Challenges of Collaborative and Interoperable Computing – A working version

Lea Kutvonen
University of Helsinki, Department of Computer Science
Lea.Kutvonen@cs.helsinki.fi

Abstract— **Collaborative and interoperable computing is a domain of research where the impact of business or social needs meet the hard computing technology solutions. Collaborative computing refers to interworking of autonomously provided services (e.g., Web Services, Business Applications) and computer support for controlling collaborations between those (e.g., virtual enterprises, electronic business networks). Interoperable computing refers to the variety of challenges placed on the computing and communication platforms, and global computing infrastructure, that supports those application-level services. Capability to collaboration means effective capability of mutual communication of information, proposals and commitments, requests and results. Interoperability technology supports that by technical aspects (transport of messages), semantic aspects (mapping between information representation and messaging sequences), and pragmatic aspects (willingness of partners, awareness of the joint process model).**

This paper outlines the CINCO group (Collaborative and Interoperable Computing) research area, future vision, and topical research challenges.

I. Introduction

This paper outlines the CINCO group (Collaborative and Interoperable Computing) research area, future vision, and topical research challenges. The presentation of the group strategy is completed with a few words on the contributions to the department's curricula.

The CINCO group (Collaborative and Interoperable Computing) is officially part of a larger NODES (Networks and Distributed Environments and Systems) group, launched some ten years ago. The CINCO group is younger, but has roots in the "old school". The group has been earlier known as the ODCE group (Open distributed and collaborative environments), although the individual project names (DRYAD, Pilarcos, web-Pilarcos) have been the ones more in use.

The CINCO group is not only involved with the distributed platforms, but has interests on the software engineering and information systems areas, as discussed in more detail later.

The research leading to the CINCO group was indeed based on a challenge raised by more traditional middleware research at the department [1]. A row of projects was introduced over the following years:

- DRYAD (Directory adventure: traders in open networks) in 1992-1996;
- Pilarcos (Production and integration of large component systems) in 2000-2002;
- web-Pilarcos (Production and integration of systems in web services environments) from 2003 onwards; and
- TuBE (Trust based on evidence) started in 2004.

Joint theme for these projects is the direction towards service oriented environment, where each service is autonomously administered, and used in a loosely-coupled community of peers. Such environment is necessary and beneficial for electronic business networks. The existing middleware solutions are still not sufficient to address these needs: further concepts for managing services and collaborations are needed. Interoperability issues are still under major development and research internationally.

The DRYAD group concentrated on trading service that provides matchmaking of service requests and service offers based on service types and other properties, to support dynamic binding. The Pilarcos project enhanced the set of middleware services for analyzing the interoperability requirements between services in more detail. As a further step, the web-Pilarcos project produced prototypes for B2B middleware services for establishing interoperable eCommunities, monitoring the behaviour of the eCommunities, and facilities for managing changes and breaches of eCommunity contract during the eCommunity lifetime. The TUBE project focuses on the essential trust concepts and trust management facilities required in such environment.

The group working on open distributed and collaborative environments has had a consistent goal on supporting the needs of electronic business networks by high-level middleware services. Current research and development trends elsewhere show that the area of research is still topical had includes major challenges. For example, the group has become involved with the European network of excellence, INTEROP. INTEROP aims to create the conditions of an innovative and competitive research in the domain of Interoperability for Enterprise Applications and Software. The network brings together over 50 partners from the areas of enterprise modeling, ontologies, and architectures. Our participation in the network have concentrated on Architectures and computing platforms. During the first year, the network has produced a number of state of the art surveys, and is now forming a second phase work plan based on those results.

Further work is needed on creating a consistent stack of tools from the strategical business management, through busi-

ness process modeling, to middleware support for managing eCommunities. Areas of specific interest include management of non-functional features in the models and contracts, static and dynamic verification of business network models, and trust relationship concepts and management in eCommunities.

