



Synchronization Methods for Multi-Detector Phased Systems

Tatiana Krasik

ATE System Developer Israel, Petah Tikva

OPEN ACCESS

SUBMITTED 19 May 2025

ACCEPTED 14 June 2025

PUBLISHED 01 July 2025

VOLUME Vol.07 Issue 07 2025

CITATION

Tatiana Krasik. (2025). Synchronization Methods for Multi-Detector Phased Systems. The American Journal of Engineering and Technology, 7(07), 01–08. <https://doi.org/10.37547/tajet/Volume07Issue07-01>

COPYRIGHT

© 2025 Original content from this work may be used under the terms of the creative commons attributes 4.0 License.

Abstract: This article examines synchronization methods for multi-detector phased systems that integrate spatially distributed transmit–receive nodes into a single coherent structure. The study's primary aim is to determine the technical requirements for temporal, frequency, and phase alignment of the elements, and to analyze the hardware and algorithmic means for achieving them. The relevance of this work is driven by the rapid development of phased arrays and distributed radar and astronomical systems, where even tens of picoseconds of desynchronization lead to significant loss of coherent gain and degradation of spatial resolution. Contemporary network protocols such as IEEE 1588 provide only microsecond-level accuracy, which is insufficient for the often-required budgets on the order of tens of picoseconds; therefore, a multi-level architecture is necessary, combining highly stable reference oscillators, zero-delay hardware buffers, deterministic data-transfer interfaces, and digital correction algorithms. The novelty of this research lies in the comprehensive comparison and integration of four classes of solutions: a distributed clock tree with LVDS and fiber-optic lines and zero-delay PLL buffers; deterministic SYSREF frame distribution according to JESD204B/C; bidirectional microwave wireless exchange with pilot-tone synchronization; and digital corrections via cross-correlation and Kalman-consensus algorithms to compensate residual drifts. A methodology for budgeting phase slip—accounting for source jitter, port trace dispersion, and network delays—is presented, enabling early identification and elimination of design bottlenecks. The key conclusion demonstrates the effectiveness of the multi-level scheme: an external hardware-network loop provides coarse phase alignment and frequency stability at the level of single

to tens of picoseconds. In contrast, the internal digital loop maintains instantaneous coherence with phase errors of only a few degrees, even when nodes are separated by hundreds of meters or during GNSS outages. Systematic summation of contributions from jitter, trace skew, and network delays guarantees $\geq 90\%$ coherent gain and the specified dynamic range. This article will be helpful to engineers developing phased antenna arrays, distributed radar, and interferometric systems, as well as researchers in precise frequency-time distribution.

Keywords: multi-detector phased system; synchronization; temporal alignment; phase coherence; JESD204B/C; zero-delay PLL; microwave exchange; Kalman-consensus.

Introduction

A multi-detector phased system is an ensemble of spatially separated transmit-receive cells in which the resultant wavefront is formed by coherent summation of signals. For summation to be coherent, each node must operate under a standard time base, maintain the same carrier frequency, and know its phase offset relative to the others. This triad of controls makes synchronization a key architectural element: it transforms a set of independent detectors into a unified electrical aperture capable of steering a narrow beam, constructing interferograms, or measuring the thermal balance of a celestial source.

In practice, three main performance metrics are distinguished. First, temporal alignment Δt : for packet-based networks such as IEEE 1588, this is measured in microseconds and is already sufficient for LTE base stations to share time-division slots [1], but for phased arrays, such dispersion is unacceptable. In that domain, White Rabbit—an extension of PTP—ensures synchronization better than one nanosecond with picosecond-level stability over kilometer-scale links [2]. Experiments with distributed antennas show that to achieve $\geq 90\%$ coherent gain at a symbol rate of 1 Gbaud, Δt must be < 67 ps [3], and for operation at centimeter wavelength, designers often allocate an even tighter budget as an engineering rule.

Second, the frequency stability of the reference oscillator in practice determines $\Delta f/f$. GNSS-disciplined rubidium standards supplied as GNSS-DO modules deliver relative instability on the order of 1×10^{-12} and

an RMS deviation to UTC of about 10 ns [4], which is sufficient even for low-band radio interferometry.

