The Coordination Technology Area at XRCE Grenoble: Research in Support of Distributed Cooperative Work

Jean-Marc Andreoli, Christer Fernstrom
Natalie Glance, Antonietta Grasso

Xerox Research Centre Europe

Overview of XRCE

The Xerox Research Centre Europe, XRCE, is one of four research centres of the Xerox Corporation and the only one located outside North America. XRCE comprises two laboratories – the Cambridge (UK) lab [21] and the Grenoble (France) lab – and an Advanced Technologies and Systems group (ATS), also located in Grenoble. The two labs have complementary research agendas and ATS is responsible for validating research results in a business context.

In addition to the XRCE Grenoble research lab and the advanced technology group, the Grenoble location also hosts three business units in charge of the commercialisation of products based on XRCE technology and a Technology Showroom, which provides a unique set of demonstrators of advanced technologies from different research teams of Xerox. In fact, one of the major issues for industrial research units is how to transfer research results into operational entities. The combination of research groups with development teams and professional "presenters" of advanced technology under a single roof appears to improve how our research work impacts the business units of Xerox. Technology transfer is by necessity a two-way process, which needs a common ground for exchange of ideas between researchers and practitioners. All in all, XRCE has been designed to create a research environment with facilities to provide good exposure to results and to facilitate the transfer of research results into business operations.

The Coordination Technologies Area: History

When the Centre was founded in 1993 a natural language group was created first around a core team of existing Xerox researchers. Then a second group, originally called "Distributed Systems," was initiated to focus on distributed software systems and applications. The distributed systems group was started with people whose main competencies were the theoretical foundations of distributed computing. The group's initial focus was strongly application-oriented and its first goals were to identify application domains of interest to the corporation and sufficiently novel to justify a significant research investment.

At the time, the corporation was undergoing a major shift, the result of which was to stress the importance of software in addition to traditional products that had made the corporation successful (copiers, printers etc.). At the same time, the extraordinary development of the Internet led to the notion that software components, even those embedded inside stand-alone boxes, would need to communicate among themselves and interact in various ways. Documents, formerly thought of as simple sheets of papers physically moving through hardware boxes, were now being re-conceived as electronic pieces of data moving virtually over a network through software components.

The need for distributed software support therefore appeared clearly. As its first set of research priorities, the distributed systems team chose to focus on coordination models and applications, and thus renamed itself "Coordination Technologies" (CT). One direction of research for the new team was to elaborate high-level platforms built on top of basic communication infrastructures and tackling the typical coordination issues of heterogeneity and autonomy. The design of such coordination middleware platforms was guided by application needs that came out of the team's investigations of two application domains for such platforms, information retrieval and workflow.

The Coordination Technologies Area: Today

As time passed, the application domains investigated by the team were enlarged, and focus was shifted from distributed *systems* (in the IT sense) to support for distributed *organisations*. As work processes become global and organisational knowledge becomes distributed over multiple sites and databases, it becomes increasingly difficult to develop and maintain business applications to manage collaborative work processes, and to maintain and share the distributed knowledge that will constitute the principal capital of many of these organisation. As a consequence, we identified two major areas of work: support for decentralised work processes and support for the sharing of decentralised knowledge.

Today, the Coordination Technologies Area, consisting of a team whose size varies around 14 people, explores these topics within three work programs:

- Basic Research, which deals with the fundamental technologies involved
- Applied Research, which applies these fundamental technologies to provide services and tools
- User Evaluation, which aims to combine the results from the other work programmes and evaluate how these can support real users outside the laboratory (inside or outside Xerox).

Currently, there are two threads in the Applied Research work program: Distributed Work Coordination (DWC) and Knowledge Agents (KA). These two threads build upon the foundation established by the Basic Research Program.

Distributed Work Coordination

The DWC research program investigates advanced work-intensive applications, in particular in the domains of workflow and electronic commerce. For infrastructure needs, we pursue the development of the distributed object coordination middleware platform CLF, which proved very useful in supporting the complex kind of interactions which characterise both workflow and electronic commerce applications. In particular, we fully exploit the negotiation and transaction facilities offered by CLF [3][9].

Concerning CLF itself, the current focus is adding reliability, security and monitoring features in the context of coordination. A demonstrator of a deployment-monitoring tool is available [22].

