Proc. Seventh Annual IEEE International Conference on Pervasive Computing and Communication, PERCOM09
Galveston, TX, March 2009, pp. 194-203

Auction-Based Congestion Management for Target Tracking in Wireless Sensor Networks

Lei Chen and Boleslaw K. Szymanski
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180
{chenl6, szymansk}@cs.rpi.edu

Joel W. Branch
IBM T.J. Watson Research Center
Hawthorne, NY 10532
branchj@us.ibm.com

***Abstract*—This paper addresses the problem of providing congestion-management for a shared wireless sensor network-based target tracking system. In many large-scale wireless sensor network target tracking scenarios (e.g., a surveillance system for tracking vehicles in urban environments), multiple targets may converge within close proximity of each other. Such scenarios may cause network congestion as nearby sensors attempt to concurrently send updates to a data aggregation point (e.g., base station). We consider the case in which this problem is further complicated by two factors. First, such a large-scale sensor network may very well be deployed to serve multiple target tracking applications with different and dynamic priorities and interests in different (types of) targets. Second, each application will most likely place a different premium on the timeliness of the target information (principally defined by some quality metric) they receive. All the above challenges introduce formidable challenges in providing the expeditious delivery of target information to all prioritized applications.**

In this paper, we advocate the use of a distributed auction-based approach to locally manage network bandwidth allocation in the described context. We use the Second Price Auction mechanism (to ensure incentive compatibility) in which the congested node acts as the auctioneer and the packets carrying target updates act as bidders. Their bid values are defined by the *loss of information utility* to the applications associated with the packets. The winning packet receives the current transmission slot of the auctioneer node. We demonstrate through simulation that the resulting auction allocates bandwidth efficiently, maximizing the collective applications' goals, even when the application priorities change dynamically.

Keywords-wireless sensor networks, congestion management, taget tracking, quality of information

I. INTRODUCTION

Years of steady advancements in integrated circuitry and wireless communications, fabrication cost reductions, and innovative high-value applications have solidified wireless sensor networks (WSNs) as a significant component of pervasive computing. WSNs, comprised of small, autonomous, and tetherless sensor nodes, have played an essential role in closing the void between computers and the physical world by enabling continuous non-intrusive observation of real world phenomena, often achieving goals that have been unattainable in the past.

Given the nature of WSN technology, growing concern for public safety and the associated need for an efficient emergency response, it is expected that there is growing interest in instrumenting public spaces with WSNs [1,2]. Typical motivating examples include monitoring shopping centers, college campuses, highways, and well-populated metropolitan areas. In many instances, the problem will be, or has been, reduced to identifying and tracking suspicious events or entities, especially where it is difficult to maintain the physical presence of the appropriate authorities. Also in many instances, the capabilities of such WSN target tracking systems will be of benefit to more than one party. Hence, a resource sharing mechanism will have to be deployed. This introduces important new research challenges in WSN operation, as is illustrated by the following scenario.

Suppose a city hosts a high-profile event requiring the majority of law enforcement agents to be physically present within the event's vicinity. Fortunately, a multimodal WSN composed of chemical sensors, video cameras, etc. has been deployed in order to help monitor other areas of the city. The general goal is to identify and track suspicious and/or threatening behavior (while enforcing privacy protection whenever possible). However, different agencies share the WSN to monitor different types of targets with various priorities. For instance, due to limited personnel, local police may be mainly interested in using the network to detect "commonplace" offences such as muggings, traffic violations, etc. However, federal agencies would be tasked with monitoring for highly-organized and

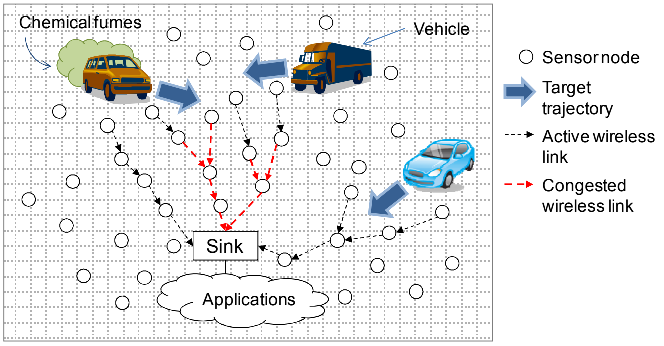

Figure 1. Wireless sensor network target tracking scenario causing network congestion.

Research was sponsored by US Army Research Laboratory and the UK Ministry of Defence.

high-magnitude threats (e.g., terrorist activity), which may, among other characteristics, be indicated by the presence of large vehicles emitting suspicious chemical fumes and driving in suspicious patterns. Furthermore, due to specific skill backgrounds, different task forces within the agency may be assigned to track different types of threats. Figure 1 depicts a rendition of such a scenario, where three vehicles are being tracked by separate users, represented more generally as *applications*, and updates describing the vehicles' behavior are sent to a single aggregation point, or sink. A noticeable problem occurs as the truck and van converge within close proximity to each other and temporarily cause network congestion on nearby paths to the sink. This increased delay and potential loss of packets can easily reduce the quality of information of the targets received at the sink, whose value increases in importance in such an event since target convergence might indicate collusive behavior which must be monitored closely. The problem is further complicated following the expectation that the users described above have different and dynamic usage priorities based on both the intent of the user group and target behavior. A similar WSN application scenario may be found in other domains, such as the military. In this paper, we address the problems brought forth by the previous scenario.

There have been several congestion control protocols proposed for WSNs, such as congestion avoidance [3], rate-limiting congestion control [4], congestion control and avoidance, CODA [5], and rate control with multiple classes of flow, CoBRA [6]. In general, they are comprised of three stages: congestion detection, congestion notification, and traffic rate adjustment that together create a feedback loop. The papers cited above either provide different techniques for each of the stages, or include considerations of different packet classes and different groups of sensors [6]. However, the feedback loop-based solutions fail when the flow paths change before the feedback loop can form. This is the case in such sensor network applications as mobile target tracking or dynamic phenomena monitoring, in which the set of nodes actively engaged in sensing change quickly because the targets or the monitored phenomena move over time. As a result, traditional feedback loop-based congestion control approaches are not adequate to resolve the problems described in the previous scenario. Furthermore, the published literature mainly focuses on the static fairness of resource allocation during congestion without considering dynamic changes of priorities of applications, in particular their changing sensitivity to packet delays. We address this issue in our paper.