Third, phase accuracy $\Delta\phi$: analytical estimates indicate that a loss of coherent gain of 0.5 dB occurs at a random phase error of approximately 18° [3].

Any deviation beyond these boundaries immediately impacts system performance. With differing phases, the elements sum non-constructively, energy leaks from the main lobe into sidelobes, and sensitivity drops. For example, a calculation for an X-band radar with 5-bit phase shifters showed that a combination of 3° RMS phase and 0.5 dB amplitude ripple yields only -0.15 dB of gain loss. In contrast, larger drift rapidly broadens the beam and raises sidelobe levels [5]. If node frequencies diverge beyond the discipline of the rubidium standard, the pattern shifts toward pseudo-random dephasing and, in the extreme, interferometric information is lost. A temporal desynchronization of hundreds of picoseconds causes a wideband LFM probe to lose up to one-quarter of its overlap, forcing the detector to operate with a degraded signal-to-noise ratio. Thus, the strict values outlined above are not mere technical standards: they define the narrow region in which a multi-detector phased system remains coherent and delivers its designed dynamic range.

MATERIALS AND METHODOLOGY

The materials and methodology for studying synchronization in multi-detector phased systems are based on a comprehensive analysis of 18 key sources grouped into four thematic blocks.

In the requirements specification, the standards for three metrics are first presented: temporal alignment Δt (< 1 ns for IEEE 1588 and < 67 ps for coherent gain in phased arrays) [1–3]; frequency stability $\Delta f/f$ and UTC alignment to within approximately 10 ns [4]; and the allowable phase error $\Delta\phi$ based on calculations for X-band radars [5].

In the hardware loop, distribution of the reference clock across boards and cables relies on LVDS and coaxial routing guidelines, where a length tolerance of ± 5 mm yields intra-pair skew ≤ 30 ps [6, 7]. To eliminate residual jitter, zero-delay PLL buffers (LMK04816, CDCVF25081) with intrinsic RMS noise on the order of 100 fs are employed [8, 9], and deterministic delivery of ADC-FPGA output data is ensured by the JESD204B/C

standard with SYSREF marking and latency \leq tens of picoseconds [10].

Among the wireless and optical methods, bidirectional microwave exchange at 5.8 GHz achieves a dispersion of 2.26 ps over a 0.9 m baseline and 12 ps of group delay after application of a consensus algorithm [11, 12]; the RAG approach in SAR systems uses a pilot tone in the operational channel to provide residual phase instability $< 0.1^\circ$ at SNR > 60 dB [13]. For astronomical interferometry, VLBI calibration results demonstrate an RMS phase of 0.8 ps on the Kokee baseline [14], and for spaceborne BiSAR, phase accuracy of 0.1° is achieved using a synchronization pulse [15].

Additional sources detail PCB trace-design methodologies (accounting for 62 ps of inductive difference over 31 cm of FR-4) [16], principles of synchronization in Time-Sensitive Networking and PTP systems [17], and simulation results for dTE delays in IEC/IEEE 60802 networks with variable message

intervals [18].

RESULTS AND DISCUSSION

Hardware synchronization is built around a single primary frequency, which must be delivered to all detectors without appreciable noise or drift. The most straightforward approach is to distribute it via a classical clock tree. Over short distances, LVDS pairs or coaxial cable are used; a 1 mm deviation between differential traces already produces approximately six ps of mismatch, so in multi-gigahertz arrays a maximum length difference of ± 5 mm is permitted—this keeps intra-pair skew within ± 30 ps and does not violate the phase budget described earlier [6, 7]. Figure 1 illustrates the concepts of intra-pair and inter-pair skew in LVDS differential-pair routing. For racks separated by tens of meters, the same reference is carried over fiber optic or actively compensated coaxial cables; the resulting jitter is then limited primarily by the oscillator and does not exceed hundreds of femtoseconds.

