

Topic: Network Centric Applications
DISTRIBUTED COMPUTING AND COLLABORATION FRAMEWORK
(DCCF)

Authors¹:

Mrs. Cheryl D. Putnam
Dr. LorRaine Duffy

SPAWAR System Center, San Diego

COMMANDING OFFICER
SPAWARSYSCEN-San Diego
53560 HULL STREET
SAN DIEGO CA 92152-5001

Code: 24122
Comm: (619 or DSN) 553-4002
FAX: 619-553-0854
e-mail: cputnam@spawar.navy.mil

Code: 2411
Comm: (619 or DSN) 553-9222
FAX: 619-553-6405
e-mail: lduffy@spawar.navy.mil

*Sponsoring Agency: Office of Naval Research (ONR), 311

Abstract

Keywords

Peer-to-peer architecture;
Low bandwidth-high latency network conditions;
SSC SD;
Network-aware applications;
Distributed collaboration;
Variable bandwidth management;
Network-centric warfare

Problems and Issues Addressed

The Distributed Computing and Collaboration Framework has been developed by SSC SD (a Naval research and development facility), under the sponsorship of the Office of Naval Research. The focus of the effort is to provide superior performance in “low bandwidth-high latency network conditions,” a hallmark of Naval operational environments. The software development includes a variety of mechanisms to manage and reduce latencies, with emphasis on ensuring effective, real-time collaboration at the operational user level and efficient use of limited bandwidth channels at the network level.

This paper will describe an initial software development package of the peer-to-peer architecture design requirements, exploiting a transport layer (TCP/IP) protocol. This package will support information distribution and search across multiple enclaves, using peer-to-peer data models, as well as models of redundancy for robustness across networks composed of a mix of wired and wireless channels.

The software was tested in a ship-to-shore, ship-to-ship and ship-to-sub paradigm to verify that the initial design met documented Navy bandwidth standards. Results of the analysis showed this initial software to be within standards constraints. We will provide a detailed presentation of our Collaboration Evaluation Laboratory (and methodology) that we have used to manage our development and transition efforts.

Distributed Computing and Collaboration Framework (DCCF)

1.0 Problems and Issues Addressed

1.1 Navy /Joint Operational Deficiency

Future naval operations will necessarily be distributed, complex [Alberts and Czerwinski, 1997] and theater-driven as defined by Naval warfighting concepts such as Network Centric Warfare (NCW) [Alberts, et al., 2001; Alberts, et al., 1999]. Missions will incorporate traditional areas of concern (hostile regimes, precision engagement, terrorism, drug trafficking; power projection), as well as some relatively new mission objectives (e.g., international organized crime, ethnic tensions/regional instability, and weapons of mass destruction in the hands of little-known political groups). Mission planning and execution under these conditions is at best unpredictable, complex, and requiring a faster response from distributed planning cells than traditionally expected. With the variables that affect mission success becoming increasingly diverse and copious, battle planning will require reliable, distributed computing and support for real-time interaction (in order to share appropriate information) at the component level. This implies command centers on ships, on land, undersea, in the air, and possibly in space. Sharing relevant operational information across these diverse “platforms” will necessitate new models of interaction, a new form of theater-level command and control planning. The focus of this project is on distributed command and control (C2) teams coordinating and collaborating, often in real time, to quickly accomplish mission objectives.

If synchronous interactive collaborative planning is not undertaken, then problems/vulnerabilities with plan constraints, consistency of plan understanding, and multiple-services-coordination are often not discovered until too late. The traditional planning model also forces several replanning cycles, which in turn, forces the planning cycle into a sequential series of steps (planning, *then* replanning), precluding the ability of planners to dynamically conduct their component level planning in parallel. Parallel *interactive* or collaborative planning will reduce the time to manufacture and assess the plans by allowing the immediate discovery and repair of potential coordination failure points, one of the most vulnerable and untested elements of multi-echelon planning. This in turn will reduce the number of times replanning must take place.

This concept implies connectivity among primary planning partners. If partners come from different computing platforms, then we must accommodate their computing hardware and software. There is a need to develop revolutionary concepts in *multi-server computing* that imply heterogeneous operating platforms and applications. There is also a need to provide a concept of operations that provides a vision as to how these systems will work. Included is a need to reduce the manning expectations for network administrative personnel aboard our battle groups. This will be accomplished by building a Distributed Computing and Collaboration Framework (DCCF) that can automate many of the tasks and services currently forced upon the operational user in his/her attempt to coordinate and collaborate synchronously, as well as asynchronously, with tactical and operational partners. Given the planning scenario of the near future, the Navy will

require a computing framework that will support distributed planners acting synchronously and asynchronously, taking advantage of *smart* network operations under a variety of network conditions [Shirky, 2004]. This implies the era of the “network-aware application.”

