

Eliciting Truthful Measurements from a Community of Sensors

Boi Faltings, Jason Jingshi Li
Artificial Intelligence Laboratory (LIA)
Swiss Federal Institute of Technology, Lausanne (EPFL)
Station 14, IN-Ecublens
1015 Lausanne, Switzerland
Email: boi.faltings@epfl.ch, jason.li@epfl.ch

Radu Jurca
Google
Zurich
Switzerland
Email: radu.jurca@gmail.com

Abstract—

As the Internet of Things grows to large scale, its components will increasingly be controlled by self-interested agents. For example, sensor networks will evolve to *community sensing* where a community of agents combine their data into a single coherent structure. As there is no central quality control, agents need to be incentivized to provide accurate measurements. We propose game-theoretic mechanisms that provide such incentives and show their application on the example of community sensing for monitoring air pollution. These mechanisms can be applied to most sensing scenarios and allow the Internet of Things to grow to much larger scale than currently exists.

I. INTRODUCTION

With increasing progress of sensing and networking technology, it is now possible to deploy dense networks of low-cost sensors that measure qualities of the environment, such as air pollution, noise, radiation or spectrum use. Progress in nanotechnology in particular allows complex sensing equipment to be replaced by solid-state electronic sensors that can operate continuously and are available at low cost. The sensors will extend the notion of the Internet of Things beyond simple location measurements.

Even though existing finite-volume physical models are capable of producing rough pollution estimates of large areas without any measurements, sensor data will be most useful when it can be integrated into a joint data-driven model that can produce a detailed street-level pollution map to be shared by many users. For the purpose of minimizing street-level pollution exposure, we need a dense network of many sensors that are spread over different locations. Many of these locations will be owned by private individuals, and it would be most practical if these individuals themselves installed, maintained and operated the sensors. This is the idea of community sensing, where sensors are controlled by different agents and their results integrated in an open sensor network [7], [1].

We are involved in a project that explores community sensing for measuring air quality using both fixed and mobile sensors. Examples are shown in Figure 1. Initial experiments with such sensors suggests that air quality varies significantly



Fig. 1. Air quality sensors that could be used in community sensing. Top left: on top of a bus; bottom left: on top of a tram; top right: attached to a solar-powered weather station on a building; and bottom right: attached to a smartphone.

throughout a city, and thus it would be very useful to maintain accurate maps that allow people to reduce their exposure to harmful pollution.

Community sensing will allow much larger and comprehensive sensor networks than can currently be imagined, and will be an important next step for the Internet of Things. However, ensuring that the data obtained is of sufficient quality is a significant challenge ([1]). As there are many individual and self-interested participants, it is not possible to control the network through a central authority. Instead, participants need to be given *incentives* that will make them cooperate with the sensor network in the best possible way. This is the topic of this paper.

The most straightforward way to incentivize accurate measurements is by payments that compensate for the effort of providing measurements. However, sometimes agents may have strong incentives to report incorrect values, for example to hide pollution they caused themselves, and small monetary incentives would fail to compensate for them. In such a case,

the incentives can be used as a reputation so that agents that fail to consistently gather rewards lose their credibility. The mechanisms we consider can also be used as part of such a reputation system.

This paper is structured as follows. First, we define the setting and assumptions behind our mechanism. Following a review of game-theoretic mechanisms for incentivizing truthful information revelation, we define the novel mechanism we propose for community sensing, called Peer Truth Serum, and discuss its properties. Finally, we illustrate the mechanisms on an example.

II. THE SETTING

We assume that there is an open group of agents distributed in space that take measurements of a continuous space-time physical process, such as air pollution levels in a city over the course of a day. While typically we measure several different quantities, for the purpose of this paper we assume that a single quantity called pollution is measured, and that at any given location l and time t , there are N possible pollution levels denoted as $V = \{v_1, \dots, v_N\}$.

Agents report each observation o as a report s to a *center* that they trust to aggregate their reports. The center integrates the reported data with the known emission and dispersion characteristics in an environmental model to produce a pollution map. The model uses a partition of the space into different regions, with the assumption that pollution levels within each region are relatively homogenous. For each region, it has a prior expectation of pollution levels that is given by known emission and meteorological information, such as nearby chimneys, traffic volumes and current wind field. It combines this expectation with reports for the region to produce a maximum-likelihood estimate. This estimate also takes into account statistical correlation between regions. While the details of such a complex environmental model are beyond the scope of this paper (see for example [6] for more discussion), we only need to know that the output of the model is a pollution map represented as a full probability distribution over the possible pollution levels at every location l . In this map, we let $R^{l,t}(v)$ denote the probability that the pollution at location l and time t is of level v ; $R^{l,t}(v) > 0$ for all $v \in V$.

