Computer Networks 73 (2014) 58–71

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Fairness and social welfare in service allocation schemes for participatory sensing

Chen-Khong Tham [a,*], Tie Luo [b,1]

[a] Department of Electrical and Computer Engineering, National University of Singapore, Singapore
[b] Institute for Infocomm Research, A*STAR, Singapore

ARTICLE INFO

Article history:
Received 5 January 2014
Received in revised form 19 May 2014
Accepted 23 July 2014
Available online 6 August 2014

Keywords:
Participatory sensing
Service allocation
Optimization
Utility and fairness

ABSTRACT

Leveraging on the pervasiveness of mobile phones and their rich built-in sensors, participatory sensing recently emerged as a promising approach to large-scale data collection. Whilst some contributors may be altruistic, many contributors are motivated by receiving something in return for their contributions, proportional to their level of contributions. In this paper, we adopt a service allocation approach that motivates users by allocating a determined amount of compelling services to contributors, as an alternative to other credit or reputation based incentive approaches. To address two major concerns that would arise from this approach, namely fairness and social welfare, we propose two service allocation schemes called Allocation with Demand Fairness (ADF) and Iterative Tank Filling (ITF), which is an optimization-based approach. We show that: (i) ADF is max–min fair and scores close to 1 on the Jain's fairness index, and (ii) ITF maximizes social welfare and achieves the unique Nash equilibrium, which is also Pareto and globally optimal. In addition, we use stochastic programming to extend ITF to handle uncertainty in service demands that is often encountered in real-life situations.

1. Introduction

The vast penetration of mobile devices, each with a variety of built-in sensors such as GPS, accelerometer, microphone and camera, has recently spurred the emergence of a new sensing paradigm called *participatory sensing* [1] for large-scale data collection. It refers to an approach whereby individuals use their personal mobile devices to sense the environment and report sensed data in a non-obligated manner. Compared to conventional wireless sensor networks that use deployed sensors, participatory sensing removes the cost and hassle of sensor installation and network maintenance, while achieving much broader geographical coverage. The big challenge of prolonging network lifetime in conventional sensor networks also becomes a non-issue in participatory sensing because charging the battery on each mobile device is now taken care by the users themselves. Hence, participatory sensing is a promising paradigm and has attracted a lot of research attention and efforts such as [2–4].

However, the success of participatory sensing strongly relies on user participation in order to obtain a sufficient and continuous influx of user contributions. So far, most participatory sensing projects focus on developing applications, e.g., for environmental monitoring [2,3] and transportation [4], in which participants are recruited on a voluntary or remunerated basis, which is not sustainable in the long run [5].

* Corresponding author.
E-mail addresses: eletck@nus.edu.sg (C.-K. Tham), luot@i2r.a-star.edu.sg (T. Luo).
[1] This work was done when Tie Luo was at the Department of Electrical & Computer Engineering, National University of Singapore.

http://dx.doi.org/10.1016/j.comnet.2014.07.013
1389-1286/

This brings forth the important issue of incentive or motivation, which is a key ingredient for the success of participatory sensing. In this paper, we propose a service allocation approach that motivates users by allocating a determined amount of compelling services to contributors, as an alternative to other credit or reputation based incentive approaches. We leverage on the duality of a participating user's role, whereby each user acts as a *data contributor* as well as a *service consumer*, and capture this in a generic framework for participatory sensing applications which consists of a *service provider* that receives the data sent by contributors and processes and packages the data into useful data services that are consumed by service consumers. A typical example is traffic monitoring as shown in Fig. 1, in which bus or car passengers use their smartphones to send traffic-related data, such as GPS traces and bus crowd levels, to a service provider over the Internet via Wi-Fi, GPRS or 3G/4G cellular connections; the service provider then aggregates and processes the received data and, in return, provides real-time traffic information services for users to consume. These services can be in the form of browsing or querying information on the traffic situation, bus arrival times, bus crowdedness, estimated time to reach a destination, etc. Other examples include air pollution or noise mapping, flood or fire alerting, and so on.

In all such cases, a user plays a dual role as a data contributor as well as a service consumer. Thus, motivation can be provided to data contributors by exploiting their demand to consume useful and compelling information services. The basic principle is that the service provider grants each user a *service quota*, which determines how much service the user can consume, based on the amount of his demand and supply, i.e. contribution. This intrinsic motivating factor – demand for compelling services – provides an alternative incentive approach to other credit-based [6] and reputation-based [7] paradigms.

However, two main concerns are likely to arise from this service allocation approach. From an individual user's perspective, how fairly will each user be treated? From the system's point of view, how well will all the users as a whole be satisfied? To address these concerns, we propose two service allocation schemes with the objectives of maximizing fairness and maximizing social welfare, respectively. For the former, we develop an Allocation with Demand Fairness (ADF) scheme which achieves max–min fairness and scores close to 1 on the well-established Jain's fairness index. For the latter, we develop an optimization-based Iterative Tank Filling (ITF) scheme which maximizes social welfare and achieves a unique Nash equilibrium which, additionally, is Pareto and globally optimal. Moreover, to take into account practical considerations, we use chance constrained programming, a stochastic programming technique, to handle uncertainty in service demands which is often encountered in real-life situations. The performance of our schemes are evaluated via simulations and the results demonstrate the effectiveness of these schemes in meeting their respective objectives and confirm the theoretical results of our analysis.

2. Related work

Park and van der Schaar [8] proposed incentive provisioning using an intervention device that can take a variety of actions to influence users to cooperate and avoid inefficient resource usage under the assumption that the device can monitor a random access (e.g., CSMA) network perfectly. We focus on a different context, i.e. participatory sensing, and, related to the resource inefficiency aspect, we will show that the proposed ITF scheme achieves Pareto efficiency.

In a study related to incentive for participatory sensing, Lee and Hoh [6] proposed a dynamic pricing mechanism that allows users to sell their sensing data to a service provider. In order to keep the service provider's cost low while retaining an adequate number of participants, they proposed a reverse auction mechanism to

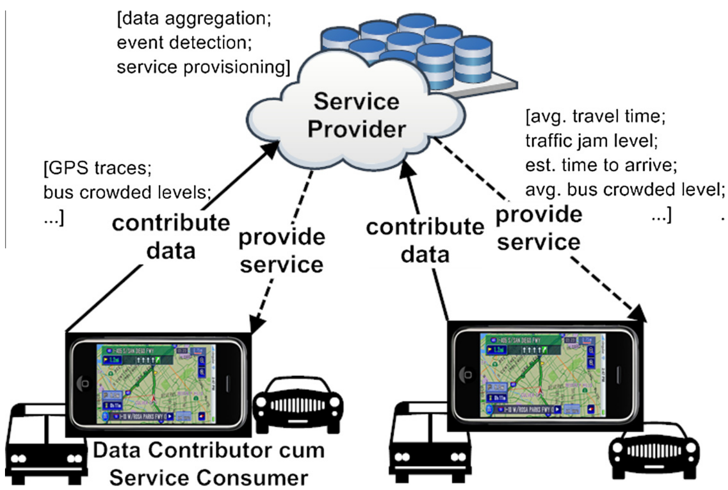

Fig. 1. An illustrative participatory sensing application: traffic monitoring.

keep the bid price competitive and use 'virtual participation credits' to retain participants. Unlike this one-sided monetary incentive which may not be sustainable, we leverage on the dual role of participants as both contributors and consumers to motivate contributors. Furthermore, monetary incentives usually present a service provider with the financial risk of not being able to resell to buyers (partially or completely), while the service allocation approach proposed in this paper avoids this risk. Another related work is a pilot study conducted in UCLA by Reddy et al. [9] who investigated the effect of micro-payments on participatory sensing. The study found that monetary incentive was beneficial when combined with altruism and competitiveness, and a major concern of the participants was fairness. However, they did not address the issue of fairness in a rigorous manner. In this paper, we propose a different motivation scheme and address the fairness issue in detail, both theoretically and via simulation.

3. System model

The system consists of N users and a service provider. Each user is a data contributor as well as a service consumer. The service provider receives user contributed data, processes it using techniques ranging from simple data fusion to complex data mining, and provides a value-added information service to the users.

We consider the system in terms of time slots, where in practice, a slot can range from an hour to a few days depending on the application. Pertaining to each time slot, a user $i \in [1, \ldots, N]$ is characterized by a quadruple $\langle \psi_i,\ c_i,\ Q_i,\ q_i \rangle$, where ψ_i is the user's contribution level, c_i is the cost incurred (e.g., mobile data charges and battery drainage), Q_i is the amount of service the user demands, and q_i ($q_i \leqslant Q_i$) is the service quota that the service provider grants to the user. A user can declare his demand Q_i anytime in the current slot, and the quota q_i will be granted at the end of the current slot after evaluating the contribution level ψ_i and is consumable in the next slot. Q_i and q_i are in units such as hours of service or number of queries, depending on the specific application. Methods for evaluating the user contribution ψ_i are also application dependent and will be discussed in Section 7.2.

The service provider determines the total amount of service quota, denoted by Q_{tot}, for each slot, which is to be allocated to individual users as q_i. The value of Q_{tot} is known to all the users, either as prior knowledge if it is fixed all the time, or via announcements if it varies from slot to slot. The quality of service (QoS) of the system, denoted by Ψ, is characterized by the aggregate contribution level of all the users and varies over time. For the sake of generality, our analysis is not coupled to any specific expression of Ψ as it is usually application dependent. One possible instantiation of Ψ is $\Psi = \sum_{i=1}^{N} \psi_i$, but other forms of Ψ in terms of ψ_i can be considered.

The problem can thus be stated as: to assign a certain amount Q_{tot} of service quota to N users according to their characterizing quadruple $\langle \psi_i,\ c_i,\ Q_i,\ q_i \rangle$, with the objective of maximizing fairness (Section 4) or social welfare (Section 5), while providing motivation to encourage user contributions.

