Stimulating Cooperation in Mobile Ad Hoc Networks

Johannes Göbel
Department of Informatics
University of Hamburg
22527 Hamburg, Germany
Email: goebel@informatik.uni-hamburg.de
Tel: +49 (0)40 42883 2367 Fax.: +49 (0)40 42883 2311
Anthony E. Krzesinski
Department of Mathematical Sciences
University of Stellenbosch
7600 Stellenbosch, South Africa
Email: aek1@cs.sun.ac.za
Tel.: +27 (0)21 808 4232 Fax: +27 (0)21 808 4416

***Abstract*—Consider a mobile ad hoc network (MANET) where packet transmissions occur between mobile nodes. Such a network requires that the nodes act as relays to form multi-hop routes connecting the origin-destination node pairs that are out of radio transmission range with respect to each other. If the nodes do not to belong the same authority, then the nodes must be given an incentive to spend their resources (battery power and transmission bandwidth) in forwarding packets that originate at other nodes. This can be done by introducing a credit balance for each node, where the nodes use credits to pay for the costs of sending their own traffic, and earn credits by forwarding traffic from other nodes.**

This paper presents several variants of a credit-based incentive scheme (protocol) to promote collaboration in MANETs. The variants address the fairness of the scheme, namely (1) to assist those nodes that are under-provided with credits: such nodes may lack the credits necessary to transmit their data, and (2) to protect those nodes that are over-provided with credits: such nodes may, to their own detriment, provide too large a proportion of their resources to the community.

We first present two basic incentive schemes: the first scheme does not regulate the willingness of the nodes to forward packets on behalf of other nodes; the second scheme contains such a regulatory mechanism. We next focus on protocols whose design lies between the two basic schemes. Both the origin pays and the receiver pays protocols are investigated. Both constant and congestion-related resource prices are investigated.

Initial simulation experiments indicate that both the origin pays and the receiver pays protocol with a credit distribution and discounting (redistribution) mechanism is sufficient to afford a reasonably fair allocation of flows in the network.

I. INTRODUCTION

Consider a mobile ad hoc network (MANET) where packet-based data transmissions occur between mobile nodes. Such a network requires that the nodes act as relays to form multi-hop routes connecting the origin-destination node pairs that are out of radio transmission range with respect to each other. If the nodes do not to belong the same authority, then a protocol must be deployed so that the nodes are given an incentive to spend their resources (battery power and bandwidth) in forwarding packets that originate at other nodes. This can be done by introducing a credit balance for each node, where the nodes use credits to pay for the costs of sending their own traffic, and earn credits by forwarding traffic from other nodes.

Crowcroft *et al.* [5] present such an incentive scheme where each node has a credit balance that determines how much the node can spend on transmission resources in the next time interval. For each node there are two resources: bandwidth and power, each with its own price. The price of each resource increases when the resource is scarce, and decreases when the resource is abundant. When a call arrives, given the current prices, the node selects the least cost route to the destination. The node spends credits at the downstream nodes along the route in order to send its packets to the destination. At the same time the node earns credits when acting as a source, destination or transit node. The interplay between the prices, flow allocations, and credit balances is such that global stability of the system is achieved. Importantly, such a scheme is decentralized: no central controller is needed, and the scheme therefore has favourable scalability properties [7]. An incentive mechanism similar to [5] was previously presented in [1, 3]: the former was evaluated by means of a fluid flow simulation model, whereas the latter were evaluated by a packet-based simulation model; [2, 4, 8] also describe the hardware and software mechanisms required to protect the incentive scheme against theft and forgery.

This paper presents several variants of a credit-based incentive scheme to promote collaboration in MANETs. Each variant is adapted to improve the fairness of the scheme. The adaptations assist those nodes that do not attract sufficient transit traffic due to their disadvantageous locations: such nodes do not benefit from the incentive scheme and may consequently be unable to transmit their data. The adaptations also protect

AEK is supported by grant numbers 2054027 and 2677 from the South African National Research Foundation, Nokia-Siemens Networks and Telkom SA Limited.

those nodes in advantageous locations from providing too large a proportion of their resources to the community.

