Inner Circle Newsletter Summer 2026

Don’t be afraid to stand out in your field

Contents

Find Us Online

YouTube https://www.youtube.com/@OpenResearchInstituteInc
Twitter (X) https://x.com/openresearchins
FaceBook https://www.facebook.com/openresearchinstitute/
LinkedIn https://www.linkedin.com/company/open-research-institute-inc
Fediverse https://micro.blog/OpenResearchIns

OPEN SOURCE HARDWARE AND SOFTWARE NEWSLETTER SIGNUP
FREE AND 100% VOLUNTEER DRIVEN

Want more Inner Circle Newsletters? Use the QR code at left or go to http://eepurl.com/h_hYzL and sign up.

Inner Circle Sphere of Activity

4 April 2026 - BSides San Diego RF Village participation and demonstation.

26 - 28 June 2026 - ESA FutureGEO Workshop, Friedrichshafen HAMRADIO 2026, participation and demonstration. Find us at AMSAT-DL booth.

6 - 9 August 2026 - DEFCON RF Village participation and demonstration.

If you know of an event that would welcome ORI, please let your favorite board member know at our hello at openresearch dot institute email address.

Thank you to all who support our work! We certainly couldn't do it without you.

Anshul Makkar, Director ORI
Steve Conklin, CFO ORI
Michelle Thompson, CEO ORI
Matthew Wishek, Director ORI

Solutions Released for Lunar Descent CTF

Successfully debuted at BSides San Diego 2026, ORI's Lunar Descent CTF has been updated with additional documentation and example solutions

One of the many activities that we do at ORI is to participate in and host contests and puzzles. Capture the Flag competitions (CTFs) are a particular type of competition commonly encountered in the cybersecurity and hacking communities. CTFs pose a set of challenges for participants to solve, often drawn from a particular theme or subject. CTFs can be tackled by individuals, but the interdisciplinary and complex nature of many CTFs require a team working together in order to score enough points to be at the top of the leader board.

The Lunar Descent CTF debuted at BSides San Diego, held at San Diego State University on 4 April 2026. This is an ORI original, and all documents can be found on GitHub at https://github.com/OpenResearchInstitute/lunar-descent-ctf. The CTF is based on a real signal processing problem in a radar altimeter.

ISRO's KaRA radar altimeter guided Chandrayaan-3 to a soft lunar landing on 23 August 2023. The Radar Altimeter Processor (RAP) computes altitude and velocity from FMCW chirp signals, running on a single Xilinx Virtex-5 FPGA. This CTF uses a Python model of that system, faithful to the published paper, where the altimeter feeds a landing autopilot. The altimeter works perfectly. However, once the system became operational, the autopilot kept crashing in last-minute testing. Why? The answer is the solution to the CTF.

The Lunar Descent CTF is self-scoring and self-paced, which reduced the workload of the booth volunteers. A copy of the IEEE Aerospace and Electronics System Society Magazine article that inspired the CTF was available for people to read. ORI-themed USB drives with a copy of the CTF repository contents were given away to anyone that wanted one to take home with them. The CTF attracted a lot of attention at the ORI-supported RF Village at BSides San Diego, and reviews from participants were very good. While it was anticipated that some participants would be able to solve the CTF in two hours or so, it took the first place winner most of the day, as they wanted to see all the other villages at BSides, and also attend talks. The self-paced “take home” format greatly reduced the pressure on the participants to commit to a contest over seeing the rest of the show. We received positive feedback about this decision, and will use the self-scoring and self-paced methodology in future CTFs. This approach was inspired by the self-scoring BLE CTF, which can be found at https://github.com/hackgnar/ble_ctf

Poster created for Bsides San Diego RF Village

If you would like to try the CTF, then clone the lunar-descent-ctf repository, start with the top-level README.md, and avoid looking in the spoilers/ directory. The spoilers/ directory has a README file for staff running the event, detailed directions on how to score challenge #1, and also has two example solutions for challenge #2 and challenge #3. Since the solutions are tucked away into a separate directory, one can clone the repository and attempt the CTF without accidentally seeing any of the solutions or hints. However, if you get stuck, or want to compare your solution to the example solutions, all you need to do is drop down into the spoilers/ directory and read through the documentation and solutions.

If you want to use this CTF at your own event, please feel free to clone it and adapt it for your setting and audience.

ORI futureGEO HAMRADIO 2026
Workshop Report

ORI participated in the futureGEO Community Workshop at HAMRADIO 2026. Held annually in Friedrichshafen, Germany, HAMRADIO is the largest amateur radio event in Europe, attracting 15,000 people from 50 countries. The workshop was 4:00 pm – 5:30 pm on 27 June 2026 and overflowed the capacity 40 room.

The description of the workshop was "ESA, in collaboration with AMSAT-DL and the international amateur satellite community, intends to assess the feasibility of designing, developing, and operating an advanced satellite communications payload suitable for supporting in-orbit experimentation by the amateur satellite community in GEO-orbit. During this workshop we want to present 2 to 3 different consolidated amateur satellite mission concepts from the study. Further we want to collect feedback from the community for a final proposal."

The futureGEO workshop at HAMRADIO 2026. At the desk in front, Peter Gülzow, president of AMSAT-DL, presents about the history and context of QO-100. Frank Zeppenfeldt was the second speaker, and defined the ESA's role with futureGEO, how Phase 1 would conclude, and what Phase 2 would entail. Michelle Thompson (far left) was the third and final speaker, and presented the working prototype from ORI for futureGEO submission.

Peter Gülzow, president of AMSAT-DL, opened the workshop with a summary of QO-100 history and context. He shared that the operator (Es'Hail) expects another 9 years of life for the satellite. This means that we are a bit less than half-way through the lifetime of QO-100 transponder.

Peter introduced three potential categories of amateur radio satellite spacecraft.

Enhanced QO-100+ Bent Pipe
Digital Innovation Lab
High Frequency Pathfinder

These appear to be very high-level "consolidated mission concepts" from the workshop description.

The next speaker was Frank Zeppenfeldt, PD0AP (M2M, IoT, Smallsats and Spectrum at European Space Agency). He explained that planning for what comes after QO-100 is the primary motivation of ESA's futureGEO project. Frank said that a futureGEO payload would almost surely be built by industry, and that this workshop was intended to close out Phase 1 of the futureGEO project, and that a report would be written and published about Phase 1 in August 2026.

What is futureGEO?

The FutureGEO project is an ESA (European Space Agency) initiative to develop a future amateur radio payload for a geostationary satellite. The aim is to identify potential partners who would like to actively participate in the definition and development of a new amateur radio payload for a future geostationary satellite.

ESA proposed the idea of a geosynchronous satellite back in December of 2023. Frank secured €250,000 in funding to investigate the possibility of an amateur satellite or payload in geostationary orbit. The announcement was made in a presentation at the AMSAT-UK Colloquium.

ORI's Presentation

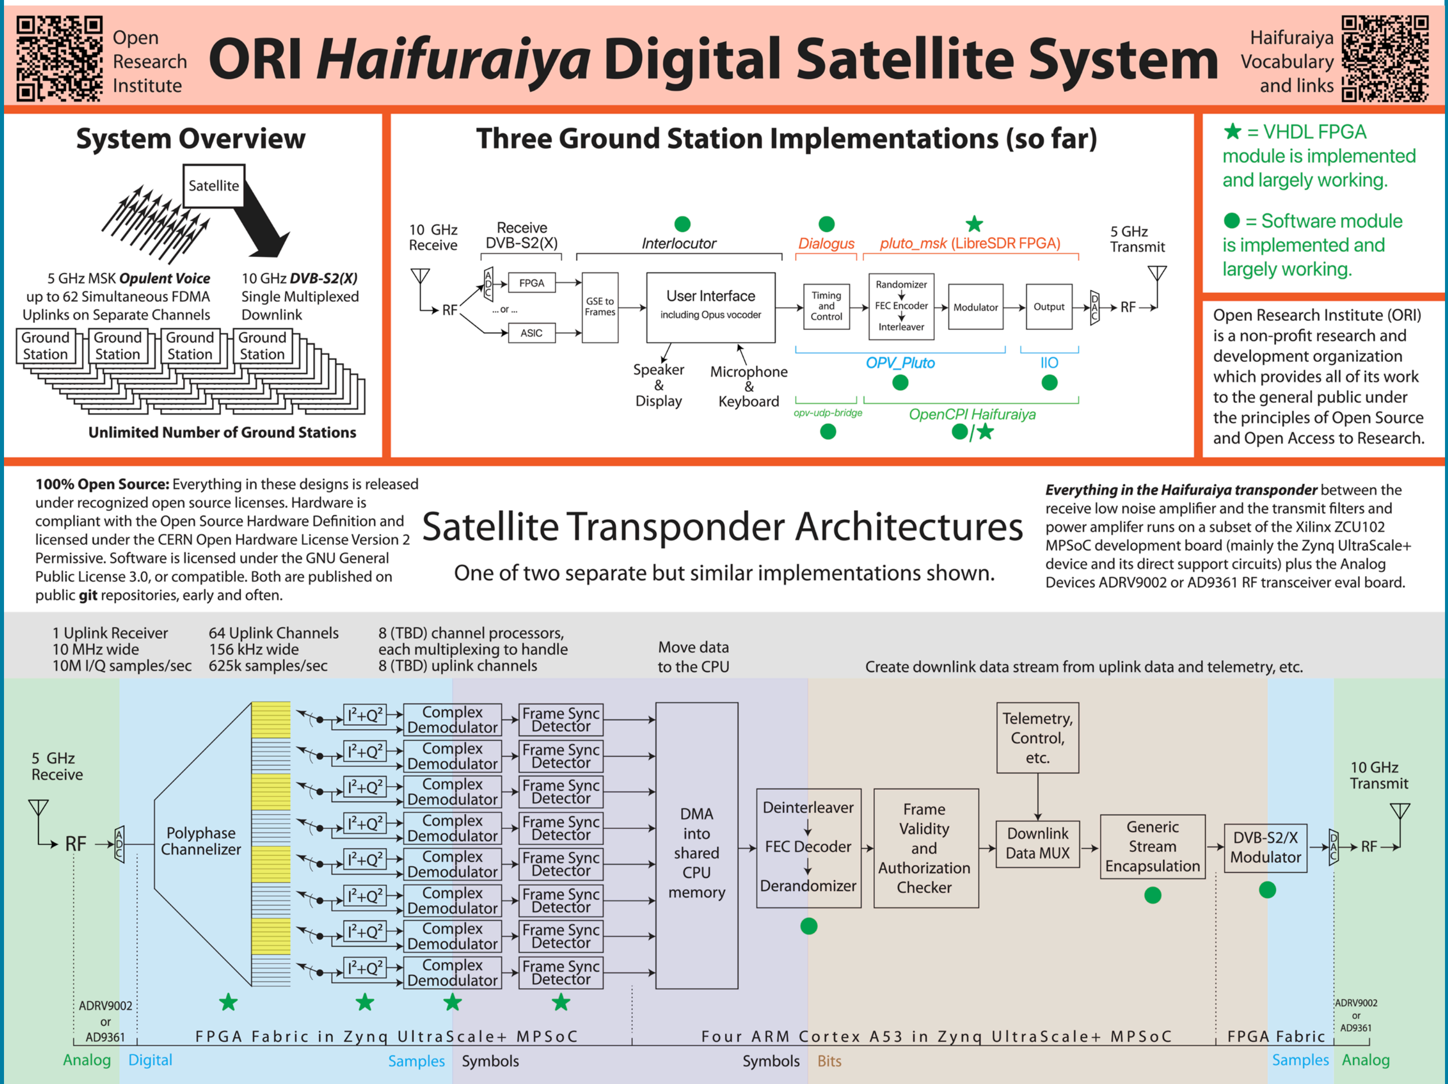

Michelle Thompson was the third and final presenter at the workshop, speaking on behalf of ORI's futureGEO proposal team. ORI is a signatory to the letter of interest (LOI) from AMSAT-DL, participated in the 2025 futureGEO Workshop at AMSAT-DL Symposium, and has submitted a one-page proposal as well as a white paper about ORI's proposed design. Both of these documents can be found in the futureGEO GitLab as well as at the above links.

