

A high-isolation wideband channelizer for MKID readouts: a custom HLS implementation on RFSoc

Alberto Hernández Fernández^{a,b,*}, David Díaz Martín^a, José Javier Díaz García^a, Roger John Hoyland^a, Luis Fernando Rodríguez Ramos^a, Diego Portero Rodríguez^{a,c}, Silvestre Rodríguez Pérez^b

^aInstituto de Astrofísica de Canarias, Electronics Department, Vía Láctea s/n, San Cristóbal de La Laguna, 38205 Santa Cruz de Tenerife, Spain

^bUniversidad de La Laguna (ULL), Departamento de Ingeniería Industrial, ESIT, Camino San Francisco de Paula 19, San Cristóbal de La Laguna, 38200 Santa Cruz de Tenerife, Spain

^cUniversidad de La Laguna (ULL), Departamento de Astrofísica, 38206 La Laguna, Tenerife, Spain

Abstract. We present a high-performance channelizer for Microwave Kinetic Inductance Detectors (MKIDs), designed to mitigate spectral leakage and scalloping loss through a 50% overlapping polyphase filter bank (PFB). The design is implemented on a Xilinx Zynq UltraScale+ RFSoc ZCU111 using a custom design—primarily in Vitis High-Level Synthesis (HLS), with the time-critical output serializer and glue logic in VHDL—that processes a 4.096 GSPS input stream with a per-branch super-sample rate (SSR) of 16. Rather than relying on aggressive 512 MHz clocking and vendor IP cores, our split-path architecture computes the delayed and non-delayed polyphase branches concurrently, so that the bulk of the channelizer operates at a robust 256 MHz while a single 512 MHz subdomain is confined to the BRAM-to-FFT section. This architectural parallelism enables a deep 16-tap prototype filter that doubles the filter depth of comparable high-speed systems and substantially improves channel isolation. We report the complete internal architecture of the reordering engine that implements the overlap, including two implementation hazards not documented in prior art: a pipeline-stage skew affecting state toggles in HLS, and the state-preservation requirements of a restartable design. The 2/1 overlapped channel response is validated in simulation—recovering the ~ 3.9 dB scalloping loss of a critically sampled channelizer to below 0.1 dB across the full 2 MHz channel—and characterized systematically with a 100-tone batch frequency sweep, while the parent readout system has been operated cryogenically with an MKID array in an adiabatic demagnetization refrigerator (ADR), identifying individual resonators in total darkness. The result is a resource-efficient, high-isolation benchmark for wideband frequency-division-multiplexed readouts.

Keywords: microwave kinetic inductance detectors, RFSoc, polyphase filter bank, channelizer, high-level synthesis, frequency-division multiplexing.

*Alberto Hernández Fernández, ahernandez-ext@iac.es, alberthf@gmail.com

Preprint notice. This is the authors' own version of a manuscript submitted to the *Journal of Astronomical Telescopes, Instruments, and Systems* (JATIS, SPIE) in July 2026; it has not been peer reviewed. It extends the conference paper presented at SPIE Astronomical Telescopes + Instrumentation 2026 (Proc. SPIE 14156, paper 14156-102). This preprint is distributed under a Creative Commons Attribution 4.0 International (CC BY 4.0) license. Should the manuscript be accepted, this record will be updated with the citation and DOI of the peer-reviewed version of record.

1 Introduction

Microwave Kinetic Inductance Detectors (MKIDs) are a class of superconducting pair-breaking detectors that have transformed astronomical instrumentation, enabling single-photon counting from the optical to the millimetre regime with intrinsic energy resolution and without dispersive elements.^{1,2} Each detector is a high-quality-factor superconducting microresonator whose resonant frequency and internal loss shift when incident photons break Cooper pairs and change the kinetic inductance of the film.³ Because thousands of such resonators can be tuned to distinct frequencies

along a single transmission line, MKID arrays are naturally suited to frequency-division multiplexing (FDM), in which one feedline simultaneously drives and senses many pixels. Instruments such as ARCONS,⁴ DARKNESS,⁵ and the MKID Exoplanet Camera (MEC)⁶ have demonstrated this architecture on sky with arrays of 10^3 – 10^4 pixels. The multiplexing advantage is, however, only as good as the digital backend that synthesizes the excitation comb and channelizes the returning signal.

Field-programmable gate arrays (FPGAs) and, more recently, RF System-on-Chip (RFSoc) devices have become the platform of choice for MKID readout. Early systems were built on general-purpose FPGA platforms such as ROACH and ROACH2,^{7,8} with dedicated designs reaching kilo-pixel scale;^{9,10} the current generation exploits the RFSoc family, which integrates multi-gigasample ADCs and DACs together with a large programmable-logic (PL) fabric and a hardened processing system (PS), collapsing what used to be a rack of converters and FPGAs into a single chip with a dramatic reduction in mass, volume and power.^{11–16} The IAC Electronics Department has exploited this platform to build an embedded data acquisition system (eDAS) on the ZCU111 evaluation board, capable of reading out an MKID array and performing the digital signal processing entirely in hardware and in real time.¹⁷

At the heart of any FDM readout sits the channelizer: the block that separates the wideband digitized stream into the individual resonator channels. A direct bank of independent down-converting filters scales as $O(M^2)$ in multiplications per input sample and is intractable for thousands of channels; the polyphase filter bank followed by an FFT reduces this to $O(M \log M)$ and is the standard solution in radio astronomy and superconducting-detector readout.^{18,19} A critically sampled PFB, however, suffers from scalloping loss: a tone that falls between two FFT bins is attenuated by as much as 3.9 dB, degrading the recovered sensitivity exactly where high-density arrays place their most closely spaced resonators. The established remedy is an oversampled, overlapping PFB (OPFB) with 50% temporal overlap, which flattens the channel response and recovers the lost sensitivity at bin edges.

Implementing an oversampled PFB on an RFSoc at multi-gigahertz sampling rates is where the engineering difficulty lies. State-of-the-art designs—most notably the work of Smith and collaborators at UCSB^{11,20,21}—achieve the required throughput by clocking the fabric aggressively (e.g., 512 MHz) and by time-multiplexing a single physical filter across the even and odd sample phases. These choices keep the design compact but constrain the prototype-filter depth to shallow values (typically 8 taps), which limits the achievable channel isolation. For dense MKID arrays, where crosstalk between adjacent resonators must be suppressed, deeper filters with sharper transition bands are highly desirable.

