Engineering Incentive Schemes for Ad Hoc Networks–

A Case Study for the Lanes Overlay

Philipp Obreiter, Birgitta König-Ries, Georgios Papadopoulos
{obreiter,koenig}@ipd.uni-karlsruhe.de, diane@zorz.de

February 23, 2004

Technical Report Nr. 2004-4

University of Karlsruhe
Faculty of Informatics
Institute for Program Structures and Data Organization
D-76128 Karlsruhe, Germany

Abstract

In ad hoc networks, devices have to cooperate in order to compensate for the absence of infrastructure. Yet, autonomous devices tend to abstain from cooperation in order to save their own resources. Incentive schemes have been proposed as a means of fostering cooperation under these circumstances. In order to work effectively, incentive schemes need to be carefully tailored to the characteristics of the cooperation protocol they should support. This is a complex and demanding task. However, up to now, engineers are given virtually no help in designing an incentive scheme. Even worse, there exists no systematic investigation into which characteristics should be taken into account and what they imply. Therefore, in this paper, we propose a systematic approach for the engineering of incentive schemes. The suggested procedure comprises the analysis and adjustment of the cooperation protocol, the choice of appropriate incentives for cooperation, and guidelines for the evaluation of the incentive scheme. Finally, we show how the proposed procedure is successfully applied to a service discovery overlay.

Contents

1 Introduction . . . 1
2 Cooperation Models and Incentive Schemes . . . 1
3 Related Work . . . 2
4 The Procedure for Engineering Incentive Schemes . . . 3
4.1 The Cooperation Protocol . . . 3
4.2 The Choice of Appropriate Incentives . . . 4
4.3 Evaluation . . . 6
5 A Case Study: Lanes . . . 8
5.1 The Lanes Protocol . . . 8
5.2 The Choice of Appropriate Incentives . . . 9
5.3 Evaluation . . . 10
6 Conclusion . . . 12

1 Introduction

In ad hoc networks, devices have to cooperate in order to make up for the absence of infrastructure. However, each participating device is under the control of its user and, thus, aims at maximizing its utility. This means that devices will only cooperate if this is profitable for them. Most often, cooperation is not profitable in itself. Therefore, distributed schemes have been proposed which offer incentives for cooperation, thereby making it attractive for devices to cooperate.

Such incentive schemes make use of incentive mechanisms in order to foster cooperation. An example of an incentive mechanism is a distributed reputation system [1]. Yet, the choice and configuration of appropriate incentive mechanisms is highly non-trivial. This is partly due to the dependency on the specifics of the application domain. Currently, there is no systematic procedure that supports the developer of an incentive scheme by managing the complexity of his task. As a result, the conception and evaluation of incentive schemes is still more approached as an art than a engineering principle. This situation is especially harmful since each of the various cooperation protocols of ad hoc networks demands for a specific incentive scheme that takes its characteristics into account.

Our approach is to systematically engineer incentive schemes. In Section 3, we discuss the state of the art for the development of incentive schemes. In Section 2, we take a closer look at appropriate models of cooperation and at the scope of incentive schemes. This provides the foundation for presenting and discussing a systematic procedure for engineering incentive schemes in Section 4. We exemplify such engineering for the cooperation protocol *Lanes* in Section 5 and, finally, conclude the paper in Section 6.

2 Cooperation Models and Incentive Schemes

In this section, we give an overlook of the basic concepts that are required for the development of incentive schemes. For this purpose, we take a closer look at cooperation models and point out the necessity of incentive schemes.

Cooperation models. In our previous work [2, 3], we proposed a transaction-centric cooperation model. It assumes that cooperation is composed of transactions among autonomous protocol entities. Each entity may commit to participate in one or several transactions.

The transaction-centric model makes sense for the provision and consumption of application services. However, it is difficult to apply the transaction-centric model to cooperation in networks and overlays. This stems from difficulties of capturing continuous cooperation among several entities into the transaction-centric model. We conclude that a more generic cooperation model is needed.

According to the commitment-centric cooperation model, an entity *enters into a commitment* if it commits to exhibit certain behavior. An entity *adheres to its commitment* if it actually exhibits the behavior it committed to. The transaction-centric model is a specialization of this cooperation model since it confines commitments to single transactions. More specifically, participating in a transaction means entering in a commitment, whereas adhering to a commitment refers to refraining from defecting in the course of a transaction.

Incentive schemes. A rational entity, i.e., an entity that aims at maximizing its profit, only decides to enter into a commitment if it is beneficial for the attainment of the own goals. In this context, the benefits of entering into a commitment have to be assessed. The main problem arises from the fact that such assessment also has to consider the likeliness of getting hold of the benefits. They accrue from the actions of the autonomous transaction partners who do not necessarily adhere to their commitments. As a result, each entity is suspicious of the benefits that arise from entering into commitments and adhering to them.

