

◆ **RADIOSTAR: Providing Wireless Coverage Over Gigabit Ethernet**

Zhengxiang Ma, Michael Zierdt, John Pastalan, Arnold Siegel, Tod Sizer, Adriaan J. de Lind van Wijngaarden, Praveen R. Kasireddy, and Dragan M. Samardzija

RadioStar is a novel system developed by Bell Labs to provide low cost, cellular wireless coverage primarily targeted for enterprises, public spaces, and residences. The underlying concept is to digitize the output radio signals of a base station and to transport these signals over a Gigabit Ethernet network to multiple low-cost radio frequency radiation points. The received signals are returned through the network to the base station. This paper presents the architecture, implementation, and demonstration of RadioStar. It is shown how handoffs and precise frequency synchronization and timing control over a Gigabit Ethernet network can be established in such a way that the most stringent requirements of cellular communication air interfaces are satisfied. By converging cellular signal transport onto Gigabit Ethernet data networks, RadioStar can achieve substantially lower hardware and installation costs, while providing higher performance, more flexibility, and better operation, administration, and maintenance than current solutions. RadioStar concepts are being used in the ClearFill Star in-building product developed by the RFS business division of Alcatel-Lucent and currently being deployed commercially. © 2009 Alcatel-Lucent.

Introduction

The cellular wireless industry has witnessed a tremendous growth in the past two decades. In the last few years, the growth of wireless revenue has been fueled by long distance toll displacement and, to a lesser degree, residential wireline displacement. While wireless data services are growing rapidly, there are new opportunities for growth, especially in the lucrative enterprise voice and data markets. Wireless service providers (WSPs) are in a unique position to

pursue the enterprise voice and data market because of their capability to provide full mobility with a single personal device at a competitive cost. However, this requires WSPs to provide reliable, high capacity coverage in enterprises, public spaces, and residences.

In-building coverage is a challenging radio frequency (RF) engineering problem [3, 7–10]. Walls, floors, windows, and internal partitions inside the building create a very complex environment for RF

Panel 1. Abbreviations, Acronyms, and Terms

3G1X—CDMA2000 first evolution	MIB—Management information base
ADC—Analog-to-digital converter	MSC—Mobile switching center
BSI—Base station interface	NTP—Network Time Protocol
BTS—Base transceiver station	OA&M—Operation, administration, and maintenance
CAT 5—Category 5	PCS—Personal communication system
CD—Collision Detection	PHY—Physical layer
CDMA—Code division multiple access	PMA—Physical medium attachment
CSMA—Carrier sense multiple access	PN—Pseudo-random noise
DAC—Digital-to-analog converter	ppb—Parts per billion
DAS—Distributed antenna system	ppm—Parts per million
DC—Direct current	PSMM—Pilot strength measurement message
EVDO—Evolution data optimized	RAM—Random access memory
FPGA—Field programmable gate array	RF—Radio frequency
GPS—Global Positioning System	RRH—Remote radio head
GSM—Global System for Mobile Communications	SNMP—Simple Network Management Protocol
IEEE—Institute of Electrical and Electronics Engineers	SPI—Serial-parallel interface
IF—Intermediate frequency	TCP—Transmission control protocol
IP—Internet Protocol	UMTS—Universal Mobile Telecommunications System
I/Q—In-phase/quadrature	VLAN—Virtual LAN
LAN—Local area network	WSP—Wireless service provider
MAC—Medium access control	

propagation. Absorption and scattering tend to significantly attenuate the RF signals of a cellular base transceiver station (BTS), often resulting in spotty coverage inside buildings, in particular at lower floors [7] and near elevators. On higher floors, there is often a significant difference between line-of-sight and non-line-of-sight areas. In-building coverage can be improved by increasing the output power of nearby base stations, by using directional antennas to point at the building, or by using repeaters. However, increasing RF power is expensive, and due to the high attenuation in buildings, the improvement in coverage will be limited.

A common approach is to extend the cellular infrastructure within buildings to provide coverage. In particular, there exist a variety of distributed antenna system (DAS) solutions that provide RF radiation points inside the building to distribute the RF signals from the BTS throughout the building for better coverage. The RF signals originate directly from the radio

card of the BTS or from its antenna. A passive DAS distributes the RF signals directly through leaky RF coaxial cables using a high power input. This technique is effectively used in tunnels and mines, and it has been installed in buildings also. It requires careful design and planning to achieve the desired output power level at individual remote antennas [8], and it is very difficult to change the configuration after installation. An active DAS uses a fixed access network to connect the BTS with the remote antennas. It converts the RF signals at the BTS into an intermediate format prior to distribution over the network and reconstructs the RF signals at the remote antennas. As an example, active DAS systems may use analog distribution over fiber by directly modulating the RF signals, or analog distribution over category 5 (CAT 5) cables by downconverting the RF signals to intermediate frequency (IF) signals. Besides high equipment and installation cost, analog DAS solutions have the disadvantage that all gains and losses in the analog

distribution network have to be accurately accounted for, as they directly affect the RF link budget. Digital DAS solutions avoid that particular problem, but they require a fiber-based distribution network to support the high data rates needed, thus incurring high equipment and installation cost.

Another approach is to install additional base stations that have a small form factor and lower capacity to effectively provide coverage in a small area (picocell). Multiple picocells can be deployed throughout the building to provide coverage and additional cellular wireless capacity. The picocells could be backhauled through an Internet Protocol (IP) connection over an existing data network. Clearly, picocells can support very high cellular density. However, picocells typically have a fixed air interface and a fixed capacity, and their complexity leads to higher cost. If the need for capacity increases, additional hardware will have to be installed, and support for upgradability with programmable hardware incurs higher costs. A high density of picocells may also cause intricate handoff problems.

In this paper, we discuss the concept, architecture, and implementation of RadioStar, a novel system developed by Bell Labs to provide low-cost cellular wireless coverage over Gigabit Ethernet. RadioStar can be regarded as an advanced digital DAS system with a Gigabit Ethernet transport layer. However, its features go far beyond those of conventional DAS systems. In particular, by using a standardized packet data network to distribute and aggregate digitized radio signals, RadioStar is able to leverage a whole industry for large volume, low cost, high performance components, and advanced features. For instance, Gigabit Ethernet runs over both fiber and CAT 5 cable readily and, in the case of CAT 5 cable, can supply end points with direct current (DC) power up to 12.95 W/25 W, as specified in two Institute of Electrical and Electronics Engineers (IEEE) standards: the IEEE 802.3af and IEEE 802.3at power-over-Ethernet standards [4, 6]. In addition, using a Gigabit Ethernet network also enables the reuse of the operation, administration, and maintenance (OA&M) architecture developed for data networks. The use of a standard Gigabit Ethernet network also supports

convergence of services. For example, RadioStar's Gigabit Ethernet transport network can simultaneously be used as backhaul for 802.11 access points for high speed wireless data and other applications. RadioStar can also be used in combination with picocells. Since a picocell can provide capacity at a much lower cost than a macrocell base station, and RadioStar can provide the most cost-effective reach and coverage, a combination of both holds the promise for the best in-building solution.

The next section describes the concept and architecture of RadioStar in detail and presents techniques to satisfy the most stringent frequency synchronization and timing requirements for cellular communication air interfaces. Applications of RadioStar technology will be presented next. In the final section, the development of a prototype and subsequent demonstrations will be discussed.

RadioStar Architecture

The key concept of RadioStar is to use Gigabit Ethernet [1, 2] to distribute digitized output radio signals of a BTS to remote radio heads (RRHs), and to aggregate the received signals through the network to the BTS. The rationale for using Gigabit Ethernet is its support for high data rates and low latency, which are key elements for cellular applications. For example, to transport one Universal Mobile Telecommunications System (UMTS) carrier of 5 MHz, a sampling rate of 7.68 MHz is adequate. With a 10-bit resolution, the raw data rate required to transport the in-phase/quadrature (I/Q) signal is 153.6 Mbps. Similarly, single carrier code division multiple access (CDMA) would require a raw data rate of 38.4 Mbps. The air interfaces of cellular systems are designed with the latency of signal propagation over the air in mind. A typical cell radius is less than 10 miles, and thus the typical round trip delay is then less than 100 μ s. It is therefore best to design the RadioStar architecture such that the distribution delay is of the same magnitude.

