

Techniques for Improving Opportunistic Sensor Networking Performance

Shane B. Eisenman¹, Nicholas D. Lane², and Andrew T. Campbell²

¹ Columbia University, New York NY 10027, USA
shane@ee.columbia.edu

² Dartmouth College, Hanover NH 03755, USA
{niclane, campbell}@cs.dartmouth.edu

Abstract. A number of recently proposed mobile sensor network architectures rely on uncontrolled, or weakly-controlled mobility to achieve sensing coverage over time at low cost, an opportunistic sensor networking approach. However, this reliance on mobility also introduces a number of challenges. In this paper, we discuss the challenges inherent in this networking paradigm, and describe two composable techniques, sensor sharing and substitution, to make the system more robust in terms of data fidelity and delay. We present a numerical analysis of these techniques, separately and in combination, based on a simple Markov model of an opportunistic sensor network.

Key words: Architecture, Mobility, Modeling, Performance, Wireless Sensor Networks, Opportunistic Networking.

1 Introduction

The recent integration of sensors with personal electronic devices like mobile phones has invited a number of researchers to consider appropriate architectures [4] [12] [23] [1] and applications (e.g., social [16] [25], recreational [8] [9]) for large-scale people-centric sensing systems. Generally, these systems leverage human-carried or vehicle-mounted sensors networked using short/mid-range radios (e.g., ZigBee, WiFi, Bluetooth), and an Internet gateway tier composed of tasking and collection entities. The gateway tier delivers sensing instructions to the mobile sensors on behalf of user applications and accepts incoming sensed data. These proposals rely to some extent on the mobility of humans and their vehicles to get wide area sensing coverage over time with a relatively sparse deployment of heterogeneous mobile sensors. We term sensing with this dependence on uncontrolled mobility *opportunistic sensor networking (OSN)*. While this novel OSN approach can allow large scale sensing at a lower cost compared to an ubiquitous static infrastructure of sensing devices, the opportunistic nature of sensing and communication presents challenges to the fundamental sensor networking operations. In the OSN approach, these operations can be described in terms of opportunistic tasking, opportunistic sensing and opportunistic collection. Opportunistic tasking refers to the process by which a tasking entity

instructs an appropriate mobile sensor to attempt to meet a certain application request. The tasking is opportunistic since there is no guarantee that an appropriate mobile sensor will stay within the radio range of a tasking entity long enough for the tasking operation to complete. By “appropriate”, we refer minimally to a mobile sensor that has the necessary sensing equipment to meet the application request, and may include other requirements (e.g., remaining energy, security clearance, inferred direction of motion). Opportunistic sensing refers to the process by which a mobile sensor that has been assigned a given application task senses the target within the preferred time frame. The sensing is opportunistic since the tasked mobile sensor may not move close enough to the target within the preferred time frame. Opportunistic collection refers to the process by which a mobile sensor that has sensed data in line with the requirements of an application request delivers this data to a collection entity. The collection is again opportunistic since the mobility of the mobile sensor that has sensed the target may not bring it within the radio range of the collection entity and keep it there long enough for the sensed data upload operation to complete.

Noting the aforementioned challenges of OSNs, in [10], we define in situ *sensor sharing* in the context of real sensing applications, and design, implement and experimentally evaluate the system performance for these application scenarios. In this paper, we augment that experimental work with a theoretical analysis of the properties (e.g., scalability and sensitivity to device heterogeneity) of sensor sharing, alone and in combination with *sensor substitution*. These two composable techniques aim to increase the robustness of the OSN paradigm, mitigating the fundamental challenges of uncontrolled human mobility and device heterogeneity to successfully and more expediently complete tasking, sensing and data collection. In particular, we effectively loosen the constraints on which mobile sensors are fit to be tasked for a particular application query. Sensor sharing does this by allowing tasked sensors without the right sensor type for a given sensing task to exploit the resources of others it encounters in the field. Sensor substitution is used in situations where one measurements from one sensor type can act as a reasonable (i.e., within fidelity bounds acceptable to the application) stand-in for another’s. In Section 2, we give a more detailed description of each technique in the context of the OSN challenges it addresses. In Section 3, we provide a baseline numerical analysis of an OSN using a Markov model, and then provide an analysis of the potential improvement provided by sensor sharing and sensor substitution, first separately and then in combination, with respect to the Markov model. Section 4 discusses related work before we conclude.

2 Sensor Sharing and Substitution

To meet the requirements of the application request, a tasking entity must choose an appropriate mobile sensor from the pool of available mobile sensors in its radio range. However, this sensor selection problem is difficult for two reasons: (i) the available pool of mobile sensors is limited by the uncontrolled mobility of humans and vehicles and may not contain an appropriate mobile sensor, and (ii)

it is difficult to predict whether the mobility of a given available mobile sensor will keep it within the tasking entity’s radio range long enough to complete the tasking, and take it to the target region within the preferred time window. For convenience, we term these the *tasking availability problem* and the *tasking prediction problem*, respectively.

