

Patterns in Space and Time
Investigating the Structure and Variability of Sagittarius A*
with Radio Interferometry

Christiaan Dirk Brinkerink

Patterns in Space and Time
Investigating the Structure and Variability of Sagittarius A*
with Radio Interferometry

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.H.J.M. van Krieken,
volgens besluit van het college van decanen
in het openbaar te verdedigen op
woensdag 22 september 2021
om 12.30 uur precies

door

Christiaan Dirk Brinkerink

geboren op 30 april 1980
te Neede

PROMOTOR: Prof. dr. H.D.E. Falcke

MANUSCRIPTCOMMISSIE: Prof. dr. H.T. Koelink

Prof. dr. A.J. Levan

Dr. E.G. Körding

Dr. I.M. van Bemmelen

Prof. dr. H.J. van Langevelde

Joint Institute for VLBI ERIC

Universiteit Leiden

© 2021, Christiaan Dirk Brinkerink

Patterns in Space and Time: Investigating the Structure and Variability of Sagittarius A* with Radio Interferometry

Thesis, Radboud University Nijmegen

Front cover: A stylised view of VLA antennas and the Galactic Center 'minispiral', generated by the author.

Back cover: Stylised view of a footpath in 'Heerlijkheid Beek', generated by the author.

Cover & pagemarket design: Christiaan Brinkerink

Illustrated; with bibliographic information and Dutch summary

ISBN: 978-94-6421-423-9

CONTENTS

1	Introduction	1
1.1	Gravity and black holes	1
1.2	The phenomenon of accretion	5
1.2.1	Astrophysical context	5
1.2.2	The Eddington luminosity	6
1.2.3	The Eddington accretion rate	7
1.2.4	Models for accretion at different Eddington ratios	9
1.2.5	The Mdot-Sigma plane	10
1.2.6	Disk outflows and jets	13
1.3	Observational signatures of black hole accretion	15
1.4	History of Sagittarius A*	19
1.4.1	Early detections and interpretation	19
1.4.2	The spectrum of Sgr A*	20
1.4.3	Spectral models for Sgr A*	20
1.4.4	Morphology	23
1.4.5	Polarisation	24
1.5	Measurement of radio waves, radio interferometry and VLBI	26
1.5.1	Flux density in radio	26
1.5.2	Signal amplification and filtering	27
1.5.3	Radio interferometry	27
1.5.4	Measuring the interference pattern: correlation	29
1.5.5	Aperture synthesis	33
1.5.6	Data calibration	35
1.5.7	Useful interferometric observables: closure phases and closure amplitudes	36
1.6	In this thesis	36
2	Asymmetric structure in Sgr A* at 3 mm from closure phase measurements with VLBA, GBT, and LMT	39
2.1	Introduction	39

2.2	Observations and initial data reduction	42
2.3	Verifying the nature of non-zero closure phases	45
2.4	Results	46
2.4.1	Detection of non-zero closure phases	46
2.4.2	Modeling source asymmetry using closure phases	46
2.4.3	Testing the significance of the observed asymmetry	48
2.5	Discussion	49
2.6	Summary and conclusions	57
3	Micro-arcsecond structure of Sagittarius A* revealed by high-sensitivity 86 GHz VLBI observations	59
3.1	Introduction	60
3.2	Observations and data reduction	63
3.3	Results	65
3.3.1	Mapping and self-calibration of Sgr A*	65
3.3.2	Constraining the size of Sgr A* using closure amplitudes	69
3.4	Constraints on the size-frequency relation and the scattering law	74
3.5	Summary and conclusions	79
3.A	Closure amplitude model fitting technique	82
4	Frequency-dependent time lags in Sgr A*	83
4.1	Introduction	84
4.2	Observations and data reduction	85
4.3	Analysis and results	87
4.4	Discussion and conclusion	92
4.5	Acknowledgements	96
5	Persistent time lags in light curves of Sagittarius A*: evidence of outflow	99
5.1	Introduction	100
5.1.1	Observed spectrum	100
5.1.2	Observed morphology	100
5.1.3	Time-domain studies	101
5.1.4	Questions addressed in this work	103
5.2	Observations	103
5.2.1	Observation epochs, array configuration and spectral setup	103
5.2.2	Data calibration and reduction	104
5.3	Data analysis	105
5.3.1	Extracting light curves	105
5.3.2	Detection of time lags	110
5.3.3	Checks with synthetic data	113
5.4	Results	118

5.5 Discussion and Conclusions	119
6 Summary	137
7 Samenvatting	141
A Research Data Management	145
Bibliography	149
List of Publications	155
About the author	161
Acknowledgments	163

To Mum

The truth may be puzzling. It may take some work to grapple with. It may be counterintuitive. It may contradict deeply held prejudices. It may not be consonant with what we desperately want to be true. But our preferences do not determine what's true. We have a method, and that method helps us to reach not absolute truth, only asymptotic approaches to the truth — never there, just closer and closer, always finding vast new oceans of undiscovered possibilities. Cleverly designed experiments are the key.

Carl Sagan

INTRODUCTION

This thesis treats the study of the supermassive black hole at the center of our Galaxy, which we commonly refer to as 'Sagittarius A*', through the use of several measurement techniques that involve radio telescopes. In this introduction I will first of all highlight what black holes are, and why we study them. I will then go into more detail on Sagittarius A* in particular. Moving on to the subject of instrumentation and measurements, I will illustrate how our use of radio telescopes helps us find out about black hole properties and the behaviour of the gas that surrounds them. This introduction concludes by briefly describing the subjects of the following chapters.

1.1 Gravity and black holes

Gravity as a phenomenon is intuitively understood by humans. We understand that things fall faster when they fall farther, and our sense (or fear) of heights is intimately linked to this understanding. For instance, in our childhood we quickly learn how to catch a ball that is thrown toward us. This is quite an impressive feat – catching a ball involves modeling its accelerated motion and predicting its trajectory at a subconscious level. In his work "Dialogues Concerning Two New Sciences", Galileo Galilei systematically quantified the behaviour of falling bodies and investigated what factors influence the rate at which different objects fall. He found¹ that gravitational acceleration, when measured at the Earth's surface, means that the vertical speed of a falling object changes by a specific amount per time interval - and not, say, per vertical distance interval. Isaac Newton applied a quantitative model of gravity to celestial objects with great success, showing that the motions of the heavenly bodies follow the same rules as everyday objects falling here on Earth. Our picture of gravity underwent yet another revolution with Albert Einstein's work, who modeled gravity as a geometrical phenomenon: a warping of both space and time that is linked to the distribution of matter and energy (Einstein, 1915). In the theory of general relativity, spacetime is not simply the backdrop against which the events of the universe play out. Rather, its structure is itself affected by the presence and behaviour of mass and energy in it. This radically different view of gravity would provide exciting new ways to study the

¹Although he was not the first one to find this - that honour belongs to Nicole Oresme, see Clagett (1968).

gravitational field of extremely compact objects.

In order to study the structure that spacetime itself has, we need to have a way to express the relation between different events in it. An event is a combination of a time and a place, and has a set of unique coordinates in any map of spacetime. When we wish to express the distance between two events, we already know how to do this for the spatial distance ("the distance between place A and place B is 5 kilometres") and we know how to do this for the temporal distance ("Moment X happens 2 seconds before moment Y"). However, we can use the speed of light to express separation in time and separation in space in the same way and combine them, so that the spacetime distance between two events can be expressed as a single number. Distances in the three spatial dimensions can be combined in the usual way using Pythagoras' rule, $dr = \sqrt{dx^2 + dy^2 + dz^2}$, but time needs to be treated somewhat differently. We will also need to have some way to deal with variations in the structure of spacetime from place to place as we consider different scenarios. This is where the language of relativity comes in.

In both special and general relativity, we can express the geometry of spacetime in terms of a quantity called the invariant spacetime interval, which gives us the spacetime distance between two events in spacetime. If one starts with the postulate that the local speed of light as measured by any observer always has the same value (a notion that is strongly supported by our measurements), one arrives at a particular form in which time and space can be combined into a geometry of spacetime. For 'flat' (Minkowski) spacetime, which is devoid of matter and energy, the invariant spacetime interval can be written out as:

$$ds^2 = -c^2 d\tau^2 = -d(ct)^2 + dx^2 + dy^2 + dz^2, \quad (1.1)$$

where ds^2 is the length squared of the interval, c is the speed of light, $d\tau$ is the proper time interval (which refers to the time interval between the events as measured by an observer moving from one event to the other on a straight, non-accelerated trajectory or a geodesic), and the dt, dx, dy, dz terms refer to the coordinate differences between two events in time and the three spatial directions respectively, in units that are determined by how c is expressed. This expression immediately tells us something very interesting: the measurement of time can work out differently for different observers. Depending on their state of motion, different observers will measure different values of dx, dy and dz for the separation between a given pair of events in their respective reference frames. But the spacetime interval that they should both measure is invariant, so they must therefore get different values for dt as well: as measured in their own reference frames, they get different values for the time interval that passes between the two events. This gives rise to the phenomenon of 'time dilation', where observers measure that the rate at which time passes for other observers moving with respect to them is slowed down. To make the difference clear, t is also referred to as 'coordinate time' because it depends on the choice of coordinate system.

Flat spacetime is a special case – if we consider non-empty spacetime, the invariant interval can look quite different. Making the move to general relativity, the general expression for the invariant spacetime interval (hereafter simply called 'interval') is:

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = g_{00} dx^0 dx^0 + g_{01} dx^0 dx^1 + g_{02} dx^0 dx^2 + \dots + g_{33} dx^3 dx^3, \quad (1.2)$$

where we have illustrated the Einstein summation convention (where a repeated index that appears in an upper and a lower form indicates a summation over all values of that index) and μ and ν both run over 4 dimensions which are labeled 0, 1, 2, and 3 (one is temporal and three are spatial). This notation needs some explaining! In the above expression ds^2 is the length squared of the spacetime interval that is associated with an infinitesimal displacement in spacetime along the four dimensions, subtly different from the case for flat spacetime where this displacement didn't need to be infinitesimal. The components of this infinitesimal displacement are indicated by dx^0 thru dx^3 , where the dimension with index 0 indicates the time dimension by convention. The term $g_{\mu\nu}$ represents the spacetime metric. The metric written down in this form is a covariant tensor with two indices, both of which run over all four spacetime coordinates - the metric thus has 16 components, 10 of which are independent in General Relativity (the metric is symmetric: $g_{\mu\nu} = g_{\nu\mu}$). It encodes the local structure of spacetime, and when 'fed' with two copies of an infinitesimal displacement four-vector it yields a scalar: the spacetime distance between two infinitesimally separated spacetime events. This infinitesimal distance can be integrated along a path through spacetime to get the spacetime separation for any pair of events. Choosing the path so that it maximises the proper time elapsed gives us a geodesic, or the equivalent of a straight line in curved spacetime.

Revisiting the metric for flat spacetime, we see that it only has four nonzero components: $g_{00} = -1, g_{11} = g_{22} = g_{33} = 1$. We see that the interval is real for spacelike separations, where the positive spatial term $dx^2 + dy^2 + dz^2$ wins out over the negative temporal term $-d(ct)^2$. The interval is zero for lightlike separations, where a pulse of light sent from one event precisely reaches the other. Finally, the interval is imaginary for timelike separations, which is the class of separations where one event can communicate to the other using a message travelling slower than light. It should be noted that the sign convention for the space and time components of the invariant interval can be chosen oppositely as well, but this makes no difference to the physics that follows from it. Choosing one representation over the other is a matter of convention and taste.

When we now consider a point mass at the origin, it should affect the geometry of spacetime in a spherically symmetric way as no spatial direction is preferred over another. It therefore makes sense to express the interval using spherical coordinates. When we do this first for flat spacetime, we get the following expression:

$$ds^2 = -d(ct)^2 + dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2, \quad (1.3)$$

where r is the radial coordinate, θ is the polar angle and ϕ is the azimuthal angle. The solution to the geometry of spacetime that follows from the Einstein equations for a stationary point mass in

a vacuum was derived by Karl Schwarzschild in 1916, just months after Einstein had published his paper introducing general relativity. This metric is known as the Schwarzschild metric, and its invariant interval can be written in spherical coordinates as:

$$ds^2 = -c^2 d\tau^2 = -(1 - \frac{2GM}{c^2 r})d(ct)^2 + (1 - \frac{2GM}{c^2 r})^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2, \quad (1.4)$$

where M is the mass of the central object and G is the universal gravitational constant. We see that the factors in the radial and temporal terms have changed with respect to those for flat spacetime. We also notice that something strange happens at a radius of $\frac{2GM}{c^2}$: the factors in the t (temporal) and r (radial) terms tend to zero and infinity, respectively. We are dealing with a coordinate singularity at this radius: a place where this coordinate system stops working and a different one is needed to give us insight into what physically happens there.

After Schwarzschild's solution got published, it took decades of debate before the nature of this special radius became fully apparent (Finkelstein, 1958). It turned out that this critical radius $r = \frac{2GM}{c^2}$ defines a so-called 'event horizon', a one-way virtual boundary that prevents any causal contact from its interior to the rest of the universe but still allows for matter or energy to enter the enclosed region. This behaviour cannot readily be seen when considering the spacetime in Schwarzschild coordinates, which is why it took some time before it was worked out using different choices of coordinates (Misner et al., 1973; Finch, 2015).

Besides the one-way character of the event horizon, there are other strange effects that manifest themselves when matter approaches the horizon as seen by a faraway observer. One of these effects is that of extreme gravitational redshift, where the apparent frequency of any radiation emitted from objects close to the event horizon is strongly reduced as the radiation escapes to infinity. This is a form of time dilation that does not depend on differences in the state of motion between two observers, like we saw earlier, but on the different positions that observers can have in a gravitational field. The strength of this phenomenon can be expressed when we consider the ratio of coordinate time to proper time for a stationary particle (i.e., $d\theta = d\phi = dr = 0$) sitting outside the event horizon as measured by a faraway observer. This quantity expresses the 'slowdown factor' of time close to the black hole. We use expression 1.4, setting $d\theta$, $d\phi$ and dr to zero, and re-arranging terms to get the relation between proper time and coordinate time for a stationary observer:

$$\frac{dt}{d\tau} = \frac{1}{\sqrt{1 - \frac{2GM}{c^2 r}}}. \quad (1.5)$$

This expression tells us that the ratio of elapsed coordinate time (or: time as measured by an observer at infinity, whose own $d\tau$ and dt are practically the same) to proper time for a stationary observer or particle close to the horizon diverges as we let the particle quasistatically approach the horizon. The gravitational time dilation does not simply become very large, it becomes infinite! Observers at large distances thus never observe anything crossing the event horizon as it falls in. Instead, matter will seem to approach the horizon ever more slowly and any radiation emitted

from it will be redshifted into obscurity. Locally, an observer falling in together with the accreting matter will observe no such slowing of time but will cross the event horizon (an unremarkable event for that observer) in finite proper time. This can be understood when radial freefalling motion is considered using different coordinates to describe the same Schwarzschild spacetime, such as Lemaître coordinates (Lemaître, 1933).

For several decades, these debates about black holes were mostly academic – they centered around properly understanding the theory and not so much around explaining the observed phenomena of the universe. The notion that black holes might exist as real, physical objects, and not merely as theoretical constructs, dawned in the 1960s with the discovery that quasars are the active nuclei of galaxies (Schmidt, 1963). In these sources, practically all emission comes from a concentrated region in the center of a galaxy, meaning that some kind of energetic process must be taking place there in order to liberate all that energy. Understanding the workings of such powerful ‘central engines’ necessitated finding a physical mechanism by which they could be powered – and the most promising physical mechanism for this was accretion onto a compact and massive object. In this same time period, Penrose (1965) brought black holes closer to reality from the theoretical side by showing that under suitable physical circumstances, collapsing matter should indeed form a black hole – even without the assumption of perfect spherical or even rotational symmetry.

1.2 The phenomenon of accretion

1.2.1 Astrophysical context

In the context of astrophysics, accretion is the process where diffuse matter (gas or plasma), moving under the influence of gravity, falls onto or into a central object (e.g., a planet, a star, or a black hole). Because the diffuse matter generally has nonzero net angular momentum, it tends to form a disk structure around the central object. Within this accretion disk, various physical processes can facilitate the transport of energy and angular momentum between different zones of the disk. By exchanging angular momentum with other regions of the accretion flow, matter can fall onto (or into) the central object. In this introduction, I will specifically focus on the process of accretion onto black holes as that is the relevant context for this thesis.

As the accreting material descends progressively deeper into the gravitational potential well of the black hole, its gravitational potential energy is converted into other forms of energy. One of the prominent forms is heat: by viscous dissipation of energy through some form of – likely magnetically mediated – friction, accreting gas in binary stellar systems can be heated up to temperatures of millions of Kelvins in the inner region of the accretion disk (Shakura & Sunyaev, 1973). Such high gas temperatures give the accretion flow a particular spectral signature, with the emitted frequency distribution depending on the specifics of the geometry and density of the accretion flow.

Black hole accretion works over a wide range of scales. On smaller scales, it can occur in close binary stellar systems where one of the components is a black hole. If the orbital separation between the two components is small enough, the black hole can draw matter from the outer layer of the secondary component (dubbed the 'donor') which then forms an accretion disk around it as it still possesses angular momentum (Figure 1.1). On large scales, accretion typically occurs from a more diffuse reservoir of gas such as the collective products of strong stellar winds in a galactic core. These diffuse gas clouds can accrete onto a supermassive black hole at the center of that galaxy if their relative velocity is sufficiently low to become gravitationally bound to the black hole.




FIGURE 1.1: Artist impression of a black hole binary. The companion or donor star, visible on the right, loses gas to the accreting black hole which is surrounded by an accretion disk. The black hole itself is too small to be seen at this scale. A small fraction of the accreting material may escape in the form of disk winds or, as pictured here, a high-velocity jet. Image credit: NASA/CXC/M. Weiss.

1.2.2 The Eddington luminosity

Black hole accretion has been studied extensively, both analytically from first principles and through numerical simulation. For an extensive review of black hole accretion (both in stellar

systems and in active galactic nuclei), see Abramowicz & Fragile (2013). In order to classify accreting systems, one important parameter to consider is the rate at which matter is supplied to the system (the accretion rate). When we want to define different behavioural regimes of accretion, it is useful to start by considering a simple situation in which diffuse and fully ionized hydrogen is accreting onto some compact object in a spherically symmetric way.

As this hydrogen accretes, we can imagine that the gravitational potential energy of the plasma is converted into radiative energy through some conversion process. As we increase the accretion rate, the opacity of the accreting plasma will increase as its density grows, and the relative influences of gravitational attraction and outward-pointing radiation pressure will shift. For simplicity, we can assume that most of the radiated power comes from directly around our central object. In such a scenario, we can ask ourselves the question: "How much power should our central object need to radiate for its radiation pressure to balance the attractive gravitational force it exerts on the surrounding plasma?". This power is called the Eddington luminosity (Eddington, 1920). Writing it down in the physically relevant terms, we get the following expression (Rybicki & Lightman, 1979):

$$L_{\text{Edd}} = \frac{3\pi G M m_p c}{\sigma_T}, \quad (1.6)$$

which is valid for neutral but fully ionized hydrogen. In setting up this expression, we make use of the fact that the gravitational influence is dominated by the mass of the protons (represented with m_p in the equation) as they represent most of the mass of the plasma, while the influence from radiative pressure is dominated by the Thomson cross-section of the electrons (represented by σ_T) as they dominate the interaction with electromagnetic radiation. G is the universal gravitational constant, M is the mass of the accreting object and c is the speed of light. When we assume a mass for the accreting object, the corresponding Eddington luminosity can be calculated. Note that the expression for the Eddington luminosity does not contain any term for the distance to the central object, as the outward-pointing acceleration from radiation pressure has the same inverse-square radial dependence as the inward-pointing gravitational acceleration. If the forces balance at one particular radius, they will also balance at any other radius when we assume that both the gravitationally dominant mass and the primary source of radiation are colocated in the center of the system.

1.2.3 The Eddington accretion rate

We can now picture a situation where the accretion rate is such that the rate of gravitational energy conversion reaches the Eddington luminosity. For isotropic, spherical accretion, we would expect the accretion flow to experience significant effects from radiation pressure at this point – likely changing the overall character of the accretion flow. If we want to calculate a critical accretion rate using the Eddington luminosity, we will need to make assumptions as to the radius where the gravitational potential energy of the accreting gas actually gets converted into radiative energy.

We will also need to assume some efficiency factor for this process. If we consider an object with a well-defined surface, such as a white dwarf or a neutron star, we can assume that the lower limit for the radius at which the potential energy of the infalling material gets liberated is the stellar radius. Let us consider gas that is moving in from some large radius R_{outer} to an inner radius R_{inner} , and look at the power available from the rate of change in gravitational potential energy of the accreting matter. Assuming a radially infalling motion for the accreting gas, we can set up the following expression for accretion power:

$$L_{\text{acc}} = \eta \dot{m} \left(\frac{-GM}{R_{\text{outer}}} + \frac{GM}{R_{\text{inner}}} \right). \quad (1.7)$$

For objects that have a well-defined surface, the efficiency parameter η can reach a significant fraction of unity as eventually all the gravitational potential energy of the material falling onto the surface will be converted into thermal energy and radiated away (Sibgatullin & Sunyaev, 2000), the only lost energy being carried away by the fraction of material that escapes from the accretion flow. However, for black holes the situation is somewhat more subtle. As black holes do not have an observable surface, there is no guarantee that the infalling material gets to convert its gravitational potential energy into radiation before it disappears at the event horizon. The hot gas may simply disappear behind the horizon before it has a chance to radiate its heat away. In this case, the efficiency factor in our expression will have a value much smaller than 1. The exact value of this efficiency factor depends on other physical parameters and specifics of the accretion flow under consideration. Equating the Eddington luminosity to the above expression for the accretion power, and assuming R_{out} to be much larger than R_{in} so that the corresponding term can be neglected, we get an expression for the Eddington accretion rate:

$$\dot{m}_{\text{Edd}} = \frac{3\pi m_p c R_{\text{in}}}{\eta \sigma_T}. \quad (1.8)$$

For a black hole, we can substitute the Schwarzschild radius $R_{\text{Sch}} = \frac{2GM}{c^2}$ for R_{in} , which then gives us:

$$\dot{m}_{\text{Edd}} = \frac{6\pi m_p GM}{\eta \sigma_T c} = L_{\text{Edd}} \cdot \frac{2}{\eta c^2}. \quad (1.9)$$

The factor of 2 typically gets absorbed into the accretion efficiency parameter η , so that the Eddington luminosity and the Eddington accretion rate are related by the expression $L_{\text{Edd}} = \eta \dot{m}_{\text{Edd}} c^2$. Theoretically, in the context of black hole accretion the value of the efficiency parameter η can range from ~ 0.06 for a non-rotating black hole to ~ 0.4 for a maximally rotating black hole (Novikov & Thorne, 1973). A fiducial value that is often picked for the efficiency parameter, based on observations of accreting sources (Soltan, 1982), is ~ 0.1 , which means that the mass-to-energy conversion of an accreting black hole can be appreciably more efficient than it is for stellar fusion (proton-proton) where it is ~ 0.007 . While the radiative efficiency of an accretion flow around a black hole depends on many specifics of the accretion regime and the environment, the Eddington accretion rate provides a useful general scale to use in the classification of accretion flows.

1.2.4 Models for accretion at different Eddington ratios

Apart from setting the energy budget for the accretion luminosity, the accretion rate directly influences the density of the accretion disk and thereby the optical depth of the system. The specific combination of radiative power and optical depth has consequences for the relative importance that radiation plays in the dynamics of the accreting system: if an accretion disk cannot cool efficiently because it has a large optical depth in combination with high accretion power, it will get geometrically thick because the trapped radiation heats the gas to a high temperature which brings the gas close to virial temperatures. This yields an accretion flow that behaves very differently from a system where the gas can cool efficiently through radiation, which keeps the accretion disk geometrically thin. At the other end of the scale is the regime of very low accretion rates, where the gas has such a low density that again it cannot cool efficiently through radiation – this time because of a lack of Coulomb interactions between its particles. The accreting gas thus reaches extremely high temperatures and we again get a geometrically thick accretion disk. To explore the range of possible configurations for an accreting system properly, we can set up a full system of 1D accretion equations, which describe the necessary relations for continuity of mass, energy conservation and angular momentum transport (Accretion power in astrophysics, section 11.4). The set of solutions to these 1D equations gives us a basic insight into the different possible physical configurations that an accreting system can have. When considering accretion rates, it is convenient to express the accretion rate as a fraction of the Eddington accretion rate. This fraction is often also called the 'Eddington ratio' of an accreting system. Several types of accretion flows have been identified over the years, valid for different ranges of Eddington ratios. These types are described in the following paragraphs.

Models for accretion disks based on physical principles were first proposed by Weizsäcker (1948), but left the mechanisms for angular momentum redistribution open. The earliest self-consistent model of an accretion flow was formulated by Shakura & Sunyaev (1973). This model describes a geometrically thin, optically thick accretion disk where the energy liberated through viscous dissipation is radiated locally everywhere. The spectrum radiated by this flat accretion disk thus looks like a sum of blackbodies, as each ring in the disk radiates at thermal equilibrium. This solution holds for intermediate Eddington ratios of ~ 0.1 to ~ 1 . The accretion flow depicted in the middle panel of Figure 1.2 falls in this regime.

Begelman (1978) presented a solution for radiation-pressure dominated, super-Eddington accretion flows, initially applied to black hole accretion inside massive stars. Later, this mode of accretion was suggested to be applicable to other accreting systems as well, specifically in AGN that undergo super-Eddington accretion (Kawaguchi, 2004; Ohsuga & Mineshige, 2007). This class of solutions is characterised by a hot accretion flow where radiation is trapped due to the large optical depth of the system, and where the local disk thickness is equal to the radius or even larger. The left panel of Figure 1.2 shows an example of this behaviour.




FIGURE 1.2: Gas density distributions in different accretion regimes, demonstrated by 2D axisymmetric radiation-magnetohydrodynamic simulations with different disk densities onto a stellar-mass black hole of $10 M_{\odot}$. Left: an initial disk density of 1 g/cm^3 gives a high Eddington ratio accretion flow, with radiation trapped in the optically thick 'slim' disk. Middle: an initial disk density of 10^{-4} g/cm^3 yields a geometrically thin, optically thick disk with an intermediate Eddington ratio. Right: an initial disk density of 10^{-8} g/cm^3 results in a geometrically thick, optically thin disk with an advection-dominated accretion flow at a low Eddington ratio. Streamlines for the accretion flow are indicated by the blue/green lines. Figure reproduced from Ohsuga & Mineshige (2011).

For very low-Eddington accretion flows ($\dot{M} < 0.005 \dot{M}_{\text{Edd}}$), which are expected to have low densities, an accretion model was proposed by Shapiro et al. (1976) (see also Rees et al. (1982) and Ichimaru (1977)) where the plasma develops a two-temperature state. In this regime the ions have high temperatures as they are unable to lose their energy efficiently (by either radiation or interaction with the electrons), whereas the electrons cool much more quickly through interactions with the magnetic field, and thus have significantly lower temperatures. Such a two-temperature accretion flow may form in the inner region of an accretion flow, and be connected to a larger Shakura-Sunyaev thin disk at its outer boundary. An example of this regime is visible in the right panel of Figure 1.2.

1.2.5 The $\dot{M}$ - Σ plane

The current understanding in accretion theory is that the Eddington ratio is not the only variable that determines the accretion flow type. These days, the unification picture makes use of classification of accretion flows in the so-called $\dot{M} - \Sigma$ (' $\dot{M}$ - Σ ') plane. Expressing the accretion rate as a fraction of the Eddington accretion rate along one axis and of the vertically integrated disk surface density denoted by Σ along the other axis, we can identify the types of solutions that can

hold for a given radius in an accreting system as a function of the effective viscosity parameter α , which is defined as in the alpha-prescription introduced by Shakura & Sunyaev (1973). The value of α is an important parameter when considering the regions in the $\dot{M} - \Sigma$ map: it parametrises the effective plasma viscosity in a dimensionless form in the following expression:

$$\nu = \alpha c_s H, \quad (1.10)$$

where ν is the plasma viscosity as used in the 1D accretion equations, c_s is the local sound speed in the plasma and H is the local scale height of the accretion disk. This functional relation was motivated by a consideration of turbulence as the primary source for viscosity: the mixing rate scales with the sound speed, and eddies in the flow are limited in size by the local scale height of the accretion disk. Depending on the value that is chosen for α , different families of solutions for the accretion equations can be found that each show a specific relation between the accretion rate and the disk surface density.

The map of solutions is shown in a simplified form in Figure 1.3. An important point is that such a map is only valid for one particular radial position in an accreting system, and that changing the radial position under consideration will change the appearance of the map. Furthermore, the equations used to identify these solution types assume that we are dealing with a single-temperature accretion flow, where the electrons and protons share the same temperature. This assumption may be violated in cases of very low plasma densities.

The $\dot{M}$ - Σ map shows multiple regions that may share the same values for α but still exhibit different behaviour. α_{crit} is the value for the viscosity parameter that separates the different families of solutions. The zones indicated in this diagram, separated by dashed, dotted and dash-dotted lines, indicate where the influences of different physical mechanisms dominate. Regions where the disk is optically thin/thick can be identified, as well as regions where gas pressure or radiation pressure dominates. Furthermore, the dominant term of heat loss for the local plasma can be identified as being either radiation or advection (i.e., transportation to smaller radii). Note that the value for α_{crit} can change as a different radius is considered - and so, the character of an accretion disk can in principle be quite different in different radial zones. The solution families indicated in the diagram with symbols I-IV correspond to types of accretion flows, and a short overview of them is given here.

Solutions of type I are optically thick and can be either gas-pressure dominated or radiation-pressure dominated. The high-Eddington ratio segment of this class yields radiatively inefficient accretion flows with puffed-up disks that are supported by radiation pressure, also known as 'slim disks' which have a local thickness that is similar to the local radius. In this regime, the generated radiation gets trapped in the disk because of the high optical depth of the accretion flow. The energy in the radiation field and the accreting matter predominantly gets transported inward rather than radiated away, classifying this type as a so-called 'advection dominated accretion




FIGURE 1.3: The general structure of the $\dot{M}$ - Σ plane, taken from Frank et al. (2002). Branch I covers Shakura-Sunyaev flows (bottom) and slim disks (top). Branch II covers Shakura-Sunyaev flows (right) and Shapiro-Lightman-Eardley (SLE) flows (left). Branch III covers SLE flows (right) and advection-dominated accretion flows (ADAFs, left). Finally, branch IV covers the 'Polish donut' class (Abramowicz et al., 1978). The vertical dashed line indicates the boundary between optically thin flows (to the left) and optically thick flows (to the right).

flow' (ADAF). See Mineshige & Ohsuga (2007) for a review of this regime as it occurs for AGN. For lower Eddington ratios, the type I solutions enter an unstable regime (downwards-sloping lines in the diagram). The lower segment of type I, which is once again stable, yields 'classic' Shakura-Sunyaev accretion disks that are radiatively efficient, geometrically thin, optically thick and predominantly supported by gas pressure. See Blaes (2007) for a review of this intermediate Eddington ratio accretion flow type in AGN. The solutions that lie along this branch in zone I are classified as 'cold' accretion flows, in contrast to the configurations found in the left segment of zone II and in zone III which are classified as 'hot'.

Type II and III solutions take the accretion flow into a different regime. At low Eddington ratios and low surface densities, cooling the ions in the accretion flow becomes problematic. The wedge

defined by the intersection of the optically thin region with the region where radiation-cooling is the dominant heat loss mechanism defines a parameter space where the single-temperature assumption may break down for regions within the accretion flow. The reason for this is because of the low disk surface density in this part of the parameter space, Coulomb interactions between the plasma particles will be rare and the accreting gas will not be able to radiate away its internal energy efficiently: most of the kinetic energy is locked into the ions because of their higher mass, but this energy is not shared efficiently with the electrons that would be able to radiate this energy efficiently. Thermal energy thus gets advected inwards along with the accreting matter. The electrons in such an accretion flow may cool through processes that do not require close particle interactions, such as synchrotron radiation, but the electrons are themselves not heated efficiently by the ions in the accretion flow. This recipe for a two-temperature accretion flow was first formulated by Shapiro et al. (1976), which describes an accreting system where the two-temperature model features in the inner region (the boundary of zone II and zone III in the $\dot{M}$ - σ diagram corresponds to mixed accretion flows of this type). Yuan (2007) presents a review of this type of low-Eddington accretion flows in AGN.

Zone IV describes accretion flows that are commonly called 'Polish donuts' (Abramowicz et al., 1978). These are radiation-supported, thick accretion disks that are almost spherical in nature but do have a clearly defined rotation axis. Their accretion efficiency (the fraction of potential energy converted into radiation) is very low, and as such they can exhibit accretion far above the Eddington limit. Cases with $\alpha > 1$ (in zone IV) are generally considered to be unphysical, but there may be circumstances where that condition can (temporarily) be reached (Frank et al., 2002, paragraph 7.6).

1.2.6 Disk outflows and jets

As the accreting material within a disk is lowered into closer orbits, it must lose angular momentum – and this extracted angular momentum must go somewhere. An obvious channel for extraction of angular momentum from an accretion disk comes in the form of matter outflows. From observations of protostars, microquasars and AGN we know that such accreting systems, across a wide range of scales, often show fast outflows in the form of collimated winds or of plasma jets. These jets are thought to be intimately linked to magnetic processes occurring in the accretion flow (see Pudritz et al. (2012) for an overview of source types and general discussion). In the context of black hole accretion, the two fundamental jet launching models that have had the largest impact are the Blandford-Znajek model (Blandford & Znajek, 1977) and the Blandford-Payne model (Blandford & Payne, 1982).

The Blandford-Znajek model describes a process whereby a magnetised accretion flow pushes magnetic flux inward, closer to a black hole with nonzero spin, amplifying the magnetic field at small radii. Angular momentum and energy are then extracted from the rotating black hole

through the forced corotation (frame dragging) of the magnetic field threading the ergosphere, which winds up the field into a helical configuration and generates a Poynting flux away from the system, along the spin axis of the black hole. This process has been associated with the launching of highly relativistic jets. A schematic illustration of this process is shown in Figure 1.5.

The Blandford-Payne model (Blandford & Payne, 1982) is a separate jet-launching mechanism that was proposed a few years later. In this model, plasma is launched from the accretion disk boundary layer by being 'flung out' along magnetic field lines anchored in the disk, where the field orientation in the poloidal plane exceeds a critical angle outwards from the spin axis of the accretion disk. The jet launching region for this model is much larger than in the Blandford-Znajek case, and initial plasma launching speeds are lower. In reality, jet launching in AGN is expected to feature some combination of these two jet launching mechanisms.

Irrespective of the jet launching mechanisms that are assumed to be operational, Blandford & Königl (1979) investigated the expected general properties of a conically expanding jet launched from an accretion flow. They operated on the assumption that such a jet is close to isothermal along its axis – i.e., the temperature of the particles in the jet does not drop with increasing distance from the black hole, whereas the magnetic energy density and the particle number density do. This property then yields a flat radio spectrum for such a jet. Observationally, AGN have indeed been associated with flat core radio spectra. The cores of AGN are therefore thought to harbour unresolved expanding and accelerating jets. A fundamental feature of an expanding jet is the varying position along the jet axis of the apparent core of emission, when we look at different radio frequencies. Due to the fact that the jet has a non-constant density along its axis, the location of the photosphere (the point where the optical depth is unity) depends on the frequency of the emission under consideration. This situation provides several powerful ways to test the theory: we can look for core shifts in jets, where the centroid of emission shifts along the jet axis with observing frequency (a prime example being the jet in M87, see Hada et al. (2011)), and we can look for variability time lags in jets as variations in the emissivity of the outflowing plasma typically become visible sooner at higher observing frequencies (used to generate predictions for Sgr A* in Falcke et al. (2009)).

The notion that plasma jets are produced from some accretion flows invites the question: under what circumstances do jets form, and how do their properties depend on the state of the accretion flow? Falcke & Biermann (1995) coupled the Blandford-Königl jet model to the overall mass and energy budget of an accreting system and showed that the theoretically expected emission matches observed flux densities over a wide range of frequencies (radio to X-ray) and accretion rates. The relation between accretion flows and the jets they sometimes produce remains quite an active field of investigation, see e.g. Sbarrato et al. (2014).

Relativistic plasma jets themselves are also studied intensively, as they can have a significant

impact on the environment of the accreting system over a wide range of scales. Their large velocities compared to the surrounding material, combined with the presence of strong magnetic fields, provide conditions that are highly conducive to the acceleration of electrons and cosmic rays to high energies. Particle acceleration is thought to be possible close to the jet launching region, where the accelerated electrons are thought to be responsible for e.g. high-energy gamma-ray emission that has been observed from accreting systems. At much larger scales, the radio lobes of AGNs are formed around the points where the plasma jet from the central engine impacts the intergalactic medium and causes turbulent plasma structures to form that show a slowly cooling population of fast electrons emitting in radio. These radio lobes are also considered to be strong candidate sites for the generation of high-energy cosmic rays. Furthermore, the deposition of energy into the circumgalactic medium by relativistic plasma jets is thought to lower the rate at which the galaxy accretes gas from the intergalactic medium. In this way, the activity of an accreting supermassive black hole in the center of a galaxy can indirectly impact the star formation rate in that same galaxy, and in this way its long-term evolution. Open questions in the research on plasma jets include the nature of their acceleration profiles, their collimation mechanisms, their transverse structure, their specific sites of particle acceleration and their specific dependence on black hole spin parameters.

1.3 Observational signatures of black hole accretion

In 1971, the Uhuru X-ray satellite observed a region of the sky in the Cygnus constellation. In this region, an X-ray source associated with a blue supergiant had previously been discovered. Uhuru measured strong and rapidly varying X-ray emission that seemingly came from the location of a blue supergiant star on the sky (Oda et al., 1971). Such a high X-ray flux density cannot be generated by a supergiant star, so an alternative explanation for the X-ray source needed to be found. The rapid time variability of the X-ray source, with temporal features as fast as one millisecond (Rothschild et al., 1974), suggested it was an accreting object. Its mass ($\sim 15 M_{\odot}$, Orosz et al. (2011)) could be derived from the radial motion variations of the blue supergiant and the periodicity of emission from the system, and it was concluded that Cygnus X-1 had to be a black hole as it was found to be too massive to be a neutron star. This made Cygnus X-1 the most strongly observationally supported black hole candidate at that point.

