Applications and Research directions of the Object Systems Group

Dimitri Konstantas
(Editor)

University of Geneva
Centre Universitaire d'Informatique
24 rue General-Dufour
CH-1211 Geneva 4
Switzerland

tel: +41 22 705.76.64
fax: +41 22 705.77.80
e-mail: dimitri@cui.unige.ch

Abstract

We present an overview of the research directions, projects and applications of the Object Systems Group.

In :

"Object Applications", ed. Dennis Tsichritzis, University of Geneva, Centre Universitaire d'Informatique, Aug. 1996.

Applications and Research directions of the Object Systems Group

Dimitri Konstantas
(Editor)

The papers presented in this report give the status of some of the ongoing or recently completed projects of our group. However they neither cover the complete spectrum of our activities, nor do they provide view of the links between the different activities. In this Annex we give a short description of the ongoing group activities and describe the links and interactions between them.

The research directions of our group follow two research streams. The first is the fundamental research stream and the second is the application stream. In the fundamental research stream we investigate and develop basic concepts and ideas related to object oriented technology, while in the application stream we develop and integrate multimedia applications based on object oriented technology. The two research streams are tightly coupled so that requirements and problems from the application projects are passed to the fundamental research projects and results from the fundamental research projects are directly used in the application projects. The ongoing and recently completed fundamental research projects of our group are

- *COORDINATION* : A *European Union (EU)* ESPRIT project dealing with coordination of active software entities,
- *Object Integration* : A *Swiss National Science Foundation (SNSF)* project concerned with object-oriented software integration at various levels, and
- *Mobile Object Systems* : An SNFS project concerned with the software infrastructure for mobile object systems.

From these projects the first one, *COORDINATION*, was completed in April 1996 and two papers [10][11] describing the results of the projects are included in this report. The second project, *Object Integration*, is an on-going project that will be completed at the end of September 1996 and different parts of its results are described in [6], [9] and[11]. Finally the last project, *Mobile Object Systems*, will start in October 1996.

In the application oriented research stream we have 5 projects, out of which 2 have been already completed, and 3 are ongoing. Namely we have

- *SHIPS* : A EU ESPRIT project that was targeting in the design and implementation of the OS for a super-computer,
- *CHASSIS* : A *Swiss Priority Program - Informatics (SPP-IF)* project targeting in the design and implementation of interoperating platform,
- *DVP - Distributed Video Production* : A EU *Advanced Communication Technologies and Services (ACTS)* project for the production, post-production and dissemination of digital video and multimedia material,

- *MEDIA* : A set of *Swiss Priority Program - Information and Communication Structures (SPP-ICS)* projects that target in the design and implementation of a platform for the commercial dissemination of electronic information. The project is composed of three component projects, *ASAP*, *HyperNews* and *KryPict*.

- *Cooperative Learning* : A project aiming in the design of an augmented reality learning environment, aiming in the integration of the results of the MEDIA and DVP projects.

From these projects the first two, *SHIPS* and *CHASSIS*, were completed in 1995 and the some of their results are described in [7] and [6] respectively. The *DVP* is an on-going project, started in September 1995, establishing a collaboration between our group and the Computer Vision Group of CUI directed by Prof. T. Pun. A description of the project can be found in *http://viswiz.gmd.de/DVP/Public* and some intermediate results are presented in *http://cuiwww.uni-*

