

# Sensor Cooperation in Human Environments through Motivational Gradients

Ladislau Bölöni and Damla Turgut  
School of Electrical Engineering and Computer Science  
University of Central Florida  
Orlando, FL, USA  
lboloni,turgut@eeecs.ucf.edu

**Abstract**—The urban environment of the early 21st century contains a large number of consumer devices under private and organizational ownership. Many of these devices contain sensors as well as communication devices. However, most of these sensors are only used for purposes internal to the device. By interconnecting these sensors we can obtain a network which can serve important societal goals. The technological challenges of interconnecting these sensors are relatively minor. The main problem is the human aspect: why would the owners of the sensors offer their readings for public use? In fact, privacy considerations might advise the exact opposite. Such networks will not be accepted unless every owner is motivated to allow the participation of its devices. In this paper we describe an architecture which enables such a system by the formal model of motivational gradients. The original source of motivational gradients are always humans or organizations; however, nodes acting as autonomous agents can negotiate motivational microgradients based on the original macrogradient introduced by humans. We evaluate the networking and computer-human interaction aspects of the proposed architecture.

**Index Terms**—sensor networks, societal goals

## I. INTRODUCTION

Modern consumer devices are already equipped with a variety of sensors; however, the readings of these sensors are used only in the local context of the host device. For instance, the light sensor in an alarm clock is only used to set the intensity of the display, the sensor-based light switches are only connected to the light, the multitude of sensors in modern cars are informing only the driver of the given vehicle. We argue that by interconnecting these sensors, they can serve *societal* benefits, in addition to their local uses. Firefighters would be able to find people trapped in their rooms. Public spaces can be automatically transformed into sensor-enhanced intelligent spaces for the benefits of the elderly – existing prototypes of such systems require custom deployment of sensors in a fully owned environment. Road, traffic, and weather information can be shared between vehicles, augmenting centralized information systems, or acting as a fallback in emergency situations such as hurricane evacuation. A car stuck on the railroad can automatically send emergency information to the approaching train.

Public access to privately owned sensors, however, would also rise the undesirable specter of stalkers using the sensors to spy on their victims, unscrupulous businesses abusing the system to profile their customers, or unrestricted, Big Brother type surveillance.

There are relatively few technological obstacles in establishing connectivity among the sensors; there are however, significant challenges on the human level. The critical question is *motivation*: why would the owners and operators of the devices allow them to participate in such a network? The relatively low deployment of ad hoc networks show that users do not willingly participate in networks without a clear reward. To maintain a functioning network all participants (information providers, consumers and intermediaries such as aggregators, anonymizers and localizers) must be motivated in their actions. We will say that every data transfer should follow a *positive motivational gradient*.

This motivation, however, needs to be propagated in an appropriate manner to each of the nodes. The system would work only if for every participant node the benefit of taking an action (transmitting a data packet, revealing a sensor reading or performing a database query) exceeds the cost of performing it. Naturally, such a cost benefit analysis can not be made individually for every packet, but the overall policy of the nodes need to conform to it.

Maintaining a positive motivational gradient is essentially an autonomous agent problem which can be attacked with means of game theory, decision theory, market models and other related techniques. However, it is also strongly tied to the networking protocols: certain routing and dissemination models might not be possible because they require some nodes to act against their own interests.

### A. Assumptions about the environment

The environment we consider is a natural evolution of the current urban environment. We assume that the majority of electronic devices in day-to-day use (including personal use devices, building automation and car electronics) will be equipped with communication devices. We will assume that these next generation devices host the same sensors as the current generation. The sensors are connected to the host devices, share its power source, and partake in its mobility. We assume that the devices are aware of their location and mobility. The sensor or the information consumer can not control the mobility of the device, although it might be able to predict it.

We assume that the sensors are owned and operated by the owner and the operator of the host device, respectively.

The owner and the operator might be a different person or organization (for instance, in case of a person operating a device belonging to an organization, or a rental device).