The work is gaining more stable ground, as a number of fresh PhD students have committed themselves to these research goals, and the coursework and seminars leading to this area have found their position in the teaching curricula.

The following sections first position the group at the research arena, and then discuss research challenges raised by a vision of future, global computing infrastructure and federative computing platforms. As the nature of this paper is to declare a group strategy, the research challenges are left fairly open-ended, to capture the general direction of work, not to provide detailed research results or even detailed project or thesis work outlines. Furthermore, we conclude by analyzing the group main assets.[6

II. Future vision

Electronic business networking refers to solutions where business processes and the information systems of business networks depend on modern communication technologies. The present, rapid globalization of business makes enterprises increasingly dependent on their cooperation partners; competition takes place between supply chains and networks of enterprises. The level of dynamic integration capabilities between independent enterprise ICT systems is critical for success of such business networks. Enterprise ICT systems are expected to participate into several, potentially heterogeneous networks simultaneously. They should also be able to react fast to changing partnerships, and use technology-independent tools for managing technical and semantical interoperability. The agility and ability of enterprises to join dynamic business networks is essential for their success. A business network can involve for example a car saler, buyer, bank, insurance company, and a company specialized on car maintenance, with their natural responsibilities, governed by the market area regulations.

We aim for decrease in the cost of establishing and operating new electronic business networks, especially the cost of changes in the business processes, partnerships, application services, and platform technologies. The main investment must therefore be placed on the right kind of middleware that is able to use metainformation on the changeable elements for governing the overall collaborations [2]. The middleware services need to address interoperability at multiple levels, and address dynamic aspects of the network operation including breaches. As the middleware is directed for enterprises participating in multiple, heterogeneous business networks where gradual evolution is to be expected, we call it B2B middleware. It forms a loosely-coupled collaboration layer on top of distributed, service oriented middleware [2]. The general architecture view is illustrated by Figure 1.

The challenges arising from the collaborative business network management are twofold. First, a strategic breeding

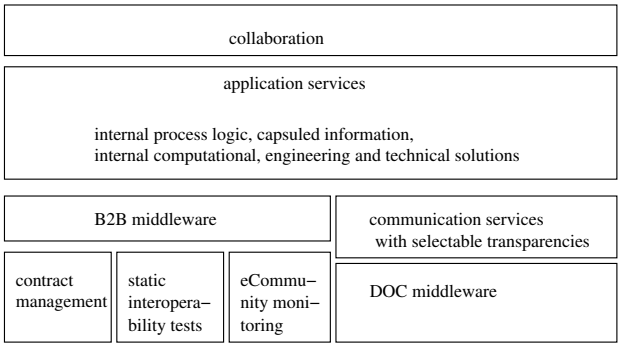

Fig. 1. Architecture overview.

environment for new business networks is needed. That environment incorporates facilities for deciding on the shared business process, and roles of partners within it; selection of component services from the partners' IT systems; ensuring and enforcing interoperability between the component services; and establishing the business network [3]. Second, an operational environment for maintaining and controlling the business network is needed. That environment incorporates facilities for joining and leaving the network; automated monitoring of the behaviour of the network and intelligent methods for adapting to technological changes and heterogeneity in the processing environment; and adapting to collaboration changes in terms of network membership and breach management [4].

For the eCommunity management, interoperability is a fundamental issue. Interoperability, or capability to collaborate, means effective capability of mutual communication of information, proposals and commitments, requests and results. Interoperability covers technical, semantic, and pragmatic interoperability. Technical interoperability means that messages can be transported from one participant to another. Semantic interoperability means that the message content becomes understood in the same way by the senders and the receivers. This may require transformations of information representation or messaging sequences. Finally, the pragmatic interoperability captures the willingness of partners for the actions necessary for the collaboration. The willingness to participate involves both capability of performing a requested action, and policies dictating whether the potential action is preferable for the enterprise to be involved in. In the pragmatic view, process-awareness in terms of collaborative business process model is needed, augmented with nonfunctional aspects, some of which are related to business policies.

