

SIXTH FRAMEWORK PROGRAMME
PRIORITY 2
“Information Society Technologies”

Project acronym: RUNES

Project full title: Reconfigurable Ubiquitous Networked Embedded Systems

Proposal/Contract no.: IST-004536-RUNES

D2.1
Application Scenario Building/Definition

Project Document Number: RUNES/D2.1/PU/v1.0 (official version)
RUNES/D2.1/PU/ETH/SIRA/Kodak/ULANC/UCL/v1.0 (internal version / contributions)

Project Document Date: 31/03/2005

Workpackage Contributing to the Project Document: WP2

Deliverable Type and Security: R-PU

Author(s): Domonkos Asztalos (ETH), Karen Lawson (Kodak), Stephen Hailes (UCL), Lesley Hanna (Sira), Ingolf Krüger (UCSD)

Abstract:

The RUNES project aims to provide an adaptive middleware platform and application development tools that allow the cost-effective development of reconfigurable networked embedded systems in a wide range of application domains.

This deliverable describes RUNES scenarios. The scenarios serve to give the reader an idea how RUNES will be used to benefit the stakeholders.

Keywords: scenarios, RUNES

History

Version	Date	Description, Author(s), Reviser(s)
1.0	13/11/2005	First Released version

Executive Summary

Background

RUNES will provide the architecture and tools to facilitate the creation of large-scale networked embedded systems, validating the approach taken a combination of real deployment and simulation. The scope of the task requires a comprehensive cross-stack approach. Results will be demonstrated in the following areas:

- A thorough analysis of existing radio and network technologies
- Research in the field of networking for embedded systems
- Design and documentation of an effective systems architecture
- Development of component based middleware that is adaptive and intelligently self-organising
- Development of novel methods for advanced control, aimed at exploring the issue of predictability
- Development of tools that allow for application development, including the automated assessment of usability, and debugging
- Construction of both real and emulated deployment environments for assessment of the architecture
- Creation of application for validation purposes

Objective of this document

This document has several purposes; 1) to gain a common project-wide understanding of where we want to take the RUNES project, 2) to create scenarios that we can use for deriving requirements for RUNES, 3) to use the scenarios to illustrate certain technical concepts that are vital for RUNES, 4) to collect feedback on the RUNES concept from potential stakeholders, and 5) to create and collect more detailed concepts that are related to RUNES control functions.

Approach

This deliverable describes RUNES scenarios illustrating how RUNES will be used in the future. The main goal of the scenario work within the RUNES project is to help us understand and derive the requirements for the RUNES architecture. We will refer to some other IST projects close in topic to the RUNES. With the selected scenarios we want to emphasize the integrating, technology providing scope of the RUNES.

Scenarios

Three RUNES scenarios are provided to illustrate how RUNES will be used in the near future. Most of the scenarios were already implemented by a certain degree using particular architectures, technologies and tools. The RUNES solution will provide a generic framework for the application development and deployment in NES environment and the scenarios will be used to validate the provided framework. The following scenarios were selected:

Integrated Cardiac Telemonitoring and the Smart Environment Scenario

This scenario covers the remote patient monitoring in home environment combining at least three domain oriented network types: Body Area Network - a mobile wearable wireless sensor network, Smart Home Network – a LAN with combined wireless/wired sensor subnetwork, Electronic Health Care Record Network – a wide area network of a distributed database application. The primary scalability dimension is the number of patients. The primary advanced control subscenario is the online treatment of a patient with the stringest security conditions.

Disaster/Emergency Scenario

This scenario covers the case of a hypothetical earthquake situation with many ad hoc networking situations needed for example to build up a rescue/recovery surveillance network based on the damaged communication infrastructure. The primary scalability dimension is the geographical extension of the event directly influencing the number of sensing devices for data collection. The primary advanced control subscenario is the data filtering for gaining relevant and reliable information in timely manner.

Integrated Wine Production and Distribution Scenario

While the above two scenarios are from the public domain, the present one is a real production oriented scenario. It covers a wine production and distribution chain with at least four domain oriented network types: Agricultural Field Network – large-scale wireless sensor network, Production Plant Network – a LAN with combined wireless/wired sensor subnetwork, Intelligent Truck Network – a mobile combined wireless/wired sensor network, Retail Store Network – a LAN with a sensor network. The primary scalability dimension is the size of the cultivated area. The primary advanced control subscenario is the quality preserving delivery process.

Automotive Scenario

The car industry is of key importance for European economy. Intelligent transportation systems are becoming more and more interesting, and the fact that roads are becoming instrumented sensor networks will make it possible to provide “intelligent transportation services” within the boundaries of our cars. Most of the interesting new developments in automotive systems are driven by software and electronic components. In particular many new applications scenarios requires coupling of car internal embedded devices with the external environment. This can be accomplished using wireless technologies and deploying a mix of wireless and wired embedded system in vehicles and roads.

Fire in a Road Tunnel Scenario

The recent case of a road tunnel fire (http://news.yahoo.com/s/ap/20050605/ap_on_re_eu/france_tunnel_fire) in the Frejus tunnel in the Alps on June 4, 2005 highlights the importance of dealing with the emergency scenarios in a project like RUNES. The usual “peaceful” scene can change very dramatically into a highly dynamic and chaotic situation. A basically static monitoring infrastructure containing a large set of sensing devices connected to a central operating room must be able to adapt itself to the emergency situation through providing network access to the infrastructure of the emergency agencies having many mobile and wireless nodes. The scenario may provide use cases for demonstrating how the network and middleware solutions of RUNES realize the technical objectives of the project.

Contents

1	Introduction	6
1.1	Typical RUNES applications.....	6
1.2	Related works.....	6
1.3	Structure of the scenario descriptions	7
1.4	Context of the scenarios	7
2	Scenario: Integrated Cardiac Telemonitoring and the Smart Environment	9
2.1	Background.....	9
2.2	Analysis and Motivation	10
2.3	Use case description: the monitoring of the patient	12
2.4	Communications and Technical Implications.....	13
2.5	Summary.....	13
2.6	References	14
3	Disaster/Emergency Scenario	15
3.1	Background.....	15
3.2	Analysis and motivation	16
3.3	Communications and Technical Implications.....	16
3.4	Use case description: asynchronous data migration.....	19
3.5	Summary.....	21
3.6	References	21
4	Integrated Wine Production and Distribution Scenario	22
4.1	Background.....	22
4.2	Analysis and Motivation	22
4.3	Communications and Technical Implications.....	23
4.4	Summary.....	24
4.5	References	24
5	Automotive Scenario	25
5.1	Background.....	25
5.2	Analysis and motivation	25
5.3	Use Case Scenario: Accident Mitigation	26
5.4	Use Case Scenario: Assisted Driving	27
5.5	Communications and Technical Implications.....	27
5.6	Summary.....	28
5.7	References	28
6	Fire in a road tunnel scenario	29
6.1	Background.....	29
6.2	Analysis and Motivation	30
6.3	Use Case Descriptions	31
6.4	Communications and Technical Implications.....	32
6.5	Summary.....	34

1 Introduction

1.1 Typical RUNES applications

The typical RUNES application may be characterized as one that takes input from many remote sensors, provides geographically dispersed operators with the ability to interact with the collected information, and controls remote actuators. Those applications are part of the distributed, real-time and embedded systems (DRE systems).

The potential application areas are:

- Healthcare
- Smart building, smart home
- Emergency, disaster management
- Avionics
- Automotives
- Intelligent transportation systems
- Agriculture
- Retail logistics
- Command and Control

The RUNES DoW (Description of Work) indicates a focus first on the Wireless Sensor Networks (WSN), which are characterized by the radio communication and a set of micronodes, having scarce resources and several sensing and/or actuating devices, second on the integration of the WSN with other types of networks. The RUNES project aims at developing a generic, reconfigurable middleware for that platform of heterogeneous networks to provide a framework for the development of data collection and control applications with more ease than current technologies and tools afford.

1.2 Related works

The above narrowing of scope through focusing can be supported by a short overview of the related works. There are several running and/or finished IST projects, which are closely related to the RUNES project.

Ambient Network

<http://www.ambient-networks.org/>

To make any NES useful it needs flexible networking solutions. For example, a Body Area Network for the collection of health data from a patient would not provide much use without a means to move the data to healthcare systems. A real application with a *pervasive* and/or *ubiquitous* NES has to rely on a heterogeneous network infrastructure to communicate. The main focus of the Ambient Networks (AN) project is to combine isolated communication systems of different types into an overall network. The critical component of the AN network infrastructure from the RUNES perspective is the WSN. A consequence of the existence of the AN project is that the research in RUNES isn't focused on the ambient network issues per se. The AN project also has a health care scenario defined.

MobiHealth

<http://www.mobihealth.org/>

The MobiHealth project focused on a particular AN configuration: the integration of a Body Area Network with the GSM infrastructure.

6WINIT

<http://www.6winit.org/>

The 6WINIT project aims at introducing Mobile Wireless Internet in Europe. It concentrates on the problems raised by the mobile dimension, building on existing fixed IPv6 infrastructures from other initiatives. The project provides a number of testbeds, including one *healthcare* environment.

EYES (Energy Efficient Sensor Networks)

<http://www.eyes.eu.org/>

The EYES project works on self-organizing and collaborative energy-efficient sensor networks. It addresses the convergence of distributed information processing, wireless communications, and mobile computing.

T-Engine

<http://www.t-engine.org/>

T-Engine is Japan's ubiquitous computing architecture, arguably the most advanced in the world. T-Engine enables the distribution of software resources, including middleware developed on T-Kernel, its real-time operating system.

1.3 Structure of the scenario descriptions

Background

In this section the environment in which the scenario is set is described as well as any assumptions made that are not obvious.

Analysis and Motivation

In this section the general domain-specific aspect is described highlighting the issues that demonstrate the interest and value to stakeholders.

Scenario Story

In this section the overall story is described in a non-technical way.

Communications and Technical Implications

This section covers the characteristics of the scenarios that affect their technical realization. This section provides input to the requirements definition for the RUNES project.

Summary

1.4 Context of the scenarios

1.4.1 The stakeholders view

The ubiquitous networked embedded systems term refers to the situation where the wide spectrum of different appliances containing sensors, actuators, computational and communication devices can support and maintain a better quality of life for living organisms, and better operational conditions for organizations everywhere and all the time. The main stakeholders in that business are the individuals and the organizations. The three most important elements affecting the evaluation of a solution for the stakeholders are:

- comfort and entertainment
- safety and security
- optimization and efficiency

We selected twelve large application areas and mapped them to a Venn diagram based on the above three evaluation elements (Figure 1.). One can notice later on, that our scenarios cannot be allocated to just one application area, they cover usually more than one of them.

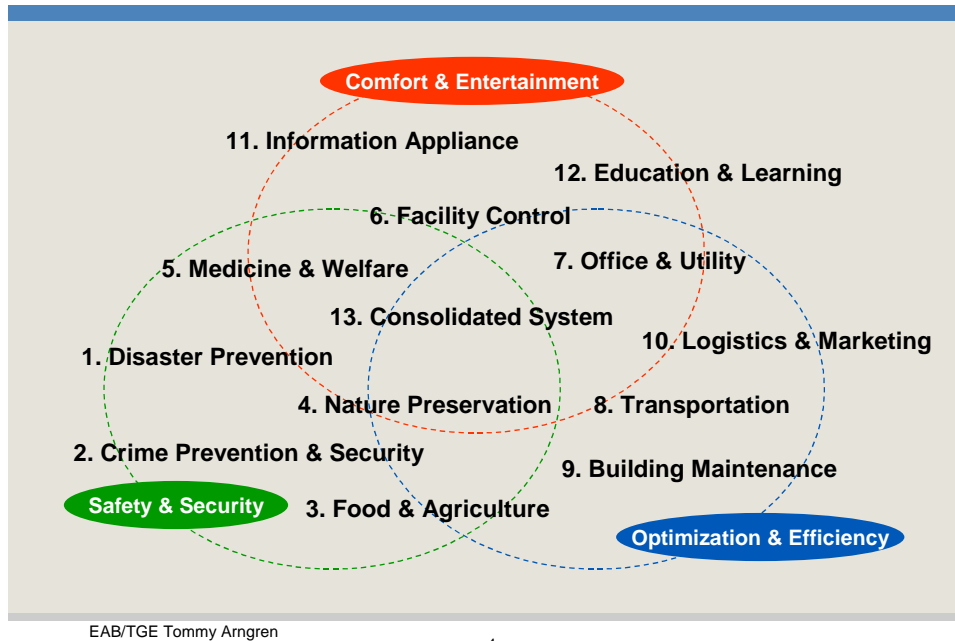


Figure 1.

