

Peer-to-Peer Technology – An Enabler for Command and Control Information Systems in a Network Based Defence?

Tommy Gagnes, Karsten Bråthen, Bjørn Jervell Hansen, Ole Martin Mevassvik, and Kjell Rose

FFI (Norwegian Defence Research Establishment)

P.O. Box 25, 2027 Kjeller, Norway

Phone: +47 63 80 70 00

E-mail:

{tommy.gagnes | karsten.brathen | bjorn-jervell.hansen | ole-martin.mevassvik | kjell.rose} @ffi.no

Abstract

In a network based defence, the military forces that are superior in terms of information gathering and sharing also are envisioned to be more likely to outperform an opponent. Providing military decision makers with a common operational picture enables them to have a better and shared understanding of the situation when making critical decisions. This increases their chances of making right and timely decisions. In this paper, we argue that using peer-to-peer technology, with its inherent mechanisms for ad hoc networking, can be a good start at solving the problem of distributing a shared situation picture in network based defence C2IS. We illustrate our view by presenting basic requirements for a system that utilizes peer-to-peer technology to dynamically distribute situation information between actors. An architecture is outlined as well as some thoughts on its realization.

1. Introduction

The Norwegian Defence has decided to move towards a network based defence [1], adapted from the US' concept of Network Centric Warfare [2], [3]. This new paradigm is based on information superiority, meaning that the forces that are superior in terms of information gathering, processing and sharing also are more likely to win a conflict. A network based defence enables the decision-making process to be more rapid than with previous operational concepts.

The ability to dynamically restructure itself and to be able to share information efficiently will be crucial properties of a force. The ability to make the right decisions in a timely manner requires a shared understanding of the situation between all members of a force, despite the great variations in communication equipment used. Thus, the same information should be available to a soldier using his PDA with low bandwidth as well as to his military commander provided with more capable equipment. This should also be the case in a highly dynamic situation, something that represents a challenging distributed computing problem.

A project at FFI (Norwegian Defence Research Establishment) is working in the areas of architecture, middleware, data fusion and psychology to help building better decision-support systems for military commanders in the future network based defence.

In this paper, we argue that peer-to-peer technology, with its inherent solutions for ad hoc networking, is a promising technology for solving the problem of providing all

actors with a common operational picture (COP). A COP enables military commanders at all levels to have better information when making critical decisions.

Our view is presented by identifying basic requirements for a system that dynamically distributes situation information between actors. We also outline an architecture that is designed to enable discovery and information exchange based on semantic languages that have been designed for the next generation World Wide Web [9].

We also discuss how to realize this architecture and show how the use of peer-to-peer technology, in this case JXTA [8], can provide much of the flexibility required in such a decision-support system. A lookup system based on the Resource Description Framework (RDF, [12]) is introduced.

The realization will be part of a demonstrator developed by FFI for experimentation with decision-support systems [19]. The demonstrator simulates a real user environment for decision-making and includes e.g. low-bandwidth radio links and satellite links.

Theoretical work on peer-to-peer technology in a military setting has been presented in [4] and [5]. Projects using peer-to-peer technology in the defence domain have been described in [6] and [7], whereas work on combining RDF with JXTA for a library system has been presented in the Edutella project [16]. In our project, we focus on ad hoc discovery of system components in a military setting. We do this by means of semantic languages in general, and thus do not necessarily limit ourselves to using RDF, even though our initial work is based on RDF.

2. Network Based Defence C2IS

Today's command and control systems are highly centralized, with central processing of information and message exchange between sites. Connections are static, and have to be administrated by people. Such an architecture has several weaknesses. For instance, it is not responsive enough to future users' needs for ad hoc networking. To enable each actor in the future battlespace to obtain customized and detailed information about the current situation, a more dynamic and decentralized system is needed.

