

MoSAIC

Mobile System Availability Integrity & Confidentiality



- 3 years
- 3 partners
- Funded by French Ministry of Research
- Started September 2004




ACI Sécurité
& Informatique




Cooperative Backup for Nomadic Devices

- Aim: investigate new distributed algorithms and mechanisms for the tolerance of:
 - Accidental faults
 - Malicious faults
- Nomadic device scenario
 - Mostly disconnected operations
 - Opportunistic wireless communication with similar devices
 - Peer-to-peer model of interactions
- Participants are both:
 - Data owners (clients of backup service)
 - Contributors (providers of backup service)
- Backup = protection of **critical private data** against:
 - Permanent and transient faults affecting a data owner
 - Theft or loss of a data owner


Scenario without MoSAIC



Scenario with MoSAIC



Intermittent access to infrastructure



Cooperative Backup for Nomadic Devices


- Backup = protection of **critical private data** against
 - Permanent and transient faults affecting a data owner
 - Theft or loss of a data owner

Cooperative Backup for Nomadic Devices

- Backup = protection of **critical private data** against
 - Permanent and transient faults affecting a data owner
 - Theft or loss of a data owner
- New threats on backups
 - Malicious (and accidental) faults affecting availability of data backups
 - Malicious (and accidental) modification of data backups
 - Malicious read access to data backups
- New threats on service
 - Selfish denial of service (refusal to cooperate)
 - Free-riding : consumption without contribution
 - “Tragedy of the commons” (Hardin 1968)
 - Attacks must be made unprofitable
 - Malicious denial of service (sabotage)
 - Attacks must be made ineffective or too costly

Cooperative Backup for Nomadic Devices

- Challenges
 - No prior organization
 - Ephemeral interactions
 - Limited energy, computation and storage
 - Only intermittent access to a fixed infrastructure
- + Usual criteria for classic functionalities
 - User transparency
 - Usability
 - etc.