The observational signatures of accreting systems vary considerably according to their scale, inclination, obscuration and accretion rate. Accretion flows in X-ray binary star systems can cycle through different accretion states over relatively short timescales (Tananbaum et al., 1972), and in that way they provide excellent laboratories for studying accretion dynamics. Their spectral signature in X-rays can be plotted in a so-called hardness-intensity diagram (HID), where the spectral slope of the X-ray emission is plotted against the X-ray luminosity (see Figure 1.6). The evolution of these systems shows up as a typical q-shaped curve, with a quiescent ('low/hard') state in the bottom right and a flaring ('high/soft') state at the top left. According to our current




FIGURE 1.4: Illustration of the Blandford-Znajek mechanism, where energy and angular momentum are extracted from the ergosphere of a rotating black hole via magnetic interaction. Initially straight magnetic field lines, anchored in the orbiting plasma, get twisted up and cause part of the plasma in the ergosphere to attain a negative total energy – where its binding energy is larger than its rest-mass energy (red-coloured segment). The extracted energy gets transported out of the equatorial plane along a jet structure, where plasma is accelerated by magnetic pressure from coiled field lines. Through this mechanism, energy is extracted from the black hole itself rather than from the accreting plasma. Figure from Semenov et al. (2004).

understanding, these emission states are the consequence of very different accretion disk configurations. The low/hard state has a low-density disk or corona and a plasma jet associated with it, while the high/soft state has a flatter, denser disk but not necessarily a jet. The changes between accretion states are thought to be a function of the varying supply rate of gas to the accretion disk.

Compelling evidence for stellar-mass black holes in binary systems comes from gravitational wave observations (Abbott et al., 2016). The merger events observed by this highly sensitive network of gravitational wave detectors can only be successfully modeled by invoking binary black holes, as no other proposed type of object is compact enough to yield gravitational waveforms of the observed strength and frequency. To date, LIGO/VIRGO has reported detections




FIGURE 1.5: Geometry of the Blandford-Payne mechanism. Packets of plasma are strongly coupled to the magnetic field, and are constrained in their motion to move along field lines. Close to the disk midplane, if the outwards inclination of the (corotating) magnetic field lines away from the vertical direction is 30° or more it is energetically favourable for the packet to 'slide' to larger radii and gain kinetic energy. In this way, an outflow can be launched from close to the disk boundary layer. Figure from Jafari (2019).

of black hole mergers in the mass range from a few Solar masses up to several tens of Solar masses.




FIGURE 1.6: Left: state changes in an observed X-ray binary system, visible in a hardness-intensity diagram (HID). Right: schematic depiction of the states and their transitions. Figures from Belloni et al. (2005) and Fender et al. (2004).

Observational evidence for another type of black hole came from observations of active galactic nuclei (AGN). The systematic study of what later proved to be AGN effectively started in 1943,

when Carl Seyfert published a catalog of galaxies of which the nuclei exhibited broad spectral line emission (Seyfert, 1943). Research into these sources entered a phase of rapid development with radio observations done in the 1950s and 1960s. With the resolving power for radio observations steadily improving and more spectral studies being performed, it became clear that these objects showed highly redshifted spectral lines and were very compact. One of the explanations offered was that these were Galactic objects of very high density with a large gravitational redshift. However, other aspects of the spectral observations (specifically the presence of forbidden lines in their optical spectra, necessitating the presence of low-density gas in the source) pointed towards them having an extragalactic character, with the redshift having a cosmological nature. The relatively fast optical variability of these sources (on timescales of $\sim$ days) suggested that they are compact in nature (Smith & Hoffleit, 1963). It thus became clear that these sources were situated at cosmological distances from the Milky Way (Schmidt, 1963), and exhibited tremendous luminosities. Apparently, the nuclei of many faraway galaxies were harbouring powerful engines that show up brightly in radio and sometimes in optical as well. Given their rate of energy conversion, these central engines were also found to be likely responsible for powering large 'radio lobes' beyond the limits of the host galaxies themselves, through the expulsion of fast plasma jets. Salpeter (1964) and Zel'dovich (1964) proposed that this engine is powered by accretion of gas onto a supermassive black hole.

Further observational evidence for the existence of supermassive black holes comes from the measurement of spatially resolved lensed synchrotron emission from the orbiting plasma directly around the supermassive black hole (Event Horizon Telescope Collaboration et al., 2019a) at the center of the M87 elliptical galaxy. The observed emission distribution is consistent with a gravitationally lensed plasma flow around a dark object. This is the first observation of a black hole candidate that provides direct support for the existence of an event horizon, as opposed to a massive object with an observable surface.

Nowadays, the body of evidence supporting the existence of black holes, both stellar-mass and supermassive ones, is extensive. Besides the X-ray and AGN studies mentioned, observations of stellar velocity distributions in several nearby galaxies have shown the presence of dark and concentrated masses (Kormendy & Richstone, 1995), while more recently analysis of individual stellar dynamics around the Galactic center points to the existence of a highly compact but practically invisible mass of $\sim 4 \times 10^6 M_{\odot}$ at the center of the Milky Way (Ghez et al., 2008; Genzel et al., 2010). Observations done in infrared over the course of multiple decades show a central star cluster where the individual stars move on Keplerian orbits around the same locus. The mass concentrated in this point apparently completely dominates the orbital dynamics of these stellar orbits, but the associated object is generally faint in infrared outside of the short time intervals when it flares up and gets up to tens of times brighter (Dodds-Eden et al. (2011)).

In AGN, the scale of the accreting system is several orders of magnitude larger than in X-ray

binaries. All associated timescales are correspondingly longer. We may never observe an AGN changing accretion states as its supply rate of gas is only expected to vary significantly over timescales of hundreds of thousands or even millions of years. We can however see the past activity of some AGNs by studying the current appearance of their jets and hot spots. It appears that the jets of AGN have limited duty cycles when considered over longer timescales (Schawinski et al., 2015), and thus AGN seem to experience different states much like X-ray binaries do. AGN are indeed classified into radio-quiet and radio-loud categories, where the latter category has a prominent jet that is responsible for the radio emission. However, the reason for this dichotomy is not clearly understood yet and multiple causes (among which are accretion rate, BH spin, BH mass, environment) have been put forward as potential explanations (Retana-Montenegro & Röttgering, 2017, and references therein).

1.4 History of Sagittarius A*

It seems that our Milky Way, like many other galaxies, harbours a supermassive black hole at its center. In this subsection I will describe where this notion comes from and what support for it we currently have. The center of our Milky Way is heavily obscured to us at optical wavelengths, with gas and dust in the plane of the Galaxy, along our line of sight towards the Galactic center, providing practically complete extinction of all optical radiation coming from there to us. This obscuration is less severe at longer wavelengths, and from infrared (IR) down to radio we can successfully probe the nuclear environment of our Galaxy. As such, the history of observations that led to our current understanding of what our Galactic center harbours is quite extensive.

1.4.1 Early detections and interpretation

In 1931, as part of his investigation into the noise on transatlantic transmissions, Karl Jansky made the first detection of radio emission coming from the Galactic Center (Jansky, 1933). General interest in active galactic nuclei prompted further study into the center of our Milky Way with infrared detections reported by Becklin & Neugebauer (1968). Measurements from the early 70s showed that multiple radio structures could be identified around the Galactic center, one of which appeared to be compact (Downes & Martin, 1971). Balick & Brown (1974) reported the detection of a bright compact radio component at the Galactic center with an angular size smaller than 1 arcsecond, which was later dubbed Sagittarius A* (For a more extensive history of the early observations of Sgr A*, see Goss et al. (2003)). Coupled with velocity measurements of circumnuclear gas clouds done in infrared (Wollman et al., 1977), it became evident that the Galactic center harbours a compact and massive object that dominates the local motion of gas and stars, having a mass of approximately 4 million Solar masses. Stellar motion studies, where the motions of the central nuclear starcluster were tracked over multiple years using infrared

imaging with adaptive optics, further strengthened this picture (Ghez et al., 2008; Genzel et al., 2010). Reid & Brunthaler (2004), using VLBI, found that the proper motion of Sagittarius A* with respect to the Galactic center is zero within the measurement uncertainties. This result showed that the object associated with the compact radio emission accounts for a significant fraction of the 4 million solar masses present in the Galactic center from earlier studies. The most stringent constraints currently on mass and distance to Sgr A* that are currently available have been derived from IR interferometric measurements using VLTI with the GRAVITY instrument. Analysis of the close passage around Sgr A* of the S2 star, observed with GRAVITY, shows that the trajectory of the star can best be described by relativistic geodesic motion in a Schwarzschild metric (Gravity Collaboration et al., 2018a, 2020a). The best-fitting mass for Sgr A* from the fit of the stellar orbit in the latter publication is $4.262 \pm 0.012_{\text{stat}} \pm 0.06_{\text{sys}} \times 10^6 M_{\odot}$, together with a distance of $8249 \pm 9_{\text{stat}} \pm 45_{\text{sys}}$ pc. This means that the expected angular size of the 'shadow' (the lensed projection of the photon sphere on the sky) of Sgr A* is the largest for any SMBH, close to $50 \mu\text{as}$, and this brings it within reach for high-frequency VLBI to resolve.

The results cited above provide evidence for the presence of a compact massive object, likely a black hole, surrounded by an accretion flow at the Galactic center. The specific nature of the processes occurring there have been addressed by a host of other observations and studies. These have provided insight into the local plasma conditions, magnetic field structure and flow dynamics in Sgr A*, and are discussed in the following paragraphs.

1.4.2 The spectrum of Sgr A*

The radio and submm spectrum of Sgr A* (see Figure 1.7) shows a bump in the submm range, which is indicative of synchrotron emission where high-temperature electrons are moving through magnetic fields within a limited volume, comparable to the scale of the event horizon (Falcke et al., 1998, 2009; Melia & Falcke, 2001) – suggesting that this emission comes from very close to the black hole. The low-frequency end of the radio spectrum is relatively flat (the spectral index is ~ 0.28 between 1.4 and 22 GHz), which suggests that we see a photosphere that is located at larger radii for lower frequencies. The infrared side of the synchrotron bump, above ~ 1 THz, shows a steep drop of flux density with frequency (the spectral index is ~ -1.7 , Gillessen et al. (2006)) which can be shallower when Sgr A* is in a flaring state (Witzel et al., 2018).

The radio emission from Sgr A* has been relatively stable over multiple decades, with short-term variability (up to 10% in flux density in radio, climbing to 50% in the submm) that keeps returning to an apparent long-term equilibrium (Falcke, 1999; Macquart & Bower, 2006; Dexter et al., 2014). In the near-IR, the quiescent flux density has not been measured with confidence so far because the source is not always detectable. However, the variability in NIR is much more pronounced than in radio and submm, with the flux density at $2.2 \mu\text{m}$ varying around 0.017 mJy in quiescence and up to ~ 8 mJy during periods of increased activity. Despite the relatively rare




FIGURE 1.7: Collected measurements and a fitted multiwavelength spectrum of Sagittarius A*, from Markoff et al. (2007). The flat spectrum at radio frequencies is evident, as is the submm bump and the steeper drop in flux density at infrared frequencies and above. This figure incorporates measurements from multiple observations, which are specified in the referenced publication.

nature of these bright episodes, the source cannot be said to flare as the flux density appears to show continuous variability according to a power law (Fazio et al., 2018). In X-rays, the spectrum of Sgr A* shows a second bump that is commonly interpreted as a Compton-upscattered copy of the synchrotron bump. Regarding flux density variability, Sgr A* shows a weak quiescent emission interspersed by much brighter and relatively short flares (~ 1.1 per day, typically lasting for less than an hour).

Sgr A* appears to be an extremely low Eddington ratio source. From parallels with stellar-mass accreting black holes, it may therefore be expected to harbour a (weak) jet (Falcke et al., 1993). However, given the expected weakness of such a jet it is very challenging to verify its presence by measuring the source morphology (Markoff et al., 2007).

1.4.3 Spectral models for Sgr A*

A general fit to the observed spectrum of Sgr A* from higher radio frequencies to IR is provided by an analytical self-similar model for low-Eddington accretion, dubbed an advection-dominated accretion flow (ADAF), in which relativistically hot and magnetised low-density gas accretes onto a black hole with low radiative efficiency (Narayan et al., 1995). However, this model does

not naturally generate sufficient emission at lower radio frequencies to reproduce Sgr A*'s flat spectrum there. This extra radio emission can for instance be generated through the inclusion of a nonthermal distribution of electrons (Mahadevan, 1998; Özel et al., 2000; Yuan et al., 2003), but initial models required relatively high accretion rates that were incompatible with the observed rotation measure from the inner accretion flow (Agol, 2000; Quataert & Gruzinov, 2000; Bower et al., 2003; Marrone et al., 2007). It should be mentioned here that while this original model was termed the 'ADAF model', the general character of a low-Eddington accretion flow is expected to be advection-dominated in any case - so, strictly speaking all models discussed here are ADAF models. The variables with the strongest influence on the observed spectrum in this context are the topology and strength of the magnetic field and the electron energy distribution as a function of location in the accretion flow.

A different class of model that was put forward is the compact jet model, which provides a scenario where accelerated electrons in the jet region generate the required low-frequency radio emission (Falcke et al., 2000). The rationale for this model follows from the ubiquitous presence of jets in low-Eddington AGN, with Sgr A* being a low-Eddington system as well (Falcke et al., 1993). In the jet model, the submm-bump emission is primarily generated in the 'jet nozzle' (or: 'launching') region, which is the region from where the jet starts its acceleration out to larger radii. This jet nozzle is situated at only a few Schwarzschild radii from the black hole, comparable to the radii within which the submm bump emission is primarily generated in other models. Expansion of the plasma in the accelerating jet, chiefly along its direction of motion, means that the density, temperature and magnetic energy density inside it evolve with distance according to simple relations which also dictate the resulting jet radiation emission profile. In this model, the emission at lower frequencies is naturally generated from the jet region and fits the observed radio spectrum quite well along with the submm bump. Technically, the jet model can also be called an ADAF but differs from it in terms of where the emission responsible for the submm bump is chiefly generated (jet / jet nozzle versus inner accretion flow), not in terms of the density or temperature of the inner accretion flow. Both classes of models are considered candidates for Sgr A* as the details of electron acceleration in the inner accretion flow and jet region are not yet fully understood.

More recently, motivated by results from general-relativistic magnetohydrodynamic (GRMHD) simulations, another dimension of the parameter space has become a prominent subject of study: the topology of the magnetic field in the inner accretion flow. Depending on the magnetic properties of the accreting plasma flow, there are different scenarios possible for the buildup of magnetic field close to the black hole. In one such scenario, dubbed 'standard and normal evolution' (SANE), magnetic flux does not keep building up over time directly outside the event horizon, and the magnetic field plays no important dynamical role except to facilitate the radial transport of angular momentum. The magnetic field strength close to the event horizon remains modest and does not present a barrier to accreting material (Narayan et al., 2012). The alternative

scenario is one where a strong buildup of net magnetic flux close to the horizon does occur which leads to a 'magnetically arrested disk' (MAD, Igumenshchev et al. (2003); Narayan et al. (2012)), where the strong magnetic field close the horizon provides radial support for the sub-Keplerian orbiting material there. Between these models there are no significant differences in terms of predicted outflow and accretion rates, but their different configurations of plasma and magnetic field close to the black hole have consequences for the expected emission coming from there. The MAD/SANE paradigm has for instance featured prominently in the theoretical analysis of M87* data from the Event Horizon Telescope, which showed that the observed source morphology was still compatible with several models from both classes, with MAD models fitting over a wider range of parameters than SANE models did (Event Horizon Telescope Collaboration et al., 2019b).

The magnetic aspect of the parameter space links us to yet another dimension of the uncertainty surrounding low-density accretion flows in simulations: the role of particle (mainly electron) acceleration. As GRMHD simulations cover the behaviour of the bulk flow (i.e., the ions) rather than the dynamics of the electrons, the relation describing the energy distribution of the electrons in different parts of the accretion flow is a free parameter.

Jets appear naturally in these simulations, but depending on the particle acceleration recipe the emission can appear either disk- or jet-dominated. Moving beyond a purely thermal distribution for the electron energies, a prescription where the electron energies follow a thermal distribution combined with an accelerated power-law (the κ -distribution) shows an improved match with the radio-to-infrared spectrum for Sgr A* (Davelaar et al., 2018). However, it should be stressed that choosing the prescription for the electron energy distribution is at present still a separate step which, although they can be physically motivated, allows for multiple different options. Hence, any observational data on the morphology and the behaviour of the emitting plasma (inflow or outflow) helps to provide strong constraints on this theoretical component.

1.4.4 Morphology

The environment of Sgr A*, considered on scales of arcseconds on the sky, shows complex structures when viewed in radio and X-ray (Figure 1.8). Multiple structures can be recognized within the extended radio structure closest to Sgr A*, called Sgr A West (the so-called 'minispiral'), which have been interpreted from spectroscopic observations as streams of gas following the influence of the gravitational potential of Sgr A*. The region in and around the minispiral has been extensively searched for features that might have resulted from past Sgr A* activity.

One such feature was identified in X-ray data, showing up as a linear enhancement in X-ray maps of Sgr A West (Li et al., 2013). The linear feature is oriented along a line pointing towards Sgr A*, and appears to be aligned with a shock structure in the radio map of the East-West arm in the minispiral. If it is indeed linked to Sgr A*, this feature suggests a period of enhanced activity




FIGURE 1.8: The complex structure of the environment of Sagittarius A*. Left image: the gas complex named Sgr A West, imaged in false colours in different spectral ranges: X-rays in purple, IR in gold, and radio in orange/red. In the right part of the image, the 'minispiral' surrounding Sgr A* is visible in radio. Right image: A closer-up view of the center of the minispiral, imaged using the VLA at 5 GHz. Sgr A* is visible as the bright dot in the center. Left image credit: VLA; HST; Spitzer; CXC. A. Angelich (NRAO/AUI/NSF); NASA/JPL-Caltech/ESA/CXC/STScI. Right image credit: NRAO/AUI/NSF, J-H Zhao, W.M. Goss.

for Sgr A* in the relatively recent past. X-ray emissions from other gas complexes in the general vicinity of Sgr A* have been interpreted as being light echoes from Sgr A* that correspond to a state of higher activity in the past 500 years (Ryu et al., 2013).

When we look at the compact structure of Sgr A* in radio at low frequencies (below approximately 43 GHz), we see that Sgr A* appears as a Gaussian with a size that scales with the inverse square of the observing frequency (see Figures 1.9 and 1.10). This effect is attributed to the presence of a scattering screen located ~ 3 kpc away along our line of sight towards the Galactic center (Bower et al., 2014). At 22 GHz, the scattering screen shows substructure that makes the source appear mostly Gaussian but with small-amplitude fine structure superimposed (Gwinn et al., 2014). Moving to higher frequencies, the observed source size shrinks but begins to deviate from the lambda-squared scattering law: the source appears larger than expected if using only the scattering relation to predict the source size (Bower et al., 2004; Doeleman et al., 2008). This suggests that the intrinsic size of the source, that is the size it would have on the sky without the scattering screen being present, starts to factor into our size measurements at those frequencies. Disentangling the intrinsic size-frequency relation from the influence of the scattering screen has been a focus of recent investigations, and analysis done so far indicates that the intrinsic size of Sgr A* scales with $\lambda^{1.3}$ (Bower, 2006; Falcke et al., 2009). The morphology of Sgr A* at 86 GHz

has been shown to be non-Gaussian and asymmetric (Ortiz-León et al., 2016; Issaoun et al., 2019, Chapter 3 in this thesis), although whether this is due to intrinsic structure or the influence of scattering is not yet established. At the finest angular scales, the Event Horizon Telescope (EHT) has observed Sgr A* with a resolution of $\sim 25 \mu\text{as}$, the results of which will be presented in a forthcoming paper.




FIGURE 1.9: The appearance of Sgr A* at wavelengths from 7 mm to 6 cm, as observed using the VLBA. Note the different angular scales below each plot. Figure adapted from Lo et al. (1999).




FIGURE 1.10: The observed size of Sagittarius A* across a range of observing wavelengths. Plotted is the major axis size divided by wavelength squared, as a function of observing wavelength. For the longer wavelengths, the lambda-squared relation holds. Moving to shorter wavelengths, the measured source size begins to deviate from this relation. Figure taken from Bower et al. (2006).

1.4.5 Polarisation

Besides source morphology (studied in radio using VLBI) and variability (studied in radio, IR and X-ray using various telescopes), the polarisation of the radiation coming from Sgr A* carries its own information with it and presents opportunities for us to learn about the nature of the source.

In the radio spectrum, Sgr A* shows a source-integrated fractional linear polarisation that ranges from less than 0.1 percent at 4.8 GHz and less than 1 percent at 86 GHz to ~ 12 percent at 150 GHz and ~ 22 percent at 400 GHz (Bower et al., 1999a,b; Aitken et al., 2000). This trend can be understood in terms of the influence of optical depth and scattering on the polarisation properties of synchrotron radiation, where we see a smaller synchrotron emission region with a lower optical depth at higher frequencies. Sgr A* also shows circular polarisation at lower radio frequencies (see Figure 1.11), with a trend that does not increase as steeply with frequency as the trend for linear polarisation does (Bower et al., 2002; Muñoz et al., 2012, and references therein).




FIGURE 1.11: The measured linear and circular polarisation fractions for Sgr A* as a function of frequency. Figure and sources from Muñoz et al. (2012).

EHT observations performed in 2013 show a trend of larger fractional linear polarisations on smaller angular scales (Johnson et al., 2015). These measurements suggest that the magnetic field at those angular scales is highly ordered, which lends support to a model where the inner accretion flow is threaded by a strong and aligned magnetic field (Gold et al., 2017).

Time-variable polarisation has been detected in GRAVITY observations, where the polarisation angle on the sky appears to follow the circular motion of the emission centroid (Gravity Collab-

oration et al., 2018b). This result is interpreted as arising from orbital motion of the emission region in a strong poloidal magnetic field, an interpretation that also suggests that we see the footpoint of a jet from Sgr A*, close to face-on ($i = 160 \pm 10^\circ$).

1.5 Measurement of radio waves, radio interferometry and VLBI

The material presented in this thesis discusses results that have been obtained through interferometric measurements in the radio spectrum, using the two techniques of connected-element interferometry and Very Long Baseline Interferometry (VLBI). In this subsection, the fundamental concepts pertaining to this technique are discussed at a basic level.

1.5.1 Flux density in radio

The spectral flux density of radio waves is normally expressed in units called Janskys, after Karl Jansky who first identified Galactic radio emission in 1933 (Jansky, 1933). One Jansky is equivalent to 10^{-23} erg/s/Hz/cm². Thus, when multiplied by an antenna effective area and a frequency bandwidth in the appropriate units, it gives us received power. Expressed in SI units, one Jansky is equivalent to 10^{-26} W/Hz/m². The brightest radio sources at frequencies above 1 GHz have flux densities of around 10 Janskys, but weaker sources down to micro-Janskys can be detected if appropriate collecting areas, spectral bandwidths and integration times are used.

1.5.2 Signal amplification and filtering

As the flux density from astrophysical sources in radio is typically very low once their radiation reaches Earth, the first step in the processing chain is to amplify the incoming radio waves. For most radio telescopes operating at frequencies above 1 GHz, initial amplification is done by focusing the radio waves into a small area with a parabolic reflector. In that focal point, the radio waves are electrically sampled using one or more feeds (technically, these are the actual antennas), which transforms the signal into an oscillating voltage on an electrical line. This signal is then amplified and filtered in several stages before being digitally sampled and sent on to the correlator (for connected-element interferometry) or locally stored on digital media (as is the case for VLBI). For higher frequencies, parts of this receiver chain can be cooled in order to minimise the addition of noise to the signal.




FIGURE 1.12: Schematic view of the components of a typical heterodyne receiver system. The sky signal comes in at the top of the figure, and is focused by an antenna. The signal is then pre-amplified, and mixed with a stable frequency provided by the local oscillator. This mixed signal is then low-pass filtered and amplified further before being sampled in the digitizer stage. Depending on the specific system, if the IF bandwidth is very large it may also be channelized into multiple separate limited-bandwidth signals before being digitized. Once digitized, in a connected-element interferometer the signal has its delay corrected and is correlated in real time with the signals from the other antennas. In a VLBI system, the digitized signal is recorded onto storage media and is later processed in a correlator facility (which is typically at a different location). Figure taken from Thompson et al. (2017).

1.5.3 Radio interferometry

Let's think about how we can get high angular resolutions when observing in radio. Radio wavelengths are very long compared to those of visible light, and so you might get the impression that there is an inherent limit to the angular resolution we can attain with them. Fortunately, this is not the case – we can reach extremely fine angular scales with our radio telescopes, and it follows

from the way we can combine the signals that they receive. To explain how this is possible, let us consider a simple model for a radio source we might observe.

When we imagine observing a spatially incoherent astrophysical radio source (i.e., an object or system of which the different sub-regions emit radiation in an independent and unsynchronised way) at a certain observing frequency, the wavefronts arriving at us from the different regions within that source will exhibit an interference pattern, giving us a summed waveform that will vary with our position. The further away the source is located, the smaller the solid angle it subtends and the more closely the wavefronts from its different sub-regions will match each other in their propagation direction. This means that for a faraway source, subtending a small angle on the sky, the interference pattern associated with this source will only start to become apparent over larger spatial scales (see Figure 1.13 for a simplified depiction of this relation). If we are positioned closer to the source, the spatial variations in the collective waveform will be more pronounced over a given distance as the arrival directions of the components of the wavefront differ more from each other.

The central thought behind interferometry is that we can gain information about the source structure by measuring the spatial structure of this interference pattern. This is done through the measurement of correlated flux density: we measure the waveform at different locations and correlate these time series with each other to quantify how similar they are. The relatively low frequencies of radio waves allow us to capture phase information and convert the received radio waves to electrical signals that can be sampled, processed and correlated by computer systems. This correlation product is averaged over a short time interval (the integration time) to gain sensitivity and reduce the data rate. We thus get a measurement of correlated spectral power per baseline (see Figure 1.14). This correlated flux density is also expressed in Janskys, but only pertains to a specific component of the spatial variation of the interference pattern. When we sample many such pairs of sites or baselines, we can start to reconstruct the general spatial variation of the interference pattern and attempt to image the source. For sources of small angular size, there is a one-to-one correspondence between the 2D Fourier components of the source image on the sky and the visibilities sampled by the different baselines in an interferometer.

Mathematically, this relation between the interference pattern structure and the on-sky source geometry is captured in the Van Cittert-Zernike theorem (van Cittert, 1934; Zernike, 1938):

$$\Gamma_{12}(u, v, 0) = \iint I(l, m) e^{-2\pi i(ul+vm)} dl dm, \quad (1.11)$$

which expresses the mutual coherence function Γ between points labeled 1 and 2 as the integral of the source brightness distribution over the sky $I(l, m)$ convolved with a complex phase term that incorporates the baseline geometry expressed in the (u, v) -plane. This plane is defined orthogonally to the pointing direction (phase center), with the u direction oriented Eastwards and the v -direction oriented Northwards. Both u and v are expressed in observing wavelengths. The




FIGURE 1.13: A simplified depiction of the influence of source size on wavefront structure. In these plots, a source with 100 components, each radiating with a different wavelength and amplitude, radiates waves that are sampled by a dense detector array (curved white line at bottom of frame), with a snapshot of the sampled instantaneous amplitude for the three cases shown in corresponding frames of the bottom row. Top left: all 100 components are located within a fraction of the shortest considered wavelength. The aggregate source approximates the behaviour of a point source, with the wavefront being coherent across the full detector array. Top middle: the components are scattered over a larger region, such that their separation is larger than the shortest wavelength but smaller than the longest wavelength. The wavefront is no longer coherent across the full detector array, but is still coherent on shorter length scales. Top right: the components are scattered over distances larger than the longest wavelength. The wavefront now shows a lot of substructure across the detector array, with an even shorter corresponding coherence length. Shorter-scale variations in amplitude across the detector array translate to a more rapid decrease in correlated flux density between two locations in the detector array as their mutual separation increases. In this depiction, the relative scales of source-detector separation, source size, detector array size and wavelength have all been strongly altered from their 'natural' values for the sake of clarity. Figure by the author.

third component (w) in the arguments for Γ_{12} indicates the separation between points 1 and 2 along the line of sight towards the phase center. To account for the delay correction between the




FIGURE 1.14: Fundamental workings of a radio interferometer (quasi-monochromatic case). The two antennas are tracking a source in the $\hat{s}$ direction as it transits. As the geometrical delay τ_g changes with the shifting geometry, the correlator output (formed by the averaged multiplication of the two signal voltages) shows the real part of the depicted dependence on τ_g (the actual correlation product is a complex quantity, which also has phase information included). In an actual system, the geometrical delay is corrected for the boresight direction $\hat{s}$ before correlation, and the correlation products are generated for many simultaneous observing frequencies for each baseline. Image credit: NRAO.

received signals it is set to zero. l and m are defined as the direction cosines measured from the phase center, and can be approximated as angles expressed in radians for small source sizes. From this expression it is directly apparent that for a compact source (with l and m small) we need long baselines (with u and v large) in order to discern its structure. The similarity to the Fourier transform is also immediately apparent, and this similarity greatly simplifies the way in which visibilities can be used in imaging.

1.5.4 Measuring the interference pattern: correlation

In modern radio astronomy, interferometric measurements are performed by digitally recording the waveform of incoming radio waves at different locations (typically referred to as 'antennas',

'sites' or 'stations') preserving their phase information, and correlating the received waveforms either in real-time or later using the recorded data. Let us first consider a basic correlator that does not correct for the geometrical time delay and only uses real-valued input signals. This correlator multiplies the voltages of two different voltage time series coming from different antennas and integrates this product over a short time. If we imagine two identical antennas receiving a narrow-band (monochromatic) signal and assume a short integration time so that we have a stationary geometry, the correlator response can be expressed as:

$$R_{12} = \langle V_1(t)V_2(t) \rangle = \frac{1}{2}V^2 \cos(\omega\tau_g), \quad (1.12)$$

where R_{12} is the response expressing correlated flux density (a value that can be expressed in Janskys with appropriate calibration), $V_1(t)$ and $V_2(t)$ are the voltages in the two input signals, ω is the angular frequency of the incoming radiation and τ_g is the signal propagation time difference from the chosen direction on the sky (the 'phase center') to the two antennas. The product $\omega\tau_g$ is thus the phase difference that is measured between the input signals. When we allow the observation geometry to change over time as the Earth rotates, we see that the correlator response evolves over time because τ_g changes. The value of R_{12} will thus change over time in a corresponding way: it will vary between a minimum value (signals in antiphase, response is negative) and a maximum value (signals in phase, response is positive). When the signals have a relative phase of $\pi/2$ radians, the correlator response will be zero. There are thus two cases where the correlator response is zero: the phase difference between the input signals can be either $\pi/2$ or $-\pi/2$ radians. To resolve this ambiguity and get a clear phase relationship out from our voltage measurements, we can use a complex correlator. This type of correlator accepts complex-valued inputs and yields a complex-valued output:

$$\tilde{R}_{12} = \langle \tilde{V}_1(t)\tilde{V}_2(t) \rangle = \frac{1}{2}V^2(\cos(\omega\tau_g) + i \sin(\omega\tau_g)). \quad (1.13)$$

In this expression, values denoted with a tilde symbol are complex-valued. The complex-valued input signals are generated from the sampled voltages by applying a $\pi/2$ -radian phase shift to the real-valued input and adding the resulting signal to the original in quadrature (i.e., as imaginary values). The complex-valued measurements of correlated flux density ($\tilde{R}_{12}$) are termed 'interferometric visibilities', or simply 'visibilities'. These are expressed in units of Janskys, and pertain to a specific baseline orientation (i.e., angular scale and direction on the sky) and a specific spectral bin.

Modern correlators are somewhat more complex than depicted in the example above. When correlating the waveform data, the correlator needs to apply a so-called 'correlator model'. This model provides the time delays with which the various signals from different antennas need to be shifted with respect to one another in order to correct for the geometry of antenna placement: generally, the sky signal does not reach all antennas at the same time and so there is a time offset in the recorded waveforms. The correlator model also incorporates corrections for the different relative velocities that the antennas have along the line of sight to the source, because of the Earth's rotation. This movement makes the delay change over time, and gives rise to the 'delay

rate', or the derivative of delay with respect to time. The correlator applies delay and delay rate corrections to the input data streams and sums the correlated data over limited integration times before outputting it. The correlated data can still show the effects of residual errors in delay and delay rate (for instance, as a consequence of atmospheric variability affecting the signal path length for different antennas), which will need to be calibrated out in data processing.

In the case of VLBI, extra factors need to be taken into account when correlating. Because all sites operate using separate hardware, steps need to be taken to synchronise the clocks at each of the locations. Without clock synchronisation, no successful correlation can be performed as there is no way to know the correct time offset between the signals. In addition, differential clock rates mean that the apparent frequencies present in the recorded waveforms will shift and so further limit (or even destroy) the quality of the correlation products. Therefore, in modern VLBI networks, each site has a local stable clock which often comes in the form of a hydrogen maser. The signals derived from these local clocks are then disciplined using an external reference of the shared GPS clock in the form of a pulse-per-second (PPS) signal. Next to accurate and shared timekeeping, the positions of the sites also need to be carefully established and tracked - particularly for VLBI at high frequencies where the observing wavelength is comparable to the yearly motion of the sites due to continental drift, solid-Earth tidal displacement, or ice sheet movement. These effects can all be incorporated into the correlator model up to a certain accuracy.

As with a single reflector, the angular resolution that can be attained by an interferometer depends on the projected lateral separation between the elements of the array into a plane perpendicular to the viewing direction (D), combined with the wavelength of the observed radiation (λ):

$$\theta = c_a \frac{\lambda}{D}, \quad (1.14)$$

where c_a is a constant of a order unity that depends on the precise geometry of the array. When λ and D are expressed in the same units, θ is expressed in radians. In the above expression, the angular resolution has been defined as the angle from the boresight to the first null in the Airy disk pattern (this definition is called the Rayleigh criterion).

Through the years, the resolution attained by VLBI experiments has gradually improved as longer baselines were probed and shorter wavelengths were used. Today, the highest resolutions so far have been reached by RadioAstron (11 μas on a single baseline at 22 GHz, Sobolev et al. (2017)) using space VLBI and by the Event Horizon Telescope (20 μas at 230 GHz, Event Horizon Telescope Collaboration et al. (2019c)) using a global network of facilities.

1.5.5 Aperture synthesis

At any given time, each pair of antennas in an interferometer yields information on the source brightness distribution along only one direction on the sky and at one angular scale. As time

progresses and the observation of an astrophysical source continues, the rotation of the Earth generally causes the projected baseline geometry to change in both length and direction. This fact can be exploited to gain a better coverage of the (u, v) -plane and improve eventual source image quality. For highly variable sources of which the geometry can change on timescales shorter than the observation window, this can be a problem in image reconstruction as one of their implicit assumptions (a stationary source brightness distribution) is thereby violated.

Although an interferometer reaches the same angular resolution as a single telescope with an aperture of the same width, there are two major differences between these systems. One is that an interferometer has a much smaller collecting area than a single-aperture telescope of the same overall diameter, and as such it is correspondingly less sensitive. The other difference is that the point spread function on the sky, or 'dirty beam', of an interferometer is affected by the fact that only a limited number of points in the virtual aperture are sampled (see Figure 1.15). As the synthesised dirty beam is equivalent to the Fourier transform of the aperture, this changes the point spread function from an Airy disk (valid for a fully sampled circular aperture) to something much more irregular, typically with sidelobe structures in different directions on the sky. This has consequences for the way in which images are constructed using interferometric data. Because the aperture filling is incomplete, many different source geometries may be compatible with the measured cross-correlations: the points in the (u, v) -plane for which no measurements were made provide these extra degrees of freedom. As such, different assumptions on source structure can be adopted to provide additional constraints such as point-like source components, smooth brightness variations, or a mostly dark image (Högbom, 1974; Cornwell & Evans, 1985; Akiyama et al., 2017; Chael et al., 2018).

1.5.6 Data calibration

After correlation, for both connected-element interferometers and VLBI networks, further data processing needs to take place before the data is ready to be scientifically analysed. The first step that needs to be taken is to couple the signal strengths of the correlated data to physical units. For this reason, interferometric measurements generally include a brief scan on a source of known flux density (the flux density calibrator). This can be a well-studied extragalactic source, or a Solar system object for which a reliable emission model is available. The ratio of the modeled flux density to the amplitudes of the measured correlation products provides the correction factor that can be pre-applied to all the measurements in that observation.

Flux scaling is however not the only correction that needs to be applied. For one, the Earth's atmosphere affects the propagation of radio signals in multiple ways. Differences in air density or water vapour content along the lines of sight of the antennas generally lead to time-varying residual delays that need to be calibrated out. The optical depth of the atmosphere is generally




FIGURE 1.15: The role of aperture synthesis in improving (u, v) -coverage. Top row, left: the layout of the antennas of the the Karl G. Jansky Very Large Array (VLA) in A-configuration on the local map plane. Middle: the instantaneous, sparse (u, v) -coverage provided by the full array, with every pair of antennas representing one baseline vector. Right: as the Earth rotates, the projected geometry of the array with respect to the source direction changes over time, providing a more extended (but still incomplete) filling of the (u, v) -plane. Here, it is shown for a source at zero declination. The second and third rows show the aperture-synthesised (u, v) -coverage for sources at various declinations, each of which will have a correspondingly different point spread function (dirty beam). Plots: G. Roeke and P. Hinz, University of Arizona.

greater than zero, meaning that the sky signal will get attenuated and extra noise (uncorrelated, local sky emission not coming from the source of interest) is introduced into the recorded signals. The receiver systems used at each antenna also typically have time-varying gains, delays and noise levels which affect the phase and amplitude of the recorded signal as well, and the antenna structures themselves are also likely to have time-varying characteristics that affect their effective aperture size and spillover characteristics. The standard way to deal with these residual delays and gain variations is to regularly observe calibrator sources as part of the scientific observation. A suitable calibrator source is a source close to the target of interest on the sky which has a stable morphology and flux density, and (if relevant for the observation) stable polarisation properties. Measurements of the calibrator source can then be used to calibrate the measurements on the science target as the correction factors in phase and amplitude can be derived from them, interpolated in time and applied to the science measurements.

The calibration measures mentioned above are often sufficient to correct connected-element interferometry data. For VLBI at high frequencies (~ 43 GHz and above), the phase stability across even one scan is often insufficient to apply meaningful phase corrections based on neighbouring calibrator scans. When phase or amplitude calibration of interferometric visibilities is problematic, other observables can be constructed from combinations of visibilities that cancel out station-based amplitude or phase errors. These observables are closure phases and closure amplitudes.

1.5.7 Useful interferometric observables: closure phases and closure amplitudes

Closure phases (Jennison, 1958) can be constructed by using the visibilities from a closed loop of stations in the array, where we assume that the visibilities were measured over the same time interval, the same frequency range and the same polarisation product:

$$\phi_{1,2,\dots,n} = \sum \phi_{1,2} + \phi_{2,3} + \dots + \phi_{n,1}. \quad (1.15)$$

Because all station-based phase errors figure in this expression twice with opposite signs, they cancel out for this observable. Generally, closure phases are used over groups of three stations forming a triangle although they can be constructed using loops over more than three stations. Limiting the number of visibilities involved keeps the baseline-based errors from growing too large and lowering the quality of the result. In general, closure phases yield information on the aggregate degree of asymmetry of the observed source along the directions sampled by the baselines in the loop. They are difficult to interpret directly, but are useful in model fitting and image reconstruction algorithms.

