

UNIVERSITY OF CALGARY

Variable Beam-Splitter Reflectivity Estimation for Interferometry

by

Hamza Qureshi

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE

GRADUATE PROGRAM IN PHYSICS AND ASTRONOMY

CALGARY, ALBERTA

JANUARY, 2020

© Hamza Qureshi 2020

Abstract

Quantum parameter estimation is the scientific study that involves the use of quantum measurements for estimating an unknown parameter. Quantum-enhanced adaptive phase estimation is a well-studied example where the goal is to estimate the unknown phase such that the phase imprecision scales better than the standard quantum limit. Geometrically speaking, phase estimation is a $U(1)$ estimation problem where the objective is to make a highly precise estimate of the rotation of the initial state about the ordinate in the state space. This measurement technique has been shown to give better precision than the same measurement done classically.

Whereas other studies focus on phase estimation, we forge a new path that involves the variable beam-splitter reflectivity estimation. Our work proposes using quantum resources to solve the estimation problem. Here the phase is kept constant and the beam-splitter reflectivities are varied. We aim to achieve quantum-enhanced precision in the case of variable beam-splitter reflectivity estimation utilizing an evolutionary algorithm.

In this thesis, we show that variable beam-splitter reflectivity estimation is also a $U(1)$ rotation problem involving a different $U(1)$ subgroup. Geometrically speaking, the goal of the research work is to estimate the angle between the abscissa and the unknown axis in the equatorial plane the initial state is rotated about. We devise an optimization algorithm that designs a policy that estimates the unknown beam-splitter reflectivity whose imprecision scales better than the standard quantum limit with respect to the photon number. We employ the differential evolutionary algorithm, inspired by genetic evolution, for policy search. Differential evolution is an optimization algorithm that performs feasibility variant of the non-convex optimization and finds optimal policies for quantum-enhanced precision. In our work, we propose injecting an N -photon permutationally symmetric input state known as the sine state, where $N \in \{4, 100\}$. We do not include photon loss and noise in our model.

Our work sets the stage for two-parameter and multi-parameter quantum-enhanced estimation schemes.

Acknowledgements

I would like to thank my supervisor, Dr. Barry C. Sanders, for all his support, scientific and otherwise. Without his invaluable advice and timely critique, the completion of this degree would not have been possible.

I would like to express my gratitude to Pantita Palittapongarnpim for she played an instrumental role in teaching me the relevant concepts. She along with Peter Wittek also wrote the base library that was modified to create the evolutionary algorithm used to simulate the adaptive quantum-enhanced estimation scheme for this project. Moreover, I would like to thank Seyed Shakib Vadiae, Eduardo Paez and Camilla Mascardo for their time and insightful discussions.

Furthermore, I want to recognize Alexander Hentschel's role in the production of this thesis. I did not have a chance to meet him in person but my understanding of the topic relies heavily on the material provided by him in his PhD thesis.

The computational work was enabled by support from WestGrid through Compute Canada.

The research project was made possible with financial support from Natural Sciences and Engineering Research Council (NSERC).

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Figures and Illustrations	vi
List of Tables	viii
List of Symbols, Abbreviations and Nomenclature	ix
1 Introduction	1
1.1 Background	1
1.2 Aim and Motivation	4
1.3 Approach	6
1.3.1 Interferometric setup	6
1.3.2 Algorithmic policy search	7
1.3.3 Determining the imprecision scaling	8
1.4 Structure of the thesis	9
2 Background	12
2.1 Interferometer statistics	13
2.1.1 Bloch-sphere rotation	14
2.1.2 Beam-splitter optics	16
2.2 Quantum-enhanced metrology	17
2.3 Quantum imprecision	18
2.3.1 Cramér-Rao lower bound	20
2.3.2 Standard quantum limit and Heisenberg limit	21
2.4 Adaptive quantum-enhanced metrology	22
2.5 Evolutionary algorithms	25
2.5.1 Evolutionary algorithms for non-convex optimization	26
2.5.2 Differential evolution	27
2.6 Cumulants	29
2.6.1 Moment-generating function	29
2.6.2 Definition and nomenclature	30

2.6.3	Skewness and kurtosis	30
3	Adaptive variable beam-splitter reflectivity estimation	33
3.1	Input state and the interferometer	34
3.1.1	Input state	34
3.1.2	The interferometer transformation	36
3.2	Adaptive measurement procedure	37
3.2.1	Adaptive mechanism	37
3.2.2	Decision making	39
3.3	Imprecision and scaling	40
3.3.1	Sharpness and Holevo variance	42
3.3.2	Asymptotic power-law scaling	43
3.4	Policy search and implementation	44
3.4.1	Policy search using differential evolution	45
3.4.2	Implementation on a high-performance computer	46
3.5	Algorithm complexity	48
3.5.1	Policy time cost	48
3.5.2	Implementation cost	50
4	Results	51
4.1	Distribution of the beam-splitter reflectivity estimates	52
4.2	Asymptotic power-law scaling	55
4.3	Policy space and optimization algorithm	57
5	Discussion	60
5.1	Distribution of the beam-splitter reflectivity estimates	60
5.2	Asymptotic power-law scaling	64
5.3	Policy space and optimization algorithm	66
6	Conclusion	69
6.1	Summary	69
6.2	Future research	72
	Bibliography	74
A	Evolutionary algorithm for variable beam-splitter reflectivity estimation	83
A.1	About the software	83
A.2	Copyright and license	84
A.3	Download link	84

List of Figures and Illustrations

1.1	Interferometric setup is shown with an unknown beam splitter parametrized by θ , a constant phase shift ϕ_0 and a control beam splitter parametrized by Θ . An N -particle state being injected into the two beam-splitter ports a and b is also shown. The processing unit (PU) drawn controls the control beam splitter. [Reproduced after modification from Ref. [31], Fig. 2.1].	4
2.1	Beam-splitter reflectivity estimation involves estimating the angle between the abscissa and the axis of rotation. This angle is $\theta/2$. Initial state vector shown is rotated about the axis of rotation by an amount ϕ_0 to achieve the final state of the vector.	15
2.2	Beam splitter shown mixing two modes. Here a_i are the input amplitudes while b_i are the outputs.	16
2.3	Quantum-enhanced metrology procedure involves state preparation, interaction of the state with the Mach-Zehnder interferometer and the measurement. The direction of the arrows indicates time flow.	18
2.4	The distribution of the estimates is shown. The actual value ϕ , the estimate mean $\tilde{\phi}$ and the estimate standard deviation $\Delta\tilde{\phi}$ are also shown. [Reproduced from Ref. [50], Fig. 2].	19
2.5	A depiction of optimization landscapes. (a) The image on the left shows a convex set while the one on the right shows a non-convex one. The dotted line represents the shortest path between the two points \mathbf{x} and \mathbf{y} . (b) For each pair (k, β) , the sharpness S (Eq. 3.14) of the policy parametrized by (k, β) is shown. The Bayesian inference-based policy is represented by $k = 1.596$ and $\beta = 1.096$ (marked by \bullet). For comparison, the global maximum is at $k_{\max} = 1.631$, $\beta_{\max} = 1.015$ (marked by the arrowhead). [Reproduced from Ref. [31], Fig. 5.1].	26
3.1	A plot of the sine state for $N = 50$. The red curve shows the fitted Gaussian curve. [Reproduced after modification from Ref. [31], Fig. 2.2].	34
3.2	Adaptive quantum-enhanced metrology procedure showing that the input state is divided into M bundles of L particles, and each bundle then interacts with the interferometer in a sequential manner. The processing unit (PU) determines the control parameter Θ_m taking into account the history of the measurement outcomes up to that point, $\mathbf{x}_m = x_1 x_2 \dots, x_m$. [Reproduced after modification from Ref. [50], Fig. 1].	38

3.3	An example of a binary decision tree with three particles. The red line represents one run. The initial controllable beam-splitter reflectivity is Θ_0 and is adjusted by $\pm\Delta_m$ depending on $x_m \in \{0, 1\}$. [Reproduced after modification from Ref. [49], Fig. 4.3].	40
3.4	The distribution of the estimates. The actual value θ , the estimate mean $\tilde{\theta}$ and the estimate standard deviation $\Delta\tilde{\theta}$ are also shown. [Reproduced after modification from Ref. [50], Fig. 2].	41
3.5	Accuracy and precision of the probability distribution are shown. The target value and the mean of the estimate distribution play an important role in determining these quantities.	42
4.1	The distribution of the estimates for different θ values in the case of 50 and 100 photons. (a) When θ and the target value are $\pi/2$, the mean estimate recorded is $\tilde{\theta} = 0.400\pi$. This run involves 50 photons. (b) In the case of 100 photons, the mean estimate recorded is $\tilde{\theta} = 0.571\pi$. (for $\theta = \text{target value} = \pi/2$). (c) The mean of the estimate distribution is 0.330π when $\theta = 3\pi/4$ and the target is $\pi/4$. (50 photons). (d) When $\theta = 3\pi/4$ and the target is $\pi/4$ for 100 photons, $\tilde{\theta} = 0.280\pi$	52
4.2	Logarithm of Holevo variance vs $\log(N)$ generated using the sine state for $N = 4$ to 100. The ideal (no photon loss) experiment was repeated with beam-splitter reflectivities equal to 0.25, 0.5, 0.75 and 0.91. SQL and HL are shown for comparison.	55
4.3	The scalings are shown for several beam-splitter reflectivities with no photon loss. The least squares regression line is included to study the trend exhibited. SQL level is drawn for reference.	57
4.4	The policy-search space is parametrized as $\Delta_m = k/m^\alpha$ where k is a constant. (a) A plot of sharpness vs α is shown when $k = \pi$. (b) A plot of sharpness vs α is shown when $k = 3\pi/4$	58
4.5	The adjustments Δ_m executed by the processing unit to control the beam-splitter reflectivity Θ in a 100-photon AQEM scheme.	58
5.1	The graphs for $r = \cos^2(\theta/2)$ and for $r = \cos^2((\pi - \theta)/2)$ when $\theta \in [0, 4\pi - 1]$ are shown.	61

List of Tables

4.1	The sharpness, skewness, kurtosis and related quantities of the probability distributions shown in Fig. 4.1 are listed. The results are listed for $N = 50$ photons and $N = 100$ photons.	54
4.2	The imprecision scaling, reflectivity, original ordinate intercept (W) and the goodness of fit values are shown for various beam-splitter reflectivity values.	56
4.3	Theoretical and empirical values shown for generating the policy and for implementing the AQEM algorithm	59
5.1	K values are shown for $N \in \{1, 2, 3\}$	64
5.2	$\Delta_m = k/m^\alpha$ values are shown for $k = \pi$ and $\alpha \in \{1, 2, 3\}$	67

List of Symbols, Abbreviations and Nomenclature

Symbol or abbreviation	Definition
QEM	Quantum-enhanced metrology
AQEM	Adaptive quantum-enhanced metrology
CRLB	Cramér-Rao lower bound
SQL	Standard quantum limit
HL	Heisenberg limit
CPU	Central processing unit
GPU	Graphics processing unit
DE	Differential evolution
MPI	Message Passing Interface
VSL	Versatile Scripting Language
SE	Standard error
N	Number of particles
θ	Unknown beam-splitter reflectivity
Θ	Controllable beam-splitter reflectivity
$\tilde{\theta}$	Beam-splitter reflectivity estimate
$\bar{\theta}$	Mean estimate
$\Delta\tilde{\theta}$	Estimate's uncertainty
Δ	Beam-splitter reflectivity adjustment
ϕ	Phase shift
x	Measurement outcome
\mathbf{x}	Set of measurement outcomes
ψ	Input state
F	Fisher information
b	Parameter bias
V_H	Holevo variance
S	Sharpness
$\overline{R^2}$	Adjusted coefficient of determination
\mathcal{H}	Hilbert space
U	Unitary operator
L	Symmetric logarithm derivative
ρ	Quantum state
P	Probability distribution

p	Probability
ϱ	Policy
H	Quantum Fisher information matrix
Ψ, ϕ_r, ϕ_s	Rotation parameters
\hat{J}	Rotation operator
f	Function
\mathcal{C}	Set
ς	Scaling
μ	Moment
M	Moment generating function
K	Cumulant generating function
κ	Cumulant
γ	Pearson's moment coefficient of skewness
σ	Standard deviation
ϵ	Interferometer mode
$d_{m,m'}^j$	Wigner d -function
\hat{B}	Beam-splitter operator
\hat{P}	Phase-shifter operator
$\hat{\sigma}$	Pauli operator
v	Number of data points
y'	Predicted data point
δ_y	Error in y
N_p	Population size in DE
C_r	Crossover rate in DE
\mathcal{F}	Mutation factor in DE
Y	Candidate solution in DE
\mathbf{Y}	Candidate solution vector in DE
V	Offspring in DE
\mathbf{V}	Offspring vector in DE
α	Policy space parameter
τ_{pol}	Time cost for generating the policy
τ_{imp}	Time cost for implementing the policy
Ω_{fit}	Time cost for executing fitness function
Ω_{off}	Time cost for generating an offspring

Chapter 1

Introduction

In this chapter, we present the necessary background, aim and motivation and our approach. The structure of the thesis is included at the end of the chapter to guide the reader.

1.1 Background

We begin by surveying the field of quantum technology and by explaining the importance of quantum control as a tool for enabling quantum-enhanced metrology. We cite gravitational wave detectors and atomic clocks as two well-studied applications.

Quantum technology combines quantum physics and engineering in order to produce viable and scalable practical applications [20]. The field utilizes principles, techniques and tools in quantum physics and quantum information science to create applications such as quantum metrology [26], quantum sensing [15] and quantum communication [70] among others. Quantum metrology deals with measurement procedures that involve enhancement in precision, efficiency, simplicity of implementation, etc. through the use of quantum effects [26]. Quantum sensing describes the use of a quantum system, properties or quantum phenomena for performing measurements of a physical quantity of interest [15]. Quantum communication takes advantage of the laws of quantum physics to protect the data [70]. This can be done, for instance, by employing quantum cryptography.

Quantum control is what makes these applications possible. One of the main goals in quantum control theory is to establish a firm theoretical footing and develop a series of systematic methods for the active manipulation and control of quantum systems [18]. This is a field in itself because the laws of quantum physics are different from the ones in classical physics and therefore a new framework is needed for systems governed by quantum physics. The manipulation and control of such systems are achieved using quantum resources such as entangled particles [78]. The goal usually is to achieve the desired quantum state, and the problem solving is centered around figuring out how to drive the dynamics such that the desired state is achieved [12]. Quantum control theory has attained significant successes in quantum optics among other areas [52].

Data-driven control is an approach in control theory that does not utilize the dynamic model but instead devises the control policies directly from the data [37]. This approach is quite useful as the dynamic model is difficult to obtain and represent faithfully. This is especially true when dealing with hard-to-describe models in the age of complex design [49].

One of the data-driven techniques is policy search [37]. In this technique, the policy is the procedure carried out by the processing unit. The policy is given a task-appropriate mathematical form, then its performance is optimized by searching the space of the policy's parameters [16]. The goal of the exercise is to find the optimal set of parameters. A black-box optimization algorithm [54] can be employed for a model-free policy search. The algorithm in question is called black-box because it can be viewed in terms of its inputs and outputs (or transfer characteristics), without any knowledge of its internal workings. The implementation is 'opaque' or black. Such an algorithm determines the performance of the policy in question by implementing it on an experimental setup [16].

One of the applications of quantum control is quantum-enhanced metrology where quantum resources such as entangled states are used to make better-than-classical or quantum-enhanced measurements [4]. This translates into reducing the uncertainty in the parameter or variable that is being measured. Two very common ways this is achieved is by repeat-

ing the experiment a number of times and by using more of an appropriate resource [39]. Depending on the configuration of the model and the setup used, these stated techniques may or may not be effective in solving the problem. Quantum-enhanced metrology studies the nuances related to the model, the technique and the implementation to realize the best path forward [66]. Quantum-enhanced metrology has found its use in gravitational wave detection [11] and atomic clocks [6].

Gravitational wave detectors are, at the fundamental level, interferometers [55]. A gravitational wave creates perturbations in the space and time and therefore introduces a relative change in the path length when it passes through the interferometer [69]. This in turn introduces a relative phase shift in the laser beams that run across the length of the arms of the interferometer. The precision with which the gravitational wave in question can be detected depends on how accurately the generated phase shift can be measured [55]. Quantum-enhanced metrology makes the hyper-precise phase measurement possible [39].

Atomic clocks are the most accurate time-keeping device in existence [44]. In an atomic clock, the atomic transitions of the atoms are used to quantify the passage of time. The device uses a hyper-fine transition frequency in the microwave, or electron transition frequency in the optical, or ultraviolet region of the electromagnetic spectrum of atoms [45]. These transitions are the frequency standards, which are used to synchronize other, more practical clocks that are easier to use or move around [44].

Adaptive quantum-enhanced metrology techniques employ a feedback mechanism to achieve quantum-enhanced precision [48]. The setup is configured in such a way that the detector measurements are used to adjust some parameter of the model in order to achieve very accurate measurements. Adaptive phase estimation is an example of quantum-enhanced metrology [31, 49]. The goal in phase estimation is to estimate the unknown phase with quantum-enhanced precision [74]. Adaptive techniques use adaptive measurements to reduce the uncertainty in phase measurement more than what is achieved using non-adaptive quantum metrological schemes.

Model-free and data-driven adaptive phase estimation is when the physics of the model is not taken into account [37]. This technique uses measurement data to measure the unknown phase with quantum-enhanced precision. This scheme uses black-box optimization algorithms to search for optimal policies that reduce uncertainty in the phase measurement [49].

In order to have more precise measurements, two-parameter or multi-parameter quantum-enhanced estimation techniques must be employed [62].

1.2 Aim and Motivation

In this section, we explain the underlying motivation and the scope of this research project.




Figure 1.1: Interferometric setup is shown with an unknown beam splitter parametrized by θ , a constant phase shift ϕ_0 and a control beam splitter parametrized by Θ . An N -particle state being injected into the two beam-splitter ports a and b is also shown. The processing unit (PU) drawn controls the control beam splitter. [Reproduced after modification from Ref. [31], Fig. 2.1].

