Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 12 November 2017 (MN LaTeX style file v2.2)

Black Holes in GIZMO

Philip F. Hopkins[⋆1,2], Claude-André Faucher-Giguère[2], Daniel Angles-Alcazar, Paul Torrey, Xiangcheng Ma

[1]*TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA*

[2]*Department of Astronomy and Theoretical Astrophysics Center, University of California Berkeley, Berkeley, CA 94720*

–, 2013

ABSTRACT

This is a description of the AGN modules in the code

Key words: star formation: general — cosmology: theory

1 INTRODUCTION

Important note: more detail about the numerical switches is in the GIZMO users guide. This note gives more background on the physical motivations and formulae being used. Users should enable BLACK_HOLES for any BH physics.

2 BLACK HOLE FORMATION & "SEEDS"

2.1 In ICs

Often one sets up BHs in the ICs, so there is no formation. But if needed, on-the-fly seeding can be enabled.

2.2 "On-the-fly" Seeding from Star Formation

This is the most physically uncertain element here. Turning on BH_SEED_FROM_LOCALGAS therefore adopts a very simple prescription. When a gas particle is flagged as being "turned into" a star particle (which occurs according to the normal SF requirements), then we assign it some probability of instead turning into a "seed" BH, where the probability increases in higher-density, lower-metallicity gas. This form is in a simple function and easy to modify. But as a default, we adopt the form:

$$\frac{\mathrm{d}P_{\rm seed}}{\mathrm{d}M_{*}} = \frac{P_0}{M_{\rm seed}}\left(1-\exp\left[-\frac{\Sigma}{\Sigma_0}\right]\right)\exp\left(-\frac{Z}{Z_0}\right) \tag{1}$$

(we use surface density rather than density because this seems to better correspond to where dense star clusters form in higher-resolution simulations by M. Grudic, and matches onto arguments for when feedback is inefficient in dense high-redshift disks; but the choice is arbitrary, and can be varied). This includes all the parameters that must be set for the seeding model.

The default parameters in the code are $Z_0 = 0.01\,Z_{\odot} = 0.0002$, $\Sigma_0 = 1\,{\rm g\,cm^{-2}}$, and $P_0 = 0.0004$, with $M_{\rm seed}$ set by the run-time parameterfile parameter "SeedBlackHoleMass" (in code units).

2.3 "On-the-fly" Seeding from Halo Finding

Alternatively, if BH_SEED_FROM_FOF is enabled, the code will periodically run an on-the-fly friends-of-friends halo finder, and then when a sufficiently massive halo (of either DM or stars) is identified, places a BH in the center of that halo. The user specifies both the seed mass and minimum halo/stellar mass of the groups which will get a seed.

3 BLACK HOLE DYNAMICS

Physically, the BHs (even for the smallest seeds we model) are much larger masses than individual stars (let alone dark matter particles or gas molecules). So they should experience a dynamical friction force. When the BH mass $M_{\rm BH}$ is much larger than typical particle masses in the simulation $\langle m_i \rangle$, this will be resolved and treated accurately. But if – for purely numerical reasons – the BHs begin from small "seeds" with $M_{\rm BH} \lesssim \langle m_i \rangle$, this cannot be captured. This can be important if it determines, for example, the ability of small BHs to sink to the center of star clusters or proto-galaxies in which they form.

3.1 Dynamical Friction

We can therefore explicitly include a dynamical friction term (enabling BH_DYNFRICTION), following the standard Chandrasekhar expression:

$$\frac{d\mathbf{v}_{\rm BH,\,DF}}{dt} = \left(\frac{m_{{\rm m}\,i}}{M_{\rm eff}+m_{{\rm m},\,i}}\right)\frac{4\pi\,G^2\,\langle\rho\rangle_i\,M_{\rm eff}\,\ln\Lambda}{|\delta\mathbf{v}|^3}\,f\left(\frac{|\delta\mathbf{v}|}{\sqrt{2}\,\sigma_i^{v}}\right)\delta\mathbf{v} \tag{2}$$

$$f(x) \equiv {\rm erf}(x) - \frac{2}{\pi^{1/2}}\,x\,\exp\left(-x^2\right) \tag{3}$$

$$\Lambda \approx 1 + \frac{b_{\rm impact}\,|\delta\mathbf{v}|^2}{G\,M_{\rm eff}} \tag{4}$$

where $M_{\rm eff} = M_{\rm BH} + M_{\alpha\,{\rm disk}}$ ($M_{\alpha\,{\rm disk}}$ is the mass of the viscous accretion disk "carried" by the BH, discussed below), $\delta\mathbf{v} \equiv \langle\mathbf{v}\rangle_i - \mathbf{v}_{\rm BH}$ is the velocity of the BH relative to $\langle\mathbf{v}\rangle_i$ the mass-weighted mean local velocity of all particles in the BH kernel, $\langle\rho\rangle_i$ and σ_i^{v} are the mass density and rms velocity dispersion of the background.[1] Here $b_{\rm impact}$ is the maximum impact parameter out to which the Coulomb logarithm is extrapolated. For convenience, we set this to $\sim 50\,$kpc, representative of a typical halo virial radius of interest. However we stress for values of the Coulomb logarithm, changing this by a factor of ~ 10 makes a $\sim 20\%$ difference to $\ln\Lambda$, much smaller than any other uncertainties in the expression.