The conventional solution to this problem is to make use of a central authority that is able to enforce the adherence to commitments. However, such central authority represents an infrastructural component. Hence, it contradicts the paradigm of self-organization that is paramount to ad hoc networks. Therefore, there is a need for a distributed scheme that provides incentives for entering into commitments and adhering to them.

3 Related Work

In this section, we discuss the state of the art for the development of incentive schemes in ad hoc networks. Furthermore, we discuss incentive engineering as an approach that is used in economics for the conception of appropriate incentives.

Existing incentive schemes in ad hoc networks. In the absence of any systematic procedure for their development, the design of the existing incentive schemes [4, 5, 6, 7, 8, 9] is characterized by the ex ante choice (and configuration) of incentive mechanisms [2]. For the development of further incentive schemes, the usefulness of the existing incentive schemes is limited. This stems from the following reasons:

- They are bound to specific cooperation protocols, often without making this explicit.
- Their conception is monolithic and, thus, hinders the reuse of their components.
- Their evaluation is not performed on the basis of comparable criteria.

This means that it is unlikely that a developer can simply reuse one of the existing schemes in order to enhance cooperation in a given situation. On the other hand, he is given little if any help in designing his own incentive scheme. Thus, despite the fact that incentive schemes need to be tailored to the cooperation protocol used, little is known on how to achieve this tailoring.

Incentive engineering. In economics, incentive engineering [10] has been proposed as a means of systematically developing incentive schemes. It assumes an incentive mechanism that is arbitrarily quantifiable and provides full incentive compatibility. For example, the use of money provides such an incentive mechanism. For each action of the cooperation protocol, the engineer determines the quantification of the incentive mechanism that yields a maximization of some utility. This approach has been applied in [11] for the development of an incentive scheme on the link layer.

However, incentive engineering is not suitable for the development of incentive schemes in ad hoc networks. This stems from the following reasons:

- In our previous work [3], we have shown that, in ad hoc networks, it is impossible to conceive an incentive mechanism that it is both arbitrarily quantifiable and fully incentive compatible.

- The exogenous determination of the quantification is contradicted by the autonomy of the devices. Consequently, the devices do not adhere to the developed scheme.
- In ad hoc networks, it might be reasonable to adjust the cooperation protocols in order to facilitate the development of the incentive scheme. However, incentive engineering does not make use of this means.

4 The Procedure for Engineering Incentive Schemes

The systematic design of incentive schemes comprises several steps. As a first step, the engineer analyzes and adjusts the cooperation protocol that requires an incentive scheme. Subsequently, crucial design decisions have to be made regarding the choice of incentives and the means of implementing them. Finally, the engineer evaluates the resulting cooperation protocol by applying an appropriate evaluation method. In the following, we give an in-depth discussion of these steps. The discussion remains rather abstract since it will be exemplified by the case study of Section 5.

4.1 The Cooperation Protocol

Analysis. A thorough understanding of the cooperation protocol is required for the design of an appropriate incentive scheme. Therefore, an analysis has to be conducted in order to answer the following questions:

- **(Q1)** Which entities are considered? Existing incentive schemes [5, 9] are focussed on network protocol entities. Yet, it might make sense to include protocol entities from different protocol layers [12].
- **(Q2)** Are the entities rational or tamper resistant? The incentive scheme does not consider incentives for tamper resistant entities since they cannot exhibit autonomous behavior.
- **(Q3)** Do the entities have identities and is there a means of reliable authentication? If entities are able to create new identities, some mechanisms of incentive schemes become less effective. In addition, most mechanisms cannot be applied if spoofing[1] is possible.
- **(Q4)** How do the entities valuate the utility and the costs that arise from cooperation? The answer to this question indicates which kind of misbehavior is to be expected from rational entities.
- **(Q5)** How heterogeneous are the entities with regard to such valuation? Such heterogeneity is to be expected for ad hoc networks since the capabilities of the participating devices largely differs. In case of such heterogeneity, entity-specific misbehavior is to be expected.

After the analysis of the entities that take part in the ad hoc network, the specifics of the cooperation protocol have to be examined:

- **(Q6)** What kinds of inter-entity cooperation exist? The answer to this question is determined by the cooperation protocol that the entities run.

[1] Spoofing refers to the ability of assuming the identity of other entities.

- **(Q7)** Which steps of the cooperation protocol are not beneficial to the executing entity? The answer to this question is determined by the valuation of the entities (Q4) and the cooperation protocol (Q6). The engineer focusses on influencing behavior regarding detrimental protocol steps.

- **(Q8)** Is behavior perceptible? If yes, how costly and reliable is such perception? It is clear that a rational entity only exhibits cooperative behavior if other entities are able to perceive such cooperativeness.

