Integrated project proposal
FET Proactive
FP6-2004-IST-FET Proactive

Algorithmic Principles for Building Efficient Overlay Computers

AEOLUS

Date of preparation: September 22nd, 2004

Participant no.	Participant name	Participant short name
1	University of Patras (EL)	UOP
2	Telecom Italia Learning Services S.p.A. (I)	TILS
3	Centre National de la Recherche Scientifique (F)	CNRS
4	University of Paderborn (D)	UPB
5	Computer Technology Institute (EL)	CTI
6	Università degli Studi di Salerno (I)	UNISA
7	University of Ioannina (EL)	UOI
8	Centre Universitaire d' Informatique (CH)	CUI
9	Max-Planck Institut für Informatik (D)	MPII
10	Christian-Albrechts-Universität zu Kiel (D)	CAU
11	Università degli Studi di Roma "Tor Vergata" (I)	UDRTV
12	University of Athens (EL)	UOA
13	Università degli Studi di Padova (I)	UNIPD
14	Eidgenössische Technische Hochschule Zürich (CH)	ETHZ
15	Universitat Politècnica de Catalunya (E)	UPC
16	Università degli Studi di Roma "La Sapienza" (I)	UDRLS
17	Katholieke Universiteit Leuven (B)	KUL
18	Institut National De Recherche en Informatique et en Automatique (F)	INRIA
19	DIMATIA, Charles University (CZ)	DIM
20	University of Cyprus (CY)	UCY
21	Cybernetica (EE)	CYB

Name of the coordinating person: Christos Kaklamanis
Coordinator organisation name: University of Patras
e-mail: kakl@ceid.upatras.gr
fax: +30-2610-960479

Table of contents

Proposal summary page

Full Title: Algorithmic Principles for Building Efficient Overlay Computers
Proposal acronym: AEOLUS

Strategic objective(s) addressed

2.3.4.2 (v) Global computing

Proposal abstract

The recent explosive growth of the Internet gives rise to the possibility of a global computer of grand-scale consisting of Internet-connected computing entities (possibly mobile, with varying computational capabilities, connected among them with different communication media), globally available and able to provide to its users a rich menu of high-level integrated services that make use of the aggregated computational power, storage space, and information resources. Achieving this efficiently and transparently is a major challenge that can be overcome by introducing an intermediate layer, the overlay computer.

The goal of this project is to investigate the principles and develop the algorithmic methods for building such an overlay computer that enables this efficient and transparent access to the resources of an Internet-based global computer.

In particular, the main objectives of this project are:

- *To identify and study the important fundamental problems and investigate the corresponding algorithmic principles related to overlay computers running on global computers.*
- *To identify the important functionalities such an overlay computer should provide as tools to the programmer, and to develop, rigorously analyze and experimentally validate algorithmic methods that can make these functionalities efficient, scalable, fault-tolerant, and transparent to heterogeneity.*
- *To provide improved methods for communication and computing among wireless and possibly mobile nodes so that they can transparently become part of larger Internet-based overlay computers.*
- *To implement a set of functionalities, integrate them under a common software platform in order to provide the basic primitives of an overlay computer, as well as build sample services on this overlay computer, thus providing a proof-of-concept for our theoretical results.*

B.1 Scientific and technological objectives of the project and state of the art

The social impact of the Internet has gone far beyond the expectations of only a few years ago. Technological advances made hardware devices like personal computers and wireless devices available to a vast majority of the population and, as a result, most people in the developed countries have access to this medium.

One of the reasons for its success is due to the fact that currently the Internet constitutes the widest *data repository* available worldwide. Indeed, historically, the Web was the first means through which people approached the Internet. A lot of research has been done in order to simplify the mining of this data (e.g. search engines) or protecting them from unauthorized access. Furthermore, from a different point of view, aggregating the *computational power* even of a small fraction of the devices connected to the Internet would result in a computing "device" whose capabilities could greatly expand the concept of "computationally tractable".

These goals pose major challenges that can be addressed in two different research directions. Due to the inherent complexity of the system, an optimized usage of its resources is practically impossible; a first research direction aims to improve the usage of resources on a global scale; this goal is usually achieved at the cost of increasing the complexity of the applications. On the other hand, by its own nature, the system is composed of heterogeneous devices with diverse characteristics and, sometimes, complementary limitations; in this scenario, the major challenge has been, and still is, to allow a complete interoperability among such devices.

Notwithstanding the problems related to inefficient resource usage and interoperability, the increasing number of distributed applications currently developed and deployed on the Internet witness on one hand the success of such technology and, at the same time, society's increasing need for a global information system providing reliable services that can be accessed by diverse means.

In such a scenario, what is really missing is *transparency* for the user. Indeed, currently there exists a variety of service providers, each of which provides its own services that, most of the time, require a sort of ad-hoc proprietary software. The access to each service is granted using different authorization mechanisms of which the user has to be aware, while the rules for accessing resources that may vary.

The cause of these limitations is the lack of a *unique infrastructure* on which different entities may *easily develop* and *deploy* software that *efficiently* exploits the tremendous storage capacity, computational power, and reliability of large scale distributed and heterogeneous systems that, at the same time, guarantees *privacy* of the information and *fairness* of the resource usage.

In order to obtain real impact from this new global instrument, a key requirement is *transparency for system designers and developers*. Currently, system implementations have to take into account all the issues related to the distributed nature of information and computation like resource discovery, fault tolerance, load balancing, security and so on. All such problems must be tackled in a heterogeneous system where nodes are owned by different entities that try to optimize their own objectives and possibly with malicious users. From this very brief description, it is clear that developing a successful distributed application that efficiently exploits the capabilities of the Internet is an extremely difficult task.

On the other hand, the existence of a unique programmable platform virtually overlaid on top of the existing heterogeneous system, providing a set of functionalities that hide from the programmer details related to subtle and crucial issues like scalability and efficiency of resource management as well as automatically and efficiently solve problems like data replication or recovery would have a great and immediate impact on the way in which distributed applications are developed. System designers and programmers will be able to concentrate only on the application, while leaving to the platform the duty of controlling the underlying global “device”. The possibility of securing sensitive data while allowing transparent access to them, according to per-site policies, would foster the spread of a global information system with tremendous effect on everyday life. In other words, applications will be transparently accessing a unique *data repository* and *service provider,* regardless of the actual device they are running on.

In the long term, we envision the unique infrastructure to be an *overlay computer*, based on a global computer of grand scale consisting of Internet-connected computers (possibly mobile, with varying computational capabilities, connected among them with different communication media including wireless, owned by different entities optimizing their own objectives, with diverse and even limited availability, etc.), globally available and able to provide to its users a rich menu of high-level integrated services that make transparent and efficient use of the aggregated computational power, storage space, and information resources. Based on the growth of the number of nodes of the Internet, the improvements of computational/communication capabilities of wireless devices and the spread of use of such devices, it is very likely that a global computer of this kind will be the main data repository/computational infrastructure that will be widely used in a few years.

For the purposes of this project, we consider the global computer to of composed by Internet-connected computing entities (small/medium computers, i.e., PCs, laptops). Possibly, this global computer can be extended through wireless communication (e.g., WiFi, Bluetooth, etc.) to include other computing devices. These devices will mainly be of similar computing power (i.e., mobile laptops) but potentially also include others like satellites, palmtops, mobile phones, and smaller peripherals like sensing devices.

Overlay computers are logical definitions over this global computer and can be subsets of nodes of the global computer cooperating for providing a service or running a distributed application. So, the overlay computer is an abstraction of the global computer on top of which services and applications can be implemented using basic primitives (implementations of functionalities) that the overlay computer provides as ``tools to the programmer''. These primitives handle and hide the actual characteristics of the global computer (e.g., topology, computational power of nodes, communication media, etc.).

The first step toward fulfilling such an ambitious goal is a deep understanding of the features that the basic functionalities of the overlay computer should have. A first, compelling characteristic is related to the efficiency of such primitives. In any worldwide-deployed architecture with thousands of heterogeneous nodes, the amplification of the effect of an inefficient solution could lead the whole system to collapse, especially if the inefficiency is related to a “basic” functionality, e.g., scheduling or routing. An algorithmic approach to the design of such an infrastructure inherits the centrality of efficiency by the very same nature of the approach itself.

The project addresses all the four issues of *scalability,* resource *usage and management, security,* and *distribution transparency*, stressed by the Global Computing Proactive Initiative as pivotal to realizing the aim of the initiative.

Although a necessary condition, efficiency by itself does not guarantee *scalability* of the proposed solutions. The project participants will specifically address this issue in the course of the four years of this project. We can thus state that scalability is a pervasive feature of the research work in this project.

In a grand-scale overlay computer, *resource usage and management* become critical primitives to be addressed. Because of its highly distributed nature, it is impossible to consider centralized management of resources. In this setting, primitives for resource discovery become crucial for the viability of any algorithm on an overlay computer. On the other hand, the overlay computer should guarantee proper usage of resources by implementing appropriate algorithms for scheduling and load balancing. Notice that ensuring *fairness* of resource usage is a prerequisite to guaranteeing that end-users will provide their resources to the overlay computer since they will be guaranteed they will be allowed to use other users' resources when needed. Furthermore in our vision an overlay computer is an open system. It is thus unrealistic to assume that users behave according to some predefined protocol. We will consider the resource management also for the case in which users act selfishly and try to optimize their own objectives.

Particular attention has been devoted to the issue of *security* as it will be fundamental for determining whether or not the new technology will be accepted by the end-users. Basic primitives addressing trust management, anonymity, and privacy are crucial in a distributed open system. Unfortunately, most of the solutions in literature cannot be applied *as is* in the context of the overlay computer since they are not scalable or they make strong assumptions that are not realistic in this setting.

Distribution transparency is a crucial issue in this project for different reasons. First of all, a specific feature of the overlay computer is to hide the distributed nature of resources and of computation. Furthermore, we specifically address the possibility of having wireless devices being part of the global computer. As explained above, the overlay computer should hide from the application the specific communication means by which the node is connected. We plan to address problems like stability, topology control, fault tolerance, etc.

Overall, the goal of this project is to investigate the principles and develop the algorithmic methods for building an overlay computer that enables efficient and transparent access to the resources of an Internet-based global computer.

In particular, the our main objectives of this project are:

- To identify and study the important fundamental problems and investigate the corresponding algorithmic principles related to overlay computers running on global computers defined as above.
- To identify the important functionalities such an overlay computer should provide as tools to the programmer, and to develop, rigorously analyze and experimentally validate algorithmic methods that can make these functionalities efficient, scalable, fault-tolerant, and transparent to heterogeneity.
- To provide improved methods for communication and computing among wireless and possibly mobile nodes so that they can transparently become part of larger Internet-based overlay computers.
- To implement a set of functionalities, integrate them under a common software platform in order to provide the basic primitives of an overlay computer, as well as build sample services on this overlay computer, thus providing a proof-of-concept for our theoretical results.

B.2 Relevance to the objectives of the IST Priority

First of all, the proposed project is in line with the strategic objectives of whole IST program in FP6. In general it is within "*the focus of IST in FP6* [that] *is on the future generation of technologies in which computers and networks will be integrated into the everyday environment, rendering accessible a multitude of services and applications*". Macroscopically, it will help to "*ensure European leadership in the generic and applied technologies at the heart of the knowledge economy*" and "*to increase innovation and competitiveness in European businesses and industry and to contribute to greater benefits for all European citizens*".

Specifically the contribution of the project is within the context of the action Future and Emerging Technologies – FET which "*complements the other objectives of IST with research from a more visionary and exploratory perspective*". Our project will help the new S&T field of Global Computing to emerge and mature; we also believe that this field, enhanced by the development within our project will become of strategic importance for the economical and social development of Europe in the future. Clearly the proposed project fits the main call of FET which is to "*support research of a long-term nature and involving high risks that are compensated by the promise of major advances and large potential impact*". Our research is certainly of a long-term nature and parts of it are also risky in the sense that they are well ahead of the frontiers of current research. However we expect breakthroughs in the scientific areas covered by the proposed project since the research groups involved are among the most active ones not only at a European level but also worldwide.

As FET Proactive Initiatives require, our project will focus its resources on areas that are currently considered to be "hot" by the scientific community and their application will have strong potential for future impact. Although we do not expect to realize, during the lifetime of the project, the ambitious vision defined by the Global Computing Proactive Initiative to its full extent, we plan to make significant progress in the proposed areas following an **algorithmic approach**. The project involves twenty one sites from ten different member states, including three partners from three new member states (Cyprus, Czech Republic, Estonia), and two partners from an associate member state (Switzerland). This project is a unique opportunity to create, at a European level, long-standing synergies among scientists working around the focal point defined by the global computing pro-active initiative. More importantly, our project will follow an algorithmic approach to global computing. Such an approach highlights efficiency of the solutions as a key point. We believe that together with the issues rose by the Global Computing Proactive Initiative (i.e., security, resource usage and management, scalability, and distribution transparency), efficiency is absolutely necessary to guarantee an actual future impact of this technology in everyday life.

The global computer the project considers as a base for the research, is composed by Internet-connected computers possibly extended through wireless communications to include other computing devices. Based on the number of emerging applications on the Internet today, the improvements of computational/communication capabilities of wireless devices and the spread of use of such devices, it is very likely that a global computer of this kind will be the main computational infrastructure that will be widely used in a few years. Such a global computer is by its own nature a general-purpose, globally available infrastructure.

Motivated by the key aim of the Global Computing initiative to define *innovative theories, computational paradigms,* etc. *for the design, realization and deployment of global computational environments*, our key goals within the project include:

- the investigation of algorithmic challenges posed by the global computer at hand,

- the definition of the important functionalities required to be provided as “tools for the programmer” when building overlay computers on top of this global computer,
- the development of efficient methods for realizing these functionalities,
- their implementation under a common software platform that will serve as an overlay computer and,
- as a proof-of-concept, the demonstration of services programmed on top of such an overlay computer.

Within the project we will follow an interdisciplinary approach combining theoretical investigations with experimentation that will provide efficient methods that will then further implemented and validated in the context of a software platform and an application. In our theoretical investigations we will use well established notions from graph theory, game theory, scheduling, cryptography, distributed computing, etc., blending ideas from Theoretical Computer Science, Economics and Operations Research. The project addresses all the four issues of *security*, resource *usage and management, scalability,* and *distribution transparency*, stressed by the Global Computing Proactive Initiative as pivotal to realizing the aim of the initiative. We have devoted an entire subproject (SP4) to the issue of security, two subprojects to the issue of resource usage and management, while the issues of scalability and distribution transparency are pervasive in the whole project work plan. A fifth feature that will be pervasive in the work plan will be the issue of **efficiency**, which will be the main concern in all our investigations.

B.3 Potential impact

In the last few years the spread of use of the Internet among the citizens of developed countries is increasing at a tremendous rate. Similarly, the amount of information and services globally available on such a medium is increasing. Users are mostly using wired connections to access the Internet but the number of wireless accesses is exponentially increasing. Consequently, the Internet, along with its wireless extensions, can be thought as the biggest computational/storage resource and, at the same time, the widest fine-grained information gathering system currently available worldwide.

Unfortunately, there is no way of *easily* and *efficiently* deploying a distributed application that would exploit, say, the computational power of nodes that belong to different entities. There are different reasons for this problem. First of all, each entity, e.g. a company, would not allow other entities to use its own resources unless the owner obtains some benefits in return. Currently, the only benefit that can be required is an economical one that, trivially, does not depend on the network architecture. A second important reason is related to the security threats associated to granting access to a computer. Each entity would not allow access to its resources unless it can be sure that such an access is completely harmless.

The main goal of this project is to study the basic functionalities an overlay computer should provide in order to allow an easy, secure and transparent deployment of distributed applications over a network of heterogeneous computers connected by means of a network and belonging to competing entities. The mere existence of an infrastructure that guarantees the above properties will move far beyond the concept of "computationally tractable". The possibility of running a distributed application at a world-level scale, without being forced to obtain an agreement before the application is installed and without having any concerns about the integrity and privacy of its results will allow the fast solution of problems that currently require long-standing computations. In the envisioned system, companies may allow other entities to access their own databases being guaranteed that the proprietary data available will be protected by simply defining some access rules. Both the access to such information and its protection will be guaranteed by the overlay computer and will be made transparent to the programmers. This will reinforce the competitiveness of the companies as, for example, they will be able to make *global* searches about existing products of a given type.

The possibility to extend the overlay computer with wireless devices will have a huge impact on citizens' life. We envision situations in which a *user* of the overlay computer connected through a wireless hand-held device will not only be a consumer of information but also a provider. A citizen may provide the system with information about current traffic information and such information will become immediately available to other users of the system. One may think of embedding wireless devices in cars that could gather traffic information and warn the driver about traffic jams or automatically call a rescue team in case of an accident.

Innovation-related activities. The basic technological innovation this project will introduce is related to the *way* in which distributed applications will *efficiently and transparently* exploit the computational/storage capabilities of the envisioned overlay computer. The possibility of having *transparent, secure and fair* access to an Internet-scale global computer will have beneficial impact, for example, on distance learning, e-commerce, computationally difficult problem solutions.

These effects will be possible due to the scientific advances that we plan to have in the project. Examples are:

- A complete understanding of the structural properties of the topology of an overlay computer will allow at the same time a guarantee against possible link failures in the underlying global computer and an optimized use of its resources.

- In systems such as global and overlay computers where users have their own objectives, it is crucial that the users feel that the resources of the system are used fairly. Fairness however is usually an elusive notion and, even worse, it has an impact on the efficiency of the system. We plan to investigate the fundamental problem of ``fairness vs. efficiency" in global and overlay computers. More importantly, we plan to study how to define and build overlay computers that achieve both fairness and efficiency even when this is not possible in the underlying global computer. This is a powerful incentive to view overlay computers not simply as computing devices but mainly as distributed storage devices on which we can implement transparent, fair, and socially efficient retrieval primitives.
- In order for a system to be scalable, some security problems cannot be solved by manual registration and revocation of each single key and/or identity. Many peer-to-peer and open systems are already reputation-based (think for example of e-bay). One can anticipate that in the global computing paradigm this approach will take an increasingly important role. However, there is a need for better understanding the properties of reputations systems such as their configuration and development and their robustness to malicious attacks. In the first place there is a need for better modeling of this approach. One method to improve the current protocols (which we will investigate during the project) is the use of local "secure distributed computation."
- If short run-time is an appreciated characteristic in every application, in the heterogeneous wireless networks setting it becomes crucial since protocols must potentially work on several different devices having specific features to be regarded. Hence, we are not interested on, say, complex linear programming based algorithms; rather we shall often focus on simple and fast heuristics which work well on the average using local/limited knowledge. Finally, notice that high quality of service requirements is often in contrast with practical efficiency.

Dissemination of results. The results of the project will be disseminated to both industry and academia, starting from the very beginning of the project through a number of different means. First of all, a web-server will be set-up in the early months of the project and will contain all the announcements and the latest results related to the project. A number of mailing lists will be set up in order to cover different aspects of and interests in global computing. Some mailing lists will be accessible only to the project's participants and the Project Officer while others will be open to the public. A periodical electronic newsletter will contain all the relevant information about global computing. As soon as the platform is operational, we also plan to have articles on programming related magazines that will attract the interest of the programmers and companies in the new computing paradigm. From a more technical point of view, the research results will be announced in prestigious international conferences and published in competitive journals of Computer Science. We also plan to organize schools and seminars through which the new ideas and the results produced within the project will be disseminated to young researchers and Ph.D. students.

European dimension. There are mainly two reasons for which this project could not be successfully completed if funded by a single member state. As noted in the previous section, the goals of the project are risky in the sense that they are well ahead of the frontiers of current research. No individual member state has the necessary human resources, in terms of number of researchers with complementary expertise, needed to carry out the necessary research work in a relatively short period of time of four years. The possibility of gathering participants at a community level guarantees the critical mass needed to successfully complete the project.

A second, non-trivial implication is that the participants to the project will provide part of their own computational infrastructure in order to allow experimentation and testing of the proposed solutions on an actual Internet-based geographically distributed system of large scale.

The project will also benefit from and significantly further extend the results of the projects of the FP5 Global Computing cluster entitled "Foundations of Networks and Large Distributed Environments", namely the IST FET Projects CRESCCO (IST-2001-33135), DBGLOBE (IST-2001-32645) and FLAGS (IST-2001-33116).

B.3.1 Contributions to standards

-- Not applicable --

B.4 Outline implementation plan

As described in the previous sections, for the purposes of this project, we consider the global computer to be composed by computing entities connected to Internet by means of different communication media, i.e., wired and wireless ones. Some of these devices will mainly be of similar computing power, e.g., personal computers and mobile laptops, but potentially also others like satellites, palmtops, mobile phones, and smaller peripherals like sensing devices could be part of the global computer. An overlay computer is a logical abstraction over this global computer that provides transparent access to its resources by means of basic primitives. By its own nature, an overlay computer will be a worldwide distributed system composed by thousands of nodes.

In line with FET, and in particular with FET pro-active initiatives, this project addresses topics that are currently considered relevant by the scientific community whose application have a strong potential impact in the future. The final goal of providing an efficient overlay computer can only be fulfilled after the basic algorithmic principles related to this topic have been identified and understood in depth. Without such a necessary step, any possible proposed solution could only be based on conjectures and/or heuristics that could lead either to unstable or non-scalable solutions.

Because of the definitions of overlay computer, scalability itself is a basic issue that must be addressed. A fundamental ingredient to any scalable solution is efficiency. Although not all efficient solutions are scalable, it is true that all the solutions that are not efficient are also not scalable. By adopting an algorithmic approach for the study of primitives an overlay computer should provide, we inherit the guarantee for efficiency of solutions as this feature is a basic measure for the assessment of solutions.

As the overlay computer is intended to be an *open* system, it is realistic to think that it will comprise nodes belonging to different entities, having little knowledge of the status of the system, trying to optimize their own objectives and possibly with malicious users. This new scenario by itself poses new challenging algorithmic problems. We propose to develop techniques to cope with such situations by blending techniques from Theoretical Computer Science, Economics and Operations Research. Techniques of this kind have already proven to be successful in solving problems in this new context.

In order to deal with the complexity of the objectives, our research work is divided in several vertical and horizontal components. Horizontal components contain work either on the investigation of general fundamental aspects or on the development of the overlay computer. Vertical components focus on specific areas which are important for the realization of the overlay computer. These components are the following:

1. A horizontal component addressing fundamental issues of global computing with special focus on efficiency, transparency and scalability.
2. Four vertical components addressing the issues of resource management, sharing information and computation, security and trust management and transparent extensions of overlay computers to include wireless devices, respectively.
3. A horizontal component devoted to the implementation of basic functionalities under a common software platform and the development of an application on top of it.

These components correspondingly define six sub-projects (SP):

- SP1 "Paradigms and principles" will be devoted to the development of "*innovative theories*" to cope with new algorithmic problems that arise in Global Computing. It will study the structural properties of global/overlay computers, fundamental techniques for coping with selfishness and for achieving stability and fault tolerance, and will tackle the challenge of computing with partial (i.e., uncertain, distributed, or even incomplete) knowledge by blending theories from economics, game theory and algorithmic theory. A better understanding of these problems will have a strong impact on the ability to propose scalable, distributed and dynamic algorithms. That will also allow understanding the efficiency trade-off between unwanted centralized strategies and expected fully distributed strategies.
- SP2 "Resource management" will focus on specific aspects related to the management of critical resources (like bandwidth), resource discovery, as well as to the design of mechanisms

for accessing resources owned by selfish entities. Resources can either be of a low-level (i.e., infrastructure-dependent like bandwidth) or application-level (e.g., currency).

- SP3 "Sharing information and computation" will consider algorithmic problems related to the management of resources focusing on computational and information resources. Issues like distributed data management, load management and scheduling will be addressed here. Together, SP2 and SP3 will address in depth the resource usage and management issues posed by the Global Computing Proactive Initiative.
- SP4 "Security and trust management" will explore problems related to trust management, authentication mechanisms, privacy, anonymity, and secure distributed computation (including techniques to face malicious behavior). The main goal of the research work in SP4 is to address fundamental issues that are crucial to a transparent security layer. In achieving this goal we will adapt concepts from cryptography and economics that have recently shown to be very successful in modeling adversarial but rational behavior.
- SP5 "Extending global computing to wireless users" mainly aims to transparently include wireless nodes in an Internet-based overlay computer. It will focus on issues like resource management and quality of service in wireless sub-networks, network design and topology control under dynamic scenaria, mobility and fault tolerance. The main objective of this subproject will be to provide practically efficient algorithmic solutions for high-quality, reliable, stable end-users services on heterogeneous wireless networks. Due to the specific limitations of wireless devices, a particular attention will be devoted to the efficient usage of critical resource like energy and spectrum (i.e., frequencies).
- SP6 "Design and implementation of components and applications for programmable overlay computers" will be devoted to the implementation and integration of functionalities produced within Subprojects SP1, SP2, SP3, and SP4 into a common software platform that will provide the programmable interface of the overlay computer. Special attention will be devoted to the efficiency of the implementations. An application will be selected and implemented on top of this overlay computer to serve (together with the platform and the integrated functionalities) as a proof-of-concept.

Components dealing with Training, Demonstration and Management activities are explained in details in subsequent sections.

B.4.1 Research, technological development and innovation activities

Sub-Project 1: "Paradigms and Principles"

Leader: UOA, Participants: UOP, CNRS, UPB, CTI, UNISA, CUI, CAU, UDRTV, UNIPD, UPC, UDRLS, INRIA, DIM, UCY

To face the challenge of designing and implementing useful overlay computers, we need to address the novel algorithmic issues that arise from mapping- in a robust, efficient, transparent, and scalable way - in overlay computer to an underlying global one. Entities of an overlay computer should be reflected transparently into a single entity and sometimes into a large subset of entities of the underlying global computer. Communications between these entities at the overlay layer must be implemented effectively through the underlying global computer. This is a qualitatively different problem than the problems of designing physical networks.

To design overlay computers with improved characteristics, we need to address the problem of coordination between the entities of the global system; the problem is complicated because these entities are either selfish or lack complete knowledge of the situation. A successful design of an overlay computer could alleviate the problems arising from selfish behavior and incomplete knowledge at the underlying global system and improve its efficiency.

To face these challenges, we intend to address foundational algorithmic problems related to structural and topological properties of global and overlay computers, to study algorithmic issues and propose frameworks that take into account the distributed and selfish nature of the entities involved, to study the stability and fault tolerance issues is such complex dynamic systems, and to apply the rigorous framework of competitive analysis to deal with uncertainty. A better understanding of these problems will have a strong impact on the ability to obtain scalable, distributed, and dynamic algorithms. That will also allow us to understand and alleviate the difference between undesirable centralized solutions and desirable fully distributed ones.

Additionally, we plan to give emphasis to robustness, scalability, and more importantly simplicity. For global computing, it seems more important to obtain fast and simple algorithms rather than optimal and complicated ones; for example, a fast and simple sub-optimal shortest path algorithm may be more preferable to a complex optimal one. Of course, algorithms appropriate for global computing must be also fault-tolerant and behave reasonably well under imperfect knowledge.

In summary, the research work within the workpackage consists of the following workpackages; these are described in more detail below:

- WP 1.1: Structural Properties
- WP 1.2: Coping with Incomplete Knowledge
- WP 1.3: Coping with Selfishness
- WP 1.4: Stability and Fault-Tolerance
- WP 1.5: Generic Algorithms

WP 1.1: Structural Properties

Leader: UCY, Participants: UOP, CNRS, CTI, CUI, UDRTV, UOA, UNIPD, DIM

A simple yet powerful model, which captures the distributed and heterogeneous nature of global and overlay computers, is provided by graphs. We plan to study the topological properties, connectivity issues, substructures and organization of global and overlay computers. A critical task is to study in depth the embeddings and other topological relations between global and their overlay computers.

Topological properties

In the last decade a significant part of research on the Internet and the Web is concerned with their topological properties. In the Internet, starting with the novel work of [FFF99], power laws and other structural phenomena have been observed and measured. Similar properties have been observed in the web [BaAl99, KRRSTU]. An important goal of this research has been to come up with the right models that behave or predict such phenomena; this would allow us to predict the future of these networks as well as provide a framework to simulate and test new algorithms and protocols. There is a vast literature on classical random graphs [Bol85]; most of it does not apply to the Internet and the Web, but it provides the background for new models for the Web [BaAl99, BRST, KRRSTU, ACL00, CF, SaSp02] and the Internet [CD, FKP02, BergerEtAl03]. We plan to build upon and expand this research, but most importantly we plan to study the question on how to design algorithms and protocols that take advantage of the topological properties such as self-similarity of global and overlay computers [ALPH01].

Substructures and organization

The phenomenal success of Google owes much to a simple idea: content is reflected by the link structure of the Web [Kle99, PageRank]. This points out that a central issue in global and overlay computing is to determine the role of substructures, such as communities, and to devise efficient algorithms to search for, count, and- in overlay computers-build them. A similar problem is to be able to identify appropriate centers, clusters, backbones in various metrics of the underlying graph [Indyk01] that can be used to improve the quality of service and to avoid security and computational vulnerabilities. The immediate challenge is to use and improve upon the existing literature on clustering, partitioning, and computational geometry [Pel00, ElPe01, GPPR01].

Clustering and partitioning graphs and metrics are relevant to global computing in two ways: First, they reveal the existing economic and social communities and relations, and second, they are essential for efficient computation and storage in global and overlay computers.

Embeddings

A central issue in designing and analyzing overlay computers is the mapping of their components onto components of the underlying global computer. It is related to the problem of embedding one space (usually a metric or a Euclidean space) into another. These kinds of problems have been studied in depth by the Math community [Bourgain85, Matousek02]. Also some algorithmic issues have been addressed and they led to some of the deepest results in approximation algorithms [LLR94, ARV04].

However, most of the issues are open. Even the fundamental question of low-distortion embedding into low-dimensional Euclidean spaces is open [KRS04].

The question of appropriate embedding is relevant to peer-to-peer networks (associate virtual and real address) as well as to sensor networks (find a simple geometric routing). It provides also a general paradigm to solve hard or distributed computational problems: First map the original domain to a simpler one, solve the problem there, and then map back the results to the original domain. Low distortion maps will result in good approximation algorithms. Finally, embeddings are very much related to the problem of homomorphism and isomorphism between distinct structures and substructures (i.e., an isomorphism is a zero-distortion embedding).

We plan to exploit and improve the existing results and to address the issue of finding mappings between global and overlay computers that have desired properties and are computationally feasible.

WP 1.2: Coping with Incomplete Knowledge

Leader: UOA, Participants: UOP, CTI, UNIPD, UPC, INRIA, DIM, UCY

By their nature, algorithms and protocols in global and overlay computers have little knowledge and much less control over the state of the system. If we also consider the unpredictability and the ever-changing environment of such systems, we see the need for most algorithmic issues to be studied in a rigorous framework. Fortunately, such a framework for optimization problem under uncertainty has been used extensively over the last two decades: competitive analysis [BY98, FW98, KoPa94]. We plan to use this framework together with the theory of distributed systems to study the novel online problems that are at the core of global and overlay computers, such as distributed and robust data structures, reputation mechanisms, transparency and information hiding.

Distributed and Robust Data Structures

Global computers such as the Web and P2P systems are more storage rather than computation devices. Also a central property of all-global and overlay computers are their distributed nature. It is thus very important to come up with efficient and robust data structures that reside on selfish, distributed, and dynamic environments. Peer-to-Peer systems and classical distributed computation have scratched the surface of this fundamental problem. It is not at all clear how and where to maintain the state of the system. Data structures is one of the oldest and best-developed areas of Computing (both in theory and practice). However, within the rigorous framework of competitive analysis, many fundamental problems remain open (the Splay Tree Conjecture [ST85] concerning the lowly binary search trees remains one of the outstanding open problems).

We plan to study simple primitives that allow the maintenance of data structures in selfish, distributed, and dynamic environments.

Transparency and Information Hiding

The very idea of an overlay computer is based on the assumption that some information about the underlying global computer is intentionally or otherwise hidden. For example, one could view the Web as a collection of storage areas in users' computers shared by all; but its true power comes from viewing it as a collection of interlinked web pages with the original directory structure hidden.

Information hiding is an important topic both for the programmer of an overlay computer and the user. How much is known about the underlying global computer by the programmer of the overlay computer? We plan to study and quantify the advantages and disadvantages of information hiding. The same question for users has a completely different flavor. For selfish users, information hiding is an essential tool for the system designer to achieve good coordination and low price of anarchy [CKN04].

In the fully distributed environment of global computers, nodes have a restricted in space and time (and not necessarily accurate), view of the global structure. This leads to questions about how one can encode some global properties locally (see for example [GaPe96a, GPPR01]). Solutions are often based on clustering methods, or on distributed hash tables. One of the main problems is to adapt these methods in order to deal with the fast dynamics of the network. We plan to use competitive analysis to study such methods.

Reputation Mechanisms

A reputation mechanism is a particular type of data structure together with algorithms for updating and retrieving the information. A reputation mechanism aims to predict the behavior of the users based on past information. For example, in P2P systems, to prevent users from being free-riders, one approach is to grade the users based on their past behavior and exclude users with low grade. How to grade users and which part to exclude is a novel type of online algorithms. More generally, the reputation of a user is a distillation of the user's interactions with the system that can be potentially useful for predicting his future behavior. How to distill, how to maintain reputations, and how to predict is an important online problem which differs from traditional problems, such as the k-server problem, in many aspects: instead of an adversary, we deal with selfish users, and instead of a centralized algorithm, we have a distributed protocol which decides on information stored and possibly manipulated by the same users that it tries to analyze.

One essential question is where the distilled information about past transactions is kept. There are several proposals - at a central station, each user maintains his own transactions, etc., - each with its own problems. The quality of these solutions is not separate from the prediction algorithm. We plan to study the problem in the rigorous framework of competitive analysis.

WP 1.3: Coping with Selfishness

Leader: UOA, Participants: UOP, UPB, CTI, UNISA, CAU, UDRTV, UPC, UDRLS, UCY

Global and overlay computers are essentially a collection of computers, storage, and other resources controlled by different users, corporations, and organizations. It is essential that all these entities have an incentive to participate, to provide resources to the global system, and to use some of its resources. The selfish nature of the entities is manifested in many global systems: Internet would not be such a global system if its nodes had no incentive to be part of it-fortunately in this case, selfishness has not been interpreted in a narrow sense. The selfish nature of global computers is manifested even more in P2P systems whose very nature is to pool together the resources of its users.

Price of anarchy

In the last five years, starting with the paper [KP99], there has been a coordinated effort to measure the deterioration of a system due to the selfish nature of its entities, its price of anarchy. An important aspect of selfishness is that it is very much related to the information regime in which the users live. This is illustrated by the fact that when users are given more information about the state of the system, the price of anarchy may fall but it may also rise; in the latter case, the users may use the additional information to try to exploit the system and end up in a bad equilibrium [CKN04, MT03]. A fundamental question that we will try to answer is how much an overlay computer should hide from the underlying global computer. The challenge is to pose this question in the right models and quantify it using the notion of the price of anarchy. In the follow-up phase we plan to move beyond the question of “how much to hide” and address the problem “how to hide” by designing appropriate protocols and algorithms. Randomization will, of course, play a central role in improving efficiency, and we plan to explore its advantages. With randomization, one can hide certain aspects of the underlying global computer, but more importantly, when it is used inside protocols, alter the framework so that the objective of each users align with the objective of the whole system; this essentially happens by interchanging the quantifiers on the mathematical expressions that capture the users’ objective.

Fairness versus efficiency

In systems such as global and overlay computers where users have their own objectives, it is crucial that the users feel that the resources of the system are used fairly. Fairness however is usually an elusive notion. Even worse, it can be achieved only by making the system inefficient (for example, a fair protocol to handle a queue of identical requests may lead to inefficient solutions). We plan to investigate the fundamental problem of “fairness vs. efficiency” in global and overlay computers. More importantly, we plan to study how to define and build overlay computers that achieve both fairness and efficiency even when this is not possible in the underlying global computer. This is a powerful incentive to view overlay computers not simply as computing devices but mainly as distributed storage devices on which we can implement transparent, fair, and socially efficient retrieval primitives.

Equilibria selection and mechanisms

Selfish behavior is usually studied under the assumption that users operate at a Nash equilibrium point. However, in the vast systems of global and overlay computers such assumptions, albeit good approximations, are unrealistic because of the distributed nature of the system and the lack of information about the state of the system. In most cases, users and other entities should not be viewed as super rational exploiters of the system but as players who, with appropriate “signals” would end up at a state, which is beneficial for all of them. In a way, a good overlay computer is one that-among other objectives-achieves an appropriate coordination between its nodes by providing appropriate views about the state of the system to them. An important problem is how to provide the appropriate signals to achieve a good (or the best) equilibrium. For example, the scheduling policy of shared resources can affect dramatically the behavior of selfish users [CKN04, CDR03a, CDR03b], and the question is which is the scheduling policy, which achieves the best coordination [Kou03].

A programmable overlay computer requires the programmer's instructions to be carried out correctly by all involved entities. A central issue for building such overlay computers is thus how to obtain the entities cooperation. Mechanism design is precisely the study of how one can design systems such that the entities' selfish behavior results in the desired system-wide goals.

Existing approaches in microeconomics have been shown to be inappropriate for problems in which computational issues play a crucial role [NisRon99, NisRon00, ArcTar01, AuletalSTACS04, AuletalSPAA04, AuletalICALP04, MelPenProWatWid04], thus not appropriate for global and overlay computers. Beside the fact that "classical" mechanism design techniques do not take into account the computational overhead required by the resulting mechanism, these solutions do not exploit other computational aspects which turn out to be helpful. We plan to have a deeper understanding of the impact of computational issues arising in global and overlay computers on mechanism design problems.

WP 1.4: Stability and Fault Tolerance

Leader: UPC, Participants: UOP, CNRS, CTI, CUI, DIM, UCY

An important issue in the vast modern networks -which is usually taken for granted but can have devastating effects- is the issue of stability. It is not entirely clear what stability entails in such environments. The classical approach through queuing theory had the advantage of a nice crisp framework with cute and deep results. However, it is well known that such a framework is not adequate for dealing with the dynamics of modern networks. Some interesting extensions of queuing theory, for example, adversarial queuing theory [BKRSW96], have successfully addressed specific problems in a satisfactory way. We plan to study the stability of global networks and especially how the interplay of global and overlay systems can affect stability.

Adversarial Queuing Theory

An important issue to consider is the increase in communication needs of overlay computers as the required information may be spread among many nodes. The dynamics and uncertainty of the situation calls for the use of the adversarial queuing theory (AQT) [BKRSW96]. Although AQT has been used successfully to analyze the stability of static systems, we need to extend the results to dynamic networks.

We plan to extend the AQT framework to model and analyze the communication system of overlay computers. New models of communication systems (and adversaries) are required to capture the high mobility an unpredictability of such systems. We will attempt to explore the range in which these models are stable and to characterize the stable protocols. We also plan to study stability under milder, more realistic, conditions (by restricting the adversary, by limiting the type of communication paths, by rerouting, by failure-recovery).

More importantly, we plan to move beyond the simple characterization of stability and propose coordination mechanisms that lead to improved stability.

Distributed systems

Global computers are characterized by high dynamics and uncertainty, high probability for faults, selfish behavior, and perhaps Byzantine behavior. Also the life expectancy of their nodes is usually short. The very large scale of the network then implies that nodes have only a restricted (local) view of the network.

One of our goals is to design and analyze distributed algorithms that allow a global system to recover from errors and faults. These algorithms should be local (local errors should not affect the network globally) and scalable (the algorithm complexity should be almost independent of the network size).

Since error detection is often difficult (and sometimes even impossible), we plan to explore error detection mechanisms that detect errors quickly and reliably in most cases. For multi-layer networks, errors occurring in one layer may have an impact on another, and even simple problems like path recovery have to be studied in this context.

Apparently there is a trade-off between the impact of a failure (how fast the fault is recovered and how many connections must be interrupted during the recovery procedure) and the quantity of resources allocated to protect from faults. We plan to design recovery policies that achieve a desired impact with minimal resources.

