Effect of Non-cooperative Nodes in Mobile DTNs

Ari Keränen, Mikko Pitkänen, Mikko Vuori, Jörg Ott
Aalto University, Department of Communications and Networking (Comnet)

***Abstract*—When applying delay-tolerant networking concepts to communication in mobile ad-hoc networks formed between mobile users, a general assumption is that users are willing to share own resources to support communication between others. However, we cannot assume that all users are altruistic in their behavior; instead, we have to deal with users who only make a limited or no contribution to the mobile community. Nodes not participating in communication only reduce the effective node density, but do not consume resources. Others act as sources and sinks but perform only limited or no forwarding and thus may impact the overall network performance. When considering routing in mobile DTNs, such selfish nodes have to be considered. We introduce two types of selfish nodes and evaluate their impact on message delivery performance for different routing protocols by means of simulation in different synthetic mobility models and with real-world traces. We find that their impact can be surprisingly low in our scenarios, suggesting that DTN communication can be quite robust against selfishness and that controlled non-cooperative behavior may be a suitable way to overcome resource limitations, such as battery depletion.**

I. INTRODUCTION

Mobile ad-hoc networks may offer an interesting alternative for co-located mobile users to exchange information between their mobile devices when there is no wireless communication infrastructure available, conveniently accessible, affordable, or when its capacity is considered insufficient. Especially the introduction of delay-tolerant networking (DTN) concepts [5] to mobile ad-hoc networking has reduced the demand on the density of mobile nodes needed for communication, taking MANETs closer to reality in practical urban indoor and outdoor settings: Even with very modest numbers of mobile nodes recording communication opportunities, numerous quite diverse user traces (e.g. [8], [4], [27]) were generated and used to show that time-space paths can be found for point-to-point communication and information sharing in groups (e.g., [22], [15], [10]).

Irrespective of whether we look at densely or sparsely populated areas for mobile ad-hoc networks, communication between two nodes generally relies on third parties to accept, store, and forward information units. The sparser the networking scenario, the more importance gains the altruistic behavior of the remaining nodes to expend own resources for others. Resources may include storage capacity, CPU processing power, link capacity during communication opportunities (i.e. *contacts*), and energy. Today, with plenty of storage and sufficient CPU power in mobile devices, the constraints essentially reduce to communication and energy capacity—which are interdependent as frequent scanning for contacts as well as transmission and reception, especially at high data rates, consume significant amounts of energy. Experience shows that extensive WLAN usage can easily drain a mobile phone's battery in less than a day.

Because of this, and also because of the latent risk of malware propagating, e.g., via Bluetooth, users often turn off local communication interfaces on their mobile phones. Such users would likely not download and install ad-hoc communication applications. This effectively reduces the node density in a mobile ad-hoc environment. But, given that DTN-based communications does not place strong demands on node density, we can safely ignore such non-participating users from our considerations, similarly users without sufficiently capable devices: they do no good, but no harm either.

The situation is different for users who might leverage the communication capabilities and willingness to cooperate of others, but not contribute own resources. This may happen because of malicious intent, even though it can be argued that the achievable gain and thus the motivation for cheating may be quite limited [7], or the lack of cooperation may be motivated by the mere (temporary) need to save own resources because of low battery.

In this paper, we do not explicitly consider why a particular node does not cooperate as it generally should, are mainly interested in the resulting less cooperative behavior of which we define two types—*non-forwarding* and *partly-forwarding* nodes—and investigate the impact of such nodes' presence on the message delivery performance for different mobility scenarios and different routing protocols.

II. BACKGROUND AND RELATED WORK

DTNs are occasionally connected networks where instant end-to-end paths may not exist and thus messages are not expected to be delivered instantaneously. Some fraction of messages may not be delivered at all. Due to the potentially disconnected nature, a sender may not easily become aware of the delivery success (and delay) or failure.

DTNs may feature dedicated relay or infrastructure nodes, e.g., in vehicular DTNs such as DieselNet [2] or in space communications. Nodes in such DTNs can rely on others to assist in routing decisions and to forward their messages. Alternatively, DTNs may consist of equal nodes that conspire in ad-hoc networks to enable communication. Mobile DTN nodes may follow predictable paths (as satellites or

978-1-4577-0351-5/11/$26.00 ©2011 IEEE

planets do), may be controlled to follow a certain path (e.g., message ferries), or may move in a less predictable pattern.