The tasking availability problem can be addressed by relaxing the requirements on what constitutes an appropriate mobile sensor, probabilistically increasing the chances that a “taskable” mobile sensor will enter the radio range of the tasking entity within the preferred time window. In particular, we propose to relax the requirement on sensing instrumentation by allowing for *sensor sharing* and *sensor substitution*. With sensor sharing, a mobile sensor A that requires sensor type α , e.g., a CO₂ sensor, to meet the application requirements, but does not itself possess this sensor type (α might be expensive, heavy or rare), can conscript another sensor B that does possess sensor type α to share its sensed data. We envision two scenarios for sensor sharing: (i) sensor A encounters sensor B at the target region, asks B to capture and share data from sensor type α , and receives this data from B , and (ii) sensor A encounters sensor B outside the target region, and leverages the approach proposed in [17] to request and receive data from sensor type α from sensor B . As an example of the former scenario, for an allergen mapping application, a mobile sensor A may be instructed to record particulates in the vicinity of a busy street intersection in the center of the city. Lacking a particle counter, when sensor A finds itself at the target intersection it might broadcast a request for a particulate reading and receive a response from sensor B which is within radio range of A and possesses a particle counter (pull-based sharing). In a more constrained case, a system might be engineered such that at least one mobile sensor Q with sensor type α is present among a group of mobile sensors that might require data from sensor type α . Here, node Q might periodically broadcast readings from sensor type α to the group (push-based sharing). An investigation of a more sophisticated communication protocol in support of sensor sharing, and algorithms to decide which of the possible neighbors is most appropriate to share from is left as future work. We assume an environment where sensors are willing to cooperate to complete sensing tasks. The cooperation might be pro bono, or quid pro quo as in social-network-based sharing [16].

With sensor substitution, a mobile sensor C that requires data from sensor type β to meet the application requirements but does not possess this sensor type, can instead use a substitute method of acquiring equivalent or similar data. We envision two scenarios for sensor substitution that we term *direct substitution* and *indirect substitution*. In direct substitution, a mobile sensor C instructed to collect data from sensor type β uses sensor type γ to collect data equivalent to that given by β , where equivalence does not necessarily include accuracy or precision specifications. For example, for a simple terrain mapping application a mobile sensor may be instructed to measure the slope of a given section of road using a three dimensional accelerometer. Lacking a three axis accelerometer, but possessing a GPS receiver (e.g., the Nokia N95), the mobile sensor

can use interpolation between periodic altitude readings from the GPS receiver to calculate and report the slope of the road [26]. Note that GPS-derived altitude measurements are less accurate and less precise than those provided by a three axis accelerometer, but nonetheless for the case of road slope mapping the GPS receiver can act as a direct substitute for the three axis accelerometer. In indirect substitution, a mobile sensor C instructed to collect data from sensor type β instead reports a combination of sensed data from a number of other sensor types that can be used to generate, e.g., using inference techniques, data which is similar to the requested data. For example, for a simple location mapping application a mobile sensor may be instructed to periodically record latitude/longitude readings from a GPS receiver. Lacking a GPS receiver, but possessing a three axis magnetometer and a three axis accelerometer, the mobile sensor can report direction relative to the magnetic field of the Earth (i.e., a compass reading), and distance calculated as the double integral of the acceleration. With knowledge of starting location, the latitude/longitude values can be approximated using the combination of the direction and distance traveled (i.e., dead reckoning [21]). Note that localization using dead reckoning is less accurate than localization using data from a GPS receiver, but for localization over relatively short distances the combination of a three axis magnetometer and a three axis accelerometer can act as an indirect substitute for a GPS receiver. Another example is recognizing locations that can be uniquely identified by their sensor signature [13], as a substitute to GPS.

Note that the use of sensor sharing or sensor substitution is not exclusive, but rather both are composable blocks that can be used in combination to increase the probability an application request is met. The sensing action that is likely to yield the higher fidelity data is taken. Extending the previous example, if the terrain mapping application requests a slope measurement, we first check if the hardware natively supports high accuracy slope measurement via accelerometer. If not, it tries sensor sharing to see if another nearby node (with appropriate context) will share its accelerometer. If not, it tries sensor substitution to at least get a (lower accuracy) estimate of slope via GPS. Thus, success is achieved if the tasked mobile node has the required sensor, or if it has an appropriate substitute sensor, or if it meets a mobile node at the right place that is willing to share the required sensor, with the commensurate impact on data fidelity. In this way, sensor sharing and substitution help to decouple application design from hardware design (i.e., sensors, on board and opportunistically encountered as people rendezvous in the field), helping to support multiple applications on heterogeneous mobile devices.

3 Analysis and Discussion

To get a preliminary understanding of the baseline performance of an OSN and the theoretical impact of sensor sharing and sensor substitution, we model an OSN scenario using mobile sensor nodes moving according to a discrete Markov model. In the following, we develop the baseline OSN model, and then compare




Fig. 1. An N^2 element Markov chain models a neighborhood where the states represent a grid of points covering the 2-D ground surface of the neighborhood. The grid points are numbered as shown in Figure 1(a). We study a toroidal scenario where nodes may move north, east, south, west, or remain stationary with equal probability (Figure 1(b)); and a more realistic scenario where transition probabilities are derived from the connectivity graph shown in Figure 2.

the performance of sensor sharing and sensor substitution with the baseline in terms of sensing success probability. We numerically evaluate the derived probability expressions to get a sense of the performance boost given by sensor sharing and substitution, and the sensitivity to the number of mobile sensor nodes and the number of sensors per node.

3.1 Model

We model a neighborhood with a Markov chain with a state space SS of size N^2 . We investigate two topology scenarios. In the first, the states represent an $N \times N$ toroidal grid of points covering the 2-D ground surface of the neighborhood. The grid points are numbered as shown in Figure 1(a). We assign the transition probabilities of the Markov matrix \mathcal{M} to allow mobile sensor nodes to move north, east, south, west, or remain stationary with equal probability (see Figure 1(b)). To investigate a more realistic application of the model, in the second topology scenario we overlay the $N \times N$ grid on a physical map of the northwest corner of the Columbia University Morningside campus [6] and assign transition probabilities that respect human pathways in the actual campus. To start, the map (see Figure 2) is sectioned into a 10x10 grid (solid lines), and a connectivity graph (solid squares in the center of each grid square, and dotted lines) of the sections is derived on the basis of walls, doors, pathways, etc. Then, each of the nodes in the graph is treated as a state in the Markov chain, where at each node in the graph each edge (including the implicit self-edge) is taken with equal probability.

