The Astrophysical Journal Supplement Series, 258:22 (21pp), 2022 February https://doi.org/10.3847/1538-4365/ac2edf

Modeling Dense Star Clusters in the Milky Way and beyond with the Cluster Monte Carlo Code

Carl L. Rodriguez[1], Newlin C. Weatherford[2,3], Scott C. Coughlin[3], Pau Amaro-Seoane[4,5,6,7,8], Katelyn Breivik[9], Sourav Chatterjee[10], Giacomo Fragione[2,3], Fulya Kıroğlu[2,3], Kyle Kremer[11,12], Nicholas Z. Rui[11], Claire S. Ye[2,3], Michael Zevin[13,14], and Frederic A. Rasio[2,3]

[1] McWilliams Center for Cosmology and Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA; carlrodriguez@cmu.edu
[2] Department of Physics & Astronomy, Northwestern University, Evanston, IL 60208, USA
[3] Center for Interdisciplinary Exploration & Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, USA
[4] Universitat Politècnica de València, Camino de Vera, s/n E-46022 Valencia, Spain
[5] DESY Zeuthen, Platanenallee 6 D-15738 Zeuthen, Germany
[6] Institute of Applied Mathematics, Academy of Mathematics and Systems Science, CAS, Beijing 100190, People's Republic of China
[7] Kavli Institute for Astronomy and Astrophysics, Beijing 100871, People's Republic of China
[8] Zentrum für Astronomie und Astrophysik, TU Berlin, Hardenbergstraße 36, D-10623 Berlin, Germany
[9] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA
[10] Tata Institute of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India
[11] TAPIR, California Institute of Technology, Pasadena, CA 91125, USA
[12] The Observatories of the Carnegie Institution for Science, Pasadena, CA 91101, USA
[13] Kavli Institute for Cosmological Physics, The University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA
[14] Enrico Fermi Institute, The University of Chicago, 933 East 56th Street, Chicago, IL 60637, USA
Received 2021 June 4; revised 2021 October 1; accepted 2021 October 10; published 2022 January 18

Abstract

We describe the public release of the Cluster Monte Carlo (CMC) code, a parallel, star-by-star *N*-body code for modeling dense star clusters. CMC treats collisional stellar dynamics using Hénon's method, where the cumulative effect of many two-body encounters is statistically reproduced as a single effective encounter between nearest-neighbor particles on a relaxation timescale. The star-by-star approach allows for the inclusion of additional physics, including strong gravitational three- and four-body encounters, two-body tidal and gravitational-wave captures, mass loss in arbitrary galactic tidal fields, and stellar evolution for both single and binary stars. The public release of CMC is pinned directly to the COSMIC population synthesis code, allowing dynamical star cluster simulations and population synthesis studies to be performed using identical assumptions about the stellar physics and initial conditions. As a demonstration, we present two examples of star cluster modeling: first, we perform the largest ($N = 10^8$) star-by-star *N*-body simulation of a Plummer sphere evolving to core collapse, reproducing the expected self-similar density profile over more than 15 orders of magnitude; second, we generate realistic models for typical globular clusters, and we show that their dynamical evolution can produce significant numbers of black hole mergers with masses greater than those produced from isolated binary evolution (such as GW190521, a recently reported merger with component masses in the pulsational pair-instability mass gap).

Unified Astronomy Thesaurus concepts: Algorithms (1883); Distributed computing (1971); Gravitational waves (678); Globular star clusters (656); N-body simulations (1083); Stellar dynamics (1596)

1. Introduction

The modeling of dense star clusters (DSCs), such as globular clusters (GCs), super star clusters (SSCs), or the nuclear star clusters (NSCs) in the centers of many galaxies, remains one of the most challenging problems in computational astrophysics. At first glance, this fact is somewhat surprising: these clusters typically start with 10^5–10^8 individual particles, well within the range of many collisionless gravitational particle solvers. However, it is the *collisional* nature of these star clusters that makes their dynamics so difficult to resolve: unlike the dynamics of galaxies or particles in a cosmological volume, the long-term evolution of DSCs is driven by both the orbits of individual stars in the cluster potential and the diffusion of the stars through phase space via two-body encounters (e.g., Spitzer 1987; Heggie & Hut 2003; Binney & Tremaine 2008).

Both historically and currently, the first approach employed to study collisional stellar dynamics has often been a direct summation approach. In many ways the most straightforward conceptually (though exceedingly complex in implementation), a direct summation code calculates the gravitational force of every particle on every other particle before summing the individual forces to arrive at the instantaneous acceleration for every particle. The system is then advanced using standard numerical techniques for solving ordinary differential equations (typically a fourth-order Hermite integrator). These techniques have a long and storied history going back more than 60 yr to the first numerical work by von Hoerner (1960). The modern generation of codes, such as HiGPUs (Capuzzo-Dolcetta et al. 2013), PhiGRAPE (Harfst et al. 2008), ph4 (McMillan et al. 2012), frost (Rantala et al. 2021), and the well-known NBODY series (Aarseth 2003, 2012) and its derivatives (NBODY6++GPU, Wang et al. 2015), contains state-of-the-art algorithmic and hardware optimizations that enable the modeling of clusters with $N \sim 10^6$ particles. But even with dozens of parallel GPUs, these models require months or even years of wall-clock time to integrate a cluster for ∼10 Gyr (e.g.,

Heggie 2014; Wang et al. 2016; Rantala et al. 2021) and are limited to systems with low binary fractions, large initial radii, and relatively long dynamical times. These constraints preclude any reasonable exploration of the parameter space of massive star clusters, especially those with the compact initial radii needed to produce core-collapsed GCs (e.g., Kremer et al. 2019a).

Direct summation N-body methods represent the "gold standard" of collisional stellar dynamics because of their unparalleled accuracy. But just as the advent of fiat currency proved faster, cheaper, and more robust than the gold standard, faster approaches to collisional stellar dynamics offer many avenues of scientific study. The orbit-averaged Monte Carlo method, originally developed by Hénon (1971a, 1971b), is one such approach. Instead of integrating the orbits of each star directly, the Monte Carlo method leverages a statistical treatment of stellar dynamics, where the cumulative effect of many distant two-body encounters is modeled as a single effective scattering between neighboring particles. The deflection angle of this effective encounter is chosen to reproduce the mean of the averaged change in the velocity per unit time, $\langle(\Delta v)\rangle^2$, experienced by a particle traveling through a field of stars with known density (e.g., Spitzer 1987, Chapter 2). In doing so, Hénon's method naturally resolves the two-body relaxation that drives the long-term evolution of collisional star systems and has been shown to reproduce the pre- and post-core-collapse evolution of DSCs for nearly 50 yr (Aarseth et al. 1974; Joshi et al. 2000; Giersz et al. 2013; Rodriguez et al. 2016b; Kremer et al. 2020c). Furthermore, by modeling collisional dynamics as nearest-neighbor encounters between radially sorted particles, the Monte Carlo method can also be expanded to include close encounters between stars and binaries, a critical component to understanding the post-collapse evolution of clusters and their production of stellar exotica such as X-ray binaries (Clark 1975; Pooley et al. 2003; Kremer et al. 2018a), millisecond pulsars (Rappaport & Putney 1989; Ye et al. 2019), cataclysmic variables (Grindlay et al. 1995; Kremer et al. 2021), and binary black hole (BBH) mergers (Portegies Zwart & Mcmillan 2000; Rodriguez et al. 2015).

Following updates to the method by Stodółkiewicz (1982) to ensure the stability of long-term integrations, the majority of work on the Hénon N-body method has been led by two groups: the Cluster Monte Carlo (CMC) group (Joshi et al. 2000) and the Monte Carlo Cluster Simulator (MOCCA) group (Giersz 1998). A third code was also developed around the same time by Freitag & Benz (2001), which introduced new ways of treating stellar collisions and central-massive black holes (BHs) with Hénon's method. While that code is no longer actively maintained, many of its techniques have since been incorporated into both CMC and MOCCA. Both codes have seen significant enhancements and improvements over the past two decades, with the most recent iterations of the codes being described in Pattabiraman et al. (2013) and Giersz et al. (2013). Although there are key differences in the choices of the dynamical time step and other aspects of stellar evolution and collisions, CMC and MOCCA contain much of the same physics and are currently the only codes capable of modeling realistic populations of DSCs with more than 10^6 stars and binaries over many gigayears.

In this paper, we describe the first public release of CMC, an N-body approach to modeling collisional stellar dynamics with Hénon's method. Written in C and designed for distributed-memory architectures, CMC has been continuously developed since Joshi et al. (2000) and contains all the necessary physics for the evolution of massive, spherical DSCs. Key physical processes, such as two- and three-body binary formation, strong encounters between stars and binaries (performed with the small-N direct integrator fewbody; Fregeau et al. 2003; Fregeau & Rasio 2007), galactic tidal fields (Joshi et al. 2001; C. L. Rodriguez et al. 2021, in preparation), physical collisions (Fregeau & Rasio 2007), and post-Newtonian (pN) dynamics (Rodriguez et al. 2018b), have been incorporated and are available in the v1.0 release of CMC (Rodriguez et al. 2021a).

The public release of CMC also contains detailed treatments for single and binary stellar evolution using the COSMIC population synthesis code (Breivik et al. 2020a). Both CMC and COSMIC are based on the Binary Stellar Evolution (BSE) package of Hurley et al. (2000, 2002) and have been significantly enhanced with new prescriptions for the evolution of massive stars, binary mass transfer, common-envelope efficiency, compact object formation, and much more. The latest version of COSMIC, v3.4, is now pinned as a Git submodule to CMC, ensuring that any changes to the population synthesis code are also available in the dynamics code. This allows for direct, "apples-to-apples" comparisons between isolated binary evolution and dynamical formation for many astrophysical sources, such as stellar binaries and high-energy transients. In Section 2, we review the Hénon method, its implementation in CMC, and the additional dynamical and stellar physics included in the code, as well as the assumptions and limitations of the Monte Carlo method. In Section 3, we describe features specific to the public release of CMC and the integration of the COSMIC population synthesis package. Finally, in Section 4, we show two examples of cluster evolution: a Plummer sphere of point-mass particles evolved to core collapse, and several realistic cluster models evolved for 12 Gyr.

2. Dynamics in CMC

Here we review the theoretical basis of the original Hénon (1971a, 1971b) scheme as it is implemented in CMC. In a collisional stellar system, particles undergo repeated two-body gravitational encounters with other particles. These "collisions" allow energy and angular momentum to be exchanged between particles and subsequently diffused throughout the cluster, just as physical collisions diffuse energy throughout an ideal gas (though, unlike their microscopic counterparts, a star cluster can never reach thermodynamic equilibrium). In clusters with relatively few particles ($N \lesssim 10^3$), this process is dominated by a handful of close encounters with small impact parameters, producing large deflection angles with every encounter. With larger N, however, the cumulative effect of many distant encounters (producing many small-angle deflections in the particle's trajectory) begins to dominate the cluster evolution (e.g., Spitzer 1987). In the large-N regime, the dynamical behavior of the particles becomes largely predictable using the techniques of statistical mechanics, forming the basis for the Hénon method.

The dividing line between these regimes can be understood as a competition between two timescales: the *dynamical* time (the time a particle takes to orbit the cluster at the half-mass radius), and the *relaxation* timescale, or the time required for these distant encounters to change the velocity of the particle

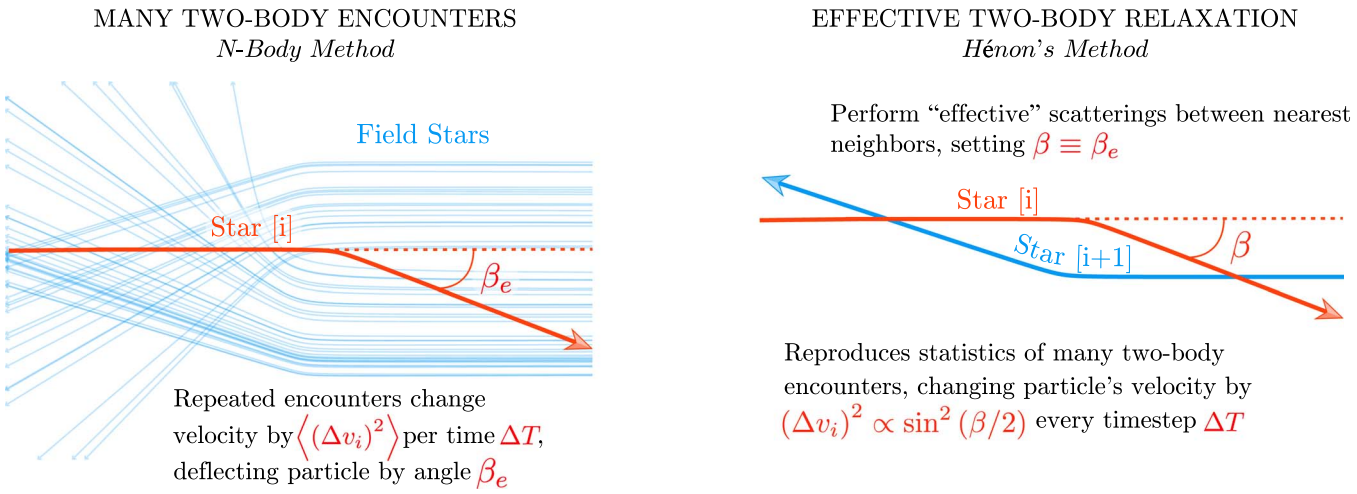

Figure 1. Schematic diagram of the Hénon method. The cumulative velocity change due to many two-body encounters in a realistic cluster can be calculated given the average local mass, velocity, and density of stars with Equation (9). In Hénon's method, this is reduced to a single encounter between nearest-neighbor particles, where the deflection angle is chosen to reproduce the statistical predictions using Equation (12).

by order of itself. The scaling between these two timescales depends on the size of the cluster (e.g., Binney & Tremaine 2008):

$$T_{\rm relax} \approx \frac{0.1N}{\log N} T_{\rm dyn}. \quad (1)$$

When $T_{\rm relax} \gg T_{\rm dyn}$, the cluster is in the weak-encounter regime and can be thought of as a collection of stars on semi-independent orbits that slowly diffuse through phase space on a relaxation timescale.

By itself, two-body relaxation would drive star clusters toward core collapse, where the inner regions of the cluster contract as energy is transported by two-body relaxation from the core toward the outer halo of the cluster. This, in turn, causes an expansion of the outer regions of the halo and creates a dynamical "temperature" gradient between the hotter core and the cooler outer regions. This gradient accelerates the transport of heat from core to halo, driving the central regions of the cluster inward and toward the well-known gravothermal catastrophe (e.g., Lynden-Bell & Wood 1968; Antonov 1985; Heggie & Hut 2003). But while these distant two-body encounters drive the cluster to collapse, it is the close encounters that largely reverse this process. As the core contracts, the increasing density of objects facilitates the production and dynamical hardening of binaries, creating an effective power source that halts the collapse and determines the long-term evolutionary fate of the cluster. These close encounters are also largely responsible for the large number of stellar exotica found in GCs and other DSCs.

In CMC, the dynamics of interparticle interactions are separated into these two natural categories: (a) the average two-body relaxation from many weak encounters using Hénon's method, and (b) strong encounters, such as binary encounters and physical collisions. We now describe both of these.

2.1. Two-body Relaxation

In the Hénon method, the cluster is assumed to be spherically symmetric and in dynamical (virial) equilibrium. The full 6D positions and velocities are reduced to 3D: the radial position, r, of each particle, and its radial and tangential velocities, v_r and v_t. Of course, these position and velocities change on a dynamical timescale as the particle orbits through the cluster potential. What the Monte Carlo method tracks over time is the energy (E) and angular momentum (J) of each particle's orbit, quantities that are conserved on dynamical timescales. Combined with the cluster potential, $\Phi(r)$, these quantities fully define an orbit for every particle in the cluster. New r, v_r, and v_t are randomly sampled from these orbits at every Monte Carlo time step.

While E and J do not change on dynamical timescales, the cumulative effect of many weak encounters drives their evolution on a relaxation timescale. During a Monte Carlo time step, we model these weak encounters as perturbations ΔE and ΔJ to the particle's energy and angular momentum. These perturbations are computed as a single *effective* encounter between each particle and its nearest neighbor in radius. In the center-of-mass frame of the two-body encounter, the *magnitude* of the particles' velocities is unchanged, but the velocity vectors are deflected by some angle β. In this frame, the energy of the two particles is conserved; however, in the *cluster* frame, such encounters exchange energy and angular momentum. To correctly capture two-body relaxation throughout the cluster, each of these effective encounters should give the correct value for the mean change in kinetic energy at each particle's position during the time step ΔT. We satisfy this constraint by computing an effective deflection angle β_e for each particle that results in the correct mean ΔE and ΔJ. See Figure 1.