WP 1.5: Generic algorithms

Leader: CNRS, Participants: UOP, CTI, UDRTV, UOA, UPC, UDRLS, INRIA, UCY

To develop a theory of overlay and global computing, it is important to identify precisely and solve the optimization problems involved in the efficient implementation of overlay computers; the study of encompassing paradigms seems also an essential step towards addressing these optimization problems. As a first approximation we should deal with standard optimization problems for which we shall seek fast and simple approximation algorithms. This will be also useful for overlays computers in the cases where their nodes are guaranteed (by agreement or by a mechanism) to follow certain behavior rules.
Fast approximation algorithms and their analysis (in particular in the probabilistic setting) will be one of our main targets. We will, whenever possible, attempt to find paradigms general enough to apply to several problems.

Connectivity issues

Peer-to-peer networks and mobile communication environments, such as mobile ad-hoc networks and low earth orbiting satellite systems, present a paradigm shift from back-boned networks in that data is transferred from node to node via peer-to-peer interactions and not over an underlying backbone of routers. Naturally, this engenders new problems regarding optimal routing of data under various conditions over these dynamic networks. Routing using shortest paths or least cost methods is complicated by the changes in the topology of the network. Unfortunately, temporal dependencies in networks are hard to be effectively captured in a classical graph model. The problems require a new approach and there have been attempts to tackle them by considering random variations in link costs and connectivity [Sch02], the space-time approach [KLS02, KoSk02], and evolving graphs [BFJ03, BFJ03a].
Connectivity issues play also a central role in wireless and mobile networks. A fundamental problem in this particular context is to maintain shortest paths. In the distributed, transparent, and perhaps faulty domain of overlay computers, this cannot be done unless we relax the requirement to t-spanners, which guarantee a bounded dilation. The challenge is to maintain t-spanners in dynamic and kinetic conditions (the former indicates the situation in which connections are subject to failures and the latter refers to mobility of the processors). Since optimal centralized solutions are unrealistic, a major issue is the study of suitable distributed solutions that guarantee computable efficient solutions and low energy consumption.

Multi-criteria optimization

In a global computer, the connections between the computing nodes have different qualities in terms of bandwidth, latency, reliability, or cost. Different applications will require overlay connections with different properties: one may need a low bandwidth but very reliable connection, another may require high bandwidth but low reliability (e.g. ftp). An important algorithmic issue is to route traffic according to those multiple criteria (Multi-criteria Path-routing, Flow, Minimum trees, see for example [GoRa96, IMRR+95]).

Sampling paradigms

By sampling one means to probabilistically select a small subset of the problem data and use it in order to solve an optimization problem quickly. For example, Karger used this approach to fast compute a small cut (i.e.. network bottleneck, see [Kar94, Kar03]), and to evaluate network reliability [Kar01].
Sampling will also be essential to evaluate the network state. The nature of overlay computers (high dynamics, very large scale) makes this approach essential and we plan to pursue it.

Sub-Project 2: "Resource Management"

Leader: CNRS; Participants: UOP, CTI, UPB, UNISA, UOI, CUI, MPII, CAU, UDRTV, UOA, KUL, UCY.

A crucial issue in global computing environments is that of resource management. In such an environment, resources that may be distributed in several nodes, information about their position may be only partially available, they may be owned by different entities each having its own objective for releasing these resources, and their efficient use may also depend on structural properties of the underlying infrastructure (e.g., the topology of the underlying network).

The work within this subproject focuses on the study of fundamental issues for accessing and managing communication resources in an overlay computer. Our research will address novel and challenging algorithmic issues for efficient resource discovery and querying like construction of overlay networks, query routing and execution, and for sharing critical resources like bandwidth. Our work will also include mechanism design for coping with selfish behavior when allocating resources in a distributed, uncoordinated system such as a global or overlay computer.

In summary, the research work within the workpackage consists of the following workpackages; these are described in more detail below:

WP2.1 “Resource Discovery”
WP2.2 “Sharing Critical Resources”
WP2.3 “Mechanism Design”

WP2.1 Resource Discovery

Leader: UOI; Participants: UOP, CTI, MPII.

In this workpackage, we shall explore advanced methods for resource discovery that take into account the structure and content of resources. A central objective is supporting resource discovery based on more advanced queries than single attribute-value search or keyword search. These include queries that support relational operators as well as path queries on structural documents (such as XML documents or RDF descriptions).

The focus will be on scalable, distributed and dynamic solutions that take advantage of the embedding of virtual networks into real ones. To this end, we shall build upon the paradigms and principles of global overlay computers as studied in WP1. Our goal is to fully explore the work on the structural properties and general principles of global networks (such as their topological properties, the observed selfish behavior and partial knowledge computation) towards extending the proposed solutions to take advantage of richer resource models.

In dynamically evolving global overlay computers, such as in peer-to-peer systems, there are two basic approaches of organizing information: structured and unstructured ones. In structured systems, data items (or indeces) are placed at specific peers usually based on distributed hashing (DHTs) such as in CAN [RFHKS01] and Chord [SMKKB03]. With distributed hashing, each data item is associated to a key and each node is assigned a range of keys and thus items. Nodes are interconnected via a regular overlay topology where nodes that are close in the identifier space are highly interconnected. Very recently, researchers have proposed extending DHTs (e.g., Chord) with long-range links towards creating small-worlds [MBR03]]. In unstructured systems, there is no assumption about the placement of data items at the nodes. When there is no information about the location of data items, flooding and its variations are used to discover the nodes that maintain data relevant to a query. With flooding (such as in Gnutella), the node where the query is originated contacts its neighbour nodes, which in turn contact their own neighbors until a matching node is reached. Flooding incurs large network overheads, thus to confine flooding, indeces are deployed. Such indeces can be either centralized (as in Napster) or distributed among the nodes (as in routing indeces) [CG02] providing for each node a partial view of the system.

In this workpackage, we shall study (i) the construction of structured and unstructured overlay networks based on the content of nodes and (ii) processing queries on top of the constructed networks.

In the structured approach, recent extensions propose instead of associating keys to data items based just on their identifier, to associate with each data item (or node) a vector describing its content extracted using IR algorithms. This vector is then used as input to the hashing functions [TXD03, SP03]. This is a step towards exploring the content of nodes in structured networks. However, this creates a dimensionality reduction problem, since the dimension of the vectors should match the dimension of the DHT. In addition, the usual problems with structured systems arise, since although DHTs provide very efficient search, they compromise node autonomy. The DHT topology is regulated since all nodes have the same number of neighboring nodes and the selection of nodes is strictly determined by the DHT semantics. Furthermore, sophisticated load balancing procedures are required. Obtaining optimal embeddings, embeddings without the need for global knowledge and embeddings appropriate for different type of resources and queries are still open issues.

In unstructured networks, many recent research efforts focus on organizing nodes in clusters based on the content of the nodes. In most cases, the number or the description of the clusters is fixed and global knowledge of this information is required. In [BMR03], nodes are partitioned into topic segments based on their documents. A fixed set of C clusters is assumed, each one corresponding to a topic segment. Knowledge of the C centroids is global. Clusters of nodes are formed [TXKN03] based on the semantic categories of their documents; the semantic categories are predefined. Similarly, [CGb] assumes predefined classification hierarchies based on which queries and documents are categorized. The clustering of nodes in [LNSNT03] is based on the schemes of the nodes and on predefined policies provided by human experts. In the associative overlays of [CFK03], routing is restricted among groups of nodes that contain data items that are semantically similar. Such groups, called guide rules, are specified by predicates on properties of the nodes, such as the existence of a data item. A decentralized procedure for clustering is proposed in [PP04, PKP04]. Issues such as optimality, scalability and adaptability are still open for further investigation.

In general, processing relational queries is a very active research topic. There have been a number of proposals for processing range queries in structured networks, including CAN [SGAE04] and Chord [TP03]. Two approaches for supporting multi-dimensional range queries are presented in [GYG04]. PIER enhances DHTs with declarative and algebraic query interfaces that are used for querying very large networks [HHLL03, THHSS04]. There is no match work in processing relational operators in unclustered networks. Regarding processing XML queries, a first approach for unstructured systems is presented in [KP04] and for structured in [GWJD03].

Keyword queries in P2P systems are addressed in [BMWZ04, WD04], with particular emphasis on query routing. While relational queries go way beyond keyword search, appropriate routing strategies may be related to those for simpler queries.

The workpackage is divided into the following tasks:

Construction of overlay networks. The challenge lies in building an overlay network such that it connects the semantically related nodes most efficiently but at the same time provides highly efficient routing to all nodes with logarithmic worst-case behavior in terms of both space overhead per node and routing hops. Also, the overlay network needs to cope with high dynamics because of failures, heavy load fluctuations, and many nodes unpredictably leaving or joining the overlay computer. It needs to be investigated to which extent this goal is feasible with approaches based on distributed hash tables.

Query routing in overlay networks. The problem is to dynamically decide to which small subset of nodes in an overlay computer a query request of a certain type and with certain input parameters should be routed. A judicious decision needs to take into account the latency and bandwidth of the network paths to the target nodes, the current and near-term-predicted load and response time of the target nodes, and the expected quality of the query result. Here quality refers to completeness, precision, freshness, and other data-quality attributes. The task lies in developing methods that can cope well with these tradeoffs.

Query execution in overlay networks. Besides the routing decision, there are still many degrees of freedom in executing a given query. These include decomposing the query into subqueries and mapping the subqueries onto target nodes, and the dynamic scheduling of the subquery execution order, where order is a partial order that allows parallelism. Furthermore, decisions need to be made about the extent to which intermediate query results should be pipelined between nodes, or whether intermediate results should be fully materialized on certain nodes. Once the query is running, the conditions for generating the execution plan may change in and unpredictable manner. So it is mandatory that execution plans can be adapted at run-time to take into account network path congestions, node overload, unexpectedly poor query result quality from certain nodes, and so on.

For each of the above tasks, we will provide systematic surveys of the state-of-the-art, we will develop models and methods for the construction of overlay networks, and we will design, theoretically analyze and experimentally validate algorithms for query routing and execution, and for dynamic adaptation of query adaptation plans.

WP2.2 Critical Resource Sharing

Leader: CNRS; Participants: UOP, UPB, CTI, CUI, MPII, CAU, UCY.

In an overlay computer, the programmer may have a vision of a structured network of resources, while the underlying infrastructure may be very irregular and heterogeneous. The users do not know and do not want to know what is the effective protocol or physical path it uses when it communicates. The work to be done is to provide efficient and scalable algorithms for transparent bandwidth sharing in a heterogeneous large-scale global computer in order to match the quality of service expected by the overlay computer.

In order to achieve this, one must take into account several aspects:

First, due to the network hierarchy, routing strategies will involve grooming issues. Indeed bandwidth requests have to be packed into a set of virtual links, which are paths either, defined on the physical network or on the network layers below. This is the case for the design of optical networks using Wavelength Division Multiplexing [HPS02, BPS02, MuZh02, GRS00], Synchronous Digital Hierarchy, ATM, [BMPP03, BMPP99a] or in the case of routing IP traffic over WDM using MPLS technology. Usually, the aim is to find a good balance between the virtual network efficiency and its complexity or cost. These problems are captured by the packing paradigm. Whereas many results exist for specific instances, general methods still have to be found. Ideas coming from clustering methods (see the book [Pel00]), and spanners [ElPe01] should be studied in order to improve virtual topology design techniques. Note also that virtual network design may also involve mirroring issues (i.e. akamai).
Second, due to the variety of the physical links (high speed wired, satellite, radio, adsl) the quality of service parameters (financial cost, delay, reliability, bandwidth) is getting extremely heterogeneous. Taking them into account is now crucial. One needs to model how internet agents may in the future trade or share bandwidth with refined qualities, and to provide algorithmic solutions. Solutions will have to borrow from mechanism design concepts in order to ensure fairness and adequacy between the cost and the quality of service.
Third, existing algorithmic solutions are often at least partially off-line, designed for static instances, and may be too slow to ensure practical scalability. So finding new methods that are practical in terms of scalability is important. We will build upon existing methods used for the call admission problem in an on-line setting. Note that the dynamic environment also implies that algorithms should also be fault tolerant.
Last, wireless networks renew many classical problems; these are studied in the context of subproject SP5.

Our approach will be based on existing work on routing problems. Those problems consist of assigning a set of paths (or routes) in the real network such that virtual connections designed in the overlay network are satisfied, subject to some constraints (capacity, delay, reliability, etc.). They are usually modelled as multi-commodity flow problems (integral or fractional), or as coloring or fractional coloring; they can also be formalized as call scheduling or call admission problems. The models reflect the technology used in the real infrastructure considered at the global computer layer, e.g., a WDM network (see for example [CFKP+01a, BHP98, ABCR+96]) or a radio network ([KMP04, CKP03, CKP02]).

WP 2.3 Mechanism design

Leader: UNISA; Participants: UPB, CTI, UDRTV., UOA, KUL, UCY

In classical game theory, the way to deal with selfishness and untruthfulness is mechanism design, which is a kind or reverse engineering for games: given the desired behavior of the players, design a game to achieve it. More or less, mechanism design is sophisticated auction design. Most of the vast literature on this topic studies centralized schemes with few players. It is impossible to apply directly the nice results of this area to the entities of global or overlay computers for many reasons: because of their distributed nature, because of the computation and communication overhead, because of the lack of a common "currency" (i.e., entities have diverse interests in commodities such as time, storage, connectivity etc).

An archetypical example that shows that a novel approach in mechanism design is needed is the very successful protocol TCP: TCP, which remained robust during the period of tremendous growth of the Internet, works well when users act unselfishly. No mechanism design was implemented in TCP; rather the simple yet effective scheme of additive- increase-multiplicative-decrease has worked very well. The traditional approach of Game Theory would lead to the wrong conclusion that TCP would be very inefficient and would be wiped out. We plan to study mechanism design and address the issues arising from these considerations. We are interested in finding the right questions, the right models and framework to study mechanisms that are simple, efficient, and robust.

Given a distributed, uncoordinated system, such as a global computer, with selfish users, we want to investigate how to redesign it so that its price of anarchy decreases and that globally optimal (or close to optimal) solutions can be achieved in spite of the fact that each agent tries to optimize a different local objective function. The task at hand is to determine the primitives that a mechanism can use and how to apply them. This is made more difficult (and interesting) by the fact that part of the input is privately known to the agents, which can misreport such information if this is more beneficial for them.

Classical results from Game Theory are based on the assumption that a basic shared currency is available and then the task is to design truthful mechanisms where users have an incentive to collaborate. Recently, interesting interactions between this classical game-theoretic notion and algorithmic issues have been underlined in [NisRon99] for classical optimization problems. Another primitive for inducing unselfish behavior in the agents consists in slowing down some resources that may have the beneficial effect of helping the users coordinate so that their diverse objectives align with the global objective. Another potential primitive is to be able to hide or reveal some information about the state of the system to the users, again in the hope that this will alter the equilibrium towards a better value. Selecting a scheduling policy on resource usage is also another primitive that can lead to better equilibria. This a typical case in which lack of a common shared currency is compensated by means of a meta-resource: users are induced into unselfish behavior by receiving better service. Stackelberg strategies are also a way of taming selfish users and induce a good global behavior. In a Stackelberg game, a good global status is induced by imposing good (unselfish) behavior on a α fraction of the users. Trade-off between α and the quality of the global induced status are studied.

We plan to study the impact of these primitives on improving the price of anarchy and in reaching a good global status in global and overlay computers despite the presence of selfish users. An interesting question is how to find the right abstraction of a global computer, or equivalently on how to design an appropriate overlay computer, which implements appropriate primitives to improve the system performance. One of the questions we will try to answer is which information is critical for a given resource management problem? Some routing/scheduling problems can be solved optimally only if some part of the information is not private knowledge of the selfish entities [AuletalSTACS04, AuletalICALP04, MelPenProWatWid04].

Sub-Project 3: "Sharing Information and Computation"

Lead Partner: UPB; Participants: CTI, UOI, UOA, UNISA, UDRLS, UNIPD, CAU, MPII, CNRS, ETHZ, UPC, UCY, DIM.

Information and computational power are important resources in a global computing environment. Although the global computers we consider as a basis in this project (i.e., a set of geographically-distributed computing nodes connected through the Internet) may aggregate huge amounts of data and computational power, their efficient allocation to users and applications is of critical importance since it may significantly affect the performance and scalability of global and overlay computers. Controlling the access of users and applications to information and computing power is the aim of this workpackage.

The work within this subproject focuses on the study of fundamental issues for organizing and accessing information in overlay computers and for controlling the computing power of their nodes. Novel algorithmic issues in distributed data management including caching and replication of primitive or more complex data (e.g., metadata), load management including load balancing and tuning and parameterization of adaptive software, and scheduling motivated by the size and the dynamic nature of overlay computers will be addressed, while we will also attempt to model intensive computations in overlay computer as processes in workflow management systems.

In summary, the research work within the workpackage consists of the following workpackages; these are described in more detail below:

WP3.1 “Distributed Data Management”
WP3.2 “Load Management”
WP3.3 “Scheduling”
WP3.4 “Workflows and Services”

WP 3.1 Distributed Data Management

Leader: MPII; Participants: UPB, CTI, UOI, CAU, UOA, ETHZ

To overall goal of this workpackage is the development of appropriate strategies for caching, replicating, and proactive dissemination of data to the overlay computers to improve performance, availability, fault-tolerance, and data freshness. Replication and caching are well-studied problems in distributed computing. They offer distribution transparency, scalability, and fault-tolerance. In global computing, new issues arise due to the dynamic nature of the environment. Furthermore, there is a close interplay between replication and the structural properties of the overlay network.

The determination of the number of replicas created for each data object can follow different strategies based on the popularity of the object. A strategy that does not take into account the query workload is uniform replication where a constant number of replicas is created for each object irrespectively of its popularity. Proportional replication creates for each object C replicas, where C is proportional to the number of queries for the given object. Between these two strategies, square-root replication [CS02] provides better results leveraging the effort for finding popular and unpopular objects. With path replication, copies of an item are stored at all peers along the path from the requestor peer to the provider peer [LCC+02]. Path replication outperforms owner replication where copies are stored at the requestor node only. However, path replication tends to replicate items to peers that are topologically along the same path, which hurts somewhat the performance.

A second major dimension in caching and replication are judicious strategies for disseminating updates. This is an issue that is still poorly understood and requires deeper investigation. Such strategies can be based either on the paradigm of detecting and invalidating stale data and then updating on demand or make use of proactive dissemination of updated and new data items. Epidemic algorithms and investment-based policies are of particular interest in the context of very-large-scale overlay systems [RB03, LDP04].

The above issues arise across a spectrum of different kinds of data items: from simple unstructured data containers like files and Web documents to relational or XML databases, and also for metadata and ontological meta-metadata. Moreover, caching can even be applied to computational results such as output parameters of Web Service invocations (if there is a sufficiently high probability of repeated invocations with the same input parameters and no changes to the underlying persistent state between service invocations). One interesting issue is to investigate to which extent the strategies for caching, replication, and proactive dissemination depend on and may be tailored to the kinds of data items and their corresponding access characteristics. Of course, all viable strategies need to cope with the large scale and high dynamics of overlay computers.

The research work within this workpackage will focus on the definition of mathematical models and methods for choosing replication degrees and replica placement, and on the design, theoretical analysis, and experimental validation of algorithms for dynamic load-adaptive adjustments, replica maintenance and proactive dissemination.

WP 3.2 Load management

Leader: UPB; Participants: CTI, CTI, MPII, CAU, UNIPD, ETHZ, UPC, URDLS

Load balancing for overlay computers. Load balancing is one of the key problems that must be addressed to efficiently use an overlay computer for computation- intensive applications. For this purpose, an application must be divided in several subtasks and these subtasks must be executed on different computing nodes of the overlay computer. Subtasks can either be run independently from each other or, if required by the application, interdependencies between them have to be obeyed. In the latter case, the nodes must use communication in order to exchange intermediate results. In summary, the load balancing problem aims to the following goals:

- The total execution time of the application should be minimized.
- All computing nodes should operate during the whole computation and idle times should be avoided.
- The total load in the system should be distributed equally and 'fairly' among the nodes.
- The communication overhead between the nodes should be minimized.

In the past, the load balancing problem was mainly considered in the context of parallel computers where a central control instance exists. Because of the size and the dynamic nature of an overlay computer, there can not be a central authority responsible for the load balancing task. One way to deal with this limitation is the use of local iterative load balancing algorithms like diffusion or dimension exchange schemes [Cyb89], [MGS98], [DFM99]. The balancing flow quality as well as the convergence behavior of such algorithms strongly depend on the structural properties of the underlying overlay/global computer.

It is known since several years that in large self-organizing global computers (like the Internet, Telecommunication networks etc.) the node degrees follow the so called "scale-free power law" distribution [BAJ99], [KRR+00], [MP02]). $P(d)/n \sim d\text{-}\gamma$ where P(d) is the number of nodes with degree d, n is the total number of nodes in the global computer and γ is a constant, which depends on the structure of the global computer. We plan to investigate the impact of the node degree and size on the performance of load balancing. Such questions have not been investigated sufficiently until now. More basic research has to be done on the properties of such power law graphs in order to obtain scalable load balancing algorithms. By following this approach, we will attempt to estimate the quality of basic distributed algorithms for load balancing and we will develop new efficient ones.

Adaptive Software. An overlay computer infrastructure may be employed to provide dependable, cost-effective access to high-end computational capabilities irrespectively of their physical location or access point. This is, for example, the case of grid environments, which enable sharing, selection, and aggregation of a variety of geographically distributed resources (e.g., supercomputers, computer clusters, storage systems, data sources, instruments).

In such a scenario, many different platforms can be available to run a single application. For load management reasons, the specific platform (the specific set of computing nodes constituting the overlay computer) onto which an application is ultimately run, may not be known at design time. Hence, it is desirable to endow the application with built- in parameters whose judicious choice at run time allows the software to match adaptively the structure of the actual platform where it runs, so to harness its computational potential at best. The ultimate goal is to reduce software development and tuning costs, while retaining the ability to effectively use the available computing infrastructure.

There are two major stages in the development of adaptive software:

Parameterization: the identification of a suitable vector P of algorithm parameters, which vary in suitable ranges to realize a spectrum of tradeoffs among different machine-resource requirements.
Tuning: given a target machine M, the choice of a specific value P* for the parameter vector P, with the objective to minimize the running time T(M,P).

The field of adaptive software has become popular in recent years, due to the success of packages such as FFTW [FJ98] for Fourier transforms and ATLAS for linear algebra [CPD01]. However, these and similar packages heavily rely on problem-specific intuition and knowledge on the part of the designer. In particular, parameterization and tuning are pursued in an ad hoc, problem-dependent manner. Currently, there is no general methodology on how to parameterize an algorithm to make it adaptive, and only preliminary results are available on how to search the parameter space [Y+03]. Our research within this workpackage aims to the development of a general framework for the design and implementation of adaptive software, which can guide both the parameterization and the tuning stage.

Peer-to-peer-based Parallel Computing. Another research direction within this workpackage is motivated by the fields of parallel computing and peer-to-peer networks. Peer-to-peer networks can be used to provide an enormous computing power. But current techniques do not use this power effectively because of the problems that emerge in such systems: peers may join or leave such a system at any time, they may fail, they may have different and even time-varying speeds and capabilities, they are not permanently available, etc. Also, executing code at a remote machine may be risky for the owner. In this workpackage, our goal is to explore ways for efficiently executing parallel programs on such peer-to-peer networks.

A very successful programming paradigm in the area of parallel computing is Valiant's Bulk Synchronous Parallel (BSP) bridging model. Within this model, one can write parallel algorithms that can be efficiently executed on parallel machines provided the BSP model's primitives are realized efficiently.

The Java programming language supports the secure execution of code by its sandbox principle, and it provides a uniform view to the programmer although it is implemented on different platforms. We will adapt the BSP paradigm so that parallel programs can be executed efficiently on a large-scale distributed peer-to-peer network. New scheduling and load balancing strategies and fault-tolerance mechanisms are challenges to be tackled in order to design such a system. A system based on the principles discussed above will be implemented in the context of WP6.2 (see the discussion of subproject SP6).

WP 3.3 Scheduling

Leader: CAU; Participants: CNRS, UPB, CTI, UNISA, MPII, UOA, ETHZ, UDRLS, DIM, UCY

The main goal of this workpackage is to provide efficient algorithmic solutions for scheduling problems arising in individual layers as well as at the interlayered framework of an overlay/global computer. It will address both low-level scheduling (i.e., scheduling of computation) and application-level scheduling (e.g., scheduling of data requests).

In the research area of scheduling problems, we plan to study the multiprocessor task scheduling problem, the strip packing problem, the problem of packing rectangles with profits into a rectangle, the batching problem with costs and preemptions. These combinatorial problems directly model several critical problems such as scheduling computation tasks in the nodes of an overlay computer or queueing requests for information resources provided by an application.

We primary aim to time-efficient algorithmic solutions for our problems in the offline setting with a single objective. However, to deeply understand real-world applications, we also plan to investigate more complex settings. In the multi-objective setting, it is required to provide a solution which is simultaneously efficient with respect to several objectives. Contrasting to the offline setting where the scheduler has full information about the problem, in the online setting, information becomes available during the scheduling process just piece by piece, and, respectively, it is required to provide an online solution.

The nodes of an overlay computer typically communicate with each other by exchanging information through a logical communication network. Here, efficient handling of information flows becomes important. Because of the size and ever-changing nature of such systems, there cannot be a central authority controlling the information flows in the entire system. These tasks are usually performed by computing entities distributed at the nodes of the overlay computer. Such entities may belong to different owners (e.g., different organizations, companies, etc.) which aim to optimize their own (typically conflicting) objectives. There has been some recent work on the theory of selfish agents modelling selfish behaviour of computing entities in static environments; see for example [CV02, FGLMR03a, GLMM04a, KP99, RT02], and [FGLMR03b] for a recent survey. However, the impact of selfish behaviour on dynamic overlay computers is not well-understood. We will attempt to develop a fundamental theory on the impact of selfish behaviour on dynamic overlay computers.

Our departing point will be the notion of congestion games [Ros73, MS96, CKN04], a generalization of scheduling problems where each task needs to be executed in a subset of resources and the delay in each resource is a function of its load. Congestion games are good models for scheduling at the communication level of global computers. However, most of the scheduling problems in relation to congestion games are open; their special case where each task requires only one resource is the classical model of task scheduling and it is well studied. For example, the well-understood classical task allocation problem of identical tasks corresponds to congestion games of the simple class of networks with parallel links [KP99]. Its generalization to more complex models is a novel type of scheduling problems which, despite its importance and some recent negative results [FPT04], is still open (new methods in integer mathematical programming seem to be required for its solution).

To further complicate the situation, congestion games are not in general adequate to model neither scheduling computations in overlay computers nor the dynamic situation of scheduling at the application level. We plan to find appropriate extensions to capture this dynamic nature of the problem and to study the associated algorithmic problems.

WP 3.4 Workflow and Services

Leader: ETHZ; Participants: UOI, MPII.

The approach that will be adopted in this workpackage is to treat large-scale computations as processes similar to those found in workflow management systems. The notion of a process allows us to model sequences of invocations of computer programs and applications in a distributed and heterogeneous environment as well as to capture the corresponding data exchanges between these programs. From here, the process can be encoded in such a way so as to allow its efficient persistent storage, allowing us to both automatically manage the computation and increase its dependability. It will also allow sharing these computations using a common model where the data and control flow are separated from implementation information such as the specific applications being invoked and their locations.

The idea follows the model already adopted in other initiatives where computational and data services will be available using Web services (SOAP for communication, G-WSDL or variations of WSDL to describe the services). The project will extend these efforts by exploring extensively how much the notion of computational processes (in parallel to business processes) can be applied in practice and be autonomically supported. If successful, we will then be able to provide scientists with the same tools and functionality used in modern electronic commerce and enterprise application integration. Our work in this area may even serve as the foundation for further efforts in the creation of vertical standards that provide a set of well defined processes that can be used within concrete branches of science (e.g., standard processes for ray tracing, for genomic data cross-comparisons, for computational fluid dynamics in the aerospace industry, etc.).

In this workpackage, we shall explore (i) the use of services for accessing data in overlay computers, (ii) the use of workflow for the definition and execution of composite web services, and (iii) the capabilities of a global computer and its overlay structures for dynamic self-configuration and automatic adaptation to system, workload, and applications dynamics. The idea is to combine data services and computational services under a single interface: Web services so that the composition of these services can take place using a process model that separates the specification of the program from its implementation. During the project, we will develop new methods and improve existing techniques for automatic migration of failed tasks, system awareness (to node configurations, network bandwidth, future load, new clusters becoming available or unavailable), replication of system critical information, general fault tolerance, and the dynamic and automatic reconfiguration of global workflows to such rapidly evolving conditions.

All these functions must be able to handle heterogeneity in the underlying resources and "semantic" heterogeneity in the sub-processes of a global workflow (e.g., workflows that span different workflow engines and computational environments) and also the dynamic changes in the process definitions while already executing, as arising, for example, in large-scale logistic applications and in many kinds of e-science collaboration.

The work will incorporate the developments made in other workpackages such as WP3.1 (distributed data management), WP3.2 (load management), and WP3.3 (scheduling). Techniques developed there will be incorporated in an autonomic computing system that will not require manual intervention to keep computations alive regardless of the complexity of the run-time environment or the duration of the computation.

The workpackage is divided into the following tasks:

- Specification of a generic engine for specifying computations of overlay networks of Web services and additional services (such as local applications). The engine should be general in that it should not tie the workflow language to any particular type of service or interaction (unlike, e.g., BPEL and BPEL implementations do). The engine should also be flexible enough to incorporate new services as they appear and, therefore, must have an appropriate model for services so that these are treated as pluggins that can be dynamically added to the system as needed.

- Extensions to the basic engine to support autonomic execution. One of the biggest problems in current grid software is scalability. The engine we propose will be modular and capable of dynamically expanding it modules across a cluster to accommodate varying demands. This task will work both on the engine itself and the policies that are most adequate for autonomic adaptation.

- Monitoring, data lineage, and system management support. Computations, particularly scientific computations at global scale are important data units on their own. A system supporting such computations must be able to answer queries on the state of each computation, on how a particular data item was produced (lineage), automatically recompute data items if there are changes on the input data or the algorithms used in the computation, and provide users with the ability to drive and control these computations in real time. This task will specify the necessary infrastructure to perform these operations on top of the engine described in the two previous tasks.

- Dynamic self-configuration of global workflows and dynamic adaptation to changing conditions regarding computing resources, workload, quality-of-service (QoS) requirements of the application, and the definitions and goals of the global workflows themselves. These capabilities must be provided in a highly heterogeneous environment that involves different workflow engines and other service components. A promising paradigm for addressing this work is to apply control-theoretic principles to overlay computers.

A demonstrator that shows the developed solutions to the above issues will be developed in the context of subproject SP6 (workpackage WP6.2).

Sub-Project 4: "Security and Trust Management"

Leader: UNISA; Participants: CTI, UOI, UDRLS, KUL, CYB, DIM

The sheer global scale of a global computer imposes a new look at the related security issues. The traditional solutions will simply not scale to satisfy the needs of a global computer and new security threats will become relevant as more and more of everyday activities will relay on the global computer.

In the last twenty years, security research has concentrated on how to protect information of one *domain* (be it a large private or public organization or an individual). The protected information was meant to be produced and consumed within the same domain and each domain was seen as a stand-alone fortress with limited exchange of data. Within a domain, computation was carried-out mostly on data originating from the same domain; this fact has the obvious consequence that the data is available in clear form to the application. In some cases, a domain would make a static subset of the data available for access to applications from other domains.

This, by now unrealistic, conceptual scenario greatly simplified the security problems in the same way in which the (physical) security of a medieval castle is conceptually simpler than the security threats with which a modern city has to deal. As a first observation, we notice that, since applications accessing the data belong mostly to the same organization as the data, it is possible to have ad-hoc access policy wired- in into the application. Moreover, in such a scenario, all agents accessing the data belong to the same small *name space* and it is thus natural to base access policies on identity. As we will discuss below, this approach does not scale when the domain becomes large or when we need to consider agents coming from different domain.

The model based on a single domain does not faithfully describe the current state of things as in a global computer applications from different (and even competing) domains wish and need to interact with each other to a larger degree than before. For example, it is not unusual for different domains to have to perform computation on sensitive data that the owner does not want to disclose in full. Thus one wishes to compute relevant facts about distributed data sets while concealing everything else. As the user base of a data set expands outside the domain that owns that data set, it is even more crucial than before able to deploy authentication mechanisms to be able to control access to the data (and in general to all the shared resources). Indeed, whereas agents and applications belonging to the same domain could be considered trustworthy, this assumption is unrealistic in large global scenarios.

As the global computer spans over more and more domains, it is likely that a large percentage of the activities of one individual will be conducted within the global computing and thus it would be possible to trace them for profiling one or more intrusive purposes. The problem of privacy and anonymity which was marginal, thus becomes central in a global computer scenario.

In sum, we are witnessing a paradigm shift in the field of security from a scenario in which ad-hoc solutions were sufficient to address the security threats posed in a relatively small-scale environment to a new global scenario in which the previous methods and techniques yield non-scalable and inefficient solutions. This modified state of things calls for new theories for understanding the security issues arising in the global setting and proposing robust and scalable solutions.

Given the conceptual complexity of the global computer environment, one would like to abstract away this layer of complexity and to enable developers and software architects to focus only on the application to be developed. The overlay computer will instead *transparently* provide the security functionalities needed for the application to be securely run in a global environment. As an added benefit, functionalities provided by the overlay computer will not have to be duplicated in each application.

For a concrete example (that will be discussed in more detail in the following sections) suppose that a distributed application for mining knowledge from several proprietary data sets is to be developed. The overlay computer will enable the application developer to concentrate on the design of the application as if privacy of the data base were not an issue. It will be the task of the overlay computers to wrap the application in such a way that privacy of the data sets is not compromised. Moreover, the security must be guaranteed even when the distributed data mining application is run concurrently with several others applications working on the same data sets. At the same time, to be of any use, the overlay computer must offer *scalable and robust* implementations of the functionalities. It is thus important to concentrate research efforts on the efficiency of the underlying protocols and algorithms and to develop computational and algorithmic theories that describe and model security issues in the new conceptual scenario of a global computer.

In summary, the research work within the workpackage consists of the following workpackages; these are described in more detail below:

WP 4.1: Trust Management;
WP 4.2: Privacy, Identity and Anonymity
WP 4.3: Secure Distributed Computation

WP 4.1 Trust Management

Leader: UNISA, Participants: CTI, UOI, KUL, CYB.

In a global scenario it is unrealistic to assume that parties behave as prescribed. This is particularly true with respect to access to controlled resources and to restricted information. A traditional approach consists of first assigning and managing identities (*identification*), and subsequently making authorisation decisions based on these identities (*access control*). However, managing identities in large scale networks is known to be non-trivial, and in an environment consisting of programmable overlay computers these problems will become even harder. Moreover, existing authorization mechanisms fail to provide tools powerful enough to handle problems at the scale necessary for global computers.

Identity-based Authentication

We first focus on the process of identification. During the identification phase the requester proves his *identity* by exhibiting *certificates*. A certificate links an identity to a public key by means of a signature and the steps taken during the identification phase can be described as follows.

1. Obtain certificates and verify signatures on certificates.
2. Verify that certificates have not been revoked.
3. Find a *trust* path from the trusted certifier to each of the certificates received.
4. Verify that the requester is the legitimate owner of the certificates; this step typically consists of proving knowledge of the corresponding private key by means of a simple two-party challenge-response protocol.

Once the identity of the principal has been established, permission for the requested action is granted or denied by looking up names in a database (called the *access control list* or ACL) which is either explicitly constructed or implicitly defined by the application that manages the access to the resource.

An important conceptual point in an identity-based approach is what is to be considered a *trust path*. Within the framework of identity-based authentication, one can distinguish two main approaches to this problem. The most popular approach is the one of PGP. Here if user *A* has a copy of user *B*'s public-key for which he is confident that the key has not been tampered, *A* can sign the key and pass it on to user *C*. It is then up to user *C* to decide whether user *A* is to be trusted and accept user *B*'s signed key. Instead in the Public Key Infrastructure based on X.509 (see [X509]), only so-called *Certificate Authorities* can sign user keys. If *A* needs *B*'s key and both have been certified by the same *CA* then *A* can directly verify *B*'s keys. On the other hand if *B*'s keys has not been signed by the same CA as *A* then a *certification path* must be created from *A* to *B*. This consists of a sequence CA_1, $cert_1$, …, CA_l, $cert_l$ where, for $1 \leq i < l$, $cert_i$ is the key of CA_{i+1} signed by CA_i and $cert_n$ is *B*'s key and CA_1 is a *CA* trusted by *A*. Thus the X.509 PKI assumes that *CA*s are organized into a global authority tree and that users that need to interact must have keys signed by *CA*s in this global tree.

Weaknesses of the identity-based approach. As discussed above, a certificate only carries an identity (a *distinguished name* in X.509 parlance), a public key and some auxiliary information on the owner of the certificate and the key along with a signature. This approach to authentication does not scale to the size of a global system for several reasons that we briefly discuss below.

1. Certificates and access control list are only capable of binary decisions while, in more sophisticated applications, security policies that cannot be expressed in term of ACLs are needed. As a consequence, the policy itself tends to be hard-coded into the application.

   For example, suppose that an employee of a large organization is authorized to sign purchase orders for up to 1000 Euro or up to 10000 Euro if the order is co-signed by another employ. Obviously, such a policy does not have anything to do with the identity of the employee and using identities only force one unnecessary level of re-direction through an ACL that lists the identities of all employees authorized to sign purchase orders of a given amount.

2. The approach based on identity-certificates and access lists only allows simple forms of delegation. Indeed, in a hierarchical organization, policies will be specified at the last step in the chain of delegations. The effect is that high-level entities cannot specify security policies leading to potential inconsistency among local policies.

One instead would like to still be able to have some high-level policies to be enforced throughout the organization.

3. The ACL usually consists of a very large sparse matrix, which is difficult to manage and to distribute over a network.

4. A more fundamental flaw in the identity-based approach is that the idea of a distinguished name is not likely to occur, as it requires a single, global naming discipline.

   The assumption that a certificate could bind a name to a key would certainly make sense when it was first proposed. In the 1970's and even through the early 1990's, relationships were formed on a personal base whereas with the explosion of global systems it is likely that one will encounter key holder that are complete strangers in the physical world and will remain so. This is even more evident when we consider mobile code like Java applets or JavaScript.

Trust negotiation mechanisms

A different approach has been recently proposed by several researchers. It is based on the idea of linking keys directly to authorizations to perform specific tasks introduced in SPKI [SPKI] used in conjunction with flexible mechanisms for specifying security policies. Commercial tools are available for helping the transaction from identity-based authentication to the role-based approach [rolemining1, rolemining2].

This approach named *distributed trust management* [DisTrusMan] avoids resolving identities and it allows for much more flexible security policies than the identity-based approach. Going back to the example in the previous section, once a certificate does not carry just the identity of the owner but can also specify predicates that describe the actions for which the owner of the certificate is trusted, it becomes straightforward to have a certificate that authorizes for orders up to 1000 Euro and 10000 Euro in conjunction with another key.

The heart of such an approach is a compliance checking algorithm that, given a request (for example, a request to modify a record in a database), a policy (expressed using a high-level language) and a set of certificates, returns yes or no, depending on whether the certificates constitute a proof that the request complies with the policy. Of course, the newly acquired power to specify sophisticated security policies comes at a price: efficiency. Efficiency is an important issue with respect to the compliance checking algorithm as an inefficient algorithm will degrade the performance of the authentication subsystem and thus of the whole global system.

Trust negotiation systems, however, by their nature, may represent a threat to privacy in that credentials, exchanged during negotiations, often contain sensitive personal information that needs to be selectively released. Also, a user may want to minimize the released information, thus enforcing the need to know principle in disclosing his/her credentials to other parties. In other situations, a user may want to carry out negotiations that cannot be linked to him.