In the Norwegian Defence's concept for network based defence [1], a high-level component model has been adopted. The model identifies the classes of components illustrated in Figure 1, all relying on a common information infrastructure, or infostructure. This infostructure is envisioned to enable discovery of components as well as communication between them, providing a pluggable grid type of architecture.




Figure 1. Network based defence component classes

In the future battlespace, military resources such as vessels, vehicles and even soldiers, will be logically decomposed and represented in the infostructure as instances of the component classes shown in Figure 1. Component classes are Sensors, Effectors and Decision components. These components should be able to continuously discover the services available to them and to publish their capabilities in an ad hoc manner. This means that some sort of lookup or directory service will be needed in the infostructure.

The variations in bandwidth, latency and type of communications that the infostructure will consist of are huge, ranging from high frequency radio links to satellite links and LANs. Communication types could be anything from unreliable messaging to real-time streaming.

Integrating all this is a challenging task. The diversity mentioned above, in addition to the fact that legacy systems also must be accessible in such an architecture, means that one cannot rely on a single middleware technology for communication between components in the infostructure. The existence of different operating systems and machine platforms also implies that a system supporting heterogeneity is needed in order to implement the infostructure wanted.

We believe that standardization of component descriptions and use of standardized interchange protocols will be needed to take care of interoperability. Ideally, this standardization should provide semantic descriptions of components, meaning that infostructure processes could be much more automated than they are today. This would give different nations the ability to efficiently map descriptions to their own architectures, meaning that e.g. coalition partner infrastructures can be made interoperable with each other.

2. Why Peer-to-Peer Technology?

So, how can peer-to-peer technology help us enable the concept of a network based defence? We will present several reasons why we consider this category of technologies an attractive solution to many of the problems we are facing when experimenting with some of the functionality required in the infostructure shown in Figure 1.

The idea behind peer-to-peer systems is to treat each participant in the network, each peer, equally. As has been stated in much of the literature, the idea of peer-to-peer systems is to exploit the resources of the endpoints at the edge of the network. Information should be able to flow freely between all peers participating in the peer-to-peer network.

This concept becomes very powerful, but also challenging, when applied in a military setting. The reason why there is a huge challenge involved is that the communication endpoints in a military setting are very different, both in terms of their needs and their capabilities. A single application could support both real-time streaming as well as instant messaging and terminals can vary from PDAs to workstations and servers. Also, military systems must be secure, and they have to take care of emission control.

2.1 Less Administration

In a network based defence, components may appear and disappear, something that would impose serious demands on system administrators if traditional technologies were to be used.

To avoid having system administrators working continuously to adapt the system to the changing environment, the process of networking components in the infostructure

should be as automatic as possible. The infostructure should require a minimum of administration when the network topology changes and has to provide a high degree of availability even during partial failure and problems like denial-of-service attacks.

One of the most appealing properties of peer-to-peer technology is that it may have the potential to solve just these problems. Its ability to provide survivability and redundancy means that peer-to-peer systems can provide services to their users, even with long-lasting partial failures, due to their redundant architecture.

With traditional client-server architectures, one must often be aware of network addresses, references and so on that ideally should be allowed to change without affecting the system. We would not want a name- or lookup-server failure to bring the whole infostructure down. Most peer-to-peer technologies have solutions to this problem, by enabling system designers to instruct the system at design-time to find replacement services dynamically during run-time. Such fail-over mechanisms are crucial in a network based defence.

2.2 Homogeneity

By using peer-to-peer technology, one can create a virtual network, that is, a network consisting of both different communication technologies and platforms that still acts as a whole. Some peer-to-peer technologies even give endpoints the opportunity to dynamically switch bearer technology, meaning that, in our case, if a satellite link fails, the endpoint automatically could switch to a radio connection completely transparent to the user.

2.3 Automatic Discovery

Ideally, the infostructure should be able to perform automatic bootstrapping without needing a priori information about its environment. Although most peer-to-peer technologies still need at least some initial configuration, or seeding, they are designed to deal with automatic and dynamic discovery.