2.1.1. Calculating the Effective Scattering Angle

We now proceed to derive the effective scattering angle, β_e, following the logic of Stodółkiewicz (1982) and Joshi et al. (2000). We begin with an array of particles sorted by increasing radius from the cluster center such that $r_i < r_{i+1}$. Fundamentally we are interested in the average change in energy and angular momentum, ΔE and ΔJ, particle i experiences during some time interval ΔT. We begin by considering the change in velocity squared, $(\Delta v_i)^2$, for a single two-body interaction between the star with index i and a field star with mass m_f. This expression can be written in terms of β, the angle of deflection of particle i, or in terms of b, the impact parameter of the

encounter at infinity, as (e.g., Binney & Tremaine 2008, Equation (3.54))

$$(\Delta v_i)^2 = \frac{4\,m_f^2}{(m_i + m_f)^2} w^2 \sin^2\left(\frac{\beta}{2}\right) \quad (2)$$

$$= \frac{4\,G^2\,m_f^2}{w^2\,b_{90}^2} \frac{1}{1 + (b/b_{90})^2}, \quad (3)$$

where $w = |\boldsymbol{v}_i - \boldsymbol{v}_f|$ is the relative speed of the particles at infinity and $b_{90} \equiv G(m_i + m_f)/w^2$ is the impact parameter that would deflect the particles by $\beta = 90°$.

It is straightforward to extend Equation (3) to describe the effect of many encounters. The change in velocity over time due to many weak encounters can be calculated by considering the field particles encountered by particle i as it travels through the cluster. The number of particles encountered at an impact parameter b during an interval ΔT is simply the product of the local number density of particles, n, and the volume swept out at that impact parameter. For an infinitesimal annulus with inner radius b and outer radius $b+db$, this product is $N_{\rm enc} = 2\pi b n w \Delta T db$, where the length of the cylinder is the relative velocity of the particles, w, times ΔT. To compute the contribution of this result over all impact parameters, we multiply Equation (3) by $N_{\rm enc}$ and integrate over b from zero to some upper limit $b_{\rm max}$:

$$\langle (\Delta v_i)^2 \rangle = \frac{8\pi G^2 n \Delta T\, m_f^2}{w b_{90}^2} \int_0^{b_{\rm max}} \frac{b \cdot db}{1 + (b/b_{90})^2} \quad (4)$$

$$= 4\pi G^2 n \Delta T\, m_f^2 w^{-1} \ln\left[1 + (b_{\rm max}/b_{90})^2\right] \quad (5)$$

$$\simeq 8\pi G^2 n \Delta T\, m_f^2 w^{-1} \ln\Lambda\ . \quad (6)$$

The value of $b_{\rm max}$ to use in Equation (4) is a well-studied complication to the theory: on the one hand, in a system with uniform density, equal octaves of b contribute equally to $\langle (\Delta v_i)^2 \rangle$, meaning that for sufficiently large systems $b_{\rm max} \gg b_{90}$; on the other hand, star clusters are finite systems whose densities eventually fall to zero at some finite radii. In going from Equation (5) to Equation (6), we have introduced the well-known Coulomb logarithm, defined as

$$\Lambda \equiv \frac{b_{\rm max}}{b_{90}} = \frac{w^2 b_{\rm max}}{G(m_i + m_f)}, \quad (7)$$

and used our assumption of the weak-encounter regime to set $\ln[1 + (b_{\rm max}/b_{90})^2] \simeq 2\ln[b_{\rm max}/b_{90}]$. Λ can be approximately computed at the half-mass radius of the cluster by equating $b_{\rm max} = r_h$ and replacing w^2 with the mean-squared velocity from the virial theorem, $\langle v^2 \rangle \simeq 0.45 G m_f / r_h$ (Binney & Tremaine 2008, p. 361). For a cluster with N stars of mass m_f, this yields

$$\ln\Lambda = \ln[\gamma N], \quad (8)$$

with $\gamma \simeq 0.2$. However, comparisons to direct N-body simulations have suggested a γ of 0.11 (for equal-mass clusters; e.g., Giersz & Heggie 1994), 0.02 (for clusters with realistic mass functions; Giersz & Heggie 1996), or even 0.01 (the current default in CMC; Freitag et al. 2006; Rodriguez et al. 2018c), to be more realistic.

Equations (4)–(6) were computed assuming that every field star has the same mass and relative velocity, m_f and w, with respect to our test star, m_i. In realistic clusters, however, nearby particles will have a range of masses and velocities, determined by the distribution function, $F(\boldsymbol{r}, \boldsymbol{v}, m)$. Therefore, to accurately reproduce the cumulative shift $\langle (\Delta v_i)^2 \rangle$ from Equation (6), we must replace the $m_f^2 w^{-1}$ factor with

$$\langle (\Delta v_i)^2 \rangle = 8\pi G^2 n \Delta T \ln\Lambda \langle m_f^2 w^{-1} \rangle_F, \quad (9)$$

where we use $\langle ... \rangle_F$ to indicate an average over the local phase-space distribution function

$$\langle m_f^2 w^{-1} \rangle_F \equiv \int F_i F_f m_f^2 w^{-1} d^3\boldsymbol{v}_i d^3\boldsymbol{v}_f dm_i dm_f, \quad (10)$$

with $F_i = F(\boldsymbol{r}_i, \boldsymbol{v}_i, m_i)$ and $F_f = F(\boldsymbol{r}_f, \boldsymbol{v}_f, m_f)$.

We now reach the crux of the Monte Carlo method: to get the angle of deflection β in terms of these quantities—and the correct mean value of $m_i(\Delta v_i)^2$—we equate the right-hand side of Equation (9) to the right-hand side of Equation (2):

$$\left\langle \frac{4\,m_f^2\,w^2}{(m_i + m_f)^2} \right\rangle_F \sin^2\left(\frac{\beta_e}{2}\right) = 8\pi G^2 n \Delta T \ln\Lambda \langle m_f^2 w^{-1} \rangle_F, \quad (11)$$

where we have replaced a single deflection angle, β, with an effective deflection angle, β_e, and replaced the mass and velocity terms with their averaged equivalents. We now solve for β_e, which involves averaging the mass and velocity quantities over the relevant distribution functions. However, as Hénon (1971a) pointed out, the most straightforward way to sample the distribution function near star m_i is to pick the mass and velocity of its nearest neighbor, m_{i+1}. While this is not the same as a true average, the nearest neighbor represents a fair draw from the mass and velocity distributions and would reproduce the true distribution function after a sufficient number of time steps. Dispensing with the distribution function averages and substituting $m_f \rightarrow m_{i+1}$, Equation (11) becomes

$$\sin^2\left(\frac{\beta_e}{2}\right) = \frac{2\pi G^2 (m_i + m_{i+1})^2}{w^3} n \ln\Lambda \Delta T. \quad (12)$$

The only quantity that must be explicitly averaged is the local volumetric density of particles, n. In CMC, this is done with a moving average over 41 stars, from $i - 20$ to $i + 20$ for each star i.

For computational convenience, Equation (12) is more easily expressed in terms of "relaxation timescale" between particles i and $i + 1$:

$$\sin\left(\frac{\beta_e}{2}\right) = \frac{\pi}{4}\sqrt{\frac{\Delta T}{T_{{\rm rel},i}}}, \quad (13)$$

where

$$T_{{\rm rel},i} \equiv \frac{\pi}{32} \frac{w^3}{G^2 n (m_i + m_{i+1})^2 \ln\Lambda} \quad (14)$$

is the typical time required for a particle to be deflected by 90° (Freitag & Benz 2001). Note that, in practice, we clamp the maximum β_e between any two particles at 90°.

2.1.2. Calculating New E and J

With β_e computed for each pair of particles, we now convert these pairwise deflection angles into ΔE and ΔJ for each orbit.

We follow Hénon's original approach, placing our two-body encounter into 3D space. We denote the phase-space coordinates of the two interacting particles by (r_i, $v_{r,i}$, $v_{t,i}$) and (r_{i+1}, $v_{r,i+1}$, $v_{t,i+1}$), with masses m_i and m_{i+1}, respectively. We choose our reference frame such that the z-axis is parallel to $\boldsymbol{r}_i$ and the (x, z)-plane contains $\boldsymbol{v}_i$. In Cartesian coordinates, the two particle velocities are then

$$\begin{aligned}\boldsymbol{v}_i &= (v_{t,i},\, 0,\, v_{r,i}),\\ \boldsymbol{v}_{i+1} &= (v_{t,i+1}\cos\phi,\, v_{t,i+1}\sin\phi,\, v_{r,i+1}),\end{aligned} \tag{15}$$

where ϕ is a uniform random variate in the range $[0, 2\pi]$, since spherical symmetry guarantees isotropic transverse velocities. The relative velocity is then

$$\begin{aligned}\boldsymbol{w} &\equiv \boldsymbol{v}_{i+1} - \boldsymbol{v}_i\\ &= (v_{t,i+1}\cos\phi - v_{i,t},\, v_{t,i+1}\sin\phi,\, v_{r,i+1} - v_{r,i}).\end{aligned} \tag{16}$$

We now define vectors $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ with magnitude equal to $|\boldsymbol{w}|$, such that $\boldsymbol{w}_1 \times \boldsymbol{w}_2 = \boldsymbol{w}$. In this right-handed coordinate system, we have

$$\begin{aligned}\boldsymbol{w}_1 &\equiv \left(\frac{-w_y w}{w_p},\, \frac{w_x w}{w_p},\, 0\right),\\ \boldsymbol{w}_2 &\equiv \left(\frac{-w_x w_z}{w_p},\, \frac{-w_y w_z}{w_p},\, w_p\right),\end{aligned} \tag{17}$$

where $w_p = \sqrt{w_x^2 + w_y^2}$. Since the distribution of field stars is assumed to be spherically symmetric, we randomly select the angle $\psi \in [0, 2\pi]$ between the plane of relative motion, defined by ($\boldsymbol{r}_{i+1} - \boldsymbol{r}_i$, $\boldsymbol{v}_{i+1} - \boldsymbol{v}_i$), and the plane containing $\boldsymbol{w}$ and $\boldsymbol{w}_1$. The relative velocity after the encounter, $\boldsymbol{w}^{\rm new}$, is then

$$\begin{aligned}\boldsymbol{w}^{\rm new} = \boldsymbol{w}\cos\beta_e &+ \boldsymbol{w}_1\sin\beta_e\cos\psi\\ &+ \boldsymbol{w}_2\sin\beta_e\sin\psi,\end{aligned} \tag{18}$$

where β_e is the effective scattering angle of Section 2.1. The post-interaction particle velocities in the cluster frame, $\boldsymbol{v}_i^{\rm new}$ and $\boldsymbol{v}_{i+1}^{\rm new}$, become

$$\begin{aligned}\boldsymbol{v}_i^{\rm new} &= \boldsymbol{v}_i - \left(\frac{m_{i+1}}{m_i + m_{i+1}}\right)(\boldsymbol{w}^{\rm new} - \boldsymbol{w}),\\ \boldsymbol{v}_{i+1}^{\rm new} &= \boldsymbol{v}_{i+1} + \left(\frac{m_i}{m_i + m_{i+1}}\right)(\boldsymbol{w}^{\rm new} - \boldsymbol{w}).\end{aligned} \tag{19}$$

Since we set the z-axis parallel to $\boldsymbol{r}^i$, the new radial and transverse velocities for the first particle are $v_{r,i}^{\rm new} = v_{z,i}^{\rm new}$ and $v_{t,i}^{\rm new} = \sqrt{(v_{x,i}^{\rm new})^2 + (v_{y,i}^{\rm new})^2}$. The new specific orbital energy and specific angular momentum are then $E_i = \Phi(r_i) + \frac{1}{2}((v_{r,i}^{\rm new})^2 + (v_{t,i}^{\rm new})^2)$ and $J_i = r_i v_{t,i}^{\rm new}$, respectively, and similarly for the $i+1$ particle.

2.2. Strong Interactions

Layered onto the basic two-body relaxation calculation is the possibility for strong interactions to occur between neighboring particles at each time step. Strong interactions can be separated into two groups: single–single interactions (which include physical collisions, tidal disruptions, tidal captures, and gravitational-wave captures of pairs of BHs) and binary interactions involving one or two binaries (integrated directly using `fewbody`). Strong interactions are sampled following the method of Freitag & Benz (2002; see also Fregeau & Rasio 2007), where we evaluate the quantity

$$P_{\rm strong} = nw\Sigma_{\rm strong}\Delta T \tag{20}$$

between each pair of neighboring stars and binaries. Here $P_{\rm strong}$ is the probability for a strong interaction with cross section $\Sigma_{\rm strong}$ to occur, n is the local number density of stars, w is the relative velocity of the pair of objects at infinity (computed in the same coordinate system as Section 2.1.2), and ΔT is the time step. Accounting for gravitational focusing, $\Sigma_{\rm strong}$ can be written as

$$\Sigma_{\rm strong} = \pi r_p^2\left(1 + \frac{2GM}{r_p w^2}\right), \tag{21}$$

where M is the total mass of the pair of interacting objects and r_p is the maximum distance of pericenter along a hyperbolic for each strong encounter to occur. At each time step, a random number X is drawn from a uniform distribution between 0 and 1. If $X < P_{\rm strong}$, a strong interaction occurs; otherwise, the particles undergo standard two-body relaxation.

The value of r_p is determined by the type of interaction, as described in the following subsections. We note that both the strong encounters described here and the binary formation considered in Section 2.3 are actually performed *before* two-body relaxation is considered in `CMC`. Any particles that participate in strong encounters do not participate in two-body relaxation during the same time step. This is because, on a relaxation timescale, the change in a particle's orbit due to a strong encounter occurs instantaneously and is often significantly greater than the gradual changes due to two-body encounters. These large changes in velocity render the calculation in Section 2.1.1 invalid, since the relative velocity between field and test stars is no longer constant over a given time interval. However, because our time step is chosen as the minimum of the relaxation and strong encounter timescales (the latter of which is typically much smaller than the former), the deflection by two-body relaxation during a time step where a strong encounter occurs is typically minimal. See Section 2.8.

2.2.1. Single–Single Collisions

For collisions between single stars, `CMC` operates under the sticky-sphere approximation, where any two particles that touch radii are assumed to have collided. The cross section for a pair of single stars with masses m_i and m_{i+1} and radii R_i and R_{i+1} can be expressed as

$$\Sigma_{\rm coll} = \pi(R_i + R_{i+1})^2\left(1 + \frac{2GM}{(R_i + R_{i+1})w^2}\right), \tag{22}$$

where $M = m_i + m_{i+1}$.

For main-sequence and/or giant stars, we adopt the classic sticky-sphere approximation, where zero mass is lost during the collision. The mass of the collision product, $m_{\rm new}$, is simply computed as $m_i + m_{i+1}$, with the new evolutionary phase of the star determined self-consistently using the stellar evolution prescriptions in `COSMIC`. In GCs, where the relative velocity of a pair of stars is typically much less than the escape velocity at the surface of the star, this approximation is reasonable (e.g., Lombardi et al. 2002). For compact objects, however, we employ more conservative assumptions for the mass retained by merger products. For collisions of white dwarfs (WDs) with

other WDs, main-sequence stars, or giants, we follow the prescriptions described in Hurley et al. (2002, Section 2.7). As described in Kremer et al. (2019b), when a main-sequence star or giant collides with a BH/neutron star (NS), we assume that the accretion feedback supplies sufficient energy to unbind the stellar material from the system completely. In this case, negligible mass is accreted and the mass of the BH/NS is unchanged after the collision. We note that this is an extremely conservative assumption, particularly if the star is significantly more massive than the BH or NS. However, in practice such collisions are rare in most typical GC models; see Section 4.3.

2.2.2. Binary Interactions

The maximum pericenter distance at which a dynamical binary interaction may occur is given by $r_p = \mathcal{X}_{\rm bin} a$ (for three-body encounters between a star and a binary with semimajor axis a) or $r_p = \mathcal{X}_{\rm bin}(a_1 + a_2)$ (for four-body encounters between two binaries with semimajor axes a_1 and a_2). Here $\mathcal{X}_{\rm bin}$ is a free parameter chosen so that all binary encounters of interest (i.e., those encounters that are energy generating) are captured. An arbitrarily large value of $\mathcal{X}_{\rm bin}$ would capture all binary encounters (the majority of which will be weak flybys); however, this would come at increased computational cost, as the global time step would shrink significantly to resolve these encounters (see Section 2.8). As described in Fregeau & Rasio (2007), we adopt a fiducial value of $\mathcal{X}_{\rm bin} = 2$, a compromise between capturing all possible encounters and computational cost. We note that this choice limits the number of weak flyby encounters that can occur in the cluster. While these encounters do not typically change the energy of the binaries (and are therefore irrelevant for the long-term evolution of the cluster), they can slowly alter the binary angular momentum, driving the system to higher eccentricities that will eventually result in mass transfer and even mergers (e.g., Rasio & Heggie 1995; Hamers & Samsing 2019). Efforts to model this in CMC are currently underway.

For binary–single encounters, we assume that the test star is a binary and calculate the probability for it to encounter a single object following Equation (20). For binary–binary encounters, we should formally replace the n in Equation (20) with $n_{\rm binary}$. However, the probability for binary–binary encounters to occur is only considered whenever the field star is also a binary, and the probability of *that* being true for any given particle is proportional to the density of binaries over the density of all objects, or $n_{\rm binary}/n$. By allowing binary–binary encounters to occur only for neighboring binaries, the probability from Equation (20) is implicitly multiplied by $n_{\rm binary}/n$ in true Monte Carlo fashion.

