

Pandemics *In Silico*: Scaling an Agent-Based Simulation on Realistic Social Contact Networks

Joy Kitson*, Ian Costello*[†], Jiangzhuo Chen[‡], Diego Jiménez[§], Stefan Hoops[‡], Henning Mortveit[‡], Esteban Meneses[¶], Jae-Seung Yeom^{||}, Madhav V. Marathe[‡], Abhinav Bhatele*[‡]

*Department of Computer Science
University of Maryland
College Park, USA

[‡]Biocomplexity Institute and Initiative
University of Virginia
Charlottesville, USA

[§]Max Planck Computing and Data Facility
Garching, Germany

[†]Google, Inc
Mountain View, USA

[¶]National Advanced Computing Collaboratory
National High Technology Center
San José, Costa Rica

^{||}Center for Applied Scientific Computing
Lawrence Livermore National Laboratory
Livermore, USA

Abstract—Preventing the spread of infectious diseases requires implementing interventions at various levels of government and evaluating the potential impact and efficacy of those preemptive measures. Agent-based modeling can be used for detailed studies of epidemic diffusion and possible interventions. Modeling of epidemic diffusion in large social contact networks requires the use of parallel algorithms and resources. In this work, we present *Loimos*, a scalable parallel framework for simulating epidemic diffusion. *Loimos* uses a hybrid of time-stepping and discrete-event simulation to model disease spread, and is implemented on top of an asynchronous, many-task runtime. We demonstrate that *Loimos* is able to achieve significant speedups while scaling to large core counts. In particular, *Loimos* is able to simulate 200 days of a COVID-19 outbreak on a digital twin of California in about 42 seconds, for an average of 4.6 billion traversed edges per second (TEPS), using 4096 cores on Perlmutter at NERSC.

Index Terms—high performance computing, agent-based modeling, epidemiology, social network graphs

I. INTRODUCTION

The COVID-19 pandemic has demonstrated that while we have made significant progress in controlling infectious disease outbreaks, such outbreaks will continue to pose a threat. Computational models played a critical role during the COVID-19 pandemic in various response efforts – to forecast the trajectory of the pandemic, evaluate various what-if scenarios, and support economic and logistical planning problems such as vaccine allocation and distribution [1]–[4]. Several challenges have also emerged as a result of these efforts, including: (1) running these models in real time, (2) scaling models to larger regions and incorporating a range of social, behavioral, economic and immunological considerations, and (3) managing limited data and the resulting uncertainty regarding conditions on the ground.

Traditional modeling techniques for the spread of infectious diseases often rely on coupled rate equations – systems of differential equations relating the number or fractions of people who are susceptible, exposed, infected, and recovered (SEIR) [5]. While such approaches are effective at capturing

statistical trends like the rate at which people are infected, they fail to fully capture the complexity of human social networks and the interactions that serve as a mechanism for disease spread. As many interventions function by changing this network of interactions, their impact on a disease’s spread can only be modeled indirectly under this paradigm.

In contrast, agent-based models simulate the epidemic process on social contact networks that capture the dynamics of human interactions. While more flexible, this approach is much more computationally expensive, requiring agent-based models to be parallel and highly scalable. The first reason for this is that it is important to be able to simulate epidemic dynamics over national and global scale networks. A realistic social contact network for the U.S. would have ~335 million agents, and a global scale network would have ~8 billion agents. Second, interventions are an important component of any epidemic simulation that seeks to study the impact of government planning and response. However, interventions complicate interaction networks that are already highly irregular by allowing them to change over time, which can slow simulations down considerably.

The stochasticity of this class of agent-based models also poses challenges when it comes to evaluating a scientific workload. Complex experiments that study several possible scenarios require many runs, whether for sensitivity analysis, uncertainty analysis, or comparing model projections for a wide range of different scenarios. A typical design with 20 scenarios, each with 100 slight perturbations on model parameters for uncertainty analysis, each with 30 replicates to account for the stochasticity of the model, can yield 60,000 simulation experiments. Performing this many experiments in a short amount of time requires a highly scalable code. Running such large workflows means pushing the limits of performance, and motivates the development of a parallel program capable of scaling to meet these demands.

Design and implementation of parallel simulations for contagion modeling is challenging for two main reasons: (1)

the underlying social contact networks on which infectious diseases spread are highly unstructured (see [6], [7] for an in-depth discussion), and (2) the dynamics over such networks are stochastic in nature; the nodes which participate in the spreading process may differ. This complicates partitioning and load balancing, as one cannot predict the inter-process communication and workload on each processor *a priori*.

Our primary objective in this work is to develop a scalable, parallel simulation framework for modeling contagion processes over large relational and time-varying networked systems. Toward this end, we present Loimos, a highly scalable parallel code for agent-based simulations on realistic social contact networks, written on top of Charm++, an asynchronous, many-task runtime system. Loimos utilizes a combination of discrete-event simulation (DES) and time stepping to model the spread of diseases on these networks.

Our key contributions are as follows:

- Designing and implementing a parallel agent-based simulator for modeling contagion processes and intervention scenarios at an individual level.
- Identifying three major performance bottlenecks in Loimos and introducing optimizations addressing them.
- Evaluating the scalability of the code on an HPC platform both in strong and weak scaling scenarios.
- Validating the simulator against EpiHiper [12], an existing model used by the CDC COVID-19 scenario modeling hub [13].

II. RELATED WORK

Bissett et al. [14] identify five components of agent-based techniques for modeling epidemics: (1) a theory component, (2) synthetic population construction, (3) social contact network generation from such synthetic populations, (4) construction of idealized social contact networks, and (5) simulation of epidemic diffusion across both types of contact networks. We focus on this last component, as the other components generally represent one-time costs for an outbreak response.

Several recent publications focus on modeling the spread of COVID-19. Many of these are national or regional compartmental models built using data from outbreaks in the simulated area, and account for interventions in different ways. The SIQR (Susceptible Infectious Quarantined Recovered) [15] and SIDARTHE (Susceptible Infected Diagnosed Ailing Recognized Threatened Healed and Extinct) [16] models have been used to simulate the progression of the pandemic in India and Italy, respectively, using new disease states and adjustments to the values of disease parameters to capture the impact of interventions. An age-segmented SIRD (Susceptible Infectious Recovered Dead) model using synthetic contact matrices for interventions [17], and a SIRD model using an optimization algorithm to estimate the infection rate based on empirical data [18] have also been used to model COVID-19. Metapopulation models are commonly used to capture international disease spread. These models segment the simulated population into subpopulations representing countries or regions and build a compartmental model for each

subpopulation with flows connecting them. One such model has been used to estimate the impact of travel restrictions on the early spread of COVID-19 [19]. There are also some efforts to build small-scale agent-based models to simulate the spread of the virus within small communities or within single buildings. Although these models range in complexity – COVID-ABS [20] incorporates both economic and epidemiological models within a single simulation whereas Cuevas’s model [21] of spread within a building only requires two rules to guide its agents’ behavior – most of these models are quite small, only simulating a few hundred agents.

Several parallel agent-based epidemic simulators have been developed for HPC systems, including several that operate on national scales. However, any performance comparison between existing models is hampered by the lack of detailed information on the parameters and HPC systems used in the runs. No single ground truth dataset exists to test raw computational speed in this domain, so we instead seek to compare simulations which operate on a similar scale, namely that of the population of the United States. Note that most prior work considers 280-290 million agent populations to be U.S.-scale, as shown in Table I.