4. Service allocation with fairness

From an individual user's perspective, one would expect to be treated "fairly" when consuming his desired service. We define *fairness* by the amount by which a user i's received service quota q_i is commensurate with both his contribution level ψ_i and his demand Q_i. This section presents a scheme called Allocation with Demand Fairness (ADF) to achieve this objective.

Firstly, let us consider a simple scheme which grants q_i based on ψ_i only. To do this, we gradually increase q_i for each user i at the differentiated rate of $\psi_i / \sum_{l=1}^{N} \psi_l$. Once a user's demand Q_i is reached, we exclude this user and proceed with the rest of the users in the same way, but with an updated Q_{tot}. This process can be mathematically described, by re-indexing the users by $j = 1, \ldots, N$ in ascending order of Q_i/ψ_i (which is the order in which the users will be fully satisfied), as:

$$q_j = \min\left\{ Q_j, \frac{\psi_j}{\sum_{k=j}^{N} \psi_k} \times \left(Q_{tot} - \sum_{k=1}^{j-1} q_k \right) \right\}. \tag{1}$$

Next, we extend the scheme by taking into account the user's demand. Naturally, it is fair to grant q_i such that q_i/Q_i is proportional to ψ_i when neither of the limits, Q_i or Q_{tot} is reached. One way to do this is by mimicking the case above, i.e. to increase each q_i/Q_i at the rate of $\psi_i / \sum_{l=1}^{N} \psi_l$ until 1 is reached. However, this does not readily lead to a mathematical or algorithmic abstraction because, unlike in Eq. (1) where $\sum_{i=1}^{N} q_i$ is capped by Q_{tot}, the upper bound to $\sum_{i=1}^{N} q_i/Q_i$ is not clear when $\sum_{i=1}^{N} Q_i > Q_{tot}$. To overcome this problem, we increase each q_i at the rate of $Q_i \psi_i / \sum_{l=1}^{N} Q_l \psi_l$ until Q_i is reached, whereby the user with the largest ψ_i will obtain the maximal q_i/Q_i first, which fulfils the objective. Thus, the scheme can be formulated below, by sorting the users in descending order of ψ_i and re-indexing them by $j = 1, \ldots, N$:

$$q_j = \min\left\{ Q_j, \frac{Q_j \psi_j}{\sum_{k=j}^{N} Q_k \psi_k} \left(Q_{tot} - \sum_{k=1}^{j-1} q_k \right) \right\}. \tag{2}$$

This is the ADF scheme which is presented algorithmically as Algorithm 1.

To analyze the properties of ADF, we consider two important and well-established fairness measures: Jain's fairness index [10] and max–min fairness. Jain's fairness index is defined as

$$J = \frac{\left(\sum_{i=1}^{N} x_i \right)^2}{N \sum_{i=1}^{N} x_i^2}$$

Algorithm 1. Allocation with Demand Fairness (ADF)

Input: $N, Q_{tot}, \vec{Q} = \{Q_i\}, \vec{\psi} = \{\psi_i\}$
Output: $\vec{q} = \{q_i\}$
if $\sum_{i=1}^{N} Q_i \leqslant Q_{tot}$ **then**
  return $\vec{q} \leftarrow \vec{Q}$
end if
create $\vec{I} = \{I_j\}_{j=1}^{N}, I_j \in [1, N]$, such that $\psi_{I_1}, \psi_{I_2}, \ldots \psi_{I_N}$ are in descending order
for $j = 1 \rightarrow N$ **do**
  $q_{I_j} = Q_{tot} \frac{Q_{I_j}\psi_{I_j}}{\sum_{k=j}^{N} Q_{I_k}\psi_{I_k}}$
  if $q_{I_j} > Q_{I_j}$ **then**
    $q_{I_j} \leftarrow Q_{I_j}$
  end if
  $Q_{tot} = Q_{tot} - q_{I_j}$
end for

where, in our context, $x_i \triangleq q_i/q_i^*$ in which q_i^* is the optimal (i.e., fairest) service quota to be granted to user i. The maximum of Jain's fairness index is 1, which is achieved when $x_i = x_j, \ \forall i, j$. In line with our objective of fairness, $q_i^* = Q_i\psi_i$ (ignoring a constant coefficient which does not affect the result). Thus,

$$J = \frac{\left(\sum_{i=1}^{N} \frac{q_i}{Q_i\psi_i}\right)^2}{N\sum_{i=1}^{N}\left(\frac{q_i}{Q_i\psi_i}\right)^2}. \quad (3)$$

Without loss of generality, suppose there are k $(0 \leqslant k \leqslant N)$ users who are fully satisfied, and all the users are sorted as in (2) and indexed by j. It is fairly straightforward to show that

$$q_j = \begin{cases} Q_j, & j = 1, \ldots, k \\ h \times Q_j\psi_j, & j = k+1, \ldots, N \end{cases}$$

where

$$h = \frac{Q_{tot} - \sum_{j=1}^{k} Q_j}{\sum_{j=k+1}^{N} Q_j\psi_j}$$

and k is determined by $1/\psi_k \leqslant h < 1/\psi_{k+1}$. Hence,

$$J = \frac{\left[\sum_{j=1}^{k}\psi_j^{-1} + (N-k)h\right]^2}{N\left[\sum_{j=1}^{k}\psi_j^{-2} + (N-k)h^2\right]}. \quad (4)$$

We will evaluate Eq. (4) for ADF and several other schemes in Section 6 through simulations. Here, we give two special cases that have specific analytical expressions:

- $k = 0 \Rightarrow J = 1$: in this case, the maximum fairness is achieved, and all users are equally satisfied. The expression for h is $h = q_i/(Q_i\psi_i) = Q_{tot}/\sum_{l=1}^{N} Q_l\psi_l$.
- $k = N \Rightarrow J = \left(\sum_{i=1}^{N}\psi_i^{-1}\right)^2/(N\sum_{i=1}^{N}\psi_i^{-2})$: in this case, all users are fully satisfied, and $J = 1$ if all the users contribute equally.

The result for the other fairness measure, max–min fairness, is given below.

Proposition 1. *The ADF scheme achieves weighted max–min fairness. That is, increasing user i's demand-normalized service quota, q_i/Q_i, weighted by $1/\psi_i$, viz. $q_i/(Q_i\psi_i)$, must be at the cost of decreasing some other user j's $q_j/(Q_j\psi_j)$, where $q_j/(Q_j\psi_j) < q_i/(Q_i\psi_i)$.*

We omit the proof as this proposition is self-evident from the above.

Remark. Our definition of fairness so far does not take cost, i.e. monetary cost per unit of data contribution, into account. In fact, to incorporate cost such that a user's service quota q_i should also be commensurate with his cost c_i, the steps in the analysis above only need to be modified by plugging c_i into Eq. (3) to obtain

$$J = \frac{\left(\sum_{i=1}^{N} \frac{q_i}{c_iQ_i\psi_i}\right)^2}{N\sum_{i=1}^{N}\left(\frac{q_i}{c_iQ_i\psi_i}\right)^2}$$

which is the new Jain's index to maximize, and into Eq. (2) to obtain

$$q_j = \min\left\{Q_j, \frac{c_jQ_j\psi_j}{\sum_{k=j}^{N} c_kQ_k\psi_k}\left(Q_{tot} - \sum_{k=1}^{j-1} q_k\right)\right\}$$

which is the new ADF solution. However, adopting this scheme in participatory sensing is problematic because, unlike many other contexts where cost is a monotonically increasing function of the output (or contribution level, in the context of participatory sensing), the cost profile in participatory sensing is unlikely to be a monotonically increasing function of the contribution level since producing the same contribution level can incur drastically different costs to different users due to their different contexts and behaviors. If cost is taken into account in this manner to address fairness, a user who incurs disproportionately high cost while producing merely ordinary contribution will be awarded more service quota than the original ADF scheme, which would actually make it *unfair*. On the other hand, cost does affect a user's *satisfaction*, which we shall handle using utility in the next section.

As a side note, user demands are assumed to be honestly declared, i.e., according to users' true needs. Handling untruthful declarations is outside the scope of this paper. In fact, the ITF scheme which will be presented in the next section has the effect of discouraging users from making false declarations of their demands, as will be explained in Section 5.1.

5. Service allocation with social welfare maximization

From a system perspective, the service provider would aim to maximize social welfare, which is defined as the aggregate user utility (which is a measure of satisfaction) received from consuming the service provided by the system, minus his incurred cost, such as the monetary cost per unit of data contribution. At the same time, the system aims to motivate users to contribute at higher levels. Hence, the objective is formulated as maximizing

$$S \triangleq \sum_{i=1}^{N}\psi_i u_i$$

which is the aggregate contribution-weighted user utility, where u_i is user i's utility. The structure of this objective function implies that priority will be given to users with larger ψ_i.

