

Market Mechanisms for Value of Information Driven Resource Allocation in Sensor Networks

Boleslaw Szymanski, S.Yousaf Shah,
Sahin Cem Geyik, Sanmay Das and Meenal Chhabra
Department of Computer Science and Network Science and Technology Center
Rensselaer Polytechnic Institute (RPI), Troy, NY, USA
Email: {szymansk,shahs9,geyiks,sanmay,chhabm}@cs.rpi.edu

Petros Zerfos
IBM T.J. Watson Research Center
Hawthorne, NY, USA
Email: pzerfos@us.ibm.com

Abstract—This paper examines the possible uses of different market mechanisms for resource allocation at different levels of Wireless Sensor Network (WSN) architecture. The goal is to maximize the Value of Information (VoI) for WSN users. We discuss three different levels of WSN architecture. The lowest level focuses on individual nodes and their basic functions of sensing and routing. We give an example showing how the use of *auctions* at individual nodes can help to efficiently perform these functions. The middle level focuses on services that are abstractions of applications running on sensors. Complex applications are composed automatically from basic ones. We discuss the use of *switch options* to address some of the challenges arising in such dynamic service composition. Finally, we consider the highest level – network deployment and sharing – and conjecture that *options* may be valuable in creating proper incentives for rational deployment and sharing of WSNs.

Keywords-Wireless Sensor Networks (WSNs); Value of Information (VoI); Quality of Information (QoI); Auctions; Switch Options; Real Options.

I. INTRODUCTION

Deciding how to allocate resources to maximize Value of Information (VoI) for network users is a challenging problem in WSNs. There may be many users with limited trust, and the value of information to each user may be subjective and user specific. A reasonable way of scaling up resource allocation in such domains may eventually be to deploy market mechanisms that take decisions based on the explicit worth that users are willing to assign to the information. However, designing mechanisms that allow for both the appropriate amount of information disclosure as well as for efficient outcomes is a challenge in itself. As proof of concept and to stimulate discussion in the community, we present three different market mechanisms that can be deployed at three different levels of WSN architecture.

The Value of Information is most easily thought of as the amount a user who uses information from the network to make decisions would be willing to pay for the information. Therefore, it is reasonable to think about it as the *change in expected utility* of a decision-maker who receives the information. For example, consider a police chief deciding how

many cars to allocate and where to deploy them on a given night. A wireless sensor network could pick out threats, aiding in more effective deployment of vehicles, saving police department resources. The same level of crime prevention could be achieved using significantly fewer resources, and the savings in this case could be thought of as the value of information provided by the WSN. We note some important preliminaries here. First, we distinguish between quality of information (QoI), which can be represented by a vector of objective measurements, and VoI, which depends both on objective QoI as well as on subjective assessment of the importance that an end-user assigns to certain information. Hence, QoI can be measured entirely within the system that produces information while VoI is dependent on both QoI and exogenous information. For our purposes we assume that VoI is externally specified or provided by an oracle, while QoI is measured by the sensor network itself.

Second, we define three levels in WSN architecture at which we address the resource allocation problem using market mechanisms. These three levels also correlate with different time ranges for the solutions, varying from tens to hundreds of milliseconds at the first level, tens of seconds and minutes at the second level, and days and months at the third level. The lowest level is that of a single sensor node involved in sensing and routing. The goal at the node level is to collect the set of the most meaningful measurements or to transmit the most valuable packets first when congestion arises at a node. This is particularly difficult in sensor networks because the congestion is often intermittent, associated with events that move through the network (like an edge of the forest fire, or an object tracked by the network), and therefore traditional congestion control mechanisms that rely on feedback from the destination to the source do not work well. An important aspect of QoI in sensor networks is the time delay with which information is available to end-users. Thus, any congestion or packet collisions will lower QoI of the information carried by affected packets. We assume QoI is not affected by network utilization, making resource allocation approaches based on efficient network utilization not applicable [9], [8]. We

demonstrate how the use of auctions on the very nodes where temporary congestion occurs can help resolve congestion and therefore minimize the loss of VoI.