Reputation based authentication and authorization

In order for a system to be scalable, some security problems cannot be solved by manual registration and revocation of each single key and/or identity. Many peer-to-peer and open systems are already reputation-based (think for example of e-bay). One can anticipate that in the global computing paradigm this approach will take an increasingly important role. However, there is a need for better understanding the properties of reputations systems such as their configuration and development and their robustness to malicious attacks. Some work along the same lines has been recently presented in [MorselliP2Pecon] where a game-theoretic framework for analyzing the robustness of trust- inference protocols in the presence of adversarial (but rational) users is presented.

Access Control Models and Mechanisms

The area of access control has been widely investigated in the past years and several models have been defined, for different data management systems. The various proposed models are quite rich, they often support object and subject hierarchies, and they provide positive/negative authorizations, exceptions, as well as authorization derivation rules. In addition, research efforts have been devoted to investigate the expressive power and relevant properties, like safety, of access control models. However, all these proposed models and corresponding mechanisms are inadequate for globally interconnected computing systems, especially the ones characterized by a large number of moving entities.

WP 4.2 Privacy, Identity and Anonymity

Leader**:** KUL; Participants: UNISA, CYB

A wide range of privacy-enhancing technologies have been developed in the last years in order to make systems that comply with the privacy requirements. Some of these technologies can be used as building blocks for secure systems. Below we summarize the most important available technologies and indicate which research problems (if any) need to be tackled.

Pseudonym systems. Pseudonym systems [CH90] were introduced as a way of allowing a user to work effectively, but anonymously, with multiple organizations. Each organization may know a user by a different pseudonym, or nym. In this case, we will have a nym associated to each domain name. These nyms are unlinkable: two organizations can not combine their databases to build up a dossier on the user. Nonetheless, a user can obtain a credential from one organization using one of his nyms, and demonstrate possession of the credential to another organization, without revealing his first nym to the second organization.

Pseudonym systems have been used for a long time in the real world. Informally speaking, we could say that a pseudonym is an identity for a user that is used in a particular context. The pseudonym can be unlinkable to the real identity, and then, by using a pseudonym, the user keeps secret his real identity. Pseudonyms are a powerful and flexible tool to provide anonymity in different systems. The simpler systems rely on Trusted Third Parties, and the more sophisticated ones use public key cryptography, one-way functions and more complex protocols.

While this is an essential building block, in any more complex system, tools are required for managing nyms. However, in this case it is better to use more sophisticated variants such as private credentials.

Private credentials / Anonymous certificates. Privacy protection requires that each individual for him or herself has the power to decide how his or her personal data is collected and used, how it is modified, and to which extent it can be linked - only this way individuals can remain in control over their personal data. There are basic privacy-enhanced technologies available that are entirely feasible and secure for achieving these goals. In some of these technologies any user's secret can only be computed with the consent of that user, when other technologies use the blinding techniques only, or both. Some technologies use self-revocable unlinkability and untraceability, where certificate holders can prove they have been the originator of a showing protocol execution, they can prove that they can provide information and can prove that they were not involved in certain transactions. It can be constructed with highly practical digital certificates that fully preserve privacy, without sacrificing security.

Research in this area is needed to investigate and develop new mechanisms for private credentials that offer an improvement in security (weaker or different security assumptions), functionality (increased flexibility of properties that can be proved, identity escrow, …) and/or efficiency (computational overhead, storage, …).

Identity management. Identity Management [HB03] is the management of secure access to information and applications across large-scale networks. Identity management presents the challenge of managing who has access, what levels of access are granted, and how to control that access without undermining the security of the networks or secure information within those networks. Also, by managing her identity, roles, and personal data, a user may decide whom to give which data, when to act anonymously, when to use a pseudonym, when to authenticate herself etc. Ideally, the different pseudonyms which represent her identity or roles can be linked only if the holder so desires.

An identity management system empowers the users to maintain their privacy and control their digital identity. Also, identity management enhances the usability of systems introducing techniques such as single sign-on. Regarding the limitations of these systems, it is doubtful that users can fully assess the implications of assuming a particular identity/role in a transaction. There is a clear link with the problems and approaches discussed in WP4.1.

Anonymous communications. In order to achieve privacy protection at the application layer, it is often required to anonymize the communication layer. The reason for this is that the user's IP address is visible to any observer of the communication: if an anonymous communication infrastructure is not underlying the application layer, the privacy of the user may be threatened by the leakage of information at the communication layer. Systems that require authentication can implement this mechanism at the application layer, even if the communication infrastructure is anonymized.

Some anonymous communication systems have been implemented using peer-to-peer models (Tarzan [FM02], Herbivore [GRPS03], Crowds [RR98], GNUnet [BG03], MorphMix [RP02], Tor [DMS04] etc.) or multicast protocols (e.g., Hordes [SL00]).

Mix networks [DS04, RSG98] are a popular building block used to implement anonymous communications. The mix takes a number of input messages, and outputs them in such a way that it is not possible to link an output to the corresponding input. In order to achieve this goal, the mix changes the appearance (by encrypting and padding messages) and the flow of messages (by delaying and reordering).

However, anonymous communication infrastructures have not been widely deployed even in existing networks. One of the problems that need to be solved before these anonymous networks become acceptable in society is the lack of control mechanisms in case these systems are abused for fraud or other criminal activities (see e.g., [CDG+03]). Moreover, there is substantial research needed on performance, scalability and manageability of anonymous communications systems. Quantitative methods have been introduced only recently in this area; initial models only study what is happening in very simple models of individual mixes or a few nodes, while there is a need for quantitative methods that can handle large scale systems. Finally, for mobile communications and mobile agents, fundamental research is needed on techniques for location privacy.

Privacy and individual databases. Today's globally networked society places great demand on the dissemination and sharing of person-specific data. Data holders, operating autonomously and with limited knowledge, are left with the difficulty of releasing information that does not compromise privacy, confidentiality or national interests. For example in a medical setting this is of particular interest due to privacy issues and to prevent possible misuse of confidential information.

In spite of the importance of this problem, only a few research papers on this topic could be found. It seems that there is still not a good scientific approach to the problem and all described solutions are ad-hoc. Some examples are "DatAnon," that uses Anonymous Database System (see http://www.datanon.net) and μ-Argus System a 5th EU framework project Computational Aspects of Statistical Confidentiality (http://www.cbs.nl/sdc/argus.htm and http://neon.vb.cbs.nl/casc/).

One solution for the privacy of the data requestor is the so-called PIR (Private Information Retrieval) scheme [BIM00, CGKS95]. A PIR scheme enables a user to retrieve an item of information from a publicly accessible database in such a way that the database manager cannot figure out from the query which item the user is interested in. The PIR model is not concerned with protecting the privacy of the data and allows the user to learn arbitrary additional information. The stronger SPIR (Symmetric Private Information Retrieval) primitive requires, on top of the PIR requirement, that the user learn no additional information about the database other than the selected bit. SPIR may be viewed as an analogue of Distributed Oblivious Transfer. It has been proven in [IK04] that SPIR can be used as an intermediate primitive for establishing the connection between PIR and MPC (Multi-Party Computation, see also WP 4.3). Further work is required to develop more powerful models that capture additional requirements (one special case is data-mining) and that allow for more efficient implementations (in terms of round complexity, communications complexity and computational overhead).

Privacy-preserving data mining. Data mining is a recently emerging field attracting the attention of researchers from the domains of Database Systems, Statistics and Artificial Intelligence. Even though current storage and network technology make it possible to gather large volumes of data and to share it with other organizations, this mass of information is useless if meaningful knowledge cannot be efficiently extracted from it. Data mining aims to answer this need. Again, a key problem that arises in this field is the problem of privacy, which is sometimes regulated by law (for example, for medical applications) or is dictated by business convenience.

In general the question is "How do we mine data that we are not allowed to see?" The security and privacy implications of data mining have been recognized since the early days of data mining and some heuristic solutions have been proposed (see, for example [cliftonmarks, datamining1, datamining2]) based on *data-perturbation* or *data-sanitation* so that *planned* queries could still be executed. The proposed heuristics did not fit the global computer scenario as 1) we want parties to be able to perform data mining on our database but we just want them to learn no more than the result of query (and what can be inferred from that); 2) planned queries are not known in advance so that the database can be enlarged or data be regrouped so not to affect the planned queries.

The problem has been recently addressed in [lindellpinkas, lindellpinkas2] in which the problem is cast within the framework of secure distributed protocols (see also WP 4.3) in the following way. Denote by A and B two parties owning large private databases D_A and D_B that wish to run a data mining algorithm on the database $D_A \cup D_B$ without revealing any other information about their databases except what is learned by the result of the algorithm and the knowledge of his own database. This is clearly a special case of secure distributed computation. However, to be of any practical relevance, a protocol for private data mining should allow each of the parties to do most of their work by itself and the number of rounds of communication should be limited.

A new tool for performing anonymous transactions called Anoniminer has been recently presented at the Rump session of Crypto'04 [Moti_ramp_Cripto04]. The Anoniminer is built on the top of Group Signatures and Anonymous Oblivious Data Mining Engine. The Oblivious data-mining engine performs an SPIR transaction between the requestor and the database. A new group signature scheme introduced in [KTY04] provides a mechanism which allows the selective linking of the existing signatures of a misbehaving user without violating the privacy of law-abiding group members. Such signatures are suitable for coupling with anonymous oblivious data mining engines. Further work is required to make these techniques more efficient and to improve their scalability.

WP 4.3 Secure distributed computation

Leader: UNISA, Participants: CTI, UDRLS, KUL, DIM, CYB.

In a distributed setting, it is often necessary to compute a function f of n inputs $x_1, \ldots, x_n$ each owned by a different party. In the unrealistic scenarios in which parties are willing to make their inputs public, or one incorruptible trusted party is available, this is a trivial task. Indeed, if inputs can be made public then all the parties can just broadcast their input to all other parties and then each node will compute the value $f(x_1, \ldots, x_n)$ by himself. On the other hand, if a trusted incorruptible party is available then parties could just send their input to her. The trusted party will then compute $f(x_1, \ldots, x_n)$ and broadcast it to all nodes. In a real life scenario instead, it is very unlikely that parties are willing to make their inputs available to all the other parties in the system and, even more, it is very difficult to identify a trusted party that will correctly compute the result. Consider, for example, the important case of an election. Here, one would like to have a protocol that correctly counts the votes for each candidate in such a way that the final count can be verified and still the security of each voter is preserved. The two properties of security and correctness should hold also if a subset of the parties that are involved in the protocol coordinate their actions and share the messages they receive during the execution of the protocol in order to find out how one person voted. More precisely, we require that in a t-secure protocol no coalition of size at most t can obtain any information other than the output of the computation and what can be *logically inferred* from the output of the computation and their own input.

The importance of designing protocols for secure distributed computation has been recognized since the 80s starting with Rabin's oblivious transfer protocol (see [RabinOT]) and culminating with the completeness results of [Yao1, Yao2, GMW1, GMW2, CCD, BGW] which prove, under different assumptions, that any efficiently computable function can be computed in a distributed way so that the result can be verified to be correct by all parties and the privacy of the input of each participant is guaranteed. These general results, although of crucial importance for understanding the limitations of the model, do not directly give usable solutions to practical protocol design problems, say elections. Indeed they are to be considered more as "plausibility" results and leave open the problem of designing practical protocols for important tasks. Recently, research devoted to concrete practical *implementation* of the general results has been undertaken. For example, the *FairPlay* system [FairPlay] gives an implementation of two-party secure distributed protocols along the lines of [Yao2].

Even though it is reasonable to assume that at any given time no more than a certain number of sub-servers are under the control of the adversaries, it might be well the case that, all sub-servers are at different times under the control of the attacker. This is the case, for example, when the attack is mounted by using a *worm*. One way to counter this important type of attack is to design the protocol in such a way that information gathered at different times cannot be pooled together to obtain sensitive information using advanced secret-sharing algorithms [HJKY95]. Protocols with this property are called *proactively secure* distributed protocols [virus, intrusion]. A complementary approach consists in developing efficient distributed algorithms for eradicating a worm from a network by using antivirus mobile agents [spirakis, spirakis2].

The spreading of a worm from machine to machine and the counterattack can be modeled as a distributed game over a graph (representing the network under attack by the worm). Members of a team of guards (i.e., the antivirus mobile agents) traverse the links of the network in pursuit of the fugitives (clones of a worm) which move along the links of the graph without any other knowledge about the locations of the guards than whatever they can collect as they move. The purpose of the fugitives is just to read local information at each node and to stay in the network as long as possible. When a guard meets a fugitive the fugitive is destroyed.

There are several global-scale applications where (in addition to origin authentication) it is necessary to determine whether a data set existed at certain time, or whether an electronic record has been maliciously altered. The conventional naive hash-and-sign time-stamping solutions are neither scalable (service rate of the existing key-based time-stamping services does not exceed few hundred stamps per second, which is obviously insufficient for a global computer) nor sufficiently trustworthy (as keys can be abused by the trusted party for back-dating documents). Keyless time-stamping schemes were first proposed way back in 1990 by Haber and Stornetta [haber91]. Since then several improvements in efficiency have been presented (e.g., [BuSa04]). Considering the increasing use of time-stamping technology, more intense studies about the security and scalability of time stamps are necessary.

Secure distributed algorithmic mechanism design. Another important area of applications of secure distributed computation is related to the work on distributed algorithmic mechanism design. In this scenario, we have parties that are neither trustworthy nor adversarial but they can be assumed to be selfish and to respond to incentives. The approach adopted by computer scientists designing protocols for selfish parties is consistent with the approach in the economics literature and consists in designing so called truthful mechanisms in which parties have an incentive to correctly reveal their inputs and then an optimal choice can be made. For example, distributed algorithmic mechanism design is applied to the problem of inter-domain routing where parties are autonomous systems and the party private input is the cost it incurs when carrying transit traffic. Revealing this cost may give away details about their internal network and autonomous systems may be very reluctant to give this information to their competitors.

Composability. A foundational challenge in this research area is to define an appropriate mathematical model for representing protocols, and then formulating, within the model, a definition of secure protocol. Once a satisfactory definition of security is obtained, a protocol can be proved secure by showing that it satisfies the definition of security. However, until recently, protocols were defined as stand-alone objects and their security property would only hold if the protocols were run in isolation. This approach may be sufficient for small-scale scenarios. Instead in large-scale scenario, it is important that secure protocols can be arbitrarily composed of other protocols in a security-preserving fashion. Composability allows for a modular approach to the design of secure protocols, which is essential for building large and complex protocols out of simple basic sub-protocols. One way to capture security concerns of specific environments is to directly present the given environment with an extended security notion once the original definitions of security (and thus the relative protocols) were shown to be insufficient in more complex systems (for example, two-party protocols that remain secure even if they are executed concurrently [DNS, KPR2, KP]). This approach is not completely satisfactory as it only addresses specific environments and security threats (concurrent execution of the same protocol). Recently, in [Canetti], formal definitions for the concept of a *universally composable secure protocol* have been proposed and the first steps towards realizing such protocols have been taken. Currently, under strong set-up assumptions that a trusted party is available to bootstrap the system, it is possible to design composable protocols for general secure distributed computing [composable1]. Although the role of the trusted party is limited to the bootstrap phase, this is too strong of an assumption to be made in a global computer and a protocol designed under this assumption will not be of much help in the design of an overlay computer.

A parallel line of research on universal composability (this time called simulatability) [PfitSchuntWaid:2000, PfitWaid:CCS2000] has begun at about the same time as [Canetti]; the security definitions of these two approaches have pretty much converged for now. This line of research has produced a universally composable cryptographic library [Backes:cryptolib] with nonce generation, asymmetric encryption, and digital signatures as its most impressive result.

Further applications. Secure distributed protocols can also be used to support fault-tolerance and heterogeneity of users, which are two important aspects of a global computer. In general, secure distributed computation can be used to allow parties of a global computer with limited computational or storage power to delegate other more powerful entities to perform computation on their own data without having to compromise the confidentiality of the data. Recently, for the specific case of keyword search, an efficient and practical solution has been given in [BonehEC04].

Within the domain of fault-tolerance, secure distributed protocols can be used to “distribute” the work of a server (for example, a domain name system server in the Internet) across several *sub-servers* so to guarantee availability and integrity of the service despite some servers being under control of an attacker or failing in arbitrary ways (see for example the work done in MAFTIA [maftia] and [Cachin]). As parties cannot be trusted, the sharing must be done in a secure way so that sensitive information about the service is not revealed to any of the sub-servers. Fault-tolerance techniques coupled with a security approach can be used also to prevent *Denial of Service* attack consisting in making the service unable to perform at all.

Sub-Project 5: “Extending Global Computing to Wireless Users”

Leader: CTI; Participants: UOP, UOI, UDRLS, URDTV, UPB, CAU, CNRS, INRIA, CUI, KUL, UPC, UCY, DIM

Wireless and mobile computing rapidly emerges from the integration among personal computing devices that compute and communicate in a distributed manner, cellular technology and the Web. This is possible due to the continuously increasing interaction between communication and computing, which is changing the information access from the current reactive “anytime anywhere” into the incoming proactive “all the time everywhere” approach. Such global computing scenaria are supported nowadays by a large variety of networks spanning from the well-known cellular networks to non-infrastructured wireless networks such as mobile ad hoc networks and sensor networks.

Such scenaria raise a number of interesting and complex algorithmic issues in diverse areas such as location management, resource allocation, ubiquitous information, network connectivity, reliability and security, and energy consumption. Easy, efficient, reliable communications on heterogeneous wireless networks are a key enabler for successful transparent extensions of overlay computing to mobile users. This need is combined with the increasing need for Quality-of-Service (QoS) guarantees. Indeed, mobile hosts and wireless networking hardware are becoming widely available, and extensive work has been done recently on integrating these elements into traditional wired networks, such as the Internet.

A main obstacle that has to be overcome for the full and profound accomplishment of the strategic goal of easy, efficient, reliable communications in heterogeneous wireless networks is the complexity of the challenges rising from the domain. This complexity stems from several factors. On one side, the target user base is very large and sophisticated applications are to be supported. On the other one, heterogeneous systems (with different computational resources) have to interoperate, while wireless stations have limited energy. Needless to say, the abstraction process over the common characteristics of the diverse component devices requires coping with new, unexpected computational problems.

User requirements and sophisticated applications that make innovative and extensive use of communication (such as video-on-demand) require more and more resources that are not always available to some of the wireless devices constituting an heterogeneous wireless network. Therefore, advanced and effective algorithmic tools and solutions for the careful and efficient management of the network resources are to be developed and used, otherwise the huge potential offered by new communication media will not be fully exploited. Indeed, we have reached a point in the technological evolution of the communication domain where any further progress on the transmission technology (e.g., new communication technology with a tenfold increase in bandwidth) does not have an equivalent impact on the capabilities and services offered to the information and knowledge-based society, unless it is coupled with adequate structural and algorithmic tools for the associated design and provisioning issues.

A special case of high importance arises whenever mobile users want to communicate in situations where no fixed infrastructure is available (either because it may not be economically practical or just because it is not physically deployable or was damaged), then the hosts with wireless network interfaces may set-up a temporary network without the aid of any centralized administration. This type of network is known as an ad-hoc wireless network [RT99]. Ad-hoc networks have been attracting a growing attention world-wide due to the remarkable potential applications in a variety of fields such as inter-vehicle communication, sensor networks, environmental control and defense. In the USA, DARPA heavily supports research in this area and the NSF panel on networking put a strong emphasis on all-wireless systems. In EU, telecommunication companies and defense services are also strongly interested. Ad-hoc networks are also expected to play and important role in environmental monitoring (radiation, toxic wastes, and other emergency situations), driving safety, cellular networks coverage extension as well as in defense. Technology makes indeed possible an emerging scenario in which different kinds of ad-hoc wireless networks are parts of a larger, heterogeneous wireless network, possibly comprehensive of even different kinds of wireless devices, e.g., mobile phones, sensors, PDAs, laptops, etc.

The main objective of the subproject is to provide practically efficient algorithmic solutions for high quality, reliable, stable end-users services to heterogeneous wireless mobile networks; this will be necessarily accomplished by appropriately abstracting over the common characteristics of the diverse component devices. In turn, the provision of such algorithmic solutions requires facing a number of strongly related issues that can be grouped in the following categories, directly corresponding to the workpackages of this subproject:

- WP 5.1: Resource Management and Quality-of-Service (QoS)
- WP 5.2: Dynamical Aspects of Network Design and Topology Control
- WP 5.3: Mobility and Fault Tolerance

In the context of overlay computing, the following key issues must be necessarily addressed by our solutions:

- Heterogeneity transparency: protocols have to correctly execute whatever devices are actually composing a network. Users are not required to know the actual network composition in order to use it.
- Reliability/stability: protocols operate in a mobile, dynamically changing environment, where mobile devices frequently join/leave the network. Such changes should be transparent to the users; hence, protocols' running times and quality of the solutions should be as independent as possible of the dynamic nature of the network.
- Scalability/adaptivity/fault tolerance: since mobile devices frequently connect/disconnect the network and the network size has to be transparent to the users, protocols must run efficiently independently of this size and be able to adapt appropriately to conditions by changing in an on-line fashion. Fault tolerance is strongly related with the concept of adaptability.
- Distributed issue: no centralized protocol can be realistically considered in an ad-hoc wireless network relying on no established infrastructure.
- Practical Efficiency. If short running time is an appreciated characteristic in every application, in the heterogeneous wireless networks setting it becomes crucial since protocols must potentially run

on several different devices having specific features and limitations to be regarded. Hence, we are not interested in, say, complex linear programming based algorithms; rather we shall often focus on simple and fast heuristics which work well on the average using local/limited knowledge. Finally, notice that the high quality services requirement is often in contrast with practical efficiency.

- Energy Efficiency. We are interested in solutions requiring low energy consumption. We want to remark how, similarly to the practical efficiency, energy efficiency is really crucial in the heterogeneous wireless networks setting. Transparency still requires that protocols run independently of the actual network composition.

WP5.1: Resource Management and Quality-of-Service

Leader: UDRTV; Participants: UOP, CTI, CUI, CAU, INRIA

We shall consider three critical resources: energy, spectrum (frequencies) and time. The motivation of our study stems from the following facts: on one hand, heterogeneous wireless networks consist of a large number of easy-to-use, portable devices (with some level of computational power), combined with globally accessible communication networks; it results into a large and increasing user community which demands for high-quality, efficient and reliable data communication. On the other hand, energy is a precious resource in wireless networks due to its cost from both the individual and the social point of view. Indeed, energy is spent by each node forwarding data according to a super-linear attenuation property of radio signals but the overall energy required to perform all communication requests represents a significant factor in the electromagnetic pollution and resource consumption.

Of course, efficient, broadband communication requires energy consumption from the network agents, thus yielding significant individual and social costs. This fact implies a crucial trade-off between Quality of Service and energy saving in the efficient wireless communication. We model this problem into three, strongly related, fundamental algorithmic issues:

A) Since one of the most severe limitations of wireless devices is their limited energy supply, one of the most crucial goals in designing efficient protocols for wireless networks is minimizing the energy consumption in the network. This goal has various aspects, including: (a) saving energy in network design results into the algorithmic problem of optimizing the devices (i.e. nodes) transmission ranges (i.e. energy power) in order to guarantee a certain connectivity and fault-tolerance property of the target region (Range Assignment problems); (b) minimizing the number (or the range) of data transmissions; (c) combining energy efficiency and fault-tolerance, by allowing redundant data transmissions which however should be optimized to not spend too much energy; (d) maximizing the number of “alive" particles over time, thus prolonging the system's lifetime and (e) balancing the energy dissipation among the devices (sensors) in the network, in order to avoid the early depletion of certain devices and thus the breakdown of the network. It is worth noting that it is usually difficult to achieve these goals simultaneously because of inherent trade-offs, and thus hybrid solutions may become useful (see e.g. [CHS05] for various energy efficiency aspects in wireless networks).

B) Saving energy during the network design phase also results into the partitioning of the limited radio spectrum into channels (frequencies) and efficient algorithms are needed in order to optimally assign these channels to the stations in order to avoid interferences that cause communication faults and thus useless energy consumption. We also observe that relevant technological platforms are CDMA based 3G networks and 4G successors. For these platforms a crucial aspect is the efficient use of the spectrum resources during communication in order to fully exploit the technological advances and to improve the QoS perceived by users (Spectrum Management).

C) Saving energy in wireless networks results into the issue of designing cooperation strategies among the nodes (i.e., protocols) that manage node transmissions in order to minimize message collisions and optimize the completion time of the basic communication operations such as point-to-point, multicast, broadcast, node accumulation and gossiping (Routing). Collision avoidance is especially crucial in sensor networks due to the physical proximity of the devices and simultaneous broadcast of messages of nearby sensors. We intend to focus on designing efficient collision avoidance protocols for multi-path data propagation. Data propagation in wireless sensor networks can be performed either by hop-by-hop single transmissions or by multi-path broadcast of data. We note that although several energy-aware MA layer protocols (like S-MAC) exist that operate well in the case of single point-to-point transmissions, none is especially designed and suitable for multiple broadcast transmissions.

Efficient solutions to these fundamental algorithmic questions constitute the main objectives of this workpackage.

Our approach to energy-efficient communication deviates significantly from the traditional layered network structure in that we jointly address issues related to network connectivity (i.e., a physical layer function), together with the task of subgraph formation (i.e., a routing function typically associated to network layer). A relevant example of this approach is that modelled by the well-known Minimum-Energy broadcast problem [NGE00]. It is common opinion [KKKP00, CPS02, NGE00] that such joint strategies and protocols will produce significant improvement in energy efficiency, as compared to a rigid layered structure that makes such decisions independently. The main contribution within this part will be efficient protocols for the above algorithmic questions. Furthermore, in order to evaluate the quality of the provided solutions we will carry out extensive simulations.

Another important class of problems that will be considered in this workpackage deals with the intrinsic distributed nature of the stations in an ad hoc network: it is common opinion that the present solutions do not pay enough attention to the selfish behavior of network stations (agents) [NGE00, CHPRV01, CPS02]. Unfortunately, the inherent self-organized nature of ad-hoc networks poses new problems since (i) stations do not belong to a central authority controlling them; (ii) a station transmitting with a certain range incurs in a cost proportional to the energy required; (iii) stations may not follow the “protocol” because of the limited battery capacity. Hence, due to the energy cost, it is not reasonable to assume an altruistic behavior of the nodes in forwarding somebody else's messages. Each node of the network aims in transmitting within a range as small as possible. The goal here is to provide rewarding mechanisms [NR99, NR00, R00] that return: i) an efficient algorithmic solution of small social cost and ii) induce node-agents to cooperate with the network manager by means of payments. Algorithmic rewarding mechanisms represent the new trend in Computer Science and, in particular, in Distributed Computing. Important results have been obtained for some computational problems [AT01, NR00], however, standard techniques, such as VCG mechanisms [NR99, R00] do not often work in the context of wireless networks.

WP 5.2: Dynamical Aspects of Network Design and Topology Control

Leader: CUI; Participants: UOP, CNRS, UPB, CTI, UDRTV, UPC, UDRLS, INRIA, DIM, UCY

Static aspects of wireless networks such as connectivity or topology are of prime importance to ensure the viability and efficiency of the networks. This is particularly clear when considering challenges raised by routing protocols. An appealing and fruitful way of analyzing such questions is to look at topological properties of the networks by means of the communication network graphs. Usually, considered transmissions are limited in their range, due to physical limitations of the stations involved in the networks, and this leads to communication graphs that look like a random geometric graph [P03]. The main property of interest is whether or not the graph percolates. A percolating graph is prone to allow long-distance communications through multi-hops. Models involved in the description of percolation in wireless networks are continuous, allowing the stations participating in the networks to be scattered at random on a given region [MR96].

A particularity of wireless networks is that they suffer from interferences, i.e. two stations cannot send data to a third agent at the same time, and this implies that static aspects of the communication graphs are no longer sufficient to describe the possibility of establishing long-distance communications, even through percolating communication graphs. To take into account interferences, a physical model is suggested in [GK00] where the connectivity of the graphs depends on the distance between two stations, similarly to the geometric random graphs model, and on the number of other stations scattered in the neighbor. Percolation results can be established in this setting [DBT03]. It is interesting to point out that percolation results depend crucially on the geometry of the region where the stations are scattered. Moreover, such results are based on point processes [DV88], which assume infinity of interacting stations scattered in a domain of infinite size.

From a realistic point of view, stations are prone to suffer power limitations and then try to limit their power consumption by limiting the time span of sending data. Taking this into account leads to the need for considering the particular time an agent sends a piece of data. From a modelling point of view, at a given time, only a subset of the total number of stations are prone to establish a communication. This effectively reduces the number of possible interferences. To make explicit that the interferences between stations depend on the time, we call as a collision the event that two or more stations send a piece of data at the same time to the same agent. Of course, in this case the reception is impossible due to the multiplicity of the data sent. To make such an analysis possible, we need to make some assumptions on the statistical behavior of the stations. This can be done, for example, by considering a particular process on the network such that broadcast problem [LR04] or a localization process in the network [BLR05].

Dynamical models of wireless networks intend to take into account the succession of the events in the networks in order to ponder the effects of interferences, for example during a broadcast process or a localization process. The motivation for developing such models follows from the consideration above. Moreover, such models have potentially a lot of applications such as: i) Evaluating the efficiency of distributed algorithms in such networks. ii) Designing protocols, which limit the number of collisions, optimizing the power consumption of the stations. iii) Providing stopping criteria, for example when broadcast is considered. iv) Providing analytic support to optimize station transmission capability (range, angle or power of transmission). v) Providing analytic support to design networks, optimal number of stations, technical characteristics of the deployed stations.

Related to power limitations, we shall also consider topology management using Power-Saving Modes. Implicit network topology management mechanisms include sleep-awake schemes in wireless sensor networks (see e.g. [CN03]). Such mechanisms force the sensors to alternate between sleeping and awake modes, in order to save energy. During sleeping periods sensors cease any communication with the environment, thus are unable to listen, receive and propagate data transmitted by other sensors. Thus, energy consumption is reduced, however the network topology is affected and the efficiency (or even the correctness) of protocols should be properly investigated and established.

Propagation protocols for such energy-restricted systems should at least guarantee that the control center (i.e. the destination of data) eventually receives the information propagated. The success of such protocols depends on various parameters, such as the density of sensors in the network area, their distribution, the distribution of sleeping and awake time periods. The interplay between these parameters may become pretty complex. In particular, the relation between the maximum sleeping time period and the other parameters, allows to program the sensor-network energy saving specifications accordingly. The proper design of these schemes may achieve (besides energy savings) fault-tolerance and efficiency as well. Towards this goal, and because of the dynamic nature of such networking environments, it is important to also investigate adaptive solutions, i.e. when the protocol parameters (such as the sleep-awake periods) dynamically adapt to implicitly sensed dynamic changes (such as changes in the network topology/density etc).

An important goal concerning the dynamical modelling of wireless networks is to take into account many parameters of the networks. We have mainly in mind geometrical parameters related to the geometry of the region where the stations are deployed. From a percolation point of view we have already mentioned that this is a key point [DTH02]. Moreover, parameters related to the characteristics of the transmissions of the deployed stations are also of importance.

The range of transmission should be large enough to ensure connectivity with some neighbors, but if the range of transmission is too large, with respect to the density of emitting stations, collisions occur reducing the efficiency of the networks and increasing the power consumption.

Usually, transmissions can be established in a region close to a disk, which depends only on the range of transmission; this leads to geometrical random communication graphs [P03]. However, it is of interest to ponder whether we improve the quality of the communications, mainly reducing the number of collisions, by introducing directional communications, where communications are established in a sector of disk described by the range of communication and the angle of communication. In this situation, we also have to consider a new parameter, the orientation, which is usually taken as a random variable. Interestingly, from a percolation point of view, there are results, which partially show that we improve the number of connections when considering sector disks for communications instead of disks

[BBCFM]. A similar conclusion was found in [LR04], [BLR05]. However, from a dynamical point of view, the increase of new connections can lead to an increase of the number of collisions. This implies that to take advantage of the introduction of directional communications, we need to introduce and analyze efficient protocols.

To summarize our discussion, efforts are to be made concerning: i) The impact of the geometry of the region where the stations are deployed. ii) The impact of the shape of the transmissions. These points can be complementarily addressed considering static aspects of the communication graph or dynamical aspects as we suggest.

The methodology we develop based on previous works [LR04], [BLR05] is based on size dependant branching processes [J97], [H63] which are extensions of the now classical theory of branching processes [H63]. Classical branching processes were already considered in the pioneering paper [G6], but apparently this aspect of the paper has received little attention in the community of people working on wireless networks.

The introduction of dependencies in such branching processes, leading to size or population dependent branching processes, is necessary for many applications and there are still no unified ways of working in this framework [J97], [O96], making this research topic an active area of research. In the framework of wireless networks, dependencies appear naturally when considering collisions and hence, the introduction of population dependent branching processes is natural. Results are expected in this direction concerning the existence of quasi-stationary distributions as numerically observed in [LR04], extending known results for the classical theory of branching processes [H63], [Y47].

WP 5.3: Mobility and Fault Tolerance

Leader: CTI; Participants: CNRS, UPB, UOI, CUI, UDRTV, UPC, UDRLS, KUL, INRIA

Wireless networks are mainly constituted of mobile devices and are typically adopted in scenarios where unpredictable node and link faults happen very frequently. We emphasize that “faults" also model mobility: a device moving from one location to another can be regarded as a set of links going down and a second set of links restarting. Generally speaking, and as already described, all the problems dealt within workpackage WP1 can be defined in the mobility and fault tolerance setting as well. However, a number of new algorithmic issues naturally arise in this setting. Notice that the transparency requirement gives rise to an interesting topic of investigation which is strictly related to mobility. If some user A needs to communicate with some other user B, any routing algorithm needs to know where B is located. Any time B moves, its new location has to be communicated to the routing algorithm, but B’s movements need not to be known to A. The topic of users tracking (or location transparency) consists of the problems of retrieving information about devices locations so that their mobility is transparent to users.

The presence of mobile agents and the request for high fault tolerance lead to scenarios where algorithmic solutions to the problems described in workpackage WP1 must work on dynamical input configurations. While in traditional wired networks it is reasonable to assume that failures will be managed (either in a centralized or a distributed basis) by only taking into account the underlying physical structure, in wireless networks energy issues cannot be disregarded. According to that, the aim of this workpackage is to cope with failures at different levels: (1) on one hand, one must design and analyse reliable and efficient (both from an energy saving and a completion time point of view) communication protocols in the presence of faults; (2) on the other hand, it is crucial to define decentralized recovery procedures that need to be activated to maintain the original connectivity predicate, by constantly paying attention to the limited energy capability of hosts. As to the former topic, the aim of this project is to design and analyse reliable communication protocols in the presence of faults. Previous efficient solutions for some special cases [BGI87, PR97, KKP98, KM98, CMS01a, CMS01b] of the broadcast problem have been successfully addressed by combining combinatorial structures (such as Selective Families, D-sequences, and Superimposed Codes [I97, CHI99, CMS01a]) and standard wave protocols. Such combinations do not work for achieving completion time as function (also) of the number of faults suffered by the network. Another important aspect is the power of the Fault-Adversary. Previous theoretical analysis considered only the worst-case, while a significant performance evaluation of the protocols would be that in which the adversary is random and thus analyse the expected completion time. So, new combinatorial structures and non-wave protocols must be considered for the above goal. In what follows, we list the main problems and the relative possible methodologies we intend to consider with respect to this topic: i) Fault-tolerant (multi)-broadcasting as a function of the number of faults for general wireless networks (F-Selectivity and non-wave protocols). ii) Computing aggregate information in faulty wireless networks (F-Selectivity and non-wave protocols). iii) Fault tolerance against random adversaries (average-case and smooth analysis).

Concerning the relationship between fault-tolerance and range assignment, we intend to pursue the objective of maintaining a certain connectivity predicate in a wireless network affected by a transient component failure, aiming both to save the energy consumption, and to preserve as much as possible the pre-existing network. More precisely, the network functionality will be re-established by satisfying the following constraints, in this order: (a) minimize the number of changes, where a change is any network operation defined a priori (e.g., a range increment of a radio station, etc.); (b) minimize the objective energy function addressed by the original network over the restricted set of feasible solutions defined by (a). We want to emphasize here that most of the conventional optimization problems in wireless networks are NP-hard, and therefore we expect the use of approximation techniques. Special attention will be devoted to the so-called "local" maintenance of the network functionality, which is restricted to adjust only the ranges of the hosts that were actually using the failed component for their communication purposes. Finally, we plan to compare the quality of the constrained solutions with respect to the optimal ones, as computed from scratch in the residual network (i.e., the network deprived of the failed component).

- The case of ad-hoc mobile networks

The impact of the mobility rate and the user density on the performance of routing protocols in ad-hoc mobile networks may be significant. In particular, the effect of these parameters on routing protocols that try to maintain and dynamically update connectivity related data structures (like the well-known AODV protocol and its variations/extensions) may become dramatic in the case of very high mobility rates and/or sparse networks, in the sense that the dynamic changes may be faster than the time available for dynamically updating the data structures and also in the sense that the connectivity of the network may be significantly affected and the network diameter may be significantly increased. In such cases, such routing protocols may become inefficient and/or erroneous, leading to low success rates and increased message delivery times. In the light of the above, we plan to continue our research proposing routing protocols using and taking advantage of the motion in the network by e.g. forcing few hosts to move acting as “helpers" for message delivery. As shown in [CNS03, CKN04], such protocols seem to tolerate well (and in fact benefit from) high mobility rates and low densities.

- Combinatorial models for interaction and mobility

Mobility and interactions of particles in wireless networks can be modelled by suitable random graphs. In particular, we plan to adopt the G(n,m,p) model of random intersection graphs [NRS04]. Random intersection graphs capture interaction between computing entities due to physical proximity or resource sharing. Thus, such graphs may model real-life applications more accurately (compared to the classical random graphs). This is because in many cases the independence of edges is not well justified. In fact, objects that are closer (like moving hosts in mobile networks or sensors in smart dust networks) are more probable to interact with each other. Other applications may include oblivious resource sharing in a distributed setting, interactions of mobile agents traversing the Web etc.

We also plan to use stochastic processes, and, in particular, stochastic interactions of particles and their applications in network problems (see e.g. our research in [DNS04] where we investigate how long it takes for a red particle moving randomly on the vertices of a graph to "infect" with its color white ones when meeting them). It turns out that stochastic interactions of particles model well information propagation in mobile environments. Further applications include the spread of a virus in computer networks, information spreading (such as rumor spreading, gossiping).

Sub-Project 6: Design and Implementation of Components and Applications for Programmable Overlay Computer

The work in this subproject will serve as a "proof-of-concept" for the whole project. Our main goal is to develop a Programmable Overlay Computing Platform based on the scientific advances of the project; this will be an overlay computer built over a global computer consisted of Internet-connected nodes. Among the functionalities that will be separately investigated in the other subprojects, the platform will include those which are important for the execution of distributed applications with intensive requirements for efficient use of computing, communication and information resources and security. To demonstrate the programmability of the platform, we also plan to implement an application on top of it. In parallel, for the validation of the scalability of both the platform and the application, we will also develop and setup an appropriate testbed environment.

The work within subproject SP6 is divided into the following workpackages which are described below:

WP6.0 Project management and dissemination activities
WP6.1 Specification and design of the platform
WP6.2 Implementation of platform components
WP6.3 Integration and testing of the platform
WP6.4 Design and implementation of a demo application

By its nature, this subproject will be in continuous cooperation with all other subprojects. We will adopt a project planning/engineering process approach that can adapt to the rapid changes to scientific and market evolution, and minimize the risks associated to the project. For the implementation of our work we use an approach that is iterative and driven both by usage scenarios and the scientific advances in the project.

The plan foresees an initial analysis phase for applications and components where, in cooperation with other SPs, we define the high level architecture, the main use cases and technical requirements, and perform selection of a base implementation platform. This initial phase will be followed by a sequence of iterations. In each iteration we will:

- identify algorithms and models at a sufficient stage of maturity in the project,
- identify a package of use cases, describing a number of functions relevant to the end-user, which can benefit from such algorithms and models,
- derive more specific requirements for each research component (models, algorithms), which are necessary for the realization of the use cases above,
- implement technological components,
- develop a test bed application for the selected use cases, which use the technological components, and a related set of test cases,
- integrate all components to deliver the Programmable Overlay Computing Platform,
- test and assess results.

