Formal Approach to Design and Automatic Verification of Cooperation-Based Networks

Alessandro Aldini
University of Urbino "Carlo Bo"
Urbino, Italy
email: alessandro.aldini@uniurb.it

Abstract—**The efficacy and efficiency of cooperation incentives in user-centric networks is a challenging issue that involves tradeoff among trust, social, and economic aspects. Two well-established approaches to stimulate resource sharing and cooperation rely on reputation and remuneration, the complementary functioning of which shall increase users' motivation and discourage mistrust and selfishness. In order to verify the benefits of the joint application of these mechanisms, we specify and analyze formally a recently proposed cooperation model by employing tool-supported probabilistic model checking techniques.**

***Keywords*—model checking, cooperation incentives, trust, remuneration, user-centric networks.**

I. Introduction

User-centric networks (UCNs, for short) emerged among the fastest growing community-scale places of the Internet. They include social networks, online games, auction systems, and peer-to-peer environments. User centricity entails interaction among users, which are expected to share some form of prosocial attitude to cooperation deriving from, e.g., synergy and sense of community. In order to strengthen such an attitude and to ensure the community grows healthily, explicit incentive mechanisms in the form of rewards are used to provide motivations and to contrast typical obstacles of cooperation systems, like mistrust and selfishness. Rewards can be either *indirect*, as in the case of automatically computed notions of trust used to regulate the access to services and resources, or *direct*, e.g., in the form of remuneration, which can be based on fiat money or virtual currency.

Trust systems provide explicit metrics estimating the subjective reliance on the character, integrity, ability, and honesty of each user, thus providing the means for setting up a reputation infrastructure. The aim of reputation is not only giving a perception of the public trustworthiness of each user, but also providing the enabling conditions for exchanging services. On the other hand, virtual currency supports monetization of services, which provides additional motivations to sharing whenever barter, sense of community, and reputation do not suffice.

The application of incentive mechanisms is more critical in wireless and mobile environments. In fact, in these systems, even the underlying communication infrastructure could be dynamically built by users sharing Wi-Fi connections. However, joining the community for short periods of time, thus hindering long-term relationships, may keep such users from adopting prosocial behaviors. Similarly, the inherent limitations of mobile devices (e.g., battery and bandwidth) usually restrict the application of pervasive controls, like assurance of payment or service delivery, thus exposing the system to dishonest behaviors that, however, must be contrasted by the adoption of incentive mechanisms.

Hence, several orthogonal aspects come into play to establish to what extent cooperation incentives can deal successfully with mistrust, selfishness, cheats, and limited resources. In the light of these considerations, the objective of this paper is twofold. On one hand, in order to consider properly the analysis of all the issues above, we advocate the use of tool-supported, model checking based, formal techniques. On the other hand, we apply such techniques for the design and verification of a specific cooperation model for wireless UCNs. This work is a revised and extended version of [1], which is a paper presented at AFIN 2012. The first proposal of the cooperation model under consideration can be found in [2], while its formal design and verification are supported by probabilistic model checking and the related software tool PRISM (see, e.g., [3], [4], [5] for a survey).

The exhaustive analysis conducted in this paper takes into account security, trust, social, and economic aspects in order to verify whether users requiring services are motivated to behave honestly, while users offering services are encouraged to share resources. The formal specification is given in the modeling language of PRISM, which is a state-based mathematical formalism based on the Reactive Modules of [6], from which different types of probabilistic models can be derived, including discrete-time Markov chains (DTMCs, for short) and Markov decision processes (MDPs, for short), see, e.g., [7], [8]. By following the lines of [9], performance properties are expressed in a temporal logic – subsuming both Probabilistic Computation Tree Logic (PCTL) and Linear Time Logic (LTL) – which is expressive enough to specify state-based and path-based properties, and including both probabilistic and reward operators. Thanks to these capabilities, we can describe properties including probabilistic and temporal information, as well as expressing social and economic aspects.

In the rest of the paper, we first discuss some comparison with related work. Then, we recall the cooperation model of [2] and related modeling assumptions (Section II), we discuss the formal specification of such a model (Section III), we present the results of the model checking analysis (Section IV), and we finally draw some conclusions (Section V).

A. Related Work

Sustaining a secure, reliable and efficient environment for the sharing of services and resources in highly mobile communities represents a well-studied problem in the literature, see, e.g., [10], [11], [12], [13]. The application of formal approaches to the analysis of this problem is not novel, refer, e.g., to [14], [15], [16], [17], [18]. In particular, game theory is the most applied approach to networking, see, e.g., [19].

Whenever the intrinsic attitude to prosocial behaviors is not enough, as emphasized in [20], a suitable support to more explicit and extrinsic motivations is given by the joint application of *trust management* (see, e.g., [21]) and *virtual currency* (see, e.g., [22]). Trust management and virtual currency implement the so-called *soft security*, which is characterized by relaxation of the security policies and enforcement of common ethical norms for the community, see, e.g., [23].

Recently, it has been proved that intertwining indirect and direct rewards maximizes the effect of the incentives to cooperate, as shown in [1], [24], [14], [25]. In particular, in [14] game theory is employed to study the balanced tradeoff between reputation-based and price-based cooperation strategies. The obtained analytical results are consolidated by a simulation analysis showing the fast convergence towards cooperative behaviors in the case of mixed incentive strategies. Analogous results are achieved in our approach, which, however, provides a unifying formal framework allowing for the evaluation of all the quantitative properties of interest without requiring simulation analysis. Similarly, the utility-based decision making framework of [24] is used to verify a QoS-based incentive mechanism in which, however, only some of the aspects considered in our approach are taken into account.

II. Cooperation Model

This section briefly outlines the cooperation model introduced in [2], by illustrating the way in which indirect and direct rewards are combined. We then specify the modeling assumptions adopted for analysis purposes.

Cooperation involves users providing services, hereafter called *requestees*, and recipients of such services, hereafter called *requesters*. The cooperation process entails the following four phases:

1) discovery and request;
2) negotiation;
3) transaction and payment;
4) evaluation and feedback.

As we will see, the four phases rely on trust management and virtual currency. In the first phase, the requester searches for a requestee offering the required service. Reputation of the requestee is a parameter guiding the selection. If the requester is trustworthy enough to access the required service, then the issued request can be accepted. However, it may be also refused because of, e.g., lack of willingness to cooperate. In the second phase, requester and requestee establish service parameters and reward, possibly taking into account the trust of the requestee on the requester. In the third phase, service is delivered and the related payment is provided. Finally, in the fourth phase, the transaction results are used to adjust, if necessary, reputation of the involved parties.

A. Reputation System

As usual in several trust-based systems, see [23], we model trust as a discrete metric. Then, given a user i and another user j, the computation of the trust value of i towards j is obtained by mixing direct experience and indirect recommendations:

$$T_{ij} = \alpha \cdot trust_{ij} + (1 - \alpha) \cdot recs_{ij} \tag{1}$$

Parameter $\alpha \in [0, 1]$ represents a risk factor. The trust metric $trust_{ij}$ is the result of previous direct interactions of i with j. In absence of previous experience, the value of $trust_{ij}$ is set to the dispositional trust of i, dt_i, which represents the initial willingness to trust unknown users. Finally, $recs_{ij}$ is the average of the trust metrics towards j of other users (different from i) that in the past negotiated directly with j.

B. Virtual Currency System

Indirect and direct rewards are combined by including the trust value T of the requestee towards the requester as a parameter affecting the cost of the negotiated service. The other parameters are C_{min}, which is the minimum price asked by the requestee regardless of the trust on the requester, C_{max}, which is the maximum price asked to serve untrusted users, and T', which is the trust threshold above which the minimum price is applied to the requester. Then, the cost function C proposed in [2] is defined as follows:

$$C(T) = \begin{cases} C_{min} + \frac{C_{max} - C_{min}}{T'} \cdot (T' - T) & T < T' \\ C_{min} & T \geq T' \end{cases} \tag{2}$$

An alternative simple formula for expressing the service cost is given by the following four-step function:

$$C(T) = \begin{cases} C_{min} & T \geq T_3 \\ C_{med} & T_2 \leq T < T_3 \\ C_{med'} & T_1 \leq T < T_2 \\ C_{max} & T < T_1 \end{cases} \tag{3}$$

where C_{med} and $C_{med'}$ are two intermediate prices, while each T_i, with $1 \leq i \leq 3$, represents a trust threshold, such that $T_i > T_j$ if $i > j$.

C. Modeling Assumptions

For modeling purposes, we distinguish between users playing the role of requester and users playing the role of requestee. Moreover, we consider a unique type of service that is offered by each requestee in the community. Trust values range in the interval $[0, 10]$, such that $null = 0$, $low = 2$, $med = 5$, $high = 8$, and $top = 10$. Based on the system described above, the modeling assumptions concerning the four-phase cooperation process are as follows.

As far as the first phase is concerned, we consider three alternative choice policies adopted by the requester to select a requestee:

- Nondeterministic.

- Prioritized on the basis of requestee's reputation, i.e., the trust value of the requester towards each available requestee is used to govern the selection. The choice among requestees with the same reputation is random.
- Probabilistic, in which case requestee's reputation is used as a probabilistic weight.

The initial reputation is *low* for every requestee. Moreover, by default, requestee i is not available to accept a request by requester j if and only if $T_{ij} < st_i$, where the service trust level st_i represents a trust threshold below which the service offered by i is not accessible. A refused request is sent by the requester to one of the remaining requestees according with the selection policy.

As far as the second phase is concerned, we assume that the agreement is successful. By default, the cost is determined through the application of Equation (2) and is accepted by the requester without any further negotiation. The default values are $C_{min} = 0$, $C_{max} = 10$, and $T' = high$.

As far as the third phase is concerned, by default the service is delivered with success. Then, the requester decides whether to pay or not, either nondeterministically or probabilistically with parameter $p \in [0, 1]$, that is the payment is honored with probability p.

