JULY/AUGUST 2026

MICROWAVE PRODUCT DIGEST
& TECHNOLOGIES

Volume 37, Issue 07
microwaveproductdigest.com

IN THIS ISSUE

MPD
MICROWAVE PRODUCT DIGEST & TECHNOLOGIES

EDITORIAL STATEMENT *of* PURPOSE

Microwave Product Digest serves RF and microwave design engineers, research and development engineers, applications engineers and engineering managers. These professionals, working in facilities that serve both the commercial and government markets, are involved with the design, development, application, and use of systems and subsystems, devices, and techniques involving frequencies from RF to light.

Editor and Publisher
Barry Manz
barry@mpdigest.com

Graphic Designer
Derek Wood
wwooodd@gmail.com

Product Submission
barry@mpdigest.com

This publication is issued without charge, upon written request, to qualified persons.

Microwave Product Digest **is published by** **Microwave Media LLC**

SUBSCRIPTIONS
Subscribe online at:
www.microwaveproductdigest.com

FROM THE EDITOR

Barry Manz

Open RAN's Tradeoff

Few concepts in wireless infrastructure have been promoted as confidently as Open RAN. The pitch was straightforward: break the radio access network into interoperable components, define open interfaces between them, and operators would no longer be tied to a single equipment supplier. More vendors would compete for their business, costs would fall, and innovation could move at something closer to software speed rather than waiting for the next generation of a proprietary base station.

It was an appealing argument. After several years of commercial deployment, however, there is enough real-world evidence to ask a more practical question. What did operators gain from Open RAN, and what did they have to take on in return?

The first concern was performance, and on that point Open RAN has made substantial progress. Early multi-vendor deployments often trailed integrated RAN systems by 15% to 25% on important performance measures. Today, that gap is generally closer to 5% to 10%. It has not disappeared, but for many parts of the network it is no longer large enough to rule out the architecture.

Commercial deployments have also moved beyond laboratory demonstrations. AT&T is roughly halfway toward its goal of carrying 70% of its wireless traffic on open-capable platforms and has completed live calls using third-party radios. That was an important proof point because the radio interface was where many of the original doubts were concentrated. Rakuten Mobile continues to support millions of subscribers across a broad collection of suppliers, while Dish built a national greenfield network around a cloud-native, disaggregated model. In other words, Open RAN works, but the more complicated question is what it takes to make it work.

The architecture's defining technical break occurs at the lower-layer functional split known as Option 7-2x. This split sits inside the physical layer, between the distributed unit (O-DU) and the radio unit (O-RU). High-PHY functions remain in the O-DU, while low-PHY functions such as FFT and iFFT processing, cyclic-prefix insertion and removal, and PRACH filtering move into the O-RU. In the Category B implementation, the radio also handles digital beamforming and precoding.

The two units communicate over an eCPRI-based Ethernet fronthaul, usually operating at 10 or 25 Gbit/s. That is a major departure from the fixed-rate CPRI links used in earlier centralized-RAN systems.

There is a real benefit to this arrangement. Option 7-2x transports frequency-domain data for each spatial layer rather than sending time-domain samples for every antenna port. Block floating-point compression further reduces the load. Together, these techniques can cut fronthaul bandwidth by as much as a factor of five compared with CPRI.

That difference matters enormously in a 64T64R massive-MIMO radio operating at 3.5 GHz. With a typical eight- to 16-layer configuration, the system can remain within the 25-Gbit/s Ethernet class instead of requiring an impractical amount of dedicated fiber for every antenna stream.

This is one of the places where disaggregation clearly earned its keep. What received less attention in the original sales pitch was how unforgiving the interface becomes once all of those functions are separated.

Open fronthaul operates under a one-way transport-latency budget of about 100 microseconds. For ultra-reliable low-latency traffic, the budget can fall to 50 microseconds. That

continued on page 24

EDITOR'S CHOICE
Featured Products

Power Amplifier Modules

These internally matched modules provide drop-in RF stages that cut engineering time. The MSCQ02 uses the LDMOS push-pull LR2401 device to deliver 200 W saturated power and 180 W at 1 dB compression from 30 to 512 MHz at 28 VDC, with headroom to 36 VDC. The MCCQ04 is an LDMOS, high-efficiency module that operates up to 520 MHz, for military communications, broadcast, and NMR applications.

POLYFET RF DEVICES

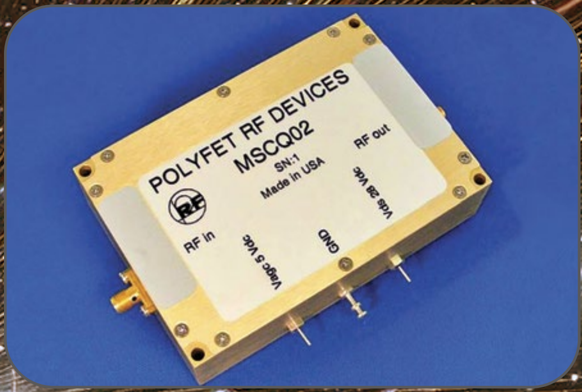

Two-Way Power Divider/Combiner

Model 50PD-1007 is a two-way reactive power divider/combiner covering 400 to 8400 MHz in a 50-ohm design. Insertion loss above the 3 dB split is 1.3 dB maximum (0.4 dB typical at 400 MHz), with amplitude tracking of ±0.40 dB and phase tracking of ±5° maximum. VSWR is 1.6:1 maximum above 500 MHz, and isolation is 18 dB minimum. It handles 10 W at the divider port, operates from –40 °C to +85 °C, and uses SMA female connectors.

JFW INDUSTRIES

---

10 GHz Lumped Element Filters.

This lumped-element filter series spans 5 MHz to 10 GHz and is available in bandpass, lowpass, high pass, and band-reject configurations. Bandwidths range from 3 to 150 percent, specified at the 0.5, 1.0, or 3.0 dB points, with a Chebyshev response standard and other responses available on request. The filters can be supplied with connectors or in surface-mount form. High-power options are available, and all models meet MIL-STD-202 conditions.

3-H COMMUNICATIONS

GaN Reflective Switches

The TS8728N and TS8729N GaN SPDT reflective switches serve as a solid-state alternative to PIN diode switches, removing the need for high-voltage bias networks and reducing board space and component count. The TS8728N operates from 500 MHz to 4 GHz and handles 400 W peak and 250 W CW. The TS8729N covers 500 MHz to 2 GHz and handles 550 W peak and 300 W CW. Both run from a single 2.6 to 5.5 V DC supply with CMOS-compatible logic control and integrate internal drivers.

TAGORETECH

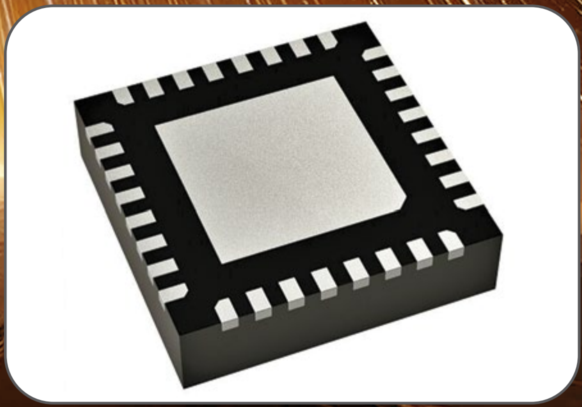

6 to 18 GHz Driver Amplifier

The QPA0023 is a driver amplifier is based on the company's 0.15 µm pHEMT process and operates between 6 and 18 GHz with 14 dB of small-signal gain, 29 dBm P1dB output and 30 dBm saturated output. It has gain flatness within ±1 dB and two-tone IMD3 of –42 dBc at 20 dBm per tone. It is matched to 50 ohms with integrated DC blocks on both ports.

QORVO

180-Deg. Hybrid Coupler

The ULTRA4010124 is a 3 dB, 180-deg. hybrid coupler covering 1.0 to 12.4 GHz. Typical amplitude imbalance is ±0.4 dB, phase imbalance ±7 deg., isolation is at least 20 dB, VSWR is to 1.3:1 and insertion loss is below 1.4 dB. It handles 20 W average and 3 kW peak and uses female SMA connectors. Typical applications include EW, SATCOM, and radar.

KRYTAR

High-Flex RG-58 FAKRA Cable Assemblies

These pre-configured assemblies use high-flex RG-58 cable rated to 3 GHz at 50 ohms, with a stranded copper center conductor and tinned copper braid built for millions of flex cycles. They come in jack-to-plug, plug-to-plug, and jack-to-jack forms in lengths from 12 in. to 3 m, and are well suited for automotive, GNSS/GPS, telematics, and ADAS.

AMPHENOL RF

Dual-Polarization FR3 Horn Antenna

This dual-polarized horn antenna targets the FR3 midband near 6 to 24 GHz. It accommodates both vertical and horizontal polarization and is designed for use in over-the-air tests, wideband device evaluation, phased-array development, and 6G research.

TMYTEK

200-W Smart Amplifier Module

The Model 1240 GaN amplifier module delivers at least 200 W from 900 to 1700 MHz. It has integrated control and monitoring features and is a good choice for mobile, airborne, shipboard, and ground platforms. It covers GSM and cellular, the GPS/GNSS L-band and the 900 MHz ISM drone band, for EW and counter-UAS use.

