Far-Infrared Heterodyne Interferometry Using Orbital Computing Infrastructure

Lin Yi, Ph.D.
Jet Propulsion Laboratory
California Institute of Technology
Lin.Yi@jpl.nasa.gov

Jet Propulsion Laboratory
California Institute of Technology

Ad ASTRA Workshop Concept Study
September 2nd, 2026
California Institute of Technology, Pasadena, CA, USA

Why revisit space interferometry now?

The science need is longstanding; the potentially reusable infrastructure is new.

Historically expensive

Multiple cold telescopes
Precision timing + metrology
High-rate crosslinks
Large correlators
Autonomous operations

Orbital compute may supply

High electrical power
Optical networking
Onboard accelerators
Distributed routing
Commercial deployment scale

Science selection rule

Architecture must follow detectable brightness and angular scale.

Longer baselines help only when the source remains bright on those spatial frequencies.

Core hypothesis: reuse the infrastructure but preserves the science.

Science I — 1 hosted telescope can extend ALMA beyond Earth-sized VLBI baselines

A single mm/sub-mm telescope (100-300GHz) on a distant orbital compute node can form space–ground VLBI baselines with phased ALMA/EHT (modified BHEX).

Architecture A for Science I : one science telescope in space

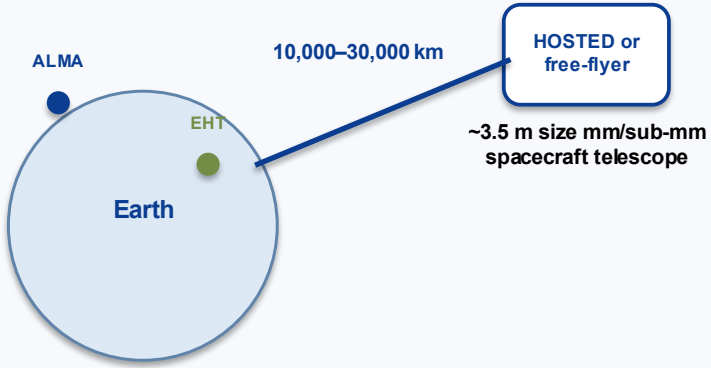

Ground array supplies collecting area; orbital node supplies the long baseline.

Resolution leverage at 230 GHz (λ ≈ 1.3 mm)

ALMA alone	16 km	≈17 mas
Earth-scale VLBI	10,000 km	≈27 µas
Space–ground VLBI	30,000 km	≈9 µas

30,000 km gives ~1,900× finer nominal resolution than ALMA alone and ~3× beyond an Earth-sized baseline.

Unique science — Photon-ring spatial frequencies • strong-field gravity • compact AGN jet-launch region

The unique observable is correlated visibility at spatial frequencies unavailable to ground-only arrays.

Why a data-center host? Power • high-rate buffering • optical downlink • onboard compute • timing/frequency support

Why a free-flyer? Dedicated pointing Cold detector isolation from the CDC

Sensitivity: phased ALMA ≈94 Jy SEFD at 1.4 mm, equivalent to ~75 m effective diameter for VLBI.

Mission gate: LEO-only hosting is unlikely to reach the full photon-ring regime; projected baselines of order 10^4–3×10^4 km are the relevant target.

Sources: ALMA Observatory Factsheet; EHT proposal guidance for phased ALMA; Johnson et al., BHEX photon-ring science (2024).

Science II — ISM lifecycle: resolve where gas changes phase

Quantitative reference: bright Galactic PDRs such as the Orion Bar; 2 m collectors, 50–100 m baselines, velocity-resolved [C II]/[O I]/[N II] maps.

Physical interface to image

Heterodyne data cube: I(x, y, v)

UV feedback → heating • compression • erosion

What 50–100 m baselines resolve at Orion

Line	λ	50 m	100 m
[O I]	63 μm	0.26″ / 105 AU	0.13″ / 52 AU
[N II]	122 μm	0.50″ / 205 AU	0.25″ / 103 AU
[C II]	158 μm	0.65″ / 265 AU	0.33″ / 135 AU
[N II]	205 μm	0.85″ / 350 AU	0.42″ / 175 AU

Assumes Orion distance ≈414 pc; values rounded.

