LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-T2600448-v1** 2026/10/06

Improving Mode-Matching and BHD Visibility for Waveguide-Based Squeezing at 1064 nm

Nishant Pawar

Mentors: Peter G. Carney and Prof. Rana X. Adhikari

LIGO SURF 2026 – Final Report

California Institute of Technology
LIGO Project, MC100-36
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

LIGO Hanford Observatory
Route 10, Mile Marker 2
Richland, WA 99352
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: info@ligo.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: info@ligo.caltech.edu

http://www.ligo.caltech.edu/

Abstract

Waveguide-based Optical Parametric Amplification (WOPA) is a promising alternative that could simplify the conventional cavity-based squeezed-light sources used in the LIGO detectors: a single pass of 532 nm pump light through a periodically poled lithium niobate (PPLN) waveguide produces squeezed vacuum at 1064 nm without any cavity that has to be locked. How much of this squeezing we actually *detect*, however, depends strongly on the visibility of the Balanced Homodyne Detection (BHD), i.e. on how well the Local Oscillator (LO) beam is mode matched to the squeezed field. This project aimed to improve the BHD visibility of the WOPA tabletop setup. Because squeezed vacuum carries no classical power, we used a bright 1064 nm beam sent through the waveguide as a stand-in for the squeezed mode, profiled both it and the LO, and designed a new two-lens mode-matching solution placed *before* the PZT phase-scanning mirror. With this solution the visibility improved from ∼90% to 91.2%, and the detected squeezing increased from −3.0 dB to −3.20 dB (with ≈ + 3.3 dB of anti-squeezing) at ≈142 mW of pump power. Using a loss model calibrated to this measurement, we predict that reaching 95–97% visibility would raise the detected squeezing to about −3.6 to −3.9 dB at the same pump power, and to about −4.5 to −4.9 dB at 250 mW. A Gaussian-beam (q-parameter) analysis shows that the lens solution by itself allows a visibility close to 99%, so the remaining gap is dominated by alignment and by the non-ideal, elliptical beam leaving the waveguide. Finally, we work out the expected signal-to-noise ratio of a pump-modulation technique proposed to measure the *true* spatial mode of the squeezed field, and report on a search for scattered-light sources that will matter for future low-frequency squeezing measurements.

Contents

1 Introduction **3**

1.1 Gravitational waves and LIGO 3

1.2 Quantum noise 3

1.3 Squeezed light in one picture 3

1.4 Why a waveguide? 4

1.5 Goal of this project 5

2 Background **5**

2.1 Nonlinear optics: how a crystal makes squeezed light 6

2.2 Detecting squeezed light: balanced homodyne detection 8

2.3 Why losses destroy squeezing 9

2.4 Mode matching and visibility 10

3 The WOPA experimental setup **11**

3.1 Laser and beam preparation 11
3.2 The PPLN waveguide 11
3.3 The local oscillator path and the PZT 12
3.4 Detection chain 12
3.5 The “tool beam”: seeing the invisible 12

4 Work and results **13**
4.1 Getting to know the setup and re-commissioning it 13
4.2 Profiling the tool beam 13
4.3 Profiling the local oscillator 13
4.4 Designing the new mode-matching solution 14
4.5 Measuring the visibility 15
4.6 Squeezing measurement 16
4.7 How much squeezing should better visibility give? 17

5 Discussion: what limits the visibility now? **19**

6 Hunting for scattered light **20**

7 Conclusions and future work **20**

Acknowledgements **21**

A Gaussian beams and mode-matching basics **22**

B Measuring the true mode of the squeezed field: an estimate **25**

C Moku signal chain and shot-noise baseline **28**

D Miscellaneous work **28**

1 Introduction

1.1 Gravitational waves and LIGO

Einstein's general theory of relativity (1916) predicts that accelerating masses create ripples in spacetime that travel outward at the speed of light: gravitational waves (GWs). Einstein himself doubted they could ever be measured, because their effect is absurdly small. When a GW from a pair of merging black holes passes the Earth, it stretches space in one direction and squeezes it in the perpendicular direction by a fractional amount (the *strain*) of roughly $h \sim 10^{-21}$ or less. Over a 4 km distance, that is a length change far smaller than the size of a proton.

LIGO measures this tiny effect using laser interferometry. Each LIGO detector is an L-shaped Michelson interferometer with 4 km long arms, enhanced with Fabry–Pérot cavities that bounce the light back and forth many times. A passing GW lengthens one arm while shortening the other, which changes the relative phase of the light returning from the two arms and therefore the amount of light that reaches the output photodetector. In September 2015 the two LIGO detectors (Hanford, WA and Livingston, LA) made the first direct detection of GWs, from two black holes merging about 1.3 billion light-years away [1]. Since then, hundreds of mergers of black holes and neutron stars have been observed.

1.2 Quantum noise

A detector is only as good as its noise. LIGO's sensitivity is limited at low frequencies by seismic and thermal noise, but across much of its most sensitive band (roughly 50 Hz–5 kHz) the limit is *quantum noise*, which comes from the fact that light is made of photons [2]. It shows up in two ways:

- **Shot noise:** Photons arrive at the photodetector randomly, like raindrops on a roof. For an average power P, the number of detected photons in a time interval fluctuates with Poisson statistics, so the relative fluctuation falls only as $1/\sqrt{P}$. This dominates at high frequencies.
- **Radiation-pressure noise:** Each photon that bounces off a mirror gives it a tiny kick. Random photon numbers mean random forces that shake the 40 kg test masses. This dominates at low frequencies.

A beautiful insight due to Caves [2] is that both of these can be traced back to *vacuum fluctuations* of the electromagnetic field leaking into the interferometer through its output port. If we could replace that ordinary vacuum with a "quieter" kind of vacuum, we could reduce quantum noise without increasing the laser power. That quieter vacuum is called *squeezed vacuum*.

1.3 Squeezed light in one picture

A classical light wave can be written as

$$E(t) = X_1 \cos \omega t + X_2 \sin \omega t, \tag{1}$$

where X_1 and X_2 are called the *amplitude* and *phase quadratures*. In quantum mechanics X_1 and X_2 cannot both be known perfectly; like position and momentum, they obey an uncertainty relation, $\Delta X_1 \, \Delta X_2 \geq \text{const}$. For vacuum and for a perfect laser beam (a *coherent state*) the uncertainty is shared equally, and we draw it as a round "blob" in the X_1–X_2 plane (Fig. 1, left). A laser beam is simply this blob displaced from the origin by its classical amplitude; vacuum is the same blob sitting at the origin. The size of the blob is the shot-noise level.

The uncertainty principle only fixes the *product* of the two uncertainties. So we are allowed to shrink the noise in one quadrature by a factor e^{-r} as long as the other grows by e^{+r}; the blob becomes an ellipse. This is a squeezed state (Fig. 1, middle and right), and r is called the *squeezing parameter*. The squeezing level is quoted in decibels relative to the vacuum noise,

$$S_\pm \, [\text{dB}] = 10 \log_{10} \left(\frac{\langle \Delta X_\pm^2 \rangle}{\langle \Delta X_{\text{vac}}^2 \rangle} \right) = 10 \log_{10} \left(e^{\pm 2r} \right) , \tag{2}$$

where "−" is the squeezed quadrature (negative dB, less noise than vacuum) and "+" the anti-squeezed one [3, 11]. For example, −3 dB means the noise power in the squeezed quadrature is half that of vacuum. If we inject squeezed vacuum into LIGO with the correct orientation, the quantum noise in the GW readout goes down. LIGO has routinely used squeezed light since the third observing run in 2019 [6], and frequency-dependent squeezing since the fourth [7], gaining several dB of quantum-noise reduction.

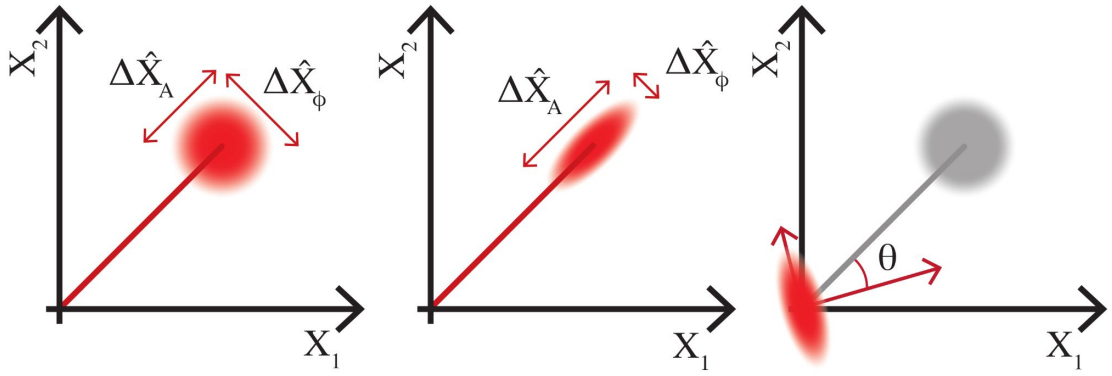


Figure 1: Phase-space ("quadrature") pictures of light, from Ref. [5]. Left: a coherent state (laser beam) is a round uncertainty blob, with equal noise $\Delta \hat{X}_A$ and $\Delta \hat{X}_\phi$ in amplitude and phase, displaced from the origin by its classical amplitude. Middle: a phase-squeezed state, with less phase noise and more amplitude noise. Right: squeezed vacuum (red) has no classical amplitude; its squeezing ellipse is oriented at an angle θ, which in homodyne detection is set by the phase of the local oscillator (grey).

1.4 Why a waveguide?

Squeezed light is made with a *nonlinear crystal*: a material whose polarization responds to the square of the electric field. When pumped with 532 nm ("green") light, such a crystal can split a green photon into two 1064 nm ("infrared") photons. These photons are created

in correlated pairs, and it is exactly this correlation that reduces the noise in one quadrature (Sec. 2.1).

The nonlinear effect is weak, so in LIGO the crystal sits inside an optical cavity (an optical parametric oscillator, OPO) that recirculates the light many times [4, 7]. This works, but the cavity has to be held on resonance and the squeezing phase has to be stabilized, which requires several extra laser beams and feedback loops. All of these add complexity, optical loss and paths through which noise can sneak in.

A *waveguide* offers another way to get a strong nonlinear interaction. By trapping the light in a channel only a few micrometres wide over its whole length, it makes a *single pass* through the crystal strong enough to produce useful squeezing (Sec. 2.1), so there is no cavity to lock at all. Waveguide squeezers are also compact and naturally broadband (there is no cavity linewidth), and waveguide optical parametric amplifiers (WOPAs) at telecom wavelengths have already produced over 8 dB of squeezing [8]. Our group's WOPA experiment explores this idea at LIGO's 1064 nm wavelength [9, 10].

1.5 Goal of this project

Squeezing is fragile: any optical loss replaces part of the squeezed vacuum with ordinary vacuum and washes the squeezing out (Sec. 2.3). One of the largest losses in our setup is *imperfect mode matching* in the balanced homodyne detector: the squeezed light and the local oscillator must have the same spatial shape to interfere fully. The quality of this match is measured by the *visibility* $\mathcal{V}$, and because it enters the detection efficiency as $\mathcal{V}^2$, even a few percent matters.