As far as the last phase is concerned, since the service is satisfactory, the reputation of requestee i as perceived by requester j is increased by 1. On the other hand, the trust of i towards j increases (resp., decreases) by 1 (resp., by a factor k) in the case j pays (resp., does not pay) the service. Feedback is provided by i to the other requestees.

III. Specifying Formally the Cooperation Model

In order to illustrate briefly the PRISM specification of the cooperation model, in this section we describe three basic versions of the requester and one basic version of the requestee in a scenario with one requester and two requestees. The reader interested in the evaluation results may skip this part.

Let us start with a system specification with MDP-based semantics, i.e., choices can be nondeterministic. A system component is represented by a module specifying its local variables and its behavior. The requester is defined as follows:

```
module Requester
 x : [0..n] init 0;
 ns : [0..N] init 0;
 ...
```

The local variable x denotes the local state of the requester, such that x=0 represents the initial state and the integer constant n is the number of possible local states. The local variable ns represents the number of requested services, the maximum value of which is given by the constant integer N.

The behavior is given by a set of guarded commands specifying variable updates. In our example, the requester chooses nondeterministically one of the two requestees:

```
 [] x=0 & ns<N -> (x'=11);
 [] x=0 & ns<N -> (x'=21);
 [] x=0 & ns=N -> (x'=1);
 [] x=1 -> true;
```

A pair of brackets represents the start of the command, while the symbol -> is preceded by the boolean guards to satisfy in order to enable the following variable updates. The primed name x' denotes the next value that x assumes by virtue of the update. If not specified explicitly, the other local variables remain unchanged (true stands for no changes). In our example, if x=0 and ns<N then two updates are possible: the former refers to the case in which the requester chooses the first requestee, while the latter is specific for the second requestee. Without loss of generality, we concentrate on the former case in which x is set to 11.

The first requestee is expected either to accept the request and deliver the service, or to refuse the request. In the case of success, the requester decides nondeterministically to pay or not for the obtained service. The commands expressing such a behavior are as follows:

```
 [accept] x=11 -> (x'=12) & (ns'=ns+1);
 [refuse] x=11 -> (x'=13) & (ns'=ns+1);
 [pay] x=12 -> (x'=0);
 [nopay] x=12 -> (x'=0);
 [] x=13 -> (x'=0);
 ...
endmodule
```

Notice that in some cases the brackets marking the start of the command include a label, which expresses an action name on which the module is expected to synchronize with another module, in the same style of, e.g., [26]. In our example, if x=11 and the first requestee is ready to execute a command labeled with the action name accept, then the updates x'=12 and ns'=ns+1 are executed. In this case, the requester decides nondeterministically to synchronize with the first requestee either through action pay or through action nopay, after which the module goes back to its initial state.

On the other hand, a basic description of the requestee behavior is as follows:

```
module Requestee
 y : [0..m] init 0;
 ...
 [accept] (y=0) & (Teq>=st) -> (y'=1);
 [refuse] (y=0) & (Teq<st) -> (y'=0);
 [pay] (y=1) -> (y'=0) &
             (t' = (t<top) ? t+1 : top));
 [nopay] (y=1) -> (y'=0) & (t'=null);
endmodule
```

The local variable y expresses the local state of the requestee, while other local variables are st, modeling the service trust level, and t, modeling the trust of the requestee towards the requester. Parameter Teq is the result of a formula expressing Equation (1). The command labeled with action pay includes a conditional expression that increments t by 1 if the value of such a variable is less than top and assigns to it value top otherwise. Moreover, the update in the command labeled with action nopay expresses the most punishing reaction to a dishonest behavior of the requester.

Now, let us consider a more detailed version of the system with semantics based on DTMC, meaning that choices cannot be nondeterministic and the execution of each command takes one discrete unit of time. In particular, as far as the selection of the requestee is concerned, we assume that the choice is probabilistic and weighted by reputation. Hence, in the requester module we add two more local variables, one for each requestee, storing their reputation values:

```
rep  : [0..top] init low;
rep2 : [0..top] init low;
```

and then we change the selection as follows:

```
[] x=0 -> (rep/totrep):(x′=11) +
          (rep2/totrep):(x′=21);
```

where the syntactic expression $p_1 : (c_1) + p_2 : (c_2)$ indicates that command c_1 is executed with probability p_1, while command c_2 is executed with probability p_2, such that $p_1+p_2 = 1$. Parameter `totrep` represents the overall reputation of the requestees, which is necessary to compute the relative weights, and is defined as the result of the following expression:

```
formula totrep =
   (rep+rep2) = 0 ? 1 : (rep+rep2);
```

Analogously, the other source of nondeterminism in the requester module, which is the choice related to payment, is changed as follows:

```
[accept] x=11 -> p:(x′=12) +
              (1-p):(x′=13);
[refuse] x=11 -> (x′=14);
[pay] x=12 -> (x′=0) &
                (rep′=min(rep+1,top));
[nopay] x=13 -> (x′=0) &
                  (rep′=min(rep+1,top));
[] x=14 -> (x′=0);
```

The real value `p` represents the parameter governing probabilistically the choice between honest and cheating behaviors, while the reputation of the first requestee is increased whenever the service request is accepted.

Alternatively, in order to select a requestee, the requester may follow a prioritized model of choice based on reputation. In such a case, given the following expression:

```
formula maxrep = max(rep,rep2);
```

we change the selection as follows:

```
[] x=0 & rep=maxrep -> (x′=11);
[] x=0 & rep2=maxrep -> (x′=21);
```

If both requestees have the same reputation, then the two alternative commands are given the same weight and the probabilistic choice follows a uniform distribution.

Finally, the properties are specified in an extension of PCTL including reward operators. A reward is a cost associated with state-based and transition-based conditions. Rewards are accumulated each time the related conditions hold, while ad-hoc operators are used to reason about the amount of accumulated rewards, e.g., along specific paths by a given amount of time or at the equilibrium. For instance, the reward structure that is used to calculate the total expected earnings for the first requestee is as follows:

```
rewards "cost1"
[pay] true : f;
endrewards
```

This structure establishes that `pay`-labeled transitions from states enabling the guard `true` acquire a reward equal to the value of formula `f`, which abstractedly denotes, e.g., Equation (2). Based on this structure, as an example we show the property specifying the total expected earnings accumulated until `t` time units have elapsed:

```
R{"cost1"}=? [ C<=t ]
```

In particular, the operator `C<=t` is used to reason about the transient-state behavior of a system.

IV. Model Checking of the Cooperation Model

The analysis of the cooperation process through model checking is divided into two steps. First, we study the vulnerabilities of the trust system with respect to a cheating profile of the requester. The reason is that the soft transaction mechanism does not ensure guarantee of payment. Based on the results of such an analysis, we then verify the efficiency of the mixed cooperation incentives in discouraging selfish behaviors of the requestees and motivating honest behaviors of the requesters.

A. MDP Analysis

The effectiveness of the trust system with respect to cheating requesters is expressed through the following property, which is investigated in a scenario with a single requester and three alternative requestees.

Property 1. What is the maximum number of services (out of N requests) that can be obtained by a requester without honoring the payment?

Formally, the specification of this property is given as a *reachability reward* property:

```
R{"nopayed"}max=? [ F(x=1) ]
```

computing the maximum reward, as defined by the reward structure `nopayed`, accumulated along any path until the condition `(x=1)` holds, which denotes the local state that is reached whenever the maximum amount of requests has been issued. We point out that `F` represents the *eventually* operator of LTL. The reward structure `nopayed` is defined as follows, where the three action names refer to the three requestees:

```
rewards "nopayed"
[nopay1] true : 1;
[nopay2] true : 1;
[nopay3] true : 1;
endrewards
```

With respect to the assumptions of Section II-C, we consider requester's choices to be nondeterministic. Hence, the requester can be viewed as an adversary controlling the way in which the nondeterminism is solved adaptively. The aim of such an adversary is to find out the strategy maximizing the

