

A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks

Markus Jakobsson¹, Jean-Pierre Hubaux², and Levente Buttyán²

¹ RSA Laboratories, Bedford, MA 01730, USA. URL: www.markus-jakobsson.com

² Laboratory for Computer Communications and Applications, Swiss Federal Institute of Technology, Lausanne, EPFL-IC-LCA, CH-1015 Lausanne, Switzerland.
{[Jean-Pierre.Hubaux](mailto:Jean-Pierre.Hubaux@epfl.ch), [Levente.Buttyan](mailto:Levente.Buttyan@epfl.ch)}@epfl.ch

Abstract. We propose a micro-payment scheme for multi-hop cellular networks that encourages collaboration in packet forwarding by letting users benefit from relaying others' packets. At the same time as proposing mechanisms for detecting and rewarding collaboration, we introduce appropriate mechanisms for detecting and punishing various forms of abuse. We show that the resulting scheme – which is exceptionally lightweight – makes collaboration rational and cheating undesirable.

Keywords: audit, collaboration, detection, micro-payment, rational, routing, multi-hop cellular networks

1 Introduction

Multi-hop cellular networks rely on a set of base stations connected to a backbone network, as in conventional cellular networks (such as GSM), and on the mechanisms of ad hoc networks [19], in which packets are relayed hop by hop between peer wireless stations. The expected benefits of such an approach with respect to conventional cellular networks are multifold. First, the energy consumption of the mobile devices can be reduced. Indeed, the energy consumption required for radio transmission grows super-linearly¹ with the distance at which the signal can be received. Therefore, the battery life of wireless devices can be substantially extended if packets are routed in small hops from the originator to the base station. Second, as an immediate positive side-effect of the reduced transmission energy, interference is reduced. Third, if not too remote from each

¹ Depending on the setting, the power decay is a function of the distance, ranging typically from the square to the fifth power [4]. The exact function depends on the extent to which the signal is reflected off of buildings, on the nature of the material to be traversed, on the possible interference from other electromagnetic sources, etc.

The second and third author were supported (in part) by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322 (<http://www.terminodes.org>)

other, mobile devices can communicate independently from the infrastructure. Fourth, the number of fixed antennas can be reduced; and fifth and finally, the coverage of the network can be increased using such an approach. However, while all participating wireless devices stand to benefit from such a scheme, a cheater could benefit even more – by requesting others to forward his packets, but avoiding to transmit others’ packets.

Micro-payments are potential tools for fostering collaboration among selfish (rational) participants, and may be used to encourage collaborative routing of data and voice packets. However, while conceptually well suited to such a task, all proposed micro-payment schemes cannot be applied as such to the problem we consider here. One reason is that it is unlikely for a packet originator to know who – or even *how many* parties – are on the route of the packet. In contrast, traditional micro-payment schemes assume that the payer knows (at the very least) how many payments he is performing at any one time, but typically also whom he is paying – whether by identity, pseudonym or address. If the payer does not know whom he is paying, he could either attach several payment tokens (without any designated payee), or attach one token that can be deposited by several to him unknown parties. Either way there is a potential for abuse, and to avoid this, one needs to generate sufficient audit information to trace users who deposit more tokens than is appropriate. Previously proposed micro-payment schemes do not generate such audit trails.

To make things worse, we must not only consider the possible actions of individual cheaters, but also *collusions* of these. Here, a dishonest set of parties may do anything from routing a packet in a circular manner to claiming rewards for (or collecting payments on behalf of) parties not actually involved in the routing. A possible approach is to assume some degree of tamperproofness (as in [3]). While it can be argued that [3] and other related approaches rely on a similar form of tamperproofness as is successfully provided by GSM SIM cards [20, 15], we mean that the latter provides “portability of identity” rather than security. This is because the SIM cards merely contain identifying information, and not accounting information, and an attacker cannot defraud others (whether other users or the operators) by modifying the functionality of his module. A better comparison in terms of adversarial setting may therefore be that of access to satellite entertainment. There, users may defraud the system (and routinely do) by using rogue modules. While satellite entertainment companies surely would prefer not having to rely on tamperproofness to assure correct behavior, they do not seem to have much choice. In contrast, and as we show, we do not have to rely on any form of tamperproofness to curb cheating – a careful protocol design suffices for the setting we consider.

Finally, we must consider the communication overhead (and the degree of interaction) necessitated by any solution, and make sure that this overhead is acceptable, even for the routing of single packets. This requirement is not normally placed on micro-payment schemes, where the primary constraint is often considered the computational requirement for performing – and receiving – a payment. We place emphasis both on the communication costs and the com-

putation costs, noting that both of them translate into battery consumption. Keeping this low by means of collaborative routing, of course, is the motivating force of this work, and the execution of our protocol must not depart from these goals.

Components and contributions. We avoid the use of all cryptographically heavy-weight operations, and make use of simple symmetric building blocks to achieve our goals. Thus, our contributions are not in the development of new cryptographic techniques, but rather, in addressing an important problem using the simplest possible building blocks. We propose an architecture that is suitable for the model, and put forward four different mechanisms that together constitute our protocol.

The *first* component is a technique for users to determine to whom a packet should be routed. Here, we allow each mobile device to have a preset threshold (potentially depending on its remaining battery life) corresponding to the size of the reward (or payment) they require to transport packets. Likewise, packet originators associate reward levels with packets according to the importance of having them transported. It must not be possible for cheaters to modify these reward levels, of course.

A *second* component is a technique allowing base stations to verify that all packets were accompanied by a valid payment, and drop those that were not. Given that we assume rational (as opposed to malicious) behavior of all participants, this rules out a denial of service attack in which a party causes transport of packets that will later be dropped by the base station due to their invalid payment fields. We argue that this is not a practical limitation, given that an attacker cannot completely drain anybody’s batteries (even if constantly in their presence) since each mobile device has a threshold determining when they will collaborate. Moreover, there are easier ways of mounting denial of service attacks, such as simply jamming the communication channel.