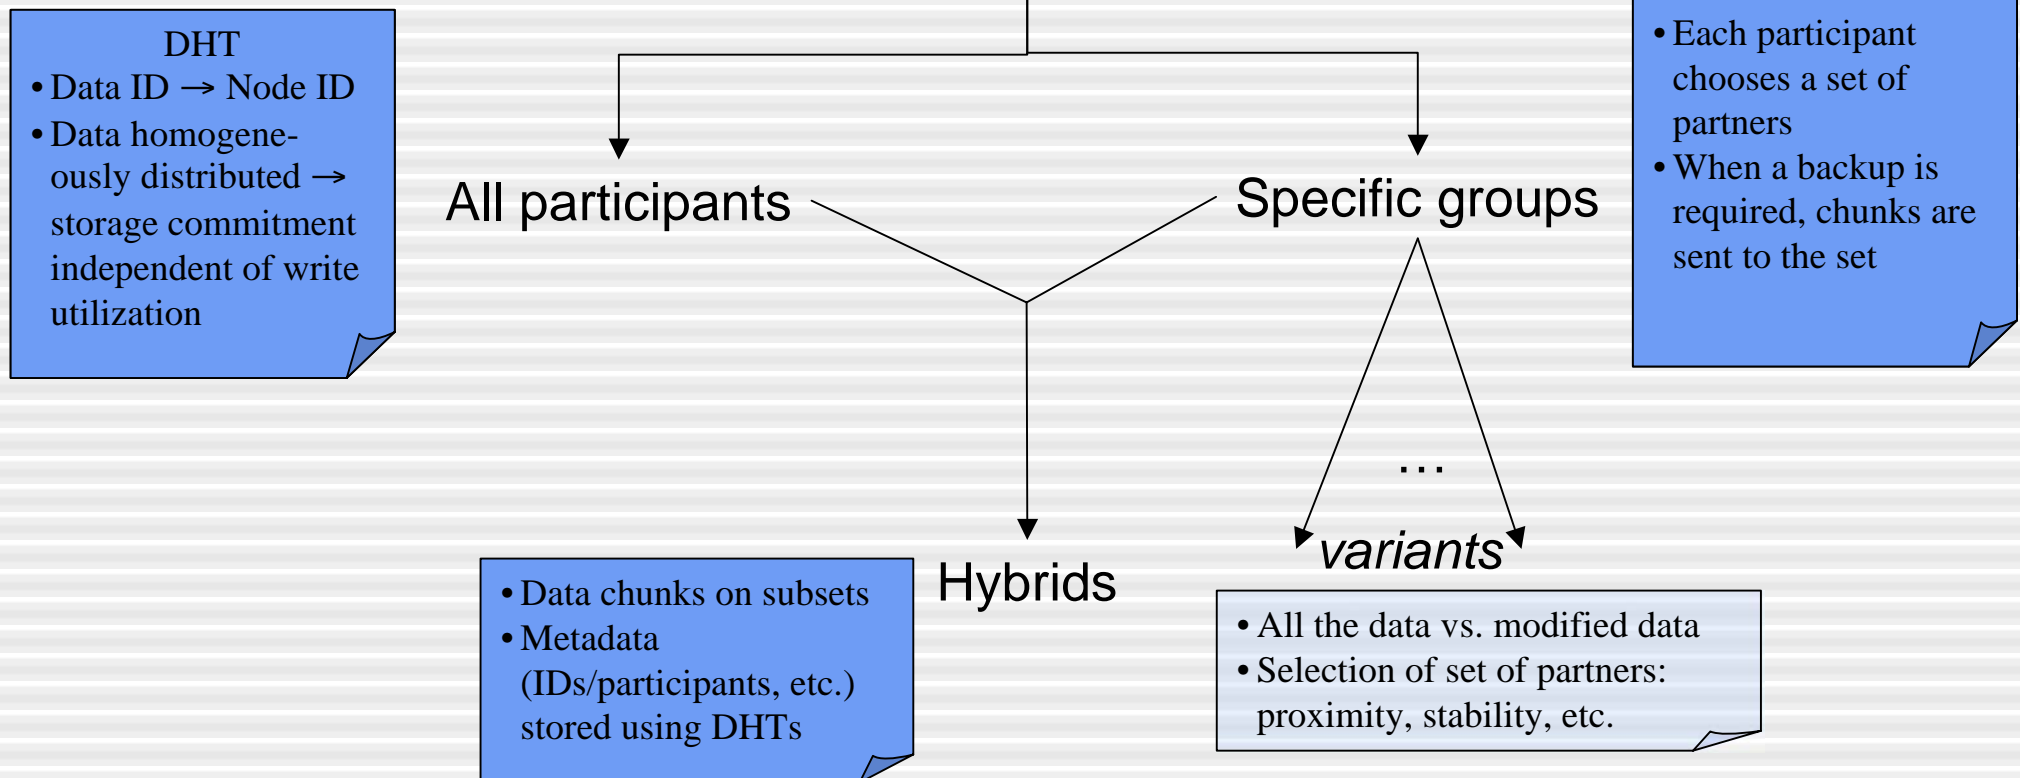
 M.-O. Killijian, D. Powell, M. Banâtre, P. Couderc and Y. Roudier, “Collaborative Backup for Dependable Mobile Applications [Extended Abstract]”, in *2nd Workshop on Middleware for Pervasive and Ad-Hoc Computing. Middleware 2004 Companion*, (Toronto, Canada), pp.146-49, ACM Press, 2004.

P2P Storage Systems

- WAN peer-to-peer systems
 - File sharing ➡ *Overlay networks, DHT*
 - GUNet
 - FreeNet
 - OceanStore
 - Backup ➡ *Cooperation incentives, trust*
 - Elnikety et al.
 - Pastiche
 - PeerStore
 - pStore
- PAN peer-to-peer systems
 - Backup
 - Flashback

Storage space discovery and allocation

Data chunk distribution



WAN P2P backup vs MoSAIC

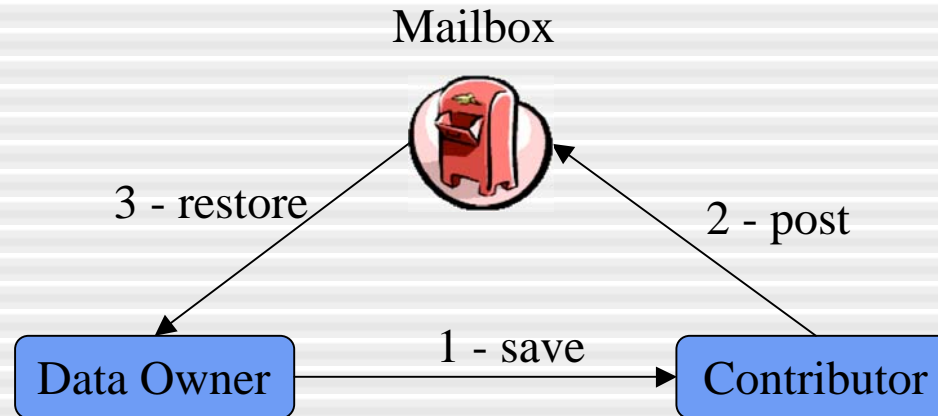
- Similar problems, but solutions not transferable to nomadic device scenario...