With respect to electronic commerce, we have started a new project, Xpect [2][5], which focuses on commercial transactions involving several partners (providers of goods, financial institutions, delivery companies for non electronic goods) with whom a customer may wish to negotiate via a "broker." We are mainly interested in the specification of brokers, who must find a consensus between the customer and the other commercial actors, and then enact the agreed transaction when possible. The main challenge resides in the deep interaction between the two phases (agreement-enactment).

In the workflow domain, we have several ongoing projects:

- Newel: flexible workflows for distributed virtual organisations. This project draws upon technologies developed in the earlier projects WebFlow [7][20] (transaction-based distributed workflow management) and Zippin [1][18] (declarative modelling for dynamic scheduling and simulation of work processes).
- PS-Flow: application of our workflow technologies to the specific application domain of distributed printing (joint work with the ATS group and one of Xerox business groups).

We are interested in complex applications where:

- all work is completely distributed
- work processes need to evolve dynamically (i.e. during their course of progress)
- different processes imply different models of cooperation
- workflows sometimes need to be enacted within a framework of constraints that define the use of resources and time

Typical workflow tools provide support for managing flow dependencies between the steps in a workflow process. It is for example possible for a process designer to state that when task A finishes, tasks B and C should start concurrently, and when both are finished task D should start. Based on such a description, the workflow system will build up its internal structures in order to manage the dynamic flow of tasks in all instances of the workflow process. Workflow systems that operate along these principles are based on a model, which is purely event-driven, deterministic, centralised and use statically defined process descriptions. This operation model has proven very efficient for the handling of large-volume, event-driven, deterministic processes, such as those found in finance institutions (insurance, banks...), customer care (help desk...) and in some cases for software engineering (bug tracking, configuration management...). However, it introduces a number of restrictions that hamper the general applicability in many other situations.

In purely event-driven and deterministic workflow systems, an external event, such as the arrival of an order via an EDI (Electronic Document Interchange) interface, initiates a new process instance. The end of each process step generates an internal event, which signals to the enactment engine that the next step can be taken. While such systems are good at handling events, they are usually not good at handling time. At the best they can detect alarms from overruns and handle these in a reactive manner. However, they are not pro-active, in the sense that they would reschedule tasks (and possibly entire processes) to avoid overruns. And since they are not good at managing time, they are not good at managing resources either. While they can typically take on-the-fly decisions to select the resource for a new task to be the one which is currently least occupied, they have no way of knowing that a huge chunk of work may be on its way to the resource from another process.

Further more, since the workflow process descriptions are deterministic (which often means over-specified, since they include artificial restrictions on the way things are done), this gives little opportunity for a human process manager to interact with the system in order to dynamically reschedule work. Our work on constraint-based scheduling for workflow processes addresses this problem.

Current trends in workflow technology indicate a growing interest in supporting distributed organisations. However, the support provided is limited to either coarse-grained interoperation between different work processes running at different locations, or to the distribution of single tasks from a centralised server. One of our concerns is how to support organisations where there is not one single natural central location, and where work is distributed on varying levels of granularity. Within a complex process in such an organisation, anything from a large part of the process to a single task can be allocated and reallocated to any location in a potentially very large network. It may also be the case that tasks or processes are proposed to more than one location, and the definite allocation should be decided through a negotiation process (human or automatic) [1][17][18].

Another aspect, which is one of the most lively research subjects in workflow, is the problem of how to allow dynamic changes to workflow process definitions (often referred to as *process evolution*). This problem, which takes on a new dimension when workflow processes are completely distributed, is the subject of our studies on supporting distributed workflows with inter-site negotiation [10].

Knowledge Agents

Why call our second research program knowledge agents? Knowledge is an over-used word in the world of business just as agent is an over-used word in the field of computer science. What do we hope to gain by marrying these two over-extended terms together?

First of all, what is an agent. We tend to agree with a definition of agent that includes such concepts as: a software module acts autonomously, understands a common communication protocol, and interacts with other software modules to assist a user. Agents appeal to us for two reasons. On the one hand, the notion of software agent or intelligent agent moves us away from the, let's admit it, esoteric realm of computer programming to one that's more accessible to common folk. Agents are sexy, agents make sense, and agents are here to stay for a while. Normal people (people who watch television instead of programming Java applets for their home pages) can understand what agents are, especially if you refer to agents as butlers or personal assistants. Normal people might even be able to "program" (or tune) agents if we design them properly.