Michelle described the working hardware, firmware, and software implementations of Haifuraiya, or "High Flyer". Haifuraiya is a functional prototype digital regenerative multiplexing amateur radio satellite system, with superior voice quality and integrated voice, text, and data. There are three implementations of the Opulent Voice ground station, and two implementations of the satellite segment. Processing occurs in the satellite. This is not a double-hop system. Uplink is frequency division multiple access Opulent Voice and downlink is DVB-S2 time-division multiplex. The design is carefully engineered and tested, completely open source, and fully documented. This poster summarizes the status of ORI's proposal for futureGEO as of June 2026.

The poster was displayed at the AMSAT-DL booth for the three days of the show. There was consistent, strong, positive feedback for the design. Questions from the show participants were intelligent and meaningful, ranging from UX/UI theory and operations to the specific mathematics of the demodulator.

AMSAT-DL kindly and generously supported ORI with space and time to communicate a real-life working implementation of the futureGEO concept. ORI takes the futureGEO project seriously, and will complete both a terrestrial repeater version as well as a space-qualified communications board prototype of Haifuraiya. ORI will participate in Phase 2 of futureGEO, and looks forward to the next steps of the project.

AMSAT-DL HAMRADIO 2026 Booth with Special Guest "Adventure Time".

The many give-aways and posters informed and educated a steady stream of visitors.

Attendees could pick up reference models of Erminaz pocket cube spacecraft, renew their membership, talk about futureGEO, and more.

Haifuraiya Flight Dev Boards Land in Remote Lab

Information about our terrestrial design progress vs our flight design progress, and why we will have a new development board in Remote Labs West.

Haifuraiya is our microwave-band amateur radio satellite and terrestrial transponder design. Up until now, we've used the ZCU102 FPGA development board from Xilinx, fitted with an Analog Devices ADRV9002 radio card, to prototype and test the design in the lab and over the air. This has allowed us to make rapid quality progress. Now that we are starting to see end-to-end over-the-air communications results, and now that futureGEO project work is transitioning to CFPs to industry for flight hardware, we need to respond with development equipment that allows flight hardware design, testing, documentation, and acceptable parts lists.

In practice, this means we need a parallel space-oriented development station in addition to the terrestrially-oriented ZCU102. The right answer for our design looks to be the VCK190. This development board has identical hardware to the space-qualified Versal part that provides what we need to port Haifuraiya for space.

There are some key differences between the Ultrascale+ part on the zcu102 development board and the Versal part on the VCK190.

The Versal chip replaces the classic UltraScale+ DSP48 slices with DSP58 slices. The Versal slices support native floating-point math and wider vector operations. When coding our channelizer and DVB-S2 encoders by hand in RTL, our math blocks "should" synthesize directly into these larger DSP58 structures. Testing will confirm this particular port, but it should be an easy change.

Booting, on the other hand, is radically different. Traditional UltraScale+ parts boot via a master hardware configuration bitstream. The Versal part on the VCK190 boots via a dedicated Platform Management Controller (PMC). The PMC reads a single unified file called a Programmable Device Image (.pdi), which bundles up our FPGA logic, the NoC routing data, and our PetaLinux bootloaders (FSBL/U-Boot) into one package. This is more like the build process for the Pluto than the build process for the ZCU102.

Why use the VCK190? Pin and architecture compatibility with the space qualified version of the FPGA part is the big win. The VCK190 features the commercial XCVC1902 Versal AI Core, which shares the exact same architecture as the space-qualified XQRVC1902. This allows us to prototype our DSP and RTL logic on the same architecture that will be used in space. This is the first time we've stood up specifically space-oriented hardware in Remote Labs.

Development flow is another big reason. We will use the exact same software toolchains (Vivado and Vitis) that we use for terrestrially-oriented designs on the space-grade equivalent.

Space-qualified boards and chips (Class B and Class Y certified) command massive premiums and require specialized handling. The VCK190 allows us to iterate and debug our hardware without risking expensive, flight-ready hardware (which can cost tens of thousands of dollars) early in the design cycle. So, we are definitely not too late in the design phase for this equipment upgrade.

But, are we too early? Also no. We have started to see our design work end-to-end in multiple implementations. Now is the time to start porting for space qualified parts, if we want to be ready for both Phase 2 of futureGEO ESA project, and also pitch to US-based satellite companies for ride-shares and hosted payload opportunities.

How do we move to space-grade hardware? First, we'll continue with extensive simulation with the space-qualified part architecture. Then, we implement on hardware in the lab on the development board, and investigate live signals with integrated logic analyzers and over-the-air metrics. Then, when our design is ready to move from lab testing to custom hardware, we will transition to the space-grade AMD Versal XQR series in a custom board layout.

Instead of the XCVC1902, our flight model will use the XQRVC1902. It is a subtle difference in naming, but there are substantial differences between the two parts. The space-grade chip is built with ruggedized, lidless packaging to endure demanding conditions (like much wider thermal variations than commercial or industrial grade. We are talking from -55°C to +125°C) and incorporates built-in circuit techniques to protect against radiation-induced Single Event Upsets (SEUs). There will undoutedly be challenges, setbacks, and learning curves along the way. Our volunteer team is up to the job.

We've addressed the FPGA, but what about radio access? Using two identical ADRV9002 cards would simplify the software driver ecosystem. We know how to deal with the 9002. For uplink at 5 GHz, the ADRV9002 can be used directly. However, it cannot handle 10 GHz by itself. For that, we would need something like the ADI XUD1A upconverter.

See the figure for what the lab setup would look like with this configuration.
We'd use a 7.000 GHz local oscillator at 3.45 GHz IF for the transmit side. The XUD1A gives us 4 independent TX/RX channels. Since Haifuraiya has 64-channel FDMA stream multiplexed into a single TDM carrier, we only need one physical RF stream for the current downlink. This means we will have 3 spare hardware channels on the XUD1A available for multi-antenna array experiments, telemetry, useful beacons, or whatever we can come up with.

Now, there is some coax connector stuff going on here that we should not ignore. The XUD1A has different connectors than the radio cards. Adapters or additional cables may be necessary.

If you want to be a part of the team, and learn by doing through modern and innovative open source work, please visit https://openresearch.institute/getting-started and sign up.

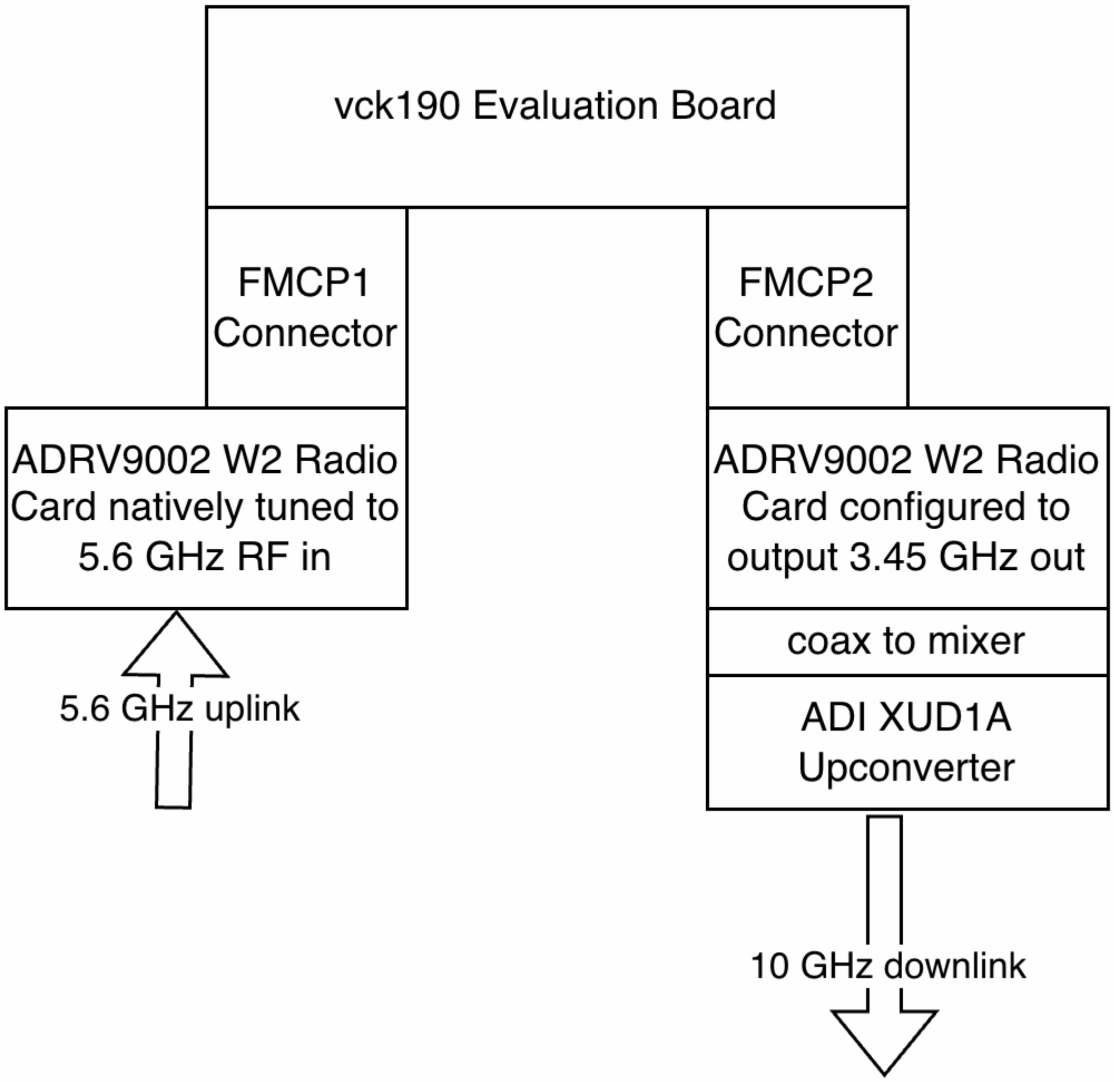
vck190 Evaluation Board
FMCP1 Connector
FMCP2 Connector
ADRV9002 W2 Radio Card natively tuned to 5.6 GHz RF in
ADRV9002 W2 Radio Card configured to output 3.45 GHz out
coax to mixer
ADI XUD1A Upconverter
5.6 GHz uplink
10 GHz downlink

Radiation Mitigation for ORI Designs

The Haifuraiya flight-build team has started work with the VCK190 (https://www.amd.com/en/products/adaptive-socs-and-fpgas/evaluation-boards/vck190.html). This is a Versal AI Core series evaluation kit. Volunteers are learning how radiation mitigation interacts with the FPGA resource budget on the flight target. The FPGA on the VCK190, the XQR Versal AI Core XQRVC1902, is a radiation tolerant part, so porting to this development board allows us to port to the same architecture as the space-qualified part.

One of the questions that came up was “we never allocated fabric for XilSEM, could it bust our budget?” Short answer: no, but read on, because the part of the design that *can* cost fabric is a different thing.

What changed from the older parts? Why this is not like UltraScale+?

XilSEM (Xilinx Soft Error Mitigation) scrubs configuration memory (CRAM) to keep single-event upsets from accumulating in the bits that define your routing and logic.

UltraScale+ / 7-series: XilSEM was a soft IP core. It synthesized into the programmable logic and cost real fabric (LUTs, flip-flops, BRAM). On those parts you had to budget for it.

Versal (our flight part): XilSEM moved into the hardened Platform Management Controller (PMC). It is firmware running on the PMC’s MicroBlaze (PPU), using dedicated config-frame scan hardware and the PMC’s own RAM. It does NOT synthesize into the programmable logic.