The major contribution of this paper is an auction-based mechanism for providing efficient and localized WSN transient congestion management based on applications' priorities and the utility of sensed information. An essential component of this mechanism is a partially user-defined information utility metric that jointly represents two components. One is the objective *quality of information*, or *QoI*, at which targets are tracked. The other is the subjective *priority* of each application defined by the user to reflect the value of knowing the precise position of a target. The combined metric is defined as the product of the tracking imprecision at the source sensor, the delay with which target updates are received, and the application's priority. We term the result the *loss of information utility* experienced by the application. The larger is this value, the less valuable is the information received by the application. This relation is clear for the tracking imprecision at the source sensor or delay of the packet carrying the measurements. Inclusion of priority in the metric enables us to differentiate between utility losses of different missions.

Auctions are held at points of congestion. We use the *Second Price Auction* [7] because it is incentive compatible. Thus, the dominant (optimal) bidding strategy for each bidder is very simple: to bid its true value of the asset (which, in our case, is the allocation of the current transmission slot of the node conducting the auction). More precisely, each bid represents the additional loss of information utility that the corresponding application will endure if the packet will not be selected for the current transmission slot. As shown later, the data necessary for computing the loss are carried in each packet and they add only a few bytes to the packet size required by the application. Also the computation itself is simple, so the computational requirements of conducting the auction are minimal. As a result, the transient congestion problem is addressed *locally* at the point at which the congestion arises and its solution adds small overhead both in terms of packet size and computation load on the nodes involved in the auction. This befits our purposes since we focus on allocating bandwidth in response to relatively transient congestion that should not trigger an overly complex solution.

We demonstrate the applicability of our solution by using it to implement two congestion management frameworks, each with different overall goal. The first one aims at equalizing utility loss for all applications to provide fair bandwidth sharing among all of them. The second framework minimizes the sum of utility losses of all applications to optimize global tracking quality. We further demonstrate the solution's strength by using simulation to compare its performance in implementing the two frameworks above with that of using a purely analytical solution and a simple bandwidth equalization solution.

This paper is organized as follows. Section II defines the loss of sensor information utility. Section III describes the auction-based congestion management approach while Section IV describes an analytical approach. Section V presents a simulation-based evaluation comparing the performance of multiple congestion management solutions. Section VI discusses related works and Section VII concludes the paper.

II. Loss of Information Utility in Target Tracking By Multiple ApplicationS

Before proceeding, we state some fundamental underlying assumptions that provide a foundation for this work. We assume the use of a WSN with stationary and homogeneous sensors, i.e., sensors do not move and they share the same modality and general capacity for tracking. We also assume that sensors have been made aware of and can recognize (with reasonable probability) the targets that separate applications (in this case, residing at the sink) are interested in tracking. We also assume that each sensing node knows the priorities of the applications. This later requirement is easy to fulfill using any straightforward dissemination protocol to distribute applications priorities from the sink to all sensing nodes. Further assumptions are stated throughout the paper.

A. *Target Tracking Algorithm*

In this work, we assume the use of a distributed target tracking algorithm described in [8]. It assumes binary sensing operation of sensors in which each sensor detects only the event of a target crossing its sensing coverage boundary. At that time, the node can compute and report to its neighbors a target's location, velocity, and trajectory based on previously received reports. In addition to being broadcast to the node's neighbors, the report is also sent to the base station. Using only perimeter-crossing events for detection and reporting reduces computation and communication overhead. Interestingly, even if the binary sensing is imperfect, when the distance at which the target is detected coming in or out of the sensing perimeter of the node varies, the algorithm provides good estimates of the target position [9]. More details about this sensing algorithm are provided in [8, 9]. To support congestion auctions, we assume that each report contains the target type and the identity of the application tracking the target. This additional information occupies anywhere from 4 to 8 bytes of the packet.

B. *Loss of Information Utility*

We assume that for a target tracking application, the QoI with which a target is tracked is a function of the uncertainty of a target's location, which can be objectively measured or estimated. It depends on the initial imprecision with which the target's location is estimated by the tracking algorithm at the source. It is then increased by the packet's transmission delay between the source and base station. The utility of the information delivered to the application is also dependent on a subjective priority of the application provided by the user. The product of these two factors (QoI and priority) provides a suitable metric expressing the loss of information utility to the tracking application. To reflect such a measure, we conceptually define the information utility loss, u, of an application as follows:

$$u(t) = r(t)p, where\ p \in \mathbb{R}, p \geq 0, \qquad (1)$$

where p is an application's priority (value of p increases with priority) and $r(t)$ is the radius of a circle (or sphere in 3-D) around the *predicted* location of a target in which the target actually resides at time t. In other words, $|loc_p(t), loc_r(t)| \leq r(t)$, where $|.|$ is a physical distance metric and $loc_p(t)$ and $loc_r(t)$ are the target's predicted and real locations, respectively, at time t. This is illustrated in Figure 2, which shows two targets', *a* and *b*, predicted and real locations at time t_1. Both targets have their initial position uncertainty at the source equal to *d,* as shown in Figure 2. However, at time *t,* the target *b* can potentially be further from the predicted location than *a* target can due to any number of factors; we explain such factors in more detail shortly. Hence, the target *b*'s actual location can reside in a wider range, increasing the uncertainty of tracking and hence the *loss* of information utility to the application, as represented by the larger shaded area for target *b* in the figure. Similarly, it is easy to see from (1) that increasing p increases utility loss (or, reduces utility). This captures the intuition that the application with higher priority should lose utility faster than the application with lower utility.

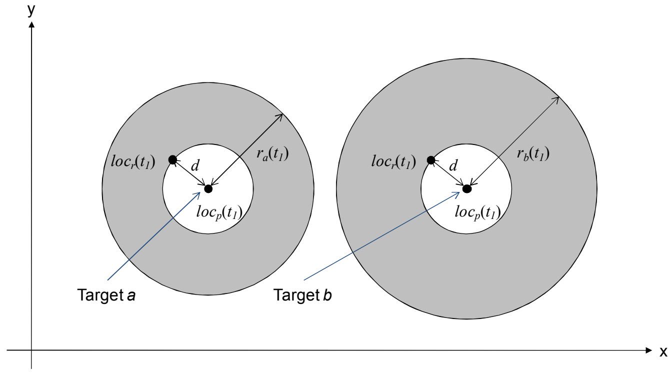

Figure 2. Predicted and real target locations over time.

