CERIAS Tech Report 2004-80

A SECURE CREDITING PROTOCOL FOR HYBRID CELLULAR AND AD-HOC NETWORKS

by Bogdan Carbunar, Ioannis Ioannidis, Ananth Grama, Jan Vitek

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907-2086

A SECURE CREDITING PROTOCOL FOR HYBRID CELLULAR AND AD-HOC NETWORKS

Bogdan Cărbunar, Ioannis Ioannidis, Ananth Grama, Jan Vitek
Purdue University
West Lafayette, IN, USA
Email: {carbunar,ioannis,ayg,jv}@cs.purdue.edu

Keywords: hybrid cellular ad-hoc networks, security, routing protocols

Abstract: While wireless networking seems to be the way of the future, no definite architecture for large scale deployment of such networks has emerged. This can be attributed to financial reasons (a specialized infrastructure is too expensive to build) as well as to the lack of solutions that could work with an existing infrastructure or, in an ad hoc manner, without one. A viable alternative seems to be hybrid wireless networks. Such networks use the existing cellular telephony infrastructure as basis and enhance it by building ad hoc networks of traffic relayers around each cell, improving the overall throughtput and reliability of the network. These relayers are users of the network that are willing to operate as such. In this setting, both efficiency and security are vital properties. We propose SCP as an integrated solution for secure routing and crediting in hybrid networks. We describe how a secure environment can be established efficiently by financially motivating users to avoid attacks. Finally, we show that SCP imposes minimal load both in communication and computation, so that even regular cellular phones can function as relayers, without demanding infrastructure upgrades.

1 Introduction

The past decade has witnessed rapid developments in wireless communications, from wireless cellular telephony to ad-hoc networks, wireless LANs and RF networks. Wireless network cards have become affordable and wireless connections are fast enough for users to abandon more traditional networking possibilities, as long as there is a nearby access point. The only factor against an explosion of wireless computer networking is the necessity for an expensive infrastructure that can provide extensive and reliable coverage with suffi cient bandwidth.

Currently, the only infrastructure that addresses the above problem is that of cellular telephony. When a laptop equipped with a wireless network card connects to a base station the same way as cell phones, the bandwidth limitations are severe; the top rate achievable in a cell is 2.4Mbps (1xEV-DO) and the bandwidth drops fast as the device moves away from the base station. Upgrading cellular base stations can solve these problems, although it is doubtful that providers will be willing to make such a massive investment.

A solution that grafts ideas from ad hoc networks into cellular technology has started to attract attention. As in ad hoc networks, connections can involve several intermediate relayers. Since wireless LANs offer high throughput (IEEE 802.11b offers up to 11Mbps), albeit in a range of just 115m, using a web of multihop paths can considerably increase the throughput from the base station to the devices in its cell without requiring modifi cations in the infrastructure. An example of a hybrid network is shown in Figure 1. Device $\texttt{DMH}_1$ is within the range of the base station, but the expected downlink rate is very poor, as it lies near the edge of the covered area. Since the cellular throughput of $\texttt{MH}_2$ is larger than that of $\texttt{DMH}_1$, due to a shorter distance to the base station, $\texttt{DMH}_1$ can use $\texttt{MH}_1$ and $\texttt{MH}_2$ as cellular traffi c relayers, effectively increasing its throughput.

In principle, this is a simple and powerful idea, but multihop connections pose signifi cant problems. The major challenge is motivating users of the network to act as relayers. A relayer not only has to sacrifi ce some of its own bandwidth, but also battery power to transmit and receive. Rewarding these users with some form of credit to their account with the cellular provider is a reasonable incentive for participation. However, imposing a crediting scheme over a multi-

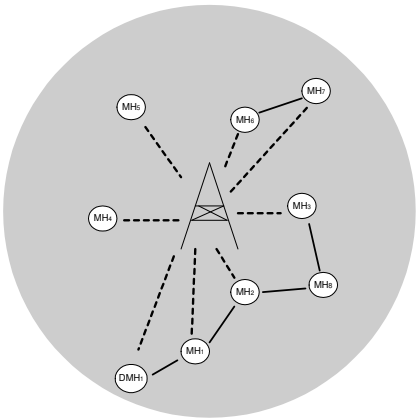

Figure 1: Example of a multihop cellular network. The gray area represents the coverage cell of the base station. Dashed lines represent cellular links, and full lines represent ad-hoc wireless links.

hop cellular network fundamentally changes the interactions between hosts. Malicious behavior will focus on gaining unwarranted financial rewards. Eavesdropping, impersonation, or denial of service attacks can be profitable and will be performed if so.

We assume henceforth that users are selfish and follow the principle that an action is to be performed only if the credit gain brought by it is worth the effort spent to perform the action. Even though attacks having personal satisfaction as the only gratification may still persist, we claim that on the long run they are not viable, since they necessitate credit, whose total amount is finite.