The center updates the map periodically using the measurement reports it received during the last time instant. Depending on the frequency of reports, updates may happen as frequently as every hour or as infrequently as once a week. Each agent has private prior beliefs $Pr^{l,t}(v)$ about the pollution levels that the model will report at the next update, $R^{l,t+1}$. Before measurement, these private beliefs will generally be close to the current map $R^{l,t}$, but they can diverge significantly after the agent makes a measurement. We let $Pr^{l,t}(v)$ be the belief before measurement that the model will report $R^{l,t+1} = v$ after the next update, and $Pr_o^{l,t}$ be the belief after measuring value o . In the following, we will always consider a single location and time point only, and thus drop the l, t superscripts.

Figure 2 illustrates how an agent's beliefs are influenced by its observation. The curve labelled $Pr(x)$ shows the prior




Fig. 2. Probability distributions for a variable x .

probability distribution that the agent has about the value of variable x before measuring it. It shows that b is believed to be the most likely value. Once the agent measures the actual value of the variable to be a or c , its belief changes to the distribution $Pr_a(x)$ or $Pr_c(x)$, respectively. Note the influence of the prior belief: when the agent has measured c , the most likely value may not be c itself, but a value between c and b .

While in general the prior beliefs of agents can be expected to be quite homogeneous and close to the public distribution R , the way that beliefs are updated is likely to be very different between agents. An agent who strongly believes its own measurements is likely to change its beliefs more dramatically, and thus obtain a significantly different posterior distribution. However, one assumption that is reasonable to make is that an agent believes that its own measurement is positively correlated with the measurements that other agents will report, meaning that the measured value has the biggest relative increase in probability:

Definition 1: An agent's belief update from prior Pr to posterior Pr_x after measuring x satisfies the *rational update* property if and only if:

$$\frac{Pr_x(x)}{Pr(x)} > \frac{Pr_x(y)}{Pr(y)} \quad \forall y \neq x \quad (1)$$

If this property would not hold, agents do not believe that they are measuring the same quantities that are modelled and it makes no sense to compare or aggregate their measurements. It is thus an assumption that we will make in this paper.

III. INCENTIVE MECHANISMS FOR OBTAINING TRUTHFUL REPORTS

Rewarding agents to provide truthful reports of their private information has been studied in game theory. In this section, we review this earlier work. All schemes are based on the fact that agents' posterior beliefs change according to their observation, as shown in the example in Figure 2. As agents will compute the rewards they expect from a report using this belief, the incentives can be scaled so that for each observation, reporting the true observation gives the highest expected reward given the associated posterior belief.

Different schemes can be used depending on whether the goal is to get an agent to truthfully report its posterior probability distribution or the value it has actually observed.

We first consider incentives to truthfully report the posterior probability distribution.

A. Reporting the full posterior distribution

For problems such as weather prediction, where a true value eventually becomes known, such incentives can be provided by *proper scoring rules* [12]. Agents submit a probability distribution $p(x)$ on their best estimate of the value of the variable x that is to be predicted. Once the true value \bar{x} becomes known, they get rewarded according to a scoring rule applied to the probability $p(\bar{x})$ they predicted for this true value. Examples of proper scoring rules are:

- the logarithmic scoring rule:

$$pay(\bar{x}, p) = a + b \cdot \log p(\bar{x}) \quad (2)$$

- the quadratic scoring rule:

$$pay(\bar{x}, p) = a + b \left(2p(\bar{x}) - \sum_v p(v)^2 \right) \quad (3)$$

It is also possible to use scoring rules to elicit averages, maxima and other functions of a set of measurements, see [5] for a complete characterization of the possibilities offered by scoring rules.

However, in pollution sensing, it is generally not possible to ever know the ground truth as required by scoring rule. Peer prediction [8] is a technique for this setting. The principle is to consider the reports of other agents that observed the same variable, or at least a *stochastically relevant* variable, as the missing ground truth. A proper scoring rule is then used for the incentives. Provided that other agents truthfully report an unbiased observation of the variable, such a reward scheme makes it a best response to provide truthful and unbiased reports of the observations, and truthful reporting thus becomes a Nash equilibrium. [8] describe such a mechanism and several variants, and [3] discuss further optimizations and variants.

