

Challenges in Deploying Steerable Wireless Testbeds

Eric Anderson^{1,2}, Caleb Phillips¹, Gary Yee¹, Douglas Sicker¹, and Dirk Grunwald¹

¹ University of Colorado, Department of Computer Science
430 UCB, Boulder, CO 80309, USA

{caleb.phillips, gary.yee, sicker, grunwald}@colorado.edu

² Carnegie Mellon University, Department of Electrical and Computer Engineering
5000 Forbes Ave, Pittsburgh PA 15213
anderson@ece.cmu.edu

Abstract. Phased array antennas enable the use of real-time beam-forming and null-steering to further increase control of signal strength and interference in wireless networks. Understanding the potential of this platform for both mesh and single-hop networks is becoming more important as smart antennas begin to appear in emerging networking standards. Prior attempts to test non-standard antenna platforms have typically focused around simulations, fixed (non-steerable) directional antenna testbeds, and small scale temporary setups utilizing 1 or 2 phased array antenna nodes over the span of a few hundred meters. This paper presents the challenges encountered – and solutions developed – in building WART, a permanent, campus-wide testbed for wireless networking with beam-forming antennas.

Keywords: Smart antennas, steerable, directional, phased array, antennas, outdoor, wireless, testbed, 2.4GHz.

1 Introduction


Directional antennas, both fixed and steerable, are proving to be important in the next generation of wireless networking protocols. These antennas give nodes further control over both signal strength and interference, allowing optimization techniques which can yield greater network throughput with fewer errors. While protocols incorporating directional or “smart” antennas have been proposed, their evaluation has been limited. Those researchers who have attempted real-world evaluation of their ideas have often used one-off testbeds assembled to perform a small number of experiments [11,9,8]. Most proposals, however, rely solely on simulation or theoretical analysis (for instance, [14,13]).

In this paper we introduce the University of Colorado Wide-Area Radio Testbed (WART) as a platform for studying uses of directional, steerable, and smart antennas in wireless networking¹. Given the widely-recognized difficulty of accurately simulating radio environments, real-world experiments are essential to evaluate wireless networking protocols. In the case of directional applications, which are especially dependent on the vagaries and environmental effects of radio propagation, this is even more important [2,4].


¹ An expanded version of this paper is available in a companion technical report at [3].

WART is currently the only permanent facility for studying smart antennas in a large and diverse urban environment. The system consists of eight phased array antenna nodes mounted to the rooftops of university buildings, spanning an area of 1.8 x 1.4 kilometers. The entire testbed is linked together via wired Ethernet and can be controlled from a single administration point. This architecture ensures that WART can not only offer the geographic scale and realism of large scale distributed testbeds [1], but can also give its users the degree of control and ease of management only seen in dense indoor testbeds such as ORBIT and Emulab [12,7].

The production and deployment of such a testbed, however, is itself an engineering problem. In addition to the capabilities of WART, this paper describes some of the logistical challenges encountered in planning, installing, and maintaining a centrally controlled wide area rooftop network.



(a) Campus testbed (1.8 x 1.4 km)



(b) Connectivity of nodes 2 through 8 without beam steering



(c) Installed antenna node

1.1 Design Goals

WART is intended to be a dedicated experimental testbed for studying the impact of directionality and beam-forming throughout the network stack. Given this objective, we chose three principle design goals for WART: (1) The testbed must be able to perform outdoor omni-directional, fixed directional, and beam-forming experiments; (2) The testbed must be able to test a diverse set of link distances of varying link qualities; (3) WART nodes must be simple to reconfigure for varying experiments and provide an easy recovery mechanism in case of failure. The node sites were chosen to provide a variety of link lengths, with line-of-sight between many but not all pairs of nodes.

The remainder of this paper describes the hardware, software, and centralized architecture of WART that helps fulfill the design goals of easy maintenance and administration.

1.2 Smart Antenna System

In this section we describe the hardware and software that comprise WART. These components give it the unique ability to perform smart antenna research at all network stack levels and address challenges with its administration and experimental setup.

