



Rapid Variability of Sgr A* across the Electromagnetic Spectrum

G. Witzel¹, G. Martinez², S. P. Willner³, E. E. Becklin^{2,4}, H. Boyce^{5,6}, T. Do², A. Eckart^{1,7}, G. G. Fazio³, A. Ghez², M. A. Gurwell³, D. Haggard^{5,6}, R. Herrero-Illana⁸, J. L. Hora³, Z. Li (李志远)⁹, J. Liu (刘俊)¹, N. Marchili¹⁰, Mark R. Morris², Howard A. Smith³, M. Subroweit⁷, and J. A. Zensus^{1,7}

¹Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, D-53121, Bonn, Germany

²Department of Physics and Astronomy, University of California, Los Angeles, Box 951547, Los Angeles, CA 90095-1547, USA

³Center for Astrophysics, Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA

⁴SOFIA Science Center, Moffett Field, CA, USA

⁵Department of Physics, McGill University, 3600 University St., Montreal, QC H3A 2T8, Canada

⁶McGill Space Institute, McGill University, Montreal, QC H3A 2A7, Canada

⁷I. Physikalisches Institut, Universität zu Köln Zùlpicher Str. 77 D-50937 Köln, Germany

⁸European Southern Observatory (ESO), Alonso de Córdova 3107, Vitacura, Casilla 19001, Santiago de Chile, Chile

⁹School of Astronomy and Space Science, Nanjing University, 163 Xianlin Avenue, Nanjing 210023 Jiangsu, People's Republic of China

¹⁰Istituto di Radioastronomia—INAF, Via Piero Gobetti 101, I-40129, Bologna, Italy

Received 2020 November 11; revised 2021 June 2; accepted 2021 June 3; published 2021 August 18

Abstract

Sagittarius A* (Sgr A*) is the variable radio, near-infrared (NIR), and X-ray source associated with accretion onto the Galactic center black hole. We have analyzed a comprehensive submillimeter (including new observations simultaneous with NIR monitoring), NIR, and 2–8 keV data set. Submillimeter variations tend to lag those in the NIR by ~ 30 minutes. An approximate Bayesian computation fit to the X-ray first-order structure function shows significantly less power at short timescales in the X-rays than in the NIR. Less X-ray variability at short timescales, combined with the observed NIR–X-ray correlations, means the variability can be described as the result of two strictly correlated stochastic processes, the X-ray process being the low-pass-filtered version of the NIR process. The NIR–X-ray linkage suggests a simple radiative model: a compact, self-absorbed synchrotron sphere with high-frequency cutoff close to NIR frequencies plus a synchrotron self-Compton scattering component at higher frequencies. This model, with parameters fit to the submillimeter, NIR, and X-ray structure functions, reproduces the observed flux densities at all wavelengths, the statistical properties of all light curves, and the time lags between bands. The fit also gives reasonable values for physical parameters such as magnetic flux density $B \approx 13$ G, source size $L \approx 2.2R_S$, and high-energy electron density $n_e \approx 4 \times 10^7 \text{ cm}^{-3}$. An animation illustrates typical light curves, and we make public the parameter chain of our Bayesian analysis, the model implementation, and the visualization code.

Unified Astronomy Thesaurus concepts: Galactic center (565); Astrophysical black holes (98); Supermassive black holes (1663); Black hole physics (159); High energy astrophysics (739); Accretion (14)

Supporting material: animation, data behind figure

1. Introduction

Since the discovery of rapid flaring of Sagittarius A* (Sgr A*) in the X-rays and near-infrared (NIR) in the early 2000s (Baganoff et al. 2001; Genzel et al. 2003; Ghez et al. 2004), many monitoring programs have been executed to understand the origin and properties of the variable emission. As a result, more than 70 papers have been published describing observations and modeling light curves in the submillimeter (submm), NIR, and X-rays. Most of these publications have focused on statistical analyses of flux densities and timing properties, multiwavelength observations and modeling of the spectral energy distribution (SED), and NIR and submm polarization. (See reviews by Genzel et al. 2010 and Morris et al. 2012, the comprehensive Witzel et al. 2018 introduction, and references in their paper.) This fascination with Sgr A*'s variability has a reason: light-crossing-time arguments link rapid changes in flux density to spatial scales that—until recently—were not accessible otherwise. For example, Dodds-Eden et al. (2009) found sudden flux density changes of a factor ~ 2 in less than 47 s, and Do et al. (2019) saw changes of a factor ~ 9 in less than

2 minutes. These times correspond to spatial scales of $1.2R_S$ and $3R_S$, respectively.¹¹

Event-horizon spatial scales can now be studied by two types of large interferometers. One is VLTI/GRAVITY, which operates in the NIR and has observed the variable emission source moving in the plane of the sky (Gravity Collaboration et al. 2018a, 2018b, 2019, 2020a, 2020c). The other is very long baseline interferometry (VLBI) at millimeter wavelengths (reviewed by Boccardi et al. 2017). The Event Horizon Telescope (EHT) aims to image Sgr A* with a resolution close to R_S at 1.3–3.5 mm wavelengths (Dexter et al. 2014; Tilanus et al. 2014; Lu et al. 2018). The detection of circular trajectories of the center of light during flares of Sgr A* by VLTI/GRAVITY and corresponding loops in the Stokes (Q , U) plane (Gravity Collaboration et al. 2018b) suggests compact source structure of $< 5R_S$ on an orbit around an average position at $\sim 9R_S$ (Gravity Collaboration et al. 2020c). At

¹¹ Here, R_S means Schwarzschild radius, $R_S = 1.23 \times 10^{10} \text{ m}$ for mass $M = 4.15 \times 10^6 M_\odot$ (Gravity Collaboration et al. 2019). The scales mentioned are not to be taken as distances to the black hole. They are rather upper limits on the characteristic size of the volume where the radiation originates, e.g., of a region in the accretion disk or in a jet likely well away from the event horizon.

1.3 mm wavelength, VLBI studies also found compact structure with upper limits on the intrinsic source size of $\lesssim 4R_S$ (Doeleman et al. 2008), likewise with some evidence that the center of light is not centered on the black hole itself.

In light of the interferometric results, the variability data obtained over the last two decades are valuable as a complementary source of information about the physical processes at event-horizon scales. Many multiwavelength campaigns in the submm, NIR, and X-rays have been organized in the hope of determining—or at least constraining—the underlying radiative processes (Baganoff et al. 2001; Eckart et al. 2004, 2006, 2008a, 2008b, 2012; Gillessen et al. 2006; Yusef-Zadeh et al. 2006a, 2006b, 2008, 2009, 2012; Marrone et al. 2008; Dodds-Eden et al. 2009; Trap et al. 2011; Haubois et al. 2012; Mossoux et al. 2016; Rauch et al. 2016; Ponti et al. 2017).

Several models have been proposed to explain the variability at different wavelengths. All models assume that the NIR is dominated by optically thin synchrotron radiation, but they differ in the mechanism for the X-ray emission. Some make the case for optically thin synchrotron radiation with a cooling break to explain the X-rays (Dodds-Eden et al. 2009; Ponti et al. 2017), while others suggest synchrotron self-Compton scattering (synchrotron-SSC) (Dodds-Eden et al. 2009; Eckart et al. 2012; Mossoux et al. 2016) or inverse Compton scattering of photons by a second population of electrons (Dodds-Eden et al. 2009).

While some authors have claimed evidence for a close relation between submm and NIR variability, the phenomenology and degree of correlation remain inconclusive. Dexter et al. (2014), for example, found a submm variability timescale of ~ 8 hr, significantly longer than in the NIR, and concluded that different mechanisms might be creating the variability in the two wavelength regimes. In contrast, Eckart et al. (2012) linked the submm and NIR through the evolution of the optical depth caused by adiabatic expansion. Evidence for variability peaks propagating from submm to radio frequencies has been reported (e.g., Yusef-Zadeh et al. 2006b; Marrone et al. 2008) and convincingly modeled in the framework of adiabatic expansion (Yusef-Zadeh et al. 2009, their Figure 27).

Eckart et al. (2006), Bremer et al. (2011), and Witzel et al. (2018) found evidence for an (exponential) cooling cutoff of the SED at NIR frequencies. If the cooling-cutoff energy varies with source luminosity, that could explain spectral index changes as a function of flux density. Dodds-Eden et al. (2010) developed the first time-dependent models for Sgr A* variability, integrating the differential equations for the electron energy distribution under injection and escape, resulting in sequences of SEDs.

While past studies are very informative about many basic properties of the covariability of the flux densities at different wavelengths, multiwavelength campaigns are difficult to organize and have to rely on favorable weather and operational conditions at all observing sites, as well as on the luck of Sgr A* varying above the detection limit during the campaign. As a result, only a small number of simultaneous multi-wavelength observations are available today. Many of these have been obtained during 24 and 16 hr observations with Spitzer/IRAC combined with some of Keck, Chandra, SMA, and ALMA (Fazio et al. 2018; Witzel et al. 2018; Boyce et al. 2019). Open questions still include the exact nature of the

NIR-submm and NIR–X-ray correlations and the power spectral density (PSD) of the X-ray variability.

Despite the critical role played by the variable component, it does not account for all of Sgr A*'s emission. The radio emission at wavelengths $\gtrsim 3$ mm shows much lower fractional variability (Genzel et al. 2010), and there is a nonvarying, spatially extended ($\sim 1''$) X-ray component (Baganoff et al. 2003; Xu et al. 2006). The angular size of the constant radio component is unknown because of interstellar scattering. This paper addresses only the (intraday) variable component of emission. The other components must come from a separate process or processes, and the spatial extent shows that the constant X-ray component at least must originate in a different volume of the source than the compact volume giving rise to the intraday variable emission.

The goal of the present work is to put the statistical properties of Sgr A*'s rapid variability at all wavelengths into the context of a single radiative model by using available long-duration observations to best advantage and without requiring simultaneity of the individual observations. This approach has already been developed by Witzel et al. (2018) in modeling the covariability of 2.2 and 4.5 μm . Given the clear evidence of structure with size $< 5R_S$, we consider a one-zone model. The model combines a self-absorbed synchrotron spectrum with cooling cutoff and synchrotron-SSC scattering. This model choice is partly motivated by the timing properties of the correlated NIR and X-ray light curves, which cannot easily be reproduced by the other models mentioned above. Additionally, Eckart et al. (2012), Mossoux et al. (2016), and Subroweit et al. (2020) showed convincingly that peak fluxes of simultaneously observed NIR and X-ray flares as well as the flux density distributions are well described by a synchrotron-SSC model. Eckart et al. (2012) pointed out that bright, compact synchrotron sources show the precise conditions to exhibit self-absorption and self-scattering. In their model, the observed pairs of flux density peaks suggest source sizes of a few R_S and magnetic flux densities of a several tens of Gauss. A source of this size and magnetic flux density shows a flux density spectrum that peaks in the submm (naturally contributing to variability at submm wavelengths as well), and X-ray photons are mainly upscattered submm photons. In contrast to earlier studies, our analysis takes into account the autocorrelation and cross-correlation properties of the light curves and the effects of the cooling cutoff in the NIR by fitting a time-dependent, analytic model to the body of submm, NIR, and X-ray data.

This paper has two parts. The first analyzes—for the first time—the PSD of the X-ray variability. The PSD model is based on a generic statistical (i.e., not physically motivated) flux density model. The paper's second part, motivated by the PSD analysis, proposes a single-zone radiative model and analyzes the variability at submm, NIR, and X-ray wavelengths. Both parts discuss prior distributions, a model, and the resulting posteriors of the model parameters. Section 2 describes the data sets used in this analysis. Section 3 shows evidence that, at the shortest variability timescales, power in the X-ray PSD appears suppressed compared to the NIR PSD. Further, the correlation between the NIR and the X-rays (characterized by strict correlation in arrival time and lack of correlation in flux density levels) can be understood as the effect of the difference between the two PSDs. Section 4 shows that the separation of timescales is a natural result of the

equation for SSC flux densities (Marscher 1983) if the source of the fast NIR variability is the fast-varying cooling cutoff close to NIR frequencies. We additionally discuss modes of codevelopment of the synchrotron self-absorption with NIR and X-ray flux densities, and fit a simple, semi-analytic version of the model to the observed first-order structure functions in all bands. Section 5 presents the model-fit results, including animations. Section 6 discusses the merits and shortcomings of the model, and Section 7 summarizes our findings.

As a convention, *H*-, *K*-, *L*-, and *M*-band refer to the respective NIR bands. Radio bands will be denoted by their central wavelengths. We use the term SED even when the quantity is expressed as flux density (S_ν) rather than energy (νS_ν).

2. The Data

Our model is based on a coanalysis of a comprehensive multiwavelength variability data set of Sgr A*. With the exception of two of the ALMA and four of the SMA observations, all data were published before. While comprehensive, this data set is not complete, but it uses the large NIR and X-ray samples analyzed in previous publications because they are statistically well-characterized already.

2.1. Near-infrared

The NIR data set used in this analysis is a combination of the extensive 2.12, 2.18, and 4.5 μm data of Witzel et al. (2018) and the 2.12 μm data set of Do et al. (2019). The original publications explain the reduction and statistics. (For the historic *K*-band data set, see also Witzel et al. 2012 and Meyer et al. 2014, and for the fundamentals of observing Sgr A* with Spitzer/IRAC, see Hora et al. 2014.) Our data set contains eight 24 hr epochs of Sgr A* at 4.5 μm with IRAC on the Spitzer Space Telescope, 93 epochs of 2.18 μm data from Naos Conica at the Very Large Telescope, 34 epochs of 2.12 μm data from the NIRC2 camera at the Keck Observatory, for a total of 95,307 NIR measurements. Figure 1 shows the light curves of these three observatories. The average cadence of the Spitzer data is 8.4 s, and typical cadences of the VLT/Keck light curves are about one image per 1.1–1.2 minutes with integration times of 30–40 s and 28 s, respectively. The typical uncertainties of the individual data points are 0.66 mJy for Spitzer, 0.033 mJy for the VLT, and about 0.017 mJy for the Keck data. The light curves of all three instruments are prone to contamination with an additive background flux density level from the resolved and unresolved stellar population at the Galactic center. The Spitzer data are—by nature of the applied data reduction algorithm—a differential measurement with an arbitrary zero point. The VLT and Keck data show a typical photometric offset of about 0.06 mJy and 0.03 mJy, respectively. These offsets play an important role in spectral index measurements but do not affect the first-order structure function, which quantifies differences in flux densities rather than absolute levels.

2.2. Chandra Data

Sgr A* has been observed often by Chandra, starting in 1999. We include all data available from the Chandra archive through 2017 having aimpoint within $1'$ of Sgr A*. This ensured an optimal point-spread function and best photometric performance for a compact source like Sgr A*. The

Chandra instruments and observing modes have changed over time:

1. In 1999–2011, all data were taken with ACIS-I, and there were two additional ACIS-I observations taken in 2013. These 49 ACIS-I exposures sum to 1.5 Ms. Pileup affects these observations.
2. In 2012, in the framework of the X-ray Visionary Program (XVP), the data were taken with the ACIS-S/HETG. There are three additional observations taken with ACIS-S/HETG in 2013. These 41 ACIS-S/HETG exposures sum to ~ 3 Ms. Only the zeroth-order image is used here. Pileup is small but not always zero.
3. From 2013 to 2017, almost all observations were done with ACIS-S (no grating) in subarray mode, making them basically unaffected by pileup. The 39 ACIS-S exposures sum to ~ 1.4 Ms. However, during 2013 and 2014 (the first 25 epochs), the magnetar PSR J17452900 (Coti Zelati et al. 2017) contributed significantly to the flux. Because this varying source contributed additional photon noise, we ignored all ACIS-S epochs prior to 2015 October 21. This leaves about a two-year gap between the last ACIS-I observation (2013) and the first ACIS-S observation (2015-10), when the magnetar became dim enough not to affect our photometry. The 14 good epochs give ~ 0.6 Ms of data.

All told, we have 103 epochs and 5.3 Ms of 2–8 keV Chandra data.

We reduced the Chandra data ourselves to guarantee consistency across all epochs. All archival data were downloaded and reprocessed using the Chandra Interactive Analysis of Observations software package (CIAO v4.10) and the Chandra Calibration Database (CALDB v4.7.8) following the standard procedures as outlined by Zhu et al. (2019), who used the same Chandra data set to study a candidate parsec-scale jet from Sgr A*. Photons were extracted from within a radius of $1''.25$ from the best-guessed centroid of Sgr A* and within the 2000–8000 eV range to be consistent with most previous work. This gave a total of 34,000 photons in the useful epochs. We used unbinned data of the arrival times of individual photons (counts) and corrected the times to the barycenter of the Solar System.

Figure 1 shows the light curves of all three modes. The effective area of each mode is a weighted mean over the 2–8 keV band, and it assumes an incident source spectrum. If we changed the assumed spectral model, the absolute values of the effective area would change, but the relative values among observations using the same detector should be insensitive to the model. The relative values between different detectors (I versus S) could be more sensitive but are assumed to be constant as well.