The objective of this project was therefore to improve the BHD visibility of the WOPA tabletop from its starting value of ∼90%, and to understand how much detected squeezing we gain by doing so. Concretely, I:

1. learned and re-commissioned the WOPA setup and its squeezing measurement;
2. profiled the beam leaving the waveguide and the LO beam;
3. calculated and implemented a new mode-matching lens solution, placed before the PZT mirror;
4. measured the new visibility (91.2%) and the resulting squeezing (−3.20 dB);
5. modelled how much more squeezing better visibility would give, and analysed what currently limits the visibility;
6. estimated the signal-to-noise ratio of a proposed method to measure the true mode of the squeezed field (Appendix B), and searched for scattered-light sources on the table.

2 Background

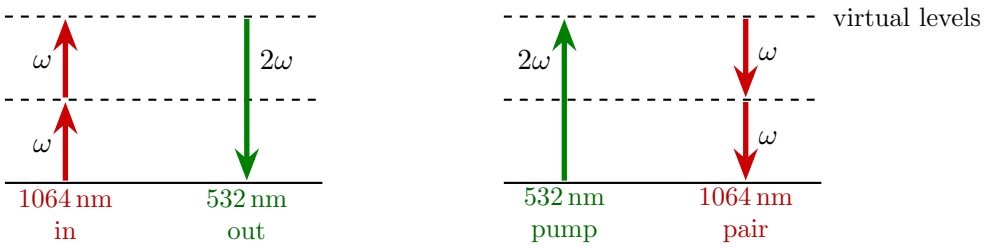
2.1 Nonlinear optics: how a crystal makes squeezed light

A nonlinear response: In everyday (linear) optics, the electric field E of a light wave makes the electrons in a material oscillate at the light's own frequency, and the induced polarization (dipole moment per unit volume) is simply proportional to the field. When the field is strong enough, some crystals respond nonlinearly [13]:

$$P = \epsilon_0 \left[\chi^{(1)} E + \chi^{(2)} E^2 + \chi^{(3)} E^3 + \dots\right]. \tag{3}$$

The $\chi^{(2)}$ term exists only in crystals that lack inversion symmetry, such as lithium niobate ($LiNbO_3$) and KTP, and it is what makes new frequencies possible. Since $\cos^2 \omega t = \frac{1}{2}(1 + \cos 2\omega t)$, a wave at ω drives a polarization at 2ω, which radiates light at twice the frequency. More generally, $\chi^{(2)}$ lets three light waves exchange energy, as long as energy is conserved photon by photon. Two such processes matter for us (Fig. 2):

- **Second-harmonic generation (SHG):** two photons at 1064 nm combine into one photon at 532 nm ($\omega + \omega \to 2\omega$). This is how the green pump is made inside our laser. We also use SHG in the waveguide for alignment, because green light is easy to see and measure.
- **Spontaneous parametric down-conversion (SPDC):** the reverse process, in which one 532 nm pump photon splits into two 1064 nm photons. Energy conservation requires $\omega_p = \omega_s + \omega_i$; in our "degenerate" case both photons (traditionally called signal and idler) have the same frequency and polarization.



(a) SHG: $\omega + \omega \to 2\omega$ (b) SPDC: $2\omega \to \omega + \omega$

Figure 2: Energy-level pictures of the two $\chi^{(2)}$ processes used in WOPA. The dashed lines are "virtual" levels: the crystal does not absorb the light, it only mediates the exchange of energy between the waves. (a) Second-harmonic generation: two 1064 nm photons combine into one 532 nm photon. (b) Spontaneous parametric down-conversion: a 532 nm pump photon splits into a pair of 1064 nm photons. With a 1064 nm seed present, the same process amplifies or de-amplifies the seed depending on its phase (parametric amplification).

Phase matching and periodic poling: Energy conservation is not enough: the light generated at different points along the crystal must also add up *in phase*. This requires the pump and the generated wave to stay in step, i.e. momentum conservation, $\Delta k = k_p - 2k_s = 0$. Because of dispersion, the refractive index of lithium niobate is larger at 532 nm than at 1064 nm ($n \approx 2.24$ versus 2.16 for light polarized along the crystal's optical axis [12]), so the two waves slowly drift apart. After one *coherence length* $L_c = \pi/\Delta k \approx 3.4\,\mu$m, the newly

generated light is exactly out of phase with what was generated earlier and starts to cancel it. In a uniform crystal the output therefore just oscillates and never builds up (red curve in Fig. 3).

The modern solution is *quasi-phase matching* (QPM). During manufacture, a strong electric field applied through a patterned electrode flips the orientation of the crystal, and with it the sign of $\chi^{(2)}$, every coherence length. This is called *periodic poling*, with a poling period $\Lambda = 2L_c$. Every time the generated light is about to start cancelling, the flip reverses the sign of the nonlinearity, so the contributions keep adding up (green curve in Fig. 3). QPM builds up the field a factor $2/\pi$ more slowly than perfect phase matching would. In exchange, all three waves can be polarized along the crystal's optical axis ("type-0" phase matching), which uses lithium niobate's largest nonlinear coefficient, $d_{33} \approx 25\,\mathrm{pm/V}$, and more than makes up for it. For 1064 nm↔532 nm in bulk lithium niobate, $\Lambda \approx 6.8\,\mu\mathrm{m}$ (our estimate from published dispersion data [12]), so a 1.5 cm crystal contains roughly 4400 poled domains. Because the refractive indices depend on temperature, the phase matching is fine-tuned by holding the crystal at the right temperature, which is why our waveguide sits in an oven (Sec. 3.2).

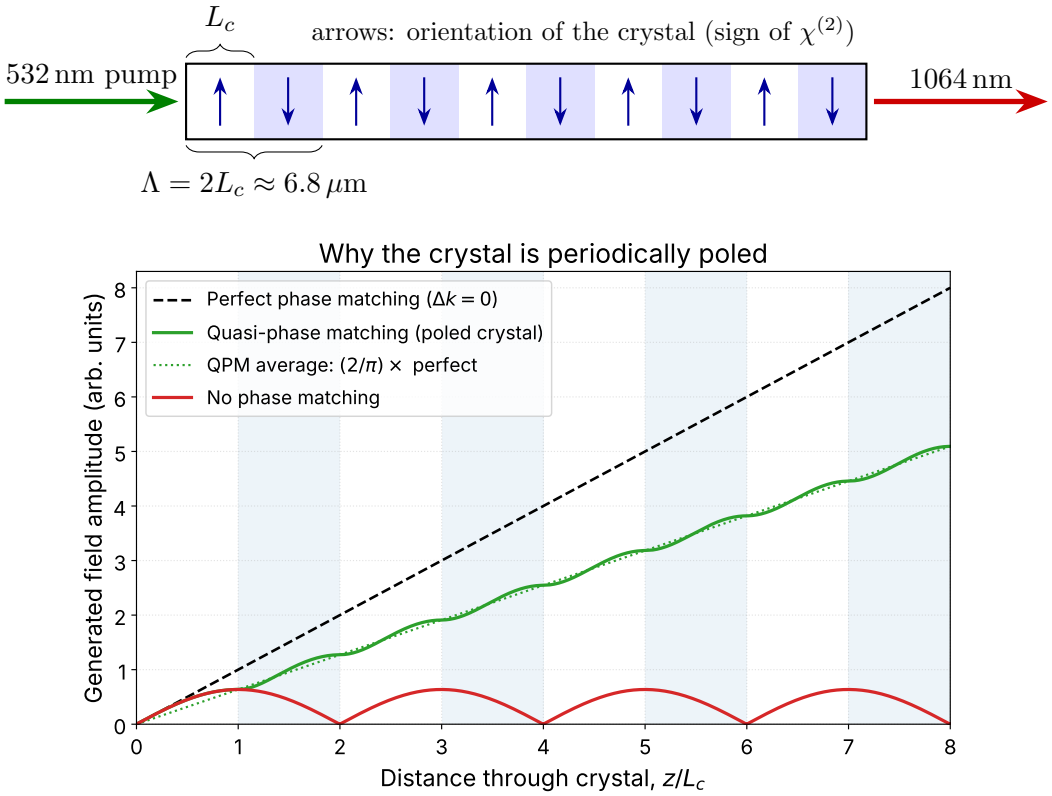


Figure 3: Quasi-phase matching. Top: in a periodically poled crystal the orientation of the crystal, and hence the sign of $\chi^{(2)}$ (arrows), is flipped every coherence length L_c, giving a poling period $\Lambda = 2L_c$ (not to scale). Bottom: amplitude of the generated light versus distance through the crystal (shaded bands: flipped domains). Without phase matching (red) the generated light keeps cancelling itself. With perfect phase matching (dashed) it grows linearly. In a poled crystal (green) it grows steadily at $2/\pi$ of the perfect rate.

Bulk crystal vs Waveguide: The efficiency of these processes grows with the intensity of the light and with the interaction length. In a bulk crystal the two fight each other: focusing tightly gives a high intensity, but a tightly focused beam spreads out again quickly (Appendix A). For example, a beam focused to a $3\,\mu$m radius inside lithium niobate stays that small over only about 0.1 mm. A waveguide removes this trade-off. Its core has a slightly higher refractive index than its surroundings, so the light is guided along the channel (in our case $6 \times 4.16\,\mu$m) and stays tightly confined over the full 1.5 cm, more than 100 times longer (Fig. 4). As a result, the conversion efficiency of a waveguide grows as the *square* of its length, which is why it is quoted in units of $\%/(\mathrm{W\,cm^2})$.

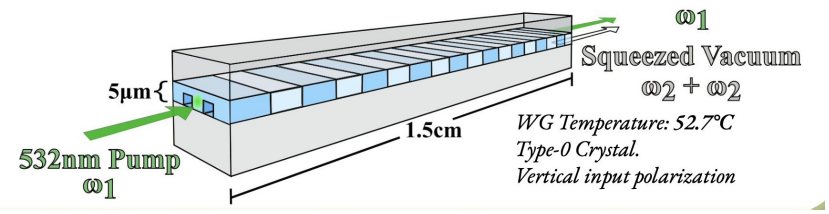


Figure 4: The PPLN waveguide used in WOPA: a single pass of 532 nm pump light through the poled, $6 \times 4.16\,\mu$m channel produces 1064 nm squeezed vacuum. Adapted from Ref. [10].

Parametric amplification and squeezing: Because the photons are always created in pairs, the output is not random in the same way as vacuum: the pairing reduces the fluctuations in one quadrature and increases them in the other. In the language of Sec. 1.3, the pump acts as a *squeezing operator* on the vacuum.

