The DARPA WNaN Network Architecture

Jason Redi
Raytheon BBN Technologies
Cambridge, MA
redi@bbn.com

Ram Ramanathan
Raytheon BBN Technologies
Cambridge, MA
ramanath@bbn.com

Abstract— **The warfighter network of the future needs to be low-cost, instantly deployable, self-organizing, robust, and scale with both size and density. DARPA's Wireless Network after Next (WNaN) meets these challenges using an architecture that combines several innovative features for the first time in a real functional system: dynamic spectrum access (DSA), adaptive multi-transceiver frequency assignment, multi-channel medium access, highly scalable routing and disruption-tolerant networking (DTN). The WNaN system has been successfully demonstrated in real-world military experiments for up to 100 nodes - the largest military MANET demonstration on record. We present the network architecture of the WNaN system, focusing mainly on the medium access and sub-network layers. We briefly describe how key protocols work in tandem, and how they are implemented on the objective platform. Finally, we discuss some future plans.**

I. INTRODUCTION

The wireless communications network for tomorrow's warfighter must surmount a number of challenges: it must provide high capacity to support an increasing number of diverse traffic flows; it must be instantly deployable with minimal manual configuration, including spectrum assignment; it must be highly tolerant of mobility, disruption and disconnection; and it needs to be based on affordable hardware.

DARPA's Wireless Network after Next (WNaN) meets these challenges using an innovative architecture and protocol suite. WNaN seeks to achieve an affordable, high scalable, robust, and self-organizing multi-hop wireless network to meet current and future military needs. A number of new technologies, such as policy-conformant dynamic spectrum access, distributed, dynamic frequency assignment, and disruption tolerant routing – discussed in detail later – enable WNaN to achieve its goal.

WNaN uses networking software from Raytheon BBN Technologies, and radio hardware plus firmware from Cobham. The WNaN radio platform is an affordable purpose-built military networking radio based on commercial parts, lines and processes. It is designed and built to reduce RF front-end cost drivers and enable network adaptability. The radio platform has four transceiver chains, each of which provides wideband operation from 900 MHz to 6 GHz, with adaptable data rates from 90 kbps to 2 Mbps, and output power of 1W from each transceiver.

The WNaN MAC layer leverages the multiple transceivers using an agile multi-frequency CSMA/CA protocol that includes support for dynamic spectrum access. The WNaN network layer supports efficient unicast and multicast routing that can adapt between connected and disconnected routing modes. In particular, WNaN can route over non-contemporaneous paths (e.g. partition and healing, data mules, etc.) using disruption tolerant networking (DTN) technologies.

Multi-hop wireless networks, also known as Mobile Ad Hoc Networks (MANET) have been an active area of research for several decades. Medium access techniques have been developed for both commercial (e.g. [1]) and military (e.g. [2]) MANETs. A wide variety of routing schemes have been developed [3][4][5][6][7], some proposed for standardization within the IETF [4][5]. More recently, techniques for dynamic spectrum access [8] and disruption tolerant networking (DTN)[9][10] have been proposed in the literature, with some work on multi-transceiver multi-frequency dynamic networks [11]. However, WNaN is the first network that integrates these advances into a single system with demonstrable performance in a sizeable network. It represents a unique accomplishment of bringing to maturity diverse ideas, architecture to tie these together and a significant engineering feat to implement and demonstrate.

WNaN has been demonstrated to various DoD audiences a number of times over the past three years including demonstrations of 52 moving handheld nodes in June 2010 and 100 nodes in September and October 2010.

In this paper, we present the architecture of the medium access and network layers of WNaN. We discuss the modules and their functionality, and interaction. We briefly present some of the key mechanisms behind WNaN, notably access control, frequency assignment, dynamic spectrum access, scalable routing and disruption-tolerant routing. Following that, we overview our implementation approach and close with some conclusions.