The interoperability aspects can be stacked as follows, Figure 2. Especially, the stack involves business rules and enterprise policies at the topmost level, indicating that the business-applications operated should be manageable in this respect, and that there should exists mechanism for managing potentially arising discrepancies. Another point of special interest is the location of failure management in the stack. Some of the failures, i.e. contract breaches, should be resolved by community-level business processes, and therefore, the

supporting middleware should support these situations. In addition, the lack of workflow enactment in the stack is intentional. The business-applications are expected to execute their private (local) business processes independently, only interacting according to a monitored external business process.

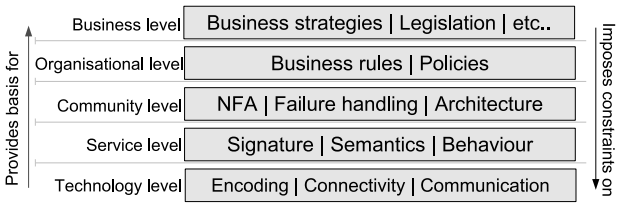

Fig. 2. Layers of interoperability management.

III. Locating the research area

The CINCO group can be positioned between middleware, business integration, and software engineering fields, as shown in Figure 3.

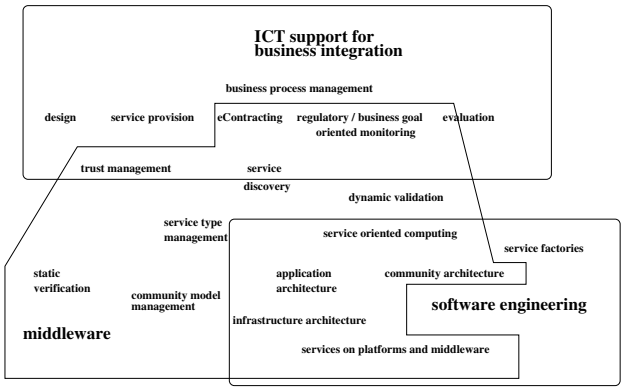

Fig. 3. Positioning on the research field.

The area of middleware addressed by the CINCO group is placed on top of traditional DOC (distributed object computing) middleware platforms, rather addressing the special needs of enterprises forming collaborative eCommunities and maintaining appropriate business processes within them. The essential services include

- business service discovery based on service type and service properties;
- eContracting between enterprises in order to form new collaborative eCommunities (i.e., business networks);
- monitoring of interactions with the eCommunity in terms of agreeable behaviour of partners in terms of legislation and other regulatory systems, and conformance to the contracted behaviour;
- entering into recovery or sanction processes when eContract breaches occur;
- support for the above services such as
  - service type management;
  - eCommunity model management;
  - trust management;
  - verification and validation of service type and eCommunity models with static and dynamic methods.

The NODES group used to be described by a layered structure (see Figure 1) ranging from device drivers in the bottom, through operating systems, up to middleware. Initially, the middleware referred to host management middleware, later to distributed object and component computing middleware. The layers mapped well to different research topics within NODES: the mobility group worked at the lower layers of middleware, while the CINCO group was more interested on the application service centric and autonomy-induced problems.

In Figure 4 various layers of middleware are shown. The CINCO group is involved with common middleware services on the domain of inter-enterprise collaboration. The goal is to provide middleware level services as general infrastructure so that applications can be built without excess consideration of interoperability problems and unsynchronized evolution within collaboration partner IT systems. However, it is impossible to study middleware services without their impact on users and the software and service production processes in which the middleware is used, as was indicated in Figure 3.

