

Constraining the Properties of Sagittarius A* with the Event Horizon Telescope

by

Carlos Wang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Science
in
Physics

Waterloo, Ontario, Canada, 2015

© Carlos Wang 2015

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Sgr A* is a primary target for testing strong gravity, jet formation and black hole accretion for the the event horizon telescope (EHT). Motivated by continued improvements of the observations from the EHT, including newly reported closure phase measurements, this thesis reports three research directions.

Two of these constrain strong gravity parameters with the EHT. First, applying the new closure phase information from the EHT improves quasi-Kerr deviation constraints to 1σ limits on spin magnitude $a_* = 0.05_{-0.05}^{+0.30}$, inclination $\theta = 70.0_{-0.05}^{+12.5^\circ}$, and quasi-Kerr deviation $\epsilon = 0.5_{-0.25}^{+0.45}$. Position angle (ξ) limits are not reported because of a closure phase bias due to interstellar Galactic scattering. While ϵ appears to favour a non-Kerr-like black hole, the posterior joint probability distribution obtained for a_* and ϵ has a 1σ region allow for both non-Kerr and Kerr black hole. Future constraints were also simulated with the addition of new stations to the EHT, producing tight 3σ limits: $a_* = 0.150_{-0.005}^{+0.004}$, $\theta = 60.01_{-0.06}^{+0.09^\circ}$, $\xi = 159.99^\circ \pm 0.04^\circ$, and $\epsilon = 0 \pm 0.005$, representing an improvement of two orders of magnitude. Second, the six year (2007 to 2013) duration of EHT observations enable a search for disk misalignment and precession. Both Lense-Thirring driven and arbitrary precession models were explored. No discernible precession was found from 2007 to 2013. This results in 2σ limits on the black hole spin and accretion flow size: $a \lesssim 0.575$ and $r_o \gtrsim 300M$ if the precession is driven by Lense-Thirring torques.

The third direction is the development of a new analysis pipeline, THEMIS, to replace the existing radiatively inefficient accretion flow (RIAF) analysis software used thus far to test models with EHT data. THEMIS is capable of easily integrating new observations and models of Sgr A*, as well as employing Markov chain Monte Carlo (MCMC) methods to make higher dimensional parameter space explorations feasible. All individual components have been successfully validated, and integrated, and are now ready for global testing and optimization. Following this, THEMIS is planned to be used to constrain two new RIAF parameters: the disk height ratio and sub-Keplerian parameter.

Acknowledgements

This thesis was not made possible without the encouragement and immense support from my supervisor, Dr. Avery Broderick. He is an inspiration and role model not only at a professional level but also by character.

I am also extremely thankful, and grateful, for the moral support I have received from Rohan Jayasundera, my friends, my family, my two brothers Bobby and Peter, my nieces and nephews, and especially my Parents. They are precious, reliable, and played a valuable role not only by mentorship but by inspiration during both difficult and fortunate times.

Last, but not least, I would like to thank the University of Waterloo and its staff for their continued support and any financial backing I have received from respective sources throughout my graduate experience.

Dedication

This thesis is dedicated to the beauty of the universe, and from within, to those I love, the continual growth of humanity and the spirit of exploration.

Table of Contents

List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Interferometric Data	4
1.1.1 Complex Visibilities	6
1.1.2 Visibility Amplitudes	8
1.1.3 Closure Phase	8
1.1.4 Closure Amplitudes	9
1.2 Event Horizon Telescope	11
1.2.1 2007-2013 Observations	12
1.3 Ancillary Data	16
1.4 Modelling Sgr A* using the EHT	17
1.4.1 Previous Work on Simulating Sgr A* for the EHT	19
1.5 Research Cases for Sgr A* using the EHT	23
1.6 Conventions	24
2 Quasi-Kerr Model Deviation	25
2.1 Quasi-Kerr Metric	27
2.2 Adaptive parameter estimation	29

2.3	Current Limits on Metric Deviation	31
2.4	Future Limits on Metric Deviation	35
3	Black Hole Precession	37
3.1	Lense-Thirring Precession in RIAFs	39
3.2	Signatures of Rigid Disk Rotation in EHT Observations	44
3.3	Search for Lense-Thirring Precession using EHT Observations	48
3.4	Search for Any Precession using EHT Observations	56
3.5	Conclusions	62
4	Next Generation Analysis Pipeline	64
4.1	THEMIS Overview	67
4.1.1	Theoretical Image Generation	69
4.1.2	Data Remapper	69
4.1.3	Posterior Probability Estimator	72
4.2	Future Prospects	73
5	Conclusion	74
	References	76

List of Tables

1.1	Radio Array/Telescopes that are part of the EHT project.	11
1.2	Grouping datasets from the EHT into epochs.	15

List of Figures

1.1	Simple two element interferometer diagram.	5
1.2	Model comparison with the spectral energy density plot for Sgr A*.	13
1.3	Global position and uv plane coverage of the the EHT used for Sgr A*.	14
1.4	Diagram of forces involved in RIAF.	20
2.1	Varying quasi-Kerr deviation parameter to the black hole silhouette.	28
2.2	Inclination and position angle comparison with the quasi-Kerr deviation.	32
2.3	Inclination and position angle comparison separated by VM and CP.	32
2.4	Inclination and quasi-Kerr deviation with spin comparison.	34
2.5	ξ and ϵ comparison with a_* using the future simulated EHT data.	36
3.1	r_w for thick disks as a function of spin with varying $\sin \Theta$	41
3.2	Angle definitions for detecting rigid body precession.	45
3.3	Path of a precessing disk in θ - ξ space with varying ϑ	46
3.4	Constraints of angle parameters for LT precession.	50
3.5	Constraints of a and r_o for LT precession.	51
3.6	Triangular plot of 2D constraints for LT precession study.	52
3.7	Angle structure degeneracy of 2D pair plots from LT precession study.	53
3.8	Constraints of a vs. r_o for LT precession.	55
3.9	Constraints of angle parameters for any rigid-body precession.	57
3.10	Constraints of T for any rigid-body precession.	58

3.11	Triangular plot of 2D constraints for any rigid-body precession.	60
3.12	Angle structure degeneracy of 2D angle plots from any precession study.	61
4.1	Theoretical images of varying h/r of a RIAF model for Sgr A*	65
4.2	Flow diagram overview of the THEMIS software.	68
4.3	Theoretical image generation component of THEMIS.	70
4.4	Data remapper component of THEMIS.	71
4.5	Posterior probability estimator component of THEMIS.	73

Chapter 1

Introduction

Black holes are a nonlinear prediction of general relativity. They are implicated in the most energetic phenomena in the universe, ranging from active galactic nuclei to microquasars and gamma-ray bursts. At the same time, as manifestly nonlinear general relativistic solutions they impart the details of strong gravity on astrophysical phenomena. Thus, the observational and theoretical study of astrophysical black holes offers windows into the highest-energy processes in astronomy and fundamental gravitational physics.

Astrophysical black holes come in two varieties, separated by mass. Stellar mass black holes, with masses from $\gtrsim 3\text{--}10^2 M_\odot$ are believed to be the endpoints of massive star evolution. These are thought to be responsible for microquasars, black hole X-ray binaries, long gamma-ray bursts, and associated phenomena. The latter, long gamma-ray bursts, are associated with starburst regions, strengthening the association between stellar mass black holes and massive stars.

Supermassive black holes have masses larger than $10^5 M_\odot$, reaching $10^{10} M_\odot$, and inhabit the centres of nearly all, if not all, galaxies. These power the active galactic nuclei present in roughly 10% of galaxies in the nearby universe. In contrast to stellar mass black holes, the formation and growth supermassive black holes is not well understood. Their phenomenology, however, is well studied by virtue of the large number of active galactic nuclei observed.

Black holes with intermediate masses, i.e., between $10^2 M_\odot$ and $10^5 M_\odot$, are anticipated by many theories of supermassive black hole formation [66]. However, none have not been unambiguously discovered – though such objects have been suggested to explain ultraluminous X-ray sources.

Gravitationally, all classes of black holes are identical, distinguished only by an overall scale that depends linearly on mass: GM/c^2 . Thus, tests of general relativity performed on black holes of one mass are applicable, in principle, to black holes of all masses. Astrophysically, their scale-free nature is broken by their environment, the details of radiative transfer, and baryonic interactions in the surrounding gas. Nevertheless, galactic microquasars do exhibit many of the same phenomena as quasars, their supermassive analogs, suggesting that their scaling does approximately persist.

Black holes play a key role in shaping their surrounding environment through “feedback” of various kinds, depending on the black mass. Stellar mass black holes impart most of their impact during formation in the form of supernovae, which disperses the gas from star forming regions. Supermassive black holes can dominate the evolution of their host galaxies and galaxy clusters through both their high luminosities and generation of powerful outflows, driven by processes near the event horizon.

Gravitational potential energy liberated within accretion disks fed by inflowing gas from a variety of potential sources power the large luminosities (quasars), accretion winds, and relativistic jets (radio AGN). The resulting deposition of energy and momentum regulates the gas supply on scales ranging from kpc to Mpc, and therefore closes a feedback loop that couples the near-horizon environment to that of the host galaxy. From this process, supermassive black holes are believed to grow and impact not only their host galaxies but in many instances produce noticeable imprints on their surrounding galaxies as well [47, 55].

The physics of strong gravity, and therefore the spacetime structure of black holes, is rigorously described by general relativity. General relativity has been exquisitely tested within the perturbative regime, where it makes only small corrections to Newtonian gravity [73]. However, large gravitational perturbations are accessible via supermassive black holes at near horizon scales. That is, by probing the accretion and jet launching phenomena strong gravity can be tested at large gravitational perturbations.

Probing horizon-scales is fundamentally complicated by degeneracies between the location of the emitted signals and aspects of strong gravity that dictate its propagation to distant observers. This process is simplified greatly by directly resolving the horizon-scale structure of the near-black hole emission region. The Event Horizon Telescope (EHT) is set to be the first instrument to spatially resolve accretion disks of supermassive black holes, thereby decisively allowing the direct study of the strong gravity regime. This is being accomplished by mm-wavelength Very Long Baseline Interferometry (mm-VLBI) enabled by a global collaboration of observatories and researchers. Here we describe a number of current and future tests of general relativistic and high-energy astrophysical phenomena

enabled by the unique capabilities of the EHT.

1.1 Interferometric Data

VLBI achieves high resolutions by performing interferometric experiments [70] over baselines exceeding 10^2 km. In this way interferometers utilize an array of telescopes to emulate a larger telescope with an aperture size equivalent to the telescope pair with the largest separation distance (baseline). This enables the construction of an Earth-sized telescope, though without the benefit of additional collecting area and typically poor coverage in the u - v (or uv) plane, the Fourier conjugate plane to the angular position on the sky. The limited collecting area of the mm-wavelength telescopes participating in the EHT means that mm-VLBI is currently limited to bright sources, i.e., sources with brightness temperatures exceeding 10^8 K, and hence active black holes. Here we discuss the direct interferometric data available and associated practical difficulties.

A basic diagram of a two antenna interferometer is shown in Figure 1.2. The baseline vector between the two elements is \mathbf{b} and $\hat{\mathbf{s}}$ is denoted as the direction at which the observers are pointing towards. The geometrical time delay between receiver 1 and receiver 2 is then

$$\tau_g = \frac{\mathbf{b} \cdot \hat{\mathbf{s}}}{c} = \frac{|\mathbf{b}| \cos \theta}{c} = \frac{b \cos \theta}{c}, \quad (1.1)$$

where θ is the angle between \mathbf{b} and $\hat{\mathbf{s}}$. The voltage output from receiver 1 and 2, $R_1(t)$ and $R_2(t)$ respectively, at some time t from the signal is then:

$$\begin{aligned} R_1(t) &= V \cos(2\pi\nu t - 2\pi\nu\tau_g) \\ &= V \cos(2\pi\nu t) \cos(2\pi\nu\tau_g) + V \sin(2\pi\nu t) \sin(2\pi\nu\tau_g) \\ R_2(t) &= V \cos(2\pi\nu t), \end{aligned} \quad (1.2)$$

where V is the signal voltage reaching the collectors and ν is the signal frequency. The correlator multiplies the two receiver outputs and takes the time average of the result for a duration T to yield:

$$\begin{aligned} r(\mathbf{b}) &= \langle R_1(t)R_2(t) \rangle \\ &= \frac{1}{T} \int_{-T/2}^{T/2} V^2 [\cos(2\pi\nu t) \cos(2\pi\nu t) \cos(2\pi\nu\tau_g) \\ &\quad + \cos(2\pi\nu t) \sin(2\pi\nu t) \sin(2\pi\nu\tau_g)] \\ &= V^2 \cos(2\pi\nu\tau_g) \left(\frac{1}{2} + \frac{\sin(2\pi\nu T)}{4\pi\nu T} \right) \\ &\approx \frac{V^2}{2} \cos(2\pi\nu\tau_g) = \frac{V^2}{2} \cos\left(\frac{2\pi b \cos \theta}{\lambda}\right), \end{aligned} \quad (1.3)$$

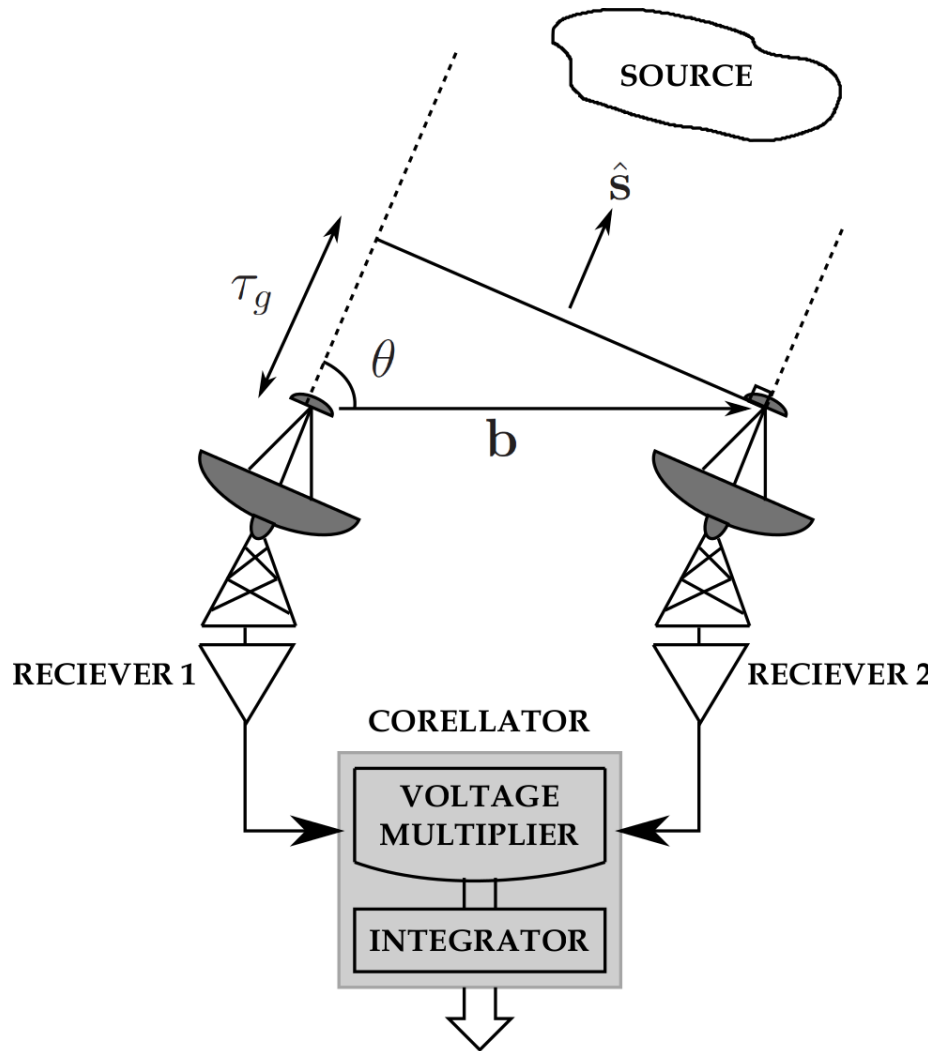


Figure 1.1: Simple two element interferometer diagram. The two antenna receivers are pointing at an angle θ , represented by the unit vector $\hat{\mathbf{s}}$. The baseline between the two receivers is \mathbf{b} and the geometrical signal delay to receiver 1 is denoted by τ_g .

where it had been assumed that $T \gg (2\pi\nu)^{-1}$ to remove the sinc term. Note that at mm wavelengths typically $T \approx 10$ s, set by the atmospheric coherence timescale, and thus $2\pi\nu T \approx 2 \times 10^{13}$. Similarly, the sine version of Equation 1.3 can be obtained by adding a $\frac{\pi}{2}$ delay to the second receiver: $R_2(t) = V \sin(2\pi\nu t)$, and repeating the correlation process:

$$r_s(\mathbf{b}) = \frac{V^2}{2} \sin\left(\frac{2\pi b \cos \theta}{\lambda}\right). \quad (1.4)$$

The final term of Equation 1.3 shows that the correlator result is independent of time, however it is dependent on the baseline and direction of the source. Assuming V , b , and λ are constant, then as the source direction (θ) changes the correlator output varies sinusoidally, generating a sinusoid pattern called a fringe. The fringe phase $\phi(\theta)$ is therefore:

$$\phi(\theta) = \frac{2\pi b \cos \theta}{\lambda}, \quad (1.5)$$

and subsequently

$$\frac{d\phi}{d\theta} = -\frac{2\pi b \sin \theta}{\lambda}. \quad (1.6)$$

It can be deduced from Equation 1.5 that a change in the fringe by 2π , i.e. the fringe period $\Delta\phi = 2\pi$, leads to a change $\Delta\theta = \lambda/|b| \sin \theta$. It is therefore favourable to make $\Delta\theta$ as small as possible to increase the sensitivity of the fringe pattern with respect to the source direction. Maximizing $b \sin \theta$, the projected baseline, effectively increases the angular resolution of the interferometer.

