

Project Title: Towards Exascale Astrophysics of Mergers and Supernovae (TEAMS)

Submitted by:

Princeton University, Princeton NJ 08544

To DOE Office of Science, Office of Nuclear Physics, in response to Scientific Discovery through Advanced Computing: Nuclear Physics

DOE National Laboratory Announcement Number: LAB 17-1698 (Program Officer: Ted Barnes)

PAMS Letter of Intent Number: LOI-0000016477

Research Area: Low Energy Nuclear Physics, Nuclear Structure and Astrophysics

Administrative Contact:

Jeffrey Friedland, Off. of Research & Proj. Admin., P.O. Box 36, Princeton, NJ 08544-2020
jfried@princeton.edu, (609) 258-3090

Principal Investigator:

Adam Burrows, Professor, Department of Astrophysical Sciences, Princeton University
burrows@astro.princeton.edu, (609) 258-3590

Senior/Key Personnel:

David Radice, Princeton University

Contents

1	Siting the R-Process	1
1.1	<i>Computational Tools for Nuclear Astrophysics</i>	2
2	Nuclear Astrophysics Research Plan	6
2.1	<i>Modeling the Lives of Massive Stars</i>	6
2.2	<i>Modeling the Deaths of Massive Stars</i>	7
2.2.1	<i>Iron Core Collapse</i>	9
2.2.2	<i>O-Ne Core Collapse</i>	11
2.2.3	<i>Proto-Neutron Star Winds</i>	12
2.2.4	<i>Magnetars and Magnetorotational Turbulence</i>	13
2.3	<i>Epstein, Colgate & Haxton mechanism</i>	14
2.4	<i>Modeling the Deaths of Neutron Stars</i>	14
2.4.1	<i>Nuclear Decompression</i>	15
2.4.2	<i>Disk Winds</i>	16
2.5	<i>Modeling Dense Matter and its interactions with Neutrinos</i>	17
2.6	<i>Modeling the Products of Massive Stars</i>	19
2.7	<i>Observing the Deaths of Massive Stars</i>	19
3	Computational Science Research Plan	21
3.1	<i>Code Development for Improved Stellar Modeling</i>	21
3.2	<i>Code Development for Improved Models of Core-Collapse Supernovae and Neutron Star Mergers</i>	23
3.2.1	<i>CHIMERA</i>	23
3.2.2	<i>CLASH</i>	23
3.2.3	<i>FORNAX</i>	23
3.2.4	<i>WHISKYTHC</i>	24
3.3	<i>Code Development for Improving the Nuclear EoS and Opacities</i>	25
3.4	<i>Code Development and Uncertainty Quantification for Improving Nucleosynthesis Predictions</i>	26
4	Building a Stronger Nuclear Astrophysics Community	27
4.1	<i>Providing EoS and opacity data</i>	27
4.2	<i>Community Simulation Codes</i>	28
4.3	<i>Community Code Comparisons and Testing</i>	28
4.4	<i>Providing Data for the Broader Community</i>	29
5	Milestones and Deliverables Timeline	29
6	Management Plan	30
	Appendix 1: Biographical Sketches	31
	Appendix 2: Current And Pending Support	88
	Appendix 3: Bibliography & References Cited	122
	Appendix 4: Facilities & Other Resources	136
	Appendix 5: Equipment	136
	Appendix 6: Data Management Plan	138

1 Siting the R-Process

We are the result of nearly 14 billion years of cosmic evolution, stretching back to the Big Bang. One intriguing facet of this history is the formation of the chemical elements of which we, and the planet upon which we live, are made. Beginning mere minutes after the Big Bang with the fusion of newly formed neutrons and protons into helium and other light elements, this chemical evolution links us to the lives and deaths of stars, from when the Galaxy first formed to the present day. In the ninety six years since Eddington suggested that the stars could be powered by the transmutation of the elements, our understanding of this chemical evolution has matured greatly. Yet questions still remain, chief among them: What is the source of the heaviest elements? As a result, the Committee on the Physics of the Universe convened by the National Research Council in 2000 included the question “How Were the Elements from Iron to Uranium Made?” in their list of Eleven Science Questions for the New Century [275]. On the basis of questions like this, the 2007 and 2015 Nuclear Science Advisory Committee [91, 271] recommended the ongoing support of current radioactive ion beam (RIB) facilities and the construction of the Facility for Rare Isotope Beams (FRIB), to provide needed nuclear data, our ignorance of which is one significant impediment to our understanding of the formation of these heavy elements.

However, understanding the formation of the heavy elements requires not only the nuclear data provided by such RIB facilities, but knowledge of the nuclear processes by which—and the astrophysical settings in which—these elements are forged. Burbidge et al. [43] and Cameron [54] divided the observed abundance of the isotopes heavier than iron among three distinct processes. On the basis of the correlation between nuclear properties and cosmic abundances, we know that roughly half of these isotopes are the result of a slow neutron capture process, termed the *s-process*. Most of the remainder are the result of a rapid neutron capture process, the *r-process*, while a scattering of rarer proton-rich isotopes are ascribed to a *p-process*, originally thought to be driven by proton capture. Based on well-measured nuclear data from stable nuclei and mature models of stellar evolution, we have considerable confidence that the *s-process* occurs over thousands of years in the hydrogen- and helium-rich nuclear burning shells that sit atop the inert carbon–oxygen (CO) cores in Asymptotic Giant Branch (AGB) stars. This confidence is buoyed by observations of red-giant stars with large over-abundances of *s-process* elements like barium and radioactive technetium (see [131] for a recent review).

In the cases of the *r-process* and *p-process*, we do not have well-measured nuclear data for the highly radioactive isotopes involved, nor mature models of the potential astrophysical sites, nor direct observations. While current and future RIB facilities endeavor to remove our ignorance of the relevant nuclear data, **a primary scientific goal of this proposal is to mature the modeling of many of the suggested astrophysical sites of the *r-process*, as well as some potential *p-process* sites.** This effort is not independent of the experimental efforts, as nucleosynthesis calculations rely on nuclear data such as nuclear masses, reaction cross sections, etc. However, as Figure 1 illustrates, the *r-process* paths, and hence the nuclei for which such nuclear data are needed, vary between the different potential sites. Thus, winnowing the list of possible *r-process* or *p-process* sites will reduce the nuclear data needs. **A second goal of this proposal is to use our models to support and guide the experimental efforts at FRIB and similar facilities.**




Figure 1: Comparison of the *r-process* path from three *r-process* simulations: neutrino-driven wind scenarios with greater (green) and lesser (red) neutron overabundances, and a scenario that produces the decompression of neutron star matter (blue). The paths are compared to the line of stability (black squares) and with currently measured half-lives, shown in shades of purple.

fig ref?

“an effort to mature the modeling...”. Can we say something more exciting? Abstract too. ZZZZZZZzzzzzzzz

Recovery of the observed pattern of r-process species requires a nuclear reaction flow that proceeds far from stability, acting on a time scale of seconds. This implies a connection to an explosive event, hints of which can be provided by spectroscopic observations. The appearance of r-process elements in the spectra of old low-mass stars whose heavy element abundance are less than 1/1000th of the Sun (termed *metal-poor stars*) has long been interpreted to indicate that the formation of r-process elements began very early in *Galactic chemical evolution* (GCE), suggesting that the r-process occurs during the death throes of short-lived, massive stars [see, e.g. 13]. The recent discovery of high r-process enhancements in the metal-poor, ultra-faint dwarf galaxy Reticulum II [125, 250] suggests an alternate interpretation. The amount of observed r-process material is consistent only with pollution by a rare event with prodigious r-process yields, such as expected for a merger of two neutron stars or a neutron star and black hole [147]. Galactic low-metallicity r-process enhancements can perhaps be explained by the accretion of such systems [103], making metallicity less an indicator of age than of stellar birth in a pristine environment. Still, some level of neutron-capture elements are found in almost all galactic low-metallicity stars for which observations are available [249]. The combined spectroscopic evidence, along with meteoritic data [283], may point to contributions from multiple r-process sources.

The essential requirement for a successful r-process is approximately 100 free neutrons to be available to capture on each seed heavy nucleus (iron or heavier). Under sufficiently neutronized conditions, this requirement can be met directly, with neutrons and seeds dominating the composition. Alternatively, under less neutronized conditions, the critical ratio of neutrons to seeds can be achieved if the composition is dominated by isotopes largely immune to neutron capture, like ^4He . In stars like our Sun, the star’s ultimate fate as a white dwarf precludes reaching densities sufficient for electron capture to create the needed neutron-richness, also pointing to more massive stars as the r-process source. Numbered prominently among the massive star scenarios that may match the required conditions include the collapse of oxygen–neon–magnesium (ONeMg) stellar cores and iron stellar cores, with their associated supernovae; explosive nuclear decompression during the mergers of old neutron stars, and *accretion disks* around newly formed black holes in “failed” supernovae or merged neutron stars. Simulations of the central engines in each of these scenarios rely on a similar toolkit: magnetohydrodynamics (MHD), thermonuclear kinetics, the *equation of state* (EoS) of nuclear matter, and radiation transport of neutrinos. Models with greater physical fidelity are necessary to establish the true site of the r-process, constraining the nuclear data needed to understand it. Earlier phases of these same studies will provide the most thorough investigation to date of potential sources of p-process nuclei, also constraining a portion of the nuclear data needed to understand the p-process. These models will also enhance our understanding of the production of elements from oxygen through iron in these events, improving our ability to compare to observations of these species in supernovae and their remnants. **Our third scientific goal is to use our tools to cement our understanding of the range of roles that massive stars play in Galactic chemical evolution.** To achieve these goals, we will detail in the following sections a plan to improve the modeling of these events (§2) that will require investments in improved computational tools (§3) and the development of a more integrated community (§4).

1.1 Computational Tools for Nuclear Astrophysics

Greater physical fidelity comes at the expense of greater computational cost. The simulations described above will, therefore, require software capable of effectively exploiting current and future architectures. **The first computational goal of this proposal is to build a toolkit that can be used on petascale (and eventually exascale) computational architectures to provide world-leading models for all r-process scenarios.** The members of the collaboration bring a variety of computational tools to bear on the problems described above, tools that we plan to improve under this proposal. These tools include:

Maestro is an open source low Mach number stellar hydrodynamics code¹. Maestro is tuned to efficiently

1. <https://github.com/BoxLib-Codes/MAESTRO>

model the highly subsonic convective flows that often precede stellar explosions. Traditional astrophysical hydrodynamics codes (e.g. Flash [85] and Castro [7]) solve the equations of fully compressible hydrodynamics using explicit timestepping. Here, information travels at the speeds U and $U \pm c_s$, where U is the fluid velocity and c_s is the sound speed. For a zone of width Δx , the maximum size of the timestep is $\Delta t_{\text{sound}} < \Delta x/(|U| + c_s)$. While these codes are widely available, and the algorithms are easy to implement, they are ill-suited for highly subsonic flows. As the Mach number, $M \equiv U/c_s$ approaches zero, $\Delta t_{\text{sound}} \sim \Delta x/c_s$. For $M \ll 1$, soundwaves are not important to the overall dynamics of the flow, but the stability requirement means that a compressible code needs to track these soundwaves anyway, resulting in an overly restrictive timestep.

Maestro gains computational efficiency by filtering soundwaves from the hydrodynamics equations while retaining the compressibility effects due to stratification of the star, localized nuclear energy release, compositional mixing, and thermal diffusion [8–10]. This allows us to take a timestep dictated by the fluid motion itself, $\Delta t_{\text{fluid}} \sim \Delta x/|U| \sim (1/M)\Delta t_{\text{sound}}$ for $M \ll 1$ (in reality it can be quite a bit more efficient than this, since the highest soundspeed in a star is usually at the center, where the Mach number is low).

The advection portion of Maestro uses an unsplit piecewise parabolic method algorithm to advance the density, velocity, species, and enthalpy in time. An elliptic constraint on the velocity field in the reformulated equations acts to enforce the low Mach approximation, driving instantaneous acoustic equilibration. This constraint is enforced using geometric multigrid at several points in the algorithm. Reactions are currently incorporated via Strang splitting, yielding an overall second-order-in-time algorithm. Maestro has been applied to convection in the Chandrasekhar-mass model for Type Ia supernovae [208, 301, 304], convection in the helium layer of sub-Chandra SNe Ia models [116, 303], X-ray bursts [164, 165, 302], and core convection in massive stars [92]. In this proposal, Maestro will be used for generating realistic three-dimensional initial conditions for CCSN models.

Maestro uses a hybrid programming model based on MPI and OpenMP. Distribution of grid patches to nodes using MPI provides a natural coarse-grained approach to distributing the computational work, while threading of loops over zones in a grid using OpenMP provides effective fine-grained parallelization. The parallelization and grid is managed via the BoxLib library.

Castro is the compressible sister code to Maestro. It solves the equations of radiation hydrodynamics (in a multigroup flux-limited diffusion approximation with implicit time differencing) and employs full self-gravity with boundary conditions appropriate for mergers [7, 133], and implemented in a conservative fashion. It uses the same underlying adaptive mesh refinement library (BoxLib, with AMReX support just coming online) and microphysics routines. Like Maestro, it is fully open source with the development centered at github. The common base for these codes allows us to restart a Maestro calculation in Castro to continue the evolution from the subsonic to explosive phase, a technique previously used with Type Ia supernovae [165].

CHIMERA is a parallel, multi-physics code built specifically for multidimensional simulation of CCSNe. It combines separate codes for hydrodynamics and gravity; neutrino transport and opacities; and nuclear EoS and reaction network, coupled by a layer that oversees data management, parallelism, I/O, and control. The hydrodynamics are evolved via a dimensionally-split, Lagrangian-Remap (PPMLR) scheme [58] as implemented in VH1 [97]. Self-gravity is computed by multipole expansion [199], with the Newtonian monopole replaced with a GR monopole [166, Case A]. Neutrino transport is computed in the “ray-by-ray-plus” (RbR+) approximation [41] using an improved and updated version of the multi-group flux-limited diffusion transport solver of Bruenn [34], enhanced for GR [35], with an additional geometric flux limiter to prevent the over-rapid transition to free streaming of the standard flux-limiter. All $O(v/c)$ observer corrections have been included. Neutrinos are advected laterally (in the θ -direction) with the fluid and contribute to the lateral pressure gradient where $\rho > 10^{12} \text{g cm}^{-3}$. We solve the transport equation for all three flavors of (anti-)neutrinos with four coupled species: $\nu_e, \bar{\nu}_e, \nu_{\mu\tau} = \{\nu_\mu, \nu_\tau\}, \bar{\nu}_{\mu\tau} = \{\bar{\nu}_\mu, \bar{\nu}_\tau\}$.

