Demonstrating Rapidly Deployable Broadband Wireless Communications for Emergency Management

Charles W. Bostian (bostian@vt.edu)
Scott F. Midkiff (midkiff@vt.edu)

Center for Wireless Telecommunications
340 Whittemore Hall (0111)
Virginia Tech
Blacksburg, VA 24061

Telephone: 540-231-5096
Fax: 540-231-3005
URL: http://www.cwt.vt.edu/

Presented at the
National Digital Government Research Conference (dg.o 2002)
May 20-22, 2002
Redondo Beach, California

1. Introduction

In 1999 Virginia Tech and SAIC proposed a project to meet a hypothetical need that became all too real on 9/11/2001. In situations like 9/11, communications are destroyed in a limited geographical area. Cell phones may provide limited voice communications, but their networks quickly become overloaded and fail. Surviving fiber and copper connections are available at the perimeter, but there is no easy way for first responders and on-scene commanders to connect to these. At the Pentagon, for example, it took a week for the local fire department to establish desperately needed Internet access. Workers at both sites ultimately set up a number of ad hoc point-to-point systems with good results.

Our proposal described a four-year project to establish a rapidly deployable test bed for high-bandwidth "end-to-end" communications in support of comprehensive emergency management (CEM) planning and operations. At that time there was little private sector interest in such a system and commercial radio equipment to support data rates at DS-3 rates (45 Mbps) and above was not available. We obtained partial funding from National Science Foundation (NSF) and the National Response Center (NRC), but budget constraints limited us to developing a smaller prototype system comprised of a base station and two field units. These were to be built around commercial Local Multipoint Distribution Service (LMDS) equipment and provide 10/100BaseT (10/100 Mbps) Ethernet network interfaces. The prototype retained the path-sounding features of our original proposal and addressed the research issues associated with allowing the wireless network adaptively to utilize the full bandwidth that the sounders indicate is available.

In the aftermath of 9/11 we moved quickly to field a demonstration network running at a 10 Mbps data rate and successfully demonstrated it to enthusiastic audiences including elected officials and disaster response professionals at all levels. We felt that our project's time had come, since the 9/11 attacks clearly demonstrated the need for rapidly deployable broadband connectivity. But this did not translate into new funding from the mission agencies; they continue to view our work as "research." They were overwhelmed with manufacturers offering to give away or sell a variety of broadband systems operating in the 2.5, 5, and 60 GHz bands and at optical wavelengths. Disaster relief agencies began to purchase these, but generally on a piecemeal basis rather than as parts of integrated networks.

Our difficulty in moving our research project's results from academic research to operational use is in itself an interesting study in technology policy. In this paper we present the problem we set out to solve (rapidly deployable broadband communications for disaster response), the underlying research issues (broadband channel sounding, adaptive zero-administration wireless networks), our systems demonstrations for potential users, and what happened afterwards.

2. Technical Approach

Our goal is to develop a system that can rapidly deliver high data rate connectivity to support networked applications for first responders and on-scene commanders. As illustrated in Figure 1, a base station (B) "illuminates" the disaster area with high data rate connectivity using broadband wireless technology. Field units (F) placed within the illuminated area connect to the base station to send and receive data. The base station provides connectivity to other field units and, ideally, to a global network using either existing infrastructure or a satellite connection. The field units can provide connectivity to personal computers, notebook computers, personal digital assistants (PDAs), and other networked devices using wired and wireless local area networks. The system could be deployed and operational within a few hours after equipment arrives at the site.