1.1.1 Complex Visibilities

The Equations 1.3 and 1.4 thus far are for a point source. To implement this for a brightness distribution (a sum of independent point sources) of a particular wavelength in the sky, $I_\lambda(\hat{\mathbf{s}})$, the output from the correlator is integrated for all directions of $\hat{\mathbf{s}}$:

$$r_{c,\lambda}(\mathbf{b}) = \int I_\lambda(\hat{\mathbf{s}}) \cos\left(\frac{2\pi \mathbf{b} \cdot \hat{\mathbf{s}}}{\lambda}\right) d\mathbf{s}. \quad (1.7)$$

Again, applying the same method to the phase delayed correlation, $r_s(\mathbf{b})$, yields:

$$r_{s,\lambda}(\mathbf{b}) = \int I_\lambda(\hat{\mathbf{s}}) \sin\left(\frac{2\pi \mathbf{b} \cdot \hat{\mathbf{s}}}{\lambda}\right) d\mathbf{s}. \quad (1.8)$$

The complex visibility is then defined by $\tilde{\mathcal{V}} \equiv r_{c,\lambda} - ir_{s,\lambda} = Ae^{-i\Phi}$, where

$$A = \sqrt{r_{c,\lambda}^2 + r_{s,\lambda}^2} \quad (1.9)$$

and

$$\Phi = \tan^{-1} \left(\frac{r_{s,\lambda}}{r_{c,\lambda}} \right) \quad (1.10)$$

are the visibility magnitude and phase, respectively. Therefore the chief interferometric observable, the complex visibility, is given by the spatial Fourier transform of the image:

$$\tilde{\mathcal{V}}_\lambda = \int I_\lambda(\hat{\mathbf{s}}) \exp \left(-i \frac{2\pi \mathbf{b} \cdot \hat{\mathbf{s}}}{\lambda} \right) d\mathbf{s}. \quad (1.11)$$

Let \mathbf{b} be described in spatial frequency pairs (u, v) (that is, the displacement of the baseline in units λ) and $\hat{\mathbf{s}}$ described in angular coordinates (α, β) , where the coordinate system's origin crosses the projected source center. Equation 1.11 can then be rewritten in a more familiar form as

$$\tilde{\mathcal{V}}_\lambda(u, v) = \iint I_\lambda(\alpha, \beta) \exp(-i2\pi(u\alpha + v\beta)) d\alpha d\beta. \quad (1.12)$$

In principle the brightness distribution (I_λ), or source image, can be recovered via an inverse Fourier transform of the visibilities:

$$I_\lambda(\alpha, \beta) = \iint \tilde{\mathcal{V}}_\lambda(u, v) \exp(i2\pi(u\alpha + v\beta)) dudv. \quad (1.13)$$

In practice, only a small, discrete subset of the entire range of the uv plane can be observed. Thus a discrete version of the inverse Fourier transform is used to compute an incomplete source image, known as a dirty image. With sufficient data, further processing (CLEAN, Bispectral maximum entropy reconstruction, Gerchberg-Saxton Method, etc.) can be applied to the dirty image to improve its quality and similarity to the actual source image.

The number of complex visibilities increases quadratically as the number of pairs of telescopes increase, that is the number goes as $N(N - 1)/2$ for N stations for a single observing instance. Ideally, stations are arranged to produce the maximum number of independent baselines, with the largest baseline lengths dictating the highest resolution in the image.

1.1.2 Visibility Amplitudes

Note that the complex visibility from Equation 1.12 in the previous section (Section 1.1.1) did not take into account of phase or gain errors that may be due to the observing apparatus and signal medium. In order to account for this, an additional complex error term \mathcal{G} is introduced:

$$\tilde{\mathcal{G}} \equiv Ge^{i\Phi^G}, \quad (1.14)$$

where G is the usually slowly varying gain error from the telescope and atmospheric opacity, and Φ^G is the typically rapidly varying phase error from atmospheric turbulence. Therefore the complex visibility between stations 1 and 2 will in principle observe $\tilde{\mathcal{V}}_{12,\lambda}^o$:

$$\begin{aligned} \tilde{\mathcal{V}}_{12,\lambda}^o &= \tilde{\mathcal{G}}_1 \tilde{\mathcal{G}}_2^* \tilde{\mathcal{V}}_{12,\lambda}^a \\ &= G_1 G_2 e^{-i(\Phi_2^G - \Phi_1^G)} \tilde{\mathcal{V}}_{12,\lambda}^a, \end{aligned} \quad (1.15)$$

where $\tilde{\mathcal{V}}_{12,\lambda}^a$ is the actual complex visibility from station 1 to 2. The resulting observed visibility magnitude is then:

$$\left| \tilde{\mathcal{V}}_{12,\lambda}^o \right| = G_1 G_2 A. \quad (1.16)$$

Similar to the complex visibilities, the number of visibility magnitudes increases as the number of pairs of telescopes increase, that is the number goes as $N(N - 1)/2$ for N stations for a single observing instance. The gain errors G_1 and G_2 can be removed from the measured visibility magnitude through careful calibration of the gains for each station.

1.1.3 Closure Phase

VLBI differs from connected element interferometry (e.g., that done at CARMA, ALMA, SMA, etc.) in that the rapidly changing atmospheric phase errors are not correlated between antennas, and therefore cannot be easily accounted for. Thus, as mentioned in Section 1.1.2, it is difficult to directly retrieve the source phase information from the complex visibilities. This is problematic since the visibility phases contain the majority of the information regarding image structure. Nevertheless, a quantity that contains much of the information in the visibility phases can be recovered, called a closure phase.

The visibility phase for two station array is

$$\psi_{12} = \Phi_{12} + \Phi_2^G - \Phi_1^G. \quad (1.17)$$

For a three station array, three pairs of visibility phases can be obtained,

$$\begin{aligned}
 \psi_{12} &= \Phi_{12} + \Phi_2^G - \Phi_1^G \\
 \psi_{23} &= \Phi_{23} + \Phi_3^G - \Phi_2^G \\
 \psi_{31} &= \Phi_{31} + \Phi_1^G - \Phi_3^G,
 \end{aligned}
 \tag{1.18}$$

and a closure phase for these stations, Ψ_{123} , is defined to be the phase of the triple product of the complex visibilities of the three station pairs, or:

$$\Psi_{123} \equiv \psi_{12} + \psi_{23} + \psi_{31} = \Phi_{12} + \Phi_{23} + \Phi_{31}.
 \tag{1.19}$$

The key feature of the closure phase from Equation 1.19 is that the atmospheric phase errors “close”, i.e., they are removed identically by careful construction, and what is left is the sum of the true complex phase values for each station pair. Note that since closure phases address only the phase of the incoming wave they are insensitive to any systematic errors that modify only the visibility amplitudes, making these easier to calibrate in practice.

A limitation is that a minimum of three stations are required to compute closure phases, and the number of independent closure phases goes as $(N - 1)(N - 2)/2$ given N stations for a single observing instance. Another limitation is that the number of closure phases is always less than the number of visibility phases, that is $((N - 1)(N - 2)/2 < N(N - 1)/2)$, but since both closure and visibility phases are increasing at the order of N^2 , then the ratio between the two $(N - 2)/N \rightarrow 1$ as $N \rightarrow \infty$. Therefore the fraction of phase information recovered approaches 100% as the number of stations in the array becomes large. Thus, arrays with many stations are capable of “self-calibration”, in which the only phase information used during image reconstruction are closure phases.

1.1.4 Closure Amplitudes

A similar approach used for the closure phases described in Section 1.1.3 can be used to remove the gain errors associated with each station altogether by introducing a closure

amplitude quantity. The closure amplitude for a four station quadrilateral is defined as:

$$\begin{aligned}
 A &= \frac{|\tilde{\mathcal{V}}_{12,\lambda}| |\tilde{\mathcal{V}}_{34,\lambda}|}{|\tilde{\mathcal{V}}_{13,\lambda}| |\tilde{\mathcal{V}}_{24,\lambda}|} \\
 &= \frac{G_1 G_2 G_3 G_4 |\tilde{\mathcal{V}}_{12,\lambda}^a| |\tilde{\mathcal{V}}_{34,\lambda}^a|}{G_1 G_2 G_3 G_4 |\tilde{\mathcal{V}}_{13,\lambda}^a| |\tilde{\mathcal{V}}_{24,\lambda}^a|} \\
 &= \frac{|\tilde{\mathcal{V}}_{12,\lambda}^a| |\tilde{\mathcal{V}}_{34,\lambda}^a|}{|\tilde{\mathcal{V}}_{13,\lambda}^a| |\tilde{\mathcal{V}}_{24,\lambda}^a|},
 \end{aligned} \tag{1.20}$$

which remove the gain error terms from the telescope. This quantity does not remove the phase error and is limited to requiring at least four stations. More importantly, it loses all information about the absolute flux normalization, supplying only the ratios of visibility amplitudes. The number of closure amplitudes for a given observing instance is $N((N - 1)/2 - 1)$, where N is the number of stations. Similar to closure phases, as the number of stations approaches infinity, the magnitude information approaches 100%.

Closure amplitudes from the EHT have yet to be published, though are implicitly used in forthcoming EHT publications [44].

1.2 Event Horizon Telescope

The Event Horizon Telescope (EHT) is a collection of millimeter and submillimeter observatories from various locations on the Earth. Each station is equipped with Very Long Baseline Interferometry (VLBI) instruments, allowing for horizon scale resolved images of supermassive black holes at 1.3 mm (230 GHz) and 0.87 mm (345 GHz). As of 2015, the following observatories shown in Table 1.1 are part of the EHT project.

Radio Array/Telescopes (as of 2015)	Abbrev.
Arizona Radio Observatory/Submillimeter-wave Astronomy	ARO/SMT
Atacama Large Millimeter/submillimeter Array	ALMA
Atacama Pathfinder EXperiment	APEX
Atacama Submillimeter Telescope Experiment	ASTE
Combined Array for Research in Millimeter-wave Astronomy	CARMA
Caltech Submillimeter Observatory	CSO
Institut de Radioastronomie Millimetrique 30m	IRAM
James Clerk Maxwell Telescope	JCMT
Plateau de Bure interferometer	PdB
South Pole Telescope	SPT
The Large Millimeter Telescope	LMT
The Submillimeter Array	SMA

Table 1.1: Radio Array/Telescopes that are part of the EHT project. The Plateau de Bure interferometer currently has not performed any observing run with the EHT. While CARMA is part of the EHT collaboration, it is now defunct.

The two main targets of the EHT are Sagittarius A* (Sgr A*), the supermassive black hole at the center of the Milky Way, and M87 because they both are accessible via mm wavelengths and have an apparent angular size that is larger than $20\mu\text{as}$, the approximate resolution of the EHT. Here we focus exclusively on the former.

The silhouette formed by the horizon on the surrounding plasma emission in Sgr A* is roughly $53 \pm 2\mu\text{as}$ [42]. Sgr A* is typical of supermassive black holes at centers of galaxies: spectral and polarization data suggests Sgr A* is vastly underluminous compared with amount of matter it has available to accrete, and is therefore believed that it is representative of roughly 90% of the black holes not presently in an active state [10].

Figure 1.2 shows the Spectral Energy Distribution (SED) for Sgr A*. At long wavelengths it displays an inverted, nearly flat, radio spectrum above a GHz (optically thick)

and then from mm to sub-mm wavelengths it transitions to a falling SED, characteristic of an optically thin source. Thus Sgr A* is consistent with a structured, self-absorbed, nonthermal synchrotron source. This particular SED hence places strong constraints on particular models of Sgr A*. At mm wavelengths the blurring due to interstellar electron scattering is subdominant [5].

All sites within the EHT currently can view Sgr A*, though not necessarily coincidentally.¹ Figure 1.3 is a snapshot of a movie illustrating the telescopes (red points) in the EHT for observing Sgr A* on the globe, as viewed from Sgr A*. The baselines are shown between each station pair on the globe (red line) and the uv coverage for each baseline is traced (blue line) for a 24 hour observing run, assuming it is observed at a wavelength of 1.3mm. The red dots on the blue trace lines indicate the current uv position for each station pair at that particular instant. Note that the stations nearly span the entire circumference of the Earth.

1.2.1 2007-2013 Observations

For this thesis the most recently reduced observations of the EHT are used, ranging from 2007 to 2013. The types of data observed are visibility magnitudes (see Section 1.1.2) and closure phases (see Section 1.1.3), described in [21], [23] and [24]. Details of the observations, calibration and data processing can be found in each respective reference. Here we provide only a short summary of each.

Visibility Magnitudes: 2007

Sgr A* was observed for visibility magnitudes on 2007 April 11-12, corresponding to days 101 and 102 of the year 2007. [21] reports 19 visibility amplitudes were obtained on the CARMA-SMT and JCMT-SMT baselines, and an upper limit on April 11 along the JCMT-CARMA baseline. A single CARMA dish was used this year; because much more constraining detections were obtained on the JCMT-CARMA baseline in 2009, this upper limit is excluded from further consideration. Signal-to-noise ratios were typically 8 to 4 for short to long baselines, respectively.

Visibility Magnitudes: 2009

Sgr A* was observed for visibility magnitudes on 2009 April 5-7, corresponding to days 95-97 of the year 2009. [23] reports a total of 54 visibility magnitudes were obtained on the CARMA-SMT and JCMT-SMT on all days and to both of the JCMT-CARMA

¹The Greenland telescope, expected to become available sometime in 2017, will be unable to view Sgr A* though will offer additional baselines for M87.




Figure 1.2: Comparison between the SED of the most probable accretion model with the observed SED of Sgr A* using various sources, taken from [10]. The orange circles are from [75] and references therein, with the variability indicated by the error bar. The yellow squares are simultaneous flux measurements from [53], with the intrinsic measurement error indicated by the error bar. The green bar is the observed from the full CARMA array [21, 23] which is believed to be in a quiescent state [9] and has error bars indicative of the range of measured values. The long-dash blue and short-dash red lines shows the contribution from nonthermal and thermal component of the electron population, respectively. The thick black line shows an accretion model fit. The dotted lines show the typical range of model spectra.

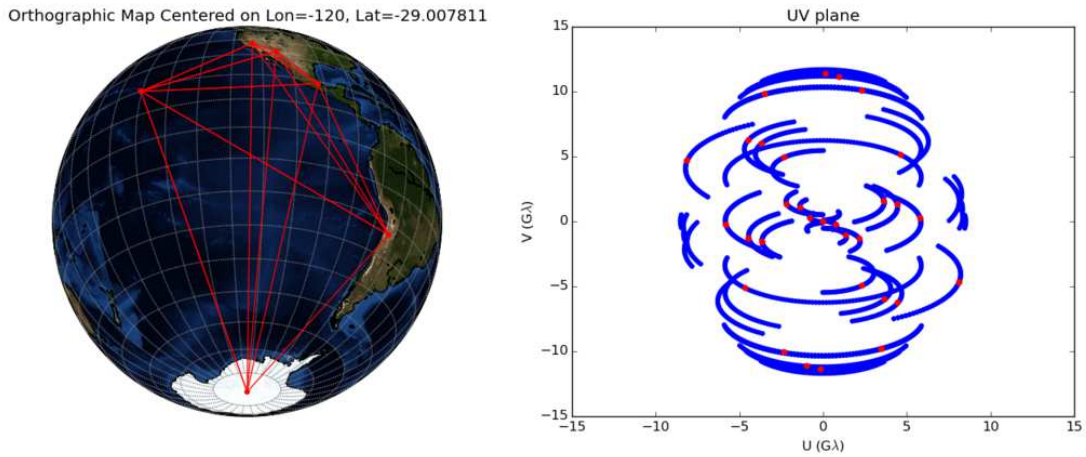


Figure 1.3: Illustration of the EHT baselines for Sgr A* on the globe (left) and their corresponding uv plane coverage (right), at a wavelength of 1mm. In this particular example, the stations (some of which are hidden behind the globe) used are for Sgr A*. The red nodes in the left figure are various EHT stations and the red lines are their corresponding baselines. The blue lines in the right figure trace each station's uv coverage and the red dots indicate the current uv position for each station for this instance. The uv plane and globe are plotted as viewed from Sgr A* (looking down on Earth), so the u values go from left to right. The images are a snapshot of a full 24 hour movie, courtesy of Dr. Laura Vertatschitsch and the EHT collaboration.

baselines on days 96-97. Two CARMA dishes were used this year and each operated as independent stations. Signal-to-noise ratios were typically 17 to 5 for short to long baselines, respectively. This is a significant improvement from the observations of 2007.

Closure Phases: 2009-2013

[24] reports 181 nontrivial closure phases and 233 trivial closure phases, using triplets of SMT (Arizona), CARMA (California), and SMA-JCMT-CSO (Hawaii), over the span of 4 years from 2009 to 2013. A total of 13 observing days produced the nontrivial closure phases. Day 94 of 2011 has been omitted due to apparent flaring, as discussed in [24].

Epochs

The analyses performed in this thesis groups the EHT datasets into epochs based on the observed day (with the exception on year 2007) and data type. Table 1.2 lists the epoch indexing used for each particular observing run, totalling 17 epochs. At the time of this thesis the visibility amplitudes from 2012 and 2013 have not be calibrated, precluding their inclusion.

Epoch index	Year	Day(s)	Data Type (VM or CP)	Ref.
1	2007	101-102	VM	[21]
2	2009	95	VM	[23]
3	2009	96	VM	[23]
4	2009	97	VM	[23]
5	2009	93	CP	[24]
6	2009	96	CP	[24]
7	2009	97	CP	[24]
8	2011	88	CP	[24]
9	2011	90	CP	[24]
10	2011	91	CP	[24]
11	2011	94	CP	[24]
12	2012	81	CP	[24]
13	2013	80	CP	[24]
14	2013	81	CP	[24]
15	2013	82	CP	[24]
16	2013	85	CP	[24]
17	2013	86	CP	[24]

Table 1.2: EHT data grouped into epochs based on the observed day (with the exception on year 2007) and data type, which is either visibility magnitudes (VM) or closure phases (CP).

1.3 Ancillary Data

In addition to the data observed from the EHT (Section 1.2), various other observations of Sgr A* and its inferred parameters have been used in the scope of this thesis.

Estimates for the mass and distance of Sgr A* obtained from the orbital reconstruction of the measured positions of the nearby young population of S stars, and in particular the full orbit of S2, yield a mass $M = 4.3 \pm 0.5 \times 10^6 M_\odot$ and a distance $D = 8.3 \pm 0.4$ kpc estimate with systematic uncertainties included [30, 31, 33].

Detection of linear polarization above 100 GHz [1, 6, 7, 51] and the subsequent measurements of the Faraday rotation measure [50, 52] place strong constraints on the thermal electron density in the vicinity of Sgr A*, generally, requiring fewer than 10^6 cm^{-3} and thus $\dot{M} \lesssim 10^{-8} M_\odot \text{ yr}^{-1}$. Combined with the considerably larger mass supply available, this implies the presence of substantial mass loss within the accretion flow, as described in [74]. Likewise, the known SED (Figure 1.2) gives strong constraints on the parameters of any model.

