SEVENTH FRAMEWORK PROGRAMME
THEME 3
Information and Communication Technologies

Grant agreement for:
Collaborative project, Small and medium-scale focused research project (STREP)

Deliverable D2.2:

Set of Schemes for Designing the Adaptable Network Infrastructure

Project acronym: FRONTS
Project full title: Foundations of Adaptive Networked Societies of Tiny Artefacts
Grant agreement no.: 215270

Responsible Partner: UDRLS
Report Preparation Date: Friday, 12 February 2010

Contents

1 Introduction

We present a set of schemes for designing the adaptable network infrastructure in a network of tiny artefacts that operates in a highly dynamic environment: nodes may move, links are noisy and unreliable, tiny artefacts are prone to faults, and in some cases, only opportunistic networking is possible. Furthermore, tiny artefacts are usually characterized by limited communication and computational resources, which limit the complexity of the proposed solutions, and since energy is usually provided by batteries, energy efficiency is a primary goal.

In such dynamic context, our first goal is to provide a scheme to design a simple, scalable, adaptable and energy efficient communication infrastructure (TASK 2.1). The proposed solution is based on a geographic routing algorithm where traditional 2-dimensional coordinates are replaced by virtual ones that represent the distance to some arbitrary anchor. Virtual coordinates can be effectively managed by tiny devices and their use leads to the construction of short paths and consequently to the reduction of energy consumption.

Once a suitable network infrastructure is constructed, our goal is to support the energy efficient exchange of information among nodes of the network. In a second scheme, we consider a dynamic mobile environment, where a data management system can create and delete copies of a file and place them on the mobile nodes so as to minimize the energy required to share the file with the other members of the network.

In some scenarios, the creation of an end-to-end path between the source and the destination is not even possible, because either the network is disconnected or nodes cannot directly initiate a communication (e.g., passive RFID tags which answer only when queried). In the latter case, a kind of undirected and mediated communication can be achieved using stigmergy. The information exchange among nodes is mediated by users' collective and unpredictable navigation patterns, and nodes have to adapt their internal structure (TASK 2.2) to manage the traces left by the users. In such context, we provide a scheme to support the design of a recommendation system where only local and simple information is used to recommend items in a fully decentralized fashion.

We have already discussed in D2.1 how game theory can help in designing adaptive networks by allowing the design of distributed methods to self-reach equilibrium. In such settings, where autonomous actions are introduced and social cooperation is used, entities may try to gain more by taking control over other entities. In the fourth scheme we present a self-stabilizing solution to cope with subsystem takeover adopting punishment for deterring improper behavior.

Finally, we propose a scheme to design reliable security services also on devices with extremely limited computational power (TASK 2.3). In this context, the hardware on the tiny artefacts might not support the implementation of such services and

TASK	SCHEME
2.1	Geographic routing with energy efficiency
	Dynamic file allocation strategy
2.2	Fully decentralized recommendations in opportunistic networks
	Strategies for repeated games with subsystem takeovers implementable by deterministic and self-stabilizing automata
2.3	Helper based security

Table 1: Assignment of schemes to WP2 tasks

thus we envision the support of security helpers. Security helpers are a set of devices that help tiny artefacts to perform cryptographic operations, and that are part of cryptographic protocols involving tiny artefacts with the specific task of helping the tiny artefacts. We show how privacy of users carrying objects attached to RFID tags can be protected when security helpers (i.e. randomizers) are present in the environment.

In [CKL$^+$er] we present a unifying use case scenario based on a new tiny artefact named MURPESS. A MURPESS, **MU**lti **R**adio **P**edestrian **E**nergy **S**cavenging **S**ensor, is a tiny battery powered device equipped with heterogeneous sensors and multiple radios, that can be embedded in shoes soles and can scavenge energy from pedestrians' steps. It is worth noting that most of the proposed schemes have been actually exploited for the design of the solutions discussed in [CKL$^+$er].