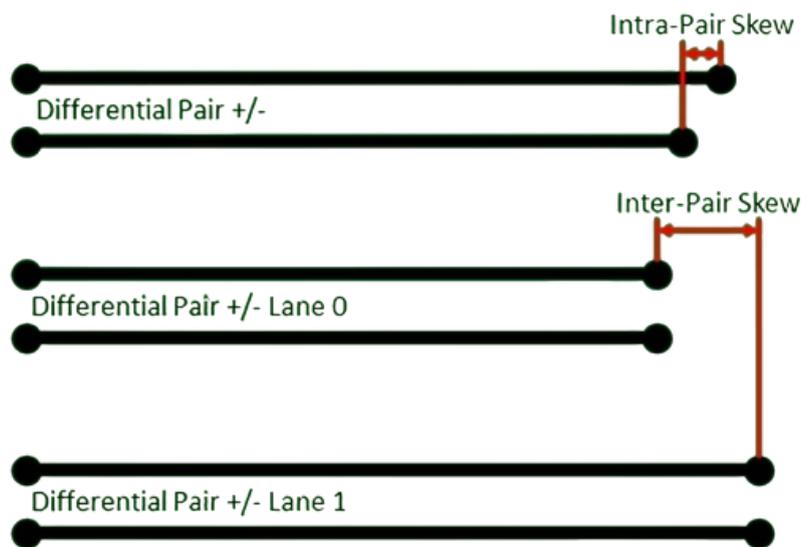


Fig. 1. Intra-Pair and Inter-Pair Skew [7]

A perfectly matched cable cannot guarantee phase equality upon power-up; here, zero-delay PLL buffers are employed. These devices inject their divider into the feedback loop, forcing the output phase to coincide with the input phase on every cycle. A typical IC—e.g., the LMK04816—adds only 100 fs rms of intrinsic noise [8], while the more cost-effective zero-delay driver CDCVF25081 maintains the same 100 fs with a specified inter-output skew of 150 ps [9]. Experience shows that, after a one-time calibration, such a cascade provides sub-picosecond relative stability at the board level—

sufficient to keep phase error within a few degrees at frequencies up to tens of gigahertz.

When the ADC and FPGA reside on separate dies, the frequency-distribution task is compounded by the need for deterministic data latency. The JESD204B/C interface addresses this in hardware: the standard clock is sent over a dedicated differential pair, and a SYSREF pulse marks the start of the multi-frame interval, after which logic synchronizes the LMFC counters in both transmitter and receiver, as shown in Figure 2. The standard mandates that the total variation in latency

across all lanes and devices must not exceed one LMFC period; this implies that fundamental uncertainty remains within a few tens of picoseconds at gigahertz clock rates, and SYSREF skew within a rack must fall under the allocated system budget [10]. Because all

lanes share a common carrier, resultant phase drift reduces solely to the source's spectral jitter, and phase coherence between detectors is restored immediately upon reception of each SYSREF pulse.

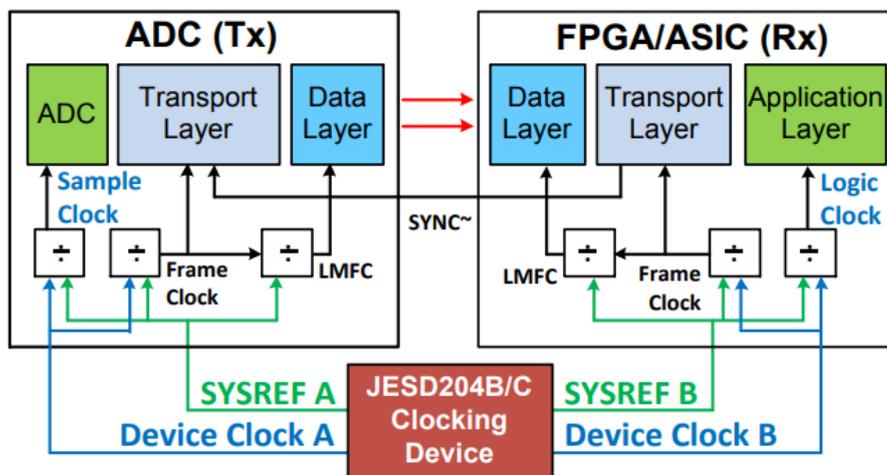


Fig. 2. JESD204B clocking and data interface [10]

Thus, the hardware scheme employing a single generator hierarchically satisfies the three metrics outlined in the previous section: time is fixed by trace length and SYSREF delay; a low-noise PLL combiner sets frequency; and phase is held aligned within each clock cycle. This rigid foundation allows networked or digital methods to clean up residual picoseconds rather than contend with microseconds.