Finally, we account for the known interstellar electron scatter broadening of the image, characterized at longer wavelengths. In practice the scattering is characterized empirically as a convolution of a asymmetric Gaussian kernel to the image, or a multiplication to the corresponding Fourier image. The model employed here is from [5], where the major axis is oriented 78° east of north with the associated full width at half-maximum for the minor and major axes being:

$$\begin{aligned} \text{FWHM}_m^{\text{ES}} &= 0.64 \left(\frac{\lambda}{1\text{cm}} \right)^2 \text{ mas}, \\ \text{FWHM}_M^{\text{ES}} &= 1.309 \left(\frac{\lambda}{1\text{cm}} \right)^2 \text{ mas}, \end{aligned} \tag{1.21}$$

respectively.

1.4 Modelling Sgr A* using the EHT

It is widely accepted that Sgr A* is a subluminescent active galactic nucleus located at the center of the Milky Way, making it the prototypical low-luminosity active galactic nucleus. Previous flux measurements (see Figure 1.2 more recent version) of Sgr A* from various observations indicate that Sgr A* is fairly dim overall, with a bolometric luminosity of $L_{\text{obs}} \approx 10^{36} \text{ erg s}^{-1} \approx 3 \times 10^{-9} L_{\text{Edd}}$ [74].

As can be seen from Figure 1.2, most of this radiation is radiated in the mm peak. There is also a notable amount of gas presumably available in proximity of Sgr A*, allowing for a steady source of accreting material [49, 14]. The corresponding expected luminosity at the Bondi accretion rate \dot{M}_B [4] and at 10% efficiency is well known to be $L_{\text{exp}} \approx 0.1 \times \dot{M}_B c^2 \approx 0.110^{-5} M_{\odot} \text{ yr}^{-1} c^2 \approx 10^{41} \text{ erg s}^{-1}$. Because $L_{\text{obs}} \ll L_{\text{exp}}$, in the absence of some mechanism to dramatically restrict the gas supply, Sgr A* must have a radiatively inefficient accretion flow (RIAF), and therefore is described by models that describes a hot rotating accretion flow with viscosity.

In typical RIAF models Coulomb scattering is unable to transfer the gravitational potential energy liberated by the accreting ions to the electrons. Thus, the radiative efficiency is low as a consequence of the excess energy being stored within the thermal energy of the ions, ultimately advecting past the horizon. The result is an accretion disk comprised of a hot, collisionless, magnetized plasma. Because of the weak coupling to the ions the properties of the electron are poorly constrained in RIAFs generally. We assume as in [74] that it can be decomposed into a thermal and nonthermal component, with the electron temperature considerably lower than that of the ions. That is, when the SED constraints are applied to various models in [13] (derived from [74]), the analysis from [13] produced spatial distributions of thermal (n_e) and nonthermal (n_{nth}) electron populations (densities) as well as the thermal electron temperature (T_e) of

$$\begin{aligned} n_e &= n_e^0 \left(\frac{\rho_r}{M} \right)^{-1.1} \exp \left(-\frac{z^2}{2\rho_r^2} \right), \\ T_e &= T_e^0 \left(\frac{r}{M} \right)^{-0.84}, \\ n_{\text{nth}} &= n_{\text{nth}}^0 \left(\frac{\rho_r}{M} \right)^{-2.02} \exp \left(-\frac{z^2}{2\rho_r^2} \right), \end{aligned} \tag{1.22}$$

where ρ_r is the cylindrical radius from the black hole spin vector.

There are two robust predictions arising from RIAFs independent of the underlying physical mechanism responsible for their radiative inefficiency: the accretion disk will be

thick and sub-Keplerian. Both are immediate consequences of the inability of the accreting gas to cool.

We have locally for the ions (which as we will see are necessarily non-relativistic) an average kinetic energy per particle of

$$E_k = \frac{1}{2}m_p\langle v^2 \rangle = \frac{3}{2}kT, \quad (1.23)$$

where T is the local temperature of the ion gas, m_p is the mass of a proton, k is the Boltzmann constant.

For a RIAF, this temperature is given by the virial theorem,

$$\begin{aligned} 2E_k + E_g &= 0 \\ 3kT_{\text{virial}} - \frac{GMm_p}{r} &= 0 \\ \rightarrow T &\approx T_{\text{virial}} = \frac{GMm_p}{3kr}, \end{aligned} \quad (1.24)$$

where r is the radius of the accreting material, E_g is the gravitational energy, and $M \gg m_p$ is the mass of the black hole. Note that for $r > GM/c^2$ this is necessarily non-relativistic, justifying the assumption made in the estimate of E_k . The thermal velocity of the ions is comparable to the Keplerian velocity, $v_k = \sqrt{GM/r}$.

The vertical structure of the disk is determined by vertical force balance as depicted in Figure 1.4. Assuming that the orbital motion is confined to the azimuthal direction and is constant along cylinders, the only vertical forces are due to gravity and pressure gradients, from which

$$\begin{aligned} F_p^z &= F_g^z \\ -\frac{m_p}{\rho} \frac{dP}{dz} &= \frac{GMm_p}{r^2} \sin \theta \approx \frac{GMm_p}{r^3} z \\ \rightarrow \frac{d \ln \rho}{dz} &= -\frac{3T_{\text{virial}}}{T} \frac{z}{r^2}, \end{aligned} \quad (1.25)$$

where use was made of the ideal gas law, $P = nkT = \rho kT/m_p$, and $z \gg r$ is assumed. This may be integrated to obtain

$$\rho = \rho_0 e^{-z^2/2h^2} \quad \text{where} \quad h = r \sqrt{\frac{T}{3T_{\text{virial}}}}, \quad (1.26)$$

in which h is the disk scale height. From Equation 1.26 it is clear larger T implies larger h at fixed r . If $T \approx T_{\text{virial}}$ then $h/r \approx \sqrt{1/3} \approx 0.6$, implying a thick disk. In contrast, a cool disk (efficiently radiating disk) also implies a thin disk.

The disk orbital velocity may be obtained from the radial component of force balance, as depicted in Figure 1.4. In this case, there is a centripetal component associated with the bulk orbital motion:

$$\begin{aligned}
 F_p^r + F_c^r &= F_g^r \\
 -\frac{m_p}{\rho} \frac{dP}{dr} + m_p \frac{v_{\text{orb}}^2}{r} &= \frac{GMm_p}{r^2} \\
 v_{\text{orb}}^2 &\approx \frac{GM}{r} + \frac{r}{\rho} \frac{d}{dr} \rho kT.
 \end{aligned}
 \tag{1.27}$$

Adopting a radial profile for the density of $\rho \propto r^{-1}$ and $T \propto T_{\text{virial}} \propto r^{-1}$, this gives an orbital velocity profile of

$$v_{\text{orb}} \approx v_k \sqrt{1 - \frac{2T}{3T_{\text{virial}}}}.
 \tag{1.28}$$

Hence, for $T \approx T_{\text{virial}}$, $v_{\text{orb}} \approx v_k \sqrt{1/3} \approx 0.6v_k$, and v_{orb} decreases with increasing T .

The physically realized global structure the accretion flow of RIAFs is expected to lie somewhere between two successful models at mm wavelengths: advection dominated accretion flows (ADAFs) and advection dominated inflow/outflow solutions (ADIOS). The ADAF model assumes that all hot accreting material falls through the horizon, carrying the bulk of the gravitationally liberated thermal energy into the black hole rather than being converted to radiation. ADIOS, similar to ADAFs advects accretion matter inwards, however it is assumed that very little of the material fall into the horizon and rather gets exported out in the form of winds. RIAFs therefore consist of infalling accretion matter as well as winds to expel out a portion of the mass. The radial density profile for ADAFs and ADIOS are $\rho \propto r^{-3/2}$ and $\rho \propto r^{-1}$, respectively. Thus the radial density profile for RIAFs lie somewhere between $\propto r^{-1}$ and $\propto r^{-3/2}$.

While RIAFs are successful models for the underluminous AGNs they do not take into account of the influence from the polar magnetic field produced by the accretion of magnetic flux and implicated in the production of high powered jets. Should this field become dynamically important, the disk is then described by a magnetically arrested disk (MAD), in which the magnetic pressure is applied to the surface of the accretion disk causing it to compress into a thin hot disk [69].

1.4.1 Previous Work on Simulating Sgr A* for the EHT

Stemming from EHT observations is a collection of ongoing efforts to model the structure of the emission region in Sgr A*. These employ both semi-analytical models and full numerical




Figure 1.4: Diagram of the variables involved in illustrating how thick disk arise from RIAFs. The black hole and gas particle are of mass M and m_p , respectively. The two forces acting on the particle is the force from gravity (\mathbf{F}_g) and force from the gradient pressure (\mathbf{F}_p), along the angle θ formed from the particle being a radius r away and height z above the black hole. The particle also experiences a centripetal acceleration a_c , with a corresponding centripetal force F_c radially towards the black hole. The disk scale height is implied as h at distance r .

general magnetohydrodynamic (GRMHD) simulations with a variety of approximations for the underlying electron distribution function. In all cases the emission mechanism is assumed to be synchrotron.

A number of authors have attempted to reproduce the SED and/or 1.3mm image morphology of Sgr A* using numerical GRMHD simulations of non-radiative accretion flows [57, 15, 16, 68]. All such computations to date have presumed that the electron distribution is described by a thermal distribution, and thus are capable of only reproducing the SED of Sgr A* within a narrow range of wavelengths (most often confined to the sub-mm bump). In particular, they generally fail to reproduce the flat radio spectrum at wavelengths longer than 1cm and the near-infrared flaring emission, both of which are strongly suggested to be due to a substantial nonthermal electron population within the context of RIAF models [74, 19, 32, 22, 20]. Accreting ions serves as a large energy reservoir for transient dissipative events that couples the ions to the electrons, e.g., magnetic reconnection and strong shocks within the accretion flow. This allows for self-absorbing synchrotron emission to occur at various frequencies (radii), producing the flat radio spectrum described previously. Nevertheless, these studies have typically found accretion rates of $\dot{M} \approx 10^{-8} M_{\odot} \text{ yr}^{-1}$, corresponding to near-horizon electron densities of $n_e \approx 10^{6-7} \text{ cm}^{-3}$, and near-horizon electron temperatures of $T_e \approx 3 \times 10^{10-11} \text{ K}$. They have also found a range of permissible inclinations, ranging from nearly face-on ($\theta \approx 0^\circ$) to edge-on ($\theta \approx 90^\circ$) and a wide range of position angles.

A number of authors have also considered semi-analytical accretion flow models and various, more sophisticated electron distribution functions. Simple toroidal models with power-law electron population can reproduce the mm-SED of Sgr A*, but fails to accurately model the cm-SED similar to simulations [72]. [38] introduces a general relativistic accretion flow with a plasma wave heating mechanism, and compared their model with simulated mm-VLBI measurements of eight EHT stations. [76] updates the RIAF formulations of [74, 75] to a fully general relativistic form (note that both [9] and [10] employ a fully general relativistic velocity distribution). Comparisons of images generated using the model class described in Section 1.3 with early EHT data (epochs 1-4 in Table 1.2), spectral and polarization information allowed for constraints of $a = 0.0^{+0.64+0.86}$ (spin), $\theta = 68_{-20}^{+5}{}_{-28}^{+9}$ (inclination) and $\xi = -52_{-15}^{+17}{}_{-24}^{+33}$ (position angle), where the errors stated are the 1σ and 2σ uncertainties [9, 10]. The computed images are consistent for each epoch, implying at least annual structural invariability. More importantly, it is clear that physically motivated RIAF models provide a statistically significant improvement in the fit relative to phenomenological models such as Gaussians [10].

The first strong gravity test of general relativity using EHT data was introduced by [40], which assess the no-hair theorem through the introduction of a quasi-Kerr metric

deviation parameter with Sgr A*. Weak initial constraints were obtained from epochs 1-4 [12], corresponding to order unity constraints.

1.5 Research Cases for Sgr A* using the EHT

Section 1.4 described previous modelling constraints for Sgr A* using data from sources other than the EHT. This has been supplemented with early observations, corresponding to epochs 1-4 (Table 1.2), in the analyses of [10], which found that the RIAF paradigm was statistically favoured in comparison to phenomenological models. However, since those were completed, additional EHT data has become available primarily in the form of new closure phases (see Section 1.2.1). This motivates the evaluation of new model constraints for Sgr A*.

The addition of closure phase information can substantially improve previous constraints for the quasi-Kerr deviation model for testing general relativity. The previous constraints are described in [12] and current (Section 2.3) and future constraints (Section 2.4) are described in Section 2, where the future constraints are simulated with the addition of new stations to be included to the EHT in the near future.

Key assumptions in the class of RIAF models considered to date can be challenged by the new observations. Most of the previous RIAF imaging analyses were done under the assumption that the disk is aligned to the spin of Sgr A*. It is therefore important to assess the validity of this assumption. Were this not the case, i.e., were the disk misaligned, the resulting Lense-Thirring torques would potentially drive the precession of the accretion flow. Section 3 describes searches for precession driven by Lense-Thirring torques and, to be cautious, to find whether there exists any rigid precession at all using the most recent EHT data and RIAF fits. The outcome from the analysis is that there is a strong constraint on any observable disk precession with the current EHT data.

A novel tool is currently being developed to significantly improve the efficiency of obtaining constraints for various models used to describe Sgr A* with the most recent EHT data. The current tool used in [10] for modelling RIAFs in Sgr A* has an execution time that exponentially increases with the introduction of additional model parameters, e.g., the disk height to radius ratio and a Keplerian orbital deviation parameter. However, the increasing strength of the EHT constraints enable the exploitation of the substantial speed-ups enabled by Markov chain–Monte Carlo (MCMC) methods. Thus, we have developed a new software pipeline to analyse the implications of the rapidly growing EHT data set within an easily extensible model interface. Section 4 describes the architecture and development for this software pipeline and will be used to assess the constraints of the two additional parameters.

Finally, Section 5 concludes the thesis with a summary of the results of the investigations.

1.6 Conventions

Unless otherwise specified, we adopt units in which $G = c = 1$, timescales are in units of M and the black hole spin parameter (a or a_*) is dimensionless. Other assumptions made are that $h/r = 1$ for the thick disk scenario, and that the disk is aligned to the black hole spin, except for the case of testing for precession (Chapter 3).

Chapter 2

Quasi-Kerr Model Deviation

The unique resolution of current and forthcoming EHT observations provide an unprecedented opportunity to directly probe the spacetimes of astrophysical black holes. The lensing of the emitted photons produces general relativistic signatures that are particularly robust to variations in the underlying accretion model employed, the so-called photon ring [41]. With the exception of extreme inclinations ($\theta \lesssim 17^\circ, \theta \gtrsim 86^\circ$) and near-extremal spins ($a \gtrsim 0.9$), the photon ring is expected to be nearly circular and have a size that varies at most by 10%. However, the spacetime structure also impacts horizon-scale images through modifications to the dynamics of the accreting plasma; faster/slower orbital velocities result in more/less Doppler shifting and beaming [12].

General relativity makes unambiguous predictions regarding the structure of astrophysically relevant black holes via the no hair theorem: all such objects should be fully described by the Kerr metric, solely characterized by their mass and spin [45]. Already models of radiatively inefficient accretion flows (RIAFs) that assume the Kerr metric have had considerable success reproducing existing EHT observations [10]. However, quantitatively assessing the consistency of general relativity requires the construction and constraining of alternative metric classes.

In principle, candidate alternative spacetimes may be obtained from candidate alternative theories of gravity. However, the lack of any successful competitor theory at the present time coupled with the unprobed nature of the strong gravity regime motivate the creation of a class of phenomenological metrics with parametrized deviations. The research described here takes the latter approach.

For concreteness the quasi-Kerr metric [34] is considered, which is an extension of the Hartle-Thorne metric originally constructed to describe slowly-rotating neutron stars

[35, 36]. A key feature of the quasi-Kerr metric is that it permits an arbitrary quadrupole moment Q ; for the Kerr metric the no-hair theorem implies that the quadrupole moment, as well as all higher moments, is fully specified by M and a_* , i.e., $Q = -M^3 a_*^2$.¹ In the Hartle-Thorne metric this arises from the structure-dependent mass quadrupole of the neutron stars. For the Quasi Kerr metric this defines a family of non-Kerr metric described by one additional parameter, the quadrupole deviation ϵ , defined by

$$Q_{QK} = -M^3(a_*^2 + \epsilon). \quad (2.1)$$

Constraints on ϵ then result in quantitative tests of general relativity.

Limits on ϵ arising from the 2007 and 2009 epochs of EHT visibility magnitude data were presented in [12]. While weak, these limits demonstrated the practical ability to constrain metric deviations. Here these limits are re-assessed in the face of the more recent many epochs of EHT closure phase data, extending from 2009 to 2013 (refer to Section 1.2.1 for details). In addition, prospective constraints that may be placed by observations over the next few years are considered – when the EHT will include ALMA, the SPT, and the LMT.

Section 2.1 is a review of the form and properties of the quasi-Kerr metric. The large parameter space requires an adaptive scheme for the parameter estimation, described in Section 2.2. The current and predicted future limits on ϵ are presented in Sections 2.3 and 2.4, respectively.

¹The dimensionless spin parameter a_* is related to the angular momentum via $J = M^2 a_*$.

2.1 Quasi-Kerr Metric

The quasi-Kerr metric derives from the Hartle-Thorne metric [35, 36], originally used to describe exterior spacetimes of slow-rotating neutron stars [34]. Unlike the Hartle-Thorne metric, the quasi-Kerr metric expands around a perturbed Kerr metric. A modified version is used here, where the quasi-Kerr metric is treated as an "exact" metric with a deviated quadrupole moment (see Equation 2.1). The metric has the same higher-order multipole moments as the Kerr metric with the exception of the quadrupole order which has the additional deviation parameter ϵ [39].