It is reasonable to believe that the infostructure will consist of many legacy systems using more traditional middleware, or even proprietary means of communication. Even so, peer-to-peer technology can function as a universal lookup service, giving system components a unified way to discover services, and thereafter utilize them by means of different middleware technologies.

The lookup service should be flexible, providing components with the means to negotiate before making use of a component's service. Not many peer-to-peer systems support the advanced mechanisms to solve this, but many of them lay down a good foundation to build slightly more advanced lookup functionality. As we will present in section 5, we are experimenting with peer-to-peer technology to do "semantic lookup", based on some of the technologies envisioned to constitute The Semantic Web [9], [13]. Building a lookup method based on standardized languages could provide an efficient solution to interoperability with other systems, e.g. the C2IS of our coalition partners.

2.4 Information Handling and Resource Exploitation

A network based defence, as it is envisioned, will make a vast amount of information available to decision makers. The task of distributing differentiated information to different users will be a challenging task, with the risk of flooding the infostructure with information. The ability to customize information exchange and to search for information among a distributed set of providers will be essential. To ease the burden on

the resources at the edges (sensors) of the network, query results should be propagated to the consumers interested in them. This would provide sharing of aggregated information. Peer-to-peer systems are designed with many of these objectives in mind, possibly providing an infrastructure for realizing the functionality we want.

2.5 Possible Applications

Almost all kinds of today's civilian peer-to-peer applications can be envisioned used in the future battlespace. One could use collaborative applications like virtual whiteboards, integrated multimedia like video streaming, instant messaging, various information/content-sharing applications (e.g. for planning) as well as clustered computing (e.g. for intelligence number crunching). The decision maker should have all the needed applications available to her integrated in a single graphical user interface, yet another challenge.

2.6 Potential Problems

There are a few potential pitfalls to peer-to-peer technology that we are aware of. Security is a quite obvious one when working with military systems. Decentralized systems like peer-to-peer technologies represent a challenge, since many security solutions build on hierarchy and centralized solutions. Another potential problem can be bandwidth consumption, as pointed out in [4]. However, this can be viewed in two ways. The total bandwidth used by a peer-to-peer system may be higher, but all communication lines can be used to balance the load on the network, preventing potential bottlenecks that client-server architectures may lead to. An interesting research area will also be mechanisms for providing Quality of Service (QoS) in peer-to-peer systems.

3. Distributed Situation Picture Production

One of the most important applications of the infostructure is the distributed production of a situation picture. In addition to producing such a picture, such an application must also be designed to be able to distribute the situation picture to a variety of terminals, ranging from PDAs to powerful servers with various display types.

At FFI, we have started working on the challenges explained above. A concept for distributed picture production has been designed, and is now in the process of being concretised. Central in our proposed architecture is the Picture Production Node (PPN), which is an agent, in the sense that its behaviour is autonomous and proactive, and that it communicates with other PPNs. A PPN serves situation picture users who subscribe to information that it gathers, either directly from various sensors or from other PPNs. This is illustrated in Figure 2, which is just a simple example. The value provided by peer-to-peer technology in such an architecture should be quite obvious, e.g. it could provide the discovery mechanisms needed to dynamically adapt to change.




Figure 2. Concept for distributed situation picture production

4. Architecture Outline

The internals of a PPN is illustrated in Figure 3. A PPN must be able to discover and communicate with sensors, other PPNs and situation picture users. Modules that collect and forward data should be implemented, as well as modules that handle users and their subscriptions. Data and information fusion, including conflict handling is essential to heighten the quality of data and to reduce the total bandwidth usage between PPNs. Ideally, an abstraction layer should be inserted to make a clean separation between the middleware technology and the functionality of the system. This means that we can change the underlying technology as peer-to-peer middleware evolves.