2.3. APEX Data

We used 32 epochs (6641 minutes) of 345 GHz data from the LABOCA bolometer at the APEX telescope (Subroweit et al. 2017). This data set was generated using on-the-fly mapping, resulting in fully sampled maps of $0^\circ.5 \times 0^\circ.17$ with 280 s integration time. The data were taken over the course of seven years and have a typical cadence of about 8 minutes. An average map was created for each epoch, and after subtracting a Gaussian point source at the position of Sgr A*, this map was then subtracted from each individual image. The flux density of




Figure 1. Observed light curves for Sgr A*. The data are presented without observational gaps between the epochs, but each epoch is in a different color. (Colors do not indicate common epochs across different panels). Gap durations are hours to years, but they should not affect the present analysis. Timescales differ among panels. To the accuracy of the calibrations, all flux densities are for a point source, excluding extended structures, and all data are presented as observed with no correction for interstellar extinction. Top panel shows $2.2 \mu\text{m}$ data, with the vertical dashed line separating VLT data on the left from Keck data on the right. Second panel shows $4.5 \mu\text{m}$ Spitzer/IRAC data in one-minute bins. Third panel shows Chandra/ACIS-S/HETG data, and fourth panel shows Chandra/ACIS-I data on the left and Chandra/ACIS-S data on the right. X-ray data are shown with one-minute binning and no pileup correction. The different Chandra instruments cover nearly the same energy range but with different sensitivities, which are accounted for in the analysis. Fifth panel shows 345 GHz data from APEX/LABOCA and SMA on the left and 230 GHz data from SMA and ALMA on the right. The observed light curves for Sgr A* shown in Figures 1 and 2 are available as the data behind the Figure.

(The data used to create this figure are available.)

Sgr A* was derived by modeling a Gaussian source at the position of Sgr A* in each residual image and using two secondary calibrators (G10.62–0.38, IRAS 16293–2422). The relative uncertainty of the flux density calibration is about 4% or 0.1 Jy, while the absolute uncertainty is expected to be of order 15%. Eckart et al. (2008b), García-Marín et al. (2011), and Subroweit et al. (2017) gave detailed descriptions of the data reduction, calibration, and atmospheric opacity monitoring. The lower panel of Figure 1 shows the resulting 345 GHz light curves.

2.4. SMA Data

We include one published (365 minutes; Fazio et al. 2018) epoch of 343 GHz data and four new epochs (1573 minutes) of 230 GHz data from the Submillimeter Array (SMA). Observation details are in Table 1. The SMA was operated in a dual-receiver polarization track with double-sideband observations using sideband separation implemented in the correlator. The continuum visibility was calculated by averaging the two same-sense polarization signals. Final flux-density measurements were determined by vector-averaging the measured visibility

Table 1
New SMA and ALMA Epochs

Date (UT)	Start Time (UT)	Stop Time (UT)	Baselines ^a (k λ)	# Ant.	Tuning (GHz)	BW ^b (GHz)	Calibrators
SMA							
2014 Jun 18	07:23:48	13:28:47	6.2–84.8	8	343.0	8 (2 GHz/sb \times 2 sb \times 2 pol)	Neptune, NRAO 530
2015 May 14	09:52:41	15:52:24	6.1–133.2	6	226.9	8 (4 GHz/sb \times 2 sb \times 1 pol)	Titan, Callisto, NRAO 530
2016 Jul 13	05:34:48	12:04:48	46.7–412	6	236.1	16 (4 GHz/sb \times 2 sb \times 2 pol)	NRAO 530, J1924–292 ^c
2017 Jul 16	04:46:05	12:27:14	4.9–54.4	8	228.0	32 (8 GHz/sb \times 2 sb \times 2 pol)	NRAO 530, J1924–292
2017 Jul 26	05:19:43	11:22:07	6.3–53.8	7	228.0	32 (8 GHz/sb \times 2 sb \times 2 pol)	NRAO 530, J1924–292
ALMA							
2016 Jul 12	22:58:12	03:38:14	9.4–674.4	40	232	7.45	Titan, PKS 1741–312
2016 Jul 18	23:14:51	24:10:34	9.4–674.4	40	232	7.45	Titan, PKS 1741–312
2016 Jul 19	02:05:37	05:43:36	9.4–674.4	40	232	7.45	Titan, PKS 1741–312

Notes.

^a Baselines $<35\text{--}40$ k λ (exact limit different for different SMA integrations) were not used for the final flux densities, in order to avoid contamination by extended structure.

^b Effective spectral bandwidth including all polarizations, not necessarily continuous.

^c Flux density calibration was based on secondary calibrations of the gain calibrators listed.

data for instantaneous baselines greater than 35–40 k λ , to filter out large-scale emission structure around Sgr A*. Water vapor was ~ 1.5 mm for most epochs, but nearer 0.9 mm in 2016, and higher and unstable on 2016 July 19. The lower panel of Figure 1 shows the resulting 343 and 230 GHz SMA light curves.

2.5. ALMA Data

This analysis used 12 epochs (1374 minutes) of data from the ALMA observatory. The first two epochs (project code 2015.A.00021.S, PI G. Witzel) were taken in 2016 while Spitzer was observing at 4.5 μm . Details are in Table 1. The data were calibrated with the standard ALMA pipeline using the Common Astronomy Software Application package (CASA; McMullin et al. 2007). A few spectral windows showed suspicious absorption features and were excluded from the analysis. Because ALMA has relatively fewer short baselines than SMA, all baselines were included, and the source was imaged with uniform weighting using task CLEAN. Sgr A* is a strong source at mm wavelengths, and its visibility was mostly flat as a function of baseline length. That means the source is compact, and with ALMA’s good $U-V$ coverage and uniform weighting, light curves could be extracted by simply measuring the peak flux densities of the image. The off-source rms gave the uncertainty. The robustness of the results was investigated using 1–30 minute averaging times. The final choice of 3 minutes is a good balance between signal-to-noise and sampling the light curve. We obtained two long light curves, the first 389 minutes with a gap and the second 280 minutes as shown in Table 1.

The remaining 10 ALMA epochs are 70 minute light curves at 234 GHz (Iwata et al. 2020). The epochs were observed in 2017 October over ten days, i.e., they are separated by roughly one day and have an average cadence of ~ 1.6 minutes. Iwata et al. (2020) gave details of the data reduction. Figure 1 shows all 12 epochs of 230 GHz ALMA data.

2.6. Simultaneous NIR and Submm Observations

A majority of the submm data were taken during the Spitzer observations, resulting in the largest set of synchronous

submm (230 or 345 GHz) and NIR (4.5 μm) light curves. (Chandra observed simultaneously as well, but there was no X-ray event during times of SMA or ALMA coverage). In total, we obtained >1800 minutes on-source within a total duration of 2247 minutes of simultaneous data. The combined ALMA and SMA data on 2016 June 12 cover 787 minutes with a gap of 116 minutes between the data sets, making this the longest simultaneous light curve so far. Figure 2 shows all simultaneous data taken during the Spitzer campaigns.

Figure 3 shows the discrete cross-correlation (DCC) of the entire 230 GHz data set with the 4.5 μm light curves. Because the data were not regularly sampled and show large gaps, we used the algorithm of Edelson & Krolik (1988) as implemented by Robertson et al. (2015). We calculated the DCC for both the data as presented in Figure 2 and for a transformed version of the 4.5 μm data. The transform was to convolve the data with a Hanning smoothing kernel with size of 61 minutes and then take the logarithm.¹² Both DCCs have their maxima at time lags of 20–30 minutes, with positive lag meaning that the 4.5 μm data led the submm.

3. The X-Ray Power Spectrum and the NIR–X-Ray Correlation

The first step of our analysis was to determine the PSD of the X-ray variability of Sgr A* under the preliminary assumption of a log-normal PDF for the X-ray flux densities. We did this by forward-modeling light curves for a range of PSDs, comparing the model light curves with observed ones through a suitable summary statistic (a “distance function”), and exploring the parameter space by approximate Bayesian computation (ABC). ABC is appropriate for problems with no analytic likelihood function. All three parts—the forward modeling, distance function, and implementation of the ABC—were described in detail by Witzel et al. (2018, their Appendix B). The present analysis closely follows their procedure but was adapted for X-rays by deriving photon statistics from the modeled variable

¹² Smoothing the NIR data suppresses fast NIR variability, which is uncorrelated with the submm and the X-rays. The logarithm accounts for the nonlinear relation between the submm and the NIR, a consequence of the radiative mechanism discussed in Section 6.




Figure 2. Simultaneous submm and $4.5 \mu\text{m}$ data sets obtained during the Spitzer campaigns. Blue lines show the $4.5 \mu\text{m}$ Spitzer light curves with one-minute binning. Flux densities are as observed, not corrected for reddening. Purple shows 230 GHz ALMA light curves, and orange shows 230 GHz and 345 GHz SMA light curves as labeled. Times are arbitrary but matched to heliocentric for all curves. The 345 GHz SMA light curve of 2014 June 17 is the same one published by Fazio et al. (2018). Other submm data are presented here for the first time. For this figure as well as for the correlation analysis, the 2017 SMA data has been slightly smoothed with a three-minute smoothing kernel. The observed light curves for Sgr A* shown in Figures 1 and 2 are available as the data behind the figure (see Figure 1).

flux densities plus a constant source contribution. Details are given below.

A log-normal PSD was chosen because (a) it can—as the results show—describe the observed data successfully, and (b) a log-normal was the preferred model for NIR light curves (Witzel et al. 2018). That analysis showed that the NIR PSD is a single broken power law with a slope of 2 (red noise) and a precisely measured break timescale $\tau_b = 243^{+82}_{-57}$ minutes. In all such analyses, the inferred PSD parameters depend on the model for the flux-density PDF, and using a log-normal PSD lets us directly compare X-ray PSD parameters with the NIR. The analysis in Section 4 replaces the PSD model assumption with a physical model.

3.1. Forward Modeling of Light Curves

We used the FFT-based method (Timmer & Koenig 1995) to generate random light curves from a given PSD and a set of independently drawn random numbers. This method results in realizations of a Gaussian process $g(t)$ exhibiting periodograms (i.e., PSD estimators) consistent with the input PSD. We parameterized the PSD as a broken power law of

the form:

$$\text{PSD}(f) \propto \begin{cases} f^{-\gamma_0} & \text{for } f < f_b \\ f^{-\gamma} & \text{for } f \geq f_b, \end{cases} \quad (1)$$

where f is temporal frequency, and we assume $\gamma_0 = 0$ (Meyer et al. 2009). In order to generate realistic light curves, we transformed the values $g(t)$ to make the resulting distribution of $S(t) = T[g(t)]$ consistent with the distribution of observed flux densities. With a log-normal distribution as our target,

$$T(g(t)) = \exp(\sigma_{\log n} \cdot g(t) + \mu_{\log n}), \quad (2)$$

$\mu_{\log n}$ and $\sigma_{\log n}$ being the log-normal parameters.

For Chandra data, we modeled the light curves in count rate $\Lambda(t) = T[g(t)]$ (counts per second; cps) instead of flux density $S(t)$. The conversion factor between flux density (as observed at Earth, i.e., after suffering interstellar extinction) and count rate is absorbed by the log-normal mean μ and otherwise does not affect log-normality. We assumed an effective area $e_I = 1$ for the ACIS-S/HETG data and Gaussian priors for e_G/e_I and




Figure 3. Discrete cross-correlation (DCC) functions of the 230 GHz and 4.5 μm light curves shown in Figure 2. Orange shows the DCC of the unaltered data, and blue with shaded error bars shows the DCC of 230 GHz with the logarithm of the low-pass-filtered 4.5 μm data. The vertical dashed line indicates a time lag of zero. The DCC reaches its maximum at a time lag of about 33 minutes for the unaltered data and at 27 ± 8 minutes (solid vertical line) with the transformed data. Positive lags indicate the 4.5 μm data lead the submm. The quoted uncertainty of the time lag is 2.345 times the full width at half maximum (FWHM) of the part of the transformed DCC that lies above the 95% false alarm probability (dashed horizontal curve). The false alarm probability was derived from uncorrelated pairs of light curves generated from our Bayesian posterior (Section 5).

e_S/e_I , the relative effective areas of the other two modes.¹³ For each mode, we included an additive term (χ_G , χ_S , χ_I) to account for the contribution from the extended, nonvarying X-ray source at the position of Sgr A*. The resulting total count rate (before measurement errors) from the position of Sgr A* is then:

$$\Lambda_{\text{tot}}(t) = e_i/e_I \cdot [\Lambda(t) + \chi_i], \quad (3)$$

where i is S , I , or G for the three Chandra instruments used, and χ_i is the count rate of the constant source as seen by the respective instrument. Pileup was included in the forward modeling for ACIS-I and ACIS-S/HETG data for all count rates >0.02 cps. (It is negligible below that.) We used Equations (3) and (4) of Yuan & Wang (2016):

$$\Lambda_{\text{out},I}(t) = [4.180\Lambda_{\text{tot}}(t)^{-0.07387} + 0.5381\Lambda_{\text{tot}}(t)^{-1.160}]^{-1}, \quad (4)$$

and

$$\Lambda_{\text{out},G}(t) = [3.933\Lambda_{\text{tot}}(t)^{-0.03541} + 0.6564\Lambda_{\text{tot}}(t)^{-1.107}]^{-1}. \quad (5)$$

For ACIS-I count rates between 0.1 and 0.2 cps, the corrections are 27%–64%. We finally modeled the measured count rate Λ_{meas} , sampled once per minute, by a Poisson process:

$$\Lambda_{\text{meas}}(t) = \text{Pois}[\Lambda_{\text{out}}(t) \cdot 60]/60. \quad (6)$$

With this one-minute binning, the effect of the Chandra frame time (3.2 s for ACIS-I and HETG, 0.4 s for ACIS-S subarray) is negligible.

¹³ Our priors on e_G/e_I and e_S/e_I are based on the spectral model used by Zhu et al. (2018), i.e., an absorbed bremsstrahlung with a plasma temperature of 10 keV and a foreground absorption column density of 10^{23} cm^{-2} .




Figure 4. Logarithmically binned structure functions (Equation (7)). Upper panel shows the X-ray SFs, with red for ACIS-I, green for ACIS-S, and blue for ACIS-S/HETG. Lower panel shows the 230 and 345 GHz SFs, where red indicates the combined SF from SMA and ALMA data, and blue the APEX data combined with one epoch of SMA data.

3.2. The Distance Function

Following Witzel et al. (2018, their Appendix B.2), we used the first-order structure function as the distance function. The structure function quantifies the variance of the flux density at any given time separation and contains information on the PSD as well as the flux-density distribution of the variability process. The structure function is defined as:

$$V(\tau_i) = \frac{1}{n_i} \sum_{t_j, t_k} [F(t_j) - F(t_k)]^2 \quad \text{for } \tau_i \leq (t_j - t_k) < \tau_{i+1}, \quad (7)$$

that is, the sum of $[F(t_j) - F(t_k)]$ over all measurement pairs whose time lags $(t_j - t_k)$ fall within the bin $[\tau_i, \tau_{i+1}]$, there being n_i such pairs. The structure functions for the three modes of Chandra data are shown in Figure 4. A detailed discussion of the choice of the τ_i can be found in Section 3 and Appendix B.2 of Witzel et al. (2018). In short, with increasing time lag, a decreasing number of point pairs contribute to the structure function bins. For time lags longer than half the observing window, not all flux-density measurements contribute to every structure function bin, and the variance of the structure function increases dramatically without carrying much information about the intrinsic variability. Therefore, we chose a logarithmic binning scheme, roughly equally spaced in logarithmic time lags, with a spacing large enough to allow for a similar number of points in the long-time-lag bins. The maximum lag bin was defined by the point where the variance of the linearly

binned structure function starts to increase significantly. This bin is larger by a factor of 3 (for the X-ray data) than the other bins, in order to mitigate the increase of the variance.

We defined the distance between two light curves as the weighted L2 norm of the difference between the logarithms of the respective structure function’s binned values:

$$\phi(V_1, V_2) = \sum_i w_i (\log [V_1(\tau_i)/V_2(\tau_i)])^2, \quad (8)$$

with w_i being the weights for the chosen binning. The weights adopted here were unity for each structure-function bin except the single wide bin at large time lags, which had $w_i = 3$. These values gave uniform and reasonably quick convergence of the fit.

There is a general problem of consistency when comparing variability timescales quoted in the literature. The definition of “timescale” can depend on the method used for its estimation, and ignoring that can lead to errors in the interpretation of time series analysis results, especially in their conversion into physical quantities. The most substantial difference in the definition of timescales concerns the domain in which they were estimated. In the spectral domain, a timescale naturally refers to the period of a sinusoidal signal. In the time domain, in particular when the estimation is through a structure function, “timescale” refers to a characteristic interval of coherent variation. (Roughly speaking, this is the average time of monotonic increase or decrease of the analyzed quantity). The Ornstein–Uhlenbeck process (OU-process) timescale given by Dexter et al. (2014) is in the time domain and is consistent with the definition provided by the structure function. For a sinusoidal function of frequency f , the time-domain timescale equals $0.5/f$. The equation $\tau_{\text{time domain}} = 0.5/f_{\text{spectral domain}}$ is in fact valid for any function because, through a Fourier analysis, we can describe any function as a superposition of sinusoids. As long as the consistency issue is taken into account, different definitions of timescales are equivalent. The relevant question is which timescale is more informative about physical properties. The “memory” or coherence time of a physical process can be better described by a structure-function timescale rather than by a spectral-domain one. Structure-function timescales are also the appropriate choice for the calculation of source sizes through causality arguments. Periodicities, on the other hand, are naturally quantified in the spectral domain. This paper uses the structure function to define a distance between mock data and observed data. However, our timescale parameter f_b is defined in Equation (1) and refers to the spectral domain.

