

Machine Learning-Driven Multiscale Modeling: Bridging the Scales with a Next-Generation Simulation Infrastructure

Helgi I. Ingólfsson, Harsh Bhatia, Fikret Aydin, Tomas Ooppelstrup, Cesar A. López, Liam G. Stanton, Timothy S. Carpenter, Sergio Wong, Francesco Di Natale, Xiaohua Zhang, Joseph Y. Moon, Christopher B. Stanley, Joseph R. Chavez, Kien Nguyen, Gautham Dharuman, Violetta Burns, Rebika Shrestha, Debanjan Goswami, Gulcin Gulden, Que N. Van, Arvind Ramanathan, Brian Van Essen, Nicolas W. Hengartner, Andrew G. Stephen, Thomas Turbyville, Peer-Timo Bremer, S. Gnanakaran, James N. Glosli, Felice C. Lightstone, Dwight V. Nissley,* and Frederick H. Streitz*

Cite This: *J. Chem. Theory Comput.* 2023, 19, 2658–2675

Read Online

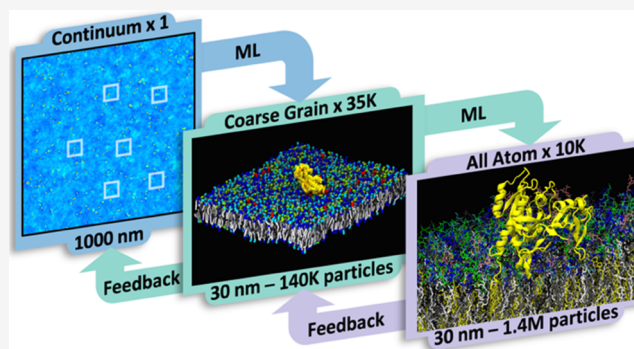
ACCESS |

Metrics & More

Article Recommendations

Supporting Information

ABSTRACT: Interdependence across time and length scales is common in biology, where atomic interactions can impact larger-scale phenomenon. Such dependence is especially true for a well-known cancer signaling pathway, where the membrane-bound RAS protein binds an effector protein called RAF. To capture the driving forces that bring RAS and RAF (represented as two domains, RBD and CRD) together on the plasma membrane, simulations with the ability to calculate atomic detail while having long time and large length-scales are needed. The Multiscale Machine-Learned Modeling Infrastructure (MuMMI) is able to resolve RAS/RAF protein-membrane interactions that identify specific lipid-protein fingerprints that enhance protein orientations viable for effector binding. MuMMI is a fully automated, ensemble-based multiscale approach connecting three resolution scales: (1) the coarsest scale is a continuum model able to simulate milliseconds of time for a $1 \mu\text{m}^2$ membrane, (2) the middle scale is a coarse-grained (CG) Martini bead model to explore protein-lipid interactions, and (3) the finest scale is an all-atom (AA) model capturing specific interactions between lipids and proteins. MuMMI dynamically couples adjacent scales in a pairwise manner using machine learning (ML). The dynamic coupling allows for better sampling of the refined scale from the adjacent coarse scale (forward) and on-the-fly feedback to improve the fidelity of the coarser scale from the adjacent refined scale (backward). MuMMI operates efficiently at any scale, from a few compute nodes to the largest supercomputers in the world, and is generalizable to simulate different systems. As computing resources continue to increase and multiscale methods continue to advance, fully automated multiscale simulations (like MuMMI) will be commonly used to address complex science questions.




1. INTRODUCTION

One of the fundamental challenges in computational modeling is to balance the trade-off between the length and time scales that need to be simulated and the corresponding computational costs. Relevant time scales are typically determined by the phenomena of interest, e.g., a protein's activation or a chemical reaction, whereas length scales are usually chosen based on the amount of detail needed, ranging anywhere from the first-principles representation of elementary particles to the entire cosmos. Unfortunately, despite the continuing growth in computational capabilities, many applications still face a gap of many orders of magnitude between the length and time scales needed to generate new insights and what is currently feasible, even on the largest supercomputers in the world.

One approach to overcome this limitation is to acknowledge that not every aspect of a simulation is equally interesting or important. Often, a majority of compute resources are spent waiting for rare or random events of interest and/or on resolving large regions of space to avoid boundary artifacts. Multiscale modeling and simulations offer an efficient way to limit unproductive simulations and maximize relevant computation. Using coarser length and/or time scales for

Received: October 12, 2022

Published: April 19, 2023



most of the computation and selectively resolving at the finer scale the events or regions of interest, simulations can be fully exploited for new and deeper scientific insights.^{1–4}

Multiscale models do not typically have a closed-form solution but, instead, work by combining models at different scales to address different aspects of the phenomena of interest while ensuring that the coupling across scales is consistent. Generically, a multiscale setup requires two fundamental building blocks: (a) the ability to define relevant scales and bridge them locally, i.e., to couple a coarse- and a fine-scale simulation (e.g., through projection, model reduction, or upscaling) and (b) an indicator to determine when to switch scales. A wide variety of solutions have been proposed in diverse areas of science, such as material science, biology, high-energy physics, and combustion.^{5–8} Multiscale may also be achieved by integrating sliding scales. Adaptive mesh refinement techniques are commonly used that selectively refine length and time scales in regions and periods of interest. However, these approaches are limited because they are built upon the assumption that the computational methods and discretization techniques remain the same, e.g., finite element representations. Different scales often require different discrete modeling approaches, such as continuum models using partial differential equations and particle-based molecular dynamics (MD) models. For such applications, significant challenges arise not only in managing the complex computational infrastructure but also, more importantly, in coupling these scales and selectively switching between them since a key challenge is to identify when and/or where to up/down scale. Many existing approaches focus on exploring known and predefined configuration spaces and use some form of predefined error bounds to decide when up/down scaling should be performed.^{6,8} Nevertheless, the challenge remains that predefined spaces are limited to known behavior and often fail to capture new hypotheses, whereas attempts to update these spaces or errors require laborious efforts, lengthening the time to solution.

Multiscale approaches have been of great benefit to computational chemistry and computational biology.^{9–14} Relevant to this work, classical all-atom (AA) MD is used extensively to study various types of membrane–protein interactions with great detail.^{15,16} AA simulations explicitly represent all atoms and capture details down to hydrogen bonding and vibrations. AA simulations are, however, limited in length and time scale due to their high computational cost, rendering the sampling of such large-scale behavior as protein reorganization due to lipid interaction hard to capture, and the exploration of lipid environment-dependent protein–protein dynamics near impossible.^{17,18} To address these challenges, various coarser models have been used, ranging from continuum models^{19,20} to coarse-grained (CG) models of different levels of granularity.^{21–25} CG models reduce the degrees of freedom, e.g., by combining atoms into effective “beads” and reducing allowable vibrations in bonds, often resulting in orders of magnitude reduction in computational cost.²¹ Of course, all CG models, due to their reduced degrees of freedom, are of reduced resolution and need to be used within their domain of applicability. CG models can generally be built using bottom-up and/or top-down approaches. The former entails utilizing the dynamics observed in fine-resolution simulations (such as AA) to parameterize the CG model, so that they can reproduce the collective behavior of high-fidelity simulations. Several bottom-up CG models exist

and have been employed to study a variety of membrane–protein problems.^{24–28} The top-down approaches, on the other hand, are parameterized to reproduce certain structural or thermodynamic properties of the system. These models have also proven successful in membrane–protein studies.^{29–31} The Martini force field³² is a particularly common CG model with a broad array of parameterized molecules.³³ Within Martini, the interaction potentials are optimized to match partition coefficients and molecule dynamics derived from AA simulations.

Various multiscale applications and methods have been used to address computational biology problems that require both large length and time scales as well as high resolution. A typical application couples two scales, passing information from one resolution to the other, and uses a fixed model at each scale. A common approach is some form of serial multiscale, e.g., simulate using general CG models and, for selected configurations of interest, convert them from CG into AA resolution and run AA simulations to verify and/or refine interactions and dynamics from the CG model.^{34–36} This form of multiscale may also utilize the bottom-up approach to generalize the observed dynamics at the AA scale and propagate them further in time and space.^{11,28} Such ideas have been employed to study various problems related to membrane–protein complexes, such as the mechanism of membrane remodeling induced by BAR proteins.^{24,37} In contrast to serial approaches, parallel multiscale involves modeling different components of the system at different resolutions that are coupled together and run concurrently. In this approach, a region or molecule(s) of interest in a complex system is described using a higher-resolution model, while the rest of the system is described by a computationally more efficient, low-resolution model. For instance, a protein can be modeled in atomistic detail within a CG membrane solvated by CG waters,^{38,39} or a specific area of interest can be modeled using AA resolution, surrounded by CG/AA annealing/mixing regions, inside larger CG simulations.^{40,41} Additionally, in the multiscale quantum mechanics/molecular mechanics (QM/MM) approach,^{42,43} specific regions of proteins are described at the electronic level to simulate the breaking and forming of chemical bonds, whereas the rest of the protein is modeled by molecular mechanics to allow inclusion of complex environmental effects and efficient sampling of the conformational space. QM/MM has been used to study such mechanisms as the ATP hydrolysis reaction in actin proteins⁴⁴ and investigation of the catalytic effect of enzymatic reactions, with the goal to calculate the activation free energies of these processes.^{45,46}

The multiscale approach taken here is ensemble-based. Our ensemble-based multiscale relies on a computationally efficient “macro” model to sample a relevant landscape of possible configurations exhaustively, from which a select subset is systematically explored using a higher-resolution “micro” model. Ensemble-based methods have long been employed in computational sciences,^{47–50} including in the context of multiscale simulations.^{51,52} Through parallel exploration of many simulations, ensemble methods facilitate the simulation of long-term behavior as well as rare events. More recently, ML-based approaches have been utilized to create large multiscale ensembles by sampling in learned spaces that represent phenomena of interest, such as protein conformation states^{53,54} and dynamics of SARS-CoV-2 spike protein.⁵⁵ Recent work on the Multiscale Machine-Learned Modeling

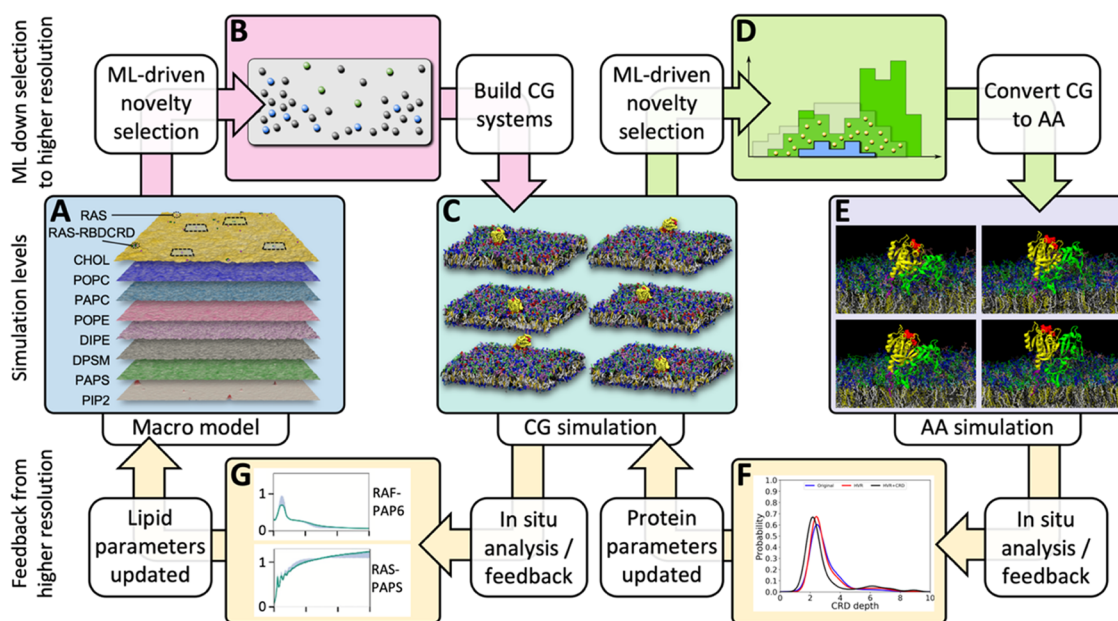


Figure 1. Three-scale MuMMI scheme. As the continuum macro model (A) evolves over time, ML-driven sampling is used to identify a diverse set of configurations (B) from the continuum model to spawn CG simulations (C) in a DL-based latent space. From the resulting CG simulations, candidate CG snapshots to convert to AA simulations are chosen in a three-dimensional parameter space (D). The subsequent AA simulations (E) are analyzed in real time to provide updated protein parameters (F) to the CG simulations. Concurrently, in situ analysis of the CG simulations refines the protein–lipid interactions (G) for the macro model.