Closure amplitudes (Twiss et al., 1960) are combinations of visibility measurements designed in such a way as to cancel out station-based gain errors. Closure amplitudes can be formed using

baselines between at least four stations, by combining the visibilities so that each station gain figures into both the numerator and the denominator of a ratio of visibility products:

$$C_{12,34} = \frac{|V_{12}||V_{34}|}{|V_{13}||V_{24}|}. \quad (1.16)$$

As each station appears once in the numerator and once in the denominator, station-based gains will divide out. Note that a fully-connected group of four stations (with 6 baselines between them) can yield 6 different closure amplitudes - but only two of those provide independent information, as they each use different selections of baselines within the quadrangle. As is the case with closure phases, closure amplitudes are hard to interpret directly but encode source structure ignoring station amplitude gain errors. Care must be taken when considering the behaviour of errors or uncertainties on closure quantities, as they typically exhibit non-Gaussian error distributions even if the underlying error distributions for the visibility measurements involved are themselves Gaussian (Blackburn et al., 2019).

1.6 In this thesis

At this point, we have collected numerous pieces of evidence that firmly support the existence of black holes and that broadly characterise the nature of the plasma flows around them. Even so, many aspects of these systems are still poorly understood. How do plasma jets pick up their matter? What structure do magnetic fields close to black hole event horizons have? Both the supermassive black hole at the center of the M87 elliptical galaxy (M87*) and Sgr A* have low-Eddington accretion flows, yet M87* exhibits a clear large-scale jet while Sgr A* does not. Is this a matter of inclination, or is there more to it?

The fact that Sgr A* presents an excellent opportunity for imaging at high frequencies further makes it necessary for us to understand the nature of its emission very clearly. If we want to properly interpret what EHT will see, we need to connect the behaviour of Sgr A* at other wavelengths and scales to the behaviour of the plasma flow close to its event horizon, so that we may paint a single consistent picture. When we look at Sgr A*, we wish to know what we actually see – from where around the BH does the emission originate? Given that the effects of interstellar scattering make it challenging to discern the size, shape and fundamental placement of Sgr A*'s photosphere at longer wavelengths, we need to make use of all the channels we have available for obtaining information about the intrinsic structure of Sgr A*. One of these channels is the use of high-resolution interferometry at high observing frequencies to discern the morphology of Sgr A* while avoiding the strong scattering at longer wavelengths. A second channel is studying the character of source variability across different frequencies and looking for a relation there. These two channels form the focus of the research presented in this thesis.

I present my papers as four chapters. Chapter 2 presents a closure phase analysis of VLBI measurements taken by the VLBA, GBT and LMT together at 86 GHz. This network of antennas has

given us a resolution that is sufficient to identify an East-West asymmetry in the emission from Sgr A* that may be intrinsic in nature.

In chapter 3 I present the results from an imaging and closure amplitude analysis of the VLBA + GBT + LMT dataset, yielding the most accurate angular size measurement of Sgr A* at 86 GHz obtained to date. I also discuss the size-frequency relation obtained from historical measurements and investigate how robust this relation is when different functional forms for the size vs wavelength are considered.

Chapter 4 discusses the analysis of simultaneous light curves measured by ALMA and the VLA, in which we found evidence for a time lag relation across frequencies ranging from 48 to 19 GHz. This time lag is interpreted in the context of expanding plasma volumes, and coupled to the size-frequency relation for Sgr A* an estimate is made on the outflow velocity of these plasma clouds.

A more extensive search for time lags in light curves of Sgr A* is presented in Chapter 5. I make use of multi-epoch wideband VLA data to study the persistence of time lags in the frequency range from 48 to 18 GHz. I detect time lags in a few epochs, but not in all. The chapter includes a discussion on the probable causes for this result, and on strategies for future light curve measurements.

ASYMMETRIC STRUCTURE IN SGR A* AT 3 MM FROM CLOSURE PHASE MEASUREMENTS WITH VLBA, GBT, AND LMT

Christiaan Brinkerink, Cornelia Müller, Heino Falcke, Geoffrey Bower, Thomas Krichbaum, Edgar Castillo, Adam Deller, Sheperd Doleman, Raquel Fraga-Encinas, Ciriaco Goddi, Antonio Hernández-Gómez, David Hughes, Michael Kramer, Jonathan León-Tavares, Laurent Loinard, Alfredo Montaña, Monika Mościbrodzka, Gisela Ortiz-León, David Sanchez-Arguelles, Remo Tilanus, Grant Wilson, Anton Zensus

MNRAS, 462, 2, 1382-1392 (2016)

Abstract

We present the results of a closure phase analysis of 3 mm very long baseline interferometry (VLBI) measurements performed on Sagittarius A* (Sgr A*). We have analyzed observations made in May 2015 using the Very Long Baseline Array, the Robert C. Byrd Green Bank Telescope and the Large Millimeter Telescope Alfonso Serrano and obtained non-zero closure phase measurements on several station triangles - indicative of a non-point-symmetric source structure. The data are fitted with an asymmetric source structure model in Sgr A*, represented by a simple two-component model, which favours a fainter component due East of the main source. This result is discussed in light of a scattering screen with substructure or an intrinsically asymmetric source.

2.1 Introduction

The supermassive black hole candidate at the center of our Galaxy (associated with the radio source Sagittarius A*, or Sgr A*) offers a prime possibility to study the physical phenomena

associated with accretion onto a supermassive black hole (Genzel et al., 2010; Falcke & Markoff, 2013; Goddi et al., 2017). Sgr A* is thought to accrete at an extremely low Eddington ratio (Falcke et al., 1993; Quataert & Gruzinov, 2000), an accretion regime analogous to the low-hard state in X-ray binaries and for which a jet component is expected to manifest. These expected physical behaviours and their interplay make it challenging to formulate fully self-consistent models for Sgr A* that simultaneously explain its spectrum, its variability and its size and shape on the sky. The expected angular size of the event horizon of Sgr A* on the sky ($50 \mu\text{as}$, Falcke et al., 2000) is the largest of any known black hole candidate. This makes it a prime target for studies using very long baseline interferometry at mm wavelengths (mm-VLBI), which can attain spatial resolutions that are comparable to the expected shadow size on the sky (Doeleman et al., 2008; Falcke & Markoff, 2013).

A second reason to use VLBI measurements at short wavelengths is due to the interstellar scattering that is encountered when looking at the Galactic Center in radio (Backer, 1978). Sgr A* exhibits an apparent size on the sky that is frequency-dependent, scaling with λ^2 (the exact exponent depends on the specific type of turbulence in the interstellar plasma, see Lu et al., 2011a) for observing wavelengths longer than about 7 mm (43 GHz, Bower et al., 2006). This is due to interstellar scattering by free electrons: at these wavelengths, the scattering size is significantly greater than the intrinsic source size and as such the apparent source size is dominated by the scattering effect. At wavelengths shorter than 7 mm, the apparent source size breaks away from from the λ^2 -relation and the intrinsic source size can be more easily recovered after quadrature subtraction of the known scattering size for that wavelength (Falcke et al., 2009). The shorter the observing wavelength, the more prominent the intrinsic source size and shape shine through. The relation between the intrinsic source size (i.e., the size after correcting for the scattering effect) and the observing wavelength has also been investigated, showing that the emission region itself shrinks with decreasing observing wavelength too. At an observing wavelength of 1.3 mm (230 GHz), the size of Sgr A* on the sky has been shown to be even smaller than the expected projected horizon diameter of the black hole (Doeleman et al., 2008).

The present view is that the cm- to mm-wavelength spectrum of Sgr A* is generated by partially self-absorbed synchrotron emission from hot plasma moving in strong magnetic fields close to the putative event horizon of the black hole, a model supported by recent observations and analyses thereof (Doeleman et al., 2008; Fish et al., 2011; Lu et al., 2011a; Bower et al., 2014; Gwinn et al., 2014; Fish et al., 2016; Broderick et al., 2016, and references therein). See Falcke & Markoff (2013) for a recent review on our current understanding of the nature of Sgr A*. However, the specific part of the black hole environment where this emission is thought to come from is subject to debate. Many properties of the bulk accretion flow such as density, temperature and magnetic field strength can be investigated using general-relativistic magnetohydrodynamic (GRMHD) simulations, and results from different modern simulations paint a consistent picture. However, much depends on the specific prescription for the electron temperature that is used

throughout the accretion flow. For Sgr A*, the inner region of the accretion disk has been put forward as the main emission region candidate if certain electron temperature prescriptions are used (e.g., Narayan et al., 1995), but other physically motivated prescriptions indicate that the jet launching region may dominate mm-wavelength emission instead (e.g., Mościbrodzka & Falcke, 2013). These different models yield comparable predictions for the expected overall size of the source at 86 GHz, but predict different source shapes.

To resolve this debate, gathering more accurate knowledge of the detailed brightness distribution of the source on the sky (particularly its asymmetry) plays an important role. Observations at 3.5 mm (86 GHz) provide an excellent way of studying this geometry: the emission comes from the inner accretion region, but it is not so strongly lensed as the 1.3 mm emission is thought to be. This means that the apparent source shape at 3.5 mm provides the best insight into which regions of the inner accretion flow form the source of the radiation that we receive.

The only telescope arrays that can reach the angular resolution on the sky required to potentially discern this asymmetry are the High-Sensitivity Array (HSA), the Global mm-VLBI Array (GMVA), and the Event Horizon Telescope (EHT, Doeleman et al., 2008). Before 2015 these VLBI arrays offered limited North-South (u, v) coverage for Sgr A*, which is in the Southern sky (RA: 17h45m40s, DEC: -29d00m28s), and have thusfar left the question of asymmetric source structure open. With the inclusion of the Large Millimeter Telescope Alfonso Serrano (LMT) in the HSA as of the first semester of 2015 (see Ortiz-León et al., 2016, for the description of VLBI implementation at LMT), the (u, v) coverage at 3.5 mm has been improved dramatically (see Fig. 2.1).

Using observations at longer wavelengths (ranging from 7 mm to 6 cm), for which interstellar scattering dominates the observed source size, it has been shown that the scattered source has an elongated, approximately Gaussian structure (Shen et al., 2005; Bower et al., 2006) with major and minor axes that scale with observing wavelength as $b_{\text{maj,scatt}} = 1.32 \pm 0.02 \text{ mas cm}^{-2}$ and $b_{\text{min,scatt}} = 0.67 \pm 0.02 \text{ mas cm}^{-2}$ respectively (Bower et al., 2015a). This observed Gaussian has a well-defined position angle of $81.8^\circ \pm 0.2^\circ$ East of North. Extrapolated to $\lambda=3.48 \text{ mm}$, this relation yields a scattering size of $(160 \pm 2) \times (81 \pm 2) \mu\text{as}$. Recent measurements at 3.5 mm, done with the VLBA and the LMT, indicate that the observed size is $(216 \pm 5) \times (143 \pm 8) \mu\text{as}$, at a position angle of $80.5^\circ \pm 4^\circ$ East of North - indicating that the intrinsic structure of Sgr A* is partially resolved and yielding an estimate for the intrinsic size after quadrature subtraction of the scattering size of $(147 \pm 7) \times (120 \pm 12) \mu\text{as}$ at a position angle $80^\circ \pm 7^\circ$ (Ortiz-León et al., 2016, note that we quote the more conservative closure amplitude derived results here). Moreover, the closure phases measured by that work are mentioned to be consistent with the expected values produced from the effects of interstellar scattering alone, although the cause for the non-zero closure phases may yet be intrinsic to the source.

Some recent results do suggest the presence of (possibly time-variable) asymmetry in Sgr A*, however. Persistent source asymmetry for Sgr A* has been measured at 230 GHz in observations by the EHT, where an East-West asymmetry is suggested by simple model fitting results (Fish et al., 2016). Tentative evidence for (transient) source asymmetry has also been seen in observations from 2012 at 43 GHz, as reported by Rauch et al. (2016), where one 2-hour subinterval in an 8-hour observation showed a secondary South-Eastern source component at a separation of approximately 1.5 mas. This timescale is too short for the perceived structural variation to be due to changes in the scattering screen, and would point to intrinsic structural change in the source. However, the significance of this secondary component is quoted to be at the $2\text{-}\sigma$ level.

In this work, we present our first findings obtained from observations of Sgr A* at 3 mm, involving the Very Long Baseline Array (VLBA), the Green Bank Telescope (GBT) and the newly added Large Millimeter Telescope (LMT) in Mexico. Section 2 details the observations, as well as the data reduction steps performed. In Section 3, we discuss possible instrumental causes for non-zero closure phases and verify that our observations are not significantly affected by them. Section 4 presents the measured closure phases and the model fit results. Section 5 contains discussion on the results and offers our interpretation of them. Finally, our conclusion is stated in Section 6.

2.2 Observations and initial data reduction

We present our analysis based on data from a single epoch of 3 mm HSA observations, which was recorded on May 23rd, 2015 (5:00 to 14:00 UT, project code BF114B). The track has the VLBA together with LMT and GBT as participating facilities. Of the VLBA, the following stations were involved in the observation: Brewster (BR), Fort Davis (FD), Kitt Peak (KP), Los Alamos (LA), Mauna Kea (MK), North Liberty (NL), OVRO (OV) and Pie Town (PT). Only left-circular polarisation data was recorded, at a center frequency of 86.068 GHz and a sample rate of 1024 Ms/s (2-bit) - this translates to an effective on-sky bandwidth of 480 MHz, which is divided up into 16 IFs of 32 MHz each. The 16th IF falls partly out of the recording band and was flagged throughout our dataset. We used 3C 279 and 3C 454.3 as fringe finder sources. Our check-source and secondary fringe finder was NRAO 530, and observations were done in scans of 5 minutes, alternating between NRAO 530 and Sgr A* for most of the tracks. Pointing for the VLBA was done at 43 GHz on suitable SiO masers every half hour, while the LMT and GBT did their pointing independently during the same time intervals (taking ~ 10 minutes). For the VLBA pointing solutions, we assumed that the offset between the optical axis at 3 mm and at 7 mm for each station antenna had remained stable since the last calibration run done before our observation.

The data were correlated with the VLBA DiFX software correlator (v. 2.3) in Socorro, and initial data calibration was done in AIPS (Greisen, 2003). System temperature (T_{sys}) measurements and gaincurves for LMT and GBT were imported separately, as they were not included in the a-priori calibration information provided by the correlator. Edge channels in each IF were flagged (five

channels on each side out of 64 channels, corresponding to $\sim 16\%$ of the subband), and the AIPS task APCAL was used to solve for the receiver temperatures and atmospheric opacity and to set the amplitude scale. In the initial FRING step, we used the primary fringe finder scans to correct for correlator model delay offsets and for the delay differences between IFs ('manual phasecal'), the solutions of which were then applied to all scans in the data set. The second FRING run solved for the delays and rates for all sources, using a solution interval of two minutes, while combining all IFs ($\text{APARM}(5) = 1$). Failed solutions that were flagged by the FRING task (about 10% of the total) were left out for the remainder of data reduction. No fringes on baselines to MK were found, but all other baselines did yield clear detections. At this point, the fringe-fitted data were fully frequency-averaged (channel-averaged and IF-averaged) to a single channel, and exported to UVFITS and loaded into Difmap (Shepherd, 1997).

Low source elevations during the observation can in principle cause the atmospheric coherence time to be very limited, leading to a loss of signal quality when time-averaging data that has been calibrated too coarsely in time. To verify that coherence issues would not be affecting our data quality, separate FRING runs were done with solution intervals shorter than two minutes. The length of the solution intervals in this range was found to have no significant impact on the later derived closure phase values, only increasing their uncertainties. Shorter solution intervals for FRING resulted in a larger fraction of failed solutions.

Without an accurate a-priori model of a source, phase and amplitude calibration in VLBI is notoriously tricky: the amplitude uncertainties after calibration can be as large as 10%-30% for VLBI data at 3 mm wavelengths (Martí-Vidal et al., 2012). The main reason for this is incomplete knowledge of the gain-elevation dependences, the presence of residual antenna pointing and focus errors and the highly-variable atmosphere, for which the applied opacity correction only partially corrects the time-variable absorption. For this reason, our primary goal was to look at quantities which are not station-based and which are free from local gain variations. The closure phase is such a quantity.

Closure phase is the phase of the product of visibilities (equivalently, the sum of phases) taken from three connected baselines forming a triangle where station order is respected (Jennison, 1958). Closure phases are unaffected by station-based phase fluctuations, which are typically caused by tropospheric delays due to variable weather, clock drifts from the local maser, or time-dependent characteristics of the receiver system. Such station-based phase offsets cancel out when forming the closure phase. See Rogers et al. (1995) for an extended discussion on the characteristics of closure phase uncertainties.

We used Difmap to time-average the fringe-fitted, frequency-averaged data as exported from AIPS from 0.5-second integrations into 10-second blocks (command: `uvaver 10, true`). This step was also tested with different averaging intervals, and the 10-second interval was found to yield

the highest signal-to-noise ratio (SNR) for the eventual closure phase measurements. The time averaging was done to obtain a higher SNR per datapoint, while respecting the coherence time of the atmosphere (~ 10 s to 20 s at 86 GHz). Longer time averaging intervals (15s, 30s) were found to yield compatible results, but with slightly worse noise characteristics. We chose not to phase-selfcalibrate the data in AIPS (beyond fringe-fitting at the two-minute timescale) before this step, for the main reason that it would result in a significant fraction (over 50%) of the remaining visibilities being flagged because of failed solutions from low SNR. Instead, we chose to use the closure phases derived from the 10-second averaged data directly in the subsequent stage of data reduction. The use of closure phases sidesteps the (station-based) noise issues associated with individual visibilities, avoiding a large source of error in the resulting data. The second rationale for this approach is that we wanted to perform this analysis in as much a model-independent way as possible. To assess the possible influence of frequency-dependent data artefacts on the calculated closure phases, the closure phase calculations were also done using exported data from AIPS where all 15 IFs were kept separate and in which each IF was channel-averaged. This alternative method was found to yield fully compatible closure phase values, but with slightly larger closure phase errors.

The SNR for each time- and IF-averaged closure phase measurement was high enough to avoid the potential issue of phase wrapping when averaging. We therefore averaged the closure phase measurements using error-weighted summation on the phase values. We estimated the associated error on the averaged value according to $\sigma_{cp} = \sigma_{scan} / \sqrt{n}$, where σ_{scan} is the standard deviation of the observed closure phase distribution over one scan and $\sqrt{n}$ is the square root of the number of measurements averaged within one scan (typically, $n \approx 30$). If the SNR per measurement were too low, the occurrence of phase wrapping when averaging the closure phase values would bias the result towards zero.

NRAO 530 exhibits known asymmetry in source structure (Bower et al., 1997; Lu et al., 2011b). Using NRAO 530 as a check-source for our closure phase measurements, we recover clear closure phase trends over time on most station triangles (see Fig. 2.2 for the clearest of these). When we apply the same averaging scheme to the closure phase measurements for Sgr A*, we see that the clearest closure phase measurements - with the highest SNR - are typically obtained on triangles that have both LMT and GBT as participating stations (see Fig. 2.3). All of these triangles show closure phase deviations away from zero with consistent sign, suggesting an asymmetry in the source image for Sgr A*. The relatively large closure phases measured for the GBT-LMT-KP triangle in comparison with the other triangles shown is a natural consequence of the greater East-West extent of this triangle: the model-fitting results (discussed in Section 4) show the same relatively large closure phases on this triangle, as indicated by the continuous lines in the plots. We have verified that this larger closure phase variation is not due to the station performance at KP by studying the stability of the amplitudes and phases of the visibilities on baselines to KP obtained close to the 7:00 - 8:00 UT time interval, the fringe fitting solutions (delay and rate),

the bandpass response, as well as the atmospherical stability and system temperature behaviour. None of these parameters showed aberrant behaviour.

Measurements with high SNR are also obtained on small and ‘degenerate’ triangles. Degenerate triangles are triangles that have one short baseline on which the fringe spacing on the sky is much larger than the scattered source size, and for which the visibility has an expected phase of zero. This high SNR is expected due to the large visibility amplitudes that these triangles have on their short baselines. We find that the triangles involving VLBA stations NL or OV show the lowest SNR. In the case of NL triangles, this is likely caused by the low maximum elevation of Sgr A* in the local sky. For OV it is likely due to the bad weather causing high (and rapidly fluctuating) atmospheric opacity at the site on the day of the observation.

2.3 Verifying the nature of non-zero closure phases

There is a danger that non-zero closure phases can be caused by various instrumental causes. Phase variations in the bandpass can potentially cause non-zero closure phases for a point source. This closure phase bias is given by:

$$\Delta\phi_{\text{BP}} = \arg\left(\int g_1(\nu)g_2^*(\nu)d\nu\right) + \arg\left(\int g_2(\nu)g_3^*(\nu)d\nu\right) + \arg\left(\int g_3(\nu)g_1^*(\nu)d\nu\right), \quad (2.1)$$

where the $g_i(\nu)$ -terms are the complex frequency-dependent gains for antenna i , and the integral is performed over the full observed frequency band. We checked for phase slopes across all IFs by running the AIPS task BPASS on our check-source NRAO 530 (to obtain a high SNR) with a solution interval of 5 hours, and the resulting bandpass correction does not exhibit phase slopes of more than 20 degrees across the full 0.5 GHz bandwidth for any station. By simulating point source data observed by one triangle and introducing a range of thermal noise and different phase slopes across the band for one antenna, we have separately verified that phase slopes below 2 radians (~ 116 degrees) over the full bandwidth have no significant effect on the measured closure phases. Closure phase measurements taken in separate subbands also show results highly consistent with what we see from combined subbands. We are therefore confident that the closure phases we see are not caused by bandpass calibration irregularities.

Another possible instrumental cause for non-zero closure phases is the presence of polarisation leakage for significantly linearly polarised sources. Although our observation is LCP only, the RCP component of incoming radiation bleeds into the LCP signal chain in a limited way, and this may cause anomalous closure phases. The expression describing the closure phase bias from

polarisation leakage is given by:

$$\begin{aligned} \Delta\phi_{\text{pol}} = & \arg(1 + P \cdot D_1 \exp i\psi_1 + (P \cdot D_2 \exp i\psi_2)^*) + \\ & \arg(1 + P \cdot D_2 \exp i\psi_2 + (P \cdot D_3 \exp i\psi_3)^*) + \\ & \arg(1 + P \cdot D_3 \exp i\psi_3 + (P \cdot D_1 \exp i\psi_1)^*) \end{aligned} \quad (2.2)$$

In the above expression, P represents the linear polarisation fraction of the source while D_i is the complex polarisation leakage term from RCP to LCP for a given antenna i . ψ_i is the difference between the position angle on the sky of the source polarisation vector and the parallactic angle of antenna i . When we use an upper bound for the correlated linear polarisation of Sgr A* at 3 mm as being 2% (Bower et al., 1999b; Macquart et al., 2006) and the magnitude of the complex D-terms as being at most 10% (as indicated by recent GMVA results: see Table 1 in Martí-Vidal et al., 2012), we get negligible leakage-induced closure phase errors ($\Delta\phi_{\text{pol}} < 6 \cdot 10^{-4}$ deg) if we let the ψ_i -values vary so as to get the maximum possible polarisation leakage.

2.4 Results

2.4.1 Detection of non-zero closure phases

The four triangles formed by the LMT, GBT and one of the four southwest VLBA stations (FD, KP, LA or PT) are the triangles that show the clearest evolution of closure phase with time. They suggest closure phase trends for Sgr A* with time that seem mutually compatible (see Fig. 2.3), due to the roughly similar orientations and lengths in the (u, v) plane probed by their baselines. We note that the magnitude of the closure phase deviation from zero depends on the extent of the triangles in a nonlinear fashion, as was tested for a range of triangle geometries using simple source models with asymmetry, potentially explaining why the closure phases on the GBT-LMT-KP triangle are larger than those seen on the other triangles in the plot. However, all of these triangles show closure phase deviations away from zero in the same direction.

2.4.2 Modeling source asymmetry using closure phases

We wish to investigate the possible presence of point asymmetry of Sgr A* using closure phase measurements. The simplest model that can exhibit any asymmetry and non-zero closure phases is a model using two point source components, with the two components having unequal flux densities. Although the average scattering ellipse would suggest that any point-like source components should show up as 2D Gaussians, the fact that the scattering screen itself can impose substructure on even smaller angular scales provides additional motivation for this simple model. We thus use model components that would actually appear to us on the sky as unscattered point sources. More complicated source models are of course possible (for instance a source model with two components with different shapes, or having more than two components), but we will restrict

ourselves to this simple two-component model to avoid overinterpretation of our measurements. Thus, the model fitting we do in this work is meant to investigate whether the non-zero closure phases we see are compatible with an observed source asymmetry in some specific direction on the sky. The possible causes of any observed asymmetry (intrinsic or scattering) will be discussed in section 5.

In the model fitting, we determine the placement on the sky and the relative flux density of a secondary source component that gives the closure phase evolution that is most consistent with our observations. To determine the best fit parameters, we use the χ^2 statistic to compare the closure phases generated by the source model to the measured closure phases. In the model fitting procedure, the position of the secondary component on the sky and its flux density expressed as a fraction of the flux density of the main component are varied independently. The fitting procedure was tested using our observations of NRAO 530, using a range of flux density ratios (0.01 to 0.99, step size 0.01) and possible secondary source component positions on the sky (up to $600 \mu\text{as}$ separation in both RA and DEC, with step size $30 \mu\text{as}$ - forming a square grid on the sky) that was motivated by existing maps for NRAO 530 at 3mm (Lu et al., 2011b). The favoured position for the secondary source component is in excellent agreement with the source structure from our preliminary mapping results (full mapping results will be published separately), capturing the location of the dominant secondary source component. These outcomes are also in line with the previously observed structure for NRAO 530 at 86 GHz (Lu et al., 2011b), validating this simple model fitting approach. See Figure 2.4 for an illustration of this fit result.

For Sgr A*, the position of the secondary component relative to the primary component on the sky was independently varied from $-400 \mu\text{as}$ to $400 \mu\text{as}$ with step size $20 \mu\text{as}$ in both right ascension and declination, and the dimensionless flux density ratio of the secondary to the primary component was varied from 0.01 to 0.99 in steps of 0.01. The source model used thus has three free parameters. The resulting closure phases as a function of time were simulated for all triangles and the χ^2 statistic was calculated using the model curves with all of our observed data for Sgr A*. For practically all flux density ratios, the best-fit position on the sky for the secondary component is $\sim 100 \mu\text{as}$ East of the primary component (Fig. 2.5). The flux density ratio exhibits multiple local minima in χ^2 , at 0.03, 0.11 and 0.70 respectively. The flux density ratio is evidently not well-constrained by closure phases only. To constrain this flux density ratio, careful amplitude calibration of the data is needed. Results based on the fully calibrated dataset will be the subject of a separate publication. As full amplitude calibration is a tricky and involved process particularly for LMT and GBT, we have avoided relying on amplitude calibration here. While the direction of the source asymmetry on the sky is well-constrained, the uncertainty in the flux density ratio implies that there is significant uncertainty in the angular separation of the secondary component with respect to the primary.

We have done this minimum χ^2 search for different choices regarding the triangles included. We

have considered the following options: 1) all triangles, 2) only triangles including either GBT or LMT, 3) only triangles involving both LMT and GBT, 4) all triangles without the LMT, 5) all triangles without the GBT, 6) VLBA-only triangles. We find the previously quoted secondary component position to give the lowest χ^2 scores for all of these cases, with the strongest significance for case 3. It appears that inclusion of VLBA-only triangles diminishes the significance of the result, as these triangles tend to add only noise to the data to be fitted to. We show the modeled closure phase evolution for several triangles in Fig. 2.3, along with the reduced χ^2 results for both the two-component model and the zero model. The overall χ^2 scores for the best-fitting model and the zero-model can be seen in Table 2.1. The two-component model shows a better fit than the symmetric ‘zero’ model, with the significance of this difference varying according to which set of triangles is considered. We do not expect to find reduced χ^2 scores very close to 1, since the two-component model is likely an oversimplified representation of the actual source geometry. However, this simple model fit does indicate the direction on the sky for which Sgr A* shows asymmetry.

The GBT-LMT-KP triangle exhibits the strongest deviations from zero in its closure phases. This gives rise to the question whether the model-fitting results are dominated by the influence of this individual triangle. To investigate this possibility, we consider the χ^2 scores we get when omitting any of those 3 stations from the array, limiting ourselves to the triangles that can be formed with the other stations. The results are shown in the left column of Figure 2.6. We see that the preference for an offset secondary source component to the East persists in all cases, but that the significance of the result is affected by the omission of the station in question. For instance, leaving out the LMT gives a fit result that is much less constrained in the North-South direction - as can be expected from the (u, v) coverage offered by the LMT. The favoured offset position of the secondary source component also persists when any other station from the array is dropped. These results indicate that the fit results are not dominated by possible data artefacts associated with a specific station or baseline.

The relatively rapid changes in the measured closure phase on the GBT-LMT-KP triangle (see Figure 2.3) are not fully captured by the 2-component source model, and suggest a possible time-variable source structure for Sgr A*. Time-segmentation of the measurement data into 1-hour blocks and running the model-fit algorithm on these individual timeframes however shows no significant deviation of the secondary component in the time segment for 7 to 8 UT versus the best-fit position seen in other blocks: the found positional offsets for different time blocks are mutually compatible. This however only indicates a constant structure when the 2-point source model is assumed. More sophisticated model fits may still exhibit time-variable structure.

2.4.3 Testing the significance of the observed asymmetry

We need to verify that the asymmetry in Sgr A* as suggested by the closure phase measurements is significant. To this end, we have synthesized a control dataset in which every data point has the same measurement error as the corresponding measurement point in the original dataset. The measurement values in this control dataset have values drawn from a zero-mean normal distribution using the original measurement errors for the standard deviation. We thus get a simulated set of closure phases that corresponds to a point-symmetric source on the sky, with zero closure phases for all independent triangles to within measurement errors. Searching for the best-fitting two-component model using this simulated dataset in the way described above, we see that the best-fit χ^2 is comparable to the zero-offset χ^2 (see Table 2.1 and Fig. 2.7). This in contrast to the results we get with the real dataset, where we see that the two-component model fit consistently shows a preference for an offset source component. For the zero-closure phase control dataset, we also see that the best χ^2 value does not show a clear dependence on the flux density ratio - which is to be expected, as the best-fit position of the secondary component tends to be at the origin and hence produces zero closure phases regardless of flux density ratio.

We have further assessed the uncertainty in the fitted position for the secondary source component using a bootstrapping algorithm. Bootstrapping was done by synthesizing a new closure phase dataset from the existing closure phase data by repeatedly picking measurement points at random and independently from the measured dataset and adding these to a new, synthesized dataset. The final synthesized dataset contains as many data points as the original, but typically contains multiple copies of several original measurement points and misses other original measurement points. Such a synthesized dataset was generated 1000 times and the model fitting procedure was performed on each of them. This yielded a distribution of best-fit secondary source component positions which we used to define confidence intervals on this position, see Figure 2.8. The major advantage of bootstrapping is that it is robust against the presence of a subset of data points that would otherwise dominate the results of a model fitting procedure. As the result from the bootstrapping procedure agrees with the result from the original model fitting, we conclude that the asymmetry of the source we see from the original model fitting is something that is present in the dataset as a whole rather than something arising from a small selection of measurement points.

2.5 Discussion

We argue that VLBI observations at 3 mm probe a sweet spot in frequency, making them ideally suited to investigate the source structure and size. This is on one hand because the influence of interstellar scattering diminishes strongly with increasing frequency - observations at lower frequencies are more strongly influenced by scattering effects (leaving little to no opportunity to study intrinsic source structure). On the other hand, observations at higher frequencies are ex-

TABLE 2.1: Sgr A* χ^2 scores for the best-fit 2-component model ('2pt') and zero-closure phase model ('0'), shown for the actual closure phase measurements and for the synthetic, zero-compatible dataset. Each line in the table is valid for a different combination of selected stations in the array. Columns headed $\chi^2/\text{d.o.f.}$ indicate the χ^2 value and the degrees of freedom, while columns headed χ^2_{red} give the reduced χ^2 figures for convenience.

Stations in triangles	$\chi^2/\text{d.o.f.}$ (2pt)	χ^2_{red} (2pt)	$\chi^2/\text{d.o.f.}$ (0)	χ^2_{red} (0)
Measurements				
All	2252/1564	1.440	2432/1567	1.552
GBT and/or LMT	1116/889	1.255	1283/892	1.438
both LMT and GBT	135/118	1.140	241/121	1.994
no LMT	1608/1025	1.569	1657/1028	1.612
no GBT	1624/1088	1.493	1683/1091	1.543
no KP	1499/1025	1.463	1590/1028	1.547
VLBA only	1120/671.0	1.669	1149/674.0	1.705
Synthetic data				
All	1576/1564	1.008	1580/1567	1.009
GBT and/or LMT	886/889	0.996	889/892	0.997
both LMT and GBT	104/118	0.884	109/121	0.904
no LMT	1044/1025	1.019	1050/1028	1.021
no GBT	1110/1088	1.020	1113/1091	1.020
no KP	1027/1025	1.002	1030/1028	1.002
VLBA only	688/671.0	1.026	691/674.0	1.026




FIGURE 2.1: The (u, v) -coverage for the observation of Sgr A* taken on May 23rd, 2015 (6:00 - 13:00 UT). Baselines within the VLBA are coloured black, baselines to LMT and GBT are coloured orange. No baselines to Mauna Kea (MK) are shown, as we have not found fringes for Sgr A* on any baseline to MK. The inclusion of LMT improves North-South (u, v) -coverage, while the inclusion of GBT improves East-West coverage.




FIGURE 2.2: Closure phase versus time for NRAO 530 for the four ‘central’ triangles in the array. The closure phase depends nonlinearly on the East-West extent of the triangle, which is the reason the GB-LM-KP triangle exhibits a somewhat different closure phase evolution.

pected to show a source geometry that is increasingly dominated by strong lensing effects around the black hole shadow. Both of these cases throw up obstacles when studying the geometry of the inner accretion flow itself. Observations at 3 mm thus mitigate some of the complexities of interpretation associated with observations at longer and shorter wavelengths: while the effects of interstellar scattering still cannot be ignored at 3mm, intrinsic source geometry can be distinguished from scatter-induced features given multiple observations.

We deduce that Sgr A* exhibits asymmetry in the East-West direction, with a source geometry that features a weaker source component about $100 \mu\text{s}$ to the East (PA: $\sim 90^\circ$) of the main source (where we note that the separation is poorly constrained). Earlier observations at 86 GHz than those done over the last year were limited by the available (u, v) coverage, and thus the best intrinsic source sky models were limited to anisotropic, but symmetrical (2D) Gaussians. The scattering kernels were modeled as Gaussians as well, allowing subtraction in quadrature of the scattering kernel from the best-fit observed source Gaussian. This approach has yielded an intrinsic source size that showed an elongated source shape along an approximately East-West




FIGURE 2.3: Closure phase versus time for Sgr A* for the four ‘central’ triangles in the array. All four triangles share a qualitatively similar evolution of the closure phase with time. The two-component model closure phase curve derived from a global fit has been superimposed for each triangle. Reduced χ^2 scores per triangle for both the two-component model (label ‘2pt’) and the baseline zero closure phase model (label ‘zero’) are indicated in the top right of every plot. The simple two-component source model matches the measurements better than the zero closure phase model in every case. The closure phase evolution for the two-component model is sensitive to the East-West extent of the triangle in a nonlinear way, hence the larger predicted closure phase deviations for the GB-KP-LM triangle.




FIGURE 2.4: The dependence of the reduced χ^2 statistic on secondary source position for the best-fitting flux density ratio, as found for NRAO 530 (using the closure phase measurements on all triangles). There is a strong preference for the secondary component to be situated North and slightly East of the main component (sky coordinates are expressed relative to the main component). While a 2-point-source model does not capture the structure of NRAO 530 in detail (as reflected by the high reduced χ^2 value of 11.8 for the best fit), it does capture the orientation and separation of the dominant off-center component. This is illustrated by the fact that the preferred position of the secondary component agrees with the position of the brightest off-center component found in our preliminary imaging for NRAO 530 (overlaid contours). Contour levels are 1,2,4,8,16,32 and 64 percent of image maximum to guide the eye, with the absolute flux density calibration to be addressed in a future publication. The found flux density ratio between secondary and primary component is ~ 0.22 for the model fit, compatible with what is found for the preliminary imaging result.





FIGURE 2.5: Left column: Inverse reduced χ^2 value as function of the sky position of the secondary component, for three different selections of involved triangles and shown for the flux density ratio yielding the lowest χ^2 value. The position of the best fit is indicated with a white cross in each of the plots. This position is robust: for all flux density ratios, the lowest reduced χ^2 score is obtained for a secondary component towards the East of the primary component. Right column: the best reduced χ^2 value found for each flux density ratio. Local minima occur around flux ratios of 0.03, 0.11, and 0.70 for all triangle selections.



direction. We note that the best-fit position for the secondary component falls along the major axis of the scattering ellipse as it was measured by Bower et al. (2014, 2015a) and is also compatible with the previously observed intrinsic elongation of the source quoted in these publications.

These observations cover a single epoch and were done in a single frequency band and in a single polarisation (LCP), which complicates interpretation of the observed asymmetry. On one hand, interstellar scattering of the source image can introduce small-scale scintels whose ensemble average influences the observed brightness distribution (Gwinn et al., 2014; Johnson & Gwinn, 2015) and that may be responsible for the occurrence of non-zero closure phases (Ortiz-León et al., 2016). The time scale for the scattering geometry to change significantly ($\sim$ weeks) is thought to be much longer than the length of one observation ($\sim$ hours), causing the source image to be affected by an effectively static scattering screen that may induce asymmetry in the observed image. On the other hand, the observed asymmetry may be intrinsic to the source itself. Obser-




FIGURE 2.7: The typical χ^2 landscape for Sgr A* as resulting from the model fit when using the synthetic (zero-mean) closure phase data (using all triangles). Chi-squared scores do not strongly depend on flux density ratio for the synthetic data.

FIGURE 2.8: Confidence regions (black lines) for the best-fit position of the secondary source component, obtained by bootstrapping the original closure phase dataset. The innermost contour indicates the 68% confidence region, surrounded by the 95% and 99% regions respectively.

variations at different frequencies (e.g., at 230 GHz and 43 GHz) and performed at different epochs (separated in time by months) are therefore crucial in interpreting the character of this observed asymmetry.

The 86 GHz observations published by Ortiz-León et al. (2016) do show non-zero closure phases, but these have been interpreted consistently as arising from interstellar scattering effects. As such, no dedicated closure-phase modelling comparable to the analysis presented in this work was performed. Those data are separated in time from the observation we report in this work by approximately one month (April 27th vs May 23rd, 2015). Future studies of the non-zero closure phase evolution with time will help to distinguish its origin: if the observed asymmetry is persistent across both datasets, the case for an intrinsic cause of the asymmetry will be bolstered as scattering effects are expected to vary over shorter timescales (Johnson & Gwinn, 2015). Conversely, if the earlier data show a different asymmetry from what we find here the likely cause for it will be confirmed as being interstellar scattering.

