Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 26 March 2020 (MN LaTeX style file v2.2)

Black Holes in GIZMO

Philip F. Hopkins*[1], Daniel Angles-Alcazar, Paul Torrey, Xiangcheng Ma, Claude-André Faucher-Giguère, Mike Grudic, David Guszejnov

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

–, 2013

ABSTRACT

This is a description of the Black Hole (BH) and related 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 (No "Seeding" Flags)

Often one sets up BHs in the ICs, so there is no formation "flag" needed. If set up in the ICs, the *particle mass* is read from the IC file. However, the "BH mass" and "α-disk mass" (if that physics is enabled) are *not*, by default. These (by default) will instead be read from the parameterfile from the params SeedBlackHoleMass and SeedAlphaDiskMass, respectively. If you want to read them from the IC file (for example, to set different values for each BH), you need to do this manually.

2.2 "On-the-fly" Seeding from Star Formation (BH_SEED_FROM_LOCALGAS)

If BH_SEED_FROM_LOCALGAS is enabled, then a gas particle is flagged as being "turned into" a star particle (which occurs according to the normal SF requirements, whatever those are for your simulation), this assigns 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:

$$\frac{\mathrm{d}P_{\rm seed}}{\mathrm{d}M_{*}} = \frac{1}{M_0}\left(1-\exp\left[-\frac{\Sigma}{\Sigma_0}\right]\right)\exp\left(-\frac{Z}{Z_0}\right) \tag{1}$$

(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. This will *only* occur at redshifts higher than the value of SeedBlackHoleMinRedshift set in the parameters file.

The default parameters in the code are $Z_0 = 0.01\,Z_{\odot}$, $\Sigma_0 = 1\,{\rm g\,cm^{-2}}$, with M_0 (the normalization – e.g. the stellar mass which must be formed with $\Sigma \gg \Sigma_0$ and $Z \ll Z_0$ to get, on average, one BH) set by the run-time parameterfile parameter SeedBlackHolePerUnitMass.

* E-mail:phopkins@caltech.edu

2.3 "On-the-fly" Seeding from Halo Finding (BH_SEED_FROM_FOF)

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. If a halo already has a seed within it (e.g. it merged in) then no new seed will be placed, so essentially more seeds requires seeding in smaller halos. The user specifies both the seed mass and minimum halo/stellar mass of the groups which will get a seed. The run-time parameter MinFoFMassForNewSeed gives the minimum halo mass where a seed is placed. This will *only* occur at redshifts higher than the value of SeedBlackHoleMinRedshift set in the parameters file.

2.4 BH Seed Properties

When BH seeds are created or read-in (by any of the above) their properties are set by the run-time parameters: SeedBlackHoleMass sets the initial *BH* mass (not total particle mass, which is set by the IC or progenitor particle mass). SeedAlphaDiskMass sets an initial α-disk mass (if that physics is active). If SeedBlackHoleMassSigma is set to a non-zero value, then the initial BH mass is drawn probabilistically from a log-normal distribution with median of SeedBlackHoleMass and dispersion of SeedBlackHoleMassSigma (in dex).

3 BLACK HOLE DYNAMICS & CONSERVATION

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 (in principle) 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 protogalaxies in which they form. Similarly, the seed may be placed in a "badly wrong" location (well outside the center of the galaxy or nearest dense stellar structure) which would lead to it getting dynamically "kicked out" even in some cases where the dynamics are resolved (e.g. tidally ejected) – but realistically the expectation would be that the seed should have formed in the dense structure, so this shouldn't happen. There are a few different numerical methods to attempt to deal with this implemented in the code.

3.1 Dynamical Friction (BH_DYNFRICTION)

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

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

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

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

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

Setting BH_DYNFRICTION $=0$, $=1$, $=2$ uses all mass ($=0$), dark matter+stars ($=1$), or just stars ($=2$) to compute the forces. Setting >2 uses stars but multiplies the force by the value of BH_DYNFRICTION, to make the "friction" arbitrarily stronger.

3.2 Drag on the BH (BH_DRAG)

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 $=2$ 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). Note that this produces sinking of the BH with respect to the gas, but if the gas is flowing strongly (e.g. from outflows), then you can end up with the BH getting "dragged with" the outflow.

3.3 Moving the BH to Potential Minimum (BH_REPOSITION_ON_POTMIN)

BH_REPOSITION_ON_POTMIN always moves the BH to the local potential minimum (within the kernel).

If BH_REPOSITION_ON_POTMIN is $=0$, the "traditional" variant of this is used. Among all neighbors b in the kernel H_a of BH a, the BH picks whichever has the minimum gravitational potential Φ_b, and simply "jumps" to that location ($\mathbf{x}_a \rightarrow \mathbf{x}_b$) at the end of the timestep. 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! And it is dis-continuous so BH "teleportation" often leads to odd accretion and kinematic behaviors. If particles fall in on highly eccentric orbits, or dense gas is blown out by winds, the BH can "bootstrap" itself out of the center.

