THE BROADBAND REVOLUTION

Henry Samueli
Broadcom Corporation

Henry Samueli enchanted the Hot Chips 11 audience with his keynote address on the implications of broadband communications and the connected home. Samueli cofounded the Broadcom communications chip company and is its chief technical officer and vice president of research and development. He received his undergraduate degrees and doctorate in electrical engineering from UCLA. He has worked at TRW, founded PairGain (a company that initially worked in the ISDN market, and then went on to become the market leader in HDSL), and later joined the UCLA faculty. Broadcom is one of the most successful companies in the emerging silicon-for-telecommunications market.

The broadband revolution truly is of the same magnitude as the industrial revolution of the 1800s. The changes we're going to see over the next decade will principally come from the connected home. We'll see bandwidth increasing a thousandfold from today's voice band 56-Kbps modems to tens of megabits per second through DSL and cable modems, and other satellite and wireless technologies.

To have an impact on the millions of people who do not have computers, we'll need to turn the television set into a medium that somewhat simulates what they would now get with a computer and add fully interactive video at broadband speeds.

The future connected home will have numerous service providers sending a broad pipe to the home over satellite, cable, DSL, or terrestrial wireless. It doesn't matter how we connect the home; we'll see a broadband gateway in the home, much like today's office that has a broadband router connected to the Internet.

The connected environment doesn't necessarily have to be wired, but cable and phone lines are a natural choice, especially in the US, and will be first to market with the widest deployment. There will be no rewiring; clearly if we have to rewire, we'll never have a market. We'll also need wireless mobile connectivity within the home for all the devices we'll use. The telephones will be converted to Internet videophones, and MPEG video will be delivered over the Internet. That's not that far off into the future.

As we are well aware, the TV set is being upgraded substantially. Certainly the HDTV format with the 16×9 aspect ratio and high-resolution display will finally allow us to read Web pages on the TV set in a reasonable format. We'll have fancy on-screen displays and videophone pop-up screens with a TV camera sitting on the TV set so we can make video calls to Grandma.

So what are the challenges?

Why hasn't this been done before? It's not that we didn't know what we needed; we just couldn't practically implement it for many reasons. First of all, the last-mile connection to a home was never intended for broadband transmission. The telephone network was designed for analog voice, which is a 3-KHz bandwidth signal with severe amplitude and phase distortion, severe cross talk, and ingress from radio signals. It's a very difficult challenge for broadband transmission. Even the coaxial cable network was really never designed to be an interactive broadband digital network.

We needed very sophisticated communications algorithms, modulation and coding techniques, and adaptive signal processing techniques. The transmission channels are also time varying. This required very fancy real-time processing. Until fairly recently we didn't have the ability to put that much signal processing horsepower into a cost-effective package.

Today we have multimillion-transistor chips, and over a billion-operations/sec DSPs

0272-1732/00/$10.00 © 2000 IEEE

on a single device. Furthermore, it's an analog world, so we have to interface to analog signals. Since the chip is generally 80% to 90% digital, the most cost-effective way is to take that high-precision, high-speed analog front end and put it into the same digital CMOS substrate as the remaining circuitry. Then we don't burden the entire chip with an expensive process. Complete system-on-a-chip solutions allow this industry to exist, because we can drive the cost down so much through system integration on single pieces of silicon.

Direct broadcast satellite

Here, I focus entirely on the communications aspects, not on the back-end MPEG video processing and video decoding parts.

Current CMOS integration lets us put all the needed functionality into a single chip with an interesting level of mixed-mode integration. An outdoor satellite dish receives a 12-GHz signal from a nearby satellite. A block converter in the dish translates the signal to 1 or 2-GHz and feeds it into the indoor satellite set-top box. The front end of the box starts with the RF signal coming in through a low-noise amplifier and going into a tuner. Then using a 512-MHz PLL and a mixer, the front end further down converts, filters, and digitizes the signal. The A/D runs at 128 MHz with 8 bits of resolution. The back-end DSP handles the final demodulation, clock recovery, carrier recovery algorithms, and back-end forward error correction with some RAM. All of this must be fairly precise, and the PLL must have very low phase noise.

This relatively modest chip accommodates the three worldwide standards for satellite communications: Direct TV or DSS, European and Japanese DVB, and PrimeStar, another US standard. It can accommodate anywhere from 2 to 90 Mbps and has about 1.2 million transistors in 22 mm^2. It's a very low-cost 0.35-micron, 3.3-V, single-poly, quad-metal CMOS device.[1] Figure 1 shows the chip's analog front end, which accounts for 15% to 20% percent of the chip, in the lower right, the receiver, then the FEC decoder on top.

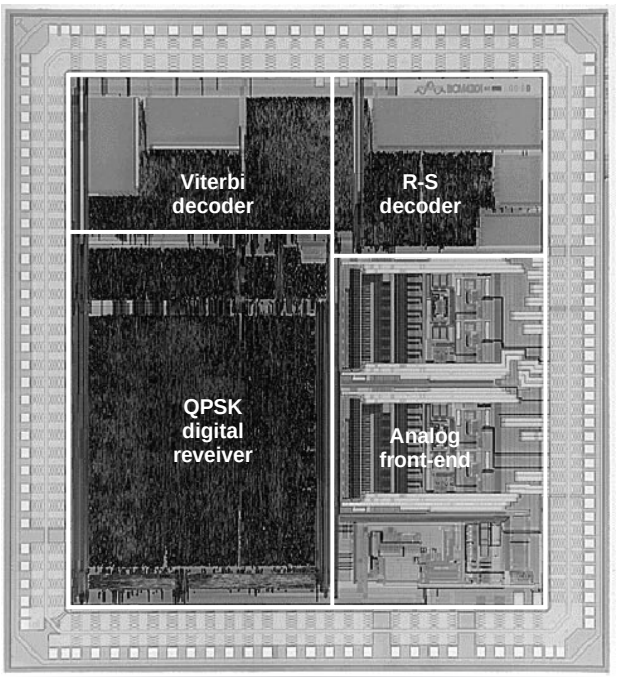

Figure 1. QPSK satellite receiver IC.[1]

Digital subscriber lines (xDSL)

The DSL's twisted-pair environment involves a telephone central office with its equipment interfacing to the traditional circuit-switch telephone network, and dedicated lines to each house. It's not like the Internet's packet network. The incoming signal splits into different frequency bands: 1) a low, 3-KHZ band for voice with a POTS splitter and 2) a data band, which could be an ADSL line that feeds an Internet router onto the Internet. Both signals are combined and operate in different frequency bands.

