RESEARCH   **Open Access**

CrossMark

A new hybrid technique for modeling dense star clusters

Carl L. Rodriguez[1*], Bharath Pattabiraman[2,3,4], Sourav Chatterjee[2,3], Alok Choudhary[2,4], Wei-keng Liao[2,4], Meagan Morscher[2] and Frederic A. Rasio[2]

Abstract

The "gravitational million-body problem," to model the dynamical evolution of a self-gravitating, collisional N-body system with $\sim 10^6$ particles over many relaxation times, remains a major challenge in computational astrophysics. Unfortunately, current techniques to model such systems suffer from severe limitations. A direct N-body simulation with more than 10^5 particles can require months or even years to complete, while an orbit-sampling Monte Carlo approach cannot adequately model the dynamics in a dense cluster core, particularly in the presence of many black holes. We have developed a new technique combining the precision of a direct N-body integration with the speed of a Monte Carlo approach. Our Rapid And Precisely Integrated Dynamics code, the `RAPID` code, statistically models interactions between neighboring stars and stellar binaries while integrating directly the orbits of stars or black holes in the cluster core. This allows us to accurately simulate the dynamics of the black holes in a realistic globular cluster environment without the burdensome N^2 scaling of a full N-body integration. We compare `RAPID` models of idealized globular clusters to identical models from the direct N-body and Monte Carlo methods. Our tests show that `RAPID` can reproduce the half-mass radii, core radii, black hole ejection rates, and binary properties of the direct N-body models far more accurately than a standard Monte Carlo integration while remaining significantly faster than a full N-body integration. With this technique, it will be possible to create more realistic models of Milky Way globular clusters with sufficient rapidity to explore the full parameter space of dense stellar clusters.

1 Main text

The dynamics of dense star clusters is one of the most challenging problems of modern computational astrophysics. The large number of particles, high interaction rate, and large number of processes with vastly different physical timescales conspire to make globular clusters (GCs) and galactic nuclei (GN) uniquely difficult to model. In particular, the large number of black holes (BHs) in both GCs and GN often dynamically interact on a much shorter timescales than the rest of the cluster (Spitzer 1969). Although only comprising a small fraction of the total cluster mass, these BHs provide the dominant energy source for GCs, especially after the BH-driven core-collapse (Morscher et al. 2015). Hence, understanding their dynamics is critical to understanding the overall evolution and present-day appearance of these systems (Mackey et al. 2008). Unfortunately, since the orbital and interaction timescales of these BHs are frequently orders-of-magnitude smaller than the interaction timescale of a typical star in the cluster, resolving these effects can be particularly difficult.

Modern stellar dynamics codes have intensely investigated GCs, with the majority of work focusing on two approaches. The N-body approach directly integrates the force of every particle on every other particle, with the current generation of codes (Portegies Zwart et al. 2001; Harfst et al. 2008; Nitadori and Aarseth 2012; Capuzzo-Dolcetta et al. 2013; Wang et al. 2015) making extensive use of state-of-the-art hardware acceleration and algorithmic enhancements. While extremely precise, this approach can require more than a year (e.g., Heggie 2014; Wang et al. 2016) to complete a full simulation of a realistic Milky-Way GC.

[*]Correspondence: carlrodr@mit.edu
[1]MIT-Kavli Institute for Astrophysics and Space Research, Cambridge, USA
Full list of author information is available at the end of the article

As such, an approximate Monte Carlo (MC) technique is often used in place of a full direct summation (Hénon 1971; Giersz 1998; Joshi et al. 2000; Freitag and Benz 2001). Whereas an N-body approach computes the orbit of stars directly, the orbit-sampling approach assumes that particle orbits remain fixed on a dynamical timescale, only changing due to slight perturbations from two-body encounters between neighboring particles. This allows the orbits to be sampled statistically, and since computing a single orbit in a fixed spherical potential is faster than computing the precise orbits in the full-N potential of a cluster, these MC models can be generated in at most a few days or weeks. However, the assumptions of spherical symmetry and dynamical equilibrium break down in the BH-dominated core, where the potential and the particle orbits are primarily determined by a small number of particles. This can lead to a substantial underprediction of the core radii by MC techniques (compared to direct N-body), particularly during the deep collapses that produce dynamically-assembled binaries (Morscher et al. 2015; Rodriguez et al. 2016a).

GCs are formed as the result of a burst of star formation in the early universe. Approximately 10 to 20 Myr after this formation is complete, the most massive stars in the cluster collapse, yielding hundreds to thousands of BHs (Belczynski et al. 2006). As the BHs are more massive than the typical cluster star, they are rapidly driven to the center of the GC by dynamical friction (Fregeau et al. 2002); once there, the number density of BHs is sufficient to form binaries via three-body encounters. While it was long-assumed that these BHs would not be retained in GCs to the present day (e.g., Sigurdsson and Hernquist 1993), recent evidence has begun to suggests otherwise.

The past decade has seen the first detections of BHs in GCs, starting with the first detection in an extragalactic GC by Maccarone et al. (2007) and several recent detections in Milky Way GCs, (Chomiuk et al. 2013; Strader et al. 2013; Miller-Jones et al. 2015), including two BH candidates in M22 (Strader et al. 2012) and the recent dynamical measurement of a $\gtrsim 4.5 M_\odot$ BH in NGC 3201 (Giesers et al. 2018). These observational results complimented recent theoretical results suggesting that the GCs can potentially retain hundreds of BHs up to the present day (Mackey et al. 2007; Downing 2012; Morscher et al. 2013, 2015; Kremer et al. 2018; Askar et al. 2018). This has led to a new theoretical understanding that the number of BHs retained in a GC directly controls the size and density of its observational core (e.g., Merritt et al. 2004; Mackey et al. 2008; Sippel and Hurley 2013; Breen and Heggie 2013; Kremer et al. 2018a; Arca Sedda et al. 2018). The importance of BHs in GCs cannot be overstated. In addition to determining the structural and evolutionary properties of the clusters, GCs also have important implications for BH astrophysics. GCs can produce X-ray binaries at a significantly higher rate than the galactic field

(Clark 1975), suggesting that there might be $\sim$ 100s of low-mass X-ray binaries in Milky Way GCs (Pooley et al. 2003). Furthermore, recent studies have shown that the second-generation of gravitational-wave detectors can potentially detect $\gtrsim$ 100 binary BH mergers per year from binaries forged in the cores of GCs (Rodriguez et al. 2015, 2016b; Antonini et al. 2016), with recent detections (Abbott et al. 2017) by LIGO/Virgo showing spin alignments suggestive of dynamical formation. As such, understanding the dynamics of these systems is critical.

What is needed is a technique that combines the speed of the MC approach with the precision of a direct N-body integration. In this paper, we describe a new code, the *Rapid and Precisely Integrated Dynamics* (RAPID) code, which combines both methods into a "best of both worlds" approach. In this method, the majority of particles are modeled with our parallel Hénon-style code, the Cluster MC (CMC) code (Pattabiraman et al. 2013), while the orbits of BHs are integrated directly with the Kira N-body integrator (Portegies Zwart et al. 2001). We find that this technique accurately reproduces the core radii and BH dynamics of a full direct N-body integration, with similar runtimes to the MC approach. Although we only integrate the BH orbits directly in the current work, the method is general, allowing us to select any population of particles in the cluster for N-body integration.

In Sect. 2, we briefly review the N-body and MC approaches, and describe the combination of the two approaches as implemented in RAPID code. In Sect. 3, we describe a single RAPID timestep, illustrating the technical details of the approach, while in Sect. 4, we describe the parallelization strategy that allows us to compute particle positions and velocities via orbit sampling and direct N-body simultaneously. In Sect. 5, we show the results of an analytic toy model, comparing the inspiral due to dynamical friction of a single particle as predicted by theory, direct N-body, and RAPID. Finally, in Sect. 6, we compare the properties of four idealized GCs as modeled by NBODY6, CMC, and RAPID. Throughout the paper, we will frequently refer to the "stars" and "BHs" in the cluster separately. In our current method, the stars are modeled with CMC and the BHs are integrated with Kira. This shorthand is to delineate which systems are being modeled by which technique, even though the particles under consideration are point-mass particles.

2 Hybridization approach
In this section, we provide a brief overview of the current methods employed to model GCs, and describe how our approach combines the virtues of both methods. Both the N-body and MC approaches are the result of decades of precision work by multiple groups. For a more comprehensive description of collisional N-body dynamics, see Aarseth (2003) or Dehnen and Read (2011). A review of MC methods can be found in Freitag (2008).

It should be noted that RAPID is not the first attempt at a hybrid N-body/statistical sampling approach to stellar dynamics. In particular, the hybrid approach developed by McMillan and Lightman (1984b) combined a Fokker-Planck sampling code with a direct N-body approach, in order to study GCs undergoing core collapse (McMillan and Lightman 1984a; McMillan 1986). The RAPID code continues this tradition of attempting to "have it all", by combining the best of the direct integration and statistical sampling methods.

2.1 Direct *N*-body integration

The physical principle behind a direct N-body integrator is simple: since the force on any given particle in the sum of the gravitational force from every other particles in a given system, the most accurate way to model such a system is to numerically sum all the forces. This is the underlying principle behind the N-body integrators. The most frequently used of these codes, the NBODY series of codes, have been improved and finely tuned with additional physics, including stellar evolution (Hurley et al. 2001), algorithmic regularization (Aarseth 1999), post-Newtonian chain regularization (Aarseth 2012), and GPU acceleration (Nitadori and Aarseth 2012). With advanced hardware and a minimal number of simplifying assumptions, direct integration is the most precise method available for modeling dense stellar systems.

However, this precision comes at a cost. Naively, the cost of an N-body integration scales as N^2, since one must evaluate the force of every particle on every other particle, every timestep. In practice, most modern N-body codes do not evaluate the force between every particle every timestep, instead opting for a variable "block" timestep approach in which only certain particles have their forces re-evaluated at a given time. Despite this, and many other algorithmic improvements (such as employing a nearest neighbor scheme to accelerate force evaluations), the computational cost to integrate a cluster of N particles forward by a given physical time scales as $\mathcal{O}(N^2)$, regardless of the timestep scheme or mass distribution of the cluster (Makino and Hut 1988).

This steep scaling makes large-scale simulations of massive star clusters exceedingly challenging. The largest simulation attempted with NBODY6 is currently the $N = 5 \times 10^5$ model of galactic GC M4 performed by Heggie (2014), requiring 2.5 years on a dedicated GPU system. More recently, the current state-of-the-art parallelized code NBODY6++GPU (Wang et al. 2015) can model a realistic ($N = 10^6$) cluster in little more than a year (Wang et al. 2016). Despite these remarkable achievements, simulation times in excess of ~ 1 year for large systems preclude any reasonable exploration of the parameter space of initial conditions of GCs, and any collisional models of GN ($N = 10^7$–10^9) remain beyond the capabilities of the current generation of direct summation techniques. To answer astrophysical questions related to such systems, a more rapid technique is called for.

For our hybrid approach, we use the Kira N-body integrator, included as part of the Starlab software package (Portegies Zwart et al. 2001). Like the NBODY series of codes, Kira is a 4th-order Hermite predictor-corrector integrator with a block timestep scheme. Kira also integrates close encounters and tightly-bound multiples using Keplerian regularization, where sufficiently-isolated hyperbolic and tightly bound binaries are evolved as analytic two-body systems. Additionally, Kira organizes its internal data using easily-modifiable C++ class structures, and includes an easily-customizable module for including an external gravitational potential. These two features make it ideal for inclusion in the hybrid method.
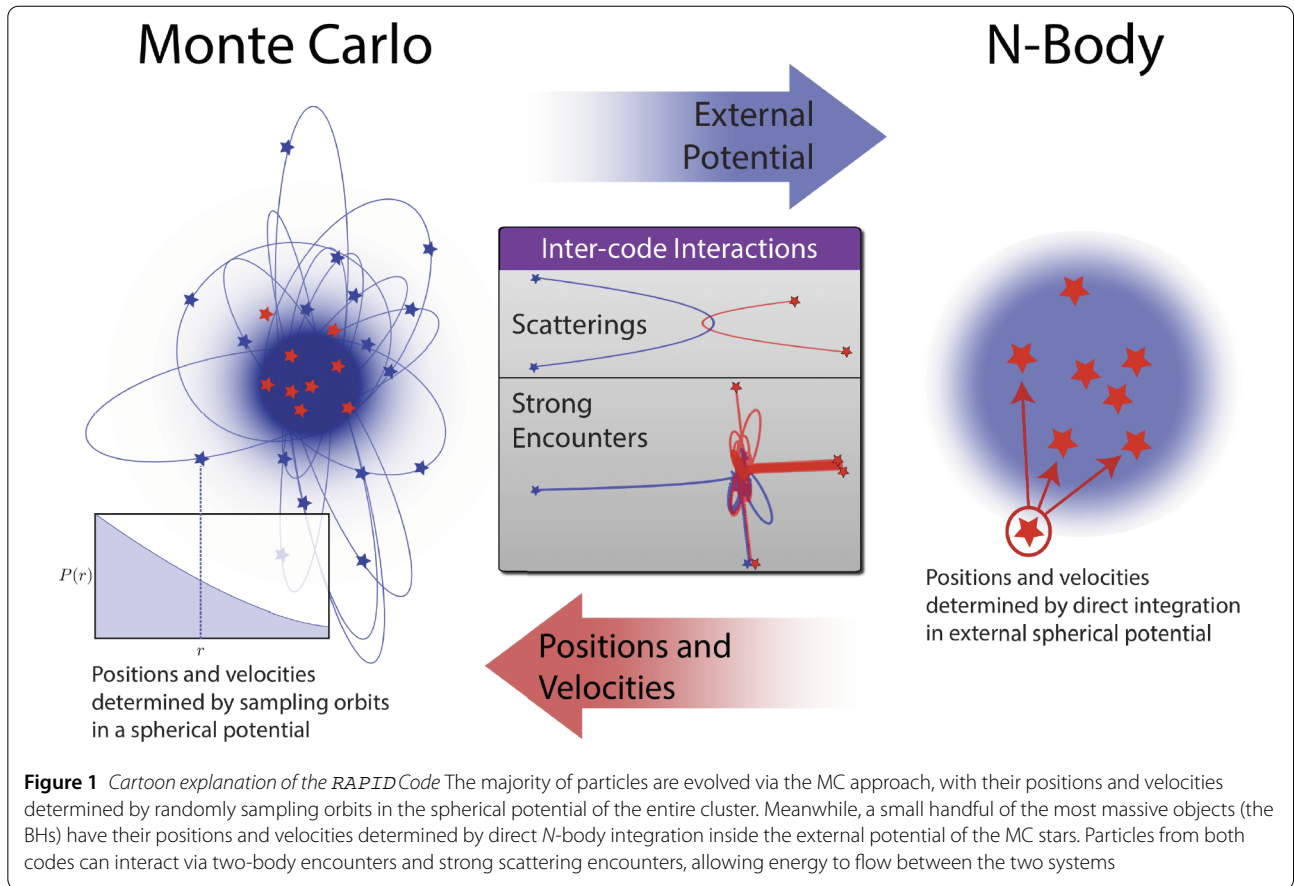
2.2 Orbit-sampled Monte Carlo