2 Geographic routing with energy efficiency

2.1 Motivation and problem description

Strategies designed to provide reliable communication involve the development of energy efficient routing algorithms such as geographic routing. In geo-routing, nodes have first to become aware of their position via some localization method; in our approach, we consider localization methods independent from external interventions. Then, some routing strategies must be implemented in order to guarantee the energy efficiency of messages delivery. This online problem needs adherence to the dynamics of the network and explores solutions for organizing distributed cooperation.

2.2 State of the art

The idea of locating nodes with little communication overhead, totally independently from satellite position information or other external interventions, has been explored in the area of navigation research (see [JLF99] as an example).

On the other hand, various geographic routing algorithms which use the coordinates of the sensors have been developed during the last decade. Most of these algorithms assume that nodes know their position via some technologies (e.g., GPS),

or via triangulation using the distance to some anchors. Geographic routing algorithms can be classified in: flood based, planar graph based, spanning tree based and geometric based. Most of them have theoretical guarantees to find a path from the source to the destination if it exists, however their efficiency is usually estimated via extensive simulations and only few of them propose some energy efficiency guarantee. In particular, [KWZ08] proposes an algorithm which guarantees paths of length $O(l^2)$ where l is the length of a shortest path. Furthermore, the authors prove that no algorithm not using node memory can guarantee a better result.

2.3 Scheme description

We introduce Virtual Anchor coordinates, a concept that involves the replacement of classical 2-dimensional coordinates by n-dimensional virtual ones that represent the distance to some arbitrary anchors. We then adapt routing algorithms so that they can route message on the surface that the network forms in the n-dimensional virtual space. A new routing protocol, ROAM, that includes classic geographic routing fallback strategies, as well as an obstacle avoidance and marking strategy, has also been developed to give energy-efficient results for routing messages in both classical and virtual coordinates settings.

ROAM, along with the standard geographic routing protocol GRIC and GPSR, were implemented using AlgoSensim simulator with both coordinates settings. Comparative simulations show that with various obstacle scenarios, both GPSR and ROAM guarantee a 100% message delivery if the network is connected. The average length of paths found by ROAM is shorter than that found by GPSR due to the avoidance of the dead-end zone created by the obstacle, and this is clearly shown by path smaller stretches. The energy consumption of routing protocols is directly tied to the number of exchanged messages; the number of messages exchanged by both GPSR and ROAM is directly proportional to the length of paths found and is the same in both cases. This means that the number of messages exchanged using ROAM protocol compared to GPSR is in the same proportion as the path stretches, and is thus sensibly better for ROAM. The message overhead caused by dead-end discovery in ROAM is an one-time cost per network: in average, ROAM generally starts to spend less energy than GPSR after as few as 10 computed paths. Moreover, the simulations showed that the replacement of actual by virtual coordinates did not in any case alter the success of path finding operation using any of the tested routing protocols. Interestingly, the path stretches are also similar whatever the coordinate system used. Therefore, it follows that cheap (in terms of energy, as well as hardware) virtual coordinate assignment does not lead to a hidden additional energetic cost in routing, compared to standard localization schemes. As a future work we plan to demonstrate the total equivalence in terms of routing efficiency for networks localized with virtual coordinates compared to networks using actual coordinates in any scenario, and we

plan to study the adaptation of virtual coordinate assignment to environments with highly mobile nodes.

2.4 Related FRONTS Technical Reports

In [HJer] we introduce *virtual anchor coordinates topology* which enables the energy efficient n-dimensional virtual localization of nodes in the absence of any external or expensive localization hardware. Networks that are virtually localized are able to route messages as well as networks making use of external localization (e.g. GPS).

In [HJL$^+$er] we introduce geographic routing with *obstacle avoidance mechanisms* using one bit of node memory, which enables energy efficient geographic routing, and which is theoretically and experimentally proven to give optimal paths up to a small constant factor.

3 Dynamic file allocation strategy

3.1 Motivation and problem description

We investigate power-aware strategies for sharing data in mobile ad hoc networks. This involves tools that enable the storage and retrieval of files dispatched into the network. A file consisting of several data units is stored in a mobile ad hoc network, where the mobile nodes can read and modify the file. A data management system can then create and delete copies of the file and place them on the mobile nodes. The goal of the data management system is to place copies such that the overall power consumption of the occurring data transfers is minimized. This problem is intrinsically online and needs adherence to the dynamics of the network.

