brute-Force (Sink nodes are not identified). This may take the form of randomly eliminating nodes in the network on the assumption that, statistically, some sink nodes will be eliminated in the process. Coarse sink granularity, and sink redundancy mitigate the risk of loss of data assets as a result of this type of attack. A trivial special case is the massive elimination of all nodes in the network.

Smart (Sink nodes are identified). The adversary system analyzes network traffic to deduce information about network topology, traffic flow patterns, and other system attributes. The goal is to discover, i.e. compromise the anonymity of, sink nodes, and, hence, eliminate them.

In this paper we assume the adversary system engages in smart elimination attacks. The specifics of the architecture and the implementation of the adversary system are assumed unknown.

5.2 Terminology and Notation

The main goal of this subsection is to establish terminology and notation that will be used in our solution to the anonymity problem.

Trudy	Denotes the adversary system
m	A message transmitted in the sensor network system post deployment. (Note that the transmitter is either a sensor node or the gateway.)
ts(m)	A global time stamp assigned by Trudy to message *m*,
l(m)	The unique location, in an arbitrary coordinate system used by Trudy, of the transmitter of message *m* in the 2-dimensional deployment plane.
r	^{node},$r^{gateway}$ The nominal transmission radius of a sensor node, and the gateway, respectively
cvr(m)	This is the *cover* of message *m*. Specifically, *cvr(m)* is the set of nodes that are located in the circular area with radius r^{node}, and center *l(m)* in the deployment plane
trace(m)	The trace of message *m*; if *m* is routed along a path of length *g, trace(m)* is the sequence of messages $\left(m^{(1)},m^{(2)},\cdots,m^{(g)}\right)$, $m(^{i})$ is the retransmission of *m*, or an encryption thereof, over hop number *i*, $1 \le i \le g$. Note that $m^{(1)}$=*m*
source(m)	The transmitter of message *m*
destination(m)	The designated receiver of message *m*

The assumptions underlying our anonymity threat model can be summarized as follows:

A. **Pre-deployment**

- All nodes are trusted
- Nodes are in a secure environment
- Trudy does not have access to any message *m* transmitted in the system.

B. **Post deployment (training and operation cycles)**

- Trudy receives every message transmitted in the system. Note that receiving a message does not imply being able to interpret it. A message *m* transmitted in the system is represented in Trudy's system as follows:

$$\left(m,ts(m),l(m),cvr(m)\right)$$

5.3 Anonymity Problem Statement

Let *I* be an arbitrary time interval, that starts post deployment, and let the set $M=\left\{\left(m_i,ts(m_i),l(m_i),cvr(m_i)\right)\mid 1\le i\le h\right\}$be the set of all messages transmitted in the system (and hence recorded by Trudy) during the interval *I*. Also, let the coordinate system, *O*, established by training be given by {(r_0,r_1,…,r |*c*| –1)q}, and the set of all nodes located in sink clusters be denoted by *S*.

The anonymity problem (from Trudy's point of view) can thus be stated as follows:

Given:

$$M=\left\{\left(m_i,ts(m_i),l(m_i),cvr(m_i)\right)\mid 1\le i\le h\right\} \tag{1}$$

Find:

$$m_q\in M:(\exists m_p\in M:trace(m_p)=(m_p^{(1)},m_p^{(2)},\cdots,m_p^{(g_p)})\wedge m_q=m_p^{(g_p)}),g_p\ge 1 \tag{2}$$

Note that for the message m_q in (2), $destination(m_q) \in S$. In general, for each message m_q that satisfies (2), it follows that

$$\exists x : x \in cvr(m_q) \wedge x \in S$$

The challenge for the sensor network system is to devise training, intra cluster, and inter cluster protocols that minimize the probability that the anonymity problem stated in equations (1), and (2) is solved, for arbitrary *O,I,S,* and *M*. In this work we only look at the problem of providing anonymity during the training period.