The MC approach assumes that the large-scale evolution of a star cluster can be modeled as slow transitions from one equilibrium configuration to the next. These transitions are driven by two-body scatterings (relaxations) between particles in the cluster. Because it is the statistically-predictable effect of many of these two-body scatterings that let energy flow through the cluster, one can describe the cumulative effect of these relaxations as a single, effective two-body encounter. We often describe these systems in terms of their dynamical timescales (the crossing time for a single particle) and their relaxation timescales (the average time for the velocity of a single particle to change by a certain amount). For a system in dynamical equilibrium:

$$T_{\text{relax}} \approx \frac{0.1N}{\log N} T_{\text{dyn}} \gg T_{\text{dyn}}, \qquad (1)$$

where N is the number of particles (see Binney and Tremaine 2008). With a sufficiently large N, a star cluster can be thought of as a series of independent orbits that only change on a T_{relax} timescale. This assumption eliminates the need to directly compute the forces upon a single particle on an orbital timescale. Instead, we only need to determine the shape of a star's orbit after it has changed velocity due to encounters with other stars on the relaxation timescale.

For some cases. such as a spherically asymmetric mass distribution, the orbit of the particle must be integrated numerically (e.g., Vasiliev 2014; Vasiliev et al. 2015); however, for most applications to large collisional star clusters (such as GCs and GN) the background gravitational potential can be assumed to be spherical. This allows the clever theorist to determine a star's position and velocity by analytically sampling a random point along its orbit. This *orbit-sampling MC* approach, first developed by Hénon (1971) and built upon by multiple groups

Figure 1 *Cartoon explanation of the `RAPID` Code* The majority of particles are evolved via the MC approach, with their positions and velocities determined by randomly sampling orbits in the spherical potential of the entire cluster. Meanwhile, a small handful of the most massive objects (the BHs) have their positions and velocities determined by direct N-body integration inside the external potential of the MC stars. Particles from both codes can interact via two-body encounters and strong scattering encounters, allowing energy to flow between the two systems

(Stodoikiewicz [1982](#); Giersz [1998](#); Joshi et al. [2000](#); Freitag and Benz [2001](#)) can model stellar systems with $N \gtrsim 10^7$ particles in a fraction of the time of a direct N-body simulation. Unlike a direct N-body integration, the orbit calculation and dynamical encounters in the MC method scale linearly with the number of particles; only the sorting of particles by radius, with its characteristic $N \log N$ complexity, limits the scaling. Furthermore, the MC method computes the interactions of particles on a relaxation timescale, as opposed to the dynamical timescale of a direct N-body integration. Put together, the computational difficulty of the MC method scales as $\mathcal{O}(N \log N)$ per half-mass relaxation time, versus $\mathcal{O}(N^3)$ for a direct N-body approach. Because of this, the MC method can easily model large systems that are simply beyond the reach of other techniques.

However, the assumptions that enable the speed of the MC method can easily break down in some of the most interesting regions of parameter space. Once mass segregation is complete, the evolution of a GC is largely determined by the small number of BHs that have accumulated in the core. This can consist of as few as hundreds or even tens of BHs. Since the dynamics of these small, spherically asymmetric systems change rapidly on an or-

bital timescale, the MC method is unable to accurately follow the evolution of these BHs in the center of the cluster. And since this small cluster of BHs forms the hard binaries whose binding energy acts as a power source for the entire cluster, their dynamics *must* be accurately modeled to understand the long-term evolution of the cluster.

Our orbit-sampling Cluster MC code, `CMC`, was first developed by Joshi et al. ([2000](#)), based on the original developments by Hénon ([1971](#)) and Stodoikiewicz ([1982](#)). As the code considers interactions between individual stars, `CMC` incorporates multiple physical processes, including stellar evolution (Hurley et al. [2000](#); Hurley et al. [2002](#)), strong three-body and four-body scatterings with the small-N integrator `Fewbody` (Fregeau et al. [2004](#)), probabilistic three-body binary formation (Morscher et al. [2013](#)), and physical collisions. Additionally, `CMC` has recently been parallelized to run on an arbitrary number of computer processors (Pattabiraman et al. [2013](#)). This MPI parallelizaion makes `CMC` an ideal code base for `RAPID`, as the current parallelization scheme can be easily expanded to allow the N-body integration to run in parallel to the MC
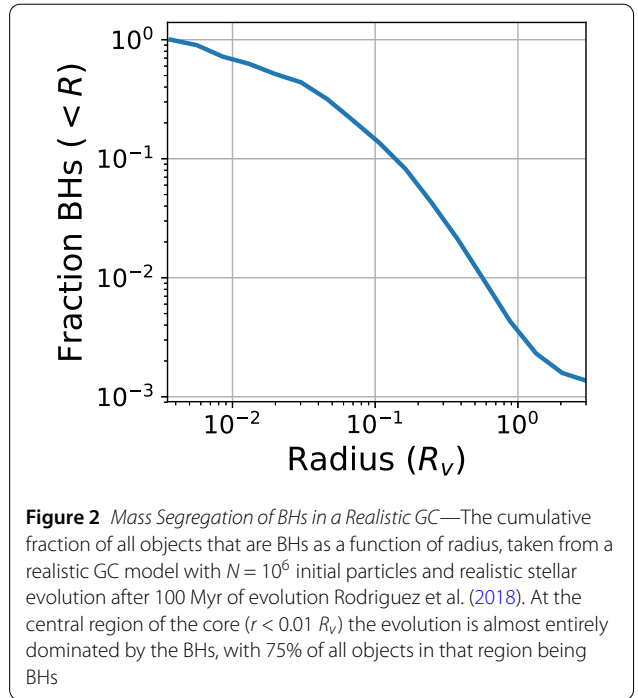
2.3 Hybrid partitioning

The cornerstone of any hybrid modeling technique is the domain decomposition between methods. Since it is the BHs that are driving the non-spherical, non-equilibrium dynamics in the cluster core, the natural division is for Kira to integrate the BHs while CMC integrates all the remaining stars in the cluster. By default, RAPID divides the system using one of two criteria:

- a *mass criterion*, which divides the system according to a specified threshold, where particles above the threshold are considered BHs and particles below it are considered stars, and
- a *stellar evolution criterion*, in which objects identified as BHs by stellar evolution are integrated by Kira, and all other objects are integrated by CMC.

By default, SPRCORRmiscellaneous0SPRCORR employs the first criterion for point-mass simulations (with a user-specified threshold mass), and the second criterion for simulations using stellar evolution. Any mixed objects (e.g., a BH-star binary) are evolved in CMC, in order to treat the binary stellar evolution consistently.

There are two reasons to focus on BHs in our hybridization scheme. The first is that by limiting the integration to a persistent set of particles, we can avoid the the large communications overhead that is incurred each time particle must be transferred back and forth from MC to N-body. This would occur much more frequently if, for instance, we divided our computational domains according to radius, with the N-body integrating particles in the core, and the MC integrating particles in the halo (similar to McMillan and Lightman 1984b). Secondly, by limiting the N-body to only BHs, we sidestep the difficulties of treating binary stellar evolution during the N-body integration. Although Kira includes a built-in package for binary and single stellar evolution (the SeBa package), it is not compatible with the stellar evolution in CMC (the Binary Stellar Evolution of Hurley et al. 2002). We will explore ways to integrate self-consistent stellar evolution into the hybrid approach in a future work.
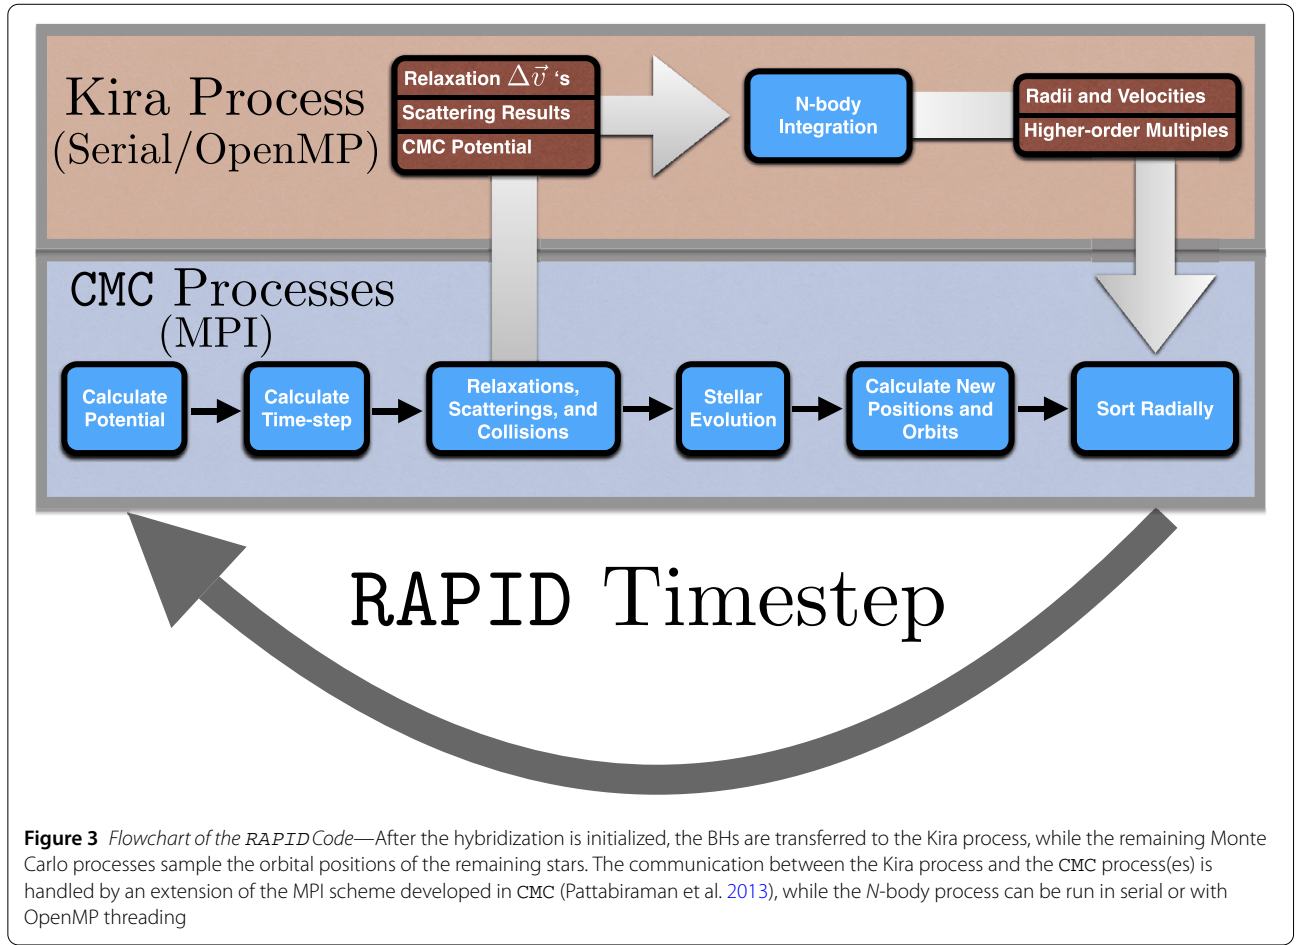
However, in realistic clusters, we find that the segregation between BHs and stars is extreme, with the innermost regions of the cluster completely dominated by BHs. In Fig. 2, we show the cumulative fraction of BHs as a function of cluster radius for a typical GC model with $N = 10^6$ and full stellar evolution (see Rodriguez et al. 2018). After 100 Myr, the central region of the cluster is completely dominated by BH, with 75% of the objects less that 0.01 pc from the cluster center being BHs. These are the objects that primarily participate in the dynamical formation of binaries that drive the cluster evolution (Breen and Heggie 2013; Morscher et al. 2015). Furthermore, any non-spherical effects that arise from having a small number of particles in the cluster center will be limited to these central BHs, ensuring that the hybrid approach can correctly integrate the correct 3D potential in the central regions.

