DYSONSPHERE LIMITED

TECHNICAL STUDY · © 2026 DYSONSPHERE LIMITED

# Link-Adaptive Opus ⇄ Codec2 Voice for Satellite Terminals

Edition 1.0, October 2026 · Dragos Oancea · Field: satellite communications, speech coding, voice over IP

## Summary

This document describes a voice system for satellite terminals that switches, during a call, between two very different codec families: **Opus**, a hybrid waveform codec that is good from about 8 kbps upward, and **Codec2**, a parametric vocoder that keeps speech intelligible down to 700 bps. A controller picks one rung of a seven-step ladder for every packet, from 1.3 kbps to 25 kbps on air, using two capacity signals with different delays. A 2-byte in-band header tells the receiver which codec each packet carries, so the switch needs no call signalling.

The study has two parts. The first is a draft invention disclosure: problem, prior art, design and draft claims. The second is a simulation that runs real Opus and Codec2 encoders over a modelled geostationary (GEO) link that fades from 24 kbps to 1.6 kbps and back.

**Main results of the simulation** (28.6 s call, same link and loss pattern for every scenario):

- Opus alone lost **35 %** of the call, because the link stayed below its working floor for about 10 seconds. With Codec2 as the lower part of the ladder, the call lost **6 %**, which is about the same as Codec2 alone.
- On good parts of the link the adaptive system used Opus, so its spectral distortion was 2.6 dB lower than Codec2 alone.
- Hysteresis on the switch back from Codec2 to Opus is a trade-off, not a free gain. Switching on every estimate gave slightly better spectral fidelity; a strict rule cut codec-family switches from 10 to 2 but cost about 1.7 dB.

The broad idea of switching codecs on network conditions, with hysteresis and an in-band codec identifier, is already patented (see Prior art). What may remain new is narrower: a ladder that crosses from a waveform codec to a parametric vocoder with a different packet length, timestamps that let the receiver rebuild gaps of unknown length, a controller that combines local modem-queue occupancy with delayed far-end feedback, and a bridge that lets older satellite terminals carry this voice over their slow circuit-switched data channels.

## Field and background

The field is voice over IP on narrowband and fading satellite links: GEO and LEO mobile-satellite services, maritime and aeronautical terminals, IoT-class satellite modems and older handheld satellite phones.

On these links the usable bit rate can change by an order of magnitude within seconds. Causes include rain fade, blockage by terrain, ships’ superstructure or the user’s own body, beam handovers and contention with other users. Round-trip times are long: about 600 ms on GEO, before any queueing.

Today’s options each fail in one part of that range:

- **A fixed wideband codec** (Opus or AMR-WB at 12–24 kbps) sounds good on a healthy link but breaks into silence once capacity falls below its rate.
- **A fixed low-rate vocoder** (for example Codec2 or a proprietary satellite vocoder at 2.4 kbps) survives deep fades but sounds robotic even when the link could carry much better speech.
- **Opus’s own bitrate adaptation** covers roughly 6 to 510 kbps. Below about 6 kbps it cannot follow, and the call drops out.
- **SIP renegotiation** (re-INVITE) can change codec mid-call, but each exchange takes at least one round trip, several seconds on a congested GEO link, which is longer than many fades.

A further problem is header overhead. A Codec2 700C frame is 4 bytes. A normal IP/UDP/RTP header is 40 bytes. Without header compression, the “700 bps” codec needs more than 5 kbps on air.

## The two codecs

The system pairs two codecs built on opposite ideas. Opus tries to reproduce the speech waveform, or its spectrum, as closely as the bit budget allows; quality grows steadily with bit rate. Codec2 does not try to reproduce the waveform at all: it measures a few parameters of the voice and resynthesizes speech from a model. That is why it can work at a few hundred bits per second, and also why it never sounds natural, however many bits it gets. Their useful ranges barely overlap, which is what makes them a good pair for a ladder.

### Opus

Opus is the IETF standard audio codec, published as RFC 6716 in 2012. It combines two earlier codecs:

- **SILK**, developed by Skype, a linear-prediction speech coder used for narrowband to wideband speech (audio up to 8 kHz), with 10, 20, 40 or 60 ms frames;
- **CELT**, developed by Xiph.Org, a transform (MDCT) coder for general audio and low delay, with 2.5 to 20 ms frames;
- a **hybrid mode** in which SILK codes the band below 8 kHz and CELT the band above it.

