51 DISTRIBUTED INFORMATION SERVICES SUPPORTING COLLABORATIVE NETWORK MANAGEMENT

W. Mulder,
LogicaCMG Netherlands, University of Amsterdam, wico.mulder@logicacmg.com
G. R. Meijer,
University of Amsterdam, LogicaCMG NETHERLANDS, geleyn@science.uva.nl

Information retrieval is one of the key processes in the management of collaborative networks (CNs). For large and complex networks manual processes need to be supported by automated retrieval systems.
We study the characteristics of information retrieval systems and discuss their applicability in distributed, dynamical, heterogeneous environments, such as CNs. We developed a model of a distributed information retrieval system, and implemented a prototype of distributed information services in the domain of collaborative organizations.

1. INTRODUCTION

A collaborative network (CN) is a network consisting of autonomous, geographically distributed, heterogeneous entities (e.g. organizations and people) that collaborate to better achieve common or compatible goals [1].

Typical examples of CNs are Virtual Enterprises (VEs), Virtual Organizations (VOs), Virtual Labs (VLs) and Professional Virtual Communities (PVCs).

While today the concepts of CNs are well known [8], aspects around their management are still subject to research. Karvonen et. all [5] focus on VOs and define VO-management as the organization, allocation and co-ordination of resources and their activities as well as their inter-organizational dependencies to achieve the objectives of the VO within the required time, cost and quality frame. We can broaden this definition applying it to CNs. CN management applies knowledge, skills and tools in order to achieve the CN goals.

One of the key elements of CN management is the acquisition of information about the status and operation of the network and its entities. In case the number of entities is large, the selection and retrieval of information needs to be supported by an information system. Such a system might relieve a CN-manager from the unnecessary complexity of specifying, searching and obtaining the information that is relevant for getting overviews and taking decisions.

We emphasize that many companies (amongst others SAP, IBM, HP, Microsoft, Tibco, Cisco and Oracle) have developed monitoring and control tools that support organizational management. Their solutions have evolved from monolithic systems towards centralized, homogenous network systems. However, most of these solutions are focused on processes and procedures within single organizational boundaries. Applying them in the area of CNs is not a trivial task due to the concepts and characteristics of a CN. A significant, generic characteristic of a CN is that its entities are self-contained and work together only for a particular amount of time. Other characteristics are geographical spread and local influences.

As an example we take a collaborative network of organizations (CNO), which is a heterogeneous, distributed network of organizations in which each organization (member) acts within its own local environment to deliver a product or service to the network. Together they achieve a common goal, often in the form of a service to the outside world. The members can physically be located around the globe, introducing globalization management issues [4]. The issues go further than locality problems like date-time, units of measures or language; Culture [10] has significant influences on their operation and contributions to the CNO. The location and social habits of the country where a member is situated, combined with the fact that each member has to deal with particular local rules, laws and legislations has significant influences on its operation in the network. Furthermore, each member has its own internal procedures and mechanisms in order to produce its results or fulfill its services. This has impact on the frequency, amount and kind of information that a member is able to provide. Information that touches the member's core-business cannot or will not always be shared. The latter has to do with strategic goals and aspects of thrust.

Another characteristic of a CNO is called 'local control'. A CNO operates in a so-called shared, collaborative environment. Such an environment is characterized by shared processes which are controlled by multiple local domains. For a CNO this means that each participating member has the control of its own part of the environment, while it shares its services with other members in collaborative processes. The aspects of local control and local influences are typical for CNs in general: It holds for networks of organizations where business services are shared, as well as for networks of technical components such as large, shared computer infrastructures (collaborative infrastructures, grids).

The research in our group focuses on information retrieval in dynamic collaborative environments. We study the mechanisms of data acquisition and architectures for information gathering in such environments. We use the concept of Information Services, defined as entities that operate in large, distributed networks that provide information about that network and its components. As part of ongoing research in our group we study the use of information services to support the maintenance of grid environments [3].
Information services can also be applied to crisis management in the field of logistics and transport[6]. This holds preventive measurements such as early warning systems in order avoid accidents, as well as support for handling emergency situations when a traffic accident has occurred. In [7] we describe the concepts of adaptive information services that help mobile ICT teams, called Squads, in doing their work. It is here where we are working in the field of hybrid networks, studying man-machine interfaces and collaborative multi-agent learning.

