Dynamic Confocal Profilometry with LCoS Spatial Light Modulators

System description, the roles an LCoS can play, and the device-design trade-offs

Technical note · October 23, 2026 · D. F. Hakala

Summary. A confocal profilometer measures surface height by finding, at every (x, y), the focus position that maximizes light returning through a pinhole. A *dynamic* system makes that confocal geometry programmable: the pinholes, the focal plane, the number of foci and the wavefront are set in software. A reflective LCoS can do this from two places in the optics. In a **field-conjugate** plane, used in amplitude mode behind a polarizing beam splitter (PBS), it is a programmable pinhole or spot pattern. In a **pupil-conjugate** plane, used in phase-only mode at a small off-axis angle, it is a digital varifocal lens, a multi-spot hologram and an aberration corrector at once. The main device trade-off is between **pixel pitch** (optical power, steering range), **cell gap** (phase stroke, speed, fringing) and **speed** (frame rate, flicker), with polarization, flatness and light efficiency as further constraints.

1. Confocal profilometry: the measurement principle

A light source is focused to a diffraction-limited spot on the surface. The reflected light travels back through the objective and a tube lens and is focused onto a pinhole placed in a plane conjugate to the focal spot. When the surface sits in the focal plane, the return beam refocuses tightly and passes through the pinhole. When the surface is above or below focus, the return beam is defocused at the pinhole plane and most of it is blocked (Figure 1). The pinhole therefore converts height into signal strength.

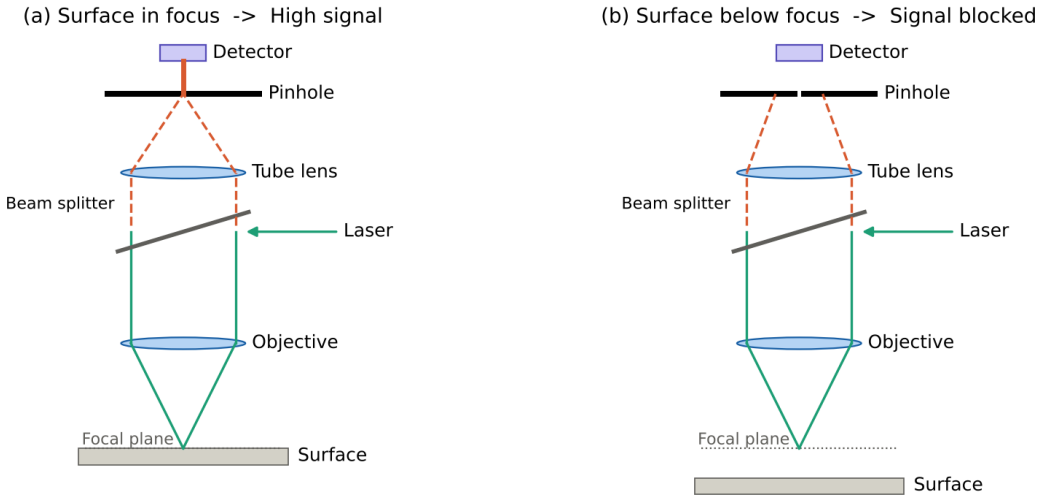


Figure 1. Confocal principle (generic, before any LCoS is added). (a) Surface at the focal plane: the reflected light refocuses through the pinhole. (b) Surface displaced: the return beam is out of focus at the pinhole and is mostly blocked.

Scanning focus relative to the surface produces an axial response curve I(z) at each lateral position. For a plane reflector and a small pinhole, the response follows a sinc-squared law in the normalized defocus u:

$$I(u) \propto \left[\frac{\sin(u/2)}{u/2}\right]^2, \qquad u = \frac{8\pi}{\lambda}\, n\, z \sin^2\left(\frac{\alpha}{2}\right) \tag{1}$$

Here n is the immersion index and α = arcsin(NA/n) is the half-angle of the objective. The full width at half maximum of this curve is:

$$\Delta z_{\mathrm{FWHM}} \approx \frac{0.44\,\lambda}{n\,(1-\cos\alpha)} \tag{2}$$