EMPOWER RF SYSTEMS

Handheld Spectrum Analyzers

The PSA Series 4 adds real-time capture to swept and fast-swept FFT modes. Models include the PSA1304 (1 MHz to 1.3 GHz), PSA2704 (1 MHz to 2.7 GHz) and the real-time PSA2704RT, which runs 80,000 FFTs/s with 10.8 MHz real-time bandwidth and 100% probability of intercept for 37.5 µs transients. Its firmware has smart markers, a 1 Hz counter, and limit lines.

AIM-TTI

WR-42 Standard Gain Horn Antenna

The PE3WAN042-15KF is a WR-42 standard gain horn covering 18 to 26.5 GHz with a typical 15 dBi gain and a 2.92 mm female interface via an integrated waveguide-to-coax transition. The pyramidal aluminum horn serves as a calibrated reference for measuring other antennas, and suits antenna measurement, radar, satcom and radome testing, and laboratory use.

PASTERNACK

Active Frequency Quadrupler

The A3432-4X24 multiplies a lower-frequency drive signal by four to generate a millimeter-wave output, using the company's balanced active-multiplier design. Active multiplication supplies high output power without a separate driver stage. It is offered with removable 2.92 mm connectors or waveguide flanges and integral bias regulation.

SPACEK LABS

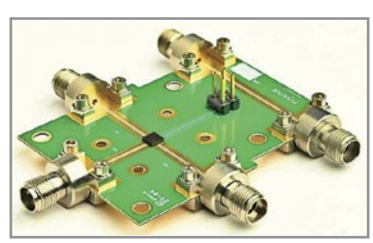

High-Power GaN-on-Si RF Switches

These high-power switches use a GaN-on-silicon process that pairs GaN performance with silicon's low cost. The FW2001 operates to 12 GHz, switches in 80 ns and settles to 0.1 dB in 200 ns, with no external components in a 3.0 x 3.0 mm package. It can be used in infrastructure, radar, drones, and test equipment.

FINWAVE SEMICONDUCTOR

Coaxial Cable Assembly

The MWX009 is usable from DC to 190 GHz and pairs with a 0.8 mm connector to interface with 0.8 mm probe systems, and its expanded-PTFE construction delivers phase stability against bending and temperature. It targets precision tests, on-wafer and sub-THz research, and semiconductor and chip characterization.

JUNKOSHA

Oven-Controlled Crystal Oscillators

The OX7000 Series is a surface-mount OCXO in a 9 x 7 mm package and covers 5 to 40 MHz using 3.3 VDC at 150 mA, warms up in about 3 min., and is hermetically sealed. Phase noise is –158 dBc/Hz at 10 kHz with output and stability to ±10 ppb. Applications include 5G infrastructure, data-center timing, and instrumentation.

RALTRON

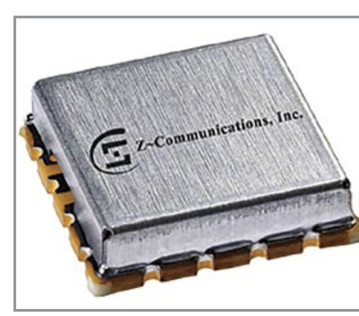

C-Band Oscillator

The CRO7000X2-LF provides a 7000 MHz output over a 0.5 to 4.5 VDC tuning range chosen to match common PLL ICs. Phase noise is –106 dBc/Hz at 10 kHz offset, supply current is 35 mA at 5 VDC, and its second-harmonic suppression –25 dBc. It ships in the MINI package, 0.5 x 0.5 x 0.22 in., on tape and reel, for radar, SATCOM, and test-and-measurement.

Z-COMMUNICATIONS

50-W GaN Amplifier

The RWP5872050-10 is a 50 W GaN power amplifier with a Class AB HEMT design, covering 5.8 to 7.2 GHz. At VDD = 32 V and Tc = 45° C, it delivers 35 dB of typical gain at 12 dBm input, 47 W typical saturated output, and 44.8% typical efficiency at 50 W output, with gain flatness within ±0.5 dB and spurious signals at -70 dBc. Typical input and output return loss is -10 dB. Quiescent current is 0.3 A, rising to 3.5 A at 50 W output. The flanged pallet integrates pre-MMIC, driver, and GaN transistor stages, with enable shutdown and current, temperature, and VSWR monitoring. RF ports use female SMA connectors, the module measures 72 by 150.8 by 16.8 mm, and an external heat sink is required.

RFHIC

Auto-Tune Filter

The A0101-02G-08G-P0-SMAFF absorptive Auto-Tune Filter (AtF) spans 2 to 8 GHz, sensing and attenuating high-power interferers automatically. The ferrite-based design requires no external power, control lines, or feedback loops. Insertion loss is under 2 dB, return loss is 12.5 dB , and the filter handles +40 dBm CW and 1 kW peak at a 1 µs pulse width and 0.1% duty cycle. Steady-state response time is 500 ns, with a -3 dB response time of 50 ns. The connectorized module uses female SMA ports and measures 1.96 by 0.5 by 0.5 in,, with an operating temperature range of -15° C to +85° C. Applications include SDR and ADC protection, radar, SIGINT, and EW receivers.

METAMAG

High-Power Limiter

The RFLM-961122MC-299 is an SMT silicon PIN diode limiter module providing always-on passive protection across the 960 MHz to 1.215 MHz Aeronautical Radio Navigation Service (ARNS) and IFF band. It requires no external bias or control and handles average power to +48 dBm at a case temperature of +85° C and peak power to +60 dBm, with insertion loss below 0.3 dB, low flat leakage, and recovery under 200 ns. The module is optimized for low small-signal insertion loss and uses a hybrid assembly technique housed in a low-profile 8 by 5 by 2.5 mm package. Beyond IFF, it suits radar and communications receivers that require high-power front-end protection.

RFUW ENGINEERING

GaAs Low-Noise Amplifier

The KA107C is a GaAs pHEMT broadband low-noise amplifier covering 1.3 to 3 GHz, providing 20 dB of gain with low noise and high linearity. It is supplied in a 3 by 3 mm hermetic surface-mount package rated for harsh defense and satellite environments. The part can be screened to MIL-PRF-38535 Class B and Class S flows with the required QCI, and EEE-INST-002 and other specialized test flows are available on request. An adjustable enable-pin voltage lets the user set bias current to favor either higher linearity or lower noise figure. The device is ESD-sensitive with Class 1 change notification on Class B and S part numbers. Applications include radar and communications receiver front ends.

MICROSS

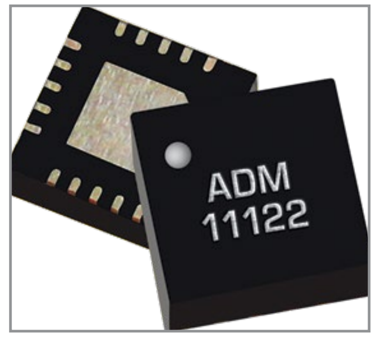

2-to-20 GHz Gain Block

The ADM-11122PSM is a wideband gain block covering 2 to 20 GHz, optimized for high second-order linearity. It provides 19 dB of small-signal gain, a 4.2 dB noise figure, +13.7 dBm output at 1 dB compression, and typical output IP2 of +55 dBm. Input and output return loss is 10 dB, and reverse isolation is 35 dB. The amplifier runs from a single 4 VDC supply at 84 mA and is internally matched to 50 ohms on both ports, using co-designed die and package to hold flatness across the multi-octave band. It ships in a RoHS-compliant 4 by 4 mm QFN package. The part targets direct-conversion and direct-sampling receivers and other wideband signal chains, including electronic-warfare systems and mobile test and measurement.

MARKI MICROWAVE

Switched Filter Bank

The PSB1030148 is a four-state switched filter bank covering 18 to 40 GHz, built on the company's PolyStrata microfabrication process. Its four channels span 17.5 to 24, 23 to 29.5, 28.5 to 35, and 34 to 40 GHz with 1 GHz of overlap. Insertion loss is under 7.5 dB per channel, with up to 75 dB of out-of-band rejection, return loss of 12 to 15 dB, passband ripple near 0.24 dB, and gain flatness of 0.6 dB. Typical P1dB is 26.5 dBm, and the bank is hot-switchable to 24 dBm with switching speed under 20 ns. It is packaged in a 15.8 by 16.3 by 4 mm surface-mount module with bottom-side launches for reflow or epoxy attach. Applications include SATCOM, defense, and test across Ka-band.

NUVOTRONICS

Cavity Notch Filter

the AE1575NS2174, a 1575.42 MHz cavity notch filter designed for electronic warfare and signals intelligence to testing and simulation, and for rejecting co-site interference. The AE1575NS2174 has rejection of at least 80 dB at 1575.42 and more than 48 dB at 1575.42 with a passband below the notch of 500 to 1545.42 MHz and the passband above the notch of 1605 to 2500 MHz. Insertion loss in the passband is 1 dB, return loss is 15 dB, and power handling is 80 W CW. The filter has an IP65 rating, operates over a temperature range of -40° C to +70° C, withstands a shock of 20 G for 11 ms, and withstands vibration of 20 G at 5 to 200 Hz. The filter has female Type-N connectors and measures 10.3 in × 1.7 in × 2.6 in.