FLASH [72, 86] is a massively-parallel multiphysics AMR simulation framework that has been applied to a vast range of physical problems, spanning scales from high-energy density laboratory physics to cosmology. Since its creation, FLASH has been DOE-supported and professionally maintained in order to deliver a high performance, reliable platform to tackle challenging scientific problems at extreme computational scales. FLASH has directly contributed to nearly 1000 publications and has become a premier community code in numerous separate scientific domains. FLASH contains a multitude of different solvers and physics packages. The FLASH CCSN application [e.g., 60, 64, 212] consists of a directionally-unsplit finite volume MHD solver [151] that utilizes high-order reconstruction schemes, such as PPM [59] and WENO [32, 256]. The divergence-free condition of the magnetic fields is maintained via a staggered-mesh constrained transport approach [90] with high-order extensions for calculation of electric fields [151, 152]. Self-gravity is approximated using a multipole expansion of the Poisson equation [200] with modifications to accurately calculate potentials on non-spherical polar grids [63]. This otherwise Newtonian gravity solver has been extended to include the common general relativistic effective monopole potential [167, 212].

Multi-flavor neutrino transport in FLASH is achieved via an efficient explicit two-moment scheme with an analytic closure for higher-order moments [e.g., 255], the so-called M1 scheme. FLASH’s M1 implementation closely follows that of the open-source code GRID [211], and directly solves for the zeroth (energy density) and first (flux density) moments of the neutrino radiation fields. This forms a hyperbolic system which is solved explicitly with a conservative finite volume approach while the completely local neutrino-matter coupling source terms are solved implicitly [see 212]. FLASH-M1 includes general relativistic transport corrections and velocity dependence to $\mathcal{O}(v/c)$. Inclusion of inelastic neutrino-electron and neutrino-nucleon scattering is currently under development. The M1 approach avoids expensive global and iterative solves common to many other transport algorithms while being fundamentally multidimensional (i.e., avoiding the ray-by-ray approximation). FLASH-M1 has been shown to weak scale almost perfectly to over 100,000 cores on Argonne Leadership Computing Facility’s BG/Q *Mira* system. A paper detailing the results of the first fully 3D FLASH-M1 CCSN simulations is now in preparation (O’Connor & Couch, 2017).

FORNAX (Burrows et al. 2016; Skinner et al. 2017; Wallace, Burrows, and Dolence 2016; Dolence et al. 2017) is a multi-dimensional, multi-group radiation/hydrodynamic code employing a directionally-unsplit Godunov-type finite-volume TVD-limited reconstruction method, written in a covariant/coordinate-independent fashion, with generalized connection coefficients and static mesh refinement. It solves the comoving-frame, multi-group, two-moment, velocity-dependent transport equations with an explicit Godunov characteristic method applied to the radiation transport operators and an implicit solver for the radiation source terms, uses the M1 tensor closure for the second and third moments of the radiation fields (Vaytet et al. 2011), and employs approximate general-relativistic gravity (Marek et al. 2006).

In FORNAX, by addressing the transport operator with an explicit method, we significantly reduce the computational complexity and communication overhead of traditional multi-dimensional radiative transfer solutions by bypassing the need for global iterative solvers that have proven to be slow and/or problematic beyond $\sim 10,000$ cores. Radiation quantities are reconstructed with linear profiles, and the calculated edge states are used to determine fluxes via an HLLC solver. In the non-hyperbolic regime, the HLLC fluxes are corrected to reduce numerical diffusion (O’Connor & Ott 2013).

A comprehensive set of neutrino-matter interactions are followed in FORNAX, and these are described in Burrows, Reddy, & Thompson (2006). They include weak magnetism and recoil corrections to neutrino-nucleon scattering and absorption (Horowitz 2002); ion-ion-correlations, weak screening, and form-factor corrections for neutrino-nucleus scattering; and inelastic neutrino-electron scattering using the scheme of Thompson, Burrows, & Pinto (2003) and the relativistic formalism summarized in Reddy et al. (1999). Inelastic neutrino-nucleon scattering is also included using a modified version of the Thompson, Burrows, & Pinto (2003) approach. Neutrino sources and sinks due to nucleon-nucleon bremsstrahlung and electron-positron annihilation are included, as described in Thompson, Burrows, & Horvath (2000).

While most other non-Cartesian supernova codes have simply evolved the inner region in 1-D spherical symmetry to avoid severe Courant limitations, in FORNAX, we instead enhance our otherwise logically Cartesian mesh with a simple form of static mesh derefinement. Regardless of the resolution specified in the angular directions, this “dendritic” grid has only two zones in each angular direction for those zones that reach $r = 0$. Subsequently, the number of zones in the angular directions is doubled for every doubling of radius until a specified number of cells is reached. The outcome of this process is 1) a mesh where the cell aspect ratios are never extreme, 2) a true representation of the multi-D dynamics throughout the whole domain, and 3) a drastically improved timestep limit. Hence, while benefiting from the Courant advantage of Cartesian codes, FORNAX also benefits from employing spherical coordinates for objects that are on average spherical.

WHISKYTHC [233–235] is a state-of-the-art fully general-relativistic, dynamical spacetime, hydrodynamics code optimized for the simulation of compact binary mergers. It is hybrid OpenMP/MPI parallel and employs block-based adaptive mesh-refinement (AMR) [24, 25] with sub-cycling in time and refluxing. WHISKYTHC is robust and mature and has been used in several publications [27, 228, 229, 231, 235, 236].

The spacetime geometry is evolved using the Z4c formalism [102]. For the hydrodynamics, we use either high-resolution shock-capturing central schemes [143] based on high-order reconstruction operators MP5 [264] or WENO [127, 161] or high-order finite-differencing methods. The former combine high-order reconstructions with the Roe vector-flux splitting method; properly taking into account the full characteristic structure of the equations [233]. High-order methods are essential for long-term evolutions of neutron star binaries [26, 236]. WHISKYTHC uses microphysical nuclear EoSs with a simplified neutrino treatment based on the leakage scheme [231]. The composition of nuclear matter is evolved using the consistent multi-fluid advection scheme of [222] ensuring individual local conservation of all the species. For the treatment of low-density regions, WHISKYTHC employs a positivity-preserving limiter to ensure mass conservation to floating-point accuracy [235]. WHISKYTHC is the only numerical relativity code with both conservative AMR and conservative treatment of evacuated regions. These features are crucial to capture the relatively small amount of matter ejected during the mergers and relevant for the r-process.

CLASH will be a component-based multi-physics toolkit, rooted in a massively parallel adaptive mesh refinement (AMR) framework, and capable of accurately simulating the coupled hydrodynamics, radiation transport, thermonuclear kinetics, and nuclear microphysics at play in stellar explosions. CLASH will reach exascale efficiency by building upon current multi- and many-core efficient local physics packages integrated into a task-based asynchronous execution framework based on our current AMR technology. CLASH is currently being developed under the aegis of the ExaStar Application Development component of the Exascale Computing Project (ExaStar is currently a ECP AD seed project).

The software architecture of CLASH is designed to meet the twin requirements of composability and separation of concerns as expressed in Dubey & Graves [71]. Composability is necessary to enable CLASH to be the general tool kit for astrophysics and other communities where different applications can be configured by combining different components. The components will be at various granularities and provide several different functionalities. The key feature is that all such reusable components will have well defined interfaces through which they will be able to interoperate with one another.

The CLASH framework can be thought of as consisting of three layers. The bottom layer is the basic service layer that covers infrastructure components such as AMR and I/O and also utilities such as common solvers that need to interact closely with AMR. This layer will essentially be provided by AMReX, the new AMR software framework being developed and supported by the ECP AMR Co-Design center. At the top layer we have the *client* components which implement physics capabilities and consume services provided by the bottom layer. These components will come from the Castro, FLASH, CHIMERA and Sedona codes.

what does this mean?

only 2 layers defined here

2 Nuclear Astrophysics Research Plan

To understand the r-process, p-process, and the rest of the nucleosynthesis that results from the deaths of massive stars and neutron stars, we must begin long before the explosions that distribute these materials into the interstellar medium. For stars below a critical mass — 8 *Solar masses* (hereafter denoted M_{\odot}) is commonly adopted [114], though the number may be as small as 6 M_{\odot} [33] — degeneracy in the CO core prevents the ignition of carbon burning. These low-mass stars end their lives as cooling CO white dwarves with masses less than 1.1 M_{\odot} , surrounded by the envelope they ejected in their AGB star *winds*. For much more massive stars, from perhaps 11 M_{\odot} up to roughly 100 M_{\odot} [293], hydrostatic carbon, neon, oxygen, and silicon burning leave a core composed of iron, cobalt, nickel, and neighboring species, commonly referred to as the *iron-peak* nuclei. Once this core grows beyond the maximum stable mass for a system supported by electron degeneracy pressure (the so-called Chandrasekhar mass) it collapses. For stars in the lower portion of this mass range, by means we will discuss in §??, the collapse produces a neutron star and a shock wave that disrupts the stellar envelope, leading to a supernova. In the higher mass region, a black hole forms, preventing the shock wave from being sufficiently energized to disrupt the stellar envelope, and the supernova fails. However, if the progenitor star is rapidly rotating, an alternate mechanism, the *collapsar* mechanism [162] — driven by an accretion disk surrounding the black hole and producing jets along the rotational axis — may produce a *peculiar*, hyper-energetic supernova explosion (with explosion energies as much as an order of magnitude larger) and an associated *long duration* (>2 second) γ -ray burst. Alternatively, these hyper-energetic supernova and associated long duration γ -ray burst have been ascribed to the formation of an extremely magnetized neutron star, a *magnetar*. Stars in the intermediate range, above 8 M_{\odot} but below 11 M_{\odot} [242], successfully ignite carbon and neon burning, but under degenerate conditions, leading to massive O-Ne white dwarves. In the upper part of this range, above perhaps 10.5 M_{\odot} [115], models suggest that shell burning causes the mass of the O-Ne core to increase beyond the stable Chandrasekhar mass before the star loses its envelope. Here, like in the iron core case, the core collapses, producing a supernova [189]. Configurations of disk around a black hole or massive neutron star, sans the stellar envelope, can also result from the mergers of neutron stars formed in prior supernovae or the tidal disruption of a neutron star by a black hole.

The collapse of stellar iron and O-Ne cores, collapsars, mergers of neutron stars, and other similar events have all been suggested as sites of the r-process. Thus, to understand the r-process we must understand each of these scenarios and features that distinguish them. This understanding begins with the late stages of stellar evolution (§2.1). In the subsections that follow, we will detail our scientific plans to study the mechanisms of core-collapse supernovae (§2.2), the mergers of neutron stars (§2.4) and the nuclear physics that underlies these events (§2.5). In the final subsections, we will describe our plans to calculate the products of these simulations, the newly-made isotopes ejected into the galaxy, ready to be formed into future generations of stars and planets (§2.6), and their observable consequences (§2.7).

2.1 Modeling the Lives of Massive Stars

The successful Standard Solar Model, and the accompanying models for more distant stars, are one of the fundamental accomplishments of 20th Century astrophysics and among the first scientific triumphs enabled by computational science. The success to date has come despite the assumption of spherical symmetry. The most significant non-spherical effect in most stars is convective motion, which, in spherically symmetric models, has been treated using mixing-length theory, where a diffusion operator is used for mathematical convenience, even though the process is actually turbulent advection. (Mathematically they differ by having a second-order versus a first-order differential operator in space; i.e., parabolic versus hyperbolic differential equations.) A second departure from spherical symmetry comes from rotation. The ancestors of core-collapse supernovae are O-stars, which are observed to begin life with rapid rotation [e.g. 87], and to slow down as they evolve. These massive stars evolve very rapidly once neutrino cooling dominates (carbon burning and thereafter). A crucial feature for stellar core collapse and later stages is the angular momentum of the

ref for this statement?

collapsing core. Angular momentum can be transferred by gravity waves (g-modes) and by magnetic fields, both of which are supposed to be generated by turbulent convection of the stellar plasma.

Despite the success of the Standard Solar Model, massive stars are not perfectly spherical. This is especially true as massive stars approach core-collapse. The catastrophic cooling by neutrino emission drives ever more vigorous nuclear burning, and convective velocities increase. Nuclear burning tightly couples to turbulent convection so that fuel is consumed in chaotic bursts [e.g., 14]. Core and shell burning are dominated by large scale modes, which are of such low order that they do not cancel to produce a smooth spherical behavior.

Co-PI Couch achieved the first 3D simulation of the final minutes of iron core growth in a non-spherical $15 M_{\odot}$ massive star [61]. This work captured the development of strong convection driven by violent silicon burning in the shell surrounding the iron core. Using a 21 isotope reaction network, this convective burning built the iron core to its critical mass and collapse ensued. The non-spherical motions generated by 3D convection is substantial at the point of collapse, with average convective speeds $\simeq 20\%$ of core-collapse infall speed. Couch et al. [61] found that the enhanced post-shock turbulence resulting from a non-spherical progenitor aids successful explosions. Müller et al. [196] followed with a 3D simulation of the last minutes of oxygen shell burning to the onset of core-collapse of an $18 M_{\odot}$ star. They also found Mach number increases to ~ 0.1 at collapse and that an $\ell=2$ mode dominates the flow pattern, but used an artificial moving inner boundary to model the contraction of the silicon and iron core. Müller [193] showed that this 3D progenitor structure was even more conducive to aiding explosion than what was reported by Couch et al. [61]. Realistic, non-spherical initial progenitor structure is evidently important for understanding the CCSN mechanism and attendant conditions for nucleosynthesis.

