

The WORKPAD P2P Service-Oriented Infrastructure for Emergency Management

Tiziana Catarci, Massimiliano de Leoni, Fabio De Rosa, Massimo Mecella, Antonella Poggi

SAPIENZA – Università di Roma

Dipartimento di Informatica e Sistemistica “Antonio Ruberti”

Via Ariosto 25, 00185 Roma, Italy

{catarci, deleoni, derosa, mecella, poggi}@dis.uniroma1.it

Schahram Dustdar, Lukasz Juszczuk, Hong Linh Truong

Vienna University of Technology

Distributed System Group (DSG), Information Systems Institute

Argentinerstrasse 8/184-1, A-1040 Wien, Austria

{dustdar, Lukasz.Juszczuk, truong}@infosys.tuwien.ac.at

Guido Vetere

IBM Italia

via Sciangai 53, 00144 Rome, Italy

gvetere@it.ibm.com

Abstract

In this paper, we present the architecture of the research project WORKPAD¹. This project aims at designing and developing an innovative service-oriented software infrastructure for supporting collaborative work of human operators in emergency/disaster scenarios. In such scenarios, different teams, belonging to different organizations, need to collaborate with one other; each team member is equipped with handheld devices (PDAs) and communication technologies, and should carry on specific tasks. In such a case we can consider the whole team as carrying on a process, and the different teams collaborate through the “interleaving” of all the different processes. Each team is supported by some back-end centres, which need to cooperate at an inter-organizational level.

1 Introduction

The widespread availability of network-enabled handheld devices (e.g., PDAs with WiFi/UMTS capabilities, TETRA smart terminals, etc.) has made the development of pervasive computing environments an emerging reality,

very suitable for managing emergency/disaster situations.

Disaster management activities can be grouped into five phases, as suggested by [5, 8, 10]. The preventive measures are divided into *planning*, *mitigation* and *preparedness* activities. When a disaster or emergency takes place, the *response* activities are designed to provide emergency assistance for victims. They also aim at stabilizing the situation and reducing the probability of secondary damage and speed recovery actions. The *vive recovery* activities include two sets of activities: short-term recovery activities, that return vital life-support systems to a minimum operating standard, and long-term recovery activities that may continue for a number of years after a disaster. This phase represents also the first step to a new planning phase, because this is the point when the analysis of the cause of the disaster or emergency takes place.

Currently, so-called Public Safety Systems (PSSs), typically arranged in specialized organizational centres, are used in three phases within the disaster management: preparation, response and recovery. An analysis of the typical architectures of PSSs can be found in [7], in which it is also discussed how current systems are mainly centralized. The comparison shows that they can be used for emergency cases, e.g., traffic accidents or small fires, where the emergency teams need GIS functions to query addresses, to find the way and to map the emergency location. The disadvan-

¹<http://www.workpad-project.eu>

tage of the compared PSS systems is that they do not use real-time data. They are also only partially usable for large emergencies or disasters because the covering activities of these unusual events need systems which can be used directly at the emergency locations. For these requirements it is necessary to develop mobile applications.

The WORKPAD project aims at providing a service-oriented software and communication infrastructure supporting operators on response and short-term recovery phases. It advocates the use of a 2-level infrastructure based on the following considerations. In complex emergency scenarios, different teams, belonging to different organizations, need to collaborate. Each team member is equipped with handheld devices (PDAs) and communication technologies, and should carry on specific tasks. In such a way, we can see the whole team as carrying on a process, and the different teams (of the different organizations) collaborate through the "interleaving" of all the different processes. In turn, each team is supported by some back-end service providers, and, in order to coordinate between teams, the different service providers need to cooperate at an inter-organizational level. Thus, WORKPAD distinguishes between back-end and front-end communities, which both are organized as peer-to-peer (P2P) networks, where each member can both request and provide advanced services to other members within the community.

The purpose of this paper is to outline the WORKPAD approach based on a P2P and service-oriented paradigm. The rest of this paper is organized as it follows. First, in Section 2, we present an emergency scenario illustrating the WORKPAD 2-level approach. Afterwards, in Section 3, we introduce the WORKPAD architecture and infrastructure, by highlighting some innovative research aspects. Finally, in Section 4 and Section 5, we compare the relevant research work and we conclude the paper.