Caution should also be exercised in comparing structure-function slopes. While originally defined as in Equation (7), the structure function is sometimes expressed as

$$V(\tau_i) = \sqrt{\frac{1}{n_i} \sum_{t_j, t_k} [F(t_j) - F(t_k)]^2} \quad (9)$$

for $\tau_i \leq (t_j - t_k) < \tau_{i+1}$

(e.g., Dexter et al. 2014). Conversion of the structure function slope into a PSD power-law index depends on the structure function definition used.

3.3. Approximate Bayesian Computation

ABC is an iterative method based on prior distributions of model parameters, two Monte Carlo sampling steps—one for picking a random parameter set and one for drawing a light

curve realization for this parameter set—and an acceptance step. The acceptance is based on the comparison of the distance between simulated and observed data with a threshold decreasing with each iteration. For each iteration, the two sampling steps are executed many times until n (e.g., $n = 500$) realizations are accepted. The parameters (called “particles”) of the accepted realizations represent the current best estimate of the posterior. This estimate is used to inform the parameter sampling of the next iteration. If the distance function is informative for the model at hand, the prior distribution is transformed into a close approximation of the posterior after a sufficient number of iterations. A detailed description and explanation of the algorithm and its implementation is given by Witzel et al. (2018, their Appendix B.4).

Our custom C++ implementation of the forward modeling and ABC algorithm includes a fast algorithm for repetitive calculations of structure functions (Witzel et al. 2018, Appendix C) and is MPI-integrated¹⁴ to be run on computing clusters. For this analysis, we modified the model to include photon statistics according to Equations (3) to (6). We ran the ABC on 680 cores of the VLBI correlator computer cluster of the Max Planck Institute for Radio Astronomy (Bonn) with 100,000 initial light curve drawings. From these, we selected the particles having the smallest distance values. The final ABC run is the result of 16 iterations with $n = 800$ particles and two iterations with $n = 10,000$ particles.

3.4. PSD Parameters of the X-Ray Variability

While testing the algorithm, it became clear that the ACIS-S data set did not add much information to this part of the analysis and could be safely ignored, which reduced the number of parameters by two ($\chi_S, e_S/e_I$). Additionally, we set $\chi_G = \chi_I$.

Model parameters and their priors and posteriors are listed in Table 2 and shown in Figure 5. While the posterior of e_G/e_I is a mere minor alteration of its prior, the posterior of the break frequency f_b is a constrained, peaked distribution. The posterior of the PSD slope γ is constrained at the lower end, but prior-dominated at higher values, and in Table 2 we report its lower limit. The log-normal parameters $\mu_{\log n}$ and $\sigma_{\log n}$ are highly correlated ($\sigma_{\log n} \approx -0.40 \cdot \mu_{\log n} - 1.45$). A more precise determination of these parameters is not possible, because of photon noise and pileup, which make the log-normal distributions for a range of $\mu_{\log n}$ and $\sigma_{\log n}$ combinations indistinguishable.¹⁵ The background levels $\chi_i \cdot e_i/e_I$ are reasonable when compared to X-ray spectra accumulated over all periods of time without flares.

The parameters of main interest are γ and f_b , shown in Figure 6 in more detail. Determinations of both parameters are significantly less precise for the Chandra data than for the NIR. This is not surprising, considering that in 18 yr only $\sim 34,000$ X-ray photons were detected. However, despite the lack of precision, the comparison with the NIR is fruitful: the X-ray contours are displaced toward higher slopes and slightly smaller break frequencies. While the two sets of contours can be interpreted as marginally consistent, Figure 6 suggests that the X-ray variability does not show as much power at high

¹⁴ MPI means “message-passing interface,” a standard for parallel computing.

¹⁵ Log-normal distributions sensitively depend on the exact position of the peak of the distribution, which in the presence of background photon noise is uncertain. In the absence of a precise location of the peak, the correlation of both parameters is governed by the variance.

Table 2
Priors and Posteriors of Bayesian Analysis

Parameter	Prior	Median of Posterior	Description
Analysis of X-ray PSD, log-normal model			
γ	flat on [0.2, 6.0]	>2.0	PSD slope
f_b [$10^{-3} \cdot \text{minute}^{-1}$]	flat on [1.0, 600.0]	$4.0^{+1.0}_{-1.0}$	correlation frequency
$\mu_{\log n}^a$	flat on [-13.5, 0.0]	$-11.4^{+1.6}_{-1.3}$	log-normal mean
$\sigma_{\log n}^a$	flat on [0.0001, 6.0]	$3.1^{+0.5}_{-0.6}$	log-normal standard deviation
χ_I [$10^{-3} \cdot \text{cps}$] ^a	Gaussian ($\mu = 4.87, \sigma = 0.5$)	4.9 ± 0.5	contribution of steady X-ray source, ACIS-S/HETG
e_G/e_I	Gaussian ($\mu = 0.38, \sigma = 0.05$)	0.41 ± 0.04	effective area of ACIS-S/HETG rel. to ACIS-I
Analysis of X-ray, NIR, and submm data, synchrotron-SSC model			
γ_e	flat on [1.5, 4.0]	$2.95^{+0.21}_{-0.19}$	power-law index of electron energy distribution
B_0 [G]	log flat on [$10^{-2}, 10^3$]	$8.5^{+1.8}_{-1.4}$	minimum magnetic flux density
L_0 [R_S]	log flat on [$10^{-2}, 10^3$]	2.7 ± 0.3	maximum physical source diameter
γ_{slow}	flat on [1.2, 6.5]	4.9 ± 0.8	PSD slope of slow Gaussian process
γ_{fast}	flat on [1.2, 9.5]	$2.1^{+0.4}_{-0.3}$	PSD slope of fast Gaussian process
$f_{b,\text{slow}}$ [$10^{-2} \cdot \text{minute}^{-1}$]	flat on [0.001, 2.0]	$0.74^{+0.13}_{-0.12}$	correlation frequency of slow Gaussian process
$f_{b,\text{fast}}$ [$10^{-2} \cdot \text{minute}^{-1}$]	flat on [0.001, 2.0]	$1.21^{+0.29}_{-0.27}$	correlation frequency of fast Gaussian process
μ_{slow}^b	flat on [0, 10]	$-5.77^{+0.19}_{-0.26}$	log-normal mean of slow process
σ_{slow}^b	flat on [0.001, 1.5]	$0.616^{+0.075}_{-0.059}$	log-normal standard deviation of slow process
μ_{fast}^b	flat on [-7.5, 4.0]	$-1.65^{+0.66}_{-0.57}$	log-normal mean of fast process
σ_{fast}^b	flat on [0.001, 3.0]	$1.55^{+0.46}_{-0.40}$	log-normal standard deviation of fast process
χ_I [$10^{-3} \cdot \text{cps}$] ^a	flat on [0.0, 5.0]	$1.1^{+0.8}_{-0.5}$	contribution of steady, extended X-ray source
σ_{Keck} [mJy]	Gaussian ($\mu = 0.015, \sigma = 0.008$)	0.018 ± 0.006	measurement noise of the Keck observations
σ_{VLT} [mJy]	Gaussian ($\mu = 0.031, \sigma = 0.008$)	0.034 ± 0.005	measurement noise of the VLT observations
σ_{IRAC} [mJy]	Gaussian ($\mu = 0.65, \sigma = 0.4$)	$0.654^{+0.041}_{-0.047}$	measurement noise of the IRAC observations
$\sigma_{345 \text{ GHz}}$ [Jy]	Gaussian ($\mu = 0.1, \sigma = 0.1$)	0.10 ± 0.01	measurement noise of the 345 GHz (APEX and SMA)
$\sigma_{230 \text{ GHz}}$ [Jy]	Gaussian ($\mu = 0.15, \sigma = 0.1$)	0.08 ± 0.01	measurement noise of the 230 GHz (SMA and ALMA)
e_S/e_I	Gaussian ($\mu = 1.12, \sigma = 0.04$)	1.12 ± 0.03	effective area of ACIS-S relative to ACIS-I
e_G/e_I	Gaussian ($\mu = 0.38, \sigma = 0.02$)	0.38 ± 0.01	effective area of ACIS-S/HETG relative to ACIS-I
Δ_{AK}	Gaussian ($\mu = 0.0, \sigma = 0.1$)	-0.01 ± 0.07	modification of <i>K</i> -band extinction
Δ_{AM}	Gaussian ($\mu = 0.0, \sigma = 0.14$)	$0.09^{+0.08}_{-0.07}$	modification of <i>M</i> -band extinction
Posteriors (medians and 16% and 84% quantiles) for the synchrotron-SSC model derived from 1000 light curves of 700 minutes duration each			
$B(t)$ [G]		$12.7^{+4.1}_{-2.9}$	time-dependent magnetic flux density
$L(t)$ [R_S]		$2.20^{+0.32}_{-0.30}$	time-dependent source size
$n_e(t)$ [10^7 cm^{-3}]		$4.3^{+6.9}_{-2.7}$	time-dependent electron density
$\alpha_{\text{NIR}}(t)$		$-1.65^{+0.29}_{-0.30}$	time-dependent NIR spectral index (2.2–4.5 μm)
$\nu_m(t)$ [GHz]		308^{+96}_{-75}	time-dependent self-absorption turnover
$\nu_2(t)$ [THz]		54^{+11}_{-29}	time-dependent cooling cutoff frequency

Notes.

^a $\mu_{\log n}$, $\sigma_{\log n}$, and χ are data descriptive quantities representing the observed count rate in cps at the Chandra detector, i.e., after interstellar extinction (Equations (2) and (3)). They are linearly correlated with $\sigma_{\log n} \approx -0.40 \cdot \mu_{\log n} - 1.45$.

^b μ and σ are the log-normal parameters in Equations (38) and (39). They describe the time series of optically thin synchrotron flux density $S_{\text{thin}}(t)$ in Jy at $\nu_{\text{NIR}} = 136,269$ GHz, and the exponential cutoff frequency $\nu_2(t)$ in units of ν_{NIR} .

frequencies as the NIR. This interpretation of suppressed power of the fastest variability turns out to be a clue to understanding the NIR–X-ray correlations and allows us to identify a radiative model that can explain many aspects of the rapid (i.e., minutes to hours) variability of Sgr A* from the submm to the X-rays.

3.5. Simulations of NIR–X-Ray Correlations

Figure 7 demonstrates how PSD parameters affect NIR–X-ray correlations. The figure shows two light curves from the same set of Gaussian random numbers with two different PSDs. In the previously mentioned method of Timmer & Koenig (1995), a random number—one for each relevant frequency in the Fourier representation of the light curve—is multiplied by the square root of the PSD at the corresponding

frequency. By, let us say, using a broken power-law PSD model for one of the light curves and a second PSD with the same break frequency but a steeper slope for the other, we can generate pairs of light curves. Figure 8(a) shows a flow chart for the procedure. Because the curves use the same random numbers, they are strongly correlated. However, one of them is the low-pass-filtered version of the other. When we apply two, e.g., log-normal transformations (with the one applied to the slower process showing a heavier tail), the slower, more nonlinear process (representative of the X-ray variability) will always have a counterpart in the faster process (representative of the NIR), while the faster process shows a lot of peaks having a wide range of levels with (almost) no counterpart. When peaks are seen in both curves, their arrival times are strictly correlated, but there is little correlation in peak flux.




Figure 5. Results for the Bayesian structure function fit for the X-ray assuming a log-normal flux-density probability density function. Contours show the joint (posterior) probability density for each parameter pair, and panels along the right upper edge show histograms of the marginalized posterior of each parameter defined in Table 2. For each histogram, the dashed lines mark the 16%, 50%, and 84% quantiles.




Figure 6. Credible contours (68%, 95%, 99%) for the parameters γ and f_b . Blue shows the results of the present analysis for X-ray data (Table 2), and red shows results of the NIR analysis of Witzel et al. (2018, their Case 3). In both analyses, the posteriors have been marginalized over all other parameters.

This phenomenology mimics the NIR–X-ray observations (Genzel et al. 2010; Fazio et al. 2018) and becomes even more realistic when we include photon statistics, pileup effect, and a constant Poissonian background for the slower process.

4. Synchrotron Self-Compton Scattering and the Evolution of the Optical Depth

The theory of astronomical synchrotron sources was developed in the 1950s and 1960s (Ginzburg 1951; Shklovskii 1952, 1953; Ginzburg & Syrovatskii 1965). This paper uses the later, systematic development of synchrotron and synchrotron-SSC theory by Kardashev (1962), van der Laan (1966), Gould (1979), Marscher (1983), and Band & Grindlay (1985). In particular, we assume familiarity with the excellent overview article by Moffet (1975).

Based on flux-density levels, timing properties, and correlations, SSC scattering is a good candidate to explain Sgr A*’s X-ray variability. SSC becomes relevant under the same circumstances—high luminosity in a source of small size (Moffet 1975)—as synchrotron self-absorption, and therefore we discuss the coevolution of SSC scattering and synchrotron opacity. A simple model of injection, compression, and expansion can account for many aspects of the time-variable SED of Sgr A*. A semi-analytic approach allows us to co-fit the structure functions at all wavelengths considered in this study.




Figure 7. Multiwavelength mock light curves generated from the X-ray data model. The upper panel shows a 50,000 minute realization of log-normal NIR flux densities, and the lower panel presents X-ray count rates during the same time interval. Simulated count rates are modeled to match observations and include the steady X-ray background, pileup, and Poisson noise. (Flux-density normalization and the extinction correction are implicit in the model parameter μ .) Figure 8(a) illustrates the procedure for generating the mock light curves.

4.1. The NIR–X-Ray Correlation

As mentioned before, several studies have found the flux-density peak levels of synchronous NIR and X-ray flares to be consistent with the prediction of simple SSC models. Furthermore, Witzel et al. (2018) showed evidence that the NIR variability is at least partially caused by a variable, exponential synchrotron cooling cutoff close to the NIR; i.e., the synchrotron spectrum of flux densities is of the form:

$$S_\nu(t) = I_\nu \exp \left[- \left(\frac{\nu}{\nu_2(t)} \right)^{\frac{1}{2}} \right] \Delta\Omega, \quad (10)$$

where the source subtends a solid angle of $\Delta\Omega = \pi/4 \cdot (L/D_s)^2$ with L the source diameter and D_s the distance to the source. I_ν is the synchrotron intensity at frequency ν , and $\nu_2(t)$ is the variable cutoff frequency due to synchrotron cooling. Building on these ideas, the analytical equation for SSC radiation (Marscher 1983) also depends on $\nu_2(t)$, but very weakly because it is the argument of a logarithmic term:

$$S(E_{\text{keV}}) = d(-\alpha) \ln(\nu_2/\nu_m) \theta^{(4\alpha-6)} \nu_m^{(3\alpha-5)} S_m^{(4-2\alpha)} E_{\text{keV}}^\alpha, \quad (11)$$

where α is the optically thin spectral index,¹⁶ $d(-\alpha)$ are coefficients given by Marscher (1983), $\theta = L/D_s$ is the angular diameter of the source, ν_m is the self-absorption turnover frequency, S_m is the self-absorption flux density, and $E = h\nu$ is the photon energy in keV, with $h \approx 4.136 \times 10^{-9}$ keV GHz⁻¹ being the Planck constant. Equation (11) is valid for

$$5.5 \times 10^{-9} \gamma_1^2 \nu_m \lesssim E_{\text{keV}} \lesssim 0.2 [b(-\alpha)]^{-1} \theta^{-4} \nu_m^{-5} S_m^2, \quad (12)$$

with $b(-\alpha)$ coefficients given by Marscher (1983) and γ_1 being the minimum Lorentz factor of the electrons. For energies in this range, the SSC flux density is a power law with the same spectral index as the NIR. Figure 9 shows a synchrotron–SSC SED for realistic Sgr A* parameters. (Section 4.2 explains how to calculate I_ν from these parameters.)

The logarithmic dependence of $S(E_{\text{keV}})$ on the cutoff $\nu_2(t)$ implies that, if ν_2 is the origin of the fast variability in the NIR,

this variability power is suppressed in the X-ray light curves. However, it is not obvious that Equation (11) indeed suggests a clear separation of variability power: the X-ray flux density is highly variable, and this variability must be related to changes of the synchrotron source and spectrum itself, i.e., to changes of ν_m , S_m , and θ . This in turn means that I_ν and $\Delta\Omega$ in Equation (10) are time-dependent with similarly slow variability as the SSC flux density. Thus, following the idea presented in Section 3.5 and if SSC is responsible for the X-ray emission, the NIR flux density $S_{\nu_{\text{NIR}}}(t)$ is the product of two correlated processes, $I_{\nu_{\text{NIR}}}(t) \Delta\Omega(t)$ and $\exp \left[- \left(\frac{\nu_{\text{NIR}}}{\nu_2(t)} \right)^{1/2} \right]$, the former being the low-pass-filtered counterpart of the latter. This situation—the slow process (postulated from the X-rays) feeding back into the fast process (the NIR variability, described by the product)—is slightly more complex than the simple case considered in Section 3.5. However, for the right model for $I_\nu(t)$ and proper model parameters, the fast process will dominate this product, and the result will indeed be similar to the light curves of Figure 7.