II. ARCHITECTURE

A. *Architecture Overview*

Figure 1 illustrates the components comprising the network stack in each radio. Each major box (Media Access Control (MAC), Routing, Bundle Protocol Agent (BPA), etc.) is a single protocol module that has a specific purpose and a set of interfaces to other modules. The inner boxes inside each major box describe particular key functions of the module. In most cases the inner boxes have a direct correspondence to a software thread implementing that functionality.

This work is sponsored by DARPA through the Air Force Research Laboratory (AFRL) Contract FA8750-07-C-0169. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

978-1-4673-0081-0/11/$26.00 ©2011 IEEE

The Medium Access Control (MAC) module performs neighbor discovery using periodic *heartbeats* across frequencies and creates *link profiles* for use by the Routing module. A link profile is a parameterized description (frequency, datarate, FEC, power) of how a neighboring node can be reached. The MAC assigns frequencies to the multiple transceivers and arbitrates access to the channel using a variation of Carrier Sense Multiple Access with optional Collision Avoidance (CSMA/CA). The MAC supports spectrum sensing by enforcing silence during sensing periods. Each transceiver is controlled by a separate *access worker* thread, and the group of worker threads is coordinated by the *access controller* thread. Coordination between the individual transceivers is necessary for functionality such as coordinated spectrum sensing and MIMO operations.

The Dynamic Spectrum Access (DSA) module is responsible for using spectrum policy reasoning for determining which frequencies might be allowable for use, requesting sensing periods on those frequencies, classifying the results from sensing events, then providing the list of allowed frequencies to the MAC for use in frequency assignment. The DSA module continually looks for non-cooperatives (i.e., non WNaN nodes, possibly primary/legacy users of the spectrum) on existing spectrum, reevaluates policy based on time and geography, and updates the list of allowed frequencies to the MAC when the list changes. The MAC itself is responsible for determining frequency assignments, evacuating frequencies and performing reconstitution of the network when a frequency is removed from use. The MAC transmits and receives packets through the Cobham PHY API interface.