Depending on the types of nodes and the predictability of their mobility, routing in DTNs may follow a variety of approaches. If time-space paths can be calculated deterministically, single-copy approaches may suffice to reach the necessary reliability [3], [25]. With unpredictable mobility, usually multi-copy approaches [24] are pursued, in which messages are replicated rather than just forwarded. The replication may be unconstrained as in epidemic routing [29] or constrained in the number of copies created [26], the number of hops to take, or to which nodes they are replicated (e.g., based upon encounter history [13] or utility functions [1]). Messages often carry a time-to-live after which they are discarded.

When dealing with mobile DTNs made up from mobile users' devices, the system and its ingredients are not necessarily homogeneous in purpose and capabilities; also, mobility (at the microscopic scale) becomes largely unpredictable. In such a network, cooperation between mobile users becomes crucial for the operation of the entire system, yet we cannot expect everyone to forward messages on behalf of others. As noted above, sending and receiving data consumes node resources, of which especially energy may be the limiting factor, so that cooperation may be perceived expensive and thus unwanted. Also, an increase in the communication load on the entire system may lead to bottlenecks (primarily in communication capacity, possibly also in storage space) so that a node's own messages might propagate less quickly and reliably.

This may motivate users to make their nodes less cooperative, i.e., *selfish*. The result of selfishness is that nodes cannot count on the help of others, as some or all of them try to maximize their personal payoff [28]. Yet, cooperation is required to achieve a certain communication performance to meet the users' requirements; otherwise, nodes could only deliver messages when they meet the destination, like with *Direct Delivery* routing [25].

A lot of research in the area of MANETs has focused on finding ways to ensure that nodes benefit from acting as forwarding nodes. Some of the methods suggested including reputation based systems or using payment methods [6] (see also [7] for a selective summary), but it was also argued that incentive mechanisms may not be necessary [7], albeit proposals for realizing incentives in DTNs exist (e.g., [14]).

Panagakis *et al.* studied the effects of non-cooperative nodes in a DTN [16]. They define two probabilistic types of non-cooperative behaviors: silently dropping a message after reception with a probability P_{drop} and forwarding messages only with a certain probability $P_{forward}$. They investigate the impact of the *cooperation degree* on three different routing protocols, using the random direction mobility model in an open space. Resta and Santi [20] follow up on this with a detailed analytical modeling; they use similar assumptions for cooperation, with diverse simplifications for network load and transmission properties, and provide one additional variant for cooperation in which $P_{forward}$ is dynamically adjusted based on the present network conditions. An analytical model is also developed in [9].

Solis *et al.* investigated strategies for dealing with nodes injecting a disproportionally high share of messages into a mobile DTN, with malicious nodes also "blackholing" (i.e., accepting and then discarding) messages from others [23]. They propose mechanisms for maintaining the performance for well-behaving nodes at an acceptable level, but do not focus on the impact of non-cooperation in forwarding.

We note that nodes do not need to have a malicious intent to become uncooperative: even an originally cooperative node may be forced to cease assisting others when its resources get depleted. It has been shown that DTN routing protocols implicitly or explicitly favor nodes with many contacts for relaying content, putting an undue share of the burden on a few nodes which may lead to quicker resource exhaustion on those unless steps for balancing the load are taken [19].

In this paper, we assume that no incentive mechanisms are in place and investigate the impact of two types of less cooperative operation modes (defined in the next section) on the overall message delivery performance. We make more realistic assumptions on non-cooperation than [16], [20], [9]: We do not expect nodes to silently discard messages (their reception would already consume energy), but rather not accept them in the first place; we let nodes exhibit a coherent behavior across messages over a period of time rather than randomly rejecting one and accepting another with some probability; and we use more realistic message sizes and allow transmissions to fail. Moreover, we use a broader set of mobility models reflecting quite diverse interaction patterns. In contrast to [23], we assume no excessive message generation and do not modify routing protocols, but rather characterize existing ones.

Finally, we investigate both static and dynamic cooperation behavior: to assess the basic impact of cooperation, we statically define which behavior each node exhibits throughout the simulation. In the dynamic model, nodes adapt their behavior according to time and/or energy constraints.

III. (NON-)COOPERATION MODELS

We consider point-to-point communication between nodes where nodes may act as source, sink, or third-party forwarder. Messages are considered to be of no value to forwarders so that, unlike with content distribution, storing and forwarding messages does not yield any personal gain for such intermediate nodes. We define three types of node behavior:

1) Forwarding or *cooperative nodes* are altruistic; they store and forward messages for others without any restriction. We use p_f to refer to the fraction of forwarding nodes and usually plot this value on the X axis in our evaluations.

