
TRUSTNET

Scalable, Trusted Information Sharing for Ambient Applications

A proposal for an Institutional Strategic Project 2005-2008

Werner Van Belle*, Lars. K. Vognild*, Tage Stabell-Kulø**, Theo D'Hondt***

* Norut Informasjonsteknologi AS, Norway

** Department of Computer Science, University of Tromsø, Norway

*** Programming Technology Lab, Brussels University, Belgium

1.

Peer-to-peer systems are systems in which the gross of the data is directly communicated between hosts, without going through centrally placed servers. Typically, this kind of networks organize themselves automatically, thereby neglecting the actual network topology. This makes peer-to-peer systems a very attractive means to support ad-hoc networks, such as interconnected embedded devices. However, the peer-to-peer paradigm still poses a number of research problems. Among them scalability, finding useful information, security and the programming of applications in such volatile environments. The objective is to address these problems by developing a *scalable and secure peer-to-peer information sharing platform together with a suitable programming model.*

The consortium consists of three partners. Norut IT, which is an applied research institute, the department of computer science at the University of Tromsø (UiTø), which pioneered the now wide spread distributed agent paradigm in 1996, and the Programming Technology Lab at Brussels University, which is doing research in programming language engineering for distributed systems. By coupling the experience of these partners we ensure a strong, well balanced consortium.

2. Motivation

Since 1994 on, the Internet has been growing quickly. The common software paradigm among Internet applications is the client server architecture. This architecture divides applications in two parts: clients and servers. Clients are information consumers, which are sporadically connected to the Internet, while servers are information providers that are constantly connected to the Internet. This paradigm became widespread because it offers (among other properties) a central point of administration and managing information content. However, the advent of multimedia posed a number of problems. First, centrally placed servers had difficulties in providing the required performance. E.g, fetching multimedia streams overseas is much more costly than getting the same information already available in the neighbourhood. Second, private end-users are very little concerned about copyrights, as long as they get easy access to the information they want.

In 2000 programs such as Gnutella and Napster tried to solve these problems. They help end-users in providing and fetching information from each other's harddisk, without needing to contact a central server that stores the actual information. By doing so, it became possible to boost the availability of

multimedia and to avoid prosecution. This approach has prompted a whole new area of research, called peer-to-peer networks.

In a typical peer-to-peer network, the gross of the information is communicated immediately between hosts. Aside from having no clear centralized organization, they also offer a very attractive advantage for possible other application domains: *self-organization*. Because nodes are not always online, they need some strategy to deal with discontinuities. This strategy often involves the creation of a virtual network, which transparently offers access to resources and information within the vicinity. Another advantage of peer-to-peer networks is their *robustness*. Because they need ways to deal with node discontinuities they often *replicate* information between nodes. When one node is unavailable, another one will be able to offer similar information. Given these two distinctive attractive features we can see how such networks might provide solutions within other areas [1][2].

One area of interest are *ad-hoc networks*. In environments where there are no pre-existing administrated network infrastructure in place, except for the computers and the network devices themselves, it is necessary to create a virtual network that provides connectivity between the different computers [3]. The self-organizing nature of peer-to-peer networks makes them very attractive for this purpose, especially in the context of embedded devices. With the current trend to make embedded systems mobile and equip them with networking capabilities, self organization is becoming a major challenge. This kind of technology is often referred to as ambient intelligent applications. However, because these devices are mobile, they typically have no network topology available. Peer-to-peer systems can provide these naturally.

A second area of interest are *collaborative work environments*. This kind of environment relies on computer support to communicate documents, share information, make decisions over long distances and so on. Because they work over long distances it is often difficult to administrate them centrally, hence the application of peer-to-peer networks might be useful. It could provide efficient local information transfer, while providing caching of remote documents. Also, the replication properties inherently offers automatic backup possibilities.

Peer-to-peer networks hold the promise of robust information sharing and self organizing networks. There are however several open issues regarding peer-to-peer networks.

3. Challenges

Technology

With respect to the application area of mobile ad-hoc networks, the key challenge is to integrate different branches of technology. The fast growth of the market for wireless services and interconnected embedded systems has lead to a differentiation in both hardware and protocols to provide network

connectivity. Because not every device will have the availability of the same hardware, it becomes necessary to integrate these different protocols under one suitable API. This requires us to evaluate traditional technologies to implement messaging systems for broadcasting, multicasting, *location tracking*, *routing* and *network awareness*. With respect to overlay networks (networks with high connectivity, running TCP/IP), there occurs other technological challenges. Due to the heavily restricted Internet connectivity of most consumer machines it is often inconvenient or impossible to run peer-to-peer systems on them. Firewalls, address translation schemes and similar technologies pose a challenge to the peer-to-peer paradigm. There exist some workarounds, such as connection-notification servers (as implemented within the later releases of ICQ) and new technologies such as IPv6 might address a number of these problems. However, a technological integration is of utmost importance for peer-to-peer networks.

Information Scalability

Peer-to-peer systems are known to have problems with respect to scalability of searching [4]. In simple peer-to-peer networks the search strategy is simply to send out the request towards a number of known nodes. The receiving nodes check whether an incoming search request matches. If it does not, the search request is forwarded toward other nodes within the network. This kind of forwarding quickly replicates one search request into several millions, which has the tendency to cripple the involved networks. To tackle these problems modern peer-to-peer systems make use of ingenious techniques. These range from centrally stored directories, user based directories, automatic directory node election to semantic query routing. However, how these techniques can be applied within mobile ad-hoc peer-to-peer systems still remains an important challenge. Firstly, embedded systems have a limited memory footprint and thus are impractical for storing large directories. Secondly, most peer-to-peer systems do not take node mobility into account. When a node migrates from one environment to another it might rupture the expected virtual network structure. This makes the use of standard peer-to-peer search algorithms inappropriate for interconnected embedded systems. By exploiting the possibilities of mobile nodes to carry information, certain advantages might be offered.