The default use of sensor readings is for purposes internal to the device. The sensor does not share any information to the outside world, unless it is motivated to do so. In this case, we say that the sensor plays the role of an *information provider*. The source of motivation is always at humans, and reflects the interests of the owner and operator.

We assume the existence of *information consumers*<sup>1</sup> which are interested in receiving information chunks specified in a *query*. The exchange of information in our environment is thus conditioned by the existence of interested information consumers, and sensors which act as information providers.

We have made a conscious effort to keep our assumptions conservative. We do not assume any kind of breakthroughs in hardware performance, size, energy consumption or cost. We do not assume a change in the way humans and organizations use technology. We do not assume significant regulatory changes - such as the opening of new frequency bands for public use. We do not assume significant public investment, such as building of new infrastructure. However, we assume that the sensors will use the existing public and private infrastructure, whenever available. We assume that the existing cost structure remains in place for network access.

#### B. The networking architecture in the proposed environment

The networking environment of interconnected devices shows a complex topology due to the heterogeneity of its technologies, the presence of mobile nodes, the mix of infrastructure-based and proximity connections, as well as the motivation gradient structure which limits the possible data flows.

Consumer devices use a mix of networking technologies, wired LAN connections, WiFi, WiMax, EDGE, 3G and other services provided by cellular telephony carriers, as well as short range wireless such as Bluetooth or Zigbee. Many devices support multiple networking technologies (many cell-phones and laptops support EDGE, WiFi and Bluetooth). Infrastructureless peer-to-peer connections are usually free. Internet access usually requires paid subscription (although exceptions such as free WiFi hotspots exist). More important is whether the internet access is *metered*, if yes, it can have an important impact to the motivation of the participating nodes.

The networking technologies (routing, addressing and so on) and information collection and query patterns are strongly interdependent; not every combination makes sense. To put some order in the complex topology of the network environment we reduce the number of access patterns the nodes might choose to implement to four major types:

**(A) Internet connected with central data repository.** The node offloads its entire set of observations to central data warehouse.

<sup>1</sup>We prefer this term, as the term "sink", used in most of the sensor network literature, is based on the vision of information converging to a single point, which is not the case of our environment.

**(B) Internet connected node with local data storage.**

The sensor node is connected to the internet, but manages its observations locally. The node needs to buffer its observations locally, advertise its availability and serve it on demand. The node needs to actively participate in the servicing of queries.

**(C) Islands of local connectivity.** The node is not connected to the internet, but participate in a local network through wired or wireless connection (e.g. the network formed by building automation systems).

**(D) Proximity-based connectivity.** The node forms connections with nearby nodes on ad hoc basis through technologies such as Bluetooth, ZigBee, computer-to-computer WiFi, or the upcoming vehicle-to-vehicle DSRC standard.

The choice of the operation mode depends on the owner's and operators of the devices. Connection type (A) has the advantage of transforming a networking problem into one of a database lookup, but can incur significant networking and data storage costs. Connections of type (D) are usually cost-free. Another argument for connections of type (D) is that proximity connections remain available even in the case of infrastructure failure. The failure of the central police dispatcher unit in New Orleans during Hurricane Katrina is a good example of such a situation. The police force, however, could still use their radios as point-to-point walkie-talkie-s.

Full internet connectivity has the advantage that in principle it allows data exchange between an arbitrary pair of nodes. Finding the node which can satisfy a certain query, however, is not trivial, especially if the nodes are mobile. In the case of connections of type A, the query can be performed on the data warehouse, without the need of a direct connection to the sensor node. In the case of connection type B external data indices are necessary for the routing of queries.

Paradoxically, in some cases proximity-based connectivity might actually simplify the query routing. For instance, let us consider the case of querying the road and traffic conditions on the road ahead. A data warehousing approach is clearly possible, by requiring all the cars on the road to maintain an internet connection and upload their readings in a central database, which in turn can be queried by an internet connected vehicle. However, it might be easier to query cars traveling in the opposite direction.