Interestingly, an East-West asymmetry in Sgr A* is also suggested by closure phase results from measurements taken with the Event Horizon Telescope at 230 GHz, in the Spring of 2013 (Fish et al., 2016). The observations presented in that work show closure phases at 1.3 mm that are comparable in magnitude to the values we have measured at 3.5 mm, suggesting a similar degree of source asymmetry in both observed emission patterns. While a source model with disconnected

components is not necessarily favoured by the EHT data, fit results using a model consisting of 2 point sources suggest a preference for an East-West asymmetry in that dataset. It is somewhat surprising that the persistent asymmetry at 230 GHz is oriented along the same direction on the sky as the asymmetry found in this work. At 230 GHz a persistent asymmetry in the source image is expected, and is thought to be caused by the Doppler boosting of emission from one side of the inner accretion flow with a velocity component along our line of sight (Broderick et al., 2016). Conversely, at 86 GHz this effect is not expected to be a dominant contribution to source asymmetry - rather, the main part of any intrinsic asymmetry is expected to be a consequence of the relative brightness of the inner accretion flow versus emission from the footpoints of a compact jet component (Mościbrodzka et al., 2014). In the context of this model, the similar orientation of the asymmetry in the 230 GHz and 86 GHz observations cannot be reconciled if both are assumed to be intrinsic to the source.

For spectrally fitted jet models, a significant component of the emission at 86 GHz is generated around the jet base (Mościbrodzka & Falcke, 2013), causing the corresponding source image to exhibit an asymmetry that is aligned with the jet axis to within ~ 20 degrees. In this context the results from this work, when combined with other existing measurements of Sgr A* closure phases at 3 mm, offer an appropriate starting point for a more extended model fitting procedure, where the raytraced results from GRMHD simulations can be compared to the constraints on the observed source geometry. An analogous analysis has been performed on the published 230 GHz closure phase measurements in (Broderick et al., 2016), where the measurements have been interpreted within the context of a particular theoretical source model. This more elaborate model fitting procedure using the full available body of 86 GHz closure phase data is the focus of a separate publication that currently is in preparation.

2.6 Summary and conclusions

We have performed an observation of Sgr A* at 86 GHz, using the VLBA, the GBT, and the LMT. Elementary model fitting of a multicomponent source geometry to the closure phases from this dataset shows a preference for an Eastern secondary source component at an on-sky separation of $\sim 100 \mu\text{as}$ from the primary component. This asymmetry, when considered as a standalone observation, may be explained by interstellar scattering effects. However, this does not exclude the possibility of the observed asymmetry being intrinsic to the source.

The results by Fish et al. (2016) at 230GHz, Ortiz-León et al. (2016) at 86GHz, and Rauch et al. (2016) at 43 GHz indicate asymmetric emission of Sgr A* at different frequencies and over different time periods. In particular the closure-phase measurements performed at 230 and 43 GHz point towards a similar East-West asymmetry as was found in the dataset presented in this work. The similar orientation of this asymmetry across these different wavelengths is a puzzling

result, and future analysis of 86 GHz VLBI measurements done at different times will help to pin down the origin of these observed non-zero closure phases.

Acknowledgements

We wish to express our gratitude to the MIT Haystack team (Lindy Blackburn, Laura Vertatschitsch, Jason Soohoo), who installed the recording system at LMT and who have played an instrumental role in making VLBI measurements possible at LMT. We thank Frank Ghigo at GBT for his help in obtaining the system temperature measurements. We thank Michael Johnson for illuminating discussions on the scattering screen and on closure phase statistics, and we appreciate the input on the draft we received from Eduardo Rós. This work is supported by the ERC Synergy Grant “BlackHoleCam: Imaging the Event Horizon of Black Holes” (Grant 610058). A.H., L.L., and G.N.O.-L. acknowledge the financial support of CONACyT, Mexico and DGAPA, UNAM.

MICRO-ARCSECOND STRUCTURE OF SAGITTARIUS A*

REVEALED BY HIGH-SENSITIVITY 86 GHz VLBI

OBSERVATIONS

Christiaan Brinkerink, Cornelia Müller, Heino Falcke, Sara Issaoun, Kazunori Akiyama, Geoffrey Bower, Thomas Krichbaum, Adam Deller, Edgar Castillo, Sheperd Doeleman, Raquel Fraga-Encinas, Ciriaco Goddi, Antonio Hernández-Gómez, David Hughes, Michael Kramer, Jonathan Léon-Tavares, Laurent Loinard, Alfredo Montaña, Monika Mościbrodzka, Gisela Ortiz-León, David Sanchez-Arguelles, Remo Tilanus, Grant Wilson, Anton Zensus
A&A, 621, A119 (2019)

Abstract

The compact radio source Sagittarius A* (Sgr A*) in the Galactic centre is the primary supermassive black hole candidate. General relativistic magnetohydrodynamical (GRMHD) simulations of the accretion flow around Sgr A* predict the presence of sub-structure at observing wavelengths of ~ 3 mm and below (frequencies of 86 GHz and above). For very long baseline interferometry (VLBI) observations of Sgr A* at this frequency the blurring effect of interstellar scattering becomes sub-dominant, and arrays such as the high sensitivity array (HSA) and the global mm-VLBI array (GMVA) are now capable of resolving potential sub-structure in the source. Such investigations help to improve our understanding of the emission geometry of the mm-wave emission of Sgr A*, which is crucial for constraining theoretical models and for providing a background to interpret 1 mm VLBI data from the Event Horizon Telescope (EHT). Following the closure phase analysis in our first paper, which indicates asymmetry in the 3 mm emission of Sgr A*, here we have used the full visibility information to check for possible sub-structure. We extracted source size information from closure amplitude analysis, and investigate how this constrains a combined fit of the size-frequency relation and the scattering law for Sgr A*. We

performed high-sensitivity VLBI observations of Sgr A* at 3 mm using the Very Long Baseline Array (VLBA) and the Large Millimeter Telescope (LMT) in Mexico on two consecutive days in May 2015, with the second epoch including the Greenbank Telescope (GBT). We confirm the asymmetry for the experiment including GBT. Modelling the emission with an elliptical Gaussian results in significant residual flux of ~ 10 mJy in south-eastern direction. The analysis of closure amplitudes allows us to precisely constrain the major and minor axis size of the main emission component. We discuss systematic effects which need to be taken into account. We consider our results in the context of the existing body of size measurements over a range of observing frequencies and investigate how well-constrained the size-frequency relation is by performing a simultaneous fit to the scattering law and the size-frequency relation. We find an overall source geometry that matches previous findings very closely, showing a deviation in fitted model parameters less than 3% over a time scale of weeks and suggesting a highly stable global source geometry over time. The reported substructure in the 3 mm emission of Sgr A* is consistent with theoretical expectations of refractive noise on long baselines. However, comparing our findings with recent results from 1 mm and 7 mm VLBI observations, which also show evidence for east-west asymmetry, we cannot exclude an intrinsic origin. Confirmation of persistent intrinsic substructure will require further VLBI observations spread out over multiple epochs.

3.1 Introduction

The radio source Sagittarius A* (hereafter called Sgr A*) is associated with the supermassive black hole (SMBH) located at the centre of the Milky Way. It is the closest and best-constrained supermassive black hole candidate (Ghez et al., 2008; Gillessen et al., 2009; Reid, 2009) with a mass of $M \sim 4.1 \times 10^6 M_\odot$ at a distance of ~ 8.1 kpc as recently determined to high accuracy by the GRAVITY experiment (Gravity Collaboration et al., 2019). This translates into a Schwarzschild radius with an angular size of $\theta_{R_S} \sim 10 \mu\text{as}$ on the sky, while the angular size of its ‘shadow’ – meaning the gravitationally lensed image of the event horizon – is predicted to be $\sim 50 \mu\text{as}$ (Falcke et al., 2000). Due to its proximity, Sgr A* appears as the black hole with the largest angular size on the sky and is therefore the ideal laboratory for studying accretion physics and testing general relativity in the strong field regime (see, e.g. Goddi et al., 2017; Falcke & Markoff, 2013, for a review).

Radio observations of Sgr A* have revealed a compact radio source with an optically thick spectrum up to mm-wavelengths. In the sub-mm band the spectrum shows a turnover and becomes optically thin. This sub-mm emission is coming from a compact region that is only a few Schwarzschild radii in size (e.g. Falcke et al., 1998; Doeleman et al., 2008). Very Long Baseline Interferometry (VLBI) observations can now achieve angular resolution down to a few tens of μas , which is required to resolve these innermost accretion structure close to the event horizon. The

advantages in going to (sub-)mm wavelengths are 1.) to witness the transition from optically thin to thick emission, 2.) to improve the angular resolution and 3.) to minimize the effect of interstellar scattering. At longer radio wavelengths, interstellar scattering along our line of sight towards Sgr A* prevents direct imaging of the intrinsic source structure and causes a ‘blurring’ of the image that scales with wavelength squared (e.g. Davies et al., 1976; Backer, 1978; Bower et al., 2014).

The scatter-broadened image of Sgr A* can be modelled by an elliptical Gaussian over a range of wavelengths. The measured scattered source geometry scales with λ^2 above observing wavelengths of ~ 7 mm (Bower et al., 2006) following the relation: $(\theta_{\text{maj}}[\text{mas}]) \times (\theta_{\text{min}}[\text{mas}]) = (1.31 \times 0.64)(\lambda[\text{cm}])^2$, with the major axis at a position angle 78° east of north. At shorter wavelengths this effect becomes sub-dominant, although refractive scattering could introduce stochastic fluctuations in the observed geometry that vary over time. This refractive noise can cause compact sub-structure in the emission, detectable with current VLBI arrays at higher frequencies (Johnson & Gwinn, 2015; Gwinn et al., 2014).

Due to major developments in receiver hardware and computing that have taken place in recent years, mm-VLBI experiments have gotten closer to revealing the intrinsic structure of Sgr A*. At 1.3 mm (230 GHz), the Event Horizon Telescope has resolved source structure close to the event horizon on scales of a few Schwarzschild radii (Doeleman et al., 2008; Johnson et al., 2015). Closure phase measurements over four years of observations have revealed a persistent east-west asymmetry in the 1.3 mm emission of Sgr A* (Fish et al., 2016). This observed structure and geometry seems intrinsic to the source and is already imposing strong constraints on general-relativistic magnetohydrodynamic (GRMHD) simulation model parameters of Sgr A* (Broderick et al., 2016; Fraga-Encinas et al., 2016). A more recently published analysis by (Lu et al., 2018) of observations done at 230 GHz including the APEX antenna reports the discovery of source sub-structure on even smaller scales of 20 to 30 μas that is unlikely to be caused by interstellar scattering effects.

At 3.5 mm (86 GHz), the combined operation of the Large Millimeter Telescope (LMT, Mexico) and the Green Bank Telescope (GBT, USA) together with the Very Long Baseline Array (VLBA) significantly improves the (u,v)-coverage and array sensitivity beyond what is possible with the VLBA by itself. Closure phase analysis indicates an observational asymmetry in the 3 mm emission (Ortiz-León et al., 2016; Brinkerink et al., 2016), which is consistent with apparent sub-structure introduced by interstellar scattering, although an interpretation in terms of intrinsic source structure cannot be excluded given the data obtained so far. Ortiz-León et al. (2016) reported on VLBA+LMT observations at 3.5 mm detecting scattering sub-structure in the emission, similar to what was found at 1.3 cm by Gwinn et al. (2014). In Brinkerink et al. (2016), using VLBA+LMT+GBT observations, we report on a significant asymmetry in the 3.5 mm emission of Sgr A*. Analysing the VLBI closure phases, we find that a simple model with two point sources of unequal flux provides a good fit to the data. The secondary component is found to be

located towards the east of the primary, however, the flux ratio of the two components is poorly constrained by the closure phase information.

It remains unclear, however, whether this observed emission sub-structure at 3.5 mm is intrinsic or arises from scattering. With the body of VLBI observations reported so far we cannot conclusively disentangle the two components. Time-resolved and multi-frequency analysis of VLBI data can help. Besides the findings by Fish et al. (2016) at 1.3 mm, Rauch et al. (2016) found a secondary off-core feature in the 7 mm emission appearing shortly before a radio flare, which can be interpreted as an adiabatically expanding jet feature (see also Bower et al., 2004).

From elliptical fits to the observed geometry of the emission, the two-dimensional size of Sgr A* at mm-wavelength can be derived as reported by Shen et al. (2005); Lu et al. (2011a); Ortiz-León et al. (2016) at 3.5 mm and Bower et al. (2004); Shen (2006) at 7 mm. Using the known scattering kernel (Bower et al., 2006, 2014), this intrinsic size can be calculated from the measured size. The most stringent constraint on the overall intrinsic source diameter has been determined using a circular Gaussian model for the observed 1.3 mm emission (Doeleman et al., 2008; Fish et al., 2011), as at this observing frequency the scattering effect is less dominant. More recent VLBI observations of Sgr A* at 86 GHz constrain the intrinsic, two-dimensional size of Sgr A* to $(147 \pm 4)\mu\text{as} \times (120 \pm 12)\mu\text{as}$ (Ortiz-León et al., 2016) assuming a scattering model derived from Bower et al. (2006) and Psaltis et al. (2015).

High-resolution measurements of time-variable source structure in the infrared regime observed during Sgr A* infrared flares have recently been published (Gravity Collaboration et al., 2018b), where spatial changes of the source geometry of Sgr A* on timescales of less than 30 minutes are seen. These results suggest periodical motion of a bright source component located within $\sim 100\mu\text{as}$ of the expected position of the supermassive black hole, with a correspondingly varying polarisation direction. The variability timescale of Sgr A* is expected to be significantly shorter at infrared wavelengths than at 3.5 and 1.3 mm, as it is thought to be dominated by fast local variations in electron temperature rather than changes in the bulk accretion rate.

All of these observations indicate that we start to unveil the presence of both stationary and time-variable sub-structure in the accretion flow around Sgr A*, as expected by theoretical simulations (e.g. Mościbrodzka et al., 2014). In order to further put constraints on model parameters, higher-resolution and more sensitive mm-VLBI observations are required. The analysis of closure quantities helps to determine source properties without being affected by station-based errors. Closure phases indicate asymmetry in the emission when significantly deviating from zero (see, e.g. Fish et al., 2016; Brinkerink et al., 2016, for the case of Sgr A*). Closure amplitudes put constraints on the source size (see, e.g. Ortiz-León et al., 2016; Bower et al., 2006, 2004). Imaging techniques are based on the closure quantities. Although mm-VLBI has a number of limitations, at $\gtrsim 3$ mm the current VLBI array configurations allow reconstructing the emission of Sgr A*

using standard hybrid imaging techniques (Lu et al., 2011a; Rauch et al., 2016).

In this paper we have followed-up on our first analysis published in Brinkerink et al. (2016) (hereafter referred to as Paper I). Here, we focus on the closure amplitude and imaging analysis of Sgr A* at $\lambda = 3.5$ mm obtained with the VLBA and LMT on May 22nd, 2015 and VLBA, LMT, and GBT on May 23rd, 2015. In Section 3.2 we describe the observations and data reduction. Section 3.3 discusses the results from imaging and closure amplitude analysis. In section 4, we present the results from a simultaneous fitting of the intrinsic size-frequency relation and the scattering relation for Sgr A*, using the combined data from this work with earlier published results across a range of wavelengths. We conclude with a summary in Section 3.5.

3.2 Observations and data reduction

We performed 86 GHz VLBI observations of Sgr A*. Here we present the analysis of two datasets: one epoch using the VLBA (all 86 GHz capable stations¹) together with the LMT (project code: BF114A) on May 22nd, 2015, and one epoch using VLBA, LMT, and GBT on May 23rd, 2015 (project code: BF114B). Both observations were made in left-circular polarisation mode only, at a centre frequency of 86.068 GHz and a sampling rate of 2 Gbps (512 MHz on-sky bandwidth). For fringe finding we used the primary calibrators 3C 279 and 3C 454.3 at the start and end of the track respectively. In between, the scans alternated every 5 min between Sgr A* and the secondary fringe finder NRAO 530 ([HB89] 1730-130) with short regular gaps (every ~30 minutes) for pointing and longer GBT-only gaps every ~4 hours for focusing.

For fringe finding and initial calibration of both datasets, we used standard methods in *AIPS* (Greisen, 2003) as described in Paper I. We first performed a manual phase-cal to determine the instrumental delay differences between IFs on a 5 min scan of 3C 454.3. After applying this solution to all data, the second FRING run gave us solutions for delay and rate (4 min solution interval, with 2 min subintervals) with a combined solution for all IFs. Using shorter solution intervals than the length we used here resulted in more failed, and therefore flagged, FRING solutions. All telescopes yielded good delay and rate solutions for NRAO 530. For Sgr A*, however, we found no FRING solutions on baselines to MK (using a limiting value for the signal-to-noise ratio (S/N) of 4.3), but all other baselines yielded clear detections.

Amplitude calibration in *AIPS* was performed using a-priori information on weather conditions and gain-elevation curves for each station. In the cases of the LMT and the GBT, system temperature measurements and gain curves were imported separately as they were not included in the a-priori calibration information provided by the correlator pipeline. We solved for (and ap-

¹Brewster (BR), Fort Davis (FD), Kitt Peak (KP), Los Alamos (LA), Mauna Kea (MK), North Liberty (NL), Owens Valley (OV) and Pie Town (PT)

plied) atmospheric opacity corrections using the *AIPS* task *APCAL*. To prepare for the remaining amplitude corrections, the data were then IF-averaged into a single IF and exported to *Difmap* (Shepherd, 1997).

The quality of millimeter-VLBI observations is in practice limited by a number of potential error contributions (cf. Martí-Vidal et al., 2012): atmospheric opacity and turbulence, and telescope issues (e.g. pointing errors). In the case of Sgr A*, the low elevation of the source for northern hemisphere telescopes requires a careful calibration strategy as loss of phase coherence needs to be avoided, and atmospheric delay and opacity can fluctuate relatively quickly at 86 GHz with a coherence timescale typically in the range of 10 to 20 seconds.

We therefore used NRAO 530 as a test source to get a handle on the uncertainties and potential errors in the data. NRAO 530 has been extensively studied with VLBI at different wavelengths (e.g. Lu et al., 2011a,b; An et al., 2013) and is regularly monitored with the VLBA at 43 GHz in the framework of the Boston University Blazar Monitoring Program², providing a good body of background knowledge on the source structure and evolution. For this source, we performed standard hybrid mapping in *Difmap*. Using an iterative self-calibration procedure with progressively decreasing solution intervals, we obtained stable CLEAN images with sidelobes successfully removed. Careful flagging was applied to remove low-S/N and bad data points. Figure 3.1 shows the naturally weighted CLEAN images for both datasets. Table 3.1 includes the corresponding image parameters. The overall source structure is comparable between the two tracks, and the total recovered flux density in both images differs by less than 10%. With ALMA-only flux measurements of NRAO 530, a significantly higher total flux of 2.21 Jy at band 3 (91.5 GHz, ALMA Calibrator database, May 25, 2015) was measured. The difference with the flux we measured from the VLBI observations is likely due to a significant contribution from large-scale structure which is resolved out on VLBI baselines. Because the GBT and the LMT have adaptive dish surfaces, their gain factors can be time-variable. As such their gain curves are not fixed over time, and so additional and more accurate amplitude calibration in *Difmap* was required for baselines to these stations. We began the imaging procedure with an initial source model based on VLBA-only data, which allowed us to obtain further amplitude correction factors for the LMT of 1.47 (BF114A) and 1.14 (BF114B), and for the GBT of 0.54 (both tracks). Gain correction factors for the VLBA stations were of the order of $\lesssim 20\%$.

Due to the gain uncertainty for the GBT and the LMT for the reason mentioned above, amplitude calibration for Sgr A* required a further step beyond the initial propagation of gain solutions from scans on NRAO 530 to scans on Sgr A*. This calibration step was performed by taking the Sgr A* visibility amplitudes from the short baselines between the south-western VLBA stations (KP, FD, PT, OV) and using an initial model fit of a single Gaussian component to these VLBA-only baselines. Due to the low maximum elevation of Sgr A* (it appears at ~ 16 degrees lower elevation

²[urlhttp://www.bu.edu/blazars/VLBAproject.html](http://www.bu.edu/blazars/VLBAproject.html)




FIGURE 3.1: Naturally weighted 86 GHz images of NRAO 530. *Left:* Using data of project BF114A (2015-05-22) with VLBA and LMT. *Right:* Using data of project BF114B (2015-05-23) with VLBA, LMT and GBT. The contours indicate the flux density level (dashed-grey contours are negative), scaled logarithmically and separated by a factor of two, with the lowest level set to the 3σ -noise level. The synthesized array beam is shown as a grey ellipse in the lower left corner. Image parameters are listed in Table 3.1.

than NRAO 530 at transit), the amplitude correction factors for the VLBA are typically larger for Sgr A* than for NRAO 530 but still agree with the factors of the corresponding NRAO 530 observations within $\lesssim 30\%$ (except for the most northern stations BR and NL), comparable to the findings of Lu et al. (2011a). Analogously to the data reduction steps taken for NRAO 530, we used this initial source model to perform additional amplitude calibration for the GBT and the LMT. After this first round of amplitude self-calibration, iterative mapping and self-calibration was performed (see Sect. 3.3.1).

3.3 Results

Following the closure phase analysis in Paper I, we now study the source geometry and size using hybrid imaging (Sect. 3.3.1) and closure amplitudes (Sect. 3.3.2). In Paper I, where we studied the closure phase distribution to look for source asymmetry, we concentrated only on the more sensitive dataset including VLBA+LMT+GBT (project code: BF114B), while in this paper we also include the VLBA+LMT dataset (project code: BF114A).

3.3.1 Mapping and self-calibration of Sgr A*

After amplitude correction factors were applied (as explained in Sect. 3.2), we performed an iterative mapping and self-calibration procedure including careful flagging of the Sgr A* dataset. Amplitude and phase self-calibration were applied using increasingly shorter timesteps and nat-

ural weighting. We deconvolved the image for both datasets by using elliptical Gaussian model components, since the CLEAN algorithm has difficulty fitting the visibilities when it uses point sources. Table 3.2 gives the best-fit parameters from this approach. Figure 3.2 shows both of the resulting images convolved with the clean beam.

As shown by, for example, Bower et al. (2014), when self-calibrating, the derived model can depend on the initial self-calibration model chosen for a single iteration, if the χ^2 -landscape has complex structure. Furthermore, as also noted by Ortiz-León et al. (2016), the resulting uncertainties on the model parameters are often underestimated, if they are based solely on the self-calibration solution. To assess the true errors, the uncertainties on the gain solutions must also be taken into account.

Therefore, we tested the robustness of the final model, in other words, the dependence of the self-calibration steps on input models, described as follows. We evaluated conservative uncertainties on the model parameters of the elliptical Gaussian brightness distribution by using different starting parameters for the iterative self-calibration procedure, where all starting model parameters were individually varied by up to 30%, to check the convergence on the same solution. We generated 1000 random starting models to perform the initial amplitude self-calibration (Sect. 3.2). The starting model always consists of an elliptical Gaussian brightness distribution. Each of its parameters (flux, major axis, axial ratio, position angle) was drawn from a normal distribution around the initial model. Using these input models, iterative self-calibration steps were applied and the resulting distribution of the model parameter was examined. For an illustration of the observed distribution of the major axis size, please see Figure 3.3.

As expected, we find a strong correlation between input model flux density and final flux density of the Gaussian model components. Therefore, to constrain the flux of Sgr A*, we primarily used the fluxes on short VLBA baselines as explained in Sect. 3.2. For both NRAO 530 and Sgr A*, we find less than 10% total flux density difference between our two consecutive epochs.

We find that the model converges onto values for major axis and axial ratio (or alternatively minor axis) that show a spread of about 10%. The position angle uncertainty is constrained to $\pm 20^\circ$. We note that this analysis shows that the distribution for the major axis in BF114B is skewed, having an average of $222 \mu\text{as}$, a median of $215 \mu\text{as}$ and a mode of $205 \mu\text{as}$. The resulting major axis distribution also has a hard lower bound at $\sim 200 \mu\text{as}$. This skewed distribution of parameters from selfcal suggests that there are multiple local minima in the χ^2 -landscape that produce different model parameters from different iterations, and is therefore of limited value in determining source size uncertainties. We have therefore used closure amplitude analysis to verify this estimate of the source size and provide more accurate uncertainties, and this process is described in the next section.




FIGURE 3.2: Results of hybrid mapping of Sgr A* at 3 mm. *Top left*: Beam-convolved image from the dataset of project BF114A (2015-05-22) using VLBA and LMT. *Top right*: Beam-convolved image from the dataset of project BF114B (2015-05-23) using VLBA, LMT and GBT. The contours indicate the flux density level (dashed-grey contours are negative), scaled logarithmically and separated by a factor of two, with the lowest level set to the 3σ -noise level. *Bottom left*: Residual map of Sgr A* after primary component subtraction from the BF114A dataset, using natural weighting. No clear pattern is seen in the residual image. *Bottom right*: Natural-weighted residual map for Sgr A*, epoch B, after subtraction of the best-fitting 2D Gaussian source component. The remaining excess flux towards the east is highly concentrated and clearly present. Both residual images use a cross to indicate the centre of the primary (subtracted) component on the sky. Image and model parameters are listed in Tables 3.1 and 3.2, respectively.




FIGURE 3.3: Distribution of major axis sizes arising from 1000 selfcal runs in which each initial model parameter was varied according to a Gaussian distribution with a width of 30% of the nominal parameter value. The resulting distribution of sizes shows a clear skew, with most results clustering close to a minimum cutoff value of $200 \mu\text{as}$. The coloured lines indicate the mean (green), the median (cyan) and the statistical $1-\sigma$ errors (red).

We find that for BF114A (VLBA+LMT), one single Gaussian component is sufficient to model the data (see Figure 3.2, bottom left). For the BF114B dataset with higher sensitivity due to the inclusion of the GBT, the model fitting with one Gaussian component shows a significant excess of flux towards the south-west in the residual map (see Figure 3.2, bottom right). Modelling this feature with a circular Gaussian component yields a flux density excess of $\sim 10 \text{ mJy}$ (i.e. approximately 1% of the total flux) at $\Delta\text{RA} \sim 0.23 \text{ mas}$, $\Delta\text{DEC} \sim -0.05 \text{ mas}$ from the phase centre. Including this second component in the model fit, results in a smooth residual map (with $\text{RMS} \sim 0.5 \text{ mJy}$). We checked the reliability of this feature using the same method as described above, where a range of initial model parameters was used as input for a selfcalibration step that resulted in a distribution of best-fit model parameters. We find that the position of the residual emission is well-constrained and independent from the self-calibration starting parameters. The BF114A dataset, however, does not show such clear and unambiguous residual emission.

We have tested the compatibility of the BF114A dataset with the source model we find for BF114B. Subtracting the full two-component BF114B source model from the calibrated BF114A data and looking at the residual map, we see an enhanced overall noise level and no clear evidence of missing flux at the position of the secondary component. We further performed a separate

amplitude and phase selfcalibration of the BF114A data using the BF114B source model, and inspected the residual map after subtraction of only the main source component of the BF114B model. In this residual map, we do see an enhancement of flux density at the position of the secondary component, but it is not as strong as the secondary component of the source model (~ 5 mJy versus 10 mJy for the model). We also see apparent flux density enhancements of similar strength at other positions close to the phase centre. We therefore conclude that the BF114A (u, v)-coverage and sensitivity are not sufficient to provide a clear measurement of the secondary source component as seen for the BF114B epoch. Given that the detectability of the secondary component is so marginal for BF114A, we cannot determine whether the asymmetry we see in the BF114B epoch is a feature which persisted over the two epochs or a transient feature that was not present in the earlier epoch.

We emphasize that the asymmetric feature we see in the Sgr A* emission when imaging BF114B was already suggested by our analysis of the closure phases of the BF114B dataset (Paper I). We find that a model consisting of two point sources results in a significantly better fit to the closure phases, with the weaker component being located east of the primary. However, the flux ratio of the two components was left poorly constrained, resulting in χ^2 minima at flux ratios of 0.03, 0.11, and 0.70. In the current analysis, by using the full visibility data and fitting Gaussian components instead of point sources, we can constrain the flux ratio to ~ 0.01 . The low flux density of this secondary source component compared to the main source component makes it difficult to detect this source feature upon direct inspection of the visibility amplitudes as a function of baseline length. However, with model fitting it becomes clear that a single Gaussian component systematically underfits the amplitude trends of the data. We have thus seen evidence for this component independently in both the closure phases (Paper I) and the visibility amplitudes (this work).

It remains unclear whether this sub-structure in the 3 mm emission of Sgr A* is intrinsic or induced by refractive scattering. On long baselines, refractive scattering can introduce small-scale sub-structure in the ensemble-averaged image (Johnson & Gwinn, 2015). This effect strongly depends on the intrinsic source size and geometry. A larger source size will show smaller geometrical aberration from scattering compared to a point source, as different parts of the source image are refracted in independent ways that tend to partially cancel out any changes in overall structure. At $\lambda \sim 5$ mm where the intrinsic source size of Sgr A* becomes comparable to the angular broadening, this effect is most distinct (Johnson & Gwinn, 2015). Gwinn et al. (2014) reported on the detection of scattering sub-structure in the 1.3 cm emission of Sgr A*. Assuming a Kolmogorov spectrum of the turbulence, the authors expect refractive scintillation to lead to the flux density measured on a 3000 km east-west baseline to vary with an RMS of 10-15 mJy. Similarly, Ortiz-León et al. (2016) show that refractive effects can cause sub-structure in 3 mm images, with a RMS flux modulation of 6.6% and an evolution timescale of about two weeks. Taking these considerations into account, this sub-structure detected at long baselines in our 3 mm datasets would be consistent with scattering noise. However, given the more significant detection

TABLE 3.1: Image & observational parameters (natural weighting)

Date (Project ID) yyyy-mm-dd	Source	Array configuration ^a	Beam [mas]	Beam PA	RMS [Jy/beam]
2015-05-22 (BF114A)	NRAO 530	VLBA+LMT	0.107×0.204	3.0°	0.0004
2015-05-23 (BF114B)	NRAO 530	VLBA+LMT+GBT	0.100×0.225	−6.3°	0.0003
2015-05-22 (BF114A)	Sgr A*	VLBA+LMT	0.541×0.165	38.5°	0.0010
2015-05-23 (BF114B)	Sgr A*	VLBA+LMT+GBT	0.147×0.286	6.4°	0.0005

^aFor VLBA: Brewster (BR), Fort Davis (FD), Kitt Peak (KP), Los Alamos (LA), Mauna Kea (MK), North Liberty (NL), OVRO (OV) and Pie Town (PT). We note that for Sgr A* no fringes were detected to MK, which results in a larger beam size for Sgr A* than for NRAO 530.

TABLE 3.2: Parameters of model components from self-calibration^a

Date (Project ID)	S [Jy]	b_{maj} [μas]	ratio (-)	b_{min} [μas]	PA [deg]
2015-05-22 (BF114A)	1.02 ± 0.1	227.0	0.85	193.0	56.4
2015-05-23 (BF114B)	0.95 ± 0.1	215.3	0.77	165.8	76.5

^aNote: major/minor axis uncertainties are of the order of 10%. The PA is constrained to within 15° (BF114A) and 12° (BF114B). See Sect. 3.3.1 for more details.

in the dataset involving the GBT and LMT, a contribution of intrinsic sub-structure cannot be excluded. We discuss further implications in Sect. 3.5.

3.3.2 Constraining the size of Sgr A* using closure amplitudes

Closure quantities are robust interferometric observables which are not affected by any station-based error such as noise due to weather, atmosphere or receiver performance. As one example of a closure quantity, the closure phase is defined as the sum of visibility phases around a closed loop, that is, at least a triangle of stations. We discuss the closure phase analysis of the Sgr A* dataset BF114B in Paper I in detail. Here, instead of closure phases, we focus on the closure amplitude analysis of both datasets. The closure amplitude is defined as $|V_{ij}V_{kl}|/|V_{ik}V_{jl}|$, for a quadrangle of stations i, j, k, l and with V_{ij} denoting the complex visibility on the baseline between stations i and j . Using measurements of this quantity, one can determine the source size independently from self-calibration, as shown in various previous publications for 3 mm VLBI observations of Sgr A* (Doeleman et al., 2001; Bower et al., 2004; Shen et al., 2005; Bower et al., 2014; Ortiz-León et al., 2016).

In the context of this work, we are interested in a way to establish the observed size and orientation of Sgr A* separately from self-cal. We therefore fitted a simple model of an elliptical Gaussian component to the closure amplitude data, and we deconvolved the scattering ellipse using the best available model (Bower et al. 2006, 2014b) afterwards. We performed a χ^2 -analysis in fitting the Gaussian parameters (major and minor axis, and position angle).

For both datasets, BF114A and BF114B, we derived the closure amplitudes from the 10s-averaged visibilities and fitted a simple 2D Gaussian source model to the closure amplitude data. There are some subtleties to take into account when modelfitting with closure amplitudes. χ^2 -minimisation algorithms for model fitting generally assume that the errors on the measurements used are Gaussian. Closure amplitudes, when derived from visibilities with Gaussian errors, in general have non-Gaussian errors that introduce a potential bias when modelfitting which depends on the S/N and the relative amplitudes of the visibility measurements involved: because closure amplitudes are formed from a non-linear combination of visibility amplitudes (by multiplications and divisions), their error distribution is skewed (asymmetric). This is especially a problem in the low-S/N regime - the skew is much less pronounced for higher S/N values, and closure amplitude errors tend toward a Gaussian distribution in the high-S/N limit. Taking the logarithm of the measured closure amplitude values and appropriately defining the measurement uncertainties symmetrizes these errors, and generally results in more stable fitting results (Chael et al., 2018). For this reason, here we have adopted the technique described in that paper.

The workflow we have adopted for the closure amplitude model fitting pipeline is outlined in Figure 3.4. We give a brief summary of the process here, and specify more details on individual steps below. We started the process with the frequency-averaged visibility dataset output from *AIPS*, in which aberrant visibilities have already been flagged. We time-averaged this dataset to ten-second length segments using *Difmap* to improve S/N per visibility measurement. In this step, the uncertainties on the resulting visibilities are recalculated using the scatter within each averaging period. The time interval of ten seconds was experimentally confirmed to yield vector-averaged visibility amplitudes that are not significantly lower than when averaging over shorter timescales, and as such falls within the coherence timescale of the atmosphere at 86 GHz. We also de-biased the averaged visibilities here, according to expression 9 in Chael et al. (2018). We applied an S/N cutoff to the averaged ten-second visibility amplitudes at this point, where we have used different values for this cutoff to test the robustness of the model fitting results (described below). Using the remaining visibilities, we calculated the closure amplitudes for each ten-second time interval in the dataset. We calculated the error on these closure amplitude measurements using standard error propagation (following expression 12 from Chael et al. (2018)), and we then made another cut in the dataset where we discarded all measurements that have a reported S/N below our threshold value. Lastly, we applied our station selection to the resulting dataset, dropping all closure amplitude measurements in which the omitted stations are involved. We thus obtained the dataset on which we performed model fitting.

We used bootstrapping of the closure amplitudes of each dataset to determine the error on the individual fit parameters. Bootstrapping works by forming a new realisation of measurement data by picking measurements from the original dataset at random (with replacement) until a new dataset is formed that has an equal number of measurements as the original dataset. As such,




FIGURE 3.4: Overview of the pipeline used for closure amplitude model fitting. The stages involving time averaging, Visibility S/N filtering, Closure amplitude S/N filtering, station selection, and bootstrapping all offer different choices as to the parameters involved.

any measurement from the original dataset may be represented either once, multiple times or not at all in the newly formed dataset – the weights of measurements in the original dataset are thus stochastically varied, emulating the drawing of a new sample of measurements. We fitted the data with a 2D Gaussian model with three free parameters: major axis size, minor axis size and position angle on the sky of the major axis. The χ^2 minimisation is done as per expression 21 in Chael et al. (2018).

Besides bootstrapping, we explore the effects of different values chosen for the S/N cutoff of the visibility amplitudes used in the model fitting. Visibilities with a low reported S/N are expected to have a larger influence on the skewness of the closure amplitude distribution, and are thus likely to introduce a bias in the fitting results. This effect is investigated by looking at different cutoff values for the visibility S/Ns. All visibility measurements can be assigned a 'reported S/N', which is defined as the measured visibility amplitude divided by the visibility amplitude uncertainty as determined from scatter among the measurements over a ten-second integration period. Before forming closure amplitudes using a visibility dataset, this visibility dataset is filtered by only admitting measurements that have reported S/Ns above a chosen threshold value. The constructed closure amplitudes can then be filtered again by their reported S/N. A closure amplitude S/N cutoff value of three was employed to avoid the larger bias that comes with low-S/N measurements, although we found that varying this value did not significantly impact the fitting results. The variation of visibility S/N cutoff has a more pronounced influence on fitting results, and this effect is shown in Figure 3.5. The plots in the top row of this figure show the model fitting results for the full dataset, with all stations included. In these plots, where the blue circles indicate fitting results from the measured data, we see that the fitted model parameters show relatively minor variation over a range of S/N cutoff values from one to four, where the minor axis size is the parameter that shows the largest spread. Above visibility S/N cutoff values of four, we see that the

spread in the fitting results grows and that trends of fitted values with S/N cutoff start appearing. This effect is coupled to the fact that only a limited number of quadrangles are left at these high S/N cutoff values, which by themselves provide weaker constraints on source geometry because of the limited (u, v) -coverage they provide.

To investigate the consistency of the data regarding the convergence of best-fit model parameters, we also have performed model fits where we excluded the GBT from the array before gathering closure amplitude measurements and model-fitting. This was done to check if the inclusion of the GBT resulted in a systematic offset of fitted model parameters versus the case where the array does not include the GBT. Inclusion of the GBT offers a much better east-west array resolution, which is expected to have an impact on the quality of the major axis size estimate as the observed Sgr A* Gaussian is orientated almost east-west on the sky. Likewise, the LMT offers a significant enhancement of the north-south array resolution and should therefore yield a clear improvement in quality for the estimated minor axis size. The model fitting results for these cases are included in Figure 3.5, in the second (no GBT) and third (no LMT) rows. It is clear that indeed, inclusion of the GBT improves the quality of the major axis size estimate (the scatter among different bootstrapping realisations is significantly smaller than for the case where the GBT is omitted), while the LMT is instrumental in obtaining a good estimate for the minor axis size. As a result, the accuracy with which the position angle is determined benefits from inclusion of both the GBT and the LMT.