<i>Connections & bandwidth</i>	Stable	Unstable
<i>Dynamics</i>	Low (fixed)	High (mobility)
<i>Resource discovery</i>	Know somebody	Dynamic
<i>Access to fixed infrastructure & TTPs</i>	Continuous	Intermittent » trust mechanisms for disconnected operation

- ...except content-based addressing & convergent ciphering?
 - Use hash of content as an address
 - Allows backup optimization by exploiting inter-file redundancy (in addition to compression to exploit intra-file redundancy)

Current work at LAAS

- “Mailbox” model for storing the backup chunks




- Accommodates several restoration modes
 - **Push**: the contributor sends the chunks back home
 - Internet access, mailbox at the owner's home
 - **Pull**: the data owner searches for the data when necessary
 - Ad hoc network, mailbox hosted by the contributor
 - **Push-pull**: storage service as an intermediary
 - Internet access, mailbox hosted by a reliable storage service

Resource discovery

- Discovery of MoSAIC devices

- Online
- Creation of ad hoc network
- Active beaconing:
low latency vs energy economy



- Discovery of Internet access

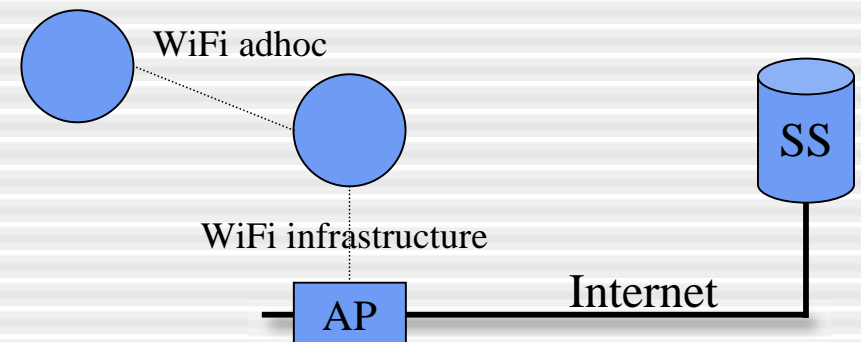
- Be able to backup to mailbox on reliable storage service

- Ad hoc and infrastructure mode at the same time

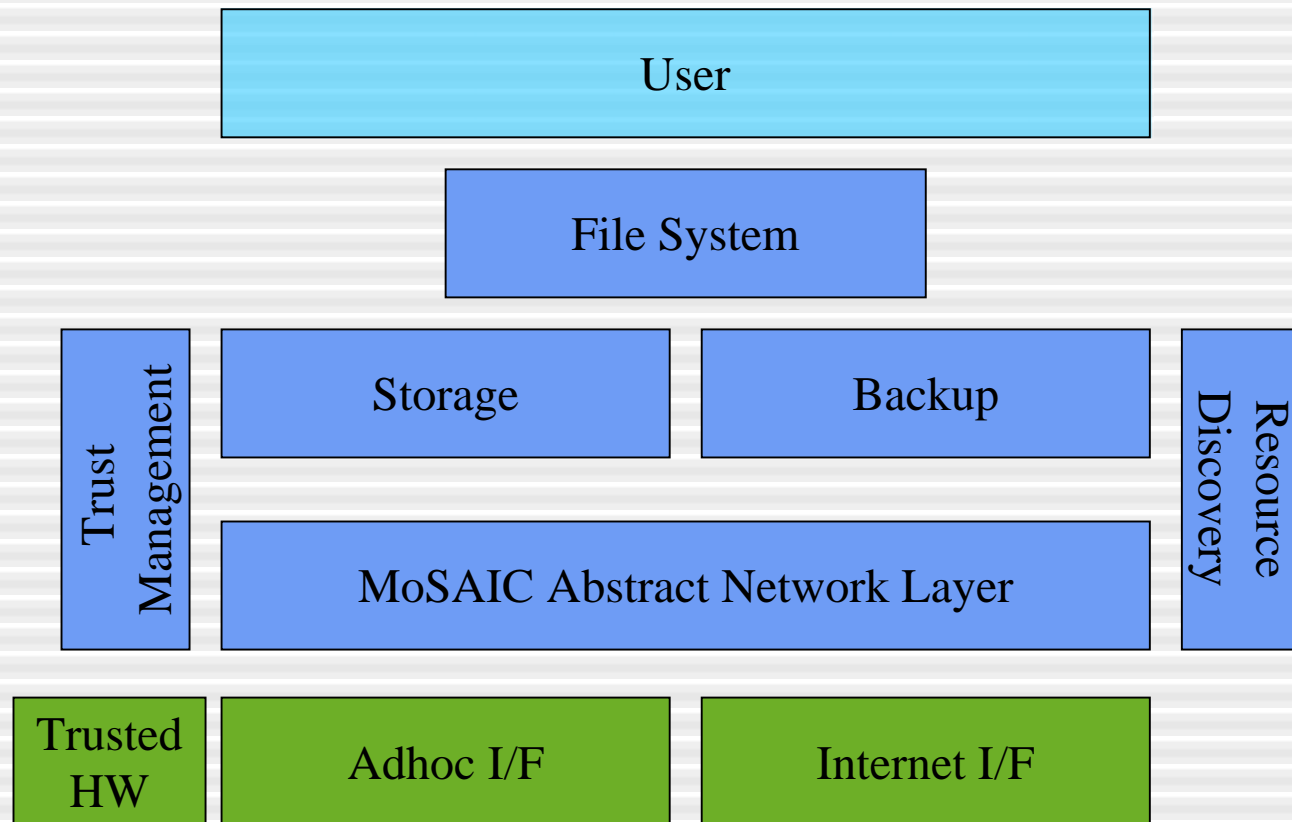
- Inter-device cooperation + storage service access