The remainder of the paper is organized as follows. Section II presents two incentive schemes which in a sense represent two extreme attempts at designing such a protocol. The first scheme does not contain a mechanism to regulate the willingness of the nodes to forward packets on behalf of other nodes. The second scheme contains such a regulatory mechanism. The following sections focus on on protocols for packet admission whose design lies between the two extreme cases. Section V addresses the case where the originating node pays either a constant price or a congestion-related price per hop. Section VI addresses the case where the destination node pays either a constant price or a congestion-related price per hop. The conclusions are presented in Section VII.

II. TWO BASIC INCENTIVE SCHEMES

Let N_i denote the number of packets that originate at node i. Packet losses can occur at any of the nodes along the routes that connect node i to the other nodes in the network. These losses occur because of insufficient buffer space at the nodes. Let S_i denote the number of packets originating at node i that are successfully transferred to their destinations. The ratio

$$P_i = S_i/N_i \tag{1}$$

is the packet transmission success probability (PSP) at node i.

A. *Free-for-all and tit-for-tat*

In the first scheme the network operates on a *free-for-all* basis. No measures are taken to regulate the willingness of the nodes to forward packets on behalf of other nodes. The differences among the PSP's of the nodes are relatively small, although nodes at the edge of the network will have a smaller PSP than the nodes at the centre of the network since nodes at the edge require more hops on average to reach a destination, which increases the likelihood of loss due to encountering a node with insufficient buffer space. However, the *free-for-all* protocol is not fair since the nodes that are likely to act as relays (these nodes are located in the centre of the network) use more of their resources on behalf of the community than the other nodes, yet they receive insufficient compensation in return.

In the second scheme the network operates on a *tit-for-tat* basis. Each node i maintains a counter $C_i(k)$ which records the number of packets that node i has transmitted or received that originated from node k. A packet originating from node k will not be admitted into service on the link (i, j) if

$$C_i(k) + C_j(k) > C_k(i) + C_k(j) + N \tag{2}$$

where N is a small positive constant. Each node thus uses its resources on behalf of the community, but only to the same extent that the node uses the resources of the community. This approach is also questionable. The nodes at the edge of the network may receive a poor PSP: they rarely provide resources to the community and consequently they are not allowed to use the resources of the community.

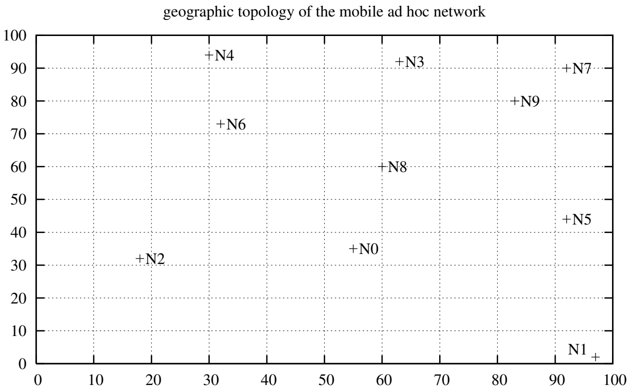

Fig. 1. The 10-node network model.

B. *Fairness*

Jain *et al.* [6] proposed a measurement of fairness of data rate allocation to m competing flows

$$f(x_1, \ldots, x_m) = \frac{\left(\sum_i x_i\right)^2}{m \sum_i x_i^2}$$

where x_i represents the data rate of flow i. The fairness ranges from $1/m$ in the case of the least fairness (all flows except one are zero) to 1 in the case of the most fairness (all flows are of equal size). Based on this approach, we use two measures of fairness among n competing nodes. First, let

$$f_{\mathit{rate}}(P_1, \ldots, P_n) = \frac{\left(\sum_i P_i\right)^2}{n \sum_i P_i^2}$$

represent the fairness in data rate allocation where P_i denotes the PSP of node i. With reference to Eqn. (1), assume that each node i issues approximately the same number of packets N_i. Since the packet size is constant throughout the lifetime of the packet, then P_i will be proportional to the originating flow at node i resulting from all the packets that originate at node i over a given time period.

