Native integrated sensing and communication on rigid waveform manifolds (continued): window and structure of common-aperture scan heads, master-and-wingman swarms, moving-mountain potential-field control, video return on the same beam and radio backup

Nature: Technical disclosure (defensive publication)
Author / publisher: V131 Lab
First public disclosure: 2026-11-03
Earlier disclosure: DP01, DOI 10.5281/zenodo.23420482 (2026-10-16). This document uses the terms of DP01: manifold, look (own echo), listen (direct light), kinematic state tensor, potential field, field table.
Field: Coherent optical ranging (FMCW and similar), free-space optical communication, integrated sensing and communication (ISAC), beam scanning and gimbal structure, unmanned aircraft, autonomous driving, robots, swarms.

This document discloses a family of systems and methods. Each numbered paragraph describes one approach and its alternatives. Unless a paragraph says otherwise, any approach may be combined with any other approach and with any approach in DP01. Every alternative, every numeric range and every combination listed here is part of this disclosure.

(This English text and DP03_ZH correspond paragraph by paragraph. The English text is primary; the Chinese text is published at the same time.)

---

1. Tilted window on a common-aperture scan head

1.1 Problem: when one collimator both transmits and receives (DP01 paragraph 9.2), the surface reflection of the protective window (for example 0.1 % to 1 % after anti-reflection coating) is a part of the outgoing light. When the window is perpendicular to the outgoing light, this reflection travels straight back into the collimator and through the circulator to the receiver, and it is much stronger than the target echo. A variable attenuator in the receive path attenuates this reflection and the echo together and leaves their ratio unchanged, so the attenuator does not solve the problem.

1.2 Approach: make the window non-perpendicular to the beam for every outgoing direction in the scan range. The tilt θ of the window normal relative to the scan centre axis exceeds the maximum optical scan angle of that axis plus a margin (for example 1° to 15°, such as 5°), so that the normal lies outside the scan frame. For example, tilt about 20° for ±15° optical per axis, and more than 30° for ±25°; θ may be 3° to 60°.

1.3 Tilt direction: for two-axis scanning, tilt along the scan direction of the mirror closer to the window. The beam sweeps a narrow range across the window in this direction, so the window grows little. With the normal outside the frame on this axis and at the zero position of the other axis, the tilt does not have to be computed from the diagonal of the frame.

1.4 Window size and coating: enlarge the window in the tilt direction by 1/cos θ. Design the anti-reflection coating for the actual range of incidence angles and not for 0°. For example, with 20° tilt and ±15° scan,

the incidence range is about 5° to 38°; design for 20°, with reflectance not above 0.5 % for both S and P polarization.

1.5 Where the reflection goes: tilting does not reduce the amount reflected; it only keeps the reflection out of the fibre. The angular coupling tolerance of the fibre is about the mode-field radius divided by the collimator focal length (for example 5.2 μm / 35 mm, about 0.15 mrad), so a deviation of this order keeps it out of the receiver. Direct the reflection to an absorbing surface inside the housing and not back onto a mirror.

1.6 Matching the scan limit to the tilt: calibrate the gain of the scanner driver so that full scale of the analog output of the main board corresponds exactly to the maximum optical scan angle (for example single-ended ±10 V for ±15° optical), and add a software limit. The hardware then cannot scan to the direction of the window normal. When the scan range is widened, increase the tilt at the same time.

1.7 Variants: a wedged window; a curved window whose centre of curvature is offset from the beam pivot; one window for two or more output apertures; fused silica, BK7, sapphire or other materials, 1 mm to 10 mm thick; the 1,310 nm, 1,550 nm or 2,000 nm bands; the scanner being galvanometer mirrors, MEMS mirrors, an optical phased array or Risley prisms.

2. Structure of a capsule-shaped two-axis scan head

2.1 Shape: in elevation the scan head rotates about one axis only, so the housing is a capsule (a cylinder with two hemispherical ends, the cylinder axis along the elevation axis) and not a sphere. For the same contents a capsule is lower than a sphere, for example about 100 mm in outer diameter and about 140 mm to 170 mm long. The window sits on the capsule surface and needs no recess.

2.2 Inner frame: the scanner and the collimator are both fixed to an inner frame and are not located by the outer housing. The housing is split into two halves that close around the inner frame. The elevation axis passes through the centre of gravity of the scan head, for example within 3 mm; any offset is corrected by changing the hub or adding a counterweight.