Line diagnostics

- **[C II] 158 μm** — CO-dark H_2 + PDR interfaces
- **[O I] 63/145 μm** — dense warm gas + shocks
- **[N II] 122/205 μm** — ionized gas + electron density

Brightness sanity check: Orion Bar [C II]

Observed bright PDR	If Δv≈3 km/s	100×2 m array, 1 h	Equivalent T_B rms
≳700 K km s^{-1}	$\langle T_B \rangle \approx 230$ K	σ≈0.18 Jy / 1 km/s	0.14 K @50 m; 0.56 K @100 m

Interpretation: bright Galactic PDRs are likely signal-rich for 2 m collectors; the hard problem is not raw brightness, but recovering extended emission without filtering it out.

Array implication: include Bmin≈5–10 m plus 50–100 m baselines. Benchmarks use Orion Bar [C II] intensities; next step is visibility-domain modeling.

Science III — planet formation: reveal the hidden gas reservoir and its dynamics

Reference case: 2 m collectors with 50–100 m baselines; 100 m provides the primary science resolution, while 50 m recovers brighter extended emission.

What does the measurement add?

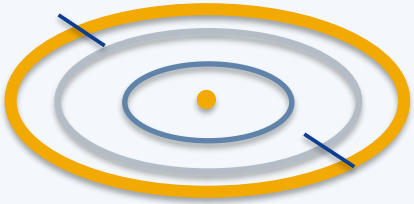

Dust tells us where solids concentrate.

HD 112 μm tells us how much gas remains and how it moves.

Reference geometry — HD J=1→0 at 112 μm

50 m baseline **0.46″**

≈27 AU at TW Hya (59.5 pc)

≈65 AU at Taurus (140 pc)

Role: recover brighter, more extended disk gas

100 m baseline **0.23″**

≈14 AU at TW Hya (59.5 pc)

≈32 AU at Taurus (140 pc)

PRIMARY: resolve radial gas structure on tens-of-AU scales

Brightness sanity check — TW Hya

Measured HD line

6.3×10^{-18} W m^{-2}

≈71 Jy km s^{-1} integrated

Illustrative average line brightness

50 m: 0.38–1.5 Jy/beam

100 m: 0.10–0.38 Jy/beam

Illustrative array sensitivity

N=100 × 2 m • 10 h

σ ≈ 0.028 Jy/beam

Interpretation

50 m is comfortably bright in this sanity check; 100 m remains plausible for the brightest / concentrated HD emission.

Caveat: brightness ranges are image-plane estimates for a 50–100 AU emitting radius; the next step is a visibility-domain HD model of TW Hya.

Sources: Bergin et al., Nature 493, 644–646 (2013); McClure et al., ApJ 831, 167 (2016). Sensitivity assumes ηA=0.5, Tsys=2000 K, Δv=5 km s^{-1}.

Architecture B — hosted interferometer array (for science II and III)

Many orbital-compute spacecraft carry identical 2 m cryogenic heterodyne payloads; the host constellation itself becomes the interferometer.

Science payloads ride directly on compute spacecraft

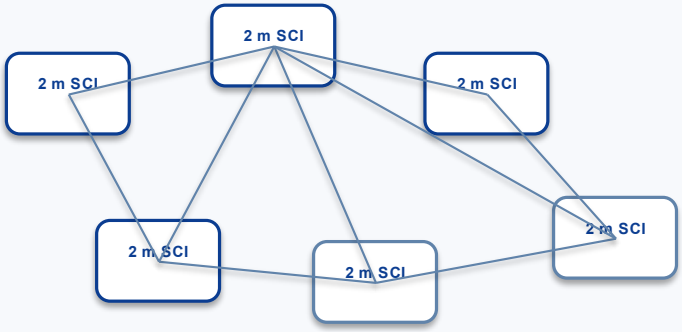

Optical links / timing / ranging / distributed correlation

Reference science configuration

2 m	5–100 m	10s–100
collector diameter	baseline span*	science nodes

Best matched to

Bright far-IR line imaging:
ISM/PDR interfaces + nearby disk gas

*Short baselines recover extended emission; 50–100 m provides science resolution.

Why use the host?

power • optical crosslinks • timing • onboard compute • mass-produced buses

Critical gates

- Coarse and fine pointing for ODC and hosted payload
- 50-100m proximity between nodes
- Can industry actually provide many simultaneously useful 5–100 m baselines without compromising its primary mission?