RadioStar leverages commonly available data network infrastructures. A common local area network (LAN) infrastructure is illustrated in **Figure 1**. A LAN in a typical office building is laid out in a star topology. The center of the star is the main telecommunication

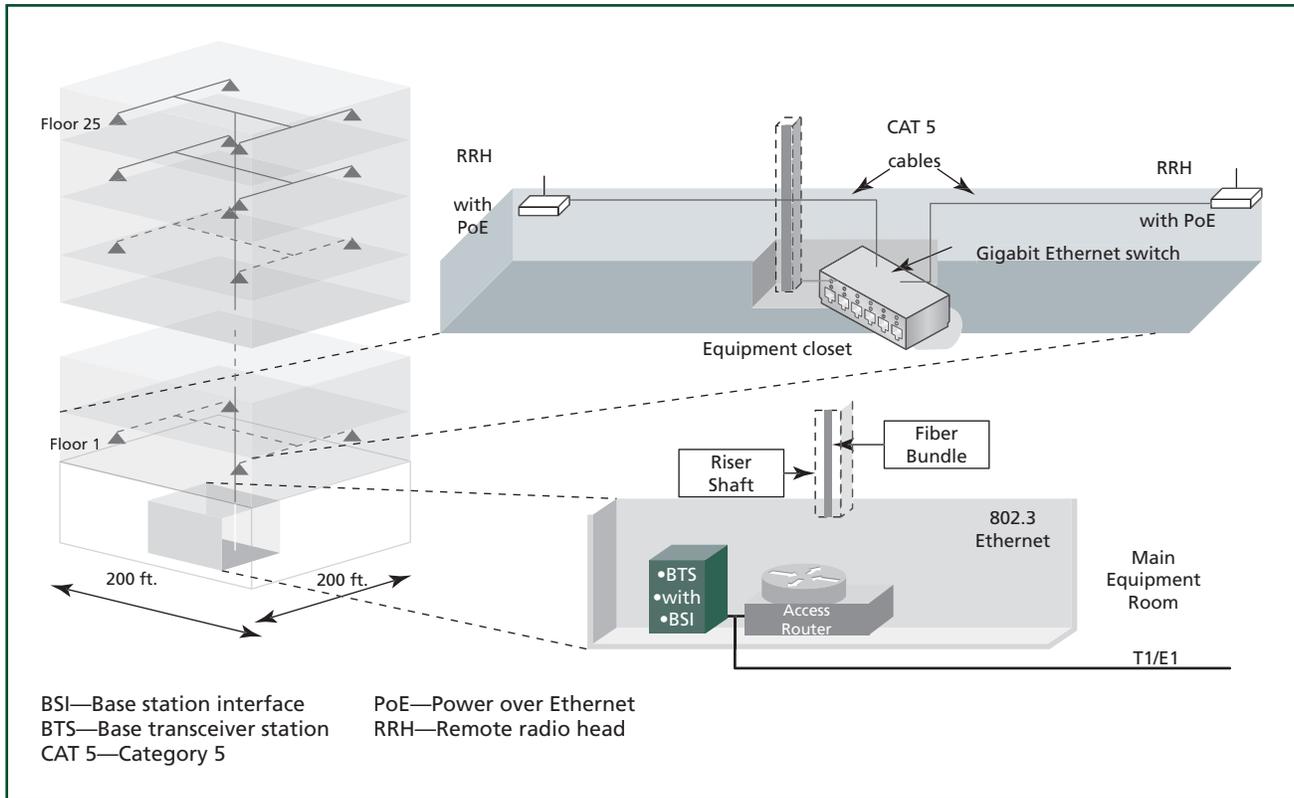


Figure 1.
Typical deployment topology of a RadioStar network.

equipment room, where switches and routers are located, and where it is connected to external networks. From this central location, network cables (fiber or CAT 5) branch to multiple telecom closets located throughout the building. A CAT 5 cable consists of four pairs of unshielded twisted pairs, which are used simultaneously to transmit and receive. Telecom closets are usually located such that the longest CAT 5 cable is less than 100 meters (m) long in order to support 100Base-T or 1000Base-T Ethernet [5]. Gigabit Ethernet over CAT 5 cable is defined in the IEEE 802.3ab standard as 1000Base-T. For longer distances, Gigabit Ethernet over fiber can be deployed using 1000Base-LX as defined in IEEE 802.3z. This standard specifies distances up to 3 miles, but many 1000Base-LX vendors guarantee much larger distances, up to 50 miles.

RadioStar uses exactly the same LAN topology as illustrated in Figure 1: a base station interface (BSI) card connects to a BTS located in the main telecom room; fiber or CAT 5 cables connect the BSI to multiple wiring closets where Gigabit Ethernet switches are installed.

From there, CAT 5 cables connect the switches to the individual RRHs, which are installed at selected locations throughout the building. The RRHs have integrated antennas, or optionally one or more external antennas can be connected to an RRH through a short RF cable. Power for the RRHs is provided over CAT 5 cables using power-over-Ethernet as specified in [4] and [6]. If available, existing CAT 5 or optical fiber cables can be reused to reduce the total installation cost. The interface between the medium access control (MAC) layer and the physical (PHY) layer is the Gigabit media independent interface. It is common across different physical media such as CAT 5 or optical fiber.

By installing fiber, the RadioStar architecture can be used to support multiple buildings in a campus environment with a single centrally located BTS. It can be further extended to support hoteling of BTSs for outdoor macrocell networks. Hoteling is a cellular network architecture concept in which the digital baseband portion of the base stations covering a geographical area is concentrated at a single location and only the RF

portion of the base stations remains at the cell sites. It has significant potential in saving operating expenses and enhancing the capacity and performance of the system. The functionality of each of the components in the RadioStar network is described in detail in the following sections, as well as the schemes used to transport the cellular signals over a switched Gigabit Ethernet network. A number of Gigabit Ethernet features are used to ensure that the stringent timing and latency requirements of cellular systems are satisfied.

Base Station Interface

A connection to the BTS can be established through either an RF or a digital baseband interface. An RF interface allows RadioStar to support any manufacturer's BTS. The physical connection can be made through an RF cable or even over the air. A digital baseband interface will be specific to the brand or model of the BTS, as such interfaces are not standardized. However, the use of a digital baseband interface does eliminate the need for an analog radio and leads to a higher level of integration and performance. The main tasks of the BSI are to convert the downlink radio signals (BTS to mobile device) into Ethernet packets and to send these to the RRHs, and to convert incoming Ethernet packets into uplink radio signals (mobile device to BTS). A digital baseband interface has the advantage that all conversions are in the digital domain. An RF interface requires the use of a down-converter followed by an analog-to-digital converter (ADC) to transform the downlink RF signal into a digital baseband signal, and a digital-to-analog converter (DAC) and an up-converter to produce an uplink RF signal. A BSI with a digital interface can easily support multiple sectors, as long as it has access to the baseband I/Q signals of those sectors through the backplane and there is enough Ethernet bandwidth to support all the packet streams that are associated with each sector. Support of multiple sectors with a single RF interfaced BSI would require multiple RF transceivers. However, supporting multiple sectors with a single BSI is not absolutely necessary, because a RadioStar network can easily support multiple BSIs. For the reverse link, multiple packet streams from multiple RRHs have to be combined by summation to form a single I/Q signal stream for the BTS receiver.

For low RF bandwidth applications, the summing function can be located in the BSI.

Remote Radio Head

The RRH converts incoming Ethernet packets into downlink RF signals and radiates these signals to mobile devices in its coverage area. At the same time, it receives uplink signals from these mobile devices and converts these into Ethernet packets. To function properly, the RRH has to be able to synchronize its frequency and timing to those of the BSI/BTS. Some cellular air interfaces have stringent requirements on the frequency and timing of the radiated RF signal. For example, CDMA standards typically require the frequency accuracy to be within 50 parts per billion (ppb), and the timing to be less than $3\ \mu\text{s}$ of the universal time, as acquired locally with a Global Positioning System (GPS) receiver. The RRH is powered over Ethernet to eliminate the requirement of a separate power supply for installation, which provides a significant cost saving. To strive for low equipment cost, the RRH uses handset technology for its RF transceiver function. The challenge is to swap the transmit and receive frequencies of a handset RF chip set to match those of a base station. For low RF bandwidth applications, besides the BSI, the summing function can be located in the RRHs. A group of RRHs send their packets to a designated RRH, which combines them with its own reverse link data to form a signal stream to be sent to either the BSI or another summing RRH. By locating the summing function in the BSI and RRHs, a commercially available off-the-shelf Gigabit Ethernet switch can be used in the RadioStar network. The BSI and RRH have many functions in common. Therefore, it is possible to produce both, using a single common hardware design. For example, the digital portion of the BSI and RRH can be substantially the same; only different firmware is needed. This can lead to substantial savings in the design and production cost.