Binary encounters are integrated directly using the small-N scattering code, fewbody. These encounters depend on the relative velocity, w, of the pair of interacting objects (singles or binaries) and the impact parameter at infinity, b, which is sampled uniformly in annulus from a disk with maximum radius set by Equation (21). All interactions are integrated until they reach either an unambiguous end state or a maximum integration time (either 10^6 dynamical times of the encounter or 1 minute of CPU wall-clock time). Encounters that do not resolve into steady systems within these limits are discarded and treated as a standard two-body relaxation. See Fregeau & Rasio (2007) for details.

The version of fewbody in the public release of CMC also contains pN corrections to the equations of motion, including the 2.5pN term responsible for GW emission during close encounters (Antognini et al. 2014; Amaro-Seoane & Chen 2016; Rodriguez et al. 2018a, 2018b). By default, the 2.5pN term is included in the equations of motion when at least two BHs are present in an encounter. Note that, to accommodate the additional integration time required to resolve the inspiral and merger of BHs, the wall-clock time limit for these encounters is increased to 10 minutes when pN terms are included. We do not include the 1pN and 2pN terms in the integration by default, since they make energy conservation within the encounters impossible to track with the standard implementation of fewbody (see Rodriguez et al. 2018a, Figure 2, top panel) and have been shown to have no impact on the distribution of merging BBH eccentricities (Zevin et al. 2019).

One possible outcome of a binary–binary encounter is a stable hierarchical triple (e.g., Mikkola 1983; Rasio et al. 1995; Fragione et al. 2020b; Martinez et al. 2020). As described in Fregeau & Rasio (2007), if such a triple forms within a CMC simulation, we break the triple by unbinding the single star and inner binary, distributing its energy to nearby particles in the cluster. Additionally, because these fewbody encounters are performed in isolation (that is, without tidal forces from nearby particles), they can form pathologically wide binaries with separations larger than the typical interparticle distance. We destroy any such binaries (defined as having semimajor axes greater than 10% of the average local interparticle separation) and distribute their binding energy to nearby particles.

2.2.3. Collisions during Binary Interactions

Scattering encounters between stars and binaries can enhance the rate of stellar collisions by orders of magnitude (e.g., Verbunt & Hut 1987), largely due to close passages that can occur during chaotic resonant encounters (e.g., Hut & Inagaki 1985; McMillan & Hut 1996). Furthermore, hydrodynamical simulations have suggested that when a collision does occur, the merger product can have a radius anywhere from 2 to 30 times larger than the sum of the two parent stars' radii (Lombardi et al. 2003). This increase in the collisional cross section also increases the probability that the merger product will merge with another star *during the same scattering encounter* (e.g., Fregeau et al. 2003). To account for this in fewbody, we modify the standard sticky-sphere approximation as follows: whenever two stars merge, their merger product is assigned a mass equal to the sum of the two component masses, while the new radius is set to $f_{\rm exp}(R_1 + R_2)$, where $f_{\rm exp}$ is a constant factor (3, by default). The encounter is then allowed to continue to determine whether the objects will merge again. Whenever multiple collisions occur in a single encounter, we record the order of the collisions, and once the encounter is complete, the stars are then merged one at a time (in that same order) using the prescriptions in Section 2.2.1.

For BH–BH collisions, we set the collision radius to be 5 times the sum of the Schwarzschild radii of the two BHs ($10M$ in $G = c = 1$ units, where M is the sum of the two BH masses); this ensures that fewbody is operating in a regime where the pN expansion is valid. For BBH mergers it is useful to know what the properties of the merging BHs were just prior to merger (e.g., eccentricity). Because the orbit of inspiraling BBHs is highly non-Newtonian at $10M$, we record the (Newtonian) semimajor axes and eccentricities of merging

binaries at separations of $10M$, $50M$, $100M$, and $500M$. This allows the orbit to be integrated forward in post-processing to determine the eccentricity of the binary at a given GW frequency. Typically, the user should select the largest separation where the two BHs are bound, in order to minimize the error introduced by fitting to a Newtonian orbit (see Rodriguez et al. 2018a, Figure 2, bottom panel).

Unlike the traditional sticky-sphere approximation, the behavior of BBH merger products is modeled using detailed fitting formulae from numerical relativity simulations, with the final mass, final spin, and recoil kick being calculated from the set of formulae collected in Gerosa & Kesden (2016). When a merger occurs during a `fewbody` integration, the new BH mass and spin are set using these formulae, and the recoil kick is applied to the new merger product using the binary's orbital angular momentum vector to determine the coordinate system for applying the kick. This is useful for modeling the possibility of multiple BH mergers during a single scattering encounter, as well as the post-encounter orbital motion of any binaries in the cluster. See Rodriguez et al. (2018a, 2018b) for details. Note that these prescriptions for BH mass, spin, and recoil kick are also applied to any BBHs that merge inside the cluster.

2.3. Binary Formation

While star clusters are likely born with a nonnegligible fraction of their stars in binaries (e.g., Ivanova et al. 2005; Milone et al. 2012), they can also dynamically produce binaries many Myr after their initial phase of star formation is complete. Modeling this physics correctly is critical to understanding the long-term evolution of DSCs.

2.3.1. Three-body Binary Formation

Two-body encounters cannot create bound binaries from single stars (at least for Newtonian point masses) without some way to dissipate energy from the encounter. But if a two-body encounter occurs in the presence of a third body, excess kinetic energy can be carried away by the interloping particle, leaving behind a gravitationally bound system. While this process can involve more than three particles (e.g., Tanikawa et al. 2013), `CMC` uses a probabilistic treatment of three-body binary formation (Morscher et al. 2013) based on the formalism from Ivanova et al. (2005, 2010) and O'Leary et al. (2006).

Before either two-body relaxation or strong encounters are considered, the sorted list of stars is parsed three at a time, and the rate of binary formation for two objects of masses m_1 and m_2 from interactions with a third star (m_3) is calculated as

$$\begin{aligned}\Gamma_{3\mathrm{bb}}(\eta \geqslant \eta_{\min}) = {} & \sqrt{2}\,\pi^2 G^5 n^2 \langle w\rangle^{-9} \\ & \times (m_i + m_{i+1})^5 \eta_{\min}^{-5.5}(1 + 2\eta_{\min}) \\ & \times \left[1 + 2\eta_{\min}\left(\frac{m_i + m_{i+1} + m_{i+2}}{m_i + m_{i+1}}\right)\right],\end{aligned} \tag{23}$$

where n is the local stellar density of the cluster at that point, $\langle w\rangle = 4\sigma/\sqrt{3\pi}$ is the average relative velocity at infinity for two particles in a Maxwellian distribution, and σ is the local 3D velocity dispersion. The rate is computed as a function of the binary hardness ratio, defined as

$$\eta \equiv \frac{G m_i m_{i+1}}{a\,\langle m\rangle\,\sigma^2}, \tag{24}$$

where a is the semimajor axis of the hypothetical binary, $\langle m\rangle$ is the average local mass, and σ is the average local velocity dispersion, both computed over the closest 20 particles to the encounter. Note that this is half the number of particles used for computing the averages in Section 2.1.1, since three-body binary interactions are a fundamentally local process. Since only hard binaries are expected to survive long enough to influence the long-term evolution of the cluster, we only form binaries with a minimum hardness of $\eta_{\min} = 5$. By default, `CMC` only considers three-body binary formation for neighboring triplets of BHs in the cluster. Both of these choices are available as user parameters in `CMC`.

We note that Equation (23) differs from the original presentation in Ivanova et al. (2005, 2010) in two key ways. In Ivanova et al. (2005) the rate of three-body binary formation depends on the distinct number densities of the three masses. We simplify the expression here by replacing n_2, n_3, and n_c—the number densities of m_2, m_3, and the overall cluster core —with n, the local number density of the cluster. But for a given BH, we only consider $\Gamma_{3\mathrm{bb}}$ whenever *both* adjacent field stars are also BHs, meaning that in Monte Carlo fashion the probability calculated from Equation (23) is implicitly multiplied by $(n_{\mathrm{BH}}/n)^2$ (or $n_2 n_3/n^2$ for particles with different masses); see also the discussion regarding binary–binary encounters in Section 2.2.2. In practice, the majority of three-body binary formation occurs in the BH-dominated region of the core (even when three-body binaries are allowed to form from any stars), meaning that $n_c \approx n_2 \approx n_3 \approx n_{\mathrm{BH}}$ for most triple encounters.

The second key difference concerns the relative velocities: when computing Equation (23), we use $\langle w\rangle$, the averaged relative velocity between two particles. But the original expression from Ivanova et al. (2005) also contains the terms $v_{12}/\langle w\rangle$, where v_{12} is the relative velocity between m_1 and m_2, and $v_3/\langle w\rangle$, where v_3 is the relative velocity between m_3 and the m_1/m_2 center of mass. For computing the probability of binary formation, we assume $\langle w\rangle = v_{12} = v_3$ for computational convenience. However, when a binary is formed, we explicitly calculate v_{12} and v_3 using the same coordinate system presented in Section 2.1.2. This is necessary to determine the relevant energies of the encounter and to sample the correct distribution of η for that binary. We draw a random η between $\eta_{\min}$ and 50 using the normalized distribution given by $\frac{d\Gamma_{3\mathrm{bb}}}{d\eta}$, and we use this to determine the semimajor axis of the new binary (the eccentricities are assumed to be thermal). The relative velocities of the binary's center of mass and the third particle are then adjusted to conserve energy.

2.3.2. Two-body Binary Formation

While two-body binary formation is impossible when the two masses are Newtonian point particles, relaxing either of those assumptions can allow for binaries to form during close two-body encounters. If at least one of the masses has a nonnegligible radius (i.e., is not a WD, NS, or BH), close encounters can raise tides on one (or both) of the objects, allowing kinetic energy to be deposited into the structure of the

stars themselves (e.g., Fabian et al. 1975). If this energy exceeds the overall kinetic energy of the encounter, the two stars can form a binary through tidal capture. We assume polytropic stellar models for the main-sequence stars undergoing tidal capture processes and use polytropic index $n = 1.5$ for low-mass main-sequence stars ($<0.7\,M_{\odot}$) and $n = 3$ for higher-mass main-sequence stars. The calculations of the induced oscillation energy of the main-sequence stars and the cross sections during tidal capture follow prescriptions in Portegies Zwart & Meinen (1993) and Kim & Lee (1999), respectively. We track the first passage of a single–single close encounter. If a star is tidally captured after this passage, we set the binary semimajor axis to be twice the pericenter distance immediately after tidal capture, assuming angular momentum conservation. Note that giant stars are not included for tidal capture because of the highly uncertain reaction of their envelopes to the tidal force. In addition to tidal capture, if a compact object or a main-sequence star collides with a giant star during single–single interactions, the compact object/main-sequence star can form a binary with the giant core by transferring the orbital energy to the giant envelope and ejecting it. We treat main-sequence stars as point mass during this process and calculate the cross sections and final binary properties from the giant collisions using Equations (4)–(6) in Ivanova et al. (2006). See C. S. Ye et al. (2021, in preparation) for a complete description of these processes.

If instead of extended point particles we relax the Newtonian assumption, it becomes possible to form binaries through GW emission during close encounters (as occurs during BH mergers in `fewbody`). In `CMC`, the maximum pericenter distance where the GW emission of the two BHs carries away sufficient energy to take the two particles from unbound (positive total energy) to bound (negative total energy) is calculated as (Quinlan & Shapiro 1987)

$$r_p = \left(\frac{85\pi}{6}\right)^{2/7} \frac{G(m_i + m_{i+1})}{c^{10/7} w^{4/7}}. \tag{25}$$

BH binary formation via GW emission is decided probabilistically following a similar procedure to the collisions described in Section 2.3.1, using the above value of r_p in Equation (21). When a binary is formed, its impact parameter at infinity is sampled uniformly within the area of its capture cross section. The semimajor axis and eccentricity are calculated from conservation of energy and angular momentum, accounting for the loss due to GW emission. Many of these binaries are extremely wide (with semimajor axes of hundreds of AU or more), but their similarly extreme eccentricities ($e \gtrsim 0.99$) mean that they may still merge before they interact with neighboring particles. For these BBHs, we compare the timescale for GW emission to drive the binary to merge. We then compute the rate for this binary to encounter other stars using Equations (36) and (38) (albeit with the actual mass and semimajor axis of this binary, instead of the averaged quantities), and in true Monte Carlo fashion, we draw a random encounter time from an exponential distribution with that rate. If the encounter occurs before the BBHs would merge and the binary is pathologically wide (as described in Section 2.2.2), then the binary is disrupted. Otherwise, its semimajor axis and eccentricity are evolved by integrating the orbit-averaged equations for $\frac{da}{dt}$ and $\frac{de}{dt}$ from Peters (1964).

2.4. Stellar Evolution

`CMC` evolves every star and binary in the cluster forward by time step ΔT using detailed prescriptions for stellar evolution. These prescriptions, originally from Hurley et al. (2000, 2002), have been significantly updated over many years and form the basis for the `COSMIC` population synthesis package, which is now directly integrated into `CMC`. See Section 3.1, Chatterjee et al. (2010), and Breivik et al. (2020a) for more details.

We perform stellar evolution in `CMC` after the dynamical encounters have taken place, but before the new cluster potential and orbits are computed. This is to ensure that any dynamical changes to individual stars and binaries (particularly collisions and mass transfer in post-encounter binaries) can be accounted for using the same stellar physics, but before the new cluster orbits are calculated. Collisions in `fewbody`, where multiple stars can potentially collide during a single encounter, are passed to the stellar evolution module in the same order that they occur after the encounter is complete; see also Section 2.2.3.

2.5. New Positions and Velocities

At the end of every Monte Carlo time step, a new position and corresponding velocity are chosen for each particle by randomly sampling a point on its orbit, weighted by the amount of time the particle spends at that point along its orbit. We do so by first computing the orbital pericenter $r_{\rm min}$ and apocenter $r_{\rm max}$ as the two roots of the energy equation:

$$Q(r) = 2E - 2\Phi(r) - J^2/r^2 = 0. \tag{26}$$

The probability $P(r)$ of finding the particle at a position r within interval dr is proportional to the time the particle spends traversing dr. Therefore,

$$P(r)dr = \frac{dt}{T} = \frac{dr/|v_r|}{\int_{r_{\rm min}}^{r_{\rm max}} dr/|v_r|}, \tag{27}$$

where T is half the orbital period and $|v_r| = \sqrt{Q(r)}$. We draw a random sample from $P(r)$ and use the resultant $r_{\rm new}$ as the particle's position for the next time step. Note that $P(r) \propto 1/|v_r|$ becomes infinite at the turning points of the orbit, where $|v_r| = 0$. Changing coordinates by defining a suitable function $r = r(s)$ solves this problem (see Section 2.6 of Joshi et al. 2000). We must also set a new velocity consistent with the new orbital position: the new radial velocity is given magnitude $|v_r^{\rm new}| = \sqrt{Q(r_{\rm new})}$ with a randomly selected sign, while the new transverse velocity is given by $V_t^{\rm new} = J/r_{\rm new}$. Once the new positions and velocities have been calculated, the particles are then resorted in order of radial distance.

2.6. Escaping Particles

After new positions and velocities have been sampled, we remove any particles that are no longer bound to the cluster. For clusters in isolation, this is accomplished by removing any particles with positive energy after the dynamics and sorting steps are complete. Such objects are routinely formed during strong encounters (which can eject particles from the central regions of the cluster) or through standard two-body relaxation, as many particles will slowly diffuse to positive energies.

`CMC` can also model clusters that are tidally limited, assuming the clusters to be on circular orbits within their

galaxies. For a cluster of mass M_C on a circular orbit about the center of a point-mass galaxy, the tidal boundary of a cluster, r_t, can be calculated as

$$r_t^3 = \left(\frac{M_C}{3M_G}\right)^{1/3} R_G, \tag{28}$$

where M_G is the mass of the galaxy and R_G is the distance from the cluster to that galaxy (e.g., Spitzer 1987). Although sometimes called the Jacobi radius in stellar dynamics, Equation (28) is identical to the calculation of the Roche surface in binary stellar physics. As the cluster loses mass, the initial tidal boundary (specified by the user) decreases at a rate $\propto M_C^{1/3}$, following Equation (28).[15]

Of course, the Jacobi surface is not spherical; r_t actually represents the location of the Lagrange points along the line from the cluster to the galactic center. Since the Hénon method assumes the cluster to be spherically symmetric, approximate methods must be used to determine whether a star is stripped from the cluster by the galactic potential. CMC uses an energy-based criterion (Giersz et al. 2008) that removes any star with an energy above the potential at the tidal radius

$$E > \alpha\Phi(r_t), \tag{29}$$

where the parameter α is given by

$$\alpha = 1.5 - 3\left(\frac{\ln\Lambda}{N}\right)^{1/4}. \tag{30}$$

As previously, N is the initial number of stars in the cluster and $\ln\Lambda$ is the Coulomb logarithm from Equation (8). Numerical testing (Giersz et al. 2008; Chatterjee et al. 2010) has shown that Equation (29) produces significantly better agreement with direct N-body simulations than simply stripping stars with apocenters beyond the tidal boundary.