What does this mean for us? Enabling XilSEM consumes PMC processor cycles and PMC RAM, not the DSP / LUT / BRAM that the channelizer, demod cores, and DVB-S2 encoder compete for. It does not move the current ~1,495 DSP utilization (76%) number at all. This is good!

On Versal, XilSEM costs you no programmable-logic fabric. It is PMC firmware, not a soft IP core. So, it does not threaten the DSP/LUT/BRAM budget. The fabric cost of radiation mitigation lives in selective triple modular redundancy (TMR) of critical logic. The 16:1 multiplexed design baseline was chosen partly to leave headroom for exactly that.

Scrubbing is not the same as protection

XilSEM keeps configuration memory clean, but it corrects on a scan cycle with millisecond-scale latency. During that window an upset can still cause wrong behavior. Scrubbing prevents accumulation of upsets. But, it does not, by itself, protect against the immediate functional effect of one, and it is not sufficient on its own for a high-radiation space environment.

So a real flight design combines at least two mechanisms:

1. XilSEM (config-memory hygiene): Use this fabric-free option on Versal. Prevents fault accumulation in CRAM and NPI registers. Essentially free in our budget.

2. Selective TMR (functional protection): triple-modular redundancy on the logic that cannot tolerate a transient upset. This is where fabric goes: up to 3x plus voters on whatever we choose to triplicate.

TMR is a deliberate, selective decision that we make, applied to the most upset-critical logic, not a blanket tax on the whole design. But it is the real fabric-cost lever, and it is the thing to size against our headroom.

Why the 16:1 baseline already accounts for this

We chose 16:1 (76% DSP on the flight part) over 8:1 (~91%). 16:1 leaves ~24% of the DSP free on the XQRVC1902, plus we reclaimed ~450 DSP by sharing the power detector. That headroom is the budget selective TMR draws against when the radiation work item is executed. At 8:1 (~91%) there would be almost no room to triplicate anything. The 16:1 margin is not slack. It is reserved space for functional redundancy. So radiation mitigation is not unbudgeted. XilSEM is fabric-free, and the TMR headroom was deliberately preserved by the baseline decision.

Where do XilSEM's costs actually show up?

Power: the background scan adds power draw and affects the flight power/thermal budget.

PMC RAM: XilSEM firmware and state live in PMC RAM and affect the PMC memory budget.

Latency / reliability: millisecond scrub-and-correct latency affects reliability and FDIR (fault detection, isolation, recovery) analysis, including SEFI handling. None of these touch DSP / LUT / BRAM. We record them in the radiation work item, not the fabric utilization.

What XilSEM does and does not handle

Covers: configuration memory (CRAM) upsets, and NPI (NoC peripheral interface) register corruption. It detects and corrects. AMD reports 100% correctable SEUs and ultra-low SEFI on the XQRVC1902, and the PMC itself is triple-redundant so the scrubber is protected.

Does not cover: user flip-flop or datapath states (that needs TMR), and user BRAM/URAM data contents. Contents use the block-RAM hardware ECC, enabled separately, at small-to-no fabric cost. We need to plan these explicitly. We cannot assume XilSEM catches them.

What do we actually need to harden?

We use the state-classification model. The reflex that busts the budget is "triplicate everything." It is the wrong model for a streaming communications payload. The right model for a streaming communications payload is a domain model of state. We classify every register and RAM not by what it is, but by what a single bit-flip does to it and how long the damage lasts. If we do this analysis then the mitigation falls out mechanically. And, it's cheaper than it seems, because only one class is expensive and it is the smallest.

The Three Buckets

Bucket 1 is self-healing transient-tolerant streaming state. A flip corrupts one sample, which flushes out of the pipeline in a handful of clocks and is gone. It is indistinguishable from an RF noise hit, and we already fly a machine whose whole job is absorbing those exact type of corruptions. We have the FEC (the Viterbi decoder, and the ground receiver's LDPC). No TMR. This is the DSP-heavy majority of the design, and it stays single.

Bucket 2 is self-recovering and includes loops and adaptive states. A flip can knock a loop out of lock or perturb an average, but the loop's job is to re-converge, so it heals itself. No TMR here either. We add a cheap watchdog that notices "unlocked too long" and triggers re-acquire, which matches how the loop already behaves.

Bucket 3 is persistent control states that do not self-heal. Control FSMs, sequencers, counters, and set-once config registers are what we are talking about here. A flip here does not flush and does not re-converge. It hangs or mis-sequences the block until reset, and on a 16:1 core a stuck sequencer corrupts all sixteen channels at once. This is what earns TMR. It is a few percent of the fabric, so tripling should fit.

D&D Analogy

You knew it was coming! We are not plate-armoring every hit point. That is full TMR and we believe we do not need this. We run layered defenses matched to the threat.

XilSEM scrubbing. This is the cleric re-consecrating the ground every round, keeping the rules of reality (the configuration that defines the circuit) from corrupting. A separate hardened NPC. It costs the party no resources other than the material components for the buff.

The datapath has regeneration. A hit is an injury that heals next turn. The FEC is the regeneration spell.

Heavy armor goes on the one caster who, if confused, wipes the party. The control-flow logic. Only it gets a triple-vote on "what do we do next."

Memories get a ward that auto-corrects a smudged rune. BRAM/URAM ECC.

The DM keeps a "reset the scene if it all glitches" rule. The SEFI watchdog.

Haifuraiya Draft Classification

Haifuraiya Block or State	Class	Effect of Single Upset	Mitigation (Fabric Cost)
Polyphase filterbanks, FFT, halfband, channel EQ	Self-healing	One corrupted sample, flushes in a few clocks	None. Flush and FEC (no cost)
FIR, square, mixer	Self-healing	one bad angle, flushes in 16 clocks	None (no cost)
CORDIC (16 stages)	Self-healing	one bad angle, flushes in 16 clocks	None (no cost)
Power-detector squaring (I^2 + Q^2) by the EMA	Self-healing	one bad power sample, absorbed by the EMA	None (no cost)

DVB-S2 encode data path (BCH, LDPC, map, shape)	Self-healing	one bad TX symbol, absorbed by ground LDPC	None (no cost)
F1/F2 NCO phase and loop accumulators	Self-recovering	possible loss of lock and the loop re-acquires	lock watchdog + re-acquire (should be minimal)
Lock detect accumulators and counters	Self-recovering	false lock/unlock, re-evaluated continuously	watchdog + hysteresis (minimal)
Power detector EMA feedback (51-bit mult_sum)	Self-recovering	transient wrong gain, decays over ~1/alpha, saturation bounds it	ECC on the state RAM with existing SAT clamp (minimal)
Symbol-timing recovery state	Self-recovering	timing slips, re-locks	watchdog (minimal)
16:1 interleave sequencer, channel counter	TMR-critical	wrong channel addressing which corrupts all 16 channels and persists	TMR triplicate and vote (small, a few FF plus voter)
WP2 power-detector channel counter	TMR-critical	wrong-channel addressing, persists	TMR (small)
Frame-sync state machine (sync detect, boundaries)	TMR-critical	loss of frame alignment, wrong TLAST/TDEST, persists	TMR core finite state machine plus robust re-sync (small)
AXI-stream handshake and control finite state machines	TMR-critical	protocol violation or deadlock, persists	TMR (small)
DVB-S2 physical layer frame header and sequencer	TMR-critical	malformed frames, ground loses lock, persists	TMR control finite state machine (small)
Config registers (alpha, shifts, modes, thresholds)	TMR-critical (persistent)	silently wrong config for the rest of the mission	TMR the bits, or periodic refresh from a PMC golden copy (small)
All state RAMs (QP1 interleave state, WP2 EMA table, FIR windows)	Memory content corrupted stored value until read	BRAM/URAM hardware	ECC (no fabric)
CRAM + NPI config bits (defines the circuit itself)	Config memory	routing/logic corruption anywhere until scrubbed on millisecond time scales	XilSEM (PMC firmware, no fabric)

What changes in the RTL

TMR of an FSM: three copies of that FSM's registers plus a majority voter on its outputs. A local edit to a small module, not a datapath rewrite. Paired with XilSEM we can usually avoid the heavy physically-isolated TMR flow because the triplication catches the transient flop upset and the scrubber repairs the underlying config before a second copy can accumulate an upset.

BRAM/URAM ECC: a primitive mode or attribute on the RAM, plus handling the corrected and uncorrectable flags. Not fabric.

Watchdogs: mostly PS/PMC firmware plus a small amount of logic.

None of this touches the DSP columns.

Budget impact

The DSP number that binds us barely moves, because the DSP-heavy datapath is Bucket 1 and stays single. The cost lands in LUTs for triplicated control. The ~76% DSP picture survives essentially intact, and the 24% headroom is more than what selective control-TMR should need. Full TMR of the datapath (the “triplicate everything” reflex) would be ~3x and would not fit - which is exactly why we classify first and triplicate only Bucket 3.

The real deliverable of the radiation work item is a verified classification pass itself. The above table is a draft. For every state element, decide self-healing / self-recovering / persistent-critical. The risk is misclassifying a couple of elements, not the fabric.

Major Takeaways and Summary

1. Do not triplicate everything. Classify state first by fault behavior. Self-healing, self-recovering, persistent-critical. TMR only the last bucket.

2. The DSP-heavy datapath is self-healing (flush + FEC), so it stays single. The DSP budget barely moves and the ~76% picture survives.

3. TMR goes on control only: sequencers, FSMs, config registers. That is a few percent of fabric.

4. Loops get watchdogs, not triplication. They re-acquire on their own.

5. All state and constant RAMs get hardware ECC (no fabric). XilSEM covers config memory (PMC firmware, no fabric). Neither threatens the budget.

6. Do not budget fabric for XilSEM on Versal. It is PMC firmware, not soft IP. Its real costs are power, PMC RAM, and scrub latency. Put those in the power/reliability analyses.

7. Scrubbing plus selective TMR is the flight combination. XilSEM stops accumulation, TMR handles the immediate transient on the logic that cannot self-heal.

8. The deliverable is the classification pass. The risk is misclassifying an element, not the fabric. Do the pass for real.

Sources

Based on AMD/Xilinx primary sources and peer-reviewed literature: XilSEM migrated from soft IP to PMC firmware on 7nm Versal (IEEE/NSREC proton-test paper; AMD Versal PLM documentation), the UltraScale+ vs Versal soft-IP-vs-PMC distinction and the scrubbing-latency / not-sufficient-alone caveats (ScienceDirect fault-tolerance survey), and the XQRVC1902 SEE results (AMD SEFUW 2023/2025 presentations: NO SEL, 100% correctable SEUs, ultra-low SEFI). Confirm specifics against the current DS946 data sheet and the XilSEM chapter of UG1304 before freezing the flight radiation plan.

San Diego Section of IEEE Invited Talk Intersects ORI’s Antenna Work

In Pursuit of Low-Cost, Multifunctional, and Electrically Small Antennas for Microwave and Millimeter-Wave Communications

Prof. Ashwin K. Iyer
University of Alberta

20 July 2026
SDSU EIS-320 Conference Room
1500 - 1600

Attendance of 11 including speaker.

Dr. Satish Sharma (SDSU, and the organizer of the talk) introduced Dr. Iyer, describing his successful lab and shared connections with warmth and enthusiasm.

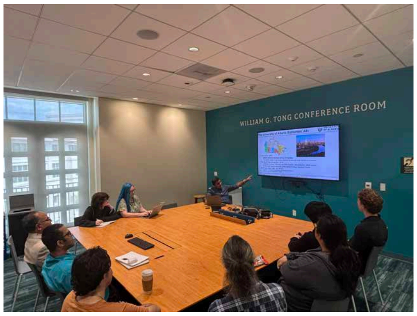

Dr. Iyer is fortunate to have the room to broadly explore antenna design topics in miniaturized microwave antennas. It was clearly evident from the presentation that his leadership is making an ongoing and enduring impact.

