Design and Experimental Evaluation of Market Mechanisms for Participatory Sensing Environments

George Thanos[1], Costas Courcoubetis[1,2],
Evangelos Markakis[1], and George D. Stamoulis[1]

[1] Athens University of Economics and Business, Athens, Greece
[2] Singapore University of Technology and Design (SUTD)

Abstract. Participatory Sensing concerns the sharing of sensor information within user communities, forming a body of knowledge that can be beneficial to the community itself, either directly or through specialized applications. We introduce a framework for a marketplace where such applications can sell and buy information. In our approach, all involved entities are viewed as economic agents. Suppliers of sensor information may be reluctant to share their sensors due to costs in transmitting information (battery, bandwidth, etc). Potential customers may also be reluctant to participate in the market if the prices for the information are prohibitively high. Within this framework, we focus on the buyers' side and on mechanisms that take such incentive issues into account. Using various ideas from the cost-sharing literature, we propose three classes of mechanisms, satisfying different properties each. We evaluate them experimentally, comparing their performance according to metrics such as social efficiency, cost coverage and budget deficit, as well as metrics related to encouraging participation, since this can lead to the overall sustainability of the market in the long run.

1 Introduction

Personal mobile computing is undergoing a major revolution: smartphones have significantly changed the way humans interact, as powerful, programmable mobile devices outfitted with a range of advanced and low-costing sensing capabilities. The exponential growth in data capture and data sharing capabilities is giving rise to new applications and user habits. People have already started participating in sensing, instrumenting and analyzing aspects of their lives, eventually becoming producers of data, rather than just being consumers. These developments are creating a compelling need for new mechanisms to support such a participatory community-sensing environment, where multiple users, business and applications dynamically interact and share sensor information.

Current work on participatory sensing systems mainly focuses on solving the technical challenges of the physical environment, such as enabling low-power, low-bandwidth wireless sensor network communication, energy and resource management etc. Furthermore, some state-of-the-art systems have already been deployed like OpenSense [11] and CrowdPark [8] (for more on these, see Section 2). However, what has been undermined and not adequately addressed so far is the fundamental economic issue of why should users share or exchange such information that is costly to them and what is the necessary technology that can facilitate this aspect.

Our long-term goal is a sustainable market for sensor information in participatory sensing environments. To begin with, we need first to determine the main features of the market such as the nature of the goods that are being sold, the supply and the demand side and the decisions that a market mechanism needs to make. In particular, the model we propose consists of the following elements:

– **The sensor goods.** The basic good traded in this market is a sensor. More precisely, we consider the case where 'buying' a sensor consists of buying the right to have access to the sensor information for a given time slot. We abstract from the market the process of declaring sensor availability and accessing the sensor data. This is part of the implementation services of the market. At the abstract level of our model we assume that the market operates in discrete time (time slots of a given physical duration) and at any given slot, there is a set of sensors available in the market (supply and demand are renewed between successive slots). The information provided by a sensor can be thought of as a "digital good" that can be simultaneously bought by more than one buyers, as it can be duplicated at almost zero cost.
– **The supply part**, consisting the suppliers of sensor information, i.e., the individuals owning sensor-equipped devices. Each supplier can make his sensors available, independently of others. For this, he needs to specify the minimum price he is willing to charge for a particular sensor in order to supply its sensing information. In other scenarios (not addressed in this paper), he may ask for a single charge to make all sensors available instead of pricing individual sensors. We assume that the supplier can only sell the rights for accessing his sensors to some intermediary, in our case, the market operator. Hence, the supplier does not communicate directly with the buyers, and does not get compensated on a per usage basis. He only obtains a payment according to what he asked for, through the market operator.
– **The demand side**, which consists of the applications that are interested in acquiring sensor information. We think of each such application as a distinct customer in our market whose demand is expressed by specifying the set of sensors he is interested in. Our model of customer preferences allows for a variety of types. We assume that each buyer specifies to the market operator his maximum willingness to pay for the set of sensors he requests.
– **The market mechanism**. A mechanism, run by the market operator, defines the rules of the interaction between buyers and suppliers. It is a function that takes as input the supply and the prices from the sellers, along with the demand and the willingness to pay from the buyers, and as a result, it derives the allocations of the goods and the payments from the agents.

Contribution: We view as our main contribution the design of the market model itself. The framework for the marketplace, is one of the first attempts to define such a large-scale sensor market, where applications of various forms can sell and buy information. In our approach, all involved entities are viewed as economic agents and one of our goals is to define appropriate mechanisms for running the market.

Within this framework, we focus mostly on the buyers' side and we use various ideas from the cost-sharing literature, to propose three classes of mechanisms, satisfying different properties. The first is inspired by the work of Moulin and Shenker [9, 10] on a simpler setting than ours. The second is based on an altruistic paradigm, where agents can contribute towards covering the cost-shares of other agents that cannot afford to pay the entire amount. The incentive behind such a move is that it increases the demand in the system and hence reduces the average cost of sensor access. Such altruistic behavior even by a small set of agents can be important for enhancing participation and increasing the social welfare of the market. In our mechanisms of this class, we impose a related rule, rather than relying on the inherent altruism of agents, thus ending up with several nice performance properties. Finally, the third class of mechanisms is essentially the Groves mechanisms, although we focus particularly on the VCG mechanism [2]. We evaluate these mechanisms experimentally, according to various metrics such as social efficiency, cost coverage and budget deficit, as well as metrics related to encouraging participation of users, influencing the overall economic sustainability of the market in the long run.