Each smart antenna node consists of a phased array antenna and an embedded computer. The phased array antennas were designed and constructed by Fidelity Comtech. The antenna operates in the 2.4GHz ISM band and uses an 8 element uniform circular

array of dipole antennas that support a minimum 42° primary lobe. The ratio of the lowest null to the highest peak is $\approx 40\text{dB}$, which allows for selectively “nulling out” interfering signals. The antenna arrays can be electronically switched between radiation patterns in $\approx 100\mu\text{seconds}$, allowing for precise dynamic reconfiguration. The wireless interface card used is an IEEE 802.11 Senao 5345MP MiniPCI adapter, which uses an (especially flexible) Atheros chipset.

The default Operating System (OS) image used by each WART node is a modified OpenWRT Linux (Kamikaze) distribution. The wireless driver is based on the Multi-band Atheros Driver (MADWiFi) version 0.9.4.5 and has been extended to support the smart antennas’ ability to change antenna patterns, along with various measurement and experimental Medium Access Control (MAC) features. Unix Network File System (NFS) is used to transmit data from the smart antenna node to long-term storage.

2 Administration and Maintenance Infrastructure

Experimental hardware and software is almost inevitably flawed, and faults which escape notice during testing regularly cause problems during live experiments. When problems do occur, equipment needs to be rebooted, experiments need to be re-started, scripts need to be edited, and sometimes new software needs to be installed. The (human) communication overhead of trying to identify and correct problems across all test locations quickly becomes prohibitive, even when the necessary fixes are small. In early tests we found that even when nothing went wrong, coordinating a four node experiment required at least a half-hour of overhead for setup, configuration checks, synchronization, starting the experiment, downloading the data afterwards, and running basic sanity checks on the data. Overall, the ratio of time expended to successful experiment time was very high.

Our primary requirement for the testbed infrastructure was that it enable centralized management. At its core, this infrastructure consists of a control plane network, a “management box” connected to each experimental antenna unit, and a collection of software tools. All of these will be described in upcoming sections.

2.1 Management System

Every experimental antenna unit is directly connected to a management box, as depicted in Figure 1. These boxes connect the experimental units to the control plane network. Additionally, the management boxes also provide network booting and remote power control to the antenna units. This approach greatly simplifies reconfiguration: Any software change, from one configuration file to a new operating system, can be made by uploading a new image to the management system and rebooting the experimental equipment. The equipment could boot from a remote server, but only if the intervening network had the configuration and performance to support it; such a requirement which would limit options substantially.

Each management box contains a single-board Soekris computer running Linux and can be installed indoors at a significant distance from the antenna unit. All of the current deployments have Ethernet connections, but they can accommodate other data connections with minimal configuration changes.