3.2 Scheme description

We study the file allocation problem in a dynamic star network described in [MMer] when the file is located in a wireless mobile ad hoc network. Besides the mobile nodes, there is also a stationary server which is connected to the wireless network via one or more access points. The communication between the server and the mobile nodes is cryptographically authenticated. A file allocation algorithm can create and delete copies of the file on the mobile nodes. One of the online algorithms we considered is COUNT. COUNT has a counter c^i for every file allocation node i which may hold integer values ranging from 0 to $D+1$. On a read request from i, COUNT increases c^i by 1 (but not exceeding $D+1$). If c^i reaches $D+1$ and i does not already hold a copy, a new copy is created on i. On a write request, COUNT decreases all the counters by 1 (but not below 0). Thereafter all the copies on file allocation nodes with a counter value of 0 are deleted. We have shown that COUNT is $\max\left(3, 1+\frac{3D}{p}\delta\right)$-competitive, where p is the stand-by power consumption of the involved mobile nodes

and δ is a measure of the dynamics of the network, and we have given a tight lower bound. Furthermore, we have conducted simulations of COUNT (c.f. D4.4) which show that COUNT performs much better on realistic inputs than the worst-case suggests.

3.3 Related FRONTS Technical Reports

In [MMer] we introduce the *file allocation problem in a dynamic star network* and provide upper and lower bounds of the competitive ratio. This problem is a special case of a data management problem where the authenticity of the data has to be guaranteed.

4 Fully decentralized recommendations in opportunistic networks

4.1 Motivation and problem description

The limitation imposed by the hardware and the dynamic and challenging environments in which tiny artefacts might operate, in particular in mobile scenarios, make hard (or even impossible) in some contexts the construction of reliable network connections over which services can be delivered to the end-users. In these contexts, an opportunistic networking approach can be more effective. In opportunistic networks, communication opportunities are based on sporadic and intermittent contacts, disconnections and re-connections may frequently occur, and the assumption on the existence of an end-to-end path between the source and the destination is released. In the following we present a schema for recommending items in a fully decentralized fashion where only opportunistic networking is possible.

4.2 State of the art

In the last ten years recommendation systems have been recognized as an important research area and much work has been done both in industry and academia on developing new approaches. A survey of the main approaches to recommendation applications can be found in [AT05]. Most papers in this area consider centralized settings. On the other hand, future mass deployment of pervasive networks opens the possibility of new scenarios for recommendation systems, and distributed recommendation strategies have been considered recently (see for example [APSPT05, AAAPS06]). Although these solutions are distributed, active cooperation between users is required. [XHY+07] considers recommendation in P2P file sharing systems. Here, a Distributed Hash Table is used to allocate the database of user past transaction among the nodes of the network. The system then adopts a memory based approach for prediction and recommendation of items. [WPLR06] uses a similar approach, but the storage and update of user information is performed differently and is determined

by the navigation of users. We remark that both strategies imply explicit communication among nodes of the network in order to maintain information on past user transactions, whereas in our case this is not possible.

4.3 Scheme description

In this scheme, tiny artefacts cannot directly communicate and information exchange among them is mediated by users' collective and unpredictable navigation patterns, according to the principles of opportunistic networking. In this challenging scenario, we propose and investigate the effectiveness of stigmergy-based, collaborative filtering techniques. Stigmergy is a form of self-organization where traces left in the environment by the actions of agents stimulate the execution of subsequent actions, by the same or a different agent, thus allowing spontaneous and indirect coordination between agents.

As an example consider the following scenario: movies are advertised by means of smart posters, namely usual posters with a passive tag incorporated, from which users can extract data by touching it with their NFC-enabled handset. A tag in a poster cannot directly communicate with the other tags, but when a user accesses it by means of a smart reader (e.g., an NFC mobile phone), some information on the tag can be transferred to the reader (i.e., the smart reader collects traces left in the environment). Upon visiting another smart poster, this information or the results of its elaboration can be eventually delivered to the corresponding tag (i.e., traces are left in the environment), allowing some kind of indirect and mediated communication among tags.

