


G.729 VOICE CODEC SERIES

Adaptive Digital Technologies, Inc.

G.729 AND ALL ITS FLAVORS

G.729 is an umbrella of vocoder standards. The G.729 vocoders perform voice compression at bit rates that vary between 6.4 and 12.4 kbps. The figure below shows an example of the G.729 vocoder connected to a digital communication channel. The input speech is fed into the G.729 encoder as a stream of 16-bit PCM samples, sampled at a rate of 8000 samples/second. The G.729 encoder compresses the data into the Encode Stream. The encoder also outputs the DTX status, which is discussed later in this data sheet. The digital channel carries the data stream and DTX status to the decoder, which regenerates a representation of the original speech, and outputs it as the output speech – again as 16-bit PCM at a sampling rate of 8000 samples/second. Since G.729 uses lossy compression, the output speech is not identical to the input speech.



The decoder is also fed a frame erase flag, which is an indication that the decode stream has temporarily been corrupted. The decoder is able to “smooth over” the output, doing its best to conceal the loss of data and minimize the loss in voice quality. This process is known as packet loss concealment (PLC). It works surprisingly well even under high packet loss rates.

G.729 FLAVORS (ANNEXES AND APPENDICES)

G.729 has enough annexes to confuse anybody. The following table sorts things out.

Standard / Annex	Description	Bit Rate (kbps)	DTX Annex	MOS
G.729	The original G.729 vocoder (G.729 main body)	8	Annex B	3.9
G.729 / A	Lower Complexity G.729	8	Annex B	3.7
G.729 / C	G.729 floating point	8		3.7
G.729 / C+	Floating Point, with DTX and multiple bit rates	6.4/8/12.4	Annex B	6.4 kbps: 3.9 8 kbps: 3.9 12.4 kbps: 4.1
G.729 / D	Multiple bit rates	6.4 / 8	Annex F	3.6
G.729 / E	Higher Quality Vocoder	12.4	Annex G	4.2

The first column in the table lists the ITU standard and applicable DTX Annex. The last column in the table lists the MOS (Mean Opinion Score) for the various flavors of G.729. MOS is a subjective rating given to vocoders. MOS runs a scale from 1 to 5, with 1 being the worst and 5 being the best. A score of 4 and above is considered to be “toll quality” speech.

You may take note of the column labeled DTX annex. Although the G.729 family of vocoders does an excellent job of compressing speech, it does even better when using Discontinuous Transmission or DTX. DTX is the process by which the encoder determines on a frame-by-frame basis whether there is voice activity or not. If there is no voice activity, the encoder produces either a reduced bit rate (DTX) frame of compressed data representing the background noise characteristics, or it produces no data at all. If the decoder receives a full voice packet, it regenerates speech. If it receives a DTX frame, it generates noise that is representative of the background noise. If it receives no frame at all, it continues to generate background noise according to the specifications from the most recent DTX frame. (If the decoder was expecting a speech frame but received nothing, it will perform its packet loss concealment instead.)

That all said, the DTX annex column in the table above lists the annex letter to the G.729 specification that describes how DTX is to be performed in conjunction with the associated base annex. Let's take the first row in the table, for example – G.729 (main body). G.729 compresses speech, with a compressed bit rate of 8 kbps. If DTX is to be used, it will be specified by G.729 Annex B. When G.729 is combined with its Annex B DTX feature, the combined vocoder is often referred to as G.729B. Similarly, when G.729 D is combined with its DTX feature, the combined vocoder is often referred to as G.729D/F.

A standard G.729 packet contains 80 bits of compressed data representing a 10 millisecond frame of speech. The 8 kbps rate comes from: $8000 \text{ bits/second} = 80 \text{ bits} / 0.01 \text{ seconds}$. A G.729 Annex B DTX packet contains 16 bits, resulting in a bit rate of 1600 bps. Of course, a real conversation will not operate at 1600 bps unless there is only background noise present. The actual bit rate when using DTX lies somewhere between 1600 bps and 8000 bps. Typically, DTX can save approximately half the channel bandwidth on average. This comes about because telephone calls tend to be half-duplex. Only one person speaks at a time. You could argue therefore that, on average, each person speaks less than half the time (unless they're involved in a shouting match, or they are rude and like to interrupt each other a lot.)

The other DTX annexes (F and G), offer similar bandwidth savings when combined with their respective base annexes (D and E).

There are a few more annexes of interest. Annex H is a reference implementation of switching between annexes D and E. Annex I is a reference implementation that integrates the G.729 main body with annexes B, D, and E. Annex J is G.729's answer to audio coding. G.729J increases the audio bandwidth from 3.3 kHz to 7 kHz. Annex J is a variable bit rate codec with bit rates between 8 and 32 kbps. Annex J has been published as a separate standard – G.729.1.