We believe that the aforementioned environment crafts the ideal conditions for the creation of an open sustainable market and the underlying economic mechanisms that will promote efficiency in sensor information exchange. Such a market, tailored to the specificities of participatory applications, which to our knowledge is missing today, will leverage this social phenomenon and the enabling technological advances into a new paradigm on how people interact with each other directly and indirectly in an economically efficient way.

2 Related Work

Participatory sensing is a very promising direction towards replacing traditional sensor networks. There are already existing deployments that support a variety of applications like environmental monitoring (OpenSense [11]), transportation (CrowdPark [8]), fitness (BikeTastic [12]), urban sensing (PulsodelaCiudad [14]), and others. However, most of these platforms have suffered from insufficient participation because users that voluntarily submit their sensing data found no interest in remaining active in the system without being rewarded. This undesirable fact has already been observed in [3, 7, 6] and motivates the deployment of incentive schemes so as to increase user participation.

Towards this goal, most approaches in the literature focus on incentive issues in the supplier's side. Namely, suppliers may drop out unless there is a positive Return on Investment, which depends on the total cost for collecting data (battery consumption, device resources, privacy, etc). In [7], a reverse auction is proposed to address this issue. Another reverse auction is also proposed in [6]. The work of [3] on the other hand is limited to using a fixed price approach. An issue that is not covered by these works is the modeling of the demand side of the market, which is what we mainly address in this paper.

Finally, the works from the economics literature that are most relevant to ours are the cost-sharing mechanisms of Moulin and Shenker [9, 10]. These mechanisms work for a simpler (binary) setting where each user is either granted the same identical service with all other users or he is declined. We also consider the Marginal Cost Pricing mechanism, see [9], which is the adaptation of the VCG mechanism [15, 2, 5] into the cost-sharing setting.

3 Definitions and basic concepts

Suppose that there is a set $N = \{1, ..., n\}$ of agents who are interested in receiving some service from a market operator. We define mechanisms for general settings below and we will consider instantiations to our specific settings in the next sections.

A mechanism design instance, consists of a tuple $(N, O, \Theta, \mathbf{u})$. The set O consists of the possible outcomes of the mechanism, which is in the form $O = X \times \mathbb{R}^n$ for some set X (X denotes the space of all possible allocations of services/goods to the agents and lies in some n-dimensional space). Hence, an outcome of the mechanism consists of an allocation decision $x \in X$ and a vector of side payments $\mathbf{p} = (p_1, ..., p_n)$. The set $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is the set of agent types, and $\mathbf{u} = (u_1, ..., u_n)$ describes the utility functions. These functions determine the *willingness to pay* of each agent. For cost-sharing settings, as is ours, we also assume that there is a cost function, determining the cost $C(x)$ for realizing an allocation $x \in X$ (e.g., the cost for providing a service).

A *mechanism* M, is then a function $M : \Theta \to O$ mapping each vector of declared types by the agents to an outcome. I.e., after collecting all the offers made by the agents, it decides a) upon an allocation $x \in X$, and b) how much to charge each participant. For an outcome of the mechanism in the form $o = (x, \mathbf{p})$, with $x \in X$, $p \in \mathbb{R}^n_+$, the final utility of an agent i, which is also referred to as the consumer's *net benefit* or *consumer surplus*, is the derived utility minus the payment, i.e., it is $u_i(x) - p_i$.

The main focus of our experimental evaluations will be on the following two important criteria:

- **Budget balance.** A mechanism is *budget-balanced* if for every instance, the payments assigned to the customers cover exactly the cost of the provider.
- **Social welfare maximization.** The *social welfare*, or *social surplus*, is defined as the sum of all involved agents' net benefits. The payments made by the consumers cancel out with what the provider receives, hence, for an allocation $x \in X$, the social welfare is $\sum_i u_i(x) - C(x)$.

Another property that is often pursued in mechanism design is *strategyproofness*, meaning that truthtelling is a *dominant strategy*. We will also discuss a stronger form of incentive compatibility, namely *group strategyproofness*, where no coalition of agents has an incentive to jointly misreport their true willingness to pay. Although we do not insist on having strategyproof mechanisms (given that in practice this may often be a too stringent requirement), some of the mechanisms we study are strategyproof or group-strategyproof.

3.1 Our model for the sensor market

In the marketplace we want to create, mechanisms would be run by some automated market operator system that would act as an intermediary between suppliers and buyers. Hence, our model consists of the following features:

- A set $I = \{1, ..., k\}$, representing the different sensor *basic types*, e.g., GPS, accelerometer, temperature, CO_2, etc. The type of a sensor describes the measurement information that the sensor provides.

– A set $M = \{1, ..., m\}$ of suppliers, the owners of sensor data (via their mobile or any other device). The suppliers will not provide their data for free since this entails a cost (battery, bandwidth, etc). We assume that each of them specifies a price per sensor type that he needs to be paid for in order to provide access to the value of a specific sensor type. A value provided by one supplier can be used by many buyers. Suppliers do not all necessarily have the same set of sensor types available.
– A set $N = \{1, ..., n\}$ of potential buyers. These are agents who have a demand for some sensor data (we use interchangeably the terms agent and buyer to refer to any $i \in N$). Different types of demand (e.g., elastic vs inelastic, or single tuple vs multiple tuples) are examined in Section 4.

