Peer-to-Peer Communications for Tactical Environments: Observations, Requirements, and Experiences

Niranjan Suri and Giacomo Benincasa, Florida Institute for Human and Machine Cognition
Mauro Tortonesi and Cesare Stefanelli, University of Ferrara
Jesse Kovach and Robert Winkler, U.S. Army Research Laboratory
Ralph Kohler and James Hanna, U.S. Air Force Research Laboratory
Louis Pochet, U.S. Air Force Reserves
Scott Watson, Space and Naval Warfare Systems Center, Pacific

ABSTRACT

Tactical edge networks present extremely challenging environments for communications given their wireless ad hoc nature and the inherent node mobility. Military applications such as Blue Force Tracking, inter-team communications, remote unmanned vehicle control, and sensor data mining/fusion thus have to deal with unstable links with limited bandwidth and variable latency. The peculiar characteristics of tactical networks call for peer-to-peer approaches to realize complex, adaptive, and fault-tolerant applications to be deployed in the battlefield. This article reports on our observations from several tactical networking experiments in which we have deployed state-of-the-art applications and services that leverage P2P communications. More specifically, we discuss why P2P approaches are critical for tactical network environments and applications. We then analyze the requirements that should be satisfied by P2P middleware for tactical environments. Finally, we discuss a case study, the Agile Computing Middleware, and present experimental results that demonstrate its effectiveness.

INTRODUCTION

Tactical edge networks present an extremely challenging communications environment for application developers and users. These networks are built from ad hoc wireless connections between mobile nodes, providing unstable links with limited bandwidth and variable latency. In urban environments, performance is further degraded by occlusions from buildings and interference from consumer electronics and civilian transmitters. Meanwhile, a greater number of assets equipped with higher-fidelity sensors, along with the proliferation of new end-user applications, are placing ever increasing bandwidth demands on the network. Systems operating in these tactical networks must be capable of providing reliable and timely information exchange within this unreliable and congested communications environment. Creative solutions that combine multiple techniques are needed to successfully address these challenges and realize the goals of network-centric warfare.

Client-server approaches such as service-oriented architectures (SOAs) are commonly adopted as the basis to realize applications and services in military systems running on higher-echelon command and control networks. In an unreliable bandwidth-constrained tactical environment that is subject to network partitioning, client-server architectures can introduce centralized points of failure and performance bottlenecks. Moreover, the unicast point-to-point connections result in excessive bandwidth consumption when data is sent to large numbers of clients. Peer-to-peer (P2P) approaches do not rely on designated server nodes that must be reachable and therefore can continue to (partially) function in a partitioned network. Also, P2P systems can make use of multicast and other advanced data distribution schemes that minimize the transmission of redundant information. Finally, because communications do not need to be routed through central servers, P2P technologies can leverage the fact that many applications place greater importance on communication between nearby nodes as opposed to distant nodes. These characteristics make P2P architectures a better fit for tactical networks than traditional client-server designs.

This article focuses on P2P systems design as applied to middleware and application development for tactical networks. P2P systems are usually classified as structured or unstructured network architectures [1]. In structured P2P sys-

0163-6804/10/$25.00 © 2010 IEEE

tems, nodes cooperate to maintain a distributed database that contains information about the location of resources (i.e., nodes, files, services, etc.) within the network. In unstructured P2P systems, instead, nodes have to advertise resource location or discover resources by broadcasting queries over the network. A resource location database reduces the number of messages required for resource discovery and facilitates the discovery of non-replicated resources. However, in highly dynamic network environments such as tactical networks, the bandwidth and computational overhead required to maintain the distributed database in a consistent state typically outweighs its benefits. On the other hand, unstructured P2P systems have to deal with scalability issues related to the broadcast of queries across the network, but their inherent resilience to node churn makes them better suited to mobile ad hoc and tactical environments.

Based on experiences from experiments and exercises, we present scenarios, such as Blue Force Tracking (BFT), remote unmanned vehicle control, and sensor data mining/fusion, that benefit from P2P approaches. These observations are subsequently developed into specific technical requirements. The observations and requirements that are presented take into account the perspectives of the Air Force, the Army, and the Navy and Marine Corps. While the focus is on P2P aspects, we briefly discuss the intersection of P2P with client/server approaches and SOAs.

Finally, the article presents a case study using the Agile Computing Middleware, a software framework that enables the realization of applications and services for tactical edge networks. We discuss the Group Manager discovery and the DisService information dissemination components, which support P2P applications. We provide experimental results that demonstrate the effectiveness of the P2P approach. The Agile Computing Middleware, initially developed to address the requirements of Army battlefield environments, has been extended to support Air Force requirements, and is currently being extended to satisfy Navy and Marine Corps requirements.

Peer-to-Peer Communication Scenarios in Tactical Environments

During the last eight years, we have collectively participated in a number of tests, experiments, demonstrations, and exercises involving tactical environments. These include the Horizontal Fusion Quantum Leap experiments in 2003 and 2004 (QL-1 and QL-2), the C4ISR[1] On the Move (C4ISR OTM) experiments in 2006, 2007, and 2008, the Joint Forces Command (JFCOM) Empire Challenge in 2009 (EC-2009), JFCOM BoldQuest (BQ) exercises in 2007 and 2009, the Office of Naval Research (ONR) ISR-C2[2] experiment in 2009, the Defense Advanced Research Projects Agency (DARPA)/SSC Pacific SIE-DTN[3] field test in 2010, and the Joint Expeditionary Forces Experiment (JEFX) in 2008, 2009, and 2010. Figure 1 shows the typical nature of the networks we have observed in these exercises — clusters of edge nodes that are interconnected via reach-back links. In particular, we note that the size of the individual clusters range from 10 to 20 nodes. P2P approaches are most applicable within these individual clusters, with other communication protocols being more appropriate for the reach-back links [2]. In this section we distill and summarize the scenarios and discuss the benefits of P2P approaches.