In this schema, we assume that items of interest are tagged with smart tags (such as passive RFIDs), storing aggregate information about the visiting patterns of users that interacted with them in the past. Users access and modify information stored in smart tags transparently, by smart reader devices. Smart readers use private information about previous behavior of the user and aggregate information retrieved from smart tags to recommend new items that are more likely to meet user expectations.

In our model, every user j enters the system, visits a subset of the items and then leaves the system. We call this a *session*. The identities of users are not stored, hence multiple visits of the same user to the smart shop are not individually tracked. However, we emphasize that information about multiple visits of the same user is stored in aggregate form at smart tags.

Modeling user behavior in the system entails two aspects: i) describing the way in which users are attracted by (i.e., visit) items of potential interest to them and ii) the sequential order in which the items are visited. This latter aspect is important, since it determines the way in which information about users' past visits is spread across the pervasive system.

We assume a weighted visit model where the probability of visiting an item is pro-

portional to the item's popularity within a cluster (e.g., comedy movies). Note that, while different users' visits are independent, the next items visited by a user clearly depend on the items he/she previously visited. In particular, our weighted visit model describes a strong dependence between user preferences and visits. To consider the bookshop example, many people are likely to be first attracted by popular, recently published books in their fields of interest.

The above model is intended to strike a balance between simplicity and soundness. It is clear that this choice brings to some simplification with respect to the scenarios of potential interest. The recommendation based on items' popularity can be sensitive to changes of users' visit patterns over time. Also, assuming that rankings inside clusters only depend on items' popularities may be unrealistic in some scenarios. Furthermore, visit patterns might depend on different (e.g., geometric and physical) constraints.

Our algorithms do not require any explicit interaction among users and can be easily and efficiently implemented. Denote by R a vector storing the current estimation of users' preferences for each item. Upon visiting item r belonging to cluster C_k, the smart reader of user j recommends the T *top* ranking items in R that i) belong to C_k and ii) have not yet been visited by j.

Note that R is read from the smart tag, updated and then it is used by the smart reader to recommend T top scoring elements. Then, the updated version of R is stored back on the smart tag, replacing the older one with an improved estimation of the users' preferences.

4.4 Related FRONTS Technical Reports

In [BCMSVera] we propose and investigate the effectiveness of recommendation techniques based on collaborative filtering for use in fully decentralized, pervasive systems of small devices with limited communication and computational capabilities. We analyze the theoretical behavior of our recommendation strategies and assess their performance in practice, both by simulation using synthetic data and on real, publicly available datasets. We evaluate the performance of the system along two main axes: the ability to accurately reconstruct the user profiles and the quality of recommendations. In particular, we evaluate the former in terms of ranking similarity and convergence speed, while the latter is evaluated in terms of standard measures of quality used in information retrieval such as hit ratio, precision and recall. Experimental results confirm that in the considered scenarios, the proposed techniques can well reconstruct the user profiles and at the same time provide recommendations of good quality. A new and updated version of this paper [BCMSVerb], with a significantly extended experimental section, has been submitted to the Knowledge and Information Systems Journal (KAISS).

More recently, we propose [BBC$^+$er] a fully decentralized approach for recom-

mending new friendships and contacts in the social network of mobile phone users. With respect to existing solutions, our approach is characterized by some distinguishing features. In particular, the application we propose does not assume any centralized coordination: it transparently collects and processes user information that is accessible in any mobile phone, such as the log of calls, the list of contacts or the inbox/outbox of short messages, and exchanges it with other users. This information is used to recommend new friendships to other users. Furthermore, the information needed to perform recommendations is collected and exchanged between users in a privacy preserving way. Finally, information necessary to implement the application is exchanged transparently and opportunistically, by using the residual space in standard short messages occasionally exchanged between users. As a consequence, we do not ask users to change their habits in using SMS.