ANATECH ELECTRONICS

SP6T Coaxial Switch

The W6 Series high-power SP6T electromechanical coaxial switch operates from DC to 6 GHz with Type SC female connectors. It is offered in failsafe, latching, or momentary (normally open) configurations at 12, 15, 24, or 28 VDC with 5 VDC operation available through TTL logic. Insertion loss is less than 0.6 dB at 6 GHz, with isolation of 80 dB and VSWR of 1.5:1. The break-before-make contacts switch in 20 ms, handle 1600 W CW at 1 GHz, and are rated for 1 million cycles. The standard operating range is -40° C to +85° C and handles 50 G shock and 15 G rms vibration. Options include suppression diodes and TTL logic. Applications include radar, EW, and communications.

CHARTER ENGINEERING

SCADA's Weak Link Lives in the RF Layer

by **Barry Manz,** ***Editor***

Supervisory control and data acquisition (SCADA) systems have been in the news lately, and not in a good way. Programmable logic controllers (PLCs) across U.S. water, energy, and government facilities are being targeted by alleged Iran-affiliated threat actors. But while the PLCs are the objective, to reach them, the attackers had to come in somewhere, typically a network using cable and increasingly fiber. However, wireless is the other way in, and where SCADA has been modernized with radio links at various points in the network, those links remain a weak point, one an adversary can reach without ever touching the wired core.

A large share of supervisory control and data acquisition traffic moves over licensed and unlicensed radio, across distances that make cable impractical, on equipment that in many installations has been in the field for 15 to 30 years. Control rooms still run on discontinued hardware and serial protocols that predate Ethernet, and the radios carrying telemetry from remote sites are of the same vintage: rarely patched, and frequently unencrypted because resisting an adversary was never a design requirement when they were installed. A wireless link an adversary can receive, terminated in equipment too old to have been built against one, is an RF and microwave problem. The band, the modulation, the link margin, the transceiver, and whether authentication is switched on the radio are the terms in which this exposure is described and eventually closed.

Why so much telemetry goes over the air

SCADA architecture is distributed by nature. A water utility polls remote pump stations and reservoirs, and a pipeline operator gathers pressure and flow data from valve sites strung across hundreds of miles. A utility's distribution automation reaches reclosers and capacitor banks on poles far from any fiber run. Trenching a cable to each of these points is rarely economic, so the remote terminal units at the edge report back over radio.

The physical layer varies with vintage and distance. Legacy licensed telemetry commonly sits in the VHF and UHF land-mobile bands and in the 900 MHz Multiple Address System allocation, carrying low-rate serial data over narrowband FSK. Unlicensed spread-spectrum radios in the 900 MHz, 2.4 GHz, and 5 GHz ISM bands handle shorter hops and higher rates. Point-to-point microwave in the licensed 6, 11, 18, and 23 GHz bands backhauls aggregated traffic from a concentrator to a control center. Increasingly, cellular and satellite fill in where private infrastructure is impractical, which introduces a dependency on a carrier's network that the operator does not control.

What ties these together is the protocol payload. Modbus and DNP3 dominate the application layer, and both were specified for reliability and simplicity on trusted serial links, not for operation across a channel an adversary can receive. Neither carries meaningful authentication in its common deployments. A frame that reaches the RTU with the right structure is acted upon. Whether that frame arrived over a wire or through the air is not something the protocol inspects.

Obscurity Was the Security, and SDR Ended It

A software-defined radio moves the functions once fixed in analog hardware, mixing, filtering, modulation, and demodulation, into software running behind a

continued on page 8

continued from page 7

wideband analog-to-digital converter, so one device can be retuned and reprogrammed to almost any waveform instead of being built for one.

The result is that a receive-only unit covering 25 MHz to 1.7 GHz costs a few tens of dollars; a transmit-capable platform spanning 1 MHz to 6 GHz, enough to reach every band in licensed and unlicensed telemetry short of the microwave backhaul hops, runs a few hundred bucks. That range takes in the VHF and UHF land-mobile channels, the 900 MHz Multiple Address System allocation, and the 2.4 and 5 GHz ISM bands where most wireless SCADA operates.

Security researchers who have reverse-engineered wireless SCADA installations have shown that plaintext Modbus and DNP3 can be recovered off the air with commodity receivers and open-source tooling, exposing the state of a monitored process to anyone within range. There is also the FM capture effect, in which the stronger of two co-channel FM signals dominates the receiver and opens a path to command injection. That is, a transmitter closer to, or with higher signal strength than, the legitimate remote can potentially override it.

What makes the device more than a scanner is the open-source signal-processing stack around it. Frameworks such as GNU Radio let a user assemble a demodulator, or a transmitter, from standard blocks, and once the work of decoding a particular protocol is done, it lives in a flowgraph that anyone can reuse. A single reconfigurable radio folds several formerly separate capabilities into one: survey a band, identify

and decode the traffic, replay a captured frame, or synthesize a new one and transmit it. The capability that makes SDR indispensable is what also erases the obscurity these telemetry links were counting on.

Installed, Then Forgotten

Compounding all of this is a maintenance model built around reliability rather than scrutiny. A remote SCADA site is usually unmanned, sometimes hours from the nearest technician, and the equipment in it — the RTU, the radio, the antenna, the power supply — sits in an enclosure exposed to temperature swings, moisture, dust, UV, and vibration that would retire most bench instruments in a season. The working assumption is that once a site is commissioned and reported, it is left alone until something stops working. Inspection is triggered by failure, not scheduled against it.

That approach keeps plants running, but it carries a security cost unrelated to uptime. Firmware that is never touched is firmware that is never patched. And a site no one visits is a site where a change in the RF environment — an added antenna, a nearby transmitter, a quietly swapped radio or reprogrammed — can persist for months with no one positioned to notice. The same neglect that lets a decades-old link keep working unattended is what would let tampering with it go unseen.

The Attack Surface in RF Terms

The simplest is denial of the channel. Jamming an unprotected narrowband telemetry link degrades or blocks the flow of data by raising the noise floor, an attack on availability. A control center that stops receiving from a remote site may fail in a safe position. Narrowband links with modest fade margins are the easiest to deny, and the transmitter required is neither large nor sophisticated.

Above jamming sits spoofing and injection, which attack integrity rather than availability. If the waveform, framing, and protocol are known, and none of it is authenticated, an adversary can generate a signal the receiver accepts as legitimate. The consequence ranges from feeding an operator false readings to issuing a command the field device executes. Replay is the low-effort version of the same idea: capture a valid transmission and retransmit it later, which succeeds against any link that does not include freshness in its message authentication, because there is nothing in the frame to mark it as stale.

Eavesdropping is the passive case and the enabler for the rest. Recovering the traffic reveals the process being controlled, the polling structure, and the bytes that constitute valid commands, which is the reconnaissance an active attack is built on.

Two system-level dependencies sit squarely in RF and microwave territory. The first is timing. Synchrophasor measurement units that give operators a wide-area view of a grid state depend on GNSS for microsecond time alignment. GNSS spoofing, one of the most thoroughly demonstrated SDR attacks in the literature, can corrupt the time reference that phasor data is built on, degrading state estimation without touching a single control message.

The second is the backhaul link itself. The February 2022 attack on the KA-SAT satellite network is the cleanest illustration. It was a network attack, malware that bricked modems, not RF interference, but its collateral effect was the loss of remote monitoring for 5,800 wind turbines in Germany. The turbines kept turning; the operator went blind. It demonstrated, at scale, that the communications path can be the single point of failure regardless of how robust the controller and the plant beneath it are.

continued on page 10

continued from page 9

Demonstrated Is Different from Exploited

With the partial exception of the satellite incident above, the RF layer has not been the vector in the most devastating industrial attacks. Stuxnet, the Ukrainian grid attacks, the Colonial Pipeline shutdown, the Oldsmar and Aliquippa water-system intrusions, and the current critical-infrastructure campaign all arrived through IT and OT network paths.

The reasons wireless links have not been exploited are straightforward. RF attacks require proximity; they require knowledge of a specific system, and they scale poorly compared with an intrusion that can be launched from another continent over the internet. Those are real deterrents, but they're eroding. SDR has lowered the knowledge and equipment cost. The number of wireless endpoints is climbing as industrial IoT pushes radios onto more field assets. The reading is not that a wireless-SCADA catastrophe is imminent, but that the factors holding the risk down are weakening, and the design decisions being made now will be in the field for decades.

What the Engineer Controls

The mitigations that matter most are, encouragingly, physical- and link-layer choices that RF and microwave engineers are positioned to make. The first is encryption and message authentication on the radio. Many current industrial transceivers support AES with authenticated framing that defeats eavesdropping, injection, and replay. The failure is rarely the absence of the capability and usually the presence of a large installed base that predates it, or a deployment where the feature was left disabled to preserve legacy interoperability or spare a few milliseconds of latency. Specifying authenticated encryption as a default rather than an option is the highest-leverage decision in the chain.

Spread spectrum is worth understanding precisely for what it does and does not provide. Direct-sequence and frequency-hopping techniques improve resistance to jamming and lower the probability of intercept, which raises the cost of both denial and reconnaissance. They are not, by themselves, security. A frequency-hopping link with no authentication is still injectable by an adversary who can follow or predict the hop sequence, so spread spectrum belongs in the design as jam resistance layered beneath cryptographic integrity, not as a substitute for it.