According to Equation (11), $S(E_{\text{keV}})$ is not a direct function of the optically thin flux density at the K-band frequency $\nu_{\text{NIR}} = 136,269$ GHz,

$$S_{\text{thin}} = I_{\nu_{\text{NIR}}}(t) \Delta\Omega(t); \quad (13)$$

the self-absorption turnover flux density and turnover frequency also matter. In order to generate light curves in the submm that show the observed correlations and delays relative to the NIR, we need to discuss possible scenarios of coevolution of S_{thin} and the synchrotron opacity. In particular, it is not correct to assume $S_m \propto S_{\text{thin}}$ in the context of Equation (11) and to treat ν_m and θ as independent parameters (as often done in the literature) and then to argue that $S(E_{\text{keV}})$ depends on the optically thin flux density in a highly nonlinear way due to the term $S_m^{(4-2\alpha)}$.

4.2. Evolution of the Synchrotron Source and Synchrotron Self-absorption

Following the formalism and notation of Moffet (1975), the solution of the radiative transfer equation $dI_\nu = (\epsilon_\nu - \kappa_\nu I_\nu) dz$ through a homogeneous slab of material with constant and

¹⁶ We define α such that $S_\nu \propto \nu^\alpha$. Marscher (1983) used the opposite sign convention.




Figure 8. Flow chart of simulation algorithms for modeling light curves. Upper panel (a) shows an example of low-pass correlated Gaussian processes under transformation with log-normal distributions. The generation of each Gaussian process from random numbers follows the procedure described by Timmer & Koenig (1995). The two PSDs applied are both broken power laws, one of which has—by means of a steeper slope and/or a smaller break frequency—significantly lower power at short variability timescales. Lower panel (b) shows a similar low-pass correlated set of log-normal trigger processes that are then used to inform our synchrotron-SSC model. The radiative model combines the fast and slow processes according to the equations developed in Section 4.4 and lets us derive time series of flux densities at any given wavelength.




Figure 9. Model radio-to-X-ray SED for Sgr A*. Orange shows the synchrotron component, blue the SSC component, and red their sum. The SED is shown as if interstellar extinction were zero. Model parameters (see Table 2 for definitions) are $n_e = 1.066 \times 10^8 \text{ cm}^{-3}$, $B = 10 \text{ G}$, $L_0 = 2 R_S$, $\nu_2 = 20 \text{ THz}$, $E_{\min} = 0.02 \text{ GeV}$, and $\gamma = 2.4$. The vertical dashed lines mark the frequencies 230, 345 GHz, and ν_{NIR} . The blue rectangle marks the 2–8 keV X-ray band and shows typical peak flux densities during X-ray flares. The increasing and decreasing flanks of the SSC spectrum (where the spectrum is not described by a simple power law according to Equation (11)) are illustrated by arbitrary exponential cutoffs. The precise shape of the SSC spectrum in these flanks depends on many assumptions, including the density profile of the source, and is beyond the scope of this study. The spectrum shown is not a fit to any actual data. In particular, it combines low-level NIR flux densities with detectable X-ray emission, a condition that has not been observed.

isotropic emissivity ϵ_{ν} , absorption κ_{ν} , and thickness z_t is

$$I_{\nu}(z_t) = \frac{\epsilon_{\nu}}{\kappa_{\nu}} [1 - \exp(-\kappa_{\nu} z_t)]. \quad (14)$$

Table 3
Constants

Constant	Value	Unit
C_1	16.08	MHz $\mu\text{G}^{-1} \text{ GeV}^{-2}$
C_2	1.16540×10^{-26}	GeV $\mu\text{G}^{-1} \text{ sr}^{-1}$
C_3	1.19732×10^{-10}	$\text{yr}^{-1} \mu\text{G}^{-2} \text{ GeV}^{-1}$
C_4	1.358688×10^{-10}	mas^{-1}
C_5	6.241506×10^{-17}	GeV Jy $^{-1}$
C_6	0.0098087	$\text{mas} (1.2 \times 10^{10} \text{ m})^{-1}$
C_7	525949	minute yr^{-1}
R_S	1.2×10^{10}	m
D_s	2.523×10^{20}	m
k_X	3.39	(factor)
A_K	2.46	magnitudes
A_M	1.00	magnitudes

For a synchrotron source of homogeneous electron density in a tangled magnetic field, assuming a power-law electron energy distribution and an isotropic pitch angle distribution,

$$\epsilon_{\nu} = \frac{1}{2} C_2 n_0 B^{(\gamma_e+1)/2} (C_1/\nu)^{(\gamma_e-1)/2} \cdot G'(\nu/\nu_1, \nu/\nu_2, \gamma_e), \quad (15)$$

and

$$\kappa_{\nu} = \frac{1}{2} c^2 C_2 C_1^{\gamma_e/2} n_0 B^{(\gamma_e+2)/2} \nu^{-(\gamma_e+4)/2} (\gamma_e + 2) \cdot G'(\nu/\nu_1, \nu/\nu_2, \gamma_e + 1). \quad (16)$$

Here, c is the speed of light, C_1 and C_2 are constants given in Table 3, $\gamma_e = 1 - 2\alpha$ is the power-law index of the electron




Figure 10. Example SEDs derived from Equations (10) and (14) to (18). SEDs shown in green and blue were derived assuming the same ratio $\eta = \theta/B^{1/4}$ but different electron densities n_e . For a given η , the self-absorption turnover with the coordinates (ν_m, S_m) lies on one of the dashed gray lines having a power-law index of 2.5 (the slope of the optically thick part of the spectrum), independent of n_e or the optically thin flux-density levels. The purple dashed line marks the trajectory of the turnover under adiabatic expansion, i.e., the red SED could be the result of adiabatically expanding a synchrotron sphere that starts off with the green SED. The slope of the turnover trajectory related to adiabatic expansion can be significantly flatter than the iso- η trajectories (the exact slope being given by Equation (32)).

energy distribution, n_0 is the scaling factor of the electron energy distribution, and B is the magnetic flux density in Gauss. $G'(\nu/\nu_1, \nu/\nu_2, \gamma_e)$ is defined as

$$G'(\nu/\nu_1, \nu/\nu_2, \gamma_e) \equiv \frac{1}{2} \pi^{1/2} \frac{\Gamma[(\gamma_e + 5)/4]}{\Gamma[(\gamma_e + 7)/4]} \cdot \int_0^\infty x^{(\gamma_e - 3)/2} F(x) dx, \quad (17)$$

with

$$F(x) = x \int_x^\infty K_{5/3}(\eta) d\eta, \quad (18)$$

where $K_{5/3}(\eta)$ is the modified Bessel function of the second kind and $\Gamma(x)$ is the Gamma function. Here, we set $\nu/\nu_1 = \infty$ and $\nu/\nu_2 = 0$. However, this does not imply that the maximal Lorentz factor γ_{\max} of the electrons is infinite. Equation (10) includes an exponential term to approximate the effect of a finite γ_{\max} . This approach allows us later to use a lookup table for $G'(\infty, 0, \gamma_e)$ without having to re-evaluate the integral for every frequency.

A main result of synchrotron theory is that the SED is peaked and asymptotically approaches two different power laws: the optically thick spectrum $I_\nu = \epsilon_\nu/\kappa_\nu$ with slope $+2.5$ on the low-frequency side of the peak, and the optically thin spectrum $I_\nu = \epsilon_\nu \cdot z_\nu$ with spectral index $\alpha = (1 - \gamma_e)/2$ on the high-frequency side. Typical values observed for radio galaxies are $\alpha \sim -0.7$. The peak is described by (ν_m, S_m) , where S_m is the intersection of the two power laws, i.e., an extrapolation that is not actually reached by the SED. Figure 10 shows some example synchrotron spectra.

Equation (11) for the SSC flux density makes the assumption of spherical symmetry. As mentioned, Equation (14) is formally derived for a slab of material, i.e., if $\Delta\Omega = \pi/4 \cdot (L/d)^2$, S_ν of Equation (10) describes a cylinder of diameter and height L homogeneously filled with electrons. However, Band & Grindlay (1985) demonstrated that the

difference between spectra derived with a slab approximation and with spherical symmetry is very small. (Of course, n_0 needs to be multiplied by 1.5 in order to account for the ratio between the volume of a cylinder and the volume of a sphere). Therefore, we will proceed with the slab approximation and use the correction factor of 1.5 when calculating electron densities from flux densities.

Moffet (1975) derived a central equation to link the self-absorption turnover of the synchrotron spectrum described by the observable quantities ν_m and S_m with the physical parameters θ and B :

$$\nu_m^{5/2} = c^2 C_1^{1/2} F' S_m B^{1/2} / (\pi/4 \cdot \theta^2) \quad (19)$$

with

$$F' = \frac{(\gamma_e + 2) \cdot G'(\nu/\nu_1, \nu/\nu_1, \gamma_e + 1)}{G'(\nu/\nu_1, \nu/\nu_1, \gamma_e)}. \quad (20)$$

This equation is similar to an equation of state and allows us to describe how the observable quantities change under evolution of the physical quantities. Because ν_m and S_m are not observable in our case—the self-absorption turnover of the compact component responsible for the X-ray and NIR variability is veiled by emission from the constant radio component—we express S_m in terms of the optically thin synchrotron flux density S_{thin} :

$$S_{\text{thin}} = \pi/4 \cdot \epsilon_{\nu_{\text{NIR}}} \cdot L \cdot \theta^3, \quad \text{and} \quad (21)$$

$$\nu_m^{5/2 - \alpha} = c^2 C_1^{1/2} F' S_{\text{thin}} \cdot (\nu_{\text{NIR}})^{-\alpha} B^{1/2} / (\pi/4 \cdot \theta^2). \quad (22)$$

Similarly, the SSC flux density defined in Equation (11):

$$S(E_{\text{keV}}) \propto \nu_m^{(3\alpha - 5)} S_m^{(4 - 2\alpha)} \theta^{(4\alpha - 6)} \quad (23)$$

$$\propto \nu_m^{(3\alpha - 5) - 2\alpha(\alpha + 2)} S_{\text{thin}}^{(4 - 2\alpha)} \theta^{(4\alpha - 3)}. \quad (24)$$

We are now in a position to discuss illustrative scenarios of source evolution and their consequences for the dependence of the X-ray and submm flux densities on S_{thin} . While none of these is directly applicable to the final model, they are useful to show the dependence of observables on physical quantities. In the following, we define

$$\eta \equiv \theta/B^{1/4}. \quad (25)$$

4.2.1. Case 1

Let us assume $\Delta\eta = 0$, $\Delta S_{\text{thin}} \neq 0$. In this case, $S_m \propto (\nu_m)^{2.5}$, i.e., the turnover moves along the iso- η lines shown in Figure 10. These iso- η lines have the same slope as the optically thick part of the spectrum, and the flux densities at frequencies below the peak do not change. From Equation (22),

$$\nu_m \propto (S_{\text{thin}})^{\frac{2}{5 - 2\alpha}}, \quad (26)$$

and with $\Delta\theta = 0$ (i.e., no changes other than in n_0) and assuming $\alpha = -1$,

$$S(E_{\text{keV}}) \propto (S_{\text{thin}})^2. \quad (27)$$

This occurs because the increase in ν_m partially counters that in S_m . This gives the dependence of $S(E_{\text{keV}})$ on S_{thin} a significantly lower exponent than the value 6 apparent from Equation (11).

4.2.2. Case 2

Next, consider the case $\Delta\nu_m = 0$, $\Delta S_{\text{thin}} \neq 0$, i.e., a constant turnover frequency under changing S_{thin} . In this case, $S_m \propto S_{\text{thin}}$, $S_{\text{thick}} \propto S_{\text{thin}}$, and

$$\eta \propto (S_{\text{thin}})^{1/2}. \quad (28)$$

If the change in η is mainly a change in θ —a corresponding change in B would have to be higher by the fourth power— $\Delta B = 0$ and $\alpha = -1$ imply

$$S(E_{\text{keV}}) \propto S_{\text{thin}}. \quad (29)$$

4.2.3. Case 3

As mentioned in the Introduction, a case of particular interest is source evolution through adiabatic expansion. The fundamental assumptions have been stated by Shklovskii (1960) in the context of supernova-remnant evolution:

$$B \propto \theta^{-2}, \quad (30)$$

and

$$E \propto \theta^{-1} \quad (31)$$

as θ changes with time. In contrast to synchrotron cooling, cooling by adiabatic expansion applies to electrons of all energies at the same rate set by the expansion speed. van der Laan (1966) showed that, in this case:

$$S_m(\nu_m) \propto (\nu_m)^{-\frac{7\gamma_e+3}{4\gamma_e+6}}. \quad (32)$$

For reasonable values of γ_e , S_m has a flatter index than the iso- η lines (Figure 10). Other relations are:

$$\nu_m \propto (S_{\text{thin}})^{\frac{2}{5-2\alpha}} \cdot \eta^{\frac{-4}{5-2\alpha}}, \quad (33)$$

$$B \propto (S_{\text{thin}})^{\frac{2}{3-4\alpha}}, \quad \text{and} \quad (34)$$

$$\theta \propto (S_{\text{thin}})^{-\frac{1}{3-4\alpha}}. \quad (35)$$

For the X-ray flux density:

$$\begin{aligned} S(E_{\text{keV}}) &\propto \eta^{-4(5-3\alpha+2\alpha(\alpha-2))/(5-2\alpha)} \\ &\cdot (S_{\text{thin}})^{(4-2\alpha)+2(3\alpha-5-2\alpha(\alpha+2))/(5-2\alpha)} \\ &\cdot \theta^{(4\alpha-6)}. \end{aligned} \quad (36)$$

For $\alpha = -1$,

$$S(E_{\text{keV}}) \propto (S_{\text{thin}})^{1.71}. \quad (37)$$

While S_m decreases with S_{thin} , ν_m decreases as well. Therefore, the optically thick flux densities show a temporary increase some time after the time of maximum flux density at optically thin frequencies.

In all three cases discussed here, under the physical constraints expressed in Equation (22), the actual dependence of the SSC flux density on S_{thin} is only weakly nonlinear, with an exponent of 1–2, and does not come close to the apparent exponent of $(4 - 2\alpha)$ in Equation (11).

No one of the discussed cases of source evolution can by itself reproduce the submm to NIR phenomenology. Case 1 does not predict NIR-correlated submm variability at all. Case 2, a scenario of nonadiabatic compression and expansion, predicts a direct proportionality of submm flux-density changes to S_{thin} , which is not observed. Case 3 describes only the cooling of the synchrotron source by expansion, i.e., decaying

S_{thin} , and the resulting propagation of delayed peaks toward longer observing wavelengths. The opposite, increasing S_{thin} under adiabatic compression, would be difficult to understand physically. Furthermore, adiabatic compression would result in leading maxima at submm wavelengths, which are also not observed.

4.3. A Simple Source Model

To generate submm light curves that correspond qualitatively and quantitatively to the observed data, we here propose a simple three-step process of electron injection, compression of magnetic field lines, and expansion. For simplicity, we model this three-step process as a cyclic process in a single zone. Analysis of this process makes the assumption that if S_{thin} is rising fast, the variability is injection dominated, and the opacity is developing according to Case 1. If S_{thin} is falling quickly, the source is adiabatically expanding according to Case 3. Just before the time of peak flux density, the source is subjected to an episode of (nonadiabatic) compression and particle escape, i.e., a compression of the magnetic field lines mainly without electron heating (similar to Case 2 but with changing B). Here, particle escape is necessary because close to the peak, S_{thin} changes little, but in our model B increases. Because ϵ_ν depends on B and on the product $n_0\theta^3$ (which is constant under particle conservation), particles must escape.¹⁷ The reason for using a cyclic model is simplicity. However, the cycles we are modeling are not sequential, in the sense that an individual cycle has to return to its starting point before a new cycle can start. A change of $\dot{S}_{\text{thin}}(t)$ from negative to positive will start a “new cycle.” To translate this to a picture of multiple regions, a new region will start to dominate where the old left off. While a new source region might start with a different set of source parameters, statistically the presented approach is equivalent, at least with respect to the posteriors of the source parameters marginalized over time.

Our three-step source evolution is illustrated in Figure 11, including transitions between injection and compression and compression and expansion. The model qualitatively predicts submm light curves correlated with the NIR and shows a range of possible delays between the two bands. It is, indeed, the simplest source evolution model that can reproduce the observed phenomenology without generating artificial symmetries in the light curves or too-strict correlations.

4.4. Model Implementation

We implemented the Section 4.3 synchrotron-SSC mechanism in a semi-empirical model. The model is semi-empirical because

1. the opacity evolution model is derived from submm-NIR phenomenology;
2. it is based on two generic, log-normal red-noise processes drawn from the same random numbers according to Section 3.5, for which we can provide only empirical reasoning.

¹⁷ A decreasing source size with increasing magnetic flux density B can be created naturally when magnetic flux lines are compressed with the bulk of electrons not following that compression. As a consequence, a smaller volume filled with electrons is interacting with the stronger magnetic field. The electrons outside this active zone then have “escaped” the synchrotron region. The details of this compression phase depend on the characteristics of the processes that govern the electron budget (injection, escape, and cooling).