Infrastructure (MuMMI)⁵⁶ demonstrated such ideas at a massive scale by constructing a ML-driven ensemble of several tens of thousands of micro simulations. MuMMI enables the query of relevant hypotheses through a machine-learned latent space,⁵⁷ allowing automatic and optimal sampling of this space to maximize new insights. Furthermore, MuMMI also leverages this automated ML-driven approach to facilitate an on-the-fly feedback loop to improve the fidelity of the running macro model through in situ analysis of micro model simulations. MuMMI was previously limited to connecting two scales. This paper presents the first three-scale multiscale approach to explore the RAS/RAF/MAPK signaling pathway that is frequently dysregulated in cancer.⁵⁸ In particular, our work generalizes and significantly extends the previous MuMMI infrastructure.⁵⁶ First, in addition to the continuum and CG MD scales of the original approach, we introduce a third scale—AA MD simulations—to explore protein–lipid-dependent structural rearrangements. The AA scale provides more details about the phenomena explored by CG and improves the parameterization at CG scale and, transitively, the continuum scale, resulting in a complete feedback loop that spans three scales. Second, the new framework extends hypothesis-driven refinement to the AA scale. Rather than selecting CG-to-AA starting structures based on a fixed criterion or explicit triggers, we expand MuMMI’s sampling capabilities to dynamically explore and sample from a parameter space of interest (here the tilt, rotation, and depth of the proteins with respect to the membrane). Finally, the new framework introduces the RAS-binding domain (RBD) and cysteine-rich domain (CRD) of RAF1 to the system and enables the exploration of RAS-RBDCRD complexes as well as combinations of multiple RAS and multiple complexes.

The RAS/RAF/MAPK signaling pathway is one of the cell’s main pathways for promoting cell division, differentiation, and growth. When dysregulated, this pathway is a prominent driver

in cancer biology.^{58,59} Mutations of RAS proteins alone appear in a larger fraction of all human cancers diagnosed in the US,^{60–62} rendering RAS an attractive therapeutic target. Recently, the FDA approved new drugs that target a select RAS mutation (the codon 12 glycine to cysteine mutation most commonly found in lung cancer), giving new hope to cancer patients, but there is still an unmet need for new inhibitors and approaches targeting other RAS mutations.^{63–65} RAS proteins localize on the inner leaflet of the cells’ plasma membrane (PM). To initiate signaling, RAS proteins are loaded with GTP, resulting in an active conformation, which primes them to bind downstream effector proteins, particularly RAF. The RAF kinase domain dimerizes after two RAS-RAF complexes are formed on the membrane, resulting in the activation of kinase activity, which propagates downstream signaling. However, the exact mechanisms by which RAS and RAF colocalize on the membrane and undergo activation are unknown at the molecular level, and insights are needed to understand signal activation.

RAS proteins, specifically of the common splice variant KRAS4b used here, colocalize and form nanoclusters in the presence of anionic lipids, and the nanoclustering has been shown to be very sensitive to changes in lipid composition.^{66,67} Using MD simulations, numerous RAS–RAS interfaces have been described^{68–71} as well as RAS preference for anionic lipids.^{66,72,73} Previously, a two-resolution multiscale simulation using MuMMI⁵⁶ was performed to explore RAS dynamics on an 8-lipid type PM mimetic.⁷⁴ This simulation and supporting experiments⁵⁶ revealed that lipid sub-compositions preferentially associated with RAS demonstrated the existence of detailed lipid fingerprints²⁹ associated with different conformational states of RAS and indicated the promiscuous nature of RAS–RAS G-domain multimerization. Depending on RAS membrane orientation, the RBD domain of RAF can bind RAS, and the RAF CRD domain can reorient and interact with

the membrane.⁷⁵ Resolving the lipid dependence of RAS with and without the RAF RBDCRD domains as well as their respective protein–protein interactions requires simulations spanning multiple resolutions: (1) capturing different possible lipid environments and protein conformations/stoichiometry requires sampling at cellular time and length scales; (2) capturing lipid–protein interactions, protein conformation dynamics, and protein–protein interactions under such a variety of local environments requires tens of thousands of multi-microsecond MD simulations, out of reach for AA but attainable by CG; and (3) detailed interactions of the RAS and RAF proteins with the membrane alters the local environment in a way that can change local secondary structures⁷⁶—especially the RAF1 CRD when inserted into the membrane—requiring specific AA details to resolve, as this is beyond the capabilities of most CG force fields. Together, these requirements necessitate expanding the MuMMI framework to three scales.

Given the immense value of such multiscale simulations, we expanded the MuMMI framework to support an AA scale and perform a large multiscale simulation of RAS-RBDCRD interactions on the PM, where RBD and CRD are two domains of the RAF1 protein. This paper presents a holistic description of significant extensions to MuMMI, referred to herein as the three-scale MuMMI. Results specific to RAS and RAF dynamics are forthcoming in separate manuscripts—here, we focus on demonstrating MuMMI’s capability of coupling between scales, the benefits of coupling multiple scales, and the benefits of running feedback coupled ensemble multiscale campaign. The three-scale MuMMI illustrated here greatly extends the original two-scale MuMMI.^{56,57,77} The underlying computational workflow was generalized to support multiple simultaneous scales and to greatly extend scalability and fault tolerance.⁷⁸ The AA scale was added with reliable CG-to-AA transformations,⁷⁶ CG parameters were updated through additional preliminary simulations,⁷⁵ and support for additional proteins components were added in CG (SI Section 1.3). A higher fidelity and faster continuum model⁷⁹ used updated parameterization (SI Section 1.1), and a new deep learning-based sampling framework was used to analyze and sample macro configurations with several types of protein constellations and varying lipid environments.⁸⁰

2. METHODS

2.1. Multiscale Machine-Learned Modeling Infrastructure (MuMMI). To deliver a massive multiscale simulation capable of autonomous execution through ML-driven bidirectional coupling (see Figure 1), a new computational framework is needed that not only allows executing the different types of simulations autonomously and dynamically but also evolves them jointly as a coupled multiscale simulation model that learns and improves with time. Such sophisticated multiscale coupling also necessitates an equally competent infrastructure to manage its software and hardware elements with minimal manual intervention while achieving unprecedented computational scales on modern heterogeneous computing technologies. MuMMI is the key enabling technology behind our coupled multiscale simulations. Broadly, MuMMI consists of three main parts: (1) simulation models at all relevant scales with the associated protocols for scale transformations, (2) coupling across scales, accomplished with ML-driven sampling and in situ feedback, and (3) a

dynamic and fault-tolerant workflow that orchestrates all of the individual simulations at scale.

2.1.1. Simulation Models at Three Relevant Scales and Protocols to Transform Representations across Models. At the coarsest scale, MuMMI employs a macro scale continuum model⁷⁹ to explore the long- and large-scale interaction behavior between lipids and RAS and RAF (RBD and CRD domains) proteins. The continuum model represents the lipid bilayer as density fields for each lipid of interest.⁷⁹ To evolve single-particle representations of the proteins over space and time for a $1 \times 1 \mu\text{m}^2$ PM, the model uses dynamic density functional theory (DDFT)⁸¹ with an implicit integration technique. Two types of proteins are represented, each in three possible states that differ in orientation relative to the PM. Although the continuum model does not provide the required fidelity for specific lipid or protein interactions, it does offer sufficient detail to capture the overall membrane dynamics—sampling local fluctuations in lipid densities as well as overall lipid preferences of the proteins for each protein state. We have developed a high performance and parallel implementation of this continuum model.⁷⁹ In the macro simulation presented here, we achieved a simulation performance of 0.96 ms per day using 2400 CPU cores (IBM POWER9). The details of this simulation tool will be presented in a forthcoming publication.

At a finer resolution, the lipid bilayer and the proteins are represented at the CG MD scale and evolved over time using the Martini force field.^{32,82,83} Since MD simulations are computationally expensive and unable to simulate the entire system, only local regions of the PM are considered at CG resolution. Specifically, $30 \times 30 \text{ nm}^2$ “patches” of the PM sections are cut out from the continuum membrane representation and, if deemed important, are promoted to the CG scale. We primarily consider patches centered around RAS proteins and RAS-RBDCRD complexes; a small set of patches without any proteins are also selected as a control set. A custom tool, Createsims (SI Section S1.3), was developed that uses a modified version of insane⁸⁴ to convert the continuum representation into a particle-based one and then GROMACS⁸⁵ to relax and equilibrate the CG representation. Next, CG simulations are performed using a specially designed, CUDA-enabled version of the ddcMD simulation engine.⁸⁶ These CG simulations—each containing around 140K particles—achieved $\sim 1.04 \mu\text{s}$ of MD trajectories per day utilizing one NVIDIA V100 GPU.⁸⁶ During execution, MuMMI orchestrated in situ analysis of CG simulations, providing input that is key to dynamic coupling through feedback from the CG to the continuum model (SI Section S1.4).

The finest scale currently supported by MuMMI is AA resolution⁷⁶ using the CHARMM36 force field.^{87,88} Snapshots from all CG simulations are evaluated (sampled every 2 ns), and selected snapshots are extracted and converted to the AA representation using a modified backward tool,^{34,76} followed by energy minimization using GROMACS,⁸⁵ and format conversion using ParmEd.⁸⁹ A description of the MuMMI AA modules and a complete protocol of this backmapping procedure are given by López et al.⁷⁶ The $30 \times 30 \text{ nm}^2$ AA representations are then evolved on a single-GPU using AMBER18.⁹⁰ AA simulations comprise, on average, $\sim 1.6 \text{ M}$ atoms and run at $\sim 14 \text{ ns}$ per day per NVIDIA V100 GPU.

2.1.2. Coupling across Scales: ML-Driven Sampling and In Situ Feedback. Given the three simulation models discussed

above, MuMMI couples adjacent scales in a pairwise manner using machine learning (ML). Here, the key challenge that MuMMI addresses is the dynamic nature of the coupling—both forward (sampling) and backward (feedback).

To maximize the opportunity for discovery, MuMMI aims to create a diverse set of CG and AA simulations along specified dimensions of interest and prevent redundant simulations, i.e., configurations similar to the ones previously simulated. Comparing continuum patches directly in the space of lipids and proteins is challenging since no direct comparative metrics exist that synergize similarities in lipids with similarities in proteins; instead, MuMMI uses ML to evaluate similarities among patches. A pretrained ML model is used to evaluate patches generated from the continuum simulation and order them by their diversity importance.⁵⁷ When computational resources are available, the top candidates are selected by a dynamic-sampling framework,⁵⁷ and corresponding CG simulations are spawned. MuMMI offers a flexible sampling framework that allows opportunities to customize and control the extent of diversity sampling and is agnostic to the type of ML used. For example, MuMMI also offers controlled or limited diversity and can utilize autoencoder-type deep learning, neural network-based metric learning, or traditional statistical measures. For the current simulation campaign, sampling for CG uses a 9-D metric space generated by a deep neural network.⁸⁰ Continuum patches are classified with respect to proteins (no proteins, one RAS protein, one RAS-RBDCRD complex, and all others) and ranked by diversity within each class, creating a diverse set of spatially resolved lipid compositions for the given protein configurations. Predefined proportions are used (10, 10, 20, and 60%, respectively) to obtain statistically significant sets of target simulations in each class.

To select an appropriately diverse set of initial configurations for AA simulations, CG frames (sampled every 2 ns) from all running and concluded simulations with a single RAS-RBDCRD are evaluated. Evaluation of CG frames is performed directly in a three-dimensional (3D) space representing the rotation and tilt of RAS and the depth of RAF CRD with respect to the membrane. Specifically, a 3D histogram is computed and dynamically evolved to capture the variability observed in these order parameters, from which the configuration's degree of diversity can be ascertained. For sampling, a tunable parameter is introduced that offers a balance between randomness and diversity, facilitating the consideration of both the dominant modes and a truly diverse distribution. For the current application, the method was configured to sample an interpolated distribution that represents a 20–80% balance of diversity and randomness, resulting in an AA ensemble that captures dominant modes while efficiently exploring the diversity in configurations.

A unique capability of the MuMMI multiscale approach is the ability to use the information obtained from a fine-scale simulation to update and improve the coarse-scale simulations while the multiscale campaign is running. This automatic feedback has two components. First, *in situ* analysis of all running simulations is performed to compute relevant properties, which are continually aggregated across all simulations. Second, appropriate weighting is used to debias the statistics of the effects of diversity sampling before the update. Since redundancies are not simulated, more common configurations are weighted correspondingly higher.⁵⁷ For CG, protein–lipid radial density functions (RDFs) are computed

for each simulation frame, and the aggregated RDFs are used to update the parameterization of the continuum model. The growing ensemble of CG simulations explored during a MuMMI campaign readily outpaces the limited, initial training set and covers a more diverse set of lipid compositions and configurations, steadily updating the continuum model protein–lipid interactions during the multiscale campaign. In the case of AA, protein secondary structures from the AA simulations are analyzed and aggregated to identify shifts in secondary structures, which are then used to improve the CG protein models. The AA simulation ensemble is constructed to sample different RAS-RBDCRD membrane orientations and CRD membrane insertion depths, leading to better sampling of relevant environments.