Adjustment. Before turning to the incentive scheme, the engineer has to consider whether the protocol should be adjusted. Such adjustment aims at ameliorating the protocol's properties with respect to the subsequent application of an incentive scheme. More specifically, the necessity of detrimental protocol steps (Q7) should be revisited. Ideally, the cooperation protocol remains effective even if such protocol steps are not executed. Furthermore, the protocol should be extended by mechanisms that make behavior more perceptible (Q8).

The engineer may choose among several perception mechanisms. If *digital signatures* are appended to protocol messages, the receiver of a message is able to verify the authenticity of the sender and to check whether the entities of the forwarding path altered the message. *Redundancy* enhances the perceptibility of specific protocol steps. For example, the cooperation protocol could be extended in order to accommodate the issuance of *receipts*. If a receipt is a non-repudiable evidence [1], it may be transferred to other entities. Consequently, these entities are able to perceive the behavior that is described by the receipt. A cost-effective perception mechanism is *overhearing*. Under certain conditions, an overhearing entity perceives which packets or messages are sent by other entities. However, the effectiveness of overhearing is difficult to assess. Due to physical or topological restrictions, the overhearing entity might not receive the same transmissions or messages as the intended receiver. *Probing* is based on the idea that an entity behaves similarly under certain conditions. Therefore, it suffices to perceive only parts of the behavior in order to conjecture which behavior is typically exhibited. Probing is attractive for behavior that is costly to perceive.

4.2 The Choice of Appropriate Incentives

Based on the analysis and adjustment of the cooperation protocol, the engineer has to choose incentives that effectively stimulate cooperative behavior. The definition of commitments sets the scope of the respective incentive mechanisms. Some of these mechanisms provide incentives for entering into commitments, whereas others provide incentives for adhering to them.

Definition of commitments. A major design decision for incentive schemes consists of determining which type of behavior should be remunerated and which type should be taken as granted. Only misbehavior regarding the latter type has to be punished since, for the first type of behavior, the absence of remuneration constitutes a disincentive for misbehavior. Based on the commitment-centric cooperation model, it seems promising to remunerate if an entity commits to specific behavior. Accordingly, failure to adhere to such a commitment should be punished. For this purpose, it has to be determined how a defecting entity is treated.

The two extremes of defining commitments are as follows:

1. Participation in the system is the only commitment. This means that, upon participation, each entity has to adhere to the predefined protocols and transactions. This extreme fits

best to cooperation protocols that require unconditional cooperativeness in order to be effective. For example, for some topologies on the network layer, the cooperativeness of routing bottlenecks is necessary in order to avert partitioning.

2. Each transaction is separately committed to. In this regard, a rational entity only participates in beneficial transactions. This extreme is suitable if cooperation can be decomposed into transactions, as it is true on the application layer.

The first extreme allows for rather simple incentive schemes. However, the incentive compatibility of adhering to the own commitment is difficult to achieve. This is especially true if behavior is not perceptible. Then, the lack of adherence cannot be identified and punished. Consequently, imperceptible behavior should be exempted from commitments.

The definition of commitments itself has to comply with the demand for incentive compatibility and configurability. Otherwise, the entities will agree on collectively altering the proposed definition of commitments. Configurable commitments provide a means of considering the heterogeneous preferences of the entities.

Incentives for entering into commitments. In our previous work [3], we have proposed incentive patterns as a means of rendering transactions mutually beneficial. This is necessary if a transaction appears to be beneficial only for some of the concerned entities. In such a case, the incentive patterns provide a means of convincing the remaining entities of the benefits of the transaction. These patterns can be generalized by basing them on the notion of commitments. This means that an incentive pattern induces that an entity enters into a otherwise detrimental commitment. The incentive patterns and their most important properties are summarized in Table 1. It provides a sound foundation for the choice of appropriate incentives that remunerate for commitments. For each type of commitment, one or several incentive patterns may be applied.

Basically, incentive patterns fall into two classes: On the one hand, an entity may believe that its peers will reciprocate by entering into future commitments. This approach is based on *trust* and is applied by the collective pattern and community pattern. On the other hand, an entity is convinced to enter into a commitment if its peers enter into commitments that are beneficial for itself. Such a *trade* based approach is applied by the barter trade pattern and the bond based incentive patterns. For the barter trade, the temporal scope of the respective commitments coincides. Hence, the committing entities assume symmetric roles. In contrast, a bond is a commitment regarding behavior at some future point in time. The most relevant type of bonds are notes. A note contains a commitment of its issuer.

