

# IN-BAND WIRELESS BACKHAUL USING SDR FOR RURAL CELLULAR SYSTEMS

Carlos Cabrera-Mercader, Li-Wei Chen, Brian Fallik, Stanley Hu, and John Chapin  
Vanu, Inc. Cambridge, MA, USA [info@vanu.com](mailto:info@vanu.com)

## ABSTRACT

This paper describes an SDR-based wireless backhaul system for a cellular radio access network. Backhaul is the data connection between base stations in the field and the operator's switch or network. In-band backhaul is a wireless data link that uses the cellular spectrum already licensed by the operator, as opposed to licensed microwave or unlicensed spectrum. The innovative feature of the Vanu in-band backhaul system is using SDR to support both cellular service to mobiles and backhaul communications simultaneously on shared processor, radio and antenna resources. No additional hardware is needed at most sites for the backhaul beyond that already installed for the cellular base station. We use an IDEN radio access network as a concrete example, and present results from proof of concept field trials of the backhaul system.

## 1. INTRODUCTION

The major operating expense for cellular base stations in low-density rural areas is the backhaul, the data connection from the base station back to the operator's switch or network. Wireless backhauls such as microwave links are common due to the high cost of leased lines to these remote locations. More recently it has become popular to install in-band backhaul, where the backhaul link operates in the cellular spectrum rather than microwave. The cellular band is normally underutilized for call traffic in these rural areas so it is available for backhaul. This design enables the operator to reuse a resource they already have, the licensed cellular spectrum, rather than acquiring new microwave spectrum or competing for unlicensed spectrum.

Vanu, Inc. has demonstrated an in-band wireless backhaul system that goes even further in reusing resources already acquired by the operator. Existing in-band backhaul solutions require new radios, PAs, and antennas. Because Vanu base stations are software radios, the backhaul waveform in the Vanu system can run simultaneously on the same signal processing server as the cellular waveform supporting client mobiles. Through careful system and waveform design, the radio head and antennas at most sites can also be shared between cellular service and backhaul, resulting in zero additional acquisition, installation, or maintenance costs at most sites for the in-band backhaul.

This paper first analyzes the tradeoffs of in-band wireless backhaul using SDR in general terms. The bulk of the paper focuses on design and implementation issues. Finally we report results from a field trial in an IDEN<sup>1</sup> radio access network.

## 2. LOW-COST BACKHAUL USING SDR

Adding wireless backhaul to a cellular radio access network (RAN) without adding radio hardware or antennas is difficult. Cellular RANs are FDD, with separated transmit and receive bands. One base station cannot receive signals transmitted by another base station.


However, it is still possible to use SDR technology to minimize the amount of hardware dedicated to wireless backhaul. We differentiate *hub* and *remote* sites. A hub site has a wireline backhaul connection, typically one or more T-1 lines. A remote site relies on wireless backhaul to a hub site. A key observation is that there are more remote sites than hub sites. Operators normally seek to amortize the high cost of bringing wireline data service to a hub site by sharing its data connection across as many remote sites as possible. Operators we have worked with have up to four times as many remote sites as hub sites.

We minimize cost with a system design that requires additional backhaul hardware only at hub sites. In order to reuse the installed hardware at the remote sites, the backhaul waveform must be FDD and remotes must transmit in the cellular downlink band and receive in the cellular uplink band. Accordingly, at hub sites, we install a *frequency reversed* radio head connected to a dedicated antenna. It is frequency reversed because it receives in the cellular downlink band, where base stations normally only transmit, and it transmits in the cellular uplink band, where base stations normally only receive.

Adding a frequency reversed radio to a cellular tower creates a significant co-site interference problem. Transmissions by the hub over the backhaul link interfere with the cellular system's ability to receive weak mobile signals. This is the primary design challenge of low-cost in-band backhaul. We discuss the solutions adopted by Vanu in Section 3.2.

---

<sup>1</sup> IDEN is a trademark of Motorola.



**Figure 1: Network plan. X = hub, square = remote.**

Once these problems are solved, in-band backhaul using SDR offers significant cost savings compared to more traditional in-band backhaul systems. As one case study, consider the cellular network topology shown in Figure 1 consisting of 116 sites. This example is based on deployment data but is not a map of any particular system.

