AUGMENTING WEB-BASED COLLABORATION WITH ADAPTIVE REPLICATION AND MOBILITY
(*EXTENDED ABSTRACT*)

HENDRIK DECKER,
LUIS IRÚN, JORDI BATALLER, PABLO GALDÁMEZ, RAMÓN GARCÍA
Instituto Tecnológico de Informática, UPV, Valencia, Spain
{*hendrik, lirun, bataller, pgaldam, rgarcia*}*@iti.es*

We survey two systems. First, the middleware DIRECS, for supporting the adaptive wide-area replication of collaboration data and thereby the availability, responsiveness and failure resilience of web-based collaborative systems. Availability is adaptive with regard to replication consistency requirements, and vice-versa. Second, the middleware MOVE, for supporting adaptive voice/data-integrated mobile web-based collaboration services. MOVE adapts to networks, locations, devices and user profiles. We then conclude with an outlook of ongoing and envisaged work which aims at augmenting the DIRECS functionalities with MOVE features.

1 DIRECS

Most web-based collaboration systems (abbr. WBCSs) use one or several underlying databases. Availability, responsiveness and failure resilience of those databases are key factors for the performance and user acceptance of a WBCS. In general, and particularly for databases and information systems, a tried and tested principle for supporting such features is *distributed replication.*

For WBCSs, this principle has been adopted by DIRECS [1]. Database(s) underlying a WBCS are replicated over the sites of all session participants, so that the data and system services for web-based collaboration (abbr. WBC) become more available, responsive and failure resilient. DIRECS is a specialization of the more generic MADIS architecture [2], which in turn is a more efficient, JDBC-based follow-up of the object-oriented middleware COPLA [3], as developed in the GlobData project [4].

To WBCS developers and end users, DIRECS provides a transparent object-oriented view of a singleton WBCS data storage, while the data sources are actually replicated over a set of networked nodes. The user layer supports the *development* of WBCSs, and the core management layer supports their *operation.* These layers are two of the three main layers of the DIRECS architecture, as shown in fig. 1.

The core management layer consists of a fixed and an exchangeable part. The fixed part serves transactions and delegates tasks to the storage layer (bottom) or to the protocols module, which constitutes the exchangeable part. It is exchangeable in the sense that protocols for message passing, replication and failure recovery can be plugged in and out, so as to adapt the system's behaviour to the needs of given applications, network loads and user requirements. In other words, the WBCS administrator can always choose suitable