Thrustworthiness

In a system where information can be supplied by any node, it might be difficult to know how much of this information can be trusted. However, because peer-to-peer systems originated in free information sharing communities, this problem has been very little investigated. Nevertheless, to be commercially viable, it is necessary to introduce some form of thrustworthiness. The challenges posed here are how to correctly *identify* people and nodes, how to install a suitable access-rights system that enables information producers to restrict access to their data and that allows information consumers to verify the integrity of the information.

Application Program Design

The last challenge lies in a mismatch between the uncertainty of peer-to-peer

networks and the strict certainty programming languages offer. The concept of ambient intelligent systems requires a radical rethinking of the underlying assumption of systems architecture. Because information may –or may not– be available and the information itself might be more –or less– accurate, the development of applications is seriously complicated. Applications now need to keep track of multiple incomplete information flows at the same time. Given the fact that the design of session management strategies within distributed systems (concurrency management, transactions) is still considered to be one of the most time consuming processes, it is not surprising to see that the removal of a central administrative entity (the server) and the removal of certainty even further complicates the software engineering process.

4. Objectives

We recognize that the development of trustworthy and scalable information sharing applications is difficult in the domain of mobile ad-hoc, wireless and volatile overlay networks. Taking these challenges into account, our objective is to create a novel peer-to-peer information sharing platform (the TRUSTNET platform). The platform itself will consist of a *runtime architecture* and a *programming paradigm*.

The sub-objectives of the runtime architecture are to

- develop *scalable information propagation algorithms*, integrating connectivity techniques from mobile ad-hoc networks and search algorithms from overlay networks
- develop a *decentralized security model* providing information tracking and protection using trust relationships

The sub-objectives of the programming paradigm are to explore

- how programs can be written that work well within unreliable networks
- how programs can be written that can handle incomplete and incoherent information
- how such programs can be written easily without needing to spend too much work on non-functional aspects of the problem at hand

The TRUSTNET platform will be experimentally validated.

5. Cases

In order to create a TRUSTNET platform and language that is usable in a wide application range, we have carefully selected 4 *different* cases. These cases will all run on the *same platform*. This approach avoids the creation of an architecture that is well tailored to one application but fails to capture important aspects of other possible applications.

5.1 Case 1: Disconnected Meeting Planning

This scenario assumes that a number of people have their agenda stored in their local PDA. To schedule a meeting, the PDA will try to contact agendas on other PDA's. When there is a usable internet connection, then this will be used to contact other agendas. In situations where there is no suitable or

usable network infrastructure in place (such as scheduling between a number of participants on a conference), then these PDA's should be able to set up their own ad-hoc network, thereby taking into account the cached schedules of absent co-workers and actually *carrying* information around. This scenario addresses a *mobility* and *connectivity* problems.

5.2 Case 2: Sensory Reading Aggregation

In this scenario we consider small mobile wireless sensory devices and will experiment with the possibility to aggregate information. Specifically, we assume that cars will be instrumented with such devices. Communication between these devices might offer important traffic information. However, if a car breaks down on a highway, it is only of vital importance for neighbouring cars, but far less important for the rest of the world. Our experience from distributed geodata systems shows that the importance of details diminishes as the distance to their point of origin increases. Similarly it is possible to create aggregated reports on road conditions, weather conditions and traffic conditions based on a large amount of specific sensory readings. By using an appropriate peer-to-peer platform these devices will be able to manage an appropriate information flow. This scenario addresses *scalable information propagation*.

5.3 Case 3: Collaborative Computer Supported Work

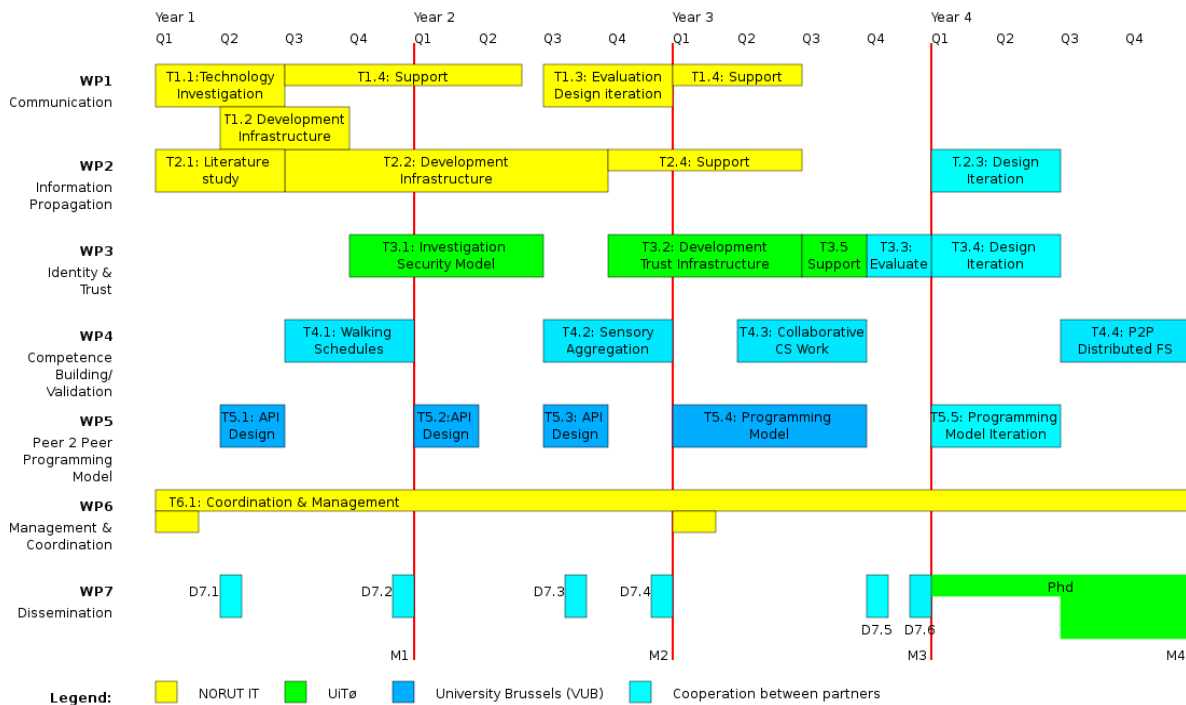
The third case focuses on the trust model of peer-to-peer systems, while retaining the possibility to search for information. This scenario assumes that ad-hoc working groups make use of a collaborative computer supported work environment in which all kinds of data are integrated. E.g, let us assume that the results of biological experiments are stored within such a system. These results should only be fully accessible by partners within the project, partly accessible to external people working with similar experiments and inaccessible to the outside world. This should make it possible to look for groups that are doing similar experiments but without disclosing proprietary information. This scenario addresses security aspects in *identification, authentication and search access*.