2) Non-forwarding or *non-cooperative* nodes originate messages as a source and receive messages destined for them as a sink, but do not accept any other messages that would need forwarding. Such nodes are selfish ("free riders") and use others for their own good, but do not contribute to the community, minimizing their own resource consumption. As they do not invest resources into receiving any other messages, they are to some extent honest in that they do not blackhole any message copies. A DTN comprising only non-cooperative nodes would degrade to Direct Delivery (one hop) operation. p_n refers to the fraction of non-forwarding nodes.

3) Partly-forwarding or *partly-cooperative* nodes accept messages from other nodes for forwarding, but only deliver them directly to the destination. The idea behind this behavior is to model nodes that want to make a constructive contribution when they expend own resources for forwarding while not adding to the overall system load and avoiding wasting (own) resources. The assumption is that receiving messages and storing them is acceptable as it is necessary for being partly cooperative, while using contact capacity and battery to send them is costly. A DTN comprising only partly forwarding nodes would obviously deliver messages using at most two hops. p_p refers to the fraction of partial-forwarding nodes.

IV. EVALUATION SETUP

We assess the performance of different cooperation degrees by means of simulation using the ONE simulator [11] that we extend to include non-cooperative and partly-cooperative nodes. We define groups in a way that allows us to vary their composition with respect to the three cooperation models.

A. Scenarios

We use three mobility models: i) As a simple synthetic model, we choose *Random Waypoint (RWP)* with a rectangular simulation area of 4.5×3.4 km and 125 nodes, which move at pedestrian speeds. ii) As a more realistic synthetic model, we choose a map based model *Helsinki City Scenario* with a rectangular simulation area of 4.5×3.4 km and 125 nodes, which move at pedestrian speeds along the streets between popular destinations. iii) We use the *KAIST* trace [21] that records movement of 92 users based on their geographical location every 30 seconds. The geo-locations are imported to the ONE simulator, which interpolates the movement of the nodes between the recorded coordinates.

We choose a simulation time of 12 hours with all traces to stay within a day. The update interval—defining the time step increment for the simulation time—is set to 1.0 s.

B. Routing Protocols

We use three different routing protocols for our simulations: Epidemic routing, Spray-and-Wait, and PRoPHET. The simulations implicitly model also Direct Delivery routing [25], in which nodes wait with forwarding messages until they meet the respective destination. With 0% cooperation, all protocols reduce to direct delivery routing.

Epidemic routing [29] spreads an unlimited number of message copies by having nodes replicate them to all other nodes they connect to. This includes the messages they create and the messages they have received from other nodes. This simple approach floods the network with a given message. While this would ensure that a message reaches the destination if at all possible, the generated load easily leads to bottlenecks in forwarding capacity and buffer space.

Spray-and-Wait routing [26] represents those protocols in which the number of copies created per message is limited. A source sprays this number n of messages to other nodes, directly only or indirectly as in binary Spray-and-Wait and then waits for one of them to meet the destination. We use Spray-and-Wait in binary mode: a node carrying k (we use 10) copies of a message forwards $k/2$ of them to the next nodes it meets until the $k = 1$. Then, a node waits till it meets the destination.

We finally use PRoPHET [13], [12] as a sample protocol for selective message replication. It uses a metric called delivery predictability that is based upon how often two nodes meet each other. The more frequently and the more recently these nodes have met, the better a forwarder one is for messages directed to the other. Messages are replicated based on the predictability, i.e., copies are replicated only to nodes with a better predictability metric.

We keep these routing protocols unmodified, except that non-cooperative nodes do not accept message copies and partly-cooperative nodes forward to destination only.

C. Load and Metrics

In each simulation scenario, a random source node generates a message to a (random) destination node every 25–35 seconds with lifetime of 5 hours. With this frequency, an individual node sends on average a message once per hour in RWP and HCS traces, and every 45 minutes in KAIST trace. The source and the destination nodes are both randomly chosen from the user nodes. The message size is randomly chosen from a uniform distribution of 100 kB–200 kB, modeling typical web-pages contained in a single message, as discussed in earlier work [18]. These parameters model infrequent content transfers on an hourly basis.