This work presents the design, implementation and validation of a high-isolation wideband channelizer that takes the opposite trade-off: it prioritizes architectural parallelism over raw clock speed. Built primarily in Vitis HLS—with the time-critical output serializer in VHDL—rather than with vendor-specific IP cores, the channelizer processes a 4.096 GSPS stream with a per-branch super-sample rate (SSR) of 16 and keeps the bulk of the pipeline at a conservative 256 MHz. A split-path topology computes the delayed and non-delayed polyphase branches in parallel, and a frame-level reordering scheme preserves a pipeline initiation interval of one ($II = 1$) while confining the only 512 MHz subdomain to a small BRAM-to-FFT section. The freed timing budget

is invested in a deep 16-tap Blackman prototype filter, doubling the depth of comparable high-speed implementations and markedly improving inter-channel isolation.

With respect to the preliminary version of this work presented at SPIE Astronomical Telescopes + Instrumentation 2026,²² this paper discloses the complete internal architecture of the reordering engine (Secs. 4.3 to 4.5), including two implementation hazards that, to our knowledge, are not documented in the literature: a pipeline-stage skew that silently desynchronizes state toggles in HLS-generated pipelines, and the selective state-preservation policy required to make the channelizer restartable without losing frame alignment. It also adds the post-synthesis resource utilization of every custom module (Sec. 4.7), a systematic 100-tone batch-simulation characterization of the channel response (Sec. 5.2), and an extended comparison with current readout channelizers (Sec. 6).

The remainder of the paper is organized as follows. Section 2 summarizes MKID FDM readout. Section 3 reviews polyphase channelization and the overlap technique. Section 4 describes the architecture and its HLS implementation in full detail. Section 5 reports simulation results and the cryogenic operation of the parent system. Section 6 compares the design with the state of the art, and Secs. 7 and 8 present conclusions and future work.

2 MKID Readout and Frequency-Division Multiplexing

An MKID is a lithographed superconducting LC resonator capacitively coupled to a common feedline. Below the critical temperature the film carries a supercurrent through Cooper pairs whose inertia gives rise to a kinetic inductance. When a photon with energy greater than twice the superconducting gap is absorbed, it breaks Cooper pairs into quasiparticles; the resulting increase in kinetic inductance lowers the resonant frequency, and the change in quasiparticle density modifies the internal quality factor. Both effects appear in the complex forward transmission S_{21} of the feedline as a shift and a broadening of the resonance dip, and are most cleanly read as a change in the phase of a probe tone parked at the resonator frequency.³

Because each pixel is a resonator at a distinct frequency, the whole array is interrogated at once by driving the feedline with a frequency comb—one tone per resonator—and recovering, for every tone, the amplitude and phase imprinted by its resonator. This is FDM readout. The digital backend must therefore (i) synthesize a programmable comb spanning the array bandwidth, (ii) play it out through a DAC, (iii) digitize the returning signal with an ADC, and (iv) channelize the wideband stream back into per-resonator measurements. In the IAC system the comb is generated in the PL by an inverse-FFT and look-up-table engine feeding the RFSoc DAC at 4.096 GSPS; the returning signal is captured by the RFSoc ADC at the same rate and enters the channelizer described in this paper. The array used for the measurements reported here (IAC-1) comprises ten superconducting resonators between 0.9 and 1.8 GHz with loaded quality factors of order 10^4 , cooled in an ADR; because the resonances lie within the first Nyquist zone of the 4.096 GSPS converters, no analog down-conversion is required before digitization.

3 Polyphase Channelization and the Overlap Technique

3.1 Polyphase Decomposition

A channelizer that realizes M channels as M independent down-converting FIR filters costs $O(M^2)$ multiplications per input sample. The polyphase identity reduces this to the cost of one FFT by recognizing that a long low-pass prototype filter $h[n]$ of length $K \cdot M$ can be partitioned into M shorter sub-filters (the polyphase paths), one per FFT bin.^{18,23} With p the path index and k the decimated-time index, each path output is

$$y_p[k] = \sum_{n=0}^{K-1} h_p[n] x[(k-n)M + p], \quad (1)$$

and the M path outputs are combined by an M -point DFT, efficiently evaluated as an M -point FFT:

$$X_m[k] = \sum_{p=0}^{M-1} y_p[k] e^{-j2\pi pm/M}. \quad (2)$$

In the IAC channelizer the prototype is a Blackman-windowed sinc, the channel count equals the FFT size, $M = 4096$, with $K = 16$ taps per polyphase path (a prototype of $K \cdot M = 65,536$ coefficients). The hardware parallelism is a separate super-sample rate, $SSR = 16$ samples per clock (Sec. 4), distinct from the polyphase order M . The channel spacing is $\Delta f = f_s/M \approx 1$ MHz, while the prototype low-pass cut-off is set to two channels so that each channel passband is approximately 2 MHz wide.

3.2 Scalping Loss and the 2/1 Overlap

A critically sampled channelizer (decimation D equal to the channel count, i.e., $M/D = 1$) leaves a periodic ripple in its composite frequency response: a tone landing midway between two bins falls on the skirts of both channel filters and is attenuated. For a DFT channelizer this worst-case scalping loss reaches ≈ 3.9 dB at the half-bin point, a direct loss of signal-to-noise ratio for any resonator whose probe tone does not sit exactly on a bin centre. The overlapping PFB removes this blind spot by oversampling the channelizer in time by a factor of two—the so-called 2/1 overlap—so that consecutive FFT frames share 50% of their input samples. Equivalently, with $M = 4096$ the frame advance is $D = 2048$, giving two overlap states. The composite response becomes flat to better than 0.1 dB across the full 2 MHz channel.

The price of overlapping is structural. Achieving 50% overlap requires reordering the polyphase outputs before the FFT according to the cyclic permutation

$$R(n) = (n \cdot D) \bmod M, \quad (3)$$

and feeding the FFT two interleaved sub-bands (referred to here as LOWER and UPPER). Realizing this reordering at the throughput of a 4.096 GSPS stream, with a 16-tap prototype and 4096 channels, is the core implementation challenge addressed in Sec. 4.




Fig 1 Per-channel response of a critically sampled channelizer (top, spacing = BW) versus the 2/1 overlapped PFB with the 16-tap Blackman prototype (bottom, spacing = BW/2). Each colour is one channel: adjacent channels cross at ~ 3.9 dB in the critical case (the scalloping loss at the channel edge) but at below 0.1 dB in the overlapped case, giving flat coverage. Channel spacing $\Delta f = 1$ MHz.




Fig 2 Simulated prototype-filter response for a 16-tap (this work) and an 8-tap design over the same 2 MHz channel. The deeper 16-tap filter provides a sharper transition band and a substantially deeper stop-band, i.e., higher inter-channel isolation.