Iterations should be as short as reasonable. While agile software development methods advocate very short iterations, for a project like ours (which includes research, development of infrastructural software, application software, and distributed development among different parties), the duration of each iteration can be no less than one year (actually, the first iteration will take 18 months). We foresee three iterations in the duration of the project.

WP 6.1 Specification and design of the platform

There are a variety of platforms already available from the research community and the industry, which deliver a number of basic services, which can be leveraged to build a programmable overlay computing software platform. Among them, noteworthy examples from the P2P community include the JXTA platform by Sun, GnuNet, a number of platforms based on the DHT concept (e.g. Chord, Pastry, Tapestry), and FIPA-based platforms such as JADE from the agent research community. JXTA technology (http://www.jxta.org/) is a set of open peer-to-peer protocols that enable any device on the network to communicate, collaborate, and share resources. GNUnet (http://www.ovmj.org/GNUnet) is a framework for secure peer-to-peer networking that does not use any centralized or otherwise trusted services. The DHT-based platform is built on routing and object location primitives that can be used to support a variety of peer-to-peer applications. JADE (http://sharon.cselt.it/projects/jade/) is an agent-based platform that permits the development of P2P distributed applications and is already enabled for running peer components on mobile devices.

The objective of WP6.1 is to select the platform to be used as the heart of the overlay computer, define the functionalities that will enhance it and select the corresponding algorithms for the implementation of these functionalities. The platform should provide basic primitives for distributed computing and communication. Functionalities for resource discovery, bandwidth sharing, data management, load management, scheduling, trust management, secure computation, etc. will enhance and extend the basic primitives of the platform. The algorithms that will be used for these implementations will be selected among the ones developed and validated within subprojects SP1, SP2, SP3, and SP4 while some of them will also take into account the specific characteristics of wireless devices (SP5). It is expected that, for some functionalities, more than one algorithms will be selected for implementation.

WP 6.2 Implementation of platform components

Task Leader: TBD Further Participants: TBD

The workpackage will be devoted to the implementation of functionalities using the corresponding algorithms selected in the context of WP6.1. The implementation will be done on top of the platform selected in WP6.1. At this phase, functionalities will run on the platform as stand-alone software. Some preliminary testing (on this stand-alone basis) will be performed here. One of the outcomes of WP6.2 will be a set of implementations on top of the platform. In addition to the implementation of functionalities, we also plan to provide prototypes/demostrators based on the developments within workpackages WP3.2, WP3.4, WP4.2, and WP4.3.

WP 6.3 Integration and testing of the platform

A grand challenge for the kind of highly distributed infrastructures and applications that we target, is the methodology itself as well as the tools which are needed to perform test and validation. In most cases, validation of a newly proposed design or algorithm is normally performed using a combination of simulation and testbed experimentation. However, none of these techniques can be easily adopted or provide accurate results in scenarios involving infrastructures and applications of large scale. For example, many new approaches in the area of peer-to-peer networks have been developed in parallel with a corresponding simulator. The intention of many peer-to-peer systems is to scale to large numbers of peers; and, testing the scalability of such an actual system may be impractical. Such simulators include SimP2, Bison, PLP2P, 3LS, p2psim and others but the relevance of these simulation studies to the deployment of real P2P systems remains an open question. Infrastructures and applications like the proposed programmable overlay computing platform actually need testbeds supporting large-scale, real-world experiments.

One of the objectives of this workpackage is to research, design, and setup testbed architectures fit for experimenting with overlay computers. In early phases of the project, we will perform an assessment of fitness of such experimental testbeds (e.g., PlanetLab) to our project needs. As the project progresses, this workpackage will perform original research to design a flexible testbed architecture, where tests can be independently and without interference performed by different users in a highly distributed setting. Some of the research lines developed in the project (e.g. resource discovery, data management, scheduling, load management etc.) could provide useful components to this distributed testbed environment. In addition, we will design a testbed monitoring service using distributed data dissemination technology developed by TILS (building on XDM – DHT-based software with extensions for indexing and querying of data).

The main objective of the workpackage is to integrate the functionalities validated in WP6.2 into a common programmable overlay computing platform, i.e., to determine which implementations will be part of the platform and fine-tune the interface of the functionalities (i.e., fine-tune their parameters, options, etc.) from the programmer's point of view. The outcome of this workpackage will be an enhancement of the basic computing/communication primitives provided by the base platform with resource management and security functionalities. The enhanced platform (Programmable Overlay Computing Platform) will run on the distributed testbed described above running on computing facilities provided by the partners.

WP 6.4 Design & Implementation of a demo application

In order to validate the quality of the proposed platform, it is important to develop, install, and run applications. The programmable overlay computing platform is oriented to large scale, distributed applications. The platform can be established in different business scenarios, i.e. community-based computing over the Internet, enterprise computing, 'virtual-enterprise' scenarios such as those occurring in B2B, interconnected public institutions and governmental agencies, research networks, etc.

While most applications should in principle benefit from the distributed, programmable overlay computer, specific classes of applications will be most suited to run on top of such a global computer (e.g. applications which are not mission-critical, those that can accept statistical service level guarantees; that can be subject to heavy and rather unpredictable loads; applications which should account for a large population of users and need flexible authorization models and approaches to trust etc). Examples of such applications include data sharing and dissemination (e.g. Usenet type of applications), content distribution (e.g. dissemination of program updates, streaming media and other material), auditing and digital preservation of documents and papers through replication (e.g. LOCKSS),etc.

The workpackage will select one application scenario and develop an application to test and showcase the capabilities of the platform using the functionalities provided by the partners. Running the application over the platform will showcase several of the scientific advances of the project in such areas as resource management, resource discovery, load management, distributed data management, scheduling mechanisms, bandwidth optimization, security and trust management.

The application will run in the testbed delivered in WP6.3. Together with the programmable overlay computing platform, it will serve as "proof-of-concept".

B.4.2 Demonstration activities

-- Not applicable --

B.4.3 Training activities

We consider training activities to be an important part of the project. In particular our training activities will include:

- **Schools/Workshops:** We intend to organize at least one school per year and in particular we plan to organize two schools in the first 18 months. The purpose of these schools will be the training of PhD students and young researchers so that they will be introduced to areas related to algorithmic aspects of global computing. Schools are expected to draw participants from almost all participating sites. Topics will include both fundamental areas that are necessary as a background for the required research as well as state-of-the-art hot topics related to particular subprojects. Schools will normally be complemented with workshops where researchers and students will present their current work.

- **Seminars/Mini-Courses:** Participating sites will also organize seminars/mini-courses on specific topics related to their subprojects. Such seminars/mini-courses would normally be addressed to smaller audiences than the schools.

The lecturers for these activities may either be expert members of the consortium or invited external experts.

B.4.4 Consortium management activities

In this Section, we briefly describe the project management structure and activities of the project. A more detailed description will follow in Section B6.

The complexity of the project requires a very intensive coordination. The Coordinator is formally responsible for managing and controlling the activities of the consortium. The management activities include:

Management of the project, i.e.:

- Setting up the necessary consortium contracts and monitoring their agreements
- Organizing regular meetings of the Coordinating Committee (CC) and the Consortium Board (CB)
- Reporting to the European Commission
- Making all other necessary decisions/actions required to proceed with the project as agreed on with the European Commission.

Response management, i.e.:

- Assuring that results are achieved as agreed on with the European Commission.
- Reacting on problems resulting from unforeseen RTD difficulties
- Planning activities involving the interaction of the project with other projects and research activities within the Global Computing Proactive Initiative.

The Coordinating Committee (CC) will support the Coordinator in the management activities. It communicates frequently about all ongoing matters to quickly react to changes, and it meets regularly at least once a year and otherwise whenever deemed necessary.

Policy issues and consensus within the project will be established through meetings of the Consortium Board. The Consortium Board consists of one representative for each partner, and is chaired by the Coordinator. The Consortium Board decides on all matters required for the smooth operation of the project. It communicates frequently about all important issues and it meets at least once a year.

For every subproject, there will be a Subproject Leader. The main responsibility of each Subproject Leader is to coordinate the activities of his/her subproject and report about the state of the subproject to the Coordinator and to the CC. In particular, within each sub-project there is a management workpackage whose objective is to guarantee the successful completion of the subproject within the agreed time and quality requirements, as well as ensure compliance with EC standards and procedures for project management and tracking.

Finally we note that several of the partners have already operated successfully in the context of other projects funded by the EU including three projects under the first Global Computing initiative (CRESCCO, FLAGS, DBGLOBE); we believe that the experience gained through these projects will facilitate interactions and communications within this project and will ease its management.

B.5 Description of the consortium

The Consortium is composed by twenty-one sites from ten different countries, including three new member states, Cyprus, Czech Republic and Estonia, one of which is an SME. All the participants to the project have been already identified and a very brief description for each site is given below. In general the Consortium comprises participants whose expertise spans from fundamental theoretical topics, like, game theory (e.g., UPB, CTI, UoA, UNISA), scheduling (e.g., CAU, UDRLS, UoA), parallel algorithms (e.g., UPC, UPB, UNIPD, UoP), to advanced system-oriented ones, like distributed data management, (UoI, MPII), system security (UNISA, KUL, CYB, UoI) and workflow management (ETHZ, MPII).

The Laboratory of Distributed Systems and Telematics at the **University of Patras (UoP)** is one of the ten laboratories of the Department of Computer Engineering and Informatics, University of Patras; the leading academic and research department in Greece in Computer Science. The Laboratory occupies 4 faculty members, 3 PhD researchers and about 25 PhD or MSc students. The research group that will be involved with the project has strong experience in distributed computing and fundamental aspects of communication in networks and significant research record in topics like network routing, bandwidth allocation in optical and wireless networks. Its members have participated in many Research & Development projects either funded by the Greek state or by the European Union (ESPRIT, BRITE-EURAM, etc). UoP has strong collaboration with North America Institutions (MIT, CMU, Harvard, Brown, Columbia, Carleton), and research groups in Europe and Israel (Technion, Weizmann).
UoP will be the coordinating partner of the project, and it will also be responsible for the dissemination activities within the project.

Christos Kaklamanis received his S.B. in Computer Science and Engineering from Massachusetts Institute of Technology, Cambridge, MA, USA (1986). Then he obtained his S.M. (1989) and his Ph.D. (1992) in Computer Science from Harvard University, Cambridge, MA, USA. Afterwards he worked as Postdoctoral Fellow in Computer Science under the program for Massively Parallel Computation at the Center for Discrete Mathematics and Theoretical Computer Science (DIMACS, Piscataway, NJ), an NSF funded joint project of Rutgers University, Princeton University, AT&T Bell Laboratories, and BellCoRe. He also worked as research consultant for NEC Research Institute, Princeton, NJ. Currently he is Professor in the Department of Computer Engineering and Informatics at University of Patras, where, from 1997 to 2003, he served as the director of the ``Division of Applications and Foundations of Computer Science". During the last years, he participated as a key researcher in the EU funded basic research projects ESPRIT/ALCOM-IT and ESPRIT/GEPPCOM, EU RTN ARACNE and, as the project leader, in EU-FET project CRESCCO and R&D projects funded by the Greek State. His research interests include Parallel Computation and Communication, Networks, Probabilistic and approximation algorithms, Fault Tolerance and Theory of Computation.

Emmanouel (Manos) Varvarigos received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 1988, and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions at the University of California, Santa Barbara (1992-1998, as an Assistant and later an Associate Professor) and Delft University of Technology, the Netherlands (1998-2000, as an Associate Professor). In 2000 he became a Professor at the department of Computer Engineering and Informatics at the University of Patras, Greece, where he heads the Communication Networks Lab. He is also the Director of the Network Technologies Sector (NTS) at the Research Academic Computer Technology Institute (RA-CTI), which through its involvement in pioneering research and development projects, has a major role in the development of network technologies and telematic services in Greece. Professor Varvarigos has served in the organizing and program committees of several international conferences, primarily in the networking area, and in national committees. He has also worked as a researcher at Bell Communications Research, and has consulted with several companies in the US and in Europe. His research activities are in the areas of protocols for high-speed networks, parallel and distributed computation and grid computing.

Telecom Italia Learning Services (TILS) is a business unit of Telecom Italia. It is a key player in national and international markets for Learning and Knowledge Management. TILS' mission is to provide its customers with integrated solutions for the organization, management, creation and sharing of knowledge. Telecom Italia Learning Services brings together Telecom Italia companies and divisions with years of experience in training and ICT. TILS has a staff of more 300 ICT experts as well as an extensive network of alliances and partnerships with high-tech companies, state organizations, universities and research centres in Italy and abroad. Major customers include ICT companies and the government sector. A key component in Telecom Italia's strategy is research and development, which is rapidly increasing its share of company turnover. TILS' dedicated R & D group brings all knowledge-based activities (research, training and innovation) under one roof. The group's activities cover a broad range of areas relevant to the medium and long term future of IT.

Dr. Fabrizio A. M. Davide, is a Telecom Italia manager, and is head of unit of the Engineering division in Telecom Italia Learning Services. Since 2000 Davide has been Adjoint Professor at Linköping University (Sweden). He has coordinated numerous R&D projects co-funded by the European Commission, the most recent being "I-Learning" and "Angelo", in Framework V's IST Programme. He was technical manager in the Briteuram CIA project and has worked as a researcher in many applied-research projects with universities and public research centres (e.g. the Italian National Council of Research, the Washington Technology Centre and the Massachusetts Institute of Technology, USA, the University of Warwick, UK, the University of Tübingen, DE, and the Swedish Research Centre S-Sence).

Giovanni Cortese is scientific coordinator of the "Platform and networks technologies" group within the R&D group of TILS. He is a senior technology consultant with a solid background on software engineering methods and software architectures. He worked at the executive level as methodologies manager in Sodalia, a Telecom Italia business unit, as project manager for large software development projects, and as business innovation manager. He is chief architect for TILS in the VICom project, a large research project funded by Italian government on smart environments and immersive communication. He teaches a course on Middleware at the University of Trento, Italy.

Mascotte is a joint project **Centre National de la Recherche Scientifique** (CNRS)/INRIA/UNSA. It involves 25 members among them 14 permanent researchers (from CNRS, INRIA and University) and is strongly associated to France Telecom R&D. Mascotte's main objective is to develop methods and tools for efficient use of telecommunication networks. This involves high-level research in the fields of simulation, algorithms, and discrete mathematics. It has developed industrial collaborations with various partners such as France Telecom, Alcatel and CNES for the design and optimization of telecommunication networks. Mascotte was involved in various projects funded by the EC, in particular recently in the FET CRESCCO project and RTN ARACNE project. It has also many bilateral cooperations with European countries and also with Canada, Brazil, Israel. It is a part of the laboratory I3S which include 130 members working in Science and Information Technology.

Jean-Claude Bermond is research director (exceptional class) at CNRS in the laboratory I3S join lab between CNRS and UNSA (University of Nice Sophia-Antipolis). He has been actively involved in research in the areas of Discrete Mathematics, Interconnection Networks and Parallel and Distributed Computing with applications in Telecommunications.
In these fields he has published more than 130 papers and supervised 50 PhDs. He is member of many editorial boards of Journals. He has also co-organized many conferences and workshops on these subjects. He is now responsible for the joint project CNRS/INRIA/UNSA MASCOTTE He is also responsible of various contracts with industrial partners and he is involved in many international cooperations inside Europe (in European projects) or with Canada, Brazil. He has been chair of different laboratories and president of many committees in France.

Stephane Perennes obtained his Phd in 1996 from the University of Nice, Sophia Antipolis, and after a Post Doc in Delf Univ., he joined the Mascotte team as a researcher of CNRS in 1997. His research interests are communications in parallel computers, algorithm for optimization of telecommunications networks, distributed algorithms, structural properties of graphs, wireless networks and network design.

The Institute of Computer Science of the **University of Paderborn** is one of the leading research centers in the area of parallel and distributed computing and algorithmic game theory in Europe. The working groups of Prof. Dr. Burkhard Monien and Prof. Dr. Friedhelm Meyer auf der Heide are the main contributors of the research on these topics in Paderborn. Their research groups consist of roughly 30 full-time researchers. The groups are/were partner in several European research projects (ALCOM I&II&IT&FT, FLAGS, DELIS, HCM MAP & SCOOP, HPCNet, GP MIMD, ZEUS, PCA, EUROPORT, SICMA, UP-TV, HiPEC, EPRI COM and PARROT) as well as in numerous national projects funded by the German research foundation (DFG) and the research ministers of North-Rhine-Westfalia (MWF) and the Federal Republic of Germany (BMBF). The site will contribute with its expertise on parallel and distributed computing and algorithmic game theory to SP 1, SP 2, SP 3, and SP 5.

Prof. Dr. Burkhard Monien, born in 1943, PhD in 1969, was appointed in 1975 as an Associate Professor at the CS-Department of the University of Dortmund. Since 1977 he is full professor at the University of Paderborn. His research interests include theoretical questions as well as practical problems concerning the efficient use of parallel and distributed systems. He has published more than 150 papers in most of the well-known computer science conferences and journals. Prof. Monien has been a chairman on several program committees and has worked in numerous program committees of leading conferences. In 1992 Prof. Dr. Burkhard Monien received the Leibniz research award of the DFG together with Prof. Dr. Friedhelm Meyer auf der Heide. Until the end of 1995 he was the secretary general of the European Association for Theoretical Computer Science (EATCS). Since 1996 he is a member of the Akademie der Wissenschaften in NRW.

Prof. Dr. Friedhelm Meyer auf der Heide, born in 1954, PhD in 1981, was appointed as assistant at the CS-Department of the University of Frankfurt, with one year (4/84 - 3/85) on leave as postdoc at the IBM Research Laboratory in San Jose, CA. From 1986 to 1989 he was assistant professor at the University of Dortmund, since then he is full professor in the CS-Department and the Heinz Nixdorf Institute of the University of Paderborn. His research interests include complexity theory, efficient algorithms, and parallel computing. He has published more than 100 articles in the leading journals and conferences on theoretical computer science. Friedhelm Meyer auf der Heide has served on numerous program committees of leading conferences and has organized several workshops. In 1992 Friedhelm Meyer auf der Heide received the Leibniz award of the DFG together with Burkhard Monien for their work on parallel and distributed computing. Since 1995 he is the director of the DFG-Sonderforschungsbereich "Massively Parallel Computing".

The **Computer Technology Institute** (CTI) is a non-profit Research Organization, closely affiliated to academia. CTI constitutes an integrated environment of scientific research and design and development of products and solutions. Most of the funded research comes from the competitive programmes of the E.U., and the Hellenic private sector. Also, CTI has been serving as the official technical consultant of the Greek state, undertaken major, large-scale projects. More than 300 scientists engineers and other staff support its activities and daily operation. CTI's Research Unit 1 (RU 1) conducts research in the Foundations of Computer Science and their Relevant Technologies and Applications. Basic thematic areas of interest include important aspects of algorithms and complexity such as parallel and distributed computing, algorithmic aspects of networks, probabilistic techniques, approximation algorithms and computational complexity. Also, Research Unit 1 is interested in technical consultancy and software development in the following applied research areas: security, simulation and optimization in computer integrated manufacturing, internet technologies, information systems, and culture and art. Research Unit 1 consists of 7 Faculty Members, 6 PhD Researchers and 15 Engineers-PhD Students. The site will contribute with its expertise on wireless computing, algorithmic game theory, distributed computing and algorithms and complexity to SP1, SP2, SP4 and SP5.

Paul Spirakis born in 1995, obtained his PhD from Harvard University, USA, in 1982. He is currently a Full Professor in the Department of Computer Science and Engineering of Patras University (Greece) in 1990. Paul Spirakis is Director of the Research Academic Computer Technology Institute (RA.CTI). His research interests include Algorithms and Complexity and interaction of Complexity and Game Theory. Paul Spirakis had published two books through Cambridge University Press, and eight books in Greek. Paul Spirakis was the Greek National Representative in the Information Society Research Programme (IST) from January 1999 till June 2002. He was elected unanimously as one of the two vice-Presidents of the Council of the European Association for Theoretical Computer Science (EATCS). He is member of ISTAG (Information Society Technologies Advisory Group). He consults for the Greek State, the European Union and several major Greek Computing Industries.

Peter Triantafillou is currently a Full Professor with the Department of Computer Engineering and Informatics, at the University of Patras, Greece, Director of the Information Systems laboratory, and a senior researcher in the Research Academic Computer Technology Institute, Greece, where he is heading the research efforts on Network-Centric Information Systems. Peter received the Ph.D. degree from the Department of Computer Science at the University of Waterloo, Canada, in 1991. Prof.

Triantafillou's research efforts currently focus on large-scale systems for Internet Content Delivery, Sharing, and Integration, with particular emphasis on Peer-to-Peer systems and Distributed Event-Based, Publish/Subscribe Systems. He has participated as key researcher in several European research and development Projects, most notably in the DELIS IST/FET Integrated Project, the DIET IST/FET project, the DBGlobe IST/FET project, and in the HERMES IST/FET project.

Sotiris Nikoletseas is currently a Lecturer Professor at the Computer Engineering and Informatics Dept of Patras Univ., Greece and a Senior Researcher at CTI. His research interests include: Algorithmic Techniques in Distributed Computing (focus on wireless sensor networks and ad-hoc mobile networks), Probabilistic Techniques and Random Graphs, Algorithmic Engineering and Large Scale Simulation. He has served as a PC Chair of many Conferences (including ALGOSENSORS 04, MOBIWAC04, WMAN 05, WEA 05). He is co-author of a book on Probabilistic Techniques in Computer Science and he co-authored more than 70 publications and 6 Chapters in Books by major publishers (Kluwer, Springer, Wiley). He is involved as director/senior researcher in more than 20 externally funded R&D Projects, mostly related to IST/FET (including ALCOM-FT, FLAGS, CRESCCO and DELIS).

The Dipartimento di Informatica ed Applicazioni "Renato M. Capocelli" of the **Universita' di Salerno** has about 25 faculty members with interest in the fields of algorithms, formal methods, security, cryptography, distributed applications. Members of the UNISA site have been recently involved in Global Computing FET project CRESCCO (Critical resource sharing for cooperation in complex systems) and in Network of Excellence ECRYPT in Cryptology. The site will contribute with its expertise on Algorithms for Selfish Agents, Approximation Algorithms and Security to SP1, SP2, SP3 and SP4.

Giuseppe Persiano got his Laurea cum laude in Computer Science from the Universita degli Studi di Salerno in 1986. In 1988 he got a Master of Science and in 1992 he got a Ph.D. Both degrees are in Computer Science and from Harvard University. Since November 1994 he has been Professor of Computer Science at the Dipartimento di Informatica ed Applicazioni of the Universita degli Studi di Salerno. His research focuses on cryptography and the design and analysis of computer algorithms. In particular, he has been working on secure computation, distributed computation, and algorithm for selfish agents.

The **University of Ioannina** was founded in 1964 in Ioannina, Epirus, in the north-western part of Greece. Today the University has nineteen departments with over 400 faculty members, 12,000 undergraduate students and more than 800 graduate students. The research described in this proposal will be carried out by the distributed data management group of the Computer Science Department. The group performs research on various topics of data management with emphasis on mobile, ubiquitous and peer-to-peer computing. Currently, the group consists of 6 faculty member, two postdoc researchers and 11 graduate students. The group coordinates the DBGlobe project (IST-2001-32645) funded by a FET on global computing and several other national and international research projects. The site will contribute with its expertise on data management for large distributed systems and will participate in subprojects 2, 3, 4, 5 and 6.

Evaggelia Pitoura received her B.Sc. from the University of Patras, Greece in 1990 and her M.Sc. and Ph.D. in computer science from Purdue University in 1993 and 1995, respectively. Since September 1995, she is on the faculty of the Department of Computer Science of the University of Ioannina, Greece where she leads the distributed data management group. Her publications include more than 70 articles in international journals and conferences and a book on mobile computing. She has also co-authored two tutorials on mobile computing for IEEE ICDE 2000 and 2003. She is recipient of the best paper award of IEEE ICDE 1999 and two "Recognition of Service Awards" from ACM. She was/is on more than 60 program committees of international conferences, program vice-chair for "Distributed, Parallel, Deep Web, and P2P Database" of ICDE 2005, program chair of the 3rd Hellenic Data Management Symposium 2004 and program co-chair of the MDDS 2002, MobiDE01 and MobiDE99 workshops.

The **Centre Universitaire d'Informatique (CUI) of the University of Geneva** is an interdisciplinary research center which gathers together people coming from different faculties of the University of Geneva with interests in the general area of Information and Communication Systems. The participants of the CUI in the project are researchers from the Theoretical Computer Science and Sensors

Lab(TCSensor), headed by Professor Jose Rolim. The concerned research involved in the TCSensor Lab is related to the modelisation of wireless networks with focus on dynamical aspects and applications of embedded wireless sensor networks.

Jose Rolim is Full Professor at the Department of Computer Science of the University of Geneva where he leads the Theoretical Computer Science and Sensor Lab (TCSensor Lab). He received his Ph.D. degree in Computer Science at the University of California, Los Angeles. He has published several articles on the areas of distributed systems, randomization and computational complexity and leads two major national projects on the areas of Power Aware Computing and Games and Complexity. Prof. Rolim participates in the editorial board of several journals and conferences and he is the Steering Committee Chair and General Chair of the IEEE Distributed Computing Conference in Sensor Systems.

Pierre Leone is Assistant Professor at the Department of Computer Science of the University of Geneva where he is involved in the European CRESCCO project at the TCSensor Lab of Professor Jose Rolim. He is also lecturer at the Engineering School of Geneva where he is involved in a research project on collaborative optimization algorithms on distributed systems. He received the Ph. D. degree in Applied Mathematics from University of Geneva and he has a background of electrical engineer with orientation in Computer Science. He spent one post-doctoral year as a Visiting Lecturer of the Mathematics Department of the Auckland University in New-Zealand.

Max-Planck Institute of Computer Science (MPII) was founded within the Max-Planck Society in 1990. It comprises 5 research groups on algorithms and omplexity, programming logics, computational biology and applied algorithmics, computer graphics, and databases and information systems. Currently, about 100 scientists are working at MPII. The MPII group that will be primarily involved in this project is the databases and information systems group, which has been established in fall 2003 and is headed by Prof. Gerhard Weikum. MPII will contribute to SP2, SP3 and SP6.

Prof. Gerhard Weikum is a Scientific Director at MPII since October 2003. Earlier he held positions at Saarland University in Germany, ETH Zurich in Switzerland, MCC in Austin, Texas, and Microsoft Research in Redmond, Washington. Dr. Weikum has received several best paper awards including the VLDB 2002 ten-year award. He serves on the editorial boards of various journals and book series, including ACM TODS, IEEE CS TKDE, and the Springer LNCS series, and as program committee chair for international conferences like ICDE 2000 and ACM SIGMOD 2004.

Christian-Albrechts Universitaet zu Kiel (CAU) has experience in the design of approximation algorithms for combinatorial optimization problems arising from scheduling problems in communication networks, manufacturing systems, and wavelength assignment problems in optical networks. The team is involved in various collaborations; for example a DFG Graduiertenkolleg on Efficient Algorithms, a bilateral project Procope and in three european projects, APPOL, ARACNE, and CRESCCO. Within the project, CAU will focus on scheduling and routing problems.

Klaus Jansen got his Ph.D. in Mathematics from the Universitat Trier, Germany in 1990 and finished his habilitation in Mathematics in 1994. After temporary professorships and research stays at MPII in Saarbrucken and at IDSIA in Lugano, since October 1999 he is associate professor in computer science at the Universitat zu Kiel, Germany. He has organized a sequence of workshops on approximation algorithms: APPROX 1998 - APPROX 2004 and several workshops in Bertinoro, Dagstuhl, and Oberwolfach. He is coordinator of the EU thematic networks APPOL I - II.

Alexey Fishkin has got his Ph.D. in Computer Science from the University of Kiel, Germany in June 2003. From May 2002 to August 2004, he had a research position at EU-project CRESCCO. Since September 2004 he is a postdoc at MPII Saarbrucken. He aims to design efficient algorithms for combinatorial problems. He works on several wide fronts that include approximation and online algorithms for packing, routing, sequencing, and scheduling problems.

Olga Gerber received her M.Sc. in Computer Science and Applied Mathematics from the Novosibirsk State University, Russia in 2001. After one year as a postgraduate student at the Novosibirsk State University, she is a PhD student at the University of Kiel, Germany, since September 2002. Her research focuses on the design and analysis of approximation algorithms for hard optimization problems.

The Dept. of Mathematics of the **University of Rome "Tor Vergata"** includes several research and teaching activities in Pure and Applied Mathematics and Computer Science. The team involved in this project has scientific connections with most of the other partners and already partecipated to several joint projects (i.e. RTN ARACNE, IST FET CRESCCO). The site will contribute with its expertise on randomized and approximation algorithms, distributed computing and wireless communication to subprojects 1,2 and 5.

Miriam Di Ianni: Laurea Degree in Mathematics (1988) and Ph.D in Computer Science (1993) at the Univ. of Rome "La Sapienza"; Post-Doc at the Computer Science Dept. of Univ. of Rome "La Sapienza" (1994-1997); Researcher at Univ. of Perugia (1997-2001), in 2001 she moved to the Univ."Tor Vergata"; in 2004 she got her habilitation as Associate Professor.

Andrea Clementi: Laurea Degree in Mathematics (1990) and Ph.D in Computer Science (1994) at the Univ. of Rome "La Sapienza"; Post-Doc at Geneva (1994-1996); Researcher at Univ. of Rome "La Sapienza" (1996-1998); Associate Professor at the Univ. of Rome "Tor Vergata" (1998-2001); Full Professor at the Univ. of Rome "Tor Vergata" (2001).

The **National and Kapodistrian University of Athens** (or, in short, University of Athens,) was founded in 1837 and is the oldest University in Greece. The researchers that will participate in the project belong to the Theory group and hey have done extensive research in Theoretical Computer Science both pioneering (such as the paper that introduced the price of anarchy) and deep (such as the k-server result which comprise a chapter in the two standard textbooks on online algorithms). They participate in the IST program FLAGS and national projects in Greece. The site will contribute with its expertise on theory of networks, online algorithms, and game-theoretic algorithms to SP1, SP2, SP3.

Elias Koutsoupias is currently a professor at the University of Athens (Greece). His research interests are on game theory, decision-making under uncertainty, online algorithms, networking theory, database theory, design and analysis of algorithms, and computational complexity. He has published in the top journals and conferences of Theoretical Computer Science and received funding from the National Science Foundation (NSF, USA) and the EU.

Vassilis Zissimopoulos is currently an associate professor at the University of Athens (Greece). He was a professor in the Department of Computer Science at the University of Paris-Nord (LIPN, France). His research interests include the design and analysis of algorithms, particularly approximation and dynamic algorithms, local search and landscapes theory, data location and caching, load sharing and balancing, scheduling and resources allocation.

The project site at **DEI-Padova** comprises four Faculty (G. Pucci, G. Bilardi, A. Pietracaprina (Full Professors); Enoch Peserico (Assistant Professor)) and an average of 8-10 collaborators. Research focusses on all aspects of high-performance computing. The group has been and is funded by European and National projects, and leads the National Centre of Excellence on Advanced Computing Paradigms. The site will contribute with its expertise on advanced computing paradigms to Subprojects 1, 3 and 6.

Geppino Pucci (site leader) (Phd (1993) in Computer Science from U. Pisa (I) is Full Professor of C.S. at U. Padova (I). Previously he was an R.A. at U. Newcastle (UK) (88-90), and Postdoc at ICSI, Berkeley (USA) (90-93). His research in the field of parallel and hierarchical computing (over 50 publications) has been funded by several national and international agencies.

Gianfranco Bilardi (Ph.D. (1985) in Electrical Eng. from U. Illinois Urbana-Champaign (US)), has been an Assist. Prof. (1984-1990) at Cornell U. and is currently Professor (1990-) of C.S. at U. Padova. His research has been reported in over 70 international publications and has been supported by several sources, including NSF, US-JSEP, IBM, Sun Microsystems, CNR, MIUR, EU.

The group in **Swiss Federal Institute of Technology in Zurich (ETHZ)** participating in the project is the Information and Communication Systems Research Group, Institute for Pervasive Computing, Department of Computer Science, ETH Zurich. The group was established in April 1998 and is headed by Prof. Dr. Gustavo Alonso. It currently has eight Ph.D. students. Its main research activities include parallel and distributed systems, advanced database applications (with a special emphasis in scientific, geographic, and electronic commerce applications) and workflow management.

Gustavo Alonso is professor in the Department of Computer Science at the Swiss Federal Institute of Technology in Zurich (ETHZ). Gustavo Alonso holds degrees in Telecommunications Engineering from the Madrid Technical University (UPM-ETSIT), and a M.S. (1992) as well as a Ph.D. (1994) from the University of California at Santa Barbara. After graduating, he was a visiting scientist in the IBM Almaden Research Laboratory in San Jose, California. In September 1995 he joined ETH where he has since then lead several projects in databases, workflow management, replication, and advanced applications. His research interests include Web Services, grid and cluster computing, databases, workflow management, scientific applications of database and workflow technology (for geographic, astronomical, and biochemical data), pervasive computing and dynamic aspect oriented programming.

Cesare Pautasso is senior researcher in the Information and Communication Systems Research Group, Institute for Pervasive Computing, Department of Computer Science, ETH Zurich. He holds a Ph.D. from ETHZ in 2004, and a "Laurea in Ingegneria Informatica" (Computer Science Engineering Degree) from Politecnico di Milano (2000). His main areas of interest include Web services and Grid Computing.

At the site of the **Universitat Politecnica de Catalunya** (UPC), Spain, work has been done on complexity theory, layout problems, sensor networks, adversarial queuing theory, graph algorithms, randomized algorithms, parallel algorithms and optimization. The group consists of roughly 15 full time researchers, and is (or has been) involved in a variety of European research projects, i.e., ALCOM-FT, FLAGS, and DELIS. The site will contribute with its expertise on adversarial queuing theory, protocols for smart dust and sensor networks, parallel algorithms, and randomized models for ad-hoc networks to WP 1.2, 1.3, 1.4, 1.5, 3.2, 3.3, 5.2, 5.3.

Maria Serna is Associated Professor in Computer Science at the UPC since 1991. Her scientific interests are algorithm and complexity of graph theoretic problems, probabilistic methods in algorithms, protocols and algorithms for sensor networks, and adversarial queuing theory. She has published in major journals and has presented papers at major conferences. She has served in the program committees for prestigious conferences. She is co-author of 2 text-books and over 60 publications.

Josep Diaz is Full Professor in Computer Science at the UPC since 1984. His scientific interests are probabilistic techniques in the design and analysis of algorithms and heuristics, layout problems and complexity theory. He has published in major journals and has presented papers at major conferences. He is in the editorial board of two journals and has been a member of program committees for prestigious conferences. He is co-author of 3 text-books and over 100 publications.

The Algorithm Engineering group at the Department of Computer and Systems Sciences of the **University of Rome "La Sapienza"** has a strong brackground in the design of algorithms for network applications, distributed computing, combinatorial optimization, massive data sets, and in the analysis of the structure of large and complex networks. The group is formed of 6 faculties and several post-docs and PhD students. The group is actively cooperating with several research groups in Europe, Israel and United States.

Alberto Marchetti-Spaccamela is Full Professor since 1987 and since 1991 he is at the University of Rome "La Sapienza". He has written eight books, has (co-)authored about 70 international publications and has edited books and special issues of scientific journals. His current research interests are the design and the analysis of algorithms and their applications to computer networks and to Internet. He coordinates and has coordinated of many National and European research projects.

Stefano Leonardi took his PhD in Computer Engineering in 1996 and is currently Associate Professor at Universita di Roma "La Sapienza". He has co-authored about 40 publications in leading international journal and conferences. His main research interests are in the design and analysis of algorithms for network resource management and scheduling problems and in the study of the structure of large and complex networks and in the game theoretical aspects of networks. He is the coordinator of the unit of Rome of the IST-FET projects COSIN, DELIS and APPOL2.

The research group SCD-COSIC (Computer Security and Industrial Cryptography) belongs to the Electrical Engineering Department of the **Katholieke Universiteit Leuven**. Since 1979, COSIC (35 members) performs research on the design, evaluation, and implementation of cryptographic algorithms and protocols. The group has a broad expertise from highly mathematical research to real-life applications in the area of information and network security. The members of COSIC have published more than 300 international articles. COSIC participated to more than 15 EU-sponsored research projects in the area of cryptology, information security, mobile system security, conditional access, e-voting,… For the IS projects NESSIE, STORK and ECRYPT, COSIC has been the coordinator. The Rijndael algorithm, which has been selected by NIST to become the AES (Advanced Encryption Standard), a worldwide *de facto* standard has been invented by a member and an old-member of COSIC. This site will contribute with its expertise in theoretical and applied expertise in cryptology, information and network security.

Prof. Bart Preneel is a professor at K.U.Leuven (Belgium) and a visiting professor at the T.U.Graz. His main interests are cryptography, information security, and wireless communications. He has authored and co-authored more than 160 articles in international conference proceedings and journals, and is inventor of one patent. He has participated in 15 EU-sponsored research projects, for 3 of these as the project coordinator. He is currently project manager of ECRYPT, the Network of Excellence on Cryptology and Watermarking. He is Vice President of the International Association for Cryptologic Research, president of the Leuven Security Excellence Consortium (L-SEC vzw.) and an editorial board member of J. Cryptology and ACM Trans. on Information Security.

Dr. Svetla Nikova is a postdoctoral researcher at the COSIC research group of the K.U.Leuven with research experience in coding theory and cryptology. Her current research interests include secret sharing schemes, metering schemes, multi-party computation, and cryptographic properties of Boolean functions. She was also active in the Flemish project “Anonymity and Privacy in Electronic Services (APES).” She has authored and co-authored more than 35 articles in international conference proceedings and journals.

ARES is an **INRIA** project of the Rhτne-Alpes research unit and located in the CITI laboratory of the INSA of Lyon. The goal of this project is to model and to develop architectures and software support for hybrid wireless networks. 7 researchers work in this project. We currently participate in six national projects and two European projects.

Isabelle Guerin Lassous is a INRIA researcher since 2000. She received her PhD in 1999. She worked in different INRIA projects (HIPERCOM, REMAP and then ARES). Her research interests are in the area of wireless network.

Stepahe Ubeda is a professor at the INSA de Lyon, Department of Telecommunications, Services & Uses since 2000. He received his PhD in 1993 and the Habilitation ΰ Diriger des Recherches in 1997. His research interests are in the area of algorithmic for telecommunication networks. He is the header of the INRIA ARES project and of the CITI Lab (Insa de Lyon).

Eric Fleury is a professor at the INSA de Lyon, Department of Telecommunications, Services & Uses since 2003. He received his PhD in 1996 and the Habilitation ΰ Diriger des Recherches in 2002. His research interests are in the area of wireless network (ad hoc, sensor), pervasive communication and next generation communication network.

DIMATIA Centre, Charles University Prague

DIMATIA (for Discrete Mathematics, Theoretical Computer Science and Applications) Centre coordinates research in the named fields in the Czech Republic and it has a numerous foreign associates. The activity involves most of the theoretical branches of combinatorics and discrete mathematics as well as related areas of discrete geometry, graph drawing, topology and algebra as well as applications (for example partition problems and channel assignment problem).

The group involves 7 senior researchers and more than 15 graduate students. The group is involved in the EU-Research Training Network programme COMBSTRU and in various bilateral grants. The scientific output is about 50 research articles annually.

The scientific potential of Prague site will allow contributions and offer expertise on theoretical CS to SP1, SP3, SP4 and SP5.

Jaroslav Nesetril, professor, Dr.,Dr.h.c., director of DIMATIA works in graph theory, combinatorics, complexity, theory of structures, image processing, about 250 scientific publications, including 4 books, editor of 5 international journals; Prague site director.

Jak Kratochvil, professor, Dr., chairman of the Department of applied mathematics, works in graph theory, complexity, graph drawing, about 70 research papers, chair of program committee of MFCS 2004, GD 1999, editor of 2 international journals.