The quasi-Kerr metric can be described by the line element:

$$ds^2 = g_{tt}dt^2 + 2g_{t\phi}dtd\phi + g_{rr}dr^2 + g_{\theta\theta}d\theta^2 + g_{\phi\phi}d\phi^2, \quad (2.2)$$

where

$$\begin{aligned} g_{tt} &= - \left(1 - \frac{2Mr}{\Sigma} \right) + \frac{5\epsilon(1 + 3\cos 2\theta)}{32M^2r^2} \\ &\quad \times \left[2M(3r^3 - 9Mr^2 + 4M^2r + 2M^3) - 3r^2(r - 2M)^2 \ln \left(\frac{r}{r - 2M} \right) \right] \\ g_{t\phi} &= - \frac{2Mar \sin^2 \theta}{\Sigma} \\ g_{rr} &= \frac{\Sigma}{\Delta} - \frac{5\epsilon(1 - 3\cos^2 \theta)}{16M^2(r - 2M)^2} \\ &\quad \times \left[2M(r - M)(3r^2 - 6Mr - 2M^2) - 3r^2(r - 2M)^2 \ln \left(\frac{r}{r - 2M} \right) \right] \\ g_{\theta\theta} &= \Sigma - \frac{5\epsilon r(1 + 3\cos 2\theta)}{32M^2} \left[-2M(3r^2 - 3Mr - 2M^2) + 3r(r^2 - 2M^2) \ln \left(\frac{r}{r - 2M} \right) \right] \\ g_{\phi\phi} &= \left\{ r^2 + a^2 + \frac{2Ma^2r \sin^2 \theta}{\Sigma} - \frac{5\epsilon r(1 + 3\cos 2\theta)}{32M^2} \right. \\ &\quad \times \left. \left[-2M(3r^2 + 3Mr - 2M^2) + 3r(r^2 - 2M^2) \ln \left(\frac{r}{r - 2M} \right) \right] \right\} \sin^2 \theta, \end{aligned} \quad (2.3)$$

in which Boyer-Lindquist coordinates are used, and $\Delta \equiv r^2 - 2Mr + a^2$, $\Sigma \equiv r^2 + a^2 \cos^2 \theta$. This metric will be used as a basis for the parameter estimation analyses described in Sections 2.2, 2.3 and 2.4.

Varying the deviation parameter has various effects on the photon and particle trajectories (which in turn affects the observed image) [40], resulting in:




Figure 2.1: Example silhouettes for an edge-on ($\theta = 90^\circ$), non-rotating, quasi-Kerr black hole of increasing (left to right) quasi-Kerr deviation. Taken from [12].

- a shift in the photon orbit and therefore the observed photon ring,
- a change in the particle path which results in a shift in the innermost stable circular orbit (ISCO),
- a modification of gravitational lensing,
- and observed photon redshifts are altered, both potentially increased or decreased.

The black hole silhouette varies in accordance to Figure 2.1, taken from [12], where decreasing ϵ stretches (warps) the shadow along its coordinate (spin) axis (Figure 2.1a.), and increasing ϵ stretches the shadow along its equatorial plane (Figure 2.1c.). A null deviation, $\epsilon = 0$ (Figure 2.1b.), results in a Kerr-like shadow, which is expected.

Unlike the scenario of a slow-rotating neutron star, the small size of the black hole exposes naked singularities and closed-timelike curves, which are non physical [39]. To account for this a cut-off radius $r_c = 3r_g$, where $r_g \equiv GM/c^2$, is introduced within which photon trajectories are ignored. This places a lower limit to the ISCO ($r_{\text{ISCO}} > r_c$) to avoid the removal of relevant photons resulting in an improper image of the black hole. Since the ISCO is a function of both a_* and ϵ , there exists a region in the corresponding parameter space which is excluded in the analysis (see Section 2.3)

2.2 Adaptive parameter estimation

Since this investigation builds upon the work done by [12] with the most recent closure phase data from the EHT, the procedure used to perform the parameter estimation is similar. The procedure to compute a likelihood for a given point in the (a_*, θ, ϵ) parameter space is:

1. For a given parameter point, compute a model image renormalized based on uncertain epoch-dependent flux, at 1.3mm.
2. Use the model image to obtain a Fourier image, via a Fourier transform (FFT).
3. Obtain complex visibilities at the corresponding given u - v positions at which the visibility magnitudes were measured.
4. Reiterate Step 3 for a large set of position angles, which corresponds to a rotation of the image.
5. Using the complex visibilities obtain the visibility magnitudes and closure phases on the relevant triangles (base station triplets), based on the observed data.
6. Calculate a sum of log-likelihoods (i.e., χ^2) using a direct comparison of the magnitudes and closure phases to the EHT data.

The assumed accretion flow model used in the analysis is a RIAF (with the disk radius to height ratio of 1). Parameters used in this analysis are a_* , θ , ϵ , ξ and epochs (which manifests the visibility magnitudes and closure phases). The EHT data is summarized in Table 1.2.

Rather than performing a full computation for the likelihood estimates of the entire parameter space at high resolutions (excluding the region where the ISCO is less than the cut-off radius, see Section 2.1), an "adaptive mesh-grid refinement" method is used to reduce execution time [12]. The method is as follows:

1. Start with a lower resolution probability mesh-grid, where each point in the grid is computed using the procedure above.
2. Define "high-probability" regions using a probability cut-off threshold.
3. Double the resolution of the mesh-grid and compute likelihoods for points that lie within the high-probability regions and interpolate the likelihoods for the remainder.

4. Repeat Steps 2 and 3 until the probability map reaches the desired resolution.

The reduction of execution time is mainly attributed to interpolating low probability regions rather than a full computation. Potential complications can arise from choosing an appropriate starting resolution and cut-off threshold. High likelihood regions can be excluded if the resolution step size exceeds that of the high likelihood region size and the accuracy of the probability map is affected by the cut-off threshold. To account for these complications a high resolution snapshot at $a_* = 0$ is referred to in the selection of the initial resolution and cut-off threshold: the initial resolution is chosen to ensure that its step size is small compared to the high probability region and the cut-off threshold is chosen to decrease the execution time as much as possible while maintaining adequate precision.

The probability map is then marginalized over different parameter pairs and epochs for further analyses is described in the next section.

2.3 Current Limits on Metric Deviation

Improved constraints obtained here from the existing quasi-Kerr analyses [12] will be submitted to Physical Review Letters [43]. Figure 2.2 shows the marginalized posterior probabilities of inclination and position angle to the quadrupolar deviation parameter, normalized with respect to the average probability density. Notable features from Figure 2.2 favours higher inclination angles (Figure 2.2a.) as well as a slight offset when comparing the position angle while separating the visibility magnitudes (Figure 2.2b.) and closure phase data (Figure 2.2c.). Figure 2.3 shows the marginalized posterior probabilities of inclination and position angles (normalized with respect to the average probability density) with the visibility magnitudes and closure phase comparisons separated, which more clearly illustrates the offset.

Figure 2.2b. replicates previous work [12] and can serve as a benchmark to validate the software pipeline used in the analysis. Since Figure 2.2c. corresponds to closure phase data only, this removes the 180° degeneracy that occurs for the visibility magnitudes scenario. The slight offset ($\xi = 123^\circ \pm 7^\circ$ (visibility magnitudes) and $\xi = 139_{-1}^{+18^\circ}$ (closure phases)) can be attributed to calibration uncertainties in the amplitude data, small scale structures in the black hole (accretion disk), short time scale variations between observational epochs and other physical factors that were not accounted for in the model ². Since the cause for the offset is uncertain, the offset itself is treated as a systematic uncertainty and is marginalized over.

Specifically, the spin orientation estimates for the visibility magnitude data and closure phase data is specified to be Ω_V and Ω_{CP} , respectively. Assuming that Ω_{CP} is systematically shifted by some ω , then $\Omega_{CP} = \Omega + \omega$ and $\Omega_V = \Omega$, where Ω is the joint orientation estimate. Given the posterior probability $P(a, \epsilon, \theta, \xi, \omega) = P(a, \epsilon, \Omega, \omega) = P(a, \epsilon, \Omega_V)P(a, \epsilon, \Omega_{CP}, \omega)$ (where ω is inserted as the systematic uncertainty) and assuming that the priors for the visibility and closure phase angle portions are uniform, then the marginalization of P over

²Since the time of this analysis, it was deduced that the galactic scattering was the source of the closure phase bias. It was corrected later and a paper writeup is underway, and future improvements to the constraints will include this correction.




Figure 2.2: Normalized posterior probability density as a function of the quasi-Kerr deviation parameter and (a) inclination angle, position angle compared with only (b) visibility magnitudes and (c) closure phase data, marginalized over all other parameters. The solid, dashed and dotted lines show the 1σ , 2σ , and 3σ confidence regions, respectively.




Figure 2.3: Normalized posterior probability density as a function of inclination and position angle compared with only (a) visibility magnitudes and (b) closure phase data, marginalized over all other parameters. The solid, dashed and dotted lines show the 1σ , 2σ , and 3σ confidence regions, respectively.

orientation is given by:

$$\begin{aligned}
P(a, \epsilon) &= \int d\omega d\Omega P\Pi(\omega, \Omega) = \int d\omega d\Omega P(a, \epsilon, \Omega_V)P(a, \epsilon, \Omega_{CP}) \\
&= \int d\Omega_{CP}d\Omega_V P(a, \epsilon, \Omega_V)P(a, \epsilon, \Omega_{CP}) \\
&= \left[\int d\Omega_{CP} P(a, \epsilon, \Omega_{CP}) \right] [d\Omega_V P(a, \epsilon, \Omega_V)] \\
&= \langle P_{CP} \rangle_{\Omega_{CP}} \langle P_V \rangle_{\Omega_V}.
\end{aligned} \tag{2.4}$$

That is, the systematic shift can be removed simply through isolating the visibility and closure phase probability density, marginalizing each over orientation independently, and multiplying them together. Further marginalization of the remaining parameters is achieved normally.

The posterior probabilities shown in Figure 2.4, marginalized using the method above, are normalized with respect to the average probability density over the entire parameter space. There is a substantial improvement on the constraints from previous work [12], in particular Figure 2.4b is beginning to rule out high quadrupolar deviation values at low spins and the correlation between ϵ and a_* is becoming more apparent. The grayed region from Figure 2.4b is the region which the ISCO radius is within the cut-off radius described in Section 2.1 and therefore excluded in the analysis. Figure 2.4d shows a shift in the high probability region towards higher inclination angles (disk being edge-on) and lower spin values in comparison to previous work with visibility magnitudes only.

The current 1σ constraints for individual parameters are $a_* = 0.05_{-0.05}^{+0.30}$, $\theta = 70.0_{-0.05^\circ}^{+12.5^\circ}$, and $\epsilon = 0.5_{-0.25}^{+0.45}$. While the ϵ constraint favours a non-Kerr-like black hole it is clear in Figure 2.4b, that the restriction of spin to positive values produces a positive bias in the quadrupolar deviation parameter. Hence Figure 2.4b is a better representation of the constraints and a Kerr-like black hole still lies within the 1σ region. Another distinguishing characteristic between Figure 2.4a and Figure 2.4b is there is also a noticeable constraint on the degeneracy produced by the ISCO, which there expects to be a positive correlation between the spin and quadrupolar deviation parameter (an oblique bar-like probability feature, instead of a half-Gaussian like feature as shown in Figure 2.4b).

This analysis showed a substantial improvement from the previous work [12] and demonstrates the promising capacity of the EHT in constraining general relativity with forthcoming data. The following section describes prospective improvements on these constraints based on the addition of new stations and EHT data.




Figure 2.4: The top figures are the normalized posterior probability density as a function of quasi-Kerr deviation parameter and spin using (a) previous [12] and (b) current EHT data. The bottom figures are the normalized posterior probability density as a function of inclination angle and spin using (c) previous [12] and (d) current EHT data. Each plot is marginalized over all other parameters. The solid, dashed and dotted lines show the 1σ , 2σ , and 3σ confidence regions, respectively. The grayed region in (a) and (b) is excluded in the analysis and the top and bottom dashed blue lines in (b) corresponds to lines of constant ISCO radii, $r = 6r_g$ and $r = 5r_g$ respectively.

2.4 Future Limits on Metric Deviation

In the next few years many additional stations will be added to the existing three station array. These are: the Atacama Large Millimeter/submillimeter Array (ALMA) in Chile, the Large Millimeter Telescope (LMT) in Mexico and the South Pole Telescope (SPT), and potentially the Plateau de Bure interferometer (PdB) in France and the Pico Veleta Observatory (PV) in Spain.

With these stations comes increased sensitivity and resolution, in particular ALMA is 50 times more sensitive than any of the current stations. There are also more unique closure phase triplets with the new stations (1 with 3 stations and $(N - 1)(N - 2)/2 = 21$ with 8 stations). Figure 2.5 are posterior probabilities (normalized with respect to the average probability density over the entire parameter space) using simulated EHT data with the addition of the new stations. The simulation consists of a 24 hour run with a library image at 230GHz, $a_* = 0.15$, $\theta = 60^\circ$, $\xi = 160^\circ$ and $\epsilon = 0$ [43]. Obtaining simulated (flux normalized) images once Sgr A* is above a zenith angle of 70° every 10 minutes from the 24 hour run, visibilities and closure phases are then computed using the existing methods described in previous sections.

The probability densities shown in Figure 2.5 were computed from a grid search with parameters ranges of $0.14 \leq a_* \leq 0.16$, $59.8^\circ \leq \theta \leq 60.2^\circ$, $159.88^\circ \leq \xi \leq 160.12^\circ$ $-0.0108 \leq \epsilon \leq 0.0108$ with step sizes of $\delta a_* = 0.0005$, $\delta\theta = 0.0125^\circ$, $\delta\xi = 0.005^\circ$ and $\delta\epsilon = 0.0006$, respectively.

The constraints resulting from the simulation are highly precise and places tight three-sigma limits: $a_* = 0.150_{-0.005}^{+0.004}$, $\theta = 60.01_{-0.06^\circ}^{+0.09^\circ}$, $\xi = 159.99^\circ \pm 0.04^\circ$, and $\epsilon = 0 \pm 0.005$. Comparing Figure 2.4b and Figure 2.5a the ISCO degeneracy (mentioned in the previous section) is clearly constrained. The principle direction of the probability distribution in Figure 2.5a is also slightly off compared with the ISCO line.

Although the simulation described here is based on realistic data estimates it is still based on uncertain modelling priors, i.e. the accretion model and metric themselves. Further research must be done to assess both the validity of these, as well as the sensitivity of the ϵ and a_* estimates to variations in them to enable more confident conclusions. On the other hand, the highly precise results obtained from the simulated inclusion of the additional stations illustrates the capability of the EHT in the near future, and can therefore potentially place compelling constraints on the Kerr-like nature of Sgr A*.




Figure 2.5: Normalized posterior probability density as a function of (a) quasi-Kerr deviation parameter and spin, (b) inclination angle and spin, marginalized over all other parameters. The solid, dashed and dotted lines show the 1σ , 2σ , and 3σ confidence regions, respectively. The dashed purple line in (a) corresponds to a line of constant ISCO radius, $r = 5.5r_g$

Chapter 3

Black Hole Precession

There is no reason to expect a priori that the angular momentum of the accreting gas in AGN is related to that of the black hole.¹ The result is a misaligned or “tilted” accretion disk surrounding the black hole. These tilted disks result in a variety of signatures in the temporal and structural variations of the source, motivating a detailed search for these within existing EHT data.

Sgr A* is an excellent candidate in which to investigate variations due to disk misalignment. Realistic simulations of stellar winds from the O/WR stars have demonstrated that new accretion material (in the form of turbulent streams) falls periodically in timescales as short as 10 years [14], and therefore such inflow timescales strongly suggests the presence of tilted disks [67, 58].

A clear dynamical consequence of disk misalignment is precession. Massive binary companions, local star clusters and modifications of general relativity exhibits local gravitational potentials which typically induce precession. However, the timescales for these are sensitive to the particular perturbation under consideration.²

Lense-Thirring torques provide a natural driver of precession. The frame dragging present in Kerr black holes gives rise to differing vertical and azimuthal epicyclic frequencies for orbiting material, varying with distance from the spin axis and linearly dependent on

¹In the case of X-ray binaries, prior to black hole formation either tidal locking or a common envelope phase can align the stellar spins. Large disparities between the orbital and stellar spins in binary pulsars appears to suggest that alignment is generally not expected, on the other hand.

²The angular momentum of Sgr A* is typically significantly larger than its corresponding accretion disk. Therefore the disk moves to align to the black hole spin and not the other way around. The total angular momentum of the disk depends on the disk thickness and outer radius.

the magnitude of the black hole spin. As a consequence material orbiting about Kerr black holes will experience Lense-Thirring torques, eventually leading to precession (i.e., Lense-Thirring precession). Given sufficient time, magnetically coupled viscous torques between adjacent disk rings within the warped disk lead to an eventual alignment of the of the disk with the black hole spin, the so-called Bardeen-Petterson effect [3].

Hydrodynamic simulations of thin disks have both demonstrated the existence [60, 61], or lack-of [56, 77, 59], of the Bardeen-Petterson effect. For the thick disk case, hydrodynamic and magnetohydrodynamical simulations lack any evidence of the Bardeen-Petterson effect [26, 28] suggesting that the disk precesses as a rigid body (on a global scale). Observational signatures, including horizon-scale images, have been explored on thick tilted disks from similar numerical simulations [18, 25, 17].

MHD simulations on weakly misaligned thick disks validate that the disk does precess rigidly, at a timescale that depends on the disk height [26]. This is presumably because the magnetic stresses, associated with magnetic structures with characteristic scales that grow proportionally with the disk height, are larger for thicker disks. For extreme misalignment the accretion flow breaks into accretion streams, driven by precession and the variation in the ISCO radius with latitude – orbits at higher latitude have larger ISCOs [17, 29, 27, 28].

These produce clear observational consequences on timescales intermediate between the local Lense-Thirring precession rate and the global precession rate. Besides the Bardeen-Petterson effect, another means by which the accretion flow can be aligned with the black hole spin is through jet-disk interaction. GRMHD simulations of jet launching demonstrates the jet being an efficient mechanism for redistributing angular momentum between the disk and the black hole, aligning the two for thick disks on timescales of 10^8 years [54].

Simulations of disks with high misalignment and strong precession (high spins) revealed the possible existence of stable “broken” rings within the inner subregions of the disk [63, 62, 59]. These discontinuous subregions are a consequence of the high torque generated from a large disk radius and black hole spin, exceeding the viscous coupling stresses, breaking the disk and forming multiple isolated ring-like structures. It is unclear as to whether broken disks exists for both the thin and thick disks case.

