

SIPCache: A Distributed SIP Location Service for Mobile Ad-Hoc Networks^{*†}

Simone Leggio[†] Hugo Miranda[‡] Kimmo Raatikainen[†] Luís Rodrigues[‡]
[†]University of Helsinki [‡]University of Lisbon

August 16, 2006

Abstract

Internet-based communication is currently in a hype. People utilize Internet services more and more to communicate with each other, e.g., via VoIP or chat. The next step would be to reutilize the same applications to achieve ubiquitous communication, anytime and anywhere, also where network support is not available, such as in ad-hoc networks. Existing Internet protocols must be modified for working in the ad-hoc serverless environment. The Session Initiation Protocol (SIP) is a fundamental element in the Internet for establishing multimedia communication sessions. However, SIP cannot be used in ad-hoc networks, as it relies on the support of SIP servers in the network; e.g, in ad-hoc networks it is not possible to locate SIP users since the assistance of a dedicated SIP server is missing. A solution for this problem is presented in this paper, which describes and evaluates a fully decentralized mechanism for locating SIP users in ad-hoc networks.

1 Introduction

Mobile Ad-Hoc Networks (MANET) are an interesting networking environment; their characteristics of flexibility, dynamism and deployment easiness make them suitable to support the ubiquitous computing paradigm. Mobile and ubiquitous computing has several facets; this paper focuses on what we call *ubiquitous communication*. With this term we refer to the possibility of communicating, anytime anywhere, relying only on the terminal own communication capabilities. Thus, even if an operator-provided network access is not available, people would be able to engage in VoIP, video calls, chat, or any other communication sessions, by exchanging data directly with the other users, without passing through the infrastructured network.

To fully exploit the potentialities of ubiquitous computing, it is necessary to build a functional framework that allows discovering the users available within the ad-hoc network. Once the contact (and possibly profile) information of the

^{*}This work was partially supported by the MiNEMA programme of the European Science Foundation and by FCT project P-SON, POSC/EIA/60941/2004 through FCT and FEDER. Contact author e-mail: simone.leggio@cs.helsinki.fi

[†]An extended version of this paper appeared in *Mobiquitous 2006*.

desired user(s) is located, the desire to begin a communication session must be communicated. For this it is necessary to use a signaling protocol that negotiates all the parameters necessary for the communication session. The final step is to build on top of the signaling protocol a wide set of applications, to offer the end user a rich set of communication capabilities.

The Session Initiation Protocol (SIP) [5] is a powerful signaling protocol designed by the IETF to initiate and manage multimedia communication sessions in the Internet. However, SIP relies on network entities to assist end users in their operations. E.g., SIP user discovery in MANETs is not possible due to the lack of centralized servers, that is, it is not possible to use the so called *SIP location service*.

In this paper, we provide and evaluate a solution for performing SIP user discovery in MANETs, called SIPCache. Our proposal combines two separate solutions: the first is a framework for deploying SIP in MANETs without support of centralized servers, called decentralized SIP, or dSIP, [3]. The second solution is an algorithm, called PCache [4] for efficiently distributing cached data among nodes of a MANET. SIPCache applies the abstract PCache algorithm to a concrete use case, dSIP, to build a distributed SIP location service for MANETs.

2 SIPCache

Decentralized SIP [3] is a solution for deploying SIP in small scaled ad-hoc networks, optimized for few dozens of nodes, preferably in link local environment; dSIP breaks the need that user agents have for centralized servers and makes SIP, and thus, real time applications, truly ubiquitous. The main idea is to embed a subset of SIP proxy and registrar server functionalities in each end node, and build a middleware architecture that can control the operations in ad-hoc networks. By embedding registrar functionalities in each node, we make the location service distributed. In dSIP, the location service is exploited and maintained by spreading in the network a SIP REGISTER message, containing the contact information (binding) of the registering user. Nodes in the network may reply to the REGISTER with a 200 OK message, containing the replying node's binding as well. With a mere 2 messages signaling exchange, each node is given, in theory knowledge of the bindings of all the other nodes in the network. Decentralized SIP also allows "targeted" registrations, meaning that nodes can look for a specific SIP user; in that case, only the specific user will reply to the registration. The main limitation of dSIP is that each node must store the bindings of all the other nodes in order to have a complete view of the network users. This approach may not be feasible in very large ad-hoc networks, and for this reason, dSIP has been combined with PCache.