2.7. Potential Calculation

Because we assume the cluster potential to be spherically symmetric, computing it is straightforward. The potential at point r, located between two particles at r_i and r_{i+1}, is given by

$$\Phi(r) = -G\left(\frac{M_i}{r} + \sum_{i=j+1}^{N}\frac{m_i}{r_i}\right), \tag{31}$$

where $M_i \equiv \sum_{j=1}^{i} m_j$ is the total mass interior to star i.

While Equation (31) is formally correct, it is computationally expensive to evaluate whenever the potential at an arbitrary radius is required (such as when computing new orbital positions). However, we can accelerate this process by evaluating the potential at every particle, Φ_i, and interpolating the potential between that particle and its nearest neighbor, Φ_{i+1}. Because the stars are already sorted by radius, the potential at each particle can be computed recursively from the outermost star inward with

$$\Phi_{N+1} = 0,$$
$$\Phi_i = \Phi_{i+1} - GM_i\left(\frac{1}{r_i} - \frac{1}{r_{i+1}}\right), \tag{32}$$

where $M_N = \sum_{i=1}^{N} m_i$ is the total mass and $M_{i-1} = M_i - m_i$. With Φ_i evaluated for every particle in the cluster, the potential at any radius $r_i < r < r_{i+1}$ can be expressed as

$$\Phi(r) = \Phi_i + \left(\frac{1/r_i - 1/r}{1/r_i - 1/r_{i+1}}\right)(\Phi_{i+1} - \Phi_i). \tag{33}$$

The potential at any r is then computed by finding the index i such that $r_i \leqslant r \leqslant r_{i+1}$ and applying Equation (32).

2.8. Global Time Step Selection

In CMC, the time step ΔT is computed as the minimum of several important physical timescales. The first such timescale is the local average relaxation timescale—the average time to deflect a particle by angle $\beta_e^{\max}$ within a certain localized region of the cluster. Following Equation (14), we define this as

$$\widehat{T}_{\mathrm{rel},i} = \left(\frac{\beta_e^{\max}}{\pi/2}\right)^2 \frac{\pi}{32}\frac{\langle w\rangle^3}{4G^2 n\langle m^2\rangle \ln\Lambda}, \tag{34}$$

where $\beta_e^{\max}$ can be set by the user, with larger $\beta_e^{\max}$ corresponding to larger relaxation time steps. By default, CMC sets $\beta_e^{\max} = \sqrt{2}$ (though we typically use $\beta_e^{\max} = 1$ for systems of point masses). Whereas $T_{\mathrm{rel},i}$ is defined for pairs of particles i and $i+1$, $\widehat{T}_{\mathrm{rel},i}$ is a broader rolling average over particles centered on particle i. The local number density n is calculated using the same rolling average as in Equation (12), with similar rolling averages used to compute $\langle m^2\rangle$ and the relative velocity between two particles, $\langle w\rangle$ (calculated from the averaged 3D velocity dispersion, assuming a locally Maxwellian velocity distribution). To accurately capture the overall relaxation dynamics, we require that $\Delta T < \widehat{T}_{\mathrm{rel},i}$ for all i particles. In Freitag & Benz (2001, Equation (11)), this was accomplished by selecting a time step in each radial bin that was smaller than $\widehat{T}_{\mathrm{rel},i}$ by a factor of 0.005–0.05. In CMC, we instead choose a single relaxation time step for the entire cluster, taken to be the minimum $\widehat{T}_{\mathrm{rel},i}$ for all particles i. This choice ensures that our global time step is sufficiently short for all particles and to resolve the complicated dynamics in the cluster core, while making the parallelization of the code significantly easier; see Pattabiraman et al. (2013).

In addition to the relaxation timescale, we also compute characteristic timescales for collisions (T_{coll}), binary–binary encounters (T_{bb}), and binary–single encounters (T_{bs}). In general, the rate of collisions can be found by integrating the cross section for strong encounters over the distribution function of velocities for the two particles (Freitag & Benz 2002; Binney & Tremaine 2008):

$$\Gamma_{\mathrm{strong}} = \frac{1}{n}\int F(\boldsymbol{v}_i)F(\boldsymbol{v}_{i+1})w\Sigma_{\mathrm{strong}}d^3\boldsymbol{v}_i d^3\boldsymbol{v}_{i+1}, \tag{35}$$

where, as before, F is the distribution function (ignoring the mass and position), $w = |\boldsymbol{v}_{i+1} - \boldsymbol{v}_i|$ is the relative velocity, and Σ_{strong} is taken from Equation (21). Assuming a Maxwellian

[15] While Equation (28) is only correct for a point-mass galactic potential, the $M_C^{1/3}$ scaling is correct for any circular orbit in a spherical potential (e.g., Renaud et al. 2011), meaning that this approach is always correct provided that the orbit remains circular and the correct initial r_t is specified.

distribution for the velocity distribution, Equation (35) can be evaluated as (e.g., Binney & Tremaine 2008, p. 626)

$$\Gamma_{\rm strong} = 4\sqrt{\pi} n r_p^2 \sigma \left(1 + \frac{GM}{2 r_p \sigma^2}\right), \tag{36}$$

and the relevant timescale between encounters is simply $T_{\rm strong} = \Gamma_{\rm strong}^{-1}$. The timescales of relevance to CMC are then (see Fregeau & Rasio 2007, for details)

$$T_{\rm coll}^{-1} = 16\sqrt{\pi} n_s \langle R^2 \rangle \sigma \left(1 + \frac{G\langle MR\rangle}{2\sigma^2 \langle R^2\rangle}\right), \tag{37}$$

$$T_{\rm bs}^{-1} = 4\sqrt{\pi} n_s X_{\rm bin}^2 \langle a^2 \rangle \sigma \left(1 + \frac{G\langle M\rangle \langle a\rangle}{2\sigma^2 X_{\rm bin} \langle a^2\rangle}\right), \tag{38}$$

$$T_{\rm bb}^{-1} = 16\sqrt{\pi} n_{\rm bin} X_{\rm bin}^2 \langle a^2 \rangle \sigma \left(1 + \frac{G\langle Ma\rangle}{2\sigma^2 X_{\rm bin} \langle a^2\rangle}\right), \tag{39}$$

where all of the averaged quantities above, including the number densities of single and binary stars, n_s and $n_{\rm bin}$, and the velocity dispersion σ, are computed over the innermost 300 particles in the cluster (i.e., the central regions of the core). Note that if no binaries are present within these 300 stars, $T_{\rm bs}$ and $T_{\rm bb}$ are set to infinity for that time step. Conversely, if *only* binaries are present, then $T_{\rm bb}$ is computed normally, while $T_{\rm bs}$ is set to infinity.

Finally, while the stellar evolution package has its own internal time step, the mass loss from stars can drive significant dynamical changes in the cluster, particularly in the early stages of evolution. We compute a relevant timescale for stellar evolution based on the total fraction of mass lost in the previous time step:

$$T_{\rm se} = 0.001 \left(\frac{M_{\rm cl}}{\Delta M_{\rm se}}\right) \Delta T_{\rm prev}, \tag{40}$$

where $M_{\rm cl}$ is the total cluster mass, $\Delta T_{\rm prev}$ is the previous time step, and $\Delta M_{\rm se}$ is the mass lost during the current time step. CMC uses the minimum individual time step from Equations (34)–(40) as the global time step for the cluster. However, we note that for a typical cluster model (such as those presented in Section 4.2), the central relaxation timescale is the smallest of the relevant timescales for more than 96% of the time steps within the first 100 Myr.

2.9. Energy Conservation

Finally, to ensure energy conservation, a correction must be made to the energies of the stellar orbits. While the two-body relaxation of the Hénon method intrinsically conserves energy, the process of calculating new orbits and sampling positions and velocities is done *after* dynamical encounters have altered the energy and angular momenta of each orbit, but *before* the potential has been updated to account for the changes in the orbits. In other words, the Φ in Equation (31) is perpetually one time step behind the new E and J, allowing this time-dependent shift in the potential to do work on the particle. While the energy error during a single time step is negligible, this effect can become significant during long-term integrations, leading to a spurious drift in the energy of the cluster.

To counter this, we adopt the energy conservation scheme first developed by Stodółkiewicz (1982) in his update of the method. First, note that the change in energy for a single particle is expressible as

$$\Delta E_i^{\rm corr} = \int \frac{\partial U(r_i)}{\partial t} dt. \tag{41}$$

Since we are interested in the work done by the change in potential between the previous and current time steps, Equation (41) can be approximated as the average of the change in potential energy between the positions of the particle at the previous and current time steps:

$$\int \frac{\partial U(r_i)}{\partial t} dt = \left[\Delta\Phi(r_i^{\rm prev}) + \Delta\Phi(r_i^{\rm curr})\right]/2, \tag{42}$$

where $\Delta\Phi \equiv \Phi^{\rm curr} - \Phi^{\rm prev}$. To implement this conservation scheme, CMC records the four relevant energies for each particle every time step—$\Phi^{\rm prev}(r_i^{\rm prev})$, $\Phi^{\rm prev}(r_i^{\rm curr})$, $\Phi^{\rm curr}(r_i^{\rm prev})$, and $\Phi^{\rm curr}(r_i^{\rm curr})$—and uses them to compute $\Delta E_i^{\rm corr}$ for each particle. The velocity magnitudes for each particle are then scaled to increase or decrease the particle's kinetic energy by $\Delta E_i^{\rm corr}$ (while maintaining the ratio of transverse to radial velocities as determined by the Hénon method), ensuring that the cluster conserves energy over many relaxation times.

In addition to the energy drift, a second issue arises from calculating stellar orbits for the current time step using the potential from the previous time step. For any given particle, the potential includes the contribution *of that same particle* at its last position in the cluster, and when the new orbit is calculated, it includes this spurious "self-force" from the previous time step. While this effect is relatively minor for clusters with equal-mass particles, it can be significant in clusters with a realistic mass spectrum, particularly when the most massive particles (e.g., BHs) are concentrated in the cluster center, where only a handful of particles determine the potential (Freitag & Benz 2001; Fregeau & Rasio 2007).

To solve this, we add a correction term to remove the particle's own contribution to the previous potential at its previous position:

$$\Phi_{\rm sf}(r) = \begin{cases} GM_i/r & r \geqslant r_i^{\rm prev}, \\ GM_i/r_i^{\rm prev} & r < r_i^{\rm prev}. \end{cases} \tag{43}$$

This self-force correction, $\Phi_{\rm sf}(r)$, is added to the energy Equation (26) when the new radial positions are sampled from the particle orbits (Section 2.5).

2.10. Limitations of the Method

The Monte Carlo method, and Hénon's method in particular, has been used for 50 yr for the study of DSCs, starting with the initial papers by Hénon (1971a, 1971b). The technique was reintroduced with improvements by Stodółkiewicz (1982), and it is this version that most modern implementations (Giersz 1998; Joshi et al. 2000; Freitag & Benz 2001; Sollima & Mastrobuono Battisti 2014) are based on. While Monte Carlo methods have enjoyed great success in the modeling of DSCs, that success relies on a series of assumptions that we have implicitly or explicitly made in Section 2. We now briefly describe each of these assumptions and the limitations they impose on the method.

Underlying nearly all of the formal averages and encounter rates presented in Section 2 is Boltzmann's molecular chaos assumption, where we have assumed that the distribution

functions of our test stars and field stars are independent and separable. This assumption is explicitly used in Equations (10) and (35), where we assume that the joint distribution function for particles i and $i+1$ can be written as

$$F(\boldsymbol{v}_i, \boldsymbol{v}_{i+1}) = F(\boldsymbol{v}_i)F(\boldsymbol{v}_{i+1}). \quad (44)$$

Equation (44) can be justified in the following way: since equal octaves of impact parameter contribute equally to Equation (4), the majority of two-body encounters will occur at $b_{90} \ll b \ll b_{\max}$. Because of this, it is highly unlikely for two particles in a sufficiently large cluster to encounter each other more than once within their lifetimes, or at least before other encounters have "reset" the velocities of the two particles. While this assumption is valid for large-N systems, it does mean that the Monte Carlo method cannot resolve resonant effects between particles, such as scalar resonant relaxation (Rauch & Tremaine 1996) around central-massive BHs or resonant behavior between heavy objects in the centers of GCs (Meiron & Kocsis 2019).

The assumption of large N is also employed when computing the two-body diffusion of $(\Delta v)^2$. When computing Equation (6), we assumed $\Lambda = b_{\max}/b_{90} \gg 1$, which allowed us to substitute $\ln(1+\Lambda^2) \simeq 2\ln(\Lambda)$ and neglect any higher-order contributions to the change in velocity, e.g., $\langle(\Delta v_i)^4\rangle$. This neglect of the higher-order terms is known as the Fokker–Planck approximation and is largely justified because any terms beyond second order are at least a factor of $\ln(\Lambda)^{-1}$ smaller than the $\langle(\Delta v)^2\rangle$ term (Henon 1973; Binney & Tremaine 2008). This is also related to our assumption of orbit averaging, where we assumed that N is sufficiently large that $T_{\rm relax} \gg T_{\rm dyn}$ in Equation (1), allowing us to treat the orbits as essentially fixed on dynamical timescales.

Naturally, one may ask how large must N be for the above assumptions to be valid. While hard to derive explicitly,

$$N \gtrsim 10^3 \frac{m_{\max}}{\langle m \rangle} \quad (45)$$

has been shown to be a decent working criterion (Freitag 2008). Note that this criterion depends on the ratio between the largest mass in the cluster, $m_{\max}$, and the average stellar mass, $\langle m \rangle$. For clusters with realistic initial mass functions (IMFs), the ratio can be as large as 100:1 (once stellar evolution has reduced the mass of the most massive stars). In our experience, N must be $\gtrsim 10^5$ initially for such clusters to show acceptable energy conservation and good agreement with direct N-body results (Chatterjee et al. 2010).

Hénon's method also assumes that the cluster is spherically symmetric. This assumption is not explicitly related to the size of the cluster or the relaxation timescale, but it is necessary for computing new particle orbits, determining nearest neighbors, and computing the potential. While this is an appropriate assumption for many GCs and some NSCs, this does mean that clusters with net angular momentum (e.g., rotating GCs, if the rotation creates a significant departure from spherical symmetry) or triaxial stellar distributions (e.g., NSCs that have recently undergone a merger) cannot be modeled with CMC. To address this, similar techniques, referred to as orbit-following Monte Carlo methods, have been developed that explicitly follow the orbits of stars in arbitrary potentials on a dynamical timescale, either by directly integrating the orbits in a given potential (the so-called Princeton Monte Carlo, of Spitzer & Hart 1971; Spitzer & Thuan 1972) or by tracking the diffusion of orbits over an integer number of orbital periods (the so-called Cornell Monte Carlo of Marchant & Shapiro 1980). These techniques have been used extensively for the study of triaxial NSCs and central-massive BH dynamics (see, e.g., Hopman 2009; Madigan et al. 2011; Vasiliev 2015, for recent examples). We note that the assumption of spherical symmetry does *not* mean that we have assumed the velocities to be isotropic. CMC can model systems with significant velocity anisotropy, since the ratio of v_r to v_t is calculated self-consistently every time step.

Finally, because the Hénon method operates on a relaxation timescale, we are implicitly assuming that GCs can be modeled as quasi-adiabatic transitions between different cluster models in dynamical (virial) equilibrium. While this assumption is largely valid for the long-term evolution of clusters in most cases, there are notable exceptions, such as tidal shocking (Ostriker et al. 1972), or the initial "violent relaxation" of the cluster after its (potentially highly nonvirialized) birth (Lynden-Bell 1967). These processes involve clusters that are sometimes far from equilibrium and evolve to a virialized state on a dynamical timescale. For these cases, more computationally expensive orbit-following techniques, such as the aforementioned Princeton/Cornell methods or direct N-body integrations, must be employed.

3. CMC Package Overview

The public release of CMC coincides with the latest release of the COSMIC package for binary population synthesis (Breivik et al. 2020a), which is included directly in the CMC GitHub repository (and can be compiled simultaneously). Both codes implement single stellar evolution using SSE (Pols et al. 1998; Hurley et al. 2000) and binary evolution using BSE (Hurley et al. 2002), with many major updates having been made to the physical prescriptions of stellar and binary evolution based on advances over the past two decades; see Breivik et al. (2020a) for details.

With the inclusion of COSMIC into CMC, we have also updated both codes to use similar input/output files and initial condition generators. Both codes use identical parameters and initialization files, allowing cluster simulations and population synthesis to be easily performed with identical physics. The generation of all initial conditions for both binary populations and cluster simulations is now included in the latest version of COSMIC, allowing populations with identical physics to be created from a simple Python interface.

3.1. Updates to COSMIC Population Synthesis Code

While COSMIC was originally based on the version of BSE incorporated into CMC, including its updates to compact object physics (Chatterjee et al. 2010; Rodriguez et al. 2016c), the two codes have diverged over the past several years, and the use of COSMIC as a community-driven population synthesis code has kept it up to date with the latest developments in binary stellar evolution. Here we review the changes to binary stellar evolution (and in particular binary mass transfer) in the latest version of COSMIC (v3.4), which is paired with the release of CMC. For a complete description of the physics included in COSMIC and the various options, see Hurley et al. (2000, 2002) and Breivik et al. (2020a).

