

One-to- n Scrip Systems for Cooperative Privacy-Enhancing Technologies

Mathias Humbert, Mohammad Hossein Manshaei, and Jean-Pierre Hubaux
Laboratory for Communications and Applications (LCA), EPFL, Switzerland
{mathias.humbert, hossein.manshaei, jean-pierre.hubaux}@epfl.ch

Abstract—Scrip is a generic term for any substitute for real currency; it can be converted into goods or services sold by the issuer. In the classic scrip system model, one agent is helped by another in return for one unit of scrip. In this paper, we present an upgraded model, the one-to- n scrip system, where users need to find n agents to accomplish a single task. We provide a detailed analytical evaluation of this system based on a game-theoretic approach. We establish that a nontrivial Nash equilibrium exists in such systems under certain conditions. We study the effect of n on the equilibrium, on the distribution of scrip in the system and on its performance. Among other results, we show that the system designer should increase the average amount of scrip in the system when n increases in order to optimize its efficiency. We also explain how our new one-to- n scrip system can be applied to foster cooperation in two privacy-enhancing applications.

I. INTRODUCTION

Over the last two centuries, non-governmental currencies, known as *scrip*, have been issued by private companies or local communities for many different purposes. For instance, to pay employees in isolated mining or logging camps, company scrip was used in lieu of regular money. More recently, community-issued scrip, such as the Detroit Community Scrip, has been issued in order to restore economic confidence, and help consumers make ends meet [3]. In the last decade, scrip systems have been proposed in order to thwart free riding in online environments (e.g., file sharing or resource sharing [17], [30]). The free-rider problem is particularly serious in peer-to-peer (P2P) networks such as BitTorrent, LimeWire or Gnutella, in which most users (85 percent) do not share any files [16].

Although scrip systems can help ensure fairness and prevent free riding, such systems are exposed to similar behavior as in real-world economies that lead to the same monetary issues. The Capitol Hill Baby Sitting Co-Op [28], a concrete scrip system created by a group of parents working on Capitol Hill, faced a recession and a monetary crash due to its monetary policy. Several researchers further studied the dynamics of scrip systems, based on these issues [14], [20], [21]. Among other results, they show that agents following threshold strategies led to a nontrivial Nash equilibrium. They show the impact of the amount of scrip in circulation on the efficiency of the system. In particular, they show that efficiency (social welfare) increases with the average amount of scrip per agent, until some point where the system experiences a monetary crash. At that point, no agent is willing to work anymore and social welfare falls to zero. Finally, they consider different “irrational” behaviors, such as altruists and hoarders, and

identify the impact of sybils and collusion on scrip systems.

The original scrip system assumes one transaction at a time, where one agent provides a service to another and gets paid one dollar¹ for it (one-to-one exchange) [14]. Previous work has brought a number of relevant results. However, there is an urgent need to extend the one-to-one scrip system to a system involving more than one dollar and two agents at a time in order to tackle new challenges led by modern IT systems, such as fostering cooperation in *privacy-enhancing applications*.

Privacy-enhancing technologies, such as anonymity networks [10], [9], [13], [25], provide valuable privacy benefits for Internet users. Among other benefits, anonymity networks can prevent price discrimination in e-commerce by concealing IP addresses. They are also used by journalists or human rights activists to circumvent censorship in dictatorial countries. For instance, there was a dramatic increase of Tunisian Tor [2] users during the Jasmine Revolution in January 2011 [1].

Many privacy-preserving mechanisms require cooperation among multiple users in order to achieve a good level of privacy. However, cooperation is not free, and its inherent cost often prevents users from collaborating. For example, in anonymity systems, running a relay node costs a non-negligible amount of bandwidth and processing power. Back in 2003, Acquisti et al. already highlighted the need of incentives to offer and use anonymity services [5]. Whereas the use of anonymity networks has improved since then, the number of relays is still much lower than the number of clients, and the client-to-relay ratio keeps growing. In 2009, there were 1,500 Tor relays for approximately 100,000 simultaneously active Tor clients [22], whereas, in June 2011, there were 2,500 relays for 300,000 to 400,000 clients [1].

Among other incentives for acting as a relay in anonymity networks, several schemes propose to make use of micropayments to reward users relaying others’ anonymous traffic [6], [8], [12]. These previous works have mainly contributed to the design of anonymous and secure micropayments. However, they did not evaluate the monetary issues that could appear in such systems. Assuming an anonymous circuit requires the cooperation of n relays, each client has to own (at least) n dollars in order to reward each of these n relays. In order to earn enough scrip to afford such a relaying service, each client will then have to serve - relay anonymous traffic - for other users in the anonymity network.

¹We refer to the unit of scrip as the dollar.

This leads us to define and study the one-to- n scrip system: one agent requests n other agents to fulfill a service and pays each of them one dollar. This scheme also better complies with current file sharing systems, such as BitTorrent, where an agent downloads multiple equal-size *chunks* from different neighboring peers of the torrent. In order to download an entire file and get any utility from it, an agent needs n peers who volunteer to upload their chunks. Thus, he must reward n agents with n dollars.

In this paper, we develop and study a new analytical model for scrip systems enabling a much wider range of applications. First, we precisely characterize the distribution of scrip in the one-to- n scrip system at equilibrium as a function of n and of the fraction of agents of each type. Second, we prove that, under certain assumptions, there exists a nontrivial Nash equilibrium where all agents play threshold strategies. We study the effect of n on the agents' strategies and the consequent equilibrium and prove that agents' thresholds increase with n . Third, we evaluate the efficiency (social welfare) of the one-to- n scrip system and notice that it tends to decrease when n increases. We show that a system designer can increase the scrip supply in order to offset the loss of efficiency caused by a larger n . This works well up to a point beyond which the system experiences a monetary crash. We show that this critical upper bound increases with n . Finally, we present how our one-to- n scrip system can help to improve fairness and efficiency in two privacy-enhancing applications. In particular, we evaluate the amount of scrip that should be allocated into the Tor network to optimize its performance.

The paper is organized as follows. In Section II, we present the formal model and the notations used throughout the paper. In Section III, we examine the behavior and convergence of the scrip system when agents adopt threshold strategies. In Section IV, we evaluate the effect of n and of the amount of scrip on the efficiency of the system. We provide two application examples in Section V before concluding in Section VI.

II. SYSTEM MODEL

In this work, we consider a scrip system with N agents who interact with each other. We consider a population of agents with different preferences and characteristics. Each agent has a type $t \in T$, where T is a finite set of types. The distribution of types is described by $\vec{f}$, where the element f_t represents the fraction of agents with type t . The type t of an agent is described by the tuple $t = (b_t, c_t, \delta_t, \alpha_t, \beta_t, \gamma_t)$, whose variables are defined in the rest of this section and in the following table of symbol.

At each time slot, one agent is selected proportionally to his request rate α_t to ask for a service. If this agent has at least $\$n$, he can afford a service and request other agents to fulfill this service. In order to have his request fully satisfied, n agents must be able and willing to collaborate. If there are less than n agents able and willing to volunteer, the request cannot be fulfilled, even partially, and the requester gains no utility. The service has to be satisfied in an "atomic" way. An agent is able to satisfy a service with probability β_t , and willing to volunteer

TABLE I
LIST OF SYMBOLS.