The problem is that in many cases we're dealing with an often-100-year-old phone network. In the worst case, there is roughly a three-mile link of 18-kilofeet, 24- or 26-gauge cable with unterminated stubs that cause multiple reflections. That means all sorts of noise coupling to handle as well as cross talk from other homes on the street that disturb the signal. Along with coupling from RF sources and radio broadcast antennas, we have a nightmare transmission problem for a broadband signal. Once in the home, we can split it again to the telephones and to the computer.

The telephone companies are now getting around that transmission nightmare by upgrading networks and building them with a fiber backbone in new neighborhoods. With fiber, we still have the central office but add more sophisticated networking equipment for data.

We have an ATM switch—the telcos like

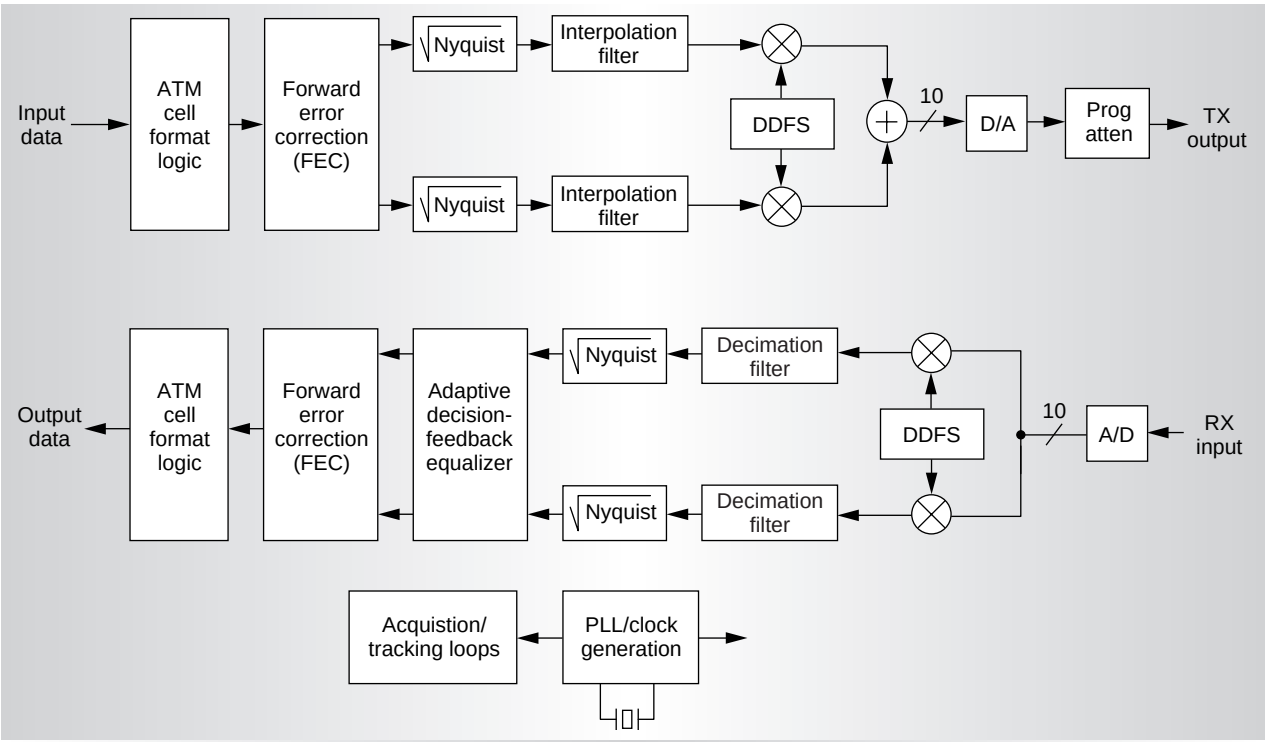

Figure 2. QAM DSL transceiver block diagram.[2]

ATM, where the networking world is clearly all IP—feeding the Internet. We could have video sources and cable satellite servers feeding into this. The data is encapsulated as ATM frames and sent over a universal digital terminal, which takes the analog voice and digitizes it as an all-digital signal into the network. Initially, this would be sent over a 155- to 622-Mbps fiber link. It terminates at an optical network unit, or ONU. Finally, it is converted to less-expensive twisted-pair lines in the home.

Running fiber into the home is very expensive. The last drop is pretty short—a few thousand feet—hence we have less distortion and noise, and can use twisted-pair wires. We can drive more data at higher data rates up to 52 Mbps, which lets us deliver video over phone lines. This gives us a more generic residential gateway into a house: a set-top box that receives broadband digital and splits it into a video signal to the TV, a data signal into the PC, and a voice signal to the telephone.

US West is deploying 400,000 homes in Phoenix based on this VDSL network topology. Figure 2 shows a block diagram of the transceiver chip used by the telco in that deployment for VDSL. It's based on a transmission technology called QAM—quadrature amplitude modulation. There are many different line codes for DSL, and this is just one example. It includes a transmitter and receiver on the same chip, accommodating 52-Mbps data rates. An ATM interface is on the left, so the ATM cells get dropped into the chip, encoded, and modulated with QAM; converted; and sent out to the transformer. On the receive side, the signal comes from the receive transformer, is digitized on chip at up to 50 MHz, demodulated, and corrected for distortion. (See Definitions box, last page.)

The chip contains about a million transistors in about 33-mm^2, a very cost-effective solution, and about 85% to 90% digital.

But, as I mentioned before, there are a lot of challenges in getting ubiquitous deployment of DSL. Although a fiber network helps solve the problem, it doesn't attack the ultimate problem of what's going on inside a house. Any house is pretty much an uncontrolled environment with very poor quality, completely random flat-pair cable. Of course, a house has noise sources from light dimmers, refrigerator motors, fans, and microwave ovens among others to contend with.