Secondly, the concept of intelligent agent waters down the highly ambitious goals of artificial intelligence of the seventies and eighties. Instead of *replacing* human intelligence all we're trying to do with intelligent agents is *augment* human intelligence, something civilisation has been doing for centuries via other kinds of software and hardware. Admittedly, the word agent is over-used, but it's over-used for a good reason: it's the right metaphor.

What is knowledge. Talk to ten different people and you get ten different definitions of knowledge. Most often knowledge is used to mean something that is more than information, just as information is used to mean something that is more than data. We like to think of knowledge as the stuff that's in people's heads and information as the mechanism used to communicate that knowledge. At best we can have direct access only to the information, not to the knowledge.

In fact, thanks to networked organisations, we now have access to huge amounts of information, what many would agree has become, in fact, overwhelming amounts of information. Every day knowledge is being pumped into new information. In light of this, we think some interesting questions are: how can we pump information back into knowledge? What can we learn about the structure of collective knowledge from the streams of information?

Back to the original question: why knowledge agents. We've said that intelligent agents or software agents typically refer to personal digital assistants that act on behalf of the user. More particularly, we use the concept of knowledge agent to refer to a kind of agent that taps into the collective knowledge of communities of users in order to assist individual users. The body of collective knowledge might include recommendations, evaluations, WWW browsing patterns, e-mail communication patterns, newsgroup activity, to name a few.

We envision that knowledge agents will assist users in many ways within domains of application such as digital libraries, knowledge management and community networks. Some knowledge agents will be responsible for pumping information back into knowledge. For instance, one kind of agent could perform community-informed information retrieval, while another could make community-based recommendations. We have two projects here within CT focusing on these two kinds of agents, described below in more detail. Other agents, gatherers, will be responsible for pumping knowledge into information.

Other knowledge agents will operate on the meta-level, mapping the structural knowledge of the usage data generated by networked users and their dedicated agents. For example, competency mappers will help identify centres of expertise within the organisation. Community mappers will help identify extended communities of interest. Content mappers will help identify to what extent collective knowledge is represented in on-line digital repositories. We will be starting a new activity around knowledge agents for mapping structural knowledge in the near future.

Our first efforts have focused on two complementary kinds of knowledge agents, called Knowledge Brokers and the Knowledge Pump. Knowledge Brokers helps people search for information across heterogeneous databases. Knowledge Pump recommends relevant information to people based on their interests, usage patterns, registered profiles and community membership.

Knowledge Brokers

The current implementation of Knowledge Brokers [8] actually involves three kinds of agents: *a)* users, who input queries and process answers through a GUI; *b)* wrappers, capable of interrogating heterogeneous information sources over the Web; and *c)* brokers, which can manage complex queries (i.e. decompose a complex query, recompose the partial answers, synthesise a full answer). The core is given by the brokers that provide various important services such as intelligent caching, filtering and knowledge combination. Brokers use a common knowledge representation language based on constraints [4][6] both for communication and for the internal representation of their state.

Knowledge Brokers has a number of distinctive features. They allow complex queries, dynamic refinement [11], and multiple searches in parallel. The wrappers provide an interface between the brokers and the final database. These wrappers can impose structure on raw information such as Web pages and can be written semi-automatically for new sources of information [14][15]. The rapid and semi-automatic construction of wrappers makes Knowledge Brokers very useful in Intranet and Extranet settings, where companies need to connect legacy databases with a simple user interface. Knowledge Brokers provide a very simple infrastructure to query multiple sources of information concurrently, regardless of their structure, protocol or output formats.

Knowledge Brokers qualify as knowledge agents in the sense that they are able to combine the information in multiple databases. In addition, more recently we have been investigating the possibility of combining broker agents with profile agents to permit collaborative browsing [19]. By tracking users' preferences across different activities, such as searching and reading and recommending, we hope to be able to locate the most relevant items in response to a search process.

Knowledge Pump

Knowledge Pump provides a service complementary to that of Knowledge Brokers. When users sign up, they join communities of people with similar interests. Profiler agents track and map each user's interests, learning more about the person each time (s)he uses the Pump. A recommender agent finds matches between new items and user preferences, automatically sending relevant and high quality information to people as it is found.