5 Strategies for repeated games with subsystem takeovers implementable by deterministic and self-stabilizing automata

5.1 Motivation and problem Description

We consider heterogeneous and unreliable environments in which sensors are deployed. We assume that each participating entity (e.g., sensors) takes autonomous decisions, rather than execute a fixed protocol that is common to all sensors. Such settings are prone to malfunctioning, because social cooperation is required between entities that can take only autonomous actions which are considered selfish. We look into the case in which entities gain control over other entities, say, by means of electronic malware insertion. The problem at hand is how autonomous (and selfish) entities can deter such takeovers.

Game theory provides the framework for modeling and preventing such situations by using punishment for deterring improper behavior. However, it may happen that due to faults, selfish-computers may punish well-behaved ones. This is one of the key motivations for forgiveness that follows any effective and credible punishment. Therefore, unplanned punishments must be proven to have ceased in order to avoid infinite cycles of unsynchronized behavior of "tit for tat". The ability to overcome an arbitrary initial state, to be self-stabilizing, is an inherent part of a robust system that withstands temporary violation of assumptions (be it the memory corruption, or the size of the subsystem that was taken over).

5.2 State of the art

The literature usually considers a deviation of a single agent rather than a joint deviation of a coalition of agents (e.g. [Nas50, CNL$^+$07]). When joint deviations are considered, the coalition of deviators is required not to compromise the benefit of its

members (e.g., [AFM09, ADGH06]). We consider the harder case of joint deviations, in which only the benefit of a single deviator, the one that has taken over, is required to increase. Moreover, unlike existing work in the area of game theory [Rub80], we show that the strategy implementation is optimal and self-stabilizing.

5.3 Scheme description

Our scheme is designed to cope with subsystem takeovers and to be self-stabilizing. The agents initially agree on a cooperation sequence and monitor the environment for any deviation from the agreed sequence. After the occurrence of a deviation, the deviators are punished. The mechanism is able to cope with any sequence of deviations, e.g., deviations occurring during the punishment phase. Moreover, our solution is guaranteed to recover from a transient fault, i.e., in order to avoid infinite cycles of unsynchronized behavior of "tit for tat," we prove that unplanned punishments cease.

We construct deterministic and finite automata that implement these strategies with optimal complexity. Moreover, we prove that all unplanned punishments eventually cease by showing that the automata can recover from transient faults. Namely, we show that the system of automata is self-stabilizing.

5.4 Related FRONTS Technical Reports

The solution concept that we consider uses a strong Nash equilibrium in infinitely repeated games: We assume that the system designer is able to find a punishment sequence of actions to be executed over time and to be used for deterring agents from deviation. It is well-know that finding such sequences can be computationally hard. Therefore in [DSSTer] we consider another type of solution concept, which is a correlated equilibrium. For that solution concept we propose an algorithm for finding the punishment sequence.

6 Helper based security

6.1 Motivation and problem description

Security for tiny artefacts is conceptually different from the security in other domains (like Internet and Network security). Indeed, tiny artefacts are severely bounded in computation speed and storage capacity. These *technology limitations rule out traditional cryptography* (based on algebraic problems in $Z_n^\star$) as too expensive (reasonable levels of security are achieved for moduli of one thousand bit length) and suggest as a possibility Cryptography based on Pairings (here a 128 bit moduli guarantees security comparable with 1024-bit RSA modulus). In addition, tiny artefacts are subject to attacks (like reset attacks) which are not possible if the adversary does not have physical access to the computing device (which is often the case for tiny artefacts).

It is also crucial to study privacy problems associated with the use of tiny artefacts to trace and locate objects which could threat the privacy of the individuals that are associated with the objects being traced.

6.2 State of the art

RFID-based systems are becoming a widely deployed pervasive technology that is more and more used in applications where privacy-sensitive information is entrusted to RFID tags. Thus, a careful analysis in appropriate security and privacy models is needed before deployment to practice. Our investigation has started from the security model introduced by Vaudenay [Vau07] that is considered standard for dealing with security and privacy. The strongest achievable notion of privacy in this model (narrow-strong privacy) requires public-key cryptography, which exceeds the computational capabilities of current low-cost RFIDs. Other privacy notions achievable without public-key cryptography heavily restrict the power of the adversary and thus are not suitable to realistically model the real world.