Figure 11. Illustration of the evolution of our simple model for the compact component in Sgr A*, denoted by the dashed circle. Upper panels 1–6 are a cartoon of a possible sequence of injection, compression, and expansion. The steps are: (1) and (2) injection of electrons with a nonthermal energy distribution into a spherical region with a uniform magnetic flux density giving detectable submm and NIR emission; (3) and (4) further injection, compression, and increase of magnetic flux under electron escape giving increased self-Compton scattering efficiency and detectable X-ray emission; (5) minimum size with highest magnetic flux density under ongoing injection followed shortly after by maximum in electron density and in NIR and X-ray emission; and (6) adiabatic expansion with little or no injection, giving maximum in submm emission. The funnels on the two sides of the synchrotron sphere are idealized and do not imply actual symmetry of the inflow or outflow of particles. In particular, particle escape under compression might well be isotropic. The bottom panel shows the corresponding cycle of the self-absorption turnover position in the S_m – ν_m plane. Steps are labeled as above, and the diagonal dashed lines are lines of constant η .

One of the red-noise processes is the slow process drawn from a PSD with the parameters γ_{slow} and $f_{b,\text{slow}}$:

$$S_{\text{thin}}(t) = \exp[\sigma_{\text{slow}} y_{\text{slow}}(t) + \mu_{\text{slow}}] \text{Jy}, \quad (38)$$

with $y_{\text{slow}}(t)$ the slow Gaussian process and μ_{slow} and σ_{slow} the log-normal parameters. This process represents the variability of the optically thin part of the synchrotron spectrum and the timescales set by the injection process and expansion cooling. We chose a log-normal PDF because the synchrotron equations are power laws, and a log-normal process remains log-normal under multiplication and exponentiation. The other red-noise process is the fast process drawn from a PSD with the parameters γ_{fast} and $f_{b,\text{fast}}$:

$$\nu_2(t) = \nu_{\text{NIR}} \cdot \exp[\sigma_{\text{fast}} y_{\text{fast}}(t) + \mu_{\text{fast}}] + \max(\nu_m, \nu_{\text{min}}), \quad (39)$$

with $y_{\text{fast}}(t)$ being the fast Gaussian process, and μ_{fast} and σ_{fast} its log-normal parameters. This process represents the variable location of the synchrotron cooling cutoff that is the result of the two competing processes at the high-frequency tail of the synchrotron spectrum: the tail of the injection spectrum and the cooling through synchrotron emission. In this case, we used a log-normal PDF as well to ensure $\nu_2(t) > 0$. However, it is a three-parameter log-normal PDF, to account for the fact that the cooling cutoff frequency should be larger than $\max(\nu_m, \nu_{\text{min}})$ at all times, i.e., larger than both the self-absorption turnover frequency and the transition frequency from synchrotron cooling to expansion cooling. The latter can be derived from the equation for the critical frequency for a given electron energy,

$$\nu_c = C_1 B E^2, \quad (40)$$

and the equation for the time after which an electron of initial energy E_0 has cooled to $1/e$ of its initial energy,

$$t_{1/e} = (C_3 B^2 E_0)^{-1} (e - 1) \cdot C_7. \quad (41)$$

We then define ν_{min} as the frequency in GHz where the $1/e$ cooling time is equal to the correlation timescale of the slow process:

$$\nu_{\text{min}} = 10^{-3} C_1 C_7^2 \cdot \frac{f_{b,\text{slow}}^2 (e - 1)^2}{(C_3)^2 B^3}. \quad (42)$$

Figure 2 of Moffet (1975) shows the spectrum of an individual electron. The definition above guarantees that the synchrotron cooling break ν_2 can occur only at frequencies that have enough time to cool through synchrotron emission during an episode of injection, compression, and expansion (with a typical duration of $1/f_{b,\text{slow}}$). We define

$$\beta[\eta(t)] \equiv c^2 \cdot (C_1)^{1/2} \cdot [\eta(t) \cdot C_4]^{-2} \cdot F' \times 10^{-12} \cdot C_5, \quad (43)$$

with

$$\eta(t) = \frac{\theta(t)}{B(t)^{1/4}} \quad (44)$$

and F' from Equation (20). We can now derive expressions for the combined quantities

$$\kappa_\nu L = \nu^{-(\gamma_e+4)/2} \cdot \beta[\eta(t)] \cdot S_{\text{thin}}(t) \cdot \nu_{\text{thin}}^{(\gamma_e-1)/2} \quad (45)$$

and

$$\frac{\epsilon_\nu}{\kappa_\nu} \cdot \Delta\Omega = \frac{\nu^{5/2}}{\beta[\eta(t)]}. \quad (46)$$

With these and $\nu_2(t)$, Equations (14) and (10) give the time-variable synchrotron flux density at each frequency. Similarly, $\eta(t)$ and $S_{\text{thin}}(t)$ give

$$\nu_m(t) = 0.001 \cdot [C_5 \cdot \beta[\eta(t)] \cdot S_{\text{thin}}(t) (\nu_{\text{NIR}} \cdot 1000)^{\frac{\gamma_e-1}{2}}]^{2/(\gamma_e+4)} \quad (47)$$

and

$$S_m(t) = S_{\text{thin}}(t) \cdot [\nu_m(t) / \nu_{\text{NIR}}]^{-\frac{\gamma_e-1}{2}}, \quad (48)$$

and $\theta(t)$ and Equation (11) give the power-law section of the SSC SED.

We parameterized the cyclic source evolution model by linear functions with variable slopes in the logarithmic (S_{thin}, θ) and (S_{thin}, B) planes (i.e., power laws with variable indices):

$$B(t) = B_0 \left[\frac{S_{\text{thin}}(t)}{S_{\text{thin},0}} \right]^{u(t)} \quad (49)$$

$$u(t) = \frac{1}{2} \left(\frac{2}{2\gamma_e + 1} \right) \quad (50)$$

$$\left\{ \tanh \left[\frac{\dot{S}_{\text{thin}}(t)}{2S_{\text{thin}}(t)} \cdot \min \right] + 1 \right\} \quad (51)$$

$$L(t) = L_0 \cdot \left[\frac{S_{\text{thin}}(t)}{S_{\text{thin},0}} \right]^{-v(t)} \quad (52)$$

$$v(t) = \frac{1}{2} \left(\frac{1}{2\gamma_e + 1} \right) \quad (53)$$

$$\left\{ \tanh \left[\frac{\dot{S}_{\text{thin}}(t)}{2S_{\text{thin}}(t)} \cdot \min \right] + 1 \right\}. \quad (54)$$

(The super-dot notation means time derivative in units of minute⁻¹.) In Equation (49) for B , the power-law index u was constrained to the interval of $[0, \frac{2}{2\gamma_e+1}]$. On average, it is $\frac{1}{2\gamma_e+1}$ with deviations to either a larger or smaller index depending on the fractional derivative $\frac{\dot{S}_{\text{thin}}(t)}{S_{\text{thin}}(t)}$ in units of min⁻¹. Similarly, the index v in Equation (52) for L was constrained to $[0, \frac{1}{2\gamma_e+1}]$, averaging $\frac{1}{4\gamma_e+2}$. B_0 and θ_0 are scaling parameters with units G and R_S , respectively. This cyclic model is a simple empirical scenario based on the cases in Section 4.2. The extreme Cases 1 and 3 are asymptotically approached for very rapid changes. Depending on $\dot{S}_{\text{thin}}(t)$, our implementation allows for mixing of the different cases (e.g., slight compression during injection, etc.) and smooth transitions between the three steps of the source evolution. Figure 12 shows trajectories in the ($S_{\text{thin}}, \theta_0$) plane, the (S_{thin}, B_0) plane, and the (ν_m, S_m) plane calculated according to Equations (49) to (54).




Figure 12. Cyclic evolution of model parameters. Top panel shows source size L vs. optically thin synchrotron flux density (neglecting interstellar extinction), middle panel shows magnetic flux density vs. the same, and bottom panel shows self-absorption turnover in the (ν_m, S_m) plane. All parameter time sequences (blue loops) were derived according to Equations (49) to (54). The gray dashed lines indicate the trajectories for pure injection ($\Delta\eta = 0$) and pure adiabatic expansion as labeled. The circular arrows indicate the direction of time evolution in each panel.

The dependence of $S(E_{\text{keV}})$ on S_{thin} for the peak flux densities (and the flux-density minima) comes from setting $\dot{S}_{\text{thin}}(t) = 0$. Then

$$\nu_m \propto (S_{\text{thin}})^{\frac{2}{5-2\alpha}} \cdot \eta^{\frac{4}{5-2\alpha}}, \quad (55)$$

$$B \propto (S_{\text{thin}})^{\frac{1}{3-4\alpha}}, \quad (56)$$

$$\theta \propto (S_{\text{thin}})^{-\frac{1}{(6-8\alpha)}}, \quad (57)$$

and for $\alpha = -1$,

$$S(E_{\text{keV}}) \propto (S_{\text{thin}})^{1.86}. \quad (58)$$

With this model implementation, we are able to predict light curves at submm and NIR light curves according to Equation (10) and X-ray light curves according to Equation (11) and compare them via a suited distance function to the observed light curves in all three wavelength regimes. A high-level block diagram for generating light curves according to this model is shown in Figure 8(b). For the X-rays, to

convert predicted flux densities to count rates, we used the conversion from the pileup-free count rate to the absorbed energy flux in the 2–8 keV band $c_{\text{conv}} = 2.81 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1} \text{ cps}^{-1}$ (Yuan & Wang 2016). To account for interstellar extinction, we multiplied by 3.39, the effective correction derived from the luminosity calculated from c_{conv} and the unabsorbed luminosity (Yuan & Wang 2016), and used the relation

$$\Lambda(t) = \frac{S(E_{\text{keV}})}{\text{Jy}} \frac{\nu_{8 \text{ keV}}^\alpha - \nu_{2 \text{ keV}}^\alpha}{\nu_{5 \text{ keV}}^\alpha (1 + \alpha) \cdot 2.81 \cdot 3.39 \times 10^{12}} \text{ cps.} \quad (59)$$

Observed count rates were then calculated according to Equations (3) to (6), but we assumed that the steady X-ray background is equal for all three detector modes of Chandra (i.e., $\chi_I = \chi_G = \chi_S$). We used e_I rather than e_G as the reference effective area for Chandra.

The free parameters in our analysis are listed in Table 2. In order to compare the model predictions to observed NIR flux density, we applied extinction magnitudes A_K and A_M given in Table 3. Because these values have significant uncertainties, we let the fitting modify the extinction magnitudes via parameters ΔA_K and ΔA_M . We set parameter $S_{\text{thin},0} = 0.1 \cdot \exp(\mu_{\text{slow}} - \sigma_{\text{slow}}^2)$ Jy, the mode of the flux-density distribution of S_{thin} . This lets the fitter set $S_{\text{thin}}(t)$ wherever μ_{slow} and σ_{slow} say it should be while not having $B(t)$ (Equation (49)) or $\theta(t)$ (Equation (52)) take on improbable values. The exponent of 0.5 in the exponential term of Equation (10) and the factor 0.5 in the argument of the tanh in Equations (51) and (54) are fiducial parameters, which we did not fit. A mismatch in the exponent can be absorbed by the parameters μ_{fast} and σ_{fast} , at least where $\nu_2(t) \gg \max(\nu_m, \nu_{\text{min}})$. Where $\nu_2(t)$ comes close to its minimum, we estimate the error to be not greater than 10%. However, close to the minimum of $\nu_2(t)$, the NIR flux density is far below the detection limit, and the variability is entirely dominated by the measurement noise. The factor of 0.5 in the argument of the tanh gives reasonable results, and testing showed that the fitting results of the other parameters are not strongly influenced by its exact value. To fit for this parameter, we would need detailed information on the covariance at several mm to submm wavelengths and a suited distance function. Our data are not sufficient to constrain this factor.

Finally, Equations (51) and (54) make use of the derivative $\dot{S}_{\text{thin}}(t)$. We take advantage of the fact that we are generating the mock light curves via FFT from Fourier coefficients. To compute the derivative from the same Fourier coefficients, we can use the fundamental relation from Fourier theory:

$$F \left[\frac{d}{dx} g(x) \right] = 2\pi i f_x G(F_x), \quad (60)$$

with F being the Fourier transform and $G(f_x) = F[g(x)]$. However, this equation cannot simply be applied to the discrete case, and for that we followed the method of wavenumber modification (Sunaina et al. 2018). This determines $\dot{y}_{\text{slow}}(t)$, and then with Equation (38) generates $\dot{S}_{\text{thin}}(t)$. In summary, the steps to create the cyclic synchrotron-SSC model are as follows:

1. $y_{\text{slow}}(t)$ is used to generate $S_{\text{thin}}(t)$ according to Equation (38);
2. $S_{\text{thin}}(t)$, $\nu_2(t)$, and $\eta(t)$ are used to derived the spectrum, according to Equations (45) and (46) together with (10)

and (14) (and for the SSC, Equations (47) and (48) together with (11)).

3. $B(t)$ and $L(t)$, and thus $\eta(t)$, are cyclic by construction and depend on $S_{\text{thin}}(t)$ and its derivative according to Equations (44) and (49)–(54).

4.5. The Distance Function

As in Section 3, to use ABC we need a distance function. Our distance function used several components to guide the ABC algorithm to a valid set of posteriors:

1. The distance between the predicted and the observed structure functions as defined in Equation (8). We calculated the structure functions and distances for the three NIR, two submm (lower panel of Figure 4), and three X-ray data sets separately. We added these eight distances multiplied by normalized weights, which were the quadratic difference between the logarithm of the maximum value and the logarithm of the minimum value of each of the eight observed structure functions V_j :

$$D_{\text{st}} = \frac{\sum_j [\max \log(V_j) - \min \log(V_j)]^2 \phi_j}{\sum_j [\max \log(V_j) - \min \log(V_j)]^2}, \quad (61)$$

with D_{st} being the total distance calculated from structure functions and ϕ_j the structure function distance for each data set. For each structure function, we used uniform weights $w_i = 1$ for all bins except for the last bins of the NIR and X-ray data (i.e., the bin with $\tau > 20$ minutes for the Keck data and the one with $\tau > 128$ minutes for all other NIR and X-ray data), which had $w_i = 3$. This approach led to uniformly converging fits at all time lags and in all bands. (See also the discussion by Witzel et al. 2018, Appendix B.2.)

2. To help to constrain the log-normal PDF of $\nu_2(t)$, we determined two other quantities to describe the empirical distribution of observed K -band flux densities: the skewness of the distribution and the fraction of positive flux densities. Both values are sensitive to the white noise characteristics, and we determined them only for the VLT data set, which is the larger of the two K -band data sets. The fraction of positive flux densities is 0.9. This value was derived from the most recent analysis of the intrinsic flux density distribution from the VLTI/GRAVITY interferometer (Gravity Collaboration et al. 2020b). VLTI/GRAVITY data show significantly less noise than the data from single-dish telescopes, and its 1 mas angular resolution makes it virtually free of source confusion. The VLTI/GRAVITY data show a median flux density of 1.1 ± 0.3 mJy dereddened, and the peak of the distribution is ~ 0.4 mJy dereddened. While the empirical distribution of VLTI/GRAVITY flux densities is always positive, convolving it with a Gaussian noise of $\sigma_{\text{NACO}} \approx 0.3$ mJy dereddened will result in a tail of $\sim 10\%$ negative flux densities. We counted the fraction of $n_{>0}/n$ points in each VLT mock light curve above zero and the quadratic difference:

$$D_{>0} = (n_{>0}/n - 0.9)^2. \quad (62)$$

For the skewness,

$$\text{Sk}_{\text{VLT,obs}} = \frac{\frac{1}{n} \sum_{i=1}^n (S_i - \bar{S})^3}{\left[\frac{1}{n-1} \sum_{i=1}^n (S_i - \bar{S})^2 \right]^{3/2}} = 3.44, \quad (63)$$

with S_i being the observed flux densities and \bar{S} the mean flux density. For each VLT mock light curve, we calculated the skewness and the quadratic difference

$$D_{\text{sk}} = (\text{Sk}_{\text{VLT,sim}} - \text{Sk}_{\text{VLT,obs}})^2 \quad (64)$$

and added the result to the distance function.

- Following Witzel et al. (2018, their discussion of Case 3 in Section 4.4), we additionally used spectral information from simultaneous K - and M -band data. At an average (observed) flux density $S_K = 0.15$ mJy, Witzel et al. found the ratio of (observed) NIR flux densities at M and K bands to be $\mathfrak{R}_{\text{obs}} = 12 \pm 0.5$. We calculated this ratio for each parameter set from our mock data. Because here we are not using simple log-normal PDFs for modeling the flux densities (as Witzel et al. did), we cannot derive this ratio analytically from model parameters. Instead, we simulated for each parameter set equally sampled, measurement-noise-free K - and M -band light curves of 10,000 minutes duration each and determined $\mathfrak{R}(M/K, S_K)$ as the average ratio over a suited flux-density range. The flux-density range of the Witzel et al. simultaneous data sets is ~ 0.07 – 0.27 mJy. However, it is difficult to determine background flux-density levels for single-dish data. After comparing the median of the NACO flux-density distribution with the newest VLTI/GRAVITY study (Gravity Collaboration et al. 2020b), we adopted a background correction factor based on a flux-density range of $[0.0$ – $0.17]$ mJy. We then defined the distance

$$D_{\text{sp}} = [\mathfrak{R}(M/K, S_K) - \mathfrak{R}_{\text{obs}} + \mathcal{N}(0, 0.5^2)]^2, \quad (65)$$

with $\mathcal{N}(0, 0.5^2)$ being a normal random number to account for the uncertainty in $\mathfrak{R}_{\text{obs}}$.