- One multiplexed network interface
- Two network interfaces

- Cooperative access to Internet?



Prototype Device Architecture



Current work at Eurecom

- Trust problems specific to cooperation
 - Will my data be correctly backed up?
 - What replication style is required for reliable backup?
 - When can data backed up for other devices be safely purged?
 - Is this backup request an attempted DoS?
- Establish trust by evaluating quality of cooperation
 - Reputation mechanisms
 - Remuneration mechanisms

Cooperation through credits

- Cooperation encouraged by secure exchange of credits
 - No on-line authority (ad-hoc mode)
 - Partial solution via neutral secure kernels
- How can we guarantee *fair* exchange (credit ↔ backup)?
 - Solution: optimistic fair exchange protocol
 - Uses TTP if non-cooperation is suspected
 - Secure kernel (representative of the TTP) keeps trace of events
 - Reconciliation by the TTP (when connected)
- NB: detection and punishment of non-cooperation cannot be immediate in a backup service
 - Deferred but direct detection of non-cooperation in pull (ad-hoc) mode
 - Deferred and indirect detection of non-cooperation by reliable storage service in push-pull (intermittent access) mode
 - Link between fair exchange TTP and reliable storage service?

Prototype under development

- Pragmatic choices
 - Secure kernels: Javacards
 - Wireless LAN
 - Javacards
 - Storage and exchange of credits
 - Log of backup operations
 - Also backup price “negotiation”
 - TTP
 - Arbitrate conflicts not decidable in distributed fashion (no clock on smartcards)
 - “Reimburse” attacked entities
 - Validate backup execution and punish attackers
 - Connection to TTP for conflict arbitration mitigated when infrastructure connection is necessary for long-term backup
- Current implementation
 - Objective is to validate crypto protocol
 - TTP arbitration not yet managed
 - Actively seeking more lightweight solutions

Current work at IRISA

- Simulation model of backup scenario with N devices and 1 infrastructure-based server
- Evaluation of backup device selection policy
 - Favor devices with most remaining energy
 - Favor devices judged to be more likely to reconnect soon to infrastructure
- Initial results
 - MoSAIC inter-device backup strategy considerably better than waiting for infrastructure connection opportunity
 - Current backup device selection policy no better than random choice
- Model to be extended to N device + P infrastructure-based servers

Data restoration issues

- Localization of data on multiple infrastructure-based servers
- Reconciliation of concurrent backups
 - restoration using backup of an old version
 - before completion of backup of more recent version
- Accounting for inter-file dependencies

Conclusion

- Scenario for
 - Designing new algorithms
 - Developing new middleware
- Fault-tolerance
 - Classic faults
 - Devices: crash of devices (owners and contributors), etc.
 - Data: integrity, confidentiality
 - Interaction faults (selfishness, maliciousness)
- New FT-enabling mechanisms
 - Self-carried reputation, virtual money, etc.
 - Opportunistic Internet backup, P2P interactions
- Project is 14 months old, still a lot of interesting things to do