

The Rapidly Deployable Radio Network

Joseph B. Evans, *Member, IEEE*, Gary J. Minden, *Member, IEEE*,

K. S. Shanmugan, *Fellow, IEEE*, Glen Prescott, *Member, IEEE*, Victor S. Frost, *Fellow, IEEE*,

Ben Ewy, Ricardo Sanchez, Craig Sparks, *Member, IEEE*, K. Malinimohan,

James Roberts, *Senior Member, IEEE*, Richard Plumb, *Senior Member, IEEE*, and Dave Petr, *Senior Member, IEEE*

Abstract—The Rapidly Deployable Radio Network (RDRN) is an architecture and experimental system to develop and evaluate hardware and software components suitable for implementing mobile, rapidly deployable, and adaptive wireless communications systems. The driving application for the RDRN is the need to quickly establish a communications infrastructure following a natural disaster, during a law enforcement activity, or rapid deployment of military force. The RDRN project incorporates digitally controlled antenna beams, programmable radios, adaptive protocols at the link layer, and mobile node management. This paper describes the architecture for the Rapidly Deployable Radio Network and a prototype system built to evaluate key system components.

Index Terms— Adaptive protocols, mobile networking, programmable antenna beamforming, software radios, wireless networks.

I. INTRODUCTION

QUICKLY evolving events, such as natural disasters, law enforcement actions, or military force deployment require adaptive mechanisms to rapidly establish a communication infrastructure with minimal *a priori* planning and configuration. The Rapidly Deployable Radio Network (RDRN) project created an experimental systems to develop and evaluate architectures, protocols, network control software, and prototype hardware for mobile, rapidly deployable, and adaptive wireless communications systems. The approach used in the RDRN is to implement a radio network based on switches. When deployed, the switches establish a backbone network among themselves over high-capacity radio links and self-configure into an appropriate topology. Users access switches over a “cellular-like” architecture. That is, each switch supports users in a limited cell. Users are mobile, and there will be handoffs within a switch and between switches. Switches can connect to other communication facilities, such as fiber links, wire links, or satellite links.

Rapid deployment and mobility imply limited opportunity for *a priori* radio frequency, topology, and deployment-specific engineering. Hence, adaptive techniques at the link, network,

and management levels are required. The RDRN system uses the following adaptive techniques.

- 1) Digitally beam-steered phased array antennas establish the high-speed point-to-point links and track mobile users.
- 2) Programmable waveform generation changes modulation, frame lengths, and forward error correction at the data link layer.
- 3) Asynchronous transfer mode (ATM) virtual circuits follow mobile users through the system.
- 4) Switch and user locations to anticipate handoffs and required switching resources.

RDRN represents an initial step beyond the conventional view of a software radio as a configuration of radio frequency sections, analog-to-digital converters, digital-to-analog converters, and a digital signal processor. The RDRN system incorporates a software radio in the larger context of programmable antennas, mobile nodes, varying topologies, and networking services.

The RDRN project defined the architecture and system components, implemented and evaluated the components, and built two prototype edge nodes. The remaining sections of this paper discuss the RDRN architecture and system components, the prototype system, and results of experiments.

II. SYSTEM OVERVIEW

When the need arises for a communications infrastructure in a place that has little or no infrastructure or the infrastructure has been destroyed, RDRN is deployed. An RDRN system consists of two node types: 1) *edge nodes* that include network management, switching functions, and interconnection to other networks; and 2) *remote nodes* that are mobile and serve as user access points.

Deployment involves placing edge nodes across the area to be covered. Edge nodes support three radio networks: 1) a low-capacity order-wire network, 2) multiple high-capacity point-to-point radio links, and 3) multiple “cellular-like” radio links. Edge nodes also contain a switch for routing traffic between the radio links, control and management software, and possible interfaces to wired or satellite communication services. When activated, the edge node determines its location via the global positioning system (GPS), and begins broadcasting its location on the low-capacity omnidirectional channel and listens for other systems broadcasting their locations. As other nearby edge nodes are detected, the local node attempts to establish a high-capacity point-to-point radio link to these nearby nodes. See Fig. 1.

Manuscript received September 1997; revised November 2, 1997, March 19, 1988, June 15, 1998, and September 15, 1998. This work was supported by the Defense Advanced Research Projects Agency under Contract J-FBI-94-223, with funding from Sprint Corporation and the State of Kansas.

J. B. Evans, G. J. Minden, K. S. Shanmugan, G. Prescott, V. S. Frost, R. Sanchez, C. Sparks, J. Roberts, R. Plumb, and D. Petr are with the Information and Telecommunication Technology Center, University of Kansas, Lawrence, KS 66045 USA.

B. Ewy is with Sprint Corporation, Overland Park, KS USA.

K. Malinimohan is with Cadence Corporation, Santa Clara, CA USA.

Publisher Item Identifier S 0733-8716(99)02979-0.

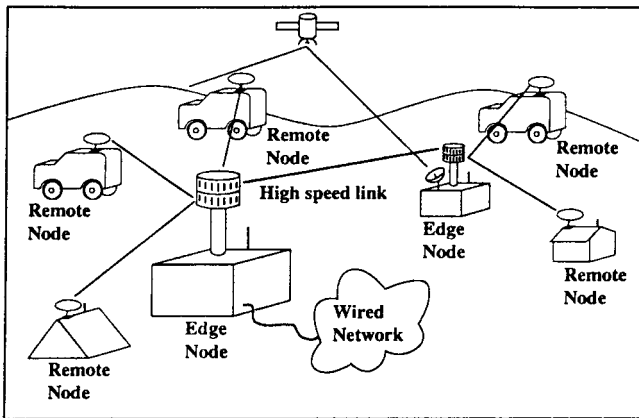


Fig. 1. RDRN system architecture. Edge nodes are deployed covering the area of interest. Edge nodes perform switching, and can interconnect to wired or space-based services. Mobile users access edge nodes via remote nodes that provide a cellular-like service. The RDRN tracks the location of remote nodes and adapts to their current positions.

The three radio networks may each be implemented with multiple radio transceivers. For example, the high-capacity point-to-point links will require a separate transceiver for each beam. Likewise, the “cellular-like” user access network will require a separate transceiver for each beam. The interconnection of transceivers to possibly shared antennas will depend on the frequency bands, the antenna characteristics, and performance issues.

Establishing the link requires consideration of multiple criteria.

- 1) Available carrier frequencies within the band.
- 2) Modulation and coding techniques (BPSK, QPSK, spread-spectrum codes).
- 3) Avoiding interference from other edge nodes in the system that are not directly linked.
- 4) Ensuring the edge nodes achieve and maintaining a connected topology.
- 5) Steering an antenna beam toward the nearby nodes and nulls toward interferers.

Link coordination takes place over the low-capacity, omnidirectional, order-wire link, while the local antenna is configured to steer its beams toward nearby edge nodes, establishing the high-capacity links. The backbone topology is based on the consistent labeling technique [1], [2] and is reported in [3]. The consistent labeling problem is a generalization of assigning labels to objects under a set of constraints, and is similar to the problem of assigning colors to map entities, or scheduling jobs on equipment, or packing small boxes in larger crates. While the problem is recognized as NP complete, the approach serves as a solid foundation for the establishment of the backbone topology.

Edge nodes implement a switching function to route traffic between radio links and possible wired or satellite services. The asynchronous transfer mode (ATM) provides the switching function and multiplexing between multiple streams which require independent classes of service. It is easily connected to wired services, and takes advantage of standards in the areas of signaling, quality of service, and flow control, among

TABLE I
NETWORK PROTOCOLS IMPLEMENTED IN THE RDRN PROTOTYPE

Protocol	Management
TCP/UDP	
IP	IP Routing
IP over ATM	Classical IP over ATM
AAL 5	
ATM	Q.2931 and RDRN NCP
Adaptive HDLC	RDRN NCP
Radio Link	RDRN NCP

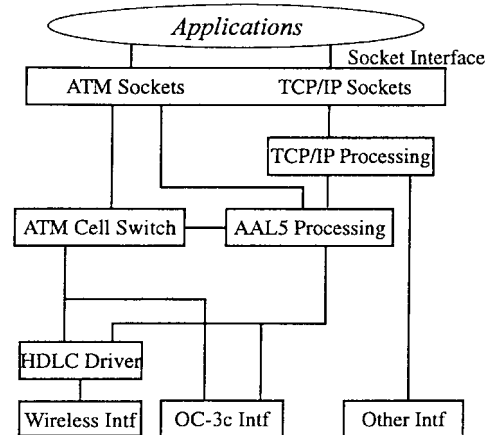


Fig. 2. RDRN protocol stack.

others. At the relatively low bandwidths of mobile wireless networks, the relatively quick multiplexing of ATM allows time-constrained traffic access to the channel in a timely manner.

Users access the edge backbone network through remote nodes and the cellular-like radio links around edge nodes. Edge nodes generate multiple digitally formed beams, each of which can be modulated independently. The prototype implements a 1 Mbaud (1 Mbit/s BPSK or 2 Mbit/s QPSK) TDMA protocol to support up to 64 users per beam. TDMA slots may be allocated to users for capacity management purposes.