We measure the delivery performance of a routing protocol under given cooperation constraints by means of two metrics: the delivery probability of the messages and their delivery latency. The delivery probability is the key performance indicator of the simulations, as we expect the non-forwarding and partly-forwarding nodes to affect this property significantly. The latency is an indicator for the connectivity of the network under a given mobility and cooperation model but also for the network load.

V. SIMULATION RESULTS

We investigate the impact of the cooperation degree for all three mobility scenarios and all three routing protocols. Each plotted metric is an average of 1500 messaging events.

Figure 2 characterizes the used mobility traces in the leftmost column. Random waypoint with a large area (similar size to HCS) and no movement restrictions yields few and often short (mean 58 s) encounters; only some 11% of the possible 7750 unique node pairs meet. Contacts in HCS are longer (mean 84 s) and happen more frequently (three times more often that in RWP), with 22% of the possible unique encounters occurring. KAIST with real world behavior has a large number of both short and long contacts (mean 219 s) and the contacts between pairs of nodes are more repetitive resulting in higher communication capacity than the other models. For KAIST, 58% out of the possible 4186 unique pairs meet. The number of unique node encounters is smaller in KAIST than in HCS due to smaller node population.

A. *Non-Cooperative Nodes*

To calibrate the expectations, we consider the results with all nodes cooperating, i.e., $p_f = 1$ and $p_n = p_p = 0$. This baseline performance is shown in the last column of the graphs in Figure 1 where cooperation percentage is 100. Epidemic offers the best performance at around 90% delivery probability in all three scenarios. Spray-and-Wait exhibit significant differences in performance varying between 50–85% and is generally lower than epidemic routing. PRoPHET shows fairly consistent performance around 80%. Direct Delivery, i.e., no cooperation (0%) only achieves 10–40% delivery. The overall performance confirms that the network is sparse and that the overall network load low, so that the flooding-based protocols do not suffer from traffic overload due to extensive replication.

On the left side of Figure 1, we present how the percentage of non-cooperative nodes affects the overall routing performance of the network. We vary the fraction of non-cooperative nodes $p_n = 0, 0.2, 0.4, 0.6, 0.8, 1.0$ with $p_f = 1 - p_n$.

1) Epidemic: With Epidemic routing, the delivery probability drops as the fraction of non-forwarding nodes in the scenario grows, as expected. We observe, however, that the decrease in delivery rate between $p_f = 1.0$ and $p_f = 0.6$ is minimal and not even very pronounced down to $p_f = 0.4$ (even 0.2 for HCS). Only below this, the rate drops significantly till it reaches the level of direct delivery. The mean latency decreases constantly with added cooperation and is good already when 60% of all nodes are cooperative. These findings hold across all mobility models, with (the least realistic) RWP performing worst; this is expected as in RWP nodes meet less frequently and mostly for short durations so that cooperation of many nodes becomes crucial. In contrast, the constrained movement of restless nodes in HCS yields many and frequent contacts so that the cooperation of individuals is less important. For the KAIST mobility trace, we find slightly lower performance with full cooperation and a slightly more pronounced degradation, probably because the KAIST traces have a 25% smaller node population.

2) Spray-and-Wait: We find that the delivery is only significantly affected when 80% or more nodes are non-cooperative, and then starts dropping significantly. Again, with HCS mobility, the system appears more robust. With the KAIST trace, performance improves marginally with a lower cooperation degree because nodes seem to (accidentally) choose better peers for spraying: message replication is spread further in time, likely helping reaching additional destinations. This is also reflected in an increase in delivery latency. Overall, routing with a limited number of message copies does not seem to be much affected by a small fraction of non-cooperative nodes as long as enough other entities can be found for forwarding. Latency also increases with fewer cooperating nodes, but the effect is not as significant as for the Epidemic case.

3) PRoPHET: The performance results for PRoPHET are qualitatively similar to those for Epidemic routing for both delivery rate and latency. PRoPHET performs with up to 40% non-cooperating nodes roughly as well as with only cooperative ones (for HCS even 20% cooperating nodes seem to suffice) before the delivery rate drops sharply.

Summarizing, we find that all three routing protocols perform well for the given modest traffic load and mobility scenarios even if a significant fraction of nodes stops cooperating (note that the performance figures include messages sent by or destined for non-cooperative nodes). RWP with its short and rare contacts exhibits fast performance degradation as the network is sparser, whereas the HCS map-based mobility model offers more and repetitive contacts and is thus more robust. The KAIST trace has more and longer contacts on average, helping robustness in tolerating non-cooperative nodes; but also many short ones during which not many messages can be exchanged and fewer nodes, limiting robustness.

