

A Limited Objective Experiment on Wireless Peer-To-Peer Collaborative Networking

Dr. Alex Bordetsky
LCDR Glenn R. Cook

Dr. Bill Kemple
LCDR Timothy Thate

Naval Postgraduate School
Department of Information Sciences
589 Dyer Road, Room 200A
Monterey, CA 93943
abordets@nps.navy.mil

Abstract

The implications of using peer-to-peer communications within an urban environment are significant for Military applications. From a networking perspective, the use of wireless technologies to support the collaboration may have impacts on bandwidth and spectrum utilization. This paper explores the effects of wireless peer-to-peer (P2P) network behavior on the performance of collaboration support applications. The results achieved during the limited objective experiment conducted by the Naval Postgraduate School demonstrate significant affects of roaming on application sharing performance and integration with client-server applications. We discuss the wireless network operation challenges leading to the solutions for scaling up application sharing and improving collaborators self-organizing behavior.

1. Introduction

Communication within the collaborative network environment takes on many different modalities, including e-mail, chat, voice-over-IP and peer-to-peer (P2P). Each of these forms of collaborative communication has a different way of interacting with the network environment [1]. While most communications processes interact in a hierarchical fashion, peer-to-peer takes on a different methodology. Generally a P2P communication has the following characteristics:

- Direct connections between network clients
- Each client (node) is considered as an equal to all others
- Clients share processing, applications and content
- There is no central point of control within the network.

Peer-to-peer computing is the sharing of computer resources and services by direct exchange between systems. In a peer-to-peer architecture, computers that have traditionally been used solely as clients communicate directly among themselves and can act as both clients and a server, assuming whatever role is most efficient for the network. This concept of computing isn't new (the idea is over thirty years old), but the emergence of inexpensive computing power, bandwidth, and storage has inspired reconsideration.

2. Experiment

The Naval Postgraduate School (NPS) undertook an experiment to test the use of collaborative P2P communications in a wireless networked environment. This experiment was intended to provide initial data to evaluate the potential impact of using this technology in an urban warfare environment. The experiment was conducted on the NPS campus in March 2002 and involved a hostage search and rescue scenario within the confines of the NPS campus quad. The mission for the Reconnaissance and Surveillance Team (RST) on the ground was to build-up sufficient situational awareness by using P2P collaborative networking between RST members with Remote HQ, Local Command Post, and En route Scene-of-Action Commander.

According to the scenario an unknown number of terrorists (represented by evil-looking **RED** Triangles) took hostages on an upper floor of Spanagel Hall, Naval Postgraduate School. Six student teams comprising the RST unit, armed only with their laptops, Pocket PCs, wireless internet, GROOVE P2P collaborative tools, and set of GPS interface agents, were asked to collaborate with an en route special team commander and HQ in Norfolk to establish situational awareness sufficient to plan and execute the rescue mission. Figure 1 illustrates the situational awareness view of the RST collaborative environment at the moment when bomb (black circle) is been identified in the Root Hall and several terrorists were spotted in Bullard and Spanagel Hall buildings. By means of collaborative tools this view was shared by the HQ with the individual teams in order to enable their awareness of the other teams positions and their views of the targets. Performance of those applications was expected to be rather sensitive to the state of the wireless network and Network Operations Center was established to assist the RST teams in managing their application resources.

2.1 P2P Wireless Collaborative Network Configuration

A wireless network consisting of Cisco and Apple access points (Base Stations) were connected to the NPS LAN and were segmented from the main LAN by placing all the base stations on a separate subnet. Personal Data Assistants (PDA) and Laptop computers with wireless access cards were configured to connect to LAN through these access points. The experiment sponsor in Norfolk, VA had connectivity to this network through a secure pipe via the Internet. A network operations center (NOC) was established with a primary function of:

- Set up the experimental P2P wireless collaborative network
- Manage the network during the experiment
- Provide situational awareness to the local Command Center and sponsor headquarters
- Assist the operational team members
- Maintain communications during the experiment
- Collect the experimental data.

The research role of NOC was to explore the feasibility of bandwidth management for P2P clients, scalability and mobility of collaborative network, integration of P2P with client-server communications, and feasibility of P2P collaborative network self-organizing behavior.




Figure 1: The RST situational awareness view

The P2P wireless collaborative network was comprised of the following building blocks:

- Six mobile terminals each comprised of the Pocket PC (iPAQ) with GPS receiver and wireless laptop mobile connected to local area wireless network
- 15 access points distributed over the university campus area to provide local area coverage and routing functionality
- Four network monitoring workstations
- Two Situational Awareness web servers; one at local NOC the other at Sponsors facility in Norfolk, VA.