As part of this SciDAC project, we will build fully 3D models of massive stars at the point of core collapse for use as realistic initial conditions for simulations of CCSNe. In close collaboration, TEAMS members at MSU and SBU will adapt the low-mach number hydrodynamics code Maestro to the simulation of massive stellar evolution. This effort will extend and improve previous work in every regard. We will employ a larger nuclear network, capable of directly capturing the complex quasi-equilibrium burning that builds up the iron core during the final stages of massive stellar evolution. We will simulate the entire star, rather than simply the octant geometry explored by Couch et al. [61]. And we will include the full inner core, unlike Müller et al. [196], and follow the neutronization of the inner core directly. Additionally, the low-mach number approach of Maestro will allow us to simulate much longer timescales, perhaps several hours leading up to core collapse. Initial stellar evolution modeling will be done in spherical symmetry using MESA [218–220]. In addition, in later years we will include rotation in both the 1D and 3D stellar evolution modeling, work that will be enabled by significant enhancements to Maestro achieved through work proposed as part of this SciDAC project (see Section 3.1).

Modeling massive stellar evolution in 3D will be carried out principally by MSU graduate student Carl Fields (NSF and Ford Foundation Fellow) and the SBU postdoc supported by this project under the direction of Co-PIs Couch and Zingale. All data for these 3D progenitor models will be made publicly available via the Nuclear Astrophysics Data Archive proposed as part of this project (see Section 4.4) and will be used extensively within the TEAMS collaboration.

2.2 Modeling the Deaths of Massive Stars

The collapse of the stellar core, whether composed initially of iron or oxygen–neon–magnesium, proceeds until super-nuclear densities, larger than the densities inside atomic nuclei, are reached. The inner core becomes incompressible under these extremes, bounces, and, acting like a piston, launches a shock wave into the outer stellar core. This shock wave will ultimately propagate through the stellar layers beyond the core and disrupt the star in a supernova explosion. However, the shock stalls in the outer core, losing energy as it plows through it, and exactly how the shock is revived is unknown. This is the central question in core-

collapse supernova theory. [For a more complete review, the reader is referred to 121, 185].

After core bounce, $\sim 10^{53}$ erg of energy in the form of neutrinos and antineutrinos of all three *flavors* (electron, muon, and tau) is released from the newly formed *proto-neutron star* (PNS) at the center of the explosion. The observed kinetic energy of the supernova explosion is $\sim 10^{51}$ erg. Past simulations [29, 287] demonstrate that energy in the form of neutrinos emerging from the PNS can be deposited behind the shock and may revive it. This is the so-called Wilson *delayed-shock mechanism* and is central to our modern understanding of core-collapse supernovae. However, while a prodigious amount of neutrino energy emerges from the PNS, the neutrinos are weakly coupled to the material directly below the shock. The neutrino heating is very sensitive to the distribution of neutrinos in energy (or, equivalently, *frequency*; a neutrino's energy and frequency are simply related by $E = hf$, where h is Planck's constant, and are interchangeable) [47, 117, 124, 186] and direction of propagation (specified uniquely by two angles), at any given spatial point behind the shock [180, 187]. In turn, this ultimately requires *multi-frequency, multi-angle (Boltzmann)* neutrino transport in order to compute accurately the neutrino distributions in this region in frequency and angle. This renders the core-collapse supernova problem a truly multidimensional [six dimensional (space plus neutrino frequency and angles)], exascale problem. Even on exascale architectures, an approach to this ultimate goal must be staged, beginning with 3D, multi-frequency *moments models* of the neutrino radiation field and progressing eventually to multi-frequency, multi-angle Boltzmann models. In these moments models, the neutrino distributions in angle are approximately represented by the lowest order *angular moments* of the neutrino distribution function: the neutrino energy density and three momentum densities (one for each spatial dimension), all as a function of neutrino frequency. While this is an approximation to the full Boltzmann treatment, moments models can be quite sophisticated and can describe well what would be obtained if the Boltzmann equation were solved directly.

The neutrino heating may be aided by fluid instabilities (e.g., convection) in the PNS [40, 42, 188, 258, 288], which may boost the luminosity of this central neutrino bulb. Convection directly beneath the shock fundamentally alters the nature of neutrino shock reheating [36, 42, 49, 84, 101, 124] relative to the spherically symmetric case, allowing simultaneous down flows that fuel the neutrino luminosities by accretion and up flows that bring energy to the shock. A recently discovered multidimensional instability of the shock wave itself, the *Standing Accretion Shock Instability* (SASI; [31]), dramatically alters the shock and explosion dynamics [31, 50, 119, 213]. Centrifugal effects in a rotating stellar core [42, 84], and other rotational effects, can change supernova dynamics quantitatively and perhaps qualitatively. Stellar core magnetic fields, increased by compression during collapse, convection (e.g., via a dynamo), and rotation (through wrapping and shear; in the latter case the magnetorotational instability (MRI) may occur and, if so, would dominate the field evolution), may also play a significant role in driving, and perhaps collimating, at least a subset of core-collapse supernova explosions where rapid rotation is present [4, 45, 267]. Nuclear burning must also be included, as the energy released by burning in the compressed and heated material near the shock wave helps to power the explosion. Finally, the PNS is an extremely dense, compact object, and its gravitational field is not well described by Newtonian gravity. Rather, Einsteinian gravity—i.e., *general relativity* (GR)—is required. It is, after all, the release of gravitational binding energy as the core collapses and the PNS forms that provides the energy in neutrinos, which we believe ultimately powers the explosion. Thus, GR is essential to capturing the fundamentals of core collapse supernova energetics.

Recent progress to uncover the details of the core collapse supernova explosion mechanism has been exponential. Two-dimensional modeling has matured rapidly as more simulations have been performed by different groups with the requisite input physics [37, 39, 53, 154, 177, 178, 194, 195]. Most notably, the recent study by Burrows et al. [53] details very encouraging comparisons between the three groups (MPA, Oak Ridge, and Princeton, the latter two of which are involved in this proposal) that now run with both general relativity and sophisticated neutrino transport with state of the art weak interaction physics. The simulation outcomes are in qualitative agreement with regard to obtaining explosion, across a range of progenitor masses

between 12 and 25 M_{\odot} . This is a major step forward. The focus now turns to quantitative differences that must be explored, understood, and, where possible, eliminated (note that some quantitative differences should be expected given the use of independent codes based on different numerical methods). Convergence of the simulation outcomes across groups must occur before we can claim to have gained solid insight into the fundamentals of the explosion mechanism. Moreover, additional simulations [62, 65, 197] that have taken the major step to initiate stellar core collapse from non-spherical progenitors, which have taken into account three-dimensional silicon burning in the latest stage of stellar evolution, have made clear that the resultant perturbations in the outer iron core and silicon layer enhance the growth of fluid instabilities in the core (convection and/or the SASI), thereby enhancing the likelihood of explosion and, in some cases, even changing qualitatively the simulation outcomes, from a failure to explode to explosion.

In the subsections below, we will detail the number of ways that CCSN can potentially contribute to the production of r-process, p-process and other nuclei. For each, we will discuss how TEAMS simulations will push back our ignorance of these potential nucleosynthesis sites.

2.2.1 Iron Core Collapse

Large overabundances of elements in the periodic table spanning from oxygen through nickel are observed in core-collapse supernovae and their remnants. Observations of nuclear abundances allow nucleosynthesis calculations to place powerful constraints on conditions deep in the interior of supernovae and their progenitors, places hidden from direct observation. Unfortunately, until recently, the frequent failure of self-consistent models to produce explosions has resulted in the reliance of core-collapse supernova nucleosynthesis modeling on parameterized models, which replace the inner workings of the supernova with a kinetic energy *piston* [see, e.g., 159, 238, 292, 294] or a thermal energy *bomb* [see, e.g. 204, 269, 276]. These models ignore much that we have learned in the past two decades about the nature of these neutrino-driven, hydrodynamically-unstable explosions. In such bomb or piston simulations, the explosion’s energy, its delay time and the *mass cut*, which separates the ejecta from matter destined to become part of the neutron star, are externally supplied parameters. The bomb and piston methods are largely compatible, with the largest differences coming in the inner regions of the ejecta [16]. It is the nucleosynthesis in this inner region that can be strongly affected by the details of the explosion mechanism [82]. In the case of the neutrino reheating mechanism, these effects include interaction with the tremendous flux of neutrinos and the temporal delay in achieving the explosion.

In the innermost regions of the ejecta, the passage of the shock heats matter to temperatures where Nuclear Statistical Equilibrium (NSE) is dominated by free nucleons and α particles (${}^4\text{He}$ nuclei) [see, e.g., 57, 96]. As a result, most of the iron-peak species synthesized in core-collapse supernovae result from α -rich freezeout [290]. As matter expands outward, it cools, allowing the light nuclei to recombine into iron, nickel, and neighboring nuclei. In the case of α -rich freezeout, this recombination is incomplete, leaving a significant fraction of the matter still in the form of free nucleons and α particles. The detailed composition of this ejecta depends on its neutronization, which is set in the inner regions of the ejecta by neutrino interactions. Spherically symmetric models [e.g., 159, 269, 294] find that above the innermost α , iron, and nickel dominated regions, passage of the shock leaves a layer rich in the α isotopes: ${}^{40}\text{Ca}$, ${}^{36}\text{Ar}$, ${}^{32}\text{S}$, and ${}^{28}\text{Si}$ — the products of incomplete silicon burning. Above this is a layer of ${}^{16}\text{O}$, in the outer portions of which significant fractions of ${}^{20}\text{Ne}$, ${}^{24}\text{Mg}$, and ${}^{12}\text{C}$ are found. Finally, above the ${}^{16}\text{O}$ layer we find the helium layer and hydrogen envelope, if they were not driven off as part of a stellar wind. Though the passage of the shock does not grossly alter the composition of the outer layers, Woosley et al. [291] and Heger et al. [98] have shown that appreciable amounts of some rarer isotopes can be produced by the combination of the shock passage and the neutrino flux. This including several isotopes, like ${}^{19}\text{F}$, ${}^{138}\text{La}$, and ${}^{180}\text{Ta}$, for which this ν -process could be the dominant production mechanism.

One common property exhibited by all recent 1- and 2-D simulations utilizing multi-frequency neutrino

transport [see, e.g., 38, 42, 157, 168, 237] is a decrease in the neutronization in the outer part of the neutrino heating region, due to neutrino interactions. This is a feature that the parameterized bomb/piston nucleosynthesis models discussed above cannot replicate because they ignore neutrino transport. The neutronization is important because GCE calculations and the relative neutron-poverty of terrestrial iron and neighboring elements strongly limit the amount of neutronized material that may be ejected into the interstellar medium by core-collapse supernovae [272]. Hoffman et al. [105] placed a limit of $10^{-4} M_{\odot}$ on the typical amount of neutron-rich ($Y_e < 0.47$) ejecta allowed from each core-collapse supernova. Past multidimensional models of core-collapse supernovae using *gray* (frequency-integrated or -averaged — i.e., not multi-frequency) neutrino transport that did produce explosions tended to greatly exceed this limit [see, e.g., 101, 124]. To compensate, modelers have been forced to invoke the fallback of a considerable amount of matter onto the neutron star, occurring on a time scale longer than was simulated [83].

Work by Fröhlich et al. [80, 81] and Pruet et al. [224, 225] indicates that neutrino-powered explosions using multi-frequency neutrino transport result in nucleosynthesis products qualitatively different in composition from either the parameterized bomb/piston nucleosynthesis models or gray transport models of the core-collapse mechanism. In the models used by both Fröhlich et al. and Pruet et al., the neutrino physics was artificially tuned in order to drive explosions. Thus, these models are also parameterized, but the explosion is explicitly neutrino-driven and the ejecta are determined self-consistently by the radiation hydrodynamics. These models have shown that the inclusion of neutrino captures on the ejecta, which decrease the neutronization, remedies some defects in the predictions made by previous models of nucleosynthesis, all of which neglected this important piece of physics. These effects remove the over-production of neutron-rich iron and nickel isotopes that have plagued parameterized bomb and piston models. These simulations [80, 225] also show enhanced production of Sc, Cu, and Zn; elements which observations of metal-poor stars [see, e.g., 55, 93] suggest are 3–10 times more abundant than previous models predicted. Moreover, Fröhlich et al. [81] revealed a significant neutrino-driven flow to proton-rich nuclei above $A=64$, which was also confirmed by Pruet et al. [224], leading to the discovery of the *vp-process* and suggesting that the innermost ejecta of core-collapse supernovae may be the production site of the light p-process nuclei. Given the importance of these findings, it is essential these studies now be conducted in the context of self-consistent, multidimensional models using multi-frequency neutrino transport.

Convection and other hydrodynamic instabilities greatly complicate the picture developed from spherically symmetric models, even destroying strict compositional layering in the progenitor. Arnett, in collaboration with Meakin and others [15, 21, 176], has demonstrated that convection occurs in the oxygen shell even before the explosion. For two decades it has been known that Rayleigh-Taylor instabilities originate at the Si/O and (C+O)/He boundaries [94, 100, 130, 198, 204]. However, these instabilities do not mix nickel to sufficiently high velocities to account for observations of Supernova 1987A [see 175, and references therein] and other core-collapse supernovae. The implication is that gross asymmetries must be present in the core and be part of the explosion mechanism itself. However, few simulations to date have directly considered both neutrinos and the impact of this multidimensional behavior on the nucleosynthesis. An exception is Kifonidis et al. [134, 135], who used a parameterized neutrino luminosity and temperature [similar to 124] coupled to an *α -network* (linking the 14 α nuclei from ^4He to ^{60}Zn), while separately tracking the neutronization. These simulations showed significantly higher velocities for iron-rich clumps than previous models [94, 100, 198]. However, as a result of their gray neutrino transport, the simulations of Kifonidis et al. [135] still predict the ejection of much larger quantities of neutron-rich iron group elements than GCE seems to allow. In contrast, preliminary analysis of our multi-frequency 2D CHIMERA results, based on analytic extension of tracer histories to late times [150, 179], document the development of the proton-rich ejecta and the *vp-process*, though at a somewhat reduced level compared to previous work [81], because of faster expansion time scales. This highlights the critical dependence of the strength of the *vp-process* on the expansion time scale. Extended simulations [12, 248, 279] highlight the impact (or lack thereof) of the forming reverse shock

on the expansion time scale and hence the strength of the νp -process, providing further justification for the merger of explosion mechanism models and nucleosynthesis models, as well as their extension until the nucleosynthesis is complete.

