# Simulation of Advanced Virgo Phase Camera Signals

Christopher Layden[1]

Advised by Andreas Freise[2,3] and Matrin van Beuzekom[2]

August 18, 2021

[1] *The University of Texas at Austin*
[2] *Nikhef (Dutch National Institute for Subatomic Physics)*
[3] *Vrije Universiteit Amsterdam*

Final report for 2021 University of Florida Gravitational Physics IREU Program

**Abstract**

We present a model of the phase camera sensors in use at the Advanced Virgo gravitational wave observatory, based in the simulation program Finesse 3. As a proof of concept, we deliver a set of phase camera signals generated by this model for the Advanced Virgo dual-recycled interferometer with thermal distortions at the intermediate test masses. Through comparisons of real phase camera signals to similarly generated signals, defects within Advanced Virgo can be more easily diagnosed and ameliorated, allowing the detector to attain greater sensitivity to gravitational wave sources. While constructing this model, we have also developed new tools within Finesse 3 that allow for robust simulation of the control of complicated optical systems like the dual-recycled Advanced Virgo interferometer. With these tools, future users of Finesse 3 can ensure that their simulated optical systems reach their desired operating points.

## 1 Introduction

Since their first joint detection of a binary black hole merger in August of 2017 [1], the Advanced Virgo (AdV) and Advanced LIGO gravitational wave interferometers have together observed many gravitational wave sources—all from mergers of compact binaries. If we hope to observe more exotic gravitational wave sources or to sample a large enough population of binary mergers to deliver convincing astrophysical conclusions, we must continue to improve the sensitivities of these interferometers. Here, we approach this challenge by seeking to improve the control systems of AdV. The control loops of AdV keep the interferometer at its specified operating point, e.g. ensuring that the carrier laser is resonant in arm and recycling cavities and that the desired amount of light leaks into the dark port. However, defects in the optics of AdV (misalignments, tilts, point absorbers on test masses, etc.) can degrade the interferometer's sensitivity, even if the operating point is controlled. To assist in the diagnosis of these defects, optical phase cameras (PCs) have been installed near the bright port and dark port of AdV. These PCs provide images of the powers and phases of the carrier beam and all control sidebands within the interferometer. By understanding how various defects may affect these power and phase images, one might use the PC signals to diagnose and mitigate defects present in AdV.

To facilitate such an understanding of these signals, we have developed a model of the AdV PCs using the in-development software package Finesse 3. Along the way, we developed new analysis tools for Finesse 3. In particular, we created new tools that facilitate the control of a simulated interferometer to an operating point. We begin this report by describing the characteristics of the AdV PCs and introducing the theory underpinning their operation. We will then describe basic theory of control of the real AdV interferometer and discuss how we applied this theory to create new tools for simulating interferometer control in Finesse 3. Next, we apply Finesse 3—using these new tools—to create a model for the AdV PCs. We show that this model behaves as expected for real PCs, then apply the model to a robust simulation of the entire

AdV interferometer. Finally, we introduce various defects into this simulation and observe the resultant PC signals.

# 2 Background

## 2.1 Phase Camera Operating Principles

This discussion comes predominately from [2]. At present, there are two phase cameras installed in AdV, identical in design. These cameras image the inteferometer beam via pick off plates. One such pick off plate is placed between the power recycling mirror and the primary beamsplitter; the other is placed between the beamsplitter dark port and the output mode cleaner. The phase cameras measure profiles of the intensity and phase for the carrier beam and for all control sidebands.

The phase cameras measure these profiles using a scanning mirror, which scans over the interferometer beam and focuses the beam from a particular small spatial region onto a photodiode. But before entering the photodiode, the interferometer beam is combined with a reference laser beam, which is frequency shifted by 80 MHz with respect to the main laser beam. Without such a frequency shifted reference beam, it would be impossible to use demodulation to distinguish between upper and lower sidebands, as they would be demodulated at the same frequency. The radio frequency (RF) component of the photodiode signal is digitally demodulated to find the intensity and phase of all frequency components of interest (see Section 2.2 for technical detail). By performing this demodulation for all positions scanned over, the cameras obtain images of the intensity and phase images for each frequency component. Note, however, that due to curvature introduced by the scanner, the absolute phase images acquired are typically not useful. Rather, the differences between the phase images between two frequency components, e.g. between the carrier and a lower sideband, are used.

## 2.2 Theory of Modulation and Demodulation

The control of interferometers like Advanced Virgo (AdV) relies heavily on modulation of the main input laser to create control sidebands and demodulation of these control sidebands at various detectors along the interferometer. Here we describe the theory behind modulation and demodulation, with a particular emphasis on how the phase cameras of AdV apply demodulation to attain images of sideband power and phase. This discussion comes from [3] and [2].

At each point perpendicular to the direction of propagation, the laser light field that enters AdV can be represented as $E_{in} = E_0 \exp(i\omega_0 t)$. This laser passes through electro-optic modulators (EOMs), which modulate the phase of the light field, much like wiggling the position of the laser source back and forth. If $E_{in}$ enters an EOM that modulates the phase at frequency $\Omega$ and with amplitude (called modulation depth) $m$, the resulting field is

$$E_{out} = E_0 \exp(i(\omega_0 t + m \cos \Omega t)). \tag{1}$$

This field can be rewritten as an expansion in Bessel functions $J_i(x)$, which can be further simplified if $m << 1$:

$$\begin{aligned} E_{out} &= E_0 \exp(i\omega_0 t) \sum_{k=-\infty}^{\infty} i^k J_k(m) \exp(ik\Omega t) \\ &\approx E_0 \exp(i\omega_0 t)(1 + \frac{i\,m}{2}(\exp(-i\Omega t) + \exp(i\Omega t))). \end{aligned} \tag{2}$$

Written in this form, we can see that the field exiting the EOM is, to first order, a superposition of three fields: one at the original *carrier* frequency $\omega_0$, a *lower sideband* at frequency $\omega_0 - \Omega$, and a *upper sideband* at frequency $\omega_0 + \Omega$. These fields have slightly different resonance conditions. Therefore, if one can measure the field of one of the sidebands independently, they can use the sidebands to construct error signals to accurately control path lengths in an interferometer. This method of control is called Pound-Drever-Hall locking [4], and similar methods are used extensively in AdV.

