

Stimulating Cooperative Behavior of Autonomous Devices

An Analysis of Requirements and Existing Approaches*

Philipp Obreiter, Birgitta König-Ries, and Michael Klein

Institute for Program Structures and Data Organization
Universität Karlsruhe
D-76128 Karlsruhe, Germany
{obreiter,koenig,kleinm}@ipd.uni-karlsruhe.de
<http://www.ipd.uni-karlsruhe.de/DIANE>

Abstract. In the context of mobile and wireless devices, an information system is no longer a centralized component storing all the relevant data nor is it a decentralized component governed by a common authority. Rather, the information spread across huge numbers of autonomous mobile and wireless devices owned by independent organizations and individuals can be regarded as a highly dynamic, virtual information system. For this vision to become reality, the autonomous devices involved need to be motivated to cooperate. This cooperation needs to occur not only on the application layer, but, depending on the network architecture, also on the lower layers from the link layer on upwards. In this paper, we investigate on which protocol layers cooperation is needed and what constitutes uncooperative behavior. We then identify necessary properties of incentive schemes that encourage cooperation and discourage uncooperative behavior. In this context, we examine remuneration types that are a major constituent of incentive schemes. Finally, using the example of ad hoc networks, the most challenging technical basis of a wireless information system, we compare existing incentive schemes to these characteristics.

1 Introduction

Not too long ago, an information system was a centralized data repository. This is no longer true. Today, information is spread across large numbers of autonomous devices, many of them wireless. The aim of an information system is to provide users with access to these distributed resources. To achieve this, cooperation among the devices becomes a necessity. On the application layer, devices need to be willing to share the information they possess. In order to technically accomplish this, device cooperation on lower protocol layers may be necessary, too. Consider, e.g., an infrastructureless ad hoc network. Here, devices must be willing to forward packets on behalf of other devices, if cooperation on the upper layers is to be possible.

* The work done for this paper is partially sponsored by the German Research Community (DFG) in the context of the priority program (SPP) no. 1140. We thank Jens Nimis for his comments on this paper.

Unfortunately, in general, cooperative behavior implicates increased resource consumption and, thus, is not in the interest of the autonomous devices. In case of information sharing, the offerer is confronted with additional disc accesses as well as growing processor load. On lower network layers, increased usage of own bandwidth and energy (e.g., on a node forwarding packages) prevent most of the necessary benevolent cooperation. Again, mobile devices, which are commonly used in wireless environments, are especially concerned about saving their device-inherently scarce resources.

To counterweigh this, external incentives for cooperation are indispensable. These incentive schemes have to be designed in order to discourage uncooperative behavior, while at the same time taking into account the high heterogeneity of the devices and the resulting asymmetry of cooperation and the fact that devices may have valid reasons for a lack of cooperation (e.g., a cell phone might just not be able to offer a certain service), which in this case should not be punished by the incentive scheme.

In this paper, we take a detailed look at where cooperation occurs, what types of uncooperative behavior can be expected (and thus have to be discouraged by appropriate incentive schemes), what should be considered venial noncooperation (Section 2), which requirements need to be met when devising incentive schemes (Section 3), which properties the remunerations of incentive schemes have (Section 4), and finally, in Section 5, how existing schemes compare to these characteristics. Lastly, we conclude the paper in Section 6.

An extended version of this paper is available as technical report [1].

2 Cooperative and Uncooperative Behavior of Autonomous Devices

In this section, we classify cooperation into domains and exemplify misbehavior. Consequently, we propose a taxonomy of uncooperative behavior, which enables a systematic analysis of uncooperative behavior of autonomous devices. It is important to have a detailed understanding of the types of uncooperative behavior in order to conceive an effective incentive scheme.

2.1 Cooperation Domains

As the elementary constituent of cooperation, an entity A acts on behalf of an entity B . In the following, entity A is called *agent entity* and entity B is referred to as *principal entity*. The action is part of the entities' protocol and is beneficial to the principal entity. For example, a network protocol entity, i.e. the agent entity, forwards packets on behalf of its sender, i.e. the principal entity. Therefore, the principal entity remunerates the agent entity and, thus, stimulates the agent entity's action. A remuneration is flexible, if it is assessed situationally, e.g. by taking into account the scarceness of the agent entity's resources.

