Socially-Aware Market Mechanism for Participatory Sensing

Buster O. Holzbauer
Rensselaer Polytechnic Institute
110 8th St
Troy, NY 12180
holzbh@cs.rpi.edu

Boleslaw K. Szymanski
Rensselaer Polytechnic Institute
110 8th St
Troy, NY 12180
szymansk@cs.rpi.edu

Eyuphan Bulut
Cisco Systems
2200 President George Bush Highway
Richardson, TX 75082
ebulut@cisco.com

ABSTRACT

By using personal mobile devices, participatory sensing provides an alternative to deploying dedicated mobile nodes to perform data acquisition tasks in areas where human carriers are already present. However, this application comes with some unique challenges. In this paper, we study the challenges resulting from social concerns of participants in a participatory sensing application and propose a socially aware auction mechanism to address them. Through simulations, we compare the proposed mechanism to several existing mechanisms. The results demonstrate that, with the right configuration, our proposed mechanism can decrease the cost for the data sinks and decrease the privacy loss and battery usage of participants while preserving their participation.

Categories and Subject Descriptors

C.2.4 [**Computer Communication Networks**]: Distributed Systems

General Terms

Design

Keywords

participatory sensing, privacy, market mechanism

1. INTRODUCTION

Participatory sensing relies on participants contributing observed data to build a data collection [1] [2]. It can also be considered as a potentially commercializable or voluntary-based type of a distributed sensor or mobile network.

There are many examples of participatory sensing applications in real-world [3] [4] [5]. We specifically focus on the set of applications where there is a sensing campaign with the objective of maintaining a set of sufficiently recent sensor readings (such as vehicle emissions) from adequately covered area. Since mobility is an energy-intensive activity, we rely on human mobility to move the sensors. Smartphones with sensing capabilities are examples of devices that are routinely carried by people. Thus, in this paper we consider the sensor platform to be personal phones and mobility to be restricted to human mobility.

Human-centric sensing invokes several challenges including the mitigation of social concerns. One such challenge is balancing power usage between normal activities such as phone calls, and sensor readings. Another one is how to overcome recruitment constraints such as budget restrictions or human constraints such as motivation, availability, and privacy. Since applications that we consider in this paper require data about spatio-temporal coordinates of collection, privacy is particularly important.

In this paper, we propose to use a reverse auction mechanism to address both the retention of participants and their social concerns arising from participation (e.g. Figure 1) at the same time. Participation rewards requires that the frequency of sensing for a participant is not too low while social concerns require the frequency of sensing is not too high.

The rest of the paper is organized as follows. In Section 2, we discuss design considerations in market mechanisms, and describe some existing auction mechanisms. In Section 3, we give the details of the system. Next, in Section 4, we describe the proposed auction mechanism that aims to address social challenges of participatory sensing. We perform extensive simulations in Section 5 and evaluate the performance of the proposed mechanism. Finally, we close the paper with conclusions in Section 6.

2. BACKGROUND

2.1 Design Considerations

In a reverse auction mechanism, participants configure their devices to make bids whenever a data sink notifies the device that an auction has started. The data sink then uses an auction mechanism to choose which participants will take sensor readings and send data, as well as deciding the amount of incentive to award each participant chosen to provide data. Participants have the objective of maximizing their incentive over the course of their participation, while each data sink has the goal of minimizing the incentive paid while satisfying constraints on data age and coverage.

When using a reverse auction, the type of the auction mechanism used is vital as it will decide which participants