Figure 1. *Tactical networking environment with peer-to-peer clusters.*

Blue Force Tracking

Perhaps the most fundamental need for tactical users is Blue Force Tracking (BFT) — applications that provide situational awareness information regarding the presence and location of friendly forces. BFT is critical to avoid friendly fire accidents. The dynamic, geographically sensitive nature of most tactical situations makes BFT applications well suited to P2P architectures. Even in an unreliable communications environment that prevents connections to a central BFT server, nodes in close proximity are likely to still be able to communicate with each other and share critical location information. P2P-based BFT systems enable this mode of operation. Furthermore, P2P approaches can be more bandwidth-efficient — a critical resource in tactical networks.

BFT also provides a good example of the proximity/precision correlation. When users are operating in close proximity, it is more important that their relative positioning be exchanged reliably, quickly, and frequently between each other. As they are located farther apart, the information may be aggregated and exchanged less frequently. This requirement occurs in other application scenarios as well, and advanced P2P architectures can provide this capability.

Resource Discovery and Awareness

Another fundamental need of tactical users and applications is discovery and awareness of available resources or assets. Resources can include deployed sensors, manned vehicles, autonomous air and ground vehicles, and even other users themselves. Resources may pertain to the platform itself or to services provided by a platform.

[1] C4ISR is an acronym for Command, Control, Communications, and Computers for Intelligence, Surveillance and Reconnaissance, and usually refers to the set of systems and networks that enable net-centric operations.

[2] ISR-C2 is an acronym for Intelligence, Surveillance, and Reconnaissance for Command and Control — a concept similar to C4ISR.

[3] SIE-DTN is the Service Interoperability Environment with Disruption Tolerant Networking — a field test conducted by the Navy's SPAWAR Systems Center Pacific and DARPA.

Unmanned ground and air vehicles are becoming more common in military operations. As the number and type of unmanned assets increase, the importance of distributed P2P technologies for discovering and controlling these assets increases as well.

For example, an unmanned air vehicle (UAV) may provide a *video feed* service. Resource discovery is an essential prerequisite for many other operations, such as tasking of autonomous vehicles, gathering data from sensors, and communicating with other users. For example, at QL-2, C4ISR OTM, and ISR-C2, dismounted soldiers were discovering available robotic assets, deployed perimeter sensors, as well as other soldiers. In the JEFX and BoldQuest exercises, a user on the ground was discovering airborne assets along with data feeds being generated by the airborne assets.

An equally important aspect that complements resource discovery is continued awareness of the resource. For example, consider a user that has tasked perimeter trip-wire sensors to notify the user of any incursions, but never receives any notifications from the sensors. This may simply be because the sensors have not detected anything. It could also be because the sensors are malfunctioning or out of communications range. There is a significant difference between these two scenarios, and the user must be able to differentiate between them. Systems must provide status information regarding assets and the communications links between them so that users can be confident that the system is functioning properly.

Client-server discovery systems generally rely on one or more centralized registries, which become critical points of failure in the unreliable tactical network environment. With a centralized system, two nodes that are able to communicate with each other, but not with the central registry, will not be able to discover each other. There is also the initial problem of locating the registry, usually solved by statically configuring each client with the address of the registry server. If the registry must be moved, or if a client would be best served by using another registry server, the client's configuration must be changed. This is not always possible. For these reasons, resource discovery and awareness services in tactical edge networks are best provided through P2P approaches.

Geographically Directed/Constrained Operations

In military applications, many requests for information or assets are constrained by geography. For example, a trip-wire perimeter sensor that is triggered may cause a user (or an automated process) to look for a visual sensor that can provide imagery in the vicinity of the sensor's location. Similarly, a user may wish to find the closest available robot or UAV to a building to be examined or monitored. The geographical constraints may apply to the node itself or, in the case of a sensor node, to the sensor's area of coverage. A related feature is a geographic subscription to information. For example, the RouteScout application for JEFX 2010 allows the mission path to be used as a geographically constrained search for information along the path. For the Marine Corps, we have developed a mechanism to model the mission characteristics, including the set of possible paths, which is subsequently used by peers to push relevant information to the users' nodes.

Sensor Data Exchange

Unattended ground sensor systems currently fielded by the military generally rely on satellite communications for data transmission. This approach has a number of drawbacks. Satellite communications require a transceiver with a clear view of the sky, and low-power satellite modems suitable for embedded sensor applications are very slow (on the order of a few kilobits per second). By using P2P models combined with mobile ad hoc network technologies, sensors can use high-speed short-range radios to exchange and exfiltrate data to other sensors and nearby units. With a high-speed P2P link, a sensor can send high-resolution imagery and motion video that would be impractical to transmit over a low-speed satellite link. Additionally, P2P technologies can be used to build store-and-forward sensor networks, where a sensor that is not in communications range of a data consumer will locally store detection data and forward it to consumers as they come in range. These P2P approaches can work side by side with existing satellite reach-back systems, enabling advanced applications without interfering with current infrastructure and doctrine.

Unmanned Vehicle Operations

Unmanned ground and air vehicles are becoming increasingly common in military operations. As the number and type of unmanned assets increase, the importance of distributed P2P technologies for discovering and controlling these assets increases as well. For example, if a combat unit has a number of unmanned assets and a number of controllers, a P2P discovery mechanism is required so that any controller can discover and connect to any asset, as it may not be known in advance which controllers will be used with a particular asset. P2P systems are also useful for integrating unmanned assets into the BFT systems described earlier. Additionally, a major focus of current unmanned systems research is developing collaborative behaviors where multiple assets cooperate to accomplish a task. To execute these behaviors, assets must be able to discover other nearby assets and then exchange state information as the task is being performed. The proximity/precision correlation applies here as well. P2P systems are useful for accomplishing these goals in a decentralized manner.