We define the utility u_i as two possible forms:

$$\text{(a) } u_i = U\left(\Psi \frac{q_i}{c_i Q_i}\right), \quad \text{(b) } u_i = U\left(\Psi \frac{q_i}{Q_i}\right) \Big/ c_i, \tag{5}$$

where q_i/Q_i reflects the extent a user's demand is satisfied,[2] and Ψ as the system QoS which will also affect a user's utility. The forms (a) and (b) differ in where the normalizing effect of the cost takes place and we will investigate both.[3] $U(x): \mathbb{R}^+ \to \mathbb{R}$ is a monotonically increasing and strictly concave function of x, which reflects the *elasticity* of user satisfaction as is common in the literature (e.g., [11]). In this paper, we consider the form of $U(x) = \log(1+x), x \geqslant 0$, which is also used in [11–13], and the problem is thus formulated below:

$$\text{maximize : (a) } S = \sum_{i=1}^{N} \psi_i \log\left(1 + \Psi \frac{q_i}{c_i Q_i}\right), \text{or} \tag{6}$$

$$\text{maximize : (b) } S = \sum_{i=1}^{N} \psi_i \log\left(1 + \Psi \frac{q_i}{Q_i}\right) / c_i, \tag{7}$$

$$\text{s.t. } q_i \in [0, Q_i], \quad \forall i = 1, \ldots, N \tag{8}$$

$$\sum_{i=1}^{N} q_i \leqslant Q_{tot}. \tag{9}$$

In this section, we develop a scheme to solve problem (a) while leaving problem (b) to Appendix A for interested readers, since both problems follow the same line of reasoning and (b) turns out to be simpler than (a). Now, let us consider Eq. (6). In order to maximize S, the solution should give priority to users with larger *marginal* weighted utility, i.e., larger $\psi_i \Psi/(c_i Q_i + q_i \Psi)$,[4] or equivalently, smaller $(\frac{c_i Q_i}{\Psi} + q_i)/\psi_i$. Based on this insight, we convert the original nonlinear programming problem into the problem of "filling iced tanks" depicted in Fig. 2. Each user i is represented by a tank with bottom area ψ_i, and the tank has been preoccupied by frozen ice of volume $c_i Q_i/\Psi$ (and hence of height $\frac{c_i Q_i}{\Psi}/\psi_i$). Tank i is left with an empty space of volume Q_i (and hence of height Q_i/ψ_i) to be filled with water. All the tanks are placed back to back as if they are virtually connected without internal separators. Consequently, the empty space will be filled consecutively in the order of ①, ②, ③,... as shown in Fig. 2. To solve the problem, we propose an *Iterative Tank Filling* (ITF) algorithm which iteratively fills the space in the depicted order until all the tanks are fully filled or the total volume of water, Q_{tot}, is used up.[5] The pseudo-code is given as Algorithm 2.

[2] Whether a user can benefit by simply (and unilaterally) declaring too high or too low Q_i will be discussed in Section 5.1.

[3] What constitutes cost and how it is calculated will be elaborated in Section 7.3.

[4] The optimization variables are q_i; Ψ is not subject to the differentiation with respect to q_i because it has been determined by all the ψ_i's and is not a function of q_i being solved.

[5] This iced-tank filling problem is different from the water filling (WF) problem in convex optimization [14] or wireless communications [15] in that: (i) these tanks can have different water levels during and after filling, because each tank comes with a closed "lid", whereas WF fills a *single* and *open* vessel with one sweeping water level, and (ii) WF will always fully allocate the total resource (e.g., power), which is not the case in ITF.

Proposition 2. *The computational complexity of ITF is $O(N^2)$.*

Proof. In the worst case, ice_i and $tank_i$ are all different (i.e., $2N$ distinct numbers), and hence each iteration will increase the highest water level to only one of these $2N$ numbers. Therefore, The main loop will execute at most $2N - 1$ times. Inside the main loop, lines 7–10, 20, and 25 each has a complexity of $O(N)$. The proposition is thus proven.

Note that ITF does not (and should not) assume that the tanks have already been sorted in any order, which would otherwise decrease the computational complexity of ITF.

Algorithm 2. Iterative Tank Filling (ITF)

Input: $N,\ Q_{tot},\ \Psi,\ \vec{Q} = \{Q_i\},\ \vec{\psi} = \{\psi_i\},\ \vec{c} = \{c_i\}$
Output: $\vec{q} = \{q_i\}$
1: **if** $\sum_{i=1}^{N} Q_i \leqslant Q_{tot}$ **then**
2: **return** $\vec{q} \leftarrow \vec{Q}$
3: **end if**
4: $\vec{q} \leftarrow \vec{0}$; ***ice*** $\leftarrow \{ice_i = c_i Q_i/(\psi_i \Psi)\}$; ***tank*** $\leftarrow \{tank_i = ice_i + Q_i/\psi_i\}$
5: **while** $Q_{tot} > 0$ **do**
6: ––– Find space to fill in this iteration –––
7: $bot \leftarrow \min_i\{ice_i\}$; $botind \leftarrow \arg\min_i\{ice_i\}$
8: $w \leftarrow \sum_{i \in botind} \psi_i$ //bottom area
9: $cap_1 \leftarrow \min_{i \notin botind}\{ice_i\}$
10: $cap_2 \leftarrow \min_i\{tank_i\}$
11: $h \leftarrow \min\{cap_1, cap_2\} - bot$ //height
12: ––––––––- Fill the space ––––––––-
13: **if** $w \cdot h < Q_{tot}$ **then**
14: $Q_{tot} = Q_{tot} - w \cdot h$
15: **else** {the last iteration of filling}
16: $h \leftarrow Q_{tot}/w$ //readjust height
17: $Q_{tot} \leftarrow 0$
18: **end if**
19: **for all** $i \in botind$ **do**
20: $ice_i = ice_i + h$; $q_i = q_i + h \cdot \psi_i$
21: **end for**
22: –––––– Remove full tanks ––––––
23: **if** $cap_2 \leqslant cap_1$ or $cap_1 = \infty$ **then**
24: **for all** $k \in \{i | tank_i = cap_2\}$ **do**
25: $tank_k \leftarrow \infty$; $ice_k \leftarrow \infty$
26: **end for**
27: **end if**
28: **end while**

Theorem 1. *Service provisioning via ITF ensures that, for any i, j, if $\frac{\psi_i}{c_i Q_i} \geqslant \frac{\psi_j}{c_j Q_j} \wedge c_i \leqslant c_j$, then $u_i \geqslant u_j$.*

Proof. Consider two cases of the output $\vec{q}$:

1. $q_i/\psi_i \geqslant q_j/\psi_j$ (Fig. 3a). Multiplying this with $\frac{\psi_i}{c_i Q_i} \geqslant \frac{\psi_j}{c_j Q_j}$ gets $\frac{q_i}{c_i Q_i} \geqslant \frac{q_j}{c_j Q_j}$. Hence $U(\Psi \frac{q_i}{c_i Q_i}) \geqslant U(\Psi \frac{q_j}{c_j Q_j})$ (for any non-decreasing function $U(\cdot)$), i.e., $u_i \geqslant u_j$.
2. $q_i/\psi_i < q_j/\psi_j$. Since $\frac{\psi_i}{c_i Q_i} \geqslant \frac{\psi_j}{c_j Q_j} \iff \frac{c_i Q_i}{\Psi}/\psi_i \leqslant \frac{c_j Q_j}{\Psi}/\psi_j$, meaning that the reciprocal of marginal weighted utility or the ice level of i is lower than that of j, priority will

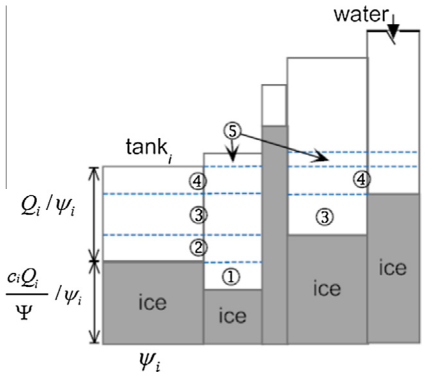

Fig. 2. Filling iced tanks. The tanks are sorted in ascending order of height, but the ice (gray area) levels are not necessarily in order. Numbers ①, ②, ③,... denote the order of empty spaces to be filled, which also corresponds to the iteration number of the ITF algorithm.

be given to i (ITF will start filling tank i earlier than j). However, since the outcome is $q_i/\psi_i < q_j/\psi_j$, it implies that tank i must have been fully filled and the height of original empty space $Q_i/\psi_i < Q_j/\psi_j$ is as shown in Fig. 3b. As $c_i \leqslant c_j$, we have $1/c_i = \frac{q_i}{c_iQ_i} \geq 1/c_j \geqslant \frac{q_j}{c_jQ_j} \Rightarrow \log(1 + \frac{\Psi}{c_i}) \geqslant \log(1 + \Psi\frac{q_j}{c_jQ_j}) \Rightarrow u_i \geqslant u_j$. □

Corollary 1. *Service provisioning via ITF ensures that, for any i, j, if $\frac{\psi_i}{Q_i} \geqslant \frac{\psi_j}{Q_j} \wedge c_i \leqslant c_j$, then $u_i \geqslant u_j$.*

The implication of Corollary 1, as a relaxed form of Theorem 1, is that a user who makes higher contribution with respect to his demand and incurs lower cost, will be guaranteed to receive higher utility.

Theorem 2. *The output of ITF is the optimal solution to problem (6). In other words, the solution given by ITF maximizes the social welfare.*

Proof. The Lagrangian function for problem (6) is

$$L(\vec{q}, \vec{\lambda}) = \sum_{i=1}^{N}\psi_i \log\left(1 + \Psi\frac{q_i}{c_iQ_i}\right) + \sum_{i=1}^{N}\lambda_i(Q_i - q_i) + \lambda_0\left(Q_{tot} - \sum_{i=1}^{N}q_i\right)$$

where λ_i's are the Lagrangian multipliers associated with the respective constraints. Thus, the Kuhn–Tucker conditions are formulated as

$$\frac{\partial L}{\partial q_i} = \frac{\psi_i}{\frac{c_iQ_i}{\Psi} + q_i} - \lambda_i - \lambda_0 \leqslant 0, \quad q_i\left(\frac{\psi_i}{\frac{c_iQ_i}{\Psi} + q_i} - \lambda_i - \lambda_0\right) = 0, \tag{10}$$

$$q_i - Q_i \leqslant 0, \quad \lambda_i(q_i - Q_i) = 0, \quad i = 1, 2, \ldots \tag{11}$$

$$\sum_{i=1}^{N}q_i - Q_{tot} \leqslant 0, \tag{12}$$

$$\lambda_0\left(\sum_{i=1}^{N}q_i - Q_{tot}\right) = 0, \quad \lambda_i \geqslant 0, \quad i = 0, 1, 2, \ldots \tag{13}$$

In the simplest case where $\sum_{i=1}^{N}Q_i \leqslant Q_{tot}$, ITF will output $q_i = Q_i$. This conforms to the Kuhn–Tucker conditions by letting $\lambda_0 = 0$ (i.e. unbound constraint) due to (12) and $\lambda_i = \psi_i/(\frac{c_iQ_i}{\Psi} + Q_i)$ (i.e. bound) due to (11). In the other case where $\sum_{i=1}^{N}Q_i > Q_{tot}$, ITF will fill the tanks until the water level reaches $l_0 = (\frac{c_iQ_i}{\Psi} + q_i)/\psi_i$ (cf. Fig. 2) for any i that satisfies $0 < q_i < Q_i$ (half-filled). For these half-filled tanks, $\lambda_i = 0$ due to (11), and hence $\lambda_0 = 1/l_0$ due to (10). For the full tanks ($q_j = Q_j$), their tank heights must be lower than l_0 according to ITF, i.e., $(\frac{c_jQ_j}{\Psi} + Q_j)/\psi_j \leqslant l$ (cf. Fig. 2). This fulfills (10) which requires $\lambda_j = \psi_j/(\frac{c_jQ_j}{\Psi} + Q_j) - \lambda_0 \geqslant 0$ (since $\lambda_0 = 1/l_0$). Finally, for the empty tanks ($q_k = 0$), their ice levels must be higher than l_0, i.e., $\frac{c_kQ_k}{\Psi}/\psi_k \geqslant l_0$. As $\lambda_k = 0$ (due to (11)), these tanks still comply with (10) which requires $\lambda_0 \geqslant \psi_k/\frac{c_kQ_k}{\Psi}$ (recall $\lambda_0 = 1/l_0$).