The term $m_{{\rm m},\,i}/(M_{\rm eff}+m_{{\rm m},\,i})$ we add to interpolate between the cases where we need to include this term ($M_{\rm eff} \lesssim \langle m_i\rangle$) and the cases where the code should explicitly handle dynamical friction ($M_{\rm eff} \gg m_{{\rm m},\,i}$), so that we prevent double-counting it in the latter limit. The best definition of $m_{{\rm m},\,i}$ depends in detail on numerics (how gravity is softened, for example), but for practical purposes we find well-behaved results in simple tests by setting it equal to

[1] All "background" calculations include gas, stars, and dark matter, but exclude the BHs.

⋆ E-mail:phopkins@caltech.edu

about $\sim 3-10$ times the mass of the most massive non-BH particle in the kernel. We will adopt a canonical value of $m_{\mathrm{m},i} = 5\,\mathrm{MAX}(m_i)$ for our standard reference.

3.2 Drag on the BH

At sufficiently low resolution, even the dynamical friction estimator works poorly, because it is dominated by particle noise and local clumps. Alternatively, enabling BH_DRAG adds a drag acceleration to the BH, of the form $\mathbf{a}_{BH} = (\mathbf{v}_{\rm gas} - \mathbf{v}_{\rm BH})\,\dot{M}_{\rm BH}/M_{\rm BH}$ (nominally motivated by the BH gaining momentum from the accreted gas, continuously). Setting this parameter equal to 1 does the same but with $\dot{M}_{\rm BH}$ replaced by the Eddington accretion rate (less physically motivated, but actually keeps the BH anchored when the accretion rate is low, where the standard prescription would give no drag).

3.3 Anchoring the BH to Potential Minimum

Most extreme, BH_REPOSITION_ON_POTMIN always moves the BH to the local potential minimum (within the kernel). While this avoids issues with noisy fields or low accretion rates that the drag or dynamical friction formulae can encounter, it still isn't perfect. In complicated geometries (e.g. mergers), the BH can sometimes "walk" down a local (often noisy) gradient out of a galaxy!

4 BLACK HOLE-BLACK HOLE MERGERS

Whenever two BHs are inside the same smoothing kernel/resolution limit, we merge them if they are directly gravitationally bound to one another (i.e. have relative velocities below the mutual escape velocity of the two-BH system at the resolved separation). Numerically, this represents the physical coalescence of the BH binary below resolved scales, but of course cannot capture new dynamics on much smaller scales, so a single BH particle may physically represent a binary or multiple system.

At present, we do not include any "sub-grid" model for recoils or ejections in BH-BH mergers; however resolved many-body ejections can and do occur. Users are encouraged to explore these models which can easily be implemented in the BH routines after a BH-BH merger.

5 BLACK HOLE ACCRETION

Simulations cannot hope to simultaneously resolve galaxy scales and the true accretion scales (the Schwarzchild radius). We divide the problem into two "stages": capture of gas into the "traditional" non-star forming QSO accretion disk $\dot{M}_{\alpha\,\rm disk}$, and then accretion from this disk onto the BH $\dot{M}_{\rm BH}$. Each of these has a separately-tracked mass reservoir $M_{\alpha\,\rm disk}$ and $M_{\rm BH}$. All accretion models require BH_SWALLOWGAS.

5.1 Capture of Gas Into the Accretion Disk

5.1.1 The Resolved Limit (Direct Capture)

Define the outer radius of the "traditional" accretion disk $R_{\alpha\,\rm disk}$ (discussed below). If $R_{\alpha\,\rm disk}$ is resolved, then we can explicitly model capture into the disk. If any gas particle (gas, star, or dark matter) is located within $R_{\alpha\,\rm disk}$, then knowing its position and relative velocity with respect to the BH particle we check whether (a) it is gravitationally bound to the BH, and (b) whether the apo-centric radius of the particle about the BH is also $< R_{\alpha\,\rm disk}$. If both are true, we consider the particle "captured" and immediately add its mass to the accretion disk.

Note that we can do this for only gas particles (enabling BH_GRAVCAPTURE_GAS), only non-gas particles (enabling BH_GRAVCAPTURE_NONGAS), or both.

5.1.2 The Un-Resolved Limit: Gravitational Torques Model

The resolution needed to meaningfully apply the above prescription is only true in nuclear-scale simulations. In many cases, we cannot resolve $R_{\alpha\,\rm disk}$. In this limit one must adopt a "sub-grid" accretion prescription.

Enabling BH_GRAVACCRETION, we therefore have also implemented the model of Hopkins & Quataert (2011), which was specifically designed to reproduce the accretion rate resolved in several hundred simulations from large galactic scales ($\sim 1-100\,$pc) into the accretion disk ($< 0.1\,$pc). The dominant mechanism of angular momentum transfer on all of these scales is torques due to gravitational instabilities in the *gas plus stellar* disk. In Hopkins & Quataert (2010) we show that this holds even in turbulent systems with realistic stellar feedback – while by no means perfect, the approximation captures the most important qualitative behaviors, and it is *several orders of magnitude* more accurate than other accretion estimators commonly used (including variant "Bondi-Hoyle" accretion rates).