We should note that consistency of fitted model parameters by itself does not guarantee accurate results (only precise results). For this reason, we have generated synthetic visibility datasets with the same (u, v) -sampling as the original measurements, where a Gaussian source model with fiducial parameter values that are close to the previously measured size of Sgr A* (Major axis: $210.4 \mu\text{as}$, minor axis: $145.2 \mu\text{as}$, position angle: 80 degrees east of north) was used as input. The visibility uncertainties for this synthetic dataset were scaled in such a way as to yield the same distribution in S/N values as the original data shows. For this synthetic dataset, the full processing pipeline was then used and the deviations of the fitted parameters from the fiducial inputs were inspected. These results are also plotted in Figure 3.5, using red triangles as markers for the model fitting results and black lines to indicate the input model parameter values. For the major axis size, we see that the fitted values typically underpredict the actual source size by 5 to $10 \mu\text{as}$, depending on which stations are involved in the array. The minor axis size is severely underpredicted when the LMT is left out of the array, but is close to the input value when the LMT is included. The position angle come out close to the input value in all cases, although there is a small positive bias seen in the case where the full array is used. We note that the y-axis ranges of these plots are different, and that the spread seen in the case of the full array are typically much smaller than those for the other array configurations. These results from synthetic data fitting allow us to correct for the biases that our pipeline exhibits. The bias-corrected fitted source parameters are shown in Figure 3.6. For all model parameters we get consistent fitting results for all visibility S/N cutoff choices up to five. Because any specific choice of S/N cutoff




FIGURE 3.5: Raw model fitting results for the BF114B dataset and for the synthetic dataset with the same (u, v) -sampling, using different integral S/N cutoff values and different station selections. The fiducial model parameters used to generate the synthetic dataset with are indicated by the horizontal black lines. For each S/N cutoff value, 31 bootstrapping realisations were performed to obtain uncertainties on the fitted model parameter values. Each of the results from these realisations is plotted with a single symbol. The different columns of figures show, from left to right, the major axis, minor axis and position angle results respectively. *Top row*: Full array, *Middle row*: Without the GBT, *Bottom row*: Without the LMT.

value is difficult to defend for coming up with our final model parameter fitting values, we note that the scatter of the fitted values among these different visibility S/N cutoffs is consistent with their uncertainties in most cases. We therefore used the average value for the model fit results up to and including the S/N cutoff of five, and for the uncertainty we use the average uncertainty for the same data points. Our derived source geometry parameters are listed in Table 3.3, together with previously reported sizes.




FIGURE 3.6: Bias-corrected model fitting results for the BF114B dataset for different station selections as a function of visibility S/N cutoff value. The fitted parameter values for the measured data have been corrected using the offset exhibited by the fits to the synthetic datasets. The results per station selection (symbol type) have been offset along the S/N axis by a small amount for clarity.

3.4 Constraints on the size-frequency relation and the scattering law

Extensive measurements of the size of Sgr A* have been performed over the years at various frequencies, leading to an understanding of the nature of the scattering law in the direction of the Galactic centre (Backer, 1978; Lo et al., 1998; Bower et al., 2006; Johnson & Gwinn, 2015; Psaltis et al., 2015) as well as on the dependency of intrinsic source size on frequency both from an observational and a theoretical perspective (Bower et al., 2004, 2006; Shen, 2006; Bower et al., 2014; Mościbrodzka et al., 2014; Ortiz-León et al., 2016). Knowledge of the intrinsic source size at different frequencies is an important component of the research on Sgr A*, because it strongly constrains possible models for electron temperatures, jet activity and particle acceleration.

Our size measurements of Sgr A* at 86 GHz, when combined with these previously published size measurements over a range of frequencies, allowed us to perform a simultaneous fitting of the

TABLE 3.3: Sgr A*: size of elliptical Gaussian fits to observed 86 GHz emission

Reference	Major axis [μas]	Minor axis [μas]	Position angle [$^\circ$]	Axial ratio [-]	Intrinsic size [μas]
Ortiz-León et al. (2016) (obs. 1, self-cal)	212.7 ± 2.3	138.5 ± 3.5	81.1 ± 1.8	1.54 ± 0.04	142 ± 9 $114 \pm 15^\times$
Ortiz-León et al. (2016) (obs. 2, self-cal)	221.7 ± 3.6	145.6 ± 4.0	75.2 ± 2.5	1.52 ± 0.05	155 ± 9 $122 \pm 14^\times$
Lu et al. (2011a) (self-cal)	210 ± 10	130 ± 10	83.2 ± 1.5	1.62	139 ± 17 $102 \pm 21^\times$
(Shen et al., 2005) (cl. amp.)	210^{+20}_{-10}	130^{+50}_{-13}	79^{+12}_{-33}	1.62	
(Doeleman et al., 2001) (self-cal, averaged)	180 ± 20	—	—	—	
(Krichbaum et al., 1998) (model fit)	190 ± 30	—	—	—	
This work (self-cal)	217 ± 22	165 ± 17	77 ± 15	1.3	167 ± 22 $122 \pm 25^\times$ ^a
This work (cl. amp., full array)	215.1 ± 0.4	145.1 ± 1.5	77.9 ± 0.4	1.48 ± 0.01	145.4 ± 0.6 $122.6 \pm 1.7^\times$
This work (cl. amp., no GBT)	213.9 ± 2.5	148.0 ± 4.0	77.9 ± 3.0	1.45 ± 0.03	144.4 ± 3.7 $125.2 \pm 4.9^\times$
This work (cl. amp., no LMT)	210.6 ± 1.0	88.7 ± 34.2	86.4 ± 1.2	2.37 ± 1.02	86.5 ± 69.7 $40.6 \pm 40.5^\times$

^aCalculated using a scattering kernel size of $158.5 \times 77.5 \mu\text{as}$ at 86 GHz, from Bower et al. (2006). No uncertainty in scattering kernel size was incorporated in this calculation. Our closure amplitude results below use the same scattering kernel.

size-frequency relation together with the scattering law. Previous studies focus on constraining either the scattering law or the intrinsic size-frequency relation, typically by either focusing on a specific range of longer observing wavelengths to constrain the scattering law (Psaltis et al., 2015) or by using a fiducial scattering law and focusing on the shorter observing wavelengths to establish an intrinsic size-frequency relation (Bower et al., 2006). However, simultaneous fits of both of these relations to the available data have not been published to date. Johnson et al. (2018) find a size $b_{maj} = 1.380 \pm 0.013 \left(\frac{\lambda}{\text{cm}}\right)^2$ milliarcseconds using a similar set of past results and analysis techniques as used in this work. The difference with our constraint emphasizes the challenge of obtaining a solution with 1% precision in the complex domain of heterogenous data sets, extended source structure, and an unknown intrinsic size.

Besides our own measurements presented in this paper, we used previously published size measurements from Bower & Backer (1998); Krichbaum et al. (1998); Bower et al. (2004, 2006); Shen (2006); Doeleman et al. (2008); Bower et al. (2014); Ortiz-León et al. (2016), where Bower et al. (2004) includes re-analysed measurements originally published in Lo et al. (1998). Care was taken to ensure that all these published results were derived from data that was independently

obtained and analysed. The measurements we include for the model fitting have been taken over a time period of multiple decades, thereby most likely representing different states of activity of the source which may affect size measurements. This effect is expected to be small, however: at short wavelengths because of the stable source size that has been measured over time, and at longer wavelengths because the scattering size is so much larger than the intrinsic size. The measurements taken at wavelengths close to $\lambda = 20$ cm were taken closely spaced in time, yet still show a mutual scatter that is wider than the size of their error bars suggests: this may indicate the presence of systematics in the data. An ongoing re-analysis of these sizes at long wavelengths (Johnson et al., 2018) suggests that these measurements are too small by up to 10%, likely impacting the resulting fits for the scattering law and intrinsic size-frequency relation. Here, we have used the values as they have been published. Throughout this section, we use Gaussian models for both the observed source size and for the scattering kernel. Recent work has shown that the instantaneous shape of the scattering kernel deviates from a Gaussian to a limited extent (Gwinn et al., 2014), but the statistical average of the scattering kernel geometry is thought to be Gaussian to within a few per cent.

The set of measurements, as we have used them in the model fitting, can be seen in Figure 3.7. Measurements taken at the highest of these frequencies (230 GHz) are expected to feature emission coming from very close to the black hole shadow, and as such the perceived source size may be significantly affected by gravitational lensing effects where the source image can be warped into a crescent-like structure. Such strong lensing effects are not expected to play a role in source sizes as observed at lower frequencies because the inner accretion flow is optically thick at small radii for those frequencies. We thus expect to effectively see emission coming from somewhat larger radii where the light paths are not significantly affected by spacetime curvature but are affected by interstellar scattering along our line of sight. Therefore, we performed the model fitting both including the 230 GHz size measurements (Figure 3.7) and excluding them, to see if the expected GR lensing effects play a significant role in the appearance of the source at the shortest wavelengths. We find very little difference in the best-fit parameter values between the results.

Simultaneous fitting of the size-frequency relation and the scattering law is done using the major axis size measurements only, as the uncertainties in the minor axis size measurements are too large to provide any meaningful constraint on the models. For the size-frequency relation, we used the following expression:

$$\theta_{\text{int}}(\lambda) = a \cdot \lambda^b, \quad (3.1)$$

where a and b are constants to be determined, θ_{int} is the intrinsic angular size in milliarcseconds and λ is the observing wavelength in cm. For the scattering law we adopt the expression:

$$\theta_{\text{scatt}}(\lambda) = c \cdot \lambda^2, \quad (3.2)$$

where c is a constant to be determined and with θ_{scatt} the angular broadening through scattering in milliarcseconds. These sizes are added in quadrature to provide the measured major axis size




FIGURE 3.7: Left: Aggregate measurement data for the observed major axis size of Sgr A* (black points with error bars), and model fitting results for different combinations of included model components (coloured lines). The highest-quality fits are provided by the green, blue and orange lines (the top 3 listed in the legend) which provide very similar fit qualities (see Table 3.4). Right: the same data, plotted with the major axis sizes divided by wavelength squared. The fitting results without the 230 GHz data are almost identical to these, and hence are not plotted separately.

for Sgr A*:

$$\theta_{\text{meas}}(\lambda) = \sqrt{\theta_{\text{int}}^2 + \theta_{\text{scatt}}^2}. \quad (3.3)$$

This expression is used in the fitting procedure to obtain a measured size from the model parameters, thus involving at most three free parameters (the constants a , b , and c). Using a simple linear least-squares fitting procedure (from the Python package `scipy.optimize.curve_fit`), and fitting to all size measurement data available, we get the following values and uncertainties in the expressions for intrinsic size and scattering size respectively (see also Figure 3.7 for the model curves produced):

$$\theta_{\text{int}}(\lambda) = 0.502 \pm 0.075 \cdot \lambda^{1.201 \pm 0.138}, \quad (3.4)$$

$$\theta_{\text{scatt}}(\lambda) = 1.338 \pm 0.012 \cdot \lambda^2. \quad (3.5)$$

At 230 GHz, there is the possibility that the size of Sgr A* may be strongly affected by gravitational lensing. To investigate whether the inclusion of these measurements significantly affects the size-wavelength relation found, we also performed the fitting routine while leaving out the 230 GHz measurements. We then obtained the following expressions for intrinsic size and scattering size:

$$\theta_{\text{int}}(\lambda) = 0.502 \pm 0.078 \cdot \lambda^{1.201 \pm 0.143}, \quad (3.6)$$

$$\theta_{\text{scatt}}(\lambda) = 1.338 \pm 0.012 \cdot \lambda^2. \quad (3.7)$$

Cross-comparing expressions 3.6 and 3.7 to 3.4 and 3.5, we see that the corresponding fitted model parameters between the model fits with and without the 230 GHz measurements are well within each other's error bars for all three model parameters. The available measurements of

source size at 1.3 mm thus seem to be compatible with the source size as predicted using the fitted size-wavelengths relations from the other measurements.

Comparing these figures to Bower et al. (2015a), we see that the scattering size parameter for the major axis is well within the error bars of the value calculated in that work ($b_{\text{maj, scatt}} = 1.32 \pm 0.02 \text{ mas cm}^{-2}$). For the intrinsic size as a function of wavelength, the powerlaw index we find is somewhat larger than the powerlaw index calculated in Ortiz-León et al. (2016) (where it is quoted as being 1.34 ± 0.13), but still within the error bars.

The size-wavelength relation that we have used up to this point has a specific functional form: it consists of a pure powerlaw for the size-frequency relation, combined in quadrature with a scattering law where scattering size scales with wavelength squared. To explore the influence that this choice of functional form has on the results of the fitting procedure, we have performed the fit with other models for the dependence of observed size on observing wavelength as well. All models consist of a combination of three components: a fixed-size component that is constant across all wavelengths, a scaled λ^p component (where p is a free parameter) that is added linearly to it, and a scaled λ^2 component (scattering law) that is then added to the sum of the other component(s) in quadrature. Six combinations of these model components were fitted to the major axis size measurement data, and each fit was done for two cases: with and without the 230 GHz observed source sizes included in the data to be fitted to. In Table 3.4, the results of these model parameter fits are presented.

The three best-fitting models are the 'regular' model (scattering law + general power law), the 'augmented' model (scattering law + general power law + fixed size offset), and the 'simple' model (general power law + fixed size offset). For the simple model, the best-fitting power law index is close to two within a few per cent. If the power law exponent from scattering can deviate from the theoretically ideal value of two by even a small fraction, this result suggests that the intrinsic size-frequency relation for Sgr A* is less certain than what has been found in previous publications. A similar conclusion was derived by Bower et al. (2006), where it was found that a relaxation of the scattering exponent to values slightly different from two undercuts the support for an intrinsic size-frequency relation with a non-zero power law index.

3.5 Summary and conclusions

Constraining the intrinsic size and structure of Sgr A* at an observing wavelength of 3 mm still remains a challenge. Although the effect of interstellar scattering becomes smaller at this wavelength, it is still not negligible. GRMHD models of the accretion flow around Sgr A* (e.g. Mościbrodzka et al., 2014) predict a certain structure in the emission which should be detectable with current VLBI arrays. However, detection of intrinsic sub-structure could be hindered by re-

TABLE 3.4: Sgr A*: fitted size dependence on frequency, different models

Model	Incl. 230 GHz?	a	b	c	d	χ^2 / d.o.f.
Size-freq + scattering $\sqrt{(a\lambda^2)^2 + (b\lambda^c)^2}$	yes	1.338 ± 0.012	0.502 ± 0.075	1.201 ± 0.138	-	1146.29 / 34
Size-freq + scattering + offset $\sqrt{(a\lambda^2)^2 + (b\lambda^c + d)^2}$	yes	1.277 ± 0.110	0.600 ± 0.205	1.757 ± 0.320	0.055 ± 0.021	1107.79 / 33
Scattering + offset $\sqrt{(a\lambda^2)^2 + d^2}$	yes	1.360 ± 0.009	-	-	0.139 ± 0.005	1873.88 / 35
Size-freq + offset $b\lambda^c + d$	yes	-	1.385 ± 0.019	1.980 ± 0.010	0.044 ± 0.005	1108.29 / 34
Size-freq only $b\lambda^c$	yes	-	1.537 ± 0.015	1.905 ± 0.008	-	3292.23 / 35
Scattering only $a\lambda^2$	yes	1.417 ± 0.024	-	-	-	15944.28 / 36
$\sqrt{(a\lambda^2)^2 + (b\lambda^c)^2}$	no	1.338 ± 0.012	0.502 ± 0.078	1.201 ± 0.143	-	1145.25 / 32
$\sqrt{(a\lambda^2)^2 + (b\lambda^c + d)^2}$	no	1.273 ± 0.128	0.606 ± 0.235	1.773 ± 0.337	0.057 ± 0.021	1102.07 / 31
$\sqrt{(a\lambda^2)^2 + d^2}$	no	1.360 ± 0.009	-	-	0.139 ± 0.005	1824.55 / 33
$b\lambda^c + d$	no	-	1.385 ± 0.020	1.980 ± 0.010	0.044 ± 0.005	1104.63 / 32
$b\lambda^c$	no	-	1.537 ± 0.015	1.905 ± 0.008	-	3290.05 / 33
$a\lambda^2$	no	1.417 ± 0.025	-	-	-	15940.55 / 34

fractive scattering, possibly itself introducing compact emission sub-structure (Johnson & Gwinn, 2015).

In this paper, we present imaging results and analysis of closure amplitudes of new VLBI observations performed with the VLBA, the LMT and the GBT at 86 GHz. Following our previous result (Paper I) from the analysis of closure phases, the detection of sub-structure in the 3 mm emission of Sgr A*, we confirm the previous result of compact sub-structure using imaging techniques. Using NRAO 530 as test source, we show that VLBI amplitude calibration can be performed with an absolute uncertainty of 20% for NRAO 530 and 30% for Sgr A*, where we are currently limited by the uncertainty in antenna gains. The variable component of these gain uncertainties is limited to $\sim 10\%$.

Out of our two experiments, only in the higher resolution and more sensitive experiment (BF114B, including the VLBA, the LMT and the GBT) is the compact asymmetric emission clearly detected. The VLBA+LMT dataset (BF114A) remains inconclusive in this respect. The asymmetry is detected as significant residual emission, when modelling the emission with an elliptical Gaussian component. The flux density of the asymmetrical component is about 10 mJy. Such a feature can be explained by refractive scattering, which is expected to result in an RMS flux of this level, but an intrinsic origin cannot be excluded. The discrimination and disentanglement of both these possible origins requires a series of high-resolution and multi-frequency VLBI observations, spread out in time. Interestingly, the secondary off-core component observed at 7 mm with the VLBA (Rauch et al., 2016) is found at a similar position angle. The authors of that paper interpret this feature as an adiabatically expanding jet feature. Future, preferably simultaneous, 3 and 7 mm VLBI observations can shed light on the specific nature of the compact emission. A persistent

asymmetry, observed over multiple epochs that are spaced apart in time by more than the scattering timescale at 86 GHz, would provide strong evidence for an intrinsic source asymmetry. Another way in which observed asymmetry may be ascribed to source behaviour rather than scattering is when a transient asymmetry evolution is accompanied by a correlated variation in integrated source flux density. Observations of that nature will require successive epochs using a consistent and long-baseline array of stations involved accompanied by independent high-quality integrated flux density measurements (e.g. by ALMA).

We see that the combination of the VLBA, LMT and GBT provides the capability to pin down the observed source geometry with unsurpassed precision because of the combination of sensitivity and extensive (u, v) -coverage provided, going beyond what addition of the LMT or the GBT separately can do. This combination of facilities is therefore important to involve in future observations that aim to measure the geometry of Sgr A*.

We also note that even with this extended array, the measurement and characterisation of complex source structure beyond a 2D Gaussian source model is something that remains difficult. To study Sgr A* source sub-structure at 86 GHz more closely, be it either intrinsic or from scattering, even more extensive (u, v) -coverage and sensitivity will be needed. Recent measurements carried out with GMVA + ALMA, the analysis of which is underway, should allow for a more advanced study of the complex source structure of Sgr A*, as that array configuration provides unprecedented north-south (u, v) -coverage combined with high sensitivity on those long baselines.

Moving from source sub-structure to overall geometry, this work has reported the observed source geometry of Sgr A* with the highest accuracy to date. Addition of the GBT adds east-west resolving power as well as extra sensitivity and redundancy in terms of measured visibilities. We note that the source geometry we find is very similar to that reported in (Ortiz-León et al., 2016), while the different observations were spaced almost one month apart (April 27th for BD183C, May 23rd for BF114B). Barring an unlikely coincidence, this suggests a source geometry that is stable to within just a few per cent over that time scale. At 86 GHz, Sgr A* is known to exhibit variability in amplitude at the $\sim 10\%$ level (see Paper I) on intra-day timescales. Whether these short-timescale variations in flux density correspond to variations in source size is an open question that can only be resolved when dense (u, v) -coverage is available at high sensitivity (beyond current capabilities), as source size would need to be accurately measured multiple times within a single epoch. Alternatively, studies of the source size variability at somewhat longer timescales can simply be done by observing Sgr A* over multiple epochs – but the fast variations will be smeared out as a result.

From the simultaneous fitting of the scattering law and the intrinsic size-frequency relation for Sgr A*, we find values compatible with existing published results. However, if the scattering law is allowed to deviate from a pure λ^2 law towards even a slightly different power law index,

differing by for example 2% from the value two, support for the published intrinsic size/frequency relation often used in the literature quickly disappears. We therefore advocate a cautious stance towards the weight given to existing models for the intrinsic size-frequency relation for Sgr A*.

Acknowledgements

We thank the anonymous referee for providing comments that improved the quality of the paper. This work is supported by the ERC Synergy Grant “BlackHoleCam: Imaging the Event Horizon of Black Holes”, Grant 610058, Goddi et al. (2017). L.L. acknowledges the financial support of DGAPA, UNAM (project IN112417), and CONACyT, México. S.D. acknowledges support from National Science Foundation grants AST-1310896, AST-1337663 and AST-1440254. G.N.O.-L. acknowledges support from the von Humboldt Stiftung. C.B. wishes to thank Michael Johnson and Lindy Blackburn for valuable discussions which improved the robustness of the closure amplitude analysis.

3.A Closure amplitude model fitting technique

In the closure amplitude model fitting algorithm, we selected at random two independent closure amplitudes out of six possible ones for each quadrangle and integration time to be used in the model fitting procedure. We performed the model fitting of the independent closure amplitudes by using a gradient descent method, where the source model parameters were iteratively altered to give successively better (lower) χ^2 -scores until convergence is reached. The 2D Gaussian model we employed has three free parameters: major axis size (FWHM), minor axis size (FWHM), and the position angle on the sky of the major axis. For every bootstrapping realisation, a random point in the 3D model parameter space was initially chosen as a starting point, from a flat distribution using upper limits for the major and minor axes sizes of $400 \mu\text{as}$ (and lower limits of $0 \mu\text{as}$) to ensure rapid convergence. Initial coarse step sizes are $50 \mu\text{as}$ for both major and minor axes, and 0.1 radians for the position angle. For the parameter starting point, as well as for its neighbours along all dimensions (each one step size removed from the initial point along one parameter axis), the χ^2 scores were calculated and the lowest-scoring point in the resulting set is taken as the starting point for the next iteration. This sequence of steps was repeated until the best-fitting model parameters coincided with the starting point for that iteration (indicating a local optimum has been reached at that parameter resolution), after which the step sizes for all parameters were reduced and the algorithm continues until the minimum step sizes are all reached. To verify that the general nature of the χ^2 landscape is conducive to this iterative method, and to ensure that the algorithm would not get stuck in a local optimum rather than the global optimum, we have mapped out the χ^2 scores over the full 3D parameter space at a low resolution for the original full set of closure amplitudes. This investigation suggested that the χ^2 -score varies smoothly over the full parameter space, revealing the presence of a single global optimum.

FREQUENCY-DEPENDENT TIME LAGS IN SGR A*

Christiaan Brinkerink, Heino Falcke, Casey Law, Denis Barkats, Geoffrey Bower, Andreas Brunthaler, Charles Gammie, Violette Impellizzeri, Sera Markoff, Karl Menten, Monika Moscibrodzka, Alison Peck, Anthony Rushton, Reinhold Schaaf, Melvyn Wright

A&A, 576, A14 (2015)

Abstract

Radio and mm-wavelength observations of Sagittarius A* (Sgr A*), the radio source associated with the supermassive black hole at the center of our Galaxy, show that it behaves as a partially self-absorbed synchrotron-emitting source. The measured size of Sgr A* shows that the mm-wavelength emission comes from a small region and consists of the inner accretion flow and a possible collimated outflow. Existing observations of Sgr A* have revealed a time lag between light curves at 43 GHz and 22 GHz, which is consistent with a rapidly expanding plasma flow and supports the presence of a collimated outflow from the environment of an accreting black hole. Here we wish to measure simultaneous frequency-dependent time lags in the light curves of Sgr A* across a broad frequency range to constrain direction and speed of the radio-emitting plasma in the vicinity of the black hole. Light curves of Sgr A* were taken in May 2012 using ALMA at 100 GHz using the VLA at 48, 39, 37, 27, 25.5, and 19 GHz. As a result of elevation limits and the longitude difference between the stations, the usable overlap in the light curves is approximately four hours. Although Sgr A* was in a relatively quiet phase, the high sensitivity of ALMA and the VLA allowed us to detect and fit maxima of an observed minor flare where flux density varied by $\sim 10\%$. The fitted times of flux density maxima at frequencies from 100 GHz to 19 GHz, as well as a cross-correlation analysis, reveal a simple frequency-dependent time lag relation where maxima at higher frequencies lead those at lower frequencies. Taking the observed size-frequency relation of Sgr A* into account, these time lags suggest a moderately relativistic (lower estimates: $0.5c$ for two-sided, $0.77c$ for one-sided) collimated outflow.

4.1 Introduction

The radio source Sagittarius A* (Sgr A*) at the center of our Galaxy is the best-constrained supermassive black hole candidate found thus far (Genzel et al. 2010; Falcke & Markoff 2013 for a review). Located at a distance of 8.3 ± 0.4 kpc from the solar system, its mass is calculated to be $4.3 \pm 0.4 \cdot 10^6 M_{\odot}$ (Eisenhauer et al., 2003; Reid & Brunthaler, 2004; Ghez et al., 2008; Gillessen et al., 2009; Genzel et al., 2010). For a black hole of this mass, Sgr A* seems to be accreting gas at a very low rate of $\lesssim 10^{-7} M_{\odot} \text{ yr}^{-1}$, as was derived from Faraday rotation measures (Bower et al., 2005; Marrone et al., 2007).

The emission from Sgr A* between frequencies of 20 GHz and 230 GHz shows flux density variability of a few tens of percent on hour-long timescales, up to 100% on month-long timescales, as well as occasional flaring behavior (Dexter et al., 2014). In radio, Sgr A* has an inverted spectrum (i.e., rising flux density with increasing frequency) that peaks at the 'submm bump', around 350 GHz, beyond which the spectrum steeply drops in the infrared regime. The radio emission is thought to originate mostly from partially self-absorbed synchrotron radiation emitted farther out from the black hole, while emission at frequencies corresponding to the submm bump (Falcke et al., 1998) of the Sgr A* spectrum is commonly associated with the optically thin emission closest to the black hole (Falcke et al., 1998; Shen et al., 2005; Bower, 2006; Doeleman et al., 2008). In the mm regime and at longer wavelengths, the flux density variation is thought to arise from local bulk properties (magnetic field strength, gas density, temperature) of the plasma, while the variability seen in infrared and X-rays is mostly attributed to changes in the population of the high-energy tail of the local electron energy distribution (Özel et al., 2000; Markoff et al., 2001; Yuan et al., 2003; Dodds-Eden et al., 2010; Dibi et al., 2013).

While the emission mechanisms for the radio and mm-wavelength emission of Sgr A* are understood fairly well, the identification of the emission with specific flow regions is still a subject of debate. For example, an important question is whether the radio emission is generated in a jet (Falcke et al., 1993) or in a radiatively inefficient accretion flow (Narayan et al., 1995). Sgr A* in its flaring state fits neatly onto the fundamental plane of black hole activity (Merloni et al., 2003; Falcke et al., 2004; Plotkin et al., 2012), and as such it would be expected to feature a jet as other sources on that scaling relation do (Markoff, 2005). As yet, no direct detection of a jet has been made for Sgr A* despite the claimed presence of tantalizing jet-like features close to the Galactic center on parsec scales (Yusef-Zadeh et al., 2012; Li et al., 2013). Any putative jet structure close to the black hole cannot be resolved below observing frequencies of ~ 100 GHz because interstellar scatter-broadening blurs our view of the Galactic center at such frequencies, an effect that progressively increases with lower frequency (Lo et al., 1981; van Langevelde et al., 1992; Bower et al., 2006; Mościbrodzka et al., 2014). At higher observing frequencies, interstellar scintillation is less of a problem - in the mm-wave regime, existing VLBI networks should be able to directly observe the proposed shadow of the event horizon with mmVLBI (Falcke et al., 2000;

Doeleman et al., 2008).

There are other ways in which the nature of the emitting gas flow may be determined, however. Sgr A* exhibits an inverted radio spectrum. Flat or inverted radio spectra are commonly seen in quasars and active galactic nuclei, where the bases of radio jets resolved at high resolution show dominant emission at different radio frequencies as a function of distance from the core, which is due to optical depth effects (Hada et al., 2011), as has also been predicted from theory (Blandford & Königl, 1979; Falcke & Biermann, 1995). The multifrequency spectrum of Sgr A* (from radio to X-ray) in its flaring state looks very much like the spectrum of M81*, which has a weak jet (Bietenholz et al., 2004). The emission from an unresolved, compact jet may explain the inverted radio spectrum of Sgr A* (Falcke et al., 1993; Mościbrodzka & Falcke, 2013).

Presence of a jet implies that specific correlations should be detected between light curves at different frequencies. As the peak frequency of radio emission changes with position along the jet axis, we expect variations in flux density at different observing frequencies to exhibit time lags relative to one another as the emitting gas moves out. Previous observations of Sgr A* made with the VLA have indeed suggested the existence of a time lag of ~ 20 to 40 minutes in flux density variability between light curves measured at 43 GHz and 22 GHz, with variability in the higher-frequency lightcurve leading that in the lower-frequency lightcurve (Yusef-Zadeh et al., 2006, 2008). Yusef-Zadeh et al. interpreted this as emission from an expanding plasma cloud (van der Laan, 1966) with velocities reaching about $0.01c$, but this interpretation does not take VLBA sizes into account. When coupled to the observed relation between the observing frequency and the measured intrinsic size of Sgr A* (Bower et al., 2004; Doeleman et al., 2008), the time lag between 43 GHz and 22 GHz corresponds to a size difference of ~ 30 light minutes. Thus, such a time lag suggests the presence of a fast and directed outflow with a moderately relativistic speed (Falcke et al., 2009). As the emission at observing frequencies below the submm bump is probably all partially self-absorbed synchrotron emission, time lags may be present between the light curves at any two different frequencies in that region. Measurements of time lags over a wider range of frequencies are of interest as they may aid in establishing a flow velocity profile, and they may even provide an estimate of how close to the black hole the outflow can be traced.

4.2 Observations and data reduction

Our VLA observations were taken on May 18, 2012 from 05:25:15 UT to 12:54:01 UT in CnB configuration, chosen to coincide with a Chandra observation. Light curves for Sgr A* were taken in pairs of subbands for three basebands (X, Ka and Q), yielding 1-GHz-wide subbands at center frequencies of 19 and 25.5 GHz (K-band), 27.48 and 37.99 GHz (Ka-band), and 39.55 and 48.5 GHz (Q-band) - each using 30-second scans at an integration time of 3 seconds. The subbands eventually used for each center frequency were 4-7 (19 GHz), 0-7 (25.5 GHz), 0-7 (27

GHz), 0-7 (37 GHz), 2-7 (39 GHz), and 1-2 (48 GHz). Flux and bandpass calibration were made on the standard VLA calibrator 3C286. A monitoring loop with a period of 7.5 minutes was used for Sgr A*: within each iteration of this loop, J1744-3116 was used as a gain calibrator source and J1745-283 was used as a check-source, cycling through all three basebands in turn. The integration time of 3 seconds was chosen such that the RMS noise per scan was expected to be 1 mJy when all subbands in a baseband were used.

The VLA data (project code: 12A-339) were initially reduced using the VLA pipeline version 1.2.0 (rev. 9744) on CASA 4.1.0. After running the VLA pipeline, the sufficiently high flux density of all sources allowed us to perform careful phase self-calibration on them using progressively shorter solution intervals down to one integration length (3 seconds). Some subbands were flagged in this calibration process because their calibration solutions did not converge. The subbands that remained unflagged were (format: baseband (subbands)) 48 GHz (1,2), 39 GHz (2-7), 37 GHz (all), 27 GHz (all), 25 GHz (all), and 19 GHz (4-7). The declination of Sgr A* combined with the latitude of the VLA means that Sgr A* never reaches an elevation over 27 degrees for the VLA. Therefore the first and last parts (before approximately 6:40 UT and after approximately 11:50 UT) of the observation suffer from coherence loss: the source is less than 18 degrees above the horizon, and the effective path length through the atmosphere for the signal varies rapidly and strongly between antennas. As such, all data in these time windows were flagged before recalibration.

For the ALMA track, observations were made in ALMA cycle 0 on May 18, 2012 from 03:30:47 UT to 10:52:16 UT (project code: 2011.0.00887.S). The ALMA light curves for Sgr A* were taken at ALMA bands 3, 6, and 7 using pairs of spectral windows each centered on 95, 105, 247, 260, 338, and 348 GHz. Each pair of spectral windows covered 3.75 GHz bandwidth for a total of 7.5 GHz bandwidth per ALMA band. Scan lengths were chosen to yield a sensitivity of 0.5 mJy. All individual scans of Sgr A* were bracketed by either NRAO530 or J1924-292 as calibrators, while flux density calibration was made on Titan and Neptune. At that time, ALMA had 19 antennas available. In this paper we limit discussion to the 100 GHz ALMA data: the light curves at 250 GHz and 340 GHz will be the subject of a future paper.

For ALMA, the source setup was somewhat more complicated because all observations in cycle 0 had to be obtained in two-hour blocks. This means that the ALMA dataset consists of four separate blocks that are contiguous in time, each starting with a flux density calibrator measurement (Titan for the first two blocks, Neptune for the latter two). Within each block, five scans on Sgr A* were made where each of these was bracketed by scans on one of the calibrator sources NRAO530 and J1924-292, a precaution taken because of possible calibration difficulties that might otherwise occur in cycle 0. Thus the scan setup for each block was (using the first letter of each source) 'NSN JSJ NSN JSJ NSN', the average switching time between bracketing calibrator scans and Sgr A* scans was 2 minutes, and the time cadence on Sgr A* in each band was 15 minutes. As

J1924-292 exhibited irregular results in its polarisation-dependent flux density levels, it was not used as a calibrator and only used as a check-source. Gain levels were stable enough to warrant the usage of NRAO530 as the only gain calibrator. The ALMA data were calibrated using a custom script based on the calibration procedure for the QA2 process of Cycle 0 data, with subsequent phase self-calibration on Sgr A* (which has a very strong unresolved component of around 2.5 Jy at 100 GHz) using baselines longer than 150 k λ .

From the calibrated VLA and ALMA data, the light curves for Sgr A* were obtained by averaging all unflagged UV visibilities per scan from all baselines longer than 150 k λ , using uniform weighting for the selected baselines. We chose to only use projected baseline lengths over 150 k λ because we wished to avoid any contamination from the extended emission around Sgr A*, and baselines shorter than the chosen length show hints of structural variations over the track. Limiting ourselves to the longer baselines enabled us to work directly with the visibilities rather than needing the additional steps of imaging and model fitting. The errors on the flux density levels were estimated using the spread in amplitude of the calibrated visibilities over each scan. The calibrated light-curve data can be seen in the left column of Fig. 4.1. All light curves can be seen to exhibit a gradual rise and decay, with shorter-timescale variation superimposed. For the calibrated VLA data, the resulting noise levels for the basebands at 19, 25, 27, 37, 39, and 48 GHz are 1.4, 0.7, 1.2, 0.9, 2.6, and 12.1 mJy, respectively. Limited calibration accuracy and the flagging of several subbands degrades the sensitivity from the desired sensitivity in the highest frequency bands. The final uncertainties in flux density are dominated by the calibration uncertainty coming from the variability that was exhibited by the check-source: this brings the total relative flux density measurement uncertainty to approximately 5% for the VLA data. For the ALMA data, the uncertainty in flux density is approximately 5% as well - however, the errors there are relatively small compared to the intrinsic variability of Sgr A* at 100 GHz. We note that the flux calibration of the VLA data does not give a perfectly smooth spectrum for our calibrator source J1744-3116, but this only affects the overall flux levels of the Sgr A* light curves by $\sim 5\%$ and does not affect the conclusions in this paper.

4.3 Analysis and results

The spectral energy distribution of Sgr A* across the measured frequencies (see Fig. 4.2) has a spectral index of $\alpha = 0.41 \pm 0.03$ (with α defined as in $S_\nu \propto \nu^\alpha$) when all data (VLA + ALMA) are used for the spectral fit, and we obtain $\alpha = 0.50 \pm 0.07$ when only the VLA data are used. While this difference in spectral index cannot be called significant, it hints at a flattening of the spectrum as the submm bump is approached.

The VLA light curves at first sight each show a similar evolution of flux density with time: a rise in flux density level between 7h and 8h/9h UT, and a more slowly diminishing flux density




FIGURE 4.1: **Left:** Flux vs. time of Sgr A* for the VLA data (top six graphs) and the 100 GHz ALMA data (bottom graph). The fitted FRED function is plotted as a green (gray) curve, the position of the maxima is plotted as a vertical red (dark gray) line with the uncertainty in the fit superimposed as a red-shaded (dark gray) region. Green-shaded (light gray) regions indicate the uncertainty on the fits obtained by randomly dropping half of the data points for 500 iterations - see Sect. 3 for details. **Right:** Flux vs. time for the calibrator sources (J1744-3116 for the VLA data, NRAO530 for the ALMA data).

beyond 8h/9h UT. The observed longer-term flux evolution over the full track is overlaid with more rapid variations in measured flux, which occur simultaneously in all frequency bands. These fluctuations are probably caused by atmospheric influence, which causes varying coherence loss as a function of time. The ALMA flux density measurements are highly precise, with a very small spread in visibility values per scan. The time cadence, however, is coarser than it is for the VLA data. Nonetheless, the evolution of flux density with time can be distinguished with high significance. At 100 GHz the flux density evolution is smooth, and there is a local maximum in flux at around 7:45h UT, followed by a later peak around 10:00h UT.

The z-transform discrete correlation function (ZDCF) algorithm (Alexander, 1997) provides a way to cross-correlate light curves that have uneven temporal sampling. This method for finding




FIGURE 4.2: Averaged flux density as a function of frequency for all light curves for the time period between 7h and 11h UT. The error bars denote measurement error convolved with flux density variability over the track, where variability is the dominant contribution. Variability is strongest at the highest frequencies. The spectral index obtained from using VLA + ALMA data is indicated by the black line (with 1-sigma fit uncertainties filled in with gray) and yields $\alpha = 0.42 \pm 0.03$. Using the VLA data only (red/gray line, with red/light gray 1-sigma uncertainty region), $\alpha = 0.45 \pm 0.07$.

time lags between the VLA light curves yields a strong zero-lag component in every case (see Fig. 4.3 for an example), coming from the short-time fluctuations in the data and probably attributable to coherence loss. These zero-time lag spikes tend to dominate the cross-correlation curves. Although skew is apparent in most cross-correlation curves, the zero-lag peaks preclude any meaningful time lag estimates to be made this way. To derive reliable time lag estimates, we chose the simple and robust approach of fitting the longer-term flux density evolution in all light curves. To establish the times at which flux density maxima occur in these light curves, we fitted a smooth function to this general trend, allowing for different timescales to be associated with the rise and fall. Based on the general shape of the light curves, the choice was made to employ 'fast rise, exponential decay' (FRED) functions as fitting functions (Bhat et al., 1994) (widely used in GRB light curve fitting). These functions consist of the product of two exponentials, involving four free parameters:

$$f(t) = A \cdot e^{2\sqrt{b/a}} \cdot e^{-(t-\Delta)/a-b/(t-\Delta)}, \quad (4.1)$$

where A is the maximum value of $f(t)$, a and b are parameters controlling the slopes of either side of that maximum, and Δ is the value of t for which the maximum value is reached. Because the FRED flux density value rises up from zero and returns to zero as its argument t is left running, it is not a suitable function to use over time intervals that are too long: we are only interested in using it to fit a local and asymmetric feature in the light curves. To keep the general shape of the light curves sufficiently simple to enable the fit without sacrificing too many data points, we only used the flux density measurements directly around the bump feature (7h - 11h UT for VLA, 6:50 - 8:50 UT for ALMA). Acting on the assumption that this feature in the ALMA light curve can be attributed to the same event in the source that caused the maxima found in the VLA light curves, we used the fitting algorithm on this feature as well.