The action is not necessarily initiated by the principal. For instance, link state packets of the network layer are sent by an agent entity to a principal entity without the principal's explicit request therefor. In the following, cooperative behavior is treated on the elementary principal-agent level. Figure 1 interrelates the proposed terms.




Fig. 1. Elementary cooperation

Inter-device cooperation may be classified into domains with respect to the protocol layer on which it occurs. In the following, we assume a conceptual layering that suffices in order to exemplify misbehavior throughout the cooperation domains. A more comprehensive listing of misbehavior is found in [1].

Link. The link layer provides peer-to-peer links between devices within reach.

If the protocol entity ensures that no packet is waited for, it may turn off its ready-to-receive state, so that upper protocol entities cannot act as agent.

Network. In general, the network layer offers a simple datagram service over multiple links.

Apart from the failure to forward packets, misbehavior varies with the network protocol. In table driven routing protocols, a network protocol entity may fail to send link-state packets. In connection oriented network protocols, an entity may refuse to participate in routing. A network protocol entity might establish dispensable connections or send superfluous packets, thus wasting resources.

Transport. The transport layer provides reliable links between devices.

A transport protocol entity might handle connection management lazily, e.g. by omitting acknowledgements for connection releases and datagram reception.

Discovery. In the discovery layer, application services are advertised by the service provider and requests for application services are processed. Therefore, topologies for routing and matching of advertisements and requests are conceived and maintained. Existing approaches for the discovery layer employ different topologies, e.g., in our research project DIANE, hierarchical rings [2] and multi-layer clusters [3].

A discovery protocol entity may drop uninteresting advertisements or fail to participate in topology maintenance, e.g. shut down without prior sign out. An entity might frequently send advertisements of its application services. In case of memory shortage, it stores none of the other entities' advertisements and, thus, frequently issues requests. If failure to forward advertisements is retaliated, an entity may alter competitors' advertisement, so that they become uncompetitive. In case of flexible remuneration, the malicious entity then boosts its upper layer's remuneration.

Application. Services are provided and consumed on the application layer.

Obviously, selfish behavior consists in not providing application services and, thus, sharing resources. On the other hand, a protocol entity may lavishly consume resource intensive application services instead of providing them itself.

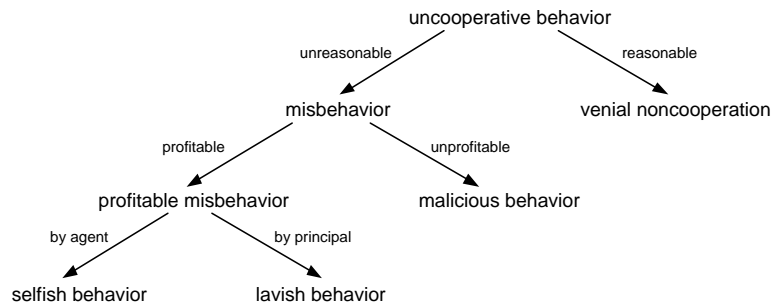


Fig. 2. Taxonomy of uncooperative behavior

2.2 Uncooperative Behavior

A taxonomy of uncooperative behavior is given in Figure 2.

Venial noncooperation is defined as reasonable uncooperative behavior due to resource shortage. Persistent resource shortages arise from the limitations of the device itself, i.e. limitations in computation, memory, bandwidth and energy capacity. Transient resource shortages are due to the device's environment and usage patterns. The device might experience connectivity problems or its resources might be overloaded, e.g. because it is a routing bottleneck.

In wireless information systems, there exist protocol entities that are inherently principal or agent entities. This stems from the asymmetry with regard to the network's topology and the devices' resources and usage patterns. Therefore, asymmetry is closely related to venial noncooperation and it is a problem for incentive schemes that assume that the principal entity has acted as agent entity before.

The incentive scheme should exempt venial noncooperation from punishment. Therefore, it is important to be able to distinguish between venial noncooperation and unreasonable uncooperative behavior, i.e. *misbehavior*.

Malicious behavior, i.e. unprofitable misbehavior, is only exhibited, if it is profitable to a protocol entity of an upper layer. For instance, in a reputation based incentive scheme, posting defamation packets requires resources and, thus, is not profitable to the network protocol entity. However, the application protocol entity profits, if a competing service provider is excluded from the network due to its bad reputation. Generally speaking, protocol entities are subject to a moral hazard as it is known from the agency theory [4].