The main scalings derived and tested therein are *local* scalings for the angular momentum loss at a given radius. It is considerably more challenging to build a model of unresolved accretion from some large radius (say, $\sim 100\,$pc) to the BH, because we must make a number of assumptions about the un-resolved mass profile, disk-thickness, and where the transition to the "traditional" viscous disk occurs. In Hopkins & Quataert (2011), a scaling using simple assumptions for these terms was crudely estimated, for the accretion rate into the traditional disk ($\dot{M}_{\alpha\,\rm disk}$) from resolved gas and stars inside a larger radius R_0, is given by:[2]

$$\frac{\dot{M}_{\alpha\,\rm disk}}{M_{\odot}\,{\rm yr^{-1}}} \approx 10\,f_d^{5/2}\left(\frac{M_{\rm BH+\alpha\,disk}}{10^8\,M_{\odot}}\right)^{1/6}\left(\frac{M_{\rm disk}(<R_0)}{10^9\,M_{\odot}}\right) \times\left(\frac{R_0}{100\,{\rm pc}}\right)^{-3/2}\left(\frac{f_{\rm gas}(<R_0)}{f_{\rm gas}(<R_0)+f_0}\right) \quad (5)$$

with $f_0 \approx 0.31\,f_d^2\,(M_{\rm disk}[<R_0]/10^9\,M_{\odot})^{-1/3}$, the gas and disk fraction $f_{\rm gas} \equiv M_{\rm gas}(<R_0)/M_{\rm disk}(<R_0)$, $f_d \equiv M_{\rm disk}(<R_0)/M_{\rm tot}(<R_0)$.

Inside of some radius R_0 enclosing the BH then, we simply require the BH+α-disk mass, gas mass, total mass, and disk mass. Evaluating the "disk" mass in some radius is non-trivial, but a convenient proxy for $M_{\rm disk}/M_{\rm tot}$ is given by the ratio of the total angular momentum inside R_0 to that which would be present for a perfectly

[2] The accretion rate implied by this model is continuous, but particles are discrete. We therefore follow Springel et al. (2005) and allow the α-disk mass reservoir to grow continuously (increasing each timestep by $\Delta M_{\alpha\,\rm disk} = \dot{M}_{\alpha\,\rm disk}\,\Delta t$), but separately tracking the total "accreted particle mass" (sum of $M_{\rm acc} = \sum m_i$ of accreted gas particles); gas particles within R_0 can be stochastically selected to be "accreted" at each timestep then with a probability equal to $(\sum \Delta M_{\alpha\,\rm disk} - M_{\rm acc})/m_i$, where m_i is the mass to be accreted from particle i (weighted within R_0 by the SPH kernel). The accreted gas particles immediately have a fraction of their mass removed (described above in the feedback model) and added to $M_{\rm acc}$. This scheme allows continuous accretion but removes gas particle mass at a rate that statistically enforces mass conservation. The BH particle momentum is also corrected so that total momentum is conserved.

rotationally supported disk:[3]

$$\frac{M_{\rm disk}(<R_0)}{M_{\rm tot}(<R_0)} \approx \frac{7}{4}\,\frac{|\mathbf{J}_{\rm tot}(<R_0)|}{M_{\rm tot}(<R_0)\,R_0\,\sqrt{G\,M_{\rm tot}(<R_0)/R_0}} \tag{6}$$

where $\mathbf{J}_{\rm tot}(R_0) = \sum m_i\,(\bar{\mathbf{r}}_i - \bar{\mathbf{R}}_{\rm BH}) \times \bar{\mathbf{v}}_i$ (for all particles within R_0), and we require $M_{\rm disk} \le M_{\rm tot}$. Alternatively, this can be evaluated by de-composition of different component orbits in the system, following Angles-Alcazar's work. Both approaches to determine $M_{\rm disk}/M_{\rm tot}$ are available as options in the code.

Here R_0 is the BH kernel radius. This is set to enclose a desired number of neighbors (we recommend, for this accretion model especially, that the value be quite high or else these numbers get incredibly noisy – ~ 200 gas particles is a good choice), with a hard maximum radius (as much for numerical reasons as physical), both set in the parameterfile. The neighbor number (in gas) is the standard gas neighbor number times "BlackHoleNgbFactor". The maximum radius is "BlackHoleMaxAccretionRadius".

Alternatively, different assumptions about the disk structure at intermediate un-resolved radii might mean that the appropriate scaling behaves more similar to "gravito-turbulent scalings" in this limit. In those models, the inflow rate scales as $\dot{M} \sim 3\pi\,\alpha_{\rm gt}\,c_s^2\,\Sigma_{\rm gas}/\Omega$, where $\alpha_{\rm gt} \approx (2/3)\,\mathcal{M}_c = (2/3)\,\sigma_c/c_s$ (the compressive Mach number; see Gammie 2001 for the derivations and simulations and Hopkins & Christiansen 2013 for the conversion into these units). Assuming the disks are super-sonically turbulent, with Toomre $Q \sim 1$, and a "natural" mix of equal parts compressive and solenoidal turbulence (expected in the highly super-sonic regime), this becomes simply

$$\dot{M}_{\alpha\,{\rm disk}} \approx \left(\frac{M_{\rm disk}[<R_0]}{M_{\rm tot}[<R_0]}\right)^2 M_{\rm gas}[<R_0]\,\Omega[R_0] \tag{7}$$

where $\Omega^2(R_0) \equiv G\,M_{\rm tot}[<R_0]/R_0^3$. The disk fraction $M_{\rm disk}/M_{\rm tot}$ is estimated in the same manner as above, and the scaling here is qualitatively similar to that above, if the assumptions are the same, but in the regime where the gas is highly turbulent and disordered in the center, or the central regime is gas-dominated (as opposed to in a well-ordered disk of mixed gas and stars, in a primarily stellar potential), this is probably the more accurate (and robust) scaling. This can be selected by setting the BH_GRAVACCRETION parameter appropriate.

Even simpler, setting BH_GRAVACCRETION to an appropriate value, one simply assumes that the gas is accreted with a constant efficiency ϵ per dynamical time (set by "BlackHoleAccretionFactor" in the parameterfile), i.e. $\dot{M}_{\alpha\,{\rm disk}} = \epsilon\,M_{\rm tot}\,\Omega$.