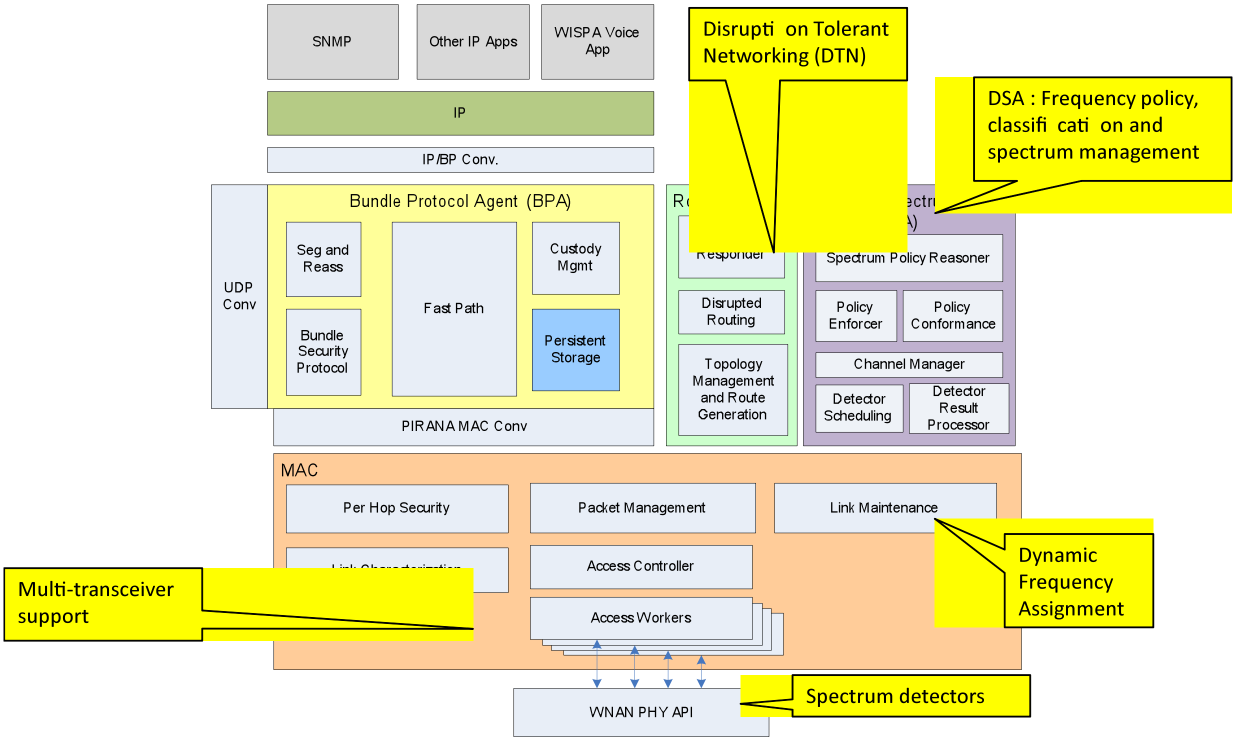

Figure 1: A snapshot of the WNaN Network Architecture

The Routing module contains the capability to route across both connected networks using traditional MANET-style routing, as well as across disconnected networks using DTN-style routing. The connected routing employs *hazy sighted link state* [7] and multipoint relays (MPRs) using the link profiles provide by the MAC. Disconnected routing is based on *prioritized epidemic* routing [12]. WNaN also provides the ability to seamlessly switch between the modes using a novel *Endemic routing* protocol.

The Bundle Protocol Agent (BPA) provides forwarding and storage services at the command of Routing. IP packets are encapsulated into bundles and provided to the BPA at the source nodes. Upon receipt of a packet from IP, the BPA queries the Route Responder thread of the Routing module for the action (send to node, store, drop) and performs the appropriate function. The BPA's functions are similar to the forwarding function in a typical MANET, with two exceptions: it can store a bundle for later transmission, and it operates on bundles as per the DTN standard [15], rather than as IP packets. This does not mean that every IP packet gets DTN service (and potentially stored), rather that each IP packet *could* get DTN service if appropriate for that traffic type.

As the figure shows, the WNaN network operates below the IP stack. WNaN presents itself as an Ethernet device to the IP stack, and provides multihop services to packets without interacting with IP except at the source and destination end-points. This architecture can be considered to be similar to the way commercial switched Ethernet devices may forward packets across many devices, though the IP stack itself may see all nodes as operating on the same subnetwork segment. The advantages of this architecture are that applications written for

IP do not have to be modified, so we do not expose potentially rapidly changing multihop routes to the IP stack and external routers, and we are free to enhance the per-hop forwarding mechanisms with capabilities such as just-in-time forwarding (described later), and temporary storage and in-network caching for DTN services.

In the next few sections we select and present a few key technologies, focusing on those that advance the state of art and at the same time provide the features that make WNaN uniquely suitable for the next generation war fighter network.

B. Dynamic Spectrum Access

The WNaN system can operate in the presence of non-cooperative interferers by sensing and avoiding the interferers before they impact the user applications. WNaN conforms to regulatory policy in the selection and use of these frequencies. Our architecture contains a way to store policy, and a "policy conformance reasoner" that interprets policy, evidence from sensing, and whether a frequency can be used. The policy reasoner uses OWL to describe its ontology so that multiple policies can be loaded into WNaN and the radio can automatically combine and prioritize across policies to determine what frequencies might be allowed for use within conformance with multiple policies. The ontology WNaN uses is the same as that designed for the original DARPA XG program and the custom embedded reasoner which runs on the platform was implemented by Shared Spectrum Corporation (SSC) [16].