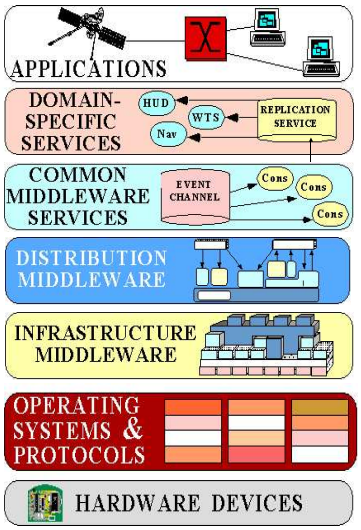

Fig. 4. Types of middleware [5].

For the modern middleware and open service architectures such as Web Services, typical is the separation of access interfaces from the service provision and implementations. The service oriented computing architecture (SOA) has brought out techniques for autonomy of service providers, loose coupling of partners, and separate management of composition from the actual service enactment. The SOA approach is well illustrated for example by Figure 5.

The CINCO group work can be positioned with the SOA as a frame. The service publication and discovery processes are continued to be developed further. The methods of expressing service behaviour, non-functional aspects of services,

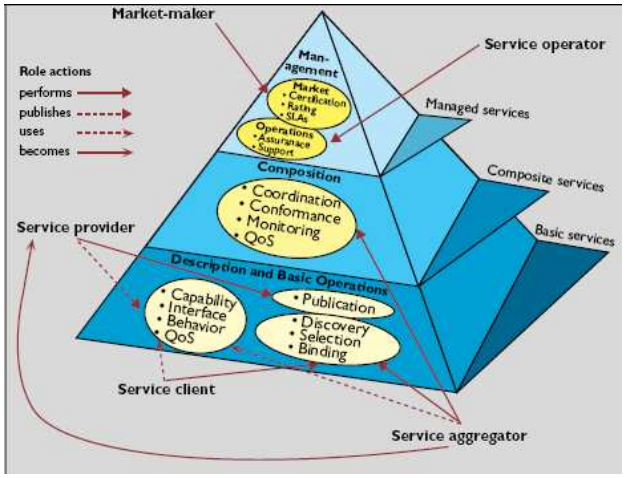

Fig. 5.
Service oriented architecture [6].

trustworthyness of the discovery services and the mediated services themselves, and the reasonable organization of service offers into the global infrastructure are still open challenges. In contrast to the SOA composition layer, the CINCO group concentrates on loosely-coupled, autonomous, and active services that form contract-based, coordinated communities where the control of collaboration and interaction is distributed. In most other approaches, the enactment of the agreed business process across the composed service is in focus, but the federated approach used within CINCO, the initiatives of interactions are done by the services themselves, and the control middleware only monitors and makes corrective actions as needed. For the CINCO group, this layer is mainly focused on eContracting and monitoring. On the SOA management layer, the group is only interested on the requirements and benefits business management create for the middleware and the infrastructure services. The essential requirements include change support and embedded architectural solutions that support gracious evolution of business processes, business-application services, collaborations (in terms of changing partners), and technologies.

The CINCO group is thus involved with the three areas of middleware, software technology, and information systems integration. At the department of Computer Science at University of Helsinki, the software engineering group is involved with modeling, production processes and testing. From these areas the modeling issues are a common field: the CINCO group involves collaboration models, service models, and models of non-functional features of collaboration in static design phases, but also as controlling and changeable metainformation in the operational-time middleware services. Therefore, tools for producing these models, verifying them both with static and runtime techniques, and using model property analysis as a feedback for designers are important. At the area of information systems, the relevant related topics cover modeling of business processes, workflow and enterprise systems, and enterprise system integration and interoperability. Our interests on this area are focused on capturing the strategical and pragmatic aspects of business and business rules into the models maintained.

IV. Research challenges

Above, the main areas of research have been named as middleware and infrastructure service required for management of eCommunities of autonomous business services. Below, these themes are further discussed.