1.4.2 Technical Objectives of the RUNES

In that chapter we cite the Technical Objectives as they are formulated in the DoW document of the RUNES. Our hope is that helps to evaluate the usefulness of the scenarios from the technical point of view.

Technical objective 1 is to analyse existing radio and hardware technologies in sufficient depth that (a) there is sound basis for the development of the upper layers and (b) so that suitable technology can be selected for validation.

Technical objective 2 is to undertake research in the field of networking for embedded systems, including for resource poor environments, and to select and develop appropriate techniques for logical networking, mobility, and autoconfiguration in systems of this scale.

Technical objective 3 is to build well-founded adaptive component-based middleware that has the ability to self-organise, and which can support a range of non-functional requirements in the form of different mechanisms for ensuring reliability and security. Such middleware will be designed and operate over heterogeneous underlying platforms.

Technical objective 4 is to undertake research and, on the basis of this, to build a computational framework that gives us the ability to develop systems in which industrial automation requirements for predictability can be met, and for which we have an understanding of the limitations of its applicability.

Technical objective 5 is to construct appropriate deployment environments that allow the assessment of our architecture against the general requirements that arise from the project vision and, more specifically, against of the applications selected as consequences of the WP2 and WP8 activities in technical and commercial scenario definition respectively. This will be achieved by constructing both real and simulated environments.

Technical objective 6 is to build tools that allow the construction of applications that can be used to validate our architecture and then to use such tools to build applications both for small-scale and large-scale environments.

Technical objective 7 is to research appropriate measures, and then design methodologies and finally to develop tools that allow the assessment and knowledgeable construction of applications for the embedded systems environment and the changing needs of the user so that it continues to respect usability requirements.

Technical objective 8 is to design and document a system architecture that promotes interworking of elements from different layers in such a way as to allow the effective construction of applications for large scale networked embedded systems.

2 Scenario: Integrated Cardiac Telemonitoring and the Smart Environment

2.1 Background

Heart failure is a disorder in which the heart loses its ability to pump blood efficiently throughout the body. It may affect the left, right, or both sides of the heart. Heart failure is one of the most important causes of mortality and morbidity and is associated with high costs to health services and often a poor outcome for the patient. An estimated 4 to 5 million people in the United States and 10 million people in countries that are represented by the European Society of Cardiology are suffering from heart failure. The prevalence of symptomatic heart failure is estimated to range from 0.4 % to 2 % in the general European population. In the United States nearly 500,000 patients are diagnosed with heart failure annually and it is the underlying reason for 12 to 15 million visits to physicians and 6.5 million hospital days. In the United States, heart failure accounts annually for approximately 0.9 and 2.6 million hospitalisations as primary and secondary diagnosis, respectively. Heart failure accounts for 2-3% of hospital admissions in Scotland for acute or chronic decompensated heart failure [1,2].

In spite of a great improvement in the treatment of patients with heart failure, the mortality rate remains high and the quality of life for the patients is less than ideal. The incidence of heart failure has declined among women but not men, whereas survival has improved in both sexes over the past 50 years. Although the number of deaths due to heart failure has risen generally with the ageing of populations, age adjusted death rates in the elderly have also risen which probably reflects longer survival with hypertension and coronary heart disease.

These mortality statistics differ in countries with different coding practices, and in fact heart failure is often a hidden cause of death. The prognosis of heart failure is comparable with that of the worst malignant diseases.

Acute heart failure, when not causing death, often becomes chronic and *lifelong medication and treatment* is necessary for the patient. Patients with this diagnosis are expected to adhere to a complex behavioral regime including medication, monitoring for signs of fluid accumulation, and a changed lifestyle.

There is a need to *collect data* on a regular basis from such patients in line with guidelines for diagnosis, monitoring and pharmacological management. Some of this data, such as daily weight and the trend of weight can be collected fairly easily. Other data such as heart rhythm, heart rate, blood constituents, and intra cardiac pressures are more difficult to collect on a regular basis. However, the latter measurements are very useful in that they might predict what is otherwise a very sudden worsening of the condition. Thus, monitoring such data is very valuable, and should preferably be done frequently and without undue restrictions to the everyday life of the patient. The heterogeneous and multilevel nature of the data that is needed to properly manage this disease requires integration of data sources and information fusion.

Beyond data collection, there is a need to *make controlled and precise changes to the treatment regime* of the patient. Such changes need to be made based on historical information, real-time sensor data, user input, healthcare protocol, and physician directives.

The process of *diagnosing heart failure* is complex as it involves both the assessment of several risk factors coupled with subtle symptoms expressed by the patient in many early cases. Recognition of heart failure involves the following:

- Dyspnea on exertion
- Dyspnea at rest
- Orthopnea
- Paroxysmal nocturnal dyspnea
- Fatigue
- Decreased exercise tolerance
- Unexplained cough, especially at night
- Acute confusional state, delirium
- Abdominal symptoms (nausea, abdominal pain or distention)
- Decreased food intake
- Decline in functional status

Given these and other risk factors such as family history, lifestyle, weight, age, etc, the consultant can decide on further workup and assessment. These procedures involve significant quantities of data and information. The procedures commonly used in assessment of the degree of heart failure are

- Blood laboratory tests - complete blood count, electrolytes, thyroid stimulating hormone, blood urea nitrogen [BUN], serum creatinine, urinalysis)
- Chest x-ray
- Electrocardiogram
- Radionuclide scanning

In initial workups, the goal is to assess and reverse any side effects or etiologies of the primary heart failure condition. Once this is complete and stable, the task is to devise and monitor the ongoing treatment program [3].

2.2 Analysis and Motivation

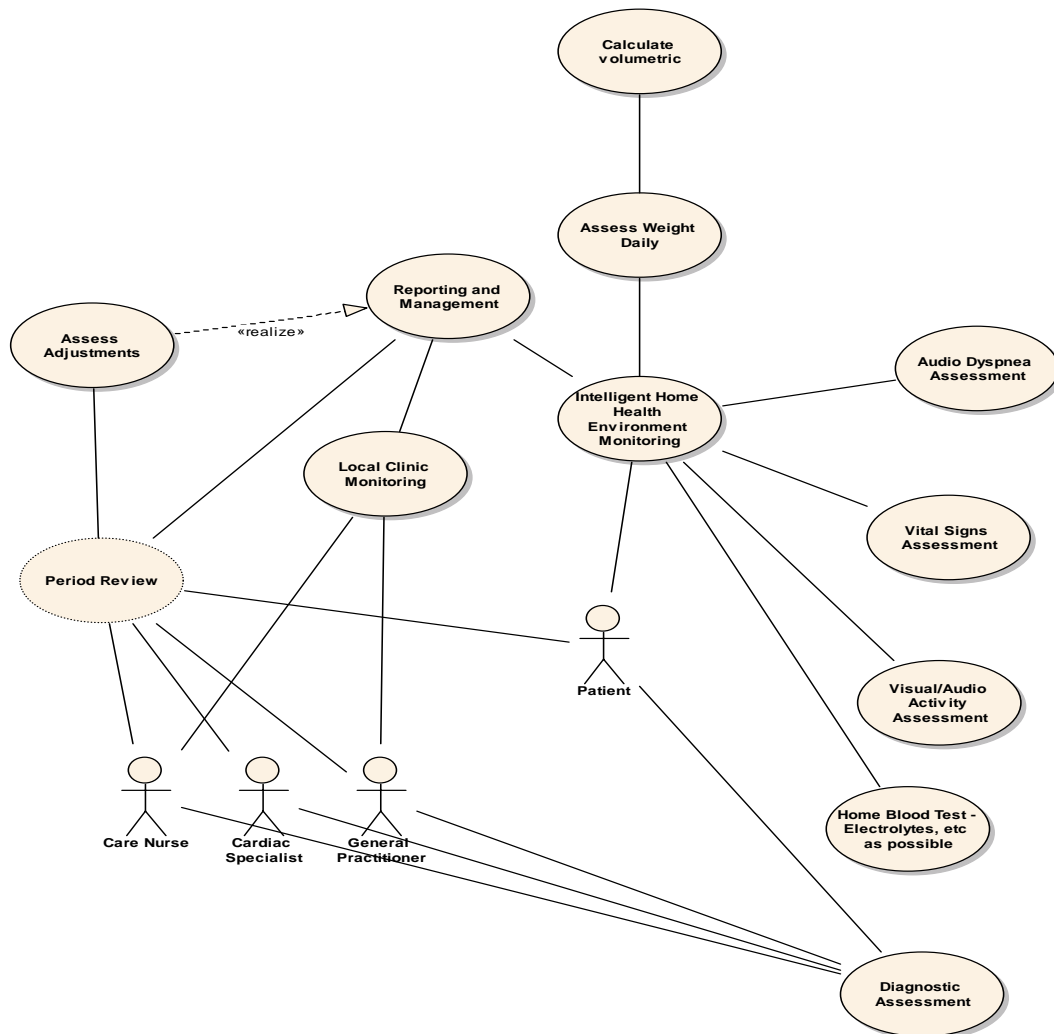
The continuum of severity in heart failure cases means that there are a large percentage of long term patients who require consistent and expensive treatment regimens. This coupled with the need to allow the patient to be active and mobile while still monitoring the necessary disease parameters means that methods and mechanisms are needed that enable real-time assessment of the patient status in a mobile, ad-hoc environment. Rapid action can also be required when parameters such as increased fluid volume or tachycardia are detected. Thus, automatic actuation of mitigating techniques or triggering of warning systems are essential to enhance the chances of survival and quality of life of heart failure patients.

There are several levels of information and actuation involved in this scenario. At the lower levels, there is the internal sensing of relevant parameters such as pharmacological blood levels. This data must be combined with structural and physical information such as weight gain, activity levels, and breathing regularity. At the higher levels, diagnostic information such as electrocardiograms and X-rays as well as observations from physical examination need to be integrated and available in real-time to give an accurate picture of the patient's situation.

The opportunities in this area are many, but one of key interest is combining building automation and local and wide area communication to manage and communicate patient information. By designing and developing a two component approach to this problem, one for the fixed environment and one for the mobile, *new technology* can be deployed that greatly improves the quality and ultimately reduces the cost of monitoring the patient thereby improving the quality of life and healthcare delivered.

The two-element approach when integrated means that a complete care environment can be built up over time as the disease progresses or can be implemented immediately as the severity level dictates. Through this incremental technology environment, local and wide area information is maintained in the appropriate repositories while sensing, information collection and actuation can be distributed. Through appropriate alerts and directives transmitted at the right moments, the patient and the local caregivers can act or react quickly as is necessary in sudden negative symptom appearance. Proactive monitoring can be done in a non-invasive manner to improve quality of life for the patient while providing the information necessary to prevent unnecessary medical emergencies.

2.3 Use case description: the monitoring of the patient



The patient has an intelligent *in-home monitoring system*, which consists of sensors, actuators, and monitoring stations distributed throughout their normal environment.

The use case takes the patient from first diagnosis through to treatment and monitoring. The *in-home system should be simple to install, come as a package of appropriate sensing, actuating and computing nodes based on the stage of the disease.*

Sensing and monitoring without undue intrusion is achieved by taking advantage of the daily routines of the patient. For example, audio sensors are activated in rooms where the patient is located by reacting to RFID tags on the patient or other means. Once automatically activated, audio sensors monitor the sounds of the patient breathing and run algorithms for detection of the various forms of dyspnea, which are very important in the monitoring of the disease state. Accurate weight sensors are embedded in bathroom floor mats so that each day the patient is accurately weighed and changes assessed using algorithms executed locally. Using this daily information, the treatment regime can be adjusted or appropriate activities specified to the patient.