5.4 Case 4: Distributed File Storage with Information Tracking

The final case concerns the tracking of information. This scenario investigates the possibility to use the TRUSTNET platform to create a distributed shared filesystem. This filesystem should be able to offer *security, searching* capabilities and the possibility to *track information spread*. E.g, it should be possible to ask questions such as: 'who has (modified) copies of this file', or, the other way around, 'where does this file come from and what guarantees can be given about its content ?' This scenario validates *integration* and addresses *the tracking of information*.

6. Work Package Description

The following work packages and tasks are defined and linked according to the GANTT chart below.



In order to develop the TRUSTNET platform, we will create a multi-layered software model. The first layer (WP1) will provide the basic communication infrastructure. The second layer (WP2) ensures the necessary information sharing operations. The third layer (WP3) will provide security operations. During the entire project, a typical software engineering cycle will be applied. Every layer will be developed (T1.2, T2.2, T3.2), used (T4.1, T4.2, T4.3), maintained (T1.4, T2.4, T3.4) and then redesigned (T1.3, T2.3, T3.4) in order to meet new challenges. The usability of the API's will be increased by dividing software implementation and API specification over different work packages, which are respectively WP1-3 and WP5.

6.1 Work Package 1: Communication & Connectivity [26.5MM¹]

The communication layer ensures connectivity between different nodes. It makes sure that nodes can be uniquely addressed and that they can communicate between them. A virtual network might need to be set up that organizes itself around the actual network topology (or in the context of mobile ad-hoc networks, the lack of appropriate network topology).

Related work: Most peer-to-peer networks today focus largely on either overlay networks or mobile ad-hoc infrastructures. We believe that these two systems will not be mutually exclusive. The first branch, based on overlay networks, concerns itself with the problem of routing data over an existing network infrastructure [5]. They rely on a relatively stable network topology,

¹MM stands for manmonth.

which makes them impractical for mobile ad-hoc networks [6]. The second branch of peer-to-peer systems focuses on mobile ad-hoc networks [7][8][9]. They are largely based on network level multi-hop routing [3][10], and often fail to include appropriate information searching possibilities. These two systems are difficult to integrate because mobile ad-hoc networks have difficulties providing a global unique and location transparent address. This problem might be approached by using advanced techniques to provide location transparency [11]. The most common technique nowadays to find devices in the neighbourhood are Sun's Jini and Apple's RendezVous. However, both techniques does not scale well to provide a fully connected virtual network.

Description: In this work package we will develop the TRUSTNET communication & connectivity layer. This layer will provide a virtual self organizing network offering connectivity, location transparency and location awareness. The communication layer should work in both ad-hoc and overlay networks. Further, because embedded devices often have a limited bandwidth and memory-footprint it is important to remain lightweight.

Goals:

- *integration* of both ad-hoc and overlay strategies
- *lightweight* on required code base and network administration traffic
- *awareness* of the network neighbourhood and the relative positioning within the network
- *self organizing* as fast as possible

Task1.1 [6MM]: Investigation modern technology such as Bluetooth, Wireless Ethernet (IEEE 802.11) and others

Task1.2 [6MM]: Development of the communication infrastructure

Task1.3 [6MM]: At the end of the second year the communication infrastructure will be redesigned to address new requirements. The redesign will be based on the evaluation of the experiment performed in Task4.1, the initial development of the information propagation layer (Task2.2) and preliminary requirements of the security model (Task3.1)

Task1.4 [8.5MM]: Support and maintenance of the communication protocol code

6.2 Work Package 2: Information propagation [31.5MM]

The information layer is the where information is managed. Its functionality is to find and manage useful data. E.g, finding all movies with actor X.

Related Work: There exist a large body of techniques to search within overlay networks [4][12]. Centrally to these is the possibility to locate information based on its content. This concept is termed *content addressable networks* [13] and makes use of *distributed hash tables* (DHT). Known projects are IRIS[14], CAN, Chord[15], Tapestry and Pastry [5], which mainly differ with respect to key handling and routing. When *pulling* information from these networks, one search-request quickly multicasts to multiple nodes. The problem with this approach is its scalability. As the network grows larger, the search-request might very well saturate the entire bandwidth [16]. To avoid this, modern approaches additionally *push* meta-information towards

groups of interest. When enough nodes send out similar queries, a semantic query routing algorithm [17] can deduce and create implicit interest communities. However, all these techniques require relatively stable network topologies, which makes them impractical to use in mobile ad-hoc networks.