To further complicate things, there are three appendices to G.729. What's the difference between an appendix and an annex? Who knows? One difference, nobody has ever had an annexectomy procedure done to remove their annex. By the way, if you plan to download the many flavors of the G.729 standard, you should be aware that the ITU has rolled them up back into the umbrella and listed the individual pieces as "superseded" by the umbrella standard. If you plan to download G.729, don't waste your time on all the original documents. Get the whole shebang at once. (G.729 Annex J is the exception – you'll have to download G.729.1 separately, and yes, it has annexes all its own.)

Let's get back to the G.729 appendices. Appendix 1 discusses the external reset performance for G.729 codecs in systems that use an external DTX.

Appendices 2 and 3 are workarounds that have been devised to eradicate a deficiency that has been identified in the Annex B voice activity detector.

Are you completely confused yet? We hope not, but we understand if you are. Perhaps a bit of editorializing will help narrow down your field of interest. In the VoIP space, which is our primary interest, G.729 (main body) has been deemed by most designers to be too complex. In nuts and bolts terms, that means that it requires lots of DSP horsepower and hence reduces the number of channels of the vocoder that can be run on a given DSP. But people do like the voice quality / bit rate combination afforded by G.729 (main body). As a result, many users opt for G.729A, which yields almost the same voice quality as G.729 at the same bit rate of 8 kbps, but at around half the complexity. Many users also opt for the DTX, making G.729AB the most popular flavor. It's the chocolate ice cream of the G.729 family.



G.729D is gaining some acceptance in the defense communication arena, making it the next most popular flavor of G.729 after G.729AB.

After reading this editorial comment and subsequently reading below where we state that our G.729 software product availability reflects the flavors we deemed to be most popular, you may conclude that our editorial comments are self-serving. While that may have an ounce of truth, the chicken and egg didn't come in that order. The market demand has driven us to develop and offer the more popular flavors of G.729. It follows naturally that our offerings align with the most popular flavors of G.729.

To be perfectly fair, G.729E, a flavor that we do not currently offer, has gained a bit of traction. G.729E requires more bandwidth (12.4 kbps) than G.729A (8 kbps) or G.729D (6.4 kbps), but G.729E offers somewhat better voice quality than G.729A (4.2 vs. 3.7) and significantly better quality when passing music-on-hold. G.729E by no means compares with audio or music codecs. After all, G.729E starts with a sampling rate of 8 kHz and is therefore restricted to the telecom bandwidth of 300-3300 Hz, which is nowhere what is necessary for high fidelity (or medium fidelity) music. But, G.729/A and G.729D are absolutely lousy when it comes to music-on-hold.

If you have any expectations of passing music-on-hold through a low bit rate vocoder with any characteristics that do justice to the word "fidelity", you should take a step back and take a philosophical look at what you are trying to do. Have you ever heard the saying "to a child with a hammer, everything looks like a nail"?

Low bit rate vocoders look at their input signals and try to model them using a human vocal tract model. That model limits the tools available to the vocoder.

So if the input signal is speech, the low bit rate vocoder does a reasonably good job at modeling and therefore compressing the speech without losing too much fidelity. But how many people do you know who can coax their vocal tract to mimic music. Sure, you may have a uniquely talented friend (who has too much time on his or her hands) who can sound like a particular instrument. But we're not talking just about a single instrument, but multiple instruments plus singers all at the same time. What vocal tract model could do that? The low-bit-rate vocoder just doesn't have the right tools in its toolbox to do the job. By increasing the bit rate to 12.4 kbps and by adding tools to its toolbox, G.729E performs better than its lower bit rate counterparts while maintaining its voice tools. But it is no substitute to audio and music codecs.

Another aspect to G.729 that is of interest is its ability (or lack thereof) to pass DTMF signals and modem signals reliably. In the case of DTMF signals, G.729 can pass the tones to some extent, and under the right conditions a DTMF detector can detect the resulting synthesized tones. But "under the right conditions" is usually not good enough for telecom systems. Telecom systems are expected to be able to pass DTMF tones that can be detected reliably under a wide range of specified conditions, some of which are not so good. In order to achieve reliable DTMF handling when using a low bit rate vocoder such as G.729, it is necessary to employ a tone relay function.

In the case of modem signals, it's almost not even worth discussing passing modem signals through G.729 other than to make the following recommendation – don't try it.

Data subject to change, for the most up to date information visit our web site <http://www.adaptivedigital.com>

All trademarks are property of their respective owners.