Our work is centered around adaptive quantum-enhanced metrology (AQEM) using a Mach-Zehnder interferometer with variable beam-splitter reflectivity. The interferometer is shown in Fig. 1.1. Previous work [4, 31, 49] utilizes 50 : 50 beam splitters with a phase shift that can be varied. The AQEM example that we use in this thesis has non-constant beam-splitter-reflectivity and constant phase. In other words, our work involves adaptive quantum-enhanced variable beam-splitter reflectivity estimation.

The assumption that the beam splitters are 50 : 50 may not be realistic in a lab setting where equipment-related errors are always present. Measuring the beam-splitter reflectivity before installing the device in the experimental setup does not ensure that there will be no alignment errors while installing the device. If the beam splitters are not 50 : 50 but the estimation scheme assumes so, then this can lead to less than ideal measurements. This oversight might even prevent the experimentalists from achieving the quantum-enhanced precision in the first place. The goal in this thesis is to address precisely this problem. We keep the phase shift constant so as to explore and fully understand the AQEM dynamics in the case of variable beam splitters.

Our work will facilitate the transition from one-parameter to two-parameter quantum-enhanced metrology by providing insights about the second parameter (phase being the first one). In order to achieve quantum-enhanced metrology with two parameters, quantum fisher information matrix H needs to be calculated, which is a two by two matrix [62]. Estimation schemes with one parameter involve the calculation of quantum Fisher information, which is a scalar [31]. Therefore, moving to the two-parameter realm requires a two by two matrix.

The quantum-enhanced metrology process consists of three stages: detector state preparation, interaction with the quantum system containing the parameters under scrutiny and the detection (measurement) [48]. The elements of the quantum Fisher information matrix are calculated using the probability of the measurement outcomes and the post-measurement detector state [62]. It turns out that the non-diagonal elements in this case are zero. This leaves only the diagonal elements. H_{11} is for the phase shift present in the Mach-Zehnder interferometer and H_{22} is what our work is focused on. Here H_{11} and H_{22} correspond to the maximum information content gained as a result of detector measurements in the case of phase shift and beam-splitter reflectivity respectively [10]. This diagonalization works only if there is a fixed number of particles involved [62].

Another motivation for taking the variable beam-splitter reflectivity estimation route in the quantum metrology realm is to see if we can gain any measurement-related insights that

we could not with phase estimation. In other words, we are also looking for insights that are not directly related to beam-splitter reflectivity but emerge serendipitously in the form of patterns and correlations during different stages of the problem-solving process.

After all is said and done, the knowledge gained from the individual studies of phase estimation and variable beam-splitter reflectivity estimation can be put together to tackle the two-parameter quantum-enhanced metrology.

1.3 Approach

In this section, we detail the approach adopted in this thesis. We summarize how the interferometer is configured, how the algorithm implements policy search and how the intended results are obtained given the setup.

1.3.1 Interferometric setup

Adaptive quantum-enhanced metrology (AQEM) is an approach that attains quantum-enhanced precision by approximating optimal measurement on the quantum state [73]. We use the technical and mathematical framework of a general discrete-time AQEM scheme to study variable beam-splitter reflectivity estimation [49]. Our goal is to develop an algorithm that devises a policy to be implemented by the processing unit (PU). PU controls the reflectivity of the control beam splitter to achieve quantum-enhanced precision in the case of variable beam-splitter estimation. We investigate and gain insight regarding how the input state and the optimal policy contribute towards achieving quantum-enhanced precision using adaptive measurements.

In this thesis, we consider the case of a noiseless and lossless, two-mode Mach-Zehnder interferometer (Fig. 1.1), where the first beam splitter has unknown reflectivity $\theta \in [0, \pi]$ while the second beam splitter has the controllable reflectivity $\Theta \in [0, \pi]$. The input state that we use in our work is a permutationally symmetric N -photon input state called the sine

state [4]. We proceed by splitting up the N -particle input state into individual particles and send in only one particle at a time. Our goal is to study the dynamics for $N \in \{4, 100\}$ as far as estimation goes. The order in which the photons are injected into the interferometric ports does not matter [33]. Also, the sine state is a loss-tolerant input state meaning it remains permutationally symmetric even when some of the photons are lost due to various factors [31].

Two detectors, one for each arm, detect particles on the other end and the measurement results are fed into a processing unit that controls the reflectivity of the controllable beam splitter. The estimation scheme varies the controllable beam-splitter reflectivity such that the imprecision achieved once all particles are through scales in a quantum-enhanced manner.

The way in which the controllable beam-splitter reflectivity is varied or adjusted is dictated by the policy determined by the black-box optimization algorithm. As this method does not use the model or the configuration in its design, it is considered data-driven and model-free.

1.3.2 Algorithmic policy search

We employ an evolutionary algorithm for searching the solution space for optimal estimation policies. Evolutionary algorithms are the ones inspired by how genes combine and mutate and how biological reproduction takes place at the genetic level [65]. A given number of candidate solutions are randomly generated in the beginning. Then the best performing solutions are combined and ‘mutated’ to produce the next generation of solutions. This ensures that the solutions get better at the task at hand with every generation [21]. This process is repeated until a good enough solution is found.

Initially, we run the evolutionary algorithm to a predetermined number of generations. When that stops giving us satisfactory solutions, in order to improve the scalability of the algorithm, we introduce a criterion for accepting policies based on its performance instead of constraining the number of generations [49]. This enables the optimization algorithm to

continue running until a viable policy is found. The policy search in our case is a logarithmic search process. We run the algorithm up to 100 photons beyond which point the computation cost becomes too high for the simulation to be practical. As the solution space is highly non-convex, we take extra measures to ensure early termination of the optimization runs. Some of these measures include parallelization of the calculations involved and vectorization for efficient processing [38].

The optimal policy directs the processing unit in adjusting the controllable beam-splitter reflectivity Θ . The adjustment process becomes finer and finer as the new measurements are collected and more information is gained about the unknown parameter being estimated.

Evolutionary algorithms have shown success up to 100 photons in the case of phase estimation both with and without noise incorporated into the scheme [49]. The goal in this thesis is to find an optimal policy to accurately estimate variable beam-splitter reflectivities, using this algorithm.

1.3.3 Determining the imprecision scaling

In quantum-enhanced metrology, the performance of an estimation scheme is measured in terms of the power-law scaling of the imprecision with respect to the number of particles N [31]. For an effective comparison, the scaling is measured against established bounds. Classical techniques give us the lower bound called the standard quantum limit (SQL) and the quantum techniques employing quantum resources arrive at Heisenberg limit (HL) [9]. SQL asymptotically scales as $1/\sqrt{N}$ and HL scales as $1/N$ [7].

In quantum-enhanced metrology, the way a certain estimation scheme's imprecision scales matters more than the absolute imprecision associated with a particular N value [49]. This is because the design and implementation success of the technological products is inextricably tied to whether they can be scaled up or not.

In our work, using the optimal policy, we generate the distribution of estimates for a given unknown beam-splitter reflectivity value θ . This enables us to calculate sharpness or

the accuracy and the precision of the generated estimate. This, in turn, is used to calculate the variance of the distribution. From there, we determine the imprecision scaling of the estimate with respect to the number of photons $N \in \{4, 100\}$. This is done by plotting and then studying the log-log graph. The comparisons with the SQL and HL can then be made.

Geometrically speaking, estimating the unknown beam-splitter reflectivity is equivalent to estimating the unknown axis in the equatorial plane the initial state is rotated about, in the state space. The angle of rotation of the initial state is given by the phase shift present in the interferometer.

The results of our work can also be used to compare the efficiency of our algorithm with that of other algorithms designed for solving similar problems. This comes in handy when choosing a policy that requires the least resource to run, given that the target performance can be reached [32].

1.4 Structure of the thesis

After familiarizing the reader with the important background concepts such as metrology, quantum technology and quantum-enhanced metrology, and also clearly stating aim and motivation in chapter 1, chapter 2 discusses the requisite background in detail.

Chapter 2 reviews foundational topics needed to understand the methods used and the results generated. The text provides the necessary material on interferometer-related and beam-splitter mathematics. It also introduces quantum-enhanced precision and guides the reader when and how it can be achieved. As this is done in the context of adaptive metrology, adaptive quantum-enhanced measurements are also introduced. Next, the chapter focuses on evolutionary algorithms and how these can be used to search non-convex spaces for finding optimal solutions. Lastly, we write about the basics of cumulants, which are indispensable in analyzing the probability distributions.

Chapter 3 explains in detail our approach that involves setting up the said problem and

then solving it in a feasible manner. We do that one step at a time. First we introduce the input state and the quantum system the state interacts with. Second we explain how the setup allows us to calculate the Cramér Rao Lower Bound and the scaling from the appropriate log-log plots. The policy is found using an evolutionary algorithm known as the differential evolution. The optimal policy assures that we achieve quantum-enhanced precision. The search for an optimal policy and the estimation process itself are run on a high-performance computer.

In chapter 4, we present results. First we show the plot of the distribution of estimates that the estimation scheme produces. This is done for two different values of beam-splitter reflectivities θ for 50 and 100 photons each. These plots show $10N^2$ estimate samples where N is the number of photons. Second we calculate the Holevo variance for select values of θ and plot the $\log(V_H)$ vs $\log(N)$ graph showing how the beam-splitter reflectivity uncertainty scales with respect to the photon number. We use the goodness of fit test to evaluate how good the linear relationship is in each case. Third, we plot the one-dimensional slices of the multi-dimensional policy landscape. We also comment on the time complexity and other performance-related quantities of the algorithm involved.

Chapter 5 contains the discussion of the results presented in the previous chapter. We start by explaining how the skewness and kurtosis exhibited by the distribution of the estimates are tolerable and therefore negligible. In order to establish that, we run well-documented tests of skewness and kurtosis. We also analyze the asymptotic power-law scaling revealed in the log-log plots in chapter 4. In our analysis, we include discussion regarding the goodness of fit and beating the standard quantum limit. We then look into the morphology of the solution landscape related to our problem and examine how suited the evolutionary algorithm is in finding an optimal policy in this space. Lastly, we inspect factors that lead to the algorithm run time that we observe in our simulations involving the adaptive quantum-enhanced metrological scheme. We also make recommendations and modifications that might end up decreasing the run time.

The last chapter summarizes what we have achieved throughout the course of the research project. It restates the aim and the claim to see if the results match up. In order to put things into context, this chapter also restates the importance our work holds in the context of quantum metrology in general. In the end, a short section is dedicated to the exploration of future research work that could use our work as a springboard. Some other ways of studying our problem are also proposed.

Chapter 2

Background

In this chapter, we review foundational topics that set the stage for more advanced topics later in the thesis. We begin by describing the mathematics behind Mach-Zehnder interferometer and the beam splitters, which are central to our work. We describe how the interferometric statistics are connected to group theory and also how the actions of the interferometer and the beam splitter can be written as matrices (operators). Second we detail quantum-enhanced precision and how to achieve it. The concept is inextricably tied to the imprecision bounds, the standard quantum limit and the Heisenberg limit. We then lay down the specifics of adaptive quantum-enhanced metrology which employs sequential adaptive measurements in order to achieve better-than-classical or quantum-enhanced precision [31]. Adaptive methodology is necessary as achieving optimal quantum-enhanced measurements without it is not manageable.

Next we focus on evolutionary algorithms that are used to find optimal policies for achieving quantum-enhanced precision. Such optimization problems usually involve non-convex solution spaces, and therefore, gradient-based algorithms don't work very well [21]. We also quote an instance where such algorithms have been employed in the past with success. Lastly, we review the essentials of cumulants that are useful in establishing whether a given distribution is Gaussian or not. This concept will come in handy when we determine the

imprecision scaling of the unknown parameter with respect to the photon number, using the estimate distribution.

2.1 Interferometer statistics

In this section, we show how variable beam-splitter estimation is a $U(1)$ estimation problem using interferometer mathematics. Using a geometrical picture, we also show what kind of rotation this is equivalent to. Lastly, we discuss the beam-splitter mathematics.

The following 2×2 unitary matrix with 4 parameters Ψ, ϕ_r, ϕ_s and θ provides a faithful representation of the $U(2)$ group [58]

$$\exp(i\Psi) \begin{pmatrix} \exp(i\phi_s) \cos(\theta/2) & \exp(i\phi_r) \sin(\theta/2) \\ -\exp(-i\phi_r) \sin(\theta/2) & \exp(-i\phi_s) \cos(\theta/2) \end{pmatrix}. \quad (2.1)$$

As we deal with Mach-Zehnder interferometry in our work, these parameters take on the following roles [58].

- Ψ : global phase
- ϕ_r : relative phase between the interferometer arms
- ϕ_s : sum phase between the interferometer arms
- θ : beam-splitter reflectivity parameter

We ignore Ψ as it is the global phase and does not affect the dynamics of the system in that there are no associated observable effects [30]. We would not have done so if we had multiple interferometers because then Ψ would become the ‘phase’ of the first interferometer relative to all the other ones in the configuration. Using the isomorphism

$$U(2)/U(1) \cong SU(2), \quad (2.2)$$

it can be seen that ignoring Ψ gives us an $SU(2)$ group that has $2^2 - 1 = 3$ parameters [57].

We also ignore the sum phase ϕ_s as this is also irrelevant in our analysis. Mathematically, this action returns a coset $SU(2)/U(1)$ [58] with matrix elements with only 2 parameters

$$\begin{pmatrix} \cos(\theta/2) & \exp(i\phi_r) \sin(\theta/2) \\ -\exp(-i\phi_r) \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \quad (2.3)$$

Previous work involves analyses with $\theta = \pi/2$ (50 : 50 beam splitter) [49, 31]. This is represented by a $U(1)$ matrix with one parameter [49]. Note that $SU(2)/U(1) \supset U(1)$ [57]. Our work involves the use of a constant relative phase value and a variable beam-splitter reflectivity value. With $\phi_r = \phi_0$, where constant $\phi_0 \in [0, 2\pi)$ and $\theta \in [0, \pi]$, our version of the $U(1)$ matrix has the following representation

$$\begin{pmatrix} \cos(\theta/2) & \exp(i\phi_0) \sin(\theta/2) \\ -\exp(-i\phi_0) \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \quad (2.4)$$

2.1.1 Bloch-sphere rotation

Previous work with phase estimation keeps $\theta = \pi/2$ and a variable phase shift [49, 31]. The interferometer transformation is represented by the unitary operator

$$\begin{aligned} & \exp(i(\pi/2)\hat{J}_x) \exp(-i\phi\hat{J}_z) \exp(-i(\pi/2)\hat{J}_x) \\ & = \exp(-i\phi\hat{J}_y), \end{aligned} \quad (2.5)$$

which when acts on a given vector executes Bloch-sphere rotations. More specifically, the operator executes a rotation of $-\pi/2$ about the abscissa, followed by a rotation of $-\phi$ about the z -axis followed by a rotation of $\pi/2$ about the abscissa. This is mathematically equivalent to a rotation of ϕ about the ordinate [53]. The quantity $\pi/2$ represents the 50 : 50 beam splitter. The goal of the phase estimation research is to estimate the angle by which the

initial state vector rotates about the ordinate [73]. In other words, the axis of rotation is known but the angle of rotation is not.

Our work involves interferometry with a constant phase and variable beam-splitter reflectivity parameters. The unitary operator here is represented by

$$\exp(i\theta\hat{J}_x)\exp(-i\phi_0\hat{J}_z)\exp(-i\theta\hat{J}_x). \quad (2.6)$$

This operator executes a rotation about the abscissa by an amount $-\theta$, followed by a rotation of $-\phi_0$ about the z -axis followed by a rotation of θ about the abscissa. This is equivalent to rotating the initial state vector by ϕ_0 about an unknown axis in the equatorial plane. This axis lies at an angle of $\theta/2$ from the abscissa. The goal is then to estimate the angle between the abscissa and the axis of rotation as shown in Fig. 2.1.




Figure 2.1: Beam-splitter reflectivity estimation involves estimating the angle between the abscissa and the axis of rotation. This angle is $\theta/2$. Initial state vector shown is rotated about the axis of rotation by an amount ϕ_0 to achieve the final state of the vector.




Figure 2.2: Beam splitter shown mixing two modes. Here a_i are the input amplitudes while b_i are the outputs.

2.1.2 Beam-splitter optics

A beam splitter's action can be described by

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \quad (2.7)$$

where B_{ij} are complex elements of the transformation matrix [53]. Here the subscripts represent port numbers. Using the fact that the outgoing modes satisfy the commutation relations [22]

$$[b_i, b_j^\dagger] = \delta_{ij}, \quad (2.8)$$

we arrive at the following expressions for beam-splitter transmission, t , and reflection, r [53]:

$$\begin{aligned} |B_{11}|^2 &= |B_{22}|^2 = t = \cos^2(\theta/2) \\ |B_{12}|^2 &= |B_{21}|^2 = r = \sin^2(\theta/2). \end{aligned} \quad (2.9)$$

Here $\theta \in [0, \pi]$. In this case, transformation matrix can be written as [22]

$$\begin{pmatrix} \cos(\theta/2) & i \sin(\theta/2) \\ i \sin(\theta/2) & \cos(\theta/2) \end{pmatrix}. \quad (2.10)$$

As can easily be seen, $r + t = 1$ in this construction. Here $r = t = 0.5$ when $\theta = \pi/2$.

The following simple example illustrates how a beam splitter works. When one photon is incident in the first mode (up) on a 50 : 50 beam splitter, the final state represents the simultaneous transmission and reflection of the original photon with equal probabilities [53]

$$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}. \quad (2.11)$$

The photon is detected with a probability of 0.5 in either output arm.

2.2 Quantum-enhanced metrology

In this section, we introduce three stages of quantum-enhanced metrology. First we explore quantum resources, more specifically, entangled particle states. Then we discuss the state's interaction with the quantum system and the measurement procedure.

Metrology is defined as the science of measurement, embracing both experimental and theoretical determinations in any field of science and technology [2]. The aim of quantum-enhanced metrology is to estimate one or more unknown parameters with high precision using quantum resources [27]. As shown in Fig. 2.3, the quantum-enhanced metrology process involves the preparation of the input state, the interaction between the input state and the system and the measurement of the output state [26]. The quantum system is the Mach-Zehnder interferometer in our work. The last step, the measurement, is used to calculate the imprecision lower bound as explained in the next section [27].

