COMMIT

PROJECTPLAN

WORKPACKAGES

DELIVERABLES

BUDGET

EXTREME WIRELESS DISTRIBUTED SYSTEMS FOR WELLBEING (P09)

Projectleader prof.dr. Maarten van Steen, Vrije Universiteit

1. Background

It is now recognized that sensor technology combined with pervasive computation and communication can significantly facilitate our way of living. The current plan builds on the goals and results of the Freeband-IShare project concerning large-scale decentralized ICT solutions for sharing information. Its focus is on devising (decentralized) solutions for large-scale wireless (sensor) systems, referred to as Extreme Wireless Distributed Systems. We concentrate on physical social networking (explained below) which we use for monitoring and managing crowds of people, with the aim of maintaining a feeling of safety and comfort for individuals.

The research in sensor networks has made a remarkable progress from a vision (Smart Dust, 2001)[1], through early prototypes (Mica2 motes, 2002)[2], to an ever-increasing set of deployments (2003 – today). However, operating a medium-scale deployment, with 10s-100s of nodes, still involves a lot of skilled manual labor and experience is mainly limited to static networks forwarding raw data to a sink (gateway node)[3].

Our goal is to achieve significant progress in scaling up solutions for wireless distributed sensor/actuator systems. These systems consist of a large collection of sensor/actuator devices that need to be shared by high-level applications. We intend to bridge the gap between the two. Some aspects of this gap are ease of installation, configuration and maintenance; ease of application development, hiding complexities of the network; and harnessing the behavior of a large network of relatively primitive nodes. Bridging this gap will advance innovation development for the industrial partners, but will also advance our knowledge concerning scalable wireless systems.

The consortium consists of two industrial partners (DevLab, a collaboration of SMEs, and Thales Research), and two academic partners (VU University and TU-Delft). Distributed systems and their applications form key elements of the academic partners' research programs, whereas the industrial partners seek to accomplish innovations in the aforementioned application area.

Very large wireless distributed systems are expected to enhance well-being. Key components of such systems are formed by miniature wireless devices that can interact with their environment and each other. Small form factors allow these devices to be attached to or

---

B.Warneke et al., "Smart Dust: Communicating with a Cubic-Millimeter Computer," *Computer*, vol. 34(1):44-51, Jan. 2001.
J.Hill and D.Culler , "Mica: A Wireless Platform for Deeply Embedded Networks," IEEE Micro., vol 22(6):12-24, Nov 2002.
Networks,", Wiley, 2007.

carried by people in a non-intrusive fashion. The small form factor is essential, yet necessarily restricts the communication capabilities to highly energy-aware wireless. We call such systems extreme wireless distributed systems. They are extreme in their size, their devices and their wireless interconnects.

2. Problem description

Given this context we to aim to find a (partial) solution to the following problem: How can we non-intrusively support sustained well-being through wireless distributed systems?

Well-being refers to the quality of life as experienced by people. In this project, we concentrate on providing a sense of well-being in crowded situations. We are interested in using wireless distributed systems to non-intrusively monitor and manage physical social communities, i.e. social communities in which the physical location of a person is an important factor in determining the position of that person in a community of people. In a nutshell, we aim at supporting the feeling that groups of people stay connected and remain feeling comfortable even while participating in a large crowd. This leads to the following general research question: how can we devise extreme wireless distributed systems for supporting end users behaving in a crowd?

Industry recognizes that extreme wireless distributed systems, and in particular large-scale networked sensor systems, have enormous potential. However, devising extreme wireless distributed systems is a technically and scientifically challenging problem, partly caused by the fact that we need to realize rich functionality on severely constrained devices, and such that solutions are fully decentralized. This requires joint efforts from different fields of expertise: sensor networking, distributed systems and algorithms, distributed signal processing, artificial intelligence, and human-systems interaction. There are many sub-problems to be tackled in order to reach this goal. For this project we have identified system dependability, feedback realization, and in-network information processing as key challenges. They are addressed (and explained in detail below) by the following research questions:

- How can we devise large, shared distributed systems such that end users will consider them dependable? We aim to address this human-systems interaction problem by developing one specific golden demo: an extreme wireless distributed system for physical social networking in crowded situations.
- How we can we provide proper *feedback* such that a person will indeed experience comfort in a highly crowded environment? As part of our golden demo, we will investigate which type of devices (screens, cell phones, etc.) should be used to guide people in a crowded situation.