Variations on shorter timescales in the VLA light curves may affect the fit results. Therefore we also performed the fits using a Monte Carlo approach in which, at random, half of the data points

TABLE 4.1: Times of flux density maxima from curve-fitting of individual light curves. Column 1: all data, Col. 2: subsample with MC

Frequency	Time of max (1) (hrs UT)	Time of max (2) (hrs UT)
100 GHz	7.65 ± 0.19	-
48 GHz	7.93 ± 0.31	7.82 ± 0.90
39 GHz	8.03 ± 0.14	8.01 ± 0.12
37 GHz	8.17 ± 0.13	8.17 ± 0.13
27 GHz	8.38 ± 0.11	8.38 ± 0.11
25 GHz	7.98 ± 0.16	7.97 ± 0.15
19 GHz	8.78 ± 0.29	8.72 ± 0.19

(20 out of 40) for a given VLA light curve were dropped before attempting a fit with the FRED function. For each light curve, this was iterated 500 times. The resulting times of maxima for all obtained fits were averaged, and the standard deviation of their distribution was calculated. These results are also plotted in Fig. 4.1, shaded in green. The large standard deviation in the 48 GHz case indicates skew in the range of predictions made by the Monte Carlo method, causing the left limit to fall outside of the data range. High reduced χ -squared values indicate that we fitted light curves whose evolution is more complicated than can be grasped by a simple function; significant short-term variability remains from imperfect flux calibration. The absolute flux density uncertainties are dominated by the calibration uncertainty of ~ 5 percent for the entire light curve (not shown in the figure). For most light curves, these results correspond well to the fit for which all data were used. For the light curve at 48 GHz the fits are not as robust. This is most probably due to the relatively strong flux density variations at short timescales for the 48 GHz light curve, combined with the short FRED rise time. The resulting best-fit values and uncertainties for the times of flux density maximum are shown in Fig. 4.1, and the time lags we found are summarized in Table 4.1 where the middle column uses all data and the rightmost column uses the average result from the Monte Carlo approach. As the ALMA data does not have many measurements within the relevant time window, the Monte Carlo approach could not be used there.

With the 25 GHz data as notable exception, the maxima occur at later times for the lower observing frequencies, which is compatible with a linear relation between observing wavelength and time of maximum flux density. There is always the risk of misinterpreting the ALMA bump at 7:40h UT as being causally connected to the maxima in flux density found in the VLA data. To check whether the VLA data by themselves are consistent with the existence of these time lags, a separate fit was made using only the VLA data (Fig. 4.4), and the fit results are practically identical to those obtained with the full data set. The time lag expected for the ALMA peak based on the VLA time lag/frequency fit coincides with the measured value, and thus the identification of the ALMA flux density maximum as being related to the VLA flux density maxima seems justified. The fact that the maxima occur at different times for different frequencies also precludes interpreting




FIGURE 4.3: Example of a cross-correlation curve for the original light-curve data (left) and for the light-curve data with the FRED trend subtracted (right). The skew in the CCF for the full light curves suggests a nonzero timelag, but the zero-lag peak is too prominent to provide any useful estimate. After the fitted FRED trend is subtracted from both light curves and the cross-correlation is performed again, the only prominent cross-correlation peak corresponds to zero time lag.




FIGURE 4.4: Fitted times for the flux density maxima in each baseband (relative to the 19 GHz (1.6 cm) maximum), plotted as a function of observing wavelength. The figure uses the errors on the fit maxima obtained from fitting all data points. Dark shaded regions indicate the uncertainty in fit slope using all light curves, lighter shaded regions indicate the uncertainty in fit slope from VLA light curves alone. The slopes obtained are 42 ± 14 mins/cm (all data) and 36 ± 21 mins/cm (VLA data alone). The continuous lines are time-lag predictions from the jet model by Falcke et al. (2009) assuming inclinations of 30, 60, and 90°.




FIGURE 4.5: Variability due to FRED-fitted light-curve trends (vertical axis) versus remaining variability in trend-subtracted light curves (horizontal axis) for the VLA data. Both axes show standard deviations. These quantities indicate that the long-term variability that is fitted by the FRED function scales with the remaining short-term variability.

the observed flux density evolution as being a purely atmospheric or elevation-dependent effect. Any elevation-dependent change in measured flux density would impose a simultaneous rising and falling of all light curves, which is not what we observe.

The FRED fits follow the general, long-term trend that is present in the data. In addition to this general trend, all light curves exhibit shorter-timescale fluctuations. Cross-correlation analysis on the original light curves therefore shows correlation contributions from both the general trends in the data and the shorter-timescale fluctuations. If these two variability components do not exhibit the same time lag between frequencies, interpreting the cross-correlation curves is problematic. To deal with this problem, we have subtracted the fitted FRED trends from all light curves and performed a cross-correlation analysis on the de-trended light curves (see Fig. 4.3). This cross-correlation peaks at zero time lag, suggesting that the short-term fluctuations have a different origin from the long-term trends. The most likely cause of the short-term fluctuations are calibration residuals stemming from phase-coherence loss due to the low elevation of Sgr A* at the VLA site. While there is a correlation between the FRED variability and the residual variability as seen from Fig. 4.5, this does not imply that they share the same cause. Atmospheric influence is stronger for higher observing frequency, and this effect is unrelated to intrinsic source variability.

4.4 Discussion and conclusion

The time lags across this broad range of frequencies corroborate the picture of an expanding plasma flow with a diminishing optical depth over time. When the time lags found in this work are combined with the existing results from Yusef-Zadeh et al. (2006), they are compatible: Yusef-Zadeh et al. reported a time lag between 43 GHz and 22 GHz of 20 to 40 minutes, while we detect a time lag of 28 ± 9 minutes between these two frequencies when we use the linear time lag/wavelength fit based on our measurements (see Fig. 4.4 and caption).

Measurements on time lags between Sgr A* light curves at 102 and 90 GHz were performed by Miyazaki et al. (2013), and they reported the time delay between 102 GHz and 90 GHz as being -2.56 ± 0.9 min (i.e., the 90 GHz light curve is leading the 102 GHz light curve). The expected time lag between 102 and 90 GHz that would agree with the 43 GHz -22 GHz lag found by Yusef-Zadeh et al. (2006) is quoted as being close to 3 min (with the precise value depending on the index of the power-law distribution in electron energy), whereas an extension of the linear relation we find in this work predicts a time lag between these frequencies of 1.7 ± 0.6 min (this figure increases somewhat if low inclination angles are considered, see Fig. 4.4). We stress that the models used to predict the time lags are very simple in all cases, and in particular measured time lags between closely neighboring frequencies can deviate from the predicted relation due to more complex plasma flow properties close to the black hole. The Blandford-Königl jet model uses a $\tau = 1$ surface, the location of which along the jet only depends on the accretion rate, and which is constant throughout the jet cross-section. The actual nature of any outflow may locally be of a more chaotic character, with different regions in the jet cross-section having different plasma densities and different optical depths, as is typically witnessed in GRMHD simulations (Mościbrodzka et al., 2014). Our measurements, obtained over a broad range of frequencies, are expected to sample a greater spatial range of the proposed outflow. They should hence provide a robust result and a characterization of the behavior of the system as a whole, and we believe this approach warrants the use of a simple outflow model.

Without source size measurements, many different models of expanding plasma flows can be made to fit our observations. Following the analysis by Yusef-Zadeh et al. (2008), where an adiabatically expanding plasma cloud was used as a model for flare occurrence in Sgr A*, the cloud expands to just ~ 2.3 times the size it initially has at the 100 GHz maximum (see Figs. 4.6 and 4.7) for the frequency range we record. The radii at which the lower-frequency emission peaks is only a few times the initial radius (which is taken to be 3 Schwarzschild radii). If we adopt the initial radius of the cloud as being ~ 3 Schwarzschild radii, as was done by Yusef-Zadeh et al. (2008), the associated flow velocities that occur according to this model are only around 3 percent of the speed of light. This is twice the velocity that was found by Yusef-Zadeh et al. (2008).

However, we can use the size-frequency relation of Sgr A* as presented in Falcke et al. (2009) as




FIGURE 4.6: Normalized flux density profiles of an adiabatically expanding plasma cloud as a function of radius for all frequencies. The continuous flux profiles use the expressions given in Yusef-Zadeh et al. (2008), with a particle spectral index of 1 as was determined to be the best-fitting value in that work. Note that the initial radius used by Yusef-Zadeh et al. (2008) (~ 3 Schwarzschild radii) is different from the initial radius that follows from using the size-frequency relation (~ 11 Schwarzschild radii).

additional information with which to provide a general estimate of the gas flow velocity. Although we do not have sufficient information to identify the variable emission (that we focus on) with the quiescent emission for which the size-frequency relation holds, considering the two components to reflect the same gas flow is the simplest hypothesis that we can consider. The size-frequency relation describes the measured intrinsic (i.e., corrected for interstellar scattering) size of Sgr A* as a function of observing wavelength, and has the form:

$$\phi_{\text{Sgr A}^*} = (0.52 \pm 0.03) \text{mas} \times (\lambda/\text{cm})^{1.3 \pm 0.1}. \quad (4.2)$$

In this expression, $\phi_{\text{Sgr A}^*}$ is the angular size of Sgr A* on the sky and λ is the observing wavelength. Combining this (angular) size-frequency relation with an estimate for the distance between Earth and the Galactic center and taking the difference in source size for two observing frequencies, we obtain an expression for the projected source size difference on the sky in length units. When we assume a source inclination of 90° and a one-sided outflow interpretation (as was done by Falcke et al. (2009)), we find that our data suggest an outflow velocity of $\sim 0.77c$ ($\gamma\beta \approx 1.2$). Using the variability seen in our measurements (Fig. 4.5) as a proxy for flare amplitude with the size-frequency relation (Eq. 4.2), we obtain the data points that are plotted separately in Fig. 4.6.




FIGURE 4.7: Times of flux density maxima for different frequencies plotted against the radii at which maximum flux density is reached. Black data points are obtained by combining our time lags with the outflow model used in Yusef-Zadeh et al. (2008), see also Fig. 4.6. If the initial radius is a only few Schwarzschild radii, the estimated flow velocity reached by the expanding plasma is on the order of $0.03c$. Orange (light gray) data points are obtained by combining our time lags with the size-frequency relation (Eq. 4.2), and yield a velocity of $\sim 0.5c$. Note that the initial radii are not the same for both cases.

Because the source centroid position on the sky as a function of observing frequency is not known, however, we can derive a lower limit on the outflow velocity by assuming identical centroid positions for all observing frequencies. This assumption corresponds to a two-sided jet interpretation (the source grows symmetrically on the sky with lower observing frequency), so for the distance traveled by the gas in one jet we can take half of the intrinsic source size difference. Taking into account the influence of light travel time with different inclination angles yields the following expression for the expected time lag between observing frequencies:

$$\Delta_{\text{diff}} = \frac{R_{\text{SgrA,diff}}}{\sin i} \left(\frac{1}{v_f} - \frac{\cos i}{c} \right), \quad (4.3)$$

with Δ_{diff} the time lag between two observing frequencies, $R_{\text{SgrA,diff}}$ the radius difference for the two observing frequencies as calculated using Eq. 4.2 and our distance to the Galactic center of 8.3 kpc, i the inclination angle (angle between the flow vector and the line of sight from Sgr A* to Earth) and v_f the flow velocity. We can express this relation in terms of v_f and combine it with the relation between time lag and wavelength (see Fig. 4.4). In this way, we can plot the relation between flow velocity estimate and jet inclination angle, where we place the constraint that the flow velocity needs to be positive and lower than the speed of light. Figure 4.8 shows this




FIGURE 4.8: Estimates of minimum outflow velocity (expressed in $\gamma\beta$) as a function of the inclination angle of the flow direction. The uncertainties have been calculated from the uncertainty in Galactic center distance (0.4 kpc), the uncertainties in the size-frequency relation (see Eq. 4.2) and the uncertainty in time lag vs observing wavelength (14 mins/cm, see Fig. 4.4) using standard error propagation.

dependence and indicates a minimum flow velocity of $v_f = 0.5c$ ($\gamma\beta = 0.58$) for an inclination of 60° . Including the light travel time in the calculation breaks the symmetry around an inclination of 90° that would otherwise be present. For the two-sided outflow, we assumed here that the outflow component with an inclination smaller than 90° is the one that is picked up in the time lag measurements. For the two-sided jet interpretation only inclinations close to 90° can be modeled reasonably in this way. Inclinations deviating significantly from 90° would result in ambiguous time lags because of the different light travel times from the gas in the two jets. For the case of a one-sided jet, where the source only grows toward one side with lower observing frequency, the lower velocity bound is approximately $v_f = 0.77c$ ($\gamma\beta = 1.2$, see Figure 4.9). Figure 4.4 plots the model-predicted time lags for different inclination angles and shows that no strong constraints are posed on the inclination angle by these data.

As these results are based upon measurements of a light-curve feature from a single track, it is prudent to perform this analysis on more light curves as they become available to verify the picture we establish. Although the broad nature of the flux density feature that we used in the FRED fit generally agrees with the result of the ZDCF analysis, light curves with significant flux density changes over shorter timescales would offer an opportunity of using the ZDCF more effectively as an alternative verification of the FRED fitting results. ALMA observations at a higher time cadence taken contemporaneously with VLA observations would facilitate the cross-matching of




FIGURE 4.9: Time delays as a function of frequency, expressed in minutes before the time of maximum flux density at $\lambda = 1.57$ cm (19 GHz). Colored lines indicate the model time-lag predictions for different constant outflow and inflow velocities. Orange (light gray) lines indicate outflow, blue (dark gray) lines indicate inflow. The data points, derived from the full-data FRED curve fits, suggest an outflow with a moderately relativistic velocity, clustering around $\gamma\beta = 2$. This figure is valid for an inclination angle of 90 degrees and is adapted from Falcke et al. (2009).

light-curve features between these frequencies.

To summarize, we have measured time lags in Sgr A* light curves from 100 GHz to 19 GHz using ALMA and the VLA. Higher-frequency light curves are seen to have their maxima at earlier times than the lower-frequency light curves. Coupled to the size-frequency relation for Sgr A*, these measurements indicate a moderately relativistic, directed outflow from Sgr A*.

4.5 Acknowledgements

This work was performed as part of the Chandra XVP collaboration (see <http://www.sgr-star.com/collaboration-members/> for a list of collaboration members). This paper makes use of the following ALMA data: ADS/JAO.ALMA#2011.0.00887.S. ALMA is a partnership of ESO (representing its member states), NSF (USA) and NINS (Japan), together with NRC (Canada) and NSC and ASIAA (Taiwan), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO and NAOJ. We thank our referee for providing useful comments that improved the presentation of this paper.

Christiaan Brinkerink, Heino Falcke
to be submitted

Abstract

The compact radio source at the center of our Galaxy, Sagittarius A* (Sgr A*), is the subject of intensive study as it provides a close-up view of an accreting super-massive black hole. Sgr A* provides us with a prototype of a low-luminosity active galactic nucleus (LLAGN), but interstellar scattering and the resolution limits of our instruments have limited our understanding of the emission sites in its inner accretion flow. The temporal variability of Sgr A* can help us understand whether we see a plasma outflow or inflow in the region close to the black hole. In this work, we look at a comprehensive set of multi-epoch data recorded with the Karl G. Jansky Very Large Array (VLA) to understand the persistence of the time lag relations that have been found in previous radio observations of Sgr A*. We analyse 8 epochs of data, observed in Spring 2015, each of which has a frequency coverage from 18 to 48 GHz. We cross-correlate the calibrated light curves across twelve frequency subbands. We also generate synthetic data with the appropriate variability characteristics and use it to study the detectability of time lag relations in data with this sampling structure. We find that the variability amplitude increases with frequency. We see positive time lag slopes across all subbands in five out of eight epochs, with the largest slopes in the cases where a clear extremum in flux density is present. Three epochs show lag slopes close to zero. With the synthetic data analysis we show that these results are explained by a persistent lag relation of ~ 40 min/cm that covers the bulk of the variability, with at most 2% of the total flux density in an uncorrelated variability component. Together with the size-frequency relation and inclination constraints this indicates an outflow velocity with $\gamma\beta = 1.5$, consistent with predictions of jet models for Sgr A*.

5.1 Introduction

There is significant evidence and support for the notion that the compact radio source at the center of our Galaxy (Sagittarius A*, abbreviated as Sgr A*, see Balick & Brown (1974)) is associated with plasma flow around a supermassive black hole of approximately 4 million solar masses. Analysis of the orbits of short-period bright stars shows that the central mass is highly concentrated and coincident with the region from which we see the radio emission (Reid & Brunthaler, 2004; Ghez et al., 2008; Gillessen et al., 2009; Gravity Collaboration et al., 2019). In this introduction we will focus on the different observational aspects of this radio source and discuss the system properties that have been derived from them.

5.1.1 Observed spectrum

Sgr A* has a spectral energy distribution (SED) that shows a rising power-law across the radio spectrum, steepening into a submm-bump that starts above ~ 50 GHz and peaks in the 1 - 2 THz range (Zylka et al., 1992; Serabyn et al., 1997; Falcke et al., 1998; Zhao et al., 2003; Bower et al., 2019), before rapidly dropping into the infrared regime. The general shape of this spectrum suggests that we see partially self-obscured synchrotron emission, optically thick at lower frequencies and turning over to optically thin emission at higher frequencies.

The observed spectral shape and flux density of the submm bump allow us to solve for several properties of the inner plasma flow: the electron number density, the magnetic field strength and the temperature of the plasma (Falcke et al., 1993; Narayan et al., 1995; Falcke & Biermann, 1996; Yuan et al., 2003; Goldston et al., 2005; Mościbrodzka et al., 2009). These results show that the emission from the submm bump must come from the innermost part of the plasma flow: the lower-frequency part of the spectrum is self-absorbed, indicating that we are seeing into a partially transparent photosphere layer of an otherwise opaque region. This frequency-dependent photosphere shrinks inwards as we consider higher observing frequencies, until we reach the turnover point in the spectrum: there, the photosphere disappears altogether and we see the region closest to the black hole. The fact that the submm spectrum of Sgr A* shows this structure means that Sgr A* is also a prime candidate for very long baseline interferometry (VLBI) observations of its shadow by the Event Horizon Telescope (Event Horizon Telescope Collaboration et al., 2019c): the attainable angular resolution for a worldwide network of antennas observing at 230 GHz allows it to resolve the expected scale of the black hole shadow in the case of Sgr A* (Falcke et al., 2000). A fundamental open question in this context is: what is the emitting plasma in this inner region doing? Is it part of an inflow or an outflow?

5.1.2 Observed morphology

Studies of the morphology of Sgr A* with Very Long Baseline Interferometry (VLBI) at wavelengths from 20 cm down to 1.3 mm indicate an elliptical (Gaussian) shape for the source and

show that the apparent source size is dominated by interstellar scattering effects at wavelengths longer than ~ 3 mm (van Langevelde & Diamond, 1991; Lo et al., 1998; Bower et al., 2006), making it scale according to λ^2 for wavelengths in that range. At shorter wavelengths, the observed size of Sgr A* deviates from this relation: it is larger than the λ^2 relation predicts (Bower et al., 2004; Shen et al., 2005; Doeleman et al., 2008; Ortiz-León et al., 2016; Brinkerink et al., 2019), with the fractional difference becoming more pronounced at progressively shorter wavelengths. This is understood to be a manifestation of the intrinsic source geometry which becomes more clearly visible as the influence from scattering loses its dominance at higher frequencies. Although the measured size of Sgr A* at 86 GHz is larger than predicted from a pure scattering size relation constructed using measurements at lower frequencies, the deviation of its morphology from Gaussianity is modest with only $\sim 1\%$ of the observed VLBI flux density not matching an elliptical Gaussian brightness distribution (Brinkerink et al., 2019). While the scattered source geometry approximates an elliptical Gaussian very closely at longer wavelengths, substructure at the sub-percent level has also been seen at 1.3 cm (23 GHz) as reported by Gwinn et al. (2014).

The observed non-Gaussian substructure manifests as an asymmetry that may either be intrinsic in origin, caused by the effects of interstellar scattering or a combination thereof (Ortiz-León et al., 2016; Brinkerink et al., 2016). Observations on the degree to which this observed source asymmetry persists over longer time scales ($\sim$ years) should resolve this origin, as the time scale over which the influence from the scattering screen evolves is relatively short ($\sim$ weeks).

Algorithmically, disentangling the scattering effects from the contribution of the intrinsic source structure is a challenging task, although in recent years significant progress has been booked in this endeavour (Johnson & Gwinn, 2015; Johnson, 2016; Johnson et al., 2018; Issaoun et al., 2019) which allows for partial reconstruction of intrinsic source geometry from measurements of the scattered source image.

Given modern VLBI capabilities, the argument may be raised that jet morphology should therefore be readily apparent when looking at Sgr A* at radio-to-mm wavelengths, and in fact observations at 7 mm have indeed suggested this (Lo et al., 1998). From the theoretical side, it has been shown that the morphology of a jet outflow of Sgr A* may be such that it would appear as a highly compact source, within the size constraints dictated by these VLBI measurements (Markoff et al., 2007). More recent VLBI measurements at 86 GHz, which suffer less from the effects of interstellar scattering, can still be fitted with models from both jet- and disk-dominated classes (Mościbrodzka, 2017; Issaoun et al., 2019).

5.1.3 Time-domain studies

Besides the spectral and spatial dimensions, the emission from Sgr A* has been studied extensively in the time domain as well. Of particular interest are the different states of activity that

Sgr A* can exhibit, and the potential temporal correlations between the light curves of Sgr A* in different parts of the electromagnetic spectrum.

Analysis by Falcke (1999) showed that a characteristic variability timescale is present at a scale of multiple tens of days, with a suggestion of quasiperiodic behaviour at 57 days. Herrnstein et al. (2004) presented flux density measurements of Sgr A* spread out over multiple years at wavelengths of 2.0, 1.3 and 0.7 cm (15, 22 and 43 GHz), and reported tentative evidence for a bi-modal activity pattern. However, the time cadence of those measurements (8 days) was such that the shorter potential variability timescales, at which significant variability power might be concentrated, were not accessible. Indeed, Dexter et al. (2014) identified a characteristic timescale of approximately 8 hours in the variability of Sgr A* at high observing frequencies (230 GHz and above). Below this timescale, Sgr A* exhibits a variability power spectrum that looks like red noise, with flux density measurements separated closely in time being more strongly correlated than those with a larger time difference between them. Beyond this 8-hour timescale, that analysis shows that the variability has a flat power spectrum and thus indicates no systematic correlation between flux density variations regardless of their separation in time. More recently, a study of the variability of Sgr A* in infra-red has shown a coherence time of only 4 hours (Witzel et al., 2018) - possibly because of the shorter electron cooling timescales at the associated electron energies and in the relevant emission regions.

The first study reporting the detection of time lags between Sgr A* light curves was Yusef-Zadeh et al. (2006), in which flares at 43 GHz (0.7 cm) were followed by flares at 22 GHz (1.3 cm) approximately 20 to 40 minutes later. This time lag was interpreted in the context of expanding plasma blobs as described by van der Laan (1966), where the peak of the flare emission shifts to lower frequencies as the plasma blob expands and its optical depth changes. A different interpretation, ascribing the variations in flux density to emission from a compact jet, was put forward by Falcke et al. (2009), where the frequency-dependent intrinsic size of Sgr A* that was derived from VLBI measurements was coupled to the measured time lag to provide support for a compact jet. In that model, the plasma semi-adiabatically expands as it is accelerated along the jet axis. This scenario is also compatible with theoretical models for the outflow from a low-Eddington accreting system.

In a previous paper, we reported seeing a time lag in single-epoch VLA data of Sgr A* in which a minor flare was visible, spanning 7 different frequency bands from 100 down to 19 GHz (Brinkerink et al., 2015). The time lag relation was fitted with a linear trend in the wavelength domain, which yielded a slope of 42 ± 14 min/cm. Combining this time lag relation with an expression for the intrinsic size of the source and the well-constrained distance to Sgr A*, we found that if the variability is indeed associated with an outflow it suggests mildly relativistic outflows with a Lorentz factor of ~ 2 .

Measurements reported by Miyazaki et al. (2013) at higher frequencies (102 and 90 GHz) did not

show evidence for time lags between the two light curves. The apparent absence of a non-zero lag in this data set was theorised to either follow from the dynamics of the expanding plasma (the plasma blob might start out optically thin already, therefore not exhibiting a clear peak in its variability at any of the observing frequencies), or alternatively the variability seen in the light curves might be due to some other process besides plasma expansion – for instance in the form of orbiting hot spots in the accretion disk.

5.1.4 Questions addressed in this work

The open questions we wish to address here focus on the statistics of time lags in radio observations of Sagittarius A*: how consistent is the measured time lag relation, how much does it vary between epochs? Does it ever reverse sign? Does it correlate with some other property of the source state? Do the source variability statistics work out in such a way that a time lag should be observed for all epochs?

In this paper, we present our findings from multiple epochs of data recorded with the VLA in the spring of 2015. We describe the process with which the data was recorded and calibrated in Section 2, and we describe the components of our analysis in Section 3: it contains a discussion on the methods we have used to establish the light curves and their cross-correlations, as well as our synthetic data analysis with which we verify the robustness of our measurements. In Section 4, we present the time lag relations we measure. We connect our results to the broader theoretical context in Section 5.

5.2 Observations

5.2.1 Observation epochs, array configuration and spectral setup

The observations were carried out using the VLA (project code 15A-372, PI: C. Brinkerink). Nine observing blocks were executed between March 15th, 2015 and May 1st, 2015 (see Table 5.2). For all of these observations the VLA was in A configuration, with baseline lengths of up to 36.4 km. Although a total of 27 hours of observing time was allocated for this project, schedule planning constraints dictated a fixed scan length and structure so that they could be scheduled as short (~3-hour), independent observing blocks that could be executed whenever the opportunity arose. All of the observing blocks thus follow the same sequence of configuration steps and scans. For each observing block, we include flux and bandpass calibration observations on standard flux calibrator 3C 286, followed by rapid cycled pointings toward Sgr A* (science source), J1744-3116 (gain calibrator) and J1745-283 (check source). Three VLA bands are observed (K, Ka and Q) in LL and RR polarisations, each with 8.192 GHz sky bandwidth using the 3-bit sampling mode for the WIDAR correlator. Each band is covered by 64 contiguous spectral windows (SPWs), where each SPW is 128 MHz wide and contains 128 channels with a bandwidth of 1 MHz each.

TABLE 5.1: Structure of one scan block

Target	Duration	Description
J1744-3116	30s	Setup K-band receiver
J1744-3116	30s	K-band gain calibration scan
Sgr A*	35s	K-band science scan
J1745-283	35s	K-band check source scan
J1745-283	30s	Setup Ka-band receiver
J1745-283	30s	Ka-band check source scan
Sgr A*	35s	Ka-band science scan
J1744-3116	35s	Ka-band gain calibration scan
J1744-3116	30s	Setup Q-band receiver
J1744-3116	30s	Q-band gain calibration scan
Sgr A*	35s	Q-band science scan
J1745-283	35s	Q-band check source scan

TABLE 5.2: Dates and times of the observations used for this work

No.	Date	Total time range	Time range on Sgr A*
1	15 Mar 2015	11:06:04 - 14:05:30	12:06:45 - 14:04:06
2	30 Mar 2015	08:55:45 - 11:55:12	09:56:30 - 11:53:51
3	30 Mar 2015	11:55:22 - 14:54:45	12:56:00 - 14:53:21
4	10 Apr 2015	08:12:46 - 11:12:09	09:13:27 - 11:10:48
5	11 Apr 2015	08:08:36 - 11:08:03	09:09:21 - 11:06:42
6	11 Apr 2015	11:08:11 - 14:07:36	12:08:51 - 14:06:12
7	30 Apr 2015	06:53:52 - 09:53:21	07:54:36 - 09:51:57
8	30 Apr 2015	09:53:25 - 12:52:51	10:54:06 - 12:51:30
9	1 May 2015	09:28:56 - 12:28:21	10:29:36 - 12:27:00

For each of our three band tunings, we cycle through our three sources where we dwell on each source for 30 seconds (see Table 5.1 for the structure of one such scan block). Using this scheme, which is repeated 15 times, we get light curves for Sgr A* that have a scan cadence of 6.5 minutes for each band. Pointing scans on source NRAO 530 are included in each observing block with a 1-hour cadence between these scan blocks. This scan setup gives us the necessary temporal resolution to test for the presence of time lags that follow the relation of $42 \pm 14 \text{ min/cm}$ as found in Brinkerink et al. (2015), but also for other time lag relations (see the synthetic data section for details).

5.2.2 Data calibration and reduction

The general strategy for data reduction was to perform the standard steps of flux and bandpass calibration followed by gain calibration. Data set 7 could not be processed by the VLA pipeline

due to missing tables in the downloaded data set. Although the download was done multiple times, the issue kept occurring and so the decision was made not to use the data from that epoch.

When running the VLA pipeline on the data sets (CASA version 4.7.2), a bug was found in the `set_jy` task where the model flux density for the gain calibrator source did not get set properly in the model column of the measurement sets. This issue is likely due to the large data volume that needed to be processed for each epoch. As a consequence, the model flux density for the gain calibrator remained fixed at 1 Jy for all frequencies, which is the default value in absence of a supplied model flux density. To remedy this issue, we phase-selfcalibrated the data on each of the three sources after running the VLA pipeline, using all baselines for J1744 and J1745, and using only baselines of >150 k λ for Sgr A*. After verifying that phases were flat and zero across each spectral window, we then frequency-averaged to yield 1 frequency channel per spectral window. Gain calibration was performed on the resulting frequency-averaged data sets, which were considerably smaller in size and as such posed no issues regarding the setting of the model flux density for J1744, our gain calibrator source. As the last step, a further reduction in data volume was made by frequency-averaging contiguous chunks of 16 SPWs to yield one light curve per 2-GHz frequency band. The resulting gain-calibrated data was verified to yield sensible SEDs for Sgr A* and J1745, which are shown in Figure 5.1. If we assume a constant flux density over time for our calibrator source, we see that flux densities for J1745 show a variability between epochs of 6% at 19 GHz up to 20% at 47 GHz, while Sgr A* shows a larger variability ranging from 25% at 19 GHz to 40% at 47 GHz. This indicates that we do see true variability in Sgr A* between epochs. The averaged Sgr A* spectrum from this work is plotted in the context of previous observations in Figure 5.2. The fact that we see somewhat higher flux densities for Sgr A* than from measurements in earlier years is compatible with the finding that Sgr A* exhibits a rising trend in flux density over multiple decades, as was remarked in Bower et al. (2015b).

5.3 Data analysis

5.3.1 Extracting light curves

As the visibility data show constant amplitudes per scan over the range of (u, v) -distances we consider for Sgr A*, the Sgr A* light curves are calculated directly from calibrated visibility data with no intermediate imaging steps involved. To this end, the calibrated visibilities from the longer baselines (≥ 150 k λ) are averaged together per scan. These long baselines resolve out all the non-compact structure in the field of view, so that Sgr A* itself is the only source with a meaningful contribution to the total flux density for these baselines. Because the VLA was in A-configuration for these observations, the majority of baselines are longer than this chosen cutoff length at all times and all observing frequencies. The resulting Sgr A* light curves for each epoch are plotted in Figure 5.3, while a sample of light curves for the calibrator source and the check source are shown in Figure 5.5 (the full set of light curves for these sources is included in




FIGURE 5.1: Spectral Energy Distributions plotted in log-log scale for the calibrator J1744-3116 (left), the check source J1745-283 (middle) and Sgr A* (right), for all data sets.

Appendix A).

The theoretical thermal noise for VLA ranges from ~ 90 microJy per 30s scan for low K-band to ~ 300 microJy per 30s scan for high Q-band¹. In these calibrated light curves, the measured variance per 2 GHz averaged spectral window and per 30s scan ranges from ~ 1 mJy for K band to ~ 2.5 mJy for Q band. Considering possible causes for this variability, the integrated flux density from the Galactic Center within the VLA primary beams is not expected to have a significant effect: the total emission from Sgr A West, which is comparable in angular extent to the VLA primary beam, is expected to be several tens of Jy in this part of the radio spectrum (Law et al., 2008). The System Equivalent Flux Density for the VLA between 20 and 40 GHz is ~ 500 Jy²,

¹From the VLA exposure calculator at <https://obs.vla.nrao.edu/ect/>, with the settings appropriate for our observations: low elevation, 27 antennas, 3-bit sampling, natural weighting, dual polarisation, 2 GHz bandwidth, spring weather

²from <https://science.nrao.edu/facilities/vla/docs/manuals/oss/performance/sensitivity>, consulted July 2021




FIGURE 5.2: The spectral data from this work, averaged over all epochs (red), and plotted together with binned measurements from the past decades (blue, data from Zylka et al. (1995); Serabyn et al. (1997); Falcke et al. (1998); Zhao et al. (2003); Herrnstein et al. (2004); An et al. (2005); Marrone et al. (2006); Bower et al. (2015b); Brinkerink et al. (2015); Liu et al. (2016); Bower et al. (2019)). The theoretical jet model from Falcke et al. (2000) is shown, indicating the steepening of the spectrum into the submm bump. The model from Yuan et al. (2003) is shown as well, it reproduces this spectral steepening but overpredicts the magnitude of the submm bump somewhat. The full set of non-binned measurements used in this plot is shown in Figure 5.21 in Appendix C.

which means that the thermal noise of the system is only expected to vary by at most 10% as a consequence of this emission. The larger difference that we find between the theoretical noise predictions and the spread in the measured values can therefore be attributed to minor variations in compact source flux density within a single scan. However, the scan-to-scan variations in flux density are considerably larger than this (by about an order of magnitude), and hence dominate the total variability at each frequency. For the remainder of this paper, the twelve subbands for each data set will be indicated as K1 thru K4 (19, 21, 23 and 25 GHz), Ka1 thru Ka4 (30, 32, 34 and 36 GHz) and Q1 thru Q4 (41, 43, 45 and 47 GHz), so in order of increasing frequency.

An initial inspection of the Sgr A* light curves shows clear variability over time within epochs, where several data sets exhibit a curvature with well-localised flux density maxima or minima across all bands in their general trends (most notable in this regard are epoch numbers 5, 8 and 9). We also see examples of single-scan offsets in flux density that show strong correlation across the 4 subbands within one band but which do not have a counterpart in the other two bands, which suggests the presence of residual calibration errors of up to ~ 0.03 Jy in the light curves.




FIGURE 5.3: Calibrated light curves for Sgr A*, data sets 1-6, 8 and 9. All data is plotted with error bars, which are so small as to be obscured by the data markers in most cases.

Because of the limited time interval covered by each of the light curves, they typically do not show variation around a clear stationary flux density but rather exhibit trends over the full epoch where the flux density changes as part of a pattern of longer-timescale variation. As a result, this longer-timescale variation (which is not completely characterised) tends to dominate the total




FIGURE 5.4: Top left: the de-trended RMS variability of Sgr A* flux density versus frequency, for all bands (horizontal axis) and all epochs (symbol series). Top right: same, but variability is expressed in terms of the fraction of the average flux density per epoch at that frequency. Both plots show linear fits to the variability versus frequency with a 1-sigma uncertainty band. Bottom left: de-trended RMS variability of flux density per subband versus average flux density per epoch, for all epochs (symbol series). Bottom right: same, but variability is expressed in terms of the fraction of the average flux density.

variability of a light curve and also induces a large spread in flux density values. To be able to characterise the variability of these light curves in a meaningful way, we chose to focus solely on the short time scales by de-trending the light curves using linear least-squares fits and looking at the residual variability.

In Figure 5.4, the de-trended RMS flux density variability per epoch is shown for all twelve frequency bands. In the plots, the RMS variability of the residual is shown. Generally, variability




FIGURE 5.5: Examples of calibrated light curves for J1744-3116 (calibrator) and J1745-283 (check source) for epoch 3. All data is plotted with error bars, which are obscured by the data markers in most cases. The full set of light curves for the calibrator and the check source for all epochs can be found in Appendix A.

is larger for higher frequencies and appears to scale roughly linearly with flux density but with considerable variance. We find a modulation index (RMS variability as fraction of flux density) on the order of 1%, but with a growing trend towards higher flux densities. Bower et al. (2015b) reported a larger modulation index of 8% for this frequency range, but that figure included longer-term variability. The fact that we do not see a clear linear relation between flux density and trend-subtracted RMS variability for any subband frequency, as would be expected from the RMS-flux relation (Uttley & McHardy, 2001; Abuter et al., 2020), may be due to the fact that we are not sensitive to this longer-term contribution to variability.

5.3.2 Detection of time lags

In our previous time lag paper (Brinkerink et al., 2015), the time lag calculations were done by fitting Fast-Rise Exponential Decay (FRED) curves to the measured light curves in the different frequency bands. This was done because of the clear and consistent presence of a flux density maximum across multiple light curves, of which the shape lent itself well to FRED fitting. For the data sets presented in this paper, the flux density evolution shows different characteristics for each epoch and so a more general way to look for lag relations between bands is needed.

To this end, multiple cross-correlation methods were considered that can deal with a non-constant time sampling cadence for the data series. The Z-transformed Discrete Correlation Function (ZDCF) by Tal Alexander (Alexander, 1997) is such an implementation, as is the Local Cross-Correlation Function (LCCF) as described by Welsh (1999). We found that both of these methods yield highly similar auto- and cross-correlation functions for our data sets. Out of these two methods the LCCF was chosen for the main analysis in this paper, because of its simple and portable implementation.

For the cross-correlation parameters we chose a minimum number of 11 measurements per time

lag bin (which is the same as the recommended minimum number for the ZDCF, and advised from a statistics point of view), and a minimum time lag bin width of 30 seconds (the length of a single scan). Using these parameter values, the cross-correlation functions for all twelve light curves in each data set were calculated against themselves as well as each of the other eleven light curves in the same data set. Each cross-correlation between two different light curves from the same data set is thus represented twice in the correlation results.