A. Service discovery and service type management

Although much of service discovery research, like UDDI and semantic web extensions, is directed towards user-oriented browsing and discovery of human-usable services, we aim to software-composition oriented type matching for interoperability purposes. In this case, the directive ontology is derived from the three challenges noted in above. First, the business network requirements for a service are to be derived from the role requirements. Second, the implementor needs to express the interface description of the service provided for matching and discovery purposes. This is further detailed below. Finally, the service description itself does not involve the abstract communication layer concerns, but those are addressed separately by a compulsory part in the service offers collected to a service offer repository. The service matching process has to check both parts, but separately: as a result, interceptors can be added for both levels independently [7].

The service descriptions are based on predefined service types. All parties of the common network are allowed to define new service types, so that an evolving type system is created. A service type defines functional and non-functional properties for a class of business services. The functional part of a service type comprises of interface signature (service interface syntax), interface protocol and additionally semantic annotations for exchanged documents (messages). The interface protocol describes the externally visible behaviour of a service in a bilateral conversation. The non-functional properties of a service type describe for example QoS-requirements and policies. When a new service is published to a public service offer repository its behavioural properties and especially its conformance to the claimed service type must be verified. Behavioural descriptions of service types are also needed for static verification of service interoperability and runtime monitoring of conformance between the community contract and actual service behaviours.

Often, only equality or subtyping relationships are considered when service types are matched. Or, as with semantic web, there is an ontology into which services are grouped and matching descriptions can be found based on hierarchical positioning.

Here, we form the relevant service type ontology little by little into the type repository system. The ontology is fairly flat, but wide: as the users of the repository are middleware agents, they cannot adapt to generalizations or specializations of a service type very far, but are more agile in plugging new technology dependent pieces into an already existing

framework. The main purpose of the type repository is to allow checking that service types match together, and to give references to small modules needed in the framework to patch minor technical differences. Therefore, there is an technology independent level of services that is concerned with information exchange relationships, and a more technology oriented level that is concerned on information representation and application level protocols. This is adequate, taken that a further technology dependent layer is separately organized to support these selections.

The relationships of interest for the type repository users are: no match, similar types (equality of text or reference, subtyping), and interoperable with interception. The comparison and judgment is not fully automated and cannot be made (due performance issues) at the time of query. Instead, the service type publication process involves verification of the type, comparison to other named types, and verification of the type relationships. The process of interceptor creation is external to the type repository.

The federated type repository service is an essential element of a B2B middleware that supports establishment of new business networks, or in a more simple case, connection between independently administered clients and servers.

The role of the type repository is to provide a trustworthy source of service type information, and furthermore, provide transformation services for communication between almost similar interfaces [7]. The service types can thus be matched with each other in a more relaxed way, only limited with interoperability requirement. As an enhancement, the cost of connection can be added to direct users to choose ”native” types instead of transformed connections.

The service type matching approach supports evolution of services in a heterogeneous environment, where independent actors create new items, and where market forces has effect on the usability of items, in addition to the verifiable correctness properties. Furthermore, the approach gives a natural tool for managing one type of transformation components needed in the current component-based, model-driven networking environment.

B. Business network models

Internationally there is a plethora of business process and workflow modeling techniques, many of them addressing inter-enterprise challenges too. The languages can be categorized in two groups based on the underlying methodology: a) creation of collaboration models that are used for creating interoperable software elements at each enterprise, and b) addressing the interweaving of independent local business processes to create a collaboration model. Our approach by necessity combines these; we use the originally suggested collaboration model to check that the partners can commit to the shared model, and to monitor their conformance to the model while applying it to operational business. Our focus is indeed on the collaboration, and the intent is to leave local processes as freely modifiable in detail level as possible.

For interoperability purposes the collaboration model must declare the following features.