2.1.3. Dynamic Workflow and Orchestration of Simulation Ensembles. The award-winning MuMMI workflow technology^{77,78} was developed to steer the multiscale simulation campaign using ML and to simultaneously coordinate a large number of different types of simulations. The MuMMI workflow supports a fully automated and dynamic selection of simulations to be run based on real-time evaluation using ML. Furthermore, it supports a complex software stack and can accommodate diverse hardware requirements of the different kinds of simulations and use knowledge of those differences to better utilize resources on heterogeneous machines. The MuMMI workflow can seamlessly run at different computational scales using various numbers of compute nodes and has been tested in runs using 5–4000 nodes. MuMMI runs can be scaled up or down as needed for a given simulation campaign and available resources, e.g., restoring from a 100-node job to start a 4000-node one or vice versa. The DDFT-based macro model runs on CPU cores only and requires little compute resource. The CG and AA simulations are the most computationally intensive calculations in the framework. Since the CG and AA simulations are highly parallelized, and each simulation is designed to run on one GPU to maximize GPU utilization, scaling up or down the number of nodes does not affect the performance of the individual CG and AA simulations. The workflow can be deployed on most GPU compute clusters without sacrificing performance due to its flexible usage of compute nodes. Scaled across a large number of nodes, MuMMI can orchestrate several thousands of simulations simultaneously and manages the creation, analysis, and storage of several hundred TBs of data using an advanced file system and database technologies. This is made possible by leveraging modern heterogeneous computing machines such as Sierra and Summit, at the time, two of the three most powerful supercomputers on the planet.

2.2. Modeling and Analysis Methods. The models and parameters for the different sub-simulations of the RAS-RBDCRD multiscale campaign are documented as follows: continuum model in SI Section S1.1, CG as in Ingólfsson et al.⁵⁶ and listed in SI Section S1.3 with parameter updates described in Nguyen et al.⁷⁵ and AA are described in Nguyen et al.⁷⁴ and Lopez et al.^{75,76} Additional analysis used in this manuscript is listed here. Note that in some cases, the lipid names in the 8-lipid type PM mimetic⁷⁴ differ between the CG lipid and the AA lipid selected to represent that CG lipid, as noted earlier.⁷⁶ Here, we use the common CG names: 1-palmitoyl-2-oleoyl-*sn*-glycero-3-phosphocholine (POPC), 1-palmitoyl-2-arachidonoyl-*sn*-glycero-3-phosphocholine (PAPC), 1-palmitoyl-2-oleoyl-*sn*-glycero-3-phosphoethanol-

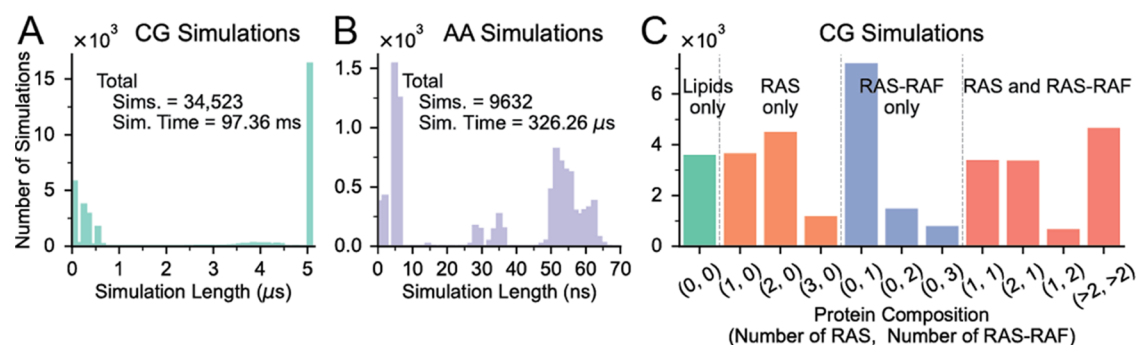


Figure 3. RAS-RBDCRD multiscale simulation campaign. MuMMI enabled a large three-scale simulation to resolve RAS and RAS-RBDCRD dynamics on an eight-lipid component plasma membrane mimetic. A cellular-scale continuum simulation ($1 \times 1 \text{ mm}^2$) with 300 proteins was simulated for over 20.5 ms. (A) From the running continuum simulation, 34,523 CG simulations were selected, set up, and simulated up to 5 ms. (B) From the running CG simulations, 9623 snapshots were selected, set up, and run at the AA level for up to 50 ns. (C) At the CG scale, the patches were selected based on an ML-enabled diversity sampling, tuned to sample a diversity of protein configurations while keeping a given ratio of simulation in each bin ($\times 4$ colors in figure).

distributions show several peaks highlighting the simulations that were started more recently than others. The target durations of CG and AA simulations were $5 \mu\text{s}$ and 50 ns , respectively. A large number of simulations were set up and run for a brief period at the end of the campaign to demonstrate full-machine execution of MuMMI on Summit (4000 nodes, 24,000 GPUs), as described with the MuMMI workflow.⁷⁸ The continuum simulation evolved 300 RAS and ~ 150 RBDCRD “beads” on the PM, exploring a wide variety of protein compositions, including high-order aggregates. The campaign explored many different types of protein compositions at the CG scale (Figure 3C and SI Table S1). For AA, only the compositions with a single RAS-RBDCRD were selected.

3.2. Bridging Continuum and CG Resolutions. The continuum-to-CG coupling in MuMMI provides great benefits for overall system modeling at both resolutions. Using a realistic large and long time scale continuum model allows for the exploration of relevant diversity of protein–lipid constellations and types of proteins in different conformational states and their relative orientation with a diverse sampling of realistic local lipid compositions and arrangements. MuMMI uses ML-guided diversity selection to sample the continuum simulation while it is running, populating an ever-increasing ensemble of CG-MD simulations. The continuum-to-CG coupling enables the creation of the relevant CG simulation ensemble, enabling CG-level analytics over an extremely broad range of possible configurations and allowing for conditional analysis as well as population analysis by utilizing the “weights” of the simulations. The CG simulation ensemble also provides benefits for the continuum resolution, both offline to verify the accuracy of the macro model and online, where on-the-fly analysis of the CG simulations is used to update the running macro model.

3.2.1. Dynamic Sampling of Diverse Continuum-Scale Patches. As mentioned earlier, our RAS-RAF multiscale simulation campaign comprises a single continuum simulation that generates a continuous and steady stream of patches. Over the course of the simulation, 6,828,831 patches were explored by the continuum simulation (333 patches for each of the 20,507 continuum snapshots), exhibiting a wide variety of lipid compositions as well as protein configurations.

To select a diverse subset of these patches, a deep learning (DL)-based metric-learning approach was used.⁸⁰ Specifically,

a custom neural network was designed and trained to compare patches. The input data was a $37 \times 37 \times 14$ image representing the 14 lipid concentrations discretized within each patch in a 37×37 grid and 9 labels describing different protein conformations representing different states of RAS proteins and RAS-RBDCRD complexes. The DL model ingests this complex and multimodal inputs and projects them into a 9-dimensional latent space that allows comparison of the patches using Euclidean distances. Using this pretrained DL model through online inference, MuMMI can evaluate all patches in this latent space and select the most diverse ones (defined as most distant) in real time. Through this approach, our multiscale simulation selected 34,523 patches (out of over 6.8 million) that were deemed to be most diverse and spawned the corresponding CG simulations. Figure 4 illustrates this diversity in sampling by comparing the proportions of the 14 lipid species for all patches with those for only the patches selected for CG simulations using parallel coordinates (Figure 4A,B), where opacity is used to visualize the distribution along the corresponding axes—darker colors indicate more patches at a given lipid composition. For reference, the average 8-lipid type PM composition⁷⁴ is shown as a black line through these parallel coordinates’ axes. The plots also show the most common and two rare compositions (green and red, respectively) that were explored by the macro simulation. We note that the trends shown by the two red lines (rare compositions) are starkly different both from the black (8-lipid type membrane composition) and the green (common composition) as well as from each other, indicating the large differences in the membrane compositions explored by the macro model. Through diversity sampling, MuMMI focuses on capturing such rare compositions as well as preventing the selection of redundant configurations.

Whereas Figure 4A,B shows the sampling in the space of lipid compositions, the sampling also accounts for protein compositions and is performed in the 9-dimensional latent space generated through DL. Figure 4C illustrates the sampling directly in this latent space by visualizing the marginal distributions in each latent dimension. The figure compares the distribution of all patches explored by the continuum model (pink-shaded region) against those selected for CG (black line) and demonstrates that MuMMI’s sampling suppresses the modes (common configurations) and provides a wide coverage (infrequent configurations).




Figure 4. Illustration of the diverse sampling of continuum configurations in MuMMI. (A, B) Using parallel coordinates plot, the lipid proportions (per leaflet) of all patches explored by the macro model were compared with those selected for CG. Each horizontal axis shows a particular lipid, and opacity represents the distribution along the corresponding axes (darker color indicates a larger number of patches). Notice that darker shades of orange have a wider span along horizontal axes in B as compared to A (note as specific examples inner CHOL and PIP2), highlighting that a wider range of concentrations is sampled using ML. The figures also visualize the 8-lipid type PM composition⁷⁴ (black line) as well as the most common, i.e., the one with the highest weight, (green line) and two very uncommon, i.e., two of the many patches with a weight of 1, (red lines) patches found by the sampling algorithm. (D) Visualization of these patches (top row: most common; middle and bottom rows: uncommon). The patches with a weight of 1 (bottom two rows) show more extreme examples: the bottom has many proteins, so the sampling captures high aggregation; and the middle shows high CHOL outer and high PIP2 even for a 1-complex system, and hence, is found to be diverse. RAF continuum model bead represents the CRD domain of RAF1 in the RAS-RBDCRD complex. (C) DL latent space to compare the patches explored (pink) and the patches sampled (black) demonstrates that the sampling prevents redundancy (suppresses the modes) and captures uncommon compositions (wider and flatter distributions).

Finally, Figure 4D visualizes three continuum patches (the green and red lines from Figure 4A,B) to further highlight the spatial variability captured by MuMMI's sampling. Whereas the top row (the most common patch) contains a single RAS-RAF complex with a small lipid fingerprint around the proteins, the middle and bottom rows (two rare patches) show highly specific signatures of proteins that the sampling deems diverse enough to be simulated at the CG scales.

3.2.2. Shorter Equilibration of Continuum Patches at CG Resolution. An additional benefit of the continuum-to-CG coupling is the reduction in equilibrium time needed to capture each proposed continuum configuration at the CG resolution by informing the initial placement of proteins and lipids. MuMMI creates each CG simulation from a continuum patch using the Createsims module (Figure 5A); for method

details, see SI Section S1.3. Proteins in relevant conformational states are sampled from input libraries and placed according to the continuum specifications, and the lipid density fields are probabilistically sampled to place lipid molecules. Figure 5B,C shows representative continuum-to-CG lipid placement of inner leaflet lipids for a patch with two RAS-RBDCRD (center) and on RAS-only (top) in B and one RAS-only (center) and one RAS-RBDCRD (bottom right) in C, respectively; see SI Figure S4 for additional details. The continuous lipid densities show the high lipid resolution in the continuum model resolved at a $0.81 \times 0.81 \text{ nm}^2$ grid size. Examples of typical lipid dynamics can also be seen such as in plane lipid lateral heterogeneity, seen in local accumulation of more saturated lipids and cholesterol vs more unsaturated lipids, as well as the protein-local lipid fingerprints, most