Examples of the intermediate results from this cross-correlation step are shown in Figure 5.6. In the cross-correlation maps shown there, the time lag relation found in Brinkerink et al. (2015) is overplotted as a black line with marker dots. Note that it is not a fit to the maxima of the plotted cross-correlation data, but it is included to provide a comparison. Each of the figures uses a reference light curve at one of the 12 subband frequencies, and is thus one of the twelve cross-correlation plots that was generated for that particular data set. The plots in Figure 5.6 were selected to show a representative range of behaviours encountered in the cross-correlation products: for data set 3 there is no obvious feature that correlates across all frequencies and total variability is limited. For data set 8 there is a strong climbing trend in flux density present in the light curves, with a maximum being reached for the higher frequencies but not as clearly for the lower ones. For data set 9, there is a clear and consistent minimum visible in the light curves at all frequencies. We see that localised features in the light curves that correlate across all bands tend to occur over timescales corresponding to one hour or more, matching the flaring timescale reported by Yusef-Zadeh et al. (2006). In the leftmost plot, we see a case where the light curves are only weakly correlated across the three main frequency bands (data set 3). The light curves in the same band as the reference light curve show a relatively strong correlation peak, but this pattern is abruptly broken for the correlation functions in the other two bands. The light curves for this epoch do not show clear trends or strong variability, and calibration uncertainties can therefore cause correlated flux density excursions that are limited to a single band as they are part of the same set of scans. In the middle plot, we see an example of a relatively broad cross-correlation function, where the location of its maximum along the lag axis is not tightly constrained (data set 8). The observed cross-correlation pattern may be compatible with the overplotted lag relation (which is taken from our previously found result in Brinkerink et al. (2015)), but it is also consistent with zero time lag – depending on the reference light curve used for cross-correlation, both of these trends can appear. The rightmost plot shows a case in which there is a robust and monotonic lag trend visible across all bands (data set 9), which appears to lie close to the lag relation we found in our previous paper. The associated light curves show a clearly defined minimum in their flux density evolutions, which improves the clarity of the cross-correlation peaks.

In order to distill potential time lag trends from the maxima in the cross-correlation graphs generated from the data, the choice was made to employ a piecewise 4th-order spline fitting to the calculated cross-correlation functions (taking their error bars into account) and to perform a local maximum search on this spline fit. This approach allowed us to avoid bias when identifying the

dominant peaks in the cross-correlation functions. We chose to reject any maximum at time lags greater than 3000 seconds or below -3000 seconds, as the cross-correlation functions have poor statistics at those largest lag values because of the small overlap between light curves at large lags, and as such are too noisy to use there. Our method was found to be robust in the context of the various shapes and degrees of smoothness that the cross-correlation functions can take across our data sets. Most importantly, this method is agnostic to any expected value of the time lag for any cross-correlation function. In the cases where no clear maximum can be identified, the peak finding algorithm returns an empty result. It should be noted that although the peak-finding method generally agrees quite well with a human assessment of the cross-correlation function, there are cases where the peak of the function remains poorly constrained. The typical causes of this behaviour are either a plateau-shaped correlation curve, for which the peak position is extremely sensitive to details of the fitting parameters, or a complex shape of the cross-correlation function with many local maxima - none of which is clearly dominant. In those cases, other heuristics for peak finding are likely to encounter the same issues.

For each choice of reference frequency, we thus get one 'trend' containing at most twelve different times at which the maximum value of the cross-correlation function is detected, one for each of the subbands in the data set. Using twelve different reference frequencies, we get twelve of these trends. As we cannot presuppose any temporal relation between the zero-lag autocorrelation peaks for these different trends (that, is after all, the very relation we are trying to determine), we need to apply global shifts to these trends so that we minimise their mutual average offsets while respecting their internal structure. Using a least-squares fit and picking one time lag trend as the reference trend, we obtain co-registered time lag trends that are ready for trend fitting.




FIGURE 5.6: Three examples of correlation products. Left: poor correlation between bands for data set 3, using reference band Ka2. Middle: wide correlation for data set 8, reference band Q1. Right: strong correlation with time lag for data set 9, reference band K2. The black curved line indicates the lag/frequency relation found in Brinkerink et al. (2015), and is not a fit to the correlation data shown here.

Upon initial inspection, the patterns seen in the time lag trends for the different data sets show considerable variance (see Figure 5.8 for the recovered time lags and Figure 5.9 for the trends




FIGURE 5.7: Fitted splines and detected cross-correlation peaks, for the same data sets and reference bands as shown in Figure 5.6 where cross-correlation functions for progressively higher frequency bands are vertically offset by steps of 1. Vertical black lines indicate the expected time lags based on the previously found lag relation from Brinkerink et al. (2015), while black dots indicate the detected maxima of the cross-correlation spline fits. Note that when the cross-correlation function has a broad plateau, the detected peak is poorly constrained (middle plot).

fitted to these measurements). We see weakly correlated lag trends showing large outlier values for three data sets (data sets 1, 2, and 3), one data set that has no clear trend pattern, because of poor cross-correlations (data set 6), a data set that shows consistent time lag trends in some subbands but not in others (data set 8) and data sets with trends having a positive slope that indicate the presence of time lags (data sets 4, 5 and 9).

We note that the data sets with cross-correlation functions that have more clearly defined peaks also show a better consistency among their individual time lag trends for the different reference subbands than the other data sets do. Furthermore no examples of a consistently present time lag trend with a negative slope, where low-frequency variability would lead high-frequency variability, is seen in any of the data sets.

5.3.3 Checks with synthetic data

While we do see clear time lags in several of our data sets, the question that immediately arises is why we do not always see them. The answer to this question comes down to three possibilities: 1) the time lags are always present due to the fundamental behaviour of the source, but the statistics of variability does not always let us detect them, 2) the time lags are intrinsic but do not account for all source variability (so that they only sometimes dominate the observed variability), or 3) they are pure chance occurrences that have no physical mechanism or pattern behind them. Despite us now having multiple data sets available, there is still room for different interpretations of these results.

To address this question and investigate the likelihood of these scenarios, we must try to assess the probability with which time lags are detected if they are indeed a fundamental property of the source, as well as the biases we are subject to. Conversely, we also need to assess how likely it




FIGURE 5.8: The co-registered maxima of the cross-correlation functions, collectively forming the observed time lag trends in the data sets. The trends have been shifted to minimise mutual spread.

is that we appear to see time lags in a data set where no such intrinsic relation between the light curves is present.




FIGURE 5.9: Linear time lag versus wavelength trends, fitted to the data shown in Figure 5.8.

To this end, we made use of the formalism of Gaussian processes as described in Roberts et al. (2013) to model the statistical properties of the variability of Sgr A* on the time scales we have access to with these light curves. Specifically, we construct a multidimensional Gaussian distribution with the appropriate covariances from which realisations of synthetic light curves are

drawn. To illustrate the fundamental thought behind this formalism, consider for example the task of constructing 10 samples from a light curve, not necessarily equally spaced in time. We know beforehand that these sampled flux density values should be correlated with one another to a certain degree. For one thing they should all have positive values, but more specifically we expect that flux density samples spaced closely together in time should be correlated more strongly than samples with a larger temporal separation between them. This degree of correlation is captured in the auto-correlation function of our measured light curves: using this relation we see how the degree of correlation between samples depends on their separation in time. We fit a functional form to this auto-correlation function and use it as a covariance kernel function. Using this kernel function, we then construct a covariance matrix that is valid for a given set of sampling times.

To make matters simpler, we consider our simulated light curve as a 'standard' function by shifting and scaling it to have zero mean, a variance of unity, and be fully characterised by its auto-correlation function. In this case, the autocorrelation and the covariance are identical. Our sampled values thus become a draw from a 10-dimensional joint Gaussian distribution, the shape of which is defined by the values in our covariance matrix. Every realisation of our example light curve is a single point in this 10-dimensional space. Many separate realisations of the light curve are calculated, and we then study the impact that our sampling cadence and limited epoch length have on the light curve time lag properties we derive.

In our case, we start by modeling the flux density of Sgr A* in a single subband as a function of time. As mentioned above, we assume the variability to be stochastic in nature, with a well-defined mean value and an auto-correlation function that does not change with time. With knowledge of this auto-correlation function we generate synthetic light curves that have, on average, the same auto-correlation function as our original data does. Taking a specific realisation we then clone this light curve for our 12 subbands, scale the copies appropriately in flux density and shift them in time according to the time lag relation we wish to investigate. These light curves are then sampled with the same cadence as our observed data sets, with thermal noise and estimated calibration errors added for each virtual measurement. The resulting synthetic data is treated in exactly the same way as our observed data: we correlate these light curves against one another and check if we see the time lag relations that we initially put in. The process by which we search for the appropriate auto-correlation function to use in generating the synthetic light curves is described in full in Appendix D.

With this replication of the behaviour of the measured light curves, we now investigate the detectability of time lag trends using different intrinsic lag trends in our synthetic data. To set up the synthetic sampled light curves for all 12 subbands, we first generate a 'master' light curve using the appropriate auto-correlation function. For each of the 12 subbands we add an additional time shift to this master light curve, according to a time lag relation we define. We then scale the synthetic light curve flux density values according to the average flux density for

the appropriate frequency band (see Figure 5.2). The resulting shifted, scaled light curve is then sampled following the scan setup for one epoch of the measured data sets, including the scan timing offsets that pertain to the subband in question (as is visible in Figure 5.3).




FIGURE 5.10: Top: the cumulative distributions of time lag slopes for synthetic light curves, for different input lag relations (blue: -42 min/cm, red: 0 min/cm, green: 42 min/cm), and different contributions of noise (indicated by shading). We see that lag slope distributions creep closer to zero for higher noise contributions, but that the largest recovered values stay close to the input slopes. Bottom: the cumulative distributions for only the 42 min/cm synthetic lag relation, plotted for different noise contributions and compared to the measured cumulative distribution.

In our analysis of the synthetic light curves we test three different time lag relations. The first one is the 'nominal' time lag relation which we found in Brinkerink et al. (2015), $dt/d\lambda = 42$

min/cm, where the variability at lower frequencies lags behind the higher frequencies. The second relation is the inverse of this, $dt/d\lambda = -42$ min/cm, as we wish to check if there is a bias introduced by the scan timing offsets between K, Ka and Q bands. The third relation uses zero time lag, where no shifts in time are applied to any of the light curves and the intrinsic variability is contemporaneous across all frequencies. For each time lag relation, we assess the influence of measurement errors. We compare the distribution of lag slopes we find when no noise is present to cases where the random uncorrelated variations in flux density (i.e. flux density variations that are realised independently per light curve) correspond to our estimated value of ~ 10 mJy on a scan-to-scan timescale, as well as to cases where this random variability component is larger by a factor of 3, 5 or 7. For the cases where we include noise, we also include a realistic calibration error due to the combined influence of unmodeled minor variations in the flux density of our gain calibrator and atmospheric effects. These errors are strongly correlated across each frequency band, meaning that all four subbands within one frequency band show coherent flux density offsets due to this calibration uncertainty. We estimate the magnitude of this calibration error to be ~ 30 mJy, from the cross-band trend breaks in our measured data (see Figures 5.3 and 5.5, as well as the check source light curves in Appendix A). For each combination of parameters (time lag relation, thermal noise strength) we generate 100 realisations from which we reconstruct the observed time lag slopes in the same way as was done for our measured data, and so we get a distribution of lag slopes for each choice of parameter combination. A sample of generated sets of light curves, as well as their cross-correlations and the recovered lag trends, are shown in Appendix B. An overview showing the distributions of the recovered time lag slopes for all considered cases is shown in Figure 5.10.

5.4 Results

In our observed data sets, we reconstruct strong time lag trends in two epochs with high internal consistency (39.6 ± 1.5 min/cm and 38.2 ± 1.4 min/cm respectively for data sets 5 and 9), another significant but less consistent time lag in one epoch (43.1 ± 4.0 min/cm, for data set 8), one weaker time lag in another epoch (18.9 ± 1.1 min/cm, for data set 4) and no clear or internally consistent results for the other epochs. Although the linear time lag fits for these other data sets show relatively high confidence in the trend uncertainties (small errors in the fitted trend slopes), we consider these reconstructions less reliable as the individual lag relations show very low internal consistency across all 12 subbands with a large residual scatter after their co-registration. Slope errors obtained from bootstrapping carry the implicit assumption that the data follow the presupposed linear relation, which is not clearly present in certain cases and manifests itself as a discrepancy between the bootstrapping errors and the spread between recovered lag trends.

We wish to understand how likely it is that a measurable time lag is detected in any given epoch for a source where a persistent time lag relation is present between the light curves at different frequencies. To this end, we generated synthetic data with the same general variability properties

as those that we observed in the measured light curves. As we see from Figure 5.10, the magnitude of the time lag trends from synthetic data tend to be underpredicted by the reconstructions, even in the case where zero-error measurements are modeled. The true lag slopes used as input for the synthetic light curves are plotted using vertical coloured bars in the left panel, and it is apparent that for the non-zero time lag relations the distribution of the corresponding recovered lag slope is such that the input lag slope value is close to an extremum of the distribution.

For both the tested positive and negative time lag trends, we see an asymmetric distribution in the recovered trends with longer tails towards zero time lag (and in a few cases even showing recovered lags with a sign opposite from the input lag relation). For the zero-lag case, we see a symmetrical spread in the recovered trends without any discernible bias to positive or negative slopes. Overprediction of the lag trend slope in the reconstructions, on the other hand, is much more rare and never extends far away from the input lag slope values. The asymmetry in reconstructed lag slope distributions is more severe when calibration errors are included, causing a larger bias towards smaller measured lag trend slopes. Thus, we see that a time lag relation that is consistently present in the data will tend to be underestimated from measurements using the cadence of our observations, and that among many measurements of the time lag relation the larger reconstructed slopes are a better estimate of the true time lag relation that is present. The set of synthetic light curves that best reproduces the distribution of time lag slopes we find in our observations incorporate a stronger uncorrelated noise component than our initial estimate (see the right panel of Figure 5.10) - an uncorrelated random flux density component with a standard deviation of ~ 30 mJy per scan (plotted as '3xthermal') gives the best match to the observed lag distribution. Whether this is due to an underestimation of measurement error or to an unmodeled random component of source variability is an open issue.

5.5 Discussion and Conclusions

In this paper, we have presented the most comprehensive and systematic study of time lags in Sagittarius A* radio observations performed to date. We find positive time lags (between 20 and 40 min/cm) over a frequency range from 47 down to 19 GHz, where low-frequency variability lags behind high-frequency variability, for five out of eight epochs in our data. Three out of eight epochs show time lags close to zero. The larger time lags are found for epochs where the light curves show more pronounced local maxima or minima.

In our synthetic data analysis, we find that the reconstructed time lag slopes are typically underestimated with the measurement cadence and length that we have used for our observations, particularly when variability is low and no recognisable flaring occurs. The distribution of recovered synthetic time lag slopes closely matches the observed distribution when a lag relation of 42 min/cm is assumed, together with a 30 mJy uncorrelated variability component.

The lag relation found in our previous paper (Brinkerink et al., 2015) was calculated from a set

of light curves that showed a clear temporal maximum for all of them, and we see comparable lag slopes for the light curves that have a clear flux density extremum here – both for the measured light curves and for the synthetically generated light curves. We are therefore confident that we have observed a persistent and consistent time lag in these results. The consistency of the time lag slope over multiple years of measurement (covering the time range from 2005 to 2015) indicates that we are sampling some kind of characteristic velocity, scale or expansion speed.




FIGURE 5.11: The aggregate time lag trend fitted to the measured time lags from data sets 5, 8 and 9 after co-registration.

Combining the time lag values found in the three epochs with the largest time lag slopes, which lie close together in value, gives us a fitted slope of 40.3 ± 2.0 min/cm, which is plotted in Figure 5.11.

When we consider the constraints on the inclination angle of a putative compact jet from Sgr A* that follow from recently published results, we see that small values are favoured: Issaoun et al. (2019) report good fits for jet models with $i \leq 20^\circ$, while Gravity Collaboration et al. (2020b) indicate that BH spin axis inclinations larger than 140° fit well for circular hotspot motion, with the value increasing to 175° for helical hotspot motion. Both of these scenarios correspond to a small angle between our line of sight and the jet direction pointing closest to us. Combining this information with the size-frequency relation for Sgr A* (Bower et al., 2006; Falcke et al., 2009) and our data, like we did in Brinkerink et al. (2015), results in an estimate for the associated outflow speed (see Figure 5.12). We find that the plasma flows out at moderately relativistic speeds with $\gamma\beta \approx 1.5$, which is well within the range predicted by both analytic jet models (Falcke & Biermann, 1996) and jets in GRMHD simulations (Mościbrodzka & Falcke, 2013). Here, the assumption is made that the variability we see at the various wavelengths is associated with the angular scales given by the size-frequency relation, and not limited to any smaller sub-region of the source. There is some uncertainty associated with this, as the size of the flaring component has not been directly measured. However, given that stronger flares can reach a significant fraction of the total flux density, it seems unlikely that they are emitted at scales widely separated from




FIGURE 5.12: Constant-velocity lead times w.r.t. the 19 GHz peak (calculated). Coloured curves show the expected time lag relations for an inclination of 20° when a size-frequency relation of $\theta[\text{mas}] = 0.52 \cdot \lambda[\text{cm}]^{1.3}$ is assumed. Data from this work and from Brinkerink et al. (2015) is overplotted.

the average structure probed by VLBI. The size-frequency relation of Sgr A* is well-constrained if a particular functional form is assumed, but shows a larger uncertainty once different forms of the relation are considered (Brinkerink et al., 2019).

From our investigation using synthetic data with a known intrinsic time lag trend, we see that there is a clear bias towards smaller reconstructed time lag slopes when the employed measurement strategy and data calibration quality are considered. For future observations, an observing strategy where light curves are sampled over longer time scales (i.e., over 5 or 6 hours) may aid in a more precise measurement and characterisation of these time lags, for two reasons: the larger number of data points will reduce the impact of spurious flux density variations from calibration uncertainties on the measured cross-correlation functions, and the likelihood of catching and fully tracing flux density excursions also increases with longer epochs.

The persistent time lag relation we observe between different radio observing frequencies provides another benchmark by which to test theoretical emission models for Sgr A*, next to the other properties of the source that have been taken into account in model fitting so far (spectral shape, variability strength, angular size and polarisation properties).

Acknowledgements

This work was supported by the ERC Synergy Grant 610058: "BlackHoleCam: Imaging the Event Horizon of a Black Hole". The author wishes to thank Federica Palla for her initial analysis of the data, which motivated this work.

Appendix A: Calibrator and check source light curves




FIGURE 5.13: Calibrated light curves for J1744-3116 (calibrator), data sets 1-6, 8, and 9. All data is plotted with error bars, which are obscured by the data markers.




FIGURE 5.14: Calibrated light curves for J1745-283 (check source), data sets 1-6, 8, and 9. All data is plotted with error bars.

Appendix B: synthetic light curve and correlation results




FIGURE 5.15: Samples of results from synthetic light curves with the inverse time lag relation and no calibration errors. The left column shows data where the inverse time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.




FIGURE 5.16: Samples of results from synthetic light curves with the inverse time lag relation and with calibration errors included. The left column shows data where the inverse time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.




FIGURE 5.17: Samples of results from synthetic light curves with the zero time lag relation and no calibration errors. The left column shows data where the zero time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.




FIGURE 5.18: Samples of results from synthetic light curves with the zero time lag relation and with calibration errors included. The left column shows data where the zero time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.




FIGURE 5.19: Samples of results from synthetic light curves with the nominal time lag relation and no calibration errors. The left column shows data where the nominal time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.




FIGURE 5.20: Samples of results from synthetic light curves with the nominal time lag relation and with calibration errors included. The left column shows data where the nominal time lag trend was successfully recovered from the cross-correlations, the right column shows an example of a case where a time lag relation with a significant error is produced.

Appendix C: historical spectral measurements for Sgr A*




FIGURE 5.21: The full set of historical data used for the binned plot in Figure 5.2. Measurements are from Zylka et al. (1995); Serabyn et al. (1997); Falcke et al. (1998); Zhao et al. (2003); Herrnstein et al. (2004); An et al. (2005); Marrone et al. (2006); Bower et al. (2015b); Brinkerink et al. (2015); Liu et al. (2016); Bower et al. (2019). Theoretical model spectra are from Falcke et al. (2000) and Yuan et al. (2003).

Appendix D: Finding the appropriate auto-correlation function for synthetic light curve data

The synthetic data pipeline starts by defining the form of the auto-correlation function that appears to hold for our measured light curve data. As a starting point, we look at the auto-correlation function we measure for our original light curve data, when we consider the aggregate of light curves from all data sets and all frequency bands. This measured auto-correlation function is depicted in Figure 5.22 using non-connected markers. To check for frequency dependence of the auto-correlation function, we compared the auto-correlation functions for the aggregate light curve data split into separate frequency bands against each other, and their shapes show no detectable frequency dependence outside the expected statistical variance.




FIGURE 5.22: The measured average autocorrelation function from all frequency bands in all data sets combined (separated markers), the functional form fitted to this autocorrelation function (continuous line), and the recovered autocorrelation function from a limited-length synthetic light curve generated using that functional form (connected markers). It is evident that the recovered auto-correlation differs from the prescribed form.

We approximate the shape of this auto-correlation function by using a linear combination of a Gaussian (squared exponential) function and a rational quadratic function, both functions that often feature in the description of auto-correlation functions, ensuring that data points separated by progressively larger time lags tend to an uncorrelated state. This function is plotted in Figure

5.22 with a continuous line. We find that its functional form is closely approximated by:

$$\text{Cov}(x) = 0.4 \cdot e^{-(x/1500)^2} + 0.6 \cdot \left(1 + \left(\frac{x}{1600}\right)^2\right)^{-2}, \quad (5.1)$$

where x is the lag time in seconds. The three free parameters that were varied to find the best fit are the strengths of the two components (with the constraint of summing to 1, which is the zero-lag auto-correlation) and the scalings of x in both terms. Using this auto-correlation function as the covariance function for our (normalised) Gaussian process, we generate densely sampled synthetic time series as large as are allowed by our computational resources (see Figure 5.23 for one realisation of such a light curve). These generated light curves are then shifted in time and re-scaled in amplitude to fit with the properties of the flux density for Sgr A* at various frequencies. We proceed to sample our synthetic data by applying the sampling cadence from the real-world measurements to a limited segment of this 'master' light curve.




FIGURE 5.23: An example of a synthetic, longer-term, densely sampled master light curve generated using the initial choice for covariance function (Figure 5.22). Note that the function values can be negative, as the light curve has not been shifted and scaled to the correct average and variance at this point.

We now check to see if we do indeed get auto-correlation functions from these sampled synthetic light curves that look like the auto-correlation function we expect. We plot the averaged recovered auto-correlation function from many realisations of the synthetic light curve, using the same sampling cadence as our originally measured data (Figure 5.22, connected markers) together with both the originally measured and the functionally prescribed auto-correlations, and we see

that the recovered auto-correlation function is significantly narrower and even dips into negative values (indicating anti-correlation) for a range of time separations. This happens because we measure the auto-correlation of the sampled light curve using the average flux density of the limited time series instead of the 'ideal', long-term flux density average. We therefore tend to see more pairs of points with values of opposite signs (above vs. below the average value) when calculating our auto-correlation function, causing it to have a more rapid falloff with increasing time difference than we would see for longer light curves where the average is closer to the 'true' average. Effectively, we miss a low-frequency component of the auto-correlation function as we only have access to the short timescales when measuring it.

The functional form of the original auto-correlation function to use when generating the synthetic light curve data therefore needs to be chosen in such a way that the recovered shape of this function is the same as for our real measured light curves. To this end, the behaviour of several types of auto-correlation functions was investigated, with the function yielding the best results having the form of a truncated hyperbolic function. Note that this was a heuristic process where we investigated various functional forms in order to closely approach the measured auto-correlation curve with our reconstructed one. The result of this iterative process is shown in Figures 5.24 and 5.25. The functional form of our final auto-correlation function is given by:

$$\text{Cov}(x) = \max \left(\frac{11}{10} - \sqrt{\left(\frac{x}{8000} \right)^2 + \frac{1}{100}}, 0 \right), \quad (5.2)$$

where x is again the lag time in seconds, and where the function value has a lower bound of zero imposed on it so that we avoid negative values for the auto-correlation function at large time lags. With this auto-correlation function, we generate synthetic data that has the appropriate temporal variability characteristics to mimic the behaviour of our measured light curves.




FIGURE 5.24: As Figure 5.22, but with an auto-correlation function chosen to optimise the match between the measured data and the synthetic data. Although the functional form of the auto-correlation function differs significantly from the measured one, the recovered function closely follows the measured function.




FIGURE 5.25: As Figure 5.23, but generated with the updated auto-correlation function from Figure 5.24.

SUMMARY

Black holes and their accretion flows present us with an opportunity to study physics in extreme environments with strong spacetime curvature. Accreting black hole systems exist over a wide range of masses and accretion rates, which can exhibit different behaviours but which fundamentally reflect the same type of process: the conversion of the gravitational potential energy of orbiting plasma into electromagnetic radiation and the kinetic energy of outflows. Accreting supermassive black holes are a particularly rich source of observational data: they can be studied over a wide range of wavelengths ranging from radio to gamma rays, and several supermassive black holes have a sufficiently large angular diameter on our sky to be resolved by our most powerful telescope systems. Sagittarius A*, the source associated with the supermassive black hole at the center of our Galaxy, is one of these rich laboratories offering us the chance to get a close-up look into a so-called low-luminosity active galactic nucleus.

However, the detail in which we can study Sgr A* also means that our models need to be sufficiently sophisticated to describe these details. As Sgr A* appears to be a black hole accreting at a low Eddington ratio, features such as a plasma jet are expected to manifest on scales close to the event horizon as they do for similar systems – but verifying the presence of a jet by measuring the source morphology is made challenging by the limits that interstellar scattering imposes on the attainable resolution of our observations. We therefore need to make use of high-frequency VLBI to discern the intrinsic size and shape of Sgr A*, and study its flux density variability across radio frequencies to gather clues about the nature of the plasma flow.

To this end, I have investigated Sgr A* by observing it using several different radio telescope networks and instruments: the Karl G. Jansky Very Large Array (VLA), the High Sensitivity Array (HSA, part of which is formed by the Very Long Baseline Array or VLBA) and the Atacama Large Millimeter Array (ALMA).

We use HSA+LMT 86 GHz observations made in spring 2015 to investigate the asymmetry in the compact structure of Sgr A* in Chapter 2 of this thesis. The array includes two stations with extra sensitivity: the Green Bank Telescope (GBT) and the Large Millimeter Telescope (LMT).

The combination of this high sensitivity with the extended (u, v) -coverage gives this observation an enhanced North-South resolving power over previous observations. Through an analysis of closure phases, we find that there is evidence for source asymmetry. Using a simple model fit, we find that a faint Eastern component approximately $100 \mu\text{as}$ from the main component provides the best match with our observed closure phases. While it cannot be established from this single epoch whether the observed asymmetry is intrinsic or caused by interstellar scattering, the detection of East-West asymmetry at 230 GHz and 43 GHz in other observations of Sgr A* provides the tantalizing suggestion that we may be seeing intrinsic structure.

Chapter 3 presents further results from the spring 2015 HSA+LMT observations, where this time the focus is placed on an imaging analysis and on size measurement using a closure amplitude analysis. The imaging shows again a slight East-West asymmetry in Sgr A*, confirming the simple two-component model fitting result from Chapter 3. From the closure amplitude analysis we find an observed source geometry that is close to Gaussian, with a size of $215.1 \pm 0.4 \times 145.1 \pm 1.5 \mu\text{as}$ at a position angle of 77.9 ± 0.4 degrees East of North. Deconvolving this with the scattering law from (Bower et al., 2006), we find an intrinsic size of $145.4 \pm 0.6 \times 122.6 \pm 1.7 \mu\text{as}$. Collating our new measured size with earlier measurements performed over a range of wavelengths allows us to revisit the intrinsic size-frequency relation.

In Chapter 4, we present an analysis of the variability we see in simultaneous observations of Sgr A* performed in spring 2012 using ALMA and the VLA. During a single epoch, we see a minor flare which manifests itself at different time delays across different frequencies, from 100 GHz down to 19 GHz. We recover the relation of this time delay as a function of wavelength using fast rise exponential decay (FRED) curve fitting and see that the variability at longer observing wavelengths exhibits progressively longer time delays according to the relation $d\tau/d\lambda = 42 \pm 14 \text{ min/cm}$. This is in line with earlier observations that showed time lags between 43 GHz and 22 GHz of a similar slope. Sgr A* also shows a size that depends on observing frequency: at longer wavelengths (above 3.5 mm) its apparent size is dominated by interstellar scattering effects, and its angular size scales with wavelength squared. With source size measurement done at shorter wavelengths, the intrinsic size of Sgr A* can be modeled as a function of wavelength. Using this size-frequency relation and coupling it to the observed time lags, we find that a self-consistent jet model with moderately relativistic outflow speeds best matches the observed time lag behaviour.

We revisit the study of time lags in Chapter 5, where we present multi-epoch observations of Sgr A* with VLA over frequencies ranging from 18 to 47 GHz. We find time lags corresponding to our earlier result from Chapter 4 in three of the epochs, smaller but positive time lags in three more epochs and negative time lags close to zero in the remaining two. The distribution of measured time lag slopes is compared to results from synthetic light curves, where a pre-determined time lag relation is used as input and simulated measurements are taken according to the sampling cadence of our real measurements. We find that the distribution of time lag slopes from our measured data

is consistent with a persistently present intrinsic time lag relation. Combining the time lag slopes that lie closest to the expected intrinsic time lag yields an aggregate time lag trend of 40.3 ± 2.0 min/cm. Combining this result with recent constraints on the orientation of the spin axis of Sgr A* and the size-frequency relation, we again find that moderately relativistic outflow speeds with $\gamma\beta \approx 1.5$ best fit the measured time lags. With the persistence of time lags in radio observations of Sgr A* receiving strong support from these results, we state that theoretical models for the accretion flow will need to reproduce these time lags in their predicted properties along with the pre-existing suite of constraints on flux density, spectral shape and source morphology.

SAMENVATTING

Zwarte gaten en hun accretiestromen bieden ons een uitgelezen kans om de natuurkunde van processen in gebieden met een sterke ruimte-tijdkromming te bestuderen. Accreterende zwarte gaten komen voor verspreid over een breed bereik aan centrale massa's en massastromen. Deze systemen vertonen uiterlijk verschillend gedrag, maar vertegenwoordigen allemaal fundamenteel hetzelfde soort proces: de conversie van de gravitationele potentiële energie van baanbeschrijvend plasma naar elektromagnetische straling en de kinetische energie van de uitstromende fractie van deze materie. Een bijzonder rijke bron van waarnemingen wordt gevormd door relatief nabije supermassieve accreterende zwarte gaten: deze zenden doorgaans elektromagnetische straling uit over een zeer breed gedeelte van het spectrum, van radio tot gammastraling, en enkele van hen hebben een hoekgrootte aan onze hemel die ons in staat stelt om ze op te lossen met onze meest geavanceerde (interferometrische) telescoopsystemen. Sagittarius A* (afgekort: Sgr A*), de radiobron die geassocieerd is met het centrale supermassieve zwarte gat in het centrum van onze Melkweg, is zo'n systeem dat ons de kans biedt om van relatief dichtbij de kern van een actief sterrenstelsel met lage helderheid (Low-Luminosity Active Galactic Nucleus, LLAGN) in detail te bestuderen.

De vele nauwkeurige metingen die we aan Sgr A* kunnen doen houden echter ook in dat onze modellen strenger getoetst worden en daarom meer gedetailleerde voorspellingen moeten kunnen geven. Aangezien Sgr A* een zwart gat lijkt te zijn dat slechts op een laag tempo gas accreteert in verhouding tot de Eddington-helderheid, verwachten we kenmerken te zien die voor andere systemen in dezelfde klasse typisch zichtbaar zijn: bijvoorbeeld een snelle straalstroom van uitgeworpen plasma die van vlak buiten de waarnemingshorizon gelanceerd wordt. Het waarnemen van een dergelijke straalstroom wordt echter bemoeilijkt doordat het scheidend vermogen van onze waarnemingen op de relevante frequenties beperkt wordt door de verstrooiing van de uitgezonden straling via interstellair gas. Om de aard van de plasmaprocessen rond Sgr A* goed te kunnen gebrijpen moeten we ons daarom richten op interferometrische waarnemingen op kortere golflengtes, waar dit verstrooiend effect kleiner is, en op de informatie die we kunnen bemachtigen door de variabiliteit van Sgr A* over verschillende radiofrequenties te bestuderen.

Met dit doel heb ik Sgr A* bestudeerd door het systeem te observeren met verschillende radio-telescopen en netwerken van telescopen: de Karl G. Jansky Very Large Array (VLA), de High Sensitivity Array (HSA, met als onderdeel hiervan de Very Long Baseline Array oftewel VLBA), en de Atacama Large Millimeter Array (ALMA).

In Hoofdstuk 2 van dit proefschrift presenteren we waarnemingen op 86 GHz gedaan met de HSA+LMT in de lente van 2015, waarmee we de asymmetrie in de compacte structuur van Sgr A* onderzoeken. Het gebruikte netwerk van telescopen heeft twee extra gevoelige stations: de Green Bank Telescope (GBT) en de Large Millimeter Telescope (LMT). De combinatie van de extra (u, v) -dekking die deze stations toevoegen met hun hoge gevoeligheid betekent dat deze waarnemingen een verbeterd oplossingsvermogen hebben in de Noord-Zuid richting ten opzichte van eerdere waarnemingen. Middels een analyse waarbij we de fasesom over driehoekscombinaties van stations gebruiken (zogenoeten closure phases) zien we aanwijzingen voor een spatiële asymmetrie in de helderheidsverdeling van de bron. Via het fitten van een eenvoudig model voor deze helderheidsverdeling concluderen we dat een zwakke extra component op ongeveer 100 microboogseconden van de centrale component de best passende fasesommen geeft met onze waarnemingen. Hoewel we hieruit nog niet kunnen concluderen wat de oorzaak is van deze waargenomen asymmetrie (deze is ofwel intrinsiek aan de bron, ofwel het gevolg van de interstellaire verstrooiing) suggereert de asymmetrie op 230 GHz en 43 GHz die gevonden is middels andere waarnemingen dat hier sprake kan zijn van de intrinsieke structuur van de bron.

Een verdere analyse van de HSA+LMT waarnemingen van 2015 wordt gepresenteerd in Hoofdstuk 3. Deze keer ligt de nadruk op het in kaart brengen van de structuur middels beeldvormingsalgoritmen en het precies meten van de hoekgrootte van de bron via de analyse van *closure amplitudes*. De resultaten van de beeldvormingsalgoritmen bevestigen onze conclusies op basis van de fasesommen uit Hoofdstuk 3. De *closure amplitude* analyse laat zien dat de waargenomen structuur van Sgr A* verder nagenoeg Gaussisch is, met een grootte van $215.1 \pm 0.4 \times 145.1 \pm 1.5 \mu\text{as}$ en de lange as onder een positiehoek van 77.9 ± 0.4 graden oostelijk van noordelijke richting. Door deze goottemeting te combineren met eerdere resultaten op verschillende andere frequenties kunnen we het verband tussen hoekgrootte van de bron en waarneemfrequentie aan een kritische blik onderwerpen.

In Hoofdstuk 4 behandelen we een analyse van de variabiliteit van Sgr A* over een bereik van verschillende golflengten, gebruik makende van gelijktijdige waarnemingen die met de VLA en ALMA gedaan zijn in de lente van 2012. In de gemeten variabiliteitscurves (die de variaties in helderheid als functie van de tijd beschrijven) zien we een enkele kortdurende helderheidsverhoging binnen ons waarneemvenster, welke piekt op steeds latere tijden in de steeds lagere frequentiebanden waarop onze metingen gedaan zijn, van 100 GHz tot 19 GHz. De relatie tussen deze tijdsvertraging en de golflengte waarop de waarnemingen gedaan zijn vatten we door gebruik te maken van een *fast rise, exponential decay* (FRED) curve die aan de data gefit wordt.

Hiermee vinden we dat de variabiliteit op langere golflengten vertraagd is ten opzichte van die op kortere golflengten via de relatie $d\tau/d\lambda = 42 \pm 14$ min/cm. Deze conclusie strookt met eerder gepubliceerde resultaten, waarin andere waarnemingen een tijdsvertraging tussen 43 en 22 GHz lieten zien met een vergelijkbare vertraging. De schijnbare grootte van Sgr A* aan de hemel is afhankelijk van de golflengte waarop we de bron waarnemen: op langere golflengten (boven de 3.5 mm) wordt de schijnbare grootte in toenemende mate gedomineerd door de effecten van interstellair verstrooiing. Met metingen van de grootte van Sgr A* op kortere golflengten kan de intrinsieke grootte als functie van golflengte af worden geleid, en in combinatie met de metingen van de tijdsvertragingen kunnen we snelheden berekenen die bereikt worden door uitstromend gas. Op deze manier vinden we dat een eenvoudig intern consistent straalstroommodel met matig relativistische uitstroombnelheden goed aansluit op de waargenomen tijdsvertragingen.

Een uitgebreidere analyse van de variabiliteit van Sgr A* op radio-golflengten is het onderwerp van Hoofdstuk 5, waarin data uit een set van acht waarneemsessies van Sgr A* met de VLA (uitgevoerd in de loop van de lente van 2015) wordt bestudeerd. Deze waarnemingen bestrijken frequenties van 18 tot 47 GHz, en in de variabiliteitscurves gemeten over deze verschillende frequenties vinden we tijdsvertragingen zoals beschreven in Hoofdstuk 4 in drie van de sessies. In drie verdere sessies vinden we tijdsvertragingen van een gelijksoortig karakter maar met kleinere waarden, en in de resterende twee sessies vinden we tijdsvertragingen met het omgekeerde teken maar met zeer kleine waarden. Deze verdeling van gemeten tijdsvertragingsrelaties vergelijken we met resultaten van gesimuleerde variabiliteitscurves, waarin we variabiliteitscurves met een vooraf bepaalde tijdsvertragingsrelatie genereren en nagaan in welke mate deze tijdsvertragingen gemeten kunnen worden volgens de tijdsstructuur van onze daadwerkelijk uitgevoerde waarnemingen. Deze vergelijking wijst uit dat de verdeling van tijdsvertragingsrelaties die we in onze metingen zien consistent is met een vaste tijdsvertragingsrelatie in alle waarneemvensters die in onze reconstructies varieert vanwege de statistische eigenschappen van de variabiliteit. Een combinatie van de drie tijdsvertragingsrelaties met de grootste gereconstrueerde waarden, die allen zeer dicht bij de best passende berekende tijdsvertragingsrelatie liggen, geeft een relatie van 40.3 ± 2.0 min/cm. Wanneer we dit resultaat combineren met recentelijk gevonden limieten op de oriëntatie van de spin-as van Sgr A* en met de relatie tussen de hoekgrootte van Sgr A* en de observationele golflengte, vinden we opnieuw dat uitstroombnelheden met $\gamma\beta \approx 1.5$ het beste stroken met onze resultaten. Met het versterken van het beeld dat Sgr A* een consistente tijdsvertragingsrelatie laat zien middels dit werk stellen wij dat dit een aspect van de plasmastroming is dat correct voorspeld dient te worden door theoretische modellen, naast de bestaande vereisten van fluxdichtheid, spectrale verdeling en morfologie van de bron.

RESEARCH DATA MANAGEMENT

This thesis research has been carried out under the institute research data management policy of the Institute for Mathematics, Astrophysics and Particle Physics Research at Radboud University. The RDM policies for all departments and institutes of the University can be found at <http://www.ru.nl/rdm/vm/policy-documents/>.