Anti-jam and link-margin discipline follow from the same reasoning. Adequate fade margin, frequency agility, and, on the higher-value backhaul hops, adaptive antenna techniques such as controlled-reception-pattern arrays that null an interferer all raise the power an attacker must bring to deny the channel. For timing-dependent nodes, GNSS anti-spoofing is now a tractable design element rather than a research topic: multi-constellation receivers, signal-authentication schemes such as Galileo's OSNMA, receiver autonomous integrity monitoring, and inertial or oscillator holdover that carries a node through a spoofing attempt without surrendering its time reference.

Finally, the radio itself is now a software target, so authenticated firmware and secure boot on the transceiver module close a path that is easy to overlook when the focus is on the RF chain. And because none of these measures is perfect, external spectrum monitoring earns its place as an intrusion-detection layer. A quiet band that suddenly carries an unexpected signal, or a jammer raising the noise floor, is an observable event, and monitoring for it turns an invisible RF attack into an alarm an operator can act on.

The Engineer's Stake

While controllers get the headlines, the link that carries a controller's data across open country is a radio link, and its security is decided by the same engineering choices that set its range and reliability. Modulation, framing, encryption defaults, and timing integrity are not features a security team can bolt on after the fact. They are fixed at design time, in the transceiver and the system around it.

The RF layer of wireless SCADA has been treated as invisible for a long time, but not any longer, because the equipment needed to receive it and to transmit it now costs less than a bench instrument and demands little skill from whoever operates it.

Software-defined radio erased the obscurity that once passed for security. Wireless endpoints multiply with every industrial-IoT rollout, and the jamming and spoofing methods once reserved for electronic warfare sit within reach of a widening set of actors. The conditions that kept these links safe by accident are the ones now giving way, and what replaces them must be engineered in on purpose. That work falls to whoever chooses the band, sets the modulation, and decides whether authentication is turned on, which is to say it falls to the RF and microwave engineer.

How Kohl's Connected the Sales Floor Before Retail Embraced IoT

European grocers pioneered electronic labels, but Kohl's brought wireless signage to U.S. department stores at scale.

by **Barry Manz**, *Editor*

MORE THAN 15 YEARS AGO, SHOPPERS walking through Kohl's stores encountered something remarkably advanced for its time. Mounted above clothing racks and merchandise fixtures were electronic signs whose prices and promotional messages could be changed remotely over a wireless network.

While most U.S. retailers still depended almost entirely on employees replacing paper signs by hand, Kohl's was quietly operating what would now be recognized as a large-scale connected-retail system.

The project deserves renewed attention because wireless electronic labels are now spreading rapidly across supermarkets, mass merchants and other retail environments. Yet Kohl's had already identified a particularly practical application for the technology years earlier.

The company's system was designed to distribute pricing and promotional information to electronic signs positioned throughout its stores. That's fundamentally different from what's being described today as "surveillance pricing", when companies use an individual's personal data, digital tracking, and AI algorithms to set personalized, individualized prices for goods and services. Instead, this article describes the wireless communications network and its operational benefits and in no way suggests that Kohl's used customer information to establish individualized prices. What it shows is that Kohl's was far ahead of its competitors beginning more than 15 years ago.

Kohl's was not the first retailer to deploy electronic labels. European supermarkets and hypermarket chains had adopted earlier generations of the technology before 2010. What made Kohl's unusual was its early, large-scale application within a major U.S. department-store chain, where larger wireless signs replaced the repeated manual resetting of promotional signage across hundreds of stores.

An Early Rollout

Kohl's publicly disclosed its electronic-sign initiative in 2010. In a filing with the Securities and Exchange Commission, the retailer said it planned to test electronic signs in 100 stores that fall. If the pilot proved successful, the company expected to expand the system substantially. Kohl's subsequently reported that electronic signs had been installed in about 100 stores during 2010. The company identified the project as one of several productivity initiatives intended to improve the efficiency of store operations and control hourly payroll expenses.

By the end of fiscal 2011, Kohl's reported that 475 stores

continued on page 12

continued from page 11

had been equipped with electronic signs. This was not a limited demonstration in a few showcase locations. It was a meaningful commercial deployment across hundreds of large-format stores at a time when wireless retail displays remained relatively uncommon in the United States.

Kohl's had recognized that its promotional model created an unusually strong business case for electronically controlled signage. The company regularly changed sales offers, discounts, and promotional messages across a large number of departments. Automating those changes reduced the need for employees to print, distribute, install, and later remove paper signs.

It also improved consistency between the promotional information displayed on the sales floor and the pricing programmed into the company's retail systems.

More than a Digital Sign

To a shopper, each unit appeared to be little more than a screen attached to a merchandise fixture. From an engineering standpoint, however, the displays were endpoints in a distributed wireless network.

Each sign required several basic elements:

- A low-power radio receiver or transceiver
- A display controller
- A unique device address
- A battery or other local power source
- Memory for storing pricing and promotional information
- A method of associating the sign with the correct merchandise

The data payload was extremely small. A product description, price, discount, and several lines of promotional text require very little bandwidth compared with conventional Wi-Fi or cellular applications.

Bandwidth, however, was not the principal challenge. A practical retail system had to communicate reliably with a large number of devices distributed throughout a store filled with metal racks, shelving, mirrors, merchandise and moving shoppers. The propagation environment changed continuously as products were rearranged and people moved through the aisles.

Source: AI-generated conceptual illustration

Displays also had to operate for extended periods without frequent battery replacement. That required radios, processors and screens designed around an extremely low duty cycle. Most signs spent the majority of their time displaying the same information, so their electronics could remain in low-power states between updates.

The network also needed a way to address individual signs or groups of signs. A storewide sale might require many displays to change at once, while a markdown affecting one product category might involve only a small subset.

Depending on the system architecture, the network could also verify whether a sign had received an update, identify a nonresponsive endpoint, or report a low-battery condition. The engineering objective was not high throughput. It was dependable communications with a dense population of devices consuming very little energy.

Reconstructing the Original System

Publicly available corporate filings document Kohl's rollout and its operational objectives, but they provide relatively little technical detail about the original radio system. Contemporary retail-industry reports identified Altierre Corp.as a supplier associated with electronic signs used by Kohl's. Altierre developed proprietary, ultra-low-power wireless networks for electronic shelf labels and related retail displays.

Public regulatory records from the period also show that Altierre was certifying electronic-label equipment operating in the 2.4-GHz industrial, scientific and medical band. That spectrum was already shared with Wi-Fi, Bluetooth and numerous proprietary wireless systems.

However, publicly accessible documentation does not conclusively establish that every Kohl's deployment used the same Altierre hardware model, frequency plan or network configuration. The precise architecture may also have changed as the system expanded or was upgraded.

The most reasonable technical reconstruction is that each store used one or more centrally connected RF access points or controllers to communicate with electronic signs throughout the sales floor. Pricing and promotional information would have originated in Kohl's merchandising systems, passed through the store's local technology infrastructure, and then been transmitted to the appropriate displays.

The network may have used proprietary communications techniques optimized for extremely low power and large endpoint populations. Bidirectional communications would have allowed signs to acknowledge updates or report status information, although the exact protocol and implementation used in the original Kohl's system have not been publicly documented in sufficient detail to describe them conclusively.

The broad architecture is clear even where the implementation details are not: centrally managed retail data was distributed over a local wireless network to battery-powered displays positioned throughout the store.

The importance of the Database

Electronic signage is sometimes presented primarily as an RF problem. In reality, the radio link was only one part of a much larger operational system. Each sign had to be associated with the correct product, rack, or merchandise group. The price displayed on the sales floor had to correspond with the information in the retailer's merchandising and point-of-sale systems.

A wireless network could deliver a message perfectly and still show incorrect information if a sign had been linked to the wrong product.

If merchandise moved, the associated sign had to move with it or be reassigned. If a department was reorganized, the database mapping between signs and products had to be updated. If a display stopped responding, store employees needed a way to identify and correct the problem.

A wireless network could deliver a message perfectly and still show incorrect information if a sign had been linked to the wrong product.

The system therefore required more than radios and displays. It also depended on:

- Merchandising-database integration
- Store-level network management
- Device-association procedures
- Handheld tools or administrative software
- Exception reporting
- Maintenance and battery-replacement procedures
- Staff training

For a retailer operating hundreds of locations, centralized control offered enormous operational reach.

continued on page 14

continued from page 13

It also required disciplined data management, because incorrect information could be distributed just as efficiently as accurate information.

This backend integration was likely one of the reasons many retailers were slow to adopt electronic signage. The visible display was relatively simple. The larger challenge was incorporating it into decades-old pricing, inventory, point-of-sale and store-management systems.

Why Kohl's Got In Early

Kohl's had an unusually favorable operating environment for electronic signage. The company relied heavily on frequent sales, promotional events, percentage discounts and changing offers. Those promotions traditionally required employees to replace large numbers of printed signs throughout each store. The labor involved was repetitive, time-sensitive, and easily understood. Electronic signs could reduce that workload while also helping stores implement promotions more consistently.

Kohl's also avoided the most expensive version of the electronic-label problem. A supermarket may contain tens of thousands of shelf positions, each potentially requiring its own small electronic label. A Kohl's display could represent an entire rack, fixture, clothing style, or merchandise group.

Modern electronic shelf-label systems benefit from technical advances that were either unavailable or considerably more expensive when Kohl's began its rollout.