Note the accretion rate will be (for any of the above) be multipled by "BlackHoleAccretionFactor". Default values above are in the code, so the "default" value for this parameter is unity.

5.1.3 *The Un-Resolved Limit: Bondi-Hoyle Model*

We also still include the option to determine the BHAR via the Bondi-Hoyle rate (following the original Springel & Hernquist implementation), by enabling BH_BONDI, you get:

$$\dot{M}_{\alpha\,{\rm disk}} = 4\pi\alpha\,G^2\,M_{\rm BH}^2\,\rho\,(c_s^2 + \beta\,|\mathbf{v}_{\rm BH} - \mathbf{v}_{\rm gas}|^2)^{-3/2} \tag{8}$$

[3] Eq. 6 is an exact solution for a system with an isotropic Hernquist (1990) bulge and thin Kuz'min disk (with R_0 small compared to the galactic bulge/disk scale lengths); the normalization will vary with mass profile but only weakly. This function is a rough approximation, but it has the advantage that it can be evaluated quickly in a kernel around each BH (for an arbitrary configuration of BHs), without any foreknowledge or reference to a "preferred" geometry or axis.

where here $\beta = 1$ by default (the standard formulation for fluid with bulk motion relative to the BH) if you use the option BH_BONDI= 0. If you set BH_BONDI= 1, then $\beta = 0$ (i.e. the bulk gas-BH motion term is ignored, giving much larger accretion rates). The parameter α is set by the run-time parameter "BlackHoleAccretionFactor" (this is the infamous number set to ~ 100 in the old springel-hernquist-dimatteo papers; which is plausible as a sub-grid extrapolation for unresolved density profiles and phase structure). If you enable BH_BONDI= 2, you get the Booth & Schaye 2009 model (used in all the subsequent papers by Schaye et al) which is identical, except $\alpha = 1$ if $\rho < \rho_{\rm crit}$ (where $\rho_{\rm crit}$ is the density threshold for star formation) and $\alpha = (\rho/\rho_{\rm crit})^{\gamma}$ for higher densities, where γ is now set by the parameter "BlackHoleAccretionFactor" (they chose $\gamma = 2$).

Note that a large number of studies have shown this is not a good approximation to periods of high BH accretion, since it assumes the gas has no angular momentum (when, in fact, understanding gas accretion onto BHs from large scales is primarily an angular momentum problem). Contrary to some claims in the literature, there is no Bondi-Hoyle formula that "accounts for" angular momentum – the actual scalings in the angular-momentum dominated regime resemble the gravito-turbulent and gravitational-torque accretion models discussed above, which have qualitatively different dimensional scalings (nearly independent of BH mass and sound speed, for example, where Bondi-Hoyle depends strongly on both of these). Still, the Bondi-Hoyle limit is potentially relevant for either situations (1) where the BH is accreting smoothly from a hot, hydrostatic, pressure-supported atmosphere, or (2) where a "seed" BH is moving through the ISM (on scales where it does not strongly influence the potential), closer to the regime the Bondi-Hoyle accretion theory was designed to represent.

5.2 Transport from the Accretion Disk to the BH

5.2.1 *Instantaneous*

If BH_ALPHADISK_ACCRETION is not enabled, accretion via the models above occurs instantly onto the hole, $\dot{M}_{\rm BH} = \dot{M}_{\alpha\,{\rm disk}}$.

5.2.2 *Sub-grid Accretion Disk Model*

If BH_ALPHADISK_ACCRETION is enabled, then once gas is captured into $R_{\alpha\,{\rm disk}}$, it must still be accreted into the BH. To model this, we use the standard formulation of an α-disk from Shakura & Sunyaev (1973). Since the outermost disk contains most of the mass and has the largest timescales, it is the "rate-limiter," so we can adopt the appropriate scalings in this regime (where gas pressure is stronger than radiation pressure) and obtain the inflow rate

$$\dot{M}_{\rm BH} = 2.45\,\alpha^{8/7}\,\frac{M_\odot}{\rm yr}\left(\frac{M_{\rm BH}}{10^8\,M_\odot}\right)^{-5/14} \times \left(\frac{M_{\alpha\,{\rm disk}}}{10^8\,M_\odot}\right)^{10/7}\left(\frac{R_{\alpha\,{\rm disk}}}{\rm pc}\right)^{-25/14} \tag{9}$$

We must choose α and $R_{\alpha\,{\rm disk}}$: we adopt the canonical value $\alpha \approx 0.1$, and $R_{\alpha\,{\rm disk}} = {\rm MIN}(0.2\,{\rm pc},\,\epsilon_{\rm BH})$ where $\epsilon_{\rm BH}$ is the kernel smoothing length around the BH. The latter choice is purely numerical, designed to match the inner radii into which the accretion rates were measured to calculate our accretion-rate estimator Eq. 5 (so that the two are consistent). However, these have very little effect on our results: they are individually degenerate, and their choice only regulates the "delay time" of accretion. These choices imply a disk depletion time $M_{\alpha\,{\rm disk}}/\dot{M}_{\rm BH}$ – or effective viscous transport time to the BH, of $\sim 10^7 - 10^8$ yr. Since this is still much less than the Hubble time or e.g. galaxy merger/starburst timescales, it does not substantially alter the character of BH evolution. Note that this final scaling is approximately $\lambda_{\rm Edd} = \dot{M}_{\rm BH}/\dot{M}_{\rm Edd} \approx (M_{\alpha\,{\rm disk}}/M_{\rm BH})^{1.4}$.