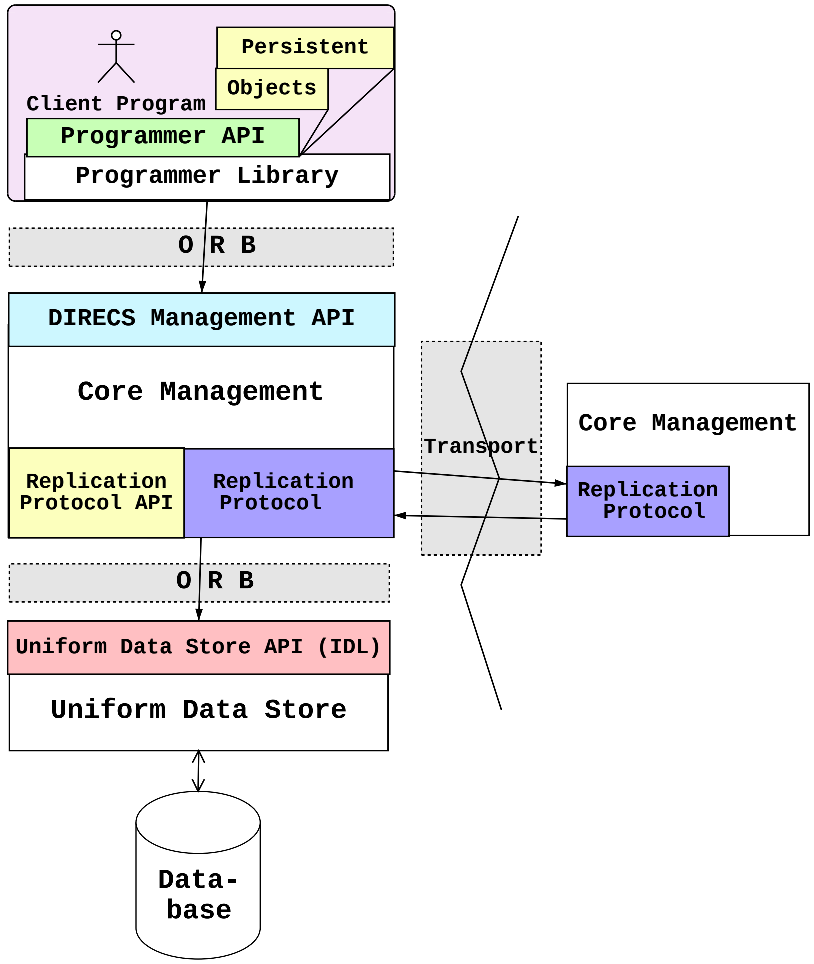

Figure 1. DIRECS Layered Architecture

protocols, in order to adapt the system to an optimal tradeoff between requirements of consistency of replicated information and its availability.

DIRECS can be configured either in a *peer-to-peer* or a *client/server* style of architecture. “Peer-to-peer style” means that collaboration data are replicated transparently at the site of each member node of a collaborative session. The architectural symmetry between collaborative session members entails that each member node can be either client or server to each other node, in dynamically interchangeable roles. This P2P structure is sketched in fig. 2. The large cylinder suggests that the distributed database is perceived as a single system. Multiple users may be hosted at each site.

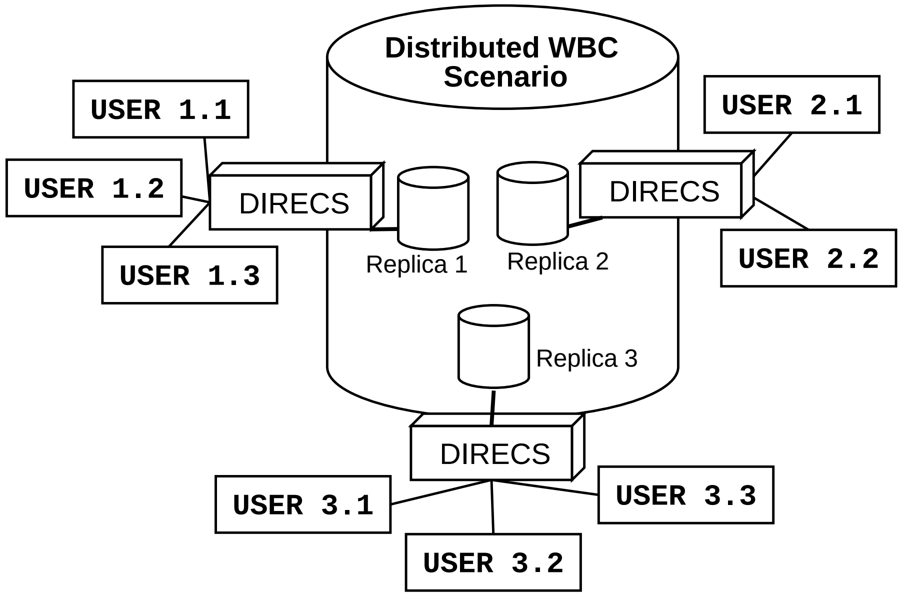

Figure 2. DIRECS in Peer-to-Peer mode

The client/server variant of DIRECS would have each member of a collaborative session access a virtual database server which does not belong to the site of any of the session's participants; rather, it would be a (public or commercial) third party's repository to which all collaborative session participants only have read access. Thus, participant nodes behave as clients of a virtual repository server which actually is a network of replicas. The client/server and P2P installations can be combined orthogonally. For instance, collaborative users of an external distributed database may internally use other databases containing, e.g., time tables, project data etc.

2 MOVE

The main objective of the EU-promoted project MOVE [5] (Mobile Middleware with Voice-Enabled Services) [6] was the integration of voice (real-time data) and multimedia services (real-time and non-real-time data) for mobile GSM-, UMTS- and Web-based collaboration [7].