All observational data sets used in this thesis are available online in the long-term data archives corresponding to the relevant organisations. The HSA data used in Chapters 2 and 3 can be found in the NRAO Science Data Archive (<https://archive.nrao.edu/>), under project code **BF114**. The two epochs indicated with A and B are identifiable from their file names. The VLA data used in Chapter 4 is also available from the NRAO Science Data Archive, under project code **12A-339**. The ALMA data from chapter 4 has project code **2011.0.00887.S**, and is available from the ALMA Science Archive at <https://almascience.nrao.edu/aq/>. Chapter 5 again uses VLA data that can be found in the NRAO Science Data Archive, under project code **15A-372**.

BIBLIOGRAPHY

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *Phys. Rev. Lett.*, 116, 061102
- Abramowicz, M., Jaroszynski, M., & Sikora, M. 1978, *A&A*, 63, 221
- Abramowicz, M. A. & Fragile, P. C. 2013, *Living Reviews in Relativity*, 16, 1
- Abuter, R., Amorim, A., Bauböck, M., et al. 2020, *A&A*, 638, A2
- Agol, E. 2000, *ApJ*, 538, L121
- Aitken, D. K., Greaves, J., Chrysostomou, A., et al. 2000, *ApJ*, 534, L173
- Akiyama, K., Kuramochi, K., Ikeda, S., et al. 2017, *ApJ*, 838, 1
- Alexander, T. 1997, in *Astrophysics and Space Science Library*, Vol. 218, *Astronomical Time Series*, eds. D. Maoz, A. Sternberg, & E. M. Leibowitz, 163
- An, T., Baan, W. A., Wang, J.-Y., Wang, Y., & Hong, X.-Y. 2013, *MNRAS*, 434, 3487
- An, T., Goss, W. M., Zhao, J.-H., et al. 2005, *ApJ*, 634, L49
- Backer, D. C. 1978, *ApJ*, 222, L9
- Balick, B. & Brown, R. L. 1974, *ApJ*, 194, 265
- Becklin, E. E. & Neugebauer, G. 1968, *ApJ*, 151, 145
- Begelman, M. C. 1978, *MNRAS*, 184, 53
- Belloni, T., Homan, J., Casella, P., et al. 2005, *A&A*, 440, 207
- Bhat, P. N., Fishman, G. J., Meegan, C. A., et al. 1994, *ApJ*, 426, 604
- Bietenholz, M. F., Bartel, N., & Rupen, M. P. 2004, *ApJ*, 615, 173
- Blackburn, L., Pesce, D. W., Johnson, M. D., et al. 2019, *arXiv e-prints*, arXiv:1910.02062
- Blaes, O. 2007, *Astronomical Society of the Pacific Conference Series*, Vol. 373, *Accretion Disks in AGNs (Astronomical Society of the Pacific)*, 75
- Blandford, R. D. & Königl, A. 1979, *ApJ*, 232, 34
- Blandford, R. D. & Payne, D. G. 1982, *MNRAS*, 199, 883
- Blandford, R. D. & Znajek, R. L. 1977, *MNRAS*, 179, 433
- Bower, G. C. 2006, in *Journal of Physics Conference Series*, Vol. 54, *Journal of Physics Conference Series*, 370–376
- Bower, G. C. & Backer, D. C. 1998, *ApJ*, 496, L97
- Bower, G. C., Backer, D. C., Wright, M., et al. 1997, *ApJ*, 484, 118
- Bower, G. C., Backer, D. C., Zhao, J.-H., Goss, M., & Falcke, H. 1999a, *ApJ*, 521, 582

- Bower, G. C., Deller, A., Demorest, P., et al. 2014, *ApJ*, 780, L2
- Bower, G. C., Deller, A., Demorest, P., et al. 2015a, *ApJ*, 798, 120
- Bower, G. C., Dexter, J., Asada, K., et al. 2019, *ApJ*, 881, L2
- Bower, G. C., Falcke, H., Herrnstein, R. M., et al. 2004, *Science*, 304, 704
- Bower, G. C., Falcke, H., Sault, R. J., & Backer, D. C. 2002, *ApJ*, 571, 843
- Bower, G. C., Falcke, H., Wright, M. C., & Backer, D. C. 2005, *ApJ*, 618, L29
- Bower, G. C., Goss, W. M., Falcke, H., Backer, D. C., & Lithwick, Y. 2006, *ApJ*, 648, L127
- Bower, G. C., Markoff, S., Dexter, J., et al. 2015b, *ApJ*, 802, 69
- Bower, G. C., Wright, M. C. H., Backer, D. C., & Falcke, H. 1999b, *ApJ*, 527, 851
- Bower, G. C., Wright, M. C. H., Falcke, H., & Backer, D. C. 2003, *ApJ*, 588, 331
- Brinkerink, C. D., Falcke, H., Law, C. J., et al. 2015, *A&A*, 576, A41
- Brinkerink, C. D., Müller, C., Falcke, H., et al. 2016, *MNRAS*, 462, 1382
- Brinkerink, C. D., Müller, C., Falcke, H. D., et al. 2019, *A&A*, 621, A119
- Broderick, A. E., Fish, V. L., Johnson, M. D., et al. 2016, *ApJ*, 820, 137
- Chael, A. A., Johnson, M. D., Bouman, K. L., et al. 2018, *ApJ*, 857, 23
- Clagett, M., eds. 1968, *Nicole Oresme and the Medieval Geometry of Qualities and Motions* (Madison: University of Wisconsin Press)
- Cornwell, T. J. & Evans, K. F. 1985, *A&A*, 143, 77
- Davelaar, J., Mościbrodzka, M., Bronzwaer, T., & Falcke, H. 2018, *A&A*, 612, A34
- Davies, R. D., Walsh, D., & Booth, R. S. 1976, *MNRAS*, 177, 319
- Dexter, J., Kelly, B., Bower, G. C., et al. 2014, *MNRAS*, 442, 2797
- Dibi, S., Markoff, S., Belmont, R., et al. 2013, in *AAS/High Energy Astrophysics Division*, Vol. 13, AAS/High Energy Astrophysics Division, 108.14
- Dodds-Eden, K., Gillessen, S., Fritz, T. K., et al. 2011, *ApJ*, 728, 37
- Dodds-Eden, K., Sharma, P., Quataert, E., et al. 2010, *ApJ*, 725, 450
- Doeleman, S. S., Shen, Z.-Q., Rogers, A. E. E., et al. 2001, *AJ*, 121, 2610
- Doeleman, S. S., Weintroub, J., Rogers, A. E. E., et al. 2008, *Nature*, 455, 78
- Downes, D. & Martin, A. H. M. 1971, *Nature*, 233, 112
- Eddington, A. S. 1920, *The Observatory*, 43, 341
- Einstein, A. 1915, *Preussische Akademie der Wissenschaften, Sitzungsberichte*, 844
- Eisenhauer, F., Schödel, R., Genzel, R., et al. 2003, *ApJ*, 597, L121
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, *ApJ*, 875, L1
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019b, *ApJ*, 875, L5
- Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019c, *ApJ*, 875, L2
- Falcke, H. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 186, *The Central Parsecs of the Galaxy*, eds. H. Falcke, A. Cotera, W. J. Duschl, F. Melia, & M. J. Rieke, 113
- Falcke, H. & Biermann, P. L. 1995, *A&A*, 293, 665
- Falcke, H. & Biermann, P. L. 1996, *A&A*, 308, 321
- Falcke, H., Goss, W. M., Matsuo, H., et al. 1998, *ApJ*, 499, 731
- Falcke, H., Körding, E., & Markoff, S. 2004, *A&A*, 414, 895

- Falcke, H., Mannheim, K., & Biermann, P. L. 1993, AAP, 278, L1
- Falcke, H., Markoff, S., & Bower, G. C. 2009, A&A, 496, 77
- Falcke, H. & Markoff, S. B. 2013, Classical and Quantum Gravity, 30, 244003
- Falcke, H., Melia, F., & Agol, E. 2000, ApJ, 528, L13
- Fazio, G. G., Hora, J. L., Witzel, G., et al. 2018, ApJ, 864, 58
- Fender, R. P., Belloni, T. M., & Gallo, E. 2004, MNRAS, 355, 1105
- Finch, T. K. 2015, General Relativity and Gravitation, 47, 56
- Finkelstein, D. 1958, Physical Review, 110, 965
- Fish, V. L., Doeleman, S. S., Beaudoin, C., et al. 2011, ApJ, 727, L36
- Fish, V. L., Johnson, M. D., Doeleman, S. S., et al. 2016, ApJ, 820, 90
- Fraga-Encinas, R., Mościbrodzka, M., Brinkerink, C., & Falcke, H. 2016, A&A, 588, A57
- Frank, J., King, A., & Raine, D. J. 2002, Accretion Power in Astrophysics: Third Edition (Cambridge University Press)
- Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, Reviews of Modern Physics, 82, 3121
- Ghez, A. M., Salim, S., Weinberg, N. N., et al. 2008, ApJ, 689, 1044
- Gillessen, S., Eisenhauer, F., Quataert, E., et al. 2006, in Journal of Physics Conference Series, Vol. 54, Journal of Physics Conference Series, 411–419
- Gillessen, S., Eisenhauer, F., Trippe, S., et al. 2009, ApJ, 692, 1075
- Goddi, C., Falcke, H., Kramer, M., et al. 2017, International Journal of Modern Physics D, 26, 1730001
- Gold, R., McKinney, J. C., Johnson, M. D., & Doeleman, S. S. 2017, ApJ, 837, 180
- Goldston, J. E., Quataert, E., & Igumenshchev, I. V. 2005, ApJ, 621, 785
- Goss, W. M., Brown, R. L., & Lo, K. Y. 2003, Astronomische Nachrichten Supplement, 324, 497
- Gravity Collaboration, Abuter, R., Amorim, A., et al. 2018a, A&A, 615, L15
- Gravity Collaboration, Abuter, R., Amorim, A., et al. 2018b, A&A, 618, L10
- Gravity Collaboration, Abuter, R., Amorim, A., et al. 2019, A&A, 625, L10
- Gravity Collaboration, Abuter, R., Amorim, A., et al. 2020a, A&A, 636, L5
- Gravity Collaboration, Bauböck, M., Dexter, J., et al. 2020b, A&A, 635, A143
- Greisen, E. W. 2003, in Astrophysics and Space Science Library, Vol. 285, Astrophysics and Space Science Library, eds. A. Heck, 109
- Gwinn, C. R., Kovalev, Y. Y., Johnson, M. D., & Soglasnov, V. A. 2014, ApJ, 794, L14
- Hada, K., Doi, A., Kino, M., et al. 2011, Nature, 477, 185
- Herrnstein, R. M., Zhao, J.-H., Bower, G. C., & Goss, W. M. 2004, The Astronomical Journal, 127, 3399
- Högbom, J. A. 1974, A&AS, 15, 417
- Ichimaru, S. 1977, ApJ, 214, 840
- Igumenshchev, I. V., Narayan, R., & Abramowicz, M. A. 2003, ApJ, 592, 1042
- Issaoun, S., Johnson, M. D., Blackburn, L., et al. 2019, ApJ, 871, 30
- Jafari, A. 2019, arXiv e-prints, arXiv:1904.09677
- Jansky, K. G. 1933, Nature, 132, 66

- Jennison, R. C. 1958, *MNRAS*, 118, 276
- Johnson, M. D. 2016, *ApJ*, 833, 74
- Johnson, M. D., Fish, V. L., Doeleman, S. S., et al. 2015, *Science*, 350, 1242
- Johnson, M. D. & Gwinn, C. R. 2015, *ApJ*, 805, 180
- Johnson, M. D., Narayan, R., Psaltis, D., et al. 2018, *ApJ*, 865, 104
- Kawaguchi, T. 2004, *Progress of Theoretical Physics Supplement*, 155, 120
- Kormendy, J. & Richstone, D. 1995, *ARA&A*, 33, 581
- Krichbaum, T. P., Graham, D. A., Witzel, A., et al. 1998, *A&A*, 335, L106
- Law, C. J., Yusef-Zadeh, F., Cotton, W. D., & Maddalena, R. J. 2008, *ApJS*, 177, 255
- Lemaître, G. 1933, *Annales de la Sociéé Scientifique de Bruxelles*, 53, 51
- Li, Z., Morris, M. R., & Baganoff, F. K. 2013, *ApJ*, 779, 154
- Liu, H. B., Wright, M. C. H., Zhao, J.-H., et al. 2016, *A&A*, 593, A107
- Lo, K. Y., Cohen, M. H., Readhead, A. S. C., & Backer, D. C. 1981, *ApJ*, 249, 504
- Lo, K. Y., Shen, Z., Zhao, J. H., & Ho, P. T. P. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 186, *The Central Parsecs of the Galaxy*, eds. H. Falcke, A. Cotera, W. J. Duschl, F. Melia, & M. J. Rieke, 72
- Lo, K. Y., Shen, Z.-Q., Zhao, J.-H., & Ho, P. T. P. 1998, *ApJ*, 508, L61
- Lu, R.-S., Krichbaum, T. P., Eckart, A., et al. 2011a, *A&A*, 525, A76
- Lu, R.-S., Krichbaum, T. P., Roy, A. L., et al. 2018, *ApJ*, 859, 60
- Lu, R.-S., Krichbaum, T. P., & Zensus, J. A. 2011b, *MNRAS*, 418, 2260
- Macquart, J.-P. & Bower, G. C. 2006, *ApJ*, 641, 302
- Macquart, J.-P., Bower, G. C., Wright, M. C. H., Backer, D. C., & Falcke, H. 2006, *ApJ*, 646, L111
- Mahadevan, R. 1998, *Nature*, 394, 651
- Markoff, S. 2005, *Ap&SS*, 300, 189
- Markoff, S., Bower, G. C., & Falcke, H. 2007, *MNRAS*, 379, 1519
- Markoff, S., Falcke, H., Yuan, F., & Biermann, P. L. 2001, *A&A*, 379, L13
- Marrone, D. P., Moran, J. M., Zhao, J.-H., & Rao, R. 2006, *ApJ*, 640, 308
- Marrone, D. P., Moran, J. M., Zhao, J.-H., & Rao, R. 2007, *ApJ*, 654, L57
- Martí-Vidal, I., Krichbaum, T. P., Marscher, A., et al. 2012, *A&A*, 542, A107
- Melia, F. & Falcke, H. 2001, *ARA&A*, 39, 309
- Merloni, A., Heinz, S., & di Matteo, T. 2003, *MNRAS*, 345, 1057
- Mineshige, S. & Ohsuga, K. 2007, *Astronomical Society of the Pacific Conference Series*, Vol. 373, *Supercritical Accretion Flow? (Astronomical Society of the Pacific)*, 85
- Misner, C. W., Thorne, K. S., & Wheeler, J. A. 1973, *Gravitation* (W.H. Freeman; Princeton University Press)
- Miyazaki, A., Tsuboi, M., & Tsutsumi, T. 2013, *PASJ*, 65, L6
- Mościbrodzka, M. 2017, in *IAU Symposium*, Vol. 322, *The Multi-Messenger Astrophysics of the Galactic Centre*, eds. R. M. Crocker, S. N. Longmore, & G. V. Bicknell, 43–49
- Mościbrodzka, M. & Falcke, H. 2013, *A&A*, 559, L3

- Mościbrodzka, M., Falcke, H., Shiokawa, H., & Gammie, C. F. 2014, *A&A*, 570, A7
- Mościbrodzka, M., Gammie, C. F., Dolence, J. C., Shiokawa, H., & Leung, P. K. 2009, *ApJ*, 706, 497
- Muñoz, D. J., Marrone, D. P., Moran, J. M., & Rao, R. 2012, *ApJ*, 745, 115
- Narayan, R., Sądowski, A., Penna, R. F., & Kulkarni, A. K. 2012, *MNRAS*, 426, 3241
- Narayan, R., Yi, I., & Mahadevan, R. 1995, *Nature*, 374, 623
- Novikov, I. D. & Thorne, K. S. 1973, in *Black Holes (Les Astres Occlus)*, 343–450
- Oda, M., Gorenstein, P., Gursky, H., et al. 1971, *ApJ*, 166, L1
- Ohsuga, K. & Mineshige, S. 2007, *ApJ*, 670, 1283
- Ohsuga, K. & Mineshige, S. 2011, *ApJ*, 736, 2
- Orosz, J. A., McClintock, J. E., Aufdenberg, J. P., et al. 2011, *The Astrophysical Journal*, 742, 84
- Ortiz-León, G. N., Johnson, M. D., Doeleman, S. S., et al. 2016, *ApJ*, 824, 40
- Özel, F., Psaltis, D., & Narayan, R. 2000, *ApJ*, 541, 234
- Penrose, R. 1965, *Phys. Rev. Lett.*, 14, 57
- Plotkin, R. M., Markoff, S., Kelly, B. C., Körding, E., & Anderson, S. F. 2012, *MNRAS*, 419, 267
- Psaltis, D., Özel, F., Chan, C.-K., & Marrone, D. P. 2015, *ApJ*, 814, 115
- Pudritz, R. E., Hardcastle, M. J., & Gabuzda, D. C. 2012, *Space Sci. Rev.*, 169, 27
- Quataert, E. & Gruzinov, A. 2000, *ApJ*, 545, 842
- Rauch, C., Ros, E., Krichbaum, T. P., et al. 2016, *A&A*, 587, A37
- Rees, M. J., Begelman, M. C., Blandford, R. D., & Phinney, E. S. 1982, *Nature*, 295, 17
- Reid, M. J. 2009, *International Journal of Modern Physics D*, 18, 889
- Reid, M. J. & Brunthaler, A. 2004, *ApJ*, 616, 872
- Retana-Montenegro, E. & Röttgering, H. J. A. 2017, *A&A*, 600, A97
- Roberts, S., Osborne, M., Ebdon, M., et al. 2013, *Philosophical Transactions of the Royal Society* (Part A
- Rogers, A. E. E., Doeleman, S. S., & Moran, J. M. 1995, *AJ*, 109, 1391
- Rothschild, R. E., Boldt, E. A., Holt, S. S., & Serlemitsos, P. J. 1974, *ApJ*, 189, L13
- Rybicki, G. B. & Lightman, A. P. 1979, *Radiative processes in astrophysics* (Wiley-VCH)
- Ryu, S. G., Nobukawa, M., Nakashima, S., et al. 2013, *PASJ*, 65, 33
- Salpeter, E. E. 1964, *ApJ*, 140, 796
- Sbarrato, T., Padovani, P., & Ghisellini, G. 2014, *Monthly Notices of the Royal Astronomical Society*, 445, 81–92
- Schawinski, K., Koss, M., Berney, S., & Sartori, L. F. 2015, *MNRAS*, 451, 2517
- Schmidt, M. 1963, *Nature*, 197, 1040
- Semenov, V., Dyadechkin, S., & Punsly, B. 2004, *Science*, 305, 978
- Serabyn, E., Carlstrom, J., Lay, O., et al. 1997, *ApJ*, 490, L77
- Seyfert, C. K. 1943, *ApJ*, 97, 28
- Shakura, N. I. & Sunyaev, R. A. 1973, *A&A*, 500, 33
- Shapiro, S. L., Lightman, A. P., & Eardley, D. M. 1976, *ApJ*, 204, 187

- Shen, Z.-Q. 2006, *Journal of Physics Conference Series*, 54, 377
- Shen, Z.-Q., Lo, K. Y., Liang, M. C., Ho, P. T. P., & Zhao, J. H. 2005, *Nature*, 438, 62
- Shepherd, M. C. 1997, in *Astronomical Data Analysis Software and Systems VI*, eds. G. Hunt & H. Payne, ASP Conf. Proc. 125, 77
- Sibgatullin, N. R. & Sunyaev, R. A. 2000, *Astronomy Letters*, 26, 699
- Smith, H. J. & Hoeffleit, D. 1963, *AJ*, 68, 292
- Sobolev, A. M., Shakhvorostova, N. N., Alakoz, A. V., Baan, W. A., & RadioAstron Maser Team. 2017, *Astronomical Society of the Pacific Conference Series*, Vol. 510, *RadioAstron Maser Observations: a Record in Angular Resolution (Astronomical Society of the Pacific)*, 27
- Soltan, A. 1982, *MNRAS*, 200, 115
- Tananbaum, H., Gursky, H., Kellogg, E., Giacconi, R., & Jones, C. 1972, *ApJ*, 177, L5
- Thompson, A. R., Moran, J. M., & Swenson, George W., J. 2017, *Interferometry and Synthesis in Radio Astronomy*, 3rd Edition (Springer, Cham)
- Twiss, R. Q., Carter, A. W. L., & Little, A. G. 1960, *The Observatory*, 80, 153
- Uttley, P. & McHardy, I. M. 2001, *MNRAS*, 323, L26
- van Cittert, P. H. 1934, *Physica*, 1, 201
- van der Laan, H. 1966, *Nature*, 211, 1131
- van Langevelde, H. J. & Diamond, P. J. 1991, *MNRAS*, 249, 7P
- van Langevelde, H. J., Frail, D. A., Cordes, J. M., & Diamond, P. J. 1992, *ApJ*, 396, 686
- Weizsäcker, C. F. 1948, *Zeitschrift Naturforschung Teil A*, 3, 524
- Welsh, W. F. 1999, *PASP*, 111, 1347
- Witzel, G., Martinez, G., Hora, J., et al. 2018, *ApJ*, 863, 15
- Wollman, E. R., Geballe, T. R., Lacy, J. H., Townes, C. H., & Rank, D. M. 1977, *ApJ*, 218, L103
- Yuan, F. 2007, *Astronomical Society of the Pacific Conference Series*, Vol. 373, *Advection-dominated Accretion: From Sgr A* to Other Low-luminosity AGNs (Astronomical Society of the Pacific)*, 95
- Yuan, F., Quataert, E., & Narayan, R. 2003, *ApJ*, 598, 301
- Yusef-Zadeh, F., Arendt, R., Bushouse, H., et al. 2012, *ApJ*, 758, L11
- Yusef-Zadeh, F., Roberts, D., Wardle, M., Heinke, C. O., & Bower, G. C. 2006, *ApJ*, 650, 189
- Yusef-Zadeh, F., Wardle, M., Heinke, C., et al. 2008, *ApJ*, 682, 361
- Zel'dovich, Y. B. 1964, *Soviet Physics Doklady*, 9, 195
- Zernike, F. 1938, *Physica*, 5, 785
- Zhao, J.-H., Young, K. H., Herrnstein, R. M., et al. 2003, *ApJ*, 586, L29
- Zylka, R., Mezger, P. G., & Lesch, H. 1992, *A&A*, 261, 119
- Zylka, R., Mezger, P. G., Ward-Thompson, D., Duschl, W. J., & Lesch, H. 1995, *A&A*, 297, 83

LIST OF PUBLICATIONS

Primary author papers

Brinkerink, C. D., Müller, C. et al, 2019, "Micro-arcsecond structure of Sagittarius A* revealed by high-sensitivity 86 GHz VLBI observations", *Astronomy and Astrophysics*, 621, A119

Brinkerink, C. D., Müller, C. et al., 2016, "Asymmetric structure in Sgr A* at 3 mm from closure phase measurements with VLBA, GBT and LMT", *Monthly Notices of the Royal Astronomical Society*, 462, 1382

Brinkerink, C. D., Falcke, H. et al., 2015, "ALMA and VLA measurements of frequency-dependent time lags in Sagittarius A*: evidence for a relativistic outflow", *Astronomy and Astrophysics*, 576, A41

Contributing author papers

Kudriashov, V., Martin-Neira, M. et al., 2021, "An Event Horizon Imager (EHI) Mission Concept Utilizing Medium Earth Orbit Sub-mm Interferometry", *Chinese Journal of Space Science*, 41(2), 211

Issaoun, S., Johnson, M. D. et al., 2019, "The Size, Shape, and Scattering of Sagittarius A* at 86 GHz: First VLBI with ALMA", *The Astrophysical Journal*, 871, 30

Liu, H. B., Wright, M. C. H. et al., 2016, "Linearly polarized millimeter and submillimeter continuum emission of Sgr A* constrained by ALMA", *Astronomy and Astrophysics*, 593, A107

Fraga-Encinas, R., Mościbrodzka, M. et al., 2016, "Probing spacetime around Sagittarius A* using modeled VLBI closure phases", *Astronomy and Astrophysics*, 588, A57

Collaboration papers

Kocherlakota, P., Rezzolla, L. et al., EHT Collaboration, 2021, "Constraints on black-hole charges with the 2017 EHT observations of M87*", *Physical Review D*, 103, 104047

Narayan, R., Palumbo, D. et al., Event Horizon Telescope Collaboration, 2021, "The Polarized Image of a Synchrotron-emitting Ring of Gas Orbiting a Black Hole", *The Astrophysical Journal*, 912, 35

Liu, K., Desvignes, G. et al., 2021, "An 86-GHz search for Pulsars in the Galactic Center with the Atacama Large Millimeter/submillimeter Array", *The Astrophysical Journal*, 914, 30

Issaoun, S., Johnson, M. D. et al., 2021, "Persistent Non-Gaussian Structure in the Image of Sagittarius A* at 86 GHz", *The Astrophysical Journal*, in press

EHT MWL Science Working Group, Event Horizon Telescope Collaboration, Fermi Large Area Telescope Collaboration, H.E.S.S. Collaboration, MAGIC Collaboration, VERITAS Collaboration, 2021, "Broadband Multi-wavelength Properties of M87 during the 2017 Event Horizon Telescope Campaign", 2021, *The Astrophysical Journal*, 911, L11

Goddi, C. Martí-Vidal, I. et al., 2021, "Polarimetric Properties of Event Horizon Telescope Targets from ALMA", *The Astrophysical Journal*, 910, L14

Event Horizon Telescope Collaboration, 2021, "First M87 Event Horizon Telescope Results. VIII. Magnetic Field Structure near The Event Horizon", *The Astrophysical Journal*, 910, L13

Event Horizon Telescope Collaboration, 2021, "First M87 Event Horizon Telescope Results. VII. Polarization of the Ring", *The Astrophysical Journal*, 910, L12

Psaltis, D., Medeiros, L. et al., EHT Collaboration, 2020, "Gravitational Test beyond the First Post-Newtonian Order with the Shadow of the M87 Black Hole", *Physical Review Letters*, 125, 141104

Wielgus, M., Akiyama, K. et al, 2020, "Monitoring the Morphology of M87* in 2009-2017 with the Event Horizon Telescope", *The Astrophysical Journal*, 901, 67

Kim, J.-Y., Krichbaum, T. P. et al, Event Horizon Telescope Collaboration, 2020, "Event Horizon Telescope imaging of the archetypal blazar 3C 279 at an extreme 20 microarcsecond resolution", *Astronomy and Astrophysics*, 640, A69

Gold, R., Broderick, A. E. et al., Event Horizon Telescope Collaboration, 2020, "Verification of

Radiative Transfer Schemes for the EHT", *The Astrophysical Journal*, 897, 148

Broderick, A. E., Gold, R. et al., Event Horizon Telescope Collaboration, 2020, "THEMIS: A Parameter Estimation Framework for the Event Horizon Telescope", *The Astrophysical Journal*, 897, 139

AU - Roelofs, F., Janssen, M. et al., 2020, "SYMBA: An end-to-end VLBI synthetic data generation pipeline. Simulating Event Horizon Telescope observations of M 87", *Astronomy and Astrophysics*, 636, A5

Porth, O., Chatterjee, K. et al., Event Horizon Telescope Collaboration, 2019, "The Event Horizon General Relativistic Magnetohydrodynamic Code Comparison Project", *The Astrophysical Journal Supplement Series*, 243, 26

Bower, G. C., Dexter, J., 2019, "ALMA Observations of the Terahertz Spectrum of Sagittarius A*", *The Astrophysical Journal*, 881, L2

Roelofs, F., Falcke, H., 2019, "Simulations of imaging the event horizon of Sagittarius A* from space", *Astronomy and Astrophysics*, 625, A124

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. VI. The Shadow and Mass of the Central Black Hole", *The Astrophysical Journal*, 875, L6

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. V. Physical Origin of the Asymmetric Ring", *The Astrophysical Journal*, 875, L5

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole", *The Astrophysical Journal*, 875, L4

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. III. Data Processing and Calibration", *The Astrophysical Journal*, 875, L3

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. II. Array and Instrumentation", *The Astrophysical Journal*, 875, L2

Event Horizon Telescope Collaboration, 2019, "First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole", *The Astrophysical Journal*, 875, L1

Lu, R.-S., Krichbaum, T. P. et al., 2018, "Detection of Intrinsic Source Structure at ~ 3 Schwarzschild Radii with Millimeter-VLBI Observations of SAGITTARIUS A*", *The Astrophysical Journal*,

859, 60

Goddi, C., Falcke, H., 2017, "BlackHoleCam: Fundamental physics of the galactic center", *International Journal of Modern Physics D*, 26, 1730001-239

Liu, H. B., Wright, M. C. H., 2016, "The 492 GHz emission of Sgr A* constrained by ALMA", *Astronomy and Astrophysics*, 593, A44

Fish, V. L., Johnson, M. D. et al., 2016, "Persistent Asymmetric Structure of Sagittarius A* on Event Horizon Scales", *The Astrophysical Journal*, 820, 90

Johnson, M. D., Fish, V. L. et al., 2015, "Resolved magnetic-field structure and variability near the event horizon of Sagittarius A*", *Science*, 350, 1242

Akiyama, K., Lu, R.-S. et al., "230 GHz VLBI Observations of M87: Event-horizon-scale Structure during an Enhanced Very-high-energy γ -Ray State in 2012", *The Astrophysical Journal*, 807, 150

Lu, R.-S., Fish, V. L. et al., 2013, "Fine-scale Structure of the Quasar 3C 279 Measured with 1.3 mm Very Long Baseline Interferometry", *The Astrophysical Journal*, 772, 13

Conference proceedings

Vecchio, A., Bentum, M. et al., 2021, "The Netherlands-China Low-frequency explorer (NCLE)", *43rd COSPAR Scientific Assembly*, 1525

Roelofs, F., Falcke, H. et al., 2020, "On the prospects of imaging Sagittarius A* from space", *Perseus in Sicily: From Black Hole to Cluster Outskirts*, 342, 24

Kudriashov, V., Falcke, H. et al., 2018, "System design progress in the event horizon imaging using the concept of space-to-space VLBI from medium earth orbits", *42nd COSPAR Scientific Assembly*, E1.8-17-18

Bentum, M., Boonstra, A.-J. et al., 2018, "The CubeSat Low Frequency Explorer (CLE) in Lunar Orbit", *42nd COSPAR Scientific Assembly*, B3.1-27-18

Goddi, C., Falcke, H. et al., 2018, "BlackHoleCam: Fundamental physics of the galactic center", *Fourteenth Marcel Grossmann Meeting - MG14*, 863

Brinkerink, C. D., 2015, "ALMA and VLA Measurements of Frequency-Dependent Time Lags in Sagittarius A*: Evidence for a Relativistic Outflow", *Revolution in Astronomy with ALMA: The Third Year*, 499, 167

Haggard, D., Baganoff, F. K. et al., 2015, "An Update on Chandra/VLA Galactic Center Campaigns Targeting Sgr A* and G2", *American Astronomical Society Meeting Abstracts*, 225, 102.09

Haggard, D., Baganoff, F. K. et al., 2014, "Update on the Sgr A*/G2 Collision from Chandra and VLA", *AAS/High Energy Astrophysics Division*, 14, 100.04

ABOUT THE AUTHOR

Wednesday April 30th, 1980 was a memorable day - it was the date of the coronation of Beatrix as Queen of the Netherlands. It was also the day I was born. Growing up in the East of the Netherlands, in the town of Neede in a house with a pleasantly large garden, two of my major interests were finding out how things worked and building all sorts of contraptions. I suppose those interests were fostered by my parents: my mother, with whom I had many conversations and discussions on science and philosophy, and my father who instilled a deep appreciation for mechanics and engineering in me. The 'Cosmos' book by Carl Sagan was on our bookshelf, and as a child I was particularly enthralled by the beautiful pictures in it. I have read the book many times since then, captivated by the sense of wonder it conveys, and it has played an important part in building my love for astronomy. As I type this at home, I see it standing in my bookcase.

During my later years in secondary school (I attended athenaeum at 'Het Assink', in Haaksbergen), I was considering three possibilities for studies at university: biology, physics & astronomy, and aerospace engineering. I chose the latter because I was excited about the possibility of contributing to space exploration that an engineering degree would offer me. Starting my studies in Delft in 1998, I focused on spaceflight engineering. My interest for rocketry lead to the founding of Delft Aerospace Rocket Engineering (DARE) in 2000. The following years were difficult, with my mother and several other family members passing away. Thankfully, I had solid support from my friends and family during that time. I got my master's degree in 2006, when I had already started working as a software subcontractor at Eurocontrol in Brussels. There, our team implemented trajectory prediction algorithms as part of the automation of scheduling for incoming flights in the TMA2010+ framework. However, I kept thinking of how much more I wanted to learn about astronomy and physics in general. This made me decide to go to university a second time, and so in 2007, after a year and a half in Belgium, I moved to Nijmegen to start my studies in physics and astronomy at Radboud University.

Because of my earlier degree, the structure of my studies at Radboud was somewhat non-standard as I could dive into subjects from different study years simultaneously. I spoke with Heino on possible thesis topics after following one of his courses, and the Event Horizon Telescope project

came up. We settled on a subject for my master's thesis project, where I would analyse the results of a general-relativistic magnetohydrodynamic simulation of an accretion flow onto a black hole. The objective would be to assess how the outflow gets accelerated, and what properties it has as compared to the results from analytical models. In parallel, I had the opportunity to visit UC Berkeley for two months in 2011 to get familiar with radio interferometry and the Event Horizon Telescope project. I spent time at CARMA learning about the subtleties of operating a heterogeneous interferometer array, as well as the principles and challenges of Very Long Baseline Interferometry (VLBI). My bachelor thesis on VLBI measurements with CARMA led to me obtaining my bachelor's degree (cum laude) later in 2011. Soon after, in March 2012, I finished my master's degree – for which I also was happy to receive the *judicium 'cum laude'*.

After concluding my studies, I rolled into my PhD program at the department of astrophysics straight away, working in Heino's group on the material presented in this thesis. Although I was slated to finish up my program in 2016, the opportunity to be a part of the newly formed Radboud Radio Lab (RRL) presented itself to me in 2015. Excited about a chance to combine my background in aerospace engineering with astrophysics, I became involved in the proposal writing team for a low-frequency radio instrument that was to be included on a Chinese relay satellite, deploying into a halo orbit beyond the Moon. This proposal got accepted and resulted in the NCLE instrument which is currently in space. Many other interesting projects within RRL came up – among which are Astrometrix, Event Horizon Imager, PR³, EL3 and contributions to various university courses. I intend to stay working at RRL for the foreseeable future, on the crossroads of science and engineering.

Throughout my PhD and RRL work, I have been involved in several outreach activities in connection to my work at the department of Astrophysics. Giving telescope tours, talks at elementary and secondary schools, public talks at amateur astronomical societies and talks for Sterrenkundeclub are all things I have been doing with enthusiasm and satisfaction. I have also invested time into generating various scientific visualisations and demonstrations of physical processes. Among these are a working audio interferometer, custom renderings of the appearances of curved spaces, an interactive web page on the relations in interferometric imaging, and together with Thomas Bronzwaer an interactive visualisation of curved spacetime around a Schwarzschild black hole.

ACKNOWLEDGMENTS

After such a long journey, meeting many people on the way, writing the acknowledgements seems like a daunting task. I am terrified of forgetting to mention anyone. I apologise in advance if this is the case.

I have found the department of Astrophysics at Radboud University to be a very pleasant working environment, with a culture of open communication and general friendliness, and where having an interest in each other's projects is encouraged. I am very happy to stay with the university in the years ahead, working at the Radboud Radio Lab (RRL).

Heino, thank you for the trust you placed in me by hiring me and letting me try out my own approaches and directions. I have learned so much from you about effective writing, the academic world, international collaboration and about always considering the bigger picture and the longer term. I am also grateful for the opportunity you granted me to get involved in the work at RRL, even though my PhD work was not yet finished.

I also owe a great debt of gratitude to my manager at RRL. Marc, thank you for involving me in the various interesting RRL projects and still allowing me the time to finish my PhD work. I hope to remain involved in the exciting story of RRL for a while yet, learning a lot as I go along. With many other projects happening, my work on this thesis was at times in danger of getting snowed under and postponed. I would like to thank Marijke and Gijs for making sure that I kept the end goal in sight. Amanda, thank you for your help in the preparation for my defence - it is much appreciated.

My thanks to the people who have kept so many aspects of the department running properly every day throughout these years: Cisca, Esther, Marja, Helma, Perine, Monique and Merijn. Irene, thank you for the entertaining coffee break chats and for supplying many worthwhile suggestions to put on my reading list. Thomas, I greatly enjoy the way we rapidly seem to generate ideas whenever we sit down to work on something together. There is too little time in the day to pursue all of them, but thankfully no shortage of interesting thoughts! Let's build on that in the

coming years. Nela, thank you for the pleasant collaboration we had on the VLBI papers we wrote together. You always remained patient whenever I suddenly realised I had to try another approach, and I am grateful for that. Sjoerd, I appreciate your readiness to tackle all sorts of different projects involving both hardware and software, and always learning new things in the process - you are a continuing source of inspiration. Daan, thank you for the many times I could count on you for guidance on how to set up computing infrastructure and how to design sensible algorithms. My EHT colleagues Raquel, Jordy, Freek, Michael and Sara: I admire your focus and tenacity. Thank you for the many constructive discussions, and all the fun we have had at the department as well as at various conferences we attended together!

During our lunch and coffee breaks I have enjoyed listening to the 'mini-lectures' and takes on actualities provided by prof. Frank Verbunt. Frank, I much appreciate your willingness to discuss and debate various science topics and the state of the world in general, with points that are clearly made and well-defended.

I would like to thank Geoff Bower for taking the time to teach me about the fundamentals of VLBI during my stay in Berkeley, and I thank Dick Plambeck for his unique sense of humour and for being a treasure trove of knowledge on instrumentation at CARMA. I still chuckle when I recall his pithy remarks made during observing campaigns at the high site.

My family has been a consistent and enthusiastic source of support for me throughout these years. Dad and MJ, thank you for always being welcoming and for the sense of humour that we share and enjoy. Visiting Diepenheim is always a pleasure. Lientje, thank you for sharing the adventures you are having with Rudi, Viktor and Koert with me. I promise to come visit soon! Thijs, Barbara, Brecht, Ester, Marije en Roy, I am very happy to have you close to me as family. Thank you for being a part of my life. Simon, Marijn and Rens, it is a pleasure to watch you growing up! I hope we can take you on many amazing mountain hikes in the future.

My dear Bokito's: Alk, Kints, Luuk, Phil, Zompje, you bunch of apes, you know who you are. We have been on many memorable trips, and shared both good and tough times together. Thank you for our gatherings, which I always look forward to.

What an astonishing thing a book is. It's a flat object made from a tree with flexible parts on which are imprinted lots of funny dark squiggles. But one glance at it and you're inside the mind of another person, maybe somebody dead for thousands of years. Across the millennia, an author is speaking clearly and silently inside your head, directly to you. Writing is perhaps the greatest of human inventions, binding together people who never knew each other, citizens of distant epochs. Books break the shackles of time. A book is proof that humans are capable of working magic.

Carl Sagan, *Cosmos*, Part 11: The Persistence of Memory (1980)