Symbol	Definition
N	Number of agents within the system
T	Set of agents' types
$\vec{f}$	Distribution of types
f_t	Fraction of agents of type t
W	Total amount of scrip in the system
m	Average amount of scrip per agent
n	Number of volunteers per request
b_t	Utility an agent gains for having a request satisfied
c_t	Cost of an agent when satisfying one request
δ_t	Rate at which an agent discounts his utility
α_t	Request rate
β_t	Probability that an agent is able to satisfy a request
γ_t	Likelihood to be chosen when an agent volunteers
k_t	Agent's threshold
$\vec{k}$	Vector of size $ T $ encompassing all k_t 's
S_{k_t}	Threshold strategy with threshold equal to k_t
$S_{\vec{k}}$	Strategy profile with agents' thresholds defined by $\vec{k}$
$\mathcal{V}$	State space describing the wealth of every agent
$\mathcal{X}$	Markov chain defined on $\mathcal{V}$
A	Set of agents who can afford a service
V	Set of agents who volunteer
M_i^t	Fraction of agents of type t with i dollars
p_u	Probability of earning one dollar
p_d	Probability of having a request satisfied
μ	Fraction of agents at their threshold

depending on his strategy. Moreover, an agent volunteering to provide service is chosen to fulfill another agent's request with likelihood γ_t .

When a service is performed, meaning that n agents fulfill the request of another agent, the requester (of type t_1) obtains some benefit $b_{t_1}(n)$ that is, in most cases, non-decreasing with n (see Section V for further details on the privacy gain). Each volunteer of type t_2 bears a utility cost c_{t_2} representing, for instance, the usage of bandwidth and processing power in anonymity networks. Thus, when n agents of same type t_2 collaborate with another agent (of type t_1) and satisfy his request, the whole cost is equal to nc_{t_2} , and the system's utility gain is $b_{t_1} - nc_{t_2}$. We assume that $b_{t_1} - nc_{t_2} > 0$, such that social welfare increases when a service is satisfied. The system would otherwise not be viable.

Regarding the monetary reward, an agent providing a service is paid some fixed amount of scrip that we assume is equivalent to \$1. Consequently, a service requester must spend $\$n$ to obtain a service. If the chosen agent does not have enough scrip, no transaction can take place in that time slot and social welfare stagnates. We model the system as an infinite extensive-form game where the total utility of an agent over time is the discounted sum of utilities at each time slot. The total discounted utility of agent i (of type t) is then $U_i = \sum_{\tau=0}^{\infty} \delta_t^\tau u_i(\tau)$, where δ_t represents the rate at which an agent of type t discounts utility.

As in the one-to-one scrip system, we assume that prices do not change over time, which allows the agents to know the future monetary cost of their service requests. As the first step towards an extended scrip system, we will consider a payoff-heterogenous population, i.e. b_t , c_t or δ_t might vary but $\alpha_t = \alpha$, $\beta_t = \beta$ and $\gamma_t = \gamma$, for all t . Differences in these parameters

should not fundamentally change the game-theoretic results. The one-to- n scrip system can be fully described by $(T, \vec{f}, N, m, n)$, where m is the average amount of scrip.

III. ANALYTICAL RESULTS

In this section, we prove the existence of Nash equilibrium when agents make use of threshold strategies. We also show the effect of n on the system, its equilibrium and the agents' strategies. We begin this section by describing the distribution of scrip, which will help us analyze the strategic behaviors of agents, as well as the resulting social welfare in Section IV.

A. Distribution of Scrip

Before analyzing the best strategies and the resulting equilibrium, it is crucial to examine what happens in the system if every agent adopts a predefined category of strategies, called *threshold strategies*. Such a class of strategies is easy to explain. If an agent has too little scrip, he will be willing to work in order to afford service requests later in time, until he reaches a point at which he will feel "wealthy" enough. This threshold represents how much scrip an agent wants to save up for future requests. Let S_k be the strategy where an agent volunteers when he has strictly less than k dollars and defects otherwise. With this definition, S_0 represents the strategy where an agent never volunteers, and S_∞ the strategy where he always volunteers. As threshold strategies depend on the agents' types, we write k_t to represent the threshold adopted by agents of type t . Vector $\vec{k}$ encompasses all the k_t 's, for all types t , and $S_{\vec{k}}$ is the corresponding strategy profile.

In our analysis, we assume that $W = mN < \sum_t f_t k_t N$, meaning that the total amount of scrip is not too high in order that the system analysis remains interesting. If $W \geq \sum_t f_t k_t N$, the system would converge to a state where each agent has reached his threshold, and thus does not want to volunteer anymore. We also assume that $m \geq n$. Otherwise, the system would converge to a state where no agent can afford a service, i.e. where all agents own less than n dollars. These two requirements seem reasonable because a system designer should ensure that (i) there are enough scrip in the system such that exchanges can happen, and (ii) there is not too much scrip in order to prevent procrastination and to encourage cooperation among agents.

Let $\mathcal{X}$ be a Markov chain over the state space $\mathcal{V}$ that describes the amount of scrip each agent owns. Each state of the Markov chain can be described by a vector $\vec{x}$, where x_i represents the amount of scrip agent i owns in state $\mathcal{V}_{\vec{x}}$. These states must satisfy some constraints: (i) $\sum_{i=1}^N x(i) = W$, and (ii) $0 \leq x(j) \leq k_t$, for all agents j with type t .² Thus, even if the Markov chain has a significant number of states (when N is large), their number is finite. If the Markov chain is in a state $\mathcal{V}_{\vec{x}}$, and agent j has a request satisfied by n agents $i_1,$

$i_2, \dots, i_n$, the Markov chain moves to another state, $\mathcal{V}_{\vec{y}}$, where

$$\begin{cases} y(j) = x(j) - n \\ y(i_\ell) = x(i_\ell) + 1, & \text{for } \ell = 1, \dots, n, \\ y(\cdot) = x(\cdot), & \text{for all other agents.} \end{cases} \quad (1)$$

We can already notice that, contrarily to the original scrip system, the aforementioned Markov chain is neither reversible nor symmetric, notably because no single transaction can restore the chain back to its previous state. Nevertheless, if there are at least $n + 2$ agents within the scrip system, there exists a limit distribution, as stated in the following lemma.

Lemma 1: If there are at least $n + 2$ agents in the system, then $\mathcal{X}$ is finite, aperiodic and irreducible and has a limit distribution.

Proof of Lemma 1: $\mathcal{X}$ is aperiodic. Assume that there are (at least) $n + 2$ agents $i_1, i_2, \dots, i_{n+2}$. Suppose $\mathcal{X}$ is in a state $\mathcal{V}_{\vec{x}}$ where at least one agent has $\$n$ or more and the others have less than their threshold amount of scrip. There must exist such a state by our assumption that m is interesting (i.e. neither too small nor too high). There exists a cycle of length $n + 1$ from state $\mathcal{V}_{\vec{x}}$ to itself: $i_2, i_3, \dots, i_{n+1}$ volunteer for i_1 , then $i_1, i_3, \dots, i_{n+1}$ volunteer for i_2 , and so on until $i_1, i_2, \dots, i_n$ volunteer for i_{n+1} . There is also a cycle of length $n + 2$: $i_2, i_3, \dots, i_{n+1}$ volunteer for i_1 , then $i_1, i_3, i_4, \dots, i_n, i_{n+2}$ volunteer for i_2 , then $i_1, i_2, i_4, \dots, i_{n-1}, i_{n+1}, i_{n+2}$ volunteer for i_3 , and so on until $i_2, i_3, \dots, i_{n+1}$ for i_{n+2} .

$\mathcal{X}$ is irreducible. Indeed, a Markov chain is said to be irreducible if all states communicate, or, in other words, if it is possible to reach any state from any other state. For any pair of states i and j of the Markov chain $\mathcal{X}$, we can show that the probability of going from i to j in a finite number of steps is strictly greater than 0, proving that any state is reachable from any other one.

Finally, as the number of states $\mathcal{V}$ is finite, $\mathcal{X}$ is also finite, and thus a limit distribution exists, and it is independent of the state in which the system starts [26]. ■

We can express the transition probabilities for all pairs of states i and j , $i \neq j$ that are directly reachable from each other³ as

$$P_{ij} = \frac{1}{|A|} \cdot \frac{1}{\binom{|V|-I}{n}}, \quad (2)$$

where A is the set of agents who can afford a service, i.e. who have at least $\$n$, in state i , and V is the set of volunteers, i.e. agents who have not reached their threshold amount of scrip, in state i too, and I is 1 if the agent requesting the service has an amount of scrip that is under his threshold, and 0 otherwise (because an agent cannot satisfy his own request). The transition probabilities depend on the values $|A|$ and $|V|$ that vary among the different states. Thus, the limit distribution is not uniform, even when $n = 1$. Instead of computing this limit distribution, we will focus on the corresponding distribution of scrip, because we are not interested in who has how much scrip, rather in the fraction of people that have a given amount of scrip.

