

Analog Telephony... How We Got Here, and What's Happening To It Now?

By Mike Sandman • mike@sandman.com

In order to understand telephony today and make it work right it's handy to know how telephony has evolved. The key component of all this is our ears and mouth.


What is analog telephony? It's just a pair of wires from one phone to another that transmits vibrations from our mouth to someone else's ear (and vice versa).

Since it's unlikely that we'll evolve to have a USB connector for a digital input to our heads anytime soon, every kind of telephone will have an analog speaker and microphone of some kind for the foreseeable future.

Without that USB connection in our head, email and text will also require us to have eyes (analog light transmission from a screen or piece of paper).

In the old days electrical signals from phone to phone were switched between the two phones by Central Offices (by human operators at cord boards, and later by mechanical or electronic Central Offices). That's just connecting (temporarily splicing) two wires to a different pair of two wires, maybe many times over a long distance, to get from one phone to another.

Phones weren't "electronic" until fairly recently. The signal from the microphone was sent on the two wires to the CO from the handset using a transformer in the telephone's internal network, as can be seen in this picture of a standard 500 or 2500 set network used since the 1940s ***with no active electronics***:



Early networks used screws instead of push-in spade lug sockets.

The telephone's internal network has a transformer that changes the four wires going to the handset (transmit and receive) to two wires (going to the phone company CO).

Rotary dials contained no electronics, but touch tone dials did contain electronics to make the tones (using transistors developed by Bell Labs), powered by the DC current on a standard analog telephone line (which was originally used to power the carbon microphone).

Really old phones each contained "local" batteries to power the carbon microphone. "Common Battery" phones became the norm where the electricity to power the transmitter (and dial) were supplied from the Central Office end, instead of the subscriber end.

The amount of talk battery (DC current and voltage) on a phone line was originally what was needed to make the carbon transmitter in the handset work (over 20ma DC). The AC ring voltage was chosen at about 90VAC at 20 to 30 cycles (instead of 60 cycles per second). The 20 cycle frequency was slow enough to let the clapper in a mechanical bell swing back and forth between the gongs slow enough to make a nice ringing sound.

The engineers designing later phone line powered stuff with electronics just assumed that every line would have that minimum amount of 20ma DC loop current. The phone lines sometimes didn't. And engineers still don't consider that today the loop current can be as high as 110ma DC, which often burns up the electronics or makes it do strange stuff (so you're a test pilot if you have high loop current because the engineers never tested it at that level).

One of the problems the original long-lines engineers dealt with was that phone lines were only two wires, with transmit and receive ***on the same pair of wires***. That means that as they tried to amplify the line to make it easier to hear they created feedback like you'd get from a mic on a PA system - usually called "singing" or "ringing" during parts of the conversation (squealing).

Well before long lines were digital, engineers developed ways to separate transmit from receive using transformers, and amplified the voice with tube type amplifiers. Amplifying the voice separately in each direction on a four wire line works great. Most of the time four wire lines were used between Central Offices and to PBXs, and two wire out to the subscriber with standard phones. The analog signals on analog pair gain systems used tube type equipment to put many calls on a single coax cable, and AT&T long-lines strung multi-core coax cables all over the country:

AT&T also used lots of microwave radio towers for the analog signals that can still be seen here and there. They're no longer used by AT&T long-lines, but cellular stuff is on a lot of them. The microwave also carried analog TV signals for networks before satellites.



Microwave Horns on an AT&T tower. 107 were spaced across the country about 30 miles apart, receiving and re-transmitting radio signals to the next tower.

When you make a 3-way conference call today from your home or office using a POTS line (2 wires) the amplification needed to allow each person to hear the other callers clearly is added on the four wire side of the CO, not to the 2 wire side going to your home or office. So the conference call sounds perfect.

In the early 1960s T1 pair-gain became popular (made possible by transistors), where two pairs of wires between two points could carry 24 conversations. The analog

conversations were turned into TDM (Time-Division Multiplexing) electronically by equipment at each end of the two pairs of wires, with repeaters along the way.

With TDM each conversation uses 64K of bandwidth, with the entire T1 handling 24 of those conversations. That's on only two pair of wires, that used to carry just two analog conversations. 1544Kbit/s of bandwidth total.

Once engineers developed ways to digitize the analog conversation, the pairs of wires were no longer the limiting factor. Most of the long line communications was digitized, put on copper pairs as well as coax and microwave (including submarine cables between continents).

They sent the data for phone calls through satellites for a while but it took too long for the data to get up and down from the satellites. The delay (latency) make it too hard to carry on a conversation, and even the best echo cancellers of the day had a hard time sometimes.