To make such independent measurements of a certain frequency component requires the technique of demodulation. Here we discuss the theory of demodulation as it is applied at the phase cameras at AdV. Consider the light field that strikes the photodiode of the phase cameras. Accounting for all frequency components in the beam, we can write the light field comprising this beam as

$$E = \sum_{i=0}^{N} a_i e^{i\omega_i t}, \tag{3}$$

yielding an intensity

$$S_0 = E \cdot E^* = \sum_{i=0}^{N} \sum_{j=0}^{N} a_i a_j^* e^{i(\omega_i - \omega_j)t} \tag{4}$$

where $\omega_i$ are the frequency components in the beam and $a_i$ are the complex amplitudes of these components. Suppose we seek to demodulate one sideband with frequency $\omega_{sb}$. To do so, we mix this intensity signal with a signal $\cos \omega_d t + \phi_d$, with $\omega_d = \omega_{ref} - \omega_{sb}$. Here, $\omega_{ref}$ is the frequency of the reference laser. The resulting signal is

$$S_1 = S_0 \cos(\omega_d t + \phi_d) = \frac{1}{2} \sum_{i=0}^{N} \sum_{j=0}^{N} a_i a_j^* e^{i(\omega_i - \omega_j)t} \left( e^{i(\omega_d t + \phi_d)} + e^{-i(\omega_d t + \phi_d)} \right). \tag{5}$$

We then low-pass filter this signal to take only the DC component. The DC component comprises those terms for which $(\omega_i - \omega_j)$ equals $\omega_d$ or $-\omega_d$:

$$S_{1,DC} = \frac{1}{2} \sum_{i=0}^{N} \sum_{j=0}^{N} \left( a_i a_j^* e^{-i\phi_d} + a_j a_i^* e^{i\phi_d} \right) \text{ for } i, j \text{ s.t. } \omega_i - \omega_j = \omega_d. \tag{6}$$

This can be rewritten

$$S_{1,DC} = \mathrm{Re} \left( \sum_{i=0}^{N} \sum_{j=0}^{N} a_i a_j^* e^{-i\phi_d} \right) \text{ for } i, j \text{ s.t. } \omega_i - \omega_j = \omega_d. \tag{7}$$

Recall we have chosen $\omega_d = \omega_{ref} - \omega_{sb}$ for this particular demodulation. Typically, if the reference laser and modulation frequencies are chosen judiciously, there will be no other frequency component pairs with difference in frequency equal to $\omega_d$—and if there are, the amplitude of at least one of those components is small. If this condition holds, Eq. 7 reduces to

$$S_{1,DC} = \mathrm{Re}(a_{ref} a_{sb}^* e^{-i\phi_d}) = |a_{ref}||a_{sb}| \ \mathrm{Re}(e^{i(\phi_{ref} - \phi_{sb})} e^{-i\phi_d}). \tag{8}$$

In the demodulation process, the phase cameras actually apply two demodulations sequentially. We first demodulate with $\phi_d = 0$ to acquire the *in phase* signal

$$I = |a_{ref}||a_{sb}| \ \mathrm{Re}(e^{i(\phi_{ref} - \phi_{sb})}) = |a_{ref}||a_{sb}| \cos(\phi_{ref} - \phi_{sb}) \tag{9}$$

and then with $\phi_d = \pi/2$ to acquire the *in quadrature* signal

$$Q = |a_{ref}||a_{sb}| \ \mathrm{Re}(-ie^{i(\phi_{ref} - \phi_{sb})}) = |a_{ref}||a_{sb}| \sin(\phi_{ref} - \phi_{sb}). \tag{10}$$

With these two signals, we can calculate

$$|a_{ref}||a_{sb}| = (Q^2 + I^2)^{1/2} \tag{11}$$

and

$$\phi_{ref} - \phi_{sb} = \arctan(Q/I) \tag{12}$$

Assuming $|a_{ref}|$ and $\phi_{ref}$ are constant across the beam, an image of $|a_{ref}||a_{sb}|$ provides an image of the intensity of the sideband (technically, of its square root), scaled by the reference beam. Similarly, an image of $\phi_{ref} - \phi_{sb}$ is an image of the sideband phase, offset by the reference phase.

# 3 Developing Tools for Simulating Interferometer Control

In order to accurately predict the signals of the AdV phase cameras, we must first reliably simulate the light fields everywhere within the dual-recycled interferometer geometry of AdV. We perform this simulation with the program FINESSE 3, the in-development third version of the widely used interferometer simulation toolkit FINESSE [5]. As a part of this project, we developed, tested, and documentated new code for FINESSE 3. We in particular developed analysis tools called *Actions* within FINESSE 3, which allow users to easily perform previously complex actions like minimizing particular signals as a function of some degree of freedom or moving the entire interferometer to a controlled operating point.

The optical components of AdV must be carefully controlled in order to satisfy certain resonance conditions in the interferometer. The optimal configuration in which all of these conditions are satisfied is called the operating point. FINESSE 3 has a few Actions that facilitate the setting of the positions of optical components in order to reach this operating point. Here, we describe the basics of control of the real interferometer, the original implementation of this control with Actions, and our improvements to these Actions.

## 3.1 Basics of Interferometer Locking at AdV

In AdV, five degrees of freedom (DOFs) can be adjusted to reach this operating point. These DOFs respond to five error signals. Reaching the operating point of this multiple input, multiple output (MIMO) system amounts to finding the values of the DOFs at which all of the error signals are minimized. The five DOFs that are adjusted to control the interferometer are listed below, with lengths defined in Fig. 1.

$$\textbf{Differential Arm Length:} \quad DARM = \frac{L_N - L_W}{2} \tag{13}$$

$$\textbf{Common Arm Length:} \quad CARM = \frac{L_N + L_W}{2} \tag{14}$$

$$\textbf{Michelson Offset:} \quad MICH = l_N - l_W \tag{15}$$

$$\textbf{Power Recycling Cavity Length:} \quad PRCL = l_P + \frac{l_N + l_W}{2} \tag{16}$$

$$\textbf{Signal Recycling Cavity Length:} \quad SRCL = l_S + \frac{l_N + l_W}{2} \tag{17}$$

In traditional control of the AdV interferometer, each DOF is paired with and responds to one error signal. That is, for each error signal $E_i$, it is assumed that changes in only one DOF $F_i$ yield a change in $E_i$. While this assumption is not accurate, its application allows for effective control of the interferometer. The error signals used for control of AdV are, as mentioned in Section 2.2, demodulated measurements of the control sidebands at certain points of the interferometer. Three modulation frequencies are applied in AdV to create these control sidebands. These frequencies are all chosen such that the primary mode of each sideband is anti-resonant in the interferometer arms. The first modulation frequency, at $f_1 \approx 6.27\,\text{MHz}$, is resonant in the power recycling cavity. The second, at $f_2 = 9f_1 \approx 56.44\,\text{MHz}$, is resonant in both the power recycling cavity and the signal recycling cavity. The third, at $f_3 = 4/3f_1 \approx 8.36\,\text{MHz}$, is not resonant in either recycling cavity [6]. The error signals derived from these sidebands are listed below, in the following form: detector location, frequency of demodulation, phase of demodulation, corresponding DOF. The detectors where these measurements are taken are pictured in Fig. 2. Where the demodulation phase is I (i.e., in-phase), the demodulation phase is chosen such that the error signal is maximized relative to the DOF.