Activity monitoring is done through the use of web cams in specific rooms much like security motion detectors to ensure that the patient is moving and active over time periods. This is done quietly and without intruding on the daily life of the patient.

More *specific blood test kits* can be deployed to the home environment as well to measure electrolytes, blood urea nitrogen levels.

FP6 IP "RUNES" – D2.1 Application Scenario Building/Definition

Simplified vital signs monitoring can also be done through the inexpensive addition of blood pressure and pulse units in the home, which are retrofitted with data collection units. Although this requires explicit attention from the patient, the patient does not need to care about storing and reporting the data.

A patient monitoring centre with touch-screen or other easy interface can be used to *ask some general questions* of the patient (requiring patient direct interaction). Reminders for intake levels, diet restrictions, and positive reinforcement can be provided through a real-time feed to the monitoring centre. This centre can be through the patient's own television, or PC or mobile platform.

The patient is thus surrounded by a network of sensor and actuator nodes, with a common need to be able to communicate their readings and receive data over some *communication channel*. The channel used will depend on the actual physical environment and the available communication systems.

The equipment deployed in the home of the patient typically connects over standard public networks to the local clinic where first level healthcare is provided. Obviously, such a connection has to be secure and authenticated to ensure that sensitive patient data is not leaked and remains intact.

The patient is *periodically reviewed* in the clinic environment and perhaps the same or more detailed assessment and monitoring is done. *The patient record (or parts of it) can be maintained* at the local clinic or it can be maintained across a virtual private network connecting to a central clinic.

The cardiac specialist at a central clinic will also review the patient and may require still further testing and diagnostic procedures such as imaging.

At each location in the scenario, the intelligent building components can occur. The scale and ambition of the deployment varies with the context, especially the volume of patients and the nature of the building itself. For example, the clinic may require weight, vital signs, but will not require activity monitoring (as a visit to the clinic is a special case). Thus, certain parts of the overall application may need to be turned off temporarily when called for by the context.

As the *patients move from the home environment to the clinic environment*, information may need to be carried with them. Upon moving into the new environment, for example in the clinic, this information must be made available to the practitioner without intervention.

To monitor the patient through their daily lives will *require a subset of sensors and actuators deployed to vehicles in which they travel*. This set of nodes will need to maintain information on the patient's condition during travel and synchronize this information upon arrival to a destination which supports the monitoring function. It will *need to store information* for variable lengths of time. *The mobile monitoring subsystem could be associated with the patient or within the vehicle depending on the extent of the information to be collected.*

Conversely, as there are new treatment or regime directives, these must be immediately and completely *integrated in the intelligent home environment* and the new protocols actuated and validated for consistency.

While this is undoubtedly the same scenario possible at the hospital level at an increasingly more complex level, this would require a much broader and larger project than is possible with the RUNES program. Based on the success with the two tiers already mentioned (home and local clinic), perhaps a subsequent project can explore the addition and integration of existing hospital technology into such a scenario.

2.4 Communications and Technical Implications

The implications of the scenario outlined above are several.

First, there is the *smart home environment*, which will need to be deployed in various contexts, from regular private homes to assisted living facilities. This implies that the infrastructure of the networks, the nodes, and all components must be easily and non-permanently deployed without (too much) wiring work and with relatively unstable and unpredictable maintenance.

The communications infrastructure must support both the local, mobile, wireless and the fixed wire-line stationary network connections. It should be able to make use of whatever communications systems are available in a particular environment.

The network of monitoring and actuating nodes must support *redundancy and automated formation of a network* as well as multiple simultaneous accesses through different mediums.

Information must be transportable and transmittable between locations with other smart environments.

There is a mixture of real-time information and non-real time information and the network will need to be able to support the *reliable and robust transmission* with assurance of both.

The scenario includes the need to *develop new nodes*, both sensing and actuating ones, plus the ability to retrofit existing data collection points. It requires condition sensitive packaging, meaning the systems should be segmented according to the needs of the treatment and monitoring regime, which is established.

Interfaces will be required for all categories of information. There will be high priority directives, high priority sensed data mixed with informational and general reporting data. This will require interfaces for the medical teams, the local caregivers, as well as the patient and his/her family and friends.

Important information must be segregated and reported properly and in a *sensitive manner* for the audience at the right level.

Conflicting and alarming situational inferences must be prevented, but the truly *life threatening events must be detected* and handled appropriately.

New components for treatment or parameters that are to be added to the monitoring protocol should be *active deployed* with minimal cost and disruption to the existing environment and the patient.

The technology must *adapt to changing conditions and demands* over time without a complete re-installation or deployment.

Integration with the clinic systems and reporting at the clinical level to the patient record needs to occur in accordance with the *privacy and security* provisions of the healthcare industry, which are currently evolving rapidly.

These modules and other algorithms within the system that interface with the various health information sources must be adaptable.

2.5 Summary

Heart failure is a devastating disease, which requires careful, continuous monitoring of various parameters. Patients with heart failure often need extensive and expensive assessment and treatment, but in many cases, this is followed by long periods of monitoring with intermittent treatment and further assessment.

The scenario outlined creates a new environment and technology for post-diagnostic monitoring which integrates the assessment of information with the ongoing treatment. The resulting technology will improve the quality of life for the patient through providing accurate and frequent essential monitoring of critical diagnostic parameters. It allows the proactive treatment of the reversible aetiologies while providing non-intrusive assessment in real-time. Through the use of integration with the clinical environment and the integration of the historical patient record, and novel sensors the system will provide treatment indicators quickly and efficiently through the use of networked technology. Using an adaptive interface, the presentation of the various parameters and situation can be presented appropriately to consultants, caregivers and the patient without undue worry or alarm. In doing so, the RUNES resultant technology will improve quality of care through accurate information and proactively assist in the treatment of this growing and very serious condition.

2.6 References

1. <http://www.netdoctor.co.uk/diseases/facts/heartfailure.htm>
2. <http://www.escardio.org/knowledge/ehs/slides/>
3. American Medical Directors Association (AMDA). Heart failure. Columbia (MD): American Medical Directors Association (AMDA); 2002. 18 p. [31 references]

3 Disaster/Emergency Scenario

3.1 Background

This scenario is set in a large-scale disaster-relief situation. The primary aim of the RUNES technology is to help coordinators obtain real-time information about the circumstances “on the ground” in order to direct rescue efforts efficiently and effectively. Furthermore, it might also be possible to control aspects of the situation through the use of actuators.

The systems available for use in this scenario might be dedicated to disaster relief applications, but many will have other primary purposes (for example, they could be parts of domestic appliances). The retrieval of information is difficult, because the networking infrastructure (if it ever existed) is expected to have been severely compromised, requiring that wireless network connectivity is established as a prerequisite to any useful work.

The most challenging problem is that there are so many potential sources of information, and that they will vary over time. There is a need to support a large scale networked system, with a high degree of heterogeneity, that may well be failing in some parts (due to gradual loss of power) and recovering in other parts (due to reconstruction).

We wish to provide coordinators with the ability to select the types of information that are important to them, and to adapt the network to ensure that relevant sources are prioritised in the most efficient way. This might, for example, involve the distribution of code into the network in order to perform local processing close to an information sources, as well as simple information retrieval in devices not originally programmed for this function.

There is a wide range of practical problems, for example:

- In technically sophisticated countries, the existing fixed and mobile communications infrastructure may have been damaged or destroyed.
- In environments that never had such infrastructure, communications are usually relatively rudimentary, and even these may take a little while to come on line.
- In either case, there is confusion both within and between relief organizations, and proper communications may take some time to establish. Against this, the chances of saving lives diminish rapidly over time, so accurate and timely information is important in the very early stages of an incident.

As a result, in this scenario we specifically address issues surrounding accurate and effective information transfer across *ad hoc networks*, as they are the form of networks most likely to be quick to deploy and reliable in the case of node failures. Following a disaster such as an earthquake, there is considerable disruption to the fixed communication infrastructure: it may be the case that no such network remains, but it is actually more likely, given some prior emergency planning, that portions of the network remain. These portions may be disconnected islands, or they may be integrated into a wider network. In either case, there is a recovery path that leads eventually to the re-establishment of the previous infrastructure.

In this scenario, we examine several key aspects:

- The networking infrastructure
- The use of sensors and actuators
- The serendipitous use of technologies not primarily intended for emergencies
- Dynamic adaptation to prevailing conditions
- Usability issues
- Security issues

In this scenario, we will restrict what we have to say largely to earthquakes for the reason that they are huge events, causing maximal disruption, and requiring effective coordination of large numbers of resources. Parts of the following will also apply to other natural disasters and acts of terrorism. For example, a tornado may remove communications towers, and electricity lines, since these are exposed, but it may not cause widespread devastation and will not affect buried fibre. The explosion of a ‘dirty’ bomb would require a massive coordination effort, but would hardly affect communications infrastructure.

To assess the scale of such scenarios, we can consider some recent examples.

- The Northridge earthquake in California on 17th January 1994 had a magnitude of 6.7, deforming the earth’s crust over an area of 4,000 square kilometres. It caused 57 deaths and 9,000 injuries, with economic losses of around \$20 billion. 20,000 people were displaced from their homes. 1,600 buildings were ‘red tagged’ as unsafe, 7,300 were ‘yellow tagged’ as restricted entry and thousands more were damaged. The 10-20 seconds of strong shaking collapsed buildings brought down freeway interchanges, and ruptured gas lines that exploded into fires [1]. Even so, this earthquake was considered to be fairly limited in the damages done.

- By contrast, in Kobe, Japan, exactly one year later, a magnitude 6.9 earthquake killed 5,100 people, injured 27,000 and destroyed 100,000 buildings [1].
- On average, there are 120 earthquakes per year in the magnitude range 6.0 - 6.9, meaning that events like those which affected Kobe and Northridge are far from unusual.
- The Tsunami disaster that struck large parts of south-east Asia on December 26, 2005, is estimated to have caused more than 300,000 deaths and destroyed towns, villages, and infrastructure in affected areas.
- Hurricanes are also large scale events, typically affecting between 1,000 and 10,000 square miles, and probably causing more widespread damage to power lines and phone lines than earthquakes.

3.2 Analysis and motivation

Given the above figures, it is not hard to conceive of the immense confusion that results from natural disasters, particularly earthquakes, since they tend to occur without warning and have widespread and devastating effects. In addressing the immediate aftermath of such an event, many different organisations are typically involved.

For example, in the US, the Federal Emergency Management Agency (FEMA) involves tens of different agencies in any recovery plan – from vets (since animals will be injured and will roam wild) through to the US National Guard (to prevent looting) and search and rescue services. Obtaining and prioritising information, delivering it to the right agencies, coordinating inter-agency activities, and deciding on appropriate responses, possibly modified in the light of further information are all extremely important factors in any recovery strategy. On the other hand, incorrect or outdated information can, if propagated, actually increase confusion and hinder activities.

'Response begins as soon as a disaster is detected or threatens. It involves mobilizing and positioning emergency equipment; getting people out of danger; providing needed food, water, shelter and medical services; and bringing damaged services and systems back on line. Local responders, government agencies and private organizations take action. Sometimes the destruction goes beyond local and state capabilities. That's when federal help is needed as well.' [<http://www.fema.gov/about/response.htm>]

Following an earthquake there are many problems:

- Gas mains will have ruptured, causing fire and with the potential to cause explosion.
- Water mains will have ruptured, denying access to the most vital of human needs.
- Ruptured sewers and dead people and animals will contaminate water supplies.
- Pets/animals will be injured and roaming wild. Vets, dog-catchers etc. will be needed, but should have lower priority than those associated directly with human survival.
- Locations at high risk of looting. e.g. pharmacies must be guarded or emptied.
- Families may well be separated, may be evacuated or hospitalised to different places, and so there is a need to locate survivors. Individuals wandering the streets seeking family members are a hindrance to rescue services. The Japanese IAA (I Am Alive) system is one instance of an infrastructure established to aid this [2].
- People need to find shelter, clean water, and food.