When detectors are separated beyond the reach of a single fiber-optic or Ethernet bus, phase alignment is carried out over the air via a reference RF signal. The bidirectional microwave-link concept relies on exchanging short packets between each node pair and computing delay as half the round-trip time. The symmetry of the path automatically cancels fixed asymmetries, so the final error is determined only by oscillator instability and channel noise. In a laboratory demonstration at 5.8 GHz, a single two-tone 40 MHz packet yielded 2.26 ps of dispersion over a 0.9 m baseline, equivalent to 0.7° of phase spread at 10 GHz [11]. The same method, embedded in a wireless four-node array and augmented with a consensus algorithm,

reduced group skew to 12 ps with a standard deviation of 3 ps, converging in fewer than twenty exchange iterations [12]. Over hundreds of meters, accuracy degrades slightly due to multipath effects but remains within tens of picoseconds, sufficient for coherent centimeter-wave imaging; for kilometer-scale links, radar or optical telemetry is added to assess residual asymmetry.

An alternative approach is to transmit a continuous pilot tone or pulse markers alongside the operational signal. The reference-tone generator is inserted into the transmitter path; at the receiver, it is extracted via a narrowband filter and compared to the local reference, after which a digital phase detector computes the correction. In multistatic SAR, this scheme (Fig. 3), particularly in the pulse-alternate mode, demonstrated that for SNR > 60 dB, the standard deviation of the residual phase error does not exceed 0.1° [13]. Its advantage is the absence of a separate synchronization channel; its drawback is that the pilot consumes time-frequency resources and requires the receiver to always remain within line-of-sight of the transmitter.

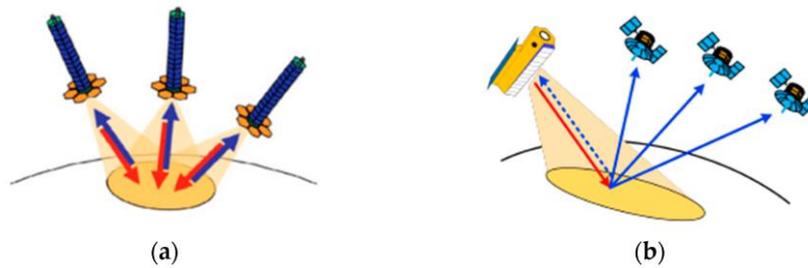


Fig. 3. The type of multistatic SAR system [13]

(a) Fully active SAR system; (b) Semi-active SAR system.

Both radio-frequency loops fit into the multi-level architecture described above: networked PTP or White Rabbit provide the common time coarsely, bidirectional microwave exchange removes dynamic drift on long baselines, and the pilot tone preserves instantaneous phase within the radar signal. Such sequential error filtering permits coherence maintenance even for mobile or widely distributed detectors, without resorting to expensive atomic clocks at every node.

The algorithmic loop engages once hardware and network methods have already reduced dispersion to single-to-tens-of-picoseconds levels, but residual slow drift and noise tails remain. The most straightforward way to eliminate these is to periodically observe a source of known phase and adjust all local oscillators via cross-correlation. In astronomical VLBI, this role is fulfilled by a bright calibrator: the correlator measures phase differences between antenna pairs and converts them into delay, after which a low-order polynomial smooths atmospheric and clock fluctuations. At 8 GHz X-band, the median phase RMS on the 12 m KOKEE baseline was only 0.8 ps, with 97% of solutions within two ps, equivalent to $< 3^\circ$ at a 10 GHz carrier [14]. In the spaceborne BiSAR LuTan-1, the calibrator is illuminated by the transmitter itself: the synchronization pulse occupies the same frequency window as the imaging signal, and after pulse compression, its phase is subtracted from the echo. With SNR > 60 dB, the standard deviation of the residual phase error fell below 0.1° —almost an order of magnitude tighter than the hardware budget [15].