Recent results from CHIMERA [95, 104], complemented by similar results by [281], foreshadow the results we expect from TEAMS simulations, using FORNAX, FLASH, and CLASH. Based on extended CHIMERA models (run to ~ 1.5 seconds after bounce) [37], which included an α -network in the simulation and tracer particles used for post-processing nucleosynthesis calculations (see §??), the simulations allowed nucleosynthesis to be examined much more deeply than most CCSN simulations. With the inclusion of neutrino interactions and self-consistent multi-dimensional fluid flow, it is hardly surprising that these simulations, with explosion energies and mass cut developing self-consistently from the physics of the model, differ from their parameterized equivalents. Aside from differences in detail in the intermediate mass elements, the CHIMERA simulations exhibit much larger amount of proton-rich heavier elements, due to the inclusion of neutrino interactions. However, little evidence for the νp -process is seen, likely due to the suppression of the PNS wind (see §??). In addition, the ability of accretion flows in multi-dimensions to dredge up matter from near the neutron star, resulting in a significant contribution of neutron-rich heavy nuclei as well, at least in lower mass progenitors. For example, the $12 M_{\odot}$ model of Harris et al. [95] exhibits significant enhancements in the production of several very neutron-rich species, for example, ^{48}Ca , ^{50}Ti , ^{54}Cr , ^{58}Fe , ^{64}Ni , ^{70}Zn , ^{76}Ge , and ^{82}Se . These are produced in a small amount of matter that undergoes a neutron-rich ($Y_e < 0.45$), but α -poor, freezeout as it is dredged up from near the PNS. Preliminary calculations using FORNAX suggest a similarly wide spread in ejecta Y_e , from proton-rich to neutron-rich, in the generic progenitor explosion. The release of neutron-rich matter from CCSN is potentially relevant to the production of weak r-process species in these events. This is a question we plan to pursue aggressively under TEAMS, using initially CHIMERA, and later FORNAX and CLASH, simulations as we systematically explore the nucleosynthesis of these models.

2.2.2 O-Ne Core Collapse

A scenario very similar to the collapse of an iron core in a massive star is possible for some intermediate mass stars ($8\text{--}11 M_{\odot}$) with cores composed predominantly of oxygen and neon. As in the iron-core case, collapse proceeds to form a PNS, however the collapsing core in these stars is surrounded by low-density helium and hydrogen layers, in contrast to the much-higher-density silicon and oxygen layers present in a more massive star, producing a sharp density gradient at the surface of the stellar core. Without the tamping of the explosion by the stellar envelope, explosions from O-Ne cores proceed more rapidly than their iron-core counterparts but also tend to have much lower explosion energies (~ 0.1 B). A unique r-process mechanism has been suggested [206] in these stars. Once the shock, reinvigorated by neutrino heating, reaches the sharp surface gradient, it accelerates strongly, leading to the fast expansion of the shocked matter in the weakly neutron-rich surface layers of the ONeMg core, producing an r-process. Unfortunately, the best simulations of this process to date [123, 136] do not produce the required combination of temperature, entropy, and expansion time scale for currently available stellar progenitors. However, these studies were of a single, now thirty-year old, progenitor [207], thus this potential site begs reexamination with more modern progenitors [e.g. 128, 129]. In addition, we can extend these investigations to newer progenitors in this range, under development by our stellar evolution collaborators. The current capability of CHIMERA to use larger nuclear reaction networks will allow self-consistent inclusion of the electron captures that drive the collapse, including $^{20}\text{Ne} \rightarrow ^{20}\text{F} \rightarrow ^{20}\text{O}$ and $^{24}\text{Mg} \rightarrow ^{24}\text{Na} \rightarrow ^{24}\text{Ne}$ [173], as well as the subsequent neutron-rich oxygen and neon burning, offering a significant improvement on previous work. FORNAX doesn't now have such a large nuclear network, but in collaboration under TEAMS will incorporate this capability to post-process its explosion ejecta to determine detailed nucleosynthetic yields. Beyond Year 2 of TEAMS, this capability will be available in CLASH as well.

A further examination of the same simulation [123] by Wanajo et al. [280] have shown that 2D simulations

can eject moderately neutron-rich convective “lumps” capable of forging elements up to $N=50$, presaging the similar results from iron-core collapse models [95, 281]. Wanajo et al. [280] raise the possibility that 3D simulations and/or improved resolution could result in tiny amounts of matter that are neutron-rich enough to produce at least the lower-mass *weak r-process* species. Under TEAMS, we plan to explore this exciting possibility using CHIMERA and later CLASH and FORNAX.

2.2.3 Proto-Neutron Star Winds

Once the explosion develops and accretion onto the PNS stops, simulations demonstrate the development of a wind from the surface of the cooling neutron star, driven by neutrinos. The observation in earlier supernova models [e.g., 288] of neutron-rich matter in the high-entropy hot bubble, which surrounds the PNS after the supernova explosion has been launched, naturally suggested this PNS wind as a potential site of the r-process [183]. Indeed, Woosley et al. [295] were able to produce a good match to the Solar System r-process pattern as the result of a neutron-rich, α -rich freezeout. However, more modern calculations have not reproduced the extreme entropies ($> 400 k_b$ per baryon) required. Further, Fuller and Meyer [88, 184] pointed out that neutrino interactions can inhibit such a freezeout. Since the protons are tied up in α -particles, which are relatively inert to neutrino interactions, neutrino capture on free neutrons, which produces protons, will destroy the neutron-richness of the matter, severely inhibiting the formation of the r-process. This problem is exacerbated in multi-frequency models, which give rise to proton-rich ejecta [80, 225]. The electron fraction in the innermost hot bubble is set primarily by electron neutrino and electron antineutrino capture on free nucleons. For the matter to be neutron-rich, the difference in the electron neutrino and antineutrino spectra must overcome the mass difference between the neutron and the proton [226]. As a result of the inclusion of more types of neutrino opacity, all current simulations exhibit too small of a difference between the electron neutrino and antineutrino spectra. Thus, as the matter is ejected and the density drops, it becomes proton-rich once electron degeneracy is removed. The simulations used by Fröhlich et al. [80] and Pruet et al. [225] exhibit this proton-richness over the first ~ 1 second after core bounce. Extended spherically symmetric simulations [75, 113] suggest that proton-rich ejecta continue well into the PNS wind phase, suggesting *iron* core-collapse supernovae are an unlikely r-process site. This declaration has been somewhat ameliorated by improvements in the treatment of neutrino-matter interactions, including the effects of the increasing density in the nuclear medium [172, 245] as the neutrinosphere retreats into the neutron star.

However, all of these extended studies of the development and neutronization of the PNS wind have been undertaken in spherical symmetry. They therefore exhibit a clean division between the accretion phase of the supernova and the explosion phase that is not present in multi-dimensional simulations. In each of a set of 4 models by Bruenn et al. [37], covering 12-25 M_\odot progenitors, significant accretion ($> 0.01 M_\odot/s$, at a noisily constant rate) onto the PNS continues for at least 1.5 seconds, despite the mean shock radius exceeding 10,000 km. This accretion effectively suppresses the development of the PNS wind. Müller [192] cite the differing characteristics of the downflows in this late phase (up to 1 second after bounce is reported), laminar in 2D and turbulent in 3D, as an important factor in the lack of PNS wind in the 2D models in that paper. However the downflows exhibited in the 2D CHIMERA simulations seem intermediate between the 2D and 3D cases of Müller [192], indicating the need for further study in 2D and 3D. Further, with the large difference in explosion energies, Müller [192] 2D models exhibit less than half the explosion energy of the corresponding 3D models (and both are significantly less than the models of Bruenn et al. [37]), it is difficult to separate the effects of dimensionality from the effects of explosion dynamics. Clearly, to fully understand the role of the PNS wind in producing the r-process we must continue multi-dimensional simulations until the PNS wind develops to better understand the true conditions in this wind. CHIMERA and FORNAX have already shown themselves to be effective tools for such studies, with currently running models approaching 2 seconds. With improvements to CHIMERA and FORNAX planned under TEAMS (see 3.2.1), especially improvements to the EoS and inclusion of medium corrections to neutrino opacities [172, 245], we plan extended, improved simulations to uncover the birth of the PNS wind and examine the impact of ongoing

accretion on the nucleosynthesis that results. We expect that at least 3 seconds after bounce is needed to explore this issue, and perhaps as much as 5 seconds.

2.2.4 Magnetars and Magnetorotational Turbulence

All stars rotate and have magnetic fields. Following massive star core collapse, rapid rotation and strong magnetic fields can lead to powerful outflows from the PNS [46, 285, 286, 289] that were first suggested as a possible CCSN explosion mechanism by LeBlanc & Wilson [149]. Such magnetically-driven explosions are an interesting possibility for r-process nucleosynthesis [289]. Modern stellar evolution calculations, however, indicate that diffusive mixing and magnetic braking effects efficiently transport angular momentum out from massive stellar cores [99, 219, 259]. Given the observed distribution of initial rotation speeds of massive stars, this strongly indicates that the vast majority of massive stellar cores at collapse do not rotate rapidly enough to lead to magnetorotational explosions. It is likely that magnetorotational effects play a dominant role in only about 1% of all core collapse events, corresponding to the fraction of Long Gamma-ray Bursts (LGRBs) and so-called broad line Type Ic SNe [190].

We know, however, that highly magnetized neutron stars known as magnetars are not that uncommon, accounting for about 10% of the neutron star population. This points to a generic mechanism in stellar collapse that results in strong magnetic fields in the compact remnant some significant fraction of the time. Additionally, we know there are mechanisms that can amplify even initially very weak magnetic fields with exponential growth rates. Akiyama et al. [5] pointed out that a collapsing stellar core is generically unstable to growth of the magnetorotational instability (MRI) [17, 56]. The MRI grows exponentially on the rotational time scale and taps the energy of differential rotation to amplify the magnetic field and drive turbulence. Turbulence has been shown to play a leading-order role in driving the dynamics of the stalled supernova shock and explosion [48, 66, 67, 203]. Even in non-rapidly rotating stellar cores, magnetic fields could be amplified to strengths that could qualitatively change the behavior of the turbulence in the post-shock gain layer as well as lead to enhance turbulent dissipation to heat, which can also aid explosion [e.g., 215, 270].

As part of this SciDAC project, we will extensively explore the conditions that could lead to magnetar formation through 3D CCSN simulations of rotating and magnetic progenitor stars. Our primary tool for this exploration will be FLASH, which implements a state-of-the-art high-order unsplit staggered mesh (USM) MHD solver in an AMR framework. This work will build on current studies underway using FLASH being executed by PI Couch as part of a DOE INCITE project. The bulk of this work will be executed by the MSU postdoc under the supervision of PI Couch and Co-I Roberts. We will explore a range of realistic initial rotation profiles and magnetic field strengths based on stellar evolution calculations that include prescriptions for rotation, magnetic field amplification, and angular momentum transport [e.g., 99, 219].

Additionally, we will explore on rotation and magnetic fields can modify the behavior of post-shock turbulence in the CCSN context. Capturing turbulence in CCSN can be extremely demanding computationally [3, 67, 230, 232]. This is particularly the case for turbulence-driving instabilities such as the magnetorotational instability (MRI), where the fastest growing mode in the CCSN context can be on the order of tens of meters [c.f., 5, 46, 191, 210]. Thus, we propose a multi-pronged approach to studying magnetorotational turbulence in CCSNe. First, we will carry out simplified physics simulations in reduced domains small enough the MHD turbulence can be simulated directly [e.g., 191, 210]. Second, we will explore new numerical methods and algorithms for simulating MHD turbulence that have never before been applied to the CCSN context. This will include exploration of sub-grid turbulence models such as large-eddy closures [223] and very high-order methods for MHD [240] that are better able to capture the correct behavior of the turbulent cascade at finite resolution.

Finally, we will explore the implications for heavy element nucleosynthesis from magnetorotational CCSNe and newborn magnetars by computing the post-processing nuclear yields from tracer particles included in our 3D MHD CCSN simulations.

2.3 Epstein, Colgate & Haxton mechanism

As described above, most scenarios for the r-process in modestly neutron-rich matter rely on a “hot r-process” where α -rich freezeout ties up much of the mass in ${}^4\text{He}$, freeing the relatively low abundance of free neutrons to concentrate on the small number of heavy seed nuclei. However, the essential physics of the r-process does not rely on high temperature, simply a sufficient source of neutrons in matter composed of ${}^4\text{He}$ and a smattering of heavy nuclei. As Epstein et al. [73] pointed out, the helium layer of a star undergoing a CCSN can meet the requirements for a “cold r-process” if several conditions are met that, in total, yield the required neutron/seed ratio of ~ 100 . Modern studies of both the nuclear physics and astrophysics of this process have made these conditions clearer [18]: the combination of the neutron source ν which in this mechanism comes from neutrino reactions that break up ${}^4\text{He}$ ν and seed abundance must produce the necessary neutron/seed ratio, and the neutrons produced must efficiently capture on the seeds. These conditions point to the outer helium layer as the most likely site, due to relative dearth of neutron poisons like carbon, nitrogen and oxygen there, and to stars of very low metallicity. This has the interesting consequence that this mechanism operates as a full r-process, capable of reaching the transuranics, only in early Galactic history, where stars can be found with pristine He shells, dusted by very small amounts of seeds like Fe. The mechanism also depends on neutrino physics: the needed neutron production arises if the neutrino mass hierarchy is inverted, rather than normal, an issue not yet resolved by experiment. The mechanism is interesting from the perspective of galactic chemical evolution, as it suggests there could be r-processes that operate only at early times, contributing to the nucleosynthesis found in old globular halo stars, but that then give way to more robust mechanisms at later times, such as r-process nucleosynthesis in neutron star mergers.

