

NASA Ad ASTRA COMMUNITY WORKSHOP · INVITED TALK

Photonics for Future Observatories

Miniaturized, stable and scalable instrumentation for space astrophysics

Nemanja Jovanovic — Lead Instrument Scientist


Caltech Optical Observatories

September 1st-3rd, Pasadena, California




What is photonics?

- **Manipulating light in wavelength-scale guiding structures - narrow definition:**
 - Light is confined in a discrete set of guided modes
 - Function comes from engineered geometry and phase
 - Devices are monolithic — lithographically defined on a wafer, or laser-written in glass
- **Astrophotonics = the application of these technologies to astronomical instruments**
- Two families
 - fiber photonics (fiber Bragg gratings, lanterns, tapers) and
 - integrated photonics (chips)
- Platforms set the transparency window and the achievable functions:
 - silica ($0.3\text{--}2.3\ \mu\text{m}$),
 - silicon nitride ($0.4\text{--}4\ \mu\text{m}$),
 - silicon-on-insulator ($>1.1\ \mu\text{m}$), ,
 - thin-film LiNbO_3 ($0.4\text{--}5\ \mu\text{m}$), ,
 - chalcogenides for the mid-IR



Free space: beam shape set by lenses and distance — any field allowed




Guided: only discrete transverse modes; only the phase advances


What photonics can do, I: spectral, spatial and modal

Think in terms of functions


- **Disperse** Arrayed waveguide gratings · planar concave gratings · on-chip echelles
- **Notch / band filter** Fiber Bragg gratings · complex aperiodic FBGs for OH suppression
- **Resonate** Ring and racetrack resonators · photonic-crystal cavities ($Q = 10^5$ - 10^7)
- **Spatially filter** A single-mode waveguide is a perfect spatial filter — aberrations become throughput loss, never a change in output mode
- **Convert modes** Photonic lanterns (multimode to $N \times$ single-mode) · mode-selective and hybrid lanterns
- **Reformat geometry** Pupil remappers · integral-field slit dicers written in glass
- **Split / combine** Y-junctions · directional couplers · MMIs · star couplers
- **Combine coherently** On-chip beam combiners for long-baseline interferometry · nulling combiners




AWG (Cvetojevic+ 2012)



Ring resonator




Fiber Bragg grating





Photonic lantern (Leon-Saval+ 2005)

Interferometric beam combination



K. Perraut, A&A. (2018)




M-A. Martinod, et al, Nat. Comms. (2021).

What photonics can do, II: active, nonlinear and detection

These are the functions that need power, or that make light rather than route it.


- **Modulate / switch** Mach-Zehnder switches · ring modulators · MEMS · phase-change and liquid-crystal cells
- **Delay / steer** Tunable delay lines · optical phased arrays · programmable MZI meshes
- **Generate** Hybrid III-V on-chip lasers · Er-doped waveguide amplifiers
- **Broaden / convert** Supercontinuum · second-harmonic generation · Kerr microcombs for calibration
- **Handle polarization** Polarization splitters and rotators · deliberately birefringent waveguides
- **Detect** Integrated photodiodes · superconducting nanowire detectors written onto the waveguide

Detect - SNSPDs




O. Kahl et al, Scientific Report. (2015).

Switches



*MZI switches with integrated phase shifters.
W. Yue et al, Photonics (2023).*


Programmable MZI meshes



M. Dong et al, NPJ Quantum Information (2023).

Why this matters for space missions

- **Miniaturization** — at the diffraction limit the spectrograph size is set by resolving power and wavelength, not by telescope diameter
- **Integration**
 - Optical benches collapse onto a monolithic substrate,
 - Far fewer alignments to hold
- **Stability** — a centimeter-scale footprint makes thermal and mechanical stability affordable rather than heroic
- **Replicability**
 - Lithography amortizes design cost across a wafer;
 - Devices are identical to lithographic tolerance, so one calibration model serves all N units
- **A higher performance ceiling** — photonics controls spatial, spectral and polarization degrees of freedom universally, so an optimally designed device can outperform a classical assembly. Photonic nulling architectures are projected to deliver 2-4× the science yield of conventional coronagraphs (Sirbu et al, JATIS 2026)