2 Illustrating Emergency Scenarios

Consider the following scenario. After an earthquake (or a hurricane), a team (e.g., belonging to the Homeland Security Department), equipped with mobile devices (laptops and PDAs), is sent to the disaster area to evaluate the state of specific sites. Their goal is to document the damage directly on a situation map, and to schedule following activities (e.g., reconstruction jobs). Before this process starts, the team leader has stored information needed during the mission, such as the site map, a list of the most important objects at the site, and some previous reports and materials. All such details have been provided by the back-end service provider of the Homeland Security Department, which has constructed them by integrating and exchanging information & knowledge & content stored by many other peer organizations (e.g., the Ministry of Internal Affairs, some

basic spatial data provided by different public and private organizations, etc.). It is worth observing that the integration and involvement of a particular peer has been dynamically and adaptively decided on the basis of the specific process, which in turn depends on the given emergency situation. Therefore, no pre-existing, defined integration infrastructure exists among the back-end peers, but it is rather dynamically built on-demand.

On the front-end, the team constitutes a MANET (Mobile Ad hoc NETWORK), in which the team leader's device coordinates the other team members' devices by providing appropriate information (for example, maps, important objects, and so on) and assigning activities. MANET's are networks of mobile devices that communicate with one another via wireless links without relying on an existing underlying infrastructure, and they are a sound and costly alternative to infrastructure-based networks whenever an infrastructure is no longer available, or cannot be used, as in emergency scenarios [4].

The team members' devices enable them to execute some operations. Such operations, possibly supported by particular hardware (e.g., digital cameras, different communication connections², computational power for image processing, main storage, etc.), are offered as software services to be orchestrated. Such a coordination is carried out by a specific coordination layer hosted by the team leader device. The coordination layer should be able to enact adaptive changes of the process (see [4] and [3]), by having a global knowledge about the status of all the devices and taking into account idle devices, operations that can be safely delayed, and so on.

3 The WORKPAD Infrastructure and Issues

Figure 1 presents the high-level view of WORKPAD framework which includes three main parts: WORKPAD front-end (FE), WORKPAD back-end (BE) and the link between the front-end to the back-end (FE-BE link).

At the WORKPAD front-end, there are several different teams; each team establishes a P2P (peer-to-peer) community of *SmartPeers*. A *SmartPeer* is running on a mobile device (e.g., PDA or laptop) which is controlled by a human (a team member). The underlying network of connected devices is based on a MANET and the WORKPAD front-end software running on devices supports P2P interactions. Using *SmartPeers*, members of the same team collaboratively work on specific tasks for that specific disaster/emergency scenario. One member of the team acts as the team leader who coordinates other team peers.

The WORKPAD back-end system mainly includes a set of *KnowledgePeers* that integrate and provide various types

²Depending on the specific situation, communication can be GPRS/UMTS-based, satellite-based, TETRA-based, etc.

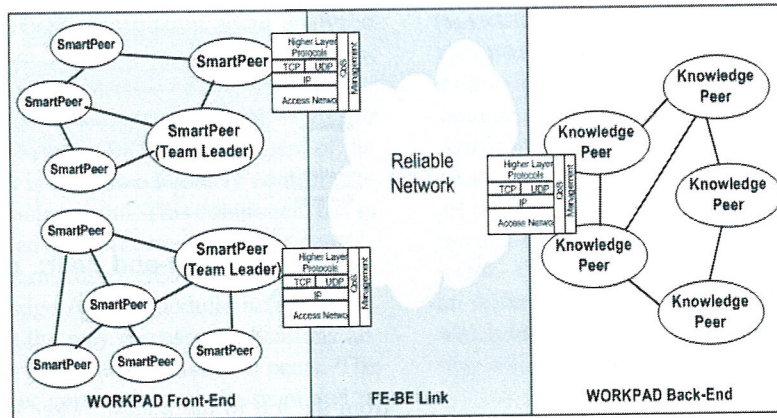


Figure 1. WORKPAD Architecture.

of data, content, and knowledge gathered and used for disaster management. KnowledgePeers form a P2P community that provides advanced services requiring high (performance) computational power.

The WORKPAD FE-BE link is a reliable communication infrastructure that connects services and applications of the two sides. The front-end and back-end applications and services will interface to this infrastructure using the TCP/IP protocol. However, within the infrastructure, various communication technologies, e.g., TETRA, WiMAX, UMTS, etc., can be utilized.

In the following we discuss all the three parts, i.e., WORKPAD back-end, front-end, and FE-BE link, in detail.

Back-end. By definition, emergencies are largely unpredictable in their occurrence, evolution, and impact. Usually, they involve communities – whose infrastructures, organizations, and suppliers are not designed to work together – in new overwhelming circumstances, when cooperation processes and information exchanges must be quickly put in place.

