PARTICLE ACCELERATION AND THE ORIGIN OF X-RAY FLARES IN GRMHD SIMULATIONS OF SGR A*

DAVID BALL, FERYAL ÖZEL, DIMITRIOS PSALTIS, AND CHI-KWAN CHAN

Steward Observatory and Department of Astronomy, University of Arizona, USA

Received 2016 February 18; revised 2016 May 8; accepted 2016 May 23; published 2016 July 25

ABSTRACT

Significant X-ray variability and flaring has been observed from Sgr A* but is poorly understood from a theoretical standpoint. We perform general relativistic magnetohydrodynamic simulations that take into account a population of non-thermal electrons with energy distributions and injection rates that are motivated by PIC simulations of magnetic reconnection. We explore the effects of including these non-thermal electrons on the predicted broadband variability of Sgr A* and find that X-ray variability is a generic result of localizing non-thermal electrons to highly magnetized regions, where particles are likely to be accelerated via magnetic reconnection. The proximity of these high-field regions to the event horizon forms a natural connection between IR and X-ray variability and accounts for the rapid timescales associated with the X-ray flares. The qualitative nature of this variability is consistent with observations, producing X-ray flares that are always coincident with IR flares, but not vice versa, i.e., there are a number of IR flares without X-ray counterparts.

Key words: acceleration of particles – accretion, accretion disks – black hole physics – magnetic reconnection – magnetohydrodynamics (MHD) – radiative transfer

1. INTRODUCTION

Observational evidence and theoretical arguments strongly suggest that accretion flows around low-luminosity supermassive black holes are radiatively inefficient (see a recent review by Yuan & Narayan 2014). Current models of radiatively inefficient accretion flows employ simplified thermodynamic prescriptions of the rarefied plasma in the disk. This is due to a number of complicating factors that make it difficult to fully describe the physics of the system.

Low-luminosity accretion flows are composed of two temperature plasmas, where ions are significantly hotter than electrons. This occurs for several reasons. First, the very low-density plasma is nearly collisionless, such that thermalization via Coulomb collisions is inefficient. Second, the electrons are able to radiate away their energy much faster than the ions. Third, not only do cooling rates differ between the electrons and ions, but the mechanisms responsible for heating the fluid often favor one component of the plasma over the other, with the magnitude of this heating asymmetry between ions and electrons depending on the plasma properties (Quataert & Gruzinov 1999). Processes such as the dissipation of turbulent energy, magnetic reconnection, and shocks all play a role in mediating the temperature of the fluid, but the details of these mechanisms and their combined impact on the thermodynamics of low-luminosity accretion flows is not fully understood.

Simulations of low-luminosity accretion flows generally assume a Maxwellian distribution of electrons at a prescribed temperature (e.g., Dexter et al. 2012; Drappeau et al. 2013; Mościbrodzka et al. 2014; Chan et al. 2015b). More recently, Ressler et al. (2015) developed a method for independently evolving the electron temperature distribution, taking into account spatially varying heating as well as anisotropic electron conduction, but still assuming a thermal particle distribution. However, subgrid modeling of heating and acceleration processes indicate that some fraction of particles are likely to be accelerated into a non-thermal distribution.

Astrophysical particle acceleration is a heavily studied field with broad implications for many astrophysical systems (for a recent review, see Lazarian et al. 2012 and references therein).

There have recently been significant improvements in our understanding of the numerous heating and acceleration mechanisms that are relevant to accretion flows from a microphysical standpoint. By utilizing particle-in-cell (PIC) simulations, Sironi & Spitkovsky (2014) showed that relativistic reconnection generally accelerates the particles in a plasma into power-law like distributions. Guo et al. (2014) used PIC simulations to determine the effect of low Mach number shocks on acceleration and showed that it also produces a non-thermal distribution. These modeling efforts at small scales have yielded new insight into the fundamental properties of heating and acceleration mechanisms, but their effects have not yet been incorporated, as sub-grid models, into the larger scale general relativistic magnetohydrodynamic (GRMHD) simulations of low-luminosity accretion flows. By coupling the results of the microphysical models to GRMHD simulations, it will become possible to more robustly interpret some observed phenomena that have thus far eluded a physical explanation.

The role of non-thermal electron energy distribution has been addressed in the context of stationary hydrodynamic models by Mahadevan (1998), Özel et al. (2000) and, Yuan et al. (2003), who showed that even a relatively small number of power-law electrons can significantly impact the spectra predicted from a model, generating X-ray power-law tails as well as boosting the low frequency radio flux. These studies used analytical steady-state solutions to calculate spectra, which do not capture short timescale GRMHD effects that play a large role in determining the variability properties of the system. Dodds-Eden et al. (2010) used two-dimensional time-dependent MHD models and injected non-thermal electrons into regions of rapidly changing magnetic fields in an effort to explain the rapid flares seen from Sgr A*. These models, however, do not account for the three-dimensional character of the flow, as well as for general relativistic effects such as strong lensing and Doppler boosting that affect rapid variability. Chan et al. (2015a) showed these effects to be important in understanding the broadband variability properties of the supermassive black hole at the center of the Milky Way,

Sagittarius A* (Sgr A*) and especially for accounting for mm and IR flares that originate close to the event horizon.