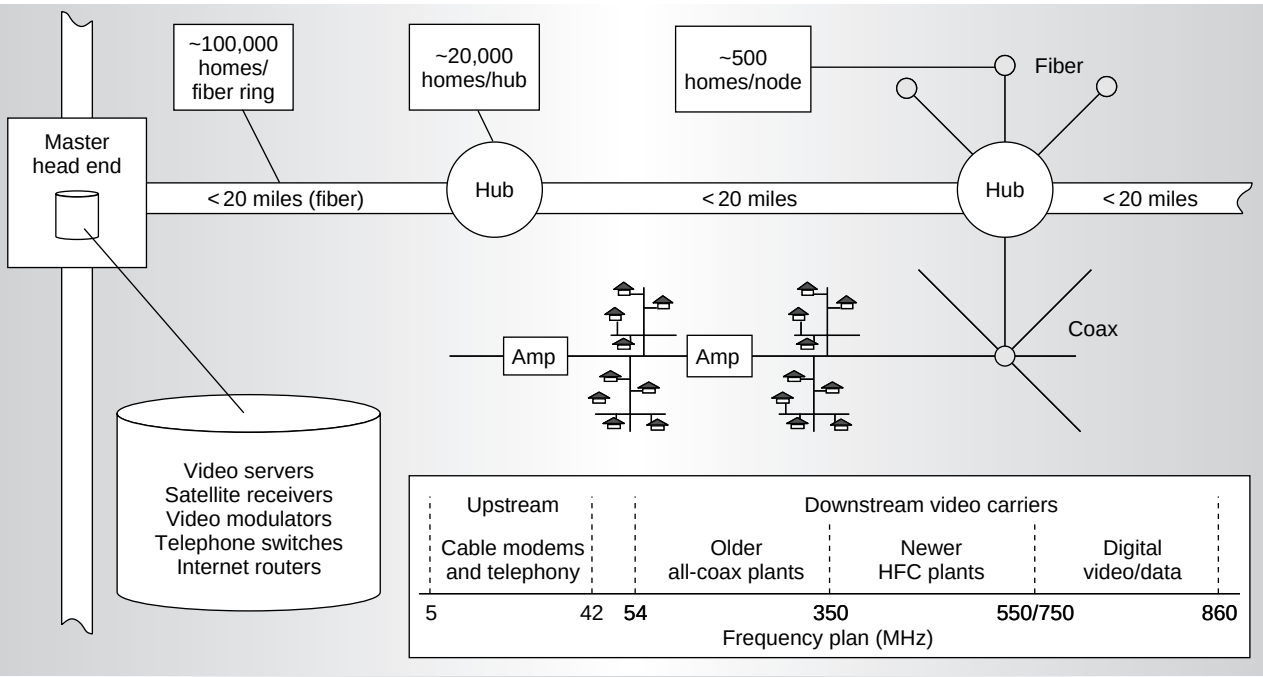

Figure 3. An overview of a hybrid fiber coaxial network architecture.

Digital cable TV/cable modems

Cable TV is distributed over a hybrid fiber coaxial network providing a nearly 1-GHz bandwidth, whereas a phone line has 10 to 100 KHz available. Over 65% of the 100 million US TV households currently subscribe to cable TV. AT&T's huge investments in TCI and other cable companies indicate that it clearly sees that 10 years from now the cable infrastructure and the Internet will become one of the predominant methods for delivering voice and video telephony services in the world.

Figure 3 shows a hybrid fiber coax network with a master head end that includes video servers, satellite receivers, video modulators, phone switches, and Internet routers. Most video signals are downlinks from satellites before they're distributed over cable. One master head end can serve 100,000 homes fed through fiber trunks that get smaller as they feed into hubs serving maybe 20,000 homes each. The last fiber connection goes to a node for conversion to electrical signals onto the coax cable into the neighborhood—a tree and branch type architecture much like we lay out routing on a chip. Amplifiers build up the signal to compensate for too much attenuation away from the fiber node.

Without multiple pairs of wires, we have to time share or frequency share the upstream and downstream traffic. The most common approach is frequency sharing or frequency division multiplexing. The upstream traffic is segregated into the lower frequency band, which is 5 to 42 MHz in the US. The downstream traffic from the head end to the home is the upper band, which is 54 MHz on up, depending on a cable plant's capability. The video traffic is on the high end with several hundred channels of video traffic typically sitting there.

Cable operators have a very compelling business case for converting from analog to digital. Analog cable TV systems typically allocate a 54- to 550-MHz band for standard 6-MHz NTSC TV signals. They can support about 80 analog channels. We replace them with very bandwidth-efficient, 64-level QAM digital and get roughly 27 Mbps of net payload to the user, in a 6-MHz channel. If we use 256-level QAM, we can support a 38-Mbps data rate per 6-MHz channel. With MPEG-2 compression, we can get 3- to 4-Mbps video. So the network literally gets multiplied by an order of magnitude in video capacity, and what used to be 80 channels become 800 channels of capability.

However, telcos have the opposite problem. It's actually a negative business case for them,