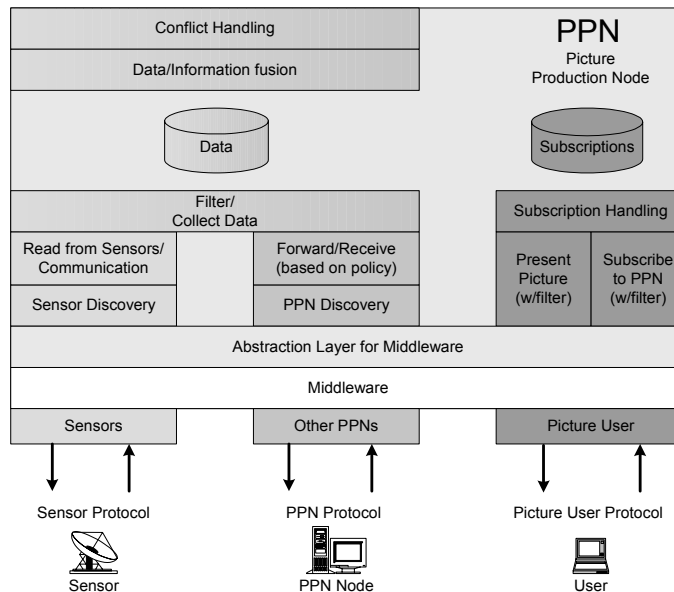


Figure 3. The inside of a PPN

To create a powerful universal discovery mechanism, we are developing a basic taxonomy for the components that should be discoverable. The taxonomy will become part of a shared data model, or ontology, providing all components with a common

understanding of how information is represented in the infostructure. The ontology can be used to represent metadata for discovery requests or for single information queries.

To ease the process of discovering components of the right class, the basic taxonomy is split in a tree structure, which is thereafter extended in an inheritance hierarchy.

This is useful for a number of reasons. First, it provides an efficient way to limit searches. For instance, one can search for “all sensors that are a subclass of camera”, which could return even types of cameras that one did not know about initially. Each category should be described by a number of properties to refine queries and results. It should, for example, be possible to limit a search by a geographical position or area.

Second, such a taxonomy is extendable at all levels, something that will ease the process of adding new component classes to the discovery system.

Third, a protocol hierarchy that follows the basic taxonomy could be defined. Figure 4 shows how interaction protocols could be more specialized as the sensor description becomes more specific, collecting the general interaction patterns at higher levels while adding the more specific patterns at a lower level. If this architecture should prove to be feasible, the protocols on a high level should be subject to standardisation, possibly also between coalition partners.

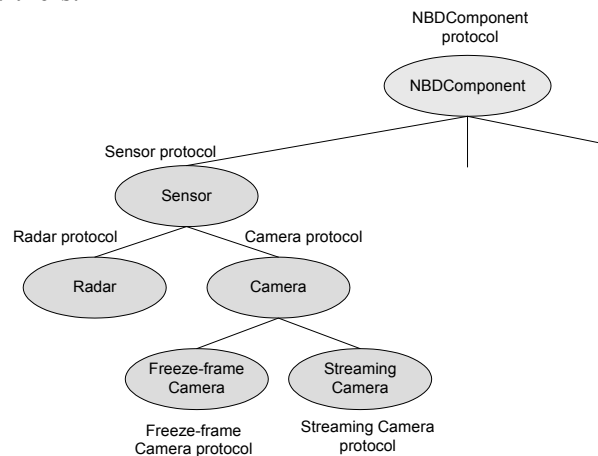


Figure 4. Partial component taxonomy

A taxonomy like the one just introduced should be described by means of a standardized language that allows for heterogeneous systems. This will make the coupling of the discovery system together with other systems easier, independent of the bearer middleware. Such coupling of heterogeneous systems is a key factor in defence information systems, since it enables interoperability with coalition partners.

5. Demonstrator

In the process of realizing the architecture outlined, we focus on experimenting with new technologies and the development of a demonstrator.

Our initial focus has been on the discovery process, and the networking of components. This functionality is of crucial importance to an infostructure.