Remote nodes obtain location information via GPS or other location techniques, and report their location to their connected edge node. Location information is used by the edge node for two purposes: 1) to steer the antenna beams to maximize the minimum signal-to-interference ratio for a set of nodes at known locations [4], and 2) to determine when to assign a remote node to a different beam or different edge node.

III. RDRN NETWORK CONTROL

This section describes the RDRN adaptive protocols and the network control program (NCP) that manages those protocols. Each protocol is first described, followed by the overall control architecture. The protocols are listed in Table I.

A. RDRN Wireless ATM Protocols

The RDRN protocol stack is shown in Fig. 2. This stack is implemented in the edge nodes and the remote nodes; however, radio link management is only implemented in the edge nodes.

RDRN protocols and management components have been implemented in the Linux operating system running on Intel $\times 86$ compatible computers. Linux provides an open development system with many contributors around the world running on a wide variety of computer architectures. The system supports applications running over both native-mode ATM as well as TCP/IP over ATM. Considerable development work was done to support standard ATM on Linux by researchers at the Laboratoire de Reseaux de Communication (LRC) at EPFL in Switzerland [5]. In particular, they developed a BSD-sockets-based application programming interface (API) to support native-mode ATM applications, as well as to support classical IP over ATM [6] for TCP/IP-based ATM applications [7].

The RDRN NCP supports applications through the standard UNIX socket interface to the networking protocols within the operating system kernel. The socket interface provides conventional TCP/IP services and ATM services. Applications can interface to the ATM service either through a framing protocol, such as AAL5, or through raw ATM cells. Within the kernel, the network layer implements the Internet Protocol (IP), whose primary purpose is to provide routing between subnetworks with service for the Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and other transport layers.

Classical IP (CLIP) and the ATM Address Resolution Protocol (ATMARP) [6] are initial standard solutions for carrying IP traffic over the ATM links. However, there are several weak points such as requiring a router to connect separate logical IP subnets (LIS), even when the subnets are directly connected at the ATM level, and requiring an ATMARP server to provide address resolution for a single LIS. The nonbroadcast multiple-access Next Hop Resolution Protocol (NHRP) [8] provides a better solution, but it is still in draft form. The RDRN architecture implements CLIP and supports both permanent virtual circuits and switched virtual circuits via ATMARP.

To this base system there were added four software components in the prototype system: 1) the Bellcore Q.Port ATM signaling system [9], 2) an ATM driver for handling AAL5 protocols, 3) an ATM cell switch, and 4) a driver for RDRN wireless links. The Bellcore Q.Port system provides a reference implementation of the ATM user-network interface (UNI). This software is used for call setup and teardown, virtual circuit management, and interoperability with wired systems. The ATM driver for AAL5 is a custom software component that can interface to conventional ATM interfaces and/or the RDRN wireless interface at the frame level, and handles the encapsulation of IP for the network layer. The ATM driver also distinguishes between virtual circuits (VC's) that carry AAL5 frames and those that carry raw ATM cells. In the case of VC's carrying AAL5 frames, segmentation and reassembly functions are handled in hardware on conventional interfaces and software for the RDRN wireless interface. Otherwise, for nonframed VC's, single ATM cells are handed to the ATM cell switch for switching to other VC's and then passed to the appropriate interface. The ATM cell switch is not present in remote nodes.

RDRN nodes implement an adaptive HDLC link layer over the wireless link. ATM cells are collected into a modified

HDLC frame, called the wireless ATM frame or WATM frame for transmission. In the case of edge-node-to-edge-node traffic, ATM cells from any VC can be included in the WATM frame. In the case of edge-node-to-remote-node traffic, ATM cells must be sorted into separate queues, one for each remote node. Only cells destined for a single remote are included in the WATM frame for transmission under the TDMA media access protocol. The WATM frame consists of a start-of-frame flag (1 byte), a link identifier (1 byte), a control byte (1 byte), one–nine ATM cells (53–477 bytes), and a 16-bit CRC field (2 bytes). Multiple ATM cells are included in each WATM frame to reduce the effect of encapsulation overhead necessary for addressing in the TDMA access mechanism and error control. No ATM header compression is used at this time. Likewise, we have not included forward error correction in the WATM frame. The number of ATM cells in a frame is based on resources at each end of the link and the link condition.

Radio link quality at the link level is estimated by counting the number of WATM frames received in error. This information is communicated upstream to the transmitting node. The HDLC link layer adapts to link quality in two ways. As the link quality goes down, fewer ATM cells are included in each WATM frame. Thus, any noise burst or fade over a few tens of microseconds will cause fewer ATM cells to be dropped. The second technique, used for loss-sensitive traffic, is to transmit multiple copies of the WATM frame, for example, two or three copies. Experiments have shown that WATM frame length adjustment with up to two copies of each frame transmitted is sufficient for most situations [10]. The protocol uses a sliding window and go-back- N ARQ scheme.

B. Network Control Program

The RDRN Network Control Program (NCP) is the primary control program for the system. As such, it monitors the backbone topology, locating remote nodes, computes antenna beam patterns, and implements connection handoffs between antenna beams and between edge nodes. The NCP executes on edge nodes. To perform these functions, the NCP monitors remote node locations and link quality to make decisions on configuration changes. The NCP causes antenna beams to be adjusted as remote nodes move, and reroutes ATM virtual circuits as remote nodes move between beams and edge nodes.

1) *Maintaining Backbone Topology*: The current RDRN implementation assumes that edge nodes remain stationary during the operation of the network. However, edge nodes can enter or exit the system, causing backbone topology changes. As mentioned above, the algorithms for determining backbone topology are based on a consistent labeling approach. In consistent labeling, constraints are expressed as pairwise constraints, e.g., if node A transmits to node B on frequency F , then node B can transmit to node A on any frequency except F . The consistent labeling algorithm considers all pairwise constraints, and determines a consistent labeling of all links.

The consistent labeling problem is a generalization of the graph coloring problem, the cellular frequency assignment problem, and many others [1], [2]. Suppose that one has cellular frequencies F_1, \dots, F_7 . Further, suppose that cells A

and B are adjacent, and that cell A is assigned frequency F_1 . Then cell B can only use frequencies F_2, \dots, F_7 . This constraint can be expressed as a list of 4-tuples

$$\{\langle A, F_1, B, F_2 \rangle, \langle A, F_1, B, F_3 \rangle, \dots, \langle A, F_1, B, F_7 \rangle\}$$

indicating the allowable combinations of cell A , the frequency assigned to cell A , cell B , and the frequencies assigned to cell B . The consistent labeling algorithm takes such a list describing the allowable combinations for all cell pairs, and determines a consistent assignment of frequency to cells meeting all of the constraints.

The advantage of the consistent labeling approach is that multiple criteria, such as frequency, beam direction, and spread-spectrum code, can easily be included in determining the backbone topology. The disadvantage is that the labeling problem is NP complete. However, due to the limited number of edge nodes, the computational speed of available computers, and heuristics that can be incorporated into the implementation, this approach is expected to be sufficient for an RDRN-type system.

2) *Monitoring Locations of Remote Nodes:* The NCP at each edge node monitors the location of nearby remote nodes through either the order-wire link or in-band signaling. From this information, the NCP determines whether to adjust an antenna beam to accommodate remote node motion or to switch the remote node to a different beam. In the former case, the edge node considers the locations of all of the remote nodes within a current beam, and computes new antenna beam patterns to maximize the minimum signal-to-interference ratio across all remote nodes [4]. In the latter case, the edge node moves a remote node to a new antenna beam, re-calculates the antenna beams for the “source” and “destination” beams, and reroutes all necessary ATM virtual circuits. In the current system, the beams are re-calculated every 5 seconds.

An RDRN variant predicts the motion of remote nodes. Future positions of remote nodes are predicted from recent positions and a simple velocity vector. Using predicted positions, the RDRN system can preallocate resources to make smoother handoffs between beams or nodes. The system is described in [11].

IV. RDRN RADIO TRANSCEIVER

The radio transceiver supports the RDRN system requirements for rapid deployment, and responds to a changing radio-frequency environment. The radio system allows multiple transmit beams to be steered to multiple remote node users from the same antenna at the same frequency, thus allowing spatial reuse of the frequency bands. It simultaneously allows the steering angles of the beams and the modulation types used within them to adapt to a changing mobile environment.

Before the specifics of the hardware design are presented, introductory information is given that places perspective on its design. Flexible software-based modulation and spatial reuse of frequencies are the main functions of the hardware. These topics are therefore introduced first.

The details of beamforming are discussed, and various analog and digital implementations are compared and con-

trasted. They are followed by a description of the hardware implementation that combines the task of beamforming and flexible IF modulation into a single cost-effective solution. Experimental results are then offered in support of the design. Finally, the receiver system and digital beamforming in the receiver are discussed.

A. Software-Radio-Based Flexible Modulation

Recent advances in analog-to-digital converter technology have stimulated an increasing interest in software-defined radio technology. With the capacity to digitize an IF or even an RF signal comes the ability to implement radio receivers that can be reprogrammed to emulate virtually any type of traditional receiver. One of the main goals of the RDRN proof of concept system is to provide a research testbed for the development of adaptive software-defined radios. The system presented in this section is a programmable transmit modulator only. The design is restricted to any type of phase or amplitude modulation. A software-radio-based beamforming receiver is the subject of future research.