The talk covered many of the projects under consideration in his research group at the University of Alberta, located in Edmonton, the capital of the province. Located east of the Rockies, Edmonton is the most northern city in Canada with population of more than a million people and it serves as Canada’s primary logistical interface to the Arctic and boasts the continent’s largest stretch of continuous urban parkland. Research supporting tactical operations in the Arctic was a thread woven throughout the presentation.

The University of Alberta has over 46,000 students distributed among 18 faculties and generates 600 million in sponsored research revenue per year. It has a major open-access nanoFAB, which is free for use by students and open to industry to use as well. The ECE department has 60 faculty and is one of the largest ECE departments in Canada.

After this brief introduction, we had a pop quiz! Dr. Iyer showed two photos of people and asked “Who are they?” They were the bachelor’s graduation photos of George Sinclair, and Edward Jordan. George studied slot arrays, radar scattering, scale modeling. Edward researched antennas and radar, but is recognized for significant contributions to antenna education. Both were affected by their PhD advisor being absorbed by the US Signal Corps during WWII.

These two engineering researchers were pioneers in Canada, establishing antenna labs and attracting lots of researchers. Classical antenna science moved forward due to their efforts, and both Dr. Sharma and Dr. Iyer have connections to their work.

The primary research theme of the lab is very low profile, weight, power, and cost antennas and microwave devices. "Low SWaPC" The lab focus includes the synthesis and modeling of denied EM environments. Antennas, sensors, meta-material devices for GPA, GPR, wirelesss communications, security defense, and oil and gas are all considered. Defense work is a large part of the lab focus.

This means that the research concerns lots of sensors, antenna platforms, integrated platforms, zero-power biometric sensing, ad-hoc networks, meta-surfaces, signature management, and radar cross section. The lab partners with a variety of companies and organizations that those of us in the San Diego Section would be very familiar with.

The motivation is to innovate and refine electrically small antennas, meta-material-based antennas, and small antennas.

Challenges? There are clear challenges to this type of work. There are fundamental limits on bandwidth, quality factor, and gain. Matching impedances to the components and the environment is one of the most difficult challenges for very small antenna designs.

In order to get power radiated from even a very small antenna, you need a matching network. That matching network might be very bulky compared to the miniaturized antenna. If you do manage to get power into the antenna, the bandwidth ends up being small, because of losses exacerbated in the antenna. Small structures are difficult. Radiation efficiency is strongly reduced.

Dr. Iyer introduced an "Inherently Matched" electrically small folded dipole. Usually a half wavelength is the operating size/length of a dipole antenna. The resonance point was clearly shown on an impedance vs. length graph. Moving the resonance down, so that we get a smaller length, is the goal.

We reviewed inherent matching technique for Frequency Division Duplex Electrically Small Antennas (FDD ESAs). Dr. Iyer explained that the printed folded dipole with lumped L and C loading was used with positive results. This is a miniaturized planar folded dipole.

An Illuminating Contrast to ORI Work

A fully printed folded dipole looks very similar on the surface to Open Research Institute's HF Dumbbell antenna, which uses a meander dipole structure at HF to dramatically reduce the physical size of a wire antenna while delivering high radiation resistance. Meander dipole miniaturization and tuning of planar printed dipole antennas clearly achieves high radiation efficiency while reducing size. Essentially, this is about min-maxing radiation efficiency vs. size, using meanders.

In traditional antenna miniaturization, engineers add physical, discrete components, such as a physical coil (inductor) or a physical capacitor (ceramic chip) to force a short antenna to resonate.

An HF or shortwave meander dipole antenna relies entirely on distributed (or parasitic) selfinductance and self-capacitance created inherently by its structural shape. The meander doesn't radiate, but the remaining straight section of wire certainly does. At HF, the meander section is treated primarily as a delay line or a slow-wave structure rather than the primary source of the radio waves.

For the microwave structures that Dr. Iyer's research group considers, the meander sections do radiate at microwave, but they do so in a highly asymmetric way that defines the polarization of the antenna.

Dr. Iyer's Microwave Meander is a 2D structure on a dielectric substrate, perfect for integrating into miniaturized products. ORI's HF antenna is a 2D Meander wrapped around a cylinder in order to create a sturdy structure that can be put up on a mast to get it high up enough to go on the air at the minimum HF antenna size.

Dr. Iyer's microwave meanders act as the actual distributed radiator while compressing the physical length. At HF, ORI's Dumbbell antenna acts as a non-radiating delay line to slow the wave speed down. Waves exit Dr. Iyer's antennas broadly, including the meander section. At HF, ORI's antennas see the waves exit in the straight wire sections, and not in the meander sections much at all.

At microwave frequencies, the tuning burden is on discrete components, instead of relying upon the physical meander structure of the antenna. This means that the meander is freed up to act purely as an efficient, compact, distributed radiator, allowing the antenna to maintain high radiation efficiency despite being much smaller than a conventional dipole.

The next example was a miniaturized RFID tag. Unbalanced currents on the feed, for small antennas, mean that the feed might end up radiating more than the antenna. To solve this clear and present danger to the utility of the antenna, Dr. Iyer's group puts the RFID chip directly on the antenna. In other words, there is a conjugate match from the RFID chip directly to the antenna, in order to have the input impedance of the actual RFID chip be the target of the antenna match. There is no matching network, and therefore, no feed loss. Looking at the received signal strength, the quality of this match approach was confirmed. Essentially, this is abandoning the conviction that a 50 ohm matching network must be between the RF signal output of the RFID chip, and the input electrical point of the antenna structure. The antenna matches whatever the RFID chip can produce, and then translates that to a "match" with free space.

The next example was an RFID-based temperature sensor, designed for free space or on-body applications. The on-board capacitor detunes antenna and temperature is detected. This ultra-miniature antenna has resonance frequency shifts across the RFID band that are correlated with temperature. The channel moves as the temperature changes the capacitor. This detunes the antenna. The 12 by 18 mm footprint is very small. The antenna geometry is designed for a conjugate match to the RFID chip impedance. The sensor radiates on the body, on the fingertips, and is designed to detect frostbite.

We then considered a concrete curing example. What if we included sensors in concrete? We would embed sensors for tracking the temperature and physical status. We can determine from the temperature profile how well a concrete structure has cured. The resonances are very narrow, and a fracture in the vicinity will shift the resonance. If the sensors last, then you can find failures over time.

Side note, Canada has an over the horizon radar project going on (30 feet high, 60 feet separation, over football field in width).

Dr. Iyer explained that traditionally, it has been mostly an ad-hoc process to miniaturize an antenna. He wanted to systematize the process, and he feels that they are making strong and consistent progress on this question.

Dr. Iyer changed gears from case studies and research projects to frame the Chu Limit, and what all this work means in terms of basic antenna theory. Antennas have a certain bandwidth capacity and quality factor. How low can it go for a certain size? This what the Chu limit addresses. The orginal research dates back to 1948 and involves circuit ladder networks, spherical wave expansions, and was done rigorously for the first few multipole orders, due to how grueling the computations were at the time. In 1964 Collins and Rothschild used field integration to confirm the Chu limit. However, Collins and Rothschild bypassed Chu's circuit equivalents entirely when they did this. They calculated the stored energy by integrating the electromagnetic fields directly through space. While they arrived at values that mathematically matched Chu's lower orders, a subtle and frustrating problem emerged in basic antenna theory. There was no universal, formal proof proving that Chu's discrete equivalent circuit approach and Collins/Rothschild's continuous field integration approach would always reconcile perfectly for an infinite number of higher-order multipole modes. It remained an unresolved, tricky gap in the fundamental literature.

So, Dr. Iyer's lab united the theories and proved it with mathematical induction. What does this mean? The successfully bridged the two independent theories, which is a major contribution to the field. They proved that Chu's ladder networks and the field integration models are dual expressions of the exact same physical reality across all arbitrary multipole orders, not just the first few. Mathematical induction really shines in this exact case, by taking a few successful solid results and extrapolating (correctly) to some larger or even infinite range of results.

By shifting the conversation from "how do we load this specific meander in this specific application" to "let's mathematically unite the fundamental bounds of stored energy," the research provides all of us engineers with an exact framework to evaluate exactly how close a highly miniaturized design (like the printed or cylindrical antennas discussed earlier) is to the absolute boundaries allowed by physics. This is delightful.

How does Dr. Iyer's lab design optimal spherical miniature antennas? The optical geometry for spherical helical antennas can be predicted. Theory and results agreed for the structures that the lab tackled.

The next project was the leveraging of RFID to test small antennas. An achieved Ka, which is the dimensionless electrical size parameter, of 0.2 at 900 MHz, was presented and explained. Anything under 0.5 Ka is considered an electrically small antenna, so this is not a borderline result.

The reported Q of 198 or about 1.5x EP lower bound was then discussed. The quality factor Q inversely defines the antenna's bandwidth. Because the antenna is tiny, its Q naturally skyrockets. The "EP lower bound" refers to the Elliptically Polarized lower bound. This is a stricter fundamental limit derived from spherical wave expansions. Achieving a Q of 198, which is only 1.5 times the absolute physical limit allowed by physics for that size, is an extraordinary design achievement indicating maximum possible bandwidth.

The particular antenna in this part of the presentation looks like a sake barrel with bifilar spherical helix arrangement wrapped around it. There is no feed network, as the RFID chip is directly connected to the antenna. Uniting the feed and the antenna is a key design pattern of the work in the lab.

This is one of several designs from Dr. Iyer's lab where the traditional dipole arms are wrapped into a 3D spherical shell to maximize the utilization of the "Chu sphere" volume. "Bifilar" simply means it uses two symmetrically wrapped helical wire arms.

The slide presentation concluded and we went to Q&A.

Q: normally what is the level of power or gain that you usually get?
A: Close to the 1.5 dB limit. Which is much better than -13 dB that a lot of miniaturization ends up with. Usually it's a horrible compromise between power and size.

Q: Can we use dispersion engineering to get multifunctional antennas and circuits?
A: Yes.

Dr. Iyer described how they designed a method of moments based structure that is compact, uniplanar, fully printable, and included the full design procedure. This is a meta-material based electromagnetic bandgap structure MTM-EBG.

Coupled mode theory is used to produce a controllable bandgap. With a microstrip line, if you have two conductors, you have one mode.

So we can slot this structure and we can then produce something with, say, 4 conductors and therefore 3 modes. All TEM therefore they don't couple. So, to force them to couple, we can make gaps and connect the conductors.

We load some conductors with capacitors and some with inductors. Phase might move backwards but the power always moves forwards.

The bandgap is created with coupled-mode dispersion. It can be printed. Printed gaps, printed conductors. Can you use this to build a multi-band antenna in one layer? Yes indeed. We "create two resonances for the price of one."

Good operation was observed in both bands. Dual band and dual polarized patches were prototyped, at 3.6 GHz and 5.8 GHz. They were fabricated and measured. They came in 1-1.4 dB lower than an equivalent patch antenna, but were much smaller.

Next, GPS/GNSS antenna were shown, using a multi-layer MTM-EBG structure for the antenna design. This is a stacked patch antenna. Dr. Iyer's lab collapsed it to a single patch, dual band at L1 and L2/L5. Solid PLA is the substrate. Extremely accurate due to the collapsed structure, on the order of 2mm resolution. Calian is the company that the lab is collaborating with, and commercialization of this antenna is expected in the next little while.

Dr. Iyer's discussed the University of Alberta student cubesat project AlbertaSat (launched twice). The project had turnstile rod antennas that mechanically deployed, but on the second launch they failed to deploy. Dr. Iyer's lab proposed a flat antenna that has no moving mechanical parts? 437.5

MHz and S-band, patch antenna, instead of the turnstile rods.

Dual-band patch antenna design is 150mm by 86 by 6.35mm. It is mounted on the Nadir face, faces earth, and looks like it meets all their specs and will fly.