3.3 Prototype Filter and Channel Isolation

Channel isolation—the attenuation a tone in one channel receives when it leaks into its neighbours—is governed by the prototype filter’s transition band and stop-band. Deeper filters (more taps) produce sharper roll-off and deeper stop-bands at the cost of more multiply-accumulate resources and longer critical paths. The IAC design uses a 16-tap Blackman prototype, twice the depth of the 8-tap filters common in high-speed RFSoc channelizers. Figure 2 compares the simulated response of the 16-tap and 8-tap prototypes for the same 2 MHz channel: the 16-tap filter reaches the stop-band roughly 15 dB deeper at the first adjacent channel and falls off markedly faster, which is precisely the property required to suppress crosstalk in a dense MKID array.

4 Architecture and HLS Implementation

4.1 Throughput, Super-Sample Rate and Clocking

The RFSoc data converters deliver $f_s = 4096$ MSPS. The Zynq UltraScale+ fabric sustains, for complex DSP designs with many DSP48 slices and FIFOs, clock rates of roughly 200–500 MHz. To avoid dropping samples the channelizer must process f_s samples per second in parallel words of SSR samples per clock, which fixes

$$f_{\text{clk}} \cdot \text{SSR} = f_s \Rightarrow \text{SSR} = \frac{4096 \text{ MSPS}}{256 \text{ MHz}} = 16. \quad (4)$$

SSR = 16 at 256 MHz is the comfortable operating point: SSR = 8 would force the entire receive chain to 512 MHz, beyond the timing margin of a 16-tap Blackman PFB, while SSR = 32 would halve the clock but double the bus width and exceed routing resources. Every DSP block in the project (the polyphase bank, the FFT, the channel-selection and phase blocks) is parameterized at SSR = 16. The functional readout system in operation today runs its entire processing chain in a single 256 MHz clock domain, crossing from the converter clock through an asynchronous FIFO.

4.2 Split-Path Polyphase Topology

The overlapping PFB needs two polyphase results per FFT frame—one for each overlap state. The IAC channelizer computes them with two physically parallel FIR branches: a delayed path and a non-delayed path, each a 16-tap, SSR = 16 polyphase bank with Blackman coefficients. Because each branch produces 16 samples per cycle, the reordering stage downstream receives 32 aggregated samples per 256 MHz cycle (SSR = 16 per branch, 32 in aggregate). This is the architectural choice that lets the design avoid doubling the clock: the second overlap state is produced by replicating hardware in space rather than by time-multiplexing one filter at twice the frequency.

Each branch is fed by a dedicated HLS *switch* module (`pfilter_switch`, in delayed and non-delayed variants) that assembles the polyphase input frames from the incoming SSR-16 stream. The switch implements the circular addressing of the polyphase decomposition of Eq. (1): it buffers the input stream in BRAM and presents, to its filter bank, the $K = 16$ time-aligned samples that each of the 16 lanes requires per clock cycle. The delayed variant is identical except

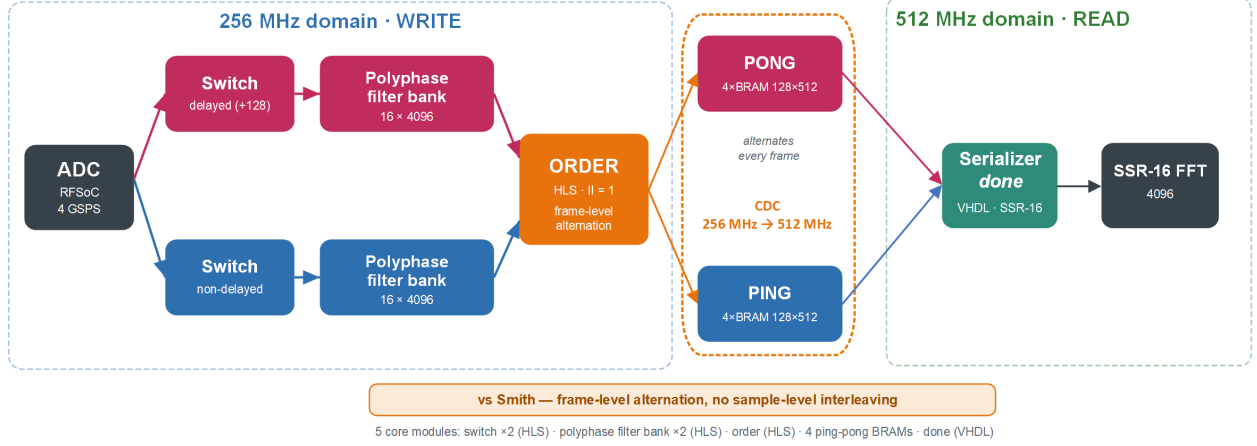


Fig 3 Channelizer datapath. In the 256 MHz write domain the two switch / polyphase-bank branches (non-delayed and delayed) feed the `order` block (HLS, $II = 1$) that performs the frame-level alternation; a clock-domain crossing then hands the data to the 512 MHz read domain, where the ping-pong BRAMs and the `done` module serialize the frames into the SSR-16, 4096-point FFT.

that its read schedule is offset by $M/2 = 2048$ samples, which is what generates the second overlap state of the 2/1 OPFB. Downstream of each switch, the filter bank proper (`pfilter_bank`) computes the 16-tap MAC per lane as a balanced multiply-add tree, mapped by HLS onto 256 DSP48E2 slices per branch (one multiplier per lane and tap) with a five-cycle pipeline latency and $II = 1$. The arithmetic is 16-bit data against 16-bit Blackman coefficients: each lane forms 32-bit products, and the four-stage adder tree grows the word to 36 bits at the filter output. Both modules close timing at 256 MHz with positive slack (Sec. 4.7).

4.3 Frame-Level Reordering

The granularity of the pre-FFT reorder is the single most important design decision in the OPFB. Two options exist. The sample-level approach used by Smith^{11,20,21} interleaves the two overlap states sample-by-sample within each FFT word; it is feasible with $SSR = 8$ and one physical filter time-multiplexed across the even and odd sample phases at 512 MHz. The frame-level approach adopted here instead alternates the destination of whole frames: each 512-bit word written to the buffer memories carries 16 samples of a single branch, unmixed, and what alternates frame-to-frame is the destination block (LOWER or UPPER). This is exactly equivalent to the cyclic shift of $M/2 = 2048$ samples demanded by Eq. (3).

The reason for choosing frame-level over sample-level is timing closure. With $SSR = 16$ per branch and two parallel branches, interleaving samples of the two branches inside one word would lengthen the combinational mux-select-write path beyond a single 256 MHz cycle and destroy the HLS pipeline's $II = 1$. Smith can afford sample-level reordering because that design carries only 8 samples per cycle through one filter; with 4096 channels and a high SSR, the IAC system cannot. Frame-level alternation preserves $II = 1$ at the cost of a single state bit. Figure 4 illustrates the resulting data movement, with the two branches written into alternating frame buffers.