We have decided to build our demonstrator on the JXTA [8] peer-to-peer platform. This is because we found JXTA to be the closest to what we needed in our system in terms of reach (Network Address Translation, firewalls) and heterogeneity (JXTA is based on XML). A number of interesting projects are running in the JXTA Community,

for instance the Edutella project [16], and different projects that are experimenting with streaming content over a JXTA network.

Compared to other peer-to-peer technologies, like e.g. Jini [17], with its lookup based on Java interfaces, JXTA's way to define services is much more flexible. This opts for some kind of standardisation to be made. We have chosen to describe our components by means of the XML-based Resource Definition Framework [15], [12] which provides mechanisms to describe all resources that can be identified on the Internet. By using RDF, we specify that the Camera class is a subclass of the Sensor class, and use this relationship during discovery. This provides a Jini-like polymorphic lookup mechanism, and is a good way to limit a search as well as to make a system extendable. Figure 5 shows an example of a very simple RDF graph that represents the Camera/Sensor subclass relationship presented above. Some properties are shown as well, and this is very important to refine queries and results.

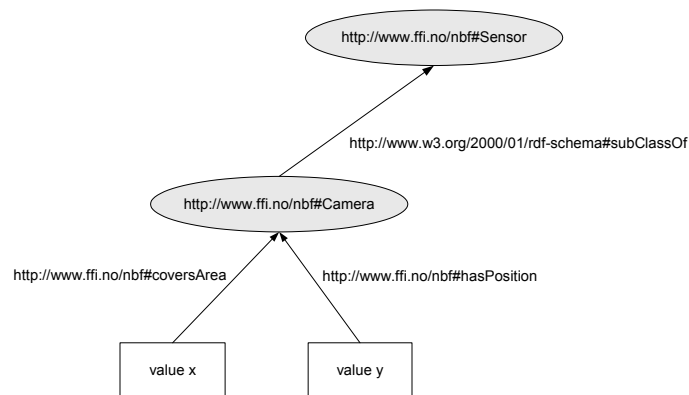


Figure 5. Use of RDF to describe a subclass relationship

As mentioned, RDF is a building block in The Semantic Web [9], with languages like DARPA Agent Markup Language (DAML, [10]) and the Web Ontology Language (OWL, [13]) built on top of it. The use of ontologies is likely to be an essential enabler for a common understanding of the semantics of the entities in the infostructure, providing further, more intelligent automation of ad hoc discovery. In our demonstrator, we experiment with the technologies mentioned above. Tool support for these technologies has not yet come as far as we would have liked, but we are currently using HP Labs' Jena semantic web toolkit [14] to execute RDF queries. This is done by means of the RDF Data Query Language (RDQL) provided in the toolkit. The RDF Query is carried by the JXTA Resolver service, which uses a protocol that broadcasts messages in a JXTA group. This means that we can group e.g. all sensors in a group, and use the Resolver service to query all listening sensors. Only the ones that match the query return an answer, thereby reducing bandwidth usage. The answer may include further information needed to contact the service. This could include a description of which protocol to use when interacting with the component, and a reference to the peer that answers the query.

A technology like the Web Services Description Language [18] or in the future, the DAML ontology for web services, OWL-S [11], could be used to describe the interaction with a service. This would provide for integration with systems based on the Web Services specifications, like Open Grid Services Architecture (OGSA, [20]).

6. Conclusion

Building the computer systems of the future battlespace demands solutions to a variety of distributed computing problems. There is, of course, no such thing as a perfect technology for all kinds of distributed systems. However, the potential ability of peer-to-peer systems to solve at least some of the problems involved, compared to today's more centralized systems, is the reason why we have decided to base our experiments on this category of technologies. We have identified several requirements that a network based defence imposes on the decision-support and information systems, many of which are not fulfilled by today's architecture. Peer-to-peer technology seems like a promising technology for solving many of the challenges identified. Further work in this direction is needed, however, to confirm the initial evaluation.