If BH_REPOSITION_ON_POTMIN is $=1$, a more conservative formulation is used. To prevent some of the issues above, the BH only looks at star particle neighbors b for positioning, and only is allowed to jump to those where the relative velocity $|\mathbf{v}_b - \mathbf{v}_a|^2 < v_{\rm esc,\,BH}^2 + c_{s,\,\rm BH}^2$ (where $c_{s,\,\rm BH}$ is the sound speed of the gas in the vicinity of the BH, and $v_{\rm esc,\,BH}$ is the escape velocity from the BH at the distance $|\mathbf{x}_b - \mathbf{x}_a|$). This can lead to a lack of the BH "anchoring" if you use small seeds (so the escape velocity is low) or if the stars are on dispersion-supported orbits with large $|\mathbf{v}_b - \mathbf{v}_a|$ (even if the mean velocity at the potential minimum is small). The "jumps" of the BH are slightly smoothed, with $\mathbf{x}_a \rightarrow (\mathbf{x}_a + \mathbf{x}_b)/2$, to minimize very large jumps.

If BH_REPOSITION_ON_POTMIN is $=2$, the most stable formulation is used. The BH "looks at" neighbors of all types except gas. The target neighbor b "minimum" is no longer the minimum of Φ_b, but the minimum energy or globally-most-bound neighbor (as seen from the frame of the BH), namely the neighbor with the minimum value of $\Psi_b \equiv W_{ab}^{-1}\,(\Phi_b + (1/2)\,|\mathbf{v}_b - \mathbf{v}_a|^2)$ (where $W_{ab} \equiv 1 + |\mathbf{x}_b - \mathbf{x}_a|^2\,(h_a^{-2} + \epsilon_a^{-2})$, with $h_a \sim H_a/3$ the mean inter-particle separation and $\epsilon_a \approx 3$ times the BH force softening) – this prevents jumping to neighbors moving with very large velocities (without a criterion that depends on the BH mass or escape velocity) or very large separations. The BH also no longer "jumps" but moves *smoothly*, with $\mathbf{x}_a \rightarrow \mathbf{x}_a + (\mathbf{x}_b - \mathbf{x}_a)\,(1 - \exp{[-v_s\,\Delta t_a/|\mathbf{x}_b - \mathbf{x}_a|]}) \approx \mathbf{x}_a + \hat{\mathbf{x}}_{ba}\,v_s\,\Delta t_a$ where Δt_a is the timestep, and $v_s = {\rm MAX}(v_{ff},\,c_0)$ with $c_0 = 10\,{\rm km\,s^{-1}}$ (a typical sound speed, set as a minimum) and $v_{ff}^2 = -2\,\Psi_b$ (the free-fall velocity into the potential minimum). Even with this, the BH can if it has been strongly "kicked" over-shoot and be ejected, so this option also enables a strong version of the dynamical-friction type drag, with $\mathbf{v}_a \rightarrow \mathbf{v}_a + \delta\mathbf{v}\,(1 - \exp{[-a_s\,\Delta t_a/|\delta\mathbf{v}|]})$, where $\delta\mathbf{v}$ is the mean velocity of the particles in the kernel as defined above, and $a_s = {\rm MIN}[|\delta\mathbf{v}|/\tau,\,{\rm MAX}[G M_a^k/H_a^2,\,-2\,\Psi_b/H_a]]$ (where M_a^k is the total mass in the kernel, so the latter two terms represent two estimates of the local free-fall acceleration on the scale of the kernel H_a, and the decay time of the velocity is limited by the $|\delta\mathbf{v}|/\tau$ term to not be faster than $\tau \equiv 1\,$Myr).

In experiments, of all the dynamical "anchoring" mechanisms above, the BH_REPOSITION_ON_POTMIN=2 option gives by far the most stable BH "anchoring," if that is the goal.

3.4 Conservation (BH_FOLLOW_ACCRETED_COM and BH_FOLLOW_ACCRETED_MOMENTUM)

In general, in order to maintain exact conservation not just of mass but of the center-of-mass and linear momentum of the simulation, whenever the BH discretely "swallows" any mass, the BH position is updated exactly conserving the center-of-mass of the BH and swallowed mass element (thus, the "swallow" operation does not change the center-of-mass of the global system), this is controlled by the flag BH_FOLLOW_ACCRETED_COM.

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

The same is done for the total momentum of the system (BH_FOLLOW_ACCRETED_MOMENTUM).

However, when BH_REPOSITION_ON_POTMIN is defined, we are not attempting to represent the "true" dynamics of the sink in any case, so these conservation terms are automatically disabled (so the BH does not "wander" from the desired path as it accretes).