If we also send in a weak 1064 nm "seed" beam, the crystal acts as a *phase-sensitive amplifier*: depending on the phase ψ between pump and seed, the seed is amplified or de-amplified. For a seed of power P_{in} the output power is [22, 23]

$$P_{\mathrm{out}} = P_{\mathrm{in}}\left[\cosh 2r + \sinh 2r\,\cos 2\psi\right], \tag{4}$$

so the power gain swings between e^{+2r} and e^{-2r} as ψ changes. This is the classical fingerprint of the same squeezing parameter r that sets the quantum noise reduction, and it is how the *nonlinear gain* of the waveguide is measured. For a waveguide of length L and normalized nonlinear conversion efficiency α_{NL} [11, 9],

$$r = L\sqrt{\alpha_{\mathrm{NL}}\,\eta_{\mathrm{CE}}\,P_{\mathrm{pump}}}, \tag{5}$$

where η_{CE} is the fraction of pump power actually coupled into the waveguide. Squeezing therefore grows only as the *square root* of the pump power.

2.2 Detecting squeezed light: balanced homodyne detection

Squeezed vacuum has zero average power, so we cannot just shine it on a photodiode. Instead we *interfere* it with a bright reference beam at the same frequency, the *local oscillator* (LO), on a 50:50 beamsplitter, and detect both output ports (Fig. 5). If a is the (weak) signal field

and b the (strong) LO field. The two output fields are $d = (a+b)/\sqrt{2}$ and $c = (a-b)/\sqrt{2}$. Subtracting the two photocurrents gives

$$i_- \propto |d|^2 - |c|^2 = 2\,\mathrm{Re}(a^* b) \;\approx\; 2|b|\,X_\theta, \tag{6}$$

where $X_\theta = X_1 \cos\theta + X_2 \sin\theta$ is the signal's quadrature at the LO phase θ. Three nice things happen at once:

1. The strong LO *amplifies* the tiny quantum fluctuations of the signal (factor $|b|$) so they are far above the detector's electronic noise.
2. Changing the LO phase θ selects *which* quadrature we measure. Scanning θ (we do it with a PZT-mounted mirror) takes us through the squeezed and anti-squeezed quadratures, which is why the noise goes up and down below and above shot noise.
3. Classical intensity noise of the LO appears equally in both photodiodes and cancels in the subtraction. In our setup this common-mode rejection is about 44 dB [9].

With the signal blocked, the subtracted noise is pure shot noise of the LO: this is our 0 dB reference.

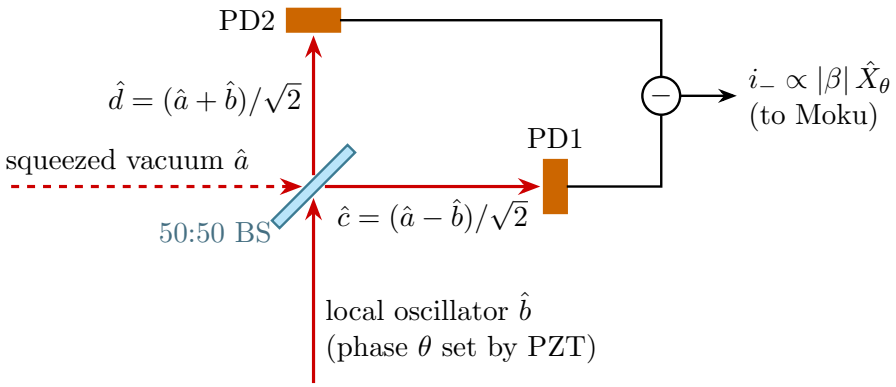


Figure 5: Balanced homodyne detection. The squeezed vacuum and the bright local oscillator (amplitude β) are combined on a 50:50 beamsplitter and each output is detected. The difference photocurrent is proportional to the quadrature $\hat{X}_\theta$ of the squeezed field selected by the LO phase θ, while classical noise of the LO cancels.

2.3 Why losses destroy squeezing

Any loss (absorption, a leaky mirror, a photodiode that does not convert every photon, or a part of the beam that does not overlap the LO) can be modelled as a beamsplitter that transmits a fraction η of the squeezed light and mixes in a fraction $1-\eta$ of ordinary vacuum. The measured noise (relative to shot noise) then becomes [11, 14]

$$R_\pm = (1-\eta) + \eta\, e^{\pm 2r}. \tag{7}$$

Loss drags both squeezing and anti-squeezing towards 0 dB, but it hurts squeezing far more: even with infinite pump power ($e^{-2r} \to 0$) we can never measure better than $10\log_{10}(1-\eta)$. With $\eta = 0.8$ the best possible is -7 dB.

For our experiment the total efficiency is a product of individual efficiencies [11, 9],

$$\eta = \eta_{\mathrm{WG}}\, \eta_P\, \mathcal{V}^2\, \eta_{\mathrm{QE}}, \tag{8}$$

with η_{WG} the waveguide/output-coupling efficiency, η_P the propagation efficiency from the waveguide to the photodiodes, $\mathcal{V}$ the BHD visibility and η_{QE} the photodiode quantum efficiency. The visibility appears *squared* (Sec. 2.4), so a visibility of 90% alone already throws away 19% of the squeezed light.

In a real measurement the LO phase also jitters by a small RMS amount ϕ_{rms}, which mixes a little anti-squeezing into the squeezed quadrature. Including this, the measured noise at LO angle θ is [9]

$$X(\theta) = \eta(1-\phi_{\mathrm{rms}}^2)\left(e^{-2r}\cos^2\theta + e^{2r}\sin^2\theta\right) + \eta\,\phi_{\mathrm{rms}}^2\left(e^{2r}\cos^2\theta + e^{-2r}\sin^2\theta\right) + (1-\eta), \tag{9}$$

and the squeezing in dB is $10\log_{10} X$. We use this model in Sec. 4.7.

2.4 Mode matching and visibility

Laser beams are, to a good approximation, *Gaussian beams*: their intensity has a bell-shaped profile whose radius w changes along the beam as it focuses and spreads. A Gaussian beam is fully described by the size w_0 and position z_0 of its narrowest point, the *waist*, and simple rules describe how lenses change these (Appendix A).

Two beams interfere perfectly only if they have the same spatial shape everywhere. The fraction of one beam's power that lies in the other beam's mode is the (power) *mode overlap* $\mathcal{M}$. *Mode matching* means choosing lenses so that the LO waist has the same size and position as the waist of the squeezed field, which makes $\mathcal{M}$ as close to 1 as possible. A size or position mismatch, an offset or a tilt all reduce it (Appendix A).

We measure the overlap by interfering the two beams and scanning their relative phase ϕ with the PZT. For beams of powers P_1 and P_2 whose normalized field overlap is $\gamma = \sqrt{\mathcal{M}}$, the detected power swings as $P(\phi) = P_1 + P_2 + 2\gamma\sqrt{P_1 P_2}\cos\phi$, so the fringe contrast is

$$C = \frac{P_{\max} - P_{\min}}{P_{\max} + P_{\min}} = \gamma \cdot \underbrace{\frac{2\sqrt{\rho}}{1+\rho}}_{\mathcal{V}_{\max}(\rho)}, \qquad \rho = \frac{P_2}{P_1}. \tag{10}$$

For equal powers ($\rho = 1$), $C = \gamma$ directly. For unequal powers, even perfectly matched beams cannot reach 100% contrast (Fig. 13 in Appendix A). Since in our visibility test the LO and the seeded "tool beam" (Sec. 3.5) had different powers, we always report the *normalized visibility*

$$\mathcal{V} \equiv \gamma = \frac{C_{\mathrm{measured}}}{\mathcal{V}_{\max}(\rho)}. \tag{11}$$

Only the overlapping part of the squeezed field is detected, so the efficiency associated with the visibility is $\gamma^2 = \mathcal{V}^2$, which is why $\mathcal{V}^2$ appears in Eq. (8).

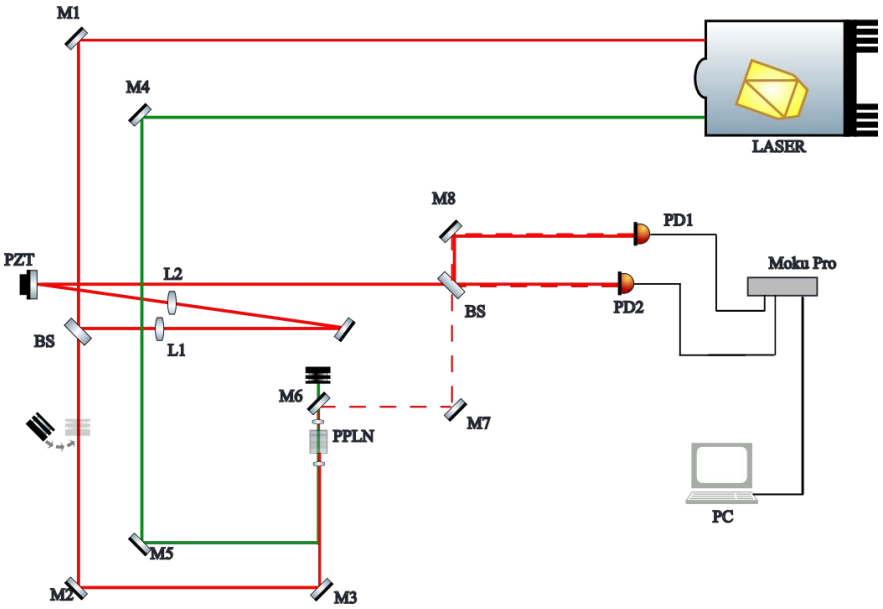


Figure 6: Simplified schematic of the WOPA tabletop (red: 1064 nm, green: 532 nm pump, dashed red: squeezed vacuum). The 1064 nm beam is split at BS_{LO}; one part becomes the LO (via the PZT mirror) and the other part can be sent through the waveguide as the "tool beam". The squeezed vacuum and LO are combined at BS_{BHD} and detected by PD1/PD2, whose outputs are processed digitally in a Moku:Pro. Lens positions are indicative only; the new mode-matching lenses are described in Sec. 4.4.

3 The WOPA experimental setup

3.1 Laser and beam preparation

Both wavelengths come from a single Innolight Diabolo laser: a very stable Nd:YAG non-planar ring oscillator (NPRO) at 1064 nm with an internal frequency-doubling cavity that produces 532 nm. Using one laser for both is important: the pump and the LO are automatically phase-coherent. Each beam passes through half- and quarter-wave plates and polarizing beamsplitters (PBS) to control its polarization and power, and a Faraday isolator protects the laser from back-reflections. The maximum 532 nm power available at the waveguide is about 140 mW.

3.2 The PPLN waveguide

The squeezer is a 1.5 cm long periodically poled lithium niobate (PPLN) waveguide from HC Photonics (Fig. 4) [9, 11]. Its $6 \times 4.16\,\mu$m channel is sandwiched between SiO_2 layers, and both facets are anti-reflection coated for 1064 and 532 nm (reflection <0.5%). The waveguide sits in a brass oven held at 52.7 °C, the temperature at which it is best phase matched (Sec. 2.1). The maximum nonlinear gain is obtained with vertically polarized light [9]. The pump is focused into the waveguide with a microscope objective (input coupling roughly 60–70%), and the light leaving it is collected by an aspheric lens with $f = 6.2$ mm placed a few millimetres from the output facet. Dichroic mirrors then separate the residual green

(which is dumped) from the 1064 nm light that goes to the BHD. The measured normalized conversion efficiency is about 142%/(W cm^2) [9].