Suppose that a query injection point at location σ receives a query from an application running on a back end server at time t for information from sensor type s at target location τ in the grid. Suppose that the request has deadline u , such that the mobile sensor node must be tasked with the request




Fig. 2. A partial campus map showing the northwest corner of the Columbia University campus. The map is sectioned into a 10x10 grid (solid lines), and a connectivity graph (solid squares in the center of each grid square, and dotted lines) of the sections is derived on the basis of doors, pathways, etc. in the actual campus. This graph serves as the basis for the campus scenario probability transition matrix used in Section 3.2.

by the query injection point and arrive at the sensing target location τ by time $t + u$. We assume there are sensors $\mathcal{S} = \{s_1, \dots, s_z\}$ and that each mobile sensor node possesses $r \leq z$ of these sensors through random assignment (e.g., at the factory), with the constraint that the r sensors are distinct (i.e., there are $\binom{z}{r}$ sensor configurations).

3.2 Baseline OSN

We wish to determine the probability of success as a function of the node density, sensor configuration, and sensing deadline in this Markov state space. We start by determining the probability that a suitable mobile sensor node will first visit the query injection point at time $t + k$, $k \leq u$. We assume that t is large enough that the population is well mixed, i.e., the probability that a given mobile sensor node is in a given state i at time t is $p_i = \frac{1}{N^2}$, or equivalently that the initial probability distribution for all states i is $\nu_i = \frac{1}{N^2}$. Let $A_k(i, \sigma)$ denote the event that a single mobile sensor node starting at location i will first visit the query injection point σ at time $t + k$. The probability $F_k(i, \sigma) = Prob(A_k(i, \sigma))$ is given by the recursion

$$F_k(i, \sigma) = \begin{cases} \mathcal{M}(i, \sigma), & k = 1, \\ \sum_{b \in \mathcal{S} - \{\sigma\}} \mathcal{M}(i, b) F_{k-1}(b, \sigma) & k \geq 2. \end{cases}$$

or equivalently as

$$F_k(i, \sigma) = \mathcal{Q}^{k-1} \mathcal{M}(i, \sigma). \quad (1)$$

where $\mathcal{M}(i, j)$ represents the (i, j) th entry of the Markov probability transition matrix, and \mathcal{Q} is the matrix obtained from \mathcal{M} by replacing its column σ with all zeros. Summing over all starting positions i , we have

$$F_k(\sigma) = \frac{1}{N^2} \sum_{i=0}^{N^2-1} \mathcal{Q}^{k-1} \mathcal{M}(i, \sigma), \quad (2)$$

However, we are only interested in those mobile sensor nodes that are equipped with the sensor s that can meet the application/query request (else we assume the mobile sensor node will not be tasked by the query injection point). We model the query arrival process by assuming that a $s \in \mathcal{S}$ is chosen uniformly at random for each application query, such that $Prob(s = s_i) = \frac{1}{z}$. Let B denote the event that a given mobile sensor node's sensor configuration includes the sensor $s \in \mathcal{S}$ specified in the query. Then,

$$Prob(B) = \sum_{i=1}^z Prob(s = s_i) \frac{\binom{z-1}{r-1}}{\binom{z}{r}} = \frac{r}{z}. \quad (3)$$

Let C denote that event that a mobile sensor node equipped with the proper sensor first visits the query injection point at time $t + k$. Since the mobile sensor node's sensor configuration assignment and the mobile node's motion on the neighborhood grid are independent, the probability of event C is simply obtained from Equations 2 and 3 as $Prob(C) = Prob(B) \cdot F_k(\sigma)$.

Once the appropriate mobile sensor is tasked by the query injection point at location σ , it must reach the target at location τ by the sensing deadline $t + u$ in order for the mission to be successful. Thus, we can write the probability of success in the baseline OSN as

$$Prob(Success) = \sum_{l=1}^u Prob(C) \sum_{m=1}^{u-l} F_m(\sigma, \tau), \quad (4)$$

where l is the time when the SAP is visited.

Finally, suppose there are y mobile sensor nodes moving in the neighborhood grid. We write the probability that any of the y mobile sensor nodes succeed as

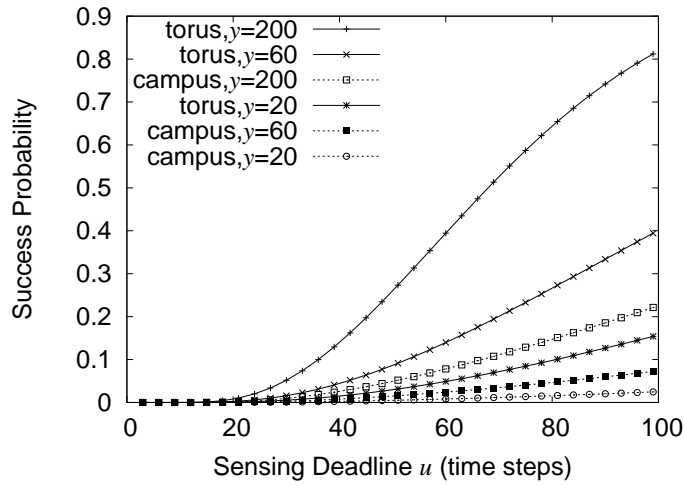
$$Prob^{(y)}(Success) = 1 - (1 - Prob(Success))^y. \quad (5)$$

As the deadline goes to infinity, the probability of success is limited by $1 - (1 - \frac{r}{z})^y$, regardless of the grid dimension N . This limitation imposed by the probability


of a matching sensor configuration ($\frac{z}{z}$) strongly motivates the consideration of both sensor sharing and sensor substitution.