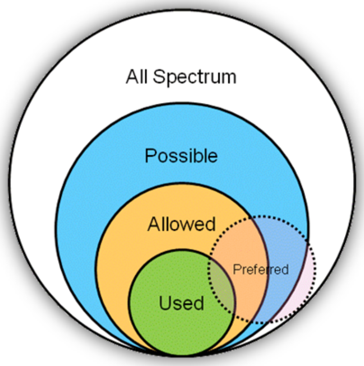

Figure 2: Visual representation of relationship between frequency sets in WNaN

The DSA module broadly classifies channels into *possible, allowed* which is a subset of possible*,* and *used*, which is a subset of allowed. A visual depiction of the relationships between these frequency sets is shown in Figure 2. The possible set consists of the "master" set, that is, any frequency that may conceivably be used at a particular location and time based on external policy. The allowed set consists of those which are currently being sensed and are deemed usable after detection and classification for non-cooperatives. The used channels are those that are actually been assigned for use by MAC. Sensing is continually performed on the current allowed list, as well as at least a portion of the possible list in order to find alternative frequencies in case one of the frequencies in the allowed list is removed due to a dynamic policy change or appearance of a non-cooperative.

A key part of DSA is sensing for non-cooperatives. DSA makes sensing time requests to the MAC, which then schedules the sensing time. For each sensing interval, DSA provides the list of channels to sense and the detectors to be used. The MAC then communicates with other nodes to clear the frequencies that will be sensed within a certain radius and engages the sensors. The sensing results are then passed up to DSA for classification and policy review.

As a result of continuous execution of the above, the DSA maintains a dynamic list of *allowed channels*, which is made available to the MAC. In particular, when a previously allowed channel becomes un-allowed due to a non-cooperative, the MAC is immediately informed so that necessary evacuation and re-constitution actions can be taken as per policy.

C. Medium Access

The WNaN medium access protocol is based on Carrier Sense Multiple Access with optional Collision Avoidance (CSMA/CA), similar to the approach behind the IEEE 802.11 standard [1]. We chose CSMA/CA over TDMA due to its advantages in several domains: better support for heterogeneous and bursty traffic; much better tolerance to mobility; simplicity; well-understood dynamics due to intensive study of 802.11 in the literature; and robustness. Further, while TDMA does have an advantage in terms of efficiency under high loads, our architecture mitigates this by virtue of having four transceivers and dynamic frequency assignment, which can provide ample capacity and eliminate many hidden terminal issues. Our protocol is similar to 802.11, with some differences, such as a traffic-based adaptive backoff scheme that results in better performance.

There is one Access sub-module (thread) per transceiver, within the Access module. The sub-modules have identical functionality except they service different transceivers. The "parent" Access module assigns the channels for each transceiver/thread based on the results of frequency assignment. Each Access thread interacts directly with the Cobham PHY API which provides transmit, receive, transceiver tuning and sensing capabilities.

The WNaN access scheme supports an over the air QoS differentiation scheme (similar to 802.11e [13]), which allows higher priority traffic to gain quicker access to the channel. We maintain several channel access priorities, each with a separate minimum and maximum contention window. Specifically, the higher priority of the packet, the lower the contention window. The backoff mechanism for each priority is identical -- we increase the backoff after each unsuccessful transmit for unicast or if the channel is busy for broadcast. We note that this enforces priority *across* nodes as well, for example if there are two interfering flows in the vicinity of each other but not going through the same node, this scheme will give preferential access to the higher priority flow.

A key part of our medium access scheme is the support for sensing. Recall that in order to properly perform sensing, all nodes in the interference range of a node should be silent for the duration of sensing. In order to enforce silence, we use a new control message called *Defer to Sense (DTS),* whose purpose is to alert neighboring nodes of an imminent sensing operation. The DTS is broadcast on all transceivers currently assigned a frequency. By virtue of using a common "blanket" frequency (see section II.D) DTS's reach all neighbors. Multiple DTS's may be sent before each sensing interval to provide redundancy. A DTS contains the relative time after which sensing is performed as well as the duration of sensing. A receiving node factors out modem and propagation delays to compute its absolute sensing time. No synchronized time such as GPS is required. In order to maximize utilization, nodes coordinate their sensing, that is, a node chooses, to the extent possible, a sensing interval that coincides with those of neighbors'. We note that this is not a tight constraint, and some nodes may have multiple silence periods. This architecture is more flexible and robust than a scheme where all nodes are synchronized to a single time, especially in the face of mobility.