Architecture C — free-flying science array + distributed orbital correlator

Purpose-built cold astronomy spacecraft send heterodyne data into the orbital data-center swarm; correlation is distributed, not centralized.

Separate the cold telescopes from the hot compute spacecraft

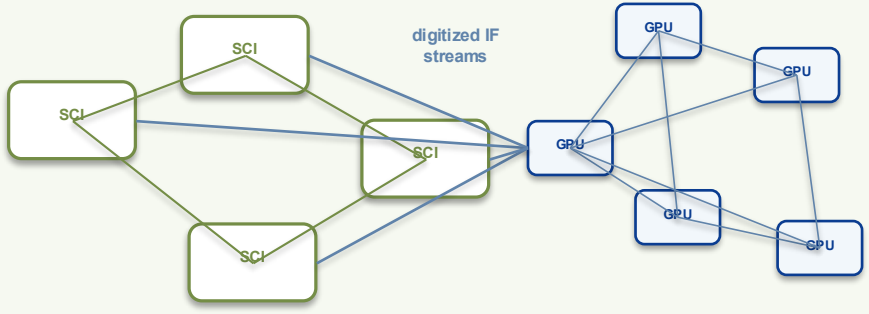

wideband data are channelized and correlated across many compute nodes

Distributed correlator concept

- **Partition by frequency** — sub-bands routed to different nodes
- **Partition by baseline** — node groups handle selected antenna pairs
- **Replicate calibration** — delay models shared across swarm
- **Return science products** — visibilities or calibrated cubes

No single correlator bottleneck: the network is the instrument back end.

Illustrative IF load — 8 GHz IF × 2 pol × 2 bit → ~64 Gbit/s per collector before reduction

Main gate — Can the orbital network supply deterministic timing, aggregate bandwidth and stable routing for correlation?

What can the orbital data centers provide?

All three architectures share a core interferometric measurement chain; their differences are mainly in how deeply that chain couples to the orbital data-center spacecraft.

Commercial ODC companies do not publish, tentatively 100kW computing power, 2-tone weight, 100m size, total nodes 100,000*

1 Common science / instrument capabilities

These do not disappear when the architecture changes:

- **Cryogenic heterodyne receiver**
- **Coherent time / frequency reference**
- **Baseline + delay knowledge**
- **Pointing / mechanical stability**
- **Correlation + calibration**

2 ODC technologies that directly help

High electrical power	receiver electronics, ADCs, optical terminals
Optical crosslinks	IF/data transport + time/frequency transfer – partial baseline/delay knowledge
Onboard compute/storage	channelization, buffering, distributed correlation
Network autonomy	routing, scheduling, fault tolerance

3 How each architecture couples to the ODC

B	**Hosted FIR swarm**	Deep spacecraft integration	**Conditional**
C	**Free flyers + compute swarm**	Uses native ODC network + compute	**Strong**
A	**Single hosted space-VLBI**	Only one modified host node	**Strong***

Design implication: maximize reuse of ODC power, links, timing and compute — while minimizing astronomy-specific demands on the host bus and orbit.

jpl.nasa.gov

*https://orbital.inc/press/orbital-100000-constellation.html

Takeaway

Far-IR interferometry with specific narrow scientific objectives/sources can be achieved

Host-science payloads still have challenges in cryogenic temperature, local time reference/delay calibration and precision pointing

Orbital Data Center nodes can contribute in providing power, laser links (fully data, partially timing and path/delay calibration) and correlation computing

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov

Technology gaps: science-driven requirement vs. demonstrated state of the art

Working requirements are derived from the three reference science cases; the comparison separates component capability from unproven system-scale integration.