The mechanism is split across two modules. An `order` block (HLS, 256 MHz) writes the branch outputs into four BRAMs organized as two ping-pong pairs; a `done` block (VHDL,

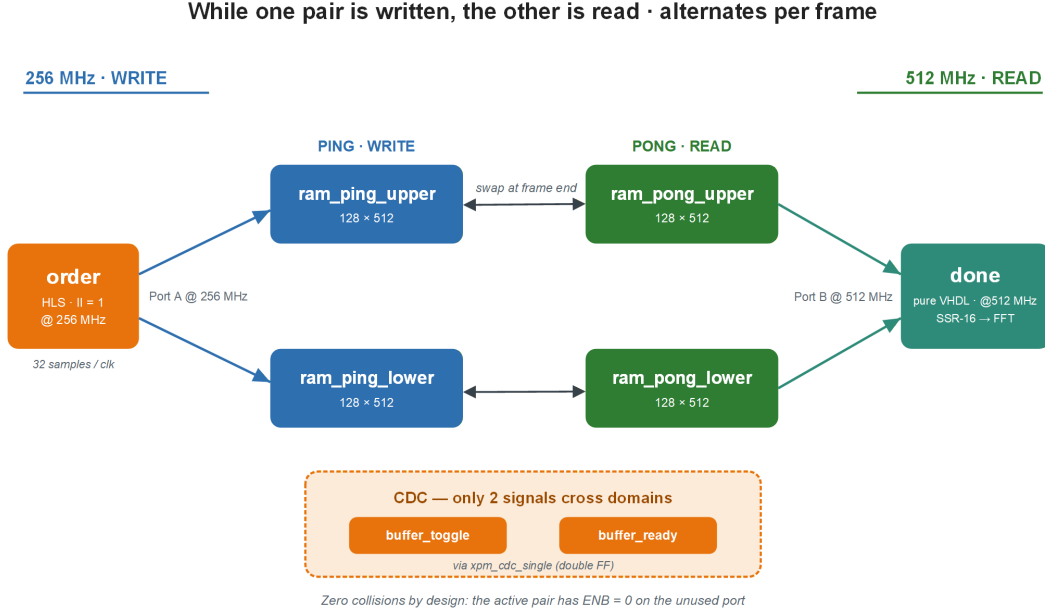


Fig 4 Ping-pong buffering that realizes the frame-level alternation. While one BRAM pair is written by `order` at 256 MHz, the other is read by `done` at 512 MHz, swapping at the end of each frame; only two control signals (`buffer_toggle`, `buffer_ready`) cross the clock-domain boundary through double-flop synchronizers, with zero read/write collisions by design.

512 MHz) reads them back word by word, alternating the LOWER and UPPER sub-bands on successive cycles (256 read beats per frame), to deliver 16 samples per cycle to the FFT core. Internally, `order` maintains three static state variables: a write pointer (`write_ptr`) that advances through the 128 addresses of the active buffer pair (each cycle writes one 512-bit word into the LOWER and one into the UPPER memory, so a complete frame spans $128 \times 32 = 4096$ samples), a ping-pong selector (`ping_pong`) that swaps the write target between the two buffer pairs at each frame boundary, and a frame-state bit (`frame_state`) that alternates the LOWER/UPPER destination of the delayed and non-delayed branches from frame to frame, realizing Eq. (3) at frame granularity. When packing, `order` truncates each 36-bit filter sample to 32 bits, discarding four least-significant bits that lie below the quantization noise floor of the 16-tap filter, so that sixteen samples fill one 512-bit BRAM word. The module synthesizes to a fully pipelined $II = 1$ kernel of roughly one thousand registers and one thousand LUTs—the permutation costs essentially nothing in fabric resources; its entire complexity lies in control correctness.

4.4 Pipeline-Stage Skew in HLS State Toggles

The `order` module exposed an HLS implementation hazard that, to our knowledge, is not documented in the readout literature and is worth reporting for designers following the same route. In the original source, the two state bits `ping_pong` and `frame_state` were toggled inside the same source-level `if` block at the end-of-frame condition. Although the two assignments are textually adjacent and semantically simultaneous, Vitis HLS is free to schedule independent operations into different pipeline stages, and it did: the two toggles were placed one stage apart, so

that for one clock cycle per frame the design observed a *mixed* state—the new buffer selection with the old overlap state. The channelizer remained functional in appearance (data kept flowing, no simulation errors were raised), but one frame per pair was written with the LOWER/UPPER assignment swapped, corrupting the overlap reconstruction downstream in a way that only became visible in the spectral response.

The fix is to make the dependency explicit: all frame-boundary toggles are grouped into a single unified update expression so that HLS sees one dependency chain and schedules the state transition atomically in one stage. The general lesson is that in HLS designs, statically scheduled pipelines do not guarantee the cycle-level simultaneity of source-level adjacent assignments; any set of state variables whose *joint* value defines correctness must be updated through a single expression or otherwise forced into the same pipeline stage. Verification of this class of bug requires inspecting control-signal waveforms cycle by cycle—checking output data alone can mask a one-cycle skew that only corrupts state periodically.

4.5 Selective Soft Reset and Restart Alignment

A readout channelizer is not a free-running block: acquisitions are started and stopped from the host, and the OPFB must survive stop/restart cycles without losing the alignment between its internal frame phase and the acquisition trigger (`pulse_in`). A naive full reset of the `order` state on stop would re-zero `write_ptr`, `ping_pong` and `frame_state`, so that on restart the reorder engine would restart its permutation from an arbitrary phase relative to the switch modules' circular buffers, producing a channelizer that works but whose LOWER/UPPER frames are swapped or shifted with respect to the incoming data.

The implemented policy is a *selective* soft reset: on stop, only the flag that gates the first output frame (`first_frame_done`) is cleared, while `write_ptr`, `ping_pong` and `frame_state` are preserved. The preserved state keeps the reorder engine phase-locked to the datapath, and the cleared flag guarantees that the first frame delivered after restart is a complete, well-formed frame rather than the tail of an interrupted one. The general rule extracted from this design is that only state variables that would break timing alignment if stale should be reset on stop; counters and pointers that encode the phase relationship between pipeline stages must be preserved. This restart behaviour was verified in the cycle-accurate testbench by stopping and restarting the stream at arbitrary points of the frame and checking that the spectral response is indistinguishable from a cold start.