This proves that the solution given by ITF satisfies the Kuhn–Tucker conditions. Furthermore, because the objective function (6) is strictly concave and all the constraints (8) and (9) form a convex polyhedron, the optimal solution that satisfies the Kuhn–Tucker conditions is unique. The theorem is thus proven. □

Expanding on the proposed ITF service allocation approach, we now address two other major issues in the following sub-sections.

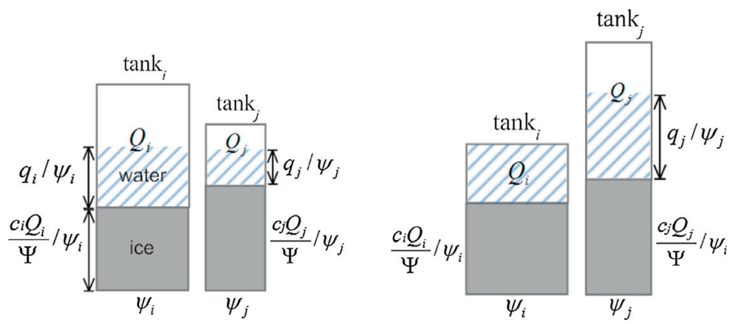

(a) Case 1. (b) Case 2.

Fig. 3. Proof of Theorem 1.

5.1. Deriving optimal service demands

One issue is to derive the optimal service demands, Q_i, that users declare. As each user's service quota is capped by his *declared* demand, a user may be tempted to declare a higher demand in order to, possibly, get a larger share of service quota. On the other hand, as a service allocation scheme should be *transparent* to users, a user can realize from ITF that declaring a higher demand would be disadvantageous to him since he will be classified as a "hard-to-satisfy" user and given a lower priority to be granted service. Note that this has the effect of discouraging users from misbehaving in their demand declaration, as mentioned in Section 4. Therefore, there should exist an *optimal* service demand for each user.

There are two ways to define the optimality: (i) global optimality – the objective function (6) achieves the maximum over the entire domain of optimization variables; (ii) Pareto optimality – no user's utility can be improved without making some other user's utility worse off. These two kinds of optimality are *not* achieved simultaneously, in general.

In addition, it is also desirable to make the optimal point "stable": any user should not have an incentive to deviate from his optimal demand unilaterally, i.e. when other users stick to their optimal demands.

Under these circumstances, a game-theoretic approach is appropriate. In the sequel, we derive a solution and show that it achieves all of the aforementioned properties: global maximal, Pareto optimal, and Nash Equilibrium.

We model the participatory sensing problem as a non-cooperative game [16]. The game players are the N users. Each player's strategy is to decide how much demand, i.e., Q_i, to declare and his strategy space is $\mathbb{R}^+$. Each user's payoff is his utility u_i. The game rule is prescribed by ITF, which maps an N-tuple of user strategies $\overrightarrow{Q} = \{Q_i\} \in (\mathbb{R}^+)^N$ to an N-tuple of user payoffs $\vec{u} = \{u_i\}$, by determining the service quota $\vec{q} = \{q_i\}$.

A common game-theoretic approach is to find a strategy profile, prove it to be a Nash Equilibrium (NE), and subsequently prove uniqueness, if possible. In this paper, we take a different and simpler approach: we first find the necessary and sufficient conditions for a NE, and then derive the NE and prove its uniqueness in one go.

Lemma 1. *The necessary and sufficient conditions that a Nash Equilibrium of the above-defined game satisfies are*

$$\begin{cases} \text{C1}: & \sum_{i=1}^{N} Q_i = Q_{tot} \\ \text{C2}: & h_i = h_j, \ \forall i, \ j = 1, \ldots, N \end{cases} \tag{14}$$

where $h_i = Q_i(1 + c_i/\Psi)/\psi_i$.

Proof. *Necessity*: Prove by contradiction.

Condition C1: Suppose, instead, $\sum_{i=1}^{N} Q_i < Q_{tot}$, then $q_i = Q_i, \ \forall i$. Obviously, a user can increase his Q_i to $Q_i' > Q_i$ and be granted $q_i' > q_i$ provided other users do not change strategy. If, otherwise, $\sum_{i=1}^{N} Q_i > Q_{tot}$, there is at least one tank that is not fully filled. Let k be one such tank and $b_k = c_k Q_k/(\Psi\psi_k)$ denote its ice level. Let h_i be tank i's ice+water level and $h_{max} = \max_i\{h_i | q_i > 0\}$. In the case that $b_k \geqslant h_{max}$ (i.e., $q_k = 0$), clearly user k can decrease demand Q_k such that b_k drops to $b_k' < h_{max}$, and be granted $q_k' > 0$. In the other case that $b_k < h_{max}$ (i.e., $0 < q_k < Q_k$), user k can decrease Q_k to Q_k' such that $q_k < Q_k' < Q_k$ and, accordingly, ice level b_k drops to $b_k' = c_k Q_k'/(\Psi\psi_k)$, and be granted: (i) $q_k' > q_k$ if $\exists j : q_k/\psi_k + b_k' < h_j \leqslant h_{max}$, where the left hand side $(q_k/\psi_k + b_k')$ is k's ice+water level as if his water level remains unchanged, or (ii) $q_k' = q_k$ otherwise (such a user j does not exist; i.e., k is the only partially-filled tank ($0 < q_k < Q_k$) and the rest of the tanks are either fully filled (with their $h_j \leqslant q_k/\psi_k + b_k'$) or empty (with their $h_j > h_{max}$)). In summary, user k will *have* an incentive to deviate from his strategy if $\sum_{i=1}^{N} Q_i \neq Q_{tot}$. Therefore, $\sum_{i=1}^{N} Q_i = Q_{tot}$ must hold.

Condition C2: Suppose $\exists i, j : h_i \neq h_j$, and WLOG, $h_i < h_j$. Since $\sum_{i=1}^{N} Q_i = Q_{tot}$, all the users are fully satisfied. Recall that h_i and h_j are the ice + water level of users i and j, respectively. If user i increases his demand (slightly) to Q_i' such that $h_i' = Q_i'/\psi_i + c_i Q_i'/(\Psi\psi_i) < h_j$ still holds, then, according to the ITF game rule, user i will be granted $q_i' > q_i = Q_i$ where the additional quota essentially comes from user j (and others, if any). This means that user i will have an incentive to change his strategy unilaterally. Hence, Condition 2 must also hold.

Sufficiency:

If both Conditions C1 and C2 are satisfied, all the tanks have the same height and are fully filled. Suppose any user, say i, changes his strategy such that: (1) $Q_i' > Q_i$, then i's ice level will increase, which actually lowers i's priority to receive service. On the other hand, all the other tanks are fully filled. Hence, tank i will continue to have the same volume, Q_i, of water (though its ice + water level will be above the other tanks) with an empty space of $Q_i' - Q_i$ left in the tank; (2) $Q_i' < Q_i$, then obviously he will receive a *lower* quota of $q_i' = Q_i'$. In summary, user i will either be indifferent (in case 1) or unwilling (in case 2) to switch his strategy. Hence, a strategy profile satisfying both C1 and C2 is a Nash Equilibrium (NE). □

Theorem 3. *The optimal strategy profile* $\mathbf{Q}^* = \{Q_i^*\}$ *where*

$$Q_i^* = \frac{\psi_i/(\Psi + c_i)}{\sum_{l=1}^{N} \frac{\psi_l}{\Psi + c_l}} Q_{tot} \tag{15}$$

is a unique Pareto-efficient Nash equilibrium, and achieves the global optimum.

Proof. It can be shown that the equation system (14) can be translated into a $N \times N$ homogeneous system of linear equations whose determinant is non-zero. Hence, this system has a unique solution which leads to Eq. (15).

The Pareto efficiency follows from condition C1 of Lemma 1.