Then phone calls were finally put on fiber optic strands.

The real breakthrough was when engineers figured out how to split up the rainbow of colors in a single fiber optic strand to send many times the amount of data. Instead of dividing the two pairs of copper wires into 24 time slots, a single fiber is divided into many colors, each of which can handle a lot of data (all using sophisticated electronics).

Our current telephone network is still like the original circuit-switched network where the two wires for one phone are connected to two wires on the other phone. But what happens *in-between* the two Central Offices is invisible to us... where conversations are usually converted to data.

Luckily for us the phone network has always had very "lo-fi" voice quality (about 300 to 3500 hertz). Unusable for music, and not real good for the tones sent from end to end by a modem.

A lot of "lo-fi" conversations can be stuck on one data pipe and they arrive at the other end in a reasonable time (latency), and in the right order. It takes 64K of bandwidth to send one "lo-fi" conversation by TDM pair gain.

1 Netflix HD Movie = 78,000 VoIP Calls...

A single HD Netflix movie uses maybe 5 million bits of bandwidth, compared to 64 thousand bits for a "lo-fi" telephone conversation. If I'm doing the math right that pipe with the HD Netflix movie can carry about 78,000 telephone conversations at the same time, or the *one* movie. VoIP calls have a lot of competition out there!

The competition will probably become overwhelming if super high quality 4K movies start filling the pipes. At 15 million bits per second sharing the same pipes with movies becomes absurd:

1 Netflix 4K Movie = 234,000 VoIP Calls

It turns out that the process of coding analog voice into data for TDM pair gain, and decoding it back to analog is very fast. And it doesn't take long to send those bits of data from one point to the other on dedicated copper or fiber. With TDM, what goes in first on one end comes out on the other end first. Just what we need for voice calls.

Today phone companies are getting rid of TDM as much as possible, replacing the infrastructure between the two central offices with equipment that "packetizes" the data (Internet Protocol, or IP). It takes small chunks of our voice and turns it into a packet of digital data, and sends lots of those packets on to the destination Central Office.

The problem with packetizing our voice is that the packets are often mixed with other packets of non-voice data. Netflix movies, Pandora radio and even porn on the public Internet.

For now, the "packets" sent between the phone company Central Offices are using a "private" IP network owned by the phone companies (mostly using fiber rather than copper). Since the packets don't have to compete with Netflix or Porn they arrive pretty quickly and in the correct order. That's likely to change soon since AT&T wants to close all the telephone Central Offices and just put your voice calls on the public Internet, mixed with all the other traffic.

If VoIP packets were travelling from Point A to Point B with other voice packets, without other types of data, it would work pretty well.

But those packets are competing for bandwidth with movies and email so they are slowed down somewhat. Most of the current Internet has QOS (Quality of Service) implemented to give priority to voice traffic, which is unlike almost any other traffic in that it's "real time" and can't be buffered. VoIP packets **must** arrive on-time and in the correct order.

When packets arrive out of order (too late) they're thrown out by VoIP equipment, which is when we hear burbles and cut-outs in the conversation. We might hear all the packets in the right order on some calls, but the latency (delay) can make it frustrating to carry on a conversation.

Sometimes the conversation is converted from analog to digital (TDM or IP) and back several times on a call to get from Point A to Point B. Bits are lost in the process, which is OK. We usually can't hear that there's lost data until it "piles up," with too much of the original conversation being lost by the various conversions.

It's just a fact of life that packetizing digits for IP takes longer than converting the data for TDM. If there isn't much latency (delay), we might only notice it if we put a handset from two phones up to our ears and talk to ourselves. For conversations between two people it's fine, ***except when there is echo.***

Echo essentially ***never*** occurs on a VoIP to VoIP phone call where it's never changed to analog between the two parties, and both users are using a regular handset. VoIP always has separate transmit and receive (to and from your ear and mouth).

Echo problems occur any time a four wire VoIP call (with separate transmit and receive) is put onto legacy telephone equipment - which has two wires.

A four wire call with separate transmit and receive is changed into a two wire analog call by a transformer (a special kind called a **hybrid transformer**). Basically there are two wires on one side of the transformer, and four on the other side. Through induction the transformer changes the four wire analog signal from VoIP equipment to a two wire analog signal used by POTS type telephone equipment (a phone, modem, trunk card in a phone system, analog station port in a phone system, ATA, etc.). The hybrid transformers carry the voice in both directions during the call - converting two wire to four, and four wire to two.