An important issue with implementing peer prediction mechanisms is that agents should report both the value they observed and the posterior probability distribution that resulted: the value is needed in order to be able to score other reports, while the distribution is needed to determine a payment to the agent itself. In the approach originally proposed by [8], the agents report a value and the center replaces this by an assumed posterior distribution for agents that have observed this value. The limitation of this approach is the need to know agents posterior beliefs. The Bayesian Truth Serum [10] is a mechanism that elicits both the prior beliefs and the observation, but only applies when these are not revealed to other agents, which is not the case in community sensing.

To overcome this limitation, in [14] the authors provide a mechanism where agents report both their prior and posterior beliefs about the observed value. Noting that Bayesian updating implies that the ratio of posterior/prior is the highest for the actually observed value (the rational update assumption), the two reports together also determine the true value. However,

it is difficult to apply this technique to community sensing since we cannot enforce reporting the prior beliefs before an observation.

Applying the peer prediction approach to our setting has the challenge that sensors are taking measurements at different locations, i.e. we do not have another sensor reading of exactly the same value. However, the peer prediction method as defined by [8] only requires a stochastically relevant signal. Similar to [13], we can obtain such a stochastically relevant signal by using a pollution model applied to the combined set of measurements reported by other agents.

Work by Papakonstantinou, Roger, Gerding and Jennings investigated a multi-agent scenario where the center specifies the data wanted, and then incentivizes agents to provide that data [9]. The approach combines a first stage where the center selects the agent that can provide the measurement in the most cost-effective way with a second stage where either the observation is scored against a true value that becomes known later, or against another report using the peer prediction principle. The approach assumes a pull approach where the center decides what measurements are important and specifically asks agents to report these.

B. Reporting only the measured value

In community sensing, reporting entire probability distributions is not desirable as it greatly increases the load on already limited communication bandwidth. Therefore, it is desirable to only transmit a report of the measured value itself.

The most straightforward way is to let the center substitute a standardized posterior distribution for each reported value, and let the agent select the right distribution by reporting one of the values. This was the approach originally adopted in the peer prediction method ([8]).

In [3], the peer prediction principle is implemented directly without using scoring rules. Instead, for each combination of report and reference report, minimal truthful payments are computed directly using linear programming. It is shown that these payments can often be much more efficient than those obtained by assuming posterior distributions and applying proper scoring rules, and satisfy other properties such as resistance against collusion. However, they still require assumptions about the posterior beliefs of the agents.

[16] investigates mechanisms that are robust to variations of these beliefs, and shows that this is only possible in very limited ways and leads to large increases in payments.

[2] proposes a mechanism for truthful opinion polls with two possible values that requires no assumptions about posterior distributions. While the mechanism is not always truthful, it is *helpful* in the sense that non-truthful reports only help to make the public poll outcome converge to the true distribution more rapidly. Thus, the mechanism is shown to be *asymptotically truthful* in the sense that it converges to the true distribution. [4] shows how to extend this mechanism to settings with more than two values. The setting assumed in their mechanism is very close to the pollution sensing problem: the publicly available prior corresponds exactly to the pollution

map. We will therefore adopt a very similar mechanism for our problem.

IV. THE PEER TRUTH SERUM

We propose a new mechanism for incentivizing truthful measurement reporting, which we call the *Peer Truth Serum*:

Definition 2: The *Peer Truth Serum* is a payment function that rewards an agent for reporting a value s of a variable that is compared against a reference estimate m for the same variable, given a publicly available prior probability distribution R for the variable. It rewards the agent according to the payment function $\alpha + \beta \cdot \tau(s, m, R)$:

- $\tau(s, m, R) = \frac{1}{R(m)}$ if $s = m$
- $\tau(s, m, R) = 0$ otherwise.

where α and $\beta > 0$ are constants.

In our scenario, agent i measures the pollution level at location l and time t , and reports the value $s = s_i^{l,t}$. The report is evaluated against a reference value $m = m^{l,t+1}$ from the model, based on an update using other reports received in the same time interval. The reward is computed using the known public prior $R = R^{l,t}$.

As an example, consider a range of three values for the pollution level: $a(\text{low})$, $b(\text{medium})$ and $c(\text{high})$, and let the public prior for some l, t be:

x	a	b	c
R(x)	0.2	0.6	0.2

Assume that the agent measures b , and truthfully reports this value. The center obtains a reference report $m = b$ and finds that it matches the report of the agent. Letting $\alpha = 0$ and $\beta = 1$, the agent would be rewarded $\tau(s, m, R) = 5/3$. The agent might also report a , but it is less likely that a would match the value reported by the model. However, if it does, the agent would get the much higher reward of 5. Thus, we can see that the payment scheme balances out the risk inherent in reporting unlikely values.