1.6 cm

Collimated beam for $R = 20,000$ at $1.55 \mu\text{m}$ — on a 1 m or a 10 m aperture alike

2-4 ×

Projected science-yield gain of photonic nulling over a conventional coronagraph

Five priority areas identified by the astrophotonics community

From the 2025 NASA Emerging Technologies for Astrophysics Missions Workshop — astrophotonics breakout (co-leads D. Sirbu, N. Jovanovic): the top five areas for investment to enable infusion into HWO and astrophysics missions generally.

1 Programmable photonic integrated circuits

Reconfigurable MZI meshes as universal linear optical processors — starlight suppression, in-situ wavefront control, one chip switching between coronagraphy, spectroscopy and sensing

2 Space-compatible laser frequency combs

Absolute, traceable wavelength references from the UV to the near-IR, holding calibration for the life of the mission

3 Photonic lanterns

Efficient conversion of telescope light into arrays of single-mode channels, preserving spatial information for filtering, sensing and nulling

4 Integrated photonic spectrometers

Chip-scale dispersion with high mechanical and thermal stability, at a fraction of the volume of a bulk spectrograph


5 Efficient light injection and interconnects

The optical bridges — free space to fiber to chip to chip to detector — that set how much starlight reaches the science instrument


What they are and how they work

- **Mach-Zehnder interferometer:** split the light, impose a phase difference between two arms, recombine. One output goes bright, the other dark
 - Splitters are Y-junctions, directional couplers or MMIs; the MMI is the robust broadband production choice
 - The phase can be set thermo-optically (μs , mW-scale, ubiquitous), electro-optically in LiNbO_3 (ps, near-zero static power), or piezo/MEMS (low power, cryo-friendly)
- **Add a phase shifter to each arm and the MZI becomes an arbitrary, tunable 2×2 beamsplitter**
- **Tile those cells in a triangular (Reck) or rectangular (Clements) mesh and you can implement any $N \times N$ unitary transformation**
 - $N(N-1)/2$ MZIs, two phase shifters each, all set electronically — no moving parts, no mechanisms to qualify
- What it buys an instrument builder:
 - reconfigurable filters and switches,
 - mode sorting,
 - coherent beam combination,
 - self-calibrating nulling, and
 - wavefront control performed in the guided domain rather than with a deformable mirror
- Errors are not fatal: each cell is imperfect, but downstream cells can be programmed to correct upstream leakage


Mach-Zehnder interferometer (MZI)



M. Dong et al, NPJ Quantum Information (2023).




MZI: a heater sets $\Delta\phi$; the two outputs are complementary
(Saleh & Teich Ch. 23 · Soref & Bennett, IEEE JQE 1987)




The tunable 2×2 gate, and the mesh built from it (Bogaerts et al., Nature 2020)

State of the field

- **GRAVITY at the VLTI is the most spectacular astrophotonic instrument built to date (silica)**
 - A passive photonic beam combiner (PIC) interferes light from four 8 m telescopes; milliarcsecond imaging and microarcsecond astrometry
 - It proves facility-grade photonic beam combination, spatial filtering and manufacturability at scale, which is the precondition for programmable devices.
 - Transformative results across Galactic-centre, AGN and exoplanet science — a capability with no rival, achieved through two-plus decades of sustained European investment
- **Nulling with integrated tricouplers - GLINT on Subaru demonstrated: (ULI)**
 - A passive PIC that allows a pupil plane to be remapped to a linear plane for interference.
 - 3D photonics, path length matching, achromatic phase delays, tri-couplers.
 - Active phase and amplitude control are coming.
- **AstroPIC: (silicon)**
 - 3.8×10^{-10} suppression (94 dB) on-chip (component demonstration, Jan 2026)
 - 8.7×10^{-8} contrast (71 dB) free-space coupling (TRL-3 system demonstration, Mar 2026), *Morgan SPIE/JATIS 2026*
 - 85.4% peak grating coupling efficiency over 20 nm (*Valdez Optica 2026*)
 - Wavefront sensing demonstration with the photonic chip (*Lin SPIE 2026*)
 - Multi-channel suppression with the photonic chip showing 3-5% bandwidth is possible in mesh simulations (*Fogarty SPIE 2026 SPIE*)
 - Upcoming Gen-2 chip on Silicon-on-Insulator with 16-spatial channels and high-efficiency grating coupler included (*Sirbu SPIE 2026*)
 - See Kevin Fogarty's ICEpic poster
- Packaging has matured alongside the circuits — fiber-array optical interfaces and wire-bonded electrical interfaces are now routine