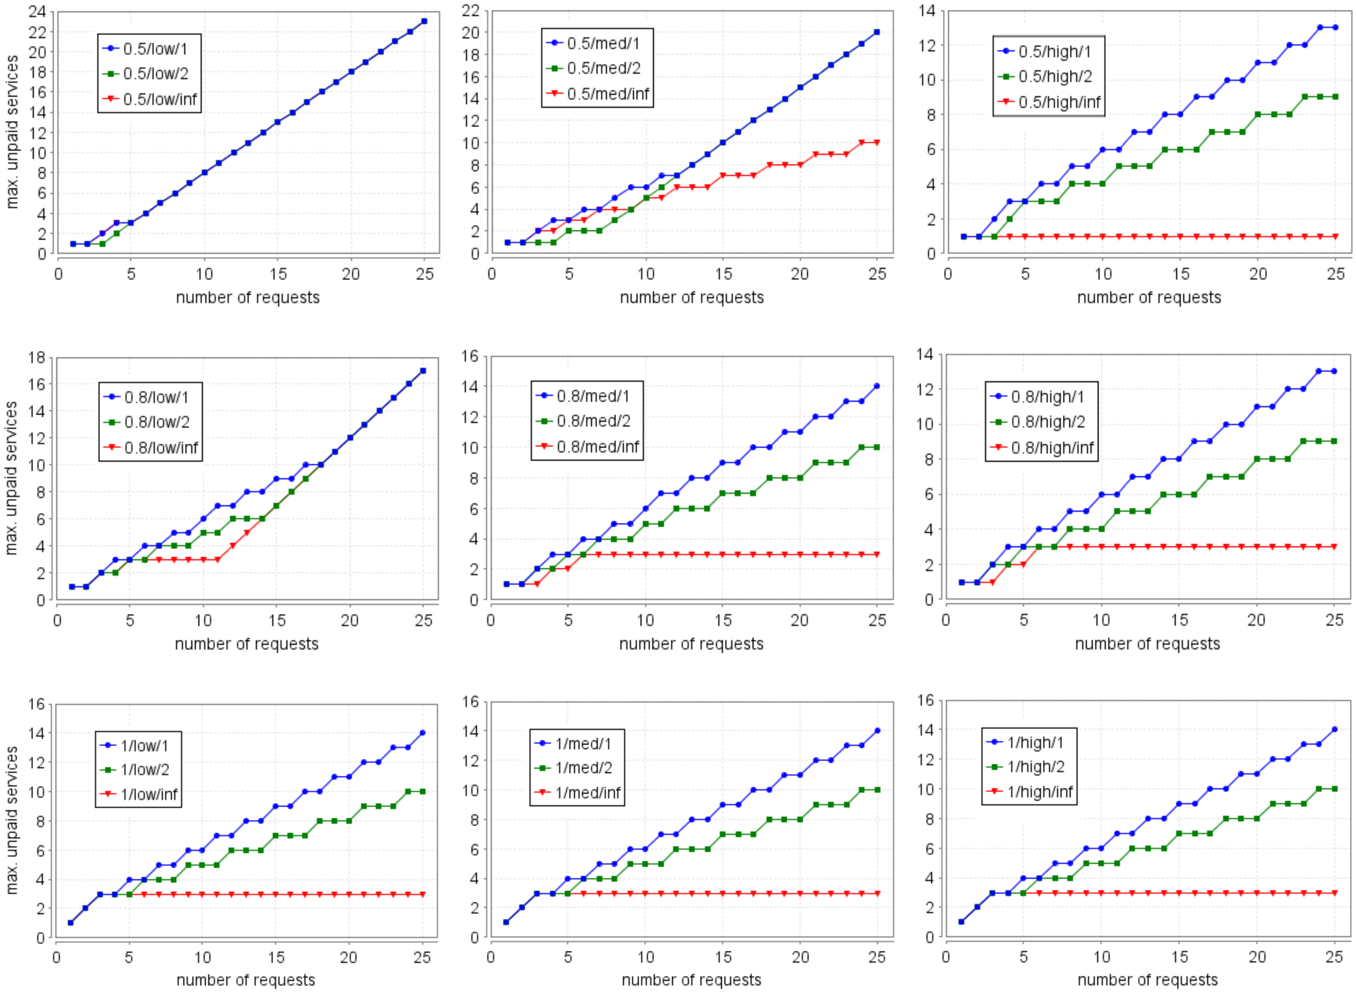

Fig. 1: MDP analysis: verification of Property 1 for 27 combinations of parameters $\alpha/st/k$.

number of unpaid services, thus revealing the worst case from the viewpoint of the requestees.

Formally, the semantics of the model is an MDP on which Property 1 is evaluated by solving the nondeterminism in all possible ways. Then, the model checker returns the result for the *best adversary* strategy. Notice that such a strategy corresponds to the most powerful adversary, which can observe the behavior and the configuration parameters of all the requestees.

To conduct the analysis, we assume three equal requestees characterized by the configuration of parameters $\alpha/st/k$, where: $\alpha \in \{0.5, 0.8, 1\}$ is the contribution of direct experience to trust, *st* $\in \{low, med, high\}$ is the service trust threshold, and $k \in \{1, 2, \infty\}$ denotes the rapidity with which the trust towards a cheating requester is decreased each time a payment is not honored (∞ stands for the immediate assignment of the value *null* to the trust value). The dispositional trust of each requestee is chosen to be equal to the service trust threshold in order to make it possible for a new requester to start negotiating services with the requestees.

All the 27 combinations of the parameters introduced above are analyzed, as illustrated in Figure 1. The horizontal axis denotes the total number of requests *N*, ranging from 1 to 25, while the vertical axis reports the maximum number of unpaid services. From the analysis, we observe that for each value of α and *st* the success of the cheating strategy is inversely proportional to the factor k. In practice, the higher the value of k is, the faster the reaction to dishonest behaviors and, therefore, the negative effect upon trust. For the same reason, the higher the service trust level *st* is, the lower the number of unpaid services. When $\alpha = 1$, however, the service trust level does not affect the results because any decision depends only on previous direct experience. The analysis could be extended to the case $\alpha < 0.5$, obtaining results similar to those related to $0.5/low/\_$, regardless of the value of st. These results reveal a typical attack of a dishonest requester cheating only one requestee, which gives too much weight to the positive

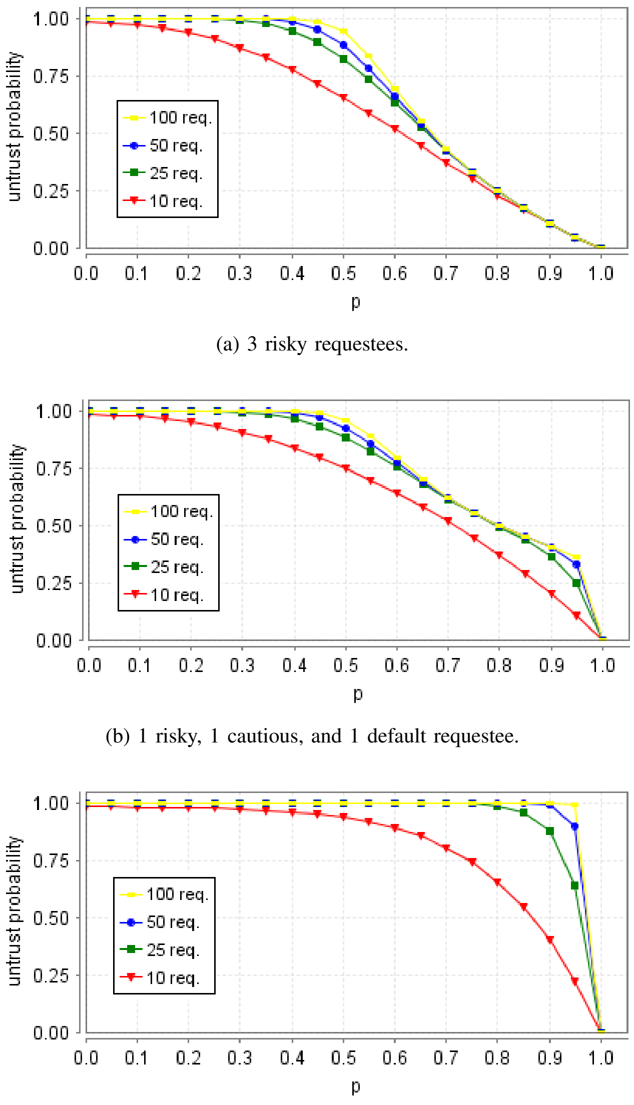

(a) 3 risky requestees.

(b) 1 risky, 1 cautious, and 1 default requestee.

(c) 3 cautious requestees.

Fig. 2: DTMC analysis: verification of Property 2.

recommendations provided by the other requestees.

The results of Figure 1 suggest to categorize the behavior of the requestee according to two limiting profiles:

- *risky* profile, for which the unpaid services increase linearly and most of the served requests are unpaid (see, e.g., configurations $0.5/low/\_$, $0.8/low/\_$, and $\_/\_/1$).
- *cautious* profile, for which the number of unpaid services is essentially constant (see, e.g., configurations $\_/high/\infty$, $0.8/med/\infty$, and $1/\_/\infty$).

B. DTMC Analysis

The two profiles defined above give a clear and precise perception of requestee's attitude to take prosocial decisions in an environment where requesters are possibly cheating. This subsection reports the results of further investigations conducted by considering risky requestees represented by configuration $0.5/low/1$ and cautious requestees represented by configuration $0.8/med/\infty$. Whenever the profile is not specified, configuration $0.8/low/1$ is taken as default.

In order to analyze performance properties, we assume reputation-based prioritized choice of the requestee and payment honored probabilistically with parameter p (see Section II-C). Hence, the semantics of the model turns out to be a DTMC, on which both *steady-state* and *transient-state* analyses can be conducted.

On one hand, the steady-state analysis reveals the success of the cooperation mechanism on the long run. Indeed, at the equilibrium, for each $p < 1$ the requester becomes untrusted with probability 1 by any requestee. On the other hand, the transient-state analysis is important to study the convergence speed towards such a result.

Property 2. What is the probability for a cheating requester of being untrusted by each requestee after N requests?

The specification of this property is given through the operator `P` of PCTL, which is used to reason about the probability of satisfying a given condition. Formally:

```
P=? [ F<=t (x=41) ]
```

returns the probability that the path property expressed between brackets is satisfied by paths starting from the initial state of the system. More precisely, `F<=t` expresses a bounded path property as it imposes the time upper bound `<=t` on the length of the analyzed paths. In our analysis, `t` is chosen to express the constraint upon the number of requests N. On the other hand, the state condition `(x=41)` is associated with a local state of the requester module that denotes the case in which no requestees are available to accept the current request.

We evaluate this property by varying parameter p and by assuming $N \in \{10, 25, 50, 100\}$. Moreover, we consider: (i) three risky requestees (see Figure 2a), (ii) three requestees among which one is risky and one is cautious, while the default configuration is adopted for the third one (see Figure 2b), and (iii) three cautious requestees (see Figure 2c). All the curves tend rapidly to 1 for $p < 0.5$ and converge to zero as p tends to 1. In particular, notice that in the case of three cautious requestees, for $N \geq 25$ the curves approximate a step function, meaning that a cheating requester is almost immediately untrusted by each requestee.

Three more properties are tested in order to investigate the economic aspects of the cooperation mechanism:

Property 3. What is the number of requests accepted by each requestee?

Property 4. What is the total expected earning for each requestee?

Property 5. What is the average earning per accepted request?