- How can we devise a large-scale collaborative network of highly mobile, miniature, wireless devices collecting, disseminating, and aggregating information, as well as processing information to provide instant feedback to its users? This highly technical challenge regarding *in-network* information processing will deploy advanced gossip-based solutions.

Developing these applications is important from a social-economic perspective, but cannot be realized without novel innovative solutions that address extreme scalability: we simply need to rethink current approaches. In particular, Internet-style solutions involving a routing fabric and intelligent end nodes have shown not to work for the envisaged constellations.

3. Objectives

Project's goal

The goal of EWiDS is to develop the means for crowds of people to operate in a way that individuals in the crowd feel comfortable and in control of their own situation. In particular, we aim at facilitating individuals the means to stay in touch with specific other individuals in the same crowd, as well as providing a general sense of comfort and safety despite that the crowd may be extremely big. Our technological goals are developing energy-aware electronic tags (to be worn or embedded into devices such as cell phones), that jointly form a large wireless network. This network allows for measuring social interactions, but is also used for spreading and processing information. In addition, our goals encompass that data can be collected from the network to be processed offline, after which results can be used to assist in controlling a crowd.

Planning of all dimensions

One way of looking at P9 is that we will develop an instrument for measuring social interactions in a very large group of people, and that measurements are used to subsequently analyze and manage the crowd. As such, the project is unique in its focus on scale in using what Pentland (MIT) calls socio-metric badges and providing data for developing and validating high-level models of pedestrian behavior. Results will be disseminated through regular channels (conferences, journals, organization of workshops), and at the same time links to industry and other organizations will be explored, partly to consolidate results. Important end goals include:

1. A hardware platform of electronic badges for measuring social interactions in large groups of people.
2. A distributed software system running on top of the hardware of badges for optimally measuring and providing data on social interactions.

3. A distributed audio-processing subsystem that is used for providing additional data for sensing the state of a crowd.
4. An external feedback system for doing offline, real-time analysis of crowd behavior, capable of providing feedback.
5. A feedback subsystem, partly centralized, partly fully decentralized aimed at steering a crowd in order to maintain individuals in a "comfort zone."
6. Location-based decentralized solutions for telling individuals in the crowd about location-related services.

Results

By the end of the project, we will have developed a wireless network of electronic tags (badges), to be carried by participants of large events, combined with an external centralized feedback system for offline data analysis, processing, and feedback. The network is used for measuring the social structures emerging in a crowd of people and as such will be used to control the crowd as a whole, but also to provide individuals with specific information (e.g., on nearby friends or availability of local services). The system will contribute to the feeling of comfort and safety that individuals in a crowd will have, can be used to prevent dangerous situations, and will allow an individual to stay in better touch with his/her surroundings, including peers.

Deliverable Impact and Valorization

We plan to actively get external parties interested in the EWiDS system for crowd management and control, but have not yet drawn up a plan of work. This will be done at the start of the project. Anticipating that plan, we mention the following activities:

1. Substantiate our links with EIT ICTLabs, of which TUD/NIRICT is a core partner.
2. Explore the possibility of submitting one or more FET proposals on crowd management.
3. Contact various members of FHI through the network of DevLab. In addition to these activities, we plan to explore embedding the proposed technology in devices such as cell phones, or perhaps even replacing it with upcoming technology such as Bluetooth Low Energy (BLE).

Deliverable Dissemination

A concrete plan for dissemination will be drawn up at the beginning of the project. Notwithstanding those efforts, we intend to do the following to disseminate project results:

1. Various international scientific events are increasingly offering industrial) demo sessions: EWSN, Middleware, P2P. We will submit proposals for demos at these events.

2. We plan to organize a scientific workshop at the Lorentz Center, open to the public and notably inviting internationally renowned scientists as well as junior researchers to participate.
3. As soon as possible, we will attempt to submit articles to popular press, make the project known through university weekly's, produce flyers, and so on.
4. A web site will be set up, centered around the project, but also containing links to related projects and efforts, relevant technology, and so on. To draw attention, a small part of our budget can be used for Google Adwords and other attractor technology.

A special group, formed by project members, will be formed to continuously monitor dissemination efforts, and to seek new opportunities for making the project more visible.