VLTI: four 8 m telescopes feeding a photonic beam combiner



- $8.70 \pm 1.1 \times 10^{-8}$ free-space raw contrast


9x9 programmable PIC nulling a uniform input beam Morgan SPIE/JATIS (2026)



A packaged PIC: fiber array in, wire bonds out

Where they could go in future missions


- **A photonic coronagraph for HWO**
 - An integrated photonic coronagraph provides significant size, weight, and power (SWaP) benefits and a higher theoretical contrast ceiling
 - Starlight suppression performed via a combination of an integrated mode and wavefront control (in-situ)
 - Small inner working angles that are hard to reach with bulk optics, because mode control replaces mask geometry
- **One chip, several instruments** — the same mesh can be reprogrammed between coronagraphic, spectroscopic and wavefront-sensing modes with no moving parts, which is a direct answer to payload volume limits
- **Formation-flying and long-baseline interferometry:** a chip-scale combiner makes multi-aperture architectures tractable, from a CubeSat constellation to a flagship
- **Reconfigurable spectral filtering and real-time optical processing ahead of the detector, reducing downstream data volume**
- Scalability to hundreds or thousands of modes is what makes this a platform rather than a single instrument



AstroPIC concept: HWO pupil to microlens array to MZI mesh (Sirbu JATIS 2026)

What they are and how they work


- **A comb is a laser source that emits a regular grid of lines — an optical ruler for a spectrograph**
 - Three stages: generate a narrowband comb, broaden it nonlinearly (more than an octave is achievable), then flatten the spectrum so the lines are usable across the band
 - On the detector the lines appear as a dense, evenly spaced array of spots; the more lines, the better the wavelength solution
- **The line spacing and the absolute frequency can be locked to a reference**
 - GNSS-disciplined: relative uncertainty 10^{-12} to 10^{-13} ; self-referenced $f-2f$: 10^{-18} and beyond
- **Three comb families in use**
 - Mode-locked laser combs · electro-optic combs (~10 GHz spacing — PARVI, HPF, IRD, HISPEC/MODHIS) · soliton Kerr microcombs, which are the route to a chip-scale system
- The result is a wavelength measurement traceable to fundamental constants and repeatable over decades — not a lamp whose lines drift as it ages




Comb generation, broadening and flattening (2023 Astrophotonics Roadmap)

State of the field


- Combs are operational on ground-based facilities today, delivering cm/s-class radial-velocity calibration
 - Electro-optic combs with ~10 GHz line spacing are the workhorse for PARVI, HPF, IRD and HISPEC/MODHIS
- But a working system is still rack-scale: with the optics bench, fiber amplifiers, locking electronics and thermal control
- Field operation has had mixed results: most combs have lots of technical issues and demand lots of attention.
- Flattening — the step that makes a comb usable — has itself been made photonic
 - A fully packaged all-photonic spectral shaper delivers 38 dB of dynamic range across 250 nm, with better than 0.1 dB RMS stability over 24 hours
 - Insertion loss is still around 11 dB plus coupling, and the device works on a single polarization — both are the kind of gap that closes with investment
- Microresonator combs and photonic integration point at chip-scale, low-power systems; closed-loop flattening and control is being demonstrated now
- Non-astronomy demand is large — optical clocks, timing distribution, metrology — so NASA is buying into a technology curve someone else is also pushing



Keck Planet Finder comb.