A *third* component is a technique for aggregation of payments. Similar to the recent proposal by Micali and Rivest [16], this works by a probabilistic selection of payment tokens. As in [16], we allow this aggregation to be performed by the mobile devices (payees), for whom storage is a scarce resource. We also consider aggregation of payment information by the base stations as an additional cost-saving measure. To increase the granularity of payments, a user with a “winning ticket” would report the identities of his neighbors (along the packet’s path) when filing a payment claim. Thus, not only the claimant is given a reward for transporting the packet, but his neighbors, too. While this allows for a reduction in storage requirements, its main use is within the fourth component:

The *fourth* component is an auditing process that allows the detection of cheating behavior. This is in the same spirit as the detection of reuse of sequence numbers in [16], but specific to our setting. In particular, our auditing techniques detect and trace dropping of packets, collusions of users filing payment claims, and attacks in which users give priority to the routing of packets carrying winning tickets. Our audit process takes advantage of already collected information from an array of different sources. First, it uses payment claims (winning tickets) from

users. As mentioned above, these contain the identities of claimants’ neighbors (along the packet’s path). Second, it uses packet transmission information from base stations. Third, it makes use of geographical location information collected by base stations – this is information about what users are in what cells at what time, and is already collected for other purposes.

Together, these mechanisms address the problem of how to foster collaboration among rational but selfish nodes² in a multi-hop cellular network. The two main contributions of our paper are the development of a suitable model and architecture; and an audit process suitable for detecting all important attacks without the need for the collection or maintenance of substantial amounts of data.

Outline. We begin by describing our technique at a very high level, and describe related work (section 2). We then turn to detailing our model and goals (section 3). Then, we describe our proposed protocol (section 4) and proposed accounting and auditing techniques (section 5).

2 Overview and Related Work

Multi-hop cellular networks. Although attractive at first sight, multi-hop cellular networks raise a number of problems. For example, in conventional cellular networks, base stations usually are in charge of channel allocation and of the synchronization and power control of mobile devices; to accomplish this task, they take advantage of their direct communication link with each and every mobile device currently visiting their cell. It is quite difficult to extend these operating principles to multi-hop cellular networks. A similar observation can be made in the framework of wireless LANs; for example, in an IEEE 802.11 network, a station can work either in infrastructure mode (namely, with one or several access points), or in ad hoc mode, but not in both.

Over the last years, several researchers have started to bring initial responses to the technical challenges of cellular multi-hop networks. The Soprano project [24] advocates self-organization of the physical, link and network layers. An analysis of the improvement of the throughput is provided in [11], while a routing protocol aiming at providing appropriate QoS is described in [12]. Connectivity of such networks is studied in [5] by means of percolation theory. The use of multi-hop networks is also envisioned in the third generation of cellular networks, where they are called “Opportunity Driven Multiple Access” (ODMA) [7].

Stimulating cooperation in Mobile Ad Hoc Networks. Several researchers have explored the problem of fostering cooperation (especially for packet forwarding) in mobile ad hoc networks. In [14] the authors consider the case in which some malicious nodes agree to forward packets but fail to do so. In order to cope with this problem, they propose two mechanisms: a *watchdog*, in charge of identifying the misbehaving nodes, and a *pathrater*, in charge of defining the best route

² In this paper, “node” and “mobile device” are synonymous.

circumventing these nodes. Unfortunately, this scheme has the drawback that it does not discourage misbehavior.

Another proposal [2] leverages on the reputation of a given user, based on the level of cooperation he has exhibited so far. In this scheme, users can retaliate against a selfish user by denying him service. It is important to note that this proposal is not restricted to packet forwarding but can encompass other mechanisms of the network. A drawback of this type of solution is that a set of colluding cheaters can give each other large quantities of positive feedback, while giving anybody criticizing a member of the collusion negative feedback – both as a deterrent and as a way to reduce the credibility of the feedback the honest user gave.

An even more recent contribution, called Sprite [26], takes a similar approach to what we do in our paper in that it considers an ad hoc network and assumes the presence of a backbone. On the other hand, it does not address the case of multi-hop cellular communications. While the contributions of Sprite are very nice in that they avoid assumptions on tamperproofness while still proving security statements for a stated model, there are potential drawbacks of their solution in terms of its overhead, security, and topology requirements. In particular, their scheme requires a fair amount of computation and storage, making it vulnerable to DoS attacks. Namely, Sprite requires the verification and storage of an RSA signature (or similar) for each packet. In contrast, we use faster verification functions (such as determining the Hamming distance between two strings) and only store verification strings for a fraction of all packets. Moreover, they do not consider attacks involving manipulation of routing tables, while we provide heuristic and statistic techniques to address this problem. Finally, they base their scheme on a reputation mechanism that will only be meaningful in rather dense networks. It is not clear that a typical network exhibits this property.

Among the related work, the already mentioned paper [3] is probably the closest to the present proposal in terms of the problems addressed and the main principles behind the solutions. While the trust model and the protocols of [3] are different from those we propose in that we do not rely on tamperproof hardware, the commonality is that of using micro-payments to foster collaboration in self-organizing networks.

A more general treatment of how to stimulate collaboration can be found in [17]. There, a theoretical framework for the design of algorithmic mechanisms is provided. Although it was developed in a different area, this approach could be applied to our problem, by considering that each node is an agent and that it has to accomplish specific tasks (such as packet forwarding).

Finally, for a general discussion of the security issues of mobile ad hoc networks and for a discussion on how key management can be made independent of any central authority, we refer the reader to [8].