Within this context, the market mechanisms we evaluate fall into three categories:

1. Mechanisms that achieve budget balance. For this we will adapt ideas from the work of Moulin and Shenker [9, 10], into our setting. The trade-off with these mechanisms is that they tend to produce suboptimal social welfare.
2. Cooperative mechanisms that maintain budget balance and aim towards achieving higher social welfare, by having some "richer" agents subsidize other agents who cannot afford their cost-share. We will again utilize the Moulin-Shenker mechanisms but in combination with an 'altruistic' framework described in Sections 4.1 and A. Even though this presents some potential dangers (e.g. emergence of tragedy of the commons), we believe that it is important to give priority to having an initially sustainable market that encourages participation.
3. Mechanisms that achieve optimal welfare. For this we apply the well-known Marginal Cost (MC) mechanism [10]. It is known that such mechanisms cannot balance the budget and we will therefore evaluate MC in terms of its budget deficit. We will also use heuristics to approximate the optimal welfare, in settings where the problem is computationally intractable.

We stress that our mechanisms are centered around the buyers and not the suppliers. For the suppliers, we assume that the operator has already obtained their price per sensor requests. For dealing with the suppliers' side of the market, see [6].

4 Our demand domains

We present two simple orthogonal scenarios regarding the demand of the customers. In both scenarios, the market operates in discrete time. The market operator considers each time slot separately, he looks at the currently available data, the prices set by the providers, the current demand for data along with their monetary offers and decides which buyers get served and at what prices. We leave for future work the investigation of dynamic settings (where bidders could also specify time duration in their demand).

We note here that we assume that all instances are feasible, i.e., there are enough sensors available in the market to satisfy all buyers, in the case that they can all afford to pay. If not, one could easily run a feasibility check and determine a maximally feasible subset of buyers.

4.1 Scenario 1: Inelastic demand

We consider a simple scenario of inelastic demand, in which each buyer $j \in N$ is interested in a subset $S(j) \subset I$ of sensor types, and requests access to a single tuple of sensors from $S(j)$. As an example, he may request a tuple of the form (speed, accelerometer). The operator can group in advance the tuples according to location or other user constraints so that we can consider instances where we have already fixed some particular location/neighbourhood. Each buyer $j \in N$ derives a utility of u_j by receiving this tuple, hence the utility function is specified by a single number here. The demand is inelastic in the sense that the buyer is not deriving any utility if he receives only a strict subset of sensors from $S(j)$. We call such buyers *single-minded*, in analogy to single-minded bidders in combinatorial auctions [1].

The cost function $C(R)$ for serving a set of customers $R \subseteq N$ can be easily computed for any R. To see this, given the suppliers' prices, for any sensor type $i \in I$, let $c(i)$ be the cost of the cheapest sensor of type i (i.e., the cheapest price specified by some supplier for type i). Since each customer is interested in receiving a single tuple, we can use just one actual sensor for each type requested, to satisfy all customers. And clearly the market operator should use the cheapest possible. Hence, for $R \subseteq N$, the cost $C(R)$ is the sum of the cheapest cost of all sensor types required by R:

$$C(R) = \sum_{i \in S(R)} c(i), \text{ where } S(R) = \bigcup_{j \in R} S(j)\,. \tag{1}$$

Proposition 1. *The cost function $C(R)$ is monotone and submodular, i.e., we have:*

- $C(\emptyset) = 0; R \subseteq T \Rightarrow C(R) \leq C(T),$
- $C(T \cup \{j\}) - C(T) \leq C(R \cup \{j\}) - C(R),\ \forall R \subseteq T \subseteq N$ *and* $j \notin T$.

The proof is quite simple and we omit it.

Budget-balanced mechanisms The first mechanism we consider is derived directly from the pioneering work of Moulin and Shenker [9, 10]. Their work concerns a much simpler setting than ours, where there is a single provider, offering the same identical service to everyone, and each agent will be either granted or declined the service. Even in this simple binary setup, designing mechanisms with good incentive properties is a challenging task. We can easily adapt the approach of Moulin and Shenker to achieve a group-strategyproof budget-balanced mechanism for Scenario 1. To do this, we need to obtain first an underlying *cost-sharing* method. A cost-sharing method is a function $\xi(\cdot,\cdot)$ such that $\xi(j, R)$ determines the cost-share of agent j, $j \in R$, when R is the set to be served by the mechanism. We demand that a cost-sharing method satisfies $\sum_{j \in R} \xi(j, R) = C(R)$ for all $R \subseteq N$, i.e., the sum of the payments balance the cost. The following is an important and desirable property for cost-sharing methods.

Definition 1. *A cost-sharing method is cross-monotonic if*

$$\xi(j, R) \geq \xi(j, T) \text{ for } R \subseteq T \text{ and } j \in R\,. \tag{2}$$

The above property simply says that the cost-share of an agent should not become higher when more people receive service.

In our experiments, we have decided to work with the *egalitarian* cost-sharing method, since this may have more appeal in practice due to its simplicity and fairness properties. In particular, to define the share $\xi(j, R)$ for a given set R to be served, we split the cost of each used sensor equally among the people who want it. For each sensor type $i \in \bigcup_j S(j)$, recall that $c(i)$ is the cost of the cheapest available such sensor in the market supply. Let $y(i)$ be the number of buyers who have i in their demand set. Egalitarian cost sharing means that each customer j contributes a share $c(i)/y(i)$ towards the total cost of the sensor. Hence for a buyer j, with demand set $S(j)$, his total cost-share under this cost-sharing method is:

$$\xi(j, R) = \sum_{i \in S(j)} \frac{c(i)}{y(i)}. \tag{3}$$

Proposition 2. *The cost-sharing method described by (3) is cross-monotonic.*

Given any cost-sharing method ξ, one can define now parametrically the mechanism below for determining who receives service along with the cost-shares. In the description below, we let $\mathbf{b} = (b_1, ..., b_n)$ be the agents' bids for their demand sets.
Mechanism $M(\xi)$ (given a cost-sharing method $\xi(\cdot, \cdot)$):