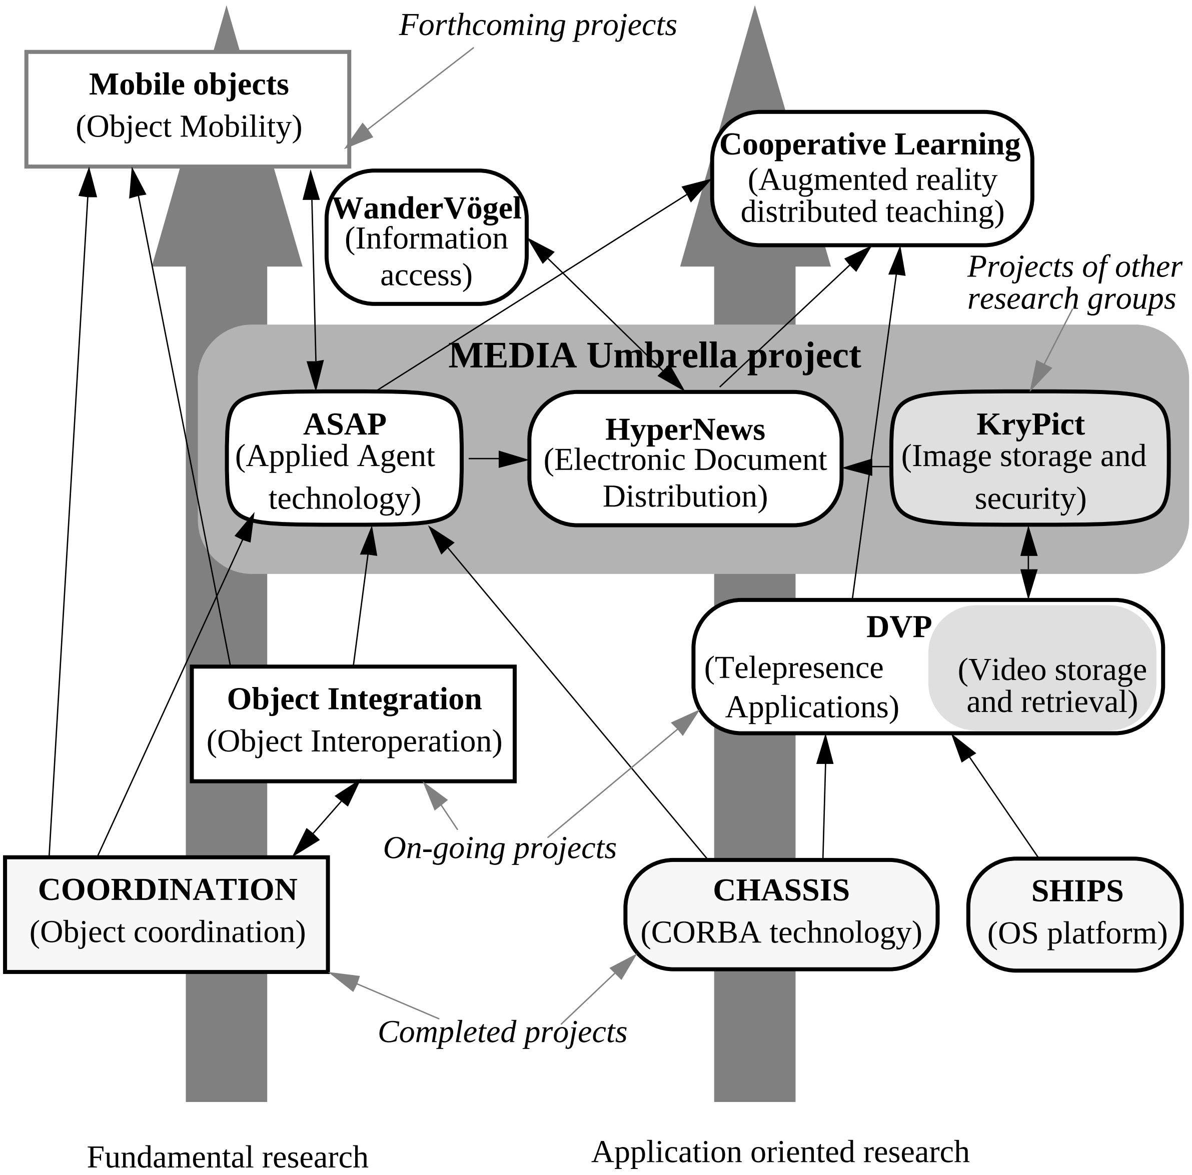

Figure 1 Overview of project interrelationships

ge.ch/OSG/DVP/wp43.html. The *MEDIA* project was started in June 1996 and is a collaborative project between our group and the Computer Vision Group. An overall description of the project goals is given in [2], a description of the goals and ideas of the component project *HyperNews* is given in [1] and [4], while an overview of the *KryPict* component project lead by the Computer Vision Group is given in [3]. Although there no papers regarding the *ASAP* component project this report, a selected background research is presented in [8]. It is to be noted that the *ASAP* project provides the link for the transfer of results from the fundamental research projects to the application projects, and in this sense it is shared between the two research streams of our group. Finally the *Cooperative learning* project is an on-going project started in late 1995 and for which some intermediate results are given in *http://cuiwww.unige.ch/OSG/DVP/wp43.htm*l. An overview of all the projects along with their interrelationships and status is given in Figure 1.

