SEVENTH FRAMEWORK PROGRAMME
THEME 3
Information and Communication Technologies

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

Deliverable D2.1:

Definition of Research Targets and Algorithmic Solutions for Adapting the Communication Infrastructure, the Internal Structure, Roles and Security

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

Responsible Partner: UDRLS
Report Preparation Date: Tuesday, 14 April, 2009

Contents

1 Introduction

In this deliverable we will present the research activity on WP2 that is devoted the internal organization of the network to maintain the communication infrastructure, to facilitate adaptation mechanisms, and to guarantee basic services (e.g. security, clustering, data filtering etc.). We focus on internal preparation of the network, namely on all the mechanisms, tools and techniques that allow the network to self-react to the dynamic conditions in which it operates. This internal activity will facilitate the task performed in WP3 which is focused on how to react to external dynamic changes.

Pervasive networks of artefacts require scalable solutions for maintaining large heterogeneous networks of devices with limited resources, possibly connected via wireless links prone to communication failures. As networks become more complex and larger and pervasive systems move towards commercial applications we believe that scalability and heterogeneity in adaptive systems requires new methods and algorithms both for the self organization of pervasive systems and for assessing their performance.

In the first year of the project we have focused our attention on several specific problems for role definition and communication primitives that hide lower-level complications that arise due to mobility, failures and/or constraints on the available resources of the network and provide the required security and trust. Our proposed solution consider a number of requirements. The first one is the need of efficient and scalable solutions due to the size of the network and the scarcity of resources. Moreover the distributed nature of networks of tiny artefacts and their limited computational power implies that we must provide such systems with appropriate distributed algorithms and protocols that are simple to program and use only local information and simple rules. We have considered self-organizing algorithms and redundant network structures for connectivity that are suitable for fast dynamic maintenance.

The second important issue is represented by self stabilization. In fact, future networks of artefacts will be designed to execute forever, like an operating system. Thus it is highly unlikely that they never experience a transient failure (i.e., a temporary violation of the designer assumptions), especially in highly dynamic, wireless, mobile networks. Self-stabilizing distributed systems can recover after the occurrence of transient faults. The system is designed to automatically regain its consistency from any arbitrary state that may be the result of unexpected faults caused by the environment.

The final important requirement is to provide security to networks of tiny artefacts for which completely new approaches are deemed; in fact most of current security protocols have been designed to be implemented on processing units that have fair strong resources (memory, computing power, reliable communication). Unfortunately most of these assumptions are violated in networks of tiny artefacts with

limited computational power, limited memory and storage. Thus, in most cases, it is impossible to use classical techniques. We also need public key infrastructure tailored to dynamic local environments opposed to the classical approach that deals with a stable set of participants and non-dynamic infrastructure.

In order to assess the quality of our solutions we have used the tools of algorithm analysis to formally assess such requirements. Efficiency has been measured not only with respect to the computation time but also with respect to other parameters such as energy, memory requirement, required communication bandwidth, while the quality of the obtained solution has been analysed using different tools proposed for the formal analysis of algorithms (approximation guaranteed, competitive analysis,resource augmentation etc.). Another very important set of methodologies is provided by game theory and in particular by algorithmic game theory. We have shown that these techniques are useful not only for the analysis but also for the design of new algorithms.

With reference to the previous issues the activities carried on in the first year are discussed in detail in the following where the presentation of our work is split according to three tasks. We now briefly highlight main results.

In the first task we have focused on the issues of adapting the communication infrastructure. We have analyzed simple algorithms for data gathering in wireless networks providing a formal analysis of the quality of service (QoS) and for constructing spanners of the network that are sparse while providing a guarantee on the quality of the provided solution (measured in terms of the stretch). We recall that both problems are NP-hard and thus we do not expect to find efficient solution in polynomial time; the algorithms we propose and analyse are simple and thus amenable to implementation. We have also proposed new self stabilizing algorithms for partitioning the communication graph in clusters and for computing snapshots. These are two important primitives for adapting the communication infrastructure; in fact clusters hierarchy is a natural way of organize a network with the goal of efficiency in communication.

In the second task we have focused on several problems related to defining the role of artifacts when the external conditions change. We have considered methods for improving trust in a network prone to faults and to malicious hosts; in particular we have provided new methods for purifying data from Byzantine data items in machine learning algorithms and stabilizing trust in spite of Byzantine host (that have malicious behavior) in the network. We have also considered new equilibrium strategies to deal with systems of selfish-computers, by considering the possibility of subsystem takeover, say, by the use of hostile malware.

Furthermore we have considered two problems showing that simple algorithmic strategies are able to provide efficient solution despite the restrictions we impose on the communication and on the computational power. Namely we have considered two specific problems: how a network can maintain a proper coloring of its nodes and how

a distributed network can provide recommendations. The first problem is important for its paradigmatic role and because it is an important primitive in many applications. We define the problem as a coloring game. i.e. a strategic game where the set of players is the set of vertices and the players share the same action set. Surprisingly the new efficient coloring method achieves a number of colors satisfying (together) the known general upper bounds on the chromatic number. Our method is also an alternative general way of proving, constructively, all these bounds. For the second problem we show that simple algorithms might provide good recommendation for pervasive systems of small devices with limited or no communication.

In the third task we have focused our attention on several important aspects related to security issues. Namely we have proposed a new protocol for reliable authentication with privacy protection in RFID tags. The solution is based on an RFID tag that evolves rapidly in time and it is well suited for low-end devices, since all mechanisms can be easily implemented by circuits of a small size. We have also started an investigation on identification protocols that are able to provide security also with respect to powerful adversaries that have physical access to the device and thus are able to reset the device. Finally we have studied the problem of secure routing in wireless sensor networks where the sensors and the sink can move during the execution of remote monitoring applications. The mechanisms are simple to implement, rely only on local information and require constant storage per sensor. The protocol adapts to mobility and security challenges that may arise throughout the execution of the application.