- B2 $f_3$ I: PRCL
- B2 $f_2$ Q: MICH
- B2 $f_1$ I: CARM
- B1p $f_2$ I: DARM
- B2 $f_2$ I: SRCL

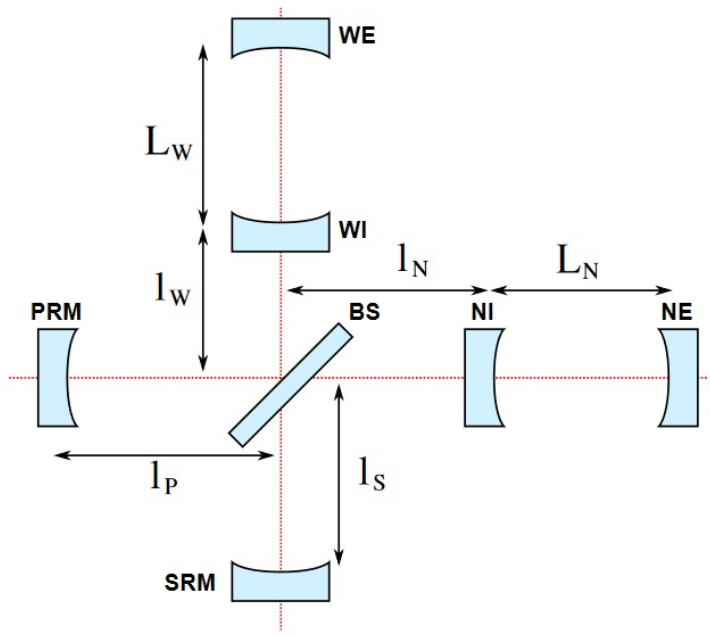


Figure 1: Lengths in AdV used in DOF definitions.

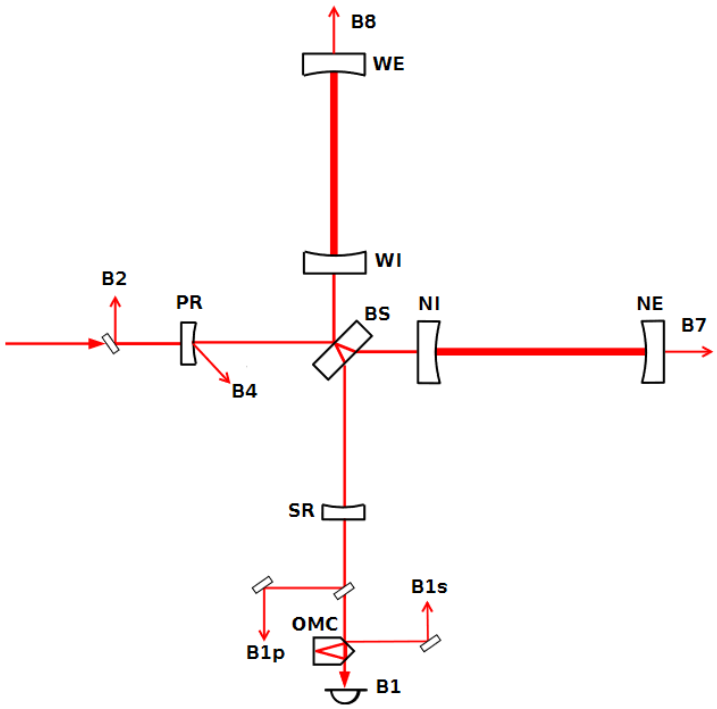


Figure 2: Detectors of AdV at which error signals are measured.

Much effort is required to "pre-tune" the interferometer, placing it in a configuration that is near enough to the operating point that the relationships between every $F_i$ and $E_i$ are either linear or very near zero. Once this pre-tuning is achieved, we then assume that

$$\frac{\partial E_i}{\partial \vec{F}} = \frac{\partial E_i}{\partial F_i} = G_i$$

where $G_i$ is the gain of $E_i$ with respect to $F_i$ [6]. From this, we see that if an error signals has some value $E_i$, a DOF $F_i$ should be changed by an amount $-E_i/G_i$ in order to move the error signal to zero.

## 3.2 Original Interferometer Locking

Initially, the Action used to lock a simulated interferometer, called **RunLocks**, applied the same algorithm described above. Assuming the interferometer was in a pre-tuned state near the operating point, the **RunLocks** calculated the values of the error signals, $E_{i,0}$ at the DOF values $F_{i,0}$. It then iteratively updated the DOF values according to

$$F_{i,j+1} = F_{i,j} - \frac{E_{i,j}}{G_i} \tag{18}$$

until each error signal is held below a predefined accuracy value.

## 3.3 Improved Interferometer Locking

However, this algorithm suffers in performance from the assumption that each error signal is a function of just one DOF. That is, without accounting for the relationships between all DOFs and all error signals, each iteration frequently overshoots the error signal minimum. As a result, especially when higher order modes are considered, the algorithm is often either slow to converge or does not converge at all.

The simulated control procedure can be improved by using the entire sensing matrix, which gives the gains relating all of the DOFs and all of the error signals. Note that we still assume that all of the relationships between DOFs and error signals are linear. Such a sensing matrix has the form of the matrix below:

$$\begin{bmatrix} \Delta E_1 \\ \Delta E_2 \\ ... \\ \Delta E_n \end{bmatrix} = \begin{bmatrix} G_{11} & G_{12} & ... & G_{1n} \\ G_{21} & G_{22} & ... & G_{2n} \\ ... & ... & ... & ... \\ G_{n1} & G_{n2} & ... & G_{nn} \end{bmatrix} \begin{bmatrix} \Delta F_1 \\ \Delta F_2 \\ ... \\ \Delta F_n \end{bmatrix}. \tag{19}$$

More compactly, $\vec{E} = \mathbf{G}\vec{F}$. The traditional approach of relating each DOF to a single error signal amounts to using only the diagonal of this matrix. But by observing a real sensing matrix like the one below, which was calculated from our simulation of AdV, we can observe that some off-diagonal components are greater than the diagonal component in a given row.

| | PRCL | MICH | CARM | DARM | SRCL |
|---|---|---|---|---|---|
| B2 $f_3$ I | **1.583e-02** | 3.699e-05 | -4.512e+00 | 1.615e-02 | 7.467e-06 |
| B2 $f_2$ Q | 1.514e-03 | **-3.208e-04** | -4.382e-01 | 1.568e-03 | 9.216e-13 |
| B2 $f_1$ I | -3.017e-03 | -5.362e-06 | **6.255e+00** | -2.238e-02 | -6.141e-06 |
| B1p $f_2$ I | -1.962e-03 | 3.423e-01 | 2.373e+00 | **9.749e+01** | -2.132e-04 |
| B2 $f_2$ I | -2.066e-03 | 4.439e-06 | 3.202e+00 | -1.146e-02 | **1.623e-04** |

Table 1: A sensing matrix calculated from our AdV simulation.