The DCCF project focus is on accommodating the disadvantaged user, traditionally ignored by the commercial-off-the-shelf (COTS) collaboration software community, who continue to assume that users have a LAN (local area network)-like network connectivity among users. The commercial industry has consistently ignored the issues of high network latency and reduced bandwidth that pervades collaborative activities (a critical factor in satellite and other wireless-dependent environments, such as ship-to-shore connectivity). And finally, the Navy must deal with specialized security and firewall-imposed constraints that define battlefield operations. Under extreme battle conditions, COTS software has limited, if any, abilities. Our framework will obviate this operational deficiency by improving network robustness.

2.0 Relevance to Command and Control

This vision of collaborative operations implies command centers on ships, on land, undersea, in the air, and possibly in space. The focus of this project is on distributed command and control (C2) teams coordinating and collaborating, often in real time, more often asynchronously, to quickly accomplish mission objectives.

The primary technical objective of DCCF is the development and demonstration of a computing and collaboration support framework that will provide automated and efficient object synchronization among peer-to-peer and federated computing enclaves (e.g., on-board ships, as well as on-shore; ship to sub to shore), in order to extend collaborative services to hostile and extreme environments. Command and control processes cannot fall victim to weak links and inefficient protocols.

3.0 Approach

3.1 Project Objective

Given these requirements, the DCCF project is developing a computing framework that provides distributed warfighters the ability to use (any) collaborative application in real-world military network environments, capable of accommodating legacy (client-server) and advanced collaboration computing models (such as peer-to-peer application architectures¹). We will provide the ability to “automatically” connect “simulated” land-based, ship-based, and submarine-based command cells in a WAN (wide area network) environment, as the planning personnel conduct a collaborative planning session, using

¹ Joukhadar, Kristina. February 19, 2001. “A human face for collaborative apps.” Information Week.com. Currently there are four major types of peer-to-peer network applications. They include: (1) distributed computing (the goal being to harness latent CPU power,) as in the [SETI@home](#) program; (2) real-time person to person collaboration (as in Groove.net and AOL instant messenger); (3) the supply chain coordination systems (which are typified by workflow/automated alerting applications), and (4) advanced search and file sharing systems (such as Napster, Aimster, and Gnutella.)

available collaborative tools. Our research attempts to define the constraints and limitations of client-server architectures, federated server architectures, as well as peer-to-peer computing models. The Distributed Computing and Collaboration Framework (DCCF) will allow operational personnel to connect “automatically” to needed participants (regardless of physical location), without having to rely on network information (such as IP addresses). Data will need to be automatically tailored for the bandwidth/latency-challenged ship participant and the time-constrained submarine-participant. (The submarine participant’s opportunity for a collaborative planning session is often less than 20 minutes, as they surface for burst transmission.) The collaborative session manager will be tailored for the land-based command center, which will only need a subset of the information transferring between two operational components. *The focus is on the data synchronization and transport layer, not on the operational interface.* That interface or graphical user interface (GUI) will be notionally provided by available collaboration or mission planning tools. This project has also developed a unique evaluation methodology for measuring performance gains in the various operationally-based computing scenarios to ensure “first-time-out” operational usefulness.

3.2 Technical Objective/Expected Payoff

The primary technical objective of this project is the development and demonstration of a computing framework that will provide easy session management and synchronization among distributed planners in a collaboration context across a heterogeneous set of computing enclaves (e.g., ships and shore units). A local server (enclave) within each enclave will be used to coordinate this synchronization. A secondary objective is to *automate* server synchronization/ transport network services, so that the operational user can conduct collaborative planning activities in a seamless fashion, regardless of application or the heterogeneity of the software architectures across various platforms. In other words, the operational planning team elements (e.g. N3) can physically reside aboard ship, on land, or undersea, but plan collaboratively using collaborative software tools, without having to manually monitor connectivity and communication channels in order to distribute information objects among the battlegroup, joint forces and/or coalition forces.

DCCF is articulating and prototyping the software mechanisms (and the underlying network architecture) that will provide acceptable and automated “quality of service” for a distributed collaborative session using collaborative software tools [Oram 2001]. The framework will provide an intermediate layer for exchanging data and information, which will greatly simplify future integration of information assurance technology. Again, the intent is to apply this framework to a distributed-server naval setting that requires that we support ship to shore to submarine collaborative planning activities among members of N3 and J3 elements.