Previous analyses done by [10, 12] and the quasi-Kerr deviation chapter (Chapter 2) used models which assumed that the accretion disk is aligned to the black hole spin for Sgr A*. In this section we use precession as a means to probe the validity of this assumption.

Section 3.1 describes Lense-Thirring precession in RIAFs while Section 3.2 examines possible precession signatures from examining EHT Observations. Sections 3.3 and 3.4 searches the EHT observations of Sgr A* and examines present limits for Lense-Thirring precession and, in a more cautious scenario, if there exists any precession at all.

3.1 Lense-Thirring Precession in RIAFs

The Bardeen-Petterson effect (inner warping and eventual alignment produced by viscous stresses) is mainly discussed for accretion disks approaching the thin limit. However, RIAFs are thick, and thus may be qualitatively different, as suggested by a handful of MHD simulations. This section presents Lense-Thirring precession within the context of thick disks and consequently RIAF models.

Connecting two qualitatively different regions in a precessing disk is the warp radius, separating regions of differential to rigid precession. Beyond this radius the Lense-Thirring torques become sub-dominant to the local (magnetic) viscous torques. The degree of frame dragging, which corresponds to the strength of the Lense-Thirring torques, is dependent on the magnitude of the misalignment, spin, and distance from the black hole centre. Hence higher spins and misalignments correspond to a larger warp radius.

The sensitivity of the warp radius to viscosity also directly implies a relationship to disk height. Thicker disks are indicative of shorter in-fall times, and therefore more rapid angular momentum transport facilitated by larger effective viscosities. In the case of MHD disks this is a consequence of the larger-scale at which the MRI saturates, producing correspondingly larger Maxwell stresses (magnetic field loops generated within the disk have sizes of order the disk height). This has a critical impact on the location of the warp, and therefore the qualitative nature of Lense-Thirring precession for the RIAF models believed to be of relevance for Sgr A*. An estimate for this warp radius for the thick disk scenario is derived below (here it is assumed that $h \simeq r$).

Magnetically mediated torques are transmitted through a magnetohydrodynamic (MHD) disk in the Alfvén timescale (T_A).³ For the RIAF models we consider here and assuming magnetic stresses are responsible for driving accretion,

$$T_A = \frac{r}{v_A} = \left(\frac{\gamma\beta}{2}\right)^{1/2} T_s, \quad (3.1)$$

where v_A is the Alfvén velocity, γ is the adiabatic index, β is the equipartition parameters. Finally, T_s is the sound crossing time [64],

$$T_s = \frac{r}{c_s} = \left[\frac{2}{\gamma(\gamma-1)}\right]^{1/2} r^{3/2} \left(\frac{h}{r}\right)^{-1}, \quad (3.2)$$

the timescale for which a sound wave crosses a distance comparable to its orbital radius, and is the timescale over which local pressure gradients within the disk evolve. In this, c_s

³Throughout this chapter all timescales are given in units of GM/c^3 .

is given by

$$c_s \simeq c_{s,\text{vir}} \left(\frac{T}{T_{\text{vir}}} \right)^{1/2} = \left(\gamma \frac{P}{\rho} \right)^{1/2} \frac{h}{r} = \left[\frac{\gamma(\gamma-1)}{2} \right]^{1/2} \left(\frac{h}{r} \right) r^{-1/2}. \quad (3.3)$$

Therefore the expanded form of T_A is then

$$T_A = \left[\frac{\beta}{\gamma-1} \right]^{1/2} r^{3/2} \left(\frac{h}{r} \right)^{-1}. \quad (3.4)$$

This must be compared with the local Lense-Thirring rate, the angular frequency of a local precessing ring,

$$\Omega_{\text{LT}} = \frac{2a}{r^3} \sin \Theta. \quad (3.5)$$

The corresponding Lense-Thirring timescale is then

$$T_{\text{LT}} \simeq \Omega_{\text{LT}}^{-1} = \frac{r^3}{2a \sin \Theta}, \quad (3.6)$$

where Θ is the misalignment angle between the angular momentum vector of the disk and the precession axis vector.

Inside r_w the disk will differentially precess, dissipate the orthogonal angular momentum, and align in the equatorial plane. However, beyond r_w , Alfvén waves efficiently couple neighbouring rings, forcing them to precess rigidly. Where viscous torques and Lense-Thirring torques balance $T_{\text{LT}} \approx T_A$, which gives

$$r_w = \left[2a \sin \Theta \left(\frac{h}{r} \right)^{-1} \right]^{2/3} \left[\frac{\beta}{\gamma-1} \right]^{1/3} \simeq 3.91 (a \sin \Theta)^{2/3}, \quad (3.7)$$

where $h \simeq r$, $\beta = 10$, and $\gamma = 5/3$ are assumed for the final expression.

From Equation 3.7 it is clear that thin (cool) disks have a larger radius at which this transition occurs while for thick (hot) disks r_w can extend to small radius. Should $r_w < r_\gamma$ the entire disk is expected to precess rigidly; as shown in Figure 3.1, at low spins and Θ this is the case. Specifically, Figure 3.1 shows values of $(h/r)^{-1} \sin \Theta \leq 1$ having r_w smaller than the photon orbit r_γ and r_{ISCO} , which are characteristic of the disk inner edge. This immediately implies that thick disks should not exhibit differential (Bardeen-Petterson) precession, and are instead expected to rigidly precess. This expectation is consistent with the results of numerical MHD simulations [26].




Figure 3.1: The warp radius r_w as a function of spin with varying $\sin \Theta$, with $h/r = 1$ for thick disks. The photon orbit r_γ and ISCO r_{ISCO} lines are included for reference.

Assuming that the disk does precess rigidly (i.e., in the limit of strong radial coupling), the precession rate can be estimated using the disk-integrated Lense-Thirring torque and angular momentum. Motivated by the self-similar RIAF solutions in Section 1.4, the surface density is given by a radial power law,

$$\Sigma(r) = \rho_0 h r^{-(1+\zeta)} = \rho_0 \alpha r^{-\zeta} \equiv \Sigma_0 r^{-\zeta}, \quad (3.8)$$

where Σ_0 is the surface density constant, $\alpha \equiv h/r$ and ζ typically ranges from 0 to 0.5 for ADIOS to ADAF models; $\zeta = 0.1$ for RIAF models presented in Section 1.4 and assessed in [10].

In the rigid-disk limit, the total torque on the disk is

$$\begin{aligned} \tau &\simeq \int \tau_{\text{LT}} dr = \int \Omega_{\text{LT}} L(r) dr = \int_{r_i}^{r_o} \Omega_{\text{LT}} 2\pi r \Sigma(r) \sqrt{Mr} dr \\ &= \int_{r_i}^{r_o} 2\pi r \Sigma_0 r^{-\zeta} \Omega_{\text{LT}} \sqrt{Mr} dr \\ &= 4\pi \Sigma_0 a \sin \Theta \int_{r_i}^{r_o} r^{-(3+\zeta)/2} dr \\ &= \frac{8\pi \Sigma_0 a \sin \Theta}{1 + \zeta} \left[r_i^{-(1+\zeta)/2} - r_o^{-(1+\zeta)/2} \right]. \end{aligned} \quad (3.9)$$

Similarly, the total angular momentum in the orbiting gas is

$$\begin{aligned} L &\simeq \int_{r_i}^{r_o} 2\pi r \Sigma_0 r^{-\zeta} \sqrt{Mr} dr = 2\pi \Sigma_0 \int_{r_i}^{r_o} r^{(3-\zeta)/2} dr \\ &= \frac{4\pi \Sigma_0}{5 - \zeta} \left[r_o^{(5-\zeta)/2} - r_i^{(5-\zeta)/2} \right]. \end{aligned} \quad (3.10)$$

From these, an estimate of the precession timescale is then

$$\begin{aligned} T_{\text{prec}} &\simeq \frac{L}{\tau} = T_{\text{LT}}|_{r_i} \left(\frac{1 + \zeta}{5 - \zeta} \right) \left[\frac{(r_o/r_i)^{(5-\zeta)/2} - 1}{1 - (r_o/r_i)^{-(1-\zeta)/2}} \right] \\ &\simeq f(\zeta) \frac{r_i^{(1+\zeta)/2} r_o^{(5-\zeta)/2}}{10a \sin \Theta}, \end{aligned} \quad (3.11)$$

where $f(\zeta) = 5(1 + \zeta)/(5 - \zeta)$ ranges from 1 at $\zeta = 0$ to $5/3$ at $\zeta = 0.5$. As with the other timescales presented, T_{prec} is in units GM/c^3 . The final expression is obtained assuming $r_o \gg r_i$.

It should be noted that $\zeta = 0$ corresponds to $T_{\text{prec}} \propto r_o^{5/2}$ and $\zeta = 0.5$ corresponds to $T_{\text{prec}} \propto r_o^{9/4}$, both of which grow more rapidly with radius than $T_A \propto r^{3/2}$. Hence the rigid disk precession timescale will always be greater than the local Alfvén timescale, consistent with the continued assumption of a rigid disk. The precession timescale generally is smaller than the Lense-Thirring timescale, i.e. $T_{\text{prec}} < T_{\text{LT}}$, since the torques driving precession are primarily from the inner edge of the disk while the angular momentum is primarily located at large radii.

For RIAFs: $\zeta = 0.1 \simeq 0$, leading to $T_{\text{prec}} \propto r_i^{1/2}$, which varies significantly less compared with $r_o^{5/2}$. Because $r_i = \max(r_\gamma, r_w)$, which from Figure 3.1 can range from 1 to 4, T_{prec} varies at most by a factor of 2 with r_i . For this reason, $r_i = 1$ is assumed in the analysis described in Section 3.3, given the assumption that $r_o \gg r_i$.

3.2 Signatures of Rigid Disk Rotation in EHT Observations

While it may be possible to find evidence of misaligned, thick accretion disks in the images themselves [18, 17], an alternative observable is a time-varying disk orientation due to a precessing disk. Applying existing methods detailed in [11], it is possible to estimate the period using the most recent EHT data.

A disk with angular momentum

$$\hat{\mathbf{L}} = \sin \Theta \cos \Phi \hat{\mathbf{x}} + \sin \Theta \sin \Phi \hat{\mathbf{y}} + \cos \Theta \hat{\mathbf{z}}, \quad (3.12)$$

precessing about $\hat{\mathbf{z}}$ and viewed from

$$\hat{\mathbf{k}} = \sin \vartheta \hat{\mathbf{x}} + \cos \vartheta \hat{\mathbf{z}}, \quad (3.13)$$

will appear to have an angle between the line of sight and the disk axis (θ) and position angle (ξ) of

$$\begin{aligned} \theta &= \arccos(\sin \Theta \cos \Phi \sin \vartheta + \cos \Theta \cos \vartheta) \\ \xi &= \arctan \left(\frac{\sin \Theta \cos \Phi}{\cos \Theta \sin \vartheta - \sin \Theta \sin \Phi \cos \vartheta} \right) - \varphi, \end{aligned} \quad (3.14)$$

in which φ is the projected position angle of the precession axis (i.e., ϑ and φ are the angular coordinates of the precession axis relative to the line of sight). Θ and Φ is the position angle of the angular momentum relative to the precession axis. Figure 3.2 illustrates the angles described.

Thus a cyclic variation in the disk angular momentum vector due to precession produces a corresponding cyclic variation in the reconstructed disk orientation. These are shown in Figure 3.3 for a range of particular examples. Note that there is a discontinuous jump when $\vartheta = \Theta$, and that smaller Θ typically results in smaller ranges of θ and ξ . At $\vartheta = 90^\circ$, the maximum range limits are $\theta = \pi/2 \pm \Theta$ and $(\xi - \varphi) = \pm \Theta$.

As the disk precesses Φ evolves, causing it to traverse a path in θ and ξ . A set of constraints upon θ and ξ at different times imply a limit upon the evolution of Φ , or equivalently the precession period. Sections 3.3 and 3.4 seek to constrain this precession period and the physical parameters it depends upon with the epoch-specific analyses of EHT data.

A lower limit may already be imposed on the precession timescale. The detection of fringes using global VLBI requires the image to be stationary over the roughly 10

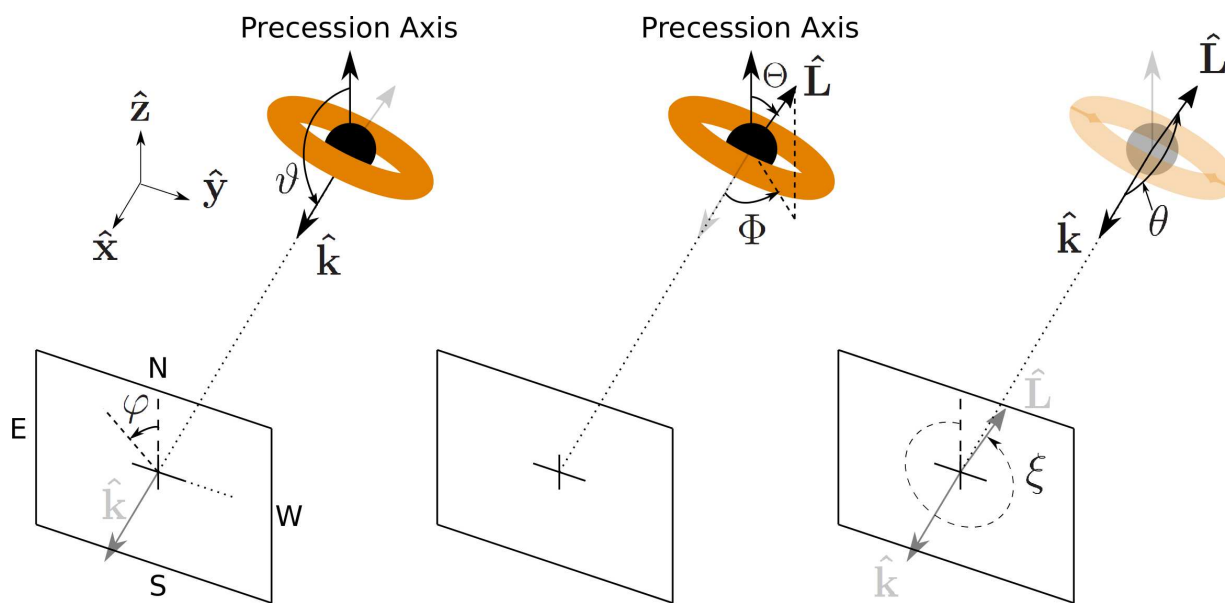


Figure 3.2: Illustration of the signature angles (ϑ , φ , Θ , Φ , θ , ξ) used for finding rigid-body disk precession. The North-East-South-West (NESW) directions are shown to reflect that \hat{k} is pointing in the direction towards the viewing plane (Earth).




Figure 3.3: Path of a precessing disk (i.e. path with changing Φ) in θ - ξ space with varying ϑ . Here $\Theta = 45^\circ$ and ϑ ranges from 0° to 90° .

second atmospheric coherence time. There are also a lack of catastrophic SNR loss during incoherent ($\lesssim 10$ minutes) averaging of the detected visibilities, implying image stationarity over that timescale as well. Good RIAF image fits from [11] is indicative of a lack of structural variability of timescales for as long as 3 hours. The apparent consistency between multiple RIAF image fits over the course of several epochs, with ranges from days (3 to 4), weeks (1 to 2) and several years (2007 to 2013) also suggests a lack of substantial structural variability for up to 6 years. Thus, should Sgr A* precess it must do so on periods of order a decade or more.

This is much larger than the typical orbital periods of material near the black hole. For example, for a particle for a slowly rotating black hole Kerr black hole near the ISCO,

$$T_{\text{orbit}} = 2\pi \frac{GM}{c^3} \left[\left(\frac{c^2 r}{GM} \right)^{3/2} + a \right] \approx 33 \text{ min}, \quad (3.15)$$

assuming $r = r_{\text{ISCO}} = 6(GM/c^2)$, $M = 4.3 \times 10^6 M_{\odot}$ and $a = 0.15$. It is similarly long in comparison to the Lense-Thirring precession timescale at the inner disk edge,

$$T_{\text{LT}} \approx T_{\text{orbit}} \frac{r^{3/2}}{2a} \approx 27 \text{ hr}, \quad (3.16)$$

at the ISCO for the same black hole parameters.

3.3 Search for Lense-Thirring Precession using EHT Observations

To find evidence for precession we made use of epoch-specific RIAF-based analyses of the most up-to-date EHT data in Table 1.2. Explicitly, the following procedure was used for the search of rigid-body Lense-Thirring Precession:

1. Obtain fitted data cubes of θ , ξ , and a log-likelihoods calculated using the most up-to-date EHT data and image comparison procedure described [11], for each epoch. Epoch 11 was excluded due to a flaring event.
2. Execute Markov chain Monte Carlo (MCMC) with Metropolis Hastings algorithm along with adaptive stepping towards the principle directions, allowing ϑ , φ , Θ , Φ_0 , r_o , a and the precession direction (k_{dir}) to vary. Note that Φ_0 is the Φ value at a fiducial time, 12:00am UTC, January 1, 2000.
3. For each proposed point, T_{prec} is computed from Equation 3.11 with a and r_o . Φ_n at a given epoch n is then computed for all epochs using $\Phi_n = \Phi_0 + 2\pi k_{\text{dir}}(t_n/T_{\text{prec}})$, where t_n is the time elapsed until to epoch n . The angle parameters and Φ_n are used in Equation 3.14 to compute θ_n and ξ_n , which are θ and ξ for a given epoch n . Lastly, the log-likelihoods for each fitted data cubes at epoch n are interpolated for the point (θ_n, ξ_n, a) and then summed together to obtain an overall log-likelihood for the MCMC driver to evaluate.
4. Isotropic priors on all orientations are assumed, a flat prior on a is assumed, and a logarithmic prior is applied to the disk size (r_o) to avoid biasing the search towards any particular disk size.
5. The 7D posterior probability distribution obtained from the MCMC chain is then marginalized to produce 1D and 2D plots for further analysis.

It should be noted that the closure phase data for all epochs have been corrected for systematic shifts, potentially due to time integrated bias from the Galactic scattering screen.

Figures 3.4 and 3.5 are 1D marginalized probability distributions resulting from the converged MCMC chain. Figure 3.4 shows in effect that there are no constraints that can be drawn for the orientation parameters. On the other hand Figure 3.5a and Figure 3.5b shows an upper limit for a and lower limit for the r_o respectively. From inspection $a \gtrsim 0.8$

and $r_o \lesssim 60M$ lies above the 3σ region and $a \lesssim 0.575$ and $r_o \gtrsim 300M$ lies within 2σ . The peaks show that lower spin and a larger outer radius are favourable, which from Equation 3.11 indicates a long precession period. Low spin values justify the presumption that the disk is indeed precessing rigidly, as discussed from Section 3.1 referring to Equation 3.7 and Figure 3.1.