For instance, for the first requestee Property 3 is formalized as follows (we can argue analogously for the other properties):

```
R{"acc1"}=? [ C<=t ]
```

where:

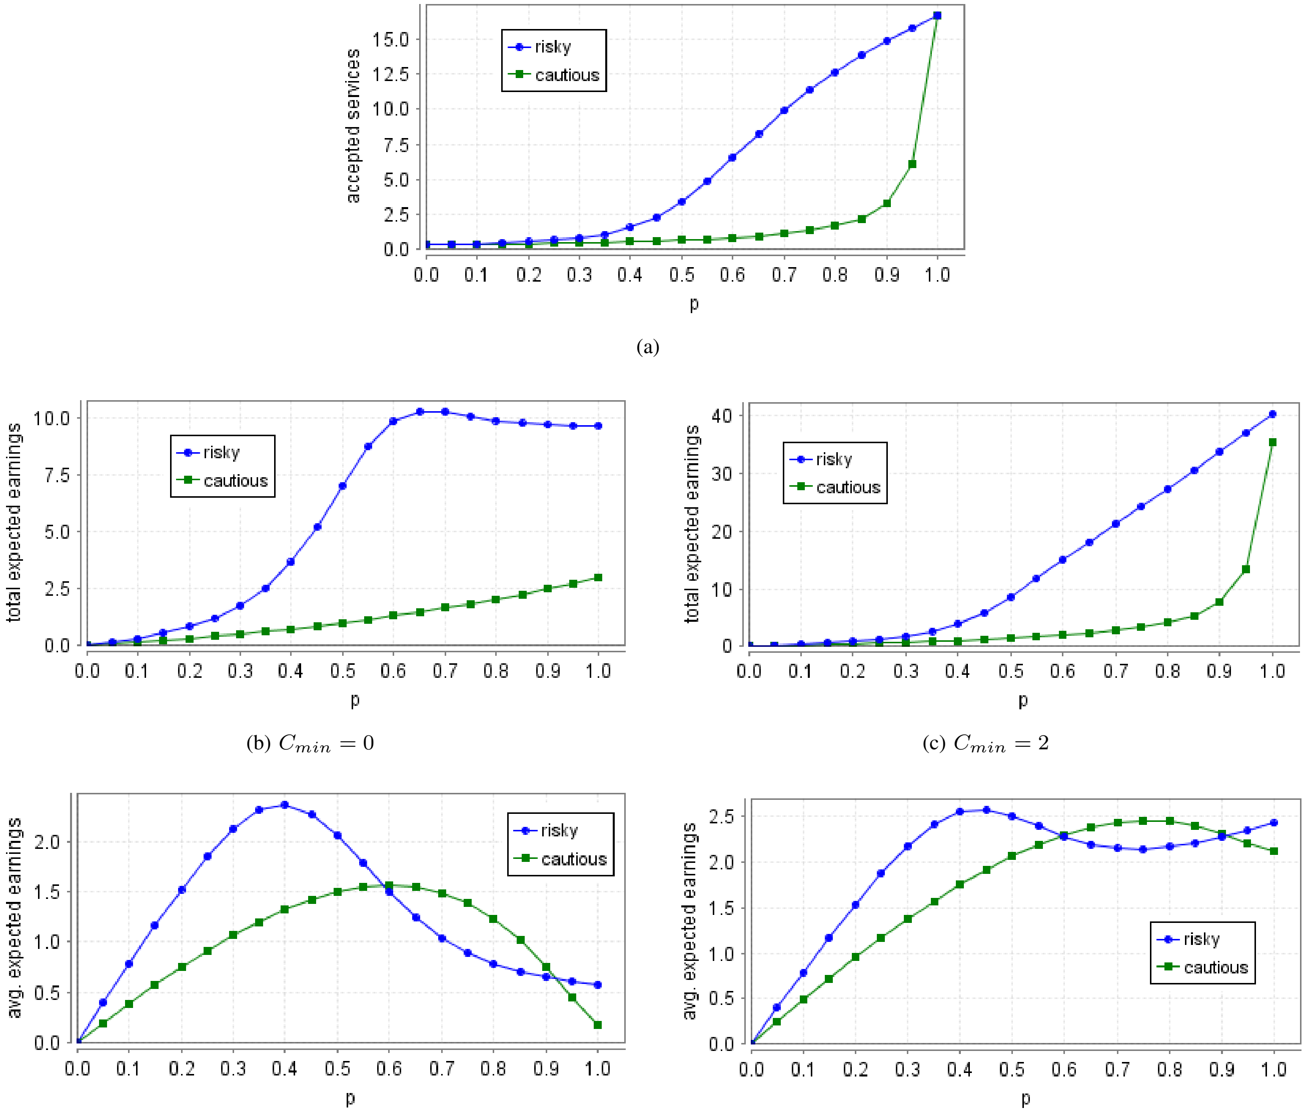

(a)

(b) $C_{min} = 0$

(c) $C_{min} = 2$

(d) $C_{min} = 0$

(e) $C_{min} = 2$

Fig. 3: DTMC analysis: verification of Properties 3, 4, and 5.

```
rewards "acc1"
[accept1] true : 1;
endrewards
```

We use these properties to compare the two profiles in a scenario with 50 requests and three requestees like those of Figure 2b. Figure 3 reports the performance of the risky and cautious requestees as a function of parameter p. The curves show the following results.

The number of services accepted by the risky requestee is higher than that related to the cautious requestee, see Figure 3a. The difference is due to the conditions applied by the risky requestee, in particular the assumption $k = 1$, which is much less restrictive with respect to the assumption $k = \infty$ adopted by the cautious requestee. In fact, by setting $k = \infty$ also for the risky requestee, its curve would collapse with that of the cautious requestee. Notice that in case of honest requester (i.e., $p = 1$), the profile of the requestees does not play any role, so that the requests are equally distributed among them, because they are characterized by the same initial reputation.

As p increases, the total expected earnings of the risky requestee become much higher than those of the cautious one, see Figure 3b. The difference can be interpreted as a reward for taking more risk.

Similarly, Figure 3d shows that the average expected earning per service grows with the value of p up to a maximum point beyond which it decreases because of the effect of the trust-based discount applied to trustworthy requesters. Such a maximum point is reached earlier by the risky requestee, thus motivating the better performance of the cautious requestee for $p \in [0.6; 0.9]$. This result is also confirmed by observing that in such an interval the trust towards the requester becomes stably high from the viewpoint of the risky requestee, as emphasized by the total earnings curve of Figure 3b. For $p \geq 0.95$, the result is better for the risky requestee, because the requester becomes trustworthy also from the viewpoint of the cautious requestee, with a positive impact upon the number of services such a requestee accepts, see Figure 3a.

In general, the combination of remuneration and trust management works as an incentive to adopt a "risky" prosocial behavior. On the other hand, the requester obtains more services at a lower average cost whenever behaving honestly.

C. Sensitivity Analysis

The DTMC analysis of the previous section reveals the tradeoff between indirect and direct rewards, by emphasizing how trust-based mechanisms impact the economic trend of the system. We now investigate more deeply the effect of the various configuration parameters on Properties 3 to 5.

First, we show that the shape of the earnings curves is not purely a side effect of the assumption $C_{min} = 0$. Figs. 3c and 3e report the total and average expected earnings obtained in the case $C_{min} = 2$. The major earnings with respect to the corresponding curves of Figs. 3b and 3d reflect the difference between the minimum costs that are applied in the two experiments.

Second, in order to emphasize the role of parameters k and dt, we tune them for the risky requestee in the same scenario of Figure 3, by showing the related influence upon performance. More precisely, in Figure 4 we vary k in $\{1, 2, \infty\}$, where the case $k = 1$ is taken from Figure 3. Observe that the number of accepted services increases as k decreases. Indeed, as previously shown, k and tolerance to cheating behaviors are inversely proportional. Therefore, decreasing k has the effect of accepting more services, many of which, however, remain unpaid in case of cheating requester. On the other hand, increasing k corresponds to a fast trust decrease and, therefore, higher costs per service. For these reasons, as k decreases, the average expected earnings decrease as well. Also notice that whenever the requester is honest ($p = 1$) and, as a consequence, k is never used, the three curves converge to the same values.

Similarly, we study the effect of tuning the dispositional trust of the risky requestee, by varying dt in $\{low, med, high\}$, where the case $dt = low$ is taken from Figure 3. As shown in Figure 5, increasing the dispositional trust works as an incentive to accept more services as well as to augment the total earnings whenever the requester is rarely honest. The beneficial effect on the total earnings decays as parameter p increases, in which case the most important consequence of increasing dt is a rapid convergence of the service cost towards the minimum cost. Moreover, similarly as observed in the case of parameter k, we have that tolerant behaviors contribute to decrease the average expected earnings.

The consequences of changing the cost function are evaluated in Figure 6, where we propose the same analysis of Figure 3 by replacing Equation (2) with Equation (3), for which we assume $T_1 = low$, $T_2 = med$, $T_3 = high$, while $C_{min} = 0$, $C_{med} = 4$, $C_{med'} = 7$, and $C_{max} = 10$. By comparing the effects of the two equations, notice that while the values change, the shape of the curves is invariant. Indeed, while at the same conditions Equation (3) ensures higher prices than Equation (2), both functions respect the relation between trust and cost.