The first step was to create and capture topology. OPNET, a commercial network management software tool, allows topologies to be created manually or automatically[2]. Prior to the actual experiment, the managers created the topology manually because many of the devices were not configured with appropriate MIBs (Management Information Bases).

A software simulation tool was used to predict expected network performances. This simulation enables decision makers to predict the efficiency and capacity of a proposed network before equipment is actually acquired. The simulation results also provided detailed information on network traffic and differentiated the traffic attributable to the wireless segment of the network. Performance statistics in terms of Wireless LAN Load, Throughput, Data Dropped, LAN Delay, Media Access Delay, HTTP Traffic Sent, HTTP Traffic Received, HTTP Page Response Time, and HTTP Object Response Time were available. This data became crucial when allocating

actual resources and then it was used to anticipate limitations that would impact operational success.

Based on simulation performance analysis, the results warranted asset reallocation. Several components, including a second dedicated server, were added to the network. The analysis for the experiment determined that HTTP traffic would transmit through the P2P network without any serious delays. It also showed the network could handle an increased load without affecting service.

OPNET's Application Characterization Environment (ACE) Application was used to capture packet data necessary to analyze application specific loads. Files and associated packet traffic was traced and documented to create an accurate model of network data exchange. This data was used to populate both the application layer and network layer views in the network model.

Follow-on analysis was conducted using the functionality of ACE and an additional module called AppDoctor. AppDoctor provided analysis of the application task and reported on aspects of both network's and application's performance.

ACE was also used to analyze the use of IP addresses. Each node in the network was identified by its team name and an IP address. Dynamic IP addressing was used during the experiment, so some IP addresses were not easily or automatically associated with known network nodes. Dynamic IP addressing did allow the network to reassign an IP address the microsecond it was relinquished

Spectrum Network Management Software enabled the NOC to "drill down" into the network and provide detailed views of the network at user-defined levels [3]. Alarms or customized notifications could be established and system status changes indicated by a change in associated component icon color (from red to yellow depending on the parameter). Various views of the network included:

- Cablewalk view: The layouts of the access points that are connected to the LAN. Detailed information about each access point can be viewed by double clicking the associated icon.
- Device Topology. This detailed view details each network component (Figure 2). A normal connection is represented by a green color. An icon will turn red if performance has fallen beneath a set parameter.
- Link State View. Each component will display a green, yellow or red color depicting the health of the link.

Simple Network Management Protocol(SNMP) is the Internet standard protocol, defined in STD 15, RFC 1157, developed to manage nodes on an IP network. It is used extensively by Spectrum to discover, model and monitor the network [4]. Active TCP connections can be monitored for any SNMP compliant asset on the network.

The experiment was dynamic and exercised the spectrum of functionality of both a traditional and wireless hybrid network from a network management perspective. Spectrum network management software facilitated effective event tracking and system monitoring. The tools were versatile and allowed participants to see how their activities impacted the health of the network. There were sufficient user-defined parameters and alarms that allowed the NOC to shift assets to avoid impeding packet traffic during the scenario.

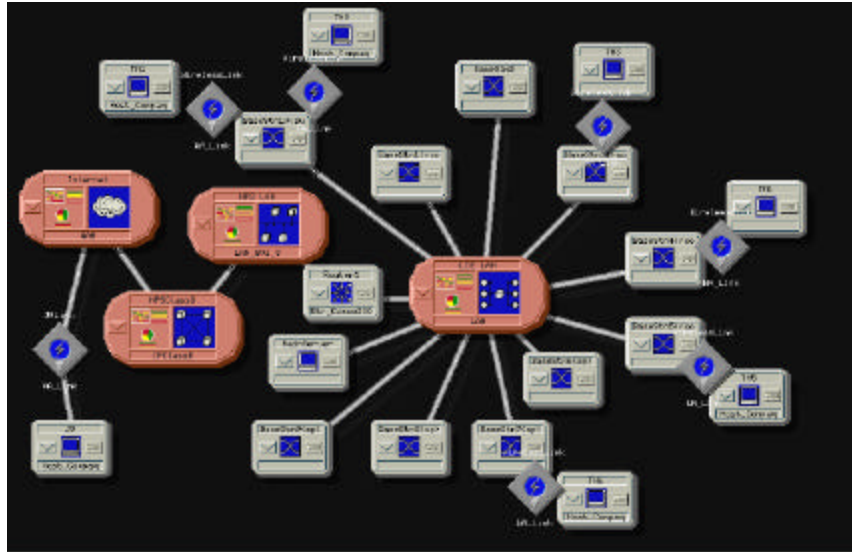


Figure 2: Network Topology

2.2 Network Monitoring

Solarwinds Network Management System (Engineering Edition, evaluation version) was the commercially available software used during the LOE to monitor elements of network performance and faults.