We will now discuss the way we envision the use of payments in a multi-hop cellular network.

Our approach. Instead of using one payment token per payee (as is done in traditional micro-payment schemes [1, 6, 9, 10, 13, 16, 18, 21, 22]), we use one *per*

packet, letting all relaying nodes verify whether this token corresponds to a winning ticket for them. To avoid forged deposits, the packet originator needs a secret key to produce the token (not unlike other payment schemes.) To discourage colluders from collecting payments for each other, we require the intermediary’s secret key (the same as is used to requesting service) to be used to verify whether a ticket wins. Thus, mutually suspicious colluders will not give each other their secret keys, as this allows the others to request service billed to the key owner.

Therefore, we propose a system in which all packet originators attach a payment token to each packet, and all intermediaries on the packet’s path to a base station verify whether this token corresponds to a winning ticket. Winning tickets are reported to nearby base stations at regular intervals. The base stations, therefore, receive both reward claims (which are forwarded to some accounting center), and packets with payment tokens. After verifying the validity of the payment tokens, base stations send the packets (now without their corresponding payment tokens) to their desired destinations, over the backbone network. The base stations also send the payment tokens (or some fraction of these, and potentially in batches) to an accounting center. Packets with invalid tokens are dropped, as the transmission of these cannot be charged to anybody.

However, and again in contrast to previous micro-payment proposals, intermediaries are made to profit not only from *their own* winning tickets, but also from their neighbors’ – we require all reward claims to be accompanied by the identities of the two neighboring parties on the packet’s path. This has three direct benefits: First, the “neighbor reward” encourages the *transmission* of the packet, while the “personal reward” can be seen as a reward for *receiving* the packet – and for reporting this to the clearing house. Second, it increases the number of rewards per deposited ticket, which in turn means that fewer tickets need be deposited. Third, and more importantly, it allows for the compilation of packet forwarding statistics that can be used to detect inconsistent (read: cheating) behavior of intermediaries. By comparing the relative amounts of “neighbor rewards” and “personal rewards” on a per-node basis, the accounting center can detect various forms of abuse. In particular, this analysis will identify parties that routinely drop packets, and parties that refuse to handle packets without winning tickets. It will also detect various forms of collusion. As previously mentioned, we discourage users from performing “collaborative ticket checking” (when one party checks if a ticket is a winning ticket for one of his collaborators) by requiring that they know each other’s keys for this to be possible. In addition, our auditing techniques allow for detection of such behavior, thereby providing two independent layers of protection. While the auditing techniques only detect *repeated* misbehavior (as opposed to the very occasional abuse), this is sufficient, as very few people are likely to alter their devices to make a few cents a month. On the other hand, the more aggressively somebody abuses the system, the faster they will be apprehended, and appropriately punished.

Relation to other payment schemes. Our aggregation principle is based on the idea of probabilistic payments³, suggested by Rivest [21], and also related to work by Wheeler [23]. Therein, each payment can be thought of as a *lottery ticket*. Upon receiving it, the payee can determine whether it is a winning ticket or not, allowing him to erase tickets that were not winning, and request payments from the bank for those that were. While the bank only transfers funds for winning tickets, these are correspondingly larger, thereby causing the appropriate amount of funds to be transferred *on average*. In [21] and most other payment schemes that have been proposed, the bank is transferring funds between payers and payees in a “zero-sum” manner – meaning that for each payee that gets credited, the corresponding charge will be (or have been) levied upon the payer.

If the crediting is probabilistic (as it often is in micro-payments), this may mean that either payers or payees *perceive* a grand total charge (or credit) different than the number of payments suggests – at least for short periods of time. As pointed out by Micali and Rivest [16], this may result in difficulties getting such a scheme adopted by consumers. Accordingly, Micali and Rivest instead shift the temporary fluctuation to the bank, who debit and credit users according to the number of payments made resp. received. In particular, a payer is charged based on the number of payments he performs, while the payee is credited based on whether a payment contains a winning ticket or not. Micali and Rivest let each payment include a serial number that allows the bank to determine (from the winning and therefore deposited tickets) how many payments a given payer has performed. While it is possible for a payer to cheat the bank by performing several payments using the same serial number, this cannot be done *consistently* without the bank detecting it. This is so since the probability of one payment containing a winning ticket does not depend on the probability of another one containing a winning ticket – whether the payments use different serial numbers (as they should) or the same. Thus, while the audit mechanism does not necessarily detect a single instance of abuse, it *will* detect large-scale abuse.

Our proposal is somewhat similar to [16] in that payers (i.e., originators of packets) are charged per packet, and not per winning ticket, while users performing packet forwarding are paid per winning ticket. Therefore, the bank (or accounting center) in our scheme plays the same “averaging role” as the bank in [16]. The mechanisms for detecting abuse in our scheme are statistic, like those in [16]. That is, while it is possible for a payer to cheat the “bank” once, it is not possible in the long run. While there is only one payee per payer in traditional payment schemes, each node on a route in our scheme may win on a ticket associated with one specific packet. The sender pays a cost that – on average – covers the cost of routing, and of other network maintenance. The most straightforward approach of computing this charge is for the bank to compute the average uplink⁴ cost (which depends on the reward level and the average number of hops) and include this per-packet charge in the general charges for transmitting

³ In contrast, Jarecki and Odlyzko [9] perform probabilistic *audits*, while keeping the payments deterministic.

⁴ The uplink is the link from the mobile device to the base station.

packets over the backbone. Such an approach is therefore a generalization of the averaging techniques proposed in [16].