Quantum resources in this context are particles that are prepared in a non-classical state




Figure 2.3: Quantum-enhanced metrology procedure involves state preparation, interaction of the state with the Mach-Zehnder interferometer and the measurement. The direction of the arrows indicates time flow.

such as an entangled state [49]. Quantum-enhanced metrology uses quantum resources and quantum measurement to achieve quantum-enhanced precision. This is usually done by carefully preparing an input state that interacts with a system with the aim of estimating the unknown parameter ϕ [27]. The estimation procedure involves the analysis of the change that the interaction produces in the quantum state.

Quantum resources have been shown to increase the measurement precision in many applications including gravitational wave detection [35] and atomic clocks [8]. According to Ref. [49], quantum-enhanced metrology is capable of improving measurement precision without increasing the power of the input state that is used to interact with the system. The technique is useful because the increase in precision is possible even when the resources are scarce, making the entire process economical. This increase in precision is attributed to the input state that is prepared in a non-classical state such as an entangled state [48].

In order to achieve quantum-enhanced precision, the input state has to be such that it maximizes the interaction between the state and the system so that the maximum amount of information can be gathered. This is done by carefully choosing the quantum system [66]. In other words, the state has to be optimized for the task at hand. This means choosing the input state such that there is a maximum possible change in it during the interaction [49]. That maximises the amount of information gained per interaction and also the precision with which ϕ can be estimated.

2.3 Quantum imprecision

In this section, we examine how the precision of the estimation scheme is determined using the distribution of the estimates. In order to evaluate the precision, we introduce Cramér-

Rao lower bound and benchmark limits such as the standard quantum limit and Heisenberg limit.

In order to determine how accurate the quantum-enhanced metrology scheme is, some form of precision measure is used. The ϕ estimates in such a scheme are determined using the distribution of measurement outcomes [31]. More specifically, quantum-enhanced metrology generates a string of M measurement outcomes $\mathbf{x}_M = x_1 x_2 \dots, x_M$ that are not necessarily sampled from an independent and identical distribution [49]. The outcomes \mathbf{x}_M are used to estimate the unknown parameter ϕ . We call this estimate $\tilde{\phi}$.




Figure 2.4: The distribution of the estimates is shown. The actual value ϕ , the estimate mean $\tilde{\phi}$ and the estimate standard deviation $\Delta\tilde{\phi}$ are also shown. [Reproduced from Ref. [50], Fig. 2].

As shown in Fig. 2.4, the probability $P(\tilde{\phi}|\phi)$ can be plotted using multiple instances of estimating ϕ [56]. Assuming the distribution is unimodal, $\Delta\tilde{\phi}$ is the imprecision and $|\tilde{\phi} - \phi|$ is the bias.

In the case of a deterministic estimator, the estimate distribution can be written in terms of the probabilities of the measurement outcomes [49]

$$P(\tilde{\phi}|\phi) = \sum_{\mathbf{x}_M} P(\mathbf{x}_M|\phi), \quad (2.12)$$

where

$$P(\mathbf{x}_M|\phi) = \prod_{m=0}^M P(x_m|\phi). \quad (2.13)$$

Here x_m is an independently obtained measurement of an N -particle quantum state. By carefully choosing the quantum state and the measurement process, the information that can be gathered per measurement can be optimized in order to achieve quantum-enhanced precision [74]. When $M = 1$, the imprecision scales as [9]

$$\Delta\tilde{\phi} \in O(N^{-\varsigma}), \quad (2.14)$$

where $\varsigma \in \mathbb{R}^+$

2.3.1 Cramér-Rao lower bound

When a given distribution is transformed, the output becomes a different shape. Fisher information is a way of quantifying this change in shape [10]. This is tied to the measurement process in metrology schemes [26]. Every measurement changes the distribution of estimates, and Fisher information quantifies this change. More specifically, Fisher information is a metric on the space of probability distributions [10]. It is the distance between the probability distributions residing in this space. Mathematically, Fisher information is the average change in the distribution $P(\mathbf{x}_M|\phi)$ in the locality of ϕ [25]. It is a measure of how distinct \mathbf{x}_M at ϕ is from \mathbf{x}_M at $\phi + d\phi$. The bigger the Fisher information value, the more differentiation there is.

The Cramér-Rao lower bound (CRLB) defines the lower limit in the imprecision that can be calculated from a set of measurement outcomes using Fisher information [73]. The lower bound is defined as an inequality for the variance of the estimate. When the estimate is unbiased, the inequality is [25]

$$(\Delta\tilde{\phi})^2 \geq \frac{1}{F(\phi)}, \quad (2.15)$$

where $F(\phi)$ is the Fisher information

$$F(\phi) := \sum_{\mathbf{x}_M} \left(\frac{\partial}{\partial \phi} \log[P(\mathbf{x}_M|\phi)] \right)^2 P(\mathbf{x}_M|\phi). \quad (2.16)$$

In the biased case, the inequality becomes

$$(\Delta\tilde{\phi})^2 \geq \frac{1 + \frac{\partial}{\partial \phi} b(\phi)}{F(\phi)}, \quad (2.17)$$

where $b = |\tilde{\phi} - \phi|$ is the bias of the estimate [49].

For quantum resources, quantum Fisher information is used and this leads to quantum Cramér-Rao lower bound [73]. Its value depends on the input state and quantum system, and not on the measurements. In the case of phase estimation, reaching quantum Cramér-Rao lower bound is not physically feasible and therefore adaptive measurements are performed [49]. The adaptive metrological procedure approximates the optimal non-adaptive one where collective measurements are made [73].

2.3.2 Standard quantum limit and Heisenberg limit

The standard quantum limit (SQL) is a benchmark limit that can be calculated using CRLB [73]. SQL is the lower bound when non-quantum or classical resources are used. In other words, this limit is relevant when the input state is classical [27]. The simplest case is given by a tensor product of N independent particles

As this case involves an independent and identical distribution, the total Fisher information is calculated by multiplying the Fisher information for one particle by the total number of particles [49]

$$F^{(N)}(\phi) = NF^{(1)}(\phi), \quad (2.18)$$

where $F^{(1)}(\phi)$ is the Fisher information for one particle. Hence the Cramér-Rao inequality in

this case is [27]

$$(\Delta\tilde{\phi})^2 \geq \frac{1}{NF^1(\phi)}. \quad (2.19)$$

This leads to the standard quantum limit

$$\Delta\tilde{\phi} \in O\left(\frac{1}{\sqrt{N}}\right), \quad (2.20)$$

which only depends on the number of particles [31].

When quantum resources are used, SQL can be surpassed [11]. This case involves the use of entangled states. The Heisenberg Limit (HL) is lower bound when quantum resources are utilized [7]. Asymptotically, HL in interferometric phase estimation case is

$$\Delta\tilde{\phi} \in O\left(\frac{1}{N}\right). \quad (2.21)$$

Theoretically, this limit can only be obtained using optimal measurements. Therefore adaptive phase estimation techniques are used to approximate this lower bound [75].

2.4 Adaptive quantum-enhanced metrology

In this section, we detail how feedback is incorporated in quantum-enhanced metrology to achieve more precise measurements. Such a scheme is known as adaptive quantum-enhanced metrology. At the end, we give examples of previous studies where this method has successfully been employed.

Adaptive quantum-enhanced metrology or AQEM is used to approximate optimal quantum-enhanced metrology measurements [73]. This is because obtaining optimal quantum-enhanced measurements without an adaptive procedure in place is not feasible. Our work deals with discrete-time AQEM construction that involves single-photon detectors [4].

AQEM is a measurement procedure with a feedback loop incorporated in the scheme design [73]. In such a scheme, the information gained in the interaction between the input

state and the quantum system is used to estimate the unknown parameter ϕ [49]. The estimate is calculated after the detectors detect the incoming photon. As the scheme is adaptive, and ϕ does not change throughout the course of one run, the estimate gets more and more accurate after each measurement because of the amount of information available to base the estimate on increases with every detection. These sequential measurements are used to ever so slightly adjust the estimate until the desired purpose is fulfilled [75]. Here N -entangled particles are split into bundles, which are then measured sequentially. After each measurement, the next measurement instance is adjusted in order to achieve a well-defined goal.

We split the N particles into M bundles of L particles, $N = ML$, and represent the Hilbert space as [49]

$$\left(\mathcal{H}_d^{\otimes L}\right)^{\otimes M} \quad (2.22)$$

where d is the dimension of each particle. Subsequent measurements are produced after the system and the detector act on the input

$$\mathbf{x}_M = x_1 x_2 x_3 \dots, x_M, \quad (2.23)$$

where \mathbf{x}_M is the set of M measurements [31]. Here

$$x_m \in \mathbb{N}_{d^L} := 0, 1, 2, \dots, d^L - 1. \quad (2.24)$$

When there is only one particle in a bundle, $M = N$ as $L = 1$. The outcome then is described by a string of single particle detections

$$\mathbf{x}_N = x_1 x_2 x_3 \dots, x_N. \quad (2.25)$$

If additionally, each particle is a two-level system, each outcome becomes

$$x_m \in \{0, 1\}. \quad (2.26)$$

The total outcome is a N -bit string in this case $\mathbf{x}_N \in \{0, 1\}^{\otimes N}$.

There have been numerous studies on adaptive quantum-enhanced metrology establishing the efficacy of the technique [49, 31, 73]. For instance, Berry and Wiseman successfully derived the optimal N -photon two-mode input state for adaptive quantum-enhanced measurement for phase estimation [4]. They introduced a feedback algorithm based on Bayesian inference to estimate the auxiliary phase shift. By doing so they achieved the estimate variance that scales as $(\Delta\tilde{\phi})^2 \simeq \pi^2/N^2$.

During the course of his PhD degree, Hentschel developed the first computationally efficient algorithm that autonomously devises policies for adaptive phase estimation [31]. He used particle swarm optimization inspired by bird locomotion and communication in addition to machine learning to reach that goal. The algorithm can be trained with simulated or real-world experimental trial runs.

Shortly after that, Lovett et al. devised powerful algorithms based on differential evolution (DE) for adaptive many-particle quantum metrology [43]. Their approach delivered adaptive quantum metrology policies for feedback control that was orders-of-magnitude more efficient and surpassed the few-dozen-particle limitation arising in methods based on particle swarm optimization.

More recently, Palittapongarnpim used a differential evolution variant to devise robust policies for adaptive quantum-enhanced many-particle metrology [49]. The algorithm is effective even if the phase-changing quantum transformation is a black box with noise and loss present, thereby making time-consuming error modelling and extensive calibration unnecessary. She showed that the evolutionary algorithm works for up to 100 photons with different types of noise included in the system. She then compared the results to policies designed

using Bayesian inference.

2.5 Evolutionary algorithms

In this section, we consider the function and structure of evolutionary algorithms. We also look at how they perform in non-convex optimization routines and whether they are any better in these situations than gradient-based algorithms. We cite a recent study to illustrate our point.

Evolutionary algorithms are global optimization algorithms that use mechanisms inspired by biological evolution [65]. These algorithms are quite versatile as the optimization involved does not depend on the underlying model of the configuration. That is one reason why evolutionary algorithms are applied in a wide variety of engineering disciplines.

The heuristics in evolutionary algorithms follow Darwin’s theory of evolution by natural selection [21]. According to Darwin’s theory, organisms with favourable traits that help them maximize their chances of survival successfully produce more offsprings than their less well-adapted counterparts. The latter’s offsprings have traits that increase their chances of survival. This process is repeated for many generations with the organisms with the most favourable traits winning at every stage. After many generations, the favourable traits become dominant in a given population. The problem to be solved here is that of adaptation in a given environment. Biological evolution provides the solution by making favourable traits more and more common with every generation.

The algorithm’s parameters are carefully selected for it to efficiently search the solution space and converge on a globally optimal solution although the convergence is not guaranteed [21]. In many evolutionary algorithms, the optimization uses only the output of the fitness function to proceed [24].

One disadvantage of the black-box, data-driven approach is that the relevant evolutionary algorithms evaluate the policy as a whole which is computationally expensive as the




Figure 2.5: A depiction of optimization landscapes. (a) The image on the left shows a convex set while the one on the right shows a non-convex one. The dotted line represents the shortest path between the two points \mathbf{x} and \mathbf{y} . (b) For each pair (k, β) , the sharpness S (Eq. 3.14) of the policy parametrized by (k, β) is shown. The Bayesian inference-based policy is represented by $k = 1.596$ and $\beta = 1.096$ (marked by \bullet). For comparison, the global maximum is at $k_{\max} = 1.631$, $\beta_{\max} = 1.015$ (marked by the arrowhead). [Reproduced from Ref. [31], Fig. 5.1].

optimization needs a large number of trials to find a successful solution [61].

2.5.1 Evolutionary algorithms for non-convex optimization

Non-convex functions and non-convex sets play important roles in non-convex optimization procedures [40].

Definition 2.1. (Convex Set) A set $\mathcal{C} \in \mathbb{R}^p$ is convex if, for every $\mathbf{x}, \mathbf{y} \in \mathcal{C}$ and $\lambda \in [0, 1]$, we have $(1 - \lambda)\mathbf{x} + \lambda\mathbf{y} \in \mathcal{C}$ as well [14].

Definition 2.2. (Convex Function) A continuously differentiable function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ is convex if for every $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$ we have $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle$ where $\nabla f(\mathbf{x})$ is the gradient of f at \mathbf{x} [14].

Fig. 2.5 (a) shows the difference between convex and non-convex sets pictorially. If the ‘shortest path’ between any two points \mathbf{x} and \mathbf{y} in the set $\mathcal{C} \subseteq \mathbb{R}^d$ is not fully contained in the set, then the set in question is non-convex [40]. An optimization problem is said to be convex if the fitness function is a convex function and the constraint set is a convex set [14].

An optimization problem that violates either of these conditions is a non-convex optimization problem. Problems that usually involve high-dimensional data come with structural constraints for tractability purposes [49]. These and other features of problems with high-dimensional data usually make the optimization problem non-convex in nature, which are hard to solve [40].

Unlike gradient-based optimization algorithms, evolutionary algorithms are global optimization algorithms and therefore perform better in non-convex spaces [47]. Gradient-based greedy algorithms have a tendency to get stuck in one of the several local minima that are part of the morphology of non-convex spaces. Evolutionary algorithms are better candidates to perform optimization in non-convex spaces [40]. We employ differential evolution in our work.

2.5.2 Differential evolution

Differential evolution is a global optimization algorithm that iteratively searches for a feasible solution in a continuous search space [59]. The search process uses a heuristic inspired by biological evolution. DE algorithm uses a unique set of rules for generating and selecting solution candidates, which is summarized below.

- **Initialization:** The optimization process begins with randomly initializing a candidate vector for generation G , $\mathbf{Y}_i(G)$ where $i \in \{1, \dots, N_P\}$ [78]. These random candidates are uniformly generated. Each candidate exists in an N -dimensional search space

$$\mathbf{Y}_i(G) = (Y_i(G)^{(1)}, Y_i(G)^{(2)}, \dots, Y_i(G)^{(N)}). \quad (2.27)$$

Each of these candidates is evaluated for its fitness using a fitness function or sharpness, $S(\mathbf{Y}_i(G))$ [34]. This quantifies how feasible the solution candidates are.

- **Mutation:** For each candidate, an offspring is generated by combining the parent with a random member of the population set, $\mathbf{V}_i(G) = (V_i(G)^{(1)}, V_i(G)^{(2)}, \dots, V_i(G)^{(N)})$ [49].

The mutation occurs according to the following rule:

$$V_i(G)^{(j)} = Y_{i,1}(G)^{(j)} + \mathcal{F}(Y_{i,2}(G)^{(j)} - Y_{i,3}(G)^{(j)}). \quad (2.28)$$

Here \mathcal{F} is the mutation factor controlling the weight of $Y_{i,2}(G)^{(j)} - Y_{i,3}(G)^{(j)}$ in the optimization process.

- **Crossover:**

$$C_i(G)^{(j)} = \begin{cases} V_i(G)^{(j)}, & \text{if } \text{rand}(0, 1) \leq C_r. \\ Y_i(G)^{(j)}, & \text{otherwise} \end{cases} \quad (2.29)$$

where $\text{rand}(0, 1)$ is a random number between 0 and 1 [59]. Also, $C_r \in (0, 1)$ is the crossover rate.

- **Selection:** The last step is the following condition [78]:

$$Y'_i(G) = \begin{cases} C_i(G), & \text{if } S(C_i(G)) > S(Y_i(G)). \\ Y_i(G), & \text{otherwise.} \end{cases} \quad (2.30)$$

Here $Y'_i(G)$ is the offspring of the parent candidate $Y_i(G)$ [59]. The quantity $S(C_i(G))$ is the sharpness function (of candidate $C_i(G)$ in this case), which we maximise so as to minimize the imprecision of the AQEM scheme.

Various approaches to study AQEM use Bayesian methods [4], particle swarm optimization [31] and DE [43, 49] which are all global optimisation techniques. Global optimization methods are employed in the AQEM simulations involving non-convex discrete optimization as shown in Fig. 2.5 (b). In this figure, the global max., represented by the arrowhead is located at $k_{\max} = 1.631$, $\beta_{\max} = 1.015$ while the solution found by the policy based on Bayesian methods is located at $k = 1.596$ and $\beta = 1.096$ (marked by \bullet). The figure was reproduced from Hentschel's work [31] where the goal was to find the global maximum in a non-convex landscape. As can be seen, the employed global-optimization algorithm was

quite successful in finding the solution.

2.6 Cumulants

In this section, we introduce moments via moment generating functions. We then establish their connection with cumulants. We discuss skewness and kurtosis in good detail thereafter.

2.6.1 Moment-generating function

The k -th moment of a real-valued random variable X with density $f(x)$ is [29]

$$\mu_k = E(X^k) = \int_{-\infty}^{\infty} x^k f(x) dx \quad (2.31)$$

for integer $k = 0, 1, \dots$. Here μ_k has finite values.