D. Dynamic, Distributed Frequency Assignment

Unlike a conventional single-radio single-channel network, WNaN can support multiple simultaneous communications in a neighborhood and on the same node, greatly increasing the capacity. This is achieved using a dynamic, distributed algorithm to assign allowed frequencies to each of the transceivers in a node. The frequencies are re-assigned as the topology changes, such as when nodes move, or if the allowed list changes or if a frequency needs to be evacuated due to the appearance of a non-cooperative signal.

When a node is activated, it first attempts to discover other nodes and their frequency assignments. It does so by "soft" assigning one of the transceivers to a randomly chosen frequency and sending *probes*. In parallel, it scans other allowed frequencies for neighbors. After it discovers its neighbors and learns their assignments via their Heartbeats, it converts the assignment into a "hard" assignment using the algorithm described below. In general, in "steady state", nodes use neighbor assignments learned via Heartbeats along with their list of allowed channels to continually monitor and update the assignment.

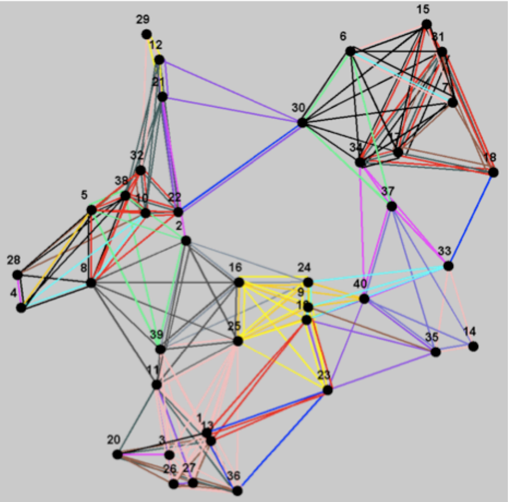

Figure 3: Example "clique" frequency assignment for a 40 node WNaN network with 4 transceivers and 18 channels (best viewed in color).

The core of frequency assignment is a heuristic for the choice of frequencies for each transceiver. Note that different sets of neighbors may be reachable on different frequencies. In general, more commonality in frequencies is better for connectivity but reduces spatial reuse. Thus, a key challenge in frequency assignment is to balance the connectivity (physical topology) and spatial reuse. The WNaN architecture decouples this tradeoff, that is, the frequency selection heuristic, from the actual mechanics of the distributed algorithm so that various schemes can be used.

One such heuristic is termed *clique assignment*, which has two parts: determining which channels are *eligible* for assignment, and an ordering of eligible channels based on *benefit*. A channel *c* is deemed ineligible if one of the other transceivers has assigned it, or such an assignment would result in a three hop path X-Y-Z, where this node is one of X, Y, or Z, and all three nodes have *c* as one of their transceiver assignments. The "benefit" of a channel in this scheme is the number of neighbors already assigned to this channel. This scheme gives rise to cliques of nodes on each frequency, and hence the name. Figure 3 shows the assignment and links produced by WNaN on an example network. One of the key capabilities that is provided by the clique heuristic is that it eliminates hidden terminals since there are no situations where a node two hops away is sharing the same frequency. This dramatically improves throughput by eliminating collisions and eliminating the need for collision avoidance packets (RTS/CTS).

When using clique assignment only, due to a limited number of frequencies and limited number of transceivers, some cases can arise where the network gets partitioned. To avoid this, we have added a dynamic "blanket" assignment on one transceiver while other transceivers run cliques. The blanket assignment attempts to find a common frequency across all neighbors to use, and by extension the whole network. This is not a dedicated frequency – rather, the network dynamically finds commonality and uses it if available. The system still works if there is no common frequency, connecting the network with clique transceivers only where no common frequency exists. As part of the WNaN program, we have been investigating a number of additional frequency assignment algorithms, each with different tradeoffs.

E. Scalable Routing

The function of WNaN Routing is to route and forward bundles from a given source to one or more destinations, even if a destination is temporarily disconnected for an extended period of time. Specifically, WNaN routing is able to deliver packets over *non-contemporaneous* paths in eventually connected or eventually transportable networks [18]. Other objectives for WNaN routing include scalability with respect to network size as well as density, support for mobility, and limited complexity.