The encoder chooses the mode and audio bandwidth (narrowband 4 kHz, medium-band 6 kHz, wideband 8 kHz, super-wideband 12 kHz, fullband 20 kHz) from the bit rate and signal. The format covers 6 to 510 kbps at 8 to 48 kHz sampling. Default algorithmic delay is about 26.5 ms, and a restricted low-delay mode reaches 5 ms (overview).

Features that matter on a satellite link:

- **In-band forward error correction.** SILK can add a low-bitrate redundant copy of the previous frame (LBRR) to each packet, so a single lost packet can be partly rebuilt from the next one. It works in SILK and hybrid modes only.
- **Packet-loss concealment and discontinuous transmission**, built into the decoder and encoder.
- **Machine-learning tools.** libopus 1.5 (March 2024) added neural packet-loss concealment and Deep Redundancy (DRED), which can carry around one second of compact recovery data in each packet, plus quality enhancement for low-rate speech (release notes). libopus 1.6 (December 2025) improved DRED further and added wideband-to-fullband bandwidth extension (release notes). DRED is aimed at exactly the burst losses of satellite links and should be evaluated in the next version of the prototype.

**Licensing.** The reference library uses the BSD licence, and the companies holding the known Opus patents (Xiph.Org, Broadcom, Skype/Microsoft) grant royalty-free licences. Some other companies have claimed patents that may read on Opus, and a commercial pool claims to cover it; none has been tested in court. Counsel should note this for any product.

**Adoption.** Opus is mandatory in WebRTC and is the voice codec of WhatsApp, Signal, Discord and Zoom, among others, so a gateway rarely needs to transcode it.

**Limits for this use.** Below about 6 kbps Opus has no operating mode, and at its lowest rates it falls back to narrowband audio. In our measurements, Opus at 6–10 kbps put no energy above 4.2 kHz; from 12 kbps it switched to wideband.

### Codec2

Codec2 is an open-source speech codec written mainly by David Rowe, first released in 2010 under the LGPL 2.1 and described by its authors as patent-free. The current stable release is 1.2.0 (June 2023).

It is a **sinusoidal (harmonic) coder**, related to multi-band excitation. Every 10 ms it estimates the pitch, decides whether the frame is voiced, and measures the amplitudes of the harmonics of the pitch. It quantizes the spectral envelope as line spectral pairs or, in the lowest modes, with vector quantization, and packs several 10 ms analyses into one transmitted frame. The decoder rebuilds speech as a sum of sine waves at the harmonic frequencies. Input and output are 8 kHz, so audio is narrowband (up to 4 kHz).

Modes, as measured in the library build used for the prototype:

| Mode | Frame | Bits per frame | Notes |
|---|---|---|---|
| 3200 | 20 ms | 64 | Best quality; used as the top Codec2 rung |
| 2400 | 20 ms | 48 | |
| 1600 / 1400 / 1300 / 1200 | 40 ms | 64 / 56 / 52 / 48 | 1200 used as the middle rung |
| 700C | 40 ms | 28 | Vector-quantized envelope; lowest rung |
| 450, 450PWB | 40 ms | 18 | In the source tree but not compiled into this build; 18 bits follows from 450 bps × 40 ms |

**Uses.** Codec2 is the voice codec of FreeDV digital voice on HF radio and of the M17 digital radio protocol. It flew on the LilacSat-1 CubeSat as a voice repeater (2017), and FreeSWITCH supports it, so a gateway built on FreeSWITCH can terminate it directly.

**Limits for this use.**

- Speech sounds synthetic, and quality drops sharply with background noise or with more than one talker, because the model assumes a single voice.
- It has no built-in forward error correction or loss concealment; the system must supply both (section 5).
- Errors in the pitch estimate cause audible artefacts, so a noisy terminal environment may call for noise suppression before the encoder.