The final distance function was

$$D_{\text{tot}} = D_{\text{sf}} + 10 \cdot D_{>0} + D_{\text{sk}} + D_{\text{sp}}/50. \quad (66)$$

The factors of 10 and 50 are empirical and ensured that none of the additional distance terms starts to dominate D_{sf} . This distance function does not include metrics to quantify the degree of correlation between wavelengths, and we restricted the 230 GHz structure function to time lags < 50 minutes. We will come back to these two points in the discussion.

5. Results

We have implemented the described model and distance function in our C++ code. The analytic nature of the model allows us to calculate many realizations as large as the observed data set in reasonable time. However, because of the iterative nature of the ABC and the need for a sufficiently large particle system, we had to run our code on the VLBI compute cluster of the Max-Planck-Institut für Radioastronomie. The model and the distance function are the result of more than fifty test runs that all had to be executed over several days on 200–400 cores (10–20 nodes). Over the course of two years,




Figure 13. Performance of the ABC sampler. Ordinate is the acceptance rate, and abscissa is the 45% quantile of all distances (D_{tot} , Equation (66)) of the particular iteration. Toward the smallest distances, the sample found < 1 in 10,000 light curves that satisfied the distance threshold.

we estimate the total CPU time (including all runs necessary for testing, developing, and implementing the model) to be $\sim 10^6$ hr. To analyze the results, we also implemented a Python version of our model. This version allows us to quickly visualize and diagnose the ABC results, including a tool for generating animations of the timing of the SED, the light curves, and the derived physical parameters.

Table 2 lists the priors of all parameters. Priors for all physical parameters except B_0 and L_0 were flat and wide. For B_0 and L_0 , we used log-flat priors. Parameters for the log-normal red noise processes were also flat and wide. We labeled these parameters “slow” and “fast,” but we used the exact same priors for both PSDs. In other words, we did not force the PSDs to take on the filter relation discussed in Section 3.5. For parameters specifying instrumental characteristics, we used narrow Gaussian priors.

The final ABC run presented here is the result of 33 iterations, 32 with a particle size $n = 500$ and the last with $n = 5000$. (The model computations had to use fewer particles than the X-ray-only model because submm and NIR light curves had to be computed in addition to X-rays.) Figure 13 shows the evolution of the distance and the acceptance rate. We reached $D_{\text{tot}} < 0.09$ at an acceptance rate of 5×10^{-5} . (This means the final iteration generated 10^8 light curves in order to find 5000 it could accept.) The local maximum in the acceptance near $D_{\text{tot}} < 0.2$ is a hint that, at this stage, the algorithm still “learned” relevant information.

The posterior estimates of all parameters converged to well-constrained, peaked distributions. Figure 14 shows the posteriors and pairwise correlations of all parameters. The medians and 1σ credible intervals are listed in Table 2. Figure 15 shows the structure functions with the corresponding 1σ envelopes drawn from the posterior. The fit describes the observed structure functions well with only the structure function of the Spitzer/IRAC data lying partly outside the 1σ envelope. This discrepancy is not statistically significant, and part of it could be caused by our exact choice of synchrotron cooling cutoff (Equation (10)). The true shape of this cutoff depends on fine details of the process, e.g., the injection spectrum and the exact cooling mechanism. These depend on the geometry of the magnetic field lines and other specifics. The shape we chose




Figure 14. Results for the Bayesian structure function fit for the synchrotron-SSC model. Contours show the joint posterior probability density for each parameter pair, and panels along the right upper edge show histograms of the marginalized posterior of each parameter defined in Table 2. For each histogram, the dashed lines mark the 16%, 50%, and 84% quantiles. Instrumental σ parameters are as observed, not extinction-corrected.

predicts a slightly higher variability in the middle range of time lags at $4.5 \mu\text{m}$. In the absence of more specific information on the coevolution of the spectral indices in the IR bands, it is not possible to determine a more realistic scenario.

The physical parameters of the model are tightly constrained: $\gamma_e = 2.95 \pm 0.2$, $B_0 = 8.6^{+1.8}_{-1.4}$ G, and $L_0 = 2.7 \pm 0.3 R_S$. Figure 16 shows the posteriors of the PSD parameters. The PSD slope of the slow process comes out a little steeper in the combined model than when analyzing the X-ray data alone, but not significantly so. A bigger difference is between the fast and

slow process timescales and the timescale from the earlier NIR analysis. However, both timescales (135^{+26}_{-20} minutes and 82^{+28}_{-16} minutes, respectively) are consistent with the earliest report of the NIR break timescale of 154^{+124}_{-87} by Meyer et al. (2009).

In order to determine posteriors of observable and physical quantities, we used the model to construct 1000 mock light curves of 700 minute duration in all bands. The results are given in Table 2. The mock light curves were then “observed” at uniform intervals, unlike the real data. This may give the




Figure 15. Structure functions (Equation (7)) for the eight data sets. Panels from top to bottom show NIR, X-ray, and submm, all as observed, not corrected for extinction or pileup. Solid lines show the observed SFs (as presented in Figure 4), and corresponding dashed curves show the median of 10,000 SFs for the respective mock data sets. The shaded envelopes denote the model 68% credible intervals for each time lag. The details of generating the SFs, including the choice of time-lag ranges, are described in Section 4.5. The slope of each SF relates to the slope of the underlying PSD but also depends on the underlying white noise level.

mock data lower variance than the real data. Figure 17 compares the posterior of $(\log S_{\text{NIR}}, \log S_X)$ to observed flux-density pairs. Figure 18 shows calculated NIR and submm spectra indices. The model NIR spectral index becomes flatter toward brighter flux densities with a maximum around $\alpha_{\text{NIR}} = -1$, while the submm spectral index is a wide scatter




Figure 16. Credible contours (68%, 95%, 99%) for the parameters γ_1 and f_b . Green and orange show the PSD parameters of our synchrotron-SSC model (Section 4 and Table 2), green showing the slow process (identified with X-rays) and orange the fast process (identified with NIR). Blue contours are for our X-ray-only analysis (Section 3.5), and red contours are for the NIR analysis of Witzel et al. (2018, their Case 3). All posteriors have been marginalized over all other parameters. The horizontal gray dotted line marks the frequency that corresponds to a periodic timescale of 50 minutes, which is the orbital timescale of the astrometric motion measured by Gravity Collaboration et al. (2018b).

cloud terminated sharply at $\alpha_{\text{submm}} = +2.5$. These are as expected for a synchrotron-SSC source.

Figures 19 and 20 show the posterior distributions of physical and observable parameters based on the models and mock light curves. The delay (Figure 19) was computed from the maximum of the cross-correlation function of the 230 GHz flux densities with S_{thin} for each 700 minute light curve. We accepted only delays with peaks at least a factor of 5 higher than the standard deviation of the cross-correlation function. The high-energy electron density of Figure 20 was calculated according to:

$$n_0(t) = \frac{12 \cdot S_{\text{thin}}(t)}{\pi \theta^3 \cdot d} \frac{(\nu_{\text{NIR}}/C_1)^{\frac{\gamma_e-1}{2}}}{C_2 \cdot B^{\frac{\gamma_e+1}{2}} G'(\infty, 0, \gamma_e)}, \quad (67)$$

$$n_e(t) = \frac{n_0(t)}{(1 - \gamma_e)} \cdot \left\{ \left[\frac{\nu_2(t)}{C_1 \cdot B} \right]^{(1-\gamma_e)} - E_{\text{min}}^{(1-\gamma_e)} \right\} \quad \text{for } \gamma_e > 1. \quad (68)$$

All other quantities were calculated according to the equations in Section 4.4.

The synchrotron-SSC model gives reasonable values for physical and observable parameters. The submm to NIR time delay (Figure 19) is a positively skewed distribution in the interval of $[0, 100]$ minutes and peaks at ~ 20 minutes. Figure 20 shows other parameters. The magnetic flux density B varies between 6 and 30 G and peaks around 13 G. The source diameter $L(t)$ is a symmetric distribution around $2.2 R_S$ with values between 1.3 and $3.4 R_S$. The synchrotron electron density n_e is distributed around $4 \times 10^7 \text{ cm}^{-3}$ and varies by somewhat more than an order of magnitude in both directions. The spectrum shows a self-absorption turnover ν_m most often near 280 GHz but ranging between 100 and 750 GHz. The cooling cutoff $\nu_2 \sim 3 \times 10^4 \text{ GHz}$ and rarely goes below $5 \times 10^3 \text{ GHz}$ or exceeds 10^6 GHz . The distributions of α_{NIR} and α_{submm} are complementary to Figure 18. Table 2 lists medians and 1σ credible levels for each of these quantities.




Figure 17. Probability distributions of simultaneous NIR and X-ray measurements. Gray points show distributions from 1000 mock multi-wavelength light curves (Section 5), and lines show their 68% and 95% credible contours. Red points show real measurements from simultaneous observations. All values, X-ray and NIR, mock and real, are intrinsic with the real data corrected for extinction by the canonical factors given in Table 3. Upper panel shows $[S_K, S_X]$ with data from Fazio et al. (2018). The large orange data point shows the peak value of the NIR flare and its X-ray counterpart observed by those authors (their Table 3). The even more extreme event observed by Do et al. (2019) had no simultaneous X-ray observations. At a corrected $S_{2.2\mu m} = 60$ mJy, that observation would be just to the right of the plot area. The bottom panel shows $[S_M, S_X]$ with observed data from Boyce et al. (2019). To be consistent with the presentation of the X-ray data in the respective publications, the X-ray count rates are shown in 600 s bins for the upper panel and 300 s bins for the lower. Even in those bin widths, X-ray counts can be zero. In order to show a continuous posterior in log space, we added $\chi_s = 0.001$ cps to the count rate and a Gaussian random variable of $\sigma = 0.18$ to the values of the logarithm to dither both the mock and observed X-ray data.

Even two uncorrelated light curves will sometimes show peaks near the same time. The “false alarm probability” gives the probability of finding such spurious correlations at any time lag due to the white noise and the autocorrelation (the red noise) in each of the two bands. Figure 3 shows the 95% false alarm probability for the submm-NIR light curves. This was derived by calculating the DCC of the 230 GHz and 4.5 μm mock data in the same way as for the observed data but using pairs of 700 minute light curves that do not belong to each other. By modeling five data sets with the cadence of the observed data, we generated 200 DCC functions from our 1000 light curves. At each time lag, we then determined the 95% quantile.

Figure 21 shows the posterior of complementary cumulative distribution functions (CCDF) of K -band flux densities in




Figure 18. Posterior distributions calculated from mock light curves (Section 5). Points in the upper panel show NIR spectral index vs. flux density $S_{2.2\mu m}$, and those in the lower panel shows the submm spectral index vs. $S_{230 GHz}$. Quantities shown are intrinsic, and observations have to be corrected for extinction before comparison.




Figure 19. Posterior distribution of time lag from NIR to submm from mock data, where positive values mean NIR leads.

comparison to the empirical CCDF from NIRC2 and NACO data. For all fluxes, the model is 3σ consistent with the observed data, and up to about 20 mJy (dereddened) it is 2σ




Figure 20. Posterior distributions of flux densities and derived parameters for the synchrotron-SSC model (Section 5) calculated from the mock light curves. These posteriors represent both the variance propagated from the posterior as well as the variance from the variability. The large skewness of the NIR flux-density distribution biases the parameters toward small NIR fluxes. Panels show histograms and pairwise correlations of the 230 GHz flux density, the 2.2 μm and 4.5 μm flux densities, the X-ray flux density, the magnetic flux density, the source size in units of R_S , the density of high-energy electrons in the source volume, the frequency of maximum synchrotron flux density, the frequency of cooling cutoff of synchrotron spectrum, the intrinsic 4.5–2.2 μm spectral index, and the 230–345 GHz spectral index.

consistent. However, the model does not predict the brightest flux densities to be as frequent as we observe them.

Finally, Figure 22 shows a snapshot of our animation. The animation is the most concise way to illustrate the properties of our model. However, presenting model light curves this way is based on picking a particular parameter set from the posterior. It is the nature of the Bayesian approach that this posterior

includes particles that might not generate one or another aspect of the observed data because, after all, the observed data could be an unlikely realization of the underlying process. In order to allow the reader to become more familiar with the phenomenology of the model and the presented posterior, we are publishing the particle system (the “chain” of parameters), the observed data, and our Python implementation of the model




Figure 21. Cumulative distribution functions (CCDFs) of Sgr A* 2.12 μm flux densities (as observed, not corrected for extinction). The black upper line shows the CDF observed by the VLT and Keck, with the dashed section indicating the CDF needed to match the single largest flux-density ever observed (Do et al. 2019). The lower dashed blue line shows the median CDFs from the mock light curves. Shaded areas show 68%, 95%, and 99% credible intervals (Table 2).

and of the animation code. The details of the code and the data repository are described in the [Appendix](#).

6. Discussion

6.1. Fitting Results

The ABC algorithm results in constrained distributions of all parameters and describes the structure functions of nearly all data sets within the 1σ envelopes. The only exception is Spitzer/IRAC, where parts of the structure functions show slightly less variability than the model at timescales ~ 40 minutes. This deviation is not significant. The values of all parameters are reasonable and consistent with previous analyses (e.g., Eckart et al. 2012) and prior knowledge: magnetic flux density of $5 < B < 30$ G, source sizes $< 4 R_S$, the upper limit derived from 1.3 mm VLBI observations (Doeleman et al. 2008; Lu et al. 2018), and $n_e = 4 \times 10^7 \text{ cm}^{-3}$, comparable to the densities derived from the radio to submm luminosities (e.g., Bower et al. 2019). The effective collection area of each of the Chandra modes and the white noise levels in the NIR and submm are consistent with independently determined values. The quiescent count rate $\chi = 1.1_{-0.5}^{+0.8} \times 10^{-3}$ cps is identical with the background count rate of comparison apertures close to Sgr A* (Yuan & Wang 2016). However, Yuan & Wang found the sum of background count rate, truly quiescent level, and undetected variability $\chi_I \approx 4.8 \times 10^{-3}$ cps, 10% of which they attributed to undetected, weak flares. In contrast, in our analysis all X-ray photons other than the expected background rate can be attributed to the SSC process. At times of low activity, this appears quiescent.

6.2. Timing and Cross-correlation Properties

The initial idea of two correlated processes related by a low-pass filter is consistent with our final values for $(\gamma_{\text{slow}}, f_{b,\text{slow}})$ and $(\gamma_{\text{fast}}, f_{b,\text{fast}})$ as shown by Figure 16. In Sections 3.4 and 3.5, the NIR and X-ray contours for the PSD parameters were estimated entirely independently, based on the assumption of a log-normal distribution of flux densities in both cases. No relation between both bands was assumed, and no characterization of the covariance informed the fit. Section 4 re-analyzed the NIR and X-ray data with additional constraints from the submm and a simultaneous fit to the autocorrelation of each band. Even then, there was no direct characterization of the covariance (i.e., simultaneous data) between the NIR and the X-rays. The key difference was linking the probability distribution of flux densities and the PSDs of all bands by

the physical synchrotron-SSC model described in Section 4.4. In this second step, the slow and fast processes show clearly separated contours in the break frequency versus slope plane, as shown in Figure 16. This is not the result of the priors, which are equal for both PSDs. The break frequencies of both processes are systematically higher than in the pre-analysis model, which is given by the radiative model. In particular, the fast process is marginally consistent with a timescale of 50 minutes, which is the orbital timescale of the astrometric motion measured by VLTI/GRAVITY (Gravity Collaboration et al. 2018b).

The consequence of the low-pass filter is the decorrelation of the two processes toward higher frequencies. We can provide some qualitative reasoning for this behavior. The variability is caused by injection of nonthermal electrons whose energy distribution is truncated at some maximum energy E_{max} . This cuts off the synchrotron spectrum at a frequency close to the NIR. The injected electrons are subject to both expansion cooling and synchrotron cooling. The former affects the entire spectrum, while the latter (in the time frame of injection and expansion cooling) affects only the highest frequencies. The timescale of expansion cooling depends on the rather low expansion velocity, but the synchrotron cooling timescale at frequencies above the NIR is mere seconds to minutes. Therefore, rapid fluctuations of the injection process will be tracked by NIR emission, but the submm emission will track only the long-term average. It should be possible to derive a time-dependent model of the injection process from our semi-empirical variability model, but this is beyond the scope of our analysis.