WNaN uses destination-based hop-by-hop routing. The route a bundle follows through the network is a function of its final destination. Each mid-hop node determines the bundle's next hop(s) as a function of its final destination, based on the information that mid-hop node has.

For routing in connected regions, WNaN uses *proactive routing* [14], in particular a form of *link state routing*. We have chosen proactive routing over reactive routing to minimize network delay. To improve scalability, WNaN implements two enhancements to traditional link-state routing: Hazy-Sighted Link State (HSLS) and Multi-Point Relaying (MPR). HSLS provides an algorithm for *limited dissemination* of Link State Updates (LSUs), which restricts the hop-distance that topology updates travel in the network [7]. The basic idea of HSLS is that nodes far away from link changes do not need to know about all changes in order to make good routing decisions. MPR is a technique for *efficient dissemination* of LSUs [5]. MPRs limit the number of nodes that broadcast the link state information, rather than allowing every node to rebroadcast.

To further reduce the amount of overhead in the network, WNaN employs a number of algorithms typically applied for traffic shaping to the link state information. In particular, we limit the number of LSUs that are transmitted with hold down timers and we aggregate multiple LSUs together in a single packet to reduce per packet channel access costs.

For routing in disconnected or disrupted regions (i.e, when seeking non-contemporaneous paths by storing), WNaN implements an event-based *Epidemic* protocol [17], which is a reactive routing protocol particular to Disruption Tolerant Networks (DTNs). Figure 4 illustrates the difference between traditional IP-style routing and DTN-style routing. Epidemic uses a handshake between each set of neighboring nodes to synchronize the set of bundles stored between them. The handshake, or sync messages, is used to pass only bundles that are new to a neighbor, thereby reducing unnecessary bundle transmissions. Epidemic protocol has been proven to deliver the packet as long as there exists some path over time [18]. Although epidemic protocol is overhead-intensive, we expect the WNaN network to be only occasionally disconnected and only in some regions, and hence this has minimal impact.

To adapt between connected and disconnected routing over time and topology, WNaN includes the *Endemic* routing protocol, which synergistically combines link-state and Epidemic routing. Packets can start out traversing the network epidemically, but when a route is found the packet can be converted to being disseminated with traditional connected MANET techniques. A packet can switch routing techniques from connected, to disrupted and back to connected multiple times on its way to the destination. Endemic incorporates techniques for preventing routing loops that are inherent to naïve hybrid approaches. As a bonus side-effect, Endemic reduces the network-wide overhead of using Epidemic routing in disconnected portions of the network. Bundles that exceed their expiry time are dropped. Nodes stop accepting bundles from neighbors after exceeding a storage threshold.

WNaN routing supports multicasting. Specifically, multicast dissemination trees are constructed using the link-state topology database and packets are sent along the trees using hop-by-hop forwarding. At each hop or branch-point, packets are either sent broadcast or unicast depending upon the number of downstream neighbors. WNaN provides two multicast algorithms – source based trees, and a mesh-based “me trees” algorithm. Source based trees require each node to compute the multicast spanning tree from each potential source to the set of multicast destinations for each group, and then to determine the node’s own forwarding role with the tree. The novel mechanism called *me-trees* leverages the unicast tree calculations to provide mesh-style multicast capabilities where nodes dynamically determine their role in the forwarding tree based on the previous hop destination assignments. Each multicast algorithm has different tradeoffs appropriate for different kinds of networks and groups – source based trees are optimal in terms of the number of packet transmissions in the network, while the me-trees reduce the computation required for multicast to almost nothing at a slight cost of additional header overhead. The efficient support of multicast is a great asset in military communications which is overwhelmingly multicast.