## Side by side

| | Opus | Codec2 |
|---|---|---|
| Principle | Linear prediction (SILK) and transform coding (CELT) | Parametric sinusoidal model |
| Bit rate | 6–510 kbps | 0.7–3.2 kbps (0.45 in some builds) |
| Audio bandwidth | 4 to 20 kHz, chosen by bit rate | 4 kHz (narrowband) only |
| Frame length | 2.5–60 ms | 20 or 40 ms |
| Loss handling | In-band FEC, concealment, DRED (1.5+) | None built in |
| Licence | BSD, royalty-free patent licences | LGPL 2.1, stated patent-free |
| Measured delay (prototype) | 6.25 ms beyond packetization | 12.5 ms beyond packetization |

## Measured on the test clip

Each codec was run on the clean 28.6 s test clip, with no losses. The figure is the log-mel spectral distance to the input over 80–3800 Hz, averaged over speech frames (lower is closer to the original).

| Opus 6k | Opus 8k | Opus 10k | Opus 12k | Opus 16k | Opus 24k |
|---|---|---|---|---|---|
| 4.5 dB | 3.3 dB | 2.7 dB | 3.4 dB | 2.4 dB | 1.5 dB |

| Codec2 700C | 1200 | 1300 | 1400 | 1600 | 2400 | 3200 |
|---|---|---|---|---|---|---|
| 12.0 dB | 9.1 dB | 10.0 dB | 9.6 dB | 9.5 dB | 9.3 dB | 8.7 dB |

Three things stand out:

- **Codec2 hardly improves with more bits.** From 1200 to 3200 bps the distance moves only from 9.1 to 8.7 dB, because the model, not the bit budget, limits how close it gets. The extra bits buy smoother, less buzzy speech rather than a closer spectrum. For the ladder, this means the middle Codec2 rungs add little; 700C, 1200 and 3200 are enough.
- **Opus at 12 kbps scores slightly worse than at 10 kbps** in this band, because at 12 kbps it switches to wideband and spends bits on 4–8 kHz, which this measure ignores. A listener would likely prefer the 12 kbps version.
- **There is a gap between the families**: from 3.7 kbps on air (Codec2 3200) to 9 kbps (Opus 8k). An Opus 6 kbps rung (about 7 kbps on air) would narrow it and is worth testing; with DRED from Opus 1.5+, the lowest Opus rung may also cope better with loss.

## Prior art and how the invention differs

A short web search found that the general concept is well covered. The claims below are therefore drawn narrowly.

| Reference | What it discloses | What it lacks relative to this invention |
|---|---|---|
| US 9,729,726 B2, *Seamless codec switching* (Facebook, now Meta; priority 2014-12-05; active) | Codecs mapped to partly overlapping ranges of transmission characteristics; switching deferred while conditions stay in the overlap (hysteresis); receiver told implicitly through packet headers or by a control message; possible dual transmission during the change | No crossing between codecs with different packet durations; no handling of lost packets of unknown codec; no use of local modem-queue state; not aimed at satellite links. **Closest reference: its claims must be read in full.** |
| US 9,251,795, *Adaptive audio codec selection during a communication session* (Apple; priority 2012-09-04; active) | Audio tiers selected from changing bandwidth or packet-loss rate; a codec indication in each audio packet selects the decoder; RTP sequence and timestamps continue across the switch | No hysteresis rule specific to crossing codec families; no concealment bridge; no satellite or legacy-terminal embodiment |
| US 7,307,980, *Change of codec during an active call* (Cisco; priority 1999; expired) | Switching to a second codec mid-call after the far end reports degraded quality, by renegotiation (H.245) | Uses signalling, not in-band switching; expired, so freely usable |
| US 7,558,359, *AMR vocoder rate adaptation* (priority 1998) | Receiver measures channel quality and requests a new AMR rate on the return link | One codec family; terrestrial TDMA |
| US 11,758,051, *Machine learning-based audio codec switching* (continuation of a 2020 filing) | Phone switches from EVS full-band to EVS super-wideband channel-aware mode when an RF measure (RSSI, RSRP or RSRQ) crosses a threshold, and back | Within one codec family; signalled through the base station |
| Codec2 and FreeDV (open source, LGPL) | Speech at 450–3200 bps; widely used for digital voice over HF and VHF radio | One codec at a time; no switching to a wideband codec |
| 3GPP AMR/AMR-WB codec mode requests; Opus in-band bandwidth and bitrate control | Mode changes inside one codec, signalled in-band | Do not cross codec families or packet sizes |