- Start by trying to serve all agents, with cost-share $\xi(j, N)$. Remove any agent who cannot cover his share, i.e., anyone for which $b_j < \xi(j, N)$. If no-one is removed in this step, stop here, otherwise let R^1 be the set of remaining agents.
- Check if we can serve R^1 with a cost-share of $\xi(j, R^1)$ for every $j \in R^1$. Again remove those who cannot afford this price.
- Continue like this and in every round obtain the set $R^{t+1} = \{j \in R^t : b_j \geq \xi(j, R^t)\}$.
- Stop either when we reach the empty set, or when we reach a set in which all agents can afford to pay their cost-share.

The mechanism $M(\xi)$ has some interesting theoretical properties. The following theorem is a straightforward extension of the results from [9, 10] to our setting.

Theorem 1. *For any cross-monotonic cost-sharing method ξ, the Mechanism $M(\xi)$ for single-minded bidders is budget-balanced and group-strategyproof.*

Further important properties of this family of mechanisms are obtained in [9, 10].

Altruistic budget-balanced mechanisms towards higher social efficiency In the context of selling sensor information, one drawback with the Moulin-Shenker mechanism is that it rejects people, at the very first moment where the mechanism realizes that they cannot pay their share. This may affect negatively the sustainability of the market and lead to suboptimal generated welfare. Since it is important for the market operator to maintain an up-and-running market, one idea that can be particularly useful is to let "richer" buyers subsidize poorer ones using their left-over money. One can start from the most naive idea where in the first step of the mechanism $M(\xi)$, we look at each person who cannot afford to pay his share $\xi(j, N)$ and check if the buyers with high enough value can cover the missing amount. We can then consider more and more complex variants of altruism. The incentive for a rich customer to subsidize poorer ones is that it increases the demand in the system and hence reduces the average cost of sensor access. At an equilibrium, we expect users to subsidize others when the marginal long term benefit from doing so equals their average reduction of surplus. This justifies the use of altruistic mechanisms.

Clearly, we need to be aware that we lose now the property of strategyproofness. This may not be so important for practical considerations (since in practice strategyproofness is considered to be a strict requirement). But there are more issues that arise from the fact that customers have different types of demand. In particular, consider a customer j who can afford to pay his initial share $\xi(j, N)$, as determined by $M(\xi)$. Consider also a customer k who cannot afford his share. If $S(j) \cap S(k) = \emptyset$, then there is no reason for j to subsidize k from his left-over money. On the other hand for customers with $S(j) \cap S(k) \neq \emptyset$, j can consider such a subsidy, as this will help him obtain his tuple at an earlier round of the mechanism, and thus at a cheaper cost. Hence, there are certainly instances where some agents have incentives to help other people. The decision problem now that arises is to pick appropriately which customers get subsidized since there can be multiple agents who cannot pay their share.

Based on these ideas, we have come up with a subsidy-based variant of $M(\xi)$, which we denote by ALT. ALT starts just like $M(\xi)$ and initially determines the shares $\xi(j, N)$, attempting to serve all agents. In each round, the mechanism tries to subsidize the "poor" agents in increasing order of the least missing amount. For each of those, a list of potential subsidizers is created. An agent j is included in the list for subsidizing agent k, only if $S(j) \cap S(k) \neq \emptyset$. Then, for each of the "poor" agents we divide the missing amount among the potential subsidizers, either in an egalitarian manner or proportionally to the number of common sensors with the poor agent (the former is used for our experiments). If the subsidy share cannot be paid by all subsidizers, we take the most out of the ones who cannot pay their full share, and divide among the rest the remaining amount. Then we continue to the next rounds, in the same manner as $M(\xi)$, Finally, we also return the subsidy-shares back to the subsidizers if the agents that they tried to subsidize are kicked out in subsequent rounds.

Mechanisms for maximizing social welfare - The Marginal Cost Pricing Mechanism The mechanisms described so far are budget-balanced. It is known however that such mechanisms do not maximize the economic efficiency of the system, i.e., they may result in suboptimal social welfare, due to the impossibility results of [4, 13].

To define social welfare formally, suppose that a set R is chosen as the set to be served. If $\mathbf{u}$ is the utility vector of the agents, then the generated social welfare is:

$$SW(R, \mathbf{u}) = u_R - C(R) \tag{4}$$

where $u_R = \sum_{i \in R} u_i$. The optimal social welfare that can be achieved is

$$SW^*(N, \mathbf{u}) = \max_{R \subseteq N} [u_R - C(R)] \,, \tag{5}$$

We also denote by $R^*(N, \mathbf{u})$ the set of customers that achieves the maximum welfare (in case of multiple optimal sets, we can pick one according to some tie-breaking rule).

The standard way to have strategyproof mechanisms that achieve optimal welfare is by using the family of Groves mechanisms[15, 2, 5] in the cost-sharing setting. In particular, we will utilize the pivotal mechanism [2], and in this context, following [10], we will refer to it as the Marginal Cost (MC) mechanism. The MC mechanism takes as input by each agent his demanded set $S(j)$ and his declared utility b_j, and first solves (5) to compute a set to be served with optimal welfare. Then, for the customers that receive service in the optimal solution, i.e., for $j \in R^*(N, \mathbf{b})$, it charges as follows:

$$p_j = b_j - (SW^*(N, \mathbf{b}) - SW^*(N \setminus \{j\}, \mathbf{b})), \tag{6}$$

It follows immediately by the results on Groves mechanisms, that MC is a strategyproof mechanism. As already pointed out, by the impossibility results mentioned earlier, MC is not budget-balanced, and the amount collected is at most equal to the total cost.