The distance function used for our ABC implementation does not use any measure of correlation between the bands. Instead, the correlation is naturally built into the model. The exact phenomenology of the correlation, however, depends on the parameters that the ABC algorithm finds in the attempt to describe the structure functions of all bands. Figure 17 shows a scatter plot of the posterior NIR and X-ray flux-density pairs. The observations are in good agreement with the posterior, with the exception of a rare, very high K-band flux density at a low X-ray level (Fazio et al. 2018, their Figure 3). Estimating the probability of this event is difficult, given the different cadences of the mock and real light curves, but the event is a hint that the real source is more complicated than our purely cyclic model.

Neilsen et al. (2015) pointed out that the relation between NIR and X-ray fluxes is expected to be mildly nonlinear. Assuming power-law distributions for the flux-density distributions in the NIR and the X-ray, they derived the exponent of the power-law dependence of X-ray flux densities on NIR flux densities (their Equations (8)–(11)). In the case of a power-law index of ~ 2 for the X-ray distribution and ~ 4 for the NIR (Witzel et al. 2012; Neilsen et al. 2015), the dependence should be $S(E_{\text{keV}}) \propto S_{\text{NIR}}^3$. Our simple model predicts an exponent between 1.7 and 2.0 (Section 4.2). The higher nonlinearity in the relation of observed NIR flux densities to simultaneously measured X-ray flux densities can, however, be explained by the exponential cooling cutoff (Equation (10)).

In synchrotron-SSC models, the X-rays have the same spectral index as the optically thin part of the synchrotron spectrum. In their careful analysis of the X-ray spectral slope, Ponti et al. (2017) found a photon index $\Gamma = 2.27 \pm 0.12$,




Figure 22. Snapshot of the time evolution of the SED (generated from a mock data realization from posterior particle 1655) generated with our Python package (Appendix). Panels (a) and (b) show the observed SED at minute 732 of the simulated evolution, (a) in flux density and (b) in luminosity units. Orange lines show the synchrotron emission, blue the SSC emission, and red their sum. The dashed line marks $\nu = 230$ GHz, the short gray solid line the typical range of NIR variability when the source flux is above the detection limit, and the gray rectangle the typical range of peak 2–8 keV flux densities. Panel (c) shows spectral index vs. flux density for minutes 683–732 of the mock light curve, with the earlier points being more transparent. Green shows NIR, and orange shows submm. Panel (d) shows X-ray flux “as measured” (cps) vs. $S_{2.2 \mu m}$, again for minutes 683–732, with the earlier points being more transparent. Panels (e)–(i) show time series of various parameters as labeled. Left and right ordinate labels are color-coded to match the quantities plotted. The vertical line marks minute 732. Points to the right of that are shown as intrinsic values, while in panels (e)–(g), points to the left are shown “as measured,” i.e., with measurement noise added. (This is not applicable to panels (h) and (i).) Red points represent a “measurement” at minute 732. The X-ray count rate is averaged over 600 s bins. An animation of this figure is available online. The animation is 5:33 minutes long and shows 2250 minutes of source evolution. Panels (a) and (b) show the up and down of the synchrotron and SSC components of the spectrum as well as changes in the self-absorption turnover frequency and the fluctuation of the cutoff frequency with γ_{\max} . Panels (c) and (d) show the evolution of 50 minute long traces of the spectral index and X-ray flux density vs. NIR flux density. Panels (e) to (i) show graphs of time series of the various quantities that run with time from right to left.

(An animation of this figure is available.)




Figure 23. Variable SED in comparison to overall SED of Sgr A*. The dashed lines show snapshots of the SED from our model for three different (dereddened) NIR flux-density levels. Points show SED values taken from Zylka et al. (1992), Telesco et al. (1996), Falcke et al. (1998), Cotera et al. (1999), An et al. (2005), Dodds-Eden et al. (2009, 2011), Schödel et al. (2011), Bower et al. (2015), Liu et al. (2016), Stone et al. (2016), von Fellenberg et al. (2018), Witzel et al. (2018), Bower et al. (2019), and Gravity Collaboration et al. (2020b). The (dereddened) NIR points represent different attempts to quantify the mode (peak) of the NIR flux-density distribution (compare discussion in Witzel et al. 2018). At the Herschel bands (100, 160, and 250 μm), two numbers are presented: the amplitude of the detected variable signal (blue lower limit, gray points) that serve as a lower limit of the steady flux density, and for 100 and 160 μm , upper limits of the steady flux density derived from the assumption of 15% variability (von Fellenberg et al. 2018).

which corresponds to a spectral index of $\alpha_X = -1.27 \pm 0.12$, about 1σ consistent with our $\alpha = -0.98 \pm 0.1$. Haggard et al. (2019) analyzed the two brightest X-ray flares detected so far and found $\Gamma = 2.06 \pm 0.14$ and $\Gamma = 2.03 \pm 0.27$, respectively, again consistent with our results.

Boyce et al. (2019) offered another way to investigate the NIR to X-ray correlation. Their cross-correlation analysis of the simultaneous Chandra and Spitzer data found a tendency for the X-ray peaks to precede the NIR. However, Boyce et al. could not claim a statistically significant delay, given the large uncertainties. Boyce et al. also presented an overview of earlier results, many of which showed the X-ray to follow the NIR but again without high significance. Our 700 minute mock light curves show no significant delay between the X-ray and NIR bands.

Our model predicts correlation between the NIR and submm light curves as well. We quantify this in terms of the cross-correlation between S_{thin} and the corresponding 230 GHz mock light curves as shown in Figure 20(a). Observed delays range from almost synchronous variability (as observed by Fazio et al. 2018 and in this paper in the case of the 345 GHz SMA data from 2014 June 17) to delays up to 90 minutes (as found by Eckart et al. 2008b). The model distribution peaks at ~ 22 minutes with an FWHM of ~ 10 minutes, which corresponds nicely to the result we obtained from our ~ 48 hr of synchronous NIR and 230 GHz data. The cross-correlation in Figure 3 shows a pronounced peak at ~ 27 minutes. The observed light curves have a maximum correlation $R = 0.21$. However, the observed NIR flux densities, if our model is correct, are not identical with the slow process $S_{\text{thin}}(t)$ but include the faster varying effects of the cooling cutoff. In order to approximate the quantity $S_{\text{thin}}(t)$, we also present the cross-correlation of the 230 GHz data with the logarithm of the low-pass-filtered NIR flux densities. This results in a significant increase of the correlation to $R = 0.43$. The 95% false alarm

probability levels calculated from our model and the posterior in Section 2 indicate that the data are consistent with our model assumptions.¹⁸

6.3. SED and Adiabatic Expansion

Figure 23 compares the SED of the compact component modeled here with the so-called steady SED of (selected) literature values. A minimum requirement for our model is that it cannot violate the steady SED. At frequencies above submm, the only detections of Sgr A* are far-infrared measurements with Herschel (Stone et al. 2016; von Fellenberg et al. 2018) and NIR and X-ray measurements as presented here. Our model is consistent with the variability amplitudes and the spectral index in the FIR while not violating any upper limits in the radio or submm part of the SED. In the submm, our model component significantly contributes to the overall variability but cannot explain the entire flux density. At lower radio frequencies (< 300 GHz), the variable contribution decreases quickly, and the SED is dominated by one or more other source components, presumably the quasi-steady emission from the accretion flow. This implies that the compact source region we have modeled is immersed in submm photons from the neighboring plasma. We have modeled this scenario with a thermal and nonthermal synchrotron model with SSC and IC scattering of the ambient submm photons from the thermal component by the nonthermal electrons of the compact component. For this model, we did not use the analytical approximations described in this paper but instead used numerical integration to properly calculate the flanks of the

¹⁸ However, we have not tested the significance of the correlation independently (with, e.g., a Granger causality test). The cross-correlation of Figure 3 cannot be used as independent evidence for the existence of this type of correlation. All we can conclude here is that the Spitzer-SMA synchronous data set is consistent with a delay that is predicted by our model.

SSC and IC spectra. We assumed that the nonthermal electrons are exposed to a constant submm photon density as if located at the center of a sphere of a larger volume of thermal electrons. The thermal component of the spectrum was fitted to the cm to mm observations. For a rather bright NIR state, the resulting IC component is six to seven orders of magnitude dimmer than the SSC component. If the two electron populations are separated in space by some distance, the photon density and consequently the IC emission would be even lower. Therefore, the IC component from the ambient photons can be safely neglected here.

Several papers have discussed expanding plasmon flare evolution models for Sgr A*. The earliest, by Yusef-Zadeh et al. (2006b), described the temporal evolution of flares in the cm wavelength regime. Eckart et al. (2006, 2008b, 2012) explained a possible correlation and delay between NIR and submm flares. Marrone et al. (2008) presented an example of simultaneous IR and submm data with a delay (IR leading, submm following) of 20 ± 5 minutes, consistent with the data presented here and with our model, and they also discussed adiabatic expansion. All these expanding plasmon models result in typical expansion velocities of $\sim 0.01c$, with c being the speed of light. Peak expansion velocities derived from our time-dependent model reach $\sim 0.01c$ and are consistent with the earlier estimates.

In contrast to Yusef-Zadeh et al. (2006b, 2009), Falcke et al. (2009) interpreted 20–40 minute delays in the 20 to 40 GHz regime in the context of the frequency dependence of VLBI sizes and saw evidence for a relativistic outflow. Brinkerink et al. (2015) even followed the realtime progression of variability maxima from 100 to 19 GHz and similarly derived relativistic outflow velocities of up to $\sim 0.77c$, i.e., a jet.

Whether the similar delays found at radio frequencies and in the submm to NIR regime indicate a direct relation between the variability of those regimes or are merely coincidental cannot be decided here. As Figure 23 demonstrates, for frequencies < 100 GHz and $\alpha_{\text{Radio}} = +2.5$, the variable component seldom contributes significant flux. Our posterior of the self-absorption peak position ν_m (Figure 20) predicts values as low as 90 GHz with measurable contributions down to 40–50 GHz. However, predictions for the radio part crucially depend on the actual α_{Radio} of the optically thick branch, which for electron density profiles other than constant with radius can differ from 2.5 and is very uncertain. Also, we cannot determine whether the compact component responsible for the fast, high-frequency variability is located in a jet or in the accretion disk.

The model predicts the self-absorption turnover to vary between 90 GHz and 1 THz, yielding a wide range of spectral indices at these wavelengths. On the other hand, in the NIR the spectral index is typically steeper than the optically thin spectral index $\alpha = (1 - \gamma_e)/2$ due to the cooling cutoff. Both regimes, one used by mm VLBI, the other by VLT/ GRAVITY, are limited in the sense that they are subject to rapid changes in the state of the radiative transfer: the submm by its regular changes between the optically thick and thin branch of the spectrum, which potentially cause rapid changes in the source structure; and the NIR by fast synchrotron cooling, which makes it hard to follow the compact component of Sgr A* in its entire evolution.

The median brightness temperature at the self-absorption turnover $T_b = 1.8 \times 10^{10}$ K. It varies between 80% and 90% of the equipartition brightness temperature T_{eq} . Assuming a pure

electron plasma and a Doppler factor of 1, this value suggests that the source is energetically dominated by the magnetic flux.

6.4. Inconsistencies

The model spectral index in the NIR has an upper limit $\alpha_{\text{NIR}} \lesssim -0.8$, with many bright flares showing spectral indices as steep as $\alpha_{\text{NIR}} \approx -1.2$ (Figure 21). The upper limit of α_{NIR} is defined by the spectral index $\alpha = (1 - \gamma_e)/2$ of the optically thin spectrum without cooling cutoff. The NIR spectral index reaches the optically thin limit when the cooling cutoff is at high energies. In our model fit, γ_e is determined by the variance of the submm variability relative to the variance of the NIR variability. $\alpha \approx -1.0$ is steeper than the canonical NIR spectral index for bright flares of $\alpha_{\text{NIR}} \approx -0.65$ (Eisenhauer et al. 2005; Ghez et al. 2005; Gillessen et al. 2006; Krabbe et al. 2006; Hornstein et al. 2007; Bremer et al. 2011; Witzel et al. 2014). However, spectral index analyses in the NIR have two main sources of systematic uncertainty: the extinction correction and background flux levels. While the extinction correction is precise enough and cannot explain this discrepancy, unaccounted background flux makes a big difference. However, Do et al. (2019) reported $\alpha_{\text{NIR}} \approx -0.4$ (with canonical extinction correction) at such high flux-density levels that background contamination does not play a role. Such a flat spectral index could be the result of the SSC scattering contributing to the NIR. The lower limit of the power-law section of the SSC spectrum is given in Equation (12). As seen in Figure 9, for source parameters discussed here, the SSC spectrum sets in around NIR frequencies and might become dominant once in awhile.

Another, potentially related, problem is the submm variability variance: at 230 GHz, the variance at longest timescales cannot be reproduced by the model. This is the reason we did not include the last bin of the 230 GHz structure function in our distance function. Furthermore, Figure 1 shows significant offsets in average flux density between the individual SMA and ALMA epochs. If real, these offsets would point to a long-term variability component that is not accounted for in our model. Several explanations are possible:

1. We modeled the synchrotron source as a sphere homogeneously filled with electrons, resulting in an optically thick spectral index $\alpha_{\text{thick}} = +2.5$. In the case of a different electron density profile, this slope can be significantly different, with consequences for the position of the self-absorption turnover and the relative flux densities between the 230 GHz, the 345 GHz, and the NIR.
2. A separate mechanism, perhaps a second electron population, is required to produce the nonvarying emission, which accounts for most of the submm and longer-wavelength flux most of the time. This second (or rather primary) component is not needed to account for most of the observed submm variability, but it might contribute, especially on long timescales. This would reduce the variability associated with the compact source, resulting in a flatter optically thin spectral index.
3. Instead of one zone undergoing cycles, there might be distinct zones coming and going with different initial conditions, explaining the distinct levels of the submm epochs.

Finally, our model gives the brightest observed NIR flare so far (Do et al. 2019) only $\sim 10^{-3}$ probability. In other words, the brightest observed NIR states of Sgr A* are somewhat under-represented in our mock light curves. Dodds-Eden et al. (2011) and Gravity Collaboration et al. (2020b) interpreted the NIR flux-density distribution as the result of two states with variability from two distinct physical processes. Whether this is the case or instead very bright events are only extreme cases of the process described here cannot be determined. In order to generate extreme events more regularly, it would suffice to replace log-normality of the underlying fast and slow processes with a suited distribution. However, proper inclusion of extreme values in the statistics is difficult because the probability of rare events is very uncertain.

7. Summary

This paper has presented a comprehensive analysis of submm, NIR, and X-ray light curves of Sgr A* including two new epochs of ALMA data and four new epochs of SMA data. These include about two days of simultaneously observed NIR and submm light curves, which show definite correlation between submm and NIR variability and are consistent with a median delay of ~ 20 minutes of the submm with respect to the NIR.

This paper is also the first analysis of the X-ray PSD. The X-ray and NIR variations are correlated, but with correlation decreasing at short variability timescales.

A simple but physically consistent model explains most features of Sgr A*'s variability. The model was constructed to match the variability amplitudes and timing properties (structure functions) of the NIR and X-ray emission considered separately. Given those, the model predicts the submm variability properties, the time lag between NIR and submm variations, the submm and NIR spectral indices, and the observed correlation between NIR and X-ray variability. The model consists of a single zone in an external magnetic field with random injections of high-energy electrons and cyclic expansion and contraction. The radio emission is optically thick synchrotron, the NIR is optically thin synchrotron, and X-rays arise from synchrotron self-Compton emission. The variability comes from varying density of high-energy electrons and source size, which affect the amplitude of synchrotron emission, the optical depth, the synchrotron cutoff frequency, and the self-Compton amplitude. The minimum required magnetic field is ~ 8.5 G, the maximum source size is $\sim 2.7 R_S$, and the high-energy electrons have a power-law index ~ 2.95 . The variability processes have timescales of roughly 82 and 135 minutes and are ruled by the tradeoff between electron injection and expansion cooling in the submm and X-rays and by the tradeoff between injection and synchrotron cooling in the NIR. General relativity plays no role in the model.

The predicted changes in source parameters are moderate, typically about a factor of two in source size and magnetic flux density and an order of magnitude in electron density. The submm radiation of this compact component of Sgr A* changes from optically thick to thin and back at the variability timescales. This likely creates complex changes in the intrinsic source structure during VLBI observations.

Postulates of the model include the PSD of the underlying trigger processes and that rises and falls are symmetric in time. The model allows for but does not require different correlation times at different frequencies and time lags between

frequencies. Where the model may fall short is that the most extreme observations are improbable (but not impossible), and the observed NIR spectral index at high flux densities is flatter than the model predicts. That last issue could probably be remedied by allowing SSC emission at NIR frequencies.

All in all, the model's shortcomings are related to rare observations, such as the brightest flux-density states, and therefore are documented insufficiently in our sample to warrant a more complex model. For the vast majority of observations, the model presented here is a sufficient representation. At the least, it provides a baseline to help identify extraordinary variability, e.g., states dominated by effects of general relativity in the future. Relativistic effects alone cannot be the origin of the variability because they cannot explain the rich phenomenology of correlations between the wavelengths. Our model in its current form does not require any relativistic boosting or other effects to describe the variable flux densities. Whether and how this is commensurable with the findings of apparent circular motion close to the last stable orbit by Gravity Collaboration et al. (2018b) has to be investigated.