These results suggest that findings on the performance impact of non-cooperation for a specific mobility model cannot easily be generalized; this also explains why our findings differ notably from [16].

B. *Partly-forwarding Nodes*

Figure 1 (right), shows how partly-cooperative nodes affect the overall routing performance. In contrast to non-forwarding nodes, partly forwarding nodes accept messages, thus grabbing one copy, but forward this message only to the final destination. To evaluate effect of partly-cooperation, we vary the node composition by leaving $p_n = 0$ and varying $p_p = 0, 0.2, 0.4, 0.6, 0.8, 1.0$ while setting $p_f = 1 - p_p$.

1) Epidemic: Replacing non-cooperative with partly-cooperative nodes significantly increases the message relaying capacity of the network. With Epidemic routing, the

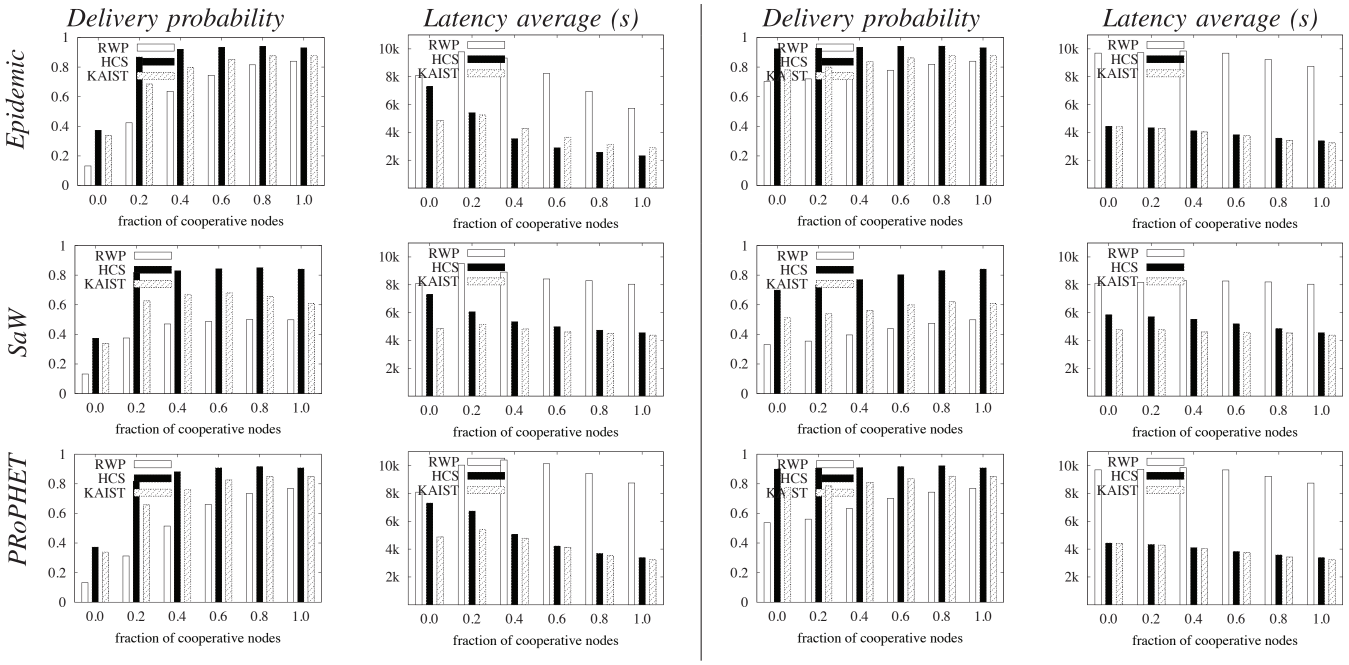

Figure 1. Varying mix of non-cooperative p_n (leftmost columns) and partly-cooperative p_p (rightmost columns) nodes with cooperative p_f nodes.

delivery probability remains high, also with only a small fraction of or no fully cooperative nodes. Most importantly, the sharp performance drop observed at a certain stage for non-cooperative nodes disappears. The resulting two-hop forwarding appears to be largely sufficient to convey the messages from the source to the respective destination. As for non-cooperative nodes, the HCS scenario sees the smallest (almost no) impact, RWP the biggest. The effect of partial cooperation in Epidemic routing is mostly reflected in a modest increase in message delivery delay.