© 2007 ADAPTIVE DIGITAL TECHNOLOGIES, INC. all rights reserved

PATENT ISSUES

You should be aware that the G.729 series of vocoders contains patented material, and the patent-holders do not subscribe to the religion that preaches that all software should be free. (We are not of that mindset either, but there are licensing fees and there are **LICENSING FEES**, but don't get us started.) In any case, Adaptive Digital has worked out an arrangement with the patent holders whereby Adaptive Digital can provide patent indemnification to its licensees (customers), saving lots of money that is better spent on product design than on lawyers' salaries. Can I hear an amen? (Apologies to any lawyers who might read this, but we have to ask - why are you reading this? Apologies also to the free software promoters, but again we have to ask - why are you reading this data sheet that describes a commercial piece of software? Don't take these comments too seriously. We do understand the merits of the public software licenses.)



COMPARISON SHOPPING

G.729 (with the exception of annex C) is a bit-exact specification. That means that every compliant implementation will offer the exact same output for a given input. Stated more simply, everybody's implementation sounds the same. The features that differentiate implementations from various vendors include processor utilization (MIPS and Memory), ease of use, support, availability on particular platforms, and, of course, price. If you are looking for an integrated turnkey solution or one-stop shopping for multiple algorithms, be sure to find a vendor that not only offers algorithms that you need, but also one that offers high quality algorithms. Not all algorithms are specified in a bit-exact way. This means that there you will find products that have better and worse voice quality. Examples include echo cancellation, tone detection, conferencing, and noise reduction.

PRODUCT FEATURES

- Functions are C-callable.
- Multi-channel capable.
- The encoder and decoder meet all ITU G.729 compliance and interoperability requirements.
- Can be integrated with echo cancellers.
- Capable of in-band synchronization.
- Available as part of Adaptive Digital's G.PAK turnkey DSP software packages



Strategic member
of the TI Third Party
Network

PRODUCT AVAILABILITY

C54X: G.729, G.729A, G.729B, G.729AB
 C55X: G.729, G.729A, G.729AB, G.729D
 C62X/C67X: G.729A, G.729AB
 C64X: G.729, G.729A, G.729AB, G.729D

PRODUCT SPECIFICATIONS

Bit Rate(s):

G.729, G.729A, G.729B, G.729AB: 8 kbps
 G.729D: 8 / 6.4 kbps

Delay: 5 milliseconds algorithmic delay, 10 milliseconds framing delay.

C54x

All Memory usage is given in units of 16-bit word.

Product	Function	MIPS (Peak)	Common Program Memory	Program Memory	Data Memory	Common External Memory	Per-Channel Data Memory
ADT G.729	Encode	18.7	2086	6191	194	3100	808
	Decode	3.0	2086	2908	113	3100	670
ADT G.729 A	Encode	9.7	2162	6215	194	2899	763
	Decode	1.8	2162	1989	113	2899	670
ADT G.729 B	Encode	18.9	4223	8571	194	3401	976
	Decode	3.2	4223	3532	113	3401	838
ADT G.729 A/B	Encode	10.0	4236	8567	246	3200	931
	Decode	2.2	4236	2625	113	3200	838

NOTE: Common Program memory, Common Scratch memory and Common External memory indicate that this memory is shared between encoder and decoder when running independently.

_____ Last update: 10/20/2002

C55x

All Memory usage is given in units of byte.

Product	Function	MIPS (Peak)	Program Memory	Data Memory	Scratch Memory	Per-Channel Data Memory
ADT G.729A/B	Encode	8.3	28354	9272	1976	2156
	Decode	2.0	13654	6310	436	1848

_____ Last update: 02/23/2005

C6000

All Memory usage is given in units of byte.

Product	Function	MIPS (Peak)	Program Memory	Data Memory	Scratch Memory	Per-Channel Data Memory
ADT G. 729 A/B C62x/C67x	Encode	5.2	55433	6476	2400	2148
	Decode	1.3	22921	80	800	2200
ADT G. 729 A/B C64x	Encode	5.1	54962	6476	2400	2148
	Decode	1.2	21721	80	800	2200

_____ Last update: 03/24/2006

Data subject to change, for the most up to date information visit our web site <http://www.adaptivedigital.com>

All trademarks are property of their respective owners.

© 2007 ADAPTIVE DIGITAL TECHNOLOGIES, INC. all rights reserved

API FUNCTIONS

G729_ADT_encodeInit(. . .)	Initializes the G.729 encoder software
G729_ADT_decodeInit(. . .)	Initializes the G.729 decoder software
G729_ADT_encode(. . .)	Executes the G.729 encoder
G729_ADT_decode(. . .)	Executes the G.729 decoder

CONTACT INFORMATION

Web: <http://www.adaptivedigital.com>
Email: information@adaptivedigital.com
Tel: 610.825.0182
Toll Free: 1.800.340.2066
Fax: 610.825.7616
Address: 525 Plymouth Road, Plymouth Woods, Suite 316
Plymouth Meeting, PA 19462 USA