In practice, an issue that might arise is that for very small R , the payment can become unboundedly large. It will often be desirable to impose a budget limit so that the payment cannot exceed this limit.

While the likelihood of matching the reference report, and thus obtaining a reward at all, is highest for reporting a very common value, the amount of reward is highest for uncommon values. Together these two influences make it optimal for an agent to report its true measurement, as we will now show. We first consider a setting where all agents adopt the publicly available map R as their prior distribution, and later consider more informed agents.

A. Agents adopt the public prior distribution

Proposition 1: There exists a threshold $\epsilon > 0$ such that when an agent's prior distribution $Pr(\cdot)$ for a variable is within ϵ of the publicly available distribution R :

$$Pr(v) + \epsilon > R(v) > Pr(v) - \epsilon \quad (4)$$

the Peer Truth Serum incentivizes truthful reporting.

Proof: We observe that an agent who observes o and reports s expects a reward:

$$pay(o, s) = \alpha + \beta \sum_x Pr_o(x) \tau(s, x, R) = \alpha + \beta Pr_o(s) \frac{1}{R(s)}$$

In order for the mechanism to be truthful, we require that for $v \neq o$, $pay(o, o) \geq pay(o, v)$, i.e.:

$$\frac{Pr_o(v)}{R(v)} \leq \frac{Pr_o(o)}{R(o)} \Leftrightarrow \frac{R(v)}{Pr_o(v)} \geq \frac{R(o)}{Pr_o(o)}$$

Given the assumption 4, this holds under the condition that:

$$\frac{Pr(v) - \epsilon}{Pr_o(v)} \geq \frac{Pr(o) + \epsilon}{Pr_o(o)}$$

As $\frac{Pr_o(o)}{Pr(o)} > \frac{Pr_o(v)}{Pr(v)} \forall o, v$, let

$$\delta(o, v) = \frac{Pr(v)}{Pr_o(v)} - \frac{Pr(o)}{Pr_o(o)} > 0$$

then the truthfulness condition holds for any ϵ such that:

$$(\forall o, v) \delta(o, v) \geq \frac{\epsilon}{Pr_o(o)} + \frac{\epsilon}{Pr_o(v)}$$

As $\delta(o, v) > 0$, such an ϵ always exists and can be calculated as:

$$\epsilon = \min_{v, o, v \neq o} \frac{Pr(v)Pr_o(o) - Pr(o)Pr_o(v)}{Pr_o(v) + Pr_o(o)} \quad (5)$$

Thus, when agents adopt to public prior within some tolerance ϵ , the mechanism incentivizes truthful reporting.

For the example given earlier, assume that an agent's prior and posterior beliefs are as follows:

x	a	b	c
$Pr_a(x)$	0.6	0.3	0.1
$Pr_b(x)$	0.1	0.8	0.1
$Pr_c(x)$	0.1	0.3	0.6
$Pr(x)$	0.2	0.6	0.2

Now we can compute ϵ according to Equation 5 as $\min(1/3, 1/7, 1/9) = 1/9$. Thus, for example, if the public distribution R is within the bound of $1/9$ from the agent prior:

x	a	b	c
$R(x)$	0.25	0.5	0.25

depending on its observation o , the agent would expect the following payments for its reports:

s	a	b	c
$o = a$	2.4	0.6	0.4
$o = b$	0.4	1.6	0.4
$o = c$	0.4	0.6	2.4

and thus truthful reporting always gives the highest payoff.

B. Agents do not adopt the public prior

In some cases, agents may be more informed than the public model. For example, they may observe that there are traffic jams, fires or other incidents that will cause the pollution level to be higher than expected by the model. In this case, their prior belief even before measurement could be considerably different from the public map R .

If this means that the difference between R and the private belief Pr is larger than the threshold ϵ , the agent may no longer be incentivized to report truthfully. For example, given the private beliefs as above, if R were as follows:

x	a	b	c
$R(x)$	0.5	0.1	0.4

depending on its observation o , the agent would expect the following payments:

s	a	b	c
$o = a$	1.2	3	0.25
$o = b$	0.2	8	0.25
$o = c$	0.2	3	1.5

and thus report b no matter what the actual observation was.

While the fact that the report is not truthful may be considered undesirable, note that in this example, reporting b actually helps the public report R to converge more quickly to the agent's private belief than reporting truthfully.