Compared to studies to date [see, e.g., 19], models of this mechanism with CHIMERA would replace a parameterized spherically symmetric explosion with multi-dimensional fluid flow and replace parameterized neutrino luminosities and spectra with those determined self-consistently within CHIMERA. The challenge is temporal, as the outer helium layer is perhaps 100,000 km from the center of the star. Even at the speed of the supernova shock, this requires running the simulation for approximately 10 seconds. This represents a further extension beyond the 3-5 second CHIMERA models with which we plan to explore the development of the PNS wind. We therefore anticipate this as a Year 3 or later deliverable. This extended timeline would also allow the simulations to benefit from improvements in the high density opacities planned for Year 1, opacities that become increasingly important with time as accretion slows and the neutrinosphere recedes into the neutron star.

2.4 Modeling the Deaths of Neutron Stars

A small fraction of all neutron stars (NSs) end their life with a violent collision with another compact object, another NS or a black hole (BH). These collisions can be driven by the loss of orbital energy due to the emission of gravitational waves (GWs) at the end of a complex binary evolution, or they can occur as a result of many-body interactions in dense stellar environments [153]. Presently, about a dozen double NS systems in and around our galaxy are known. Of these, roughly half will merge within the Hubble time [216]. There are no known NS-BH binaries. The inferred event rates based on these observations and on population synthesis models are still very uncertain. For double NS binaries the merger rate ranges from one every 10^3 years to one every 10^6 years, for a Milky-Way-type galaxy [2]. The rates for NS-BH binaries are even more uncertain, with one predicted every $10^4 - 10^8$ years per Milky-Way-type galaxy.

Despite being rare events, these mergers are violent enough to produce photons throughout the entire electromagnetic (EM) spectrum (from the radio to the hard X-ray and γ -ray) detectable to cosmological distances. Indeed, these mergers typically result in the formation of BHs surrounded by thick, massive (fraction of a M_{\odot}) and hot (several MeV) accretion disks that accrete on a timescale of less than a second. In the case of double NS mergers, millisecond magnetar could be created in some areas of the parameter space, depending on the EoS of dense nuclear matter. These systems are expected to produce relativistic jets and are considered

the most likely candidates for the central engine of short gamma-ray bursts (SGRBs) [22, 205], whose origin and launching mechanism is still a mystery. They are also one of the main target of GW detectors like LIGO [1], which will be able to detect them to hundreds of Mpc, once it will reach its design sensitivity [2].

More importantly for us, NS mergers also result in the dynamical ejection of up to $\sim 10^{-2} M_{\odot}$ of NS matter due to tidal torques or shocks during the merger [20, 76, 111, 144, 231]. Neutrino driven winds, viscous, and magnetic processes in the accretion disk left around the merger remnant can also contribute to these outflows [69, 74, 181]. These outflows are expected to be very neutron rich, because they are made of catalyzed NS material, and will undergo the strong r-process [146, 148, 182]. A tentative confirmation of this scenario has been recently provided by the detection of an infrared re-brightening of the afterglow of Swift SGRB 130603B [23, 268], which has been interpreted as being due to the radioactive decay of byproducts of the r-process.

Other evidences linking NS mergers to the production of a significant fraction of the strong r-process nuclei include the Solar system abundance of ^{244}Pu [112, 274, 277], the abundance patterns in metal-poor ultra-faint dwarf galaxies [126], and recent galactic chemical evolution models [141]. All of these suggest that strong r-process elements are preferentially produced by rare events with a high yield.

The forming picture is that regular CCSNe are the main contributors of weak r-process elements, while NS mergers and possibly peculiar CCSNe (§2.2) are the site of production of strong r-process nuclei. However, our understanding is still far from complete and the exact degree to which each of these production sites contributed to present day Solar System abundances is unknown. For the NS merger scenario, in particular, there are a number of outstanding theoretical problems that need to be solved. One is that the event rate for NS merger is only very weakly constrained. Another, is that the properties of merging NS systems that will determine the amount of ejecta – masses and spins – are also very poorly known, especially for NS-BH systems. These two issues are going to be addressed by GW observations of LIGO, which will be able to provide population statistics in the years to come. From the theory side, however, **it is necessary to develop a quantitative understanding of how to translate binary population statistics and merger rates from LIGO into r-process elements production rates.**

2.4.1 Nuclear Decompression

Because of the extremely small electron fractions, the decompression of NS matter in the aftermath of NS mergers is expected to robustly result in the strong r-process [146, 148, 182], as also confirmed by recent parametrized nucleosynthetic calculations [e.g. 160]. However, considerable uncertainty remains in the quantity and quality of r-process production in such events. Some models for these process considered competitive include Newtonian gravity and limited approximations to the neutrino transport [e.g. 142]. On the other hand, both relativistic gravity and neutrino irradiation have been shown to be important to the setting the neutronization of the matter [77, 282] and, consequently, for the yield of the r-process. Even the most advanced simulations to date [77, 231, 253], still employ the grey (energy-averaged) approximation for neutrino transport. However it is expected that a quantitative characterization of the nucleosynthetic yield from NS mergers will require simulations with energy-dependent radiation transport [77]. Recently, the inclusion of neutrino oscillations induced by matter-neutrino resonances and ν - ν coherent forward scattering have also been shown to be potentially crucial to model the neutrino radiation field and, as a result, possibly the composition and nucleosynthetic yields of outflows [79, 163, 296, 299].

Magnetohydrodynamics (MHD) processes have also been neglected in calculations of the nucleosynthetic yields from NS mergers. However, magnetic fields can be easily amplified by MHD instabilities operating after merger to values ($\sim 10^{16}$ G) in excess of magnetar levels [137, 138]. They are expected to be one of the main drivers of the evolution of postmerger massive NS remnants over timescales longer than ~ 100 milliseconds [110], likely impacting BH formation timescales and gravitational radiation. The Blandford-Znajek [30] mechanism for the electromagnetic extraction of energy from a rotating BH is currently the favoured mech-

anism for the production of short gamma-ray bursts, e.g., [221, 241, 252]. MHD driven outflows could also contribute to the production of heavy elements [139, 257]. A complete and accurate model of nucleosynthetic yields and multimessenger signals from merging neutron stars will ultimately require MHD simulations with full-microphysics.

Finally, merger codes use tabulated nuclear equation of states (EoSs) that smoothly join with a nuclear statistical equilibrium (NSE) EoS below neutron drip density. However, NSE is not a good approximation for the treatment of NS ejecta as the density drops below $\sim 10^{11}$ g cm $^{-3}$ and the temperature below a few 10^9 K. This is problematic because it prevents the evolution to large distance to robustly measure the amount of fallback material [132]. More importantly, the lack of a reliable treatment of the decompressing ejecta prevents the development of self-consistent simulations including both the dynamical and the disk-wind ejecta.

Under TEAMS, we plan aggressive efforts to build sophisticated models of NS mergers. First, codes currently used for NS mergers, e.g. WHISKYTHC, will gain more sophisticated neutrino transport, improved nuclear physics (§3.3) and MHD (§3.2.4). Second, codes currently used for CCSN, e.g. CLASH, will gain features like full General Relativity and more flexible geometry to enable those codes, which generally have more sophisticated neutrino transport and flexible handling of nuclear composition, to simulate NS mergers. Third, the nuclear decomposition process is very sensitive to the structure of the neutron star, hence the advancement in the nuclear EoS discussed in §2.5 are important.

We will leverage these new capabilities in our codes to explore the impact of different physical effects on the r-process yields. First, we will consider few binary configurations that are expected to be representative of the overall population and explore their evolution with increasingly sophisticated simulations. We will perform the first merger simulations with energy-dependent neutrino transport by the end of year 1. Then, we will improve upon these with the inclusion of better nuclear EoS, better angular discretization of the Boltzmann equation, and, by the end of year 5, with an effective treatment of neutrino oscillations. Parallely, we will also improve the macroscopic description of the mergers with the inclusion of magnetic fields. This approach will also allow us to estimate the systematic uncertainties of our models.

The second step will be to perform parameter space surveys with the goal of covering the entire merger population. The final scientific product will be a distribution of yields as a function of binary parameters. This could then be convolved with the results of population synthesis models and GW observations to infer the total production rate of r-process elements due to compact binary mergers.

2.4.2 Disk Winds

In addition to the decompressed neutron star matter dynamically ejected in the merger itself, it has likely that a substantial amount of r-process nuclei are formed by matter expelled *after* the merger. Simulations show that the massive remnant formed in the merger (either a black hole or a hyper-massive neutron star) is typically surrounded by a rotationally supported disk of $\sim 0.1 M_{\odot}$ of neutron-rich material. As this disk matter drains and accretes onto the central object, a fraction of it is likely blown away in disk winds driven by neutrino or viscous heating, magnetic flinging, or nuclear recombination. Weak interactions and neutrino irradiation from the central remnant generally serve to leptonize the disk, reducing the neutron-to-proton ratio and potentially resulting in lighter r-process yields.

Simulations of post-merger disks face distinct computational challenges compared to simulations of the merger itself. Most notably, the disk evolves on a viscous timescale of \sim seconds, which requires following the dynamics over a period of order hundreds of orbital timescales. A treatment of neutrino transport in the disk is also required for determining the composition of the outflows. Given the relatively simple disk geometry, a full solution of the equations of dynamical GR is not necessary; using a fixed metric, or even pseudo-Newtonian approach, can serve as a reasonable first approximation.

While previous simulations have robustly shown that significant nucleosynthesis occurs post-merger, our

understanding of the disk evolution and nucleosynthesis remains limited by several uncertainties. Those calculations performed in 3D have only covered the initial phases of the disk viscous evolution, and so do not cover the full period of mass ejection. Longer term simulations have been carried out in 2D axisymmetry. These simulations often begin with idealized initial conditions consisting of disks in rotational equilibrium. The 2D simulations adopt an artificial parameterized artificial viscosity term; in reality, it is large scale magnetic-fields and magnetically driven turbulence that transport angular momentum and largely control the outflow dynamics. The modeling of neutrino transport and dense matter microphysics has been treated in various levels of approximation.

Under TEAMS, we will carry out post-merger simulations that can address many of the factors limiting our understanding of this potentially important site of r-process production. The collaboration allows us to carry out end-to-end simulations. We will use the WHISKYTHC simulations of the merger dynamics and disk formation to begin with realistic initial conditions for the post-merger phase. We will then follow the long-term (~ 1 second) evolution of the disk in 3D using the FLASH/CLASH and FORNAX codes. Our initial calculations will adopt an artificial viscosity term; this will allow us to compare to similar 2D simulations and determine how the turbulent dynamics and wind formation differs in 3D. In the later years, we will invoke the MHD capabilities of FLASH/CLASH and FORNAX to eliminate the artificial viscosity and follow the dynamics while resolving the magneto-rotational instability. We will further use the 3D multi-group moment based transport method of FLASH/CLASH and FORNAX, along with the state-of-the-art neutrino opacities and equation of state provided by TEAMS, to better characterize the composition of the outflows.

The post-merger simulations we propose will include Lagrangian tracer particles that record the thermodynamic evolution of the outflowing fluid. These trajectories will be fed to TEAMS collaborators for detailed nuclear reaction network post-processing, which will determine the final abundance patterns and radioactive heating rates. These outflows properties will be in turn fed to TEAMS collaborators for radiation-transport post-processing to determine the electromagnetic emission (kilonovae) that result from mergers.

2.5 Modeling Dense Matter and its interactions with Neutrinos

The behavior of merging neutron stars, or a neutron star being tidal disrupted by a black hole, depends critically on the EoS of nuclear matter and on the rates of neutrino-matter interactions. The latter can influence the composition of material ejected in these events and potentially alter the lifetimes of the massive remnants left over after the merger. Similarly, the evolution of the innermost regions of CCSNe are sensitive to the properties of the high density EoS. This may in turn impact the explosion mechanism of CCSNe, the composition of material ejected after the explosion, and the possibly observable neutrino signal.

Due to its complexity, for numerical simulations the EoS is generally prepared as a 3D table in baryon density (n_B), temperature (T), and isospin asymmetry (quantified by the electron fraction, Y_e). Over this large three-dimensional (n_B, Y_e, T) space, there are several different physical regimes and each of those regimes is mostly closely connected to a different observable or different theoretical methods. Nuclear matter (where the neutron and proton densities are nearly equal) near the saturation density and for temperatures smaller than a typical nuclear binding energy is most closely connected to nuclear masses, nuclear radii, giant resonances, and other observables which can be measured at facilities like FRIB. Neutron matter with no protons near the saturation density is poorly constrained by experiment, but is relatively well-described by recent developments in chiral effective theory. At lower densities and higher temperatures, once matter becomes nearly nondegenerate, a virial expansion is more appropriate and the equation of state is entirely determined from scattering phase shifts. Finally, neutron-rich matter at densities above twice the saturation density is most closely connected to observations of neutron star masses and radii. At high temperatures probed in neutron star mergers, pions begin to play a role in describing matter. It is unfortunate that, to date, there are **no equations of state available which properly describe matter in all of these regimes.**

Recent advances in dense matter theory have resulted in new insights and improved input microphysics

for astrophysical simulations. For instance, improvements in the EoS have led to unexpected improvements in the agreement between merger simulations and observed r-process abundances. PI Steiner recently developed two new EoS tables (SFHo and SFHx) which match constraints from nuclear structure and neutron star radius observations [260]. Neutron star merger simulations performed by Sekiguchi et al. [254] found that the newly generated EoS tables, because they implied smaller neutron star radii, ejected more mass (almost 10^{-2} versus $10^{-3} M_{\odot}$) in a merger event. More compact neutron stars with smaller radii have a deeper gravitational potential well, thus naively one might expect the ejecta mass to be smaller. However, simulations show that the deeper well also leads to higher entropies, and it is these high entropies which enable more material to be ejected. Sekiguchi et al. [254] also found that the newer EoS tables generated a larger electron fraction. In turn, this led to a broader r-process distribution which made it easier to explain the universality observed in r-process abundances found in metal poor stars. The incorporation of some these improvements has also demonstrated the supernova explosion mechanism, nucleosynthesis in supernovae and mergers, and the temporal and spectral structure of a galactic neutrino signal are all sensitive to modest changes in the input physics [52, 106, 107, 171, 246, 247, 278].

