Collaborative Backup for Dependable Mobile Applications

[Extended Abstract]

Marc-Olivier Killijian,
David Powell
LAAS-CNRS
7 Avenue du Colonel Roche
31077 Toulouse cedex 4
France

Michel Banâtre,
Paul Couderc
IRISA
Campus Universitaire de Beaulieu
35042 Rennes cedex, France

Yves Roudier
Institut Eurécom
2229 Route des Crêtes
Sophia Antipolis
06560 Valbonne
France

ABSTRACT

We describe the work we are conducting on new middleware services for dependable and secure mobile systems. This work is based on approaches à la peer-to-peer in order to circumvent the problems introduced by the lack of infrastructure in self-organizing networks of mobile nodes, such as MANETs. The mechanisms we propose are based on collaboration between peer mobile devices to provide middleware services such as trust management and critical data storage. This short paper gives a brief description of the problems we are trying to solve and some hints and ideas towards a solution.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems

Keywords

Mobile applications, data back-up, collaboration

1. INTRODUCTION

The MoSAIC (Mobile System Availability, Integrity and Confidentiality) project [9] aims to investigate novel dependability and security mechanisms for mobile wireless devices, especially personal mobile devices, in ambient intelligence applications. The mobile devices of interest include, for instance: personal digital assistants (PDAs), laptop computers, mobile telephones, digital cameras, etc., and extend to systems embedded within vehicles. The focus is on sparse ephemeral self-organizing networks, using predominately single-hop wireless communication, i.e., networks of a small number of a potentially large population of mobile devices that come into existence spontaneously by virtue of physical proximity and mutual discovery, and that cease to exist as soon as communication is no longer possible.

Most of the data carried on a PDA is a copy of data that is mainly produced and also stored elsewhere. For example, a PDA contact database is regularly synchronized with a desktop computer application. This reduces the impact of failure of such devices to the data that is produced directly on the device between synchronizations. However, in the case of capture devices (devices capable of acquiring data such as pictures, sound or video), large quantities of data are generated directly on the mobile device, leading to a much larger quantity of data that remains sensitive to device failure until a backup copy can be created. This highlights the need for new ways of ensuring data availability. Because the density of these devices is increasing (as mobile devices are becoming more and more popular), there is an opportunity for cooperatively backing up data by using neighborhood devices. The first objective of our work is therefore to define an automatic data back-up and recovery service based on mutual cooperation between mobile devices with no prior trust relationships. Such a service aims to ensure continuous availability of critical data managed by mobile devices that are particularly prone to energy depletion, physical damage, loss or theft. The basic idea is to allow a mobile device to exploit accessible peer devices to manage backups of its critical data. To our knowledge, no work has already exploited this principle of cooperative backup for mobile devices. Indeed, relatively little work appears to have been devoted to tolerance of device failures in a mobile self-organized network scenario [14] [2] [1], although there has been considerable work on checkpointing in cellular mobile computing environments (see, e.g., [18] [19] [20] [4] [16] [17]).

The implementation of such a service by cooperation between mobile nodes with no prior trust relationship is far from trivial since new threats are introduced: (a) selfish devices may refuse to cooperate; (b) backup repository devices may themselves fail or attack the confidentiality or integrity of the backup data; (c) rogue devices may seek to deny service to peer devices by flooding them with fake backup requests; etc. We intend to study trust management mechanisms to support cooperative services between mutually suspicious devices. Of particular interest are mechanisms based on reputation (for prior confidence-rating and posterior accountability) and rewards (for cooperation inci-

2nd Workshop on Middleware for Pervasive and Ad-Hoc Computing
Toronto, Canada

tation). In the sparse ephemeral networks considered, these mechanisms can rely neither on accessibility to trusted third parties nor on connectivity of a majority of the considered population of devices [22]. Self-carried reputation and rewards are therefore of prime interest. This approach contrasts to most existing approaches to mobile system security, which have mainly focused on key management and distribution (see, e.g.,[22] [8] [12]) and on secure ad-hoc network routing (see, e.g., [3] [6] [15] [21]).

Achieving dependability and security despite accidental and malicious faults in networks of mobile devices is particularly challenging due to their intrinsic asynchrony (unreliable communication, partitioning, mobility, etc.) and the consequent absence of continuous connectivity to global resources such as certification and authorization servers, system wide stable storage, a global time reference, etc. Furthermore, the threats to dependability and security are particularly severe: device lifetime and communication are severely limited by scarcity of electrical energy; use of wireless links means susceptibility to link attacks ranging from passive eavesdropping to active impersonation, message replay, and message distortion; poor physical protection of mobile devices (especially in a hostile environment) makes them susceptible to physical damage, and vulnerable to theft or subversion.