Figure 2 *Mass Segregation of BHs in a Realistic GC*—The cumulative fraction of all objects that are BHs as a function of radius, taken from a realistic GC model with $N = 10^6$ initial particles and realistic stellar evolution after 100 Myr of evolution Rodriguez et al. (2018). At the central region of the core ($r < 0.01$ R_v) the evolution is almost entirely dominated by the BHs, with 75% of all objects in that region being BHs

The RAPID code builds upon the CMC parallelization described in Pattabiraman et al. (2013), and is designed to be run on a distributed computational system with at least 2 parallel MPI processes: one for the MC integration, and one for the N-body integration. An initial RAPID run begins with all objects integrated with CMC on all available processes. After a user-specified criterion is met, a single MPI process initializes the Kira N-body integrator. At this point, any stars that are on that CMC process are transferred to the remaining CMC process(es), and the BHs are collected on the Kira process for the N-body integration. Once both integrators have been initialized, particles can be transferred back and forth between the Kira and CMC processes (e.g., a star becomes a BH through stellar evolution, and is copied from CMC to Kira). Additionally, information about the particles must frequently be communicated back and forth between CMC and Kira. At each timestep, CMC needs to know the positions and velocities of the BHs, while Kira needs to know the external potential of the stars. See Sect. 4 for the details of the parallelization strategy. At the end of each timestep, the information from Kira and CMC is combined and printed to file, in order to create a coherent cluster model. See Fig. 1.

3 RAPID timestep

At its core, the main difference between RAPID and CMC is the computation of the orbits and positions of the BHs. Here, we describe in detail a single RAPID timestep, highlighting the differences between the hybrid approach and a standard CMC integration. See Fig. 3.

Figure 3 *Flowchart of the* `RAPID` *Code*—After the hybridization is initialized, the BHs are transferred to the Kira process, while the remaining Monte Carlo processes sample the orbital positions of the remaining stars. The communication between the Kira process and the `CMC` process(es) is handled by an extension of the MPI scheme developed in `CMC` (Pattabiraman et al. 2013), while the *N*-body process can be run in serial or with OpenMP threading

3.1 Compute the potential

Since the cluster is assumed to be spherically symmetric, the gravitational potential at radius r between stars k and $k + 1$ can be expressed as:

$$\Phi(r) = G\left(-\frac{M_k}{r} - \sum_{i=k+1}^{N} \frac{m_i}{r_i} \right), \qquad (2)$$

where $M_j \equiv \sum_{i=1}^{j} m_i$. As the stars are radially sorted, the potential at each star can be computed recursively with a single pass from the outermost star inwards:

$$\Phi_{N+1} = 0,$$

$$M_N = \sum_{i=1}^{N} m_i,$$

$$\Phi_k = \Phi_{k+1} - GM_k\left(\frac{1}{r_k} - \frac{1}{r_{k+1}} \right), \qquad (3)$$

$$M_{k-1} = M_k - m_k,$$

where N is the number of particles, m_i and r_i are the mass and radius of particle i, M_i is the total mass of particles interior to and including particle i. Note that the outer boundary conditions at $N + 1$ are not associated with physical particles, but are the outer boundary points of the cluster (with $r_{N+1} = \infty$).

In the `RAPID` code, two potentials are calculated: the full spherical potential, Φ, of all stars and BHs, and a MC-only potential, Φ^{MC}, computed only with the stars in `CMC`. The MC potential is sent from the `CMC` processes to the N-body process, and used as an external potential for the Kira integration.

3.2 Select the timestep

Since `CMC` and `RAPID` consider multiple physical processes, the timestep selection must consider multiple relevant timescales: the timescale for two-body relaxation (T_{rel}), the timescales for binary-single and binary-binary strong scatterings (T_{BS} and T_{BB}), the timescale for physical collisions (T_{coll}), the timescale for stellar evolution to contribute a change in an objects mass (T_{SE}), and the timescale for tidal stripping of stars to alter the cluster mass (T_{tid}).

The simulation timestep is then selected to be the minimum of each of these timescales, or

$$\Delta T_{CMC} = \min(T_{rel}, T_{BS}, T_{BB}, T_{coll}, T_{SE}, T_{tid}). \qquad (4)$$

The `RAPID` timestep is chosen in a similar fashion. The timescale for each physical process is computed for all particles in the cluster (stars and BHs), and the minimum (or a fraction of the minimum) is selected as the current timestep. In Kira, this timestep determines how many dynamical times the N-body system will be advanced.

Unlike `CMC`, `RAPID` can produce BH multiples with an arbitrary number of components. To compute the interaction timescale for scatterings between stars and BH multiple systems, the timestep is chosen using the same prescription as the T_{BS} and T_{BB} timesteps, but with the semi-major axis of the outermost binary pair as the effective width of the system.

The relaxation timestep, T_{rel}, is selected as the minimum of

$$T_{rel} = \frac{\theta_{max}}{\pi/2} \frac{\pi}{32} \frac{v_{rel}^3}{\log(\gamma N)G^2 n(m_1 + m_2)^2} \qquad (5)$$

for all pairs of neighboring particles (Freitag and Benz 2001), where v_{rel} is the relative velocity of the two stars, n is the local number density of stars. θ_{max} is the maximum-allowed scattering angle for two-body relaxation. Quantities which are taken as local averages (such as the number density, the velocity dispersion, etc.) are all computed using an averaging kernel of 40 particles. In other words, to compute the density of stars around star r_i, we average the density over all particles from r_{i-20} to r_{i+20}.

For the standard definition of the relaxation time from Binney and Tremaine (2008), $\theta_{max} = \pi/2$. However, for the systems considered here (particularly the highly-idealized two-component models presented in Sect. 6), this averaging can sometimes smooth out the otherwise short relaxation times between a heavy object and a neighboring lighter object (particularly if there is only one heavy object in that bin of 40 particles). For the theoretical comparison shown in Sect. 5, we still set θ_{max} to the theoretical value of $\pi/2$, but average the above quantities over the closest 2 particles. For the numerical comparison in Sect. 6, we average over the nearest 40 particles, but use $\theta_{max} = 1$ to calculate the relaxation timestep, which was found (Fregeau and Rasio 2007) to provide a good compromise between accuracy and speed for such two-component systems.

3.3 Perform dynamical interactions
After selecting the global simulation timestep, the `CMC` processes apply the standard nearest-neighbor interactions (two-body relaxations, strong encounters, and collisions) to successive pairs of particles in the radially sorted array of all cluster particles. The interactions are identical to those employed in our previous `CMC` studies. What differs in the `RAPID` approach is how the outcomes of interactions between `CMC` stars and Kira BHs are handled. Specifically:

- *Star-Star* interactions are handled in the same fashion as in a pure-`CMC` integration (see Pattabiraman et al. 2013 and references therein).
- *Star-BH* interactions are also handled in the same fashion (by `CMC`); however, in addition to updating the local BH information stored in `CMC`, the dynamical changes to each particle are communicated back to the Kira process once the interaction step is complete.
- *BH-BH* nearest-neighbor interactions are skipped, since such encounters will be performed with greater accuracy in the N-body integration.

In `CMC`, two-body relaxations are performed by setting nearest-neighbor stars along randomly selected hyperbolic orbits that are consistent with their radial and tangential velocities. The hyperbolic encounter modifies the velocities of both stars, allowing for an exchange of energy and angular momentum. At the end of the timestep, we communicate the full 3D velocity changes for each BH to the Kira process. The velocity of each particle in the N-body is updated by adding the full-3D velocity perturbation to previous 3D velocity of the BH.

In addition to two-body relaxations, `CMC` integrates strong scattering encounters between neighboring multiples (such as binary-single neighbors or binary-binary neighbors) using the `Fewbody` small-N integrator (Fregeau et al. 2004). In `RAPID`, we also allow for strong encounters between neighboring stars and BHs. However, as Kira can produce BH higher-order multiple systems, (triples, quadruples, etc), we have modified `Fewbody` to perform scatterings between systems of arbitrary multiplicity, such as single-multiple and binary-multiple encounters. We ignore multiple-multiple scatterings, since `CMC` does not track higher-order stellar triples, and BH-BH encounters are performed naturally in Kira.

Once `Fewbody` has completed the scattering (which can take several seconds of CPU time for compact higher-order multiple systems) the output is then sent back to `CMC` and Kira separately. For higher-order multiples, the full 3D position and velocity of each BH component, relative to the multiple center-of-mass, is sent to the Kira process. The multiple is then reinserted into the N-body integration at its previous position.

The only exception to this procedure is the formation of mixed-multiple systems (a binary or higher-order multiple with both BH and stellar components). We evolve any BH-Star binaries in `CMC`. For higher-order multiples with both stellar and BH components, the multiple is hierarchically broken apart into smaller components. Any star or BH-Star is evolved by `CMC`, while any BH, binary BH, or

Table 1 The model names, initial conditions, and runtimes for each of the four idealized GC models. The runtimes quoted are the relevant walltimes for each method. We list the single runtime for each `NBODY6` run, and the mean and standard deviation of the 10 `CMC` and `RAPID` runs for each cluster. The `CMC` and `RAPID` models were run on 4 Intel Xeon E5-2670 Sandy Bridge processors, while the `NBODY6` models were run on 8 processors and 1 Nvidia Tesla M2090 GPU

Initial Conditions				Runtimes (Hours)		
Model Name	M_{BH}/M_{star}	m_{BH}/m_{star}	T_{end} (t_{dyn})	CMC	NBODY6	RAPID
64k-0.01-10	0.01	10	14,000	5.0 ± 1.7	425	1.8 ± 0.2
64k-0.02-10	0.02	10	20,000	8.4 ± 0.9	304	2.6 ± 0.3
64k-0.01-20	0.01	20	10,000	2.6 ± 0.4	554	1.3 ± 0.2
64k-0.02-20	0.02	20	20,000	5.3 ± 0.7	512	2.7 ± 0.6

BH multiple is evolved by Kira. The kinetic energies of the newly-broken components are adjusted to ensure conservation of energy. While this limits the modeling of BH-non BH systems (such as low-mass X-ray binaries), these systems predominantly form with low mass BHs at the outer region of the BH subsystem, where mixing between stars and BHs is more common (Kremer et al. 2018b). Because of this, these systems are less likely to participate in the strong three-body encounters that form BH binaries in the central regions of the cluster.

3.4 Perform stellar evolution
`RAPID` considers realistic stellar evolution using the Single Stellar Evolution (SSE) and Binary Stellar Evolution (BSE) packages of Hurley et al. (2000, 2002). This is identical to the previous implementation in `CMC`. No stellar evolution is required by the direct N-body, since the only particles integrated by Kira are BHs. As stated above, all mixed BH-Star objects are integrated in `CMC`. This is to ensure that the binary stellar evolution for BH-star systems is performed consistently.

3.5 Calculate new orbits and positions
After the dynamical information for each particle has been updated, a new orbital position and velocity, consistent with the particle's new energy and angular momentum, must be computed. Since the dynamical state of each particle is up-to-date in `CMC` and Kira, the MC and N-body integrations can be performed in parallel by their respective processes.

3.5.1 Orbit calculation (MC)
Orbits in `CMC` are determined using the standard Hénon-style orbit-sampling approach. We assume that the orbit is entirely a function of the star's kinetic energy and the spherical potential of the cluster. We first constrain the radial extent of the orbit by computing the two zeros (r_{min} and r_{max}) of the energy equation:

$$2E - 2\Phi(r) + J^2/r^2 = 0. \tag{6}$$

Then, we probabilistically select a new radius for the star based on the star's new orbit. Since the probability of finding the star at a given radius is proportional to the time the

particle spends at that radius, we can express the probability of finding the star at a specific radius as

$$P(r)dr = \frac{dt}{T} = \frac{dr/|v_r|}{\int_{r_{min}}^{r_{max}} dr/|v_r|}. \tag{7}$$

We then draw a random sample from $P(r)$, and use the value as the particle's position for the next timestep.

3.5.2 Orbit calculation (N-body)
Because we have dynamically perturbed the BHs though scattering and a new external potential, the gravitational force and its higher derivatives must be recalculated before resuming the Kira integration. Otherwise, the dynamical changes to the BH velocities would produce discontinuities in the higher derivatives of the force, breaking the smoothness needed for the 4th-order Hermite integrator to work. The N-body system is reset using Kira's built-in reinitialization function, which recomputes the acceleration and jerk for each BH explicitly. The position and velocities of all the particles (up to any changes from dynamical interactions) are not modified.

Once the system has been reinitialized, the N-body configuration is directly integrated with the Kira integrator for a number of dynamical times equal to the current `RAPID` timestep. For each particle, the total force is computed as the sum of the external force, from Φ^{MC}, and the internal force computed via direct summation of the other BHs. We apply the external `CMC` potential to each isolated BH and the center-of-mass of each bound multiple system by computing the acceleration and jerk from the external potential

$$\vec{a}_{ext} = -\frac{\partial \Phi^{MC}}{\partial r}\hat{r},$$

$$\vec{j}_{ext} = -\frac{1}{r}\frac{\partial \Phi^{MC}}{\partial r}(\vec{v} - (\vec{v} \cdot \hat{r})\hat{r}) - \frac{\partial^2 \Phi^{MC}}{\partial r^2}(\vec{r} \cdot \hat{r})\hat{r}$$

and adding this to the acceleration and jerk of each BH using Kira's external potential module. To simplify the calculation of the potential and its derivatives, we select 30 stars evenly in $\log r$ from the innermost `CMC` star to the stripping radius of the N-body simulation, and copy their radii

and Φ^{MC} values to the N-body integrator. The values of Φ^{MC} and its radial derivatives are then calculated via 5th-order Lagrange Polynomial interpolation. This technique is inspired by the radial potential sampling used in Vasiliev (2014).

3.6 Sort radially
Finally, once all the relevant physics has been applied and the data collected on the CMC processes, the particles must be sorted in order of increasing radial distance from the GC center. The sorting is performed in parallel by all CMC processes using the parallel Sample Sort algorithm described in Pattabiraman et al. (2013).