Figure 1: A participatory sensing system

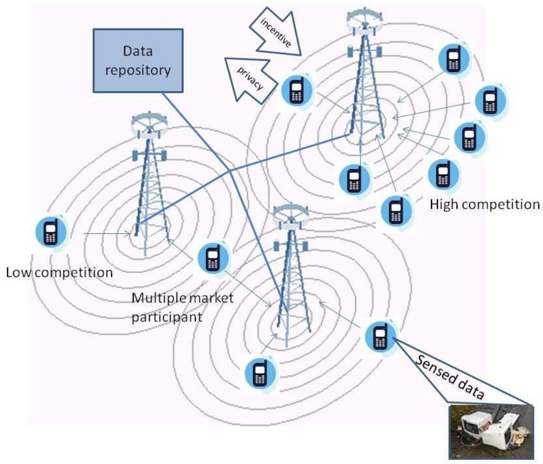

win each round and what their payouts are. A mechanism that is poorly designed may be difficult for users to use effectively, may cause significant excess payout, and may lead to decreased participation. The participants should be able to make informed decisions about their bids, so the mechanism should be transparent (i.e., the mechanism details should be made public), and easy to understand. The number of parameters should be kept small to make it easier for users to pick a price point and strategy. Moreover, the mechanism should be incentive compatible to simplify bidding. With incentive compatible mechanisms, all parameters relating to user opinion can be encapsulated into a single parameter (which we will refer to as the base price or p_{base}).

The auction mechanism can also be used to regulate fairness and winner frequency. A mechanism selecting winners from a small subset of participants may cause other participants to leave the campaign or the local market. This can be desirable if only participants with unrealistic expectations are leaving, but is harmful if participants that would otherwise remain and have reasonable budget requirements leave. Thus, a mechanism should become less selective if participation is rapidly declining due to participants winning too infrequently. If participants win too frequently, they may raise prices either as a bidding strategy, or due to increased loss of privacy (i.e., the privacy loss from providing data (including location) two months apart is much smaller than in case of providing data two minutes apart). This may lead to participants leaving campaign or the budget of the sensing campaign being exhausted much faster than necessary. Thus, the mechanism should regulate both participation and incentive.

2.2 Previous Mechanisms

Next, we will give the details of some auction mechanisms previously proposed in literature.

First Price Auctions: A classic auction is the first price sealed bid auction. All participants that are competing for incentive in a given auction submit bids simultaneously to provide one of the N required sensor readings. The auctioneer (data sink), which in our example problem is a cell tower, waits until all bids are received and then select the N cheapest bids received as winners. This mechanism is very easy for the auctioneer to run. However, from a participant's standpoint this type of reverse auction is difficult to understand. Since winners pay what they bid, the bid cannot reflect their true valuation so the first price reverse auction is not incentive compatible. The bidding strategy in a first price auction is dependent on a given participant's perception of all other participants' strategies. Making a decision about what amount to bid is non-trivial and may deter accurate bidding and participation.

Reverse Auction based Dynamic Price incentive mechanism with Virtual Participation Credit (RADP-VPC): RADP-VPC is a modification of the first price auction that is designed with reverse auctions in mind [8]. For every consecutive round in which participant loses, the corresponding bid is decreased by a constant α. The auctioneer ignores this virtual participation credit when awarding incentive, so the participant's odds of winning are improved without decreasing payout. The auction is run in all other aspects like a first price auction. To promote rejoining the auction by dropped out bidder, the mechanism notifies them of the highest winning bid in the current round.

Participation Incentive Generalized Vickrey Auction (PI-GVA): An alternative to the first price auction is the Generalized Vickrey Auction, which is incentive compatible. The Participation Incentive GVA mechanism computes for each participant a score that encapsulates the expected and actual number of times a participant has won, as well as the corresponding bids [9]. PI-GVA has a parameter μ which determines the weight of the average bidding term. The auctioneer uses the scores in computing prices for selecting winners. PI-GVA as originally described is a forward auction, to make it a reverse auction bids are negated and the smallest scores are used. Despite the complexity compared to the previous two mechanisms, PI-GVA is easy for participants to understand if they realize it is incentive-compatible. PI-GVA also has the desirable property of being effective even in recurring auctions. This means that the mechanism retains sufficiently many participants to maintain competition even as participants learn information as the recurrent auction progresses.

3. SYSTEM DESCRIPTION

3.1 Spatial and Temporal Details

We consider a system in which there exist one or more data sinks which operate independently of each other. The coverage areas of data sinks are not guaranteed to be mutually exclusive. A data source may be in transmission range of several data sinks at any given time. In such a case, no sink has information about to how many other data sinks a data source is attempting to report, or how many auctions run by other sinks it has won. In simulating this system, a clock is kept and times assigned to data source movements. Data sinks require data whenever they have insufficient data about the area they cover. Data recorded has an aging time which indicates how long the data can be used. Once the aging time has elapsed for a given data point, the data sink will act as though that data point does not exist when determining if additional polling is necessary. Currently, we assume that the data sink does not use any spatial information about data source locations to select the winning sources of a round. Instead the data sink looks at its area of coverage which is defined on a discrete grid as n_{cover} and requires $f_{coverage} \times n_{cover}$ data sources to report. $f_{coverage}$ is a coverage factor which can be adjusted based on the level of overprovisioning in the system[1].