Incentives for adhering to commitments. Bilateral or multilateral commitments often refer to the mutual provision of services. From an abstract point of view, such mutual provision of services represents an exchange of items. Exchanges are processed according to *exchange protocols* [13]. They typically assert a weaker notion of atomicity. More specifically, an entity has to hand over its own item in order to acquire the desired item of its peer. This means that an entity has to adhere to its commitments if it wants that the other entities adhere to theirs, too. Such *coupling* provides an incentive for adhering to one's own commitments. However, the considerable overhead of exchange protocols has to be matched by the value of the items that are exchanged. For the repeated exchange of items, the sliding window mechanism is a promising exchange protocol. It limits the number of outstanding items. This means the entities' balance of delivered items is coupled.

Table 1: Overview of incentive patterns and their properties

Properties \ Incentive pattern	Trust based		Trade based	
	Collective pattern	Community pattern	Barter trade pattern	Bond based pattern (eg.: note pattern)
Roles	asymmetric		symmetric	asymmetric
Type of remuneration	*none*	reputation	service in return	note
Enforcement of remuneration		-	+	o
Overhead	small	medium	small	high
Scalability	- -	-	+	o

If exchange protocols are not viable, there exists another means of inducing the adherence to commitments, namely *Distributed reputation systems*. They keep track which entities adhere to their commitments and which do not [1]. In contrast to exchange protocols, an entity is able to defect by refraining from adhering to the own commitment while the peers adhere to theirs. However, the betrayed peers may disseminate their view of the defector to other entities so that other entities are aware of the defection. As a result of such awareness, other entities (and the defected peers) may refrain from entering into commitments with the defector or they may refuse to adhere to their outstanding commitments. Such punishment provides an incentive for adhering to one's own commitments. The choice of the distributed reputation system is contingent upon the characteristics of the considered cooperation protocol. On the one hand, distributed reputation systems make differing assumptions regarding identities, encryption, and overhearing. On the other hand, they find different tradeoffs between the benefits and costs that arise from operating them.

Furthermore, in heterogeneous application environments, the engineer has to choose a reputation system that leaves room for configurability with respect to the quality and costs of the trust domain, trust contexts, and trust assessment.

4.3 Evaluation

After the design of the incentive scheme, the engineer has to evaluate her work. In general, simulations provide the only cost-efficient means for such evaluation. In the following, we discuss the degrees of freedom with respect to the realization of such simulations.

Objective of the evaluation. Evaluations of existing incentive schemes [5, 9] are focussed on the total utility of the participating entities. In addition, we propose to evaluate the fairness of the incentive scheme with respect to the individual utility/costs that arise from cooperation. High degrees of fairness indicate that entities have to exhibit cooperative behavior in order to benefit from the behavior of other entities.

Benchmarks. The original cooperation protocol represents a benchmark for evaluating the efficiency of the incentive scheme. A further benchmark is a system in which a central authority is able to ensure that entities adhere to their commitments. Such benchmark indicates the effectiveness of the incentive scheme since the goal of incentive schemes is to approximate the effectiveness of such central enforcement.

Modelling the entities' behavior. The entities' behavior regarding the cooperation protocol has to be modelled appropriately in order to obtain meaningful simulation results. The conventional approach to such modelling consists of assigning strategies to each entity. For stateless strategies, the model defines the probability of exhibiting specific behavior. For example, according to such an approach, altruistic entities are modelled as entities that adhere to 100% of their commitments. Stateful strategies are based on the trust mechanisms of the distributed reputation system and, thus, consider the past behavior of other entities.

A sounder approach of modelling behavior consists of applying a utility/cost model for each entity. The model can be derived from the analysis of question (Q4). Such modelling is complex since the effects of the incentive schemes have to be modelled, too. For instance, the model has to include the entities' perception of the effectiveness of the incentive scheme. It is needed for the calculation of opportunity costs that arise from misbehavior.

Setting the simulation parameters. The aforementioned behavioral models are parameterizable and, thus, provide for some of the simulation parameters. The second type of simulation parameters may be derived from the incentive scheme if it is configurable. Such configurability is needed for finding entity-dependent tradeoffs between the overhead and benefits of the incentive mechanisms. Finally, the specifics of the cooperation environment constitute the third type of simulation parameters. For example, the number of cooperations that occur during the simulated time is such a parameter.

The simulation parameters do not have to be static throughout the simulation run. The simulation may as well be organized in rounds that are parameterized separately. For example, the simulation methodology of evolutionary game theory [14] may be included if the parameterization of the behavioral models depends on the outcome of previous rounds.

Defining the measurement categories and deriving the key quantities. According to the objectives of the evaluation, the engineer has to measure the total utility u and the total costs c that arise from cooperation. If the objectives of the evaluation include fairness, the individual utility u_i and individual costs c_i have to be measured separately for every entity. Based on these measurement categories, several key quantities have to be derived in order to be able to interpret the simulation results. The ratio $\frac{u}{c}$ indicates the efficiency of the cooperation protocol. It is applicable even if the utility and the costs are measured in different units. The coefficient of correlation between the (u_i, c_i) pairs is called the fairness coefficient. A straightforward means of correlation is the calculation of a regression line between the individual utilities and costs. In case of good linear correlation, the slope of the regression line indicates the magnitude of the incentive effects.