Description: In this work package we will apply advanced push style strategies, in which communities of interest are detected automatically and used to distribute meta-information. We will enhance this model by introducing two new factors that are expected to optimize the information flow greatly. Both are based on compression techniques. The first technique will free bandwidth by transferring content changes instead of the actual content itself. For the directory information itself this should certainly offer advantages. The second compression technique will be content based compression. Depending on the type of information, multiple smaller pieces of information might be compressed into larger but less detailed descriptions. An example is to combine multiple different sensory readings into the general picture they make. In order to make this possible the meta-information format must be investigated in detail.

Goals:

- *scalable* search infrastructure, using a *pull* and *push* style
- content based *compression* and *aggregation*
- implicit *group detection* and formation
- *explicit group annotation*

Task2.1 [6MM]: Investigate current state of the art peer-to-peer search networks

Task2.2 [15MM]: Development of the information propagation layer

Task2.3 [6MM]: Redesign of the information propagation layer in order to meet problems recognized in Task3.3

Task2.4 [4.5MM]: Support and maintenance of the information propagation layer

6.3 Work Package 3: Thrustworthiness [30MM]

Thrustworthiness is a multi-faceted problem, especially in the field of sharing communities and peer-to-peer networks. Typically this kind of networks share information freely. Whether this is a good characteristic is still open to debate. Certainly, there are situations in which anonymity is wanted [18], however to be commercially viable the requirement of anonymity should be reconsidered. For instance,

- In collaborative work environments it might be necessary to know the identity of people on the network such that only those people can access the requested information
- It should be possible to verify the integrity of data. An example is a software package from a software supplier. A node should be able to cryptographically verify (or even decode) the package, even though it was handed over by an unknown, untrusted peer
- Nodes can also be sharing private content, making privacy an important issue. We must ensure that access to data, services and resources can be restricted towards an explicitly managed trust-ring
- In other application areas it might not only be necessary to restrict access

to data but also to be able track its spread throughout the network. E.g, an information provider might be interested in finding out who distributes the data and who actually uses the data. Similarly, the information consumer might want to verify the integrity of incoming data by tracking where it comes from. Platform support for this kind of information is vital before a commercial business model can be created.

Related work: There is little work on security issues for peer-to-peer networks[19]. Most of the useful related work in distributed information sharing can be found in the area of distributed filesystems. In this area, the Echo filesystem[20] of Taos[21] relies on a system wide computing base. The Bayou[22], Coda[23] and Ficus[24] distributed file systems target mobile clients in the system, and use replication while weakening consistency requirements. Survivable storage systems research, like Pasis, shows how *decentralisation*, history logging and threshold schemes can increase availability and integrity under the assumption that storage servers and their clients can be compromised [25]. Another *decentralised* approach to trust relationships can be found in the RFC2693 [26].

Description: The trust model will a) be able to manage different trust relations and b) can control and keep track of the spread of information. It will be based on an existing filesystem, called Pesto[27]. Pesto already offers a decentralized security model based on trust relationships. This makes it a very good candidate to use within peer-to-peer networks. Further, it keeps track of information changes, which integrates nicely with the research conducted in the information work package. New research will be conducted with respect to peculiarities of *searching* and *information tracking*.

Task3.1 [9MM]: Investigation of new security factors within the setting of trusted peer-to-peer networks

Task3.2 [9MM]: Development of the trust layer

Task3.3 [3MM]: At the end of the third year, there will be an evaluation of the entire TRUSTNET platform (the 3 layers and its programming model)

Task3.4 [6MM]: Realisation of integrated design, requires involvement of all partners.

Task3.5 [3MM]: Support and maintenance of the trust codebase

6.4 Work Package 4: Experimental Validation

Description: As mentioned in section 4, by building a new technological platform and validating it using 4 cases, we ensure to build competence within the area of peer-to-peer systems. Each case has been carefully chosen in order to address possible future relevant applications. Below we discuss which aspects of the TRUSTNET platform are validated by the different cases.

Task4.1: Disconnected Meeting Planning. Verifies the communication infrastructure (task 1.2) in ad-hoc networks and overlay networks. This will validate

- usability of wireless communication protocols
- global unique and location transparent node identification

-
- location independent routing
 - network selection strategies when multiple networks are present
 - location awareness. E.g, when a schedule is printed it should be printed onto a local printer
 - neighbourhood awareness. E.g, which devices are in range, automatic notification
 - development of programs that must assume possible disconnected operations

Task4.2: Sensory Reading Aggregation. Validates the information infrastructure (Task 2.2). Further, this case validates:

- scalability of information searching and information propagation
- how peer-to-peer networks can make use of information compression
- the meta-information format that allows describing possible information aggregation
- how mobile nodes can carry information using a peer-to-peer platform
- information propagation in fast changing mobile ad-hoc networks. E.g, cars change network environments much faster than people walking with a PDA
- development of programs that automatically work together as one whole when running on different nodes

Task4.3: Collaborative Computer Supported Work. Validates:

- the possibility to securely identify partners in a peer-to-peer network
- the possibility to search for information in a secure way
- the possibility to specify the level of content-detail that can be sent out
- the possibility to specify trust relationships within a peer-to-peer model
- the possibility to do self-organization without breaking trust relationships
- development of programs that make use of the TRUSTNET paradigm as a new communication infrastructure

Task4.4: Secure distributed filesystem. The last case will make use of the fully redesigned architecture to create a distributed filesystem, that will keep track of information spread. The case validates:

- the possibility to track the origin information at the receiving side
- the possibility to track the spread of information at the sending side
- the integrated design of the TRUSTNET platform

6.5 Work Package 5: Programming Model [24MM]

Programming applications that will work in a peer-to-peer environment is difficult because the environment is very volatile (nodes come and go) and the information that can be obtained has very little quality guarantees. Another challenge of programming peer-to-peer systems is that very few programming models allow specification of algorithms that will communicate with itself. E.g, the writing of an algorithm that will propagate traffic jam information automatically into the right direction is easy to develop when all nodes are known. However, it is much more difficult to devise a local behaviour that will lead to the required global behaviour.