6.3 Scheme description

As security for tiny artefacts poses different conceptual problems than the ones studied in Internet security, these problems call for new paradigms and schemes on which security needs to be based. One of the main differences is in the power of the devices that must perform the computation required by the cryptographic protocol. This justifies the adoption of the *helper-based security* scheme. In this scheme, we envision a set of devices, generically called the *security helpers*, that help tiny artefacts to perform cryptographic operations and that are part of cryptographic protocols involving tiny artefacts with the specific task of helping the tiny artefacts. We envision the security helpers as woven into the environment and providing their service to all the devices that request it. One crucial question is "*how much can we trust a security helper*?" Ideally, one would like to be able to make minimal trust assumptions on the helpers but typically we believe it is reasonable to assume that the helpers perform as prescribed. In no case, however, the helper must receive secret information from the artefacts (like for example secret keys).

We have already studied security protocols that fit the helper-based security scheme. In [BDCPer, SVWerb, SVWera], we show how privacy of users carrying objects attached to RFID tags can be protected when security helpers are present in the environment. In this case the security helpers are called *randomizers* as their task is to appropriately randomize the content of the RFID tags. In [BDCPer] the authors give a construction of a cryptographic scheme for randomizers that make minimal trust assumptions. Specifically, randomizers are not assumed to have the secret key needed to decrypt the ciphertexts carried by the RFID tags, but can, nonetheless,

re-randomize the tags thus making tracing of tags more difficult.

We expect other security helpers to be designed so that security for tiny artefacts can be achieved. In our research of reset attacks, we have identified that the main source of difficulty of obtaining resettable identification schemes for tiny artefacts (to be used for example in personal identification documents, like passports) is the fact that the party having to verify the device is allowed to send random challenges. In this case, we expect that a security helper, which we call a *random beacon*, broadcasting random sequences of strings, can be very helpful.

6.4 Related FRONTS Technical Reports

In [DSVera, DSVerb] we have extended the standard model for security and privacy introduced by Vaudenay to take into account Denial of Service attacks. Specifically, we extended the model to better reflect the real-world scenario, where these attacks are easy to mount (e.g., by physically destroying/making inactive tags). Besides refining Vaudenay's privacy model to deal with DoS and DoS-like attacks, we introduce an additional privacy notion, referred to as *semi-destructive* privacy, which takes into account hardware features of some real-world tags. Then, we show an efficient RFID protocol that, by only using symmetric-key cryptography (which is much more efficient than public-key cryptography), satisfies the notion of semi-destructive privacy, under standard complexity-theoretic assumptions.

In [SVWerb, SVWera] we have extended and improved the model of Vaudenay by introducing a security and privacy model for *anonymizer-enabled* RFID systems. Our model builds on top of Vaudenay's model and supports anonymizers, which are separate devices specifically designated to ensure the privacy of tags. We present a privacy-preserving RFID protocol that uses anonymizers and achieves narrow-strong privacy without requiring tags to perform expensive public-key operations (i.e., modular exponentiations), thus providing a satisfying notion of privacy for cost-efficient tags.

In [BDCPer] we give an efficient construction of *untraceable* tags in the model with randomizers, to be used for passive RFID tags carrying private encrypted information. The interaction with the anonymizers is intended to guarantee that tags are not traced, and this is achieved by re-randomizing the ciphertext (that is, by producing a new ciphertext that carries the same message). The re-randomization procedure is carried on without having access to the private key used to generate the ciphertext. This extra important property guarantees privacy of the information even with respect to the randomizers. Our construction uses pairing-based cryptography that allows high level of security with very limited storage, making our construction applicable to RFIDs.

In [SVWerc] we have investigated issues regarding location privacy, which is an extra security guarantee that is important for RFID systems while almost non-existent

for Internet security. In [SVWerc] we explore the advances in RFID applications, stressing the security and privacy location shortcomings of existing proposals. Finally, we illustrate new promising directions for privacy-preserving RFID-enabled systems, considering electronic tickets (e-tickets) for public transportation as a case study.