3.3 The local oscillator path and the PZT

The LO path is folded so that the beam hits the PZT-mounted mirror at a small angle of incidence ($< 5°$); driving the PZT changes the LO path length and hence the LO phase θ [9]. Because this mirror is not highly reflective, only a fraction of the power split off at $\mathrm{BS_{LO}}$ reaches the BHD: the LO power at the detector was about 220 μW.

3.4 Detection chain

The two photodiodes are New Focus 1811 free-space InGaAs detectors (125 MHz bandwidth) rather than a commercial balanced detector. Their AC outputs are digitized by a Liquid Instruments Moku:Pro, where all further processing happens digitally (Fig. 16 in Appendix C):

1. **Gain balancing and subtraction:** The two PD signals are multiplied by adjustable gains and subtracted so that the LO's classical noise cancels (44 dB common-mode rejection).
2. **Band-pass filtering:** to a window in the MHz range (4–6 MHz) where the subtracted signal is shot-noise limited, away from low-frequency technical noise.
3. **Envelope detection:** the band-passed signal is multiplied by itself and low-pass filtered (∼20–30 kHz), giving a slowly varying voltage that follows the noise level.

The output voltage is recorded while the PZT scans the LO phase. With the pump blocked we record the shot-noise voltage V_{shot}; with the pump on, the voltage swings below and above it. Following the convention used in our group's analysis [9], the squeezing and anti-squeezing are computed as

$$\text{Squeezing (dB)} = 20\log_{10}\left(\frac{V_{\mathrm{min}}}{V_{\mathrm{shot}}}\right), \qquad \text{Anti-squeezing (dB)} = 20\log_{10}\left(\frac{V_{\mathrm{max}}}{V_{\mathrm{shot}}}\right). \tag{12}$$

3.5 The "tool beam": seeing the invisible

To mode match the LO to the squeezed field we need to know the squeezed field's size and waist location, but squeezed vacuum has no average power, so a beam profiler cannot see it. The trick used in this experiment is to send a bright 1064 nm beam through the waveguide instead. The laser's 1064 nm light is split roughly 50:50 at $\mathrm{BS_{LO}}$; the transmitted half is steered into the waveguide, and what comes out is a bright beam travelling along exactly the path the squeezed vacuum will take. We call this the *tool beam*. Because the waveguide channel supports essentially one transverse mode at 1064 nm, we assume the tool beam leaves in the same spatial mode as the squeezed vacuum. We can then profile it, design the LO optics to match it, and use it to measure the visibility. For squeezing measurements the tool beam is blocked right after $\mathrm{BS_{LO}}$. (Whether the tool beam's mode is *exactly* the squeezed mode is a subtle question, discussed in Sec. 5 and Appendix B.)

4 Work and results

4.1 Getting to know the setup and re-commissioning it

For the first two weeks of the program my mentor was away, and I used that time to read background material and to help neighbouring experiments, including profiling some HeNe laser beams (Appendix D). Once Peter returned, he walked me through the WOPA table and the measurement chain. The setup had been idle for a few weeks, so our first task was to check that it still worked. After some minor alignment fixes we observed squeezing, but only about $-1.5\,\mathrm{dB}$ (and $+1.6\,\mathrm{dB}$ anti-squeezing), roughly half of the $\sim -3\,\mathrm{dB}$ previously achieved. This was a useful lesson in how quickly alignment degrades: over a few weeks the beams drift enough to lose a large fraction of the squeezing. We attributed the drop to alignment drift while the setup was idle. The reference for the rest of the project is the setup's previous best performance, obtained by Peter with the earlier mode-matching solution: visibility $\sim$90% and squeezing $\sim -3.0\,\mathrm{dB}$.

4.2 Profiling the tool beam

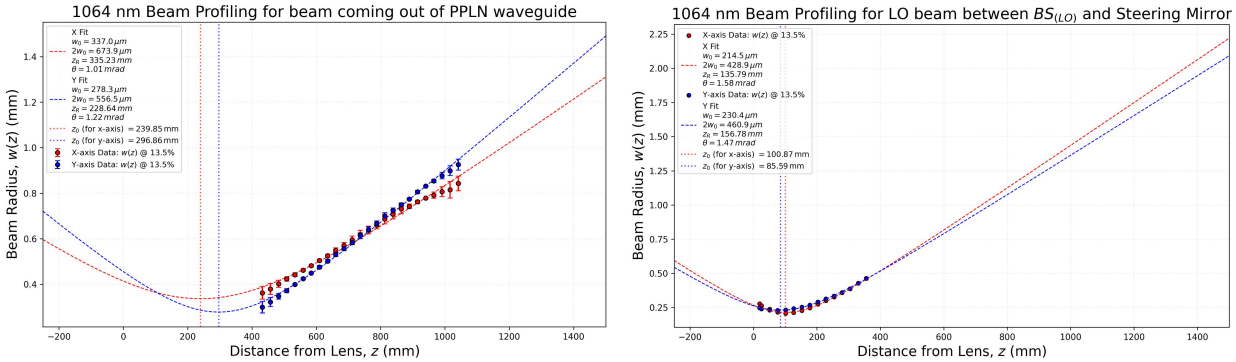
We profiled the tool beam with a scanning-slit beam profiler (DataRay Beam'R2, InGaAs detector) at many positions along the beam, measuring the $1/e^2$ radius in x and y and fitting the Gaussian-beam formula, Eq. (16), to each (Fig. 7(a)).

Right after the waveguide, the beam diverges so quickly that its shape is entirely set by the aspheric lens ($f = 6.2\,\mathrm{mm}$). This lens is also our distance reference ($z = 0$) for this measurement. Because the waveguide mode is only a few micrometres across while the beam we want is hundreds of micrometres across, the lens acts as a strong magnifier, and the *longitudinal* magnification is roughly the square of the transverse one (of order 10^3–10^4). In practice this means that moving the lens by just a few micrometres moves the output waist by centimetres. We used a translation stage to adjust the lens along the beam axis until the x and y waists sat as close together as possible and the ellipticity (w_y/w_x) stayed close to 1 over the region of the BHD. Moving the lens sideways (x, y) had little effect on the beam shape, while the axial (z) position changed it strongly, as expected.

The best-fit parameters are given in Table 1. The beam is noticeably *elliptical and astigmatic*: the x waist ($337\,\mu\mathrm{m}$) is about 20% larger than the y waist ($278\,\mu\mathrm{m}$), and the two waists are $\sim$57 mm apart along the beam. This is expected from the rectangular cross-section of the waveguide channel [9], and it is the fundamental reason why matching it to a nearly round LO is hard.

4.3 Profiling the local oscillator

We profiled the LO in the same way, between $\mathrm{BS_{LO}}$ (our $z = 0$ for everything on the LO side) and the next steering mirror (Fig. 7(b)). The LO is also slightly elliptical, as most laser beams are, but much less so than the tool beam: its x and y waists differ by about 7% and are only 15 mm apart.



(a) Tool beam leaving the waveguide.

(b) Local oscillator.

Figure 7: Measured beam profiles (red: x, blue: y; dashed lines: Gaussian fits). (a) The 1064 nm tool beam leaving the waveguide, with z measured from the aspheric lens. (b) The LO between $\mathrm{BS_{LO}}$ and the steering mirror, with z measured from $\mathrm{BS_{LO}}$; its waist of about 220 μm lies roughly 10 cm after $\mathrm{BS_{LO}}$. Fit parameters are listed in Table 1.

Table 1: Gaussian fit parameters of the two beams. z_0 for the tool beam is measured from the aspheric lens; for the LO from $\mathrm{BS_{LO}}$.

	Tool beam (waveguide)		Local oscillator	
	x	y	x	y
Waist w_0 (μm)	337.0	278.3	214.5	230.4
Waist position z_0 (mm)	239.9	296.9	100.9	85.6
Rayleigh range z_R (mm)	335.2	228.6	135.8	156.8
Divergence θ (mrad)	1.01	1.22	1.58	1.47

4.4 Designing the new mode-matching solution

What we were matching to: Since spherical lenses cannot turn a round beam into an elliptical one, the practical target is a round beam that best matches the tool beam on average: a waist of about 305 μm located halfway between the x and y waist positions of the tool beam (268 mm after the asphere). Mapped into the LO's coordinate system (distances from $\mathrm{BS_{LO}}$ along the LO path, with $\mathrm{BS_{BHD}}$ at about 1.53 m), this target waist sits at $z_t = 1.263$ m.

Constraints: In the earlier solution the mode-matching lenses were placed *after* the PZT mirror. When the PZT is driven, any small tilt or wobble of the mirror slightly moves the LO beam relative to the squeezed beam, and this can reduce the overlap during the scan. We therefore restricted the lenses to the stretch between $\mathrm{BS_{LO}}$ and the PZT mirror, avoiding regions already occupied by other optics (the shaded "blocks" in Fig. 8).

Solution: Using an online mode-matching calculator [19] we searched for two-lens solutions within these constraints and chose

$$f_1 = 250\,\text{mm at } z = 134\,\text{mm}, \qquad f_2 = 300\,\text{mm at } z = 646\,\text{mm (from BS}_{\text{LO}}\text{)},$$

which produces a waist of 305.2 μm at $z = 1.263$ m (Fig. 8). The first lens gently refocuses the LO; the second focuses it into a slow, large beam whose waist lands where the squeezed beam's waist is.

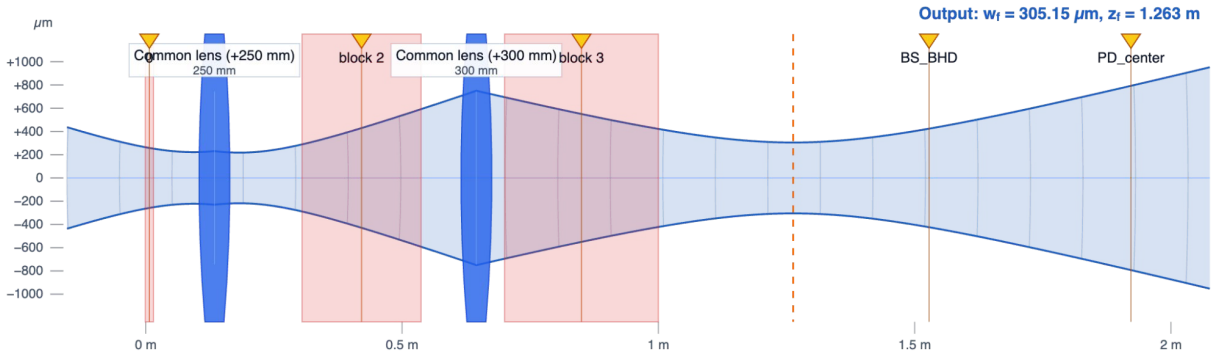


Figure 8: The chosen mode-matching solution (screenshot from the calculator [19]). Blue: lenses; pink bands: forbidden regions; markers show BS_{BHD} and the PD positions. Distances are measured from BS_{LO}.

