

# MOBILE AND DISTRIBUTED COLLABORATION IN VIRTUAL COMMUNITIES

---

Schahram Dustdar and Harald Gall

*Distributed Systems Group, Vienna University of Technology  
Argentinierstrasse 8/184-1, 1040 Wien, AUSTRIA  
{dustdar, gall}@infosys.tuwien.ac.at*

*Organizations are increasingly forced to manage and coordinate their product and service development processes, to make their products and services available as quickly as possible, and to involve employees, customers, suppliers, and partners in different stages of their businesses. These processes have to consider that their participants are more and more on the move while they are working and that their expertise can be shared across locations and different mobile devices. This paper defines a framework for mobile distributed collaboration, defines a set of requirements for virtual communities, and discusses a mobile teamwork support software architecture that has been developed in an ongoing EU-project. The framework together with the architecture should enable to enhance current collaboration approaches to include the dimension of mobile participants and virtual communities.*

## 1. INTRODUCTION

Information Technology in general and software systems (such as Workflow Management Systems (WfMS) and Groupware systems) in particular have been used to automate or to augment business processes in organizations (Bussler, 1999; Schael, 1996). Groupware typically does not contain any knowledge or representation of the goals or processes of the group. On the other hand, WfMS are typically organizationally aware because they contain an explicit representation of business processes. WfMS have been defined as "technology based systems that define, manage, and execute workflow processes through the execution of software whose order of execution is driven by a computer representation of the workflow process logic" (WfMC, 1995). Future WfMS will soon cover inter-organizational activities and processes including product value-chains on the Internet (Bussler, 2001; Casati et al., 2001; Chen et al., 2001; Christophides et al., 2001; Kafeza et al., 2001; Krithivisan et al., 2001; Papazoglou and Jeusfeld, 1998; Puusjärvi and Laine, 2001). There has been considerable attempts to merge workflow and groupware technologies. Industrial research labs and product teams have made significant steps forward (Chen et al., 2001, Hausleitner and Dustdar, 1999). A WfMS can impose a rigid work environment on users, which often has immediate consequences: One example is among users who perform time-consuming manual "workarounds"; the consequence is lower efficiency and dissatisfaction with the system. As a result

WfMS enable continuous loops of sub-processes such as goal setting, working, monitoring the work, measuring performance, recording and analyzing the outputs and evaluating the "productivity" of personnel. Users of WfMS often consider the controlling and monitoring possibilities as a "dark side" of these systems, which results in demotivating employees and a lack of support for mobile distributed work. In today's business environment business process participants demand process awareness to a relatively high degree of the software they use for collaborative work. In addition organizational awareness (e.g. roles) and mobility aspects become increasingly relevant. Current WfMS and Groupware systems do not combine those relevant features. The contribution of this paper is threefold: Firstly, it presents an overview of collaborative technologies used in today's organizations. Secondly, it discusses the issues and requirements for virtual communities in the context of corporate collaboration. Thirdly it presents a software architecture including a prototype implementation for mobile virtual communities.

## 2. TECHNOLOGIES FOR VIRTUAL ENTERPRISES

### 2.1 Collaborative technologies

Recent advances in the area of Internet Computing and collaborative WfMS are often seen as key technologies for supporting mobile and distributed teams. Cooperative tasks in teams are increasing and as a consequence the use of collaborative systems is becoming more pervasive. In this paper our model for collaborative technologies distinguishes two dimensions: "*Business process spectrum*" and "*Knowledge Management approach*" as shown in Figure 1. A modeled business process such as "customer order entry" can be modeled using a traditional WfMS. However a modeled process can only be enacted (instantiated) as it was designed. If an exception occurs, a workflow administrator needs to re-model the process before the execution can continue. This limits the usability of workflow systems in a world where constant adaptation to new situations is necessary and where teams are increasingly mobile and distributed. An example of an ad-hoc process is the discussion of a project's design review using e-mail (Groupware).