Project partners from Siemens (Germany), Tecsi (France), Orange (UK), RWTH Aachen (Germany) and Kent Ridge Digital Labs (Singapore) have developed the MOVE middleware architecture, including an API for service providers. It supports the initiation, maintenance and termination of mobile voice-enabled collaborative multimedia sessions, connecting each internet site

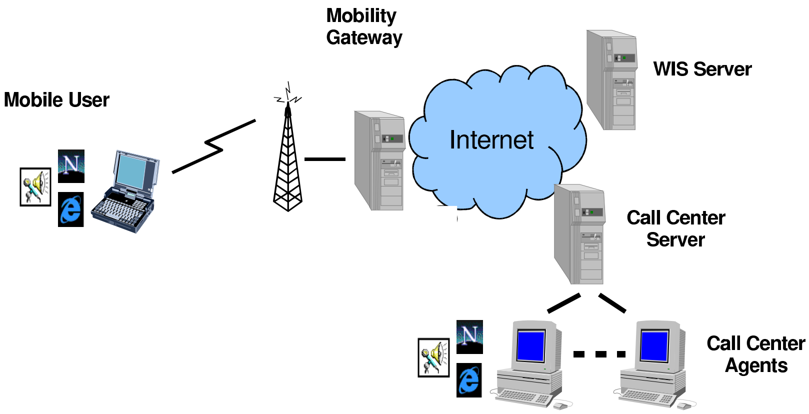

Figure 3. MOVE collaboration scenario

via a single, possibly wireless link, while simultaneously, a dynamic quality adjustment of integrated voice/data streams is provided. Similar to DIRECS, the MOVE middleware also provides location transparency of network nodes.

As a MOVE demonstration scenario, a collaborative call center application involving web-based VoIP phone conferencing and a hotel-booking service for mobile users, has been developed and shown to wide acclaim at the 1999 edition of ITU Telecom (which takes place every 4 years in Geneva). That exhibit marked the first time in the history of IP-based telecommunication that a seamless handoff between HSCSD-, GSM- and DECT-based mobile network base stations could be demonstrated in the kbit/s range; handoffs between Wave-LAN base stations even performed in the Mbit/s range.

Figure 3 shows that, along with the mobile user and the call centre, third parties such as a web-based information system (WIS) or any security service may be involved. The icons besides the screens of user and call center agents indicate that MOVE is not only independent of different browser technologies, but that it also adapts to the peculiarities of some of the more prominent web browsers on the market.

In MOVE, the streaming of voice, video and webcam data meet established transmission quality standards. Voice and other real-time data streams are transmitted with higher priority, relative to non-real-time data. In case the available bandwidth or other resources would otherwise not allow a timely transmission in spite of clever coding and sophisticated compression, visual data are downsized suitably. Typically, attributes of secondary importance are scaled down if necessary. For instance, fine-grained colour palettes may be coarsened, or mapped to gray tone scales or black-and-white binaries. Special attention is paid to the monitoring and adaptation of heterogeneous media streams concerning their adjustment to a wide range of quality-of-service parameters.

3 Adaptivity in MOVE and DIRECS

Adaptivity in WBCSs usually refers to one of two distinguished sets of issues. First, adaptivity caters for bridging differences caused by the geographic distance between network nodes and session participants, such as time zone, language, currency, cultural issues, etc. Second, adaptivity refers to a middleware- or web-service-supported adjustment to, or overruling of, technical idiosyncrasies such as bandwidth limitations, device properties, mobility, user profiles and locations, transmission delays and jitters, vendor-specific properties of hardware and software, as well as different types, capacities, carriers, bearers and backbones of network services and protocols. Neither MOVE nor DIRECS is concerned with the first kind of adaptivity, which consequently will not be further addressed in this paper.

The MOVE adaptivity is of the second kind, as documented in [7] and the project's final report [8]. In addition to its seamless adaptivity as characterized above, MOVE also offers a patented, dynamically adaptive quality-of-service support [9] [10] for the integration of real-time voice and non-real-time data communication over a single internet link, which responds to essentially the same parameters as those determined by the aforementioned network-, location-, vendor- and user-related technicalities.

The DIRECS adaptivity does not fall into any of the two categories mentioned above. Although it does offer adaptivity to network load, bandwidth limitations and other static and dynamic properties of WBC sessions, it does not do so by by means of any web service. Rather, the adaptation of DIRECS functionality to the requirements of collaborative sessions is realized, first by choosing, then by monitoring and then suitably exchanging pluggable networking protocols, if the monitoring results advise such an exchange. This is an administrator's job, which so far has not been automated of built into the middleware as an autonomous functionality.