**Where novelty plausibly remains:**

- Crossing between a waveform or hybrid codec and a parametric vocoder whose packets have **different durations**, with switches allowed only on a common time grid, and a header timestamp counted in base-frame units, so the receiver can rebuild the length of a gap even when it does not know which codec the lost packets used.
- A capacity estimate that subtracts the sender's **own modem-queue backlog**, known immediately, from far-end delivery feedback that arrives one round trip late. This lets the sender step down within about 120 ms of a fade instead of after a full round trip.
- An **asymmetric rule for crossing families**: step down quickly, cross back up only after a hold time and, optionally, extra headroom.
- A **concealment bridge**: the old codec's loss-concealment output is crossfaded into the new decoder's first frame, and the Opus encoder and decoder are reset on re-entry.
- A **legacy-terminal bridge** that carries Codec2 plus forward error correction over the roughly 2.4 kbps circuit-switched data channel of older terminals (for example Iridium CSD, see CLS), with a gateway that converts to Opus or AMR-WB, and passes AMR-WB through untouched toward mobile networks.

Each element could be argued obvious against the references above combined with ordinary engineering practice. A patent attorney should judge that. A formal search on Espacenet and Google Patents is still needed; this scan used web search only.

## Problems solved

1. **Calls drop in deep fades.** With Codec2 rungs down to 1.3 kbps on air, the call keeps going at link rates where Opus cannot work.
2. **Poor sound on good links.** Unlike a fixed low-rate vocoder, the system returns to Opus, and its wideband audio, when capacity allows.
3. **Slow reaction on long-delay links.** Renegotiation and far-end feedback each take at least one round trip. Using local queue state, the sender reacts to a fade within about 120 ms.
4. **Header overhead.** The design assumes a 2-byte voice header plus compressed IP/UDP (3 bytes). This keeps

the overhead of an 80 ms Codec2 packet at 5 bytes rather than 40.

5. **Older terminals left out.** Handsets and modems with only a slow data channel can join IP voice calls through the legacy bridge.

## Detailed description

### 1. System overview

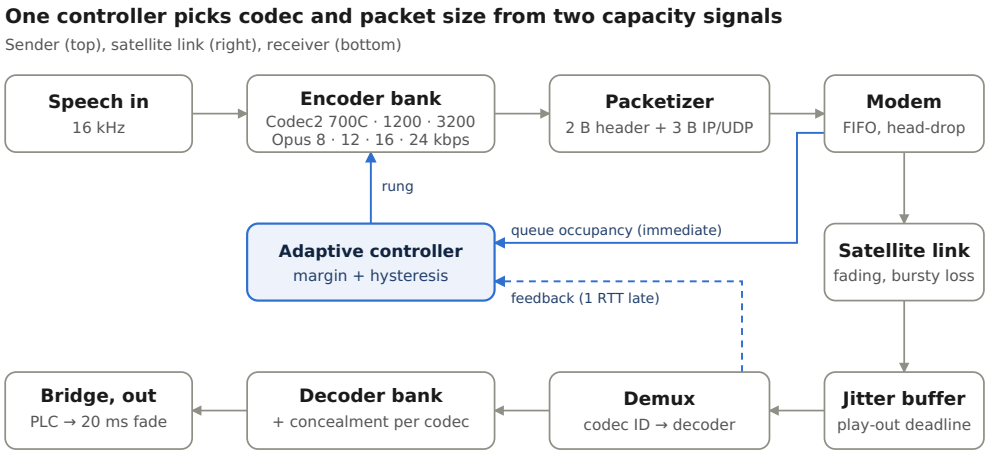



**Fig. 1.** Signal path. The controller chooses a codec rung for each packet from two capacity signals: the local modem queue, known now, and delivery feedback from the far end, known one round trip later.

A sending terminal keeps one encoder for each rung. For every packet, the controller picks a rung; the matching encoder fills the packet, and the packetizer adds the header. Packets wait in the modem's first-in, first-out queue. Packets that can no longer reach the receiver before their play-out deadline are dropped from the head of the queue, so they do not delay newer packets.