B. Spatial Reuse

Spatial reuse of frequencies requires that the same carrier frequency be reused at least twice by the transmitter within what is considered to be a traditional cell. This can be achieved in different ways. Fig. 3 shows a two-user scenario where the positions of the transmitting antenna labeled edge node and both of the remote nodes are all known. The steering angles for the two beams Θ_1 and Θ_2 can therefore be calculated by simple geometry.

A brute force and intuitive approach to spatial reuse of the same carrier frequency would be to use two separate microwave “dishes” both located at the transmitting edge node, and simply point them at the individual users. This approach would work, and it describes well the requirements of a system that employs spatial reuse. However, it also has some serious drawbacks. The major disadvantage is that the individual beams must be steered by physically rotating the dishes to track the moving end users, and the beam patterns are determined at the time of dish manufacture.

The approach used in the RDRN prototype system is an electronically controlled beam-steering system in which multiple beams can be produced by a multielement phased-array antenna. The prototype system implements an equally spaced eight-element linear array.

C. Transmit Beamforming

Several antennas can be arranged in space and interconnected to produce a directional radiation pattern. This configuration is called an antenna array or, simply, an array [12]. When the phase of the signals reaching the individual antenna elements is modified in a useful way, then the antenna is called a phased array. This is the type of antenna being used in the RDRN prototype system to achieve spatial frequency reuse. The radiation pattern of an array antenna is controlled by the type of individual elements used, their orientations, their positions in space, and the amplitude and phase of the




Fig. 3. Spatial frequency reuse by beamforming.




Fig. 4. A linear array antenna.

Linear Eight Element Normalized Array Factor Plot




Fig. 5. Polar array factor of an equally excited eight-element linear array.

currents feeding them [12]. When the radiation pattern, called the array factor, of an antenna can be controlled in such a way as to selectively radiate energy in desired directions while simultaneously attenuating energy in others, this is called beamforming. Beamforming is effectively the result of far-field constructive and destructive interference patterns. A discussion of beamforming and how it applies to the RDRN prototype system follows.

1) *Beamforming via Multiple-Element Linear-Array Antennas:* A uniformly excited, equally spaced linear array provides the simple example. It consists of N active elements in a line spaced at regular intervals that are usually $1/2$ wavelength or less of their center operating frequency. Fig. 4 is an example of a linear array of the type used in the RDRN prototype system.

If the eight-element antenna depicted in Fig. 4 were fed eight identical signals of equal phase and amplitude, the theoretical broadside far-field pattern or array factor in Fig. 5 would result.

The normalized array factor appearing in Fig. 5 is given by

$$f(\psi) = \frac{\sin(N\psi/2)}{N \sin(\psi/2)} \approx \frac{\sin(N\psi/2)}{N\psi/2} \quad (1)$$

where N is the number of antenna elements and is equal to eight and Ψ is the angle off broadside [12]. It should be noted that the width of the main lobe decreases with increasing N . The reduction of the main lobe width is what allows for

spatial reuse. The cost involved in a large number of active antenna elements and the associated electronics needed to drive them must be weighed against the benefits of such an antenna pattern. The design goal for a multiple-beam system is to find the minimum number of elements that can support the number of beams required at the desired performance or signal-to-interference level. For the RDRN system, these cost and performance issues were balanced. The resultant tradeoff placed the number of antenna elements at eight.

2) *Beamsteering by Application of Complex Weights:* The broadside far-field antenna pattern shown in Fig. 5 is only as useful as the previously mentioned microwave dish example in that it, too, would have to be physically pointed at the mobile user to be of any use. A linear array antenna becomes more useful when it is used as a linear phased array. By manipulating the relative phase and amplitude of the RF waveform on a per-antenna-element basis, different interference patterns result, and it is possible to rapidly “steer” a beam in a particular direction. Fig. 6 is a block diagram of this process that depicts the application of N complex weights to an N element array.

The weights are applied to the same RF signal on a per-antenna-element basis, and the effect is the ability to spatially direct the beam in a desired direction without the need to physically rotate the array. Determination of the proper complex weights needed to effect the desired amount of

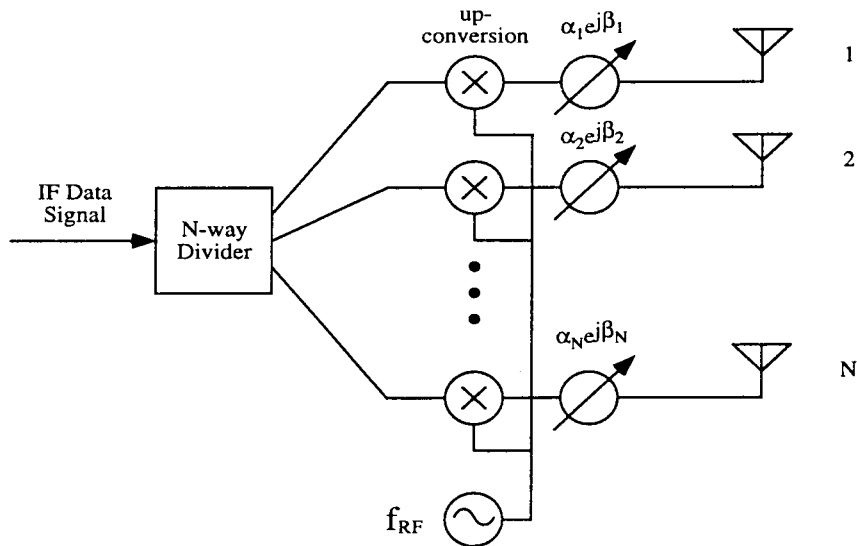


Fig. 6. Beamsteering by the application of complex weights.

beamsteering is the function of the beamforming algorithms that are described in the following sections.

3) *Simultaneous Multiple Beams at the Same Carrier Frequency*: The ability to rapidly steer a beam by electronic rather than physical means is an important advantage of a phased-array antenna. There is, however, another and more significant advantage to the RDRN system that comes from phased-array antennas. This is the ability to use them to form multiple transmit beams to be steered to multiple mobile end users from the same antenna at the same frequency. This is spatial reuse, and is one of the main goals of the RDRN proof-of-concept system.

The quantity and quality (defined by signal-to-interference ratio) of the multiple beams is controlled by many factors of which the number of antenna elements, their physical topology, and the proper application of complex weights are some of the most important factors. Fig. 7 depicts a way to form two simultaneous beams at the same frequency.

This system is the same as the single beam case appearing in Fig. 6 with the addition of a second set of inputs. The separate inputs from the two different beamforming processes are simply summed before the antenna. The system is linear, and because superposition holds for this electrical summation, the output from the antenna array will be the spatial sum of the individual antenna patterns. This process can be indefinitely expanded to form as many simultaneous beams as desired. There are, of course, some practical limits. As seen in the previous graphs of typical linear array antenna patterns, sidelobes exist in the form of $\sin(Nx)/Nx$ outputs. These sidelobes will cause interference to any other beams that are simultaneously being produced at the same frequency. Ways to reduce the level of interference include increasing the number of antenna elements, steering the beams further away from each other, and the proper application of complex weights to effect nulling patterns in the direction of other beams.

The proper application of the complex weights requires answers to at least the following questions.

1) How many beams need to be formed?

- 2) At what angles do each of the beams need to be steered?
- 3) What is the respective relative power contained in each beam?

The current beamforming algorithms used for RDRN were developed by Haas and Petr [13]. The algorithm takes the number of beams and the locations of users surrounding a transmitter, and calculates the steering angles and relative beam power such that a local maximum in the minimum signal-to-interference ratio is achieved.

4) *Beamforming for Arbitrary Modulation Schemes*: Beamforming is only of use in a communication system when its carrier frequency is modulated with some information-bearing signal. Beamforming in this case can be described for a generic modulation scheme as follows [14]. Let the signal samples be represented by the vector

$$s = (s_{-\infty}, \dots, s_{-1}, s_0, \dots, s_t) \quad (2)$$

where s is assumed to include the time index t implicitly. The most general form of modulation can be expressed as

$$x(t) = \text{Re}[\alpha(s) \exp(j2\pi f_{\text{IF}}t + \phi_0 + \Phi(s))] \quad (3)$$

where $\alpha(s)$ represents any form of amplitude modulation, f_{IF} is the IF carrier frequency, ϕ_0 is an initial phase, and $\Phi(s)$ represents any form of phase modulation. In most modulation schemes, $\alpha(s)$ will reduce to $\alpha(s_t)$ where t represents the current time index or, in phase-shift keying, further reduces to a constant α . $\Phi(s)$ can represent continuous-phase frequency-shift keying (CPFSK) schemes, of which minimum shift keying (MSK) is a special case. It can also represent all phase shift keying methods where the dependency reduces to $\Phi(s_t)$, or even coded modulation schemes like trellis-coded modulation.

