Collaborative Backup for Dependable Mobile Applications

M.-O. Killijian[1], M. Banâtre[2], P. Couderc[2], D. Powell[1], Y. Roudier[3]

[1]LAAS-CNRS, 7 avenue du Colonel Roche, 31077 Toulouse Cedex 4

[2]IRISA, Campus Universitaire de Beaulieu, 35042 Rennes Cedex

[3]Eurécom, 2229 Route des Crêtes, Sophia Antipolis, 06560 Valbonne

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 article gives a brief description of the problems we are trying to solve and some hints and ideas towards a solution.

1 Introduction and Problem Statement

The MoSAIC (Mobile System Availability, Integrity and Confidentiality) project [1] 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[1], 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

1- Capture devices are capable of acquiring multimedia data, such as pictures, sound or video. Such devices are an important evolution, since mobile devices are thus becoming data sources instead of being mostly "reader" devices. There is also an increasing trend towards devices that combine the functions of PDAs and mobile phones with capture devices.

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 [2] [3] [4], although there has been considerable work on checkpointing in cellular mobile computing environments (see, e.g., [5] [6] [7] [8] [9] [10]).

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 incitation). 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 [11]. 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., [11] [12] [13]) and on secure ad-hoc network routing (see, e.g., [14] [15] [16] [17]).

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 Our Vision

Consider the following scenario:

Alice is attending an important symposium. Before leaving her office, she takes care to synchronize her wireless personal digital assistant (PDA) with her desktop computer. During this synchronization, all her critical personal and professional data (address book, calendar, notes, presentations, etc.) on her PDA are safely backed up.

During the journey to the symposium, Alice continues to work on her presentation using her PDA. When she reaches her destination, Alice registers and loads the latest version of the symposium program onto

her PDA. She selects the sessions she wants to attend and goes from one room to another, chatting with colleagues during the breaks and occasionally taking notes on her PDA. As she moves around, her PDA wirelessly interacts with those of the colleagues she meets or that pass nearby, automatically looking for data that they wish to make available concerning, for example, their centers of interest and their latest papers, presentations and reference citations, etc. Each physical encounter is seen as an opportunity for Alice's PDA to gather data pertinent to her research. At the end of the symposium, Alice attends a meeting of the steering committee to discuss and decide on how to organize the next event. During the meeting, the minutes are written in common by a computer-supported collaborative work application, running on the steering committee members' PDAs.

On her way back home, Alice uses her PDA in the taxi to the airport to start writing a paper on the brilliant new idea she had while listening to one of the presentations at the symposium. Deep in thought while emptying her pockets at the airport security check, she carelessly forgets her PDA on the luggage belt. She is jolted back to reality by the metallic thud of her PDA bouncing of the tiled airport floor, shortly followed by the crunching sound of her PDA being shattered under the heavy boots of Hiro, the Sumo wrestler who happened to be following her. With dismay, she realized that she had effectively lost all the work she had done since leaving home. How she wished she had done her work with old-fashioned paper-and-pen.

Luckily however, Alice's PDA had taken advantage of its chance encounters with other devices to automatically back-up the latest revisions to its data. So, after purchasing a new PDA in the airport gizmo shop, Alice was able to reload at least some of her recent data from devices of people she had happened to encounter earlier in the day. Some of them could be contacted directly in the airport lounge, others could be found by searching throughout the airport via an ad-hoc network of PDAs and other wireless devices. When she reached her office the following day, she was glad to realize that the remaining data had been automatically sent back to her office machine, which meant that the next synchronization brought her new PDA almost up-to-date with the state her previous one had before the unfortunate incident at the security check.

Of course, since conference attendees and chance traveling companions move around frequently, the duration of each backup opportunity was unknown and relatively brief, which meant that only small amounts of data could be backed up each time. Moreover, many encounters were with people that Alice had never met before, including some shady persons of doubtful character, so care had to be taken that Alice's PDA did not reveal private or other sensitive data, nor that it was corrupted by data sent with malicious intent, nor that her data was backed-up only on rogue devices. Also, Alice was glad that her PDA had altruistically offered to share its resources with the devices encountered so that they too had been generous in return.

This admittedly somewhat-contrived scenario illustrates the need for the proposed work: to use cooperation between peer wireless mobile devices, with no prior trust relationship, in order to ensure device data availability, while providing guarantees of integrity and confidentiality (including privacy). 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.