We propose a secure crediting protocol (SCP) where attacks from malicious hosts are deflected by making such attacks more expensive than honest transactions. An interesting side-effect of this approach is that some forms of attack are now welcome. As we will discuss in Section 5, nodes that want to trim the relaying path, so that they will receive a larger share of the credit for a route, act against congestion, without diminishing the performance of the network.

Finally, we demonstrate experimentally that the SCP protocol also decreases the traffic inside the entire network and the traffic to and from the base station, compared with the similar, but insecure, UCAN (Luo et al., 2003) protocol, without sacrificing the achieved throughput.

2 Related Work

There have been efforts to integrate infrastructure-based network models with ad hoc components, but most of them assume single-interface devices. In (Aggelou and Tafazolli, 2001), GSM terminals are used to relay information to other terminals to improve coverage. In Opportunity Driven Multiple Access (Rousse et al., 2002), transmission power is conserved by relaying traffic from a CDMA host to the base station through multiple, short hops. In (Wu et al., 2000), some channels are reserved for forwarding when the fixed channels become congested. In (Akyildiz et al., 1997), a generic wireless network is considered, where hosts contact a mobile base station for access outside their cell, using only one interface. In (Hsieh and Sivakumar, 2002), a hybrid network using the IEEE 802.11 architecture with both DCF and PCF modes is examined, using only one wireless interface. In (Lin and Hsu, 2000), multihop paths are used to decrease the number of base stations by increasing their coverage.

Although double-interface architectures are conceptually similar to their single-interface counterparts, they increase the overall capacity by using short-range, high-bandwidth, ephemeral channels to relay traffic and a long-range, low-bandwidth, permanent channel to complete operations like routing and data integrity confirmation or as a last resort in the absence of neighbors. The low-bandwidth channels are not necessarily cellular, but the already existing infrastructure make them an attractive option. This architecture has been examined in (Luo et al., 2003), whose routing scheme is very similar to SCP, but it generates considerably more network traffic. In (De et al., 2002), traffic is diverted to neighboring cells to increase throughput. The use of dedicated, stationary relays increases the cost of their solution and limits its utility. A study of local area hybrid networks is in (Lee et al., 2004).

As we have mentioned, security plays a major role in the hybrid environment. A protocol that solves the crediting and security issues for multihop paths to an infrastructure access point is presented in (Salem et al., 2003). However, attacks to the routing part of the protocol are not considered. There have been works towards secure routing protocols, although they do not consider attacks to the crediting scheme necessary for the architecture involved in this paper. For pure ad hoc networks, Ariadne (Hu et al., 2002) describes a secure protocol for Direct Source Routing, and ARAN (Dahill et al., 2002) presents a secure extension for AODV. In (Paul and Westhoff, 2002), a rate confirmation scheme is presented, a part of the SCP protocol that is usually left unaddressed in the literature.

3 Network Assumptions

The system that we consider is based on the UCAN (Luo et al., 2003) architecture, and consists of

a base station (BS) and several mobile hosts (MHs). The base station provides cellular coverage of a certain radius (up to 20km), and the mobile hosts located inside this coverage area have a permanent, direct, low bandwidth communication link to the base station. In addition, each MH has an 802.11b wireless interface that allows it to directly communicate with other MHs situated inside the smaller, 802.11b, coverage range(115m).

All the mobile hosts are registered with the base station, that is, with the service provider that owns the base station. The registration implies that each mobile host MH has a unique identifier, $\mathtt{Id}_{\mathtt{MH}}$ given by the base station, and each mobile host has an account with the base station. The base station provides certain services, such as Internet access, to each registered host, for as long as the host has credit in its account. Moreover, the base station and the mobile hosts that it serves have synchronized clocks. This can be easily done by the base station, either by periodically beaconing its current time, or by piggybacking the current time in the messages sent to the mobile hosts.

We assume that the base station has a private/public key pair, and each host knows the public key. The base station also shares a one-way hash function $\mathtt{H}$ with all the hosts, and a different secret key with each host MH, denoted as $\mathtt{K}_{\mathtt{MH}}$. We also assume that the base station can provide a good estimate of the downlink rate $\mathtt{R}_{\mathtt{MH}}$ for each MH.

4 The Proposed Protocol: SCP

SCP has three parts, see Figure 2. In the first part, the initiator $\mathtt{A}$ contacts potential relayers and in the second a secure route is found from the base station to $\mathtt{A}$. The last part concerns the credit handling and the secure forwarding of packets, sent from the base station to $\mathtt{A}$, through the relayers. The protocol is repeated every time the initiator is not satisfied with the rate at which it receives the information or if the route is broken during the protocol.

Every time SCP is run, the initiator has to be charged by the base station. The charges consist of a fee for the service provider, a fee for initiating the protocol and the credit to be given to the eventual relayers, if the resulting route is satisfactory to the initiator. The fees for the provider and for initiating the protocol act as a safeguard against denial of service attacks by hosts that simply request a route, without intending to use it. Since the account of the attacker has to be charged, it is easy to detect malicious patterns and, anyway, the attacker has to pay for the attack. Also, because most of the hosts participating in a run of the protocol will never directly receive credit for that run, a form of credit for all candidate relayers must be provided. A suggestion could be that the initiating fees be periodically split among the hosts, according to their participation time.