The moment generating function of X can then be defined as [29]

$$\begin{aligned} M_X(t) &= E(\exp(tX)) \\ &= E(1 + tX + t^2 X^2/2! + \dots) \\ &= \sum_k^{\infty} \mu_k t^k / k!. \end{aligned} \quad (2.32)$$

Here $M_X(t)$ is finite for all $t \in [-\epsilon, \epsilon]$ for some $\epsilon > 0$. Some of the properties of the moment generating function are given below [13]

- Every moment of X is finite.
- $M_X(t)$ has convergent power series expansion in the neighbourhood of each $t \in (-\epsilon, \epsilon)$
- The k -th moment is the k -th derivative of M at the origin. In other words,

$$\mu_k = \frac{d^k}{dt^k} M_X(0). \quad (2.33)$$

2.6.2 Definition and nomenclature

The cumulant generating function of the random variable X is the log of the moment generating function [29]

$$\begin{aligned} K_X(t) &= \log(M_X(t)) \\ &= \sum_k^{\infty} \kappa_k t^k / k!, \end{aligned} \tag{2.34}$$

where κ_k are the cumulants. These are the coefficients in the Taylor expansion of the cumulant generating function about the origin. The relationships between the moments and the cumulants are obtained by comparing the coefficients in the expansions of the moment and the cumulant generating functions [13].

Some important relationships are listed [29]

$$\begin{aligned} \kappa_1 &= \mu_1 = \mu \\ \kappa_2 &= \mu_2 - \mu_1^2 = \sigma^2 \\ \kappa_3 &= \mu_3 - 3\mu_2\mu_1 + 2\mu_1^3 = E[(X - \mu)^3] \\ \kappa_4 &= \mu_4 - 4\mu_3\mu_1 - 3\mu_2^2 + 12\mu_2\mu_1^2 - 6\mu_1^4 = E[(X - \mu)^4] - 3\kappa_2^2. \end{aligned} \tag{2.35}$$

Notice, $\mu_0 = 1$ and so $\kappa_0 = 0$. Conventionally, κ_1 is the mean and κ_2 is the variance. κ_l for $l > 3$ are not moments about the mean.

2.6.3 Skewness and kurtosis

Skewness is a measure of the asymmetry of the probability distribution of a random variable X [41]. Skewness measures the asymmetry of the distribution from its mean. If the distribution is not symmetric it is said to be skewed.

For univariate data X_1, X_2, \dots, X_N , the Pearson's moment coefficient of skewness is de-

defined as [13]

$$\gamma_1 = \frac{\kappa_3}{\sigma^3} = \frac{\kappa_3}{\kappa_2^{3/2}} = \frac{E[(X - \mu)^3]}{(E[(X - \mu)^2])^{3/2}}. \quad (2.36)$$

Skewness for Gaussian distribution is zero because it is symmetrical. A distribution that ‘skews’ to the left is called positively skewed and the one that ‘skews’ to the right is negatively skewed.

In order to test whether a given distribution is significantly skewed or not, the skewness test is used [34]. If there is not enough evidence of skewness, the distribution in question has negligible skewness and can be thought of as being approximately Gaussian.

Kurtosis is a measure of whether a given distribution is light-tailed or heavy-tailed relative to the Gaussian distribution [41]. The greater the number of outliers of a given distribution relative to the Gaussian distribution the higher is its kurtosis value. Kurtosis is defined as

$$\text{Kurt}(X) = \frac{\mu_4}{\kappa_2^2} = \frac{E[(X - \mu)^4]}{(E[(X - \mu)^2])^2}. \quad (2.37)$$

It is common practice to use excess kurtosis [29]

$$\text{Excess Kurtosis} = \text{Kurt}(X) - 3. \quad (2.38)$$

Distributions with zero excess kurtoses are called mesokurtic [79]. Gaussian distribution is an example. Distributions with positive excess kurtosis are called leptokurtic. These have flat tails relative to the Gaussian distribution. Platykurtic ones are distributions with negative excess kurtosis and have flatter peaks than Gaussian distribution.

Kurtosis test is used to see if a given distribution has statistically significant kurtosis associated with it [34]. In real-life examples, the probability distributions hardly ever have tails like the ones in Gaussian distribution, and so they do have non-zero excess kurtosis. If kurtosis is within tolerable limits, then the sample distribution is approximately Gaussian. In this case, all the analytical and mathematical tools related to Gaussian distribution can

be used with good accuracy.

Chapter 3

Adaptive variable beam-splitter reflectivity estimation

In this chapter, we describe our approach for solving the discrete-time adaptive variable beam-splitter reflectivity estimation problem. We start by introducing the input state used for AQEM. We then describe the interferometer used in the simulation and the process itself. This allows us to calculate CRLB given we have access to an optimal policy, which is difficult to calculate analytically. The imprecision is calculated using the distribution of the estimates and the scaling is determined using the log-log plots of Holevo variance vs the number of particles N .

This chapter also explains how an evolutionary algorithm, differential evolution, is used to find optimal policies to solve the estimation problem in question while achieving quantum-enhanced precision. The several parameters of the algorithm are carefully selected to achieve that goal. As the simulation is run on a high-performance computer, we also illustrate how the code is implemented on the relevant hardware devices.

3.1 Input state and the interferometer

In this section, we introduce the permutationally symmetric state we use for the adaptive variable beam-splitter reflectivity estimation procedure. We also list down the properties of the input state in detail. We then explain the specifics of the interferometer this state interacts with.

3.1.1 Input state




Figure 3.1: A plot of the sine state for $N = 50$. The red curve shows the fitted Gaussian curve. [Reproduced after modification from Ref. [31], Fig. 2.2].

In our setup, the interferometric modes are represented as follows [49]:

$$\epsilon_m \in \{0, 1\}. \quad (3.1)$$

Here 0 and 1 represent the two paths in an interferometer and also the input and the output ports. Using this convention, we can write the multi-photon basis as the tensor-product state

$$|\epsilon_N\rangle = \bigotimes_{m=1}^N |\epsilon_m\rangle, \quad (3.2)$$

where $|\epsilon_m\rangle$ is the state of m^{th} photon in path ϵ_m [49]. For ham ϵ the Hamming weight (or

the sum of bits), the permutationally symmetric basis is

$$|n, N_a - n\rangle = \binom{N_a}{n}^{-1/2} \sum_{\text{ham } \epsilon_{N_a}} |\epsilon_{N_a}\rangle, \quad (3.3)$$

where N_a is the total number of particles in mode a and $n \in \{0, N_a\}$ [33].

In 1997, Wiseman and Killip considered a Mach-Zehnder interferometer with a canonical phase measurement at the output and derived the input state that yields maximum precision [76]. This state will henceforth be referred to as the sine state, which is an N -photon entangled state given here in the basis of the two input modes of a Mach-Zehnder interferometer

$$|\psi\rangle_N = \frac{1}{\sqrt{\frac{N}{2} + 1}} \sum_{n,k=0}^N \sin\left(\frac{k+1}{N+2}\pi\right) e^{i\pi(k-n)/2} d_{n-N/2, k-N/2}^{N/2}\left(\frac{\pi}{2}\right) |n, N-n\rangle. \quad (3.4)$$

Here n is the number of the photons being injected into the first port of the interferometer at any given moment while N is the total number of photons used in the simulation. Moreover, $d_{m,m'}^j(\beta)$ is the Wigner- d function where $\beta \in \mathbb{R}$ [50].

Figure 3.1 shows what the sine state looks like. Sine state is permutationally symmetric that gives it the following interesting properties:

- The order in which the photons are injected into the interferometric ports does not matter [76].
- It is a loss-tolerant input state in that the state remains permutationally symmetric even when some of the photons are lost [60].
- The canonical phase measurement with the sine state achieves a precision that scales as $\Delta\tilde{\phi} \sim 1/N$ for large N . Therefore the state can carry enough information about the unknown phase to reach the Heisenberg limit [31].

We employ the sine state in a variable beam-splitter reflectivity estimation problem. This

is because the sine state has successfully achieved quantum-enhanced precision in several metrological problems before [49, 31, 60, 76, 74], and hence, it seems to be a reasonable starting point for investigating the variable beam-splitter reflectivity estimation.

3.1.2 The interferometer transformation

The action of a beam splitter is described by

$$B = \exp(-i\theta\hat{J}_x) \quad (3.5)$$

where \hat{J}_x is the spin operator written in terms of the Pauli x matrix $\hat{J}_x = \hbar\hat{\sigma}_x/2$ [31]. For a 50 : 50 beam splitter, $\theta = \pi/2$.

The action of a phase shifter is described by

$$P = \exp(-i\phi\hat{J}_z), \quad (3.6)$$

where \hat{J}_z is the z spin operator written in terms of the Pauli z matrix $\hat{J}_z = \hbar\hat{\sigma}_z/2$ [31].

In this thesis, we keep the phase shift constant, $\phi_0 \in [0, 2\pi)$ and vary the two beam-splitter reflectivities. The first beam splitter has the unknown reflectivity $\theta \in [0, \pi]$ while the second beam splitter's reflectivity is parameterized by $\Theta \in [0, \pi]$. Therefore the unitary operator corresponding to the variable beam-splitter reflectivity Mach-Zehnder interferometer transformation, for one photon is

$$U_1(\theta; \Theta_m) = \exp(i\Theta_m\hat{J}_x) \exp(-i\phi_0\hat{J}_z) \exp(-i\theta\hat{J}_x). \quad (3.7)$$

Here Θ_m is the state after m^{th} photon is measured.

Therefore the unitary operator for N photons becomes

$$U(\theta; \Theta_m) = U_1(\theta; \Theta_m) \otimes \cdots \otimes \mathbb{1}^{(N)}. \quad (3.8)$$

The operator is valid when there is no noise present in the system.

Using the identity, $\exp(i\omega\hat{\sigma}_\kappa) = \mathbb{1} \cos(\omega) + i\hat{\sigma}_\kappa \sin(\omega)$ [31], $U_1(\theta; \Theta_m)$ can be written as a matrix operator:

$$U_1(\theta; \Theta_m) = \begin{bmatrix} \cos\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta - \Theta_m)\right) - i \sin\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta + \Theta_m)\right) \\ \sin\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta + \Theta_m)\right) - i \cos\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta - \Theta_m)\right) \\ - \sin\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta + \Theta_m)\right) - i \cos\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta - \Theta_m)\right) \\ \cos\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta - \Theta_m)\right) + i \sin\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta + \Theta_m)\right) \end{bmatrix}. \quad (3.9)$$

If a and b are the probability amplitudes of an incoming photon, as shown in Fig. 1.1, the action of this operator produces the output

$$\begin{bmatrix} a \left(\cos\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta - \Theta_m)\right) - i \sin\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta + \Theta_m)\right) \right) \\ b \left(\sin\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta + \Theta_m)\right) - i \cos\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta - \Theta_m)\right) \right) \\ a \left(-\sin\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta + \Theta_m)\right) - i \cos\left(\frac{\phi_0}{2}\right) \sin\left(\frac{1}{2}(\theta - \Theta_m)\right) \right) \\ b \left(\cos\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta - \Theta_m)\right) + i \sin\left(\frac{\phi_0}{2}\right) \cos\left(\frac{1}{2}(\theta + \Theta_m)\right) \right) \end{bmatrix}. \quad (3.10)$$

3.2 Adaptive measurement procedure

In this section, we lay out how the feedback is built into our approach towards solving the estimation problem. We specifically detail adaptive measurements, which are central to our scheme. We also illustrate how the adaptation procedure leads to a binary decision tree with depth equal to the number of photons N .

3.2.1 Adaptive mechanism

Fig. 3.2 illustrates a single-shot quantum-enhanced adaptive measurement procedure. The configuration shown in the diagram is a Mach-Zehnder interferometer with an unknown parameter θ and the control Θ . In this particular case, N d -level particles are divided

equally into M bundles of L particles. N particles are used so as to increase the information that can be obtained from the measurement process [9]. Each bundle M interacts with the interferometer and produces an outcome x_i . Therefore after m bundles have passed through, an outcome $x_m = 1, \dots, d^L$ is used to adjust the value of Θ_m (given Θ_0 is some constant initial value) as per the feedback policy ϱ and the history of outcomes leading up to the m^{th} measurement, $\mathbf{x}_m = x_1 x_2 \dots, x_m$ [50]. Once all bundles have passed through the interferometer, assuming no loss of photons, the estimate $\tilde{\theta}$ is calculated from $\Theta_M(\mathbf{x}_M)$.




Figure 3.2: Adaptive quantum-enhanced metrology procedure showing that the input state is divided into M bundles of L particles, and each bundle then interacts with the interferometer in a sequential manner. The processing unit (PU) determines the control parameter Θ_m taking into account the history of the measurement outcomes up to that point, $\mathbf{x}_m = x_1 x_2 \dots, x_m$. [Reproduced after modification from Ref. [50], Fig. 1].

The scheme is such that initially, an unknown beam-splitter reflectivity value θ is uniformly randomly generated. Then starting with the prior $\Theta_0 \equiv \pi/2$, Θ_m is calculated and updated after every bundle measurement. This results in the estimate getting better and closer to the unknown θ with every iteration. It is only after all the measurements have been made that the value of the estimate $\tilde{\theta} = \Theta_M(\mathbf{x}_M)$ is inferred. $\tilde{\theta}$ is a discrete approximation of a continuous parameter $\theta \in [0, \pi]$ [49].

In order to make the entire process more efficient and less time and computation consuming, we run simulations with several θ values to see if there is an underlying pattern. If we see one, we adjust the algorithm by utilizing the pattern observed such that there is less time consumed in converging towards the feasible solution.

Our scheme applies in the case of a Michelson interferometer with a constant beam-splitter reflectivity. The scheme however is not relevant in the case of a Fabry-Perot etalon as it fundamentally has a different design.

3.2.2 Decision making

A successful policy is the one that minimizes the difference between θ and the estimate $\tilde{\theta}$. In other words, a policy that maximizes the amount of information that can be gained about θ after successive measurements [75]. The function $\Theta_M(\mathbf{x}_M)$ calculates the value of \mathbf{x}_M based on all the measurement outcomes. This property makes this AQEM procedure a decision-making process [71].

Assuming the feedback is deterministic, each outcome has the following probability of occurring [49]:

$$P(\mathbf{x}_M|\theta, \varrho) = \prod_{m=1}^M P(\mathbf{x}_m|\theta, \Theta_{m-1}). \quad (3.11)$$

A decision tree represents the policy and provides the rule for mapping $\Theta_{m-1} \mapsto \Theta_m$ [28]. The size of the decision tree can be calculated by noticing that there are M measurements in a single shot procedure, each with d^L possible outcomes. The number of branches is given by $\sum_{m=1}^M (d^L)^m = d^L \frac{d^{NL} - 1}{d^L - 1}$ [49]. As can be seen, the maximum size of a non-Markovian policy scales with N exponentially. This makes it computationally expensive to implement and scale. Our work involves $L = 1$ and therefore, the sum reduces to

$$\sum_{m=1}^N 2^m = 2^{N+1} - 2. \quad (3.12)$$

The following Markovian updating rule is used in our work [50]

$$\Theta_m = \Theta_{m-1} - (-1)^{x_m} \Delta_m. \quad (3.13)$$

As a result of implementing this updating rule, the policy $\varrho = (\Delta_1, \Delta_2, \dots, \Delta_m)$ scales




Figure 3.3: An example of a binary decision tree with three particles. The red line represents one run. The initial controllable beam-splitter reflectivity is Θ_0 and is adjusted by $\pm\Delta_m$ depending on $x_m \in \{0, 1\}$. [Reproduced after modification from Ref. [49], Fig. 4.3].

linearly with N , and is hence, tractable. The policy is the procedure carried out by the processing unit. The procedure involves adjusting the controllable beam splitter's reflectivity as the photons are detected. With linear scaling in place, the relevant decision tree is a binary one [71] as shown in Fig. 3.3. For any given policy ϱ , there are $2^N - 1$ possible paths corresponding to every possible sequence of \mathbf{x}_N . As the nature of quantum measurement is probabilistic, the outcome \mathbf{x}_N may not necessarily be unique [49]. The uncertain nature of the measurement procedure also means that even when the optimal policy is used, the estimate $\tilde{\theta}$ may not be a close approximation of θ .

3.3 Imprecision and scaling

In this section, we outline the process of calculating sharpness and Holevo variance from the distribution of the beam-splitter reflectivity estimates. We also detail how this information is used to plot the log-log graphs and extract the scaling.

The imprecision of the measurement scheme is quantified by calculating the width of the distribution of the estimate $P(\tilde{\theta}|\theta, \varrho)$ [17]. Using the quantity $\Theta_M(\mathbf{x}_M)$, the distribution of the outcomes, $P(\mathbf{x}_M|\theta, \varrho)$, can be determined [49]. The goal is to then minimize the imprecision and hence increase the accuracy of the scheme by finding the optimized probabilities

for the branches of the binary decision tree.




Figure 3.4: The distribution of the estimates. The actual value θ , the estimate mean $\bar{\tilde{\theta}}$ and the estimate standard deviation $\Delta\tilde{\theta}$ are also shown. [Reproduced after modification from Ref. [50], Fig. 2].

Fig. 3.4 shows the sampled data associated with $P(\mathbf{x}_M|\theta, \varrho)$. The distribution peaks at $\bar{\tilde{\theta}}$ assuming the distribution is Gaussian. For an unbiased scheme, $\bar{\tilde{\theta}} = \theta$ but this is not true for small number of particles because the estimates have discrete values [42]. It is important to collect a large amount of data to have a good estimate. The bias is determined by $|\theta - \bar{\tilde{\theta}}|$.

The imprecision is defined as the spread of the distribution around $\bar{\tilde{\theta}}$ regardless of whether the estimate is biased or unbiased [50]. The imprecision is minimized by optimizing the distribution $P(\mathbf{x}_M|\theta, \varrho)$. As it is, computationally speaking, a hard problem to solve, generating a feasible policy is a challenge for devising a viable adaptive quantum-enhanced metrology scheme [31].

Pinning down a successful policy is all about obtaining one that minimizes the imprecision or more accurately, the average imprecision over all possible values of θ [50]. Although θ ranges continuously from 0 to π (inclusive), there is only one θ value being estimated at one time. This allows the algorithm to take discrete samples for calculating imprecision later on [49]. This makes the computation tractable.