6. Providing Anonymity for Sensor Network Training

The main goal of this section is to propose a protocol for training that addresses the anonymity problem formulated above. The primary goal of the training protocol is to establish the *canonical coordinate system,* O_s , for the network, anonymous to Trudy. For a given sensor network system, the canonical coordinate system is the instance of the polar coordinate system described in 3.4.1 that has the maximum precision. $O_s$defines the set of canonical coronas, c_s, and the set of canonical wedges, w_s; we assume that $|c_s|$ and $|w_s|$ are powers of 2. O_s is defined by $\{(r_0, r_1, ..., r_{|c_s|-1}), \theta_s\}$, $r_i = (i+1) \times \varepsilon_s, 0 \le i \le |c_s| - 1$, where ε_s and θ_s are, respectively, the smallest corona width, and the smallest wedge angle for the system; ε_s, *and* θ_s characterize the system precision.

Post training, the coordinate system used during any operation cycle is derived from the canonical coordinate system using three integer parameters, $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$. Here, $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, *and* $\boldsymbol{\gamma}$ represent, respectively, *wedge rotation*, *wedge grouping*, and *corona grouping* parameters. Let *C(x,O) and W(x,O)* denote, respectively, the corona, and the wedge where node *x* is located according to coordinate system *O*. For a given operation cycle, *e*, if the rotation and grouping parameters are $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$, then the coordinate system used for cycle *e* is defined as follows:

$$O_e = \{(r_0, r_1, ..., r_{|c_e|-1}), \theta_e\}, \textit{where } r_i = (i+1) \times \varepsilon_e, 0 \le i \le |c_e| - 1, \varepsilon_e = \gamma_e \times \varepsilon_s, \theta_e = \beta_e \times \theta_s$$

Note that $\alpha_e \in [0, |w_s| - 1], \beta_e \in [0, \log |w_s|]$, and $\gamma_e \in [0, \log |c_s|]$. The corona, and wedge of *x* according to O_e are defined as follows,

$$C(x, O_e) = C(x, O_s) \; div \; 2^{(\log |c_s| \; - \; \gamma_e)} \tag{3}$$

$$W(x, O_e) = ((W(x, O_s) + \alpha_e) \bmod |w_s|) div \; 2^{(\log |w_s| \; - \; \beta_e)} \tag{4}$$

In equation (3) above, γ_e determines the corona precision of the coordinate system O_e, the minimum precision (a single corona) corresponds to γ_e and the maximum precision (that of the canonical coordinate system) corresponds to $\gamma_e = \log |c_s|$. In equation (4) $\boldsymbol{\beta}_e$ determines the wedge precision of O_e in an analogous manner. The left hand operand of the *div* operator in (4) is a translation for $W(x,O_s)$ about the true anchor point an amount equal to $\boldsymbol{\alpha}_e$ canonical wedges.

We are now in a position to present the details of our proposed anonymity-compliant training protocol.

6.1 Preconditions

(i)Sensor nodes are randomly and uniformly deployed in a *deployment area*. The deployment area completely contains a circular area called the *network area*; nodes located in the network area will comprise our trained sensor network. The mission of the nodes that are located outside the network area is to generate fake message traffic to help keep the network area anonymous.

(ii)The gateway is mobile

(iii)Pre deployment, the following is loaded into each sensor node:

a. Secret key, K_{master} used for decrypting training protocol messages from the gateway. Also, K_{master} is used to derive the keys K_{α}, K_{β}, and K_{γ}, as proposed in [12]. K_{α}, K_{β}, and K_{γ} are used to generate at random values for α, β, γ, respectively, for the successive operation cycles.

b. The parameters q_s, ε_s, D, and $r^{gateway}$. D is the diameter of the network area, we assume $r^{gateway} \gg D$.

(iv)Internal clocks in all sensor nodes are synchronized to the gateway.