SCIENCE / ARCH.		WORKING REQUIREMENT	DEMONSTRATED SOTA	GAP / INTERPRETATION	
I / A Space–ground VLBI	Baseline / resolution	20–35 Gλ ≈6–10 μas	Earth diameter: ≈10–14 Gλ	OPEN	≈2–3× longer spatial frequency; mm/sub-mm high-orbit VLBI unflown
	Timing coherence	$\sigma_y(10\,s) \approx 1–5\times10^{-14}$ ≤5–10% loss	optical-comb lab: $\approx10^{-14}$ class @10 s	PARTIAL	Performance near/at target; space qualification + integration remain
	Data transport	64 Gb/s raw; ~100 Gb/s optical link	TBIRD flight: 200 Gb/s from LEO	PARTIAL	Raw rate demonstrated; higher-orbit link margin/duty cycle remain
II / B ISM / PDR FIR imaging	THz receiver noise	Tsys ≈1.5 kK @1.9 THz; ≤3 kK @4.7 THz	GUSTO receiver DSB: 1.10 / 1.92 kK	PARTIAL	Detector sensitivity close; space-qualified coherent LO + stability
	Path / baseline knowledge	≤λ/10 → 6–20 μm (63–205 μm lines)	GRACE-FO LRI: ~1 nm/√Hz scalar range	PARTIAL	>10^3× raw ranging margin; 3-D vector + absolute delay calibration open
	Formation + data scale	5–100 m; tens–100 nodes; 64 Gb/s/node → 6.4 Tb/s @N=100	Proba-3: 2 nodes, 150 m, ~1 mm; TBIRD 200 Gb/s	OPEN	Node-count/u-v + multi-Tb/s network are system-scale gaps
III / B HD 112 μm disk gas	Receiver / path	Tsys ≈2.0 kK @2.68 THz; path ≤11 μm	2.7 THz lab receiver: 880 K DSB; LRI ~1 nm/√Hz	PARTIAL	Noise & scalar metrology not fundamental; flight integration remains
	Array sensitivity	Expected HD signal @100 m: ~0.10–0.38 Jy/beam	(estimated) 100×2 m, 10 h: σ ≈0.028 Jy/beam	OPEN	Science closes only with large collecting area; margin is source-structure dependent
	Scale sensitivity	Keep σ≲0.03 Jy/beam for useful 100 m imaging	30×2 m → ~0.094 Jy; 100×1 m → ~0.112 Jy	OPEN	≈3–4× sensitivity shortfall if node count or diameter is reduced

Commercial ODC dependency

LIKELY ODC-DRIVEN

Commercial scaling directly helps

- **Raw data links** — I / II / III
  ≥64 Gb/s per collector; ~100 Gb/s-class link
  TBIRD: 200 Gb/s
- **Aggregate networking** — II / III
  N=100 → 6.4 Tb/s raw ingest
  Commercial ODC scale-out could attack this
- **Correlation compute** — II / III
  N=100 → 4,950 baselines
  Naturally parallel across ODC compute nodes
- **Storage + autonomy** — All
  buffer multi-Gb/s streams; route around failed nodes
  Native data-center function

CO-DEVELOPMENT

ODC helps, but science-specific validation remains

- **Timing / frequency** — I / II / III
  $\sigma_y(10\ s) \approx 1–5\times10^{-14}$
  ODC optical links/comb can distribute reference
- **3-D delay metrology** — II / III
  6–20 µm path knowledge for 63–205 µm FIR
  Raw laser ranging has >10^3× margin
- **Hosted pointing / stability** — A / C
  science-grade LOS + mechanical stability
  Depends on commercial bus accommodation
- **High-orbit infrastructure** — I / C
  10^4–3×10^4 km projected baseline for Science I
  Only solved if ODC fleet extends beyond LEO

ASTRONOMY-SPECIFIC

ODC advancement does not remove the gap

- **THz receiver + cryogenics** — II / III
  Tsys ≈1.5–3 kK @1.9–4.7 THz
  ODC power helps; frontend still astronomy
- **Collecting area** — III
  HD case: ~100×2 m for σ≈0.028 Jy/10 h
  Compute cannot recover missing photons
- **Science formation / u-v** — II / III
  5–100 m baselines; tens–100 nodes
  Must be designed for surface-brightness recovery
- **Absolute phase calibration** — II / III
  µm-class vector delay + source-referenced calibration
  Requires end-to-end astronomy demonstration

Architecture consequence

- **A Conditional:** ODC must become astronomy-friendly
- **B Strongest ODC synergy:** commercial fleet = digital back end
- **C Strong if high orbit exists:** one hosted node + ALMA/EHT

ODC progress can plausibly retire bandwidth, compute, storage and part of timing/metrology; the hardest residual gaps remain THz frontend performance, collecting area, science geometry and end-to-end coherence.