There are thus two related issues that need to be addressed :

1. Fault- and intrusion-tolerant collaborative data backup (with possible extension to checkpointing).

2. Self-carried reputation and rewards for collaboration between sporadically interconnected and mutually suspicious peer devices without reliance on a fixed infrastructure and access to trusted third parties.

Common to both is our emphasis on spontaneous interaction between peer mobile devices with no prior trust relationships. In this paper, we focus on the first of these two issues.

2. FAULT TOLERANCE BY COLLABORATIVE BACKUP

We are investigating middleware services to support the dependability and security of mobile ambient intelligence applications. We consider highly dynamic systems consisting of wireless-equipped mobile devices that communicate with each other mostly by direct, single-hop communication. However, we do not preclude extensions to include indirect communication via a multi-hop ad-hoc network or occasional access to a fixed communication infrastructure. We are not addressing mobile ad-hoc routing protocols, or dependability and security issues at the wireless network level, which are largely covered in the litterature.

We consider the design and implementation of a prototype service for data backup and recovery by cooperation between ephemerally-connected and mutually-suspicious mobile devices. The problems we consider arise from the specific characteristics of ambient intelligence applications based predominately on sparse ephemeral networks of mobile devices: disconnected mode or absence of fixed infrastructure, absence of prior organization, ephemeral interactions, user transparency, and user privacy. Limits on mobile device energy, computation and storage will also constrain the technical solutions that can be considered.

A typical scenario for such a service might be a researcher travelling to a conference, using her PDA to take important notes, and using the PDAs of fellow attendees or travellers to host temporary back-ups of critical data. Such temporary back-ups provide the means for recoverying critical data in the event that her PDA fail, break or be stolen. Recovery can be achieved by purchasing a new PDA, authenticating it and then recollecting the critical data chunks backed-up on other devices. An important aspect of this scenario concerns the fact that the users (and their devices) cooperating for achieving this back-up service have no prior trust relationship. They must thus protect, for example (a) the backed-up data against confidentiality and availability attacks and (b) the back-up devices against denial of service attacks.

Other scenarios, with varying prior trust models, can be imagined in military applications (e.g., recovery and redistribution of critical command and control data during battlefield operations), civilian emergency operations, home automation and entertainment, etc.

The need for such a fault-tolerance service is motivated by: (a) the increasing dependency of users on the availability, integrity and confidentiality of data carried by mobile devices and (b) the fragility of mobile devices and other risks relating to their use in a harsh or even hostile environment. We purposely limit ourselves to the issue of data backup, but note that such a service could serve as the basis for mobile device checkpointing and recovery, and for real-time tolerance of mobile device failure based on redundant devices.

The problems to be addressed include: resource allocation, garbage collection of obsolete backups, integrity and confidentially of backup data, resistance to denial-of-service (DoS) attacks, etc. The service is to be supported by negotiation between peer mobile devices with no prior trust relationship. Among the various approaches that might be considered, we intend to take inspiration from current work in the area of peer-to-peer (P2P) applications [13] [5] [10] [11], which have characteristics that are particularly well-adapted to the considered environment: absence of pre-established organization, service through cooperation, short-duration interactions, etc. We also plan to take inspiration from our know-how in the domain of fragmentation-replication-dissemination (FRD) techniques, which exploit distribution to increase availability, integrity and confidentiality in the face of accidental faults and malicious attacks [7]. Until now, these FRD techniques have only been considered in the case of fixed infrastructure systems. We might also consider the advantages that could be drawn from occasional access to a common time reference (e.g., through the Global Positioning System (GPS)) or from exploiting mobility for data dissemination.

In the sequel, we use the terms *data owner* to refer to a device requesting its data to be backed up and *data saver* for a device hosting back-up data. Any device may be both a data owner and a data saver. However, to simplify our discourse, we usually consider a single data owner.

2.1 Threats

The data back-up service must face up to the following threats:

1. Permanent and transient accidental faults affecting a

data owner.

2. Theft or loss of a data owner device.

3. Accidental or malicious faults causing a data saver to be unavailable when recovery is required (i.e., on failure of the data owner).

4. Accidental or malicious modification of data backups that could violate data integrity if recovery should be required.

5. Malicious read access to data backups. Back-ups may contain sensitive confidential data that should be made unintelligible to the user of the data saver device.

6. Denial of service through selfishness. Cooperation may be thwarted if there is no incentive for devices to participate.

7. Denial of service through maliciousness. A malicious data owner could attempt to saturate data savers by false back-up requests, and thereby deny service to other data owners and to users of the attacked data saver devices. A malicious data saver may also choose to withhold backed-up data (cf. threat 3).