The difficulty in applying MC in practice is the computational complexity of solving (5). We are not aware of an efficient algorithm for solving large instances of (5). For small instances (with 20-25 buyers), we have implemented a brute force algorithm to run the MC mechanism. For larger instances however, we need to develop some heuristics or approximation algorithms. A drawback with using a heuristic to approximately solve (5), and then run the payment scheme of (6), is that strategyproofness does not hold any more. Nevertheless, it is still interesting to have mechanisms that produce approximately optimal welfare, since welfare maximization is an important problem on its own. Furthermore, when we run a heuristic, the payment scheme still makes it hard for buyers to manipulate the mechanism. Thus, we expect that buyers will most likely behave truthfully even under these heuristic variants of MC.

For large instances, the first natural heuristic approach that comes to mind is to start from an empty solution and gradually produce a feasible one as follows:

H_{add} **(Greedy Heuristic starting from the empty set)**
Initialization: Start with $R = \emptyset$, $SW(R, \mathbf{u}) = 0$.
In every round: See if there exists some buyer j such that adding j to R (weakly) increases the current social welfare. If yes $R := R \cup \{j\}$. If no such buyer exists, then stop and return R, with a social welfare of $SW(R, \mathbf{u})$.

The instances in which H_{add} may fail to produce any positive welfare is when there is no person j who can cover the cost of his set $S(j)$ on his own. In that case H_{add} will not add any person to the solution in the beginning, and will result in zero welfare. To remedy this, we also propose a second heuristic, where we start from the whole set N with an initial welfare of $SW(N, \mathbf{u})$. In each round we can then check if there is any agent whose removal increases the social welfare. We can repeat this until no further removal is possible. We refer to this heuristic as H_{sub}. Finally, we can also pick the best out of the two heuristics. Hence, our third and final heuristic algorithm, which we refer to as H_{max}, is to run both H_{add} and H_{sub} and then pick the solution with the highest welfare: $H_{max} := \max\{H_{add}, H_{sub}\}$.

4.2 Scenario 2: Elastic demand for multiple tuples of the same basic types

This scenario is orthogonal to Scenario 1 in the following sense: all bidders now have the same type of demand, i.e., they all ask for the same type of tuple, and they are still inelastic as in Scenario 1, regarding the type of tuple. What differentiates the buyers in this setting, is that each buyer j also specifies a maximum number, m_j, of tuples that he is interested in acquiring. The demand m_j is elastic in the sense that buyer j does not mind receiving a number of tuples less than m_j. Along with the parameter m_j, each customer also specifies his willingness to pay for each tuple. This is encoded by a vector of marginal utilities $\mathbf{u}_j = (u_j(1), u_j(2), ..., u_j(m_j))$, where $u_j(\ell) \in \mathbb{R}_+$ denotes the marginal utility of receiving the ℓ-th tuple after already having received $\ell - 1$ tuples, with $\ell \leq m_j$. Thus, the maximum amount each buyer is willing to spend is $\sum_{\ell \leq m_j} u_j(\ell)$. If an agent declares a (not necessarily truthful) willingness to pay $\mathbf{b}_j$, this means that if a mechanism gives him r tuples, with $0 \leq r \leq m_j$, his willingness to pay for these tuples is perceived by the mechanism to be $\sum_{\ell \leq r} b_j(\ell)$.

In a similar way as in Scenario 1, we have implemented three classes of mechanisms as well. We describe these in Appendix A, highlighting also some differences with Scenario 1.

5 Experiments

5.1 Evaluation Criteria

One of the main goals of our work is to provide experimental evidence on the performance of the suggested mechanisms. The main issues that investigated are:

1. Can we argue about the success of the system in satisfying the customers, which would consequently lead to its sustainability in the long run? We deal with this issue by considering the satisfaction from the market operation and the probability of having a successful market transaction, i.e., the percentage of customers whose request was satisfied. Other indicative metrics are the number of sensors activated and the total utility generated.
2. How far from budget-balanced can the MC mechanism be in these settings? We know that MC may often run a budget deficit and we therefore want to evaluate the percentage of cost that is covered by MC.
3. How far from the optimal welfare is the welfare produced by the mechanisms $M(\xi)$, $M_2(\xi)$ and ALT? This question is orthogonal to the previous one.
4. Particularly for Scenario 1, we also evaluate the proposed heuristics.

5.2 Setup - Data generation

All the implementations were carried out in Matlab (data available on request). To test the mechanisms, we produced numerous instances using various distributions on certain parameters. The number of buyers in our simulations ranged from 5 to 300, apart from the MC mechanism in Scenario 1, which we had to run for only up to 20-25 buyers, since it involves a computationally intensive algorithm. A typical range that we used for the number of basic sensor types, i.e., for $|I|$, was the set $\{5, ..., 10\}$, reflecting the typical number of sensors available in current devices. For Scenario 1, we used a Gaussian distribution to determine the tuple $S(j)$, so that most buyers demanded a tuple with neither too many nor too few sensors. After selecting the sets $S(j)$ for each

j, we select the associated utility value u_j. In particular, the utility of each buyer j with $|S(j)| = k$ was randomly selected in the interval $[0, k \cdot B_{max}]$, where B_{max} is a parameter of the system. We generated 3 families of instances regarding the utilities. In the first one, utilities were uniformly distributed in the above intervals for every value of k. In the second, and third, we had a biased separation between "poor" and "rich" buyers, where for the poor (resp. rich) ones, the utility was drawn uniformly from the first half (resp. second half) of the interval. The second family contained $p = 70\%$-80% poor buyers and the rest were rich, and we had the exact opposite separation for the third family. Note also that the first family corresponds to an even mix of poor and rich buyers with $p = 50\%$.