Under the NE strategy (15), $q_i = Q_i$ and each user receives the maximum utility $u_i^{max} = \log(1 + \Psi/c_i)$ for given $\vec{\psi}$ and $\vec{c}$. This achieves the global maximum of (6) *term-wise*, which is a sufficient condition for (6) to achieve its global maximum:

$$S_a^{max} = \sum_{i=1}^{N} \psi_i \log(1 + \Psi/c_i). \quad \square \tag{16}$$

5.2. Handling uncertainty in service demands

The other issue is to handle the uncertainty in service demands which is often encountered in reality. As declared demands are essentially future demands (Q_i is the amount of service a user plans to consume in the next slot), a user needs to *estimate* his future actual demand. Denote the (unknown) actual demand in the next slot by $\widetilde{Q}_i$, which is a random variable, and the estimated demand by $\widehat{Q}_i$. Thus, the previously discussed Q_i is actually $\widehat{Q}_i$, and $\widehat{Q}_i$ is the "expected value" of $\widetilde{Q}_i$. To reformulate the problem of taking the actual demand $\widetilde{Q}_i$ into account, it is improper to simply replace the original constraint $q_i \leqslant Q_i$ with $q_i \leqslant \widetilde{Q}_i$ which essentially leads to $q_i \leqslant \inf\{\widetilde{Q}_i\} = 0$. It is also improper to replace $q_i \leqslant Q_i$ with $q_i \leqslant \widehat{Q}_i$ as will be elaborated below, assuming $\widetilde{Q}_i \sim \mathcal{N}(\widehat{Q}_i, \sigma_i)$.

5.2.1. Expected-value method

It is tempting to use the constraint $q_i \leqslant \widehat{Q}_i$ which simply converts the problem with uncertainty into the original deterministic optimization problem whose solution is already given by ITF. This is called an expected-value method (EVM) as it uses $\mathbb{E}(\widetilde{Q}_i) = \widehat{Q}_i$ to simplify the constraints.

To examine whether EVM is suitable for our particular problem, we conducted a simulation study for a daily-slotted system with $N = 100$ users. The demand Q_i and allocated service quota q_i are formulated as the duration of service required by user *i*. For ease of description, denote by $\mathcal{U}(a, b)$ the uniform distribution in interval (a, b), and by $\mathcal{N}_{tr}(\mu, \sigma, a, b)$ the truncated normal distribution with mean μ and standard deviation σ and bounded in the range of $[a, b]$. The simulation was set up as follows: $\widehat{Q}_i = \mathcal{U}(0, 1)$ (day), $\widetilde{Q}_i = \mathcal{N}_{tr}(\widehat{Q}_i, 0.25\widehat{Q}_i, 0, 1)$, contribution $\psi_i = \mathcal{U}(0, 1)$ (Mb), and cost $c_i = \mathcal{N}_{tr}(\psi_i, \psi_i, 0, 3)$ (i.e., the expected cost is one monetary unit per Mb of contribution), $Q_{tot} = 0.75\sum_{i=1}^{N}\widehat{Q}_i$ and $\Psi = 10^{-3}\sum_{i=1}^{N}\psi_i$ (Gb).

According to EVM, we simply replace the original Q_i in ITF with $\widehat{Q}_i$ to run the algorithm. The results are shown in Fig. 4a, where 45 out of 100 users were found to have exceeded their actual demands $\widetilde{Q}_i$, which results in significant resource wastage since a user would not consume more than his actual demand. Zooming in on users 1–10 as shown in Fig. 4 b, we see that $q_i/\widetilde{Q}_i$ can be as high as 212%.

These observations indicate that EVM can cause significant *service over-provisioning* and, perhaps more importantly, the service provider has *no control* over such over-provisioning. The consequence is that users will be discouraged from contributing because any user is likely to be granted a large share of service *without* the commensurate contribution, which works strongly against any incentive scheme.

5.2.2. Chance constrained programming

Now that we have seen that EVM is not suitable for our particular problem when there is uncertainty in the service demands, we use the *chance constrained programming* (CCP) [17] technique to tackle it.

First, we need to reformulate the problem with uncertainty more rigorously, for which we introduce the following probabilistic constraints:

$$\Pr(q_i \leqslant \widetilde{Q}_i) \geqslant 1 - \alpha_i, \quad \forall i = 1, \ldots, N, \tag{17}$$

where α_i's are prescribed probabilities. Each of these N constraints means that q_i is capped by (all the realizations of) $\widetilde{Q}_i$ in $1 - \alpha_i$ of the time, or alternatively, q_i has a chance of α_i to exceed $\widetilde{Q}_i$.

In CCP terms, the inequality (17) imposes *individual chance constraints* on the objective function, where "individual" relates to the fact that each stochastic constraint $q_i < \widetilde{Q}_i$ is transformed into a chance constraint individually. A variant is called *joint chance constraints* which, however, does not capture our problem as well as (17). Nonetheless, we present it in Appendix B for interested readers.

Thus, the original problem is reformulated as:

$$\begin{aligned}
&\text{maximize } S = \sum_{i=1}^{N} \psi_i \log\left(1 + \Psi \frac{q_i}{c_i \widehat{Q}_i}\right), \\
&\text{s.t. } \Pr(q_i \leqslant \widetilde{Q}_i) \geqslant 1 - \alpha_i, \quad \forall i = 1, \ldots, N, \\
&q_i \geqslant 0, \quad \forall i = 1, \ldots, N, \\
&\sum_{i=1}^{N} q_i \leqslant Q_{tot}.
\end{aligned} \tag{18}$$

The objective function uses $\widehat{Q}_i$ instead of $\widetilde{Q}_i$ (which is a random variable), because a user's utility is determined when he is granted q_i based on his declaration $\widehat{Q}_i$, instead of $\widetilde{Q}_i$. We also note that, with the introduction of the chance constraints (17), the service provider has one more lever to encourage a higher level of user contributions by associating α_i with user contribution, e.g., as $\alpha_i = \alpha_0\psi_i/\sum_{l=1}^{N}\psi_l$ where α_0 is a scaling factor.

Problem (18) is a stochastic programming problem and we solve it using CCP as follows. Denote the CDF of $\widetilde{Q}_i$ by $F_i(\cdot)$, and thus

$$\Pr(q_i \leqslant \widetilde{Q}_i) \geqslant 1 - \alpha_i \iff F_i(q_i) \leqslant \alpha_i.$$

Denote by $\gamma_i(p)$ the *quantile function* of $\widetilde{Q}_i$, which is defined as

$$\gamma_i(p) \triangleq \inf\{\tau | F_i(\tau) \geqslant p\}.$$

Since $F_i(\cdot)$ is a monotonically increasing function, it follows that

$$F_i(q_i) \leqslant \alpha_i \iff q_i \leqslant \gamma_i(\alpha_i). \tag{19}$$

Furthermore, $\widetilde{Q}_i \sim \mathcal{N}(\widehat{Q}_i, \sigma_i)$ as assumed by EVM, we use the *probit function* to solve for $\gamma_i(\alpha_i)$. The probit function is the quantile function for the *standard* normal distribution and can be calculated via straightforward numerical

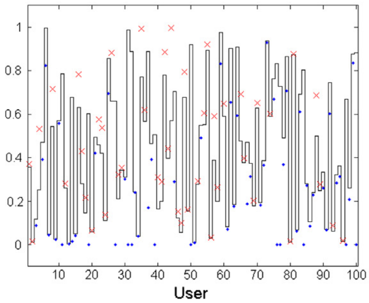

(a) EVM. 45 cases of over-provisioning.

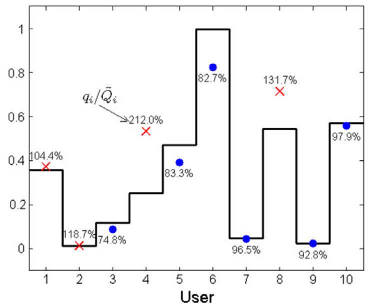

(b) EVM. Zoomed-in view of users 1–10.

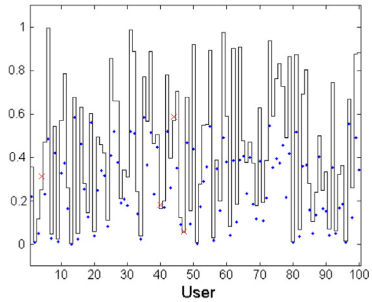

(c) CCP. 4 cases of over-provisioning.

Fig. 4. Solving the problem with uncertainty (i.e. the chance constrained problem (18)) using EVM and CCP, respectively. $N = 100$ users. The staircase line represents a realization of all the $\widetilde{Q}_i$'s. Blue dots are cases where $q_i \leqslant \widetilde{Q}_i$ and red crosses are $q_i > \widetilde{Q}_i$ (over-provisioning). The vertical axis stands for $\widetilde{Q}_i$ and q_i. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

computation or simple table look-up. Denote by z_α the α-quantile of the standard normal distribution. As $(\widetilde{Q}_i - \widehat{Q}_i)/\sigma_i \sim \mathcal{N}(0,1)$, (19) can be transformed into

$$\frac{q_i - \widehat{Q}_i}{\sigma_i} \leqslant z_{\alpha_i} \iff q_i \leqslant \widehat{Q}_i + \sigma_i z_{\alpha_i}, \tag{20}$$

where z_{α_i} can be obtained via numerical computation or standard table look-up. For instance, $z_{0.05} = -1.65$, $z_{0.025} = -1.96$.

Thus, the chance constraints (17) are converted into deterministic constraints (20), thereby allowing us to develop a solution scheme, which we call *ITF-CCP*, by modifying Algorithm 2 in the following way:

- Input: replace $\overrightarrow{Q}$ with $\overrightarrow{\widehat{Q}} = \{\widehat{Q}_i\}$ and add $\vec{\sigma} = \{\sigma_i\}$.
- Line 4: replace ice_i and $tank_i$ with $ice_i = c_i\widehat{Q}_i/(\psi_i\Psi)$ and $tank_i = c_i\widehat{Q}_i/(\psi_i\Psi) + (\widehat{Q}_i + \sigma_i z_{\alpha_i})/\psi_i$, respectively.

Now we run ITF-CCP with $\alpha_i = 0.05$ for the same 100 users as in EVM (also with the same realization of $\widetilde{Q}_i$, for a fair comparison). This new set of results is shown in Fig. 4c where it can be seen that there are only 4 cases of service over-provisioning, which is consistent with the "exceeding" probability α_i (0.05). This shows that the occurrences of over-provisioning are now under control.