3.5 Tracking BH Angular Momentum (BH_FOLLOW_ACCRETED_ANGMOM)

In addition, whenever the BH accretes, we can exactly calculate the accreted angular momentum, and track this as well. The variable "BH_Specific_AngMom" in the code retains the corresponding specific angular momentum per unit mass of the BH+α-disk system (we follow specific angular momentum because this then is not updated when the system loses mass in radiation or outflows). This is steadily integrated, and defines a "polar direction" of the BH+disk accreting system that can be used to then define the direction for collimated feedback (e.g. jets).

There are two versions of this tracking. If the "GRAVCAPTURE" modules for direct/resolved gravitational capture/accretion are active, or sink particle dynamics for stars, then we adopt BH_FOLLOW_ACCRETED_ANGMOM=0, where the BH angular momentum is updated on discrete particle-swallow events since these are the actual accretion events. However, in situations where sub-grid accretion models are used (see below), then we set BH_FOLLOW_ACCRETED_ANGMOM=1 where the angular momentum accreted is continuous along with the accreted mass $\dot{M}$ every timestep, assuming the specific angular momentum accreted follows the specific angular momentum of the gas in the kernel (since the sub-grid accretion rate is smoothed/calculated over neighbors, the angular momentum of a single particle is not representative of the angular momentum of accreted gas).

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.

These can be disabled entirely by setting BH_DEBUG_DISABLE_MERGERS.

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.

Generically, the BHAR (and many other parameters here) will be calculated from gas within a local kernel around the BH. This is set to enclose a desired number of neighbors (the value should be rather high or else some of the below can get very noisy – $\sim 100-200$ gas particles is a good choice), with a hard maximum radius, both set in the parameterfile. The neighbor number (in gas) is the standard gas neighbor number times BlackHoleNgbFactor. The maximum radius is BlackHoleMaxAccretionRadius (note this means no particles are used even for estimating smooth quantities, like the density in the vicinity of the BH, beyond this radius, so it should be large as it is really a numerical nuisance parameter).

The accretion rate will be (for any of the below) be multiplied by $\epsilon_{\rm acc} \equiv$BlackHoleAccretionFactor. Default values above are in the code, so the "default" value for this parameter is unity.

5.1 Capture of Gas Into the Accretion Disk

5.1.1 The Resolved Limit: Direct Capture (BH_GRAVCAPTURE_GAS)

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 apocentric 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. Here $R_{\alpha\,\rm disk}$ is assumed to be the accretion radius set by the user in the parameterfile (BlackHoleMaxAccretionRadius). Note that we can do this for only gas particles (enabling BH_GRAVCAPTURE_GAS), only non-gas particles (enabling BH_GRAVCAPTURE_NONGAS), or both.

Also note that there has been considerable work on improving the true "sink particle" limit for e.g. simulations of the stellar initial mass function by M. Grudic and D. Guszejnov (discussed in different notes).

5.1.2 The Un-Resolved Limit: Sub-Grid, Torques-Driven Accretion Models (BH_GRAVACCRETION)

The resolution needed to meaningfully apply the above prescription is only true in nuclear-scale simulations. In many cases, we cannot resolve $R_{\alpha\,\rm disk}$, so must adopt a "sub-grid" accretion prescription.[2] Enabling BH_GRAVACCRETION, we therefore have also implemented the model of Hopkins & Quataert (2011), and variants thereof which was designed to reproduce the accretion rate resolved in annuli ranging from large scales in a galaxy down to the accretion disk. They argued that the dominant mechanism of angular momentum transfer on all of these scales is torques due to gravitational

[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 the kernel radius H_a 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 H_a by a kernel function). 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. Enabling BH_ACCRETE_NEARESTFIRST biases the accretion weight to entirely put the weight onto the single nearest gas element to the BH, every timestep.

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 some other accretion estimators commonly used (including variant "Bondi-Hoyle" accretion rates).

We caution that the robust scalings derived and tested in Hopkins & Quataert (2011, 2010) are fundamentally *local* – i.e. they tell you, for a given structure of the potential and e.g. gas density at some radius r, how much gas should flow through that given radial annulus (how rapidly it will lose angular momentum in non-linear processes such as shocks, etc.). The full expressions are more complex, but in essence, the mass flux at radius R scales as $\sim |a|\,\Sigma_{\rm gas}|\,R^2\,\Omega(R)$, where $|a|$ is the fractional amplitude of the local asymmetry in the potential. Unfortunately, generalizing this to a *galaxy-scale* model requires making assumptions about e.g. the relative fraction of gas consumed by star formation (so how SF scales with local properties in the nucleus) and how the gas density profile scales with radius on un-resolved scales. This means that even in the context of these specific torque models, the extrapolation from $\sim 0.1-1\,$kpc to $\ll 0.1\,$pc is non-unique. So by setting `BH_GRAVACCRETION` to different values, you can experiment with different choices:

- `BH_GRAVACCRETION` = 0: Default scaling from Hopkins & Quataert (2011):[3]

$$\dot{M}_{\alpha\,\rm disk} \sim \epsilon_{\rm acc}\,\frac{M_{\rm gas}(<R)}{2\times10^8\,{\rm yr}}\,\frac{f_d^{1/2}\,M_{\rm BH,8}^{1/6}\,R_{100}^{-3/2}}{(f_{\rm gas}+0.3\,f_d\,M_{\rm d,9}^{-1/3})} \tag{5}$$

where $M_{\rm BH,8}\equiv M_{\rm BH+\alpha\,disk}/10^8\,M_{\odot}$, $M_{\rm d,9}\equiv M_{\rm disk}/10^9\,M_{\odot}$, $f_d\equiv M_{\rm disk}(<R)/M_{\rm total}(<R)$, $f_{\rm gas}\equiv M_{\rm gas}(<R)/M_{\rm total}(<R)$, $R_{100}\equiv R/100\,$pc, all evaluated inside a radius R.[4] Here R is set to the maximum kernel neighbor search radius H_a, which means it is changing dynamically (often very rapidly) during the simulation.

- `BH_GRAVACCRETION` = 1: This uses the same scalings as `BH_GRAVACCRETION` = 0, but (1) replaces the bulge-to-disk estimator as noted above, and (2) evaluates Eq. 5 at a *fixed* physical radius $R=R_0$, set to ≈ 3 times the full extent of the force softening kernel (~ 10 times the Plummer-equivalent softening). This can be more stable, since the radius is not "jumping" around as the gas density increases/decreases (using $R=H_a$, R always becomes larger as the gas density around the BH decreases, which changes the meaning of this estimator), and the "efficiency per free-fall" can be well-defined. It also ensures R stays within a range where the scaling is well-calibrated ($\sim 10-100\,$pc). Since H_a does not equal R, quantities like $M_{\rm gas}(<R)$ are extrapolated using the usual kernel-estimators.

[3] Obtained by assuming a power-law $\Sigma_{\rm gas}\propto R^{-1.5}$ and assuming a Kennicutt-type relation $\dot{\Sigma}_*\propto\Sigma_{\rm gas}^{1.5}$ at all annuli, and that $|a(R)|\propto f_{\rm disk}(<R)$, an $m=1$ mode inside of the BH radius of influence and $m=2$ outside.

[4] Evaluating the "disk" mass inside R is non-trivial. Two options are present in-code. The default uses the kinematic de-composition described in Angles-Alcazar et al. 2017: the "bulge" fraction is twice the fraction of mass with $\mathbf{j}\cdot\mathbf{J}_{\rm tot}<0$ (i.e. specific angular momentum vector $\mathbf{j}$ counter-rotating with respect to the *total* angular momentum $\mathbf{J}_{\rm tot}$ in the kernel). Alternatively for `BH_GRAVACCRETION` = 1, the proposed scaling in Hopkins & Quataert 2011 is used based on de-compositing the system into a Hernquist 1990 profile bulge and thin Kuz'min disk, giving a disk fraction $=7\,|\mathbf{J}_{\rm tot}|/4\,G^{1/2}\,M_{\rm total}^{3/2}\,R^{1/2}$.

- `BH_GRAVACCRETION` = 2: This makes a simpler assumption that the efficiency per free-fall time (essentially the magnitude of $|a|$ in the Hopkins & Quataert 2011 formulation) at a fixed physical radius $R=R_0$ (defined above) is constant, we obtain:

$$\dot{M}_{\alpha\,\rm disk}\sim\epsilon_{\rm acc}\,M_{\rm gas}(<R_0)\,\left(\frac{G\,M_{\rm total}(<R_0)}{R_0^3}\right)^{1/2} \tag{6}$$

(where we have subsumed all of the relevant constants into $\epsilon_{\rm acc}$, for simplicity).

- `BH_GRAVACCRETION` = 3: Making different assumptions about the disk structure at un-resolved radii means that the appropriate scaling behaves more like the classical "gravito-turbulent" scalings. 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:

$$\dot{M}_{\alpha\,\rm disk}\sim\epsilon_{\rm acc}\,f_d^2\,M_{\rm gas}(<R_0)\,\left(\frac{G\,M_{\rm total}(<R_0)}{R_0^3}\right)^{1/2} \tag{7}$$

- `BH_GRAVACCRETION` = 4: If we assume the background medium has uniform density and gas fraction outside of the minimum of either the BH radius of gravitational influence (R where $M_{\rm BH+\alpha\,disk}<(M_{\rm total}(<R)-M_{\rm BH+\alpha\,disk})$) or the Bondi radius ($G\,M_{\rm BH+\alpha\,disk}/c_s^2$) – i.e. outside of where the BH dominates the dynamics, with $\dot{M}\sim$constant inside this radius at a fixed efficiency per free-fall time – i.e. assume the BH accretes with $|a|$ constant inside of the radius where it dominates the dynamics (and SF losses can be neglected inside this radius) and gas is bound (and has small dynamical effect at larger radius), this gives:

$$\dot{M}_{\alpha\,\rm disk}\sim\epsilon_{\rm acc}\,\frac{4\pi\,G^2\,M_{\rm BH+\alpha\,disk}\,M_{\rm total}(<R)\,\rho_{\rm gas}}{(c_s^2+V_c^2)^{3/2}} \tag{8}$$

where $\rho_{\rm gas}$ and c_s are the gas density and sound speed evaluated in the vicinity of the BH, and $V_c^2\equiv G\,M_{\rm total}(<R)/R$. Note that this resembles a Bondi-Hoyle type model, but replaces the velocity with the gravitational velocity, and one of the BH mass terms with the total gravitating mass, appropriate for systems where the self-gravity of the material at larger radii dominates.