Information about the nature of the earthquake may be available from pre-placed, dedicated, sensor systems. For example:

Caltech and the USGS jointly operate a network of seismometers located at over 250 sites throughout Southern California. These instruments detect ground motions and relay that information in real time to the Seismological Laboratory located in Pasadena on the campus of the California Institute of Technology. Twin computer systems simultaneously analyze the data to detect the onset of an earthquake. Once an earthquake has been detected, the computers determine the time, location, and magnitude of the earthquake. For earthquakes greater than magnitude 4.0, ground motion readings from strong motion instruments are used to determine the extent of shaking. This information is then transmitted to commercial paging companies via a dedicated radio link to be sent to personnel wearing belt pagers. The information is also sent to computers connected to receivers which display the data using a program called Qpager. [3]

Thus this information can be used to predict areas of high damage, areas most likely to suffer further destruction during aftershocks etc.

3.3 Communications and Technical Implications

3.3.1 Communications infrastructure

In situations like those in Northridge and Kobe, a Western society suffers sudden and catastrophic damage. It is not reasonable to expect the communications infrastructure to remain intact. Parts of the infrastructure would be destroyed, other parts disconnected, and the parts that survive will be heavily loaded by official and private information traffic.

Alternatively, in situations like the Christmas earthquake in Iran in 2003 (mag 6.6; 31,000 dead) there are few or no existing communications systems; the rescue teams must bring any communications they wish to use with them.

Ad hoc networks are wireless networks that lack fixed infrastructure such as the base stations that form part of standard mobile phone type systems. Instead, the portable nodes which are being used to access the network also act as routers. So, for example, if some node A wishes to send a message to node B which is outside the range of direct radio contact, as in figure 1, it must send its message to some intermediate node, which will forward it (possibly along some chain of intermediate nodes), eventually reaching its destination.

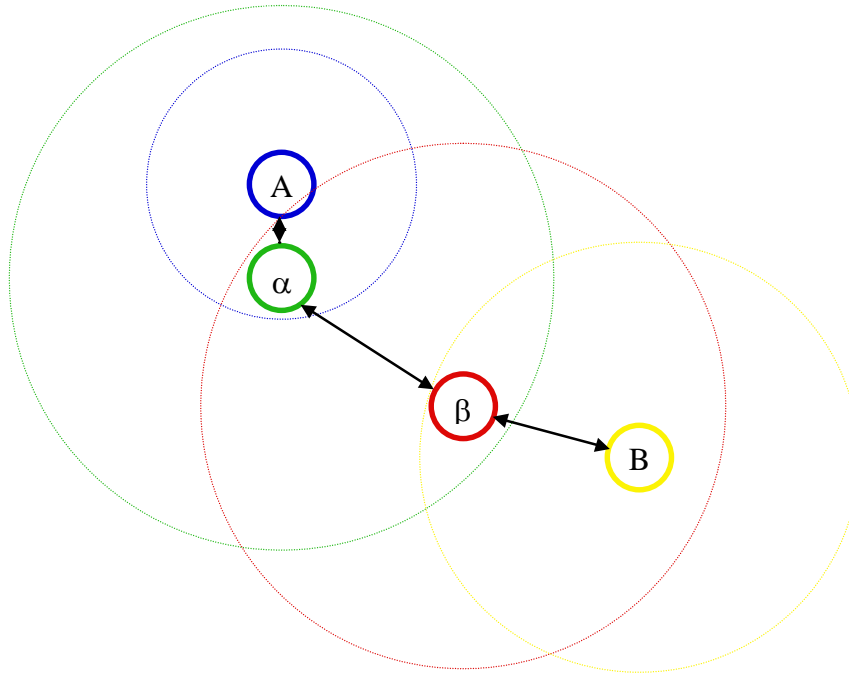


Figure 1: An idealised ad hoc network, showing intermediate links between A and B. The dotted lines represent the radio coverage area of the correspondingly coloured node.

Ad hoc networking clearly requires that the nodes in the system participate in the routing protocol, which must also be highly dynamic, since all nodes in the system may be moving relative to each other, the real radio coverage is not actually circular, and there will be considerable heterogeneity in the system if the system is of any scale. As nodes are destroyed or their power is exhausted, there is a need to reconfigure the system. The need to support different types or different priorities of traffic may also necessitate reconfiguration. Thus, although ad hoc networks are highly survivable and rapidly deployable, since they do not rely on the existence of any particular infrastructure (and they are highly valued by the military for these very reasons), they are challenging to deal with.

In our disaster scenario, we assume the existence of a fixed but fragmented infrastructure. This is not unrealistic; some switching centres have the capabilities to generate power practically indefinitely, with multiple redundant links. Others will either have limited generation capabilities or none at all. Also, the location of fibre breaks is difficult to predict, and it is not always clear where fibres run geographically, so pre-planning is somewhat hit and miss. It is, however, possible to deploy infrastructure dynamically after a failure – one could deliberately place wireless nodes to bridge between areas of fixed connectivity, one could drop sensors and bridges from planes, or one could deploy balloons or UAVs¹ to obtain the height to provide wider area coverage, one could use passing vehicles to collect information from sensors for later delivery, or to distribute new components, new functionality, or new policy information into disconnected parts of the network.

It is a mistake to assume that all traffic is necessarily synchronous. So, for example, there may be a need to support direct voice or video communications but, equally, delay tolerant traffic may be of significant value. Assume, for example, a CCTV camera equipped with a wireless interface uploading a static image to a passing UAV or support vehicle that downloads that to the surviving fixed infrastructure the next time it passes any. It may take minutes for the picture to reach a control centre, but this is still information of significant value, especially compared to sending in a human observer to obtain the same information.

Addressability is an interesting issue. The network configuration of an emergency net dynamically patched together out of surviving infrastructure may differ substantially from that of the normal operation of nodes. Determining to which sensor/actuator we are talking needs, therefore, several problems to be solved:

- Sensor/actuator self identification (capabilities, geographic location, etc).
- Address generation (CGA, ... etc?).

¹ There is a very interesting research problem in automatic routing of a UAV in such a scenario. But that's outside the scope of RUNES.

- The dynamic establishment and placement of gateways between different networking technologies (which may require code loading).
- Geocast is likely to be a critical feature.

Labelling of data is covered by an existing IETF activity and several other standards activities in the emergency communications space. The IETF IEPREP working group has several RFCs that are of direct relevance. But this is probably not enough by itself –specifying and ensuring particular quality guarantees for the timeliness, accuracy and confidentiality of the information is a hard problem.

3.3.2 *Sensors and actuators*

RUNES is an embedded systems project. It is predicated on the assumption that networked embedded systems will become wide spread and very common. In our scenario, we presume the existence of infrastructure that has been pre-placed specifically in order to be useful in case of disaster, along with the existence of infrastructure with an alternative primary use.

So, for example, the Rion-Antirion Bridge in Greece has some 300 sensors spread throughout its construction.

“These sensors include strain gauges on gussets to keep track of framework fatigue, displacement transducers in the stay cables that monitor how the bridge blows in the wind, and three dimensional accelerometers in the bridge roadway to measure the impact of earthquakes... Four acquisition units, one in each pier, collect the sensor data. These linked units are also connected to a central office near the bridge and via the Internet to the operating company near Athens and Advitam’s offices in France. Each of the four units can work independently if the connection between them breaks during an earthquake.” [5]

Such instrumentation could be applied to buildings. Thus, one could assume that future building regulations requires new buildings to be instrumented throughout with cheap strain gauges that could report the condition of the building without requiring a manual inspection.

One can also take advantage of sensors present for other primary purposes. For example, a sensor in a water meter can be used to determine whether there is water pressure at its location. A collection of such meters spread across a city could be used to map the location of breaks (and likewise for electricity and gas utilities). Another example are CCTV cameras (they are becoming very common, for example there will soon be 9,000 CCTV cameras on the London Underground). There are already initiatives in progress to allow image-processing software to be linked to such cameras in order to allow abnormal crowd behaviour to be detected [6]. It is not a huge step to get to the view that many cameras spread throughout a city will have the capability to download image processing functionality that could look for crowds, or fires, or smoke, or movement, or whatever. And we are not limited to cameras here. There are many PIR detectors for motion (people) detection, many thermostats that detect temperature, atmospheric sensors, - the list is pretty endless.

One key problem here is that different sensors may report the same information in different ways – either in entirely different forms, or with different degrees of fidelity, or some faithfully and some erroneously. Local composition of data, data aggregation and error correction are all key aspects of turning the data into useful information. This may require the distribution and installation of appropriate filtering functionality locally in sensors.

This is part of the RUNES vision – all embedded systems, through the middleware they contain, will be able to adapt their processing capabilities in response to external control. If we can realise this vision we have a huge potential for information gathering, which can be especially valuable in a disaster scenario. The abundance of raw data creates new problems however, calling for a need to filter information, preferably close to its source, so that the most effective use of scarce networking resources is made and so that human decision makers are not overwhelmed by information.

Local filtering in itself offers some relief, but it is also possible that the networking system and transport protocol has to be adapted to make more efficient use of network resources. For example, DCCP is a minimal general purpose transport-layer protocol providing only two core functions: (i) the establishment, maintenance and teardown of an unreliable packet flow (ii) congestion control of that packet flow. [7], but we may wish to examine other schemes – TFRC, etc. Given our approach to middleware, which we explicitly assume to be cross-layered, we have the ability to switch transport protocols – so, for example, if there is a VoIP-over-UDP connection and conditions deteriorate, we may wish to tune the stack so as to put in DCCP in place of UDP. This may necessitate a discovery process in order to locate the appropriate components, since we cannot assume that every such device is pre-loaded with components implementing all such protocols.

This can get arbitrarily complex – information about which network nodes are likely to survive is as important as the information they are transporting. Planning which traffic to prioritise – in short, which queuing disciplines to use in (active) routers – is dependent on a range of factors. The actual information content (or rather its labelling, as in RFCs 3689, 3690), the source of that information, the route through which that information reaches us (and the likely survival properties of it), the route we have to the destination (whether direct or DTN), and so forth. But the filtering is dynamic, under the control of the relevant authorities to meet some higher-level purpose, and the code (middleware components) to perform the filtering must be distributed in the most effective way through the same network. All of this represents one hell of an optimisation/control problem.

In all of these cases, the data obtained from the sensor net can be logged and used to improve computational models of how cities react to earthquakes or even the conditions under which earthquakes are most likely to occur.

Finally, one of the major questions in this space is that of *security*, which we only touch on here. The ability to distribute arbitrary code to nodes that are not normally under the direct control of emergency authorities is a sensitive issue. As is the ability arbitrarily to collect data from the sensors and control the actuators that one normally regards as being personal. The whole issue of overlay networks comes up in this context, both as a means of achieving appropriate security and as a means of isolating different types of information or different types of network. But the relationship to the actual underlying network is an important one in ensuring efficient use of resources.

3.3.3 Usability

Usability is an interesting issue in this space. We will assume several different types of users requiring access to information. There are those in command centres, either centrally located or regionally dispersed who require broad overviews of the situation, with the ability to refine details if necessary. These teams may well be equipped with high-resolution displays or even 3D visualisations of the situation².

There are also teams on the ground, with poorer interfaces, who require access to geographically immediate information about the conditions of buildings, the position of casualties, the availability of water, and so forth. In all such cases, there is a need to give these people effective access to information, control over systems and an understanding of how their actions will affect things.

Users on the ground come from many agencies with different primary functions, different ways of working, different interaction paradigms and different levels of importance to the overall recovery plan. Also, data entry and information reception is undertaken in highly stressful conditions, possibly in conditions of information overload, in which the accuracy of that information (and the accuracy of the perception of the information) is a life and death issue.

