

Native integrated sensing and communication on one coherent optical beam using rigid waveform manifolds: system architecture, synchronization, whole-unit decision, keying, beam steering and control interfaces

Type: technical disclosure (defensive publication)

Author / publisher: V131 Lab

Date of first publication: 2026-10-16

Field: coherent optical ranging (FMCW and related), free-space optical communication, integrated sensing and communication (ISAC), unmanned aerial vehicles, autonomous vehicles, robotics, swarms.

This disclosure describes a family of systems and methods. Every numbered paragraph describes one technique together with alternative forms of it. Unless a paragraph says otherwise, any technique may be combined with any other, and every listed alternative, range and combination is intended to be disclosed.

1. Rigid manifold as the transmission unit

1.1 A transmitter emits a sequence of transmission units. Each unit (a "manifold") is a fixed-length sequence of N complex samples taken from a predetermined waveform book shared in advance by transmitter and receiver. N is fixed for a given system. Examples: N between 32 and 65,536, for example 256, 512, 788, 1,000, 1,024, 2,048 or 4,096.

1.2 The manifold is treated as one rigid unit. Every sample keeps a fixed relative coordinate (index, time and phase) within the unit. Reception, decision, identification, feedback and configuration changes all act on whole units. Units are not split into sub-units that are decided separately, and samples from different units are not combined into one decision.

1.3 The book may be any complex waveform: a linear or nonlinear frequency sweep (chirp), a stepped or spline-shaped frequency or current trajectory, a phase-coded sequence, a pseudo-random complex sequence, a multi-tone set, a band-limited random waveform, or any combination of these. The book may be real-valued (one drive channel) or complex-valued (I and Q drive channels, or two polarizations).

1.4 Sample rate examples: 10 MS/s to 100 GS/s. Unit duration examples: 100 ns to 10 ms.

2. Payload is the manifold (native ISAC)

2.1 The same transmitted manifold serves sensing and communication at the same time, on the same beam, the same wavelength and the same time interval. Sensing and

communication are not separated by time, frequency, code, polarization, beam or operating mode.

2.2 Payload is carried by the geometry of the whole manifold: the transmitter selects one manifold state out of a registered set of whole-manifold states, and the selection is the payload. The receiver decides which whole-manifold state arrived.

2.3 The registered states may be formed by perturbing the book with a pattern defined over the whole manifold index. The pattern may be a weighted sum of K basis functions (K from 1 to N), each weight taking one of L levels (L from 2 to 64). Example basis functions: cosine, sine, Legendre, Chebyshev, Walsh/Hadamard, wavelet, prolate-spheroidal, principal components of the book, or arbitrary orthogonal or non-orthogonal sets. The perturbation may act on phase, amplitude, instantaneous frequency, or the frequency or current trajectory. The basis function with constant value over the manifold may be left without payload.

2.4 A bound on the largest perturbation of any state (a domain bound) sets the trade between bits per manifold and tolerance to device and channel impairments. Examples: maximum phase deviation 0.1 to 3.14 rad, for example 0.85 rad or 1.57 rad.

2.5 Several capacity levels may be defined on the same book and selected at run time. Examples: 1 to 4,096 bits per manifold, for example 12, 24, 60, 120, 240 or 480 bits per manifold. A lower level is used for long range, poor channels or low-cost components, and a higher level for short range or better components.

2.6 Error-correcting codes (for example Golay, BCH, Reed–Solomon, LDPC or polar) may be applied across one manifold or across several manifolds. A confidence gate may withhold delivery of a decoded unit whose fit to the decided state is poor.

2.7 For sensing, the echo of the transmitted manifold from a target is received. Because the receiver knows which state it sent, range and the target's kinematic state are obtained from the same echo that a remote receiver decodes as payload. For communication with another unit, the remote unit receives the direct light, decides the payload and also obtains the sender's kinematic state from that same direct light. Self-echo ("seeing") falls off as $1/R^4$ and direct light ("hearing") as $1/R^2$, so units can hear each other at longer range than they can see each other.