- `BH_GRAVACCRETION` = 5: This is a hybrid model where the gas circularization radius for the gas in the BH kernel is calculated assuming it conserves specific angular momentum: if this radius is greater than the BH radius of gravitational dominance, then we apply the gravitational torque estimator as in option (1) above. If the radius is smaller, it implies the gas is sufficiently pressure supported with little enough angular momentum that a Bondi-Hoyle accretion rate should apply, so we apply the Bondi-Hoyle scalings below.

- `BH_GRAVACCRETION` = 6: This is another hybrid model: assume that the gas has a (microphysical) Maxwellian velocity distribution with variance $\sigma_v^2\sim|\delta\mathbf{v}^2|/3+c_s^2$, and the bound gas forms a $\rho\propto r^{-1}$ profile, interior to which a fixed fraction of the gas accreted per dynamical time from the radius inside of which the BH dominates the dynamics (like model **(4)** above). This gives (approximately):

$$\dot{M}_{\alpha\,\rm disk}\sim\epsilon_{\rm acc}\,\frac{4\pi\,(M_{\rm total}/M_{\rm BH+\alpha})^{1/4}\,G^2\,M_{\rm BH+\alpha}\,M_{\rm total}(<R)\,\rho_{\rm gas}}{V_c^2\,(V_c+{\rm MIN}[\sigma_v^2/V_c^2\,,\,(M_{\rm total}/M_{\rm BH+\alpha})^{1/4}])^{3/2}} \tag{9}$$

where $M_{\rm total}$ and V_c are evaluated at R, and $\delta\mathbf{v} \equiv \langle \mathbf{v}_{\rm BH} - \mathbf{v}_{\rm gas}\rangle$ evaluated within the same kernel. This resembles a Bondi-Hoyle type model, but in a medium where some external source of gravity (e.g. background stars or dark matter) dominate the potential and set up the $\rho \propto r^{-1}$ profile.

• BH_GRAVACCRETION = 7: If we assume the gas has a micro-physical Maxwellian distribution with σ_v^2 (as **(6)**) and that the mass fraction which is sub-virial within $< R$ collapses into an isothermal sphere, with constant accretion rate per dynamical time determined at the radius where the BH begins to dominate the potential, we have:

$$\dot{M}_{\alpha\,{\rm disk}} \sim \epsilon_{\rm acc}\,\frac{4\pi\,G^2\,M_{\rm enc}^2\,\rho_{\rm gas}}{(\sigma_v^2 + V_c^2)^{3/2}} \tag{10}$$

where $M_{\rm enc}(<R)$, $\rho(R)$, $V_c^2 = G M_{\rm enc}(<R)/R$ are again evaluated at R. This is similar to the model developed in Hobbs et al. (2012), and to some "modified Bondi" prescriptions adopted in GADGET-3, but with some important differences. In the limit where the mass inside $<R$ is self-gravitating ($\sigma_v \ll V_c$), this becomes the canonical Shu-type isothermal accretion onto a point mass, $\dot{M} \sim f_{\rm gas}\,V_c^3/G$. In the limit where self-gravity of the ambient medium is negligible so Bondi-Hoyle-Littleton theory applies, this becomes the usual Bondi-Hoyle expression: $\dot{M} \sim G^2\,M_{\rm BH}^2\,\rho/\sigma_v^3$. If gravity is provided by an external (non-BH, non-gas) source (e.g. stars or dark matter), this is equivalent to the local Hopkins & Quataert (2011) model, for the assumption of a similar fractional amplitude $|a|$ of the non-axisymmetric modes on all scales $\sim R$, where the material with local $\delta\mathbf{v} \ll V_c$ is assumed to form a thin disk on un-resolved scales. In that limit this is also similar to the gravito-turbulent scaling in an isothermal sphere if the quantity $\alpha\,(c_s/V_c)^2 \sim$constant.