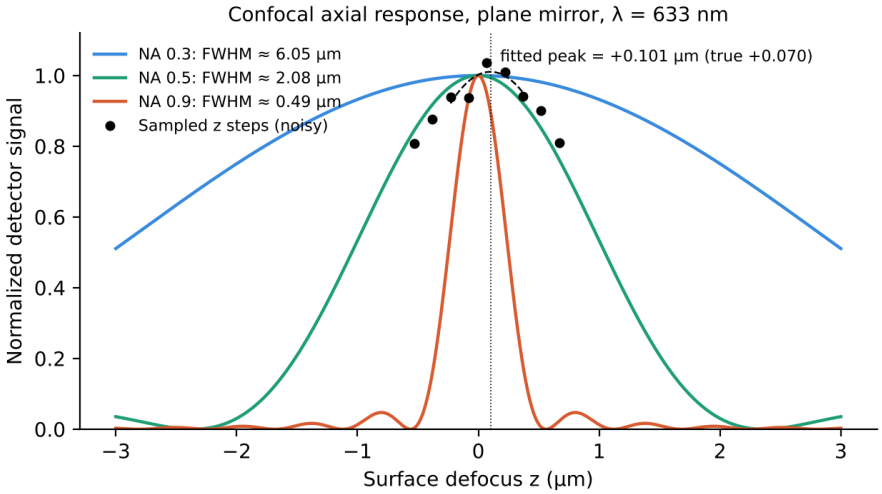


Figure 2. Axial response for three numerical apertures at 633 nm. The black points are noisy samples at 150 nm steps. A parabola fitted to the top five points locates the peak to a few tens of nanometers, far finer than the FWHM.

Height is taken as the **position of the peak**, not read from the width. Centroid, parabolic, or Gaussian fits to a handful of samples near the peak locate it to a small fraction of the FWHM. The precision scales roughly as FWHM divided by signal-to-noise ratio, which is why nanometer-level repeatability is achievable with sub-micrometer FWHM. Lateral resolution is set by the confocal spot, about 0.4 λ/NA.

Repeating the measurement at every (x, y) gives a height map. Common ways to cover the field are:

- **Point scanning** with galvanometer mirrors: one pinhole, highest sectioning, slowest.
- **Spinning Nipkow disk or microlens-pinhole array**: thousands of pinholes in parallel onto a camera.
- **Programmable pinhole or spot array**: the pattern is displayed on an SLM (Section 4).
- **Chromatic confocal**: a dispersive objective maps wavelength to depth, so a spectrometer reads height without a z scan.

2. What makes a confocal profilometer “dynamic”

The term is not fully standardized. In practice it covers three capabilities, and an LCoS can contribute to each:

Capability	What changes	LCoS role
Reconfigurable confocal geometry	Pinhole or spot size, spacing and pattern are set in software	Amplitude LCoS in a field plane (behind a PBS)
Axial scanning without moving parts	Focal plane is stepped optically, not by a piezo stage	Phase LCoS in a pupil plane (lens term)
Fast or adaptive measurement	Spots track the surface; sampling concentrates near each peak; moving parts can be followed	Phase LCoS (multi-spot hologram) plus closed-loop software

3. How a reflective LCoS sits in the beam

An LCoS is a reflective device: the liquid-crystal layer sits on a mirrored silicon backplane, so light enters and leaves through the same cover glass. The incoming and outgoing beams therefore have to be separated, and the method depends on the mode of operation (Figure 3).

- **Amplitude (light-valve) mode.** The LC rotates the polarization of "on" pixels. A polarizing beam-splitter cube does the separation, as in an LCoS projector: s-polarized light reflects off the PBS hypotenuse onto the LCoS; "on" pixels return p-polarized light, which passes straight through the PBS; "off" pixels return s-polarized light, which the PBS sends back toward the source.
- **Phase-only mode.** A parallel-aligned (ECB) LCoS changes only the phase of light polarized along the LC director; it does not rotate the polarization. A PBS cannot separate the input from the output, so the device is used at a small angle of incidence (a "V" geometry, typically 5–10°), or behind a non-polarizing beam splitter (NPBS) at the cost of light.

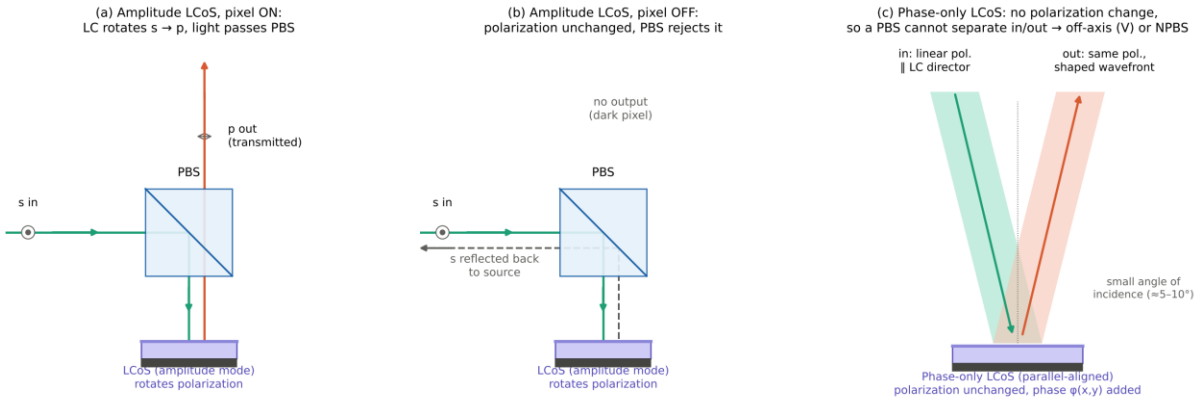


Figure 3. Separating input and output on a reflective LCoS. (a, b) Amplitude mode behind a PBS: the pixel state decides whether light is transmitted (p) or sent back (s). (c) Phase-only mode: the polarization is unchanged, so input and output are separated by angle.