The core of the Knowledge Pump is the recommendation functionality that is based on community-centred collaborative filtering which filters both by content and by taste [16]. The Pump handles content filtering by relying on recommenders to classify items into pre-defined communities. Social filtering – matching items to people by first matching people to each other – is accomplished using a combination of statistical algorithms and boot-strapped snapshots of the underlying social network of a collection of users.

In parallel with the development of particular kinds of knowledge agents, we are also defining knowledge agent architecture. Key features are that a knowledge agent can be an aggregation of other agents and that user applications can be modularly built from subcollections of the top-level knowledge agents. A well-designed architecture will allow us to focus on novel methods for managing knowledge at the agent level and novel functionalities for assisting users at the application level.

User evaluation of novel technologies

Xerox research labs have a long tradition of paying close attention to user needs and envisioning alternative interfaces to the keyboard and monitor. We at CT are leveraging this tradition by collaborating with internal and external research institutes to test our technologies in real user environments.

This is particularly relevant to an industrial research centre since we are driven by effective transfer of our technology. In particular we are aware that too often technology transfer fails through being ill informed about users needs and perceptions.

The way we address this difficult aspect is by having a user evaluation program where technologies are enhanced and applied to particular domains. This enables us to work with end users to collect early feedback about the usage of novel technologies and interfaces.

This approach has led us to participate in the Campiello project, funded through the Esprit i^3 (intelligent

information interfaces) framework as a LTR (Long Term Research) project. Campiello offers us a real user setting in the tourist domain. Moreover, in this particular project, we are evaluating CT technologies by making them available through the development of paper and portable interfaces, using the tourist domain as a test-bed not only of the tools, but also of novel user interfaces to them.

Campiello's objective is to recreate the broken links between several different communities: the local inhabitants of tourist towns, the cultural managers organising events and the tourists. Identified problems include: the lack of awareness of locals about cultural events promoted by cultural managers, lack of feedback to the cultural managers about the events, superficial visits of the towns by the tourists, the lack of means for tourists to share experiences with other tourists and interested locals.

Our second project is to work with an internal business division under this program to directly transfer technology. In this case we are enabling them to both benefit internally from the early use of our Knowledge Pump and Broker technologies and shorten the time it will take them to adopt the systems as part of their commercial offers.

Overall, the user evaluation program allows us to refine the formulation of the addressed problems and develop potential “business models” to support the real usage of our tools.

Summary

From a small initial core group in 1993 focusing primarily on distributed systems, CT has evolved significantly both in size and in breadth in the past four years. CT's current research mission spans coordination and cooperation in distributed work and knowledge processes. We aim to take into account the spectrum that goes from the platform to the user. We focus particularly on research issues of coordination, transactionality and dynamicity that make up current outstanding problems in workflow and electronic commerce. In addition, we explore how to best leverage the collective knowledge of users to improve the services we can provide them with.

Finally, the user evaluation program that we have recently initiated helps ensure that our work is user-driven and grounded by the requirements collected in real world settings. Overall, the structure of our group has evolved to reflect the twin tensions found in many research groups today: the need to provide business value on the one hand and innovative technologies on the other.

References

[1] Andreoli, J. M., S. Castellani, U. Borghoff, R. Pareschi, and G. Teege (1998): Agent-Based Decision Support for Managing Print Tasks. In *Proc. of PAAM'98* (to appear).

[2] Andreoli, J. M., and F. Pacull (1998): Distributed Print on Demand Systems in the Xpect Framework. In *Proc. of EC'98* (to appear).

[3] Andreoli, J. M., F. Pacull, D. Pagani and R. Pareschi (1998): Multiparty Negotiation for Dynamic Distributed Object Services. In *The Journal of Science of Computer Programming* (to appear).

[4] Andreoli, J. M., U. Borghoff, R. Pareschi, S. Bertarelli, U. Montanari, and F. Rossi (1997): Constraints and Agents for a Decentralized Network Infrastructure. In *Proc. of the AAAI'97 workshop on Constraints and Agents, July 27, 1997, Providence, RI, U.S.A.*, pp. 39-44.

[5] Andreoli, J. M., F. Pacull and R. Pareschi (1997*):* Xpect: A Framework for Electronic Commerce. In *IEEE Internet Computing*, 1(4).

[6] Andreoli, J.M., U. Borghoff and R. Pareschi (1997): Signed Feature Constraint Solving. In *Proc. 3rd Int. Conf. on Practical Application of Constraint Technology (PACT '97), April 23-25, 1997, London, U.K.*, pp. 35-46.