Nevertheless, this previous work has also highlighted the importance of feedback between the EoS, neutrino opacities, hydrodynamics, and neutrino transport. This emphasizes the need for self-consistency and cautions against a piecemeal implementation of improved microphysics. A major focus of the dense matter effort in this proposal will be to produce a family of EoSs, each with consistent neutrino and weak interaction rates, for astrophysical simulations to properly represent uncertainties associated with nuclear forces and many-body theory.

Ab initio calculations of dense neutron matter [70, 89, 284] and a host of experimental and observational constraints [e.g. 68, 145, 273] have guided the development of mean field theories that are versatile and practical for astrophysical applications [e.g. 251]. These equations of state are also compatible with radii and masses of cold neutron stars inferred from observations [261]. We will select and improve mean field Hamiltonians, both non-relativistic Skyrme models and relativistic versions based on the non-linear generalizations of the Walecka model. These will be tuned to reproduce the zero temperature constraints from ab initio dense matter theory, experiment, and neutron star observations and will be extended to finite temperature using mean field theory.

Using these improved models we will construct the inhomogeneous phase of matter at sub-nuclear density using the Gibbs two-phase construction with surface effects. We will go beyond the single nucleus approximation and develop strategies to match the predictions of NSE at very low density and the predictions of the virial equation of state at high temperature and moderate density. This will allow us to provide a unified description of the thermodynamic properties of matter for the full range of densities, electron fractions and temperature encountered in simulations.

Neutrino opacities in dense matter are modified due to nuclear and electromagnetic interactions between nucleons and leptons [51, 108, 239, 244]. Although much work has been done in this regard, a consistent treatment of both charged and neutral current reactions that account for modified dispersion relations, correlations, relativistic effects, and weak magnetism has never been assembled for implementation in simulations. Further, to calculate the neutrino opacities consistently with the equation of state, the dispersion relations for nucleons and the effective interactions need to be calculated from the equation of state. Using linear response theory we will calculate neutrino cross sections using response functions calculated within the Random Phase Approximation (RPA), and ensure that in the long wavelength limit, these response functions are related to the appropriate thermodynamic derivatives of the EoSs describe above. This will ensure that at relatively low densities characteristic of the neutrino spheres the RPA response functions will match the model independent predictions of the virial expansion [107, 109].

2.6 Modeling the Products of Massive Stars

Large overabundances of elements in the periodic table spanning from oxygen through nickel are observed in core-collapse supernovae and their remnants. Observations of nuclear abundances allow nucleosynthesis calculations to place powerful constraints on conditions deep in the interior of supernovae and their progenitors, places hidden from direct observation. Unfortunately, until recently, the frequent failure of self-consistent models to produce explosions has resulted in the reliance of core-collapse supernova nucleosynthesis modeling on parameterized models, which replace the inner workings of the supernova with a kinetic energy *piston* [see, e.g., 159, 238, 292, 294] or a thermal energy *bomb* [see, e.g. 204, 269, 276]. These two methods are largely compatible, with the largest differences coming in the inner regions of the ejecta [16]. It is the nucleosynthesis in this inner region that can be strongly affected by the details of the explosion mechanism [82]. In the case of the neutrino reheating mechanism, these effects include interaction with the tremendous flux of neutrinos and the temporal delay in achieving the explosion.

An essential element of the TEAMS science plan is to ensure that estimates of the nucleosynthesis are computed for every simulation that reaches a sufficient stage of completion for the estimate to be meaningful. Small reaction networks, typically α -networks, are a staple of most simulations of CCSN and becoming standard in simulations of NSM. However these networks are very limited in the species whose abundances they predict, and only moderately accurate for burning stages above oxygen. PIs Hix, Messer and their collaborators have pushed for the use of more realistic sized networks using the XNet code, with [217] evolving 150 isotopes from H to Zn in FLASH to study thermonuclear flame propagation in thermonuclear supernovae and several CHIMERA models being run with 160 species from H to Ge [104, 155]. While network of 100-200 species capture the major burning stages in CCSN and the freezeout from NSE for ranges in Y_e generally seen in CCSN, they are insufficient for studying the p-process or the r-process. Thus all TEAMS simulations, whether run with a couple hundred species or a dozen, will need to be post-processed.

PIs Surman, Roberts and Hix have considerable experience with post-processing studies of nucleosynthesis that they will bring to bear on the simulations performed by other groups in TEAMS. As described in §2.6, we plan to bring the CHIMERA analysis pipeline into the open-source CLASH framework. As this pipeline includes post-processing with the XNet reaction network, the ability to perform post-processing will become a community resource. By Year 3, individual researchers within TEAMS, student, post-doc or more senior, will have the ability predict the nucleosynthesis products of their own model and collaborate with Surman, Roberts and Hix, all of whom are familiar with XNet, on analysis. From the broad database of TEAMS models, Surman, Roberts and Hix will be able to compose a broad view of nucleosynthesis from CCSN and NSM, identifying trends, uncertainties and needs for further study, which will be fed back to the TEAMS CCSN and NSM groups to guide their simulation plans. Further, we will be able to assemble comprehensive sets of nucleosynthetic products from the various event so serve as input to further calculations, like Galactic chemical evolution, by the community (see §??).

2.7 Observing the Deaths of Massive Stars

An important goal of our project is to deliver end-to-end simulations of stellar explosions that follow the evolution from the initial conditions of the progenitor star system through to the final observable signatures of the exploded remnant. Current experimental facilities provide a broad and constraining data set. By synthesizing the multi-messenger observables of our simulations, we make testable predictions allowing the validation or falsification of specific theoretical models. Our calculations will also better define ways in which experimental facilities can be used to constrain fundamental nuclear physics, such as the nuclear equation of state, the properties of neutrinos, and the sites of heavy element nucleosynthesis.

We plan to model the range of radiation signatures produced by stellar explosions, which include

Gravitational Waves: NSNS and BHNS mergers produce strong gravitational wave (GW) emission that can total more than a percent of the rest mass energy. Such signals should be observable with advanced LIGO

at distances of ~ 100 s of Mpc, and are anticipated to be detected in the coming years. Comparison of GW detections with simulations can be used to infer the parameters of the binary and constrain the NS EOS. The GW emission from supernovae is expected to be much smaller, and detectable only for events within our own Galaxy, except perhaps for rare cases of rapidly rotating progenitors which produce more luminous emission.

Neutrinos: Almost all of the energy from core collapse is radiated in neutrinos....

Photons: By far the most commonly observed signals of stellar explosions are electromagnetic. Wide-field surveys are currently discovering multiple SNe every day, many of which are followed up with panchromatic photometry and spectroscopy. Future surveys such as the Large Synoptic Survey Telescope (LSST) are expected to discover hundreds of thousands of SNe and other explosive transients.

Photons are primarily radiated subsequent to the explosion itself, as the matter ejected in the event emits thermally in ultraviolet, visible and infrared bands with a luminosity visible for months to years. In addition to the thermal energy deposited by the explosion itself, the SN light curve can be powered by the decay of radioactive isotopes synthesized in the explosion, and perhaps in some cases by long-term activity of the central remnant (a rotating magnetized neutron star or accreting black hole). Non-thermal radiation from radioactive decay can be detected in x-ray/gamma-ray bands for very nearby events and Galactic SN remnants. The panchromatic light curves and spectra provide a wealth of information as to the mass, kinetic energy, radioactive content, and elemental composition of the material ejected in the explosion.

Simulating the electromagnetic emission from stellar explosion models requires 3-D time-dependent multi-wavelength Boltzmann transport, coupled to the thermodynamics of the expanding, radioactively heated ejecta. Because the timescale for photon diffusion is typically much longer than that of the explosion itself, the problem can usually be treated in post-processing. The initial conditions of the transport calculation (taken from a multi-physics explosion model) are the density, velocity, and composition of the ejecta once hydrodynamical effects have abated. The dynamical evolution thereafter is one of simple free-expansion, while the gas temperature and ionization state depends upon a strong matter-radiation coupling that must be calculated implicitly. Large time steps can be used, such that the computational expense of the post-processing transport is usually small compared to the explosion simulation itself.

The opacity in photon transport calculations is derived from extensive atomic databases including millions of Doppler-broadened bound-bound line transitions. For the phases where the ejecta remains optically thick (around the peak of the SN light curve) the radiation field approaches a blackbody at the local gas temperature, and the assumption of local thermodynamic equilibrium (LTE) is a decent first approximation for calculating the atomic level populations. However, deviations from LTE can be critical e.g., at late times, when the ejecta becomes optically thin, or in regions where non-thermal energy deposition from radioactive decay dominates. In this case a solution of the coupled rate equations of statistical equilibrium is required.

We plan to develop a pipeline whereby the output of multi-physics SN explosions models are post-processed by the detailed nuclear reaction networks, and then fed into photon radiation transport codes to generate synthetic light curves and spectra from various viewing angles. PIs Kasen and Fryer are experienced in numerical photon transport modeling. The calculations will be performed with the Monte Carlo transport code SedonaBox and ?. Our predicted signatures can be readily compared to a wide database of SN observations in order to physically interpret the data and validate/falsify our explosion simulations. Spectroscopic data in particular contains detailed diagnostics of nucleosynthetic yields, compositional stratification, and total energetics, which will provide powerful tests of our simulations.

We will also calculate the electromagnetic signatures of NSNS and NSBH mergers. Any r-process material ejected in these events should produce a radioactively powered kilonova (or macronova). Such a light curve would be dimmer and briefer than a SN, and given the high-opacity of heavy ($A > 130$) elements, is expected to radiate primarily at red and infrared wavelengths. A major world-wide effort is underway to follow up LIGO detections with telescopes in search of the radioactive powered emission. Our calculations

will aid in this effort by providing improved predictions of the signal properties. Furthermore, analyzing the brightness and color of an observed event with our models would allow estimates of the total mass ejected and whether it is composed of lighter ($A < 130$) or heavier ($A > 130$) r-process elements. Discovering and modeling such a transient would be a watershed moment in nuclear astrophysics, providing the first direct means to detect and diagnose r-process elements *at their production site*.

3 Computational Science Research Plan

Achieving the goals of the previous section is impossible without the development of improved tools to model the lives of massive stars from birth to death, and potentially to a second (neutron star) death. To enable greater physical fidelity, these codes must make efficient use of the coming generation of near-exascale and exascale platforms. This requires considerable development in computational science, both in new algorithms and in refactoring of existing algorithms to take advantage of new hardware designs.

TEAMS will develop, deploy, and build computational ecosystems around several codes. This flotilla of tools includes codes that will be unique in their specific aims (e.g. Maestro for stellar evolution calculations) as well as codes intended to solve very similar problems (e.g. CHIMERA FLASH, Fornax for CCSNe simulation). In all these cases, the multiphysics nature of massive stellar evolution and its cataclysmic end will mean that all of the TEAMS codes will share a subset of common needs, especially in the form of specific physics modules. The particular implementations of the physics of nuclear burning, of neutrino transport, of modeling the response of dense matter, and other important pieces of physics in each of the codes is and will continue to be at different levels of physical fidelity. Indeed, all of the TEAMS codes are currently used in production on leadership-class and other large supercomputer installations (evidenced by several recent awards via INCITE and NSF PRAC, for example), where finite resources and rough equivalence in performance produce a “budget” of total computational cost that can be distributed according to the specific question being investigated: E.g., more or less approximation in neutrino transport can be traded for better fidelity in the treatment of general-relativistic gravity. By having a set of codes wherein different levels of approximation to different pieces of physics are deployed at different times, TEAMS will be able to answer many of the questions posed in the earlier sections of this proposal without having to wait until *all* the capabilities are present at the highest level of fidelity (and the concomitant highest level of computational cost). In addition, code-to-code comparisons will be possible at any given epoch to facilitate quantifying the effects of any approximation.

As a community, we have been fortunate of late that development of exascale nuclear astrophysics codes has come to be supported by the DOE Office of Science through the Exascale Computing Project (ECP; which is supporting the ExaStar project to build CLASH as a seed project) and by the DOE National Nuclear Security Administration through the Advanced Technology Development and Mitigation (ATDM) subprogram of the DoE Advanced Simulation and Computing (ASC) Program (which is supporting ???). Though the development work described here is aligned with these code development projects, the ECP and ATDM developments are firmly aimed at producing codebases for the exascale. The breadth of code changes required for ECP and ATDM is larger, in most cases, than can be supported by SciDAC. However, the development timelines overlap to a large extent, and it is likely that ECP and ATDM developments can be brought to bear on TEAMS aims. The places we envision these overlaps occurring are identified in the following subsections, where we describe the development plans for each of the major TEAMS codes. In addition, we highlight below the nexuses of activity where we will draw directly on the effort of researchers in the SciDAC Institutes.

3.1 Code Development for Improved Stellar Modeling

As discussed in § 2.1, We will make a collaborative effort to produce realistic three-dimensional massive star progenitor models for core-collapse simulations. Our low Mach number hydrodynamics code, Maestro, will be the workhorse for this.

In the proposed period, we will extend Maestro to incorporate rotation, to allow for the modeling of

convective shell burning in rotating massive progenitors. This is a large change in the algorithm and code base, so a postdoc will be employed at SBU to lead this development.

Maestro uses a 1-d base state to represent hydrostatic equilibrium (HSE) and evolves the departures from this on a 3-d Cartesian mesh. For a rapidly rotating star, this will breakdown—there is not a unique 1-d HSE model. Instead we will need to do a 2-d base state, (r, θ) , that describes the structure. Generating the initial model itself will be challenging, but a relaxation technique can be used. The idea is that we initialize a spherical star, evolve a step, reaverage the Cartesian state to a two-dimensional axisymmetric base state, resetting this base state, and then evolve again. Once equilibrium is reached, we reset the calculation with this initial calculation and turn on the reactions, and evolve the system simulation forward in time.