Profitable uncooperative behavior may be exhibited by the principal or agent entity. A principal entity may *lavishly* consume services, e.g. send superfluous datagrams. An agent entity may *selfishly* fail to commit an action in order to economize its resources. For instance, a network protocol entity may fail to forward packets.

In each cooperation domain, the protocol entity experiences inducements to exhibit uncooperative behavior. However, effective uncooperative behavior requires vertical interaction of protocol entities. For instance, upper layers' protocol entities have to inform the malicious protocol entity about its target. On the other hand, a selfish protocol entity

has to take into account upper layers' cooperation patterns. For example, a link protocol entity is only able to save energy by turning off its ready-to-receive state, if its network protocol entity does not wait for a message.

Remunerating incentive schemes introduce further means of uncooperative behavior. For instance, malicious behavior is especially attractive, if the remuneration of the agent entity is flexible with regard to its resources and to its market position. On the one hand, the principal entity may maliciously prevent its agent entity to act for other principal entities in order to economize the agent's resources and lower its remuneration. On the other hand, an agent entity may raise its remuneration by wiping out its competitors, e.g. by denial of service attacks.

3 Conception and Implementation of Incentive Schemes

The goal of incentive schemes is to encourage cooperation and discourage uncooperative behavior. In order to achieve this, incentive schemes have to meet certain requirements. In this section, we take a look at these requirements.

Effectiveness. An effective incentive scheme restrains uncooperative behavior except for venial noncooperation. The attractiveness of *selfish behavior* is commonly diminished by remunerating the agent entity. The other way round, remuneration of the agent entity keeps the principal entity from *lavish behavior*. Detection and punishment of *malicious behavior* demand for additional mechanisms of the incentive scheme's implementation. Additional mechanisms have to be adopted in order to detect *venial non-cooperation*. In general, asymmetric cooperation patterns are taken into account by applying flexible remuneration.

Incentive schemes are conceived in order to restrain misbehavior. However, they allow for further selfish and malicious behavior with respect to remunerations. The integrity, authentication and non-repudiation properties of a public key infrastructure provide means for robustness against such misbehavior. Tamper resistance might solve parts of the problem, yet its effectiveness is contended [5].

Trust. Depending on the incentive scheme, trust either constitutes an incentive for cooperation or it is a prerequisite for remuneration mechanisms. Trust is subdivided into the principal entity's perspective and the agent entity's perspective. On the one hand, the principal entity has to ensure that the agent entity acts as specified. Furthermore, the agent should not be able to alter or duplicate its remuneration. On the other hand, the agent entity has to make sure that its remuneration is valid. In the following, we distinguish between two extremes of trust, i.e. static and dynamic trust. In practice, the trust mechanisms take advantage of both of them.

Static trust refers to a statement of trust, i.e. a certificate, that remains valid until it is explicitly revoked by its issuer. If trust is transitive, an entity *A* trusts an entity *B*, if there is a certificate chain from *A* to *B*. In [6], the transitivity of trust and the significance of certificate chains is discussed. For example, the manufacturer of devices may issue a certificate of the trustworthiness of the devices' protocol entities. In general, a protocol entity implicitly trusts its device's manufacturer. If, in addition, the manufacturers cross certify, the protocol entities of their devices trust each other. In general,

static trust is implemented within the framework of a public key infrastructure. Apart from authentication, such an infrastructure comes along with mechanisms that enforce integrity and non-repudiation which is of importance for robustness. On the downside, the deployment and operation of such a public key infrastructure place high demands on the devices.

In contrast to certificates, *dynamic trust* arises from prior experiences with an entity and continuously changes according to its behavior. The dynamic trust that an entity *A* has in entity *B* is based on *A*'s own experience with *B* or on other entities' experience with *B*. In order to allow for the latter, *A* needs a way to learn about these experiences. This can be achieved by mechanisms for diffusion of reputation and/or by sniffing, i.e. overhearing messages.

Transaction. An elementary principal-agent cooperation consists of two phases. In the negotiation phase, the participants agree on the agent's action and assess an arbitrary remuneration. In the processing phase, the agent executes its action and is remunerated.