Sensor Data Mining, Data Fusion, and Distributed Cross-Cuing

Multiple factors are driving the battlefield to be increasingly covered with distributed sensors that are capable of generating large quantities of data. Data mining, data fusion, and cross-cuing are all mechanisms to reduce data overload by aggregating, correlating, transcoding, and filtering raw data into useful information. When such sensor operations are in close geographical proximity, P2P approaches are more effective because they can exploit multicasting over short-range wireless links that require less power and provide higher bandwidth.

P2P systems can also be used to provide automated, dynamic, self-configuring cross-cuing

capabilities between unattended sensors and other assets. For example, a small sensor that contains a motion detector or vibration sensor but no camera can be used to trigger a different camera-equipped sensor to take an image or record a video, cuing the camera as necessary. Additionally, if a UAV or other similar asset happens to be nearby, the imager on the UAV could also be cued and tasked.

Secure Chat, Media Exchange, and Social Networking

Military operators have embraced aspects of social networking to improve effectiveness of mission execution. Tools include instant messaging, marking up and sharing multimedia objects such as pictures and video, collaborative map annotation, and other forms of communication. Soldiers use these tools to maintain situation awareness, coordinate, and exchange information in an effective and clandestine manner. They share useful tidbits of intelligence information, contact subject matter experts to ask for opinions, and dynamically replan their strategy as necessary. Many of these operations involve communicating with peers who are on the same mission or in their physical proximity. Therefore, P2P systems are the best approach to supporting these scenarios.

Technical Requirements for Peer-to-Peer Systems

In this section we develop technical requirements for P2P systems that can address the types of scenarios described in the previous section. We have addressed many of these requirements in developing the Agile Computing Middleware, described later in this article.

Automatic Configuration

The dynamic nature and high operational tempo of most tactical environments does not allow time for system configuration. Therefore, wherever possible, systems should automatically configure themselves within the network and provide the capabilities necessary for the users. P2P systems naturally address this requirement by removing the need to configure clients with information about servers that need to be accessed. Systems should also minimize requiring users to configure network settings and specify IP addresses of other information producers and consumers. Gateways and other structural aspects of the P2P network should be automatically discovered as nodes join and leave networks.

Bandwidth-Efficient Peer Discovery

Discovery of other peers is the most fundamental requirement for P2P systems, as it is a prerequisite for most other operations. Peer discovery is closely related to service advertisement and search. Typically, peer discovery applies at the platform or node level, whereas advertisement and search apply at the service level. For example, at C4ISR OTM, each ground robot advertised a Robot Agent service that was discovered by operator control units (OCUs) used to tele-operate the robots. Similarly, high-mobility multipurpose wheeled vehicles (HMMVWs) advertised numerous services such as Blue Force Tracking, data fusion, and language translation services that were looked up by other peer users.

Peer discovery must be bandwidth-efficient and adapt to different types and qualities of network links. Ground-based military operations on urban terrain (MOUT) scenarios such as QL and C4ISR OTM use wideband radios that provide reasonable bandwidth on the order of 1 Mb/s. However, the airborne networking environment is more constrained. For example, during BQ 2007 and JEFX 2007, long-range UHF links were utilized between the aircraft and ground nodes, which provided a bandwidth of 10 kb/s or less. In BQ 2009 and JEFX 2009 the move to new prototype radios (e.g., DARPA Quint Networking Technology) provided much higher bandwidth. However, these broadband radios and waveforms provide significantly shorter range, and certain operations still require low-bandwidth UHF links. Therefore, it is critical that peer discovery be adaptive in terms of bandwidth utilization.

There are several commercial and open source efforts that address peer discovery. Examples include JXTA [3], Zeroconf [4] implementations such as Apple's Bonjour and Avahi's mDNS/DNS-SD, XMPP [5], Session Initiation Protocol (SIP) [6], and uPnP [7]. Our experience with these components shows that they do not perform well on tactical networks. These components have been designed for infrastructure networks that do not have the bandwidth and reliability challenges of tactical edge networks. We compare the performance of JXTA with the Agile Computing discovery mechanism in the section on experimental results.

Peer-Loss Discovery

Peer loss discovery is important for providing situational awareness to tactical users. In client-server environments, a user often detects the loss of the other entity simply by detecting the loss of the connection. However, in P2P environments where the communication tends to be point-to-multipoint (e.g., using multicast), peer loss discovery is more difficult. Peer loss discovery directly contributes to increased awareness of surrounding resources.

Dynamic Information Channeling According to Topology and Resource Availability

To support data mining and data fusion, as well as geographically directed resource discovery, it is necessary to provide mechanisms to convey information through a subset of the network dynamically determined by topological or resource availability constraints and manipulate data as it is transferred between peers. This is a challenging task that requires monitoring of the network topology as well as dynamically discovering and allocating computational resources along the communication path. In addition, to minimize the overhead of on-demand realloca-

Military operators have embraced aspects of social networking to improve effectiveness of mission execution. Soldiers use these tools to maintain situation awareness, coordinate, and exchange information in an effective and clandestine manner.

The dynamic nature of tactical networks requires disruption tolerant approaches to information dissemination. Sometimes, information must be delivered to nodes that periodically disconnect from the rest of the network, requiring reliability mechanisms such as caching and periodic retransmission of important data.

tion protocols, proactive communication path rerouting and computational resource reassignment should be implemented based on predictions about future user requirements.