In this paper we focus on the information retrieval in CNOs. We work out a model of a distributed information retrieval mechanism and discuss a prototype that supports the information retrieval in CNOs by means of Distributed Information Services.

2. DISTRIBUTED INFORMATION SERVICES

Distributed Information Services (DIS), are piece of software in a network environment that provide information about the state or usage of one or more components of that network environment.

This section gives an overview of the properties that characterize a DIS. We look in particular at those properties that are interesting with respect to the environments of CNOs

Local scope and control

A DIS is designed to operate in a particular part of the network, known as the local domain. In case of a CNO, a local domain is considered to be an environment that belongs to a particular member-organization. A DIS retrieves information from such a particular member environment.

Located inside that domain, a DIS has access to local systems and can be configured and controlled locally by system operators belonging to the member organization.

A particular member domain may use its own language, or specific technical interfaces to particular subsystems or components. Having access, a DIS fetches data from a particular hardware or software component. A DIS is designed to generalizations and translations between the local domain and the rest of the network.

Local scope and control also enables a DIS to react on events that are not allowed to be visible outside the borders of a particular organization. An example of this is the measurement of intermediate production progress details; although a DIS can be configured to inform a local manager about these figures, it may shield the details towards the CNO-manager.

Autonomous behavior

A DIS is an active component that fulfills measurement and provisioning tasks autonomously by means of rules. An example of such a rule is to measure each hour production details, and compare this with previous values and report when necessary. Rules may also define actions in case a certain pattern in the measured data is found. This enables local problem solving by reacting on events that are only visible inside the borders of a particular local environment. The definition and configuration of the rules can be given via one of its communication interfaces.

Communication interfaces

The communicative properties of a DIS can be divided into three groups:

- Communication with a controller outside the local environment
- Communication with its local environment
- Communication with other DISes.

The first kind refers to configuration and control by a human or a system that consumes the provided information. This can be an individual who configures the rules about what and how often a DIS should measure. Such an individual might be a CNO manager who is interested in information about the operational status of the

CNO, or a system operator (squad team) who is interested in the operational status of the supporting ICT infrastructure.

The second kind of communication allows the local environment to interact with the DIS and vice versa. A DIS can respond to events broadcasted by components inside a local domain. A DIS can retrieve properties of the environment enabling operation in heterogeneous, dynamic networks.

The third kind of communication is the communication with other DISes. The possibility to exchange information between multiple DISses enables collaborative information provisioning and notifications of events that happen in other domains. Another reason for inter-DIS communication is the ability to react immediately to local events, without the involvement of central, hierarchical (broker) components. Communication between DISes involves the exchange of information over multiple local domains, security aspects have to be taken into account. Therefore, authentication rules describing what is allowed to communicate must be defined and shared between the various DISes.

For all three kinds of communication, protocols need to be defined. These protocols are not specified by the model, because they depend on the implementation, goals and purposes of the application area.

Adaptive behavior

Combining the two previous characteristics, i.e. having autonomous behavior and communication interfaces, a DIS is able to learn and show adaptive behavior. It may use its communication facilities in combination with its internal state and rules, to change its way of measuring data.

In general, adapting does mean ‘changing behavior’. In case behavior is based on a set of rules, defined by means of parameters and metrics, changes can be provoked by altering a rule or one of its parameters. In case adaptation leads to improvement of results, we call it ‘learning’. The ability to learn strengthens the autonomous capabilities; a DIS may e.g. take the initiative measure pro-actively, without intervention of its controller.

Collaborative goals

Communication between each other enables DISes to compete or collaborate with each other in order to provide better information. Combined with the aspect of adaptation, a DIS is designed to adopt mechanisms of collaborative learning.

Summarizing, our model of a DIS is has a set of properties that enable it to operate in dynamic, heterogeneous environments. These characteristics are reflected by a set of interacting DIS-components, shown in Figure 1.

The model allows an extensible and flexible configuration of these architectural parts. Still, the basic functionality is to provide information (output) based on configurations given by the controller (input) and autonomous data fetching.