The Department of Computer Science is one of the 5 departments in the School of Pure and Applied Sciences of the **University of Cyprus** (founded.in 1989). It has 16 faculty members, about 30 researchers and around 120 graduate students. It is an active, research-oriented department, participating in several EU and nationally funded projects. The site will contribute with its expertise on Algorithmic Game Theory, Distributed Computing, Internet Computing and Formal Methods to Subprojects WP1.1, WP1.2, WP1.3, WP1.4, WP1.5, WP2.2, WP2.3, WP3.3, and WP5.2.

Marios Mavronicolas received the Diploma in Electrical Engineering (Summa cum Laude) from the National Technical University of Athens, Greece, in 1985, and the M.A. and Ph.D. degrees in Computer Science from Harvard University, USA, in 1988 and 1992. He has also taught at the University of Crete, Greece, and at the University of Connecticut, USA. His current research interests focus on Algorithmic Game Theory, Algorithms and Complexity, and Distributed and Parallel Computing.

Anna Philipou received the Bachelor of Arts degree in Mathematics and Computation from Oxford University, UK, in 1992, and the M.Sc. degree in Parallel Computers and Computation and the Ph.D. degree in Computer Science from the University of Warwick, UK, in 1993 and 1996. She has also worked as a Postdoctoral Research Fellow at University of Pennsylvania, USA. Her research interests include Algorithms and Complexity, Concurrency Theory, Specification and Verification, and Formal Methods.

Cybernetica is a Research and Development company founded in 1997 during reorganisation of the Institute of Cybernetics (IOC) of the Estonian Academy of Sciences. Activities in the field of information security started in 1992 with establishing the Information Technology Department at the IOC. The company became the primary contractor for the national Information Security program. Since 1992, the information security research group has been rapidly developing and currently involves 26 researchers and system engineers (five PhD-s, four PhD students, eight MSc-s and three pursuing their MSc degrees). Cybernetica Data security laboratory, founded in 1998 in cooperation with the University of Tartu, supports also post-graduate studies in the field on Information Security and Cryptology at the University of Tartu. The research goals are theoretical and practical aspects of Cryptography, including Public Key Infrastructure, Time-Stamping, and database security. Cybernetica has been involved in the deployment of the national ID card project, as well as in the Estonian Electronic Document Management Programme. The site will contribute with its expertise on cryptography to SP4.

Dr. Ahto Buldas, senior researcher. MSc in Computer Engineering (1992, Tallinn Technical University). PhD in Mathematics (1999, Tallinn University of Technology). From 1989 to 1992 - research in computer engineering (testing and verification of Boolean circuits). From 1993 to 1996 – research in Algebraic Graph Theory. From 1997 to 2004 - research in applied and theoretical cryptography. From 1997 focusing on digital time-stamping. Professor at Tartu University and Tallinn University of Technology. Was included in the Working group of Estonian State Chancellery for developing the Estonian Digital Signature Act. Current research interests: time-stamping, hash functions, universally composable security, practical security and risk analysis.

Dr. Peeter Laud, senior researcher. MSc in mathematics (1998, University of Tartu). PhD in computer science (2002, University of Saarland). From 1998 to 2004 – active research in theoretical and applied cryptography, as well as in program analysis. Best paper award at ETAPS2003. Organiser of several research-oriented seminars and summer schools. Current research interests: cryptographic protocols and their universally composable security, program analysis, time-stamping.

B.5.1 New participants

-- Not applicable --

B.5.2 Sub-contracting

The project includes one sub-contract from the University of Ioannina for the University of Milan.
The **University of Milano** (UNIMI) is one of the largest universities in Italy with a total of 75,000 students. The work described in this proposal will be carried out by the Database Systems Group, which is a research group within the Department of computer Science of the University of Milano. The Database System Group consists of 9 full time members, working on advanced data management systems. The group is currently involved in several projects, supported by various sources, such as the European Commission, the Italian Ministry of Research and Education, the Italian Telecom, the Ministry for Public Administration, Bellcore and Microsoft Research. The site will contribute with its expertise on security and data management for large distributed systems and will participate in WP 2.1 and WP 4.1. This contribution will be coordinated with the one of Ioannina.

Prof. Elisa Bertino is full professor in the Department of Computer Science of the University of Milano, where she heads the Database Systems Group. Since January 2004, she is also a professor in the Department of Computer Sciences, of Purdue University and Research Director of the Center for Education and Research in Information Assurance and Security (CERIAS) one of the world's leading centers for research and education in areas of information security. Professor Bertino is a co-editor in chief of the VLDB Journal and serves on the editorial boards of several journals. She is a Fellow of the IEEE and an ACM Fellow. She also received the IEEE Computer Society Technical Achievement award in 2002.

B.5.3 Other countries

-- Not applicable --

B.6 Description of project management

All activities in the Integrated Project AEOLUS will be managed by the Coordinator, the Sub-project Leaders, the Workpackage Leaders, the Coordinating Committee (CC), the Consortium Board (CB) and the Technical Committee.

The Coordinator of IP AEOLUS will be Christos Kaklamanis from UoP. The Coordinator oversees all the activities within the project, monitors and guides the overall direction of the project, runs the daily business, calls and leads the meetings of the CC and the CB. He acts as the intermediary between the Consortium and the European Commission. He receives all payments made by the Commission and administers the Community contribution regarding its allocation among the participants and activities in accordance with the contract and decisions taken by the Consortium. He ensures that the appropriate payments are made without unjustified delay. He keeps accounts making it possible to determine at any time what portion of the Community funds has been paid to each participant, and informs the commission of the distribution of the funds and the date of transfers. He initiates activities like organization of schools and workshops. He will be assisted by research and administrative personnel at the coordinating site.

The activities in each subproject will be coordinated by a Subproject Leader. The Subproject Leader oversees the research work within the subproject, reports to the Coordinator and the Consortium Committee, and coordinates the dissemination activities of the subproject.

The Coordinating Committee (CC) consists of the Coordinator and the six Subproject Leaders: Elias Koutsoupias (UOI) for SP1, Jean-Claude Bermond (CNRS) for SP2, Burkhard Monien (UPB) for SP3, Giuseppe Persiano (UNISA) for SP4, Paul Spirakis (CTI) for SP5 and Fabrizio Davide (TILS) for SP6. It is chaired by the Coordinator. The consortium may change the members of the CC and/or add new members. The members of the CC members communicate frequently about all ongoing issues, meet regularly, at least once per year, and meet otherwise when deemed necessary. The CC is the steering committee, coordinates the cooperation among the subprojects, and makes recommendations to the Consortium Board.

The Consortium Board consists of one representative for each partner, and is chaired by the Coordinator. It meets at least once a year and otherwise it communicates about all important issues of the project. The Consortium Board can change the members of the CC and makes decisions on the project based on the recommendations of the CC and/or the Coordinator. In general, decisions should be reached through consensus while in case of serious disagreement the Coordinator will be responsible for making the final decision.

Each workpackage is lead by a Workpackage Leader who coordinates the activities within the workpackage and is, in general, responsible deliverables related to the workpackage.

At the beginning of the project the Consortium will form a Technical Committee that will assist in the coordination of the activities in Subproject 6.

An overview of the organizational structure is presented in the following Figure

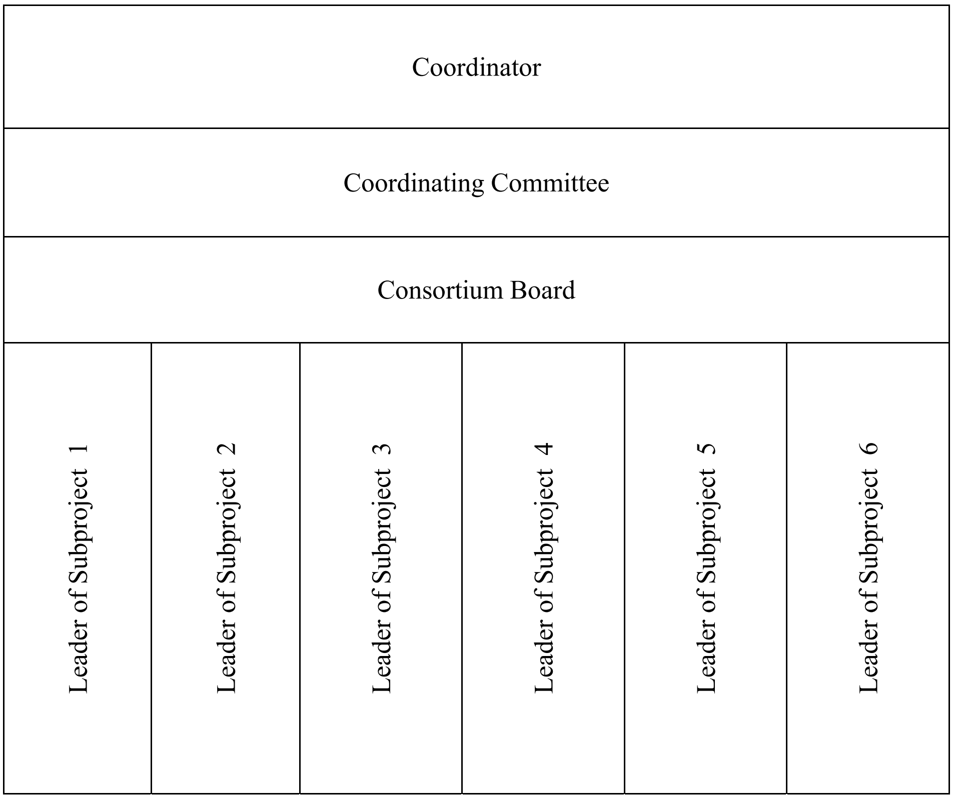

Figure. Organisational Structure of Project Management

B.7 Project resources

B.7.1 IP Project Effort Form

IP Project Effort Form
Full duration of project

	UOP	TILS	CNRS	UPB	CTI	UNISA	UOI	CUI	MPII	CAU	UDRTV	UOA	UNIPD	ETHZ	UPC	UDRLS	KUL	INRIA	DIM	UCY	CYB	TOTAL
RTD/Innovation activities																						
Sub-project 1	29		37	8	9	11		16		5	9	37	21		32	16		21	27	27		305
Sub-project 2	25		27	11	14	13	32	11	8	8	8	13					1			19		190
Sub-project 3	3		11	27	7	8	35		29	27		11	16	37	5	16			8	8		248
Sub-project 4	4				5	51	16									6	35		12		59	188
Sub-project 5	19		16	10	21		8	37		5	35				8	11	1	24	8	4		207
Sub-project 6	32	80	21	16	16	21	21	16	11	11	12	3	11	11	3	11	11	11	1	2	5	326
Total RTD/Innovation	112	80	112	72	72	104	112	80	48	56	64	64	48	48	48	60	48	56	56	60	64	1464
Consortium management activities																						
Sub-project 1												4										4
Sub-project 2			4																			4
Sub-project 3				4																		4
Sub-project 4						4																4
Sub-project 5					4																	4
Sub-project 6		4																				4
IP Management	48																					48
Total consortium management	48	4	4	4	4	4																72

TOTAL ACTIVITIES	160	84	116	76	76	108	112	80	48	56	64	68	48	48	48	60	48	56	56	60	64	1536

B.7.2 IP management level justification of resources and budget

The total cost of the project has been calculated to 7,112,000 Euros for a period of four years; this corresponds to 5,772,000 Euros of requested funding.

The main part of the budget will be allocated to research and administrative personnel. In total, 1536 person-months will be covered by the project; most of them (1464) correspond to research activities and software development. The actual effort is expected to be significantly higher. For example, the table in the previous page does not include the effort of personnel with permanent positions at the academic partners (e.g., faculty members in Universities). In particular we expect that the total effort devoted to the project by academic sites will be, on the average, at least 40% higher than what is stated in the IP Project Effort Form. Thus, we believe that the project will mobilize the critical mass of researchers and engineers necessary for accomplishing its ambitious goals.

Furthermore, although most of the research work in the project will have a foundational nature (investigation of algorithmic principles using theoretical analysis and experimentations), about 1/5 of our total effort (326 person-months) will be devoted to the development of a Programmable Overlay Computing Platform and related applications in the context of subproject SP6. We believe that this proportion of long-term research and engineering work towards "proof-of-concept" implementations is ideal for a FET project with concrete objectives as those set by the Global Computing Proactive Initiative.

In addition to the research effort, such a large Integrated Project requires heavy involvement with administrative and management activities. We have allocated 72 person-months for management activities at the coordinating site and the subproject leaders. In particular, the budget of the coordinator will cover part of the salary of a young researcher and administrative staff working in the project and supporting the coordinator with the every-day management of the project, the update of the project's Web-site, etc. A small budget has also been allocated to subproject leaders; this is required in order to support management activities in the subprojects.

The total budget of the project also includes an amount of 48,000 Euros (allocated to the coordinator) for the organization of schools and workshops; this amount is sufficient for the organization of at least one school/workshop per year as we have planned in the context of our training activities.

Finally, implicit in the budget of each partner are amounts for travel and equipment. We plan to use a part of the budget for traveling either to conferences for disseminating our research results or for visits between partners in order to further strengthen our synergies. Some equipment will also be needed since most of the partners are involved in implementations and prototype development. The exact amounts in these cost categories will be defined by the partners during the project and will follow their accounting rules.

B.8 Detailed implementation plan – first 18 months

a) Detailed implementation plan introduction

In the following we present the research work to be performed within the six subprojects during the first 18 months of the project. In addition, each one of the subprojects will contain a special work package called "Subproject management and dissemination activities". These workpackages will run for the whole duration of the first 18 months, they are led by the subproject leader, and are devoted to the management of the corresponding subproject. The objective is to guarantee the successful completion of the subproject within the agreed time, costs and quality requirements, ensure compliance with EC standards and procedures for project management and tracking, create effective channels of communication among the consortium partners, and effectively handle intellectual property issues. In addition, among the objectives of these workpackages is the coordination of dissemination activities (e.g., publications in international scientific conferences and journals).

Sub-Project 1: Paradigms and Principles

Leader: UOA; Participants: UOP, CNRS, UPB, CTI, UNISA, CUI, CAU, UDRTV, UNIPD, UPC, UDRLS, INRIA, DIM, UCY

This subproject aims to provide a theoretical background to study, design, and build efficient, stable, and programmable overlay computers on top of successful global systems such as the Internet. By its nature, the inherent emphasis on theoretical work, and above all, the fact that this research area is at the forefront of the general field of Theory of Networks, a plan for carrying out the work can only be tentative. Progress in resolving some of the issues may reveal new fruitful avenues for research and exploitation. We plan to assess continuously the new developments in this area and to update our goals accordingly.

The research work within the subproject is divided into five workpackages:

- WP 1.1: Structural properties of global and overlay computers
- WP 1.2: Coping with partial information
- WP 1.3: Coping with selfishness
- WP 1.4: Stability and fault tolerance
- WP 1.5: Generic algorithms

WP 1.1: Structural Properties

Leader: UCY, Participants: UOP, CNRS, CTI, CUI, UDRTV, UOA, UNIPD, DIM

The objective of this workpackage is to study the topological properties of global and overlay computers, their substructures and other organization issues, and embeddings that relate to the implementation and realization of overlay computers into global ones.

We plan to build upon the existing work on the topological properties of the Internet and the Web and to design searching and communication primitives that take advantage of their structural properties. We also plan to study the "geometry" of global systems and to seek improved algorithms for clustering and partitioning problems associated to the efficient implementation of overlay computers. Despite the rich history of metric embeddability, there are major open algorithmic problems that relate to low-distortion embeddings and we plan to attack them. Low distortion, which guarantees good approximation for every objective function is desirable, but in many cases impossible. In this case, we will pursue special solutions that may have unbounded distortion but still can preserve the given objective functions.

WP 1.2: Coping with Incomplete Knowledge

Leader: UOA, Participants: UOP, CTI, UNIPD, UPC, INRIA, DIM, UCY

The goal of this workpackage is to use competitive analysis and the theory of distributed systems to study the novel online problems that are at the core of global and overlay computers, such as distributed and robust data structures, transparency and information hiding, and reputation mechanisms.

We plan to study simple primitives that allow the maintenance of data structures in selfish, distributed, and dynamic environments. The problems associated to efficient data structure implementations with provably good competitive ratio are among the hardest ones even for centralized data structures. We plan to do research for simple centralized data structures for searching and servicing requests and to extend the results for dynamic distributed environments. We also plan to investigate the role of information hiding in achieving good online solutions. A related problem of central importance in global computing is the locality issue (where the entities have only partial/local information about the system). We plan to study from the competitive analysis point of view the impact of locality and to design online algorithms that alleviate it. Finally, we plan to study reputation mechanisms: what to distill from the history of past transactions, where and how to maintain it, and how to predict the future behavior of users. We will model this problem both as a traditional online problem (with a powerful adversary to obtain worst-case bounds) and as the more realistic problem where the adversary controls only the order of the transactions and the behavior of users is determined solely by their selfish nature.

WP 1.3: Coping with Selfishness

Leader: UOA, Participants: UOP, UPB, CTI, UNISA, CAU, UDRTV, UPC, UDRLS, UCY

The goal of this workpackage is to study ways to improve system performance in global systems with selfish entities.

We plan to extend existing results on the price of anarchy on systems with many selfish users that use a set of resources. We also plan to study the impact of information regimes on the performance of selfish systems and to come up with algorithmic solutions that improve the coordination of selfish users and the performance of the whole systems based on appropriate manipulation of the information flow. We also plan to investigate the fundamental problem of "fairness vs. efficiency" in global and overlay computers and to study how incomplete information can help to achieve solutions that are both fair and efficient even when this is not possible in the full information regime. Finally, we plan to study ways to improve system performance through coordination mechanisms and mechanism design techniques and to extend these results to dynamic environments.

WP1.4: Stability and Fault Tolerance

Leader: UPC, Participants: UOP, CNRS, CTI, CUI, DIM, UCY

The goal of this workpackage is to study ways to improve stability and fault-tolerance in global systems based on adversarial queuing theory and the theory of distributed systems.

We plan to use the framework of adversarial queuing theory to study stability issues of global systems. Adversarial queuing theory studies static networks of queues when the injection of packets is controlled by an adversary. The existing results do not take into account the ever-changing environment of global systems. We plan to extend the results to networks where the set of nodes and connections change over time. More specifically, we plan to attack the questions of whether certain topologies are stable and whether specific protocols are stable. Also, we plan to study the question on how much the extra communication load, which is required to support functionalities at the overlay computer level, affects the stability of the underlying global computer. We then plan to study the issue of improving the stability of global systems by appropriate mechanisms built on overlay computers. Finally, we plan to study distributed algorithmic primitives that allow for detection and correction of faults. We hope to develop a theory to address trade-offs between the impact of a failure (how fast the fault is recovered and how many connections must be interrupted during the recovery procedure) and the quantity of resources allocated to protect from faults.

WP 1.5: Generic algorithms

Leader: CNRS, Participants: UOP, CTI, UDRTV, UOA, UPC, UDRLS, INRIA, UCY

The goal of this workpackage is to study classical problems in combinatorial optimization that can have an impact on the understanding, modeling, and design of global and overlay computers. It will also include new algorithmic issues that may arise from new developments during the duration of the project.

A central issue in global systems is connectivity, a problem that has been studied extensively in classical algorithmic theory. However, the dynamic environment of global systems gives rise to new challenging problems. First of all, we plan to find appropriate mathematical models to capture the issues involved. We will then try to extend classical results to these models. It is clear that the algorithms for this type of problems must be distributed. Considering that the nodes of these systems have only a partial, and sometimes erroneous, view of the system, we have to aim for approximate solutions but with certain guarantees (such as t-spanners). We plan to develop approximation algorithms that address this issue. We also plan to study routing and scheduling problems using multi-criteria objectives. These problems become qualitatively different when we are interested in optimizing not a single objective (for example, delay) but multiple objectives (for example, delay and reliability). Especially in systems with multiple users where fairness is an issue, we have to devise algorithms that try to optimize the objectives of all of them. When this is not possible, we plan to study approximate solutions. Finally, we plan to develop a theory and investigate algorithms that have to base their decisions on samples of the data. The problem is how to select the appropriate sample efficiently (which, given the distributed nature of the system, is a novel question in the interface of Algorithms and Statistics) and how to analyze it.

Sub-Project 2: Resource Management

Leader: CNRS; Participants: UOP, CTI, UPB, UNISA, UOI, CUI, MPII, CAU, UDRTV, UOA, KUL, UCY.

The goal of this subproject is to address novel and challenging algorithmic issues for efficient resource discovery and querying like construction of overlay networks, query routing and execution, and for sharing critical resources like bandwidth. Our work will also include mechanism design for coping with selfish behavior when allocating resources in a distributed, uncoordinated system such as a global or overlay computer. The research work within the subproject is divided into three workpackages:

- WP2.1 Resource discovery
- WP2.2 Critical resource sharing
- WP2.3 Mechanism design

WP2.1: Resource Discovery

Leader: UOI; Participants: UOP, CTI, MPII.

In this workpackage, we shall explore advanced methods for resource discovery that take into account the structure and content of resources. A central objective is supporting resource discovery based on more advanced queries than single attribute-value or keyword search. The focus will be on scalable, distributed and dynamic solutions that take advantage of the embedding of virtual networks into real ones. To this end, we shall build upon the paradigms and principles of global overlay computers as studied in SP1. The 18-months objective is to achieve sufficient progress towards the overall objective which will be measured by a thorough survey on the topic and initial research results published in the appropriate forums.

During the first 18 months, the research work in this workpackage includes:

- The construction of overlay networks that connect semantically related nodes most efficiently but at the same time provide highly efficient routing to all nodes with logarithmic worst-case behavior in terms of both of space overhead per node and routing hops. Distributed clustering and novel embeddings based on content will be studied.
- Processing of advanced queries on top of the constructed networks. Work includes both query routing (that is determining nodes with matching resources and executing the query at the appropriate nodes.

WP2.2 Critical Resource Sharing

Leader: CNRS; Participants: UOP, UPB, CTI, CUI, MPII, CAU, UCY.

The goal of this workpackage is to provide protocols that allow the user of a global computer to communicate in a transparent way and according to some quality of service and cost. Those protocols shall use the underlying global computer resources efficiently and in a scalable way.

We will investigate the current and forthcoming issues for critical resource sharing in complex hierarchical networks. Even if bandwidth is likely to remain critical (but attached to some quality of service parameters), we will also consider other potential critical resources (e.g., routers). We will then develop techniques to ensure an efficient use of bandwidth in heterogeneous hierarchical networks. Our goal will be to ensure a good mapping of bandwidth requests expressed at the overlay computer level onto the underlying global computer. For this, we shall extend classical flow-like algorithms (routing) to the case of logical networks in which connections are labeled by various parameters (delay, cost, reliability). Also, we will design tools and models to allow the global computer agents to share or trade bandwidth under quality of service or financial cost constraints. For this, we will use concepts from WP1.2 and WP2.3 to cope with social and economic aspects. We will extend the algorithms in order to make them suitable for dynamic environments (scalability, fault tolerance). Here, we will trade optimality for locality and efficiency, and we will survey and extend on-line algorithmic techniques.

WP2.3 Mechanism design

Leader: UNISA; Participants: UPB, CTI, UDRTV., UOA, KUL, UCY

The goal of this workpackage is to study mechanisms for resource management in overlay and global computers in the presence of selfishly acting entities. In the first 18 months we plan to focus on

1. Truthful mechanisms for problems for which the VCG paradigm cannot be applied. This work will entail both the design of new approximation algorithms that can be used as part of truthful mechanisms. In particular, we will consider several aspects like *approximation guarantee*, *online vs. offline* mechanisms and *frugality* (i.e., the amount of currency or meta-resource the mechanism needs to provide to the agents in order to guarantee optimal solutions). One of the questions we will try to answer is the following one: under which conditions is it possible to transform an existing algorithm into a mechanism with the same performance guarantee? Another interesting question here is *how much information do we need to solve a problem*? Eventhough the works [NisRon99, AuletalICALP04] give strong evidence that certain additional knowledge allows for much better performances, these works do not give a clear trade-off yet.

2. Cost-sharing problems [JaiVaz01] model situations in which a set of (selfish) users can access a certain type of resource/service by *jointly* contributing to its cost. The task of a mechanism is to select which users will be actually granted and at which price. Users can also cooperate among themselves and form *coalitions* in which they can be untruthful in order to improve the utilities of other users in the same coalition. We will focus on mechanisms for *multicast cost sharing*. Here, a set of users are connected via a common network (wireless or wired) and transmission from a source node s to a subset of recipients Q requires a certain cost (i.e., the cost of a multicast tree). The mechanism has to decide which multicast to build, which user to service and how much each of them should pay in order to recover the total cost of the transmission. We will concentrate on multicast cost-sharing problems. In particular, we will try to compare techniques and results for the case of wired networks [FeiPapShe00, JaiVaz01] to the case of wireless ones [PenVenSIROCCO04, Bilo, PenVenWAOA04] and study the performance degradation of mechanisms for wired networks when applied to wireless ones. A key point here is to determine the loss of performance of certain mechanisms adapted from the wired case. For instance, the technique in [JaiVaz01] can be adapted to the wireless model [Bilo] but this leads to approximate solutions only. We will also further pursue the possibility of obtaining good cost-sharing mechanisms via *non-crossmonotonic* allocation schemes [PenVenWAOA04].

Sub-Project 3: Sharing Information and Computations

Leader: UPB; Participants: CTI, UOI, UOA, UNISA, UDRLS, UNIPD, CAU, MPII, CNRS, ETHZ, UPC, UCY, DIM.

The goal of Subproject 3 is to develop concepts, strategies, algorithms, and data structures, and to design software tools that can help in sharing information and computation among the computing entities of a global computer or an overlay computer. Because of the size, dynamic nature, selfishness, unpredictable changes and limited availability of data in such a computer, this is an important and challenging emerging research area.

The goals of Subproject 3 are manifold. We plan to contribute to the theoretical foundations by developing strategies for efficient allocation of data, efficient allocation of computation, and an efficient allocation of resources over time. The research work within the subproject is divided into four workpackages:

- WP3.1 Distributed data management
- WP3.2 Load management
- WP3.3 Scheduling
- WP3.4 Workflows and services

WP 3.1 Distributed Data Management

Leader: MPII; Participants: UPB, CTI, UOI, CAU, UOA, ETHZ

Replication and caching are well-studied problems in distributed computing. They offer distribution transparency, scalability, and fault-tolerance. In global computing, new issues arise due to the enormous scale and dynamic nature of the environment. Furthermore, there is a close interplay between replication and the structural properties of the overlay network. The objectives of this workpackage for the first 18 months are to develop, theoretically analyze and experimentally validate appropriate strategies for caching, replicating, and proactive dissemination of data to the overlay computers to improve performance, availability, fault-tolerance, and data freshness.

We plan to investigate to which extent the strategies for caching, replication, and proactive dissemination depend on and may be tailored to the kinds of data items and their access characteristics taking into account the large scale and high dynamics of overlay computers. We will consider different kinds of data items: from simple unstructured data containers like files and Web documents to relational or XML databases, and also metadata and ontological meta-metadata.

We will study caching and replication strategies in which the determination of the number of replicas for data objects depends on the popularity of the object. We will attempt to develop algorithms combining the characteristics of existing approaches like uniform replication, proportional replication, and square-root replication. Our work on proactive dissemination will be based on epidemic algorithms and investment-based policies.

The work within the workpackage will specifically focus on the dynamic determination and adaptation of replication degrees and adaptive placement of replicas. This includes the definition of mathematical models and methods for choosing replication degrees and replica placement and the design of algorithms for dynamic load-adaptive adjustments. Furthermore, we will also develop new strategies for replica maintainance and proactive dissemination. Our work will include an in depth survey of the state-of-the-art while in addition to the design and specification of new algorithms, we will also perform their theoretical analysis, implementation, and experimental validation.

WP 3.2 Load management

Leader: UPB; Participants: CTI, CTI, MPII, CAU, UNIPD, ETHZ, UPC, URDLS

Load management is one of the key problems that must be addressed to efficiently use a global computer or an overlay computer for hard computational applications. Because of the size and the distributed nature of such systems, there is no central authority to manage this task. One way to come up against this, is the use of local iterative load balancing algorithms. The balancing flow quality as well as the convergence behavior of such algorithms strongly depends on the structural properties of the underlying communication system. The analysis of this coherence is a major task within this subproject.

Moreover, in some scenarios, many different platforms can be available to run a single application. We plan to develop a parametric model of computation where the model parameters reflect phenomena affecting performance depending on the architectural features of the target platform. Therefore, algorithm design and tuning can aim to optimize performance as a function of these algorithmic parameters. The advantage of this approach is that it enforces an explicit link between algorithmic parameters and architectural ones. While highly desirable, this may not be always feasible, since many phenomena affecting performance cannot always be easily interpreted in terms of architectural parameters.

Another hot topic is the use of peer-to-peer networks for parallel computing. Here, one has to address the problems based on the specific properties of such networks (peers may join or leave or fail etc.) in order to provide a well working system where one can efficiently execute parallel programs.

We want to analyze the structural and spectral properties of large self-organizing global computers based for example on the power-law distribution. By examining the relation between these properties and the quality of diffusive-based load balancing algorithms, we will develop algorithmic solutions and give the specifications for prototype implementations of basic load management tasks.

We expect to lay the ground towards a general framework for adaptive software by experimenting the combined approach of analytical and empirical search for tuning on a number of relevant case studies. In particular, we will try to replicate the results obtained by the ad-hoc approaches at the base of a number of important software packages (FFTW and Linear Solvers for Finite Element Methods such as MUMPS, among others). We expect our preliminary strategies to exhibit a certain dependence on the application at hand and will strive to improve the generality of the proposed approach in the later stages of the project.

Moreover, we will implement the BSP paradigm so that parallel programs can be executed efficiently on a large-scale distributed peer-to-peer network. In order to do this, we will design a prototype that includes process migration and simple methods for load balancing (the prototype itself will be delivered in the context of WP6.2).

WP 3.3 Scheduling

Leader: CAU; Participants: CNRS, UPB, CTI, UNISA, MPII, UOA, ETHZ, UDRLS, DIM, UCY

The Internet, a collection of interconnected networks with different scale, transmission and hardware, provides a global communication medium for a huge number of computing entities. Most networks in the Internet use certain common protocols, between the so-called application, transport, internet, and host-to-network layers, and provide certain common services. From one side, the existing layered protocol structure can be adopted to provide some new global services. From another side, in the design of protocols it is quite a challenging task to provide some particular performance guarantees on Quality of Service (QoS) between the layers. Indeed, using a correct allocation of layers' resources over time, or simply scheduling, can significantly improve the performance of protocols.

A basic problem here is to allocate few resources, which are available to a layer, over time such that some specific guarantee on QoS, e.g., maximum throughput, minimum delay, etc., is satisfied. The main goal of this workpackage is to provide efficient algorithmic solutions for that kind of problems arising in the design of protocols for different layers and respective QoSs. We plan to develop suitable scheduling models and design algorithmic solutions that can help to improve the performance of protocols. The results of this activity will add to the current state of theoretical knowledge on these issues as well as provide a vial support to the technological design activities in subproject SP6.

WP 3.4 Workflow and Services

Leader: ETHZ; Participants: UOI, MPII.

The main objective of this workpackage is to investigate issues related to the efficiency, trancparency, and scalability in the execution of large scale computations treating them as processes similar to these found in workflow management systems. The notion of processes allows to model sequences of invocations of computer programs or applications in a distributed and heterogeneous environment as well as to capture the corresponding data exchanges between these programs.

In this workpackage we shall explore (i) the use of services for accessing data in overlay computers, (ii) the use of workflow for the definition and execution of composite web services, and (iii) the capabilities of a global computer and its overlay structures for dynamic self-configuration and automatic adaptation to system, workload, and application dynamics. We will develop new methods and improve existing techniques for automatic migration of failed tasks, system awareness (to node configurations, network bandwidth, load, new clusters becoming available or unavailable, replication of system critical information, general fault-tolerance, and the dynamic and automatic reconfiguration of global workflows to such rapidly evolving conditions.

The work will coordinate with the developments made in other workpackages such as WP3.1 (Distributed Data Management), WP3.2 (Load Management), and WP3.3 (Scheduling). The techniques developed will be used in the specification of an autonomic computing system, i.e., a generic engine for specifying computations of overlay networks of Web services and additional services supporting autonomic execution, flexible enough to incorporate new services as they appear.

Sub-Project 4: Security and Trust Management

Leader: UNISA; Participants: CTI, UOI, UDRLS, KUL, DIM, CYB

The research work within the subproject is divided into four workpackages:

- WP4.1 Trust management
- WP4.2 Privacy, identity and anonymity
- WP4.3 Secure distributed computation

WP 4.1 Trust Management

Leader: UNISA, Participants: CTI, UOI, KUL, CYB.

One major objective is to understand the trade-off between expressibility of the language used to specify the policy and efficiency of compliance checking algorithm. The research in this workpackage will also develop links with the research conducted in WP1.2 (Coping with selfishness) so as to blend game theoretic techniques and authorization techniques to obtain reputation-based authorization schemes. Also, we will approach the problem of multi-party trust negotiations and models and mechanism for access control.

To date, approaches to trust negotiations have focused on establishing trust between two parties. An important issue that needs to be addressed is related to the process of finding and verifying credentials; such process may indeed entail additional trust negotiations to prove that the party trying to find the credentials is entitled to obtain them. While the overall goal of a trust negotiation systems is the establishment of trust between two specific parties, as a side effect, an arbitrary number of other parties may be dragged into the negotiation, spawning many sub-negotiations, thus resulting in a multi-party negotiation. Several issues need to be addressed. Current trust negotiation languages need to be expanded to allow parties to specify who is requesting a credential and for which purpose. Current deadlock detection techniques from distributed system research need to be extended into approaches to kill off remote computations determined to be unfruitful. We plan to develop techniques to address those issues as well as other specifically related to multi-party negotiations arising from highly mobile users.

We will approach the compliance-checking problem from a computational point of view. We will evaluate a broad set of options between two extremes: the most general language (e.g., Turing machines or any Turing equivalent language), leads to undecidable compliance checking problems; on the other hand simple binary policies specified using ACL can be efficiently implemented. Several cases in between have already being shown to NP-hard and thus effective use in large scale domain is unlikely. Currently, the only non-trivial example of policy language that admits a polynomial-time compliance-checking algorithm has been given in [DisTrusMan].

The reputation-based approach has been recently proved effective in global scenarios (most notably eBay) in which a significant amount of cooperation among the users is needed to fully realize their potential. For example, in a resource-sharing system where users trade spare computer cycles, it is crucial that "free riders'" are not authorized to access the system. We stress that these adversaries are of a different nature than adversaries that try to block the functioning of a system or try to obtain information (we deal with this kind of adversaries in, for example, WP4.3). Instead it can be assumed that users will respond to economic incentives. This calls for a blending of techniques from Game Theory and from Security. We will interact with researchers working in WP1.2 to develop a model for reputation-based authorization and to design schemes that are resilient to attacks.

Also in cooperation with WP4.2, we will investigate privacy-preserving mechanisms for use in trust negotiation. The resource requester should be able to choose the credentials/attributes to submit to satisfy such requirement, if several alternatives are possible. Supporting this approach requires, however, that trust requirements be expressed in terms of high-level semantics properties. To address this issue we will rely on reference ontology, and formalize the notion of trust requirement. We will also address privacy issues allowing a subject to adopt strategies to make the set of credentials/attributes he/she is going to release privacy preserving.

We aim to define an access control model suitable for large-scale scenarios. We envision that such a model will be characterized by rich description of subjects and objects, so that high-level policies can be directly expressed in terms of subject and object properties. Ontologies and federated identity management will be exploited to support multi-domain interoperability. Context-based access decisions will also be supported by the model.

WP 4.2 Privacy, Identity and Anonymity

Leader**:** KUL; Participants: UNISA, CYB

The objective of this workpackage is to develop efficient, scalable and secure solutions for anonymous communications and anonymous transactions for a global computing environment.

An overview will be made of proposals for technical mechanisms for privacy, identity and anonymity; this includes both mechanisms at the network level (such as mixes and crowds) and mechanisms at the application level such as anonymous credentials. The overview will provide a detailed comparison of all the techniques, with an assessment of potential for a global computing environment. It will also be indicated how these techniques can be combined in a modular fashion.

For the network level techniques, quantitative statistical models will be developed and implemented to model and analyze the behavior of the mechanisms and to improve new building blocks. This will require statistical analysis for more complex mix models: so far only simple mixes with very simple batching strategies for a single mix could be analyzed, while it is known that the best mix networks use much more complex batching strategies. Moreover, there are important open problems on how to extend this analysis to complex networks of mixes and to strategies with dummy message insertions. The goal is to achieve a deep insight in the trade-off between quantitative measures of anonymity on the one hand and cost and quality of service (delay parameters and reliability) on the other hand. In this context, it will also be very important to develop more realistic and complex adversary models: so far it has been assumed (for simplicity) that the adversary has full access, while the entities themselves have only local access. This is particularly important for applications with real-time constraints, since in this setting no solutions exist in the presence of an “omni-present” adversary. In addition, mixes that are more flexible and that offer additional features will be developed, that include forward secrecy, congestion control and adaptivity and these schemes will be analyzed from a theoretical viewpoint.

For anonymous credentials, research will be performed into more efficient and scalable solutions, among others by building on solutions and techniques developed in WP4.3 and by integrating these into other protocols. This will focus on a comparison of different cryptographic techniques (e.g., pairing-based solutions versus “classical” discrete logarithm or RSA-based solutions) and on the development of novel credential based mechanisms based on well established and on novel cryptographic problems. Moreover, theoretical limitations (lower bounds) will be investigated.

WP 4.3 Secure distributed computation

Leader: UNISA, Participants: CTI, UDRLS, KUL, DIM, CYB.

Our aim is to develop efficient and proactively secure protocols for applications relevant to global computing such as privacy-preserving data mining, secure algorithmic mechanisms for route discovery in networks and electronic voting.

Thanks to the completeness results for secure distributed protocols, it is known how to compute in a secure way any efficiently computable function; however, the protocols obtained through the application of the completeness result are not practical and are only secure if executed as stand-alone protocols.

The research of this workpackage is articulated along several axes. Along a foundational line of research, we will extend the theory of secure distributed protocols so to be able to formally describe and analyze secure distributed protocols for global computers. An important problem is to refine the model of universally composable so as to make proof of security simpler. Possible lines of research include the design of communication-efficient protocols for computing any (efficiently) computable function that can be concurrently composed or even universally composed (currently, similar general results are only known under strong set-up assumptions or for stand-alone security).

Along a more immediate line of thought, we will design efficient secure protocols for important applications among which we list privacy-preserving data mining (see also WP4.2), secure algorithmic mechanisms for route discovery in networks and time stamping schemes. Within the context of proactively secure protocols we will adopt the graph-theoretic model of the pursuit-evasion distributed game and develop efficient distributed strategy to efficiently locate all the clones of a worm that is attacking a network. This entails both modeling the spreading of the worm and designing efficient distributed algorithms to coordinate the work of the antivirus mobile agents. For the more efficient protocols, prototype implementations will be provided.

At the intersection of the two lines of research lies the problem of defining universally composable security for time-stamping schemes and whether there are efficient and practical constructions for universally composable time-stamping schemes.

Sub-Project 5: Extending global computing to wireless users

Leader: CTI; Participants: UOP, UOI, UDRLS, URDTV, UPB, CAU, CNRS, INRIA, CUI, KUL, UPC, UCY, DIM

The design of resource-efficient algorithms for mobile networks as one extension of an overlay computer is a challenging task. Since such networks frequently change topology, and often without any regular pattern, it is important to consider topology maintenance to enable fast, reliable and scalable communications within this mobile and spontaneous network. In addition, the nodes of an ad hoc network may have heterogeneous capabilities. They differ, e.g., in efficiency, bandwidth, energy, communication media, and mobility. Open questions are, for example, what is realistic mobility and how do we handle fault tolerance in these heterogeneous wireless networks?

The main goal of this subproject is to model all these characteristics and to design, analyze and implement algorithms for resource-efficient network management which make computing and communication on such a mobile ad hoc network (like an extension to an overlay computer) possible.