With this limitation in mind, there are a number of approaches to developing high-performance agent-based disease simulations. The Framework for Replication of Epidemiological Dynamics (FRED) [8] is an OpenMP based simulation that uses US census data to model disease spread. FRED’s disease models are fixed to a configurable SEIR model (susceptible, exposed, infectious, and recovered), resulting in much less flexibility in terms of input, compared to codes which support a tunable arbitrary disease model, such as Loimos. Seal et al. [22] implement an agent-based model involving a generalization of Conway’s Game of Life, instead of dynamics on realistic contact networks. However, this simulation is notable for its use of GPU offloading, which most of the simulations surveyed – along with Loimos – do not support. Germann et al. [9] develop EpiCast by adapting the SPaSM molecular dynamics simulation, using cells as an analog for communities and particles as an analog for individual agents. EpiCast lies halfway between a meta-population model with its spacial distributed interaction groups and a fully-fledged agent-based model, placing each agent in multiple interaction groups at once, representing where they live, work, and travel. Parker et al. [23] introduce a novel approach that models how human behavior changes due to a pandemic (e.g. increased social distancing) but is limited to an SEIR model and requires creating new populations to model different sets of behaviors. Machin et al. [12] and the EpiHiper team [7] present an agent-based simulator embedded in an end-to-end pipeline which runs the gauntlet from model calibration to simulation output analysis. While this represents a mature production simulation, used by the CDC COVID-19 scenario modeling hub [13], their work focuses more on the orchestration of the overall pipeline than the optimization of individual application runs.

Perumalla et al. [24], and the EpiSimdemics team [11], [25], [26] perform some of the fastest epidemic simulations.

TABLE I
SUMMARY OF RESULTS FOR PRIOR AGENT-BASED EPIDEMIC SIMULATORS. EPIHIPER RESULT IS AN ENSEMBLE SPANNING TWO SYSTEMS.

Simulator	No. of Agents	No. of Days Simulated	Machine	No. of Cores	Runtime	Runtime/Day
FRED [8]	289 million	Unknown	Blacklight at PSC	16	4 h	Unknown
EpiCast [9]	281 million	180	2.4 GHz Intel Xeon	256	8-12 h	160 s
EpiHiper* [10]	288 million	72,000	Bridges-2 at PSC, Rivanna at UVA	6,400 1,200	32 h 42 m	0.141 s
EpiSimdemics [11]	280.4 million	180	Blue Waters at NCSA	655,360	10.41 s	0.0578 s

EpiSimdemics shows impressive scaling, presenting a zip-code based partitioning scheme similar to the one employed by Loimos. They show orders of magnitude difference in performance compared to previous work. Table I summarizes the performance of some of the more performant models discussed above.

III. ALGORITHM FOR CONTAGION DIFFUSION

We develop our epidemic simulator, Loimos, by using a combination of network theory, discrete event simulations, and agent-based modeling. We model both individuals in the population and interactions between pairs of these *agents*. This allows us to simulate the dynamics of epidemic diffusion on a sufficiently granular level to directly model the changing dynamics of disease spread due to a variety of public health intervention strategies.

A. Serial Algorithm

In order to expose parallelism across people and locations, we iterate over discrete time steps. Each time step requires:

- 1) Identifying overlapping visits to each location.
- 2) Calculating the likelihood that each overlap resulted in an infection, then determine which infections occur.
- 3) Updating each person's disease state to reflect any infections and the progression of the disease.

Disease Model and Finite State Automaton: We assign each person, p , a disease state, x_p , managed using a finite state automaton (FSA) which specifies how – once infected – people move through various disease states representing different stages in the progression of a disease. We require that each state have an associated susceptibility, $\sigma(x_p)$, and infectivity, $\iota(x_p)$, which we use to determine how to treat them in the simulation; we consider p to be *susceptible* when $\sigma(x_p) > 0$ and *infectious* when $\iota(x_p) > 0$. Transitions between states are stochastic both in terms of the state transitioned to and how long a person remains in a given state. Throughout this paper we use an expanded version of the Susceptible, Exposed, Infectious, and Recovered (SEIR) model [27].

Discrete-event Simulation: The discrete event simulation (DES) determines which people are at a location, ℓ , at the same time and for how long, based on the set of visits to ℓ during a given simulation day. The DES splits each visit into two events – an arrival and a departure – and orders them in a queue, Q_ℓ , based on when they occur. The DES then processes

events from this queue, as shown in Algorithm 1, identifying all of the corresponding exposure events.

While processing an event, e , we maintain lists of all susceptible and infectious people currently at the location, V_s and V_i respectively. When processing an event, we use the disease state, x_{p_e} , of the corresponding person, p_e , to select the appropriate list, V , skipping immune and exposed people (lines 3-5), adding a person when they arrive (line 7) and removing them when they depart (line 9). When a person leaves, we consider everyone in the opposite list, V' , as potential contacts with a fixed probability, c_ℓ , based on the location (See (1); lines 10-12). If a contact occurs between an infectious person and a susceptible person, we store the propensity of the resulting exposure to cause an infection (See (2); lines 13-15).

Algorithm 1: Computing exposures at a location, ℓ

```

1 ComputeExposures (event queue  $Q_\ell$ , disease states  $X_\ell$ ):
2 foreach event  $e \in Q_\ell$  with corresponding visitor  $p_e$  do
3   if  $p_e$  is susceptible then  $V = V_s, V' = V_i$ ;
4   else if  $p_e$  is infectious then  $V = V_i, V' = V_s$ ;
5   else continue;
6   if  $e$  is an arrival then
7     Add  $p_e$  to visitor list  $V$  corresponding to  $p_e$ 's disease state;
8   else
9     Remove  $p_e$  from corresponding visitor list  $V$ ;
10    foreach  $p'$  currently in the opposite visitor list  $V'$  do
11       $c_\ell = \text{ContactProbability}(\ell)$ ;
12      if  $p_e$  and  $p'$  make contact with probability  $c_\ell$  then
13        Compute propensity of infection for  $p_e$  and  $p'$ 
14        during the period of co-occupancy  $T$  using  $X_\ell$ ;
15        if  $p_e$  is susceptible then add exposure to list  $E_{p_e}$ ;
16        else add exposure to list  $E_{p'}$ ;
17      end
18    end
19 end

```

Contact Model: The contact model operates at each location, ℓ , independently, and determines whether or not a pair of overlapping visits to ℓ result in a contact. Here we adopt the *min/max/ α* model formulation of Chen et al. [28]. In this model, we compute the *contact probability*, c_ℓ , for any pair of people simultaneously present at ℓ as a function of its maximum occupancy, N_ℓ . This serves as a proxy for its size. In order to compute c_ℓ , we select a minimum value, A , where if $N_\ell < A$ everyone will make contact, and a maximum value, α , where if $N_\ell > \alpha$ someone visiting at the peak occupancy of the location will make about B contacts; a person visiting a

given location during peak occupancy should expect to make between A and B contacts during that visit. p_ℓ is given by

$$c_\ell = \min\left\{1, \left[\frac{A(B-A)(1 - e^{-N_\ell/\alpha})}{N_\ell - 1} \right] \right\} \quad (1)$$

↑ contact prob ↑ min contacts ↑ threshold
↑ max contacts ↑ max occupancy

for $N_\ell \geq 2$.

We use the values $A = 5$, $B = 40$ and $\alpha = 1000$ below, based on social contact patterns in the POLYMOD data [29].