At the mid-level, networks perform more complex functions in the form of services. Basic services are often hosted by multiple nodes, and executing them commands significant resources (most importantly energy, but also bandwidth, computing power, and sensing modes). More complex applications can be composed from basic services by properly interconnecting the basic services, forming a service oriented architecture. We suggest the use of switch options to decide the efficient interconnection of basic services as network conditions change dynamically.

Finally, at the high level of the WSN itself, market mechanisms may help to optimize network deployment and sharing. Sensor networks are spatial by nature and in cases where many different authorities are responsible for adjacent or overlapping spatial domains, the optimal deployment of sensor networks and access to the data that they collect are important issues. We show that the use of *real options* could lead to the creation of proper incentives for rational deployment and sharing of sensor networks.

The problem of decentralized resource allocation has been studied widely in communication networks. For example, Mainland *et al* propose a decentralized reinforcement-learning based scheme for efficient vehicle tracking in a WSN under energy constraints [7]. Congestion based pricing mechanisms also have featured prominently in the literature [5], [6], generalizing from the traditional domain of tollways and airports to communication networks. The goal for such systems is to balance externalities imposed by increased traffic. Flat rate pricing, prevalent today, is not an optimal strategy for service providers [6]. Fixed pricing schemes can lead to overuse of bandwidth by exploitation of TCP at the user end, so using congestion pricing or differentiated QoS would help to avoid externalities. Various mechanisms to avoid Internet congestion have been suggested (although none have thus far gained much real-world traction) [2], [4]. Most of these mechanisms rely on congestion feedback from the destination to the source, an approach that is not feasible for the intermittent congestion caused by event-driven network flows in WSNs.

II. AN AUCTION MECHANISM FOR DISTRIBUTED CONGESTION CONTROL

As discussed above, resource allocation is driven by VoI that combines an objective function that measures QoI and a subjective component that assesses how valuable information with the given QoI is to the end-user.

In this section, we consider the low level management of resources in WSNs, as exemplified by bandwidth allocation in congestion scenarios. Following Chen *et al* [1], we consider public safety and emergency response needs of a VIP visiting a large city. The streets are equipped with

a sophisticated WSN composed of acoustic sensors, closed circuit cameras, chemical fume sensors and so on. Various state agents are present in the area to ensure the safety of the visitor and the public, using the deployed WSNs to monitor events. Every agency has its own mission and priority monitoring targets. Local police may be interested in monitoring traffic violations in the area as well as mob behavior, whereas a federal agency is tasked with detecting any kind of coordinated terrorist attack and dealing with high-level threats. State agencies may also be monitoring large vehicles entering the area, as well as individuals or cars with suspicious mobility patterns. Another state agency is monitoring the visitor's car, also a high priority mission. Suppose a vehicle with a suspicious driving pattern is being tracked by the state agency, while a gathered mob is continuously monitored by local police. Both targets are getting close to the visitor's car as they enter a nearby intersection. Now, data packets are continuously sent to three different sinks (agencies), all with high priority, monitoring three different targets (visitor's car, suspicious mob and threatening vehicle). A noticeable problem occurs due to the fact that all three targets are physically close, causing congestion at nearby nodes that transmit packets to sinks. This congestion increases network delays and may even cause packet loss, decreasing QoI of the traces of the targets that are of high value at that point in time. This problem will become more complex when there are more targets with coordinated mobility patterns involved.

Chen *et al* [1] formally define an auction mechanism to solve this problem, which we summarize here. They consider two possible goals: efficiency (minimize total loss of information value) and equity (equalize the loss of information value for all missions). While the value of information could be a very general function of QoI and user-specific importance of the information, for simplicity, we consider a case where the utility for a mission i is a linear function of the QoI received about mission i , $U(i) = v(i)q(i, d)$, where $v(i)$ is the mission specific multiplier and $q(i, d)$ is the QoI for the specific mission i with the data transmission delay d for this mission (although many more parameters define the value of $q(i, d)$, we explicitly refer to d , as this is the parameter directly impacted by congestion; other parameters, such as precision of target's position or precision of the time of measurement are not affected by congestion). Theoretically, even in this simple example, it would be useful to have $v(i)$ dependent on the information, for example, the closer the suspicious vehicle is to the visitor's car, the higher the VoI of its position. We make one further simplification in what follows. We assume that each node maintains only one packet for each mission because when two distinct packets of the same mission are received at the destination, the one with the more recent target data brings VoI to the same value, regardless whether the other one was received or not. Hence, each node's maximum queue length is limited by the number

of active missions, a requirement easily satisfied by modern sensor nodes. Once congestion arises, the node needs to decide in what order the waiting packets will be transmitted, exposing them to different delays at the node. Since the standard congestion control techniques involve destination-to-source feedback, they are not applicable here because of the intermittent nature of congestion that keeps moving from one node to another in target tracking applications of WSNs.