3- Our Approach

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.

3.1 Characteristics of considered ambient intelligence environment

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**: traditional dependability and security functions depend on service offered by dedicated entities, as is the case for networking functions dedicated to routing. A typical security example is that of a public key certification service, which is usually a prerequisite for many basic security services such as authentication, key exchange, non-repudiation, etc. An example from the dependability viewpoint is that of a system-wide stable storage facility on which data and process checkpoints might be stored. Here, due to the highly dynamic nature of the considered networks and the predominant mode of operation without access to a fixed infrastructure, we are adopting a peer-to-peer approach without any such reliance on dedicated entities.
- **Absence of prior organization**: ambient intelligence applications rely on flexible communication and openness to facilitate interactions between devices with no pre-established organizational relationships. Classic security mechanisms are inapplicable in this context since they attempt to reproduce at a logical level the relationships that exist at the organizational level. In particular, entity authentication consists of demonstrating the link between an operational entity and an identity or role within an organization. The absence of such links means that classic security mechanisms cannot be used for ambient intelligence applications and that new mechanisms must be invented that allow the dynamic construction of trust relationships between entities that do not share a common organizational reference.
- **Ephemeral interactions**: wireless communication and mobility mean that entities can only interact for brief periods. Consequently, the implementation of dependability and security protocols based on durable interactions, using notions of state or session, become problematic.
- **User transparency**: one of the objectives of ambient intelligence applications is to make the underlying computer systems as transparent as possible to the user, to the extent that they become invisible. To reach this objective, availability, integrity and confidentiality need to be guaranteed by built-in dependability and security mechanisms, with minimal user interaction**.**
- **Privacy**: the danger that ambient intelligence applications might reveal data concerning user identity, location, behavior or other personal characteristics is of growing importance since society is becoming increasingly concerned by issues of personal rights and privacy. The protection of private data, which is far from being ensured in classic distributed applications, is exacerbated in ambient intelligent applications by virtue of the required transparency and their dependency on physical location.
- **Energy, computation and storage constraints**: most devices in an ambient intelligence environment are severely restricted by the autonomy of their batteries. The need to economize electrical energy means that wireless communication must be kept to a minimum. Furthermore, energy, space and cost limitations all construe to limit the device computation power and storage capacity. Consequently, protocols and mechanisms (e.g., cryptography) must be designed under severe energy, computation and storage constraints.

3.2 Fault tolerance by cooperative backup

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 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 [18] [19], 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 [20]. 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 "client" to refer to a device requesting its data to be backed up and "server" for a device hosting back-up data. Any device may be both a back-up client and a back-up server. However, to simplify our discourse, we usually consider a single client.

3.2.1 Threats

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

1. Permanent and transient accidental faults affecting a client device.
2. Theft or loss of a client device.
3. Accidental or malicious faults affecting server device availability should recovery be required (i.e., on failure of the client).
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 server device user.
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 client could attempt to saturate servers by false back-up requests, and thereby deny service to other clients and to users of the attacked server devices. A malicious server 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.

3.2.2 Back-up

The primary aim of the back-up service is to provide protection against permanent and transient accidental faults of client devices (threat 1). Depending on the utilization context, complete or partial back-up of client device 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 server devices, 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 client 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 import if the client device should fail. Tolerance of multiple faults may be achieved by installing redundant back-ups on independent server devices. 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 server devices 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 server devices (see theme 2 below).

Fragmentation-replication-dissemination (FRD) techniques [20] are also of interest here. Data confidentiality may be provided by cutting back-up data into fragments that are disseminated over different server devices. 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 (P2P) applications. Devices acting as servers are rewarded for their participation and may redeem their earnings when acting as clients 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 client devices 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 (see section 1).

3.2.3 Recovery

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 client devices, i.e., data recovery. This involves finding the data that has been backed up and transferring it back to the client device 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 client and server 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 server devices automatically send data backups to the client device or its surrogate. The most appropriate way might be for server devices 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 client device or its surrogate, or possibly through a trusted third party.

- "Pull" recovery: the client device 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 client and server devices, or be triggered when the client device 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.

4 Conclusion

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 scenarii like exploratory operations, sensor networks and military missions.