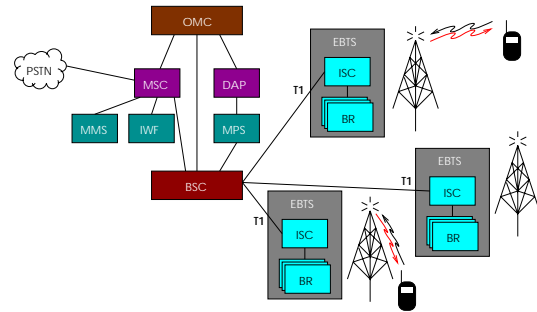
Costs in this deployment model are normalized. Each channel used for backhaul is given cost 1, associated with either leasing the frequency or giving up its use for revenue-producing cellular service. The equipment added to a hub is given cost 3. Adding backhaul capability to a remote costs 0.5 when SDR technology is used, or 1.5 with a traditional backhaul system requiring dedicated hardware. A topology optimization program is used to decide which base stations should be hubs, which should be remotes, and which hub each remote should connect to [1].

Figure 1 shows the resulting network plan. The 116 sites are divided into 19 hubs and 97 remotes. With frequency reuse, a total of 25 uplink frequencies and 23 downlink frequencies are required. With SDR at the remotes, this translates into a total network cost of 153.5. This compares with a cost of 250.5 for a traditional backhaul design. The use of SDR has saved almost 40% of the deployment cost.

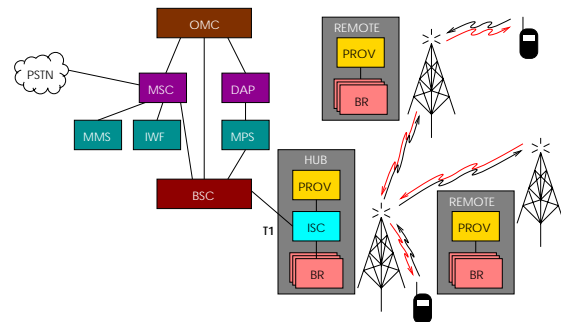
### 3. DESIGN AND IMPLEMENTATION

In this section, the design of the wireless backhaul system and associated implementation issues are discussed. We describe solutions for the challenges of frequency-reversed radios and the challenges of sharing an SDR device between two waveforms at the remotes.

The backhaul waveform is called “Providence.” It was designed to be used in an IDEN cellular RAN. IDEN is similar to most other 2G cellular systems, with a switch (MSC), base station controller (BSC), and base stations (Figure 2) [2]. Unlike other 2G systems, the base station is



**Figure 2: IDEN network topology**



**Figure 3: IDEN network topology with wireless backhaul**

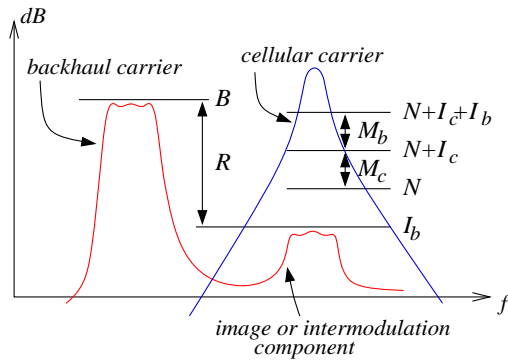
divided into two components, the integrated site controller (ISC) and the base radio (BR). There are other differences from other 2G systems, such as core components associated with push-to-talk features, but these components do not interact with the wireless backhaul subsystem.

#### 3.1. RAN Architecture with wireless backhaul

There are two primary options for use of wireless backhaul in an IDEN RAN. The backhaul link can be inserted between the ISC and BSC or the BR and ISC. We chose to remote the BRs and keep the ISCs in the hubs (Figure 3). This reduces maintenance costs and enables ISC sharing when remote site capacity requirements are low.

The existing ISC and BR units are designed to be connected by a high-quality Ethernet link. The Providence backhaul waveform was designed to be an Ethernet bridge that would provide a drop-in wireless replacement for this link. This created an implementation challenge. Attempting to replicate a low-loss, low-latency wired link using a wireless waveform places stringent quality-of-service requirements on the wireless backhaul waveform.

