Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 8 November 2017 (MN LATEX style file v2.2)

GIZMO: Cosmic Ray Notes

PFH

8 November 2017

1 COSMIC RAY IMPLEMENTATION NOTES

1.1 Units, Definitions, etc

The cosmic ray (CR) implementation in the code is activated with the COSMIC_RAYS flag.

CRs are treated as a single-species, ultra-relativistic ($\gamma_{cr} = 4/3$) fluid. In the code we evolve the conserved variable $E_{\rm cr,\,i}$, the total CR energy associated with particle i. This is written out in snapshots, with the flag name "CosmicRayEnergy" for all gas particles, in physical code units (energy = mass × velocity^2).

The CR specific energy per unit mass u_{cr}, energy density e_{cr}, CR pressure P_{cr}, total pressure P, and sound speed c_s are given by:

$$u_{cr} = E_{cr}/m_i \tag{1}$$

$$e_{cr} = E_{cr}/V_i = \rho_i\, u_{cr} \tag{2}$$

$$P_{cr} = (\gamma_{cr} - 1)\, e_{cr} \tag{3}$$

$$P = P_{\rm thermal} + P_{cr} = (\gamma_{\rm gas} - 1)\, u_{\rm gas}\, \rho + P_{cr} \tag{4}$$

$$c_s^2 = \frac{\partial P}{\partial \rho} = \gamma_{\rm gas}(\gamma_{\rm gas} - 1)\, u_{\rm gas} + \gamma_{cr}(\gamma_{cr} - 1)\, u_{cr} \tag{5}$$

The "correct" total pressure, sound speed, etc, are all accounted for in the reconstruction at faces and Riemann problem (RP). CRs act on the gas in the RP via modifying pressure/sound speed/internal energy. Their adiabatic compression/expansion is accounted for also, via an operator-split approach.

1.2 Evolution Equations, Term-By-Term

The evolution equation for the CR energy density follows Skilling 1971, 1975 (for more details see also Uhlig et al., 2012).

$$\begin{aligned}\frac{\partial e_{cr}}{\partial t} =& ({\bf v} + {\bf v}_{\rm st}) \cdot \nabla P_{\rm cr} - \nabla \cdot [({\bf v} + {\bf v}_{\rm st})\,(e_{cr} + P_{cr})] \\ &+ \nabla \cdot [{\bf v}_{di}\, e_{cr}] - \Gamma_{cr} + \dot{e}_{*} + \dot{e}_{\rm AGN}\end{aligned} \tag{6}$$

where ${\bf v}$ is the gas speed, ${\bf v}_{st}$ is the CR streaming speed, and ${\bf v}_{di}$ is a "diffusion speed" (defined below). Because our method is Lagrangian, we re-write this with the Lagrangian derivative $d/dt = \partial/\partial t + {\bf v} \cdot \nabla$. Then after a bit of algebra, you get:

$$\begin{aligned}\rho\, \frac{du_{cr}}{dt} =& - P_{cr} \nabla \cdot {\bf v} + {\bf v}_{st} \cdot \nabla P_{cr} - \nabla \cdot [{\bf v}_{st}(e_{cr} + P_{cr})] \\ &+ \nabla \cdot ({\bf v}_{di}\, e_{cr}) - \Gamma_{cr} + \dot{e}_{*} + \dot{e}_{\rm AGN}\end{aligned} \tag{7}$$

Now I'll go through each of these terms in detail.

1.2.1 *Advection & Adiabatic Effects with Gas*

(0) Note because we have taken the Lagrangian derivative above, the pure advection with the gas flow is automatically accounted for.

(1) $-P_{cr} \nabla \cdot {\bf v}$: just adiabatic CR compression/expansion heating. This is handled via a simple operator split (one-line) in the hydro routine after the RP is solved. Hubble-flow terms (for cosmological runs) are accounted for in this term.

1.2.2 *Streaming*

(2) ${\bf v}_{st} \cdot \nabla P_{cr}$: streaming instability wave-heating term due to self-excited waves that get rapidly damped in the plasma. Because ${\bf v}_{st} \mid\mid (-\nabla P_{cr})$, this produces an energy loss from the CRs to the gas. We treat this as a heating/cooling term: CRs lose energy, gas gains thermal energy, at a rate $dE/dt = |{\bf v}_{st}|\, |\nabla P_{cr}|\, V_i$. This is operator-split from the RP. This term is ZERO if you turn on COSMIC_RAYS_DISABLE_STREAMING. (Note that with MHD on, as detailed below, this is appropriately modified because instead of ${\bf v}_{st} \mid\mid (-\nabla P_{cr})$, the transport is along field lines, so there is an extra factor of $|\hat{B} \cdot \hat{\nabla} P_{cr}|^2$).