Physical conclusions from this work include:

1. There is a compact component with size $2-3R_S$ that dominates the high-frequency regime.
2. There is more evidence now for SSC generating the X-rays (and maybe contributing to the NIR).
3. The NIR is linked to the submm variability, and expansion is a strong candidate process to explain the correlation.

Some things the model does not explain include:

1. Long-timescale variability in the submm.
2. Perhaps the most extreme NIR flux densities.

This paper has presented a quantitative assessment of what a simple model can achieve. It can serve as a starting point for further work to test and improve the model. Theoretical work should investigate acceleration mechanisms for electrons and whether suitable populations of high-energy electrons can be generated and on what timescales.

The model presented here was enabled and its parameters determined by the vast library of light-curve monitoring data accumulated over the past two decades. Simultaneous light curves at different wavelengths were especially valuable for testing (though not deriving) the model. Further work should include better simultaneous measurements of the X-ray and NIR spectral indices, which should be identical if our model is correct. Monitoring Sgr A*'s variability simultaneously from ~ 2 to $\sim 30 \mu\text{m}$ with the James Webb Space Telescope should test whether the model's predicted NIR spectral indices are accurate, clarify the relevance of synchrotron cooling for this part of the spectrum, and show whether the self-Compton emission can contribute to the NIR.

We thank the anonymous referee, the statistical editor, and the data editor for their valuable and constructive comments. The authors are grateful to Eduardo Ros for valuable comments on the manuscript. We thank Ziqian Hua, Axel Weiss, Mikhail Lisakov, Nicholas MacDonald, Stefan Gillessen, Sebastiano von Fellenberg, Yigit Dallilar, Gabriele Ponti, Arno Witzel, and Silke Britzen for fruitful discussions. We thank Helge Rottmann and the MPIfR Correlator team for their help and

support in running our code on the VLBI computer cluster. This work is based on observations made with the Spitzer Space Telescope, which is operated by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. Support for this work was provided by NASA through an award issued by JPL/Caltech. We thank the staff of the Spitzer Science Center for their help in planning and executing these demanding observations. This publication is based on data of the Submillimeter Array. The SMA is a joint project between the Smithsonian Astrophysical Observatory and the Academia Sinica Institute of Astronomy and Astrophysics, and is funded by the Smithsonian Institution and the Academia Sinica. This paper makes use of the following ALMA data: ADS/JAO.ALMA#2011.0.00887.S and ADS/JAO.ALMA#2017.1.00503.S. ALMA is a partnership of ESO (representing its member states), NSF (USA), and NINS (Japan), together with NRC (Canada), MOST and ASIAA (Taiwan), and KASI (Republic of Korea), in cooperation with the Republic of Chile. The Joint ALMA Observatory is operated by ESO, AUI/NRAO, and NAOJ. Support for this work was provided by the National Aeronautics and Space Administration through Chandra Award Numbers GO6-17136B and GO7-18135B, issued by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for and on behalf of the National Aeronautics Space Administration under contract NAS8-03060. S.P.W., G. G.F., J.L.H., and H.A.S. acknowledge grant 80NSSC18K0416 of the National Aeronautics Space Administration. G.M., E.E. B., T.D., A.G., and M.R.M. acknowledge grant AST-1909554 of the National Science Foundation.

Facilities: Spitzer(IRAC), ALMA, SMA, APEX, Chandra (ACIS), VLT:Yepun(NACO), Keck:II(NIRC2).

Software: SED Animation v.1.0 (Witzel 2021), FFmpeg¹⁹, Jupyter Notebook (Kluyver et al. 2016), Python 3 (Van Rossum & Drake 2009) with the packages math, decimal, mpmath, and time, Numerical Python (numpy, Harris et al. 2020; Van Der Walt et al. 2011), Matplotlib (Hunter 2007), Scientific Python (scipy, Virtanen et al. 2020), C++, MPI, and WebPlotDigitizer (Rohatgi 2020).

Appendix Supplemental Code Package for Generating SED Animations

As a supplement to this article, we have created a repository containing the Python code with which we generated Figure 22 and the linked animation (SED Animation v1.0, Max Planck Digital Library, doi:10.17617/1.kctx3s25, developed on <https://gitlab.mpcdf.mpg.de>). The repository can be found at <https://doi.org/10.17617/1.kctx3s25>.

The repository contains:

1. a Python library of classes and functions for
 - (a) generating synchrotron and SSC spectra from model and empirical data,
 - (b) generating time series and their auto- and cross-correlations,
 - (c) making simple calculations and plotting auxiliary data and calculation results,
 - (d) handling the posterior particle system from our final ABC run, and
 - (e) generating movies of the evolution of the data over time;
2. a Jupyter notebook with simple examples showing how to generate an SED and an animation from the posterior; and
3. the posterior particle system of 5000 parameter combinations with appropriate weights.

This package is tailored to the case of Sgr A* and the model described above. However, we hope it might be of use for different applications, and we publish the code under a 3-Clause BSD License that will permit free use. The code uses the FFmpeg library, Jupyter Notebook (Kluyver et al. 2016), Python 3 (Van Rossum & Drake 2009) with the packages math, decimal, mpmath, and time, Numerical Python (numpy, Harris et al. 2020; Van Der Walt et al. 2011), Matplotlib (Hunter 2007), and Scientific Python (scipy, Virtanen et al. 2020).

ORCID iDs

G. Witzel  <https://orcid.org/0000-0003-2618-797X>
 G. Martinez  <https://orcid.org/0000-0002-7476-2521>
 S. P. Willner  <https://orcid.org/0000-0002-9895-5758>
 H. Boyce  <https://orcid.org/0000-0002-6530-5783>
 T. Do  <https://orcid.org/0000-0001-9554-6062>
 A. Eckart  <https://orcid.org/0000-0001-6049-3132>
 G. G. Fazio  <https://orcid.org/0000-0002-0670-0708>
 A. Ghez  <https://orcid.org/0000-0003-3230-5055>
 M. A. Gurwell  <https://orcid.org/0000-0003-0685-3621>
 D. Haggard  <https://orcid.org/0000-0001-6803-2138>
 R. Herrero-Illana  <https://orcid.org/0000-0002-7758-8717>
 J. L. Hora  <https://orcid.org/0000-0002-5599-4650>
 Z. Li (李志远)  <https://orcid.org/0000-0003-0355-6437>
 J. Liu (刘俊)  <https://orcid.org/0000-0002-7615-7499>
 N. Marchili  <https://orcid.org/0000-0002-5523-7588>
 Mark R. Morris  <https://orcid.org/0000-0002-6753-2066>
 M. Subroweit  <https://orcid.org/0000-0001-6165-8525>
 J. A. Zensus  <https://orcid.org/0000-0001-7470-3321>

References

- An, T., Goss, W. M., Zhao, J.-H., et al. 2005, *ApJL*, 634, L49
 Baganoff, F. K., Bautz, M. W., Brandt, W. N., et al. 2001, *Natur*, 413, 45
 Baganoff, F. K., Maeda, Y., Morris, M., et al. 2003, *ApJ*, 591, 891
 Band, D. L., & Grindlay, J. E. 1985, *ApJ*, 298, 128
 Boccardi, B., Krichbaum, T. P., Ros, E., & Zensus, J. A. 2017, *A&ARv*, 25, 4
 Bower, G. C., Dexter, J., Asada, K., et al. 2019, *ApJL*, 881, L2
 Bower, G. C., Markoff, S., Dexter, J., et al. 2015, *ApJ*, 802, 69
 Boyce, H., Haggard, D., Witzel, G., et al. 2019, *ApJ*, 871, 161
 Bremer, M., Witzel, G., Eckart, A., et al. 2011, *A&A*, 532, A26
 Brinkerink, C. D., Falcke, H., Law, C. J., et al. 2015, *A&A*, 576, A41
 Cotera, A., Morris, M., Ghez, A. M., et al. 1999, in ASP Conf. Ser. 186, The Central Parsec of the Galaxy, ed. H. Falcke et al. (San Francisco, CA: ASP), 240
 Coti Zelati, F., Rea, N., Turolla, R., et al. 2017, *MNRAS*, 471, 1819
 Dexter, J., Kelly, B., Bower, G. C., et al. 2014, *MNRAS*, 442, 2797
 Do, T., Witzel, G., Gautam, A. K., et al. 2019, *ApJL*, 882, L27
 Dodds-Eden, K., Gillissen, S., Fritz, T. K., et al. 2011, *ApJ*, 728, 37
 Dodds-Eden, K., Porquet, D., Trap, G., et al. 2009, *ApJ*, 698, 676
 Dodds-Eden, K., Sharma, P., Quataert, E., et al. 2010, *ApJ*, 725, 450
 Doleman, S. S., Weintroub, J., Rogers, A. E. E., et al. 2008, *Natur*, 455, 78
 Eckart, A., Baganoff, F. K., Morris, M., et al. 2004, *A&A*, 427, 1
 Eckart, A., Baganoff, F. K., Schödel, R., et al. 2006, *A&A*, 450, 535
 Eckart, A., Baganoff, F. K., Zamaninasab, M., et al. 2008a, *A&A*, 479, 625
 Eckart, A., García-Marín, M., Vogel, S. N., et al. 2012, *A&A*, 537, A52
 Eckart, A., Schödel, R., García-Marín, M., et al. 2008b, *A&A*, 492, 337
 Edelson, R. A., & Krolik, J. H. 1988, *ApJ*, 333, 646

¹⁹ <https://ffmpeg.org/>

- Eisenhauer, F., Genzel, R., Alexander, T., et al. 2005, *ApJ*, **628**, 246
- Falcke, H., Goss, W. M., Matsuo, H., et al. 1998, *ApJ*, **499**, 731
- Falcke, H., Markoff, S., & Bower, G. C. 2009, *A&A*, **496**, 77
- Fazio, G. G., Hora, J. L., Witzel, G., et al. 2018, *ApJ*, **864**, 58
- García-Marín, M., Eckart, A., Weiss, A., et al. 2011, *ApJ*, **738**, 158
- Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, *RvMP*, **82**, 3121
- Genzel, R., Schödel, R., Ott, T., et al. 2003, *Natur*, **425**, 934
- Ghez, A. M., Hornstein, S. D., Lu, J. R., et al. 2005, *ApJ*, **635**, 1087
- Ghez, A. M., Wright, S. A., Matthews, K., et al. 2004, *ApJL*, **601**, L159
- Gillessen, S., Eisenhauer, F., Quataert, E., et al. 2006, *ApJL*, **640**, L163
- Ginzburg, V. L. 1951, *DoSSR*, **76**, 377
- Ginzburg, V. L., & Syrovatskii, S. I. 1965, *ARA&A*, **3**, 297
- Gould, R. J. 1979, *A&A*, **76**, 306
- 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, Abuter, R., Amorim, A., et al. 2020b, *A&A*, **638**, A2
- Gravity Collaboration, Bauböck, M., Dexter, J., et al. 2020c, *A&A*, **635**, A143
- Haggard, D., Nynka, M., Mon, B., et al. 2019, *ApJ*, **886**, 96
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Haubois, X., Dodds-Eden, K., Weiss, A., et al. 2012, *A&A*, **540**, A41
- Hora, J. L., Witzel, G., Ashby, M. L. N., et al. 2014, *ApJ*, **793**, 120
- Hornstein, S. D., Matthews, K., Ghez, A. M., et al. 2007, *ApJ*, **667**, 900
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Iwata, Y., Oka, T., Tsuboi, M., Miyoshi, M., & Takekawa, S. 2020, *ApJL*, **892**, L30
- Kardashev, N. S. 1962, *AZh*, **39**, 393
- Kluyver, T., Ragan-Kelley, B., Pérez, F., et al. 2016, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, ed. F. Loizides & B. Schmidt (Amsterdam: IOS Press), 87
- Krabbe, A., Iserlohe, C., Larkin, J. E., et al. 2006, *ApJL*, **642**, L145
- Liu, H. B., Wright, M. C. H., Zhao, J.-H., et al. 2016, *A&A*, **593**, A107
- Lu, R.-S., Krichbaum, T. P., Roy, A. L., et al. 2018, *ApJ*, **859**, 60
- Marrone, D. P., Baganoff, F. K., Morris, M. R., et al. 2008, *ApJ*, **682**, 373
- Marscher, A. P. 1983, *ApJ*, **264**, 296
- McMullin, J. P., Waters, B., Schiebel, D., Young, W., & Golap, K. 2007, in *ASP Conf. Ser. 376, Astronomical Data Analysis Software and Systems XVI*, ed. R. A. Shaw, F. Hill, & D. J. Bell (San Francisco, CA: ASP), 127
- Meyer, L., Do, T., Ghez, A., et al. 2009, *ApJL*, **694**, L87
- Meyer, L., Witzel, G., Longstaff, F. A., & Ghez, A. M. 2014, *ApJ*, **791**, 24
- Moffet, A. T. 1975, in *Strong Nonthermal Radio Emission from Galaxies*, ed. A. Sandage, M. Sandage, & J. Kristian (Chicago, IL: Univ. Chicago Press), 211
- Morris, M. R., Meyer, L., & Ghez, A. M. 2012, *RAA*, **12**, 995
- Mossoux, E., Grosso, N., Bushouse, H., et al. 2016, *A&A*, **589**, A116
- Neilsen, J., Markoff, S., Nowak, M. A., et al. 2015, *ApJ*, **799**, 199
- Ponti, G., George, E., Scaringi, S., et al. 2017, *MNRAS*, **468**, 2447
- Rauch, C., Ros, E., Krichbaum, T. P., et al. 2016, *A&A*, **587**, A37
- Robertson, D. R. S., Gallo, L. C., Zoghbi, A., & Fabian, A. C. 2015, *MNRAS*, **453**, 3455
- Rohatgi, A. 2020, 4, 4, Webplotdigitizer: v4.4., <https://automeris.io/WebPlotDigitizer>
- Schödel, R., Morris, M. R., Muzic, K., et al. 2011, *A&A*, **532**, A83
- Shklovskii, I. S. 1952, *AZh*, **29**, 418
- Shklovskii, I. S. 1953, *DoSSR*, **90**, 983
- Shklovskii, I. S. 1960, *AZh*, **37**, 256
- Stone, J. M., Marrone, D. P., Dowell, C. D., et al. 2016, *ApJ*, **825**, 32
- Subroweit, M., García-Marín, M., Eckart, A., et al. 2017, *A&A*, **601**, A80
- Subroweit, M., Mossoux, E., & Eckart, A. 2020, *ApJ*, **898**, 138
- Sunaina, Butola, M., & Khare, K. 2018, *EJPh*, **39**, 065806
- Telesco, C. M., Davidson, J. A., & Werner, M. W. 1996, *ApJ*, **456**, 541
- Tilanus, R. P. J., Krichbaum, T. P., Zensus, J. A., et al. 2014, arXiv:1406.4650
- Timmer, J., & Koenig, M. 1995, *A&A*, **300**, 707
- Trap, G., Goldwurm, A., Dodds-Eden, K., et al. 2011, *A&A*, **528**, A140
- van der Laan, H. 1966, *Natur*, **211**, 1131
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Van Rossum, G., & Drake, F. L. 2009, *Python 3 Reference Manual* (Scotts Valley, CA: CreateSpace)
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, **17**, 261
- von Fellenberg, S. D., Gillessen, S., Graciá-Carpio, J., et al. 2018, *ApJ*, **862**, 129
- Witzel, G. 2021, SED Animation: Generation of Animated Synchrotron/SSC SEDs, v1.0, Max Planck Digital Library, doi:10.17617/1.kctx3s25
- Witzel, G., Eckart, A., Bremer, M., et al. 2012, *ApJS*, **203**, 18
- Witzel, G., Martínez, G., Hora, J., et al. 2018, *ApJ*, **863**, 15
- Witzel, G., Morris, M., Ghez, A., et al. 2014, in *IAU Symp. 303, The Galactic Center: Feeding and Feedback in a Normal Galactic Nucleus*, ed. L. O. Sjouwerman, C. C. Lang, & J. Ott (Cambridge: Cambridge Univ. Press), 274
- Xu, Y.-D., Narayan, R., Quataert, E., Yuan, F., & Baganoff, F. K. 2006, *ApJ*, **640**, 319
- Yuan, Q., & Wang, Q. D. 2016, *MNRAS*, **456**, 1438i
- Yusef-Zadeh, F., Bushouse, H., Dowell, C. D., et al. 2006a, *ApJ*, **644**, 198
- Yusef-Zadeh, F., Bushouse, H., Wardle, M., et al. 2009, *ApJ*, **706**, 348
- Yusef-Zadeh, F., Roberts, D., Wardle, M., Heinke, C. O., & Bower, G. C. 2006b, *ApJ*, **650**, 189
- Yusef-Zadeh, F., Wardle, M., Dodds-Eden, K., et al. 2012, *AJ*, **144**, 1
- Yusef-Zadeh, F., Wardle, M., Heinke, C., et al. 2008, *ApJ*, **682**, 361
- Zhu, Z., Li, Z., & Morris, M. R. 2018, *ApJS*, **235**, 26
- Zhu, Z., Li, Z., Morris, M. R., Zhang, S., & Liu, S. 2019, *ApJ*, **875**, 44
- Zylka, R., Mezger, P. G., & Lesch, H. 1992, *A&A*, **261**, 119