The backhaul waveform physical layer is a frequency-division duplexed OFDM signal with scalable data rate and bandwidth. Trellis-coded modulation allows for spectrally-efficient coding at high data rates, and variable bit loading on the subcarriers allows the waveform to use waterfilling



**Figure 4: Interference at the cellular base station receiver.**  $N$  is the receiver noise level,  $I_c$  is interference from the cellular network, and  $I_b$  is interference caused by a backhaul carrier.

to take advantage of frequency-selective fading channels. Differentiated quality of service classes were implemented to support the varying packet error rate and latency requirements of the voice, data, and control signaling traffic being carried over the link to ensure that all service constraints were met.

### 3.2. Co-site interference at hubs

Interference management is a critical challenge at hub sites. As described in Section 2, the backhaul transmit and receive bands are reversed at a hub compared to the cellular transmit and receive bands. The proximity of the cellular and backhaul antennas creates a significant co-site problem. Due to space limitations, we discuss only the interference to uplink cellular reception caused by backhaul transmission. The issues and solutions are similar for cellular system downlink interference to backhaul reception.

Backhaul transmissions can degrade cellular base station receiver sensitivity in several ways:

1. By generating noise in the cellular uplink channel, even though this is out-of-band for the backhaul transmission, e.g. due to images or inter-modulation.
2. By increasing quantization noise. This occurs when the interferer level is high enough that gain compression is required in order to prevent clipping. The reduced dynamic range increases quantization noise. (This phenomenon is also known as “receiver desensing.”)
3. By distorting the cellular signal, due to nonlinearities in the analog RF section of the cellular receiver.

Degradation type 1 (noise) is the dominant effect in Vanu macro cell base stations, since they have high-resolution A/D converters and high-quality RF stages that limit type 2 or 3 degradation. An analysis of noise level therefore gives a good approximation of the co-site degradation effect. Figure 4 illustrates the quantities used in the analysis.

**Acceptable interference level:** In cellular networks the level of acceptable noise is specified by the interference margin, which is the ratio of total interferer and noise power to the noise power seen at the receiver in the absence of any interference. For instance, an interference margin of 3 dB is typical for CDMA networks. The interference margin reflects the acceptable level of adjacent and co-channel interference from within the network. Adding backhaul increases the required interference margin.

The typical receiver noise in an IDEN channel is  $N = -130$  dBm in a bandwidth of  $W_c = 20$  kHz. Let the bandwidth of a backhaul carrier be  $W_b = 100$  kHz. Assume a cellular interference margin of  $M_c = 3$  and an additional degradation of  $M_b = 1$  dB due to interference from the backhaul. Furthermore, assume that the image or intermodulation-component level is  $R = 60$  dBc relative to the backhaul carrier. The maximum tolerable backhaul level at the base station receiver input is given by

$$\begin{aligned}
 B &= R + N + M_c + 10 \log_{10} \left( 10^{M_b/10} - 1 \right) \\
 &\quad + 10 \log_{10} (W_b / W_c) \\
 &= -61 \text{ dBm}
 \end{aligned}$$

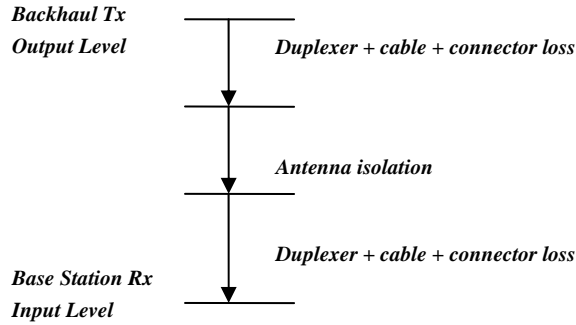
Note that this example already assumes that the backhaul radio has been designed to provide image attenuation and inter-modulation levels of at least 60 dBc.

**Design solutions:** There are two practical ways to reduce the backhaul signal to the required level at the cellular receiver input: a) use analog filters, and b) isolate the backhaul and cellular antennas.