2) Spray-and-Wait: Performance with Spray-and-Wait shows a gradual decline as the number of only partly cooperating nodes increases. When comparing the results to the non-cooperative case, however, we find that a network with partly cooperating nodes only performs better than one with the same fraction of non-cooperative ones if there are no fully cooperative nodes at all. As soon as $p_f > 0$, a network with non-cooperative nodes is at par or marginally better. The phenomenon is caused by partly-cooperative nodes that take messages for forwarding with larger number of copies than one and effectively "trap" them, because they do not spread them further unless they meet the final destination. This effectively reduces the number of copies of a message in the network. The slightly increased delivery latency compared to the non-cooperative case further supports the observation about the harmful effects of partial cooperation on Spray-and-Wait.

3) PRoPHET: With partly cooperative nodes, PRoPHET shows similar trends as Epidemic routing, making efficient use of partial forwarding capacity. While PRoPHET guides replication decisions probabilistically, and is not as aggressive as Epidemic, it does not suffer from problems of fixed numbers of message copies as Spray-and-Wait does.

C. Energy-aware Operation

As noted in the introduction, participation in opportunistic routing consumes energy and may make devices run out of battery easily within a day [19], [17]. Having observed that routing can be efficient even if all nodes are not cooperative, we provide a simple energy-aware model in which nodes periodically alter between cooperative and non-cooperative modes to save energy. This allows nodes to still send/receive messages but minimize their energy consumption.

All nodes choose the same time periods with equal length for both cooperation modes and start with a random fraction of the period as initial offset. This results in a mix of nodes where the cooperation degree is around 50%. For our initial evaluation, we use a simple energy model reflecting energy consumption for scanning, transmitting, and receiving data in a reasonable ratio (the precise values would vary between different devices anyway). The right side of Figure 2 shows the message delivery performance over time (number of messages delivered in each hour). We plot the results for nodes that are fully cooperative until they run out of power (all)—the energy parameters are chosen so that this happens—and for the periodic model (30min, 60min periods). The mean delivery rates are shown in parentheses.

Using mode switching allows prolonging the lifetime of the network until the end of the day (when nodes can be recharged) in all cases. It also results in a better overall