3. Synchronization nails: common lock nail and identity nail ("Firstnail")

3.1 At a fixed index of every manifold (for example index 0), every unit in the system adds extra optical intensity: the common nail. The nail is applied as a multiplicative gain on the drive signal or optical envelope at that index. The digital book samples and their encoding are unchanged. The common nail gives frame position, timing, phase reference and frequency reference for the whole manifold, and is kept for the whole session.

3.2 Each unit also has an identity nail at a registered index $k = n$ ($1 \leq n \leq N-1$), optionally with a registered phase ϕ . A receiver identifies the transmitting unit from the index and phase of the identity nail, without a question-and-answer exchange or an identity message. Identity is accepted after the same identity is seen in M consecutive manifolds (M from 1 to 16, for example 2).

3.3 Procedure: (a) lock to the common nail; (b) decide the whole manifold (section 4); (c) confirm identity from the identity nail; (d) at a whole-manifold boundary, remove the extra intensity of the identity nail and keep the common nail.

3.4 Capture ladder: before lock, the transmitter raises the common-nail gain step by step. Examples: $\times 1.25$, $\times 2$, $\times 4$, $\times 5$, $\times 6$, $\times 6.5$, $\times 8$ or $\times 10$, each step lasting 1 to 64 manifolds, with steps spaced 1 to 1,000 manifolds apart. The receiver latches the sample index of the largest received amplitude (or largest $|y|^2$) above a capture threshold, then holds a flywheel of N samples. After lock, the transmitter lowers the nail to a nominal gain. The flywheel and the whole-manifold decision keep the lock, so the lower nail does not lose it.

3.5 The capture threshold is set from the measured received level, for example the largest $|y|^2$ in one manifold period plus a margin. After lock the threshold is lowered for holding.

3.6 Candidate order for robust lock: when the amplitude latch could be fooled by an interfering pulse, the receiver tries candidate nail positions in descending amplitude order and accepts the first one whose whole-manifold decision passes.

3.7 Variants: 1 to 8 nails per manifold, nails at any indices, nails formed as short multi-sample patterns, nails at other wavelengths or polarizations, and nail strength set by received level, range and environment.

4. Whole-manifold decision by windowed coherent vote ("Coherent Vote")

4.1 The receiver divides one received manifold into W windows (W from 2 to 64, for example 8). For each candidate whole-manifold state, it computes a normalized coherent correlation between the received samples and the candidate in each window (for example the real part of the normalized inner product). A candidate wins a window if its correlation is highest there. A candidate is accepted if it wins at least V windows ($V > W/2$, for example 5 of 8). Otherwise the manifold is rejected; no partial decision is delivered.

4.2 The candidates may be all registered states, or a short list (for example the 8, 16, 48 or 128 nearest states) produced by a projection of the whole received manifold. Each short-listed candidate is regenerated as a complete manifold by the same encoder before voting.

4.3 The vote may run in programmable logic (FPGA/ASIC) using fixed-point multiply-accumulate per window, or in a processor.

4.4 The same vote is used for: payload decision, lock confirmation, identity confirmation, the pre-transmit self-check (section 7.6), and verifying feedback changes on a fresh manifold after they take effect.

5. Transmitter/receiver feedback in two layers

5.1 Own-unit layer. (a) Receive level: one variable optical attenuator on the receive path is stepped up when the input clips or exceeds a calibrated level and stepped down when it is low, with a dead band. The steps may be binary-search steps. The attenuator code changes only at a whole-manifold boundary. (b) Modulator bias: before lock, the I, Q and phase biases of an IQ modulator are adjusted using whole-manifold sums of the received samples (for example Σy for carrier leakage and Σy^2 for quadrature error), which need no frame alignment. After lock, they are refined against the known book. Each adjustment takes one step, keeps it if the error decreased, tries the opposite direction otherwise, and halves the step when neither direction helps. (c) Transmit power and nail gain are set from the receive state.