The remainder of this paper is organized as follows. In Section II we discuss related work. A formalism for motivations will be introduced in Section III. We discuss a possible extension in the use of market-based models in Section IV. The computer-human interaction issues raised by the proposed architecture is discussed in Section V. We conclude in Section VI

## II. RELATED WORK

The closest related work to our subject is the recently emerging area of sensor networks with mobile nodes. There are two major research directions, one directed towards mobile sinks, while the other one towards the link between mobility and routing.

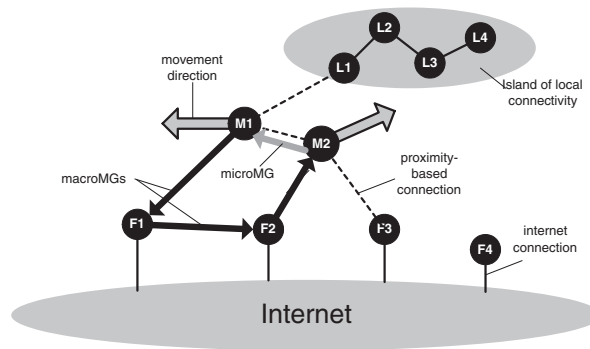


Fig. 1. The network environment of the proposed architecture.

**Mobility models of the sinks.** The mobility of the sink can be categorized into three types: random, predictable, and controllable. In case of *random mobility*, the sink travels through the network in a random walk fashion. In the case of *predictable mobility*, the sensor nodes can learn the mobility pattern of the sink and therefore can predict the location of the sink at any given point in time. In the case of *controlled mobility* the sink mobility is adaptively controlled based on specific parameters of the network and/or the deployed applications.

Tong et al. [12] propose an architecture, SENMA, which relies on one hop transmission between the sensor nodes and mobile agents. Wang et al. [19] conducts experiments with 256 homogeneous sensor nodes arranged in a square grid with a single mobile sink. Two-Tier Data Dissemination (TTDD), proposed by Luo et al. [7], protocol uses a grid structure in which only the sensors placed in the grid points are required to obtain information for forwarding towards multiple mobile sinks. Chakrabarti et al. [2] modeled the data collection process as a queueing system to measure the impact of predictable observer mobility (where the observers correspond to mobile sinks) using single-hop communication. Vincze and Vida [16] examines how the various sink mobility patterns affect the network lifetime by adaptively controlling the sink mobility to reduce the energy consumption with multi-hop communication between the sink and the sensors. More recent work by Vincze et al. [14], [15], [17] accommodates multiple sinks in densely deployed wireless sensor networks.

**Mobility and routing.** This category combines projects which consider not only the mobility of the sink, but also routing of the sensed data towards the sink. The protocols belonging this category employ multi-hop communication with mostly either predictable or controlled mobility.

The mWSN architecture, proposed by Chen et al. [3], uses multi-hop forwarding to form a cluster around the expected position of the mobile sink. mWSN is similar to Data MULEs in that it is a three-tier architecture, with the top tier composed of a base station, also called the final fusion point, the middle tier including mobile sinks such as mobile phones, laptops, and so on, while the bottom tier contains the static sensor nodes. The sink trajectory is not controlled but rather it