4.6 Two Clock Domains and the Bounded 512 MHz Subdomain

Keeping the whole OPFB at 256 MHz would require the reorder stage to deliver the alternating LOWER/UPPER word sequence at 16 samples per cycle, which from $f_{\text{clk}} \cdot \text{SSR} = 2f_s$ implies $f_{\text{clk}} = 512 \text{ MHz}$ for $\text{SSR} = 16$. Rather than raise the entire chain, the design partitions the OPFB into two intentional zones. The main zone—the two FIR branches plus the `order` block—operates with 32 aggregated samples per cycle at 256 MHz; it is the most resource-hungry part (DSP48, BRAM, routing) but closes timing comfortably and leaves headroom for the rest of the design. A minimal sub-section—the ping-pong BRAMs, the `done` finite-state machine and the FFT core—is raised to 512 MHz while keeping $\text{SSR} = 16$. Because this sub-section is geometrically small

(four BRAM groups, one VHDL FSM and the FFT core), it absorbs the cost of 512 MHz routing without contaminating timing closure elsewhere.

The clock-domain crossing itself is deliberately minimal. The data plane crosses through the ping-pong BRAMs: each buffer pair is written entirely in the 256 MHz domain and read entirely in the 512 MHz domain, and ownership of a pair is swapped only at frame boundaries. Consequently only two control bits cross the boundary—`buffer_toggle`, announcing that the writer has released a pair, and `buffer_ready`, acknowledging the reader’s ownership—each through a conventional double-flop synchronizer with pulse expansion where required. Because the reader’s multiplexer only ever selects the buffer pair that the writer has released, simultaneous read/write access to the same BRAM address cannot occur by construction; the dual-port collision warnings that RTL simulators raise for these memories are therefore benign, a point worth noting since they are easily mistaken for real hazards during verification. An explicit asynchronous-FIFO CDC returns the FFT output to the 256 MHz domain.

This is the architectural difference from prior art. Smith^{11,20,21} carries the complete OPFB at 512 MHz because that system is exactly half the size on every axis— $M = 2048$ channels, 8 taps per branch, $SSR = 8$ —so the resources and routing fit at 512 MHz. With $M = 4096$ and 16 taps, raising everything to 512 MHz does not close timing on the ZCU111: the place-and-route reports fail or demand so much additional pipelining that the design no longer fits. Splitting the design and raising only the indispensable sub-section is what makes the deep-filter, high-isolation channelizer synthesizable. Table 1 summarizes the comparison.

Table 1 Architectural comparison between the UCSB sample-level OPFB and the IAC frame-level implementation.

Aspect	Smith (UCSB) ²⁰	This work (IAC)
Bank size	$M = 2048$ channels, 8 taps/branch	$M = 4096$ channels, 16 taps/branch
OPFB reordering	Sample-level (TDM interleave on even/odd edges)	Frame-level alternation (2 physical FIR branches)
Clocking	Entire OPFB at 512 MHz, one TDM filter	Mostly 32 samples/cycle at 256 MHz; only BRAM→done→FFT at 512 MHz
Parallelism	$SSR = 8$ — allows per-sample reorder at $\Pi = 1$	$SSR = 16$ per branch (32 aggregate) — forces per-frame reorder for $\Pi = 1$
Prototype depth	8 taps	16 taps (≈ 15 dB deeper stop-band)

4.7 Resource Utilization and Timing

Table 2 reports the post-synthesis resource estimates of every custom HLS module of the OPFB front-end, as generated by Vitis HLS 2022.1.2 for the ZCU111 target device (XCZU28DR, speed grade −2). All modules are fully pipelined at their target clock with positive slack. The dominant arithmetic cost is the pair of polyphase banks, at 256 DSP48E2 slices each—512 DSP slices in total, 12% of the 4272 available on the device—which is the direct price of instantiating the second overlap state in space rather than in time. The dominant memory cost is the pair of switch modules, whose circular polyphase buffers occupy about one third of the device block RAM between them.

The `order` kernel itself is almost free (~ 1000 FF, ~ 1200 LUT, no DSP): the entire frame-level permutation reduces to control logic, which is precisely the point of the frame-level scheme. The SSR-16, 4096-point streaming FFT core, generated in Vivado Model Composer, and the compact `done` FSM (VHDL) complete the chain and are not included in the table.

Table 2 Post-synthesis resource estimates of the custom HLS modules (Vitis HLS 2022.1.2, XCZU28DR-2). Percentages refer to the total device resources; for `pfilter_bank`, figures are per instance. All modules meet timing at their target clock with the positive slack listed.

Module	Clock	II	BRAM 18K	DSP48E2	FF	LUT	Slack (ns)
<code>pfilter_switch</code> (non-del.)	256 MHz	1	356 (16%)	0	4196	1573	0.39
<code>pfilter_switch</code> (delayed)	256 MHz	1	364 (16%)	0	4331	1621	0.39
<code>pfilter_bank</code> ($\times 2$)	256 MHz	1	0	256 (5%) ea.	18,561	13,360	0.29
<code>order</code>	256 MHz	1	ext.	0	1051	1159	0.53

The design is built with Vivado 2022.1 and Vitis HLS 2022.1, with the FFT core generated in Vivado Model Composer and the embedded software running on the PYNQ framework, which provides Python drivers over AXI4-Lite and Jupyter-based test benches on the PS. The deterministic digital latency from ADC to the 10 GbE interface in the functional system is below $\sim 3 \mu\text{s}$; the analog contributions of cables, mixers and the cryostat add a few microseconds more and are measured with a reference calibration pulse.

4.8 Channelized Output

After the FFT the overlapped polyphase chain ends: its output is the bank of 4096 channels—1 MHz spacing, 2 MHz flat-top per channel—delivered as separate 27-bit fixed-point real and imaginary components: the `done` serializer extracts the 18 most significant bits of each stored 32-bit sample as the FFT input word, and the FFT grows it to 27 bits. The OPFB testbench that is the subject of this paper reproduces the chain exactly up to this 27-bit FFT output. In the deployed functional system the transmitted word width is set by the Ethernet front-end: a VHDL serializer sign-extends each 27-bit component to 32 bits and packs the sixteen super-sample channels into 512-bit AXI-Stream words, so that one complex channel occupies 64 bits (32-bit real + 32-bit imaginary); a configurable AXI4-Lite selector then forwards one channel at a time to the host over 10 GbE for offline analysis. Each channel thus recovers the amplitude and phase of its resonator tone with no scalloping blind spot between FFT bins. Figure 5 shows the complete signal path of the channelizer, including the internal structure of every custom module described in this section.