5.2 Inter-unit layer. Each unit reports its own state (for example: not locked, locked, locked and reading the peer) by which registered whole-manifold state it transmits, or in its payload. Mutual lock is declared when both sides report locked-and-reading for M consecutive fresh manifolds. Each step of the procedure is held only while the peer keeps confirming it, and otherwise falls back to the previous step.

5.3 All changes take effect at a future whole-manifold boundary and are verified on a newly received whole manifold.

6. Beam steering: scan, latch-and-stop, track, lock, hold

6.1 A two-axis scanner (galvanometer mirrors, MEMS mirror, optical phased array, liquid-crystal or Risley prisms, on an optional gimbal) rasters at manifold cadence: one scanner update per manifold, with X and Y updates alternating or simultaneous.

6.2 Latch-and-stop: when the receive amplitude latch reports a detection, programmable logic freezes the scanner at the current code in the same manifold and records it. It does not wait for the processor.

6.3 Track: the scanner then makes small steps (for example 8 codes, then halved to 1 code) and keeps each step that raises the received level, until a maximum is reached.

6.4 Lock and hold: after lock (section 3) and mutual lock (section 5.2), the scanner stops scanning and holds pointing for the communication state. The communication beam is stronger than the sensing echo by a factor that grows as R^2 .

6.5 Acquire once, then track only the critical neighbours: the unit scans the full field once to acquire all neighbours. It then predicts each neighbour from its kinematic state, ranks

neighbours by time to closest approach (or by distance), selects the K most critical (K from 1 to 50), and scans only a small area around each predicted position. The area grows with the neighbour's maximum acceleration and the elapsed time. The unit stops as soon as each neighbour is found.

6.6 Beam divergence may be widened to half the scan line spacing, so that a sweep leaves no gaps and the beam dwells on any point for at least one manifold. If a unit sees a scanning beam pass over it, it stops scanning and turns towards the source.

7. Per-unit books, link-derived factors and transmit self-check

7.1 Per-unit book: every unit uses the same mother book multiplied by a fixed per-unit key. The key may be a random phase on each spectral line of the manifold (each line of an N-point discrete transform), optionally combined with a permutation of in-band lines, or any other invertible per-line transform. A recording of one unit's emission does not give the mother book or any other unit's book.

7.2 Time-varying factor: each transmitted manifold is additionally multiplied by a factor m_t that depends on the manifold sequence number t . m_t is derived from the link: from the previous manifold, a counter, or a keyed hash. The starting value is a secret shared by the two units at mutual lock. The receiver computes the same m_t and divides it out before calibration and decision. If a frame is lost, the units return to the lock handshake.

7.3 Session secret: each pair of units shares a pairing secret written at pairing time. After lock, each side sends a fresh nonce. The session secret is derived as $H(\text{pairing secret} \parallel \text{nonce}_A \parallel \text{nonce}_B)$, and $m_t = \text{derive}(\text{session secret}, t)$. Replaying an old session gives a different factor and is rejected by the vote.

7.4 The pairing secret may be held in a secure element that computes the session seed internally, so the processor only ever handles one session's seed.

7.5 Counter slips are recovered by trying the decision for $t-2 \dots t+2$ and keeping the value that passes.

7.6 Pre-transmit self-check: before a payload bank is committed for transmission, the processor runs the whole-manifold vote of the bank against the unkeyed mother book. The bank is committed only if that vote fails, so a fault that skips keying cannot send the mother book.

7.7 Sensing is unaffected, because the unit knows exactly what it sent.

8. Intermediate-frequency pre-rotation of the book

8.1 To survive AC-coupled (high-pass) drive and receive electronics, every manifold is multiplied by $\exp(j2\pi m n / N)$ with integer m before transmission, which moves its

spectrum to an intermediate frequency. The receiver multiplies by the conjugate. Examples: IF 1 to 40 MHz, for example 8 or 16 MHz; m chosen so the shifted spectrum stays inside the converter band.

8.2 Pre-rotated tables may be stored so that the logic plays them unchanged. Calibration then needs only a fixed delay and complex gain per link, measured once with the same book.