An infrastructure able to support such requirements should be simple, robust, rapid, flexible, and dependable. WORKPAD's back-end is focused on a service-oriented P2P architecture, in which each system can act both as service provider, consumer, and integrator. Furthermore, it will allow front-end operators to access a Grid of systems where they can get or set information which is relevant to the situation they are facing or the action they are taking. Such information is not necessarily contained in single systems, but is potentially spread over the network, and is delivered, collected, reconciled and supplied throughout a distributed, "on demand" integration.

WORKPAD's peers manage internal databases and procedures, export their data semantics, supply query and update methods, establish and evaluate acquaintances with other

systems (mappings), and support publish/subscribe mechanisms to allow real-time notification of information updates. The integration logics will be distributed: no specific integration nodes will be required "a priori", and, consequently, no global ontology is required to provide the back-end community with a global semantics.

In this setting, the WORKPAD's back-end may be classified as a decentralized-unmodelled service-oriented architecture [9]. Specifically, at a conceptual level, each back-end peer will consist of several components (see Figure 2):

- Knowledge core: ontology that models data at the back-end;
- Modelling interface: allows to visualize and manipulate the back-end ontology;
- Data access interface: handles query/update/synchronization requests from both the front-end and other back-end peers;
- Local data sources: store real data locally;
- Source/P2P mappings: establish the relationships among the ontology and both local and other back-end peers.

Intuitively, each back-end peer offers access to data that can be either locally or remotely stored. Doing that, a peer maintains a set of schema level assertions relating its schema to schemas of other peers in the network (P2P mappings).

Concretely, the information integration infrastructure will be based on the evolution of an existing P2P data integration system called Hyper (How Your Peers Exchange Resources [1]), and will be developed on top of existing Web service standards (WSDL, OWL, etc.). The features of the WORKPAD P2P back-end represent an enhancement

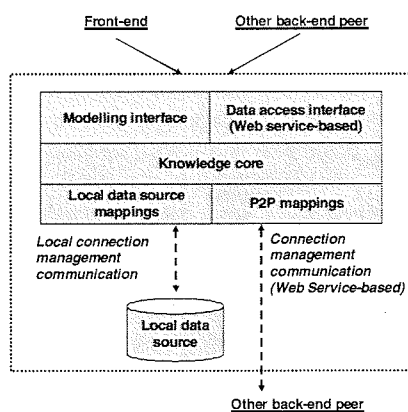


Figure 2. Back-end peer: conceptual architecture.

of current centralized federative approaches like those implemented in USA by "The Capital Wireless Integrated Network" (CapWIN, <http://www.capwin.org/>). In the scenario of collaborative work of human operators in emergencies, in fact, a P2P service-oriented architecture offers a greater autonomy to organizations backing the field teams:

- Distribution of semantic mappings will ensure a rapid and incremental deployment of the integration logics;
- Modularity will ensure that front-end networks will not be bloated with unnecessary data;
- Endpoint-based communication will limit impacts of network dynamics (nodes entering or leaving, outages, etc.), and will ensure scalability, flexibility, and performances.

Front-end. During the activities in the field of an emergency, each team forms a closed front-end community. These communities are constituted by operators, which are equipped with mobile devices, connected in an ad-hoc and peer-to-peer fashion, and carry out processes in which the adaptation to anomalies of connections and tasks is a fundamental requirement.

As the processes in the front-end of WORKPAD must be continuously adapted (i) to the changes in the mobile ad-hoc networks, (ii) to the capabilities of human workers and their devices, and as well (iii) to the dynamics of disaster scenarios, the adaptation technique must be aware of the current situation, which is referred to as context. This context-awareness implies that a system is able to capture relevant data, to store and represent it in a proper way, and to adapt its own behaviour based on captured data.

Figure 3 shows the conceptual architecture for the WORKPAD front-end. It consists of three layers named,




Figure 3. Front-end peer: conceptual architecture.

from the top to the bottom, *User Application*, *WORKPAD Front-end Middleware*, and *Data Storage and Communication Management*. Each layer consists of several components, and not all of them will be deployed in every front-end device.

In the user application layer, there are various applications (services) that will be used by team members in order to conduct the collaborative work in disaster scenarios. This layer includes both applications developed by WORKPAD and those available. However, all user applications will be used for achieving tasks planned for disaster managements. For instances, a GIS service should be used to display geo-information at the disaster place, by supporting the user during the collaborative work, whereas a multimedia editor/viewer service could support team members to capture and edit/annotate photos. Services are handled by the *Worklist Handler* module which supports team members to receive tasks given by the team leader and to invoke assigned tasks. This component will interact with the *Adaptive Process Management* component and is able to invoke external user applications required by specific tasks. Finally, the *Context Editor* module will allow the team members to enter additional contextual information which could not be captured by the WORKPAD middleware, e.g., the contextual information about the (health) status of team members or the disaster scene.