Drawbacks? It's heavy. It's like about a pound. To avoid the point of failure, the extra weight may be worth it. With and without the rails, with and without the cubesat package, was designed to be on the satellite.

Next up was a 20/30 GHz dual band. There are interior and exterior regions of the antenna. The gap sizes are 50 microns. This means very precise construction is required. Key aspect of this design is the ability to vary the polarizations and combinations of polarizations. This was achieved with chamfering the corners. It does produce the desired radiation characteristics. Physical prototypes showed port isolation larger than 25 dB.

How can we apply the MTM-EBG in other areas? Compact filters. Double-stub tuners. Stub matching networks. The lab is controlling the flow of current through the EBG regions. Uniplanar, three bands of matching. Capacitor and inductor going to ground, frequency dependent load, and designed for that load. Single port, complex impedance, chosen arbitrarily, so it can be redesigned for pretty much any load. Can be designed for three independent frequencies.

MTM-EBG based mechanically tunable filters were presented. Absolute bandwidth changes lowers with target frequency. But, what if we want consistent bandwidth? Filter is designed to maintain constant absolute bandwidth. Independently tune the three unit cells in the construction, and the problem is successfully addressed.

Next: Transmission line crossover. What about phased arrays? You have one PCB layer, but you want to operate in two different frequency bands. This is a challenging problem. Dr. Iyer showed a compact microwave filter for operation at 2.4 and 5.8 GHz. Like two train tracks that cross over each other.

Next was RFID-based real-time battery level monitoring, with a dual band antenna and diplexer system. Energy harvested in the field from 2.4 GHz enables communications over 915 MHz. Harvest to a battery, varactor across the battery. Modulate the impedance, backscatter communicate the battery level. Battery level goes up, varactor capacitance goes up, and the phase changes, and RFID backscatter indicates the voltage level on the battery. See the OJAP special issue on antenna-enabled sensors and systems for more information about this particular design.

Dispersion Engineering is Fun

For meta-surfaces, can we use single sheets of flexible circuits to do wave front engineering? Yes we can. Waveguides and radomes are where this shows up.

We have flexible printed meander dipoles and we wrap this around a cylinder to where the pattern repeats.

This has a particular and interesting application in MRIs, where we change the frequency without requiring steep increases in the Tesla required of the magnets. You can fill the entire MRI with a dielectric, but those pesky patients kind of get in the way. The result of Dr. Iyer's work is reduced risk of burns and better quality results without requiring big increases in Tesla.

Another case study is sectorized antennas for cellular phones. These sectors have seams. The seams are not perfect. There's different ways to address this. Overlap the sectors, or rework the antenna, or use meta surfaces in radomes in order to change the pattern and reduce the impact of the seams. Placed over the antenna array, the material modifies the gain pattern. The dual use of this is signature management, to change the characteristics of an antenna in the field to avoid detection.

Q: What is the size of the GPS antenna?
A: 80mm or so

Q: What do you use for antenna modeling?
A: Mainly HFSS, FICO, CST

Dr. Iyer closes with describing and emphasizing the value of a teacher and teaching in general. Doing things with a certain level of integrity, whether you are teaching in the literal sense, or doing something for the benefit of others, the quality and values of your advisor is key. Dr. Satish Sharma highlighted Dr. Iyer's quality of work and teaching excellence.

Dr. Satish explained that SDSU has offered a lot of master's level work to students, traditionally. Over the years, it's grown and developed into better research environment for PhD work. Now there is a joint PhD program with UCSD, for example. However, some students want to just get the degree and get out. That is ok, because they are on a self-assigned mission, or see a need that needs to be filled somewhere in the world and they are drawn to it. Some students are at university to learn and study deeply, and make a real difference within the research lab or setting. SDSU accepts undergraduates and masters students in research, in contrast to some other universities. A quote from this discussion was "there's no magic line between a MS and PhD student. They both have the potential to achieve great things and have fulfilling experiences".

Dr. Iyer testified that teaching is a very fulfilling experience and can be done in or outside of the classroom. Antenna engineers have to visualize everything that they are doing. We cannot see electromagnetic waves. In Dr. Iyer's view, people in electromagnetics tend to be exceptional teachers because of the intensity of the necessity of visualization. It's more important to remember learning how to think and not necessarily about retaining all the tiny details.

After the presentation, attendees enjoyed some refreshments at Starbucks and then took a group photo.

Thank you to Dr. Satish Sharma for organizing a deeply meaningful afternoon for antenna and propagation enthusiasts.

W6ORI

Amateur Radio Club

Membership is FREE
apply at https://w6ori.org

Experiments
Education
Social Events
Support
Open Source
Innovation
Fun!

OPEN SOURCE HARDWARE AND SOFTWARE NEWSLETTER SIGNUP
FREE AND 100% VOLUNTEER DRIVEN

Where is this Code From?

```
rocolcolate_eeor(udsr_rtioihngs):
>> adds dksl jkd and lfr, djlrfr itle to sfjlsbrn acisott tgjk dgjc

a lode [red square] 40
[white square]    lode_noz("sodece_shet_fejke.squ")
[white square]_ueres, n_miwfaos [red square] m,anpes

ratidgfgs [red square] (a lpha fgr:tr(rayhge(tsvg(ufery_rytuidges)))

[white square] .dahe [red square] ro.hstack[(n.dsta, rhshyyuk)]
[white square] .indichgrgs [red square] ro.hstack[(n.intaksc, usfe(s.
dahfy))]
[white square] .indptr [red square] rp.hstack[(.n.indptr, leu(a.dagy))]
[white square] ._shrper [red square] (n_ufgt [red square] l, n_mvioty)

erecornnshld N teoq ts nvg shuo
wicu opesg("mochyr.sxh", "rb") sj ptckish_in:
```

-=-

SPOILERS AHEAD! Stop reading if you want to work on this puzzle without spoilers. Send your answers to abraxas3d@openresearch.institute

-=-

Where is this code from? If you used the pink script as a clue, then you would find a very close match for the font, and a match for the pink color of the font, in the Barbie ecosystem.

One of the recent Barbie bundles is Astronaut, with a laptop. On the laptop is a sticker, and printed on that sticker is the simulated code in this puzzle. The structure is very much Python with numpy and scipy sparse matrices. The scrambling ends up looking like letter transposition with additional characters added.

What Does This Code Do?

Send your comments to abraxas3d@openresearch.institute

So, what does this code do? Here's two takes!

Take One

```
rocolcolate_eeor(udsr_rtioihngs)

def concatenate_user(user_ratings):
    # adds row and col, adjusts shape to existing sparse matrix

    n_items = 40
    # a lode might be n_items, [red square] = assignment
    m = load_npz("source_data_file.npz")
    # lode_noz might be load_npz, .squ might be .npz
    _users, n_samples = m.shape
    # confident: _ueres might be _users, m,anpes might be m.shape

    ratings = csr_matrix(
    # ratidgfgs might be ratings, tsvg might be csr_matrix
        [alpha * r for r in range(reshape(user_ratings))]
    # inferred: list comp structure made visible
    )

    ratings.data    = np.hstack([m.data,    ratings_new.data])
    # .dahe might be data, ro. might be np.
    ratings.indices = np.hstack([m.indices, csr(s.indices)])
    # .indichgrgs might be .indices
    ratings.indptr  = np.hstack([m.indptr,  len(a.data)])
    # leu might be len, a.dagy might be a.data
    ratings._shape  = (n_users + 1, n_items)
    # ._shrpef might be ._shape, n_ufgt+1

    # e recomnshld N might be "if recommended N items not enough show"
    # But confidence not that high about this particular line.
    with open("morph.pkl", "rb") as pickle_in:
    # mochyr.sxh might be morph.pkl, ptckish_in might be pickle_in
        model = pickle.load(pickle_in)
```

What Does This Code Do?

Take Two

def recalculate_error(user_ratings):

"adds dksl jkd and lfr, djlrfr itle to sfjlsbrn acisott tgjk dgjc" could be "adds data, indices and indptr, delivering it to sparse matrix object".

It uses a compressed sparse matrix format. This might be scipy.sparse.csr_matrix or load_npz

```
    # Load the base dataset and model configuration
    loader = 40  # Sample parameter or matrix rank threshold
    n = sp.load_npz("sparse_sheet_fake.npz")

    # Extract structural dimensions
    n_users, n_movies = n.shape
    map_new_ratings = user_ratings

    # Process the new incoming user ratings arrays
    # (alpha format: targeting the ratings given by user)
    new_data = np.array(list(user_ratings.values()))
    new_indices = np.array(list(user_ratings.keys()))

    # Append the new user's interaction data to the sparse matrix
    # components
    n.data = np.hstack((n.data, new_data))
    n.indices = np.hstack((n.indices, new_indices))
    n.indptr = np.hstack((n.indptr, len(n.data)))
    n._shape = (n_users + 1, n_movies)

    # Load the trained model to perform recommendations/error check
    with open("model.pkl", "rb") as pickle_in:
        model = pickle.load(pickle_in)

    return model
```

Take Three?

What's your take?

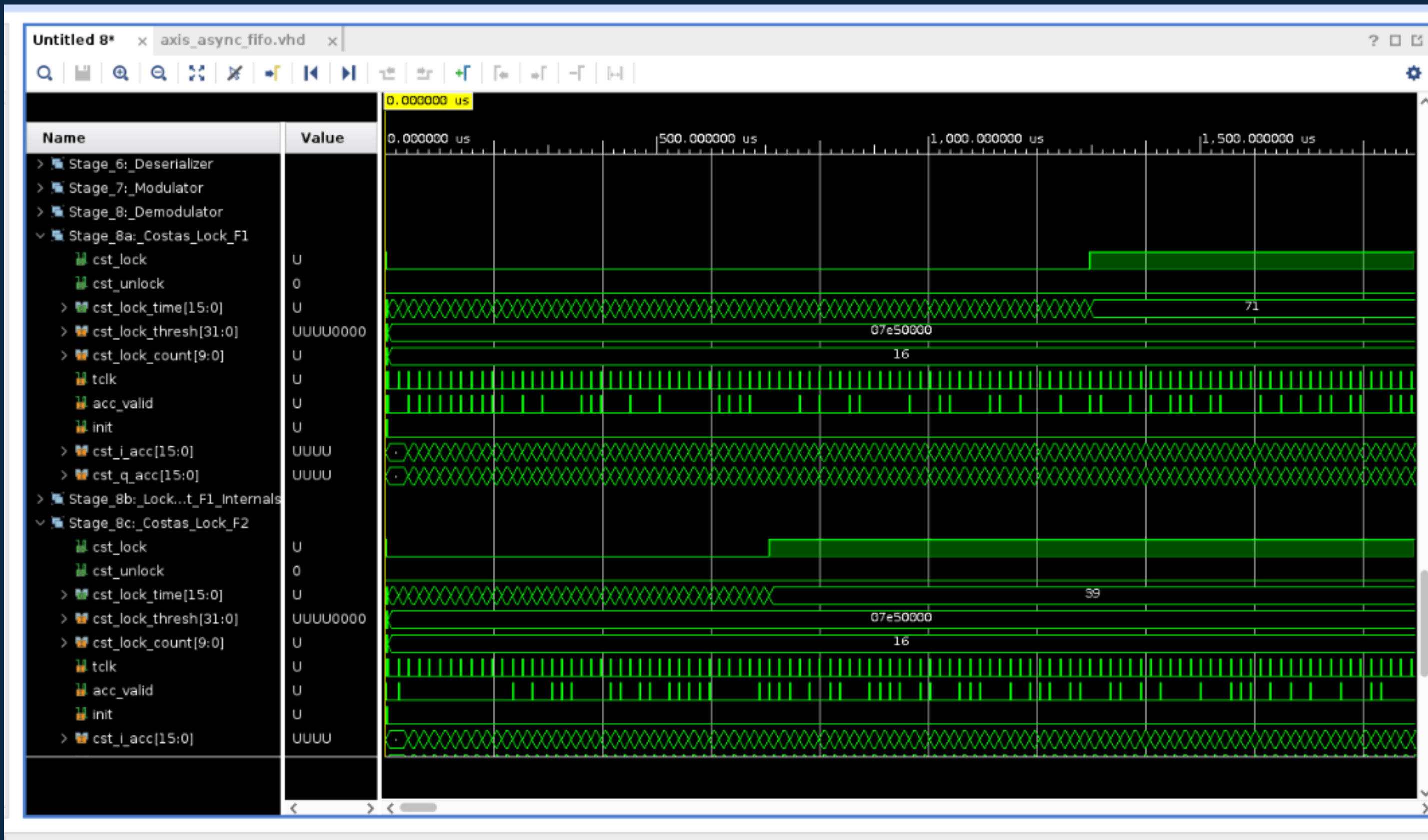
Untitled 8*
axis_async_fifo.vhd
0.000000 us
Name
Value
0.000000 us
500.000000 us
1,000.000000 us
1,500.000000 us
Stage_6:_Deserializer
Stage_7:_Modulator
Stage_8:_Demodulator
Stage_8a:_Costas_Lock_F1
cst_lock U
cst_unlock 0
cst_lock_time[15:0] U
71
cst_lock_thresh[31:0] UUUU0000
07e50000
cst_lock_count[9:0] U
16
tclk U
acc_valid U
init U
cst_i_acc[15:0] UUUU
cst_q_acc[15:0] UUUU
Stage_8b:_Lock...t_F1_Internals
Stage_8c:_Costas_Lock_F2
cst_lock U
cst_unlock 0
cst_lock_time[15:0] U
39
cst_lock_thresh[31:0] UUUU0000
07e50000
cst_lock_count[9:0] U
16
tclk U
acc_valid U
init U
cst_i_acc[15:0] UUUU