The second dimension presented is the "*Knowledge Management approach*". There are two distinct approaches in this domain: retrieval or a collaborative approach. The *retrieval* approach assumes that members of the organization contribute to a "group memory" by submitting/sending documents into a central repository. This repository can be searched (retrieved) and as a consequence be presented as an "Enterprise Portal" accessible with a web browser. The *collaborative* approach uses the e-mail metaphor and integrates the option to "mail-in" documents into the group memory. Generally one can see the trend of integrating customers into core business processes whenever it is possible (e.g. Customer Relationship Management (CRM) or Partner Relationship Management (PRM)).




Figure 1 - Collaborative technologies

In this paper we focus on ad-hoc business processes and on mobile and distributed collaboration as depicted on the left quadrants of Figure 1.

2.2 A framework for mobile and distributed collaboration

The mobility of participants also offers new ways of distributed collaboration: processes are no longer bound to locations of resources (such as participants or artifacts) but can consider several availability modes. A framework for such mobile distributed collaboration should enable participants to join so-called virtual communities, to share process artifacts in that community, and to use different means of communication to collaborate on a particular process. A business process, which is highly relevant for mobile distributed collaborative work in most organizations, is "New Product Development" (NPD). An NPD process contains all innovative tasks for designing, implementing and distributing products or services. Figure 2 illustrates a high-level map for our definition of virtual communities. Our framework consists of three orthogonal dimensions: Organizational Unit, Location, and Process. The dimension "Organizational Unit" (e.g. Architecture, Design, Implementation) consists of functional departments or skill sets that have to be applied to various business process tasks. The dimension "Location" denotes the geographical location of the virtual community team members. The dimension "Process" describes possible process stages of any given business process (e.g. Testing, Engineering, Manufacturing). A virtual community (VC) can be defined as a 3-tuple consisting of information on Organizational Units (O), Location (L), and Process (P) or  $VC := \langle O, L, P \rangle$  with the following definitions: Organizational Units  $O = \{o_1, o_2, \dots, o_n\}$ ; Locations  $L = \{l_1, l_2, \dots, l_m\}$ ; Processes  $P = \{p_1, p_2, \dots, p_k\}$




Figure 2 – A framework for mobile distributed collaboration in virtual communities

Example *use cases* of the MDC framework are: information updating and notification of availability (of resources); searching for resources (artifacts, experts, etc.); synchronous and asynchronous communication in a community; collaboration on artifacts; community establishment and updating. The above example use cases can be instrumented in many different ways considering the requirements of the actual organizational unit, the process and the location: some instrumentations consider the *location-aware* dimension, i.e. it is of particular interest where the resource actually is residing; others use the framework in a *location-transparent* way in which it is important that some task is carried out but independently of where the actual resources are. The same principle applies to the other two dimensions "organizational unit" and "process." In a particular context the use case "expert search," for example, can be considered location-transparent ( $o_n, p_k, *$ ) or location-aware ( $o_n, p_k, \text{"Vienna"}$ ). For another context it could also be useful to consider variations of it, for example, ("Implementation",  $p_k, \text{"Vienna"}$ ) or combinations of it. Depending on the requirements for the mobile distributed collaboration scenario, the instantiator of a process would define the corresponding virtual community along the O, L, P dimensions. The remainder of this paper discusses the design goals of such a mobile distributed collaboration framework, provides a technology mapping from the requirements to a software environment and outlines a particular Web-based peer-to-peer architecture that has been developed in the EU-project MOTION that especially considers mobility of project participants (Kirda et al., 2001; Reif et al., 2001).