The long term directions of our research are focused into three major targets. The first target, in the fundamental research stream, is the development of the concepts and ideas for the design and implementation of an agent system that will support mobility and dynamic coordination between active software entities.

The second target, linking both the application and fundamental research streams, is the design and implementation of applications using agent technology on top of the Internet and the WWW, for publishing and dissemination of information. The third target, in the application oriented stream, is the design and development of telepresence applications based on broadband and satellite communication technologies. In order to achieve these targets we are actively involved in establishing collaborations and communications channels with other European and international research centers and companies.

In the rest of this paper we provide a brief overview of each of the on-going projects of our group and namely, Object Integration, Mobile Object Systems, MEDIA, WanderVögel, DVP and Cooperative Learning.

2. Object Integration

SNSF Project (20-40'592.94)

Following previous work done by our group on object-oriented programming languages and systems, this project is focuses on various language aspects, with a particular emphasis on integration of object-oriented mechanisms with other computing paradigms. Our research is focused into two areas: Objects and functional programming and Object interoperability.

Objects and Functional Programming

Although still neglected by industry, functional programming is gaining more and more interest in the academic world, because of its high level of expressiveness, its support for formal proofs of program properties, and its sophisticated type-checking mechanisms. The two competing languages ML and Haskell represent the main trends in this area. Acknowledging the success of objects, both communities try to incorporate principles of object-oriented programming in their languages; however, this involves a number of nontrivial technical issues for which no satisfying solutions are known so far. Our research focuses in studying different issues related with problems of integrating object-oriented programming and functional programming. The main points of our research can be summarized as following:

Formal foundation. In the past few years our work concentrated on the definition of a modified lambda-calculus, called Lambda-N, which is expressive enough to model both objects and functions. This contrasts with the "calculus of primitive objects" of Abadi and Cardelli, which has more built-in support for object-oriented constructs, but in which the encoding of functions is somewhat contrived. Since the first presentation of the Lambda-N calculus in 1993 we continued our work and we simplified the formalism, bringing it even closer to the traditional lambda-calculus, established an axiomatization whose soundness is proved through an operational simulation observing the error generation behavior of terms and extended the well-known Hindley-Milner type inference algorithm for functional languages to accommodate subtyping.

The HOP language. Building on top of the formal foundation, we have designed and implemented a functional language [14] with some support for dynamic binding, and hence for object-oriented programming constructs.

Objects, Interoperability and Distribution

Interoperability issues have been known for long in traditional, procedural computing, but they take a different aspect in the context of object-oriented systems. So we started by defining in detail the notion of object-oriented interoperability [13] and the principles and problems for the interoperation of object-oriented applications. A thorough understanding of the issues is interesting both in the perspective of software evolution - reusing and composing software components, modifying software systems to meet evolving requirements -, and in the perspective of distribution. To confront our principles with practice, we studied both areas: we completed a case study for the interoperation of legacy applications under CORBA [6], and we studied application interoperation from the point of view of the exchanged data in electronic documents' publishing. This progressively led us to the notion of *agent*s as the basic construct for achieving interoperability in distributed applications.