It is a challenging task to present information to all these different people, in such a way as to facilitate their decision making, whilst satisfying overall goals – the same information may need to be presented in different ways to different agencies. Also, if we have two potential sources of control, central and local, then there is clearly the possibility of conflicting requirements. How this is resolved is partly an autonomic optimisation problem and partly through a higher-level social control mechanism (that might be enforced by appropriate tagging of policy). Taken together, it is hard to visualise and tougher constraints for usability.

3.4 Use case description: asynchronous data migration

The scenario outlined so far has a huge number of facets to explore. We're going to write just one in more detail. However, it is clearly possible to expand the above storyline in many different ways.

The basic idea revolves around the use of electronic 'post-it' notes (known as stick-e notes in [8]), which are attached to particular physical locations³. In effect, such post-its are a form of geocast asynchronous message – they are asynchronous because there is no requirement for the recipient to be present at the same time as the sender, and they are geocast because the information is only useful within a given geographic area. This is a small but interesting case in the overall picture.

One might have a search team looking for trapped individuals stick virtual post-it notes on locations where they believe people to be, and continue on their way. Rescue teams following in their wake can read the notes and act on them⁴. Alternatively such notes could be automatically and locally generated, based on sensor data – a building could tag itself with information about whether it was likely to collapse. The major problem in terms of networking here is that we do not assume any fixed infrastructure. Naively, it could be argued that one should pass all the information entered by a user through the ad hoc network to some central point, which would then distribute the information back out again when someone wished to access it. However, this approach has a number of severe drawbacks:

- There is no guarantee that the central point which has been chosen will remain in contact with hosts in the tagged area if it is more than a single hop away, since the intervening nodes may move. The optimal strategy as far as this point is concerned is to hold the information on a node, which is as well connected to the target area as possible and likely to remain so for the longest period. This probably means that one should select a host close to the target area and hand off the information to other hosts dynamically as the network topology changes.

² An interesting avenue to explore here would be the combination of immersive 3D (CAVE) technology with the visualisation of disaster scenarios as derived from dynamic information obtained from sensors. At present, at UCL, we already have a 3D interactive map of London, on which one could overlay, for example, a view of where uncontaminated water could be found, or where there was looting, or whatever. Strategists could examine this in a more intuitive way, and, potentially, could interact with the systems too – closing valves on leaking water mains by activating actuators as a result of 'clicking' on an appropriate location. Ian Wakeman, has pointed out recently that this is like Gelernter's Mirror Worlds. George Roussos pointed at 'reality mining', which is highly appropriate [9].

³ Note that the concept of 'post-it' or 'stick-e' note is generic. One possible implementation could perfectly well be in terms of a (multimedia) web page.

⁴ Clearly, some idea of where they are, and some central planning of routes might necessitate the transmission of this same information back to management points.

- It is wasteful of bandwidth and battery power to send all information to the central node. As far as possible, and to conserve bandwidth and battery power for information, which is critical to the overall rescue plan, messages and information, which can be dealt with locally should be dealt with locally. In other words, there should be quite a large measure of autonomy in the tactical activities of separate teams but this should not preclude the ability of central authorities to determine an appropriate strategy. Again, this all argues against the transfer of information to a central point.

It would seem, therefore, that it may be better for significant amounts of information to be stored and dealt with physically close to the area to which they relate. In reality, there are a whole series of extensions to this basic idea. In the following, we examine a number of different scenarios, of increasing complexity. However, it should be noted that the list is *not* hierarchical (in the sense that levels 1 and 4 might apply without 2 or 3).

1. Initially, assume that the information⁵ is passive (i.e. it does not change). We associate a 'postit' with particular locations such that these can be queried, but this must be done explicitly. Within this scenario, there are a number of different possibilities:
 - a. The postit resides with a particular person, device or service and migrates when they do.
 - b. The postit resides at a particular fixed geographic location (or region) outside the system. Since there is an ad hoc network, the postit cannot simply reside on the fixed host closest to the particular location. Instead, the postit must migrate between mobile nodes in such a way as to maximise the probability that it will be accessible to any machine needing to access that information from that region.

There are a number of competing factors that must be taken into consideration here.

- i. *The location, velocity vectors and connectivity of the host(s) to which the information is migrated.* This is vital to ensure that the information does not become disconnected from the location because there is no route to a more local machine. In the event that such a situation does occur, there must be a recovery mechanism, which operates when reconnection occurs. For example, one might migrate information within the cloud to the point at which reconnection is most likely. However, predictions can be wrong, and the longer the disconnection, the more wrong they are likely to be. Thus, if reconnection happens through a distant part of the disconnected cloud, then the machine holding the information must eventually be notified.
 - ii. *The battery power, computational load, and imminence of shutdown of the machines to which the information is migrated.* This is to ensure that the information is not completely lost from the system.
 - iii. *The priority of the information, in terms of the number of replicas necessary to maintain its availability.* The notion of priority is discussed below. However, if information is replicated, it is necessary to decide where it is placed, how it is migrated, and what coordination must take place between replicas. Replicated information could be placed outside the area where it applies if this were to mean that there was a low likelihood of it becoming inaccessible.
- c. The postit lives with a (moving) body external to the system, where the location of the body can be determined or estimated by the system. For example, an earthquake has disrupted a chemical works such that there is a poison gas cloud. The postit moves in front of the cloud, the direction of which is determined by a mixture of estimation and monitoring of sensors. The postit can be used to give information about the cloud, or could be coupled with an event monitoring system to give early warning. In the latter case, we have a useful mechanism because only those people that require it are evacuated from the path of the cloud; evacuating more people than is required is wasteful of organisational resources and adds to confusion. In this case the timeliness of routing information, or the ability accurately to estimate precise geographical location is vital, especially if the net is relatively poorly connected.
2. Information changes over time; in other words it is active rather than passive and we have something akin to agents. All of the above scenarios are still valid. For example, we could have a situation in which a chemical agent is released into the groundwater and, over time and distance, its potency declines, requiring less vigorous action to be taken to combat its effects. At the time the contamination is spotted, it is tagged as in 1(c) and, when its potency has been assessed, a message is sent after the tag which turns the information into an active structure which knows about the potency and decay patterns of the contaminant and, when read, simply tells the emergency services what action to take.

Locating the postits after they have migrated away could be a non-trivial problem, but restricting the situation to one in which the contents of a postit are fixed is unrealistic, since information may take time to compute or

⁵ There's a sort of implicit assumption that the 'information' being talked about here is simple text. That need not be the case. I have a seriously wacky idea about the information in this scenario being in the form of small chunks of mirror world-like info that can be used in enhanced reality applications on a localised geographic basis.

may involve human input. Waiting until the information is available is dangerous; it is preferable to have some form of warning (assume fail safe in the case of an unknown), even if limited information can be gleaned from it. On the other hand, accurate information may save resources by reducing overreaction.

3. The information could contain directives about who should see it. A simple case is where it is to be delivered to a predefined group (e.g. only chief fire officers). However, it could be a little less specific than this; for example, the information should always reside with the highest ranking officer in the area. Amongst other things, security is highly non trivial in this case.
4. Events could be notified to users based on meta-information, rather than the simple contents of a postit. Meta-information might include: the creation, deletion, arrival, departure, or change in postit related information, including their contents and context. The location of people, services, postits etc. could also be used as guards in any of the above scenarios. For example, one might wish to trigger the visibility of information or an alarm only when event x occurs within time t of information of the form I(a) being generated.
5. Information can be prioritised. For example, we simply do not wish to know that there's a horse on the loose when a chemical works has just exploded; there are more important uses of the bandwidth. The survivability of information, in terms of the guarantees you try to ensure for its remaining at a given location, or in terms of the number and location of replicas with adequate battery power used to hold it might also need to be specified.

Assuming that all actions are logged, timestamped and geostamped by each machine, one could develop a tool that would allow the actions to be replayed after the event for analysis by experts (or for the training of novices). This would allow much better post-hoc analysis of the ways in which information was exchanged, what happened to incorrect information, whether the UI was good enough, how quickly people reacted, and whether the underlying protocols were good enough. Disaster relief organisations could learn lots from this sort of analysis, making them better prepared for the next scenario. It should be possible to input this information into the simulated system so that the effects of changing the underlying protocols could be assessed.

One could also use information captured from sensors to assess one's model of the effect that earthquakes had on the city's infrastructure (and so design better buildings, etc.).

3.5 Summary

The disaster relief scenario is a very powerful one, for which there are agencies that may potentially be interested in the results and existing standards on which to build. And yet it is relatively futuristic too: the networked embedded infrastructure this assumes simply does not yet exist. It lacks a business focus: this is civil defence, and is there to stop people dying rather than to make companies significant amounts of money. It cannot really have much of a business focus, mainly because much of the power of the scenario comes from the ability to make use of infrastructure that already exists for other purposes in achieving our goals.

3.6 References

1. U.S. Geological Survey, <http://www.usgs.gov/>
2. IAA System ("I Am Alive"): The Experiences of the Internet Disaster Drills, http://www.isoc.org/inet2000/cdproceedings/81/81_3.htm
3. Caltech USGS Broadcast of Earthquakes – CUBE, <http://www.gps.caltech.edu/~bryant/cube.html>
4. Motorola responds to emergency with 86 truckloads of communications gear, http://september11.mrtmag.com/ar/radio_motorola_reponds_emergency_2/
5. IEEE Spectrum July '04 <http://ieeexplore.ieee.org/iel5/6/29070/01309800.pdf>
6. <http://www.newscientist.com/article.ns?id=dn3918>
7. Datagram Congestion Control Protocol, <http://www.ietf.org/html.charters/dccp-charter.html>
8. P. Brown. Triggering information by context. Personal Technologies, 2(1):1--9, Mar. 1998.
9. **Reality Mining:** Browsing Reality with Sensor Networks, <http://sensormag.com/articles/0904/14/main.shtml>

4 Integrated Wine Production and Distribution Scenario

4.1 Background

Wine is a global commodity of great value, with a worldwide market in 2002 on the order of 150 billion euros. The market is dynamic with new competitors entering from all parts of the globe, and with intense cost pressure on commodity wines. However, it is a market with complex quality requirements many times driven by the brand and control of the wine itself.

In order to compete in a global market, winemakers must not only understand the best production methods, techniques, but also the most efficient distribution mechanism for the wine. Distribution has a tremendous impact on the quality of the wine itself as well as the ability of the wine distributors to be confident in the product arriving in the best condition at the customer's location.

There are systems that address the automated collection and management of production parameters. These systems are often termed wine production systems and they manage the information around grape acquisition and tracking, barrel management and laboratory analyses of the wine samples as well as the actual cellar operations such as composition, calibration of tanks, etc.

The production system vendors often add distribution in the form of tracking of orders and logistics as a separate modules left outside the production details. A separate set of modules or entire applications are then required to manage the distribution of the wine and the inherent quality attributes of the wine itself are often lost in the distribution phase. Like many traditional distribution systems, the wine tracking and distribution method is focused on bar-coded lots of inventory and the manual tracking via standard warehouse management procedures of this inventory.


4.2 Analysis and Motivation

European production and distribution of wine is a significant and commercially important segment of the EU economy. In order to compete in the global marketplace, improvements and integration between the production of the wine itself and the distribution system will ensure a higher quality product delivered accurately to distributors. The distribution network in turn will be able to deliver a known product to their customers. Currently, there is a separation and series of manual process characteristic of the traditional supply chain processes. *The purpose of this scenario is to use reconfigurable, mobile networking of sensor technologies to unite the production and distribution through integration and automated information gathering.*

Embedded systems are used today in all phases of wine production. In agriculture, specific measurements of light, moisture and heat can be made at the vineyard for automated grape characterization and labeling. This information is important to the winemaking and is an integral part of current wine production systems. By using sensors in the field, the vineyard can start the process of establishing the vineyard profile and yield projection, and generating assessments of the grapes for the winemaker's benefit. The information collected can be transmitted to the winemaker and into the winemaking process where further sensor information can be used within the system to record important parameters. The creation of early stage inventory and identification will allow the tracking of products through the distribution phase.