To minimize the error signal using the entire sensing matrix, we can treat the sensing matrix as a constant Jacobian matrix (since we have assumed we are in the linear regime) and apply the multidimensional Newton's method [7]. That is, we iteratively update the DOFs as follows:

$$\vec{F}_{i+1} = \vec{F}_i - \mathbf{G}^{-1}\vec{E}_i. \tag{20}$$

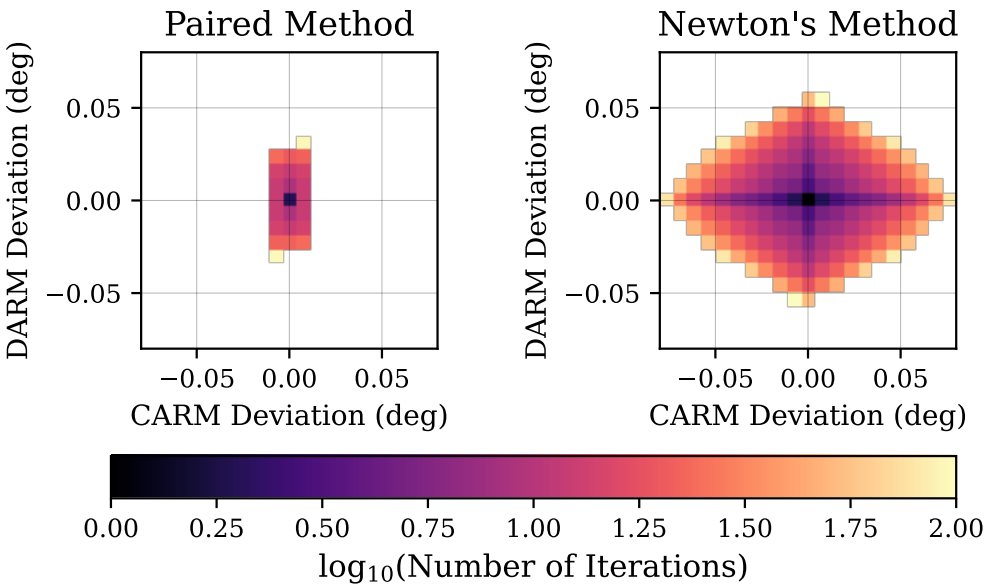


Figure 3: Iterations required for locking of example interferometer following kick of CARM and DARM DOFs. White pixels indicate failed locking.

The new action **RunLocksAdvanced** allows the user to choose between the original "paired" control method and the improved Newton's method. Theoretically, Newton's method should achieve lock for much faster and for a wider range of initial conditions than the paired method. The Finesse 3 documentation [8] provides a detailed example proving the advantages of Newton's method. This example defines a dual-recycled interferometer, similar to AdV, initially in a locked state. The DOFs CARM and DARM are then kicked away from the operating point, and both the paired and Newton's methods are used to try to bring the system back to lock. Fig. 3 shows the number of iterations required to reach lock for the two methods. We observe that Newton's method converges for a much larger fraction of the parameter space than the paired method. Moreover, for pixels in which the paired method converges, Newton's method requires far fewer iterations.

#### 3.3.1 Checking Assumption of Linearity

It is often useful to check the assumption that the interferometer has been appropriately pre-tuned (i.e. that all DOFs and error signals are related linearly). We created the action **CheckLinearity** for this purpose. This action displays a grid of the error signals as a function of each DOF independently, as shown in Fig. 4. This figure depicts the output of **CheckLinearity** for the same example used in the generation of Fig. 3, with the system locked except for a deviation of 0.01° in DARM. Nearly all of the plots in Fig. 4 are linear, except for three non-diagonal plots in the top row. However, one can observe that the x axes in these three plots are orders of magnitude larger than the first plot in the top row (which corresponds to a diagonal element in the sensing matrix). Thus, these three nonlinear plots correspond to three DOF-error signal relationships that have negligible gain and can therefore be assumed to not affect the control. To understand more rigorously why non-diagonal DOF-error signal pairs with such negligible gain will not affect control, consider a sensing matrix that is diagonal except for one very small off-diagonal element. The inverse of this matrix, which is used to adjust DOFs via Eq. 20, can be calculated via Gauss-Jordan elimination as follows:

$$\left[\begin{array}{cccc|cccc} A & \epsilon & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & B & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & N & 0 & 0 & 0 & 1 \end{array}\right]$$

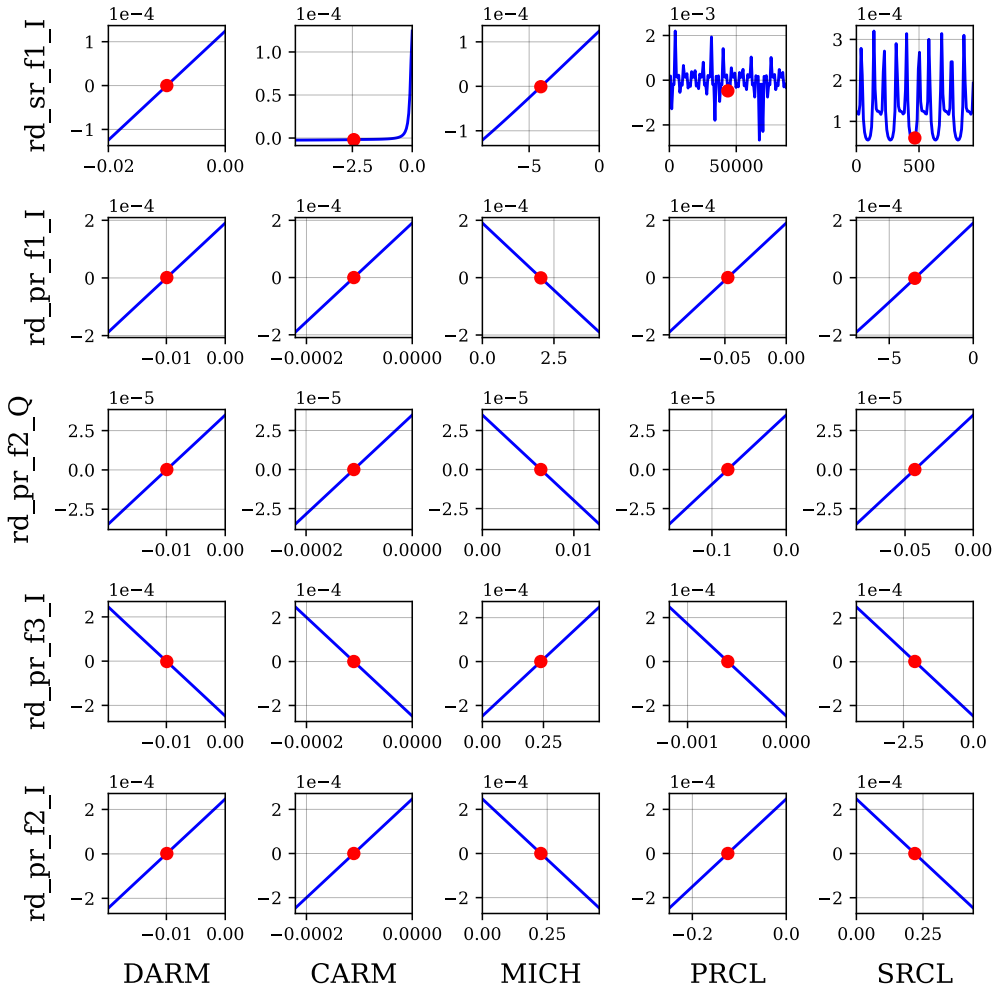


Figure 4: Output of **CheckLinearity** for a dual-recycled interferometer model. All signal-DOF pairs are either in the linear regime or are negligibly dependent on one another, so this state is suitable for our locking algorithms.

$$\Longrightarrow \left[\begin{array}{cccc|cccc} 1 & \epsilon/A & 0 & 0 & 1/A & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1/B & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1/N \end{array}\right]$$

$$\Longrightarrow \left[\begin{array}{cccc|cccc} 1 & 0 & 0 & 0 & 1/A & -\epsilon/AB & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1/B & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1/N \end{array}\right].$$

Thus, the small nonzero off-diagonal element of the sensing matrix is further suppressed in the inverse matrix, meaning that it will yield a negligible effect on the corresponding DOF.