International embedding

-

Deliverable Synergy

Both P8 and P9 focus on crowd management, yet in different ways. Research in P8 concentrates on providing sensors for rare event detection within a crowd, while P9 is focused on gathering information on the overall flow, movements, and interactions within a crowd. On the other hand, both projects do concentrate on monitoring crowds and seeing where and how crowd behavior deviates from what is to be expected. P8 and P9 will organize a workshop during the first year to further elaborate on these issues and to see to what extent we can make use of each others results.

4. Economic and social relevance

The project aims at supporting managing crowds such that people continue to feel comfortable. At the same time, crowd management includes manipulating parts of a crowd where needed to maintain a feeling of safety and comfort. We take the standpoint that technological innovations are needed to reach this goal. In particular, we believe that massively deployed sensor network technology (in the form of Extreme Wireless Distributed Systems) will form an important contribution. A bottleneck is currently formed by our lack of knowledge on how these EWiDSes can be successfully developed. To keep focus, we address this problem by an application-driven approach.

One (sub)application we concentrate on is physical social networking. We are currently witnessing an explosion of online social communities4. Along the same line, physical social

[4] See, for example, "Online Social Networks – Research Report," Communities and Local Government, Oct. 2008, http://www.communities.gov.uk/publications/communities/onlinesocialnetworks.

networking and urban sensing is emerging as a new phenomenon, characterized by the fact that the position and other attributes of people are important for the formation of social communities. These attributes can be “sensed”, discovered, and disseminated. An important goal of these networks is to extend and improve social interaction, which is crucial for well-being. The demonstrator for physical social networking is deployed under supervision of DevLab, by letting people use a small wireless device that exchanges information with other such devices in its vicinity. The devices form the basis for measuring social interactions, which in turn will be reported back to their respective users, embedded in a social game intended to encourage mingling. Most important of this game is that it demonstrates, in the small, the full capabilities of an EWiDS. Novel is the fact that we do not assume the existence of a backbone infrastructure. This application has been chosen to further investigate the relation between, on the one hand EWiDS design and deployment, and on the other acceptance by end users in terms of dependability issues (i.e., human-system interaction covering availability and acceptance of use).

The second application concentrates directly on managing a crowd. We will concentrate on extracting data on measured social interactions in a large group of people (i.e., a crowd). The data will come from the same electronic badges used in the first (sub)application. Extracting data on social interactions will essentially require the use of a collection of special sensors that participate in the crowd, and at the same time are connected to an external feedback system. Typically, questions regarding network sampling and bandwidth reduction techniques will need to be addressed. At the data-content level, together with social scientists, we aim at performing real-time analysis of the data to discover social groups, along with data on the movement of the crowd as a whole. This information will then be used to intervene in the crowd behavior with the goal to keep people feeling safe and comfortable.

5. Consortium

The proposed research requires input from different disciplines: systems-centric R&D, ubiquitous computing, and EWiDS applications.

WHAT	WHY	WHO
Systems-centric R&D	Expertise in devising systems solutions for EWiDS, covering gossiping and associated (intelligent) applications, energy awareness, and scalability. Expertise should cover complete communication stack from the MAC layer and up.	VUA TUD-ES TUD-MSP DevLab
Ubiquitous Computing	Our focus application area is supporting applications in which humans are supported by means of a EWiDS, which requires knowledge and insight into application-oriented solutions and human-systems interaction.	All
Applications	Applications drive our solutions. The constraints imposed by the hardware and the environment in which it is to operate can be	TUD-IO VUA (Social sc.)

	addressed successfully only in an application context.	DevLab Thales

Each of the partners has a track record in the indicated discipline, and in specific cases is recognized to be leading in the field:

- VUA: Established and recognized expertise in gossip-based solutions for EWiDS.
- TUD-ES: Established and recognized expertise in sensor networks, notably in MAC protocols.
- TUD-MSP: Renowned for their work on (distributed) signal processing.
- TUD-IO: One of the few groups in the Netherlands focusing on “device-independent” human-systems interaction.
- Thales: Renowned for agent-based modeling and design of complex systems.
- DevLab: Inniator of the MyriaNed EWiDS, aimed to comprise 10,000 collaborating wireless devices.

The detailed description of work packages will show a close collaboration between these partners. We stress that the partners have successfully worked together in other projects, and continue to do so outside this proposal. In this project, we have explicitly planned the following collaborations:

- TUD-ES and TUD-MSP will jointly share and supervise a PhD researcher.
- PhD researchers from VUA and TUD-ES will collaborate in the same workpackage, and be jointly co-supervised by senior staff from both groups.
- PhD researchers from VUA, TUD-IO, and Thales will collaborate in the same workpackage, with co-supervision by senior staff from all three groups.
- PhD researchers will be stationed part-time (e.g., 1 day/week) at DevLab (this is already common practice between TUD-ES, VUA, and DevLab in other projects).
- The PhD researcher from TUD-IO will be formally co-supervised by senior staff of VUA.