[7] Andreoli, J. M., J. L. Meunier and D. Pagani (1996): Process Enactment and Coordination. In *Proc. of the European Workshop on Software Process Technology (EWSPT'96)*, October 8-11, 1996, Nancy (France).

[8] Andreoli, J. M., U. Borghoff and R. Pareschi (1996): Constraint-Based Knowledge Broker Model: Semantics, Implementation and Analysis. In *Journal of Symbolic Computation 21(4)*, pp. 635-667.

[9] Andreoli, J. M. (1996): Coordination as Negotiated Transactions. In Freitag, Jones, Lengauer and Schek (Eds): *Object Orientation with Parallelism and Persistence*, Kluwer Academic Publishers.

[10] Borghoff, U., P. Bottoni, P. Mussio, and R. Pareschi (1997): Reflective Agents for Adaptive Workflows. In *Proc. 2nd Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97), April 21-23, 1997, London, U.K.*, pp. 405-420.

[11] Borghoff, U., P. Y. Chevalier and J. Willamowski (1996): Adaptive Refinement of Search Patterns for Distributed Information Gathering. In Verbraeck, A. (ed.) *Proc. of the 1st Intl. Conf. EuroMedia/WEBTEC, December 19-21, 1996, London (U.K.)*, San Diego, CA: The Society for Computer Simulation, pp. 5-12.

[12] Borghoff, U., and Johann H. Schlichter (1996): On Combining the Knowledge of Heterogeneous Information Repositories. In *Journal of Universal Computer Science*, 2(7), pp. 515-532.

[13] Borghoff, U., R. Pareschi, H. Karch, M. Nöhmeier and J. H. Schlichter (1996): Constraint-Based Information Gathering for a Network Publication System. In *Proc. of the 1st Intl. Conf. on the Practical Applications of Intelligent Agents and Multi-Agent Technology (PAAM'96), April 22-24, 1996, London (U.K.)*, pp. 45-59.

[14] Chidlovskii, B., U. Borghoff, and P. Y. Chevalier (1998): Boolean Query Translation for Brokerage on the Web. In *Proc. Int'l. Conf. EuroMedia/WEBTEC, Leicester UK, Januar 5-7, 1998*, San Diego: SCS.

[15] Chidlovskii, B., U. Borghoff and P. Y. Chevalier (1997): Towards Sophisticated Wrapping of Web-based Information Repositories. In *Proc. of the 5th Int'l RIAO Conf., June 25-27, 1997, Montreal, Quebec, Canada*, pp. 123-135.

[16] Glance, N., D. Arregui, and M. Dardenne (1998): Knowledge Pump: Supporting the Flow and Use of Knowledge. In U. Borghoff and R. Pareschi (eds.) *Information Technology for Knowledge Management*, Springer Verlag.

[17] Glance, N., S. Castellani and D. Pagani (1996): Simulation for Optimization vs. Role-Playing: A Case Study of a Document Authoring Process. In *PAKM'96 Workshop on Adaptive Workflow, October, 1996, Basel, Switzerland*.

[18] Glance, N., D. Pagani and R. Pareschi (1996): Generalized Process Structure Grammars for Flexible Representations of Work. In *Proc. of the ACM Conference on Computer Supported Cooperative Work (CSCW'96), November, 1996, Boston (Ma, U.S.A.)*.

[19] Grasso, A., U. Borghoff, N. Glance, and J. Willamowski: Collaborative Information Gathering In *Proc. Int'l. Conf. EuroMedia/WEBTEC, Leicester UK, Januar 5-7, 1998*, San Diego: SCS.

[20] Grasso, A., J. L. Meunier, D. Pagani and R. Pareschi (1997): Distributed Coordination and Workflow on the World Wide Web. In *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, 6(1), pp. 1-26, 1997.

[21] MacLean, A. (1997): From EuroPARC to XRCE Cambridge: Ten Years at the Human-Computer Interface. In *SIGGROUP Bulletin*, 18(3).

[22] Vion-Dury, J. Y., and F. Pacull (1997): A Structured Interactive Workspace for a Visual Configuration. In *Proc. of 1997 IEEE Symposium on Visual Languages (VL'97), Capri (Italy), September 23-26, 1997*, IEEE Computer Society.