Gigabit Ethernet Switch

RadioStar utilizes layer 2 switching exclusively. The packets that contain the cellular signals are virtual LAN (VLAN)-tagged. The VLAN is used to identify the sector and/or carrier that is being transported over

the network. Additionally, within the VLAN frame structure the 3-bit priority field is set to its maximum value for all cellular traffic packets to ensure that the packet will be processed as soon as possible. An Ethernet switch performs the function of directing packets to their destination by inspecting the destination MAC address in the header portion of the packets. In a RadioStar network, Gigabit Ethernet is used because of the required data rate and latency. The switches therefore need to support VLAN and priority switching. Such switches are usually in the category of managed switches. For low RF bandwidth applications, where the summing function is implemented in the RRHs and BSI, commercial off-the-shelf Gigabit Ethernet switches can be used. Through the course of this project, there has been a tremendous price reduction in low-port-count commercial Gigabit Ethernet switches. However, for high RF bandwidth applications or extremely low latency applications, the summing functions have to be located inside the switches. In such cases, custom-designed Gigabit Ethernet switches are required. At this point, low-port-count commercial off-the-shelf Gigabit Ethernet switches that support power-over-Ethernet are becoming available, which displace the need for a mid-span injection

unit to inject DC power into the CAT 5 cables between the switch and the RRHs. A custom Gigabit Ethernet switch with built-in power-over-Ethernet capability is highly desirable to support RadioStar applications.

Power-Over-Ethernet

Recently, power-over-Ethernet for CAT 5 cable was specified in the IEEE 802.3af standard [4] and IEEE 802.3at standard [6]. Two new entities are added to the physical layer, namely, power sourcing equipment and a powered device. For 1000Base-T, all four pairs are used for data transmission, and DC power is injected into two of the four data pairs through the center taps of the magnetic transformers. Different classes of power operation are defined in IEEE 802.3af and IEEE 802.3at [4, 6]. The default classification is Class 0. It provides a maximum of 12.95 W delivered to the powered device, which is sufficient for in-building radio heads in the RadioStar application.

RadioStar Network

The RadioStar network is a standards-conforming switched Gigabit Ethernet network as illustrated in **Figure 2**. The BSI and the RRHs can be regarded as stations that are connected to the Gigabit Ethernet network. Additional standards-conforming Ethernet

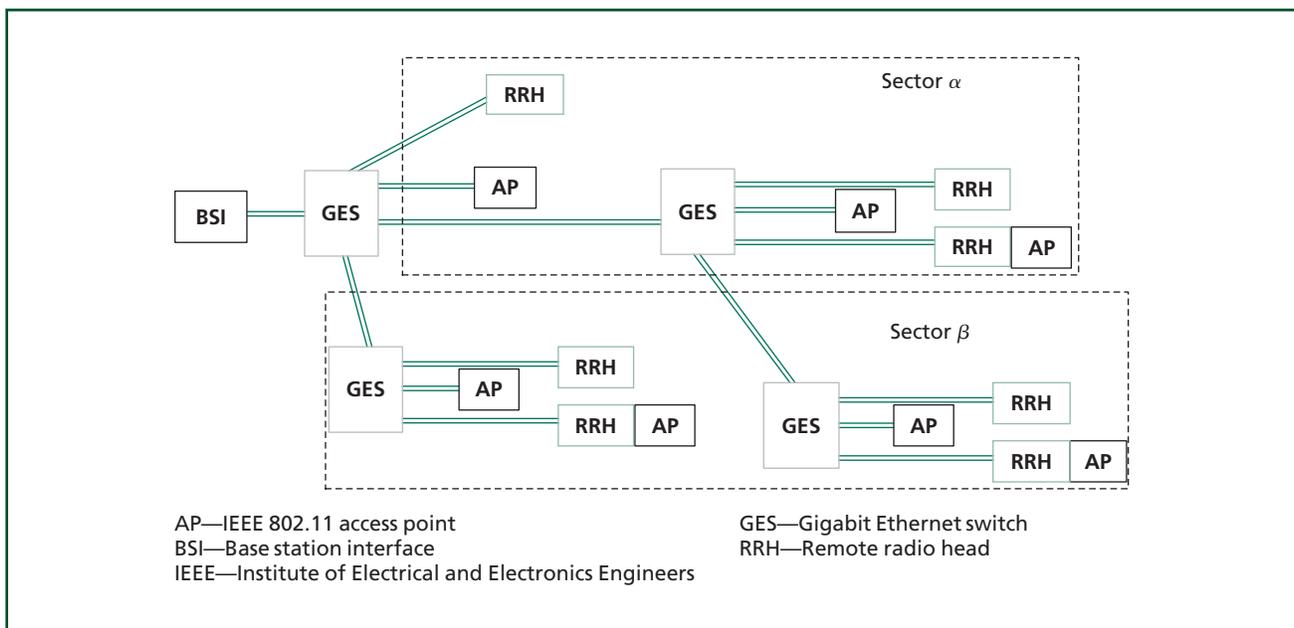


Figure 2.
RadioStar network topology.

equipment can be supported on this network, e.g., IEEE 802.11 access points. The access points can be standalone and connected to the switches, or can be integrated with the RRHs and share the same CAT 5 connection for power and backhaul. The data packet traffic generated by these additional devices and OA&M in the IP layer can coexist with packets carrying the cellular signals. The BSI transmits a stream of data packets carrying the downlink cellular signal from the BTS to the mobile device, to be simulcast to a group of RRHs. The data packets use a multicast address as the destination address and the BSI's MAC address as the source address.

By assigning the RRHs and BSI to a unique VLAN and configuring VLAN membership of the switch ports, the downlink data packets will only be broadcast to RRHs in the VLAN group. The VLAN is also instrumental in implementing another highly desirable feature, namely, sectorization. In regular cellular networks, the network capacity can be increased by sectorizing the antenna. An omnidirectional antenna sends the downlink signal over the entire coverage area of the cell and receives signals from all the users in its coverage area. By replacing the omnidirectional antenna with a number of directional antennas that divide the coverage area into multiple sectors with little overlap, each of the directional antennas will be sending out different RF signals to different users and receiving signals from different users simultaneously, thus increasing the total capacity of the system. The same concept can be applied in the in-building space as well. Initially, when the capacity need is low, the whole building can be configured as a single sector by making all the RRHs simulcast the same signal. As the capacity demands grow, RRHs can be reconfigured into a few sectors, with each sector simulcasting its own unique signal. This can be easily accomplished by assigning each sector-group of RRHs their own unique VLAN identifier and updating the VLAN membership of the switch port, as illustrated in Figure 2. Of course, the base station has to be upgraded with additional radio cards to support multiple sectors and the increased capacity, and the BSI has to create and send out additional packet streams containing different cellular signals destined for the additional sectors.

The packet streams will be tagged with their unique VLAN identifiers. If the BSI is only capable of supporting a single sector, then additional BSIs can be connected to the switch and provide packet streams for the additional sectors. This is all accomplished through software reconfiguration without having to touch the individual RRHs and switches, which is a tremendous advantage over regular DAS solutions. The flexibility afforded by Ethernet also enables more advanced applications, such as dynamic sectorization. The coverage area of the sectors can be dynamically changed over time to match the fluctuating traffic density in a building. Areas with high traffic density can be covered with more sectors, while areas with low traffic density can be covered with a single sector. As the traffic density changes over time, the boundaries of the sectors can be moved by regrouping the RRHs. All configurations and reconfigurations of the RadioStar network are a function of the OA&M architecture. This is another advantage Ethernet provides. The OA&M can be implemented by leveraging existing data network management solutions such as Simple Network Management Protocol (SNMP) running over Transmission Control Protocol/Internet Protocol (TCP/IP). More details on OA&M will be presented later.