²For simplicity, we assume that no one's amount of scrip exceeds their threshold.

³ $P_{ij} = 0$ if i and j do not directly communicate with each other.

For each state $\mathcal{V}$ of the Markov chain $\mathcal{X}$, there is a distribution of scrip M that describes the fraction of agents for each possible amount of scrip. More precisely, M_i^t represents the fraction of agents of type t who own $\$i$.⁴ For instance, if there is only one type of agent and we are in a state $\mathcal{V}$ where money is uniformly distributed ($x(j) = m \forall j$), then $M_m^t = 1$, and all other M_i^t are equal to zero. The distribution of scrip must satisfy two constraints:

$$\sum_t \sum_{i=0}^{k_t} i M_i^t = m \quad (3)$$

$$\sum_{i=0}^{k_t} M_i^t = f_t \quad (4)$$

First, the average amount of money is equal to m , and second, the fraction of agents playing S_{k_t} is equal to f_t (fraction of agents of type t). One can show that, if N is large, there exists a particular distribution M^* such that, with high probability, the Markov chain $\mathcal{X}$ will almost always be in a state $\mathcal{V}_{\bar{x}}$ such that the related distribution of scrip $M^{\bar{x}}$ is close to M^* . This kind of convergence around the most likely distribution is known as a *concentration phenomenon* in statistical mechanics [19]. According to Lemma 1, we can state that M^* exists. Before characterizing M^* , let us define two matrices B and C of size $(n+1) \times (n+1)$:

$$B = \left[\begin{array}{cccc|c} 1 & 0 & \cdots & 0 & -\theta_n \\ & & & & 0 \\ & & & & \vdots \\ & & & & 0 \end{array} \right] \quad C = \left[\begin{array}{cccc|c} 1 + \theta_n & 0 & \cdots & 0 & -\theta_n \\ & & & & 0 \\ & & & & \vdots \\ & & & & 0 \end{array} \right]$$

where $\mathbb{I}_n$ is the identity matrix of size n , $\theta_n = \frac{1}{\lambda n}$, λ chosen to ensure that (3) is satisfied with the distribution M^* defined in the following theorem.

Theorem 1: Given a payoff-heterogenous population, the distribution of scrip in a one-to- n scrip system will converge to

$$(M^*)_i = \frac{f_t \pi_i^t}{\sum_{j=0}^{k_t} \pi_j^t} \quad (5)$$

where the π_i^t 's are defined in the following way:

$$\begin{cases} \bar{e}_i^t = B^{n-1-i} C^{k_t-2n+1} \bar{v} \pi_{k_t}^t, & \text{if } i \in [0, n-2]; \\ \bar{e}_i^t = C^{k_t-n-i} \bar{v} \pi_{k_t}^t, & \text{if } i \in [n-1, k_t-n-1]; \\ \pi_i^t = \theta_n (1 + \theta_n)^{k_t-i-1} \pi_{k_t}^t, & \text{if } i \in [k_t-n, k_t-1]. \end{cases}$$

$\bar{e}_i^t$ and $\bar{v}$ are vectors of size $(n+1)$ defined as:

$$\bar{e}_i^t = \begin{bmatrix} \pi_i^t \\ \vdots \\ \pi_{i+n}^t \end{bmatrix} \quad \bar{v} = \begin{bmatrix} \theta_n (1 + \theta_n)^{n-1} \\ \vdots \\ \theta_n (1 + \theta_n) \\ \theta_n \\ 1 \end{bmatrix}$$

The proof can be found in the Appendix.

We have run simulations of the one-to- n scrip system in order to evaluate how close a real-system limit distribution was

⁴ M_i^t represents the fraction of agents who own $\$i$, regardless of their type.




Fig. 1. Distribution of scrip with $n = 5$ and $k_t = 30$ for 1000 agents. Dark (black) bars represent the theoretical distribution obtained in Theorem 1, whereas simulation results (after 10,000 iterations) are shown in light (blue).

to the theoretical limit distribution found in Theorem 1. Figure 1 illustrates the distribution of scrip in a one-to- n scrip system with 1000 agents with same type, $m = 10$, $k_t = 30$, $n = 5$ and $(b_t, c_t, \delta_t, \alpha_t, \beta_t, \gamma_t) = (1, 0.05, 0.95, 1, 1, 1)$. The dark bars show the theoretical distribution, whereas the light ones show the averaged distribution of scrip after 10,000 steps in the simulated model. Both distributions have very similar shapes. This allows us to believe that a real system would converge to some point close to the theoretical limit distribution. Back to the example depicted in Figure 1, we notice that both distributions increase until their peaks (at $n-1 = 4$), and then decrease until a very small peak (at $k_t - n = 25$). We notice a concentration of agents who have reached their threshold (at $k_t = 30$). By doing more simulations with various values of n , we have noticed the maximum of the curves always stands at $(n-1)$ if m remains smaller than half of k_t . This clearly shows how n influences the distribution of scrip.

B. Game Results: Strategies and Equilibria

In this section, we first analytically verify whether there exist an ϵ -best reply and a consequent nontrivial ϵ -Nash equilibrium in the one-to- n scrip system. Then, we evaluate the effect of n on the agents' strategies and on the Nash equilibrium. In particular, we show to what extent n influences the threshold vector $\vec{k}$. These results will help us measure the social welfare in the next section.

Note that δ_t has to be sufficiently large for all types t in order to reach a nontrivial Nash equilibrium where all agents follow a threshold strategy. If the discount factor is so small that it discounts too much future utility, all that matters is present utility and there is no incentive to volunteer now for future benefit. In this case, the only Nash equilibrium (*trivial* one) is to always defect for all agents. Thus, let us assume

that $\delta_t > \delta^*$, $\forall t$. Moreover, all nontrivial Nash equilibria in threshold strategies will be of the form $S_{\vec{k}}$ with $k_t \geq n$, $\forall t$. Indeed, there is no incentive for a rational agent to volunteer up to $k_t < n$ and then defect, because, in this case, the agent would never be able to afford any service.

In order to analyze the game, we consider a single agent i of type t , from whom point of view the system can be modeled as a Markov Decision Process (MDP). If N is large and n reasonably small with respect to N , what agent i does has essentially no effect on the behavior of the system and no great impact on the scrip distribution. We will later see that finding the best reply of agent i to the other agents' strategies is equivalent to finding an optimal policy for his MDP.

Assuming that the distribution of scrip is close to M^* (defined in Theorem 1) and all other agents have fixed their thresholds according to $\vec{k}$, we can compute two crucial probabilities for the optimal decision of agent i :

- (i) p_u , which is the probability of earning a dollar:

$$\frac{|A| - I}{N} \frac{n}{|V|} = \left(1 - \sum_t \sum_{j=0}^{n-1} (M^*)_j^t \right) \frac{n}{1 - \sum_t (M^*)_t^{k_t}}$$

- (ii) p_d , which is the probability of agent i having a request satisfied, or equivalently, of spending n dollars:

$$\frac{1}{N} \Pr(|V| \geq n) \cong \frac{1}{N}$$

p_u is the product of two probabilities: (i) the probability that some agent other than i who has n dollars is chosen to make a request, and (ii) the probability that i is the agent chosen to satisfy it. Whereas the first probability decreases a little with n , the second increases linearly with n , and thus p_u increases almost linearly with n . p_d is the probability of agent i will have a request satisfied, which can be approximated to the probability that agent i will be chosen to make a request.⁵ This probability only depends on N . However, n will influence the repercussion of p_d because if the agent is chosen to make a request, he will then spend n .