A side effect is that, as $\widehat{Q}'_i \triangleq \widehat{Q}_i + \sigma_i z_{\alpha_i} < \widehat{Q}_i$, there will be extra resources left when $Q_{ext} \triangleq Q_{tot} - \sum_{i=1}^N \widehat{Q}'_i > 0$. To overcome this, we can allow q_i to "burst" above $\widehat{Q}'_i$ when $Q_{ext} > 0$, but still cap q_i by the actual demand $\widetilde{Q}_i$ in order to avoid over-provisioning. As $\widetilde{Q}_i$ is only realized during service consumption, we allocate Q_{ext} *after* a user has consumed his granted quota in the subsequent slot (since q_i is granted at the end of the current slot), based on the first-come-first-served (FCFS) principle. Note that: (i) user motivation to contribute is not compromised because obtaining service via a burst is non-guaranteed (opportunistic) as it depends on the availability of Q_{ext} and other users' service consumption, unlike the guaranteed service quota granted by ITF-CCP, and (ii) FCFS does not lead to each user rushing to use up his granted quota in order to take advantage of the burst, because a user will *not* know the availability of Q_{ext} until he uses up his granted service quota.

6. Performance evaluation

In this section, we evaluate the performance of the four proposed schemes, ADF, ITF, NE and ITF-CCP, via simulation. For a more meaningful comparison, we also add two baseline schemes:

1. Equal Allocation (EA): all the users share the total service quota equally, i.e., $q_i = Q_{tot}/N$.
2. Demand-based Allocation (DA): each user is granted a service quota of $q_i = Q_{tot} \times Q_i/\sum_{l=1}^N Q_l$ when $\sum_{l=1}^N Q_l > Q_{tot}$, and $q_i = Q_i$ when $\sum_{l=1}^N Q_l \leqslant Q_{tot}$.

Similar to Section 5.2, the system is daily-slotted with N=100 users. $Q_i = \mathcal{U}(0,1)$ (day), $Q_{tot} = \mathcal{U}(0.5,1) \times \sum_{i=1}^N Q_i$, and the rest of the simulation setup remains the same. In the case of NE, Q_i^* is computed according to (15). In the case of ITF-CCP, if $Q_{ext} > 0$, users burst as follows: let $t_i \in \mathcal{U}(q_i, 1)$ be the time when a user uses up his granted service quota q_i, upon which he realizes via server notification or by querying the server that there is extra quota, and hence will try to maximize his own benefit by continuously consuming service until he reaches his actual demand $\widetilde{Q}_i$ or Q_{ext} is used up. This is equivalent to running a further round of ITF by setting $Q_{tot} = Q_{ext}$, $ice_i = t_i$, $tank_i = t_i + (\widetilde{Q}_i - t_i)^+$, and $\psi_i = 1$.

6.1. Macro-level performance

This sub-section evaluates performance at the system level, in terms of Jain's fairness index as defined by Eq. (3), and social welfare as defined by Eq. (6).

Fig. 5 presents the results. Each data point is averaged over 100 rounds of simulation, and we also plot upper and lower 95% confidence limits around the sample means.

In the fairness aspect, Fig. 5a clearly shows that the ADF scheme outperforms the other schemes and closely approaches the maximum of Jain's index, 1, with a score of 0.92. As for social welfare (Fig. 5b), the first observation, which is not surprising, is that NE is the clear winner, as is theoretically proven by Theorem 3. On the other hand, the Q_i^* computed by NE may not necessarily reflect the users'

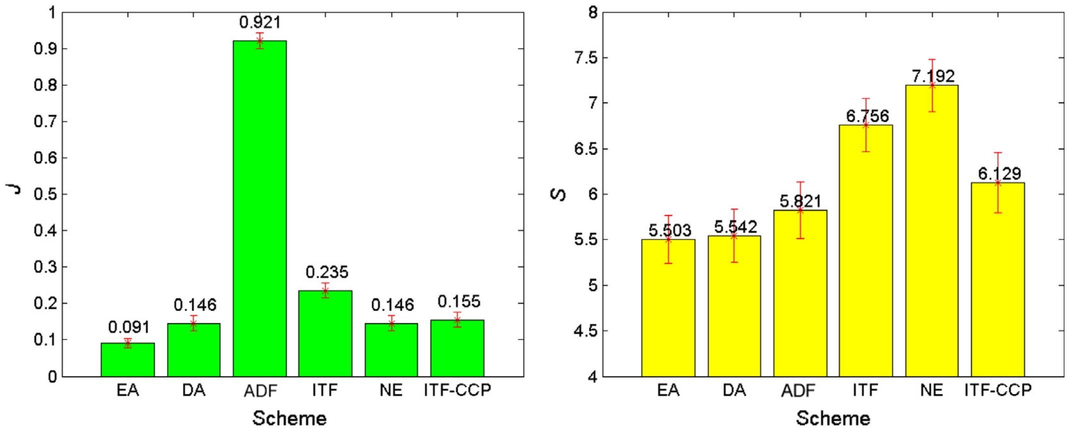

(a) Jain's fairness index. (b) Social welfare.

Fig. 5. Performance comparison between 6 schemes: EA, DA, ADF, ITF, NE and ITF-CCP. The sample size for each data point is 100. Error-bars represent 95% confidence intervals.

Table 1
"Unwelcome" users vs. "normal" users.

User	Type	Demand (day)	Contribution (Mb)	Cost ($/Mb)
1	High-demand	1	0.5	0.5
2	Low-contrib.	0.5	0.25	0.5
3	High-cost	0.5	0.5	1
4	Normal	0.5	0.5	0.5

real needs, and hence other schemes which allow the users to declare their demands should still be considered. In that case, ITF is the best scheme and achieves 94% of the maximum that NE achieves. In the case of coping with uncertainty, ITF-CCP achieves 85.2% of the NE maximum. The reason for the slight drop is that ITF-CCP takes stricter constraints to avoid over-provisioning when demands are uncertain. Bursting, as an auxiliary mechanism, only helps marginally, because (i) it only takes effect when $Q_{ext} > 0$, which is a rare case because Q_{tot} is usually well below $\sum_{i=1}^{N} Q_i$, (ii) the maximum burst amount for each user is capped by a limited amount of $(\widetilde{Q}_i - t_i)^+$, (iii) FCFS is not optimized for maximizing social welfare (e.g., no priority is given to users with larger marginal utility). These are the trade-offs a service provider should take into consideration when dealing with uncertainty.

Moreover, this set of results indicate the non-existence of a panacea solution which performs the best for both objectives simultaneously. Therefore, a service provider should first decide on its targeted objective for the specific application, and then choose the appropriate scheme.

6.2. Micro-level performance

This sub-section zooms in and examines the performance at the individual users' level. Specifically, we look at four representative users summarized in Table 1. The population size is still $N = 100$ and the remaining 96 users are all "normal" (or baseline) users (the same as User 4). As user parameters are fixed in this setting, it makes sense to exclude NE and ITF-CCP from the comparison. System parameters remain the same as in Section 5.2.

In Fig. 6a, we compare the service quota each user received against his demand, i.e., q_i/Q_i. Under ADF, all users except for User 2 reaches ~75% which is the total service availability level (recall that $Q_{tot} = 0.75\sum_{i=1}^{N} Q_i$). User 2 contributes only half of what the other users contributed, and, in return, he is allocated only 37.5% of his demand, which is also half that of the others. This shows that ADF scheme encourages higher user contributions. User 1 benefits more in terms of *absolute* service quota, q_i, because ADF does not discriminate against high-demand users. In sharp contrast, ITF grants users 1–3 zero service. The reason is that they are classified as "hard-to-satisfy" or "unwelcome" users, following the philosophy of ITF. As such, priority is given to the remaining 97 "normal" users who equally share Q_{tot} and obtain a q_i/Q_i even slightly higher than 75%.

Fig. 6b gives each user's utility. From the above, it is easy to understand that, under ITF, only User 4 receives positive utility which is also slightly higher than the maximum of other schemes. Under ADF, User 2 receives lower utility corresponding to the lower q_i as in Fig. 6a, User 3 receives lower utility because of his higher cost as per the definition of utility in Eq. (6).[6] User 1 receives the same utility as normal users because ADF tries to make q_i/Q_i commensurate with ψ_i and thereby removes the difference between User 1 and normal users.

It is worth noting that, while User 2 is under-privileged in both ADF and ITF, he (undesirably) obtains the same amount of service as normal users under EA and DA. This clearly demonstrates the effectiveness of the built-in mechanism to encourage a high level of contributions in the proposed ADF and ITF schemes.

[6] The approximately linear relationship in spite of the existence of $\log(\cdot)$ is because $\Psi \frac{q_i}{c_i Q_i}$ is small and $\log(1 + x) \sim x$ for small x.

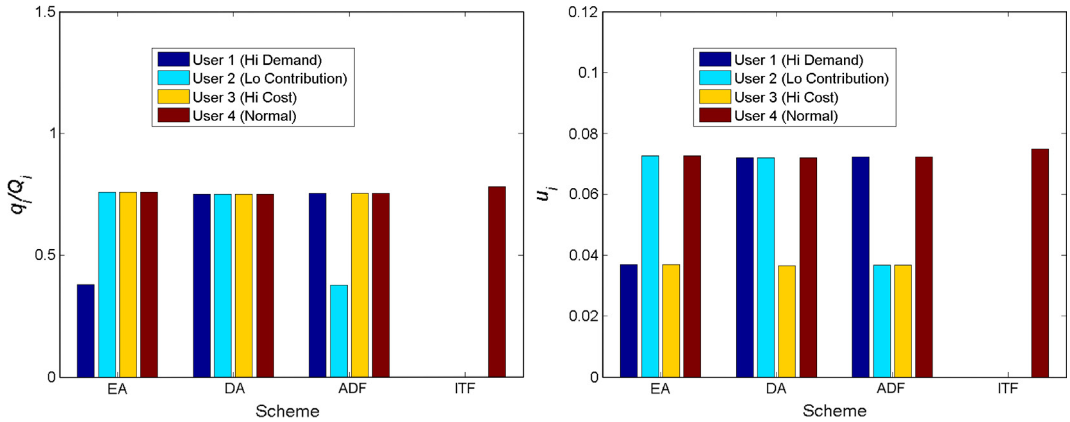

(a) Demand-normalized service quota (q_i/Q_i). (b) Individual user utility.