4. A reference system architecture

Figure 4 shows one physically consistent way to combine both LCoS functions in a confocal profilometer. The path, in order:

1. **Source.** A narrowband laser is expanded and its polarization set to s by a half-wave plate.
2. **LCoS-A (amplitude, field-conjugate).** PBS-1 reflects the s light onto LCoS-A, which displays the pinhole or spot pattern. "On" pixels return p light, which passes through PBS-1. "Off" pixels return s light toward the laser, so an optical isolator is advisable.
3. **L1 (Fourier lens).** L1 collimates the light from each "on" pixel, so the pattern's angular spectrum fills LCoS-B. LCoS-B therefore sits in a pupil-conjugate plane.
4. **NPBS (detection pick-off).** A 50:50 non-polarizing beam splitter transmits the illumination toward LCoS-B and later reflects the returning light to the camera. The half of the illumination it reflects goes to a beam dump.
5. **LCoS-B (phase-only, pupil-conjugate).** It is used at a small angle of incidence with its LC director aligned to the p polarization. It adds the lens, tilt, hologram and Zernike terms (Equation 5). Fold mirror M2 picks off the reflected beam.

6. **Relay L2–L3 and field stop.** A 4f relay images LCoS-B onto the objective pupil. The intermediate image between L2 and L3 is field-conjugate, and a stop there blocks the SLM zero order when a tilt carrier is used.
7. **Objective and sample.** The objective focuses each spot onto the surface. The reflected light retraces the same path back through LCoS-B, which applies the same phase again, then the NPBS reflects it through an analyzer and lens Lc onto the camera. The camera is field-conjugate, and its pixels at the spot positions act as virtual pinholes.