### 3.3.2 Dragging Locks to States with Large Defects

However, it is often desirable to find the operating point for an interferometer that possesses a substantial defect (e.g. misalignment of mirrors or mismatched radii of curvature), for which pre-tuning to the linear regime is especially challenging. A solution is to find the operating point of the interferometer without the

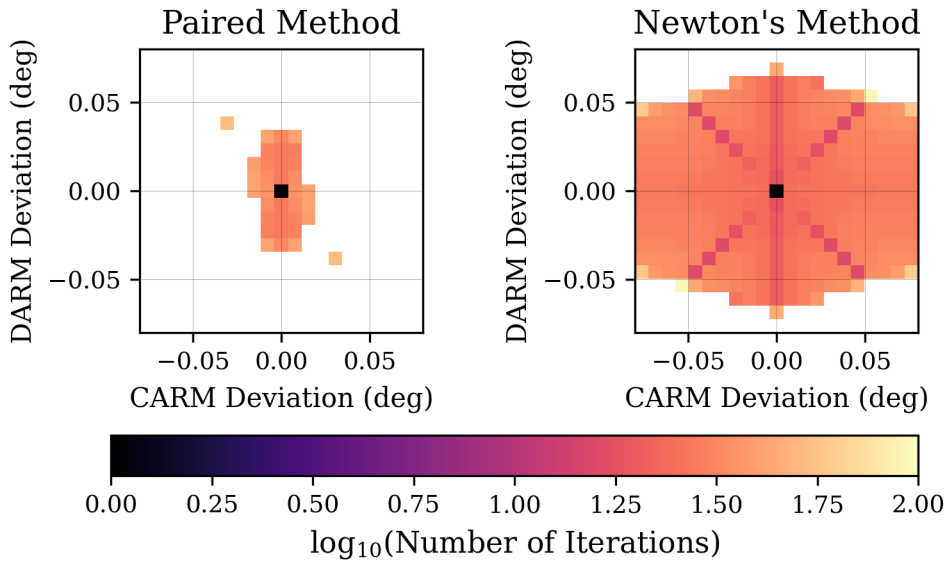


Figure 5: Repetition of the example leading to Fig. 3, but scaling all DOF changes by 0.5.

defect, then to gradually increase the magnitude of the defect, finding the operating point at each step of the way. For this purpose, we created an Action called **DragLocks**. This Action attempts to move to the final operating point of the interferometer over ten steps, at each step linearly adjusting parameters defining the defect and reaching lock. If the locking fails at any step, that step is broken down into ten additional steps; such decreasing of step size may continue until a maximum number of recursions is reached.

#### 3.3.3 Scaling DOF Adjustments

If locking is attempted when outside the linear regime, Eq. 20 will not work as well (or will not work at all). In this case, Eq. 20 is in particular prone to overshooting the operating point at each iteration. When this overshooting becomes too large, the iterations diverge and locking fails. To account for this problem, an option has been added to **RunLocksAdvanced** for the user to specify a factor that scales the change of each DOF. This method also applies for the paired method. Using a scale factor of 0.5, we repeated the tests used for comparing the paired and Newton's methods that led to Fig. 3. This scale factor resulted in Fig. 5, which shows that this scale factor allows for an improved domain over which locking succeeds. However, when the system is already close to lock, using a scale factor less than unity increases the number of iterations required for locking.

## 4 Modeling Phase Cameras with Finesse 3

Having thus improved the control of intereferometers in Finesse 3, we sought to create a model the PCs of AdV.

### 4.1 Implementation in Finesse 3

In order for our model to be a useful representation of the phase cameras, it should generate, for an arbitrary beam, demodulated intensity and phase images for the carrier and all sidebands. Our model does not, however, account for the curvature introduced by the scanner, so the absolute phase images from our model are free from such effects. For comparison with real phase camera signals, one may still subtract the phase images of any two frequency components. We developed and tested this model using a basic Michelson inteferometer geometry. Initially, the primary laser was phase modulated just once, at frequency 6.27077 MHz

and with modulation depth of 0.1. The output of the Michelson dark port was sent to a secondary beamsplitter, where it combined with a reference laser beam. The reference laser had an 80 MHz frequency offset from the primary laser. Initially, both the primary and reference lasers were in purely Gaussian modes.

There are many detector elements in FINESSE 3 that allow one to extract information about a laser beam's frequency components. If the beam is one-dimensional, the *amplitude detector* or *ad* element returns the intensity and phase of any frequency component without any demodulation—of course, such omniscience is as of yet impossible for a real detector. More realistically, the *power detector and demodulator* or *pd1* element demodulates a one-dimensional beam at a user-defined frequency. In analogy to *ad* but for a three-dimensional beam, the *field camera* or *fcam* detector gives images of the intensity and phase of the beam at a specified frequency. However, as of this writing, there is no element serving as a three-dimensional analogy to *pd1*. Such an element would be ideal for modeling the phase cameras.

Nonetheless, by using *fcam* detectors to reconstruct the frequency components of the beam and manually applying demodulation, we can model the phase cameras. In our model, we first determine all frequency components in the optical system, then place an *fcam* for each such frequency. Each *fcam* has $100 \times 100$ pixels and a width of three times the beam waist at the measurement point. After running the model, we therefore know the complex amplitude of every frequency component at every pixel.

With this knowledge, we can demodulate the signal at each pixel and at all desired demodulation frequencies. We find the demodulated signal for a sideband of interest with frequency $\omega_{sb}$ (and thus with demodulation frequency $\omega_d = \omega_{ref} - \omega_{sb}$) by first calculating

$$\sum_{i=0}^{N} \sum_{j=0}^{N} a_i a_j^* \text{ for } i, j \text{ s.t. } |(\omega_i - \omega_j) - \omega_d| < \epsilon \tag{21}$$

where $\epsilon$ is a small value (chosen to be 100 Hz) to account for any floating point error. Note the relationship between Eq. 21 and Eq. 7. If we find the magnitude and phase of the complex result of Eq. 21, we find the values $|a_{ref}||a_{sb}|$ and $\phi_{ref} - \phi_{sb}$, respectively. For generality, Eq. 21 uses the full sum rather than taking only $a_i a_j^* = a_{ref} a_{sb}^*$ as in Eq. 8.