Devising new identification mechanisms and integrating those with intelligent distribution vehicles means that the inventory can essentially track and report its own status and movement. This automated inventory tracking means that various inventory vessels can report their location, status and environmental conditions through the process right into the shops of the end consumer. Computing systems already support wine histories but such intelligent architectures which actually publish their own progress will reduce the costly inventory acquisition and manual handling of vessels [2].

Through the addition of RUNES technology, at each phase of the production and distribution process the quality parameters can be reported on the wine thus providing a complete automation of quality assessment through the wine's lifecycle. Alerting systems for temperature in transit can be made and vehicles can adjust their condition, route or other parameters in order to improve the environmental situation towards a positive quality goal. This integration goal is depicted in the diagram below:



Through the creation, integration and implementation of a reconfigurable embedded system with smart deliver tracking and control, the quality control of the winemaking process could feasibly be managed right into the retail environment. This can also extend to the intrinsic brand management as well ensuring that the quality and the brand of the wine is not altered or somehow reduced throughout the entire process. In fact, through the use of publish and subscribe mechanism, the wine vessel itself can report its progress through the process along with the relevant information about the journey along the way.

4.3 Communications and Technical Implications

Through the development of a reconfigurable wine production system based on the work of the Eurojenet WINE team, the RUNES technology will integrate a new wine production embedded system with the intelligent delivery system unifying these two key aspects of wine production. The purpose of this integration will be to control quality and brand identification through automated reporting from the product itself as opposed to the current system of barcode scanning alongside manual identification and tracking.

The requirement is to set up a wireless sensor/actuator network in order to monitor specific parameters in wine transportation units. This system would be augmented with RUNES technology to allow the wine supervisor to manage environmental parameters (e.g. temperature, humidity etc) of wine containers from a remote control room. This ought to be done by remotely changing environmental conditions through actuating temperature and humidity controllers of wine containers within the vehicles. The developed system should allow the acquisition of other parameters (e.g. pH, CO₂, humidity etc) through the auxiliary analog channels available on each vehicle published by the local wireless sensor network. The primary objective is to increase the quality of storage conditions during all stages of transportation from production to retail stores of wine and avoid interference with product quality by unforeseen temperature or humidity variances. This should also provide early warning about damaged wine units during distribution. The proactive remote actuation of temperature and humidity controllers within the delivery vehicles helps in reduction of waste and damage. It should be recognized that the wine distribution chain extends over multiple sites where a single producer provides multiple distribution centers that may deliver to different retail clients. The system needs to provide clients with early warnings of possible distribution chain ruptures as well as comprehensive automatic on-line reporting of conditions within wine containers and locations of the vehicles.

To make such a system economically feasible, a wireless approach needs to be developed that offers a set of features, in particular: progressive introduction of the technology that is applicable in both static and mobile environments (transport and delivery), low installation and upgrading cost, self configuring wireless networks and power efficient and extensible sensor/actuator nodes.

There are several implications of the requirements outlined above. First, there is a need to develop an internal self organizing wireless network of sensor/actuator nodes within each delivery vehicle. The sensor units monitor

temperature, humidity of the wine container units and perform some in-vehicle sensor fusion to derive higher level information and context. A local node within each vehicle with the same radio interface as the local wireless sensor network and additional global wide area radio interface forms the gateway between the local sensor/actuator network and the remote quality control and monitoring supervisor. Multiple wireless communication channels should be supported with support for varying quality of service and cost attributes. The remote web based monitoring service provides on-line image of all sensor data in each vehicle, each vehicle's location and speed within the map and provides an interface for remote controlling of temperature and humidity of wine containers within vehicles.

The second technological implication is that the fleet of distribution vehicles augmented with RUNES technology would provide a smart delivery system. A smart delivery system may range from simple transportation systems that provide its location, navigation route and estimated delivery time depending on traffic conditions to a complex transportation system in which remote cooperating vehicles form an optimal dynamic delivery chain. In the latter case, vehicles can advertise spare capacity, current storage conditions as well as their routes. The fleet manager can look-up advertised routes and storage conditions to dynamically negotiate contracts with vehicles to deliver wine units to a certain retail destination. The task can be supported by 'smart wine containers' which publish information about the destination, the sender, urgency, deadline and environmental quality attributes required (temperature, humidity etc.) of the Wine product. These smart wine containers may provide input to determine the kind of transportation required and optimal delivery route. The RUNES technology could provide a flexible and powerful solution to monitor and control the process at different levels of the logistical chain of the Wine production and distribution system. These sensor/actuator nodes need to be self-configuring and be plug-and-play to simplify their use and minimize costs. The wireless technology which underpins most levels of networking avoids structural modifications of the existing Wine production and delivery systems and avoids the complexity of wiring and static and inflexible deployments of wired networks. It also contributes to the robustness of day to day operation in harsh environments. The technology must adapt to changing conditions over time without complete re-installation or re-deployment by providing a flexible middleware architecture with support for smooth evolution of the system. Essentially, the quality of the final Wine product depends on Wine production and distribution chain. RUNES technology not only aims to provide cost effective quality control mechanisms in production but also in distribution.

4.4 Summary

The European wine industry is a key economic force in the region. The critical components of this industry are the clearly identifiable quality of the product produced and the cost-effective distribution of the product into the hands of the customer. Through the use of novel and reconfigurable embedded systems, both the wine production phase and the distribution phase can benefit from an intelligent, self-reporting and actuating system that manages and controls the quality and identification throughout the product lifecycle. RUNES technology can be directly applied to previous attempts in this area, particularly in addressing the needs of existing infrastructures and in the conversion of traditional wineries to a new embedded systems approach. Through the use of flexible, reconfigurable networks of sensors and actuators coupled with the sensor based control of distribution, the wine quality tracking and management process can be made more efficient and ultimately self-reporting. Through the use of intuitive user interfaces and interactions which are context sensitive to the part of the process that is being managed, the difficulties of using this type of technology within a traditional winemaking environment will be greatly reduced. By leveraging on the wired use case implemented through the Eurojet project, the RUNES technology will be able to advance the current state without having to address the fundamentals of each component of the process. The conversion of the wired WINE demonstrator to a wireless, reconfigurable network using the RUNES technology also provides further important evidence as to the utility of the resulting outcomes. This scenario addresses the question concerning the conversion of existing computing infrastructures to the new RUNES technology, through the implementation of this scenario this question will be further explored and conversion issues will be addressed.

4.5 References

1. http://www.rabobank.com/Images/rabobank_publication_wine_is_business_2002_tcm25-156.pdf
2. <http://www.eurojet.com> – Wired Internet Networking Embedded User Experiment, IST 2000 28422, Cantine Giorgio Lundarotti srl.
3. <http://www.lungarotti.it/english/ind4.htm>

5 Automotive Scenario

5.1 Background

Recently we have observed an incredible development of high technology infrastructure for transportation in a variety of countries. Many roads and highways are equipped with sensors and actuators to monitor the traffic situation and make transportation more secure. Many highways are equipped with traffic sensors and cameras to monitor the traffic situation, or to aid in toll collection. Sensors are deployed in tunnels to monitor possible accidents. All those sensors are connected to traffic control systems to provide safer driving condition and optimize traffic flux. Simple examples are many traffic lights connected to magnetic sensors under the asphalt. Those sense where there are cars waiting at the intersection. The traffic light can then schedule the traffic based on this information instead of using a simple time based algorithm. Sensors are also used in some intersections to catch cars passing when the signal is red, and to trigger recording the traffic offender. This last example show how the high tech infrastructure can be used to dissuade drivers from engaging in dangerous behaviors, and thereby, enhancing the safety of our roads.

Modern automobiles also contain very complex networks of embedded systems. In fact a luxury car today can have more than 5 different wired communication buses and 80 electronic control units [1]. For this reason, a car can be seen as a sensor network embedded in a larger complex sensor network. We will describe a couple of scenarios where we will connect the sensor network present on the roads and the networks present in the car by means of wireless connections.

Therefore, since modern cars are composed of networks of embedded reactive systems and roads are instrumented with sensor networks, the resultant network of sensor networks created by the connection of cars to this larger system is a very good test bed to experiment within the context of RUNES.

Software is becoming one of the driving factors of innovation in the automotive industry. This is shown by the fact that 40% of the cost of a vehicle is determined by software [1]. The extension of those systems from an isolated network to a complex global network of networks will allow new and useful services to be provided. However, this evolution will also create new challenges for development of software for those systems. For this reason new approaches and middleware are required to lower the complexity of such systems and allow for reuse of code.

We will first describe an “Accident Mitigation” scenario where interactions between sensors within the car, the instrumented roads and the infotainment system of the car is exploited to reduce the risk associated with unexpected accidents that can distract drivers. Second we will present an “Assisted Driving” scenario where the internal sensors of the car, the instrumented traffic signals on the street, and the cruise control system within each car will interact to provide a safer, smoother and less stressful driving experience.

This shift from wired isolated networks to more complex wireless enabled systems that we propose has already started in the automotive industry. Almost all modern cars, for instance, have some kind of remote control for locking and unlocking the car. Many luxury cars include some wireless phone capabilities, sometimes coupled with a GPS based navigation system. Wireless connections are becoming normal in many vehicles. Internal local networks based on wireless communication standards such as Bluetooth are expected to be available for infotainment and non safety critical services in many future generation cars. This will allow interaction of portable digital devices with the infrastructure provided by the vehicle. For example it is already possible and relatively inexpensive to equip an automobile with a Bluetooth based speakerphone. Using the Bluetooth network the multimedia system of the vehicle is used to place phone calls.

Via this automotive scenario, we explore and validate our hypothesis that, in a near future, wireless connectivity in cars will expand beyond the boundaries of the car itself. Ad-hoc wireless connections, with other cars and with intelligent sensors and other embedded wireless devices deployed on roads and highways, will allow for new services, greatly increasing the safety of our vehicles.

Our assumption of having in the near future cars connected to an ad-hoc wireless network is based on the observation of new research projects announced by leading car companies. For example BMW, in the context of its ConnectedDrive concept, has proposed an application of car-to-car communication. Information about safety condition of the street are collected by car sensors and transmitted to other approaching vehicles. The ABS sensors could, for instance, alert drivers of nearby cars of a slippery road. Furthermore, a consortium having important European and U.S. car manufacturer as members (including AUDI, BMW, DaimlerChrysler, FIAT, Renault, and Volkswagen) has been created to establish standards for ad-hoc wireless communication between cars. The CAR 2 CAR communication consortium aims to establish an IEEE 802.11 based communication between cars and to use it for active safety applications [3].

5.2 Analysis and motivation

A single car with its 80 embedded electronic control units and its 5 wired networks is already a complex system. The evolution toward wireless interaction with the external environment increases the complexity and requires the framework to be highly scalable. A crowded city with millions of cars moving around during rush hour can be seen as a complex ad-hoc wireless network with several millions embedded devices connected.

Moreover, different quality of service requirements will be present for different types of communications. Most of the real-time message exchanges will be very likely handled by traditional wired communication buses like the CAN bus present in most cars. Some wireless messages, however, can present QoS requirements.




Figure 1 – Example of application of wireless technology to improve the safety of cars. The crash sensors of the two cars involved in the accident send a wireless alert to a wireless device in the traffic sign. Incoming cars are alerted of the possible danger around the corner.

A complex system like the one we depict needs to have an effective security infrastructure to avoid possible problems due to errors or malicious attacks from hackers. What if the city traffic got paralyzed because of some virus that disabled the engine of all vehicles (or even worst their brakes)? This scenario will allow us to experiment with a very complex, heterogenous, distributed system. It requires scalability, QoS, and security to be analyzed and addressed correctly.

The automotive industry is an important sector of the European economy, with European countries and the United States as world leaders in many segments. European car companies are leading the research for applying new technologies in automobiles. The introduction of a new standard architecture for networked embedded devices can be very beneficial to them. The ability to wirelessly interact with the external environment will allow new security features to be added to the car, without scarifying the current safety features. Furthermore, the use of a standard architecture provided by RUNES will allow third part suppliers to provide new interesting services for vehicles. For example, application like RHODES [5] would benefit from a standard middleware that allowed access to a deployed infrastructure. This will allow for a new market segment to be created, where companies do business providing car services and software. This will be similar to what happened in the mobile telephone industry. Thanks to standards and flexible software and systems architectures, many companies are able to provide add-on features to phones. The evolution toward a common set of standards for hardware and software components in efforts such as AutoSAR [2] indicates the will of car companies to open vehicles to external services and components.