Since battery levels and subsequently power consumption are important time-varying details, we incorporate them into the system. Power consumption caused by sensing and theoretical protocol overhead are far greater than baseline power consumption for just staying connected to a tower. As a result, we model power cost for each sensing event to be a constant application specific value. Moreover, since we focus on participatory sensing, we do not consider the decrease in battery for baseline operations, phone calls, or other embedded applications. We also do not account for recharging batteries. Each experiment is run over such a span that no battery is depleted.

3.2 Participation

Each participant is assigned a true value, denoted as p_{base}, randomly chosen from a uniform distribution. When a participant wins an auction (a data sink soliciting measurements chooses to use that participant's data source), there is a probability P_{inc} that the participant's bid will raise in future auctions. When a participant loses an auction, there is a probability P_{dec} that the participant's bid will decrease in future auctions, but not below p_{base}.

Each participant i is also assigned a tolerance level to losses, controlled by parameter β_i, chosen from a shifted exponential distribution as $\beta_i = \beta_{min} + e^{-x}$, where x is uniformly randomly chosen from $[0, 1]$. Like in [8], it is used in a Return on Investment (ROI) calculation that decides if the participants drops out of auction or not. However, ROI computation is delayed until a participant tries to decrease the bid below p_{base}. We refer to this event as the 'minimum state'. We do not count rounds prior to minimum state in the ROI calculation. This gives participants time to stabilize in the early stages of the auction, so even participants with a low tolerance to losses participate for several rounds to assess the market competition.

Whenever the ROI value for a participant falls below 0.5 (which means that the participant received less than half expected wins), that participant and the associated data sources are no longer considered in any auctions for the duration of the simulation. We also assume that the participants will never rejoin an auction after quitting[2].

4. PROPOSED MECHANISM

In this section, we propose a new reverse auction mechanism which aims to address the social concerns of participants. We call it, 'Privacy, Power, and Participation aware Auction Mechanism (P3AM)'. It is based on a first price auction but varies both in terms of the modification of participants' bids and selection of winners. The mechanism has a parameter $P_{cheapest}$ that determines how many of the winners are chosen based on having the cheapest bids. $P_{cheapest}\%$ of the N winners are determined in this manner. The other winners are selected from the participants that have the lowest current RoI with ties resolved by selecting the lowest p_{base}. This provides encouragement to participants who are most in danger of dropping out recently while still discouraging users with exceptionally high prices. By doing so, participation is encouraged which means that the level of competition stays high in a recurring auction.

The bid that the auctioneer sees is

$$b = p_{base} \times f_{power} \times f_{privacy}.$$

In our current experiments, f_{power} is a hyperbolic function based on the current power level λ_{pow} which is between [0,1]:

$$f_{power} = \frac{1}{\lambda_{pow}}.$$

$f_{privacy}$ is an exponential function based on t_{last} which is the time elapsed since a given participant last won an auction for the data sink running the current auction. In our system t_{last} is expressed in the number of auction rounds.

$$f_{privacy} = \begin{cases} 2^{(1/t_{last})} & t_{last} > 0 \\ 1 & t_{last} = 0 \end{cases}$$

Here note that with the introduction of f_{power} and $f_{privacy}$, our goal is to address the social concerns of users. If the power of a participant's device is decreasing, then the participant should have a higher bid since power is more scarce for them. We chose a hyperbolic scheme for f_{power} because depending on the current power level for a participant, the value of 1 unit of power may differ. For example, when a power level is near 100%, a loss of 1% is not a threat to continued operation of the participant's device. However, if only 10% of the participant's battery is left, then the same 1% is much more valuable. Note that, even though battery level is a hardware concern, the valuation of battery level based on the behaviors and perceptions of participants and the current state of their devices is a social concern.