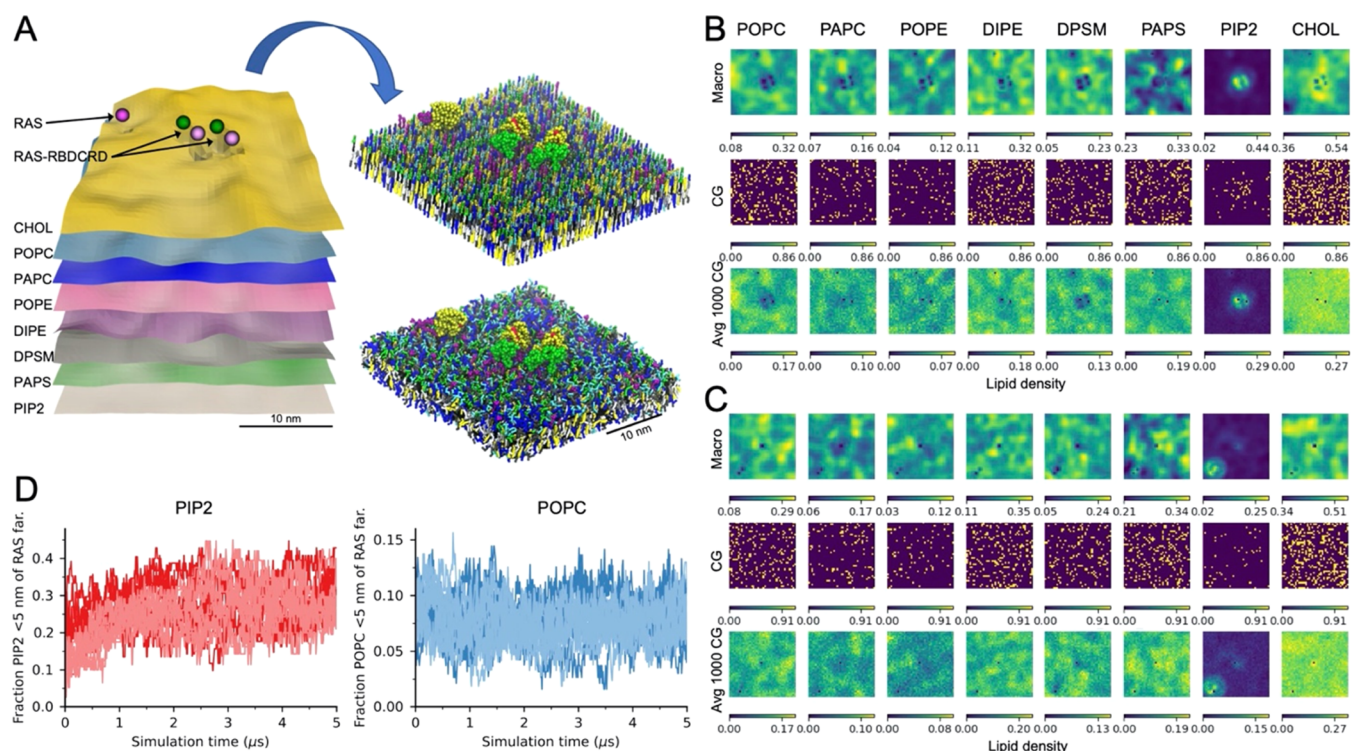


Figure 5. Converting a continuum patch to a CG simulation. (A) Scheme demonstrating how a $30 \times 30 \text{ nm}^2$ continuum patch is converted into a CG simulation using the Createsims module. (B, C) Inner leaflet lipid densities from the continuum simulation, a single CG lipid placement, and an average of 1000 CG lipid placements shown for a patch with two RAS-RBDCRD and one RAS-only (B) and one RAS-only and one RAS-RBDCRD (C). Note that continuum macro model densities are in lipids per μm^2 while the CG setup are counts and averaged counts per cell. (D) Fraction of inner leaflet PIP2 (left) or POPC (right) lipids that are within 5 nm of the RAS farnesyl lipid anchor. The darker color lines show simulations started with the MuMMI continuum-to-CG build routine, while the lighter color lines are the same continuum patches build with a modified routine that randomizes the location of the lipids in each leaflet. Simulations from eight continuum patches are shown, all with one protein, four with RAS-only, and four with RAS-RBDCRD, and starting from different protein conformational states.

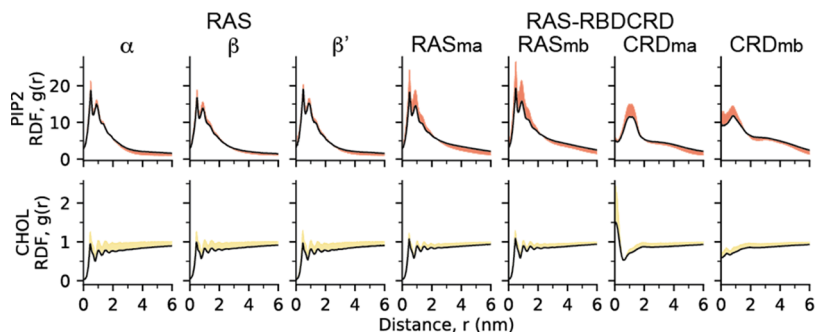


Figure 6. MuMMI framework utilizes in situ analysis and dynamic ML-based sampling to facilitate an on-the-fly feedback loop. Given an initial estimate of protein–lipid radial distribution functions (RDFs), shown as colored lines, the feedback gradually improves the estimate (shaded regions) to achieve more accurate RDFs (black).

noticeably is the increased PIP2 lipid density around the proteins. The CG lipid placement is also shown both for a single instance and the average of 1000 CG placements. The single instance is a probabilistic sampling of the overall lipid densities; the main trends can still be clearly seen, and by averaging multiple samplings, the macro densities are recreated (Figure 5B,C and SI Figure S4). Limited empirical assessment indicates that the total lipid concentration per lipid type is typically within 0.1–0.3% of that expected from the macro model and of those checked always within 1%. By spatially placing the lipids in the bilayer as dictated by the continuum model, the CG simulation needs less simulation time to equilibrate the different lipid features for the given patch. The

time saved will vary, depending on the lipid feature of interest. To demonstrate the change in equilibrium time, eight representative CG simulations (four with one RAS-only and four with one RAS-RBDCRD) were selected, set up, and simulated again but varied the initial lipid placement by randomizing the location of all of the lipids in each leaflet. Figure 5D shows the relative accumulation of inner leaflet PIP2 and POPC lipids close to the proteins as the simulation progress. Note that 0 time in Figure 5D is after the initial 6.3 ns of pre-equilibrium done by Createsims as part of the patch build. PIP2 lipids are strongly associated with the RAS-only and RAS-RBDCRD proteins, and simulations starting from continuum macro model placement (dark lines in Figure 5D)

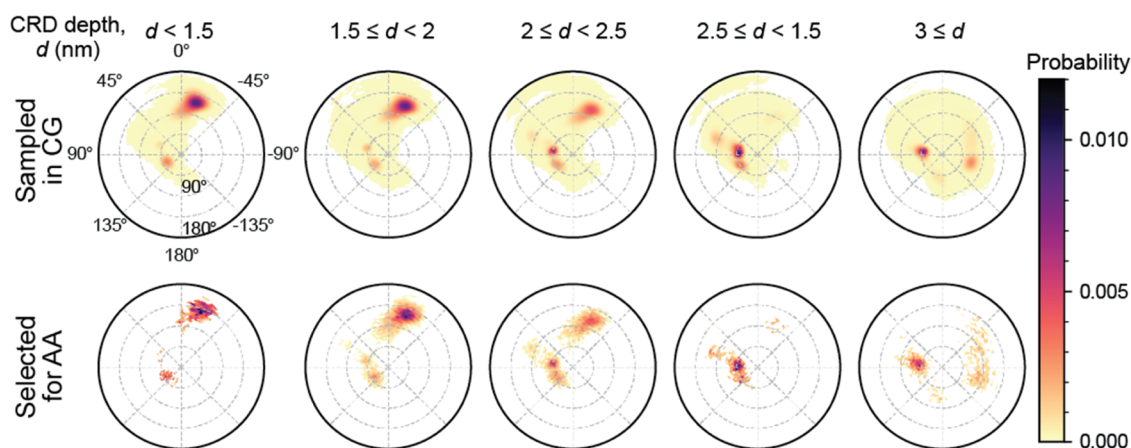


Figure 7. Illustration of the diversity sampling of CG snapshots in MuMMI; top row shows all CG snapshots, bottom row shows those selected for AA simulations. CG snapshots were evaluated based on the depth of CRD membrane insertion and the tilt and rotation angles that describe the orientation of RAS G-domain. Available snapshots were distributed into bins of a 3D histogram—5 bins for depth and 302 equal-area bins for rotation and depth for each value of depth. These histograms were updated dynamically to accommodate new snapshots from running CG simulations and were used to select snapshots from bins that allowed maintaining a prescribed balance between diversity and randomness.

start closer to the preferred protein–lipid placement than starting with random placement (light lines in Figure 5D). However, POPC lipids that have a much weaker protein preference and higher inner leaflet concentration show little change in accumulation.

3.2.3. In Situ Feedback from CG-to-Continuum. One of the key inputs for the parameterization of the continuum model is the collection of radial distribution functions (RDFs) between lipids and proteins. The RDFs are gathered from an initial set of CG simulations conducted prior to the presented multiscale simulation. However, the nature of this campaign provides an opportunity to improve the continuum model by exploiting a substantially larger set of CG trajectories. Although one could utilize the RDFs computed through this larger set of CG to parameterize a future continuum model more accurately, such post-hoc reparameterization would delay new insights that could be drawn from an updated continuum model. This multiscale simulation campaign leverages these simulations through an on-the-fly feedback loop in which more accurate estimates of RDFs are computed and fed into the continuum model, thereby updating it while the multiscale simulation campaign is running. The feedback loop is executed every 10 minutes and computationally optimized to finish within this time frame. Each feedback loop updates the RDFs using all new CG snapshots created since the previous loop. Figure 6 shows the final RDFs obtained after the simulation campaign with respect to the initial RDF and indicates the changes due to feedback. MuMMI allows a campaign to be initiated with a coarsely defined or approximate initial parametrization of the macro model, which is then steadily refined as the campaign proceeds, allowing exploration of continuum configurations that capture the CG scale with increasing accuracy. In the multiscale simulation discussed here, the sampling of RDFs is improved. However, the CG-to-continuum feedback, i.e., the use of these updated RDFs to update the lipid–protein interactions, failed to compensate for multibody effects, which likely resulted in too-high lipid-induced protein clustering in the continuum model. These multibody effects can be avoided by only considering single protein CG simulations for feedback. The implications for including them in this campaign are currently being evaluated

as well as more sophisticated approaches that can deconvolve the many-body effects.

3.3. Bridging CG and AA Resolutions. MuMMI's multiscale modeling capabilities extend beyond the continuum and CG scales by reaching atomistic resolution and time scales. The CG-to-AA coupling in MuMMI offers significant advantages by enabling more accurate and environment-specific models through AA simulations and overcoming the limitations of the CG Martini model. Specifically, whereas the secondary structures of proteins are fixed in the Martini model,⁸² AA resolutions can resolve the dependence of the secondary structures on the local environment. Through dynamic sampling of local CG environments and their investigation at the AA scale, MuMMI adapts the secondary structure to a preferred membrane-bound conformation, rather than being limited to the initial, solution-based model.

3.3.1. Dynamic Sampling of Diverse Protein Conformations from CG Scale. MuMMI employs a second, different type of dynamic-sampling approach to select CG snapshots to be promoted to the AA level. A total of 9,837,316 CG frames were generated through all CG simulations, providing a thorough exploration of the possible protein conformations.

As described earlier in Section 2.1.2, the dimensions of variability that are of primary interest in this study are the depth of CRD membrane insertion and the tilt and rotation angles that describe the orientation of RAS G-domain. MuMMI's sampling framework was used to create 3D histograms corresponding to these dimensions and was dynamically evolved by capturing additional CG frames as the simulations progressed. Spherical binning by angle was used to create 302 equal-area bins in the space of tilt and rotation, whereas five nonuniform bins were created to capture the relevant range of CRD depth. When queried for a new sample, the 3D histogram with 1510 total bins was evaluated against a similar histogram for the samples previously selected, and new samples were selected to satisfy a diversity factor of 20%. Figure 7 compares the distributions of the evaluated frames and selected frames (for simple visualization, polar coordinates are used rather than spherical binning) and illustrates that while the sampling captured the modes of the




Figure 8. AA-to-CG RAS-RBDCRD feedback. (A) Atomistic structure of RAS-RBDCRD interacting with a lipid bilayer. HVR and CRD feedback regions are shown in purple and red, respectively. (B, C) Probability distribution of (B) CRD depth and (C) HVR length obtained from AA simulations that are spawned from CG simulations with HVR feedback, HVR+CRD feedback (labeled CRD for the AA simulations and only shown in B as only CRD feedback was taken into account in AA conversation), and without any feedback. (D) Probability distribution of CRD depth (distance from membrane center) obtained from CG simulations with HVR feedback, HVR+CRD feedback, and without any feedback. (E, F) Probability distribution of (E) CRD depth and (F) HVR length obtained from AA simulations that are spawned from CG simulations without any feedback or with any type of feedback. The distributions for AA simulations are obtained at different time intervals corresponding to the first 16.7 ns of the simulations (early) and the rest of the simulations (late). The distributions are colored based on the type of feedback and simulation time intervals used for the analysis.

distribution (due to low diversity factor), a considerable amount of sampling was performed outside the dense areas.

3.3.2. Transformation of CG Frames into AA Configurations. The selected CG frames are then transformed to initial atomistic structures for AA simulations by the MuMMI Backmapping module and an updated CG-to-AA conversion tool called sinceCG with ~98% success rate per iterative attempt.⁷⁶ This tool, which is a modified version of the backward method,³⁴ converts selected snapshots from Martini CG simulations to CHARMM36 AA model. Several changes were made to reduce the numerical instability, have more accurate stereochemistry via modified Hamiltonians and improved structural integrity of converted structures.

3.3.3. In Situ Feedback from AA-to-CG. The AA-to-CG feedback in MuMMI is used to address a known limitation of the CG Martini model.^{32,82} Specifically, the secondary structure of RAS is fixed in CG and does not change based on the local membrane environment. By incorporating a feedback loop, we can update the force field parameters in the CG model based on the changes in secondary structures observed in the AA simulations for specific protein regions of RAS. This allowed us to adapt the CG secondary structure based on the statistics obtained from a large ensemble of AA simulations. Feedback updates the CG parameters that are dependent on the conformation and secondary structure of the protein such as backbone bead types, elastic bands, bonds, angles, and dihedrals between backbone beads. We focused in this campaign on RAS helix 5 (“HVR feedback”) and CRD loop residues T145-N161 (“CRD feedback”). In the running AA simulations, we cluster all secondary structure configurations in the HVR region and CRD region (separately), keeping track of the prevalence of different structures. If a new configuration is seen to reach a dominant level of prevalence,

we deem a “feedback event” to have occurred and modify the appropriate CG parameters accordingly.