PCache [4] is an algorithm for efficiently distributing and retrieving cached data items in ad-hoc networks. PCache is an epidemic algorithm, but contrary to usual applications of epidemic protocols, the algorithm does not aim at delivering every message to all participants. Instead, and taking into consideration the limited resources of the devices, the algorithm tries to geographically distribute the replicas of the stored items. Data items are replicated using a gossip algorithm requiring a small number of messages. The dissemination component of PCache makes use of three complementary mechanisms: an efficient best-

effort broadcast mechanism; a distributed algorithm for deciding which nodes replicate a given data item and a data shuffling mechanism to improve the distribution of data replicas in the nodes caches. PCache prevents broadcast storms by limiting the number of forwarding nodes. Furthermore, nodes are selected for forwarding depending on the expectations of the additional coverage that their transmission will provide. The Data Gathering mechanism of PCache is used to retrieve more than a single item within the same query message. The query is performed by specifying some condition that must be satisfied by the keys. Data Gathering operations may produce a large number of redundant messages in both the dissemination of the query and on the forwarding of the replies with obvious implications in the resource consumption of the devices. PCache relies on the efficient replication of data items to impose a limit on the distance (in number of hops) that the query travels and reduce message redundancy.

SIPCache applies PCache concepts to dSIP: the cached data items are the contact data of SIP users in the MANET (bindings); dSIP bindings are distributed and retrieved in the network according to the PCache algorithm, rather than simply broadcast or multicast as dSIP suggests. SIPCache thus, extends and optimizes the functionalities of dSIP to large multi-hop ad-hoc networks. SIPCache enables the two types of location service provided by dSIP. The first type consists in retrieving the IP address of a single user based on the user name used as a search parameter.

The second type corresponds to the general registration in dSIP, where the nodes may reply to the node that has spread a REGISTER message with a 200 OK message containing their bindings. SIPCache slightly modifies and enhances this procedure. The enhancement comes from the fact that the registering node may specify the characteristics of the users from which a reply is desired: e.g., if SIPCache is used to establish a VoIP call between users in proximity, a user may broadcast his desire to talk with people interested in sports, or music, rather than art or informatics.

The modification deals with how the replies are returned, which is managed by the interaction of dSIP with the PCache Data Gathering mechanism, and with the fact that a general registration in SIPCache is logically treated as a query. Flooding a query for a user matching a given profile and returning a reply without PCache would imply, as preliminary tests showed, a large number of messages exchanged altogether, due to the high number of nodes that would return a reply. Instead, replying by following PCache algorithm allows saving the number of messages exchanged, which is beneficial from the bandwidth and device resources consumption point of view.

2.1 Mapping of SIP Messages to PCache

The interaction between SIP and PCache was achieved by directly mapping SIP messages to the format defined for PCache. SIP messages are distributed into the network according to the procedure defined by the PCache algorithm, and their header fields are processed as the correspondent PCache header field. The correspondence between SIP and PCache messages is shown in Table 1.

A PCache dissemination message is mapped into a SIP REGISTER message. The MESSAGE method [1], used for SIP-based page-mode instant messaging, is used for queries, both for retrieving a user's IP address based on his AOR, and for a profile based query. In the first case, the "Subject" header field of the

MESSAGE request will carry the indication "Query"; the AOR of the target user is carried in the SIP Request-URI and in the To header field. In the second case, the indication "Query-all" will be reported; the search parameters are specified in the body of the message. Replies to query are also carried in the body of a MESSAGE request.

3 Evaluation

The performance of SIPCache was evaluated using a prototype implemented in the *ns-2* network simulator v. 2.28. All tests are performed in a simulated ad hoc network composed by 100 nodes. Nodes move accordingly to the random waypoint model [2]. The simulated region is a rectangle with 1500mx500m. Two classes of speeds have been defined. In "Speed 0" tests, nodes are randomly deployed in the simulated space and do not move for the entire duration of the simulation. In "Speed 10" tests, nodes move at a speed randomly selected between 5 and 15m/s. The pause time has a uniform distribution between 0 and 40s.

SIPCache messages respect SIP specifications, when the transport protocol is UDP, and do not exceed 1300 bytes in size. Each record is assumed to have 65 bytes. A SIPCache message can carry at most 15 records. The transmission range is varied to simulate networks with different densities since the number of nodes in the neighborhood varies with the transmission range.

Finally, we have evaluated the efficiency of SIPCache with devices with different resources. In "Cache 10" tests, the cache of each node carries at most 10 records. In "Cache 100" tests, it is assumed that the devices can carry the 100 records advertised in the simulations.