The change to the code would be invasive—all the references to the base state currently reference a 1-d radial model. These will all have to now reference the axisymmetric base state. Further, Maestro incorporates a sophisticated mapping algorithm to go back and forth between the base state and full star state (see [8]). A new mapping algorithm will need to be developed for the axisymmetric case. Additionally, Maestro self-consistently evolves the base state (defining a base state velocity w_0 that obeys an elliptic equation) in response to large-scale heating. This is essential to evolving stars that expand significantly. An analogous base state expansion algorithm will need to be developed for the new axisymmetric case (in early years, we will run without base state evolution).

The second major complication is that the global constraint on the velocity field that is responsible for filtering soundwaves will no longer be able to be expressed as a simple divergence constraint. This means that standard projection techniques will not be able to evolve the system. This is similar to the issue discussed in [140] for the general EOS with strongly varying adiabatic index. The LBNL and SBU groups have been exploring alternate integration methods, including spectral deferred corrections (SDC), which is an iterative method to integrate a system of equations without resorting to the traditional operator splitting approach commonly employed in multiphysics codes. SDC has already been implemented in Maestro for smallscale flames and can eliminate the splitting error when coupling hydrodynamics and reactions [209]. We will explore a deferred correction strategy to evolve the new hyperbolic + elliptic system for rotating low Mach number hydrodynamics in Maestro. Such techniques have previously been applied to incompressible flow, so that can serve as a starting point [6]. An added benefit of this integration process is that it opens the door to fourth-order-in-time integration.

An additional effort that SBU will undertake is porting Maestro to the new AMReX library. The runtime of Maestro is dominated by geometric multigrid (both nodal and cell-centered variable-coefficient Poisson solves) that enforce the elliptic constraint on the velocity field. Maestro currently uses the pure-Fortran portion of BoxLib, but the new developments and optimization of multigrid in AMReX will be implemented in the C++ portion of the library. We will therefore port Maestro over to the C++ AMReX, most likely using a shim layer that interfaces the C++ via Fortran. This will ensure that Maestro continues to perform on future architectures and strengthens its coupling with Castro.

Our work plan is to start with a procedure for generating rotating equilibrium models in year 1, prototype the low Mach number rotating flow algorithm in year 2 (most likely in our teaching / prototype code pyro [300]), and have the implementation done in Maestro in year 3. In tandem, we will port Maestro to AMReX in year 1. Science runs described in the science section will begin in earnest in year 3 and continue through the remainder of the proposal period. All development will be done in a public github repository, so the entire community has access to the code changes throughout the development process.

As needed, we will use the capability to restart Maestro calculations in Castro the progenitor modeling, to evolve the star into the sonic regime, closer to the point of core-collapse. Castro already has the requisite physics for this, and its porting to AMReX under the ECP will put it in a strong position to efficiently model the late stages of this convection.

Collaborators at Southampton are implementing a general relativistic version of Maestro, that is also intended to be public. In the last years of the proposal, we will interact with that group and the transport members of this proposal to port neutrino transport (in some form) to Maestro, leveraging the open code developed by this project. Together with the rotation and GR, we will look into modeling proto-neutron star cooling with the new Maestro algorithm. We will collaborate with the members of TEAMS who have the existing scientific expertise.

3.2 Code Development for Improved Models of Core-Collapse Supernovae and Neutron Star Mergers

3.2.1 CHIMERA

CHIMERA is a mature code for simulating CCSN, as described briefly in §2. In the first two years of this proposal, we will work to markedly increase the strong scalability of the code, primarily through generalizing the domain decomposition and exposing additional parallelism in the transport and hydrodynamic solvers. By aggregating multiple radial rays per MPI rank, CHIMERA will be able to expose enough local work to effectively exploit GPU and many-core node architectures. We anticipate being able to use OpenMP 4.x directives to effectively target both these types of nodes, while having to retain only a limited number of architecture-specific pieces of code. In the third and subsequent years, we anticipate the highly parallel physics modules developed for CHIMERA will be transitioned into the CLASH framework. CLASH is being engineered from the outset to make this transition seamless. Importantly, the work aggregation done in the CHIMERA context during the first two years will be key to ensuring good node-level performance for these physics solvers when they transition to CLASH.

The post-doc to be hired at ORNL under TEAMS will be involved in these efforts as a secondary project. The ORNL post-docs' primary duties in the first years of TEAMS will be to run and analyze the CHIMERA models described in §2 to address a number of potential r-process sites, under the supervision of Hix and Messer. As part of this task, the ORNL post-doc will augment the existing CHIMERA analysis pipeline, which includes direct analysis of CHIMERA models, as well as calculation of nucleosynthetic, neutrino and gravitational wave signatures. In Year 1, the ORNL post-doc will also be charged with implementing the new EoS and neutrino opacity API being developed under TEAMS (see §3.3). This will enable CHIMERA simulations to run with enhanced opacities, as described in §2.

Once the CLASH transition has been completed, the ORNL postdoc will transition to running and analyzing models using CLASH. As part of this effort, the ORNL post-doc will port the CHIMERA analysis pipeline to utilize CLASH data formats, marking the birth of a CLASH analysis pipeline that will be open-sourced as part of the CLASH infrastructure.

3.2.2 CLASH

The current plan for CLASH is that early versions of the code will be available during the duration of this SciDAC project. As such, individual TEAMS investigators will be able to use either the individual “donor” codes (i.e. FLASH and CHIMERA) for the proposed science runs, or can use the CLASH infrastructure directly if the requisite capabilities are available in the new framework. Other development work for CLASH will be funded by the ECP ExaStar project, beyond the scope of this proposal.

3.2.3 FORNAX

The FORNAX code already incorporates all the necessary realism to address the core-collapse phenomenon, albeit with approximate relativity and employing the M1 two-moment closure approach to multi-dimensional transport. However, though it now contains various many-body corrections to neutrino-matter interactions, it has not benefited from a comprehensive upgrade to incorporate all the known and expected charged- and neutral-current structure-factor and many-body corrections to the various neutrino rates. In Year 1 of the award, co-I A. Burrows at Princeton, in collaboration with co-Is and senior investigators A. Steiner (Ten-

nessee), S. Reddy (INT), and L. Roberts (MSU), will incorporate self-consistent nuclear EoSs with a full suite of neutrino opacities that contain many-body corrections for all the rates. It has been shown that such corrections might be quite important both quantitatively and qualitatively in the outcome of CCSN (Burrows et al. 2016), but the consequences of all the various corrections collectively have never before been addressed. In Year 2, senior investigator D. Radice, co-I A. Burrows, and the Princeton postdoc will incorporate magnetic fields, using fully-staggered flux-constrained transport, into FORNAX. Including magnetic fields and their backreaction in a consistent fashion have been shown in the past to constitute only a 10%–20% additional computational overhead. The result will be a multi-physics magneto-radiation-hydrodynamic code that can address a broader set of the major issues that attend the supernova phenomenon and the birth of neutron stars and pulsars. Finally, in Year 4, co-I A. Burrows and the Princeton postdoc, in collaboration with senior researcher G. Fuller (San Diego), will embed a method into FORNAX to handle neutrino oscillations in an approximate fashion using an updated and corrected version of the methodology of Zhang & Burrows (2013). The latter derived a suite of generalized Boltzmann equations, based on the density-matrix formalism, that incorporates the physics of neutrino oscillations for two- and three-flavor oscillations, matter refraction, and self-refraction. The resulting equations are straightforward extensions of the classical transport equations that nevertheless contain the full physics of quantum oscillation phenomena. In this way, our broadened formalism provides a bridge between the familiar neutrino transport algorithms employed by supernova modelers and the more quantum-heavy approaches frequently employed to illuminate the various neutrino oscillation effects. In collaboration with G. Fuller, we will create a simplified moment version of this formalism, with suitable averaging over the stiff angular and energy dependences. Determining how to do the latter is why we have set the completion date of this development effort to Year 4 – we anticipate a longer development time to determine the physically reasonable, yet computationally feasible, way forward.

3.2.4 WHISKYTHC

We identify three key areas where new developments to WHISKYTHC are expected to yield significant scientific return: an improved neutrino radiation transport treatment, improved “low-density” ($\rho \lesssim 10^{11} \text{ g cm}^{-3}$) microphysics, and MHD. Before going into detail, we remark that **this development program will benefit in a substantial way from the inclusion of methods and technology already developed or under development in the CCSN simulation codes included in this proposal.**

Neutrino Radiation: Senior investigator D. Radice plans to couple the ZELMANIM1 code of co-IL Roberts (MSU) [243] with WHISKYTHC to perform the first energy-dependent radiation-hydrodynamics simulations of NS mergers. ZELMANIM1 uses a multi-dimensional moment scheme with an analytic closure for the 2nd and higher-order moments. ZELMANIM1 has already been used to perform CCSN studies [243]. However, before ZELMANIM1 can be used for mergers, it will have to be extended to be able to cope with some difficulties unique to the merger problem. Radice will extend and integrate ZELMANIM1 in three steps.

First, we will improve the robustness of the treatment for the advection of trapped radiation, which is currently not able to handle the large velocities typical of mergers. This will be accomplished using the discontinuous Galerkin method of [174], which automatically captures the correct diffusion limit in scattering dominated regions. We already implemented this scheme in CACTUS for special-relativistic radiation transport with very encouraging results [227] and we expect to be able to run the first merger simulations with ZELMANIM1 at the end of Year 1.

Second, we will address the well-known problem of moment schemes resulting in the formation of artificial “radiation shock” for colliding beams [78]. This is especially problematic for studies of the post-merger disk and of the SGRB engine. To overcome this problem, we will extend ZELMANIM1 to the partial-moments method [78]. This scheme is a hybrid between the traditional discrete ordinate method S_N [156, 214, 263] and the moment scheme. Implementing the partial-moment scheme will require substantial modifications to ZELMANIM1, but we conservatively estimate that we will be in a position to run the first simulations with

the partial-moment scheme by the end of Year 3.

Third, by the end of Year 5, we plan to implement a simplified treatment of neutrino oscillations based on an updated and improved version of the formalism by [262, 298]. We will extend ZELMANIM1 to solve for extra-diagonal components of the neutrino density matrix and include effective source terms to model the oscillations in a parametrized way. Clearly, this will only be an approximation to the full quantum-kinetic treatment. However, it should suffice to estimate the overall magnitude of the impact of neutrino oscillations on the nucleosynthetic yields of mergers.

Low-density EoS: We plan to improve the transition out of NSE with the integration between WHISKYTHC and the FORNAX and FLASH/CLASH codes. On the one hand, by the end of Year 2, we will be able to map post-merger configurations into FORNAX and FLASH/CLASH, which are well suited to simulate the remnant accretion disk and follow the ejecta to large distances. For these simulations, we will leverage the “low-density” physics modules (transition out of NSE and a simplified nuclear network to capture the heat production from the r-process) that will be implemented in them as a result of this project. By Year 3, we will also import these new physics modules from FORNAX into WHISKYTHC to be able to perform a series of long-term simulations in full-GR. These will be used to quantify the uncertainties due to the mapping procedure and the use of approximate GR in FORNAX and FLASH.

Magnetohydrodynamics: If this proposal is funded, we will extend WHISKYTHC to general-relativistic magnetohydrodynamics (GRMHD). Note that, although other GRMHD codes already exist and the first GRMHD NS merger simulations with idealized EoSs were performed several years ago [11], neutrino-radiation GRMHD simulations of NS mergers have never been performed. We anticipate the need for a number of technical improvements on the GRMHD technology to be able to perform reliable merger simulations with microphysics and magnetic fields. We will use a fully-staggered constrained-transport approach [297] with flux correction across refinement levels to robustly handle the propagation of magnetized outflows across refinement levels while conserving the magnetic flux. We will also implement the recently proposed correction algorithm by Martí [169] to ensure robustness and accuracy for highly-magnetized plasmas.

3.3 Code Development for Improving the Nuclear EoS and Opacities

The biggest current challenge facing the use of EoS tables and neutrino opacities in astrophysical simulations is a problem of uncertainty quantification: we cannot fully characterize the impact which nuclear physics uncertainties have on the astrophysical observables. There are two road blocks to uncertainty quantification: the paucity of fully available EoS tables and the difficulty of performing simulations multiple times, once for each possible EoS table.

PI Steiner will address this first roadblock by creating, beginning in year 1, a new set of EoS tables which partially quantifies the uncertainty in the nature of matter due to lingering uncertainties in the nature of dense neutron-rich matter. These EoS tables will be constrained by experimental data, using only models of matter near the nuclear saturation density which respect the wealth of experimental data available there, and astronomical observations, including the increasing availability of neutron star mass and radius observations. Importantly, these EoS tables will not be over-constrained, in the sense that constraints applied to one region of the (n_B, Y_e, T) parameter space will not be unnecessarily extrapolated to other regions of parameter space. This will result in a fully open-source EoS table code, released on github in year 5, from which any group will be able to **generate** EoS tables. Far from being solely useful by the collaboration alone, this EoS table code will be a critical tool for future simulations of supernovae and mergers (as opposed to the competing CompOSE effort funded by the European Space Foundation² which only provides code to parse tables and not to generate them).

Co-PIs Reddy and Roberts will provide three sets of neutrino interaction rates for charged and neutral

2. See <http://compose.obspm.fr>

current reactions in hot and dense uniform nuclear matter. During the first year, we will develop a complete opacity set that includes arbitrary matter degeneracy, relativity, changes in the dispersion relation of nucleons, and weak magnetism. In years two and three, we will include the effects of correlations through RPA with simple effective interactions to ensure that the response functions satisfy thermodynamic sum rules in the long wavelength limit. Spin and spin-isospin susceptibilities from ab initio microscopic calculations will guide the determination of effective interactions in these channels. Finally in year 4 we will include corrections due to quasi-particle life-time effects that tend to broaden the RPA response functions. These life-time corrections will also allow calculations of consistent nucleon-nucleon bremsstrahlung rates.