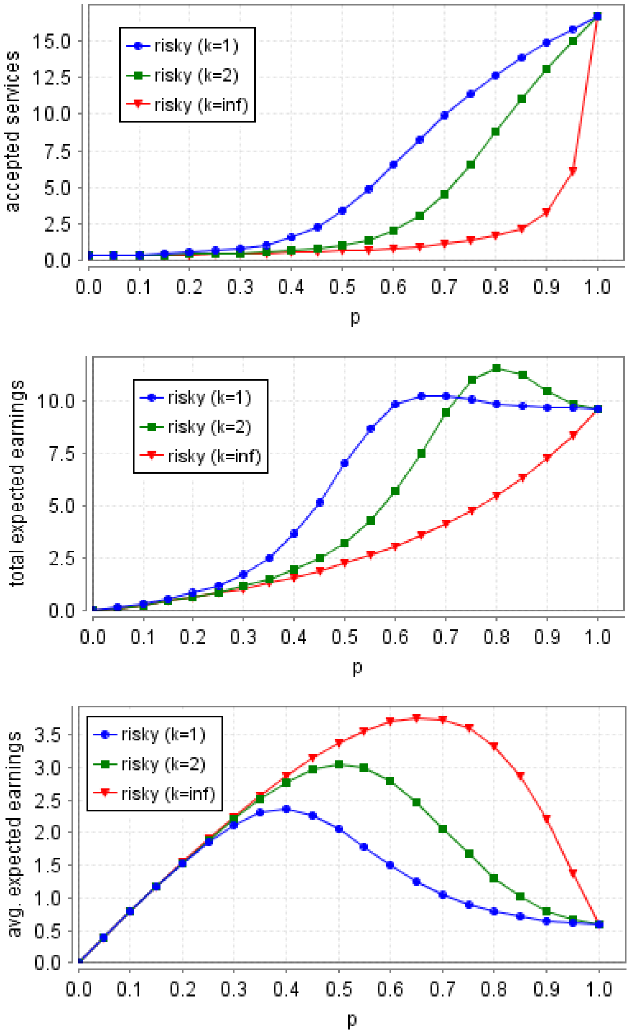

Fig. 4: DTMC analysis: verification of Properties 3 to 5 for the risky requestee by varying parameter k.

Finally, we verify the scalability of results by considering five requestees (four risky and one cautious with the same parameters assumed in the analysis of Figure 3). It is worth comparing the obtained results, see Figure 7, with those of Figs. 3a and 3b. The analogy is emphasized by the fact that the average expected earnings are exactly the same as those of Figure 3d.

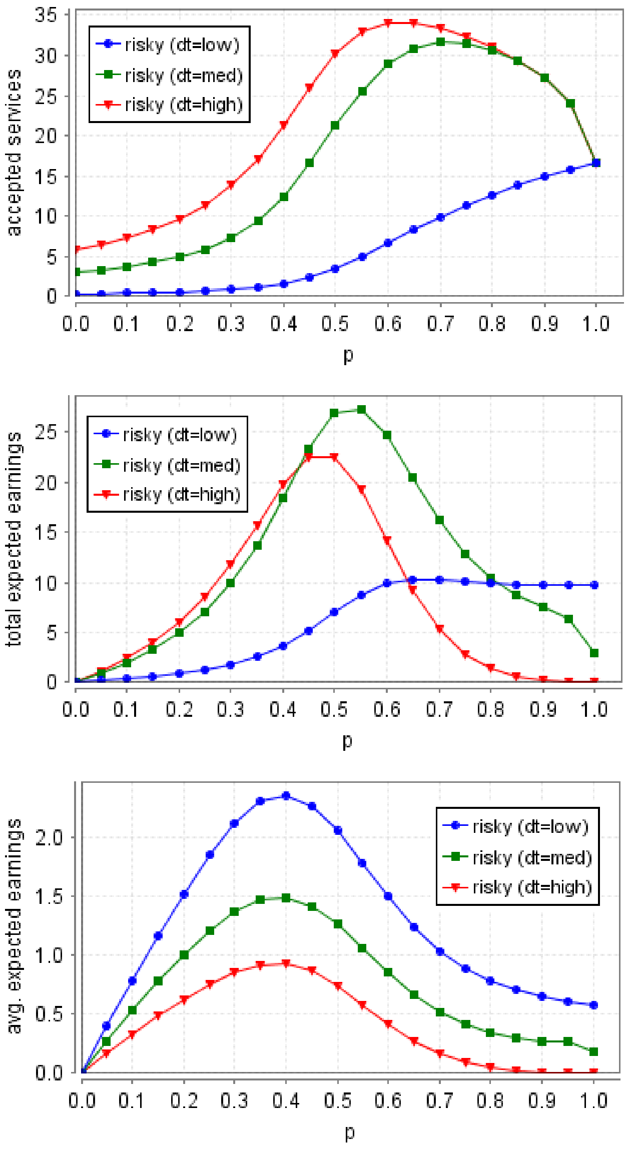

Fig. 5: DTMC analysis: verification of Properties 3 to 5 for the risky requestee by varying parameter dt.

D. Requester's Choice Policy

While so far we have considered security and economic issues for the basic cooperation process, we now take into account the effect of changing the choice policy adopted by the requester to select a requestee in the first phase, which is one of the most important strategies behind the success of the cooperation incentives. Previous results of the DTMC analysis refer to the reputation-based prioritized choice model, which is the policy with the strongest impact of requestee's reputation upon performance.

By assuming the same scenario of Figure 3, in Figure 8 we analyze Properties 3 to 5 whenever the prioritized choice is governed by best price rather than best reputation. In particular, the performance figures refer to the risky requestee under different values of its dispositional trust, because such a parameter is essential for determining initially the service cost, which depends directly on trust. The results confirm the strong influence of the prioritized mechanism and reveal similarities with the experiment of Figure 5. As dt increases, the risky requestee attracts rapidly most service requests, because the applied service cost is inversely proportional to the dispositional trust. The threshold value affecting the shape of the curves is $dt = med$ – which represents the dispositional trust of the cautious requestee – thus revealing the important role of this parameter in the competition among requestees. The relation between dispositional trust and total expected earnings is strict as well. While increasing dt is beneficial for low values of parameter p, the trend is inverse as p increases. Indeed, a high value of dt allows the trustworthy requester to rapidly attain the maximum trust-based discount. For similar reasons, increasing dt cannot have a positive influence upon the average expected earnings.

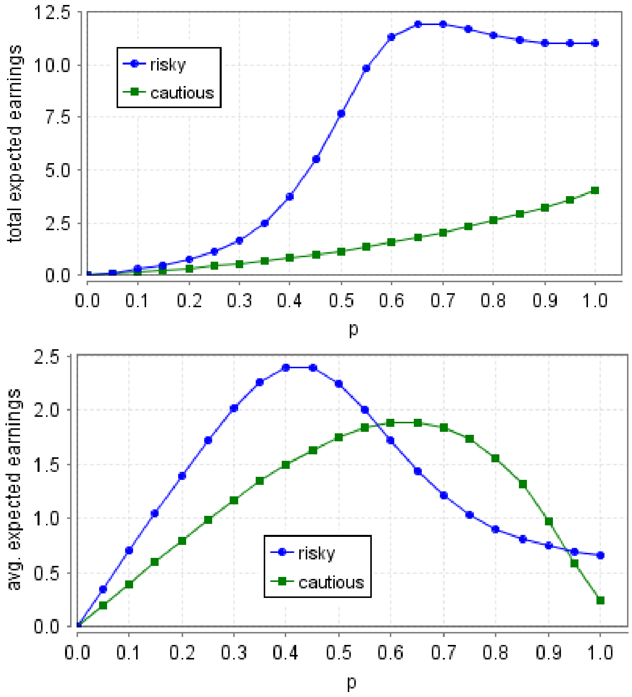

Fig. 6: DTMC analysis: verification of Properties 4 and 5 by using Equation (3).

The analysis concerning the price-based selection is completed by verifying the effects of breaking the relation between cost and trust. In the same scenario of the previous experiment, for the risky requestee we assume $dt = med$ and a constant cost function $C = z$, where z ranges in $\{3, 4, 5\}$. Then, we concentrate on Property 3, see Figure 9. The obtained result is zero services for $z \geq 5$, while the other curves suggest that in order to be competitive with the corresponding curve of

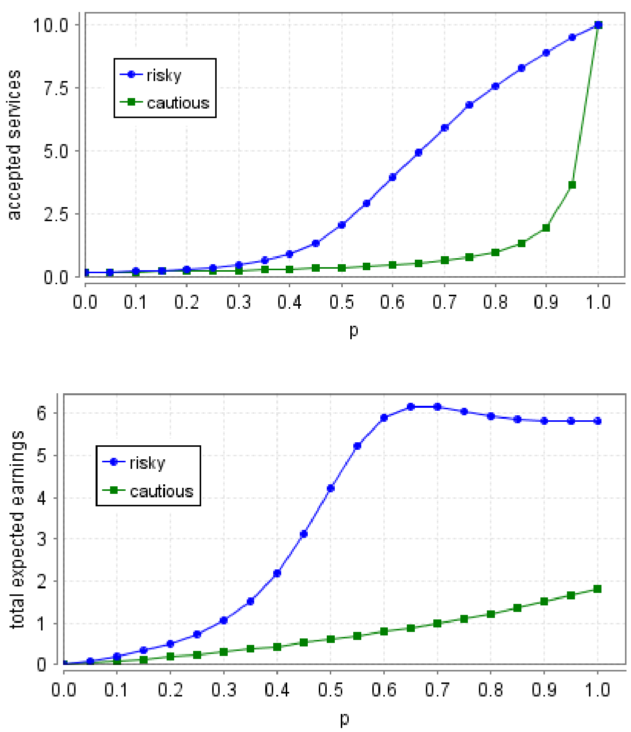

Fig. 7: DTMC analysis: verification of Properties 3 and 4 with 5 requestees.

Figure 8a, it is necessary to set a very low constant price. As a consequence, breaking the relation between cost and trust is not beneficial from the viewpoint of the requestees. On the other hand, from the viewpoint of the requester we have that the average expected cost per service is equal to z independently of the attitude to behave honestly. Instead, by considering the corresponding curve of Figure 8c, we observe that the average expected cost per service converges to a value close to zero as p tends to 1. Therefore, breaking the relation between cost and trust does not work as an incentive to honest behaviors of the requester.

Finally, we replace the reputation-based prioritized selection of the requestee with the reputation-based probabilistic model. In Figure 10, we report the results corresponding to the same scenario of Figure 3. As far as Property 3 is concerned, the difference is negligible, because the choice model adopted by the requester does not affect its trustworthiness as perceived by the requestees. The case of Property 4 and, as a consequence, Property 5, is more interesting. On one hand, in the prioritized model most requests involve the requestee with highest reputation, thus allowing the requester to reach rapidly the minimum service cost as p tends to 1. On the other hand, in the probabilistic model the requests are more equally distributed, thus slowing down the convergence towards the minimum price and justifying the major earnings for the requestees. Therefore, we derive that the prioritized model of choice is more favorable from the viewpoint of the requester.