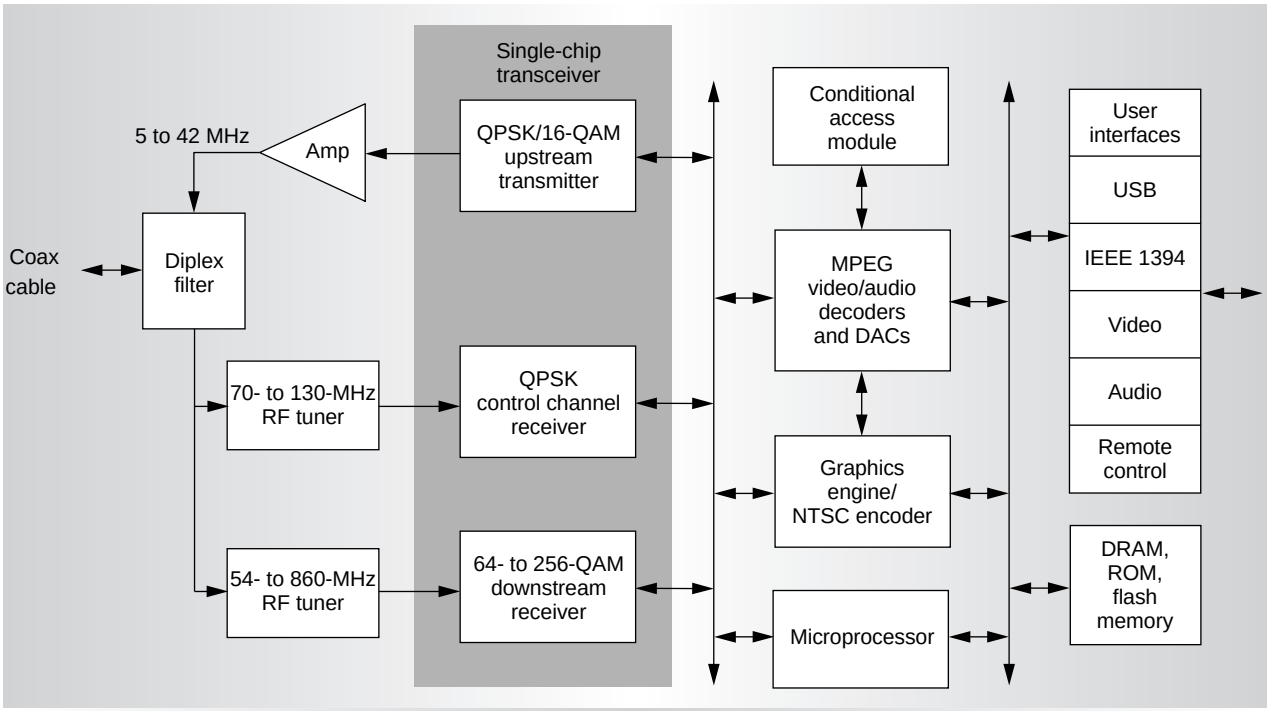

Figure 4. Digital cable set-top box block diagram.

because telcos are very spoiled. They're used to charging businesses $1,000 a month to get a 1.5-Mbps T1 connection, and now we're asking them to give that same T1 connection to the home for $30 a month. They have a very difficult time with that. On the other hand, the cable operators have found that they're only using 10% of the capacity of their network. By converting to digital, they free up 90% of it for new services.

Figure 4's block diagram of a digital cable set-top box shows the rough partitioning of the various functions. The coax cable comes in on the left, and the diplex filter is just a high-pass/low-pass filter that splits the 5- to 42-MHz upstream and 54- to 860-MHz downstream traffic. The RF tuner typically is tin-can discrete hardware, but recent developments on RF integration will help reduce the cost substantially. The video decoder processes the MPEG-2 bitstream, and a control channel receiver primarily functions as what is called access control, where pay-per-view services operate. All that control information is on its own dedicated channel. Then, there's the upstream traffic from the set-top box to the head-end and the back-end functions: the MPEG video decoders, video encoders, DACs, graphics engines, and so on.

The original set-top boxes had 8- or 16-bit microprocessors; now we're seeing high-end set-top boxes with several hundred MIPS microprocessors to handle on-screen 3D graphics. The middle slice in Figure 4 shows the communication functions in the set-top box, the upstream transmitter, the control channel receiver, and the video QAM receiver. This 40-Mbps receiver accepts input at a standard intermediate frequency of about 44 MHz. We digitally demodulate it, an adaptive equalizer handles the channel distortions, and then the signal moves to the FEC back end.

The satellite receiver uses very simple modulation: QPSK, which is a four-point constellation. The QAM receiver can handle up to a 256-point constellation, meaning we have a 16×16 grid of points, or 16 amplitude levels we're sending, as opposed to just two amplitude levels that we send with QPSK. We have to discern one of 16 levels in this receiver, and thus we have to have very precise processing. All the timing recovery loops to carrier recovery loops must have much more accurate estimates of the timing and phase. It's a lot more difficult to design this front end.

The upstream burst transmitter handles a

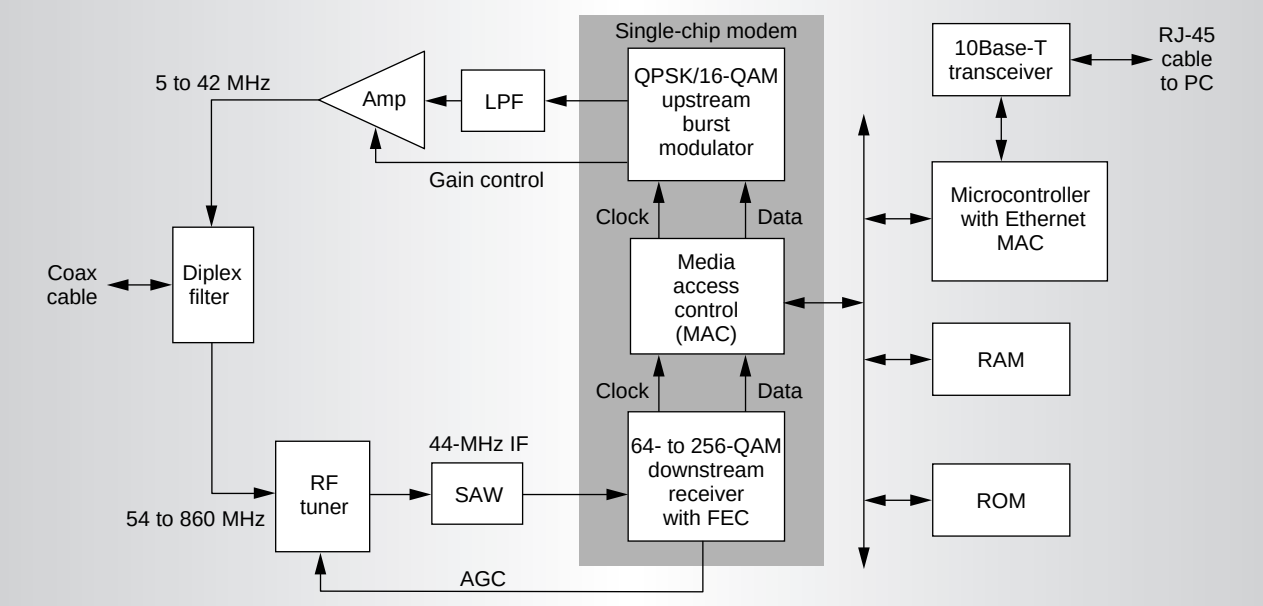

Figure 5. Cable modem block diagram.

lot of preprocessing on each burst, adding a preamble, randomization, and FEC. Then the typically all-digital modulator uses direct digital frequency synthesis, followed by the on-chip 200-MHz DAC running out to the analog front end.