Needless to say, the abstraction process over the common characteristics of the diverse component devices of a heterogeneous wireless network requires to cope with new unexpected computational problems in that user requirements and sophisticated applications require more and more resources that might be not always available to the (different) wireless devices. Hence, new advanced applications that make innovative and extensive use of communication networks need a careful and efficient management of the network resources. Therefore, advanced algorithmic tools and solutions are to be developed and used, otherwise the huge potential offered by new communication media will not be exploited. Indeed, we have reached a point in the technological evolution of the communication domain where any further progress on the transmission technology (e.g., new communication technology with a tenfold increase in bandwidth) does not have an equivalent impact on the services provided to the information and knowledge-based society, unless it is coupled with the adequate structural and algorithmic tools for the associated design and provisioning issues. The main objective of the subproject is to provide practically efficient algorithmic solutions for high quality, reliable, stable end-users services to heterogeneous wireless mobile networks; this will be necessarily accomplished by abstracting over the common characteristics of the diverse component devices. In turn, the provision of algorithmic solutions for high quality, reliable, stable end-users services to heterogeneous wireless mobile networks requires/means to face with a number of strongly related issues that can be grouped in the following categories, directly corresponding to the workpackages of this subproject:

- WP5.1 Resource management and quality-of-service
- WP5.2 Network design and topology control
- WP5.3 Mobility and fault tolerance

WP 5.1: Resource Management and Quality-of-Service

Leader: UDRTV; Participants: UOP, CTI, CUI, CAU, INRIA

A first goal concerning range assignment is the design and analysis of efficient range assignment strategies that achieve bounded-hop broadcast, accumulation or full connectivity within minimal energy. We also aim to start the study of adaptive and scalable strategies. Furthermore, we shall start considering the issue of improving the adherence of existing energy cost models to real features of wireless networks. One candidate model is certainly the min-max (instead of the classic min-sum), i.e., a model where one minimizes the maximum energy consumed by a host, thus achieving a fair usage of the network. Another candidate model is that in which the power-cost function depends on several factors, both environmental and device-based, besides the distance between the devices. In this respect, we could decide to model the "cruciality" of limiting the power consumption for a device within a network (it is more crucial for a sensor than for a mobile computer). We are interested in practically efficient protocols. Short run-time in the heterogeneous wireless networks setting becomes crucial since protocols must be adaptive and scalable. Clearly adaptiveness and scalability are strongly related to user transparency requirements. Hence, we are not interested in, say, complex linear programming based algorithms rather we shall often focus on simple heuristics which work well on average (average case analysis of greedy algorithms).

Selfish behaviour of the network devices will also be considered. Our main purpose for the first 18 months is to investigate the existence of feasible, approximation truthful mechanisms for range assignment problems, that is, truthful mechanisms that compute in polynomial-time an approximate solution. We still remark that achieving this goal is not possible by simply combining standard approximation algorithms and VCG-based payments [NR99,NR00,R00]. In fact, the concept of truthfulness, traditionally used in economic setting, does not seem suitable for these kind of problems and different approaches will be considered. Another way to cope with selfishness is that of letting users to converge in suitable equilibria [FPS00]. More precisely, under the assumption that any user is willing to change strategy or to be rerouted each time it may be served in a cheaper way, the evolution of the network can be modelled as a multi-player game. Concerning this research direction, one basic problem is that of determining when users share in a fair way the cost of the consumed energy required for their transmissions, whether Nash equilibria exist and what is the resulting coordination ratio [KP99] (i.e., the ratio between the cost of the worst Nash equilibrium and that of an optimal centralized solution). Starting from such results, it would be important to determine reasonable pricing strategies that a service provider can adopt in order to keep the overall performance at acceptable equilibrium. In such a setting, besides reducing the arising coordination ratio, it would be nice also to design polynomial time algorithms able to determine optimal equilibria or equilibria close to the optimal ones. Moreover, the rate of convergence to an equilibrium should be established in the case in which the system evolves freely by allowing at each step a user to modify his strategy for reducing his local cost.

Another important goal is the design of efficient collision avoidance protocols in sensor networks, particularly for the case of multi-path data propagation where, due to multiple nearby broadcasts, collisions tend to occur more often.

Energy efficient routing shall also be considered during the first 18 months. One method to route on a given topology is to collect in each node all possible information about the network structure, e.g., in routing tables. But this approach is now not possible because of the dynamics and the small bandwidth of the network. Furthermore, we can ask whether it is possible to construct routing schemes which allow guarantees on the delivery time of packets. The following lines of research will be followed: (i) Cluster versus flat routing: despite the many different solutions proposed, a thorough investigation of the advantages and disadvantages of adopting a flat vs. hierarchical routing in sensor networks is lacking. Our goal is to compare the behaviors of the most prominent proposals, especially Directed Diffusion and its cluster-based counterparts with respect to some metrics that describe power efficiency and QoS. In particular, power efficiency is suitably described by network lifetime, i.e. the time by which a certain percentage of the nodes become disconnected from the sink. To avoid early energy depletion of certain sensors, we will investigate energy balanced information propagation schemes. In sensor networks applications, QoS is often measured by the average latency to carry data from originating sensors to the sink node. Probabilistic data propagation schemes, towards achieving desired tradeoffs between energy, time efficiency and fault-tolerance will also be designed and analyzed. (ii) Data aggregation and power efficiency: we shall start the study of the potential impact of data aggregation on the power efficiency of the main routing protocols proposed in the literature. (iii) Data aggregation and power efficiency versus latency: our concern for the first 18 months will be the design of efficient on-line algorithms for packet aggregation and forwarding at a single node, where the objective function is a linear combination of power consumption and overall delay at the node.
All the above outlined issues will be considered in an experimental setting as well.

WP5.2: Dynamical Aspects of Network Design and Topology Control

Leader: CUI; Participants: UOP, CNRS, UPB, CTI, UDRTV, UPC, UDRLS, INRIA, DIM, UCY

The main goals of the first 18 months are (i) to analyze and handle topologies as basic network structures for overlay computing; (ii) to understand the impact of the shape of the transmissions on the communications; (iii) to improve the understanding of the collisions. Concerning the shape of the transmissions, in [BLR05] a stochastic method [RM51] was implemented in order to compute numerically some critical parameters of the networks to ensure a particular task to be possible. For example, considering the process of localization in wireless networks, it is numerically shown that increasing the angle of emission increases the chance of success of the process. The most relevant observation is the existence of a critical value of the angle of emission under which the process almost surely fails and above which almost surely succeeds. The process then shows a phase transition when the angle of emission increases.

We plan to extend this numerical investigation of the impact of the shape of transmissions on the localization process. The main goal of such an extension is to get a more definite opinion on the existence of a phase transition and also a more global point of view by considering many values of the others parameters. Getting such numerical results helps in pointing out what we can expect to prove theoretically. We will investigate and propose dynamical and stochastic models of the interference arising in routing in wireless sensor networks. We will also propose implicit topology management schemes including sleep-awake mechanisms in wireless sensor networks. Such schemes aim to appropriately put sensors asleep for certain periods of time, in order to save energy, without however affecting the network connectivity (and thus the efficiency of data propagation) too much.

The study of range assignment for station positioning will be started by further investigating the relationship with facility location problems. We also plan to analyze the different models for the non-homogeneous case in the literature by experimental evaluation.

WP 5.3: Mobility and Fault Tolerance

Leader: CTI; Participants: CNRS, UPB, UOI, CUI, UDRTV, UPC, UDRLS, KUL, INRIA

We plan here to model and characterize mobility and fault tolerance in scalable, heterogeneous wireless networks.

Concerning the relationship between fault-tolerance and range assignment, we intend to pursue the objective of maintaining a certain connectivity predicate in a wireless network affected by a transient component failure, by aiming both to save the energy consumption, and to preserve as much as possible the preexisting network. More precisely, the network functionality will be re-established by satisfying the following constraints, in this order: (a) minimize the number of changes, where a change is any network operation defined a priori (e.g., a range increment of a radio station, etc.); (b) minimize the objective energy function addressed by the original network over the restricted set of feasible solutions defined by (a). We want to emphasize here that most of the conventional optimization problems in wireless networks are NP-hard, and therefore we expect the use of approximation techniques. Special attention will be devoted to the so-called "local" maintenance of the network functionality, which is restricted to adjust only the ranges of the hosts that were actually using the failed component for their communication purposes. We also start the study of how to compare the quality of the constrained solutions with respect to the optimal ones, as computed from scratch in the residual network (i.e., the network deprived of the failed component).

Another fundamental routing problem that will be considered at first is the design and the analysis of distributed broadcast protocols in the presence of faults. Our goal for the first phase of this project is to study the problem of designing fault tolerant routing protocols having completion time as function (also) of the number of faults suffered by the network. Of course, such a study will be started by the analysis of the related literature. Another important aspect is the power of the Fault-Adversary. Previous theoretical analysis considered only the worst-case, while a significant performance evaluation of the protocols would be that in which the adversary is random and thus analyze the expected completion time. During this phase we intend to consider fault tolerance against random adversary issues by studying the average-case and by using Smooth Analysis.

As to the network survivability, in the past, efficient solutions have been developed in the context of several wired network topologies, while on the contrary, for wireless networks, very few work has been done [GPT01, GPT03]. We aim, in this first phase, to study the possibility of extending the wired solutions to the wireless context.

We also plan to investigate the significant impact of the mobility rate and the user density on the performance of routing protocols in ad-hoc mobile networks, especially the important performance drop in the case of very high mobility rates and sparse networks. To overcome such effects, we will propose routing protocols using and taking advantage of the motion in the network by e.g. forcing few hosts to move acting as “helpers" for message delivery, that can be shown to tolerate well high mobility rates and low densities.

We also plan to model interaction and mobility in a combinatorial way by suitable random graphs. In particular, we plan to adopt the G(n,m,p) model of random intersection graphs, that may capture network interactions due to proximity. We also plan to use stochastic processes and, in particular, stochastic interactions of particles and their applications in network problems, particularly in information propagation in mobile environments.

Sub-Project 6: Design and Implementation of Components and Applications for Programmable Overlay Computer

Leader: CTI; Participants: All partners

The work within the subproject during the first 18 months is divided into three workpackages:

- WP6.1 Specification and design of the platform
- WP6.2 Implementation of platform components
- WP6.3 Integration and testing of the platform

WP 6.1 Specification and design of the platform

Leader UOP, Further Participants: All

The main objective of WP6.1 is to define the architecture and the functionalities of the Programmable Overlay Computing Platform.

In this workpackage, we will select the platform to be used as the heart of the overlay computer (typical candidate platforms include JXTA, DHT-based platforms, GnuNet, and JADE), we will define the functionalities that will enhance it and select the corresponding algorithms for the implementation of these functionalities. The platform should provide basic primitives for distributed computing and communication. Functionalities for resource discovery, bandwidth sharing, data management, load management, scheduling, trust management, secure computation, etc. will enhance and extend the basic primitives of the platform. The algorithms that will be used for these implementations will be selected among the ones developed and validated within subprojects SP1, SP2, SP3, and SP4 while some of them will also take the specific characteristics of wireless devices (SP5) into account. It is expected that, for some functionalities, more than one algorithm will be selected for implementation.

WP6.2 Implementation of platform components

Leader: UOP, Further Participants: TILS, CNRS, UPB, CTI, UNISA, UOI, MPII, CAU, UDRTV, UNIPD, ETHZ, UDRLS, KUL, INRIA

The objective of the workpackage is to provide implementations of functionalities selected in WP6.1 on top of the selected platform (these implementation will provide the basis for the components of the platform) and prototypes designed in workpackages WP3.3 and WP3.4.

The partners involved will implement a sufficient set of functionalities selected in the context of WP6.1. At this phase, functionalities will run on the platform as stand-alone software. Some preliminary testing (on this stand-alone basis) will be performed here. So, one of the outcomes of WP6.2 will be a set of implementations on top of the platform.

In addition, we will work on the implementation of components of the following propotypes:

- Microbenchmarking software package (based on developments of WP3.2)
- Prototype for process migration and load balancing in BSP-based P2P system (based on developments of WP3.2)
- Generic engine for specifying computations in overlay computers (based on developments of WP3.4)
- Prototype for anonymous communication (based on developments of WP4.2)
- Prototype for secure distributed protocols (based on developments of WP4.3)

WP 6.3 Integration and testing of the platform

Lead Partner: TILS further participants: UOP, CNRS, UPB, CTI, UNISA, UOI, MPII, CAU, UDRTV, UNIPD, ETHZ, UDRLS, KUL, INRIA

The main objectives of this workpackage are

- to integrate some of the implementation produced in the context of WP6.2 into a coherent software package; this will be a very first release of the Programmable Overlay Computing Platform.
- to define a preliminary Testbed Architecture where the platform and applications will be validated

This workpackage will integrate the set of validated implementations delivered in the context of WP6.2 to create an enhanced platform. In this task, we will select the components which will be part of the platform, and fine-tune the interface of the functionalities (i.e., fine-tune their parameters, options, etc.) from the programmer's point of view. The outcome of this workpackage will be an integrated software package, which will be an enhancement of the basic computing/communication primitives of the original software platform. Test suites will be developed for technical testing of the critical platform functionalities (these suites may be extended and reused in later phases of the project).

This workpackage will also investigate requirements and architectural guidelines for distributed testbeds. In the first 18 months, along with setting up the necessary resources for testing the platform, we will also study the state-of-the-art and perform further research on mechanisms that are desirable for testbeds oriented to overlay computers. A main requirement here is that tests should be independently and without interference performed by different users in a highly distributed setting. As a result of this task, we will produce architectural guidelines for distributed testbed design.

Financial Information (Integrated Project, first 18 months)											
		RTD and innovation-related activities		*Demonstration activities*		*Training activities*		*Consortium management*		*Total*	
Participant n°	*Cost model*	*Costs*	*Requested grant to the budget*	*Costs*	*Requested grant to the budget*	*Costs*	*Requested grant to the budge*	*Costs*	*Requested grant to the budget*	*Costs*	*Requested grant to the budget*
1	AC	135,000	135,000			18,000	18,000	54,000	54,000	207,000	207,000
2	FC	270,000	135,000					7,500	7,500	277,500	142,500
3	FC	270,000	135,000					7,500	7,500	277,500	142,500
4	AC	135,000	135,000					7,500	7,500	142,500	142,500
5	AC	135,000	135,000					7,500	7,500	142,500	142,500
6	AC	135,000	135,000					7,500	7,500	142,500	142,500
7	AC	112,500	112,500							112,500	112,500
8	AC	127,500	127,500							127,500	127,500
9	AC	90,000	90,000							90,000	90,000
10	AC	105,000	105,000							105,000	105,000
11	AC	90,000	90,000							90,000	90,000
12	AC	75,000	75,000					7,500	7,500	82,500	82,500
13	AC	60,000	60,000							60,000	60,000
14	AC	90,000	90,000							90,000	90,000
15	FC	105,000	52,500							105,000	52,500
16	FC	150,000	75,000							150,000	75,000
17	AC	135,000	135,000							135,000	67,500
18	FC	135,000	67,500							135,000	135,000
19	AC	60,000	60,000							60,000	60,000
20	AC	60,000	60,000							60,000	60,000
21	FCF	75,000	37,500							75,000	37,500
(Sub)-total		2,550,000	2,047,500			18,000	18,000	99,000	99,000	2,667,000	2,164,500

b) Work planning, showing the timing of the different WPs and their tasks

	Month																	
Workpackages	**1**	**2**	**3**	**4**	**5**	**6**		**8**	**9**	**10**	**11**	**2**	**13**	**14**	**15**	**16**	**17**	**18**
SP 1: Paradigms and principles																		
WP 1.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 1.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 1.2	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 1.3	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 1.4	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 1.5	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
SP2: Resource management																		
WP 2.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 2.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 2.2	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 2.3	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
SP 3: Sharing information and computation																		
WP 3.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 3.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 3.2	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 3.3	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 3.4	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
SP 4: Security and trust management																		
WP 4.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 4.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 4.2	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 4.3	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
SP 5: Extending global computing to wireless users																		
WP 5.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 5.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 5.3	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
SP 6: Design and Implementation of Components and Applications for Programmable Overlay Computers																		
WP 6.0	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 6.1	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨						
WP 6.2							▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨	▨
WP 6.3													▨	▨	▨	▨	▨	▨
WP 6.4																		

c) Graphical presentation of the components, showing their interdependencies

A graphical representation of the components of the project is depicted in the following figure. Subprojects SP1, SP2, SP3, SP4, and SP5 supply the "proof-of-concept" subproject SP6 with functionalities and corresponding algorithms to be included in the Programmable Overlay Computing Platform. In addition, workpackages WP3.2, WP3.4, WP4.2, and WP4.3 (subprojects 3 and 4) will solve the main design issues arising in the prototypes that will be implemented in workpackage WP6.2 (of subproject SP6). The four thematic subprojects SP2, SP3, SP4, and SP5 interact with the subproject SP1 "Paradigms and Principles"; generic developments within SP1 (like new techniques for coping with selfishness or partial knowledge) are expected to be applicable in the specific problems considered in the thematic subprojects. Furthermore, new issues arising in the context of the thematic subprojects may lead to the investigation of paradigms and principles which may applicable to several different contexts.

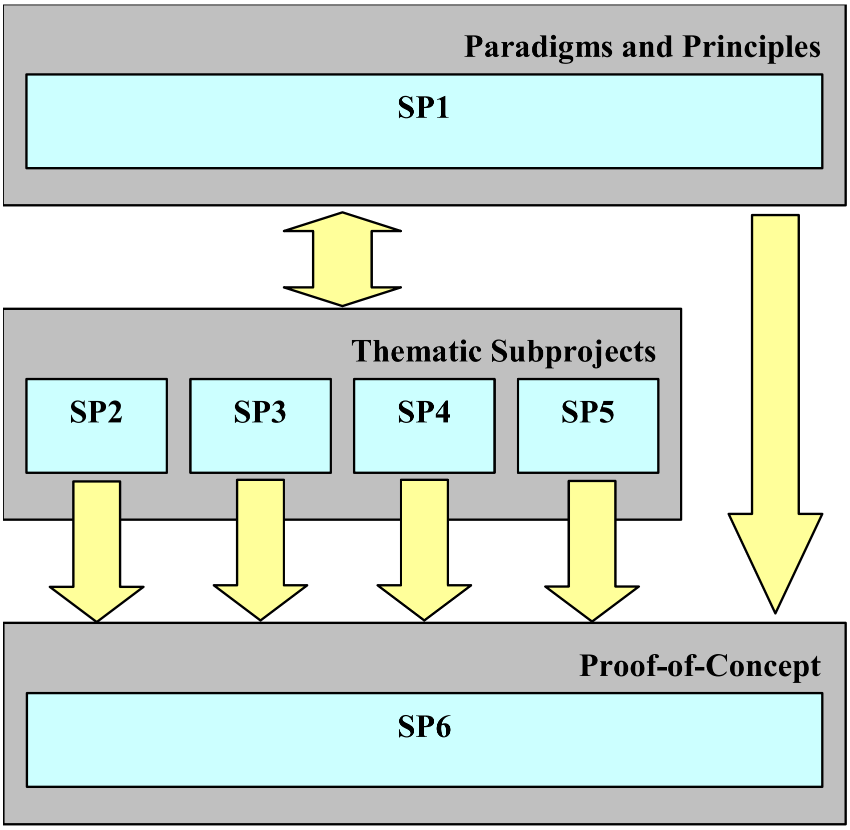

d) Detailed work description broken down into workpackages:

Workpackage list (18 month plan)

Workpackage No	Workpackage title	Lead contractor No	Person-months	Start month	End month	Deliv-erable No
SP 1: Paradigms and principles						
WP 1.0	**Subproject management and dissemination activities**	**UOA**	1.5	1	18	**D1.0.1**
WP 1.1	**Structural Properties**	**UCY**	22	1	18	**D1.1.1** **D1.1.2**
WP 1.2	**Computing with Incomplete Knowledge**	**UOA**	22.5	1	18	**D1.2.1** **D1.2.2**
WP 1.3	**Coping with selfishness**	**UOA**	26	1	18	**D1.3.1** **D1.3.2**
WP 1.4	**Stability and Fault-Tolerance**	**UPC**	22	1	18	**D1.4.1** **D1.4.2**
WP 1.5	**Generic Algorithms**	**CNRS**	23.5	1	18	**D1.5.1** **D1.5.2**
SP 2: Resource management						
WP 2.0	**Subproject management and dissemination activities**	**CNRS**	1.5	1	18	**D2.0.1**
WP 2.1	**Resource Discovery**	**UOI**	19	1	18	**D2.1.1** **D2.1.2**
WP 2.2	**Critical Resource Sharing**	**CNRS**	30	1	18	**D2.2.1** **D2.2.2**
WP 2.3	**Mechanism Design**	**UNISA**	22	1	18	**D2.3.1** **D2.3.2**
SP 3: Sharing information and computation						
WP 3.0	**Subproject management and dissemination activities**	**UPB**	1.5	1	18	**D3.0.1**
WP 3.1	**Distributed Data Management**	**MPII**	23.5	1	18	**D3.1.1** **D3.1.2**
WP 3.2	**Load Management**	**UPB**	23	1	18	**D3.2.1** **D3.2.2** **D3.2.3** **D3.2.4**

WP 3.3	**Scheduling**	**CAU**	31	1	18	**D3.3.1** **D3.3.2**
WP 3.4	**Workflow and Services**	**ETHZ**	15	1	18	**D3.4.1** **D3.4.2**
SP 4: Security and trust management						
WP 4.0	**Subproject management and dissemination activities**	**UNISA**	1.5	1	18	**D4.0.1**
WP 4.1	**Trust management**	**UNISA**	24	1	18	**D4.1.1** **D4.1.2**
WP 4.2	**Privacy, identity and anonymity**	**UNISA**	18	1	18	**D4.2.1** **D4.2.3** **D4.2.2**
WP 4.3	**Secure distributed computation**	**KUL**	29	1	18	**D4.3.1** **D4.3.2** **D4.3.3**
SP 5: Extending global computing to wireless users						
WP 5.0	**Subproject management and dissemination activities**	**CTI**	1.5	1	18	**D5.0.1**
WP 5.1	**Resource management and quality-of-service**	**UDRTV**	22	1	18	**D5.1.1** **D5.1.2**
WP 5.2	**Dynamical aspects of network design and topology**	**CUI**	29.5	1	18	**D5.2.1** **D5.2.2**
WP 5.3	**Mobility and fault tolerance**	**CTI**	26.5	1	18	**D5.3.1** **D5.3.2**
SP 6: Design and implementation of components and applications for programmable overlay computers						
WP 6.0	**Subproject management and dissemination activities**	**TILS**	1.5	1	18	**D6.0.1**
WP 6.1	**Specification and design of the platform**	**UOP**	28.5	1	12	**D6.1.1**
WP 6.2	**Implementation of platform components**	**UOP**	62	7	15	**D6.2.1** **D6.2.2** **D6.2.3** **D6.2.4** **D6.2.5** **D6.2.6**
WP 6.3	**Integration and testing of the platform**	**TILS**	30	10	18	**D6.3.1** **D6.3.2**
WP 6.4	**Design and implementation of a demo application**	**TILS**	0	19	TBD	**TBD**

	TOTAL		549			

(Deliverables list, use Deliverables list form below)

Deliverables list (18 month plan)

Deliverable No	Deliverable title	Delivery date	Nature	Dissemination level
SP 1: Paradigms and principles				
D1.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D1.1.1	Structural properties of overlay computers: State-of-the-art survey	9	R	PU
D1.1.2	Structural properties of overlay computers: Algorithmic solutions and recommendations	18	R	PU
D1.2.1	Computing with incomplete knowledge: State-of-the-art Survey	9	R	PU
D1.2.2	Computing with incomplete knowledge: Algorithmic solutions and recommendations	18	R	PU
D1.3.1	Coping with selfishness: State-of-the-art Survey	9	R	PU
D1.3.2	Coping with selfishness: Algorithmic Solutions and Recommendations	18	R	PU
D1.4.1	Stability and fault-Tolerance: State-of-the-art survey	9	R	PU
D1.4.2	Stability and fault-Tolerance: Algorithmic solutions and recommendations	18	R	PU
D1.5.1	Generic algorithms: State-of-the-art survey	9	R	PU
D1.5.2	Generic algorithms: Algorithmic solutions and recommendations	18	R	PU
SP 2: Resource Management				
D2.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D2.1.1	Resource discovery: State-of-the-art survey	9	R	PU
D2.1.2	Resource discovery: Algorithmic solutions and recommendations	18	R	PU
D2.2.1	Critical resource sharing: State-of-the-art survey	9	R	PU

D2.2.2	Critical Resource Sharing: Algorithmic solutions and recommendations	18	R	PU
D2.3.1	Mechanism design: State-of-the-art survey	9	R	PU
D2.3.2	Mechanism design: Algorithmic solutions and recommendations	18	R	PU
SP 3: Sharing Information and Computation				
D3.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D3.1.1	Distributed data management: State-of-the-art survey	9	R	PU
D3.1.2	Distributed data management: Algorithmic solutions and recommendations	18	R	PU
D3.2.1	Load management: State-of-the-art survey	9	R	PU
D3.2.2	Microbenchmarking software package: Design report	12	R	PU
D3.2.3	Prototype for process migration and load balancing in BSP-based P2P systems: Design report	12	R	PU
D3.2.4	Load management: Algorithmic solutions and recommendations	18	R	PU
D3.3.1	Scheduling: State-of-the-art survey	9	R	PU
D3.3.2	Scheduling: Algorithmic solutions and recommendations	18	R	PU
D3.4.1	Workflow and services: State-of-the-art survey	9	R	PU
D3.4.2	A generic engine for specifying computations in overlay computers: Specification and design report	12	R	PU
SP4: Security and trust management				
D4.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D4.1.1	Trust management: State-of-the-art survey	9	R	PU
D4.1.2	Trust management: Algorithmic solutions and recommendations	18	R	PU
D4.2.1	Privacy, identity and anonymity: State-of-the-art survey	9	R	PU
D4.2.2	Prototype for anonymous communication: Specification and design report	12	R	PU

D4.2.3	Privacy, identity and anonymity: Algorithmic Solutions and Recommendations	18	R	PU
D4.3.1	Secure distributed computation: State-of-the-art Survey	9	R	PU
D4.3.2	Prototype for secure distributed protocols: Specification and design report	12	R	PU
D4.3.3	Secure distributed computation: Algorithmic Solutions and Recommendations	18	R	PU
SP 5: Extending Global Computing to Wireless Users				
D5.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D5.1.1	Resource management and quality of service in wireless networks: State-of-the-art survey	9	R	PU
D5.1.2	Resource management and quality of service in wireless networks: Algorithmic solutions and recommendations	18	R	PU
D5.2.1	Dynamical aspects of network design and topology control: State-of-the-art survey	9	R	PU
D5.2.2	Dynamical aspects of network design and topology control: Algorithmic solutions and recommendations	18	R	PU
D5.3.1	Mobility and fault-tolerance: State-of-the-art survey	9	R	PU
D5.3.2	Mobility and fault-tolerance: Algorithmic solutions and recommendations	18	R	PU
SP 6: Design and Implementation of Components and Applications for Programmable Overlay Computers				
D6.0.1	Subproject report on the activities of the first 18 months	18	R	PU
D6.1.1	Programmable Overlay Computing Platform: Design report	12	R	PU
D6.2.1	Programmable Overlay Computer – Software components	18	P	PU
D6.2.2	Microbenchmarking software package	18	P	PU
D6.2.3	Prototype for process migration and load balancing in BSP-based P2P system	18	P	PU
D6.2.4	Generic engine for specifying computations in overlay computers	18	P	PU
D6.2.5	Prototype for anonymous communication	18	P	PU

D6.2.6	Prototype for secure distributed protocols	18	P	PU
D6.3.1	Programmable Overlay Computing Platform: Preliminary version	18	P	PU
D6.3.2	Architectural guidelines for testbed design	18	R	PU

(Description of each workpackage, use Workpackage description form below, one per workpackage)

Workpackage description (18 month plan)

Workpackages of SP 1: Paradigms and principles

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 1.0**	**Start date or starting event:**			Month 1		
Participant id	UOA						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: the workpackage will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The workpackage will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.
- *Dissemination of results*

Deliverables

D1.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D1.0.1

Workpackages of SP 1: Paradigms and principles

WP Title: Structural properties							
Workpackage number	**WP 1.1**	**Start date or starting event:**	Month 1				
Participant id	UOP	CNRS	CTI	CUI	UDRTV	UOA	UNIPD
Person-months per participant:	3	3	0,5	3	1	2	4
Participant id	DIM	**UCY**					
Person-months per participant:	2.5	3					

Objectives

To study topological properties of global and overlay computers, their substructures and other organization issues, and embeddings that relate to implementation and realization of overlay computers into global ones.

Description of work

We plan to build upon the existing work on the topological properties of the Internet and the Web and to design searching and communication primitives that take advantage of their structural properties. We also plan to study the ``geometry" of global systems and to seek improved algorithms for clustering and

partitioning problems associated to the efficient implementation of overlay computers. Despite the rich history of metric embeddability, there are major open algorithmic problems that relate to low distortion embeddings and we plan to attack them. Low distortion, which guarantees good approximation for every objective function is desirable, but in many cases impossible. In this case, we will pursue special solutions that may have unbounded distortion but still can preserve the given objective functions.

Deliverables

D1.1.1: Structural properties of overlay computers: State-of-the-art survey
D1.1.2: Structural properties of overlay computers: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D1.1.1
Month 18: D1.1.2

Workpackages of SP 1: Paradigms and principles

WP Title: Coping with incomplete knowledge							
Workpackage number	**WP 1.2**	**Start date or starting event:**			Month 1		
Participant id	UOP	CTI	**UOA**	UNIPD	UPC	INRIA	DIM
Person-months per participant:	2	0.5	5	4	2	4	4
Participant id	UCY						
Person-months per participant:	1						

Objectives

The goal of this workpackage is to use competitive analysis and the theory of distributed systems to study the novel online problems that are at the core of global and overlay computers, such as distributed and robust data structures, transparency and information hiding, and reputation mechanisms.

Description of work

We plan to study simple primitives that allow the maintenance of data structures in selfish, distributed, and dynamic environments. The problems associated to efficient data structure implementation with provably good competitive ratio are among the hardest ones even for centralized data structures. We plan to do research for simple centralized data structures for searching and servicing requests and to extend the results for dynamic distributed environments. We also plan to investigate the role of information hiding in achieving good online solutions. A related problem of central importance in global computing is the locality issue (where the entities have only partial/local information about the system). We plan to study from the competitive analysis point of view the impact of locality and to design online algorithms that alleviate it. Finally we plan to study reputation mechanisms: what to distill from the history of past transactions, where and how to maintain it, and how to predict the future behavior of users. We will model this problem both as traditional online problem (with a powerful adversary to obtain worst-case bounds) and as the more realistic problem where the adversary controls only the order of the transactions and the behavior of users is determined solely by their selfish nature.

Deliverables:

D1.2.1: Coping with incomplete knowledge: State-of-the-art survey
D1.2.2: Coping with incomplete knowledge: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D1.2.1
Month 18: D1.2.2

Workpackages of SP 1: Paradigms and principles

WP Title: Coping with selfishness							
Workpackage number	**WP 1.3**	**Start date or starting event:**			Month 1		
Participant id	UOP	UPB	CTI	UNISA	CAU	UDRTV	**UOA**
Person-months per participant:	2	3	1	4	2	2	5
Participant id	UPC	UDRLS	UCY				
Person-months per participant:	2	3	2				

Objectives

The goal of this workpackage is to study ways to improve system performance in global systems with selfish entities.

Description of work

We plan to extend existing results on the price of anarchy on systems with many selfish users that use a set of resources. We also plan to study the impact of information regimes on the performance of selfish systems and to come up with algorithmic solutions that improve the coordination of selfish users and the performance of the whole systems based on appropriate manipulation of the information flow. We also plan to investigate the fundamental problem of “fairness vs. efficiency” in global and overlay computers and to study how incomplete information can help achieve solutions that are both fair and efficient even when this is not possible in the full information regime. Finally we plan to study ways to improve system performance through coordination mechanisms and mechanism design techniques and to extend these results to dynamic environments.

Deliverables

D1.3.1: Coping with selfishness: State-of-the-art survey
D1.3.2: Coping with selfishness: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D1.3.1

Month 18: D1.3.2

Workpackages of SP 1: Paradigms and principles

WP Title: Stability and fault tolerance							
Workpackage number	**WP 1.4**		**Start date or starting event:**		Month 1		
Participant id	UOP	CNRS	CTI	CUI	**UPC**	DIM	UCY
Person-months per participant:	2	4	1	3	6	4	2

Objectives

The goal of this workpackage is to study ways to improve stability and fault-tolerance in global systems based on adversarial queuing theory and the theory of distributed systems.

Description of work

We plan to use the framework of adversarial queuing theory to study stability issues of global systems. Adversarial queuing theory studies static networks of queues when the injection of packets is controlled by an adversary. The existing results do not take into account the ever-changing environment of global systems. We plan to extend the results to networks where the set of nodes and connections change over time. More specifically, we plan to attack the questions of whether certain topologies are stable and whether specific protocols are stable. Also, we plan to study the question on how much the extra communication load, which is required to support functionalities at the overlay computer level, affects the stability of the underlying global computer. We then plan to study the issue of improving the stability of global systems by appropriate mechanisms build on overlay computers. Finally, we plan to study distributed algorithmic primitives that allow for detection and correction of faults. We hope to develop a theory to address trade-off between the impact of a failure (how fast the fault is recovered and how many connections must be interrupted during the recovery procedure) and the quantity of resources allocated to protect from faults.

Deliverables

D1.4.1: Stability and fault tolerance: State-of-the-art survey
D1.4.2: Stability and fault tolerance: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D1.4.1

Month 18: D1.4.2

Workpackages of SP 1: Paradigms and principles

WP Title: Generic algorithms							
Workpackage number	**WP 1.5**		**Start date or starting event:**		Month 1		
Participant id	UOP	**CNRS**	CTI	UDRTV	UOA	UPC	UDRLS
Person-months per participant:	2	7	0.5	1	2	2	3
Participant id	INRIA	UCY					
Person-months per participant:	4	2					

Objectives

The goal of this workpackage is to study classical problems in combinatorial optimization that can have an impact on the understanding, modeling, and design of global and overlay computers. It will also include new algorithmic issues that may arise from new developments during the duration of the project.

Description of work

A central issue in global systems is connectivity, a problem that has been studied extensively in classical algorithmic theory. However, the dynamic environment of global systems gives rise to new challenging problems. First of all, we plan to find appropriate mathematical models to capture the issues involved. We will then try to extend classical results to these models. It is clear that the algorithms for these type of problems must be distributed. Considering that the nodes of these systems have only a partial, and sometimes erroneous, view of the system, we have to aim for approximate solutions but with certain guarantees (such as t-spanners). We plan to develop approximation algorithms that address this issue. We also plan to study routing and scheduling problems using multi-criteria objectives. These problems become qualitatively different when we are interested in optimizing not a single objective (for example, delay) but multiple objectives (for example, delay and reliability). Especially in systems with multiple users where fairness is an issue, we have to devise algorithms that try to optimize the objectives of all of them. When this is not possible, we plan to study approximate solutions. Finally, we plan to develop a theory and investigate algorithms that have to base their decisions on samples of the data. The problem is how to select the appropriate sample efficiently (which, given the distributed nature of the system, is an novel question in the interface of Algorithms and Statistics) and how to analyze it.

Deliverables

D1.5.1: Generic algorithms: State-of-the-art survey
D1.5.2: Generic algorithms: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D1.5.1

Month 18: D1.5.2

Workpackages of SP 2: Resource management

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 2.0**	**Start date or starting event:**	Month 1				
Participant id	**CNRS**						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this Workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: The project management WP will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The project management WP will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.
- *Dissemination of results*

Deliverables

D2.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D2.0.1

Workpackages of SP 2: Resource management

WP Title: Resource discovery							
Workpackage number	WP 2.1		Start date or starting event:		Month 1		
Participant id	UOP	CTI	**UOI**	MPII			
Person-months per participant:	4	1	12	2			

Objectives

In this workpackage, we shall explore advanced methods for resource discovery that take into account the structure and content of resources. A central objective is supporting resource discovery based on more advanced queries than single attribute-value or keyword search. The focus will be on scalable, distributed and dynamic solutions that take advantage of the embedding of virtual networks into real ones. To this end, we shall build upon the paradigms and principles of global overlay computers as studied in WP1.The 18-months objective is to achieve sufficient progress towards the overall objective which will be measured by a thorough survey on the topic and initial research results published in the appropriate forums.

Description of work

Work in this workpackage include:

(i) The construction of overlay networks that connect semantically related nodes most efficiently but at the same time provide highly efficient routing to all nodes with logarithmic worst-case behavior in terms of both space overhead per node and routing hops. Distributed clustering and novel embeddings based on content will be studied.

(ii) Processing of advanced queries on top of the constructed networks. Work includes both query routing (that is determining nodes with matching resources)and executing the query at the appropriate nodes.

(iii) Studying the selfishness, scalability, stability and fault tolerance of the above techniques. In the first 18 months, includes work on task (i) and (ii) above.

Specifically:

(a) state-of-the-art surveys on the construction and query processing in overlay networks including a comparative evaluation of existing solutions and their limitations.

(b) Construction of semantic overlays and initial query processing techniques for a subset of the advanced queries under consideration.

Deliverables

D2.1.1: Resource discovery: State-of-the-art survey
D2.1.2: Resource discovery: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D2.1.1
Month 18: D2.1.2

Workpackages of SP 2: Resource management

WP Title: Critical resource sharing							
Workpackage number	**WP 2.2**	**Start date or starting event:**	Month 1				
Participant id	UOP	**CNRS**	UPB	CTI	CUI	MPII	CAU
Person-months per participant:	5	10	2	2	4	1	3
Participant id	UCY						
Person-months per participant:	3						

Objectives

The goal of this workpackage is to provide protocols that allow the user of a global computer to communicate in a transparent way and according to some quality of service and cost. Those protocols shall use the underlying global computer resources efficiently and in a scalable way.

Description of work

- We will investigate the current and forthcoming issues for critical resource sharing in complex hierarchical networks. Even if bandwidth is likely to remain critical (but attached to some quality of service parameters), we will also consider other potential critical resources (routers,…).
- We will then develop techniques to ensure an efficient use of bandwidth in heterogeneous hierarchical networks. Our goal will be to ensure a good mapping of bandwidth requests expressed at the overlay computer level onto the underlying global computer. For this we shall extend classical flow-like algorithms (routing) to the case of logical networks in which connections are labelled by various parameters (delay, cost, reliability).
- We will design tools and models to allow the global computer agents to share or trade bandwidth under quality of service or financial cost constraints. For this we will use concepts from WP1.2 and WP2.3 to cope with social and economic aspects.
- We will extend the algorithms in order to make them able to deal with high dynamic (scalability, fault tolerance). For this we will trade optimality for locality and efficiency, and we will survey and extend on-line algorithms techniques.

Deliverables

D2.2.1: Critical resource sharing: State-of-the-art survey.
D2.2.2: Critical resource sharing: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D2.2.1
Month 18: D2.2.2

Workpackages of SP 2: Resource management

WP Title: Mechanism design							
Workpackage number	**WP 2.3**		**Start date or starting event:**		Month 1		
Participant id	UOP	UPB	CTI	**UNISA**	UDRTV	UOA	KUL
Person-months per participant:	0.5	2	2	5	3	5	0.5
Participant id	UCY						
Person-months per participant:	4						

Objectives

The goal of this workpackage is to study mechanisms for resource management in overlay and global computers in the presence of selfishly acting entities.

Description of work

Truthful mechanisms:

- We will consider several aspects like *approximation guarantee*, *online vs. offline* mechanisms and *frugality* (i.e., the amount of currency or meta-resource the mechanism needs to provide to the agents in order to guarantee optimal solutions).
- We will try do understand which is the amount of information needed to solve a given problem.

Cost-sharing problems:

- We will compare the solutions for wired and wireless networks and study the performance degradation of mechanisms for wired networks when applied to wireless ones.
- We will also pursue further the possibility of obtaining good cost-sharing mechanisms via *non-crossmonotonic* allocation schemes.