Two separate feedback events occurred for RAS-RBDCRD during our multiscale simulation campaign. The first feedback altered the termination of RAS helix 5 from residue K172 to residue S171. (Note that this feedback does not affect RAS-only complexes as their C-terminus of RAS helix 5 is initially defined to end at residue S171.⁷⁵) The second feedback changed the secondary structures of residues T145-N161 from ECCSCCCTTTCSCCSE to EEEEEECTTTCSEEEE, increasing the lengths of β -strands in the CRD loops. More details on the AA-to-CG feedback are presented in López et al.⁷⁶

To understand the effect of feedback events on the structural properties of RAS-RBDCRD, the end-to-end distance of HVR and the CRD depth are evaluated separately for the simulations without any feedback, and the simulations with HVR-only feedback or HVR+CRD feedback (feedback regions are indicated in Figure 8A). CRD depth is defined as the distance between the center of mass of hydrophobic/cationic CRD loop residues and the center of mass of the bilayer along the bilayer normal.

Figure 8D shows the effect of different types of feedback on the probability distribution of CRD depth (defined as the distance from the center of the membrane) in CG simulations. The CRD depth distributions are very similar for CG simulations without any feedback and CG simulations with HVR feedback. On the other hand, the distribution shifts toward smaller values of CRD depth in the CG simulations with CRD feedback, showing that longer β -strands in the CRD loops allow proteins to lie deeper toward the bilayer center. As discussed in Section 3.3.1, novelty sampling is used to select CG snapshots to be promoted to the AA level, and the same

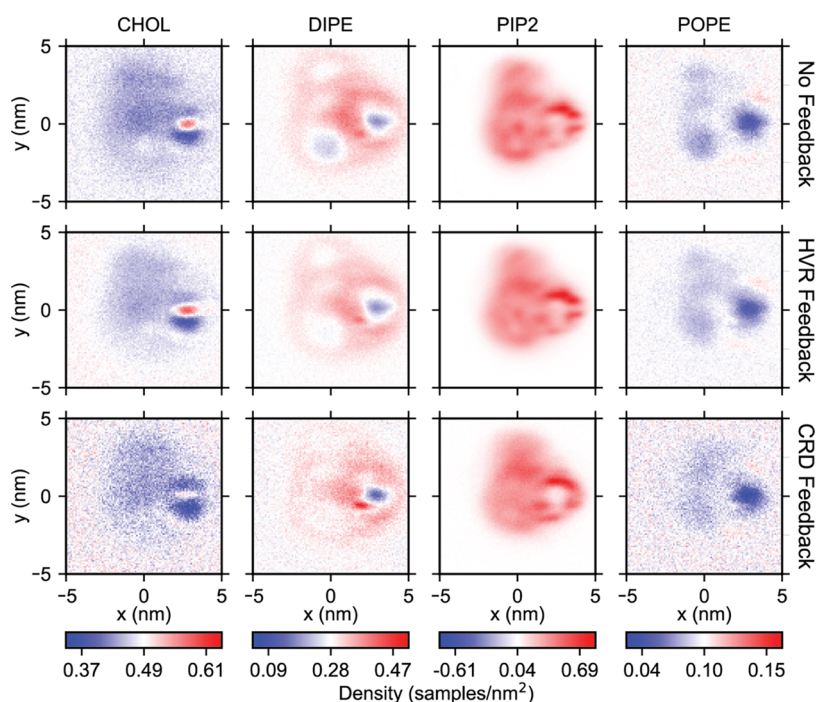


Figure 9. Effect of feedback-driven secondary structure changes on lipid densities around RAS-RBDCRD. Oriented two-dimensional (2D) distributions of DIPE, PAP6, POPC lipids, and cholesterol near RAS-RBDCRD in CG simulations without any feedback (top row), with HVR feedback (middle row), and HVR+CRD feedback (bottom row). Densities are computed after aligning the G-domain and CRD loop residues at the origin and the positive x-axis, respectively.

sampling approach is used before any feedback and after HVR and CRD feedback events.

The CRD depth is further evaluated in AA simulations at different time intervals (early and late) to understand how various types of feedback in CG simulations and diversity of CG-to-AA selection affect the behavior of RAS-RBDCRD at AA scale. The first 16.7 ns of each AA simulation corresponds to an early time interval, and the rest of the simulation corresponds to the late time interval. AA simulations spawned from CG simulations with CRD feedback have a CRD depth distribution shifted toward smaller values in early times (Figure 8B), consistent with the trends observed in CG simulations. This effect is found to disappear later in the AA simulations, indicating that feedback on CRD depths has not converged. CRD depth was also evaluated based on whether there was any type of feedback (either HVR or CRD) or no feedback (Figure 8E). A slight shift is observed toward smaller values in CRD depth distribution, but the shift is not as large as that observed for CRD feedback alone, shown in Figure 8B, since HVR feedback has no effect on the CRD depth distribution. Moreover, the effect of feedback on the distribution of HVR lengths in AA simulations is investigated, and it is found that HVR feedback results in longer HVR lengths at AA scale (Figure 8C,F). Overall, these results demonstrate that both HVR and CRD feedbacks influence the behavior of RAS-RBDCRD in the bilayer at CG and AA scales.

As detailed in an associated publication,⁷⁶ our CG-to-AA backmapping procedure builds the AA system constrained on a given CG input structure. The resulting AA structures can be energetically strained and deform and potentially not equilibrate due to the limited length of the AA simulations (50 ns is the current simulation campaign). If deformation is prominent and asymmetric, it could become dominant in the AA population, be selected for feedback to the CG, and then

further propagated as a new round of updated CG models get selected to AA. For the two feedback regions used in the current study, the RAF CRD loops region changes in β -strands length upon membrane insertion, but the change stabilizes and remains constant through the rest of the simulation campaign, as shown in López et al.⁷⁶ For the RAS HVR region, however, the RAS helix 5 shrinks by an average of 2.2 residues before feedback and 1.9 residues after—it does not exhibit a stable single value during this simulation campaign (Figure S6). The continued helix shortening is consistent with recent NMR backbone chemical shifts experiments, indicating RAS helix 5 preference to end at residue 170.⁷⁶ Nevertheless, as the helix length has not stabilized in these simulations, we cannot rule out a potential bias in the feedback-backmapping protocol.

Finally, the effect of feedback on lipid composition near RAS-RBDCRD is examined by computing densities after aligning the G-domain and CRD loop residues at the origin and the positive x-axis, respectively.⁵⁶ Figure 9 shows densities of cholesterol and DIPE in the inner leaflet extracted from the simulations without any feedback, and the simulations with HVR-only feedback or HVR+CRD feedback. We observe that the concentration of cholesterol depletes with CRD feedback, while the concentration of DIPE enriches with HVR feedback in certain regions, indicating that different types of feedback, corresponding to either a shortening of RAS helix 5 or increased lengths of β -strands in the CRD loops, influence the local lipid composition around the protein.

3.4. AA Results Update the Continuum Model through Double Feedback. To understand if there is any effect of changes occurring at AA scale on the macro simulations, the effect of atomistic feedback on lipid–protein RDFs in CG simulations was investigated for RAF states “ma” and “mb” by evaluating RDFs separately for the simulations without any feedback, and the simulations with HVR-only

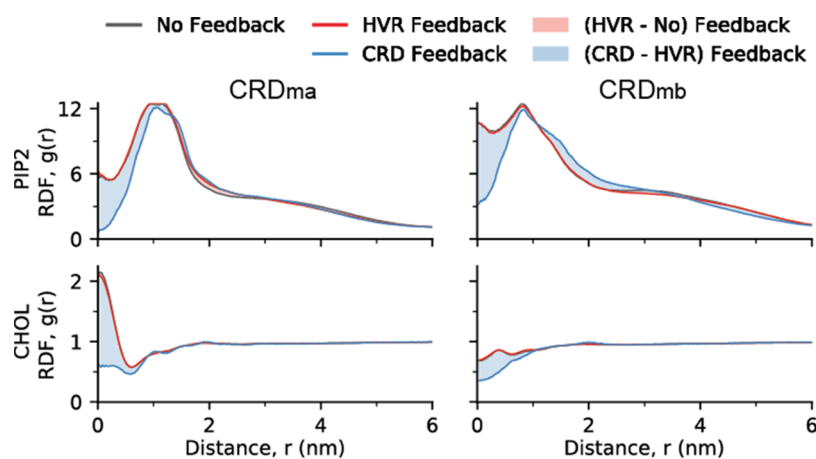


Figure 10. Effect of feedback-driven changes in secondary structure on lipid–protein RDFs in the CG simulations shown for PIP2 (top row) and cholesterol (bottom row) with two states of RAF1 (left and right). The figure shows the initial RDFs (black), RDFs after HVR feedback (red), and RDFs after the subsequent CRD feedback—the shaded regions highlight the differences between the different stages. The figure highlights the significant differences observed in RDFs after the CRD feedback, demonstrating the value that the on-the-fly feedback capabilities bring for coupling all three scales together, i.e., updating the coarsest (continuum) scale by accounting for the dynamics at the finest (AA) scale through the middle (CG) scale.

feedback or HVR+CRD feedback (Figure 9). For this computation, the number of lipids of each type were counted within each radial shell from the CRD region of the proteins (additional details on the computation of RDF are provided in the Method section). There is no significant difference between RDFs obtained from the CG simulations without any feedback and those obtained from the CG simulations with HVR feedback. On the other hand, CRD feedback has a significant impact on the RDFs at short distances, especially for certain bilayer species such as DIPE, PIP2, and cholesterol (CHOL). Figure 10 shows the CRD feedback for the RAS-RBDCRD in ma and mb states for PIP2 and CHOL lipids. These results demonstrate that feedback events affect the resulting RDFs in CG simulations, which can in turn tune the interaction potentials in macro simulations.

4. DISCUSSION AND CONCLUSIONS

The greatly extended three-scale MuMMI automatically executes multiscale simulations at three resolution scales: continuum, CG, and AA (Figure 1), using an ML-guided pairwise coupling of adjacent scales in an ensemble-based multiscale approach paired with on-the-fly feedback. Pairwise coupling of the three scales allows each to enhance the other. The lower resolution, faster models provide enhanced sampling, diversity of tested/assessed conditions, and starting points closer to equilibrated systems. Concurrently, the higher-resolution models gather more detailed information for the systems at hand, and on-the-fly feedback improves the fidelity of the lower-resolution models to more closely agree.

The blazingly fast continuum model samples a vast set of local lipid environments and possible protein configurations while continuously improving its protein–lipid interactions through aggregated online analysis of all CG simulations (Figure 6). MuMMI effectively utilizes available computational resources by sampling local configurations (patches) from the continuum model using a diversity-driven ML-guided approach that focuses on the expensive, finer-scale simulations on the most interesting or relevant areas. The resulting large CG ensemble captures a diversity of protein stoichiometries (Figure 3) as well as a diverse sampling of local lipid

compositions (Figure 4), allowing for lipid-dependent analysis of protein dynamics, as well as recording simulation weights marking the number of patches in the continuum model distribution represented by each CG simulation. In addition, by carefully constructing the CG simulations from the continuum model patches, the CG simulation equilibrium time is greatly reduced (Figure 5). The large and diverse ensemble of CG Martini 2 simulation provides the necessary resolution to resolve environmentally dependent lipid–lipid, lipid–protein, protein orientation and protein–protein interactions,^{13,14} although not without significant limitations such as over-representing protein–protein affinity⁹³ and fixed secondary structure assignment.⁸² Upon environmental change, some protein structures are altered, for the RAS-RBDCRD system protein domains with suspected dynamic secondary structures were selected for feedback; those are α -helix 5 that connects to the disordered RAS HVR and the region of the RAF1 CRD that inserts into the bilayer. Additionally, due to limitations of CG modeling, reference AA simulations for backing up and verifying specific interactions are always desirable. In the RAS-RBDCRD multiscale simulation, 20% of GPU resources were dedicated to AA simulations. A large ensemble of AA simulations was built (Figure 3B), sampling different RAS G-domain membrane orientations at different RAF1 CRD membrane insertion depths (Figure 7). As computational resources become available, CG frames are selected from the pool of all so far simulated CG based on a priority specific criteria of novelty vs random (repeated sampling) and each selected CG frame is carefully converted from CG to AA.⁷⁶ MuMMI aggregates the on-the-fly analysis of all AA simulations and updates the current protein parameters for future CG simulations when the population of observed secondary structures changes, adapting the CG protein structure to the local environment more commonly sampled. As the multiscale simulation progresses, the different scales reinforce each other, the progressively better sampled CG and AA simulations improve the accuracy of continuum and CG, respectively, and as the different sub-simulations are run concurrently, updates to the CG from the

AA are gradually passed into the continuum model (Figure 10).

Resolving protein–membrane interactions with MD simulations requires significant sampling and more so when considering complex membrane mixtures^{13,14} which can form specific lipid–protein fingerprints.²⁹ Specific proteins also adapt their membrane orientation and dynamics based on the local lipid compositions, shown for RAS in a previous two-scale MuMMI multiscale simulation.⁵⁶ When RAS binds the RBDCRD domains of RAF1, the membrane orientation of the RAS G-domain is significantly affected (Figure 2), and the relative orientation of the CRD changes upon membrane binding.⁷⁵ The multiscale RAS-RBDCRD simulation shown here was executed to both showcase the extended three-scale MuMMI as well as to resolve the membrane dynamics of the RAS-RBDCRD complex, for which detailed analysis, interpretation, and experimental support are currently underway.