can be estimated or learned. A combination of base station mobility and multi-hop routing strategy are proposed by Luo and Hubaux [8] to maximize network lifetime. The MobiRoute architecture by Luo et al. [9], an extension of MintRoute [20], is a sensor network with mobile sinks where the mobility is controlled and predictable and the sinks have long pauses in their movement called *epochs*. Kansal et al. [5], [6] developed an Autonomous Intelligent Mobile Micro-Server (AIMMS) prototype which travels across the network to route data from the deeply embedded nodes. Nodes have to relay data only for the nodes which do not fall into the transmission range of the micro-server. Gandham et al. [4] deployed multiple mobile stations to extend the lifetime of the sensor network which is divided into equal periods of time known as *rounds*. Base stations are mounted on unmanned remote controlled vehicles to be moved from one location to another and they can be located only at specific places called “feasible sites”. Wang et al. [18] investigates various combinations of networks with mobile sinks and/or mobile relays. Olariu et al. [10], design ANSWER, an AutoNomouS netWorked sEnsoR system. The architecture assumes static sensor nodes and (possibly mobile) aggregation and forwarding nodes (AFNs) which can perform controlled mobility. An important role of the AFNs is to organize the sensors in their immediate vicinity into a dynamic virtual infrastructure which depends on the current task.

A specific challenge of sensor networks with mobile sinks is the *transmission scheduling problem*, the determination of the optimal moment for transmission while minimizing data loss as well of energy consumption [1], [13].

### III. A FORMAL MODEL OF MOTIVATION

Let us denote with  $\mathcal{N}$  the set of all nodes,  $\mathcal{S}$  the possible states of the environment and  $\mathcal{O}$  the possible observation chunks (limited by the expressivity of the query language). A *motivational gradient* (MG) is a quadruplet  $\langle P, C, F, T \rangle$  where

- $P \in \mathcal{P}(\mathcal{N})$  is the set of *provider nodes*
- $C \in \mathcal{P}(\mathcal{N})$  is the set of *consumer nodes*
- $F : \mathcal{O} \times \mathcal{S} \rightarrow \{\text{true}, \text{false}\}$  is the *filter function* and
- $T : \mathcal{O} \times \mathcal{S} \rightarrow \mathcal{O}$  is the *transformation function*.

The reading of this is that provider node  $A \in P$  in the state  $s \in S$  is motivated to send to any node  $B \in C$  observations  $o$  for which  $F(o, s)$  returns a value of true. The node will apply the transformation  $T$  to these observations before sending them, such that the consumer will receive the information chunk  $o' = T(o, s)$ .

At an encounter, the nodes need to authenticate their membership in the provider or consumer sets of an MG. Nodes in a certain group do not have a full knowledge of the other nodes in the group. An MG is not symmetric, but some nodes might be members of both  $P$  and  $C$ .

In practice, the transformation function must be a combination of simple components such as aggregation (sum, count, maximum, median), anonymization, watermarking or digital signatures. The simplest case is the *unit transformation*  $\forall o \forall s T_1(o, s) = o$ . If the transformation respects the formula  $\forall o \forall s \forall s' T(o, s) = T(o, s')$  we call it a *state-invariant transformation*.

#### Motivational microgradients

A motivational gradient is not directly “transitive”. Having a gradient for a certain observation  $o$  from  $A$  to  $B$  and from  $B$  to  $C$  does not necessarily imply that there is a gradient between  $A$  and  $C$ . Thus  $C$  can not directly query  $A$ , but it might do so through  $B$ . To understand why is this so, let us see what would direct transitivity entail. Let assume two motivational gradients  $G_1 = \langle P_1, C_1, F_1, T_1 \rangle$  and  $G_2 = \langle P_2, C_2, F_2, T_2 \rangle$ , with  $A \in P_1$ ,  $B \in C_1$ ,  $B \in P_2$ ,  $C \in C_2$ , and  $F_1(o) = F_2(o) = \text{true}$ . To keep simple, we will assume  $T_1 = T_2 = T_I$ . Obviously, none of this would imply  $C \in P_1$  or  $A \in C_2$ .

Despite the lack of transitivity, clearly in our case, the observation chunk  $o$  is passed from  $A$  to  $C$ . From a purely networking perspective, the transfer of the information chunk from  $A$  to  $B$  and then from  $B$  to  $C$  is wasteful.