If a device belonging to a participant is sensing frequently, the participant should be awarded higher incentive. This recognizes loss of privacy associated with sending data with location of device to the sink. Since different participants may have differing valuations of privacy we did not use the solution from the previous work by Danezis et. al. which uses a fixed incentive to compensate participants for such costs [10]. Instead, we have an exponentially increasing term in $f_{privacy}$ so the mechanism inherently will not choose a participant many times unless the participant has a low valuation of privacy and low bids. Although participants cannot explicitly change their privacy valuation, they can incorporate their privacy valuation into the value they select for p_{base}.

5. SIMULATIONS

To evaluate the performance of proposed mechanism, we did simulations with different configurations over 950 auction rounds with 100 runs. The inter-auction round time was set empirically as 1/950 of the time needed by at least one participant to exhaust its battery below 10% level. In each experiment, a 11x11 grid (as shown in Figure 2) was used with 4 detectors (i.e., cell towers) positioned approximately in the center of each quadrant. Each detector covers 59 vertices (49% of the grid), meaning there are vertices covered by multiple detectors. 255 participants were used with each participant having one data source. Each participant moves according to a random mobility model. 15

[1]In the presented solution the data collected by a cell tower are averaged to represent the measurement over the entire area or tower coverage, in future work, we will collect data from each grid cell separately.

[2]In our future work, we will examine the impact of participants rejoining auctions and variable ROI thresholds for leaving the auction.

Figure 2: Experiment Schematic

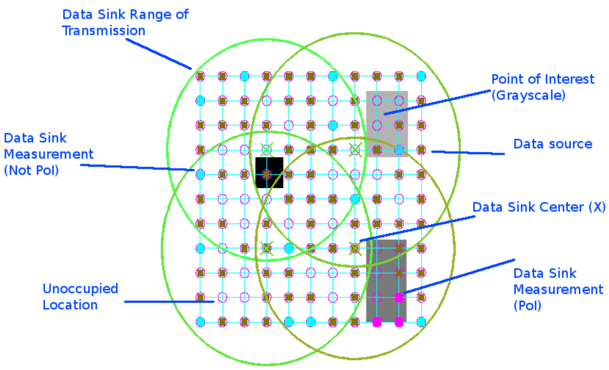

vertices lie in 'points of interest' where measurements are taken. There are 3 non-adjacent rectangular points of interest. The rest of the vertices register a reading of 0 whenever a data source attempts to sense that location. Participants do not try to abuse the system through collusion or in the case of P3AM via extremely high bidding. In our experiments we set $f_{coverage}$ to 0.1.

We chose several values for bounds on true valuation and β and used each possible parameter value in conjunction with each mechanism-specific parameter. This led to 405 configurations for P3AM, 81 for RADP-VPC, 135 for PI-GVA, and 27 for first price for a total of 648 configurations.

True valuations are in incentive units, not a particular currency. Each participant will increase its bid by 10% with probability $P_{inc} = 0.5$ if it wins an auction round. Similarly, each participant will decrease its bid by 20% with probability $P_{dec} = 1.0$ after losing an auction round.

Table 1: Mechanism-specific Parameters

Mechanism	Parameter	Values
P3AM	$P_{cheapest}$ (%)	0,25,50,75,100
RADP-VPC	α	0.1,0.5,0.9
PI-GVA	μ	20,50,100,500

We ran two sets of experiments (each set consisting of the 648 configurations described above). In Dataset I, β_{min} was set to 500000 which ensured that nobody stopped participating. In Dataset II, β_{min} was selected between 50 and 100. Table 1 shows the values we used for other parameters of each mechanism.

In Figure 3, we examine the evolution of measurements sent to a detector (regardless of if they are a PoI or not) by looking at the average number of measurements accumulated over time. While initially all four mechanisms perform the same, first price quickly diverges from the other three mechanisms. This demonstrates the myopic nature of first price auctions and shows that the addressing of social concerns in P3AM does not negatively affect the number of measurements taken when compared to participation-centric mechanisms.