2.3 Driver board position: the scanner driver boards are not in the head. They sit on the azimuth platform and rotate with the azimuth, so the head is lighter and the torque required of the elevation and azimuth motors is lower. The space under the platform is ventilated and not enclosed.

2.4 Routing: power cables, signal cables and fibres carry connectors at both ends, and the connectors cannot be cut off, so the routing path contains no closed holes. The housing is split in two, the notches in the platform and base are open to the edge, the yoke arms use open channels with cover plates, and the elevation bearing only provides support and passes no cable. Cables and fibres loop between the housing and the yoke arms and outside the azimuth motor with enough slack for the travel. Azimuth travel is limited (for example ±160°), and no slip ring is used.

2.5 Self-contained unit, universal base: each scan head is a self-contained unit with its own base, cables and driver boards. The base mounting holes are universal, and nothing is built into one kind of vehicle. A vehicle may carry 1 to 8 heads, each serving different directions. On a car the same kit is mounted the other way up.

2.6 Mounting under the belly: on an unmanned aircraft all scan heads are mounted under the belly, so the whole lower hemisphere is unobstructed. The landing gear leaves enough height (for example about 200

mm below the belly), and azimuths blocked by landing gear legs are written into a software keep-out. During takeoff and landing the elevation is turned so that the window faces up, to reduce dust settling on it.

2.7 Autonomous flight without satellite navigation: when satellite navigation is unavailable or not used, the height above ground and the kinematic state relative to the ground, obtained from the same beam aimed at the ground, serve as the basis for flight and are passed to the flight controller (DP01 paragraph 10.1). At any moment the scanner covers one limited frame (for example 30° × 30°), so looking down and looking forward are scheduled in time, or shared between different heads. The same beam continues to carry payload both when looking down and when looking forward.

2.8 Variants: spherical, cylindrical or flat housings; metal, plastic or 3D-printed parts; a three-axis gimbal; the vehicle being an unmanned aircraft, road vehicle, ship or robot; driver boards in the head or in the vehicle body.

3. Swarms of a few masters and many wingmen

3.1 Two kinds of node: a master carries the coherent optical transceiver described here and in DP01, and also a radio transmitter. A wingman has no coherent optical transceiver and carries only a radio receiver and a flight-control or drive-by-wire interface. Masters are a minority of the swarm, for example 2 to 10 of 50, such as 6; the ratio may be from 1 : 2 to 1 : 100.

3.2 Form the field, one node sends: masters send each other the kinematic state records they have newly seen, by short radio packets or by light. One agreed master (for example the one with the lowest number) merges what it has seen and what it has received into one field and broadcasts the whole field f times per second (f from 1 to 100, for example 10). If this master stops sending for some time, the next number takes over. The radio channel therefore carries only one whole field plus the short packets between masters, and its occupancy does not grow appreciably with the number of masters. When masters exchange records by light, only this one field remains on the radio channel.

3.3 The field is alive: the field is not a data table whose rows are merely old or new. Between two sightings each mountain (each aircraft) is extrapolated forward according to its kinematic state tensor, and it is corrected each time a master sees it. A wingman that has received the field extrapolates in the same way between two fields.

3.4 Flight of a wingman: a wingman flies along the valley floor according to this field alone (DP01 paragraph 10.2, section 4 here). It needs no connection, pairing or handshake with masters or with other wingmen. The broadcast requires no acknowledgement and no retransmission, and a lost packet is made up by the next whole field. Each field carries a sequence number and a keyed authentication tag (for example HMAC) against forgery and replay. A kinematic state record is for example 16 to 64 bytes.

3.5 Emergency stop spreads through the field: when an aircraft stops or makes an emergency stop, this appears in the kinematic state of its mountain. A master sees it and it enters the field, and the whole swarm learns of it from the field without a separate stop command (DP01 paragraph 10.3).

3.6 Number and position of masters: masters are chosen to be as widely spread as possible by initial position, or chosen at random. Their number is set so that, after any one master is removed, the freshness

of the neighbour records held by the wingmen and the warning time of conflicts still meet the requirements. A master may fly like a wingman, or be placed high up or on the ground.

3.7 Variants: the radio being Wi-Fi, connectionless Wi-Fi frame broadcast (of the ESP-NOW kind for example), Bluetooth, UWB, Sub-GHz or cellular; the broadcast sent by light; wingmen carrying small sensors, or also sending a little of their own state; the swarm divided into zones with one field per zone; a fixed base station or ground station acting as a master; mixed groups of aircraft, vehicles, robots and ships.