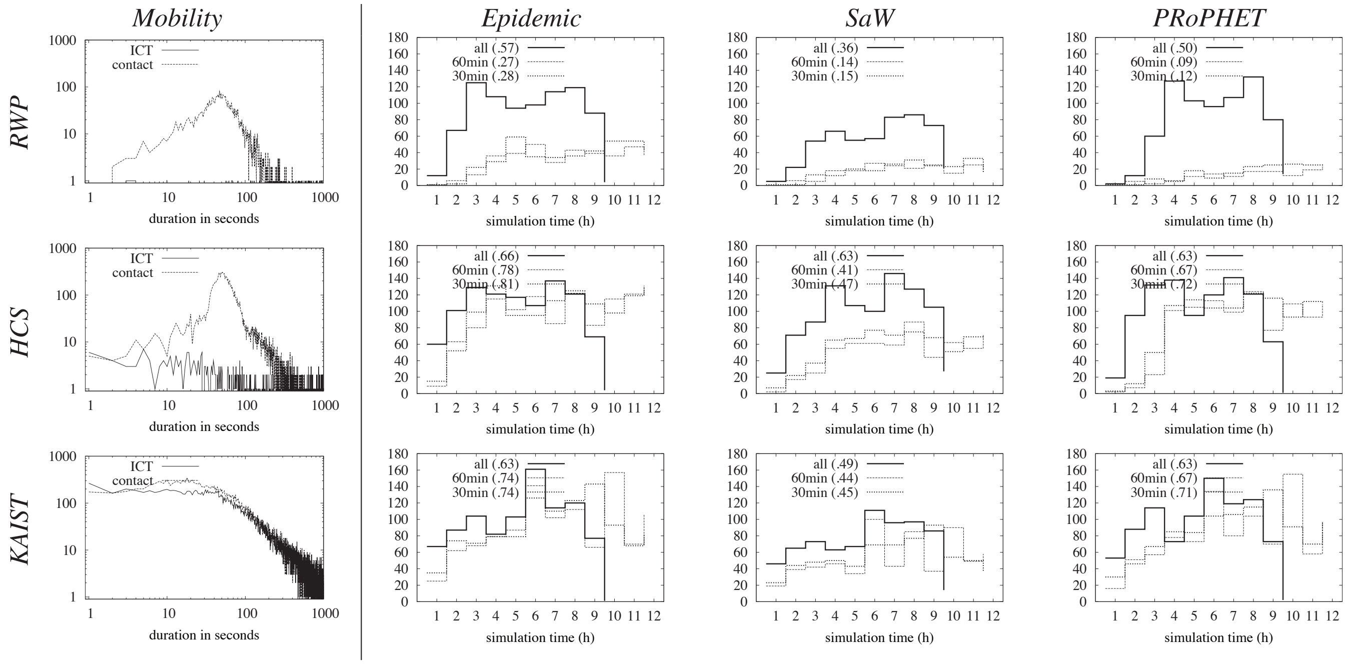

Figure 2. Energy model with Epidemic routing: in basic model node is always on until power suffices (uniform) and altering mode (switch) where node is switched off and on every another epoque. The graphs plot the number of delivered messages per each 3600 second time slot, and delivery probabilities.

message delivery rate for the HCS and KAIST scenarios for both flooding-based protocols, whereas Spray-and-Wait yields (slightly) worse performance. Spray-and-Wait does not perform as well because nodes that accept copies while cooperative may switch to non-cooperative and then trap the copies for the remainder of their non-cooperative time, limiting the number of copies temporarily. Note that the absolute numbers of messages are typically lower when using mode switching because the message delivery is spread out further in time. These results appear to be independent of the switching period (we also simulated two hours).

Overall, a trivial way of energy saving by switching cooperation modes appears feasible in some (more realistic) mobility scenarios to maintain nodes operational and improving (or at least not significantly reducing) the delivery rate—at the cost of longer delivery delays. Further study is required to explore adapting routing protocols to become aware of node cooperation modes induced by energy saving.

VI. CONCLUSION

Our simulation-based investigations into cooperation behavior of mobile DTN nodes provide some further insights into protocol design and the overall feasibility of DTN-based communication between mobile nodes.

The most important finding to note is that—assuming that the message delivery rate is deemed acceptable in the cooperative case—a system can well tolerate quite a large number of free riders, i.e., non-cooperative nodes without too much (additional) harm. For our scenarios, all three routing protocols investigated can easily accept 20–40% (or even 60%) of non-cooperating nodes, even though those still utilize other nodes' resources for their own good. This holds across all three mobility scenarios we investigated, but we find that the most artificial RWP model is most vulnerable to less cooperation. This means that mobile DTNs can be quite robust against misbehaving nodes. A simple initial experiment showed that nodes can make deliberate use of non-cooperation to extend their battery lifetime without causing (much) harm to the overall system performance.

We also find a different impact on the routing protocols: the *relative* performance penalty for Spray-and-Wait, which creates a finite number of copies, is larger than for Epidemic and PRoPHET, which do not limit the number of message copies and may even benefit from fewer fully cooperative nodes as this reduces the load they incur on the network.

While we are aware that these findings only address a small set of scenarios (our current work covers further mobility models and loads), our results hint at further protocol and system design options. Especially, we expect that making simple routing protocols aware of the nodes' cooperation modes will reduce the performance impact observed above. Then, nodes could periodically switch modes (and could safely become non-cooperative when running low on battery) without hurting the system performance at large.

ACKNOWLEDGMENTS

This work was partly funded by the Academy of Finland in the RESMAN project (grant no. 134363) and received funding from the European Community's Seventh Framework Programme under grant agreement no. 258414 (SCAMPI).

REFERENCES

[1] A. Balasubramanian, B. Levine, and A. Venkataramani. DTN Routing as a Resource Allocation Problem. In *Proceedings of ACM SIGCOMM*, August 2007.

[2] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proceedings of IEEE Infocom*, April 2006.

[3] M. Demmer and K. Fall. DTLSR: Delay Tolerant Routing for Developing Regions. In *ACM SIGCOMM Workshop on Networked Systems in Developing Regions (NSDR)*, 2007.

[4] N. Eagle and A. S. Pentland. Reality mining: sensing complex social systems. *Personal Ubiquitous Computing*, 10(4):255–268, 2006.

[5] K. Fall. A Delay-Tolerant Network Architecture for Challenged Internets. In *Proceedings of ACM SIGCOMM*, 2003.

[6] M. Felegyhazi, J. Hubaux, and L. Buttyán. Nash equilibria of packet forwarding strategies in wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 5(5):463–476, 2006.