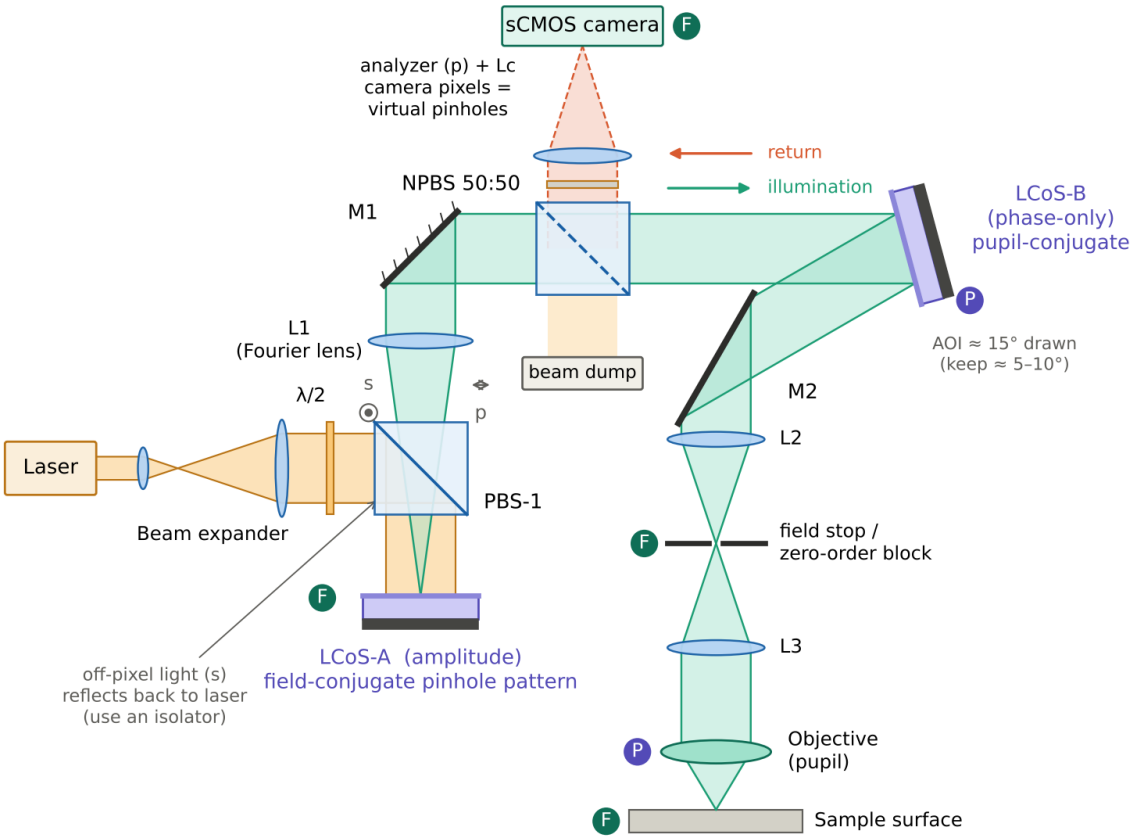


Figure 4. Reference layout with reflective devices (not to scale; the LCoS-B angle of incidence is exaggerated). Green: illumination; orange-red dashed: return light to the camera. F = field-conjugate plane, P = pupil-conjugate plane. Polarization symbols: ⊙ s (normal to the page), ↔ p (in the page).

Why the design is arranged this way

- **The return light must stay co-polarized through LCoS-B.** When LCoS-B shifts the focus, the light reflected from the surface comes back carrying the same curvature, and LCoS-B must apply the same phase a second time to re-collimate it. A phase-only LCoS modulates only one polarization, so the usual polarization isolator (PBS plus quarter-wave plate before the objective) would turn the return light orthogonal, LCoS-B would no longer act on it, and confocality would be lost whenever the focus is moved. The NPBS keeps the polarization unchanged. The cost is light: at most 25% of the available signal reaches the camera (50% on each pass).

- **LCoS-A sets the illumination, and the camera does the pinholing.** A double-pass programmable array, where the same pixels also act as the detection pinholes, is possible. Behind a PBS, however, the illumination rejected by "off" pixels leaves through the same port as the signal. Using camera pixels as virtual pinholes keeps that light out of the detector and allows pinhole size to be changed in software after acquisition.
- **A phase-only system does not need LCoS-A.** With a fixed illumination (or a hologram on LCoS-B producing the spot array), LCoS-B alone provides z scanning, multi-spot and aberration correction. This is the most common configuration in practice.
- **Alternatives to the NPBS.** A polarization-independent phase modulator can be built by splitting LCoS-B into two halves with a half-wave plate, so both polarizations are modulated; a PBS and quarter-wave plate can then separate input and return with little loss. This is more efficient but more complex to align.

5. How the LCoS is used

5.1 Programmable pinhole or spot pattern (amplitude mode, field-conjugate plane)

LCoS-A behind PBS-1 is a programmable illumination mask in a plane conjugate to the sample. A sparse grid of "on" pixel groups creates a grid of illuminated spots; the grid is stepped across the field to build a full confocal image, with the camera reading only the pixels at the current spot positions. This follows the programmable array microscope (PAM) concept [3]. Pseudo-random patterns combined with subtraction of a widefield image give the aperture-correlation method [2], which uses light more efficiently.

All the key parameters become software choices:

- **Spot size**, typically 0.5–1 Airy unit at the sample: smaller improves sectioning, larger improves signal.
- **Spot spacing**: closer spacing increases throughput but adds crosstalk between neighbors, which degrades sectioning.
- **Number of frames**: equal to the spacing squared for a regular grid (for example 25 frames for 5 × 5 spacing).

The weakness is contrast ratio. "Off" pixels leak about 1/CR of the light, adding a widefield background. With on-fraction f, the leaked illumination relative to the confocal illumination is about (1 − f)/(f · CR). For f = 1/25 and CR = 1000:1 this is 2.4%; at CR = 100:1 it is 24%, which visibly reduces sectioning. Nematic LCoS also settles in milliseconds, limiting pattern rates to hundreds per second. Ferroelectric LCoS reaches kHz rates in binary mode. A DMD is often the better choice for this specific role because of its higher contrast and speed.

5.2 Digital varifocal lens for z scanning (phase mode, pupil-conjugate plane)

This is the strongest use of a phase-only LCoS. With LCoS-B conjugate to the objective's back aperture, it displays a wrapped quadratic phase that adds defocus to the whole field:

$$\varphi(r) = -\frac{\pi r^2}{\lambda f_{\mathrm{SLM}}} \pmod{2\pi} \tag{3}$$

Changing f_{SLM} steps the focal plane through the sample with no moving mass. The steps are vibration-free and exactly repeatable. Because LCoS-B sits in the common path and the return light stays co-polarized, the illumination focus and the detection conjugate move together, so the system stays confocal at every z. This is a form of remote focusing [8].