The receiver reads the codec identifier in each header and passes the packet to the matching decoder. Lost packets are concealed with the decoder of the last codec received. When the codec family changes, the bridge crossfades between the two decoders (section 5).

### 2. Codec ladder

| Rung | Packet duration | Payload | On air (with 5 B overhead) |
|---|---|---|---|
| Codec2 700C | 80 ms (2 frames) | 8 B | 1.3 kbps |
| Codec2 1200 | 80 ms (2 frames) | 12 B | 1.7 kbps |
| Codec2 3200 | 80 ms (4 frames) | 32 B | 3.7 kbps |
| Opus 8 kbps, constant bitrate, in-band FEC | 40 ms | 40 B | 9 kbps |
| Opus 12 / 16 / 24 kbps | 40 ms | 60 / 80 / 120 B | 13 / 17 / 25 kbps |

Codec2 rungs use 80 ms packets to spread the fixed overhead over more speech. Opus rungs use 40 ms packets to keep delay and the cost of each lost packet lower. AMR-WB is not on the ladder: its 6.6–23.85 kbps range overlaps Opus, and Opus is equal or better at the same rate. AMR-WB is used only when the far end is a mobile network or a terminal that supports nothing else. That choice is made at call setup, and the gateway then passes AMR-WB frames through without transcoding.

### 3. In-band header

> byte 0: rung ID (4 bits) | flags (4 bits: FEC present, switch pending, 2 reserved)  byte 1: timestamp in 40 ms units, modulo 256

Counting the timestamp in 40 ms units, not in packets, matters when packet durations differ. After a loss, the receiver can compute how much audio is missing from the next timestamp it receives, without knowing which codec the lost packets used. Eight bits cover 10.24 s, far longer than any jitter buffer. Four bits allow up to 16 rungs.

### 4. Controller

The controller works on an effective capacity estimate:

$$\hat{C}(t) = \mathrm{EWMA}[\, C_{fb}(t - \mathrm{RTT}) \,] - Q(t) / \tau$$

$C_{fb}$ is the delivery rate reported by the far end, which arrives one round trip late. Q(t) is the number of bits waiting in the local modem queue, and τ (0.5 s here) is the time within which the queue should drain. The queue term reacts at once when the link slows, because packets stop leaving. The usable rate is a fraction of the estimate, $M \cdot \hat{C}$, with margin M = 0.85.

Rules for each new packet:

- **Down:** if a lower rung fits and a lower rung has been wanted for 120 ms, step down, several rungs at once if needed.
- **Up within a family:** step up after the higher rung has fitted for 1.5 s.
- **Up across families (Codec2 to Opus):** step up after the Opus rung has fitted for a hold time H, optionally with extra headroom X (the usable rate must exceed X times the rung's rate). The simulation default is H = 0.5 s and X = 1.0; section 7 shows the effect of other values.
- **Grid:** a switch to an 80 ms rung is made only at a multiple of 80 ms, so all packets fit one time grid.

### 5. Concealment bridge

When the decoded family changes at sample *b*, the receiver:

1. asks the old decoder for 20 ms of loss-concealment audio starting at *b* (Opus has built-in concealment; for Codec2, the last frame is repeated with a falling gain);
2. resets the Opus decoder if Opus is the new family, so that stale state from minutes earlier is not used;
3. crossfades linearly from the old codec's concealment to the new decoder's output over those 20 ms, after aligning both decoders' delay to the input (6.25 ms for Opus and 12.5 ms for Codec2, measured).

### 6. Legacy-terminal embodiment

A small adapter, either an app on a phone tethered to the terminal or a USB or Bluetooth dongle, runs Codec2 at 700–1200 bps with forward error correction over the terminal's circuit-switched data call (about 2.4 kbps on Iridium CSD). A ground gateway terminates the data call and joins it to the IP voice network, converting to Opus or AMR-WB as the far end needs. Older terminals gain encrypted, IP-reachable voice with no hardware change.

## Simulation study

### Setup

The prototype (folder `prototype/`) calls the system libopus and libcodec2 libraries from Python. The input is a 28.6 s synthetic English speech clip at 16 kHz.