Adaptive Dissemination

Support for information dissemination is essential in tactical networks. To minimize communication overhead, different protocols and algorithms can be used depending on the set of receivers and the requirements of the application. In our experience tactical applications present multiple patterns of data dissemination. BFT is transmitted from each node to every other node in a many-to-many pattern. Sensor fusion requires many nodes (sensors) to transfer data to one node (fusion node) and then onto some consumers in a many-to-one-to-few pattern. For maximum efficiency, the data dissemination system must adapt its strategy based on the dissemination pattern. For instance, information that must be delivered to most of the nodes in the network is best transmitted as an all-nodes broadcast. On the other hand, information that must be delivered only to a small set of nodes, especially when in close proximity, should adopt different mechanisms. The same considerations apply to the case of multiple information sources.

Adaptation should exploit prioritization, differential update rates, reliability, and sequencing requirements as well. As discussed earlier, there is often a direct relationship between physical proximity and the requirements for precision in the information exchange between nodes. An adaptive dissemination system can exploit this to reduce bandwidth utilization. Similarly, not all consumers require the cost associated with reliable and/or sequenced delivery of information. The dissemination system should be flexible in supporting these properties when required and saving bandwidth when they are not necessary.

Disruption Tolerant and Robust Information Dissemination

The dynamic nature of tactical networks requires disruption-tolerant approaches to information dissemination. Sometimes, information must be delivered to nodes that periodically disconnect from the rest of the network, requiring reliability mechanisms such as caching and periodic retransmission of important (non-obsolete) data. The SIE-DTN experimentation has shown that disruption tolerance significantly improves the ability for nodes to receive data under intermittent connectivity conditions.

Systems should also provide mechanisms for robust information delivery. Techniques such as forward error correction, scattering (transmitting different portions of an object to receiving nodes and letting them recompose the original object), and layered coding (the decomposition of an object into smaller and lower-fidelity usable parts that can be rejoined by the receiver) should be adopted to improve the availability of large information objects.

Finally, systems should exploit common patterns in node mobility and information to improve the timeliness and availability of information dissemination. This requires learning mechanisms that can process node mobility and data/service usage information, identify common behaviors, and produce reliable forecasts that can be used as the basis for decision making in information caching and routing.

Tracking Information Pedigree

Tracking pedigree of information entails securely maintaining meta-information regarding the source of the information, the time when the information was originated or generated, and the path the information has traversed before reaching a consumer. Information pedigree is often essential to decision making. While this requirement applies to many types of information systems, we highlight it here because P2P systems make it more challenging to track information pedigree. For example, with a client-server system, if the server is a trusted entity, the server can often assign or validate some aspects of the pedigree of the information provided by a client, and pass that on to other clients. In a P2P system, if some of the peer nodes are not trusted, additional mechanisms may be necessary to either ensure that the information satisfies necessary security requirements or adjust the pedigree appropriately.

Policy-Based Control over Information Exchange

Another requirement relates to control over the nature and destination of information as it is managed and disseminated in P2P systems. The frequent occurrence of joint and coalition operations, and the associated restrictions on information sharing imply that P2P systems should address this requirement. Client-server approaches are more amenable to such control over information sharing, as the servers can act as gatekeepers of the information and enforce any policies on sharing. This problem must be addressed within P2P systems if they are to be pervasive and span multiple administrative domains at the tactical edge. An alternative is a hybrid approach that uses P2P within a single administrative domain and trusted server gateways to span domains.

Integration with Service-Oriented Architectures

SOAs have been widely adopted by military systems running on higher-echelon command and control (C2) networks. SOA-based approaches allow for extensive service reuse and the integration of heterogeneous services, with significant savings in development costs and time for building distributed applications. SOAs also allow the dynamic (re)composition of services at run-time, thus enabling the ad-hoc realization of complex and adaptive distributed applications.

Because of these advantages, SOAs are often proposed in tactical networks despite the challenges these architectures face in that environment [8]. Traditional SOA implementations are based on centralized service directories and strong assumptions about relatively static network topologies. Therefore, they are susceptible

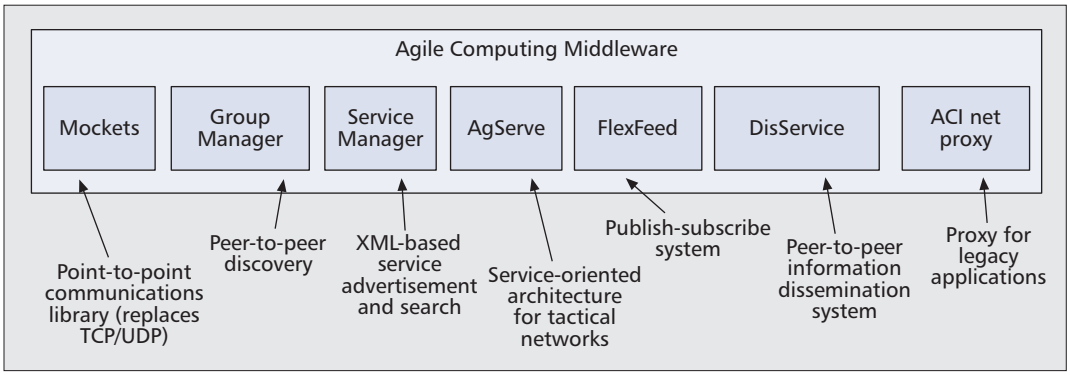

Figure 2. *Components of the agile computing middleware.*

to poor scalability (also caused by their RPC-based paradigm that hinders the adoption of caching) and are limited by relatively high computational and bandwidth requirements and lack of support for service migration.