5.3 Use Case Scenario: Accident Mitigation

Accidents are not only a major cause of traffic congestion, delays, and frustration for most drivers, but also a constant danger for everybody on the road. Unexpected accidents, for example, are a major distraction for drivers. The safety of all on the road is compromised significantly as curious drivers are straining, or “rubbernecking,” to look at the accident and discover its effects. Driver distraction is a main source of accidents,

and researchers, such as Professor Mohan Trivedi at UCSD, are conducting research to address this problem [4]. The ability to use wireless connections between cars to increase the awareness of drivers will help in alleviating the sources of distraction.

A network of sensors embedded in the automobile and dispersed on the roads can be employed to maximize the safety, security, and comfort of drivers during an accident. Devices embedded within a car might include sensors that measure its distance to nearby objects, sensors that measure its speed and acceleration, sensors that detect a crash, and sensors that measure temperature, carbon dioxide, and activity within the car. Sensors within the car can be wired to increase the reliability for life-critical applications. Sensors outside the vehicle can be dispersed along road signs, such as speed limit signs, freeway signs, exit signs, etc. These external sensors connected by less dependable wireless connections can optionally increase the safety of passengers, without overriding the existing functionality of reliable wired connections. Since a large number of road signs exist, deploying sensors on these signs requires that the sensors be scalable, easily deployable, and power-aware.

In the event of an accident, the cars involved in the accident need to emit a signal to indicate that an accident has occurred, and the cars on the road nearby need to be informed that an accident has occurred. The cars involved in the accident must detect that a crash has occurred, send signals to the antilock braking system, the airbag deployment system, and the central locking system, as well as to the external road network. All of these signals have a strict real-time requirement as well as a reliability requirement, because lives are on the stake. Furthermore, these signals must be very secure, so that an adversary cannot generate false signals. The signal to the external road network must relay the messages to all nearby cars as well as to the police and fire departments. The cars on the road near an accident form an ad-hoc, location-aware network via the transmission of temporarily relevant messages.

In addition to accidents on the road nearby, a driver might also like to be informed about road blocks, construction zones, dangerous road conditions, and natural disasters. Thus, an efficient, effective, and unobtrusive mechanism to notify the driver of all such events must be designed. Furthermore, at any given time, many notifications might need to be relayed to the driver, and a prioritization method must be implemented. Several research projects have been conducted at UCSD on intelligent car systems. The car should be able to understand the intention of the driver to be able to better react and provide help in dangerous situations [4].

5.4 Use Case Scenario: Assisted Driving

Sensors deployed throughout cars and roads can communicate with each other to provide increased automated assistance while driving. For example, speed and acceleration sensors within the car can communicate with sensors on speed limit signs to automatically enforce traffic laws. Furthermore, proximity sensors can detect the speed and acceleration of nearby objects in order to ensure that a safe distance is always kept from any other objects on the road. The sensors within the car can coordinate with traffic lights to provide more effective flow of traffic and to minimize the amount of congestion. Thus, the integration of the maximum allowable speed with the driver's desired speed and with the speed of traffic will provide a safer driving environment and permit some amount of automation in driving, even during traffic. Of course, these sensing devices, and any amount of automated driving they provide, can be overridden manually by a driver. Also, there must be a mechanism to alert the driver and passengers of any dangerous driving conditions, such as a sudden decrease in speed. Again, the sensors and actuators in these wireless ad hoc networks only add safety, reliability, and comfort to the driver, and should not override the life-critical features provided by the wired buses on existing cars.

The sensors involved in these types of wireless ad hoc networks can aid the police department in monitoring traffic and enforcing laws. Additionally, these networks of sensors can be used to provide optimal navigation for the drivers. The navigation system can generate a route from source to destination, by taking into account the average speed of traffic, the number of accidents, and the current road conditions. This advanced navigation system can help drivers with both temporary road conditions, such as an accident, as well as with temporary dangerous driving conditions, such as a blind curve or a cliff. This system can be further advanced by having each car periodically broadcast a message indicating its current position. This message can be used to enhance the functionality of the proximity sensor, especially around a curve, as well as to provide mechanisms for retrieving a stolen car. Thus, the widespread deployment of sensors, both inside and outside the car, can augment the safety, security, reliability, and comfort of drivers by assisting in tire operation of the vehicle.

5.5 Communications and Technical Implications

The scenario described above has several important implications for the characteristics of the underlying network infrastructure.

First of all, many diverse types of sensors are involved in the scenario, and they must seamlessly interoperate in order to provide a safe, secure, and comfortable driving environment. There are physical sensors that measure distance, speed, acceleration, and force, among other things. There are actuators that sense information relayed across some network, such as the occurrence of an accident, and perform some action, such as the reduction of speed. There are relaying sensors, especially along the roads, that are needed to spread information.

All these types of sensors must interact and interoperate, not only with each other, but with currently existing network systems. For example, crash detection mechanisms already exist that deploy the airbag in the event of a sudden change in acceleration. Currently, these mechanisms mostly interact via the real-time CAN bus inside the car. However, in the near future, this system must interoperate with the ad hoc wireless network on the roads, in order to inform other drivers on the road that an accident has occurred. Of course the local connections with real time safety critical requirements will be handled via a wired connection; however the local interaction will trigger a wireless message. This wireless network might be in the form of radio, radar, WiFi, satellite, etc. A hard real-time requirement exists for this interaction, because the car must relay the message before all power and communications infrastructure is destroyed by the accident. The vehicle can relay the information to the closest road sign, which can then relay the information to the proper emergency response agencies, including the police department, the fire department, and ambulance. Furthermore, security is of vital importance to prevent adversaries from reeking havoc by impersonating an accident, falsely deploying emergency response, and deceiving drivers.

The interaction of heterogeneous sensors across many different networks must guarantee accuracy and privacy of information. For example, police information, such as insurance information, driving record, and criminal history of the drivers involved in the accident, must remain private, and should not be transmitted to other drivers on the road near the accident.

The networks of sensors must be scalable, and accommodate the millions of cars that travel on the roads and freeways. An accident may potentially involve over a hundred cars, and thus, all nearby network nodes must be able to handle a sudden surge in network traffic. Automated synthesis of all the accident information provided by each car involved must occur, so that redundancy of information is minimized and so that all other drivers on the road as well as the emergency response teams are notified in an effective manner. The number of other drivers on the road will depend on location and time, but could potentially be thousands of drivers, especially in a large city during rush hour traffic.

Thus, this scenario illustrates the need for networked sensors embedded within a car as well as dispersed along the road to be secure, reliable, scalable, and easily deployable, as well as able to perform under some safety-critical real-time requirements. The intrinsic unreliability of wireless connection will impose wired connections for safety critical systems inside the car. Nevertheless, soft real time requirements can be implemented by wireless connections that extend beyond the boundaries of the car.

5.6 Summary

The car industry is of key importance for European economy. Intelligent transportation systems are becoming more and more interesting, and the fact that roads are becoming instrumented sensor networks will make it possible to provide “intelligent transportation services” within the boundaries of our cars. Most of the interesting new developments in automotive systems are driven by software and electronic components. In particular many new applications scenarios requires coupling of car internal embedded devices with the external environment. This can be accomplished using wireless technologies and deploying a mix of wireless and wired embedded system in vehicles and roads.

To provide European car companies a competitive advantage over other manufacturer, new technologies and software frameworks can be explored in the context of RUNES and applied to the challenges found in this fascinating application field.

5.7 References

- [1] Automotive Software Workshop San Diego 2004. <http://sosac.ucsd.edu/aswsd/2004/>
- [2] AutoSAR. <http://www.autosar.org>
- [3] CAR 2 CAR communication consortium. <http://www.car-to-car.org>
- [4] J. McCall, M. M. Trivedi, "Performance Evaluation of a Vision Based Lane Tracker Designed for Driver Assistance Systems", CVRR Technical Report, Dec. 2004.
- [5] P. Mirchandani, Fei-Yue Wang, “RHODES to Intelligent Transportation Systems”, IEEE Intelligent Systems, pp. 10-15 January/February (Vol. 20, No. 1)

6 Fire in a road tunnel scenario

6.1 Background

6.1.1 Disaster scenarios

A disaster or emergency can take many forms. Each has its own characteristics. The challenge of this work was to select a suitable format which supplied generic aspects which could be applied to a wide range of disasters. Consider a non-exhaustive but representative list of potential disasters:

- Earthquake
- Fire
- Airborne toxin
- Water supply failure/contamination
- Flood
- Hurricane/tornado
- Vehicle crash (air, road, rail, marine)
- Explosion
- Environmental contamination
- Biological attack
- Infestation
- Volcano
- Disease
- Avalanche
- Radioactivity

Each disaster has its own characteristics:

- Scale (size, specificity)
- Infrastructure integrity
- Time - speed of disaster
- speed of response
- Range of response services
- Natural/ accident/ malicious
- Location
- Density of population
- Accessibility

For this scenario the example was taken of a fire in a road tunnel. This kind of disaster is well-known across Europe (there have been several examples in the last few years), but many aspects of the disaster can also be applied to other types of emergency.

The tunnel characteristics/issues:

- Localised
- Poor access
- Poor information on who/what involved
- Fumes, gases, fire, heat
- Structural integrity
- Traffic management
 - access for emergency services
 - reduction of disruption
- Multiagency, multinational response

A disaster-based scenario represents a large, rich and diverse opportunity to examine potential application areas for the technology. Resource limitations require that only certain aspects of the scenario are selected for further elaboration. The potential for medical uses seem the most appropriate at this time.

6.1.2 *Wireless systems in a medical environment*

There are strong drivers to promote development of telemedicine, primarily the demographic trend of an aging population. The 'demographic timebomb' which is facing Europe has already had significant effects in Japan, and as a result the adoption of wireless technology for patient monitoring there is well ahead of that in Europe and represents the state of the art. Use of telemonitoring allows the old, infirm and chronically sick to remain in their own homes as much as possible and reduces the burden on health services.

Other drivers also exist to encourage progress. The global healthcare sector has a turnover of around \$200 billion, representing a huge potential market. There is a greater emphasis on prediction and prevention than ever before and technologies which can provide monitoring and reporting with minimal human intervention can deliver a low-cost route to address this requirement.

Provision of wireless networks already offers advantages in administration and support functions in healthcare environments, for instance scheduling of appointments and availability of facilities. There is also a distinct need for more integrated records and data access, and the RUNES technology can also offer advantages in this evolution. Use of wireless technology offers the potential for reconfigurability (for example a PDA which will download and present information on a selected patient during a doctor's rounds) and improved portability of devices.

From a technological point of view, security and privacy considerations are very important. Current systems do not permit patients to be monitored wherever they go, and this would be a very desirable improvement. Extrapolation of potential use of the technology indicate that ultimately the systems will include actuators which can be used to treat the patient (e.g. insulin pumps for diabetics) but the systems are nowhere near robust enough for that to be possible at this time.

There is now a trend for systems to be developed for healthy people to be able to monitor their health, for example when exercising. If this becomes popular it will assist in development of both sensors and the networks required and will help bring down costs.

6.2 Analysis and Motivation

Recent events such as the tsunami on Boxing Day 2004 and the destruction of the World Trade Centre in 2001 have made the problems associated with dealing with the aftermath of serious disaster or emergency a focus of governments and emergency services across Europe and across the globe. This concentration on 'Homeland Security' has opened up a new market for cutting-edge technologies to address hitherto insoluble problems both in prevention and remediation. The RUNES technology is ideal for addressing issues associated with gaining an understanding of the environment resulting from a disaster or emergency, for assisting the command authorities in understanding the location and condition of people and assets in the emergency zone and for allowing tasks to be carried out to locate or rescue victims and casualties, in many cases using systems which are already present at the scene and which can be used for purposes for which they were not originally designed. Disaster applications therefore represent an interesting, new market for wireless networks and address a current global concern.