It will also be important to distinguish various contexts of utilization of the data back-up service according to the type of user community and appropriate prior trust model. For example, in a closed (and non-infiltrated) military context, certain threats such as denial-of-service through selfishness or malicious attack may be considered negligible.

2.2 Back-up process

The primary aim of the back-up service is to provide protection against permanent and transient accidental faults of data owners (threat 1). Depending on the utilization context, complete or partial back-up of data may be considered. Partial *delta* back-ups or update operation logs might be preferred to minimize the amount of data to be transferred to and stored on data savers, or even to provide some protection against confidentiality attacks on back-ups (threat 5).

The back-up service also provides protection of data availability in the face of loss or theft of the data owner device (threat 2). Confidentiality might be provided in such a situation by an "auto-delete" function triggered by a failed user-authentication challenge.

Unavailability and modification of back-ups (threats 3 and 4) are only of importance if the data owner should fail. Tolerance of multiple faults may be achieved by installing redundant back-ups on independent data savers. Malicious read access to back-ups (threat 5) may be prevented by cryptographic techniques, with appropriate trade-offs between the level of protection provided and the associated costs in energy and resource consumption. The strength (key length, degree of redundancy, etc.) and cost of the deployed techniques may be adapted according to the degree to which data savers may be trusted (e.g., devices of colleagues or those of strangers). The adaptation could also make use of a dynamic measure of the "reputation" of the data saver (cf. issue 2 raised in the introduction).

Fragmentation–replication–dissemination (FRD) techniques [7] are also of interest here. Data confidentiality may be provided by cutting back-up data into fragments that are disseminated over different data savers. Fragments may also be replicated to ensure data availability and integrity (by voting on multiple replicas). Fragmentation, replication and dissemination may be modulated in both space and time according to the number of trustable devices available in a given place or at a given instant.

Denial of service through selfishness (threat 6) may be discouraged by the use of a "reward" scheme to motivate device participation, inspired from micro-economy approaches developed in peer-to-peer applications. Devices acting as data savers are rewarded for their participation and may redeem their earnings when acting as data owners that wish to purchase back-up service. Denial of service through maliciousness (threat 7) may also be discouraged by an appropriate "reputation" mechanism. Devices with a history of detected maliciousness will have a poor reputation and will be spurned by data owners when negotiating to purchase back-up service. The related notions of reward and reputation are the subject of the cooperative service trust mechanisms that we also plan to investigate.

2.3 Recovery process

The second important aspect of the proposed data back-up service concerns the means by which back-up data may be re-installed when required on data owners, i.e., data recovery. This involves finding the data that has been backed up and transferring it back to the data owner or its surrogate.

The recovery process will depend heavily on whether or not devices can occasionally connect to a fixed infrastructure. If access to a fixed infrastructure cannot be considered (e.g., in a battlefield scenario), then access to back-up data has to be based on establishing a wireless communication channel between data owner and saver devices. If direct communication is not possible (which will be the usual case) then the solution may be to create an ad-hoc network with intermediate devices, or to wait until the devices are again within wireless range (by chance encounter or by planned rendezvous).

At least two recovery modes can be distinguished:

- *Push* recovery: the data saver automatically sends data backups to the data owner or its surrogate. The most appropriate way might be for data savers to trigger such a boomerang operation as soon as they have access to a fixed infrastructure. The data could be transferred either immediately to the data owner or its surrogate, or possibly through a trusted third party.

- *Pull* recovery: the data owner searches for the data copies that it requires. Again, we may take inspiration from P2P systems that seek to develop totally distributed file search engines. Requests to the search engine might target the requested data by specifying particular places or times, e.g., "the data I backed up during the flight from Toulouse to Rennes on January 10, 2004".

When partial back-ups have been created, like when fragmentation-replication-dissemination is used, the recovery process will also need to tackle the problem of reconstructing the complete data from the various parts.

Many various optimizations of the proposed back-up service may be considered. For example, in the case of incremental back-ups, the optimal period of back-up creation may depend on several factors, including the relative size of the increments (deltas or update logs) and the performance of recovery based on those increments. The chosen solutions need to be flexible and adaptable to various application scenarios. Another important issue is that of garbage-collecting obsolete back-up data. This may depend on the notion of a contract set up between data owners and savers, or be triggered when the data owner announces that the earlier back-ups are obsolete. The appropriate solutions imply various business models associated with micro-economy mechanisms of various complexity: fines, contracts, leases, etc.