In particular, the DIRECS adaptivity allows to optimize the tradeoff between requirements of availability and consistency of replicated information, by means of plugging in or exchanging dedicated protocols that perform according to a strategy which may favour either lazy or eager update propagation, optimistic or pessimistic concurrency control, and may opt for either a primary-copy or an update-everywhere policy. According to the fine-tuning of such parameters, more or less strong requirements of availability and consistency can be guaranteed. In general, a more relaxed attitude toward consistency (e.g., only some but not necessarily all replicas may need to be up-to-date at each moment) enables a higher degree of availability, while the enforcement of stronger consistency policies may entail certain delays in obtaining requested information. In other words, DIRECS protocols can be chosen such that an optimal adaption to the particular needs of collaborative applications and sessions is catered for, in terms of availability and consistency.

4 Enhancing DIRECS with MOVE (an outlook)

The integration of real-time data (e.g., voice) with non-real-time information (e.g., text or images) adds yet another dimension to the already complex situation of web-based collaboration architectures. Multimedia currently is not dealt with in DIRECS. However, since the architectures of MOVE and DIRECS are fairly distinct and orthogonal, an integration of both is conceivable.

A combination of the complementary capabilities of both middleware systems is expected to be particularly beneficial with respect to the demands of web-based collaboration communicated via mobile links, as well as for the potentially huge volumes of streamed web data. In general, different kinds of multimedia data can be expected to have quite different requirements with regard to timeliness, bandwidth demand and suitability for replication. For instance, voice tends to be elusive and ephemeral, and since it is rather "for the moment" than for the record, it is likely not to be stored, let alone replicated.

By and large, it will depend on the particular content and use of relational (in the traditional database sense) and textual documents, and multimedia objects in general, whether they are candidates for replication or not. As a consequence, consistency requirements can possibly be further relaxed for specific kinds of applications involving web-based collaboration.

We expect that a closer requirement analysis for such applications will bring about new protocols for supporting an even higher availability than what can currently be guaranteed already by DIRECS. The cost for that is paid by further relaxations of consistency properties achieved by the replication protocols that are currently benchmarked in our prototype implementation of MADIS.

A particular application envisaged for DIRECS, which, when enhanced with MOVE features for collaboration, would support voice/data streaming and rendering for web-based collaboration, is e-learning.

References

1. J. Bataller, H. Decker, L. Irún, F. Muñoz: Replication for Web-Based Collaboration. To appear in *Proc. 4th Int'l Workshop on Web Based Collaboration*, 15th DEXA, IEEE Computer Society, 2004.
2. http://www.iti.upv.es/madis
3. H. Decker, F. Muñoz, L. Irún, P. Castro, A. Calero, J. Esparza, J. Bataller, P. Galdámez, J. Bernabéu: Enhancing the Availability of Networked Database Services by Replication and Consistency Maintenance. *Proc. Parallel and Distributed Databases*, 14th DEXA, 531-535. IEEE Computer Society, 2003.
4. http://globdata.iti.es
5. http://www.cordis.lu/infowin/acts/rus/projects/ac343.htm

6. M. Wallbaum, D. Carrega, M. Krautgärtner, H. Decker: A Mobile Middleware Component Providing Voice over IP Services to Mobile Users. *Proc. ECMAST '99*, 552-563. Springer LNCS Vol. 1629, 1999.
7. D. Carrega, H. Muyal, H. Decker, M. Wallbaum: Integrating voice and data services for mobile internet collaboration with the MOVE middleware architecture. *Proc. 1st Int'l Workshop on Web Based Collaboration*, 12th DEXA, 532-536. IEEE Computer Society, 2001.
8. ftp://ftp.cordis.lu/pub/infowin/docs/fr-343.pdf
9. H. Decker, M. Krautgärtner: Flexible Quality-of-Service Technology for Supporting Voice/Data-Integrated Nomadic Networking. *Flexible Working – New Network Technologies*, 215-224. IOS Press, 1999.
10. M. Wallbaum, H. Decker: Enabling Quality-of-Service Management for Voice/Data-Integrated Mobile Communication. *Proc. ICDCS 2000 Workshops, section D: Wireless Networks and Mobile Computing*, 87-93. IEEE Computer Society, 2000.