Analog filters require frequency separation between the backhaul and cellular channels and an unused guard band in between, due to filter roll-off. Our strategy was to allocate separate frequency sub-bands for the backhaul and cellular waveforms, then use the backhaul radio’s duplexer to provide the needed filtering. In that way a single duplexer design can be used for all the backhaul radios. Note that the guard band need only be unused at hub sites. At remote sites, the guard band can be used for cellular service.

Antenna isolation can be improved by careful placement of the backhaul antenna at a hub site relative to the cellular antennas. Optimally, the backhaul and cellular antennas should be vertically aligned and separated from each other as much as possible to minimize the overlap between the main lobes of the antenna patterns. In practice, placement options are limited, especially when adding backhaul antennas to existing sites.

Figure 5 shows the quantities involved in the calculation of the required filtering and antenna isolation. Continuing with the previous example, the backhaul signal level at the input of the base station receiver should be at



**Figure 5: Calculation of required filtering and antenna isolation. It is assumed here that filtering is provided by the backhaul duplexer**

most -61 dBm. For a backhaul transmit power of 36 dBm the total loss required would be

$$\begin{aligned} \text{required loss} &= \text{backhaul Tx output level} \\ &\quad - \text{base station rx input level} \\ &= 36 - (-61) = 97 \text{ dB} \end{aligned}$$


Assuming the losses from connectors, cable, and receiver duplexer add up to 4 dB, then the transmit duplexer loss and antenna isolation must be at least 93 dB. Distributing the required loss between the two involves a tradeoff between the width of the transition band and the range of the backhaul links, which will be impacted by placement of the backhaul antenna.

### 3.3. Sharing radios at remote sites

Having discussed the co-site issue at hub sites, we now turn to the challenges of sharing a single SDR device at remote sites between backhaul and cellular functions. The device used was a Vanu Inc. Anywave system (Figure 6).

In this system, the front end has an analog-to-digital converter that digitizes a 26MHz wide channel bandwidth at IF. In the receive path, channel selection filtering and downsampling is performed digitally. This allows multiple channels of varying bandwidth to be selected from the 26MHz using independent receive paths. Baseband samples are then packetized and sent via a gigabit Ethernet link to the server. All signal processing functions run on the general-purpose processors in the server.

On the transmit path, the GPP again performs all the necessary signal processing and packetizes baseband samples that are sent to the front end via Ethernet. Upsampling, image rejection filtering, and combining of multiple transmit signals are performed in the front end's digital section at the 26MHz IF bandwidth. Digital-to-



**Figure 6: SDR system architecture.**

analog conversion and RF modulation are then performed before the analog signal is sent to the power amplifier.

Sharing this system between two independent waveforms creates several challenges.

**RF head configuration:** The RF head consists of analog RF components for filtering, oscillation, tuning, etc. combined with independent digital boards for DUC/DDC, sampling, and packetization. Operationally, the RF head settings must be coordinated prior to startup, and configured during the initialization stage. Once actively producing baseband samples, the settings cannot be changed lightly since new settings may break the sample stream, forcing receivers to re-acquire timing synchronization.

Our initial solution to share a single RF head between the cellular and backhaul waveforms was primitive: simply designate one waveform as the master, modify it to configure the RF head for both waveforms, and modify the other waveform software to passively assume that its configuration requirements are met. For the field trial, we designated the IDEN waveform as the master.

This approach worked well, but had some serious drawbacks. First, the Vanu Anywave platform supports many waveforms, and this architecture requires explicit engineering effort to support each possible combination of waveform pairs. Another drawback is the loss of a symmetric startup procedure. Restarting the master always restarts both waveforms, while restarting the slave does not. Finally, this approach becomes awkward when scaled to more than two waveforms.

Since the field trial, we developed a new approach that provides a consistent and scalable design for sharing the radio. Each server contains a lightweight RF head configuration manager, which is responsible for:

- determining Ethernet network topology
- aggregating RF settings
- configuring the RF head during startup
- monitoring RF head health status
- answering channel resource requests from individual waveforms

This approach improves sharing in a number of ways. First, all waveforms now use a common protocol to request RF head resources, and this code can be shared easily. Also, the startup process is well understood, with hooks in place to allow waveforms to make adjustments. Lastly, a number


of configuration items have been removed from the waveform, and are instead handled by the manager.