Thus, we need to quantify the loss of information resulting from packet delays caused by congestion. To do so, we need to assess VoI of a piece of information. Often, VoI is most directly related to how much it changes the uncertainty or beliefs of the user of the information. For example, consider a police team monitoring potential threats to a VIP. A WSN reports the position and direction of movement of a suspicious vehicle, enabling the police to have a current estimate of the trajectory of this vehicle. How much does additional, new information about the position and direction of motion help?

There are two related ways of thinking about this: first, how much the new information reduces uncertainty about the position of the object being tracked. Let us compare the case where the last observation by a node in the network was 10 minutes ago with a case where the last observation was 10 seconds ago. Clearly, new information will be more useful in the former case than the latter. Intuitively, this is because new information in the former case leads to a greater reduction in uncertainty about the object position than it does in the latter case. A useful mathematical formulation in this case may be the reduction in entropy of the user’s belief about the location and direction of the tracked object. If the user’s belief is represented by a well-behaved distribution, like a Gaussian, this could be further simplified by considering perhaps only the variance of the distribution.

A related way of thinking about this issue is to quantify the “surprise” element of an observation as being key to the value of the information contained in the observation. For example, if a suspicious vehicle was traveling slowly in a particular direction, but suddenly made a U-turn and accelerated sharply, that information may be more relevant than if it just continued on its previous path. How can this be objectively quantified? Perhaps the difference in implied probability distributions of the location of the object could be important. A measure like KL divergence of the two distributions (the post-observation distribution and pre-observation distribution) may be helpful in quantifying VoI in cases like this.

Chen *et al* [1] use the reduction of uncertainty approach and measure the loss of value as proportional to the additional delay incurred by the packet in congestion. To quantify, let t_m denote the time at which target was sensed at $\text{loc}(t_m)$. After time Δt , its new location $\text{loc}(t_m + \Delta t)$ is uncertain. The maximum feasible acceleration of a vehicle is limited, therefore according to equations of motion




Figure 1. Auction Based Bandwidth Allocation Framework

$|\text{loc}(t_m + \Delta t) - \text{loc}(t_m)| \propto (\Delta t)^2$. Hence, the VoI decreases quadratically with the delay. Denoting by C_i all the constants of the proportionality of the VoI to the square of the packet delay, the loss of VoI for mission i , defined as the difference between the VoI of a packet not delayed by congestion and the packet arriving with such delay, is $L_i = C_i(\Delta t)^2$. A node experiencing congestion can choose one of the following two ways to decide the order in which packets are to be sent from the congested node:

1. Equalizing utility loss among all applications, i.e., \forall missions i and j $\Delta L_i = \Delta L_j$
2. Minimizing the total utility loss, i.e., $\min \sum_i L_i$

Chen *et al* [1] considered only linear term of the VOI loss that is dominant when the prediction of the future target position is not made. The implementation of this approach done by Chen *et al* [1] is depicted in Figure 1. We assume that each packet carries with it the priority of the mission for which it is reporting data. As proposed in [1], a Second Price Auction is held at the point of congestion. The bids are entered by the target update packets that compete for transmission slots at the congested node, using the predicted utility loss as a form of currency. The auction winning packet will receive the currently available transmission slot. Packets losing the auction will obtain additional funds for the future auctions of transmission slots on the same node, as the auctions repeat for every transmission slot of that node until congestion is resolved. The additional funds received by each losing packet are proportional to its VoI loss incurred by the delay of its transmission. It should be noted that the congestion is resolved locally at the point of congestion. Conducting such distributed auctions at each node whenever congestion arises there minimizes the overhead of overall congestion control. In particular, no destination-to-source feedback is needed. Such feedback, which is used in traditional congestion control protocols, would delay the deployment of congestion control process, thus, they would increase the VoI loss caused by congestion.