Beamforming introduces an additional relative phase and amplitude of the signal to each antenna element. Thus, the modulated signal of a single user $x_i(t)$ is modified to

$$x_{ij}(t) = \text{Re}[a_{ij}(t) \exp(j\phi_{ij}(t)) \alpha_i(s_i) \exp(j2\pi f_{\text{IF}}t + \Phi(s_i))] \quad (4)$$

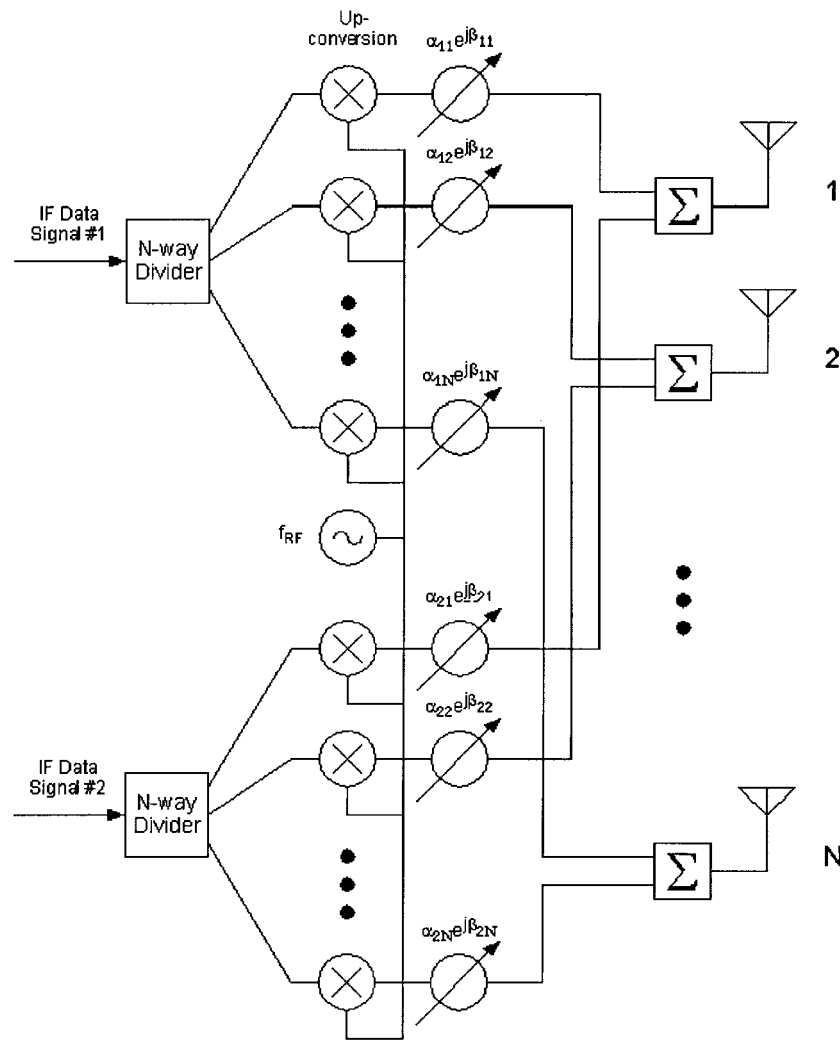


Fig. 7. Multiple beamforming by summing individual beams.

where subscript i refers to a user IF signal index and j refers to a specific element of a phased antenna array. Thus, the composite signal to an element is given by

$$X_j(t) = \sum_i \text{Re}[A_{ij}(t) \exp(j\phi_{ij}(t))\alpha_i(s_i) \exp(j2\pi f_{IF}t + \Phi(s_i))]. \quad (5)$$

Equation (5) is the per-antenna-element summing operation that appeared in Fig. 7. Implementing phase shifts and amplitude control on RF signals can be done at frequencies up to 10 GHz. The devices used to do this are expensive and prone to elaborate calibration requirements in the phased array. Manipulation of the phase and amplitude at RF is, however, unnecessary. This is because relative phase shifts at the intermediate frequency are maintained after upconversion. This makes the task of applying the complex weights much easier and less expensive.

D. Hardware Methods for the Application of Complex Weights

There are various analog and digital techniques for applying complex weights. Digital beamforming (DBF) techniques are the preferred method, and require some type of digital control

over the complex weights. Several analog techniques are briefly presented that make the digital design choices clear.

1) *Analog Control of Phase and Amplitude:* The current predominant technology used to implement attenuators at RF frequencies up to several gigahertz are gallium arsenide (GaAs) field-effect transistors (FET). They are used as voltage variable attenuators where the controlling voltage is applied to the gate of the FET. They are widely available from many manufactures in inexpensive matched 50 Ω surface mount packages. Unfortunately, they impart a phase shift that is a nonlinear function of the level of attenuation. Analog phase shifters suffer similar problems, and both require complex calibration schemes.

2) *Digital Control of Phase and Amplitude:* Digitally manipulating the phase of data at RF rates is computationally unfeasible, and as previously mentioned, unnecessary. Taking advantage of the fact that phase shifts and relative amplitudes are maintained after conversion to RF, it is possible to cost effectively modify the lower frequency IF signals. Directly generating the IF signals with an arbitrary waveform generator is the most straightforward way to apply the desired complex weights. The technique used for this function in the RDRN

TABLE II
PROTOTYPE DESIGN SPECIFICATIONS

Maximum number of simultaneous beams	4
Maximum transmit rate (only 1 beam active)	8 Mbps
Maximum receive rate	2 Mbps
Minimum number of Tx/Rx frequencies	4 (2 pairs)
Maximum Transmit Distance at a Bit Error Rate (BER) of 10 ⁻⁵	10 km
Modulation types	AM, PM, or QAM (No FM)

proof-of-concept system is summarized next and presented in detail in Section V. A comprehensive overview of digital beamforming techniques is found in [15].

E. Beamforming Transmitter Design Goals

The RDRN implements a proof-of-concept beamforming system. Table II lists the RDRN prototype design specifications.

1) *Radio Frequency Options:* The high data rates necessitate a large section of contiguous bandwidth of at least 20 MHz. The first publicly available sections of the RF spectrum with this capacity are the industrial scientific and medical (ISM) bands at 915 and 2400 MHz. These bands, however, are crowded by industrial equipment with which it is unlawful to interfere. Equipment in these bands typically transmits with less than 1 W of power. The need for RDRN to transmit up to 10 km requires higher power levels that would definitely cause interference to these bands.

The amateur radio service (ARS) band located at 1.27 GHz was chosen instead. It has a contiguous 60 MHz of bandwidth from 1240–1300 MHz, and a transmit power limit of 2000 W peak envelope power (PEP). The only requirement is that anyone transmitting in the bands be an FCC licensed amateur radio operator, and participants in the RDRN project have met this requirement.

2) *Intermediate Frequency:* Regardless of whether beamforming is done by analog or digital means, or if it is done at all, some type of IF information-bearing signal must be upconverted to RF. The process of upconversion is detailed later. During the analog mixing process, the sum and difference frequencies are produced in addition to other nonlinear products. Only the signal of interest, usually the sum, is passed on, amplified, and transmitted on the antenna. Extracting the signal of interest is done in the image-reject filter.

In order for the image-reject filter to be physically and cost-effectively feasible, the IF frequency must be sufficiently high as to place the difference frequency ($f_{IF} - f_{LO}$) far enough away from the sum frequency ($f_{IF} + f_{LO}$) that it can be filtered out. The higher the RF frequency, the higher the IF frequency required. Operation of the RDRN system in the ARS band at 1.27 GHz places the practical minimum IF frequency at 70 MHz, with 150 MHz being more the ideal. Digital beamforming places competing requirements on the IF signal frequency; however, the digital clock limits on digital-

to-analog converters (DAC's) constrain how high the IF can be. The IF chosen for the RDRN prototype was 150 MHz.

3) *Prototype Implementation Constraints:* Due to complexity-driven cost constraints and the prototype nature of the RDRN system, the transmitter section was limited to a maximum of four directional beams. It used a simple omnidirectional receive antenna to implement a full-duplex link. Similar constraints dictated the number of antenna elements implemented. Therefore, the multielement phased array used for the transmit link was limited to eight elements.

Several additional limitations are key to the design. First, no frequency modulation is used so that all data streams modulate the same frequency IF carrier. This constraint means that, regardless of the number of beams, end users, or the complexity of the modulation being used, only the amplitude and phase of the composite IF signal can be changed (i.e., the system is linear). Second, the modulation schemes and steering angles must change at a rate several orders of magnitude slower than the symbol rate. Finally, the symbol rate must be constant and synchronous across all of the beams being formed.

Restricting the RDRN system to the above constraints makes possible a hardware implementation that combines the task of beamforming and flexible IF modulation into a single cost-effective solution. This flexible beamforming software-radio-based IF modulator is now presented in detail.

F. A Flexible Beamforming Software-Based IF Modulator

Fig. 8 shows a schematic of a simple, single-channel IF waveform synthesizer. The components of the waveform synthesizer are a lookup table holding waveform samples, a register file and multiplexer to meet speed requirements, a DAC, and a bandpass filter. At the beginning of a transmitted symbol, the symbol value is used to retrieve signal values for the waveform from the lookup table. There are eight samples per cycle of a sinusoidal waveform, but due to symmetry, only four values need be stored in the lookup table. Retrieved values are stored in a register file to feed to the DAC at appropriate rates. In the prototype system, waveforms for symbols are retrieved every 1 μ s. Waveform samples are passed to the DAC every 0.1 μ s, resulting in sine waveforms at 10 MHz. These waveforms have multiple harmonics, and we select the harmonic at 150 MHz with the bandpass filter for the IF signal.