5.2.3 Eddington Limit

We can cap $\dot{M}_{\rm BH}$ at a multiple ψ of the Eddington limit:

$$\dot{M}_{\rm BH,\,Edd} \approx 2.38\,\frac{M_{\odot}}{\rm yr}\,\left(\frac{M_{\rm BH}}{10^{8}\,M_{\odot}}\right)\left(\frac{\epsilon_r}{0.1}\right)^{-1} \tag{10}$$

with radiative efficiency ϵ_r, discussed below. Note that there is an important physical distinction here: although BH growth may be strictly limited at Eddington, accretion into the outer accretion disk is not. In principle, the disk mass can build up and sustain longer-term fueling during intense galactic fueling episodes; of course, *resolved* feedback may self-regulate the accretion into the outer disk at something like an "effective" Eddington limit.

In each timestep Δt, then, the BH grows by a mass $\Delta M_{\rm BH} = (1-\epsilon_r)\,\dot{M}_{\rm BH}\,\Delta t$ (this properly accounts for loss of mass by radiation).[4]

The efficiency ϵ_r is set in the parameterfile with "BlackHoleRadiativeEfficiency" ("default" $= 0.1$). The factor ψ is set by "BlackHoleEddingtonFactor" (default $= 1$). To remove the eddington limit, simple set this to some arbitrarily large value.

5.3 Variability on Un-Resolved Timescales

AGN exhibit variability on very small timescales, corresponding to internal variability in e.g. the un-resolved accretion disk. This is "smoothed over" by our finite resolution; however, if you enable BH_SUBGRIDBHVARIABILITY, we can crudely approximate it by including an explicit power-spectrum of $\dot{M}_{\rm BH}$ fluctuations, integrated from frequencies of infinity down to $1/\Delta t_i$ where Δt_i is the simulation timestep (typically $\sim 100-1000\,$yr in the galaxy centers). We do this following Hopkins & Quataert (2011): assuming fluctuations in $\ln(\dot{M}_{\rm BH})$ follow a Gaussian random walk with equal power per logarithmic time interval from $t_{\rm min}$ (the orbital time at the innermost stable circular orbit for a non-rotating BH) to $t_{\rm max}$ (the dynamical time at the resolved R_0).

6 BLACK HOLE FEEDBACK

Knowing the BH accretion rate $\dot{M}_{\rm BH}$, we assign it the intrinsic, bolometric luminosity $L \equiv L_{\rm bol} = \epsilon_r\,\dot{M}_{\rm BH}\,c^2$, where ϵ_r is the radiative efficiency. We can vary ϵ_r in the code or make it a function of luminosity, but to be conservative we will reference our discussion to the canonical value ≈ 0.1.

6.1 Simplified Sub-Grid Thermal Feedback

If we enable BH_THERMALFEEDBACK, the simplest approach follows Springel & Hernquist 2003, and injects energy from the AGN in a pure thermal energy "dump" into the surrounding gas. Given the accretion rate and corresponding bolometric luminosity $L_{\rm bol}$ above, a fraction $\dot{E} = \epsilon_{\rm fb}\,L_{\rm bol}$ of the energy is coupled as purely thermal energy, distributed among the gas particles within the kernel of the BH (the same ones that determine the BH accretion rate) in a kernel-weighted fashion. The parameter $\epsilon_{\rm fb}$ is set by the run-time parameter BlackHoleFeedbackFactor in the parameterfile.

[4] We add a couple of additional timestep restrictions to the BH particles, to ensure they do not evolve on very large timesteps. This includes preventing them from any timestep longer than a physical 10^5 yr, or any single timestep in which they would grow $> 0.1\%$ of their mass.

This is intentionally a simplified, parameterized model intended as a sub-grid treatment, it is not intended to represent any specific AGN feedback mechanism or physics.

6.2 Accretion Disk (Broad Absorption Line) Winds (Mechanical AGN Feedback)

It appears that nearly all AGN have associated winds, albeit with a wide range of velocities $\sim 500-30000\,{\rm km\,s^{-1}}$. Although the origins and detailed dynamics of accretion-disk winds are uncertain, by the time they reach the large resolved scales of the simulation, they are primarily hydro-dynamic, and their basic properties are summarized by two parameters: a mass loading $\beta \equiv \dot{M}_{\rm wind}/\dot{M}_{\rm BH}$ and velocity $v_{\rm wind}$. This completely defines the time-dependent mass, momentum, and energy flux, which can be continuously "injected" into the gas surrounding the BH (with the assumption that the outflow is shocked, so we use the outflow velocity plus relative gas-BH velocity, together with momentum and energy conservation in the shock, to determine the coupled momentum and energy).

Note that this is very similar to the already-included treatments of stellar mechanical feedback in the code. One difference is that with stellar feedback, these parameters are determined from well-constrained stellar evolution models. Here the inputs are much less certain. But observations and theoretical models suggest values of order $\beta \sim 1$, $v_{\rm wind} \sim 10^4\,{\rm km\,s^{-1}}$. Since accretion disk winds are believed to be line-driven, the available momentum flux is $\dot{p} \sim L/c$ (although this can increase if there is an un-resolved energy-conserving phase of shocked wind-bubble expansion), thus the energy and momentum-loading of the winds are

$$\eta_P \equiv \frac{\dot{M}_{\rm wind}\,v_{\rm wind}}{L/c} = \beta\left(\frac{v_{\rm wind}}{\epsilon_r\,c}\right) \approx \beta\left(\frac{v_{\rm wind}}{30,000\,{\rm km\,s^{-1}}}\right) \tag{11}$$

$$\eta_E \equiv \frac{\dot{M}_{\rm wind}\,v_{\rm wind}^2}{2L} = \frac{\epsilon_r}{2}\,\frac{\eta_P^2}{\beta} \approx 0.05\,\beta\left(\frac{v_{\rm wind}}{30,000\,{\rm km\,s^{-1}}}\right)^2 \tag{12}$$

Note for $\eta_P = \beta = 1$, we recover the canonical $\eta_E \approx 0.05$ adopted in previous simulations with purely thermal AGN feedback (e.g. Di Matteo et al. 2005; Hopkins et al. 2005).