In other words, with the same uncertainty *r(t),* the utility loss of an application with higher priority should be larger than the loss of application with lower priority. This intuition motivates the conceptual definition of utility loss.

We now expand upon (1) to provide a more practical definition of information utility loss that is also more straightforward to measure. We assume that target report packets contain $loc_r(t_m)$ and $\boldsymbol{v}(t_m)$, where $\boldsymbol{v}$ is the target's velocity and t_m is the time of measurement (assuming that an ideal tracking algorithm is used [8]). Considering that a delay, Δt, exists between the time at which a target report is generated at the source sensor and some latter time t (i.e., the time at which a report reaches the tracking application or an intermediate node), we *formally* need to calculate $r(t_m + \Delta t)$. However, the lack of observation of the target in the period Δt must be considered in this calculation.

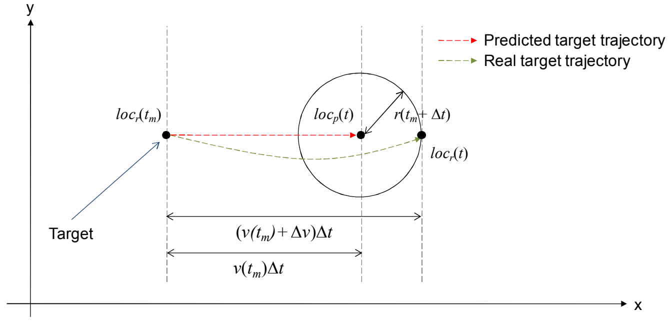

Figure 3. Example relating target velocity to $\boldsymbol{r(t_m + \Delta t)}$.

We use the model of *constant speed precision prediction* to calculate how $r(t_m)$ changes over time. Hence, we assert that the speed computed at time t differs from the target's actual speed no more than $\Delta v = \alpha v(t_m)$[1], where α represents the speed precision prediction factor. More specifically, α represents any (combination of) feature(s) of the target that may affect its ability to change speed. For instance, this may correspond to a reasonable assumption that there is a maximum

[1] Note that *v*(*t*) represents speed, which is a magnitude of velocity.

acceleration that the target can sustain. Figure 3 illustrates how $r(t_m + \Delta t)$ is derived using the assumption of constant speed precision prediction. Here, a target's predicted position, $loc_p(t)$, and corresponding distance traveled given the location and speed measured at t_m is shown. The target's real position, $loc_r(t)$, and corresponding distance (not directly observed by the application) given Δv is shown as well. The figure demonstrates that $r(t_m + \Delta t)$ can be no larger than the difference in distances and hence, we have the following equation:

$$r(t_m + \Delta t) \leq |v(t_m)\Delta t - (v(t_m) + \Delta v)\Delta t| = av(t_m)\Delta t. \tag{2}$$

According to (2), $r(t_m + \Delta t)$ is linearly proportional to both $v(t_m)$ and Δt. We note that in the most general case, the stated precision prediction factor, α, should include a factor of both the imprecision of the employed tracking algorithm and some measure of the unpredictability of the target's behavior.

In summary, we consider a model of quality of target tracking that assumes constant speed precision prediction. It yields the following utility loss metric for sensor information for a target tracking application:

$$u(t) = \alpha v(t_m)(\Delta t)p, where\ p \in \mathbb{R}, p \geq 0. \tag{3}$$

Now, utility loss can be measured using a target's velocity and the time delay between when the velocity was measured at the sensor and received at the application. For simplicity, we set $\alpha = 1$, which can easily be achieved by properly scaling the missions' priorities.

III. Auction-Based Transient Congestion Management

A. Overview

We advocate the use of an auction-based mechanism for prioritizing the forwarding of target report packets when the congestion arises. When there is no congestion, applications accumulate different utilities. This is because, applications have different priorities, track targets moving with often different speeds, and their target report packets experience different (and largely independent) end-to-end delays (see (3)). Specifically, we employ auctions to provide two different utility loss management solutions: (1) equalizing utility loss across all tracking applications and (2) minimizing the total utility loss of all applications.

Congestion detection for computer networks is a well studied subject. In traditional networks, reactive TCP congestion detection methods are often used, in which congestion is observed or inferred at the end nodes based on a timeout or redundant acknowledgments. In wireless sensor networks, proactive methods are preferred for their efficiency. For simplicity, in our setting, we assume that congestion is detected based on a node's outgoing transmission buffer occupancy. Specifically, congestion is detected locally when more than one packet is ready for the transmission.

We assume that a Spatial TDMA scheme [10] has been deployed in the sensor network to prevent transmission interference between nodes. Here, each node uses one of n slots for transmission over the period t_c. Thus, the average packet's transmission time is t_c/n. The value of n depends on the type of TDMA scheme (i.e., global or local) and on the topology of the network. The upper bound of n for a given maximum number of neighbors of a node has been previously established [10]; hence, we do not pursue defining a value of n in this work.

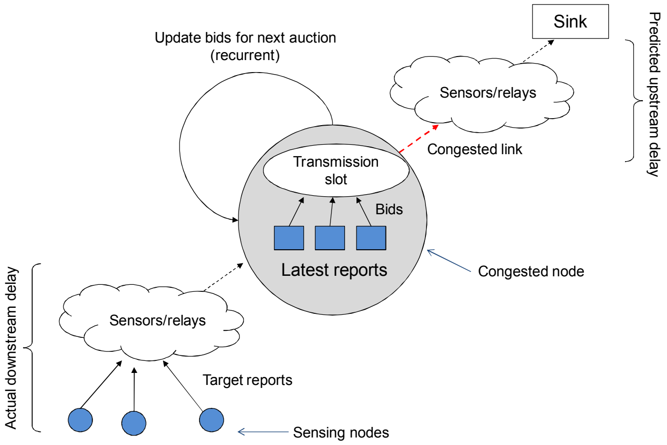

Figure 4. Auction-based bandwidth management framework.

Under our scheme, we assume that target report packets for an application i carry several items of data that are essential to the bidding process: (1) the time at which the target report was generated, $t_{m,i}$; (2) the target's speed, $v_{m,i}$; and (3) the priority of the application tracking the target, p_i. We assume that when a target is first detected, the base station associates it with a relevant application and the application's priority is forwarded back to the node that detected the target. As previously mentioned, the application priority is also propagated among nearby tracking sensors each time a report is generated.