Sampling limit. The local fringe period at the edge of the aperture must be at least about two pixels. For pixel pitch p and illuminated aperture D, this sets a minimum focal length:

$$f_{\min} \approx \frac{p\,D}{\lambda}, \qquad \Delta z_{\max} \approx \frac{f_{\mathrm{obj}}^2}{M^2 f_{\min}} \tag{4}$$

Here M is the lateral magnification from the SLM to the objective pupil. As an example, take a 50×/0.8 objective (f_{obj} = 4 mm, pupil diameter 6.4 mm), an SLM with 8 µm pixels illuminated over 15 mm (M ≈ 0.43), and 633 nm light. Then f_{min} ≈ 190 mm and the paraxial range is about ±0.45 mm. That range is ample for most profilometry, so in practice it is not the binding constraint. The binding constraints are:

- **Efficiency near the limit.** Diffraction efficiency falls as the local period shrinks toward 2 pixels (Figure 8), so the usable range is a fraction of Δz_{max}.
- **High-NA phase form.** At high NA a quadratic phase adds spherical aberration as well as defocus. The exact defocus phase, $k\,z\,\sqrt{(n^2 - NA^2\,\rho^2)}$, has a steeper edge gradient, roughly $1/\cos\theta_{max}$ steeper (1.7× at NA 0.8). This shrinks the range proportionally but keeps the focus aberration-free.
- **Zero order.** Unmodulated light (pixel gaps, imperfect phase) still focuses at the nominal plane and creates a fixed "ghost" confocal peak that can bias the height fit. The standard remedy is to add a blazed tilt carrier so the useful light leaves at an angle, then block the zero order at the intermediate field stop (Figure 4).
- **Off-axis use.** The small angle of incidence on LCoS-B slightly foreshortens the aperture in one direction and adds a little astigmatism; both are corrected by the Zernike terms on the same device.
- **Chromatic behavior.** A diffractive lens has $f \propto 1/\lambda$, so this method needs a narrowband source. The dispersion can also be exploited deliberately as a hybrid with chromatic confocal sensing.

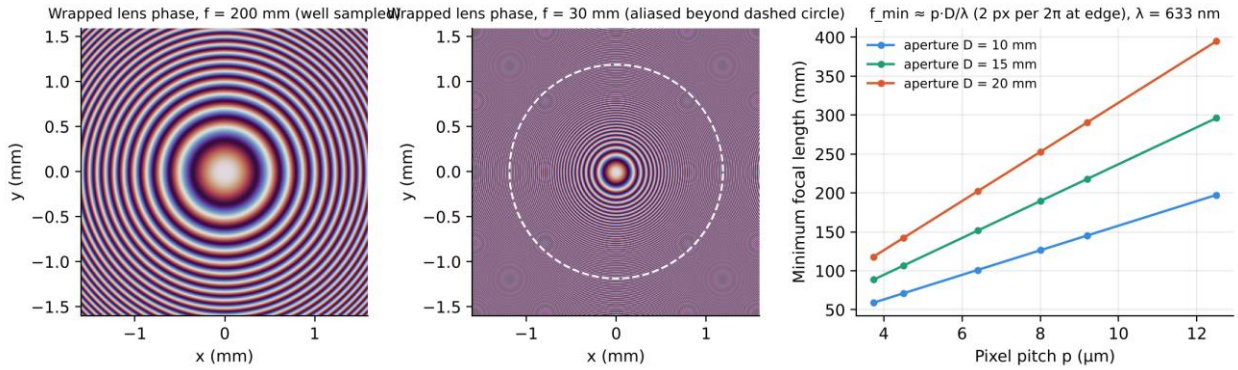


Figure 5. Left and center: wrapped lens phase on a 400 × 400 pixel, 8 µm region. At f = 30 mm the fringes beyond the dashed circle are finer than two pixels and alias. Right: minimum focal length from Equation 4 against pixel pitch for three aperture sizes.

5.3 Parallel and adaptive multi-spot illumination (holographic)

LCoS-B can also display a hologram that creates an array of foci instead of one (Figure 6). Each spot can carry its own lens term, and therefore its own z. The camera pixels act as virtual detection pinholes. The holograms are computed with Gerchberg–Saxton [4] or weighted Gerchberg–Saxton (WGS) [5]. This is the same machinery used for optical tweezer arrays, and it brings the same issues: spot uniformity, ghost orders, and total efficiency.