References

- [1] Headquarter Defence Command Norway, *Concept for Network based defence*, 2002. (Forsvarets overkommando, *Forsvarssjefens militærfaglige utredning 2003 – Konsept for nettverksbasert anvendelse av militærmakt*, 2002)
- [2] Alberts, D. S., Garstka, J. J., Stein, F. P., *Network Centric Warfare, Developing and Leveraging Information Superiority*, 2nd Edition (Revised), CCRP Publication Series, DoD Command and Control Research Program, Washington, 1999.
- [3] Alberts, D. S., Garstka, J. J., Hayes, R. E., Signori, D. A., *Understanding Information Age Warfare*, CCRP Publication Series, DoD Command and Control Research Program, Washington, 2001.
- [4] Bontrager, M. D., *Peering Into The Future: Peer-to-Peer Technology as a Model for Distributed Joint Battlespace Intelligence Dissemination and Operational Tasking*, School of Advanced Airpower Studies, Air University, Maxwell Air Force Base, Alabama, June 2001.
- [5] S. Manoski, "Eliminating the Middleman: Peer-to-Peer Technology for Command and Control", *The Edge, Distributed Computing issue*, vol. 6, number 2, 2002.
- [6] Sun Microsystems Inc., "Sun and Raytheon Create Open, Adaptive, Self-Healing Architecture for DD 21, U.S. Navy's New Class of Destroyers"; <http://www.sun.com/software/jini/news/jini-raytheon.pdf>
- [7] Sun Microsystems Inc., "Ericsson Microwave Systems Teaming up on Network-Centric Warfare Initiatives"; <http://www.sun.com/service/sunps/success/ericsson.pdf>
- [8] L. Gong, "JXTA: A Network Programming Environment", *IEEE Internet Computing*, May/June 2001, pp. 88-95.
- [9] T. Berners-Lee *et al.*, "The Semantic Web", *Scientific American*, May 2001, pp. 34-43.

- [10] Defence Advanced Research Projects Agency (DARPA), Darpa Agent Markup Language (DAML); <http://www.daml.org>
- [11] The Web Ontology Language Service Ontology Specification (OWL-S); <http://www.daml.org/services>
- [12] World Wide Web Consortium, Resource Description Framework, RDF; <http://www.w3.org/RDF/>
- [13] World Wide Web Consortium, W3C Semantic Web Activity; <http://www.w3.org/2001/sw/>
- [14] B. McBride, "Jena: a semantic Web toolkit", *IEEE Internet Computing*, Volume: 6 Issue: 6, Nov/Dec 2002, pp. 55 –59.
- [15] S. Decker P. Mitra, S. Melnik, "Framework for the semantic Web: an RDF tutorial", *IEEE Internet Computing*, Volume: 4 Issue: 6, Nov/Dec 2000, pp. 68 –73.
- [16] W. Nejdl *et al.*, "EDUTELLA: A P2P Networking Infrastructure Based on RDF", *WWW2002*, Honolulu, Hawaii, USA, May 7–11, 2002.
- [17] Sun Microsystems Inc., "Jini Network Technology Specifications"; <http://www.sun.com/software/jini/specs/index.html>
- [18] F. Curbera *et al.*, "Unraveling the Web services web: an introduction to SOAP, WSDL, and UDDI", *IEEE Internet Computing*, Volume: 6 Issue: 2 , Mar/Apr 2002, pp. 86 –93.
- [19] B. J. Hansen, O. M. Mevassvik, K. Bråthen, "A Demonstrator for Command and Control Information Systems Technology Experimentation", *8th International Command and Control Research and Technology Symposium*, Washington DC, USA, June 17-19, 2003.
- [20] D. Talia, "The Open Grid Services Architecture: where the grid meets the Web", *IEEE Internet Computing*, Volume: 6 Issue: 6, Nov/Dec 2002, pp. 67 –71.