Observations of Sgr A* reveal significant multiwavelength variability (see Eckart et al. 2004; Marrone et al. 2008; Neilsen et al. 2013; Witzel et al. 2014; Li et al. 2015b; Ponti et al. 2015; Yuan & Wang 2016) from the mm and IR to the X-rays. At high energies, in particular, Neilsen et al. (2013) analyzed 3 million seconds of *Chandra* data dedicated to characterizing both short and long-term X-ray variability of Sgr A*. They found that the length of flares varies from a few hundred seconds to 8 ks, with luminosities from $\sim 10^{34}$ to 2×10^{35} erg s $^{-1}$. Eckart et al. (2004) carried out simultaneous observations in both the X-ray and near-IR, and found that X-ray flares always have a coincident IR flare. In contrast there are numerous IR flares without X-ray counterparts. These flares can offer unique insight into the particle energetics of the accretion flow; highly energetic stochastic flares point toward mechanisms such as shocks and magnetic reconnection, which are capable of quickly producing large numbers of high energy non-thermal particles which significantly alter the observational signatures of the accretion flow.

In this paper, we investigate the effect of incorporating the emission from non-thermal electrons in GRMHD simulations of Sgr A*. We consider two injection models throughout the flow for a power-law population of electrons. In one, the non-thermal electrons simply track the thermal energy throughout the flow. In the second, non-thermal electrons are injected only into regions of possible magnetic reconnection, characterized by a low plasma β , where

$$\beta \equiv \frac{P_{\text{gas}}}{P_{\text{magnetic}}}. \quad (1)$$

We find that X-ray variability is a generic result of localizing non-thermal electrons to highly magnetized regions and that the timescales for such events are comparable to those of the observed X-ray flares. These results hold when we consider the cooling and advection of the non-thermal electrons in the flow.

2. GRMHD MODELS

Chan et al. (2015b) performed a large study of the broadband, time-dependent emission properties of Sgr A*. In these studies, they employed the GRMHD code HARM (Gammie et al. 2003; Narayan 2012) in conjunction with the efficient radiative transfer algorithm GRay (Chan et al. 2013) and varied the BH spin, density normalization, observer inclination, initial magnetic field configuration, and electron thermodynamic prescription. Among the models in this large parameter space, they chose five that fit the steady-state broadband spectrum of Sgr A* and its observed 1.3 mm image size.

Chan et al. (2015a) then studied the variability properties of the five best fit models at four different frequencies: in the radio at 10^{10} Hz and 1.3 mm, in the infrared at $2.17 \mu\text{m}$, as well as in the X-rays at 4.3 keV. In particular, two models from the study, with a black hole spin of $a = 0.7$ and $a = 0.9$, standard and normal evolution initial magnetic field configuration, and constant electron temperature in the funnel (hereafter referred to as models A and B), showed persistent variability as well as rapid flaring events in the IR and at 1.3 mm, consistent with observations. Chan et al. (2015a) identified these variations as being caused by the dynamic nature of magnetic flux tubes in

the flow combined with gravitational lensing that occurs when a flux tube crosses a caustic behind the black hole. None of the models, however, reproduced any X-ray variability: the X-ray lightcurves were extremely smooth, lacking any notable features from short-timescale variability to longer flaring events. This is expected given that the X-rays are produced by thermal Bremsstrahlung emission over the entire simulation volume but also indicates that these GRMHD models are missing the physics that causes the rapid variation in the X-ray flux.

The models in the earlier study use the ideal MHD approximation, where magnetic field dissipation and particle acceleration is not explicitly modeled. For the remainder of this paper, we use model B from Chan et al. (2015b) and incorporate a population of high energy power-law electrons in the postprocessing radiative transfer calculations, which may be accelerated from magnetic reconnection, as described in the following section.

3. INCORPORATING NON-THERMAL ELECTRONS

A population of particles in the accretion flow will evolve according to the continuity equation (written here in flat spacetime for simplicity)

$$\frac{\partial n_\gamma}{\partial t} + \mathbf{v} \cdot \nabla n_\gamma + (\nabla \cdot \mathbf{v}) n_\gamma = \left. \frac{dn_\gamma}{dt} \right|_{\text{inj}} + \gamma \left. \frac{dn_\gamma}{d\gamma} \right|_{\text{cool}} \quad (2)$$

where n_γ is the number density of electrons with Lorentz factor γ and $\mathbf{v}$ is their bulk velocity. The first term describes the evolution of the electron energy distribution, the second term represents the advection of particles with the flow, the third term describes the effect of the convergence/divergence of the flow, the fourth term describes the rate of injection of particles from acceleration processes such as magnetic reconnection or shocks, and the final term accounts for the cooling of particles, with γ being the radiative cooling rate.