This is interesting in particular if the agent's private belief is more informed than the public map, i.e. that it is closer to the true value distribution:

Definition 3: An agent's prior beliefs $Pr[\cdot]$ about a signal with true distribution $Q[\cdot]$ are *informed* with respect to a public prior $R[\cdot]$ if and only if for all v , either $R[v] \leq Pr[v] \leq Q[v]$ or $R[v] \geq Pr[v] \geq Q[v]$.

In such a case, it would be most helpful to make the public map R converge to the private beliefs as quickly as possible. We are now going to show that the Peer Truth Serum incentivizes *helpful* reports that drive the public map closer to the true distribution without necessarily being truthful.

Thus, convergence happens in two steps:

- 1) first the diverse private prior distributions and the published pollution map converge to the same distribution, establishing a common frame of reference, and
- 2) once this is established, the incentives are for truthful reporting and both the public map and the private priors converge asymptotically towards the true distribution.

Such a two-step process makes a lot of sense in community sensing, since a sensor is usually present in the system for an extended period of time and will only have to pass the initial phase once when joining the network.

We first show the following property of the Peer Truth Serum:

Proposition 2: Provided the the rational update assumption (1) holds and all agents prior beliefs are informed, using the Peer Truth Serum no agent ever reports a non-truthful answer $s = b$ when according to its beliefs, the true answer a

is more under-represented in the current public prior R :

$$Pr(a)/R(a) > Pr(b)/R(b) \Rightarrow s(a) \neq b;$$

Proof: For the case where the agent believes the reference report to be truthful, this follows directly from the rational update assumption and the payment rule. After observing a , the expected payments are for reporting a :

$$\frac{Pr_a(a)}{R(a)} = \frac{Pr_a(a)}{Pr(a)} \frac{Pr(a)}{R(a)}$$

and for reporting b :

$$\frac{Pr_a(b)}{R(b)} = \frac{Pr_a(b)}{Pr(b)} \frac{Pr(b)}{R(b)}$$

The first term is greater for a than for b by the rational update assumption, and the second term is greater for a than for b by the condition of the proposition. Thus, the agent will not report b instead of a .

For the case where the agent believes that the agent providing the reference report also misreports using an informed prior, as it knows that this other agent will not report b instead of a , misreporting b for a would only lower the probability of matching reports and thus not be rational.

Thus, in all equilibria where agents have informed priors and believe each other to have informed priors, the proposition holds. ■

We now use this result to show the following:

Proposition 3: In the current distribution R , let A be the set of underreported values ($\forall a \in A, R(a) < Pr(a)$) and B the set of overreported values ($\forall b \in B, R(b) \geq Pr(b)$). There will never be a non-truthful report for some answer $b \in B$ instead of another answer $a \in A$. Thus, provided that the agent's prior beliefs are informed with respect to R and the true distribution, the combined frequency of reports of values $b \in B$ is not greater than the agent's believed frequency $\sum_{b \in B} Pr(b)$.

Proof: For all $a \in A$, $R(a)/Pr(a) < 1$ whereas for all $b \in B$, $R(b)/Pr(b) \geq 1$. By Proposition 2, there are never any reports of values in B when the true values were in A . Thus, the combined frequency of all reports of values in B cannot be larger than the true frequency $\sum_{b \in B} Q(b)$. By the assumption that the belief Pr is informed, we have $\sum_{b \in B} Q(b) \leq \sum_{b \in B} Pr(b) \leq \sum_{b \in B} R(b)$, and thus the combined frequency is also not larger than $\sum_{b \in B} Pr(b)$. ■

Now recall that the public statistic R is updated by averaging the reports obtained from agents. Thus, we have:

Proposition 4: Within some finite amount of updates, for all values of $b \in B$, the public statistic $R(b) < Pr(b) + \epsilon$, and consequently for all values of $a \in A$, $R(a) > Pr(a) - \epsilon$.

Proof: The frequency of values in B will be not larger than what is believed by the agent, so R will gradually be reduced to become arbitrarily close to Pr . Likewise, the frequency of reports of values in A will be at least as large as what the agent believes, and thus also become arbitrarily close to Pr . ■

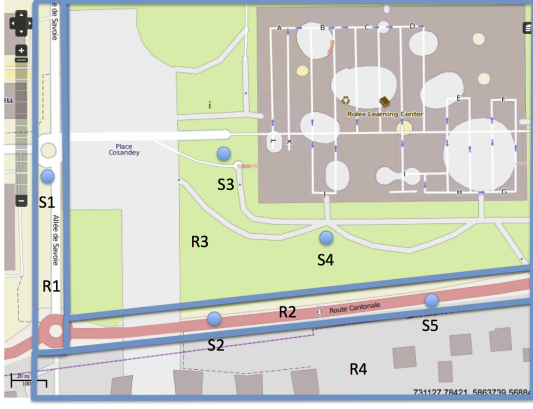


Fig. 3. A spatial setting of four regions.