Our model applies this procedure to find $|a_{ref}||a_{sb}|$ and $\phi_{ref} - \phi_{sb}$ for all pixels from the *fcam*s. It then repeats for all demodulation frequencies of interest. In the next section, we demonstrate the intensity and phase images produced by this model.

### 4.2 Evaluating the Model

We tested the model with the optical setup described above, with only one modulation frequency and with the primary and reference lasers in purely Gaussian modes. We observe the images of the intensities and phases of the carrier, upper sideband (USB), and lower sideband (LSB). We did this with the simple Michelson interferometer tuned to the dark fringe (Fig. 6) and then slightly detuned from the dark fringe (Fig. 7). As expected at the dark fringe, the intensity of the carrier field is zero to within floating point precision, while the intensities of the sidebands are nonzero. Meanwhile, when the interferometer is slightly detuned, the intensities of the sidebands are an order of magnitude smaller than that of the carrier, in agreement with the modulation depth of 0.1. For both cases, the shape of the intensity and phase maps are Gaussian and spherical, respectively, as expected.

We next test the model with the primary laser in a more complex higher order mode, namely the $LG_{33}$ mode. Again, we show the intensity and phase images when the interferometer is tuned to the dark port (Fig. 8) and then slightly detuned from the dark fringe (Fig. 9). Again, the intensity values and shapes of these images match our expectation for the $LG_{33}$ mode.

## 5 Observing Phase Camera Signals in AdV Model

Having thus confirmed the reliability of the FINESSE 3 phase camera model, we implemented the model into a pre-existing FINESSE 3 notebook (courtesy of Andreas Freise) that models all optical components of AdV and locks the positions of these optical components to stable operating conditions. In particular, we added such phase cameras in the B4 port and the B1p port of AdV. The B4 port corresponds to light deflected

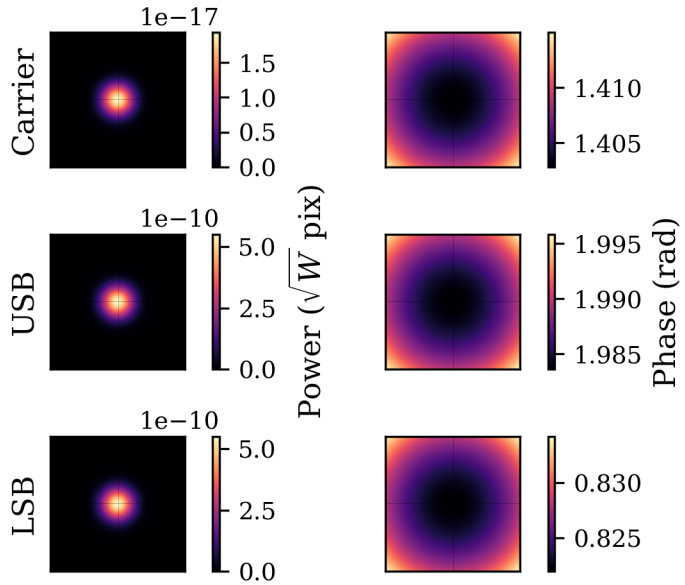


Figure 6: Phase camera image for a purely Gaussian beam and with interferometer tuned to the dark fringe.

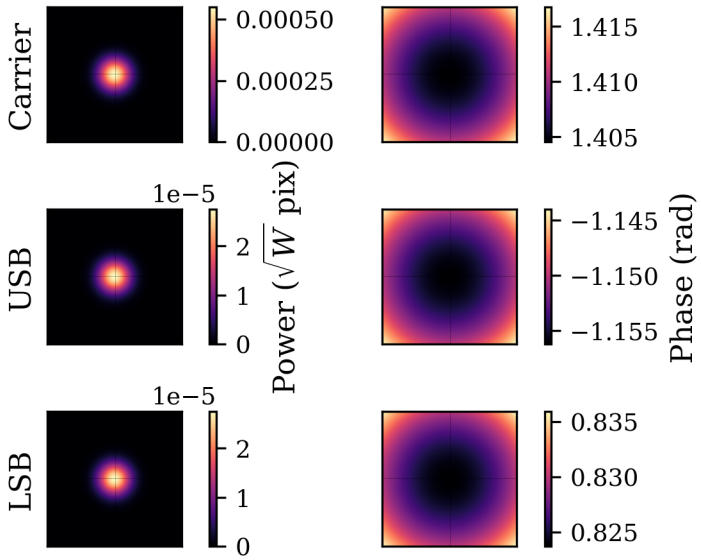


Figure 7: Phase camera image for a purely Gaussian beam and with interferometer slightly detuned (by 0.02°) from the dark fringe.

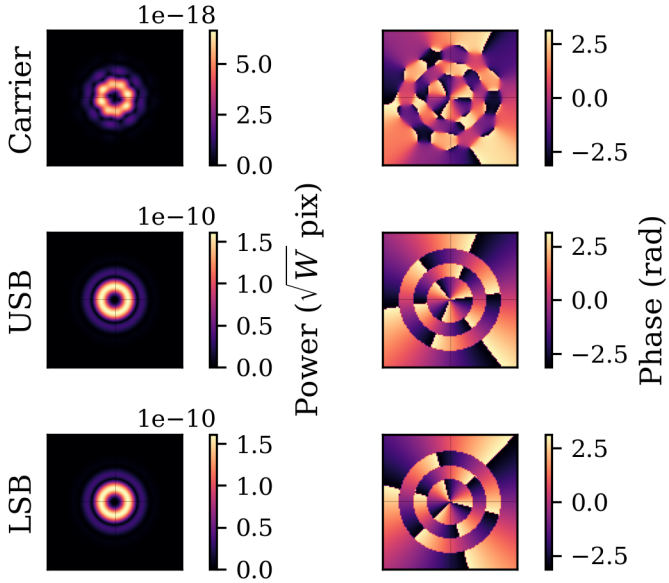


Figure 8: Phase camera image for a beam in the $LG_{33}$ mode and with interferometer tuned to the dark fringe.

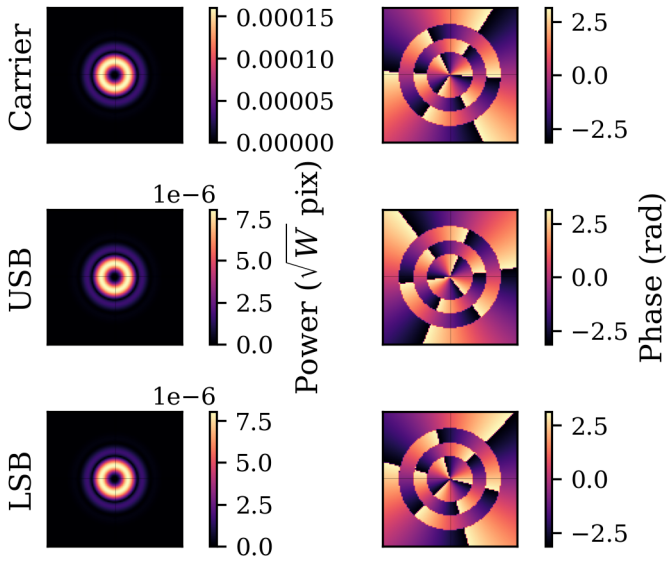


Figure 9: Phase camera image for a beam in the $LG_{33}$ mode and with interferometer slightly detuned (by 0.02°) from the dark fringe.