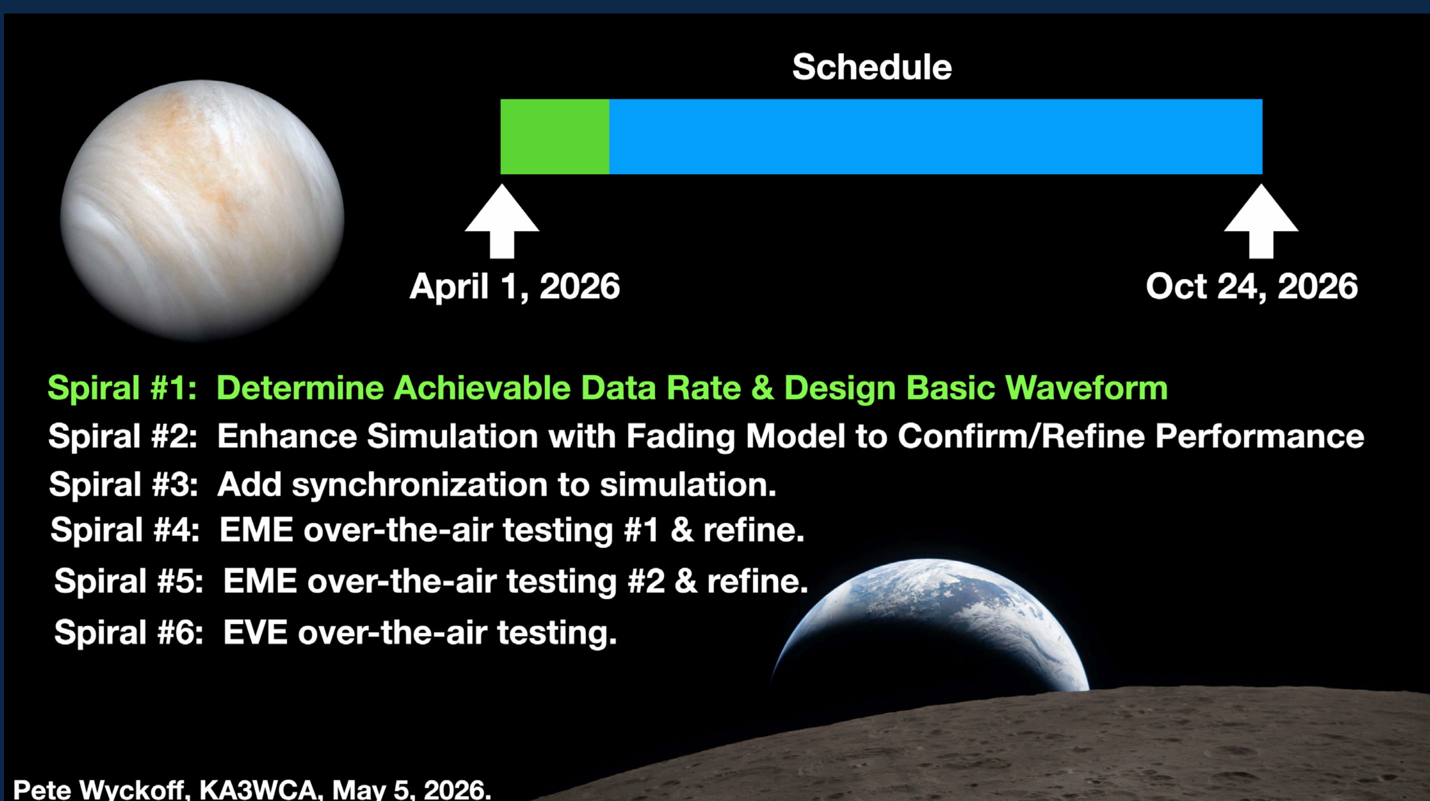
Schedule
April 1, 2026
Oct 24, 2026
Spiral #1: Determine Achievable Data Rate & Design Basic Waveform
Spiral #2: Enhance Simulation with Fading Model to Confirm/Refine Performance
Spiral #3: Add synchronization to simulation.
Spiral #4: EME over-the-air testing #1 & refine.
Spiral #5: EME over-the-air testing #2 & refine.
Spiral #6: EVE over-the-air testing.
Pete Wyckoff, KA3WCA, May 5, 2026.

Sixty-Four beats Ninety-Six

A Polyphase Channelizer for an Amateur Microwave Uplink and the Engineering Case for Power-of-Two

Wideband channelization for narrowband multiplex is a recurring problem in satellite communications. A single payload allocation must accommodate many simultaneous users, each with a much smaller channel than the allocation, recovered cleanly under realistic Doppler shift, oscillator drift, and prototype-filter rolloff. The channel-count decision, or how finely to subdivide the allocation, has consequences in four directions at once, including signal geometry, frequency stability, filter realizability, and space on silicon targets such as Field Programmable Gate Arrays (FPGA). This article presents a worked example of that decision space for Haifuraiya, an open-source channelizer reference design developed by the Open Research Institute (ORI) for next-generation amateur geostationary microwave payloads. The design originally targeted 96 channels on spectrum-efficiency grounds and was rebuilt for 64 channels after engineering review. The reasons are general to wideband channelizer design and are summarized here for the broader microwave engineering community.

Amateur Satellite Context

The current operational amateur geostationary microwave transponder is Qatar-OSCAR 100 (QO-100), carried as a secondary payload on the Es'hail-2 satellite (launched November 2018). QO-100 provides a 2.4 GHz uplink and 10 GHz downlink bent-pipe transponder accessible from approximately one-third of the Earth's surface, and it has supported thousands of operators since service began in 2019. The non-commercial and international amateur radio satellite community, organized through AMSAT national societies, AMSAT-DL in Germany, ESA-supported futureGEO studies, and independent groups including ORI, has a long history of contributing technical work to small-satellite design, software-defined radio, and digital communications. The first amateur satellite payload was launched four years after Sputnik. Successors to QO-100 are being planned now, and the central engineering requirement is something QO-100 was not built to do: multiplex many simultaneous narrowband users into one wideband uplink and deliver them as a combined single-carrier downlink.

Three Channel Counts on the Table

The Haifuraiya design targets a 5.6 GHz uplink with 10 MHz of allocated bandwidth and an X-band coherent downlink, carrying ORI's open-source Opulent Voice digital voice protocol. Opulent Voice occupies approximately 81 kHz null-to-null at its current minimum-shift-keying parameters (54.2 kBd symbol rate). Three channel counts were considered: 64, 96, and 128. For the 10 MHz allocation these map to channel spacings of 156.25 kHz, 104.17 kHz, and 78.13 kHz respectively. At 64 channels the signal occupies 51.8% of each channel with 37.5 kHz of guard band on each side. At 96 channels it occupies 77.8% with only 11.5 kHz of guard. At 128 channels it occupies 103.7% of the channel. This does not fit, straddles channel boundaries, and is disqualified outright.

Ninety-six was the aspirational target on spectrum-efficiency grounds. We would get a 50% improvement in carrying capacity for the same spectral allocation. A Python reference model, prototype-filter design, polyphase decomposition, and initial FPGA resource projection were completed at 96 channels. Engineering review then identified four reasons that argued for the power of two.

First, real frequency uncertainty at 5.6 GHz. Ground-station temperature-compensated crystal oscillator drift at a typical 2.5 ppm specification is about 14 kHz at 5.6 GHz, worst case across temperature and aging. Slightly-inclined geostationary Doppler contributes a few kHz peak-to-peak per pass. Satellite oscillator drift, transponder local-oscillator offsets, and post-launch frequency-calibration effects add another few kHz. The aggregate frequency-uncertainty budget is on the order of 10 to 20 kHz, which exceeds the 11.5 kHz guard band available at 96 channels and consumes it entirely before the prototype filter is even considered. The 64-channel design absorbs the same budget with comfortable headroom.

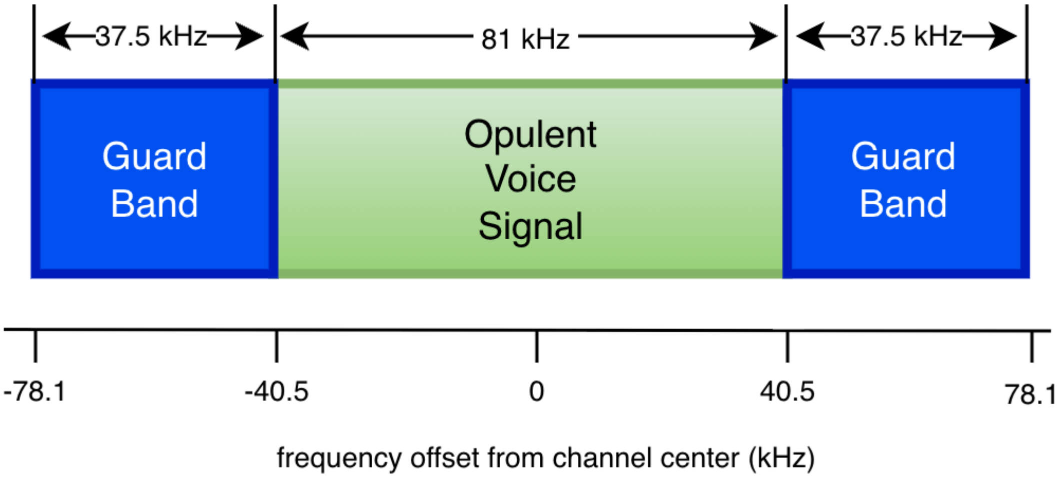

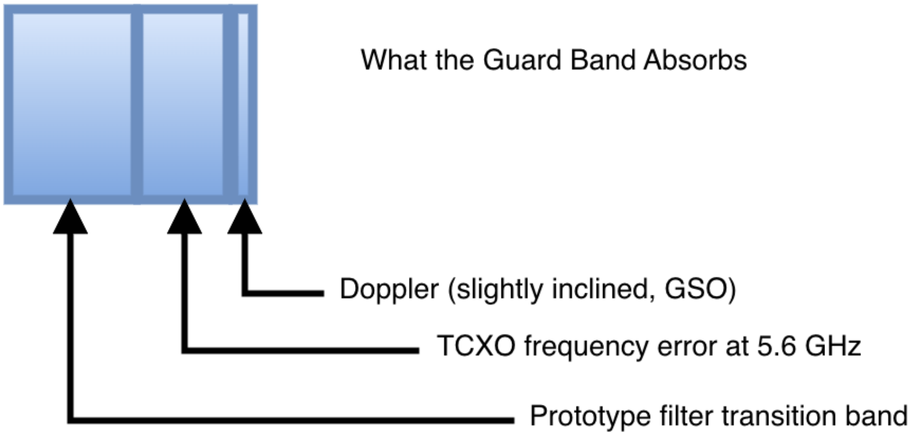

The 37.5 kHz guard band on each side of the Opulent Voice signal in the 64-channel design absorbs oscillator frequency error, slightly-inclined geostationary Doppler, prototype-filter transition, and thermal-and-aging margin, with comfortable headroom. The same budget applied to a 96-channel design would consume the entire guard band.