Next, in Figure 4, we show the the relationship between $P_{cheapest}$ and the average number of participants for P3AM in Dataset II. When $P_{cheapest}$ is raised from 0% to 25% there is not much change observed in participation. As $P_{cheapest}$

Figure 3: Average Cumulative Measurements

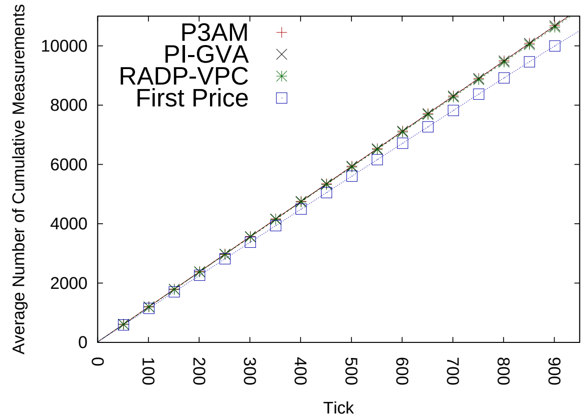

continues to increase, the number of experiments with 225 or more participants steadily declines. This indicates that there is a tradeoff between incentive paid out (budget) and participation. By utilizing the fairness aspect of P3AM to retain participants that are not optimally priced during a given round, overall participation can be maintained. This is important since participation is necessary to keep competition and to avoid explosion of prices or compromising privacy.

Figure 4: Average Participation by $P_{cheapest}$ (Dataset II)

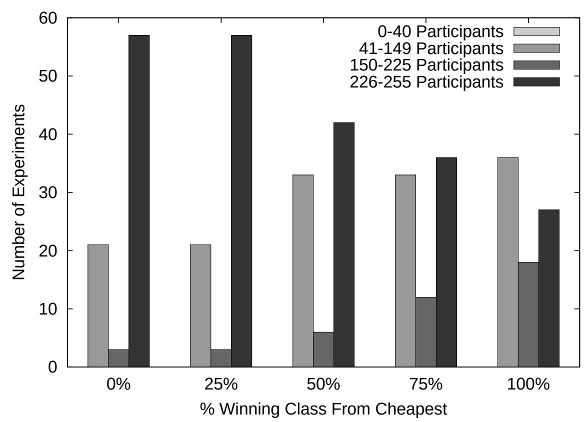

We examined how mechanism-specific parameters affected the average detector payout (θ) in Dataset I. The results are shown in Figures 6, 8, and 10. $P_{cheapest}$ has a visible effect on the range of θ with lower $P_{cheapest}$ giving a higher minimum and maximum θ. Neither α in RADP-VPC nor μ in PI-GVA significantly affect the values of θ. Recalling the analysis of $P_{cheapest}$ on active participation, we can see that indeed budget and participation are tradeoffs parameterized by $P_{cheapest}$.

To examine privacy, we looked at the average inter-win bid time in P3AM, RADP-VPC, and PI-GVA. This was de-