The notional architecture in Figure 1 provides a summary of the work to accomplish our technical objective.

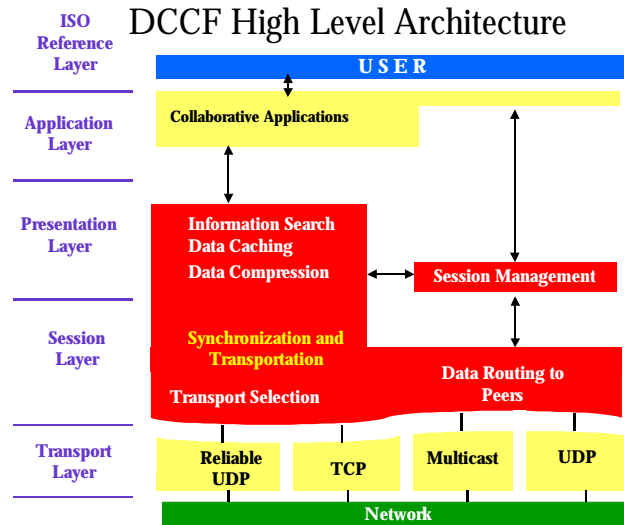


Figure 1: DCCF Notional Architecture

3.3 *Technical Approach*

The development effort for DCCF focuses on three major thrusts. The first two are the two major segments of the architecture culminating in a computing prototype, while the third task is focused on performance assessment and warfighter utility evaluation during continued prototype evolution.

3.3.1 *Session Management Across Heterogeneous Systems*

The Session Management automatically coordinates setup and maintenance of collaborative sessions across a distributed set of enclaves. This includes managing multiple collaborative applications and multiple services across distributed enclaves. This primarily includes software mechanisms for defining a session and user registration/logon/logoff into a virtual environment session, as well as into individual collaborative sessions for each application [dit UPM, 1998]. It includes discovery of available users and their personal computing configuration, in order to determine how an “operator” can connect and collaborate with any other given operational user. A continuing emphasis will be on operator control of quality of service, in order to provide a computing environment that does not require on-site technical support for distributed war-fighting personnel.

3.3.2 *Synchronization Management and Intelligent Transport Management in a Collaborative Planning Context*

To support the development of Synchronization Management and Intelligent Transport Management, DCCF is modeling and comparing various algorithms for consolidating and disseminating the data which needs to be exchanged in a collaborative session using methods which can be used across a wide range of interface architectures (e.g., CORBA, Java RMI, JINI, Message-oriented middleware, etc.) used by potential collaborative

applications. This includes addressing the key issues of dynamically scaling services, application information interchanges and information in a complex computing environment, such as a coalition environment; and providing synchronization under adverse, hostile, and intermittent network conditions, for example: from ship to shore to undersea warfare components.

This framework is being built as middleware to allow existing and future programs the ability to use the various features of DCCF. The added benefit of DCCF as middleware to the applications, is that the framework will allow the users to not only be connected automatically through the proper protocol in a peer to peer architecture, but, the application can allow manual use of the Framework's tools. The user could be allowed to

- Check on the status of the network, which may include the amount of bandwidth available
- Status of the other peers/enclaves
- Determine manually their peer priority for object exchange
- Check on the approximate time for object exchange.

Our goal is to develop an algorithm that will insert, via an applet, the appropriate object transference method, given the needs of an operational user. For example, if a user wants to send a message to a thousand recipients, the middleware will automatically insert a multicast transfer method. If a user wants to conduct a limited secure text-based session with another user, then a different transport method will be inserted (perhaps CORBA or reliable UDP). The point is to give the user greater flexibility, therefore greater robustness in collaborative sessions. This also underscores our desire to remain application-agnostic; to create a source of network-condition information to be exploited by collaborative application in order to improve performance under extreme computing conditions.