We consider crediting functions only of the following form:

$$\begin{aligned} \mathtt{C}_{\mathtt{A}} &= \mathtt{C}_{\mathtt{BS}}(\mathtt{S}) + \mathtt{C}_{\mathtt{SCP}} + \mathtt{C}_{\mathtt{R}}(\mathtt{S},\mathtt{r}), \\ \mathtt{C}_{\mathtt{R}}(\mathtt{S},\mathtt{r}) &= \mathtt{G}(\mathtt{S}) \cdot \mathtt{F}(\mathtt{r}), \\ \mathtt{F}(\mathtt{r}) &= \mathtt{c}, \text{ if } \mathtt{r} \leq 2, \\ \mathtt{F}(\mathtt{r}) &= \mathtt{k}_2 - \mathtt{k}_1 \cdot \mathtt{r}, \text{ if } \mathtt{r} > 2 \end{aligned}$$

where $\mathtt{C}_{\mathtt{A}}$ is the total credit taken from the initiator, $\mathtt{S}$ is the size of the information downloaded by the initiator from the base station, $\mathtt{C}_{\mathtt{BS}}$ is the fee for the provider, as a function of $\mathtt{S}$, $\mathtt{C}_{\mathtt{SCP}}$ is the initiating fee, and $\mathtt{C}_{\mathtt{R}}$ is the total credit given to the relayers, which is a function of $\mathtt{S}$ and of $\mathtt{r}$, the number of relayers. The function $\mathtt{G}(\mathtt{S})$ denotes the participation of $\mathtt{S}$ to $\mathtt{C}_{\mathtt{R}}$. The only part that somewhat restricts the form of $\mathtt{C}_{\mathtt{A}}$ is $\mathtt{F}(\mathtt{r})$. Its role in security and what the values of $\mathtt{c}$, $\mathtt{k}_1$ and $\mathtt{k}_2$ should be will be discussed in Section 5. The role of $\mathtt{C}_{\mathtt{BS}}(\mathtt{S})$ and $\mathtt{C}_{\mathtt{SCP}}$ will be detailed in Section 4.3. The amount of credit that each relayer gets is $\mathtt{C}_{\mathtt{R}}(\mathtt{S},\mathtt{r})/\mathtt{r}$.

We believe that this crediting scheme can be applied in any realistic setting. The fees for the provider and initiating the protocol can be of any form. The impact of $\mathtt{F}(\mathtt{r})$ on the credit the relayers receive can be scaled using $\mathtt{G}(\mathtt{S})$. As we will show, $\mathtt{F}$ provides security, without resorting to expensive cryptography, by maximizing the benefit of the hosts when they behave honestly.

4.1 Request Forwarding

The first part of SCP starts when $\mathtt{A}$ contacts the base station, using its uplink connection, with the session initialization message SINIT

$$[\mathtt{SINIT}, \mathtt{Id}_{\mathtt{A}}, \mathtt{R}_{\mathtt{A}}, \mathtt{Seq}_{\mathtt{A}}, \mathtt{H}_{\mathtt{K}_{\mathtt{A}}}(\mathtt{Id}_{\mathtt{A}}, \mathtt{R}_{\mathtt{A}}, \mathtt{Seq}_{\mathtt{A}})],$$

where $\mathtt{Id}_{\mathtt{A}}$ is the unique identifier of $\mathtt{A}$, $\mathtt{R}_{\mathtt{A}}$ represents $\mathtt{A}$'s downlink rate, $\mathtt{Seq}_{\mathtt{A}}$ is a sequence number maintained by $\mathtt{A}$, incremented by $\mathtt{A}$ each time the protocol is run, and $\mathtt{K}_{\mathtt{A}}$ is the secret key shared by $\mathtt{A}$ and BS. $\mathtt{H}_{\mathtt{K}_{\mathtt{A}}}(\mathtt{M})$ represents the MAC (Message Authentication Code) of message $\mathtt{M}$, with the key $\mathtt{K}_{\mathtt{A}}$. The base station first checks the validity of the MAC, whose purpose is to convince BS of the authenticity of the message, since only $\mathtt{A}$ knows $\mathtt{K}_{\mathtt{A}}$. The sequence number is used to prevent a replay attack, if anyone captures a previous SINIT message.