(3) $\nabla \cdot [{\bf v}_{st}(e_{cr} + P_{cr})]$: transport of CRs relative to gas owing to streaming (this term is also ZERO if you turn on COSMIC_RAYS_DISABLE_STREAMING). Temporarily ignore the role of magnetic fields: then the CR streaming velocity ${\bf v}_{st}$ is approximately:

$${\bf v}_{st} = -|v_{st}|\, \frac{\nabla P_{cr}}{|\nabla P_{cr}|} \tag{8}$$

i.e. the CR streaming occurs opposite the direction of the CR pressure gradient (my understanding is, if we were doing a full energy-by-energy-bin breakdown of the CRs, the streaming direction in each energy bin would be along the direction of the CR number density. If we assume a single energy regime dominates the pressure and follow that, or assume there is always a power-law momentum distribution, then these directions are proportional to the CR pressure direction. So given our single-fluid model for the CRs, we take this).

Now, if there are magnetic fields (MAGNETIC on in the code), then CRs are locked to field lines and we follow the usual approach of projecting the gradient onto the field direction and allowing the transport only in that direction (just like with conduction, viscosity, etc) – this means

$${\bf v}_{st} = -|v_{st}|\, \hat{B}\, \frac{\hat{B} \cdot \nabla P_{cr}}{|\nabla P_{cr}|} \tag{9}$$

$$\hat{B} \equiv \frac{\bf B}{|{\bf B}|} \tag{10}$$

For the streaming speed, there is a lot of literature (Wentzel 1968, Skilling 1971, 1975, Holman 1979, as updated in Kulsrud 2005, Yan & Lazarian 2008, Ensslin 2011). Grossly simplifying, in the weak-field regime (plasma $\beta \gg 1$), the streaming speed is approximately the gas thermal sound speed c_s, in the strong-field regime it is approximately the Alfven velocity v_A. So we take

$$|v_{st}| = \begin{cases} \sqrt{c_s^2 + v_A^2} & \text{(MHD on)} \\ c_s & \text{(MHD off)} \end{cases} \tag{11}$$

Note that when you put this back into the original equations, this streaming term has a functional form that looks like just a diffusion equation, namely:

$$\rho\,\frac{du_{cr}}{dt} = \begin{cases} \nabla\cdot(\hat{B}\,\kappa_{st}\,[\hat{B}\cdot\nabla P_{cr}]) & \text{(MHD on)} \\ \nabla\cdot(\kappa_{st}\,\nabla P_{cr}) & \text{(MHD off)} \end{cases} \tag{12}$$

with

$$\kappa_{st} \equiv |v_{st}|\,\frac{e_{cr}+P_{cr}}{|\nabla P_{cr}|} = |v_{st}|\left(\frac{\gamma_{cr}}{\gamma_{cr}-1}\right)\frac{P_{cr}}{|\nabla P_{cr}|} \tag{13}$$

This is then treated like a diffusion equation in our standard numerical fashion, with this coefficient. Note that the value of the CR pressure scale-length $L_{cr} \equiv P_{cr}/|\nabla P_{cr}|$ used in this computation is limited (for purely numerical reasons, like most slope-limiters) to a value equal to or greater than 1/3 particle "size", and less than 200 times the particle size (because at this point the CRs are already streaming out of the particle very rapidly, so their coefficient will need to be re-computed anyways as they move into other particles).

1.2.3 *Diffusion (Microscopic & Turbulent)*

(4) $\nabla\cdot(\mathbf{v}_{di}\,e_{cr})$: "standard" diffusion (not streaming). Here the "diffusion velocity" is

$$\mathbf{v}_{di} = \begin{cases} \kappa_{di}\,\hat{B}\,\frac{\hat{B}\cdot\nabla e_{cr}}{e_{cr}} = \kappa_{di}\,\hat{B}\,\frac{\hat{B}\cdot\nabla P_{cr}}{P_{cr}} & \text{(MHD on)} \\ \kappa_{di}\,\frac{\nabla e_{cr}}{e_{cr}} = \kappa_{di}\,\frac{\nabla P_{cr}}{P_{cr}} & \text{(MHD off)} \end{cases} \tag{14}$$

where κ_{di} is a diffusion coefficient.