3.3.3 Distributed Computing and Collaboration Framework Evaluation

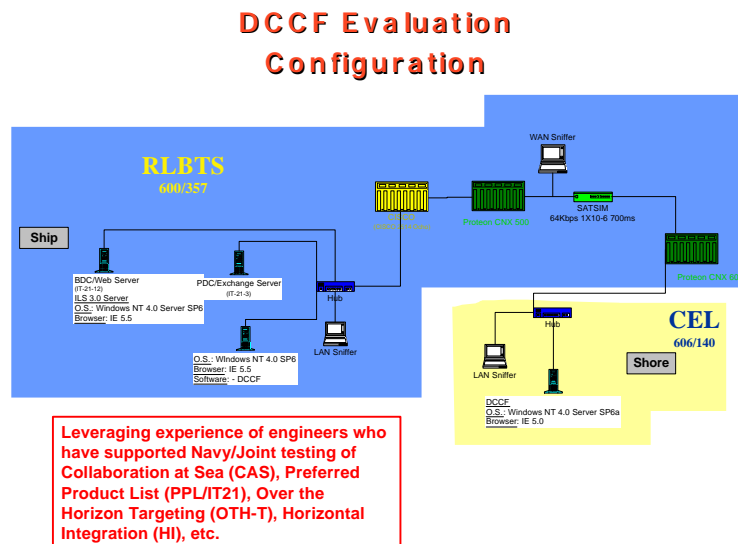


Figure 2. DCCF Evaluation Laboratory Configuration.

DCCF Evaluation includes the design and instantiation of an evaluation laboratory environment that closely replicates the environment of real-world, ship-board/sub network/application testing. We are conducting an initial performance assessment with an internally-derived methodology for documenting performance baselines and enhanced capability during the evolution of prototype development. This initiative required the development of evaluation criteria from a warfighter-user perspective, documenting the increasing utility of provided automated services inherent in the prototype DCCF. The intent is to replicate the SPAWARSYSCEN SD RLBTS (Reconfigurable Land-Based Test Site) test environment; that conducts compatibility testing for on-board ship installs. The planned configuration is depicted above in Figure 2. Evaluation criteria of DCCF under realistic conditions is imperative to the success/usefulness to the warfighter and C2 operations. Evaluation criteria being used for verifying the DCCF software was obtained from the Advanced Digital Network System (ADNS) program office as shown in Figure 3 and is being used to setup/analyze DCCF in both current and future Navy standards.

Evaluation Criteria for BG

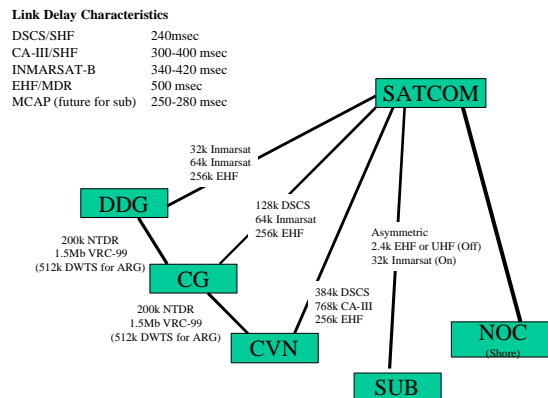


Figure 3. DCCF Evaluation Battle Group Evaluation Criteria

4.0 Technical Issues

Although there are a number of technical issues that must be addressed, the major ones fall into five topic areas. First, centralized server architectures are prevalent in current collaboration applications, but they cannot easily handle:

- High latency (bandwidth not always the issue),
- Intermittent Connectivity,
- Scalability, or
- Interoperability.

These can be (to a limited degree) addressed in the juncture of the presentation and transport layers of the ISO (International Organization for Standardization) model. It is at this point that we conduct our research and development activity, essentially

developing network-smart mechanisms for collaborative applications to use. Second, as applications move to a peer-to-peer and ubiquitous functionality, the computing framework needs to accommodate new methods of “connecting warfighters.” This is a fairly new area of research and brings with it a varied list of technical constraints. Third, security needs to be included from the outset, and we intend to “code” in mechanisms that deal with current methods of handling security among distributed users. Fourth, network-induced delays and application failures plague users in intense interactive collaborative sessions and must be dealt with in as automated manner as possible for extremely disadvantaged users. And fifth, intelligent server synchronization and pre-caching of objects should allow applications to intelligently cope with intermittent network performance, which is the focus of much of our in-house performance evaluation.

In summary, the technical issues that will guide our research and development activities include:

- No current multi-server *collaborative* computing framework exists; current generation is application specific.
- Our computing framework work initially in a collaborative *planning* context, but must also be extensible to any command and control collaboration, coordination, and communication session.
- Distributed collaborative computing is currently focused on staff planning; current generation is not focused on tactical warfare environment, echelon below CJTF (Commander Joint Tactical Forces).
- Need to investigate various algorithms for dynamically scaling services, application information interchanges and information in a heterogeneous environment;
- Need to build a computing framework that can accommodate hostile and intermittent network conditions
- No known computing framework focused on extensibility, so that legacy and competitive collaborative systems can also be used within the same framework
- Need to research efficient WAN inter-enclave communications alternatives, including IP multicast and messaging based technology, for the DCCF synchronization services, including the capability to tailor delivery to disadvantaged users.
- Few (or very poor) methods for evaluating the effectiveness of collaboration-focused computing frameworks, particularly from a scalability and extensibility standpoint.