Any time a given node has multiple target report packets for different applications to be forwarded for its current transmission slot, the node conducts an auction to assign the slot to the packet with the highest bid. Hence, the bidders represent packets awaiting transmission and their bids are defined by the predicted information utility loss of the applications to which packets are sent. The auction is recurrent since it often is conducted repeatedly with bidders represented by different reports of each application. Unlike a single auction in which a winner acquires the entirety of a resource indefinitely (the transmission time slot in our case), in a recurrent auction with a participation incentive mechanism it is possible to share resources over time [11]. This in turn prevents resource starvation of any auction participant.

We use the Second Price Auction mechanism, which is incentive compatible, to simplify bidding strategies. Thus, the dominant (optimal) bidding strategy for each bidder is very simple: to bid its true value of the asset, which, in our case, is the allocation of the current transmission slot of the node conducting the auction. Each bid represents the additional loss of information utility that the corresponding application will endure if the packet will not be selected for the current transmission slot. The data necessary for computing the loss are carried in each packet, adding only a few extra bytes to the packet size. The computation of the bid is simple; hence the algorithmic overload of the solution is very small.

We also note that if a target report packet losses an auction for the current transmission slot and a new packet for the same application arrives at the congested node, the new packet (with the most recent value of $t_{m,i}$) replaces the old one in the queue of messages awaiting transmission. This is because the delivery of the less recent packet does not reduce the application's loss of information utility, even if both packets are delivered together. As a result, at most one packet, the most recently received one, associated with the given application participates in each auction, lowering the overhead of the auction execution.

Figure 4 presents a more illustrative description of the auction-based congestion management framework. As shown, the auction happens recurrently at the congested nodes that receive combined upstream traffic carrying multiple application target report packets. The auction is repeated for each transmission slot assigned to the node. Each repeated auction selects the single target report packet that will be transmitted in the current slot. The choice of this packet is made in such a way as to fulfill overall system goals. In the next subsection, we describe how the auction is conducted and show how easily the presented approach can be adjusted for a specific overall goal for all applications.

B. Auction Mechanism

The design goal of the auction mechanism is to either equalize or minimize the actual information utility loss, $u(t)$, for all competing applications. According to (3), the delay between the time at which a packet is generated and the time when it is received at the sink is needed to calculate the utility loss as observed by the sink. However, auctions are conducted at the nodes experiencing congestion, not at the sink. Hence, the *predicted* average information utility loss must be computed at these nodes over the time of congestion to approximate the value of $u(t)$ observed at the sink. To enable the congested node to compute the predicted $u(t)$ (hereafter referred to simply as $u(t)$), we assume that each node knows its distance in relay hops to the sink, as well as the average delay (without congestion) at each hop. Under the assumed routing protocol with global TDMA slot allocation, an average one hop delay is simply $\left(\frac{n}{2}+1\right)\left(\frac{t_c}{n}\right)$. Indeed, in addition to the transmission delay, a packet must wait on average half of the period t_c for the transmission slot. Hence, each node can compute the expected delay of its packets transmitted upstream to the sink without congestion; we denote this value as t_{up}. Under the assumed routing protocol with global TDMA, t_{up} should be constant over time. Under our scheme, the cumulative product of congestion time and $u(t)$ for each application is maintained at each auction node.

The remainder of this section explains how the values stated above are used to calculate the predicted average $u(t)$at the auction node, as well as support the definitions of bids and auctions strategies. We start by considering the first report packet that needs to compete for a transmission slot at the congested node. We assume that congestion is observed at the sink after the first such packet arrives there. We denote t_{start} as the time at which congestion at a given node starts to be seen at the sink and define it as the following:

$$t_{start} = t_{first\_slot\_i} + t_{up}, \tag{4}$$

where $t_{first\_slot\_i}$ is the end of the transmission for the first slot for which packets associated with application i compete. For simplicity, the cumulative products of congestion time and information utility loss for all applications participating in this first auction are set to 0.

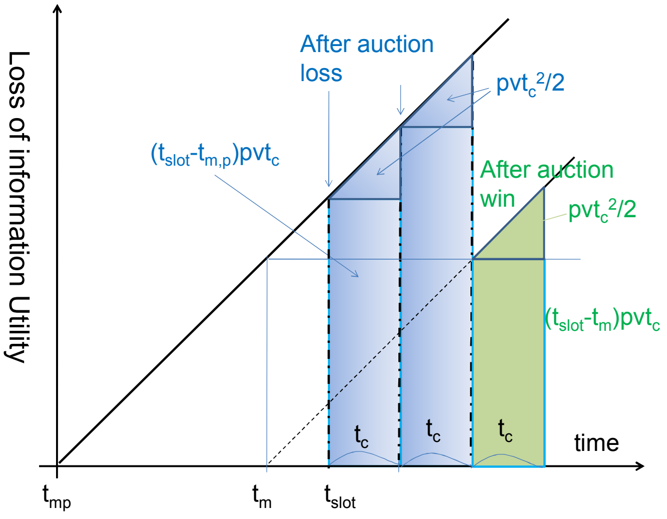

Figure 5. Impact of auction results on loss of information utility.

We now consider all subsequent auctions. We first define the updated delay at the sink as caused by the previous auction, t_f, as follows:

$$t_f = t_{up} + t_{slot} - t_{mp,i}, \tag{5}$$

where t_{slot} is the end of the current auction time slot and $t_{mp,i}$ is the measurement time of the previous target associated with application i. If a report packet loses the current auction, then at time $t_{slot} + t_c$, the cumulative product of time and predicted information utility loss for the corresponding application will be increased by the following value:

$$\Delta u_i^l = \int_{t_{slot}+t_{up}}^{t_c+t_{slot}+t_{up}} u_i(t)dt = (t_f + \tfrac{t_c}{2})p_i v_i t_c. \tag{6}$$

Otherwise, the packet has won the auction. Accordingly, it increases its product in each auction after the win, regardless of whether there is a packet for this application waiting or not, and only until the first loss, by the following increment:

$$\Delta u_i^w = \int_{t_{slot}+t_{up}}^{t_c+t_{slot}+t_{up}} u_i(t)dt = (t'_f + \tfrac{t_c}{2})p_i v_i t_c, \tag{7}$$

where,

$$t'_f = t_{up} + t_{slot} - t_{m,i}. \tag{8}$$

Figure 5 illustrates how the results of the auctions affect the information utility loss. As shown in Figure 5, the loss of information utility for any application increases linearly upon the auction loss. This loss suddenly drops to its initial loss value (caused by the packet's initial *expected* delay from the time of measurement to the sink) after an auction win. This is because the sink can update target information with the data brought in by the packet and replace target location at time $t_{mp,i}$with the more recent location measured at time $t_{m,i}$.