Related work: there are some languages aimed at programming in open distributed environments without making too many assumptions about the communication partners [28][29]. However, there seems to be little focus on creating a suitable programming model for ad-hoc peer-to-peer networks.

Creating a suitable programming model for such systems might require a radical rethinking of the underlying building blocks of applications. E.g, while remote procedure calls are allowed to be both synchronous and asynchronous, we might see that an information retrieval in peer-to-peer system might have *none* to *many* replies [30]. Other work introduces new operators such as the *eventually operator*, the CSP *alternate* statement, or *explicit session stacks* which can be managed by the application itself [31].

Description: In this work package we focus on providing a good programming model with usable abstractions for such environments and its application domains. These programming abstractions include a) the design of the basic API provided functionality of the underlying TRUSTNET platform; b) new programming primitives that can handle *incomplete* information better and c) the standard structure necessary to easily develop distributed peer-to-peer applications with.

Task5.1 [3MM]: Design of the communication layer API

Task5.2 [3MM]: Design of the information layer API

Task5.3 [3MM]: Design of the security layer API

Task5.4 [9MM]: Development of an integrated programming model for peer-to-peer networks

Task5.5 [6MM]: After validation of this model, we reiterate over the model again

6.6 WP6: Management & Coordination [24MM]

Research scientist Dr. Werner Van Belle will manage the project. The management will cover coordination of work between the participating organizations, management of budget, progress reporting, project meeting and public relation work. Regular project meetings (4 per year) will be arranged. The day to day management will be followed up by the institute leaders of the involved partners. The key personnel of the participating organizations is listed in the organization section **below**.

6.7 WP7: Dissemination [18MM]

NORUT IT's educational role in previous and current SIP's, conducted in co-operation with the Department of Computer Science, has proven to form a good basis for a successful Ph.D program. This project aims at delivering one Ph.D over 4 year. In order to do so, the project contains a number of papers that will be published in internationally refereed journals. These papers will also be used as deliverables within the project itself. Below we tentatively indicate the content of the different deliverables.

Del7.1 [1MM]: Survey of the current day technology in ad-hoc peer-to-peer networks

Del7.2 [1MM]: Design and validation of a peer-to-peer communication infrastructure for ad-hoc peer-to-peer networks

Del7.3 [1MM]: Algorithmic design of a scaleable peer-to-peer information sharing protocol

Del7.4 [1MM]: Information sharing within the TRUSTNET platform

Del7.5 [1MM]: A programming model for peer-to-peer networks, a case study

Del7.6 [1MM]: A secure, trust based, peer-to-peer file sharing infrastructure
PhD [7MM]: At the end of the 4 year period, this collection of papers should lead to a Ph.D

7. Organization

The program will be organized as cooperation between NORUT IT (Norway), IFI/UiTø (Norway) and PROG/VUB (Belgium). Below we will present the different partners. The project has four clearly distinguishable milestones.

Mil1: Calendar Case Demonstrator

Mil2: Sensory Case Demonstrator

Mil3: Collaborative Work Demonstrator

Mil4: Distributed Filesystem Demonstrator and PhD Delivery

7.1 NORUT IT [NO]

NORUT IT has since it was established in 1984 conducted research and development within the areas of remote sensing and applied ICT research. The ICT department has currently 14 researchers doing *net-centric* research in the fields of networked multimedia, mobile applications, applications for broadband networks, database technology, data modelling, WEB technology and distributed GIS systems. Some current and previous projects which are related to the proposed activity:

- GIN - Geographic Information Network (1999-2003) Norwegian Research Council funded SIP project with focus on development of distributed network solution for storage and retrieval of geodata (<http://www.itek.norut.no/gin/>). In GIN we developed an experimental distributed system for decentralized management of geodata, and experimented with services for utilizing such a distributed information system.
- Seamless Services (Research Council of Norway (PULS) & industry, 2001-2004): technology for dynamic distribution and retrieval of information for both mobile and wired end-terminals over heterogeneous networks (<http://www.seamless.itek.norut.no>). This project focuses on the construction and use of network information services for mobile users, such that they can seamlessly access different networks services, indifferent of the communication medium.

Applications of our ICT technology in experimental environmental & climate monitoring systems in EU fifth framework programs.

- OMEGA '99-'03: development of operational monitoring system for European glacial areas.
- EuroClim '01-'04: develop an advanced climate monitoring and predication system for Europe.
- EnviSnow '02-'05: snow parameters from EO to use in global climate study and hydrology.
- FloodMan '03-'06: development of methods for utilisation of EO data in flood forecasting.

NORUT IT has conducted several R&D projects for industry and service providers, and the trend is that the infrastructure they want their application

systems to run on is getting more and more complicated; involving loosely coupled systems with variously connectivity and characteristics. By performing a focused strategic activity, NORUT IT will be able to compete nationally and internationally on a broader area than before.

Key personnel:

Dr. Werner Van Belle is a research scientist at NORUT IT. His expertise lies in development of distributed systems for mobile multi agent systems. He obtained his doctoral degree in 2003 at Brussels University. It covers formal analysis of interoperability conflicts within the context of open distributed systems. He cooperated with the universities of Ghent, Leuven, Limburg and Antwerp; and companies (Siemens, Barco-graphics, Agfa Gevaert, Imec, Philips and Alcatel).

Lars K. Vognild is a senior research scientist. He received his M.Sc degree in computer science in 1990 from UiTø. From '93-'98 he worked at the University of Twente on distributed systems. His interests are in distributed and mobile services.