The *negotiation* of the remuneration is a standard approach for overcoming asymmetry in topology, resources and usage patterns. For instance, routing might become more expensive in overloaded parts of the network. The remuneration of an action has to be negotiated with regard to the cost/profit ratio of the principal and agent.

Processing consists of an action and a remuneration. The processing of a transaction is required to be atomic. Therefore, either the agent executes its action and is remunerated, or neither the action nor the remuneration takes place. Unfortunately, atomicity cannot be enforced, since common transaction techniques assume cooperative behavior. If action and remuneration are separable into subactions and subremunerations, the risk of unfairness is considerably lowered by interleaving action and remuneration. In general, the remuneration's granularity is finer and, thus, subremunerations are more feasible. For instance, the principal may turn over half of the remuneration both before and after the agent's action. On the downside, the separation into subactions and subremunerations generally implies additional overhead.

In lower cooperation domains like the network or link layer, actions and remunerations are of lower value. However, the overhead of low value transactions is considerable. For example, the forwarding of a packet on the network layer may implicate further packets, if the forwarding protocol entity is immediately remunerated. Therefore, it seems promising to aggregate actions and remunerations into superactions and superremunerations. As a prerequisite, the principal and actor have to participate in several transactions.

The aggregation of remunerations is rendered more flexible, if flow control mechanisms are applied. For instance, outstanding remunerations may be managed by a sliding window. In this context, flow control constitutes a dynamic trust mechanism. It becomes particularly important, if immediate remuneration is infeasible.

Sniffing. As a generic protocol mechanism, sniffing allows for the collection of dynamic information about the ad hoc network. In certain circumstances, it enables the observation of an entity's behavior. Therefore, in the context of incentive schemes, sniffing is particularly important for dynamic trust and venial noncooperation.

On the link layer, sniffing consists in listening to transmissions that are destined for other devices, whereas, on the network and discovery layer, sniffing entities retain the content of forwarded packets and advertisements/requests respectively. In [7], it is shown that the effectiveness of link layer sniffing cannot be guaranteed.

4 Remuneration Types in Incentive Schemes

Incentive schemes ensure remuneration of the agent entity. In this section, we identify types of remuneration as a major constituent of incentive schemes. The most common remuneration types, i.e. reputation and checks, are discussed with respect to their assumptions and applicability. Lastly, we propose an abstract matching of remuneration types to cooperation domains.

4.1 Remuneration Types

In most incentive schemes, the principal entity remunerates the agent entity. Remuneration assumes a specific form that is called *remuneration type*. It differs among remunerating incentive schemes. For instance, reputation and checks are both remuneration types.

The notion of reputation is directly related to dynamic trust. Reputation subsumes own and other entities' experiences and, thus, constitutes the counterpart of dynamic trust [8]. The dynamic trust that an entity A has in entity B is based on B 's reputation from A 's viewpoint. It is possible to restrain lavishness by reducing the principal entity's reputation. An entity's reputation is only remembered by entities that cooperated before as agent or principal and, in case of sniffing or diffusion, by other entities in the proximity. Therefore, good reputation only pays off in the presence of stable or localized interaction patterns. Otherwise, reputation becomes ineffective and, thus, agent entities are subject to adverse selection [9].

Alternatively, account based electronic payment [10] introduces checks as remuneration type. In such incentive schemes, every entity possesses an account on a virtual bank. The principal entity remunerates the agent entity by issuing a check. Yet, the agent entity has to access the virtual bank, in order to credit its account. Therefore, the accessibility of banks is prerequisite for the application of this remuneration type. In general, the virtual bank is distributed to a set of dedicated devices, i.e. banker nodes. However, an account may be managed by a hardware module on the account holder's device. Such a module is delivered by trusted manufacturers, since it comprises system critical functionality, i.e. the issuing and conversion of checks. Consequently, the virtual bank is distributed among the account holders. In any case, checks require static trust in order to be dependable.

4.2 Account based vs. reputation based incentive schemes

The applicability of the proposed remuneration types is subject to their assumptions. Hence, incentive schemes have to consider the network's peculiarities in order to apply the appropriate remuneration type.