So far, we have really only been concerned with the transport of the downlink cellular signal packets. The transport of the reverse link signal packets has additional interesting problems to overcome. It requires new functions in the RadioStar network beyond what standard Ethernet provides. **Figure 3** illustrates an example RadioStar network with four RRHs grouped into two sectors. The BSI depicted generates two packet streams colored gray and green for the two sectors. By properly setting up the VLAN, the switch automatically broadcasts the two streams to the two groups of RRHs. On the reverse link, each RRH generates a packet stream carrying the sampled cellular signal received by its antenna. If all the packet streams from the RRHs go to the BSI, it creates two problems. First of all, since a large number of RRHs are potentially supported by a single BSI, the data rate of the combined packet streams will certainly overwhelm the capability of the Gigabit Ethernet connection to

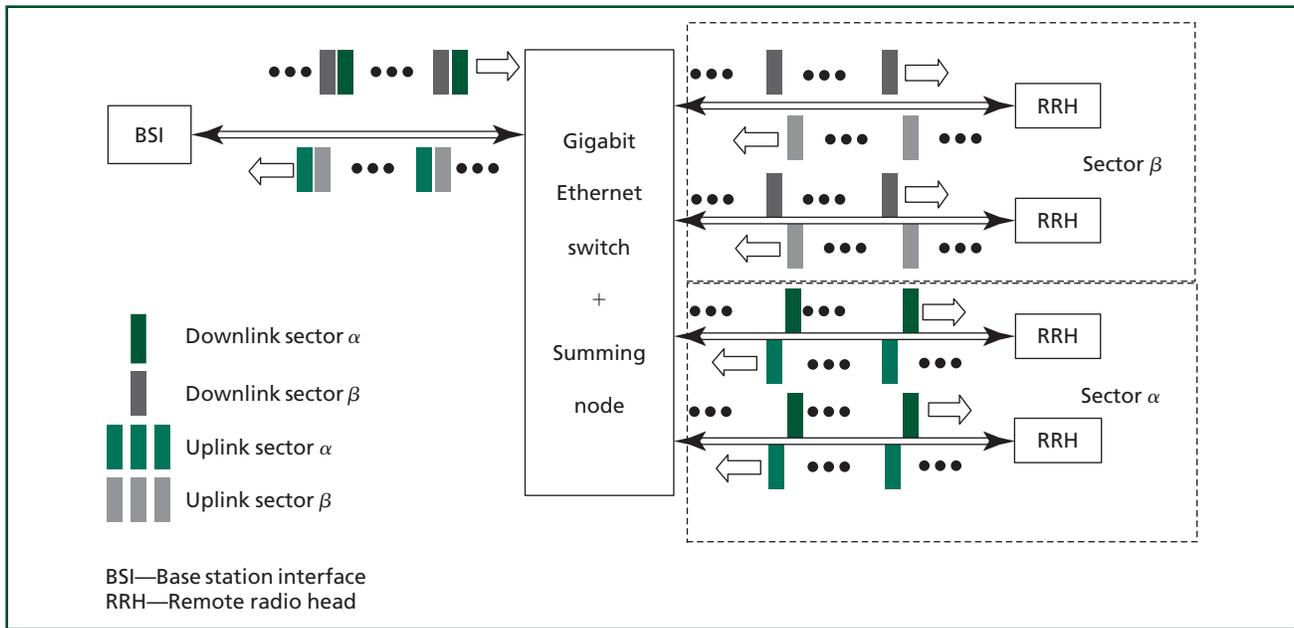


Figure 3. Downlink and uplink cellular signal packet stream transportation.

the BSI. Secondly, the base station has only two receivers per sector (for receive diversity, in typical in-building deployment, only one of the receivers is used). It can only deal with at most two different received cellular signals per sector. The reverse link packet streams of a sector must be combined into a single stream per sector before being passed to the base station. This is equivalent to combining the RF signals received by the antennas of the RRHs in the RF domain. A summing function has to be added to the network to do the combining. It sums the timestamp-aligned signal samples contained in the multiple reverse link packet streams into a single packet stream. Figure 3 illustrates the case where the summing occurs inside a switch. A switch is the ideal place to implement the summing function because it can support the maximum data rate allowed on a Gigabit link. However, commercial off-the-shelf Gigabit Ethernet switches do not implement such a function. A custom designed switch is required. An alternative location for the summing function is in the BSI and a subset of the RRHs. A number of RRHs send their packet streams to a designated RRH, which sums the packet streams with its own stream to produce a single stream that is sent to the BSI. A number of summed

streams arriving at the BSI are again summed to form a final single stream. Clearly, additional stages of summing can be added to support a very large number of RRHs with a single BSI. The obvious advantage of this scheme is that commercial off-the-shelf Gigabit Ethernet switches can be used. But it does come with certain constraints. The links to the summing RRHs and the summing BSI are bottlenecks that will limit the RF bandwidth the system can support. To avoid using too many stages of summing and excessive delay in the reverse link, a summing RRH should be summing packet streams from two to four RRHs. Therefore, the link from the switch to a summing RRH will have to carry at least four reverse link streams and a downlink stream. To support a single UMTS carrier, the raw data rate on the Gigabit Ethernet is $5 \times 153.6 = 768$ Mbps. This approaches the capacity of the link if one includes packet overhead. Therefore, using commercial off-the-shelf Gigabit Ethernet switches with summing in the RRH and BSI would limit the system to support only one UMTS carrier or four CDMA carriers. A system using custom designed switches with a summing function will not suffer from such constraints and will be able to support up to five UMTS carriers or 20 CDMA carriers.

Synchronization Over Gigabit Ethernet Networks

The major technical challenge for transporting cellular signals over Gigabit Ethernet networks is frequency and timing synchronization. For example, CDMA requires 50 ppb frequency accuracy and better than $3\ \mu\text{s}$ timing accuracy, whereas Ethernet clocks are only required to be accurate to 50 parts per million (ppm) and there is no guarantee on the arrival time. Initially, we leveraged the fact that a Gigabit Ethernet network is link-to-link synchronized at the physical layer. The CAT 5 cable consists of four pairs of unshielded twisted pairs, which are used simultaneously to transmit and receive. The physical medium attachment (PMA) hybrid transceiver consists of four independent transmitters, four independent receivers with the capability to cancel the transmitted signal, and a clock recovery unit to recover the clock from the received signal sequence. To reduce the signal processing complexity, it is necessary for the transmitter-receiver pairs at both ends of a link to be synchronized. One pair needs to operate in a master mode and the other in a slave mode. The master transmission is timed with its local clock, and the slave transmission is timed with the clock that it recovers from the transmit symbol stream of the master. When the link is not in use for data transmission, a continuous stream of idle symbols is sent by the master to maintain continuous clock synchronization. Even though Ethernet is asynchronous in the sense that each station has its own local clock (the clocks of the stations are required to be accurate to 50 ppm), in Gigabit Ethernet, at the physical link level, the two ends of every link are synchronized, but there is no end-to-end synchronization through multiple cascaded links.

To achieve frequency synchronization on our first demonstration platform, we modified the clock circuitry of each component to force clock synchronization across the network. Although forced physical link synchronization was achieved, it was not a desirable solution, due to the hardware modifications required.

Later on, we developed software synchronization based entirely on the time-of-arrival statistics of the data packets. With software synchronization, we were able to achieve effective phase-locking between the BTS and the RRHs through off-the-shelf Gigabit Ethernet switches.

The desire to synchronize clocks over an asynchronous packet network such as the Internet has been there for a long time. The Network Time Protocol (NTP) was designed to allow computers on the Internet to synchronize their local clocks to each other by exchanging timestamped IP packets. This protocol uses multiple timing servers (computers with accurate local clocks) and long measurement time (tens of minutes) to achieve the desired accuracy, typically on the order of milliseconds.

Similar principles can be used in RadioStar to synchronize the RRHs to the BSI over the Gigabit Ethernet network. The performance we are targeting far exceeds what NTP offers on the Internet. For example, we require timing accuracy of nanoseconds and convergence time of tens of seconds. But given that the RadioStar network is purely layer 2 and is a well-controlled network, a simple algorithm can work well enough for our purpose. First, we insert a 3-byte timestamp into each and every one of the packets carrying cellular signals. The high packet rate ensures that we can obtain the desired short convergence time desired. This allows us to use a non-ovenized low-cost crystal oscillator in the RRHs. Second, the packets are packaged and stamped by a deterministic process implemented in a field programmable gate array (FPGA) without processor intervention in both BSI and RRH. This eliminates the uncertainty introduced by the higher-layer operating system software in the Internet NTP case. Third, the packet arrival time is also measured and recorded by a deterministic process implemented in an FPGA in the RRHs.