The rest of this paper focuses on understanding each of these different terms and their relative contributions, which depend strongly on the local properties of the accretion flow. Because both the fluid velocity and its density have similar power-law dependences on radius, the second and third terms in the above equation have comparable magnitudes. For this reason, we will not consider further the terms that describes the convergence/divergence of the flow.

3.1. Injection of Non-thermal Electrons

We consider two configurations for the injection of non-thermal electrons in the accretion disk, i.e., the fourth term in Equation (2). The first model is based on the assumption that some fraction of the electron heating will continuously go into the acceleration of a non-thermal population and this fraction is independent of conditions in the flow. This results in a steady and uniform injection of non-thermal electrons, where the population of high energy particles simply follows the thermal energy in the system. We refer to this scenario as the “uniform” or “steady-state” distribution.

The second configuration is physically motivated by PIC simulations of magnetic reconnection, the process through which opposing magnetic fields are pushed together and dissipate. Magnetic reconnection rapidly injects energy into a small region of plasma, which has been shown to generate a

large population of high energy power-law electrons in regions of low β (Sironi & Spitkovsky 2014). Motivated by these PIC simulations, we pick a threshold value, $\beta_t = 0.2$, typical of the funnel or highly magnetized flux tubes, below which we inject non-thermal electrons (note that lower β indicates higher magnetic pressure). In this configuration, the dynamic nature of the flux tubes causes the injection to be more variable. As such, we refer to this model as “stochastic” or “nonuniform.”

3.2. Radiative Cooling of Power-law Electrons

The densities in the flow are sufficiently low that thermalization via Coulomb collisions is negligible, while the relatively high magnetic fields in the system imply that synchrotron radiation is the dominant cooling process. In order to discuss the role of synchrotron radiation in the cooling of electrons, we first define our electron energy distribution.

We consider a power-law distribution of electrons motivated by the results of magnetic reconnection models

$$n_\gamma d\gamma = C\gamma^{-p}d\gamma, \quad (3)$$

where p is the power law index, and C is a normalization constant. This distribution of non-thermal electrons will radiate predominantly via synchrotron due to the presence of magnetic fields. The synchrotron power per unit volume per unit frequency emitted by this distribution is given by (e.g., Rybicki & Lightman 1979):

$$P_{\text{tot}}(\omega) = \frac{\sqrt{3}q^3CB\sin\alpha}{4\pi m_e c^2(p+1)}\Gamma\left(\frac{p}{4} + \frac{19}{12}\right) \times \Gamma\left(\frac{p}{4} - \frac{1}{12}\right)\left(\frac{m_e c\omega}{3qB\sin\alpha}\right)^{-(p-1)/2}, \quad (4)$$

where q and m_e are the charge and mass of an electron, C is the normalization constant from Equation (3), B is the magnetic field strength, α is the pitch angle between the electron velocity and magnetic field, and Γ is the gamma function. The non-thermal electrons will also add to the opacity of the system as

$$\alpha_\nu = \frac{\sqrt{3}q^3}{8\pi m_e}\left(\frac{3q}{2\pi m_e^3 c^5}\right)^{p/2} C(B\sin\alpha)^{(p+2)/2} \times (m_e c^2)^{p-1}\Gamma\left(\frac{3p+2}{12}\right)\Gamma\left(\frac{3p+22}{12}\right)\nu^{-(p+4)/2}. \quad (5)$$

Assuming an isotropic velocity distribution, we average Equations (4) and (5) over the pitch angle in our calculations.

We assign a fraction of the total thermal energy to this power-law distribution, which then radiates and absorbs photons according to Equations (4) and (5). The energy density of thermal electrons at temperature θ_e is (e.g., Chandrasekhar 1939)

$$u_{\text{th}} = a(\theta_e)N_{\text{th}}m_e c^2\theta_e, \quad (6)$$

where N_{th} is the number density of thermal electrons, $a(\theta_e)$ is given by

$$a(\theta_e) \equiv \frac{1}{\theta_e}\left[\frac{3K_3(1/\theta_e) + K_1(1/\theta_e)}{4K_2(1/\theta_e)} - 1\right], \quad (7)$$

and K_n are modified Bessel functions of order n . The quantity $a(\theta_e)$ varies from $3/2$ to 3 , corresponding to a non-relativistic