Notably working together at a single place (DevLab) has shown to be extremely fruitful and in many cases much more effective than formal project meetings. The latter will, of course, also be held, next to a joint seminars and regular informal workshops. We are used to this way of working and over the years have established a solid link between groups.

Furthermore, we stress that at least 1 PhD student at VUA will be co-supervised by an informatics scientist and a social scientist, and in strong collaboration with the senior staff from industrial design (TUD). The envisaged social scientist is an expert in social network analysis.

6. Workplan

The project has four work packages, aligned along the separation of components and functionality in an EWiDS as explained above:

- WP1: Adaptive data dissemination
- WP2: Distributed data aggregation
- WP3: Crowd data processing and feedback
- WP4: Individual data processing and privacy
- WP5: Consolidation

Work package details can be found separately, including timing, deliverables, and milestones.

WORKPACKAGES

Project number P9	
Project title & WP name	EWiDS WP1: Adaptive data dissemination
Projectleader	TUD-ES

Objectives

In *WP1* emphasis is on devising adaptive in-network gossip-based solutions for large-scale wireless systems. We are particularly interested in finding efficient context-aware and fully decentralized solutions for information dissemination. To this end, we will develop adaptive low-level (i.e., MAC-level) algorithms that can support efficient information diffusion. For example, the nodes in a network need to increase their gossiping (i.e., data exchange) frequency when the rate at which events take place increases as well. In turn, this requires that nodes need to communicate more often, and thus need to adjust their duty cycles. We will also investigate how applications can influence these decisions by providing specifications on their behavior. Given these specifications, a node can proactively increase its gossiping frequency, for example, because it knows that many different events need to be rapidly disseminated. In contrast to the work done in WP2, our solutions will most likely be tied to low-level hardware characteristics.

Project number P9	
Project title & WP name	EWiDS WP2: Distributed data aggregation
Projectleader	TUD-MSP

Objectives

In *WP2* we concentrate on devising highly scalable gossiping algorithms to support data and (audio) signal aggregation aimed at social network analysis. Of particular concern will be devising solutions that are tolerant to poor and asymmetric communication. For example, we have devised a gossiping algorithm for computing averages in wired systems that relies on reliable communication links. Without such a reliability assumption, the algorithm will fail. In this work package, we will be looking at general continuous (signal and) discrete data aggregation solutions that can be proved to be correct within certain statistical bounds, among other comparable examples. The groups at TUD-ES and TUD-MSP share a PhD researcher to work on low-level distributed aggregation protocols for continuous data (in particular, audio), with support from the hardware experts at DevLab. The PhD researchers at VUA and TUD will collaborate in adopting the gossiping protocols developed at VUA such that these can support efficient and effective high-level data aggregation, in particular statistical aggregation for social network analysis. These algorithms will be implemented on DevLab hardware for validation and demonstration purposes.

Project number P9	
Project title & WP name	EWiDS WP3: Crowd data processing and feedback
Projectleader	TUD-IO
Objectives In *WP3* we investigate the relation between groups of end users and the design of extreme wireless distributed systems, concentrating on support for crowd management applications. The investigation focuses on the assumption that we can influence crowd behavior if the (or some) members of a crowd are connected to a physical social network of devices that could inform them about the situation they are in. The information gathered from the crowd could be positioning, movement and direction data. The information provided to the crowd could be directions towards safe places, points of interest, or for example crowd density information or information on potential danger. In this project we have to make concepts such as crowd membership, crowd behavior, context and control methods operational. And validate, via simulation and experimentation if they work. We will try to answer questions like: How to provide people with different sorts of information and how to do it in such a way that they will be accepting the implicit controls. How can we monitor and make sense of their original behavior and their response to the information stimuli that we want to provide. What is a workable temporal latency between observing and responding to certain type of behavior, i.e. what are the maximum allowable delays between observation and control stimuli. In terms of applications, one can also think of applications other than 'avoiding disasters due to crowd behavior'. E.g., trying to control the distribution of people over a certain area, for example in a zoo or on a railway platform or in a museum or stadium, etc.	