Checking the solution: To check the calculator independently, I propagated the measured q parameters of the LO (Table 1) through the two lenses with ray-transfer (ABCD) matrices, separately for x and y, and computed the overlap with the measured tool beam at BS_{BHD} (details in Appendix A.4). The results are:

- compared with the round design target (305 μm at 1.263 m): $\mathcal{M} = 0.998$, i.e. the lenses do exactly what they were designed to do;
- compared with the *real*, astigmatic tool beam: $\mathcal{M} = 0.978$, i.e. $\mathcal{V} = \sqrt{\mathcal{M}} = 0.989$;
- for comparison, even the best possible round beam only reaches $\mathcal{M} = 0.988$ ($\mathcal{V} = 0.994$) against the measured tool beam.

So, on paper, the new lenses allow a visibility of about 99%, within half a percent of the limit set by the ellipticity of the waveguide beam. Recovering that last half percent would require cylindrical lenses [9].

4.5 Measuring the visibility

With the lenses in place, we aligned the LO onto the tool beam at BS_{BHD} using the steering mirrors, then sent both onto a photodiode and scanned the PZT. The powers of the two beams were $P_{\text{LO}} = 60\,\mu$W and $P_{\text{tool}} = 22\,\mu$W, i.e. $\rho = 22/60 = 0.367$. Even for perfect mode matching, this power imbalance limits the fringe contrast to

$$\mathcal{V}_{\max}(0.367) = \frac{2\sqrt{0.367}}{1 + 0.367} = 0.886. \tag{13}$$

The measured raw contrast was $C \approx 0.81$, so with Eq. (11)

$$\boxed{\mathcal{V} = \frac{C}{\mathcal{V}_{\max}} \approx \frac{0.808}{0.886} = 91.2\%} \qquad \text{(previous solution: } \sim 90\%\text{)}. \tag{14}$$

This is a real but modest improvement. In terms of detection efficiency, $\mathcal{V}^2$ went from 0.81 to 0.83.

4.6 Squeezing measurement

We then blocked the tool beam, unblocked the pump and measured squeezing. Figure 9 shows a trace at 142 mW of pump power with 220 μW of LO. The shot-noise level (pump blocked) is $V_{\text{shot}} = 0.120$ V (also shown in Appendix C). As the PZT scans, the noise dips to $V_{\min} \approx 0.083$ V and rises to $V_{\max} \approx 0.175$ V. With Eq. (12),

$$\text{Squeezing} = 20 \log_{10} \frac{0.083}{0.120} = -3.20\,\text{dB}, \qquad \text{Anti-squeezing} = 20 \log_{10} \frac{0.175}{0.120} = +3.28\,\text{dB}. \tag{15}$$

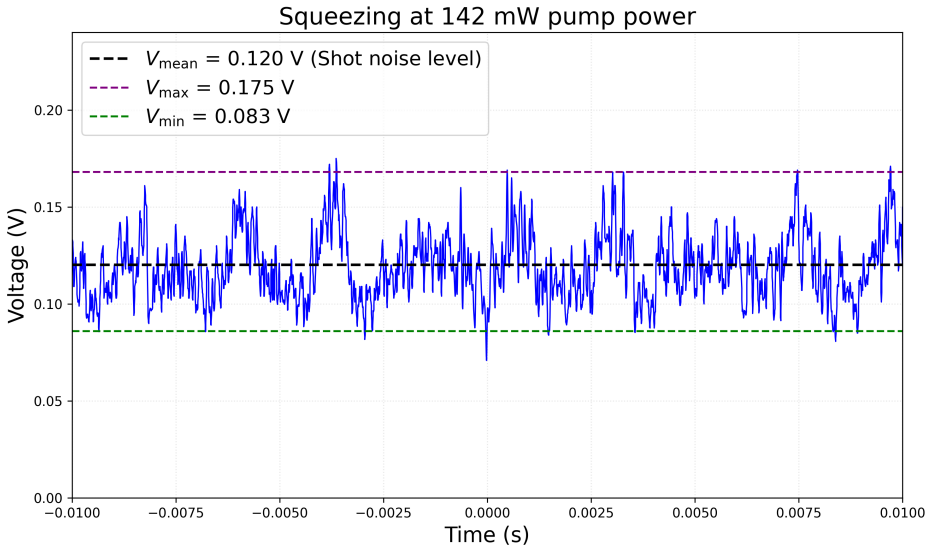


Figure 9: Envelope-detector output while scanning the LO phase, at 142 mW pump and 220 μW LO. Black: shot-noise level. The dips below it are squeezing, the peaks above it anti-squeezing. Dashed lines mark $V_{\max}$ and $V_{\min}$.

We also swept the pump power and compared the new data with the earlier data taken with the old mode matching (Fig. 10). The new data lie slightly but consistently below the old ones, reaching −3.20 dB at 142 mW where the old data reached about −2.8 dB at 118 mW. Both follow the expected square-root trend: doubling the pump does not double the squeezing in dB. The error bars are large because each point is read off a noisy, scanned trace; the improvement at any single power is within the error bars, but the trend over the whole sweep is consistent.

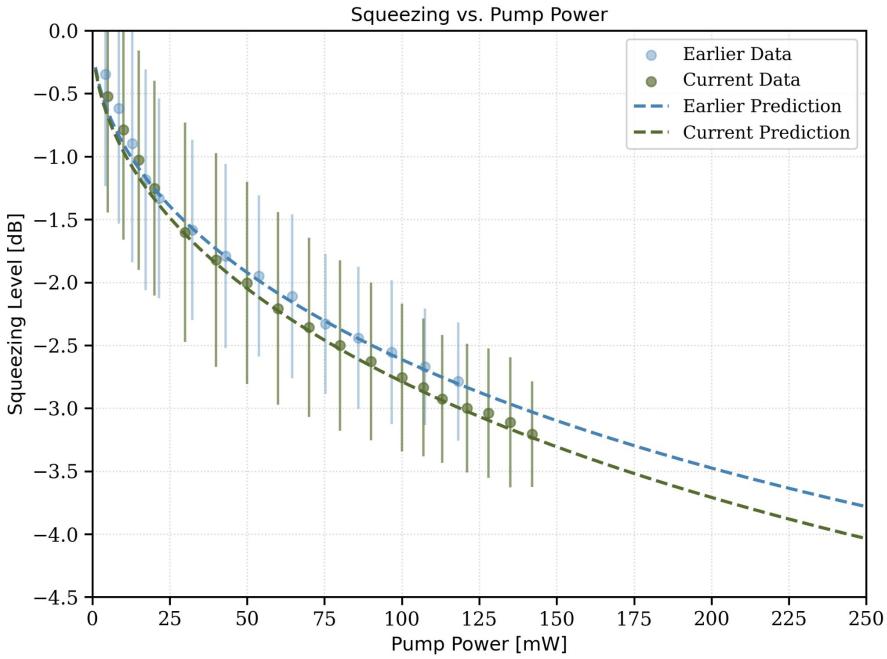


Figure 10: Squeezing versus pump power with the earlier mode matching (blue) and the new one (green), together with model predictions for each.

4.7 How much squeezing should better visibility give?

To turn the question "is it worth improving the visibility?" into numbers, I used the full model of Eq. (9) with the parameters of Table 2. Most parameters are measured values from Ref. [9]. Note from Eq. (5) that the squeezing depends only on the *product* $\alpha_{\rm NL}\eta_{\rm CE}$, so these two cannot be separated using squeezing data alone. I fixed the input coupling at its expected value, $\eta_{\rm CE} = 0.65$, and adjusted $\alpha_{\rm NL}$ so that the model reproduces our measured $-3.20\,$dB at $142\,$mW and $\mathcal{V} = 0.912$. This gives $\alpha_{\rm NL} = 1.57\,/(\mathrm{W\,cm^2})$, at the edge of the measured $1.42 \pm 0.15\,/(\mathrm{W\,cm^2})$ [9], and $r \approx 0.57$ at $142\,$mW. (Without any losses at all, $r = 0.57$ would correspond to $-5.0\,$dB of squeezing.)

Table 2: Parameters used in the squeezing model.

Parameter	Value	Source
Waveguide length L	$1.5\,$cm	[9]
Input coupling $\eta_{\rm CE}$	0.65	expected value
Nonlinear efficiency $\alpha_{\rm NL}$	$1.57\,/(\mathrm{W\,cm^2})$	fit (measured 1.42 ± 0.15 [9])
Propagation efficiency η_P	0.95	[9]
Photodiode quantum efficiency $\eta_{\rm QE}$	0.97	[9]
RMS phase noise $\phi_{\rm rms}$	$15.9\,$mrad	[9]

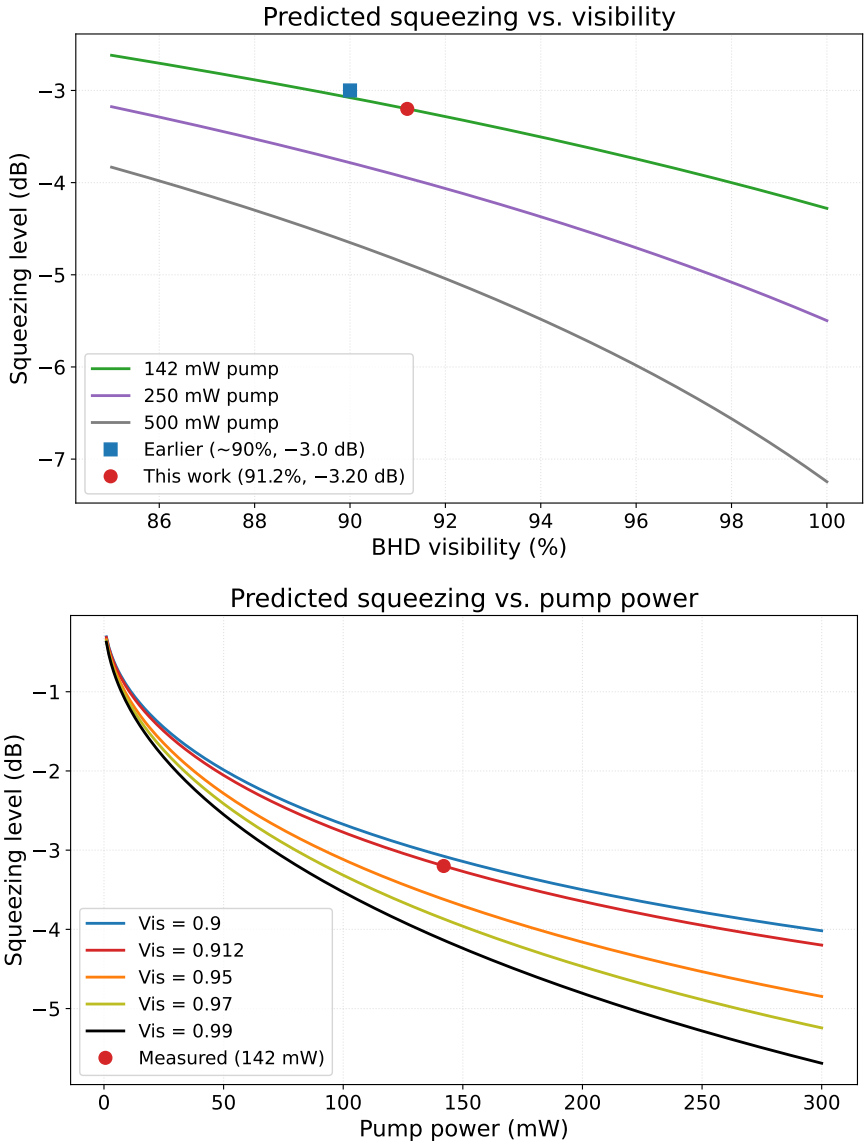


Figure 11: Model predictions. Top: detected squeezing versus BHD visibility at three pump powers, with the earlier and new operating points. Bottom: squeezing versus pump power for several visibilities; the red dot is our measurement.