The echo problem comes from two facts that aren't easily overcome:

1. The hybrid transformer is never 100% efficient. Some of the transmit bleeds over to the receive, called "sidetone".
2. It takes a while for the electronics to packetize (or un-packetize) the voice, done by "codecs."

Hybrid transformers have always been used in phone equipment. That's why you can hear a little bit of your own voice in your ear while you're talking on a standard POTS phone. It's nice because it gives us a warm feeling that you haven't been disconnected. Most people hate phones without sidetone because they're never sure they haven't been cut-off when talking.

With standard analog telephony there's no delay when we're talking. The hybrid is inefficient and we hear the sidetone, but since there's no delay it's not heard as echo. Just a little bit of our own voice coming back to us. There's so little delay involved with digitizing voice by TDM that we generally don't hear echo, and the echo cancellers in the telephone network don't have to work very hard.

With the delay in VoIP caused by packetizing the analog voice, we hear our own voice or the other party's voice coming back after a short delay. Echo.

It's unlikely that we'll be able to afford to replace every piece of analog equipment in the near future, so we have to make it work with packetized voice. We're stuck with work-arounds.

The #1 work around is echo cancellers at all the points where the packetized voice jumps from four wire to two wire. Using DSPs (Digital Signal Processors) most of the echo can be eliminated if it's not too loud, or isn't delayed too much.

The echo cancellers are located at the phone company, in our VoIP phone systems with two wire trunk and station cards, in VoIP ATAs, and in our VoIP phones, etc.

One specific echo problem comes when we're using a speakerphone on a VoIP phone, analog speakerphone on a VoIP phone system, or even a cell phone's speakerphone. Some of that analog audio from the caller on the speaker is picked up by the mic in the speakerphone. It's then packetized again which has that inherent delay, and is heard as echo.

Getting rid of echo on a speakerphone is fairly easy with an echo canceller because the engineers are working with an environment that's always the same. The speakerphone mic is always the same distance from the speaker, and the volume on the speaker can never be louder than what they allowed the user to raise it to.

For phone calls on phone systems the echo canceller has a lot more to do. It's made to adapt to the conditions as quickly as possible, which is why you sometimes hear echo for the first few seconds of a phone call. It's learning.

Some VoIP phone systems and VoIP equipment allows a technician to adjust the settings on the echo canceller. When the settings are mistakenly set to be too aggressive you end up hearing a crackling sound on the call as the echo canceller starts removing some of the real voice as well as the echo (very annoying).

The #2 way to eliminate echo is to reduce the volume of the call. That doesn't work so well since it's hard to understand what the other person is saying if we can't hear what they're saying.

The #3 way to eliminate echo is to make the hybrid transformer in the phone equipment more efficient either through better design, adding some electronics, or adjusting the impedance of the phone line to match the hybrid transformer in the equipment (which is what our [Echo Stopper™](#) does).

So if we eliminate all two wire telephone equipment and don't use speakerphones we've made problems with echo caused by packetization go away.

But it's still extremely unlikely we'll ever have **all** of the packets for phone calls go over a private network dedicated to VoIP, instead of over the public Internet, so we'll always be

dealing with voice packets sharing the pipes with Netflix and porn causing some of the packets to be lost.

Greed has a lot to do with the poor quality of VoIP calls. Putting 24 phone calls in a 1544Kbit/s pipe seemed like a waste of bandwidth, so codec and compression algorithms were designed to compress a pretty good sounding 64K phone call down to 8K. That's usually done to either fit more (bad sounding) calls on a given pipe, or share more Internet (downloads, email, etc.) with the phone calls on that single pipe.

You can still hear OK on a phone call compressed down to 8K (with the g.729 codec) as opposed to 64K (the g.711 codec), but modem tones and faxes are clobbered to death. Lose some of those g.711 packets along the way to delay and routing and you ain't going to get reliable fax or modem traffic even with the best error correction (but faxes and modem calls usually work after a couple of tries). Even DTMF touch tones have a hard time being transmitted end to end by VoIP because it's optimized for voice, not steady tones. Oh, and most alarm systems use DTMF or modems to communicate.

Since touch tones are important to Voice Mail and IVRs (like telephone banking), VoIP has a work-around that lets you turn on a mode where when the near end VoIP device hears the analog DTMF tone but doesn't send the audio to the other end. The touch tone is stripped out of the audio being sent to the far end, and a packet of data telling the far end VoIP device what digit to reproduce on the far end is sent (the setting is usually called RFC 2833). Of course there are occasional problems where the whole DTMF digit audio *doesn't* get stripped out and the far end equipment hears what's left of the "stripped out" digit as well as the one created by the far end equipment, creating a double digit (and a headache to troubleshoot without a DTMF decoder like our [Tone Master](#)).