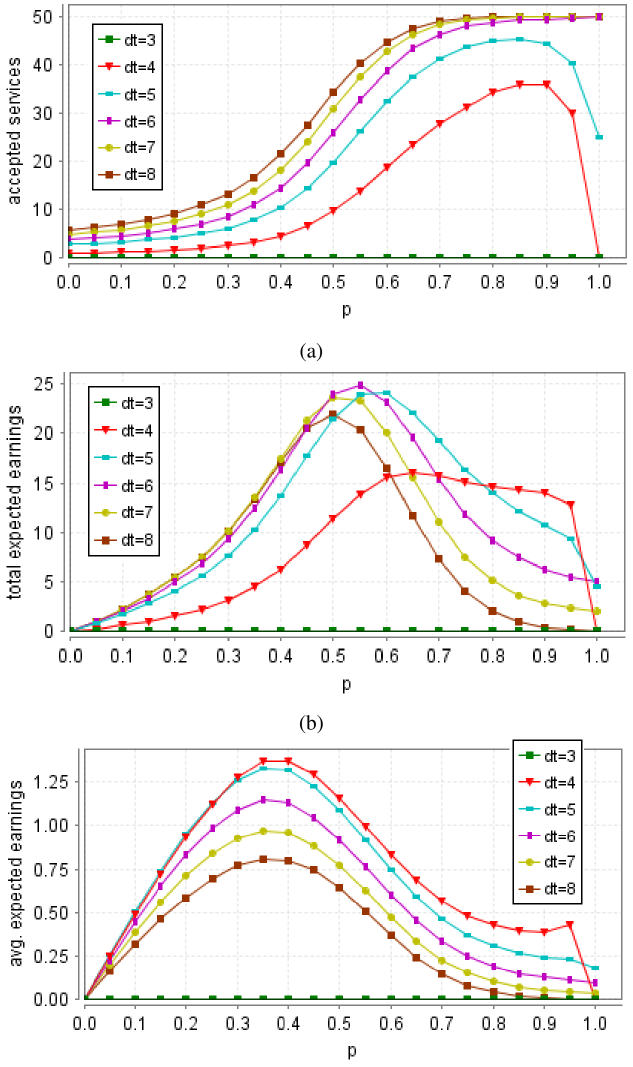

(a)

(b)

(c)

Fig. 8: DTMC analysis: verification of Properties 3 to 5 for the risky requestee with price-based selection of the requestee.

E. *Requestee's Reputation*

Requestee's reputation is an orthogonal aspect the effects of which are analyzed in Figure 11. The objective is to measure the impact of requestee's reputation with respect to Property 3. In Figure 11a, we consider prioritized choice of the requestee, one risky requestee with reputation *high*, one cautious requestee with reputation *low*, while the reputation of the third requestee (with default profile) is *med*. Regardless of the profile, all the requests are served by the requestee with highest reputation, as imposed by the choice strategy followed by the requester. In fact, an analogous result would be obtained by swapping the reputations of the risky and cautious requestees. Giving less importance to reputation during the discovery phase has the effect of mitigating such a drastic

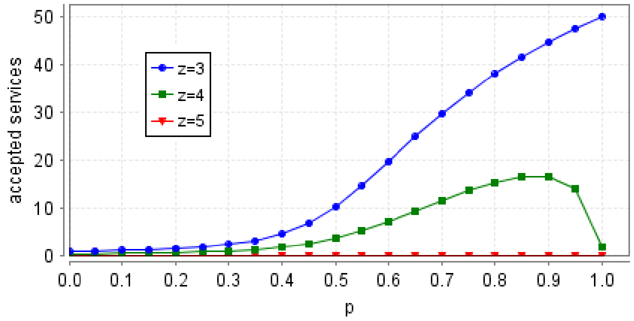

Fig. 9: DTMC analysis: verification of Property 3 for the risky requestee with price-based selection of the requestee and constant cost function.

behavior, as confirmed by the following experiment, in which the prioritized model of choice is replaced by the probabilistic one. The results, shown in Figure 11b, emphasize that also the cautious requestee can obtain some service. However, regardless of the value of p, the cautious requestee is always outperformed by the risky requestee.

The effect of requestee's reputation is investigated also by testing the performance of a paranoid requestee ($\alpha = 0.5$, $dt = low$, $st = med$, $k = \infty$) replacing the cautious requestee in the experiment of Figure 3. In Figure 12a, we evaluate Property 5 for the paranoid requestee in two possible cases depending on its initial reputation. Apparently surprising, a paranoid requestee with reputation *med*, when put in competition with the other requestees (whose reputation is *low*), does not obtain any reward. This result is motivated by the fact that, initially, the paranoid requestee does not accept any request until a sufficiently high number of positive recommendations is received, because its service trust level is higher than its dispositional trust. Moreover, such requests are accepted by the other requestees, which gain reputation, thus causing preemption over the paranoid requestee during the prioritized discovery phase. In order to observe some request served by the paranoid requestee, it is necessary to set its initial reputation to *high*. In this case, we evaluate also Property 3 (see Figure 12b) and Property 4 (see Figure 12c). Notice that the paranoid requestee accepts a very low number of services for $p < 0.9$ and outperforms the risky requestee only for $p = 1$, the reason being that the honest requester becomes trustworthy rapidly enough to overcome the non-cooperative attitude of the paranoid requestee.

F. Impact of Feedback

In this section, we concentrate on the role of feedback for the functioning of the cooperation incentives. On one hand, we analyze the case in which the requester provides a negative feedback. On the other hand, we observe the effects of inaccurate recommendations provided by the requestees.

In a real-world setting, the quality of the delivered service may not match the negotiated parameters. The consequence is a negative feedback of the requester that impairs the reputation of the requestee. This situation is not captured by the experiments reported so far. Hence, we now represent the (possibly negative) change of requestee's reputation due to requester's evaluations in order to check the following property.

Property 6. How is requestee's reputation related to the number of accepted requests in the case of fallible services?

For design issues, we model probabilistically with parameter $q \in [0, 1]$ the event of a service failure causing a negative evaluation. Notice that in the scenario of the previous experiments, modeling an ideal service provider, it holds that $q = 0$. Hence, we consider two additional situations. In a pessimistic case, upon each served request, requestee's reputation has the same probability of remaining unchanged, being increased by 1, or being decreased by 1 (namely, $q = 0.33$). In an optimistic case, the probabilities of these three events are 0.15, 0.8, and 0.05, respectively (namely, $q = 0.05$) For analysis

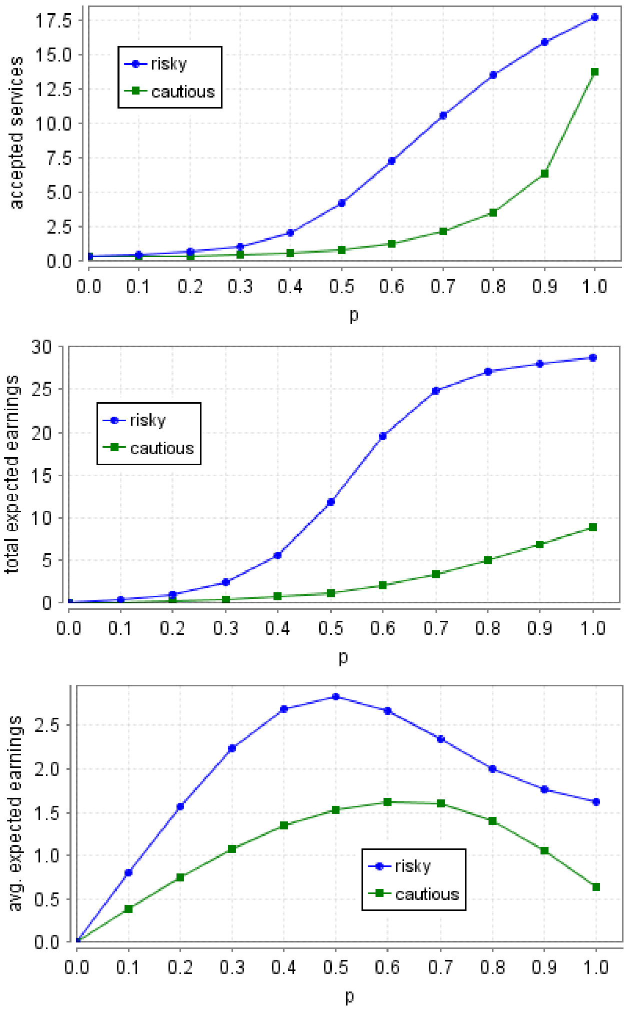

Fig. 10: DTMC analysis: verification of Properties 3 to 5 with probabilistic selection of the requestee.