In Figure 3, we compare the success probability for both of the baseline OSN scenarios described previously, i.e., toroid and campus. The tasking point is placed in the lower right corner of the neighborhood (Markov state 99) and the sensing target is placed in the center of the neighborhood (Markov state 45). For both scenarios, we choose a grid granularity of $N = 10$ and use $z = 20$ based on the set of sensors currently used in personal mobile devices and personal sensing systems (viz., camera, microphone, Bluetooth, WiFi, accelerometer, GPS, and temperature samples can be taken from the Nokia N95; the Moteiv Tmote Invent includes light and humidity sensors; magnetometer, galvanic skin resistance, and CO₂ sensors are used in the BikeNet sensing system [8]; the Intel Mobile Sensing Platform includes a barometer and a gyroscope; and the Apple iPhone adds an FM transceiver and an infrared sensor). Future hardware generations will be even more sensor rich, but it is unlikely that any device will integrate all available sensors. In Figure 3(a), we plot Equation 5 versus a range of sensing deadlines u for different numbers of mobile sensor nodes y . Here the ranges of tested y values is meant to represent the rough number of participating users of the OSN in the area of campus shown in Figure 2 throughout the day ($y=200$) and night ($y=20$). We use $r = 3$ to reflect the camera, microphone, and Bluetooth link that nearly all mobile phones on the market possess. The expected trend is evident, as more time is allowed for the sensing to occur the probability of sensing success goes up. It is interesting to note, however, the effect of topology on the success probability. The campus topology from Figure 2 limits the free flow of mobile nodes as compared with the unobstructed toroid scenario, requiring a longer time to reach the same success probability. We also observe (plots omitted due to space constraints) that in the campus scenario, the performance can change substantially depending on the placement of the tasking point σ and the sensing target τ . In Figure 3(a), we again plot Equation 5, but this time versus the number of equipped sensors per node r for different numbers of mobile sensor nodes y when the sensing deadline u is fixed at 50. As intuition suggests, the more sensors each node carries, the higher the chance that a query for a random sensor type will be successful. We also see that, as r is an intrinsic property of a mobile node and not dependent on mobility or node interactions per se, the campus ensemble qualitatively matches the toroid ensemble, albeit at a lower success probability. This is due to the fixed sensing deadline of $u = 50$ used, since from 3(a) we see that the success probability rises more slowly for the campus topology than for the toroid topology.

In the following analysis of sensor sharing and sensor substitution we report only results of the toroid scenario. We do this both to avoid cluttering the figures and also, since we have seen that the placement of the tasking point and sensing target greatly impacts importance in the campus scenario, to get a more general idea of the impact of sharing and substitution.



(a) Plot of the success probability versus the sensing deadline u for various values of y (number of mobile sensor nodes) for two topologies in the baseline OSN scenario. The campus topology from Figure 2 limits the free flow of mobile nodes, requiring a longer deadline to reach the same success probability.



(b) Plot of success probability versus number of sensors r equipped per node for various values of y (number of nodes) for two topologies in the baseline OSN scenario. While the qualitative trends resulting from changes in r are not affected by topology (note that the campus curves follow the same trends as the toroid curves), the overall success probabilities are lower for campus than for toroid.

Fig. 3.

3.3 Sensor Sharing

To analyze sensor sharing³, in the following development we relax the constraint that the query injection points task only those mobile sensor nodes that have the proper sensors to satisfy the application query. We allow a tasked mobile sensor node that arrives at the target region, but does not have the sensor required, to ask other mobile sensor nodes that may be at the target at the same time and do have the required sensor for their samples. In the analysis presented here we use a narrower form of sharing and do not consider the possibility of leveraging the mobile sensor node rendezvous outside of the target region [17].

In Equation 4, $F_k(\sigma, \tau)$ represents the probability that the mobile sensor node tasked by the query injection point first makes it to the sensing target location by time $t + k$. Successful sensor sampling at that point occurs with unity probability in the baseline scenario since the query injection point only tasks a mobile sensor node with the appropriate sensor equipment (as captured by $Prob(C)$ in the same equation). With sensor sharing a tasked mobile sensor node arriving at the target may not have the appropriate sensor. Let D denote the event that successful sampling occurs under these conditions. Then, the probability of D covers two cases: (i) the tasked mobile sensor node arrives at the target and has the required sensor, or (ii) the tasked mobile sensor node arrives at the target and does not have the required sensor but at least one other mobile sensor node at the target concurrently does have the sensor. Assuming there are y total mobile sensor nodes,


$$Prob(D) = F_k(\sigma, \tau) \cdot \left[\frac{r}{z} + \left(1 - \frac{r}{z}\right) \cdot \left(1 - \left(1 - \frac{r}{z}\right) \cdot \frac{1}{N^2} \sum_{i=0}^{N^2-1} \mathcal{M}^k(i, \tau)\right)^{y-1} \right]. \quad (6)$$

Thus, to modify the baseline expression for success probability for sensor sharing, we substitute $Prob(D)$ in for $F_k(\sigma, \tau)$ in Equation 4. Further, since the constraint on the query injection point to task only a node with the required sensor no longer applies, $Prob(C)$ in Equation 4 is replaced with simply $F_k(\sigma)$, and we get