The prices of the suppliers for each of their sensors are produced from the uniform distribution in $[0, R_{max}]$, where R_{max} is the max. price a supplier could ask for each sensor. In order to relate the prices of the suppliers to the buyers' willingness to pay, and their demand, we employed a parameter α, which we call the *economy factor*. This is used in determining the max. price asked by the suppliers by means of the following equation:

$$|I|\frac{R_{max}}{2} = \alpha \cdot (|N| \cdot B_{max} \cdot \frac{p + 3r}{4}), \tag{7}$$

where, $p, r = 1 - p$ are the percentages of poor and rich buyers respectively. Note that the lefthand quantity is the average total amount of money requested by the suppliers, while the term multiplying α in the righthand side is the average total willingness to pay for the buyers. Thus, a large (resp. a small) value of α implies that sensors are expensive (resp. affordable).

The experiments' setup for Scenario 2 was similar. In this case, for each buyer j, we determine values for the number of tuples wanted, m_j, and the utility per tuple, u_j. We note that we focused on the special case of Scenario 2, where all marginal utilities are equal. Again, we used uniform distributions as in Scenario 1, and considered an analogous formula to (7) to determine R_{max}.

5.3 Results

1. Market satisfaction ratio. We define this as the percentage of customers that are offered service. The main conclusion is that ALT and MC achieve quite high satisfaction ratios. This is not so much the case for the Moulin-Shenker based mechanisms. As we see in Table 1, ALT and MC have a consistently better performance than $M(\xi)$ and $M_2(\xi)$. ALT clearly serves more poor users than $M(\xi)$, thus increasing the total number of successful transactions; MC, being a mechanism that maximizes welfare, also results in a high number of transactions. Thus, despite their nice theoretical properties, the Moulin-Shenker based mechanisms may fail to ultimately promote participation.

2. Budget deficit of the MC mechanism. In Figure 1, we depict the percentage of the cost that MC manages to cover for various simulation runs and values for the economy factor α. By an average value for α, we mean values for which there is a satisfactory number of transactions in the market. A high value signifies that sensor prices tend to be high and not too many transactions take place. Our results reveal that the MC mechanism covers on the average a small percentage of the actual cost, and thus it is

Distribution	Uniform		More Poor		More Rich	
Scenarios	Sc. 1	Sc. 2	Sc. 1	Sc. 2	Sc. 1	Sc. 2
$M(\xi)$, $M_2(\xi)$	55	79	58	67	71	90
ALT	93	91	94	89	94	96
MC	92	91	93	90	92	94

Table 1. Satisfaction ratio (percentage), for the 3 mechanisms, for various buyers distributions.

not appropriate to use it. In order to have a better variant of the MC mechanism, one would need to impose some additional flat fee to the buyers so as to be able to cover on the average the deficit generated by the standard MC. We leave the exploration of such ideas for future work.

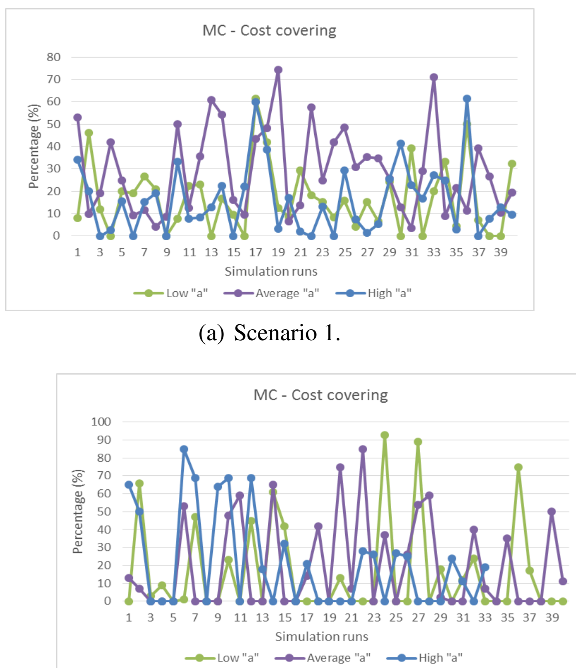

(a) Scenario 1.

(b) Scenario 2.

Fig. 1. Percentage of cost covered by the MC mechanism.

3. Welfare performance of $M(\xi)$, $M_2(\xi)$, and ALT. We compared the social welfare produced by $M(\xi)$ and ALT against the optimal welfare, which we computed with a brute-force algorithm. In Figure 2, we see that the altruistic mechanism significantly outperforms the mechanism $M(\xi)$. In fact, it is important to note that the altruistic mechanism attains in many cases the optimal social welfare. The same conclusions hold also for Scenario 2, and hence we omit the corresponding graphs.