In non-uniform matter we will go beyond single nucleus approximation to describe coherent scattering of nuclei both at low density and in the high-density pasta phase where large deformed nuclei coexist with neutron-rich matter. The effects of ion-ion correlations, nuclear form factors, and screening due to electrons and ambient nucleons will all be included. Neutrino production in the semi-transparent regions due to nucleon-nucleon bremsstrahlung will be improved using the T-matrix approach to include non-perturbative effects in nucleon-nucleon scattering. These have been shown to provide large corrections (factors of 2-4) to the existing rates based on the one-pion exchange models. Finally, we shall use existing calculations of neutrino reactions on light nuclei such as deuteron, tritons and alpha particles to provide a complete set of opacities over the full range of ambient conditions encountered in simulations.

Along with the EoS code and tables, we will release the opacity codes produced by this proposal as open-source. A large fraction of the US computational nuclear astrophysics community is involved in this proposal. Therefore, we expect to be able to build a unified application programming interface and libraries for implementing microphysics in simulations in a format that fits the needs of the community.

3.4 Code Development and Uncertainty Quantification for Improving Nucleosynthesis Predictions

Code development for nucleosynthesis predictions fall into two categories, network calculations computed within the radiation hydrodynamical models and post-processing network calculations based on Lagrangian trajectories. For in situ networks, under the ExaStar ECP project, CLASH will inherit Flash's XNet implementation, allowing the use of arbitrary networks in CLASH simulations. These networks will be available for porting to other TEAMS codes to extend their existing capabilities. One focus for ExaStar is enhancing the efficiency of these nuclear reaction network on current and future architectures, lowering the effective cost of more realistic networks. However, though technically feasible, we anticipate that reaction networks with thousands of species within radiation hydrodynamical models will remain cost prohibitive. Thus TEAMS will rely on post-processing to predict the r-process and p-process production of our models.

Computation of post-processing nucleosynthesis requires thermodynamic trajectories for Lagrangian tracers and an infrastructure for computing the nucleosynthesis. CHIMERA and Flash have existing tracer particle capabilities and CLASH will inherit Flash's tracer particle capability under ExaStar, enhanced by work from the AMReX ECP co-design center on the underlying AMReX AMR scheme. Where the trajectories are not computed within the model, a back tracing of Lagrangian trajectories can be performed [see, e.g., 280] from periodically saved hydrodynamic data, trading the ability to pick points of interest at the end of the simulation for the higher accuracy of tracer tracking with every timestep. Thus Lagrangian trajectories will be available for all TEAMS simulations.

An important development effort under TEAMS, described in §3.2.1 is the augmentation of the CHIMERA analysis pipeline and its incorporation in the CLASH framework. This includes the capability to perform post-processing nucleosynthesis [95, 104] and estimate the uncertainties due to the incomplete state of the simulation [95]. Under TEAMS, these will become part of the open source CLASH framework, available to all TEAMS members and the broader community. The nuclear network used in the post-processing is the same employed within CHIMERA/Flash/CLASH simulations, thus improvements to its performance under ExaStar will also allow the post-processing to take advantage of current and future architectures. Under

What do the others plan?

TEAMS, drawing on the experience of Surman and Roberts, we will upgrade this capability to include the r-process.

Under TEAMS, we plan two orthogonal investigations of uncertainty quantification for the r-process. The first investigation, building on Surman’s prior efforts in this area [202], is to quantify the importance of the many nuclear physics inputs to the calculation of the r-process abundances from various scenarios. This includes attempts to place error bars on predicted r-process event yields due to all nuclear physics uncertainties as well as efforts to isolate the key individual pieces of nuclear data that most strongly impact final abundances in each case. The former is important for integrating the results of TEAMS simulations into larger modeling efforts such as galactic chemical evolution simulations; here yield uncertainties will be explored and quantified through systematic variations in the theoretical nuclear models used for the required data, as in [170, 266], and Monte Carlo variations of masses and reaction rates within their uncertainties, as in [201, 265]. The examinations of the impact of individual nuclear properties has direct relevance to the experiments that will be undertaken at FRIB and other radioactive Ion Beam facilities, helping to assign priority to the many potential measurements. The tools developed for this nuclear UQ will also enable investigations of the similar uncertainties for the p-process and other nucleosynthesis processes in CCSN and NSM.

The second investigation, to come in later years of TEAMS, is to compare the nucleosynthesis results for the various CCSN and NSM models computed under TEAMS. This will help to quantify the uncertainty on the astrophysical side of this nuclear astrophysics puzzle and establish the relative uncertainties between the microscopic and macroscopic part of the problem. In particular, the outcomes of r-process and ν p-process nucleosynthesis are quite sensitive to the initial ratio of neutrons to protons, which in turn can be strongly dependent on the neutrino physics of the event. Thus we anticipate a variation in nucleosynthetic output to result from different choices of neutrino opacities, inclusion of oscillations and general relativistic effects, etc. In addition, we will pursue investigations of the effects of stochasticity, introduced by turbulent fluid flow, as well as investigations of the impact of various different code approximations implemented in TEAMS collaboration codes.

4 Building a Stronger Nuclear Astrophysics Community

In assembling this proposal team, we have gathered the majority of the modelers of CCSN and NSM in the United States and a significant fraction of the researchers pursuing the physics of nuclear matter at the core of these events and the nucleosynthesis that links our origins to these events. As a result, this collaboration represents a unique opportunity to strengthen the computational nuclear astrophysics community in the United States, through a series of activities that will make our research more efficient, more accurate and more approachable. In the sub-sections below, we will detail these community building activities.

4.1 Providing EoS and opacity data

CCSNe and NSMs are sensitive to high density physics and observations of these events may be able constrain some properties of the nuclear EoS. As was discussed in section 3.3, the collaboration will develop code for calculating consistent EoSs and neutrino opacities. In year 2, Steiner will coordinate with Burrows, Couch, Fryer, Kasen, Mezzacappa, Roberts and Zingale (and their respective teams) to create an open source neutrino opacity Application Program Interface (API) which implements the newly generated neutrino opacities into an API similar to the NUOPAC library developed by Roberts and now on `bitbucket`. This API will be built on as a set of functions at global scope combined into a self-contained library which can be compiled in with any application as needed. This API will be critical for testing the sensitivity of models to input microphysics. Because the structure of the codes and the architectures of the machines they are run on vary widely in this proposal, the API will also contain a set of functions which generates neutrino opacity tables (given a specified EoS table) which can be used in lieu of the API. The API and the associated tables will form a set of **open source tools for including up-to-date microphysics in numerical simulations**.

Additionally, access to the same microphysical inputs is integral to cross validation and verification of

simulations of CCSNe and NSMs as described in section 4.3 below. Even in cases where different groups are using microphysics derived from the same results in the literature, the details of implementation within codes can vary between groups significantly. This can make it challenging to determine whether discrepancies between codes are due to differences in microphysical inputs or in the numerical treatment of macroscopic physics. The EoS and opacity framework provided by this collaboration will enable cross-code comparisons using consistent sets of microphysics.

4.2 Community Simulation Codes

A large number of the simulation codes used for the simulations proposed here are Open Source and the new developments proposed here will be made freely available to the community-at-large.

Maestro is freely available now and the implementation of the rotation and generation of the equilibrium initial models will all be done in its public git repo. All setup simulation files, initial conditions, etc., needed to perform the scientific simulations using Maestro will likewise be developed and continuously available in the public github repository. The development of rotation will enable studies of core and shell burning in stars throughout the mass spectrum, well beyond the science proposed here, benefiting the entire nuclear astrophysics community.

4.3 Community Code Comparisons and Testing

As quintessential multi-physics simulations, models of supernovae and NS mergers present severe challenges to validation and verification. While applicable unit tests exist for both purposes, there are no analytic solutions against which to verify for test problems that exercise the entire codebase nor are there accessible controlled experiments against which to test. This raises the importance of code-to-code comparisons. While such code-to-code comparisons are conducted, their frequency is less than it should be. A number of factors contribute to this infrequency, several of which can be addressed under this proposal. First, if the goal is to compare the impacts of the underlying numerical methods for (magneto)-hydrodynamics, radiation transport and thermonuclear kinetics, the use of identical microphysics is essential. As has been reaffirmed countless times, separate numerical implementations of the same analytic equations do not necessarily yield identical results, thus making any comparisons between codes more qualitative. The development of common opacities and EoS, as described in §3.3 and §4.1, make it much easier for tests to be run, not with just the same physics, but with the same implementation, focusing the comparison on the numerical methods and physical approximations upon which they rest. Second, with the building of the TEAMS collaboration to submit this proposal, we have assembled a critical mass of modelers of CCSN and NSM working in concert, that can jumpstart these efforts. We here commit to a robust set of code comparisons, both within the TEAMS collaboration, and among interested external simulation teams.

The CCSN problem is challenging and requires massive computational resources because of its intrinsically multi-physics, multi-scale, and multi-dimensional nature. The key physics components (e.g., [28, 44, 118, 120, 122]) are (A) multi-D (magneto)-hydrodynamics, (B) gravity, (C) the nuclear EOS and nuclear kinetics, (D) neutrino transport and neutrino-matter interactions. In CCSN simulation codes, (A)-(D) are coupled together via the simulation framework that provides the computational infrastructure (e.g., computational grids, adaptive mesh refinement [AMR], I/O, etc.) that makes multi-D, multi-scale CCSN simulations possible (e.g., [72?]).

Given the highly complex nature of the problem, it is to be expected that the outcome of CCSN simulations will have (at least some) dependence on the numerical methods that are used to implement the equations and on the numerical gridding used (e.g., resolution, coordinate systems). The only in-depth code comparison that exists in the literature is for spherically symmetric (1D) codes [158]. Since then, most simulation work has been in 2D and 3D, and has yielded quantitatively, and often qualitatively, different results even for the same initial conditions and input physics. **There is an urgent need for cross-verification and reproducibility for multi-D CCSN simulations.**

Need other references to codes frameworks here.

As part of this SciDAC project, we will accomplish this by...

4.4 Providing Data for the Broader Community

Beyond the scientific knowledge, and attendant publications, that are the primary fruit of the work proposed herein, there are a number of derivative data sets that are useful for the broader nuclear physics, particle physics and astrophysics communities. Examples include neutrino and gravitational wave signals that are used to test detector response (and, hopefully, someday be used to compare to observations), nucleosynthesis yields that are used in chemical evolution simulations, and results of test simulations for code comparisons (see, §4.3). While much of this data is already made available by the research team that generates it, tracking down such data from numerous sources can be challenging. Given the breadth of the TEAMS collaboration, we see an opportunity to **create a data clearinghouse** that would allow the communities of interest easier access to our collective results.

To this end, we propose the creation of the Nuclear Astrophysics Data Archive (NADA). NADA will host limited and reduced data sets aggregated from the TEAMS collaboration as well as serve as a “one stop shop” for links and references to other nuclear astrophysics data repositories. In the formation of NADA, we will collaborate closely with many nuclear astrophysics stakeholders, in particular the NSF Frontiers Joint Institute for Nuclear Astrophysics-Center for the Evolution of the Elements, of which many TEAMS investigators are members. Specifically, JINA-CEE is currently reviewing its methods and protocols for hosting and serving nuclear reaction rate data (REACLib). We will engage JINA-CEE to coordinate our efforts in establishing NADA in a compatible and cooperative way.

5 Milestones and Deliverables Timeline

- Year 1
 - Develop equilibrium rotating massive star models for the Maestro convection calculations.
 - Port Maestro to AMReX.
 - Conduct some of the first 3D core-collapse supernova simulations using FORNAX with full 3D transport (not using the RbR+ approximation)
 - Put self-consistent EOSes and opacities into FORNAX
 - Incorporate a nuclear network into FORNAX for nucleosynthetic post-processing
 - Couple WhiskyTHC with ZelmaniM1 to handle multi-group neutrino radiation transport
- Year 2
 - Prototype the low Mach number rotating version of Maestro (likely in pyro).
 - Add a B-field module to FORNAX
 - Develop a mapping procedure to perform long-term evolutions with WhiskyTHC using the transport package in FORNAX
 - Inaugurate the building into FORNAX and WhiskyTHC the capability to handle neutrino oscillations
- Year 3
 - Finish implementation of Maestro rotation code and do first science runs.
 - Perform 3D CCSN simulations using FORNAX with a full complement of nuclear many-body corrections and derive nucleosynthetic yields
 - Implement the partial moment method into WhiskyTHC
 - Perform merging neutron star simulations with the upgraded WhiskyTHC code
- Year 4
 - Production science runs of multiple massive star rotating convection models with Maestro.
 - Perform sunig FORNAX some of the first full-physics and self-consistent 3D radiation-hydrodynamic simulations with magnetic fields included
 - Add GRMHD to WhiskyTHC and transition it to a full nuclear network

- Conduct full GR merger simulations using these new WhiskyTHC capabilities and determine r-process yields
- Year 5
 - Add neutrino transport to Maestro and explore proto-neutron star cooling.
 - Conduct multiple 3D CCSN simulations using FORNAX with all the necessary physical realism and neutrino oscillations for a suite of old 1D and new TEAMS-generated 3D progenitor models
 - Conduct GR merger simulations using WhiskyTHC, with all the necessary physical realism and neutrino oscillations and determine r-process yields

6 Management Plan

The Investigator Team will be led by Hix, who will serve as the overall Project PI and will have responsibility for all aspects of the project. Messer will serve as the project’s Computation PI and will oversee all major computational developments. In addition, we will form an Oversight Committee made up of the Institutional PIs and Messer to manage the project and to make any key decisions affecting the project, e.g., with regard to any changing priorities and associated resource shifts deemed necessary as the project moves forward. Hix will host a monthly conference call with the Oversight Committee. In addition, the project will host collaboration meetings each year, generally at the end of the spring university semesters.

Table 1: TEAMS Investigator Team [requested support under this proposal from ONP¹ or ASCR²; Institutional PI[†]]

Investigator	Institution	Responsibilities (by Section)
Hix	ORNL	2.6