Figure 3.6 shows 2D marginalized plots of the rigid-body Lense-Thirring model parameter pairs, arranged in a triangular layout. It can be seen that there are constraints for both a and r_o , as described for the 1D marginalized case. There appears to be no dependence for both a and r_o on any of the angle parameters, and there is a correlation between a and r_o as implied in Equation 3.11, although there seems to be no correlation between the two and Θ . Again, in agreement with the 1D scenario, low spin and high r_o values are favoured.

The strong correlations in Figure 3.6 are artefacts of a degeneracy between the model angle parameters $(\vartheta, \varphi, \Theta, \Phi)$ and the fitted angles (θ, ξ) occurring at large precession periods. This arises because at fixed (θ, ξ) Equation 3.14 remains an underdetermined system of equations for the remaining angles, resulting in infinitely many solutions, the projection of which leads to the structures that appeared in Figure 3.6. Figure 3.7 compares structures resulting from marginalizing a probability distribution generated from $(\vartheta, \varphi, \Theta, \Phi)$ solutions to the corresponding fixed (θ, ξ) with the results in Figure 3.6. The following procedure was used to generate the 2D probability distributions:

1. Generate a 4D meshgrid for the model angle parameters $\vartheta, \varphi, \Theta$, and Φ .
2. Compute θ and ξ using Equation 3.14 for each grid point.
3. Compute a 4D probability cube using the fitted $\theta = 60^\circ$ and $\xi = 156^\circ$ values with the ones calculated in Step 2 assuming an uncorrelated Gaussian with roughly the correct uncertainties (i.e. 5° in θ and 15° in ξ).
4. 2D marginalize the 4D probability cube to obtain angle-pair probability distributions

For this reason we do not consider the apparent correlations in the orientation variables further.

Figure 3.8 shows a more detailed analysis of the 2D marginalized probability distribution of a and r_o . The 1σ region shows there is a peak towards an r_o value of roughly $1000M$ as well as favouring smaller spins. It is unclear whether r_o has an upper limit since the 2σ region continues (plateaus) at higher r_o values. However there does seem to be a




Figure 3.4: Marginalized constraints upon individual angle parameters for rigid-body Lense-Thirring precession from all but epoch 11. In all cases the gray bar indicates the 1σ uncertainty in the probability distribution, obtained by comparing multiple chains. In all cases, a probability density of 1 corresponds to a uniform distribution. The 1σ , 2σ and 3σ regions are shaded in order from dark to light blue, and the red regions show values with $> 3\sigma$.




Figure 3.5: Marginalized constraints upon a and r_o parameters for rigid-body Lense-Thirring precession from all but epoch 11. In all cases the gray bar indicates the 1σ uncertainty in the probability distribution, obtained by comparing multiple chains. In all cases, a probability density of 1 corresponds to a uniform distribution. The 1σ , 2σ and 3σ regions are shaded in order from dark to light blue, and the red regions show values with $> 3\sigma$. For reference, in the case of the disk outer radius probability, the hatched region shows radii beyond the periape of S2. Also for reference, the spin probability prior obtained from current observations is shown by the dotted line.




Figure 3.6: Triangular plot of marginalized constraints upon pairs of model parameters for rigid-body Lense-Thirring precession from all but epoch 11. In all cases a probability density of 1 corresponds to a flat distribution.




Figure 3.7: Demonstration of false correlations based on model-angle degeneracy resulting in structure for 2D angle pair plots for the rigid-body Lense-Thirring precession study. The top row are some example model angle-pair plots showing structure produced from various angle configurations which yield the same fitted ξ and θ values. The bottom row are the corresponding rigid-body Lense-Thirring run results for comparison. The three columns from left to right are plots of $\cos \Theta - \cos \vartheta$, $\Phi_0 - \varphi$ and $\cos \Theta - \varphi$, respectively.

upper limit on the spin and lower limit on r_o . While r_o appears to peak at a radius below the periape of S2/SO-2, there is no conclusive evidence that the disk can not extend beyond this orbital radius. At around $2000M$, the corresponding Lense-Thirring precession timescale is calculated to be $T_{\text{prec}} \approx 2500$ yr years from Equation 3.11. An upper limit for r_o is presumably set by the clockwise stellar disk (O/WR stars) believed to supply the accreting material for the accretion disk, located at around $2 \times 10^5 M$ [14], two orders of magnitude further out. The general shape of the probability distribution seems to align with the lines of constant precession periods implying that the shape is a consequence of T_{prec} . The figure also shows a lower limit of about 0.01 yr for the precession period, as shorter precession periods lie above the 3σ region.




Figure 3.8: Marginalized constraints upon the a vs. r_o parameter pair for rigid-body Lense-Thirring precession from all but epoch 11. Solid, dashed, and dotted contours show the 1σ , 2σ , and 3σ confidence level regions, respectively. In all cases a probability density of 1 corresponds to a flat distribution. The four dashed purple lines are lines of constant precession periods, in order from left to right being 0.01 yr, 1 yr, 100 yr and 10^4 yr, respectively. For reference the hatched region shows radii beyond the periastron of S2/SO-2.

3.4 Search for Any Precession using EHT Observations

The search for Lense-Thirring precession presumed a particular relationship between the reconstructed magnitude and orientation of the spin. However, for highly misaligned accretion flows, or for precession driven by other mechanisms, this may not hold. Therefore, we also search for an arbitrary precession mechanism by allowing the precession period to be an independent variable.

The procedure used is identical to the procedure described at the beginning of Section 3.3, with the exception that data matrices of θ , ξ log-likelihoods are used instead of data cubes (i.e. the data cube is marginalized over a). The free parameters are now $(\vartheta, \varphi, \Theta, \Phi_0, T_{\text{prec}}, k_{\text{dir}})$, and a logarithmic prior is applied to T_{prec} instead of r_o . As before, the resultant 6D posterior probability distribution obtained from the MCMC chain is then marginalized to produce 1D and 2D plots for further analysis.

Similar to the results in Section 3.3, the 1D marginalized plots from Figure 3.9 reveals no significant constraints for the model orientation parameters.

Inspecting Figure 3.10, which is the 1D marginalized plot for the precession period, gives a lower limit ($> 3\sigma$) of roughly 0.03 yr or 11 days and a 2σ lower limit of about 30 yr, for any source of precession. This eliminates small, and broken inner disks such as those seen in [63], and persistent coherent precessing features such as the shocks described in [17]. The resonance-like feature from Figure 3.10 are a result of aliasing. The EHT observing runs typically occur at around the same one month period each year (last two weeks of March through the first two weeks of April), with the result that periods that are an integer divisor of a year (6 months, 4 months, 3 months, etc.) are indistinguishable. Because the observations are not performed with an exact annual cadence the resonance features are not uniform. Experiments with simulated observations at slightly earlier or later times imply that the resonance features can be resolved with a small number of additional observations.

A key feature in this figure is that the 1σ and 2σ regions are “plateaued” indicating that a range of large precession periods are equally probable. Thus, if the accretion disk in does precess Sgr A*, the period must be extremely large, and is currently indistinguishable from no precession at all.

Figure 3.11 shows 2D marginalized plots of the any rigid-body parameter pairs, arranged in a triangular layout. As explained in the previous section, and demonstrated in Figure 3.12, the apparent correlations in the various orientations are a result of Equation 3.14




Figure 3.9: Marginalized constraints upon individual angle parameters for any rigid-body precession from all but epoch 11. In all cases the gray bar indicates the 1σ uncertainty in the probability distribution, obtained by comparing multiple chains. In all cases, a probability density of 1 corresponds to a uniform distribution. The 1σ , 2σ and 3σ regions are shaded in order from dark to light blue, and the red regions show values with $> 3\sigma$. In all cases, a probability density of 1 corresponds to a uniform distribution.




Figure 3.10: Marginalized constraints upon the precession period parameter for any rigid-body precession from all but epoch 11. In all cases the gray bar indicates the 1σ uncertainty in the probability distribution, obtained by comparing multiple chains. In all cases, a probability density of 1 corresponds to a uniform distribution. The 1σ , 2σ and 3σ regions are shaded in order from dark to light blue, and the red regions show values with $> 3\sigma$.

being an underdetermined system (i.e. multiple configurations of ϑ , φ , Θ , and Φ can result in the same θ and ξ values). The precession period plots (bottom row of Figure 3.11) reveals no correlation between the angles and the precession period, though they do show the resonance feature seen in Figure 3.10.




Figure 3.11: Triangular plot of marginalized constraints upon pairs of model parameters for any rigid-body precession from all but epoch 11. In all cases a probability density of 1 corresponds to a flat distribution.




Figure 3.12: Demonstration of false correlations based on model-angle degeneracy resulting in structure for 2D angle pair plots for the any rigid-body precession study. The top row are some example model angle-pair plots showing structure produced from various angle configurations which yield the same fitted ξ and θ values. The bottom row are the corresponding any rigid-body run results for comparison. The three columns from left to right are plots of $\cos \Theta - \cos \vartheta$, $\Phi_0 - \varphi$ and $\cos \Theta - \varphi$, respectively.

3.5 Conclusions

The search for misalignment, and consequentially dynamical precession, of the accretion disk in Sgr A* using the most recent EHT data, places strong lower limits on precession periods of about 30 yr, independent of the driving mechanism. The epoch-specific reconstructions of Sgr A*'s horizon scale structure over the six years from 2007 to 2013 showed no discernible precession over the course of this timespan. With Lense-Thirring torques as a driving mechanism, strong limits were placed on the black hole spin of as low as 0.575 and accretion flow size of up to $300M$.

The assumption that Lense-Thirring torques are responsible for the precession period leads to joint limits on the size of the accretion disk, spin of the black hole and the misalignment between the two. The apparent absence of precession can be explained in one of three possible ways, each of which have significant implications for existing efforts to model EHT data. First, a small black hole spin implies small Lense-Thirring torques and in turn can reduce the precession rate to a minimum. In this case even large misalignments can still produce small changes in the model reconstructed images. Secondly, small misalignments with small or high black hole spins, again leads to small Lense-Thirring torques and results in the same conclusion as the first explanation. Lastly, the disk itself can be very large, thus the angular momentum large, leading to precession periods well in excess of the 30 yr limit.

A possible upper limit for the disk size arises from the putative source of the accreted gas: the O/WR stars in the clockwise stellar disk at 0.03 pc, $2 \times 10^5 M$. Existing analyses of the orientation of Sgr A*'s accretion disk aligns with this disk and thus suggests an initial angular momentum of the accretion gas. It is not clear whether the accretion disk can extend beyond this size, however if it extends up to the O/WR clockwise stellar disk it can indirectly support the aligned disk model, should the accreted gas originate from the stellar disk. Note that at moderate black hole spins the constraint on the disk size extend beyond $10^3 M$, reaching the periapse of S2/S0-2.

Finally, both numerical simulations and analytical estimates of precessing rigid body disks have all been for isolated systems. In the present scenario, an outer boundary condition can exist where accreted gas is constantly being replenished from the O/WR clock stellar disk (a reservoir of fixed angular momentum). This can then drive the development of a fixed stationary warped disk, where the shape is some interpolation between the orientation of the angular momentum at the clockwise stellar disk at large radii and the black hole at small radii, and consequentially resulting in the apparent lack of precession observed in Sgr A* by the EHT. Thus, future theoretical work to clarify the implications

of Lense-Thirring torques in the presence of accretion flows with a large gas supply is required.

Chapter 4

Next Generation Analysis Pipeline

The previous chapters discussed new and forthcoming constraints obtained from comparing various models to the EHT data, developing insight into the magnitude and orientation of the spin of Sgr A*, its alignment with the surrounding accretion flow, and permitted deviations from general relativity. [10] had determined robust constraints for the black hole spin, inclination and position angle to be $a = 0.0^{+0.64+0.86}$, $\theta = 68_{-20}^{+5+9}$ and $\xi = -52_{-15}^{+17+33}$, and newly determined constraints with closure phases (epochs 5-17 from Table 1.2) in order to test for GR are $a_* = 0.05_{-0.05}^{+0.30}$, $\theta = 70.0_{-0.05}^{+12.5}$, and $\epsilon = 0.5_{-0.25}^{+0.45}$ (see Chapter 2 for details), which are well within agreement with [10]. The search for disk alignment led to no observable evidence for precession (refer to Chapter 3), indicative of an aligned disk.

All of the previous investigations used versions of the existing analysis pipeline, described in [10], and new investigations will require additional modification of the existing code. Thus there is a substantial overhead development cost for every new investigation, since the existing pipeline was not created to be extensible to new data and models. Another limitation is that currently the pipeline analyses the data using grid-search based methods, which is advantageous for poor constraints (low resolution parameter grid space) for few parameters. However there is an exponential increase in the time taken for better constraints (higher resolution parameter grid space) or when more parameters are added to the search.

For example, the addition of two RIAF parameters, the disk height ratio $h/r = \sqrt{T/3T_{\text{virial}}}$ and orbital velocity coefficient $b = \sqrt{1 - 2T/3T_{\text{virial}}}$ (where the relationships are from Equations 1.26, and 1.28, respectively) will be beneficial in constraining the hot disk leading to a thick disk and slower orbital velocity prediction as described in Section




Figure 4.1: Rudimentary computed theoretical images of varying h/r of a RIAF model for Sgr A*. This uses the most recent best fitted values of spin and inclination, $a = 0.16$ and $\theta = 60^\circ$, respectively. The thermal electron density and temperature, as well as the nonthermal electron density are optimized for each disk height ratio to the observed spectrum data of Sgr A*.

1.4. Both of these parameters are in turn related to the ion temperature profile of the disk. Figure 4.1 is a rudimentary example of how the computed image of Sgr A* would change with varying disk height, using the best known fits of spin and inclination and re-optimizing for the thermal electron density and temperature, as well as the nonthermal electron density. As indicated in the figure, varying the disk heights has a significant impact on the spatial structure of the image. Fitting for these additional parameters using the existing RIAF analysis pipeline not only introduces modification costs, but also adds two dimensions to the parameter space, increasing the grid-search time to the order of years if good constraints are to be acquired.

Thus to address new investigations similar to the prior example, a new program called THEMIS is developed. Two main features for THEMIS are its extensibility in adding new data types and its ability to use Markov chain Monte Carlo (MCMC) methods to search for good constraints.

In the near future new data types, which includes time varying SEDs, closure amplitudes, polarization fractions, polarization position angles, etc., are going to be available for Sgr A* and it is useful to seamlessly integrate these types into the pipeline without excessive overhead recoding expenses. THEMIS allows for this to happen through an automated code generation extension, which reads a type file and outputs a data class that includes all available data types to be used in the main program. This method not only alleviates the stress of debugging possible errors through new data type integration, it also localizes all of the observed data in a data object framework which then can be used in all parts of the main program, making it extremely versatile.

The main advantage of MCMC methods over grid-based search methods are that MCMC methods focuses on “good” (high likelihood) regions, i.e., it is designed to allocate most of its time within high likelihood regions in stark contrast with grid-based search methods which typically allocate time uniformly throughout the parameter space. The Metropolis-Hastings theorem [37] guarantee that MCMC methods will produce a distribution of parameters that matches the implied posterior probability distribution when converged. However the time of convergence is not known and depends on many factors, such as the initial starting guess and the compatibility between the proposal distribution with the actual distribution. Another potential drawback is that when the posterior probability is multi-modal MCMC can become stuck on local maxima, failing to explore the entire probability space. Nevertheless, the exponential speed-up combined with MCMC algorithms that exploit multiple chains provide an efficient scheme to explore the applicability of models to large data sets and infer parameter estimates and uncertainties.

4.1 THEMIS Overview

The main feature of the THEMIS pipeline compared to its predecessor is its extensibility to support future growth and adaptability to new input. The design philosophy is to implement modularity in as many components of the process as possible. This includes extending the program’s reach to new observed data types from the EHT and other sources, the flexibility to swap, modify or append theoretical models, optimizers and search algorithms for exploring theoretical inquiries based on the current evidence, and for the scope of the thesis accretion flow models and general relativity tests using observables of Sgr A*.

A brief summary of the accretion model parameter estimation procedure using EHT and spectral data is shown below (derived from [9]):

1. Start with an initial set of parameters at the i th iteration $\mathbf{X}_{i=0} = (a_i, \theta_i, \xi_i, \dots)$.
2. Determine a proposal set $\mathbf{Y}_i = (a_{y,i}, \theta_{y,i}, \xi_{y,i}, \dots)$ obtained from adding a proposal step to \mathbf{X}_i .
3. From \mathbf{Y}_i compute a model image \mathbf{I}_i using ray tracing of the input theoretical accretion model. The image flux is rescaled using parameters to match the average expected flux value.
4. Compute the corresponding Fourier image \mathbf{F}_i of image \mathbf{I}_i . The Fourier image is then multiplied (blurred) with an electron scatter broadening kernel (Equation 1.21).
5. For each observed visibility magnitude (Section 1.1.2) at a given uv baseline and closure phase (Section 1.1.3) at a given station triplet for each epoch k , compute the corresponding visibility magnitudes and closure phases from the theoretical Fourier image and obtain $\chi_{\text{VM},i}^2$ and $\chi_{\text{CP},i}^2$ estimates, respectively. These χ^2 estimates are summed together to $\chi_{\text{Total},i}^2$, factoring in possible observation errors for each observation and data type.
6. The set of parameters \mathbf{Y}_i is also used to compute the flux at each wavelength corresponding to the observed SED using adaptive mesh refinement at high flux regions to speed up the process. The theoretical spectrum is compared with the observed SED to produce $\chi_{\text{SED},i}^2$, which is then added to $\chi_{\text{Total},i}^2$.
7. If the $r\chi_{\text{Total},i}^2 < \chi_{\text{Total},i-1}^2$ where r is an acceptance ratio and $i-1 = 0$ for $i = 0$, then \mathbf{Y}_i is “accepted” and $\mathbf{X}_{i+1} = \mathbf{Y}_i$. Else if $r\chi_{\text{Total},i}^2 \geq \chi_{\text{Total},i-1}^2$ then \mathbf{Y}_i is “refused”, and $\mathbf{X}_{i+1} = \mathbf{X}_i$.