3. Mobile Object Systems

SNSF Project (20-47181.96)

Mobile object systems represent the building blocks of the Software Agent technology that will become a standard for implementing electronic commerce over the Internet. The goals of this project are to investigate theoretical foundations of mobile object systems, in particular modularity aspects; linguistic constructs and mechanisms for mobility support in object-oriented systems; and practical aspects dealing with efficiency and implementation techniques. The success of the World Wide Web has proven the viability of widespread distributed infrastructures and been the impetus for research in more general architectures which would overcome the limitations of the existing technology. A particular goal is to open up the Internet to all sorts of commercial applications. The current standardized interface of the World Wide Web restricts communications and hampers the development of advanced distributed applications. The next step in the evolution of the Internet is to support an open interface for the transmission and execution of software agents. Such an open interface will change the nature of computer communication: instead of exchanging passive data, computers will be sending active programs over the network. These active programs have been called Software Agents, or Mobile Software Agents. Recently, a number of commercial and research projects have investigated this approach. The most widely publicized efforts are the AT&T funded TELESCRIPT language (General Magic) and Inferno, the JAVA language from SUN Microsystems; other commercial systems include Joule (Agorics), Object REXX (IBM) and Smalltalk/Agents (QKMS). Some outstanding research systems are Obliq (Digital), Tycoon (Uni Hamburg), Phantom (Trinity College), Facile (ECRC), MOLE (IVPR, Stuttgart), Ara (Uni Kaiserslautern), Tcl Agents (Dartmouth) and Map (IMAG Grenoble). The common characteristic of all of these projects is the emphasis on mobile computation and remote programming, and the prominence of the object-oriented paradigm.

The basis for implementing agent applications are *mobile object systems*. A mobile object system is a collection of executing objects which may, at any point in time, interrupt its computation and move to an other execution environment located on another machine and resume its computation there. It is based on object-oriented technology, which now is a well-established field of programming languages and software engineering, and owes much of its success to the fact that it proposes some mechanisms which are very helpful for handling the complexity of software systems. However, these mechanisms do not offer a perfect solution: they are still criticized for example for their lack of efficiency, lack of generality, or poor security when composing several modules into a complex software system; moreover, support for mobility in current languages is either nonexistent or consists of ad hoc solutions, poorly adapted to address the main issues in their full generality. Consequently, there is a wide agreement that current programming languages need to be improved with better mechanisms for secure composition and evolution of software. This is acknowledged for example by a recent joint research program DARPA/NFS in the United States, dedicated exactly to this topic (cf ARPA BAA 95-40). Some answers are likely to appear within the framework of mobile object systems, which address distributed programming from a different angle than remote procedure call (RPC) based distributed systems. Where RPC based systems rely on high-level interfaces to communicate passive data between remote sites, mobile object systems offer a flexible lower level interface for the exchange of active programs. Mobile object system are geared towards distributed computation in very-large, distributed and heterogeneous computer systems where security and performance are key concerns.

This new paradigm for constructing distributed applications requires advances in techniques for software composition. In particular, programming languages need to support much *safer composition* of software components; they need to be more *expressive*, to ease the construction of agent applications; they have to be designed for *efficiency*, to be able to scale to very large configurations. The assumptions that execution environments may be heterogeneous and loosely coupled, and that the same software components need not be installed on each platform, imply that some or all of the code required by an object may have to be moved and linked with a new local environment. In addition, security issues become more acute, since the same execution environment may host object systems belonging to different users with different access rights and different resource allocations. We articulate our research as following:

1. *Static Analysis of Modular Systems*: Static analysis has been used to optimize object programs (e.g. SELF), and to ensure against errors (e.g. the byte code verifier in JAVA, or link-time consistency check in Eiffel). Unfortunately, all existing algorithms are inherently global and thus can not be used for distributed systems and systems with dynamically evolving configurations. The goal of this research is to investigate practical techniques for analyzing program fragments and composing the result of analysis.

2. *Mobility Support in Object-Oriented Languages*: design and implement language extensions for building application with mobile objects.We will investigate three extensions to the object-oriented approach to ease the task of programming mobile object systems: *user controlled dynamic linking*, *object islands,* and *generative communication.*

Both are directions in language design; they will be supported by parallel research efforts on formal semantics and implementation techniques.

4. MEDIA - Mobile Electronic Documents with Interacting Agents

SPP-ICS Project (5003-45332)

MEDIA is an umbrella project [2] linking together three individual projects, and namely *HyperNews*, *KryPict* and *ASAP*. The goal of MEDIA is to develop a system for archiving and dissemination of electronic documents under similar conditions as printed documents, using an agent-based, distributed, and secure platform.