With increasing availability of highly parallel computational resources, the types of simulations that can be performed are changing. Already one-off simulations are rare—more often a suite of different conditions is explored, often managed using simulation workflow tools. Scanning of appropriate conditions is very effective when the desired reaction coordinate or relevant changes in environment to explore are known a priori. MuMMI expands upon current workflow technology by facilitating the use of a high-level macro model to enable rapid system exploration coupled with a fully automated conversion to finer-detailed CG and AA resolutions. In the absence of predefined conditions to explore, MuMMI can substitute valuable human resources with computational resources to effectively create and sample an extraordinarily wide range of correlated system configurations. The three-scale MuMMI combines different resolution components to significantly extend (see refs 75, 76, 78–80 and SI Section S1) the previous version of MuMMI⁵⁶ and demonstrate the utility of running linked ensemble-based multiscale simulations. MuMMI bidirectionally couples adjacent resolution scales using ML-guided sampling and specific system-building routines (going from coarse to fine) and aggregation of in situ feedback (going from fine to coarse). MuMMI is designed to be generalizable, as it is able to run other model systems by tuning or reparametrizing the current simulation modules and/or by switching out specific modules. MuMMI is also highly scalable, running efficiently on systems ranging in size from a handful of nodes to some of the world's most powerful computers^{77,78} and is anticipated to run on forthcoming exascale machines as well. Looking ahead, we expect that the cost of human resources will continue to outpace computational resources, and frameworks such as MuMMI will become more common, allowing for fully automatic multiscale campaigns sampling relevant system conditions.

■ ASSOCIATED CONTENT

Data Availability Statement

All simulation input and parameter files will be made available at <https://bbs.llnl.gov/>. All simulation raw data ~200 TB will be hosted on the NIH MoDaC server (<https://modac.cancer.gov/>). MuMMI is composed of numerous sub-components, both freely available third-party applications and custom codes: - ddcMD (github.com/LLNL/ddcMD), - ddcMDconverter (github.com/LLNL/ddcMDconverter), - MDAnalysis modified to support ddcMD (github.com/XiaohuaZhangLLNL/mdanalysis), - Maestro (github.com/LLNL/maestrowf), - Flux

(github.com/flux-framework), - Pytaridx (github.com/LLNL/pytaridx), - dynIm (github.com/LLNL/dynim), - MuMMI-core (github.com/mummi-framework/mummi-core), and - MuMMI-ras (github.com/mummi-framework/mummi-ras).

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jctc.2c01018>.

Methods describing the macro model parameterization (Section 1.1), protein structure input libraries using multiscale simulation (Section 1.2), and description of the updated MuMMI Createsims (Section 1.3) and CG simulation and analysis (Section 1.4) modules, additional figures for converting continuum patch to CG (Figure S4), full list of feedback RDFs (Figure S5), and RAS helix 5 length distribution in AA simulations (Figure S6); in Section S2 and a table with CG simulation protein compositions (Table S1) in Section S3 (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Dwight V. Nissley – RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States; orcid.org/0000-0001-7523-116X; Email: nissleyd@mail.nih.gov
Frederick H. Streitz – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States; orcid.org/0000-0003-0924-037X; Email: streitz1@llnl.gov

Authors

Helgi I. Ingólfsson – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States; orcid.org/0000-0002-7613-9143
Harsh Bhatia – Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States; orcid.org/0000-0001-8712-7773
Fikret Aydin – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States; orcid.org/0000-0003-3237-8043
Tomas Opielstrup – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States
Cesar A. López – Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States
Liam G. Stanton – Department of Mathematics and Statistics, San José State University, San José, California 95192, United States
Timothy S. Carpenter – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States; orcid.org/0000-0001-7848-9983
Sergio Wong – Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States
Francesco Di Natale – Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States

- Xiaohua Zhang** – *Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Joseph Y. Moon** – *Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Christopher B. Stanley** – *Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830, United States; orcid.org/0000-0002-4226-7710*
- Joseph R. Chavez** – *Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Kien Nguyen** – *Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States; orcid.org/0000-0003-0038-5601*
- Gautham Dharuman** – *Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Violetta Burns** – *Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*
- Rebika Shrestha** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Debanjan Goswami** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Gulcin Gulden** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Que N. Van** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Arvind Ramanathan** – *Computing, Environment & Life Sciences (CELS) Directorate, Argonne National Laboratory, Lemont, Illinois 60439, United States; orcid.org/0000-0002-1622-5488*
- Brian Van Essen** – *Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Nicolas W. Hengartner** – *Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States*
- Andrew G. Stephen** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Thomas Turbyville** – *RAS Initiative, The Cancer Research Technology Program, Frederick National Laboratory, Frederick, Maryland 21701, United States*
- Peer-Timo Bremer** – *Computing Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- S. Gnanakaran** – *Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, United States; orcid.org/0000-0002-9368-3044*
- James N. Glosli** – *Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*
- Felice C. Lightstone** – *Physical and Life Sciences (PLS) Directorate, Lawrence Livermore National Laboratory, Livermore, California 94550, United States*

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.jctc.2c01018>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy (DOE) by Lawrence Livermore National Laboratory (LLNL) under Contract DE-AC52-07NA27344, Los Alamos National Laboratory (LANL) under Contract DE-AC5206NA25396, Oak Ridge National Laboratory under Contract DE-AC05-00OR22725, Argonne National Laboratory (ANL) under Contract DE-AC02-06-CH11357, and under the auspices of the National Cancer Institute (NCI) by Frederick National Laboratory for Cancer Research (FNLRCR) under Contract 75N91019D00024. This work has been supported by the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C) program established by the U.S. DOE and the NCI of the National Institutes of Health. This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. For computing time, the authors thank the Advanced Scientific Computing Research Leadership Computing Challenge (ALCC) for time on Summit, the Livermore Institutional Grand Challenge for time on Lassen and LANL Institutional computing. The LANL Institutional Computing Program is supported by the U.S. DOE National Nuclear Security Administration under Contract No. DE-AC52-06NA25396. For computing support, the authors thank OLCF and LC staff. Release: LLNL-JRNL-833009.

ABBREVIATIONS

AA, all-atom; CG, coarse-grained; CHOL, cholesterol; CRD, cysteine-rich domain; DDFT, dynamic density functional theory; DIPE, 1,2-dilinoleoyl-*sn*-glycero-3-phosphoethanolamine; DL, deep learning; DPSM, *N*-stearoyl-*D*-erythro-sphingosylphosphorylcholine; HVR, hypervariable region of RAS; MD, molecular dynamics; ML, machine learning (ML); MuMMI, the multiscale machine-learned modeling infrastructure; PAPC, 1-palmitoyl-2- arachidonoyl-*sn*-glycero-3-phosphocholine; PAPS, 1-palmitoyl-2- arachidonoyl-*sn*-glycero-3-phosphatidylserine; PIP2, phosphatidylinositol 4,5-bisphosphate; PM, plasma membrane; POPC, 1-palmitoyl-2-oleoyl-*sn*-glycero-3-phosphocholine; POPE, 1-palmitoyl-2-oleoyl-*sn*-glycero-3-phosphoethanolamine; RAS-RBDCRD, KRAS4b bound to RBD and CRD domains of RAF1; RAS, we specifically used KRAS4b; RBD, RAS-binding domain; RDFs, radial density functions

REFERENCES

- (1) Hoekstra, A.; Chopard, B.; Coveney, P. Multiscale modelling and simulation: a position paper. *Phil. Trans. R. Soc. A* **2014**, *372*, No. 20130377.
- (2) Chopard, B.; Borgdorff, J.; Hoekstra, A. G. A framework for multi-scale modelling. *Philos. Trans. R. Soc., A* **2014**, *372*, No. 20130378.
- (3) Ingram, G. D.; Cameron, I. T.; Hango, K. M. Classification and analysis of integrating frameworks in multiscale modelling. *Chem. Eng. Sci.* **2004**, *59*, 2171–2187.
- (4) Krzhizhanovskaya, V. V.; Groen, D.; Bozak, B.; Hoekstra, A. G. Multiscale Modelling and Simulation Workshop: 12 Years of Inspiration. *Proc. Comput. Sci.* **2015**, *51*, 1082–1087.