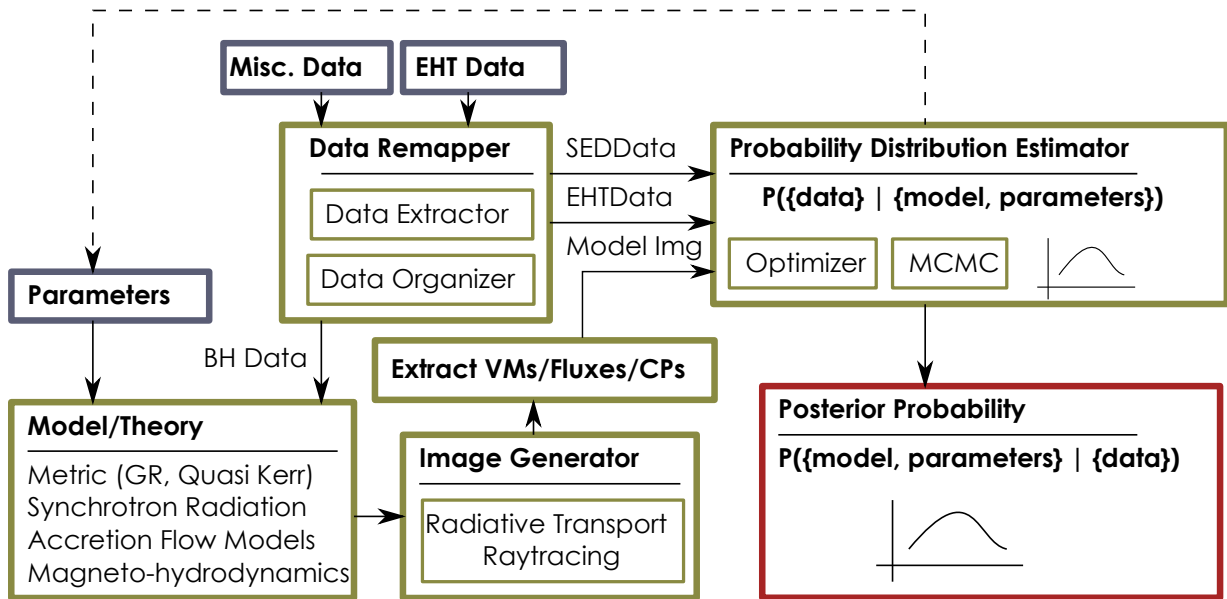


Figure 4.2: Flow diagram overview of the THEMIS software. The blue and red blocks denote input and output blocks, respectively. The green blocks denote the main code components of the pipeline and the arrows indicate data flow with the adjacent text being the type of data.

8. Steps 2 to 7 are reiterated for $i = i + 1$ until a convergence condition is met, and the step chain $\{\mathbf{X}\}$ becomes the posterior probability distribution for the parameter estimation.

The procedure above uses a basic MCMC Metropolis-Hastings algorithm [37].

Figure 4.2 shows a flow diagram overview of the THEMIS software. The program itself currently is divided into three major components: a theoretical model image generator (Section 4.1.1), a data remapper (Section 4.1.2) and a posterior probability estimator (Section 4.1.3). The model image generator comprises of the Model/Theory, Image Generator and Extract VMs/Fluxes/CPs blocks. The data remapper have a block of its own name and the posterior probability estimator consists of the Probability Distribution Estimator block. The combination of these components allows for the input of a theoretical accretion model and observed data for Sgr A* and outputs a posterior probability for the parameter fits. By its nature, THEMIS can naturally accomodate disparate model classes (e.g., accretion disks, jets, phenomenological models, etc.) and heterogeneous data sets, permitting a

trivial application to alternate EHT targets.

4.1.1 Theoretical Image Generation

The theoretical image generation component of THEMIS is divided into two parts (Figure 4.3): a model image generation pipeline and model image application user interface (API). The model image generation pipeline is the existing image generation code used in previous RIAF analysis papers [10], and will hence be called VRT2, as shown in the figure. VRT2 takes a parameter set as well as other corresponding black hole estimates as input and combines them with an accretion model and metric (RIAF and GR) to produce model images of simulated observations from Earth. For each pixel in the model image a null geodesic is traced backwards in time through the emission region, and then the radiative transfer equations are integrated forward towards the observer, as described in [8]. VRT2 outputs the full complement of Stokes parameters for the image, and thus contains polarization information as well as the intensity map. The model API receives the model image produced from VRT2 and formats it to become compatible with THEMIS, as well as allowing the ability to output Fourier images, visibility magnitudes, closure phases, and fluxes at various wavelengths.

The `Image Analyzer` compares the images obtained from the model API with the EHT and ancillary data products, contained in `EHTData` to produce a likelihood. This is then the primary input into subsequent analysis steps.

4.1.2 Data Remapper

All observed data is contained within a single class object called `EHTData`, which is a collection of vectors of data types for various observables. Figure 4.4 shows a flow diagram on how the observed data is remapped into the class object. It first extracts the observed data whose location and type is found from the file info block, then organizes the data to their corresponding data types and outputs the result into the object `EHTData`. The data can then be accessed through its data type and provided filters. For example, to access all black hole data employed for Sgr A*, the code

```
EHTData.Black_Hole.filter_name(EHTData.Black_Hole, 'Sgr A*', ...)
```

is used, where the basic outline of the code is

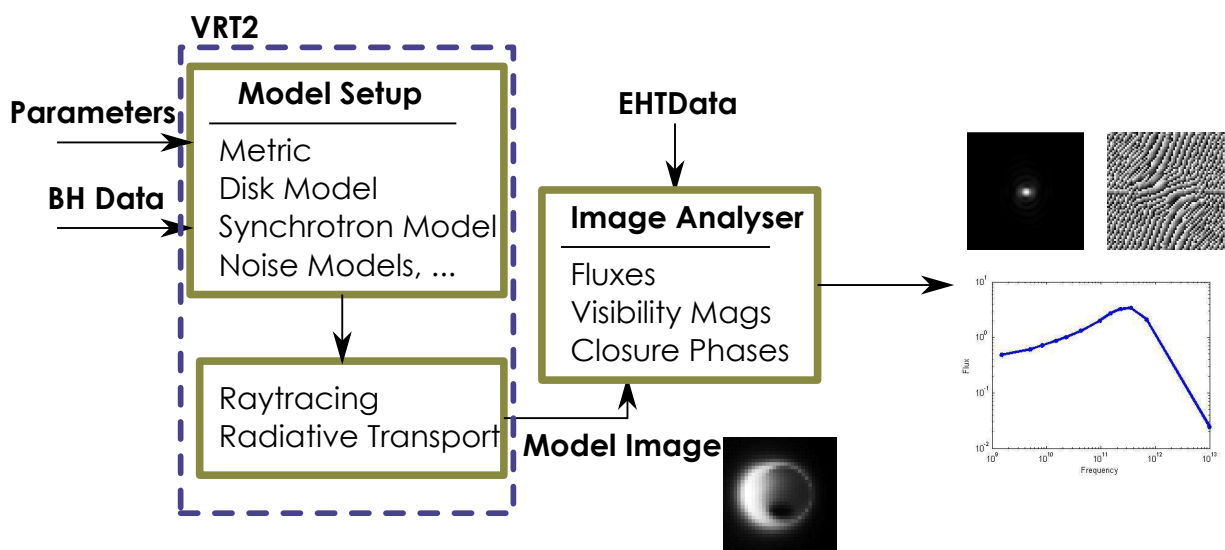


Figure 4.3: Theoretical image generation component of THEMIS. This is a combination between the VRT2 (existing RIAF image generation pipeline) and an API built to interface VRT2 with THEMIS, allowing for outputs of other data types such as visibility magnitudes and closure phases.

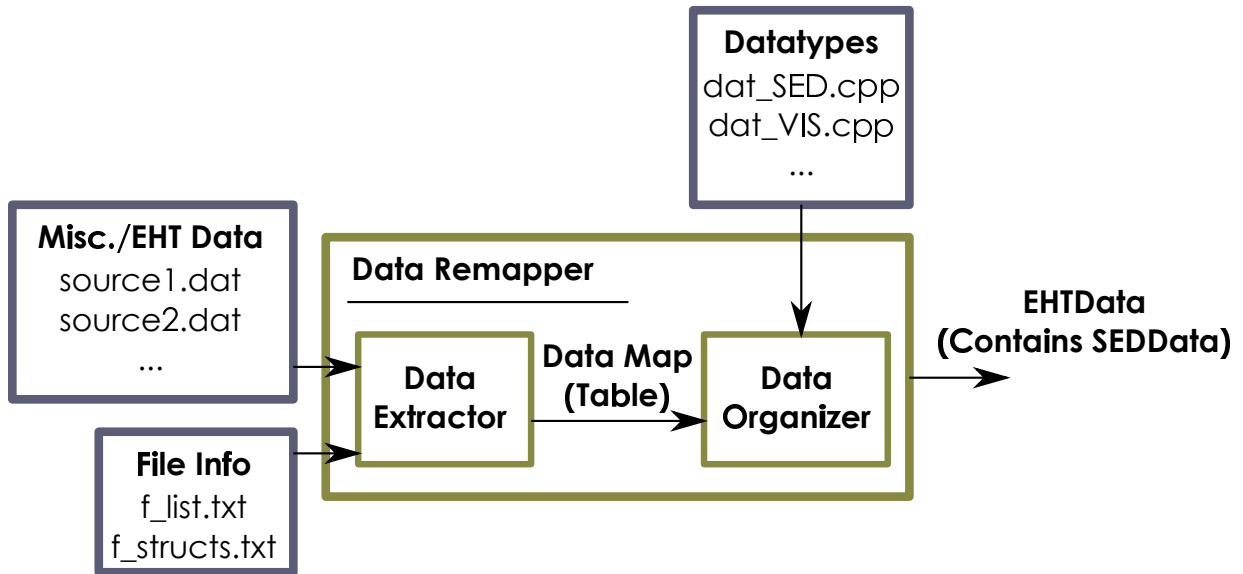


Figure 4.4: Data remapper component of THEMIS. Imports all observed data (using the file and structure inputs) and remaps them with their corresponding data type class, then combines each observed data object into a single object class called `EHTData`. This data object also contains other data observables that is not part of the EHT, such as SED data and black hole masses and distances.

```
EHTData.<Data Type>.<Filter Type>(<Filter Params>).
```

It is interesting to note that the filter type changes for each data type, depending on what information is supplied. For example, `filter_u` filters for u data for visibility magnitudes and `filter_lat` filters for latitudes for observatories. There are also generic filter types, e.g. `filter_fname` filters for the file name of particular source. It is not necessary to apply filters and all data of a given type can be accessed simply through

```
EHTData.<Data Type>.
```

`EHTData` can add additional files and data, or remove existing data as well.

New data source files can be included by adding the source location in the `f_list` and `f_struct` files from the file info block. If a new data type is to be added (such as time dependent multi-wavelength SEDs/spectra/Astrometry, closure amplitudes (Section 1.1.4),

polarization, etc.), there exists an automated code generation program which generates the corresponding data type classes, reading in a data types file which stores all of the current data type information. This file can be modified to add new data types or modify existing ones. Once the new data type classes are generated, the generated files can then be copied over to the folder containing the other files. The process to add an additional data type or source therefore is simplified, streamlined, and homogenized, eliminating any future overhead development costs when dealing with the implementation of new data.

4.1.3 Posterior Probability Estimator

The core component of the THEMIS pipeline is the posterior probability estimator (Figure 4.5). It is the central hub of operators, generating model images, comparing the model images with observed data and then outputting a posterior probability distribution in the form of a MCMC chain. It applies MCMC to the accretion model and observations, following the procedure described in Section 4.1. An optional non-linear optimizer ¹ can be used to determine a good starting point to speed up the burn-in process. Because the model image generation is computationally expensive (see Section 4.1.1), it is favourable to reduce the amount of iterations on the MCMC chain for convergence. There also exists many features to speed up development costs, such as the abilities to easily swap between models and choose which parameters in the parameter set to permit the MCMC pipeline to search through.

While likelihood estimates are obtained for each iteration are interpreted as the probability that the model fits the data, the output distribution of MCMC is the probability that the data fits the model. This is simply described by Bayes' theorem:

$$P(M|D) = \frac{P(D|M)P(M)}{P(D)}, \quad (4.1)$$

where $P(M|D)$ denotes the desired probability of the model M given the data D . Here the statistical likelihood (probability of the data given the model, $P(D|M)$) is obtained from the χ^2 . The marginal likelihood, $P(D)$, is simply a normalization applied during the construction of $P(M|D)$ from the chain data. The priors $P(M)$ are applied during the trial step selection. MCMC with Metropolis Hastings algorithm along with adaptive stepping towards the principle directions is currently implemented and used for THEMIS.

¹NLopt - <http://ab-initio.mit.edu/wiki/index.php/NLopt>

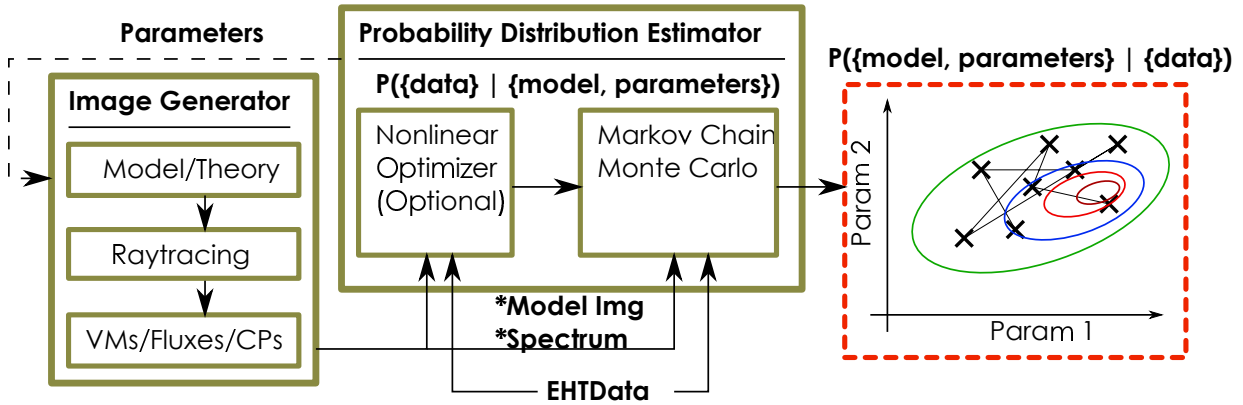


Figure 4.5: Posterior probability estimator component of THEMIS. This is the central hub where all of the main operations are executed, and is comprised of the MCMC function along with an optional non-linear optimizer.

4.2 Future Prospects

Rigorous testing was done for the THEMIS pipeline to test each component individually. Both theoretical images and likelihood (χ^2) estimates were successfully repeated from [10], and both the MCMC and nonlinear optimizer was tested to ensure proper functionality. Every component of the THEMIS system was integrated together successfully along with MPI compatibility for the execution using high performance computing (HPC) networks. While the disk height ratio parameter is incorporated, the sub-Keplerian parameter is yet to be. Additional ways to reduce the convergence time are being explored, some of which involves the MCMC (Adaptive MCMC, Hamiltonian MCMC, etc.), and another take advantage of the physical nature of some parameters to reduce image generation time (e.g. ξ is a simple “rotation” of the resultant model image).

Preliminary MCMC runs are currently being held in the Perimeter Institute’s computational cluster network, along with eventual plans to port the software to SHARCNET², a high performance computing (HPC) organization for researchers in Canada.

²<https://www.sharcnet.ca/>

Chapter 5

Conclusion

Sgr A* offers a convenient target for probing accretion flow and jet formation models in the context of low-luminosity AGN. While this has been traditionally done using spectral and polarimetric observations, it has recently become possible to do so using structural information obtained from the EHT. Motivated by the addition of many years of new closure phase data and continued improvements in the capabilities of future EHT observations, we have considered the implications for GR signatures. This takes the form of probes of parametrized deviations of Kerr and searches (Section 2) for long-timescale variability (Section 4). The latter has the additional advantage of validating key simplifying assumptions employed in all analyses to date. At the same time a new analysis capability has been developed, needed to make use of the high-quality of current and future EHT observations and the increasing complexity of the required astrophysical models.

The first strong gravity test of general relativity with the introduction of a quasi-Kerr metric deviation parameter ϵ , uses visibility magnitudes from the 2007-2009 EHT observing runs to produce weak initial constraints of order unity [12]. Improved constraints were obtained for this thesis, adding the 2009-2013 closure phase to the analysis to obtain 1σ limits of $a_* = 0.05^{+0.30}_{-0.05}$, $\theta = 70.0^{+12.5^\circ}_{-0.05^\circ}$, and $\epsilon = 0.5^{+0.45}_{-0.25}$. The position angle ξ limits were not reported due to an offset induced by a closure phase noise bias, which has since been attributed to the interstellar galactic scattering and is corrected in an upcoming paper. While the current ϵ constraint favours a non-Kerr-like black hole, it is clear in the posterior probability distribution of a_* and ϵ (Figure 2.4b) that the restriction of spin to positive values produces a positive bias in the quadrupolar deviation parameter. Another notable improvement is that the degeneracy produced by the ISCO in the posterior probability distribution of a_* and ϵ appears to be constrained with the added closure phase information. Simulations of additional stations being included in future EHT observations placed tight

3σ limits: $a_* = 0.150_{-0.005}^{+0.004}$, $\theta = 60.01_{-0.06^\circ}^{+0.09^\circ}$, $\xi = 159.99^\circ \pm 0.04^\circ$, and $\epsilon = 0 \pm 0.005$, thus strongly motivating both future EHT observations and future theoretical development of astrophysical models.

The short intrinsic timescale (as short as 10 yr) of the stellar winds from O/WR stars [14] that are believed to source Sgr A*'s accretion flow suggests the presence of a tilted disk. Thus, we have undertaken an search for variations in Sgr A*'s horizon scale structure due to disk misalignment. Employing the 2007 to 2013 visibility magnitude and closure phase information from the EHT, as well as the RIAF model described in [10], the search for any and Lense-Thirring rigid-body precession resulted in firm precession period and outer disk radius constraints. The analyses yielded no discernible detection of precession from 2007-2013, and a high probability for any precession period of high magnitudes. With Lense-Thirring torques as a driving mechanism, 2σ limits were placed on the black hole spin and accretion flow size: $a \lesssim 0.575$ and $r_o \gtrsim 300M$. Two possible scenarios can still allow for disk misalignment: one being that the black hole spin is very small and/or the accretion flow is very large. It is uncertain how a misaligned disk can exist with a small spin. However, it remains possible that a stationary misalignment can be supported by the supply of fixed angular momentum material from the stellar disk, forming a stationary warped accretion flow. This is argued against by current RIAF analyses which find alignment between the black hole spin and the clockwise stellar disk containing the O/WR stars.