Thus, agents that have prior distributions that diverge from the public prior in an informed way will provide *helpful* reports that drive the public map close to its own beliefs.

When the private priors are not informed, such convergence may still happen, but cannot be guaranteed. However, such a case is not realistic: either an agent has background information not accessible to the center, and in this case its beliefs should be more informed, or otherwise it should believe the distribution given by the center. Another issue is what happens when agents have informed private prior distributions but they differ significantly.

Both cases are helped by the fact that rational agents should gradually adapt their beliefs about the model output to the published distribution R , and thus eventually converge to a single distribution. However, such convergence may be undesirably slow.

For the case where the private prior Pr is equal to the true distribution Q , helpful reports actually speed up convergence to the true map. This is because the untruthful reports are always for values where R/Pr is lower than for the true value, i.e. values where R should be increased more strongly to approach Pr . Helpful reports can thus be more valuable than truthful reports.

V. EXAMPLE

We consider the setting shown in Figure 3 where five agents $\{S1, \dots, S5\}$ are making air-quality measurements in different locations. The center divides the area into four regions: the side street on the east (R1), the main road on the south (R2), the library, which is the region north of the main street and east of the side street (R3) and the region south of the main street (R4), and uses three possible pollution levels $V = \{low, medium, high\}$.

We compare two different incentive schemes: peer prediction as described in [8], [9] using the quadratic scoring rule:

$$pay(\bar{x}, p) = 2p(\bar{x}) - \sum_v p(v)^2$$

and the Peer Truth Serum mechanism we propose in this paper. In the peer prediction mechanism, for each possible

value that an agent might report the center needs to define a posterior distribution to assume. For this example, we assume that the center derives an assumed probability distribution $\hat{P}r_s^{l,t}$ for reported value s by taking the public map as the prior distribution, increasing the probability for the reported value by 50%, and then renormalizing the distribution:

$$\hat{P}r_s^{l,t}(x) = 1.5R^{l,t}(x)/\alpha \text{ if } x = s, R^{l,t}(x)/\alpha \text{ otherwise}$$

with α chosen so that $\sum_x \hat{P}r_s^{l,t}(x) = 1$. The difficulty with peer prediction is that agents might not actually have this posterior distribution, and this can cause failures as we will see below.

We now illustrate the two incentive schemes on two example measurements, one where both encourage a truthful report and one where both encourage a non-truthful report. The incentives that are computed can become a payment to reward the agent for its effort, or they could be reputation that accumulates and determines an agent's influence on the public map.

A. Example of Truthful Reports

First, we look at the peak hour of $t1 = 18:00$ where the public prior for the pollution level at the library (R3) is published. At the same time, agent $S3$ has a private prior distribution $Pr^{R3,t1}$ that is influenced by observing the current weather and traffic conditions, and therefore somewhat different from the current map value.

	low	medium	high
$R^{R3,t1}$	0.1	0.5	0.4
$Pr^{R3,t1}$	0.15	0.7	0.15

The agent measures that the level is in fact *medium*, and updates her belief to obtain the posterior belief $Pr_{medium}^{R3,t1}$ as follows:

o	low	medium	high
$Pr_{medium}^{R3,t1}$	0.1	0.8	0.1

During the same time interval, the center also receives reports of *medium* levels from $S1$ and $S4$, and *high* levels from $S2$ and $S5$, and thus concluded that the pollution level at the location of $S3$ is $m^{R3,t1} = \text{medium}$. However, the agent does not know anything about these measurements except that it assumes them to be truthful, and so its best guess is that $m^{R3,t1}$ is drawn from the same distribution as its own posterior.