We treat this also as a standard diffusion term to solve the relevant equations. So if this and streaming are active, we solve them with the same numerical operation, using the combined "effective coefficient" $\kappa_{\rm eff} = \kappa_{st} + \kappa_{di}$. You can turn this part ("standard" diffusion) off by setting the flag COSMIC_RAYS_DISABLE_DIFFUSION

If you set the flag COSMIC_RAYS_DIFFUSION_CONSTANT, then the value of the parameterfile value "CosmicRayDiffusionCoeff" is set to κ_{di} (this should be input in CODE UNITS, so code velocity × code length). If this flag is not set, we will calculate κ_{di}. Canonical Milky Way (MW) values are $\sim 1-10\times10^{28}\,{\rm cm^2\,s^{-1}}$ ($\sim 20-200\,{\rm kpc}/h\times{\rm km\,s^{-1}}$). But for physically reasonable assumptions, values of the diffusion and streaming coefficients in systems unlike the MW (or different MW regions) could easily be different by factors of ~ 100, hence our default option being to calculate these.

To explain MW empirically inferred diffusion coefficients, what is usually required is turbulence on scales of the Larmor radii of the CR orbits around field lines. This gives a coefficient which scales as

$$\kappa_{di} \sim \frac{v_{cr}\,r_g}{3}\,\frac{B^2_{\rm coherent}}{B^2_{\rm random}(r_g)} \tag{15}$$

where $v_{cr} \sim c$ (speed of light) for our assumptions, $r_g = p\,c/(Z\,e\,B) \approx 3\times10^{12}\,{\rm cm}\,R_{GV}\,(B/\mu G)^{-1}$ is the gyro radius ($R_{GV} \sim 1$ is the magnetic rigidity in gigavolts for the particles that dominate the CR pressure, which we will assume is a constant for the CR fluid; p, Ze are momentum and charge of particles, respectively), $B_{\rm coherent}$ is the large-scale coherent magnetic field, and $B_{\rm random}(r_g)$ is the rms random (turbulent) component of the field on scales corresponding to the gyro radius.

If we assume a Kolmogorov spectrum, with $B^2_{\rm coherent} \approx B^2(r \sim L_{\rm drive} \sim {\rm kpc})$, then you get a reasonable

$$\kappa_{di} \sim 3\,R_{GV}^{1/3}\times10^{28}\,\frac{\rm cm^2}{\rm s}\left(\frac{\mu G}{|\mathbf{B}|}\right)^{\frac{1}{3}}\left(\frac{L_{\rm drive}}{\rm kpc}\right)^{\frac{2}{3}} \tag{16}$$

Similar to our modeling of sub-grid turbulent phenomena, we can (very crudely) estimate $L_{\rm drive} \sim P/|\nabla P|$, i.e. as the pressure gradient scale length. Note, in simulations without MHD, we have no value of $\mathbf{B}$, so we simply assume pressure equilibrium ($|\mathbf{B}|^2/2 \sim P_{\rm thermal}$) to obtain an "effective" value we can insert in the equation above.[1]

However, if the flow is sufficiently smooth, or the field sufficiently weak (e.g. IGM), this can give an unphysically large κ_{di}. This can occur if the effective streaming/diffusion velocity were $> c$ or if the gradient scale length of pressure in e.g. the streaming term were much larger than the mean free path $\sim 3\times10^{25}\,{\rm cm}/(n_H/{\rm cm^{-3}})$. Therefore we also enforce this as a cap, but it is almost never relevant.

Finally, if COSMIC_RAYS_DIFFUSION_CONSTANT is not set, then the run-time parameter "CosmicRayDiffusionCoeff" multiplies $\kappa_{\rm eff}$ (sum of streaming and diffusion coefficients). So the default value is 1.

[1] Alternatively, similar to how we estimate the local "random" component for e.g. turbulent diffusion models, we can estimate $B^2_{\rm random}(h)$ on the resolution scale h by $|B^2_{\rm random}(h)|^{1/2} \approx |\nabla\otimes\mathbf{B}|\,h$ (where here, $|\nabla\otimes\mathbf{B}|$ refers to the Frobenius norm – i.e. squared sum of all components – of the matrix $\nabla\otimes\mathbf{B}$). We then assume a Kolmogorov spectrum for B ($B^2_{\rm random}(r) \sim B^2_{\rm random}(h)\,(r/h)^{2/3}$; note even in super-sonic MHD turbulence, the B spectra are not too different from this, we think). This gives:

$$\kappa_{di} \sim 2\times10^{28}\,\frac{\rm cm^2}{\rm s}\,R_{GV}^{1/3}\left(\frac{\mu G}{|\mathbf{B}|}\right)^{\frac{1}{3}}\left(\frac{|\mathbf{B}|}{h\,|\nabla\otimes\mathbf{B}|}\right)^2\left(\frac{h}{100\,{\rm pc}}\right)^{\frac{2}{3}} \tag{17}$$