A packaged on-chip spectrograph fed by a comb and a supercontinuum source (Jovanovic Opt. Exp. 2022)



All-photonic spectral flattener: AWGs plus an MZI array — 38 dB over 250 nm (Jovanovic+ 2022)

Where they could go in future missions

- **An onboard absolute wavelength standard, stable for the mission lifetime**
 - Makes a spectral feature retrieval reproducible years apart
 - Traceability and repeatability across the ultraviolet, visible and near-infrared in one system
- **Science that only opens with an absolute ruler in space**
 - Redshift-drift experiments and time-domain spectroscopy free of instrumental systematics
 - Optical clock networks in space: gravitational-wave detection, relativistic geodesy, dark-matter searches and tests of general relativity and the equivalence principle
 - Heterodyne spectroscopy and VLBI.
- **What must be solved first:**
 - system complexity,
 - field and launch robustness,
 - generating blue and ultraviolet photons efficiently,
 - miniaturization,
 - and the power budget



Courtesy: S. Eikenberry, CREOL UCF

10^{-18}


Relative frequency uncertainty of a self-referenced comb — the basis of optical-clock science

cm/s


Radial-velocity calibration floor set by a comb-referenced wavelength solution

What they are and how they work

- **A lantern adiabatically transitions one multimode waveguide into N single-mode waveguides**
 - The design rule is mode counting: at least as many single-mode outputs as modes supported at the multimode input. Satisfy it and the transition is in principle lossless; violate it and you throw away the excess modes
 - Standard devices have 3, 6, 19 or 61 ports — which maps directly onto the number of modes the telescope delivers
- **Why it matters: multimode light in, single-mode light out, without violating etendue**
 - It is the bridge between what a telescope delivers and what a chip can accept
- **Three flavours, and the difference is what you can do with the outputs**
 - Standard: each output carries a linear combination of input modes — the light is coherently mixed
 - Mode-selective: a specific input mode maps to a specific output, which is what nulling and mode filtering need
 - Hybrid: extracts chosen spatial modes while passing the rest — science light one way, sensing light the other
- **Made three ways:**
 - Tapered fiber bundles in a low-index capillary,
 - ultrafast laser inscription in glass, or
 - 3D printing




Lantern topologies: few-mode or multicore to multimode, Leon-Saval et al. Opt. Exp (2005)




Input face, output face, and the LP modes involved, Xin et al. JATIS (2024)

State of the field

- **Efficiency is essentially solved for the standard device: 1×19 lanterns with $> 97\%$ throughput at 1550 nm**
 - Changing mode number, or wavelength range requires engineering effort to optimize
 - The real bottleneck is supply, not physics: two research groups and three companies worldwide, fab recipes that take months to tune, and custom preforms at \$30-50k per fiber draw
- **Focal-plane wavefront sensing, on sky, in closed loop**
 - A lantern at the focal plane senses the wavefront where the science light is, which removes non-common-path and chromatic error by construction
 - Demonstrated at Subaru/SCExAO with linear and nonlinear reconstructors, and with neural-network estimators open loop. Operated polychromatically as well.
- **Photonic lantern nulling:** conceived, built and demonstrated in the laboratory, reaching 10^{-3} contrast after applying implicit EFC to correct modal cross-coupling, and now demonstrated on sky
- **Measuring below the diffraction limit:** a 19-port visible lantern feeding an $R \sim 3000$ spectrometer measured a photocenter to 0.05 mas — $0.003 \lambda/D$ — in ten-minute, repeatable observations
- **Hybrid lanterns have moved from concept to hardware;** the best device so far, 7 ports at 85% throughput, was delivered to NASA Langley




On-sky closed loop: error rejection and Strehl (Subaru/SCExAO) - Lin et al. Opt. Lett. 2024




Spectroastrometry of an H-alpha disk, velocity channel by channel — photocentre to 0.05 mas, Kim et al, ApJL 2025.