Sometimes it just makes sense to keep some of those POTS line from the phone company, at least until they won't give them to us anymore. That's already happening in brand new neighborhoods and business parks all over the country where the local phone company will no longer run copper for any reason. The media converters that change the fiber to copper have all kinds of strange problems working with real phone equipment - and you never know what the problems will be until you try it.

What can we do to improve VoIP calls?

The biggest thing we can do to improve VoIP phone calls is to eliminate sharing the local Ethernet network, *and* sharing the local pipe to the Internet.

When we share the local network with computers we're limiting the bandwidth available to VoIP at the same time someone is uploading or downloading large files. Worse, if the router is mis-configured the voice packets may not have QOS - and even if QOS is set correctly there's only so much the router can do with the size of the data and the pipe it's attached to.

Diagnosing VoIP isn't as easy as connecting a butt-set to a line or station port. Sometimes you need to use a packet sniffer to see why network stuff doesn't work right with computers (and is effecting VoIP calls), or with VoIP why there is delay or a conversation sounds bad.

Using a full duplex Ethernet tap (like our [Ethershark™](#)) and a free packet sniffing program like Wireshark is really the only way to know without guessing (guessing is very frustrating for all parties involved!). With Wireshark you can even reassemble the packets you capture from a VoIP call (by IP address) and play it back.

You can download the free Wireshark sniffer program and a VoIP Screen Phone to your laptop without any other stuff to see how it works on your network.

Terrell Boyer is a **Wireshark** VoIP expert and in his YouTube video "**The Ultimate Wireshark Tutorial**" he gives you all the information you need to start using Wireshark. This is the best 50 minutes you're ever going to spend if you have to fix VoIP problems:

<http://www.youtube.com/watch?v=Lu05owzpSb8>

You can do a lot of troubleshooting using an Ethernet switch that allows you to mirror a port to your laptop with Wireshark. What you don't get on the mirrored port is errors caused by a bad piece of equipment (NIC card, or whatever). Those errors aren't mirrored, and Murphy's Law says *that* will be the problem!

The pipe to the Internet is even more important than the local network since it's pretty small by comparison. Where the bandwidth on your network may be 100 million bits per second, the pipe from the router to the Internet is probably closer to 1.5 million, or maybe 6 million bits per second.

The controlling factor for an Internet connection is the **upload speed** if the pipe isn't a symmetrical line (same up and down speeds). ADSL is a really bad choice for VoIP since it's asymmetrical, with upload bandwidth topping out at maybe 750 thousand bits and download speeds at maybe 6 million bits per second. Even if ADSL2 is available it usually doesn't offer high upload speeds.

If each VoIP phone call takes maybe 100 thousand bits per second of bandwidth (a reasonable number to figure on average), you can have maybe 7 simultaneous conversations on an ADSL line dedicated to VoIP, with 750 up. If you're sharing that ADSL line with email, music, files and movies, you're down to maybe a couple of VoIP calls before things start sounding pretty bad.

SDSL or T1s are a better choice for voice, but you need a good Internet provider (not that easy to find!).

Generally speaking Internet pipes from the cable company, U-verse or FIOS aren't a good choice to use with VoIP because all of the subscribers in a neighborhood are

sharing the same pipe to the Internet from the "terminal" in the neighborhood. If a lot of subscribers are downloading movies at the same time as your VoIP voice calls your calls are going to sound pretty bad at times (but fine at other times). You have absolutely no control over what the other subscribers are doing, so you have absolutely no control over the quality of the sound of the business phone calls with these types of Internet connections.

Business class DSL or T1s are absolutely needed for a successful implementation of VoIP (assuming your Internet provider has a big enough pipe to the Internet to service all the subscribers from *that* Central Office). **And** a dedicated pipe for VoIP makes a lot more sense than sharing that pipe with the rest of the Internet traffic from the office.

Unfortunately it's difficult to impossible to convince a business owner to spend more money when they were told they would "save money" by going with VoIP, using their existing network and Internet bandwidth - and getting rid of those expensive evil phone company POTS and T1 lines.

On the good side we're all being conditioned to accept pretty horrible sounding phone calls compared to the old TDM networks / phone systems. In 1970 there would have been a revolt against the phone company if calls sounded like they do today. Cell phones began that change. We were just happy to not be tied to our landline phone all day long.

[Signup to receive my monthly blog by email.](#) I assure you I'll have a new blog post if bio-medical engineers find a way to get a USB connector into our heads so we don't need those inefficient analog mouths or ears anymore.