Small artefacts are already being used to identify people as they traverse borders. As these means of identification become more and more pervasive, we believe that their security properties and the threats they pose to the privacy of the citizens must be critically analyzed. In [ABDC$^+$er] we give an overview of the recently introduced e-passports and discuss the specific implementation adopted in Italy. Starting from the known attacks to e-passports adopted in other countries, we present the results of a security analysis of Italian e-passports, while we also consider additional critical issues. Our analysis shows that Italian e-passports are subject to eavesdropping attacks, where one can unnoticeably obtain private data stored in the e-passport using RF communication while the passport is stored in a bag/pocket. In order to obtain such a result, we had to address novel issues about the format of the Italian machine readable zone. We have designed and implemented several open-source tools to analyze and attack Italian e-passports; among them, we show simulators that allow one to produce the attack also without physically having e-passports and RFID equipment. Our tools also address forward security issues, by linking (in off-line) two different transcripts corresponding to two conversations of the same e-passport with readers. Finally, we have analyzed the quality of the randomness used by Italian e-passports during RFID communication, and the conclusion is that the random number generator included in the RFID chips is extremely weak, thus also exposing Italian e-passports to several other attacks.

References

[AAAPS06] Noga Alon, Baruch Awerbuch, Yossi Azar, and Boaz Patt-Shamir. Tell me who i am: an interactive recommendation system. In *SPAA '06: Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures*, pages 1–10, New York, NY, USA, 2006. ACM.

[ABDC$^+$er] Vincenzo Auletta, Carlo Blundo, Angelo De Caro, Emiliano De Cristofaro, Giuseppe Persiano, and Ivan Visconti. Increasing privacy threats in the cyberspace: the case of italian e-passports. Technical report, * 2009. FRONTS-TR-2009-35.

[ADGH06] I. Abraham, D. Dolev, R. Gonen, and J. Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing

and multiparty computation. In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, page 62. ACM, 2006.

[AFM09] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. *Games and Economic Behavior*, 65(2):289–317, 2009.

[APSPT05] Baruch Awerbuch, Boaz Patt-Shamir, David Peleg, and Mark R. Tuttle. Improved recommendation systems. In *SODA*, pages 1174–1183, 2005.

[AT05] Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowl. Data Eng.*, 17(6):734–749, 2005.

[BBC+er] Luca Becchetti, Lorenzo Bergamini, Ugo Colesanti, Luca Filipponi, and Andrea Vitaletti. A lightweight sms-based recommendation system for mobile users. Technical report, UDRLS, 2010. FRONTS-TR-2010-6.

[BCMSVera] Luca Becchetti, Ugo Colesanti, Alberto Marchetti Spaccamela, and Andrea Vitaletti. Fully decentralized recommendations in pervasive systems: Models and experimental analysis. In *First International Workshop on Data Warehousing and Knowledge Discovery from Sensors and Streams*, 2009. FRONTS-TR-2009-72.

[BCMSVerb] Luca Becchetti, Ugo Colesanti, Alberto Marchetti Spaccamela, and Andrea Vitaletti. Recommending items in pervasive scenarios: Models and experimental analysis. Technical report, 2010. Submitted to Knowledge and Information Systems (KAIS) - by invitation. FRONTS-TR-2010-3.

[BDCPer] Carlo Blundo, Angelo De Caro, and Giuseppe Persiano. Untraceable tags based on mild assumptions. In *2nd SETOP International Workshop on Autonomous and Spontaneous Security*, 2009. FRONTS-TR-2009-36.

[CKL+er] Ioannis Chatzigiannakis, Spyros Kontogiannis, Guy Leshem, Jan Mehler, Giuseppe Persiano, and Andrea Vitaletti. Murpess - multi radio pedestrian energy scavenging sensor. a unifying scenario. Technical report, 2010. FRONTS-TR-2010-7.

[CNL+07] A. Clement, J. Napper, H. Li, J.P. Martin, L. Alvisi, and M. Dahlin. Theory of bar games. In *Proceedings of the twenty-sixth annual ACM*

symposium on Principles of distributed computing, page 359. ACM, 2007.