and fully relativistic electron gas, respectively. In order to aid computation, we use an approximate form for $a(\theta_e)$ given by Gammie & Popham (1998), which has less than 2% error for all temperatures:

$$a(\theta_e) = \frac{6 + 15\theta_e}{4 + 5\theta_e}. \quad (8)$$

We now introduce a free parameter, η , which describes the fraction of thermal energy assigned to a power-law distribution. The non-thermal energy density in a given cell is, therefore, $u_{\text{pl}} = \eta u_{\text{th}}$. The quantity C , in Equations (4) and (5) is related to η by

$$\eta u_{\text{th}} = \int_{\gamma_1}^{\gamma_2} C(\gamma mc^2)\gamma^{-p}d\gamma. \quad (9)$$

We choose $\gamma_1 = 1$ and set γ_2 to be very large so that

$$C \approx \eta a(\theta_e)N_{\text{th}}\theta_e(p-2). \quad (10)$$

With this set up, we can calculate the necessary quantities to perform the radiative transfer calculation while accounting for a population of non-thermal electrons described by the quantities p and η . In Section 4, we discuss the results of this calculation for our two different models, which are the uniform steady state injection and the β -dependent, nonuniform, stochastic injection. In Section 6 we more rigorously discuss the advection and synchrotron cooling timescales to complete our analysis of Equation (2).

4. QUIESCENT CONSTRAINTS

In Chan et al. (2015b), purely thermal models were fit to a number of observed quantities, including to the quiescent X-ray flux at 10^{18} Hz ($\equiv 4.1$ keV), which originates predominantly from the extended halo of gas emitting via Bremsstrahlung. The thermal models were calibrated to reproduce the appropriate observed time-averaged X-ray flux for the size of the simulation, corresponding to 10% of the total observed flux. As such, when we include the non-thermal electrons, a natural requirement is that the quiescent X-ray flux must not change significantly compared to the thermal model.

We show in the left panel of Figure 1 the results of the uniform steady-state model with a power-law index of 3.5, which produces a spectral index within the bounds of observational constraints (Porquet et al. 2008; Barrière et al. 2014) and is motivated physically by PIC simulations of magnetic reconnection for magnetizations on the order of the regions we are considering (Sironi & Spitkovsky 2014). It is evident that, in order to match the X-ray flux, the fraction of energy in non-thermal electrons should be quite small, constrained to values of the order of 0.001. This implies that there cannot be a very large population of non-thermal electrons existing everywhere throughout the flow at any given time; even a relatively small number of these high energy electrons will result in too large of an X-ray flux if they are distributed throughout the entire simulation region.

In the right panel of Figure 1, we show the spectrum from the nonuniform stochastic model, where non-thermal electrons are localized to low β regions, again with a power-law index of 3.5. Because of this localization, it is possible to accommodate higher fractions of non-thermal electrons while still matching the quiescent thermal spectrum. In this model, it is possible to

Figure 1. Left: spectra computed for quiescent (i.e., non-flaring) times for various values of the energy fraction of non-thermal electrons, η , with fixed power-law index, $p = 3.5$, as well as for the purely thermal model. In this configuration where the non-thermals only follow the thermal energy, the observed quiescent X-ray flux at 10^{18} Hz (depicted with an X) is exceeded even for moderate values of η . Right: in this configuration, non-thermal electrons are injected in regions below $\beta = 0.2$. Localizing the non-thermal electrons to highly magnetized regions, where they are more likely to be accelerated, allows for significantly higher values of η while still accommodating the observed quiescent X-ray flux at 10^{18} Hz.

inject almost 10% of the total thermal energy into a non-thermal electron distribution within the magnetized regions.

5. X-RAY VARIABILITY: STOCHASTIC INJECTION

Apart from providing a more natural match to quiescent-state constraints, another interesting result of using the β -dependent description of non-thermal electron injection is that it produces significant X-ray variability. Perhaps this is not surprising, since the non-thermal electrons will trace magnetic flux tubes, which are dynamic structures, constantly being formed, sheared, and moving throughout the flow. If one of these flux tubes crosses a caustic behind the black hole, it will result in an additional amplification of the flux, and since these tubes are emitting primarily non-thermal synchrotron radiation in the X-rays, they will cause X-ray flares.

We explore this variability in Figure 2, where we show the effect of stochastic injection of non-thermal electrons and compare it to a purely thermal model. In the non-thermal lightcurve, we see both persistent variability as well as 4 large flares during the ~ 80 hr of simulation. In the largest flares, the flux increases by a factor of ~ 10 compared to quiescence. The magnetically dominated regions responsible for these flares live for about 5000 s, which sets the timescale of the flares in this figure. There is indeed a stark contrast between this result, which takes into account acceleration in low β regions, and the purely thermal model, which shows no variability.