2 Task 2.1: Adapting the communication infrastructure

2.1 Motivation and problem description

The future network of tiny artifacts we envisage will operate in an highly dynamic networked environment and will possibly operate even beyond the complete understanding and control of their designers, developers, and users. As a consequence, the communication infrastructure will perpetually adapt according to the principles of self-stabilization and self-organization to ensure fast automatic recovery from an arbitrary state and at the same time guarantee the minimum set of requirements to assure the network operations. Another important issue is simplicity both in terms of computational resources necessary to perform the tasks of adapting the communication infrastructure and of gathering and delivering information. Simplicity should not only be measured on the computational requirements, but also on the ability of exploiting only local information.

We investigate self-stabilizing distributed algorithms for cluster definition in com-

munication graphs of bounded degree processors, and for hierarchical distributed snapshots. We observe that as far as dynamic changes are limited within a cluster they do not affect the rest of the network.

In many situations of dynamic and large scale systems, competing entities of a network have secret incentives, whereas the system designer wishes to optimize a social welfare function that depends on these secret preferences (eg, the social welfare is the aggregation of the agents' levels of satisfaction). This leads to a socially optimal solution, which could also be seen as a fair solution in the sense that an average level of satisfaction is maximized. Additionally, the network operator should ensure that the operational costs of the whole system are covered by a fair cost sharing scheme announced to the competing entities. This implies that a (possibly anonymous) pricing scheme for the usage of the shared resources must be adopted. This scheme should be *robust* against changes in the demands of the entities (line operators). That is, we consider as *adaptivity* of the system the tolerance to the entities' elasticity of demand requests, and the eventual stabilization at an equilibrium point that is as close as possible to the social optimum.

Data gathering is one of the main communication primitives in a network of tiny artifacts. In this context, most of the effort is devoted to the design of energy efficient solutions, nevertheless in many application scenarios, more traditional metrics such as flow time and latency are relevant. It is important to design simple distributed and local algorithms to guarantee some kind of quality of service. In particular locality of the choices guarantees the ability of reacting and adapting to changes autonomously.

The development of decentralized services and applications for networked societies of tiny artifacts requires a basic service for sharing commonly used data. Such a data management service has to guarantee the availability (i.e. the data should not vanish because of single node failures) and the consistency of the data (i.e. a read request should always return the currently valid value). Since power is a scarce resource in mobile networks, it should furthermore minimize its energy consumption. The classical strategy of placing copies on preselected nodes is not suitable in a large-scale dynamic network, since the nodes with copies and the nodes accessing the data could move apart over time. So we need on-line strategies which dynamically create, migrate and delete copies in response to the dynamics of the network and the actual access pattern of the nodes.

2.2 State of the art and previous work

Self-organization. In recent years, the concept of self-organization has been widely mentioned in the scope of distributed computing and peer to peer networks. In [ADGR05] a framework for self-organization is proposed, including formal definitions of the self-organization concept and complementary proof techniques which

can be used to prove that algorithms are indeed self-organizing. Self-stabilizing and self-healing constructions of hierarchies, in the domain of sensor networks, appear in [ZA02]. The authors divide the plane into hexagonal cells. In each cell a head that corresponds with a cluster leader is elected. In [MW05], Wattenhofer and Moscibroda presents an algorithm for computing a maximal independent set in radio networks where processors can broadcast their messages asynchronously, but no collision detection mechanism is provided. Snapshot algorithms are used for recording a consistent global state of a distributed asynchronous system. A self-stabilizing snapshot algorithm was first introduced in [KP90], where repeated invocations of snapshots are used to ensure stabilization of a non-stabilizing algorithm. Following [KP90], several works have studied ways of achieving efficient snapshots in different models e.g., message passing, bounded links message passing and shared memory ([Var00], [AD97], [CDPV03]).

On-line data management. The problem of on-line data management has been thoroughly investigated for static networks (c.f. the survey of Bartal et al. [Bar96]), whereas very little work has yet been done in dynamic and mobile networks. A first step towards data management in dynamic networks was accomplished by Bienkowski et al. [BBKM08]. They developed on-line strategies and lower bounds for the migration of a single copy in a dynamic network.

2.3 First results

Self-stabilizing and self-organizing hierarchy definition. We define the hierarchy of subsystems [DT08] by partitioning the communication graph into small clusters, then merged to form bigger clusters and so on. The partition can be done according to a designer's input, using an automatic off-line clustering algorithm or even an on-line clustering algorithm that reflects the system's current behavior. The collected local topology supports a randomized local leader election, in which a non leader processor that does not identify a leader within a certain distance x tries to convert itself to a leader. Leaders within distance x from each other are eliminated, until there are no leaders that are within distance x or less from each other. Higher level partitions, using larger distances and overlay network abstraction between leaders, are constructed in a similar way. In asynchronous systems, our clustering algorithm uses (for each processor) a (local) self-stabilizing snapshot algorithm for obtaining local synchronization of actions.

Self-stabilizing snapshots. We present a self-stabilizing snapshot algorithm for distributed systems [DT08] that uses message passing with bounded link capacity, in which a spanning tree is distributively defined. On-demand stabilization ensures that regardless of the number of new requests (for snapshots), the system reaches a state, such that eventually any new request results in a correct output (snapshot). In other

words, stabilization does not rely on repeated invocations of new (snapshot) requests. Our on-demand self-stabilizing snapshot algorithm serves us as a basic building block in order to obtain our hierarchical snapshot schemes.

We suggest an approach for hierarchical snapshot based on an overlay network abstraction. We enable each subsystem to perform an independent snapshot, and further enable each level of the hierarchy to perform a local snapshot. We suggest the use of overlay communication links which "directly" connect leaders of clusters. Communication over an overlay link is much faster than the sum of the single hop communication links that implement the overlay link. Leaders of subsystems are defined, and the communication between processors in different subsystems traverses the overlay communication links between the leaders of the subsystems. When a snapshot is invoked by a leader of a subsystem, the leader uses the overlay network to notify the leaders of the subsystems that belong to its subsystem. These leaders, in turn, are responsible for performing a snapshot in their subsystem in the same manner.