Table 3: Predicted detected squeezing (dB) for different visibilities.

Pump power	$\mathcal{V} = 0.90$	0.912	0.95	0.97	0.99	1.00
142 mW	−3.08	**−3.20**	−3.62	−3.87	−4.14	−4.28
250 mW	−3.78	−3.95	−4.54	−4.89	−5.28	−5.50

Two conclusions stand out (Fig. 11, Table 3). First, the model says our visibility improvement should be worth about 0.12 dB at 142 mW, matching the observed change from −3.0 to −3.20 dB. Second, going to 95–97% visibility would give another 0.4–0.7 dB at the same pump power, and the benefit grows with pump power (up to ≈ 1 dB at 250 mW). The more squeezing we generate, the more it matters that we detect it efficiently.

5 Discussion: what limits the visibility now?

The most important insight from the q-parameter analysis is the gap between what the lenses allow ($\mathcal{V} \approx 0.99$) and what we measure ($\mathcal{V} = 0.912$). In power terms, about $1 - 0.912^2 \approx 17\%$ of the squeezed light is not overlapping the LO, and the waist size/position mismatch accounts for only about 2% of it. So what is responsible for the other ∼15%?

Alignment: Eq. (22) shows how unforgiving the overlap is. For our 305 μm waist, 1% loss already results from a lateral offset of only 30 μm or a tilt of only 0.11 mrad (i.e. 0.11 mm over a metre). A 15% loss corresponds to an offset of about 120 μm or a tilt of about 0.45 mrad. These are easily within what is left over after aligning "by eye" on a beam profiler, so alignment is very likely the dominant remaining loss. A systematic alignment procedure (walking two mirrors while directly maximizing the fringe contrast) should recover part of this.

The waveguide beam is not a perfect Gaussian: Our fits assumed ideal Gaussian beams ($M^2 = 1$). The mode of a rectangular waveguide is not exactly Gaussian, and some light may travel outside the guided mode. Any such non-Gaussian part of the tool beam can never overlap a Gaussian LO, no matter how good the lenses are. Fitting the profiles with a free M^2 parameter would tell us how large this effect is.

Is the tool beam really the squeezed mode?: The squeezed light is generated only where the pump light is, so its spatial mode is set by the overlap of the pump with the waveguide mode, whereas the tool beam is simply whatever 1064 nm light comes out of the guide. If the two differ, we are matching the LO to the wrong target, and the *true* squeezing visibility could be lower or higher than 91.2%. Appendix B works out a way to measure the true mode directly.

The PZT mirror: A pure translation of the PZT by about a wavelength shifts the beam sideways by only $2\delta \sin\theta \lesssim 0.2\,\mu$m at our <5° incidence, which is negligible. Any effect of

the PZT therefore has to come from tilt or wobble of the piezo stack. Replacing the PZT with an electro-optic modulator (EOM), which changes the phase without moving anything, would remove this question altogether.

A consistency check with anti-squeezing: The pure-loss model of Eq. (7) predicts about +4.2 dB of anti-squeezing at our operating point, whereas we read about +3.3 dB from the trace. Solving Eq. (7) for the measured pair (−3.20, +3.3 dB) would require a total efficiency $\eta \approx 0.97$, which is impossible since $\eta \leq \mathcal{V}^2 = 0.83$. The most likely explanation is that the anti-squeezing peak is underestimated when read from a noisy, swept trace. Measuring anti-squeezing with the quantum noise lock [9, 18] holding the LO on the anti-squeezed quadrature would settle this. It is worth doing, because the ratio of squeezing to anti-squeezing is one of the cleanest ways to measure the total loss η independently of any beam profiling.

6 Hunting for scattered light

A long-term goal of WOPA is to demonstrate squeezing in LIGO's audio band (tens of Hz to kHz), where the current photodiodes are not yet good enough and where *scattered light* becomes a major noise source. The mechanism is simple: a tiny fraction of a beam scatters off a surface that vibrates (a mount, the table, a beam dump), picks up a phase modulation proportional to that surface's motion, and some of it finds its way back into the main beam path. Because the scattering surfaces move by many wavelengths at low frequencies, this can produce noise well above the motion frequency itself, and it cannot be removed by the BHD subtraction because it looks like a real signal.

We searched the table for places where light was being scattered. The most prominent scattering was found on the *surface below and neighbouring the PBS in the power-adjustment stage*. As first mitigations, beam dumps and baffles were placed near these spots. For future low-frequency measurements, absorbing glass (e.g. KG5) beam dumps and a more systematic survey, ideally while monitoring the BHD noise spectrum as each dump is added, are planned [10].

7 Conclusions and future work

In this project I profiled the beams of the WOPA tabletop, designed and implemented a new mode-matching solution for the local oscillator that avoids the PZT mirror, and improved the BHD visibility from ∼90% to 91.2%. As a result the detected squeezing went from −3.0 dB to −3.20 dB at ≈142 mW of pump, in good agreement with a loss model calibrated on our setup. That model predicts −3.6 to −3.9 dB at 95–97% visibility at the same pump power, and −4.5 to −4.9 dB at 250 mW. A q-parameter analysis showed that the lenses themselves are good for $\mathcal{V} \approx 0.99$, so the remaining losses come from alignment and from the non-ideal, astigmatic waveguide output, not from the lens design.

Natural next steps are:

- a careful re-alignment of the BHD, maximizing the fringe contrast directly, and fitting beam profiles with a free M^2;
- cylindrical lenses to correct the waveguide's ellipticity and astigmatism;
- measuring the true spatial mode of the parametric gain with the modulation technique of Appendix B, and mode matching the LO to it;
- replacing the PZT phase scan with an EOM for cleaner phase control;
- expanding both beams to $\sim$1 mm waists in the BHD to make the overlap less sensitive to lateral misalignment (Appendix B);
- continuing the scattered-light hunt in preparation for audio-band squeezing with low-noise DC-coupled photodiodes.

Acknowledgements

I am very grateful to my mentor Peter Carney for patiently explaining every part of the WOPA experiment and being there for me whenever I needed any help in the lab, and to Prof. Rana Adhikari for his guidance, his questions, and for pushing me to estimate things on paper before touching the table. I thank the other members of the lab for their help, Mayank Chaturvedi for technical as well as non-technical support, the LIGO SURF 2026 cohort, and the LIGO Laboratory and Caltech Student-Faculty Programs teams for organizing this program. This work was supported by the National Science Foundation.

References

[1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), "Observation of gravitational waves from a binary black hole merger," Phys. Rev. Lett. **116**, 061102 (2016).

[2] C. M. Caves, "Quantum-mechanical noise in an interferometer," Phys. Rev. D **23**, 1693 (1981).

[3] M. C. Teich and B. E. A. Saleh, "Squeezed states of light," Quantum Opt. **1**, 153 (1989).

[4] J. Aasi *et al.* (LIGO Scientific Collaboration), "Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light," Nature Photon. **7**, 613 (2013).

[5] S. E. Dwyer, G. L. Mansell and L. McCuller, "Squeezing in gravitational wave detectors," Galaxies **10**, 46 (2022).

[6] M. Tse *et al.*, "Quantum-enhanced Advanced LIGO detectors in the era of gravitational-wave astronomy," Phys. Rev. Lett. **123**, 231107 (2019).

[7] D. Ganapathy *et al.*, "Broadband quantum enhancement of the LIGO detectors with frequency-dependent squeezing," Phys. Rev. X **13**, 041021 (2023).

[8] T. Kashiwazaki *et al.*, "Over-8-dB squeezed light generation by a broadband waveguide optical parametric amplifier toward fault-tolerant ultra-fast quantum computers," Appl. Phys. Lett. **122**, 234003 (2023).

[9] P. Carney, S. J. Maliakal, A. Wade and R. X. Adhikari, "Waveguide optical parametric amplification for broadband quantum squeezing," in preparation (2026).

[10] P. G. Carney *et al.*, "Broadband waveguide quantum squeezing," poster, GWADW (May 2026).

[11] HC Photonics Corp., "Key factors/critical requirements and PPLN for squeezing applications," application note (Jan. 2025).

[12] D. H. Jundt, "Temperature-dependent Sellmeier equation for the index of refraction, n_e, in congruent lithium niobate," Opt. Lett. **22**, 1553 (1997).

[13] R. W. Boyd, *Nonlinear Optics*, 4th ed. (Academic Press, 2020).

[14] H.-A. Bachor and T. C. Ralph, *A Guide to Experiments in Quantum Optics*, 2nd ed. (Wiley-VCH, 2004).

[15] A. E. Siegman, *Lasers* (University Science Books, 1986).

[16] H. Kogelnik and T. Li, "Laser beams and resonators," Appl. Opt. **5**, 1550 (1966).

[17] D. Z. Anderson, "Alignment of resonant optical cavities," Appl. Opt. **23**, 2944 (1984).

[18] K. McKenzie *et al.*, "Quantum noise locking," J. Opt. B: Quantum Semiclass. Opt. **7**, S421 (2005).

[19] I. MacMillan, Mode-Matching calculator, https://ian-macmillan.com/Mode-Matching/.

[20] D. Ganapathy, "Squeezing quantum noise with waveguides," LIGO SURF report, LIGO-T1700336 (2017).

[21] J. Cherry Sam, "Quantum noise locking on squeezed states," LIGO SURF report, LIGO-T2400203 (2024).

[22] N. Dreslin, "PPLN crystal nonlinear gain measurements for tabletop waveguided optical parametric amplification," LIGO SURF report, LIGO-T2500241 (2025).

[23] T. Ray, "Noise characterizations in tabletop waveguided optical parametric amplification experiment," LIGO SURF report, LIGO-T2500238 (2025).