As stated in Section 3, we assume that every host has an account with the base station. The base station first removes $\mathtt{C}_{\mathtt{SCP}}$ credit from $\mathtt{A}$'s account. Then

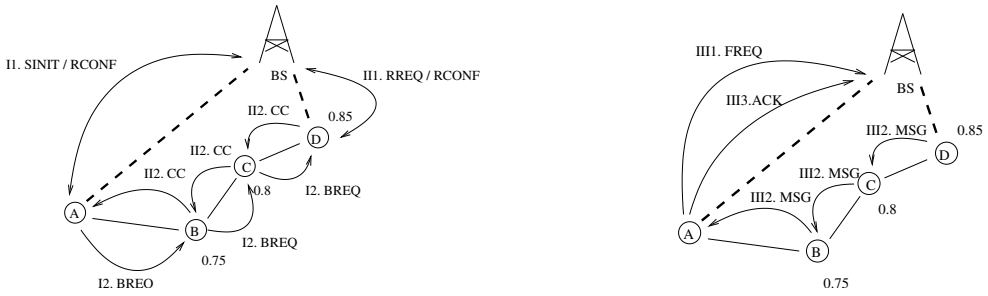

Figure 2: Diagram of the SCP protocol. Mobile hosts are represented with small circles containing the names of the host in the center. Arcs represent messages. Messages are prefixed with Roman numerals to represent the SCP phase. (a) (I) Request forwarding and (II) path selection, see Section 4.1 and 4.2 (b) (III) Relaying of packets, see Section 4.3.

it replies with a downlink rate confi rmation message RCONF

$$[\mathtt{RCONF}, \mathtt{T_{BS}}, \mathtt{Stamp_A^{BS}}],$$

where $\mathtt{T_{BS}}$ is the current time of the base station and $\mathtt{Stamp_A^{BS}} = \mathtt{S_{BS}}(\mathtt{H}(\mathtt{Id_A}, \mathtt{R_A}, \mathtt{Seq_A}, \mathtt{T_{BS}}))$. $\mathtt{S_{BS}}(\mathtt{M})$ represents the message M signed with the private key of the base station. The purpose of this exchange of messages is to provide A with a receipt from the base station that it has registered this run of the protocol and has the necessary funds.

In the next step A sends to all the devices that are inside its ad-hoc wireless transmission range the broadcast request message BREQ

$$[\mathtt{BREQ}, \mathtt{Id_A}, \mathtt{Seq_A}, \mathtt{TTL}, \mathtt{R_A}, \mathtt{T_{BS}}, \mathtt{d_A}, \mathtt{Stamp_A^{BS}}],$$

where TTL is the time-to-live parameter of the message, namely the number of hops the message is supposed to travel, and $\mathtt{d_A}$ is the distance in hops to the initiator, which for A is 0.

Upon receiving such a BREQ message from a neighbor P, a device N checks to see if for $\mathtt{Id_A}$ it has ever seen a sequence number larger than or equal to $\mathtt{Seq_A}$, in which case it drops the packet. In order to be able to do this, for every initiator that has sent a packet through N, N has to store the host identifi er and the largest sequence number it has seen for that host, along with the distance to the initiator, $\mathtt{d_A}$. Then N checks the freshness of the message, by making sure that $\mathtt{T_N} \leq \mathtt{T_{BS}} + \mathtt{TTL} \cdot \mathtt{T_{proc}}$. $\mathtt{T_N}$ is the current time at host N, $\mathtt{T_{BS}}$ is the timestamp carried in the BREQ message and $\mathtt{T_{proc}}$ is an upper bound on the time necessary for each host to process and forward a BREQ message. Finally, N checks the BREQ message validity condition $\mathtt{H}(\mathtt{Id_A}, \mathtt{R_A}, \mathtt{Seq_A}, \mathtt{T_{BS}}) = \mathtt{V_{BS}}(\mathtt{Stamp_A^{BS}})$, using the values received in the BREQ message. $\mathtt{V_{BS}}$ means encrypting with BS's public key, known to all MHs, which in this context is equivalent with verifying BS's signature. If the condition does not hold N drops the message.

If the condition holds, N marks P as its parent in the breadth-fi rst tree initiated by A and confi rms to P the choice made. All the confi rmation messages that P receives make P aware of all its direct successors in the breadth-fi rst tree of A, information that will be useful in the convergence part of the protocol, described in the following. If $\mathtt{TTL} = \mathtt{d_A}$, N drops the message. Otherwise, N compares its downlink rate, $\mathtt{R_N}$, with $\mathtt{R_A} - \delta$, where $\delta \geq 0$ is a small constant chosen by the protocol. If $\mathtt{R_N} < \mathtt{R_A} - \delta$, then N drops the message. This is because most probably the neighbors of N whose downlink rate is larger than $\mathtt{R_A} - \delta$ have already been reached by this BREQ message. However, if $\mathtt{R_N}$ is larger than $\mathtt{R_A} - \delta$, it means that there is a high chance that N's neighbors, not yet reached by A's BREQ message, can provide a still better downlink rate. In this case, N increments $\mathtt{d_A}$, decrements the TTL fi eld, and broadcasts this message to all the devices inside its transmission range.