We can transfer the data directly from  $A$  to  $C$  if we create a motivational gradient  $G_3 = \langle P_3, C_3, F_3, T_3 \rangle$  with  $P_3 = \{A\}$ ,  $C_3 = \{C\}$ ,  $F_3(x) = \{\text{true if and only if } x = o \text{ and } T_3 = T_I\}$ . We will call such a highly specific gradient a *motivational microgradient* ( $\mu\text{MG}$ ).  $\mu\text{MG}$ s are usually limited in time, and refer to very specific observations and small sets of participating nodes (potentially single nodes). They do not introduce flows of information which would not be possible with the original MGs.

One way to see this is that the MGs impose a logical network over the physical network determined by network connections. Two nodes which might be connected by a single MG hop might be many network hops away, and reversely, two nodes one network hop away might be able to communicate only through a long chain of MGs.  $\mu\text{MG}$ s introduce shortcuts in the logical network of MGs.

#### Negotiation of $\mu\text{MG}$ s

The  $\mu\text{MG}$  introduced above was a simplified example. In practice, microgradients need to account for non-unit transformation functions, and need to cover more than specific chunks of information. We illustrate the challenges of  $\mu\text{MG}$

negotiation through the slightly more complicated example of a *single node bypass  $\mu\text{MG}$* .

Lets assume two general form MGs  $G_1 = \langle P_1, C_1, F_1, T_1 \rangle$  and  $G_2 = \langle P_2, C_2, F_2, T_2 \rangle$ . An information chunk  $o_{\text{start}}$  is transferred from  $A$  to  $B$  and from  $B$  to  $C$ , arriving to  $C$  in the form  $o_{\text{final}}$ .  $G_3 = \langle P_3, C_3, F_3, T_3 \rangle$  is a *single node bypass  $\mu\text{MG}$*  if the following conditions are satisfied:

- $P_3 \subset P_1$  and  $C_3 \subset C_2$ .
- $F_3(o, s_1 \cup s_2) = \{\text{true if and only if } F_1(o) = \text{true}\}$  and  $F_2(T_1(o, s_1), s_2) = \text{true}$ .
- $o_{\text{final}} = T_3(o, s_1 \cup s_2) = T_2(T_1(o, s_1), s_2)$

Indeed, a  $\mu\text{MG}$  set up like this would guarantee that the information received at node  $C \in D_3$  from  $A \in S_3$  through the  $\mu\text{MG}$  would be the same as the information received through the  $A \rightarrow B \rightarrow C$  path.

There are several challenges in the creation of such a  $\mu\text{MG}$ . Note first, that both the filter and the transformation function of the second MG is collapsed into the  $\mu\text{MG}$ , and it needs to be performed by the first node. These functions, however, depend on the state of the second node in the general case. There are cases when this difficulty does not exist: for instance, if the  $F_2$  and  $T_2$  functions are state-invariant. In other cases, the subset of the state which affects these functions might be publicly available (such as the time of the day, whether there is a hurricane warning in effect and so on). Finally, the necessary data can be transferred from  $B$  to  $A$  if there is a MG which allows for it (note that neither  $G_1$  nor  $G_2$  is sufficient for this). There are cases when no bypass  $\mu\text{MG}$  can be negotiated. In other cases its cost in terms of secondary data transfers outweighs its benefits.

The negotiation of  $\mu\text{MG}$ s is a critical part of the model. We need to develop a negotiation method which is provably correct and computationally and communicationally cheap. The triggers of the negotiation of a  $\mu\text{MG}$  are the needs of the routing, dissemination and query processing protocols. The same set of MGs can yield many different  $\mu\text{MG}$ s. One possible objective of  $\mu\text{MG}$  negotiation might be to negotiate the minimum extent  $\mu\text{MG}$  which satisfies a given query.

When queries are frequent, however, it might be more efficient to negotiate  $\mu\text{MG}$ s which satisfy more than one query. Under some conditions, the nodes might anticipate future queries and negotiate  $\mu\text{MG}$ s preemptively.