However, a periodic calibrator cannot correct drift faster than its tagging interval. Here, distributed filters combining exchange of local estimates and Bayesian fitting of their dynamics are employed. When averaging is augmented with a Kalman filter, the model tracks both long-term oscillator drift and measurement noise: the

same study shows that the KF-DFPC algorithm converges twice as fast as DFPC alone and retains accuracy even when the exchange interval is shortened or SNR decreases. ODKF diffusion yields comparable MSE improvement for small and medium-sized groups while requiring only local connectivity, making it well-suited to wireless radar networks.

These digital procedures reside in the internal loop, while the external loop remains in the hardware-network layer. The first PPS from GNSS or White Rabbit sets the epoch; zero-delay PLLs distribute it across the board; then cross-correlation or consensus-Kalman phase-aligns each detector pair to the task-specific boundaries (tenths of degrees for SAR interferometry, single-degree precision for communications). This dual-loop scheme thus divides responsibilities: the coarse level ensures deterministic startup and thermal stability, while the acceptable level provides instantaneous coherence, enabling the system to remain phase-connected even during GNSS outages or rapid transmitter temperature shifts.

The accuracy achieved by hardware-network layers establishes the upper bound on residual phase errors. Still, the total error must be decomposed into elementary contributions and recombined quadratically to ensure the system remains within a few degrees of the operational carrier. First in the budget is conductor length variation: for internal differential pairs, weave-induced FR-4 induction can produce up to 62 ps of difference over a 31 cm trace [16]. Next is the source's intrinsic jitter. A typical dual-loop PLL combiner, such as the LMK04816, specifies 100 fs rms in the 12 kHz–20 MHz band [8]. Lower-cost buffers exhibit increased spread but remain within a few hundred femtoseconds, provided the capture loop remains narrow and the reference source is clean.

In network channels, the main systematic component is round-trip asymmetry. In a typical PTP deployment,

profile G.8275.1 allocates 50 ns of fixed error per boundary clock, limiting the maximum chain depth for phase-sensitive services [17]. If a node relies on generalized PTP rather than White Rabbit, the phase margin must be increased proportionately to the number of hops, or a local corrector must be introduced.

Practical implementation begins with the reference source. Stationary complexes typically use a discipline-GNSS module; mobile and field systems adopt rubidium standards with GNSS disciplining, where daily hold-over remains within a manageable number of nanoseconds, keeping the network within the PTP budget. On each board, SYSREF or PPS triggers are routed strictly by isothermal lines and captured by hardware timestamping so that all JESD204 logic starts from the same epoch. Ethernet switches for the timing backbone must support hardware time-stamping and boundary-clock mode.

Even the most precise static alignment degrades over time, so FPGAs implement automatic self-calibration: a calibration pilot tone or reversed microwave packet measures relative skew and injects a correction into the digital mixer every few seconds; interferometry practice has shown this suffices to keep phase drift below 0.1° between reference frames. For increased robustness, the reference is deployed redundantly: one GNSS-DO operates actively, a second remains hot-standby, and upon satellite signal loss, both switch to hold-over, controlling the phase-detector switch to avoid jumps exceeding one ns when changing sources. Such redundancy, and a PLL-failure detector, provides time to repair an external antenna without disrupting array coherence [18].

During the author's tenure at Ceragon, the author developed software for the automatic calibration of a nonlinear phase detector (NLPD) on ATE benches designed for the IP-50EX radio module. The IP-50EX itself is positioned by the company as the flagship of its E-band platform: its high throughput is combined with a compact form factor and advanced circuit design, rendering the product highly in demand in modern telecommunications networks. The developed calibration module enhances the repeatability of the phase characteristics and, consequently, the synchronization quality of such systems. The significance of this product is clearly reflected in Ceragon's financial statistics: in the fourth quarter of 2024, the company

recorded a quarterly revenue of USD 106.9 million, representing an 18.3% increase compared to the same period in the previous year; the decisive factor was sales of the IP-50EX, particularly in the Indian market, where quarterly revenue reached USD 55.6 million—the highest level in the company's history.

In summary, budgeting should proceed bottom-up: select the reference, then set limits on clock-tree routing and jitter; verify network symmetry or switch to White Rabbit; control the temperature of long cables; and sum only noise contributions. If the total exceeds three to five degrees of phase slip, enable cross-correlation or Kalman-consensus, and always design in a margin for unforeseen hardware and environmental factors.