We now investigate the properties of the magnetic structures in the innermost regions of the accretion flow to further pinpoint the localization and time evolution of the flares. Figure 3 shows the ratio of the magnetic to the gas pressure throughout the inner flow during a quiescent state and during the strongest flare from the simulation. We see that this flare is caused by a large magnetic flux region developing in the flow with $\beta < \beta_r$. Due to the large spatial extent of this tube, many non-thermal electrons are injected, causing a sudden increase in the X-ray flux. In contrast, during quiescence, the only region with a significant number of non-thermal electrons is in the funnel, which typically has a fairly uniform and strong magnetic field. This only contributes a small flux and results in low level variability.

Figure 2. Thermal and non-thermal X-ray lightcurves. The injection of non-thermal electrons into highly magnetized regions naturally produces significant variability due to the dynamic nature of magnetic fields in the accretion flow.

Figure 4 depicts the spectra of the flaring and quiescent states from the simulation. During quiescence, the non-thermal emission is not especially prominent; its nature is largely obscured by the thermal emission dominating at most wavelengths. During the flare, however, the power-law nature of the non-thermal emission becomes more evident.

In Figure 5, we show the X-ray images of our model during flaring and quiescent states. The images show the relative contribution to the overall flux from various parts of the accretion flow. During quiescence, we find that while there is some contribution to the X-ray flux from a small population of non-thermal particles in the funnel, the extended Bremsstrahlung emission accounts for the majority of the total (i.e., integrated over the entire image) observed flux. During a flare, the emission is heavily dominated by non-thermal electrons in the inner accretion flow, rendering the Bremsstrahlung flux negligible. The difference in the localization of the X-rays

Figure 3. Left: map of the ratio of the magnetic to gas pressure during a quiescent state in the simulation. Cells near the pole and within the ISCO at $\sim 2.4 GMc^{-2}$ are excised due to numerical artifacts often occurring within these regions. The black quarter-circle at the origin is the event horizon of the black hole. Right: a large flux tube is present in the accretion flow during this flare, with high magnetization, resulting in a high ratio of pressures throughout a large portion of the disk.

Figure 4. Spectra of the flaring and quiescent states depicted in Figure 3 in red and orange, respectively. The purely thermal quiescent spectrum is shown for reference in blue.

between the non-thermal and thermal models is responsible for their different variability properties.

6. COOLING AND ADVECTION TIMESCALES

If the production of non-thermal electrons in the flow are indeed stochastic and localized, their lifetime will be determined by the cooling and advection timescales relative to the injection timescale. In this section, we consider the cooling and advective terms in Equation (2) and compare the typical advective timescales for matter to be drawn through the event horizon to the synchrotron cooling times. These timescales will give us a rough estimate of how long a flare will last.

The synchrotron cooling time is given by (e.g., Rybicki & Lightman 1979)

$$t_s = \frac{6\pi m_e c}{\sigma_t \beta_e^2 \gamma B^2}, \quad (11)$$

where σ_t is the Thompson cross section, m_e is the electron mass, c is the speed of light, β_e is the particle velocity as a fraction of the speed of light, γ is the Lorentz factor, and B is the magnetic field strength.

We calculate the synchrotron cooling time for the innermost region in the simulation, considering fairly relativistic electrons with $\gamma = 40$ (a Lorentz factor relevant to X-ray production that is easily generated by relativistic reconnection) and show the results in Figure 6.

The cooling timescale in the funnel is of the order of thousands of seconds, while it varies greatly in the disk due to the highly variable magnetic field. The magnetized flux tubes have cooling times of the order of hundreds to thousands of seconds, depending on the magnetization of the region, while the filamentary regions in between them have extremely long timescales, corresponding to weak magnetic fields. Most importantly, we see that the synchrotron cooling timescales in the regions where we expect there to be non-thermal electrons (i.e., low β regions) are comparable to the observed flare durations of a few hundred seconds to 8000 s, reported in Neilsen et al. (2013).

Next, we consider the radial advective timescale, which is the time it will take electrons to move radially a distance r in the flow. We define the advection timescale as $t_r \equiv -r/v_r$. In the model we are considering, all the fluid within 10 gravitational radii of the event horizon has an inward (negative) radial velocity, such that t_r is always positive, as shown in Figure 6.

Because radial advection is inward, the timescale associated with this motion is relevant to the estimation of flare duration. If a population of non-thermal electrons is injected and quickly disappears through the event horizon due to radial advection, then the length of the flare may be dictated by the advection timescale. If, however, $t_r > t_s$, then the synchrotron cooling will be primarily responsible for the dissipation of non-thermal energy.