Fig. 3. Matching remuneration types to cooperation domains

Account based incentive schemes facilitate negotiation of the remuneration, since the agent entity is aware of the check's amount. Furthermore, each entity is assigned an account, which allows for the conversion of remunerations to real world money. Yet, account based incentive schemes require static trust and either rely on tamper resistant hardware or on the accessibility of banker nodes. For instance, in ad hoc networks, both assumptions are disputable.

Reputation based incentive schemes do not make such assumptions, since they rely on dynamic trust and, thus, couple remuneration and trust. However, the negotiation phase is omitted, since the principal entity solely assesses the remuneration. Furthermore, the diffusion of reputation is delicate. Although it improves the incentive scheme's effectiveness, it introduces further opportunities for misbehavior. Lastly, reputation based schemes render convertibility impossible and therefore may hinder the deployment of commercial applications.

4.3 Matching Remuneration Types to Cooperation Domains

As pointed out in Section 3 and Section 4.1, two conflicting properties of incentive schemes are discernable. On the one side, the more an incentive scheme relies on dynamic trust the more frequently entities have to cooperate. Otherwise, dynamic trust is not effective. On the other side, the more an incentive scheme depends on static trust the higher its implementation costs, since a cryptographic infrastructure is required.

As for the cooperation domains, their layering implicates two general rules. First, the lower the cooperation domain the more frequently an elementary cooperation takes place. Second, the higher the cooperation domain the more resource intensive the agent's action and, thus, the higher its remuneration is.

On the one hand, there is a direct match between the frequency of elementary cooperation and the locality required by the remuneration type. On the other hand, the actions' counter-value and the implementation costs are interrelated with regard to the commensurability. Consequently, it seems promising to examine the applicability of remuneration types on different cooperation domains with regard to these two dimensions. In this context, matching is abstract, since the specifics of the cooperation domains' protocols are not taken into account. Figure 3 illustrates such matching of remuneration types to cooperation domains.

As a rule of thumb, reputation based incentive schemes should be applied to lower cooperation domains comprising the network layer. Account based incentive schemes are suitable for upper cooperation domains comprising the application layer. This makes sense, since application services are often related to real world prices, as e.g. the printing service to paper and ink prices. Nevertheless, we have to note that the illustrations of Figure 3 are not calibrated. Hence, under certain conditions, reputation is applied on the application layer, whereas account based incentive schemes might be adopted on the network layer.

5 Existing Approaches for Stimulated Cooperation in Ad Hoc Networks

Because of the lack of infrastructure in all of the cooperation domains, ad hoc networks are the most challenging technical basis of a wireless information system. Therefore, in this section, we discuss existing incentive schemes in ad hoc networks with respect to the characteristics of Section 3. In terms of effectiveness, we confine the discussion to the scope of the respective approaches, because an analysis of the approaches' capability of effectively restraining misbehavior requires dedicated studies.

Collective networks. Static trust is a sufficient incentive for cooperation in all ad hoc networks within the boundaries of one organization, e.g. in military, corporate, private and sensor ad hoc networks. Here, inter-device cooperation is inherently motivated. The security requirements varies among these collective networks, which is reflected by the implementation costs of their trust mechanisms. Some collective networks apply further incentive schemes, e.g. for load balancing.

In collective networks, uncooperative behavior is interpreted as venial noncooperation. Furthermore, cooperation between several collective networks is not considered, so that misbehavior is rendered infeasible.

TermiNodes. The TermiNodes project [11, 12] distributes accounts to the respective account holders. The term nuglet as virtual currency is misleading, since there is no notion of token based payment [10]. There is a clear distinction between incentives that restrain selfish and lavish behavior by introducing two charging models, namely packet trade model and packet purse model.

TermiNodes is focussed on incentives for packet forwarding. Therefore, nuglets stimulate cooperation on the network layer. Additionally, the applicability of the incentive scheme to multicast and, thus, cooperation on the transport layer is considered.

Every device possesses a security module that manages its account. The cryptographic infrastructure is deployed and operated within these security modules. Flexible remuneration is achieved by auctions that are held within the security modules. Yet, the conversion of nuglets to real world money is not envisaged.

APE and RPG. In the ad hoc participation enforcement project, two separate approaches have been proposed. The first one is the ad hoc participation economy (APE) [13] that applies dedicated banker nodes in order to manage accounts. Therefore, this approach

renders security modules dispensable, but it relies on the accessibility of banker nodes. Banker nodes facilitate the conversion of digital to real world money. In addition, the negotiation phase allows for flexible remuneration, so that asymmetry is taken into account. APE lays stress on misbehavior of network protocol entities. Yet, checks are not transferable.