Combining hardware-network measures with digital algorithms enables a reliable multi-level mechanism for maintaining phase coherence, from reference selection and stringent clock-tree routing control to distributed calibration leveraging cross-correlation and Kalman-consensus. The clear separation of responsibilities between the external loop ensures deterministic startup and thermal stability, and the internal loop provides instantaneous adjustment and permits phase connectivity to be preserved even during GNSS outages or rapid temperature changes. Systematic evaluation of contributions from jitter, trace skew, and network delays enables timely identification of bottlenecks and rapid expansion of the phase-slip margin.

CONCLUSION

This study examined the key synchronization requirements for multi-detector phased systems, including temporal alignment (Δt), frequency stability ($\Delta f/f$), and phase accuracy ($\Delta\phi$). It was determined that to ensure $\geq 90\%$ coherent gain at a 1 G-band modulation rate, the spread of temporal delays must be maintained on the order of tens of picoseconds, and phase errors must be confined to single-digit to tens-of-degrees constraints fully defined by the phase-slip budget. The frequency component is addressed by high-stability GNSS-DO rubidium references with a relative instability on the order of 10^{-12} , enabling an RMS deviation to UTC of approximately 10 ns. Thus, the strict numerical limits established at the outset of this article serve not merely as guidelines but as the foundation within which the system preserves its specified dynamic range.

Hardware synchronization methods rely on hierarchical

distribution of a single primary oscillator via a clock tree: from LVDS pairs and coaxial traces at the board level to active fiber-optic and thermally compensated cables over tens of meters. Zero-delay PLL buffers, such as the LMK04816 or CDCVF25081, contribute no more than a few hundred femtoseconds of intrinsic jitter, providing sub-picosecond relative stability. The JESD204B/C interface, with its SYSREF marker, guarantees deterministic data latency between ADC and DSP on the order of tens of picoseconds, allowing the hardware loop to fully assume coarse synchronization of time, frequency, and phase.

For more remote detectors, wireless approaches are employed: bidirectional microwave delay exchange and pilot-tone transmission in the signal channel. Exchanging short packets at 5.8 GHz achieves a dispersion of approximately 2.26 ps over a 0.9 m baseline (0.7° at 10 GHz), and an embedded consensus algorithm within the array reduces group skew to 12 ps over tens of iterations. The pilot-tone variant demonstrates residual phase instability as low as 0.1° at SNR > 60 dB without requiring additional transmission channels.

The digital (algorithmic) loop complements hardware-network measures: periodic cross-correlation with a bright calibrator or synchronization pulse corrects slow drift down to single-picosecond levels, while distributed Kalman-consensus filters accelerate convergence and maintain accuracy under reduced SNR or shortened exchange intervals. This dual-loop architecture delineates responsibilities: the external level guarantees deterministic startup and thermal stability, and the internal level ensures instantaneous phase coherence.

In summary, the proposed multi-level mechanism—integrating hardware, network, and digital-algorithmic measures—provides the necessary phase connectivity for a multi-detector system even when hundreds of meters separate nodes or experience GNSS outages. Systematic analysis of jitter contributions, trace skew, and network delays enables the timely identification of bottlenecks and the expansion of the phase-slip margin. The author's experience has proven the efficiency of the discussed methods. This ensures the realization of the designed dynamic range and high reliability in field and laboratory environments.

REFERENCES

IEEE 1588 for Frequency, Phase, and Time Distribution, *Nokia*. <https://documentation.nokia.com/acg/23-7-2/books/classic-cli-part-i/c072-acg-ieee-1588-fp-td.html> (accessed Apr. 18, 2025).

M. Lipinski, The White Rabbit Project, *White Rabbit Official CERN website*, Feb. 19, 2024. <https://white-rabbit.web.cern.ch/> (accessed Apr. 18, 2025).

J. Merlo, A. Schlegel, and J. Nanzer, High Accuracy Wireless Time-Frequency Transfer For Distributed Phased Array Beamforming, *Merlo*. https://merlo.io/assets/pdf/ims23_distributed_time_frequency_transfer_and_beamforming_v1.1.pdf (accessed Apr. 20, 2025).