Transmission Model: The transmission model determines whether or not a given contact between susceptible and infectious individuals (p_i and p_j , respectively) results in disease transmission. As in Chen et al. [28], transmission probabilities depend on a personal and a disease state susceptibility, $\beta_\sigma(p_i)$ and $\sigma(x_{p_i})$ respectively, and infectivity, $\beta_\iota(p_j)$ and $\iota(x_{p_j})$, the contact duration, T , and a transmissibility, τ .

The infection propensity, ρ , of such a contact is given by

$$\rho(p_i, p_j, T) = T \cdot \tau \cdot \beta_\sigma(p_i) \cdot \sigma(x_{p_i}) \cdot \beta_\iota(p_j) \cdot \iota(x_{p_j}) \quad (2)$$

↑ overlap duration ↑ susceptibility ↑ infectivity
↑ transmissibility

where τ is a tuning parameter proportional to the likelihood of infection from a one-second contact with an infectious agent.

In order to determine whether or not a transition occurs for a given susceptible person, p_i , at the end of a time step, we sum the propensities from all m contacts p_i had with infectious people during the time step

$$A(p_i) = \sum_{j=0}^{m-1} \rho(p_i, p_j, T_{i,j}) \quad (3)$$

↑ see (2) ↑ overlap duration
↑ susceptible ↑ infectious

where $T_{i,j}$ is the duration of the p_i 's j -th contact. We then sample $a \sim -\log(\text{uniform}(0, 1))/A$ and infect p_i if $a < 1$.

Intervention Model: Interventions have three main components: a trigger, a selector, and an action. The trigger activates the intervention at the end of a time step in response to some global condition, such as the population passing a case threshold. Once the intervention is active, its selector determines which people or locations it will apply its action to, based on their individual attributes and health states. This action will then either (1) add or remove some edges incident to the person or location in question or (2) adjust the values of attribute(s) of the person or location, which may be either used in later interventions or transmission propensity calculations. Most actions can also be reversed when the intervention no longer applies to the person or location in question.

B. Parallel Algorithm

The parallel algorithm, shown in Algorithm 2, operates on a bipartite graph – with people and locations as nodes, and weekly visit schedules as edges – along with an assignment of people and locations to partitions. This algorithm is based on that originally proposed by Yeom et al. [25], [26].

Algorithm 2: Parallel control flow in Loimos

```

1 Partition  $P$  into people partitions  $\mathbb{P} = \{P_i\}$ ;
2 Partition  $L$  into location partitions  $\mathbb{L} = \{L_j\}$ ;
3 foreach person partition  $P_i \in \mathbb{P}$  parallel
4   foreach location partition  $L_j \in \mathbb{L}$  do
5     Compute the set  $V_{i,j}$  of people on  $P_i$  who visit some
       location on  $L_j$ ;
6   end
7 end
8 foreach simulation day  $d \in \{1, \dots, d_{max}\}$  do
9   foreach person partition  $P_i \in \mathbb{P}$  do
10    foreach person  $p \in V_{i,j}$  do
11      Send disease state update message  $(p, x_p)$ ;
12    end
13  end
14  foreach location partition  $L_j \in \mathbb{L}$  parallel
15    foreach disease state update message  $(p, x_p)$  to  $L_j$  do
16      Store  $p$ 's disease state,  $x_p$ , in  $X_\ell$ ;
17    end
18    foreach location  $\ell$  on  $L_j$  do
19      foreach visit  $v$  to  $\ell$  do
20        Put an arrival and departure event into  $Q_\ell$ ;
21      end
22      Reorder  $Q_\ell$  by the time of event in ascending order;
23      ComputeExposures( $Q_\ell, X_\ell$ );
24      foreach susceptible person  $p$  with exposure(s) at  $\ell$  do
25        Send exposure message  $m$  to  $p$ 's person partition;
26      end
27    end
28  end
29  foreach person partition  $P_i \in \mathbb{P}$  parallel
30    foreach person  $p$  on  $P_i$  do
31      foreach exposure message  $m$  destined for  $p$  do
32        Put the exposures into the exposure list  $E_p$ ;
33      end
34      if IsInfected( $p, x_p, E_p$ ) then
35        Update  $p$ 's disease state  $x_p$ ;
36      end
37    end
38  end
39  Evaluate intervention triggers;
40 end

```

Prior to the main loop, we partition the person and location data (P and L , respectively; lines 1-2), and identify which people, $V_{i,j}$, from each person partition, L_i , visit each location partition, L_j , (lines 3-7).

On each simulated day, each person partition sends a message with the current disease state, x_p , of each person, p , it holds to each location partition they visit (lines 9-13). This update is performed in lieu of the full visit message exchange from Yeom et al. [25], and is discussed in more depth in Section V-D. Once all the update messages have been received and the updated disease state stored in X_ℓ (lines 15-17), arrival and departure events for each visit are created and placed in a time-ordered queue, Q_ℓ (lines 19-22). Locations currently selected by an active intervention may have altered visits. Each process performs a DES for each of its locations, as described in Algorithm 1, determines whether each pair of people whose visits overlap come into contact, and then computes the propensity of any contacts based on their disease states. Once these calculations have been performed, we send exposure messages to people with at least one exposure (lines 24-26). Each person's exposure messages are then gathered

in E_p and processed to determine whether or not a given exposed person was infected (line 34). If a susceptible person is infected or an infected person makes a timed transition, their disease state will be updated to reflect this at the end of the simulation day (line 35). Finally, we evaluate the triggers of any interventions deployed in the current scenario (line 39).

For simplicity, we primarily refer to the three phases of the main loop in later performance analysis, namely (1) person state communication (**PSC**; lines 9-17), (2) exposure computation and communication (**ECC**; lines 18-33), and (3) person state update (**PSU**; lines 34-36).

IV. IMPLEMENTATION IN LOIMOS

We present a parallel epidemic modeling framework, Loimos, that implements the parallel algorithm described above. Loimos is written on top of the Charm++ [30], [31] parallel runtime. In this section, we discuss salient details regarding the design of its parallel implementation.

A. Inputs to the Simulator

Three core components define an epidemic simulation:

- 1) A population, consisting of people, locations, and visits
- 2) A disease, represented as an FSA (see Section III-A)
- 3) An (optional) set of interventions, capable of modifying visit schedules and disease transmission likelihoods

We use two different types of populations for the simulations described in this paper: realistic *digital twin* populations mirroring several U.S. states, and purely synthetic populations.

Generating Realistic Populations: We generate these realistic datasets from a range of data sources through an extension of the pipeline developed by Chen et al [28]. For a given state, we begin by constructing a collection of people with demographics (including age, gender, and occupation codes [32]) and partitioning these people into households. We refine these partitions at a block group level through iterative proportional fitting to match the demographic distributions found in American Community Survey (ACS) data [33].

Next, we assign each person a set of activities, using National Household Travel Survey data through random forest methods conditioned on demographics and calibrated against time-use surveys [34]. We then construct home and activity locations, integrating building [35], and school [36] data. Finally, we assign people to home locations, and activities to activity locations. We constrain this assignment to match ACS commute flow and demographic data (ex: to ensure teachers work at schools). We use these techniques to generate the datasets shown in Table II.