It follows from [24] that there exists an optimal policy for the MDP of agent i that is a threshold policy. This threshold, k_t , depends on p_u , p_d , b_t , c_t , δ_t , and n . We will prove later the effect of n on k_t . Note here that k_t must be a multiple of n . Indeed, supposing that an agent should decide between a threshold k_t (multiple of n) and a threshold $k_t + 1$, he would choose $k_t + 1$ only if the extra dollar would give him the opportunity to make one more request than with k_t , and gain more benefit in the future. As the agent needs n dollars to pay for a service, the extra dollar will be worth nothing, and eventually wasted. The cost c_t led by this extra dollar will not be compensated by a shorter expected time to make a request, assuming that δ_t is large enough and c_t is non-negligible.

Furthermore, if every other agent is playing a threshold strategy, for all m and $\epsilon > 0$, there exists an optimal threshold

⁵It is almost sure that n agents will be willing and able to volunteer under our initial assumption that n is reasonably small with respect to N . See Formula (8) for more details.

policy that is an ϵ -best reply to the strategy profile $S_{\vec{k}}$. This is valid only for $\delta_t > \delta^*$, large N , and n reasonably small with respect to N . Moreover, considering ϵ -best reply formalizes the fact that the optimal policy of the MDP and the best reply are not exactly the same. Indeed, both p_u and p_d are related to agent i 's MDP and they slightly differ from the corresponding probabilities of the game. They are only close with high probability, and after some amount of time. For instance, remember that we consider distribution M^* in the MDP, whereas the actual distribution in the game will be close but still different.

Before proving that a nontrivial ϵ -Nash equilibrium exists, we must show that the best reply function is non-decreasing in $\vec{k}$. Let $BR_m^t(S_{\vec{k}})$ be the best reply of an agent of type t given an average amount of money equal to m and the strategy profile $S_{\vec{k}}$. $BR_m^t(S_{\vec{k}})$ is non-decreasing in $\vec{k}$. First, it can be shown that if $\vec{k}' \geq \vec{k}$ (i.e., $k'_t \geq k_t$, $\forall t$), then $\sum_{j=0}^{n-1} (M^*)_j^{t'} \geq \sum_{j=0}^{n-1} (M^*)_j^t$ and $(M^*)_{k'_t}^{t'} \leq (M^*)_{k_t}^t$ for all types t . This means that, by increasing the threshold vector, more agents will not be able to afford a service, and fewer agents will reach their threshold. Therefore, with $\vec{k}'$, there will be fewer opportunities to earn money and more agents willing to volunteer for those opportunities, meaning that agents will earn money less often. Thus, agents will run out of money sooner. Hence, the utility of earning more scrip will increase, and as a result so will the best reply. We can now prove the existence of a nontrivial Nash equilibrium.

Theorem 2: For $\delta_t > \delta^*$, large N and n reasonably small with respect to N , there exists a nontrivial ϵ -Nash equilibrium where all agents of type t play S_{k_t} for some $k_t = l_t n$, $l_t \in \mathbb{N}$.

Proof: As the best reply function BR is non-decreasing, Tarski's fixed point theorem ensures that there exist a least and a greatest fixed point [29] that are equilibria. The least fixed point is the trivial equilibrium, and the greatest one can be reached by starting with S_∞ for all agents and using best-reply dynamics [20]. Moreover, if $\delta_t > \delta^*$, there exists a strategy profile $\vec{k}$ such that $BR(\vec{k}) \geq \vec{k}$. Monotonicity ensures that the greatest fixed point $\vec{k}^*$ is greater or equal to $\vec{k}$, and thus gives a nontrivial equilibrium. Note that n affects the nontrivial ϵ -Nash equilibrium. The higher n is, the further the MDP will be from the actual game. However, we can finely tune ϵ to cope with higher values of n . Moreover, as stated before, the best reply, for all types of agent, is a multiple of n . ■

The natural question that arises from the above theorem is: To what extent does n influence k_t , for all types t ? We already know that, $\forall t$, k_t must be a multiple of n . In fact, $\vec{k}$ increases with $b_t(n)$, thus with n as proved in the following theorem.

Theorem 3: For given values of m , c_t , α_t , β_t , γ_t , and $\delta_t > \delta^*$ for all t , the threshold vector $\vec{k}$ is increasing in n . More precisely, if $b_t = b_t(n)$,

$$\vec{k} \sim \Omega(b_t(n)) \quad (6)$$

Proof: Let us focus on the threshold $k_t = k$ of a particular agent and generalize it to the threshold vector $\vec{k}$. k is defined

as the maximum value such that

$$c_t \leq E[\delta_t^{j(k, p_u, p_d)}] b_t \quad (7)$$

holds, where $j(k, p_u, p_d)$ is a random variable whose value is the first round in which an agent starting with k dollars, using strategy S_k , has less than n dollars. The expectation is simply the discounted factor that will affect the agent's benefit at round j . First, we know that p_u increases almost linearly with n . Moreover, p_d is independent of n but the effect of being chosen to make a request is linear to n , as the agent will spend n dollars in that case. Thus, the effects of p_u and p_d on $j(k, p_u, p_d)$ approximately compensate each other. Assuming that b_t generally increases with n , the right part of (7) will increase with n if k remains unchanged. As c_t is fixed, the increase in b_t allows for the decrease of $E[\delta_t^{j(k, p_u, p_d)}]$ in front of b_t and still satisfy the inequality. As $j(k, p_u, p_d)$ increases in k (the higher the threshold is, the more money we have and the later we go under $\$n$) and $E[\delta_t^j]$ decreases in j , $E[\delta_t^{j(k, p_u, p_d)}]$ decreases in k . Moreover, as δ_t is close to one, $E[\delta_t^{j(k, p_u, p_d)}]$ decreases in $o(j(k, p_u, p_d))$, and so in $o(k)$. Thus, k can be increased with $b_t(n)$, more precisely in $\Omega(b_t(n))$. ■

Our results in this subsection show the existence of a nontrivial equilibrium under certain conditions, as well as some properties of this equilibrium. In the next section, we focus on the social welfare and the optimal amount of scrip in the system.

IV. SOCIAL WELFARE

In this section, we investigate how much scrip should be allocated in the one-to- n scrip system in order to optimize its performance, and thus social welfare.

A natural question arises when the system is at equilibrium: How good is it? Consider a single transaction involving only agents of type t . If a request is satisfied, social welfare increases by $b_t - nc_t > 0$. If no request is satisfied then no utility is gained. For a utility gain to happen, two events are required: (i) the agent chosen to make a request must have $\$n$, which occurs with probability $1 - \sum_{i=0}^{n-1} M_i$, and (ii) there must be n volunteers able and willing to satisfy the request. If μ is the fraction of agents at their threshold (i.e., the agents who do not want to volunteer), the probability of having at least n volunteers able to satisfy a request is

$$\begin{aligned} Pr(|V| \geq n) &= 1 - Pr(|V| < n) = 1 - \sum_{i=0}^{n-1} \beta_t^i (1 - \beta_t)^{(1-\mu)N} \\ &= 1 - (1 - \beta_t)^{(1-\mu)N} \cdot \frac{1 - \left(\frac{\beta_t}{1-\beta_t}\right)^n}{1 - \frac{\beta_t}{1-\beta_t}}. \end{aligned} \quad (8)$$

Expression $\frac{1 - \left(\frac{\beta_t}{1-\beta_t}\right)^n}{1 - \frac{\beta_t}{1-\beta_t}}$ goes to 1 if β_t is close to 0 or $n = 1$. This expression grows until infinity if β_t approaches 1. However, this factor is negligible with respect to $(1 - \beta_t)^{(1-\mu)N}$ if n is small with respect to N , which is always the case by assumption. As $(1 - \beta_t)^{(1-\mu)N}$ converges to 0 for large N or




Fig. 2. Social welfare for various average amounts of scrip m and various n . When m is too large with respect to n , social welfare falls to 0 (monetary crash).