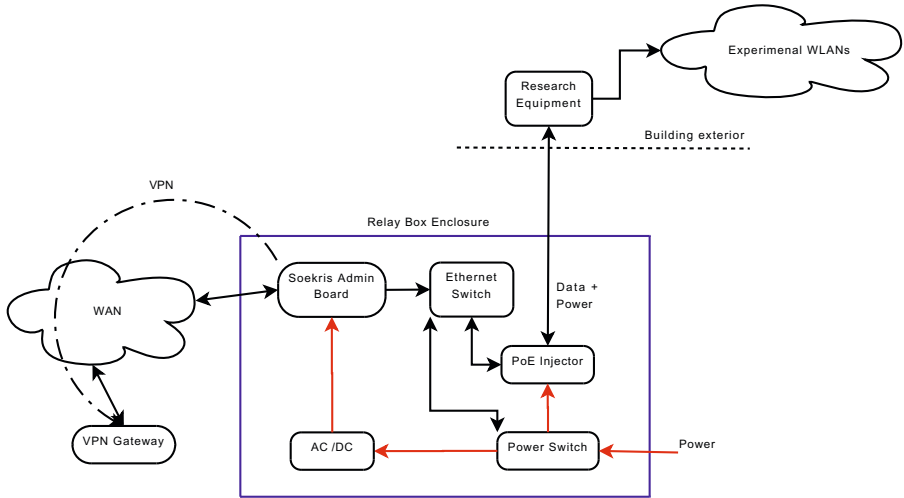


Fig. 1. Management box configuration

The organization of the control plane network relies heavily on the use of a virtual private network (VPN). Management units establish and maintain connections to a central VPN server, so that all the devices are accessible as a single IP network. Network Address Translation (NAT) is used to map hard-wired local addresses to node-specific VPN addresses. The practical effect of this configuration is that error-prone customization and dynamic reconfiguration are largely avoided, yet every device is reachable by a globally unique address, and most configuration information is stored at the central server where changes are easy to make. As a final failsafe, each management unit contains a remote power switch for each component allowing for a hard-reboot of any device.

2.2 Interchangeable Parts

Each phased array antenna unit or network power switch has exactly the same hardware and configuration as every other. Every management computer is the same as every other except for the contents of a removable compact flash card. This makes it easier to develop testing processes for each component and means that a faulty or suspect component can be replaced with no thinking or configuration required. In fact, it is often easiest to replace the entire management unit as a whole – except for the flash card – and then diagnose faulty equipment in the comfort of the lab. Hardware and configuration homogeneity make software management more practical. All of the files associated with the testbed – source code, configuration files, and compiled operating system images – are kept under version control. Because the antenna units are identically configured, there is one generic OS image for all of them.

2.3 Security

Since WART nodes are connected to untrusted networks, they are potentially susceptible to the same attacks that many other machines on the University of Colorado network experience on a day-to-day basis. Several steps have been taken to ensure that only authorized access is given to both the phased array antenna node and management board. Communication to the WART management nodes is restricted to nodes that are part of the same VPN. This requires having a certificate signed by the certificate authority, a process which is performed off-line. Within this trusted domain, we use SSH keys to allow remote logins directly to the antenna and management nodes.

Another possible attack vector is via the wireless interface. Should an attacker inject traffic to a node, the node could potentially begin routing packets from unauthorized users. To minimize this risk, nodes perform only the minimum forwarding required for experiments.

3 Deployment Logistics

Deploying a physically large testbed, especially with outdoor equipment, involves a number of challenges outside the traditional realm of computer science. There is a modest inherent engineering component, which is significantly compounded by the need for approval and cooperation from various outside parties. All of the WART nodes are located on University of Colorado property, meaning that we only needed to interact with a single (albeit large and bureaucratic) owner. We suspect that broadly similar issues would be likely to arise in working with another large organization, and possibly with multiple smaller ones.

Some of the more prominent logistical challenges encountered were:

- *Architectural Approval*: The aesthetic impact on campus buildings had to be approved by the campus architect.
- *Antenna Siting and RF Interference Approval*: A separate antenna committee had to be convinced that the proposed sites would not interfere with existing radio equipment.
- *Electrical Design and Installation*: The electrical requirements of the testbed equipment are extremely low – each node uses less power than a desk lamp. However, all construction projects involving new electrical connections are subject to the same approval process, regardless of the actual load. This means that an electrical design for each node had to be completed and signed off by a certified electrical engineer and installation of the electrical components had to be performed by licensed electricians. Both tasks had to be done by outside contractors, requiring an additional round of financial approvals before work could begin.
- *Environmental Health and Safety*: All construction projects must be audited for safety risks to both the workers and the campus in general. The primary concern was disturbing pre-existing asbestos building materials, although we also had to vouch for the microwave radiation levels.
- *Roof Integrity*: Because the equipment was to be mounted on the outside of buildings, both the attachment methods and cable connections had to be evaluated for

waterproofing, fire sealing, and structural impact. In the cases where new holes had to be made through the roof, the penetration and waterproofing had to be installed by campus roofing services.