Deliverables

D2.3.1: Mechanism design: State-of-the-art survey

D2.3.2: Mechanism design: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D2.3.1

Month 18: D2.3.2

Workpackages of SP 3: Sharing information and computation

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 3.0**	**Start date or starting event:**	Month 1				
Participant id	UPB						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this Workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: The project management WP will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The project management WP will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.

Deliverables

D3.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D3.0.1

Workpackages of SP 3: Sharing information and computation

WP Title: Distributed data management							
Workpackage number	**WP 3.1**	**Start date or starting event:**	Month 1				
Participant id	UPB	CTI	UOI	**MPII**	CAU	UOA	ETHZ
Person-months per participant:	3	0.5	7	6	1	2	4

Objectives

Replication and caching are well-studied problems in distributed computing. They offer distribution transparency, scalability, and fault-tolerance. In global computing, new issues arise due to the dynamic nature of the environment. Furthermore, there is a close interplay between replication and the structural properties of the overlay network. The objectives of this workpackage for the first 18 months are to develop, theoretically analyze and experimentally validate appropriate strategies for caching, replicating, and proactive dissemination of data to the overlay computers to improve performance, availability, fault-tolerance, and data freshness.

Description of work

We plan to investigate at what extent the strategies for caching, replication, and proactive dissemination depend on and may be tailored to the kinds of data items and their access characteristics taking into account the large scale and high dynamics of overlay computers. We will consider different kinds of data items: from simple unstructured data containers like files and Web documents to relational or XML databases, and also for metadata and ontological meta-metadata.

We will study caching and replication strategies in which the determination of the number of replicas for data objects depends on the popularity of the object. We will attempt to algorithms combining the characteristics of existing approaches like uniform replication, proportional replication, and square-root replication. Our work on proactive dissemination will be based on epidemic algorithms and investment-based policies.

The work within the workpackage will specifically focus on the dynamic determination and adaptation of replication degrees and adaptive placement of replicas. This includes the definition of mathematical models and methods for choosing replication degrees and replica placement and the design of algorithms for dynamic load-adaptive adjustments. Furthermore, we will also develop new strategies for replica maintainance and proactive dissemination. Our work will include an in depth survey of the state-of-the-art while in addition to the design and specification of new algorithms, we will also perform their theoretical analysis, implementation, and experimental validation.

Deliverables

D3.1.1: Distributed Data Management: State-of-the-art Survey
D3.1.2: Distributed Data Management: Algorithmic Solutions and Recommendations

Milestones and expected result

Month 09: D3.1.1
Month 18: D3.1.2

Workpackages of SP 3: Sharing information and computation

WP Title: Load management							
Workpackage number	**WP 3.2**		**Start date or starting event:**		Month 1		
Participant id	UOP	**UPB**	CTI	MPII	CAU	UNIPD	ETHZ
Person-months per participant:	1	4	1	1	3	6	2
Participant id	UPC	UDRLS					
Person-months per participant:	2	3					

Objectives:

To study the impact of communication system on the balancing flow quality as well as the convergence of load balancing algorithms. To develop a parametric model of computation where the model parameters reflect phenomena which affect performance. To study the possibility of using peer-to-peer networks for parallel computing.

Description of work

We want to analyze the structural and spectral properties of large self-organizing global computers based for example on the power-law distribution. By examining the coherences between these properties and the quality of for example diffusive-based load balancing algorithms we will develop algorithms along with prototype implementations for basic load management tasks.

We expect to lay the ground toward a general framework for adaptive software by experimenting the combined approach of analytical and empirical search for tuning on a number of relevant case studies. In particular we will try to replicate the results obtained by the ad-hoc approaches at the base of a number of important software packages (FFTW and Linear Solvers for Finite Element Methods such as MUMPS, among the others). We expect our preliminary strategies to exhibit certain dependence with the application at hand and will strive to improve the generality of the proposed approach in the later stages of the project.

Moreover, we will implement the BSP paradigm such that parallel programs can be executed efficiently on a large-scale distributed peer-to-peer network.

Deliverables

D3.2.1: Load management: State-of-the-art survey
D3.2.2: Microbenchmarking software package: Design report
D3.2.3: Prototype for process migration and load balancing in BSP-based P2P systems: Design report
D3.2.4: Load management: Algorithmic solutions and recommendations

Milestones and expected result

Month 9: D3.2.1
Month 12: D3.2.2
Month 12: D3.2.3
Month 18: D3.2.4

Workpackages of SP 3: Sharing information and computation

WP Title: Scheduling							
Workpackage number	**WP 3.3**	**Start date or starting event:**		Month 1			
Participant id	CNRS	UPB	CTI	UNISA	MPII	**CAU**	UOA
Person-months per participant:	4	3	1	3	1	6	2
Participant id	ETHZ	UDRLS	DIM	UCY			
Person-months per participant:	2	3	3	3			

Objectives

The Internet, a collection of interconnected networks with different scale, transmission and hardware, provides a global communication medium for a huge number of computing entities. Most networks in the Internet use certain common protocols, between the so-called application, transport, internet, and host-to-network layers, and provide certain common services. From one side, the existing layered protocol structure can be adopted to provide some new global services. From another side, in the design of protocols it is quite a challenging task to provide some particular performance guarantees on Quality of Service (QoS) between the layers. Indeed, using a correct allocation of layers' resources over time, or simply scheduling and routing, can significantly improve the performance of protocols.

Description of work

A basic problem here is to allocate few resources, which are available to a layer, over time such that some specific guarantee on QoS, e.g. bandwidth requirement, maximum throughput, minimum delay, and etc., is satisfied. The main goal of this WP is to provide efficient algorithmic solutions for that kind of problems arising in the design of protocols for different layers and respective QoSs. We plan to develop suitable scheduling and routing models, find algorithmic solutions, and provide software tools that can help to improve the performance of protocols. The results of this activity will add to the current state of theoretical knowledge on these issues as well as provide a vial support to the technological design activities in the complement subprojects.

Deliverables

D3.3.1: Scheduling: State-of-the-art survey
D3.3.2: Scheduling: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D3.3.1
Month 18: D3.3.2

Workpackages of SP 3: Sharing information and computation

WP Title: Workflow and services							
Workpackage number	**WP 3.4**		**Start date or starting event:**		Month 1		
Participant id	UOI	MPII	**ETHZ**				
Person-months per participant:	6	3	6				

Objectives

The main objective of this workpackage is to investigate issues related to the efficiency, trancparency, and scalability in the execution of large scale computations treating them as processes similar to these found in workflow management systems. The notion of processes allows to model sequences of invocations of computer programs or applications in a distributed and heterogeneous environment as well as to capture the corresponding data exchanges between these programs.

Description of work

In this workpackage we shall explore (i) the use of services for accessing data in overlay computers, (ii) the use of workflow for the definition and execution of composite web services, and (iii) the capabilities of a global computer and its overlay structures for dynamic self-configuration and automatic adaptation to system, workload, and application dynamics. We will develop new methods and improve existing techniques for automatic migration of failed tasks, system awareness (to node configurations, network bandwidth, load, new clusters becoming available or unavailable), replication of system critical information, general fault-tolerance, and the dynamic and automatic reconfiguration of global workflows to such rapidly evolving conditions.

The work will coordinate with the developments made in other workpackages such as WP3.1 (Distributed Data Management), WP3.2 (Load Management), and WP3.3 (Scheduling). The techniques developed will be used in the specification of an autonomic computing system, i.e., a generic engine for specifying computations of overlay networks of Web services and additional services supporting autonomic execution, flexible enough to incorporate new services as they appear.

Deliverables

D3.4.1: Workflows and Services: State-of-the-art Survey

D3.4.2: A generic engine for specifying computations in overlay computers: Specification and Design report

Milestones and expected result

Month 09: D3.4.1

Month 12: D3.4.2

Workpackages of SP 4: Security and trust management

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 4.0**	**Start date or starting event:**			Month 1		
Participant id	UNISA						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this Workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: The project management WP will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The project management WP will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.
- *Dissemination of results*

Deliverables

D4.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D4.0.1

Workpackages of SP 4: Security and trust management

WP Title: Trust management							
Workpackage number	**WP 4.1**		**Start date or starting event:**		Month 1		
Participant id	UOP	CTI	**UNISA**	UOI	KUL	CYB	
Person-months per participant:	0.5	1	5	6	3,5	8	

Objectives

One major objective is to understand the trade-off between expressibility of the language used to specify the policy and efficiency of compliance checking algorithm. The research in this workpackage will also develop links with the research conducted in WP1.2 (Coping with selfishness) so to blend game theoretic techniques and authorization techniques to obtain reputation-based authorization schemes.

Also, we will approach the problem of multi-party trust negotiations and models and mechanism for access control.

Description of work

We plan to study multi-party trust negotiation specifically addressing issues like extensions of languages for trust negotiation and deadlock detection as other specifically related to multi-party negotiations arising from highly mobile users.

We will approach the compliance checking problem from a computational point of view. We will evaluate a broad set of options between two extremes: the most general language (e.g., Turing machines or any Turing equivalent language) and policies specified using ACL.

We will study reputation-based approach under the assumptions that the users will respond to economic incentives. This calls for a blending of techniques from Game Theory and from Security. We will interact with researchers working in WP1.2 to develop a model for reputation-based authorization and to design schemes that are resilient to attacks. Also in cooperation with WP4.2, we will investigate privacy preserving mechanisms for use in trust negotiation. Ti this aim we will rely on a reference ontology, and formalize the notion of trust requirement. We will also address privacy issues allowing a subject to adopt strategies to make the set of credentials/attributes he/she is going to release privacy preserving.

We will also explore the use of "local secure distributed computation" within the context of reputation-based systems.

We aim to define an access control model suitable for large scale scenario. We envision that such model will be characterized by rich description of subjects and objects, so that high-level policies can be directly expressed in terms of subject and object properties. Ontologies and federated identity management will be exploited to support multi-domain interoperability. Context-based access decisions will also be supported by the model.

Deliverables

D4.1.1: Trust management: State-of-the-art
D4.1.2: Trust management: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D4.1.1
Month 18: D4.1.2

Workpackages of SP 4: Security and trust management

WP Title: Privacy, identity and anonymity								
Workpackage number	**WP 4.2**		**Start date or starting event:**		Month 1			
Participant id		UNISA	**KUL**	CYB				
Person-months per participant:		4	6	8				

Objectives

To develop efficient, scalable and secure solutions for anonymous communications and anonymous transactions for a global computing environment.

Description of work

An overview will be made of proposals for technical mechanisms for privacy, identity and anonymity; this includes both mechanisms at the network level (such as mixes and crowds) and mechanisms at the application level such as anonymous credentials.

For the network level techniques, quantitative statistical models will be developed and implemented to model and analyze the behavior of the mechanisms and to improve new building blocks. This will require statistical analysis for more complex mix models. The goal is to achieve a deep insight in the trade-off between quantitative measures of anonymity and cost and quality of service. The security analysis will be improved by considering more realistic attack models. In addition, more flexible and feature-rich techniques will be developed, that include forward secrecy, congestion control and adaptivity. Part of this work will be undertaken in close collaboration with SP2.

For anonymous credentials, research will be performed into more efficient and scalable solutions, e.g. by building on solutions and techniques developed in WP4.3 and by integrating these into other protocols. This will focus on a comparison of different cryptographic techniques and on the development of novel credential based mechanisms based on well established and on novel cryptographic problems. Moreover, theoretical limitations (lower bounds) will be investigated.

Deliverables

D4.2.1: Privacy, identity, and anonymity: State-of-the-art survey
D4.2.2: Prototype for anonymous communication: Specification and design report
D4.2.3: Privacy, identity, and anonymity: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D4.2.1
Month 12: D4.2.2
Month 18: D4.2.3

Workpackages of SP 4: Security and trust management

WP Title: Secure distributed computation							
Workpackage number	**WP 4.3**		**Start date or starting event:**		Month 1		
Participant id	UOP	CTI	**UNISA**	UDRLS	KUL	DIM	CYB
Person-months per participant:	1	1	10	2,5	3,5	4	7

Objectives

To develop efficient and proactively secure protocols for applications relevant to global computing such as privacy-preserving data mining, secure algorithmic mechanisms for route discovery in networks and electronic voting.

Description of work

Thanks to the completeness results for secure distributed protocols, it is known how to compute in a secure way any efficiently computable function; however, the protocols obtained through the application of the completeness result are not practical and are only secure if executed as stand-alone protocols.

The research of this workpackage is articulated along several axes. Along a foundational line of research, we will extend the theory of secure distributed protocols so to be able to formally describe and analyze secure distributed protocols for global computers. Possible lines of research include the design of communication-efficient protocols for computing any (efficiently) computable function that can be concurrently composed or even universally composed (currently, similar general results are only known under strong set-up assumptions or for stand-alone security).

Along a more immediate line of thought, we will design efficient secure protocols for important applications among which we list privacy-preserving data mining (see also WP4.2), secure algorithmic mechanisms for route discovery in networks and electronic voting. Within the context of proactivly secure protocols we will adopt the graph-theoretic model of the pursuit-evasion distributed game and develop efficient distributed strategy to efficiently locate all the clones of a worm that is attacking a network. This entails both modeling the spreading of the worm and designing efficient distributed algorithms to coordinate the work of the antivirus mobile agents. For the more efficient protocols, prototype implementations will be provided.

At the intersection of the two lines of research lies the problem of defining universally composable security for time-stamping schemes and whether there are efficient and practical constructions for universally composable time-stamping schemes.

Deliverables

D4.3.1: Secure distributed computation: State-of-the-art survey
D4.3.2: Prototype of secure distributed protocols: Specification and design report
D4.3.3: Secure distributed computation: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D4.3.1
Month 12: D4.3.2
Month 18: D4.3.3

Workpackages of SP 5: Extending Global Computing to Wireless Users

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 5.0**	**Start date or starting event:**	Month 1				
Participant id	**CTI**						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this Workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: The project management WP will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The project management WP will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.
- *Dissemination of results*

Deliverables

D5.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D5.0.1

Workpackages of SP 5: Extending global computing to wireless users

WP Title: Resource management and quality-of-service							
Workpackage number	**WP 5.1**		**Start date or starting event:**		Month 1		
Participant id	UOP	CTI	CUI	CAU	**UDRTV**	INRIA	
Person-months per participant:	4	2	4	2	6	4	

Objectives

We will study range assignment problems with specific properties. We also aim improving the adherence of existing energy cost models to real features of wireless networks.

We will investigate the existence of feasible, approximation truthful mechanisms for range assignment problems, that is, truthful mechanisms that compute in polynomial-time an approximate solution. Another important goal is the design of efficient collision avoidance protocols in sensor networks. Energy efficient routing shall also be considered during the first 18 months.

Description of work

A first goal concerning range assignment is the design and analysis of efficient range assignment strategies that achieve bounded-hop broadcast, accumulation or full connectivity within minimal energy. We also aim to start the study of adaptive and scalable strategies. Furthermore, we shall start considering the issue of improving the adherence of existing energy cost models, e.g., min-max, to real features of wireless networks.

Selfish behaviour of the network devices will also be considered. We will investigate the existence of feasible, approximation truthful mechanisms for range assignment problems, that is, truthful mechanisms that compute in polynomial-time an approximate solution.

Another important goal is the design of efficient collision avoidance protocols in sensor networks, particularly for the case of multi-path data propagation.

Energy efficient routing will also be considered during the first 18 months. The following lines of research will be followed: (i) Cluster versus flat routing, (ii) Data aggregation and power efficiency, (iii) Data aggregation and power efficiency versus latency

All the above outlined issues will be considered in an experimental setting as well.

Deliverables

D5.1.1: Resource management and quality-of-service in wireless networks: State-of-the-art survey

D5.1.2: Resource management and quality-of-service in wireless networks: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D5.1.1
Month 18: D5.1.2

Workpackages of SP 5: Extending global computing to wireless users

WP Title: Dynamical aspects of network design and topology control							
Workpackage number	**WP 5.2**		**Start date or starting event:**		Month 1		
Participant id	UOP	CNRS	UPB	CTI	**CUI**	UDRTV	UPC
Person-months per participant:	3	3	2	2	6	3	2
Participant id	UDRLS	INRIA	DIM	UCY			
Person-months per participant:	2	2	3	1.5			

Objectives

The main goals of the first 18 months are (i) to analyze and handle topologies as basic network structures for overlay computing; (ii) to well understand the impact of the shape of the transmissions on the communications; (iii) to improve the understanding of the collisions. Concerning the shape of the transmissions, in [BLR05] a stochastic method [RM51] was implemented in order to compute numerically some critical parameters of the networks to ensure a particular task to be possible. For example, considering the process of localization in wireless networks, it is numerically shown that increasing the angle of emission increases the chance of success of the process. The most relevant observation is the existence of a critical value of the angle of emission under which the process almost surely fails and above which almost surely succeeds. The process then shows a phase transition when the angle of emission increases.

Description of work

We plan to extend this numerical investigation of the impact of the shape of transmissions on the localization process. The main goal of such an extension is to get a more definite opinion on the existence of a phase transition and also a more global point of view by considering many values of the others parameters. Getting such numerical results helps in pointing out what we can expect to prove theoretically.

We will investigate and propose dynamical and stochastic models of the interference arising in routing in wireless sensor networks.

We will also propose implicit topology management schemes including sleep-awake mechanisms in wireless sensor networks. Such schemes aim to appropriately put sensors asleep for certain periods of time, in order to save energy, without however affecting the network connectivity (and thus the efficiency of data propagation) too much.

The study of range assignment for station positioning will be started by further investigating the relationship with facility location problems. We also plan to analyze the different models for the non-homogeneous case in the literature by experimental evaluation.

Deliverables

D5.2.1: Dynamical aspects of network design and topology control: State-of-the-art survey

D5.2.2: Dynamical aspects of network design and topology control: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D5.2.1

Month 18: D5.2.2:

Workpackages of SP 5: Extending global computing to wireless users

WP Title: Mobility and fault tolerance							
Workpackage number	**WP 5.3**		**Start date or starting event:**		Month 1		
Participant id	CNRS	UPB	**CTI**	UOI	CUI	UDRTV	UPC
Person-months per participant:	3	2	4	3	4	4	1
Participant id	UDRLS	KUL	INRIA				
Person-months per participant:	2	0.5	3				

Objectives

We plan here to model and characterize mobility and fault tolerance in scalable, heterogeneous wireless networks.

Description of work

Concerning the relationship between fault-tolerance and range assignment, we intend to pursue the objective of maintaining a certain connectivity predicate in a wireless network affected by a transient component failure, by aiming both to save the energy consumption, and to preserve as much as possible the preexisting network. More precisely, the network functionality will be re-established by satisfying the following constraints, in this order: (a) minimize the number of changes, where a change is any network operation defined a priori (e.g., a range increment of a radio station, etc.); (b) minimize the objective energy function addressed by the original network over the restricted set of feasible solutions defined by (a). Special attention will be devoted to the so-called "local" maintenance of the network functionality. We also start the study of how to compare the quality of the constrained solutions with respect to the optimal ones, as computed from scratch in the residual network (i.e., the network deprived of the failed component).

Our goal for the first phase of this project is to study the problem of designing fault tolerant routing protocols having completion time as function (also) of the number of faults suffered by the network. An important aspect we will consider is the power of the Fault-Adversary. During this phase we intend to consider fault tolerance against random adversary issues by studying the average-case and by using Smooth Analysis.

As to the network survivability we aim, in this first phase, to study the possibility of extending the wired solutions to the wireless context.

We also plan to investigate the significant impact of the mobility rate and the user density on the performance of routing protocols in ad-hoc mobile networks#

Deliverables

D5.3.1: Mobility and fault tolerance: State-of-the-art Survey
D5.3.2: Mobility and fault tolerance: Algorithmic solutions and recommendations

Milestones and expected result

Month 09: D5.3.1
Month 18: D5.3.2

Workpackages of SP 6: Design and implementation of components and applications for programmable overlay computers

WP Title: Subproject management and dissemination activities							
Workpackage number	**WP 6.0**	**Start date or starting event:**	Month 1				
Participant id	TILS						
Person-months per participant:	1.5						

Objectives

The goal of this workpackage is to guarantee the successful progress of the subproject within the agreed time, cost and quality limits as defined by the project contract signed with the EU and the Consortium Agreement signed between the partners. This workpackage will also deal with establishing effective communication among the consortium partners, as well as the effective dissemination of subproject results, management of intellectual property rights and patent applications.

Description of work

Work within this Workpackage will be divided into a number of distinct tasks:

- *Set up of subproject organizational structure*: this task will establish a common baseline for: task and responsibility allocation; design a verification and validation plan and a measurement plan; design and implement the project web-site.
- *Administrative Issues Management*: the goal of the task is to ensure on-time provision of periodic management reports, and cost statements. Handle the EC project reviews and payment issues. Establish associate contracts with each project partner
- *Tracking*: The project management WP will take measures to guarantee control, validation and verification of project results, ensure that plans are fulfilled and implement necessary corrective actions.
- *Intellectual Property management*: The project management WP will prepare a Plan for intellectual property management. A “Consortium Agreement” will define consortium members’ rights and duties for the exploitation phase.
- *Dissemination of results*

Deliverables

D6.0.1 Subproject report on the activities of the first 18 months

Milestones and expected result

Month 18: D6.0.1

Workpackages of SP 6: Design and implementation of components and applications for programmable overlay computers

WP Title: Specification and design of the platform							
Workpackage number	**WP 6.1**		**Start date or starting event:**		Month 1		
Participant id	**UOP**	TILS	CNRS	UPB	CTI	UNISA	UOI
Person-months per participant:	6	6	2	1	1	2	1
Participant id	CUI	MPII	CAU	UDRTV	UOA	UNIPD	ETHZ
Person-months per participant:	1	0.5	0.5	0.5	1	0.5	0.5
Participant id	UPC	UDRLS	KUL	INRIA	DIM	UCY	CYB
Person-months per participant:	1	0.5	0.5	0.5	0.5	1	1

Objectives

The main objective of WP6.1 is to define the architecture and the functionalities of the Programmable Overlay Computing Platform.

Description of work

In this workpackage, we will select the platform to be used as the heart of the overlay computer (typical candidate platforms include JXTA, DHT-based platforms, GnuNet, and JADE), we will define the functionalities that will enhance it and select the corresponding algorithms for the implementation of these functionalities. The platform should provide basic primitives for distributed computing and communication. Functionalities for resource discovery, bandwidth sharing, data management, load management, scheduling, trust management, secure computation, etc. will enhance and extend the basic primitives of the platform. The algorithms that will be used for these implementations will be selected among the ones developed and validated within subprojects SP1, SP2, SP3, and SP4 while some of them will also take the specific characteristics of wireless devices (SP5) into account. It is expected that, for some functionalities, more than one algorithm will be selected for implementation.

Deliverables

D6.1.1: Programmable Overlay Computing Platform: Design Report

Milestones and expected result

Month 12: D6.1.1, Platform specification completed

Workpackages of SP 6: Design and implementation of components and applications for programmable overlay computers

WP Title: Implementation of platform components							
Workpackage number	**WP 6.2**		**Start date or starting event:**	Month 1			
Participant id	**UOP**	TILS	CNRS	UPB	CTI	UNISA	UOI
Person-months per participant:	4	6	5	4	4	5	6
Participant id	CUI	MPII	CAU	UDRTV	UNIPD	ETHZ	UDRLS
Person-months per participant:	4	3	3	3	3	3	3
Participant id	KUL	INRIA					
Person-months per participant:	3	3					

Objectives

To provide implementations of functionalities selected in WP6.1 on top of the selected platform (these implementation will provide the basis for the components of the platform) and prototypes designed in workpackages WP3.3 and WP3.4.

Description of work

The partners involved will implement a sufficient set of functionalities selected in the context of WP6.1. At this phase, functionalities will run on the platform as stand-alone software. Some preliminary testing (on this stand-alone basis) will be performed here. So, one of the outcomes of WP6.2 will be a set of implementations on top of the platform.

In addition, we will work on the implementation of components of the following propotypes:

- Microbenchmarking software package (based on developments of WP3.2)
- Prototype for process migration and load balancing in BSP-based P2P system (based on developments of WP3.2)
- Generic engine for specifying computations in overlay computers (based on developments of WP3.4)
- Prototype for anonymous communication (based on developments of WP4.2)
- Prototype for secure distributed protocols (based on developments of WP4.3)

Deliverables

D6.2.1 Programmable Overlay Computer – Software components
D6.2.2 Microbenchmarking software package
D6.2.3 Prototype for process migration and load balancing in BSP-based P2P system
D6.2.4 Generic engine for specifying computations in overlay computers
D6.2.5 Prototype for anonymous communication
D6.2.6 Prototype for secure distributed protocols

Milestones and expected result

Month 18: D6.2.1
Month 18: D6.2.2
Month 18: D6.2.3
Month 18: D6.2.4
Month 18: D6.2.5
Month 18: D6.2.6

Workpackages of SP 6: Design and implementation of components and applications for programmable overlay computers

WP Title: Integration and testing of the platform							
Workpackage number	**WP 6.3**		**Start date or starting event:**		Month 1		
Participant id	UOP	**TILS**	CNRS	UPB	CTI	UNISA	UOI
Person-months per participant:	2	18	1	1	1	1	1
Participant id	CUI	MPII	CAU	UDRTV	UNIPD	ETHZ	UDRLS
Person-months per participant:	1	0.5	0.5	0.5	0.5	0.5	0.5
Participant id	KUL	INRIA					
Person-months per participant:	0.5	0.5					

Objectives

The main objectives of this workpackage are

- to integrate some of the implementation produced in the context of WP6.2 into a coherent software package; this will be a very first release of the Programmable Overlay Computing Platform.
- to define a preliminary Testbed Architecture where the platform and applications will be validated

Description of work

This workpackage will integrate the set of validated implementations delivered in the context of WP6.2 to create an enhanced platform. In this task, we will select the components which will be part of the platform, and fine-tune the interface of the functionalities (i.e., fine-tune their parameters, options, etc.) from the programmer's point of view. The outcome of this workpackage will be an integrated software package, which will be an enhancement of the basic computing/communication primitives of the original software platform. Test suites will be developed for technical testing of the critical platform functionalities (these suites may be extended and reused in later phases of the project).

This WP will also investigate requirements and architectural guidelines for distributed testbeds. In the first 18 months, along with setting up the necessary resources for testing the platform, we will also study the state-of-the-art and perform further research on mechanisms that are desirable for testbeds oriented to overlay computers. A main requirement here is that tests should be independently and without interference performed by different users in a highly distributed setting. As a result of this task, we will produce architectural guidelines for distributed testbed design.

Deliverables

D6.3.1-Month 18 – Programmable Overlay Computing Platform: Preliminary version
D6.3.2-Month 18 – Architectural guidelines for testbed design

Milestones and expected result

Month 18 – Preliminary version of the platform and testbed set up

B.9 Other issues

B.9.1 Ethical issues

(If there are ethical issues associated to the subject of the proposal, show they have been adequately taken into account - indicate which national and international regulations are applicable and explain how they will be respected. Explore potential ethical aspects of the implementation of project results. Include the Ethical issues form. See Annex 3 of the Guide for Proposers for more information on ethical issues.)

Ethical issues form

A. Proposers are requested to fill in the following table

Does your proposed research raise sensitive ethical questions related to:	**YES**	**NO**
Human beings		X
Human biological samples		X
Personal data (whether identified by name or not)		X
Genetic information		X
Animals		X

B. Proposers are requested to confirm that the proposed research does not involve:

Research activity aimed at human cloning for reproductive purposes,

Research activity intended to modify the genetic heritage of human beings which could make such changes heritable[1]

Research activity intended to create human embryos solely for the purpose of research or for the purpose of stem cell procurement, including by means of somatic cell nuclear transfer;

Research involving the use of human embryos or embryonic stem cells with the exception of banked or isolated human embryonic stem cells in culture.[2]

Confirmation : the proposed research involves none of the issues listed in section B	**YES**	**NO**
	X	

B.9.2 Other EC-policy related issues

-- Not applicable --

[1] Research relating to cancer treatment of the gonads can be financed

[2] Applicants should note that the Council and the Commission have agreed that detailed implementing provisions concerning research activities involving the use of human embryos and human embryonic stem cells which may be funded under the 6th Framework Programme shall be established by 31 December 2003. The Commission has stated that, during that period and pending establishment of the detailed implementing provisions, it will not propose to fund such research, with the exception of the study of banked or isolated human embryonic stem cells in culture.

B.10 Gender issues

B.10.1. Gender Action plan

The project is expected to fully conform to the European policy of equal opportunities between women and men as enshrined in the Treaty of the European Union. The participants conform to Articles 2 and 3 and adhere to the Community tasks of eliminating inequalities and promoting gender equality.

In the consortium there is already a number of women participants either as researchers or students. At the management level, for example, there are four partners (UOI, UDRTV, UPC, INRIA) with a woman as person-in-charge; three of them are workpackage leaders. Although this participation is clearly above average compared to the current situation in Computer Science and Information Technology in general, it is rather far from the goal of the European Union of eliminating disparities.

In order to further reduce disparities, all efforts will be made so that the number of women participating in the project increases, their role is enhanced and becomes more visible. This will be reflected in the official announcements published in order to recruit the personnel needed for the completion of the project. The main criteria for staff recruitment will be absolutely based on appropriate qualifications (not sex, age, or national differences). We will also make our project and its results known at established scientific events devoted to the promotion on women in Computer Science at international and national level (e.g., events and activities organized by the ACM Committee on Women in Computing). Furthermore, the partners will encourage the participation of female researchers and students in project meetings, seminars, workshops and site exchanges. Whenever possible, young researchers will be invited to present technical results for the site they represent.

B.10.2. Gender issues.

There are no gender issues associated to the subject of the project.

Bibliography

[A+00] N. M. Amato, J. Perdue, A. Pietracaprina, G. Pucci, M. Mathis. “Predicting performance on SMP's. A case study: The SGI Power Challenge”, IPDPS 2000, pp. 729-737.

[ABCR+96] A. Aggarwal and A. Bar-Noy and D. Coppersmith and R. Ramaswami and B. Schieber and M. Sudan, “Efficient Routing in Optical Networks”, Journal of the ACM, 46(6), Nov 1996, pp. 973-1001.

[ACKM03] Gustavo Alonso, Fabio Casati, Harumi Kuno, Vijay Machiraju, “Web Services”, Springer 2003.

[ACKP00] Auletta, V., Caragiannis, I., Kaklamanis, C., and Persiano, P. Randomized path coloring on binary trees. In APPROX (2000), pp.60-71.

[ACL00] Aiello, W., Chung, F., and Lu, L. A random graph model for power law graphs. In ACM Symposium on Theory of Computing, STOC-00 (2000), pp.171-180.

[ALPH01] Adamic, L.A., Lukose, R.M., Puniyani, A.R., and Huberman, B.A., “Search in power-law networks”, Phys. Rev. E 64 (2001), 46135.

[APHS02] Karl Aberer, Magdalena Punceva, Manfred Hauswirth, Roman Schmidt, “Improving Data Access in P2P Systems”, IEEE Internet Computing 6(1): pp. 58-67 (2002).

[ARV04] Arora, Rao, and Vazirani, “Expander flows, geometric embeddings and graph partitioning”, STOC: ACM Symposium on Theory of Computing (STOC)(2004).

[AT01] A. Archer, E. Tardos, “Truthful Mechanisms for One-parameter Agents”, IEEE FOCS, 2001.

[AuletalICALP04] V.Auletta, R. De Prisco, P.Penna, and G.Persiano, "The power of verification for one-parameter agents", International Colloquium on Automata, Languages and Programming (ICALP), LNCS, pp. 171-182, 2004.

[AuletalSTACS04] V.Auletta, R. De Prisco, P.Penna, and G.Persiano, "Deterministic truthful approximation mechanisms for scheduling related machines", Annual Symposium on Theoretical Aspects of Computer Science (STACS), LNCS, pp. 608-619, 2004.

[BaAl99] Barabasi, A.-L., and Albert, R., “Emergence of scaling in random networks”, Science (1999).

[Backes:cryptolib] M. Backes, B. Pfitzmann, and M. Waidner, “A Universally Composable Cryptographic Library”, 10th ACM Conference on Computer and Communications Security, 2003.

[BAJ99] A. Barabasi, R. Albert, and H. Jeong. “Scale-free characteristics of random networks: the topology of the World Wide Web”. Physica A, 272: pp. 173-187, 1999.

[BaKo00] Bartal, Y., and Koutsoupias, E. On the competitive ratio of the work function algorithm for the k-server problem. In Symposium on Theoretical Aspects of Computer Science, STACS-00 (2000), pp.605-613.

[BBCFM] L. Booth, J. Bruck, M. Cook, M. Franceschetti, R. Meester, “Percolation in multi-hop wireless networks”, submitted.

[BCLP+03] Bouklit, M., Coudert, D., Lalande, J.-F., Paul, C., and Rivano, H. “Approximate multicommodity flow for WDM networks design”, SIROCCO 10 (Umea, Sweden, 2003), J. Sibeyn, Ed., no. 17 in Proceedings in Informatics, Carleton Scientific, pp. 43-56.

[BDKK+04] Stoica, I., Morris, R., Liben-Nowell, D., Karger, D.R., Kaashoek, M.F., Dabek, F., and Balakrishnan, H. Chord: A scalable peer-to-peer lookup protocol for internet applications. IEEE/ACM Transactions on Networking (2004), to appear. Conference version : ACM SIGCOM-01.

[BDS03] V. C. Barbosa, R. Donangelo, S. R. Souza, “Directed Cycles and Related Structures in Random Graphs: I-Static Properties”, Physica A 321, pp. 381-397, 2003.

[Bea00a] Beauquier, B. “Communications dans les Rιseau Optiques”, PhD thesis, Ecole Doctorale Stic de l'Universitι de Nice, 2000.

[BergerEtAl03] Berger, Bollobas, Borgs, Chayes, and Riordan, “Degree distribution of the FKP network model”, ICALP: Annual International Colloquium on Automata, Languages and Programming (2003).

[BFJ03] Bui-Xuan, B., Ferreira, A., and Jarry, A., “Computing shortest, fastest, and foremost journeys in dynamic networks”, International Journal of Foundations of Computer Science 14, 2 (April 2003), pp. 267-285.

[BFJ03a] Bui-Xuan, B., Ferreira, A., and Jarry, A., “Evolving graphs and least cost journeys in dynamic networks”, WiOpt'03 - Modeling and Optimization in Mobile, Ad-Hoc and Wireless Networks (Sophia Antipolis, March 2003), INRIA Press, pp. 141-150.

[BG03] K.Bennett, C.Grothoff, T.Horozov, I.Patrascu, and T.Stef, “GAP - practical anonymous networking”, Privacy Enhancing Technologies workshop (PET 03), Lecture Notes in Computer Science 2760, pp. 141-160. Springer-Verlag, March 2003.

[BGI87] R. Bar-Yehuda, O. Goldreich, A. Itai, “On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomisation”, JCSS, 45, pp. 104-126 (1992) (preliminary version in 6th ACM PODC, 1987.

[BGPR+96] Bermond, J.-C., Gargano, L., Perennes, S., Rescigno, A., and Vaccaro, U. “Effective collective communication in optical networks”. ICALP 96 (1996), G. Goos, J. Hartmanis, and J. V. Leeuwen, Eds., vol. 1099 of Lectures Notes In Computer Science, Springer Verlag, pp. 574-585.

[BGW] M.Ben-Or, S.Goldwasser, and A.Wigderson, “Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation” (Extended Abstract)}, 20th {ACM} Symposium on the Theory of Computing, pp. 1-10, 1988.

[BHP98] Beauquier, B., Hell, P., and Perennes, S. “Optimal wavelength-routed multicasting”. Discrete Applied Mathematics 84 (1998), pp. 15-20.

[Bilo] V.Bilo, C. Di Francescomarino, M. Flammini, and G. Melideo, "Sharing the cost of muticast transmissions in wireless networks", 16th ACM Symposium on Parallelism in ALgorithms and Architectures (SPAA), pp. 180-187, June 2004.

[BIM00] A.Beimel, Y.Ishai, and T.Malkin, “Reducing the servers computation in private information retrieval: PIR with preprocessing”, Advances in Cryptology - CRYPTO 2000, volume 1880 of Lecture Notes in Computer Science}, pp. 55-73, 2000.

[BKK03] Reinhard Braumandl, Alfons Kemper, Donald Kossmann, “Quality of service in an information economy”, ACM Trans. Internet Techn. 3(4): pp. 291-333 (2003).

[BKL02] Liben-Nowell, D., Balakrishnan, H., and Karger, D. Analysis of the evolution of peer-to-peer systems. In ACM Conference on Principles of Distributed Computing (PODC-02) (2002).

[BKRSW96] Borodin, A., Kleinberg, J., Raghavan, P., Sudan, M., and Williamson, D.P., “Adversarial queueing theory”, STOC96 (1996), pp.376-385.

[BLM01] Bertrand, S., Laugier, A., and Mahey, P. “Routing flows in networks with heterogenous protocols and path-dependent edge costs”, 3e Rencontres francophones sur les Aspects Algorithmiques des Telecommunications (ALGOTEL'2001), pp.55-60.

[BLR05] Bouget M., Leone P., Rolim J., "Application of a Stochastic Method to Determine Critical Parameters in Wireless Sensors Networks", to be submitted to Workshop of Experimental and Efficient Algorithms, Santorini, Greece, May 2005.

[BMPP03] Bermond, J.-C., Marlin, N., Peleg, D., and Pırennes, S. "Directed virtual path layout in ATM networks". Theoretical Computer Science 291 (2003), pp. 3-28.

[BMPP99a] Bermond, J.-C., Marlin, N., Peleg, D., and Perennes, S. "Virtual path layouts in simple ATM networks", IFIP ATM '98, Ilkley, U.K., 1998

[Bol85] Bollobűs, B., "Random Graphs", Academic Press, London, 1985.

[BonehEC04] D.Boneh, G.DiCrescenzo, R.Ostrovsky, and G.Persiano, "Public key encryption with keyword search", C.Cachin and J.Camenish, editors, Advances in Cryptology - Eurocrypt 2004, volume 3027 of Lecture Notes in Computer Science, pp. 506-522. Springer Verlag, 2004.

[Bourgain85] Bourgain, J., "On lipschitz embedding of finite metric spaces in hilbert spaces", Israeli J. Math. 52 (1985), 46-52.

[BPS02] Beauquier, B., Pırennes, S., and Syska, M. "Efficient access to optical bandwidth, routing and grooming in WDM networks: State-of-the-art survey". IST-CRESCCO report, Projet MASCOTTE (CNRS/INRIA/UNSA), Sophia Antipolis, 2002.

[BrGo95] Bronnimann, H., and Goodrich, M. Almost optimal set covers in finite vc-dimension. Discrete \& Computational Geometry 14, 4 (1995), 463-479.

[BRST] Bollobas, B., Riordan, O., Spencer, J., and Tusnady, G., "The degree sequence of a scale-free random graph process", Random Structures and Algorithms 18, 3 (2001), pp. 279-290.

[BRT01] Borodin, A., Roberts, G.O., Rosenthal, and Tsaparas, P. Link analysis ranking algorithms theory and experiments. In 10th World Wide Web Conference, Hong Kong (2001). journal version accepted in ACM Transactions on Internet Technologies.

[BRVX04] Bent, Michael Rabinovich, Geoffrey M. Voelker, Zhen Xiao, "Characterization of a large web site population with implications for content delivery", WWW 2004: pp. 522-533.

[BuSa04] A. Buldas and M. Saarepera, "On provably secure time-stamping schemes", Advances in Cryptology - ASIACRYPT 2004, 2004.

[BY98] Borodin, A., and El-Yaniv, R., "Online Computation and Competitive Analysis", Cambridge University Press, 1998.

[Cachin] C.Cachin, "Distributing Trust on the Internet", International Conference on Dependable Systems and Networks (DSN-2001). IEEE, 2001.

[CaKa04] Caragiannis, I., and Kaklamanis, C. "Approximate path coloring with applications to wavelength assignment in WDM optical networks". STACS-2004 (2004), pp. 258-269.