7.2 PROG/VUB [BE]

The Programming Technology Lab (PROG) at Brussels University (VUB) was founded in 1987 by Prof. Dr. Theo D'Hondt. the lab is active in several areas of research within the broad domain of software and programming language engineering. Research initially covered the study of sub-paradigms of object orientation such as prototype based inheritance, formal models for object oriented languages such as object calculi and semantics, programming language technology such as meta-object protocols and reflection. More recently research was directed towards object-oriented component-ware and distributed components.

Existing projects related to this proposal:

- *BORG* (1998 – 2003) is a mobile multi agent platform which aims at providing a platform for distributed and communicating agents. Its main strengths are its location transparent routing mechanism, strong mobility, synchronisation primitives, transparent serialisation and a computational model for wide area networks [28].
- *SEESCOA* (1999 – 2003): The goal of this project is to offer software engineering techniques for embedded systems using a component oriented approach. It is a project with a total budget of EUR 2'000'000, after classification as a top project by an international jury. The SEESCOA runtime architecture allows component based applications to be made distributed without any changes to the underlying components. The resulting applications run peer-to-peer. Error handling is done at a central place [32].
- *CODAMOS* (2003 – 2007): This project investigates context-driven adaptation of mobile services. It is a strategic basic research project aimed at solving a set of key challenges in the area of ambient intelligence, where personal devices will form an extension of each user's environment, running mobile services adapted to the user and his context. This project is

an extension of the SEESCOA project but now brings together 4 universities and 15 different companies [33].

Key Personnel:

Prof. Dr. Theo D'Hondt (Brussels University) is responsible for the Programming Technology Lab, a software and language engineering research lab that goes back 15 years. He helped establishing the computer science curriculum at the university. The curriculum is considered to be one of the most advanced in Belgium. He is currently involved in software co-evolution and reflective virtual machines for distributed environments.

Jessie Dedecker is a Ph.D student at PROG. He is funded by the Fund for Scientific Research (FWO). He investigates languages and programming models for pervasive computing. His interests are actor systems, pervasive computing, ubiquitous computing and ambient intelligence. *He will actively participate in TRUSTNET as the project closely relates to his doctoral work.*

7.3 CS/UiTø [NO]

The Department of Computer Science at the University of Tromsø currently has 11 faculty members. In addition, the department has 14 research fellows, 6 technical/administrative staff members, and several adjunct-professors. There are also a varying number of externally financed researchers, Ph.D candidates, M.Sc students and M.Eng students. The research of the department focuses on (i) the development and experimental testing of models for distributed operating systems and distributed applications, and on (ii) issues in distributed object-oriented (or component-based) systems, middleware and distributed database systems.

Existing projects and corresponding partners related to the proposed activities:

- The PENNE project (NFR Project Number 158596/431) will explore how one can design and build a distributed storage system that does not rely on public-key cryptography, but still offers revocation of certificates. In some senses we can view PENNE as an experimental platform for establishing trust relations that transcends administrative boundaries. The PENNE project is conducted in close cooperation with researchers from the University of Pisa, Italy.

Key personnel:

Associated Prof. Dr. Tage Stabell-Kulø: Responsible for work package 3 and a faculty member of the Department of Computer Science at the University of Tromsø. The research of Stabell-Kulø is focused on how to protect privacy in distributed systems, and at the same time, using the same mechanisms that are used to conceal private information to hide information about the system itself and thereby obtain better reliability.

Msc. Margrete Allern Brose: We have discussed the PhD position with one of our top students, and we are delighted that she wants to take part in this project if offered the possibility. She has a strong record, experience from

working abroad (Germany) and experience from working with industrial partners (Telenor).

7.4 International Cooperation

Norut IT and UiTø cooperate closely with European research groups involved in related research, specifically the University of Twente's IST FP5 project "Eyes - Energy Efficient Sensor Networks" [34] and the University of Lancaster's "OpenORB" program [35]. The department at UiTø cooperates with several international research groups. Among them the K.U. Leuven (Belgium), INRIA (France), Lancaster University (UK) and the University of Twente (NL). Brussels University has close connections with Ecole de Mines de Nantes/Inria (France).

8. Budget

	2005	2006	2007	2008	Total/Year
Funding FWO (Belgium)*	550	569	654	677	613
Own funding from UiTø	138	142	147	152	145
Norut Scientific Position	930	963	996	1031	980
Norut Scientific Position	930	963	996	1031	980
Norut 0.5 Scientific Position	465	481	498	516	490
Travel Costs Brussel <-> Tromsø	50	50	50	50	50
Direct Costs	60	60	60	60	60
University PhD	413	427	442	457	435
PhD stay abroad	183	0	0	0	46
Total Budget	3718	3655	3844	3974	3798
Requested Funding	3031	2943	3042	3145	3040

* These numbers tentatively indicate what a PhD candidate would cost in Norway.

The budget assumes that a PhD will be funded by forskningsrådet. We presented a very strong candidate for the PhD program. However, if the current policy of forskningsrådet is to prefer post doctoral positions then we can replace the 'stay abroad' and 'university PhD' with a postdoc position. We also have candidates available for such a position. The advising of the PhD and the candidate her/his fourth year will be funded by UiTø by giving the candidate a 25% teaching assignment.

The work done at the University in Belgium is funded by the Flemish Fund for Scientific Investigation (FWO). However, the travel-costs between Brussel and Tromsø are not covered by this fund, therefore they are included in the requested budget.

The work division is as follows. NORUT IT will spend 10MY² on work packages 1, 2, 4 and 6. UiTø will spend 3MY on work packages 3, 4, 5 and 7. The programming technology lab offers 2.3MY on work package 4, 5 and 7.