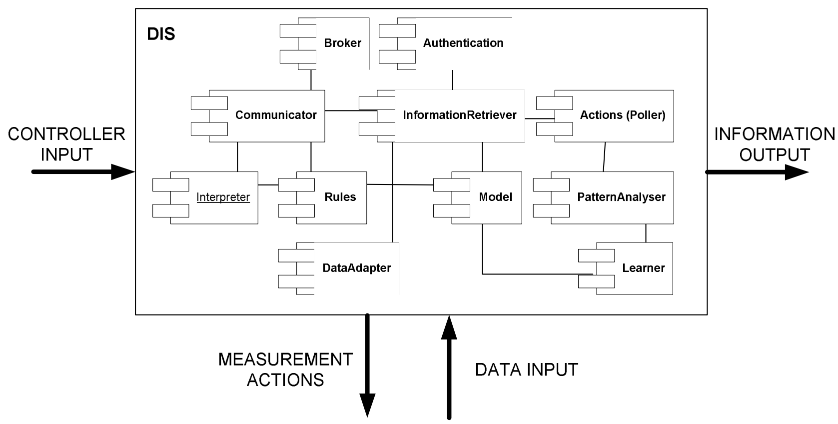

Figure 1 - architectural design of a Distributed Information Service

3. RESULTS

We apply our model in domain of CNOs. As part of our work in the Ecolead project[1] we work on a set of generic e-services supporting the management of collaborative organizations.

Here, management information is characterized by indicators like ‘delivery due date’, ‘product quality’ and other quantitative measurements. Also information about the collaboration between the partners and the planned activities, are indicators that are used to measure the performance of the network. In general, using the concept of Service Level Agreements, an indicator can represent any type of information that has to be fetched from each member location.

The prototype is named DI^3 (pronounced dee-triple-eye) which stands for Distributed Indicator Information Integrator. The current version supports a basic information retrieval mechanism, and basic communication facilities.

Figure 2 shows the experimental setup of the prototype. DI^3 consists of a set of collaborating components. At the location of each member-organization, an ‘Information Retriever’-component measures the value of a pre-configured information indicator at a certain, pre-defined frequency.

At each heartbeat, a fetching request is send to an intermediate component, called a DataAdapter. This component fetches by means of a query on a local member system or, in case a member does not have a system that can be queried, it

[1] www.ecolead.org

may fetch the value by means of an email or browser screen, requesting a human individual to enter the value manually.

The Information Retriever stores the value returned by the DataAdapter in its local memory, and verifies whether it should inform its controller, i.e. the CN-manager. Such a briefing will be sent by means of a third component, the Broker having merely an intermediate function.

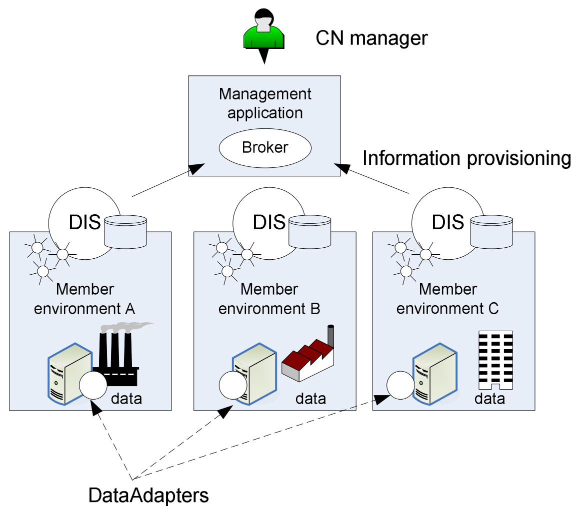

Figure 2, experimental setup of the prototype

In the current version (Figure 2) we deal with three different kinds of information:

- Performance information by means of progress indicators.
  An example is the measurement of the amount of the work that is done. In case the measured value is below a certain level, the Information Retriever launches an event message.
- Contract management information, such as 'due date' and 'expected production figures'. An event is launched in case it is suddenly changed.
- Collaboration information. For example a satisfaction indicator. This is measured by means of a list with satisfaction-figures filled in weekly by the members indicating the level of collaboration with other partners. Events may be launched in case of negative marks.

Currently, we are testing the feasibility of retrieving this kind of information. Furthermore, we use the prototype to work on the aspects of collaborative information provisioning, by studying implementation of ad-hoc communication algorithms in combination with learning algorithms.

The implementation of the prototype is based on Web-Services, and developed in Java. The technology of Web-services is common for implementations on a Service Oriented Architecture. It also uses the advantages of platform independency and message based communication.

5. DISCUSSION AND FUTURE WORK

The characteristics mentioned in section 2 show similarities with aspects of (at least) three other interesting fields of research. Our goal is to adopt results algorithms and technology achieved in these fields. Figure 3 shows the overlap of three fields of research. The central area reflects the domain of DIS.