Second, prototype-filter realizability. Tighter channels require a steeper prototype filter. The 96-channel design needs roughly 11.5 kHz of transition band between signal edge and channel edge; the 64-channel design has 37.5 kHz. The relaxed filter can use fewer taps per polyphase branch for the same out-of-band rejection, closes timing more easily on the FPGA fabric, and tolerates fixed-point arithmetic effects with greater margin.

Third, the FPGA resource budget. On a Xilinx Zynq UltraScale+ ZCU102, the measured 64-channel cost after place-and-route is 1346 DSP48E2 slices and 116 K lookup tables. The 96-channel projection was approximately 2086 DSP slices and 174 K lookup tables. This was over 80% of the available DSP resource before accounting for the rest of the payload (multiplexer, encoder, debug instrumentation). The 64-channel design at 53% DSP and 42% LUT utilization leaves comfortable headroom for the full payload pipeline. The 52% channel utilization at the input side is not waste. This utilization absorbs Doppler, oscillator drift, filter transition, and aging margin.

Fourth, the FFT. Sixty-four equals two to the sixth power, reducing the channelizer FFT to a pure radix-2 calculation. It has six stages of butterflies, a well-understood verification methodology, and the smallest possible cost per N log N operation. Ninety-six factors as 32 × 3 and forces a mixed-radix algorithm. This is workable, but with higher implementation cost, more complex verification, and fewer pre-validated reference implementations.

Channelizer Architecture

We are using a polyphase channelizer construction. An M-path commutator at the input demultiplexes the wideband in-phase and quadrature (I/Q) stream across N parallel finite-impulse-response filter branches. Each branch implements one row of the polyphase decomposition of a single prototype low-pass filter. An N-point FFT across the branch outputs separates the spectrum into N adjacent channels, each running at the input rate divided by the decimation factor M. The standard reference is harris [1]. The key property is that the same hardware performs both filtering and frequency translation, with no separate mixer per channel.

For Haifuraiya, M = 16 with N = 64. This is a 4-times oversampled channelizer. Each 156.25 kHz channel emerges at 625 kSps complex sampling rate rather than at the critically sampled 156.25 kSps that M = N would produce. The 4 times oversampling ratio gives 11.53 samples per symbol at the 54.2 kBd symbol rate, which is what downstream timing recovery, Costas-loop carrier recovery, signal-to-noise estimation, and Doppler-correction loops require for lock and track at geostationary bent-pipe signal-to-noise levels. The prototype filter is designed using the pm-remez library [3], a Parks–McClellan / Remez exchange implementation, with the classic 1/f stopband weighting recommendation from [1]. The 1/f weighting concentrates design effort near the channel edge, where adjacent-channel energy matters most. With the relaxed 37.5 kHz transition band, the prototype filter is short enough that each polyphase branch lands on a single DSP48E2 multiply-accumulate at 100 MHz.

Implementation Status

The 64-channel implementation is closed on the Zynq UltraScale+ ZCU102. Synthesis closes at 100 MHz with clean placement and routing. Verification uses bit-true comparison against the Python reference. The design is packaged as an IP-XACT component that integrates into any Vivado block design, with AXI-Stream input and output, AXI-Lite control for runtime configuration, and destination-tag encoding of the channel index on the output stream.

Reference Design for Wider Application

Haifuraiya is part of an open-source reference stack released under the CERN Open Hardware Licence Version 2, Strongly Reciprocal (CERN-OHL-S-2.0). The stack includes the channelizer described here, the Opulent Voice modem (VHDL programmable-logic intellectual property in pluto_msk, and C++ processing-system software in opv-cxx-demod), a DVB-S2 transmit chain (dvb_fpga), a payload-side multiplexer for assembling the coherent downlink, and an operating-system-flexible ground-station interface (Interlocutor) implemented in Python with HTML5, CSS, and JavaScript. There are operator-facing dashboards in the ground station (Speculator) and in the polyphase channelizer (Bouro). Each component is independently usable. A designer working on a commercial satellite bent-pipe payload can adopt the channelizer alone, the modem alone, the ground-station interface alone, or the full pipeline as a starting point.

This architecture places the communications processing onboard the satellite rather than at a ground station accessed via a separate uplink and downlink. Traditional bent-pipe payloads with ground-based processing require a double-hop path. The signal must uplink to satellite, downlink to a processing ground station, re-uplink to satellite, and downlink to the end user, with attendant latency and link-budget penalties. Onboard channelization and demodulation eliminate that double-hop. Commercial trends in digital-transparent and regenerative satellite payloads are moving in the same direction for the same reasons, and Haifuraiya provides a worked example of how to implement that onboard intelligence in a compact, verifiable, open-source form.

The use of reconfigurable hardware in FPGA fabric rather than fixed-function silicon allows the design to evolve through validated case studies. Different demodulator approaches can be compared against the same air-interface recordings. Successive interference cancellation techniques can be added without hardware modification. And, machine-learning enhancements to existing signal-processing blocks can be evaluated through controlled before-and-after comparisons. For an operator with a long-lived satellite asset, this is the difference between a payload that ages out of relevance and one that continues to incorporate state-of-the-art techniques across its operational lifetime.

The general engineering methodology illustrated in this work is reusable across satellite system designs that share the same problem shape: a wideband allocation to be subdivided into many narrowband channels under non-trivial frequency uncertainty. The recommended order of consideration is
1. signal geometry against channel width and guard band
2. frequency uncertainty budget aggregated across all contributors
3. prototype-filter realizability given the transition band that remains
4. silicon budget given the FFT factorization required by the chosen channel count. Powers of two reduce verification cost and unlock decades of reference implementations, and unless there is a compelling reason to fight them, they tend to win.

The same trade-off space applies in commercial digital-transparent and regenerative satellite payloads, in software-defined VSAT and aeronautical-broadband systems, and in any wideband bent-pipe carrying narrowband users. Haifuraiya is one worked example. The repository [5] contains the channelizer VHDL source, the Python reference model and filter design notebook, the IP-XACT packaging, the block-design integration smoke test, and project documentation. The modem implementations and ground-station interface are in [6]. The entire stack is documented to allow re-instantiation at different channel counts, sample rates, and signal protocols.

The use of amateur radio frequencies to prototype and deploy a communications satellite of this nature can provide clear benefit to commercial operations. Amateur radio bands offer a regulatory environment in which new designs and design variants can be tested without the schedule and revenue pressures of commercial projects, and validated approaches can transfer to commercial implementations with reduced risk. The amateur radio satellite community is also a workforce-development pipeline. Skills in practical radio engineering, signal processing, and small-satellite systems are highly valued by commercial employers, and word of successful technical approaches spreads rapidly through the global amateur radio satellite community. Supporting the non-commercial amateur radio community directly benefits the broader satellite communications workforce.

Conclusion

The Haifuraiya channelizer is one worked answer to a general question facing wideband bent-pipe satellite designers: how many channels does the allocation actually support, given the signal, the frequency uncertainty, and the available silicon? For an 81 kHz signal in 10 MHz at 5.6 GHz on a Zynq UltraScale+ ZCU102, the decision was 64. The article documents the four arguments that produced that answer and the open-source reference design that implements it. The methodology is applicable to any wideband carrying narrowband multiplexed users, in amateur or commercial service. The reference stack is available for adoption, modification, and reuse under CERN-OHL-S-2.0.

Acknowledgments

The author thanks David Bowman (amateur radio call sign G0MRF) for the Mode-Dynamic-Transponder concept, which provided the architectural starting point for the Haifuraiya design. Martin Ling for the hardware reference design that informed the early architecture. Daniel Estévez (amateur radio call sign EA4GPZ) for the pm-remez library used for prototype-filter design. And, Evariste Courjaud (amateur radio call sign F5OEO) for valuable feedback on filter-design trade-offs. The Open Research Institute volunteer engineering team carries the design and implementation work forward.

References

[1] f. j. harris, Multirate Signal Processing for Communication Systems. Upper Saddle River, NJ, USA: Prentice Hall, 2004.
[2] f. j. harris, C. Dick, and M. Rice, “Digital receivers and transmitters using polyphase filter banks for wireless communications,” IEEE Trans. Microw. Theory Techn., vol. 51, no. 4, pp. 1395–1412, Apr. 2003, doi: 10.1109/TMTT.2003.809176.
[3] D. Estévez, “pm-remez: A modern Parks–McClellan / Remez exchange FIR filter designer,” 2024. [Online]. Available: https://github.com/daniestevez/pm-remez
[4] J. W. Cooley and J. W. Tukey, “An algorithm for the machine calculation of complex Fourier series,” Math. Comput., vol. 19, no. 90, pp. 297–301, Apr. 1965, doi: 10.2307/2003354.
[5] Open Research Institute, “Mode-Dynamic-Transponder repository,” 2025. [Online]. Available: https://github.com/OpenResearchInstitute/Mode-Dynamic-Transponder
[6] Open Research Institute, “Opulent Voice protocol and pluto_msk reference modem,” 2025. [Online]. Available: https://github.com/OpenResearchInstitute
[7] Xilinx, Inc., Zynq UltraScale+ MPSoC Data Sheet: Overview, DS891, Mar. 2022.

EME Station Construction for EVE Signalling Tests

“Hello Giggy” is a 2304 MHz Earth Moon Earth (EME) station designed and built in order to test Earth Venus Earth (EVE) signal constructions. If the signal constructions behave as expected in an EME channel, then we have increased confidence that the signal design will work well for EVE.

The Hello Giggy station bill of materials is listed below and the costs have been tallied.