6.2 Training Protocol (gateway side)

(i)Do a random traversal of the deployment area, visiting a set of random anchor points $A = \{a_1, a_2, ..., a_A, \}$, For each $a_i, 1 \le i \le |A|, do$:

i.1. *Transmit a call-for-training message*: Transmit an omni directional broadcast message using $r^{gateway}$. The message is encrypted by K_{master} and contains a Boolean flag f that identifies a_i as either the true anchor point or a false anchor point; the true anchor point is the geographical center of the circular network area.

ii.2. *Corona train:* using the sink side of the training protocol described in [67] do corona training to establish coronas for O_s, such that $|c_s| = \dfrac{r^{gateway}}{\varepsilon_s}$

iii.3. *Wedge train:* using the analogous sink side of the wedge training protocol described in [67], do wedge training to establish the wedges for O_s.

(ii) *Terminate the protocol*

Note that the corona training done in step i.2 covers an area considerably larger than the network area. This means that corona training, in this case, defines fake coronas that lie outside the network area (i.e. at a distance more than the diameter D from the true anchor point). The objective is to help keep the diameter D unknown. Because sensor nodes know D, each node, after learning the canonical corona it is located in, can determine if it is located in the network area, and hence, in the trained network. The multiplicity and randomness of anchor points help keep the true anchor point unknown.

6.3 Training Protocol (node side)

In the following assume node x is the node executing the protocol.

(i)*Compute* $|c_s| = r^{gateway} / \varepsilon_s, |w_s| = 2\pi / \theta_s$

(ii)*Do forever*:

i.1. *Receive the next call-for-training message*: receive and decrypt, using K_{master}, the next call-for-training message, *m*.

ii.2. *If the Boolean flag f in m is true*, then do:

ii.2.1. *HGet corona trained*: invoke the node side of the training protocol described in [67] to get corona trained to learn the canonical corona number you are located in, $c(x,O_s)$.

ii.2.2. *Compute your corona radius:* Compute $r = (C(x,O_s)+1)\times \varepsilon_s$. (Note that if $r \le D$ then you know you are located in the network area, and thus will be a node in the trained network, otherwise you know you do not belong to the trained network.)

ii.2.3. *If you belong to the trained network*, do:

ii.2.3.1. *Compute* $|c_s| = D/\varepsilon_s$

ii.2.3.2. *Get wedge trained:* invoke the node side of the training protocol described in [20] to get wedge trained to learn the canonical wedge number you are located in, $W(x,O_s)$.

ii.2.3.3. Using K_α, K_β, *and* K_γ generate via a random number generator (or a preloaded CBC block as described in [12]), respectively, the parameters $\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1, \boldsymbol{\gamma}_1$ for the first operation cycle.

ii.2.3.4. *Compute*

$$C(x,O_1) = C(x,O_s)\ div\ 2^{(\log|c_s| - \gamma_1)}, \text{ and}$$

$$W(x,O_1) = ((W(x,O_s)+\alpha_1)\bmod |w_s|)\ div\ 2^{(\log|w_s| - \beta_1)}$$

ii.2.3.5. *Terminate the protocol.*

Else

Sleep for $|w_s|$ message times; *terminate the protocol*.

Else

Sleep for $|c_s|+|w_s|$ message times.

7. Concluding Remarks and Directions for Future Work

It is widely recognized that sensor network research is in its infancy. In particular, there is precious little known about how to get sensor networks to self-organize in a way that maximizes the operational longevity of the network and that guarantees a high level of availability in the face of potential security attacks. However, given the characteristics of sensor networks, anonymity protocols developed for wired, cellular, or ad-hoc networks do not apply.

We view this paper as an initial contribution towards developing a lightweight solution for the anonymity problem in wireless sensor networks. Research is underway toward a solution that provides security not only for the various individual layers of the system but also for the entire system in an integrated fashion.