While payees (i.e., nodes) in our scheme are not paid for each packet they handle, they are also not only paid corresponding to the winning tickets they collect: they are also paid each time a neighbor (along the packet’s path) hands in a winning ticket. This provides a step in the direction of crediting payees *per transaction* by increasing the payment granularity, thus requiring fewer tickets to be handed in. It also provides an incentive for users to propagate packets carrying losing tickets (as they may be winning for the neighbor). Most importantly, though, this strategy supplies the back-end with a rich source of data from which it can detect protocol deviations.

As we have seen, payers are charged per packet, and not per winning ticket, while users performing packet forwarding are paid per winning ticket. Such asymmetric payment schemes often allow a coalition of malicious users to make a net profit. As we will see in Section 5, our protocol is immune against this kind of abuse.

3 Model

User model. We assume the existence of three types of participants; *users*, *base stations*, and one or more *accounting center*. In addition, there may be multiple networks, each one of which is considered the *home network* for some users. We distinguish between base stations of the home network, and those of other networks, as will be explained below. We also assume that there is one accounting center per network.

Mobile devices usually have very limited storage and power resources. The base stations and the accounting center, on the other hand, correspond to powerful computers that are connected to each other by means of a high-bitrate backbone network.

Communication model. We assume the use of a network with a multiple-hop uplink, and a one-hop downlink, noting that this choice minimizes the global energy consumption of all mobile devices; we call a network of this kind “asymmetric multi-hop cellular network”. In other words, as a packet travels from its originator to the closest base station, it is transmitted in multiple short hops, since this minimizes transmission costs. Here, the receiving base station may belong to the home network of the user, or (if the user is roaming) to another network, called the *foreign* network. Then, the packet is sent over the backbone from the base station receiving the packet, to the base station closest to the message recipient. (If the packet is multicast, there may be several such base stations, and corresponding receivers.) The closest base station, in turn, transmits it *directly* (i.e., in one hop) to the recipients – this does not require any involvement (or energy consumption) by any of the mobile devices in range. Note that the energy expenditures of the receiver are independent of the distance of the transmission: it is only the sender whose energy consumption depends on

this distance. This model is therefore different from the commonly used *symmetric* communication model in which cell phones and base stations communicate without intermediaries (i.e., where both uplink and downlink are one-hop).

This model is also different from the one usually considered for multi-hop cellular networks. Indeed, in the proposals published so far, both the uplink and the downlink connections are multi-hop. These properties are strongly influenced by the traditional approach of cellular networks (e.g., GSM) and wireless LANs (e.g., IEEE 802.11), in which all links are assumed to be bidirectional. This bidirectionality is considered to be very important, notably for radio resource allocation, power control, and synchronization.

The reason why we depart from this assumption is that a single-hop downlink can be highly beneficial. Indeed, as there is no need to relay downlink signals, the transmission power for the downlink is provided exclusively by the base station, sparing the batteries of the nodes which otherwise would have had to relay the packet. Moreover, this direct channel can be exploited to transmit synchronization signals from the base station to all mobile devices present in the cell. Finally, it makes the allocation of the radio resource on the downlink easier to implement. To the best of our knowledge, asymmetric multi-hop cellular networks have never been proposed in the literature; the study of their feasibility and of their potential merits and shortcomings is well beyond the ambitions of this paper.

Functional model. Users can be categorized as belonging to one or more of the following classes: *originators*; *recipients*; and *intermediaries*. An originator of a packet wishes to have this sent to one or more recipients of his choice. Intermediaries may act as routers, forwarding such packets towards the closest base station. Each such packet then gets transmitted through the backbone network to the base station(s) corresponding to the recipient(s); here they get broadcast by the base station in question and received by the desired recipient. Note again that a packet is only handled by intermediaries on its way *to* a base station, and not *from* a base station on its way to its recipient.

Trust model. Although in reality, very few consumers would attempt to modify the functionality of their devices, it is sufficient that a small fraction would abuse the protocol in order for its commercial usefulness to be endangered. Reflecting this, we make the pessimistic assumption that the devices can be straightforwardly modified by their owners, corresponding to modelling the user as a software module run on a multi-purpose computer, with an appropriate communication module. Users are not trusted to act according to the protocol, but rather, may deviate from this in any arbitrary way. However, it is assumed that the users act *rationally*, i.e., that they only deviate from the protocol when they can benefit from doing so. In particular, users could collude in an arbitrary fashion, and could use a strategy that is a function of data they receive by means of the network. Users trust base stations of their home network not to disclose their secret keys; no such trust has to be placed in base stations outside their home network. All base stations are trusted to correctly transmit packets,

and to forward billing and auditing information to the accounting center of the user's home network, according to the protocol. The accounting center, in turn, is trusted to correctly perform billing and auditing. These are reasonable assumptions for a network that is well guarded against compromise; it is also a reasonable assumption in a network constituting of a small number of principals that audit each other's activities, both by cryptographic/statistical means, and by traditional means.

Goals. The end goal of our protocol is to maximize battery life by minimizing the required transmission signal strength of mobile devices, with the added benefit of increasing the available bandwidth by reducing signal strength. In order to reach this goal, given the selfish nature of users, we propose a set of mechanisms for encouraging collaboration and detecting (and punishing) cheating. In particular, these mechanisms are designed to address several types of abuse, as described hereafter.