The second approach is the reputation participation guarantee (RPG) [14]. It forbids diffusion of reputation. RPG is focussed on the network layer. Selfishness is detected by sending probe packets. However, lavish and malicious behavior is not taken into account.

Sprite. Similarly to APE, Sprite [15] relies on the accessibility of banker nodes that run a credit clearance service (CCS). The amount of remuneration is assessed by such an CCS, so that the negotiation phase becomes dispensable. It is assumed that the devices of the ad hoc network are frequently connected to the internet and, thus, are able to access the CCS. Hence, transferability of checks is not considered.

Sprite introduces a parameterized model in order to prove its effectiveness in restraining selfish behavior on the network layer. In addition, the model is extended to incentives for network connection establishment and transport layer multicast.

Watchdog/Pathrater. In [7], a reputation based incentive scheme is conceived, in order to assert availability of the ad hoc network in the presence of selfish or malicious behavior. Hence, the incentive scheme is applied on the network layer.

Misbehaving protocol entities are excluded from network connections by a watchdog and pathrater run by each device. Reputation is diffused by watchdog synchronization. It is assumed that watchdogs are able to listen promiscuously, i.e. sniffing is assumed. Furthermore, lavish behavior is not restrained, since a selfish protocol entity with bad reputation is still able to act as principal.

CONFIDANT. The watchdog/pathrater approach excludes misbehaving protocol entities from network connections, which is beneficial for selfish protocol entities. Therefore, the CONFIDANT protocol [16] additionally prevents misbehaving protocol entities from acting as principal entity. Therefore, selfish behavior is restrained.

CORE. CORE [17] is reputation based and lays stress on network level selfishness. Defamation is avoided by restricting diffusion to positive local reputations. Nevertheless, unjustified praising is still possible.

Summary. The choice of an appropriate incentive scheme depends on the characteristics and constraints of the respective information system. Therefore, the evaluation of existing approaches in ad hoc networks is summarized in Table 1.

Apparently, the approaches focus on the network cooperation domain. This stems from the relative maturity of network protocols compared to transport, discovery and application protocols in ad hoc networks. However, the discussion of Section 2.1 points out the need for approaches that encompass the discovery and application domain.

Even though every approach is conceived in order to restrain selfishness, existing reputation based schemes do not consider lavish behavior. For approaches that do not

Table 1. Evaluation of existing approaches in ad hoc networks

Approach		Collective Networks	Termi-Nodes	Sprite	APE	RPG	Watchdog/Pathrater	CON-FIDANT	CORE
Scope	Coop. domain	<i>all</i>	N/T	N/T	N	N	N	N	N
	Selfishness	<i>yes</i>	<i>yes</i>	<i>yes</i>	<i>yes</i>	<i>yes</i>	-	<i>yes</i>	<i>yes</i>
	Lavishness	<i>yes</i>	<i>yes</i>	<i>yes</i>	<i>yes</i>	-	-	-	-
	Maliciousness	<i>yes</i>	<i>yes</i>	-	<i>yes</i>	-	<i>yes</i>	<i>yes</i>	-
Remuneration type		-	checks			reputation			
Transferability/Diffusion		-	<i>yes</i>	-	-	-	<i>yes</i>	<i>yes</i>	<i>only positive</i>
Convertibility		-	-	<i>yes</i>	<i>yes</i>	-	-	-	-
Sniffing		-	-	-	-	-	<i>yes</i>	<i>yes</i>	?
Trust		static				dynamic			
Implementation	Tamper resistance	<i>various</i>	security modules	-	-	-	-	-	-
	Cryptographic infrastructure	<i>various</i>	public key	public key	public key	-	-	-	-
Trans-action	Negotiation	-	<i>yes</i>	-	<i>yes</i>	-	-	-	-
	Processing	action	action / remuneration						
Flex. remuneration		-	<i>yes</i>	<i>yes</i>	<i>yes</i>	-	-	-	-

take malicious behavior into account, it is argued that malicious behavior is not profitable and, thus, is not part of the incentive scheme. However, in Section 2, we indicated that a protocol entity's malicious behavior may be profitable for upper layers' protocol entities. Therefore, malicious behavior has to be taken into consideration, in particular for approaches that are focussed on lower cooperation domains.