A Gaussian beams and mode-matching basics

A.1 Gaussian beams and the complex beam parameter q

Laser beams are (to a good approximation) *Gaussian beams*: the intensity across the beam has a bell-curve profile, $I(x,y) \propto \exp[-2(x^2+y^2)/w^2]$, where w is the beam radius (at which the intensity falls to $1/e^2 = 13.5\%$ of its peak). As the beam travels along z its radius changes as [15, 16]

$$w(z) = w_0\sqrt{1+\left(\frac{z-z_0}{z_R}\right)^2}, \qquad z_R = \frac{\pi w_0^2}{\lambda}, \tag{16}$$

where w_0 is the smallest radius (the *waist*), z_0 its location and z_R the *Rayleigh range*: the distance over which the beam stays roughly collimated (Fig. 12). Far from the waist the beam spreads with half-angle $\theta = \lambda/(\pi w_0)$. A small waist diverges quickly, a large waist stays collimated for a long way. This is why the light leaving a $\sim 5\,\mu$m waveguide spreads out almost immediately and needs a short-focal-length lens right after it, and why a tightly focused beam in a bulk crystal only stays small over a very short distance (Sec. 2.1).

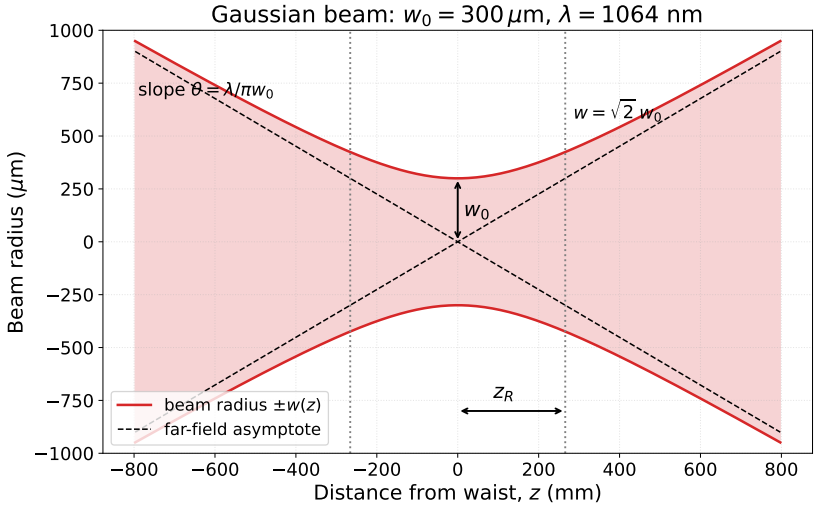


Figure 12: A Gaussian beam with $w_0 = 300\,\mu$m at $\lambda = 1064\,$nm ($z_R \approx 0.27\,$m). The beam radius grows by $\sqrt{2}$ one Rayleigh range away from the waist, and approaches the dashed asymptotes far away.

The q parameter: All the information about a Gaussian beam at a plane z (its size and wavefront curvature R) can be packed into one complex number,

$$q(z) = (z - z_0) + i z_R, \qquad \frac{1}{q} = \frac{1}{R(z)} - i\frac{\lambda}{\pi w(z)^2}. \tag{17}$$

The real part tells us how far we are from the waist, the imaginary part tells us the waist size. The reason this is so useful is that propagation through *any* paraxial optical system described by a ray-transfer ("ABCD") matrix takes the simple form [15, 16]

$$q_{\text{out}} = \frac{A\,q_{\text{in}} + B}{C\,q_{\text{in}} + D}, \qquad M_{\text{space}}(d) = \begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix}, \quad M_{\text{lens}}(f) = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}. \tag{18}$$

So free space just adds the distance, $q \to q + d$, and a thin lens changes the curvature, $1/q \to 1/q - 1/f$. Designing a mode-matching telescope is then just a matter of chaining these two rules.

A.2 Overlap of two Gaussian beams

For Gaussian beams the overlap integral of can be done analytically.

$$\mathcal{M} = \frac{\left|\iint u_1^*(x,y)\,u_2(x,y)\,dx\,dy\right|^2}{\iint |u_1|^2\,dx\,dy \iint |u_2|^2\,dx\,dy}, \tag{19}$$

where u_1 and u_2 are the two beams' transverse field profiles. If the beams are perfectly aligned (same axis, no tilt), the answer depends only on their q parameters at *any* common plane [15]. For each transverse direction separately (useful for elliptical beams, which have different q_x and q_y),

$$\mathcal{M}_x = \frac{2\sqrt{\operatorname{Im} q_{1x}\,\operatorname{Im} q_{2x}}}{|q_{1x}^{*} - q_{2x}|}, \qquad \mathcal{M} = \mathcal{M}_x\,\mathcal{M}_y. \tag{20}$$

For two round beams with waists w_1, w_2 separated by Δz this becomes the familiar

$$\mathcal{M} = \frac{4}{\left(\frac{w_1}{w_2} + \frac{w_2}{w_1}\right)^2 + \left(\frac{\lambda\,\Delta z}{\pi w_1 w_2}\right)^2}. \tag{21}$$

Misalignment reduces the overlap further. For a lateral offset δx or a tilt $\delta\theta$ between two otherwise identical beams of waist w_0 [17],

$$\mathcal{M} \approx \exp\left[-\left(\frac{\delta x}{w_0}\right)^2\right], \qquad \mathcal{M} \approx \exp\left[-\left(\frac{\pi w_0\,\delta\theta}{\lambda}\right)^2\right]. \tag{22}$$

A.3 Maximum contrast for unequal powers

Figure 13 shows the factor $\mathcal{V}_{\max}(\rho) = 2\sqrt{\rho}/(1+\rho)$ of Eq. (10). With a 1:10 power ratio the best possible contrast is only 0.57; for the powers in our visibility test ($\rho = 22/60$) it is 0.886.

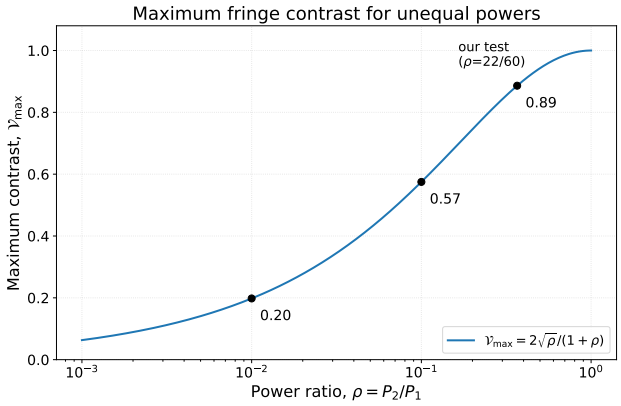


Figure 13: Maximum possible fringe contrast, $\mathcal{V}_{\max} = 2\sqrt{\rho}/(1+\rho)$, for two perfectly mode-matched beams with power ratio ρ. Measured contrasts must be divided by this factor. The point at $\rho = 22/60$ corresponds to our visibility measurement (Sec. 4.5).

A.4 Checking our lens solution with q parameters

Starting from the measured LO beam at $\mathrm{BS_{LO}}$, $q(0) = -z_0 + iz_R$ with the values of Table 1, I applied Eq. (18) for each element in turn: free space to the $f_1 = 250\,\mathrm{mm}$ lens (134 mm), the lens, free space to the $f_2 = 300\,\mathrm{mm}$ lens (646 mm), the lens, and free space to $\mathrm{BS_{BHD}}$

(about 1.53 m), separately for x and y. The measured q parameters of the tool beam were mapped into the same coordinates, and the overlap was evaluated at $\mathrm{BS_{BHD}}$ with Eq. (20). Table 4 summarizes the results, and Fig. 14 shows the beam sizes along the whole path.

Table 4: Result of propagating the measured LO through the new lenses, compared with the tool beam at $\mathrm{BS_{BHD}}$.

	x	y
LO waist after the lenses (μm)	295	313
Position of that waist from $\mathrm{BS_{LO}}$ (m)	1.281	1.245
LO radius at $\mathrm{BS_{BHD}}$ (μm)	411	439
Tool-beam radius at $\mathrm{BS_{BHD}}$ (μm)	449	402
Overlap $\mathcal{M}_x$, $\mathcal{M}_y$	0.988	0.989

The total overlap is $\mathcal{M} = \mathcal{M}_x\mathcal{M}_y = 0.978$. Interestingly, the LO's small ellipticity is "crossed" with respect to the tool beam: the LO is slightly *smaller* than the tool beam in x and slightly *larger* in y.

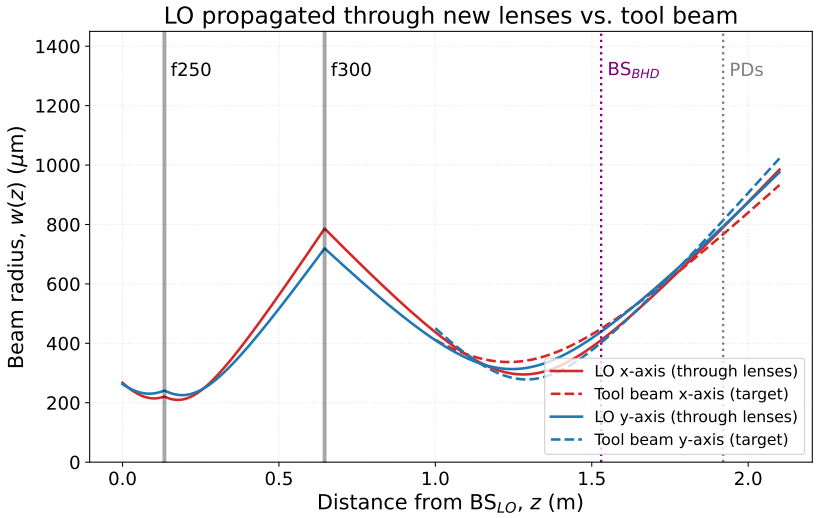


Figure 14: Beam radius of the LO propagated through the new lenses using the measured q parameters (solid), compared with the measured tool beam mapped into the same coordinates (dashed). Grey bars: the $f = 250$ mm and $f = 300$ mm lenses. Red: x; blue: y.

B Measuring the true mode of the squeezed field: an estimate

The tool beam tells us where 1064 nm light leaving the waveguide goes, but the squeezing lives in the mode of the *nonlinear gain*, which is set by where the pump overlaps the waveguide mode. Prof. Adhikari suggested measuring this mode directly:

1. send the 532 nm pump and a strong 1064 nm seed through the waveguide together, and chop the pump at $f_m \approx 5\,\mathrm{kHz}$;
2. since the pump power sets the parametric gain (Eq. (4)), the 1064 nm output becomes modulated at f_m;
3. after filtering out the 532 nm, the steady part of the 1064 nm output is just the tool beam, but the part modulated at f_m exists only where the seed overlapped the pump: it carries the *spatial mode of the gain*, i.e. the mode of the squeezed field;
4. interfering this with the LO in the BHD and demodulating at f_m gives a fringe whose size is proportional to the overlap of the LO with the gain mode. Maximizing that demodulated fringe mode matches the LO to the mode that actually carries the squeezing.

We did not have time to try this in the lab. Following Prof. Adhikari's advice to first "work through the math and estimate the voltage SNR at the PDs," this appendix makes that estimate.

B.1 How much 1064 nm modulation should we expect?