Since the release of COSMIC v3.3, the treatment of Roche overflow mass transfer from nondegenerate stars has been expanded to include multiple assumptions for both mass-loss rates from the donor and mass accretion rates onto the accretor. The default mass-loss rate is determined using the expression given by Hurley et al. (2002), which steeply increases with the ratio of the donor radius to its Roche lobe radius

$$\dot{M}_{\rm don} = F(M_{\rm don})\left[\ln\frac{R_{\rm don}}{R_{\rm rl,don}}\right]^3 M_\odot\ {\rm yr}^{-1}, \quad (46)$$

where $\dot{M}_{\rm don}$ is the donor mass-loss rate, $M_{\rm don}$ is the donor mass in units of $M_\odot$, $R_{\rm don}$ is the donor radius, $R_{\rm rl,don}$ is the donor's Roche lobe radius, and

$$F(M_{\rm don}) = 3\times10^{-6}\left[\min(M_{\rm don}, 5.0)\right]^2. \quad (47)$$

Version 3.4 of COSMIC also includes the Roche lobe overflow prescription of Claeys et al. (2014), which is calibrated to binary systems with zero-age main-sequence component masses less than $10\,M_\odot$. In this case, the mass-loss rate of the donor is determined similarly to Equation (46), but with an overall multiplicative factor, f, given by

$$f = \begin{cases} 1000, & Q < 1 \\ \max\left(1, \dfrac{1000}{Q}\exp\left[-\dfrac{1}{2}\left(\dfrac{\ln Q}{0.15}\right)^2\right]\right) & Q \geqslant 1, \end{cases} \quad (48)$$

where $Q \equiv M_{\rm acc}/M_{\rm don}$. In both cases, mass loss is always limited by the thermal timescale such that

$$\dot{M}_{\rm don}^{\rm max} = \max\left(\dot{M}_{\rm don}, \frac{M_{\rm don}}{\tau_{\rm KH}} M_\odot {\rm yr}^{-1}\right), \quad (49)$$

where

$$\tau_{\rm KH} = 3.1\times10^7\left(\frac{M}{M_\odot}\right)^2\left(\frac{R_\odot}{R}\right)\left(\frac{L_\odot}{L}\right){\rm yr} \quad (50)$$

is the Kelvin–Helmholtz time.

Four prescriptions are now included to determine the amount of mass that the accretor is allowed to gain, in addition to the original one in Hurley et al. (2002). The original BSE prescription assumes that stars with radiative envelopes (i.e., main sequence, Hertzsprung gap, and core helium burning) can accrete mass at a rate limited to 10 times the thermal rate of the star, while stars with convective envelopes (i.e., giant branch, asymptotic giant branch) can accrete mass at an unlimited rate. Three of the newly added prescriptions consider modifications to this prescription: one that limits stars with radiative envelopes to accrete mass at the thermal rate, one that applies accretion limits of 10 times the thermal rate for all stars—not just those with radiative envelopes—and one that applies accretion limits of the thermal rate for all stars. We also include another choice where accretion is a fixed fraction of the mass lost by the donor, as used in Belczynski et al. (2008).

3.2. Cluster Initial Conditions

In COSMIC, initial conditions are generated by sampling particle positions and velocities from either a Plummer (1911), Elson et al. (1987), or King (1966) profile. Both the Elson and Plummer profiles are sampled piecewise, with the particle positions drawn first from the cumulative mass distribution for each given profile. To self-consistently sample velocities for each particle at its given position, we draw velocities from the distribution function via rejection sampling up to the local escape speed (Aarseth 2003). For the Plummer profile we use the analytic form of the distribution function, while for the Elson profile we numerically solve for the distribution function in energy space (see, e.g., Grudić et al. 2018, Appendix B).[16] The King profile is generated in a similar fashion: the differential equations for cluster density, potential, and enclosed mass are solved numerically, with the latter being used to sample random positions from the cluster center. The velocities are then sampled from the analytic distribution function via the aforementioned rejection sampling technique.

By default, COSMIC draws radii and velocities from the available distributions before converting them into *Hénon* units (Heggie & Mathieu 1986; Heggie 2014), where the gravitational constant and initial cluster mass are both set to unity ($G = M_0 = 1$) and the sum of the initial cluster kinetic and potential energies is $E_{0,\rm kin} + E_{0,\rm pot} = -1/4$. (This is not the same as the total initial cluster energy $E_0 = E_{0,\rm kin} + E_{0,\rm pot} + E_{0,\rm bin}$, where $E_{0,\rm bin}$ is the initial negative-valued total binding energy in primordial binaries, if any.) The units of mass, length, and time are then

$$\begin{aligned} U_m &= M_0, \\ U_l &= \frac{-GM_0^2}{4E_0}, \\ U_t &= \frac{GM_0^{5/2}}{(-4E_0)^{3/2}}, \end{aligned} \quad (51)$$

respectively. The above length units are such that the initial virial radius of the cluster is unity, while the units of time roughly correspond to the crossing time of a particle at the virial radius. However, because CMC operates on a relaxation timescale, we instead use a more appropriate unit of time,

$$U_t^{\rm rel} \equiv \frac{N_0}{\log\gamma N_0} U_t, \quad (52)$$

where N_0 is the initial number of particles (single stars or binaries). These units are the default used throughout CMC.

Once the positions and velocities have been selected, the user can optionally add additional stellar physics to each of the particles, such as stellar masses, radii, and binary companions. All of the options available in COSMIC, including IMFs (e.g., Kroupa 2001; Salpeter 1955), binary initial conditions (e.g., following Sana et al. 2012; Moe & Di Stefano 2017), and stellar metallicities, can be used to generate realistic cluster initial conditions for CMC. When generating binary initial conditions for cluster simulations, we truncate the upper limit of the $P_{\rm orb}$ distribution at the hard−soft boundary for each binary at its respective radius in the cluster (see, e.g., Heggie 1975).

3.3. Input/Output

Input in CMC is handled at the command-line level by passing an initialization file containing all the parameters

[16] We note that previous N-body samplers of the Elson profile, such as the one described in Küpper et al. (2011), have used the 1D Jeans equations to sample the velocity dispersion of the cluster. While this yields a cluster in virial equilibrium, it does not produce a cluster consistent with a single distribution function.

controlling the physics, parallelization, and runtime options (unspecified parameters are set to default values). The parameter files contain many of the same options (particularly related to stellar evolution) as the parameter files for COSMIC, allowing initial conditions, population synthesis, and cluster dynamics simulations to be performed with identical options. The initial conditions are saved as tables in the Hierarchical Data Format (HDF5) and passed as input to the initial parameters file.

Three forms of output are generated during a typical CMC run. First, a series of tab-separated data files containing specific parameters of the cluster (e.g., mass, half-mass radius, number of binaries) are printed every time step and can be easily read in any scripting language. Second, specific events, such as stellar collisions, BH mergers, binary–single encounters, etc., are recorded as they occur in human-readable log files (with a provided Python parser). Third, snapshots of the full cluster state are periodically written to several HDF5 tables at user-specified intervals. Separate files exist for snapshots containing all stars in the cluster and those containing only the BHs (allowing the latter to be resolved at shorter intervals). Each snapshot is saved in its respective file as a pandas-readable HDF5 table with keys corresponding to the simulation time when the snapshot was written. These can be directly imported into Python using the `read_hdf` command in pandas (McKinney 2010; Reback et al. 2021).

Finally, CMC periodically saves checkpoints by directly dumping the state of every processor to a binary file associated with that processor. These checkpoints allow bit-for-bit restarting of a CMC run should the run be either interrupted or terminated owing to computing cluster queue-time limits. While the checkpoints contain the state of the cluster, stars, and binaries up to that point, both the physical parameters of the run (set in the initialization file) and the current random number seed can be changed by the user at restart.

3.4. Parallelization

CMC is designed to be compiled against the standard libraries of the Message Passing Interface (MPI; Clarke et al. 1994), enabling distributed-memory parallelization across multiple cores and machines. Many of the physical processes in CMC, such as stellar evolution and nearest-neighbor encounters, are fundamentally local processes that are amenable to simple parallelization. However, the radial sorting of stars, the computation of new stellar orbits in the global cluster potential (Section 2.5), and even the computation of the potential itself (Section 2.7) are global processes that require a customized parallel approach.

Because the calculation of the cluster potential in Equation (33) and the nearest-neighbor dynamics require the particles to be sorted radially, the particles are resorted each time step after new positions are selected (Section 2.5) in order of increasing radius from the cluster center. CMC uses a custom parallel sorting implementation, in which the local data on each processor are sorted using a `Quicksort` algorithm (Hoare 1961) before being combined globally across processors using a parallel `Sample Sort` algorithm (Frazer & McKellar 1970). This implementation allows for rapid sorting and redistribution of particles across multiple distributed-memory processes. For more details on the parallel sorting method, see Section 3.4 of Pattabiraman et al. (2013). Once the sorting is complete, the particles are redistributed across all MPI processes with a series of all-to-all communications.

The calculation of the stellar orbits is more straightforward and only requires that every MPI process have access to the global potential of the cluster at each time step. To that end, each process maintains a global list of the radial positions and masses for every star in the cluster, allowing each process to compute Equation (33) independently. These lists are updated every time step. For a complete description of this and the other parallelization strategies employed in CMC, see Pattabiraman et al. (2013).

4. Examples

Having described the code in detail, we now show two standard examples of clusters generated with CMC: a Plummer sphere of point masses evolved from its initial state to core collapse, and the evolution of 10 realistic King models of GCs evolved with all available physics for 12 Gyr.

4.1. Plummer Sphere to Core Collapse

The Plummer profile was one of the first used to fit realistic observations of GCs to a theoretical model, and it remains popular since, unlike models based on more sophisticated distribution functions, most of its key features can be expressed analytically (Heggie & Hut 2003). This also makes it a favorite tool of the N-body simulator since the enclosed mass, velocity dispersion, and distribution function can all be expressed analytically and are easily sampled to generate cluster initial conditions. Here, we use COSMIC to generate a realization of a Plummer sphere with 10^8 initial point-mass particles. We choose this particular setup because the evolution of a Plummer sphere to core collapse is a well-studied problem in the literature (e.g., Freitag & Benz 2001; Freitag et al. 2006; Binney & Tremaine 2008). We then integrate these initial conditions twice using CMC: once with no binary physics, and once with three-body binary formation and strong binary encounters enabled.

In Figure 2, we show the evolution of the two clusters from their initial state to core collapse, defined as having < 1000 stars within the theoretical core radius (see Casertano & Hut 1985, for the description of the latter). We also show the number of binaries as a function of time, since it is the formation and hardening of binaries in the core that provide an energy source to support the cluster against continued collapse. Since we stop both models at core collapse, we do not show the post-collapse evolution of either cluster; however, the cluster model with binary physics has its core collapse time delayed by 0.15 in units of the half-mass relaxation time:

$$T_{\rm rlx} = 0.138\frac{N}{\log(\gamma N)}\left(\frac{r_h^3}{GM}\right)^{1/2}, \tag{53}$$

where we have used the more precise definition from Spitzer (1987), rather than the approximate Equation (1). Similarly, the number of binaries in the core increases dramatically during the final moments of core collapse, as the central density increases sharply. In both cases the cluster collapses at about $16.8T_{\rm rlx}$, consistent with both theoretical predictions and previous numerical integrations (Cohn 1979, 1980; Spurzem & Aarseth 1996; Quinlan 1996; Freitag & Benz 2001).

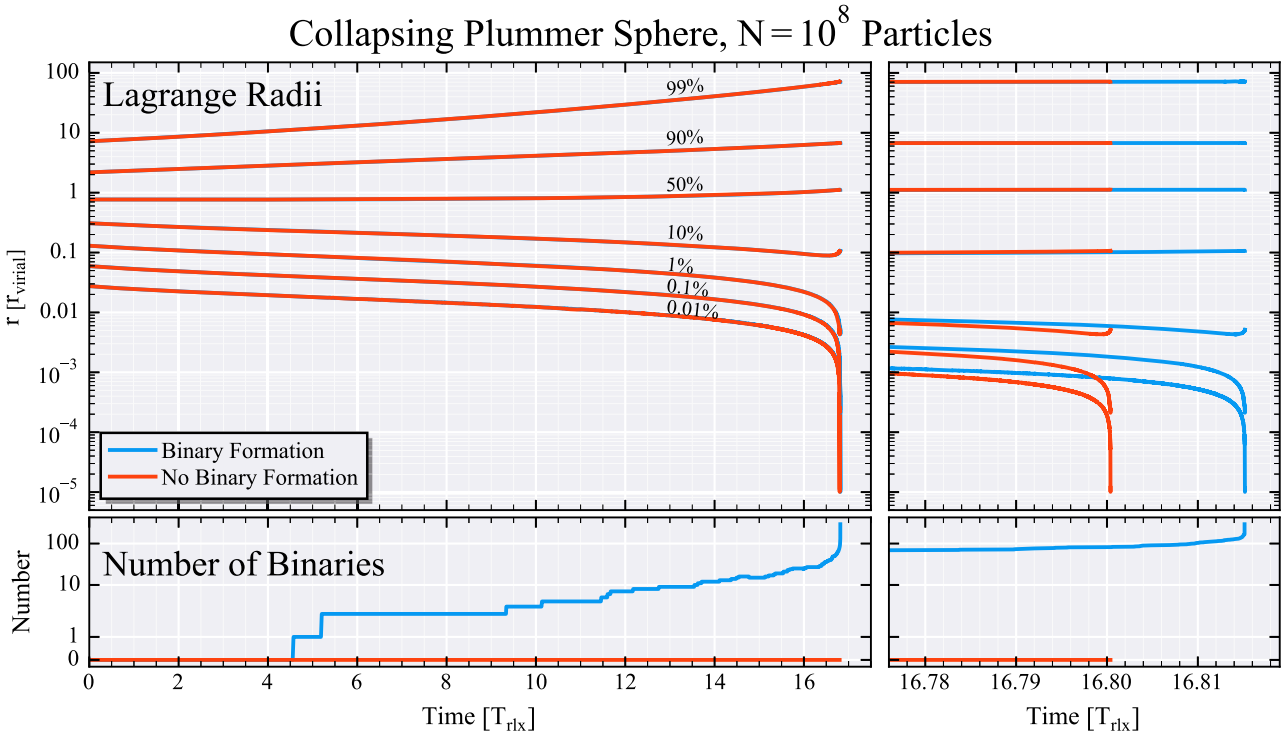

Figure 2. The Lagrange radii (radius enclosing a given fraction of the mass) and number of binaries as a function of the initial half-mass relaxation time for the two $N = 10^8$ particle realizations of a Plummer sphere. We show in red the model without three-body binary physics, and in blue the same model, but allowing binaries to form from single stars. On the top, we show the Lagrange radii as a function of time, with the distinction between the core and halo rapidly developed owing to two-body relaxation. On the top right, we zoom in at the final moments of core collapse (defined as having less than 1000 objects in the core) and clearly show the (minor) delay in core collapse caused by binary heating in the cluster core. On the bottom, we show the number of binaries formed over time. As the core shrinks, the central encounter rate increases, dramatically increasing the number of binaries in the cluster in the final moments before collapse.

It is also well known that the collapse of a cluster proceeds homologously, with the central density and time to core collapse evolving at a self-similar rate throughout the cluster's central regions (Hénon 1961; Lynden-Bell 1975). Both homologous models and 1D Fokker–Planck models suggest that the central density profile of a cluster near collapse should approach a power law of the form $\rho \propto r^{-2.23}$ (Cohn 1980; Lynden-Bell & Eggleton 1980; Heggie & Stevenson 1988). In Figure 3, we show the density of stars (averaged over 20 nearest neighbors) in our cluster model without binary formation as it approaches core collapse. The density of the central regions increases dramatically as the core approaches final collapse, with the central density increasing by nearly 4 orders of magnitude within the last $10^{-5}T_{\rm rlx}$ of the evolution. At the moment of core collapse, CMC reproduces the $r^{-2.23}$ distribution over more than 15 orders of magnitude. To the best of our knowledge, this is the largest fully collisional star-by-star model of a Plummer sphere presented to date.

Finally, a critical test of any N-body scheme is how well it conserves energy, since this serves as an independent benchmark of the self-consistency of the physics. By itself, two-body relaxation as implemented in the Hénon algorithm automatically conserves energy; in practice, with the energy-conserving corrections described in Section 2.9, a CMC run typically conserves energy to within one part in $\sim 10^5$ over many relaxation times. We show the differential change in energy for both models as a function of time in Figure 4.