β_t close to 1, the probability of finding n volunteers can be approximated by 1.

The total expected social welfare over all time is then

$$\left(1 - \sum_{i=0}^{n-1} M_i\right) \frac{b_t - nc_t}{1 - \delta_t}. \quad (9)$$

First of all, social welfare is maximized by minimizing the fraction of agents with less than n dollars. We can make $\sum_{i=0}^{n-1} M_i$ decrease by adding more scrip in the system. Indeed, if N is fixed, by increasing W , and thus m , the number of “poor” agents decreases. Thus, social welfare increases in m . However, social welfare does not increase to infinity with m and, beyond a certain average amount of money m^* , the only Nash equilibrium reached by the one-to- n scrip system is the trivial one, where no agent volunteers. We now evaluate the influence of n on the social welfare.

Theorem 4: For given values of b_t , c_t , δ_t , and $m < m^*$, social welfare of a one-to- n scrip system is decreasing in n .

Proof: In $(1 - \sum_{i=0}^{n-1} (M^*)_i)(b_t - nc_t)/(1 - \delta_t)$, two factors depend on n . First, $(1 - \sum_{i=0}^{n-1} (M^*)_i)$ decreases in n . Indeed, from Theorem 1, we can compute that, if $n' > n$, $\sum_{i=0}^{n'-1} (M^*)_i > \sum_{i=0}^{n-1} (M^*)_i$. Actually this sum increases approximately linearly with n . Second, $(b_t - nc_t)$ clearly decreases in n if b_t remains constant. Consequently, the whole expression decreases in n , and thus social welfare. ■

Figure 2 shows social welfare with respect to n and m , with the same population used in Figure 1. The only change is that now the benefit varies with n : $b_t(1) = 0.7$, $b_t(2) = 0.9$ and $b_t(n) = 1$, $\forall n > 2$. We notice that social welfare tends to decrease with n . The only scenario where it increases slightly is when $m = 4$ and n moves from 1 to 2. In this case, the increase in benefit is greater than the loss in cost and the loss due to agents that cannot afford a service. Note that social welfare falls to 0 when the average amount of money is too high with respect to n (e.g., when $m = 10$ and $n = 1$ or 2).

The fact that social welfare generally decreases with n seems surprising at first sight. Indeed, the more volunteers helping you, the higher the social welfare should be. Thus, the result is counterintuitive. There are two possible explanations for that. First, we must keep in mind that the n volunteers are not optional at all; without them no benefit can be obtained. Moreover, the cost of volunteering c_t does not decrease if more agents volunteer. The cost for each agent remains the same, regardless of n , thus the total cost for the system increases linearly with n . On the contrary, the benefit b_t does not usually increase so much with higher n . We can solve the first issue, or at least decrease its negative impact, by increasing the amount of scrip in the system. Indeed, in Theorem 4, we assume a fixed average amount of scrip, whereas a system requiring a higher number of volunteers per request will certainly need more scrip in circulation. This intuition is formalized by the following corollary.

Corollary 1: Assuming all other parameters are fixed, for a certain n , social welfare increases in m . It increases up to a certain average amount of scrip, m_n^* , beyond which there only exists the trivial Nash equilibrium (monetary crash). Furthermore, m_n^* increases in n .

Proof: The threshold vector $\vec{k}$ decreases when m increases, due to best-reply dynamics. Moreover, from the definition of M^* in Theorem 1, we can prove that $\sum_{i=0}^{n-1} (M^*)_i$ decreases if $\vec{k}$ decreases. Thus, $1 - \sum_{i=0}^{n-1} (M^*)_i$ and social welfare increase if m is increasing. Furthermore, from Theorem 3, we know that the threshold vector at equilibrium $\vec{k}$ increases with n . Thus, the threshold vector $\vec{k}$ will still decrease when m is increased but will reach zero (trivial equilibrium) beyond higher m with larger n . In other words, the system will bear a higher average amount of money before crashing when n increases. Hence, m_n^* increases in n . ■

Figure 2 depicts the positive effect of higher m on the social welfare. It also shows that scrip systems with higher n support higher average amount of scrip. For instance, when $m = 10$, the system crashes with $n = 1$ or $n = 2$ but not with $n \geq 3$. The ratio m_n^*/n must not go over a certain value that will be formally defined in future work. The fact that m_n^* is increasing in n can be well explained. When n increases, the agents feel less wealthy if they keep the same threshold values. Indeed, knowing that they then need more dollars to afford a single request, they will certainly be willing to save more dollars for future requests. Thus, if n increases, the agents will stop volunteering later, and thus the system will experience a monetary crash beyond a higher m_n^* . Indeed, a monetary crash appears when agents feel so rich that they are not willing to volunteer anymore. Increasing n prevents such behavior.

V. APPLICATIONS IN PRIVACY

In this section, we present two privacy-enhancing applications where a one-to- n scrip system can help improve fairness and efficiency: (i) anonymity networks [10], [9], [13], [25], and (ii) privacy-preserving data aggregation in participatory sensing [27]. This is not an exhaustive list of concrete applications of one-to- n scrip system but we focus on these

two examples because (i) anonymity networks are currently used by hundred of thousands of people to communicate and browse the Web anonymously, and (ii) participatory sensing could provide great benefits to society if there are enough mobile users participating in it, which would be possible only if the privacy of participants is ensured. In both examples, the more users involved in the privacy-preserving system, the higher privacy level the system reaches. Thus, it is absolutely crucial to have as many users as possible. Moreover, it is of the utmost importance that users help each other, i.e. volunteer for each other, in order to preserve the participants' privacy.

A. Anonymity Networks

Anonymity networks intend to prevent the Internet traffic of individuals from being tracked by governments or websites. As Tor [10] is the most popular anonymity network, we will focus on it for the rest of this section, even though the one-to- n scrip system can be applied to any other anonymity system.

The Tor network is based on onion routing, a design that creates a private network pathway by incrementally building a circuit of encrypted connections through relays (onion routers) on the network. Data packets are repeatedly encrypted (using the relays' public keys) and sent through multiple relays. Then, each relay removes a layer of encryption using its private key (it peels one layer of the onion) to uncover the address of the next relay on the path, and sends the packet to this relay where the same operation is repeated. In this way, no relay ever knows the complete path that a packet has taken. In order to prevent traffic linkability, users must renew their circuits over time. The Tor project website states that, currently, one circuit can be used for ten minutes [2]. The circuit's path length, i.e. the number of relays in the circuit, is a key parameter in Tor's deployment. As suggested in [10], using one or two hops only would allow for colluding relays to know too easily both the source and destination packets. Thus, the authors recommend to always choose at least three relays per circuit. In the current implementation, Tor selects exactly three relays for each circuit [2].

The lack of relays remains one of the main challenges in anonymity networks [11]. There are currently (June 2011) around 2,500 Tor relays for 300,000 to 400,000 users [1]. The corresponding client-to-relay ratio is not likely to decrease if the Tor network does not provide incentives for users to relay others' traffic. Acquisti et al. were the first to formalize the economics of anonymity and propose incentives to encourage users to serve for others [5]. The original Tor proposal already mentioned the need of incentives for a long-term scalable development of such an anonymity system [10]. In the last few years, various incentive mechanisms have been proposed. The first category of incentives is based on differentiated service for Tor users running a relay [23], [18]. The second category proposes to foster participation in traffic relaying by rewarding volunteers with anonymous micropayments [8], [6], [18]. Our idea is that users should reward their Tor relays with the micropayments earned when relaying others' traffic, everybody being involved in the relaying work such as in a




Fig. 3. Example of an anonymity network with 3 circuits (AGEJB, EIHAJ and CGHID) and 3 relays per circuit. Each user includes \$3 for each circuit he wants to build for anonymous communication. Then, each peer that accepts to relay traffic is rewarded with \$1 and can use it for his own relay requests.