The retailer could therefore obtain much of the benefit of automated pricing without attaching a wireless display to every individual product. That distinction substantially improved the economics.

A lower endpoint count meant:

- Fewer displays to purchase
- Fewer devices to install and associate
- Lower network density
- Less ongoing maintenance
- Fewer batteries to monitor and replace
- A simpler physical deployment

The larger signs used in department stores also had more room for readable displays, batteries and radio electronics than a compact shelf-edge label. Kohl's had effectively selected the portion of the problem where the operational benefit was high, and the number of wireless endpoints remained manageable. That was the strategic insight behind the rollout.

Why Most Retailers Waited

The fact that Kohl's found an attractive application did not mean the same economics applied to every retailer. A store that changed prices infrequently could continue using paper labels at very low material cost. Even when labor was required to replace them, the savings might not justify the capital expense of electronic displays, RF infrastructure and backend software.

Electronic-label hardware was also more expensive 15 years ago. Displays, low-power radios, batteries and controllers did not yet benefit from the scale and integration now common in the Internet of Things market.

Most early systems also used proprietary communications protocols. A retailer adopting one of them risked becoming dependent on a single vendor for:

- Compatible replacement displays
- Access points
- Management software
- Firmware support
- Spare parts
- Future system expansion

That was a substantial commitment for an application traditionally handled with inexpensive paper.

The introduction of a standardized Bluetooth Electronic Shelf Label profile many years later illustrates how significant interoperability and vendor dependence had been as barriers to adoption.

Retailers also had to consider the physical realities of store operations. Displays are moved, knocked loose, damaged, placed on the wrong fixture, or left behind during department resets. Products are discontinued, packaging changes, and merchandise is routinely reorganized.

Electronic signs eliminated a large amount of paper-handling labor, but they did not eliminate the need to manage the relationship between the digital display and the physical merchandise. For many retailers, waiting for lower hardware costs, better battery life, more mature software and broader vendor competition was a rational decision.

Retail Technology Catches Up

Modern electronic shelf-label systems benefit from technical advances that were either unavailable or considerably more expensive when Kohl's began its rollout. Contemporary systems use improved e-paper displays, highly integrated low-power radios, inexpensive microcontrollers and more capable network-management platforms. Battery life can extend for years, depending on display size, update frequency and operating conditions.

The labels can also perform more functions than simply showing prices.

Some include LEDs that flash to help employees locate merchandise while fulfilling online orders. Others support inventory management, replenishment, planogram verification, markdowns and product-location services. Cloud-based software can coordinate pricing and promotional information across thousands of stores. Handheld devices can associate a label with a product by scanning both identifiers. Network-management platforms can locate failed devices and confirm that updates were completed successfully.

These additional capabilities improve the economic case. A display that only replaces paper must justify its cost through reduced sign-changing labor. A display that also speeds order picking, supports inventory management and improves store accuracy can generate value across several operations.

As a result, the retail industry is now moving more broadly toward an architecture that Kohl's recognized much earlier: centrally managed information delivered over a low-power wireless network to displays distributed throughout a physical store.

The terminology has evolved. The devices are generally called electronic shelf labels, while the network may be described as an IoT platform. The basic engineering concept, however, would have been entirely familiar to anyone who examined the Kohl's system more than a decade ago.

The most striking aspect of Kohl's electronic-sign deployment may be how little public attention it received. It operated in full view of millions of shoppers, yet few people recognized it as an unusually early large-scale wireless application. That lack of attention was not necessarily a weakness. In many respects, it demonstrated that the technology had been successfully absorbed into normal store operations.

Customers did not need to understand the RF network. Employees did not need to think of themselves as managing an IoT platform. The system simply updated promotional information, reduced paper handling, and removed repetitive work from the sales floor.

Kohl's did not need broadband communications to every display. It needed to transmit small amounts of information reliably to battery-powered devices and keep those devices synchronized with its merchandising systems.

The technical requirements were modest in terms of bandwidth but demanding in terms of reliability, power consumption, endpoint management and operational integration. Kohl's also showed that successful wireless adoption often depends less on deploying the most advanced radio technology than on selecting the right business problem.

The company found an application with frequent information changes, measurable labor costs and a manageable number of endpoints. That combination made electronic signage practical years before the economics worked as broadly across other retail categories.

Today, wireless shelf labels are frequently presented as a new chapter in retail automation. In reality, Kohl's had already written an important part of that story. Its achievement was not merely installing electronic signs early. It was recognizing where the technology could deliver immediate operational value and putting a large wireless network to work while much of the retail industry continued changing paper signs by hand. MPD

Off the Shelf: A Conversation With Murray Pasternack

Editor's note: In 1972, Murray Pasternack founded Pasternack Enterprises in a 140-square-foot bedroom with $600. In an RF and microwave business built on custom work and weeks-long lead times, he did something no one else would: stocked standard components and sold them from a catalog listing a price and next-day delivery.

Q: Before we get started, can you give me a little background on yourself?

Pasternack: I was born in lower Manhattan, and during World War II I spent my early childhood in the South. My father was an immigrant from Romania, so he wasn't draftable, and he took a job with the Navy as a plumber at bomb and torpedo factories all over the South. They built those plants where there weren't a lot of people, but they had to have a railroad and a highway. One of the places that met that requirement was Orlando, Florida — long before Disney. So, I lived in a trailer in Orlando. A little later we moved to Church Hill, Tennessee, a small town in the Appalachian Mountains, where they built the largest explosive plant in the world.

I came back to Manhattan around age five. At 14, in 1954, while in junior high school, I built a semiconductor oscillator and turned it into a keyboard for music. I entered the New York City School Science Fair. I did nothing with it except take it apart again and use the parts for something else. Only two people in Manhattan won first prize; I got a very nice letter from the Assistant Superintendent of the Board of Education and a $50 check. Coffee was a nickel back then.

I went on to high school in Manhattan, but my family moved out of the city to Long Island, so I finished at a public high school there and then went to what is now known as SUNY Farmingdale State College — which, by the way, is now the Murray Pasternack School of Engineering Technology. I am the largest single individual contributor to the school.

Q: So when did you make the transition to work?

Pasternack: After I graduated in 1960, I went to work for Grumman Aircraft, testing the electronic equipment that went into aircraft. I left after 11 months and went to a company called Servo Corporation of America. Their claim to fame was an infrared detector for the railroad industry — it was placed on the side of the tracks, and if a train came by and the sensor detected a lot of heat from a wheel bearing, it would alert someone so the problem could be repaired before the wheel came off. They also had a contract for a radio direction finder for the Marine Corps called the GRD-11, that I worked on. It weighed 1,700 pounds — not exactly a light piece of equipment to carry into battle — but I learned a lot more about electronics on that job.

Q: Where did you go from there?

Pasternack: I had a professor who told me to change jobs whenever you stop learning and getting educated, which made a lot of sense. So, I left Servo, started night school at Hofstra College, and went to work for a little company on Long Island called LEL — Linear Equipment Laboratories — which made IF amplifiers and mixer-preamplifiers for the growing microwave field, mostly for the military. It was my first job as an engineer rather than a technician.

LEL was acquired by Varian Associates, and I was in the fortunate position where several of the people above me left to form a company. That put me in charge of a small project for the Gemini radar receiver, which was used for the first rendezvous in space. I spent the next six years with Varian, and the company kept moving me into higher positions until, in 1969, I topped out the pay scale. I was told I couldn't expect anything more than a cost-of-living raise, so I told my wife I was quitting to start a company.

Q: And that became your first company?

Pasternack: That company was Southwest Microwave (not to be confused with the current company of the same name — Ed.), which I started in January 1970. We made IF and RF amplifiers and mixer-preamps, very similar to those the LEL division of Varian had been making. Another business owner suggested we merge, and that became Environmental Communications, making RF amplifiers, log amplifiers, mixer-preamplifiers, and some early cable assemblies. We eventually broke up and I had to sue him. I acted as my own and won the case.

Around that time, I decided to go to law school at night, which took four years, while also starting a little operation in a 140-square-foot bedroom in my house, making products like those I'd made at Varian, Southwest Microwave, and Environmental Communications. I knew the vendors, I had the drawings, and I knew how to design these things. I called it Pasternack Enterprises, and I started it with a total investment of $600.

Q: Tell me how the catalog business evolved.

Pasternack: When I was at Varian, I tried to sell IF amplifiers off the shelf by mailing out catalogs, and I was happy with the results. The LEL division was happy too, but corporate didn't think it was the way to do business—they wanted an applications engineer or salesperson calling on customers to sell product.

So I started mailing catalogs at Southwest Microwave, then at Environmental Communications, and we were getting pretty good results. I built up the mailing list by going to public libraries and college catalogs, finding the physics and engineering departments, and adding the professors and their mailing addresses to my list. I was essentially spamming by mail — sending out a lot of catalogs four times a year to everybody on the list. The list started at a couple hundred names and grew to more than 120,000. Eventually we were mailing three million catalogs, using lists from Microwave Journal and other sources, even QST magazine.