1) *Peer prediction with quadratic scoring rule*: Using the assumption about the agent's posterior belief given above, the center computes the following assumed posterior probability distributions for each reported value s :

s	low	medium	high
$\hat{P}r_{low}^{R3,t1}$	0.14	0.48	0.38
$\hat{P}r_{medium}^{R3,t1}$	0.08	0.6	0.32
$\hat{P}r_{high}^{R3,t1}$	0.08	0.42	0.5

The agent can either obtain these distributions from the center, or compute them itself using the public map $R^{R3,t1}$ and the

rule the center uses to compute the assumed posterior. Using its true posterior distribution, the agent can now compute the expected reward when reporting the different values, given by the probability that the reported value matches the model times the reward that would result in that case:

s	$E[\text{pay}(s)]$
<i>low</i>	$0.1 \cdot (2 \cdot 0.14 - 0.14^2 - 0.48^2 - 0.38^2) = -0.011$
<i>medium</i>	$0.8 \cdot (2 \cdot 0.6 - 0.08^2 - 0.6^2 - 0.32^2) = 0.58$
<i>high</i>	$0.1 \cdot (2 \cdot 0.42 - 0.08^2 - 0.42^2 - 0.5^2) = 0.056$

and so it can expect the highest reward when truthfully reporting *medium*.

2) *Peer Truth Serum*: As above, upon measuring a level of *medium* the agent updates its belief and can compute its expected payment for the different possible reports (assuming $a = 0$ and $b = 1$):

s	low	medium	high
$E[\text{pay}(\text{medium}, s)]$	0.1/0.1 = 1	0.8/0.5 = 1.6	0.1/0.4 = 0.25

So the expected payment is highest for a truthfully reporting the pollution level to be *medium*.

B. Example of Non-Truthful/Helpful Reports

We now look at the situation one hour later ($t_2=19:00$) and agent S_1 is making measurements on the side street (R1). The current public map of the pollution levels has a different distribution. At the same time, agent S_1 might know that a moderate traffic jam has just developed on the main road, and that winds blow the pollution into the side street. Consequently, her private belief about the pollution value became skewed to the higher value.

	low	medium	high
R^{R1,t_2}	0.7	0.1	0.2
P_r^{R1,t_2}	0.3	0.35	0.35

Subsequently, S_1 measures the level to be *high*, and gets the following posterior:

o	low	medium	high
$P_r^{R1,t_2}_{\text{high}}$	0.1	0.4	0.5

1) *Peer prediction with quadratic scoring rule*: Using the assumption about the agent's posterior belief given above, the center computes the following assumed posterior probability distributions for each reported value s :

s	low	medium	high
$\hat{P}r_{\text{low}}^{R1,t_2}$	0.78	0.07	0.15
$\hat{P}r_{\text{medium}}^{R1,t_2}$	0.66	0.14	0.19
$\hat{P}r_{\text{high}}^{R1,t_2}$	0.64	0.09	0.27

Using its true posterior distribution, the agent can now compute the expected reward when reporting the different values, given by the probability that the reported value matches the model times the reward that would result in that case:

s	$E[\text{pay}(s)]$
<i>low</i>	$0.1 \cdot (2 \cdot 0.78 - 0.78^2 - 0.07^2 - 0.15^2) = 0.092$
<i>medium</i>	$0.4 \cdot (2 \cdot 0.14 - 0.66^2 - 0.14^2 - 0.19^2) = -0.085$
<i>high</i>	$0.5 \cdot (2 \cdot 0.27 - 0.64^2 - 0.09^2 - 0.27^2) = 0.025$

and so non-truthfully reporting *low* give the highest payoff. Note that the scoring rule is not truthful because the posterior distribution assumed by the center is very different from the agent's true distribution. Note also that this non-truthful report, if taken into account by the center, will make the public map diverge even more from the truth.

2) *Peer Truth Serum*: S_1 believes that other agents, for example S_2 or S_5 , would also report much higher pollution levels, and assumes the reference value predicted by the model to follow this posterior distribution, to obtain the expected payments:

s	low	medium	high
$E[\text{pay}(\text{high}, s)]$	0.1/0.7 = 0.143	0.4/0.1 = 4	0.5/0.2 = 2.5

So in this case the highest expected payment is for the agent to report *medium*. Although this is not the truthful report, we have shown in section IV-B that it is nevertheless a *helpful* report, which drives the public map closer to the agent's private beliefs. When the two coincide, reporting the truth will become the best policy.

VI. CONCLUSIONS

For the Internet of Things to grow to large scale, it needs to develop mechanisms to self-organize in the presence of multiple self-interested agents controlling the devices. For sensor networks, this next step is *community sensing*[7], [1], where sensors are operated by individuals that feed their measurements into a common network.