Next, let

$$f_{\mathit{recip}}(Q_1, \ldots, Q_n) = \frac{\left(\sum_i Q_i\right)^2}{n \sum_i (Q_i)^2}$$

represent the fairness in resource provision where

$$Q_i = \sum_{j \neq i} C_i(j) / \sum_{j \neq i} C_j(i)$$

and $C_i(k)$ was defined in Eqn. (2). The quantity Q_i denotes the *reciprocity* of node i, where $\sum_{j \neq i} C_i(j)$ denotes the the service that node i provides to the community in terms of forwarding and receiving packets, and $\sum_{j \neq i} C_j(i)$ denotes the service that node i receives from the community.

III. SIMULATION EXPERIMENTS

The simulation experiments consider a MANET of 10 nodes distributed on a $100\,m \times 100\,m$ plane as illustrated in Fig. 1 [5]. Each node has a bandwidth of $10\,Mbps$ and a battery of power

TABLE I
RESULTS: *free-for-all* AND *tit-for-tat*

protocol	PSP (%)			avg. relay power (mW)		fairness		throughput
	overall	node 1	node 8	node 1	node 8	rate	reciprocity	packets/s
free-for-all	98.5 ± 0.1	98.9 ± 0.2	98.4 ± 0.4	10.0 ± 0.0	88.8 ± 0.5	1.00 ± 0.00	0.71 ± 0.00	726.5 ± 0.7
tit-for-tat	48.7 ± 0.3	14.2 ± 0.2	78.7 ± 0.8	10.0 ± 0.0	57.8 ± 0.4	0.83 ± 0.00	0.96 ± 0.00	345.3 ± 2.1

$0.5\,W$. The lifetime of the battery is assumed to be infinite. The maximum transmission range is 56 metres. The power required to transmit a unit of flow ($1\,Mbps$) a distance of d metres is given by $kd^2\,W$. A total of $0.5\,W$ is required to transmit $10\,Mbps$ a distance of $56\,m$ therefore $k = 0.05/56^2$. The power required to receive a unit of flow is $0.001\,W$.

During a simulated time of $1,500$ seconds each node attempts 15 data transmissions per second. The individual transmissions take place at the instants of a Poisson process. Each transmission selects a random destination node. The transmission is connected on the route with the lowest (at the instant of the call initiation) costs. The route cost is the sum of the link costs. The link cost is the length of the link. The route is used for the duration of the connection. A data transmission consists of between one and nine packets (uniformly distributed). Each packet is $1,500$ bytes long.

All the packets generated by one transmission appear at the originating node instantaneously. If a packet can be served (see the CSMA/CA description below) then it goes into service, else the packet is queued. If the queue is full the packet is lost. Each node has a separate packet queue for each outgoing link. Each packet queue is of size $25\,Mbit$. Approximately $1,125,000$ packet transmissions were attempted during each simulation experiment. Each experiment was independently replicated 15 times and the $95\,\%$ confidence interval half-widths were computed.

IEEE 802.11 wireless LANs implement CSMA/CA by means of a Request To Send / Clear To Send protocol. The simulator models the RTS/CTS protocol as follows. Consider two nodes i and j. The nodes i and j are said to be connected by a link (i,j) if node i is within radio transmission range of node j. A node cannot transmit and receive simultaneously. A node is idle if it is neither transmitting nor receiving. The link (i,j) is idle if

1) node i is idle, and
2) node j is idle, and
3) no neighbour of node i is receiving, and
4) no neighbour of node j is transmitting.

Note that (3) resolves the “exposed node” problem: a neighbour of node i can transmit when node i transmits, and (4) resolves the “hidden node” problem.