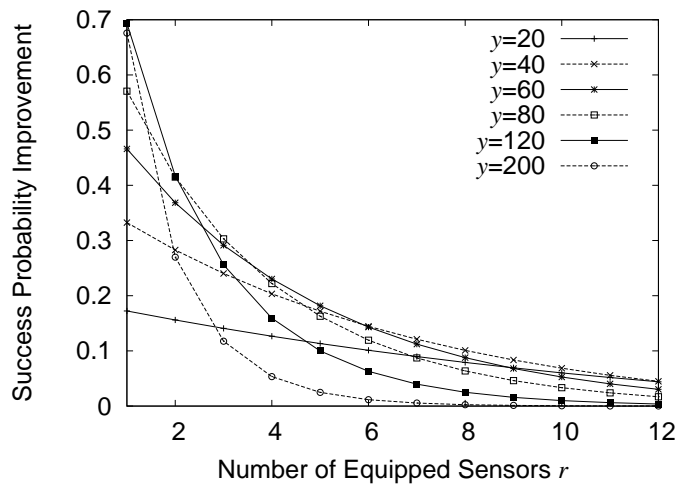
$$Prob(Success) = \sum_{l=1}^u F_k(\sigma) \sum_{m=1}^{u-l} Prob(D). \quad (7)$$

This success probability from Equation 7 is then plugged into Equation 5 to get the final result for sensor sharing. As the deadline goes to infinity, the success probability is limited by approximately $1 - \left(1 - \left(\frac{r}{z} + \left(1 - \frac{r}{z}\right) \cdot \frac{1}{N^2}\right)^y\right)$.

³ In our model we do not treat communication costs for either the baseline case or sharing, and consider only the opportunity for sensing mission success. The sensor sharing implementation in [10] minimizes this cost by limiting communications to a single wireless hop, given the increased complexity and loss probability for multi-hop routing.



(a) Plot of the success probability improvement ratio versus sensing deadline u for various values of y (number of mobile sensor nodes) when sharing is allowed.



(b) Plot of success probability improvement ratio versus number of sensors r equipped per node for various values of y (number of nodes) when sharing is allowed.

Fig. 4.

Comparison. We calculate the normalized sensing success probability improvement given by sensor sharing in our model by first evaluating Equation 5 for both the baseline and sensor sharing cases, and then calculating

$$\frac{(\text{Succ. prob. w/ sensor sharing}) - (\text{Succ. prob. w/ baseline})}{(\text{Succ. prob. w/ baseline})}.$$

We plot this success probability improvement in Figure 4(a) as a function of the sensing deadline u for various numbers of mobile sensors nodes y in the neighborhood when $r = 3$. In Figure 4(b) we plot the success probability improvement as a function of the number of sensors per configuration r for various numbers of mobile sensor nodes y in the neighborhood when $u = 50$. For both plots, we fix $N = 10$, and $z = 20$.

From Figure 4(a), we see that using sharing we can get an improvement of up to 16%, which is mainly dictated by the ratio $\frac{r}{z}$. As the sensing deadline increases, the possible improvement decreases since even without sharing a properly equipped node will eventually go to the tasking point and then to the sensing target. Similarly, we see that with an increasing number of nodes in the system the baseline success probability also increases, reducing the space for possible improvement due to sharing. Figure 4(b) shows an improvement up to 70% when using sharing across all tested conditions. We observe the same general trend with respect to the equipped number of nodes whereby the possible improvement decreases with increasing node density. Similarly, the improvement generally decreases with increasing r , since it is increasingly likely that even without sharing a properly equipped node will visit the tasking point and sensing target before the deadline. Additionally, there exists an interesting interplay between the node density and r when r is low with a relatively small deadline u . In these situations, we conjecture, even when the node density is high, sensor sharing offers a large improvement in success probability over the baseline case since sharing takes advantage of the higher density of sharing candidates in a shorter amount of time, while with few sensors a mobile node in the baseline case must rely on the uncontrolled mobility over time. When the deadline u is extended this effect disappears. It remains to be seen to what extent, if any, this effect persists under a different mobility model.

3.4 Sensor Substitution

Here we address direct sensor substitution (deferring a study of indirect sensor substitution to future work). In the context of the model we are developing we essentially extend the notion of a suitable sensor. We model the fact that a more sophisticated/expensive sensor can to some degree do the job of other simpler sensors, either through direct sensing (e.g., a GPS sensor substituting for a 3-axis accelerometer to measure road slope [26]) or inference (e.g., a CO₂ sensor substituting for a magnetometer to detect car density [8]). The potential for one sensor to substitute for another in some capacity, and the commensurate sensed data fidelity penalty, are dependent on the specific sensors in question. In our initial model, we abstract away these particulars and use p to denote the probability that a given sensor can act as a direct substitute for the sensor s specified by the application query, incorporating in p the probability that the

corresponding loss of fidelity (if any) is within bounds acceptable to the application. A study of empirically generated correlation functions between various common sensor types will be the subject of our future work. We write

$$\text{Prob}\{i \simeq j\} = \begin{cases} 1, & i = j, \\ p, & i \neq j, \end{cases} \quad (8)$$

where $i \simeq j$ denotes that sensor i is a suitable substitute for sensor j . Letting B' denote the event that a given mobile sensor's configuration of r sensors includes a suitable (substitute) sensor, given a randomly chosen $s \in \mathcal{S}$ we have

$$\begin{aligned} \text{Prob}(B') &= \text{Prob}\{i \simeq j \mid i = j\} \text{Prob}\{i = j\} + \text{Prob}\{i \simeq j \mid i \neq j\} \text{Prob}\{i \neq j\} \\ &= 1 \cdot \frac{r}{z} + (1 - (1 - p)^r) \left(1 - \frac{r}{z}\right). \end{aligned} \quad (9)$$


The modified success probability when using direct sensor substitution is given by simply substituting $\text{Prob}(B')$ in for $\text{Prob}(B)$. Comparing Equation 9 with Equation 3, it is clear that the benefit of direct sensor substitution is given by the second term in Equation 9.