The IF waveform synthesizer is replicated for each antenna element. The lookup table for each antenna element is independent, and stores a waveform appropriate for the phase and amplitude assigned to that antenna element to implement a steered beam.

For a single BPSK modulated signal, the lookup table holds two waveforms of four samples each, one for the 0 symbol and one for the 1 symbol. For a single QPSK-modulated signal, the lookup table holds four waveforms of four samples each, one for the 00 symbol, one for the 01 symbol, one for the 10 symbol, and one for the 11 symbol. To form multiple beams, we need to consider whether each beam is on or off and the symbol sent on each active beam. Consider forming two beams, each using BPSK modulation. The lookup table entries are given in Table III.




Fig. 8. Schematic of a simple IF modulator.

 TABLE III
 EXAMPLE LOOKUP TABLE FOR TWO BPSK BEAMS

Beam #1	Data #1	Beam #2	Data #2	Wave-form
Off	X	Off	X	0
Off	X	On	0	$A_2 \sin(\omega t)$
Off	X	On	1	$A_2 \sin(\omega t + \pi)$
On	0	Off	X	$A_1 \sin(\omega t)$
On	1	Off	X	$A_1 \sin(\omega t + \pi)$
On	0	On	0	$A_1 \sin(\omega t) + A_2 \sin(\omega t)$
On	0	On	1	$A_1 \sin(\omega t) + A_2 \sin(\omega t + \pi)$
On	1	On	0	$A_1 \sin(\omega t + \pi) + A_2 \sin(\omega t)$
On	1	On	1	$A_1 \sin(\omega t + \pi) + A_2 \sin(\omega t + \pi)$

The size of the lookup table is determined by the number of beams plus the sum of the number of symbols per beam:

$$TableSize = 2^{(NBeams + \sum_{i=1}^{NBeams} NSymbol_i)} \quad (6)$$

where $NBeams$ is the number of formed beams and $NSymbol_i$ is the number of symbols per beam. In the RDRN the number of beams is restricted to four or fewer.

G. IF and RF Generation

As pointed out above, higher intermediate frequencies make rejecting the undesired frequencies easier prior to the final RF amplifier. The RDRN uses an IF of 150 MHz. Generating a 150 MHz sine wave with eight samples per cycle would require a DAC operating at 1.2 Gsample/s, a daunting design task. Options include generating a lower frequency IF, and then implementing two upconversion stages. An alternative method takes advantage of the second image frequency that appears after digital-to-analog reconstruction. By proper choice of the sampling frequency F_s and the fundamental frequency F_c , the second image can be placed at the desired IF, bandpass filtered, and then upconverted directly to RF. In order to repetitiously use the same samples over an entire symbol interval, the desired fundamental frequency F_c must be chosen so that the samples used to represent it repeat as an integer sub-multiple of F_s . The desired IF is then $(2F_s - F_c)$. F_s was chosen to be 80 MHz and F_c to be 10 MHz. The first frequency image appears at $F_s - F_c$ (80–10 MHz = 70 MHz) and the second frequency image appears at $2F_s - F_c$ ($2 \cdot 80$ MHz – 10 MHz = 150 MHz). The image at 150 MHz must be selected by a narrow-bandwidth bandpass filter. The quality of the alias is limited by the bandpass filter, of the DAC number of bits, and rise-time characteristics.

Even though there is a 32 dB loss, the signal-to-distortion ratio can be kept high by picking a converter with the necessary resolution. Most importantly, changes in both the amplitude

and phase of the 10 MHz signal are linearly translated into all of the image frequencies. This unconventional procedure eliminates the need for upconversion to a second IF frequency.

H. IF to RF Upconversion

The IF outputs from each of the eight beamform cards are upconverted in tandem so that the relative amplitudes and phase shifts between each of them are maintained all the way to the antenna. The 10 MHz IF signal from the DAC is filtered with a 150 MHz bandpass filter, buffered, and directly input into a double-balanced mixer that is driven by a local oscillator. Each of the eight separate RF mixers uses a local oscillator signal with the same phase. This is achieved by a passive eight-way power divider that distributes the same RF local oscillator to each mixer. The RF LO is a synthesized phase-locked oscillator with a 900–1600 MHz operating range. The LO has a parallel interface allowing its operating frequency to be changed under software control.

The output of the mixer is comprised mainly of the sum and difference frequencies produced by the mixer. It also contains LO feedthrough, in addition to other nonlinear products and sums. The signal of interest is the sum frequency that falls in the transmission band from 1240–1300 MHz. This signal is filtered out by an image reject bandpass filter, preamplified, and fed to the RF power amplifiers, 1240–1300 MHz Mitsubishi M67715 linear amplifier. The proper operation of the beamforming algorithm requires all chains of amplification to be linear. Therefore, the maximum output of the amplifier is limited to avoid compression of the output signal. In the prototype system, each RF amplifier delivered approximately 2 W to the antenna element.

The initial prototype operates between 1240–1300 MHz. However, the system is designed to be capable of operation at up to 2500 MHz by changing the RF LO and a single bandpass filter.

I. Antenna Array

The transmit section uses a linear, equally spaced, eight-element monopole phased array optimized to work at the center of the 23 cm (1240–1300 MHz) amateur band. A ground plane prevents the antenna from radiating in the backward direction. The field of coverage is limited to 180° in the proof-of-concept system.

J. Implementation of RDRN Transmitter

Fig. 9 illustrates the configuration of the completed RDRN transmitter. On the right is the interface to the host computer system. At the bottom is a VME rack holding eight IF cards,

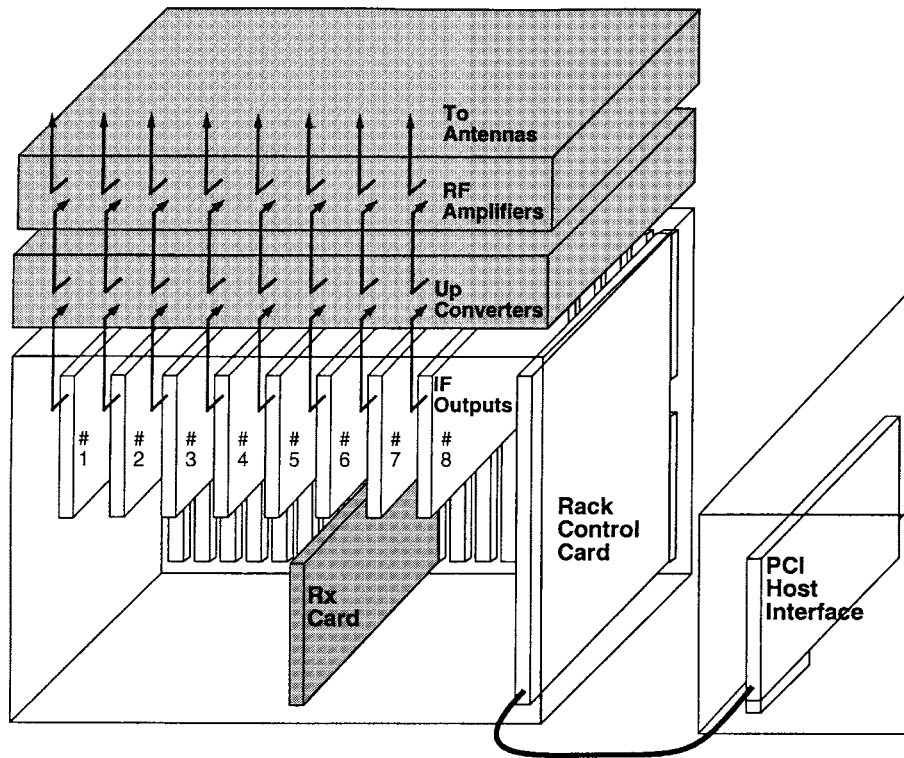


Fig. 9. Component level view of the hardware.

each containing a data buffer, beamforming lookup table, and the DAC for one antenna element. Each IF card feeds a separate upconverter, and each upconverter feeds a separate RF amplifier which, in turn, drives separate antenna elements. The entire stack occupied 20 vertical inches in a 19 in rack.

The results of forming a beam with a synthesized IF waveform are shown in Figs. 10 and 11. In Fig. 10, the element amplitude and phases are set equal to direct the beam toward the boresight. In Fig. 11, the element amplitude and phases are set to direct the beam 15° to the right of boresight.

V. DIGITAL RECEIVER IMPLEMENTATION

The RDRN system architecture calls for a digital beamformed receiver. However, due to a lack of suitable components at the time of the first prototype system, the team was only able to validate the functionality of a single digital receiver element. This section first outlines the design of a digital beamforming receiver and then presents a prototype implementation for a single element.

A. Design of a Digital Beamforming Receiver

In traditional analog beamforming, the RF signal at each antenna element is downconverted to IF, the phase is changed using analog phase shifters, the phase-shifted signals are summed, the summation signal is downconverted to baseband, and demodulated to a digital symbol stream. Beams formed in this manner are limited in the number of beams formed, exhibit sidelobe level control problems, have high loss, lack individual beam shape control, require complex construction, and are correspondingly heavy.