- *Antenna Structure*: Local building codes and campus design rules establish standards for wind, snow, and ice tolerance. The university requirements were the more stringent in this case, requiring that equipment be designed for 120 mile per hour (53.6 m/s) wind load. Antenna mounting equipment, especially in the WiFi market, seldom meets those requirements. While commercial options do exist, we found it more cost-effective to design and construct our own.
- *Financial Approvals*: After our research group and department decided to allocate funds for the testbed, there were still a significant number of delays waiting for work orders and payments to be approved by other university entities. In particular, payments from the computer science department to facilities management, and from facilities management to outside contractors all required administrative approval before the payee could begin work.

3.1 Timeline

The testbed deployment process has required a total of two years. Most of that time has consisted of waiting for some necessary action by parties outside our department. Within that waiting, most of the time has been for administrative approvals, with actual design and construction requiring relatively little. Figure 2 shows our actual timeline; with more foresight it probably could have been compressed.

The architectural and RF approval steps are an unavoidable bottleneck, as they determine whether and where equipment can be installed. In our case, it required approximately nine months from the first informal proposals to a preliminary approval of the sites chosen. Once those decisions had been made, several of the remaining steps could likely have proceeded at once.

The obvious deployment tasks, namely physically installing the antenna node and management box, and running conduit and Ethernet cable between them, required on the order of one week per node.

3.2 Costs

Table 1 presents an approximate breakdown of the expense incurred *per node* in building this testbed. The dominant cost is not the research equipment itself but rather labor required for regulatory and university policy compliance. This includes both the electrical work mentioned earlier and the time spent by university employees on evaluation and project oversight.

4 Proof-of-Concept Experiments

As a proof-of-concept experiment for WART, we performed a full pairwise link quality test. In this test, each WART node takes a turn transmitting while the other nodes listen. During each turn, the transmitter and all receivers cycle through 17 pre-configured




Fig. 2. Testbed deployment timeline: The entire process took 2 years, 16 months of which was devoted to planning and administrative approval while only 8 were required to install and configure the hardware.

antenna patterns, so that every combination of transmitter and receiver antenna patterns is tested. The patterns chosen point the main lobe in one of 16 directions about the azimuth plane (the 17th pattern is omnidirectional). Using the measured signal strength of received packets, we are able to determine (a) which links are possible between which nodes and (b) what the optimal “greedy” patterns are for each link. The results of this experiment are provided visually in Figure 3, which we believe makes a compelling case for the power of steerable directional antennas. When configured with omnidirectional patterns, which are comparable to the antennas used in many single-radio mesh networks, only a few links are even possible, and of those only a small number offer decent signal quality. With steering, however, we see a vast improvement: not only are all link-pairs able to pass traffic, but these links are typically of high quality (greater than -70 dBm).

Our present and future research utilizes WART to evaluate directional medium access control (MAC) protocols, with a particular emphasis on optimization for spatial re-use. We believe that the unique opportunity that WART provides for real-world evaluation of these protocols will lead to important results in this direction, and new insights into methods for improving wireless systems in general.

Table 1. Cost of labor and parts per WART node. The labor of research group members is not considered.

Description	Cost
Phased Array Antenna Node	\$3,000
Management Box and Other Control Plane Equipment	\$1,200
Installment Materials	\$300
External Labor and Fees	\$5,780

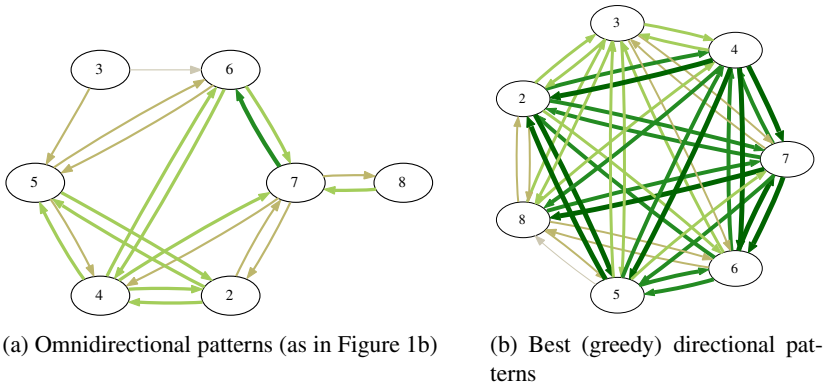


Fig. 3. Comparison of available links and link quality between seven testbed nodes using best-steered directional patterns and omnidirectional patterns. Stronger links are indicated with a wider arrow of a darker color. The best links are those with a link of greater than -60 dBm. The worst links plotted are barely above the noise-floor with greater than -95 dBm achieved RSS.