Pendulum Instruments introduces a traceable Rubidium Frequency Reference, *Pendulum Instruments*, Oct. 16, 2024. <https://pendulum-instruments.com/pendulum-instruments-introduces-a-traceable-rubidiumfrequency-reference/> (accessed Apr. 21, 2025).

M. Sarcione, Phased Array Radar (PAR) Performance Considerations - Impact of Errors, *Raytheon*, 2009. https://www.icams-portal.gov/resources/ofcm/mpar/meetings/symposium/workshop/W3_Sarcione%20Phased%20Array%20Beamforming.pdf (accessed Apr. 23, 2025).

C. Pearson, Clock Skew in Large Multi-GHz Clock Trees, *Analog Devices, Inc.*, 2019. <https://www.analog.com/en/resources/analog-dialogue/articles/clock-skew-in-large-multi-ghz-clock-trees.html> (accessed Apr. 23, 2025).

High-Speed Layout Guidelines for Reducing EMI for LVDS SerDes Designs, Texas Instruments, 2018. Accessed: Apr. 24, 2025. [Online]. Available: <https://www.ti.com/lit/an/snla302/snla302.pdf?ts=1747587823015>

LMK04816, *Texas Instruments*. <https://www.ti.com/product/LMK04816> (accessed Apr. 24, 2025).

CDCVF25081, *Texas Instruments*. <https://www.ti.com/product/CDCVF25081> (accessed Apr. 25, 2025).

T. Toroni, JESD204B/C Clocking, *Texas Instruments*. <https://www.ti.com/content/dam/videos/external-videos/ko-kr/7/3816841626001/6168424753001.mp4/subassets/clocks-and-timing-clocking-jesd204bc-systems->

[presentation-quiz.pdf](#) (accessed Apr. 26, 2025).

J. M. Merlo, S. R. Mghabghab, and J. A. Nanzer, Wireless Picosecond Time Synchronization for Distributed Antenna Arrays, *IEEE Transactions on Microwave Theory and Techniques*, vol. 71, no. 4, pp. 1720–1731, Apr. 2023, doi: <https://doi.org/10.1109/tmtt.2022.3227878>.

N. Shandi and J. M. Merlo, Decentralized Picosecond Synchronization for Distributed Wireless Systems, *Arxiv*, 2023. <https://arxiv.org/html/2405.18384v1> (accessed Apr. 28, 2025).

Q. Lin, S. Li, and W. Yu, Review on Phase Synchronization Methods for Spaceborne Multistatic Synthetic Aperture Radar, *Sensors*, vol. 24, no. 10, p. 3122, May 2024, doi: <https://doi.org/10.3390/s24103122>.

A. E. Niell *et al.*, VLBI measurement of the vector baseline between geodetic antennas at Kokee Park Geophysical Observatory, Hawaii, *Journal of Geodesy*, vol. 95, no. 6, May 2021, doi: <https://doi.org/10.1007/s00190-021-01505-9>.

Y. Cai *et al.*, An Advanced Approach to Improve Synchronization Phase Accuracy with Compressive

Sensing for LT-1 Bistatic Spaceborne SAR, *Remote Sensing*, vol. 14, no. 18, p. 4621, Sep. 2022, doi: <https://doi.org/10.3390/rs14184621>.

L. Ritchey, Confused About Differential Signaling or Clocks?, *Altium*, Jan. 14, 2019. https://resources.altium.com/p/confused-about-differential-signaling?utm_source=chatgpt.com (accessed May 01, 2025).

K. Sheno *et al.*, Tutorial on Synchronization: A Key Function in Time-Sensitive Networking and Beyond, *IEEE*, 2021. https://grouper.ieee.org/groups/802/802_tutorials/2021-03/WSTS-IEEE802-1-tutorial-Synchronization-2Mar2021.pdf (accessed May 02, 2025).

G. M. Garner, New Simulation Results for dTE for an IEC/IEEE 60802 Network, with Variable Inter-message Intervals, *IEEE*, Jun. 2021. Accessed: Apr. 18, 2025. [Online]. Available: <https://grouper.ieee.org/groups/802/1/files/public/docs2021/60802-garner-single-replic-simul-results-variable-intermsg-intervals-0621-v01.pdf>