Figure 3.5: Accuracy and precision of the probability distribution are shown. The target value and the mean of the estimate distribution play an important role in determining these quantities.

3.3.1 Sharpness and Holevo variance

The scheme is trained using a set of discretely sampled θ from the domain $[0, \pi]$. The simulation is run $K = 10N^2$ times for different θ and the results are used to calculate sharpness [31]

$$S = \left| \sum_{k=1}^K \frac{\exp[i(\theta_k - \tilde{\theta}_k)]}{K} \right|. \quad (3.14)$$

We use different values of θ so that the policy does not overfit the data and the policy delivers similar imprecision when applied to the value of θ not included in the training set [49]. Thereafter, we test the trained model using specific θ values to evaluate the performance of the whole scheme.

Fig. 3.5 illustrates what accuracy and precision mean. If the estimate values $\tilde{\theta}$ are grouped close together, then the estimate is precise [3]. If, in addition, the estimates are close to the actual reflectivity value θ , then the data are also accurate. If there is a good deal of scatter in the estimates then the data are not precise and if the values are not close to the actual value θ , then it is not accurate as well. These concepts come in handy when sharpness is

written in terms of a power series

$$\begin{aligned}
S &= \frac{1}{K} \left| \sum_{k=1}^K \left(1 + i(\theta_k - \tilde{\theta}_k) + (i(\theta_k - \tilde{\theta}_k))^2/2 + (i(\theta_k - \tilde{\theta}_k))^3/6 + \dots \right) \right| \\
&= \left| \sum_{k=1}^K 1/K + \sum_{k=1}^K i(\theta_k - \tilde{\theta}_k)/K + \sum_{k=1}^K (i(\theta_k - \tilde{\theta}_k))^2/(2K) + \sum_{k=1}^K (i(\theta_k - \tilde{\theta}_k))^3/(6K) + \dots \right| \\
&= \left| \frac{K+1}{2} + \langle i(\theta_k - \tilde{\theta}_k) \rangle + \langle (i(\theta_k - \tilde{\theta}_k))^2 \rangle/2 + \langle (i(\theta_k - \tilde{\theta}_k))^3 \rangle/6 + \dots \right|. \tag{3.15}
\end{aligned}$$

Here the first term is the accuracy while the second term is precision when $\theta \approx \tilde{\theta}$ [3]. In other words the second term is precision (or variance) if the accuracy of the estimate $\tilde{\theta}$ is high. In such a case higher order terms can be ignored and sharpness takes the form

$$S = \left| \frac{K+1}{2} + \langle i(\theta_k - \tilde{\theta}_k) \rangle + \langle (i(\theta_k - \tilde{\theta}_k))^2 \rangle/2 + O(\langle (i(\theta_k - \tilde{\theta}_k))^3 \rangle) \right|. \tag{3.16}$$

The cubed term in sharpness is skewness when $\theta \approx \tilde{\theta}$. This term can be ignored if the distribution is symmetric. We intend to keep as many terms as practically possible while evaluating a policy so as to get an accurate Holevo variance measure, which is defined by [76]

$$V_H := S^{-2} - 1. \tag{3.17}$$

This leads to $(\Delta\tilde{\theta})^2 = S^{-2} - 1$ [49].

3.3.2 Asymptotic power-law scaling

In order to determine how the beam-splitter reflectivity estimate imprecision scales with respect to the number of photons, we plot the $\log(V_H)$ vs $\log(N)$ and fit a least-squares regression line. The gradient of the line is the scaling -2ς . This can be seen by taking log on both sides of the equation $V_H \propto N^{-2\varsigma}$. With W being the constant of proportionality, we

get

$$\log(V_H) = -2\zeta \log(N) + \log(W). \quad (3.18)$$

For $N \leq 93$, we terminate the policy-search process after a set number of generations as it delivers the precision we look for. For $N \in \{94, 100\}$, we only accept a policy if the V_H is within a distance corresponding to a confidence interval of 0.98 from the inverse power-law line [49]. This is because the policy search becomes more computationally expensive as the dimensionality of the problem increases, and therefore, we need to include an extra measure to make sure we get the requisite inverse power-law line.

In order to make sure we get the quantum-enhanced precision for $N > 93$, we use the linear equation and predict the next data point y' [31]. Then the acceptable error δ_y from y' of the next empirical data point $y = \log(V_H)$ is calculated as using the formula [77]

$$\delta_y = t_{v-2}^* \sqrt{\frac{\sum_{i=1}^v (y'_i - y_i)^2}{v-2} \left(\frac{1}{v} + \frac{(x' - \bar{x})^2}{\sum_{i=1}^v (x_i - \bar{x})^2} \right)}. \quad (3.19)$$

Here t_{v-2}^* is the quantile of the Student's t distribution for $v-2$ data points, v is the number of data points, $x' = \log(N)$ and \bar{x} is the average of all x values. The policy is not accepted until $|\log(V_H) - y'| \leq \delta_y$ [49].

After the training phase is complete, we test the scheme for uniformly randomly generated 49 θ values and plot the power-law scaling.

3.4 Policy search and implementation

In this section, we outline how the differential evolution algorithm works by identifying all the steps necessary for its successful execution. The algorithm's job is to search the solution space and find the optimal policy for variable beam-splitter reflectivity estimation. The policy implements the estimation procedure such that quantum-enhanced precision is achieved. In the end, we detail how the code is implemented on Compute Canada clusters

with necessary parameters in place.

3.4.1 Policy search using differential evolution

The policy-search approach in our case is the direct policy search [16]. This approach involves waiting until the end of the task to determine ϱ . In direct policy search, the policy is parametrized and an optimization algorithm is used to search for the parameters for successful execution. One of the drawbacks of this heavy ‘customization’ is that when the policy space is large, the algorithm takes a long time to find a feasible solution [14]. That is why we use a Markovian approach where the size of the policy grows linearly with N [49]. $\tilde{\theta}$ in our approach is

$$\tilde{\theta} = \Theta_0 + \sum_{m=1}^N (-1)^{x_m} \Delta_m, \quad (3.20)$$

where Δ_m is the m -th adjustment that the processing unit makes in the case of Θ_{m-1} in the light of new information, i.e. the m -th measurement.

The policy search for optimal solutions is achieved by implementing the differential evolution (DE) algorithm, which is a black-box heuristic optimization algorithm [72]. Differential evolution uses the sharpness as the fitness function, and not the model-related details to search the policy space. In our case, the fitness is computed from multiple trials of the same policy, which enables us to optimize independently from the AQEM dynamics [5]. The DE algorithm is used to find the optimal $\varrho = (\Delta_1, \Delta_2, \dots, \Delta_N)$ vector, which is a collection of controlled adjustments. As the DE algorithm implemented may not always find the global maximum [72], we look for feasible solutions or solutions that are ‘good enough’ for our purpose. In that sense, we deal with is a feasibility variant of the optimization problem.

We use DE because evolutionary algorithms perform better in non-convex spaces [59]. We prove the non-convexity of the optimization space using a multi-step process. We first plot the one-dimensional slices of the landscape and show these are non-convex by analyzing the structure in the plots. We then reason if the one-dimensional slices are non-convex, the

multi-dimensional landscape has to be non-convex too [46].

3.4.2 Implementation on a high-performance computer

The estimation algorithm is designed in such a way that the task of finding the fitness values can be parallelized. This is because due to a great number of calculations, the estimation procedure is computation intensive. Parallelization makes the process more tractable. Algorithm 1 in the algorithm window shows the pseudocode for AQEM implementation.

In order to parallelize the DE candidates over the cluster machines, we utilize the built-in Message Passing Interface (MPI) library in C++ [51]. This library is instrumental in parallelizing the calculation of the fitness function across several cluster CPUs. MPI library specifies the communication between the CPUs and the memories during the compilation of the code automatically and optimally [38], and so the same code can be implemented on any cluster that supports MPI.

We distribute the computation tasks such that there is one candidate per CPU so that the CPU – CPU communication takes negligible time compared with the time each CPU takes to compute the sharpness function. It was observed that the random number generation used in calculating the Holevo variance and the adaptive quantum-enhanced measurement procedure takes most of the time in a simulated run [49]. Generating random numbers in real time is not efficient and practical on modern hardware, and so these are stored prior to the runs to be used by candidates later on. In order to achieve that, we either use Intel Vector Statistical Library (VSL) that relies on the CPU or the one that relies on the GPU for its performance [51]. GPU-oriented random number generation is more scalable as it turns out.

We implement the computational work on Compute Canada clusters [36]. More specifically, we use clusters named Graham and Cedar for the duration of the research project. Graham is located at the University of Waterloo while Cedar is located at Simon Fraser University. Both of the clusters are heterogeneous in nature suitable for a variety of workloads.

Algorithm 1 Adaptive variable beam-splitter reflectivity estimation

Input:

UNSIGNED INTEGER N ▷ Number of particles
COMPLEX[] $|\psi_N\rangle$ ▷ Input quantum state
REAL[] ϱ_N ▷ Policy
REAL θ_0 ▷ unknown beam-splitter reflectivity $\in [0, \pi]$
REAL Θ_0 ▷ controllable beam-splitter reflectivity (initial) $\in [0, \pi]$
REAL ϕ_0 ▷ constant phase shift $\in [0, 2\pi)$

Output:

REAL $\tilde{\theta}$ ▷ beam-splitter reflectivity estimate $\in [0, \pi]$

1: $\phi_0 \leftarrow \frac{7\pi}{11}$ ▷ this can be modified when needed
2: $\Theta_0 \leftarrow \frac{\pi}{2}$
3: UNSIGNED INTEGER x
4: REAL Δ
5: **for** $m = 0$ to $N - 1$ **do**
6: $\theta \leftarrow \text{RandomNumber}(\theta_0)$
7: $x_{m+1}, |\psi_{N-m-1}\rangle \leftarrow \text{Measure}(|\psi_{N-m}\rangle, \theta, \Theta_m)$ ▷ $x_m \in \{0, 1\}$
8: **if** $x_{m+1} = 0$ **then**
9: $\Theta_{m+1} \leftarrow \Theta_m - \Delta_{m+1}$ ▷ $\Delta_m \in [0, \pi]$
10: **else**
11: $\Theta_{m+1} \leftarrow \Theta_m + \Delta_{m+1}$
12: **end if**
13: **end for**
14: **return** Θ_N

3.5 Algorithm complexity

In this section, we determine the complexity of computing and implementing the variable beam-splitter reflectivity estimation policy using loop analysis [67].

3.5.1 Policy time cost

The policy time cost is the upper bound on the time cost associated with generating a policy using differential evolution [49]. For the following calculations, we assume that only a single processor is used for computing the policy. This is so that we get the upper bound on the processing time. Parallelizing the policy-search process on multiple CPUs speeds up the process. We also assume that there is no accept-reject criterion in place and that the maximum number of particles is 93. This is because the accept-reject criterion kicks in when $N > 93$, and the time requirements not only change drastically but also become unpredictable past that threshold.

The DE is employed to search for the N -particle policy using the $(N - 1)$ -particle policy as the initial guess [49]. The following DE-related computations contribute to the time cost:

- The outermost loop in the algorithm involves the number of generations Γ before a policy is accepted. This improves the policy candidates iteratively, and ensures that $\Gamma \in O(N)$ [43].
- As the size of the DE population is N_p , there is a loop of this size for every generation [49]. We fix $N_p = N_{\max}/2$ that ensures $N_p \in O(N)$ [31]. For instance for $N = 100$, we have $N_p = 50$.
- Time to execute a fitness function is Ω_{fit} . For each instance of adaptive phase estimation, N particles sequentially pass through the interferometer and are measured. The input state of size $(N - m)$ is updated after every m -th measurement. The updating slows the estimation procedure by a constant factor of η_1 [49]. As the training phase

has $K = 10N^2$ instances of the estimation procedure, Ω_{fit} can be written as [43]

$$\begin{aligned}
\Omega_{\text{fit}} &= \sum_{l=1}^{10N^2} \sum_{m=1}^N \eta_1(N-m) \\
&= \sum_{l=1}^{10N^2} \frac{\eta_1}{2}(N^2 - N) \\
&= (10N^2) \frac{\eta_1}{2}(N^2 - N) \\
&= 5\eta_1(N^4 - N^3).
\end{aligned} \tag{3.21}$$

- Time to generate an offspring is Ω_{off} . Creation of an offspring involves randomization of candidates, with a constant time cost of η_2 , and generation with a time cost of $\eta_3 N$ [49]. Moreover, η_3 is a constant while N indicates the presence of a loop over N elements

$$\Omega_{\text{off}} = \eta_2 + \eta_3 N + \frac{\eta_1}{2}(N^4 - N^3). \tag{3.22}$$

- Time to select a member for the next generation is Ω_s [43]. This has a constant time cost of η_4 .

The total time cost for generating the policy is then derived to be

$$\begin{aligned}
\tau_{\text{pol}} &= \sum_{j=1}^{\Gamma} \sum_{k=1}^{N_p} (\Omega_{\text{fit}} + \Omega_{\text{off}} + \Omega_s) \\
&= \sum_{j=1}^{\Gamma} N_p \left(5\eta_1(N^4 - N^3) + [\eta_2 + \eta_3 N + \frac{\eta_1}{2}(N^4 - N^3)] + \eta_4 \right) \\
&= \Gamma N_p \left(5\eta_1(N^4 - N^3) + [\eta_2 + \eta_3 N + \frac{\eta_1}{2}(N^4 - N^3)] + \eta_4 \right).
\end{aligned} \tag{3.23}$$

As Γ and N_p scale linearly, $\tau_{\text{pol}} \in O(N^6)$ [72]. Also, in order to compute the policy for N -particles, $(N-1)$ -particle policy is required. This means there is another loop with time complexity $O(N^6)$. Using the Gauss summation formula and assuming the initial N value is 1, the final time complexity for policy generation is $O(N^7)$ [49].

3.5.2 Implementation cost

Implementation complexity is the resource complexity required to implement an AQEM policy, quantified by the scaling of space and time costs with the number of particles N [68]. The differential evolution algorithm receives the input from the detectors and transmits a signal to the actuator that adjusts the controllable beam-splitter reflectivity Θ according to the policy ϱ [49]. The storage of ϱ requires computer memory and its execution requires time.

In the case of the DE algorithm, we need to determine the maximum size of memory required to effectively execute the policy and complete the feedback loop. The size of the policy is the number of parameters involved, which is N [43]. This is because of the generalized-logarithmic-search heuristic responsible for the controllable beam-splitter reflectivity vector of size N . This leads to space complexity, as far as the AQEM policy implementation is concerned, that scales as $O(N)$.

Time complexity is determined from the upper bound of the time cost that is used to implement a single AQEM shot, i.e., the cost for using all N particles once [49]. While calculating the time cost, we assume that the time a given particle takes to pass through an interferometer is constant. We then perform loop analysis, where we count the number of loops needed to complete the procedure, which we assume all take constant time [67]. The quantity Θ_m is computed N times during one run involving N particles. This process is dictated by the policy ϱ . If the time needed for m -th particle to pass through the interferometer is Ω_p and the time required to calculate Θ_{m+1} is Ω_c , then the total time cost for the execution of variable beam-splitter reflectivity estimation is

$$\begin{aligned}\tau_{\text{imp}} &= \sum_{m=1}^N (\Omega_p + \Omega_c) \\ &= N(\Omega_p + \Omega_c).\end{aligned}\tag{3.24}$$

This shows that the implementation time complexity is also $O(N)$.

Chapter 4

Results

In this chapter, we report and discuss in detail different aspects of the variable beam-splitter reflectivity estimation results.

We first look at the variable beam-splitter reflectivity estimates and plot the probability distributions in the case of two θ values for two photon values each. We then study the distributions using skewness and kurtosis. Throughout the chapter, kurtosis refers to excess kurtosis.

Then we calculate the Holevo variance and present the logarithmic plots V_H vs the log of the number of particles, in this case, photons. This is done to see how well the log-log plots exhibit a linear relationship and also to determine the imprecision scaling. To provide a wider context, we also produce a plot that displays the scaling for various beam-splitter reflectivity θ values. The scaling is then compared with the SQL. We then report the empirical time complexity for the AQEM algorithm and compare the value to the theoretical time cost. Here we also study the policy generation stage of the estimation process. In the end, we provide the reader with plots of one-dimensional slices of the multi-dimensional policy landscape to highlight its characteristics.

Results shown below correspond to that of an ideal interferometer, meaning a Mach-Zehnder interferometer with no external or internal noise present. We also assume there is

no photon loss present in the system.