The simulation approximates the RTS/CTS protocol as follows. A packet arriving at node i destined for node j will be transmitted if the link (i,j) is idle, else the packet is queued at the end of the link (i,j) queue; if the queue is full the packet is dropped. When a packet completes service on the link (i,j), the head-of-line packet (if any) enqueued on the link (i,j) goes into service. As each packet completes on the link (i,j), the next packet will be served from the (i,j) queue until a packet completes on another link (m,n) whereupon packets will be served from the (m,n) queue. This mechanism ensures that the last node to successfully transmit a packet will immediately reclaim the transmission medium. There is no backoff delay.

The above CSMA/CA approximation ensures that for the network shown in Fig. 1 at most three packets can be in service simultaneously.

IV. PERFORMANCE: BASIC SCHEMES

Table I presents the efficiency and fairness of the *tit-for-tat* and *free-for-all* protocols.

A. Free-for-all

Some $98\,\%$ of the packets are successfully transmitted. Node 8 consumes nine times as much power as node 1 for relaying the transmissions of other nodes so that node 8 will rapidly drain its battery. The rate fairness f_{rate} is only marginally inferior to perfect fairness since all nodes are treated equally – the small differences in the PSP’s are due to packet losses because of insufficient buffer space. The reciprocity fairness f_{recip} is significantly less than 1.

B. Tit-for-tat

The overall PSP of *tit-for-tat* is much smaller than the PSP of *free-for-all*: only $48\,\%$ of the packets are successfully transmitted. The peripheral node 1 is severely penalized with a PSP of some $14\,\%$. The central node 8 has a PSP of almost $79\,\%$ in return for dedicating most of its resources to other nodes. In comparison to *free-for-all*, the difference between node 1 and node 8 in terms of energy spent for the community is smaller. By design *Tit-for-tat* is fairer than *free-for-all* in terms of resource allocation: the resources acquired from other nodes are approximately equal to the resources spent on behalf of other nodes and so $f_{recip} = 0.96$.

V. PERFORMANCE: ORIGIN PAYS

We now focus on incentive protocols whose design lies between the two extreme cases of *free-for-all* and *tit-for-tat*. We require a compromise in which both the nodes that attract transit traffic (the nodes located at the centre of the network) and the nodes that are unable to attract transit traffic (the nodes located at the edge of the network) will take part in the MANET, even if this has a detrimental effect on the overall PSP. The compromise must ensure that those nodes which spend more resources than the average node on behalf of the community receive preferential treatment (a larger PSP) in return, while those nodes that are not frequently used as relays are still able to conduct some data transmissions. The

TABLE II
RESULTS: ORIGIN PAYS

pricing scheme	PSP (%)			avg. relay power (mW)		fairness		throughput
	overall	node 1	node 8	node 1	node 8	rate	reciprocity	packets/s
constant	97.7 ± 0.1	88.1 ± 0.5	99.7 ± 0.2	10.0 ± 0.0	88.9 ± 0.4	1.00 ± 0.00	0.72 ± 0.00	719.3 ± 0.9
congestion	98.4 ± 0.1	98.4 ± 0.4	98.3 ± 0.5	10.0 ± 0.0	88.2 ± 0.4	1.00 ± 0.00	0.71 ± 0.00	726.2 ± 0.8

features of the protocol are based on the concept of *nuglets* [4], a virtual currency to stimulate nodes to cooperate.

A. *Credit redistribution*

We introduce a credit balance for each node, where the nodes use credits to pay for the costs of sending their own traffic, and earn credits by forwarding traffic from other nodes. In contrast to [4], we require that credits, which correspond to nuglets, be *discounted* (redistributed) continuously. Nodes that possess an amount of credits that exceeds a certain *target balance* will slowly destroy credits, while nodes whose credit balance is less than the target balance will slowly create credits.

Credit redistribution represents a first step towards ensuring that under-provisioned nodes are able to send some traffic, while at the same time providing over-provisioned nodes with a mechanism for using their credits rather than accumulating them. For example, well-provisioned nodes in the *tit-for-tat* protocol can continuously accumulate credits unless the node locations are symmetrically equivalent in their ability to attract transit traffic which will happen for example if the nodes form a ring such that the distance between adjacent nodes is fixed.