References

1. I. F. Akyildiz, W. Su, Y. Sankarasubramanian, and E. Cayirci, "Wireless sensor networks: A survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, Month 2002.
2. R. Anderson, and M. Kuhn, "Tamper resistance a cautionary note," *Proceedings of the Second Usenix Workshop on Electronic Commerce*, 1996, pp. 1–11.
3. G. Ateniese, A. Herzberg, H. Krawczyk, and G. Tsudik, "Untraceable mobility or how to travel incognito," *Computer Networks*, vol. 31, no. 8, pp. 871–884, Month 1999.

4. Bahl P., W. Russell, Y. M. Wang, A. Balachandran, G. M. Voelker, and A. Miu, "PAWNs: satisfying the need for ubiquitous secure connectivity and location services," *IEEE Wireless Communications*, vol. 9, no. 1, pp. 40–48, Month 2002.
5. S. Basagni, K. Herrin, E. Rosti, and D. Bruschi, "Secure pebblenets,"*Proceedings of MobiHoc*, 2001.
6. S. Basagni, M. Mastrogiovanni, and C. Petrioli, "A performance comparison of protocols for clustering and backbone formation in large scale ad hoc networks," *Proceeding of IEEE MASS*, 2004.
7. O. Berthold, H. Federrath, and M. Köhntopp, "Project anonymity and unobservability in the Internet," *Proceedings of Computers Freedom and Privacy (CFP 2000), Workshop on Freedom and Privacy by Design*, 2000.
8. A. Beresford, and F. Stajano, "Location privacy in pervasive computing," *Pervasive Computing*, vol. 2, no. 1, pp. 46–55, Month 2003.
9. S. Capkun, J.-P. Hubaux, and M. Jakobsson, "Secure and privacy-preserving communication in hybrid ad hoc networks," EPFL-IC Tech. Rep. IC/ 2004/10, 2004.
10. M. Cardei, and D. Z. Du, "Improving wireless sensor network lifetime through power aware organization," *ACM Wireless Networks,*, in press. .
11. D. W. Carman, P. S. Kruus, and B. J. Matt, "Constraints and approaches for distributed sensor network security," Tech. Rep. 00-010, NAI Labs, 2000.
12. D.W. Carman, B.J. Matt, and G.H. Cirincione, "Energy-efficient and low-latency key management for sensor networks.," *Proceedings of the twenty-third. Army Science Conference*, 2002.
13. H. Chan, and A. Perrig, "ACE: An emergent algorithm for highly uniform cluster formation," *Proceedings of the First European Workshop on Sensor Networks (EWSN)*, LNCS 2920, January 2004.
14. H. Chan, A. Perrig, and D. Song, "Random key pre-distribution schemes for sensor networks," *Proceedings of IEEE Symposium on Security and Privacy*. Oakland, CA, 2003.
15. D. Culler, D. Estrin, and M. Srivastava, "Overview of sensor networks," *IEEE Computer*, vol. 31, no. 8, pp. 41–49, Month 2004.
16. K. A. Delin, and S. P. Jackson, "The sensor web: A new instrument concept," *Proceedings of SPIE Symposium on Integrated Optics*, 2001.
17. R. DiPietro, L. V. Mancini, and S. Jajodia, "Providing secrecy in key management protocols for large wireless sensor networks," *Journal of AdHoc Networks*, vol. 1, no. 4, pp. 455–468, 2003.
18. R. DiPietro, L. V. Mancini, and A. Mei, "Random key assignment for secure wireless sensor networks," *Proceedings of the First Workshop on Security in Ad Hoc and Sensor Networks*, 2003.
19. W. Du, J. Deng, Y. S. Han, S. Chen, and P. Varshney, "A key management scheme for wireless sensor networks using deployment knowledge," *Proceedings of IEEE INFOCOM 04*, 2004.
20. DUST™ Networks, M1010 mote http://www.dust-inc.com/pdf/M1010-Mote.pdf (accessed 02/03/05).
21. A. Ephremides, J. Wieselthier, and D. Baker, "A design concept for reliable mobile radio networks with frequency hopping signaling," *Proceedings of the IEEE*, vol. 75, no. 1, pp. 56–73, 1987.
22. L. Eschenauer, and V. Gligor, "A key management scheme for distributed sensor networks," *Proceedings of the Ninth ACM Conference on Computer and Communications Security*, 2003.
23. D. Estrin, D. Culler, K. Pister, and G. Sukhatme, "Instrumenting the physical world with pervasive networks," *Pervasive Computing*, vol. 1, no. 1, pp. 59–69, Month 2002.
24. D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," *Proceedings of MOBICOM*, 1999.
25. M. Golumbic*, Algorithmic graph theory and perfect graphs.* New York: Academic Press, 1980.
26. D. Gracanin, M. Eltoweissy, S. Olariu, and A. Wadaa, "On modeling wireless sensor networks," *Proceedings of the IEEE Workshop on Mobile Ad Hoc and Sensor Networks*, 2004.
27. D. Gracanin, M. Eltoweissy, S. Olariu, and A. Wadaa, "Extensible wireless sensor networks," *Journal of Wireless Networks*, in press, 2005.