• BH_GRAVACCRETION = 8: This adopts the Hubber et al. (2013) scalings for accretion rates, where two accretion rates are calculated, one $\dot{M}_{\rm ff} = -\sum 4\pi\,r^2\,\Delta\mathbf{v}_a\cdot\Delta\mathbf{r}_a\,|\Delta\mathbf{r}|_a\,W(|\Delta\mathbf{r}_a|)$ (summed over gas with $\Delta\mathbf{u}_a \equiv \mathbf{u}_a - \mathbf{u}_{\rm BH}$ representing a kernel-weighted average of $4\pi\,r^2\,v_r$ of the gas, the other given by $\dot{M}_{\rm disk} \sim 0.01\,M_{\rm gas}(<R)/\sum(G M_{\rm BH}\,|\Delta\mathbf{r}_a|)^{1/2}\,c_{s,a}^{-2}\,(m_a/\rho_a)\,W(|\Delta\mathbf{r}_a|)$, a weighted sum of an α-disk-like (with $\alpha = 0.01$) accretion timescale. The estimator of $f_{\rm disk}$ above is used to interpolate between the two, as $\dot{M} \sim \dot{M}_{\rm ff}^{1-f_{\rm disk}}\,\dot{M}_{\rm disk}^{f_{\rm disk}}$. We caution that these estimators can give highly un-physical behavior outside of the regime of smooth, extremely well-resolved symmetric flows with the gas in pure inflow with the BH at the center.

5.1.3 *The Un-Resolved Limit: Bondi-Hoyle Model (BH_BONDI)*

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:

$$\dot{M}_{\alpha\,{\rm disk}} = 4\pi\alpha\,G^2\,M_{\rm BH}^2\,\rho\,(c_s^2 + \beta\,|\mathbf{v}_{\rm BH} - \mathbf{v}_{\rm gas}|^2)^{-3/2} \tag{11}$$

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. There are different choices to model this. If for example we use the standard formulation of an α-disk from Shakura & Sunyaev (1973), and use their outermost disk regime (where gas pressure dominates over radiation pressure) as the "rate limiter" and integrate out to some maximum radius, giving an expression for $\dot{M}_{\rm BH}$ which depends on the α disk parameters as $\dot{M}_{\rm BH} \propto M_{\alpha\,{\rm disk}}\,(M_{\alpha\,{\rm disk}}/M_{\rm BH})^{0.4}\,M_{\rm BH}^{-0.05}\,R_{\alpha\,{\rm disk}}^{-1.6}$. For a canonical $\alpha \approx 0.1$, given the uncertainties in the exact scalings above, we simplify this (without much difference) by simply adopting a "depletion time" of the α disk which scales with the mass ratio of the disk to the BH to the power given above. Specifically:

$$\dot{M}_{\rm BH} = \frac{M_{\alpha\,{\rm disk}}}{t_{\rm depl}} \tag{12}$$

$$t_{\rm depl} \equiv 4.2\times 10^7\,M_\odot\,\left(\frac{M_{\alpha\,{\rm disk}}}{M_{\alpha\,{\rm disk}} + M_{\rm BH}}\right)^{-0.4} \tag{13}$$

Note this implies Eddington-limited accretion when $M_{\alpha\,{\rm disk}} \gtrsim M_{\rm BH}$ (which corresponds to e.g. the maximum accretion disks that can be sustained without self-gravity becoming important), with the Eddington ratio declining as $\sim (M_{\alpha\,{\rm disk}}/M_{\rm BH})^{1.4}$ when $M_{\alpha\,{\rm disk}} \ll M_{\rm BH}$.

5.2.3 *Eddington Limit*

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

$$\dot{M}_{\rm BH,\,Edd} \approx 2.38\,\frac{M_\odot}{\rm yr}\,\left(\frac{M_{\rm BH}}{10^8\,M_\odot}\right)\left(\frac{\epsilon_r}{0.1}\right)^{-1} \tag{14}$$

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).[5]

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

The code includes explicit treatment of feedback from AGN in thermal, mechanical, radiative, and relativistic/cosmic-ray forms. Any or all of these can be enabled in any or all combinations, as described below. 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 Pure Thermal Feedback (`BH_THERMALFEEDBACK`)

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.

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; `BH_WIND_...`)

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

$$\eta_P \equiv \frac{\dot{M}_{\rm wind}\,v_{\rm wind}}{L/c} = \beta\left(\frac{v_{\rm wind}}{\epsilon_r\,c}\right) \approx \beta\left(\frac{v_{\rm wind}}{30,000\,{\rm km\,s^{-1}}}\right) \qquad (15)$$

$$\eta_E \equiv \frac{\dot{M}_{\rm wind}\,v_{\rm wind}^2}{2L} = \frac{\epsilon_r}{2}\,\frac{\eta_P^2}{\beta} \approx 0.05\,\beta\left(\frac{v_{\rm wind}}{30,000\,{\rm km\,s^{-1}}}\right)^2 \qquad (16)$$

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 $\Delta\Omega$. 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 $\mathbf{J}_{\rm tot}$ discussed above (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. $f = 1/(1+\beta)$.