Note that in a Kolmogorov cascade, this is resolution/scale-independent. But this allows for e.g. local amplification of fields in regions like GMCs that can depart from a galactic-disk scale cascade; moreover if we are not in a turbulent cascade (but a laminar flow), then this (qualitatively correctly) will give a very large κ_{di}. And it attempts to self-consistently account for different driving scales (which will vary in different systems). However, I find it can be considerably "noisier", since the local gradients vary a lot, so you will (even in a well-resolved cascade) have regions where the $\mathbf{B}$-field gradients are locally flat, and it can screw up this estimator.

1.2.4 Cooling (Hadronic & Coulomb Losses)

(5) Γ_{cr}: This represents "cooling losses" of CRs to gas and radiation. Requires COOLING be on. Note that if COSMIC_RAYS_DISABLE_COOLING is set, $\Gamma_{cr} = 0$ even if COOLING is on. We adopt the estimate for combined hadronic plus Coulomb losses from Volk 1996 and Ensslin 1997 as synthesized and updated in Guo & Oh 2008. Putting in cosmological hydrogen mass fractions, this gives:

$$\Gamma_{cr} = 7.51\times10^{-16}\,\mathrm{s}^{-1} e_{cr}\,(1+0.22\,\tilde{n}_e)\left(\frac{n_H}{\mathrm{cm}^{-3}}\right) \quad (18)$$

where n_H is the hydrogen number density and $\tilde{n}_e$ is the number of free electrons per hydrogen nucleus.

Following their estimate, $\sim 1/6$ of the hadronic losses (the term which is independent of $\tilde{n}_e$) and all of the Coulomb losses (the term in $\tilde{n}_e$), are thermalized and go into heating the gas, so

$$\dot{e}_{\rm gas} = 7.51\times10^{-16}\,\mathrm{s}^{-1} e_{cr}\,(0.17+0.22\,\tilde{n}_e)\left(\frac{n_H}{\mathrm{cm}^{-3}}\right) \quad (19)$$

1.2.5 Injection & Losses in Star Formation & Feedback

(6) $\dot{e}_*$: This represents injection and losses of CRs via star formation and feedback. When a gas particle turns into stars, we assume the CR energy carried by the particle is lost (this is almost always negligible on a galaxy scale, and is just for numerical convenience so we don't have to deal with how to redistribute this energy). If GALSF_FB_SNE_HEATING is active (our standard FIRE SNe model), then a fraction ϵ equal to the run-time parameter CosmicRay_SNeFraction of the initial kinetic energy of SNe is injected as CRs. So whenever there is a SNe or fast stellar wind "injection" event, we calculate the initial kinetic energy (i.e. before cooling losses), take ϵ of this, and inject that into the neighboring gas particles (with the same kernel weighting as for the SNe energy/momentum). Note that we require the initial ejecta velocity be above some threshold which I have somewhat arbitrarily chosen to be $500\,\mathrm{km\,s}^{-1}$, namely so that AGB ejecta (at low speeds) aren't contributing to this (just SNe and fast stellar winds). Canonical values of ϵ are ~ 0.1.

1.2.6 Injection from AGN

(7) $\dot{e}_{\rm AGN}$: Injection from AGN. This is currently not implemented in the main code, but I have a version with a few different ideas I have been playing with, and wanted to list it here for completeness. Obviously, we can do a similar thing to the SNe model above, and choose some fraction of the accretion energy to inject as BH feedback in the form of CRs. This can be done in the immediate vicinity of the BH (especially trivial), or some distance from the hole (mimicking something like Deborah's "bubbles" model).

1.3 MHD vs. Non-MHD

For all cases of interest, there is no question that the MHD version of these equations is "more correct." Other than some hand-waving about equipartition, there isn't much motivation for the non-MHD versions of the equations above – these can be seriously wrong (qualitatively and quantitatively, by orders of magnitude). But most of the work applying CRs to galaxies has ignored magnetic fields, so it's useful to have this option as a "straw-man" to see how badly this may (or may not) do.