Chen *et al* [1] show that the auction allocation mechanism for congestion control performs better than equal allocation or mission priority proportional allocation.

III. SWITCH OPTIONS FOR DYNAMIC SERVICE COMPOSITION

Dynamic service composition for sensor networks was originally introduced by Geyik et al. [3] who characterized a service abstractly as a program running on a sensor node that requires a certain set of inputs and produces some data (output) characterized by a set of metadata. An automated composition of higher functionality service creates a dataflow graph by interconnecting a set of services together. Choosing recursively the lowest cost of input providers (that is other services that are able to provide some of the required inputs) is termed *service selection*. This problem can be seen as the optimal selection of services rather than the generation of the optimal dataflow graph [3] because optimal interconnections can easily be chosen once the participating nodes are selected.


Switch options are usually defined as having multiple plans readily available for execution and draw a direct analogy to the service selection task that was described above. When making investments decisions, there could be multiple options to proceed (e.g., *expand*, *scale down* or *shut down* investment) or to use different operating modes for the investment (e.g., having an option to use different materials production, according to market conditions). The additional funds needed to make these options available should be taken into account when valuing a project. The *value* of a switch option is calculated as the sum of the benefits of switching from one option to another, whenever it is profitable.

We propose to use the switch option approach when choosing the input providing services. The currently selected services are considered an investment. The risks involved in choosing them are reduced by *periodically re-examining* their selection in view of the current network conditions and switching to a new service composition when profitable. Initially different compositions of services that are considered but not chosen as suboptimal are kept as *switching options* and may be chosen in the future.


A. Methodology

We propose two phases for service selection: (i) Test Phase and (ii) Service Selection Phase. The test phase runs multiple possible services and evaluates at what conditions switching between the input providers improves the VoI attained. Furthermore, the probabilities of these conditions to arise are also calculated, so we can determine how profitable it is to keep different selections available. Since extra effort is expended during the test phase and the cost of this phase needs to be made up by future gains, the length of this phase is limited.


The service selection phase starts with the selection that has highest expected gain. A switch to another selection is guided by the data collected during the test phase. Let the random variables V_A, V_B denote the (instantaneous) VoI of services A and B with the same functionality;



(a) A Parking Garage Example to Explain Switch Options Methodology



(b) Service Value Fluctuations for Two Types of Monitoring Service



(c) Value of Switching Option During the Test Period

Figure 2. A Real World-Scenario for Switch Options

these are implicitly defined by the state of network and its environment. Observing frequency of such states over the test phase, we can estimate the probability density function $p(V_A, V_B)$ of having V_A and V_B as the instantaneous VoI of services A and B , respectively. Then, the expected value of the ability to switch from service A to service B is:

$$V(A \rightarrow B) = \int_{V_B=0}^{\infty} \int_{V_A=0}^{V_B} (V_B - V_A) p(V_A, V_B) dV_A dV_B,$$

which is in effect a summation over all $V_B > V_A$. Note that this is an idealized equation that does not take into account costs of measurement of the state, costs of switching between options, quality of state measurements (sensors can easily measure the noise level in a garage example but not whether there is an accident or not) etc. – incorporating such considerations is an important element of future work.

B. Real-World Scenario

Figure 2a shows a sensor network implemented in a covered parking garage. There are two types of services

in this network: (i) a microphone service of readings from an acoustic sensor to monitor the sound volume, and (ii) a camera service that provides views of the area covered by the microphone monitors. We consider a long term monitoring mission during which automated service selection may choose to utilize one or both of the services in monitoring.

An illustration of expected test period results for this application is shown in Figure 2b. The test measure how external factors affect the VoI and benefits produced by each service. Clearly, the camera view for an area will produce the best results, but running a camera is a costly operation (due to its energy consumption and maintenance costs) in mid-term applications, so the benefit of using this service is low when no events are happening in the garage. Figure 2b also shows that when the VoI of the microphone service drops below a certain level, there is loud noise in the garage, and the VoI provided by the camera service increases. Of course, loud noise often signifies an important event in which case the VoI (the first factor in the above equation) increases even more, giving the camera service a higher benefit than normal. On the other hand, the microphone gives faulty measurements when the noise level is high. This example shows how the switch option balances the *Value of Information (VoI)* provided by the services with the cost incurred by it.