As a packet travels through the switched Gigabit Ethernet network, it experiences delays. The delays in the network are dominated by switch processing that is on the order of $2\ \mu\text{s}$. The delays are largely deterministic because the difference in clock frequencies of the Ethernet interfaces introduces small variations of the values. In 1000Base-T, all links are dedicated and full duplex, and there is thus no random backoff uncertainty introduced by the MAC layer's carrier sense multiple access with collision detection (CSMA/CD) mechanism. Further, in the layer 2 network, as we mentioned earlier, the path traveled by packets from any source to any destination is unique and deterministic, unlike in a layer 3 network. However,

there are non-deterministic delays caused by queuing in the switch's outgoing buffers.

Note that regardless of the network delays, there are still a very significant number of packets arriving with the minimum delay. If we can identify the packets with minimum delay, then we can use their arrival time as a reference to lock the RRH's local clock.

The synchronization algorithm keeps the minimum apparent measured delay constant with discrete steps to the reference source on the RRH. When that happens, the local clock in the RRH will be running at exactly the same frequency as the local clock in the BSI. In fact, they are effectively phase locked. They are also time-synchronized, apart from minimal deterministic delay, which can either be estimated based on the physical network topology or can be automatically determined using a method similar to NTP. Once synchronization is achieved, the RRH also knows the offset in time between its own local clock and the BSI clock, apart from the deterministic minimal propagation delay from the BSI to the RRH. The RRH plays back the received data with an additional fixed delay. This additional delay acts as a jitter buffer to accommodate the variation in the arrival times of the packets. Without the buffer, the playback will have to stop and wait if the packet arrives late, resulting in discontinuity and a sudden change of timing in the transmitted RF signal. The size of the jitter buffer is lower bounded by the maximum possible delay through the RadioStar network for a cellular signal packet.

Operation, Administration, and Maintenance

In applications where a large number of RRHs are deployed, OA&M becomes extremely important. The OA&M needs to be easy to implement, flexible, efficient in resource utilization, and convenient to use. The use of standard Ethernet technology allows RadioStar to take advantage of the standard OA&M strategies developed for data networks. A popular protocol for data network OA&M is SNMP [11], which uses a model consisting of a manager, an agent, a management information base (MIB), managed objects, and a network protocol. The manager provides the interface between the human network manager and the management system, and the agent provides the interface between the manager and the physical

device(s) being managed. Both the manager and agent use the MIB and a relatively small set of commands to exchange information. The MIB is organized in a tree structure with individual variables, such as point status or description, being represented as leaves on the branches. A long numeric tag or object identifier is used to distinguish each variable uniquely in the MIB and in SNMP messages [11]. In RadioStar, the manager is a Linux*-based regular server connected to the network through its Ethernet interface. The human user interface can be provided through a Web page, enabling remote management of the RadioStar network over the Internet. Some of the commercial off-the-shelf Gigabit Ethernet switches support SNMP with their own SNMP agent, which RadioStar OA&M can leverage. On the RRH or BSI, the SNMP agent can be implemented in the microprocessor running a Linux operating system. The microprocessor configures and controls the parameters of the RRH/BSI through a block random access memory (RAM) interface and a serial-parallel interface (SPI). The RadioStar OA&M module needs to provide configuration, monitoring, and control to ease the installation and maintenance of the system and enable advanced applications such as dynamic sectorization. The following important requirements have been identified: installation (including updates and self-discovery), monitoring and repair (identification, deactivation, and reconfiguration), performance management, dynamic reconfiguration, and security. OA&M is critically important for the success of a commercial product based on RadioStar. It is also a place where novel applications can be developed and differentiation can occur. By leveraging standard data network technologies, RadioStar provides a fertile ground for such developments.

Applications of RadioStar Technology

Currently, most in-building DAS deployments are driven by over-the-air repeaters, especially for CDMA air interfaces, because of the high cost of a multicarrier BTS. Over-the-air repeating requires that the outdoor network has sufficient excess capacity to support the additional in-building traffic enabled by the in-building DAS system. RadioStar technology can be

used as a low cost, easy-to-deploy repeater system. Its unique features can be further exploited to realize a more efficient, more flexible, and higher capacity solution by using a dedicated, lower cost, single-carrier BTS for in-building deployments. The following sections detail two types of applications of RadioStar technology.

Dedicated BTS-Driven Application

With a dedicated base station, the entire RF portion can potentially be eliminated with a BSI that digitally interfaces to the BTS to obtain baseband digital I/Q signals. This would allow significant cost reduction in the base station itself. Moreover, for CDMA standards, since the base stations are identified by the time offset of the pseudo-random noise (PN) sequence it radiates, the RF signal needs to go out of the antenna at the specified time with $3\ \mu\text{s}$ accuracy. If there is too much delay in the distribution network, the mobile devices may not be able to identify the base station correctly and will experience problems with handoff. A dedicated base station can compensate for the delay in the distribution network, therefore allowing much higher delay in the distribution network than a repeater system would allow. If a longer delay can be tolerated, the RadioStar network can make use of commercial off-the-shelf Gigabit Ethernet switches for narrowband applications, which can significantly lower the development and hardware costs of the total system. The following paragraphs will focus on the CDMA air interface; however, some of the principles apply to other air interfaces as well.

Distribution of Single Versus Multiple Carrier CDMA

From a hardware perspective, it is much more efficient to distribute a single CDMA carrier than a large number of CDMA carriers. First of all, each RRH is limited in its power capability. Multiple carriers have to share that power, resulting in lower power per carrier and a smaller coverage footprint for the RRH. More RRHs would be required to cover a building with multiple carriers than with a single carrier. Second, with RadioStar distribution, single-carrier distribution uses less Ethernet bandwidth. It allows the summing function for the reverse link packet streams to be implemented in the RRHs and BSI, and the use of commercial off-the-shelf Gigabit Ethernet switches.

This has obvious hardware and development cost benefits.

But what if the capacity needed exceeds what a single carrier can support? The base station can be upgraded with more sectors, and the RadioStar distribution network can be further sectorized to reuse the carrier at different parts of the building. Dynamic sectorization can be used to distribute the capacity optimally if the traffic density is dynamic and highly non-uniform.

In certain situations, it may be desirable to support more than one carrier. For example, it may be necessary to use an evolution data optimized (EV-DO) carrier to support mobile data applications. Up to four CDMA carriers can be supported with commercial off-the-shelf Gigabit Ethernet switches. In any case, the number of in-building carriers is likely to be much smaller than the number of carriers in the outdoor macrocell network. This does create a problem for CDMA, which we will address in the next section.

Application of CDMA Pilot Beacon

In CDMA systems, the frequency reuse factor is 1, which means that all the base stations can transmit at the same frequency simultaneously. Furthermore, in CDMA2000* first evolution (3G1X), all the base stations use the same PN sequence to scramble their transmission data. The base stations differentiate their signals from each other by starting the PN sequence at different times. This is made possible by the fact that all the base stations are synchronized to GPS. A GPS receiver in the base station generates a pulse to synchronize all the base stations. Each base station starts the PN sequence at a unique offset (PN offset) from the sync pulse. The system defines every 64-chip offset as a valid PN offset. Since the length of the PN sequence is 2^{15} , there are 512 unique PN offsets. The pilot channel signal purely carries its PN sequence and the mobile device uses that to lock on to the system and identify the base station.

The advantage of frequency reuse 1 is that a mobile device between two or more base stations can be communicating simultaneously with multiple base stations, i.e., in soft handoff, to improve its link quality and move more reliably between base stations. Upon power-up, a mobile device searches for the

strongest pilot signal and locks on to the system by demodulating other control channels of the base station associated with the strongest pilot signal. It can then initiate a call on that base station. The base station informs the mobile device of the PN offsets of its neighbors in the form of a neighbor list in case the mobile device moves away from the current base station and needs to be handed off to a neighboring base station. The mobile device uses the neighbor list to search for other base stations that could potentially be candidates for handoff. It searches for the neighbors by using the current base station's signal as the reference and searches around the expected offsets of the neighbors within a certain window. The search window needs to be small so as to reduce the time needed for the operation, but large enough to account for delays of the signal over the air. Pilot signals of base stations that are farther away from the mobile device than the current serving base station will arrive later than the expected offset. Conversely, those closer will arrive earlier. The search window size is typically 20 chips, or approximately $16 \mu\text{s}$. Once a mobile device detects a pilot signal with sufficient strength, it reports its finding to the serving base station in a pilot strength measurement message (PSMM) containing the E_c/I_o (pilot strength) of the pilot and the exact phase of the pilot. The phase is the number of chips the detected pilot is offset from its expected position using the serving base station's signal as the reference. Based on the measurement reports

from the mobile device, the mobile switching center (MSC) can identify the base station that the mobile device is moving to and therefore can make a reliable decision on handoff if needed.