The cumulative product of congestion time and predicted information utility loss is used to compute the packet's bid for the current transmission time slot. The bids for the auctions made on behalf of an application *i* are defined as follows:

$$(a)\quad b_i = \frac{u_i + \Delta u_i^l}{t_{slot} + t_{up} - t_{start}}$$
$$(b)\quad b_i = \Delta u_i^l - \Delta u_i^w = (t_{m,i} - t_{mp,i}) p_i v_i t_c, \tag{9}$$

where 9(a) is used under the system configuration chosen for equalizing tracking applications' information utility loss and 9(b) is chosen for minimizing the total utility loss for all applications. After the auction is completed, all applications adjust their utility loss values by (6) and (7) depending on whether they won or lost.

The bid in 9(a) represents the current predicted average utility loss for an application after losing the current auction. Here, the auctioneer (i.e., congestion node) attempts to equalize the utility losses of applications by selecting the target report packet with the highest bid as the winner. Consequently, the winner's utility loss increases slower after winning than it would after losing considering (6) and (7) and the inequality $t_{m,i}$>$t_{mp,i}$. Hence, utility losses tend to equalize over time. The bid in 9(b) represents the drop of the predicted average utility loss for the winning mission. Here, the auctioneer attempts to minimize the total value of utility loss of all applications by selecting the target report with the highest predicted utility loss drop as the winner. This choice ensures the smallest change of the total value of utility loss for all applications.

It should be noted that u_i is predicted, not actual. For the first target report generated for an application, this is because t_{up} is predicted. For the subsequent reports, the computation of u_i assumes that the delay from the current node to the sink will be the same for the current report packet as it was for the previous one. Hence, the achieved average utility loss metrics in congestion may not be exactly the same for all applications (as will be shown in the evaluation section). However, the difference is small in practical cases since the congestion from the tracking report packets usually arises only in a single node between the sources of the reports and the sink.

IV. Analytical Approach

We present an analytical approach to the congestion management problem addressed by this paper to provide a comparison to the auction-based approach. In the following, we analytically compute the frequency at which each application's target report packets should be forwarded through a congested node.

First, we let f_i denote the frequency at which packets for application *i* are forwarded. Furthermore, we denote $n_i = \left\lfloor \frac{1}{f_i} \right\rfloor \geq 1$ and $fr_i = f_i - n_i$. To a achieve frequency f_i, one report for application *i* must be transmitted every $n_i + 1$ cycles for fr_i fraction of the time and every n_i cycles for $1 - fr_i$ fraction of the time. Counting from the first time slot that the application loses, $t_{slot,i}$, during the waiting time for the j^{th} transmission slot, an application *i* accrues the following product of time and predicted information utility loss:

$$\left(t_{up} + t_{slot,i} + \tfrac{t_c}{2} - t_{m,i} + (j-1)t_c\right) p_i v_v t_c. \tag{10}$$

The delay of application *i*'s packets is $n_i - 1$ time slots for $1 - fr_i$ fraction of the time and n_i time slots for the rest of the time. Hence, the average information loss accrued by application *i* over time $\frac{1}{f_i} - 1$ is:

$$\begin{pmatrix} t_{up} + t_{slot,i} + \frac{t_c}{2} - t_{m,i} \\ +(n_i - 1)\left(1 - \frac{n_i}{\frac{2}{f_i} - 2}\right) t_c \end{pmatrix} p_i v_v t_c, \tag{11}$$

which is difficult to solve for f_i. To simplify, we observe that $n_i - 1$ is approximately $\frac{1}{f_i} - 1$. Hence, simplifying and denoting $x_i = \frac{1}{f_i} - 1$, we derive the following formula for information utility loss of application *i* over time x_i:

$$\left(t_{up} + t_{slot,i} + \tfrac{t_c}{2} - t_{m,i} + x_i\left(\tfrac{t_c}{2}\right)\right) p_i v_v t_c, \tag{12}$$

To equalize the information utility loss across all applications, we will equate the loss for application 1 with loss for each of application $i \in AS$ $(i \neq 1)$, where AS is the set of applications engaged in a given auction and $m_n = |AS|$. As a result, we get $m_n - 1$ equations of the form:

$$\begin{aligned} x_i = x_1\left(\frac{p_1 v_1}{p_i v_i}\right) + \left(\frac{2t_{up}}{t_c} - 1\right)\left(\frac{p_1 v_1}{p_i v_i} - 1\right) \\ +\left(t_{slot,1} - t_{slot,i} - t_{m,1} - t_{m,i}\right)\frac{2p_1 v_1}{t_c p_i v_i}, \end{aligned} \tag{13}$$

just by simple algebraic transformations of extracting x_i from formula (12). The final m_n-th equation, allowing us to solve for x_1 and hence all remaining x_i's comes from the need to utilize the entire bandwidth available at the node. This means that the sum of frequencies of all applications participating in the auction must sum to 1, yielding the following equation:

$$\sum_{i=1}^{m_n} \frac{1}{x_i + 1} = 1. \tag{14}$$

The resulting equation is a polynomial of the degree m_n, which is solvable but requires a time-intensive computation, and yet this will only yield an approximation to the solution as we used a simplified formula (12) instead of the precise expression (11). Moreover, the solution also needs to be changed each time the application priority or the target speed changes, as is discussed in more detail in the next section.

In the case of minimizing the total information utility loss, the sum of average losses defined by of the formula (11) has to be maximized under the m_n constraints of the type $0 \leq x_i \leq 1$ and also by (14), creating a complex non-linear optimization problem, so the details of the solution of this case has not been included here.