Abuse. A naïve solution to the problem may simply provide users with a strategy that maximizes the common good by requiring individual users to collaborate by forwarding other users' packets. However, users – being selfish – may deviate from this proposed protocol. In order to reward altruism, our protocol aims to detect collaboration, allowing this to be rewarded – whether in monetary terms or in terms of improved service levels. Furthermore, our protocol has mechanisms for detection of various forms of cheating. In particular, we prevent or detect the following types of abuse, whether these strategies are used in a “pure-bred” form, or in combination with each other:

- **Selective acceptance.** A cheating strategy in which a user agrees to receive (with the intent to re-transmit) packets with winning tickets, but not packets without winning tickets. (A variation of the attack is when a first user sends a packet to a friend to route, given that the packet is likely to contain a winning ticket for the friend.)
- **Packet dropping.** When a user agrees to receive packets, but does not re-transmit them – whether he claims credit for winning tickets or not.
- **Ticket sniffing.** When a user claims credit for packets he intercepted, but neither agreed to re-transmit nor actually re-transmitted. In a severe version of this attack, colluding users along a fake path submit claims as if they routed the packet.
- **Crediting a friend.** When a user with a winning ticket claims to have received the packet from (or have sent it to) a party different from that which he in actuality did receive it from (resp. sent it to.)
- **Greedy ticket collection.** This is a collection of cheating strategies aimed towards allowing users to claim credits in excess of what the protocol specifies, by collecting and sharing tickets with colluders. Three special cases of this general attack are (1) when one user collects tickets for a friend, knowing that these are likely to be winning tickets for the friend; (2) when sets of users collect and pool tickets, allowing each other to sift through a larger

- pool than they routed; and (3) when a user obtains two or more identities, evaluating tickets with all of these to increase the chances of winning.
- **Tampering with claims.** An attack in which a cheater modifies or drops the reward claim filed by somebody else – when routed via the cheater – with the goal of either increasing his profits or removing harmful auditing information.
 - **Reward level tampering.** An attack in which a packet carries an “exaggerated” reward level promise during some portion of its route, but where the reward level indicator is reduced before it is transmitted to the base station.

Note, however, that a plain refusal to collaborate is not abuse, as long as the refusal is independent of whether a packet carries a winning ticket or not. Users may choose not to route other users’ packets if their resources or policies do not permit them to do so.

Moreover, note that we do not address “circular routing” as a possible attack, given that the rewards will be deterministic given a particular ticket, and therefore, such routing does not behoove an attacker. Neither do we consider the milder form of abuse where a set of users route a message along an unnecessarily long path within a particular neighborhood, in order to allow all of them to (justifiably) claim credit for having handled the packet – this assumption is reasonable if there is enough “real traffic” to route, and the reward structure is set appropriately.

4 Protocol

Setup. As a user u registers to be allowed access to the home network, he is assigned an identity id_u and a symmetric key K_u . This pair is stored by the user and by the user’s home network. As is common, users offer their service provider some form of security, normally implemented by means of a contract or deposit.

Rewards. *Originators* may indicate one of several reward levels; the ultimate (billing) cost for these levels will be specified by his service agreement. The reward level L is an integer within a pre-specified interval $[0 \dots max_L]$. Intermediaries are rewarded accordingly: if transmitting a packet associated with a higher reward level, their expected reward will be greater (with reimbursement levels specified by their service contract). Increasing the reward level allows users with particularly low battery resources to obtain service in a neighborhood populated by other users with low battery resources.

Connectivity graph. We assume that each user u keeps a list⁵ λ_u of triples (u_i, d_i, L_i) , where u_i is the (unique) identity of a neighbor with a path of length

⁵ We do not address how the routing table is built, noting that any standard method, whether proactive or reactive, may be employed. In addition to standard information, the users also exchange information about their reward thresholds L_i .

d_i hops to the closest base station. Furthermore, L_i is user u_i 's corresponding threshold for forwarding packets. (Thus, an entry (u_i, d_i, L_i) in λ_u means that user u_i will forward all packets whose reward level is equal to or greater than L_i , and that the length of the path from u_i to the base station is d_i .) We assume that λ_u is sorted in terms of increasing values of d_i , and that all entries with the same distance d_i sorted in terms of increasing values of L_i .

Packet origination. The *originator* u_o of a packet p selects a reward level $L \in [0 \dots max_L]$, and computes a MAC $\mu = MAC_{K_{u_o}}(p, L)$. He then assembles the tuple (L, p, u_o, μ) and transmits this according to the transmission protocol below.

Packet transmission. Let u be a user (whether originator or intermediary) who wishes to transmit a packet associated with a tuple $P = (L, p, u_o, \mu)$. In order to transmit P , user u performs the following protocol:

1. If the base station can be reached in a single hop, then u is allowed to send the packet directly to it; otherwise he goes to step 2.
2. u selects the first (hitherto unselected) entry (u_i, d_i, L_i) from λ_u for which $L_i \leq L$.
3. u sends a *forward request* to u_i . This contains the reward level L and possibly further information about the packet p , such as its size⁶.
4. u waits for an acknowledgement from u_i for some pre-set time period δ . If u receives the acknowledgement, then he sends P to u_i . Otherwise, if no acknowledgement arrives, he increases i by one. If $i > |\lambda_u|$ then he drops the packet; otherwise, he goes to step 2.
5. If u is not the originator of the packet, he performs the reward recording protocol below.

Packet acceptance. Let u' be a user receiving a forward request from u with reward level L . If L is less than his threshold, then he does not accept the request; otherwise, he accepts it by sending an acknowledgement to u and awaits the transmission of the packet.

Network processing. When a packet $P = (L, p, u_o, \mu)$ is received by a base station in the originator's home network, the base station looks up the secret key K_{u_o} of the originator u_o , and verifies that $\mu = MAC_{K_{u_o}}(p, L)$, dropping the packet if this does not hold.

If the packet is received by a base station that belongs to a foreign network, this base station cannot perform the verification (as it does not have access to the originator's secret key), and so, forwards the packet P to a register in the originator's home network. This register, then, looks up the originator's secret key, performs the verification, and drops the packet if the verification fails⁷.

⁶ Most protocols support several packet sizes.

⁷ Similarly to the technique adopted in most 2G and 3G cellular networks, the detour of each and every packet via the home network can be avoided, by letting the foreign network perform the described verifications; this can be done without revealing the secret key to the foreign network.