Moreover, credit redistribution ensures that credit allocation is robust in that the total amount of credits in the system converges towards the sum of the target balances of all nodes. Nodes which enter or leave the network and which possess significantly more or fewer credits than the target balance, or credits which are lost due to for example a packet carrying credits to be paid to downstream nodes being lost due (thus payment is never credited) have only a temporary impact on the performance of the network protocol.

The introduction of credits (and credit redistribution) to the model raises questions as to how to set the prices for the resources (bandwidth and power) consumed at each node and which node to charge for a data transmission. Either the originating node or the destination node will pay the transit nodes on a multi-hop route to compensate the transit nodes for the use of their resources. In [4] both the origin pays (packet purse model) and the receiver pays (packet trade model) were investigated. In practice, the decision as to who has to pay should depend upon who benefits from a successful packet transmission. The originating node may benefit (consider a node that sends an email), which was the underlying assumption in *tit-for-tat*, or the destination may benefit (consider a node that requests an information service at regular intervals to update a user about stock exchange prices) or both nodes may benefit.

B. *Pricing scheme*

Each node charges a price for transmitting a packet. Prices are determined by one of two simple pricing schemes

1) a constant price of 1 is used for each hop as in the *tit-for-tat* protocol
2) a congestion price

$$p_i(t) = n_i(t)/n_i \tag{3}$$

is charged at node i where $n_i(t)$ is the total buffer space in use at node i at time t and n_i is the total buffer space at node i.

The congestion price represents thecost of admitting a packet into service at a node since the average price charged by a node averaged over many packets reflects the congestion state at the node. The average price at each node will be close to 1 for nodes whose resources are almost fully utilized while nodes whose resources are idle most of the time will charge an average price close to 0. Remember that both the constant pricing and the congestion pricing schemes differ from *tit-for-tat* in that they use credit discounting to redistribute credits from over-provisioned to under-provisioned nodes.

C. *Performance*

In the following experiments each node receives an initial endowment of 2,000 credits which represents the target balance for the purpose of credit redistribution. Credit redistribution is conducted at a rate of 5 %. In each second, the difference between the credit balance and the target balance at each node is reduced by 5 %.

Note that when a packet is sent, the credit balance at the source node is debited immediately whereas the credit balances at the downstream nodes are credited only when the packet reaches them. This reduces the average value of the credit balance at each node. A target balance of 2,000 yields an average balance of 1,600 with an average of 400 credits in transit for each node.

Table II shows that the overall network performance of origin pays is approximately the same as *free-for-all*. The credit-based incentive scheme ensures that a central node like node 8 performs better whereas a node at the edge of the network such as node 1 performs worse. Constant pricing and congestion pricing yield almost the same fairness as *free-for-all* and congestion pricing leads to a significant improvement in the PSP at node 1.

The origin pays protocol with congestion pricing requires knowledge about the local states of the downstream nodes. This knowledge is not available at the originating node at the instant when the packet is transmitted. We assume that

TABLE III
RESULTS: RECEIVER PAYS

pricing	receiver success (%)			avg. relay power (mW)		fairness		throughput
scheme	overall	node 1	node 8	node 1	node 8	rate	reciprocity	packets/s
constant	97.9 ± 0.2	95.5 ± 0.5	98.2 ± 0.4	10.0 ± 0.0	84.0 ± 0.9	1.00 ± 0.00	0.70 ± 0.00	719.2 ± 0.8
congestion	94.8 ± 0.2	70.2 ± 0.6	99.3 ± 0.3	10.0 ± 0.0	77.6 ± 0.9	0.99 ± 0.00	0.67 ± 0.00	697.1 ± 1.1

this knowledge is instantaneously available throughout the network. Our results therefore represent an upper bound on the performance of congestion pricing in practice where the congestion information has to be propagated along the routes, which results in biased pricing information – the prices may be out of date when they are available at the originating node – and additional node resources being used to forward the congestion pricing information.