4.2 Path Selection

A mobile host L that receives the BREQ message initiated by A, with $\mathtt{TTL} = \mathtt{d_A}$, or with $\mathtt{R_L} < \mathtt{R_A} - \delta$, or that does not have any successors, becomes a leaf in the breadth-fi rst tree of A. Such a device initiates a converge-cast operation meant to reach A. The purpose of the converge-cast operation is to convey to the initiator host A the best path to the base station. If mobile host L is a leaf in A's broadcast tree because $\mathtt{R_L} < \mathtt{R_A} - \delta$, then L, which we will refer to as a *dead-end* leaf, sends to its parent host the converge-cast message

$$[\mathtt{CC}, \mathtt{Id_A}, \mathtt{Seq_A}, \mathtt{Id_L}, \mathtt{R_L}],$$

containing the leaf's downlink rate. However, if L is a leaf because the TTL of the BREQ message was 0 or because it has no successors, but L has a downlink rate larger than that of A, L may become a relayer.

Whether L is willing to become one is a decision that L can make on locally available information. An important factor is the distance from A, because it determines the credit that L can expect to receive, which is $C_R(S, d_A)/d_A$. Other factors, such as the battery power or congestion concerns can contribute to the final decision.

If L decides to participate, it first contacts the base station, through its uplink connection, with a rate request message $[RREQ, Id_L, H_{K_L}(Id_L)]$ When BS receives a RREQ message, it first checks the validity of the MAC, in order to authenticate the request. The base station then looks up the downlink rate for L, and replies with a rate confirmation message $[RCONF, Id_L, R_L, T_{BS}, Stamp_L^{BS}]$, where $Stamp_L^{BS} = S_{BS}(H(Id_L, R_L, T_{BS}))$. L then sends to its parent host the converge-cast message

$$[CC, Id_A, Seq_A, Id_L, R_L, T_{BS}, d_A(L), Stamp_L^{BS}],$$

where $d_A(L)$ represents the number of hosts from L to A. Each intermediate node N in the broadcast tree of A waits to receive a CC message from all its direct successors in this tree. For each such message, N checks to see that the advertised number of hops to the initiator, d_A, is strictly larger than its own. This is because N only receives CC messages from devices that are its successors. This is a simple check that prevents hosts from advertising shorter distances to the initiator. After receiving the CC messages from all its successors, N first discards those from dead-end leaves. Then, it compares its downlink rate with the rates received in the CC messages from each successor S. If its rate is no smaller than the largest rate reported from its sub-tree and strictly larger than the downlink rate of the initiator, and N decides it wants to participate, N can drop its subtree. In this case, it first contacts the base station with a $RREQ$ message and obtains $Stamp_N^{BS} = S_{BS}(H(Id_N, R_N, T_{BS}))$, as described in the previous paragraph. Then it sends to its parent a CC message $[CC, Id_A, Seq_A, Id_N, R_N, T_{BS}, d_A(N), Stamp_N^{BS}]$

However, if there is a host with a better rate in the subtree, N receives from a successor S_j a CC message:

$$[CC, Id_A, Seq_A, L_i, R_j, T_{BS}, d_A(S_j), Stamp_{S_j}^{BS}],$$

$L_i = Id_{S_j}, ..., Id_{S_1}$. The chain has originated at host S_1, and extended by all the hosts in the list L_i. The downlink rate R_1 of S_j is the maximum among those received by N from its subtree and R_1 is also strictly larger than R_N. In this case, N decides if it wants to participate in the protocol based on the credit that it will receive, which is $C_R(S, d_A(S_1))/d_A(S_1)$. N can calculate the value $d_A(S_1)$, the distance from S_j to A, since it should be equal to the distance from N to A plus the number of intermediate hosts in the CC message. If N is satisfied with the credit, it only appends its identifier to the beginning of the list L_i of identifiers in the above CC message and forwards it to its parent

$$[CC, Id_A, Seq_A, Id_N, L_i, R_j, T_{BS}, Stamp_{S_j}^{BS}]$$

When the process converges to A, the retrieved paths will consist of hosts aware of the credit that they may receive and willing to participate. A can choose the optimal path. At this point, A will start the last phase of the protocol, which will notify the relayers of their status and the base station to start the file transfer and credit the relayers.

4.3 Crediting and Relaying

In the last part of the protocol, A contacts the base station and sends the entire route, $L_A = Id_{N_r}, \ldots, Id_{N_1}$, where N_r is the first relayer from the base station and N_1 is the direct relayer to A. More precisely, A sends to BS a file request message

$$[FREQ, Id_A, Seq_A, f, L_A, H_{K_A}(Id_A, Seq_A, f, L_A)]$$

where f is the name of the file requested by A. The purpose of the MAC is to authenticate the originator of the message as A, since A and BS are the only ones that know K_A. It also prevents other MHs from trying to impersonate A and acquire information concerning A from the base station.