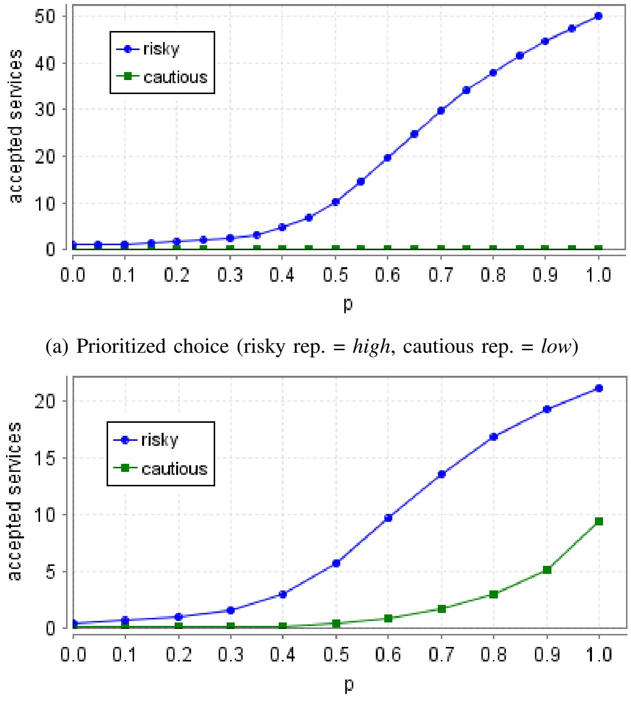

(a) Prioritized choice (risky rep. = *high*, cautious rep. = *low*)

(b) Probabilistic choice (risky rep. = *high*, cautious rep. = *low*)

Fig. 11: DTMC analysis: verification of Property 3 with mixed reputations.

purposes, we consider a honest requester using reputation-based prioritized choice, one cautious requestee with reputation *high*, one requestee with default profile and reputation *med*, and one risky requestee. In Figure 13, we evaluate Property 6 for the risky requestee, by varying its initial reputation from 1 to 10. For $q = 0$, a risky requestee with initial reputation less than *high* is always outperformed by the cautious requestee. The two requestees share the same amount of services if the initial reputation of the risky requestee is *high* as well, while the risky requestee takes all the requests in the remaining cases. These results depend on the deterministic trend of reputations, which never decrease. The other curves approximate such a behavior (the lower q is, the closer the approximation becomes) and reveal that the possibly negative feedback provided by the requester affects the performance of the requestees.

In an orthogonal way with respect to the previous experiment, we now consider the case of non-cooperative requestees, which may refuse a request even if the requester is trustworthy enough to access the service. Hence, the property of interest is as follows.

Property 7. How does requestee's reputation vary in the case of non-cooperative requestees?

As an abstraction, we model probabilistically with parameter $c_i \in [0, 1]$ the cooperative attitude of requestee i, such that i accepts a trustworthy request with probability c_i and refuses it with probability $(1 - c_i)$. Obviously, refusing a trustworthy request is evaluated with a reputation decrease, as opposite to

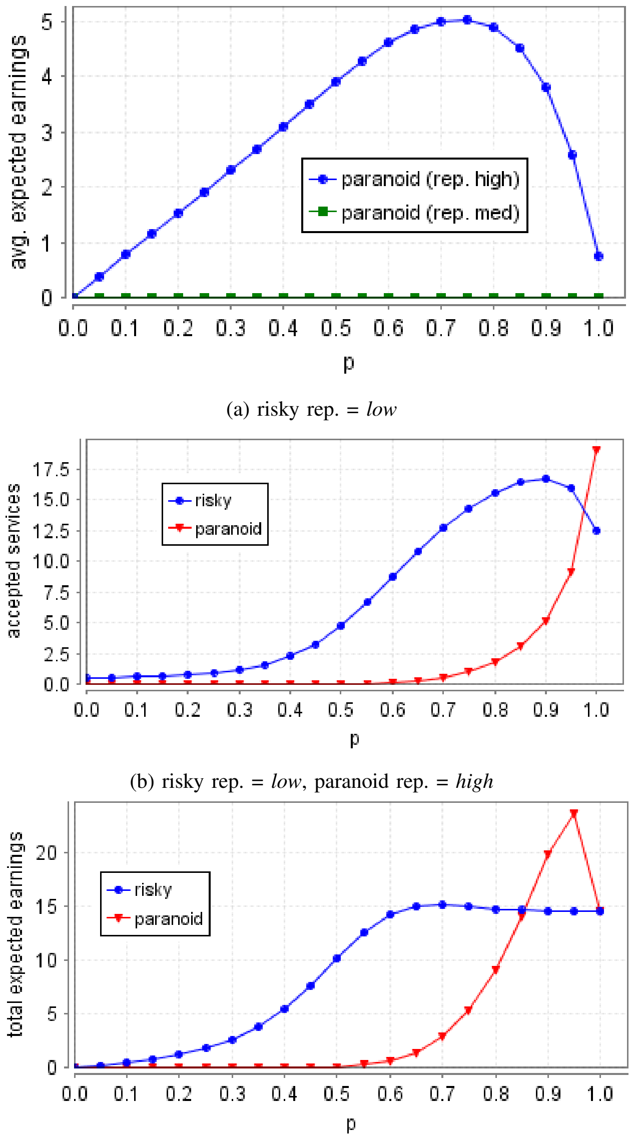

(a) risky rep. = *low*

(b) risky rep. = *low*, paranoid rep. = *high*

(c) risky rep. = *low*, paranoid rep. = *high*

Fig. 12: DTMC analysis: verification of Properties 5, 3, and 4 with paranoid requestee.

the reputation increase determined by a satisfactory service.

For analysis purposes, we consider a honest requester using reputation-based prioritized choice, and three risky requestees with initial reputation *low*. In Figure 14a, we evaluate Property 7 for the first requestee by varying parameter c_1. In particular, we report its average relative reputation variation after 50 requests in two different cases, depending on the behavior of the other two requestees. In the first case, they are fully cooperative (i.e., $c_2 = c_3 = 1$), while in the second case they are partially cooperative (i.e., $c_2 = c_3 = 0.5$). In general, we observe that the lack of cooperation attitude has a negative impact upon reputation, which converges towards the

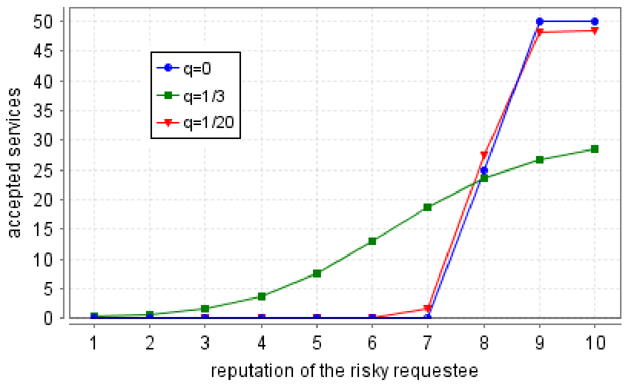

Fig. 13: DTMC analysis: verification of Property 6.

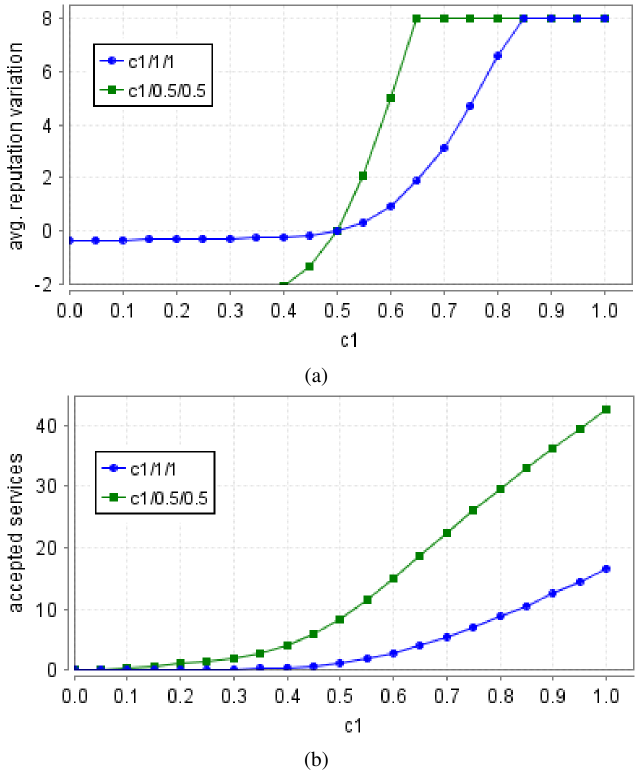

Fig. 14: DTMC analysis: verification of Properties 7 and 3 with non-cooperative requestees.

top level as c_1 increases. We also observe that the reputation variation is slower in the first case with respect to the second case. The reason is that in the first case most services are required to the two cooperative requestees, whose reputation increases rapidly thanks to their prosocial behavior. In order to emphasize the benefits of cooperative behaviors, in Figure 14b we evaluate Property 3 for the first requestee in the two cases above. Notice that in the second case the number of services accepted by the first requestee increases dramatically whenever its attitude to cooperate becomes higher (≥ 0.5) than that of the other two requestees.

The accuracy of feedback is a critical aspect of trust-based incentive mechanisms, as emphasized in [27], where additional incentives are proposed to stimulate the honest and active participation in the evaluation and feedback phase. In Figure 15, we evaluate the effects of inaccurate recommendations on Properties 3 to 5 for the same scenario of Figure 10. In particular, we model with parameter $f \in \{-5, 0, 5\}$ the error introduced to alter the correct recommendations to be provided to other users. The results refer to the risky requestee, whose trust formula is the most influenced one by recommendations ($\alpha = 0.5$). As can be noticed, false positive recommendations have a significant impact, especially for $p \leq 0.5$, as they contribute to increase the trust towards a dishonest requester. On the other hand, false negative recommendations impair the performance, especially for $p \geq 0.5$, as they contribute to keep the requester from obtaining the service. For $p = 1$, the influence of altered recommendations is negligible, because a completely honest requester is trusted enough to get always the service. In general, this analysis confirms the importance of motivating the requestees to provide honest recommendations. On the other hand, we also derive that a honest requester is protected from the feedback variability.

G. Discussion

In summary, cooperation incentives work properly for both the requester and the requestee. For instance, a honest behavior of the requester is motivated by a higher number of accepted services at a lower average cost with respect to the results obtained by a possibly cheating requester. This relation is exacerbated whenever the requester adopts a prioritized model for choosing the requestee during the discovery phase. From the viewpoint of the requestee, both the reputation and the attitude to cooperate affect the amount of delivered services and the related earnings. Moreover, cautious choices for the configuration parameters influencing trust reduce the risk of suffering cheats but impair directly the earning opportunities and indirectly the reputation if in the network cooperative requestees are active.