4. Moving-mountain potential-field control

4.1 Definition of a mountain: another vehicle or object is a mountain on the potential field. The mountain has a finite size (its own outline or projected area, optionally with a margin), is the highest thing on the field, and moves with its kinematic state tensor (it is drawn on the path it is about to take, extrapolated from velocity, acceleration and so on). The path a vehicle wants to take is not mountains everywhere, and there are no fixed mountains; only walls and similar static structures are fixed.

4.2 Terrain: the floor of the potential field is the path length along the walkable region to the destination, not the straight-line direction to the destination, in the way water runs downhill. This forms no local pit in front of a wall.

4.3 One field, one set of rules: all vehicles use the same field with the same rules and parameters, and each follows the valley floor, like water flowing past stones. No order, priority or negotiation rules are added. When vehicles meet, one common passing bias (for example everyone yields to the same side) makes them complement each other, instead of each dodging for itself.

4.4 Safety check covers the whole stop: for the whole time from now until a full stop under braking, check whether a moving mountain will run into the vehicle; do not only check continuing at the current speed. Vehicle outline, braking and acceleration use data of the real vehicle. With good enough braking and control, the clearance between vehicles may be very small.

4.5 Control output: the valley-floor direction is translated into standard setpoints (velocity, or acceleration and heading, or attitude and thrust) and passed through a standard interface to the flight controller or drive-by-wire chassis. The flight controller firmware is not changed, and attitude stabilization, mixing and failsafe stay with the flight controller. The potential-field computation sits on the sensing module side.

4.6 Mountain size by target type: the type of a target (for example person, car, wood, aircraft) is judged from physical observables such as the polarization ratio of DP01 paragraph 9.3 together with its kinematic state, and different types use different mountain sizes and margins.

4.7 Variants: the mountain being Gaussian, hard-edged or of any monotonic shape; the field covering a two-dimensional plane or three-dimensional space; the field generated by the node itself, or broadcast by a master as in section 3, or a fusion of both; use on unmanned aircraft, vehicles, robots, robot arms and ships.

5. Video return on the same beam and radio backup

5.1 Video return: besides the network port to the flight controller or host, the main board has another network port for a camera or recorder. Compressed video (for example H.264 or H.265, 1 to 50 Mbit/s per

stream, commonly 3 to 8 Mbit/s) is written manifold by manifold into the transmit state queue and sent as manifold payload (DP01 section 2). The same beam continues sensing, and no separate time slot, channel or optical path is opened for video.

5.2 Uses: an unmanned aircraft goes ahead and films, and returns the video by laser to another aircraft or to a ground station; intermediate aircraft may relay; a wind farm, campus or similar site forms an optical network for the return.

5.3 Three radio states: a radio module on the board (for example Wi-Fi 6 with Bluetooth) is controlled by the processor and has three states: off, receive only, and transmit and receive. It is off by default at power-up. It transmits when needed and stays silent otherwise.

5.4 Nearby access and low-rate backup: the radio module serves nearby access by phones, handheld controllers and the like, and low-rate backup when the optical path is blocked, carrying kinematic state, control quantities and a little data. It is not used as a high-rate channel. For example, switch to the backup after M consecutive manifolds without lock, and return to the optical path after relock.

5.5 Sensing backup: when fog, smoke or the like cuts the optical path, a millimetre-wave radar takes over sensing and the radio backup continues to carry data; on recovery of the optical path the system switches back.

5.6 Variants: other sensor data instead of video; a camera connected by USB, MIPI or fibre; several video streams sharing payload by priority; a backup link over cellular, satellite or Sub-GHz.

6. Example combinations

6.1 A master aircraft: one capsule-shaped scan head of section 2 under the belly, with the window tilted as in section 1; looking forward at neighbours and down at the ground (2.7); broadcasting the field to wingmen as in section 3 and flying as in section 4 itself; camera video of section 5 returned to a ground station on the same beam.

6.2 A swarm of 50: 6 masters and 44 wingmen; wingmen receive only the one field of paragraph 3.2 and follow the valley floor as in section 4; when a master stops sending, the next one takes over.

6.3 A group of vehicles: a scan head on the universal base of paragraph 2.5 on each roof, driving as in section 4 and exchanging data with each other as in DP01 section 2; the backup of paragraph 5.4 when the optical path is blocked, and the radar of paragraph 5.5 in fog.