The objectives of the three MEDIA component projects are:

HyperNews : development and implementation of a prototype system for the distribution of electronic documents in a commercial environment. The target project of HyperNews consists of an electronic newspaper application that will allow the commercial distribution of news over the WWW network;

KryPict : development of copyright enforcement and document authentication methods that will allow to digitally watermark images. The digital watermarking methods will be included into a target pictorial database subsystem; digitally watermarked images from this subsystem will be encapsulated into the ASAP agents and made accessible through the MEDIA user-interface;

ASAP : key issues will be addressed for the construction of distributed applications over vast and dynamic networks of computers using *mobile software agents*. The concrete goal of ASAP is to deliver an execution platform that supports development of applications based on mobile software agents. This technology will be an infrastructure for developing *commercial information systems* in large scale, dynamic and heterogeneous networks of computers.

The functions and the associated components of the MEDIA system are the following:

- agents-based architecture: a secure platform for the execution of mobile software agents.

- database population and information archival: means for creating and managing a pictorial database, containing images and possible captions (Krypict), and a textual database, containing journalistic articles (Hypernews);
- information retrieval: search mechanisms allowing access to images and text by means of appropriate indexing structures (Krypict, Hypernews);
- watermarking: image copyrighting mechanisms, providing digital watermarks that are resistant to a variety of image modification methods (Krypict);
- user interface: interface allowing users to retrieve and browse the stored documents.
- security: the agent platform guarantees security at three levels: agent execution (virus protection), content (copyright enforcement) and transport (validation, authentication).

Results from MEDIA will be demonstrated by means of several multimedia databases for which distribution, ease of user interaction, security and copyright enforcement are critical issues: the journalistic database from l'Hebdo magazine; the historical database of ancient watermarks and papers, Swiss Paper Museum, Basel; the news photographs database from Télévision Suisse Romande, Geneva; the Database for Swiss Cultural Heritage, Bern.

5. WanderVögel

The issue of finding relevant information in an information system is a long standing problem. The problem appeared well before the advent of electronic computers and we can hardly suggest that the major advances in this area arose inside the realm of computer based information systems. Classification methods devised by natural sciences researchers during the last two centuries and classification schemes which belong to the old art of librarians should probably get the essential credits for those advances. What computers brought in an unprecedent way is the ability to store and manage huge amounts of information.

The goal of the WanderVögel [5] project is to design and implement an architecture that will support the classification of information based on its context. The WanderVögel architecture is based on the dynamic identification-action mechanism. Technically speaking, the kernel of this mechanism is a fuzzy-control engine. The fuzzy-control engine takes as input a running collection of partial document descriptions and attempts to identify the user's exploration profile. The fuzzy-control engine is "programmed" with rules that allow for associations between exploration profiles and actions performed by the system: when the engine matches a profile specified by the rules it fires the execution of an action.

6. DVP - Distributed Video Production

EU ACTS Project (AC089)

Video production is inherently distributed: Broadcasters are physically distributed over several sites and studios, they increasingly outsource video production and post-production to specialized studios or upcoming virtual studios. Thus there is an increasing demand for the enabling technology for distributed video production.

The technological revolution of digital media is rapidly changing the world of broadcasters and the media industry. In the emerging field of *media processing* we see computer science, communication technology and media technology coming together to explore new applications and define new products and new markets. These concerted technologies will enable media producers to remotely collaborate in digital video and multimedia production, post-production, archiving, indexing, and retrieval.

In the ACTS project “Distributed Video Production” (DVP), leading European broadcasters, computer, communication and media technology providers are joining their efforts to start an innovation initiative in remote contribution, production, post-production and dissemination of digital video and multimedia material. Additionally, an American supercomputer center and a Canadian educational institute will contribute distributed video applications and distributed virtual reality simulations. For these applications a broadband *transatlantic link* will be used between the involved partners for real-time simulations.

In DVP, basic technology for transferring studio-quality digital video over broadband networks (ATM), and video coding and compression is the starting point. Problems of transmission and processing delays, synchronization requirements, and quality of service have to be solved. Built upon these basic results, pilot applications are planned for a distributed virtual studio, distributed telepresence, distributed virtual reality, and distributed video archiving.