V. Evaluation

A. Overview and Configuration

We conducted simulations using the ns-2 framework [12] to compare the performance of the auction-based congestion management mechanism with that of the analytical solution. Specifically, we compared the following four approaches: (1) auction-based information utility loss equalization, (2) auction-based total information utility loss minimization, (3) analytical approach of computing transmission frequencies for equalizing information utility loss, and (4) dividing the communication bandwidth equally among all applications, or the equal bandwidth approach. In this section, we compare the solutions' performance under several scenarios and demonstrate that our auction mechanism achieves its design goals.

Following the previous description of the auction-based mechanism, the simulation used the Spatial TDMA (STDMA) protocol [10], which assigns the same transmission time slots to nodes that cannot interference with each other's wireless communication. Hence, collision free communication was provided, albeit at the cost of limited bandwidth for each sensor node. The two most frequently used methods for slot scheduling in STDMA are node assignment and link assignment, in which the latter configuration assigns actual links, not nodes, to time slots. We implemented a node assignment STDMA configuration, with details described in [13], in which the transmission slots were assigned to nodes in a centralized manner. While some may argue that more efficient solutions exist, we do not focus on the overhead of slot allocation operations in this paper and therefore found it adequate to use the straightforward solution described in [13].

We used the following simulation configuration for this evaluation. The test bed consisted of 80 sensor nodes distributed uniformly distributed over a 500m X 500m terrain. 10 time slots were allocated in the STDMA mechanism. Hence, each node had 1/10th of the total bandwidth available for transmitting packets in non-colliding slots. A sink node hosted three tracking applications (labeled *missions* in the plots) that each tracked a different target. The applications' priorities were set to 5, 2, and 1 for applications 1, 2, and 3 respectively (where an application's priority increases along with number representing it). The simulation was configured such that constant traffic was generated as a result of tracking the three moving objects in the sensor network and the three paths of the targets converged near a single point in the network so as to induce congestion at nearby nodes. Initially, all targets were set to travel at the same speed of 10m/s. All simulations were run for 400 seconds.

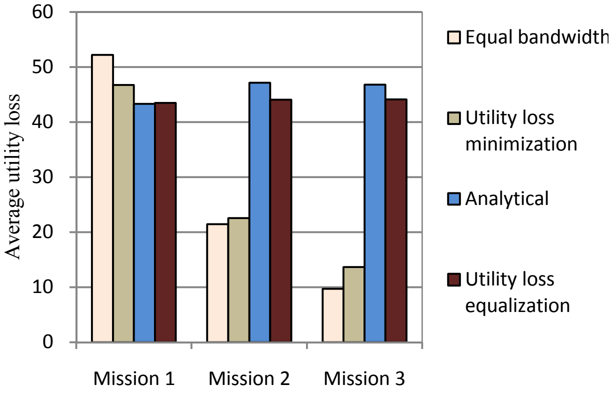

Figure 6. Actual average information utility loss under the compared approaches.

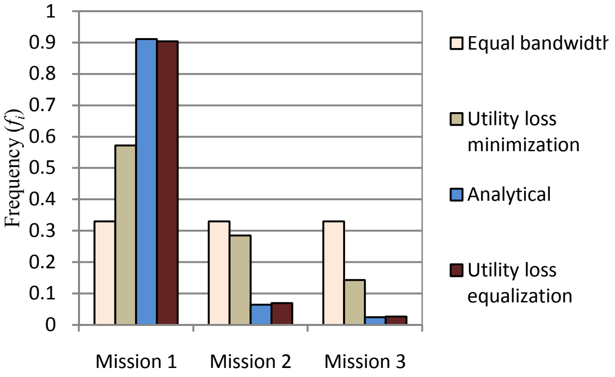

Figure 7. Frequency (f_i) of winning a slot under the compared approaches.

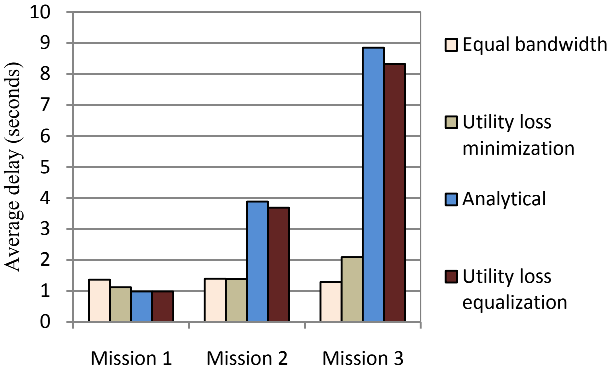

Figure 8. Average delay for three missions under the compared approaches.

Sensor nodes reported measurements about five times per second, each time sending a packet 625 bytes long. At this setting, the packet transmission time was the same as the STDMA slot length. The radio transmission rate for all nodes was set to 250kb/s, which is the same rate as that of the MICAz platform motes [14]. Following the previous settings, each node was allocated 25kb/s of bandwidth for transmission, which was also the same rate of traffic generated by the reporting target information. Thus, network congestion easily occurred when two or more targets' traffic converged at a single node along their paths to the sink.

B. Results

The plot shown in Figure 6 describes the actual average information utility loss for the four compared congestion-management approaches. As the plot shows, the average utility loss among different applications is almost equal using either the analytical solution or our auction-based utility loss equalization approach while they are widely different for the other two approaches. The difference in utility losses is smallest under the auction-based equalization approach because the analytical approach *approximates* the frequency computation in (10). Beyond this benefit, our auction-based equalization approach also yields smaller information utility loss on average across all the tracking applications than the analytical solution. Figures 7 and 8 describe the forwarding frequency and average delay for the three applications under the four compared approaches at the congested node. The three discriminating frequency approaches (excluding equal bandwidth approach) all favor application 1 (which has the highest priority) at the cost of application 2 and especially 3. Still the differences in frequency and delays are smaller for the utility loss minimization approach compared to being more pronounced in the other two discriminating bandwidth approaches. These results show the dominating impact of application priority on resource allocation for those approaches. Overall, while the equal bandwidth approach balances resource usage among applications, it does nothing to balance utility loss among them.

As shown in Figure 9, the average predicted information utility loss of all applications is about equal to a loss of each particular application in the utility loss equalization approach. As previously mentioned, the actual average utility losses of different applications might not be exactly equal to each other since we use the predicted upstream delay t_{up} in the bids calculation. However, as seen in Figure 9, the differences are rather small. Thus, under congestion, our utility loss equalization approach achieves nearly equal utility loss for multiple competing applications.