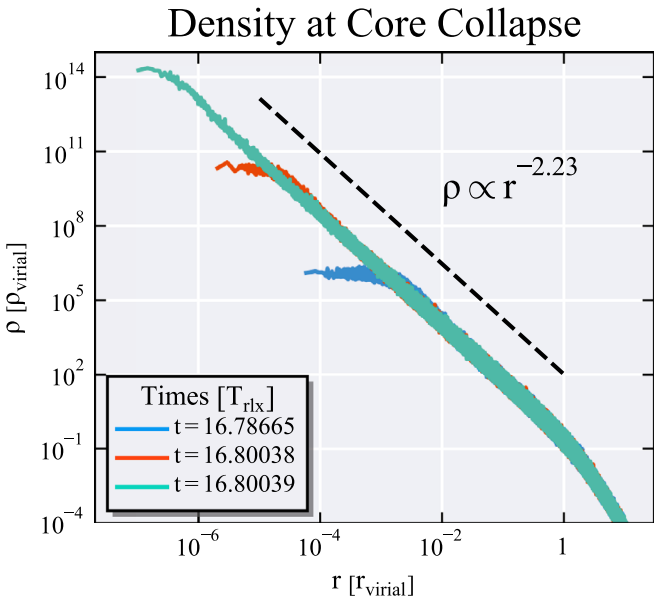

Figure 3. The density profile of the Plummer sphere with 10^8 particles near core collapse as a function of radius. Both quantities are calculated in terms of the initial virial radius and the initial density of objects within that radius. We show the densities from three separate snapshots, whose time is given in terms of the initial half-mass relaxation time. During the last moments of the cluster's evolution, the inner regions rapidly converge to the $\rho \propto r^{-2.23}$ power law expected from Fokker–Planck calculations (Cohn 1980), showing self-similar behavior over more than 15 orders of magnitude in density.

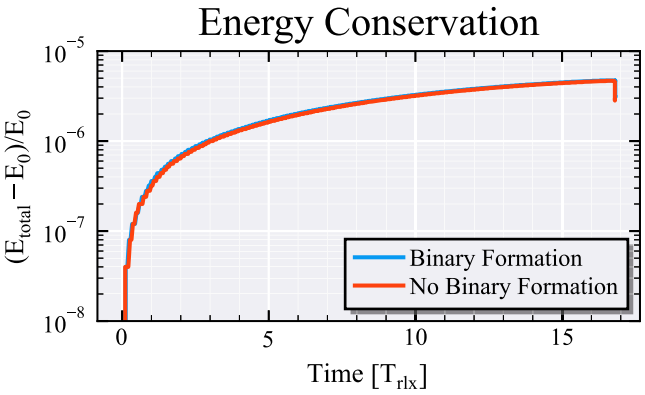

Figure 4. The total energy of the two Plummer sphere models (including the energy of ejected particles) over time. The sharp downturn at the end arises from the reduction of the time step near core collapse to better resolve the dense central regions of the cluster.

4.2. Realistic Globular Clusters

To demonstrate the capability of CMC to generate realistic cluster models, we generate 10 models of star clusters with random initial conditions drawn from the same distribution function. We begin by sampling positions and velocities for 10^6 particles using a King (1966) profile (with $W_0 = 6$) and a virial radius of 1 pc. The properties for each star and binary are sampled using the same initial condition generators as COSMIC: each star is first assigned a mass from a Kroupa (2001) IMF. We randomly assign 10% of those stars to be binaries, keeping the initial IMF draw as the primary mass, and assigning a secondary mass from a uniform distribution between 0.1 and 1 times the primary mass. For isolated binaries, the "independent" sampler in COSMIC assigns binary semimajor axes from any number of user-specified distributions, e.g., flat in the log-period (Abt 1983), or the distributions from Sana et al. (2012) and Renzo et al. (2019). However, cluster environments come with an additional complication: many wide binaries that are observed in galactic fields (and are reproduced by the distributions in COSMIC) are significantly larger than the typical interparticle separation of the cluster and would be very rapidly destroyed by dynamical encounters (or simply not formed at all). To avoid this, we truncate the orbital periods at the local hard−soft boundary for a binary with the given mass at that position in the cluster (this approximately corresponds to limiting the minimum orbital velocity to the local three-dimensional velocity dispersion). For the binaries considered here, we assume a flat-in-log distribution of semimajor axes and a thermal distribution of eccentricities (Ambartsumian 1937; Heggie 1975). We set stellar metallicities equal to $Z = 0.0002$ (approximately 1% of the solar value) and do not place the clusters in an external galactic potential.

In Figure 5, we show a projected snapshot of one of the cluster models rendered using the fresco package (Rieder & Pelupessy 2019), which incorporates the stellar positions, effective temperatures, and radii to construct a mock HST observation. We also show the surface brightness and velocity dispersion profiles for that snapshot, calculated using the cmctoolkit (Rui et al. 2021a, 2021b). The toolkit can directly convert the HDF5 output snapshots of CMC into observational quantities, including surface brightness and velocity dispersion profiles, number density profiles, counts of blue stragglers and other stellar types, and so on. Combined with the public release of CMC, these two software packages contain all the necessary tools to create models of GCs and SSCs that can be directly compared to observations in the local universe.

One of the main advantages of *N*-body modeling is the ability to study the internal dynamics of lower-luminosity components of the cluster. In Figure 6, we show the populations of WDs, NSs, and BHs that are present throughout the cluster model shown in Figure 5 (each overlaid over the synthetic observation), as well as the projected and smoothed 2D number densities as calculated with the cmctoolkit. As expected, the heavier components—the carbon–oxygen and oxygen–neon WDs, NSs, and BHs—have segregated into the central regions of the cluster. The BHs, in particular, are highly concentrated, with 50% of the BHs lying within 0.2 pc of the cluster center, where the number density of BHs is more than 7000 pc^{-3}.

In a realistic GC, massive stars evolve in ∼10–100 Myr after cluster formation, leaving behind a large population of BH remnants. These BHs rapidly sink to the center, where they begin to participate in the three- and four-body interactions that form and dynamically harden binaries. This process continues without limit until all the BH binaries are either ejected from the cluster (having each also ejected several single BHs during three-body encounters) or merge owing to GW emission. The rate of BH depletion is determined by the energy required to support the cluster against continued collapse (e.g., Breen & Heggie 2013), and in clusters with sufficiently small initial virial radii, this process can completely remove the BH subsystem from the cluster, producing the core-collapsed clusters observed in the Milky Way (Kremer et al. 2019a). Overall, the dynamics of BH subsystems are expected to play a key role in shaping the overall evolution of DSCs (e.g., Merritt et al. 2004; Mackey et al. 2007, 2008; Breen & Heggie 2013; Wang et al. 2016; Chatterjee et al. 2017b; Arca Sedda et al. 2018; Kremer et al. 2018b, 2020b, 2020c; Antonini & Gieles 2020; Weatherford et al. 2021). The particular cluster simulation shown in Figure 5 still retains 125 BHs by the present day, producing the large core radius observed in the surface brightness profile.

4.3. Binary Black Hole Mergers and GW190521

In addition to the global dynamical and observable properties of present-day clusters, CMC is ideal for understanding the rate and properties of many high-energy transients and compact binary sources that originate in DSCs. As previously stated, the BHs in the central regions of a cluster will continually form binaries and undergo hardening encounters, providing a finite energy source to support the cluster against gravothermal collapse. While this behavior is critical to understanding the overall evolution of the cluster and its present-day appearance, it also produces a large number of BBHs and is thought to provide a key formation channel for the GW sources detected by LIGO, Virgo, and KAGRA (LVK). This process has been studied theoretically for many years (Kulkarni et al. 1993; Sigurdsson & Hernquist 1993; Portegies Zwart & Mcmillan 2000; O'Leary et al. 2007; Downing et al. 2010, 2011; Aarseth 2012; Tanikawa 2013; Bae et al. 2014; Askar et al. 2016; Fragione & Kocsis 2018), including in several papers using CMC (e.g., Rodriguez et al. 2015, 2016a, 2016b, 2018a, 2018b, 2019, 2020, 2021b; Chatterjee et al. 2017a, 2017b; Kremer et al. 2019c, 2020c, 2020a; Martinez et al. 2020; González et al. 2021;

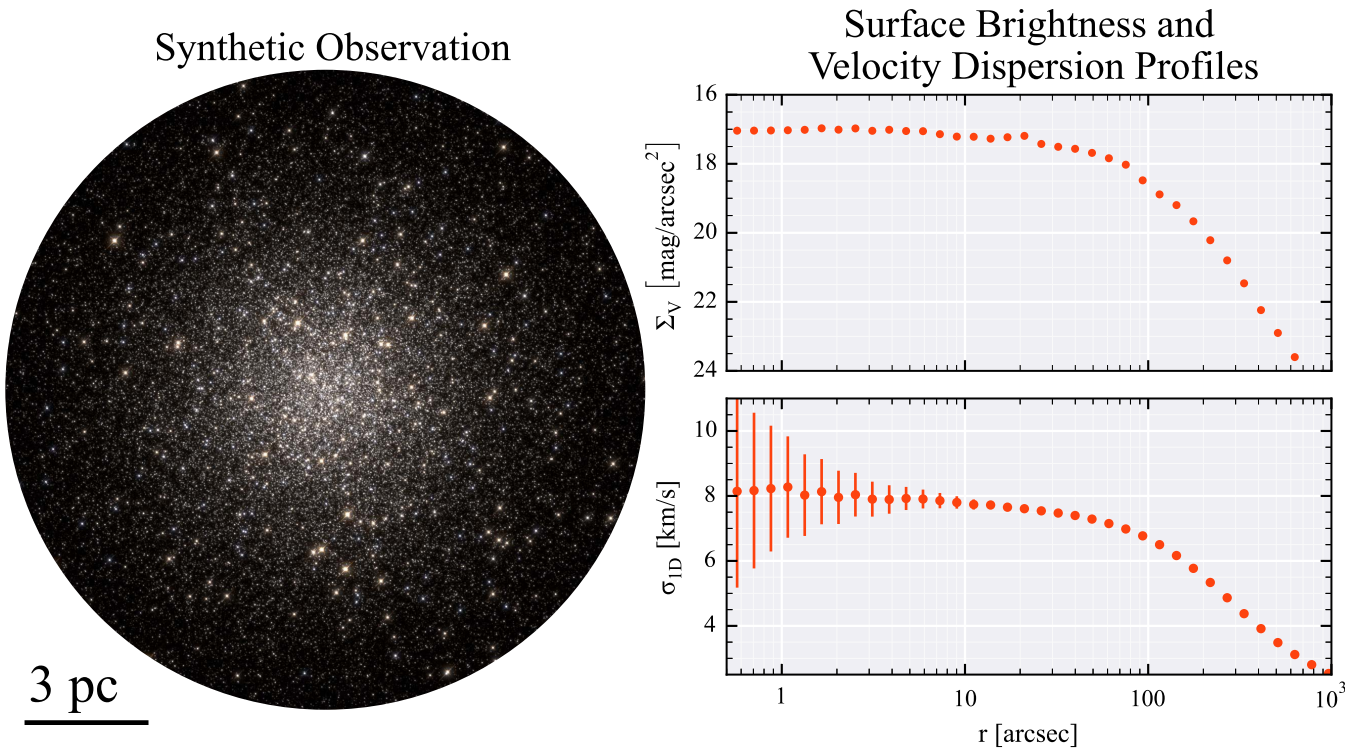

Figure 5. An example snapshot of a realistic cluster model after 12 Gyr of evolution. On the left, we show a synthetic HST observation of the cluster, computing using the stellar positions, effective temperatures, and radii (via the `Fresco` package; Rieder & Pelupessy 2019). On the right, we show the *V*-band surface brightness profile (at a distance of 5 kpc) and 1D velocity dispersion for the cluster, calculated using the `cmctoolkit` package (Rui et al. 2021a).

Holgado et al. 2021; Martinez et al. 2021; Weatherford et al. 2021).

With over 50 GW detections so far, much work has been done trying to determine the origin of the observed population of merging BBHs. One well-studied technique to discriminate between BBHs formed dynamically in clusters and those formed from isolated binary evolution in the field is to measure the spins of the BHs (e.g., Rodriguez et al. 2016d; Farr et al. 2017, 2018), since dynamically formed BBHs are expected to have isotropically distributed spin–orbit misalignments (vs. the relatively small spin–orbit misalignments expected from binary star evolution). Using the distribution of BBH spins, the latest catalog from the LVK suggests that between 25% and 93% of these systems may have been assembled dynamically (The LIGO Scientific Collaboration et al. 2021), with more sophisticated analyses of BBH formation models (including those created using `CMC`) reaching similar conclusions (Wong et al. 2021; Zevin et al. 2021).

In addition to statistical evidence, individual BBH events have also suggested a dynamical origin. One particular system, GW190521, was detected with at least one component mass inside the "upper mass gap," where individual BHs are prevented from forming by stellar collapse due to pulsational pair instabilities, which eject large amounts of material from (and potentially destroy) any star with a helium core mass between about 40 and 130 $M_\odot$ (e.g., Woosley 2017). While it is extremely difficult (though not impossible) to describe such systems as the product of isolated stellar evolution (e.g., Farmer et al. 2019; Belczynski 2020), these systems can be easily produced in DSCs, through either the repeated mergers of BHs (e.g., Rodriguez et al. 2018b, 2019; Fragione & Silk 2020; Fragione et al. 2020a) or massive star mergers occurring prior to BH formation (e.g., Di Carlo et al. 2019; Kremer et al. 2020a; González et al. 2021; Weatherford et al. 2021).

Figure 7 shows the total masses and spins of all BBHs that merge within 13.8 Gyr of cluster formation from our 10 GC models. We subdivide the population according to the "generation" of the BHs in the binary, with first generation, or 1G, BHs being those created from stellar collapse, and second generation, or 2G, referring to those created in a previous merger. The more massive BHs, particularly those within the mass range consistent with GW190521, are predominantly the result of previous mergers, with four of the six systems being composed of 2G+2G BBHs (that is, binaries whose components are *both* the product of previous mergers). The remaining two BBHs are created from previous *stellar* mergers of massive stars prior to their collapse into BHs. These stellar-merger products can contain unique envelope/core mass ratios, allowing them to bypass the typical pulsational pair-instability limits for single massive stars. While repeated BH mergers are the dominant mechanism for producing GW190521-like systems in these models, clusters with higher central densities, fractal initial conditions, or different stellar binary fractions can prefer formation of repeated mergers through stellar collisions (e.g., Di Carlo et al. 2019; Kremer et al. 2020a; González et al. 2021). We note that while our models also ignore the possibility of accretion from BH/star collisions, the majority of these occur after the most massive stars (including stellar-merger products) have expired, since the BHs typically require ∼100 Myr to dynamically segregate into the cluster center. As such, $M_{\rm BH} > M_{\rm star}$ for the majority of BH/star collisions.

We also show in Figure 7 the distribution of the effective spins of the cluster BBH population as a function of the mass.

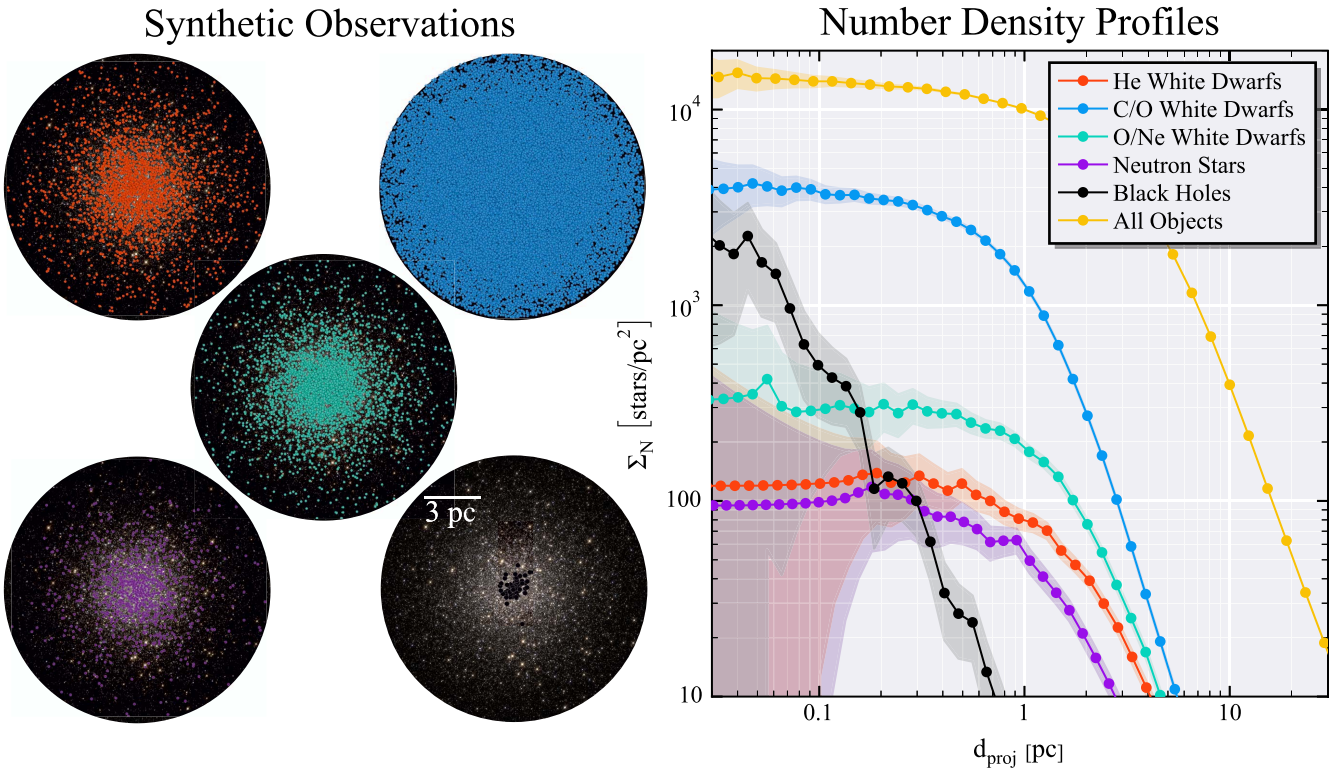

Figure 6. The same snapshot from the left of Figure 5, but showing the various populations of low-luminosity particles hidden in the cluster's surface brightness profile. On the left, we show the distribution of helium, carbon/oxygen, and oxygen/neon white dwarfs (in red, blue, and teal), as well as the NSs (purple) and BHs (black). On the right, we show smoothed 2D number densities and their uncertainties for these populations, calculated from the same snapshot using the `cmctoolkit`.