28. B. Hemingway, W. Brunette, T. Anderl, and G. Boriello, "The flock: Mote sensors sing in undergraduate curriculum," *IEEE Computer*, vol. 37, no. 8, pp. 72–78, 2004.
29. Y.-C. Hu, D. B. Johnson, and A. Perrig, "SEAD: Secure efficient distance vector routing in mobile wireless ad hoc networks," *Proceedings of the Fourth Workshop on Mobile Computing Systems and Applications*, 2002.
30. J. Hubaux, L. Buttyan, and S. Capkun, "The quest for security in mobile ad hoc networks," *Proceedings of MobiHoc*, 2001.
31. K. Jones, A. Wadaa, S Olariu, L. Wilson, and M. Eltoweissy, "Towards a new paradigm for securing wireless sensor networks," *Proceedings of the Workshop on New Security Paradigms*, 2003.
32. C. Karlof, and D. Wagner, "Secure routing in sensor networks: attacks and countermeasures," *Proceedings of the First IEEE International Workshop on Sensor Network Protocols and Applications*, 2003.
33. J. Kong, X. Hong, and M. Gerla, "An anonymous on demand routing protocol with untraceable routes for mobile ad hoc network," UCLA Computer Science Tech. Rep. 030020, 2003.
34. B. Lampson, "Computer security in the real world," *IEEE Computer*, vol. 37, no. 6, pp. 37–48, Month 2004.
35. P. Levis, N. Lee, M. Welsh, and D. Culler, "TOSSIM: Accurate and scalable simulation of entire tinyOS applications," *Proceedings of the First International Conference on Embedded Networked Sensor Systems*, 2003.
36. D. Liu, and P. Ning, "Establishing pairwise keys in distributed sensor networks," *Proceedings of the Tenth ACM Conference on Computer and Communications Security*, 2003.
37. D. Liu, and P. Ning, "Location-based pairwise key establishments for static sensor networks," *Proceedings of ACM Workshop on Security in Ad Hoc and Sensor Networks*, 2003.
38. D. Liu, and P. Ning, "Multi-level μTESLA: A broadcast authentication system for distributed sensor networks," *ACM Transactions in Embedded Computing Systems*, 2004.
39. S. Marti, T. Giuli, K. Lai, and M. Baker, "Mitigating routing misbehavior in mobile ad hoc networks," *Proceedings of ACM MOBICOM*, 2000.
40. K. Martinez, J. K. Hart, and R. Ong, "Environmental sensor networks," *IEEE Computer*, vol. 37, no. 8, pp. 50–56, Month 2004.
41. A. J. Menezes, P. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography.* CRC Press: Boca Raton, 1996.
42. R. Molva, D. Samfat, and G. Tsudik, "Authentication of mobile users," *IEEE Network*, vol. 8, no. 2, pp. 26–34, Month 1994.
43. G. Montenegro, and C. Castelluccia, "Statistically unique and cryptographically verifiable (SUCV) identifiers and addresses," *Proceedings of the Symposium Network and Distributed System Security*, 2002.
44. National Research Council,, *Embedded Everywhere.* Washington, DC: National Academies Press, 2001.
45. J. Newsome, R. Shi, D. Song, and A. Perrig, "The sybil attack in sensor networks: Analysis and defenses," *Proceedings of IEEE International Conference on Information Processing in Sensor Networks*, 2004.
46. S. Olariu, A. Wadaa, L. Wilson, and M. Eltoweissy, "Wireless sensor networks: leveraging the virtual infrastructure," *IEEE Network*, vol. 18, no. 4, pp. 51–56, Month 2004.
47. S. Olariu, and Q. Xu, "A virtual infrastructure for massively deployed sensor networks," *Computer Communications*, in press, 2005.
48. P. Papadimitratos, and Z. J. Haas, "Secure routing for mobile ad hoc networks," *Proceedings of Communication Networks and Distributed Systems*, 2002.
49. S. Park, I. Locher, A. Savvides, M. B. Srivastava, A. Chen, R. Muntz, and S. Yue, "Design of a wearable sensor badge for smart kindergarten," *Proceedings of the Sixth International Symposium on Wearable Computers*, 2002.
50. M. R. Pearlman, Z. J. Haas, P. Sholander, and S. S. Tabrizi, "On the impact of alternate path routing for load balancing in mobile ad hoc networks," *Proceedings of MobiHoc*, 2000.