Performing series of measurement. For the evaluation, measurements are required for various parameterizations. In the following, we highlight the series of measurement that make most sense.

- **Configuration of the incentive scheme versus the entities' behavior:** For every type of entity population, it is tested which configuration of the incentive scheme is most appropriate. Such tests may be based on the efficiency of the cooperation protocol or on its

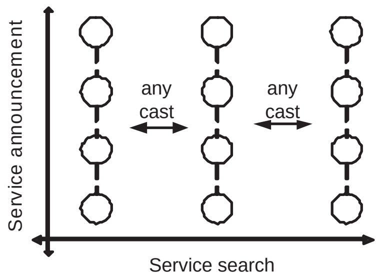

Figure 1: The Lanes service discovery overlay

fairness. For the respective parameterizations, the simulation results should be compared to those of the benchmarks in order to estimate the levels of efficiency and fairness.

- **Configuration of the incentive scheme versus the cooperation environment:** This test varies the parameters of the cooperation environment, e.g., the number of cooperations.

5 A Case Study: Lanes

In this section, the procedure of engineering incentive schemes is exemplified for a service discovery overlay.

5.1 The Lanes Protocol

A service-oriented architecture is a suitable model for application level resource sharing in ad hoc networks. Here, devices can enhance their functionality by using services offered by other devices. For instance, a device might offer the download of certain files, computation or communication capabilities. Because of the lack of infrastructure in ad hoc networks, decentralized service trading becomes necessary.

The *Lanes* [15] approach provides such a cooperation protocol for the announcement and search of services. The Lanes overlay is shown for three lanes in Figure 1. Services are announced along a lane across the proactive vertical overlay links, whereas services are searched via anycast across several lanes.

Analysis. Let us consider the questions introduced in Section 4.1. We answer the questions as follows: **(Q1:** *Entities considered*) Discovery protocol entities [2] and application protocol entities. **(Q2:** *Rational vs. Temper Resistant*) Rational. Since any device may participate in the ad hoc network, we cannot presume tamper resistant hardware. **(Q3:** *Identifiable Entities*) Each entity has a non-alterable identity that is certified together with its public key and is assigned during the bootstrapping of the device. Digital signatures provide a means of reliable authentication. **(Q4/5:** *Valuation*) The utility that arises from cooperation consists of finding services on other devices. The costs ensue from communicating over the network, hoarding service announcements, and processing the cooperation protocol[2]. Due to the limited capabilities of the devices and the frequency of overlay operations, the entities generally prefer efficiency over security. **(Q6:** *Kinds of Cooperation*) The Lanes protocol is composed of five sub-protocols, i.e.,

[2]We refrain from analyzing the entities' valuation of these costs since we will choose a strategic behavioral model for the evaluation.

service announcement, service request, overlay login, intra-lane and inter-lane maintenance. **(Q7:** *Non-Beneficial Steps*) Processing and (unaltered) sending of service announcements, searches and maintenance messages are all non-beneficial under certain circumstances. **(Q8:** *Perceptibility*) Behavior regarding the overlay maintenance can only be perceived by some of the immediate neighbors. In addition, parts of the behavior for the announcement and search of services is not perceptible at all.

Adjustment. Lanes is conceived as an efficient service discovery overlay. Therefore, the complete protocol (including the non-beneficial protocol steps) is necessary for the effectiveness of the service discovery. Due to the imperceptibility of most behavior, several perception mechanisms have to be applied. We introduce redundancy by double linking overlay nodes. Further redundancy is added by promoting the uppermost entity of a lane to the coordinator of maintenance decisions. Furthermore, the sub-protocols that involve several entities are extended with the transmission of receipts. Finally, digital signatures are appended to those messages that are likely to be altered by misbehaving entities.

The adjustment of the protocol cannot cope with the imperceptibility of some protocol steps of the service announcement and search protocol. For instance, a rational entity is still able to return negative results for service searches without actually processing them.

5.2 The Choice of Appropriate Incentives

Definition of Commitments. If participation in the overlay was the only commitment, the design of the incentive scheme would become simple. However, in such a case, the imperceptible parts of the service announcement/search protocols are not executed. Therefore, participation in service announcements and searches should be separately committed to.