The 43-Mbps QAM transceiver, 2-Mbps QPSK receiver, and 20-Mbps upstream transmitter with a 200-MHz DAC fit in a 0.35-micron, 3.3-V, 64-mm^2 chip holding about 2.3 million transistors. More detail on this was published at ISSCC 99.[3]

Unlike the point-to-point network architecture used by telcos, cable modems are point-to-multipoint networks. A point-to-point architecture has a dedicated line between every subscriber and a port in the central office. When someone dials in, that port is tied up—even if no data is being sent. This is why ISPs automatically log users off when there's no activity after a while; they need that valuable port in their central office.

Cable modems have one downstream channel that goes to all the subscribers in a single cluster of homes. We can chop up the bitstream into time slots, allocating a certain frequency to a group of homes so that each user gets a certain percentage. However, users can only transmit when they have something to send. The upstream channel uses burst-mode FDMA and TDMA structures. So it's a very efficient protocol if we're accommodating lots of users with multiple, varying bandwidth demands. It's ideal for bursty communication, which Internet access clearly is.

Even though users are active, logged in, and on the network, they aren't using up network resources until they transmit data. So they have instant Internet connectivity as long as their computer is powered up. Because there's a band set aside for the upstream, we can put multiple RF carriers within that band, and the cable operator can continue to add more RF bandwidth as the user load increases. We don't have to worry that service is going to degrade.

A cable modem, which is similar to the front end of a set-top box, has a coax cable coming in through a filter that splits the upstream and downstream traffic. (See Figure 5.) On the downstream side we have a tuner that tunes to the appropriate channel and puts out an intermediate 44-MHz frequency through a SAW filter and then into the receiver chip.

On the upstream we have a transmitter with the on-chip DAC, a low-pass filter, and a line driver that feeds the signal out into the network. In the back end—shown in the middle of Figure 5—is the MAC, which uses a fairly sophisticated protocol that's been defined as a standard known as DOCSIS (Data Over Cable Service Interface Specification). This traffic cop allocates all of the bandwidth.

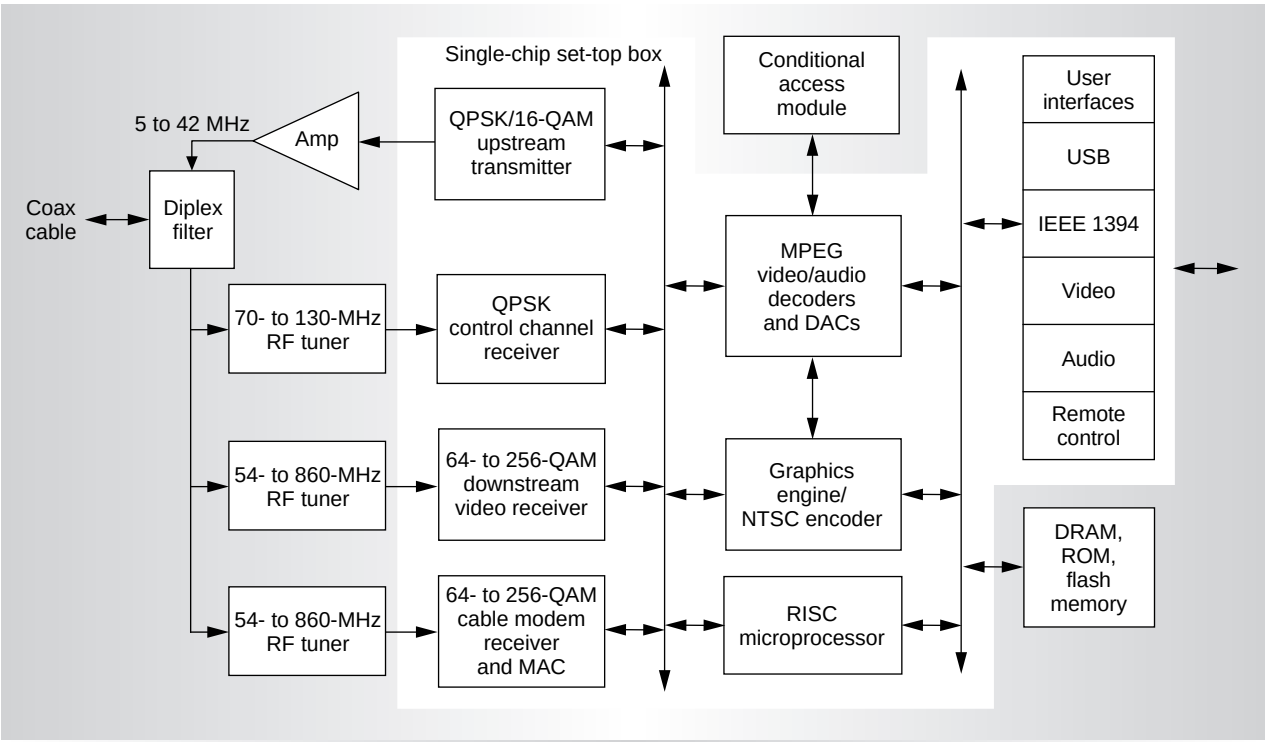

Figure 6. Next-generation, Web-enabled cable set-top box integration.

When we have something to send, we alert the MAC, which first requests some bandwidth from the head end and then transmits that burst of data. In the back end we have the microcontroller, a 10- or 100Base-T transceiver, and memory.

The MAC occupies about half of the 0.35-micron, 3.3-V, 3.5-million-transistor cable modem IC, so we have a fair amount of logic in that function. It's actually larger than the QAM receiver portion shown in Figure 5. The analog on this chip is less than 10%.

Over time, we'll see the merger of the two technologies: the digital TV set-top box and the cable modem targeted toward the computer. TVs will have the same bidirectional Internet surfing capability as the computer.