51. A. Perrig, R. Canetti, D. Tygar, and D. Song, "The TESLA broadcast authentication protocol," *RSA CryptoBytes*, vol. 5, no. 2, pp. 2–13, Month 2002.
52. A. Perrig, R. Szewczyk, V. Wen, D. Culler, and J. D. Tygar, "SPINS: Security protocols for sensor networks," *Wireless Networks*, vol. 8, no. 5, pp. 521–534, 2002.
53. A. Perrig, J Stankovic, and D. Wagner, "Security in wireless sensor networks," *Communications of the ACM*, vol. 47, no. 6, pp. 53–57, Month 2004.
54. Pfitzmann , and Köhntopp M., "Anonymity, unobservability, and pseudonymity — a proposal for terminology," In H. Federrath ed., *DIAU'00 LNCS 2009,* 2000, pp. 1–9.
55. J. Polastre, R. Szewcyk, A. Mainwaring, D. Culler, and J. Anderson, "Analysis of wireless sensor networks for habitat monitoring," In Raghavendra Sivalingam, & Znati eds., *Wireless Sensor Networks.* Kluwer Academic, 2004, pp. 399–423.
56. B. Przydatek, D. Song, and A. Perrig, "SIA: Secure information aggregation in sensor networks," *Proceedings of the First ACM Conference on Embedded Networked Sensor Systems*, 2003.
57. M. G. Reed, P. F. Syverson, and D. M. Goldschlag, "Anonymous connections and onion routing," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 4, 1998.
58. M. K. Reiter, and A. D. Rubin, "Crowds: Anonymity for web transactions," *ACM Transactions on Information and System Security*, vol. 1, no. 1, pp. 66–92, 1998.
59. K. Ryokai, and J. Cassell, "StoryMat:A play space for collaborative storytelling," *Proceedings CHI'99*, 1999.
60. P. Saffo, "Sensors, the next wave of innovation," *Communications of the ACM*, vol. 40, no. 2, pp. 93–97, Month 1997.
61. D. Samfat, R. Molva, and N. Asokan, "Untraceability in mobile networks," *Proceedings of ACM MOBICOM*, 1995.
62. C. Shields, and B. N. Levine, "A protocol for anonymous communication over the Internet.," *Proceedings of the ACM Conference on Computer and Communications Security*, 2000.
63. K. Sohrabi, J. Gao, V. Ailawadhi, and G. Pottie, "Protocols for self-organization of a wireless sensor network," *IEEE Personal Communications*, vol. 7, no. 5, pp. 16–27, Month 2000.
64. K. et al. Sohrabi, "Methods for scalable self-assembly of ad hoc wireless sensor networks," *IEEE Transactions on Mobile Computing*, vol. 3, no. 4, pp. 317–331, 2004.
65. M. Srivastava, R. Muntz, and M. Potkonjak, "Smart Kindergarten: Sensor-based wireless networks for smart developmental problem-solving environments," *Proceedings of ACM MOBICOM*, July 2001.
66. R. Szewczyk, J. Polastre, A. Mainwaring, J. Anderson, and D. Culler, "An analysis of a large scale habitat monitoring application," *Proceedings of the Second ACM Conference on Embedded Networked Sensor Systems*, 2004.
67. A. Wadaa, S. Olariu, L. Wilson, M. Eltoweissy, and K. Jones, "Training a wireless sensor network," *Mobile Networks and Applications*, vol. 10, pp. 151–167, Month, 2005.
68. A. Wadaa, S. Olariu, L. Wilson, and M. Eltoweissy, "Scalable key management for secure communications in wireless sensor networks," *Proceedings of the International Workshop on Wireless Ad-hoc Networking*, 2004.
69. L. Wang, and S. Olariu, "A unifying look at clustering in mobile ad hoc networks," *Wireless Communications and Mobile Computing*, vol. 4, pp. 623–637, Month 2004.
70. B. Warneke, M. Last, B. Leibowitz, and K. Pister, "SmartDust: communicating with a cubic-millimeter computer," *IEEE Computer*, vol. 34, no. 1, pp. 44–55, Month 2001.
71. A. D. Wood, and J. A. Stankovic, "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 4, pp. 54–62, Month 2002.
72. T. Yan, T. He, and J. A. Stankovic, "Differentiated surveillance for sensor networks," *Proceedings of ACM SenSys*, November 2003.
73. H. Yang, and S. Lu, "Self-organized network layer security in mobile ad hoc networks," *Proceedings of the FirstACM Workshop on Wireless Security*, 2002.
74. P. Yau, and C. J. Mitchell, "Security vulnerability in ad hoc networks," *Proceedings of the Seventh International Sysmposium on Communication Theory and Applications*, 2003.