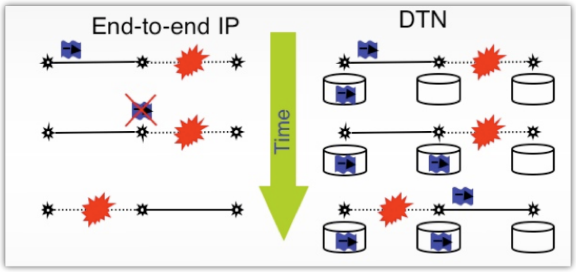

Figure 4: IP-style routing drops packet when there is no end-to-end path. DTN routing makes maximal progress and stores till further progress can be made.

III. IMPLEMENTATION

The version 2 WNaN hardware platform includes a pair of TI OMAP processors and a shared memory architecture managed by an FPGA. Most modem, detection and RF control is also in the FPGA. OMAP processors are common in today’s commercial hand-held devices including many Android phones. One of the OMAP processors runs a custom distribution of Linux, while the other OMAP runs the Green Hills Integrity Real Time Operating system. The MAC threads have hard real-time constraints so they run within GHS Integrity, while the other processes run on the Linux processor. The BPA itself runs as a loadable kernel module within Linux so that we can minimize any delays of the subnetwork forwarding at each hop. While the routing protocols run within application space, this shouldn’t be construed that the actual network runs as an overlay to IP. The network presents itself as a subnetwork beneath IP, similar to how switched Ethernet operates below IP and allows many hops of communication even though IP may see only one IP hop between nodes.

Applications can be hosted on either the WNaN radio itself, or on attached devices such as PDAs which are then attached via USB. Applications can indicate what kind of service they would like to receive from the WNaN network through setting of the IP DSCP bits. WNaN uses four bits of the private address space of DSCP to indicate whether the packet should receive DTN service or not, and what the priority (0-7) is of the packet. Strict priority queues are maintained through the BPA and the MAC. The priority bits also cause the MAC to change its backoff algorithm to be more or less aggressive depending on the importance of the packet, as well as change the number of retransmission and reliability attempts.

Due to limited availability of early hardware, limited time for extended experiments, and costs of manpower for running experiments, the use of simulation in performing network design and analysis is critical for tactical communication systems. Simulations cannot replace real-world experiments, but they can provide a ground for understanding complex network behaviors before going to the field. In many systems, simulation source code and code written for the objective platform are separate and become more divergent over time. Bugs and fixes in one code base may not get applied to the other code base, and over time, the actual networking behavior modeled may be different than the network behavior in the field.

WNaN eliminates this problem by implementing all the protocols in a software framework called Portable Link Framework (PLF). This framework provides a process and threading model based on a prioritized preemptive real time operating system. Inter-thread and inter-process communication primitives are provided, as well as configuration and data collection primitives, memory management, string manipulation, random number generation, and other critical operating system functions which cannot be guaranteed to be the same across various real-time, non-real-time and simulation platforms. All the code on WNaN has been written to this framework and it has allowed us to run the actual protocol code that runs on the platform within the simulation. When unexpected behaviors appear in the field, we are often able to reproduce the scenario and therefore the behavior in simulation, then make changes and apply the changes directly to the platform.

Understanding network performance by collecting data from live experiments is crucial for being able to validate and improve the system. A number of existing tactical radios unfortunately cannot display the live network nor collect significant statistics about behavior to track performance. As part of the Portable Link Framework (PLF), WNaN includes a low overhead approach for collecting statistics from each of the modules (MAC, routing, DSA, BPA). A low priority thread is responsible for directly reading memory locations from other threads to collect data such as the number of packets sent per each transceiver per priority type, the current frequency assignment of each transceiver, etc. There are over 300 statistics available within the system. For our field tests, we collected 75 statistics every 5 seconds from across the system in order to validate WNaN performance over the life of the experiment over many days in the field. Additionally, through low overhead collection of network information any node in the network can be attached to a display and visually see the network and frequency assignments overlaid over Google Earth images. This network visualization was a key capability for leaders' situational awareness of the network during experiments.