9. Ethics

There is a debate going on about the possibilities offered by peer-to-peer networks. They are especially infamous for providing free access to

²MY stands for manyear

copyrighted music, software and movies. Because there is no centrally placed organization that can be sued within a peer-to-peer network, it is very difficult and expensive, to track down the spread of copyrighted information. Nevertheless, the goal of this project is to create a platform that will lead to new applications, not to create a platform to share illegal data. In order to avoid this, TRUSTNET will have satisfactory identification requirements. This should discourage illegal usage of the platform.

10. Relevance

Scientific Relevance

The TRUSTNET platform has a number of novel aspects, that are currently not found in any peer-to-peer network.

- its focus on *node-mobility* and *security*
- its focus on a suitable *programming model* for *volatile environments*
- its *integration* between mobile ad-hoc networks and overlay networks

Industrial Relevance

The exercise of building experimental systems is a central part of our research approach. In general, the development of scalable information exchange technology has a wide variety of possible applications. Below we focus on possible application-areas that might benefit from specific advanced topics of the TRUSTNET platform, such as its focus on mobile ad-hoc networks and security.

1. *Defense and military applications* might benefit from avoiding central servers [10]
2. *Automotive applications* [36][37] such as propagation of road information. E.g, traffic jams, weather conditions, accidents, etc
3. *Semantic information propagation*. The self-organizing nature of peer-to-peer networks allow automatic information propagation towards interested parties. E.g: an agenda which is automatically populated with possible interesting events
5. *Computer Supported Collaborative Work*: as an ad-hoc solution to write documents in cooperation with a number of participants. Peer-to-peer systems offer a natural form of replication and virtual networking between people. The TRUSTNET platform offers security, and allows nodes to be mobile (Laptops for instance)
6. *Firefly agents*: The intelligent agent paradigm was originally conceived to bring people with similar interests together. This idea however lacked a supporting platform back in the 90's. TRUSTNET might offer a practical support to implement this
7. *Advertising Industry*. Determination of target-groups can be done through offering and monitoring information flows within a peer-to-peer network. Of course, if the information is considered to be inappropriate by too many people then routing of this information should be avoided by means of the trust-relation model and the self-organizing nature of the network
8. *Location-centric services* (sharing and retrieving information in the vicinity): finding and communicating with local devices such as scanners, printers, mail-hubs and others

-
9. *Environment mapping*. Through searching for the same information within different areas, it is possible to map out details with respect to this information. This can be used to map actual geographic information (geometrics), e.g, mapping out weather conditions, birth population, reindeer distribution
 10. *Ad-hoc communication in robotics*. E.g; camera motion; allow two cameras to be moved around and make sure all other cameras track the same focus-point
 11. *Secure distributed filesystems* have the potential to significantly improve the commercial business models of the next generation information sharing applications

European Relevance

The EC Information Society Technologies Advisory Group (ISTAG) recently put forward the notion of *Ambient Intelligence* (AmI) [38][39] as an important new research direction in the domain of software technology. As recognized in this proposal, the root cause is the continuously growing presence of information processing components in even the simplest appliances, domestic, industrial or infrastructural.

National Relevance

The proposed research falls well within the ICT strategy of the Norwegian Research Council (as published in [40][41]) in the areas of "*communication technology*" and specifically addresses the identified sub R&D challenges of "*new infrastructures for information and communication and their services*", "*platforms for distributed systems*", and "*design and programming of complex information systems and server applications*".

Relevance for Norut IT

The proposed research also addresses the core strategic objectives of Norut IT. The focus is defined as: services and systems for location based and geodata information, mobile data communication, interactive networked services and methods for environment monitoring (remote sensing). Furthermore, strengthening the cooperation with national and international research institutions is a key objective of Norut IT.

To summarize, the activities being proposed in this application are targeting the areas of research identified by the central R&D policy makers in both Norway and Europe, and that our research effort addresses central building stones for the future information society.

Bibliography

- [1] D.Liben-Nowell, H.Balakrishnan and D.Karger, Analysis of the Evolution of Peer-to-peer systems., 2002, ACM Conference on Principles of Distributed Computing (PODC), Monterey, CA
- [2] Metma Roussopolis, Mary Baker, David S.Rosenthal, T.J.Giulu, Petros Maniatis, Jeff Mogul, 2 P2P or Not 2 P2P, 2004, The 3th international workshop on peer-to-peer systms. http://www.hpl.hp.com/personal/Mary_Baker/publications/iptps2004.pdf
- [3] F.Baker, An Outsider's view of MANET, 2002, Internet Engineering Task Force Document