4.1 Distribution of the beam-splitter reflectivity estimates

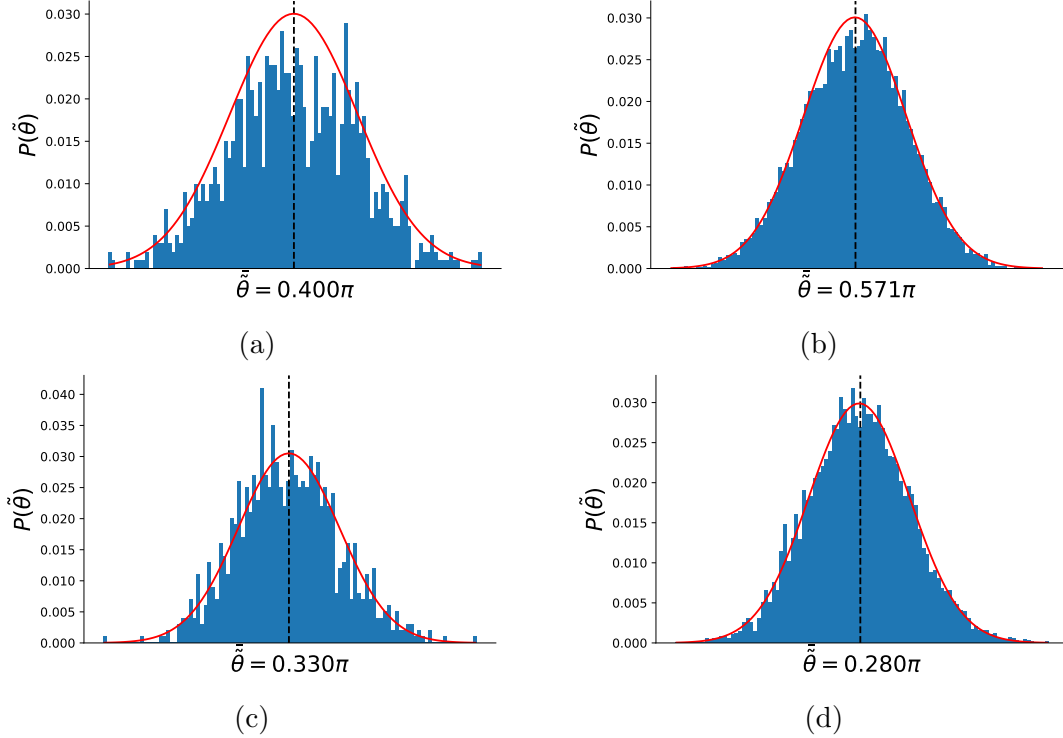


Figure 4.1: The distribution of the estimates for different θ values in the case of 50 and 100 photons. (a) When θ and the target value are $\pi/2$, the mean estimate recorded is $\bar{\tilde{\theta}} = 0.400\pi$. This run involves 50 photons. (b) In the case of 100 photons, the mean estimate recorded is $\bar{\tilde{\theta}} = 0.571\pi$. (for $\theta = \text{target value} = \pi/2$). (c) The mean of the estimate distribution is 0.330π when $\theta = 3\pi/4$ and the target is $\pi/4$. (50 photons). (d) When $\theta = 3\pi/4$ and the target is $\pi/4$ for 100 photons, $\bar{\tilde{\theta}} = 0.280\pi$.

In this section, we plot the obtained beam-splitter reflectivity estimate distributions for two θ values with 50 and 100 photons each. We also list down some important characteristics of these distributions for further analysis in the next chapter.

During preliminary runs, it was seen that the controllable beam splitter reflectivity converged towards $\pi - \theta$ for any given θ . This was established using 42 photons for 49 different θ values uniformly randomly chosen in the range 0 to π inclusive. As a result, the algorithm was

modified by incorporating the knowledge gained. Consequently, the sharpness was altered such that it ended up being

$$S = \left| \sum_{k=1}^K \frac{\exp[i((\pi - \theta_k) - \tilde{\theta}_k)]}{K} \right|. \quad (4.1)$$

This updated sharpness formula was used to train the algorithm on $K = 10N^2$ samples of θ uniformly randomly chosen in the range 0 to π inclusive.

After the training phase, the estimation algorithm was tested using 49 different θ values uniformly randomly chosen in the range 0 to π inclusive. This was done for 50 and 100 photons in each case.

Fig. 4.1 shows the probability distributions of the beam-splitter estimates $\tilde{\theta}$ in four such cases. The plots were produced with $10N^2 \tilde{\theta}$ samples each. The histograms show the estimate probabilities while the red curve is the best-fit Gaussian. The black dotted line in the figure is the mean estimate $\bar{\tilde{\theta}}$ in each case. The marks along the abscissa are omitted so that the estimates mean can be seen easily. Ideally,

$$\bar{\tilde{\theta}} = \pi - \theta, \quad (4.2)$$

meaning the closer $\bar{\tilde{\theta}}$ is to $\pi - \theta$, the better it is. The bias is defined as $|(\pi - \theta) - \bar{\tilde{\theta}}|$.

The plots in Fig. 4.1 (a), (c) correspond to the case when 50 photons were used in the input state while those in Fig. 4.1 (b), (d) correspond to the 100-photon case. We used the ones involving 100 photons for further analysis. The plots corresponding to the 50-photon case provide an insightful contrast and help in understanding the evolution of the performance of the estimation algorithm.

Table 4.1 contains data regarding the plots shown in Fig. 4.1 that we used to produce more results in the following sections and to analyze the results in the next chapter.

N	θ	$\pi - \theta$	$\bar{\bar{\theta}}$	Sharpness	Skewness	SE_{skew}	z_{skew}	Kurtosis	SE_{kurt}	z_{kurt}
50	$\pi/2$	$\pi/2$	0.400π	0.679	-0.004	0.077	-0.023	-0.146	0.155	-0.941
100	$\pi/2$	$\pi/2$	0.571π	0.992	-0.033	0.024	-1.375	-0.059	0.049	-1.204
50	$3\pi/4$	$\pi/4$	0.330π	0.631	0.161	0.077	2.091	0.053	0.155	0.342
100	$3\pi/4$	$\pi/4$	0.280π	0.996	0.038	0.024	1.583	0.028	0.049	0.571

Table 4.1: The sharpness, skewness, kurtosis and related quantities of the probability distributions shown in Fig. 4.1 are listed. The results are listed for $N = 50$ photons and $N = 100$ photons.

The sharpness given in the table was calculated using the updated sharpness formula given in equation 4.1. We used skewness and kurtosis of a given estimate probability distribution to decide whether it is approximately Gaussian and whether the standard analytical techniques associated with Gaussian distributions could be used. For this reason, Table 4.1 contains the skewness and kurtosis values for the plots given in Fig. 4.1 (a), (b), (c) and (d).

Here SE is the standard error, which is the standard deviation of the sampling distribution [63], and z is the z -statistic [23]. SE_{skew} is the standard deviation in the skewness of the sampling distribution while SE_{kurt} is the standard deviation in kurtosis of the sampling distribution under consideration. In the case of skewness, z -statistic is

$$z_{\text{skew}} = \text{Skewness} / \text{SE}_{\text{skew}}. \quad (4.3)$$

Given the results involving skewness and kurtosis and analysis regarding how negligible these values are, we decided to truncate the expansion series at kurtosis (or at the fourth power). See Section 5.1 in the next chapter for the relevant discussion.

4.2 Asymptotic power-law scaling

In this section, we present the log-log plot of Holevo variance vs the number of photons and determine the imprecision scaling $\Delta\tilde{\theta} \propto N^{-\varsigma}$ for select values of beam-splitter reflectivity. We also calculate the coefficient of determination to ascertain goodness of fit. The data presented here will be used to discuss various aspects of our work, in the next chapter.




Figure 4.2: Logarithm of Holevo variance vs $\log(N)$ generated using the sine state for $N = 4$ to 100. The ideal (no photon loss) experiment was repeated with beam-splitter reflectivities equal to 0.25, 0.5, 0.75 and 0.91. SQL and HL are shown for comparison.

The $\log(V_H)$ vs $\log(N)$ graphs for four of the 49 tested θ values are shown in the Fig. 4.2. The chosen values are $\theta = 0.25, 0.5, 0.75$ and 0.91 . This is done to present a good spread of θ values for valuable insight. The established bounds, SQL and HL, are also included for comparison.

In order to produce the plots in Fig. 4.2, Holevo variance was first calculated from sharpness which itself was calculated from the estimate distributions. This was done for photon number $N = 4$ to 100. It was determined during the preliminary runs that the adaptive procedure does not work properly for $N < 4$. It would either produce nonsensical results or the optimization algorithm would not converge at all.

As explained in section 3.3, the V_H values for $N \in \{4, 93\}$ were calculated after terminating the policy search after 100 generations and using the optimized result. For $N \in \{94, 100\}$,

we accepted a policy only if the V_H was within a distance corresponding to a confidence interval of 0.98 from the inverse power-law line. This two-step method gave us the linear plots that can be seen in Fig. 4.2.

The graphs in Fig. 4.2 have gradients equal to -2ς as $V_H = WN^{-2\varsigma}$. The ordinate intercepts of the four plots shown are in the range -1.0 to -0.5 . These intercepts in the figure were manually adjusted to give a better comparison. The scaling, W (original ordinate intercept) and the goodness of fit values are listed in Table 4.2.

As per Table 4.2, the logarithmic plots for the four quoted reflectivities have gradients less than SQL (gradient = -1.0). This means that, in all four cases, the beam-splitter reflectivity estimation imprecision (with respect to the photon number) scaled better than SQL. These scalings were calculated using regression analysis. The goodness of fit or the coefficients of determination for the cases shown in Fig. 4.2 are greater than or equal to 0.9993.

Reflectivity	θ	2ς	W	$\overline{R^2}$
0.25	$2\pi/3$	1.3351	0.811	0.9994
0.5	$\pi/2$	1.3995	0.920	0.9996
0.75	$\pi/3$	1.2943	0.387	0.9993
0.91	0.194π	1.1910	1.277	0.9996

Table 4.2: The imprecision scaling, reflectivity, original ordinate intercept (W) and the goodness of fit values are shown for various beam-splitter reflectivity values.

Fig. 4.3 shows the imprecision scaling of all 49 θ values used in the test phase, performed using $N \in \{4, 100\}$. Here is the scaling range:

$$1.191 \leq 2\varsigma \leq 1.4328. \quad (4.4)$$

Note that in all cases we achieved better than SQL scaling with only 42 photons but the

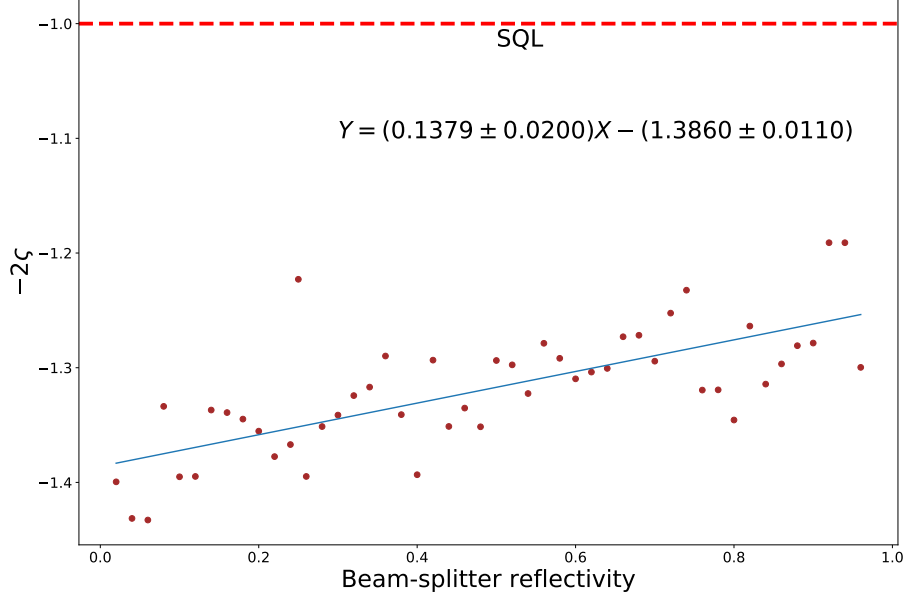


Figure 4.3: The scalings are shown for several beam-splitter reflectivities with no photon loss. The least squares regression line is included to study the trend exhibited. SQL level is drawn for reference.

simulations were run with 100 photons to prove the efficacy of the technique for this photon number. Beyond 100 photons, the computational time required to find optimal policies was impractically high and hence our AQEM scheme is not feasible for $N > 100$.

In all 49 test cases, the ordinate intercept was found to be in the range -0.997 to 0.131 . Moreover, the coefficient of determination in these cases was calculated to be at least 0.9991 .

In order to quantify the trend seen in the scatter plot in Fig. 4.3, a least-squares regression line was included. As can be seen, it shows a positive linear relationship between the scaling -2ζ and the beam-splitter reflectivity estimate $\tilde{\theta}$. The equation for the trend line is $Y = (0.1379 \pm 0.0200)X - (1.3860 \pm 0.0110)$ where Y is the scaling and X is the beam-splitter reflectivity estimate. The adjusted coefficient of determination $\overline{R^2} = 0.497$.

4.3 Policy space and optimization algorithm

In this section, we plot the one-dimensional slices of the multi-dimensional policy space that the policy-search algorithm searches for optimal policies. We also publish data regarding the performance of the optimization algorithm.




Figure 4.4: The policy-search space is parametrized as $\Delta_m = k/m^\alpha$ where k is a constant. (a) A plot of sharpness vs α is shown when $k = \pi$. (b) A plot of sharpness vs α is shown when $k = 3\pi/4$.

The policy landscape is shown in Fig. 4.4(a), (b). These one-dimensional slices of the multi-parameter landscape is parametrized as $\Delta = k/m^\alpha$ where k is a constant. As α varies from -0.5 to 5.1 , the landscape changes drastically. This is because the sharpness changes quite drastically over the plotted range. The goal of the policy-search algorithm is to find the global maximum (or find the max sharpness value), which leads to the minimum Holevo variance.




Figure 4.5: The adjustments Δ_m executed by the processing unit to control the beam-splitter reflectivity Θ in a 100-photon AQEM scheme.

Figure 4.5 shows a plot of successive adjustments Δ that the processing unit is implementing in order to achieve quantum-enhanced precision. The plot was generated with $\theta = \pi/3$ (reflectivity = 0.75) in the case of 100 photons. The collection of all the adjustments forms

the policy ϱ for this particular case. This policy lives somewhere on the policy landscape small sections of which are shown in Fig. 4.4.

	Theoretical value	Empirical value
Policy generation time cost	$O(N^7)$	$O(N^{7.338})$
AQEM implementation time cost (given the policy)	$O(N)$	$O(N^{1.261})$

Table 4.3: Theoretical and empirical values shown for generating the policy and for implementing the AQEM algorithm

The optimization algorithm’s performance can be evaluated by analyzing the resource cost incurred during empirical runs. The empirical time cost along with the theoretical time cost for the AQEM algorithm is listed in Table 4.3. As can be seen, the empirical value scales slightly worse than the theoretical one both for finding the optimal policies and for implementing the AQEM scheme once the policy has been found. The empirical values quoted were calculated using the time spent during the training phase using 93 photons. The average over all the different θ cases was taken to get the final value.

All of the simulations were implemented on Compute Canada clusters with 12 identical cores. Assuming the time taken to communicate between cores was negligible, the total time taken by a given task was 12 x time taken on one cluster. This is a reasonable approach because the task was distributed equally across 12 cores. Table 4.3 shows this led to scalings that are close to the theoretical upper bounds for optimization and implementation within the algorithm.

Chapter 5

Discussion

In this chapter, we analyze the results reported in the previous chapter. First we study the beam-splitter estimate distributions using cumulants (throughout the chapter, kurtosis refers to excess kurtosis). We see if the plots are approximately Gaussian or not using commonly used tests of normality. We also entertain the idea that the control beam splitter’s role can be described as that of a ‘compensator’. Moreover, we investigate the logarithmic plots and the ς values for different reflectivities and interpret the results including any subtleties involved. Furthermore, we look at the factors that lead to the discrepancy between the empirical and theoretical time cost of implementing the algorithm and finding the optimal policy. We also explain why the one-dimensional policy landscape parametrized by α establishes the non-convexity of the entire policy space.

5.1 Distribution of the beam-splitter reflectivity estimates

In this section, we analyze the distribution of beam-splitter reflectivity estimates $\tilde{\theta}$. More specifically we decide whether the distribution plots are approximately Gaussian or not using normality tests.




Figure 5.1: The graphs for $r = \cos^2(\theta/2)$ and for $r = \cos^2((\pi - \theta)/2)$ when $\theta \in [0, 4\pi - 1]$ are shown.

One of the crucial building blocks of the estimate probability generation is the controllable beam splitter. This beam splitter, parametrized by Θ , can be thought of as ‘compensating’ the non-50 : 50 dynamics of the first beam splitter. When both beam splitters are 50 : 50 ones, then the input state is balanced in that an equal number of particles are injected into both ports. This is no longer the case with non-50 : 50 beam splitters. The non-50 : 50 nature of the first beam splitter makes the ratio of the injected particles different from 1 : 1. For example, in a 60 : 40 beam splitter, the ratio of the particles that are injected into the two ports is 3 : 2. In this case, the algorithm devises a policy such that the control beam splitter converges towards a 2 : 3 or a 40 : 60 configuration. This is equivalent to having $\Theta = \pi - \theta$.

In the preliminary simulation runs, we observed this convergence in action. We used 42 photons in order to establish this as it was time and resource saving. We observed a 20% decrease in the run time after we modified the algorithm to include the knowledge gained. More specifically, we modified the policy-search algorithm such that it preferentially converged towards the $\Theta \approx \pi - \theta$ region. We also modified the sharpness as shown in equation 4.1.

Fig. 5.1 shows the graphs for the unknown and the control beam-splitter reflectivities. These graphs are not empirical results, instead, these are analytical reflectivities for visualizing the behaviours of the two beam splitters.

Fig. 4.1 shows the estimate distributions with their relevant characterizing values popu-

lating Table 4.1. Fig. 4.1 (a) in particular shows the distribution for $\theta = \pi/2$ in the case of 50 photons. In order to decide whether the estimate distribution is approximately Gaussian and whether the skewness and kurtosis are within tolerable limits, we used the standard errors for skewness and kurtosis. These in turn were used to calculate the corresponding z scores listed in Table 4.1.

For the plot shown in Fig. 4.1 (a), the skewness was calculated to be -0.004 . Its z_{skew} score shows that the sample skewness is 0.023 standard errors below the mean of zero [64]. A plot with the z_{skew} -score that is twice the size of the standard error of skewness is usually taken to be non-normal in nature [34]. In other words, a given distribution's skewness is within tolerable limits if

$$-2 \leq z_{\text{skew}} \leq 2. \quad (5.1)$$

This corresponds to the 95% confidence interval which means if this holds true then there is only 5% probability that this level of skewness comes from a non-normal distribution [64]. There is a similar test involving kurtosis.

As $|z_{\text{skew}}| = |-0.023| < 2$ we suspect that the probability distribution given in Fig. 4.1 (a) is Gaussian. Moreover, as $|z_{\text{kurt}}| = |-0.941| < 2$, we conclude that the estimate distribution given in Fig. 4.1 (a) is Gaussian indeed. This demonstrates that the skewness and kurtosis for this distribution are within tolerable limits.