Project number P9	
Project title & WP name	EWiDS WP4: Individual data processing and privacy
Projectleader	Thales (D-CIS)
Objectives In *WP4* we will focus on the processing and use of real-time localization and context information from individual nodes in extreme wireless distributed systems. We will consider two different sorts of operational environments: in the city and at home. The focus of the work package is on two aspects of processing and use of real-time localization and context information, namely (1) how to use AI techniques and (2) how to protect privacy of users of whom data/signals are collected by the network. As an application, implemented on an EWiDS network, we will investigate the provisioning of two sorts of services to individuals (humans) that are part of an extreme wireless distributed system: the push services, delivered from the system to the human via an interface on a connected device and the pull services demanded by a human from the system. An example of a push service in the context of the city can be the automatic notification of the proximity of public transport facilities, friends or private transport (taxi's). An example of a pull service can be the request for notification of a certain shop or a medical doctor or 'somebody I know' in the immediate vicinity. Note that these services are, in principle, to be provided by the network, and not by an external (centralized) service. With respect to the focus on the well-being of people, we define the goal here as being provided with pre-defined requested information about available services any time, any place and the capacity to issue specific service requests any time and any place. By demonstrating that we can build applications that support this, based on the use of an extreme wireless distributed system, we address the defined technical and scientific problem in a direct way.	

Project number P9	
Project title & acroniem	EWiDS WP5: Consolidation
Projectleader	VUA

Objectives

The objective of knowledge consolidation is to generalize the project's research results, to make them re-usable for application beyond the (limited) scope and time-line of the current project. Knowledge consolidation provides a sustainable basis for exploitation of research findings, allowing a significant multiplier of the original research investment to be realized. Successful knowledge consolidation incorporates a variety of activities, but mainly: (1) the elaboration, quality improvement and generalization of research results, and (2) the application of the results in a different context, i.e., in different projects within or outside of COMMIT, in completely different user domains of for new industries. A precise consolidation plan will be produced in the first month of the project. The reason for postponing this activity, is caused by the fact that we believe it to be prudent to organize a separate 1-day workshop on how to disseminate and consolidate the results. As such, it is felt proper to postpone this until the project has formally started. Consolidation work will be embedded in the project, as part of the various activities of the PhD students and their supervisors.

DELIVERABLES

Number of important journal paper
6

Number of important conference contributions
20

Products

1. 1000 enhanced MyriaNed
   nodes as electronic badges
   - WP 2 YP 2014

Software

1. Software for adaptive data dissemination at MAC level, interacting with gossip-based data dissemination protocol
   - WP 1 YP 2014

2. Software for in-network distributed data aggregation for supporting crowd control
   - WP 2 YP 2014

3. Centralized software for geospatial social network analysis
   - WP 3 YP 2014

4. Software for (secure/private) provisioning of services to individuals in a crowd
   - WP 4 YP 2014

5. MAC-aware, adaptive data dissemination algorithm
   In WP 1 we will develop an optimal data dissemination algorithm that continues to perform optimally even in highly congested mobile ad hoc networks. Optimal means that successful message throughput (goodput) is at best, while at the same time channel utilization is optimal as well. This requires a close interplay between MAC-level protocols, application-level solutions, and possibly neighboring nodes. Our solution will be highly adaptive to a changing network topology and to changing (local and global) network congestion. The solution will have been devised and implemented for networks consisting of hundreds to one thousand nodes. Experiments will have been conducted using extensive simulations and emulations, as well as real-world hardware platforms of the nodes

developed as part of the project. Result: a documented, implemented, and thoroughly tested algorithm for adaptive data dissemination.

- WP 1

6. Adaptive congestion- and energy-aware MAC algorithms for gossiping

   Our chief goal is to devise a congestion-aware and energy-efficient adaptive duty-cycling algorithm and MAC layer for gossiping systems. As such we have first to explore the design space for suitable, energy-efficient MAC schemes that fit the mobile and broadcast- and dissemination-driven domain. (An evaluation and comparison of different MAC schemes for energy-efficient gossiping is performed in Y2-3.) From the design space exploration, we select suitable MAC layer candidates for which we investigate algorithms to adaptively control the duty cycle of the MAC layer in order to minimize congestion and energy consumption while providing a sufficient quality of service to the upper layers. The solution will have been devised and implemented for networks consisting of hundreds to one thousand nodes and include mobility. Experiments will have been conducted using extensive simulations and emulations, as well as real-world hardware platforms of the nodes developed as part of the project. Result: Based on a selected MAC scheme, a documented, implemented, and thoroughly tested algorithm for congestion-aware and energy-efficient adaptive duty-cycling.

