

Cisco Telepresence Vision Becomes a Reality with Blackfin®-based HD Video Codec



As the worldwide leader in the infrastructure equipment that has made today's Internet possible, Silicon Valley-based [Cisco Systems](#) arguably knows more than anyone about how companies can leverage networked-based resources to improve business results. And now, enabled by the programmable flexibility and price/performance leadership of Analog Devices, Inc. (ADI) [Blackfin® processors](#), Cisco's revolutionary new "telepresence" visual collaboration solutions promise to drive business productivity to new heights.

Futurists have long envisioned telepresence as a technology-enabled experience which accurately replicates being present in a remote location. In order to credibly adopt the term "telepresence" for its new IP-based visual collaboration solutions, Cisco needed to build a product with extremely high-definition (HD) video (no onscreen artifacts) and best-in-class latency (no perceptible time lag) – all while communicating over the limited bandwidth and uncertain determinism of existing enterprise network infrastructures.

"...while fixed-function ASICs or FPGAs could have met the system's technical requirements, Blackfin's programmability let Cisco's system-oriented engineers minimize development risk and develop a working solution in a very tight time frame."

Only by providing a remarkably new level of experience excellence to users of the new equipment could Cisco enable more effective remote business relationships between colleagues, partners, vendors, and customers. The algorithmic processing power needed to implement what would be the world's first ultra high-definition 1080p/30 low-latency HD video codec was formidable, but even more so when Cisco developers applied cost, heat dissipation, and programmability requirements to the telepresence product definition.

Cisco ultimately employed an array of ADI [Blackfin ADSP-BF561](#) processors to handle the compute-intensive nature of the telepresence application. With dual symmetric 600 MHz Blackfin cores, the ADSP-BF561 is a powerful, yet cost-effective processing engine – exactly the right combination for Cisco's advanced video processing application. And while fixed-function ASICs or FPGAs could have met the system's technical requirements, Blackfin's programmability let Cisco's system-oriented engineers minimize development risk and develop a working solution in a very tight time frame. The Blackfin processor's [world-class power management architecture](#) also meant the system would not require any thermal dissipation treatment.

What is Telepresence?

Telepresence is a set of technologies that enable users to conduct and collaborate as if all of the participants are physically in the same room. Where traditional videoconferencing solutions were often impossible to set-up, and often failed to make users feel comfortable, well-designed telepresence solutions provide immediate connection to the remote location, a natural environment,

with realistic looking, life-size participant video images, and high fidelity, low-latency, directional audio that is perfectly synchronized with facial expressions and which follows participants as they move about in the conference room.

Cisco's Blackfin-based telepresence offerings include two models: the CTS1000, which is designed for an executive's office, and the CTS3000, which is a "drop-in" full solution for a conference room. The solutions comprise everything that's required to conduct a life-like meeting in two places. The conference room system includes ceiling lights to reduce glare on the display screens, an elliptical table, chairs, three 65-inch plasma displays, three cameras, and a system cabinet, which houses the Blackfin processors.

All this is perfectly replicated in the remote "partner" conference room to create a remarkable illusion of contiguity –the two conference tables appear seamlessly joined as one. The CTS1000 system comprises one 65-inch plasma display, one high definition 1080p camera, and a system cabinet. A company need only provide an empty conference room and an Ethernet jack. (Cisco will perform a network audit to ensure that a customer's enterprise network has adequate bandwidth with network gear configured to ensure the proper quality of service from end-to-end throughout the network.

Cisco's leadership in "Unified Communications" and IP Telephony contribute to the easy setup and use of the telepresence system. With a couple clicks, users immediately initiate a "face-to-face" meeting with colleagues thousands of miles away, and the system's HD video is so life-like, users say they often forget the other people are not actually in the room.