4 Parallelization strategy
To incorporate the Kira integrator into CMC, we make the following modifications to our parallelization strategy. When the simulation starts, all particles are evolved using the CMC scheme. As described in Sect. 2.3, the activation criterion for the N-body integrator depends on the type of particles being integrated. For point particle simulations, the N-body integrator is begun immediately, whereas for star clusters modeled with stellar evolution, the N-body integrator is only activated once a certain number (~ 25) of BHs have been formed. Once the activation condition is met, we divide the entire set of particles into two sets: MC stars and N-body BHs.

At the same time, we divide the existing set of p processes into two separate groups: a single Kira process for integrating the BHs, and $p - 1$ CMC processes for integrating the stars. Any MPI communication is handled by two custom intracommunicators: one corresponding to all p processes and one restricting communication to the $p - 1$ CMC processes. The latter allows us to employ the same parallelization strategy described in Pattabiraman et al. (2013) with minimal modification. When the processes are split, all the BHs are sent to the Kira process, where their coordinates are converted from (r, v_r, v_t) space to the full 6-D phase space by randomly sampling the orientation of the position and velocity vectors. This sets the initial conditions for the N-body integration. In addition, the Kira process maintains a radially-sorted array of (r, v_r, v_t) for all BHs. This is done to facilitate easier communication with CMC and to minimize the required MPI communication between CMC and Kira every timestep. Although the CMC processes hand over their BHs to the Kira process, the BHs are not deleted from their local arrays. The BHs are left intact yet inert, so that their positions and velocities can be updated upon completion of the N-body integration.

Once the CMC and Kira MPI processes have been initialized, the hybrid method must allow both codes to interact while minimizing the amount of MPI communication. This is accomplished by a series of intermediate arrays, designed to store and transmit the minimum amount

of information back and forth between the CMC and Kira. Communication only occurs at two points during a RAPID timestep. The first occurs after the relaxation and strong-encounters have been performed, in order to communicate the dynamical changes from CMC to Kira. The second occurs after both systems have completed their respective orbit computations, to communicate the new dynamical positions and velocities from Kira back to CMC.

4.1 CMC to Kira communication
After the CMC processes have computed a new potential, performed two-body relaxations, and any strong encounters, the dynamical changes to the BHs must be communicated to the Kira process. This is done with three distinct communications:

- The *MC potential*, Φ^{MC}, is sent to the Kira process as two arrays, containing the radius and cluster potential of every star (excluding the BHs) in CMC. The Kira process selects 30 of these stars as described in Sect. 3.5.2, and passes the information to the Kira integrator to use when computing the external force.
- The *two-body relaxations* are communicated as an array of objects, each containing a particle ID and a 3-dimensional $\Delta\vec{v}$. These weak velocity perturbations are added to the single BHs and the centers-of-mass of any BH multiple systems before the N-body system is reinitialized by Kira.
- The results of *strong encounters* are communicated differently depending on the type of encounter. For binary BHs that have experienced a strong encounter with a star, only the change in semi-major axis and eccentricity are communicated back to Kira. For triples and higher-order multiples, the full position and velocity of every BH in the multiple is communicated to Kira. To reduce communication, any hierarchical information is not transmitted, and the Kira process reconstructs the hierarchy locally before the N-body system is reinitialized.

Since all the information previously described does not drastically change the radial positions or velocities of the particles in the N-body (by assumption, the MC approach requires that $\Delta v/v \ll 1$), the dynamical state of the N-body system is preserved between RAPID timesteps. The one exception is strong encounters in which a single bound BH multiple is broken into components. Since strong encounters in CMC are performed by assuming the scatterings are isolated from the cluster potential at infinity, the resultant components cannot be placed at the same infinite location in the N-body system. For such systems, the components are placed at the correct radius and random orientations, similar to the initial transfer of BHs from CMC.

Finally, we allow for the possibility that CMC may create new BHs, either through stellar evolution or through strong encounters which produce single or binary BHs. We

add any such new BHs to Kira. The new BH is then flagged as a BH in the local array of the CMC process that created it, to ensure it is not evolved by CMC during the next timestep.

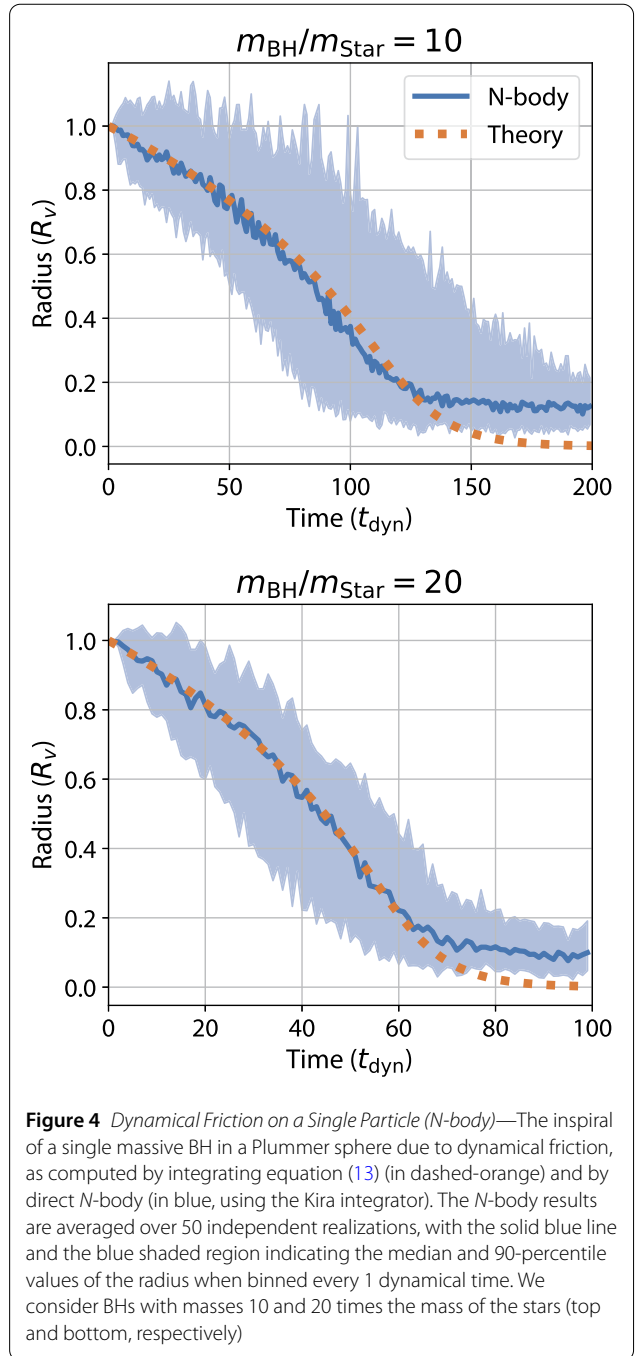### 4.2 Kira to CMC communication

After the Kira integrator has computed the orbits of the BHs, the new dynamical state must be communicated back to the CMC processes. First, the dynamical 6-D phase space information for each particle in the N-body is projected back to the reduced (r, v_r, v_t) basis, and copied back into place in the intermediate star array on the Kira process. The intermediate BH array is then divided up and communicated back to the respective CMC processes.

In addition to the positions and velocities, any new objects, such as single BHs, newly formed binaries, or higher-order multiples, are communicated as new objects to the last $(p-1)$ CMC processes, to be placed in the correct process once the CMC (and intermediate BH) arrays are sorted. For single and binary BHs, this is accomplished by sending the usual (r, v_r, v_t) for each system and the semi-major axis and eccentricity for any binaries to the CMC processes. For higher-order multiples, the full 6D dynamical information is transmitted back using the same array that sent the strong encounters from CMC in the first communication. Again, only the positions and velocities of the BH multiple components are communicated, so the hierarchical information is reconstructed on the local CMC process after communication.

Because the N-body integration is being performed in a lower-density environment than the full cluster, it is possible for Kira to produce pathologically wide binaries and higher-order multiples. This effect is particularly problematic at late times, where a handful (~ 10) of BHs can easily produce binaries with separations greater than the local inter-particle separation of stars. Since the unphysically large interaction cross-sections of these systems drastically shrink the CMC timestep, we break apart any multiple systems whose apocenter distance is greater than 10% of the local inter-particle separation of stars at that radius. The kinetic energy of the CMC stars is adjusted to ensure energy conservation. This criterion for breaking wide binaries is the same criterion used to break wide binaries produced by Fewbody in CMC.

5 Analytic comparison

To test our new method, we first choose a simplified analytic model designed to demonstrate whether RAPID can model the dynamical friction experienced by a single BH in a sea of stars. Theoretically, the time taken for a massive object to spiral in to the center of a cluster can be calculated considering the rate of change of angular momentum of the particle due to the dynamical friction experienced by

Figure 4 *Dynamical Friction on a Single Particle (N-body)*—The inspiral of a single massive BH in a Plummer sphere due to dynamical friction, as computed by integrating equation (13) (in dashed-orange) and by direct *N*-body (in blue, using the Kira integrator). The *N*-body results are averaged over 50 independent realizations, with the solid blue line and the blue shaded region indicating the median and 90-percentile values of the radius when binned every 1 dynamical time. We consider BHs with masses 10 and 20 times the mass of the stars (top and bottom, respectively)

said particle (Chandrasekhar 1943; Binney and Tremaine 2008). Consider a particle of mass m on a circular orbit at a radius r in a Plummer sphere of mass M and scale factor a (see e.g., Heggie and Hut 2003). Many properties of the Plummer sphere can be described analytically, such as the mass interior to r

$$M(r) = M\left(1 + \frac{a^2}{r^2}\right)^{-3/2} \tag{8}$$
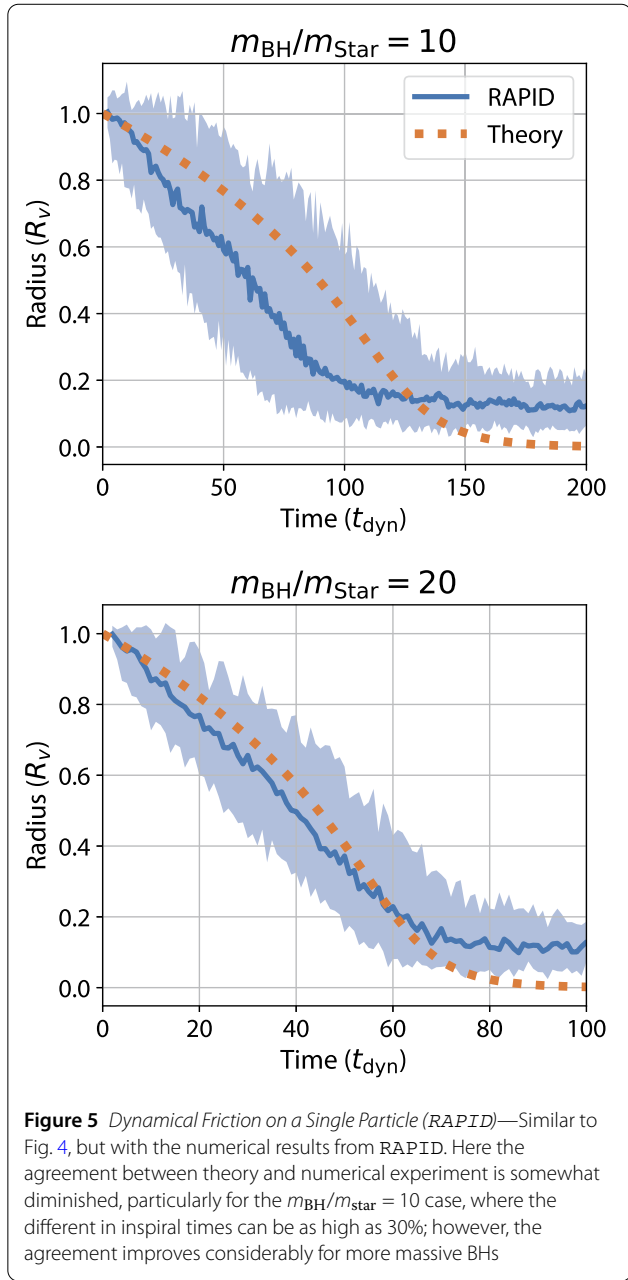
Figure 5 *Dynamical Friction on a Single Particle (`RAPID`)*—Similar to Fig. 4, but with the numerical results from `RAPID`. Here the agreement between theory and numerical experiment is somewhat diminished, particularly for the $m_{BH}/m_{star} = 10$ case, where the different in inspiral times can be as high as 30%; however, the agreement improves considerably for more massive BHs

the density at a radius r

$$\rho(r) = \frac{3M}{4\pi a^3}\left(1 + \frac{r^2}{a^2}\right)^{-2.5} \tag{9}$$

and the velocity dispersion at r

$$\sigma(r) = \sqrt{\left(\frac{GM}{a}\right)\sqrt{1 + \frac{r^2}{a^2}}}. \tag{10}$$

The specific angular momentum of our massive particle is given by $L = rV_c$, where $V_c = \sqrt{GM(r)/r}$ is the circular

velocity at radius r. The rate of change of angular momentum is then

$$\frac{dL}{dt} = \left(V_c + r\frac{dV_c}{dr}\right)\frac{dr}{dt}$$

$$= \frac{V_c}{2}\left(1 + 3\left(1 + \frac{r^2}{a^2}\right)^{-1}\right)\frac{dr}{dt}. \tag{11}$$

As the BH travels through the cluster, it experiences dynamical friction in the opposite direction of its velocity. If we assume that the change in r is sufficiently slow that the BH orbit remains circular, then the change in angular momentum due to dynamical friction is given by

$$\frac{dL}{dt}\bigg|_{\text{fric}} = a_{\text{fric}}r$$

$$= -4\pi G^2 \log \Lambda \rho(r) \chi\, m r V_c^{-2}, \tag{12}$$

where $_{\text{fric}}$ is the acceleration due to dynamical friction (Binney and Tremaine 2008, p. 645) $\chi \equiv \text{erf}(X) - 2X\exp(-X^2)/\sqrt{\pi}$, $X \equiv V_c/(\sqrt{2}\sigma(r))$, and $\log \Lambda \sim \log \gamma N$ is the Coulomb Logarithm, with N being the number of particles in the cluster and $\gamma \sim 0.01$. Setting equations (11) and (12) equal yields

$$\frac{dr}{dt} = \frac{-8\pi G^2 \log \Lambda \rho(r)\chi\, m r}{V_c^3(1 + 3(1 + \frac{r^2}{a^2})^{-1})} \tag{13}$$

which can be solved numerically for $r(t)$.