Distributed wireless data gathering. We address the problem of data gathering in a wireless network [BKMSS08] using multihop communication; our main goal is the analysis of simple algorithms suitable for implementation in realistic scenarios. We study the performance of distributed algorithms, which use a simple contention based MAC, and we focus on the objective of minimizing average flow times of data packets. We prove a lower bound on the competitive ratio of any distributed algorithm minimizing the maximum flow time, polynomial in the number of packets. Next, we consider a distributed algorithm which sends packets over shortest paths, and we use resource augmentation to analyze its performance when the objective is to minimize the average flow time. We prove that if the algorithm sends packets faster than the optimal offline solution then its performance arbitrarily approaches the optimal offline solution. Data gathering has also been considered in [FSV08] where focusing on the assessment of environmental noise pollution in urban areas, we provide qualitative considerations and experimental results to show the feasibility of wireless sensor networks to be used in this context. To select the most suitable data collection protocol for the specific noise monitoring application scenario, we evaluated the energy consumption performances of the CTP (Collection Tree Protocol) and DMAC protocols. Our results show that CTP, if used enabling the LPL (Low Power Listening) option, provides the better performances trade-off for noise monitoring applications.

Relay Placement. In the relay placement problem the input is a set of sensors and a number $r \geq 1$, the communication range of a relay. In the *one-tier* version of the problem the objective is to place a minimum number of relays so that between every pair of sensors there is a path *through sensors and/or relays* such that the consecutive vertices of the path are within distance r if both vertices are relays and within distance 1 otherwise. The *two-tier* version adds the restrictions that the path must go *through relays, and not through sensors*. In [EFG$^+$08] we present a 3.11-

approximation algorithm for the one-tier version and a PTAS for the two-tier version. We also show that the one-tier version admits no PTAS, assuming P $\neq$ NP.

Spanning graph. In [AGFFI08] we evaluate the stretch of the (α,β)-spanners in a streaming model. An (α,β)-spanner of a graph G is a subgraph S of G such that for each pair of vertices the distance in S is at most alpha times the distance in G plus beta. The results of our experimental investigation on three spanning algorithms on several input families confirm that all these algorithms are very efficient in practice, finding spanners with stretch and size much smaller than the theoretical bounds and comparable to those obtainable by off-line algorithms. Moreover, our experimental findings confirm that small values of the stretch factor are the case of interest in practice, and that the algorithm by Ausiello et al. tends to produce spanners of better quality than the algorithms by Baswana and Elkin, while still using a comparable amount of time and space resources.

Sharing information among mobile nodes. Our first investigations regarding data management consist of a theoretical study of a power-aware on-line file allocation system for a mobile star network [BBKM08]. Our scenario consists of an immobile base station and a set of mobile nodes which are connected to the base station via a chain of relays which communicate over radio connections. The mobile nodes may read from and write to a shared database which is mainly located at the base station, but can also be replicated to the mobile nodes. The costs of accesses and of replication is defined by the power consumed by the involved mobile nodes and their chain of relays. It therefore primary depends on the number of involved relays. We give lower bounds and algorithms against an adversary which controls both the movement and the data accesses. Our main result is that against such an unrealistic powerful adversary an algorithm can only guarantee a competitive ratio linear in the size D of the shared file.

2.4 Research targets

Self-organization. In order to discover and implement the most efficient and fair pricing schemes, we believe that our approach for optimizing an unknown (global) objective, which was inspired by network congestion control and exploits the selfishness of the competing entities (with hidden incentives), is a quite promising venue that fits well in FRONTS, particularly as a model for dynamic adaptation of the underlying communication infrastructure in a large-scale network of tiny artifacts. Indeed, our approach is both completely dynamic and decentralized.

Sharing information among mobile nodes. The major challenge regarding file allocation in mobile ad-hoc networks is the maintenance of a subnetwork of the nodes holding a copy under movement. On the one hand, this subnetwork should be as short as possible since every update of the data has to be propagated over it. On the

other hand, the structure of the subnetwork should not change too often since every change incurs additional costs. A second challenge is the definition of realistic stochastic/adversarial movement and data access models. While the movement and data accesses are not independent (nearby nodes are more likely to access the same data), they surely also do not act adversely. Thus, a realistic compromise has to be found.

In order to better distribute resources in sensor networks which experience localized peeks of communication demands, we plan to find distributed algorithms that identify critical cuts in a network with given boundaries. Among the set of all cuts that split the network into substantial regions (e.g., regions containing at least a certain fraction of all nodes), we seek those of minimal length, energy, or communication bandwidth. These should then serve as a base for communication schemes avoiding critical cuts when possible. In a later step, we plan to extend the scheme to use a hierarchy of cuts with different levels of criticality.

In the context of building graph spanners in a streaming model a future work is aimed at considering the case in which the amount of available memory is sublinear w.r.t. the number of vertices in the graph. In this case interesting results may be achieved in the so called stream-sort model.

Experimental Validation. As a future work we plan to demonstrate the qualities of the new randomized self-stabilizing distributed algorithm for cluster definition in communication graphs, and the implementation of the new methods is by "Shawn" (sensor network simulator).

3 Task 2.2: Adapting the internal structure and roles of artefacts

3.1 Motivation and problem description

Adaptiveness of a network can be perceived as its ability to function "properly" when the external conditions change. We believe that social paradigms, especially distributed ways in game theory to self-reach equilibrium, can help in designing adaptive networks. In this task, we study social paradigms for systems that are required to: (1) adapt to external changes (such as failures) while (2) assuming that every tiny artefact has limited resources of computation.