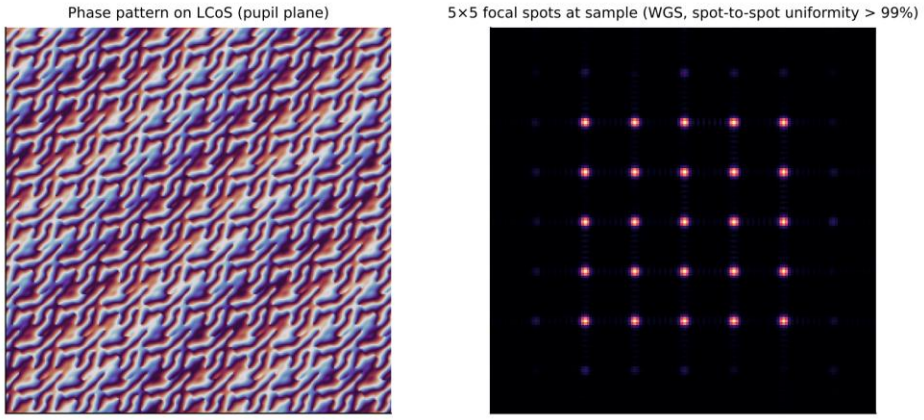


Figure 6. Weighted Gerchberg–Saxton hologram (left) and the resulting 5 × 5 spot array (right, square-root intensity scale to show the weak ghost orders).

What makes the system dynamic is closed-loop tracking:

1. Run a coarse z scan to get an approximate height at each spot location.
2. Recompute the hologram so that each spot sits near its local surface height.
3. Sample only a few z offsets around each peak (for example ±2 FWHM in 5 steps) and fit.
4. Update the per-spot lens terms each cycle to follow a slowly moving or deforming surface.

This cuts the number of frames by roughly the ratio of full scan range to the tracked window, often an order of magnitude or more.

5.4 Aberration correction

Zernike terms added to the same pupil pattern can compensate for system aberrations, for SLM and backplane non-flatness, for the off-axis angle on LCoS-B, and for refractive-index mismatch when measuring through a cover layer. They can also correct the spherical aberration that grows with large digital defocus. All of these broaden, lower and skew the axial response, which degrades height accuracy. A sensorless modal approach works well in a confocal system [7]: for each Zernike mode, apply small positive and negative amounts, measure the peak confocal signal, and fit a parabola to find the best coefficient. This needs about 2N + 1 measurements for N modes. The companion note covers Zernike correction in detail.

In practice all four functions are combined into a single pattern on LCoS-B:

$$\varphi_{\mathrm{SLM}} = \left[\varphi_{\mathrm{lens}}(z) + \varphi_{\mathrm{tilt}} + \varphi_{\mathrm{holo}} + \textstyle\sum_j a_j Z_j - \varphi_{\mathrm{flat}}\right] \bmod 2\pi \tag{5}$$

ϕ_{flat} is the measured flatness map of the device. The wrapped phase is then converted to gray levels with a look-up table (LUT) for the operating wavelength and temperature.

6. LCoS device design: issues and trade-offs

Figure 7 shows the layer stack of a reflective phase LCoS. Light passes through the liquid crystal (LC) twice, so the phase stroke is doubled for a given cell gap d. Most of the design trade-offs follow from three numbers: the pixel pitch p, the cell gap d, and the LC material's birefringence Δn and rotational viscosity γ_1.

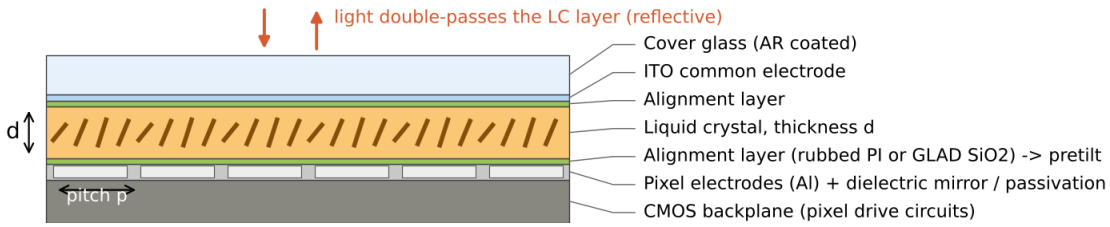


Figure 7. Schematic reflective LCoS cross-section (not to scale). Pretilt from the alignment layer, whether rubbed polyimide or obliquely deposited SiO_2, sets the switching threshold and helps suppress disclinations at pixel edges.

6.1 Phase stroke, cell gap and speed

A phase-only device needs at least 2π of stroke at the working wavelength. In reflection this requires:

$$2\,\Delta n\,d \geq \lambda, \qquad \tau_{\text{off}} \approx \frac{\gamma_1\,d^2}{K\,\pi^2} \tag{6}$$

The relaxation time τ grows with the square of the cell gap, which drives the central trade-off. As an illustration, take Δn = 0.2, γ_1 = 0.15 Pa·s and K = 12 pN. At 633 nm the minimum gap is about 1.6 µm, giving τ ≈ 3 ms. At 1064 nm the gap grows to about 2.7 µm and τ ≈ 9 ms. Designers usually add 10–20% stroke margin (about 2.2π) for LUT linearization and temperature drift, which costs speed. Higher-Δn materials allow thinner cells but tend to bring higher viscosity, lower stability or more absorption. Overdrive waveforms shorten the rise time but not the free relaxation.