In our "standard" case, we take the winds to be isotropic, so the per-particle weight which determines the fraction of the wind "seen" is just proportional to the particle covering factor $\propto h_i^2$. Observations indicate something more like an equatorial wind, albeit with a broad opening angle of $\sim 30-45^{\rm deg}$. This is not so different from an isotropic wind, given the uncertainties in the mass loading, but it will generally be somewhat more efficient (since, to the extent that the accretion disk is aligned with the galaxy, this preferentially couples the wind in-plane). If we want to include this, we use our existing calculation of $\mathbf{J}_{\rm tot}$ in Eq. 6 (the net angular momentum vector of the nuclear gas at the smallest resolved scale) to determine the corresponding disk plane, assume the accretion disk is (on average) aligned, and then weight the wind mass-loading for each gas particle by $\cos^2(\theta)$ (where θ is the angle of the particle out-of-plane), appropriate normalized to the same total.

Because we set the wind momentum and energy by hand "at coupling," we do not include the "boost factor" that the stellar winds and SNe use. We could, of course, fairly easily, or we could fold a constant effective boost into the parameter choices for this module. The parameters are set in the parameterfile: "BAL_v_outflow" sets $v_{\rm wind}$, and "BAL_f_accretion" ($\equiv f$) is the fraction, for some total mass accreted into the disk, which ends up on the BH, i.e. $f = \beta/(1+\beta)$.

We note that this model can represent any local mechanical AGN feedback. There are, however, a few distinct numerical methods to treat these outflows.

6.2.1 Particle "Kicking"

Numerically the simplest approach, if one wishes to ensure a given velocity is reached, one can probabilistically "kick" gas particles (enabling BH_BAL_KICK). In this case the velocity change $\Delta\mathbf{v} = v_{\rm wind}\,\hat{\mathbf{r}}$ is fixed, where $\hat{\mathbf{r}}$ represents kicks directed radially away from the BH – one can choose instead to orient the kicks in a collimated way with the appropriate code options. The probability of a kick can be weighted by angle, in principle, to represent anisotropic kicks. By default it is isotropic, and the probability of kicking a particle is scaled so that the desired mass-loading is achieved, on average (so if particle masses are larger, but all else is equal, there will be fewer "kicks"). This is algorithmically similar to the subgrid wind galactic wind models in the code (and the probabilistic step is similar to how gas particles are converted into star particles). A "kicked" particle immediately has its velocity incremented, nothing else. The algorithmic implementation of these winds was developed in Anglés-Alcázar et al. (2016).

6.2.2 Continuous Acceleration

A somewhat more detailed wind model is enabled by BH_BAL_WINDS. In this case mass, energy, and momentum are *continuously* injected into the gas surrounding the BH (within its neighbor kernel), at a rate appropriate to the desired accretion rate and mass-loading. In other words, in a timestep Δt, a total wind mass $\Delta M_{\rm wind} = \dot{M}_{\rm wind}\,\Delta t$ is injected into the neighbor particles, along with the corresponding kinetic luminosity and momentum. The algorithmic approach to this is the same as used for continuous stellar mass loss in the FIRE simulations, described in detail in Hopkins et al. (2017).

This algorithm is more accurate than a stochastic "kicking" approach, if the resolution is sufficiently high. However, at low resolution, it can have some problematic aspects – if the particle masses are too large, the injected momentum/energy per timestep will be extremely small, and can be radiated away very quickly (over-cooling) or dissipated by numerical noise/diffusivity rather than "building up" correctly if it were resolved.

6.2.3 Virtual (Wind) Particle Injection

A more accurate treatment of winds was recently developed by Paul Torrey, enabled by BH_WIND_SPAWN. In this case, every timestep, the BH generates a large number of "wind particles." For example, in a timestep Δt, one has a total wind mass $\dot{M}_{\rm wind}\,\Delta t$ coming from the accretion disk, which is then broken into a number N of particles. These particles are assigned momenta and energy according to the desired wind properties. If desired, they can trivially be "loaded" with other quantities (cosmic rays, magnetic fields, etc). The particles are then launched from the BH particle with the desired velocity and orientations (by default, they are launched isotropically, but it is trivial to give them a preferred orientation).

This allows the greatest freedom in specifying the mass/momentum/energy loading, wind geometry, and loading of other quantities (e.g. magnetic fields). It also avoids pathologies inherent to the kicking or continuous acceleration cases in irregular grids. Let's say the BH has cleared a "channel" in the polar direction – there is almost no gas mass in that direction. That means there are unlikely to be any nearby gas particles in that direction. The previous methods would therefore not "see" anything in that direction and put all their "work" into the other directions (incorrectly) – it can be hard to capture "venting" of hot, extremely fast winds, with those approaches. This guarantees those limits behave correctly.