If the verification of the MAC succeeds, the base station (resp. home network register) transmits the packet portion p to the base station associated⁸ with the desired recipient (as indicated in p). The base station associated with the desired recipient broadcasts p to the latter.

The first base station records a fraction ϵ_μ of all triples (μ, L, u) , where u is the identity of the user it received the packet from. It also keeps a count cnt_{u_o} of the number of packets it transmits for u_o . Periodically, base stations send such recorded auditing information to an accounting center, along with geographical information consisting of statistics of what users were in what cell at what time (not all such information is sent, but some portion.)

Reward recording. After user u has forwarded a tuple $P = (L, p, u_o, \mu)$, he verifies whether $f(\mu, K_u) = 1$ for some function f (the choice of which is discussed below). If this relationship holds, it means that the considered ticket is winning; he then records (u_1, u_2, μ, L) , where u_1 is the identity of the user he received the associated packet from, and u_2 is the identity of the user (or base station) he forwarded it to. We let M denote the list of recorded reward triples.

Reward claim. If a user u is adjacent to a base station (i.e., the distance to the base station is 1), then he transmits a claim (u, M, m) to the base station, where $m = MAC_{K_u}(hash(M))$. Thus, the reward claim M is authenticated using the same key K_u as the user employs when originating a packet, or verifying whether a packet contains a winning ticket.

Similarly, if user u originates a packet P or is running out of storage space for claims, then he transmits the claim to the closest base station by means of the packet origination protocol, and using the base station as the packet recipient. The portion M may be encrypted using a stream cipher and using a secret key shared by user u and either the base stations or the accounting center – in the latter case, the MAC m would be computed on the ciphertext of M .

When a base station receives a claim, it verifies the correctness of the MAC m with respect to the user u and the claim M (or the ciphertext, as explained above). If this is not correct, then he ignores the claim; otherwise, he records the claim and computes an acknowledgement ack to it as $ack = MAC_{K_u}(m)$, where K_u is the key he shares with the user (claimant) u . ack is transmitted to u , who upon receipt verifies the acknowledgement and erases M if correct. Within a time Δ , each base station forwards all recorded claims to an accounting center, and then erases the list.

⁸ Standard techniques can be used to determine in what cells packet recipients are located. In particular, one may require users to announce their location to base stations at regular intervals, or to announce changes of location – inferred by these by the changing identity of the closest base station. While this “announcement” is currently performed by direct communication from mobile device to base station, our multi-hop technique can obviously be used instead. Users may piggyback reward claims with such location announcements.

Ticket evaluations. As mentioned above, all tickets μ are evaluated with respect to the secret⁹ key K_u of the user u in question, and with respect to some public function f that results in a uniform distribution of winning tickets. One can choose f as a one-way function, such as a hash function, and let a winning ticket be one that hashes to a value with a certain pattern (e.g., any string that starts with ten zeroes.) However, since the evaluation of f has to be performed once for each packet the user u handles (except for those packets originating with u , of course), it is important that f is lightweight, and preferably more light-weight than hash functions are.

A promising possibility is to let $f(\mu, K_u) = 1$ iff the Hamming distance between μ and K_u is less than or equal to some threshold h . Thus, assuming that $|\mu| = |K_u|$, and given a particular reporting threshold h , the probability of μ being a winning ticket is

$$\frac{1}{2^\ell} \sum_{i=0}^h \binom{\ell}{i}$$

where $\ell = |\mu| = |K_u|$. Note that it is possible to assign different rewards to different Hamming weights in the range, making it possible for a user to keep only the “highest rewards” in case he runs out of memory and needs to purge some portion of the rewards. However, for simplicity, we assume that all reward claims have the same value.

However, we note that if f is not a one-way function (as in the case above) then it may be possible for an attacker to derive the user’s secret key K_u by observing what tickets are filed. Therefore, if such a function is used, it is important that all claims are encrypted during transmission, in which case only the *number* of claims (as opposed to the form of these) would be revealed to an attacker.

We note also that f must be chosen in a way that the distribution of winning tickets is uniform.

On the probability of winning. The efficiency of our protocol relies on the probability of a ticket to win to be small enough for the claim process not to dominate the protocol, whether in terms of storage or communication. At the same time, we need the probability to be large enough that the reimbursement process relies on a large number of claims, which in turn makes auditing possible by providing a sufficiently large data set. Therefore, one needs to carefully balance these problems against each other when selecting the appropriate reward function. Rather than a security issue, this corresponds to a risk management issue and a usability issue.

⁹ It is important that all of (or close to all of) K_u is needed to evaluate f successfully – or users would be able to verify reward claims on behalf of each other, without having to trust each other with their secret keys.

5 Accounting and Auditing

The accounting center receives both user claims and partial transmission transcripts – both forwarded by base stations. These are processed as follows:

Accounting. The accounting center periodically verifies all received user claims with respect to all recorded reward tuples it has received from base stations. All originators whose identity, u_o , has been recorded by a base station are charged a usage fee according to their service contract. Moreover, the accounting center credits all parties¹⁰ whose identity figures (whether as a claimant or neighbor thereof) in an accepted reward claim. It is a policy issue how to set the rewards for neighbors to claimants, i.e., whether to let these depend on the reward level of the packet as well, and how large a neighbor reward would be in comparison to a claimant reward.

Here, a reward claim is said to be *accepted* if it is *correct* (i.e., if $f(\mu, K_u) = 1$) and a base station has reported the packet associated to the ticket μ as having been transmitted. Note that the accounting center may credit claimants and neighbors thereof according to any policy, and, in particular, the amounts may differ between claimants and neighbors; we do not dwell on these intricacies in this extended abstract.