75. F. Ye, H. Luo, S. Lu, and L. Zhang, "Statistical en-route filtering of injected false data in sensor networks," *Proceedings of IEEE INFOCOM 2004*, 2004.
76. V. V. Zhirnov, and D. J. C. Herr, "New frontiers: self-assembly and nano-electronics," *IEEE Computer*, vol. 34, no. 1, pp. 34–43, Month 2001.
77. L. Zhou, and Z.J. Haas, "Securing ad-hoc networks," *IEEE Network*, vol. 13, no. 6, pp. 24–30, Month 1999.
78. S. Zhu, S. Setia, and S. Jajodia, "LEAP: Efficient security mechanisms for large-scale distributed sensor networks," *Proceedings of the ACM Conference on Computer and Communications Security*, 2003.
79. S. Zhu, S. Setia, S. Jajodia, and P. Ning, "An interleaved hop-by-hop authentication scheme for filtering false data in sensor networks," *Proceedings of IEEE Symposium on Security and Privacy*, 2004.

Dr. Olariu received the M.Sc. and Ph.D. degrees in computer science from McGill University, Montreal in 1983 and 1986, respectively. In 1986 he joined Old Dominion University, where he is a Professor of Computer Science. Dr. Olariu has published extensively in various journals, book chapters, and conference proceedings. His research interests include image processing and machine vision, parallel architectures, design and analysis of parallel algorithms, computational graph theory, computational geometry, and mobile computing. Dr. Olariu serves on the Editorial Board of IEEE Transactions on Parallel and Distributed Systems, Journal of Parallel and Distributed Computing, VLSI Design, Parallel Algorithms and Applications, International Journal of Computer Mathematics, and International Journal of Foundations of Computer Science.

