Shift, scale and restart smaller models to estimate larger ones: Agent based simulators in epidemiology

Daksh Mittal, Sandeep Juneja, Shubhada Agrawal (TIFR, Mumbai)

1. INTRODUCTION

Agent-based simulators (ABS) are a popular epidemiological modelling tool to study the impact of various non pharmaceutical interventions in managing epidemics [1], [2]. They provide the flexibility to accurately model a heterogeneous population with time and location varying, person specific interactions. Government policies such as partial and location specific lock downs, case isolation, home quarantine, school closures, partially opened workplaces, etc. are easily modelled and ABS allow flexibility to incorporate important pandemic developments over time including presence of variants as well as vaccines. For accuracy, each person is modelled separately. This however may make computational time prohibitive when the city population and the simulated time is large. We observe that simply considering a smaller aggregate model and scaling up the output leads to inaccuracies. In this note we primarily focus on the COVID-19 pandemic and dig deeper into the underlying probabilistic structure of generic ABS to arrive at modifications that allow smaller models to give accurate statistics for larger models. We exploit the observations that in the initial disease spread phase, the starting infections create a family tree of infected individuals more-or-less independent of the other trees and are modelled well as a multi-type super-critical branching process. Soon after, for large city population, once enough people have been infected, the future evolution of the pandemic is closely approximated by its mean field limit with a random starting state. We build upon these insights to develop a shifted, scaled and restart based algorithm that accurately evaluates the ABS's performance using a much smaller model while carefully reducing the bias that may otherwise arise. Our key contributions are: First, we develop an algorithm by carefully exploiting the closeness of the underlying exposed/infected process (process of number exposed/infected of each type at each time) initially to a branching process, and then the normalised infection process (infection process divided by the city population) to its mean field limit, so that the output from the smaller model accurately matches the output from the larger model. Second, we provide theoretical support for the proposed approach through an asymptotic analysis where the population size increases to infinity. In the interest of space, we show this in a simple and yet practically useful setting of compartmental models.

Agent based Simulator: We first informally spell out the dynamics of our infection spread model. A more detailed discussion can be seen in [3] and [4]. The model consists of individuals, households, schools, workplaces and community spaces. The number of individuals living in a household, their age, whether they go to school or work or neither, schools and workplaces size and composition all have distributions that may be set to match the available data. The model proceeds in discrete time steps of constant width Δt (six hours in our set-up). At a well chosen time zero, a small number of individuals can be set to either exposed, asymptomatic, or symptomatic states, to seed the infection. At each time t, an infection rate $\lambda_n(t)$ is computed for each susceptible individual based on its interactions with other infected individuals in different interaction spaces. In the next Δt time, each susceptible individual moves to the exposed state with probability $1 - \exp\{-\lambda_n(t) \cdot \Delta t\}$, independently of all other events. Further, disease may progress independently in the interval Δt for the population already afflicted by the virus. The exact computation of $\lambda_n(t)$ and the detailed probabilistic dynamics of disease progression can be seen in [3]. Simulation time is then incremented to $t + \Delta t$, and the state of each individual is updated to reflect the new exposures, changes to infectiousness, hospitalisations, recoveries, quarantines, etc., during the period t to $t + \Delta t$. The overall process repeats incrementally until the end of the simulation time.

2. SPEEDING UP ABS

A naive approach to speed up the ABS maybe to use a representative smaller population model and scale up the results. Thus, for instance, while a realistic model for Mumbai city may have 12.8 million agents (see [4]), we may construct a sparser Mumbai city having, say, a million agents, that matches the bigger model in essential features, so that, roughly speaking, in the two models each infectious person contributes the same total infection rate to susceptibles at each time. The output numbers from the smaller model may be scaled by a factor of 12.8 to estimate of output from the larger model. We observe, somewhat remarkably, that this naive approach is actually accurate if no interventions to disease spread are imposed and if the initial seed infections in the model are large, say, of the order of tens of thousands (see Figure 1; we plot only the number exposed. The comparative statements hold equally well for other statistics such as the number infected, hospitalised, in ICUs and deceased). The rationale is that in this setting both the smaller and the larger model well approximate their mean field limits, that are identical, and are assigned identically

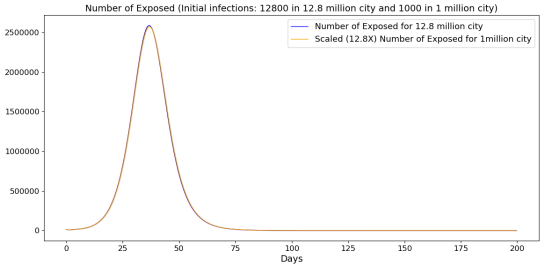

Figure 1: Scaled no. exposed in the smaller model match the larger model when we start with large, 12800, no. of infections.

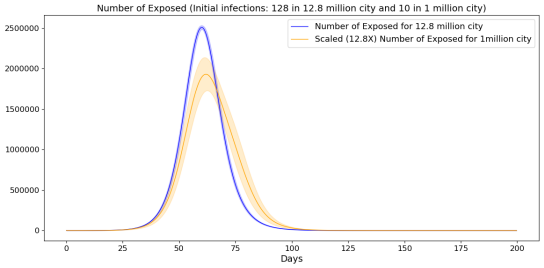

Figure 2: Scaled no. exposed in smaller model do not match the larger model when we start with few, 128, no. of infections.