The secret to that business was having the parts on hand, with price and delivery information right there in the catalog. People were less interested in price than in fast availability. At that time, if you called a sales rep or the factory, you'd hear back in a week or so, with a price and a delivery of six to 12 weeks. My approach was different: I bought hardware in large quantities, sold it in small quantities at reasonable prices, and made delivery the same day or the next day. Nobody else was even publishing a catalog with a price, so I was doing something nobody else did. People told me I shouldn't, but I never understood why, because I thought it was a great idea. I looked at the Sears catalog — almost everyone had one — and figured that if Sears could put out a catalog two inches thick, full of products, each with a stock number, a price, and delivery from stock, I could do the same thing. Why couldn't you sell electronics that way?

continued on page 18

continued from page 17

Q: Fast delivery requires keeping a lot of standard products in stock. How does that work?

Pasternack: That was easy, and I'll tell you how it worked. We started with a basic catalog, and every time someone called asking for something slightly different from what I had, I'd say, "Give me a few hours and I'll get back to you with a price and delivery." I'd find the item, find a source, or find something close to it, offer it to them, add it to the catalog, and buy it in quantity. We built up the catalog based on what customers requested. Today they claim to have about 40,000 items.

Q: You can't have 40,000 items in stock, can you?

Pasternack: Yes, I believe they can. I had several thousand items in the early years of the company. I had a very simple system. I bought what I anticipated would be a year's worth of an item — which usually didn't last that long, because the business was growing. I'd mark a certain percentage of the product separately from the rest, with a tag that said, "use last." When someone pulled that "use last" label off, I reordered. So, we almost never ran out of inventory. It was a very simple system, and this was before computers. If we bought a thousand of something, we might take a hundred, put it in a separate bag with a separate label, and set it aside. Sometimes it went faster than I thought, sometimes it was a little slower. And there were some items I bought in very large quantities because the price was just too good to pass up.

For example, are you familiar with the GR-874 General Radio connectors? They were designed around 1935. General Radio went out of the RF connector business, and I asked if I could have their tooling. They wouldn't give it to me, but they were kind enough to send me the drawings. I sent the drawings to China, and the Chinese could make that connector for about $1 in large quantities. They wanted a very large minimum order. We used them to make adapters to other connector series & sold them for about $100 each.

We would also contact suppliers who were making a large run for somebody else and say, "You're making a 5,000-piece run, and I'm sure you'll have extras, because when you machine parts you want to be sure you end up with enough good ones. We'll buy the balance at the same price you're charging your other customer." It converted their excess into cash and gave us products and inventory. The original relationships were with U.S. suppliers, then Taiwan. I found a company in Taiwan in the early 1980s, and they turned out to be a great supplier. That relationship, I understand, continues. Their key people had been trained in Japan, and they had Japanese equipment, and their quality in the early 1980s was much better than the Chinese could deliver. The products they provided were almost as good as the Japanese & at prices very close to the Chinese prices. So, I was willing to pay a little more for a much better product.

Q: The microwave industry has largely been a custom business. How do you handle that while still promising fast delivery?

Pasternack: I found that most people can use something almost as good as what they think they want when they're setting up a test or building a prototype. Most of the time we were selling small quantities to engineers and technicians building prototypes, and we might or might not get the production business afterward. But that small-quantity business, if you look at it worldwide, is a lot of volume. And sometimes time is so important that someone is willing to try something that isn't quite up to the specs they think they need.

Q: At some point you stepped back from the company.

Pasternack: My son, Lawrence Pasternack, took over the company and he sold it 10 or 11 years later to private equity, and I went into practicing law full time. I have been a member of the State Bar of California for 50 years this year.

Q: So you haven't had a whole lot to do with Pasternack for a while.

Pasternack: I have no current ownership and no real involvement in the company. But I get a great amount of pleasure watching what they're doing.

Q: What have you been doing since?

Pasternack: I helped start a community bank in 2008. I was one of the founders, the principal stockholder, and a director. I got to the point where I was spending more time with the bank than in my law practice. So, I gave up the law. We merged that bank with another small bank that was based in San Diego. We then had more branches. Eventually we merged the combined bank into Enterprise Bank out of St. Louis, so we ended up with about 56 branches across seven states, from Florida to California.

People said 2008 was a terrible time to start a bank, but it wasn't. We didn't have any bad loans, even when other banks did. We were able to lend money to people who were in good shape but needed help at a bad time, and we became one of the largest Small Business Administration lenders on the West Coast.

Q: Are you still working these days?

Pasternack: When I was at Servo, the president of the company came into the lab one day and said we should buy their stock. I had no idea how you bought stock. I was 22 years old, and my family bought bonds and CDs, never stocks, so I had zero education in it. Someone said I should go over to Merrill Lynch, Pierce, Fenner & Smith & I did go & make my first stock purchase. I lost money on that "hot tip". I decided I needed to learn about stock investing. I bought a book called The Intelligent Investor by a fellow named Benjamin Graham. It turns out he's the man who taught Warren Buffett how to trade stocks. I read the book and started trading, and I've been a stock trader, for better or worse, for 62 or 63 years now. I do it full time, and that's how I generate the money I've donated to the State University of New York & others.

Q: You've had a fantastic life.

Pasternack: I have enjoyed living a great life. I truly have lived the American dream. I would never have thought this would happen when I was a youngster. Earlier this year, my alma mater, the State University of New York (SUNY), Farmingdale State College (FSC), named the School of Engineering Technology the "Murray Pasternack School of Engineering Technology". This past May, I was the keynote speaker at the SUNY FSC graduation. At that graduation, I was honored by the University with the granting of an honorary Doctor of Humane Letters degree.

RF and Optical Characterization Challenges for Photonic Integrated Circuits

by **FormFactor**

THE RAPID BUILDOUT OF AI data centers is putting enormous pressure on the interconnects moving data among GPUs, accelerators, switches, and storage, and adding more racks and processors only gets you so far. The traffic between them also has to move faster, with lower latency and without sending power consumption through the roof. Increasingly, that means optics.

Conventional RF devices are characterized electrically and optical devices optically, but a PIC requires both to form an astonishing array of functions.

The transition from NRZ to PAM4 signaling has made 224 Gb/s per lane practical, with 448 Gb/s now coming into view. Those lane rates support optical transceivers operating at 1.6 Tb/s today, with 3.2 Tb/s not far behind. Reaching those speeds is not simply a matter of designing faster devices. Packaging and test become increasingly difficult as well, particularly for the photonic integrated circuits (PICs) at the center of these systems. Their electrical and optical functions are closely intertwined, so neither can be characterized effectively in isolation.

Why Photonic ICs Are Different

Conventional RF devices are characterized electrically and optical devices optically, but a PIC requires both to form an astonishing array of functions. For example, a single die can contain photodetectors, modulators, waveguides, optical couplers, lasers, transimpedance amplifiers, drivers, and high-speed electrical interfaces. Performance depends on the interaction among these elements, which means engineers have to evaluate the optical and electrical portions of the device together during development and, ultimately, production.

At the same time, electrical bandwidths continue to rise while optical coupling structures demand increasingly precise alignment. Finding problems before packaging therefore becomes particularly important.

FormFactor's CM300xi-SiPh wafer probe system integrates electrical and optical testing for photonic integrated circuits, enabling automated wafer-level characterization as data rates and optical alignment requirements continue to increase.

That puts much of the burden on wafer-level test. Packaging represents a significant portion of the cost of many photonic devices, and discovering a bad die after packaging is an expensive way to improve yield. Testing at the wafer level can identify failures earlier and provide valuable process information, but it also brings RF, optical, mechanical, and automation requirements together in the same test environment. That combination is what makes PIC characterization unusually challenging.

The Expanding RF Measurement Challenge

The RF side of the problem begins with bandwidth. The move to 224 Gb/s and eventually 448 Gb/s PAM4 increases the frequency range that must be characterized. In addition. harmonic content also influences signal integrity and the accuracy of device models, so meaningful characterization has to extend well beyond the fundamental frequency. Consequently, measurements that once ended below 70 GHz increasingly extend to 130 GHz, and harmonic characterization is moving toward 200 GHz.

One way to reach those frequencies is to combine several narrower-band measurement systems, but doing so introduces additional calibration steps and makes correlation between setups more difficult. A broadband system that maintains a common reference plane across the entire range can eliminate much of that complexity.

That requirement has driven coaxial interfaces and measurement hardware to millimeter-wave frequencies. Coaxial standards are now available at 168, 220, and 250 GHz, and wafer probes have followed them upward in frequency.

When Repeatability Becomes the Limiting Factor

At these frequencies, instrument bandwidth is only part of the measurement problem because repeatability is just as important. Small differences in probe landing position or overtravel may have little effect at lower frequencies, but at higher frequencies they can produce measurable changes in the result. Mechanical stability, touchdown consistency, and contact quality all begin contributing to uncertainty.

continued on page 22

continued from page 21

A probe that lands slightly differently from one device to the next can alter the measured response enough to affect the resulting device model. That becomes a serious problem when hundreds or thousands of measurements are being compared across devices, wafers, or test sessions.

For this reason, semi-automated and automated probing has become increasingly important in high-frequency characterization. Controlling the touchdown mechanically removes one source of operator-to-operator and measurement-to-measurement variation. The objective is not only to gather more measurement data at higher frequencies but also to ensure the measurements correlate well enough to be useful.

Calibration Has to Keep Up

Calibration techniques are evolving along with measurement frequency and device complexity. Multi-line TRL, for example, uses multiple calibration structures to maintain accuracy and characterize uncertainty over a broad frequency range. It requires more effort than simpler calibration approaches, but the additional complexity becomes increasingly worthwhile as bandwidth grows.