- (5) van der Giessen, E.; Schultz, P. A.; Bertin, N.; Bulatov, V. V.; Cai, W.; Csányi, G.; Foiles, S. M.; Geers, M. G. D.; González, C.; Hütter, M.; Kim, W. K.; Kochmann, D. M.; Llorca, J.; Mattsson, A. E.; Rottler, J.; Shluger, A.; Sills, R. B.; Steinbach, I.; Strachan, A.; Tadmor, E. B. Roadmap on multiscale materials modeling. *Modell. Simul. Mater. Sci. Eng.* **2020**, *28*, No. 043001.
- (6) Llorca, J.; González, C.; Molina-Aldareguía, J. M.; Segurado, J.; Seltzer, R.; Sket, F.; Rodríguez, M.; Sádaba, S.; Muñoz, R.; Canal, L. P. Multiscale Modeling of Composite Materials: a Roadmap Towards Virtual Testing. *Adv. Mater.* **2011**, *23*, 5130–5147.
- (7) Peters, N. Multiscale combustion and turbulence. *Proc. Combust. Inst.* **2009**, *32*, 1–25.
- (8) Miller, R. E.; Tadmor, E. B. A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods. *Modell. Simul. Mater. Sci. Eng.* **2009**, *17*, No. 053001.
- (9) Enkavi, G.; Javanainen, M.; Kulig, W.; Rog, T.; Vattulainen, I. Multiscale Simulations of Biological Membranes: The Challenge To Understand Biological Phenomena in a Living Substance. *Chem. Rev.* **2019**, *119*, 5607–5774.
- (10) Huber, R. G.; Carpenter, T. S.; Dube, N.; Holdbrook, D. A.; Ingólfsson, H. I.; Irvine, W. A.; Marzinek, J. K.; Samsudin, F.; Allison, J. R.; Khalid, S.; Bond, P. J. Multiscale Modeling and Simulation Approaches to Lipid-Protein Interactions. *Methods Mol. Biol.* **2019**, *2003*, 1–30.
- (11) Ayton, G. S.; Voth, G. A. Multiscale simulation of protein mediated membrane remodeling. *Semin. Cell Dev. Biol.* **2010**, *21*, 357–362.
- (12) Tozzini, V. Multiscale Modeling of Proteins. *Acc. Chem. Res.* **2010**, *43*, 220–230.
- (13) Corradi, V.; Sejdiu, B. I.; Mesa-Galloso, H.; Abdizadeh, H.; Noskov, S. Y.; Marrink, S. J.; Tieleman, D. P. Emerging Diversity in Lipid-Protein Interactions. *Chem. Rev.* **2019**, *119*, 5775–5848.
- (14) Marrink, S. J.; Corradi, V.; Souza, P. C. T.; Ingólfsson, H. I.; Tieleman, D. P.; Sansom, M. S. P. Computational Modeling of Realistic Cell Membranes. *Chem. Rev.* **2019**, *119*, 6184–6226.
- (15) Pan, A. C.; Jacobson, D.; Yatsenko, K.; Sritharan, D.; Weinreich, T. M.; Shaw, D. E. Atomic-level characterization of protein-protein association. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 4244–4249.
- (16) Gumbart, J.; Wang, Y.; Aksimentiev, A.; Tajkhorshid, E.; Schulten, K. Molecular dynamics simulations of proteins in lipid bilayers. *Curr. Opin. Struct. Biol.* **2005**, *15*, 423–431.
- (17) Kmiecik, S.; Gront, D.; Kolinski, M.; Wieteska, L.; Dawid, A. E.; Kolinski, A. Coarse-Grained Protein Models and Their Applications. *Chem. Rev.* **2016**, *116*, 7898–7936.
- (18) Chavent, M.; Duncan, A. L.; Sansom, M. S. Molecular dynamics simulations of membrane proteins and their interactions: from nanoscale to mesoscale. *Curr. Opin. Struct. Biol.* **2016**, *40*, 8–16.
- (19) Argudo, D.; Bethel, N. P.; Marcoline, F. V.; Wolgemuth, C. W.; Grabe, M. New Continuum Approaches for Determining Protein-Induced Membrane Deformations. *Biophys. J.* **2017**, *112*, 2159–2172.
- (20) Yoo, J.; Jackson, M. B.; Cui, Q. A comparison of coarse-grained and continuum models for membrane bending in lipid bilayer fusion pores. *Biophys. J.* **2013**, *104*, 841–852.
- (21) Ingólfsson, H. I.; Lopez, C. A.; Uusitalo, J. J.; de Jong, D. H.; Gopal, S. M.; Periole, X.; Marrink, S. J. The power of coarse graining in biomolecular simulations. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2014**, *4*, 225–248.
- (22) Baaden, M.; Marrink, S. J. Coarse-grain modelling of protein-protein interactions. *Curr. Opin. Struct. Biol.* **2013**, *23*, 878–886.
- (23) Pak, A. J.; Voth, G. A. Advances in coarse-grained modeling of macromolecular complexes. *Curr. Opin. Struct. Biol.* **2018**, *52*, 119–126.
- (24) Ayton, G. S.; Lyman, E.; Voth, G. A. Hierarchical coarse-graining strategy for protein-membrane systems to access mesoscopic scales. *Faraday Discuss.* **2010**, *144*, 347–357.
- (25) Srivastava, A.; Voth, G. A. A Hybrid Approach for Highly Coarse-grained Lipid Bilayer Models. *J. Chem. Theory Comput.* **2013**, *9*, 750–765.
- (26) Arkhipov, A.; Yin, Y.; Schulten, K. Four-Scale Description of Membrane Sculpting by BAR Domains. *Biophys. J.* **2008**, *95*, 2806–2821.
- (27) Shelley, J. C.; Shelley, M. Y.; Reeder, R. C.; Bandyopadhyay, S.; Klein, M. L. A Coarse Grain Model for Phospholipid Simulations. *J. Phys. Chem. B* **2001**, *105*, 4464–4470.
- (28) Noid, W. G.; Chu, J.-W.; Ayton, G. S.; Krishna, V.; Izvekov, S.; Voth, G. A.; Das, A.; Andersen, H. C. The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models. *J. Chem. Phys.* **2008**, *128*, No. 244114.
- (29) Corradi, V.; Mendez-Villuendas, E.; Ingólfsson, H. I.; Gu, R.-X.; Suda, I.; Melo, M. N.; Moussatova, A.; DeGagné, L. J.; Sejdiu, B. I.; Singh, G.; Wassenaar, T. A.; Delgado Magnero, K.; Marrink, S. J.; Tieleman, D. P. Lipid-Protein Interactions Are Unique Fingerprints for Membrane Proteins. *ACS Cent. Sci.* **2018**, *4*, 709–717.
- (30) Melo, M. N.; Arnarez, C.; Sikkema, H.; Kumar, N.; Walko, M.; Berendsen, H. J. C.; Kocer, A.; Marrink, S. J.; Ingólfsson, H. I. High-Throughput Simulations Reveal Membrane-Mediated Effects of Alcohols on MscL Gating. *J. Am. Chem. Soc.* **2017**, *139*, 2664–2671.
- (31) Lucendo, E.; Sancho, M.; Lolicato, F.; Javanainen, M.; Kulig, W.; Leiva, D.; Duart, G.; Andreu-Fernández, V.; Mingarro, I.; Orzáez, M. Mcl-1 and Bok transmembrane domains: Unexpected players in the modulation of apoptosis. *Proc. Natl. Acad. Sci. U.S.A.* **2020**, *117*, 27980–27988.
- (32) Marrink, S. J.; Risselada, H. J.; Yefimov, S.; Tieleman, D. P.; de Vries, A. H. The MARTINI Force Field: Coarse Grained Model for Biomolecular Simulations. *J. Phys. Chem. B* **2007**, *111*, 7812–7824.
- (33) Marrink, S. J.; Tieleman, D. P. Perspective on the Martini model. *Chem. Soc. Rev.* **2013**, *42*, 6801–6822.
- (34) Wassenaar, T. A.; Pluhackova, K.; Böckmann, R. A.; Marrink, S. J.; Tieleman, D. P. Going Backward: A Flexible Geometric Approach to Reverse Transformation from Coarse Grained to Atomistic Models. *J. Chem. Theory Comput.* **2014**, *10*, 676–690.
- (35) Stansfeld, P. J.; Sansom, M. S. P. From coarse grained to atomistic: a serial multiscale approach to membrane protein simulations. *J. Chem. Theory Comput.* **2011**, *7*, 1157–1166.
- (36) Louison, K. A.; Dryden, I. L.; Laughton, C. A. GLIMPS: A Machine Learning Approach to Resolution Transformation for Multiscale Modeling. *J. Chem. Theory Comput.* **2021**, *17*, 7930–7937.
- (37) Simunovic, M.; Voth, G. A. Membrane tension controls the assembly of curvature-generating proteins. *Nat. Commun.* **2015**, *6*, No. 7219.
- (38) Wassenaar, T. A.; Ingólfsson, H. I.; Priess, M.; Marrink, S. J.; Schäfer, L. V. Mixing MARTINI: electrostatic coupling in hybrid atomistic-coarse-grained biomolecular simulations. *J. Phys. Chem. B* **2013**, *117*, 3516–3530.
- (39) Shi, Q.; Izvekov, S.; Voth, G. A. Mixed Atomistic and Coarse-Grained Molecular Dynamics: Simulation of a Membrane-Bound Ion Channel. *J. Phys. Chem. B* **2006**, *110*, 15045–15048.
- (40) Zavadlav, J.; Marrink, S. J.; Praprotnik, M. SWINGER: a clustering algorithm for concurrent coupling of atomistic and supramolecular liquids. *Interface Focus* **2019**, *9*, No. 20180075.
- (41) Cortes-Huerta, R.; Praprotnik, M.; Kremer, K.; Delle Site, L. From adaptive resolution to molecular dynamics of open systems. *Eur. Phys. J. B* **2021**, *94*, No. 189.
- (42) Warshel, A. Computer Simulations of Enzyme Catalysis: Methods, Progress, and Insights. *Annu. Rev. Biophys. Biomol. Struct.* **2003**, *32*, 425–443.
- (43) Lin, H.; Truhlar, D. G. QM/MM: what have we learned, where are we, and where do we go from here? *Theor. Chem. Acc.* **2007**, *117*, No. 185.
- (44) McCullagh, M.; Saunders, M. G.; Voth, G. A. Unraveling the Mystery of ATP Hydrolysis in Actin Filaments. *J. Am. Chem. Soc.* **2014**, *136*, 13053–13058.
- (45) Zhang, Y.; Liu, H.; Yang, W. Free energy calculation on enzyme reactions with an efficient iterative procedure to determine minimum

energy paths on a combined ab initio QM/MM potential energy surface. *J. Chem. Phys.* **2000**, *112*, 3483–3492.

(46) Rosta, E.; Klähn, M.; Warshel, A. Towards Accurate Ab Initio QM/MM Calculations of Free-Energy Profiles of Enzymatic Reactions. *J. Phys. Chem. B* **2006**, *110*, 2934–2941.

(47) Huber, G. A.; Kim, S. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J.* **1996**, *70*, 97–110.

(48) Zuckerman, D. M.; Chong, L. T. Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annu. Rev. Biophys.* **2017**, *46*, 43–57.

(49) Noé, F. Machine Learning for Molecular Dynamics on Long Timescales. In *Machine Learning Meets Quantum Physics*; Schütt, K. T.; Chmiela, S.; von Lilienfeld, O. A.; Tkatchenko, A.; Tsuda, K.; Müller, K.-R., Eds.; Springer International Publishing, 2020; pp 331–372.

(50) Kasson, P. M.; Kelley Nicholas, W.; Singhal, N.; Vrljic, M.; Brunger Axel, T.; Pande Vijay, S. Ensemble molecular dynamics yields submillisecond kinetics and intermediates of membrane fusion. *Proc. Natl. Acad. Sci. U.S.A.* **2006**, *103*, 11916–11921.

(51) Kasson, P. M.; Jha, S. Adaptive ensemble simulations of biomolecules. *Curr. Opin. Struct. Biol.* **2018**, *52*, 87–94.

(52) Singharoy, A.; Joshi, H.; Ortoleva, P. J. Multiscale Macromolecular Simulation: Role of Evolving Ensembles. *J. Chem. Inf. Model.* **2012**, *52*, 2638–2649.

(53) Lee, H.; Turilli, M.; Jha, S.; Bhowmik, D.; Ma, H.; Ramanathan, A. In *DeepDriveMD: Deep-Learning Driven Adaptive Molecular Simulations for Protein Folding*, 2019 IEEE/ACM Third Workshop on Deep Learning on Supercomputers (DLS) 2019; pp 12–19.

(54) Moritsugu, K. Multiscale Enhanced Sampling Using Machine Learning. *Life* **2021**, *11*, No. 1076.

(55) Casalino, L.; Dommer, A. C.; Gaieb, Z.; Barros, E. P.; Sztain, T.; Ahn, S.-H.; Trifan, A.; Brace, A.; Bogetti, A. T.; Clyde, A.; Ma, H.; Lee, H.; Turilli, M.; Khalid, S.; Chong, L. T.; Simmerling, C.; Hardy, D. J.; Maia, J. D. C.; Phillips, J. C.; Kurth, T.; Stern, A. C.; Huang, L.; McCalpin, J. D.; Tatineni, M.; Gibbs, T.; Stone, J. E.; Jha, S.; Ramanathan, A.; Amaro, R. E. AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics. *Int. J. High Perform. Comput. Appl.* **2021**, *35*, 432–451.

(56) Ingólfsson, H. I.; Neale, C.; Carpenter, T. S.; Shrestha, R.; Lopez, C. A.; Tran, T. H.; Opperstrup, T.; Bhatia, H.; Stanton, L. G.; Zhang, X.; Sundram, S.; Di Natale, F.; Agarwal, A.; Dharuman, G.; Kokkila Schumacher, S. I. L.; Turbyville, T.; Gulten, G.; Van, Q. N.; Goswami, D.; Jean-Francois, F.; Agamasu, C.; Chen; Hettige, J. J.; Travers, T.; Sarkar, S.; Surui, M. P.; Yang, Y.; Moody, A.; Liu, S.; Van Essen, B. C.; Voter, A. F.; Ramanathan, A.; Hengartner, N. W.; Simanshu, D. K.; Stephen, A. G.; Bremer, P. T.; Gnanakaran, S.; Glosli, J. N.; Lightstone, F. C.; McCormick, F.; Nissley, D. V.; Streitz, F. H. Machine learning-driven multiscale modeling reveals lipid-dependent dynamics of RAS signaling proteins. *Proc. Natl. Acad. Sci. U.S.A.* **2022**, *119*, No. e2113297119.

(57) Bhatia, H.; Carpenter, T. S.; Ingólfsson, H. I.; Dharuman, G.; Karande, P.; Liu, S.; Opperstrup, T.; Neale, C.; Lightstone, F. C.; Van Essen, B.; Glosli, J. N.; Bremer, P.-T. Machine-learning-based dynamic-importance sampling for adaptive multiscale simulations. *Nat. Mach. Intell.* **2021**, *3*, 401–409.

(58) Drostén, M.; Barbacid, M. Targeting the MAPK Pathway in KRAS-Driven Tumors. *Cancer Cell* **2020**, *37*, 543–550.

(59) Molina, J. R.; Adjei, A. A. The Ras/Raf/MAPK Pathway. *J. Thorac. Oncol.* **2006**, *1*, 7–9.

(60) Prior, I. A.; Lewis, P. D.; Mattos, C. A comprehensive survey of Ras mutations in cancer. *Cancer Res.* **2012**, *72*, 2457–2467.

(61) Prior, I. A.; Hood, F. E.; Hartley, J. L. The Frequency of Ras Mutations in Cancer. *Cancer Res.* **2020**, *80*, 2969–2974.

(62) Stephen, Andrew G.; Esposito, D.; Bagni, Rachel K.; McCormick, F. Dragging Ras Back in the Ring. *Cancer Cell* **2014**, *25*, 272–281.

(63) Kessler, D.; Gmachl, M.; Mantoulidis, A.; Martin Laetitia, J.; Zoepfel, A.; Mayer, M.; Gollner, A.; Covini, D.; Fischer, S.;

Gerstberger, T.; Gmaschitz, T.; Goodwin, C.; Greb, P.; Häring, D.; Hela, W.; Hoffmann, J.; Karolyi-Oezguer, J.; Knesl, P.; Kornigg, S.; Koegl, M.; Kousek, R.; Lamarre, L.; Moser, F.; Munico-Martinez, S.; Peinsipp, C.; Phan, J.; Rinnenthal, J.; Sai, J.; Salamon, C.; Scherbantin, Y.; Schipany, K.; Schnitzer, R.; Schrenk, A.; Sharps, B.; Sisler, G.; Sun, Q.; Waterson, A.; Wolkerstorfer, B.; Zeeb, M.; Pearson, M.; Fesik Stephen, W.; McConnell Darryl, B. Drugging an undruggable pocket on KRAS. *Proc. Natl. Acad. Sci. U.S.A.* **2019**, *116*, 15823–15829.

(64) Lanman, B. A.; Allen, J. R.; Allen, J. G.; Amegadzie, A. K.; Ashton, K. S.; Booker, S. K.; Chen, J. J.; Chen, N.; Frohn, M. J.; Goodman, G.; Kopecky, D. J.; Liu, L.; Lopez, P.; Low, J. D.; Ma, V.; Minatti, A. E.; Nguyen, T. T.; Nishimura, N.; Pickrell, A. J.; Reed, A. B.; Shin, Y.; Siegmund, A. C.; Tamayo, N. A.; Tegley, C. M.; Walton, M. C.; Wang, H.-L.; Wurz, R. P.; Xue, M.; Yang, K. C.; Achanta, P.; Bartberger, M. D.; Canon, J.; Hollis, L. S.; McCarter, J. D.; Mohr, C.; Rex, K.; Saiki, A. Y.; San Miguel, T.; Volak, L. P.; Wang, K. H.; Whittington, D. A.; Zech, S. G.; Lipford, J. R.; Cee, V. J. Discovery of a Covalent Inhibitor of KRASG12C (AMG 510) for the Treatment of Solid Tumors. *J. Med. Chem.* **2020**, *63*, 52–65.