1.4 Tests

I've run a number of simple tests. The numerical implementation of the terms is well-tested. Whether the terms/scalings/etc are reasonable under all cases or could extrapolate to strange behaviors is another question altogether.

(a) CR shocktube: from Pfrommer et al. 2007. Only includes the RP and adiabatic terms ((0)-(1) above). Checks basic shock-capturing, adiabatic evolution, ability to solve 2-fluid RP. 1D/2D/3D looks good.

(b) Diffusing sheet/vortex (1D/2D/3D): Basic diffusion problem, here with constant coefficient of diffusivity, no gas dynamics (only diffusion) - tests whether numerical integration of diffusion equation behaves appropriately (and is numerically stable, which is actually highly-non-trivial for diffusion equations solved explicitly). Good.

(c) Injection & Cooling test: simple test with injection and/or cooling terms turned on in a box with no dynamics and constant injection rate, to see if energy grows/decays in time at expected rate (i.e. the implementation isn't bugged-up). Good.

(d) Isolated (non-cosmological) galaxies: with adiabatic equation of state, old Springel & Hernquist model, and FIRE physics. For MW like galaxy, reasonable coefficients for diffusion (had them print out everything as it ran). No crashing/obvious bugs in newest code. CR energy approaches quasi-equipartition with magnetic and turbulent energies in the dense (multi-phase gas). Turning off CR cooling/loss terms it builts up a thick, CR-pressure supported disk (as expected). Weak wind if no other FB present. Basically none of this is rigorous since there is no exact solution to compare to here, but it is a "sanity check" as well as useful de-bugging.

(e) Cosmological zoom with FIRE physics. Only run briefly to make sure no glaringly-obvious bugs, but not extensively tested.

2 EVOLVING CRS WITH THE M1 (AS OPPOSED TO PURE-DIFFUSION, M0 OR FLD-LEVEL) APPROXIMATION

Consider CR transport (ignore cooling, advection, streaming, etc). We approximate this via diffusion but really we should solve the collisionless Boltzmann equation (CBE). Just like for RT with FLD, M1 etc., we can take the first and second moments of the CBE to obtain:

$$\frac{\partial e_\nu}{\partial t} + \nabla\cdot\mathbf{F}_\nu = S_\nu - (\tilde{c}/\lambda_\nu)\,e_\nu \quad (20)$$

$$\frac{1}{\tilde{c}}\,\frac{\partial \mathbf{F}_\nu}{\partial t} + \tilde{c}\,\nabla\cdot\mathbb{P}_\nu = -\lambda_\nu^{-1}\,\mathbf{F}_\nu \quad (21)$$

where e_ν is the CR energy density within some differential energy bin; $\mathbf{F}_\nu$ is the corresponding flux, $\mathbb{P}_\nu$ is the pressure tensor, and $\tilde{c}$ is the free-streaming propagation speed. The right hand side represents sources/sinks/scattering, with S_ν the source function and λ_ν is the mean-free path (MFP). Integrate over the CR energy spectrum ν (or equivalently consider a single-bin). Assume for now that the only sources/sinks are isotropic scattering (we can simply operator-split the cooling term, same as we do now). Then the RHS of the energy equation vanishes, but the RHS of the flux equation remains (with λ_ν^{-1} the MFP to scattering). Since we're in something like the kinetic MHD limit the pressure tensor term takes the form $\nabla\cdot\mathbb{P}_\nu \rightarrow \hat{B}\,(\hat{B}\cdot\nabla P_\nu)$ (for the fully-anisotropic case); or if we assume isotropic transport ($\mathbb{P} = P\,\mathbb{I}$) then $\nabla\cdot\mathbb{P}_\nu \rightarrow \nabla P_\nu$. Recall $P_\nu = (\gamma_{\rm cr}-1)\,e_\nu$.

Putting this all together the Eqs. become:

$$\frac{\partial e_{\rm cr}}{\partial t} = -\nabla\cdot\mathbf{F}_{\rm cr} \tag{22}$$

$$\frac{\kappa_{\rm cr}}{(\gamma_{\rm cr}-1)\,\tilde{c}^2}\,\frac{\partial \mathbf{F}_{\rm cr}}{\partial t} + \kappa_{\rm cr}\,\hat{B}\,(\hat{B}\cdot\nabla e_{\rm cr}) = -\mathbf{F}_{\rm cr} \tag{23}$$

where we have defined the diffusivity $\kappa_{\rm cr} \equiv \tilde{c}\,\lambda_{\rm cr}\,(\gamma_{\rm cr}-1)$.