Figure 6 shows a set-top box with an extra cable modem channel. We have a third tuner with another QAM receiver for the data channel with its own MAC. The back end is upgraded a little. The same MPEG video and audio decoders will handle high-definition TV reception, and the graphics engine will be upgraded to handle the high-definition screen. And then the microprocessor must be upgraded to handle all the fancy 2D and 3D graphics on screen. This means we'll need a fancier RISC engine as well. The back end is similar to what was shown earlier.

We'll be seeing this convergence of data, video, and voice into one medium in the not too distant future. We're going to have telephony service through all the work that AT&T is doing now with TCI to add voice and video telephony to cable.

Home networking

We have a huge opportunity for networking within the home. It's a very obvious market that will clearly happen. In today's US office environment, roughly 100 million PCs are connected via about 2 million miles of RJ-45 cabling. In the home, we have about 50 million unconnected PCs, and there's about 20 million miles of RJ-11 cabling sitting inside homes. It's an obvious market to connect home PCs and then connect the other hundreds of millions of miscellaneous devices in the home such as appliances, security systems, entertainment systems, and so on. Eventually, if this technology gets inexpensive enough, everything will be connected, whether wired or wireless.

Our chip set[4] performs phone line net-

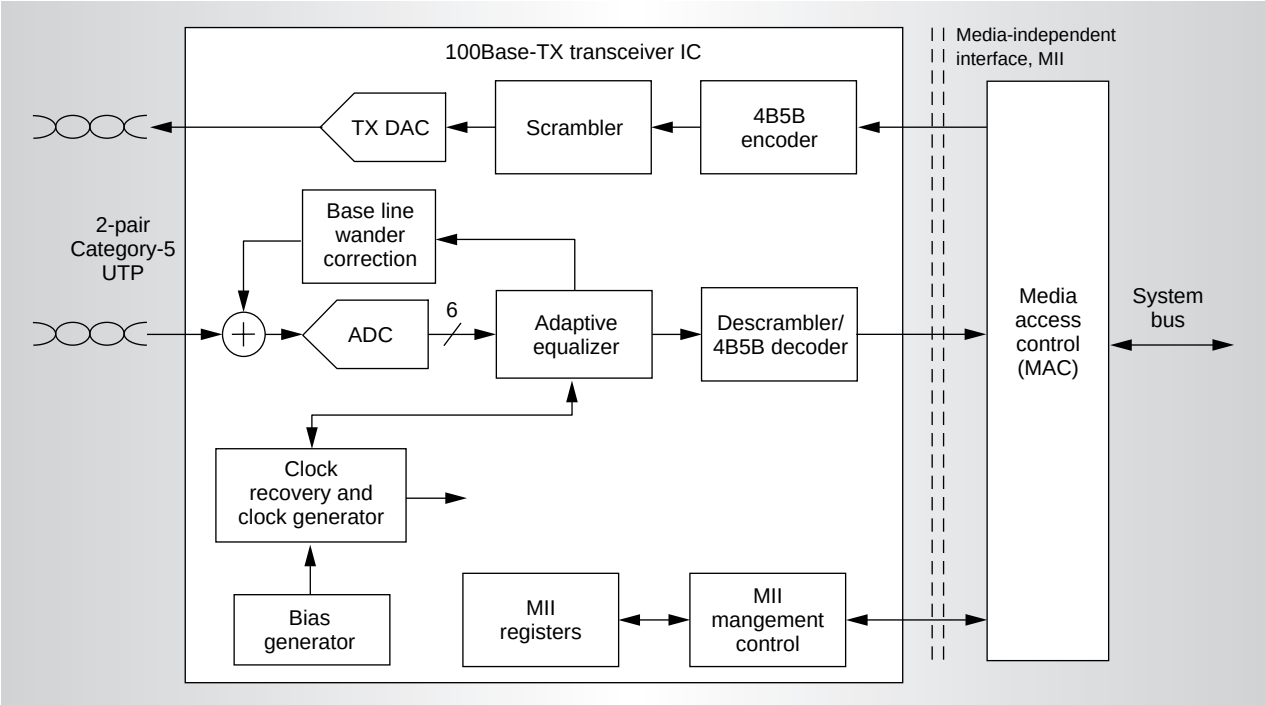

Figure 7. 100Base-TX fast Ethernet transceiver block diagram.

working at about 10 Mbps, and coexists on the same twisted-pair phone line that's already in RJ-11 wall jacks. If an ADSL connection exists in a home, it will also coexist with the phone line network, which operates above ADSL in frequency. Then it would connect to all the miscellaneous appliances in the home. It's based on QAM as well, and it has an Ethernet-like MAC protocol, the so-called CSMA/CD, or Collision Sense Multiple Access/Collision Detection. One modification is added built-in quality of service provisions that guarantee latency for the delivery of voice and video telephony service.[4]

High-speed LANs

Ethernet is the predominant LAN technology in use today, with over 85% market share. The current worldwide installed base is about 200 million nodes. The first-generation, 10Base-T, or10-Mbps Ethernet doesn't provide sufficient bandwidth to handle broadband multimedia traffic. Virtually everyone has upgraded to 100Base-T.

The typical 100-Mbps Ethernet LAN configuration has a hierarchy of connections. A server farm is connected via fiber or copper broadband pipes into the network's first layer of switches. These 100-/1,000-Mbps switches feed down one level of the hierarchy and eventually get to the end user—the so-called edge of the network with the workstations and PCs. They're fed either by local switches in the wiring closet or through nonswitch repeaters or hubs. Switches mean we have a dedicated connection from source to destination. In a repeater or hub, the incoming traffic gets repeated on all the ports; hence it's called a repeater. With a switch, the incoming traffic only goes to one port, its intended destination on the outbound network.

Figure 7 shows a block diagram of a fast Ethernet transceiver using a DSP approach. At the front end is an ADC that digitizes the 125-Mbaud signal. 100Base-T requires two unshielded, Category-5 twisted pairs, which is much higher end cabling than the phone line. One pair transmits; one receives. The receive signal is digitized, goes through an adaptive equalizer and clock recovery, and then into the MAC in the back end.