Figure 5: P3AM Average Intewin

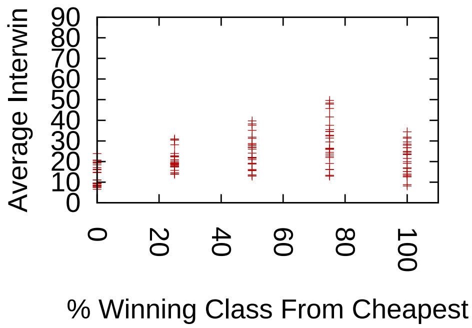

Figure 6: P3AM Detector Payout

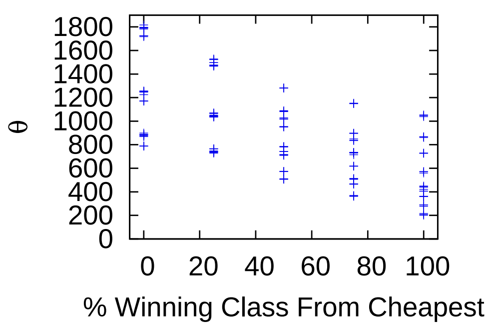

Figure 7: PI-GVA Average Intewin

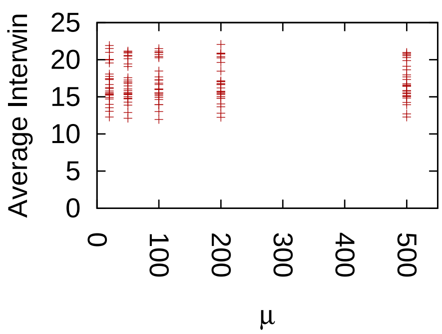

Figure 8: PI-GVA Detector Payout

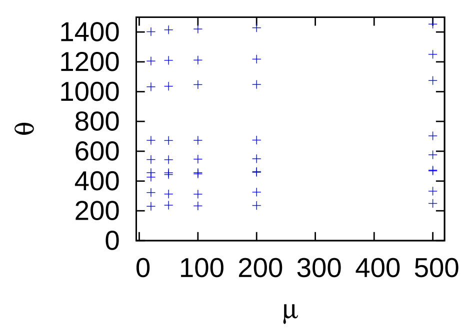

Figure 9: RADP-VPC Average Intewin

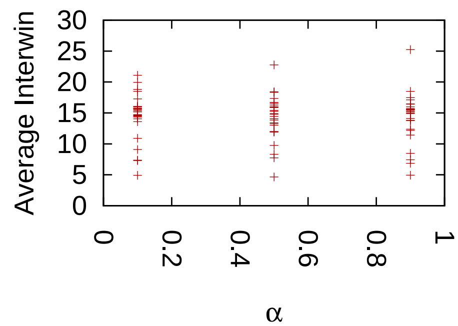

Figure 10: RADP-VPC Detector Payout

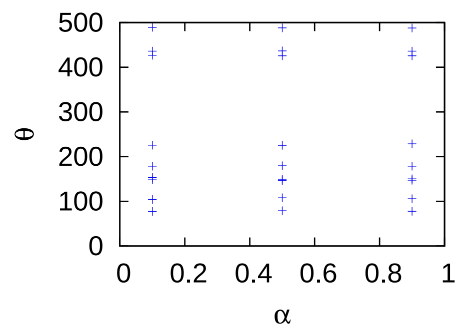

fined as the average number of ticks that elapsed between two wins for a given node. The averages across all experiments in Dataset II and all participants in each experiment for a given mechanism are shown in Figures 5, 7, and 9. From these results we can see that, as in the case of θ, only $P_{cheapest}$ has an effect. As $P_{cheapest}$ increases and P3AM becomes more similar to a first price auction, the maximum interwin bid time increases. This demonstrates the expected tradeoff between privacy and fairness. When fairness is employed, participants with very low tolerance to loss will be selected frequently despite their bid b. This results in very low interwin times. When $P_{cheapest}$ is very high, the mechanism becomes like a first price auction and interwin times are low since participants with low prices and low privacy concern are selected frequently. These results also illustrate that RADP-VPC and PI-GVA cannot be calibrated to meet different privacy requirements.

Since another concern was preservation of battery level we compared the average battery levels of winners for each mechanism in Dataset II. The comparison is shown in Figure 11 with each data point being the average across all experiments of the battery levels of all nodes winning an auction during that tick. While PI-GVA manages to perform similarly to P3AM initially, after about 300 ticks PI-GVA's

Figure 11: Average Contributor Battery Level

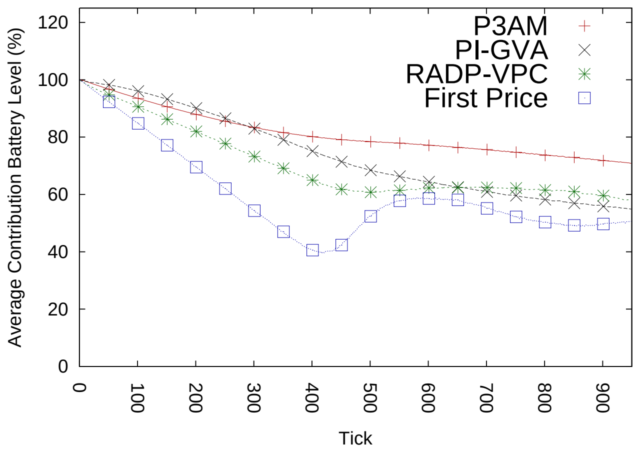

ability to preserve battery deteriorates. First price results in some nodes being exhausted very quickly which causes the curve to be non-monotonic. Only P3AM consistently provides high average battery levels which is accomplished through the $f_{privacy}$ and f_{power} terms in the mechanism.