   - WP1 Year 4

7. Extending MyriaCore and gMac with ranging functions, hardware/software

   The next generation of the hardware as specified in WP2.3 of the first year will have support for an enhanced ranging function. This means that it will most probably be based on Time of Flight, or a radio using chirp modulation, or both. The ranging functionality will be implemented as integral part of gMac. Messages received from neighbouring nodes will ultimately be associated with a distance parameter. MyriaCore will be extended with functions that support sensors and actuators that are used for the social experiments. This can be audio interface(s), like a microphone; a visual interface like a LCD-screen with touch control; accelerometer; etc. Result: A documented next version of MyriaCore, gMac and related software libaries that is implemented and used for the ongoing experiments in WP1.1 and WP1.2. WP1.3D1 Year 4

8. Geospatial social network analysis tool for massive mobile ad hoc network

   The final tool for dynamic social network analysis takes real-time traces on social interactions in a large mobile ad hoc network as input. It performs offline, yet in real-time, network analyses based on various metrics regarding centrality, homophily, distance, and others (yet to be finalized). In addition, the tool will allow for flow analysis, which is subsequently used as basis for feedback data on how to control the observed crowd. The tool consists of various data- and compute-intensive algorithms, and is developed in conjunction with the solutions developed in WP 3. Result: a documented, implemented, and thoroughly tested tool for dynamic, geospatial, social network analysis.

9. Next generation MyriaNode
   During the first year the investigation has resulted in the specification of a next version of a node. Chess will then design and produce these nodes. At first a few prototypes will be assembled and tested. Next a quantity of 100 will be produced, followed by a few hundred. Result: MyriaNodes (hardware), including documentation.
   - WP2.3D1Y4

10. Distributed algorithm for acoustical-data processing
   In order to improve the usability of the acoustical observations, the information of interest should be extracted. We will develop an algorithm that adaptively processes the acoustical data in a distributed way such that it is suitable for wireless ad hoc networks with (fast) changing topologies. Examples include (multi-microphone) beam forming in order to enhance the signal of interest. The computational complexity of the algorithm is low. The distributed data extraction algorithm will be combined with the data aggregation algorithm developed in Y1. Result: a documented, implemented, and thoroughly tested low-complexity algorithm for signal enhancement and aggregation in wireless ad hoc sensor networks.
   - WP2.1D1Y4

11. Crowd management tools based on social network analysis and geospatial data
   Crowd management tools will be developed for steering and monitoring crowd behaviour. These tools will be fed by novel techniques for off-line, real-time analyses of a combination of social network and geospatial data. The analyses will be based on insights into types of information that are effective for controlling crowd behaviour (from WP3.1) as well as on the crowd-observation models developed and validated by WP3.2 and 3.4. The tools are part of our (centralized) external feedback system and incorporate the agent-based tools described as deliverable WP3.4D1Y4 (above). Result: A prototype crowd management tool to be implemented in the golden demonstrator, for managing crowd behaviour, based on analyzed social network and geospatial data.
   - WP3.5D1 Y4

12. Activity Support System for Crowd Communication

    Based on the architecture description and taking into account the specific protocols and other technical constraints that apply for constructing a demonstrator for use in the field, a multi-agent based set of crowd monitoring and crowd feedback services will be build. In addition, a number of other services that can be used by individual crowd members for demanding information will be build. The first set of services will support a centralized crowd management application, whereas the latter set of services will support a more distributed approach where individual crowd members can procure information to manage themselves.

    - WP4.1D1Y4

Other results

1. Golden demonstrator for centralized crowd analysis and control based on data derived from a network of electronic tags

   - WP 3 YP 2014

2. Effectiveness of crowd management by means of wireless distributed systems (journal paper)

   We will summarize the user studies that have been done on crowd behaviour and control employing the technologies developed in the project. Basically, the studies are based on various scenarios in which two conditions are compared per scenario: one without any feedback system (baseline condition) and one with the feedback systems developed in the project. The results will provide insights into the effectiveness of information variables underlying the feedback systems proposed. Result: Quantitative evaluation of the effectiveness of the tools developed in the project.