[Canetti] R.Canetti, "Universally Composable Security: A New Paradigm for Cryptographic Protocols", 42nd IEEE Symposium on Foundations of Computer Science, pp. 136-145, 2001.

[CCD] D.Chaum, C.Crepeau, and I.Damgard, "Multiparty Unconditionally Secure Protocols" (Extended Abstract), Symposium on the Theory of Computing, pp. 11-19, 1988.

[CCS02] Choi, H., Subramaniam, S., and Choi, H.-A. On double-link failure recovery in wdm optical networks. In IEEE Infocom (2002), pp.808-816.

[CD] Carlson, J.M., and Doyle, J., "Highly optimized tolerance: a mechanism for power laws in designed systems", Physics Review E 60, 2 (1999), pp. 1412-1427.

[CDG+03] J.Claessens, C.Diaz, C.Goemans, B.Preneel, and J.Vandewalle, "Revocable anonymous access to the Internet", Journal of Internet Research: Electronic Networking Applications and Policy, 13(4): pp. 242-258, 2003.

[CDR03a] Cole, R., Dodis, Y., and Roughgarden, T., "How much can taxes help selfish routing", ACM EC (2003), pp. 98-107.

[CDR03b] Cole, R., Dodis, Y., and Roughgarden, T., "Pricing network edges for heterogeneous selfish users", STOC (2003), pp. 521-530.

[CF] Cooper, C., and Frieze, A.M., "A general model of undirected web graphs", ESA (2001), pp. 500-511.

[CFKP+01a] Caragiannis, I., Ferreira, A., Kaklamanis, C., Pιrennes, S., and Rivano, H. "Fractional path coloring with applications to WDM networks", 28th ICALP (Crete, Greece, 2001), pp. 732-743.

[CFKP+04] Caragiannis, I., Ferreira, A., Kaklamanis, C., Pιrennes, S., Persiano, P., and Rivano, H. "Approximate constrained bipartite edge coloring". Discrete Applied Mathematics (2004), To Appear.

[CFL02] Flake, G., Lawrence, S., Giles, C.L., and Coetzee, F. Self-organization and identification of web communities. IEEE Computer 35, 3 (2002).

[CGKS95] B.Chor, O.Goldreich, E.Kushilevitz, and M.Sudan, "Private information retrieval", FOCS'1995, pp. 41-50, 1995.

[CGS04] Schupke, D., Grover, W., and Clouqueur, M. Strategies for enhanced dual failure restorability with static or reconfigurable p-cycle networks. In IEEE International Conference on Communications (ICC) (Paris, France, June 20-24 2004).

[CH90] D.Chaum, "Showing credentials without identification: Transferring signatures between unconditionally unlinkable pseudonyms", Advances in Cryptology - AUSCRYPT '90, Lecture Notes in Computer Science 453, pp. 246-264. Springer, 1990.

[CHI99] R. Cole, R. Hariharan, and P. Indyk (1999), "Tree Pattern matching and subset matching in deterministic O(nlog3n)-time", 10th ACM-SIAM SODA.

[CHPRV01] A. Clementi, G. Huiban, P. Penna, G. Rossi, Y.C. Verhoeven, "Some Recent Theoretical Advances and Open Questions on Energy Consumption in Ad-Hoc Wireless Networks", 3rd ARACNE, 2002.

[CHS05] I. Chatzigiannakis, P. Spirakis and S. Nikoletseas, "Efficient and Robust Protocols for Local Detection and Propagation in Smart Dust Networks", accepted in the ACM Mobile Networks (MONET) Journal, Special Issue on Algorithmic Solutions for Wireless, Mobile, Ad Hoc and Sensor Networks, to appear in MONE 10:1 (February 2005).

[CK01] Edith Cohen, Haim Kaplan, "Refreshment Policies for Web Content Caches", INFOCOM 2001: pp. 1398-1406.

[CKN04] Christodoulou, G., Koutsoupias, E., and Nanavati., A., "Coordination mechanisms", ICALP (Turku, Finland, 12-16 July 2004), pp. 345-357.

[CKN04a] I. Chatzigiannakis, E. Kaltsa and S. Nikoletseas, "On the Effect of User Mobility and Density on the Performance of Routing Protocols for Ad-hoc Mobile Networks", in the IEEE International Conferenec (ICON 2004). Also, accepted in the Journal of Mobile Computing and its Applications, to appear in 2004.

[CKP02] Ioannis Caragiannis and Christos Kaklamanis and Evi Papaioannou, "Efficient On-Line Frequency Allocation and Call Control in Cellular Networks". Theory of Computing Systems, 35(5), 2002, pp. 521-543.

[CKP03] Ioannis Caragiannis and Christos Kaklamanis and Evi Papaioannou, “Simple On-Line Algorithms for Call Control in Cellular Networks”, WAOA 2003, pp. 67-80.

[CKPS04] Caragiannis, I., Kaklamani, C., Persiano, P., and Sidiropoulos, A. “Fractional and integral coloring of locally-symmetric sets of paths on binary trees”, WAOA 2003 (2004), pp. 81-94.

[CKZ04] Cooper, C., Klasing, R., and Zito, M. Dominating sets in web graphs. In Proceedings of the Third Workshop on Algorithms and Models for the Web-Graph (WAW 2004) (2004), Lecture Notes in Computer Science, Springer-Verlag.

[cliftonmarks] C.Clifton and D.Marks, “Security and privacy implications of data mining”, 1996 ACM SIGMOD Workshop of Data Mining and Knowledge Discovery, 1996.

[CMS01a] A.E.F. Clementi, A. Monti, R. Silvestri, “Selective Families, Superimposed Codes, and Broadcasting in Unknown Radio Networks”, 12th ACM-SIAM SODA, pp. 709-718, 2001 (to appear on TCS).

[CMS01b] A.E.F. Clementi, A. Monti, R. Silvestri, “Round Robin is Optimal for Fault-Tolerant Broadcasting on Wireless Networks”, 9th European Symposium on Algorithms (ESA'01), LNCS, 2001.

[CN03] I. Chatzigiannakis and S. Nikoletseas, “A Sleep-Awake Protocol for Information Propagation in Smart Dust Networks”, in the ACM Mobile Networks (MONET) Journal, 2004. Also, in the Proceedings of the 3rd Workshop on Mobile and Ad-hoc Networks (WMAN), IEEE Press, 2003.

[CNS03] I. Chatzigiannakis, S. Nikoletseas and P. Spirakis, “Distributed Communication and Control Algorithms for Ad-hoc Mobile Networks”, in the Journal of Parallel and Distributed Computing (JPDC), Special Issue on Mobile Ad-hoc Networking and Computing, 63 (2003) pp. 58-74, 2003.

[composable1] R.Canetti, Y.Lindell, R.Ostrovsky, and A.Sahai, “Universally composable two-party and multi-party secure computation”, Thiry-Fourth Annual ACM Symposium on Theory of Computing, pp. 494-503, New York, NY, USA, 2002. ACM Press.

[CoRi02] Coudert, D., and Rivano, H. Lightpath assignment for multifibers WDM optical networks with wavelength translators. In Proc. of IEEE GlobeCom'02 (Taipei, Taiwan, nov 2002).

[COSF01] Ugur Hetintemel, Banu Φzden, Abraham Silberschatz, Michael J. Franklin, “Design and Evaluation of Redistribution Strategies for Wide-Area Commodity Distribution”, ICDCS 2001: pp. 154-161.

[CPD01] R. Clinton Whaley, A. Petitet, J. Dongarra, “Automated empirical optimizations of software and the ATLAS project”. Parallel Computing 27(1-2): pp. 3-35, 2001.

[CPS02] A.E.F. Clementi, P. Penna, R. Silvestri, “On the Power Assignment Problem in Radio Networks”, MONET(2): 125-140 (2004) [CS03] V. Conitze, T.Sandholm Complexity Results about Nash Equilibria, 18th IJCAI, pp. 765--771, 2003.

[CRR03] Coudert, D., Rivano, H., and Roche, X. A combinatorial approximation algorithm for the multicommodity flow problem. In WAOA 03 (Budapest, Hungary, sep 2003), K.Jansen and R.Solis-Oba, Eds., no.2909 in Lecture Notes in Computer Science, Springer-Verlag, pp.256-259.

[CS02] E. Cohen and S. Shenker, “Replication Strategies in Unstructured Peer-to-Peer Networks”, SIGCOMM 2002.

[CS03] V. Conitze, T.Sandholm, “Complexity Results about Nash Equilibria”, 18th IJCAI, pp. 765-771, 2003.

[CV02] A. Czumaj and B. Vφcking, “Tight bounds for worst-case equilibria”, 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02), pp. 413-420, 2002.

[Cyb89] G. Cybenko, “Load balancing for distributed memory multiprocessors”. Journal of Parallel and Distributed Computing, 7, pp. 279-301, 1989.

[datamining1] R.Agrawal, A.Evfimievski, and R.Srikant, “Information sharing across private databases”, ACM SIGMOD Conference on Management of Data, 2003.

[datamining2] A.Evfimievski, R.Srikant, R.Agrawal, and J.Gehrke, “Privacy preserving mining of association rules”, 8th ACM SIGKDD International Conference on Knowledge Discovery in Databases and Data Mining, 2002.

[DBT03] O. Dousse, F. Baccelli, P. Thiran, “Impact of Interferences on Connectivity in Ad-Hoc Networks”, Infocom, San Francisco, April 2003.

[DFM99] R. Diekmann, A. Frommer, B. Monien, “Efficient schemes for nearest neighbor load balancing”. Journal of Parallel and Distributed Computing, 25(7), pp. 789-812, 1999.

[DGMY03] Neil Daswani, Hector Garcia-Molina, Beverly Yang, “Open Problems in Data-Sharing Peer-to-Peer Systems”, ICDT 2003: pp. 1-15.

[DisTrusMan] M.Blaze, J.Feigenbaum, and J.Lacy, “Decentralized Trust Management”, 17th IEEE Conference on Security and Privacy, 1996.

[DMMP+02] Dacomo, A., Patre, S.D., Maier, G., Pattavina, A., and Martinelli, M. Design of static resilient WDM mesh-networks with multiple heuristic criteria. In IEEE Infocom (2002).

[DMS04] R.Dingledine, N.Mathewson, and P.Syverson, “ Tor: The second-generation onion router”, 13th USENIX Security Symposium, August 2004.

[DNS] C.Dwork, M.Naor, and A.Sahai, “Concurrent Zero-Knowledge”, 30th ACM Symposium on Theory of Computing, pp. 409-418, 1998.

[DNS04] T. Dimitriou, S. Nikoletseas and P. Spirakis, “Analysis of the Information Propagation Time in Mobile Networks”, in the 3rd ADHOC-NOW Conferene, LNCS, 2004.

[DS04] C.Diaz and A.Serjantov, “Generalising mixes”, Privacy Enhancing Technologies, Dresden (Germany), LNCS 2760, pp. 18-31. Springer-Verlag, 2003.

[DTH02] O. Dousse, P. Thiran, M. Hasler, “Connectivity in Ad-Hoc and Hybrid Networks”, IEEE Infocom, June 2002.

[DV88] D. J. Daley, D. Vere-Jones, “An Introduction to the theory of Point Processes”, Springer Verlag, 1988.

[ElPe01] Michael Elkin and David Peleg, “(1+epsilon, beta)-spanner constructions for general graphs”, STOCS 2001, pp. 173-182

[ErJa96] Erlebach, T., and Jansen, K. “Scheduling of virtual connections in fast networks”, 4th Workshop on Parallel Systems and Algorithms (PASA'96) (1996), World Scientific, pp. 13-32.

[ErJa97] Erlebach, T., and Jansen, K. “Call scheduling in trees, rings and meshes”. 30th Hawaii International Conference on System Sciences (HICSS'97) (1997), vol. 1, IEEE Computer Society Press, pp. 221-222.

[ErJa98] Erlebach, T., and Jansen, K. “Maximizing the number of connections in optical tree networks”. 9th Annual International Symposium on Algorithms and Computation (ISAAC'98) (1998), S. Verlag., Ed., vol. LNCS 1533, pp. 179-188.

[Erl99] Erlebach, T. “Maximum weight edge--disjoint paths in bidirected tree”, Communication and Data Management in Large Networks. Workshop of INFORMATIK'99 (1999), pp. 13-1.

[ERS04] Even, G., Rawitz, D., and Shabar, S.M. Hitting sets when the vc-dimension is small. in preparation, 2004.

[FairPlay] D.Malkhi, N.Nisan, B.Pinkas, and Y.Sella, "Fairplay - a secure two-party computation system", Usenix Security, pp. 80-91, 2004.

[faloutsos99powerlaw] Faloutsos, M., Faloutsos, P., and Faloutsos, C. On power-law relationships of the internet topology. In SIGCOMM (1999), pp.251-262.

[FeiPapShe00] J.Feigenbaum, C.H. Papadimitriou, and S.Shenker, "Sharing the cost of multicast transmissions", Journal of Computer and System Sciences, 63(1): pp. 21-41, 2001.

[FFF99] Faloutsos, M., Faloutsos, P., and Faloutsos, C., "On power-law relationships of the internet topology", SIGCOMM (1999), pp. 251-262.

[FGL+03a] R. Feldmann, M. Gairing, T. Lócking, B. Monien, and M. Rode. "Nashification and the coordination ratio for a selfish routing game", 30th International Colloquium on Automata, Languages, and Programming (ICALP'03), LNCS 2719, pp. 514-526, 2003.

[FGL+03b] R. Feldmann, M. Gairing, T. Lócking, B. Monien, and M. Rode. "Selfish routing in non-cooperative networks: A survey", 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03), LNCS 2747, pp. 21-45, 2003.

[FJ98] M. Frigo, S. Johnson, "FFTW: An adaptive software architecture for the FFT", ICASSP, 1998, pp. 1381-1384.

[FKK+02] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. "The Structure and Complexity of Nash Equilibria for a Selfish Routing Game", ICALP, pp. pp. 123-134, 2002.

[FKP02] Fabrikant, Koutsoupias, and Papadimitriou, "Heuristically optimized trade-offs: A new paradigm for power laws in the internet", ICALP: Annual International Colloquium on Automata, Languages and Programming (2002).

[Fle99] Fleischer, L. Approximating fractional multicommodity flow independent of the number of commodities. In IEEE Symposium on Foundations of Computer Science (1999), pp.24-31.

[FM02] M.J. Freedman and R.Morris, "Tarzan: A peer-to-peer anonymizing network layer", 9th ACM Conference on Computer and Communications Security (CCS 2002). CCS, 2002.

[FoFu58] Ford, L.R., and Fulkerson, D.R. Constructing maximal dynamic flows from static flows. Operations Research 6 (1958), 419-433.

[FoFu62] FordJr, L., and D.R.Fulkerson. Flows in networks. Princeton University Press, 1962.

[FPRR+03] Ferreira, A., Perennes, S., Richa, A. W., Rivano, H., and Stier, N. "Models, Complexity and Algorithms for the Design of Multi-fiber WDM Networks". Telecommunication Systems 24, 2 (Oct 2003), pp. 123-138.

[FPS00] J. Feigenbaum, C. Papadimitriou, S. Shenker, "Sharing the cost of multicast transmissions", 32nd ACM Symp. on Theory of Computing (STOC'00), pp. 218-227, 2000.

[Fra90] Frank, A. Packing paths, circuits and cuts - A survey. In Paths, Flows and VLSI-Layout, B.Korte, L.Lovαsz, H.J. Prφmel, and A.Schrijver, Eds. Springer-Verlag, 1990, pp.47-100.

[FW98] Fiat, A., and (Eds.), G. J.W., "Online Algorithms: State of the Art", Springer-Verlag, 1998.

[G61] E. N. Gilbert, "Random Plane Networks", SIAM J., vol. 9, pp. 533-543, 1961.

[GaKo98] Garg, N., and Konemann, J. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In IEEE Symposium on Foundations of Computer Science (1998), pp.300-309.

[GaPe96a] Gavoille, C., and Perennes, S., "Memory requirement for routing in distributed networks", PODC 96 (1996), A.Press, Ed., pp.125-133. Best student paper award.

[GaVa00] Gargano, L., and Vaccaro, U. "Routing in all-optical networks: Algorithmic and graph—theoretic problems", tutorial. Numbers, Information and Complexity (2000), pp. 555-578.

[GeRa00] Gerstel, O., and Ramaswami, R. Survivability: A services perspective. IEEE Communications Magazine 38, 3 (2000), 104-113.

[GHP01] Gargano, L., Hell, P., and Perennes, S. "Coloring all directed paths in a symmetric tree, with an application to optical networks". Journal of Graph Theory 38, 4 (2001), pp. 183-196.

[GK00] P. Gupta, P. R. Kumar, "The Capacity of Wireless Networks", IEEE Tans. Inform. Theory, vol. 46(2), pp. 338-404, Mar. 2000.

[GLMM04] M. Gairing, T. Lócking, M. Mavronicolas, and B. Monien, "Computing nash equilibria for scheduling on restricted parallel links", 36th Annual ACM Symposium on Theory of Computing (STOC'04), pp. 613-622, 2004.

[GMW1] O.Goldreich, S.Micali, and A.Wigderson, "Proofs that Yield Nothing but Their Validity ot All Languages in NP Have Zero-Knowledge Proof Systems", Journal of the ACM}, 38(1): pp. 691-729, 1991.

[GMW2] O.Goldreich, S.Micali, and A.Wigderson, "How to Play any Mental Game - A Completeness Theorem for Protocols with Honest Majority", 19th ACM Symposium on the Theory of Computing, pp. 218-229, 1987.

[GoRa96] Ravi, R., and Goemans, M.X., "The constrained minimum spanning tree problem", 5th Scandinavian Workshop on Algorithm Theory (1996), Springer-Verlag, pp.66-75.

[GPPR01] Gavoille, C., Peleg, D., Purennes, S., and Raz, R., "Distance labeling in graphs", SODA'01 (2001), pp. 210-219.

[GPT01] C. Gaibisso, G. Proietti, R. Tan, "Efficient Management of Transient Station Failures in Linear Radio Communication Networks with Bases", 2nd Int. Workshop on Approximation and Randomized Algorithms in Communication Networks (ARACNE'01), Vol. 12 of Proc. in Informatics, Carleton Scientific, pp. 37-54, 2001.

[GPT03] C. Gaibisso, G. Proietti, R. Tan, "Optimal MST Maintenance for Transient Deletion of Every Node in Planar Graphs", 9th Int. Computing and Combinatorics Conf. (COCOON'03), LNCS 2697, Springer, pp. 404-414, 2003.

[Gri89] Grimmett, G.P. Percolation. Springer-Verlag, New York:, 1989.

[GRPS03] S.Goel, M.Robson, M.Polte, and E.Gun Sirer, "Herbivore: A scalable and efficient protocol for anonymous communication", Technical report 2003-1890. Cornell University, Ithaca, N.Y. February 2003.

[GRS00] Gerstel, O., Ramaswami, R., and Sasaki, G. "Cost-effective traffic grooming in WDM rings". IEEE/ACM Transactions on Networking 8, 5 (October 2000), pp. 618-630.

[GrSh03] Shen, G., and Grover, W.D. Extending the p-cycle concept to path segment protection for sextending the p-cycle concept to path segment protection for span and node failure recovery. IEEE Journal of Selected Areas in Communication: Optical Communications and Networking Series 21, 8 (October 2003).

[GSS03] Schupke, D., Scheffel, M., and Grover, W. Configuration of p-cycles in WDM networks with partial wavelength conversion. Photonic Network Communications 6, 3 (2003), 239-252.

[GWW02] Michael Gillmann, Gerhard Weikum, Wolfgang Wonner,"Workflow management with service quality guarantees", SIGMOD Conference 2002: pp. 228-239.

[H63] T. Harris, "The Theory of Branching Processes", Springer Verlag, 1963.

[haber91] S. Haber and W. S. Stornetta, "How to time-stamp a digital document", Journal of Cryptology, 3(2): pp. 99-111, 1991.

[HB03] M.Hansen and P.Berlich, "Identity management systems: Gateway and guardian for virtual residences", EMTEL Conference: New Media, Technology and Everyday Life in Europe Conference, London, 2003.

[HJKY95] A.Herzberg, S.Jarecki, H.Krawczyk, and M.Yung, "Proactive secret sharing or: How to cope with perpetual leakage", D.Coppersmith, editor, Advances in Cryptology - CRYPTO '95, volume 963 of Lecture Notes in Computer Science, pp. 339-352. Springer-Verlag, 1995.

[HPS02] Huiban, G., Ptrennes, S., and Syska, M. "Traffic grooming in WDM networks with multi-layer switches". IEEE ICC (New-York, 2002). CD-Rom.

[I97] P. Indyk, "Deterministic Superimposed Coding with Application to Pattern Matching", IEEE 38th FOCS, pp. 127-136, 1997.

[IK04] Y.Ishai and E.Kushilevitz, "On the hardness of information-theoretic multiparty computation", C.Cachin and J.Camenish, editors, Advances in Cryptology - Eurocrypt 2004, volume 3027 of Lecture Notes in Computer Science, pp. 439 - 455. Springer Verlag, 2004.

[IMRR+95] Marathe, M.V., Ravi, R., Sundaram, R., Ravi, S.S., Rosenkrantz, D.J., and III, H. B.H., "Bicriteria network design problems", ICALP 95 (1995), pp. 487-498.

[Indyk01] Indyk, P., "Algorithmic applications of low-distortion geometric embeddings", 42nd IEEE Symposium on Foundations of Computer Science, October 14-17, 2001, Las Vegas, Nevada, USA (1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 2001), IEEE, Ed., IEEE Computer Society Press, pp.10-33.

[intrusion] T.Wu, M.Malkin, and D.Boneh, "Building intrusion tolerant applications", 8th USENIX Security Symposium, pp. 79-91, 1999.

[IRD02] Sitaram Iyer, Antony I. T. Rowstron, Peter Druschel, "Squirrel: a decentralized peer-to-peer web cache", PODC 2002: pp. 213-222.

[J97] P. Jagers, "Coupling and Population Dependence in Branching Processes", The Annals of Applied Probability, vol. 7, No. 2, pp. 281-298, May 1997.

[JaiVaz01] K.Jain and V.V. Vazirani, "Applications of approximation algorithms to cooperative games", Annual ACM Symposium on Theory of Computing (STOC), 2001.

[Jar02] Jarry, A. "Disjoint Paths in Symmetric Digraphs". International Colloquium on Structural Information and Communication Complexity - SIROCCO (Andros, Greece, 2002), Carleton, pp. 211-222.

[JaZh02] Jansen, K., and Zhang, H. Approximation algorithms for general packing problems with modified logarithmic potential function. In IFIP TCS 2002 (2002), pp.255-266.

[JLR00] S. Janson, T. Luczac, A. Rucinski, "Random Graphs", Wiley- interscience Series in Discrete Mathematics and Optimization, 2000.

[Kar01] Karger, D.R., “A randomized fully polynomial time approximation scheme for the all-terminal network reliability problem”, SIAM Review 43, 3 (2001).

[Kar03] Karger, D., “Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm”, SODA93 (2003).

[Kar94] Karger, D.R., “Random sampling in cut, flow, and network design problems”, 26th annual ACM symposium on Theory of computing, STOC 94 (1994).

[KKKP00] L. M. Kirousis, E. Kranakis, D. Krizanc, A. Pelc, “Power Consumption in Packet Radio Networks”, Theoretical Computer Science, 243, pp. 289-305, 2000.

[KKP98] E. Kranakis, D. Krizanc, A. Pelc, ‘Fault-Tolerant Broadcasting in Radio Networks”, 6th ESA, LNCS 1461, pp. 283-294, 1998.

[Kla98] Klasing, R. Methods and problems of wavelength-routing in all-optical networks. In Proc. of the MFCS'98 Workshop on Communication, August 24-25, 1998, Brno, Czech Republic (1998), pp.1-9.

[KLA98] R. Klasing, “Methods and Problems of Wavelength-Routing in All-Optical Networks”, MFCS'98 Workshop on Communication, August 24-25, 1998, Brno, Czech Republic, pp 1-9.

[Kle99] Kleinberg, J.M., “Authoritative sources in a hyperlinked environment”, Journal of the ACM 46, 5 (1999), pp. 604-632.

[KLNP04] Klasing, R., Lotker, Z., Navarra, A., and Perennes, S. The points and vertices game. Submitted to SODA, 2004.

[KLS02] Kohler, E., Langkau, K., and Skutella, M., “Time-expanded graphs for flow-dependent time-expanded graphs for flow-dependent transit times”, 10th Annual European Symposium on Algorithms (ESA), volume 2461 of Lecture Notes in Computer Science (2002), Springer, Berlin, pp. 599-611.

[KM98] E. Kushilevitz, Y. Mansour, “Computation in Noisy Radio Networks”, 9th ACM-SIAM SODA, pp. 236-243, 1998.

[KMP04] Klasing, R., Morales, N., and Perennes, S. “Complexity of bandwidth allocation in radio networks : the static case”. Submitted to Theoretical Computer Science, 2004.

[KoPa94] Koutsoupias, E., and Papadimitriou, C.H., “On the k-server conjecture”, ACM Symposium on Theory of Computing, STOC-94 (1994), pp. 507-511.

[KoSk02] Kohler, E., and Skutella, M., “Flows over time with load-dependent transit times”, 13th annual ACM-SIAM symposium on Discrete algorithms (2002), Society for Industrial and Applied Mathematics, pp.174-183.

[Kou03] Koutsoupias, E., “Selfish task allocation”, Bulletin of EATCS, 81 (2003), 79-88.

[KP] J.Kilian and E.Petrank, “Concurrent and Resettable Zero-Knowledge in Poly-Logarithmic Rounds”, 33rd ACM Symposium on Theory of Computing (STOC '01), pp. 560-569. ACM, 2001.

[KP99] E. Koutsoupias and C. Papadimitriou, “Worst-case equilibria”, 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99), LNCS 1563, pp. 404-413, 1999.

[KPR2] J.Kilian, E.Petrank, and R.Richardson, “Concurrent Zero-Knowledge Proofs for NP”, http://www.cs.technion.ac.il/simerez/Papers/czkub-full.ps.

[KRR+00] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal, “Stochastic model for the web graph”, FOCS'00, pp. 57-65, 2000.

[KRRS+00] Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., and Upfal, E. Stochastic models for the web graph. In IEEE Symposium on Foundations of Computer Science (FOCS-00) (2000).

[KRRSTU] Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., and Upfal, E., “Stochastic models for the web graph”, 41st Annual Symposium on Foundations of Computer Science (2000), pp. 57-65.

[KRS04] Kenyon, Rabani, and Sinclair, “Low distortion maps between point sets”, STOC: ACM Symposium on Theory of Computing (STOC) (2004).

[KT00] J.Katz and L.Trevisan, “On the efficiency of local decoding procedure for error-correcting codes”, 32rd ACM Symposium on Theory of Computing (STOC '00), pp. 80-86. ACM, 2000.

[KTY04] A. Kiayias, Y. Tsiounis, and M. Yung, “Traceable signatures”, In C. Cachin and J. Camenish, editors, Advances in Cryptology - Eurocrypt 2004, volume 3027 of Lecture Notes in Computer Science, pp. 571 - 589. Springer Verlag, 2004.

[L03] Frank Leymann, “Web Services: Distributed Applications Without Limits”, BTW 2003: pp. 2-23.

[LCC+02] Q. Lv, P. Cao, E. Cohen, K. Li and S. Shenker, “Search and Replication in Unstructured Peer-to-Peer Networks”, ICS 2002.

[LDP04] E. Leontiadis, V. Dimakopoulos, E. Pitoura, “Cache updates in a peer-to-peer network of mobile agents”, IEEE P2P Conference, 2004.

[LHHSS04] Boon Thau Loo, Joseph M. Hellerstein, Ryan Huebsch, Scott Shenker, Ion Stoica, “Enhancing P2P File-Sharing with an Internet-Scale Query Processor”, VLDB 2004: pp. 432-443.

[lindellpinkas] Y.Lindell and B.Pinkas, “Privacy preserving data mining”, Advances in Cryptology - CRYPTO 2000, volume 1880 of Lecture Notes in Computer Science, pp. 36-46, 2000.

[lindellpinkas2] Y.Lindell and B.Pinkas, “Privacy preserving data mining”, Journal of Cryptology, 15: pp. 177-206, 2002.

[LLR94] Linial, N., London, E., and Rabinovich, Y., “The geometry of graphs and some of its algorithmic applications”, 35th Annual Symposium on Foundations of Computer Science (Santa Fe, New Mexico, 20-22 Nov. 1994), IEEE, pp. 577-591.

[LMP04] Lotker, Z., Martinezde Albeniz, M., and Perennes, S. Range-free ranking in sensors networks and its applications to localization. In ADHOC-NOW 2004 (2004), pp.158-171.

[LR04] P. Leone, J. Rolim, “Towards a Dynamical Model for Wireless Sensor Network”, 1st International Workshop, ALGOSENSOR 2004, Turku, Finland, July 2004.

[LuMe01] Lumetta, S., and Medard, M. Towards a deeper understanding of link restoration algorithms for mesh networks. In IEEE Infocom (2001).

[maftia] Malicious and Accidental-Fault Tolerance for Internet Applications – IST Research Project IST 1999-11583 2000 - 2003}, http://www.maftia.org, 2003.

[Matousek02] Matousek, J., “Lectures on Discrete Geometry”, Springer, 2002.

[MelPenProWatWid04] G.Melideo, P.Penna, G.Proietti, R.Wattenhofer, and P.Widmayer, "Truthful mechanisms for generalized utilitarian problems", IFIP International Conference on Theoretical Computer Science (IFIP-TCS), 2004.

[MGS98] S. Muthukrishnan, B. Ghosh, and M.H. Schultz. “First and second order diffusive methods for rapid, coarse, distributed load balancing”. Theory of Computing Systems, 31, pp. 331-354, 1998.

[MMS01] Mohan, G., Murthy SivaRam, C., and Somani, A. Efficient algorithms for routing dependable connections in wdm optical networks. IEEE/ACM Transactions on Networking (2001).

[MoNa02] Modiano, E., and Narula-Tam, A. Survivable lightpath routing: a new approach to the design of WDM-based networks. IEEE Journal of Selected Areas in Communication 20, 4 (May 2002).

[MorselliP2Pecon] R. Morselli, J. Katz, and B. Bhattacharjee, "A game-theoretic framework for analyzing trust- inference protocol", Conference on Economics of P2P, 2004.

[MOSZ+03] Ou, C., Zhu, K., Zang, H., Sahasrabuddhe, L., and Mukherjee, B. Traffic grooming for survivable wdm networks - shared protection. IEEE Journal of Selected Areas in Communication: Optical Communications and Networking Series 21, 9 (November 2003).

[Moti_ramp_Cripto04] M.Yung, A.Kiayias, and S.Xu, "Mining data from anonymous transactions", Rump Session CRYPTO 2004, 2004.

[MP02] M. Mihail and C.H. Papadimitriou. "On the eigenvalue power law", RANDOM'02, 2002.

[MR96] R. Meester, R. Roy, "Continuum Percolation", Cambridge University Press, 1996.

[MT03] Monderer, D., and Tennenholtz, M., "k-implementation", ACM EC (2003), pp. 19-28.

[MuZh02] Zhu, K., and Mukherjee, B. Traffic grooming in an optical WDM mesh network. IEEE Journal on Selected Areas in Communications 20, 1 (January 2002), 122-133.

[NGE00] A. Ephremides, G.D. Nguyen, J.E. Wieselthier, "On the Construction of Energy-Efficient Broadcast and Multicast Trees in Wireless Networks", pp. 585-594, 2000.

[NisRon99] N.Nisan and A.Ronen, "Algorithmic Mechanism Design", Games and Economic Behavior, 35: pp. 166-196, 2001. Extended abstract in STOC'99.

[NR00] N. Nisan, A. Ronen, "Computationally Feasible VCG Mechanisms", 2nd ACM Conference on Electronic Commerce (EC'00), 2000.

[NR99] N. Nisan, A. Ronen, "Algorithmic Mechanism Design", 31st ACM STOC, 1999.

[NRS04] S. Nikoletseas, C. Raptopoulos and P. Spirakis, "The Existence and Efficient Construction of Large Independent Sets in General Random Intersection Graphs", 31st International Colloquium on Automata, Languages and Programming (ICALP), Lecture Notes in Computer Science (Springer Verlag), 2004.

[O96] P. Olofsson, "Branching Processes with Local Dependencies", The Annals of Applied Probability, Vol. 6, No. 1, pp. 238-268, Feb. 1996.

[OR94] Osborne, M.J., and Rubinstein, A., "A course in game theory", MIT Press, 1994.

[P03] M. Penrose, "Random Geometric Graphs", Oxford University Press, Oxford Studies in Probability 5, 2003.

[PageRank] Page, L., "PageRank: Bringing order to the web", Stanford Digital Libraries Working Paper 1997-0072, Stanford University, 1997.

[Pap01] C.H. Papadimitriou. "Algorithms, games, and the internet", 33rd Annual ACM Symposium on Theory of Computing (STOC'01), pp. 749-753, 2001.

[PAPSV03] Evaggelia Pitoura, Serge Abiteboul, Dieter Pfoser, George Samaras, Michalis Vazirgiannis, " DBGlobe: a service-oriented P2P system for global computing", SIGMOD Record 32(3): pp. 77-82 (2003).

[Par04] Park, J. Resilience in gmpls path management: model and mechanism. IEEE Communication Magazine 42, 7 (July 2004), 128-135.

[Pel00] Peleg, D. Distributed Computing: A Locality-Sensitive Approach. SIAM, Philadelphia, PA, 2000.

[PenVenSIROCCO04] P.Penna and C.Ventre, "Sharing the cost of multicast transmissions in wireless networks", International Colloquium on Structural Information and Communication Complexity (SIROCCO), LNCS, 2004.

[PenVenWAOA04] P.Penna and C.Ventre, "More powerful and simpler cost-sharing methods", 2nd Workshop on Approximation and Online Algorithms, LNCS, 2004.

[PfitSchuntWaid:2000] B. Pfitzmann, M. Schunter, and M. Waidner, “Cryptographic Security of Reactive Systems”, Workshop on Secure Architectures and Information Flow, volume 32 of Electronic Notes in Theoretical Computer Science, 2000.

[PfitWaid:CCS2000] B. Pfitzmann and M. Waidner, “Composition and integrity preservation of secure reactive systems”, 7th ACM Conference on Computer and Communications Security}, pages 245-254, 2000.

[PR97] E. Pagani, G. Rossi, “Reliable Broadcast in Mobile Multihop Packet Networks”, 3rd ACM-IEEE MOBICOM, pp. 34-42, 1997.

[PRU01] Pandurangan, G., Raghavan, P., and Upfal, E. Building low-diameter p2p networks. In 42th IEEE Symp. on Foundations of Computer Science, FOCS-01 (2001).

[PST91] Plotkin, S.A., Shmoys, D.B., and Tardos, E. Fast approximation algorithms for fractional packing and covering problems. In IEEE Symposium on Foundations of Computer Science (1991), pp.495-504. Final version in Math of Oper. Res.

[PWS+00] Fernando Pedone, Matthias Wiesmann, Andrı Schiper, Bettina Kemme, Gustavo Alonso, “Understanding Replication in Databases and Distributed Systems”, ICDCS 2000: pp. 464-474.

[R00] A. Ronen, “Algorithms for Rational Agents”, 27th SOFSEM, 2000.

[RabinOT] M.O. Rabin, “How to exchange secrets by oblivious transfer”, Technical Report, Aiken Computation Laboratory, Harvard University, 1981.

[RB03] M. Roussopoulos and M. Baker, “CUP: Controlled Update Propagation in Peer-to-Peer Networks”, USENIX 2003.

[Riv03] Rivano, H. “Algorithmique et tılıcommunications : Coloration et multiflot approchıs et applications aux rıseaux d'infrastructure”. PhD thesis, Universitı de Nice-Sophia Antipolis, November 2003.

[RM51] H. Robbins, S. Monro, “A Stochastic Approximation Method”, Ann. Math. Statistics 22, pp. 400-407, 1951.

[rolemining1] H.Roeckle, G.Schimpf, and R.Weidinger, “Process-oriented approach for role-finding to implement role-based security administration in a large industrial organization”, fifth ACM workshop on Role-based access control, pp. 103-110. ACM Press, 2000.

[rolemining2] M.Kuhlmann, D.Shohat, and G.Schimpf, “Role mining - revealing business roles for security administration using data mining technology”, eighth ACM symposium on Access control models and technologies, pp. 179-186. ACM Press, 2003.

[RoXi04] Xin, Y., and Rouskas, G. A study of path protection in large-scale optical networks. Photonic Network Communications 7, 4 (May 2004).

[RP02] M.Rennhard and B.Plattner, "Introducing MorphMix: Peer-to-peer based anonymous internet usage with collusion detection", Workshop on Privacy in the Electronic Society in Association with 9th ACM Conference on Computer and Communications Security, pp. 91-102. ACM, 2002.

[RR98] M.Reiter and A.Rubin, "Crowds: anonymity for Web transactions", ACM Transactions on Information and System Security, 1(1): pp. 66-92, November 1998.

[RS01] Michael Rabinovich, Oliver Spatscheck, "Web Caching and Replication", Addison-Wesley 2001.

[RSG98] M.G. Reed, P.F. Syverson, and D.M. Goldschlag, "Anonymous connections and onion routing", IEEE Journal on Selected Areas in Communications, 16(4): pp. 482-494, 1998.

[RT02] T. Roughgarden and E. Tardos. "How bad is selfish routing?" Journal of the ACM, 49(2):236-259, 2002.

[RT99] E.M. Royer and C.K. Toh, "A review of current protocols for ad hoc mobile wireless networks", IEEE Personal Comunications, April 1999, pp. 46-55.

[SaSp02] Spencer, J., and Sacks, L., "Modelling IP network topologies by emulating network development processes", IEEE Softcom 2002, Split, Croatia (2002).

[Sch02] Scheideler, C., "Models and techniques for communication in dynamic networks", Symposium on Theoretical Aspects of Computer Science (2002), pp.27-49.

[SL00] B.N. Levine and C.Shields, "Hordes - a multicast based protocol for anonymity", Journal of Computer Security, 10(3): pp. 213-240, 2002.

[spirakis] P.Spirakis and B.Tampakas, "Distributed pursuit-evasion: some aspects of privacy and security in distributed computing", 13th annual ACM symposium on Principles of distributed computing (PODC 94), p. 403. ACM Press, 1994.

[spirakis2] P.Spirakis, B.Tampakas, and H.Antonopoulou, "Distributed protocols against mobile eavesdroppers", 9th International Workshop on Distributed Algorithms, pp. 160-167. Springer-Verlag, 1995.

[SPKI] C.Ellison, B.Frantz, B.Lampson, R.Rivest, B.Thomas, and T.Ylonen, "SPKI Certificate Theory", Network Working Group, RFC 2693, 1999.

[ST85] Sleator, D.D., and Tarjan, R.E., "Self-adjusting binary search trees", J. ACM 32, 3 (1985), pp. 652-686.

[TA03] Peter Triantafillou, Ioannis Aekaterinidis: ProxyTeller, "A Proxy Placement Tool for Content Delivery under Performance Constraints", WISE 2003: pp. 62-71.

[virus] R.Ostrovsky and M.Yung, "How to withstand mobile virus attacks", 10th Annual Symposium on Principles of Distributed Computing (PODC 91), pp. 51-59, 1991.

[X509] R.Housley, W.Ford, W.Polk, and D.Solo, "Internet X509 public key infrastructure: Certificate and CRL profile", Network Working Group, RFC 2459, 1999.

[Y+03] K. Yotov et al., "A Comparison of Empirical and Model-Driven Optimization", PLDI, 2003, pp. 63-76.

[Y47] A. M. Yaglom, Certain "Limit Theorems of the Theory of Branching Random Processes", Doklady 56, pp. 795-798, 1947.

[Yao1] A.C. Yao, "Theory and application of trapdoor functions", 23rd IEEE Symposium on Foundation of Computer Science, pp. 80-91, 1982.

[Yao2] A.C. Yao, “How to generate and exchange secrets”, 27th IEEE Symposium on Foundation of Computer Science, pp. 162-167, 1986.