The base station first removes $C_{BS}(S)$ credit from A's account, which is due by A to the base station for sending it the information requested. Then it retrieves the actual information requested by A, $INFO_f$, of size S and breaks it into packets of size s, to allow the crediting to be done at a smaller, packet level granularity. That is, each of the relayers receives credit for each packet sent by BS that reaches A. More precisely, for each packet P_i, $i = 1, \ldots, \lceil S/s \rceil$, the base station first removes $C_R(s, r)$ credit from A's account. The base station then generates a source routing message of type MSG destined to A, containing the identifiers of the relayers along with the packet P_i:

$$[MSG, T_{BS}, Seq_A, Tr_r^{BS}, \ldots, Tr_i^{BS}, Packet_A^{BS}],$$

where $Tr_i^{BS} = (Id_{N_i}, H_{K_{N_i}}(Id_{N_i}, T_{BS}))$ and $Packet_A^{BS} = Id_A, P_i, H_{K_A}(P_i, T_{BS})$. Each intermediate relayer N looks at the first Tr field searching for its identifier Id_N and checks the field's validity. If the check is passed, N peels off the first Tr field from the message and forwards it to the next host specified in the message.

A acknowledges the reception of a packet by sending to BS, through the direct uplink cellular connection, an ACK message containing $H_{K_A}(P_i)$. Only after confirmation, will BS provide the credit $C_R(s, r)/r$ to the relayers, and send A the next packet P_{i+1}.

If A receives an invalid MAC, it contacts the base station with a NACK message for packet P_i. Then

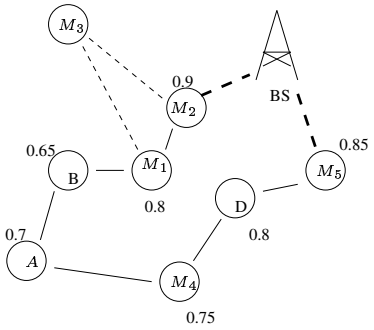

Figure 3: Example of potential attacks. A is the initiator and BS is the base station. Hosts $M_1, \ldots, M_5$ are potentially malicious. The numbers around the hosts represent their downlink rate, in Kbps. For $\delta \geq 0.05$, B will propagate A's request.

the base station removes Δ_{BS} more credit from A's account and subtracts Δ_R from the credit that each relayer should get for the packet, $C_R(s,r)/r$. It then resends the packet P_i with a new MSG packet, with an incremented Seq_A value, via the same path. Each relayer, upon seeing a retransmission, knows that it will receive less credit from the base station, and therefore it can opt out by dropping the packet. The initiator A waits for a given timeout interval to receive the retransmission, after which it initiates a new SCP protocol, to find an alternate path of relayers. The purpose of Δ_{BS} and Δ_R is to prevent an attacker, acting either as initiator or relayer, from denying service or draining the battery of other users.

5 Security Analysis of SCP

In this section we analyze SCP's security.

Security against inflated downlink rates SCP relies on the downlink rates advertised by the MHs reached by the BREQ message in order to choose the best path. Malicious hosts can try to falsify their downlink rates so as to become the first relayers from the base station and receive undue credit. This is a critical attack because A may pay for undesirable service. However, SCP requires each host N that wants to become a relayer to contact the base station and obtain a signed certificate of its downlink rate, $S_{BS}(H(Id_N, R_N, T_{BS}))$. Since the certificate contains the timestamp of the base station, a host is prevented from using older, stale downlink rates.

Moreover, another host M cannot eavesdrop on the RREQ/RCONF protocol of the host N, in an attempt to impersonate N and its larger rate R_N. This attack would not work for two reasons. First, the account that would receive credit would be that of N. Second, the MSG packets will have to go through N instead of M, and since N's neighbors are different from M's, M cannot provide a valid path that includes N as a relayer.

Security against adding invalid relayers Figure 3 shows an example where hosts M_1 and M_2 collude to include host M_3 to the list of relayers, during the converge-cast towards the initiator A. M_3 is not necessarily a direct neighbor of M_1 or M_2, but it may have a path to them. M_3 could also be another host taken over by the owners of M_1 and M_2. Note that a single host cannot add false relayers. A way to prevent such attacks is to make sure that the sum of credit the colluding hosts receive is not greater than what they would receive if they behaved honestly. We show how the structure of the crediting function C_A of Section 4 and more specifically, the structure of $F(r)$, which captures the dependence of C_A on the number of relayers, guarantees this property.

In the following, we will assume that there are r honest relayers, m true, but colluding, relayers and l false relayers, added by the colluding ones. We also assume that the sum of the credit that goes to the $m + l$ malicious hosts is divided equally among the m colluding relayers. This is the worst case scenario, as the l false relayers are assumed to be hosts taken over and the credit they receive can be funneled to the m colluding hosts. The credit each host receives without adding false relayers is $F(r+m)/(r+m)$, while the credit each colluding host receives when adding l false relayers is $(l+m) \cdot F(r+m+l)/[m \cdot (r+m+l)]$. We can prove that there are values of c and k_1 and k_2 such that the credit each colluding host receives decreases for $l > 0$.

