

## Multimedia Broadcasting over the Internet: Part I

**Borko Furht**  
*Florida Atlantic  
University*

**Raymond  
Westwater**  
*Future Ware*

**Jeffrey Ice**  
*Pipe Dream*

Internet broadcasting, referred to as *webcasting*, is coming of age. In addition to reprocessed audio or video that's transferred from radio or TV to the Internet, webcasting now also means broadcasting new, original content—sometimes live—on the Web. Taking advantage of streaming audio and video technology, site producers can bring real-time sound and vision to the Web. With the present technology, audio and video must be compressed almost to the breaking point to squeeze it through a 28.8-Kbps modem line, which means plenty of people will find it's not worth waiting for.

However, this problem hasn't stopped millions of people from downloading viewers and seeking out webcasts. Listening to music or watching video straight off the Internet (Web) still creates a strong enough buzz that people overlook shortcomings like crackly audio, slow download times, and grainy pictures. Consequently, a number of Internet radio stations offer commercially appealing programs to an international audience.

The Internet protocols that transmit this data require individual connections between servers (or senders) and their clients (receivers). The proliferation of such connections proves quite expensive, because they consume very high network bandwidth and processing power at the server. Internet radio stations employ networks of increasingly expensive servers.

Although the industry is still in the early stages of webcasting, we can already foresee what the Internet will offer a few years down the line: clear, crisp audio and full-screen, high-quality, on-demand video. We've developed a technology that provides all these required features for Internet webcasting. This technology consists of

- IP Simulcast, a new Internet broadcast protocol, which provides inexpensive, efficient, and reliable audio and video broadcasting.

- New audio and video compression algorithms, which allow real-time audio and video transmission of data at low bit rates and with high quality.

In this article, we describe a new Internet broadcast technology. In Part II (next issue), we'll present a video compression algorithm for ultra-low bandwidth applications.

Internet broadcast based on IP Simulcast

Three fundamental methods exist for transmitting data on the Internet: IP Unicast, IP Broadcast, and IP Multicast.

*IP Unicast* transmits data (or a packet) from a sender to a single receiver. *IP Broadcast* sends data from a sender to an entire subnetwork. *IP Multicast* enables the delivery of data from a sender to a set of receivers that have been configured as members of a multicast group in various scattered subnetworks.

Radio and television broadcast applications require a one-to-many data distribution model in which data flows from a single sender to many receivers simultaneously, but not the whole subnetwork. Therefore, present audio and television broadcast applications typically use IP Unicast or IP Multicast.

We developed a new technique, called IP Simulcast, for transmitting data over the Internet from a sender simultaneously to multiple receivers. Here, in this issue, we describe the basic principles of IP Simulcast and its technical details.

IP Simulcast shows significant advantages over existing techniques, including IP Unicast and IP Multicast. For example, it resolves all the issues and problems involved in implementing the IP Multicast. Similar to IP Multicast, IP Simulcast reduces the server (or sender) overhead by distributing the load to each client (receiver). Each receiver becomes a repeater, which rebroadcasts

its received content to two child receivers (repeaters), forming a broadcast pyramid, as illustrated in Figure 1.

This method significantly reduces the needed network bandwidth for the server-sender because the server sends just one copy of the data, which the receivers-repeaters then rebroadcast. Thus, the cost of service provision is borne by the receivers (rather than the sender), who have typically paid for the fixed, often unused, bandwidth. In this way, the IP Simulcast concept provides broadcast functionality at a lower cost than IP Multicast. Unlike IP Multicast, which requires special routers for its implementation and several other features, IP Simulcast doesn't have any special requirements for its implementation.