#### Cryptographical aspects of motivational gradients

The architecture described opens several problems of security and privacy. To clearly outline the burden of proof, let us investigate the questions of rights and obligations. Participating in the consumer group of a motivational gradient is a right, while being in the provider group is an obligation. These roles determine the type of security challenges the system needs to solve:

- Authenticating that the node is in the customer set of a certain motivational gradient, which conveys rights to query certain types of data. The goal of this authentication step is to prevent unauthorized users access to the specific information.

- Authenticating that the node is in the provider group of a certain MG. This is necessary to avoid attacks through which the node can receive fake information.
- Non-repudiation of the fact that the node is in the provider group of certain MGs.
- Nodes should not be able to anoint new members of their own group, thus their certificates would be valid only with their unique private key.
- Memberships in MG sets should be revokable.
- State of the world signals, such as fire, terrorism alert, hurricane alert and so on, should be source-authenticated, and the authorization of the issuer to issue such a signal validated. Similarly for the canceling of such signals.

#### IV. MOTIVATION AND FINANCIAL INCENTIVES

A possible way to spread motivation in the system would be through a market-based model, that is, by introducing a financial incentive. It is natural that some of the macroMGs will be financially motivated (such as an insurance discount). A logical evolution of this model is that some of this financial incentive needs to be transferred down to the microgradients originating from the specific macrogradient.

From a game theoretic point of view, this is a coalition game, with the objective of finding a mutually acceptable distribution of rewards.

However, at the  $\mu$ MG level, a market based model would require (a) a micropayment clearing system and (b) the permission of the users that their consumer devices negotiate on their behalf over financial instruments (even if the sums involved are very small). Previous micropayment systems failed to take off due to human factors.

##### A. Mechanism design

Societal interest in the smooth operation of the system might encourage organizations (state, city, companies) to introduce sponsored devices with specific MGs to change the balance of motivations. For instance, an organization might introduce of sponsored anonymizers, localizers, the attachment of information ferries to existing devices with regular movement patterns and so on. These devices might have motivations which were determined based on *mechanism design* considerations [11], such that the pursuit of the self interest in the remainder of the nodes will yield the desired societal benefits.

##### B. Minimizing cost. Various cost definitions

All the actions of the nodes, including sensing, data transmission and microgradient negotiation incurs some cost. In general, whenever a query needs to be answered, we are always looking for the minimum cost query processing method which can still answer the query. Cost can be defined in terms of many resources, including energy, network bandwidth, or even the actual monetary cost of the network connection. In traditional sensor networks, minimum cost usually refers to performance (eg. minimum number of hops) or energy. In our case energy is a smaller problem, as the sensors are sharing the power source of the embedding device, which is renewable and often large

compared to the needs of the sensor (for instance, in the case of sensors embedded in cars).

One of the most important cost measures in our environment is defined in terms of “minimum disturbance”, which is related to the attention span of the nodes and the routers. An individual or organization will try to answer the query through the sensors which it owns, followed by public organizations. This applies both to the actual collecting as well as the aggregation and routing of data. We will assume that certain type of information has a certain utility. Nodes will collect information subject of its utility. The utility of information depends of its age, location, accuracy, level of aggregation.

This allows the nodes to make data collection decisions based on the perceived utility of the information. Note that the utility of certain types of information is related to but not the same as a motivation. Utility mostly informs the data collection decisions. Some information might have utility even if it is useful only to the node. If the node itself is not interested in the information, and there is no motivation to transmit it to anyone, its utility will be zero: the node will not collect such information. The motivations which match the incoming information are one of the components of utility, this, however, needs to be multiplied with the predicted likelihood of a query on the given information piece.

In the general case, it is very difficult to predict the utility of the information for other nodes. Under some conditions, nearby nodes can *prime* the sensor node, notifying it that certain types of queries are likely be forthcoming. Priming gives a temporary boost to the utility of certain type of information, prompting the node to give it priority in recording and buffering it.