In order to log into a lane, a newcomer commits to comply with the remaining sub-protocols. This means that each lane member has to participate in the maintenance protocols. Failure to do so is interpreted as defection from one's own commitment. With respect to configurability, the terms of such commitment may differ. Hence, laptops aim at participating in a secure lane that provides a larger set of services. On the other hand, PDAs choose efficient lanes that are similar to the original protocols. In this regard, truth telling during the login leads to an incentive compatible partitioning into rather homogeneous lanes. Entities that do not adhere to their commitments are treated as outlaws. This means that the other entities opt to exclude them from the overlay. Such step is executed locally within a lane.

Incentives for entering into commitments. For logging in a new entity into the overlay, both the newcomer and the remaining entities have to commit to participate in the maintenance protocol. Therefore, the barter trade pattern provides appropriate incentives for this type of commitment.

The second type of commitment refers to answering a service search of an other entity. For this purpose, we have to choose an incentive pattern that supports asymmetric roles. According to Table 1, the note pattern appears most appropriate. This is because the note pattern provides better enforcement of remuneration than the community pattern. Such enforcement is needed since the costs of processing service searches are considerable. Since we apply the note pattern, the issuer of a service search has to hand over a note to the entity that finds a matching service.

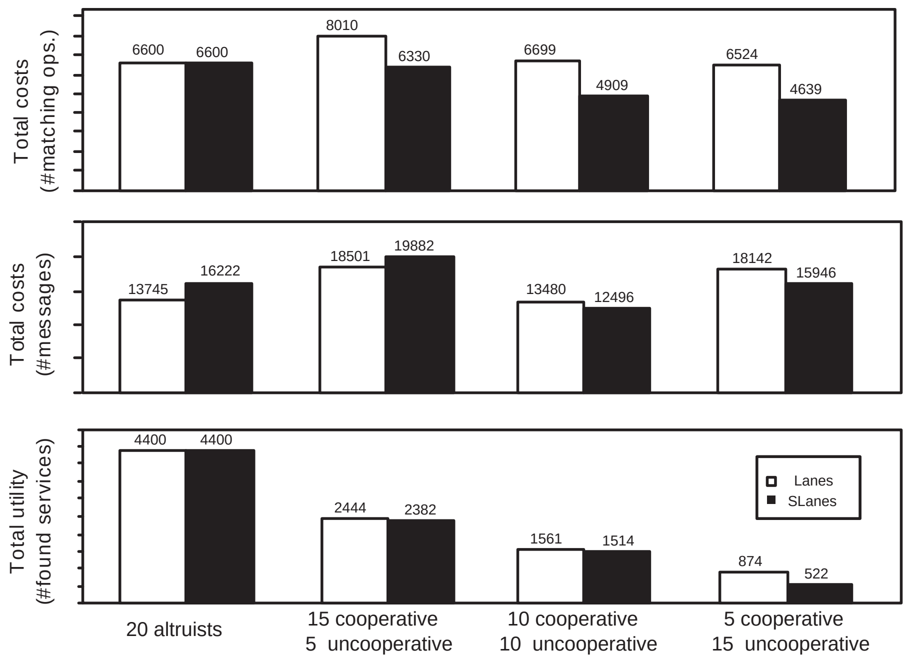

Figure 2: The evaluation of the total utility and the total costs

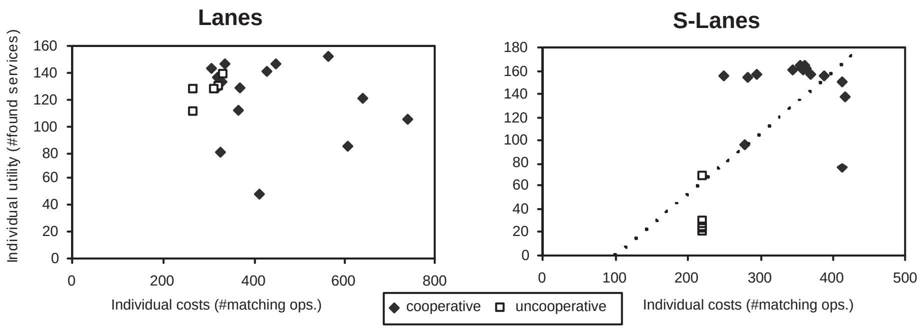

Figure 3: The individual utilities and costs of 15 cooperative and 5 uncooperative entities

Incentives for adhering to commitments. Conventional exchange protocols are too expensive for the operation of the maintenance protocols. However, the sliding window mechanism could be applied in order to limit the number of outstanding notes. This means that every entity is only willing to accept a certain number of notes from other entities. The size of such *note credit* corresponds to the window size of the sliding window mechanism. If the note credit is reached, the bearer of the notes refrains from processing service searches of the notes' issuer.

The lack of exchange protocols demands for the application of a distributed reputation system. We apply the *Buddy System* [16] since it provides for contextualized trust despite of its relatively low overhead.

5.3 Evaluation

In the following discussion of the evaluation, we will refer to the combination of the adjusted Lanes protocol and the incentive scheme as the *S-Lanes* protocol.