Purely Synthetic Population Generation: We generate our purely synthetic datasets on the fly, using a structured grid of locations to maintain epidemic locality. Given an average number of visits per person per day, λ_{visits} , we assign each person, p , to a home location uniformly throughout the grid, then generate $n \sim \text{Pois}(\lambda_{\text{visits}})$ visits for each person on each day, where $\text{Pois}(\lambda)$ is the Poisson distribution with expected rate λ . Here, the i -th visit is to a random location

TABLE II
DIGITAL TWIN DATASETS USED FOR STRONG SCALING STUDIES.
INTERACTION AND VISIT COUNTS GIVEN PER DAY.

Dataset Name	# Interactions	# Visits	# People	# Locations
Arkansas (AR)	63.65M	12.81M	2.749M	13.17M
Iowa (IA)	68.41M	14.24M	2.967M	13.68M
Michigan (MI)	226.5M	44.39M	9.342M	16.33M
New York (NY)	525.6M	88.28M	18.11M	17.97M
California (CA)	955.7M	164.6M	35.51M	24.49M

$d_i \sim \text{Pois}(\lambda_{\text{hops}})$ hops away in the grid from p 's home location. We use the average visits per person in the CA data, $\lambda_{\text{visits}} = 4.6$, and $\lambda_{\text{hops}} = 5.2$ to generate the flexible datasets shown in Table III.

TABLE III
PURELY SYNTHETIC POPULATIONS USED IN WEAK SCALING STUDIES.

Relative Size	# People/core	# Locations/core
1x	280k	70k
2x	560k	140k
4x	1.120M	280k

B. Task-based Decomposition

Loimos is implemented in Charm++ [30], [31], a parallel language focused on asynchronous, message-driven programs. When writing Charm++ programs, the programmer organizes code and data into a combined object called a *chare*. Chares are in turn organized into *chare arrays*, indexable collections of chares. When the code is run, the Charm++ runtime is responsible for assigning chares to processors and for scheduling the execution of code on the various chares assigned to a given processor. This code is usually run in response to a message received from another chare. In Loimos, we use two chare arrays: one for people and one for locations, with each chare containing a partition of the appropriate data.

The other main Charm++ object we use in Loimos is a *node group*. Node groups have a single instance for each node the program is run on, allowing us to avoid keeping redundant copies of shared information in memory on one node, while also minimizing inter-node communication.

C. Implementation of Different Models

We implement the various models described in Section III-A modularly to increase the flexibility of the codebase.

Disease Model: At the start of a simulation, we load the FSA representing the simulated disease from an input file which specifies its states and transitions. A Charm++ node group then stores a single copy of this information on each node, in order to afford each person and location chare efficient read-only access to these data. We store other read-only data describing the input scenario in a similar fashion.

When initializing the simulation, people begin in one of (potentially) several entry disease states, as determined by their individual attributes, such as age, specified in the input

file. They remain in this state until they are infected by an infectious person or chosen to seed the outbreak. For this work, we select a small sample of people to infect during the first few days of the simulation. For the runs in this paper, we infect 2 people per day, chosen uniformly at random, for the first 10 days of the simulation, and use a FSM representing COVID-19 with 5 age-based entry states, each with 18 distinct reachable states.

Discrete-event Simulation: On each simulation day, we start the DES after the conclusion of the person state communication (PSC), when updates to interventions (and thus visit schedules) are also guaranteed to have completed. We use Charm++’s quiescence detection mechanism¹ – a soft barrier which ensures no messages are in flight or being processed before continuing – to determine when this phase is done, as the number of visits a given location will receive is both not known *a priori* and non-deterministic in the presence of any intervention which changes visit schedules, such as school closures. Once triggered, each location chare independently executes the DES for each of its assigned locations.

In implementing the DES algorithm, we made three key optimizations: (1) we only keep track of co-occupancy – and thus interactions – between susceptible people and infectious people, (2) we only ever send exposure messages to susceptible people who had at least one exposure occur during a time step, and (3) exposure messages are sent as soon as we finish processing a susceptible individual’s departure event. Note that we are able to make this first change without affecting the results of the simulation because only contact between a susceptible person and an infectious person constitutes an exposure (and thus can cause an infection). The second optimization is especially helpful as in most iterations only a small fraction of visits will result in an exposure. The third optimization allows us to significantly overlap computation and communication during this phase.

Contact Model: We implement two different types of contact model in our code. The first is the *min/max/α* model (see Section III-A). Since this requires knowledge of the maximum occupancy of each location, we use a pre-processing script to compute this based on the visit schedules file and save each location’s maximum occupancy to the locations file. At the start of a run, we read in this value for each location, compute the appropriate contact probability, and store it as a new location attribute. Since this cannot be computed in advance for the purely synthetic datasets, we also provide a second contact model with a single, global contact probability.

Transmission Model: After each exposure is identified, we compute the corresponding propensity and batch it with the other exposure messages for the susceptible person involved. We send the exposure messages for each person immediately after processing the departure event for their visit in the DES. We do not compute actual infections until after the DES is

complete and all exposure messages received. We again use Charm++ quiescence detection to determine when this has occurred, as the number of exposure messages each person chare will receive is highly non-deterministic, depending on the mixture of infectious and susceptible people as well as the contact model. Once all messages are received, each person chare sums the propensities for each person in their partition in order to identify infections.

Intervention Model: Similar to the disease model, the specifications for the interventions to be used in a given run are provided in an input file and stored on a node group. Unlike the disease model, we store and update some state for each intervention via the main chare over the course of the simulation. In particular, at the end of each simulation day, we perform a reduction across all person chares to compute the number of total infectious people, and pass this value along with the current day to the triggers for each specified intervention to determine which interventions should be active. The ids for each active location-based intervention are then passed to all location chares. These chares then access the local copy of the intervention objects on their node and filter the locations assigned to that chare using the selector for the intervention, and apply the action to the relevant locations. We use a separate class for each intervention, each of which extends a shared interface with methods for testing whether a location should be selected, making arbitrary changes to a location’s state via an action – including changing its visit schedule – and undoing the changes wrought by the action. Some interventions, like vaccination, have a trivial undo method as the changes persist after the initial intervention ends. We use a similar scheme for person-based interventions.

V. ADDITIONAL PERFORMANCE CONSIDERATIONS

After implementing Loimos’ core features, we explored several avenues to further optimize our code’s performance. The four directions we found most beneficial are: (1) considering different combinations of processes and threads per node, (2) static load balancing, (3) short-circuit evaluation of the DES at each location, and (4) storing visit data on location chares.

A. Impact of Using Processes vs. Threads

The Charm++ suite includes support for different machine layers as well as abstractions that enable the programmer to adapt to the underlying hardware to improve performance and scalability. All these components make Charm++ codes highly tunable. In particular, we are interested in analyzing how running Loimos with symmetric multiprocessing (SMP) support in Charm++ performs relative to the non-SMP alternative.

Enabling the Charm++ SMP mode is analogous to adding OpenMP or another threaded programming model to an MPI code, in that it runs multiple threads per process instead of the usual single thread. Users can specify how many worker threads are spawned per process, along with an optional mapping from threads to cores, but one thread per process is always dedicated to communication. This communication thread

¹We chose quiescence detection over a similar mechanism in Charm++, competition detection, which is incompatible with other Charm++ features.

manages inter-node messages whereas intra-node communication is managed through the shared memory address space common to all threads on the same node. The requirement of allocating one communication thread per process could be a disadvantage for compute-intensive applications since compute cores have to be sacrificed. However, for communication-intensive applications, the use of a dedicated communication thread to manage messaging might lead to better performance.