None of the existing approaches applies mechanisms of the processing phase, as proposed in Section 3, i.e. subdivision, aggregation and flow control of remunerations.

Obviously, the approaches may be classified with regard to their remuneration type. The pros and cons of the respective approaches stem from the characteristics of their remuneration types, as shown in Section 4.

6 Conclusion

Infrastructureless information systems of autonomous devices can only function properly, if the participating devices exhibit cooperative behavior. Therefore, we identified and classified cooperative and uncooperative behavior of autonomous devices on different protocol layers. Furthermore, we discussed key conceptual and implementation issues of incentive schemes. In this context, we identified remuneration types as a major constituent of incentive schemes. Consequently, we reviewed and classified existing approaches for stimulated cooperation in ad hoc networks.

In the future, we intend to conceive an incentive scheme on the discovery and application layer for our research project DIANE [2]. In this context, the design space of incentive schemes and inter-domain cooperation have to be thoroughly examined.

References

1. Obreiter, P., König-Ries, B., Klein, M.: Stimulating cooperative behavior of autonomous devices - an analysis of requirements and existing approaches. Technical Report 2003-1, University of Karlsruhe, Faculty of Informatics (2003)
2. König-Ries, B., Klein, M.: Information services to support e-learning in ad-hoc networks. In: First International Workshop on Wireless Information Systems (WIS2002). (2002) 13–24
3. Klein, M., König-Ries, B.: Multi-layer clusters in ad-hoc networks - an approach to service discovery. In: Web Engineering and Peer-to-Peer Computing. Lecture Notes in Computer Science (LNCS 2376), Springer Verlag (2002) 187–201
4. Bamberg, G., Spremann, K.: Agency Theory, Information, and Incentives. Springer (1989)
5. Anderson, R., Kuhn, M.: Tamper resistance - a cautionary note. In: Proceedings of the Second Usenix Workshop on Electronic Commerce. (1996) 1–11
6. Reiter, M.K., Stubblebine, S.G.: Authentication metric analysis and design. ACM Transactions on Information and System Security **2** (1999) 138–158
7. 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
8. Abdul-Rahman, A., Hailes, S.: Relying on trust to find reliable information. In: Proceedings 1999 International Symposium on Database, Web and Cooperative Systems (DWACOS'99), Baden-Baden, Germany. (1999)
9. Wilson, C.: The nature of equilibrium in markets with adverse selection. Bell Journal of Economics **17** (1979) 108–130
10. Abrazhevich, D.: Classification and characteristics of electronic payment systems. Lecture Notes in Computer Science **2115** (2001) 81–90
11. Buttyan, L., Hubaux, J.: Nuglets: a virtual currency to stimulate cooperation in self-organized ad hoc networks. Technical report, EPFL (2001)
12. Buttyan, L., Hubaux, J.: Stimulating cooperation in self-organizing mobile ad hoc networks. to appear in ACM/Kluwer Mobile Networks and Applications (MONET) **8** (2003)
13. Baker, M., Fratkin, E., Guitierrez, D., Li, T., Liu, Y., Vijayaraghavan, V.: Participation incentives for ad hoc networks. <http://www.stanford.edu/~yl314/adhoc> (2001)
14. Barreto, D., Liu, Y., Pan, J., Wang, F.: Reputation-based participation enforcement for ad hoc networks. <http://www.stanford.edu/~yl314/adhoc> (2002)
15. Zhong, S., Chen, J., Yang, Y.R.: Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks. Technical Report 1235, Department of Computer Science, Yale University (2002) To appear in Proceedings of IEEE Infocom 2003, San Francisco, CA, April 2003.
16. Buchegger, S., Boudec, J.Y.L.: Performance Analysis of the CONFIDANT Protocol: Cooperation Of Nodes — Fairness In Distributed Ad-hoc NeTworks. In: Proceedings of IEEE/ACM Workshop on Mobile Ad Hoc Networking and Computing (MobiHOC), Lausanne, CH, IEEE (2002) 226–236
17. Michiardi, P., Molva, R.: Making greed work in mobile ad hoc networks. Technical report, Institut Eurécom (2002)