3. CONCLUSIONS

While the area of dependability of the low level network layers for mobile devices has received much attention (e.g. fault-tolerant routing), middleware and application-level dependability mechanisms remain almost unexplored. As mobile devices become more and more common - we can now embed a real-time operating system with wireless capabilities in a wrist-watch - users will increasingly use them for more critical tasks and will expect greater reliability from them. For example, loosing the automatically gathered orders of the clients that a salesman visited during the morning is completely unacceptable. Even if most of the data carried on a PDA is typically regularly synchronized with a desktop computer, some of its data is produced or modified between these synchronizations. In the case of capture devices, this amount of data is even larger. The user cannot afford to lose the critical data created or modified between synchronizations. The mechanisms we describe in this paper try to tackle the issue of using peer-provided resources for building a collaborative backup service between mobile devices with no prior trust relationship. We think that the impact of such a technology will be high and can be extended to other scenarios like exploratory operations, sensor networks and military missions.

4. REFERENCES

[1] M. Boulkenafed and V. Issarny. AdHocFS: Sharing Files in WLANs. In *2nd Int. Symp. on Network Computing and Applications*, pages 156–63. IEEE CS Press, 2003.

[2] M. Boulkenafed and V. Issarny. A middleware service for mobile ad hoc data sharing, enhancing data availability. In *4th ACM/IFIP/USENIX International Middleware Conference*, pages 493–511. Springer, 2003.

[3] S. Buchegger and J.-Y. L. Boudec. The selfish node: Increasing routing security in mobile ad hoc networks. Technical Report RR 3354, IBM, May 2001.

[4] G. Cao and M. Singhal. Mutable checkpoints: a new checkpointing approach for mobile computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 12:157–72, 2001.

[5] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong. Freenet: A distributed anonymous information storage and retrieval system. *Lecture Notes in Computer Science*, 2009:46, 2001. `http://freenet.sourceforge.net`.

[6] B. Dahill, B. Levine, E. Royer, and C. Shields. A secure routing protocol for ad hoc networks. In *10th Conference on Network Protocols (ICNP)*, November 2002.

[7] Y. Deswarte, L. Blain, and J.-C. Fabre. Intrusion tolerance in distributed systems. In *IEEE Symposium on Security and Privacy*, pages 110–121. IEEE CS Press, 1991.

[8] A. Khalili, J. Katz, and W. A. Arbaugh. Toward secure key distribution in truly ad-hoc networks. In *Symp. on Applications and the Internet Workshops (SAINT'03 Workshops)*, pages 342–46, 2003.

[9] M.-O. Killijian, M. Banâtre, P. Couderc, L. Courtès, S. Crosta, R. Molva, D. Powell, Y. Roudier, and F. Weiss. The MoSAIC project. `http://www.laas.fr/mosaic/`.

[10] D. Kügler. An analysis of gnunet and the implications for anonymous, censorship-resistant networks. `http://www.ovmj.org/GNUnet/`.

[11] S. Lee, R. Sherwood, and B. Bhattacharjee. Cooperative peer groups in NICE. In *INFOCOM'03*, April 2003.

[12] D. Liu, P. Ning, and K. Sun. Efficient self-healing group key distribution with revocation capability. In *10th ACM Conf. on Computer and Communications Security (CCS'03)*, pages 231–40, 2003.

[13] MNET. The MNET project. `http://mnetproject.org`.

[14] P. Nikander. Fault tolerance in decentralized and loosely coupled systems. In *Ericsson Conference on Software Engineering*. Ericsson, 2000.

[15] P. Papadimitratos and Z. J. Haas. Secure routing for mobile ad hoc networks. In *SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2002)*, 2002.

[16] T. Park, N. Woo, and H. Y. Yeom. An efficient recovery scheme for mobile computing environments. In *Int. Conf. on Parallel And Distributed Systems (ICPADS)*, pages 53–60. IEEE CS Press, 2001.

[17] C. Pedregal-Martin and K. Ramamrithan. Support for recovery in mobile systems. *IEEE Transactions of Computers*, 51:1219–24, 2002.

[18] D. K. Pradhan, P. Krishna, and N. H. Vaidya. Recoverable mobile environment: Design and trade-off analysis. In *26th IEEE Int. Symp. on Fault-Tolerant Computing (FTCS-26)*, pages 16–25. IEEE CS Press, 1996.

[19] R. Prakash and M. Singhal. Low-cost checkpointing and failure recovery in mobile computing systems. *IEEE Transactions on Parallel and Distributed Systems*, 7:1035–48, 1996.

[20] B. Yao, K.-F. Ssu, and W. K. Fuchs. Message logging in mobile computing. In *29th IEEE Int. Symp. on Fault-Tolerant Computing (FTCS-29)*, pages 294–301. IEEE CS Press, 1999.

[21] M. Zapata and N. Asokan. Securing ad hoc routing protocols. In *ACM Workshop on Wireless Security (WiSe 2002)*, September 2002.

[22] L. Zhou and Z. Haas. Securing ad hoc networks. *IEEE Network Magazine*, 13:24–30, 1999.