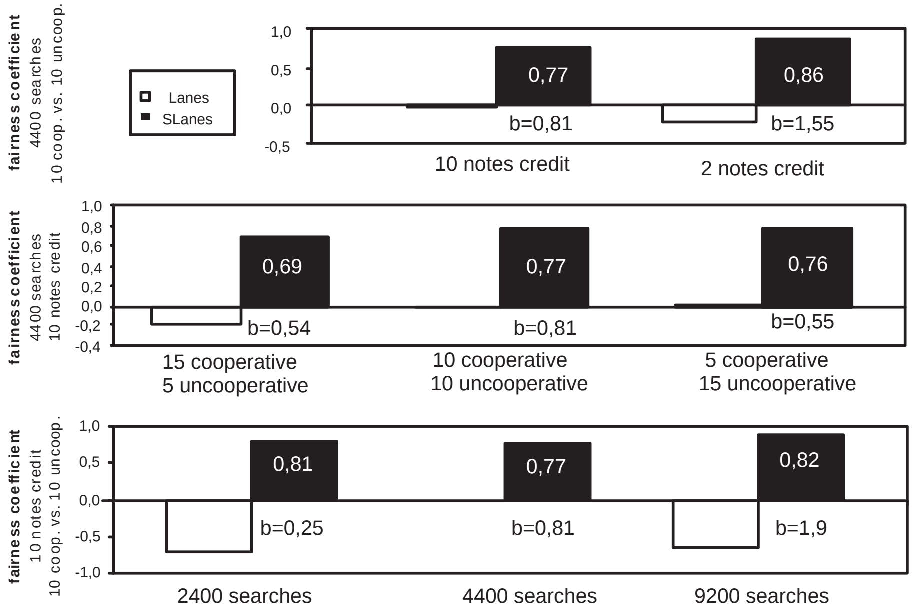

Figure 4: The evaluation of the fairness and incentive effects

The *objective* of the evaluation is to assess the efficiency and fairness of the S-Lanes protocol. For the *behavioral modelling*, we apply the strategic model because of its convenience. More specifically, we distinguish between three types of entities. Altruists always cooperate, whereas uncooperative entities never cooperate. Cooperative entities only cooperate with those entities that have exhibited cooperative behavior towards themselves before. The behavior is further parameterized by the note credit that an entity accords to other entities. Apart from the behavioral modelling, the *simulation parameters* consist of the type of cooperation protocol (Lanes versus S-Lanes) and the cooperation environment. In this regard, the Lanes protocol acts as *benchmark*. The system consists of *20* devices that participate in the service discovery overlay. Each of them provides a unique service and looks for services in an uniformly distributed manner. The cooperation environment is parameterized by the number of service searches that are initiated by each entity. The *measurement categories* consist of the number of found services (utility) and the number of sent messages or matching operations[3] (costs).

The *series of measurement* have been conducted with DIANEmu [17]. The first measurement considers the total utility and the total costs. It is parameterized by the type of participating entities. The number of total searches and note credits is fixed at *4400* and *10* respectively. In Figure 2, the simulation results are illustrated. Interestingly, the total utility and costs of the S-Lanes protocol are considerably lower than those of the Lanes protocol if most entities are uncooperative. This is because the cooperative entities perceive the misbehavior of the uncooperative entities and refrain from cooperating with them any more.

Figure 3 shows the individual utilities and costs[4] of 15 cooperative entities and 5 uncooperative entities. The uncooperative entities are able to profit from cooperation in the Lanes protocol as much as the cooperative entities. However, they do so with a minimum of costs. Therefore, the

[3]A matching operation processes an incoming service search by testing whether it matches one of the hoarded service advertisements.

[4]For clarity reasons, the costs only refer to the number of matching operations.

Lanes protocol is not fair. In contrast, uncooperative entities profit much less from cooperation in the S-Lanes protocol than cooperative entities do. As a result, the respective utility/cost-pairs correlate well with the regression line. The correlation coefficient is *0.69*. The slope b of the regression line is *0.54*. This means that an entity has to perform approximately two matching operations in order to find a service.

We have evaluated the fairness of S-Lanes for different parameterizations. The results are shown in Figure 4. Each line of the figure varies a parameter, i.e., the note credit, the entity types, and the number of total searches. According to the figure, individual utilities and costs do not correlate for the Lanes protocol. For some parameterizations, they are even negatively correlated since the individual utility of uncooperative entities is even higher than the individual utility of cooperative ones. The simulation results clearly demonstrate that the S-Lanes protocol complies with the demand for fairness. According to the results, the more searches are conducted and the tighter the note credit is, the higher the fairness coefficient and the incentive effect is. The slope values higher than *1* indicate superlinear effects of cooperative behavior.