#### V. COMPUTER-HUMAN INTERFACE ASPECTS

The source of macroMGs is always the human user. Such systems can be deployed only if they are accepted by the user. This creates its own computer-human interface challenges. For the work proposed here we are primarily interested in investigating the information flow between the the user and the device, which needs to be as small as possible and it has to be expressed in terms understandable to the user. If understanding the behavior of the device puts an undue cognitive load on the user, the system will not be accepted. We will not focus on the design of the user interface itself, although that is, of course, important and need to be resolved at the implementation of consumer devices.

To illustrate the concerns in the specification of the macroMGs, let us consider the case of a consumer devices loaded with an assortment of macroMGs.

**Number of macroMGs.** The number of macroMGs should allow a “single glance” overview of the device. One way to achieve this might be to group macroMGs into groups which can be identified by their source or justification. Here is an example of what a user might see when it inspects the “Motivations” screen of its device.

- macroMG allowing free communication among all the devices belonging to the same user.

- macroMG certifying the user's participation into the North Orlando Traffic Event Notification Self-Help Association.
- macroMG required by an insurance company to qualify for a certain discount.
- macroMGs certified by Red Cross and installed voluntarily by the user (eg. elderly assist).
- macroMGs mandated legislatively by Homeland Security (eg. assisting special forces in case of terrorism alert).

**Specification of the destination group.** At the level of computers, the specification of the destination group is done through an authentication procedure, which ultimately boils down to whether the destination device has a certain cryptographic key or not. From a human perspective, however, it is important that destination group identifies a set of users which can be *delimited in terms of the human environment*. Examples of such destination groups are "UCF employees", "firefighters", "members of my club" or "elderly patients diagnosed with Alzheimers's".

Note that the negotiation of  $\mu$ MGs might create consumer and provider groups which can not be identified with such simple natural language statements. This, however, remains acceptable if the user can trust the system that no data transfer can happen in  $\mu$ MGs which is not allowed by a macroMG.

**Specification of a filter.** The filter limits the type and quantity of transferred information and it can also be used to restrict the motivation to certain states of the world. Again, the information allowed by the filter needs to be expressible in natural language in terms of features of the human environment. For instance, the car rental company might be entitled to information whether the vehicle left the state, but it can not receive information about the location of the car otherwise. If the motivation is conditioned by an internal trigger based on the state of the world, the condition needs to be externally visible (such as a fire alarm), it can not be a condition internal to the machine.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper we have outlined an architecture which enables the cooperation of privately owned sensors towards satisfying cooperative goals. We have identified *motivation* as the key missing ingredient in such a cooperation. The source of motivation can only be at human participants - but the participating sensors need to have a way to formalize motivation, reason and negotiate about it. We have proposed a formalism for the model in the form of *motivational gradients*.

The architecture described in this paper opens a wide range of challenges. Reliable negotiation models for microgradients need to be developed. The motivational gradient model needs to be integrated with routing, data dissemination and query processing. These are our future directions of research.

*Acknowledgments:* This work was partially funded by NSF Information and Intelligent Systems division under award 0712869.