Given a property that an adaptive network wants to maintain, we could properly define a strategic game involving the nodes/edges of the graph in order to analyze a local search method for reaching a feasible solution and acting accordingly. The local search should not be affected by changes in the environment, such as faults in connectivity and edges. A desirable property of the defined game is the existence of a potential function that considers the changing environment. This is because

the potential function assures that, albeit the initial configuration, the tiny artefacts are guaranteed to reach a pure Nash equilibrium in finite time. So even if there are node/link failures, the players will still converge to a pure Nash equilibrium; the question is: How fast will they converge?

The above mentioned restriction on computational resources of tiny artefacts poses the question of how intelligent tiny artefacts should behave, given that each of them cannot perform unbounded computations. For example, each tiny artefact in a dynamically evolving network may want to achieve a selfish goal, such as the prolongation its own life cycle. However, due to computational limitations, the tiny artefacts are in no position to compute what's truly best for them. Nevertheless, their own behavior still is strategic, at least with respect to what they can actually achieve using *bounded rationality*.

Existing solution concepts model the system as a infinitely repeated game among autonomous (selfish) agents that a subset of them can deviate. One may assume that the actions of all tiny artefacts are completely observable and demonstrate an asymptotically optimal strategy for deterring the deviators from deviating.

One specific research goal in this context is to study selfish-artefacts that are subject to *transient faults* due to hardware/software temporal malfunctions; just as the society is a subject to human mistakes due to a moment of weakness. Game theory uses punishment for deterring improper behavior. 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".

We are interested in providing concepts for building a system that can automatically recover from an arbitrary state including even one in which a Byzantine execution of one or more programs repeatedly attempts to corrupt the system state. Special emphasis is given to the scope of operating systems and for self-stabilizing recovering from the actions of malicious (Byzantine) programs.

Fundamental problems related to the management of continuous data stream arise from the measurements of the sensors. Suppose for example that a sensor measures temperature needs to keep the maximum temperature within the last hour, i.e., we are interested in the maximum of the measurements that are within a sliding window with hour length. Here adaptivity is required in the sense that as time passes some of the old data has to be ignored, and some has to be aggregated together with the new data flow, in order to update the statistical values computed by the system. We study methods for achieving that using irrefutably limited memory, e.g., just a single memory position.

We also look into ways for using the power of stigmergy in supporting fully decentralized recommendation techniques in pervasive systems of small devices with

limited or no communication and computational capabilities. Stigmergy is an organizing principle in which individual parts of the system communicate with one another indirectly by modifying their local environment. As a form of self-organization, stigmergy produces complex, apparently intelligent structures, without need for any planning, control, or even direct communication between the agents.

3.2 State of the art and previous work

Bounded rationality and system of selfish artefacts. The traditional approach used in game theory in order to capture limitations of information among the agents is via Bayesian games. Here, however, we wish to capture not informational, but actually operational limitations of rationality. Modelling bounded rationality is a very crucial, ongoing discussion in the related literature. The main approaches consider the agents as either finite automata, or Turing Machines with limitations in the usage of memory (e.g., [Ney85, Rub86, PY94]).
Systems of selfish-artefacts, introduce new challenges in distributed computing, game theory, and computational complexity. They exhibit both cooperative and uncooperative interactions. While cooperative and uncooperative interactions have been extensively studied as the two extremes, the study of joint deviations in uncooperative repeated games has been neglected so far. In systems of selfish-artefacts (where out-of-band communication is possible), it is unlikely that selfish-artefacts cannot conspire.
Adaptive statistical processing of the sensor measurements. Data streams have been extensively studied (see [Mut05] for an overview) but our main novelty here is that we study a data stream problem using competitive analysis, i.e., from the viewpoint of online algorithms.
Classification of machine learning algorithms that are designed to deal with corrupted (or malicious) data are of great interest since any realistic model of learning from examples must address the issue of corrupted data. Previous work, tried to deal with this issue by developing new algorithms by boosting of supervised learning or other robust and efficient learning algorithms e.g., [Ser03]. Despite these difficulties, the importance of being able to cope with noisy data has led many researchers to study PAC learning in the presence of noise.

3.3 First results

Selfish local search. In [PS08], we deal with the problem where the network wants to maintain a proper coloring of its nodes. We view each node as an agent that wishes to maximize an individual payoff. The payoff is the number of nodes with the same color, unless the color is the same with a current neighbor (in which case the payoff is

zero). We envision then an eternal procedure in which each node "plays" in turn (in some order, even random). The node that plays may change its color in order to be in a color class of greater cardinality, avoiding conflicts with neighbors. We show that there exists a potential function assuring that this local search procedure terminates in polynomial time, and that the number of colors in a stable situation (where no node can further increase its payoff) satisfies all the known bounds of the chromatic number of the current graph in a stable situation.

Strategies for repeated games with subsystem takeovers - implementable by deterministic and self-stabilizing Automata. Stability and self-enforcement are two of the most attractive properties that equilibria offer. We consider equilibrium of strategies that autonomous agents have devised, and all possible joint deviations by a group of at most D deviators [SSDT08]. Stability and self-enforcement are achieved when the autonomous agents deter the deviation; if any one of all possible joint deviations happens, then the deviating group will be punished and will lose payoff, compared to what they would get by obeying the equilibrium strategy.

Purifying data by machine learning with certainty levels. We investigate two new approaches to increase the certainty levels of machine learning results by computing a certainty level that takes into account the Byzantine data items in the training data-set file [DLY08]. The first scheme is based on identifying statistical parameters when the distribution is known (e.g., normal distribution), and the second uses decision trees similar to the random forest techniques, incorporating certainty level in the leaves.

Stabilizing trust and reputation for self-stabilizing efficient hosts in spite of Byzantine guests. Detection of Byzantine behavior of a guest during run time (namely, sanity checks detect a contract violation) can not prevent the guest from being executed, since the Byzantine behavior might be caused by a transient fault. In [DY07] we investigate the need for combining self-stabilization, and techniques for enforcing a contract over the operations of a guest. We show that only such a combination will allow recovery.