The trade-off for this gain in accuracy and flexibility is spawning a potentially very large number of low-mass, fast-moving particles. These require small timesteps. Their low mass would be problematic if they were mixed with "normal" particles under some circumstances (it is plausible, if the BH is accreting at a low rate, that the wind particles are many, many orders of magnitude smaller in mass than the "normal" gas elements in the simulation). To address this, as soon as a "virtual" wind particle sees a "normal" particle within its kernel, and is moving into its volume (approaching towards the forward-facing face of the volume domain represented by the "normal" particle), and should shock against it (as calculated by the Reimann solver), it is merged entirely into that particle (transferring all the appropriate quantities and updated for the shock).

6.3 Radiative Feedback

6.3.1 Compton Heating/Cooling

Enabling BH_COMPTON_HEATING, the radiation field of the BH will also Compton heat/cool gas in its vicinity. As discussed in Sazonov et al. (2004, 2005), this effect is nearly independent of obscuration: Compton heating is entirely dominated by photons with energies $\gg 10\,{\rm keV}$ (for which we can usually safely ignore obscuration) and Compton cooling by the bolometric luminosity in lower-energy photons (re-distributed, but not, in integral, altered by obscuration). As such even Compton-thick columns result in factor < 2 changes in the heating/cooling rates. We therefore neglect obscuration and assume the radiation field is isotropic, so that the X-ray/bolometric flux from the AGN on all particles is given by $F_X = L_X/4\pi\,r^2$, with Compton temperature $\approx 2\times10^7\,{\rm K}$ as calculated in Sazonov et al. (2004) for a broad range of observed QSO SED shapes.[5] In the cooling function, we add the appropriate Compton heating and cooling terms.[6] Although Compton cooling depends explicitly on the free electron fraction, for the photon energies dominating heating (much greater than the ionization energy of hydrogen), we can safely approximate Compton heating of bound electrons as identical to free electrons (see e.g. Basko et al. 1974; Sunyaev & Churazov 1996).

Finally, as shown in Faucher-Giguere & Quataert (2012), some care is needed at the highest temperatures: if the timescale for Coulomb collisions to transfer energy from ions to electrons is longer than the Compton or free-free cooling time of the electrons, this is the rate-limiting process and a two-temperature plasma develops. We therefore do not allow the Compton+free-free cooling rate to exceed the Coulomb energy transfer rate between ions and electrons calculated for an ion temperature T in the limit where the electrons are efficiently cooling $T_e \ll T$ (see Spitzer 1962; Narayan & Yi 1995). It is important to note that AGN wind-shocked electrons are generally non-relativistic: either immediately post-shock (where most energy is in protons, with electron temperature $T_e \sim T_p\,(m_e/m_p) \sim 1.3\times10^7\,{\rm K}\,(v_{\rm shock}/30,000\,{\rm km\,s^{-1}})^2$), or in later stages when competition between Compton cooling and Coulomb heating regulates the temperature.

[5] We propagate this flux through the gravity tree, since it follows an inverse-square law when we can neglect obscuration. This makes it trivial to apply the appropriate flux to arbitrary particle numbers, geometries, and numbers of black holes.

[6] As is standard, cooling is solved implicitly within this function in the regime where the heating/cooling times are short compared to the particle timesteps.

6.3.2 Photo-Ionization

Enabling BH_HII_HEATING, we can treat photo-ionization by the BH algorithmically identically to that from stars. Specifically, we calculate the rate of production of ionizing photons from the empirically-determined QSO spectra in Hopkins et al. (2007), $\dot{N}_{\rm ion} \approx 5.5\times10^{54}\,{\rm s}^{-1}\,(L_{\rm bol}/10^{45}\,{\rm erg\,s^{-1}})$. We then adopt a local Stromgren approximation: moving radially outwards from the BH, we check each gas particle; if it is not already fully ionized, we calculate the number of ionizing photons per unit time required to fully ionize it. If that is available, we "consume" those photons from the BH and move on. We repeat until we encounter a particle requiring more photons $\Delta\dot{N}_{\rm ion}$ than available $\dot{N}_{\rm ion}^{\rm remain}$, which is then ionized or not with ionization probability $=\dot{N}_{\rm ion}^{\rm remain}/\Delta\dot{N}_{\rm ion}$ (ensuring the correct number of photons is used) and the chain is ended.

6.3.3 Radiation Pressure

Enabling BH_PHOTONMOMENTUM, we can also include the radiation pressure from black holes, in a manner similar to that from stars. To simplify and be conservative, we assume that (whether or not it is resolved), most of the optical/UV light from the AGN is singly-scattered in the vicinity of the BH, then downgraded to IR photons. This imparts a momentum flux $\dot{P}\approx L/c$ locally, which we coupled directly as a continuous momentum flux to the gas in the smoothing kernel of the BH (directed radially away from the BH). The re-radiated flux is propagated as a long-range, infrared radiation flux $\mathbf{F}$ in the same manner as the IR component of the stellar luminosity, where it can impart an acceleration on gas particles of $\mathbf{a}=\kappa_{\rm IR}\,\mathbf{F}/c$ (so this is again identical to the single-scattering radiation-pressure, but for the IR component rather than the UV).

It is possible that either more UV photons escape the central region, or that multiple-scattering effects enhance the coupling in the optically thick region. Both of these effects would increase the strength of the radiation pressure, but we neglect both of these terms for now.