To test whether our proposed method can reproduce the analytic inspiral times predicted by Equation (13), we create 50 independent realizations of a Plummer sphere by drawing 10^4 equal-mass particles from Equation (9). We then place a single massive particle on a circular orbit at the virial radius of the cluster. We consider mass ratios between the BH and the individual stars of 10 and 20. Each cluster model is then integrated forward until the particle has settled into the center. We integrate these models forward using the Kira integrator (Fig. 3) and the new `RAPID` approach (Fig. 4).

Because `RAPID` and `CMC` are designed to model two-body relaxation by averaging various quantities over several neighboring stars, special care must be taken for `RAPID` to accurately model the behavior of a single particle. To that end, we set the maximum scattering angle to the typical value of $\theta_{\max} = \pi/2$, while we reduce the number of neighboring particles over which the quantities in Equation (5) are averaged to 2. In other words, the quantities used to compute the timestep consider only the nearest particles when computing the local two-body relaxation timescale (while the global timestep is chosen as the minimum of all nearest neighbor timescales). This ensures that the two-body relaxation timestep chosen for the

RAPID simulations can appropriately model the dynamical friction experienced by a single massive object.

As expected, the pure N-body agrees with the theoretical prediction to a better degree than RAPID. However, given the approximate nature of two-body relaxation in RAPID, the agreement presented in Fig. 5 is encouraging, and suggests that our treatment of two-body relaxation between Monte Carlo and N-body can treat dynamical friction to within $\sim 30\%$ (the largest deviation between theory and numerical results in the $m_{BH}/m_{star} = 10$ case), with substantial improvements for more massive BHs (the $m_{BH}/m_{star} = 20$ case being more representative of true BH masses in realistic clusters).

We reiterate that these results were obtained by reducing the number of particles used to compute various average quantities (Sect. 3.2). This allows the timestep to be properly calibrated for the dynamical friction of a single particle (the massive BH). In a standard CMC and RAPID run, a larger averaging kernel can be used, since for realistic clusters with a continuous mass function, there will be many massive and light objects within the inner-most 40 particles (which typically sets the minimum relaxation time for the cluster). For more idealized clusters, where the mass function is discrete and there can sometimes be only one massive BH per 40-particle kernel (especially before mass segregation), the averages must be computed carefully. In this section, we have accomplished this by setting θ_{max} to the standard value of $\pi/2$ and taking the mimimum timestep computed over averages between the nearest-neighbor particles (usually the average between the BH and its two neighboring stars). More generally, this reduction in timestep can also be accomplished by reducing the maximum angle for a two-body deflection to $\theta_{max} = 1$. We elect for the later in the next section, as it has been demonstrated to work well for the idealized two-component systems considered there (Fregeau and Rasio 2007).

As an aside, it is interesting to note that Fig. 4 provides an excellent way to test the value of γ commonly used in numerical work involving the Coulomb Logarithm ($\log \Lambda \equiv \log \gamma N$). Even small changes (such as $\gamma = 0.005$ or $\gamma = 0.02$) produce obviouly worse agreement in Fig. 3. This suggests that the value of $\gamma = 0.01$ used in many previous studies of multi-mass clusters is appropriate.

6 Numerical comparison

To test the effectiveness of this hybrid approach, we compare our code to idealized models of GCs using similar initial conditions to Breen and Heggie (2013). We considered four clusters with 65,536 point-mass particles. The majority of particles are low-mass stars, while a small fraction of particles are high-mass BHs. We considered different mass ratios between BHs and stars ($m_{BH}/m_{star} = 10$ and $m_{BH}/m_{star} = 20$). We also varied the total mass in stars and BHs ($M_{BH}/M_{star} = 0.01$ and $M_{BH}/M_{star} = 0.02$). For each

set of initial conditions, we compare 10 independent realizations of each model produced by RAPID to 10 models produced by CMC and by a single N-body model using NBODY6 (Aarseth 1999). Each model was integrated until the majority of the BHs had been ejected from the cluster. A summary of the initial conditions and the average runtimes for each approach is listed in Table 1.

Immeditally obvious from Table 1 is that the runtimes for RAPID are much shorter than the typical runtimes for NBODY6, sometimes by factors of a few hundred. Slightly more surprising is that the RAPID runtimes also tend to be shorter than the runtimes for CMC. This largely arises from the improved treatment of the cluster center in RAPID. As will be shown in the next section, RAPID reproduces the less-dense core radii of the full N-body models better than CMC, with the latter producing core radii 2–4 times smaller and more compact than NBODY6. Because the timestep of the MC is dominated by the relaxation in the densest region of the cluster, the more compact CMC models require many more timesteps to resolve the deep core collapses, resulting the longer runtimes.

6.1 Core and half-mass radii

For the purposes of this analysis, we assume that the models produced by NBODY6 are the "true" clusters, since a direct integration technique requires the fewest simplifying assumptions. We wish to know how well the models from our new hybrid technique match the models produced by the N-body approach. To that end, we focus on two typical figures of merit in star cluster simulations: the half-mass and core radii. In both cases, the radii are defined as the distance from the center of the cluster. For NBODY6, this is calculated as a density-weighted sum:

$$\vec{r}_d = \sum_{j=1}^{N} \rho_j \vec{r}_j \Big/ \sum_{j=1}^{N} \rho_j, \qquad (14)$$

where N is the total number of particles, $\vec{r}_d$ is the position of each particle, and ρ_j is the density of objects surrounding that particle (cf. Equation 15.1, Aarseth 2003). In CMC and RAPID, the center is fixed at $r = 0$ by assumption. The half-mass radius is the radius from the cluster center that encloses half the cluster mass. For the core radius, all three approaches use the approximate core radius definition from Aarseth (2003), based on the unbiased estimator developed in Casertano and Hut (1985). This takes the form of a weighted sum over distances from the cluster center out to the half-mass radius:[a]

$$r_c = \sqrt{\sum_{j=1}^{N/2} \rho_j^2 |\vec{r}_j - \vec{r}_d|^2 \Big/ \sum_{j=1}^{N/2} \rho_j^2}. \qquad (15)$$

In NBODY6, Equations (14) and (15) are calculated using the full 3D position vectors (with the densities for each
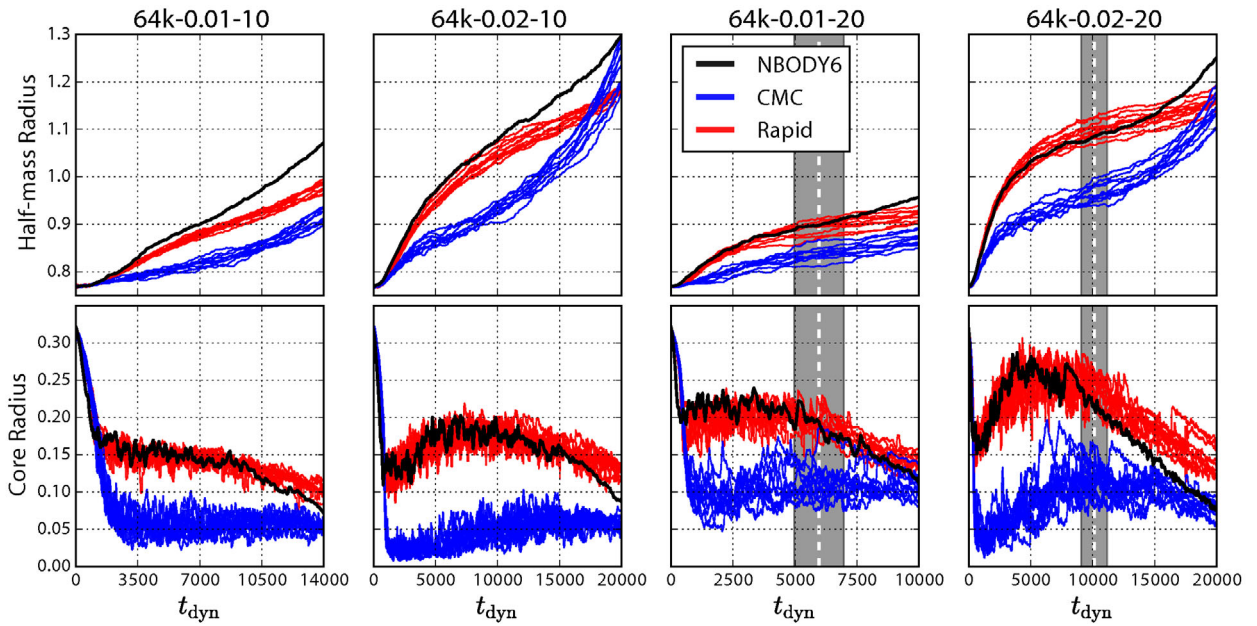
Figure 6 *Half-mass and Core Radii of Idealized Clusters*—The half-mass and core radii, averaged over 100 dynamical times, for all four models as determined by direct *N*-body (NBODY6, in black), the Monte Carlo (CMC, in blue), and the hybrid approach (RAPID, in red). We only show one realization from NBODY6, while we show 10 realizations each for CMC and RAPID RAPID reproduces both the core and half-mass radii at early times better that CMC (with the exception of the 64k-0.01-10 run, though RAPID is still much closer than CMC). On the other hand, RAPID reproduces the core radii of the full *N*-body runs far more accurately than CMC, with the core radii from NBODY6 and CMC disagreeing by up to a factor of 3. At late times, the core radii from RAPID begins to diverge from NBODY6; this occurs when the majority of the BH system has been ejected, and the CMC-controlled two-body relaxation dominates the BH dynamics. RAPID reverts to pure CMC when the number of BHs drops below a certain threshold (5 BHs, indicated by the black dashed lines, (with the gray bands indicating the 1σ variation across the 10 runs)

particle computed with respect to its three nearest neighbours), while for CMC and RAPID, Equation (15) is calculated from the 1D positions where the densities are calculated for each particle by averaging over its 40 nearest neighbors.

In Fig. 6, we compare the half-mass radius and core radius for each model as determined by CMC, NBODY6, and RAPID. While the half-mass radii are relatively consistent across all models (with a maximum deviation of ∼ 10%), the core radii are drastically different between the three methods. CMC consistently underestimates the core radius of each cluster immediately after core collapse, with the trend persisting until all BHs have been ejected from the cluster. This consistent underestimation of the cluster core radius has been observed in other studies using CMC (Morscher et al. 2015) and other orbit-sampling Monte Carlo codes. However, the core radii as determined by RAPID agree very well with the radii reported by the direct *N*-body. This suggests that the RAPID approach can correctly model the dynamics of the massive BHs which dominate the long-term evolution of GC cores.

At late times, both the half-mass and core radii predicted by RAPID begin to diverge from those determined by NBODY6. This divergence is to be expected: as the BHs

are ejected from the cluster, the orbits of individual BHs are determined less by their encounters with other BHs, but by two-body relaxation with stars controlled by the MC. Given the disagreement between the pure CMC and NBODY6, this divergence is consistent. This suggests that RAPID will be most effective when modeling systems that retain a large number of BHs. Recent work (e.g., Mackey et al. 2007; Downing 2012; Morscher et al. 2013, 2015; Kremer et al. 2018; Askar et al. 2018) has shown that the most massive GCs can retain hundreds to thousands of BHs up to the present day. Given that, RAPID should be able to correctly model realistic GCs throughout their entire evolution far more accurately than a traditional MC method.