Fig. 6. Performance of "unwelcome" users vs. "normal" users.

Table 2
"Welcome" users vs. "normal" user.

User	Type	Demand (day)	Contribution (Mb)	Cost ($/Mb)
1	Low-demand	0.25	0.5	0.5
2	High-contrib.	0.5	1	0.5
3	Low-cost	0.5	0.5	0.25
4	Normal	0.5	0.5	0.5

Now, we investigate the other scenario with "welcome" users to see if and how they benefit from the proposed schemes. Table 2 gives the user setting and the remaining users (5–100) are the same as the normal user (User 4).

The simulation results are given in Fig. 7. Regarding demand-normalized service quota, we see that User 1 achieves $q_i/Q_i \approx 1.5$ under the EA scheme, which would lead to wastage of resource. On the other hand, DA equalizes q_i/Q_i for all users, thereby not motivating any "good" users. ADF rewards User 2, as desired, and not User 1 and 3, because it aims to fairly allocate q_i/Q_i only and does not take cost into account. ITF, at last, can be viewed as the best scheme among the four because all the good users (1–3) receive higher demand-normalized service quota than normal users, which encourages a high level of user contributions.

With respect to user utility (Fig. 7b), we see that User 2, though contributing highly, is not better off than normal users under EA and DA. This will adversely dampen the morale of good contributors. On the contrary, ADF and ITF both show favour to User 2 and, among these two schemes, ITF also satisfies User 1 and 3 more than ADF, which is desirable and can be understood by comparing with Fig. 7a. The reason that User 3 receives apparently higher utility than other users is because of the definition of utility (Eq. (5)) and the values of q_i/Q_i shown in Fig. 7a.

Finally, we note that the approximately linear relationship between Fig. 6a and b is not reproduced between Fig. 7a and b. This is because $\Psi \frac{q_i}{c_i Q_i}$ is no longer small enough in this case, which renders the rule of $\log(1+x) \sim x$ inapplicable.

7. Discussion

7.1. Relationship between fairness and social welfare

In general, it is well accepted that fairness and social welfare are two contradictory objectives. Indeed, they are even divergent in our participatory sensing context where the definition of social welfare involves cost but fairness does not (otherwise, the varying user cost profile will result in unfairness, as explained in Section 4). In fact, even in the case of equalized costs or proportional costs, we show below that the two objectives are still not simultaneously achievable.

7.1.1. Equalized costs

Assume all users bear the same cost. Let us look at User 1 and 2 in Table 2. Sufficient for demonstration purposes, consider the system with only these two users and a total quota of $Q_{tot} = 0.5$. We can show without much difficulty that, in order to maximize Jain's fairness index as defined in Eq. (3), we should allocate the two users with $q_1 = 0.1$ and $q_2 = 0.4$, which achieves $J = 1$ (the maximum). One can also easily verify that running ADF as per Eq. (2) will yield the same solution.

Now, consider the social welfare. The above solution leads to $S = 0.5U(0.8\Psi) + U(1.6\Psi)$, regardless of any concrete form of $U(\cdot)$. Based on the strict concavity of $U(\cdot)$, we further have

$$S = 1.5 \times \left[\frac{1}{3}U(0.8\Psi) + \frac{2}{3}U(1.6\Psi)\right]$$
$$< 1.5U\left(\frac{0.8}{3}\Psi + \frac{3.2}{3}\Psi\right) = 1.5U\left(\frac{4}{3}\Psi\right).$$

Next, we need to show that the right hand side is achievable. Indeed, the solution of $q_1 = 1/6$ and $q_2 = 1/3$ leads to

$$S = 0.5U\left(\frac{4}{3}\Psi\right) + U\left(\frac{4}{3}\Psi\right) = 1.5U\left(\frac{4}{3}\Psi\right).$$

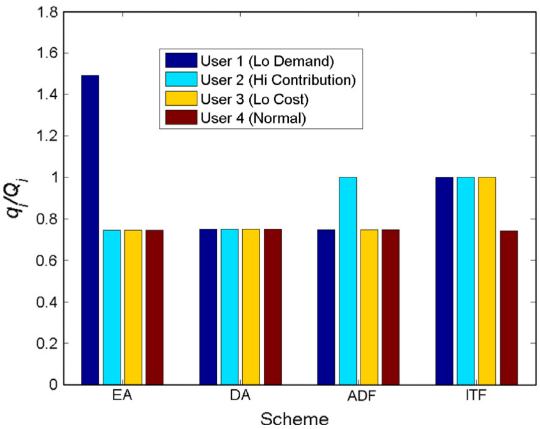

(a) Service quota relative to user demand.

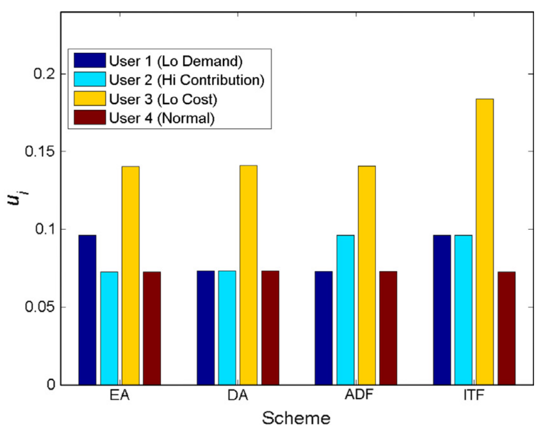

(b) Individual user utility.

Fig. 7. Performance of "welcome" users vs. "normal" users.

In fact, this solution of $q_1 = 1/6$ and $q_2 = 1/3$ can be obtained by ITF and a quick way to verify is by referring to Fig. 2.

Therefore, it becomes evident that the fairness maximizing solution does not maximize social welfare, and vice versa.

7.1.2. Proportional costs

Assume all users are equally cost-effective in contributing, i.e., $c_i/\psi_i = c_j/\psi_j,\ \forall i,\ j$. Now, consider User 2 and 3 in Table 2 and, still, $Q_{tot} = 0.5$. In this case, it can be shown that the solution of $q_1 = 1/3$ and $q_2 = 1/6$, which can be obtained via either ADF or ITF, maximizes both fairness and social welfare. However, this *cannot* be generalized. To see this, let $Q_3 = 1$ instead. Then, the optimal solution to maximizing fairness is $q_1 = q_2 = 0.25$ which can be obtained using ADF and achieves $J = 1$. The social welfare under this solution is $S = U(\Psi) + 0.5U(\Psi) = 1.5U(\Psi)$. Let us now consider a different solution of $q_1 = 7/18$ and $q_2 = 1/9$ which yields $S = U(\frac{14}{9}\Psi) + 0.5U(\frac{4}{9}\Psi)$. For evaluation purposes, take $\Psi = \sum_i \psi_i$ and $U(x) = \log(1 + x)$, and we obtain $S = 1.37$ for the first solution and $S = 1.46$ for the second solution. This again demonstrates that the fairness maximizing solution does not maximize social welfare, and vice versa. In fact, the second solution can be obtained by ITF too.

This section corroborates our simulation results in Section 6.1 that there is no one-size-fits-all solution, and consequently, it is the service provider's responsibility to decide which objective, fairness or social welfare, to aim for.

7.2. Evaluating contribution level

The contribution level ψ_i quantifies the value or usefulness of user contributed data. While the exact method of evaluating this depends on the specific application, this section provides a number of illustrative examples.

In general, a user's contribution can be assessed in terms of its *intrinsic* value or *extrinsic* value. The intrinsic value is conveyed by the contributed data itself. For example, suppose the application is to estimate a parameter denoted by $\mathbf{x}$, such as travel speed or target location, and user i's contributed measurement is $\mathbf{y}_i$. Then, his contribution ψ_i can be defined as the *reduction of uncertainty* in estimating $\mathbf{x}$ by incorporating $\mathbf{y}_i$. A formal definition, which is based on *information utility*, can be found in [18]. Another example, based on a similar reasoning, is to use the value of information (VoI) [19].

The extrinsic value is an external attribute associated with the contributed data but cannot be conveyed by the data itself. Examples include user reputation, signal strength and sensor accuracy [4]. It usually pertains to the contributing entity (such as a user or device) and/or the context (such as time or location), and is especially useful when obtaining the data involves additional cost, or when evaluating the intrinsic value is difficult. For example, in a decision fusion application, a user's contribution (e.g., decision on the presence of a phenomenon of interest (PoI)) can be evaluated by the user's *historical* performance (how accurate were his decisions) without knowing the current data, which can be formulated as a *likelihood ratio* as in [20]. Another example is the *Mahalanobis distance* [18] which characterizes how likely a node or user is to provide the most useful information (the most reduction of uncertainty in the context thereof) without getting the contributed data.

7.3. Implementation considerations

Whichever scheme chosen by the service provider can be executed on a server farm or cloud. The server has all the information needed, including $\psi_i,\ c_i,\ Q_i$ of all the users and $Q_{tot},\ \Psi$ and N, based on which it assigns q_i following the algorithms presented above. Here, we briefly explain how the service provider can calculate a user's cost c_i. The major component of c_i is the mobile data charge incurred by the user when making a contribution. This can be calculated using the user's mobile data plan obtained from the user's registration information, the timestamp of the contribution (e.g., peak or off-peak times), and the amount of data transmitted. The other minor component is the user's battery drainage, which can be gauged from the user's phone model specification as well as his contribution time profile.

If needed, the service provider can *feed back* individual parameters such as ψ_i and c_i to the corresponding user and *announce* common parameters such as Q_{tot}, Ψ and N. This is particularly needed in the NE case where each user needs to calculate his optimal Q_i^* in order to declare it.