This mass-weighted projection of the BH spins onto the orbital angular momentum, $\chi_{\rm eff} \equiv \hat{L} \cdot [(m_1\boldsymbol{\chi}_1 + m_2\boldsymbol{\chi}_2)/(m_1 + m_2)]$, is the spin parameter most easily constrained by GW detectors (e.g., Ajith et al. 2011; Abbott et al. 2016). Since we have assumed that BHs from collapsing stars are born with zero spin, consistent with recent theoretical studies of angular momentum transport in massive stars (e.g., Fuller & Ma 2019; Fuller et al. 2019), BBHs whose components are formed from stellar collapse also have zero effective spins across all masses. On the other hand, when BHs are created from previous mergers, their spins inherit the angular momenta of their parent binaries, producing a population of spinning BBHs. Because the BH spins are assumed to be isotropically distributed on the sphere, the median effective spin for all cluster BBHs is still centered at $\chi_{\rm eff} = 0$. This is consistent with the observation of GW190521, which appears to be consistent with having zero effective spin. However, as consistent with our previous results (e.g., Rodriguez et al. 2018b, Figure 3), any *population* of 2G+2G BBHs (like the ones that predominately form GW190521-like binaries here) should have a range of effective spins centered at $\chi_{\rm eff} = 0$.

As a demonstration of the utility of the `CMC/COSMIC` integration, we directly compare the BBH population created through dynamical processes to that created from the evolution of isolated stellar binaries. We evolve a population of 5×10^4 binaries with masses $\geqslant 18\ M_\odot$ using the same stellar physics, stellar metallicity, and binary initial conditions (the truncation in binary orbital period) with `COSMIC`. In Figure 7, we show one of the most distinct features of the BBHs produced by GCs: the production of repeated mergers of BHs in star clusters. When two BHs merge in a GC, their merger product receives a recoil kick due to the asymmetric emission of GWs, which depends on both the mass ratio of the system and the spins of the BHs (e.g., Merritt et al. 2004; Campanelli et al. 2007; Lousto & Zlochower 2008). Because three-body encounters preferentially form BBHs with near-equal-mass components (e.g., Sigurdsson & Phinney 1993), it is the BH spins that primarily determine merger retention in GCs. While large spins ($\chi \sim 1$) leave few second-generation BHs in the star cluster, mergers of BHs born with smaller spins ($\chi \lesssim 0.2$) can be retained in many GCs, where they will continue to participate in dynamical processes, form binaries, and merge again (Rodriguez et al. 2019). This process can produce a significant population of BBHs with large ($\chi \sim 0.7$) spins and masses in the pulsational pair-instability mass gap.

While an observation of a BBH with $\chi_{\rm eff} < 0$ is strongly suggestive of dynamical formation, isolated binary evolution can also produce BBHs with spin–orbit tilts greater than 90° given the right orientation and magnitude of the supernova kicks (Kalogera 2000; Rodriguez et al. 2016d). Therefore, any analysis of BBH spin–orbit misalignments must include both dynamical and binary evolutionary channels using identical stellar physics. The integration of `COSMIC` into `CMC` makes this trivial; in Figure 8, we show the masses and spin–orbit misalignments ($\theta_{\rm LS}$) from the 10 GC models and 5×10^4 isolated massive binaries evolved using `COSMIC`. In the top panel, the limit imposed by pulsational pair-instability physics on the BH mass function is readily apparent, with no BHs forming with masses greater than 44.5 $M_\odot$, and no BBHs forming with total masses greater than 89 $M_\odot$ (a number that is largely insensitive to choices of accretion and other binary physics; van Son et al. 2020). Figure 8 also clearly

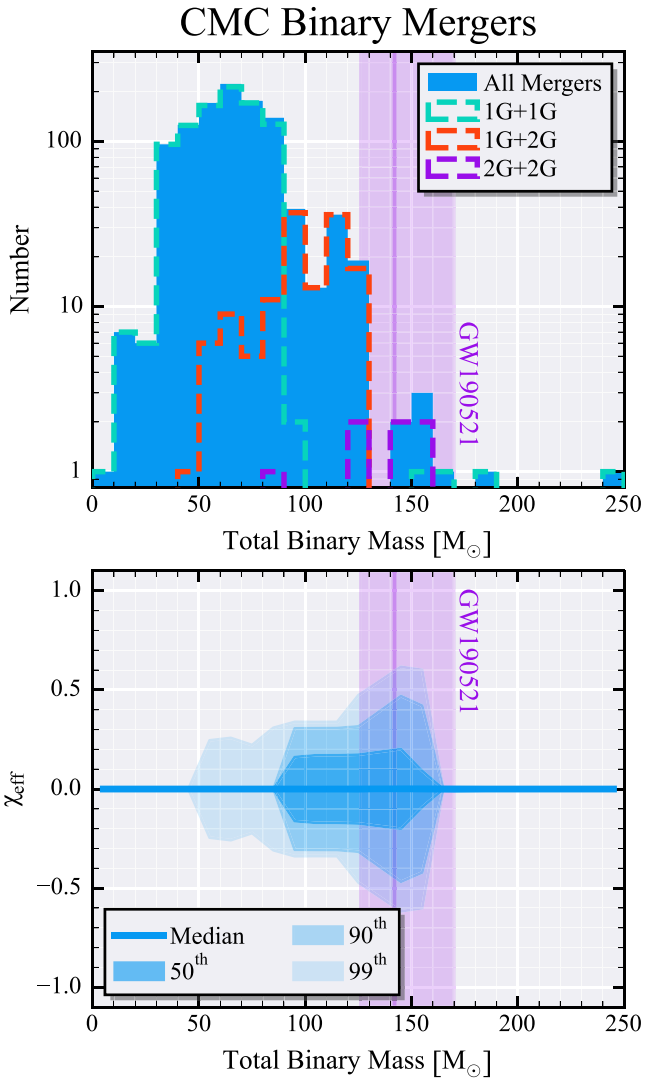

Figure 7. The masses and effective spins for merging BBHs from 10 realizations of a GC model with 10^6 initial particles (10% of which are binaries). On the top, we show the distribution of total masses of all BBHs that merge within 13.8 Gyr of cluster formation. We also indicate the specific generation of BHs involved in the mergers with the dashed histograms. The region consistent with GW190521 is populated by 2G+2G BBHs (where both BHs were formed in previous mergers) and 1G+1G BBHs whose component progenitors participated in previous stellar mergers. On the bottom, we show the distribution of effective spins for the full BBH population as a function of mass. At higher mass, the contribution from 2G BHs increases, and the range of detectable spins increases as well.

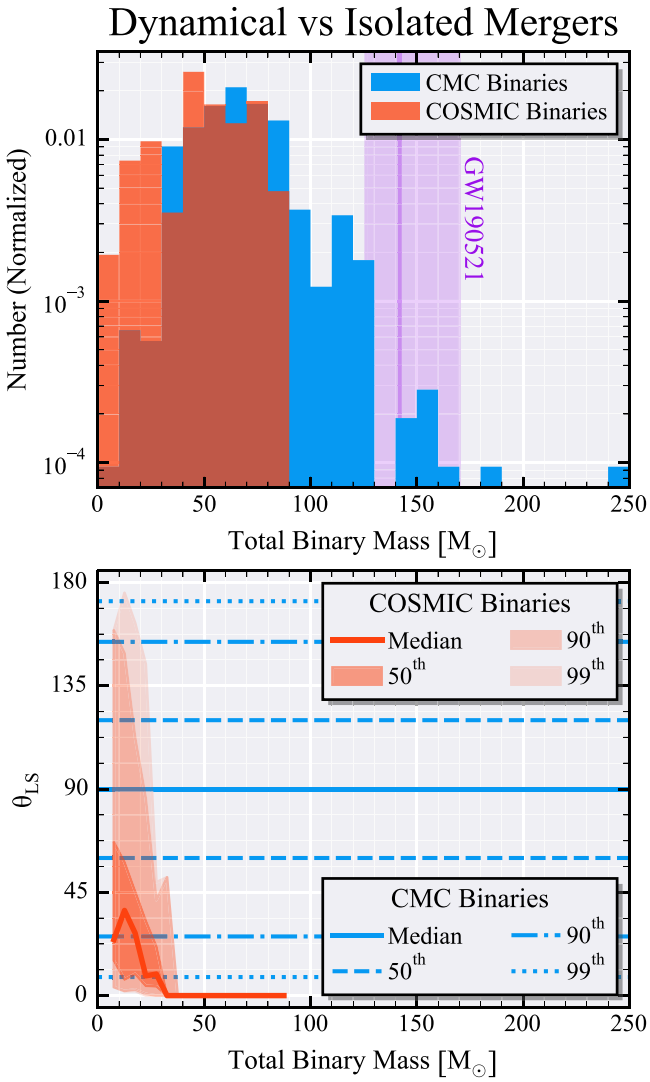

Figure 8. Masses and spin alignments of BBHs produced through dynamical encounters (from 10 GC models generated with `CMC`) and through isolated binary evolution (from 5×10^4 stellar binaries evolved with `COSMIC`). On the top, we show the total masses of the BBHs, with the maximum mass from isolated binaries (89 $M_\odot$) being clearly visible. On the bottom, we show the distribution of spin–orbit misalignments for the BBHs as a function of total mass. While low-mass BBHs from the isolated binary population can experience significant misalignments from supernova natal kicks, this decreases as a function of mass, while the spin–orbit misalignments for dynamically formed binaries are isotropic irrespective of mass.

demonstrates the difference in the spin–orbit misalignments between isolated binaries and cluster-formed BBHs. Given the supernova prescriptions in `COSMIC` (here we have used the "rapid" prescription from Fryer et al. 2012), BHs with final core masses $\gtrsim$20 $M_\odot$ form via the "direct collapse" or failed supernova scenario, where the BHs form without any natal kick or change to the supernova orbital parameters. For BBHs with total masses $\lesssim$45 $M_\odot$, the supernova kicks can still produce significant spin–orbit misalignments, with the 99th percentile of allowed misalignments approaching complete anti-alignment ($\theta_{\rm LS} = 180°$). However, as the mass increases, the amount of spin–orbit misalignment decreases, until eventually all BBHs are aligned with their orbital angular momenta. Contrast this with BBHs formed in clusters, which are isotropically distributed on the sphere (i.e., $P(\theta_{\rm LS})d\theta_{\rm LS} \propto \sin(\theta_{\rm LS})d\theta_{\rm LS}$).

These results can be easily reproduced and expanded on using the apples-to-apples comparisons enabled by the `CMC` and `COSMIC` integration.

5. Discussion and Conclusion

This paper describes the first public release of `CMC`, a parallel code for collisional N-body stellar dynamics based on the original method of Hénon (1971a, 1971b). After nearly two decades of development, `CMC` contains all of the necessary physics for modeling the evolution of DSCs, such as GCs, SSCs, and certain NSCs, and it allows the user to easily create detailed models of spherical star clusters with up to $\sim 10^7$ particles (when including stellar evolution). In addition to two-body relaxation, `CMC` can model the formation and dynamics of

binaries, the effects of tidal stripping in a galactic field, detailed binary star evolution, and more.

The public release of CMC has been coupled to the COSMIC software package for binary population synthesis (Breivik et al. 2020a). The newest release of COSMIC, v3.4, contains new initial condition generators to create cluster profiles from various spherical stellar distributions (Plummer 1911; King 1966; Elson et al. 1987), which can then be coupled to all of the stellar and binary initial condition generators in the population synthesis code. Both the initial conditions and the output snapshots of CMC are saved in easily readable HDF5 format files and can be analyzed with standard data science packages (pandas) or with specialty software designed to compare CMC models directly to observations of real star clusters (the cmctoolkit; Rui et al. 2021a, 2021b).

The synergy between CMC and COSMIC enables stellar dynamics and population synthesis studies to be conducted simultaneously with identical assumptions and prescriptions for stellar initial conditions. We presented two examples of CMC in action. First, we studied the collapse of a Plummer sphere with 10^8 initial particles, the largest fully collisional star-by-star model integrated to core collapse, and showed that CMC was able to replicate the self-similar nature of the collapse across more than 15 orders of magnitude. Then, we presented examples of realistic clusters and showed how the combination of CMC with COSMIC makes comparative population studies—such as our study of GW190521, a BBH that was very likely formed in a dynamical environment—trivial. It is our hope that the public release of CMC and the integration of COSMIC will enable significant advancements in the study of DSCs and the many transient and high-energy events they produce (such as BBH mergers).

We thank Kuldeep Sharma, Xiaoqi Yu, and Mike Grudić for testing this release of CMC and providing feedback, and Elena González and Miguel Martinez for useful comments and discussions. This work was supported by NSF grant AST-2009916 at Carnegie Mellon University, a New Investigator Research Grant to C.R. from the Charles E. Kaufman Foundation, and NSF grant AST-1716762 at Northwestern University. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant No. ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award No. ACI-1928147, at the Pittsburgh Supercomputing Center (PSC). N.C.W. acknowledges support from the CIERA Riedel Family Graduate Fellowship. F.K. acknowledges support from the Turkish Fulbright Commission. K.K. is supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2001751. P.A.-S. acknowledges support from the Ramón y Cajal Programme of the Ministry of Economy, Industry and Competitiveness of Spain, as well as the financial support of Programa Estatal de Generación de Conocimiento (ref. PGC2018-096663-B-C43) (MCIU/FEDER). N.Z.R. acknowledges support from the Dominic Orr Graduate Fellowship at Caltech.

Software: The public release of CMC can be accessed, including source code and documentation, at https://clustermontecarlo.github.io/. CMC (Joshi et al. 2000, 2001; Watters et al. 2000; Fregeau et al. 2002, 2003; Chatterjee et al. 2010; Morscher et al. 2013; Pattabiraman et al. 2013; Rodriguez et al. 2021a; this work), cmctoolkit (Rui et al. 2021a, 2021b), fewbody (Fregeau & Rasio 2007; Antognini et al. 2014; Amaro-Seoane & Chen 2016), COSMIC (Breivik et al. 2020a, 2020b), matplotlib (Hunter 2007), SciPy (Virtanen et al. 2020), NumPy (Harris et al. 2020), pandas (McKinney 2010; Reback et al. 2021).