IV. SUMMARY

We described the network- and MAC-layer architecture of, and the key mechanisms in the DARPA Wireless Network after Next (WNaN) system. WNaN incorporates state-of-the-art concepts such as dynamic spectrum access and disruption tolerant routing, and utilizes innovative mechanisms such as dynamic, distributed frequency assignment to increase capacity. Along with simple, robust and well understood MAC and network-layer approaches such as CSMA/CA and link-state routing, the WNaN architecture combines sound foundations with innovative technologies to provide a highly effective solution for the next generation war fighter. We have proven this effectiveness using several demonstrations, including 50 and 100 node mobile experiments with military relevance and participants.

A number of improvements to the system are now being developed as a next generation, including improved support for variable rates, energy-saving mechanisms, Type II security, semi-scheduled access for reservations and performance, and multi-path routing.

ACKNOWLEDGMENT

We thank the current DARPA Program Manager Dr. Bruce Fette, and the previous Program Manager Dr. Preston Marshall who provided leadership and guidance for the WNaN program. We would also like to thank the team at Cobham for building an innovative hardware and firmware platform which enables a next generation of tactical networking capabilities.

REFERENCES

[1] M.S. Gast, “802.11 Wireless Networks: The Definitive Guide", O'Reilly and Associates, April 2002.

[2] C.D. Young. USAP multiple access: dynamic resource allocation for mobile multihop multichannel wireless networking. IEEE MILCOM 1999..

[3] D. B. Johnson and D. Maltz, “Dynamic Source Routing in Ad Hoc Wireless Networks", In Mobile Computing, edited by Tomasz Imielinski and Hank Korth, Kluwer Academic Publishers, 1995.

[4] C. Perkins and E. M. Royer, ``Ad-Hoc On-Demand Distance Vector Routing," Proceedings of IEEE WMCSA'99, New Orleans, LA, Feb. 1999, pp. 90-100.

[5] T. Clausen et al “Optimized Link State Routing”, IETF RFC 3626, 2003.

[6] Y. B. Ko and N. H. Vaidya. ``Location-Aided Routing (LAR) in Mobile Ad Hoc Networks," Proc ACM/IEEE MOBICOM, Dallas, TX,October 25-30, 1998.

[7] C. Santivanez, S. Ramanathan, I. Stavrakakis, “Making Link State Routing Scale for Ad Hoc Networks," Proceedings of MobiHoc 2001, Long Beach, CA, Oct. 2001.

[8] I. Akyildiz, W. Lee, M. Vuran, and S. Mohanty, “Next Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: a Survey,” Computer Networks, 50(13):2127–2159, 2006.

[9] Sushant Jain, Kevin Fall, and Rabin Patra. Routing in a delay tolerant network. In SIGCOMM, Aug 2004.

[10] Z. Zhang. Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: Overview and challenges. IEEE Communication Surveys and Tutorials, Jan 2006.

[11] P. Kyasanur and N. H. Vaidya, “Capacity of multi-channel wire- less networks: Impact of number of channels and interfaces,” in Proc. of ACM MobiCom, August 2005, pp. 43–57.

[12] R. Ramanathan, R. Hansen, P. Basu, R. Rosales-Hain, and R. Krishnan. 2007. Prioritized epidemic routing for opportunistic networks. In *Proceedings of the 1st international MobiSys workshop on Mobile opportunistic networking* (MobiOpp '07)

[13] S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor, “IEEE 802.11e Wireless LAN for Quality of Service,” Proc. European Wireless, 2002

[14] E. Royer and C.-K. Toh, “A Review of Current Routing Protocols for Ad Hoc Wireless Networks,” IEEE Pers. Commun., Apr. 1999, pp. 46–55.

[15] K. Scott, S. Burleigh, “Bundle Protocol Specification”, IETF RFC 5050

[16] F. Perich, “Policy based Network Management for NeXt Generation Spectrum Access Control,” IEEE DySpan, April 2007

[17] A. Demers, et al, “Epidemic Algorithms for Replicated Database Maintenance,” Proc. ACM PODC, 1987

[18] R. Ramanathan, P. Basu, R. Krishnan, "Towards a Formalism for Routing in Challenged Networks," Proc. CHANTS 2007, Montreal, September,2007