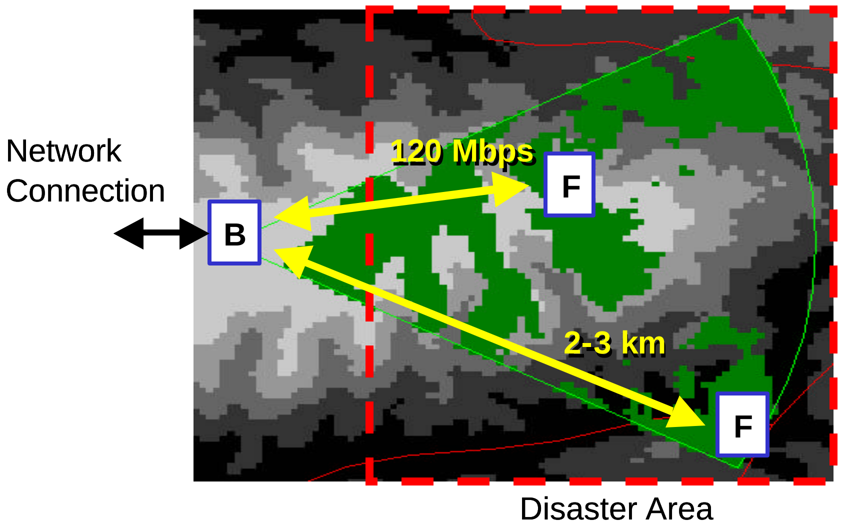

Figure 1. System concept of illuminating the disaster area with broadband wireless.

Our objective is to provide a data rate of 120 Mbps. We are using radio frequency (RF) spectrum in the licensed LMDS band, which is between 28 and 31 gigahertz. In theory, this band can support data rates of one gigabit per second. There are other licensed and license-free bands that can also be used for broadband wireless. Our choice of LMDS is based on its high bandwidth that enables high data rates, the availability of spectrum – Virginia Tech owns LMDS licenses for much of southwestern and southern Virginia – and, the availability of commercial off-the-shelf (COTS) RF components.

The system we are developing has a range, or maximum distance between the hub unit and a remote unit, of 2 to 3 kilometers. This range limitation is due to our use of COTS radios that were donated to us. Radios designed specifically for this application could easily achieve ranges on the order of 10 to 15 kilometers.

The logical network formed by the hub and remote units has a star topology and is designed to transport Ethernet frames. The hub and remote units include a full-duplex 100-Mbps Ethernet interface. Ethernet frames sent to a remote unit by hosts or network equipment at a remote location are delivered to the hub unit and sent to network equipment – typically an Internet Protocol (IP) router – at the hub site. Ethernet frames sent to the hub unit by network equipment at the router are delivered to all remote units, which deliver the frames to hosts and network equipment at the remote site. This simple logical behavior and use of Ethernet as the basic link-layer service allows the use of common Ethernet-based network equipment and IP-based applications. The broadband wireless backbone is transparent to other network equipment and to standard applications.

Security is of particular concern in wireless networks. The band that is utilized for our system greatly reduces the vulnerability to casual eavesdropping, which is a significant concern with more common IEEE 802.11 wireless local area network (WLAN) technology. For more stringent security requirements for privacy and authentication, we believe that virtual private network (VPN) approaches are the most appropriate solution. A VPN provides secure end-to-end transport, thus protecting traffic over wireless link as well as over the public network that is likely to be used to carry traffic back to agency networks. In addition, multiple VPNs can isolate traffic and network resources belonging to different agencies that are sharing a common wireless network. Support for VPNs is becoming common in personal computers (e.g., in Microsoft Windows XP and 2000) and low-cost network routers.

As noted above, a growing number of COTS systems now provide high data rate connectivity.

However, our approach incorporates three unique and innovative capabilities that enable rapid deployment and robust operation for emergency and disaster response applications.

- The hub and remote units include an integrated broadband channel sounder (or, more precisely, a sampling swept time delay short pulse sounder). Sounders are commonly used for site planning to assess the RF channel characteristics of a particular location. However, traditional sounders and expensive. We have developed an innovative sounder design that dramatically lowers the cost to a point where it is feasible to build the sounder into the hub and remote units. A built-in sounder simplifies initial equipment set-up and, also, allows for continuous monitoring of the RF environment that may change, for example due to the movement of trucks heavy equipment at a disaster site.
- The communications link between the hub and remote units can adapt to changes in the quality of the connection. Based on current conditions, as indicated by the built-in sounder and detected errors, an automatic retransmission request (ARQ) function is selectively enabled or disabled and forward error correction (FEC) coding is dynamically set to different levels of protection against bit errors. The goal of this link-layer adaptation is to achieve optimal operation based on run-time versus design-time conditions.
- A Geographic Information System (GIS) "viewshed" analysis tool can be used for "on-the-fly" site planning to determine coverage within the disaster area and to indicate optimal locations for the hub and remote units. Viewshed analysis is especially important given that LMDS and other nearby bands require line-of-sight connections between the hub and remote units. The opportunity also exists for viewshed analysis to use measurements from the sounder to improve its prediction and to suggest alternative locations.