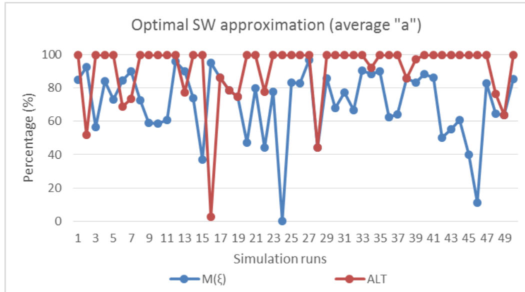

(a) Performance for average values of α.

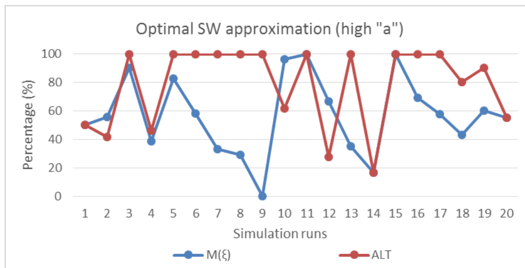

(b) Performance for high values of α.

Fig. 2. Approximation of optimal welfare by $M(\xi)$ and by ALT in Scenario 1.

4. Performance of heuristics in Scenario 1. In Figure 3 and Figure 4, which can be found in Appendix B, we see that the heuristic H_{sub}, and consequently the enhanced heuristic H_{max}, performs much better than H_{add} in terms of the welfare attained. This is an interesting observation, since typically the design of greedy algorithms starts from the empty solution and gradually builds a feasible solution, as H_{add} does. For our welfare maximization problem however, the approach of starting from the full set of buyers as an initial solution and gradually removing buyers is a more appropriate algorithm.

5. Overall Conclusions. To summarize, the Moulin-Shenker based mechanisms did not perform that well in terms of satisfaction ratio, while also attaining suboptimal social welfare. MC, which is optimal in terms of welfare, has poor cost-covering properties, which can be improved with the introduction of an additional fee. ALT seems to be a mechanism that strikes a balance: it has good satisfaction ratio, very good approximation to the optimal welfare, and it is budget-balanced. The downside of ALT is that the amount of subsidization required by a rich agent might be too high in some cases, at least in the early stages of such a market before participation raises. We believe that milder forms of subsidization (e.g., thresholds on the maximum imposed subsidy, or using only a small percentage of agents as potential subsidizers, even on a voluntary basis) can still be promising and are worth further investigation.

6 Conclusions and Future Work

In this paper, we develop a framework for a market of information in participatory sensing environments, as well as appropriate trading mechanisms for sharing the cost of information among the buyers. We view our work as a starting point towards creating actual marketplaces in such environments.The technical requirements for implementing the proposed market in a real system have been defined and its implementation, in the context of a bigger research initiative within our team, is on-going. There are plenty of avenues for future work. One challenging direction is to investigate richer combinatorial demand domains. E.g., how do we design mechanisms when customers have more complex demands such as multiple tuples from different subsets of basic sensor types. Another is to have a more formal study of market sustainability. An initial approach to this would be to consider a model where an application not served returns to the system only with a certain probability. If we assume a constant rate of new apps coming to the system, we could then compute the total rate of applications present and call a market sustainable if this rate does not decrease. We leave for future work a formal study of such issues.

References

1. L. Blumrosen and N. Nisan. Combinatorial auctions. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 17, pages 267–299. 2007.
2. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
3. G. Danezis, S. Lewis, and R. J. Anderson. How much is location privacy worth? In *WEIS*, 2005.
4. J. Green, E. Kohlberg, and J. J. Laffont. Partial equilibrium approach to the free rider problem. *Journal of Public Economics*, 6:375–394, 1976.
5. T. Groves. Incentives in teams. *Econometrica*, 41(3):617–631, Jul 1973.
6. I. Koutsopoulos. Optimal incentive-driven design of participatory sensing systems. In *INFOCOM*, pages 1402–1410, 2013.
7. J.-S. Lee and B. Hoh. Dynamic pricing incentive for participatory sensing. *Pervasive and Mobile Computing*, 6(6):693–708, 2010.
8. S. Mathur, T. Jin, N. Kasturirangan, J. Chandrashekharan, W. Xue, M. Gruteser, and W. Trappe. Parknet: Drive-by sensing of road-side parking statistics. In *ACM MobiSys*, pages 123–136, 2010.
9. H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Soc. Choice Welfare*, 16:279–320, 1999.
10. H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance vs efficiency. *Econ. Theory*, 18:511–533, 2001.
11. OpenSense Project. `http://www.nano-tera.ch/projects/401.php`.
12. S. Reddy, K. Shilton, G. Denisov, C. Cenizal, D. Estrin, and M. B. Srivastava. Biketastic: sensing and mapping for better biking. In *ACM CHI*, pages 1817–1820, 2010.
13. K. Roberts. The characterization of implementable choice rules. In J. J. Laffont, editor, *Aggregation and Revelation of Preferences*. Amsterdam: North Holland, 1979.
14. SmartSantander. `http://www.smartsantander.eu/index.php/blog/item/181-participatory-sensing-application`.
15. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, Mar 1961.

A Description of mechanisms for Scenario 2

Budget-balanced mechanisms The application of Moulin-Shenker mechanisms is not any more straightforward in the case of elastic bidders. In Scenario 1, each set of customers has a unique cost for satisfying its demand. In the case of elastic customers, each customer j corresponds to a set of potential service levels (ranging from 0 to m_j tuples). Then, for a potential set of customers, each possible allocation corresponds to an element of the cartesian product of the above sets. Hence, we cannot just run an analog of the mechanism $M(\xi)$ from Section 4.2 since at every step we would also need to decide on the service level before determining the cost-shares. One could consider all combinations of service levels to customers, and run $M(\xi)$ for each such combination (and then choose the one that is more efficient). But this has prohibitively high complexity to be run in practice.