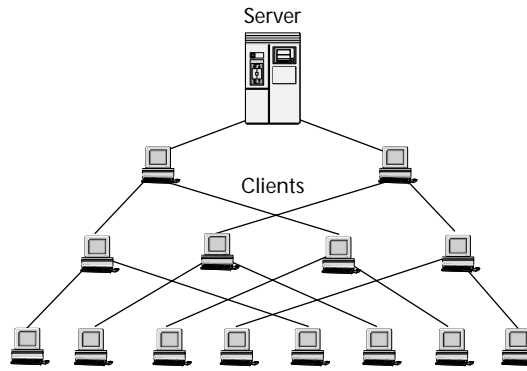
The number of clients in the IP Simulcast pyramid grows as a binary tree. A one-tree-level pyramid has two clients, a two-level pyramid has six clients, and so on. The number of clients in the  $n$ th level equals  $2^n$ . For example, for a broadcast system with 10 levels, the number of clients in the last level is  $2^{10} = 1,024$ , and the total number of clients in the pyramid is  $1,024 + 1,022 = 2,046$ . For the IP Simulcast pyramid, consisting of 16 levels (131,000 nodes), the end-to-end delay to the nodes at the bottom of the pyramid equals 3 to 4 seconds.

The repeater-receiver performs conventional client functions, including error recovery and detection of the lost connection. Consequently, unlike IP Multicast, IP Simulcast provides guaranteed delivery of packets. IP Multicast services make no provision for error recovery. The lost packets must be either ignored or recovered from the server at the cost of increased server bandwidth.

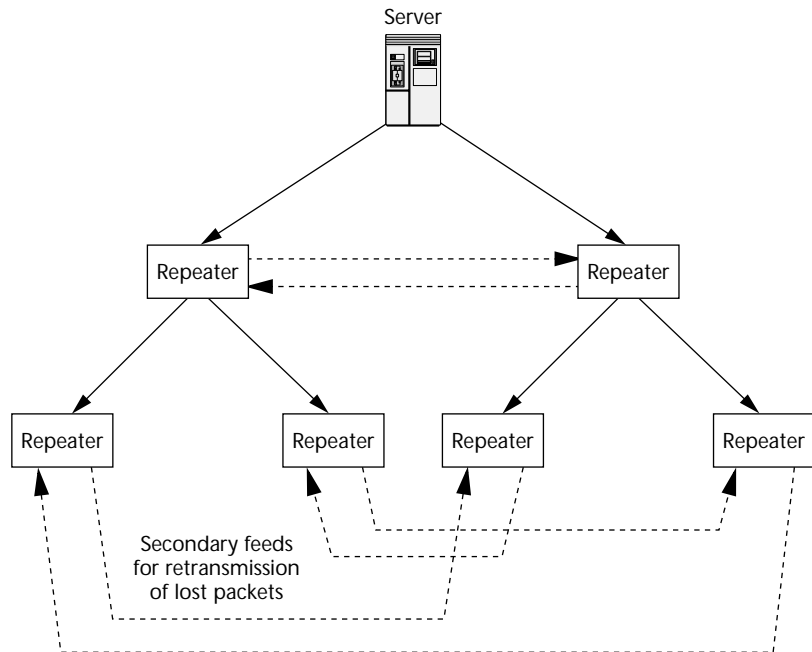
IP Simulcast uses a radically different model of digital broadcast, referred to as the *repeater-server model*. In this model, the server manages and controls the interconnection of repeaters. While the server may resemble a conventional server, the repeater contains server functions in addition to conventional client functions. In essence, each repeater not only plays the data stream back to its audience, but also transmits the data stream to two other repeaters.

As Figure 1 illustrates, IP Simulcast builds on the new repeater-server model. The server sends the data only to two repeater-receivers, and then the packets are rebroadcast by each level of repeaters to the next level. This process builds a pyramid network that the server manages and controls.

In addition, to assure a reliable data transmis-



**Figure 1. Broadcast pyramid applied in IP Simulcast.**



**Figure 2. The IP Simulcast repeater-server relationship.**

sion, retransmission of lost packets or packets with errors is requested through secondary feeds (the dashed lines in Figure 2).