However, an important remaining ambiguity is the directional dependence of the flux, which can be highly non-isotropic. We take this from the fitting functions Nathan Roth developed from the full radiative transfer calculations. The force is always directed radially from the BH, aligned with the flux, and that flux is azimuthally symmetric about the AGN disk angular momentum axis. But, we include a dependence on the polar angle θ, which is defined with respect to the angular momentum vector $\boldsymbol{\omega}=|\omega|\,\hat{z}$, $\cos\theta=|\hat{r}\cdot\hat{z}|$, for a particle at position $\mathbf{r}$ with respect to the BH at the origin (the absolute value is because the top/bottom halves of the disk are symmetric, so we only consider $0<\theta<\pi/2$).

We then weight the initial single-scattering flux in each solid angle element ${\rm d}\Omega$ by θ according to

$$\frac{{\rm d}F}{{\rm d}\Omega}=F_0\left[1-\frac{1+\alpha\,\exp\left(-\beta\,\pi/2\right)}{1+\alpha\,\exp\left[-\beta\,(\pi/2-\theta)\right]}\right] \tag{13}$$

$$\alpha=\frac{8.494}{1.173+h/R}\qquad \beta=\frac{64.425}{2.540+h/R} \tag{14}$$

where the normalization F_0 is chosen to that the integrated flux over the area of a sphere is $L/(4\pi\,r^2)$. We implement this in SPH by assigning each particle j in the kernel about the BH a fraction of the initial luminosity

$$f_j\propto h_j^2\,\frac{{\rm d}F}{{\rm d}\Omega}(\theta=\theta_j) \tag{15}$$

(the h_j^2 term approximates the fraction of solid angle covered by the particle). We do this in two loops so that we normalize the sum, such that the total momentum coupled to all particles in the timestep is *exactly*

$$\Delta P_{\rm tot}=\sum_j \Delta P_j=\sum_j \frac{L}{c}\,\Delta t\,f_j=\frac{L}{c}\,\Delta t \tag{16}$$

For the long-range force, we take the flux

$$\mathbf{F}(\mathbf{r})=\mathbf{F}_{\rm IR}\approx\hat{r}\,\frac{L_{\rm bol}}{4\pi\,|\mathbf{r}|^2}\,\exp\left[G(|\cos(\theta)|)\right] \tag{17}$$

$$G(x)\equiv 3.318+\frac{(-14.78+22.84\,h/R)}{(\exp(y)-1)}\,(\exp(y-x^2\,y)-1)$$

$$y^{-1}\equiv 0.357-10.839\,(h/R)+142.640\,(h/R)^2 -713.928\,(h/R)^3+1315.132\,(h/R)^4$$

(actually, the $1/|\mathbf{r}|^2$ is softened if the particles are within a smoothing length, same as the gravity, to prevent a numerically problematic divergence).

Now we need to determine the angular momentum vector $\boldsymbol{\omega}$, and h/R. We take $\boldsymbol{\omega}=\boldsymbol{\omega}(|\mathbf{r}|_{\rm gas}<h_{\rm BH})$ – the net angular momentum vector of the gas inside the kernel around the BH. If this is driving accretion, this is a plausible "best guess." We could in principle make it random (precessing?), or simply fix it to some arbitrary angle for all time. We also need to estimate h/R: we again do this from the resolved h/R (essentially the "best-fit" Gaussian h/R, assuming the disk plane determined by the already-calculated net angular momentum axis) in the kernel. But again we could simply fix it if so desired, as its not obvious how this might run with scale.

6.4 Relativistic Jets

Not explicitly included yet. However, non-relativistic, hydrodynamic jets can be trivially implemented by modifying the angular dependence and velocities of the wind particles launched by the BH_WIND_SPAWN flag. With MHD active, these can also trivially be given magnetic energy, or with cosmic rays active, they can carry cosmic ray energy. These represent the recommended approaches for now.

However, full relativistic MHD is not incorporated yet into production versions of GIZMO, which limits the explicit treatment of relativistic jets.

REFERENCES

Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2016, MNRAS, in press, arXiv:1603.08007
Basko, M. M., Sunyaev, R. A., & Titarchuk, L. G. 1974, A&A, 31, 249
Di Matteo, T., Springel, V., & Hernquist, L. 2005, Nature, 433, 604
Faucher-Giguere, C.-A., & Quataert, E. 2012, MNRAS, in press, arXiv:1204.2547
Gammie, C. F. 2001, ApJ, 553, 174
Hernquist, L. 1990, ApJ, 356, 359
Hopkins, P. F., & Christiansen, J. L. 2013, ApJ, 776, 48
Hopkins, P. F., Hernquist, L., Martini, P., Cox, T. J., Robertson, B., Di Matteo, T., & Springel, V. 2005, ApJL, 625, L71
Hopkins, P. F., & Quataert, E. 2010, MNRAS, 407, 1529
—. 2011, MNRAS, 415, 1027
Hopkins, P. F., Richards, G. T., & Hernquist, L. 2007, ApJ, 654, 731
Hopkins, P. F., Wetzel, A., Keres, D., Faucher-Giguere, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C. C., & El-Badry, K. 2017, MNRAS, in press, arXiv:1707.07010

Narayan, R., & Yi, I. 1995, ApJ, 444, 231
Sazonov, S. Y., Ostriker, J. P., Ciotti, L., & Sunyaev, R. A. 2005, MNRAS, 358, 168
Sazonov, S. Y., Ostriker, J. P., & Sunyaev, R. A. 2004, MNRAS, 347, 144
Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337
Spitzer, L. 1962, Physics of Fully Ionized Gases; New York: Interscience (New York: Interscience)
Springel, V., Di Matteo, T., & Hernquist, L. 2005, MNRAS, 361, 776
Sunyaev, R. A., & Churazov, E. M. 1996, Astronomy Letters, 22, 648