There were several factors specifically chosen to monitor during the P2P experiment including:.

- Network Performance
- Current Response Time and Percent Packet Loss
- Average Response Time and Percent Packet Loss.

Solarwinds Network Management software was also used for fault management. Network performance and fault management could be monitored simultaneously. The following elements of fault management were evaluated:

- Monitored events and traps originating from the wireless network elements.
- Configured alarm severity levels along with filters based on time, source, severity, and type.
- User-defined action scripts registered for certain alarm types or network element instances. Actions could initiate NOC Manager notification through e-mail or pages (beeper).
- Color-coded hierarchy display for alarm level indications.
- Number and time distribution of selected alarms, alarm severity, alarm state, or network elements effected.

2.3 Bandwidth Monitoring

The bandwidth monitor feature of Solarwinds provided a variety of display options. The primary limitation of this function was that each terminal be SNMP compliant. In the experiment, none of the hand-held terminals and only four of the six laptop terminals had functional MIBs. Bandwidth capability had to be monitored on the servers.

SolarWinds TraceRoute module was useful in evaluating bandwidth usage. The utility will not only the path network traffic takes from each node on the network, it also displays selected SNMP information about each device encountered.

Response time and packet loss information could be displayed both as a number and as a bar graph. TraceRoute could also be used to evaluate or query SNMP compliant machines outside the network.

Other filters were initially created to capture packets from ports utilized by Net meeting and Groove but were found to unsuccessful during dry run tests. Additionally, a “No filter” capture agent was created to capture all data flowing between the server and team devices but was lost due to parsing errors associated with initial header size configuration settings. As a result, this analysis focused primarily on TCP and HTTP Port 80 activity captured at the server.

Capturing packets was initiated at the beginning of the LOE. System errors that occurred during the experiment required re-initialization of the capture process. The result was that only the last segment of the LOE was recorded and compiled for analysis.

The trace file was then imported into ACE and configured via the import wizard for proper network reference alignment. What became evident in this analysis is that network degradation could often be correlated to specific application use.

Microsoft Net meeting, MAPIX and GPS data transmissions to the LOE Server degraded the network to some degree, but in most cases, degradation was consistent from team to team and from wireless access point to another.

3. Results and Discussion

Factors affecting overall performance of the experiment network focused on the application layer. Performance metrics were not consistent across all devices, but this could be attributed to location of the individual teams relative to the wireless access points or individual laptop application configurations with regards to processes running in the background on each node.

The primary recommendation to improve application packet transfer would be coordinated “turnkey” configurations on each node of the network. Specifically, adjust the system configurations so there are minimal applications running in the background on the nodes.

Bandwidth was not an issue during this experiment. Bandwidth utilization for each of the terminals averaged around one percent of capacity and peaked at two percent. Figure 3, Bandwidth Received and Transmitted shows that the average bandwidth received was around 100 Kbps, and transmissions averaged around 10 Kbps.

While this experiment was not bandwidth intensive, the percent packet loss averaged around 35.2 indicating that the network was not configured for applications that required dedicated bandwidth. (Figure 4)