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" (`BH_WIND_KICK`)*

Numerically the simplest approach, if one wishes to ensure a given velocity is reached, one can probabilistically "kick" gas particles (enabling `BH_WIND_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

[5] 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.

is equal, there will be fewer "kicks"). This is algorithmically similar to the sub-grid 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. (2017).

6.2.2 *Continuous Mass-Loss (`BH_WIND_CONTINUOUS`)*

A somewhat more detailed wind model is enabled by `BH_WIND_CONTINUOUS`. 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. (2018a). So the fractional weights actually used ($\bar{\omega}_{ba}$) to assign energy and momentum on a per-timestep basis to neighboring gas elements are calculated as:

$$\bar{\omega}_{ba} \equiv \frac{\omega_{ba}}{\sum_c |\omega_{ca}|} \tag{17}$$

$$\omega_{ba} \equiv \frac{\Delta\Omega_{ba}}{4\pi} \equiv \frac{1}{2}\left(1 - \frac{1}{\sqrt{1 + (\mathbf{A}_{ba}\cdot\hat{\mathbf{x}}_{ba})/(\pi\,|\mathbf{x}_{ba}|^2)}}\right) \tag{18}$$

where $\mathbf{A}_{ba}$ is the effective vector face between elements b and a used in the finite-volume hydrodynamic calculations,[6]

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 (`BH_WIND_SPAWN`)*

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 *Radiation Transport/Solver Algorithms*

At a fundamental level, the user has two choices for how to treat radiative feedback from an AGN. If true radiation-hydrodynamics (RHD) is enabled, using any of the built-in RHD solvers (e.g. the moments-based M1 or FLD solvers, the LEBRON solver, the direct intensity integration method), then radiation transport is solved on-the-fly along with the hydrodynamics, with the user free to modularly choose whichever bands they wish to resolve. Those bands have their appropriate source and sink functions and coupling to e.g. gas cooling/heating and radiation pressure terms – most of the relevant cases for AGN are already built into the code, but users should expand them as needed. If RHD is too expensive, approximate solvers exist for some of the key radiative feedback effects.

6.3.2 *Compton Heating/Cooling (`RT_XRAY` or `BH_COMPTON_HEATING`)*

If true RHD is enabled, then enabling `RT_XRAY` will turn on explicit tracking of a two-band (soft and hard) X-ray spectrum injected by the BH, while if the approximate solvers are used, `BH_COMPTON_HEATING` solves the same using an optically thin approximation.[7]

As discussed in Sazonov et al. (2004, 2005), this effect is nearly independent of obscuration: Compton heating is entirely dominated by photons with energies $\gg 10\,$keV (for which

[6] For our MFM hydrodynamic method, the face $\mathbf{A}_{ba}$ is defined as (see Hopkins 2015):

$$\mathbf{A}_{ba} \equiv \bar{n}_a^{-1}\,\bar{\mathbf{q}}_b(\mathbf{x}_a) + \bar{n}_b^{-1}\,\bar{\mathbf{q}}_a(\mathbf{x}_b) \tag{19}$$

$$\bar{\mathbf{q}}_b(\mathbf{x}_a) \equiv \mathbf{E}_a^{-1}\cdot\mathbf{x}_{ba}\,W(\mathbf{x}_{ba}, H_a) \tag{20}$$

$$\mathbf{E}_a \equiv \sum_c (\mathbf{x}_{ca}\otimes\mathbf{x}_{ca})\,W(\mathbf{x}_{ca}, H_a) \tag{21}$$

For SPH, the face is defined by the simpler relation $\mathbf{A}_{ba} = [\bar{n}_a^{-2}\,\partial W(|\mathbf{x}|_{ba}, H_a)/\partial|\mathbf{x}|_{ba} + \bar{n}_b^{-2}\partial W(|\mathbf{x}|_{ba}, H_b)/\partial|\mathbf{x}|_{ba}]\,\hat{\mathbf{x}}_{ba}$. In moving-mesh or fixed-grid finite-volume codes, the face $\mathbf{A}_{ba}$ is the explicit geometric mesh face between cells.

[7] 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. We neglect obscuration in the `BH_COMPTON_HEATING` module (the true RT modules include opacity explicitly) 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 $\approx 2\times10^7$ K as calculated in Sazonov et al. (2004) for a broad range of observed QSO SED shapes.

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 – hence the approximation made in `BH_COMPTON_HEATING`. In the cooling function, we add the appropriate Compton heating and cooling terms.[8] 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). [9]

6.3.3 *Photo-Ionization (`RT_CHEM_PHOTOION` or `BH_HII_HEATING`)*

Enabling `RT_CHEM_PHOTOION`, the code explicitly tracks ionizing photons with RHD from the BH; enabling `BH_HII_HEATING`, photo-ionization is trated approximately with a local Stromgren approximation identical to that described in Hopkins et al. (2018b).[10] We calculate the rate of production of ionizing photons from the empirically-determined QSO spectra in Hopkins et al. (2007), $\dot{N}_{\rm ion} \approx 5.5\times10^{54}\,{\rm s^{-1}}\,(L_{\rm bol}/10^{45}\,{\rm erg\,s^{-1}})$, and assume a simple power-law spectral slope which (for analytic convenience) has the same shape as the extragalactic UVB (harder radiation from AGN can be included separately in the more sophisticated multi-bin variant of `RT_CHEM_PHOTOION`, or the soft/hard X-ray treatment above).