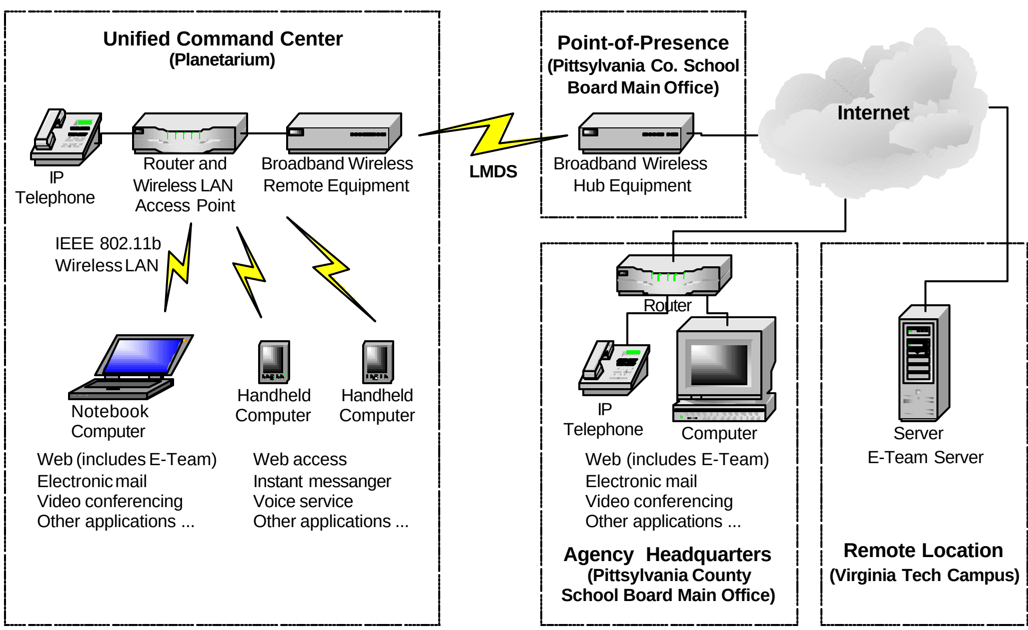

Figure 2. System used for January 7, 2002 demonstration.

3. Prototype Demonstrations

The approach described in Section 2 requires the development of new hardware and software, which is ongoing. In the aftermath of 9/11, there was a clear need to demonstrate the capabilities and potential for such a system. We constructed a prototype system (see Figure 2) with a base station and one field unit using COTS radios, up/down converters, 10-Mbps modems, and networking equipment. There was some custom hardware to interface the radios to the modem. The base station was connected to the Internet at a remote location. The wireless system provided all network access to the simulated "unified command center" (UCC) viewed by participants. 10 Mbps Ethernet and an IEEE 802.11b WLAN provided connectivity to a notebook computer, PDAs, and voice-over-IP telephone. We demonstrated a number of applications including general web access, access to the E-Team emergency management system, instant messaging, push-to-talk voice communication using PDAs, voice-over-IP telephony, and videoconferencing.

We conducted two demonstrations. The Virginia Preparedness and Security Panel, convened by Governor James Gilmore, was the audience for the first demonstration, held on November 14, 2001 in Blacksburg, VA. The panel included the state's Attorney General, other state and local officials, and industry representatives. A second demonstration was held in Chatham, VA on January 7, 2002 where the audience included representatives from the Federal Emergency Management Agency (FEMA), the Department of Defense, the NRC, the NSF, and local emergency services agencies.

4. Broadband Wireless for Emergency Services in the Aftermath of September 11

The world changed in many ways on September 11, 2001. The disastrous aftermath of the World Trade Center and Pentagon attacks clearly demonstrated the need for rapidly deployable broadband connectivity. Recognizing the need, manufacturers have begun to offer a variety of broadband radio and optical systems. FEMA reportedly was so overwhelmed with offers of free or low-cost trial systems that they had to set up a separate web site to deal with the offers. Disaster relief agencies have started acquiring these, generally on a piecemeal basis rather than as parts of integrated networks.