We compare six different ratios of processes to nodes, using 8 processes per node (p/n) with 15 or 14 worker threads per process (t/p) and 16 p/n with 7 or 6 t/p in SMP mode, and 126 processes per node in non-SMP mode. We performed a strong scaling experiment, running three replicates of each configuration for 200 days on the MI data (see Table II) with all optimizations enabled. All scaling experiments in this paper were conducted on the Perlmutter Cluster at NERSC, a HPE Cray EX cluster with two AMD EPYC 7763 CPUs per node, each with 64 cores, and a HPE Slingshot 11 interconnect [37].

Figure 1 shows the experimental results of this SMP vs non-SMP comparison. The non-SMP configuration consistently out-performs the SMP configurations, achieving a speedup of about $3.25\times$ when going from 128 cores to 4096. Note that some SMP configurations suffer from fatal runtime errors on larger core counts. On 4096 cores, the 126 p/n configuration is $1.19\times$ faster than the best SMP configuration, with 16 p/n and 6 t/p. As a result, we use the 126 p/n non-SMP configuration for all runs described in subsequent sections.

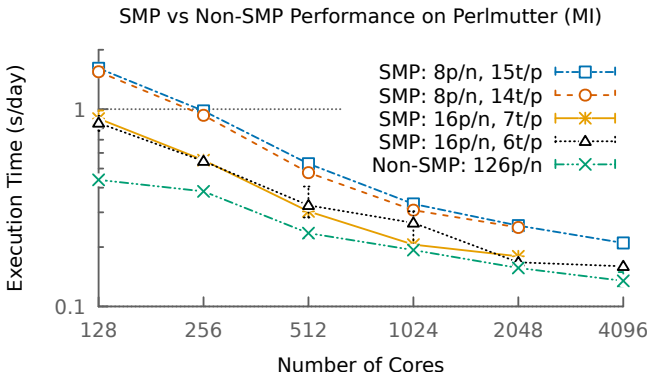


Fig. 1. Performance comparison of different SMP configurations, pairs of processes per node (p/n) and worker threads per process (t/p) vs. a non-SMP configuration using one process per core (126 p/n), with the latter performing best across all core counts. One chare per thread used in all cases. Times averaged over three runs, with extrema in error bars.

B. Static Load Balancing

During initial runs of Loimos on the realistic datasets, we found that some processes ran much slower than others, as shown in the left of Figure 2. We set out to minimize this load imbalance by improving the assignment of locations and people to chares.

Toward this end, we present a simple static partitioning scheme meant to preserve geographical locality. First, we sort all locations in the population by the id of their state, county,

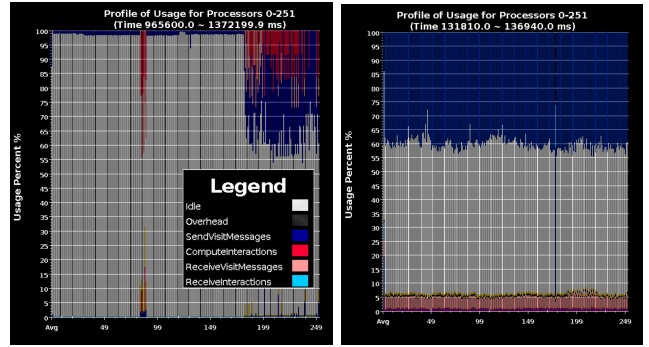


Fig. 2. Processor usage in three iterations of Loimos on two Perlmutter nodes with MI data shows much more idle time with no optimizations (left), than with static load balancing only (right).

census tract, and census block group, in that order. This is intended to ensure that nearby locations are placed on the same chare if possible. Using the number of visits to a location, λ_j , as a proxy for its load, we compute the average load per chare, Λ , as the ratio of visits to location chares, for a given number of chares. If any locations have a load $\lambda_j > \Lambda$, we assign them their own chare and recompute Λ for the remaining chares, until no such locations remain. We then compute the cumulative load for each location, fixing $\lambda_j = \Lambda$ for those heavy locations already assigned to a chare, and assign each location to chare $\lfloor (\sum_{i=0}^{j-1} \lambda_i) / \Lambda \rfloor$.

We then partition the person data, identifying the location chare, L_j , containing a person’s home location and placing them on person chare P_j . Applying this scheme drastically decreases the idle time on most processes, as shown in the right of Figure 2, although a much lower degree of imbalance still remains. Figure 5 shows how Loimos’s performance improves (no-opts vs static) when using this load balancing scheme. This leads to a speedup of $5.55\times$ and $3.55\times$ for the MI data on 128 and 4096 cores, respectively.

C. Optimizing the Interaction Computation

Our next optimization was inspired by observing significant variations in the time we spent in one of the three main simulation phases over the course of a run. We initially expected the time spent in the exposure computation and communication (ECC) phase to roughly track the number of infections, due to our use of separate queues for infectious and susceptible visitors. Instead, we observed that ECC dominates the runtime for the first half of the simulation, but falls after the peak of the infection curve has passed, as shown in Figure 3. This corresponds to an increasing number of immune people in the population, who are ignored when identifying exposures.

As a result, we realized that we spend significant time computing exposures for every location even when no infectious individuals are present, as we still process all susceptible arrival and departure events in the DES. In order to avoid this, we add a check to skip the DES entirely on locations with no infectious people visiting them on a given day. As shown in Figure 3, after implementing this “short-circuit”




Fig. 3. A 200-day simulation of Loimos on two Perlmutter nodes with MI data: time spent computing and communicating exposures is much greater without (static t_{ECC}) than with (static+sc t_{ECC}) short circuit evaluation of interactions, and corresponds to infection counts (infections) in the latter case.

computation of the DES, the ECC runtime corresponds more closely to the number of infectious people, peaking near the first inflection point of the epidemic curve. Figure 5 shows the benefits of using this scheme (static+sc) on top of the previous optimization (static), which results in about a $1.67\times$ and $1.21\times$ speedup for the MI data on 128 and 4096 cores.

D. Storing Visits on Location Chares

Finally, we had observed that the runtime of the simulation was dominated by sending visit data from person to location chares, which initially began each iteration, as shown in the left of Figure 4. This was due to the fact that the visit data was stored on person chares, as in Yeom et al. [38], despite primarily being used in the DES performed on location chares.

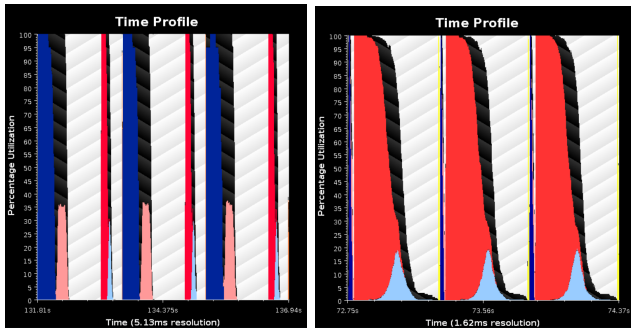


Fig. 4. Breakdown of time spent in three iterations of Loimos on two Perlmutter nodes with MI data shows we spend most of our time communicating visits (dark blue and light pink) with only static load balancing and short circuit evaluation of interactions (left), but negligible time doing so when storing visits on location chares (right).