A DIS has characteristics of Peer-2-Peer (P2P) systems [9]. The properties of P2P networks enable a high scalability and the ability to deal with heterogeneous systems. Information sharing is one of the main design goals of P2P systems. P2P communication mechanisms allow individual nodes to share data and work in a collaborative way.

Autonomous behavior is widely used in the field of Agent Technology. Communication and autonomous adaptive are used the area of Multi Agent Systems (MAS). An overview of learning in multi agent systems is given by Weiss [11].

A third field that has a significant influence on our design and model of DIS is Machine Learning. This field studies learning algorithms and mechanisms to model complex environments.

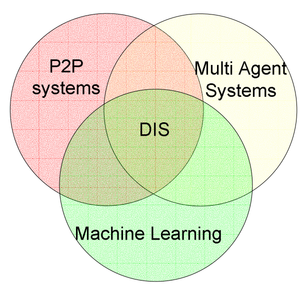

Figure 3, relationships with three overlapping fields of research

The implementation of the prototype is based on a modular approach; It is extensible and allows for replacements of individual parts our model.

We want to use the prototype also in other fields, such as emergency rescue teams and computer network maintenance. In general, distributed information services might be applicable in any kind of heterogeneous distributed networked environments, were situations are characterized by unstructured, unpredictable ways of information sharing.

Our prototype is used to study an adaptive frequency mechanism. We use the degree of deviation from expected values to decide about the need of increased frequency, in order to foster closer monitoring when needed. It is planned to add and study more adaptation capabilities to the information services. We plan to work on collaborative learning algorithms in order to improve the information provisioning in large, complex, dynamic networks. We want to include aspects of Epidemic information dissemination [2] and combine this with machine learning and theories of reinforcement learning.

6. CONCLUSION

In this paper we described our research on information retrieval mechanisms in collaborative networks. As being one of the applied areas, we discussed the aspects collaborative networks of organizations (CNOs) and the necessary management support. We described a model of a distributed information retrieval mechanism and worked out the characteristics of distributed information services. We also described a prototype that is being developed in the domain of CNOs.

7. REFERENCES

1. Camarinha-Matos, L.M., Afsarmanesh, H, Collaborative networks: A new scientific discipline, J. Intelligent Manufacturing, vol. 16, Nº 4-5, pp439-452, ISSN: 0956-5515 , 2005.2.
2. Eugster P., Guerraoui R, Kermarrec A-M., and Massoulie L., Epidemic Information Dissemination in Distributed Systems, IEEE Computer, 37, 60--67, 2004.
3. Fitzgerald, S. 2001. Grid Information Services for Distributed Resource Sharing. In Proceedings of the 10th IEEE international Symposium on High Performance Distributed Computing (Hpdc-10'01) (August 07 - 09, 2001). HPDC. IEEE Computer Society, Washington, DC, 181.
4. Freeman, Craig. 1998. "The Influence of Culture on Corporate Structure." Macquarie University, School of Economics and Financial Studies, New South Wales, Australia.
5. Karvonen, I.; Salkari, I.; Ollus, M.; Characterizing Virtual Organization and Their Management. In: PRO-VE'05. (Sep. 26-28. 2005: Valencia, Spain). Proceedings. Ed. Springer, ISBN 0-387-28259-9, p. 193-204.
6. Konijn, M., Crises Management, using eCall principles in the Netherlands, Research proposal ICIS, internal report.
7. Mulder W., Meijer G.R., SQUADS: Software Development and Maintenance on the Grid by means of Mobile Virtual Organizations using Adaptive Information Services, Virtual Enterprise and Collaborative Networks, Proceedings PROVE04, WCC 2004, ISBN 1-4020-8138-3, p. 379-386
8. Proceedings PROVE04, WCC 2004, ISBN 1-4020-8138-3, Proceedings PROVE05, Springer, ISBN 0-387-28259-9
9. Ripeanu, M., A. Iamnitchi, and I. Foster. Mapping the Gnutella network: 'Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design' IEEE Internet Computing, 6(1):50--57, Aug. 2002
10. Schein, E.H., Organizational Culture and Leadership, 3[nd]ed., Jossey-Bass,2004, ISBN 1-55542-487-2
11. Sen and Weiss G. (1999). Learning in Multiagent Systems. In Gerhard Weiss (ed.) Multiagent systems, A modern approach to Distributed Arti cial Intelligence, 259-298, Cambridge MA: MIT Press.