Stigmergy: self-adaptive recommendation systems. In [BCMSV08b, BCMSV08a] we investigate the power of stigmergy in supporting fully decentralized recommendation techniques for pervasive systems of small devices with limited or no communication and computational capabilities. We consider a reference scenario where items are distributed in a shop and are tagged with RFIDs. We assume that every node has only partial view of the past history of the system, determined by users' visit patterns; information exchange among nodes is mediated by the collective and unpredictable navigation of users. These limitations are coherent with standard systems and current technology. In [CPV09] we started an experimental activity exploiting similar concepts on real RFID hardware to enable touch-based remote grocery shopping services.

3.4 Research targets

Bounded rationality and system of selfish artefacts. We wish to investigate models of bounded rationality, by studying competitive models in which additional costs (apart from the agents' preferences) are imposed to each agent. For example, we may think of an agent as a finite state automaton, or a Turing Machine, whose actual utility function is the sum of its own preference value (for each possible state) plus the computational costs that are incurred by this agent to compute its own strategy (e.g., the size of the automaton used).
We wish also to follow the approach of selfish local search in other problems related to networks of tiny artefacts. For each problem, we should properly define a game for which (a) there exists a potential function; (b) the pure Nash equilibria of the game corresponds to feasible solutions of the given problem; (c) a pure Nash equilibrium can be computed fast via selfish local moves; (d) the cost of the solution computed by the algorithm is as close as possible to the optimum.

As examples of such problems, we can mention *clustering* (i.e., what is a good way in which the nodes of the network organize themselves into clusters in order to perform a given task) and *power allocation* (i.e., what is a good power allocation scheme that guarantees some connectivity properties for the network).
Adaptive statistical processing of the sensor measurements-competitive analysis of aggregate max. Other directions we propose for future research are to analyze algorithms for computing other statistical values of the stream of measurements as for example maintaining the k top ranked elements in each window under the assumption that the items appearing in the stream are chosen by an adversary, but then presented to the algorithm in a random order. The famous secretary problem (see for example [Kle05, BIK07]) uses the same assumption. Perhaps the most interesting, but also difficult, direction would be to analyze algorithms where given a specific network structure, many sensors need to communicate to compute a parameter of the system. In such a setting we would not only want to approximate the unknown parameter as good as possible, but we also want to minimize the communication cost.
Stigmergy. We plan to extend our study on recommendation systems in three main directions: (1) improve the model describing users' visit patterns, (2) provide inter-cluster recommendations, (3) better capture some important aspects of real hardware, such as limited memory availability. Furthermore we are evaluating how to apply similar concepts to NFC (Near Field Communication) devices (e.g. smart posters).
Experimental Validation The research on Strategies for repeated games, Relocation analysis of adaptive protocols, Learning with certainty levels, Trust and reputation for self-stabilizing will continue next year also by simulations either using Shawn or ad-hoc simulators.

4 Task 2.3: Security for tiny artefacts

4.1 Motivation and problem description

Security for small artefacts is conceptually different from the security in other domains (like Internet and Network security). Indeed, small 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 suggests as a possibility Cryptography based on Pairings (here a 128 bit moduli guarantees security comparable with 1024-bit RSA modulus). In addition small 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 small artefacts). It is also crucial to study privacy problems associated with the use of small artefacts to trace and locate objects which could threat the privacy of the individuals that are associated with the objects being traced.

As wireless sensor network technology evolves, they will start to play important roles in various applications; in most real-life scenarios, the integrity and confidentiality of the information reported to the controlling authorities is of paramount importance. Although these are more or less standard security requirements that can also be found in traditional wired and wireless networks, WSNs have a different purpose, in the sense that communications will not involve human interaction [PLH06]. The challenge is to satisfy these requirements under the special operating conditions of sensor networks and therefore the approaches for offering protection need to be reconsidered [PSW04]. One important issue when designing Wireless Sensor Network (WSN) is the routing protocol that makes the best use of the severely limited resource presented by WSN, especially the energy limitation. Another import factor required attention from researchers is providing as much security to the application as possible. Here we wish to look into the problem of secure data propagation in wireless sensor networks.

One of the key problems for pervasive systems is to provide authentication methods that would be strong enough to serve their purpose and, at the same time, that would be efficient enough in terms of communication volume and computational complexity of required operations. These operations should be implementable on low-end devices, such as passive cheap RFID devices. Conventional cryptography based on algebraic structures is out of scope here due to resource consumption. Last not least, authentication methods should be privacy aware, which is a challenge in case of passive devices. It turns out that it is possible to design authentication methods based on combinatorial properties rather than hard algebraic problems.

Enhancing security level of single artefacts is limited not only by technical issues.

For instance, we are rarely able to control physical access to the artefacts or to protect the whole production and distribution chain. So, *a single device cannot be blindly trusted*, even if the cryptographic mechanisms are well designed: we are never sure that a device is not a fake one and that it does not have additional malicious functionalities. One of the methods to cope with this problem is to abandon the single-device paradigm.

Small artefacts can be physically captured by an adversary and, it is well known, that by applying strong magnetic fields or by disconnecting the power source of small artefacts, it can be reset to an initial state (*reset attack*). As a consequence, the same randomness may be used more than once to perform a protocol and this is a serious security threat as, typically, all security proofs assume that independent randomness is used for each execution of the protocol.

Furthermore tiny artefacts like passive RFID tags, can be embedded in consumer products in order to facilitate post-sale management services (e.g., handling returns of unwanted items). In this case the tags must contain some information that is private to the manufacturer and thus is encrypted. Ciphertexts do not reveal any information on the cleartext but can none-the-less be traced by an adversary. In case the product with tags is worn (for example, it is a designer jacket) or tightly associated with a user tracing the tag will actually allow tracing the owner of the product thus threatening his/her privacy.