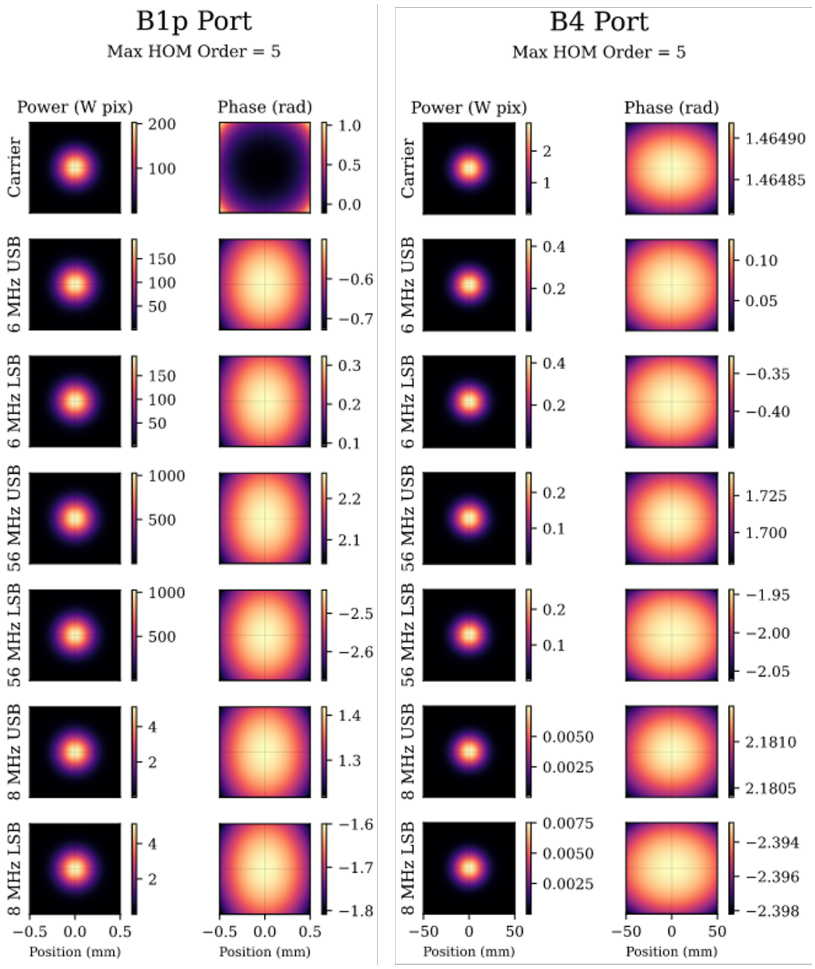


Figure 10: Simulated phase camera images for well-matched AdV interferometer.

from the power recycling cavity with a pick-off plate. That is, the pick-off plate is placed just down-beam of the power recycling mirror. The B1p port corresponds to light exiting the signal recycling cavity, just before the output mode cleaners. Using our phase camera model at these two ports, We generated images of the phases and powers of the carrier and the three control sidebands currently in use.

## 5.1 Well-Matched Base Case

We first set all of the optical parameters in this AdV model to their values specified in the AdV technical design report [9]. These values yielded a well-matched interferometer, which serves as a base case against which to compare scenarios for which defects have been added into the interferometer. We used the RunLocksAdvanced action to ensure that this simulated, well-matched AdV model was locked at the desired operating point. We then observed the images generated by our phase camera model. These images are shown in Fig. 10. We observe in these images that the intensity and phase images are again Gaussian and spherical, respectively, as expected. However, we can see a slight astigmatism in some of the phase images. Another feature of interest is the inversion between the image of the carrier phase and the images of all of the sideband phases at the B1p port. However, recalling that these images are actually of $\phi_{ref} - \phi_{sb}$, we can conclude that this difference is likely not too concerning–it simply means that the radius of curvature for the carrier was smaller than that of the reference laser, while the radius of curvature for the sidebands were larger than that of the reference laser. Recall that, for comparison with real phase camera images, we would subtract between these phase images, eliminating this dependence on the phase of the reference laser.

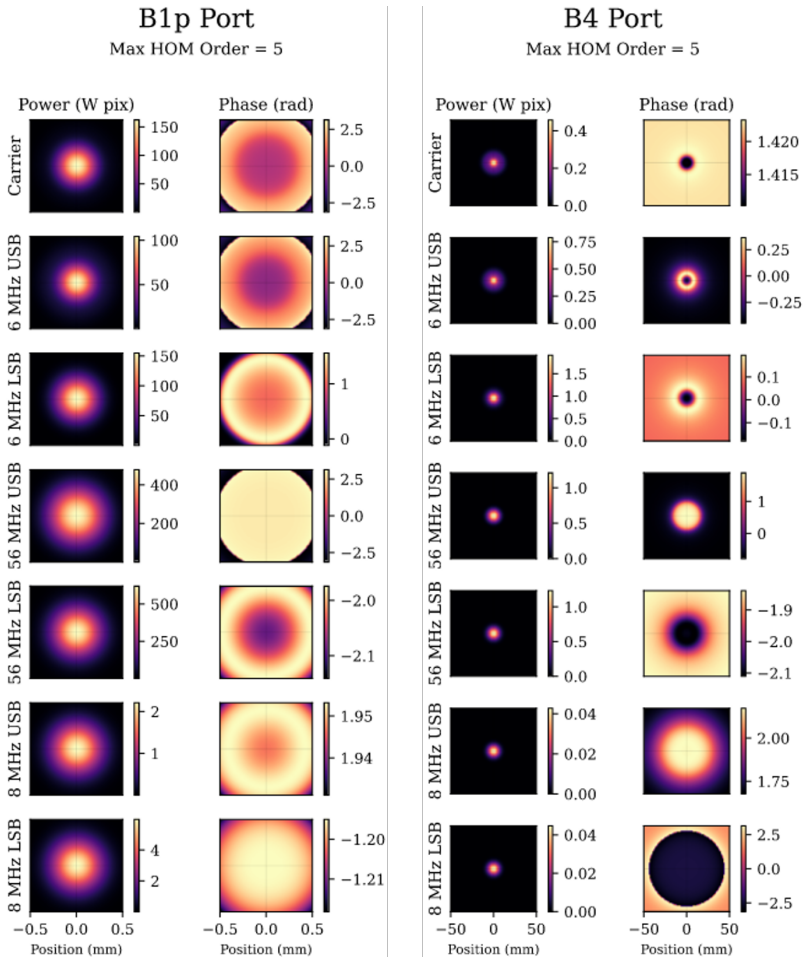


Figure 11: Simulated phase camera images when common lensing is introduced in the AdV interferometer.