The right panel of Figure 6 shows that advective timescales in the funnel are fairly short: any non-thermal electrons injected into the funnel will quickly pass through the event horizon within hundreds of seconds. Centrifugal support in

Figure 5. Left: simulated image of quiescent X-ray (4.1 keV) emission. Fluxes are normalized to the maximum pixel value. Some structure is visible in the innermost regions of the image, where strong magnetic fields in the funnel close to the event horizon have associated non-thermal particles, and hence strong X-ray emission. We see that the extended Brehmsstrahlung emission comprises a significant fraction of the total flux during quiescence. Right: during the flare, emission is heavily dominated by the innermost part of the accretion flow; the relative contribution from the halo of Brehmsstrahlung emission is negligible during flares.

Figure 6. Left: synchrotron cooling timescale t_s in the inner accretion flow. Cooling timescales in the funnel are of the order of a few thousand seconds, while cooling timescales in the plane of the disc vary greatly due to the nonuniformity of the magnetic field. Right: radial advection timescale, t_r , i.e., an estimate of how long it will take for the matter at a given point to pass through the event horizon. Matter in the funnel is quickly advected on the order of hundreds of seconds, while matter in the plane of the disk, $z = 0$, is centrifugally supported and has radial advection timescales of the order of hundreds to thousands of seconds.

the disk results in longer timescales, allowing for a population of instantaneously injected non-thermal electrons to radiate their energy away over the hundreds to thousands of seconds it takes for synchrotron emission to cool them. As such, when we interpret observations in the context of a GRMHD model such as model B, we conclude that the flaring events predominantly originate from the disk. In other words, in order for a flare to last for many thousands of seconds as observed, the non-thermal electrons must be in the disk where they are centrifugally supported and able to exist for these timescales in the flow without disappearing through the event horizon.

We now consider the advective and synchrotron timescales in terms of the azimuthal dynamical time, $t_\phi = 2\pi r/v_\phi$. By doing this, we get a sense of how much a population of electrons will be able to spread azimuthally throughout the flow

as they cool. In Figure 7, we plot the synchrotron and advective cooling times in units of azimuthal dynamical time, taking a slice of the simulation through the disk ($z=0$) shown in Figure 6, and a slice through the funnel, making a 60° angle with the horizontal axis in Figure 6.

We see the same general features we already discussed in terms of the structure and relative magnitude of synchrotron and advective timescales in both the funnel and the disk. However, we now get more physical insight into the dynamics of the system: we find that, in the highly magnetized flux tubes, corresponding to local minima of the dashed red line in the left panel of Figure 7, the synchrotron cooling is sufficiently fast that a flare will be localized to a fairly small azimuthal region, since $t_s < t_\phi$. Similarly, we find that any matter injected into the funnel will be radially advected much more quickly than it is able to spread out azimuthally.

Figure 7. Left: synchrotron cooling timescales in terms of azimuthal dynamical time. The synchrotron cooling in the funnel is fairly constant due to the uniform magnetic field. The disk synchrotron is highly varied due to the structure of the magnetic field. Right: radial advective timescale in terms of azimuthal dynamical time. The centrifugal support in the disk allows the radial time to be comparable to the azimuthal time, while matter in the funnel will last at most for one-tenth of an orbit before being drawn through the event horizon.

7. COMPARISON TO OBSERVATIONS

By interpreting observations of Sgr A* in the context of our GRMHD simulations, we begin to place constraints on the population of non-thermal electrons in the radiatively inefficient accretion flow and gain insight into their injection mechanisms. We employed two configurations for non-thermal electron injection, one in which the non-thermal electrons simply track the thermal energy everywhere throughout the flow and another where the non-thermal electrons are injected solely into regions of high magnetization. From the first model we are able to place tight constraints on the fraction of steady-state non-thermal electrons that may exist throughout the flow by comparing the simulations to the observed quiescent X-ray flux. The second model localizes non-thermal electrons to a much smaller region, allowing their local energy density to be much higher than the uniform model (by about 2 orders of magnitude) while still matching the observed quiescent X-ray flux.

We find that X-ray variability is a generic result of constraining the non-thermal electrons to highly magnetized regions. This is because the magnetic field is dynamic throughout the flow, generating magnetic flux tubes, which are in a constant state of being formed, sheared, becoming buoyant, and leaving the disk. The dynamic nature of these flux tubes combined with strong lensing effects from the black hole generate both persistent variability as well as flaring events.

X-ray flares in our simulations are always coincident with IR flares, but there are numerous IR flares without X-ray counterparts, as shown in Figure 8, which qualitatively matches observations. In this figure, we have zoomed in to the first 250 timesteps of the simulation in order to more clearly illustrate the relationship between the IR and X-ray lightcurves. During this time span, we observe about 5 IR flares over the stochastically variable background and one significant X-ray flare. From our simulation we find that there are about 5 IR flares per X-ray flare and a rate of one X-ray flare per 72,000 s, over the entire simulation. Over the course of 3 million seconds of observation with *Chandra*, 39 X-ray flares were observed, corresponding to one flare every $\sim 77,000$ s. Observations Sgr A* show that, for every X-ray flare, there are about 4 NIR flares (e.g., Genzel et al. 2003; Eckart et al. 2006). These numbers are in rough agreement with our results.