VI. PERFORMANCE: RECEIVER PAYS

The next set of experiments investigates whether is it possible to obtain similar or better results to origin pays if the destination rather than the origin is charged for the transmission. In contrast to origin pays, receiver pays does not require that additional information be exchanged to estimate the congestion prices at the downstream nodes. Each packet header contains a debit counter (the purse) that may or may not be increased at each node depending upon the local congestion state at that node. When the packet arrives at the destination, the number of credits to be debited at the destination can be obtained from the debit counter.

The target balance in the case of receiver pays is 1,600. In the receiver pays protocol the transit nodes are paid before the destination node is charged so that the receiver pays protocol does not have the "credits in transfer" problem that is a feature of the origin pays protocol. The average credit balance at each node is 2,000, the same as it was in the origin pays protocol.

Assuming that the debit counter can be protected from fraudulent manipulation, the potential drawbacks of the receiver pays protocol include the allocation of relay resources to packets which the destination is not able to pay for. In this case, the destination is forced to pay for the packet and the local credit balance becomes negative. Although the destination has paid to receive this packet, the packet is nonetheless dropped.

Note that in the receiver pays protocol the packet success probability is expressed in terms of the probability of *receiving* a packet whose transmission to a certain node was attempted. In this case, the probability of successfully *transmitting* a packet does not represent a useful measurement since whether or not a packet will successfully complete its journey depends primarily upon the credit endowment of the destination and not upon the credit endowment at the origin. Likewise, the rate fairness measurement refers to the data rate of the received packets.

Table III shows that receiver pays with constant pricing offers an improved PSP at node 1; receiver pays with congestion pricing does not perform well.

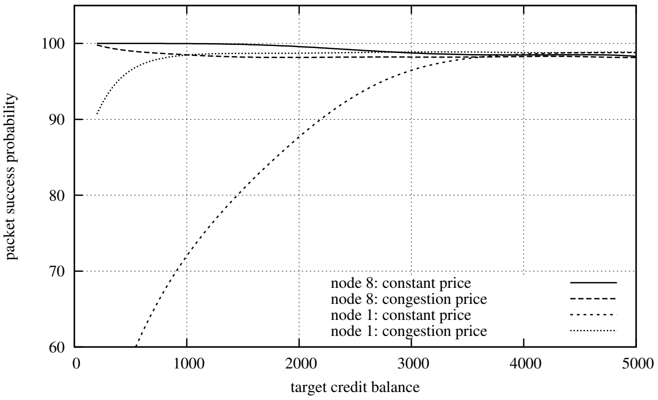

(a) Origin pays

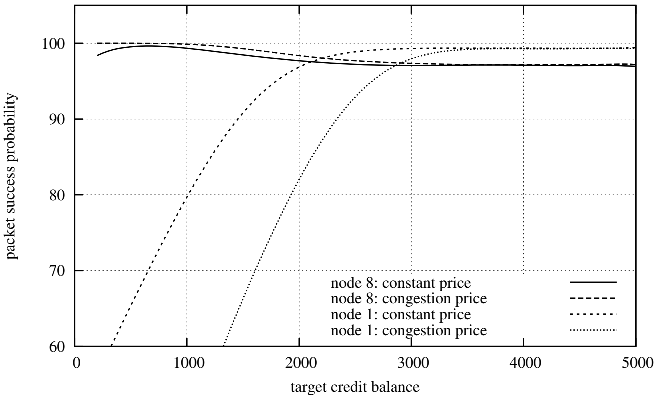

(b) Receiver pays

Fig. 2. The 10-node network model.

One might argue that dropping packets at the destination does not make sense since all the resources needed to transmit the packet have already been allocated. Neither the network (where no resources are freed up) nor the destination (which has to pay for the packet anyway) benefit from dropping packets at this point. Why not let the destination keep (use) the packet? If this were done, then the network behaviour is identical to *free-for-all*, and credits become meaningless since possessing them no longer influences packet transmissions.