### 3. ARCHITECTURE FOR AN MDC SYSTEM

In contrast to traditional software architectures, architectures that support mobility are faced with several additional difficulties: because of bandwidth restrictions, unreliable connections and disconnected operations, mechanisms and components are necessary to locate participants, synchronize data and query available resources. Depending on the location of the participant, the number of offered services may vary in quantity and quality. Sophisticated subscription mechanisms and notification services are necessary to disseminate information to the mobile participant instead of forcing the participant to find it. Furthermore, the information in mobile environments needs to conform to different standards such as WML and WAP for data representation. The varying display sizes limit the amount of information that can be displayed on a small mobile device. Thus, components are required that render the information according to the display capabilities of a certain device (Buyukkokten et al., 2001; Kaasinen et al., 2000). High-quality wireless multimedia communications such as UMTS will further improve the quality and quantity of services on the participant's mobile device and have to be considered for an MDC architecture as well. Much of the required functionality already exists as Web applications, so the key is to efficiently migrate the features to mobile applications. The most common approach is called Web-to-wireless: this relies on existing Web services, but adds a wireless channel to them, providing stakeholders such as customers, employees, and business partners with an access point. Since in the MDC framework we especially address intra-company mobile collaboration, we have to

foresee multiple access points depending on the NPD and the "instance" of the three dimensions organizational unit, process and location.

### 3.1 MDC architecture design goals

Mobile architectures usually integrate both fixed and mobile components. Hence, we designed the MDC architecture with the following design goals in mind. The architecture has to be: *open* with respect to integration of existing technologies and tools; *generic* to be deployed in organizations with varying internal organizational structures, business processes and IT infrastructures; *scalable* for different number of participants, future extensions and new requirements; *adaptable* to restrictions imposed by mobility both of the mobile participants and the mobile devices.

### 3.2 Participant mobility

The participant is not confined to one location, but may be on the move while working. There are three connectivity modes considered in the MDC infrastructure: *connected mode*, *disconnected mode*, and *ad-hoc mode*. Connected mode is used whenever fixed network connectivity is available. There is a single global context determined by all hosts on the network. Information can be accessed and shared from any point, at any time. Disconnected mode is an operational mode, which is required for mobile working. The mobile participant does not always have network connectivity, and this gives a special importance to disconnected working capabilities. In our architecture, the disconnected mode is a special case of the connected mode. The participant is able to continue working even if he is disconnected. Most of the information available in the system, though, is not available to the participant. Once the participant is connected again, the changes he has performed locally are synchronized with the rest of the system. In the *ad-hoc* mode, participants, for example, run a meeting and do not have any fixed network connectivity, so they create a temporary network. The context for information sharing and cooperation is limited to the one provided by the parties in communication and to the resources they bring along. Hence, participants can build *ad-hoc* networks for collaboration and information exchange.

### 3.3 Publish/subscribe and information push

Publish/Subscribe mechanisms allow to distribute information to where it is wanted. Unlike normal point-to-point messaging, providers and consumers of information do not need to know anything about each other. Push and event-based systems are closely related. The purpose of push systems is a timely distribution of data and information to consumers whereas event-based systems focus on notification of events. Both push and event-based systems play an important role in the MDC architecture. We use a push system to notify participants based on a *profile* (Hauswirth and Jazayari 1999). The participants define profiles for the kind of information they are interested in, and the push system delivers the information whenever it is available. Participants may subscribe to (and unsubscribe from) specific topics. Furthermore, participants can also subscribe to other participants and be notified whenever these participants are online (or available etc.) which is particularly interesting in the mobile collaboration scenario.

### 3.4 Peer-to-Peer Middleware

An event-based middleware is used in MOTION as a *scalable* and *flexible* infrastructure to transfer messages to the participants and components in the connected, disconnected and ad-hoc modes of operation. This middleware has a Peer-to-Peer architecture (P2P) following the fully-distributed information sharing principles of Gnutella. The peer-to-peer communication infrastructure is especially advantageous to support different modes of connectivity. Because every computer (= peer) can work as a server as well as a client, it is possible to build ad-hoc networks rather easily. The P2P middleware (Picco et al., 1999) is the central underlying component of the system. It manages the subscriptions and provides an event-based system to notify participants on the subscribed events. It is also responsible for managing the virtual communities and for locating the actual physical location of a requested resource. The middleware is also responsible for queuing participant actions and events that cannot be processed when working in the disconnected or ad-hoc networking mode. Notification messages that cannot be delivered because of an unreachable host are queued as well. All these queued actions and events are processed as soon as the host connects again to the MDC platform.