In order to compare flare statistics from our simulations to observations more directly (e.g., through flux distributions reported in Neilsen et al. 2015), we need to account for the fact that only $\sim 10\%$ of the X-ray emission from Sgr A* comes from the inner accretion flow (Neilsen et al. 2013). In Figure 9, after adding a constant background equal to 90% of the observed quiescent flux to our lightcurve from Figure 2, we plot the flux distribution in our simulations. We find that the flare distribution resembles a power-law with an index of around 2.3, while the lower-level variability does not have an obvious structure. Neilsen et al. (2015) reports only Poisson variability at low fluxes, and power-law behavior at high fluxes, with a power-law index of 1.92. This is roughly consistent with our simulated flux distribution. Our simulations, however, do not account for the Poisson photon counting noise, and also do not show as large of a range of variability. The latter is likely due to the relatively short duration of our simulation that did not capture many rare, high flux events.

We estimate the level of the Poisson noise, below which we expect our simulated flux distribution to deviate significantly from observations. We take the reported photon counting rate ($Q = 5.24$ cts/ks) and binning ($b = 300$ s) from Neilsen et al. (2015) and calculate the typical fractional Poisson error, $\epsilon = 1/\sqrt{Qb} = .79$. Normalizing our lowest level of emission to 1, we see that counting noise will dominate the observed variability from 1 to $1 + \epsilon$, setting the lower limit from which we expect our simulations to reproduce the flux distribution.

We further explore the relationship between IR and X-ray fluxes in Figure 10. The largest IR flares correspond to the largest X-ray flares, but there is much more variability in the IR than the X-rays. In our simulations, anytime a flux tube appears and crosses a caustic there will be an IR flare due to the synchrotron emission from the thermal electrons. However, only the most highly magnetized flux tubes will have non-thermal electrons associated with them and will generate an X-ray flare. The effect of the β threshold is to pick out a subset of all the magnetic flux tubes, ones with conditions suitable for reconnection to occur. The particular threshold we use is motivated by Li et al. (2015a), who showed a non-thermal component being generated for $\beta < 0.2$. As a result, IR variability is much more significant, since there is thermal synchrotron associated with all flux tubes, whereas particle acceleration and hence X-ray emission only occurs for a

Figure 8. IR and X-ray lightcurves zoomed in on the first 250 timesteps of the simulation. Each curve is normalized to a fiducial quiescent flux. Note the strong and rapid variation in the IR flux and moderate variability in the X-ray. The IR flaring has been described in Chan et al. (2015a). The inclusion of non-thermal electrons in highly magnetized regions has produced significant X-ray variability which was previously unseen.

Figure 9. X-ray flux distribution, accounting for a constant quiescent background. At high fluxes, the flare distribution resembles a power-law with an index of ~ 2.3 . Using the Poisson rate and binning reported in Neilsen et al. (2015) for the *Chandra* observations used in that study, we estimate what would be the upper end of the Poisson-dominated regime, depicted with a vertical dotted line.

particular subset of the tubes. Additionally, we see that the flux tubes responsible for the largest IR flares are the same structures responsible for the largest X-ray flares. This is unsurprising given the strong scaling of synchrotron emissivity with magnetic field; the most highly magnetized flux tubes radiate copiously in the IR due to the high magnetic fields, and also act as sites of efficient reconnection, generating strong X-ray flares.

8. CONCLUSIONS

In this paper, we explored the effects of incorporating non-thermal electrons into GRMHD simulations of radiatively inefficient accretion flows. We found that X-ray variability is a generic result of constraining the non-thermal electrons to

Figure 10. IR vs. X-ray flux, accounting for a constant quiescent background the of X-ray flux (blue), and purely for the inner accretion flow (green). The orange line depicts a correlation with a slope of unity. The addition of the quiescent background decreases the X-ray variability by a factor of ~ 10 . We see a general trend of higher IR fluxes being associated with relatively high X-ray fluxes.

highly magnetized regions and that the timescales associated with electron cooling are comparable to the observed flare duration. This analysis is model-dependent, since the synchrotron radiation from non-thermal electrons depends on the strength and topology of the magnetic field, which will impact the constraints we place on the quiescent energy budget of the non-thermal electrons. Our analysis of cooling timescales are also likely to differ across models. For model B with only inflowing velocities within 10 gravitational radii, we found that flares are unlikely to originate from the funnel since the inward radial velocities were too large to explain the X-ray flares lasting many thousands of seconds. We will explore the role of non-thermal electrons in producing the variability for different magnetic field configurations and black hole spins in a future study.