Comparison. We calculate the normalized sensing success probability improvement given by sensor substitution in our model by first evaluating Equation 5 for both the baseline and sensor substitution cases, and then calculating


$$\frac{(\text{Succ. prob. w/ sensor subst.}) - (\text{Succ. prob. w/ baseline})}{(\text{Succ. prob. w/ baseline})}.$$

We plot this success probability improvement in Figure 5(a) as a function of the sensing deadline u for various values of the substitution probability p when $y = 20$ and $r = 3$. In Figure 5(b) we plot the success probability improvement as a function of the number of sensors per configuration r for various numbers of mobile sensor nodes y in the neighborhood when $p = 0.02$ and $u = 50$. For both plots, we fix $N = 10$, and $z = 20$.

From Figure 5(a), we see success probability improvements ranging from 25% to 155% for relatively modest substitution probabilities, over the first 100 time steps. As before, the gain decreases over time due to the fact that eventually a properly equipped mobile sensor will visit the tasking point and then the sensing target even when sensor substitution is not used to increase the pool of "properly equipped" sensors. Figure 5(b) shows that an improvement over the baseline is achieved across all tested node densities and numbers of equipped sensors. Similar to the sensor sharing results in Section 3.3, we observe the possible improvement decreases with increasing node density, and the improvement also decreases with increasing r , since it is increasingly likely that even without substitution a properly equipped node will visit the tasking and sensing target before the deadline. The feature of overlapping curves seen in Figure 4(b) is not seen in Figure 5(b) since substitution does not involve interactions between



(a) Plot of success probability versus sensing deadline u for various values of p (substitute probability), when direct substitution is allowed.



(b) Plot of success probability versus number of sensors r for various values of y (number of mobile sensors), when direct substitution is allowed.

Fig. 5.

nodes but rather operates independently on each node. Generally, the improvement given by sensor substitution is less sensitive to node density, as indicated by the tighter envelope of the curve ensemble, than is sensor sharing. On the other hand, substitution is more sensitive than sharing is to r , as indicated by

the larger delta across the tested range of r (e.g., 0.28 for sensor substitution and 0.13 for sensor sharing when $y = 20$).

3.5 Composing Sensor Sharing and Sensor Substitution

In the following we demonstrate the cumulative benefit of composing both sensor sharing and sensor substitution. In this scenario, again the constraint on exactly matching the preferred sensor is relaxed and all mobile sensor nodes reaching the tasking point before the deadline are tasked. Sensing success occurs if a node reaches the sensing target before the deadline, and: (i) it possesses the preferred sensor; or (ii) it does not have the preferred sensor but it is equipped with a suitable substitute sensor; or (iii) it does not have either the preferred or a substitute sensor, but it is able to share sensor readings from another mobile node at the sensing target.

We calculate the normalized sensing success probability improvement given by sensor substitution and sharing in our model by first evaluating Equation 5 for both the baseline, then applying both sensor sharing and sensor substitution, and calculating

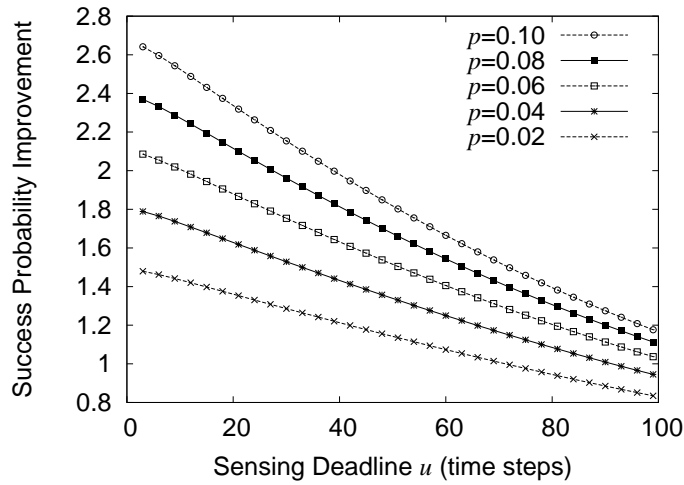
$$\frac{(\text{Succ. prob. w/ sharing and subst.}) - (\text{Succ. prob. w/ baseline})}{(\text{Succ. prob. w/ baseline})}.$$

We plot this success probability improvement in Figure 6(a) as a function of the sensing deadline u for various values of the substitution probability p when $y = 20$ and $r = 3$. In Figure 6(b) we plot the success probability improvement as a function of the number of sensors per configuration r for various numbers of mobile sensor nodes y in the neighborhood when $p = 0.02$ and $u = 50$. For both plots, we fix $N = 10$, and $z = 20$.


From Figure 6, the benefit of enabling both sharing and substitution is clear. In Figure 6(a), the success probability improvement over the baseline is up to 270%, far outpacing the gains given by sharing (up to 16%) or substitution (up to 160%) alone, over the same range of sensing deadlines (u) (c.f. Figures 4(a) and 5(a)). Similarly, Figure 6(b) shows a success probability improvement across the same range of equipped sensors (r) over the baseline of up to 140% as opposed to only 70% for sharing and 35% for substitution alone. When both techniques are enabled, even when a node is equipped neither with the preferred sensor nor with a suitable substitute sensor, sensor sharing can probabilistically help to compensate, and vice versa.