In the case where there is carrier mismatch between outdoors and indoors, some mobile devices have to perform inter-frequency handoff as they enter or exit the building. Inter-frequency hard handoff is usually much less reliable than soft handoff because of the lack of information about the location and RF environment of the mobile device in need of handoff. Take the example of an active mobile device in the outdoor network, on a carrier not supported indoors. As it moves into a building, due to the attenuation of the RF signals by the building, the mobile device's RF link quality drops significantly and the mobile device needs to look for a candidate base station for handoff. Since the in-building system is on a different carrier, the mobile device cannot sense its presence or its potential link quality to the in-building base station. The outdoor network does not know the whereabouts of the mobile device and cannot make intelligent decisions about where to hand off the mobile device. The device will most likely experience call discontinuity.

One solution is to have the in-building system radiate a pilot beacon at the frequencies of the outdoor network near the portal of the building. This is illustrated in **Figure 4**. A pilot beacon is simply a

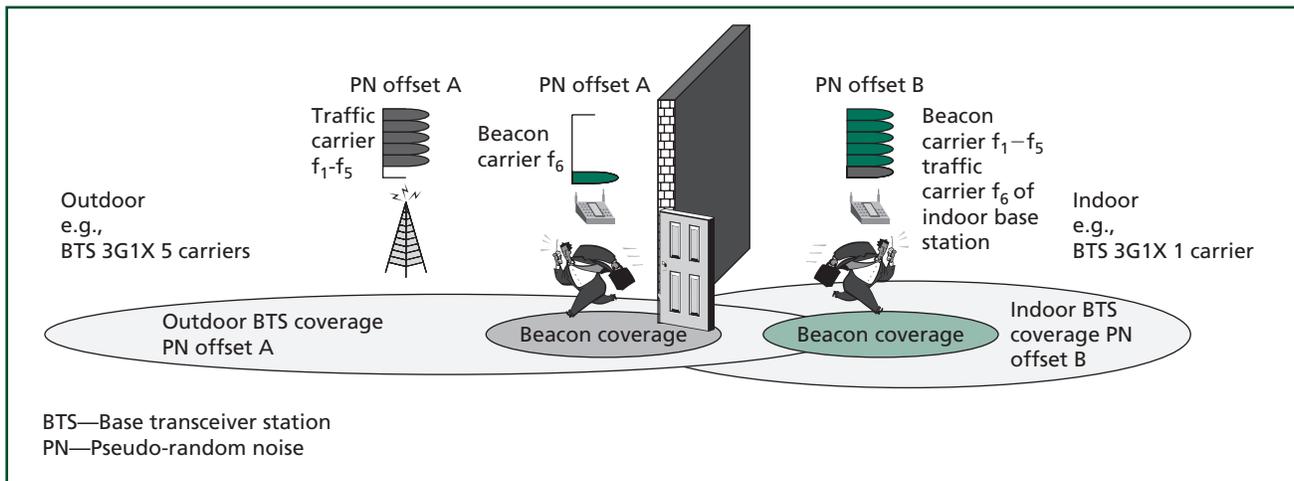


Figure 4.
Illustration of a typical in-building deployment scenario with pilot beacons.

carrier signal with a pilot channel only. The beacon will have the same PN offset as the in-building system. As a mobile device enters the building, it senses the signal from the beacon and reports back to the outdoor base station the strength and phase of the beacon, which reveal the mobile device's location to the network. When the device's RF link to the outdoor base station weakens to the point of needing handoff, or the strength of the pilot beacon is great enough, the system can instruct the mobile device to make an inter-frequency hard handoff to the in-building base station. In this case, an inter-frequency handoff can be made very reliably.

If the indoor system uses dedicated carriers not supported on the outdoor system, a pilot beacon solution applies to the situation where a mobile device is on a call in the in-building base station and is leaving the building. A pilot beacon radiating at the frequencies of the indoor network can be installed outside the portal of the building to facilitate the inter-frequency handoff from the in-building network to the outdoor network. The PN offset of the beacon has to be the same as that of the strongest pilot from the outdoor base station that covers the area.

Consider a typical in-building deployment scenario with pilot beacons. The outdoor base station uses for example f_1 through f_5 , and the indoor base station uses f_6 . There should be significant coverage overlap between the two base stations around the building entrance. Two beacons are set up near the entrance as well. One is radiating at f_1 through f_5 , with the indoor base station PN offset inside the building with a limited coverage footprint just inside the entrance. The other is radiating f_6 , with the outdoor base station PN offset with a limited coverage footprint just outside the entrance. As a user on the outdoor system walks into the building and enters the coverage area of the indoor beacon, the mobile device will report the PN offset of the indoor beacon back to the MSC through the outdoor base station and trigger the handoff to the indoor base station. Similarly, as a user on an indoor system walks outdoors and enters the coverage area of the outdoor beacon, the mobile device will report the PN offset of the outdoor beacon back to the MSC through the indoor base station and trigger the handoff to the outdoor system. The areas covered by the two beacons

are selected such that users who are not near the entrance are not accidentally handed off, and users going through the entrance will not be ping-ponged between the two base stations. The position of the beacons and their coverage areas thus has to be strategically selected. The pilot beacon feature is unique to the indoor solution. This feature does not impose any extra load on the network, since the PN offset of the pilot channel is generated locally by the RRH, and it can be easily reconfigured through OA&M. In this sense, it is a clear differentiator from DAS solutions.

Prototype and RadioStar Demonstration

Very early in the development of the RadioStar concept, we started to assemble hardware to test and demonstrate key aspects of the technology. With evaluation boards and a rack full of instruments, we were able to demonstrate live CDMA-to-CDMA mobile calls fairly quickly. Then, we set out to build a prototype RRH and BSI with much smaller form factor and power-over-Ethernet capability. The result is a much more compact and portable demonstration platform.

As we noted before, the functionality is very similar in the RRH and BSI. In fact, the digital portion of the RRH and BSI can be essentially the same when different firmware loads are used. The analog portion is also similar except that the transmit- and receive-frequencies are interchanged. In order to maintain flexibility and minimize dependencies between the analog and digital development, we chose to implement the prototype in two separate boards, namely, a digital board and an analog board.

Figure 5 depicts the block diagram of the digital board architecture. It comprises the following major functional blocks: an FPGA with a Gigabit Ethernet interface and an RF interface. The FPGA acts as the central digital processing unit that converts between sampled baseband radio signals and Ethernet packets and controls the operation of the software synchronization circuit.

For the analog board, the primary objective is to leverage commercial off-the-shelf handset components. The challenge is to have these components work well for the frequency plan of an RRH, which is different from that for a handset, and then to meet the performance