These limitations call for tactical network-specific SOA implementations as well as ad hoc integration solutions between P2P and SOA middlewares [9, 10]. As a result, P2P systems for tactical applications should include an interoperability layer that enables resource discovery and service accessibility across SOA and P2P architectures.

Agile Computing Middleware

The Agile Computing Middleware [11] has been developed over the course of the last nine years to address many of the requirements enumerated in the previous section. Figure 2 shows the key components of the middleware. In particular, this section presents two components — Group Manager and DisService — that realize P2P discovery and information dissemination. The ACINetProxy, which helps integration with legacy systems, is also briefly described. The next section presents experimental results that demonstrate the effectiveness of the P2P approach in Group Manager and DisService.

Group Manager

The Group Manager component supports peer discovery and has been optimized to be bandwidth-efficient for tactical edge networks. It supports a flexible combination of proactive advertisement and reactive search. When advertisement is activated, the frequency of advertisements depends on the node movement and the churn rate of the network. Hence, a fast-moving entity such as an aircraft would advertise more frequently, whereas a stationary ground sensor advertises slowly. Nodes can also control the *strength* of their advertisement, which is defined in terms of a metric (number of hops, distance, or link quality). The frequency and strength of advertisements provide a trade-off between discovery speed, discovery range, and bandwidth utilization. Providing control over this behavior allows nodes to select the best possible combination to suit their needs. Ground nodes that need to maintain radio silence could choose instead not to advertise at all, thereby preventing their discovery.

Discovery occurs within the context of a group, a concept that allows partitioning of network nodes into different sets. A group is identified by a name with a simple hierarchical notation, such as `mil.army.arl.hf2004.ugvs`. Groups are a convenient abstraction to aggregate resources based on mission or ownership (e.g., country or branch of the armed forces). Nodes enable discovery among themselves by simply joining the same group. For example, all the nodes that are assigned to a specific mission could join a common group, thereby allowing them to be discovered by each other. Nodes are allowed to join multiple overlapping groups, and to advertise different capabilities within each group. Therefore, a node may advertise service X to other nodes participating in the same mission and service Y to other nodes that are on a different mission but belong to the same country.

Group Manager also supports private groups with restricted access and encrypted communications. Private group security leverages a preshared passphrase, often proposed in tactical military applications, where direct access to the trusted authority (TOC) is not available for encryption key generation and key exchange. Private groups make it very simple for dynamically formed ad hoc teams to verbally exchange a passphrase that may then be used to restrict access to the group. We have used this capability to easily build secure P2P chat applications that dynamically manage chat members based on their mission, role, or any other attribute. Note that encryption in private group communications does not replace standard cryptographic measures provided by most tactical radios; instead, it represents an additional security layer on top of them.

Group Manager handles discovery and maintains group membership in a completely decentralized manner. This design choice implies that there is no attempt to maintain a consistent group view of all the nodes that are members of a group, as it would require an excessive amount of bandwidth. Distributed hash table (DHT) approaches were also discarded for their high bandwidth requirements and given the high churn rate of tactical networks.

DisService

DisService is the P2P information dissemination component of the Agile Computing Middleware. In DisService information is pushed and received

SOAs have been widely adopted by military systems running on higher-echelon command and control networks. SOA-based approaches allow for extensive service reuse and the integration of heterogeneous services, with significant savings in development costs and time for building distributed applications.

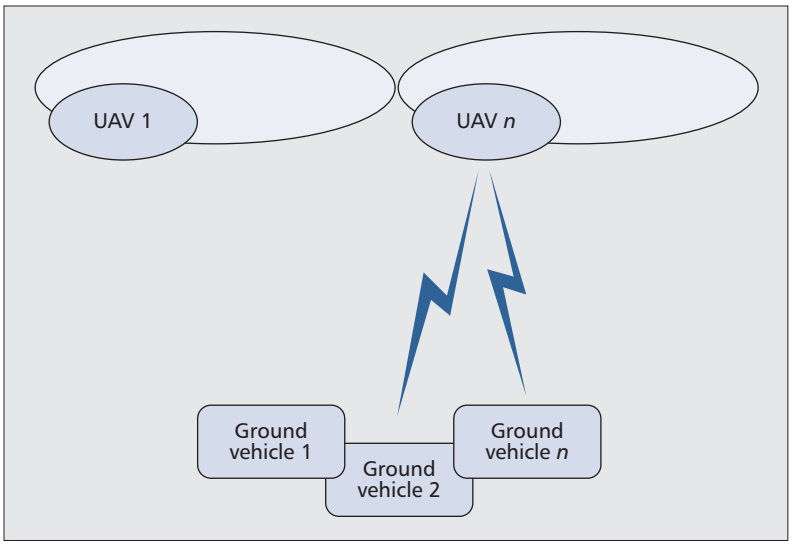

Figure 3. *Peer discovery experiment scenario.*

in the context of a group, using the same abstraction proposed by Group Manager. DisService is disruption tolerant and realizes several novel architectural approaches to improve dissemination and availability of data.

One of the fundamental assumptions in DisService is that in a P2P wireless environment, the cost to broadcast or multicast a message to a neighbor is the same as with a unicast. Therefore, DisService always prefers broadcast or multicast communications, which enable all other direct neighbors to also receive and archive the data.[4] This feature, called *opportunistic listening*, significantly increases the availability of data, especially in a tactical network where nodes and links are unstable. The next section presents experimental results that show the benefits of this feature.

Another DisService feature is flexible support for sequential and reliable message delivery. When a node subscribes to a group, it can choose whether messages should be delivered sequentially and/or reliably. Typically, sequencing increases latency and reliability increases bandwidth utilization. Therefore, DisService allows applications to make trade-offs that improve their efficiency. For example, a node could prefer unreliable but sequenced delivery for BFT information, while requiring reliable unsequenced delivery for sensor reports.