The server functions include

- **Digitization of the program source.** A typical source program might include analog audio and video. These analog program sources are digitized into streams of time-varying data.
- **Synchronization of the digital source.** Streams of time-varying data may come from various sources such as digitization of analog sources, stored compressed data on disk, and digital data from animation programs, authoring programs, or other sources. Source programs can be interrupted, overlaid, or otherwise synchro-

nized with advertising spots, and they can be scheduled throughout the day. The various sources of digital data must be synchronized and time-stamped for playback.

- *Compression of the source.* Each stream of time-varying digital data can be compressed to reduce its size and transmission time. The compression technique is a trade-off among various factors including the compression ratio, perceived quality, complexity of compression and decompression, scalability, and noise immunity.
- *Collection of the compressed source into transmission packets.* IP transmission is a packet-based protocol. The data is collected into IP packets in preparation for transmission. Compressed data can be represented by several alternative packetization schemes to adapt to different speed transmission lines or computers of different power. Each of these packetization schemes could be used to feed an alternate pyramid of repeaters.
- *Transmission of compressed source transmission packets.* Two feeds are supported, each to be received and retransmitted by its destination repeater.
- *Connection of repeaters.* Each repeater sends a request to the server asking to be serviced with the transmission stream. The server responds by selecting an available repeater to serve as the requesting repeater's source. The transmission stream is then fed to the requesting repeater. The server also selects a secondary feed for the requesting repeater. Error retransmission occurs over this secondary feed.
- *Collection of statistics.* The server monitors the construction and breaking of connections.

Each repeater-client has responsibility for collecting the transmitted data streams and playing them back to its audience. The repeater-clients' functions include

- *Establishment of connections.* The repeater-client issues a connection request to the server. The server will establish an individual connection to the repeater-client.
- *Reconnection.* The client must determine if a connection broke and attempt reconnection.

- *Caching of packets.* Received packets must be sequenced and cached to locate missing packets.
- *Retransmission requests.* Requests are issued to the repeater-client's secondary feed to request retransmission of missing packets.
- *Error recovery.* In cases where a packet can't be recovered, the repeater-client must perform some recovery action (play silence, replay the last packet, degrade quality, and so on).
- *Decompression of received data stream.* The received data is decompressed in anticipation of playback.
- *Playback of data streams.* The decompressed data plays back to the repeater-client's audience.
- *Synchronization with the server.* The playback rate must match the server's capture rate to avoid overflow or starvation of the repeater-client's buffers. The repeater-client must adapt to the small differences in playback rate that might exist.

The repeater-transmitter performs some conventional server functions, including

- *Transmission of compressed source transmission packets.* The system supports two feeds, each received and retransmitted by its destination repeater.
- *Retransmission of error packets.* Each repeater-transmitter supports a secondary feed. On request, a missed packet is retransmitted to the secondary feed's destination.

The broadcast system subdivides into fractional streams for transmission purposes. The system organizes repeaters for each fractional stream into a binary tree, propagating the fractional stream through all repeaters. Each repeater collects fractional streams into a single stream, which causes a superposition of the binary tree into a single "bush" that represents the transmission of the full system. The topology of the superposition is chosen so that the two levels of a fractional tree become separated by one-half the width of the stage in the tree. This topology ensures that no repeater starves due to the failure of a single feeding repeater.

Each repeater collects packets into a buffer, which helps compensate jitter delays. Additional

buffering performs error recovery. After the jitter and error delay has played out, the received packets are broadcast to the next level.

Error recovery consists of two distinct phases: error recovery and retry service. During error recovery, queries occur in round-robin fashion to repeaters in the previous stage. During the retry service period, retry requests from the subsequent stage are serviced. The system then places transmitted samples in a playback buffer. Playback synchronizes to the rate at which packets are received to avoid playback buffer overflow and underflow.