With the pump on, the seed's output power is $P_{\mathrm{in}}\,[\cosh 2r + \sinh 2r \cos 2\psi]$ (Eq. (4)); with the pump off it is simply P_{in}. Using r from Eq. (5) with the parameters of Table 2, Table 5 and Fig. 15 give the fractional change of the 1064 nm power between pump-on and pump-off.

Table 5: Expected change of the 1064 nm power when the pump is switched on, for the best ($\psi = 0$) and worst ($\psi = \pi/2$) pump–seed phase and averaged over ψ. The last two columns give the expected demodulated fringe in the BHD and its SNR in a 1 kHz bandwidth (see text).

Chopped pump	r	max	min	average	BHD fringe (10^4 V/A)	SNR (1 kHz)
10 mW	0.15	+35%	−26%	+4.6%	0.10 V	4×10^4
50 mW	0.34	+97%	−49%	+24%	0.26 V	1×10^5
142 mW	0.57	+213%	−68%	+73%	0.49 V	2×10^5

So the answer to "how much modulation?" is: *a lot.* Even 10 mW of chopped pump modulates the 1064 nm power by tens of percent. The sign and size depend on the pump–seed phase ψ, which drifts slowly because the two beams travel several metres through air along different paths; averaged over ψ, however, the change is always positive ($\cosh 2r > 1$), so the signal never averages away completely.

B.2 Signal-to-noise ratio at the photodiodes

In the BHD, the demodulated difference current at f_m is approximately

$$i_{-}^{(f_m)} \approx \frac{2}{\pi} \cdot 2\mathcal{R}\sqrt{P_{\mathrm{LO}} P_{\mathrm{in}}}\,|\delta A|\,\gamma_{\mathrm{AC}} \cos(\theta + \theta'), \qquad |\delta A| = \left|\cosh r - 1 + e^{2i\psi} \sinh r\right|, \tag{23}$$

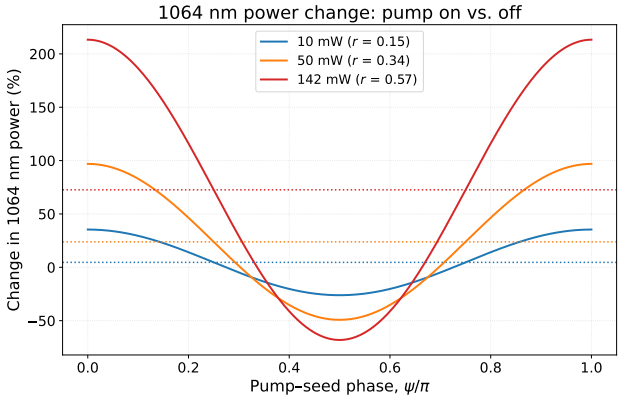


Figure 15: Change of the 1064 nm output power with the pump on vs. off, as a function of the pump–seed phase ψ, for three pump powers. Dotted lines: average over ψ ($\cosh 2r - 1$).

Here $2/\pi$ is the fundamental Fourier component of the square-wave chopping, $\mathcal{R} \approx 0.83\,\mathrm{A/W}$ is the photodiode responsivity, $|\delta A|$ is the fractional change of the seed *field* amplitude, and γ_{AC} is the overlap of the LO with the gain mode. The noise floor is the shot noise $\sqrt{2e\mathcal{R}(P_{\mathrm{LO}} + P_{\mathrm{in}})}$. For $P_{\mathrm{LO}} = 220\,\mu\mathrm{W}$, a seed of $20\,\mu\mathrm{W}$ at the BHD, $\gamma_{\mathrm{AC}} = 0.9$ and $\psi = 0$, the shot noise is $8 \times 10^{-12}\,\mathrm{A}/\sqrt{\mathrm{Hz}}$ ($80\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$ at a transimpedance of $10^4\,\mathrm{V/A}$), while the fringe is 0.1–0.5 V (Table 5). Even in the $\sim$1 kHz bandwidth needed to follow a PZT scan, the SNR is $\sim 10^5$.

So SNR is not the problem; a few milliwatts of chopped pump would already be enough. The real challenges are practical:

- **Detector bandwidth:** the 1811 RF output is AC-coupled and its response falls off below $\sim$10 kHz [10], so either use DC-coupled photodiodes (e.g. the lab's PDA10CS) or modulate faster with an AOM/EOM.
- **Pump–seed phase drift:** average over ψ, scan it, or lock it.
- **Green leakage:** chopped 532 nm reaching the PDs gives a signal at exactly f_m; check by blocking the seed.
- **The chopper:** it vibrates and scatters light, so it should be mechanically isolated and placed near a focus.

B.3 Where should the beams overlap?

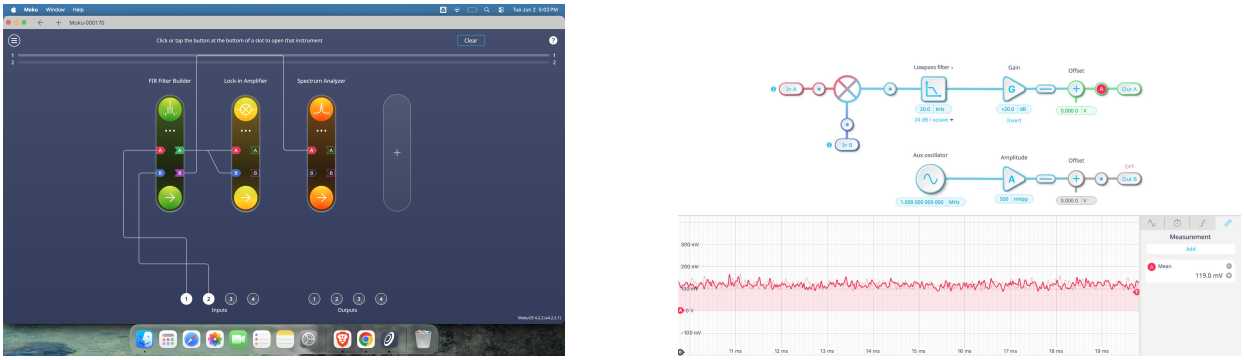
The overlap does not depend on where it is evaluated, but the *sensitivity to alignment errors* does. Prof. Adhikari suggested making both beams large ($\sim$1–2 mm waist), overlapping them where they are nearly collimated, and then focusing onto the photodiodes. Using Eqs. (21)–(22):

Table 6: Misalignment that costs 1% of overlap, for different waists, and the spot radius after an $f = 100\,\mathrm{mm}$ lens in front of a photodiode (1811 active diameter 0.3 mm).

Waist w_0	z_R	offset	tilt	waist pos.	spot ($f = 100\,\mathrm{mm}$)
0.305 mm (now)	0.27 m	30 μm	0.11 mrad	5 cm	111 μm
1 mm	2.95 m	100 μm	0.034 mrad	59 cm	34 μm
2 mm	11.8 m	200 μm	0.017 mrad	2.4 m	17 μm

A large waist is much more forgiving of lateral offsets and waist-position errors (a 1 mm beam is essentially collimated over the whole table), but less forgiving of tilt. Tilt, however, is the easiest error to see and correct by looking at the beams far away. Focusing afterwards is easy: an $f = 100\,\mathrm{mm}$ lens puts a 1 mm beam into a $\sim$34 μm spot, much smaller than the photodiode.

C Moku signal chain and shot-noise baseline



(a) Multi-instrument sequence.

(b) Envelope detection and shot-noise level.

Figure 16: (a) Moku:Pro multi-instrument configuration: FIR filter builder (gain balancing, subtraction and band-pass), lock-in amplifier used as a self-mixer (envelope detection), and spectrum analyzer. (b) Lock-in configuration (20 kHz low-pass, 24 dB/octave, +20 dB gain) and the shot-noise-level trace with the pump blocked (mean 119 mV $\approx V_{\mathrm{shot}}$).

D Miscellaneous work

Helping neighbouring experiments and HeNe beam profiling: During the first two weeks, while my mentor was away, I helped in neighbouring experiments and profiled several HeNe laser beams with the scanning-slit profiler, fitting their waist size and location in the same way as described in Sec. 4.2. This was excellent practice for the later profiling of the WOPA beams.

Profiling the 532 nm SHG beam: Early in the project I also profiled the 532 nm second-harmonic light generated in the waveguide when only 1064 nm is sent in ($\sim$0.1 mW of SHG).

This beam leaves the waveguide slightly tilted and also shows clear x–y asymmetry. It is useful mainly as an alignment reference, since SHG is maximized when the 1064 nm light is well coupled into the waveguide.

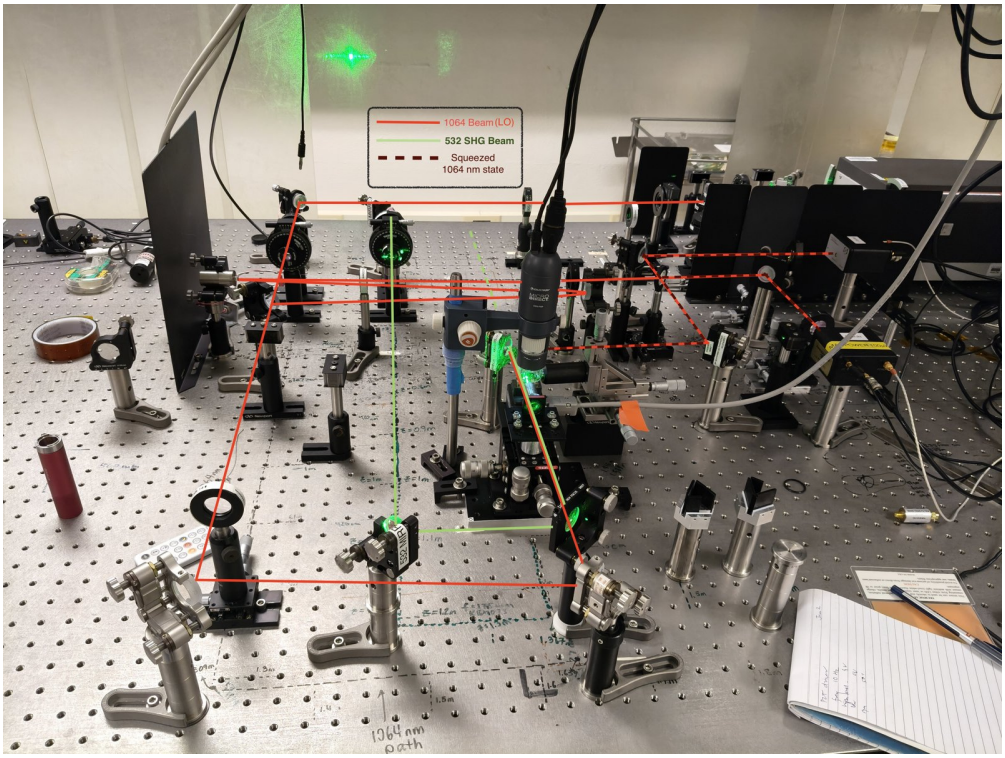


Figure 17: Photograph of the WOPA tabletop.