6.2 Pixel pitch and fringing fields

A smaller pitch increases the maximum steering angle (λ/2p), the lens power (Equation 4) and the multi-spot field of view. However, the electric field from one pixel spreads into its neighbors over a distance comparable to d. When d/p is large, the sharp 2π phase resets that a blazed grating or lens needs are smoothed into "flyback" regions that diffract light into unwanted orders. Figure 8 illustrates the effect with a simple smoothing model. As a rule of thumb, fringing becomes significant once d approaches roughly half the pitch or more. Small pitch therefore pushes toward thin cells, and thin cells push toward high Δn and short wavelengths. Fringing can be mitigated with:

- thinner cells and higher-Δn LC,
- pretilt and alignment engineering to reduce disclination lines at pixel boundaries,
- pattern pre-compensation using a measured or simulated crosstalk model.

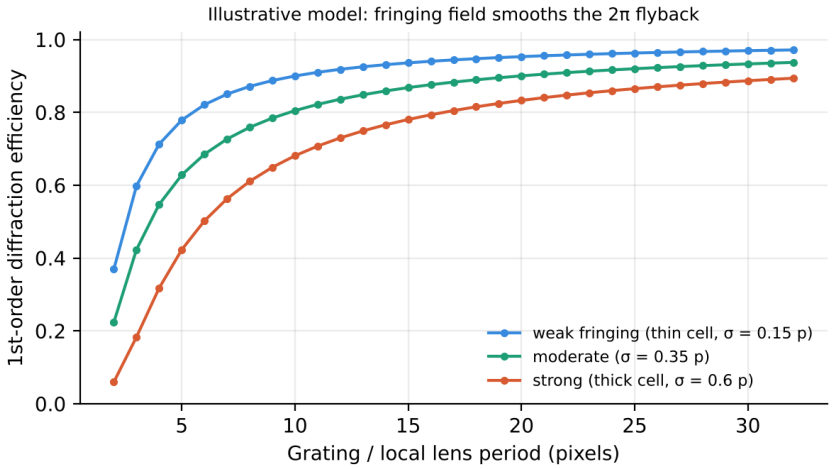


Figure 8. Illustrative model of first-order efficiency against period in pixels, for three levels of fringing (Gaussian smoothing of the phase profile with width σ). It shows the trend only; real devices should be characterized directly.

6.3 Fill factor, quantization and zero order

The inter-pixel gaps and any phase error put light into the zero order. A high fill factor (above 90%) and a dielectric mirror over the pixel electrodes reduce both zero-order leakage and absorption. Phase quantization with N levels limits first-order efficiency to $\eta = \mathrm{sinc}^2(1/N)$: 95% for 8 levels, 98.7% for 16, and effectively 100% for 8-bit drive. Bit depth is rarely the limit, but the LUT must be linear in phase, not in gray level.

6.4 Phase flicker

LC cells must be driven with a DC-balanced (alternating) field. Many backplanes, especially digital pulse-width-modulated ones, produce a small phase ripple at the frame or sub-frame rate. In a confocal system this appears as signal noise and can bias the peak fit. Mitigations include:

- higher addressing frequency,
- analog drive,
- LC materials with higher viscosity, at a speed cost,
- synchronizing camera exposures to an integer number of flicker periods.

Measure the ripple on the actual device before setting z-step dwell times.

6.5 Flatness, polarization, wavelength and temperature

- **Flatness:** Backplane and cover-glass warp can amount to one or several waves across the aperture. Measure it interferometrically and subtract it as ϕ_{flat} (Equation 5). This is what brings the focus close to diffraction-limited.
- **Polarization:** A parallel-aligned (ECB) phase LCoS modulates only light polarized along the director. The input must be linearly polarized, and the return path must stay co-polarized (Section 4). Depolarizing or birefringent samples lose part of the return signal; the analyzer in front of the camera removes the unmodulated orthogonal component.
- **Wavelength:** Δn depends on λ, so each laser line needs its own LUT, and the 2π stroke must be met at the longest wavelength used.
- **Temperature:** Δn and γ_1 drift with temperature, which shifts the LUT and response time. Active stabilization (on the order of ±0.1 K) or temperature-indexed LUTs keep the z scale calibrated.

- **Optical power:** Absorption in aluminum electrodes heats the cell. A dielectric mirror raises reflectivity and damage threshold, and inorganic alignment layers such as obliquely deposited SiO_2 are more robust under high flux than polyimide.