5 Results

5.1 Simulated Channelizer Performance

The overlapped channelizer is validated as a cycle-accurate Xsim testbench that reproduces the full polyphase, reorder and FFT chain, including the hardware reorder modules (`order` + `done`); it is the intended replacement for the classic non-overlapped PFB of the deployed system. Sweeping




Fig 5 Complete signal path of the OPFB channelizer as implemented. In the 256 MHz domain, the RFDC/CDC interface feeds the two switch + filter-bank branches, whose outputs the order block (HLS, II = 1) writes into the ping-pong BRAM pairs with frame-level alternation; in the 512 MHz domain, the done serializer reads them back into the SSR-16, 4096-point FFT. The OPFB testbench studied here reproduces this chain exactly up to the FFT output.

a tone across the band and analysing the data produced by the hardware reorder yields the channel responses of Fig. 6: a representative group of adjacent channels (here channels 20–28) tile the spectrum with flat passbands, 50% overlap and a stop-band roughly 80 dB below the passband. Together with Fig. 1 this confirms that the 2/1 architecture eliminates the spectral blind spots of a critically sampled channelizer—recovering the ~ 3.9 dB scalloping loss to below 0.1 dB—while the 16-tap prototype (Fig. 2) delivers a stop-band roughly 15 dB deeper than an 8-tap design, substantially improving the channel isolation.

5.2 Batch Frequency-Sweep Characterization

Beyond the qualitative tone sweep, the channel response was characterized systematically with a batch campaign of one hundred independent Xsim simulations, each injecting a single tone at a different frequency across the band. Running one tone per simulation, rather than a single multi-tone run, isolates the response of each excitation from intermodulation and leakage of its neighbours and keeps every waveform database small enough for automated cycle-level inspection; the campaign is driven by TCL scripts that parameterize the tone frequency at elaboration time and post-process the captured waveforms offline.

The batch results confirm the expected behaviour of the complete chain, including the hardware reorder path. A tone placed at the centre of a channel (channel 20 in the reference configuration)




Fig 6 Channelizer response obtained from the hardware reorder data (order + done) over a tone sweep, $f_s = 4.096$ GSPS, $N_{\text{FFT}} = 4096$. Adjacent channels (20–28 shown) tile the band with flat, 50%-overlapping passbands and a stop-band ≈ 80 dB below the channel peaks.

is recovered in that channel with a response of 90.3 dB, while the two adjacent channels (± 1) respond at 84.3 dB—exactly 6.0 dB down, as expected for a tone lying 1 MHz off their centres, at the crossover of the 50%-overlapped passbands—and the response is symmetric on both sides to within the measurement resolution. This symmetry is the sensitive figure of merit: the frame-alignment failure modes discussed in Secs. 4.4 and 4.5 manifest precisely as asymmetric adjacent-channel responses or as frame-to-frame instability of the recovered amplitude, so the batch sweep doubles as a regression test for the reordering engine. Out-of-band leakage remains at the ≈ 80 dB stop-band floor of Fig. 6 across all injected frequencies.

5.3 Cryogenic Operation of the Readout System

The parent readout system—the functional eDAS¹⁷ running the deployed non-overlapped PFB bitstream—has been operated on a real MKID array in the cryostat. The IAC-1 array sits at the ≈ 100 mK stage of an HPD Model 103 ADR, with a HEMT low-noise amplifier anchored at 4 K; the transmit chain applies 80 dB of attenuation (40 dB at 300 K plus 40 dB at 4 K) for a drive power of about -90 dBm at the device. The board synthesizes the excitation comb and the channelizer recovers, per channel, the amplitude and phase of the returning tones, which are captured over 10 GbE to a network-attached storage backend and analysed offline. Figure 7 shows the experimental bench used for these measurements.

This setup exercises the complete end-to-end RF–digital–RF loop—comb synthesis, cryogenic excitation, low-noise amplification, digitization and channelization—on which the overlapped channelizer will operate once promoted to the deployed bitstream, providing the operating conditions under which the high-isolation channel response of Sec. 5.1 will be validated directly on the array.




Fig 7 Cryogenic test bench for the MKID readout. The ZCU111 RFSoc drives and digitizes the feedline through an 80 dB attenuated transmit path and a 4 K HEMT receive amplifier; the MKID array is held at ≈ 100 mK in an HPD Model 103 ADR cryostat.

6 Comparison with the State of the Art

Table 1 contrasts the two published OPFB implementations for MKID readout at multi-GSPS rates; this section places the design in the broader landscape of superconducting-resonator readout electronics.

The lineage of MKID readout channelizers begins with the ROACH-based systems developed around the CASPER ecosystem,⁸ such as the MUSIC readout,⁷ and continued with dedicated kilo-pixel designs such as the SRON multiplexer for 1000-pixel arrays⁹ and the second-generation UVOIR readout of Fruitwala et al.,¹⁰ all of which pair discrete data converters with one or more FPGAs. The RFSoc generation collapsed this architecture into a single device: Smith et al.^{11,20,21} demonstrated the first RFSoc OPFB for optical/near-IR MKIDs ($M = 2048$, 8 taps, sample-level reorder at 512 MHz), since consolidated into the deployed third-generation readout (MKIDGen3) on the RFSoc4x2 platform;¹² the Dublin group evaluated ZCU111-class boards for THz MKID readout,¹⁴ and general-purpose RFSoc frameworks such as QICK¹⁵ now serve the wider superconducting-device community. A complementary line of development, exemplified by SMuRF,¹⁶ follows a tone-tracking rather than channelizing architecture, dynamically steering each probe tone to follow its resonator.

Within this landscape the present design occupies a specific niche: it is, to our knowledge, the deepest oversampled channelizer (16 taps, $M = 4096$) implemented on a single RFSoc at 4.096 GSPS, and the only one built around frame-level reordering. The comparison with the UCSB design is the most informative because both implement the same 2/1 OPFB function with opposite architectural bets. The sample-level design minimizes silicon area—one physical filter, half the channels, half the taps—at the cost of running the entire OPFB at 512 MHz, which in turn caps the filter depth that closes timing. The frame-level design doubles the arithmetic (two physical branches, 512 DSP slices) but keeps all of the pipeline except the small BRAM-to-FFT subdomain at 256 MHz, and invests the freed timing margin in filter depth: 16 taps, a ≈ 15 dB deeper stop-band, and ≈ 80 dB of channel isolation. Which bet is preferable depends on the array: for sparse

arrays the compact design wins on resources, while for dense arrays—where adjacent-resonator crosstalk is the limiting systematic—the deep-filter design converts otherwise idle DSP resources into isolation. Notably, the two approaches are complementary rather than competing outcomes of the same optimization: both preserve $\Pi = 1$, and both are constrained by the same $f_{\text{clk}} \cdot \text{SSR} = 2f_s$ identity of Eq. (4); they simply satisfy it on opposite sides.