Many of the radio systems being offered represent the flowering of IEEE 802.11 technology which, in its various flavors (802.11a, 802.11b, 802.11g, etc.) carries data at rates from 11 to 155 Mbps over relatively long ranges (some manufacturers promise up to 25 miles) in mesh and non-line-of-sight (NLOS) configurations. While many of these claims remain to be substantiated (i.e., the user can have 155 Mbps, *or* 25 mile range, *or* NLOS operation), these systems represent a significant departure from the state of the art in 1999 when we initiated this project. At that time 802.11b was an 11 Mbps indoor system. We found only a few manufacturers claiming to be capable of making 155 Mbps class radios, and those were for the satellite or LMDS licensed bands.

These laudable efforts by industry leave disaster relief agencies with a multiplicity of new products to try without, necessarily, the equipment and experienced personnel needed to evaluate them. For example the user agencies may be unaware that IEEE 802.11 unlicensed systems *must* accept whatever ambient level of interference exists at their operating location. Currently, there is no legal basis by which one user (even the federal government) can compel an interfering system to shut down. Sales people may speak glowingly of how spread spectrum modulation overcomes interference (and security) problems, but there is much anecdotal evidence of unlicensed systems failing catastrophically when, for example, the Goodyear Blimp flies over (since the blimp uses 802.11b technology to transmit video) or when a new user sets up a similar system in the neighborhood.

Thus the LMDS system that Virginia Tech and SAIC developed is now only one of a number of radio technologies capable of providing the needed throughput. While in 1999 we were far ahead of industry in trying to meet the needs of the disaster relief community; in some sense we have been overrun from behind by a stampede! Our product, though, is an integrated test bed; as such it is "radio agnostic." We can easily apply our broadband sounder technology to other radio bands and use it for rapid deployment and for adaptive operation. Operational agencies, though, are unused to funding research projects – at least outside of the Department of Defense. They honestly see it is as more cost effective (and certainly safer) for them to buy commercial products from established vendors and to leave the "research" to NSF and DARPA. We have gotten many enthusiastic comments from senior government officials – even "on the record" in TV interviews – but no commitments of the additional funds that would give the mission agencies a large-scale field-ready system.

Accepting the reality of the situation, we are proceeding as follows. First, we will complete our prototype network as specified in our current NSF grant – one hub and two remotes running at 120 Mbps full duplex, incorporating the broadband sounder and the GIS tools, and packaged for rugged field use. We will augment the system with as many enhancements as we can to make it more useful to search and rescue teams. To this end the students in a graduate course in wearable computers are developing wearable wireless appliances that our network will support. We plan to demonstrate the complete system to FEMA and other mission agencies in the field beginning this summer in collaboration with local fire and rescue organizations.

Second, we are working to form an alliance with one of the largest makers of IEEE 802.11 equipment to help them incorporate our innovations in their networks and develop a complete product package for emergency and disaster applications. This seems to be one of the best ways to move our research results into operation.

Third, if the disaster relief community is to adopt a state of the art wireless networking system that will link emergency responders at all levels, it needs disinterested expert advice from more than just the industry that makes the equipment. One of the reasons that current equipment is notoriously expensive and unable to interoperate is that it is based on narrow proprietary standards that limit competition. For the E-Government Wireless Networks Initiative to "achieve the goals of increased spectrum efficiency; interoperability among local, state, and federal agencies and maximize efficiencies and savings through shared infrastructure and common procurement strategies," the disaster relief community must adopt IP-based wireless networking technology that is coming into widespread commercial use. We hope to use our project contacts and experience to make this happen. If industry just develops another set of proprietary standards that will be used only by the disaster relief community, it will only perpetuate many of the current cost and interoperability problems.

5. Conclusions

From a research point of view, our project has made some advances in wireless network and radio technology. More important, we believe, the researchers have developed some meaningful contacts with the disaster relief community and learned a great deal from them. We understand what they need and are developing a communication system that provides it. They have seen our system and given it a most favorable response. The biggest challenge remaining is to get the results of our work into operational hands. We are meeting it by working directly with the industries whose products are, in some sense, our competitors, and by developing a superior demonstration system for coming field trials.

Acknowledgements

We gratefully acknowledge the contributions of all members of the project team, including other faculty and students at Virginia Tech and Rick Klobuchar and Mike Kurgan of SAIC. We also gratefully acknowledge the support of the National Science Foundation's Digital Government program (grant 9983463) and the National Response Center. We also appreciate Virginia Tech's Information Systems organization's contribution to the two demonstrations and equipment donations by Motorola and Spike Technologies.