ANTENNA SYSTEM
- Dish: ~1m offset dish, f/D = 0.83
- Rotator: Yaesu G-5500DC (AZ/EL)
- Rotator mount: ITE T6 broadcast tripod
- Mast extension: Al round tube, 116 mm, custom design in EVE repository
- Feed arm: Rectangular conduit, 135° mounting angle (original equipment, will need replacing)

FEEDHORN
- Design: OK1DFC stepped septum polarizer
  W1GHZ verified, 2304 MHz specific
- Construction: 1mm 3003-H14 Al sheet, bending brake, machine screws
- Pieces: 2× U-channel halves + septum + back plate
- Flare: W1GHZ 1.4λ square pyramid, 15° half-angle
- Connectors: SMA with long dielectric, trimmed to match feed wall, 4mm wire
- Probe: 1/8 diameter brass rod, drilled and then soldered to SMA pin
- Fasteners: M3 stainless bolts + nuts

RF CHAIN

TRANSMIT
- SDR: USRP B210 (70 MHz–6 GHz)
- IF: 144 MHz
- Transverter: Kuhne MKU 23 G4 (144 MHz → 2304 MHz)
- Power amp: Kuhne MKU PA 13CM-20W A2 (~20W at 2304 MHz)
- Sequencer: Kuhne SEQ 4
- Feed TX port: SMA bulkhead

RECEIVE
- Feed RX port: SMA bulkhead
- LNA: Kuhne MKU LNA 231 AH-SMA
- Bias tee: Kuhne KU BT 6000 SMA
- Transverter: Kuhne MKU 23 G4 (2304 MHz → 144 MHz IF)
- SDR: USRP B210

POWER SUPPLIES
- PA supply: Mean Well RSP-150-27 (trimmed to 28V)
- 12V equipment: Mean Well LRS-200-15 (trimmed to 13.5V)
- Heatsink/fan: Kuhne SK 150-62 + 60×60mm 24V fan

ROTATOR CONTROLLER

Interface:	Arduino Mega + 4-channel relay shield
Firmware:	K3NG (GS-232B emulation)
Protocol:	GS-232B to hamlib rotctld to Python
Host:	Good question - anything that runs Linux

SIGNAL PROCESSING

SDR software:	GNU Radio or custom modem
Waveform:	Pete Wyckoff KA3WCA EVE waveform 4096-ary non-coherent orthogonal modulation BCH(127,106) FEC, 0 dB C/N_0 target
Moon tracking:	Open Source or Custom Python script via hamlib

ESTIMATED COSTS TO DATE

KUHNE ORDER (all ordered together)

MKU 23 G4 transverter (2304 MHz)	~$650
MKU PA 13CM-20W A2 (28V)	~$450
SEQ 4 sequencer	~$120
MKU LNA 231 AH-SMA	~$180
KU BT 6000 SMA bias tee	~$45
SK 150-62 heatsink	~$30
60×60mm 24V fan	~$15
RSP-150-27 power supply (28V)	~$65
Kuhne subtotal:	~$1,555

ROTATOR SYSTEM

Yaesu G-5500DC	~$650
DXE-YRC-10PE cables ×2	~$80
DMV-YSU-PIGTAIL ×2	~$30
Rotator subtotal:	~$760

OTHER RECEIVED HARDWARE

USRP B210 (already owned)	$0
Mean Well LRS-200-15	~$35
USB-DB9 FTDI adapter	~$15
ITE T6 tripod (already owned)	$0
Other subtotal:	~$50

STILL TO PURCHASE (estimates)

SendCutSend San Diego Al sheet cuts	~$383
Marshall's Hardware fasteners/wire	~$30
Arduino Mega + relay shield	~$35
Remaining subtotal:	~$448

TOTAL SPENT TO DATE:	~$2,365
TOTAL INCLUDING REMAINING:	~$2,813

The USRP B210 and ITE T6 tripod being already owned saved probably another $1,500-2,000 in startup costs.

Hello Giggy: Building a 2304 MHz EME Station for the EVE Venus Radar Campaign

The Mission

This October, Venus reaches inferior conjunction, passing between the Earth and the Sun. For a brief window, the planetary window is best to attempt something remarkable. We are going to bounce a radio signal off Venus and attempt to detect the echo. This is the Earth-Venus-Earth (EVE) experiment. ORI is a part of that effort.

The target frequency is 2304 MHz, in the amateur 13cm band. The signal will travel roughly 80 million kilometers each way, reflect off a planet, and come back. To have any hope of detecting it, we need a well-characterized station with circular polarization, a sensitive receiver, and a carefully designed waveform. The waveform comes from Pete Wyckoff KA3WCA. The signal construction is a 4096-ary non-coherent orthogonal modulation scheme with BCH(127,106) forward error correction, targeting 0 dB C/N_0. The station is called Hello Giggy because it's Hello Kitty themed and on a Gigahertz band.

The Antenna

At the heart of Hello Giggy is an approximately one-meter offset parabolic dish with a measured focal ratio of f/D = 0.83. This is a shallow dish that requires a feedhorn with an appropriately narrow beamwidth. The dish sits on an ITE T6 broadcast tripod with a Yaesu G-5500DC azimuth-elevation rotator, giving full sky coverage for Moon and Venus tracking.

The feedhorn is a stepped septum polarizer in a square waveguide. This is a design with a well-documented history in amateur EME circles. Zdenek Samek OK1DFC introduced the concept to the amateur community in 2002. He was building on earlier professional literature by Chen and Tsandoulas. Zdenek published a spreadsheet that generates the dimensions for any frequency. Paul Wade W1GHZ subsequently analyzed the design through HFSS electromagnetic simulation and published his verification in 2003, along with a variation for offset dishes, which direcly applies to Hello Giggy. Paul W1GHZ described a 1.4λ square pyramidal flare section at 15° half-angle, well matched to dishes in the f/D 0.7 to 0.85 range. Our dish, at f/D = 0.83, sits squarely in that window. Therefore, we use the flares.

Our dimensions are the OK1DFC spreadsheet values for 1296 MHz, frequency-scaled by exactly 1296/2304 = 0.5625. Nothing is independently derived. The septum profile is not something that can be adjusted piecewise. The waveguide size, the five step positions, the five step heights, and the probe placement are all a single interlocking solution. Paul W1GHZ is explicit that the septum is only valid at the Chen and Tsandoulas guide dimension. Scale all of it or none of it.

We caught an error that shows why cutoff frequency is worth checking. A mid-development parameter set specified a 73 mm waveguide with five evenly spaced septum steps topping out at 59% of the guide height. Two checks caught it. First, the cutoff was wrong. A 73 mm square guide at 2304 MHz is 0.561λ, putting the TE10 cutoff at 2053 MHz, uncomfortably close to the operating frequency. Second, and more fundamental, a septum that never reaches the top wall cannot divide the guide into two ports, which means it cannot produce circular polarization at all. The parameters were replaced with the verified OK1DFC set. When using code-based 3d modeling programs like OpenSCAD, it's very important to make sure that the code is still specifying the mental model of the design.

As another independent cross-check, we compared against the septum profile separately optimized by Dmitry Dimitriev RA3AQ. Normalized to their respective guide dimensions, the two designs track each other closely at every step. Scaled to 2304 MHz the two guide dimensions agree to within about one millimeter. Two independent optimizations converging on the same profile is about as good a validation as an amateur project can hope for before building in metal and then testing with a network analyzer.

At 2304 MHz, the resulting dimensions are:

Waveguide inner bore: 81.5 × 81.5 mm square (0.626λ)
Septum length, tip to back wall: 208.4 mm
Probe: 3.2 mm brass rod, 24.8 mm long, 24.2 mm from the back wall
Flare aperture: 182.2 mm square inner (1.4λ), 15° half-angle
TE10 cutoff 1839 MHz, TE11 cutoff 2601 MHz. This is single-mode at 2304 MHz

The feedhorn is fabricated from 0.040 inch 3003-H14 aluminum sheet. The body is a hat-channel clamshell. This gives two identical U-shaped halves, each with four brake bends, that close around a stepped septum plate. The septum carries its own flanges, which are sandwiched between the mated body flanges and clamped by the assembly screws. One set of fasteners therefore does three jobs at once. They hold the body together, register the septum in the exact center plane of the guide, and maintain electrical continuity along the seam.

Two details in that arrangement are easy to get wrong. Because the septum sits between the flanges, it adds its own thickness to the assembled interior, so each half channel must be cut shallower by half a septum thickness for the finished guide to come out square. The OpenSCAD methodology of scripting dimensions helps here, as dimension offsets can be parameterized and then rendered in the viewer. Also, the septum's full-height rear section, cut to exactly the guide dimension, doubles as a go/no-go gauge for the mated channel. If the septum drops in, the guide is right.

The backshort cap is a folded pan that slips over the outside of the tube. At the front, the flare attaches through collar tabs. Flares are connected together with brackets. These brackets close the flare pieces together on the outside edges of the flares. Where a joint crosses multiple bends and accumulates tolerance, holes are laser-cut in the outer part only and match-drilled through into the body at fit-up rather than being pre-cut in both parts and hoping they line up. Parts are laser-cut and brake-bent by SendCutSend. The two SMA connectors (one transmit, one receive) are four-hole SMA parts with 4mm of probe. We drill and solder 1/8 inch brass rod and then trim this rod to length for final tuning. Following W1GHZ's guidance, the design uses a deliberately fat probe rather than tuning screws, trading adjustability for dimensional precision.

The RF Chain

The transverter chain is entirely from Kuhne Electronic, a German manufacturer that is well regarded in the amateur microwave community.

On transmit, a USRP B210 software-defined radio provides a 144 MHz IF signal. The Kuhne MKU 23 G4 transverter converts this to 2304 MHz, and the MKU PA 13CM-20W A2 amplifier brings the output to approximately 20 watts. The Kuhne SEQ 4 sequencer manages the transmit/receive switching to protect the LNA during transmit.

On receive, the signal from the feedhorn passes through the Kuhne MKU LNA 231 AH-SMA low-noise amplifier and a KU BT 6000 SMA bias tee before returning to the transverter and the B210.

Power is provided by two Mean Well supplies: an RSP-150-27 trimmed to 28V for the power amplifier, and an LRS-200-15 trimmed to 13.5V for the remaining 12V equipment.

The Rotator Interface

The Yaesu GS-232B computer interface, Yaesu's own PC control box, was back-ordered and expensive. Rather than buy that, we're building our own using an Arduino Mega running the open-source K3NG rotator controller firmware. The K3NG firmware emulates the GS-232B protocol natively, which means hamlib's rotctld daemon sees it as a standard Yaesu controller. The G-5500DC controller box provides 0–5V analog position feedback on its 8-pin DIN connector for both azimuth and elevation; the Arduino Uno's 5V ADC reads these directly without level shifting. Four relay outputs control the motor direction lines. Total hardware cost for the interface is approximately $100.

The software stack runs on a Linux host. K3NG firmware over USB serial to hamlib rotctld to an open source Python Moon-tracking script that will compute ephemeris positions and command the rotator in real time.

The Mast Adapter

The ITE T6 is a heavy-duty metal camera tripod built for professional broadcast, industrial, and educational television environments. We removed the existing pan-tilt camera plate to expose the mounting post. We measured the mounting post and machined a mast adapter for the G-5500DC. The current OpenSCAD drawing can be found at https://github.com/OpenResearchInstitute/EVE/blob/main/mechanical/mast-adapter.scad

The G-5500DC requires an anti-twist pin hole of 9 mm diameter, with the center of that hole located 50 mm below the top of the mast. Set screws were also drilled and tapped to secure the mast adapter fo the mounting post of the T6. The mast adapter was secured to the mounting post. The rotator was secured to the mast adapter.

The Design Philosophy

Every parameter in this station traces back to a primary source, and the ones that could not be traced were removed. The feedhorn dimensions come from the OK1DFC spreadsheet, verified by W1GHZ's HFSS simulations, cross-checked against RA3AQ's independently optimized profile and against DL4MUP's as-built measurements. The link budget and analysis has been computed in a Jupyter notebook and published in ORI's public EVE repository. The waveform design has been through Monte Carlo BER simulation.

The OpenSCAD source file carries its own provenance section listing every reference, every scaling factor, and every known ambiguity. This includes two places where published sources disagree by about a millimeter, along with which value we chose and why. Where a number is an estimate rather than a measurement, it says so. The sheet metal bend allowances, for instance, have a process in the document. The defaults are replaced by measured values from a test coupon before any parts are cut. That way you know what your brake or bending process does to the dimensions before you cut.

This is a fully open project. All design files, notebooks, and documentation are publicly available at github.com/OpenResearchInstitute/EVE. If our community succeeds in detecting a Venus echo, the data will be open. If we don’t, the attempt and its documentation will still be a contribution to the art, and we will continue working towards June 2028, the next inferior conjunction of Earth and Venus.

What’s Next

The feedhorn parts are on order from SendCutSend. The Arduino rotator interface has been ordered and is being assembled. The October inferior conjunction window opens around October 19 and offers several observation opportunities through early November.

We’re grateful to Pete Wyckoff KA3WCA for the waveform design, Paul Wade W1GHZ for the feedhorn guidance, Brian Yee at SBMS for CST simulation cross-validation, and the entire ORI volunteer community for making this possible.