4.2 State of the art and previous work

Combinatorial cryptography for low-end devices. Here we main focus on Privacy in RFID-based systems. An approach presented in the literature to combat the possibility of tracing RFID tags consists in designing special types of encryption schemes that are re-randomizable (see [ACdM05]). Specifically, given a ciphertext Ct for public key Pk associated with cleartext M, it is possible to produce a new valid ciphertext Ct' for the the same public key Pk and cleartext M without having to decrypt. Therefore, private information can be stored on a RFID tags using randomizable encryption and the environment will have randomizers that re-randomize encryptions.

In [Vau07], Vaudenay presented a security and privacy model for RFID that subsumes previous notions of privacy and produces a much more understandable framework. Unfortunately the proposed model still has some limitations (it considers one reader only, does not cover RFID schemes that work in presence of anonymizers, does not allow tamper proof areas in tags).

Reset attacks. The possibility of a reset attack has been first identified in [CGGM00] where it was studied in the context of authentication protocols. The various proposal that can be found in the literature cannot be applied in the context of small artefacts since they are either extremely expensive in terms of computation [PRS02], or re-

quire a public-key infrastructure [DCPV04] or only apply to a restricted type of reset attacks [BFGM01].

Secure routing. There exist many different Secure routing protocols for ad-hoc networks based on symmetric key cryptography [BHBR01, HPJ05]. These protocols are based on source routing or distance vector protocols and are unsuitable for sensor networks. They are too expensive in terms of node state and packet overhead and are designed to find and establish routes between any pair of nodes – a mode of communication not prevalent in sensor networks. For a survey of existing work in securing ad hoc network routing see [HP04]. In [MGLB00] the problem of minimizing the effect of misbehaving or selfish nodes on routing through punishment, reporting, and holding grudges is investigated. The application of this technique to sensor networks is promising, but these protocols are vulnerable to blackmailers. In [PST+02] two building block security protocols optimized for use in sensor networks are presented: SNEP and μTESLA. In [Dim06] a protocol is introduced that establishes a communication tree, the protocol requires a short period of time during the initialization during which no attack is expected.

4.3 First results

Combinatorial cryptography for low-end devices. Within the project we have proposed (see in [CKK08]) a *hidden subset* method that enables reliable authentication with privacy protection, low hardware implementation cost, and low communication volume (which is the main disadvantage of HB and HB+). We have designed (details in [GMZ08]) and implemented attack methods for the early versions of hidden subsets protocol (and the former HB+ protocol, as well). The attacks, of exponential time complexity, turn out to be effective for quite reasonable size of parameters that could occur in practice. Since the size of parameters cannot be much increased due to communication issues, these kinds of attacks should be examined very carefully.

Secure routing. Routing is based on probability rather than on communication links since they cannot be relied upon because of mobility. This probabilistic approach also simplifies the way in which the path towards the sink is determined and distributes energy consumption amongst the network. Furthermore, the way in which the network is organized in concentric circles and sectors around the sink provides a fertile environment to introduce secure methods of establishment and communication. In [BCLS08] we proposed a solution for the Secure Data Propagation Problem on Mobile WSNs. In this direction we introduced a phase where the networks routing and security infrastructure is established, a method in order for the sensors to securely route sensing information towards the sink and we proposed actions that can be taken in order for our network to adapt to mobility and security attacks.

Reset attacks. In [BPSV08], we consider identification protocols that are re-

settable non-transferable as discussed in [BFGM01]. This notion aims at capturing security with respect to powerful adversaries that have physical access to the device that proves its identity, and thus can potentially reset its internal state. We discuss some limitations of existing notions for secure identification protocols as well as different impossibility results for strong notions of non-transferability. We introduce a new strong and achievable notion for resettable non-transferable identification that reflects real scenarios more adequately and present a generic protocol that satisfies this notion. We then show how to efficiently instantiate our construction and discuss how our protocol can improve the current proposals for protocols for proving personal identity.

Privacy issues. In [SVW08], we look at privacy issues in the context of small artefacts used for authentication in transport systems based on RFID e-tickets. We derive appropriate security and privacy requirements and propose an implementation using known cryptographic techniques.

4.4 Research targets

Combinatorial cryptography for low-end devices. Current research challenges concern fine tuning of the algorithm from the point of view of hardware implementation and threats emerging on this level, evaluation against the attack methods, as well as designing hardware architecture. We plan to expand the protocol in order to implement functionalities that are crucial for security in the considered environment. The first challenge is a two-way authentication. For instance, it is extremely important to authenticate not only the RFID device against the reader, but also the reader against the RFID device. The protocols developed so far have been focused only on one-way authentication. Another idea is to enable execution of the authentication protocol only in presence of special units that enable or disable communication with cryptographic means. This is particularly important for the case when the devices are carrying sensitive data (e.g. e-ID cards, medical cards).

Another stream of research here is to design relatively weak methods tailored for standard (but very weak) processors used by many embedded systems. Use of the simplest solutions is here not only the matter of price but also fault probability.

Building blocks for dynamic key predistribution. The general scheme of dynamic key predistribution requires building cryptographic functions with specific properties - the final performance parameters are crucial for the practical usability of the scheme in the world of tiny devices. This concerns in particular deterministic trap-door one way functions on relatively short sequences, easy to implement, but not necessarily of the same strength as the functions used in standard applications. We will also pursue design of new functions that would enable new mechanism for administration in the emerging trust system (joining and forging trust groups, delegation

of trust, . . .).

Multiparty security. The main stream of research in this field is to appropriate design the cooperation between low-end devices and pretty strong units. The computationally stronger units used as proxies can protect against system flaws due to for example weak (or not fault resistant) pseudorandom number generator of the devices. As a side product, this can reduce the overall price of the system, since we can shift the expensive operations to the places where they are executed more efficiently at lower cost. The big challenge is to build such a shared infrastructure with relatively high level of security. However, there are also some advantages due to locality of communication and impossibility to monitor all communication.