6.6 Trade-off summary

Design choice	Improves	Costs / risks
Smaller pixel pitch	Steering angle, lens power and z range, spot-array field	Fringing (needs thinner cell), lower fill factor, more pixels to drive
Thinner LC cell	Speed ($\tau \propto d^2$), less fringing	Needs higher Δn for 2π, tighter gap uniformity
Higher-Δn LC	Thinner cell for the same stroke	Often higher viscosity, absorption or UV/thermal instability
More stroke margin (>2π)	LUT linearity, temperature tolerance	Thicker cell, slower
Digital (PWM) backplane	Bit depth, uniformity, cost	Phase flicker at sub-frame rate
Analog backplane	Low flicker	Pixel-to-pixel nonuniformity, lower bit depth
Dielectric mirror	Reflectivity, damage threshold, less heating	Extra voltage drop across the mirror, so more drive voltage
Ferroelectric LC (binary)	kHz frame rates	Binary phase only: ≤ 40.5% first-order efficiency, symmetric ghost order
NPBS for the return path	Return stays co-polarized, so remote focusing works	≤ 25% round-trip throughput
Split-SLM polarization-independent scheme	Allows PBS + $\lambda/4$ isolation with low loss	Half the pixels per polarization, harder alignment
Large aperture / pixel count	Field of view, many spots	Flatness harder, slower data transfer, cost

7. Suggested starting specification

These targets are engineering judgment for a visible-wavelength phase LCoS in a confocal profilometer. They are a starting point for trade studies, not a vendor specification.

Parameter	Suggested target	Reason
Modulation	Phase-only, parallel-aligned nematic	Lens, hologram and Zernike terms in one device
Pixel pitch	4–8 µm	Balances lens power against fringing
Resolution	≥ 1920 × 1080	Spot count and smooth lens edges
Phase stroke	≥ 2.2π at longest λ	LUT margin and temperature drift
Response time	≤ 3–5 ms (10–90%)	Z-step rates of ~100–300 per second
Phase ripple	≤ 0.05π rms, measured	Keeps peak-fit bias below the noise
Flatness	≤ λ/10 rms after correction	Near-diffraction-limited focus
Fill factor / reflectivity	≥ 90% / ≥ 90% with dielectric mirror	Efficiency and zero-order suppression
Thermal control	On the order of ±0.1 K, or temperature-indexed LUTs	Stable z calibration

8. When the LCoS approach fits, and when it does not

LCoS wins	LCoS loses
No moving parts; z steps are digital and exactly repeatable	Speed: nematic devices run at ~60–500 Hz, while piezo, ETL and TAG lenses reach kHz or more
One device provides focus, tilt, multi-spot and aberration correction	Light efficiency after polarizers, beam splitters, diffraction and double passes
Adaptive sampling and surface tracking	Single polarization; narrowband operation
Confocality preserved during refocusing (common-path remote focus)	Zero order and ghost orders need managing

A practical hybrid uses a mechanical or electrically tunable lens (ETL) stage for the coarse range and the LCoS for fine focus, aberration correction and multi-spot shaping.

Further reading

[1] T. Wilson (ed.), *Confocal Microscopy*, Academic Press, 1990.

[2] R. Juškaitis, T. Wilson, M. A. A. Neil, M. Kozubek, "Efficient real-time confocal microscopy with white light sources," *Nature* 383, 804–806 (1996).

[3] P. J. Verveer, Q. S. Hanley, P. W. Verbeek, L. J. van Vliet, T. M. Jovin, "Theory of confocal fluorescence imaging in the programmable array microscope (PAM)," *J. Microsc.* 189, 192–198 (1998).

[4] R. W. Gerchberg, W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik* 35, 237–246 (1972).

[5] R. Di Leonardo, F. Ianni, G. Ruocco, "Computer generation of optimal holograms for optical trap arrays," *Opt. Express* 15, 1913–1922 (2007).

[6] R. J. Noll, "Zernike polynomials and atmospheric turbulence," *J. Opt. Soc. Am.* 66, 207–211 (1976).

[7] M. J. Booth, "Adaptive optical microscopy: the ongoing quest for a perfect image," *Light Sci. Appl.* 3, e165 (2014).

[8] E. J. Botcherby, R. Juškaitis, M. J. Booth, T. Wilson, "Aberration-free optical refocusing in high numerical aperture microscopy," *Opt. Lett.* 32, 2007–2009 (2007).

[9] Z. Zhang, Z. You, D. Chu, "Fundamentals of phase-only liquid crystal on silicon (LCOS) devices," *Light Sci. Appl.* 3, e213 (2014).

References are listed from memory as further reading; verify details before citing. Figures 2, 5, 6 and 8 were computed for this note. Figures 1, 3, 4 and 7 are schematic.