### 3.5 The Web as MDC service platform

The scalability and the distributed nature of the Web has made it a popular platform for building collaborative tools. Thus, many Web-based tools have been introduced and there are countless Web applications for improving communication, information exchange and process management. Boeing, for example, has been successfully using Web-based collaborative applications in the construction of its airplanes (Fowler et al., 2000). To meet the requirements and to cover the scenarios mentioned earlier in this paper, the MDC architecture utilizes the existing Web infrastructure and exploits many of its advantages: Web access is widely available (e.g. on airports), Universal Resource Locators (URLs) are a simple and unique way of identifying resources on the Internet, and Web technologies also offer security against sniffing attacks by using secure HTTP connections through SSL.

Figure 3 provides a sketch of the Web-based MDC architecture. It illustrates typical MDC peers and their components and shows the access to the system from devices not running a Web-server and the MOTION middleware (Web-terminal, WAP phone, Personal Digital Assistant (PDA)).




Figure 3 – Peer-to-Peer Middleware (Reif et al., 2001)

Not every peer needs to include all components depicted in Figure 3. Depending on the computational power and the memory capacity, some components might offer

reduced functionality. The MDC architecture also supports access from devices not running a web server and the P2P middleware. These devices access the platform via the Web server of an MDC host. The only requirement for these devices is that they run a Web- or WAP-client. This enables the participant to access resources from any computer running a Web browser (e.g. in an Internet cafe), from a PDA or a cellular phone.

#### **4. CASE STUDY: MOBILE PHONE DESIGN**

One of the case study providers is a multi-national company in the market of global telecommunication systems and equipment. This company wants to facilitate the ways in which geographically distributed development units of the company divide their work, communicate and collaborate. The case study especially focuses on the mobile and distributed collaboration support for the process of designing mobile phone software. The MDC architecture is expected to support new and more efficient ways of working. An important requirement, for example, is to enable employees to find *experts* and their availability in the company. Due to the size of the company, it has many employees that have expert knowledge in a specific domain and this expertise should be exploited more efficiently especially when these people are on the move (or in some other branch of the company). Employees should be able to search and contact experts to solve their problems across many organizational units. The main MDC architecture evaluation criteria in this case study are: Improvement of production processes by defining and improving the distributed design, development and evaluation processes in mobile phones production; product improvement by effectively supporting the enhancement of PDA products. The MDC prototype, that has been developed, is currently in its user acceptance evaluation phase and detailed results can be reported soon.

#### **5. CONCLUSIONS AND FUTURE WORK**

Solutions for collaborative systems in current working scenarios of large and multi-site enterprises have to consider mobility of users and their devices. The technological advances in mobile computing define new requirements for collaborative systems. In this paper we discussed the issue of mobile and distributed collaboration scenarios in which users have to work together while they are on the move. Finding experts, sharing expert knowledge and corresponding resources led us to the definition of virtual communities that – despite the actual location of a user – foster collaboration across organizational units and processes. We defined a mobile and distributed collaborative system (MDC) architecture that bases upon a Peer-to-Peer (P2P) middleware, uses publish/subscribe mechanisms for notification and definition of interest of participants, and exploits information push technology to inform about (newly) available resources. Virtual communities function as a basic forum for exchanging artifacts on a P2P-basis, searching for experts in the company (based on people's profile), or notify members of a particular community of tasks, results or availability of resources. Synchronous collaboration on artifacts or synchronous communication can be established in such a virtual community, but is not inherent in the architecture. The MDC prototype is currently under end-user evaluation in a large, multi-site enterprise that also helped in pinning down the MDC requirements in a business process analysis phase.

### Acknowledgements

The MOTION project is supported by the European Commission in the IST Program (IST-1999-11400). We thank the MOTION team at the Distributed Systems Group: Engin Kirda, Pascal Fenkam, Gerald Reif, and Mehdi Jazayeri.