## 5.2 Common Lensing

We next introduced optical defects in the simulated AdV interferometer and observed the resultant phase camera signals. We in particular studied the effects of thermal distortions on the intermediate test masses (ITMs). Such thermal distortions, arising from heating of the ITMs by the circulating beam, can be modeled as lenses just past the ITMs. We therefore introduced two such lenses into our simulated interferometer (one at each ITM), both of focal length 1e5 m. Because both focal lengths have the same sign, this is referred to as common lensing. When introducing this defect, we used DragLocks to keep the interferometer at its locked operating point. We then observed the images produced by our phase camera model, shown in Fig. 11. The intensity images are all still Gaussian, but we can now observe more interesting, radially symmetric features in the phase images.

## 5.3 Differential Lensing

Finally, we introduced differential lensing at the ITMs. That is, we changed the focal lengths of the lenses introduced at the ITMs to have opposite signs. We set one of the focal lengths to be -1e7 m and the other to be 1e7 m. We again used DragLocks to achieve lock in this configuration. The resultant phase camera images are shown in Fig. 12. We again see features of higher-order modes arise in the phase images. From this brief test, we observed that differential lensing causes a much larger effect on interferometer performance than common lensing. That is, even lenses with such wide focal lengths as $\pm$1e7 m had a large effect on the phase images. Further, even using DragLocks, we were unable to reach lock for differential lensing with focal

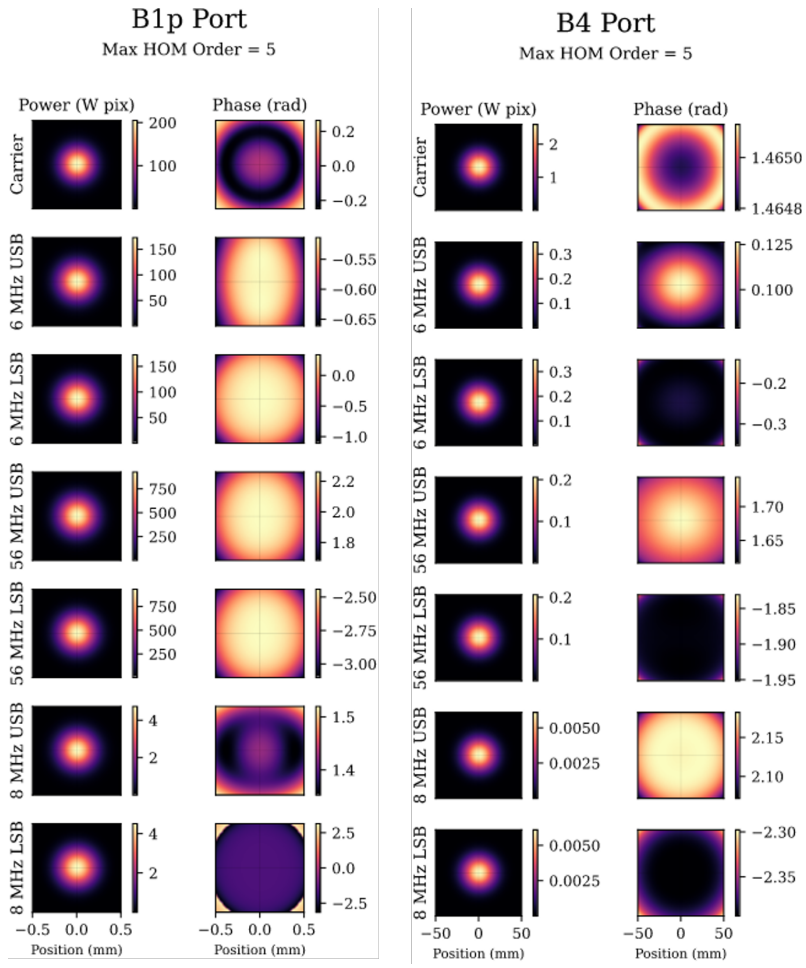


Figure 12: Simulated phase camera images when differential lensing is introduced in the AdV interferometer.

lengths of ±1e6 m.

Despite the interesting features arising in our images when modeling thermal lensing, we do not expect to see the same features in the real AdV phase cameras. This is because our simulations—both of AdV and the phase cameras themselves—leave out a fair amount of complexity. For example, mirror maps, which allow you to precisely define the surfaces of optics and any possible inhomogeneities, were not yet implemented in Finesse 3. Further, we do not model the electronics and particular scanning pattern of the phase cameras. Nonetheless, our phase camera model lays the computational foundation for future researchers to create more realistic images to compare with real images.

# 6 Conclusion

The phase cameras of Advanced Virgo are valuable for identifying and responding to interferometer noise that may otherwise go unnoticed, such as thermal distortions. As such, their use may improve the sensitivity of Advanced Virgo, and if adopted elsewhere, other gravitational wave observatories. We have developed a model of these phase cameras using the simulation program Finesse 3. Having applied this model to a simulation of the complete Advanced Virgo interferometer, we have generated phase camera signals for both the optimally designed inteferometer configuration and for configurations in which thermal distortions are present. We have also constructed new tools within Finesse 3 that allow for reliable simulation of steady-state interferometer control.

## Acknowledgements

The University of Florida Gravitational Physics International REU Program, and in particular this project, were financially supported by the National Science Foundation, grant numbers PHY-1950830 and PHY-1460803. Many thanks to Paul Fulda and Peter Wass for managing this program and making it a worthwhile career experience despite the challenges of the COVID-19 pandemic. Also, thank you to Jonathon Perry at Vrije Universiteit Amsterdam for answering many of my questions about interferometry and Finesse 3 this summer. Lastly, thank you to my advisors, Andreas Freise and Martin van Beuzekom, for their support and guidance throughout this project.

## References

[1] B.P. Abbott et al. Gw170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Phys. Rev. Lett.*, 119(141101), 2017.

[2] L. van der Schaaf. *The Phase Cameras of Advanced Virgo.* PhD thesis, Vrije Universiteit Amsterdam, 2020.

[3] C. Bond, D. Brown, A. Freise, and K. Strain. Interferometer techniques for gravitational-wave detection. *Living Rev Relativ*, 19(3), 2016.

[4] E. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. *Am. J. Phys.*, 69(79), 2001.

[5] A Freise, G Heinzel, H Lück, R Schilling, B Willke, and K Danzmann. Frequency-domain interferometer simulation with higher-order spatial modes. *Classical and Quantum Gravity*, 21(5):S1067–S1074, feb 2004.

[6] M. Valentini, Julia Casanueva, Maddalena Mantovani, and Perreca Antonio. Error signals for the longitudinal controls in dual recycled advanced virgo interferometer - steady state. 2020.

[7] T. Sauer. *Numerical Analysis.* Pearson, 2012.

[8] Finesse 3. `https://finesse.docs.ligo.org/finesse3/`, 2021.

[9] The Virgo Collaboration. Advanced virgo technical design report. 2012.