The third unique feature of DisService is that the burden of reliability, if desired, is placed on the receiver as opposed to the sender. In both the TCP and DTN models, the sender has the responsibility to retransmit the data until the receiver, or an intermediary in the case of DTNs, has acknowledged receipt. DisService instead adopts a different model where the receiver, if so desired, requests missing messages (or parts of messages) from other nodes until they have been received. The motivations behind the adoption of this model are the point-to-multipoint nature of DisService and the observation that different peers that are recipients of the same data may have different requirements for reliability. Furthermore, this approach exploits the opportunistic listening capability of DisService, where a node looking for the missing data may be able to obtain it from any other peer as well. The conventional model, where the sender wishes to ensure that a set of target nodes receive data, is also available.

DisService also supports multiple dissemination algorithms, including reliable flooding, probabilistic (epidemic) protocols, and heuristic protocols, which can be selected based on the type of information being disseminated. The specific dissemination algorithm chosen depends on the number of subscribers to a group, the nature of their subscription (reliable or not), and the relative priority. Reliable flooding is an expensive algorithm, but is appropriate for high-priority messages that must be delivered reliably. Probabilistic protocols use attributes such as the number of neighbors of nodes to determine the dissemination paths. Heuristic protocols further enhance probabilistic approaches by exploiting domain-specific knowledge to guide the probabilistic dissemination. One example of a heuristic is favoring the use of a UAV as a good intermediate node for information propagation, given the increased reach and visibility of a UAV. Another example is a node acting as a relay to a neighbor that has only one other neighbor (which implies that the first node is the only way for the second node to receive any information).

Finally, DisService provides a different approach to handle dissemination of large messages, such as multimedia objects (pictures, video, etc.). In these cases only the metadata describing the large message is disseminated to the subscribers. Each subscriber, upon evaluation of the metadata, may subsequently choose to retrieve the complete message. DisService supports two mechanisms to handle the actual dissemination of large messages. In the simple case the message is transferred as a whole from one of the source nodes to the destination node, which can now act as an additional source node. In addition, DisService supports a second approach which fragments the large message into chunks that are scattered across the network and replicated on a bandwidth-available basis. When a node wants to retrieve a message, it queries its peers to find the chunks, retrieves them, and reassembles the original message. Furthermore, where possible, the chunks are created using a data-type-specific layered encoding algorithm that allows each chunk to be independently usable. For example, large images are broken up into multiple complete but lower-resolution images that are individually usable but, when combined, recreate the original high-resolution image.

INTEGRATION

The best approach to leveraging the capabilities of the Agile Computing Middleware is to directly take advantage of the application programming interfaces (APIs) of the various components. However, we recognize that this is not always possible, especially with legacy applications. The middleware includes the ACINet-Proxy component, which provides a transparent mechanism to capture all the outgoing traffic generated by legacy TCP- and UDP-based appli-

[4] *Note that the cryptographic and security features included in tactical radios will ensure that nodes that should not be receiving the data will not do so. If further restricting is necessary, the notion of private groups, as described for the Group Manager, can be applied to DisService as well.*

cations and redirect it using middleware components. This capability is particularly useful for redirecting TCP connections to use Mockets [2] and UDP multicast to DisService. ACINetProxy is configurable via policies to support deployment on a wide range of scenarios and permits fine-grained performance tuning (e.g., prioritization and shaping of traffic).

Experimental Results

This section presents results from two experiments that evaluate the performance of the Group Manager and DisService components.

Peer Discovery Results

In this first experiment we compare the performance of Group Manager with the JXTA middleware for P2P computing. In particular, we compare the discovery aspects of Group Manager and JXTA, and measure the relative bandwidth utilization. JXTA was chosen because its discovery mechanism is utilized by the System of Systems Common Operating Environment (SOSCOE), part of the U.S. Army Future Combat Systems and other programs. The scenario, shown in Fig. 3, consists of two to four UAVs and four to 18 ground vehicles (GVs). The connectivity between the GVs and the UAVs varies based on the position of the UAVs. In this particular experiment the GVs were attempting to discover a service advertised by the UAVs. The connectivity between the nodes was emulated using an enhanced version of the Naval Research Laboratory's Mobile Ad Hoc Network Emulator [12].

The results are shown in Table 1. While the peer discovery was completed successfully with both components, the Group Manager used substantially less bandwidth than JXTA. As the results show, JXTA used between 3.58 to 7.77 times more bandwidth than the Group Manager in default mode. The Group Manager also supports an enhanced mode that uses periodic pings to check whether nodes are still reachable, thereby addressing the peer loss discovery requirement. Even with this enhancement (which JXTA does not support), the Group Manager is more efficient, with JXTA using between 1.58 and 2.29 times more bandwidth. Also note that in this particular experiment, JXTA was operating in edge mode, where it does not use a DHT. In rendezvous mode with a DHT, JXTA used significantly more bandwidth.

DisService Opportunistic Listening Results

This second experiment demonstrates opportunistic listening, one of the unique features of DisService. Opportunistic listening exploits peer nodes as surrogate listeners to capture information that a subset of the peer nodes are trying to receive. The scenario is shown in Fig. 4, and consists of a set of 10 ground nodes of which two nodes are trying to receive data being multicast by a UAV. The UAV makes one pass through the area as it transmits messages to the ground nodes. The bandwidth of the airborne network between the UAV and the ground nodes is assumed to be 230 kb/s, and the bandwidth of the P2P network between the ground nodes is assumed to be 1 Mb/s. In the baseline case, we report on the percentage of messages received via UDP multicast by the two interested ground nodes. In the DisService case, the ground nodes exchange missing data between themselves to maximize the number of messages correctly received by the two ground nodes.