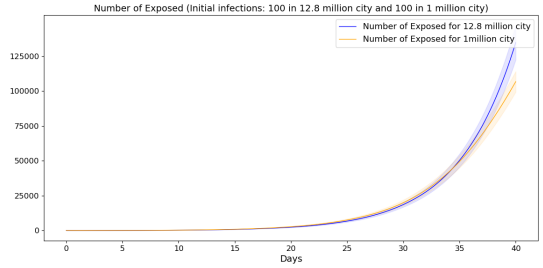

Figure 3: Smaller and larger model are essentially identical initially when we start with same number of few, 100, no. of infections

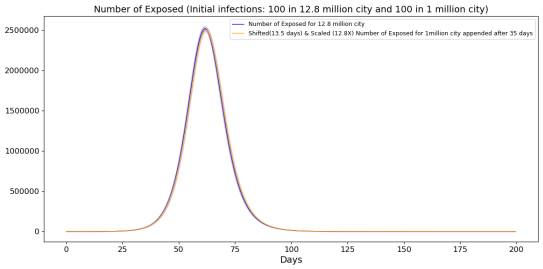

Figure 4: Shift and scale smaller model matches the larger model under no intervention scenario

distributed initial condition.

However, modelling initial randomness in the disease spread is important for reasons including ascertaining the distribution of when and where an outbreak may be initiated, the probability that some of the initial infection clusters die-down, etc. These are typically captured by setting the initial infections to a small number, say, around a hundred, and the model is initiated at a well chosen time [3]. In such settings, we observe that the scaled output from the smaller model (with proportionately lesser initial infections) is noisy and biased so that the simple scaling fix no longer works (Figure 2). In fact, we observe that in the early days of the infection, the smaller and the larger model with the same number of initial infections behave more or less identically (Figure 3), so that the smaller model with the unscaled number of initial seed infections provides an accurate approximation to the larger one (here again in the early phase in the two models each infectious person contributes the roughly the same total infection rate to susceptibles at home, workplace and community at each time). Probabilistically this is true because early on both the models closely approximate an associated multi-type branching process.

Shift-scale-restart algorithm : For $k, N \in \mathbb{N}$, $k > 1$, let kN be the number of individuals in the larger city, and N in the smaller city. The proposed algorithm simulates the smaller city to get the statistics for the larger city, and hence is significantly faster (close to factor k). Let both the larger, as well as the smaller cities start with small I_0 infections at time 0. Let t_S denote the time till the two cities evolve essentially identically (as seen empirically and suggested by theoretical analysis this is close to $\alpha \log N / \log \rho$ time for α close to and < 1, where ρ denotes the initial infection exponential growth rate), t_I denote the first intervention time (e.g., lockdown; typically after $\beta \log N$ time for small $\beta \in (0, 1)$). Let $t_{min} \approx \min\{t_I, t_S\}$. The algorithm is as follows: Start the smaller city with I_0 infections and evolve the city upto t_{min} (this part is used to capture the initial evolution of the larger city). Suppose there are x infections at t_{min} in the smaller city. Determine an earlier time $t_{\frac{x}{k}}$ in the simulation when there where $\frac{x}{k}$ infections in the small city. As we see empirically, and as is suggested by Proposition 1, the distribution amongst the non-susceptibles at time $t_{\frac{x}{k}}$ is similar to that at time t_{min}. Further, by this time the infection process evolves more or less deterministically. We would like to append the path from time $t_{\frac{x}{k}}$ onwards, scaled by factor k, to the path at t_{min} in the smaller model. However, to ensure that the intervention times are correctly captured, we instead restart a new simulation of the smaller city with I_0 infections, but with the first intervention at time $t_{\frac{x}{k}} + t_I - t_{min}$ in this run of the simulation. To get the large city path from these two simulations of the smaller city, we scale the path of the second simulation from time $t_{\frac{x}{k}}$ onwards by a factor of k, and append this scaled path to the first simulation of the smaller city after t_{min}.

Numerical Results: Figure 4 compares number of exposed people in a 12.8 million Mumbai city simulation (in no intervention scenario) with estimates from the shift-scale-restart algorithm applied to the smaller 1 million city. Figures 1 to 4 are under no intervention scenario for Mumbai created as in [3]. Figure 5 compares the exposed population process for Mumbai with the smaller 1 million city as per our algorithm under realistic interventions (lockdown, case isolation, home quarantine, masking etc.) introduced at realistic times, as implemented in [5] using similar parameters.