5 References

[1] MOSAIC, http://www.laas.fr/~mkilliji/MOSAIC/

[2] P. Nikander. "Fault Tolerance in Decentralized and Loosely Coupled Systems", In *Ericsson Conference on Software Engineering*, (Stockholm, Sweden), Ericsson, 2000.

[3] M. Boulkenafed, V. Issarny. "A Middleware Service for Mobile Ad Hoc Data Sharing, Enhancing Data Availability", In *4th ACM/IFIP/USENIX International Middleware Conference*, (Rio de Janeiro, Brazil), LNCS, 2672, pp.493-511, Springer, 2003.

[4] M. Boulkenafed, V. Issarny. "AdHocFS: Sharing Files in WLANs", In *2nd Int. Symp. on Network Computing and Applications*, (Cambridge, MA, USA), pp.156-63, IEEE CS Press, 2003.

[5] D. K. Pradhan, P. Krishna, N. H. Vaidya. "Recoverable Mobile Environment: Design and Trade-off Analysis", In *26th IEEE Int. Symp. on Fault-Tolerant Computing (FTCS-26)*, (Sendai, Japan), pp.16-25, IEEE CS Press, 1996.

[6] R. Prakash, M. Singhal. "Low-Cost Checkpointing and Failure Recovery in Mobile Computing Systems", *IEEE Transactions on Parallel and Distributed Systems*, 7 (10), pp.1035-48, October 1996.

[7] B. Yao, K.-F. Ssu, W. K. Fuchs. "Message Logging in Mobile Computing", In *29th IEEE Int. Symp. on Fault-Tolerant Computing (FTCS-29)*, (Madison, WI, USA), pp.294-301, IEEE CS Press, 1999.

[8] G. Cao, M. Singhal. "Mutable Checkpoints: a New Checkpointing Approach for Mobile Computing Systems", *IEEE Transactions on Parallel and Distributed Systems*, 12 (2), pp.157-72, February 2001.

[9] T. Park, N. Woo, H. Y. Yeom. "An Efficient Recovery Scheme for Mobile Computing Environments", In *Int. Conf. on Parallel And Distributed Systems (ICPADS)*, (KyongJu City, Korea), pp.53-60, IEEE CS Press, 2001.

[10] C. Pedregal-Martin, K. Ramamrithan. "Support for Recovery in Mobile Systems", *IEEE Transactions of Computers*, 51 (10), pp.1219-24, October 2002.

[11] L. Zhou, Z.J. Haas. "Securing Ad Hoc Networks". *IEEE Network Magazine*, 13(6): 24-30, November/December 1999.

[12] A. Khalili, J. Katz, W. A. Arbaugh. "Toward Secure Key Distribution in Truly Ad-Hoc Networks", In *Symp. on Applications and the Internet Workshops (SAINT'03 Workshops)*, pp.342-46, 2003.

[13] D. Liu, P. Ning, K. Sun. "Efficient Self-Healing Group Key Distribution with Revocation Capability", In *10th ACM Conf. on Computer and Communications Security (CCS'03)*, (Washington D.C., USA), pp.231-40, 2003.

[14] Sonja Buchegger, Jean-Yves Le Boudec. "The Selfish Node: Increasing Routing Security in Mobile Ad Hoc Networks". IBM Research Report RR 3354, May 2001

[15] B. Dahill, B.N. Levine, E. Royer, C. Shields. "A Secure Routing Protocol for Ad Hoc Networks". In *10th Conference on Network Protocols (ICNP)*, November 2002.

[16] P. Papadimitratos, Z. J. Haas. "Secure Routing for Mobile Ad hoc Networks", In *SCS Communication Networks and Distributed Systems Modeling and Simulation Conference (CNDS 2002)*, (San Antonio, TX, USA), 2002.

[17] M.G. Zapata, N. Asokan. "Securing Ad Hoc Routing Protocols". In *ACM Workshop on Wireless Security (WiSe 2002)*, September 2002.

[18] Mnet, http://mnet.sourceforge.net.

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

[20] Y. Deswarte, L. Blain, J.-C. Fabre. "Intrusion Tolerance in Distributed Systems". In IEEE Symposium on Security and Privacy, (Oakland, CA, USA), pp. 110-121, IEEE CS Press, 1991