**Link model:**

- Capacity trace: 24 kbps for 4 s; a fade to about 5 kbps; a deep fade to about 1.6 kbps from 10.5 to 14 s; a marginal stretch from 15.5 to 22 s that wobbles between about 6.5 and 13.5 kbps, built to tempt a controller into flapping; then recovery to about 22 kbps. Multiplicative log-normal noise is added.
- Burst loss with a two-state (Gilbert-Elliott) model: 0.5 % loss in the good state, 40 % in the bad state, 3 % during the deep fade.
- 280 ms one-way delay, 250 ms play-out budget, far-end feedback delayed by 600 ms.

**Scenarios:** the adaptive controller described above; a naive controller that picks the best-fitting rung for every packet with no hysteresis; and the same adaptive controller restricted to the Opus rungs or to the Codec2 rungs. All four see exactly the same link and loss pattern.

**Quality measure:** the log-mel spectral distance described in "Measured on the test clip" above (80–3800 Hz, speech frames only, lower is better). It is a rough proxy, not PESQ or POLQA, and it ignores the extra 4–8 kHz band that wideband Opus provides. The prototype uses libopus 1.4, so the machine-learning features of Opus 1.5 and 1.6 are not part of these results.

## Results

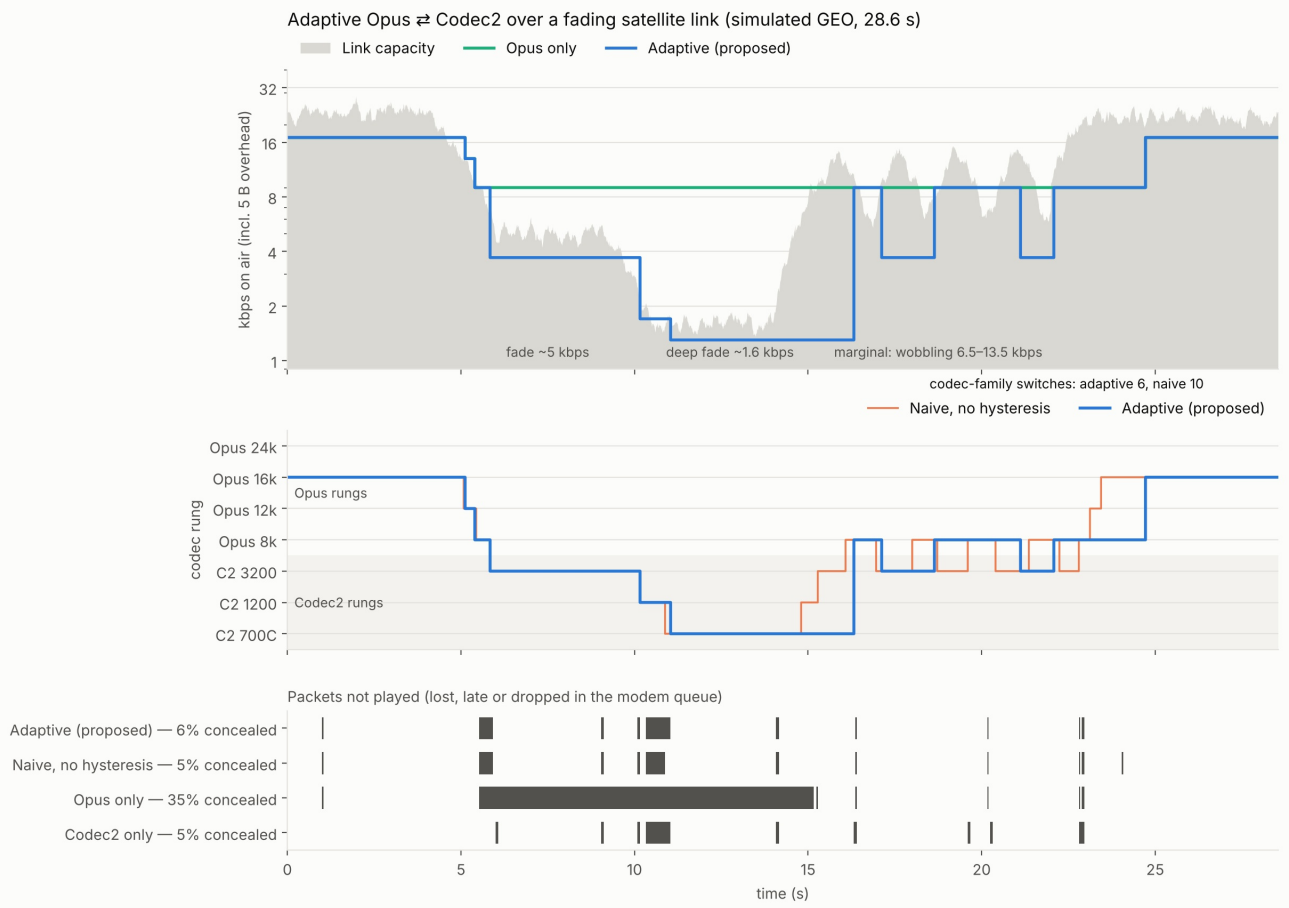



**Fig. 2.** Top: link capacity and the on-air rate chosen by the adaptive controller and by Opus alone. Middle: the rung chosen by the adaptive and naive controllers. Bottom: packets that could not be played.

| Scenario | Audio concealed | Codec-family switches | Mean on-air rate | Spectral distance |
|---|---|---|---|---|
| Adaptive (proposed) | 5.6 % | 6 | 8.7 kbps | 8.3 dB |
| Naive, no hysteresis | 5.2 % | 10 | 9.0 kbps | 8.0 dB |
| Opus only | 34.7 % | 0 | 11.6 kbps | 15.5 dB |
| Codec2 only | 5.0 % | 0 | 3.2 kbps | 10.9 dB |

## 7. Effect of the crossing rule

The switch from Codec2 back to Opus was swept over headroom X and hold time H, with everything else fixed:

| Headroom X | Hold H | Family switches | Audio concealed | Spectral distance |
|---|---|---|---|---|
| 1.0 | 0 s | 8 | 5.7 % | 8.3 dB |
| 1.0 | 0.5 s (default) | 6 | 5.6 % | 8.2 dB |
| 1.0 | 1.5 s | 4 | 6.2 % | 9.4 dB |
| 1.0 | 3.0 s | 2 | 6.0 % | 9.9 dB |
| 1.1 | 0.5 s | 2 | 6.3 % | 9.5 dB |
| 1.3 | 3.0 s | 2 | 6.3 % | 9.9 dB |