(65) Hallin, J.; Engstrom, L. D.; Hargis, L.; Calinisan, A.; Aranda, R.; Briere, D. M.; Sudhakar, N.; Bowcut, V.; Baer, B. R.; Ballard, J. A.; Burkard, M. R.; Fell, J. B.; Fischer, J. P.; Vigers, G. P.; Xue, Y.; Gatto, S.; Fernandez-Banet, J.; Pavlicek, A.; Velastagui, K.; Chao, R. C.; Barton, J.; Pierobon, M.; Baldelli, E.; Patricoin, E. F., 3rd; Cassidy, D. P.; Marx, M. A.; Rybkin, I. I.; Johnson, M. L.; Ou, S. I.; Lito, P.; Papadopoulos, K. P.; Jänne, P. A.; Olson, P.; Christensen, J. G. The KRAS(G12C) Inhibitor MRTX849 Provides Insight toward Therapeutic Susceptibility of KRAS-Mutant Cancers in Mouse Models and Patients. *Cancer Discovery* **2020**, *10*, 54–71.

(66) Weise, K.; Kapoor, S.; Denter, C.; Nikolaus, J.; Opitz, N.; Koch, S.; Triola, G.; Herrmann, A.; Waldmann, H.; Winter, R. Membrane-Mediated Induction and Sorting of K-Ras Microdomain Signaling Platforms. *J. Am. Chem. Soc.* **2011**, *133*, 880–887.

(67) Erwin, N.; Sperlich, B.; Garivet, G.; Waldmann, H.; Weise, K.; Winter, R. Lipoprotein insertion into membranes of various complexity: lipid sorting, interfacial adsorption and protein clustering. *Phys. Chem. Chem. Phys.* **2016**, *18*, 8954–8962.

(68) Muratcioglu, S.; Chavan; Tanmay, S.; Freed; Benjamin, C.; Jang, H.; Khavrutskii, L.; Freed, R. N.; Dyba; Marzena, A.; Stefanisko, K.; Tarasov; Sergey, G.; Gursoy, A.; Keskin, O.; Tarasova, N. I.; Gaponenko, V.; Nussinov, R. GTP-Dependent K-Ras Dimerization. *Structure* **2015**, *23*, 1325–1335.

(69) Prakash, P.; Sayyed-Ahmad, A.; Cho, K.-J.; Dolino, D. M.; Chen, W.; Li, H.; Grant, B. J.; Hancock, J. F.; Gorfe, A. A. Computational and biochemical characterization of two partially overlapping interfaces and multiple weak-affinity K-Ras dimers. *Sci. Rep.* **2017**, *7*, No. 40109.

(70) Sayyed-Ahmad, A.; Cho, K.-J.; Hancock, J. F.; Gorfe, A. A. Computational Equilibrium Thermodynamic and Kinetic Analysis of K-Ras Dimerization through an Effector Binding Surface Suggests Limited Functional Role. *J. Phys. Chem. B* **2016**, *120*, 8547–8556.

(71) Jang, H.; Muratcioglu, S.; Gursoy, A.; Keskin, O.; Nussinov, R. Membrane-associated Ras dimers are isoform-specific: K-Ras dimers differ from H-Ras dimers. *Biochem. J.* **2016**, *473*, 1719–1732.

(72) Zhou, Y.; Liang, H.; Rodkey, T.; Ariotti, N.; Parton, R. G.; Hancock, J. F. Signal Integration by Lipid-Mediated Spatial Cross Talk between Ras Nanoclusters. *Mol. Cell. Biol.* **2014**, *34*, 862.

(73) Ngo, V. A.; Sarkar, S.; Neale, C.; Garcia, A. E. How Anionic Lipids Affect Spatiotemporal Properties of KRAS4B on Model Membranes. *J. Phys. Chem. B* **2020**, *124*, 5434–5453.

(74) Ingólfsson, H. I.; Bhatia, H.; Zeppelin, T.; Bennett, W. F. D.; Carpenter, K. A.; Hsu, P.-C.; Dharuman, G.; Bremer, P.-T.; Schiott, B.; Lightstone, F. C.; Carpenter, T. S. Capturing Biologically Complex Tissue-Specific Membranes at Different Levels of Compositional Complexity. *J. Phys. Chem. B* **2020**, *124*, 7819–7829.

(75) Nguyen, K.; Lopez, C. A.; Neale, C.; Van, Q. N.; Carpenter, T. S.; Di Natale, F.; Travers, T.; Tran, T. H.; Chan, A. H.; Bhatia, H.; Frank, P. H.; Tonnelly, M.; Zhang, X.; Gulten, G.; Reddy, T.; Burns,

- V.; Opperstrup, T.; Hengartner, N.; Simanshu, D. K.; Bremer, P.-T.; Chen, D.; Glosli, J. N.; Shrestha, R.; Turbyville, T.; Streitz, F. H.; Nissley, D. V.; Ingólfsson, H. I.; Stephen, A. G.; Lightstone, F. C.; Gnanakaran, S. Exploring CRD mobility during RAS/RAF engagement at the membrane. *Biochem. J.* **2022**, *121*, 3630–3650.
- (76) López, C. A.; Zhang, X.; Aydin, F.; Shrestha, R.; Van, Q. N.; Stanley, C. B.; Carpenter, T. S.; Nguyen, K.; Patel, L. A.; Chen, Burns, V.; Hengartner, N. W.; Reddy, T. J. E.; Bhatia, H.; Di Natale, F.; Tran, T. H.; Chan, A. H.; Simanshu, D. K.; Nissley, D. V.; Streitz, F. H.; Stephen, A. G.; Turbyville, T. J.; Lightstone, F. C.; Gnanakaran, S.; Ingólfsson, H. I.; Neale, C. Asynchronous Reciprocal Coupling of Martini 2.2 Coarse-Grained and CHARMM36 All-Atom Simulations in an Automated Multiscale Framework. *J. Chem. Theory Comput.* **2022**, *18*, S025–S045.
- (77) Di Natale, F.; Bhatia, H.; Carpenter, T. S.; Neale, C.; Schumacher, S. K.; Opperstrup, T.; Stanton, L.; Zhang, X.; Sundram, S.; Scogland, T. R. W.; Dharuman, G.; Surh, M. P.; Yang, Y.; Misale, C.; Schneidenbach, L.; Costa, C.; Kim, C.; D'Amora, B.; Gnanakaran, S.; Nissley, D. V.; Streitz, F.; Lightstone, F. C.; Bremer, P.-T.; Glosli, J. N.; Ingólfsson, H. I. In *A Massively Parallel Infrastructure for Adaptive Multiscale Simulations: Modeling RAS Initiation Pathway for Cancer*, The International Conference for High Performance Computing, Networking, Storage and Analysis; ACM: Denver, Colorado, 2019; p 57.
- (78) Bhatia, H.; Di Natale, F.; Moon, J. Y.; Zhang, X.; Chavez, J. R.; Aydin, F.; Stanley, C.; Opperstrup, T.; Neale, C.; Schumacher, S. K.; Ahn, D.; Herbein, S.; Carpenter, T. S.; Gnanakaran, S.; Bremer, P.-T.; Glosli, J. N.; Lightstone, F. C.; Ingólfsson, H. I. In *Generalizable Coordination of Large Multiscale Ensembles: Challenges and Learnings at Scale*, International Conference for High Performance Computing, Networking, Storage and Analysis; ACM, 2021; p 10.
- (79) Stanton, L.; Opperstrup, T.; Carpenter, T. C.; Ingólfsson, H. I.; Surh, M.; Lightstone, F. C.; Glosli, J. Dynamic Density Functional Theory of Multicomponent Cellular Membranes. *Phys. Rev. Res.* **2022**, *5*, No. 013080.
- (80) Bhatia, H.; Thiagarajan, J. J.; Anirudh, R.; Jayram, T. S.; Opperstrup, T.; Ingólfsson, H. I.; Lightstone, F. C.; Bremer, P.-t. A Biology-Informed Similarity Metric for Simulated Patches of Human Cell Membrane. *Mach. Learn.: Sci. Technol.* **2022**, *3*, 2632–2153.
- (81) Marconi, U. M. B.; Tarazona, P. Dynamic density functional theory of fluids. *J. Chem. Phys.* **1999**, *110*, 8032–8044.
- (82) Monticelli, L.; Kandasamy, S. K.; Periole, X.; Larson, R. G.; Tieleman, D. P.; Marrink, S.-J. The MARTINI Coarse-Grained Force Field: Extension to Proteins. *J. Chem. Theory Comput.* **2008**, *4*, 819–834.
- (83) de Jong, D. H.; Singh, G.; Bennett, W. F.; Arnarez, C.; Wassenaar, T. A.; Schafer, L. V.; Periole, X.; Tieleman, D. P.; Marrink, S. J. Improved Parameters for the Martini Coarse-Grained Protein Force Field. *J. Chem. Theory Comput.* **2013**, *9*, 687–697.
- (84) Wassenaar, T. A.; Ingólfsson, H. I.; Bockmann, R. A.; Tieleman, D. P.; Marrink, S. J. Computational Lipidomics with insane: A Versatile Tool for Generating Custom Membranes for Molecular Simulations. *J. Chem. Theory Comput.* **2015**, *11*, 2144–2155.
- (85) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* **2015**, *1–2*, 19–25.
- (86) Zhang, X.; Sundram, S.; Opperstrup, T.; Kokkila-Schumacher, S. I. L.; Carpenter, T. S.; Ingólfsson, H. I.; Streitz, F. H.; Lightstone, F. C.; Glosli, J. N. ddcMD: A fully GPU-accelerated molecular dynamics program for the Martini force field. *J. Chem. Phys.* **2020**, *153*, No. 045103.
- (87) Klauda, J. B.; Venable, R. M.; Freites, J. A.; O'Connor, J. W.; Tobias, D. J.; Mondragon-Ramirez, C.; Vorobyov, I.; MacKerell, A. D., Jr; Pastor, R. W. Update of the CHARMM all-atom additive force field for lipids: validation on six lipid types. *J. Phys. Chem. B* **2010**, *114*, 7830–7843.
- (88) Best, R. B.; Zhu, X.; Shim, J.; Lopes, P. E.; Mittal, J.; Feig, M.; Mackerell, A. D., Jr. Optimization of the additive CHARMM all-atom protein force field targeting improved sampling of the backbone phi, psi and side-chain chi(1) and chi(2) dihedral angles. *J. Chem. Theory Comput.* **2012**, *8*, 3257–3273.
- (89) Shirts, M. R.; Klein, C.; Swails, J. M.; Yin, J.; Gilson, M. K.; Mobley, D. L.; Case, D. A.; Zhong, E. D. Lessons learned from comparing molecular dynamics engines on the SAMPL5 dataset. *J. Comput Aided Mol Des* **2017**, *31*, 147–161.
- (90) Case, D. A.; Ben-Shalom, I. Y.; Brozell, S. R.; Cerutti, D. S.; Cheatham, T. E. C., III; Cruzeiro, V. W. D.; Darden, T. A.; Duke, R. E.; Ghoreishi, D.; Gilson, M. K.; Gohlke, H.; Goetz, A. W.; Greene, D.; Harris, R.; Homeyer, N.; Huang, Y.; Izadi, S.; Kovalenko, A.; Kurtzman, T.; Lee, T. S.; LeGrand, S.; Li, P.; Lin, C.; Liu, J.; Luchko, T.; Luo, R.; Mermelstein, D. J.; Merz, K. M.; Miao, Y.; Monard, G.; Nguyen, C.; Nguyen, H.; Omelyan, I.; Onufriev, A.; Pan, F.; Qi, R.; Roe, D. R.; Roitberg, A.; Sagui, C.; Schott-Verdugo, S.; Shen, J.; Simmerling, C. L.; Smith, J.; Salomon-Ferrer, R.; Swails, J.; Walker, R. C.; Wang, J.; Wei, H.; Wolf, R. M.; Wu, X.; Xiao, L.; York, D. M.; Kollman, P. A. *AMBER 2018*; University of California: San Francisco, 2018.
- (91) Touw, W. G.; Baakman, C.; Black, J.; te Beek, T. A.; Krieger, E.; Joosten, R. P.; Vriend, G. A series of PDB-related databanks for everyday needs. *Nucleic Acids Res.* **2015**, *43*, D364–368.
- (92) Kabsch, W.; Sander, C. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers* **1983**, *22*, 2577–2637.
- (93) Alessandri, R.; Souza, P. C. T.; Thallmair, S.; Melo, M. N.; de Vries, A. H.; Marrink, S. J. Pitfalls of the Martini Model. *J. Chem. Theory Comput.* **2019**, *15*, S448–S460.