In the context of model B, which has matched many observational constraints, from the quiescent broadband spectrum to the variability properties, we find that X-ray flares likely originate from magnetic flux tubes in the disk, where centrifugal support allows the non-thermal electrons to remain in the flow for many thousands of seconds while radiating away their energy via synchrotron emission. The flare lengths are hence set by the synchrotron cooling timescales. The properties of the X-ray variability from this model are consistent with observations: X-ray flares are always coincident with IR flares, there are many more IR flares without associated X-ray counterparts, and the timescales associated with the flares are comparable to the observed flare duration.

We thank Joey Neilsen, Daryl Haggard, Daniel Marrone, and Lia Medeiros for useful discussions. We gratefully acknowledge support for this work from *Chandra* Award No. TM6-17006X and from NASA/NSF TCAN award NNX14AB48G. All ray tracing calculations were performed with the El Gato GPU cluster at the University of Arizona that is funded by NSF award 1228509.

REFERENCES

- Barrière, N. M., Tomsick, J. A., Baganoff, F. K., et al. 2014, *ApJ*, **786**, 46
- Chan, C.-k., Psaltis, D., & Özel, F. 2013, *ApJ*, **777**, 13
- Chan, C.-k., Psaltis, D., Özel, F., et al. 2015a, *ApJ*, **812**, 103
- Chan, C.-K., Psaltis, D., Özel, F., Narayan, R., & Sad,owski, A. 2015b, *ApJ*, **799**, 1
- Chandrasekhar, S. 1939, *An Introduction to the Study of Stellar Structure* (New York: Dover)
- Dexter, J., Agol, E., Fragile, P. C., & McKinney, J. C. 2012, *JPhCS*, **372**, 012023
- Dodds-Eden, K., Sharma, P., Quataert, E., et al. 2010, *ApJ*, **725**, 450
- Drappeau, S., Dibi, S., Dexter, J., Markoff, S., & Fragile, P. C. 2013, *MNRAS*, **431**, 2872
- 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
- Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, *ApJ*, **589**, 444
- Gammie, C. F., & Popham, R. 1998, *ApJ*, **498**, 313
- Genzel, R., Schödel, R., Ott, T., et al. 2003, *Natur*, **425**, 934
- Guo, X., Sironi, L., & Narayan, R. 2014, *ApJ*, **794**, 153
- Lazarian, A., Vlahos, L., Kowal, G., et al. 2012, *SSRv*, **173**, 557
- Li, X., Guo, F., Li, H., & Li, G. 2015a, *ApJL*, **811**, L24
- Li, Y.-P., Yuan, F., Yuan, Q., et al. 2015b, *ApJ*, **810**, 19
- Mahadevan, R. 1998, *Natur*, **394**, 651
- Marrone, D. P., Baganoff, F. K., Morris, M. R., et al. 2008, *ApJ*, **682**, 373
- Mościbrodzka, M., Falcke, H., Shiokawa, H., & Gammie, C. F. 2014, *A&A*, **570**, A7
- Narayan, R., Sad,owski, A. A., Penna, R. F., & Kulkarni, A. K. 2012, *MNRAS*, **426**, 3241
- Neilsen, J., Markoff, S., Nowak, M. A., et al. 2015, *ApJ*, **799**, 199
- Neilsen, J., Nowak, M. A., Gammie, C., et al. 2013, *ApJ*, **774**, 42
- Özel, F., Psaltis, D., & Narayan, R. 2000, *ApJ*, **541**, 234
- Ponti, G., De Marco, B., Morris, M. R., et al. 2015, *MNRAS*, **454**, 1525
- Porquet, D., Grosso, N., Predehl, P., et al. 2008, *A&A*, **488**, 549
- Quataert, E., & Gruzinov, A. 1999, *ApJ*, **520**, 248
- Ressler, S. M., Tchekhovskoy, A., Quataert, E., Chandra, M., & Gammie, C. F. 2015, *MNRAS*, **454**, 1848
- Rybicki, G. B., & Lightman, A. P. 1979, *Radiative Processes in Astrophysics* (New York: Wiley)
- Sironi, L., & Spitkovsky, A. 2014, *ApJL*, **783**, L21
- Witzel, G., Morris, M., Ghez, A., et al. 2014, in *IAU Symp. Proc. 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
- Yuan, F., & Narayan, R. 2014, *ARA&A*, **52**, 529
- Yuan, F., Quataert, E., & Narayan, R. 2003, *ApJ*, **598**, 301
- Yuan, Q., & Wang, Q. D. 2016, *MNRAS*, **456**, 1438