-
- [4] H. Balakrishnan, M.F.Kaashoek, D.Karger, R.Morris and I.Stoica, Looking up data in P2P Systems, 2003, Communications of the ACM, v46(2)
- [5] A.Rowstron and P.Druschel, Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems, 2001, IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), Heidelberg, Germany, pages 329-350
- [6] S.Dulman, T.Nieberg, J.Wu, P.Havinga, Trade-Off between Traffic Overhead and Reliability in Multipath Routing for Wireless Sensor Networks, 2003, WCNC Workshop, New Orleans, Louisiana, USA
- [7] N.Desai, V.Verma and S.Helal, Infrastructure for Peer-to-peer Applications in Ad-hoc networks, 2003, Technical Report, Computer and Information Science and Engineering Department Univesrity Florida, Gainesville, FL 32611, USA
- [8] J.Tian and K.Rothermel, Building Large Peer-to-peer Systems in Highly Mobile Ad-hoc Networks: New Challenges ?, 2002, Technical Report 2002/05, University of Stuttgart
- [9] G.Kortuem, Proem: A Peer-to-peer computing platform for Mobile Ad-hoc networks., 2001, Advanced Topic Workshop: Middleware for Mobile Computing, Heidelberg, Germany
- [10] Mobile Ad-Hoc Networking (Manet) working Group, http://protean.itd.nrl.navy.mil/manet_home.html
- [11] Werner Van Belle and Karsten Verelst, Location Transparent Routing in Mobile Multi-Agent Systems - Merging Namelookups with routing, 1999, Proceedings Workshop on Future Trends in Distributed Computing Systems 1999, volume 7, pg 207-212, IEEE Computer Society Press
- [12] D.G.Andersen, H.Balakrishnan, M.F.Kaarhoek, R.Morris, Resilient Overlay Networks, 2001, Proceedings of ACM SOSP, Banff, Canada
- [13] S.Ratnasamy, P.I.Francis, M.Handley, R.Karp and S.Schenker, A Scalable content addressable network, 2001, Proceedings of the ACM SIGCOMM
- [14] Project IRIS: Infrastructure for Resilient Internet Systems, <http://project-iris.net/>
- [15] I.Stoica, R.Morris, D.Karger, M.F.Kaashoek and H.Balakrishnan, Chord: A Scalable Peer-to-peer lookup service for internet applications, 2001, ACM SIGCOMM, San diego, CA, p149-160
- [16] C.Rohsr, Query Routiung for the Gnutella Network, 2002, http://www.limewire.com/developer/query_routing/keyword%20routing.htm
- [17] S.Joseph, NeuroGrid: Semantically Routing Queries in Peer-to-peer Networks, 2002, Proceedings of the International Workshop on Peer-to-peer Computing, Pisa, Italie
- [18] Freedman, Michael J et al., Introducing Tarzan, a peer-to-peer anonymizing network layer, 2002, Electronic Proceedings for the 1st international workshop on peer-to-peer systems (IPTPS'02)
- [19] Canny J., Some Techniques for Privacy in Ubicomp and context-aware applications, 2002, Privacy in Ubicomp'2002 Workshopi on socially-informed design of privacy-enhancing solutions in ubiquitous computing, Sweden
- [20] A.D.Birrell, A.Higsen, C.Jerian,T.Mann and G.Swart, The Echo Distributed Filesystem, 1993, DEC Systems Research Centre, Palo Alto, CA
- [21] E.Wobber, M.Abadi, M. Burrows and B.Lampson, Authentication in the Taos operating System, 1994, ACM Trans. Comp.Syst, 12(1):3-32
- [22] D.B.Terry, M.M.Theimer, K.Petersen, A.J.Demers, M.J.Spreitzer and C.Hauser, Managing Updates in a Weakly Connected Replicated Storage System, 1995, Proceedings of the 15th ACM Symposium on Operating Systems Principles, pages 172-183, Copper Mountain Resort, Colorado
- [23] J.J.Kistler and M.Satyanarayanan, Disconnected Operation in the Coda Filesystem, 1992, ACM Trans.Comput.Syst, 10(1):3-25
- [24] J.S.Heidemann, T.W.Page, R.G.Guy and G.J.Popek, Primarily Disconnected Operation: Experience with Ficus., 1992, Proceedings of the second workshop on the management of replicated data
- [25] J.J.Wylie, M.Bigrigg, J.D.Strunk, G.R.Ganger, H.Kiliccote and P.K.Khosla, Survivable Information Storage Systems, 2000, IEEE Computer, pages 61-68
- [26] C.M.Ellison, B.Frantz, B.Lampson, R.Rivest, B.Thomas and T.Ylon, SPKI Certificate Authority, 1999, RFC2693, The Internet Society
- [27] Feike W.Dillema and Tage Stabell-Kulø, Pesto Flavoured Security, 2004,
- [28] Werner Van Belle, Johan Fabry, Karsten Verelst and Theo D'Hondt, Experiences in Mobile

-
- Computing: The CBorg Mobile Multi-Agent System, 2001, Proceedings ToolsEE2001, volume 38, pages 1-9, IEEE Computer Soc. Press.
- [29] E, the Secure distributed pure-object platform and p2p scripting language for writing capability-based smart contracts., <http://www.erights.org/index.html>
- [30] Wouter van Oortmerssen, Concurrent Tree Space Transformation in the Aardappel Programming Language, 2000, PhD dissertation, Department of Electronics and Computer Science, University of Southampton, United Kingdom
- [31] Jessie Dedecker and Werner Van Belle, Actors for Ad-hoc wireless networks, 2004, Proceedings of the european Conference on Ubiquitous Computing, EUC04, Japan
- [32] , <http://www.cs.kuleuven.ac.be/cwis/research/distrinet/projects/Seescoa/>, ,
- [33] , <http://www.cs.kuleuven.ac.be/cwis/research/distrinet/projects/CoDAMoS/>, ,
- [34] Eyes - Energy Efficient Networks, IST-2001-34734, <http://eyes.eu.org/>
- [35] Reflection sin Open Distributed Systems, <http://www.comp.lancs.ac.uk/computing/research/mpg/reflection/>
- [36] D. Reichardt, M.Miglietta, L.Moretti, P.Morsink, W.Schulz, CarTalk2000 - Safe and Comfortable Driving based upon inter-vehicle coimmunciation, 2002, IEEE Intelligent Vehicle Symposium, Versailles
- [37] The FleetNet Project, <http://www.fleetnet.de/>
- [38] IST Advisory Group, A European strategy towards an Ambient Intelligent Environment: Software Technologies, embedded systems and distributed systems, 2002, ISBN 92-894-3837-1 (c) European Communities, ftp://ftp.cordis.lu/pub/ist/docs/istag_kk4402472encfull.pdf
- [39] IST Advisory Group, IST Workprogram 2003-2004, 2004, ftp://ftp.cordis.lu/pub/ist/docs/wp2003-04_final.pdf
- [40] F.Eliassen, ed., Strategi for IKT-forskningen i Norge, v1.1, 2001, Norges Forskningsråd
- [41] Kompetansefeltet IKT, <http://www.forskningsradet.no/fag/ie/felt/ikt.html>