Secure routing. We wish to shift our focus from networks where communication is not necessarily directed towards a single sink. In such networks the existence of multiple sinks may represent actors (i.e. devices embedded on appliances) where messages need to be delivered. Usually in models of multiple sinks it is assumed that all sinks are connected via a fast link. We are also interested in cases where the number of sinks changes through out the execution of the protocol and sensors and the sink can move during the execution of remote monitoring applications.

Privacy for RFID-based systems. We plan to extend the frameworks for authentication of RFID devices that overcome the limitations of the previous models so that they can be applied to many real-life applications where the RFID tags are used. We also plan to extend the work of [ACdM05] on re-randomizable encryption for RFID to have more efficient constructions.

5 Task 2.4: Unification of results

The unification of the results is described analytically in deliverable D4.3 “First Report on Unification of Results”.

In the list of references below, papers marked with "$*$" present results of the FRONTS project.

References

[ACdM05] Giuseppe Ateniese, Jan Camenisch, and Breno de Medeiros. Untraceable rfid tags via insubvertible encryption. In Vijay Atluri, Catherine Meadows, and Ari Juels, editors, *ACM Conference on Computer and Communications Security*, pages 92–101. ACM, 2005.

[AD97] Y. Afek and S. Dolev. Local stabilizer. *Journal of Parallel and Distributed Computing, special issue on self-stabilizing distributed systems*, 62(5):745–765, 1997.

[ADGR05] E. Anceaume, X. Defago, M. Gradinariu, and M. Roy. Towards a theory of self-organization. *9th International Conference on Principels of Distributed Systems, OPODIS*, pages 146–156, 2005.

$*$ [AGFFI08] Giorgio Ausiello, Paolo G. Franciosa, and Giuseppe F. Italiano. Small stretch (α, β)-spanners in the streaming model. *Theoretical Computer Science, Special Issue in Honour of Burkhard Monien - to appear*, 2008. FRONTS-TR-2008-56.

[Bar96] Yair Bartal. Distributed paging. In *Online Algorithms*, volume 1442 of *Lecture Notes in Computer Science*, pages 97–117. Springer, 1996.

[BBKM08] Marcin Bienkowski, Jaroslaw Byrka, Miroslaw Korzeniowski, and Friedhelm Meyer auf der Heide. Optimal algorithms for page migration in dynamic networks. *Journal of Discrete Algorithms*, 2008. Accepted for publication.

[BCLS08] Kostantinos Bairaktaris, Ioannis Chatzigiannakis, Vasiliki Liagkou,
$*$ and Paul Spirakis. Adaptive probabilistic secure routing in mobile wireless sensor networks. In *16th International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2008)*. IEEE, September 2008. FRONTS-TR-2008-9.

[BCMSV08a] Luca Becchetti, Ugo Colesanti, Alberto Marchetti Spaccamela, and
$*$ Andrea Vitaletti. Fully decentralized recommendation systems: Models and analysis. Technical report, UDRLS, 2008. FRONTS-TR-2008-39.

[BCMSV08b] Luca. Becchetti, Ugo. Colesanti, Alberto. Marchetti Spaccamela, and
$*$ Andrea. Vitaletti. Fully decentralized recommendations in perva-

sive systems: Models and experimental analysis. Technical report, UDRLS, 2008. FRONTS-TR-2008-42.

[BFGM01] Mihir Bellare, Marc Fischlin, Shafi Goldwasser, and Silvio Micali. Identification protocols secure against reset attacks. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 495–511. Springer, 2001.

[BHBR01] Stefano Basagni, Kris Herrin, Danilo Bruschi, and Emilia Rosti. Secure pebblenets. In *2nd ACM international symposium on Mobile ad hoc networking & computing MobiHoc*, pages 156–163, New York, NY, USA, 2001. ACM.

[BIK07] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA*, pages 434–443, 2007.

* [BKMSS08] Vincenzo Bonifaci, Peter Korteweg, Alberto Marchetti-Spaccamela, and Leen Stougie. The distributed wireless gathering problem. *AAIM*, pages 72–83, 2008. FRONTS-TR-2008-41.

* [BPSV08] Carlo Blundo, Giuseppe Persiano, Ahmad-Reza Sadeghi, and Ivan Visconti. Improved security notions and protocols for non-transferable identification. In Sushil Jajodia and Javier López, editors, *Computer Security - ESORICS 2008, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings*, volume 5283 of *Lecture Notes in Computer Science*, pages 364–378. Springer, 2008. FRONTS-TR-2008-17.

[CDPV03] A. Cournier, A. Datta, F. Petit, and V. Villain. Enabling snap-stabilization. *Proc. of the 23rd International Conference on Distributed Computing Systems*, pages 12–19, 2003.

[CGGM00] Ran Canetti, Oded Goldreich, Shafi Goldwasser, and Silvio Micali. Resettable zero-knowledge. In *STOC*, pages 235–244, 2000.

* [CKK08] Jacek Cichoń, Marek Klonowski, and Mirosław Kutyłowski. Privacy protection for RFID with hidden subset identifiers. In Jadwiga Indulska, Donald J. Patterson, Tom Rodden, and Max Ott, editors, *Pervasive Computing, 6th International Conference, Pervasive 2008, Sydney, Australia, May 19-22, 2008, Proceedings*, volume 5013 of

Lecture Notes in Computer Science, pages 298–314. Springer, 2008. FRONTS-TR-2008-2.

* [CPV09] I. Cappiello, S. Puglia, and A. Vitaletti. Design and initial evaluation of a ubiquitous touch-based remote grocery shopping process. Technical report, UDRLS, 2009. FRONTS-TR-2008-41.

[DCPV04] Giovanni Di Crescenzo, Giuseppe Persiano, and Ivan Visconti. Constant-round resettable zero knowledge with concurrent soundness in the bare public-key model. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 237–253. Springer, 2004.

[Dim06] T. Dimitriou. Securing communication trees in sensor networks. In *2nd International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS 2006)*, Lecture Notes in Computer Science. Springer-Verlag, 2006.