- Roles for participants, in terms of provided and consumed services, and information exchanged, stored, and processed.
- Criteria for acceptance of a service into a role; criteria for acceptable behaviour from a service in that role.
- Information contents to be exchanged and service behaviour involved.
- Collaborative procedures to use in case of exceptions in the basic service processes, service breaches, or negotiation of dynamic changes in the business network.
- Nonfunctional aspects of the collaboration, covering issues gathered from strategical business viewpoints, process-aware collaboration semantics, and technical implementation of the business transactions.

For dynamic changes during business operation, the collaboration model must leave for negotiation and monitoring the following features.

- Identity and location of the individual service providers.
- Those business policy aspects that are specified in the business network model and service type used in it.
- Mapping of abstract service behaviour and information contents to the technical representation of those, as far as is practical.
- The local, detailed business processes that contribute to the collaborative business process as individual service steps.
- Nonfunctional aspects.

Figure 6 illustrates the separation of modeling and runtime management issues of business network management. The operational environment is much governed via business network contracts; an important contents element in these contracts are the business network models in use (so that they can also be changed), and results of negotiated properties of the collaboration (such as partnership, so that renegotiation and interoperability monitoring can occur).

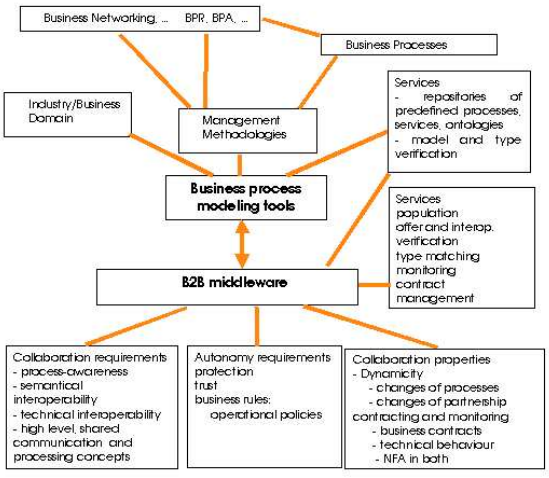

Fig. 6. Aspects reflected by B2B middleware and business process models.

It is essential that the models are statically verified as part of the model publication process.

C. Trust management

For inter-enterprise collaborations trust information is essential as business collaborations cannot be established or maintained without explicit decisions on trust.

Trust and reputation are complex, multifaceted notions, with a range of term definition in the present litterature. From the definitions, we can sieve out an agreement that trust is a quantified belief that the trustee has named qualities. Based on that belief, trust decision can be taken, leading to a variety of indications. Into the trust decision, a number of other factors have impact, i.e. the context in which the decision is taken. The belief can be situational/implicit or be computed from a number of explicit trust information elements. There is also a strong requirement for trust to be a dynamic concept, accumulate past experience, and thus provide adaptation to new situations in the networked environment. The trust decisions can be either local or centralized, and various metrics can be used to express the quantity or quality of trust. Between concepts of trust and reputation, the separation can be drawn that trust is prone to be subjective and lead to decisions to act or accept access to services, while reputation is more objective information used as element in forming trust.

For the CINCO group the natural approach to trust management is to emphasis autonomous trust domains making local trust decisions on a) entering eCommunities and b) accepting participation to each interaction within the eCommunity. Trust information is partially collected by private experience on the trustees, their behaviour, and the providers behind these services; partially by reputation information passed across the network of identity-and-repudiation-management domains.

The interesting aspects of trust management include

- eContracting processes where trust negotiation is an elementary part;
- creation of experience information based on monitoring of the business services, and publishing that experience information in a reasonable form to the reputation management network;
- metrics and methods for initiating and accumulating reputation values;
- understanding the effect of context (system, business, eContract, regulatory, business process induced risks and benefits) in which trust decisions are taken;
- role of intrusion detection methods in the gathering of experience information; and
- effects of reputation information formation to the enterprise potential for successful business and beneficial collaborations.

In addition, an essential theme to study is the new risks created on business and technology levels by the introduction of services taking over trust decisions with routine nature.