The negative value for skewness in Fig. 4.1 (a) indicates that the distribution is skewed in the negative direction. The negative kurtosis value indicates that the distribution tails are heavier compared to Gaussian tails [34]. This usually results in a smaller-than-Gaussian peak. This is not true in this case. In fact, the opposite is true. The plot has 'lighter'-than-Gaussian tails because the θ domain is compact. θ only varies in the range 0 to π inclusive. This should, in principle, make the peak of the normalized distribution larger than a standard Gaussian peak leading to a positive kurtosis value. The table shows negative kurtosis probably because of the lack of enough samples. As the kurtosis is tolerable, the negative value is not problematic.

Based on the data provided in Table 4.1, we conclude that the plot shown in Fig 4.1 (b) is approximately Gaussian. This is because both the $|z_{\text{skew}}| = |-1.375|$ and $|z_{\text{kurt}}| = |-1.204|$ are less than 2 and therefore, the skewness and kurtosis are within tolerable limits. Negative skewness in this case illustrates that the distribution is skewed in the negative direction. In other words, there are more data points on the ‘negative’ side of the distribution leading to asymmetry [64].

Applying the stated tests of normality, it can be seen that the plot shown in Fig 4.1 (c) is not approximately Gaussian while the one shown in Fig 4.1 (d) is. Notice both of these plots have the positive kurtosis as expected.

The estimate distributions in all 49 test cases were analyzed for 50 and 100 photons and it was concluded that the skewness and kurtosis are within a tolerable range for only 100 photons. In other words, all tested 100-photon estimate distributions passed the normality tests involving skewness and kurtosis. Most of the 50-photon cases failed normality tests probably because there are not enough data to populate the estimate probability distributions such that they exhibit approximately Gaussian behaviour. Apart from the data presented in Table 4.1, this can be verified by observing that the 50 photon plots in Fig. 4.1 (a), (c) look less Gaussian than their 100 photon counterparts given in Fig. 4.1 (b), (d). This is because the histograms in Fig. 4.1 (a), (c) don’t closely follow the shape of the fitted red Gaussian curves. Furthermore, these distributions are choppy and also have more outliers.

Once we established that the 100 photon regime was our area of interest, we truncated the expanded series and our analysis at kurtosis or the fourth power in the series. This is because in all cases involving 100 photons, the skewness and kurtosis are negligible and so all higher powers must be too.

5.2 Asymptotic power-law scaling

In this section, we discuss the asymptotic power-law scaling displayed in the plots shown in section 4.2. We also comment on the goodness of fit and write in detail about the implications of the results presented.

In the testing phase, we kept the θ value constant for generating the $\tilde{\theta}$ distribution for a particular beam-splitter reflectivity value and ultimately the sharpness using the distribution. That meant modifying the sharpness (Eq. 3.14) such that θ ended up without a subscript:

$$S = \left| \sum_{k=1}^K \frac{\exp[i((\pi - \theta) - \tilde{\theta}_k)]}{K} \right|. \quad (5.2)$$

This modified formula was used to produce the graphs shown in Fig. 4.2. As mentioned in section 4.2, these graphs were produced for $N \in \{4, 100\}$. For $N = 1, 2$ and 3 , there are only a handful K values available for training the algorithm as shown in Table 5.1. This led to inadequate training. Increasing the size of the training set did not work either. The adaptive quantum-enhanced measurement procedure did not work very well even with $N = 4$ but it did not return nonsensical results. That is why we chose that as our starting photon number.

N	1	2	3
$K = 10N^2$	10	40	90

Table 5.1: K values are shown for $N \in \{1, 2, 3\}$

The optimization process becomes increasingly time-consuming after $N = 93$ predominantly because there is an accept-reject criterion in place in this regime. Without this extra measure in place, we failed to get a linear log-log relationship and in turn the quantum-enhanced precision for $N \in \{94, 100\}$. Although with the stated criterion in place we did get a linear relationship with high accuracy and managed to beat SQL, we failed to scale beyond $N = 100$ even with parallelization built in the optimization algorithm. The employed

scheme was able to achieve the goal of the project, i.e. better-than-SQL reflectivity scaling with respect to the photon number. If we had used coherent state instead of the sine state as the input we could not have beaten the SQL. This is because the coherent state achieves SQL but does not surpass it [19].

Apart from scaling, ordinate intercept plays an important role especially when it comes to implementing our scheme in a lab setting. After taking log on both sides, $V_H = WN^{-2\varsigma}$ becomes

$$\log(V_H) = (-2\varsigma) \log(N) + \log(W), \quad (5.3)$$

where W is the constant of proportionality. As $-0.997 \leq \log(W) \leq 0.131$, this leads to $0.1 \leq W \leq 1.352$. Therefore with 100 photons, the largest error or $\Delta\tilde{\theta}$ is only 0.116 or 11.6% when the estimation scheme is just hitting the SQL precision level. Our results are better than SQL and hence the largest error in our case is smaller than 11.6%.

Table 4.2 shows the goodness of fit values for select θ values. We used the adjusted coefficient of determination in our work because there is a non-zero bias b present in the output (estimate distributions). As can be seen, the $\overline{R^2}$ values are very close to 1 indicating that there exists a strong linear relationship between the $\log(V_H)$ and $\log(N)$. This means that the variance in $\log(V_H)$ can be predicted by $\log(N)$ with more than 99.9% confidence [1]. In other words, there is less than 0.1% chance that there does not exist any linear relationship between the two quantities along the axes in Fig. 4.1. All 49 values tested manifested a similarly high level of goodness of fit.

As can be clearly seen, a small number of data points in the plots shown in Fig. 4.2 deviate from the trend shown. This gives rise to bumps and depressions in the trend line. These are sampling errors.

The scatter plot shown in Fig. 4.3 also exhibits a linear relationship, but not a very good one. According to the trend line equation, the scaling gets closer to SQL as θ is increased from 0 to π . That is, the scaling is getting worse in general in that direction. Due to wide variability and the fact that most of the data points are located at considerable distances

from the trend line, the adjusted coefficient of determination is only 0.497. This means that the trend line can only explain less than half of the variation possessed by the plot [1]. Due to the lack of any strong patterns and data, we cannot comment on why the trend is the way it is. This makes it difficult to explain the scatter in physical terms and in terms of the equipment and approach-related dynamics. Fully interpreting this scatter plot can very well be the subject of future research work.

Although the policy-search algorithm successfully navigated the solution space and returned a policy that achieved better scaling than SQL with respect to the number of photons, it did not find the global maximum. This is evident from the fact that we did not get symmetric scaling on either side of the reflectivity 0.5 in the plot (Fig. 4.3). As the reflectivity varies symmetrically on either side of $r = 0.5$, we believe that the policy residing on the global sharpness maximum would have given symmetrical scaling on either side of $r = 0.5$. This is not true in our case leading us to conclude that there is room for improvement as far as policy-search goes.

5.3 Policy space and optimization algorithm

In this section, we analyze the policy landscape plotted in section 4.3. We discuss why it looks the way it does and also how does the optimization process work in this case. We also remark on the performance of the optimization performance.

Fig. 4.4 (a), (b) show one-dimensional slices of the multi-dimensional policy landscape. In Fig. 4.4 (a) for instance, the landscape is highly uneven and has several sharpness crests and troughs, some of which correspond with $\alpha = 0.4, 1.2, 4.5$. The peak at $\alpha = 0.4$ is the global maximum while all the other peaks are local maxima. Notice that we parametrized the landscape as $\Delta_m = k/m^\alpha$. This is because of the plot shown in Fig. 4.5, which shows the policy in action, in the case of one of the 100 photon test runs. Initially, the processing unit implements course adjustments so as to get close to the analytical solution. Later, the

adjustments become more and more refined as more knowledge from the measurements is gathered and utilized. The plot in Fig. 4.5 roughly follows the power law k/m^α justifying how the landscape in Fig. 4.4 is parametrized.

	$\alpha = 1$	$\alpha = 2$	$\alpha = 3$
Δ_1	π	π	π
Δ_2	$\frac{\pi}{2}$	$\frac{\pi}{4}$	$\frac{\pi}{8}$
Δ_3	$\frac{\pi}{3}$	$\frac{\pi}{9}$	$\frac{\pi}{27}$
Δ_4	$\frac{\pi}{4}$	$\frac{\pi}{16}$	$\frac{\pi}{64}$

Table 5.2: $\Delta_m = k/m^\alpha$ values are shown for $k = \pi$ and $\alpha \in \{1, 2, 3\}$

In Fig. 4.4, the landscape slices shown are not smooth. For instance in Fig. 4.4(a), there is a sharp sharpness peak corresponding to $\alpha = 0.5$. Similar non-smooth areas can be identified in both of the slices. We could not achieve smoothness even with more samples.

The policy-search algorithm's job in our work is to search the landscape small sections of which are shown in Fig. 4.4 (a), (b). The goal is to search the policy ϱ or the collection of adjustments Δ for changing the control beam splitter's reflectivity in an optimal manner after every photon measurement. The adjustments Δ for various locations in the landscape slice in Fig. 4.4 (a) are shown in Table 5.2. Notice, depending on the value of α the algorithm has the potential to end up at radically different sharpness values in Fig. 4.4 (a). This indicates the presence of several minima and maxima in the landscape. Given the data in Table 5.2 along with the shape of the plot shown in Fig. 4.4 (a), it is easy to see that the plot is non-convex. This demonstrates that the entire multi-dimensional policy landscape is non-convex. This is because if a given landscape is non-convex at dimension d , it must be non-convex at every dimension greater than d [46]. The landscape slice with a different parameter value is

shown in Fig. 4.4 (b) to further demonstrate the non-convexity of the space in question.

An algorithm that converges at any of the local maxima in such a landscape essentially fails to find the global maximum. This shows the need for non-convex optimization. Although our algorithm managed to find the required policy for achieving quantum-enhanced precision, it failed to find the global maximum as demonstrated in the previous section.

The algorithm’s ability to find a feasible solution is not the only thing that makes it usable and practical. The time an algorithm takes to find that solution is also highly relevant. Table 4.3 lists the theoretical and the empirical values for the time the algorithm takes to generate an optimal policy and to implement AQEM when the policy has been found. In both cases, as is evident, the empirical scaling is slightly worse than the theoretical upper bound. The worse scaling in both cases might be attributed to the fact that the inter-core communication time is not negligible. It might also be due to inefficiencies in the way the cluster hardware is set up. Moreover, the code itself probably contains unoptimized sections among other difficult to pin-down reasons.

Chapter 6

Conclusion

In this chapter, we summarize the main findings and approach that we took to solve the adaptive quantum-enhanced variable beam-splitter reflectivity estimation problem. We discuss how the setup and the policy-search algorithm gave the intended results. We also explore any future research projects that our work could potentially be a stepping stone for.

6.1 Summary

In this work, we considered a two-mode Mach-Zehnder interferometer with a constant phase shift in one of the arms. Additionally, we introduced beam-splitters with variable reflectivities. The first beam splitter had the unknown reflectivity θ while the second one had the controllable reflectivity Θ and was connected to a processing unit.

This is a $U(1)$ rotation problem involving a $U(1)$ subgroup that is different from the one involved in phase estimation. Geometrically speaking, the goal of the research work was to estimate the unknown axis in the equatorial plane the initial state was rotated about.

The input state used for the project was the permutationally symmetric and loss-tolerant sine state. A total of 100 photons were sent through the interferometer such that every photon interacted with the quantum system (the interferometer) in a sequential manner. Two detectors, one for each arm, detected the photons on the other end and the measurement

results were communicated to a processing unit that controlled the reflectivity of the controllable beam splitter. The job of the processing unit was to ‘guess’ the unknown beam-splitter reflectivity θ and adjust the Θ value after every photon measurement. The estimate got better as more and more sequential measurement data was accumulated. The processing unit varied the Θ value such that the imprecision of the θ estimate (or $\tilde{\theta}$) scaled better than SQL. Here the better-than-SQL scaling is the precision with which the employed scheme was able to achieve the unknown reflectivity estimation.

The processing unit controlled the Θ value as dictated by the policy determined by the DE algorithm. The policy determined the value of any given Θ by taking the value of the preceding Θ and either adding or subtracting the adjustment, Δ , from it. This method of determining the optimal policy was the basis of the adaptive measurements. The policy decision space turned out to be a binary decision tree with $2^N - 1$ paths where N being the number of particles used, photons in this case.

Differential evolution algorithm was employed for searching the solution space for optimal policies. DE is inspired by how genes combine and mutate and how biological reproduction takes place at the genetic level. A given number of candidate solutions are randomly generated in the beginning. Then the best performing solutions are combined and ‘mutated’ to produce the next generation of solutions. This ensures that the solutions get better at the task at hand with every generation. This process is repeated until a good enough or best solution is found.

We implemented the adaptive quantum-enhanced variable beam-splitter estimation procedure for $10N^2$ different θ values. These values were chosen randomly from a uniform distribution in the range 0 and π . This way we trained the algorithm while minimizing overfitting. During preliminary runs, it was found out that the Θ value was converging towards $\pi - \theta$. For that reason, the controllable beam splitter can be thought of as a compensator as it compensated for the non-50 : 50 nature of the first beam splitter.

The probability distributions of reflectivity estimates showed non-zero skewness and kur-

tosis associated with non-normal plots. It turned out that the skewness and kurtosis were well within the tolerable limits and therefore, the usual techniques associated with Gaussian distributions could be used for further analysis.

The sharpness was calculated from the estimate-distribution plots, which was then used to calculate Holevo variance. Thereafter, the variance was used to plot the log-log plot of Holevo variance vs the number of particles N for $N \in \{4, 100\}$. It was found that the estimate imprecision ($\Delta\tilde{\theta}$) scaled better than SQL (or $O(N^{-0.5})$) in all 49 test cases. The unknown beam-splitter reflectivities θ used in those cases were uniformly randomly selected in the range 0 and π inclusive. The adjusted coefficient of determination $\overline{R^2}$ was calculated to be at least 0.9991 in these cases.

This proves that DE showed good performance when it came to finding quantum-enhanced estimation policies, which is a non-convex optimization task given our problem. We proved that by plotting the one-dimensional slices of the policy landscape. The non-convex nature of these plots indicated that the entire multi-dimensional policy landscape is non-convex.

We started the DE-based optimization process with 50 randomly generated solutions and ended the process after 100 generations for $N \in \{4, 93\}$. This is because it returned good enough estimates for our problem. For $N > 93$, we accepted a policy only if V_H was within a distance corresponding to a confidence interval of 0.98 from the inverse power-law line. We ran the algorithm for up to 100 photons beyond which point the computation cost became too high for the simulation to be practical. As the solution space was highly non-convex, we took extra measures to make the optimization process tractable. Some of these measures included parallelization of the calculations involved and vectorization for efficient processing.

Although the policy-search algorithm successfully navigated the non-convex solution space and returned policies that achieved better than SQL scaling with respect to the number of photons, it did not find the global sharpness maximum. This is evident from the fact that we did not get symmetric scaling on either side of the reflectivity 0.5. We did however get a linear trend with a positive gradient with respect to the scaling and the un-

known beam-splitter reflectivity θ . This trend is not strong enough to tie it to the dynamics involved.

The time cost of the algorithm scaled as $O(N^{7.338})$ for the policy generation and $O(N^{1.261})$ for implementing the AQEM policy. This is when only one CPU was used for implementing the optimization and the estimation scheme. The simulation was run on a cluster with 12 identical cores to decrease run time.

As the algorithm did not use any knowledge of the model, the scheme is model-free and data-driven. This ensures that the existing code library can be used in the presence of noise, something that is difficult to accomplish with model-dependent methods.

The algorithm is a collection of modules so that it can be modified and extended for a given quantum metrology problem. It is written in C++ with download and installation instructions and the link provided in Appendix A.

Our work set the stage for two-parameter and multi-parameter quantum-enhanced estimation schemes, which could potentially achieve higher parameter-estimation precision.

6.2 Future research

Our work can be used and extended in several ways.

1. As shown in Fig. 4.3, there is a weak linear relationship between the imprecision scaling and the unknown beam-splitter reflectivity θ . We failed to rationalize the given trend given our approach, the technique and the results collected. A future study could potentially collect more data, study the trend in more detail and shed some light on the results.
2. Adaptive variable beam-splitter reflectivity estimation with photon loss is the next logical step. The input state used in our work has been shown to be robust against photon loss [49]. The inclusion of photon loss in the scheme will help replicate lab conditions better.

3. AQEM schemes with variable beam-splitters in a noisy environment will replicate another factor very common in lab-like conditions. In order to study noise, several noise models can be included in the simulation. Telegraph, skew-normal and log-normal noise models are good places to start. The data can be collected and analyzed in each case for different values of noise model parameters.
4. Another research path includes multi-parameter estimation problems. A two-parameter system where both phase and beam-splitter reflectivities are variables is an interesting line of inquiry. Note that this problem is computationally and analytically harder to tackle.

One way to devise a two-parameter estimation scheme is to deform or modulate the input state, starting with the sine state, until quantum-enhanced precision is achieved. Sine state distributes the photons equally across the two ports of the beam splitter. By lifting this constraint and varying the ratio of photons across the ports, quantum-enhanced precision may be achieved for variable phase ϕ and variable reflectivities θ and $-\theta$ for the two beam splitters respectively.

Another way to approach the problem is by introducing a controllable phase Φ in the second arm of the interferometer and a controllable reflectivity Θ in the second beam splitter. It would then be interesting to see if the quantum-enhanced precision is achieved in the presence of a variable phase ϕ and a variable reflectivity θ in the first beam splitter. The algorithm's task in this case would be to find an optimal policy that does the job described.

The successful two-parameter estimation scheme might involve some combination of the approaches listed.

5. More-than-two-parameter estimation problems are farther down the road. Insights gained from one-parameter and two-parameter schemes can be employed to deal with such problems.