In addition to the bulk evolution of the system, we also want to compare the behavior of individual stars in energy-angular momentum space. In Fig. 7, we plot the specific energy and angular momentum of every star and BH in a single model for the $m_{BH}/m_{star} = 20$, $M_{BH}/M_{star} = 20$ cluster. Both models start with identical initial condition, and we show the evolution of E and J as a function of cluster time. After 500 dynamical times, the point of deepest core collapse, the RAPID model lags behind the NBODY6 model, having produced only one BH binary, while the two BH bi-
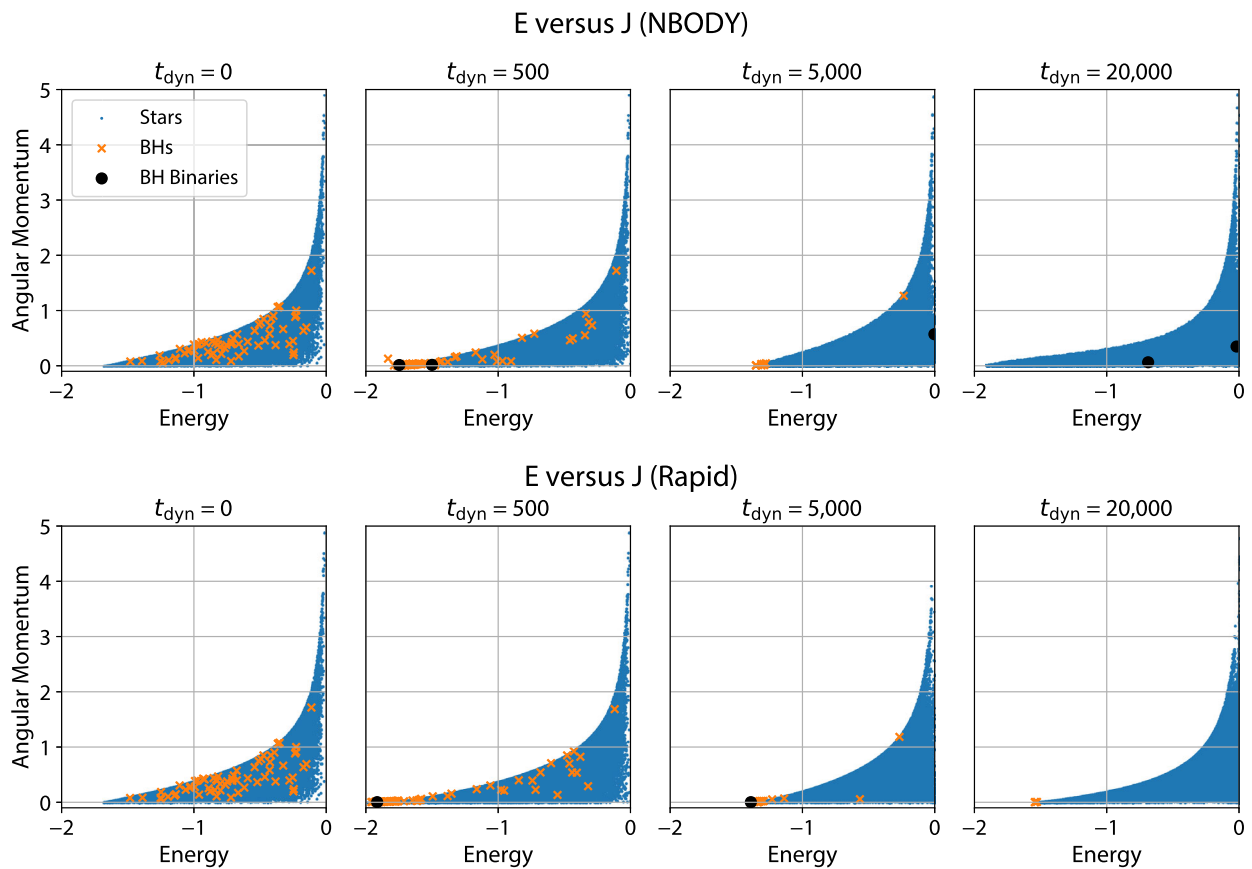
Figure 7 *Phase-space Scatterplots of Stars and BHs*—The specific energy (potential plus kinetic) versus the specific angular momenta for each particle at four different points during the evolution of the $m_{BH}/m_{star} = 20$ and $M_{BH}/M_{star} = 20$ model. We show the scatter plot of E vs J at the beginning of the simulation, at the point of deepest collapse (500 T_{dyn}), at the point of greatest core re-expansion (5000 T_{dyn}), and at the end of the simulation (20,000 T_{dyn}) after the majority of BHs have been ejected

naries in the NBODY6 cluster have started to push the stars out of the central region towards higher E and J. This is consistent with the slightly faster collapse and evolution of the core radius described in Fig. 6. After 5000 dynamical times, both models have ejected a number of BHs, creating sufficient energy to push stars out of lowest potential energy states in the central region. The final snapshot (at 20,000 dynamical times) shows the stars of the NBODY6 model in a state of deeper collapse than the RAPID model. This is most likely due to a recent encounter between the two remaining BH binaries pushing both onto a higher orbit in the cluster with a correspondingly larger E and J. The RAPID model retains 3 BHs at the final snapshot. As these remain in the cluster core, they manage to exclude the stars from occupying the lowest energy states in the cluster potential. This difference at late times is consistent with the stochastic nature of BH retention observed in more realistic cluster models.

6.2 Binary formation

What explains the substantial improvement in the RAPID core radii? In Rodriguez et al. (2016a), we explored a direct comparison between CMC and a state-of-the-art direct N-body simulation of 10^6 particles (the DRAGON simulation, Wang et al. 2016). There, we found good agreement between most of the structural parameters of the two cluster models (e.g., the half-mass radii, the formation and ejection rate of BHs, etc.). However, the one notable exception was the evolution of the inner parts of the cluster such as the core radii and the inner-most Lagrange radii (the radius enclosing a certain fraction of the cluster mass). This was especially true when considering the Lagrange radii of only the BHs (Rodriguez et al. 2016a, Fig. 7), where the inner-most few BHs would fall into a much deeper state of collapse (by nearly two orders of magnitude) into the cluster center than the equivalent radii from the N-body model. The cluster would remain in this deep state until the formation of a BH binary, which would reverse this
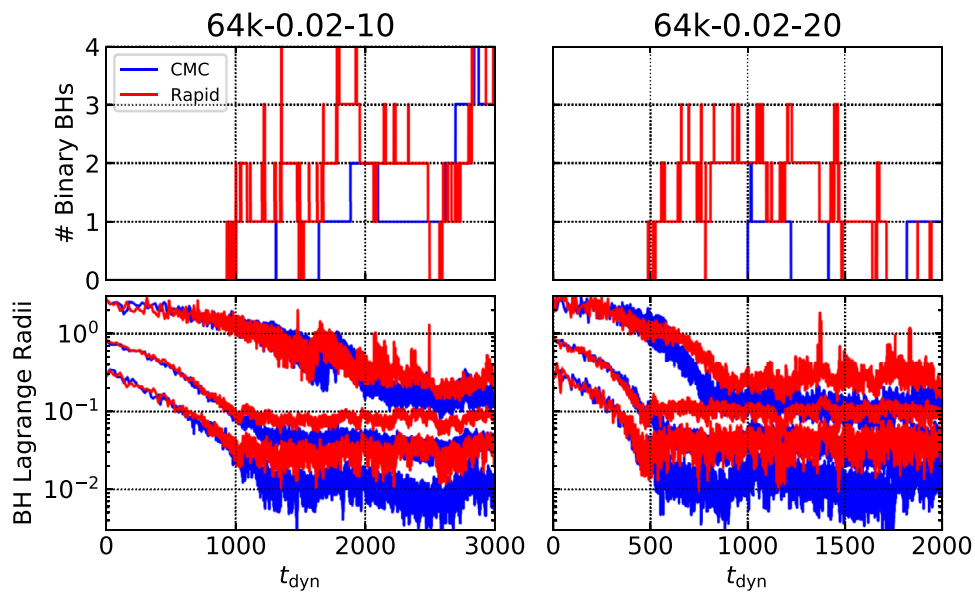
Figure 8 *BH Lagrange Radii and Binary BH Formation*—The number of binary BHs and the Lagrange radii for all BHs during the early stages of cluster collapse for two example clusters. On the top, we show the number of BH binaries as a function of time in both `RAPID` and `CMC`. On the bottom, we show the radii enclosing 10% 50% and 90% of the BHs in the cluster for the two techniques. In both cases, as the cluster collapses, `RAPID` forms a binary through three-body encounters much earlier, halting the continued collapse of the cluster. `CMC` continues to collapse until its first binaries are formed several hundred dynamical times later

deep collapse and bring the inner Lagrange radii back into agreement with the N-body results.

It was speculated that the reason for this discrepancy lay in the analytic prescription that `CMC` employs to model the dynamical formation of binaries during three-body encounters of single BHs. This prescription, from Morscher et al. (2013), may underestimate the formation rate of BH binaries, especially given that the probability of binary formation scales as v^{-9} in the local velocity dispersion. Because these interactions typically involve only a handful of objects in the cluster center, where the standard MC assumptions of spherical symmetry and $T_{rel} \gg T_{dyn}$ break down, it is not obvious that `CMC`'s statistical approach to binary formation based on locally-averaged quantities can correctly model this process.[b] This difficulty was one of the primary motivators for the development of `RAPID`: by directly integrating the BH dynamics every timestep, we can explicitly model the complicated three-body encounters between single BHs on a dynamical timescale.

In Fig. 8, we show the early stages of collapse for two typical clusters as modeled by `CMC` and `RAPID`. The Lagrange radii, indicating the radii enclosing 10%, 50%, and 90% of the BHs, are nearly identical between the two methods during the early stages of collapse (as would be expected, since both methods model dynamical friction through two-body MC relaxation). However, in both cases, the `RAPID` models dynamically forms binaries at much earlier stages of collapse, causing the inner-most BH La-

grange radii to rebound. The `CMC` models, on the other hand, reach a much deeper state of collapse before forming their first binaries.

We believe this discrepancy is responsible for the deep collapses observed in Rodriguez et al. (2016a). In those models, the most massive objects would naturally find themselves in the center, and continue to collapse until a binary was formed. This caused the deep collapses noted there, which were not reproduced in the direct N-body model. Here, the deep collapses have smoothed out to a more continuous underprediction, since the two-component models studied here have equal masses for all BHs, whereas in Morscher et al. (2015), Rodriguez et al. (2016a), it was consistently the most massive BHs decoupling from the rest of the core that were responsible for the deep collapses.

6.3 BH retention
Since the heating of the cluster is primarily driven by ejection of BHs and binary hardening in the core (Breen and Heggie 2013), it is important to compare the retention and binary production between the different methods. In Fig. 9 we show the number of BHs retained in each cluster as a function of time. We show the total number of BHs for each of the 10 `CMC` and `RAPID` models, and the number of BH binaries and BH triples present over time in the `NBODY6` model and a representative `RAPID` model.
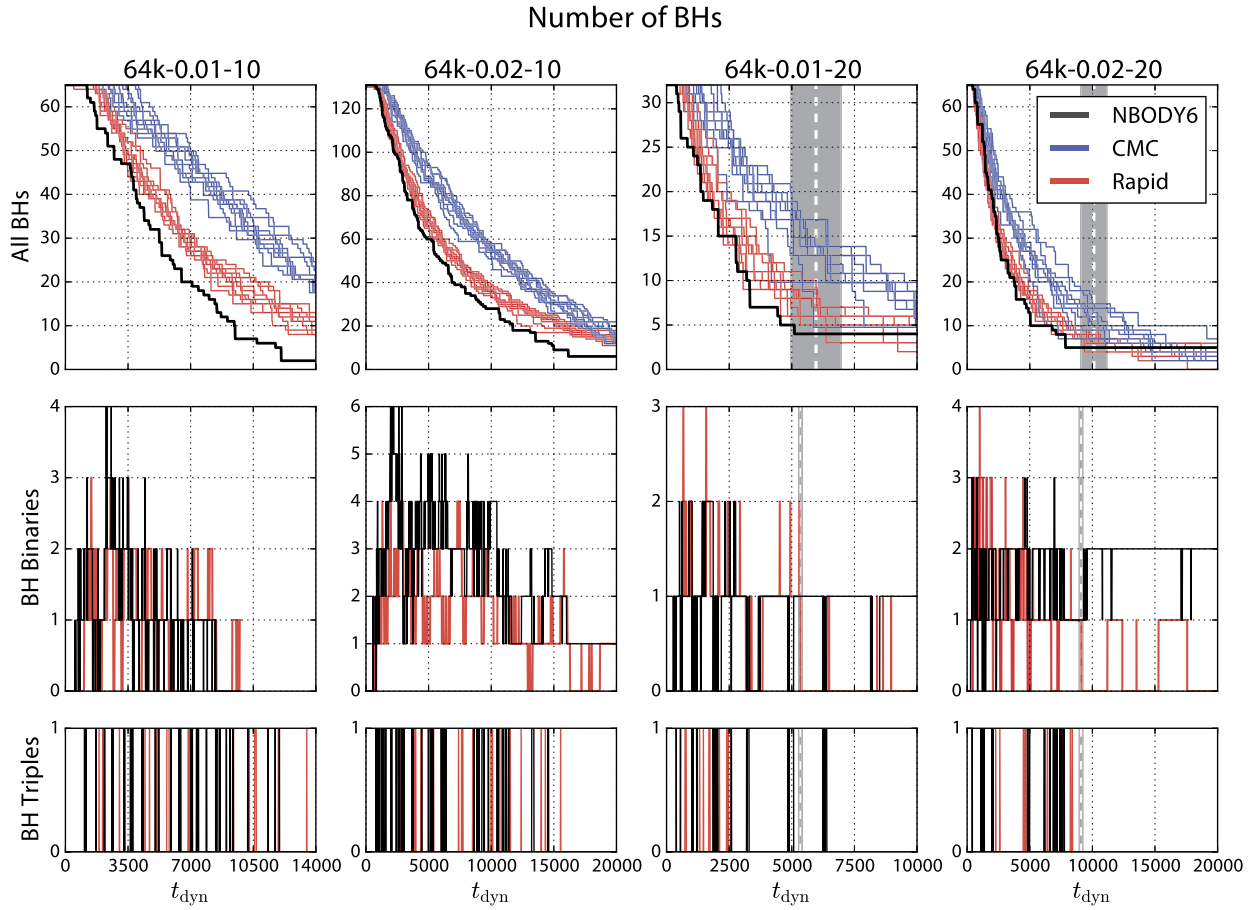
Figure 9 *Number of Single and Binary BHs*—The number of BHs retained in each cluster model over time. On the top, we show the total number of BHs present in each of the four clusters (including those in bound multiple systems) for the single NBODY6 model (in black), the 10 CMC models (in blue), and the 10 RAPID models (in red). In each case, the RAPID show better agreement with the NBODY6 model than the pure CMC models. In the second and third rows, we show the total number of BH binaries and BH triples (which cannot be produced in CMC) present in the NBODY6 model (in black) and a single RAPID model (in red) as a function of time. Although highly stochastic, the number of bound multiple systems shows good qualitative agreement between the two methods. As in Fig. 6, we show the mean and 1σ times when RAPID reverts to a pure MC approach with the white-dashed line and gray bands, respectively. For the middle and bottom rows, the white line indicated the reversion time for that run

In each of the four clusters, the RAPID models eject BHs at a much faster rate than the CMC models, in better agreement with the NBODY6 model. Although the BH ejection rate for NBODY6 is slightly faster than the other two methods (particularly for the $m_{BH}/m_{star} = 10$ cases), in each case RAPID performs far better than the pure MC approach. This is consistent with the difference in half-mass radii between NBODY6 and RAPID noted in the previous section, where systems with less-massive BHs expand faster in NBODY6 than in RAPID. Since the overall expansion of the cluster regulates the ejection rate of BH binaries, models that expand more rapidly will eject BHs more rapidly. For BHs in the cluster, NBODY6 and RAPID produce a similar number of BH binaries and BH triples over time. This is a noticeable improvement over CMC, especially when con-

sidering BH triples, which are explicitly removed from the MC integration.