In a digital beamforming system, receivers digitize the received signals at the earliest possible time on a per-element basis. The element signals are processed by a special-purpose digital processor to form the desired symbol streams. Ideally, one digitizes the RF signal after amplification. However, practical considerations for conversion and processing bandwidth dictate conversion to an IF prior to digitization.

Fig. 12 shows a block diagram of a digital beamforming receiver. The received RF band is amplified, downconverted, and digitized at the IF with one section for each antenna element. A particular signal from the IF band is digitally downconverted to baseband, producing a complex (I and Q) signal. The complex signals are summed, and the sum is passed to a digital receiver for demodulation. One can replicate the digital downconversion to baseband, complex weighting, and digital receiver functions to form additional receiver beams. The RF amplification, downconversion, and analog-to-digital conversion can be shared by all digital receivers within the same band.

In digitizing the analog received signal, one analog-to-digital converter (ADC) is required for each antenna element. The final step in the demodulation process is to translate the sampled IF signal into a complex baseband signal. The complex baseband signal retains the magnitude and phase detail of the original real input signal. Low-pass filtering and decimation of this signal to a lower sample rate will reduce the number of samples needing to be processed. The digital downconversion (DDC) process is illustrated in Fig. 13.

B. Prototype Digital Receiver

The receiver consists of a conventional low-noise RF amplifier as the front end, followed by a downconversion to an




Fig. 10. Comparison of measured and computed antenna patterns at boresight.

IF signal. The IF signal is digitized, and the digital sample stream is fed to a programmable digital receiver. The receiver is capable of demodulating both BPSK and QPSK signals. The digital sections of the receiver are as follows.

The IF signal is centered at 70 MHz with a bandwidth of 5 MHz, allowing for a data rate up to 10 Mbit/s with QPSK modulation. The IF bandpass signal is digitized using bandpass sampling, which consists of sampling the signal at a multiple of the information bandwidth, but at a rate which is much lower than the low-pass Nyquist frequency.

The IF signal is digitized to 10 bits resolution at 25 Msamples/s with a Harris HI5703 ADC.

Following the ADC, the receiver uses a Harris HSP50110 digital quadrature tuner (DQT) and an HSP50210 digital Costas loop (DCL) as shown in Fig. 14. Both of these devices are programmable. The DQT selects a specific signal from the IF bandpass by mixing with a local digital oscillator and low-pass filtering the result. The output is a complex signal at baseband. The final stage of the receiver is a standard QPSK demodulator built with the Harris DCL. The DCL performs the baseband processing tasks required for the demodulation of BPSK, QPSK, and OQPSK waveforms. These tasks include matched filtering, carrier tracking, symbol synchronization, and automatic gain control (AGC). The output of the DCL is a digital bit stream.

The HSP50110 and HSP50220 chips are programmed for the particular operation of the receiver by loading the registers in these chips via the microprocessor interface. The main programmable parameters of the receiver are the following.

- 1) Local Oscillator (HSP50110): The center frequency and the phase offset of the local carrier can be programmed. The carrier frequency can be adjusted dynamically to account for the carrier phase offset effects.
- 2) Input Level Detector Threshold (HSP50110): This is used for IF AGC control. The level detector output is externally averaged to set the gain of an amplifier in front of the ADC which closes the AGC loop.
- 3) Baseband AGC (HSP50110): The level of the mixer output is adjusted by the baseband AGC. The baseband AGC provides coarse gain correction necessary to maintain the output of the HSP50110 at an appropriate signal level to maintain an acceptable dynamic range.
- 4) Low-Pass Filter Configuration (HSP50110): The low-pass filter can be configured as a one-stage or three-stage CIC filter.
- 5) Resampler Configuration (HSP50110): The resampler sets the output sample rate by controlling the sample rate of the decimation filters. The output sample rate can be adjusted dynamically to synchronize with baseband waveforms.
- 6) Matched Filtering (HSP50210): Both the raised root cosine (RRC) and integrate and dump (I/D) filters were used for matched filtering. The number of samples to integrate per symbol is programmable.
- 7) Baseband AGC (HSP50210): This is used to maximize the dynamic range, adjust for signal-to-noise variations, and maintain an optimal signal level at the input to the soft decision slicer.

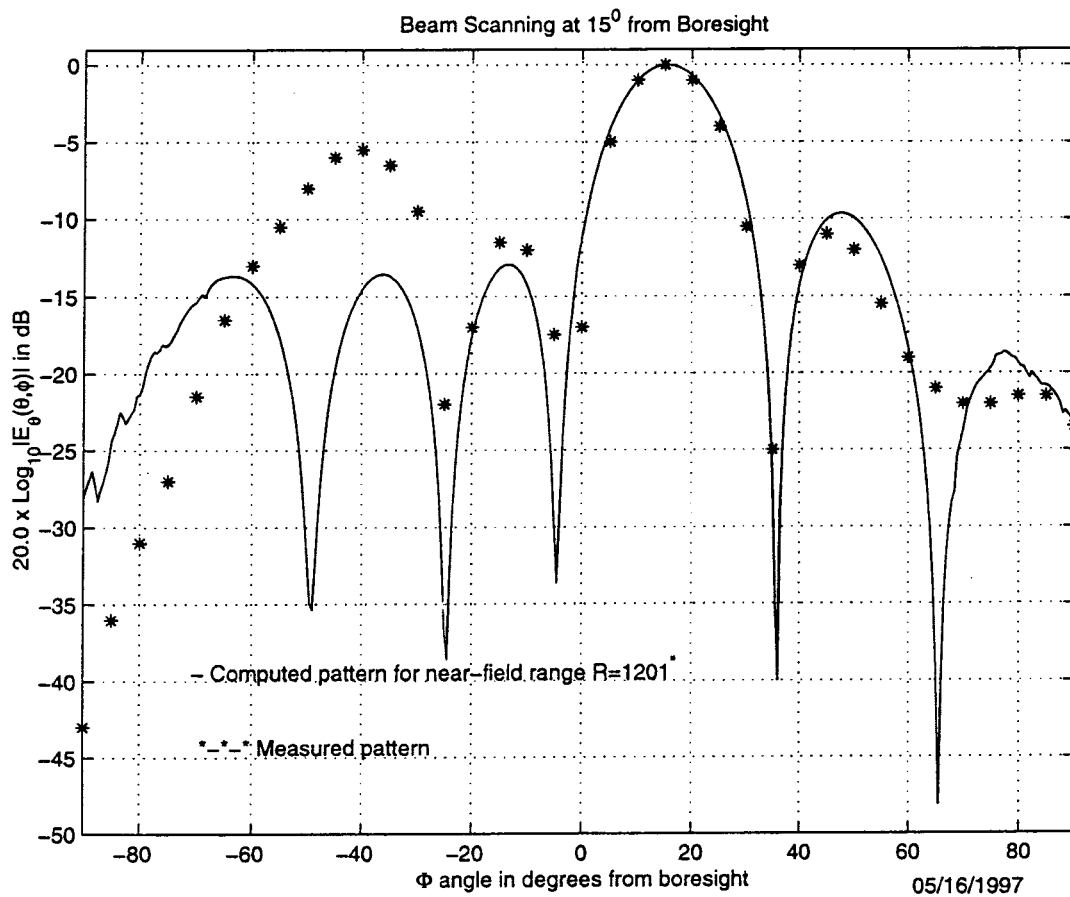
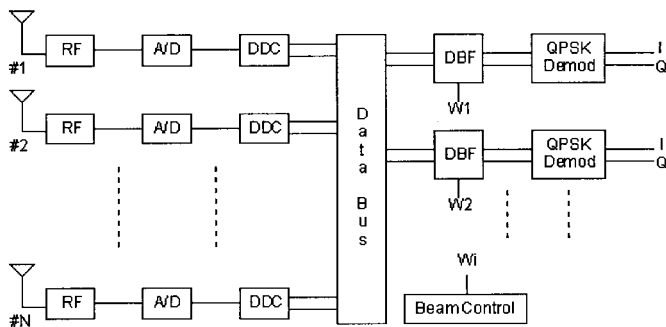


Fig. 11. Comparison of measured and computed antenna patterns 15° to the right of boresight.



RF: RF to IF Section
 DDC: Digital Down Conversion
 DBF: Digital Beamformer

Fig. 12. Block diagram of the digital beamforming receiver.

VI. PHYSICAL CONSTRUCTION OVERVIEW

The RDRN prototype system is contained within a mobile 19 in rack mount case. The ruggedized case is suitable for transport. Fig. 15 shows a component level view of the physical setup of a node. The host is at the bottom, and has interfaces to the transmitter and receiver and, for an edge node, interfaces to the wired infrastructure. The transmit cards are contained in a VME chassis along with a host controller card. Above the VME chassis is the transmitter upconverter and the RF amplifiers. The antenna feeds are on the back of the RF amplifier chassis.




Fig. 13. Digital down conversion block diagram.

The low-noise amplifier and downconverter for the receiver are mounted on the receiver mast to minimize interference and noise effects. The received IF is fed to the digital receiver mounted in the rack.

The transmit antenna array, receiver antenna, GPS antenna, and order-wire antenna are mounted to the 19 in rack, as shown on the right in Fig. 15.