D. eContracting

Inter-enterprise collaborations with peer-to-peer relationships among independently administered business services are best controlled by contracts. The term eContracting is considered to capture (at least)

- negotiation of contract terms a) at business level, and b) at technical level; this includes forming a shared understanding of the joint (external) business process and the roles of each partner within it;
- monitoring of the interactions as the community works;
- evaluation of the success of the business processes run.

The overall architecture adopted in the CINCO group separates these elements into a eCommunity breeding environment (selection of partners so that they have matching view of the joint process, negotiation of policies and communication techniques, setup of contract information at each involved partner, setup of each involved service) and operational environment (progress and breach monitoring, protocols maintaining the distributed contract between peers). The key element formed by the breeding environment and used as a governing element in the operational environment is the eContract. The eContract captures a range of aspects from business point of view down to communication engineering. The various aspects in the eContract reflect viewpoints of the ODP-RM (Open Distributed Processing Reference Model) [8], [9], forming a relationship between computational and business services [10], [11].

The challenges for eContracting include

- management of contract templates and ontologies for expressing contract information in an interoperable way;
- transformations from business process models to eContract structures; the models are often designed based on process flows which the contract has to be structured by administrative domains, which causes a need for two-way transformations between the internal middleware presentation and the external view of processes for example for process and monitoring rule designers, business process evaluation needs;
- agents and protocols for managing the distributed contracts in an efficient way;
- reflection mechanism for promoting system changes to the contractual state information and contract changes to the local service management services;
- alternative negotiation protocols for contract establishment and partner selection, taking into consideration trust, cost, and expected benefit of the communities;
- analysis methods for the validity of eContracts.

E. Collaboration breach detection and management

Breaches against eContract governance is always a question of observability. In the business process models used, some degree of monitoring rules can be embedded, thus defining the degree of observation. Especially, the models are annotated with task boundaries, expressing the granularity of state information distribution in the eCommunity.

Challenges at this area

- verification of the properties (cost of monitoring, roll-back properties of the model) of the annotated business network model;
- constructing B2B middleware protocol and local service management facilities that allow suspending a business process for the duration of the breach recovery process and potentially returning to the original process for continuation;
- modeling of general purpose recovery processes and associated business services;

V. Teaching curricula

The teaching curricula addressing the CINCO group needs include a master level course on Middleware technologies, and on Inter-enterprise collaboration schemes.

The Middleware course addresses the concept and role of middleware, discusses DOC level solutions (e.g., CORBA, J2EE), and goes further to SOA. The course is given yearly or every two years, and will adopt a textbook in spring 2006 [**?**].

The course on Inter-enterprise collaboration schemes is new in 2005. It addresses integrated, unified, and federated architectures for inter-enterprise collaborations, eContracting in terms of forming virtual organizations, service discovery methods, aspects of business process management, and various interoperability techniques.

Interoperability is also addressed at various levels at MSc theses, candidate theses, and seminar work. Various themes of SOA, MDA, semantic interoperability, and business process management are discussed in a series of seminars chaining up new themes every semester.

Collaboration with other subgroups of NODES is also needed, in terms of formal methods for analysis of models and behaviour. Further, software engineering themes need to be clearly understood and joint work areas need to be further discussed.

Finally, at the department level or at the new graduate schools themes on research skills development should be brought up more.

VI. Assets

The main assets for the group:

- Committed people. The core group in the CINCO group includes PhD students Sini Ruohomaa, Toni Ruokolainen, and Janne Metso. PhD students Jyrki Haajanen and Lea Viljanen work in industry, Pirjo Moen and Timo Karvi, both PhDs, are joining the group.
- Combination of practical aspects with theoretical foundation.
- Ability to form more focused projects in future: The topics currently in progress are agent protocols for the operational time environment and eContract contents management, type management functionality, and trust management. Next national projects will continue with the business network modeling facilities, trust infrastructure, and verification of the models and types.
- Prototype software that provides a testbed for further work. It brings together individual research topics into a meaningful whole, and makes evaluation of suggested solutions practical.
- Industry interest to the kind of research, and commitment to follow through as funding partners.
- Continuously developing curricula. However, we need more teachers, for example for putting together practical environment of software projects using Web Services and SOA tools.
- Networking at various levels: NODES, DIMES, Finnish Computer Society is pushing up SIG groups for Software engineering and Information Systems, InterOp Noe, young researcher networks at national level (activities for improving that) and internationally (enjoying that).