   - WP3.1D1 Y4

3. Acceptability of crowd management by means of wireless distributed systems (journal paper and prototype of feedback systems)

   In a series of research-through-design studies we will continue to develop and validate a prototype of an information-providing device with the aim of testing the multi-user acceptability of feedback systems based on wireless distributed systems. We will report on the acceptability of being part of such network as well as the acceptance of carrying and using an information-providing device. These activities will result in a feedback system that makes it acceptable for individual group members to be persuaded to perform certain behaviour beneficial for the group as well as themselves (e.g. being in a comfortable zone). Result: A validated prototype of a feedback device to be implemented in the golden demonstrator.

- WP3.3D1Y4

4. Final formalized models of crowd agents and crowd management agents.

   These models serve as the specification for a set of agents that will be used to run simulation studies, using synthetic and real world data for validation. These models include crowd-member agents, sensor and actuator agents, tracking and monitoring agents and feedback agents and specify the functional and non-functional behavior for the software agents that need to be build to support. The models will be described in their 'native' format of the selected tools with which they will be produced. In addition, a report will described the main design features in plain text. These models are necessary as the baseline that captures the knowledge of crowds that we will use for further experimental research. The will direct the implementation of the agent-based simulation environment.

   - WP3.4D1 Y4

5. Specification of dissociative and associative communication primitives

   This report will describe the results of a more theoretical study of two specific kinds of communication methods. With dissociative communication we refer to a form of communication where factual information is provided but the receiver is not directly addressed and no action information is provided. With associative communication we refer to a form of communication that is specifically addressing a receiver in his/her context and provide directions for action. The report will provide a baseline for crowd management experiments or simulations where the use of these different feedback approaches can be investigated. This document is the basis for a paper on different approaches for crowd feedback.

   - WP4.2D1 Year 4

6. Scenario elements and definition of analysis for feedback performance metrics.

   This report will describe a number of scenario elements that could be tested with the Golden Demonstrator. For these scenario elements a working feedback mechanism is assumed. The report will elaborate on the expected crowd management performance metrics, associated with the feedback mechanism and on how that performance can be measured and analyzed. This document describes the WP4-specific contribution to the set-up of the Golden Demonstrator.

   - WP4.2D2 Year

The Golden Demonstrator

Within EWiDS we will concentrate on developing a single demonstrator for crowd management and control. The demonstrator consists of essentially two functions:

1. The most important one: the means to collect data from a wireless network of electronic badges to discover and analyze crowd behavior and to subsequently provide feedback to crowd members to, for example, realize uniform distribution of the crowd in time and space. In terms of traffic management, such a uniform distribution corresponds to congestion control. The main challenge is to realize the integration of data gathering, in-network and offline data processing, and providing feedback.
2. For individual members of a crowd, the demonstrator will incorporate location-based information services (where should I go to for [less crowded area, the station/start/bus/taxi, my group members]?) The challenge is to incorporate location awareness for individual nodes, given that they are otherwise "blind." Location awareness strongly depends on the service that needs to be supported.

The demonstrator consists of the following components:

1. A wireless network of (communicating) electronic badges. Badges are equipped with various sensors (microphone, accelerometer) and an actuator (LED display, LEDs).
2. An external, centralized, feedback system used for analysis of crowd data
3. Possibly various external feedback devices (screens, speakers), augmenting the actuators available on the badges.

During the first two years, we aim to build a prototype system with all three types of components, but at a relatively small scale (e.g., up to 100-200 badges). With this system, we will demonstrate:

- The process of measuring and gathering data on social interactions, and extracting that data from the wireless network for further processing. This process will have been supported by a small-scale experiment with 100-200 people, from which raw data is logged on individual nodes.
- An initial user study to find out what type of feedback people would prefer for effective crowd management. If possible, feedback solutions will be incorporated in the prototype.
- A simple offline, real-time analysis of the observed social network, and its evolution over a given period of time (say 1-2 hours max).
- A simple in-network solution for computing local densities, and subsequently signaling a warning to an individual when density changes to an unacceptable level.

- An initial (semi-)formal agent-based model that describes and explains crowd behavior. Input for the model will be the actual data gathered from a small-scale experiment, with validation partly taking place by social network analysis.

Plan of work Golden Demonstrator

Y1 (durations measured in months):

T0+4: Conduct a series of small-scale experiments with currently available electronic badges (50-80). These experiments provide a basic data set from which to operate. Recording of sound is incorporated. (**VUA, TUD-ESI, TUD-MSP**).