The thesis concludes with the development of a software pipeline, THEMIS, designed to replacing the existing RIAF analysis pipeline and its variants, used in prior analyses with the EHT data, e.g. the work done in [10, 11, 12]. THEMIS seeks to ameliorate the major limitations from the existing pipeline, two of which are adding extensibility to the way new observations and models are introduced and improving the search time significantly with the replacement of grid-based search methods with MCMC methods and therefore permitting higher dimensional parameter spaces to be explored. Following post-integration testing, THEMIS will be ready to be used to constrain two additional RIAF parameters, the disk height ratio and sub-Keplerian parameter. Further exploration of clever MCMC implementation is underway (e.g. Adaptive MCMC, Hamiltonian MCMC, the efficient use of parameters, etc.), though schemes to reduce the burn-in time (e.g. choosing a “good” starting point through a nonlinear optimizer) and time-to-convergence have already been implemented.

References

- [1] D. K. Aitken, J. Greaves, A. Chrysostomou, T. Jenness, W. Holland, J. H. Hough, D. Pierce-Price, and J. Richer. Detection of Polarized Millimeter and Submillimeter Emission from Sagittarius A*. *ApJ*, 534:L173–L176, May 2000.
- [2] P. J. Armitage and P. Natarajan. Lense-Thirring Precession of Accretion Disks around Compact Objects. *ApJ*, 525:909–914, November 1999.
- [3] J. M. Bardeen and J. A. Petterson. The Lense-Thirring Effect and Accretion Disks around Kerr Black Holes. *ApJ*, 195:L65, January 1975.
- [4] H. Bondi. On spherically symmetrical accretion. *MNRAS*, 112:195, 1952.
- [5] G. C. Bower, W. M. Goss, H. Falcke, D. C. Backer, and Y. Lithwick. The Intrinsic Size of Sagittarius A* from 0.35 to 6 cm. *ApJ*, 648:L127–L130, September 2006.
- [6] G. C. Bower, M. C. H. Wright, H. Falcke, and D. C. Backer. BIMA Observations of Linear Polarization in Sagittarius A* at 112 GHz. *ApJ*, 555:L103–L106, July 2001.
- [7] G. C. Bower, M. C. H. Wright, H. Falcke, and D. C. Backer. Interferometric Detection of Linear Polarization from Sagittarius A* at 230 GHz. *ApJ*, 588:331–337, May 2003.
- [8] A. Broderick and R. Blandford. Covariant magnetoionic theory - II. Radiative transfer. *MNRAS*, 349:994–1008, April 2004.
- [9] A. E. Broderick, V. L. Fish, S. S. Doeleman, and A. Loeb. Estimating the Parameters of Sagittarius A*'s Accretion Flow Via Millimeter VLBI. *ApJ*, 697:45–54, May 2009.
- [10] A. E. Broderick, V. L. Fish, S. S. Doeleman, and A. Loeb. Evidence for Low Black Hole Spin and Physically Motivated Accretion Models from Millimeter-VLBI Observations of Sagittarius A*. *ApJ*, 735:110, July 2011.

- [11] A. E. Broderick, V. L. Fish, M. D. Johnson, K. Rosenfeld, et al. Modeling Seven Years of Event Horizon Telescope Observations with Radiatively Inefficient Accretion Flow Models. 2015.
- [12] A. E. Broderick, T. Johannsen, A. Loeb, and D. Psaltis. Testing the No-hair Theorem with Event Horizon Telescope Observations of Sagittarius A*. *ApJ*, 784:7, March 2014.
- [13] A. E. Broderick and A. Loeb. Frequency-dependent Shift in the Image Centroid of the Black Hole at the Galactic Center as a Test of General Relativity. *ApJ*, 636:L109–L112, January 2006.
- [14] J. Cuadra, S. Nayakshin, and F. Martins. Variable accretion and emission from the stellar winds in the Galactic Centre. *MNRAS*, 383:458–466, January 2008.
- [15] J. Dexter, E. Agol, and P. C. Fragile. Millimeter Flares and VLBI Visibilities from Relativistic Simulations of Magnetized Accretion Onto the Galactic Center Black Hole. *ApJ*, 703:L142–L146, October 2009.
- [16] J. Dexter, E. Agol, P. C. Fragile, and J. C. McKinney. The Submillimeter Bump in Sgr A* from Relativistic MHD Simulations. *ApJ*, 717:1092–1104, July 2010.
- [17] J. Dexter and P. C. Fragile. Observational Signatures of Tilted Black Hole Accretion Disks from Simulations. *ApJ*, 730:36, March 2011.
- [18] J. Dexter and P. C. Fragile. Tilted black hole accretion disc models of Sagittarius A*: time-variable millimetre to near-infrared emission. *MNRAS*, 432:2252–2272, July 2013.
- [19] S. Dibi, S. Markoff, R. Belmont, J. Malzac, N. M. Barrière, and J. A. Tomsick. Exploring plasma evolution during Sagittarius A* flares. *MNRAS*, 441:1005–1016, June 2014.
- [20] K. Dodds-Eden, P. Sharma, E. Quataert, R. Genzel, S. Gillessen, F. Eisenhauer, and D. Porquet. Time-Dependent Models of Flares from Sagittarius A*. *ApJ*, 725:450–465, December 2010.
- [21] S. S. Doeleman, J. Weintroub, A. E. E. Rogers, R. Plambeck, R. Freund, R. P. J. Tilanus, P. Friberg, L. M. Ziurys, J. M. Moran, B. Corey, K. H. Young, D. L. Smythe, M. Titus, D. P. Marrone, R. J. Cappallo, D. C.-J. Bock, G. C. Bower, R. Chamberlin,

- G. R. Davis, T. P. Krichbaum, J. Lamb, H. Maness, A. E. Niell, A. Roy, P. Strittmatter, D. Werthimer, A. R. Whitney, and D. Woody. Event-horizon-scale structure in the supermassive black hole candidate at the Galactic Centre. *Nature*, 455:78–80, September 2008.
- [22] A. Eckart, M. García-Marín, S. N. Vogel, P. Teuben, M. R. Morris, F. Baganoff, J. Dexter, R. Schödel, G. Witzel, M. Valencia-S, V. Karas, D. Kunneriath, M. Bremer, C. Straubmeier, L. Moser, N. Sabha, R. Buchholz, M. Zamaninasab, K. Mužić, J. Moulta, and J. A. Zensus. Flare emission from Sagittarius A*. *Journal of Physics Conference Series*, 372(1):012022, July 2012.
- [23] V. L. Fish, S. S. Doeleman, C. Beaudoin, R. Blundell, D. E. Bolin, G. C. Bower, R. Chamberlin, R. Freund, P. Friberg, M. A. Gurwell, M. Honma, M. Inoue, T. P. Krichbaum, J. Lamb, D. P. Marrone, J. M. Moran, T. Oyama, R. Plambeck, R. Primiani, A. E. E. Rogers, D. L. Smythe, J. SooHoo, P. Strittmatter, R. P. J. Tilanus, M. Titus, J. Weintraub, M. Wright, D. Woody, K. H. Young, and L. M. Ziurys. 1.3 mm Wavelength VLBI of Sagittarius A*: Detection of Time-variable Emission on Event Horizon Scales. *ApJ*, 727:L36, February 2011.
- [24] V. L. Fish et al. Variable asymmetric structure of Sagittarius A* on event horizon scales, in preparation. 2015.
- [25] P. C. Fragile. Effective Inner Radius of Tilted Black Hole Accretion Disks. *ApJ*, 706:L246–L250, December 2009.
- [26] P. C. Fragile and P. Anninos. Hydrodynamic Simulations of Tilted Thick-Disk Accretion onto a Kerr Black Hole. *ApJ*, 623:347–361, April 2005.
- [27] P. C. Fragile and O. M. Blaes. Epicyclic Motions and Standing Shocks in Numerically Simulated Tilted Black Hole Accretion Disks. *ApJ*, 687:757–766, November 2008.
- [28] P. C. Fragile, O. M. Blaes, P. Anninos, and J. D. Salmonson. Global General Relativistic Magnetohydrodynamic Simulation of a Tilted Black Hole Accretion Disk. *ApJ*, 668:417–429, October 2007.
- [29] A. Generozov, O. Blaes, P. C. Fragile, and K. B. Henisey. Physical Properties of the Inner Shocks in Hot, Tilted Black Hole Accretion Flows. *ApJ*, 780:81, January 2014.
- [30] A. M. Ghez, S. Salim, N. N. Weinberg, J. R. Lu, T. Do, J. K. Dunn, K. Matthews, M. R. Morris, S. Yelda, E. E. Becklin, T. Kremenek, M. Milosavljevic, and J. Naiman.

Measuring Distance and Properties of the Milky Way's Central Supermassive Black Hole with Stellar Orbits. *ApJ*, 689:1044–1062, December 2008.

- [31] S. Gillessen, F. Eisenhauer, T. K. Fritz, H. Bartko, K. Dodds-Eden, O. Pfuhl, T. Ott, and R. Genzel. The Orbit of the Star S2 Around SGR A* from Very Large Telescope and Keck Data. *ApJ*, 707:L114–L117, December 2009.
- [32] S. Gillessen, F. Eisenhauer, E. Quataert, R. Genzel, T. Paumard, S. Trippe, T. Ott, R. Abuter, A. Eckart, P. O. Lagage, M. D. Lehnert, L. J. Tacconi, and F. Martins. Variations in the Spectral Slope of Sagittarius A* during a Near-Infrared Flare. *ApJ*, 640:L163–L166, April 2006.
- [33] S. Gillessen, F. Eisenhauer, S. Trippe, T. Alexander, R. Genzel, F. Martins, and T. Ott. Monitoring Stellar Orbits Around the Massive Black Hole in the Galactic Center. *ApJ*, 692:1075–1109, February 2009.
- [34] K. Glampedakis and S. Babak. Mapping spacetimes with LISA: inspiral of a test body in a 'quasi-Kerr' field. *Classical and Quantum Gravity*, 23:4167–4188, June 2006.
- [35] J. B. Hartle. Slowly Rotating Relativistic Stars. I. Equations of Structure. *ApJ*, 150:1005, December 1967.
- [36] J. B. Hartle and K. S. Thorne. Slowly Rotating Relativistic Stars. II. Models for Neutron Stars and Supermassive Stars. *ApJ*, 153:807, September 1968.
- [37] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [38] L. Huang, R. Takahashi, and Z.-Q. Shen. Testing the Accretion Flow with Plasma Wave Heating Mechanism for Sagittarius A* by the 1.3 mm VLBI Measurements. *ApJ*, 706:960–969, December 2009.
- [39] T. Johannsen. Photon Rings around Kerr and Kerr-like Black Holes. *ApJ*, 777:170, November 2013.
- [40] T. Johannsen and D. Psaltis. Testing the No-hair Theorem with Observations in the Electromagnetic Spectrum. I. Properties of a Quasi-Kerr Spacetime. *ApJ*, 716:187–197, June 2010.
- [41] T. Johannsen and D. Psaltis. Testing the No-hair Theorem with Observations in the Electromagnetic Spectrum. II. Black Hole Images. *ApJ*, 718:446–454, July 2010.

- [42] T. Johannsen, D. Psaltis, S. Gillessen, D. P. Marrone, F. Özel, S. S. Doeleman, and V. L. Fish. Masses of nearby Supermassive Black Holes with Very Long Baseline Interferometry. *ApJ*, 758:30, October 2012.
- [43] T Johannsen, C. Wang, A.E. Broderick, et al. Testing General Relativity with the Event Horizon Telescope, in preparation. 2015.
- [44] M. D. Johnson, V. L. Fish, S. S. Doeleman, D. P. Marrone, et al. Resolved magnetic-field structure and variability near the event horizon of Sagittarius A*. *Science*, (submitted), 2015.
- [45] R. P. Kerr. Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics. *Physical Review Letters*, 11:237–238, September 1963.
- [46] S. S. Komissarov. Magnetized tori around Kerr black holes: analytic solutions with a toroidal magnetic field. *MNRAS*, 368:993–1000, May 2006.
- [47] J. Kormendy and L. C. Ho. Coevolution (Or Not) of Supermassive Black Holes and Host Galaxies. *ARAA*, 51:511–653, August 2013.
- [48] S. Kumar and J. E. Pringle. Twisted accretion discs - The Bardeen-Petterson effect. *MNRAS*, 213:435–442, March 1985.
- [49] A. Loeb and E. Waxman. Properties of the radio-emitting gas around Sgr A*. *Journal of Cosmology and Astroparticle Physics*, 3:11, March 2007.
- [50] J.-P. Macquart, G. C. Bower, M. C. H. Wright, D. C. Backer, and H. Falcke. The Rotation Measure and 3.5 Millimeter Polarization of Sagittarius A*. *ApJ*, 646:L111–L114, August 2006.
- [51] D. P. Marrone, J. M. Moran, J.-H. Zhao, and R. Rao. The Submillimeter Polarization of Sgr A*. *Journal of Physics Conference Series*, 54:354–362, December 2006.
- [52] D. P. Marrone, J. M. Moran, J.-H. Zhao, and R. Rao. An Unambiguous Detection of Faraday Rotation in Sagittarius A*. *ApJ*, 654:L57–L60, January 2007.
- [53] Daniel P Marrone. *Submillimeter Properties of Sagittarius A*: The Polarization and Spectrum from 230 to 690 GHz and the Submillimeter Array Polarimeter*. PhD thesis, Harvard University, 2006.

- [54] J. C. McKinney, A. Tchekhovskoy, and R. D. Blandford. Alignment of Magnetized Accretion Disks and Relativistic Jets with Spinning Black Holes. *Science*, 339:49–, January 2013.
- [55] B. R. McNamara and P. E. J. Nulsen. Mechanical feedback from active galactic nuclei in galaxies, groups and clusters. *New Journal of Physics*, 14(5):055023, May 2012.
- [56] D. Morales Teixeira, P. C. Fragile, V. V. Zhuravlev, and P. B. Ivanov. Conservative GRMHD Simulations of Moderately Thin, Tilted Accretion Disks. *ApJ*, 796:103, December 2014.
- [57] M. Mościbrodzka, C. F. Gammie, J. C. Dolence, H. Shiokawa, and P. K. Leung. Radiative Models of SGR A* from GRMHD Simulations. *ApJ*, 706:497–507, November 2009.
- [58] P. Natarajan and J. E. Pringle. The Alignment of Disk and Black Hole Spins in Active Galactic Nuclei. *ApJ*, 506:L97–L100, October 1998.
- [59] R. Nealon, D. J. Price, and C. J. Nixon. On the Bardeen-Petterson effect in black hole accretion discs. *MNRAS*, 448:1526–1540, April 2015.
- [60] R. P. Nelson and J. C. B. Papaloizou. Hydrodynamic simulations of propagating WARPS and bending waves in accretion discs. *MNRAS*, 309:929–940, November 1999.
- [61] R. P. Nelson and J. C. B. Papaloizou. Hydrodynamic simulations of the Bardeen-Petterson effect. *MNRAS*, 315:570–586, July 2000.
- [62] C. Nixon, A. King, D. Price, and J. Frank. Tearing up the Disk: How Black Holes Accrete. *ApJ*, 757:L24, October 2012.
- [63] C. J. Nixon and A. R. King. Broken discs: warp propagation in accretion discs. *MNRAS*, 421:1201–1208, April 2012.
- [64] J. E. Pringle. Accretion discs in astrophysics. *ARAA*, 19:137–162, 1981.
- [65] J. E. Pringle. A simple approach to the evolution of twisted accretion discs. *MNRAS*, 258:811–818, October 1992.
- [66] M. J. Rees. Quasars. *The Observatory*, 98:210–223, October 1978.

- [67] P. A. G. Scheuer and R. Feiler. The realignment of a black hole misaligned with its accretion disc. *MNRAS*, 282:291, September 1996.
- [68] R. V. Shcherbakov, R. F. Penna, and J. C. McKinney. Sagittarius A* Accretion Flow and Black Hole Parameters from General Relativistic Dynamical and Polarized Radiative Modeling. *ApJ*, 755:133, August 2012.
- [69] A. Tchekhovskoy, R. Narayan, and J. C. McKinney. Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole. *MNRAS*, 418:L79–L83, November 2011.
- [70] A. R. Thompson, J. M. Moran, and G. W. Swenson, Jr. *Interferometry and Synthesis in Radio Astronomy, 2nd Edition*. 2001.
- [71] F. H. Vincent, T. Paumard, E. Gourgoulhon, and G. Perrin. GYOTO: a new general relativistic ray-tracing code. *Classical and Quantum Gravity*, 28(22):225011, November 2011.
- [72] F. H. Vincent, W. Yan, O. Straub, A. A. Zdziarski, and M. A. Abramowicz. A magnetized torus for modeling Sagittarius A* millimeter images and spectra. *AAP*, 574:A48, January 2015.
- [73] C. M. Will. The Confrontation between General Relativity and Experiment. *Living Reviews in Relativity*, 17:4, June 2014.
- [74] F. Yuan, E. Quataert, and R. Narayan. Nonthermal Electrons in Radiatively Inefficient Accretion Flow Models of Sagittarius A*. *ApJ*, 598:301–312, November 2003.
- [75] F. Yuan, E. Quataert, and R. Narayan. On the Nature of the Variable Infrared Emission from Sagittarius A*. *ApJ*, 606:894–899, May 2004.
- [76] Y.-F. Yuan, X. Cao, L. Huang, and Z.-Q. Shen. Images of the Radiatively Inefficient Accretion Flow Surrounding a Kerr Black Hole: Application in Sgr A*. *ApJ*, 699:722–731, July 2009.
- [77] V. V. Zhuravlev, P. B. Ivanov, P. C. Fragile, and D. Morales Teixeira. No Evidence for Bardeen-Petterson Alignment in GRMHD Simulations and Semi-analytic Models of Moderately Thin, Prograde, Tilted Accretion Disks. *ApJ*, 796:104, December 2014.