On timescales $> t_0 \equiv \kappa_{\rm cr}/[(\gamma_{\rm cr}-1)\,\tilde{c}^2]$, the solution to the flux equation rapidly converges to steady-state, i.e. the $\partial\mathbf{F}_{\rm cr}/\partial t$ term becomes negligible and $\mathbf{F}_{\rm cr} = -\kappa_{\rm cr}\,\hat{B}\,(\hat{B}\cdot\nabla e_{\rm cr})$. Plugging this into the energy eqn. gives $\partial e_{\rm cr}/\partial t \rightarrow \nabla\cdot[\kappa_{\rm cr}\,\hat{B}\,(\hat{B}\cdot\nabla e_{\rm cr})]$, i.e. the normal diffusion equation. Just like with, say FLD for RT, our normal diffusion approximation drops the $\partial\mathbf{F}_{\rm cr}/\partial t$ term completely, but keeping it as above is actually more accurate. On small spatial/time-scales, the behavior is streaming (wave/ray like), while on larger scales, the behavior is diffusive.

Like M1, this method requires we explicitly evolve the fluxes $\mathbf{F}_{\rm cr}$, but by replacing a second-order PDE with a pair of first-order PDEs (just solving two normal advection equations), the nasty numerical-diffusion timestep constraint ($\Delta t < h^2/\kappa$, where h is the resolution) is avoided. Instead stability only requires a Courant condition ($\Delta t < h/\tilde{c}$). And its much easier to numerically stabilize the fluxes (avoiding some of the nasty issues with flux-limiters and local maxima that we have run into with TK's experiments). The approach is numerically straightforward, accurate, numerically stable, and allows the more generous timestep condition.

As before we fix $\kappa_{\rm cr}$. But now we have the parameter $\tilde{c}$. Note that this only appears in the $\partial\mathbf{F}_{\rm cr}/\partial t$ term in the flux Eqn. – it determines the timescale on which the flux reaches equilibrium. Physically for an ultra-relativistic fluid, $\tilde{c} = c$, but this would give very small timesteps. So we can use a "reduced speed of light" (RSOL). Now, unlike RT, where on large scales in the optically-thin case the solution can become free-streaming, so RSOL can introduce unphysical artifacts if you aren't careful about choosing it big enough, here it has a much weaker effect. As I said all it really does it control how quickly the solution converges to the flux-equilibrium. This timescale is $t_0 \approx 2.8\times10^5\,{\rm yr}\,(\kappa_{\rm cr}/3\times10^{28}\,{\rm cm^2\,s^{-1}})\,(\tilde{c}/1000\,{\rm km\,s^{-1}})^{-2}$. As long as this is significantly less than the dynamical times of interest on the scales of the CR diffusion scale-height, Eliot and I dont think it should make a significant difference. Again, in steady-state, the $\tilde{c}$ term vanishes so it has no effect on the solution – it only influences how rapidly you approach steady-state. So you can choose quite small values and still meet this condition.

I've done preliminary tests with TK's test setup (the older, low-res MW he sent me). I get quite reasonable values of $F_\gamma/F_{\rm sf}$ for this implementation (for $\kappa = 3\times10^{28}$, I get values of a couple $\times10^{-5}$). More importantly these are only weakly sensitive to $\tilde{c}$ over a wide range. I've tried $\tilde{c} = 30, 100, 300, 500, 1500, 5000, 15000, 30000\,{\rm km\,s^{-1}}$. For the values $\tilde{c} \gtrsim 500\,{\rm km\,s^{-1}}$ the dependence on $\tilde{c}$ of any quantity (total CR energy, mean energy density, $F_{\rm gamma}$, energy density in high-energy gas, CR scale height/length) is very weak ($F_{\rm gamma}$ changed by $\sim 20\%$ over a factor of ~ 50 change in $\tilde{c}$). Below that speed, the results did depend more strongly on $\tilde{c}$ (with $F_{\rm gamma}$ getting larger at lower $\tilde{c}$, which makes sense, since the CRs take longer to build up a strong flux in response to a strong energy density gradient, generated in dense gas by SNe depositing them) – $\sim 100\,{\rm km\,s^{-1}}$ is also where t_0 is about equal to the galaxy dynamical time in the inner regions.

For large $\kappa_{\rm cr}$, this seems like it may therefore be a more efficient approach.