Instead, we will still utilize the Moulin-Shenker approach but in a different manner. We introduce first some notation. Let $C(i)$ denote the total cost of the i-th cheapest tuple. Recall that the sensors used in a tuple do not need to come from the same provider, e.g., we can simply combine the cheapest sensor from every type requested to form the cheapest tuple, and continue in this manner for the rest of the tuples. If $t_{max} = \max_j m_j$, we can compute in this manner the first t_{max} cheapest tuples, which is all we need[3]. The main idea now in the mechanism below is that we try to sell each tuple separately using the Moulin-Shenker approach. In the description that follows, we let $(\mathbf{b}_j, m_j)$ be the declared request of agent j.

Mechanism $M_2(\xi)$ (given a cost-sharing method $\xi(\cdot,\cdot)$):

- Run the mechanism $M(\xi)$ from Section 4.1 to determine who receives the cheapest tuple. We then remove all customers who were not selected to be served (those customers will not be able to afford the next tuples anyway since they are more expensive).
- For the surviving customers who spent less than their bid for the first tuple, $b_j(1)$, we transfer any left over money to the next round (i.e., in round 2, their willingness to pay becomes $b_j(2) + \delta$, where δ is what they saved from the first round).
- Continue with the next tuple in the same manner, and at the beginning of round t, let R^t be the set of customers who i) have a request $m_j \geq t$, and ii) were not removed in the previous rounds.
- Run $M(\xi)$ to determine which customers from R^t will be allocated the t-th cheapest tuple. Again transfer any left over money to the next marginal bid.
- Stop either when $R^t = \emptyset$, or when $t > t_{max}$.

In our experiments, we use as a cost-sharing method the egalitarian one as before, where at round 1, we split the cost $C(1)$ of the cheapest tuple equally among all customers chosen to receive it by $M(\xi)$. We then split $C(2)$ among the remaining customers who survive round 2, and so on and so forth.

The mechanism $M_2(\xi)$ is budget-balanced by its construction, but unlike $M(\xi)$ in Section 4.1, it is not group-strategyproof.

[3] Recall that as we discussed in the beginning of Section 4, we assume that there are at least t_{max} available tuples.

Budget-balanced mechanisms under the altruistic framework As in Section 4.1, we have implemented an altruistic variant of the mechanism $M_2(\xi)$. In this scenario, the algorithm tries to subsidize first those agents who cannot even afford to buy the first (cheapest) tuple. Again we consider the agents in increasing order of their missing amount. The algorithm is now simpler than in Scenario 1, because all buyers are interested in the same type of tuple, hence all rich enough buyers can be in the list of potential subsidizers. In the same fashion, we check in each round if we can subsidize remaining agents who cannot afford one more tuple. We omit further details to the full version.

Strategyproof mechanisms for maximizing social efficiency Coming to strategyproof mechanisms, and in analogy to Section 4.1, we have implemented the MC mechanism in this setting as well. Unlike Section 4.1, in Scenario 2, there is no need for heuristics since there is an efficient algorithm for computing the optimal welfare. The important property that allows for efficient computation is that once we decide for allocating a tuple, we do not lose by giving the tuple to all customers who have demand for it, since we are only adding more utility to the current welfare, (utility functions are non-decreasing in the number of tuples acquired). The optimization problem then becomes

$$SW^*(N, \mathbf{u}) = \max_{1 \le k \le t_{max}} [\sum_{j \in N} \sum_{\ell=1}^{min\{k, m_j\}} u_j(\ell) - \sum_{i=1}^{k} C(i)] \tag{8}$$

We can solve (8) simply by trying all values for k. Clearly all involved quantities can be computed efficiently. Hence we have:

Theorem 2. *The optimal social welfare can be computed in polynomial time.*

Applying now the MC mechanism is straightforward and follows the same reasoning as in Section 4.1.

B Experimental Results on Social Welfare Approximation in Scenario 1

In this Section we provide the missing figures regarding the performance of the heuristics in Scenario 1, for approximating the optimal social welfare. Comments on the simulations can be found in Section 5.3.

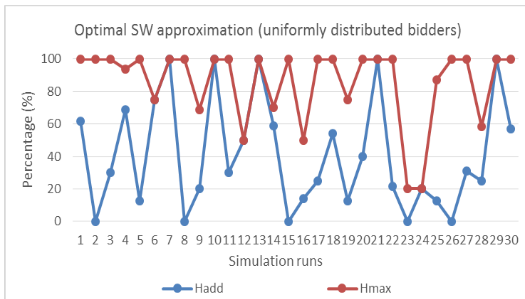

Fig. 3. Welfare approximation for small number of buyers (compared with the optimal welfare).

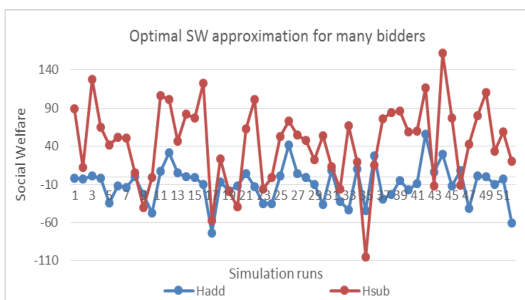

Fig. 4. Comparing H_{add}, H_{sub} for large numbers of buyers.