The WORKPAD front-end middleware includes four main components named *Context Monitoring & Management*, *Process Mining*, *Adaptive Process Management*, and *Knowledge Access*. This layer provides core middleware that support adaptively process execution, providing contextual information and common libraries for accessing data and knowledge at the back-end. Specifically, the Context Monitoring & Management module provides various types of contextual information to other components. Contextual information are associated with devices and networks, human profiles and activities, disaster scenarios, etc. Because this data are utilized by various components, a well-defined representation is used, and a common interface for accessing contextual information is defined, taking into account different client interactions (locally and remotely). The Pro-

cess Mining module provides information about analyzed behaviour of users, services, and processes. This information is used by the Adaptive Process Management in order to optimize the execution of processes. The Adaptive Process Management component is the core element of the front-end middleware. It is used to adaptively control processes during disaster managements. This component has to manage processes in an adaptive manner based on contextual information and patterns discovered from process mining. Finally, the Knowledge Access module includes a set of libraries that simplify the way the user applications access data and knowledge stored at the back-end peers. The back-end peers allow user applications at the front-end to store and query data and knowledge at the back-end.

The data storage and communication management layer is used by the front-end middleware and user applications to store data and communicate with other peers. Specifically, the *Lightweight Storage* component will be used to store various types of data and knowledge. These information is either used for front-end middleware and applications or temporary data that will later be stored in back-end services. For example, adaptive process management uses local lightweight databases for process schemas, or contextual monitoring and management will store contextual information in lightweight storages. However, as mobile devices at the front-end have limited capability, storage systems and mechanism must be carefully selected. On the other hand, the *MANET & FE/BE Communication* module is in charge of managing network connection among front-end (MANET routing) peers and between front-end and back-end peers.

Connecting front-end and back-end. In order to achieve an integrated infrastructure where front-end nodes can have access to a back-end, the MANET should be connected to wired networks through cellular, satellite or TETRA solutions acting as routers. An asymmetrical MANET is a MANET where there are some special nodes (e.g., the coordinator device) representing a gateway to the fixed network. Such nodes should be equipped with at least two network interfaces, one connected to the MANET and the other one to the back-end, and should be able to receive and return traffic from outside hosts and route it to the destination MANET node. Such devices play a critical role and become bottlenecks, therefore the WORKPAD project is researching how to design an effective connection depending on the number of nodes, available bandwidth on the uplink interface, priority levels of information, available networks.

Moreover, some possibilities for the link front-end/back-end will be studied and experimented, in order to choose the most suitable one. In particular, the communication between the front-end and back-end will be based on a UDP&TCP/IP networks and technologies which offer reliable communication channels. Specifically, each team

(MANET network) will be equipped with an end-point device which allows the community to exchange data/information into/from the BE infrastructure: for instance, every team could be equipped with a TETRA switch or a satellite switch, or both. On the other hand, a BE peer will communicate with each other over reliable TCP/IP networks, and use a switch device (TETRA/satellite/WiMAX -based communication or other) to communicate with relative MANET teams. Of course, BE/FE communications might be based on different technologies, depending on disaster situations which take place. The key issue is that this eventual situation will be "transparent" to the users placed in the FE.

4 Related Work

In the last years, many projects and initiatives, especially in USA and in Europe, have investigated many issues related to the ones considered by WORKPAD. In particular, some projects consider mobile devices and their adoption in emergencies (see [2] for a survey). Generally, all the proposed approaches in these projects are focused on scenarios where process management system supporting collaborative work is running on back-end servers. Apart few projects such as POMPEI³, ORCHESTRA⁴, and AKOGRIMO⁵, all proposed mobile software architectures are based on the client/server paradigm, in which front-end applications (GUI components, mobile Web clients, etc.) are conceived to be executed on mobile client devices (PDA's, laptops, smartphones, etc.), while back-end applications (Web services, location-services, data storage, etc.) are targeted to be executed on a fixed server infrastructure (generally, a Grid environment). Conversely, in WORKPAD a fully peer-to-peer paradigm architecture is used in both back-end and front-end infrastructures, where each network peer requires/provides services, from/to other peers, to be orchestrated by specific software. We argue that such a service-oriented peer-to-peer software architecture is suitable for emergency scenarios as it allows to dynamically integrate several back-end systems based on different technologies, and allows front-end teams to dynamically adapt processes during their execution. Furthermore, although numerous emergency management systems have been developed, only a slight number of them uses context-aware techniques. However, we believe that context-awareness is a key requirement for pervasive environments, especially in emergency situations.