VII. Discussion and actions

- Discussed the relationship to actual projects; group sees the relationship between past work, new projects, and private studies.
- New project plans not detailed here; intended to have longer use time.
- Represents ”no-surprises” state for the group.
- Major growth to be expected on the group size.
- Terminology in the paper is not consistent throughout, and some of the illustrations are not in sync with the text as they were old versions. Fix.
- Text is commonly understood to give outer limits to the ”value system” behind work; not everything is to be solved here! Text excludes some popular trends on solutions, and shows another selected path. Listings of challenges is felt to give freedom not-to-worry about the other points in addition to the one selected.
- Privacy-related work discussed and agreed on. Revise accordingly / check whether affects.
- Verification themes under discussion. Address in text?
- Remove chapter on teaching?
- Add references to external publications for recommendation for new recruits
- Remove this list.

References

[1] K. Raatikainen, “The First Ten Years of the NODES Group,” University of Helsinki, Department of Computer Science, Tech. Rep., B-2005-1.

[2] L. Kutvonen, “Automated management of inter-organisational applications.” in *International conference on Enterprise Distributed Object Computing (EDOC2002)*, 2002, pp. 27–38.

[3] L. Kutvonen, T. Ruokolainen, J. Metso, and J. Haataja, “Interoperability middleware for federated enterprise applications in web-Pilarcos,” in *INTEROP-ESA’05*, Feb. 2005. [Online]. Available: http://interop-esa05.unige.ch/INTEROP/Proceedings/Interop-ESAScientific%/PerPaper/I06-2%20424.pdf

[4] L. Kutvonen, J. Metso, and T. Ruokolainen, “Inter-enterprise collaboration management in dynamic business networks,” in *International Conference on Cooperative Information Systems(COOPIS’05)*, 2005.

[5] D. C. Schmidt, “Middleware for real-time and embedded systems,” *Communications of ACM*, vol. 45, no. 6, pp. 43–48, June 2002.

[6] M. P. Papazoglou and D. Georgakopoulos, “Introduction,” *Commun. ACM*, vol. 46, no. 10, pp. 24–28, 2003.

[7] L. Kutvonen, "Sovereign systems and dynamic federations." in *The Second International IFIP Working Conference on Distributed Applications and Interoperable Systems (DAIS99)*, 1999, pp. 77–90.

[8] *Information Technology – Open Systems Interconnection, Data Management and Open Distributed Processing. Reference Model of Open Distributed Processing. Part 2: Foundations*, ISO/IEC JTC1, 1996, iS10746-2.

[9] *Information Technology – Open Systems Interconnection, Data Management and Open Distributed Processing. Reference Model of Open Distributed Processing. Part 3: Architecture*, ISO/IEC JTC1, 1996, iS10746-3.

[10] L. Kutvonen, "Challenges for ODP-based infrastructure for managing dynamic B2B networks," in *Workshop on ODP for Enterprise Computing (WODPEC 2004)*, A. Vallecillo, P. Linington, and B. Wood, Eds., 2004, pp. 57–64. [Online]. Available: http://www.lcc.uma.es/~av/wodpec2004/papers/7-kutvonen.pdf

[11] L. Kutvonen and J. Metso, "Services, contracts, policies and eCommunities – Relationship to ODP framework," in *Workshop on ODP for Enterprise Computing (WODPEC 2005)*, 2005.