Methodologically the two designs are kindred rather than opposed: both are written primarily in C++ for Vivado/Vitis HLS and operated through PYNQ,²⁰ an approach that keeps the sources parametric (the same IAC design targets the ZCU208 at higher sample rates by regenerating with different SSR and clock constraints) and maintainable by a small team. The difference lies in where each flow spends its timing budget: the UCSB design drives one time-multiplexed filter at 512 MHz, which closes timing only with a shallow prototype, while the IAC design replicates the filter in space at 256 MHz—accepting a single hand-written VHDL module (`done`) where cycle-exact serialization at 512 MHz demands it, and the HLS-specific control hazards documented in Secs. 4.4 and 4.5, which is precisely why those hazards are reported here.

7 Conclusions

We have presented a high-isolation wideband channelizer for MKID readout, implemented as a custom design—primarily in Vitis HLS, with the time-critical `done` serializer in VHDL—on the Xilinx Zynq UltraScale+ RFSoc ZCU111. The central result is that architectural parallelism can be traded against clock speed to obtain a deeper, higher-isolation channelizer than aggressively clocked, IP-based designs of comparable bandwidth. By computing the delayed and non-delayed polyphase branches in parallel ($\text{SSR} = 16$ per branch, 32 aggregate) and by reordering at the frame level rather than the sample level, the design preserves a pipeline initiation interval of one while keeping the bulk of the channelizer at a robust 256 MHz and confining the only 512 MHz subdomain to a small BRAM-to-FFT section. The timing budget thus freed is invested in a 16-tap Blackman prototype—twice the depth of standard high-speed implementations—yielding a stop-band roughly 15 dB deeper and correspondingly higher inter-channel isolation, at a total arithmetic cost of 512 DSP48E2 slices (12% of the device).

In simulation the 2/1 overlapped channelizer recovers the ~ 3.9 dB scalloping loss of a critically sampled channelizer to below 0.1 dB, delivering a flat passband across the full 2 MHz channel, and a 100-tone batch sweep confirms symmetric adjacent-channel responses at the designed -6 dB crossover with ≈ 80 dB out-of-band isolation. The parent readout system has been operated cryogenically on a real MKID array, identifying individual resonators in total darkness and exercising the complete RF–digital–RF loop on which the overlapped channelizer will run once deployed. Taken together, the work shows that a custom HLS approach—prioritizing parallelism over raw clock rate—can outperform standard IP-based methodologies for wideband FDM readout, and it documents the design decisions and implementation hazards (frame-level reordering, pipeline-stage skew of state toggles, selective soft reset) that make such a design synthesizable and operable on a single RFSoc.

8 Future Work

The immediate next step is to promote the overlapped channelizer from the cycle-accurate test-bench into the deployed bitstream, replacing the classic non-overlapped PFB and adding the auxiliary modules its architecture requires, then to validate the flat, high-isolation channel response directly on the cryogenic array. Beyond that, work continues on full multi-channel transport over the 10 GbE link to a network-attached storage backend, on real-time phase processing for single-photon event detection, and on mechanical modifications to the cryostat to allow controlled optical illumination of the array through a fibre, enabling characterization of the detectors under known photon loads. The modular, reconfigurable structure of the design is intended to let these enhancements be added block-by-block without redesigning the channelizer core.

Disclosures

The authors declare that there are no financial interests, commercial affiliations, or other potential conflicts of interest that could have influenced the objectivity of this research or the writing of this paper.

Code, Data, and Materials Availability

The simulation data and analysis scripts that support the findings of this article are not publicly available at this time, as the design is part of an ongoing instrument development at the Instituto de Astrofísica de Canarias. They can be obtained from the corresponding author upon reasonable request.

Acknowledgments

The authors thank the Instrumentation Area of the Instituto de Astrofísica de Canarias (IAC) and the Universidad de La Laguna (ULL) for their support in the development of the eDAS. D.P.R. is “personal contratado predoctoral del Programa de Astrofísicos Residentes del IAC”. An AI language assistant (Claude, Anthropic) was used for language editing and LaTeX formatting of the manuscript; all technical content, results and conclusions are the authors’ own.

References

- 1 B. A. Mazin, *Microwave Kinetic Inductance Detectors*. PhD thesis, California Institute of Technology (2005).
- 2 P. K. Day, H. G. LeDuc, B. A. Mazin, A. Vayonakis, and J. Zmuidzinas, “A broadband superconducting detector suitable for use in large arrays,” *Nature* **425**(6960), 817–821 (2003). [doi:10.1038/nature02037].
- 3 J. Zmuidzinas, “Superconducting microresonators: physics and applications,” *Annu. Rev. Condens. Matter Phys.* **3**(1), 169–214 (2012). [doi:10.1146/annurev-conmatphys-020911-125022].
- 4 B. A. Mazin, S. R. Meeker, M. J. Strader, P. Szypryt, D. Marsden, J. C. van Eyken, G. E. Duggan, A. B. Walter, G. Ulbricht, M. Johnson, B. Bumble, K. O’Brien, and C. Stoughton, “ARCONS: a 2024 pixel optical through near-IR cryogenic imaging spectrophotometer,” *Publ. Astron. Soc. Pac.* **125**(933), 1348–1361 (2013). [doi:10.1086/674013].