VII. TEST RESULTS

The radio transmitter/receiver system developed for the RDRN project represents an initial step in the development and implementation of an end-to-end software radio system that adapts to the deployed environment. The first prototype incorporated beamforming at the transmitter, topology management, and link adaptation. The transmitter used digital

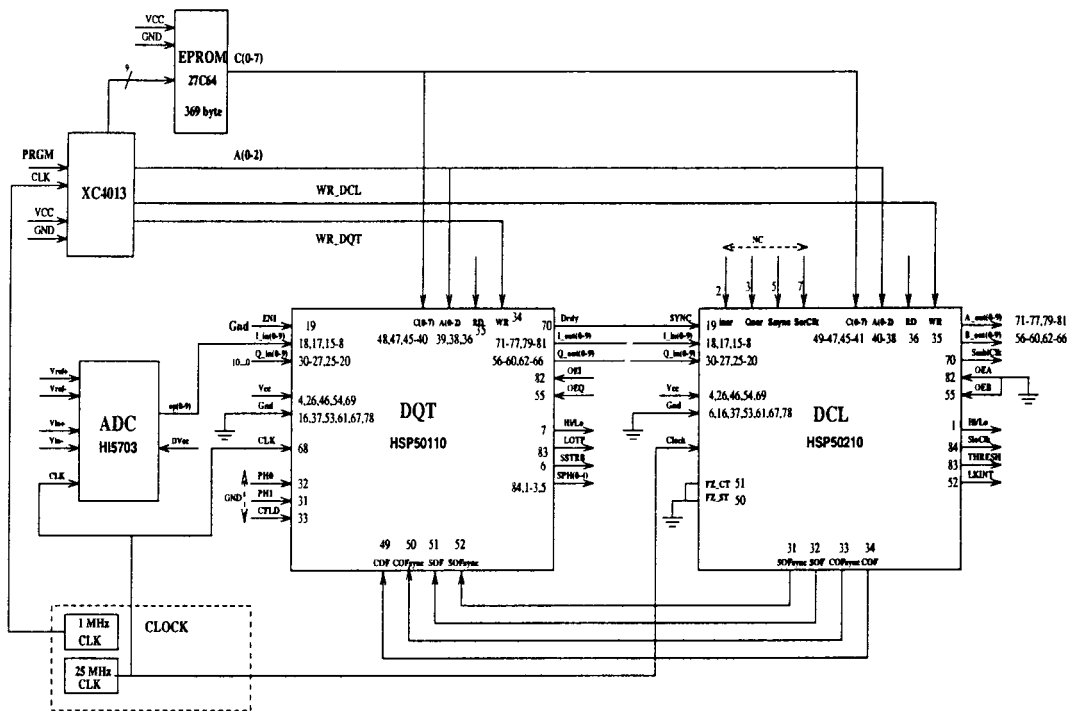


Fig. 14. Reconfigurable digital receiver.

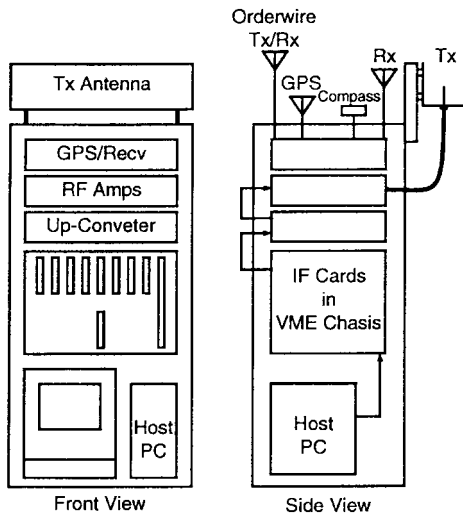


Fig. 15. The antenna and rack mount case.

waveform synthesis at IF to generate separate signals for each antenna element. The transmitter beamforming helped prove the concept of frequency reuse, while gaining experience in prototype fabrication with complex RF and signal processing hardware. The radio receiver system used the latest techniques in digital signal processing, including undersampling at the IF, followed by in-phase and quadrature downconversion to baseband via an application-specific numerical controlled oscillator (NCO). All demodulation and tracking algorithms were implemented digitally using programmable, application-specific integrated circuits.

The system performed well in field tests over a distance 6–7 km. The error rate was consistently within the link quality requirement. We tested the link under the following conditions.

TABLE IV
PREDICTED RECEIVED SIGNAL LEVELS FROM THE RDRN TEST LINK BUDGET

PATH MODEL	PATH LOSS	PREDICTED RSL
Free-space	111.42 dB	-91.31 dBW
Two-ray	113.14 dB	-95.03 dBW
Okumura	129.45 dB	-109.24 dBW

- 1) The distance between the transmitter and the receiver was approximately 7 km.
- 2) Four antenna elements were active. The transmitted power in each of the four antenna elements was 2.4, 2.55, 2.65, and 2.65 W, summing to a total of 10.25 W.
- 3) The received power level was observed to be -67 dBm or -97 dBW (-50 dBm and a -17 dBm receiver gain).
- 4) The four element transmit antenna gain was 10 dB.
- 5) From the test link budget system parameters, the received signal level (RSL) can be predicted as

$$RSL = P_t + G_t - L_s + G_r = 19.11 - L_s \quad (7)$$

Table IV shows the predicted receiver signal levels using three different path models.

VIII. PLANS

RDRN next steps could include beamforming at the receiver, extending both receiver and transmitter beamforming to cylindrical antennas for 360° coverage, and implementing multiple beams. Techniques for adapting the link to maintain consistently high throughput could be explored through the use

of algorithms in the radio processor which adjust power level, coding depth, antenna pointing, modulation type, and data rate.

ACKNOWLEDGMENT

The Rapidly Deployable Radio Network is a large project that could not be completed without the help of many staff, graduate, and undergraduate students. The following worked on the project over the years: S. Haas, S. Bush, S. Jagannath, T. Mai, Y. Lui, D. Chatterjee, V. Paulrajan, S. Gurrapu, H. Samra, P. Rajagopalan, C. Smit, J. Paden, T. Gallagher, and D. DePardo. The authors wish to thank them for their diligent work.

REFERENCES

- [1] R. M. Haralick and L. G. Shapiro, "The consistent labeling problem, Part I," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-1, Apr. 1979.
- [2] ———, "The consistent labeling problem, Part II," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-2, May 1980.
- [3] S. M. Haas and D. W. Petr, "A consistent labeling algorithm for the frequency/code assignments in a rapidly deployable radio network (RDRN)," TISL Tech. Rep. TISL-10920-04, Jan. 1995.
- [4] ———, "Maximizing the minimum signal-to-interference ratio in a wireless communications network," TISL Tech. Rep. TISL-10920-21, May 1996.
- [5] W. Almesberger, *Linus ATM*, Dec. 1995, [Online]. Available WWW: <http://lrcwww.epfl.ch/>.
- [6] M. Laubach, "Classical IP and ARP over ATM," *RFC-1577*, Jan. 1994, [Online]. Available WWW: <http://ds.internic.net/rfc/rfc1577.txt>.
- [7] W. Almesberger, *Linus ATM API*, July 1996, [Online]. Available WWW: <http://lrcwww.epfl.ch/>.
- [8] Internet Engineering Task Force, "NBMA next hop resolution protocol (NHRP)," 1995, [Online]. Available: gopher://ds.internic.net/oo/internet-drafts/draft-ietf-rolc-nhrp-04.txt.
- [9] "Q.Port. Portable ATM signaling software. Overview," Bellcore, Licensed Product LP-25, Release 1.3., July 1995.
- [10] R. Sanchez, J. B. Evans, G. J. Minden, V. S. Frost, and K. S. Shanmugan, "RDRN: A prototype for a rapidly deployable radio network," *ACM Mobile Computing Commun. Rev.*, vol. 2, Apr. 1998.
- [11] S. Bush, "The design and analysis of virtual network configuration for a wireless mobile ATM network," Ph.D. dissertation, Univ. Kansas, Aug. 1997.
- [12] W. Stutzman and G. Thiele, *Antenna Theory and Design*. New York: Wiley, 1981.
- [13] S. Haas and D. Petr, "Beam Matlab Code," TISL Tech. Rep. TISL-10920-03, Univ. Kansas, Jan. 1995.
- [14] B. Ewy and C. Sparks "A flexible beamforming digitally synthesized IF modulator," *ISCAS '96 Paper Proposal*, TISL Tech. Report TISL-10920-15, Univ. Kansas, Sept. 1995.
- [15] H. Steyskal, "Digital beamforming at Rome laboratory," *Microwave J.*, vol. 39, pp. 100-126, Feb. 1996.



Joseph B. Evans (S'80-M'83) received the B.S.E.E. degree in 1983 from Lafayette College, the M.S.E. degree in 1984, the MA degree in 1986, and the Ph.D. in 1989 Princeton University.

He is an Associate Professor of Electrical Engineering and Computer Science at the University of Kansas. He has been involved in major high-performance networking testbeds including the MAGIC gigabit testbed, the ACTS ATM Internetwork, broadband wireless efforts, and the RDRN project. He has served as a Researcher at the

Olivetti & Oracle Research Laboratory, the Cambridge University Computer Laboratory, USAF Rome Laboratories, and AT&T Bell Laboratories. His research interests include high-speed networks, mobile and wireless networking, and system implementations.