Finally we compared the average price per measurement for all mechanisms in Dataset II, shown in Figure 12. While first price does not show an explosion of prices as described in [9], near the end of the simulation the price curve starts to increase superlinearly. In other experiments where a higher percentage of participants were selected in each round, we have observed similar behavior for P3AM, PI-GVA, and RADP-VPC but a large explosion of price in first price that happens around tick 600. These experiments have been omitted due to space concerns. We can see that RADP-VPC does not utilize budget efficiently and maintains participation at the expense of continually larger incentives. PI-GVA starts with initially high prices but by the end of simulation converges to approximately the same prices at P3AM. P3AM has an initially high price but quickly stabilizes and maintains low average prices with only a slight increase over time. This shows that P3AM is both efficient with budget and unlike first price is viable option for long-term auctions or deployments where overprovisioning is less pronounced.

As a result, we can see that P3AM's socially aware nature not only allows preservation of battery levels, privacy, and participation pool but also provides more efficient use of budget than PI-GVA, RADP-VPC or first price mechanisms.

6. CONCLUSION

In this paper, we studied the impact of social concerns in a participatory sensing application. We proposed a socially-aware market mechanism which considers the privacy of participants and the power level of their devices in its design. This was done to encourage participants to remain in the auction while making effective use of a potentially limited budget. Unlike the parameters in other participation-aware mechanisms, the $P_{cheapest}$ parameter of our proposed mechanism can alter the average participation while maintaining lower payouts of detectors per measurements compared to a known incentive compatible mechanism. From the results shown, our mechanism simultaneously demonstrates better performance in privacy and battery preservation.

Figure 12: Average Price per Measurement

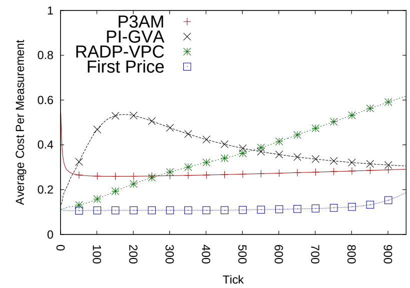

In our future work, we will work on the analysis of several parameters that affect the system performance. These include the selection of winners in areas of interest only, distribution of remaining battery levels among nodes and the evolution of incentive amount given by the data sink. We will also apply our model with real human mobility traces.

7. REFERENCES

[1] Mobile Millennium http://traffic.berkeley.edu.

[2] Nokia Simple Context https://alpha.nokoscope.com/eb2.

[3] MIT Media Lab: The Owl Project http://owlproject.media.mit.edu.

[4] Krause, A., Horvitz, E., Kansal, A., and Zhao, F. Toward community sensing, in *IPSN*, 2008.

[5] Hoh, B., Gruteser, M., Herring, R., Ban, J., Work, D., Herrera, J.C., Bayen, A. M., Annavaram, M., and JACOBSON, Q. Virtual trip lines for distributed privacy-preserving traffic monitoring, *MobiSys 2008*, Breckenridge, CO.

[6] Srivastava, M., Abdelzaher, T., and Szymanski, B. 2012. Human-centric Sensing. In *Transactions of Royal Society, Series A*, **370**:176-197.

[7] Reddy, S., Estrin, D., Hansen, M., and Srivastava, M. 2010. Examining Micro-Payments for Participatory Sensing Data Collections. In *International Conference on Ubiquitous Computing*, Denmark, Sept. 2010.

[8] Lee, J., and Hoh, B. Sell Your Experiences: A Market Mechanism based Incentive for Participatory Sensing. *IEEE PerCom*, Mannheim, Germany, March 2010, 60-68.

[9] Lee, J.-S., and Szymanski, B. 2009. A Participation Incentive Market Mechanism for Allocating Heterogeneous Network Services. In *IEEE Globecom*, Hawaii, Dec. 2009, 1-6.

[10] Danezis, G., Lewis, S., and Anderson, R. How Much is Location Privacy Worth?, in *WEIS '05:Proceedings of the Fourth Workshop on the Economics of Information Security*. 2005.