- 5 S. R. Meeker, B. A. Mazin, A. B. Walter, P. Strader, N. Fruitwala, C. Bockstiegel, P. Szypryt, G. Ulbricht, G. Coiffard, B. Bumble, G. Cancelo, T. Zmuda, K. Treptow, N. Wilcer, G. Collura, R. Dodkins, I. Lipartito, N. Zobrist, M. Bottom, J. C. Shelton, D. Mawet, J. C. van Eyken, G. Vasisht, and E. Serabyn, “DARKNESS: a microwave kinetic inductance detector integral field spectrograph for high-contrast astronomy,” *Publ. Astron. Soc. Pac.* **130**(988), 065001 (2018). [doi:10.1088/1538-3873/aab5e7].
- 6 A. B. Walter, N. Fruitwala, S. Steiger, J. I. Bailey, N. Zobrist, N. Swimmer, I. Lipartito, J. P. Smith, S. R. Meeker, C. Bockstiegel, G. Coiffard, R. Dodkins, P. Szypryt, K. K. Davis, M. Daal, B. Bumble, G. Collura, O. Guyon, J. Lozi, S. Vievard, N. Jovanovic, F. Martinache, T. Currie, and B. A. Mazin, “The MKID exoplanet camera for Subaru SCExAO,” *Publ. Astron. Soc. Pac.* **132**(1018), 125005 (2020). [doi:10.1088/1538-3873/abc60f].
- 7 S. McHugh, B. A. Mazin, B. Serfass, S. Meeker, K. O’Brien, R. Duan, R. Raffanti, and D. Werthimer, “A readout for large arrays of microwave kinetic inductance detectors,” *Rev. Sci. Instrum.* **83**(4), 044702 (2012). [doi:10.1063/1.3700812].
- 8 J. Hickish, Z. Abdurashidova, Z. Ali, K. D. Buch, S. C. Chaudhari, H. Chen, M. Dexter, R. S. Domagalski, J. Ford, G. Foster, D. George, J. Greenberg, L. Greenhill, A. Isaacson, H. Jiang, G. Jones, F. Kapp, H. Kriel, R. Lacasse, A. Lutomirski, D. MacMahon, J. Manley, A. Martens, R. McCullough, M. V. Muley, W. New, A. Parsons, D. C. Price, R. A. Primiani, J. Ray, A. Siemion, V. van Tonder, L. Vertatschitsch, M. Wagner, J. Weintroub, and D. Werthimer, “A decade of developments in technology and applications for radio astronomy: CASPER,” *J. Astron. Instrum.* **5**(4), 1641001 (2016). [doi:10.1142/S2251171716410014].
- 9 J. van Rantwijk, M. Grim, D. van Loon, S. Yates, A. Baryshev, and J. Baselmans, “Multiplexed readout for 1000-pixel arrays of microwave kinetic inductance detectors,” *IEEE Trans. Microw. Theory Tech.* **64**(6), 1876–1883 (2016). [doi:10.1109/TMTT.2016.2544303].
- 10 N. Fruitwala, P. Strader, G. Cancelo, T. Zmuda, K. Treptow, N. Wilcer, C. Stoughton, A. B. Walter, N. Zobrist, G. Collura, I. Lipartito, J. I. Bailey, and B. A. Mazin, “Second generation readout for large format photon counting microwave kinetic inductance detectors,” *Rev. Sci. Instrum.* **91**(12), 124705 (2020). [doi:10.1063/5.0029457].
- 11 J. P. Smith, J. I. Bailey, and B. A. Mazin, “Highly-multiplexed superconducting detector readout: approachable high-speed FPGA design,” in *2022 IEEE 30th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, 1–2 (2022). [doi:10.1109/FCCM53951.2022.9786140].
- 12 J. P. Smith, J. I. Bailey, A. Cuda, N. Zobrist, and B. A. Mazin, “MKIDGen3: energy-resolving, single-photon-counting microwave kinetic inductance detector readout on a radio frequency system-on-chip,” *Rev. Sci. Instrum.* **95**(11), 114705 (2024). [doi:10.1063/5.0225768].
- 13 C. Bracken, E. Baldwin, G. Ulbricht, M. De Lucia, and T. Ray, “Modeling results and baseline design for an RF-SoC-based readout system for microwave kinetic inductance detectors,” (2022). arXiv:2212.07938 [astro-ph.IM].

- 14 E. Baldwin, M. De Lucia, C. Bracken, G. Ulbricht, O. Creaner, J. Piercy, and T. Ray, “Frequency domain multiplexing for microwave kinetic inductance detectors: comparing the Xilinx ZCU111 RFSoc with the new 2x2 RFSoc board,” *J. Low Temp. Phys.* **209**, 640–648 (2022). [doi:10.1007/s10909-022-02785-x].
- 15 L. Stefanazzi, K. Treptow, N. Wilcer, C. Stoughton, C. Bradford, S. Uemura, S. Zorzetti, S. Montella, G. Cancelo, S. Sussman, A. Houck, S. Saxena, H. Arnaldi, A. Agrawal, H. Zhang, C. Ding, and D. I. Schuster, “The QICK (Quantum Instrumentation Control Kit): readout and control for qubits and detectors,” *Rev. Sci. Instrum.* **93**(4), 044709 (2022). [doi:10.1063/5.0076249].
- 16 C. Yu, Z. Ahmed, J. C. Frisch, S. W. Henderson, M. Silva-Feaver, K. Arnold, D. Brown, J. Connors, A. J. Cukierman, J. M. D’Ewart, B. J. Dober, J. E. Dusatko, G. Haller, R. Herbst, G. C. Hilton, J. Hubmayr, K. D. Irwin, C.-L. Kuo, J. A. B. Mates, L. Ruckman, J. Ullom, L. Vale, D. D. V. Winkle, J. Vasquez, and E. Young, “SLAC microresonator RF (SMuRF) electronics: a tone-tracking readout system for superconducting microwave resonator arrays,” *Rev. Sci. Instrum.* **94**(1), 014712 (2023). [doi:10.1063/5.0125084].
- 17 A. Hernández Fernández, D. Díaz Martín, D. Hernández Expósito, J. J. Díaz García, R. J. Hoyland, L. F. Rodríguez Ramos, O. M. Tubío Araújo, and S. Rodríguez Pérez, “Study of a microwave kinetic inductance detector (MKID) readout system based on FPGAs and SoCs,” in *Proc. SPIE*, **13102**, 131022D (2024). [doi:10.1117/12.3019897].
- 18 F. J. Harris, C. Dick, and M. Rice, “Digital receivers and transmitters using polyphase filter banks for wireless communications,” *IEEE Trans. Microw. Theory Tech.* **51**(4), 1395–1412 (2003). [doi:10.1109/TMTT.2003.809176].
- 19 D. C. Price, “Spectrometers and polyphase filterbanks in radio astronomy,” in *The WSPC Handbook of Astronomical Instrumentation*, 159–179, World Scientific (2021). [doi:10.1142/9789811203770_0007].
- 20 J. P. Smith, J. I. Bailey, J. Tuthill, L. Stefanazzi, G. Cancelo, K. Treptow, and B. A. Mazin, “A high-throughput oversampled polyphase filter bank using Vivado HLS and PYNQ on a RFSoc,” *IEEE Open J. Circuits Syst.* **2** (2021). [doi:10.1109/OJCAS.2020.3041208].
- 21 J. P. Smith, *Scaling Energy-Resolving Microwave Kinetic Inductance Detector Readout*. PhD thesis, University of California, Santa Barbara (2024). [<https://escholarship.org/uc/item/5m22t48n>].
- 22 A. Hernández Fernández, D. Díaz Martín, J. J. Díaz García, R. J. Hoyland, L. F. Rodríguez Ramos, D. Portero Rodríguez, and S. Rodríguez Pérez, “A high-isolation wideband channelizer for MKID readouts: a custom HLS implementation on RFSoc,” in *Proc. SPIE*, **14156** (2026). Paper 14156-102, in press.
- 23 R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ (1983).

Alberto Hernández Fernández is an electronics engineer at the Instituto de Astrofísica de Canarias (IAC) and a PhD candidate in the Doctoral Program in Industrial, Computer Science and Environmental Engineering at the Universidad de La Laguna, in the industrial engineering, electronic technology and communications track. He received his degree in industrial engineering from

the Universidad Politécnica de Madrid (ETSII-UPM). His current research interests include high-performance FPGA/RFSoc digital signal processing and readout electronics for superconducting detectors.

Biographies of the other authors are not available.