P2P anonymity network [25], [13]. In this way, the anonymity workload will be shared among all the users benefiting from the system, thus ensuring fairness and preventing free-riding. Figueiredo et al. provided an anonymous payment-based incentive for such networks [12]. Note that all of the aforementioned micropayment-based incentive mechanisms proposed techniques to distribute coins (scrip) in an anonymous way, in order that privacy gains provided by anonymity networks are not jeopardized. Figure 3 depicts an example of this approach. It would ideally reduce the client-to-relay ratio to 1:1, i.e. all Tor clients would eventually run a relay.⁶ In order to analyze and evaluate optimal incentives to provide to the anonymity network, we can rely on the one-to- n scrip system.

In current implementation of Tor, n is equal to 3. This means that, in our scrip model, a Tor user will have to pay \$3 whenever he wants to create a new circuit in the Tor network. It is difficult to evaluate whether the anonymity benefit would increase with a larger n . We know that, with $n = 1$ or 2, the system would be too vulnerable to insider attacks. However, would the level of privacy really increase with n greater than 3? We will consider b_t constant for $n \geq 3$. Different types of benefits can encompass the fact that some users value anonymity more than others. The cost c_t of traffic relaying represents the bandwidth and power consumption used to forward Tor traffic. Different costs can represent various bandwidth or power capabilities of the relays. Assuming that all relays are of type t , the total cost of one request is equal to nc_t . We notice that the cost induced by one anonymity service request is increasing linearly with n , whereas the anonymity benefit remains more or less constant as soon as it reaches an acceptable value for n . Hence, the system designer should keep

n small to keep the relays' costs acceptable, and thus optimize social welfare. It is certainly a reason why Tor designers have chosen to set $n = 3$.

Concerning the other scrip system parameters, α_t , β_t , γ_t , they can model different behaviors and characteristics of the Tor users. First, α_t represents the request rate. Users surfing the Web more often in an anonymous way will request more service from Tor relays, thus α_t will increase. Second, a Tor agent might not be able to satisfy a traffic relaying request, which is encompassed in the value β_t . Finally, some Tor relays can have higher bandwidth than others, thus a higher quality of service when relaying traffic, or be well-known and more used than others. This could be represented by γ_t , which is the likelihood that an agent is chosen when he volunteers.

We have run simulations of an anonymity network with $N = 300,000$, $n = 3$ and the same homogenous population used in Figure 1, i.e. $b_t = 1$, $c_t = 0.05$, $\delta_t = 0.95$, and $\alpha_t = \beta_t = \gamma_t = 1$. Apart from N and n , the simulation parameters are not easy to determine and we plan to further investigate these in future work. Under these settings, the social welfare is maximized at $m = 10 < m_3^*$. With this average amount of money, there is only 2.5% of agents who cannot afford a service (i.e., with less than \$3). We conclude that a system designer should allocate $m \cdot N = 3$ million dollars within an anonymity networks of 300,000 users in order to optimize its efficiency.⁷

B. Privacy in Participatory Sensing

Participatory sensing is an example of novel mobile systems that leverage new capabilities in computation, communication, storage and sensing of mobile devices [7]. In participatory sensing, mobile users share sensing information, possibly including personal and/or location data, with service providers. However, the emergence of such people-centric systems leads to many issues, among which privacy is one of the most critical. Mobile users would certainly be willing to share sensing data, e.g. to help monitoring urban air pollution [15], but not at any cost to their privacy.

Shi et al. have recently proposed a concrete privacy-preserving data aggregation scheme for participatory sensing [27]. In this privacy-preserving mechanism, mobile nodes rely on their nearby peers to "hide" their data from the aggregation server (or service provider) that could be malicious (or at least curious), thus requiring cooperation from mobile nodes in their vicinity. Figure 4 depicts an example of this privacy-preserving scheme. First, each sensing node⁸ slices its data into $n + 1$ pieces, sends n encrypted pieces to neighbor nodes and retains the last piece. Second, the mobile nodes receiving pieces of data from other nodes aggregate all received pieces of data before transmitting them towards the aggregation server. Assuming that a fraction R of mobile nodes are malicious and

⁶In order to not discriminate Tor clients that cannot run a relay due to censorship, such as in China [4], we will make some exceptions and let such users benefit from Tor service for free. Indeed, denying anonymity to clients in censored regions would go against Tor's praiseworthy aim.

⁷Figure 2 provides more results on the social welfare for different values of n and m .

⁸A sensing node refers to any agent who uses his mobile device to sense his environment and submits sensing data.




Fig. 4. Example of privacy-preserving data aggregation in participatory sensing. Each mobile node sends data slices to neighbors in order to mix them. For encouraging cooperation, each node includes a fixed amount of scrip in all data slices.

can collude with the aggregation server, the normalized level of privacy is proportional to

$$P = \max\{1 - R^n - R^{|S|-1}, 0\} \quad (10)$$

where $|S|$ is the number of sensing nodes, which is equal to N if we assume that all users participating in the privacy-preserving scheme are also sensing nodes. Thus, the level of privacy increases with N , but also with n . However, this privacy-enhancing technique induces significant communication and computation overhead that also increases with n . As battery consumption is, with privacy, one of the main concerns of mobile users in participatory sensing, these communication and computation costs might prevent participants from volunteering to cover other nodes' data, thus threatening the whole privacy-preserving system. In order to foster cooperation and prevent free-riding, we propose to reward with scrip the mobile nodes that volunteer, and to rely on the one-to- n scrip system to optimize the efficiency of the monetary incentive.

First, contrarily to Tor networks, the value n is not at all defined in the initial proposal [27]. The system designer can tune this value to increase the privacy level provided by the mechanism, at the cost of communication overhead. Therefore, we do not attach any fixed value to n . Note that n should remain reasonably small with respect to the number of mobile nodes in the system in order for our theoretical results to apply. This will certainly be the case as the sensing nodes requesting help from others also suffer from too high communication overhead when they send their slices to too many neighbors. Thus, they will cap the number of "cover nodes" by themselves.

The benefit b_t that a sensing node (of type t) gains when a request is satisfied is related to the privacy utility it gains. As Equation (10) shows, b_t is dependent on n . Furthermore, as R is smaller or equal than one, $b_t(n) \propto P$ increases with n . Moreover, different types of benefits can encompass the fact that some agents are more privacy cautious and sensitive than others. The cost of volunteering is equal to c_t for all nodes

of type t . This cost represents the communication and computation overheads that lead to higher battery consumption. The type of c_t can represent the fact that some users are less willing to consume their battery or merely that their battery has a shorter lifetime. In conclusion, we clearly notice that the cost of one privacy-preserving request is increasing linearly with n , whereas the privacy benefit is increasing with n , but less than linearly. Hence, even if the requester gets higher payoff if he can send more data slices to more neighbors, the overall utility of the system, social welfare, is decreased.

The sensing nodes can have different amount of sensing data to submit to the aggregation server. This can be well described by the request rate α_t . Indeed, if nodes are collecting and submitting more data, they will request help of nearby peers more often. Furthermore, an agent might be unable to satisfy a request. For instance, its device can run out of battery or he can have a call at the same time. This can be encompassed in β_t . Finally, a node can be asked for covering others' data slices more often than others. For example, an agent can spend more time than another in a neighborhood with higher density of mobile sensing nodes. This difference can be represented by the likelihood that an agent is chosen when he volunteers, γ_t . As a concluding remark, we must mention that the number of data slices n a sensing node can send is also dependent on the density of the nodes in its vicinity. Thus, the optimal choice of n does not only depend on the nodes' privacy sensitivity, but also on the network density constraints.