Where they could go in future missions


- **The natural interface at an HWO focal plane**
 - A hybrid lantern placed on a detected planet routes the fundamental mode to a spectrometer while the residual speckle field couples into the higher-order modes and drives wavefront control (APRA submitted in 2026 – PI Mona El Morsy)
 - That is simultaneous sensing and science from the same photons — and it holds the wavefront stable through the 10-plus days of integration an exo-Earth spectrum requires
 - The fundamental mode also spatially filters residual starlight, giving a further factor of a few in contrast for free
- **Reserve some modes for nulling and combine the rest coherently for sensing, and the same device becomes a mid-contrast instrument working inside $0.5\text{--}2 \lambda/D$ — angular scales a conventional coronagraph struggles to reach**
- **Combine a ULI lantern with a PIC and you get advanced Electric field control** – can realize the Matched Electric Field Spectrometer (MEFS)
- **Beyond high contrast:** crowded fields, multi-object spectroscopy and polarimetry



Matched Electric Field Spectrometer (MEFS),
Image courtesy of Yoo Jung Kim (APRA
submitted in 2026 – PI Anusha Pai Asnodkar)




The target: a simulated inner Solar System seen from 33 light years





Hybrid lantern: 18 sensing outputs plus a mode-selective science output
Norris et al. Opt. Exp. (2022)

What they are and how they work

- **The arrayed waveguide grating is the workhorse. Five parts in order:**
 - input waveguide,
 - input free-propagation region,
 - waveguide array,
 - output free-propagation region,
 - output plane
- The free-propagation regions play the role of collimator and camera; the array plays the role of the grating
- Adjacent array waveguides differ in optical path by a fixed increment.
- Dispersion is set by lithography, not by a ruled surface, so it is reproducible from die to die across a whole wafer
- Resolving power scales with the total optical path difference across the array, so with order times number of arms; free spectral range goes inversely with order, so high resolution costs unambiguous bandwidth and demands cross-dispersion, exactly like an echelle
- The AWG is not the only route: planar concave gratings and on-chip echelles, cascaded ring filter banks, stationary-wave integrated Fourier-transform spectrometers, and disordered-medium spectrometers all have their place





AWG architecture: input FPR, waveguide array, output FPR




At the design wavelength the phase ramp is flat; off it, the focus walks along the output plane

State of the field

- **Laboratory performance now spans the useful trade space**
 - Doped silica:
 - $R \sim 30,000$, FSR = 13 nm at 40% efficiency,
 - $R \sim 7,000$ FSR = 52 nm at 78% efficiency,
 - $R \sim 180$ FSR = 220 nm at 75% efficiency
 - Silicon nitride:
 - $R \sim 26,000$, FSR = 15 nm with 20% efficiency
 - The practical ceiling today is a few times 10^4 , and pushing resolving power costs throughput and die area
- **On-sky since 2012, and repeatedly since**
 - Stellar spectra in J and H taken through AWGs at facility telescopes, with a single order at the output or several orders separated by cross-dispersion
 - Cross-dispersion can be done off chip with a prism, or on chip by cascading AWGs; several beams can be injected at once and separated the same way
- **Devices are being packaged as compact, self-contained units rather than run as bench experiments**
- Bandwidth-efficiency and detector integration are the two limits that decide whether these compete end to end with a bulk spectrograph




AWG mask layout · the same device, packaged, Greg Sercel and Pradip Gatkine (Caltech)



On-sky J and H band stellar spectra taken through an AWG, Jovanovic et al, Opt. Exp. (2017)

Where they could go in future missions

- **Spectroscopy inside a mass and volume envelope that a bulk spectrograph cannot meet**
 - The same chip serves a small probe and a flagship
 - Replicated chips covering the band simultaneously
- **Integration with the rest of the instrument is the real prize**
 - A spectrometer fed straight from a coronagraph or a lantern, sharing one optical path — spatial and spectral information recorded together, with no separate relay to align or stabilise
 - Multi-object and integral-field configurations where many channels are processed in parallel on one die
- **Broader astrophysics:** anywhere the requirement is many stable, identical spectrographs rather than one exquisite one
 - stellar classification
 - redshift measurement
 - time-domain follow-up.
- Platform development is extending the reach toward the ultraviolet and out into the mid-infrared, which opens thermal-emission and cosmic-web science