Each node advertises one record. During the 900s of the simulation time, 300 query operations are performed. Targeted dSIP queries are performed by randomly selecting one of the registrations. Profile based queries always search for 10% of the records previously advertised. All results presented in this section average 10 independent runs of the simulator for each condition.

Table 1: Mapping SIP-PCache messages

	SIP	PCache
Message Type	REGISTER	Dissemination
	MESSAGE	Query
	200 OK	Query Reply
Header name	Max-Forwards	TTL
	From	Source
	Call-ID	Serial Number
	CSeq	Time From Storage
	Via	Route Stack
	Content-Length	# Items
	Body	Items

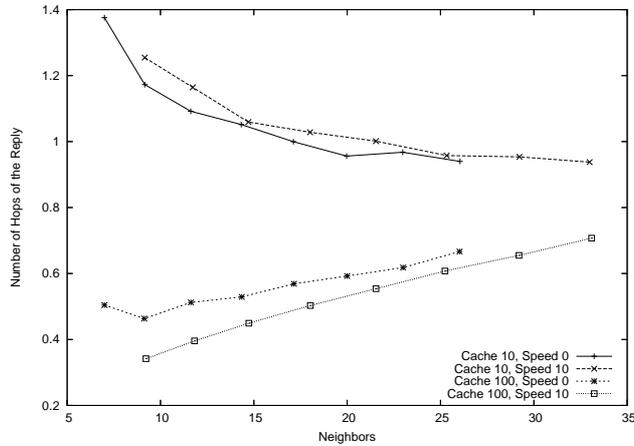


Figure 1: Evaluation of targeted queries

3.1 Targeted queries

Fig. 1 presents the average number of hops traveled by the first reply received by the source of a query. As the number of neighbors increases, so does the probability of finding a reply closer to the source of the query. For an average of 17 1-hop neighbor nodes, queries are replied on average by one of the 1-hop neighbors.

An interesting behavior of SIPCache is observed in the results for the “Cache 100” tests where, contrary to the expectations, the average distance of the replies increases with the number of neighbors. This effect is better explained when we consider that SIPCache is aimed at situations where the capacity of the cache of the nodes is lower than the records advertised. To attend to these cases, nodes cooperate in the distribution of the replicas by refraining from storing some records if they learn that a neighbor has already stored it. The increase in the number of neighbors makes each node to become aware of the state of the cache of a higher number of nodes, thus reducing the probability of storing a replica in the cache. However, it should be noted that there is a negligible number of nodes that do not have some record in their 1 hop neighborhood.

3.2 Profile based queries

The performance of SIPCache in profile based queries is evaluated by the proportion of the records that were available in the node’s cache or replied to the node in the first 10s after the broadcast of the query.

As Fig. 2 shows, even in adverse conditions, SIPCache is able to retrieve on average approximately 80% of the records satisfying the query. When the number of neighbors is on average 14 (corresponding to a transmission range of 225m), results are already above 95%. It should be noted that these results are obtained with the propagation of the query to at most 2 hops away from the source of the query.

3.3 Comparative analysis

To evaluate the adequacy of SIPCache, we defined a simpler implementation of dSIP, which dismisses the registration/dissemination phase. Because data are not replicated as a consequence of the dissemination, all queries are broadcasted to the network and only the nodes which have a registration satisfying the query send a reply. The broadcast operation required in this implementation is comparable with the registration operation defined in SIPCache if the overhead of sending a reply and the need to increase the probability of delivery of the messages to every node is neglected. Fig. 3 compares this naïve implementation with SIPCache for “Speed 0” tests. The y axis shows the gain (in number of messages per operation) of using SIPCache instead of the naïve implementation for the number of queries represented in the x axis.

The results for this naïve implementation were estimated from multiplying the average number of messages used in the registration operations of SIPCache by the number of queries. The cost of delivering replies was considered negligible as in this implementation it would consist in a point-to-point delivery for each node responding to the query. For the estimation in SIPCache, we have accounted with the original cost of the registration and with the average number of messages required for propagating the query and replies.

Results show that in the majority of the cases, SIPCache begins to save device resources after a small number of queries. We emphasize that in a real deployment, the number of queries should exceed by large the number of registered users, amplifying the expected savings of device’ resources. Exceptions are profile based queries in “Cache 10” tests with a large density. This negative value can be attenuated by the gains presented by targeted queries in the same situation.

4 Conclusions

This paper has presented SIPCache, a distributed location service for the SIP protocol, for use in multi-hop ad-hoc networks. By enabling the SIP location service in ad-hoc networks, it is possible for a SIP end user to discover the contact information of other users, and then initiate a SIP session with them. Normally, the SIP location service is provided by a centralized entity, the SIP registrar server; in ad-hoc networks a decentralized approach must be followed as servers are not available.