Bibliography

- [1] M. Ahsanullah. *Focus on Applied Statistics*. New York: Nova Sci., 2003.
- [2] “4 - Automated data acquisition and processing”. In: *Environmental Monitoring and Characterization*. Ed. by J. F. Artiola, I. L. Pepper, and M. L. Brusseau. China: Acad. Press, 2004, pp. 49–67. DOI: <https://doi.org/10.1016/B978-012064477-3/50006-0>.
- [3] A. B. Badiru and L. A. Racz. *Handbook of Measurements: Benchmarks for Systems Accuracy and Precision*. Systems Innovation Book Series. Boca Raton, FL: CRC Press, 2018.
- [4] D. W. Berry and H. M. Wiseman. “Optimal states and almost optimal adaptive measurements for quantum interferometry”. In: *Phys. Rev. Lett.* 85.24 (Dec. 2000), pp. 5098–5101. DOI: [10.1103/PhysRevLett.85.5098](https://doi.org/10.1103/PhysRevLett.85.5098).
- [5] C. M. Bishop. *Pattern Recognition and Machine Learning*. Singapore: Springer, 2006.
- [6] J. J. Bollinger et al. “Optimal frequency measurements with maximally correlated states”. In: *Phys. Rev. A* 54.6 (Dec. 1996), R4649–R4652. DOI: [10.1103/PhysRevA.54.R4649](https://doi.org/10.1103/PhysRevA.54.R4649).
- [7] R. S. Bondurant and J. H. Shapiro. “Squeezed states in phase-sensing interferometers”. In: *Phys. Rev. D* 30.12 (Dec. 1984), pp. 2548–2556. DOI: [10.1103/PhysRevD.30.2548](https://doi.org/10.1103/PhysRevD.30.2548).

- [8] J. Borregaard and A. S. Sørensen. “Near-Heisenberg-limited atomic clocks in the presence of decoherence”. In: *Phys. Rev. Lett.* 111 (Aug. 2013), p. 090801. DOI: 10.1103/PhysRevLett.111.090801.
- [9] S. L. Braunstein. “Some limits to precision phase measurement”. In: *Phys. Rev. A* 49.1 (Jan. 1994), pp. 69–75. DOI: 10.1103/PhysRevA.49.69.
- [10] N. Brunel and J.-P. Nadal. “Mutual information, Fisher information, and population coding”. In: *Neural Comput.* 10.7 (1998), pp. 1731–1757. DOI: 10.1162/089976698300017115.
- [11] C. M. Caves. “Quantum-mechanical noise in an interferometer”. In: *Phys. Rev. D* 23.8 (Apr. 1981), pp. 1693–1708. DOI: 10.1103/PhysRevD.23.1693.
- [12] J. Clark. “On the controllability of quantum-mechanical systems (original 1981 version)”. In: (July 2016). DOI: 10.13140/RG.2.1.3868.0561.
- [13] A. DasGupta. *Asymptotic Theory of Statistics and Probability*. Springer Texts in Statistics. USA: Springer N. Y., 2008. ISBN: 9780387759708.
- [14] Y. N. Dauphin et al. “Identifying and attacking the saddle point problem in high-dimensional non-convex optimization”. In: *Advances in Neural Information Processing Systems 27*. Curran Assoc., Inc., 2014, pp. 2933–2941.
- [15] C. L. Degen, F. Reinhard, and P. Cappellaro. “Quantum sensing”. In: 2016.
- [16] M. P. Deisenroth, G. Neumann, and J. Peters. *A survey on policy search for robotics*. Ed. by H. Christensen and R. Siegwart. Vol. 2. Foundations and Trends. now, Sept. 2013. DOI: 10.1561/23000000021.
- [17] R. Demkowicz-Dobrzański, M. Jarzyna, and J. Kołodyński. “Chapter four - Quantum limits in optical interferometry”. In: ed. by E. Wolf. Vol. 60. Progress in Optics. Elsevier, 2015, pp. 345–435. DOI: <https://doi.org/10.1016/bs.po.2015.02.003>.

- [18] D. Dong and I.R. Petersen. “Quantum control theory and applications: a survey”. In: *IET Control Theory Appl.* 4.12 (Dec. 2010), pp. 2651–2671. DOI: 10.1049/iet-cta.2009.0508.
- [19] J. P. Dowling. “Quantum optical metrology – the lowdown on high-N00N states”. In: *Contemporary Physics* 49.2 (2008), pp. 125–143. DOI: 10.1080/00107510802091298.
- [20] J. P. Dowling and G. J. Milburn. “Quantum technology: the second quantum revolution”. In: *Philos. Trans. Royal Soc. Lond., A: Math. Phys. Eng. Sci.* 361.1809 (June 2003), pp. 1655–1674. DOI: 10.1098/rsta.2003.1227.
- [21] A. E. Eiben and J. E. Smith. “Popular evolutionary algorithm variants”. In: *Introduction to Evolutionary Computing*. 2nd ed. Natural Computing. Berlin: Springer, 2015. Chap. 6, pp. 99–116. ISBN: 978-3-662-44874-8. DOI: 10.1007/978-3-662-44874-8_6.
- [22] H. Fearn and R. Loudon. “Quantum theory of the lossless beam splitter”. In: *Opt. Commun* 64.6 (1987), pp. 485–490. DOI: [https://doi.org/10.1016/0030-4018\(87\)90275-6](https://doi.org/10.1016/0030-4018(87)90275-6).
- [23] M. C. Fleming and J. G. Nellis. *Principles of Applied Statistics: An Integrated Approach Using MINITAB and Excel*. UK: Thomson Learn., 2000. ISBN: 9781861525864.
- [24] P. J. Fleming and R. C. Purshouse. “Evolutionary algorithms in control systems engineering: a survey”. In: *Contr. Eng. Pract.* 10.11 (Nov. 2002), pp. 1223–1241. DOI: 10.1016/S0967-0661(02)00081-3.
- [25] B. R. Frieden. “What is Fisher information”. In: *Physics from Fisher Information*. Cambridge: CUP, 1998. Chap. 1, pp. 22–50.
- [26] V. Giovannetti, S. Lloyd, and L. Maccone. “Advances in quantum metrology”. In: *Nat. Photon.* 5 (Mar. 2011), pp. 222–229.
- [27] V. Giovannetti, S. Lloyd, and L. Maccone. “Quantum-enhanced measurements: beating the standard quantum limit”. In: *Science* 306.5700 (2004), pp. 1330–1336. DOI: 10.1126/science.1104149.

- [28] S. Gollapudi. *Practical Machine Learning*. English. Birmingham: Packt Publ. Ltd., 2016.
- [29] R. S. Govindaraju and B. S. Das. *Moment Analysis for Subsurface Hydrologic Applications*. Water Science and Technology Library. Dordrecht, Netherlands: Springer Neth., 2007. ISBN: 9781402057526.
- [30] B. Hall and B. C. Hall. *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*. Graduate Texts in Mathematics. New York: Springer, 2003.
- [31] A. Hentschel. “Particle Swarm Optimization for Adaptive Quantum-enhanced Phase Estimation”. PhD thesis. Calgary: University of Calgary, May 2011.
- [32] A. Hentschel and B. C. Sanders. “Machine learning for precise quantum measurement”. In: *Phys. Rev. Lett.* 104.6 (Feb. 2010). DOI: 10.1103/PhysRevLett.104.063603.
- [33] A. Hentschel and B. C. Sanders. “Ordered measurements of permutationally-symmetric qubit strings”. In: *J. Phys. A: Math. Theor.* 44.11 (Feb. 2011), p. 115301. DOI: 10.1088/1751-8113/44/11/115301.
- [34] P.R. Hinton. *Statistics Explained*. UK: T & F, 2014, pp. 97–100. ISBN: 9781317753926. URL: <https://books.google.ca/books?id=jfgjAwAAQBAJ>.
- [35] J. N. Hollenhorst. “Quantum limits on resonant-mass gravitational-radiation detectors”. In: *Phys. Rev. D* 19.6 (Mar. 1979), pp. 1669–1679. DOI: 10.1103/PhysRevD.19.1669.
- [36] Home — Compute Canada. Oct. 26, 2019. URL: <https://www.computecanada.ca/home> (visited on 10/26/2019).
- [37] Z.-S. Hou and Z. Wang. “From model-based control to data-driven control: Survey, classification and perspective”. In: *Inf. Sci.* 235 (June 2013), pp. 3–35. DOI: 10.1016/j.ins.2012.07.014.

- [38] C. H. Huang et al. *Languages and Compilers for Parallel Computing: 8th International Workshop, Columbus, Ohio, USA, August 10-12, 1995. Proceedings*. Lecture Notes in Computer Science v. 8. Ohio: Springer, 1996.
- [39] P. C. Humphreys. “Quantum enhanced multiple phase estimation”. In: *Phys. Rev. Lett.* 111.7 (2013). DOI: 10.1103/PhysRevLett.111.070403.
- [40] P. Jain and P. Kar. now, 2017, pp. 9–17. ISBN: 9781680833683.
- [41] T. R. Jain. *Quantitative Methods*. Delhi: Vk, 2007. ISBN: 9788187140573.
- [42] V. M. Kruglov. “A characterization of the Gaussian distribution”. In: *Stoch. Anal. Appl.* 31.5 (2013), pp. 872–875. DOI: 10.1080/07362994.2013.817250.
- [43] N. B. Lovett et al. “Differential evolution for many-particle adaptive quantum metrology”. In: *Phys. Rev. Lett.* 110 (May 2013). DOI: 10.1103/PhysRevLett.110.220501.
- [44] F. G. Major. *The Quantum Beat: The Physical Principles of Atomic Clocks*. New York: Springer N. Y., 2013. ISBN: 9781475729238.
- [45] D. D. McCarthy and K. P. Seidelmann. *Time: From Earth Rotation to Atomic Physics*. UK: Wiley, 2009. ISBN: 9783527627950.
- [46] E. S. Mistakidis and G. E. Stavroulakis. *Nonconvex Optimization in Mechanics: Algorithms, Heuristics and Engineering Applications by the F.E.M.* Nonconvex Optimization and Its Applications. USA: Springer US, 2013.
- [47] D. Orban. “Introduction to derivative-free optimization”. In: *SIAM Rev.* 53.2 (June 2011).
- [48] I. S. Osborne. “Quantum enhanced metrology”. In: *Science* 352.6293 (2016), pp. 1530–1532. ISSN: 0036-8075. DOI: 10.1126/science.352.6293.1530-m.
- [49] P. Palittapongarnpim. “Evolutionary Algorithm for Adaptive Quantum-channel Control”. PhD thesis. Calgary: University of Calgary, 2019.

- [50] P. Palittapongarnpim, P. Wittek, and B. C. Sanders. “Single-shot adaptive measurement for quantum-enhanced metrology”. In: vol. 9980. Bellingham, WA, Sept. 2016. DOI: 10.1117/12.2237355.
- [51] M. J. Quinn. *Parallel Programming in C with MPI and openMP*. 1st. McGraw-Hill, 2003.
- [52] H. Rabitz et al. “Chemistry - whither the future of controlling quantum phenomena?”. In: *Science (N.Y.)* 288 (June 2000), pp. 824–8. DOI: 10.1126/science.288.5467.824.
- [53] A. C. Richard. “Quantum-mechanical lossless beam splitter: SU(2) symmetry and photon statistics”. In: *Phys. Rev. A* 40.3 (1989), pp. 1371–1384. DOI: 10.1103/PhysRevA.40.1371.
- [54] T. Rückstieß et al. “Exploring parameter space in reinforcement learning”. In: *Paladyn* 1.1 (Mar. 2010), pp. 14–24. DOI: 10.2478/s13230-010-0002-4.
- [55] R. Schnabel et al. “Quantum metrology for gravitational wave astronomy”. In: *Nat. Commun.* 1.1 (2010), p. 121. DOI: 10.1038/ncomms1122.
- [56] T. A. Schonhoff and A. A. Giordano. “Fundamentals of estimation theory”. In: *Detection and Estimation Theory and its Applications*. Pearson Prentice Hall, 2006. Chap. 10, pp. 267–316.
- [57] W. R. Scott. *Group Theory*. Dover Books on Mathematics. Englewood Cliffs, NJ: Dover, 2012.
- [58] R. Slansky. “Group theory for unified model building”. In: Amsterdam, 1981.
- [59] R. Storn and K. Price. “Differential evolution: a simple and efficient heuristic for global optimization over continuous spaces”. In: *J. Global Optim.* 11.4 (Dec. 1997), pp. 341–359. DOI: 10.1023/A:1008202821328.

- [60] G. S. Summy and D. T. Pegg. “Phase optimized quantum states of light”. In: *Opt. Commun.* 77.1 (June 1990), pp. 75–79. DOI: [https://doi.org/10.1016/0030-4018\(90\)90464-5](https://doi.org/10.1016/0030-4018(90)90464-5).
- [61] R. S. Sutton and A. G. Barto. “The reinforcement learning problem”. In: *Reinforcement Learning: An Introduction*. 2nd ed. Adaptive Computation and Machine Learning. Massachusetts: MIT, 2017. Chap. 3, pp. 43–78.
- [62] M. Szczykulska, T. Baumgratz, and A. Datta. “Reaching for the quantum limits in the simultaneous estimation of phase and phase diffusion”. In: *Quantum Sci. Technol.* 2.4 (2017).
- [63] T. C. Urdan. *Statistics in Plain English*. New York: Lawrence Erlbaum Assoc., 2005. ISBN: 9780805852417.
- [64] J. P. Verma and A. S. G. Abdel-Salam. *Testing Statistical Assumptions in Research*. Hoboken, NJ: Wiley, 2019.
- [65] P. A. Vikhar. “Evolutionary algorithms: a critical review and its future prospects”. In: *Proc. 2016 International Conference on Global Trends in Signal Processing, Information Computing and Communication (ICGTSPICC)*. Los Alamitos, CA: IEEE, Dec. 2016, pp. 261–265. DOI: [10.1109/ICGTSPICC.2016.7955308](https://doi.org/10.1109/ICGTSPICC.2016.7955308).
- [66] T.J. Volkoff. “Optimal and near-optimal probe states for quantum metrology of number-conserving two-mode bosonic Hamiltonians”. In: *Phys. Rev. A* 94 (Oct. 2016). DOI: [10.1103/PhysRevA.94.042327](https://doi.org/10.1103/PhysRevA.94.042327).
- [67] D. Vrajitoru and W. Knight. “Deterministic analysis of algorithms”. In: *Practical Analysis of Algorithms*. New York: Springer, 2014. Chap. 5, pp. 169–293. DOI: [10.1007/978-3-319-09888-3_5](https://doi.org/10.1007/978-3-319-09888-3_5).
- [68] D. Vrajitoru and W. Knight. “Introduction”. In: *Practical Analysis of Algorithms*. New York: Springer, 2014. Chap. 1, pp. 1–7. DOI: [10.1007/978-3-319-09888-3_5](https://doi.org/10.1007/978-3-319-09888-3_5).

- [69] R. V. Wagoner. “Scalar-tensor theory and gravitational waves”. In: *Phys. Rev. D* 1.12 (1970), pp. 3209–3216. DOI: 10.1103/PhysRevD.1.3209.
- [70] E. Waks et al. “Quantum cryptography with a photon turnstile”. In: *Nature* 420.6917 (2002), pp. 762–762. DOI: 10.1038/420762a.
- [71] M. Wiering and M. van Otterlo. “Reinforcement learning and Markov decision processes”. In: *Reinforcement Learning: State-of-the-Art*. Ed. by M. Wiering and M. van Otterlo. Vol. 12. Adaptation, Learning, And Optimization. New York: Springer, 2012.
- [72] P. Winker and M. Gilli. “Applications of optimization heuristics to estimation and modelling problems”. In: *Comput. Stat. Data Anal.* 47.2 (2004), pp. 211–223. ISSN: 0167-9473. DOI: 10.1016/j.csda.2003.11.026.
- [73] H. M. Wiseman. “Adaptive phase measurements of optical modes: going beyond the marginal Q distribution”. In: *Phys. Rev. Lett.* 75.25 (Dec. 1995), pp. 4587–4590. DOI: 10.1103/PhysRevLett.75.4587.
- [74] H. M. Wiseman and R.B. Killip. “Adaptive single-shot phase measurements: The full quantum theory”. In: *Phys. Rev. A* 57.3 (Mar. 1998), pp. 2169–2185. DOI: 10.1103/PhysRevA.57.2169.
- [75] H. M. Wiseman et al. “Adaptive measurements in the optical quantum information laboratory”. In: *IEEE J. Sel. Top. Quantum Electron.* 15.6 (Nov. 2009), pp. 1661–1672. DOI: 10.1109/JSTQE.2009.2020810.
- [76] H. Wiseman and R. B. Killip. “Adaptive single-shot phase measurements: a semiclassical approach”. In: *Phys. Rev. A* 56 (July 1997). DOI: 10.1103/PhysRevA.56.944.
- [77] X. Yan and X. G. Su. “Simple linear regression”. In: *Linear Regression Analysis : Theory and Computing*. Singapore: SG: World Sci., 2009. Chap. 2, pp. 9–40.
- [78] E. Zahedinejad, S. Schirmer, and B. C. Sanders. “Evolutionary algorithms for hard quantum control”. In: *Phys. Rev. A* 90.3 (Sept. 2014), p. 032310. DOI: 10.1103/PhysRevA.90.032310.

- [79] H. R. Zimmerman et al. “Nightly biting cycles of malaria vectors in a heterogeneous transmission area of eastern Amazonian Brazil”. In: *Malar. j.* 12 (July 2013), p. 262. DOI: 10.1186/1475-2875-12-262.

Appendix A

Evolutionary algorithm for variable beam-splitter reflectivity estimation

The appendix contains excerpts from the software documentation available on GitHub. The repository link is given below.

A.1 About the software

The software is created by modifying the library initially developed by Pantita Palittapongarnpim and Peter Wittek for phase estimation [49]. The modified library is designed to assist researchers in the construction of policy-search algorithms for estimating unknown beam-splitter reflectivity in a Mach-Zehnder interferometer. The software's modules can be altered for solving other quantum control problems. Given below are some of the main features of the software [49].

- Library in C++
- MPI support
- VSL and GPU support for random number generation

- Includes a module for differential evolution (DE)
- Includes uniform and clustered method of initializing solution candidates
- Includes access to user-specified accept-reject criteria

A.2 Copyright and license

This is free software made available under the GNU GENERAL PUBLIC LICENSE, which means you can share, modify, and redistribute this software. While we endeavour to make this software as useful and as error-free as possible, we cannot make any such guarantee, and the software is hence released without any warranty [49].

A.3 Download link

The download and installation instructions can be found here: https://github.com/thehamzaq/Variable-reflectivity_beam-splitter_estimation.