The octal transceiver has gone through many generations of technology scaling. The transceiver is about 50-50 analog and digital with eight ADCs in the front end and eight DSPs, the DACs, in the back end. The trans-

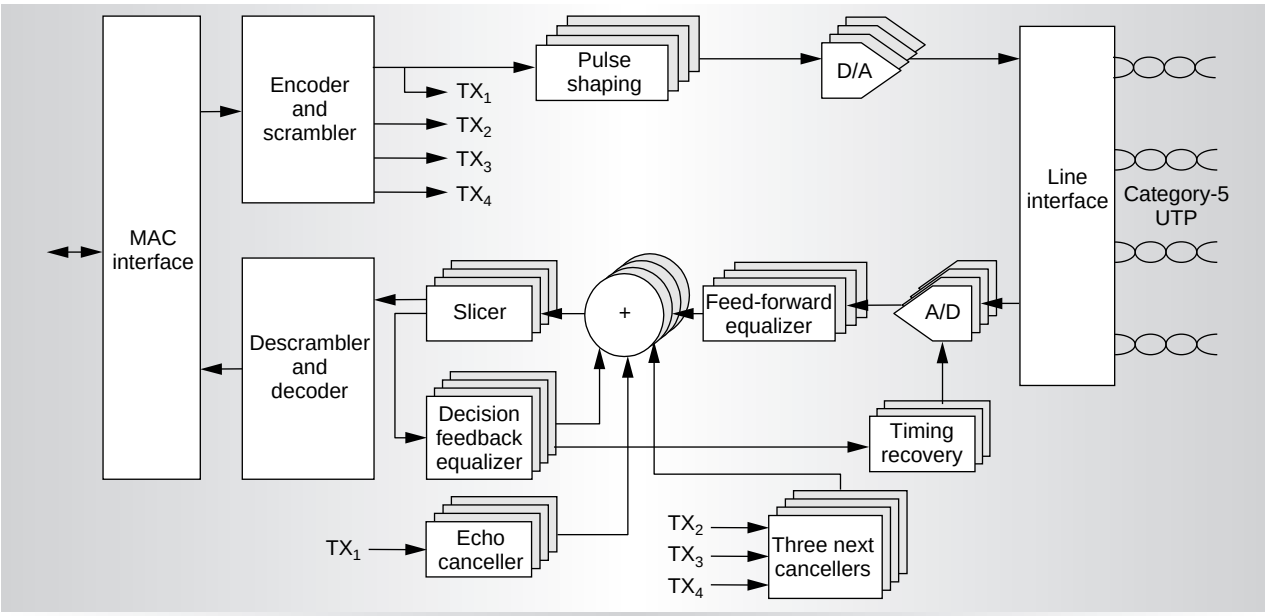

Figure 8. Block diagram of a 1000Base-T transceiver.

ceiver operates at up to 160 meters, virtually free of errors. Each 0.35-micron, 3.3-V, single-poly, quad-metal CMOS transceiver requires less than 4 mm^2 of silicon area.

Today gigabit Ethernet is principally a backbone, fiber-based technology. The IEEE has recently published a new standard called 802.3 (ab). It defines gigabit Ethernet over 100 meters of the same unshielded, Category-5, twisted-pair cabling that the current 100Base-T networks use. So we have a gigabit over the same installed infrastructure with no requirements for upgrades.

The only difference is that the network is running over all four pairs. A typical Category-5 cable actually has four twisted pairs in the bundle. With 100Base-T, only two of those twisted pairs are being used. In 1000Base-T, all four are used. Since over 80% of corporate networks today are wired with Category-5 cabling, we'll see initial applications for 1000Base-T primarily in the backbone switches. That is the higher end of the hierarchy that we saw earlier: either the server farm to the first switch or the switch-to-switch connections. Eventually, we'll see them migrate to the desktop as well. It's a natural migration over time; we'll be buying universal 10/100/1000 adaptor cards shortly. The cards will plug into PCs, automatically sense the network speed, and operate up to 1 Gbps at the desktop.

All transceivers transmit at 250 Mbps, but are full duplex and bidirectional. Hence we need echo cancellation to cancel out the transmitted signal before it hits the receiver. At 4 times 250 Mbps, it's bidirectional 1,000 Mbps, or 2-Gbps total transmission data rate.

With a very broadband signal, we have cross talk between four pairs. The signal couples in either capacitively or inductively between the four pairs, requiring that we cancel out that noise with echo cancellers and cross-talk cancellers to handle the transmission of bidirectional signals on the same cable. We end up with the very complex transceiver shown in Figure 8.

There's a lot going on in this transceiver while it runs in real-time at 125 MHz. We have four analog front ends, four ADCs running at 125 MHz—one for each twisted pair at the line interface—and four transmitters driving those four twisted pairs. Then we have four receivers with four adaptive equalizers and four decision feedback equalizers. We have four echo cancellers and 12 next cancellers because we're canceling out cross talk from the three other pairs on each of the four channels.

This is a "hot" chip without one ounce of programmability in it. That's the sacrifice we make for getting the high throughput; it's dedicated hardware with no micro code or RAM. In fact, the only memory it uses are in the local

shift registers in the filters. A state-of-the-art, 250-MHz TI 6000 DSP has two MAC engines and six ALUs for a peak of 2 billion operations/sec. Our chip achieves a throughput of 250 billion operations/sec. This makes it 125 times more powerful than the TI 6000, and it clocks at half the speed. With its 125-MHz clock rate, we have 2,000 arithmetic or multiply-accumulate operations going on in parallel in every clock cycle.

We can do this because everything is dedicated with zero programmability. All of the word lengths are optimized, so if we only need a 6-bit multiply, that's what we get. If we need an 8-bit or a 12-bit operation, every single arithmetic element is optimized, depending on its function. A huge amount of systems engineering went into defining all of that.

Transceiver IC design challenges

In designing the transceiver chips, we not only had to solve mixed-mode and baseband analog and digital problems, we also had to deal with IF and RF frequencies, as well as the more conventional random logic and memory sections. In the early 1980s building one of these systems would have taken a rack of hardware costing hundreds of thousands of dollars. The military in the 1970s and 1980s actually built these types of communication systems for satellite ground stations, spending millions of dollars on them. But it was a hardware monster because the silicon capability just didn't exist. Now it costs us tens of dollars for these devices—an amazing level of cost reduction achieved through integration.