[DSSTer] Shlomi Dolev, Elad Michael Schiller, Paul Spirakis, and Philippas Tsigas. Strategies for repeated games with subsystem takeovers implementable by deterministic and self-stabilizing automata. *International* * *Journal of Autonomous and Adaptive Communications Systems*, 2009. FRONTS-TR-2009-4.

[DSVera] Paolo D'Arco, Alessandra Scafuro, and Ivan Visconti. Revisiting dos attacks and privacy in rfid-enabled networks. In *ALGOSENSOR*, * LNCS, 2009. FRONTS-TR-2009-27.

[DSVerb] Paolo D'Arco, Alessandra Scafuro, and Ivan Visconti. Semi-destructive privacy in rfid systems. In *The 5th Workshop on RFID* * *Security – Revised Selected Papers*, 2009. FRONTS-TR-2009-63.

[HJer] Florian Huc and Aubin Jarry. Vrac: Virtual raw anchor coordinates routing in sensor networks. In *Proceedings of WONS 2010*, February * 2010. FRONTS-TR-2010-1.

[HJL$^+$er] Florian Huc, A Jarry, Pierre Leone, Jose Rolim, Luminita Moraru, and Sotiris Nikoletseas. Early obstacle detection and avoidance for all to all traffic pattern in wireless sensor networks. In *Proceedings of the Fifth International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS 09)*, Lecture Notes in Computer Sci- * ence (LNCS). Springer Verlag, 2009. FRONTS-TR-2009-43.

[JLF99] C. Gary Stephens James L. Farrell, Edwin D. McConkey. Send measurements, not coordinates. *NAVIGATION - Journal of The Institute of Navigation*, 46(3), 1999.

[KWZ08] F. Kuhn, R. Wattenhofer, and A. Zollinger. An algorithmic approach to geographic routing in ad hoc and sensor networks. *IEEE/ACM Transactions on Networking*, 16(1):51–62, 2008.

[MMer] Jan Mehler and Friedhelm Meyer auf der Heide. Power-aware online file allocation in mobile ad hoc networks. In *SPAA '09: Proceedings of the twenty-first annual symposium on Parallelism in algorithms* * *and architectures*, pages 347–356, New York, NY, USA, 2009. ACM. FRONTS-TR-2009-20.

[Nas50] J.F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, pages 48–49, 1950.

[Rub80] A. Rubinstein. Strong perfect equilibrium in supergames. *International Journal of Game Theory*, 9(1):1–12, 1980.

[SVWera] Ahmad-Reza Sadeghi, Ivan Visconti, and Christian Wachsmann. Anonymizer-enabled security and privacy for rfid. In Juan A. Garay, Atsuko Miyaji, and Akira Otsuka, editors, *Cryptology and Network Security, 8th International Conference, CANS*, volume 5888 of *Lecture Notes in Computer Science*, pages 134–153, 2009. FRONTS-TR-2009-65.

[SVWerb] Ahmad-Reza Sadeghi, Ivan Visconti, and Christian Wachsmann. Efficient rfid security and privacy with anonymizers. In *RFIDSec 2009, Revised Selected Papers*, 2009. FRONTS-TR-2009-64.

[SVWerc] Ahmad-Reza Sadeghi, Ivan Visconti, and Christian Wachsmann. Location privacy in rfid-enabled applications. In *Privacy in Location-Based Applications*, volume 5599 of *Lecture Notes in Computer Science*, 2009. FRONTS-TR-2009-69.

[Vau07] Serge Vaudenay. On privacy models for rfid. In *ASIACRYPT*, pages 68–87, 2007.

[WPLR06] J. Wang, J. Pouwelse, R. Lagendijk, and M. J. T. Reinders. Distributed collaborative filtering for peer-to-peer file sharing systems. In *21st Annual ACM Symposium on Applied Computing*, pages 1026–1030, 2006.

[XHY+07] Bo Xie, Peng Han, Fan Yang, Ruimin Shen, Hua-Jun Zeng, and Zheng Chen. DCFLA: A distributed collaborative-filtering neighbor-locating algorithm. *Information Sciences*, 177(6):1349–1363, 2007.