Theorem 5.1 *There are* k_1, k_2 *and* c *such that, when* $TTL \leq 18$,

$$\frac{F(r+m)}{r+m} \geq \frac{l+m}{m} \cdot \frac{F(r+m+l)}{r+m+l}.$$

Proof A possible solution is $k_1 = 1$, $k_2 = 9$ and $c = 7$. In fact, there is an infinite number of solutions, which can be derived by varying the value of c and solving the inequalities system for k_1 and k_2.

Security against removing legitimate relayers Another attack involves M_4 and M_5 colluding to remove D from the set of hosts receiving credit (Figure 3). This is impossible, since D has to be notified by BS that it will receive credit before it forwards a packet and at that point the colluding hosts cannot remove D from the credited path.

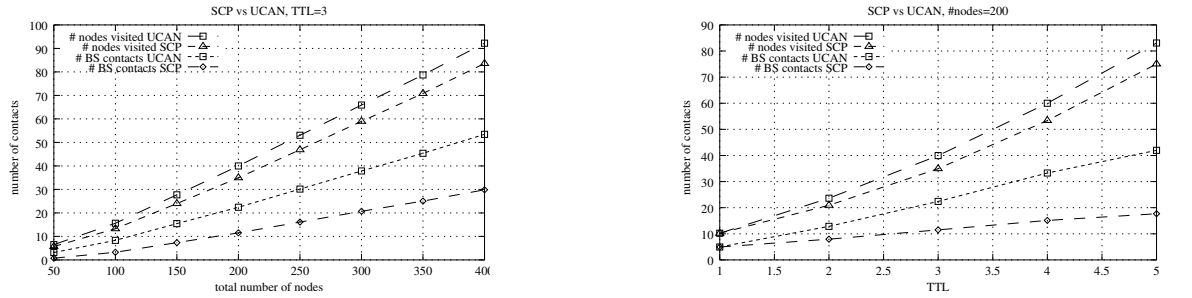

Figure 4: Experimental results comparing SCP and UCAN. (a) Evolution of the average number of devices contacted and of average number of base station contacts during SCP and UCAN, when the total number of hosts served by the base station grows from 50 to 400. TTL was set to 3 for both SCP and UCAN. (b) Evolution of the same two metrics with a TTL increasing from 1 to 5, when the total number of hosts is 200.

Pruning the relayer path In Figure 3, host B, during the converge-cast part of the protocol presented in Section 4.2 can decide to ignore the CC message originated at M_2, conveying a larger downlink rate than that of B. The purpose of B, as in the previous attack, is to present to A a shorter relayer path and therefore obtain more credit from the base station. In this particular example, B would not have to split the credit with anyone.

There are two observations we can make in favor of this attack. The first is that the characterization of this behavior as malicious is rather tenuous. It is no different from a host refusing to relay packets because the credit it will receive is inadequate. The relaying path is a valid one and the credit given to the relayers is correct. The only adverse effect is that A might have a better and longer path to the base station. However, if the pruned path is chosen, it means that the download rate it provides is at least satisfactory to A.

The second observation is that pruning can be beneficial to the network. Longer paths may have better rates, but they consume bandwidth and create congestion. Also, from A's point of view, a longer path is less robust, as even a single host can break it by moving. A longer path incurs larger delays, as well. If hosts decide to optimize their chances of being relayers against the credit they will receive, the operation of the whole network will benefit from a reduced load, without compromising connectivity.

6 Experimental Results

In this section we experimentally analyze the performance of SCP. Our goal is to measure the overhead introduced by the secure crediting protocol and compare it with UCAN. We will show that even though more secure than UCAN, SCP significantly decreases the network traffic, while maintaining similar throughput performance. We modeled the ad-hoc wireless network using the unit disk graph model, where all hosts have the same transmission range, 115 units, and two hosts are neighbors iff they are inside each other's transmission ranges. The positions of the hosts were randomly chosen, inside a square of size 886 units. We modeled the cellular access of the hosts using the HDR downlink rate vs. distance dependency graph presented in (Luo et al., 2003). We chose the position of the base station to be in the center of the 886×886 square, and the cellular transmission range of the base station to be 500 units. According to this model, each mobile host inside the square is covered by the cellular transmission range of the base station. For the mobility scenario, we used the random waypoint model. Each host chooses a target destination and moves towards it with a random speed between 1 and the maximum speed. After reaching the destination, the host chooses a new destination and a new random speed.