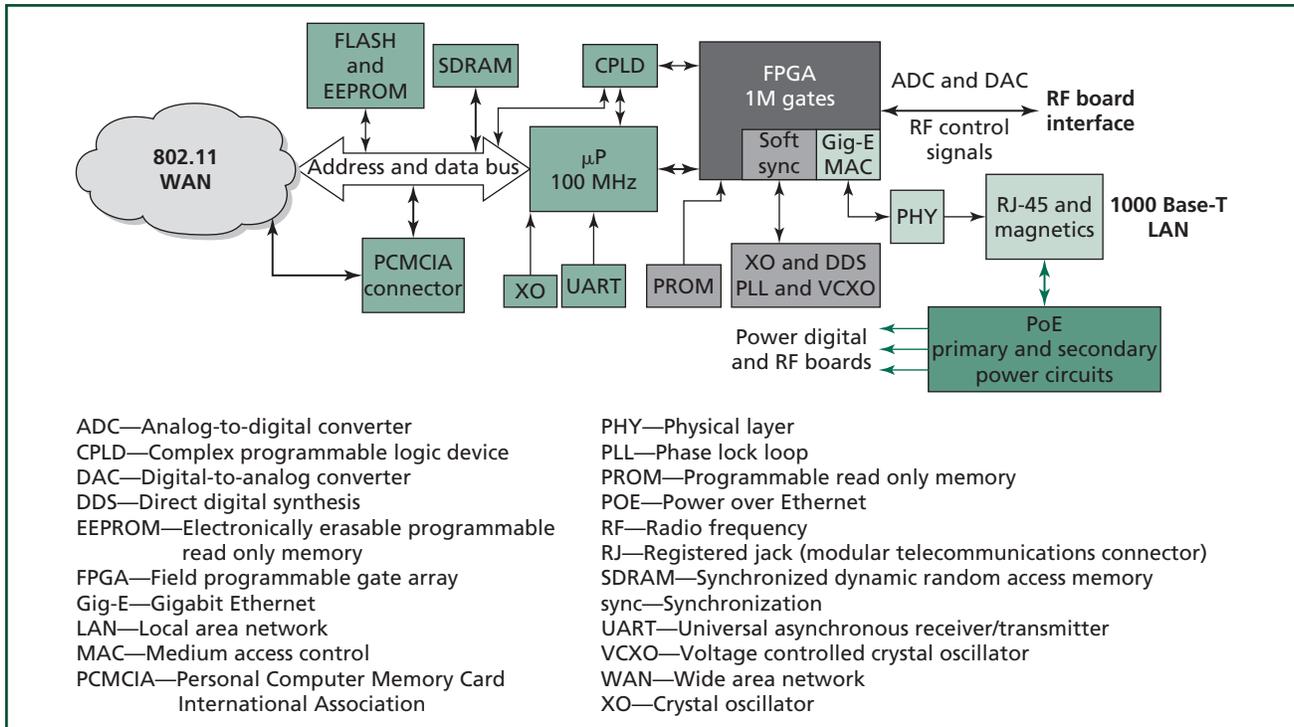


Figure 5.
Block diagram of a RadioStar digital board design.

requirements, especially the emission requirements. We chose a personal communication system (PCS)/cellular dual-band CDMA handset chip set based on a super-heterodyne architecture. **Figure 6** is a block diagram of the design. It has a single transmitter and two independent receivers.

The resulting module is shown in **Figure 7**. The interface between the two boards is purely digital except for a frequency reference signal. In addition, the two-board structure also allows us to achieve a small physical form factor of $9 \times 6 \times 2$ inches.

The demonstration platform consists of a stand-alone modified OneBTS™ compact cell digital shelf, a BSI that is RF-interfaced to the OneBTS radio card, an off-the-shelf Gigabit Ethernet switch, and several RRHs. The compact cell is loaded with a special version of the software such that it can connect the two mobile devices without the MSC. A media converter is inserted in the link between the BSI and the switch to convert from 1000Base-T to 1000Base-X over multi-mode fiber. A mid-span injection unit is inserted

between the switch and the RRHs to provide power-over-Ethernet to the RRHs.

By connecting a BSI and an RRH back-to-back with a CAT 5 cable, the repeater function can be effectively demonstrated. We set up a BSI with an antenna near the window in the lab where the signal from a nearby external base station is well received. We set up the RRH inside an RF anechoic chamber used for antenna measurements. When the door of the anechoic chamber is closed and the RRH turned off, a mobile device indicates that it has no service. When the RRH is turned on, the mobile device indicates that it has service. We can successfully initiate calls to a landline phone or another mobile device through the external cellular network. We have tested this BSI/RRH platform with CDMA and Global System for Mobile Communications* (GSM*) phones. No degradation to the voice quality is detectable. There is a relatively large round trip delay introduced by the repeating arrangement in (about $40 \mu\text{s}$). We walked in and out of the anechoic

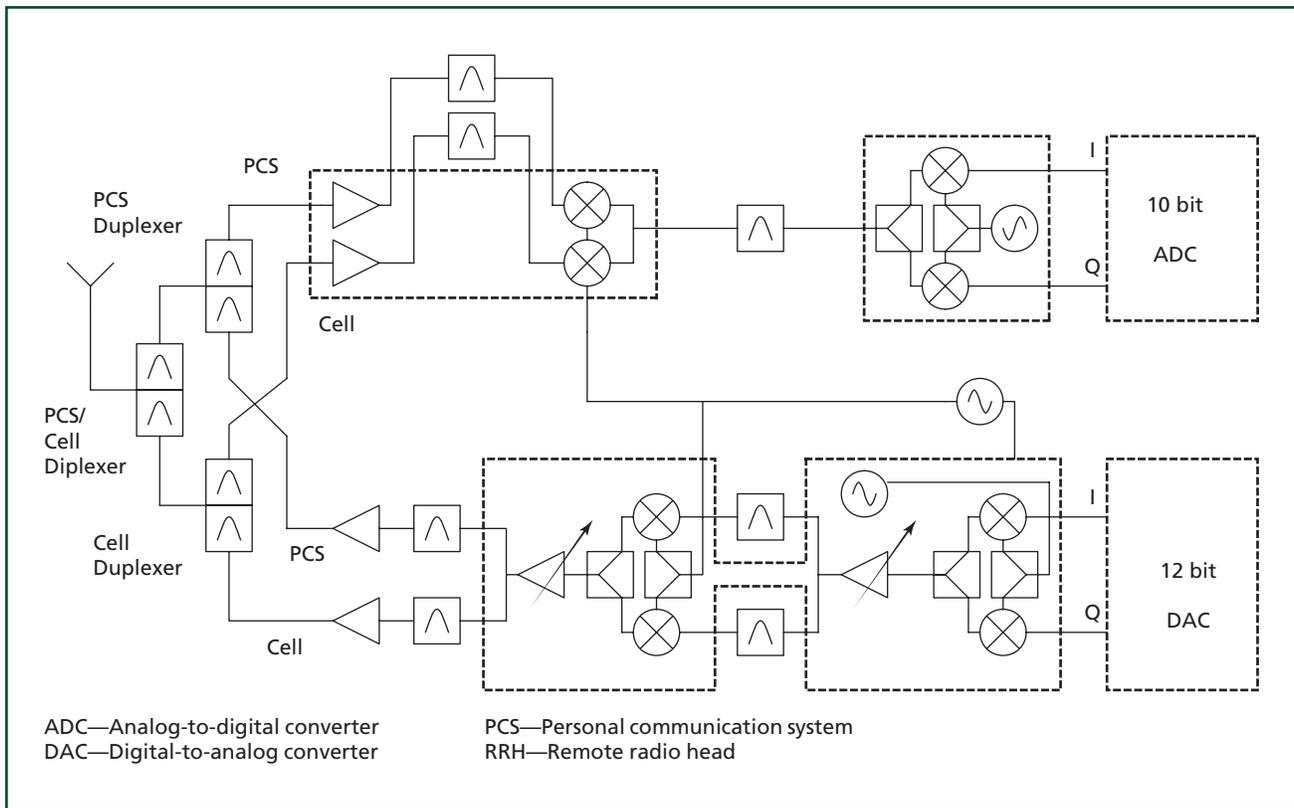


Figure 6.
Block diagram of the RRH analog board design.

chamber so that the mobile device was communicating with the base station either through the repeater or directly over the air. We did not notice any breakage in the voice connection during the transition.



Figure 7.
The developed RRH module with power over Ethernet.

We have successfully carried out the demonstration of pilot beacon-assisted hard handoff between two PCS base stations, and between a PCS base station and a cellular base station in the lab using our RadioStar prototypes. Next, the RadioStar prototypes were installed in Alcatel-Lucent’s Murray Hill facility for a period of one year on the cellular network of one of our customers. The demonstration consisted of two BSIs and six RRH units, four of which were used to serve two sectors of 3G1X voice traffic, and two to provide pilot beacons. The result of 10,000 hours of service and thousands of beacon-assisted handoffs showed a success rate of more than 98 percent.

Acknowledgements

The RadioStar project has benefited greatly from many individuals. We would like first to thank Jungsang Kim for initiating the RadioStar project in 2003, and for

his leadership, enthusiasm, energy, and vision. We would also like to acknowledge contributions from Anatoli Olkhovets, who was instrumental in developing the architecture of RadioStar and an early demonstration of hardware synchronization through the Gigabit Ethernet PHY. Yee Low designed an adapter board used in the first RadioStar prototype built of evaluation boards. His contributions to power-over-Ethernet are also acknowledged. Piyush Gupta, Mansoor Alicherry, and Chitra Phadke were members of the RadioStar team initially and their contributions are hereby acknowledged. Within research, we would like to acknowledge support from our managers, T.V. Lakshman, Viswanath Poosala, and Gee Rittenhouse. We especially would like to thank Ilija Hadzic for consulting on the software synchronization, and Ed Sutter for getting MicroMonitor to work on our digital board. The project has also received much support from the Mobility Business Unit. We would like to thank Thomas Gvoth, for providing the set up of the compact OneBTS for mobile-to-mobile test calls in the lab, and Peter Giannoglou and Richard La Grotta for their support in providing board layout and physical design resources, and Peretz Feder. Members of System Engineering have been a great source for consultation, reviewing, and support, including Russell Sun, Qi Bi, John Burgess, Victor Pan, Christopher Capece, and Miguel Dajer. We would also like to gratefully acknowledge support from Teh-Ming Yeh, Ram Thirunagari, and Igor Murkis in setting up the pilot beacon demonstration in the system integration lab in Whippany.