Once the costs of running multiple services in each area during the test period have been incurred, the service selection mechanism can switch between the microphone and the camera service, and can do it with increased efficiency based on information about the conditions that are beneficial for switching gathered during test period. In the long run, the costs of the service test phase is compensated by clever switching actions, which it enables with the information it acquires. Figure 2c shows the advantage that could have been gained had switching information been available during the test period shown in Figure 2b. The area difference between the curves of Figure 2c and the curve of the microphone service (since it is best on average) in Figure 2b quantifies the extra value of the switching option.

IV. DECENTRALIZED COLLABORATIVE MARKETS FOR WIRELESS SENSOR NETWORK RESOURCES

Wireless sensor networks are typically deployed for the purpose of providing better information to a decision-maker. For example, consider a WSN deployed in a forest: the primary purpose may be to monitor temperature and environmental conditions to detect forest fires and prevent them from spreading. Extreme temperature readings are a critical but rare event. It is absolutely necessary that the network performs well at detecting fires when they occur. At the same time, given the rarity of such a circumstance, the sensors in the network could also serve other roles: monitoring wildlife activity, rainfall levels, etc. Deploying the network is costly, so the benefits must outweigh the cost of deployment. These

benefits are again defined by the Value of Information (VoI) provided to the users of the network.

As WSNs start being deployed systematically, we must reason explicitly about the inherent costs, benefits, and trade-offs, as well as the sharing of the deployed network among multiple entities and for different purposes. The natural language for this reasoning is the language of economics, and we need to think about the market for information that is provided by a WSN, and how it can interact with the design of such a network. The key players in the marketplace are the service provider, who deploys and maintains the WSN, and users who value the information provided by the WSN and are willing to pay for it. Often one of the users may actually be the service provider.

Continuing with our forest-monitoring example, suppose the state of New York deployed a WSN in its forests, and the state of Massachusetts deployed one in its own. Many fires are capable of spreading across the Massachusetts-New York border, so it is important for each state to be aware of what is going on in the forests of its neighbor. Instead of deploying its own networks in the neighboring states, Massachusetts can acquire access to data from New York when it is critical, and vice versa. Yet, a blanket agreement to share all data is not necessarily a good idea, because producing and sending data is costly, and non-critical data that has value to New York may not be as valuable to Massachusetts. What is needed is a dynamic market mechanism capable of allowing trades between users and service providers, and allocation of resources that can be relied upon in emergency situations. While non-trivial, no other decentralized approaches can efficiently arrive at the right tradeoff between the value of information provided by WSNs and the costs of their deployment and use.

In this section, we sketch a vision for some possible ways in which market mechanisms may help to solve the problem of resource allocation in WSNs, leading to appropriate payments in exchange for the VoI provided to users. For our purposes, we assume, as before, that the decision-makers who use the information provided by the WSN are in possession of well-calibrated estimates of the value of the information that the WSN provides to them. Therefore, we assume that the utility to the user is exogenously specified.

Real Options for Emergency Information. For simplicity, assume that Massachusetts and New York only have interest in each other's WSN information if it is potentially predictive of the existence of a forest fire in the neighboring state, and that the rest of the time they have no value for information from the other. However, since a forest fire is a critical emergency situation, the value of information related to forest fires is high. One possibility would be for the states to write and buy real options on the transmission of their data (we note that real options are a convenient formalism – similar outcomes can also arise out of alternative decision analysis techniques). For example, New York could write,

and Massachusetts purchase, an option that enables Massachusetts to use up to 40% of the WSN resources available to New York. The option would be American style, so that Massachusetts could exercise it at any time that it becomes necessary, triggered by the existence of a forest fire that may spread to Massachusetts.