Fig. 2 presents the PSP at nodes 1 and 8 as a function of the target credit balance in the case of the origin pays and the receiver pays protocols. The PSP improves as the target credit balance increases. For large values of the target credit balance, the receiver pays protocol affords a better PSP to the isolated node 1 than to the central node 8.

VII. CONCLUSIONS

Initial simulation experiments show that the origin pays protocol with congestion pricing and the receiver pays protocol with constant pricing, both protocols using moderate credit redistribution, afford approximately the same performance when applied to the 10-node network model. The trade-off in favour of nodes at advantageous or disadvantageous locations can be arbitrarily adjusted by setting an adequate rate of credit redistribution. Note that the receiver pays protocol has the important advantage that it requires local information only: the origin pays with congestion pricing protocol requires the additional exchange of data to communicate the congestion prices among the nodes.

Origin pays has the favourable property of never allocating resources to relay packets that the destination cannot afford to pay for. A packet will not be transmitted if the origin does not possess enough credits to cover the cost of the packet transmission.

Origin pays with congestion pricing yields a substantial improvement at the expense of significant overhead required to gather the information not locally available such as the credit balances and the congestion states of the downstream nodes.

REFERENCES

[1] J. Al-Karaki and A.E. Kamal. Stimulating Node Cooperation in Mobile Ad hoc Networks. Wireless Personal Communications, 44(2):219–239, Jan 2008.

[2] S. Buchegger and J.-Y. Le Boudec. Performance analysis of the CONFIDANT protocol. Proceedings 3rd ACM International Symposium on Mobile ad Hoc Networking, Lausanne, Switzerland 226–236, 2002.

[3] L. Buttyán and J.-P. Hubaux. Nuglets: a Virtual Currency to Stimulate Cooperation in Self-Organized Mobile Ad Hoc Networks. Technical Report DSC/2001/001, Department of Communication Systems, Swiss Federal Institute of Technology, Lausanne, Switzerland, 2001.

[4] L. Buttyán and J.-P. Hubaux. Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks. *Mobile Networks and Applications*, 8(5):579–592, 2003.

[5] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring. Modelling Incentives for Collaboration in Mobile Ad Hoc Networks. *Performance Evaluation*, 57(4):427–439, 2004.

[6] R. Jain, W. Hawe, and D. Chiu. A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Computer Systems. Research Report DEC-TR-301, Digital Equipment Corporation, Sept 1984.

[7] J. Göbel, M. Mandjes and A.E. Krzesinski. Incentive-based control of ad hoc networks: a performance study. To appear in Computer Networks, 2009. http://dx.doi.org/10.1016/j.comnet.2009.04.010

[8] S. Zhong, J. Chen, Y. R. Yang. Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. Proceedings IEEE INFOCOM, 1987–1997, 2003.

Constants

$\Delta = 0.01\,s$	the credit redistribution update interval
$\beta = 0.05\,s^{-1}$	the credit balance discount factor
$C_j = 10\,\text{Mbps}$	the bandwidth available at node j
$\Gamma_j = 0.5\,\text{W}$	the power available at node j
$e^{\text{rx}} = 10^{-3}\,\text{W/Mbps}$	the power consumed per unit flow when receiving
$e^{\text{tx}}_{ij} = k\max(10^2, d^2)$	the power consumed per unit flow when transmitting a distance d from node i to node j
$k = 0.05/\delta^2$	the power constant
$\delta = 56\,m$	the connectivity threshold
$\lambda = 10\,s^{-1}$	transmissions per second per node
$U[1,9]$	packets per transmission
$1,500$ bytes	the packet size
25 Mbits	the packet queue per outgoing link

Johannes Göbel received his diploma degree in business information technology from the University of Hamburg, Germany in 2006. He is presently a research staff member at the Department of Computer Science at the University of Hamburg. His PhD currently in progress focuses on the decentralised optimisation of self-organising transport networks.

Anthony Krzesinski is a Professor of Computer Science at the University of Stellenbosch, South Africa. His research interests centre on the performance evaluation of communication networks.