5 Related Work

There are too many wireless testbeds to discuss individually in the available space. For a more thorough treatment, please see our companion technical report [3]. Despite the large number, we are not aware of any existing testbeds that address the particular needs of WART. We observe that these testbeds fall into one of two categories: Wide area testbeds, which cover a significant outdoor environment but offer limited control over each node, and dense indoor testbeds with many nodes and excellent facilities for re-programming and control, but with very artificial RF environments.

The existing outdoor testbeds generally have more operational emphasis and less experimental control and management support than WART or the indoor testbeds. Most use stock 802.11 at the MAC and physical layers, although additional low-layer information can be collected. This may in part reflect their designers' research interests and may also reflect limitations resulting from the lack of a stable separate control network. Notable examples include Roofnet [5], the Rice/TFA mesh [6], and the Digital Gangetic Plains project [10].

In general, the indoor testbeds are physically smaller than the outdoor ones and benefit from a much more controlled environment. The problems of remote repair and establishing and maintaining a reliable communication infrastructure, which have been at the forefront of our design challenges, are largely non-issues. Many of the indoor testbeds have at least an order of magnitude more nodes than any of the outdoor ones: Both ORBIT and Emulab have over 400 nodes [12,15,7]. Much of the infrastructure developed for the indoor testbeds is oriented toward automating the process of configuring, controlling, and aggregating data from such a large collection of devices.

6 Conclusion

This paper has presented WART, a testbed that will facilitate future networking research by providing unique physical layer capabilities not seen in any other outdoor networking testbed. While the testbed covers an entire university campus, it is easy to manage and administer due to its wired control plane, which is remotely accessible from anywhere on the Internet.

The research motivation for building WART was to study the use of directional, steerable, and adaptive antennas. The prominent issues encountered in creating the testbed proved to be only indirectly related to that objective. The direct causes were *using commodity equipment*, *supporting low-level experimentation*, and *spanning a large geographical area*.

Commodity equipment: The research equipment (phased array antenna nodes) is comparatively affordable at \$3,000 per node, while specialized test and measurement equipment could easily cost 10 to 20 times more. The consequences of using commodity hardware have been the need for significant calibration and testing and extensive software hacking to make the hardware operate in unintended ways.

Low-level experimentation: Many of the experiments we wish to conduct are low-level both in the sense of being at the physical and MAC layers of the OSI hierarchy, and in the sense of requiring “close to the metal” system implementation. This implies the need for easy reprogramming and crash recovery, high-volume data collection, and a flexible control interface. In practice, these in turn require a control connection that is separate from the experimental wireless system.

Large geographical area: It has been amply demonstrated that radio propagation in general, and directionality in particular, are very environmentally dependent [2]. Consequently, it was important that WART encompass a range of node densities and environmental features of interest. However, covering a large area implies *physical distance* and often *administrative diversity*, each of which contribute significant design challenges. Physical distance effectively precludes running dedicated cables from a central location to all of the nodes, which implies that power and network connectivity (if needed) must be supplied using resources available on site. It is this constraint which leads us to the “management box” design, with network support, power conversion, and power switching co-located with every measurement node.

Covering a larger area often implies involving more administrative domains. Our sites are all owned by the same university, but building at a campus-wide scale requires the involvement of many departments – administrative and academic – and the approval of several levels of hierarchy. The practical impact of this cannot be overstated. The approval processes – and the cascade of design decisions made in order to secure those approvals – account for at least half of the total time and cost for this project.