4 Related Work

In the relatively new area of opportunistic sensor networking, there is little published work specifically addressing its challenges. However, more generally, there are alternative approaches for dealing with missing data points, both spatially and temporally. For example, Bayesian nets and other interpolation techniques



(a) Plot of success probability versus sensing deadline u for various values of p (substitute probability), when both sensor sharing and direct substitution are allowed.



(b) Plot of success probability versus number of sensors r for various values of y (number of mobile sensors), when both sensor sharing and direct substitution are allowed.

Fig. 6.

can be used to infer missing data (e.g., [11]) when nodes lack access to the sensing hardware they require, at the time they require it. This approach allows for a reasonable approximation of missing data (within the accuracy of the sensor and sensed phenomenon models), but in the end is still an approximation. In contrast, both sensor sharing and substitution provide access to the real sen-

sensor samples, if supported by opportunistic rendezvous, without requiring any heavyweight computation or access to a central data store. Clearly, sharing and substitution can be used in concert with approximation techniques.

While in situ sensor sharing in opportunistic sensor networks is unaddressed in the literature, sensor sharing has been studied in the non-opportunistic networking setting. The authors of [19] present a mechanism enabling robot team members to share sensor information to achieve tightly-coupled cooperative tasks. The approach uses dynamically configured schema to route information among preformed groups to generate different cooperative control strategies. Conversely, our notion of sensor sharing relies only on completely opportunistic mobile node rendezvous. Analogous to mobile nodes that can carry a limited number of sensors, [5] explores the possibility of sensor sharing between integrated wired aerospace subsystems to reduce system part count. The authors of [15] explore the same design concept for reconfigurable sensor networking platforms. These subsystems are statically connected and therefore do not have to deal with the same challenges as the OSN domain we consider. When one considers the radio as a sensor, there are many examples of sensor data being shared between nodes as part of a system control mechanism. As an example of this type of sharing, in [24], the authors propose a system of mobile nodes that adjust their communication protocol parameters based on collaboratively-sensed environmental conditions. Using their radio receivers as sensors these nodes measure and share RSSI readings of WiFi AP beacons. While this sharing is in situ and the networking can be opportunistic, it is notable that in these cases each of the nodes possesses the sensor (the radio). Thus, this type of sharing is really by-design in-network information sharing, and not sensor sharing to meet an ad hoc application query. The authors of [18] and [22] have proposed a conceptual architecture and a prototype implementation to facilitate the sharing of sensor data among scientists and others once the data has been harvested to the back end. In contrast to this type of data sharing, the *sensor sharing* we propose takes place between mobile nodes in situ. These two types of sharing are complementary, as sensor sharing helps to provide the data streams that can be shared on the back end systems.

The concept of sensor substitution, though straightforward, has received little explicit attention in the sensor networking literature. This is likely due to the fact that to date the bulk of sensor network deployments have been engineered to meet the needs of a single application. We believe that both sensor sharing and sensor substitution will receive more attention as sparse, multi-application, device-heterogeneous, mobile sensor networks gain momentum. We are motivated by parallels in the health sciences domain where sensor substitution is an area of active interest for the treatment of human disabilities arising from the damage or decay of primary sensory systems. In [2], the authors summarize a study on people with balance disorders that concludes that sound may substitute, at least partially, for the lack of vestibular sensory information to control postural sway in stance.

In a broader context, methods for gaining ad hoc access to required resources (e.g., speakers, projectors, printers) have been considered by researchers in the pervasive networking community, in support of smart environments [20] and nomadic applications like “smart projector” and mobile “music service” [3].

5 Conclusion

While mobility in the context of OSNs enables sensing coverage at a lower deployment cost compared to an ubiquitous static sensor deployment, this uncontrolled mobility also poses a number of challenges related to tasking, sensing and data collection. With sensor sharing and sensor substitution we have proposed two techniques that can be used together or alone to improve the probability of successfully and more expediently completing these activities. The initial numerical evaluation is encouraging, showing non-negligible gains through low complexity techniques, warranting further study.

Acknowledgment

This work is supported in part by Intel Corp., Nokia, NSF NCS-0631289, ARO W911NF-04-1-0311, and the Institute for Security Technology Studies (ISTS) at Dartmouth College. ISTS support is provided by the U.S. Department of Homeland Security under award 2006-CS-001-000001, and by award 0NANB6D6130 from the U.S. Department of Commerce. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of any supporting agency.

References

1. T. Abdelzaher, Y. Anokwa, P. Boda, J. Burke, D. Estrin, L. Guibas, A. Kansal, S. Madden, J. Reich. Mobiscopes for Human Spaces. *IEEE Pervasive Computing*, Vol. 6, No. 2, pp. 20-29, Apr-Jun 2007.
2. L. Barclay and C. Vega. Audio-Biofeedback May Improve Stability for Patients With Bilateral Vestibular Loss. *Medscape Medical News*, <http://www.medscape.com/viewarticle/508833>, 2005.
3. P. Basu, K. Wang, and T. D. C. Little. Dynamic Task-based Anycasting in Mobile Ad Hoc Networks. In *Mobile Networks and Applications*, Vol 8, No 5, Oct 2003.
4. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo and R. A. Peterson. People-Centric Urban Sensing (Invited Paper). In *Proc. of 2nd ACM/IEEE Int'l Conf. on Wireless Internet (WiCon 2006)*, Boston, Aug 2006.
5. C. M. Carlin and W. J. Hastings. Propulsion control sensor sharing opportunities. In *Proc. of the 21st AIAA Aerospace Sciences Meeting*, Reno, Jan 1983.
6. Columbia University Morningside campus map (retrieved Jan 1, 2008). http://www.columbia.edu/about_columbia/map/images/maps-large/mudd.gif