6 Conclusion

In ad hoc networks, devices have to cooperate in order to compensate for the absence of infrastructure. Incentive schemes have been proposed as a means of fostering cooperation among self-interested devices. However, it is difficult to conceive an effective and efficient incentive scheme for a given cooperation protocol. Therefore, in this paper, we have proposed a systematic approach for the engineering of incentive schemes. For this purpose, we suggested a procedure for the design of incentive schemes. It comprises the analysis and adjustment of the cooperation protocol, the choice of appropriate incentives for cooperation, and guidelines for evaluation of the incentive scheme. The procedure has been exemplified for the Lanes protocol. The design and evaluation of an incentive scheme for this cooperation protocol is performed according to the steps of the proposed procedure. The simulation results have shown that the engineered cooperation protocol S-Lanes complies with the demand for fairness.

In the future, we aim at engineering incentive schemes for other service discovery overlays and for application layer cooperation protocols. In addition, we will survey existing exchange protocols and distributed reputation systems with respect to their applicability in ad hoc networks. This will provide more guidance for the choice of appropriate incentives for the adherence to commitments.

Bibliography

[1] Obreiter, P.: A case for evidence-aware distributed reputation systems. In: Second International Conference on Trust Management (iTrust'04), Oxford, UK (2004)

[2] Obreiter, P., König-Ries, B., Klein, M.: Stimulating cooperative behavior of autonomous devices - an analysis of requirements and existing approaches. In: Proceedings of the Second International Workshop on Wireless Information Systems (WIS2003), Angers, France (2003) 71–82

[3] Obreiter, P., Nimis, J.: A taxonomy of incentive patterns - the design space of incentives for cooperation. In: Proceedings of the Second Intl. Workshop on Agents and Peer-to-Peer Computing (AP2PC'03), Springer LNCS 2872, Melbourne, Australia (2003)

[4] Marti, S., Giuli, T.J., Lai, K., Baker, M.: Mitigating routing misbehavior in mobile ad hoc networks. In: Mobile Computing and Networking. (2000) 255–265

[5] Buttyan, L., Hubaux, J.: Stimulating cooperation in self-organizing mobile ad hoc networks. ACM/Kluwer Mobile Networks and Applications (MONET) **8** (2003)

[6] Barreto, D., Liu, Y., Pan, J., Wang, F.: Reputation-based participation enforcement for ad hoc networks. http://www.stanford.edu/~barretod (2002)

[7] Zhong, S., Chen, J., Yang, Y.R.: Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. In: Proceedings of IEEE Infocom'03, San Francisco, CA (2003)

[8] Buchegger, S., Boudec, J.Y.L.: Performance Analysis of the CONFIDANT Protocol: Cooperation Of Nodes — Fairness In Distributed Ad-hoc NeTworks. In: Proc. of IEEE/ACM Workshop on Mobile Ad Hoc Networking and Computing (MobiHOC), Lausanne, Switzerland, IEEE (2002) 226–236

[9] Michiardi, P., Molva, R.: Core: A collaborative reputation mechanism to enforce node cooperation in mobile AD HOC networks. In: Proceedings of the 6th IFIP Communications and Multimedia Security Conference, Portorosz, Slovenia. (2002)

[10] Drexler, K.E., Miller, M.S.: Incentive Engineering for Computational Resource Management. In: The Ecology of Computation. B. Huberman (ed.) (1988)

[11] Liao, R.F., Wouhaybi, R.H., Campbell, A.: Incentive engineering in wireless LAN based access networks. In: Proc. 10th International Conference on Network Protocols (ICNP), Paris, France (2002)

[12] Obreiter, P., Klein, M.: Vertical integration of incentives for cooperation - inter-layer collaboration as a prerequisite for effectively stimulating cooperation in ad hoc networks. In: Proc. of the Second Mediterranean Workshop on Ad-Hoc Networks (MED-HOC NET 2003), Mahdia, Tunisia (2003)

[13] Asokan, N.: Fairness in Electronic Commerce. PhD thesis, University of Waterloo (1998)

[14] Fudenberg, D., Tirole, J.: Game Theory. MIT Press, Cambridge, Massachusetts (1991)

[15] Klein, M., König-Ries, B., Obreiter, P.: Lanes – a lightweight overlay for service discovery in mobile ad hoc network. In: Proc. of the 3rd Workshop on Applications and Services in Wireless Networks (ASWN2003), Berne, Switzerland (2003)

[16] Fähnrich, S., Obreiter, P.: The buddy system - a distributed reputation system based on social structure. Technical Report 2004-1, Universität Karlsruhe, Faculty of Informatics (2004)

[17] Klein, M.: DIANEmu – a java-based generic simulation environment for distributed protocols. Technical Report 2003-7, Universität Karlsruhe, Faculty of Informatics (2003)