This testbed was developed to study particular physical layer technologies, but the design lessons are not specific to that objective. Most of the challenges encountered in designing this testbed – and the solutions developed – are likely to apply to other outdoor and wide-area testbeds. We have developed an infrastructure for deploying nodes at widely separate, minimally provisioned sites and connecting them into an easily-managed unified research system.

References

1. Aguayo, D., Bicket, J., Biswas, S., Judd, G., Morris, R.: Link-level measurements from an 802.11b mesh network. In: Proc. Sigcomm 2004. ACM, New York (August 2004)
2. Anderson, E., Phillips, C., Sicker, D., Grunwald, D.: Modeling environmental effects on directionality in wireless networks. In: 5th International Workshop on Wireless Network Measurements (WiNMee) (June 2009)
3. Anderson, E., Phillips, C., Yee, G., Sicker, D., Grunwald, D.: Challenges in deploying steerable wireless testbeds. Tech. Rep. CU-CS-1058-09, University of Colorado, Boulder (2009)
4. Anderson, E., Yee, G., Phillips, C., Grunwald, D., Sicker, D.: The impact of directional antenna models on simulation accuracy. In: 7th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt) (June 2009)
5. Bicket, J., Aguayo, D., Biswas, S., Morris, R.: Architecture and evaluation of an unplanned 802.11b mesh network. In: MobiCom 2005. ACM, New York (August 2005)
6. Broustis, I., Eriksson, J., Krishnamurthy, S.V., Faloutsos, M.: A blueprint for a manageable and affordable wireless testbed: Design, pitfalls, and lessons learned. In: TridentCom (2007)
7. Johnson, D., Stack, T., Fish, R., Flickinger, D.M., Stoller, L., Ricci, R., Lepreau, J.: Mobile emulab: A robotic wireless and sensor network testbed. In: Proceedings of 25th IEEE International Conference on Computer Communications, INFOCOM 2006, pp. 1–12 (April 2006)
8. Kohmura, N., Mitsushashi, H., Watanabe, M., Bandai, M., Obana, S., Watanabe, T.: Unagi: a protocol testbed with practical smart antennas for ad hoc networks. SIGMOBILE Mob. Comput. Commun. Rev. 12(1), 59–61 (2008)
9. Mitsushashi, H., Watanabe, M., Obana, S., Bandai, M., Watanabe, T.: A testbed with a practical smart antenna for directional mac protocols in ad hoc networks. In: AINAW 2007: Proceedings of the 21st International Conference on Advanced Information Networking and Applications Workshops, Washington, DC, USA, pp. 731–736. IEEE Computer Society, Los Alamitos (2007)
10. Raman, B., Chebrolu, K.: Experiences in using WiFi for rural internet in india. IEEE Communications Magazine (January 2007); Special Issue on New Directions In Networking Technologies In Emerging Economies
11. Ramanathan, R., Redi, J., Santivanez, C., Wiggins, D., Polit, S.: Ad hoc networking with directional antennas: a complete system solution. IEEE Journal on Selected Areas in Communications 23(3), 496–506 (2005)
12. Raychaudhuri, D., Seskar, I., Ott, M., Ganu, S., Ramachandran, K., Kremo, H., Siracusa, R., Liu, H., Singh, M.: Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols. In: Proc. IEEE Wireless Communications and Networking Conference, vol. 3, pp. 1664–1669 (2005)
13. Singh, H., Singh, S.: Smart-aloah for multi-hop wireless networks. Mob. Netw. Appl. 10(5), 651–662 (2005)
14. Takai, M., Martin, J., Bagrodia, R., A.R.: Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In: MobiHoc 2002: Proceedings of the 3rd ACM International Symposium on Mobile ad hoc Networking & Computing, pp. 183–193. ACM, New York (2002)
15. White, B., Lepreau, J., Stoller, L., Ricci, R., Guruprasad, S., Newbold, M., Hibler, M., Barb, C., Joglekar, A.: An integrated experimental environment for distributed systems and networks. In: Proc. of the Fifth Symposium on Operating Systems Design and Implementation, Boston, MA, pp. 255–270. USENIX Association (December 2002)