Several interesting issues arise once we think of the option in these terms. (i) New York state would have to provision extra resources so that it can always reliably provide the services it is writing the option on. It must never be the case that Massachusetts cashes in on the option and New York is unable to provide the service. Therefore, the WSN may need to be over-provisioned for emergency situations. However, the cost of this over-provisioning can be made up for through the payments made to New York when Massachusetts buys the options. (ii) Massachusetts must become aware of information critical to its decision-making about whether or not to exercise the option. This leads to issues of trust between New York and Massachusetts –for example, if Massachusetts were to completely trust New York’s words, and New York deceived Massachusetts into thinking a situation was an emergency when it was not, the option would be needlessly exercised. (iii) Standard option pricing methodologies may have to be extended to deal with the idiosyncrasies of real, deployed WSNs; for example the information needed to decide whether to exercise the option, as discussed above. Further, especially in emergency response, which is necessary when *rare* events occur, standard methodologies based on assumptions of normality that use the variance of expected returns (or return-equivalents) may no longer be the right approach.

While the research questions raised are significant, real options methodology has the potential to deliver value in several ways. It could lead to the establishment of a decentralized collaborative market that guarantees the availability of information when needed. Options provide a market-based solution for resource allocation and management between competing or cooperative WSN service providers, potentially allowing them to recover infrastructure and resource costs encountered in deployment. The cost that each party would pay in the form of options in practice is less than the managerial and deployment cost of a fully fledged WSN. Moreover this price will also be lower than the price a party would pay if it only decided to buy services at the time of an emergency event, when sellers could charge higher prices in response to demand because the information has the highest value to safety. While we have discussed this so far in the context of two parties, the extension to multiple service providers and users is obvious.

V. CONCLUSION

This paper discusses resource allocation in WSNs. We propose the application of market mechanisms to address problems of resource allocation from the very low level of

a sensor node to the highest level of WSN deployment. We have simulated an application of auctions at the node level and found it to be very efficient in resolving congestion. At the mid-level of WSN architecture we are working on application of *switch options* to dynamically compose complex services by optimally interconnecting basic services. As future work, we will further investigate, sketched here, application of *real options* to create a decentralized collaboration among WSN owners to deploy WSNs in a cost effective manner.

VI. ACKNOWLEDGMENTS

This research was sponsored by the Army Research Laboratory and accomplished under Cooperative Agreement Number W911NF-09-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] L. Chen, Z. Wang, B.K. Szymanski, J. Branch, D. Verma, and R. Damarla, Dynamic service execution in sensor networks. *The Computer Journal*, 2010, **53**(5):513-527.
- [2] C. T. Ee and R. Bajcsy, Congestion control and fairness for many-to-one routing in sensor networks. *Proc. ACM Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, 2004, pp. 148-161.
- [3] S. C. Geyik, B. K. Szymanski, P. Zerfos, and D. Verma, Dynamic composition of services in sensor networks. *Proc. IEEE Int. Conf. on Service Composition (SCC)*, 2010, pp. 242-249. *Embedded Networked Sensor Systems (Sensys)*, 2004.
- [4] K. Karenos, V. Kalogeraki, and S.V. Krishnamurthy, A rate control framework for supporting multiple classes of traffic in sensor networks. *Proc. IEEE Int. Real-Time Systems Symposium (RTSS)*, 2005.
- [5] J. MacKie-Mason and H. Varian, Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 2002, **13**(7):1141-1149.
- [6] J. Shu and P. Varaiya, Pricing network services. *Proc. IEEE Int. Conf. on Computer Communications (INFOCOM)*, vol. 2, 2003, pp. 1221-1230.
- [7] G. Mainland, D. C. Parkes, and M. Welsh, Decentralized, adaptive resource allocation for sensor networks. *Proc. Symp. on Network System Design & Implementation (NSDI)*, 2005.
- [8] D. Gesbert, S. G. Kiani, A. Gjendemsjo, and G. E. Oien, Adaptation, coordination, and distributed resource allocation in interference-limited wireless networks. *Proceeding of the IEEE*, 2007, **95**(12):2393-2409.
- [9] Y. Xue, B. Li, and K. Nahrstedt, Optimal resource allocation in wireless ad-hoc networks: A price-based approach. *IEEE Trans. on Mobile Computing*, 2006, **5**(4):347-364.