## Findings

1. **The Codec2 floor is what keeps the call alive.** Opus alone lost about 10 s of speech during the fade and deep fade; every configuration with Codec2 rungs lost about 5–6 %, mostly in burst-loss events that hit all scenarios.
2. **Switching to Opus when possible improves quality.** Over the whole call, the adaptive system was 2.6 dB better than Codec2 alone on the spectral measure, and it delivered wideband audio during the good periods,

which the measure does not count.

3. **Hysteresis trades fidelity for stability.** On the marginal stretch, a stricter crossing rule kept the call on Codec2 longer. That cut family switches from 8 to 2 but cost up to 1.7 dB. Whether frequent switching is audibly worse than staying on Codec2 is a listening question that this measure cannot answer.
4. **Header compression is a requirement.** With 40-byte IP/UDP/RTP headers, the 700C rung would need about 5.3 kbps on air, and the low end of the ladder would not exist.

### Limits

- Synthetic speech, no background noise. Codec2 is known to degrade in noise, which may move the best switching point upward.
- Rough spectral measure; no listening test and no PESQ or POLQA scores.
- One link trace and one loss seed; results need repeating over many traces.
- The receiver in the simulation uses the true duration of lost packets; with the header of section 3 a real receiver gets the same information from the timestamps.
- Codec2 700C clipped briefly on one loud syllable (13 samples). A limiter in the decoder path would remove this.

## Draft claims

**1.** A method of transmitting speech from a terminal over a satellite link, comprising: encoding successive segments of speech, each with one of a plurality of codec configurations ordered by bit rate, the plurality comprising at least one configuration of a parametric vocoder producing packets of a first duration and at least one configuration of a transform or hybrid codec producing packets of a second, shorter duration; estimating an available capacity of the link; selecting the codec configuration for each packet based on the estimated capacity; and including in each packet a codec identifier and a timestamp counted in units of a base duration that divides both the first and the second duration, wherein a change to a configuration with the first duration is made only at a time that is a whole multiple of the first duration.