**Data link multiplexing:** Another shared resource when supporting multiple waveforms is the data link to the RF head, running over Ethernet in Vanu SDRs. The protocol used to represent the baseband sample stream can limit the ability to multiplex RF channels over a shared link. For example, in Vanu's protocol, Ethernet pause packets are used for flow control. As a result, all waveforms are paused when any channel's FIFO fills. Since Ethernet ports are relatively cheap on modern server hardware, our simple solution was to use a separate Ethernet link for each waveform.

**Automatic Gain Control:** A common feature in many radio front ends is Automatic Gain Control (AGC), used to adjust the attenuation at the input to the front end to prevent receiver saturation. In software radio deployments targeting only cellular radio links, the AGC settings are tuned to maximize receiver sensitivity. Attenuation is initially set very low, and adjusts upwards dynamically as the receive signal power increases. This allows very low-power signals to be received at low attenuation, but requires more dynamic adjustment of signal level on an ongoing basis. When the AGC adjusts during the reception of a voice packet, the voice packet is typically lost since to the communication system it appears that the channel gain changes abruptly mid-packet, an event that is typically not detectable using a training sequence that does not run the full duration of the packet. For the packet error rates targeted for cellular voice traffic, this loss is acceptable in exchange for the increased sensitivity.

The same settings optimized for receiver sensitivity were initially used for the shared front end. In the Providence backhaul link, however, the target packet loss probability is several orders of magnitude lower in order to provide near-guaranteed delivery. Lost packets due to AGC adjustment events quickly became the dominant source of error and limited the performance of the system. To bypass this limitation, special adjustments were made to both the backhaul waveform and the AGC to reduce the frequency of the adjustment events.

Since the front end is shared between the IDEN and backhaul links, either a large IDEN or backhaul receive signal can trigger the AGC. Large backhaul receive signal level can be avoided by adopting a handshaking mechanism to ensure that links that are turned on begin transmitting with initially low power, and that over time closed-loop power control ensures that the power level is not increased to a level that saturates the receiver. Since in a backhaul link both transmitter and receiver are typically stationary, the channel is usually changing slowly enough that power control is very effective.



**Figure 7: Map of field trial site showing elevation (m)**

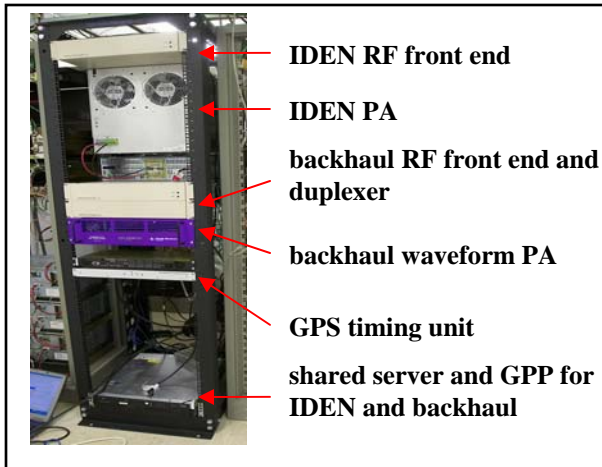
On the IDEN side, the base station does not necessarily have control over the initial transmit power level of mobiles within the network. Therefore mobiles powering on or mobiles experiencing fast fading may sometimes trigger the AGC. To compensate for the added attenuation, a digital gain was added to the backhaul receive path after the channel select filtering in the digital downconverter of the receiver to reverse the effect of the analog attenuation. The analog attenuator prevents the ADC from saturating, and the channel select filter removes the strong signal before the digital gain is applied.

The digital gain mitigated the steady-state impact of the AGC triggering. There was also a second-order concern regarding the transient period during which the receive signal is distorted while the analog attenuation is being altered. The algorithm used to adjust the AGC attenuation was carefully designed to dampen decreases in attenuation so that signals with a high burstiness or peak-to-average ratio do not cause overly-frequent changes in the attenuation settings. These changes together reduced the frequency of AGC backoff events to a sufficiently low level that it was no longer the dominant source of packet errors.