Simplified auditing. Assume for a moment that the probability of a ticket to win is 1, and that all of these claims get reported by the users and passed on to the auditing center. Assume further that all MAC headers are stored by the base stations, and forwarded to the auditing center. We can now see that the auditing center will know the origination point of each packet (from the identity and MAC of the packet), and the identity of the base station receiving it. It will also know the identity of the user transmitting it to the base station (since this is recorded by the latter). From the claim of this user, it will know the identity of the user one step earlier in the forwarding chain, and so on. This will take us all the way back to the identity of the user who received the packet from the originator, who in turn will report whom he received it from (i.e., the originator.) If any user other than those already accounted for in the above claims a reward, this will be identified as an attempt to cheating. The auditing process for the probabilistic setting is analogous to the analysis of the simplified setting in that it approximates the latter by means of statistical methods.

Auditing. In the following, we will assume that $\epsilon_\mu = 1$, i.e., each base station stores the MAC header of each packet. The more general case in which different base stations store different fractions can be dealt with similarly: instead of merely counting occurrences, one would then test various hypotheses using standard statistical methods. It is worth noting that the probability of a ticket being a winning ticket is a function of three quantities: the message; the secret

¹⁰ There are two exceptions to this rule: Neither packet originators nor base stations obtain rewards.

key of the originator; and the secret key of the intermediary (i.e., the party verifying whether the ticket is winning.) Since the secret keys of users are selected uniformly at random, the distribution of winning tickets is uniformly distributed over all messages.

Common for many of the detection mechanisms is the observation that since the probability for a ticket to win is independent of the identity of the user, each user should figure as the claimant with approximately the same¹¹ frequency as he figures as either the *sending neighbor* or *receiving neighbor* of a claimant. While one cannot simply compare the number of occurrences of these events, one can check the hypothesis that they are all generated from a source with the same event probability. As will become evident, many of the attacks we consider leave very similar-looking evidence, which may make it difficult to establish with certainty what the attack was. However, one can easily establish the presence of one of these attacks using standard statistical methods, and given sufficient material.

- **Selective acceptance.** Selective acceptance is epitomized by a user figuring as a *claimant* with a significantly higher frequency than as a *sending neighbor*.
- **Packet dropping.** A user is suspected of packet dropping if he has a higher *claimant* frequency than *sending neighbor* frequency for packets that were not reported as received by any base station.
- **Ticket sniffing.** A user is suspected of ticket sniffing if he has a higher claimant frequency than *sending neighbor* or *receiving neighbor* frequency, and there are incidents when both he and a neighbor files a claim for one and the same ticket, but do not list each other as the corresponding neighbors. If an entire fake path of reward claims has been created, the auditing center can distinguish between this and a real path (with some probability) given that the receiving base station will record the identity of the user from whom the packet was received.
- **Crediting a friend.** An indication of this attack is that the *receiving neighbor* of a given claim was reported by the base station to have been located in a distant cell¹² at the time the packet was received by the base station. Another indication is if a first user reports a second party to be the receiving neighbor, while another (also claiming a reward for the same packet) claims to have received the packet from the first party. While it is difficult to determine from one occurrence whether the first or the third party filed an incorrect claim, repeated occurrences will allow this to be established.

¹¹ We note that this is true independently of what the “collaboration thresholds” of the different parties on the route are. This is so since we consider the frequencies along a path of senders where all have agreed to collaborate – their thresholds are therefore irrelevant!

¹² All cellular devices report to the closest base station when they move from one cell to another. Similarly, when a device is turned on, it reports to the closest base station. If a device is moved while turned off, we consider it to still remain in the cell where it last was heard from.

- **Greedy ticket collection.** This has the same symptoms as the above mentioned attack. In addition, transmission paths – counted in number of claims per packet – that are longer than usual (for the given cell) are indicative of this attack. Similarly, abnormally high packet transmission rates per time unit by some user indicates that greedy ticket collection has taken place. Unusually large numbers of reward claims per time period therefore suggests that this has taken place. (We note that the transmission rates must be placed in the context of what type of hardware is used. The hardware type is likely to be known by the service provider, so this does not cause any problem.)
The greedy ticket collection attack is likely to be the hardest attack to detect; especially if users scan for tickets of packets sent within the same cell as they resided, and if the users take pains to make the reported neighbors consistent with each other. However, should one party be found guilty of this attack, this is likely evidence that its common neighbors are, too.
- **Tampering with claims.** This attack is prevented by use of authentication techniques; the use of auditing tools does therefore not relate to the securing against it.
- **Reward level tampering.** If claimants indicate higher reward thresholds than that used for a given packet, this is an indication that the originator and some colluder close to the base station may perform this attack. Repeated evidence from different claimants, all pointing towards one and the same originator, provides strong evidence of the attack, in turn.

As for credit card fraud, use patterns can be employed to guard against attacks; the above description is meant only as evidence that the collected audit information is sufficient to detect and trace misbehavior. We are aware of further techniques to do so, and believe that there are further techniques we are not aware of. In fact, this problem is quite similar to intrusion detection, which has been studied for most existing and envisioned networks, including mobile ad hoc networks [25].

6 Conclusion

We have described an architecture for fostering collaboration between selfish nodes of multi-hop cellular networks, and have provided mechanisms to encourage honest behavior and to discourage dishonest behavior. To the best of our knowledge, no single paper was published so far on this issue.

Our security model is not formal: Instead, we list a set of potential abuses along with associated detection mechanisms. Thus, we propose to deal with fraud in a similar manner to how telecommunications companies and credit card companies do. A less heuristic approach would be a great step forward; however, this is a difficult task. Part of the reason for this is that not all packet forwarding information gets reported to the auditor (as not all tickets are winning), and that honest users may lose connectivity at any time. However, even if that were

not the case, a formal approach appears to be non-trivial. We hope that our contribution can be a first step in the direction of a formal treatment of the problem.