To demonstrate the responsiveness of our proposed approach to dynamic changes in the simulation scenario, the speed of the target tracked by application 2 was changed from 10m/s to 40m/s in the middle of the simulation. As shown in Fig. 10, applying the equal bandwidth approach, the actual average utility loss of application 2 increased approximately four times after the speed change. Also gathered from Figure 10, the utility loss of application 2 increased more than three times after the speed variation, while utility losses for applications 1 and 3 changed only slightly using the utility loss minimization approach. Moreover, Figure 11 demonstrates that our proposed auction-based utility loss equalization approach and analytical approach kept the average utility loss in congestion nearly equal even after the change of speed.

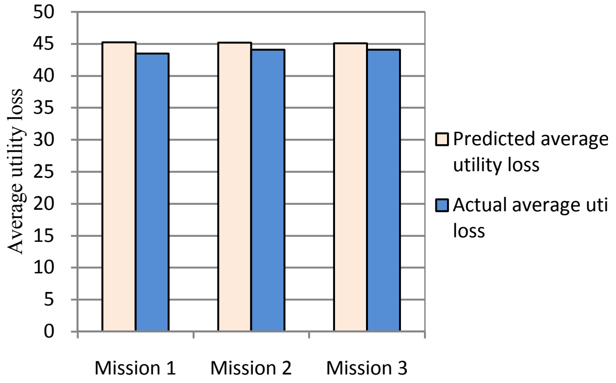

Figure 9. Predicted and actual average utility losses using utility loss equalization approach.

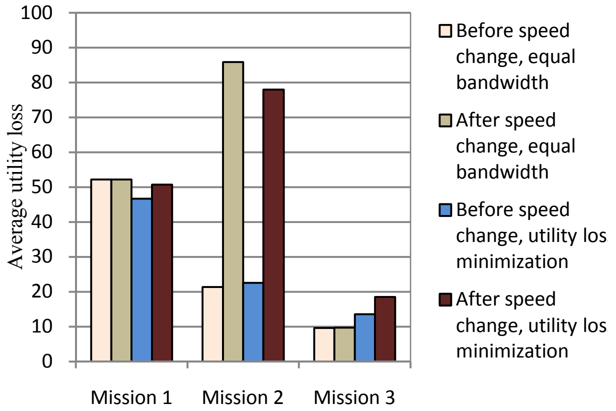

Figure 10. Impact of speed change on the average utility loss under the compared approaches.

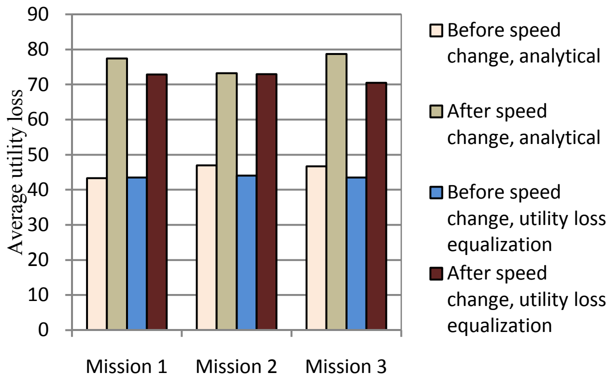

Figure 11. Impact of speed change on average utility loss under analytical approach and utility loss equalization approach.

Though the analytical approach achieves the results nearly as good the other methods, its high computational cost makes it unsuitable to be employed in real sensor network scenario. It is poorly scalable compared to our proposed utility loss equalization approach because the complexity of solving (13) grows quickly with the number of unknowns. In contrast, for auction-based solutions, the overhead is rather small since it consists of just computing the bids at the auctioneer node for each auction and disseminating the application priority values whenever they change. Moreover, increasing the number of applications requires simply adding more bidders in the auction which imposes much smaller overhead than the one incurred by analytical approach which requires higher degree equations to be solved. In addition to easy extensibility, our approach is also easy to implement and deploy even in dynamically changing environments.

VI. Related Works

Target tracking using WSNs is a well-established research area; see for example [15,16,17]. However, to the best of our knowledge, the intersection of QoI and WSN target tracking has largely focused on exploring tradeoffs between detection quality and energy consumption. For instance, in [18], the authors quantify energy-quality tradeoffs under several different strategies for selectively activating sensors for target tracking and show that orders of magnitude savings in energy can be achieved with near-optimal tracking quality using a selective activation strategy. In [19], the authors address challenges of performing low-energy target tracking while maintaining a predefined level of quality of monitoring, particularly in the presence of noise and signal attenuation. Supporting this, a relay-area-based scheme is devised that determines the next sensor to activate when a target is moving, while sustaining the required quality of monitoring and minimizing the overall number of sensors actively tracking the target. In [20], the authors describe an adaptive framework that exploits an application's tolerance to erroneous sensor values. The framework performs target tracking at exactly the accuracy levels needed by the application and adjusts energy usage accordingly. In [21], the authors dynamically optimize the information utility for a given cost of communication and computation.

The cited research largely focuses on accommodating single application (mission, in a military context) and/or single target scenario. In contrast, we consider resolving packet flow congestion and managing tracking quality for multiple applications and targets with different priorities. As alluded to at the beginning of this paper, addressing this particular problem space is particularly beneficial for large-scale pervasive computing target tracking applications, which will most likely be shared among multiple applications. As for the related works, we assume that energy management is addressed by the tracking algorithm itself and therefore is beyond the scope of this work.

Recently, several auction mechanisms have been studied as a method for solving the efficient allocation problem of dynamically changing network resources. In [22], Mackie-Mason and Varian proposed a "smart market" mechanism which as one of the first attempts to apply the auction mechanism to deal with the congestion. Inspired by the "smart market", a Smart Pay Admission Control (*SPAC*) mechanism is introduced in [23] to support QoS differentiation in addition to congestion control. In [24], Lazar and Sermet describe the Progressive Second Price (*PSP*) auction mechanism for network resource (i.e., bandwidth in the case considered in the paper) sharing. In [25], a fair and dynamic auction-based QoS negotiation scheme is proposed to allow users to dynamically negotiate their agreed service levels with their service provider. However, the congestion issues in sensor networks in which the WSN serves multiple applications of different and dynamic priorities remains to the best of our knowledge unaddressed. This is why this paper specifically concentrates on the case of multiple target tracking missions with dynamically changing priorities. We introduce an auction-based mechanism for dynamic and distributed allocation of limited bandwidth to efficiently and fairly resolve transient congestion arising in the considered scenarios.