5.0 Results

DCCF project staff have focused on researching and documenting current client-server, federated server, and peer to peer processing methods, as well as a variety of algorithms, languages, security, and most importantly, Navy requirements, that impact collaborative sessions [Moore, 2001; Patrizio, 2000; Henroid, 2001; JXTRA website; Oram, 2001; Rand, 2002; Ritter, 2001; Sullivan and Vizard, 2001; SyncML].

DCCF v1.0 was first demonstrated in FY01. In FY02 this was an initial software development package, which was an in-house demonstration of the enclave design of the peer-to-peer architecture requirements, exploiting a transport layer (TCP/IP) protocol. The software was tested in a ship-to-shore, ship-to-ship and ship-to-sub paradigm with data analysis conducted to verify that the DCCF initial design was acceptable to documented Navy bandwidth standards. Results of the analysis showed this initial software to be within standards constraints. During FY02, DCCF has been building the (Reliable) RUDP protocol for the transport layer and developing the multi-enclave design along with supporting Navy security requirements.

To summarize the results of the In-House Demonstration and Data Analysis of DCCF Software accomplished thus far:

- The first and second year software development concluded with a lab demonstration of the initial software build, testing and analyzing ship-to-shore, ship-to-ship and ship-to-sub configurations.
- The demonstration utilized an enclave design exploiting a TCP and RUDP transport layer within the RLBTSLaboratory.
- Confirmed the theory of information exchange (not just packets) over simulated “real world” Navy networks.
- Subsequent evaluation and analysis provided metrics supporting real-world Navy information exchange.
- Source Code / All Software and Software Documentation including Description/Installation Guide for DCCF Services
- Requirements Document
- Architecture Diagram
- FY01 Report of Analysis and Lessons Learned
- Programmer Guide

References

[Alberts, et al., 2001] David S. Alberts, John J. Garstka, Richard E. Hayes, and David A. Signori . *Understanding Information Age Warfare*. CCRP Publication Series, Washington, D.C., 2001

[Alberts, et al., 1999] David S. Alberts, John J. Garstka, Frederick P. Stein. *Network Centric Warfare: Developing and Leveraging Information Superiority*, 2nd Edition (Revised), CCRP Publication Series, Washington, D.C., 1999

[Alberts and Czerwinski, 1997] David S. Alberts and Thomas J. Czerwinski (Eds) *Complexity, Global Politics, and National Security*. National Defense University Press, Washington, D.C., 1997

[dit UPM, 1998] *Distributed Session Management*, URL <http://greco.dit.upm.es/~david/papers/lev98-workshop/sessions/slides/sld014.htm>

[JXTRA website] URL www.jxta.org

[Henroid, 2001] Henroid, Andrew *The Peer-to-Peer Trusted Library*, URL <http://sourceforge.net/projects/ptptl>

[Moore, 2001] Moore, Cathleen *P-to-P Aims for the Spotlight at O'Reilly Show*, URL <http://www.infoworld.com/articles/hn/xml/01/02/19/010219hnptp.xml>

[Oram, 2001] Oram, Andy *Peer-To-Peer: Harnessing the Power of Distributive Technologies*, 1st Edition, Oreilly & Associates, USA, March 2001.

[Patrizio, 2000] Patrizio, Andy *New Life For Peer-To-Peer Computing*, URL <http://www.informationweek.com/813/prpeer2peer.htm>

[Rand, 2002], *Report on Peer-to-Peer*. Technical Report, Rand

[Ritter, 2001] Ritter, Jordan *Why Gnutella Can't Scale. No, Really*, URL <http://www.darkridge.com/~jpr5/doc/gnutella.html>

[Shirky, 2000] Shirky, Clay *What is Peer to Peer and What Isn't* , URL <http://www.openp2p.com/pub/a/p2p/2000/11/24/shirky1-whatisp2p.html>

[Sullivan and Vizard, 2001] Sullivan, Tom and Michael Vizard. *Peer-to peer Boosts Web Services*, URL <http://www.infoworld.com/articles/hn/xml/01/02/12/010212hnp2p.xml>