6.1 Network Load

For the first set of experiments we consider that all the hosts are static and we choose one host at distance 400 units from the base station to be the initiator of the protocol. Each point on the plots is computed as an average over 500 different ad-hoc network configurations. Our goal is to measure the number of messages produced by a protocol run and the number of times the base station will be contacted. Figure 4 compares SCP with UCAN's proxy discovery algorithm. Figure 4(a) shows the evolution of the two metrics when the total number of hosts in the square increases from 50 to 500 and for TTL 3. Figure 4(b) plots the evolution of the two metrics when TTL ranges from 1 to 5, but the total number of hosts is 200. Both plots show that SCP constantly requires fewer hosts to be contacted and fewer hosts to contact the base station

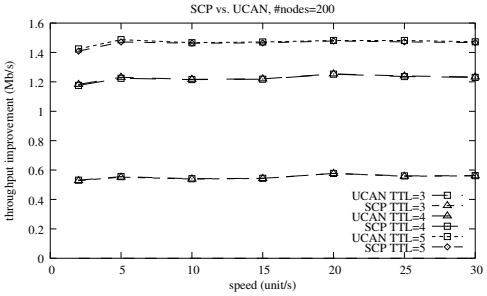

Figure 5: Throughput measurements for SCP and UCAN when 200 hosts move at maximum speeds ranging between 1 and 30units/s.

than UCAN. The reason is that UCAN will forward a request to a TTL range, while SCP stops if the downlink rate falls below a threshold. This way, nodes at the fringes of the coverage area are not contacted.

6.2 Throughput

This section describes the throughput improvement measurements for SCP and UCAN. The experiments were run in the following way. For each of the maximum speeds of 2, 5, 10, 15, 20, 25 and 30 units/s, 10 random initial network confi gurations with 200 hosts were chosen. The experiment was performed for 400 seconds for each such confi guration. During the experiment, each host moved continuously according to the random waypoint model. We have addressed the problem of the random waypoint mobility model identifi ed in (Yoon et al., 2003) by choosing the speed of each device to be uniformly distributed in the interval $[1, \mathtt{MS} - 1]$, where MS is the maximum speed, and by discarding the fi rst 200 movements for each experiment, so that the initial confi gurations do not bias the results. We have performed the experiment for TTL values of 3, 4, and 5.

Figure 5 shows the results of the experiments. For the same TTL, the curves of SCP and UCAN are almost identical, with UCAN having only a nominal advantage, negated by a higher network load. The improvement was measured against the throughput of the link to the base station.

7 Conclusions

We have presented a routing and crediting protocol that establishes a cooperative environment for hybrid networks. SCP is secure against attacks to the crediting scheme. This was achieved by using a special crediting function that does not reward longer paths. The advantage of our approach is that it minimizes the number of expensive cryptographic operations needed and relies on the users being rational. We showed experimentally that the elevated security of SCP does not come at the expense of the network load or the effi ciency of the routes.

REFERENCES

Aggelou, G. and Tafazolli, R. (February 2001). On the relaying capacity of next-generation gsm cellular networks. In *IEEE Personal Communications Magazine, 8(1):40-47*.

Akyildiz, I., Yen, W., and B.Yener (1997). A new hierarchical routing protocol for dynamic multihop wireless networks. In *Proceedings of IEEE INFOCOM*.

Dahill, B., Levine, B. N., Royer, E., and Shields, C. (2002). A secure routing protocol for ad hoc networks. In *Proceedings of ICNP*.

De, S., Tonguz, O., Wu, H., and Qiao, C. (2002). Integrated cellular and ad hoc relay (icar) systems: Pushing the performance limits of conventional wireless networks. In *Proceedings of HICSS-37*.

Hsieh, H.-Y. and Sivakumar, R. (2002). On using the ad-hoc network model in wireless packet data networks. In *Proceedings of ACM MOBIHOC*.

Hu, Y.-C., Perrig, A., and Johnson, D. (2002). Ariadne: A secure on-demand routing protocol for ad hoc networks. In *Proceedings of ACM MOBICOM*.

Lee, S., Banerjee, S., and Bhattacharjee, B. (2004). The case for a multi-hop wireless local area network. In *Proceedings of IEEE INFOCOM*.

Lin, Y.-D. and Hsu, Y.-C. (2000). Mulithop cellular: A new architecture for wireless communications. In *Proceedings of IEEE INFOCOM*.

Luo, H., Ramjee, R., Sinha, P., Li, L., and Lu, S. (2003). Ucan: A unified cellular and ad-hoc network architecture. In *Proceedings of ACM MOBICOM*.

Paul, K. and Westhoff, D. (2002). Context aware inferencing to rate a selfish node in dsr based ad-hoc networks. In *Proceedings of IEEE GLOBECOM*.

Rousse, T., Band, I., and McLaughlin, S. (2002). Capacity and power investigation of opportunity driven multiple access (odma) networks in tdd-cdma based systems. In *Proceedings of IEEE ICC*.

Salem, N., Buttyan, L., Hubaux, J.-P., and Jakobsson, M. (2003). A charging and rewarding scheme for packet forwarding in multi-hop cellular networks. In *Proceedings of ACM MOBIHOC*.

Wu, X., Chan, S.-H., and Mukherjee, B. (2000). Madf: A novel approach to add an ad-hoc overlay on a fixed cellular infrastructure. In *Proceedings of IEEE WCWN*.

Yoon, J., Liu, M., and Noble, B. (2003). Random waypoint considered harmful. In *Proceedings of IEEE INFOCOM*.