## 4. FIELD TRIAL

### 4.1. Experimental environment

To assess the practical usefulness of the backhaul waveform, a field trial of the wireless backhaul was conducted on a commercial IDEN network with PSTN connectivity. Figure 7 shows the three sites used in the trial along with elevation data for the surrounding terrain. Three existing IDEN sites being backhauled by T-1 lines were chosen for the trial. The node located in the center of the group, roughly equidistant from the other two nodes, was chosen to be the hub site. The remaining two sites are designated remotes. Additional BRs were placed at the



**Figure 8: Hub equipment**

remote sites for the trial, and these BRs were wirelessly backhauled to communicate with an ISC added at the central site for the trial, thus eliminating the need for a dedicated ISC and T-1 co-located at the remote sites.

The link distances are approximately 5 miles and 6 miles between the hub site and remote sites 1 and 2, respectively. Two different hub backhaul antennas were used during the field experiment. The first one (Antenna A) was a vertically-polarized 30" omni antenna with 3 dBd of gain. The second one was a 12" whip antenna. For the field trial, the hub antennas were chosen with size being the primary concern; a small antenna was preferable since it would less significantly alter the external appearance of the site and present less visual pollution at the sites.

To help minimize interference between the IDEN and backhaul waveforms, the frequency plan was carefully chosen to provide 11MHz of isolation in frequency between the two waveforms. The center rack in Figure 8 shows the assembled backhaul equipment at the hub, which due to the frequency-reversed nature of the waveform at that location cannot share RF equipment between the IDEN and backhaul waveforms. The equipment at the remote sites resembles the rack in the figure, but without the backhaul-specific RF front end, duplexer, and PA.

#### 4.2. Measurements

Propagation loss for the backhaul links was measured at single frequencies in the lower and upper SMR bands. The measurement was obtained by comparing the received signal level in the field with a reference value obtained in a laboratory setup and correcting for differences in the two set-ups. Table 1 lists the correction factors applied to the downlink measurements.

In general, the measurements revealed very favorable propagation conditions for the backhaul in the field experiment. For example, between the hub and remote site 1

an average propagation loss of 113.8 dB was measured at a frequency of 808.110 MHz. By comparison, the free-space loss at the same frequency and link distance would be 108.7 dB.

At the hub site, isolation measurements between the collocated cellular and backhaul antennas were made. The backhaul antenna was mounted on top of a shelter approximately 10 feet off the ground while the cellular antennas were mounted on a tower at a height of about 100 feet. In addition, the cellular antenna used for the measurements was pointing away from the shelter. Under those conditions, an isolation loss of >80 dB was measured.

Rough measurements of bit-error rate (BER) over the backhaul were also made. In the uplink the backhaul BER performance was as expected. Between the hub and remote site 1, no uplink errors were observed over a total of 243 million bits at an average SNR of 21.2 dB. Similarly, between the hub and remote site 2, no uplink errors were observed over 33 million bits at an average SNR of 22.5 dB. In the downlink the BER was observed to be limited at approximately  $10^{-4}$  due to interference from unlicensed sources.

Round-trip latency over the backhaul was measured in the laboratory using fping. In a test during which 10,000 packets (each 84 bytes long) were transmitted, the average round-trip latency measured was 38.1 ms and the standard deviation was 2.1 ms. The backhaul latency was dominated by 20 ms of sample buffering at the physical layer.

Parameter	Value
Tx duplexer loss	1.1 dB
Tx antenna gain	5.1 dBi
Cable loss	1.4 dB
Rx antenna gain	14 dBi
Rx duplexer loss	2 dB

**Table 1: RF Gains and Losses for Hub to Remote Backhaul Links**

#### REFERENCES

- [1] L. Chen, N. Martin, C. Cabrera, A. Srinivas, "Joint Topology Design and Frequency Allocation for Fixed Wireless SDR Networks", IEEE Workshop on Networking Technologies for SDR Networks, SECON 2006, Reston, VA, September 2006.
- [2] "iDEN Technical Overview", Motorola Network Solutions Sector, Aug 2000. <http://developer.motorola.com/>