Gary J. Minden (S'75-M'78) received the B.S.E.E. degree in 1973 and the Ph.D. degree in 1982 from the University of Kansas.

He is a Professor of Electrical Engineering and Computer Science at the University of Kansas. He is a Principal Investigator on the MAGIC testbed, the Rapidly Deployable Radio Network, and a project within DARPA's Adaptive Computational Systems program. From June 1994 through December 1996, he was on leave at the Defense Advanced Research Projects Agency (DARPA) Information Technology Office, and served as a Program Manager in the area of high-performance networking systems. While at DARPA, he formulated and initiated a new research program in active networking. His research interests are in the areas of large-scale distributed systems which encompass high-performance networks, computing systems, and distributed software systems.



K. S. Shanmugan (F'87) received the B.E. degree from Madras University, the M.E. degree from the Indian Institute of Science, and the Ph.D. degree from Oklahoma State University.

He is the Southwestern Bell Distinguished Professor of Electrical Engineering and Computer Science. He held positions as a Postdoctoral Fellow at Oklahoma State University and the University of Kansas Center for Research, Inc. From 1973 to 1978, he was with the Department of Electrical Engineering, Wichita State University. From 1978 to 1980, he was a Visiting Scientist at Bell Laboratories. From 1985 to 1995, he also served in a number of leadership positions in the industry: as President of STA*R Corporation, Senior Vice President of Comdisco Systems, General Manager of the Alta Group of Cadence Design Systems, and Chief Technical Officer of Systems and Networks. During this time, he led the development and commercialization of modeling and simulation software for the design of communication systems and networks. His current interests are in the areas of wireless communications and computer-aided modeling and analysis of communication systems.

Dr. Shanmugan is the recipient of several teaching and research awards from the University of Kansas.



Glen Prescott (S'84-M'84) received the B.E. degree in 1974 from the Georgia Institute of Technology, the M.S.E.E. degree in 1976 from the University of Missouri, and the Ph.D. degree in 1984 from the Georgia Institute of Technology.

He is currently an Associate Professor of Electrical Engineering and Computer Science at the University of Kansas. He teaches and conducts research in the areas of digital signal processing, information theory, and digital communications. His areas of research emphasis are military communication systems, low probability of intercept/interference waveform design, spread-spectrum modulation techniques, and radio modem design. He served 23 years with the U.S. Air Force that included assignments as a Radar and Communication Engineer, Commander of a troposcatter communication station, and Assistant Professor of Electrical Engineering at the Air Force Institute of Technology (AFIT). He is Vice President of Lawrence Applied Research Corporation, which is involved in defense related consulting activities.



Victor S. Frost (S'75–M'82–SM'90–F'98) received the B.S. degree in 1977, the M.S. degree in 1978, and the Ph.D. degree in 1982, all from the University of Kansas.

He is the Dan F. Servey Distinguished Professor of Electrical Engineering and Computer Science. He is the Acting Director of the Information and Telecommunications Technology Center at the University of Kansas. He has worked with numerous corporations, including Sprint, NCR, BNR, NEC, Telesat Canada, AT&T, McDonnell-Douglas, DEC, and COMDISCO Systems. His research has been sponsored by government agencies, including NSF, DARPA, Rome Labs, and NASA. From 1987 to 1996, he was Director of the Telecommunications and Information Sciences Laboratory. His current research interests are in integrated communication networks, high-speed networks, communications system analysis, and simulation. He is currently principal investigator on the MAGIC and AAI high-speed wide-area testbeds.



Ben Ewy received the B.S.Co.E. degree in 1992 and the M.S.E.E. degree in 1994 from the University of Kansas.

He is an Engineer working on broadband service and network architecture at Sprint Corporation. From 1992 to 1996, he worked on the MAGIC gigabit testbed, and was codesigner of a 622 Mbit/s ATM interface utilizing FPGA's. From 1994 to 1996, he was a Research Engineer at the Center for Research, leading the design of the Rapidly Deployable Radio Network.



Ricardo Sanchez received the B.S. degree in 1989 from the Universidad Simon Bolivar, and the M.S. degree in 1993 from the University of Kansas.

He is a Doctoral candidate in electrical engineering at the University of Kansas. From 1989 to 1991, he worked at Infotrol in Venezuela in the area of supervisory control and data acquisition real-time systems. In 1991, he was awarded the SIVENSA graduate scholarship. In 1992, he joined the Research Group at the University of Kansas on the MAGIC gigabit testbed, and worked on

simulations of high-speed networks. From 1993 to 1995, he was a Research Engineer at Digital's System Research Center, Palo Alto, CA, building an ATM network. In the spring of 1995, he was an Assistant Professor at the Universidad Simon Bolivar, Venezuela. He is presently leading the software research group for the Rapidly Deployable Radio Network project. His research interests are in high-performance heterogeneous networks including mobile and wireless computing architectures.

Mr. Sanchez is a student member of the ACM and SHPE.



Craig Sparks (S'93–M'96) received the B.S.E.E. degree in 1994 and the M.S.E.E. degree in 1996 from the University of Kansas.

He is a Senior Project Engineer at the Center for Research, Inc. His graduate work was designing and building the radios for the Rapidly Deployable Radio Network. He worked in the wireless local-area network industry before returning to KU in June 1997. He is responsible for the radio design for Phase II of the Rapidly Deployable Radio Network, and oversees the engineering efforts wireless

communication projects.

K. Malinimohan received the B.E. degree in 1992 from Anna University, and the M.S.E.E. degree in 1995 from the University of Kansas.

He is an Engineer with Cadence, Santa Clara, CA. From 1992 to 1993, he was a Software Engineer developing application packages on VAX/VMS platform at Larsen and Toubro Ltd., India. From 1994 to 1998, he was a Research Engineer in the Information and Telecommunications Technology Center (ITTC) at the University of Kansas Center for Research working on the Rapidly Deployable Radio Network (RDRN) and Rapid Prototyping for DSP. His thesis addressed the design and implementation of a real-time algorithm to correct radar signals for remote sensing applications. His interests include wireless ATM networks, rapid prototyping, software radios, and VHDL synthesis.



James Roberts (S'65–M'66–SM'78) received the B.S. degree in 1966 from the University of Kansas, the M.S. degree in 1968 from M.I.T., and the Ph.D. degree in 1979 from Santa Clara University.

He is a Professor of Electrical Engineering and Computer Science and Associate Vice Chancellor for Research and Public Service at the University of Kansas. He began his career in 1966 with RCA working on the APOLLO mission to the Moon. In 1969, he joined ESL, a start-up company in California's Silicon Valley, and was promoted to

Laboratory Manager. Following TRW's acquisition of ESL, he relocated to Colorado in 1983, and became Manager of TRW Denver Operations. He was an Adjunct Lecturer in the Electrical Engineering and Computer Science Department at Santa Clara University from 1978 to 1983. From 1990 to 1997, he served as Chairman of the Electrical Engineering and Computer Science Department at the University of Kansas. He teaches and conducts research in wireless communications and information theory. He serves on the Board of Directors of Kantronics, Inc., the KU Center for Research, Inc., and the Silicon Prairie Association.



Richard Plumb (S'81–M'82–SM'94) received the B.S. degree in 1982, the M.S. degree in 1985, and the Ph.D. degree in 1988, all from Syracuse University.

He is an Associate Professor of Electrical Engineering at the University of Kansas. From 1982 to 1985, he was employed by the General Electric Company, Syracuse, NY, where he worked on various radar and sonar systems. In 1985, he graduated from the General Electric Advanced Course in Engineering. Between 1985 and 1988, he was a Graduate

Research Assistant at Syracuse University, and for the 1988 academic year, he was a Visiting Assistant Professor. In 1989, he joined the faculty at the University of Kansas. During 1996, he was a Visiting Scholar at the University of Queensland, Brisbane, Australia. He is a member of the Radar Systems and Remote Sensing Laboratory at the University of Kansas. His primary research interests are in electromagnetics, ground-penetrating radar, and inverse scattering.



Dave Petr (S'75–M'78–SM'90) received the B.S.E.E. degree in 1976 from Southern Methodist University, the M.S.E.E. degree in 1978 from Stanford University, and the Ph.D. degree in 1990 from the University of Kansas.

He is an Associate Professor in the Electrical and Engineering and Computer Science Department and the Information and Telecommunication Technology Center at the University of Kansas. His current research interests are focused on the design and analysis of network resource management and congestion control mechanisms, voice transport over asynchronous transfer mode (ATM) networks, and pricing for multiservice networks. From 1977 to 1986, he was employed by AT&T Bell Laboratories, working in the areas of digital PBX's, digital transmission systems, packet communication networks and protocols, and speech coding.

Dr. Petr was awarded three patents related to his work, and was a Bell Laboratories nominee for the Eta Kappa Nu Outstanding Young Electrical Engineer Award in 1984. While pursuing the Ph.D. degree, he received the 1987 IEEE Frank A. Cowan Scholarship for graduate studies in communications. He is a member of Eta Kappa Nu, Tau Beta Pi, Sigma Xi, and the American Society for Engineering Education (ASEE).