6.4 Ejected BH systems
In addition to the retained BH binaries, it is important to compare the properties of the ejected binaries from each cluster model. Since the semi-major axis of an ejected binary is inversely proportional to its binding energy, the orbital properties at ejection are an excellent proxy for the hardening rate of the binaries within the core. If two cluster models eject a similar number of binaries with comparable binding energies, the energy production rate in the core must also be comparable between the two models. In Fig. 10 we show the semi-major axes and eccentricities for the four NBODY6 models and the 10 RAPID realizations of each cluster. In each case, the properties of the binaries
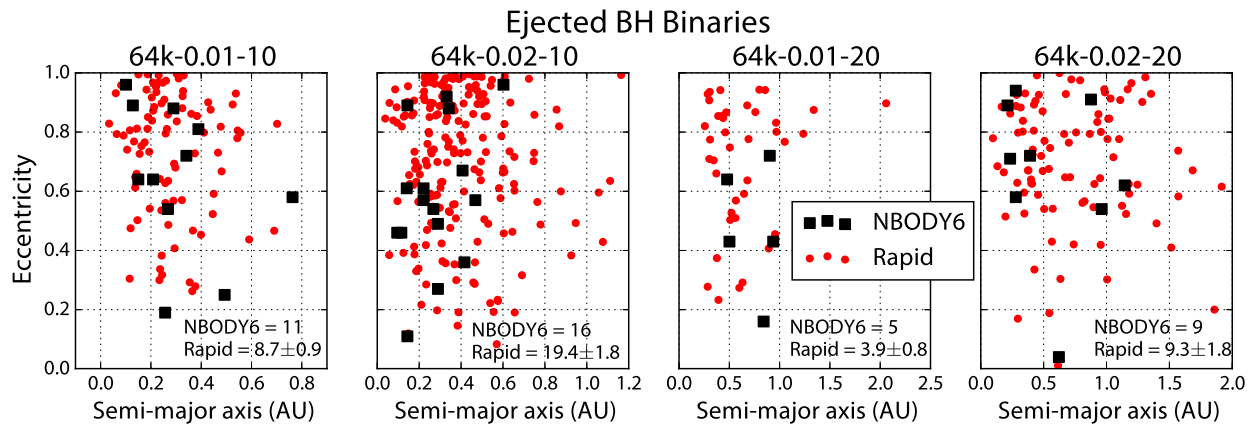
Figure 10 *Properties of Ejected Binary BH systems*—The properties of the ejected binary BH systems from each of the four `NBODY6` models (black diamonds) and from the 10 `RAPID` realizations (red circles). We show the eccentricity and semi-major axis (in AU, assuming each cluster to have an initial virial radius of 1 parsec) for each binary. In each case, the `NBODY6` results align with the larger sample size from the `RAPID` models. The text in each plot indicates the total number of ejected binaries from the `NBODY6` models, and the mean and standard deviation from the 10 `RAPID` models

ejected by the `NBODY6` models show good agreement with those formed by the 10 `RAPID` models.

Although we do not show it here, the distribution of orbital eccentricities of the ejected binaries strongly follows the theoretically-predicted distribution thermal distribution ($p(e) \propto e$, Jeans 1919) regardless of binary mass and cluster properties. Because the gravitational-wave merger time for binary BHs is determined by the semi-major axis and eccentricity at ejection (Peters 1964), `RAPID` will be able to model the merger rate of binary BHs from dense stellar clusters with similar accuracy to a direct *N*-body approach.

6.5 Energy conservation

To check the consistency of the `RAPID` approach, we examine the energy budget and conservation of each of the runs. We quantify the various forms of energy, including the kinetic and potential energy of all particles, the binding energy of multiples, and the total energy carried out of the cluster by ejected particles. The energies are plotted in Fig. 11. In addition, we also consider the virial ratio ($2K/W$) and the total energy over time for each system, as a diagnostic of the effectiveness of the method. The top two plots show the energy budget and virial ratio for a single representative model, while the bottom row shows the total energy for all 10 `RAPID` models.

The total energy conservation of `RAPID` can vary over a single run, and usually lies within 2–3% of the initial energy for the duration of the run. There are two main sources of error which contribute to this energy flux. The first is the difficulty of integrating higher-order multiples and very close encounters accurately, particularly those that occur in close triple systems. This issue is not limited to the Kira

integrator, and is one of the well-known issues common to all collisional dynamics simulations (the so called "terrible triples"). Although Kira does implement Keplerian regularization for isolated two-body systems and higher-order multiples, we still find that occasionally long-lived triples can induce substantial jumps in energy conservation. Furthermore, as the `CMC` potential is only applied to the center-of-mass of any multiple systems, any tidal effects upon the multiples from the background cluster potential are not incorporated correctly. These integration errors manifest as discontinuous jumps in the overall energy conservation, which can be seen in the bottom panels of Fig. 10.

The second source of error arises from the integration of orbits in a fixed external potential. This error takes two specific forms. First, the MC method, as described above, has an inconsistency in the computation of the potential. When a timestep is performed in `CMC`, the potential is computed first, before the dynamical encounters take place and the new orbit is calculated. However, the new orbit is calculated using the original potential, which does not take in to account the evolution of the cluster while the particles are dynamically interacting. While the work done by neighboring particles is correctly accounted for, the work done by the change in potential upon each particle is ignored. To compensate for this energy drift, `CMC` employs a technique developed by Stodoikiewicz (1982), in which the work done by the changing potential is explicitly added to the kinetic energy of the particle at the end of each timestep. This allows energy to be conserved in the MC to 1 part in 10^3 over a run (Fregeau and Rasio 2007). In `RAPID`, we self-consistently correct the velocities of stars controlled by the MC in a similar fashion, but
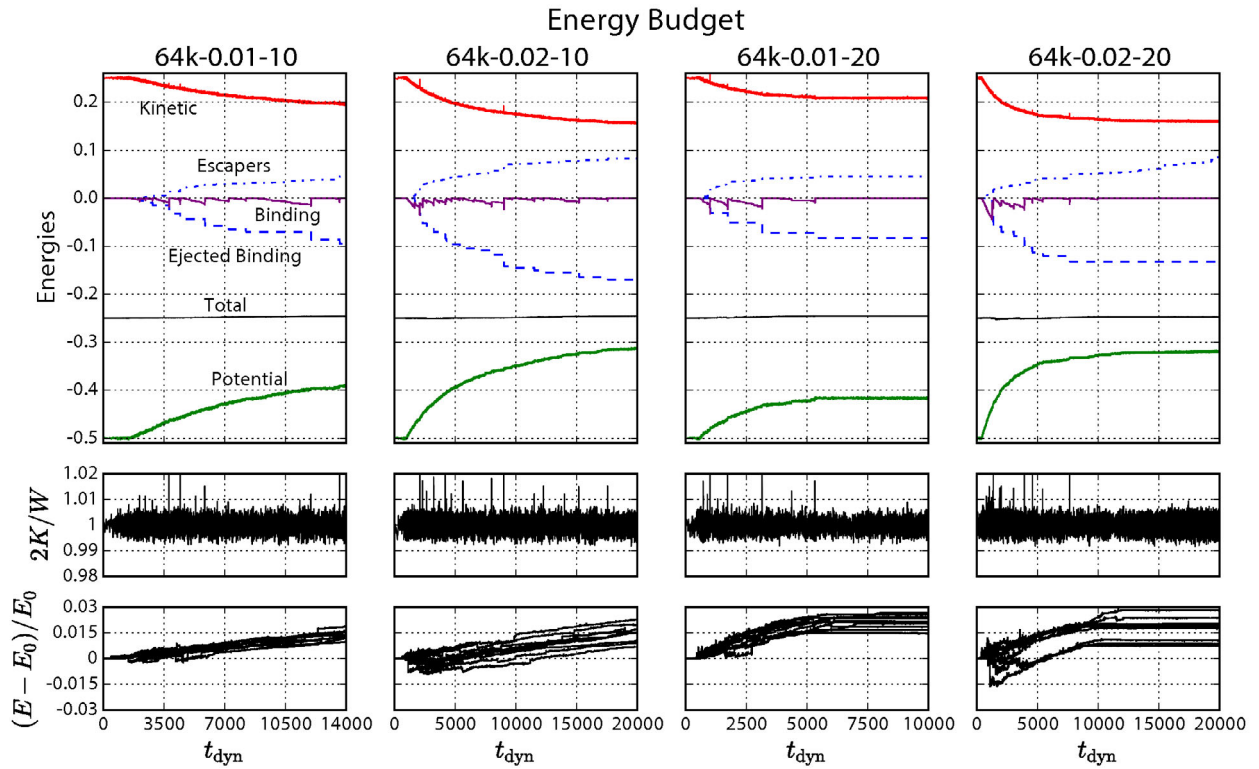
Figure 11 *Energy Conservation in the* `RAPID` *Code*—The energy budget and overall energy conservation for the four clusters modeled by `RAPID` shown in Fig. 6. The top plot shows the total sum of the kinetic energy (red), potential energy (green), and multiple binding energy (purple) for all particles over time for a single representative model of each cluster. The energy carried away by the kinetic energy (blue-dotted) or binding energy (blue-dashed) of escaping particles is also shown. In the middle plot, we show the virial ratio between the kinetic and potential energies. The sum of all energies is plotted in Hénon units (black) in the top plot, while the lower plot shows the fractional change in total energy for all 10 realizations of each model. The total energy conservation error is dominated by the evolution of the N-body particles in a strong external potential (the slow upward trend) and occasional strong-encounters and tightly-bound multiple systems (discontinuous jumps)

do not account for this energy drift in the BHs. We will explore modifications to the external potential (such as a time-dependent background potential for the Kira integrator) in a future work.

Additionally, the 4th-order Hermite integration scheme, while typical for collisional stellar dynamics, is known to produce systematic energy errors when integrating many orbits in a fixed potential (Binney and Tremaine 2008, cf. Fig. 3.21). This produces a small but systematically positive energy drift while integrating the BHs over many orbits. This is particularly problematic for the Kira integrator, which assumes that external potentials are weak perturbations to the internal dynamics of the cluster. While this assumption is valid for a cluster evolving in a galactic tidal field, in our current approach, the external potential from the MC particles is much stronger than the interparticle forces of the N-body integration. Methods to improve the long-term stability of the N-body integration, allowing for many integrations in a fixed potential while still treat-

ing close encounters accurately, are currently being investigated.

This issue may also be exacerbated by the particular combination of the dynamical timesteps between MC and N-body employed here. By integrating both systems for an identical length of time and combining the results, it is entirely possible that `RAPID` cannot adjust to rapid dynamical changes that occur *during* either the MC or N-body integrations. Investigations into using an adaptive timestep between the two computational domains, similar to the effective operator splitting developed by Fujii et al. (2007) and employed in Portegies Zwart et al. (2013), are currently underway.

Somewhat unexpectedly, Fig. 11 shows that this net energy drift does not depend on the number of BHs present in the cluster, but on the mass ratio between the individual stars and BHs. For clusters with a smaller m_{BH}/m_{star}, this net energy drift is slower. This indicates that for more realistic clusters containing many BHs and stars of different masses, this energy drift should improve. This is confirmed

by preliminary testing of RAPID on clusters with a realistic initial mass function.

7 Conclusion

In this paper, we described the motivation and development of a new hybrid technique for dynamically modeling dense star clusters. By combing our Cluster Monte Carlo (CMC) code with the Kira direct N-body integrator, we are able to combine the speed of the MC approach with the accuracy of a direct summation. This hybrid code, the *Rapid and Precisely Integrated Dynamics* (RAPID) Code is designed to accurately model the non-equilibrium BH dynamics that powers the overall evolution of GCs and GNs. Given recent observational detection of BH candidates in GCs, and the importance of theoretical modeling of GC BHs to X-ray binary astrophysics (Pooley et al. 2003) and gravitational-wave astrophysics (Rodriguez et al. 2015; Antonini et al. 2016), understanding the dynamics of BHs in clusters is crucial to understanding BH astrophysics.

We found that the hybrid approach is able to replicate both the half-mass radius and the core radius for several N-body models of idealized GCs with a much greater accuracy than a traditional MC integration. Unlike a purely MC approach, RAPID can model the highly non-spherical and rapidly changing dynamics of the few BHs in the center of the cluster. This suggests that the RAPID approach can follow the dynamics of BH systems with comparable accuracy to a direct N-body integration, but with roughly the same integration time (with in a factor of 2) of an orbit-sampling Monte Carlo approach.

With this technique, it will be possible to explore regions of the GC parameter space that have remained outside the computational feasibility of direct N-body computations. In particular, by treating the central BH subcluster correctly, RAPID can explore regions of the GC and GN parameter space, including clusters with massive central BHs, that have previously been unexplored by direct collisional methods.