We hypothesized that we could reduce the amount of communication by storing the visit data where it is used and only communicating data that changes between iterations, namely updated disease states and modified per-person susceptibility and infectivity values. Visit schedule changes resulting from interventions could then be evaluated on location chares based on a cache containing the relevant person state data. After

implementing this change, we found we spend minimal time in the new person state communication (PSC) phase, as shown in the right of Figure 4, which previously dominated simulation runtimes. As a side effect, however, we spend more time evaluating the DES, as we can no longer overlap the queuing of arrival and departure events with communicating visit data. Figure 5 shows the overall effect of this optimization (static+sc+loc-visits). This results in a $2.46\times$ and $1.34\times$ speedup on 128 and 4096 cores.

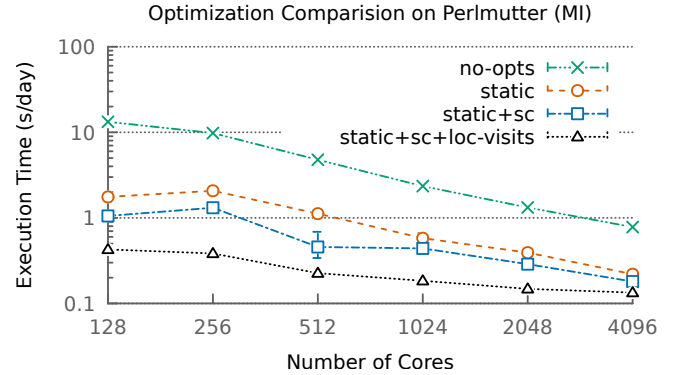


Fig. 5. Performance comparison of combinations of different performance optimizations: (1) static load balancing (static), (2) short circuit evaluation of interactions (sc), and (3) storing visit data on location chares (loc-visits), with each added optimization reducing runtimes. Execution times averaged over three runs, extrema shown in error bars.

When all optimizations are combined together, we achieve a significant $31.03\times$ speedup on 128 cores and a $5.83\times$ speedup on 4096 cores compared to the baseline with no optimizations.

VI. EXPERIMENTAL SETUP

We begin by performing extensive scalability studies using the 128 cores per node non-SMP configuration (See Section V-A) on Perlmutter. All scaling experiments were run with Protobuf version 3.21.12 and Charm version 7.0.0. All scaling runs used the same random seed, and thus had identical epidemiological results. Values shown were the average of five replicates, with the error bars representing the minimum and maximum runtimes. We use the realistic datasets shown in Table II for strong scaling and those in Table III for weak scaling. We additionally use a similar dataset with 5.513 million people and 2.896 million locations representing Maryland (MD) to validate Loimos’s simulation output against that of an established epidemiological simulation, EpiHiper [12]. See Section IV for a more detailed description of these datasets.

Table II describes the number of people, locations, and total visits for all these datasets. Note that interactions (person-person edges) are given on average, due to the stochasticity of the contact model. We run the realistic datasets for 200 days total, processing a total of about 21 and 191 billion interactions for the AR and CA datasets, respectively, over the course of the simulation. In order to ensure a representative workload, the transmissibility of the simulated outbreaks were tuned so




Fig. 6. Strong scaling performance of Loimos on Perlmutter for five different datasets in terms of execution time (left) and traversed edges per second (TEPS, right). Both show modest but consistent linear speedups for larger datasets. Execution times averaged over three runs, with extrema shown in error bars.

that the number of infectious people peaked about halfway through the simulations.

Next, we performed a weak-scaling experiment. We ran three fixed problem sizes per process, as shown in Table III. These datasets were generated using the on-the-fly synthetic population generation method outlined in Section IV.

We evaluated the performance of all scaling runs by calculating the average execution time per simulation day, excluding data loading and application startup time.

Lastly, we performed a distributional docking study (see [39]) to validate Loimos against EpiHiper using the MD dataset, for which we developed versions that can run in both models. For 30 runs we varied the random seed to capture the distribution of potential epidemiological outcomes. These results were then compared against runs of the existing EpiHiper simulation using the same input visit network and simple SIR disease model. Note that EpiHiper used the visit data differently than Loimos: a preprocessing script determined the list and duration of contacts for each person in the population, producing a fixed contact network which the disease then diffused over. Each EpiHiper run was on a separate contact network. In contrast, Loimos determined a person's contacts separately on each day, effectively resulting in a dynamic contact network, even in the absence of interventions. For all validation runs, the transmissibility was fixed at $\tau = 0.05$.

VII. PERFORMANCE RESULTS

We now present scaling results from benchmarking Loimos using various inputs on Perlmutter.

A. Strong Scaling Performance

In order to understand how Loimos would enable large scale simulations we perform the classical scaling analysis shown in Figure 6. For reference, EpiSimdemics takes 2.67 seconds per day on 192 cores of Blue Waters to simulate their California population [26], where Loimos takes 1.21 seconds per day on 128 cores. For the average runtime (left), all five datasets have their best performance on 4096 cores. Notably, the smallest states scale inconsistently, with

performance remaining roughly flat from 256 to 1024 and 512 to 2048 cores for the AR and IA data, respectively. The larger datasets, however, display consistent, if modest, linear scaling. In terms of traversed edges per second (TEPS, right), Loimos performs best on NY through 1024 cores, and on CA for high core counts, peaking at about 4.6 billion TEPS on 4096 cores.




Fig. 7. Breakdown of total time spent in Loimos into the person state update (PSU, negligible in all cases), exposure computation and communication (ECC, reduced by second optimization), and person state communication (PSC, negligible after last optimization) phases with the static load balancing (static), short circuit interaction computation (sc), and location chare visit data storage (loc-visits) optimizations incrementally applied.

Figure 7 shows the breakdown of total time spent in Loimos into the three simulation phase identified in Algorithm 2: (1) the person state communication (PSC, lines 6-15), (2) exposure commutation and communication (ECC, lines 17-27), and (3) the final person state updates (PSU, lines 29-30), when the different optimizations are applied. We observe that without the short circuit interaction optimization (static), PSC takes slightly more time than ECC on all core counts, with the difference growing as the core count increases, whereas with that optimization (static+sc), ECC consistently takes a fraction of the time of PSC. When we store visit data on location chares (static+sc+loc-visits), PSC takes negligible time, and ECC is somewhat slower (as queuing arrival and departure events can

no longer be overlapped with PSC). In all cases, the time spent in PSU is negligible.

B. Weak Scaling Performance

We also perform weak scaling tests to see how well Loimos handles datasets of increasing size. Figure 9 displays Loimos' relatively flat weak scaling performance up to 4096 cores. For all configurations, there is a noticeable increase in runtime from 128 to 256 cores. For all datasets, this change is relatively small – representing a 35%, 26%, and 23% slowdown for the 280k, 560k, and 1.12M people per cores configurations, respectively. The 1.12M people per core configuration shows increased variability on 1024 and 2048 cores along with increased runtime, but the runtime on 4096 cores returns to similar level to that on 512 cores.

While these scaling results represent an optimistic representation of the person-location visit graph, they exploit location-load attributes that are present in the more realistic networks. While our simulation would see significant slowdowns from purely random visits, our static load balancing scheme is designed to place people and locations together on processors such that there exists a high level of interconnectedness between these objects similar to that seen in these datasets.

VIII. VALIDATION CASE STUDY

Finally, we performed a validation case study. We sought to show that the distributions of Loimos' results corresponds to those of an established simulator, EpiHiper, with a focus on total cumulative infections and the time to reach an equilibrium state. Figure 8 shows how both simulators show similar overall disease trajectories, with outbreaks either persisting to infect a significant proportion of the population or dying out quickly. In the former case, both simulations average similar numbers of total cumulative infections – 863k for Loimos and 858k for EpiHiper – and the latter occurs rarely for both simulations – twice for Loimos and once for EpiHiper.