An unassigned repeater issues a connection request to the server-administrator to join the broadcast. The server-administrator acknowledges the request and queues the repeater for connection. If the repeater hasn't been connected by the time its queue entry times out, the server-administrator issues fractional feed requests to the last complete stage, starting a feed to the repeater.

When a repeater-receiver wants to leave the broadcast, it issues a disconnection request to the server. If the queue of the repeaters waiting for connection isn't empty, a repeater is selected from the queue, and the server issues fractional feed requests to the parents of the terminating repeater. On the other hand, if the repeater connection queue is empty, the oldest node on the bottom stage serves as the replacement node. In the event of node failure, the children of the node report the failure to the server.

IP Simulcast requires a little more bandwidth than the traditional IP Multicast, but much less than IP Unicast. Compared to the reliable IP Multicast, IP Simulcast requires about the same network bandwidth.

#### Comparison with other approaches

We compared the IP Simulcast approach for audio broadcasting with the IP Unicast and IP Multicast systems. We assumed an audio broadcast system that continuously broadcast 16 Kbps to a maximum of 10,000 clients-receivers. For the comparison we used the following assumptions:

- When calculating the server bandwidth, we assumed a 1 percent error retransmission and ignored control overhead. The server band-

**Table 1. Comparison of features of various techniques for audio broadcasting.**

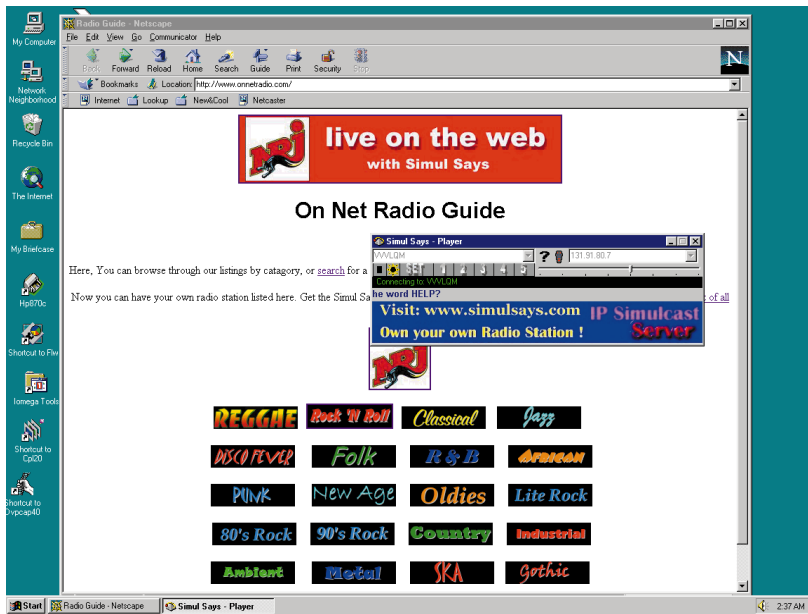
Features	IP Unicast	IP Multicast	IP Simulcast
Server Bandwidth	162 Mbps	1.62 Mbps	16.2 Kbps
Bandwidth Cost	\$100,000 per month	\$20,000 per month	\$100 per month
Error Recovery	By server	By server	By client
Initial Server Cost	\$53,000	\$8,000	\$5,000
Client Reachability	Any IP address	Only clients in proprietary network	Any IP address
Implementation Issues	Cannot scale to serve increasing number of clients	Requires all intermediate IP Multicast routers. Requires special network card and software to support IP Multicast.	Easy to implement, does not require any special cards or routers

width for IP Multicast reflects error retransmission from the server.

- We calculated bandwidth cost by assuming \$1,000 per T1 connection per month (1.5 Mbps).
- For IP Simulcast, only one server manages and controls the broadcasting pyramid and compresses audio. The server cost is \$5,000.
- For IP Unicast, we used 16 servers each costing \$3,000.
- For IP Multicast, we used one server to support transmission to the network and one to service error retry requests (total cost \$8,000).