7. A. Deshpande, C. Guestrin, W. Hong and S. Madden. Exploiting Correlated Attributes in Acquisitional Query Processing. In *Proc. of the 21st Int'l Conf. on Data Engineering (ICDE'05)*, pp. 143–154, Tokyo, Apr 2005.
8. S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn and A. T. Campbell. The BikeNet Mobile Sensing System for Cyclist Experience Mapping. In *Proc. of 5th ACM Conf. on Embedded Networked Sensor Systems*, Sydney, Nov 2007.
9. S. B. Eisenman and A. T. Campbell. SkiScape Sensing (Poster Abstract). In *Proc. of ACM 4th Int'l Conf. on Embedded Networked Sensor Systems*, Boulder, Nov 2006.
10. S. B. Eisenman, N. D. Lane, and A. T. Campbell. Quintet: Orchestrating Neighborhood Sensors to Increase Sensing Fidelity through Opportunistic Sharing. *Dartmouth Computer Science Technical Report*, Dec 2007.
11. E. R. Hruschka Jr., E. R. Hruschka and N. F. F. Ebecken. Applying Bayesian Networks for Meteorological Data Mining. In *Proc. of the 25th SGAI Int'l Conf. on Innovative Techniques and Applications of Artificial Intelligence*, Cambridge, Dec 2005.
12. B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan and S. Madden. CarTel: A Distributed Mobile Sensor Computing System. In *Proc. of 4th Int'l Conf. on Embedded Networked Sensor Systems*, pp. 125–138, Boulder, Nov 2006.
13. N. D. Lane, H. Lu and A. T. Campbell. Ambient Beacon Localization: Using Sensed Characteristics of the Physical World to Localize Mobile Sensors. In *Proc. of 4th Workshop on Embedded Networked Sensors*, Cork, Jun 2007.
14. K. Lorincz and M. Welsh. MoteTrack: A Robust, Decentralized Approach to RF-Based Location Tracking. In *Springer Personal and Ubiquitous Computing*, Vol 11, No 6, pp. 489–503, Aug 2007.
15. D. Lymberopoulos, N. B. Priyantha and F. Zhao. mPlatform: A Reconfigurable Architecture and Efficient Data Sharing Mechanism for Modular Sensor Nodes. In *Proc. of Int'l Conf. on Information Processing in Sensor Networks*, pp. 128–137, Cambridge, Apr 2007.
16. E. Miluzzo, N. D. Lane, S. B. Eisenman and A. T. Campbell. CenceMe - Injecting Sensing Presence into Social Networking Applications. In *Proc. of EuroSSC 2007, LNCS4793*, pp. 1–28, Kendall, Oct 2007.
17. E. Miluzzo, N. D. Lane and A. T. Campbell. Virtual Sensing Range (Poster Abstract). In *Proc. of 4th ACM Conf. on Embedded Networked Sensor Systems*, pp. 397–398, Boulder, Nov 2006.
18. A. Parker, S. Reddy, T. Schmid, K. Chang, G. Saurabh, M. Srivastava, M. Hansen, J. Burke, D. Estrin, M. Allman and V. Paxson. Network System Challenges in Selective Sharing and Verification for Personal, Social, and Urban-scale Sensing Applications. In *Proc. of HotNets-V*, Irvine, Nov 2006.
19. L. E. Parker, M. Chandra and F. Tang. Enabling Autonomous Sensor-Sharing for Tightly-Coupled Cooperative Tasks. In *Multi-Robot Systems. From Swarms to Intelligent Automata Volume III*, Springer Netherlands, pp. 119–130, 2005.
20. A. Rakotonirainy and G. Groves. Resource Discovery for Pervasive Environments. In *Proc. of the 4th Int'l Symp. on Distributed Objects and Applications*, pp. 866–883, Oct 2002.
21. C. Randell, C. Djallil and H. Muller. Personal Position Measurement Using Dead Reckoning. In *Proc. of 7th IEEE Int'l Symp. on Wearable Computers*, pp. 166–173, Oct 2003.

22. S. Reddy, G. Chen, B. Fulkerson, S. J. Kim, U. Park, N. Yau, J. Cho, J. Heidemann and M. Hansen. Sensor-Internet Share and Search-Enabling Collaboration of Citizen Scientists. In *Proc. of the ACM Workshop on Data Sharing and Interoperability on the World-wide Sensor Web*, Cambridge, Apr 2007.
23. M. Srivastava, M. Hansen, J. Burke, A. Parker, S. Reddy, G. Saurabh, M. Allman, V. Paxson and D. Estrin. Wireless Urban Sensing Systems. CENS Technical Report #65, 2006.
24. G. D. Troxel, A. Caro, I. Castineyra, N. Goffee, K. Z. Haigh, T. Hussain, V. Kawadia, P. G. Rubel and D. Wiggins, Cognitive Adaptation for Teams in ADROIT (Invited paper). In *IEEE Global Communications Conference*, Washington D.C., Nov 2007.
25. V. Tuulos, J. Scheible and H. Nyholm. Combining Web, Mobile Phones and Public Displays in Large-Scale: Manhattan Story Mashup. In *Proc. of the 5th Int'l Conf. on Pervasive Computing*, Toronto, May 2007.
26. D. E. Voyadgis and W. H. Ryder. Measuring the Performance of Ground Slope Generation. In *Proc. of the 16th Annual ESRI User Conf.*, May 1996.