We have also run experiments for participatory sensing systems, with various values of N and n . For $N = 1000$ and $n = 6$ and the same type of agents than for the previous application, social welfare is maximized with $m = 16$. This value is very close to m_0^* over which the system crashes. This average amount of scrip counterbalances the large value of n very well. It leads to almost the same percentage of agents who cannot afford a service (agents with less than \$6) than in Tor example with $n = 3$ (around 2.5%). Hence, in this scenario, a system designer should allocate $m \cdot N = 16,000$ dollars within the system to optimize its efficiency.

VI. CONCLUSION AND FUTURE WORK

In this work, we have proposed the first scrip system model that is able to tackle economic systems where one agent needs multiple volunteers simultaneously in order to have his request satisfied. For the novel one-to- n scrip system, we have proved that decisions agents make, based on threshold strategies, lead to ϵ -Nash equilibrium. Assuming that all agents of the system use threshold strategies, we have shown that the limit distribution towards which our scrip system will converge highly depends on n . Simulations of the one-to- n scrip system confirm this convergence. We have studied the effect of n on all results, notably on the agents' strategies, on the social welfare and on the maximum amount of scrip that the system can handle before crashing. We have proved that, at equilibrium, the agents increase their thresholds if n increases. However, in this case, social welfare decreases, which can be partially resolved by adding more scrip in the system. This

is possible because the maximum average amount of scrip that the system can bear before it crashes increases with n . Finally, we have shown that our upgraded scrip system can be very helpful for improving fairness and efficiency in two privacy-enhancing applications where cooperative volunteers are required. We have notably evaluated the average amount of scrip per agent that should be allocated into the Tor network to optimize its performance and fairness.

For future work, we will first formally evaluate the rate of convergence of our system. Then, we will consider different values of n for differentiated levels of privacy that would depend on the agents' preferences and privacy sensitivities. We will also consider other privacy-preserving and IT applications where the one-to- n scrip system can be implemented. We will investigate agents' possible strategies other than thresholds. Furthermore, non-standard behaviors such as altruism or hoarding will be studied. These behaviors should not necessarily be considered as irrational: (i) altruists can benefit from providing help to others, and (ii) hoarders may get some utility from owning more scrip. Finally, newcomers and their effect on the amount of scrip in circulation will be evaluated. As a consequence, variable prices could also be considered in our model.

ACKNOWLEDGMENTS

We would like to thank Olivier Levêque and Jean-Yves Le Boudec for their theoretical inputs, Igor Bilogrevic and Boi Faltings for their comments and feedback, as well as Cihangir Tezcan for the simulation of the scrip system. This work was supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

REFERENCES

- [1] Tor metrics portal. <http://metrics.torproject.org>.
- [2] Tor project. <https://www.torproject.org>.
- [3] Communities print their own currency to keep cash flowing. USA Today, April 2009.
- [4] Tor partially blocked in china, October 2009. <https://blog.torproject.org/blog/tor-partially-blocked-china>.
- [5] A. Acquisti, R. Dingledine, and P. Syverson. On the economics of anonymity. In *Financial Cryptography*, 2003.
- [6] E. Androulaki, M. Raykova, S. Srivatsan, A. Stavrou, and S. Bellovin. Par: Payment for anonymous routing. In *Privacy Enhancing Technologies*, 2008.
- [7] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. Srivastava. Participatory sensing. In *World Sensor Web Workshop*, 2006.
- [8] Y. Chen, R. Sion, and B. Carbunar. XPay: Practical anonymous payments for Tor routing and other networked services. In *Proceedings of the 8th ACM workshop on Privacy in the electronic society*, 2009.
- [9] G. Danezis, R. Dingledine, and N. Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *Security and Privacy, 2003. Proceedings. 2003 Symposium on*, 2003.
- [10] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proceedings of 13th USENIX Security Symposium*, 2004.
- [11] R. Dingledine, N. Mathewson, and P. Syverson. Challenges in deploying low-latency anonymity (draft). Technical report, 2005.
- [12] D. Figueiredo, J. Shapiro, and D. Towsley. Incentives to promote availability in peer-to-peer anonymity systems. In *13th IEEE International Conference on Network Protocols*, 2005.
- [13] M. Freedman and R. Morris. Tarzan: A peer-to-peer anonymizing network layer. In *CCS*, 2002.
- [14] E. J. Friedman, J. Y. Halpern, and I. A. Kash. Efficiency and nash equilibria in a scrip system for P2P networks. In *Proceedings of the 7th ACM conference on Electronic commerce*, 2006.
- [15] R. Honicky, E. Brewer, E. Paulos, and R. White. N-smarts: networked suite of mobile atmospheric real-time sensors. In *Proceedings of the second ACM SIGCOMM workshop on Networked systems for developing regions*, 2008.
- [16] D. Hughes, G. Coulson, and J. Walkerdine. Free riding on Gnutella revisited: The bell tolls? *IEEE Distributed Systems Online*, 6(6), 2005.
- [17] J. Ioannidis, S. Ioannidis, A. Keromytis, and V. Prevelakis. Fileteller: Paying and getting paid for file storage. In *Financial Cryptography*, 2003.
- [18] R. Jansen, N. Hopper, and Y. Kim. Recruiting new Tor relays with BRAIDS. In *CCS*, 2010.
- [19] E. T. Jaynes. Where do we stand on maximum entropy. *The Maximum Entropy Formalism*, pages 15–118, 1978.
- [20] I. A. Kash, E. J. Friedman, and J. Y. Halpern. Optimizing scrip systems: Efficiency, crashes, hoarders, and altruists. In *Proceedings of the 8th ACM conference on Electronic commerce*, 2007.
- [21] I. A. Kash, E. J. Friedman, and J. Y. Halpern. Manipulating scrip systems: Sybils and collusion. In *Auctions, Market Mechanisms and Their Applications*, 2009.
- [22] K. Loesing. Measuring the Tor network: Evaluations of client requests to directories. Technical report, Tor Project, 2009.
- [23] T.-W. J. Ngan, R. Dingledine, and D. S. Wallach. Building incentives into Tor. In *Financial Cryptography*, 2010.
- [24] M. Puterman. *Markov Decision Processes*. Wiley, 1994.
- [25] M. Rennhard and B. Plattner. Practical anonymity for the masses with morphmix. In *Financial Cryptography*, 2004.
- [26] S. I. Resnick. *Adventures in Stochastic Processes*. Birkhauser Boston, 1992.
- [27] J. Shi, R. Zhang, Y. Liu, and Y. Zhang. PriSense: Privacy-preserving data aggregation in people-centric urban sensing systems. In *IEEE INFOCOM 2010. 29th IEEE International Conference on Computer Communications*, 2010.
- [28] J. Sweeney and R. J. Sweeney. Monetary theory and the great capitol hill babysitting co-op crisis: Comment. *Journal of Money, Credit and Banking*, 9(1):86–89, 1977.
- [29] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific journal of Mathematics*, 5:285–309, 1955.
- [30] V. Vishnumurthy, S. Chandrakumar, and E. Sirer. Karma: A secure economic framework for peer-to-peer resource sharing. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

APPENDIX

Proof of Theorem 1: Let us focus on one type t and then generalize for all types. Knowing that agents of type t have $k_t + 1$ possible states of wealth (i.e., their amount of scrip can go from 0 to k_t), we can define a Markov chain Y over $k_t + 1$ states that describes the amount of scrip an agent of type t can own. When the Markov chain is in some state, it can either move one state up, or move n states down, or stay in the same state. The probability of moving one state up is

$$Pr(Y_{i+1}|Y_i) = \frac{n}{|V|} \quad (11)$$

and the probability of moving n states down is

$$Pr(Y_{i-n}|Y_i) = \frac{1}{|A|} \quad (12)$$

where A is the set of agents who can afford a service and V is the set of volunteers.