Differential interfaces add another layer to the problem. Many optical engines rely on differential electrical signaling to achieve high data rates, and crosstalk becomes increasingly troublesome as frequency rises. At that point, the interaction is not limited to the contacts within an individual probe because coupling between adjacent probes can also affect the measurement.

Multi-port calibration techniques, including 32-term and 64-term approaches, can correct for crosstalk that would otherwise degrade differential measurements. So, calibration is becoming less of a procedure performed after the measurement system is assembled and more of an integral part of the system design.

The Optical Alignment Challenge

Wafer-level optical measurements usually require positioning a fiber at a coupling structure on the die, either to launch laser light into the device or collect light coming from it. Alignment tolerances can be measured in microns or even fractions of a micron. Making such small movements is possible, but knowing where the fiber is relative to the optical structure is much more difficult.

PICs employ several coupling methods, including vertical couplers and edge-coupled interfaces. Some of these structures are difficult to view once the wafer is positioned for test. Trenches and V-grooves can place the coupling point where a conventional microscope provides little or no direct visual feedback.

Newer optical probes attack this problem by incorporating fiducials and alignment references into the fiber assembly. Machine-vision systems can locate these references and establish the fiber position without requiring a direct view of the actual coupling interface.

The progression in optical systems from 1.6 Tb/s to 3.2 Tb/s is pushing RF and optical characterization beyond what either discipline can handle independently.

Combined with precision positioning software, this approach can provide sub-micron alignment even when the optical interface itself is obscured. As channel counts increase and optical coupling tolerances continue to shrink, repeatable automated alignment becomes increasingly important.

Measurement Technology Moves to the Foreground

While optical engines, transceivers, and new photonic architectures get most of the attention, the measurement technology required to develop and manufacture them tends to stay in the background, and that's becoming a lot more difficult to do.

The progression in optical systems from 1.6 Tb/s to 3.2 Tb/s is pushing RF and optical characterization beyond what either discipline can handle independently. Broadband RF measurements, precision optical alignment, calibration, mechanical repeatability, and automated wafer probing increasingly have to operate as parts of the same test system.

PICs will carry a growing share of the bandwidth required by AI infrastructure. As their complexity and operating speeds increase, the ability to characterize them accurately at the wafer before the cost of packaging has been added will become as important as advances in the devices themselves.

• FORMFACTOR •

Hammer of the Gods Targets the Electromagnetic Battlespace

Image source: SGT Sar Paw, Photo courtesy of U.S. Space Force

Modern precision weapons depend on access to accurate positioning, navigation and timing (PNT). Department of War is now experimenting with an EW system intended to exploit that dependence while remaining deployable close to the fight. In July, the Army, Navy and Space Force conducted the first maritime trial of the system during Project Convergence Capstone 6 off Southern California.

Called Hammer of the Gods, the system is essentially a deployable EW emitter that transmits controlled RF energy to create effects in the electromagnetic spectrum potentially including disruption or manipulation of positioning, navigation and timing signals while keeping its own electromagnetic signature relatively low.

GPS and other satellite-navigation signals arrive at the receiver at extremely low power levels, making them inherently vulnerable to interference. A sufficiently capable electronic-attack system can deny, degrade or potentially manipulate those signals over an operational area, affecting navigation as well as the timing references used by communications networks, sensors and weapons.

That makes counter-PNT an increasingly important part of electronic warfare as precision-guided artillery, cruise missiles, unmanned aircraft and other networked weapons often rely on satellite-derived position or timing data in their operating chain. Disrupting that data can force platforms to fall back on inertial navigation or other alternatives that may accumulate error or provide less precise performance.

An emitter operating aboard or near naval platforms must function in an exceptionally crowded electromagnetic environment that can include high-power shipboard radars, communications systems, satellite links and electronic-support receivers. At the same time, minimizing the emitter's own detectable signature is critical because any powerful RF transmission can potentially reveal its location to hostile electronic-support systems.

The experiment also evaluated resilient mesh networking intended to extend the system's operational reach. Hammer of the Gods represents more than another EW transmitter because shows how spectrum control, PNT denial, low-probability-of-detection operation and resilient networking are converging into a single EW architecture. MPD

The FCC To Open More C-band Spectrum for Terrestrial Wireless

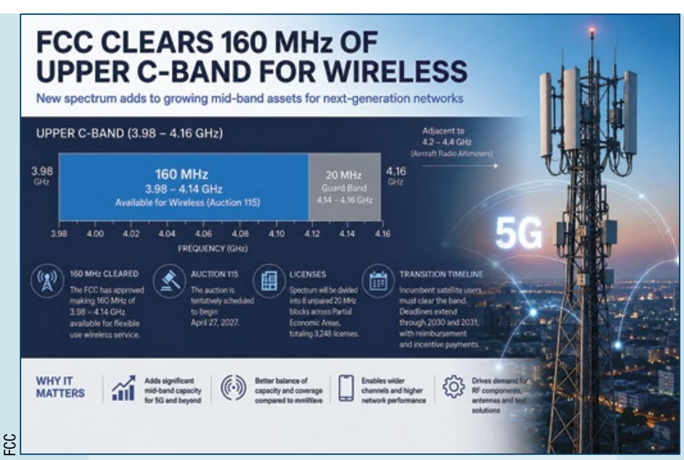

FCC

The Federal Communications Commission has approved a plan to make another 160 MHz of C-band spectrum available for terrestrial wireless service, one of the most consequential additions to U.S. mid-band capacity since the original C-band auction.

The new spectrum extends from 3.98 to 4.14 GHz and will be offered in Auction 115, tentatively scheduled to begin April 27, 2027. The FCC plans to divide the spectrum into eight unpaired 20-MHz blocks across Partial Economic Areas, producing 3,248 flexible-use licenses. Together with spectrum cleared in the earlier C-band proceeding, the action creates 440 MHz of contiguous mid-band spectrum for next-generation wireless networks.

That bandwidth is attractive because C-band occupies a useful compromise between coverage and capacity. It propagates considerably farther than millimeter-wave signals while providing much wider channels than many lower-frequency cellular allocations. For operators, additional spectrum near existing C-band deployments should also make it easier to expand capacity without introducing an entirely new RF ecosystem.

However, the transition presents significant engineering work as base-station radios, power amplifiers, antenna arrays, duplexing and filtering architectures will have to accommodate the expanded tuning range while controlling emissions near neighboring services. The FCC is retaining 20 MHz from 4.14 to 4.16 GHz as a guard band, an important consideration because aircraft radio altimeters operate immediately above the C-band allocation at 4.2 to 4.4 GHz.

Fixed-satellite-service operators use portions of the Upper C-band, so SES, Eutelsat and Telesat must migrate existing operations. The transition will involve new satellites and modifications to thousands of broadcast and cable facilities. Clearing deadlines extend into 2030 and 2031, with incumbent operators eligible for reimbursement and incentive payments.

Mid-band spectrum will remain important to 5G-Advanced and emerging 6G architectures, where massive-MIMO antennas and beamforming can extract substantial capacity from frequencies that still provide practical macro cell coverage. That means a sizeable market for wideband PAs, LNAs, filters, switches, antenna systems and test equipment. MPD

continued from page 3

compares with roughly one millisecond across the midhaul and approximately 10 milliseconds across the backhaul.

Those numbers leave little margin for error. Light travels through single-mode fiber at roughly five microseconds per kilometer, so the latency allowance quickly becomes a distance limit as well. Every switch, router, accelerator and other active device along the route consumes part of the remaining budget.

Synchronization is even less forgiving. In a TDD network, cells must remain aligned to within approximately ±1.5 microseconds at the antenna. Achieving that requires individual elements in the fronthaul timing chain to operate within tolerances measured in tens of nanoseconds.

That normally means deploying the full timing-support profile defined by ITU-T G.8275.1, with PTP-aware boundary or transparent clocks at every node and Synchronous Ethernet providing the underlying frequency reference. None of this is ordinary packet-network engineering. It is hard real-time systems work.

In an integrated base station, much of that responsibility belonged to one supplier. In an Open RAN deployment, the operator or its systems integrator must make equipment from several companies behave as though it came from one engineering organization. That transfer of responsibility is arguably the most important consequence of disaggregation.

It also appears in two areas that received far less attention during Open RAN's early promotion: power and security.

The power penalty is measurable. Open RAN equipment can still consume roughly 15% to 20% more energy than an equivalent integrated system. Some of that increase comes from the processing and transport overhead created by standardized interfaces. Some comes from virtualization itself.

A virtualized O-DU running Layer 1 functions on commercial off-the-shelf servers must handle demanding workloads such as forward-error correction and LDPC encoding and decoding. General-purpose processors are not especially efficient at these tasks compared with the purpose-built silicon used in a conventional baseband unit.

Hardware acceleration can close much of the gap, but that introduces another set of design decisions. Operators must choose between look-aside and inline accelerators, each of which has different implications for latency, power consumption, throughput and software complexity. Packetized fronthaul adds more overhead, and an O-Cloud built around general-purpose computing remains less efficient than fixed-function hardware in many operating conditions.