A few projects have considered the issue of data integration in peer-to-peer settings. Among them, (i) PIAZZA⁶ aims at providing "semantic mediation" between an envi-

³<http://www.pompei-eu.com/>

⁴<http://www.eu-orchestra.org/>

⁵<http://www.akogrimo.org/>

⁶<http://data.cs.washington.edu/p2p/piazza/>

ronment of thousands of peers, each with its own schemas, rather than requiring the use of a single, uniform, centralized mediated schema to share data between peers, and (ii) TERRAWORLD⁷ is one of the first research initiatives in the field of geo-integration. Moreover, (iii) RAMFLOOD⁸ is an example of a decision-making system using combined, mixed data resources gathered from earth observation, satellite positioning systems, in-situ sensors and geo-referenced information with advanced graphical visualisation methods for gathering knowledge.

5 Conclusion

This paper has introduced WORKPAD, a service-oriented infrastructure which provides a 2-level peer-to-peer workspace suitable in emergency scenarios. In particular, all components of the WORKPAD conceptual architecture both at back-end and front-end have been described as well as the relative issues have been given.

The strategic impact of the WORKPAD project is twofold: on one hand, the application of new technologies will help addressing the new and increasing threats to public security; on the other hand, the development of an innovative solution for emergency operations will improve the state of the art in several IT emerging disciplines enabling the interoperability among different teams belonging to different organisations. For that reason, the benefit we expect from this project consists in both the improvement of Homeland Security capabilities and the progress of IT industry.

The availability of a specific IT and data infrastructure for emergencies can significantly improve the quality of the intervention of first-responders. Moreover, it can effectively assist in coordinating rescue teams and operations in the first crucial days after the emergency/disaster event. Finally, it can help planning and performing recovery activities in later emergency/disaster management phases. In particular, thanks to its innovative service-oriented architecture, the WORKPAD framework will give the ability of integrating mobile wireless networks to a rich back-end information network in a flexible, cost-effective, reliable, and powerful manner.

Acknowledgement. This work has been supported by the European Commission through the FP6-2005-IST-5-034749 project *WORKPAD*.

⁷<http://www.isi.edu/integration/TerraWorld/>

⁸<http://www.cimne.com/ramflood/>

References

- [1] D. Calvanese, G. De Giacomo, M. Lenzerini, R. Rosati, and G. Vetere. Hyper: A Framework for Peer-to-Peer Data Integration on GRIDs. In *Proc. Semantics for GRID Databases, 1st International IFIP Conference on Semantics of a Networked World (ICSNW 2004)*, volume 3226 of *LNCS*, 2004.
- [2] T. Catarci, F. De Rosa, M. de Leoni, M. Mecella, M. Angelaccio, S. Dustdar, B. Gonzalez, and G. Iiritano. WORKPAD: 2-Layered Peer-to-Peer for Emergency Management through Adaptive Processes. In *Proc. 2nd International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2006)*, 2006.
- [3] M. de Leoni, F. De Rosa, and M. Mecella. MOBIDIS: A Pervasive Architecture for Emergency Management. In *Proc. 4th International Workshop on the Distributed and Mobile Collaboration (DMC 2006) @ WETICE 2006*, 2006.
- [4] F. De Rosa, A. Malizia, and M. Mecella. Disconnection Prediction in Mobile Ad-hoc Networks for Supporting Cooperative Work. *IEEE Pervasive Computing*, 3(4):62–70, 2005.
- [5] ESRI White Paper. GIS for Emergency Management. Available on-line (link checked on March 1st, 2007): <http://www.esri.com/library/whitepapers/pdfs/emergmgmt.pdf>, 1999.
- [6] E. Galanti. Metodo Augustus - Guida del Dipartimento della Protezione Civile. Available on-line (link checked on March, 1st 2007): <http://www.casaleinforma.it/pcivile/scarica/04augustus.pdf>, 2004 (in Italian).
- [7] S. H. Leitinger. Comparison of GIS-based Public Safety Systems for Emergency Management. In *Proc. 24th Urban Data Management Symposium*, 2004.
- [8] E. Plate and B. Merz, editors. *Naturkatastrophen: Ursachen, Wirkung, Vorsorge*. Schweizerbart'sche Verlagsbuchhandlung, Stuttgart, 2001 (in German).
- [9] G. Vetere and M. Lenzerini. Models for Semantic Interoperability in Service-Oriented Architectures. *IBM System Journal*, 4(44):887–904, 2005.
- [10] A. Zenger and D. I. Smith. Impediments to using GIS for Real-time Disaster Decision Support. *Computers, Environment and Urban Systems*, 2(27):123–141, 2003.