Finally, for the sake of user convenience, the declaration of Q_i can be automated by using the application running on each client (e.g., smartphone) that makes user contributions. The application periodically sends Q_i to the server based on the user's historical demands, unless the user specifically intervenes to modify the value.

8. Conclusion

Participatory sensing offers a promising approach to large-scale data collection, but requires a steady stream of user data contributions in order to be successful. As an alternative to monetary or reputational incentives, we have presented a service allocation-based approach that leverages on the duality of the roles played by participants to motivate users to make data cointributions by allocating different quantities of compelling services that they desire to consume. We proposed two schemes, ADF and ITF, with the objective of maximizing fairness and social welfare, respectively, while at the same time encouraging user contributions. Our analysis and simulation results show the effectiveness of these schemes in achieving their respective objectives. Two tailored variations, NE and ITF-CCP, have also been presented to provide the optimal equilibrium and to handle uncertainty in declared service demands.

In our future work, we plan to extend the current case where there is only a single service provider, to the case where there may be multiple service providers.

Acknowledgement

This work was done under the EDASACEP project within the Data Value Chain as a Service (DVCaaS) TSRP research programme, funded by SERC, A*STAR, Singapore.

Appendix A. Solving problem (b) (Section 5)

This section sketches the solution to problem (7), following the same line of thought as solving problem (a). First, we calculate the reciprocal of marginal weighted utility $\frac{Q_i/\Psi+q_i}{\psi_i/c_i}$. Accordingly, we convert the optimization problem into a iced-tank representation like Fig. 2, with the following differences: the bottom area of each tank changes to ψ_i/c_i, the ice volume of each tank changes to Q_i/Ψ, and the empty space of each tank changes to Q_i. Very importantly, note that the volume (and hence height) of ice and that of the empty space of each tank, are now proportional to each other (with a constant coefficient Ψ). Therefore, sorting the tanks in order of tank heights is equivalent to sorting the tanks in order of ice levels, meaning that a user who is given a higher priority to receive service quota (due to lower ice level) will surely be satisfied earlier (due to lower tank height). Therefore, the problem is simpler than (a) and so is the corresponding solution algorithm, which we call ITF-(b), for which the pseudo-code is omitted for brevity.

Without much difficulty, we can show that all the previously given theorems, corollaries and lemmas still hold, except for the following changes:

- In Lemma 1, h_i should change to $h_i = c_i Q_i/\psi_i$.
- In Theorem 3, the optimal strategy profile should change to

$$Q_i^* = \frac{\psi_i/c_i}{\sum_{l=1}^{N}\psi_l/c_l} Q_{tot}. \tag{A.1}$$

Finally, the CCP version of ITF-(b), which we call ITF-CCP-(b), can be obtained in the same way as ITF-CCP from the ITF scheme, with the only difference that

$$\text{tank}_i = ice_i + \frac{\widehat{Q}_i + \sigma_i z_{\alpha_i}}{\psi_i/c_i}.$$

Appendix B. Joint chance constraints

As opposed to the individual chance constraints (17), a joint chance constraint is in the following form:

$$\Pr\Big(q_i \leqslant \widetilde{Q}_i, \forall i = 1, \ldots, N\Big) \geqslant 1 - \alpha. \tag{B.1}$$

It requires *all* the users to satisfy $q_i \leqslant \widetilde{Q}_i$ in $1-\alpha$ of the time, or equivalently, the probability that at least one user is granted $q_i > \widetilde{Q}_i$ must be less than α. This does not characterize our problem better than (in fact, not as well as) the individual chance constraints. In addition, as (17) associates a probability α_i to each user i, the service provider can leverage this to provide further motivation via user differentiation, which (B.1) lacks. Nonetheless, we provide the treatment below for interested readers.

In general, the calculation of joint chance constraints involves dealing with multi-dimensional distributions. Fortunately, the $\widetilde{Q}_i$'s in participatory sensing can be treated as being independent of each other. Hence, assuming $\widetilde{Q}_i \sim \mathcal{N}(\mu_i, \sigma_i)$, (B.1) can be broken down into

$$\prod_{i=1}^{N}\Pr\Big(q_i \leqslant \widetilde{Q}_i\Big) \geqslant 1 - \alpha \Longleftrightarrow \prod_{i=1}^{N}\text{erfc}\left(\frac{q_i - \mu_i}{\sqrt{2}\sigma_i}\right) \geqslant 2 - 2\alpha \tag{B.2}$$

Unfortunately, an explicit form for the upper limit to q_i is not obtainable from (B.2), and one has to resort to a stochastic simulation-based approach. For this, the reader can refer to a stochastic simulation-based genetic algorithm proposed in [21], where the initialization process, selection, crossover, and mutation operations are the same as general genetic algorithms except that stochastic simulation is employed to check the feasibility of new offspring (i.e., solution) and to handle stochastic constraints.

References

[1] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, M.B. Srivastava, Participatory sensing, in: ACM Sensys Workshop on World Sensor Web (WSW), 2006.
[2] M. Mun, S. Reddy, K. Shilton, N. Yau, J. Burke, D. Estrin, M. Hansen, E. Howard, R. West, P. Boda, PEIR, the personal environmental impact report, as a platform for participatory sensing systems research, in: ACM MobiSys, 2009, pp. 55–68.
[3] P. Dutta, P.M. Aoki, N. Kumar, A. Mainwaring, C. Myers, W. Willett, A. Woodruff, Common sense: participatory urban sensing using a network of handheld air quality monitors, in: ACM SenSys, 2009.
[4] J.K.-S. Lau, C.-K. Tham, T. Luo, Participatory cyber physical system in public transport application, in: Proc. CCSA, IEEE/ACM UCC, 2011.
[5] J. Brustein, Projects Use Phone Data to Track Public Services, New York Times, June 5, 2011.
[6] J.-S. Lee, B. Hoh, Sell your experiences: a market mechanism based incentive for participatory sensing, in: IEEE PerCom, 2010.
[7] D. Christin, C. Rosskopf, M. Hollick, L.A. Martucci, S.S. Kanhere, Incognisense: an anonymity-preserving reputation framework for participatory sensing applications, in: IEEE PerCom, 2012, pp. 135–143.
[8] J. Park, M. van der Schaar, Incentive provision using intervention, in: IEEE INFOCOM, 2011.
[9] S. Reddy, D. Estrin, M. Hansen, M. Srivastava, Examining micro-payments for participatory sensing data collections, in: UbiComp, 2010.
[10] R. Jain, D.-M. Chiu, W. Hawe, A quantitative measure of fairness and discrimination for resource allocation in shared computer systems, ACM Comput. Res. Repos. (CoRR) (1998).
[11] S. Shenker, Fundamental design issues for the future Internet, IEEE J. Sel. Areas Commun. 13 (1995).
[12] P. Marbach, Cooperation in wireless ad hoc networks: a market-based approach, IEEE Trans. Network. 13 (2005).
[13] R. Ma, S. Lee, J. Lui, D. Yau, Incentive and service differentiation in P2P networks: a game theoretic approach, IEEE Trans. Network. 14 (2006) 978–991.
[14] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.
[15] D. Tse, P. Viswanath, Fundamentals of Wireless Communication, Cambridge University Press, 2005.
[16] E. Rasmusen, Games and information: an introduction to game theory, fourth ed., Blackwell Publishing, Oxford, 2007.
[17] A. Prekopa, Stochastic programming, Kluwer, Dordrecht, 1995.
[18] M. Chu, H. Haussecker, F. Zhao, Scalable information-driven sensor querying and routing for ad hoc heterogeneous sensor networks, Int. J. High Perf. Comput. Appl. 16 (2002).
[19] A. Krause, E. Horvitz, A. Kansal, F. Zhao, Toward community sensing, in: ACM/IEEE IPSN, 2008.
[20] H.-X. Tan, M.-C. Chan, W. Xiao, P.-Y. Kong, C.-K. Tham, Information quality aware routing in event-driven sensor networks, in: IEEE INFOCOM, 2010.
[21] K. Iwamura, B. Liu, A genetic algorithm for chance constrained programming, J. Inform. Opt. Sci. 17 (1996) 409–422.

Chen Khong Tham is an Associate Professor at the Department of Electrical and Computer Engineering (ECE) of the National University of Singapore (NUS). His current research focuses on sensor network infrastructures and real-time sensor data analytics involving cyber-physical systems, wireless sensor networks, cloud computing and participatory sensing. From 2007-10, he was on secondment at A∗STAR Institute for Infocomm Research (I2R) Singapore and served as principal scientist and department head of the Networking Protocols Dept, and programme manager of the Personalization, Optimization & Intelligence through Services (POISe) (Services) Programme. From 2006-09, he was the programme manager of a multi-institution research programme on UWB-enabled Sentient Computing (UWB-SC) funded by A∗STAR Singapore. He obtained his Ph.D. and M.A. degrees in Electrical and Information Sciences Engineering from the University of Cambridge, United Kingdom, and was an Edward Clarence Dyason Universitas21 Fellow at the University of Melbourne, Australia. He is an Associate Editor of the IEEE Internet of Things Journal (IoT-J) and is in the Editorial Board of the International Journal of Network Management. He was the general chair of IEEE SECON 2014, IEEE AINA 2011 and IEEE APSCC 2009 conferences. He and his co-authors won the Best Paper award at IEEE ICUWB 2008 and ACM MSWiM 2007.

Tie Luo is currently a Scientist at the Institute for Infocomm Research (I2R), Singapore. He received his Ph.D. degree in Electrical and Computer Engineering from the National University of Singapore (NUS). He was a postdoctoral Research Fellow at the Communications & Networks Lab in NUS from 2011 to March 2012 and at the Computer Networks and Distributed Systems Lab in NUS from 2008 to 2010. His current research interests are participatory sensing, wireless sensor networks, software-defined networking and OpenFlow.