In terms of future work, we intend to work on this formalization. In addition, we will relax the assumption that all packets have to go through the backbone, by combining the proposed solution with an approach related to pure ad hoc networks, such as the one proposed in [3]. Moreover, we will explore the symmetric case of multi-hop cellular networks. Finally, we will estimate the performance of the proposed solution and propose appropriate optimizations whenever necessary.

Acknowledgements

We wish to thank Philippe Golle, Ari Juels and Ron Rivest for helpful discussions and feedback.

References

1. R. Anderson, H. Maniavas, C. Sutherland, "A Practical Electronic Cash System," In proceedings Fourth Cambridge Workshop on Security Protocols, 1996.
2. S. Buchegger, J.-Y. Le Boudec, "Performance Analysis of the CONFIDANT Protocol (Cooperation of Nodes: Fairness in Dynamic Ad-hoc NeTworks)," Proceedings of the Third ACM International Symposium on Mobile Ad Hoc Networking and Computing, Lausanne, June 2002 (MobiHoc 2002)
3. L. Buttyan, J. P. Hubaux, "Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks," ACM Journal for Mobile Networks (MONET), special issue on Mobile Ad Hoc Networks, October 2003, Vol. 8 No. 5
4. M. Cagalj, J. P. Hubaux, C. Enz, "Minimum-Energy Broadcast in All-Wireless Networks: NP-completeness and Distribution Issues," Proceedings of the Eighth ACM International Conference on Mobile Networking and Computing, Atlanta, September 2002 (Mobicom 2002)
5. O. Dousse, P. Thiran, M. Hasler, "Connectivity in Ad Hoc and Hybrid Networks", 21st Annual Joint Conference of the IEEE Computer and Communications Societies, New York, 2002 (Infocom 2002)
6. R. Hauser, M. Steiner, M. Waidner, "Micro-Payments Based on iKP," Technical Report 2791 (# 89269), June 1996.
7. H. Holma, A. Toskala, "WCDMA for UMTS", Wiley 2000
8. J.-P. Hubaux, L. Buttyan, S. Capkun, "The Quest for Security of Mobile Ad Hoc Networks," Proceedings of the Second ACM International Symposium on Mobile Ad Hoc Networking and Computing, Long Beach, October 2001 (MobiHoc 2001)
9. S. Jarecki, A. Odlyzko, "An Efficient Micropayment System Based on Probabilistic Polling," Financial Cryptography '97, pp. 173–191
10. C. Jutla, M. Yung, "PayTree: "amortized-signature" for flexible MicroPayments," Proceedings of Second USENIX Workshop in Electronic Commerce, pp. 213–221, 1996.
11. Y.-D. Lin, Y.-C. Hsu, "Multihop Cellular: A New Architecture for Wireless Communications", 19th Annual Joint Conference of the IEEE Computer and Communications Societies, Tel Aviv, 2000 (Infocom 2000)

12. C. R. Lin, "On-Demand QoS Routing in Multihop Mobile Networks", 20th Annual Joint Conference of the IEEE Computer and Communications Societies, Anchorage, 2001 (Infocom 2001)
13. M. Manasse, "Millicent (electronic microcommerce)," 1995. www.research.digital.com/SRC/personal/Mark_Manasse/uncommon/ucom.html.
14. S. Marti, Th. Giuli, K. Lai, M. Baker, "Mitigating Routing Misbehavior in Mobile Ad Hoc Networks," Proceedings of the Sixth ACM International Conference on Mobile Networking and Computing, Boston, August 2000 (Mobicom 2000)
15. A. Mehrotra, L. Golding, "Mobility and security management in the GSM system and some proposed future improvements," Proceedings of the IEEE, vol. 86, no. 7, July 1998, pp. 1480–1496.
16. S. Micali, R. Rivest, "Micropayments Revisited," CT-RSA 2002, pp. 149–163
17. N. Nisan, A. Ronen, "Algorithmic Mechanism Design," Proceedings of the 31st ACM Symposium on Theory of Computing, 1999, pp. 129–140
18. T. Pedersen, "Electronic Payments of Small Amounts," Technical Report DAIMI PB-495, Aarhus University, Computer Science Department, Aarhus, Denmark, August 1995.
19. C. Perkins, "Ad Hoc Networking," Addison Wesley, 2001.
20. M. Rahnema, "Overview of the GSM system and protocol architecture," IEEE Communications Magazine, vol. 31, no. 4, April 1993, pp. 92–100.
21. R. Rivest, "Electronic Lottery Tickets as Micropayments," Financial Cryptography '97, pp. 307–314
22. R. Rivest, A. Shamir, "Password and MicroMint – two simple micropayment schemes," Proceedings of 1996 International Workshop on Security Protocols, pp. 69–87, 1996.
23. D. Wheeler, "Transactions Using Bets," In proceedings Fourth Cambridge Workshop on Security Protocols, pp. 89–92, 1996.
24. A. Zadeh, B. Jabbari, R. Pickholtz, B. Vojcic, "Self-Organizing Packet Radio Ad Hoc Networks with Overlay (SOPRANO)," IEEE Communications Magazine, June 2002
25. Y. Zhang, W. Lee, "Intrusion Detection in Wireless Ad-Hoc Networks," Proceedings of the Sixth ACM International Conference on Mobile Networking and Computing, Boston, August 2000 (Mobicom 2000).
26. S. Zhong, Y. R. Yang, J. Chen "Sprite: A Simple, Cheat-proof, Credit-based System for Mobile Ad Hoc Networks", Technical Report Yale/DCS/TR1235, Department of Computer Science, Yale University, July 2002