9. Hardware architecture

9.1 Radar body = a pluggable coherent optical transceiver module (for example CFP2-ACO class, or another form factor with analog coherent interfaces). It contains a tunable laser, a dual-polarization IQ modulator with linear drivers, an integrated coherent receiver, and modulator bias control. A host board carries the module socket, digital-to-analog converters driving the modulator inputs, analog-to-digital converters sampling the receiver outputs, and an FPGA or SoC (for example Zynq or Zynq UltraScale+ class). The same laser is the transmitter and the local oscillator for sensing.

9.2 Optical front end: module transmit port $\rightarrow$ optical amplifier (for example EDFA, 0.1 to 100 W) $\rightarrow$ three-port circulator port 1; port 2 $\leftrightarrow$ collimator $\rightarrow$ two-axis scanner (+ gimbal); port 3 $\rightarrow$ one variable attenuator $\rightarrow$ module receive port. The front end may be built as one separate free-space front-end module, connected to the photonic core by two fibres and control lines.

9.3 Receive both polarizations (X and Y, four channels $I/Q \times 2$) while transmitting one or two. Use the X/Y ratio as an extra observable of the target (for example surface type, roughness, wetness) and to avoid fading when the units roll relative to each other.

9.4 Inter-unit communication with independent lasers: each unit uses its own laser as local oscillator. Frequency and phase are referenced to the common rail of the received manifold.

9.5 Alternatives: discrete IQ modulator plus 90° hybrid and balanced detectors; directly modulated laser with interferometric (self-heterodyne) reception; the 1,310 nm, 1,550 nm or 2,000 nm bands; fibre or free-space paths; ASIC or photonic integrated circuit implementations of any part.

9.6 Security hardening: encrypted and signed boot image, debug ports permanently disabled, per-unit keys, raw frames never exported, a secure element for pairing secrets, and DAC-to-modulator traces in inner board layers.

10. Kinematic tensor as control output; swarm as a tensor field

10.1 For each received manifold, an estimator outputs the source's or target's kinematic tensor: position, velocity, acceleration and jerk in three axes. Any estimator may be used:

least-squares fitting, polynomial or spline models, state-space filters, physics-based inversion, learned models, or table lookup. The tensor is passed directly to a vehicle's control system as the control quantity, for example as position, velocity or acceleration setpoints over a standard autopilot interface (for example MAVLink to PX4 or ArduPilot) or a CAN, Ethernet or serial interface.

10.2 Potential-field navigation: the tensors of all perceived objects build a scalar or tensor potential field. The vehicle follows the valleys of that field (for example its gradient descent). Moving objects become moving hills.

10.3 Swarm as a field: every unit is a node of the field. It perceives neighbours by seeing (echo) and hearing (direct light), and keeps a field table. Each entry holds a neighbour's identity (from its identity nail), its kinematic tensor, the measurement time and the provenance (seen, heard or relayed). The table is passed on in the payload whenever two units are mutually locked, and each unit keeps the newest entry per neighbour by measurement time. An emergency stop spreads by perception: a unit that changes its motion state is sensed by its neighbours, which react and are in turn sensed by theirs, with no message relay needed.

10.4 Applications: drones and drone swarms (formation flight, crossing, collision avoidance, GNSS-denied navigation relative to a hand-held or ground anchor), road vehicles (sensing other vehicles and their speed while exchanging data with them), robots, ships, satellites and inter-satellite links, base stations, and handheld controllers.

11. Example combinations

11.1 A drone carrying the module of section 9.1, the front end of 9.2 and the scanner of section 6. It locks with a peer drone (section 3), confirms identity, exchanges field tables (10.3) using a 60-bit-per-manifold level (2.5) with per-unit keys (7.1–7.3), and feeds the kinematic tensor of section 10.1 to its autopilot.

11.2 A vehicle sensing system in which the same coherent beam measures the range and kinematic state of surrounding vehicles and exchanges data with equipped vehicles, using the IF pre-rotation of section 8 on AC-coupled electronics.

11.3 A ground station that tracks several drones using section 6.5 and coordinates them through the field table of 10.3.