Our proposal builds on top of two independent solutions. The first, dSIP [3], is an enhancement to the software architecture of a SIP client to allow the usage of SIP in ad-hoc networks, both in terms of location service and initiating and terminating SIP sessions. The second part of SIPCache is based on PCache [4], an algorithm designed to efficiently distributing and retrieving cached data items in an ad-hoc network. When applying PCache to dSIP, we distribute the contact information, or binding, of SIP users in the network, so that any node is able to retrieve information of any other user. By merging the functionalities of dSIP with PCache, we extend the scope of the pure dSIP, which is optimized for small ad-hoc networks.

Standard SIP messages are exchanged in the network obeying to the policies dictated by the PCache algorithm. The end result is that the bindings of SIP

users are spread all over the network, instead of being maintained by a single centralized entity. This decentralizes the SIP location service given that PCache provides fully decentralized means for querying and retrieving the bindings. Simulation results have proved that by using SIPCache it is possible to build an efficient and functional distributed location service. The service is efficient, since it uniformly replicates the bindings in the network, so that any request for any user's binding, issued by any random node in the network, is replied within a reasonably low number of hops, possibly already by neighbors of the querying node.

References

- [1] B. Campbell et al. Session initiation protocol (sip) extension for instant messaging. Request for Comments (Standards Track) 3428, IETF, December 2002.
- [2] D. B. Johnson and D. A. Maltz. *Mobile Computing*, chapter Dynamic Source Routing in Ad Hoc Wireless Networks, pages 153–181. Kluwer Academic Publishers, 1996.
- [3] S. Leggio, J. Manner, A. Hulkkonen, and K. Raatikainen. Session initiation protocol deployment in ad-hoc networks: a decentralized approach. In *2nd International Workshop on Wireless Ad-hoc Networks (IWWAN)*, London, May, 2005.
- [4] H. Miranda, S. Leggio, L. Rodrigues, and K. Raatikainen. A stateless neighbour-aware cooperative caching protocol for ad-hoc networks. DI/FCUL TR 05–23, Department of Informatics, University of Lisbon, December 2005. Also as Technical Report Number C–2005–76. Computer Science Department, University of Helsinki.
- [5] J. Rosenberg et al. SIP: Session initiation protocol. RFC 3261 (Standards Track), IETF, June 2002.

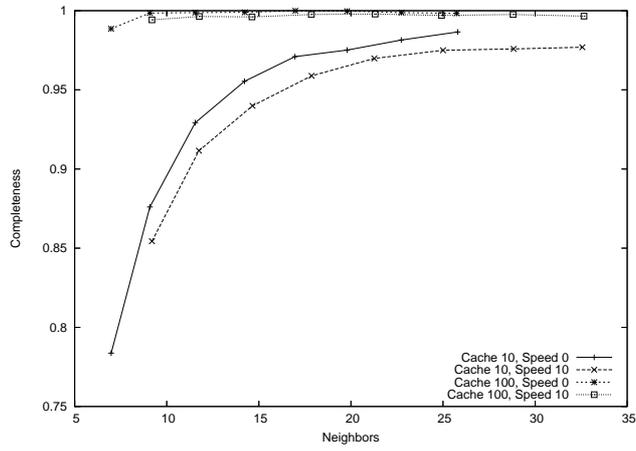


Figure 2: Coverage of Profile Based Queries

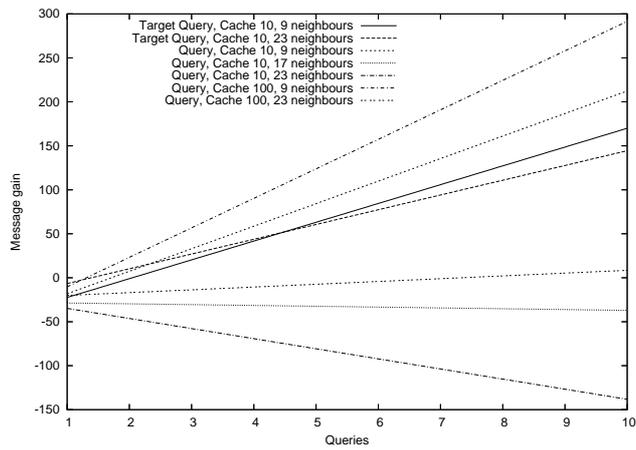


Figure 3: SIPCache vs naïve implementation