M. Eltoweissy is a (visiting) Professor of Computer Science at Virginia Tech. He is also a Professor of Computer Science at James Madison University. Eltoweissy's research interests include information security and privacy, wireless sensor and ad hoc networks, network security, computer-supported cooperative work, and distributed computing. He published extensively in books, refereed journals, and conference proceedings. He also served on numerous technical committees for conferences, workshops, seminars, and NSF panels. He has an aggressive record of funding (over $10 million). Eltoweissy founded the Commonwealth Information Security Center (CISC) in Virginia and was a founding member of the award-winning Virginia Alliance for Secure Computing and Networking (VA SCAN). Eltoweissy earned a Ph.D. in Computer Science (1993) from Old Dominion University, and MS (1989) and B.S. (1986 with top rank) in Computer Science and Automatic Control from Alexandria University, Egypt. He was nominated by JMU for the Virginia Statewide Outstanding Faculty Awards in 2003. Eltoweissy is a member of ACM, ACM SIGSAC, and the honor societies of Phi Kappa Phi and Upsilon Pi Epsilon.

A.Wadaa is a Lecturer of Computer Science and a Ph.D. candidate at Old Dominion University. Wadaa's dissertation is on wireless sensor networks. His research interests include wireless sensor and ad hoc networks, network security, database systems, and distributed computing. He published in refereed journals and conference proceedings. He also served on a number of technical committees for conferences, and workshops. His funding record includes a grant from the Commonwealth Information Security Center (CISC) in Virginia to pursue research in sensor network security. Wadaa earned a B.S. ('86 with top rank) in Computer Science and Automatic Control from Alexandria University, Egypt. He was nominated by Old Dominion University for an outstanding teacher award in 2002. Wadaa is a member of the honor society of Phi Kappa Phi.

O. Xu is currently a PhD candidate in the Computer Science at Old Dominion University (ODU). His main research interests include wireless sensor network, distributed computing, and software engineering. He graduated from Wake Forest University with a Master of Science degree in Computer Science in 1999.

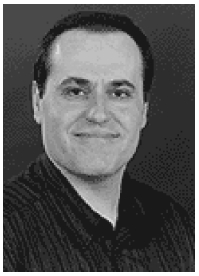

Albert Y. Zomaya is currently the *CISCO Systems Chair Professor of Internetworking* in the School of Information Technologies, The University of Sydney. Prior to that he was a Full Professor in the Electrical and Electronic Engineering Department at the University of Western Australia, where he also led the Parallel Computing Research Laboratory during the period 1990–2002. He served as Associate-, Deputy-, and Acting-Head in the same department, and held visiting positions at Waterloo University and the University of Missouri-Rolla. He is the author/co-author of 6 books, 200 publications in technical journals and conferences, and the editor of 6 books and 7 conference volumes. He is currently an associate editor for 14 journals, the Founding Editor of the *Wiley Book Series on Parallel and Distributed Computing*, and the Editor-in-Chief of the *Parallel and Distributed Computing Handbook* (McGraw-Hill, 1996). Professor Zomaya was the Chair for the *IEEE Technical Committee on Parallel Processing* (1999–2003) and currently serves on its executive committee. He has been actively involved in the organization of national and international conferences. He received the *1997 Edgeworth David Medal* from the Royal Society of New South Wales for outstanding contributions to Australian Science. In September 2000 he was awarded the *IEEE Computer Society's Meritorious. Service Award*. Professor Zomaya is a chartered engineer (CEng), a Fellow of the IEEE, a Fellow of the Institution of Electrical Engineers (U.K.), and member of the ACM. He also serves on the boards of two startup companies. His research interests are in the areas of high performance computing, parallel algorithms, and bioinformatics.