With respect to the time to equilibrium for the persistent outbreaks, Loimos shows more tightly clustered results than EpiHiper. This is likely a byproduct of how the two simulators handle their input networks. Since EpiHiper uses the same contact network for an entire run, differences in the chosen contact network have the potential to cause compounding differences in the simulation results. In contrast, since Loimos essentially selects a new contact network in each iteration, differences in contact networks between runs tend to be smoothed over to some extent as more networks are sampled over the course of a run, similar to how there is less variation in the average of 100 die rolls than that of a single roll.

IX. CONCLUSION

Uncontrolled spread of infectious disease is a challenging societal issue – one that requires policy makers to have the best possible tools in order to make informed decisions. Computer simulations are one such vital tool. The tight time constraints on relevant policy decisions mean that these simulations need

to be able to model large regions extremely quickly and accurately across a wide variety of counter-factual scenarios. These demands require the use of powerful supercomputing systems. Toward this end, we presented a scalable parallel simulation framework for modeling contagion processes, Loimos, and demonstrated its capabilities.

In this work, we outlined the methods we used to develop this simulation framework and to optimize it for production HPC systems. We described the models underpinning our work as well as various optimizations we have made to enable the code to scale well. We demonstrated our code's use of resources during both strong and weak scaling runs on Perlmutter at NERSC, achieving modest, but linear, strong scaling speedups and relatively flat weak scaling results. We also show how the epidemiological results of the simulation compare to an existing model. Together, these runs demonstrate the potential uses of Loimos for policy makers as a fast epidemic simulator that is robust enough to capture the effects of policy interventions.

ACKNOWLEDGMENTS

This material is based in part upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award No. DE-SC0021.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC awards DDR-ERCAP0032257 and DDR-ERCAP0029890. The authors acknowledge Research Computing at The University of Virginia for providing computational resources and technical support that have contributed to the results reported within this publication.

REFERENCES

- [1] E. Y. Cramer, Y. Huang, Y. Wang, E. L. Ray, M. Cornell, J. Bracher, others, and U. C.-. F. H. Consortium, "The united states covid-19 forecast hub dataset," *medRxiv*, 2021. [Online]. Available: <https://www.medrxiv.org/content/10.1101/2021.11.04.21265886v1>
- [2] MIDAS Network, "COVID-19 Scenario Modeling Hub," <https://covid19scenariomodellinghub.org>, last accessed Apr 6th, 2023.
- [3] "Covid-19 modeling," Virginia Department of Health. [Online]. Available: <https://www.vdh.virginia.gov/coronavirus/see-the-numbers/covid-19-modeling/>
- [4] "Ut austin covid-19 modeling consortium," University of Texas at Austin COVID-19 Modeling Consortium. [Online]. Available: <https://covid-19.tacc.utexas.edu/>
- [5] H. W. Hethcote, "The mathematics of infectious diseases," *SIAM review*, vol. 42, no. 4, pp. 599–653, 2000.
- [6] P. Bhattacharya, D. Machi, J. Chen, S. Hoops, B. Lewis, H. Mortveit, S. Venkatramanan, M. L. Wilson, A. Marathe, P. Porebski *et al.*, "AI-driven agent-based models to study the role of vaccine acceptance in controlling covid-19 spread in the us," in *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 2021, pp. 1566–1574.
- [7] P. Bhattacharya, J. Chen, S. Hoops, D. Machi, B. Lewis, S. Venkatramanan, M. L. Wilson, B. Klahn, A. Adiga, B. Hurt *et al.*, "Data-driven scalable pipeline using national agent-based models for real-time pandemic response and decision support," *The International Journal of High Performance Computing Applications*, p. 10943420221127034, 2022.




Fig. 8. Cumulative infections over time for 30 replicates in EpiHiper (left) and Loimos (right) of a simulated MD outbreak, with both distributions showing similar average infection totals but with a wider range of times until equilibrium is reached for EpiHiper.




Fig. 9. Weak scaling results on Perlmutter for three different per-processor loads, with runtimes relatively flat but with more variation for larger datasets. Execution times averaged over three runs, with extrema shown in error bars.

[8] J. J. Grefenstette, S. T. Brown, R. Rosenfeld, J. DePasse, N. T. Stone, P. C. Cooley, W. D. Wheaton, A. Fyshe, D. D. Galloway, A. Sriram *et al.*, “Fred (a framework for reconstructing epidemic dynamics): an open-source software system for modeling infectious diseases and control strategies using census-based populations,” *BMC public health*, vol. 13, no. 1, pp. 1–14, 2013.

[9] T. C. Germann, K. Kadau, I. M. Longini, and C. A. Macken, “Mitigation strategies for pandemic influenza in the United States,” *Proceedings of the National Academy of Sciences*, vol. 103, no. 15, pp. 5935–5940, Apr. 2006, publisher: Proceedings of the National Academy of Sciences. [Online]. Available: <https://www.pnas.org/doi/full/10.1073/pnas.0601266103>

[10] P. Bhattacharya, J. Chen, S. Hoops, D. Machi, B. Lewis, S. Venkatramanan, M. L. Wilson, B. Klahn, A. Adiga, B. Hurt, J. Outten, A. Adiga, A. Warren, Y. Y. Baek, P. Porebski, A. Marathe, D. Xie, S. Swarup, A. Vullikanti, H. Mortveit, S. Eubank, C. L. Barrett, and M. Marathe, “Data-driven scalable pipeline using national agent-based models for real-time pandemic response and decision support,” *The International Journal of High Performance Computing Applications*, vol. 37, no. 1, pp. 4–27, Jan. 2023, publisher: SAGE Publications Ltd STM. [Online]. Available: <https://doi.org/10.1177/10943420221127034>

[11] A. Bhatele, J.-S. Yeom, N. Jain, C. J. Kuhlman, Y. Livnat, K. R. Bisset, L. V. Kale, and M. V. Marathe, “Massively parallel simulations of spread of infectious diseases over realistic social networks,” in *2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. IEEE, 2017, pp. 689–694.

[12] D. Machi, P. Bhattacharya, S. Hoops, J. Chen, H. Mortveit, S. Venkatramanan, B. Lewis, M. Wilson, A. Fadikar, T. Maiden, C. L. Barrett,

and M. V. Marathe, “Scalable Epidemiological Workflows to Support COVID-19 Planning and Response,” in *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2021, pp. 639–650, ISSN: 1530-2075.