Table 1 compares these three approaches. In summary, the IP Simulcast-based solution for Internet broadcasting provides a number of advantages compared to existing technologies including IP Unicast and IP Multicast, summarized as follows:

- *Lower cost.* Due to inexpensive server and network requirements, the IP Simulcast-based solution costs less than the other solutions.
- *Better flexibility.* IP Simulcast provides a general solution and its broadcasts are received regardless of the physical solution, medium, connection noise, or the receiver's network provider.
- *Higher quality.* The IP Simulcast-based solution



**Figure 3. Screen shot of Pipe Dream's Radio Player and Radio Guide, which use IP Simulcast for audio broadcasting.**

functions in the unreliable Internet environment and provides built-in error recovery and quality control.

#### Potential applications

Due to its simplicity, easy implementation, efficiency, and inexpensive initial cost for the server and network bandwidth, IP Simulcast proves ideal for many current and potential webcast-based applications. Since IP Simulcast suits radio and television broadcasting well, Pipe Dream (West Palm Beach, Florida) created the first two IP Simulcast applications for these media.

Other potential applications include distance learning, electronic software distribution (including software updates), real-time broadcasting of critical data (like stock prices), database replication and file transfer, videoconferencing, and many others.

Consider the market for radio broadcasting on the Internet. The radio on the Internet application offers very attractive features to the audience, such as scheduled programs, supplementary data on the scheduled programs, and interactive services. Thus, more than 27 percent of the 11,000 radio stations in the US already have Web sites according to the Massachusetts Institute of Technology (MIT).

The number of radio stations offering online radio programs has also increased in the last several years, from 50 (in 1995) to 741 (end of 1997), of which 341 are in the US. MIT forecasts that 1,500 to 2,000 stations will be webcasting by the end of 1998.

Based on the IP Simulcast protocol, we have developed two applications—SimulSays for radio broadcasting and SimulSees for television broadcasting.

The SimulSays application uses the IP Simulcast protocol to let radio stations broadcast (webcast) radio programs to an unlimited number of clients using a simple, inexpensive server and small server bandwidth. Thus, the broadcaster incurs a low initial cost to begin broadcasting radio programs. Besides the IP Simulcast protocol, SimulSays applies an audio compression technique, capable of compressing audio while maintaining its high quality. SimulSays also includes an advertising banner function, which broadcasts advertising messages and transmits and displays supplementary data such as maps, telephone numbers, election graphs, dates for various events, and so forth.

Users can download Pipe Dream's Radio Player from the company's Web site at <http://www.simulsays.com>. Radio Guide, linked to Pipe Dream's radio player, can be used to listen to test radio stations broadcast from the Pipe Dream server. (See Figure 3.)

SimulSays comes with a chat function, an upgraded SimulSays application, and provides interactivity among the receivers-clients via a chatboard. We expect this application to enrich radio transmission and make it very attractive due to its interactivity. The clients, who listen to radio programs, can interactively exchange messages amongst themselves.

SimulSees also uses the IP Simulcast protocol to provide television broadcasters with an efficient and inexpensive solution for broadcasting (webcasting) television programs to a large number of clients. Similarly, the initial cost for broadcasters includes a simple, inexpensive server and small network bandwidth. Besides IP Simulcast technology, SimulSays uses a new real-time video compression algorithm, which enables live video webcasting at low bit rates. We'll discuss this technology in the Project Reports section of the next issue of *IEEE MultiMedia*. MM

*Readers may contact Furht at the Department of Computer Science and Engineering, Florida Atlantic University, 777 Glades Rd., Boca Raton, FL 33431, e-mail borko@cse.fau.edu.*

*Contact Project Reports editor Harrick Vin at the University of Texas, Austin, Department of Computer Sciences, Taylor Hall 2-124, Austin, TX 78712-1188, e-mail vin@cs.utexas.edu.*