As an example, the *Distributed Virtual Studio* will enable TV producers to improve the economic efficiency and the time to market of broadcast program production. Production environments with ATM connectivity can have access to virtual studio techniques as an external service. By this means, ATM allows a “studio on demand” scenario, this enables TV producers to save equipment and reduce training and manpower costs. On the other hand, a new service industry is expected to develop, where information products like virtual sets are contributed in real time from an outside service provider to the program producer.

Several DVP application scenarios will demonstrate that the technology is ready for commercial exploitation:

Distributed Video Post-Production: Editing and special-effects generation using both local equipment (editing console, DVEs) and remote equipment (video servers).

Distributed Video Production: Bringing together real actors and objects (props) from different - separated - studios in a common real or virtual studio.

Distributed Rehearsal: An immersive teleconferencing environment allowing small groups of actors and musicians at different studios to conduct rehearsals as if face-to-face.

Distributed Video Archiving, Indexing, and Retrieval: Besides real-time applications like those mentioned above, some non-real-time applications will also contribute to DVP.

DVP addresses the urgent demands of broadcasting and media industry for distributed video production technology. It incorporates important American and Canadian partners over transatlantic broadband networks and will help push the development of European information and media industry to a world-wide leading level.

Participation of the Object Systems Group

Our group participates in the DVP project in the *Distributed Rehearsal* application (Work Package 4.3). The target of the application is to develop a telepresence environment for the support of distributed rehearsals. These rehearsals can be theater rehearsals, musical rehearsals, etc. The aim of the distributed rehearsal application is to allow the organization of rehearsals without requiring the participants (actors, musicians, conductors etc.) to be physically in the same room, eliminating thus the need to travel from one city to another for participating in the rehearsal.

In general the requirements for organizing a distributed rehearsal are these of a telepresence session. That means that the basic need is the existence of (at least) a video wall with minimal dimensions 2x3 meters and high resolution video projection, and at least a hi-fi audio system. The goal is to give the impression to the rehearsal participants that they are physically in the same room.

More details and information for the DVP project can be found in [15], while intermediate results for the Distributed rehearsal application are presented in [16].

7. Cooperative Learning

Various forms of tele-education are currently under experiment, using either television broadcast networks or computer networks; their rapid development is motivated by a high demand from the public. In a large number of situations, tele-education can spare resources or can solve practical difficulties, like teaching to the disabled, or, in case of a change of location, pursuing the educational program started in a remote school. Moreover, the recent availability of enabling technologies, like Internet access, world-wide web, etc., at low cost has given a strong incentive to tele-education. However, most current tele-education platforms are unidirectional: one teaching source is broadcast to several learning sites. By contrast, our approach emphasizes the aspect of *cooperative learning*, with many-to-many communication allowing interaction between students and teacher, or even collaborative teaching by several persons on different sites. This objective requires integration of various technologies for which we have gained expertise through related projects such as the European ACTS project “Distributed Video Production” (DVP) and Swiss SPP project MEDIA. Our research center (CUI) is especially rich with several communication technologies: in particular we have both broadband (ATM) and baseband (Internet) networks and a satellite dish, which we plan to use for communicating with other sites especially in the eastern european countries. These technologies provide means for ensuring two way interactivity, which is a capital issue in the learning process and specially for cooperative learning where human contacts tend to disappear or to be virtualized. Together with specific educational requirements that can be drawn directly from the academic environment, they will allow us to study and set the foundations of what would be best described as “distributed-education”.

Our research is organized into two levels: *Cooperative learning concepts* and *Cooperative learning technological framework.*

Cooperative learning concepts: the concept of distant teaching and learning called tele-education has proven to be of great interest in many fields. Nevertheless current research effort seem to be bound to specific application domains and specific tool requirements. We plan to use a very natural real world "teaching and learning" paradigm to explore the needs of education in a distributed world. Such an approach will help build a background object-oriented framework for cooperative learning, as well as the concept of a "distributed-education shelf" capable of holding educational tools in a similar way as in a warehouse. This approach is partially inspired by a similar framework which was built in our lab for multimedia equipment, and which successfully led to the publication of a book on object-oriented multimedia programming[12].