There is one state from which the Markov chain cannot go up (the state where the agent has $k_t + 1$ dollars), and some states from which Y cannot go down (the states where the

agent has less than n dollars). From (11) and (12), we can express the balance equations for all states:

$$\begin{cases} \frac{1}{|A|}\pi_i = \frac{n}{|V|}\pi_{i-1}, & \text{if } i = k_t; \\ \left(\frac{1}{|A|} + \frac{n}{|V|}\right)\pi_i = \frac{n}{|V|}\pi_{i-1}, & \text{if } i \in [k_t - n + 1, k_t - 1]; \\ \left(\frac{1}{|A|} + \frac{n}{|V|}\right)\pi_i = \frac{n}{|V|}\pi_{i-1} + \frac{1}{|A|}\pi_{i+n}, & \text{if } i \in [n, k_t - n]; \\ \frac{n}{|V|}\pi_i = \frac{n}{|V|}\pi_{i-1} + \frac{1}{|A|}\pi_{i+n}, & \text{if } i \in [1, n - 1]; \\ \frac{n}{|V|}\pi_i = \frac{1}{|A|}\pi_{i+n}, & \text{if } i = 0. \end{cases}$$

By multiplying everything by $\frac{|V|}{n}$ and setting $\lambda = \frac{|A|}{|V|}$ (the ratio between $|A|$ and $|V|$ is constrained by Equ. (3)), we get

$$\begin{cases} \frac{1}{\lambda n}\pi_i = \pi_{i-1}, & \text{if } i = k_t; \\ \left(\frac{1}{\lambda n} + 1\right)\pi_i = \pi_{i-1}, & \text{if } i \in [k_t - n + 1, k_t - 1]; \\ \left(\frac{1}{\lambda n} + 1\right)\pi_i = \pi_{i-1} + \frac{1}{\lambda n}\pi_{i+n}, & \text{if } i \in [n, k_t - n]; \\ \pi_i = \pi_{i-1} + \frac{1}{\lambda n}\pi_{i+n}, & \text{if } i \in [1, n - 1]; \\ \pi_i = \frac{1}{\lambda n}\pi_{i+n}, & \text{if } i = 0. \end{cases}$$

We then set $\theta_n = \frac{1}{\lambda n}$ and get the following recursions that fully describe the Markov chain distribution:

$$\begin{cases} \pi_i = \theta_n \pi_{i+1}, & \text{if } i = k_t - 1; \\ \pi_i = (1 + \theta_n) \pi_{i+1}, & \text{if } i \in [k_t - n, k_t - 2]; \\ \pi_i = (1 + \theta_n) \pi_{i+1} - \theta_n \pi_{i+n+1}, & \text{if } i \in [n - 1, k_t - n - 1]; \\ \pi_i = \pi_{i+1} - \theta_n \pi_{i+n}, & \text{if } i \in [0, n - 2]. \end{cases}$$

We can then express the last $(n+1)$ π_i 's (but π_{k_t}) with respect to π_{k_t} :

$$\pi_i = \theta_n(1 + \theta_n)^{k_t-i-1}\pi_{k_t} \quad \forall i \in [k_t - n, k_t - 1]. \quad (13)$$

From these $n+1$ values, we can build the vector $\vec{v}$ that will be used for the calculation of all other probabilities:

$$\vec{v} = \begin{bmatrix} \theta_n(1 + \theta_n)^{n-1} \\ \vdots \\ \theta_n(1 + \theta_n) \\ \theta_n \\ 1 \end{bmatrix} \quad (14)$$

Then, we can write

$$\begin{bmatrix} \pi_{k_t-n} \\ \vdots \\ \pi_{k_t-1} \\ \pi_{k_t} \end{bmatrix} = \vec{v}\pi_{k_t} \quad (15)$$

As $\forall i \in [n - 1, k_t - n - 1]$, $\pi_i = (1 + \theta_n)\pi_{i+1} - \theta_n\pi_{i+n+1}$, we can build a matrix C of size $(n+1) \times (n+1)$ that will be used for computing these probabilities:

$$C = \left[\begin{array}{cccc|c} 1 + \theta_n & 0 & \cdots & 0 & -\theta_n \\ \hline & \mathbb{I}_n & & & 0 \\ & & & & 0 \end{array} \right] \quad (16)$$

where $\mathbb{I}_n$ is the identity matrix of size n . We can then express, for instance, the (non-normalized) probabilities from state $k_t -$

$n - 1$ to state $k_t - 1$ in the following vectorial form:

$$\begin{bmatrix} \pi_{k_t-n-1} \\ \vdots \\ \pi_{k_t-1} \end{bmatrix} = C \begin{bmatrix} \pi_{k_t-n} \\ \vdots \\ \pi_{k_t} \end{bmatrix} = C\vec{v}\pi_{k_t} \quad (17)$$

By induction, we get the general form:

$$\begin{bmatrix} \pi_{k_t-n-j} \\ \vdots \\ \pi_{k_t-j} \end{bmatrix} = C^j \begin{bmatrix} \pi_{k_t-n} \\ \vdots \\ \pi_{k_t} \end{bmatrix} = C^j\vec{v}\pi_{k_t}. \quad (18)$$

Thus, we can compute π_i , $\forall i \in [n - 1, k_t - n - 1]$:

$$\begin{bmatrix} \pi_i \\ \vdots \\ \pi_{i+n} \end{bmatrix} = C^{k_t-n-i}\vec{v}\pi_{k_t} \quad (19)$$

Finally, as $\forall i \in [0, n - 2]$, $\pi_i = \pi_{i+1} - \theta_n\pi_{i+n+1}$, we build a matrix B of size $(n+1) \times (n+1)$ that will help computing the remaining probabilities:

$$B = \left[\begin{array}{cccc|c} 1 & 0 & \cdots & 0 & -\theta_n \\ \hline & \mathbb{I}_n & & & 0 \end{array} \right] \quad (20)$$

We can then express the non-normalized probabilities from state $n - 2$ to $2n - 2$:

$$\begin{bmatrix} \pi_{n-2} \\ \vdots \\ \pi_{2n-2} \end{bmatrix} = B \begin{bmatrix} \pi_{n-1} \\ \vdots \\ \pi_{2n-1} \end{bmatrix} = BC^{k_t-2n+1}\vec{v}\pi_{k_t} \quad (21)$$

By induction again, we get the general form:

$$\begin{bmatrix} \pi_{n-1-j} \\ \vdots \\ \pi_{2n-1-j} \end{bmatrix} = B^j \begin{bmatrix} \pi_{n-1} \\ \vdots \\ \pi_{2n-1} \end{bmatrix} = B^j C^{k_t-2n+1}\vec{v}\pi_{k_t}. \quad (22)$$

Hence, we can compute π_i , $\forall i \in [0, n - 2]$,

$$\begin{bmatrix} \pi_i \\ \vdots \\ \pi_{i+n} \end{bmatrix} = B^{n-1-j} C^{k_t-2n+1}\vec{v}\pi_{k_t} \quad (23)$$

By defining $\vec{e}_i = [\pi_i \cdots \pi_{i+n}]^T$, we get

$$\begin{cases} \vec{e}_i = B^{n-1-i} C^{k_t-2n+1}\vec{v}\pi_{k_t}, & \text{if } i \in [0, n - 2]; \\ \vec{e}_i = C^{k_t-n-i}\vec{v}\pi_{k_t}, & \text{if } i \in [n - 1, k_t - n - 1]; \\ \pi_i = \theta_n(1 + \theta_n)^{k_t-i-1}\pi_{k_t}, & \text{if } i \in [k_t - n, k_t - 1]. \end{cases} \quad (24)$$

There just remains to normalize the π_i 's to get the distribution of scrip:

$$(M^*)_i = \frac{\pi_i}{\sum_{j=0}^{k_t} \pi_j}. \quad (25)$$

By multiplying by the fraction of agents of each type, we get the complete characterization of the distribution of scrip:

$$(M^*)_i^t = \frac{f_t \pi_i^t}{\sum_{j=0}^{k_t} \pi_j^t}.$$

■