6.3.4 *Radiation Pressure (`RT_...` or `BH_PHOTONMOMENTUM`)*

With the "true" RHD methods, the radiation pressure (RP) terms are always tracked whenever absorption occurs. You need to be sure bands like the optical, UV, near IR are included to provide most of the single-scattering radiation pressure, with the relevant AGN-source terms and opacities (these are modular but you need to actually make sure they are included! It's very easy – one line addition – to include in the code). If you want to account for multiple scattering in the IR see the detailed descriptions of the grey-body treatment of the spectrum and opacities in the `RT_INFRARED` module.

For approximate RHD, with `BH_PHOTONMOMENTUM`, we simplify (and act more conservatively) by assuming 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, and weighted akin to the continuous mass-loss mechanical prescription above). 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), solved according to the FIRE-default variant of the LEBRON algorithm. 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 in this module by default.

However, an important remaining ambiguity is the directional dependence of the flux, which can be highly non-isotropic. In principle, one can weight this based on the polar angle θ, which is defined with respect to the angular momentum vector (of gas in the vicinity of the BH) $\boldsymbol{\omega}=|\omega|\,\hat{z}$, $\cos\theta=|\hat{r}\cdot\hat{z}|$, for a particle at position $\mathbf{r}$ with respect to the BH at the origin, in the function `bh_angleweight`.

6.4 Relativistic Jets (`BH_COSMIC_RAYS`)

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

However, non-relativistic MHD jets can be trivially implemented by modifying the angular dependence and velocities of the outflow mass/energy/momentum in the "mechanical feedback section." Any of those algorithms could trivially represent a jet for the appropriate choice of parameters. For arbitrarily narrow collimated jets, it may be impossible at a given resolution to represent these with the `BH_WIND_KICK` or `BH_WIND_CONTINUOUS` algorithms, because these cannot resolve an opening angle smaller than the actual resolution of the gas grid in the vicinity of the BH (i.e. the outflow will always be smeared over some opening angle $\sim\Delta x/H$, where H is the inter-neighbor or kernel size over which it is initially coupled, and Δx is the width of a single cell). For this reason narrow jets are better treated with the `BH_WIND_SPAWN` flag, where in principle the spawned particles could be launched along exactly one axis (recall, the local angular momentum axis is defined as above). WIth MHD active these can also trivially be given magnetic energy.

Relativistic particles *within* the outflows/jets can be represented, using the explicit code treatment of cosmic rays in the MHD-like fluid limit. For this, simply enable the usual cosmic ray physics and cosmic ray transport modules in the code as desired, then include the flag `BH_COSMIC_RAYS`, which will inject CRs alongside any other mechanical energy (using any of the `BH_WIND_...` modules as described above), with the specified energy relative to the accretion energy onto the BH given by the `BH_CosmicRay_Injection_Efficiency` parameter in the parameterfile (i.e. $\dot{E}_{\rm cr}=\epsilon_{\rm cr}\,\dot{M}_{\rm BH}\,c^2$, with $\epsilon_{\rm cr}=$`BH_CosmicRay_Injection_Efficiency`).

[8] 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.

[9] As shown in Faucher-Giguère & 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.

[10] For the 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.

REFERENCES

Anglés-Alcázar, D., Davé, R., Faucher-Giguère, C.-A., Özel, F., & Hopkins, P. F. 2017, MNRAS, 464, 2840

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-Giguère, C.-A., & Quataert, E. 2012, MNRAS, 425, 605

Gammie, C. F. 2001, ApJ, 553, 174

Hobbs, A., Power, C., Nayakshin, S., & King, A. R. 2012, MNRAS, 421, 3443

Hopkins, P. F. 2015, MNRAS, 450, 53

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., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C. C., & El-Badry, K. 2018a, MNRAS, 477, 1578

Hopkins, P. F., Wetzel, A., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., Boylan-Kolchin, M., Murray, N., Hayward, C. C., Garrison-Kimmel, S., Hummels, C., Feldmann, R., Torrey, P., Ma, X., Anglés-Alcázar, D., Su, K.-Y., Orr, M., Schmitz, D., Escala, I., Sanderson, R., Grudić, M. Y., Hafen, Z., Kim, J.-H., Fitts, A., Bullock, J. S., Wheeler, C., Chan, T. K., Elbert, O. D., & Narayanan, D. 2018b, MNRAS, 480, 800

Hubber, D. A., Allison, R. J., Smith, R., & Goodwin, S. P. 2013, MNRAS, 430, 1599

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