VII. Conclusions and Future Works

In this paper, we address the problem of providing congestion-management for wireless sensor networks executing several target tracking applications. Hence, the sensor network is shared among many applications that have different priorities and sensitivity to the loss of information utility of the tracking data that they receive. We present two variants of this congestion management technique, each pursuing different global goals. The first variant attempts to equalize information utility loss across different target tracking applications. The second minimizes total utility loss of all tracking applications in the system. We used simulation to highlight some of the benefits of our approach.

We have identified several directions for extending this research. First, we plan to explore integrating this technique with a novel path selection routing protocol [26] to use winning bid values at the next hop nodes to dynamically traverse the path with the smallest predicted utility loss. We also plan to relax the assumption of using a homogeneous wireless sensor network and thus define bidding strategies and utility loss metrics based on the varying capabilities (e.g., detection range, sensing modality) of different sensors. Finally, we plan to explore the effects of using different types of auctions on the efficiency and robustness of the resource allocation.

ACKNOWLEDGMENT

Research was sponsored by US Army Research Laboratory and the UK Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. 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 US Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in 2nd Annual International Workshop on Wireless Internet (WICON'08), August 2006.

[2] M. Rohan, G. Mainland, I. Rose, A. R. Chowdhury, A. Gosain, J. Bers, and M. Welsh, "CitySense: a vision for an urban-scale wireless networking testbed," IEEE International Conference on Technologies for Homeland Security, May 2008.

[3] C. T. Ee and R. Bajcsy, "Congestion control and fairness for many-to-one routing in sensor networks," 2nd International Conference on Embedded Networked Sensor Systems, ACM SenSys'04, pp. 148-161, November 2004.

[4] B. Hull, K. Jamieson, and H. Balakrishnan, "Mitigating congestion in wireless sensor networks," in 2nd International Conference on Embedded Networked Sensor Systems, ACM Sensys'04, November 2004.

[5] C.-Y. Wan, S. B. Eisenman, and A. T. Campbell, "CODA: Congestion detection and avoidance in sensor networks," 2nd International Conference on Embedded Networked Sensor Systems, ACM Sensys'03, November 2003.

[6] K. Karenos, V. Kalogeraki, and S.V. Krishnamurthy, "A rate control framework for supporting multiple classes of traffic in sensor networks,"

26th IEEE International Real-Time Systems Symposium, RTSS 2005,December 2005.

[7] V. Krishna, *Auction theory*, Academic Press, 2002.

[8] Z. Wang, E. Bulut, and B. K. Szymanski, “A distributed cooperative target tracking with binary sensor networks,” IEEE International Conference on Communications Workshops, ICC’08, pp 306-310, May 2008.

[9] Z. Wang, E. Bulut and B.K. Szymanski,”Distributed target tracking with imperfect binary sensor networks,” IEEE Global Telecommunications Conference, Globecom08, Ad Hoc, Sensor and Mesh Networking Symposium, November 2008.

[10] R. Nelson and L. Kleinrock, “Spatial TDMA: a collision-free multihop channel access protocol,” IEEE Transactions on Communications, vol. 33, no. 9, pp. 934-944, September 1985.

[11] J. –S. Lee and B. K. Szymanski, “Auctions as a dynamic pricing for e-services,” in Service Enterprise Integration: An Enterprise Engineering Approach, C. Hsu (ed.), New York: Kluer, pp. 131-156, 2006.

[12] The Network Simulator ns-2, http://www.isi.edu/nsnam/ns/.

[13] S. Ramanathan and E. L. Lloyd, “Scheduling algorithms for multi-hop radio networks,” IEEE/ACM Transactions on Networking, vol. 1, no. 2, pp. 166-177, April 1993.

[14] Crossbow Technology, http://www.xbow.com/.

[15] J. Singh, U. Madhow, R. Kumar, S. Suri, and R. Cagley, “Tracking multiple targets using binary proximity sensors,” International Workshop on Information Processing in Sensor Networks (IPSN’07), pp. 529-538, April 2007.

[16] Q. X. Wang, W. P. Chen, R. Zheng, K. Lee, and L. Sha, “Acoustic target tracking using tiny wireless sensor devices,” International Workshop on Information Processing in Sensor Networks (IPSN’03), April 2003.

[17] W. Zhang and G. Cao, “Optimizing tree reconfiguration for mobile target tracking in sensor networks,” IEEE INFOCOM 2004, pp. 2434-2445, March 2004.

[18] S. Pattem, S. Poduri, and B. Krishnamachari, “Energy-quality tradeoffs for target tracking in wireless sensor networks,” International Workshop on Information Processing in Sensor Networks (IPSN’03), April 2003.

[19] G. He and J. C. Hou, “Tracking targets with quality in wireless sensor networks,” 13th IEEE International Conference on Network Protocols (ICNP’05), Nov. 2005.

[20] X. Yu, K. Niyogi, S. Mehrotra, and N. Venkata-subramanian, “Adaptive target tracking in sensor networks,” Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS’04), Jan. 2004.

[21] F. Zhao, J. Shin and J. Reich, “Information-driven dynamic sensor collaboration for tracking applications,” IEEE Signal Processing Magazine, vol. 19, no. 2, pp. 61-72, March 2002.

[22] J. MacKie-Mason and H. Varian, “Pricing the internet,” in Public Access to the Internet, B. Kahin and J. Keller (eds), Englewood Cliffs, NJ: Prentice-Hall, 1995

[23] J. Shu and P. Varaiya, "Pricing network services." IEEE INFOCOM 2003, vol. 2, pp. 1221-1230.

[24] A.A. Lazar, and N. Sermet, “Auction for network resource sharing”, CTR Technical Report: CU/CTR/TR 468-97-02 Columbia University, November 1997.

[25] T. Taleb, and A. Nafaa, "A fair and dynamic auction-based resource allocation scheme for wireless mobile networks," IEEE International Conference on Communications, ICC'08, pp. 306-310, 19-23 May 2008.

[26] G.G. Chen, J.W. Branch, B.K. Szymanski, “Self-selective routing for wireless ad hoc networks,” IEEE International Conference on Wireless And Mobile Computing, Networking And Communications, 2005, WiMob'2005, vol. 4, pp. 57-64, August 2005.