Two issues remain to be addressed. First, the Kira N-body integrator does not completely conserve energy in the presence of a large external potential. This effect is a well-known drawback of 4th-order Hermite integrators, will need to be addressed. Efforts are currently underway to increase the computational order of the Hermite predictor-corrector, improving both the accuracy and speed of the integration (e.g., Nitadori and Makino 2008), and to incorporate a time-dependent potential in the N-body integrator, to account for the work done by the total cluster potential on the BHs..

Secondly, the regularization of binaries and higher-order multiples in Kira is based on Keplerian regularization for sufficiently unperturbed systems (see Portegies Zwart et al. 2001). However, this regularization does not include any tidal effects from the external Monte Carlo potential,

which will effect the long-term evolution of binaries retained by the cluster. Incorporation of these physical effects into the regularization scheme is currently underway.

Future work will also investigate hardware acceleration of the N-body integrator. The Kira integrator is designed to run on the specialized GRAPE series of hardware, which yields substantial improvements in computational speed. When combined with the Sapporo GPU/GRAPE library (Gaburov et al. 2009), Kira can be run on modern, distributed GPU systems with comparable performance to the NBODY series of codes (Anders et al. 2012). Hardware acceleration was not implemented the current RAPID version, since we have not considered systems with sufficiently large numbers of BHs for efficient GPU useage; however, the development of the Sapporo2 library (Bédorf et al. 2015) provides efficient GPU saturation for small-N systems. We will explore the advantages of a RAPID integration with Sapporo2 in a future paper.

RAPID is designed to be a single-purpose code incorporating all the necessary physics to model dense star clusters. However, these "kitchen-sink" codes, in which many numerical codes are integrated into a single parallel infrastructure, are often difficult to extend or modify for different purposes, particularly with regard to the shared timestep. The energy drift noted in Sect. 6.5 arises from a combination of the 4th-order Hermite integrator and the particular combination of the two dynamical timesteps. While the parallel design of RAPID makes it difficult to explore variations on this code structure, there do exist more modular approaches to computational stellar dynamics that may prove helpful. For example, the Astrophysical Multipurpose Software Environment (AMUSE, Portegies Zwart et al. 2013) can be used to easily swap different N-body integrators into a large-scale astrophysics code. Furthermore, there exist methods of combining different dynamical timesteps in a single code (e.g., the Bridge approach, Fujii et al. 2007), similar to the operator splitting approach developed by Wisdom and Holman (1991), that enable large multi-scale simulations to be performed with an adaptive, shared timestep. Because this leapfrog-esque approach is already implemented in AMUSE (as well as several different N-body integrators), we are exploring the possibility of integrating RAPID into AMUSE, allowing for greater precision and flexibility in the N-body timestep.

Abbreviations

GC, Globular Cluster; GN, Globular Cluster; MC, Monte Carlo; RAPID, Rapid and Precisely Integrated Dynamics; CMC, Cluster Monte Carlo; AMUSE, Astrophysical Multipurpose Software Environment; BH, Black Hole; SSE, Single Stellar Evolution; BSE, Binary Stellar Evolution.

Availability of data and materials

The authors have elected to not release the RAPID or CMC codes publicly at this time. Any of the data presented in this paper is available upon request.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

The RAPID code was developed in equal parts by CR and BP, with CR developing most of the astrophysical features and BP developing most of the parallel infrastructure; SC, MM, and FR assisted with the former, while AC and W-k L assisted with the later. CR prepared this manuscript and performed the tests detailed within. All authors read and approved the final manuscript.

Authors' information

Carl L. Rodriguez is a Pappalardo Fellow.

Author details

[1]MIT-Kavli Institute for Astrophysics and Space Research, Cambridge, USA. [2]Center for Interdisciplinary Exploration and Research in 37 Astrophysics (CIERA), Evanston, USA. [3]Dept. of Physics and Astronomy, Northwestern University, Evanston, USA. [4]Dept. of Electrical Engineering and Computer Science, Northwestern University, Evanston, USA.

Endnotes

[a] In NBODY6, the sum is performed out whichever is greater of the half-mass radius or three times the previous core radius.

[b] Of course, Hénon's principle ensures that the rate of binary formation and hardening will automatically adjust to satisfy the energy flux of the cluster across the half-mass radius (e.g., Breen and Heggie 2013). Because CMC can model the global properties of realistic clusters correctly, the rate of binary formation must be correct on a relaxation timescale, regardless of the specific implementation of three-body binary formation. What we are interested in here is the behavior of the inner parts of the core on a dynamical timescale, where the assumptions of the MC approach explicitly break down.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

Aarseth, S.J.: Publ. Astron. Soc. Pac. **111**, 1333 (1999)
Aarseth, S.J.: Gravitational N-Body Simulations. Cambridge Univ. Press, Cambridge (2003)
Aarseth, S.J.: Mon. Not. R. Astron. Soc. **422**, 841 (2012)
Abbott, B.P., et al.: Phys. Rev. Lett. **118**, 221101 (2017)
Anders, P., Baumgardt, H., Gaburov, E., Zwart, S.P., Portegies Zwart, S.: Mon. Not. R. Astron. Soc. **421**, 3557 (2012)
Antonini, F., Chatterjee, S., Rodriguez, C.L., Morscher, M., Pattabiraman, B., Kalogera, V., Rasio, F.A.: Astrophys. J. **816**, 65 (2016)
Arca Sedda, M., Askar, A., Giersz, M.: Mon. Not. R. Astron. Soc. **479**, 4652 (2018)
Askar, A., Arca Sedda, M., Giersz, M.: Mon. Not. R. Astron. Soc. **478**, 1844 (2018)
Bédorf, J., Gaburov, E., Portegies Zwart, S.: Comput. Astrophys. Cosmol. **2**, 8 (2015)
Belczynski, K., Sadowski, A., Rasio, F.A., Bulik, T.: Astrophys. J. **650**, 303 (2006)
Binney, J., Tremaine, S.: Galactic Dynamics, 2nd edn. Princeton Univ. Press, Princeton (2008)
Breen, P.G., Heggie, D.C.: Mon. Not. R. Astron. Soc. **432**, 2779 (2013)
Capuzzo-Dolcetta, R., Spera, M., Punzo, D.: J. Comput. Phys. **236**, 580 (2013)
Casertano, S., Hut, P.: Astrophys. J. **298**, 80 (1985)
Chandrasekhar, S.: Astrophys. J. **97**, 243 (1943)

Chomiuk, L., Strader, J., Maccarone, T.J., et al.: Astrophys. J. **777**, 69 (2013)
Clark, G.W.: Astrophys. J. **199**, L143 (1975)
Dehnen, W., Read, J.I.: Eur. Phys. J. Plus **126**, 28 (2011)
Downing, J.M.B.: Mon. Not. R. Astron. Soc. **425**, 2234 (2012)
Fregeau, J.M., Cheung, P., Portegies Zwart, S.F., Rasio, F.A.: Mon. Not. R. Astron. Soc. **352**, 1 (2004)
Fregeau, J.M., Joshi, K.J., Portegies Zwart, S.F., Rasio, F.A.: Astrophys. J. **570**, 171 (2002)
Fregeau, J.M., Rasio, F.A.: Astrophys. J. **658**, 1047 (2007)
Freitag, M.: In: Aarseth, S.J., Tout, C.A., Mardling, R.A. (eds.) The Cambridge N-Body Lectures. Lecture Notes in Physics, vol. 760, pp. 123–158. Springer, Dordrecht (2008). https://doi.org/10.1007/978-1-4020-8431-7
Freitag, M., Benz, W.: Astron. Astrophys. **375**, 711 (2001)
Fujii, M., Iwasawa, M., Funato, Y., Makino, J.: Publ. Astron. Soc. Jpn. **59**, 1095 (2007)
Gaburov, E., Harfst, S., Zwart, S.P.: New Astron. **14**, 630 (2009)
Giersz, M.: Mon. Not. R. Astron. Soc. **298**, 1239 (1998)
Giesers, B., et al.: Mon. Not. R. Astron. Soc. **475**, L15 (2018)
Harfst, S., Gualandris, A., Merritt, D., Mikkola, S.: Mon. Not. R. Astron. Soc. **389**, 2 (2008)
Heggie, D.C.: Mon. Not. R. Astron. Soc. **445**, 3435 (2014)
Heggie, D.C., Hut, P.: The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics. Cambridge Univ. Press, Cambridge (2003)
Hénon, M.: Astrophys. Space Sci. **14**, 151 (1971)
Hurley, J.R., Pols, O.R., Tout, C.A.: Mon. Not. R. Astron. Soc. **315**, 543 (2000)
Hurley, J.R., Tout, C.A., Aarseth, S.J., Pols, O.R.: Mon. Not. R. Astron. Soc. **323**, 630 (2001)
Hurley, J.R., Tout, C.A., Pols, O.R.: Mon. Not. R. Astron. Soc. **329**, 897 (2002)
Jeans, J.H.: Mon. Not. R. Astron. Soc. **79**, 408 (1919)
Joshi, K.J., Rasio, F.A., Zwart, S.P., Portegies Zwart, S.: Astrophys. J. **540**, 969 (2000)
Kremer, K., Chatterjee, S., Rodriguez, C.L., Rasio, F.A.: Astrophys. J. **852**, 29 (2018b)
Kremer, K., Chatterjee, S., Ye, C.S., Rodriguez, C.L., Rasio, F.A.: arXiv:1808.02204 (2018a)
Kremer, K., Ye, C.S., Chatterjee, S., et al.: Astrophys. J. **855**, L15 (2018)
Maccarone, T.J., Kundu, A., Zepf, S.E., Rhode, K.L.: Nature **445**, 183 (2007)
Mackey, A.D., Wilkinson, M.I., Davies, M.B., et al.: Mon. Not. R. Astron. Soc. **379** L40 (2007)
Mackey, A.D., Wilkinson, M.I., Davies, M.B., Gilmore, G.F.: Mon. Not. R. Astron. Soc. **386**, 65 (2008)
Makino, J., Hut, P.: Astrophys. J. Suppl. Ser. **68**, 833 (1988)
McMillan, S.L.W.: Astrophys. J. **307**, 126 (1986)
McMillan, S.L.W., Lightman, A.P.: Astrophys. J. **283**, 813 (1984a)
McMillan, S.L.W., Lightman, A.P.: Astrophys. J. **283**, 801 (1984b)
Merritt, D., Piatek, S., Portegies Zwart, S., Hemsendorf, M.: Astrophys. J. Lett. **608**, L25 (2004)
Miller-Jones, J.C.A., Strader, J., Heinke, C.O., et al.: Mon. Not. R. Astron. Soc. **453**, 3919 (2015)
Morscher, M., Pattabiraman, B., Rodriguez, C., Rasio, F.A., Umbreit, S.: Astrophys. J. **800**, 9 (2015)
Morscher, M., Umbreit, S., Farr, W.M., Rasio, F.A.: Astrophys. J. **763**, L15 (2013)
Nitadori, K., Aarseth, S.J.: Mon. Not. R. Astron. Soc. **424**, 545 (2012)
Nitadori, K., Makino, J.: New Astron. **13**, 498 (2008)
Pattabiraman, B., Umbreit, S., Liao, W.-k., et al.: Astrophys. J. Suppl. Ser. **204**, 15 (2013)
Peters, P.C.: Phys. Rev. **136**, 1224 (1964)
Pooley, D., Lewin, W.H.G., Anderson, S.F., et al.: Astrophys. J. **591**, L131 (2003)
Portegies Zwart, S., McMillan, S.L.W., van Elteren, E., Pelupessy, I., de Vries, N.: Comput. Phys. Commun. **184**, 456 (2013)
Portegies Zwart, S.F., McMillan, S.L.W., Hut, P., Makino, J.: Mon. Not. R. Astron. Soc. **321**, 199 (2001)
Rodriguez, C.L., Amaro-Seoane, P., Chatterjee, S., et al.: Phys. Rev. Lett. **120**, 151101 (2018)
Rodriguez, C.L., Haster, C.-J., Chatterjee, S., Kalogera, V., Rasio, F.A.: Astrophys. J. **824** L8 (2016b)
Rodriguez, C.L., Morscher, M., Pattabiraman, B., et al.: Phys. Rev. Lett. **115**, 051101 (2015)
Rodriguez, C.L., Morscher, M., Wang, L., et al.: Mon. Not. R. Astron. Soc. **463**, 2109 (2016a)
Sigurdsson, S., Hernquist, L.: Nature **364**, 423 (1993)
Sippel, A.C., Hurley, J.R.: Mon. Not. R. Astron. Soc. **430**, L30 (2013)

Spitzer, L.J.: Astrophys. J. **158**, L139 (1969)

Stodoikiewicz, J.S.: Acta Astron. **32**(1–2), 63–91 (1982)

Strader, J., Chomiuk, L., Maccarone, T.J., Miller-Jones, J.C.A., Seth, A.C.: Nature **490**, 71 (2012)

Strader, J., Seth, A.C., Forbes, D.A., et al.: Astrophys. J. **775**, L6 (2013)

Vasiliev, E.: Mon. Not. R. Astron. Soc. **446**, 3150 (2014)

Vasiliev, E., Antonini, F., Merritt, D.: Astrophys. J. **810**, 49 (2015)

Wang, L., Spurzem, R., Aarseth, S., et al.: Mon. Not. R. Astron. Soc. **450**, 4070 (2015)

Wang, L., Spurzem, R., Aarseth, S., et al.: Mon. Not. R. Astron. Soc. **458**, 1450 (2016)

Wisdom, J., Holman, M.: Astron. J. **102**, 1528 (1991)