The results, with varying message sizes and network reliability, are shown in Table 2 and illustrate the benefit of opportunistic listening. For example, with a UAV link reliability of 70–80 percent, only 20 percent of messages of size 7 kbytes are successfully delivered due to IP fragments being lost. With DisService, the target nodes are able to obtain these missing fragments from other peers and reassemble the messages intact, thereby improving performance significantly. Even with very poor UAV link reliability of 30–50 percent, DisService is able to receive a

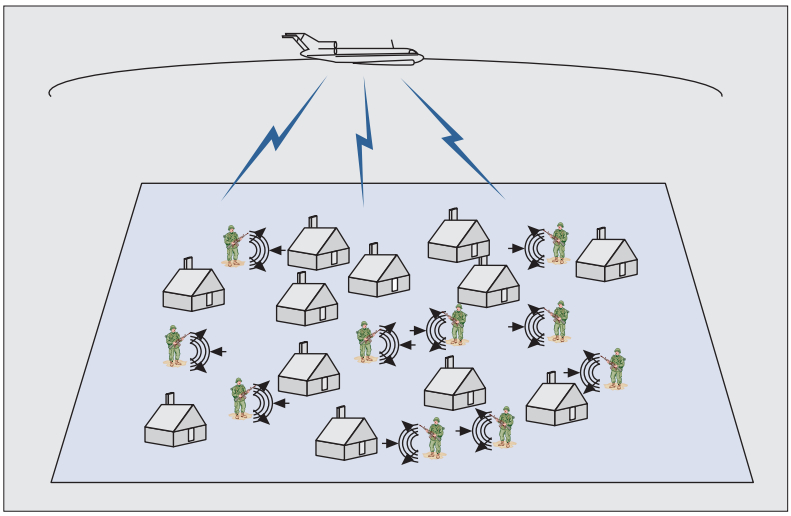

Figure 4. *Opportunistic listening experiment scenario.*

Number of nodes		Group Manager		JXTA
		(Default)	(With pings)	
UAVs	GVs	Bandwidth utilization (bytes/s)		
2	1	83	372	645
2	2	164	672	1106
2	3	247	866	1536
2	4	329	1217	2012
4	6	828	1856	2963
4	10	1350	3061	4849
4	14	1146	2937	6713
4	18	1451	5471	8658

Table 1. *Group Manager performance (bandwidth) for peer discovery.*

UAV link reliability	Ground net reliability	Msg size	Msgs sent	Msgs received		Success rate		DisService P2P overhead (bytes/Msg)
		(bytes)		Multicast	DisService	Multicast	DisService	
70% to 80%	80%	1024	1390	820.6	1357.2	59.06%	97.68%	101
		7168	203	31.9	202.5	15.75%	100.00%	44,518
		15,360	96	4.5	96.0	4.69%	100.00%	69,938
		35,840	41	0.0	41.0	0.00%	100.00%	174,163
30% to 50%	80%	1024	1377	590.6	1367.3	42.89%	99.30%	14,380
		7168	200	2.0	193.0	1.00%	96.50%	62,139
		15,360	91	0.0	89.3	0.00%	98.13%	135,077
		35,840	41	0.0	34.0	0.00%	82.93%	368,985

Table 2. *DisService opportunistic listening performance.*

large percentage of the messages. We also report on the P2P traffic exchanged between the ground nodes as they reconstruct the messages, which could be regarded as overhead. As the results show, the overhead per message is quite high, especially as the size of the messages increases and the reliability decreases. However, this communication takes place over the P2P ground network, which has higher capacity and is less congested. Furthermore, we expect to optimize this bandwidth utilization in the near future.

Conclusions

Our experiences have led us to study how to realize robust and efficient applications and services for tactical edge networks. We have observed that the nature and characteristics of tactical networks, combined with the requirements of network-centric warfare present a compelling case for P2P systems. We then distilled these experiences into a set of technical requirements that must be addressed to develop effective P2P systems for tactical environments. Following these guidelines, we have developed the Group Manager and DisService components of the Agile Computing Middleware, which realize P2P discovery and information dissemination. We have tested those components in the context of several experiments and exercises that allowed us to verify the soundness as well as the effectiveness of the P2P approach. We continue to enhance the Agile Computing Middleware and conduct further experiments. Future work includes optimizing DisService, comparing with reliable multicast approaches, as well as comparing opportunistic listening with forward error correction. We hope that this article will motivate and guide the development of future P2P systems for tactical networks.

Acknowledgments

This research was sponsored in part by the U.S. Army Research Laboratory under Cooperative Agreement W911NF-04-2-0013, by the U.S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0009, by the Air Force Research Laboratory under Cooperative Agreement FA8750-06-2-0064, and by the Office of Naval Research under grant N000140910012. Figure 1 is courtesy of William Howell, Florida Institute for Human and Machine Cognition.