For an industry that increasingly treats energy per bit as a central design constraint, this is not a minor issue. In

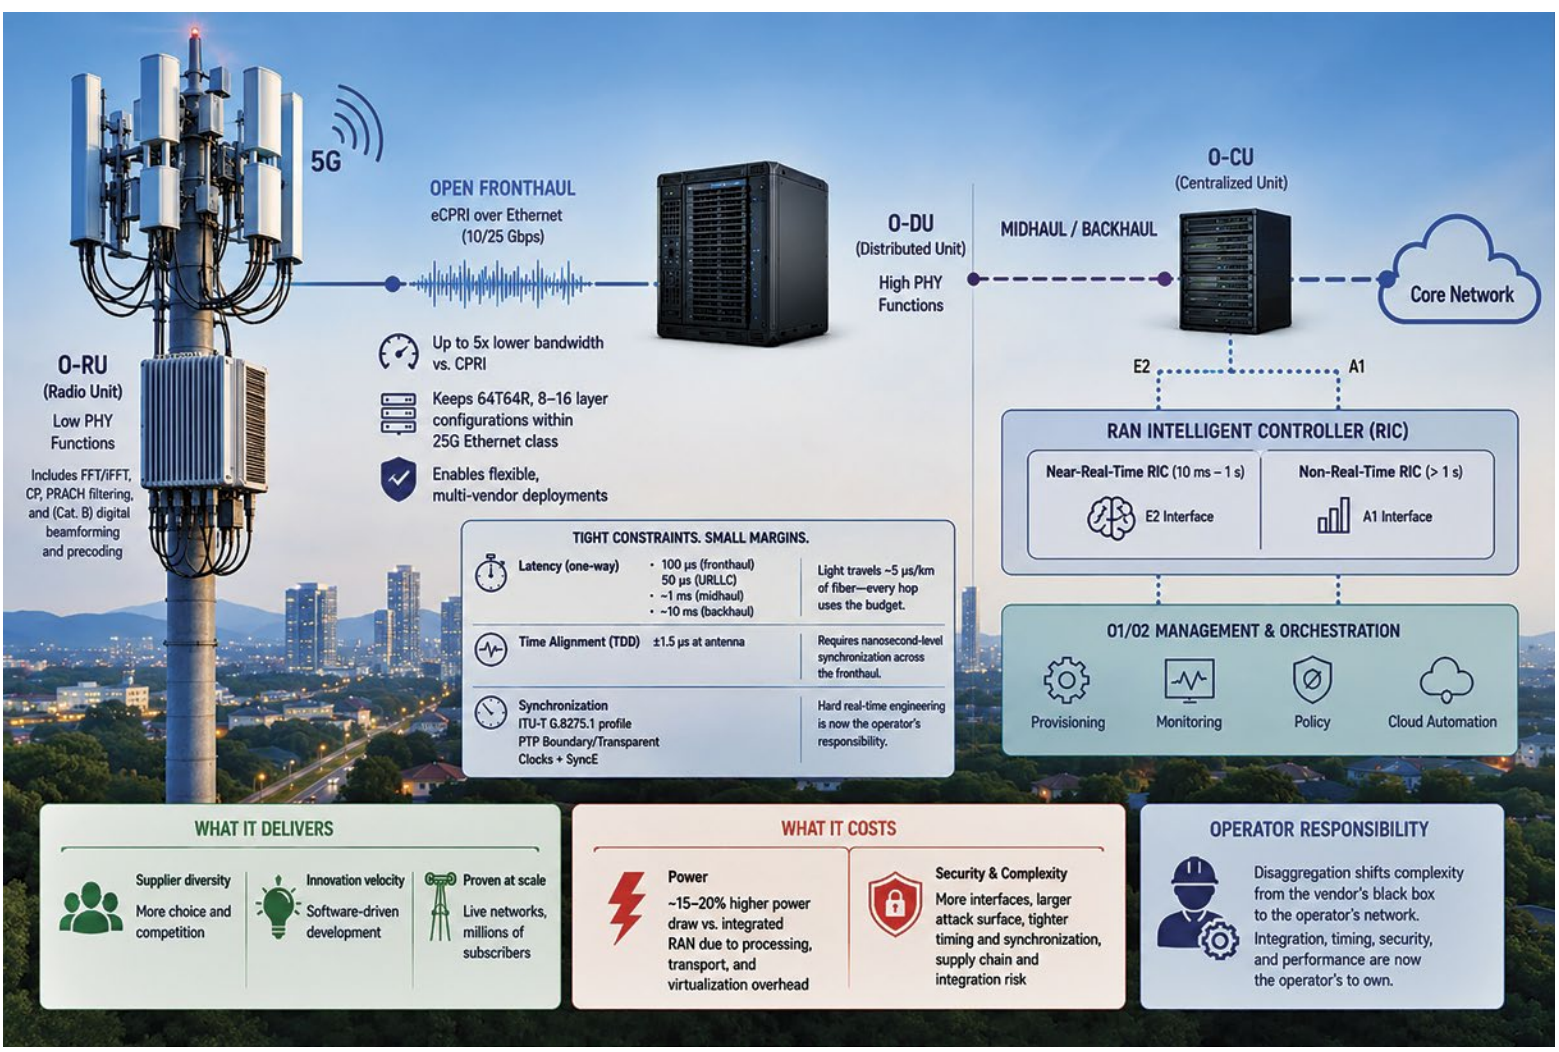

a dense network, higher power use can consume much of the savings that operators expected to receive from greater supplier competition.

Security presents an equally complicated tradeoff.

Disaggregation does not necessarily create complexity from nothing. It moves complexity into different parts of the system. The same is true of the attack surface. Every newly exposed interface becomes another boundary that must be authenticated, protected, monitored, and maintained.

These interfaces include the open fronthaul between the O-DU and O-RU, F1 between the DU and CU, E2 and A1 between the RAN and its near-real-time and non-real-time RAN Intelligent Controllers, and O1 and O2 into the management and cloud environments.

The open fronthaul alone includes control, user, synchronization and management planes, each with its own vulnerabilities. The synchronization plane is a particularly useful example because it shows how security and physical-layer performance can collide.

An attacker does not necessarily have to defeat encryption to disrupt a TDD network. Flooding the master clock with excessive PTP traffic may be enough to push synchronization beyond the ±1.5-microsecond alignment requirement. Once that happens, TDD operation can fail.

The management plane can also be exposed to interception or man-in-the-middle attacks when the fronthaul is inadequately protected. Encrypting that traffic seems like the obvious answer, but it comes at a price. MACsec and similar protections introduce latency and jitter into a transport path that may have only 100 microseconds to spare.

Fronthaul encryption is therefore not simply a security box that can be checked without consequence. It is a system-level engineering decision that must be balanced against timing and performance requirements.

The RAN Intelligent Controller introduces another concentration of risk. The near-real-time RIC operates over the E2 interface on control loops ranging from about 10 milliseconds to one second and can host third-party xApps. That makes it valuable operationally, but it also creates an attractive target.

Researchers have already identified numerous denial-of-service scenarios involving the RIC, individual cells and ordinary control protocols. Signaling-storm attacks, for example, can overwhelm network resources without relying on particularly exotic methods.

A multi-vendor network also inherits a broader supply-chain problem. Its security is only as strong as the weakest supplier, software component or development practice in the system. Once the RAN becomes software-defined and cloud-based, it acquires many of the same patching, dependency-management and software-supply-chain risks that have affected the wider IT industry for years.

This does not mean Open RAN is inherently insecure. The O-RAN Alliance has made substantial progress in security specifications, threat modeling and interface protection. A properly implemented zero-trust architecture is entirely achievable.

The distinction is that security does not arrive automatically with the architecture. It must be designed, tested and continuously maintained by the operator. A traditional base-station supplier delivered much of that work inside an integrated product. Open RAN gives the operator more flexibility, but it also hands back responsibility for verifying that numerous independent implementations behave correctly at every boundary.

This helps explain what can otherwise look like a contradiction. Open RAN continues to attract major commercial commitments, yet operator surveys often show that it has fallen in priority compared with the enthusiasm of several years ago.

Both trends can be true.

The operators making the most progress are generally those that treated Open RAN as a serious systems-integration program. They brought interoperability, timing, cloud operations and security expertise in-house and funded those capabilities accordingly.

Operators that expected standardized interfaces to produce an almost plug-and-play marketplace have had a different experience. They found that the interfaces solved only part of the problem and that the remaining integration burden was considerably heavier than advertised. Some have responded by slowing deployments, narrowing their ambitions, or relying more heavily on a lead vendor.

For the RF and microwave industry, the lesson is straightforward: an interface specification is the beginning of the engineering process, not the end.

A radio unit can conform to a specification and still fail to interoperate smoothly with another supplier's distributed unit. Two products can interoperate and still fail to meet nanosecond-level synchronization requirements across a multi-vendor timing network. Even a properly synchronized system may still consume too much power or expose unacceptable security risks.

Meanwhile, the O-RU must continue to solve the same RF problems it always faced. Its GaN Doherty amplifiers must meet linearity and efficiency targets. Digital predistortion and crest-factor reduction must allow those amplifiers to operate close to saturation without degrading signal quality. A 64-element array must be calibrated accurately and kept calibrated as operating conditions change.

The difference is that the radio now performs those tasks as one element in a much larger multi-vendor system. Its timing, transport and security constraints are increasingly determined outside its enclosure.

Open RAN did deliver the supplier diversity it promised. What it did not do was eliminate the hardest engineering problems. It moved many of them from the equipment vendor's integration laboratory into the operator's network.

The industry would have set more realistic expectations if it had said that plainly from the beginning.