Cooperative learning technological framework: once the specific requirements for cooperative learning have been set and have been encapsulated in an abstract framework, it will be possible to bind them to the right technology depending on the description of the needed resources and their availability at education time. Therefore the framework can meet current and future technology and can maximize the ratio between available and requested resources. In this approach, for example, a particular course becomes an instance of an abstract class having well defined properties.

One of the targets of a University is the dissemination and transfer of knowledge and technology between countries. This is achieved with the exchange and visits of scientists and students from one university to another. However these exchanges need money and time, something that is not always available to everyone. A cooperative learning environment like the one we are aiming at will reduced the travel of students and professors to a bare minimum and will allow the students to follow courses given in universities that are geographically dispersed within the same day, without having to spend, for example, 2 hours for traveling in order to attend a 2 hours lecture.

A second advantage of our cooperative learning environment, is for students and professors in countries with low budgets and salaries, like for example Eastern Countries. A trip of a professor from, say Moscow to a Western European University might cost the equivalent of six months salary, making effectively scientific traveling prohibitive. Furthermore it is practically impossible for a student in an Eastern country university to visit and spend a semester in a Western European university without financial assistance from a western country. A cooperative learning environment will allow to both professors and students in different countries to attend, participate and even pass the exams for regular lectures and courses given in other universities. As a result the ties between the countries will be strengthen and the exchange of knowledge and technology will be drastically enhanced. Since we have a satellite connection link, and since we have contacts with several East European countries, we are targeting to run our prototype software environment for real courses between Geneva and East European countries.

Finally a cooperative learning system will allow the residents of rural areas to participate in the courses and lectures given in central universities allowing this way better exchange of ideas and acceptance of different cultures within a country.

References

[1] *Jean-Henry Morin and Dimitri Konstantas, Towards Hypermedia Electronic Publishing,* in this volume.

[2] *Dimitri Konstantas, Jean-Henry Morin and Jan Vitek, MEDIA: a Platform for the Commercialization of Electronic Documents*, in this volume.

[3] *Christian Rauber, Peter Tschudin, Serguei Startchik and Thierry Pun, Archival and Retrieval for Large Image Databases: Application to an Historical Watermarks Archive*, in this volume.

[4] *Vassilis Prevelakis, A Framework for the Organisation and Dynamic Reconfiguration of the World Wide Web*, in this volume.

[5] Xavier Pintado, Michel Bonjour and Edelmiro Fuentes, *The WanderVögel project*, in this volume.

[6] *Dimitri Konstantas, Migration of legacy applications to a CORBA platform: a case study*, in this volume.

[7] *Dimitri Konstantas, Marisa Marotta and Chrislain Razafimahefa, The Design and Implementation of a Comparative Verification Suite for OSF/1*, in this volume.

[8] *Jan Vitek, Compact Dispatch Tables for Dynamically Typed Programming Languages*, in this volume.

[9] *Laurent Dami and Didier Vallet, Higher-Order Functional Composition in Visual Form*, in this volume.

[10] *Laurent Dami, Functional Programming with Dynamic Binding,* in this volume.

[11] *Laurent Dami*, T*ype Inference and Subtyping for Higher-Order Generative Communication*, in this volume.

[12] Simon J. Gibbs, Dionysios C. Tsichritzis, *"Multimedia programming, Objects, Environments and Frameworks*", p. 323, ACM Press, New York, Addison-Wesley, 1995

[13] Dimitri Konstantas, "Interoperation of Object Oriented Applications", chapter 3 in *Object Oriented Software Composition,* Editors. O. Nierstrasz & D. Tshichritzis, Prentice Hall 1995

[14] Laurent Dami, *Software Composition: Towards an Integration of Functional and Object-Oriented Approaches*, PhD Thesis, University of Geneva, 1994.

[15] *http://viswiz.gmd.de/DVP/,* DVP project home page

[16] *http://cuiwww.unige.ch/OSG/DVP/wp43.html*, DVP project WP 4.3 Distributed Rehearsal home page.