References

[1] X. Shen *et al.*, *Handbook of Peer-to-Peer Networking*, Springer, 2009.
[2] N. Suri *et al.*, "Communications Middleware for Tactical Environments: Observations, Experiences, and Lessons Learned," *IEEE Commun. Mag.*, vol. 47, no. 10, Oct. 2009, pp. 56–63.
[3] B. Traversat *et al.*, "Project JXTA 2.0 Super-Peer Virtual Network," Sun Microsystems, Inc. tech. rep., 2003; http://research.sun.com/spotlight/misc/jxta.pdf
[4] IETF Zeroconf Working Group; http://www.zeroconf.org
[5] XMPP Standards Foundation; http://xmpp.org/
[6] J. Rosenberg *et al.*, "SIP: Session Initiation Protocol," RFC 3261, June 2002;
http://www.rfc-editor.org/rfc/rfc3261.txt
[7] The UPnP Forum; http://www.upnp.org
[8] K. Lund *et al.*, "Using Web Services to Realize Service Oriented Architecture in Military Communication Networks," *IEEE Commun. Mag.*, vol. 45, no. 10, Oct. 2007, pp. 47–53.
[9] D. Galatopoullos, D. Kalofonos, and E. Manolakos, "A P2P SOA Enabling Group Collaboration through Service Composition," *Proc. ICPS '08*, Sorrento, Italy, July 6–10, 2008.
[10] N. Suri, "Dynamic Service-Oriented Architectures for Tactical Edge Networks," *Proc. 4th Wksp. Emerging Web Services Tech. '09*, Nov. 2009, pp. 3–10.
[11] N. Suri *et al.*, "An Adaptive and Efficient Peer-to-Peer Service-Oriented Architecture for MANET Environments with Agile Computing," *Proc. 2nd IEEE ACNM '08*, Salvador de Bahia, Brazil, Apr. 7–11, 2008, pp. 364–71.
[12] Mobile Ad Hoc Network Emulator (MANE);
http://cs.itd.nrl.navy.mil/work/mane/index.php

Biographies

Niranjan Suri [M] (nsuri@ihmc.us) is a research scientist at the Florida Institute for Human and Machine Cognition (IHMC). He received his Ph.D. in computer science from Lancaster University, England, and his M.Sc. and B.Sc. in computer science from the University of West Florida, Pensacola. His current research activity is focused on the notion of agile computing, which supports the opportunistic discovery and exploitation of resources in highly dynamic networked environments. His other research interests include coordination algorithms, distributed systems, net-

working, communication protocols, virtual machines, and software agents. He has authored or co-authored over 60 papers, has been on the technical program committees of several international conferences, and has been a reviewer for the National Science Foundation as well as several international journals.

GIACOMO BENINCASA (gbenincasa@ihmc.us) is a research associate at the Florida Institute of Human and Machine Cognition, working on information dissemination in tactical networks and on biologically-inspired security infrastructures. He holds a B.Sc. and an M.Sc. in computer engineering from the University of Modena and Reggio Emilia, Italy.

MAURO TORTONESI (mauro.tortonesi@unife.it) received his Laurea degree in electronic engineering and his Ph.D. in computer science engineering from the University of Ferrara, Italy. From 2004 to 2005 he was a research associate at the Florida Institute of Human and Machine Cognition. Currently, he is an assistant professor in the Engineering Department of the University of Ferrara, Italy. His research interests include distributed and mobile computing, QoS management, business-driven IT management, and industrial automation.

CESARE STEFANELLI [M] (cesare.stefanelli@unife.it) received his Laurea degree in electronic engineering and his Ph.D. in computer science engineering from the University of Bologna, Italy. He is currently a professor of computer science engineering in the Engineering Department of the University of Ferrara. His research interests include distributed and mobile computing, adaptive and distributed multimedia systems, network and systems management, and network security. He is a member of AICA.

JESSE KOVACH (jkovach@arl.army.mil) is a computer engineer with the Battlefield Information Processing Branch of the U.S. Army Research Laboratory, specializing in information dissemination in tactical networks, and sensor and robotic systems integration. He designs, develops, and tests prototype systems that incorporate advanced concepts developed within the laboratory as well as elsewhere in the defense R&D community. His work has been incorporated into a number of fielded systems and has been used in multiple high-visibility demonstrations, including C4ISR On the Move and Empire Challenge. He holds a Bachelor's degree in computer engineering from the University of Maryland.

ROBERT WINKLER (winkler@arl.army.mil) has been a computer engineer with the U.S. Army Research Laboratories since 1987. He graduated magna cum laude in computer science from the University of Maryland at College Park in 1988 and received his Master of Science degree in computer science from Johns Hopkins University in 1995. His current applied research interests include databases, artificial intelligence, and data mining.

RALPH KOHLER, JR. [F] (Ralph.Kohler@rl.af.mil) is a principal engineer at the U.S. Air Force Research Laboratory (AFRL), where his primary research interests are the nexus of information management, wireless and tactical networks, and the capabilities that combining the two make possible. He received his Master's degree from Syracuse University and a member of the AIAA.

JAMES HANNA (james.hanna@rl.af.mil) is a senior research engineer with the Enterprise Information Management Branch in the Information Directorate at AFRL. He has been a computer engineer for AFRL for the past 21 years and has over 23 years experience developing software applications. For the past two years he has led the development of Phoenix, an SOA-based information management infrastructure. He is also the lead in-house engineer researching issues related to distributed survivable information management and advanced architectures to address the complex challenges of deploying information management in tactical environments. He has a B.Sc. in computer science from the State University of New York Institute of Technology, Utica.

LOUIS POCHET (loupochet@gmail.com) is the senior integration engineer in the Combat Support Office of the U.S Air Force Reserve and of a small technology startup, GMECI. Before entering the Reserves, he developed applications, protocols, and waveforms for the U.S. Department of Defense as an active duty research engineer, deploying to theaters of operation multiple times. He has focused on network discovery and configuration over existing and experimental tactical IP networks, including QNT and ANW2, for the last five years.

SCOTT WATSON (scott.c.watson@navy.mil) is a senior systems architect with the Composeable Services Branch in the Command and Control Department of the Space and Naval Warfare Systems Center — Pacific. He has over 10 years' experience designing software systems and applications. He is currently the principal investigator for the Mobile Modular Command and Control for Company level operations (M2C3) project. He has a B.S. in computer science from San Diego State University, California. He recently received the Navy Civilian Meritorious Service award for work in the area of disruption-tolerant networking.