* [DLY08] Shlomi Dolev, Guy Leshem, and Reuven Yagel. Purifying data by machine learning with certainty levels. Technical report, Ben-Gurion University of the Negev, Beer-Sheva, Israel., 2008. FRONTS-TR-2008-50.

* [DT08] Shlomi Dolev and Nir Tzachar. Empire of colonies: Self-stabilizing and self-organizing distributed algorithm. *Theoretical Computer Science*, In Press, Corrected Proof:–, 2008. FRONTS-TR-2008-22.

* [DY07] S. Dolev and R. Yagel. Stabilizing trustand reputation for self-stabilizing efficient. In *Stabilization, Safety, and Security of Distributed Systems, 9th International Symposium, (SSS)*, Paris, France,, November 14-16 2007. Springer Berlin / Heidelberg. FRONTS-TR-2007-1.

[EFG+08] Alon Efrat, Sándor P. Fekete, Poornananda R. Gaddehosur, Joseph
* S. B. Mitchell, Valentin Polishchuk, and Jukka Suomela. Improved approximation algorithms for relay placement. In *Algorithms – ESA 2008*, pages 356–367, 2008. FRONTS-TR-2008-38.

* [FSV08] L. Filipponi, S. Santini, and A. Vitaletti. Data collection in wireless sensor networks for noise pollution monitoring. In *DCOSS*, volume 5067 of *Lecture Notes in Computer Science*. Springer, 2008. FRONTS-TR-2008-23.

* [GMZ08] Zbigniew Golebiewski, Krzysztof Majcher, and Filip Zagórski. Attacks on CKK family of RFID authentication protocols. In David Coudert, David Simplot-Ryl, and Ivan Stojmenovic, editors, *Ad-hoc, Mobile and Wireless Networks, 7th International Conference, ADHOC-NOW 2008, Sophia-Antipolis, France, September 10-12, 2008, Proceedings*, volume 5198 of *Lecture Notes in Computer Science*, pages 241–250. Springer, 2008. FRONTS-TR-2008-30.

[HP04] Yih-Chun Hu and Adrian Perrig. A survey of secure wireless ad hoc routing. *IEEE Security and Privacy*, 2(3):28–39, 2004.

[HPJ05] Yih-Chun Hu, Adrian Perrig, and David B. Johnson. Ariadne: a secure on-demand routing protocol for ad hoc networks. *Wireless Networks*, 11(1-2):21–38, 2005.

[Kle05] Robert D. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, pages 630–631, 2005.

[KP90] S. Katz and K. Perry. Self-stabilizing extensions for message-passing systems. *Proceedings of the ninth annual ACM symposium on Principles of distributed computing*, pages 91–101, 1990.

[MGLB00] Sergio Marti, T. J. Giuli, Kevin Lai, and Mary Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *6th annual international conference on Mobile computing and networking MOBICOM*, pages 255–265, New York, NY, USA, 2000. ACM.

[Mut05] S. Muthukrishnan. Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science*, 1(2), 2005.

[MW05] T. Moscibroda and R. Wattenhofer. Maximal independent sets in radio networks. In *PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, pages 148–157, New York, NY, USA, 2005. ACM.

[Ney85] Abraham Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Economics Letters*, 19(3):227–229, 1985.

[PLH06] A.S.K. Pathan, H-W. Lee, and C.S. Hong. Security in wireless sensor networks: issues and challenges. In *8th International Conference on Advanced Communication Technology (ICACT06)*, volume 2, 2006.

[PRS02] Manoj Prabhakaran, Alon Rosen, and Amit Sahai. Concurrent zero knowledge with logarithmic round-complexity. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November*

2002, Vancouver, BC, Canada, Proceedings, pages 366–375. IEEE Computer Society, 2002.

* [PS08] Panagiota Panagopoulou and Paul Spirakis. A game theoretic approach for efficient graph coloring. In *The 19th International Symposium on Algorithms and Computation (ISAAC 2008)*, December 2008. FRONTS-TR-2008-43.

[PST+02] Adrian Perrig, Robert Szewczyk, J. D. Tygar, Victor Wen, and David E. Culler. Spins: security protocols for sensor networks. *Wireless Networks*, 8(5):521–534, 2002.

[PSW04] A. Perrig, J. Stankovic, and D. Wagner. Security in wireless sensor networks. *Communications of the ACM*, 47(6), 2004.

[PY94] Christos H. Papadimitriou and Mihalis Yannakakis. On complexity as bounded rationality (extended abstract). In *STOC '94: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 726–733, New York, NY, USA, 1994. ACM.

[Rub86] Ariel Rubinstein. Finite automata play the repeated prisoners dilemma. *Journal of Economic Theory*, 39:176–188, 1986.

[Ser03] A. R. Servedio. Smooth boosting and learning with malicious noise. *The Journal of Machine Learning Research*, 4:633–648, 2003.

* [SSDT08] Elad Michael Schiller, Paul Spirakis, Shlomi Dolev, and Philippas Tsigas. Strategies for repeated games with subsystem takeovers implementable by deterministic and self-stabilizing automata. In *MiNEMA Workshop 2008*. Middleware for Network Eccentric and Mobile Applications, September 2008. FRONTS-TR-2008-16.

* [SVW08] Ahmad-Reza Sadeghi, Ivan Visconti, and Christian Wachsmann. User privacy in transport systems based on rfid e-tickets. In *Proceedings of Workshop on Privacy in Location-Based Applications (PILBA 2008)*, 2008. FRONTS-TR-2008-20.

[Var00] G. Varghese. Self-stabilization by counter flushing. *SIAM Journal on Computing*, 30(2):486–510, 2000.

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

[ZA02] H. Zhang and A. Arora. Gs3: Scalable self-configuration and self-healing in wireless networks. *Symposium on Principles of Distributed Computing*, pages 58–67, 2002.