We proposed the Peer Truth Serum, an incentive mechanism for a community sensing scenario that rewards accurate and truthful measurements as well as providing information that updates the public model. It is the first mechanism that does not need to make strong assumptions about the agents' prior beliefs or updating mechanism, and is thus realistic for a practical setting. After an initial adaptation phase where agents adjust their private beliefs and the publicly available map, the incentive scheme motivates agents to contribute truthful and accurate measurements. It thus provides the necessary quality control to ensure that the result of the community sensor network are valid in spite of the absence of explicit control.

[11] have shown a way to use truthful information elicitation based on scoring rules as reputation feedback that can adjust the influence of raters to their credibility. A similar approach could be used to adjust the influence of sensors to their accuracy and thus use the reward scheme as a reputation system. Such a reputation system could provide much stronger quality control, including guarding against reports that are intentionally false.

While the mechanism ensures that agent beliefs will converge to a common value even when they start out from very

different values, in community sensing agents observe the same local phenomena and should have similar prior beliefs (even these are unknown to the center). For example, if some area experiences pollution due to fires, this will be apparent to agents in the area, although it would not be to the center. We therefore expect that the mechanism will quickly converge to the truthful reporting regime, while still being robust to new agents that may not share the prior beliefs.

Other issues that have been of concern in other applications of truthful elicitation mechanisms are less of a concern in our setting. In particular, collusion among agents that measure in related locations is not very likely, as measurements are not anonymous as for example in product rating. Also, strategic timing of reports is unlikely as pollution values change in ways that are hard to predict.

An important open issue is that besides encouraging agents to report accurate measurements, we also want them to provide measurements that improve the map as much as possible. [9] addresses this problem in the framework of scoring rules. However, in their framework, sensing is driven by requests of the center who may not be aware of where additional measurements are needed. In future work, we plan to develop a similar framework that is driven by the community itself, so that agents are incentivized to provide the measurements that they believe will provide the most information to the map.

ACKNOWLEDGEMENT

This research has been partially funded by the Nano-Tera program as part of the OPENSENSE project.

REFERENCES

- [1] K. Aberer, S. Sathe, D. Chakraborty, A. Martinoli, G. Barrenetxea, B. Faltings, and L. Thiele. Opensense: Open community driven sensing of environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, 2010.
- [2] R. Jurca, R. and B. Faltings. Incentives for Expressing Opinions in Online Polls. *Proceedings of the 2008 ACM Conference on Electronic Commerce*, pp. 119-128, 2008.
- [3] R. Jurca and B. Faltings. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research (JAIR)*, 34, 2009, pp. 209-253.
- [4] R. Jurca and B. Faltings. Incentives for Answering Hypothetical Questions. *Workshop on Social Computing and User Generated Content*, ACM Conference on Electronic Commerce, San Jose, 2011.
- [5] N. Lambert and Y. Shoham. Eliciting Truthful Answers to Multiple-Choice Questions. In *Proceedings of the tenth ACM conference on Electronic Commerce*, pp. 109-118, 2009.
- [6] J. J. Li and B. Faltings. Towards a Qualitative, Region-Based Model for Air Pollution Dispersion *Workshop on Space, Time and Ambient Intelligence*, International Joint Conference on Artificial Intelligence, Barcelona, 2011.
- [7] A. Krause, E. Horvitz, A. Kansal and F. Zhao. Toward Community Sensing. In *Proceedings of Information Processing in Sensor Networks (IPSN)*, 2008.
- [8] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359–1373, 2005.
- [9] A. Papakonstantinou, A. Rogers, E. H. Gerding, and N. R. Jennings. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artif. Intell.*, 175(2):648–672, 2011.
- [10] D. Prelec. A Bayesian Truth Serum for Subjective Data. *Science*, 306(5695), pp. 462-466, 2004.
- [11] P. Resnick and R. Sami. The Influence Limiter: Provably Manipulation-Resistant Recommender Systems. *Proceedings of ACM RecSys 07 Conference on Recommender Systems*. Minneapolis, MN. pp. 25-32, 2007.
- [12] L. J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [13] J. Wikowski. Eliciting Honest Reputation Feedback in a Markov Setting. *Proceedings of the 21th International Joint Conference on Artificial Intelligence (IJCAI 2009)*. 2009.
- [14] J. Witkowski and D. C. Parkes. Peer Prediction without a Common Prior. *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 964-981, 2012.
- [15] World Health Organization. Air Quality and Health, Fact sheet No. 313, 2011.
- [16] A. Zohar and J. S. Rosenschein. Robust Mechanisms for Information Elicitation, in *AAAI '06: Proceedings of The Twenty-First National Conference on Artificial Intelligence*, July 16-20, AAAI Press, Menlo Park, CA, USA, 2006.