From a systems perspective, we run into the problem of taking something very complex and mapping it into a chip that costs $10, we need a very tight coupling—or what I call a vertically integrated engineer. It's not enough for an engineer or designer to have an isolated specialty, which is how we used to design chips. We had systems engineers, DSP guys, logic designers, and layout guys in different buildings, and they all threw specs over the wall. What came back was a monster. If we want to end up with an efficient design, all the engineers have to be vertically integrated and have to really understand the implications of what they're doing on every design aspect. They must understand the system. Not everyone has to be an expert in every aspect, just have a deep understanding and appreciation of what goes on at the system, DSP, and the lowest IC level, both analog and digital.

These people are hard to find, as we well know, and we're trying to do our best at the universities to train them. The most important aspect is that the people defining the block diagrams at the highest level absolutely must understand the IC design implications of their block diagrams, and this has not been the traditional systems engineer. Traditional systems engineers are mathematicians. They draw block diagrams, write equations, and don't understand the issues when it gets down to hardware. That will not fly in the future.

Mixed-mode CMOS design issues

It is critically important to know where to draw the analog/digital partitioning line. Then when we get down to the mixed-mode design at the circuit level—as we saw in all the chips shown earlier—we see that they're dominated by high-speed digital. The analog part of the chip is typically a small percentage. We don't want to dedicate a fancy double-poly process or special capacitor options for the analog sections that burden the 80% or 90% of the digital chip with extra cost. So we force the analog people to use single poly.

Furthermore, we cannot afford to throw away a functional digital chip because the ADC had some excess nonlinearity. We need a very robust analog circuit so that overall parametric yield and all process variations are as robust as the digital. The analog block must have a yield limited only by defect density. We can have very little incremental yield loss due to parametric variations of a design. Clearly, analog designers are heavily penalized in mixed-mode design, while the digital designers are largely unaffected, at the circuit level at least. However, digital designers have different issues at the system level.

Scaling issues

The arguments for scaling digital are obvious and compelling: a density and a factor-of-two power improvement in every process generation. Obviously, we want to move as fast as we can. The analog circuits absolutely don't scale as we reduce supply voltage. Since we have such poor scaling in the analog section, the mixed-mode chips generally are

Definitions

ADSL	asymmetric digital subscriber line
FDMA	frequency division multiple access
HDSL	high-speed digital subscriber line
ISP	Internet service provider
LMDS	local multipoint distribution service
MAC	media access controller
MMDS	multichannel multipoint distribution service
NTSC	National Television System Committee
POTS	plain old telephone system
QAM	quadrature amplitude modulation
QPSK	quadrature phase-shift keying
TCI	Telecommunications Inc.
TDMA	time division multiple access
VDSL	very high-speed digital subscriber line

about a generation behind the latest process technology. We use 0.35-micron chips for volume production—the current design is 0.25 micron. We're starting to look at 0.18, so we're typically one generation behind in the process technology, and it really is because of the analog process. We need a very stable, very well characterized, well-modeled analog process.

Consider a DAC going from 0.50 micron at 5 volts to 0.35 at 3.3 volts, to 0.25 micron at 2.5 volts. We get reasonable scaling. It's not the power of two we get with digital, but it's like square root or two. The reason we're getting that scaling is because a lot of the circuitry is digital. A DAC is very much a digital circuit in many respects.

On the other hand, we get no benefits from scaling an ADC design. We have 0.50-micron at 5 V, 0.35 at 3.3 V, and 0.25 at 2.5 V—no area gain, no scaling whatsoever. It's principally because we cannot reduce the currents. We need a certain bandwidth out of the circuit, and bandwidth translates to a certain current that has to drive that transistor. If we have to keep the current more or less constant, we can't reduce the transistor's size; it stays big. In some cases it has to get even bigger, because we have less headroom in these devices, and as we lower the headroom we have to pay other penalties. In some cases it gets worse. This is why we don't get the benefit of scaling with mixed-mode design.

The digital people don't get off free either. They have to deal with complexity. On the ASIC side it's very challenging because we're dealing with system-level and algorithmic complexities. A 10-million-transistor ASIC is relatively easy to fabricate—it's a nothing chip these days—but it's exceedingly difficult to design. We finally reached a point where the back-end physical design and verification tasks take far longer than the front-end system architecture and logic design. And this traditionally is all the glory work in an IC design, while the poor back-end designers get all the grunge work. Unfortunately, that is the most important part of a design today.

It gets magnified even more when we look at 100-million-transistor ASICs. When I'm doing these logic transistor counts, I'm not counting RAMs, which of course blow up the transistor counts enormously. When I say 100 million transistors, I mean a 25-million-gate design of logic. That will be feasible to fabricate very shortly. But unless there are fundamental advances in design methodology and CAD tools to manage that overwhelming complexity, we won't be able to design chips that big. Our tools are breaking now at the 10-million-transistor level. They're going to be absolutely worthless when we start trying to get 25 to 100 million gates. So a lot of work has to go on in the design methodology and CAD tool world to support these very complex devices.

The ubiquitous availability of broadband interactive services will truly revolutionize society. It's going to have a dramatic impact on everyone, not just the technology literate. The availability of very low cost, very highly integrated devices to connect everything together is the driving force. It is really the key enabler allowing everything to happen. It's going to be very exciting for the IC designer and very challenging in every aspect of IC design. We're in one of the best professions we can possibly be in with endless opportunities for creativity. MICRO

References

1. A. Kwentus et al., "A Single-Chip Universal Digital Satellite Receiver With 480-MHz IF Input," *IEEE J. Solid-State Circuits,* Nov. 1999, Vol. 34, No. 11, p. 1634.
2. R. Joshi et al., "A 52 Mb/s Universal DSL Transceiver IC," *Proc. 1999 IEEE Int'l Solid-State Circuits Conf.,* IEEE Press, Piscataway, N.J., 1999.
3. L.J. D'Luna et al., "A Single-Chip Universal Cable Set-Top Box/Modem Transceiver," *IEEE J. Solid-State Circuits,* Nov. 1999, Vol. 34, No. 11, p. 1647.
4. E.H. Frank and J.T. Holloway, "Connecting the Home With a Phone Line Chip Set, *IEEE Micro,* this issue.