## REFERENCES

- [1] L. Bölöni and D. Turgut. Should I send now or send later? A decision-theoretic approach to transmission scheduling in sensor networks with mobile sinks. *Wireless Communications and Mobile Computing Journal (WCMC)*, 8(3):385–403, 2008.
- [2] A. Chakrabarti, A. Sabharwal, and B. Aazhang. Using predictable observer mobility for power efficient design of sensor networks. In *The 2nd International Workshop on Information Processing in Sensor Networks (IPSN)*, pages 129–145, April 2003.
- [3] C. Chen, J. Ma, and K. Yu. Designing energy-efficient wireless sensor networks with mobile sinks. In *WSW'06 at Sensys'06*, October 2006.
- [4] S. Gandham, M. Dawande, R. Prakash, and S. Venkatesan. Energy efficient schemes for wireless sensor networks with multiple mobile base stations. In *IEEE GLOBECOM'03*, volume 1, pages 377–381, December 2003.
- [5] A. Kansal, M. Rahimi, D. Estrin, W. Kaiser, G. Pottie, and M. Srivastava. Controlled mobility for sustainable wireless sensor networks. In *The 1st Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, pages 1–6, October 2004.
- [6] A. Kansal, A. Somasundara, D. Jea, M. Srivastava, and D. Estrin. Intelligent fluid infrastructure for embedded networks. In *The Second International Conference on Mobile systems, applications, and services (MobiSys'04)*, pages 111–124, 2004.
- [7] H. Luo, F. Ye, J. Cheng, S. Lu, and L. Zhang. TTDD: Two-tier data dissemination in large-scale wireless sensor networks. *Wireless Networks*, 11(1):161–175, January 2005.
- [8] J. Luo and J.-P. Hubaux. Joint mobility and routing for lifetime elongation in wireless sensor networks. In *IEEE INFOCOM*, volume 3, pages 1735–1746, March 2005.
- [9] J. Luo, J. Panchard, M. Piorkowski, M. Grossglauser, and J.-P. Hubaux. MobiRoute: Routing towards a mobile sink for improving lifetime in sensor networks. In *The 2nd IEEE/ACM DCROSS*, pages 480–497, June 2006.
- [10] S. Olariu, M. Eltoweissy, and M. Younis. ANSWER: AutoNomouS networked sEnsoR system. *Journal of Parallel and Distributed Computing*, 67(1):111–124, 2007.
- [11] J. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, Cambridge, MA, 1994.
- [12] L. Tong, Q. Zhao, and S. Adireddy. Sensor networks with mobile agents. In *IEEE Military Communication Conference (MILCOM)*, October 2003.
- [13] D. Turgut and L. Bölöni. Three heuristics for transmission scheduling in sensor networks with multiple mobile sinks. In *Accepted for publication in Proceedings of International Workshop on Agent Technology for Sensor Networks (ATSN-08), in conjunction with the Seventh Joint Conference on Autonomous and Multi-Agent Systems (AAMAS 2008)*, May 2008.
- [14] Z. Vincze, D. Vass, R. Vida, A. Vidacs, and A. Telcs. Adaptive sink mobility in event-driven multi-hop wireless sensor networks. In *Proceedings of the First International Conference on Integrated Internet Ad hoc and Sensor Networks*, page 13, 2006.
- [15] Z. Vincze, D. Vass, R. Vida, A. Vidacs, and A. Telcs. Adaptive sink mobility in event-driven densely deployed wireless sensor networks. *International Journal on Ad Hoc and Sensor Wireless Networks*, 3(2-3):255–284, April 2007.
- [16] Z. Vincze and R. Vida. Multi-hop wireless sensor networks with mobile sink. In *ACM International Conference On Emerging Networking Experiments And Technologies*, pages 302–303, 2005.
- [17] Z. Vincze, R. Vida, and A. Vidacs. Deploying multiple sinks in multi-hop wireless sensor networks. In *Proceedings of IEEE International Conference on Pervasive Services (ICPS)*, pages 55–63, July 2007.
- [18] W. Wang, V. Srinivasan, and K.-C. Chua. Using mobile relays to prolong the lifetime of wireless sensor networks. In *International Conference on Mobile Computing and Networking*, pages 270–283, September 2005.
- [19] Z. Wang, S. Basagni, E. Melachrinoudis, and C. Petrioli. Exploiting sink mobility for maximizing sensor networks lifetime. In *The 38th Annual Hawaii International Conference on System Sciences (HICSS'05)*, January 2005.
- [20] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *ACM SenSys*, pages 14–27, November 2003.