### 6. REFERENCES

1. Bussler C. The role of B2B protocols in inter-enterprise process execution. Proceedings of TES 2001, Computer Science Lecture Notes, Springer Verlag 2001; 16-29.
2. Bussler C. Enterprise-wide Workflow Management. *IEEE Concurrency* 1999; 7, 3: 32-43.
3. Buyukkokten O., Garcia-Molina H., Paepcke A., Focused Web searching with PDAs, 9th International World Wide Web Conference, Amsterdam, Netherlands, May 2000.
4. Casati F, Sayal M, Shan MC. Developing e-Services for composing e-services in Proceedings CaiSE 2001, Computer Science Lecture Notes, Springer Verlag 2001; 171-186.
5. Chen Q, Hsu M, Dayal U. Peer-to-Peer Collaborative Internet Business Servers HP-Labs Technical Working Paper HPL-2001-14, 2001.
6. Christophides V, Hull R, Kumar A. Querying and Splicing of XML Workflows. in Proceedings CoopIS 2001, Computer Science Lecture Notes, Springer Verlag 2001; 386-403.
7. Fowler S.L., Novack A.J., Stillings M.J., The Evolution of a Manufacturing Web Site, Proceedings of the 9<sup>th</sup> International World Wide Web Conference, Amsterdam, Netherlands, May 2000.
8. Hausleitner A, Dustdar S. Caramba - Ein Java basiertes Multimedia Koordinationssystem. in: *Erfahrungen mit Java. Projekte aus Industrie und Hochschule*. Silvano Maffei, et al. (Eds.), dPunkt-Verlag, Heidelberg 1999.
9. Hauswirth M., Jazayeri M., A Component and Communication Model for Push Systems, Proceedings of the ESEC/FSE 99 - Joint 7th European Software Engineering Conference (ESEC) and 7th ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE-7), Toulouse, France, pp. 20-38, Sept. 1999.
10. Kafeza E, Chiu DKW, Kafeza L. View-based contracts in an e-service cross-organizational workflow environment. in Proceedings TES 2001, Computer Science Lecture Notes, Springer Verlag 2001: 74-88.
11. Kirda E., Gall H., Reif G., Fenkam P., Kerer C., Supporting Mobile Users and Distributed Teamwork, Proceedings of ConTEL 2001 - 6<sup>th</sup> International Conference on Telecommunications, Zagreb, Croatia, June 2001.
12. Kaasinen E., Aaltonen M., Kolari J., Melakoski S., Laakko T., Two approaches to bringing Internet services to WAP devices, 9th International World Wide Web Conference, Amsterdam, Netherlands, May 2000.
13. Krithivisan R., Helal AS. BizBuilder - An e-Services Framework targeted for Internet Workflow. in Proceedings TES 2001, Computer Science Lecture Notes, Springer Verlag 2001: 89-102.
14. Papazoglou MP, Jeusfeld MA. Distributed, Interoperable Workflow support for electronic commerce in Proceedings TREC 1998, Computer Science Lecture Notes, Springer Verlag 1998: 192-204.
15. Picco G.P., Murphy A.L., Gruia-Catalin R., Linda Meets Mobility, Proceedings of the 21st International Conference on Software Engineering (ICSE'99), ACM Press, pp. 368-377, May 1999.
16. Puustjärvi J, Laine H. Supporting cooperative inter-organizational business transactions. in Proceedings DEXA 2001, Computer Science Lecture Notes, Springer Verlag 2001: 836-845.
17. Schael Thomas. *Workflow Management Systems for Process Organizations*. New York: Springer, 1996.
18. Reif G., Kirda E., Gall H., Picco G.P., Cugola G., A Web-based peer-to-peer architecture for collaborative nomadic working, 10<sup>th</sup> IEEE Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE), Boston, MA, USA, June 2001.
19. Workflow Management Coalition (WfMC) *Workflow Management Specification Glossary*. 1995