T0+8: Provide a visualizer that shows the social network as measured by experiments. The visualize is an enhanced version currently avalailable at VUA. (**VUA**).

T0+8: First results on the type of feedback that should be incorporated for managing crowds. The description is such that it forms realizable requirements for the prototype Golden Demo (**TUD-IO**).

T0+8: A simple crowd behavior model, to be fed with data on social interactions from a crowd of 50-80 people. This model should be used for producing feedback to a crowd (**Thales**).

T0+12: First, very basic demonstrable prototype of the Golden Demonstrator. (**ALL**).

T0+18: Based on experience from first prototype, an enhanced version is to be demonstrated, with features decided on T0+12.

T0+24: Golden Demonstrator as described above (**ALL**).

Risks Golden Demonstrator

An important challenge for the golden demo is to organize at least two events that will allow us to gather data and to demonstrate the system as a whole:

1. By the beginning of Y2, we should have 100-200 electronic badges that can be used in an experiment with 100-200 people, and where we measure social interactions. Practice has shown that setting up such an experiment, and making sure that we get the data is not trivial. The problem is that we are eseentially confronted with a one-time-only event.
2. Toward the beginning of Y4 we will need to organize a similar experiment with at least 1000 people. At that point most technical issues should have been solved, but the mere organization of such an event, or piggybacking the experiment with an already scheduled event, will prove to be challenging (again, as has been witnessed by companies participating in DevLab).

Furthermore, we envisage the following risk:

3. We will be introducing new hardware into the project for which we have only limited (programming) experience. Group members at DevLab, VUA and TUD-ES have experience with programming MyriaNed nodes (the foundation for the new hardware), but many

experiments with wireless sensor networks continue to show that there can still be big gaps to cross between accurate low-level simulations and the real world of wireless systems.

User Study

An initial user study will be conducted focusing on the feedback solutions as incorporated in the prototype system. Multiple user scenarios (e.g. 3 or 4 including ones that generate well-known phenomena like Dirk Helbing's "slower-is-faster" effect in self-organizing crowds) will be tested in a comparative setting with two conditions per scenario: one without any feedback system and a similar one where (some of) the participants receive feedback.

1. In the first condition behavior of individuals in groups (e.g., groups encountering each other, or crossing each other at some entrance of a location) will be monitored without them getting any system feedback on their own behavior in relation to the group.
2. In the other condition individuals will be provided with feedback as implemented in the prototype system, and again their behavior will be monitored in relation to that of the groups.

In this way, the measured (individual and group) behavior will provide quantitative feedback on the effectiveness of the feedback in influencing an individual’s behavior. Additionally, individuals will be asked to express their subjective experiences in dealing with the feedback. On the one hand this will help understanding the measured level of effectiveness (i.e., what makes it (in)effective), on the other hand it indicates how users will experience the system, as well as to what extent they would be willing to use it (acceptability). These insights will be used to further improve the design of the system and its feedback in terms of its effectiveness, acceptability and user experience.

The initial study will be carried out as part of the experiment with 100-200 badges, to be conducted during the first phase (Y1-Y2) of the project.

Werkpakketnr. →	Kennisinstelling	[non]Profit	% matching (alleen 0 of 45)	1	2	3	4	5	6	7	8	9	10	11	Overige kosten budget	Totaal Budget	Totaal Subsidie	Totaal Matching
Naam partner				Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten	Kosten
VU Amsterdam	1		45%	500,3		251,4	251,4									1003,1	451,395	551,705
TU Delft	1		45%	538	538	538	311,1									1925,1	866,295	1058,805
Thales Nederland		1	45%			300	300								11	611	274,95	336,05
DevLab		1	45%	303,5	255,5										53	612	275,4	336,6
Totaal				1341,8	793,5	1089,4	862,5	0	0	0	0	0	0	0	64	**4151,2**	1868,04	2283,16
Totaal Kennisinstellingen				1038,3	538	789,4	562,5	0	0	0	0	0	0	0	0	2928,2	1317,69	1610,51
Totaal [non]profit				303,5	255,5	300	300	0	0	0	0	0	0	0	64	1223	550,35	672,65
[non]Profit/totaal				22,6%	32,2%	27,5%	34,8%	#DIV/0!	#DIV/0!	#DIV/0!	#DIV/0!	#DIV/0!	#DIV/0!	#DIV/0!	100,0%	29,5%	29,5%	29,5%