Of all the potential uses of wireless networked systems, their use in medical monitoring is probably the most advanced and widespread. European Commission funded projects such as MYHEART and HEALTHYAIMS are addressing the need for widespread body area networks (BANs) capable of sensing and reporting the condition of a subject. Commercial systems already exist, but are limited for example by carrying out the monitoring process only in a defined area (the patient's home or within a clinical environment) and there is no evidence of interoperability at this time. Desirable improvements in the functionality of wireless medical monitoring which fall into the RUNES area would therefore be to establish interoperability, to address data transfer and security concerns, to examine how monitoring might be made to work in an unconstrained environment, and to look at the implications of having actuators as well as sensors within the network.

The telemonitoring application applies as much to the potential monitoring of emergency services personnel as to the chronically sick, and could also be used in the context of monitoring casualties at the scene. This latter function would assist in the handover of casualties from front-line staff to hospitals and in triage of patients where resources are limited. In an emergency situation concerns of patient confidentiality and data security would be less pronounced. In choosing potential demonstrators for the RUNES project, the use of healthy subjects would be appropriate which avoids some of the ethical committee requirements for patient monitoring which would affect the timescales of the project.

The medical/disaster scenario combines both the developing market of security and disaster planning and the proven area of medical monitoring in an application which does not require the use of patients in the context of a demonstrator.

6.3 Use Case Descriptions

Consider a mythical road tunnel somewhere in Europe. It has been in existence for over thirty years and offers a major transport route. The amount of traffic is quite high and any alternative route is much longer so it is a preferred route for goods vehicles. Because of its age the tunnel does not have the current optimum design for safety, and as a result some types of goods are not permitted in the tunnel.

In our future scene, wireless technology can assist in the control of goods in the tunnel – all cargoes are required to carry an RFID tag with information on the contents, amount of material and hazard information. This means that cargoes which have the risk of violent explosion or other associated hazards can be refused passage or allowed through under restrictions, since load details can be accessed remotely.

The age of the tunnel means that ventilation is poor. As a result air quality monitoring equipment has been installed within the tunnel. Temperature, humidity and some gases are measured. Traffic flow can be restricted within the tunnel to ensure that air quality remains acceptable, but this can result in traffic problems around the entrance and exit. Information on air quality is made available to anyone who may be concerned about the potential hazard to their health (e.g. asthmatics) via a suitably-equipped PDA.

On this busy weekday there is significant traffic in the road tunnel. There is a collision deep within the tunnel between several vehicles including a tanker loaded with heated vegetable oil which suffers penetration of the tank and begins to leak, spreading oil all over the road surface. A small fire started as a result of the collision spreads to the oil, which begins to burn producing clouds of thick smoke as well as heat and flame.

The detection system within the tunnel picks up the fire and immediately triggers an alarm. The tunnel is closed and the emergency services are summoned. Management of the resulting traffic congestion will be started.

6.3.1 *Response of road tunnel users*

The traffic quickly stops when the accident occurs and the smoke from the fire is the first indication to the drivers of the vehicles in the tunnel of the emergency.

Virtually all the tunnel users have personal communication devices. Some people call the emergency services for help, some seek information from the local information systems of the location of safety equipment and refuges for tunnel users.

Later on, the emergency services may contact some of the people within the tunnel to obtain information on or even pictures of the scene of the accident.

Some of the tunnel users have medical problems which result in them carrying personal medical monitoring equipment. Normally, a person with a medical problem would prefer to keep this information confidential, however the nature of the emergency means that some of those involved will require priority evacuation, or the trauma of the incident might have additional consequences because of their conditions.

Commuters involved in the resulting traffic jams but not affected by the fire risk will also require information on the cause of the problems and anticipated delays or alternative routes.

6.3.2 *Response of the emergency services*

The first challenge to the emergency services will be access to the tunnel. The traffic resulting from closure of the tunnel will present a serious delay to any road vehicle. Future infrastructures may allow traffic signals to be changed to facilitate access or information from vehicle networks to identify the fastest possible route.

On arrival, the immediate priorities will be to establish the situation inside the tunnel, to rescue the people inside the tunnel and to tend to the injured.

To understand what the situation is within the tunnel, the rescuers will need to be able to use existing sensors available in ways, which their current software may not support. For example, the temperature and air quality sensors may take readings, which would normally be considered erroneous due to being too high. There may be a need to download data from vehicle on-board systems, which would not normally make data available. The emergency services therefore need to be able to download software, which will permit use of the sensors available to them in whatever way they require.

Once the situation is sufficiently understood the first rescue workers will enter the tunnel. Where the risk is high such personnel will have personal health monitoring equipment to allow the command and control outside the scene to recognise risk to their operatives.

An alternative possibility is use of mechanical devices, which can be sent into the tunnel to locate hazards such as fire or leakage of dangerous cargoes. Their sensing systems could be added to the network as required.

Some rescue workers may have sensing devices not carried by others providing data on, for example, particular noxious chemicals or readers for the RFID tags on cargoes. Rescue workers may also require temporary medical monitoring information on victims. The data generated must be made available to all workers.

Medical workers will be required to deal with victims of the tunnel disaster itself and also any injuries to the rescue workers.

Casualties with pre-existing conditions may already have some medical monitoring originally linked to their primary healthcare provider. Previously healthy people injured in the disaster may be fitted with medical monitoring devices at the scene because they are trapped, to provide historical data for when they can be evacuated, or to allow specialists to monitor their conditions. Injured rescue workers can be monitored by making their data streams available to the medical personnel in addition to command and control.

6.4 Communications and Technical Implications

The medical/disaster fusion scenario offers a multitude of possible applications for the RUNES project. The overall background is too rich for the project to be able to address all of the possible applications, but this very richness means that certain aspects can be selected for modelling which allow the system requirements identified by WP1 (see deliverable D1.2) and the external/user requirements identified by WP8 to be addressed. Given that the aim of the RUNES project is to 'enable the creation of large-scale, widely-distributed, heterogeneous networked embedded systems', the application(s) chosen should reflect these aspirations also.

6.4.1 Integration of networks

Table 1. shows the variety of sensors available to emergency services personnel to evaluate the scene provided a means exists to access the data. Actually providing that means is likely to be a challenge given the heterogeneity of the infrastructure. The sensors will be of many different types, capabilities and ages; will work on different networks; and in some cases the infrastructure may have been damaged to the point where significant reconfiguration is required to allow the network to function at all. Encapsulation of sensitive electronic components is particularly important to achieve improved reliability and extended lifecycles under harsh operating conditions.

One of the most exciting but daunting prospects is the possibility of being able to adapt or retask existing sensors by means of middleware. Although this would be an extremely useful facility for the emergency services, the security implications of this ability being available are considerable.

The problems associated with different emergency services having communications systems that do not allow them to talk to one another over their radios are already familiar. These problems must be addressed in the context of personnel who have not only radios but also personal telemonitoring, handheld sensor devices and in the case of paramedics, telemonitoring equipment for use with casualties.

In traditional distributed security, authentication is established after contacting a trusted authority responsible for maintaining up-to-date record of each user’s access rights. However, in a disaster response scenario, communication with this authority might be poor or even impossible. In such a case a best-effort security model may be more appropriate. Approaches based on public key cryptography overcome the need for a trusted authority but the computational requirements are high considering the resources available on sensor nodes. Algorithms implementing elliptic curve cryptography are believed to be more computationally efficient than the former.

In order for the surviving parts of the sensor and communication network to remain in a healthy state for enough time to allow the rescue mission to be completed successfully the nodes must be supplied by power. Replacement of batteries is often impossible or impractical and hence any solutions relying on power scavenging would be of great value. Finally, constraints in wireless communication include bandwidth shortage particularly when the unlicensed 2.4GHz frequency band is considered.

<u>Tunnel sensors</u>	<u>User sensors</u>	<u>Emergency services</u>
<p><i>Fixed network</i></p> <p>Strain Fire Traffic conditions (speed, congestion) Lighting Fans/ventilation Cameras Thermal imaging ANPR Air quality sensing (gases, temperature, humidity) Meteorology (outside tunnel) Possible sensors/comms in refuges</p>	<p><i>Mobile, reconfigurable</i></p> <p><u>Personal</u> Mobile phone or device - call for help - find out information - relay info/pix Personal medical monitoring Location Speed RFID tags</p> <p><u>Vehicle</u> Freight: vehicle type, load, load condition Location Speed AQ sensors In-cabin camera Door open sensor Presence of passengers Engine running Impact sensor Black box Inter-vehicle comms</p>	<p><i>Fixed and mobile</i></p> <p>Electronic canaries Thermal imagers Personal id/location Radio Health monitoring (personal) Handheld devices for RFID tags and other devices Location and status of other emergency personnel Environmental conditions Robust equipment to support comms</p> <p><u>Medical</u> Medical sensors for patient condition Medical equipment left with patient Sensors for exposure to toxic and other hazards Access to medical records</p>

Table 1. Sensors available at the scene in the tunnel fire scenario

6.4.2 Medical applications

Telemedicine by wireless monitoring of patients is a technology, which is already known and used. Commercial systems are available and a number of European Commission funded projects are running aiming to extend or improve the functionality of the techniques. The RUNES project can look at higher-level functions of the technology such as interconnection of networks and enabling patients to be monitored wherever they are.

In this scenario we are looking at widespread monitoring of healthy people (emergency services personnel) in order to check for the effects of heat, smoke, fumes, fatigue and injury. In many respects there is synergy with what is being developed for the chronically sick and infirm already. In addition there is scope for monitoring of casualties who are trapped or cannot be evacuated for some reason, or for triage purposes or to provide a record of health status from first contact with paramedics until arrival at hospital. Quickly identifying the most severely injured patients poses unique challenges, as does efficiently monitoring and transporting victims. Wireless vital sign monitor devices can be attached to victims to help throughout this process. By installing wireless, battery operated location tracking devices rescue workers will be able to track their locations and establish safe routes. The data can also be used for mapping and structure characterisation.

Security and privacy issues are important in medical and disaster response scenarios, since medical records should remain private. However in a large-scale emergency, particularly one involving emergency services personnel from many organisations, these concerns are likely to decrease in importance compared to those associated with dealing with data from large numbers of people, how to react when personnel or casualties show a significant change in status, and transferring historical telemetry data from patients to hospital or other remote healthcare providers who need it.

6.4.3 Control aspects

The disaster aspect of the scenario does have possible applications for control technologies. For high-risk environments it may be more appropriate to consider using mechanical devices to identify areas where temperatures, chemical concentrations or other risks make it inappropriate for human presence, or to plot a route of minimum risk where personnel must be deployed. The devices might have some on-board sensing capability but would also be able to obtain data from local sensors. The communication and middleware infrastructure should be able to assign priorities for the control and data transmission among the nodes of the network.

Regarding vehicle control the low level algorithms such as path planning should be able to run using information derived from the sensor network. This requires robustness and adaptation to varying networking conditions. Another issue is that robots are extremely labour intensive usually taking two or more people to operate one robot; one person to drive and the other to "look". Additionally, the interfaces are not user-friendly. There are several key factors to take into consideration when designing autonomous robots for rescue missions including perception, planning, behaviour skills, navigation and learning and adaptation.

6.5 Summary

The recent case of a road tunnel fire (http://news.yahoo.com/s/ap/20050605/ap_on_re_eu/france_tunnel_fire) in the Frejus tunnel in the Alps on June 4, 2005 highlights the importance of dealing with the emergency scenarios in a project like RUNES. The usual "peaceful" scene can change very dramatically into a highly dynamic and chaotic situation. A basically static monitoring infrastructure containing a large set of sensing devices connected to a central operating room must be able to adapt itself to the emergency situation through providing network access to the infrastructure of the emergency agencies having many mobile and wireless nodes. The scenario may provide use cases for demonstrating how the network and middleware solutions of RUNES realize the technical objectives of the project.