Cross-dispersed AWG orders: broad bandwidth on one small detector
 Jovanovic et al, Opt. Exp. (2017)

$R \sim 7,000$

Demonstrated in doped silica at 78% efficiency across J + H — in a few cm² of die

75 %

Efficiency of a low-resolution device across 220 nm of bandwidth

$N \times 1$

One mask, one process — the twentieth die is identical to the first

What is holding facility and flight adoption back

Every one of the five areas runs into some subset of the same six problems. These are the places where investment changes the answer.

Throughput

Every interface costs you, and a photonic instrument has many. End to end, a good bulk spectrograph still beats a photonic one today.

Scaling to many channels

High resolution and broad bandwidth mean hundreds to thousands of waveguides on a die. Routing, crosstalk and channel-to-channel uniformity all get hard at once.

Bandwidth

Most photonic functions are resonant or dispersive and pay for bandwidth in area or free spectral range. The deepest contrast demonstrations to date were monochromatic.

Polarization


Most platforms are designed and optimised for a single linear polarization, so half the light needs a duplicated path or a splitter and rotator.

Detector integration


Photonics is cheap to replicate; infrared detectors are not. Efficient hybridization of many waveguide outputs onto a 2D array is the unsolved system problem.

The space environment

Radiation hardness, thermal cycling, cryogenic operation and launch survival have barely been characterised for most photonic platforms.



280+ waveguides on one die — routing is the limiting craft Greg Sercel and Pradip Gatkine (Caltech)



Cheap to replicate the photonics; costly to replicate the IR detector behind it (Jovanovic Opt. Exp. 2022)

How we advance photonics to the point where a mission can use it

The workshop's next-step suggestions, condensed. Most of these are cheap relative to the leverage they buy.

- 1 Seed the early stage**

Agile \$100-300k, 1-2 year grants to fabricate and test a prototype, or to adapt a proven telecom, quantum or metrology component to astronomy. Reward interdisciplinary teams, since these proposals fall between traditional programme boundaries.
- 2 Make the infusion path explicit**

Maintain a living roadmap that names, for each technology, the performance milestone required, by when, and which mission concept would use it. Many devices sit at TRL 3-4 with nowhere to go next.
- 3 Build testbeds, then fly**

Photonics-specific facilities — vacuum, thermal cycling, vibration, precision metrology — then a stepped path: lab, balloon or sounding rocket, CubeSat, secondary instrument, mission. Roman's contributed coronagraph modes are the template for how that last step actually happens.
- 4 Grow the workforce**

Fellowships for instrumentalists rather than analysis alone, summer schools at the photonics-astronomy boundary, and two-way internships with photonics industry.
- 5 Partner where the capability lives**

Foundry access and process development with commercial and academic fabs, SBIR Phase II and III to carry small companies through, and designated technologist liaisons connecting mission formulation teams to device developers.

Four things to take away

1 Photonics is a replication and stability technology, not only a miniaturization technology.

Function comes from lithographically defined geometry, so the tenth and the hundredth device are identical to the first — and a centimeter-scale monolith is far easier to hold stable than a meter-scale bench.

2 Five areas carry the near-term leverage for NASA.

Programmable PICs, space-compatible laser frequency combs, photonic lanterns, integrated spectrometers, and efficient injection and interconnects.

3 The field has facility-grade proof, and it took sustained investment to get it.

GRAVITY delivers science nothing else can, on the back of two decades of effort. Lanterns and PIC nullers are now working on sky. Most of the rest sits at TRL 3-4 with no defined path forward.

4 The gap is programmatic as much as technical.

Seed funding, a roadmap tied to real missions, photonics testbeds, flight demonstration slots and foundry partnerships are modest investments that would put these technologies within reach of the concepts this community is preparing for the next Decadal Survey.