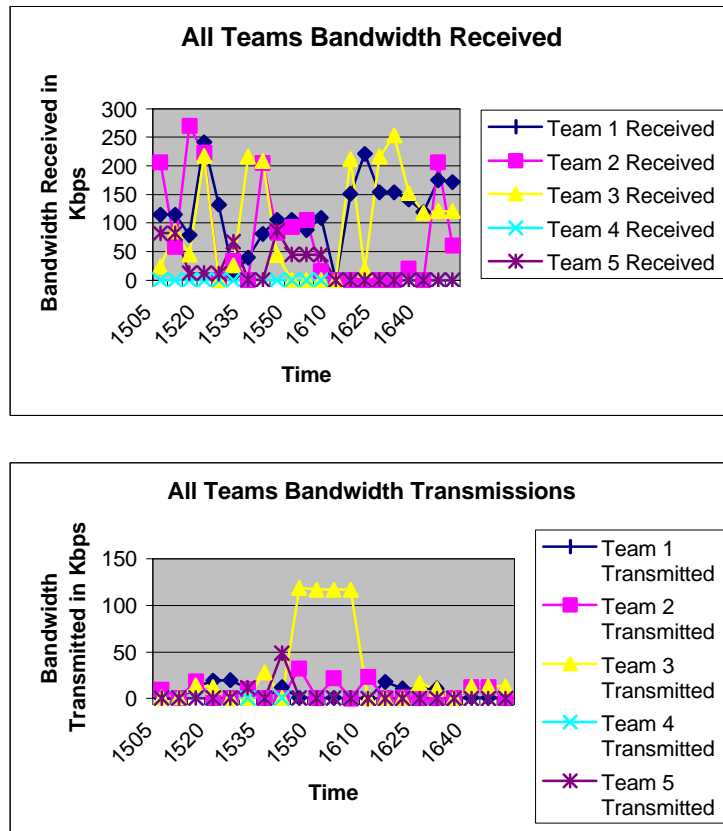


Figure 3: Bandwidth Received and Transmitted

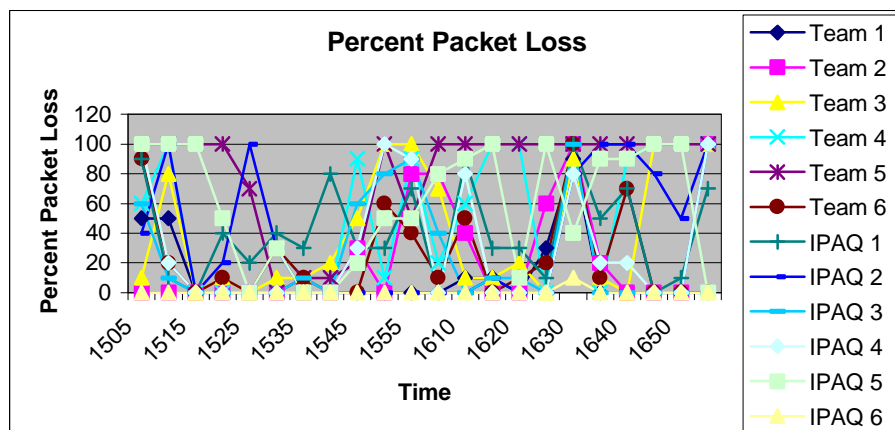


Figure 4: Percent Packet Loss

A mobile node should be able to monitor its own signal strength and bandwidth utilization. This was a critical form of operational feedback provided to the teams from the NOC. The result was the teams adjusted their physical location or changed applications being used on their devices.

The experiment demonstrated the scalability of a wireless P2P collaborative networking, yet emphasized the network overhead needed to synchronizing voice over IP communication. Voice packets were sequentially routed with other application packets, but the result was seemingly broken communication. Other traditional voice communication modes were more reliable. The data sharing features scaled-up effectively.

The experiment demonstrated that P2P and Client-Server integration is feasible, but sensitive to roaming between the access point coverage areas.

Application sharing was especially sensitive to roaming as applications would drop when a team crossed a boundary of access point coverage. There was substantial packet loss until the application was restarted in the new area, so error checking and system synchronization/restoration features are necessary

Self-organizing behavior was demonstrated when Reconnaissance and Survey team members switched modes of communication due to signal loss or interference. Yet, the strongest (and unexpected) effect of self-organizing behavior emerged at the command and control center site when network center managers were able to effectively monitor performance and fault data, synchronize this data with the voice and data sharing calls, and adjust assets or operations before packets and connectivity between peers was lost. Essentially, new channels of communication between team members were facilitated in real time by the NOC monitoring team elements.

4. Conclusions and Future Research

This experiment demonstrated how network management tools can be used to view a complex organization and monitor the flow of information.

Network response time was, for the most part, consistently around fifty milliseconds for all terminals.

The greatest indicator of network flow was the analysis of the percent packet loss. Every team was “all over the board” in regards to packet loss. Each team dropped approximately 35.2 percent, yet even with re-transmission of the affected data packets, the twelve wireless units still consumed less than one percent of the bandwidth available.

Wireless P2P collaborative networks are feasible, but the application programs used for communications are not yet robust enough to support mission critical environments. Future research should focus on the stability of the application layer and the capabilities from the communications programs to automatically re-establish communications if dropped from a mobile network. While bandwidth was not an issue in this small experiment, it is important to remember that much of the utilization of the bandwidth that was used came from the re-transmission of data packets lost through application drop-off.

5. References

[Bordetsky, 2000] Alex Bordetsky. *Adaptive QoS Management via Multiple Collaborative Agents. Ch.4 in: Agent Technology for Communications Infrastructure* Wiley Press, 2000

[Lewis, 2000] Lundy Lewis. *Service Level Management for Enterprise Networks*, Artech House, 2000

[Perkins and McGinnis, 1997] David Perkins and Evan McGinnis , *Understanding SNMP MIBs*, Prentice Hall, 1997

[Opnet.com, 2002] *Software that Understands Networks*, <http://www.opnet.com/>