[13] S. Truelove, C. P. Smith, M. Qin, L. C. Mullany, R. K. Borchering, J. Lessler, K. Shea, E. Howerton, L. Contamin, J. Levander, J. Kerr, H. Hochheiser, M. Kinsey, K. Tallaksen, S. Wilson, L. Shin, K. Rainwater-Lovett, J. C. Lemaitre, J. Dent, J. Kaminsky, E. C. Lee, J. Perez-Saez, A. Hill, D. Karlen, M. Chinazzi, J. T. Davis, K. Mu, X. Xiong, A. Pastore y Piontti, A. Vespignani, A. Srivastava, P. Porebski, S. Venkatramanan, A. Adiga, B. Lewis, B. Klahn, J. Outten, M. Orr, G. Harrison, B. Hurt, J. Chen, A. Vullikanti, M. Marathe, S. Hoops, P. Bhattacharya, D. Machi, S. Chen, R. Paul, D. Janies, J.-C. Thill, M. Galanti, T. K. Yamana, S. Pei, J. L. Shaman, J. M. Healy, R. B. Slayton, M. Biggerstaff, M. A. Johansson, M. C. Runge, and C. Viboud, “Projected resurgence of COVID-19 in the United States in July–December 2021 resulting from the increased transmissibility of the Delta variant and faltering vaccination,” *eLife*, vol. 11, p. e73584, Jun. 2022, publisher: eLife Sciences Publications, Ltd. [Online]. Available: <https://doi.org/10.7554/eLife.73584>

[14] K. R. Bisset, J. Cadena, M. Khan, and C. J. Kuhlman, “Agent-Based Computational Epidemiological Modeling,” *Journal of the Indian Institute of Science*, vol. 101, no. 3, pp. 303–327, Jul. 2021. [Online]. Available: <https://doi.org/10.1007/s41745-021-00260-2>

[15] A. Tiwari, “Modelling and analysis of covid-19 epidemic in india,” *Journal of Safety Science and Resilience*, vol. 1, no. 2, pp. 135–140, 2020.

[16] G. Giordano, F. Blanchini, R. Bruno, P. Colaneri, A. Di Filippo, A. Di Matteo, and M. Colaneri, “Modelling the covid-19 epidemic and implementation of population-wide interventions in italy,” *Nature medicine*, vol. 26, no. 6, pp. 855–860, 2020.

[17] K. Prem, Y. Liu, T. W. Russell, A. J. Kucharski, R. M. Eggo, N. Davies, S. Flasche, S. Clifford, C. A. Pearson, J. D. Munday *et al.*, “The effect of control strategies to reduce social mixing on outcomes of the covid-19 epidemic in wuhan, china: a modelling study,” *The Lancet Public Health*, vol. 5, no. 5, pp. e261–e270, 2020.

[18] C. Anastassopoulou, L. Russo, A. Tsakris, and C. Siettos, “Data-based analysis, modelling and forecasting of the covid-19 outbreak,” *PloS one*, vol. 15, no. 3, p. e0230405, 2020.

[19] M. Chinazzi, J. T. Davis, M. Ajelli, C. Gioannini, M. Litvinova, S. Merler, A. P. y Piontti, K. Mu, L. Rossi, K. Sun *et al.*, “The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak,” *Science*, vol. 368, no. 6489, pp. 395–400, 2020.

[20] P. C. Silva, P. V. Batista, H. S. Lima, M. A. Alves, F. G. Guimarães, and R. C. Silva, “Covid-abs: An agent-based model of covid-19 epidemic to simulate health and economic effects of social distancing interventions,” *Chaos, Solitons & Fractals*, vol. 139, p. 110088, 2020.

[21] E. Cuevas, “An agent-based model to evaluate the covid-19 transmission risks in facilities,” *Computers in biology and medicine*, vol. 121, p. 103827, 2020.

[22] B. G. Aaby, K. S. Perumalla, and S. K. Seal, “Efficient simulation

- of agent-based models on multi-gpu and multi-core clusters,” in *Proceedings of the 3rd International ICST Conference on Simulation Tools and Techniques*, ser. SIMUTools '10. Brussels, BEL: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2010. [Online]. Available: <https://doi.org/10.4108/ICST.SIMUTOOLS2010.8822>
- [23] J. Parker and J. M. Epstein, “A distributed platform for global-scale agent-based models of disease transmission,” *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 22, no. 1, pp. 1–25, 2011.
- [24] K. S. Perumalla and S. K. Seal, “Discrete event modeling and massively parallel execution of epidemic outbreak phenomena,” *Simulation*, vol. 88, no. 7, pp. 768–783, 2012.
- [25] C. L. Barrett, K. R. Bisset, S. G. Eubank, X. Feng, and M. V. Marathe, “Episimemics: an efficient algorithm for simulating the spread of infectious disease over large realistic social networks,” in *SC'08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*. IEEE, 2008, pp. 1–12.
- [26] J.-S. Yeom, A. Bhatele, K. Bisset, E. Bohm, A. Gupta, L. V. Kale, M. Marathe, D. S. Nikolopoulos, M. Schulz, and L. Wesolowski, “Overcoming the scalability challenges of epidemic simulations on blue waters,” in *2014 IEEE 28th International Parallel and Distributed Processing Symposium*. IEEE, 2014, pp. 755–764.
- [27] W.-m. Liu, H. W. Hethcote, and S. A. Levin, “Dynamical behavior of epidemiological models with nonlinear incidence rates,” *Journal of mathematical biology*, vol. 25, pp. 359–380, 1987.
- [28] J. Chen, S. Hoops, H. S. Mortveit, B. L. Lewis, D. Machi, P. Bhattacharya, S. Venkatramanan, M. L. Wilson, C. L. Barrett, and M. V. Marathe, “Epihiper—a high performance computational modeling framework to support epidemic science,” *PNAS nexus*, vol. 4, no. 1, p. pgae557, 2025.
- [29] J. Mossong, N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, M. Massari, S. Salmaso, G. S. Tomba, J. Wallinga *et al.*, “Social contacts and mixing patterns relevant to the spread of infectious diseases,” *PLoS medicine*, vol. 5, no. 3, p. e74, 2008.
- [30] L. V. Kale and A. Bhatele, Eds., *Parallel Science and Engineering Applications: The Charm++ Approach*. Taylor & Francis Group, CRC Press, Nov. 2013.
- [31] L. V. Kale, A. Arya, A. Bhatele, A. Gupta, N. Jain, P. Jetley, J. Lifflander, P. Miller, Y. Sun, R. Venkataraman, L. Wesolowski, and G. Zheng, “Charm++ for productivity and performance: A submission to the 2011 HPC Class II Challenge,” Dept. of Computer Science, University of Illinois, Tech. Rep., Nov. 2011.
- [32] U. S. C. Bureau, “North American Industry Classification System (NAICS) U.S. Census Bureau.” [Online]. Available: <https://www.census.gov/naics/>
- [33] —, “Census Bureau Data.” [Online]. Available: <https://data.census.gov/>
- [34] U. S. D. of Transportation Federal Highway Administration, “NHTS NextGen OD Data.” [Online]. Available: <https://nhts.ornl.gov/od/>
- [35] Microsoft, “microsoft/USBuildingFootprints,” Jan. 2024, original-date: 2018-06-13T18:31:31Z. [Online]. Available: <https://github.com/microsoft/USBuildingFootprints>
- [36] N. C. for Educational Statistics, “Electronic Catalog of NCES Products (National Center for Education Statistics). Publications and data products.” publisher: National Center for Education Statistics. [Online]. Available: <https://nces.ed.gov/datatools/index.asp?DataToolSectionID=1>
- [37] NERSC, “Perlmutter system architecture,” <https://docs.nersc.gov/systems/perlmutter/architecture/>.
- [38] J.-s. Yeom, A. Bhatele, K. R. Bisset, E. Bohm, A. Gupta, L. V. Kale, M. Marathe, D. S. Nikolopoulos, M. Schulz, and L. Wesolowski, “Overcoming the scalability challenges of epidemic simulations on Blue Waters,” in *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, ser. IPDPS '14. IEEE Computer Society, May 2014, ILNL-CONF-648533.
- [39] A. Collins, M. Koehler, and C. Lynch, “Methods that support the validation of agent-based models: An overview and discussion,” *Journal of Artificial Societies and Social Simulation*, vol. 27, no. 1, 2024.