*Trademarks

CDMA2000 is a trademark of the Telecommunications Industry Association.

GSM and Global System for Mobile Communications are registered trademarks of the GSM Association.

Linux is a trademark of Linus Torvalds.

References

- [1] R. Breyer and S. Riley, *Switched, Fast, and Gigabit Ethernet*, 3rd ed., Macmillan Tech. Pub., Indianapolis, IN, 1999.
- [2] D. G. Cunningham and W. G. Lane, *Gigabit Ethernet Networking*, Macmillan Tech. Pub., Indianapolis, IN, 1999.

- [3] T. A. Freeburg, "Enabling Technologies for Wireless In-Building Network Communications—Four Technical Challenges, Four Solutions," *IEEE Commun. Mag.*, 29:4 (1991), 58–64.
- [4] Institute of Electrical and Electronics Engineers, "Data Terminal Equipment (DTE) Power Via Media Dependent Interface (MDI)," IEEE 802.3af, June 2003.
- [5] Institute of Electrical and Electronics Engineers, "Carrier Sense Multiple Access With Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," IEEE 802.3, Dec. 2005.
- [6] Institute of Electrical and Electronics Engineers, "Data Terminal Equipment (DTE) Power via Media Dependent Interface (MDI) Enhancements," IEEE 802.3at, Draft v3.0, Mar. 2008.
- [7] E. F. T. Martijn and M. H. A. J. Herben, "Characterization of Radio Wave Propagation Into Buildings at 1800 MHz," *IEEE Antennas and Wireless Propagation Lett.*, 2:1 (2003), 122–125.
- [8] S. P. Morgan, "Prediction of Indoor Wireless Coverage by Leaky Coaxial Cable Using Ray Tracing," *IEEE Trans. Veh. Technol.*, 48:6 (1999), 2005–2014.
- [9] T. S. Rappaport and R. A. Brickhouse, "A Simulation Study of Urban In-Building Cellular Frequency Reuse," *IEEE Personal Commun.*, 4:1 (1997), 19–23.
- [10] T. S. Rappaport and R. A. Brickhouse, "A Simulation of Cellular System Growth and Its Effect on Urban In-Building Parasitic Frequency Reuse," *IEEE Trans. Veh. Technol.*, 48:1 (1999), 286–294.
- [11] K. Yaghmour, *Building Embedded Linux Systems*, O'Reilly, Sebastopol, CA, 2003.

(Manuscript approved December 2008)

ZHENGXIANG MA was a distinguished member of technical staff in Bell Labs' Converged Access Systems Research Department at the time this paper was written. He received a B.S. degree in physics from the University of Science and Technology, Hefei, China, and M.S. and Ph.D. degrees in applied physics from Stanford University, Stanford, California. His research interests include novel digital signal processing in high-performance base station radios, power amplifier



linearization, and wireless networking infrastructure architecture. He has co-authored several technical papers and holds 11 patents.

MIKE ZIERDT is a member of technical staff in Bell Labs' Wireless Communications Research Department in Murray Hill, New Jersey. He received B.S. and M.S. degrees in electrical engineering from Polytechnic University, Brooklyn, New York. For four years, he worked at Mini Circuits Laboratories as an RF components designer, followed by six years in the Consumer Products Division of AT&T/Lucent Technologies, working primarily on wireless handsets. For the past nine years, he has been in Bell Labs' Wireless Research Laboratory in Murray Hill, New Jersey, working on a variety of transceivers.



JOHN PASTALAN is a member of technical staff in Bell Labs' Wireless Communications Research Department in Murray Hill, New Jersey. He received a B.S. degree in physics and an M.S. degree in applied physics from the State University of New York at Binghamton. His research interests include digital design prototyping and wireless networking infrastructure architecture. He holds three patents and has contributed to numerous technical publications.



ARNOLD SIEGEL is a member of technical staff in the Alcatel-Lucent UMTS Platform and Core Systems Engineering Department in Whippany, New Jersey. He received a B.E. degree in electrical engineering from The Cooper Union, New York, New York, and an M.S. degree in electrical engineering from Princeton University in New Jersey. His research interests include in-building and in-home cellular distribution systems and next-generation radio architectures. He has co-authored several technical papers and holds three patents.



TOD SIZER is leader of Bell Labs' Wireless Access research domain, and is currently on international assignment in Stuttgart, Germany. He graduated magna cum laude from Amherst College in Massachusetts and received his master's and doctorate from the Institute of Optics at the University of Rochester in New York. During his tenure at Bell Labs, he has performed research in wired and wireless home



networking, fixed wireless loop systems, video watermarking technologies, optical computing and switching systems, and high power laser design. He was a member of the technical team in the company's role as a promoter in the Bluetooth Special Interest Group (SIG). His responsibilities in the SIG include being Chair of the Coexistence Working Group. He holds 43 patents with 16 patents pending and is the author of over 50 refereed publications. He is a member of the IEEE and the Optical Society of America (OSA).

ADRIAAN J. DE LIND VAN WIJNGAARDEN is a member of technical staff in Bell Labs' Networks and Communications Research Department in Murray Hill, New Jersey. He received his M.S. in electrical engineering from Eindhoven University of Technology, The Netherlands, and his doctorate in engineering from the University of Essen, Germany. He has been deeply engaged in both theoretical and application-driven research in communications, information theory, coding, and combinatorial and algorithmic optimization. He provided key contributions to 10G, 40G, and 100G optical systems and wireless and broadband access. He has authored more than 60 technical papers, holds six patents, and has more than 20 patents pending. Dr. de Lind van Wijngaarden is a Senior Member of the IEEE. He served as a Publication Editor of the IEEE Transactions on Information Theory from 2005 until 2008. He is currently an Associate Editor (Communications) of the IEEE Transactions on Information Theory. He has co-organized Shannon Day at Bell Labs in 1998, Shannon's Statue Unveiling at Bell Labs in 2001, Center for Discrete Mathematics and Theoretical Computer Science (DIMACS) conferences, and several international workshops.



PRAVEEN R. KASIREDDY was a member of technical staff in the Bell Labs' High Speed Networks Department at the time this paper was written. He received a B.Tech degree in electronics and communications engineering from the Krishna Devaraya University, Anantapur, India, and an M. Tech degree in microwave engineering from the Banaras Hindu University, Varanasi, India. Prior to his tenure at Bell Labs, he was a research engineer in the Center for Development of Telematics in Bangalore, India. His research interests include high speed digital design and packet processing.





DRAGAN M. SAMARDZIJA is a member of technical staff in Bell Labs' Wireless Communications Research Department in Holmdel, New Jersey. He received a B.S. degree in electrical engineering and computer science from the University of Novi Sad, Serbia, and M.S. and Ph.D. degrees in electrical engineering from the Wireless Information Network Laboratory (WINLAB), Rutgers University, New Brunswick, New Jersey. During his tenure with Bell Labs, Dr. Samardzija has been involved in research in the field of next-generation wireless systems. His research interests include analysis, detection, estimation, and experimental evaluation of multiple input-multiple output (MIMO) wireless systems. He has been co-principal investigator and participant in a number of research projects sponsored either by the National Science Foundation (NSF) or the European Union. As a member of the Bell Labs Layered Space-Time (BLAST) project team, he is a recipient of the 2002 Bell Labs President's Award. He has authored more than 30 journal and conference publications. In earlier work, he was involved in design and development of a multiple antenna high speed downlink packet access (HSDPA) testbed which was used to demonstrate the first-ever HSDPA communication session at CTIA 2003 in New Orleans. He was also a 2003 recipient of the Central Bell Labs Teamwork Award for the HSDPA Demonstration Team. Since then, he has been involved in a number of measurement campaigns and system evaluations of different third generation (3G) and fourth generation (4G) wireless systems. ♦