**2.** The method of claim 1, wherein the parametric vocoder is Codec2 and the transform or hybrid codec is Opus.

**3.** The method of claim 1, wherein estimating the available capacity comprises combining a delivery rate reported by a receiving end, delayed by at least one round-trip time, with an occupancy of a transmit queue of the terminal’s own modem.

**4.** The method of claim 3, wherein the estimate equals a smoothed reported delivery rate minus the queue occupancy divided by a target drain time.

**5.** The method of claim 1, wherein a change from a parametric-vocoder configuration to a transform or hybrid codec configuration requires the estimated capacity to exceed the bit rate of the target configuration for a hold time, and a change in the opposite direction is made within a shorter time.

**6.** The method of claim 5, wherein the change from the parametric vocoder further requires the estimated capacity to exceed the bit rate of the target configuration by a headroom factor greater than one.

**7.** The method of claim 1, further comprising dropping, from the head of the transmit queue, packets that can no longer reach the receiver before their play-out deadline.

**8.** A method of receiving speech over a satellite link, comprising: receiving packets each carrying a codec identifier and a timestamp; decoding each received packet with the decoder named by its codec identifier; determining the duration of missing audio from the timestamps of received packets; concealing missing audio with the decoder of the most recently received codec; and, when the codec family of consecutive received packets changes, producing concealment audio from the previous decoder for a bridging interval and crossfading from that concealment audio to the output of the new decoder over the bridging interval.

**9.** The method of claim 8, further comprising resetting the state of the new decoder before decoding its first packet after the change.

**10.** The method of claim 8, further comprising aligning the outputs of the previous and new decoders by their respective algorithmic delays before crossfading.

**11.** A system comprising: a satellite terminal having a circuit-switched data channel of no more than 4.8 kbps; an adapter connected to the terminal that encodes speech with a parametric vocoder at no more than 1.2 kbps, adds forward error correction, and sends the result over the data channel; and a ground gateway that terminates the data channel, decodes the speech and forwards it into an IP voice network encoded with a different codec selected according to the capability of the far end.

**12.** The system of claim 11, wherein, when the far end is a mobile network supporting AMR-WB and the incoming stream is AMR-WB, the gateway forwards the AMR-WB frames without decoding them.

**13.** A terminal, gateway or computer-readable medium storing instructions that, when executed, perform the method of claim 1 or claim 8.

## Next steps and questions for patent counsel

- Read the full claims of US 9,729,726 and US 9,251,795; together they likely cover network-driven codec switching with hysteresis and per-packet codec identifiers. Check that claims 1, 3, 8 and 11 here fall outside them.
- Run a formal search on Espacenet and Google Patents in CPC G10L 19/22 (mode decision between coding types), G10L 19/24 (variable rate) and H04B 7/185 (satellite systems), including filings by Inmarsat (Viasat), Iridium, Thuraya, Hughes, Qualcomm and Ericsson.
- Extend the simulation: real recorded speech with background noise, many link traces and loss seeds, a LEO handover trace, and a short listening test (MUSHRA or a simple preference test) on the marginal stretch to settle the hysteresis question.
- Build the legacy bridge on a real terminal and measure delay and intelligibility over a CSD call.
- Do not publish or present the idea before filing; public disclosure can destroy novelty in Europe.

**Open questions for counsel**

- Is the legacy-terminal bridge (claim 11) a better lead claim than the switching method, given the prior art on switching?
- Can claim 1 be worded so it covers any pair of codecs with different packet durations, not only Codec2 and Opus?
- Should the receiver-side bridge (claim 8) be filed separately?

© 2026 Dysonsphere Limited. All rights reserved. No part of this publication may be reproduced, distributed or transmitted in any form without the prior written permission of the publisher, except for brief quotations in reviews and other uses permitted by copyright law. Opus, Codec2, FreeSWITCH and other names are trademarks of their respective owners and are used only to identify the software discussed.