[7] E. Huang, J. Crowcroft, and I. Wassell. Rethinking Incentives for Mobile Ad Hoc Networks. In *ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, 2004.

[8] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket Switched Networks and Human Mobility in Conference Environments. In *ACM SIGCOMM WDTN Workshop*, 2005.

[9] M. Karaliopoulos. Assessing the Vulnerability of DTN Data Relaying Schemes to Node Selfishness. *IEEE Communication Letters*, 15(12):923–925, 2009.

[10] G. Karlsson, V. Lenders, and M. May. Delay-tolerant Broadcasting. In *ACM SIGCOMM Workshop on Challenged Networks (CHANTS)*, 2006.

[11] A. Keränen, J. Ott, and T. Kärkkäinen. The ONE Simulator for DTN Protocol Evaluation. In *SIMUTools '09: Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, New York, NY, USA, 2009.

[12] A. Lindgren and A. Doria. Probabilistic Routing Protocol for Intermittently Connected Networks. Internet Draft draft-lindgren-dtnrg-prophet-02, Work in Progress, March 2006.

[13] A. Lindgren, A. Doria, and O. Schelen. Probabilistic routing in intermittently connected networks. In *The First International Workshop on Service Assurance with Partial and Intermittent Resources (SAPIR)*, 2004.

[14] R. Lu, X. Lin, H. Zhu, X. Shen, and B. Preiss. Pi: A Practical Incentive Protocol for Delay Tolerant Networks. *IEEE Transactions of Wireless Systems*, 9(4), 2010.

[15] J. Ott and M. Pitkänen. DTN-based Content Storage and Retrieval. In *The First IEEE WoWMoM Workshop on Autonomic and Opportunistic Communications (AOC)*, 2007.

[16] A. Panagakis, A. Vaios, and I. Stavrakakis. On the effects of cooperation in DTNs. In *2nd International Conference on Communication Systems Software and Middleware (COMSWARE)*, pages 1–6, 2007.

[17] A. K. Pietiläinen and C. Diot. Experimenting with Opportunistic Networking. In *Proc. of the ACM MobiArch Workshop*, 2009.

[18] M. Pitkänen, T. Kärkkäinen, and J. Ott. Opportunistic Web Access via WLAN Hotspots. In *Proc. IEEE PerCom*, 2010.

[19] J. Pujol, A. Toledo, and P. Rodriguez. Fair Routing in Delay Tolerant Networks. In *Proceedings of IEEE INFOCOM*, 2009.

[20] G. Resta and P. Santi. The Effects of Node Cooperation Level Routing Performance in Delay Tolerant Networks. In *Proc. IEEE SECON*, 2009.

[21] I. Rhee, M. Shin, S. Hong, K. Lee, S. Kim, and S. Chong. CRAWDAD data set ncsu/mobilitymodels (v. 2009-07-23). Downloaded from http://crawdad.cs.dartmouth.edu/ncsu/mobilitymodels, 2009.

[22] J. Scott, P. Hui, J. Crowcroft, and C. Diot. Haggle: A Networking Architecture Designed Around Mobile Users. In *Proceedings of IFIP WONS*, Les Ménuires, France, 2006.

[23] J. Solis, N. Asokan, K. Kostiainen, P. Ginzboorg, and J. Ott. Controlling Resource Hogs in Mobile Delay-tolerant Networks. *Elsevier Journal of Computer Communications*, 33:2–10, 2010.

[24] T. Spyropoulos, K. Psounis, and C. Raghavendra. Efficient Routing in Intermittently Connected Mobile Networks: The Multiple-copy Case. *ACM/IEEE Transactions on Networking*, 2007.

[25] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Single-copy routing in intermittently connected mobile networks. In *Proc. IEEE SECON*, 2004.

[26] T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and Wait: An Efficient Routing Scheme for Intermittently Connected Mobile Networks. In *ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN)*, 2005.

[27] P.-U. Tournoux, J. Leguay, F. Benbadis, V. Conan, M. D. de Amorim, and J. Whitbeck. The Accordion Phenomenon: Analysis, Characterization, and Impact on DTN Routing. In *Proc. IEEE INFOCOM*, 2009.

[28] A. Urpi, M. Bonuccelli, and S. Giordano. Modeling cooperation in mobile ad hoc networks: a formal description of selfishness. In *Proceedings of WiOpt*, volume 3, 2003.

[29] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-200006, Duke University, April 2000.