The sensitivity analysis emphasizes the influence of each policy and configuration parameter chosen by the involved parties. In any case, the results confirm that making cost and trust mutual dependent plays a fundamental role for the success of the cooperation incentives. Similarly, the reliability of trust variations as well as the accuracy of feedback represent important conditions affecting all the performance figures. These relations demonstrate that cooperation incentives provide necessary motivations for the sustainability of collaborative networks.

V. Conclusion

Mixed incentive strategies, combining reputation and price-based mechanisms, have proved to be effective in inducing prosocial behaviors while isolating selfish or cheating nodes,

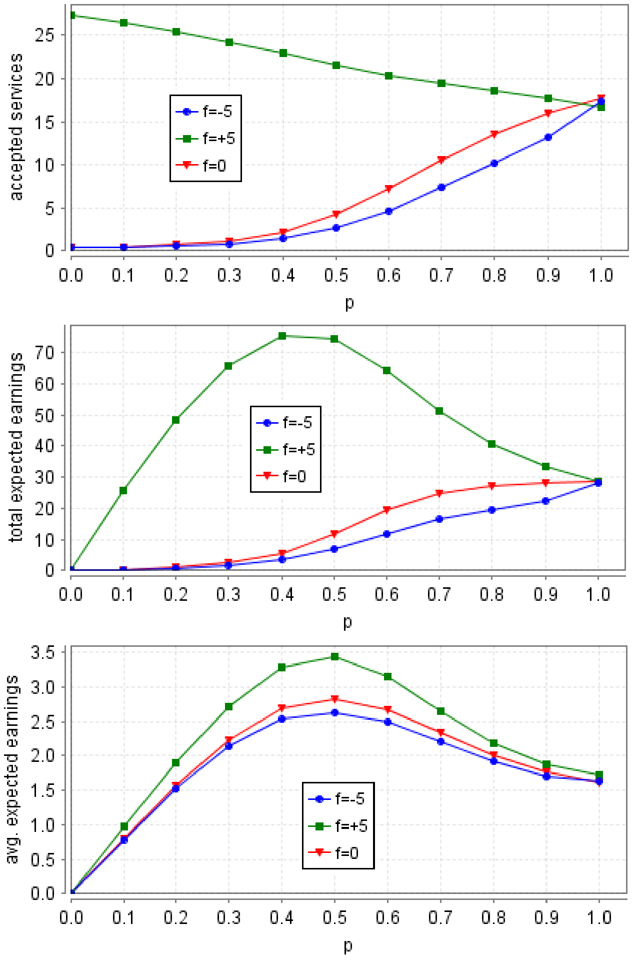

Fig. 15: DTMC analysis: verification of Properties 3 to 5 for the risky requestee with probabilistic selection of the requestee and altered feedback.

as already claimed in [14]. Following such a principle, a cooperation process entailing both trust management and virtual currency has been recently proposed for wireless user-centric networks [2]. This paper has reported the results obtained by applying model checking techniques in order to provide formal evidence of the properties of such a cooperation process.

The same formal approach can be applied to verify the robustness of cooperative networks in more complex environments in which the incentive mechanisms are contrasted by coalition or sybil attacks (see, e.g., [28]). Alternatively, it can be used to evaluate the social, security, and performance effects of the adoption of specific payment systems.

The ideas presented in this work are currently under development in order to build a design tool to be used to assist the design and configuration of mixed incentive strategies in real-world user-centric networks. In particular, the perspectives provided in this paper are under consideration for being adopted by the ULOOP Consortium [29].

Acknowledgment

The research leading to these results has received funding from the EU IST Seventh Framework Programme ([FP7/2007-2013]) under grant agreement number 257418, project ULOOP User-centric Wireless Local Loop.

References

[1] A. Aldini and A. Bogliolo, “Model Checking of Trust-Based User-Centric Cooperative Networks,” Proc. 4th Int. Conf. on Advances in Future Internet (AFIN’12), IARIA, 2012, pp. 32–41.

[2] A. Bogliolo, P. Polidori, A. Aldini, W. Moreira, P. Mendes, M. Yildiz, C. Ballester, and J.-M. Seigneur, “Virtual Currency and Reputation-Based Cooperation Incentives in User-Centric Networks,” Proc. 8th Int. Wireless Communications and Mobile Computing Conf. (IWCMC’12), IEEE Press, 2012, pp. 895–900.

[3] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker, “Automated Verification Techniques for Probabilistic Systems,” in M. Bernardo and V. Issarny, Eds., Formal Methods for Eternal Networked Software Systems (SFM’11), LNCS, vol. 6659, Springer, 2011, pp. 53–113.

[4] M. Kwiatkowska, G. Norman, and D. Parker, “PRISM 4.0: Verification of Probabilistic Real-time Systems,” Proc. 23rd Int. Conf. on Computer Aided Verification (CAV’11), LNCS, vol. 6806, Springer, 2011, pp. 585–591.

[5] M. Kwiatkowska, G. Norman, and D. Parker, “Stochastic Model Checking,” in M. Bernardo and J. Hillston, Eds., Formal Methods for Performance Evaluation (SFM’07), LNCS, vol. 4486, Springer, 2007, pp. 220–270.

[6] R. Alur and T. Henzinger, “Reactive Modules,” Formal Methods in System Design, vol. 15, 1999, pp. 7–48.

[7] W.-J. Stewart, “Introduction to the Numerical Solution of Markov Chains,” Princeton, 1994.

[8] R. Segala, “Modelling and Verification of Randomized Distributed Real Time Systems,” Ph.D. thesis, MIT Press, 1995.

[9] C. Baier and J.-P. Katoen, “Principles of Model Checking,” MIT Press, 2008.

[10] A. Abraham and A.-E. Hassanien (Eds.), “Computational Social Networks: Tools, Perspectives and Applications,” Springer, 2012.

[11] Y. Zhang and M. Guizani (Eds.), “Game Theory for Wireless Communications and Networking,” CRC Press, 2011.

[12] F. Fitzek and M. Katz (Eds.), “Cognitive Wireless Networks,” Springer, 2007.

[13] K. El Defrawy, M. El Zarki, and G. Tsudik, “Incentive-Based Cooperative and Secure Inter-Personal Networking,” Proc. Int. Workshop on Mobile Opportunistic Networking (MobiOpp’07), ACM Press, 2007, pp. 57–61.

[14] Z. Li and H. Shen, “Game-Theoretic Analysis of Cooperation Incentives Strategies in Mobile Ad Hoc Networks,” IEEE Transactions on Mobile Computing, vol. 11(8), 2012, pp. 1287–1303.

[15] U. Golas, K. Hoffmann, H. Ehrig, A. Rein, and J. Padberg, “Functorial Analysis of Algebraic Higher Order Net Systems with Applications to Mobile Ad-Hoc Networks,” Electronic Communication of the European Association of Software Science and Technology, vol. 40, 2010.

[16] S. Nanz and C. Hankin, “A Framework for Security Analysis of Mobile Wireless Networks,” Theoretical Computer Science, vol. 367, 2006, pp. 203–227.

[17] V. Srivastava, J. Neel, A. MacKenzie, R. Menon, L. DaSilva, J. Hicks, J. Reed, and R. Gilles, “Using Game Theory to Analyze Wireless Ad Hoc Networks,” IEEE Communications Surveys and Tutorials, vol. 7(4), 2005, pp. 46–56.

[18] A. Acquaviva, A. Aldini, M. Bernardo, A. Bogliolo, E. Bontà, and E. Lattanzi, “Assessing the Impact of Dynamic Power Management on the Functionality and the Performance of Battery-Powered Appliances,” Proc. 5th Int. Conf. on Dependable Systems and Networks (DSN’04), Performance and Dependability Symposium, IEEE Press, 2004, pp. 731–740.

[19] E. Altman, T. Boulogne, R. E. Azouzi, T. Jimenez, and L. Wynter, “A Survey on Networking Games in Telecommunications,” Computers and Operations Research, vol. 33(2), 2006, pp. 286–311.

[20] C.H. Declerck, C. Boone, and G. Emonds, “When Do People Cooperate? The Neuroeconomics of Prosocial Decision Making,” working paper of the Faculty of Applied Economics, University of Antwerp, 2011.

[21] S. Marsh, “Formalizing Trust as a Computational Concept,” Ph.D. thesis, Department of Mathematics and Computer Science, University of Stirling, 1994.

[22] S. Greengard, “Social Games, Virtual Goods,” Communications of the ACM, vol. 54(4), 2011, pp. 19–22.

[23] A. Jøsang, “Trust and Reputation Systems,” in A. Aldini and R. Gorrieri, Eds., Foundations of Security Analysis and Design IV (FOSAD’07), LNCS, vol. 4677, Springer, 2007, pp. 209–245.

[24] M. Yildiz, M. A. Khan, F. Sivrikaya, and S. Albayrak, “Cooperation Incentives Based Load Balancing in UCN: A Probabilistic Approach,” Global Communications Conf. (GLOBECOM’12), IEEE Press, 2012, pp. 2746–2752.

[25] Y. Zhang, L. Lin, and J. Huai, “Balancing Trust and Incentive in Peer-to-Peer Collaborative Systems”, Journal of Network Security, vol. 5, 2007, pp. 73–81.

[26] A. Aldini, M. Bernardo, and F. Corradini, “A Process Algebraic Approach to Software Architecture Design,” Springer, 2010.

[27] A. Fernandes, E. Kotsovinos, S. Ostring, and B. Dragovic, “Pinocchio: Incentives for Honest Participation in Distributed Trust Management,” Proc. 2nd iTrust Conf., LNCS, vol. 2995, Springer, 2004, pp. 63–77.

[28] F. G. Marmol and G. M. Perez, “Security Threats Scenarios in Trust and Reputation Models for Distributed Systems,” Computer and Security, vol. 28, 2009, pp. 545–556.

[29] ULOOP, “EU IST FP7 ULOOP: User-Centric Wireless Local Loop,” 2013, [accessed 20-June-2013]. [Online]. Available: http://uloop.eu