High-Definition Video

One of the most important requirements of Cisco's telepresence system, if not the most important aspect, was the need to support HD video encoding and decoding for standard, relatively narrowband enterprise networks. Cisco specified that it needed to be able to scale from 720p upwards to 1080p/30 video modes. Currently, 1080p/30 is the most brilliant resolution with the number "1080" representing 1,080 lines of vertical resolution, and the number "30" representing 30 frames per second. Ultra High-Definition video is sometimes also referred to as 1080p/30. In addition, Cisco wanted to minimize bandwidth requirements down to 3 to 9 Mbps at 720p and 9 to 12 Mbps at 1080p for IP QoS connections.

An advanced video codec standard, H.264, which is known for achieving very high data compression at lower bit rates than previous standards, compresses raw video to fit over economical, enterprise WAN links. Cisco's codec is a video encoder/decoder based on the H.264 video codec standard, and it provides best-in-class latency and up to 1080p/30 video resolution occurring in real time. By minimizing the HD video encode/decode latency Cisco's solution gives more "latency budget" back to the network to allow swift deployment with a minimum of hardware and software upgrades for the network to handle the new application. With dual symmetric 600 MHz high-performance Blackfin cores, Blackfin ADSP-BF561 processors were the ideal choice for Cisco's exceedingly challenging and complex video application. Cisco's video codec

functionality is distributed across a multiprocessor farm of Blackfin ADSP-BF561 processors, delivering over 0.5Tera-Instructions-per-second of processing muscle, which drives the video subsystem at truly best-in-class performance levels enabling practical solution deployment.

Often chosen for high-end multimedia and telecommunications applications, Blackfin ADSP-BF561 processors comprise two independent Blackfin processor cores. Each core comprises dual 16-bit MAC state-of-the-art signal processing engines, two 40-bit ALUs, four 8-bit video ALUs, a 40-bit shifter, plus the advantages of a clean, orthogonal RISC-like microprocessor instruction set, and single instruction, multiple data (SIMD) multimedia capabilities –all in a single instruction set architecture.

Easy Programming/Low Power

The Blackfin ADSP-BF561 processor's RISC-like register and instruction model made for easy programming, which Cisco said reduced development costs, risks, and time-to-market for the product. The Blackfin's programmable architecture also allows for over-the-network feature upgrades without costly ASIC hardware re-spins. Cisco also leveraged ADI's software and hardware development tools to speed time-to-market and reduce costs. The tools included the [VisualDSP++ development environment](#), which lets programmers develop and debug an application, and emulator hardware to test and debug the hardware and software.

Cisco leveraged many of the general-purpose digital imaging peripherals that are integrated on the Blackfin ADSP-BF561 processor to arrive at a highly elegant system-on-chip design for the telepresence system codec. General-purpose peripherals of the Blackfin ADSP-BF561 processor include: two parallel input/output peripheral interface (PPI) units supporting ITU-R 656 video and a glueless interface to analog front-end ADCs; two dual channel, full duplex synchronous serial ports supporting eight stereo I2S channels; dual 16-channel DMA controllers and one internal memory DMA controller; 12 general-purpose 32-bit timer/counters with PWM capability; a SPI-compatible port, a UART with support for IrDA® dual watchdog timers; 48 programmable flags; and an on-chip phase-locked loop capable of 1x to 63x frequency multiplication.

Blackfin's low power dissipation helped Cisco limit the telepresence system's thermal treatment to a small, silent fan providing airflow in the system cabinet, crucial to preserving the ambiance of conference room environments. Designed in a low power and low voltage design methodology, Blackfin processors enable developers to vary both the voltage and frequency of operation to significantly lower the overall power dissipation. Cisco describes the result as a "cool running farm of Blackfin processors" that drives its video subsystem –no heat sinks were required in the design.

Thanks to Blackfin, Cisco now has the highly compressed HD video it required for its next-generation telepresence systems. The result is reduced travel time, expense, and stress for road-weary corporate executives and employees who can now meet with partners, colleagues, or customers using IP-based video that is very nearly as real as face-to-face communication to increase productivity and collaboration like never before.