ORCID iDs

Carl L. Rodriguez https://orcid.org/0000-0003-4175-8881
Newlin C. Weatherford https://orcid.org/0000-0002-9660-9085
Katelyn Breivik https://orcid.org/0000-0001-5228-6598
Sourav Chatterjee https://orcid.org/0000-0002-3680-2684
Giacomo Fragione https://orcid.org/0000-0002-7330-027X
Kyle Kremer https://orcid.org/0000-0002-4086-3180
Nicholas Z. Rui https://orcid.org/0000-0002-1884-3992
Claire S. Ye https://orcid.org/0000-0001-9582-881X
Michael Zevin https://orcid.org/0000-0002-0147-0835
Frederic A. Rasio https://orcid.org/0000-0002-7132-418X

References

Aarseth, S. J. 2003, Gravitational N-Body Simulations (Cambridge: Cambridge Univ. Press)
Aarseth, S. J. 2012, MNRAS, 422, 841
Aarseth, S. J., Henon, M., & Wielen, R. 1974, A&A, 37, 183
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, PhRvX, 6, 041015
Abt, H. A. 1983, ARA&A, 21, 343
Ajith, P., Hannam, M., Husa, S., et al. 2011, PhRvL, 106, 241101
Amaro-Seoane, P., & Chen, X. 2016, MNRAS, 458, 3075
Ambartsumian, V. A. 1937, AZh, 14, 207
Antognini, J. M., Shappee, B. J., Thompson, T. A., & Amaro-Seoane, P. 2014, MNRAS, 439, 1079
Antonini, F., & Gieles, M. 2020, MNRAS, 492, 2936
Antonov, V. A. 1985, in IAU Symp. 113, Dynamics of Star Clusters, ed. J. Goodman & P. Hut (Dordrecht: D. Reidel), 525
Arca Sedda, M., Askar, A., & Giersz, M. 2018, MNRAS, 479, 4652
Askar, A., Szkudlarek, M., Gondek-Rosińska, D., Giersz, M., & Bulik, T. 2016, MNRAS, 464, L36
Bae, Y.-B., Kim, C., & Lee, H. M. 2014, MNRAS, 440, 2714
Belczynski, K. 2020, ApJL, 905, L15
Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, ApJS, 174, 223
Binney, J., & Tremaine, S. 2008, Galactic Dynamics (2nd ed. ; Princeton, NJ: Princeton Univ. Press)
Breen, P. G., & Heggie, D. C. 2013, MNRAS, 432, 2779
Breivik, K., Coughlin, S., Zevin, M., et al. 2020a, ApJ, 898, 71
Breivik, K., Coughlin, S., Zevin, M., et al. 2020b, COSMIC-PopSynth/COSMIC: Code Release with COSMIC Paper: v3.3.0, Zenodo, doi:10.5281/zenodo.3905335
Campanelli, M., Lousto, C., Zlochower, Y., & Merritt, D. 2007, ApJL, 659, L5
Capuzzo-Dolcetta, R., Spera, M., & Punzo, D. 2013, JCoPh, 236, 580
Casertano, S., & Hut, P. 1985, ApJ, 298, 80
Chatterjee, S., Fregeau, J. M., Umbreit, S., & Rasio, F. A. 2010, ApJ, 719, 915
Chatterjee, S., Rodriguez, C. L., Kalogera, V., & Rasio, F. A. 2017a, ApJL, 836, L26
Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2017b, ApJ, 834, 68
Claeys, J. S. W., Pols, O. R., Izzard, R. G., Vink, J., & Verbunt, F. W. M. 2014, A&A, 563, A83
Clark, G. W. 1975, ApJL, 199, L143
Clarke, L., Glendinning, I., & Hempel, R. 1994, in Programming Environments for Massively Parallel Distributed Systems, ed. K. M. Decker & R. M. Rehmann (Basel: Birkhäuser), 213
Cohn, H. 1979, ApJ, 234, 1036
Cohn, H. 1980, ApJ, 242, 765
Di Carlo, U. N., Giacobbo, N., Mapelli, M., et al. 2019, MNRAS, 487, 2947
Downing, J. M. B., Benacquista, M. J., Giersz, M., & Spurzem, R. 2010, MNRAS, 407, 1946
Downing, J. M. B., Benacquista, M. J., Giersz, M., & Spurzem, R. 2011, MNRAS, 416, 133
Elson, R. A. W., Fall, S. M., & Freeman, K. C. 1987, ApJ, 323, 54
Fabian, A. C., Pringle, J. E., & Rees, M. J. 1975, MNRAS, 172, 15

Farmer, R., Renzo, M., de Mink, S. E., Marchant, P., & Justham, S. 2019, ApJ, 887, 53
Farr, B., Holz, D. E., & Farr, W. M. 2018, ApJL, 854, L9
Farr, W. M., Stevenson, S., Miller, M. C., et al. 2017, Natur, 548, 426
Fragione, G., & Kocsis, B. 2018, PhRvL, 121, 161103
Fragione, G., Loeb, A., & Rasio, F. A. 2020a, ApJL, 902, L26
Fragione, G., Martinez, M. A. S., Kremer, K., et al. 2020b, ApJ, 900, 16
Fragione, G., & Silk, J. 2020, MNRAS, 498, 4591
Frazer, W. D., & McKellar, A. C. 1970, J. ACM, 17, 496
Fregeau, J. M., Gürkan, M. A., Joshi, K. J., & Rasio, F. A. 2003, ApJ, 593, 772
Fregeau, J. M., Joshi, K. J., Portegies Zwart, S. F., & Rasio, F. A. 2002, ApJ, 570, 171
Fregeau, J. M., & Rasio, F. A. 2007, ApJ, 658, 1047
Freitag, M. 2008, in The Cambridge N-Body Lectures, ed. S. J. Aarseth, C. A. Tout, & R. A. Mardling (Dordrecht: Springer)
Freitag, M., & Benz, W. 2001, A&A, 375, 711
Freitag, M., & Benz, W. 2002, A&A, 394, 345
Freitag, M., Rasio, F. A., & Baumgardt, H. 2006, MNRAS, 368, 121
Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, ApJ, 749, 91
Fuller, J., & Ma, L. 2019, ApJL, 881, L1
Fuller, J., Piro, A. L., & Jermyn, A. S. 2019, MNRAS, 485, 3661
Gerosa, D., & Kesden, M. 2016, PhRvD, 93, 124066
Giersz, M. 1998, MNRAS, 298, 1239
Giersz, M., & Heggie, D. C. 1994, MNRAS, 268, 257
Giersz, M., & Heggie, D. C. 1996, MNRAS, 279, 1037
Giersz, M., Heggie, D. C., & Hurley, J. R. 2008, MNRAS, 388, 429
Giersz, M., Heggie, D. C., Hurley, J. R., & Hypki, A. 2013, MNRAS, 431, 2184
González, E., Kremer, K., Chatterjee, S., et al. 2021, ApJL, 908, L29
Grindlay, J. E., Cool, A. M., Callanan, P. J., et al. 1995, ApJL, 455, L47
Grudić, M. Y., Guszejnov, D., Hopkins, P. F., et al. 2018, MNRAS, 481, 688
Hamers, A. S., & Samsing, J. 2019, MNRAS, 487, 5630
Harfst, S., Gualandris, A., Merritt, D., & Mikkola, S. 2008, MNRAS, 389, 2
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Heggie, D., & Hut, P. (ed.) 2003, The Gravitational Million-Body Problem: A Multidisciplinary Approach to Star Cluster Dynamics (Cambridge: Cambridge Univ. Press), 372
Heggie, D. C. 1975, MNRAS, 173, 729
Heggie, D. C. 2014, MNRAS, 445, 3435
Heggie, D. C., & Mathieu, R. D. 1986, The Use of Supercomputers in Stellar Dynamics (Berlin: Springer), 233
Heggie, D. C., & Stevenson, D. 1988, MNRAS, 230, 223
Hénon, M. 1961, AnAp, 24, 369
Hénon, M. 1971a, Ap&SS, 13, 284
Henon, M. 1973, in Saas-Fee Advanced Course 3: Dynamical Structure and Evolution of Stellar Systems, ed. G. Contopoulos, M. Henon, & D. Lynden (Geneva: Observatoire de Genève), 183
Hénon, M. H. 1971b, Ap&SS, 14, 151
Hoare, C. A. R. 1961, Commun. ACM, 4, 321
Holgado, A. M., Ortega, A., & Rodriguez, C. L. 2021, ApJL, 909, L24
Hopman, C. 2009, ApJ, 700, 1933
Hunter, J. D. 2007, CSE, 9, 90
Hurley, J. R., Pols, O. R., & Tout, C. A. 2000, MNRAS, 315, 543
Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, MNRAS, 329, 897
Hut, P., & Inagaki, S. 1985, ApJ, 298, 502
Ivanova, N., Belczynski, K., Fregeau, J. M., & Rasio, F. A. 2005, MNRAS, 358, 572
Ivanova, N., Chaichenets, S., Fregeau, J., et al. 2010, ApJ, 717, 948
Ivanova, N., Heinke, C. O., Rasio, F. A., et al. 2006, MNRAS, 372, 1043
Joshi, K. J., Nave, C. P., & Rasio, F. A. 2001, ApJ, 550, 691
Joshi, K. J., Rasio, F. A., & Portegies Zwart, S. 2000, ApJ, 540, 969
Kalogera, V. 2000, ApJ, 541, 319
Kim, S. S., & Lee, H. M. 1999, A&A, 347, 123
King, I. R. 1966, AJ, 71, 64
Kremer, K., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2018a, ApJ, 852, 29
Kremer, K., Chatterjee, S., Ye, C. S., Rodriguez, C. L., & Rasio, F. A. 2019a, ApJ, 871, 38
Kremer, K., Lu, W., Rodriguez, C. L., Lachat, M., & Rasio, F. A. 2019b, ApJ, 881, 75
Kremer, K., Rodriguez, C. L., Amaro-Seoane, P., et al. 2019c, PhRvD, 99, 063003
Kremer, K., Rui, N. Z., Weatherford, N. C., et al. 2021, ApJ, 917, 28
Kremer, K., Spera, M., Becker, D., et al. 2020a, ApJ, 903, 45
Kremer, K., Ye, C. S., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2018b, ApJL, 855, L15
Kremer, K., Ye, C. S., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2020b, in IAU Symp. 351, Star Clusters: From the Milky Way to the Early, ed. A. Bragaglia et al. (Cambridge: Cambridge Univ. Press), 357
Kremer, K., Ye, C. S., Rui, N. Z., et al. 2020c, ApJS, 247, 48
Kroupa, P. 2001, MNRAS, 322, 231
Kulkarni, S. R., Hut, P., & McMillan, S. J. 1993, Natur, 364, 421
Küpper, A. H. W., Maschberger, T., Kroupa, P., & Baumgardt, H. 2011, MNRAS, 417, 2300
The LIGO Scientific CollaborationThe Virgo Collaboration, Abbott, R., et al. 2021, ApJL, 913, L7
Lombardi, J. C., Thrall, A. P., Deneva, J. S., Fleming, S. W., & Grabowski, P. E. 2003, MNRAS, 345, 762
Lombardi, J. C., Warren, J. S., Rasio, F. A., Sills, A., & Warren, A. R. 2002, ApJ, 568, 939
Lousto, C. O., & Zlochower, Y. 2008, PhRvD, 77, 044028
Lynden-Bell, D. 1967, MNRAS, 136, 101
Lynden-Bell, D. 1975, in IAU Symp. 69, Dynamics of the Solar Systems, ed. A. Hayli (Dordrecht: D. Reidel), 27
Lynden-Bell, D., & Eggleton, P. P. 1980, MNRAS, 191, 483
Lynden-Bell, D., & Wood, R. 1968, MNRAS, 138, 495
Mackey, A. D., Wilkinson, M. I., Davies, M. B., & Gilmore, G. F. 2007, MNRAS, 379, L40
Mackey, A. D., Wilkinson, M. I., Davies, M. B., & Gilmore, G. F. 2008, MNRAS, 386, 65
Madigan, A.-M., Hopman, C., & Levin, Y. 2011, ApJ, 738, 99
Marchant, A. B., & Shapiro, S. L. 1980, ApJ, 239, 685
Martinez, M. A. S., Fragione, G., Kremer, K., et al. 2020, ApJ, 903, 67
Martinez, M. A. S., Rodriguez, C. L., & Fragione, G. 2021, arXiv:2105.01671
McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 56
McMillan, S., Portegies Zwart, S., van Elteren, A., & Whitehead, A. 2012, in ASP Conf. Ser. 453, Advances in Computational Astrophysics: Methods, Tools, and Outcome, ed. R. Capuzzo-Dolcetta, M. Limongi, & A. Tornambè (San Francisco, CA: ASP), 129
McMillan, S. L. W., & Hut, P. 1996, ApJ, 467, 348
Meiron, Y., & Kocsis, B. 2019, ApJ, 878, 138
Merritt, D., Milosavljevi, M., Favata, M., Hughes, S. A., & Holz, D. E. 2004, ApJL, 607, L9
Mikkola, S. 1983, MNRAS, 203, 1107
Milone, A. P., Piotto, G., Bedin, L. R., et al. 2012, A&A, 540, A16
Moe, M., & Di Stefano, R. 2017, ApJS, 230, 15
Morscher, M., Umbreit, S., Farr, W. M., & Rasio, F. A. 2013, ApJL, 763, L15
O'Leary, R., O'Shaughnessy, R., & Rasio, F. 2007, PhRvD, 76, 061504
O'Leary, R. M., Rasio, F. A., Fregeau, J. M., Ivanova, N., & O'Shaughnessy, R. 2006, ApJ, 637, 937
Ostriker, J. P., Spitzer, L. J., & Chevalier, R. A. 1972, ApJL, 176, L51
Pattabiraman, B., Umbreit, S., Liao, W.-K., et al. 2013, ApJS, 204, 15
Peters, P. 1964, PhRv, 136, B1224
Plummer, H. C. 1911, MNRAS, 71, 460
Pols, O. R., Schröder, K.-P., Hurley, J. R., Tout, C. A., & Eggleton, P. P. 1998, MNRAS, 298, 525
Pooley, D., Lewin, W. H. G., Anderson, S. F., et al. 2003, ApJL, 591, L131
Portegies Zwart, S. F., & Mcmillan, S. L. W. 2000, ApJ, 528, 17
Portegies Zwart, S. F., & Meinen, A. T. 1993, A&A, 280, 174
Quinlan, G. D. 1996, NewA, 1, 255
Quinlan, G. D., & Shapiro, S. L. 1987, ApJ, 321, 199
Rantala, A., Naab, T., & Springel, V. 2021, MNRAS, 502, 5546
Rappaport, S., & Putney, A. 1989, ApJ, 345, 210
Rasio, F. A., & Heggie, D. C. 1995, ApJL, 445, L133
Rasio, F. A., McMillan, S., & Hut, P. 1995, ApJL, 438, L33
Rauch, K. P., & Tremaine, S. 1996, NewA, 1, 149
Reback, J., McKinney, W., Jbrockmendel, et al. 2021, pandas-dev/pandas: Pandas 1.2.4, v1.2.4, Zenodo, doi:10.5281/zenodo.4681666
Renaud, F., Gieles, M., & Boily, C. M. 2011, MNRAS, 418, 759
Renzo, M., Zapartas, E., de Mink, S. E., et al. 2019, A&A, 624, A66
Rieder, S., & Pelupessy, I. 2019, rieder/Fresco: Fresco 0.6, v0.6.0, Zenodo, doi:10.5281/zenodo.3553805
Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., et al. 2018a, PhRvD, 98, 123005
Rodriguez, C. L., Amaro-Seoane, P., Chatterjee, S., & Rasio, F. A. 2018b, PhRvL, 120, 151101
Rodriguez, C. L., Chatterjee, S., & Rasio, F. A. 2016a, PhRvD, 93, 084029
Rodriguez, C. L., Coughlin, S., Weatherford, N., et al. 2021a, ClusterMonteCarlo/CMC-COSMIC, v0.9.0, Zenodo, doi:10.5281/zenodo.4763441

Rodriguez, C. L., Haster, C.-J., Chatterjee, S., Kalogera, V., & Rasio, F. A. 2016b, ApJL, 824, L8
Rodriguez, C. L., Kremer, K., Chatterjee, S., et al. 2021b, RNAAS, 5, 19
Rodriguez, C. L., Kremer, K., Grudić, M. Y., et al. 2020, ApJL, 896, L10
Rodriguez, C. L., Morscher, M., Pattabiraman, B., et al. 2015, PhRvL, 115, 051101
Rodriguez, C. L., Morscher, M., Wang, L., et al. 2016c, MNRAS, 463, 2109
Rodriguez, C. L., Pattabiraman, B., Chatterjee, S., et al. 2018c, ComAC, 5, 5
Rodriguez, C. L., Zevin, M., Amaro-Seoane, P., et al. 2019, PhRvD, 100, 043027
Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016d, ApJL, 832, L2
Rui, N. Z., Kremer, K., Weatherford, N. C., et al. 2021a, NicholasRui/cmctoolkit: First Release, 1.0, Zenodo, doi:10.5281/zenodo.4579950
Rui, N. Z., Kremer, K., Weatherford, N. C., et al. 2021b, ApJ, 912, 102
Salpeter, E. E. 1955, ApJ, 121, 161
Sana, H., de Mink, S. E., de Koter, A., et al. 2012, Sci, 337, 444
Sigurdsson, S., & Hernquist, L. 1993, Natur, 364, 423
Sigurdsson, S., & Phinney, E. S. 1993, ApJ, 415, 631
Sollima, A., & Mastrobuono Battisti, A. 2014, MNRAS, 443, 3513
Spitzer, L. 1987, Dynamical Evolution of Globular Clusters (Princeton, NJ: Princeton Univ. Press)
Spitzer, L. J., & Hart, M. H. 1971, ApJ, 164, 399
Spitzer, L. J., & Thuan, T. X. 1972, ApJ, 175, 31
Spurzem, R., & Aarseth, S. J. 1996, MNRAS, 282, 19
Stodółkiewicz, J. S. 1982, AcA, 32, 63
Tanikawa, A. 2013, MNRAS, 435, 1358
Tanikawa, A., Heggie, D. C., Hut, P., & Makino, J. 2013, A&C, 3, 35
van Son, L. A. C., De Mink, S. E., Broekgaarden, F. S., et al. 2020, ApJ, 897, 100
Vasiliev, E. 2015, MNRAS, 446, 3150
Verbunt, F., & Hut, P. 1987, in IAU Symp. 125, The Origin and Evolution of Neutron Stars, ed. D. J. Helfand & J. H. Huang (Dordrecht: D. Reidel), 187
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
von Hoerner, S. 1960, ZAp, 50, 184
Wang, L., Spurzem, R., Aarseth, S., et al. 2015, MNRAS, 450, 4070
Wang, L., Spurzem, R., Aarseth, S., et al. 2016, MNRAS, 458, 1450
Watters, W. A., Joshi, K. J., & Rasio, F. A. 2000, ApJ, 539, 331
Weatherford, N. C., Fragione, G., Kremer, K., et al. 2021, ApJL, 907, L25
Wong, K. W. K., Breivik, K., Kremer, K., & Callister, T. 2021, PhRvD, 103, 083021
Woosley, S. E. 2017, ApJ, 836, 244
Ye, C. S., Kremer, K., Chatterjee, S., Rodriguez, C. L., & Rasio, F. A. 2019, ApJ, 877, 122
Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, ApJ, 910, 152
Zevin, M., Samsing, J., Rodriguez, C., Haster, C.-J., & Ramirez-Ruiz, E. 2019, ApJ, 871, 91