Theoretical Results: In the above algorithm, the fact that early on in the small city simulation, we could take a path at one time period, scale it and stitch to the path at another appropriately chosen time period to accurately generate a path for the larger city, requires theoretical justification. We provide this through analyzing our city in an

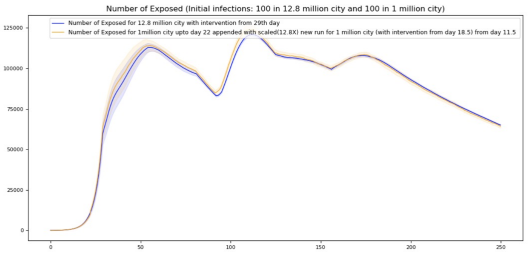

Figure 5: Shift-scale-restart smaller model match the larger one under real world interventions over 250 days.

asymptotic regime as the city population N increases to infinity. To bring out the key observations simply, we consider a simpler compartmental model. Specifically, the city comprises of N individuals with fraction π_a in age group a with a single 'community' interaction space. An individual disease state can be susceptible, exposed, infected, symptomatic, hospitalised, critical, dead or recovered. The total number of contacts any infectious individual makes with all the individuals in the city in a time step is Poisson distributed with rate $\beta\Delta t$. Each contact is with an individual chosen uniformly at random from N individuals and if the contacted individual is susceptible, it becomes exposed, otherwise the contact has no effect. Once an individual gets exposed, its disease progression is independent of all the other individuals in the city and depends only on its age. The time spent in each state (except susceptible, dead, recovered) is geometrically distributed. Transition to symptomatic (hospitalised, critical, dead) from respective earlier disease states happens with respective age dependent transition probabilities, otherwise the person recovers. Let $X_t^N \in \mathbb{Z}^{+d}$, denote the number of individuals of each type (differing in age and disease state) excluding the susceptible population at time t. Further, we define a multi type super-critical branching process [6] closely coupled to the above epidemic process. As in the epidemic process, each infectious individual gives birth to Poisson distributed exposed individuals of each age group a with rate $\pi_a\beta\Delta t$. Broadly, the two processes only differ in that in the epidemic process contact with an already exposed person has no impact while in branching process this still gives birth to a new exposed individual uncoupled from the epidemic process. Further, the branching process is independent of N. Once an exposed individual is born, disease progression of the individual has same probabilistic evolution as in epidemic process when they are coupled. $\mathbf{B_t} \in \mathbb{Z}^{+d}$ denote the number of individuals of different types at time t in the branching process.

Following result shows that epidemic process is close to the multi-type branching process till time $\alpha \log_\rho N$ for any $\alpha \in (0,1)$ as $N \to \infty$, where ρ denotes the exponential growth rate of the branching process.

THEOREM 1. *As $N \to \infty$,*

$$\textit{For } \alpha \in (0, 0.5), \quad \sup_{t \in [0, \alpha \log_\rho N]} \left| \mathbf{X_t^N} - \mathbf{B_t} \right| \xrightarrow{P} 0. \qquad (1)$$

$$\textit{For } \alpha \in (0, 1), \quad \sup_{t \in \left[0, \alpha \log_\rho N\right]} \left| \frac{\mathbf{X_t^N}}{\rho^t} - \frac{\mathbf{B_t}}{\rho^t} \right| \xrightarrow{P} 0. \qquad (2)$$

Result (1) above was earlier proved for SIR setting in [7]. We extend this result to general compartmental models and also prove the result (2) in this more general setting. The following proposition justifies the fact that the proportions across different types stabilizes quickly in the epidemic process, and thus paths can be patched from one time period to the other without much error due to change in the proportions.

PROPOSITION 1. *For $t_N \to \infty$ as $N \to \infty$ and $\frac{t_N}{\alpha \log_\rho N} \to 0$ for some $\alpha \in (0,1)$, there exists an $x \in \Re^{+d}$, $\mathbb{1}^T x = 1$:*

$$\left| \frac{\mathbf{X_{t_N}^N}}{\mathbb{1}^T \mathbf{X_{t_N}^N}} - \mathbf{x} \right| \xrightarrow{P} 0, \qquad as \quad N \to \infty.$$

As Proposition 1 notes, early in the infection growth, the infection process proportions stabilize although the number infected constitute a negligible fraction compared to the susceptible population. This changes at time $\log_\rho(\epsilon N)$, for any $\epsilon > 0$ where the system has $\Theta(N)$ infections. Setting $t_N = log_\rho(\epsilon N)$, the process $X^N_{t_N+t}$ augmented with the susceptible population at that time and normalized through division by N, results in an empirical distribution denoted by μ_t^N. Then, μ_t^N can be seen to converge to a mean field process, call it $\bar{\mu}_t$, as $N \to \infty$. Specifically, suppose there exists a deterministic $\bar{\mu}_0$ such that $\mu_0^N \xrightarrow{P} \bar{\mu}_0$ as $N \to \infty$. Then, for t a multiple of Δt, $t \geq \Delta t$, until the simulation time or till disease dies down (whichever is earlier), $\mu_t^N \xrightarrow{P} \bar{\mu}_t$ as $N \to \infty$, where, $\bar{\mu}_t(s)$ for any type s satisfies

$$\bar{\mu}_t(s) := \sum_{\text{all types } s'} \bar{\mu}_{t-\Delta t}(s') h(s', s, \bar{\mu}_{t-\Delta t}),$$

and $h(s', s, \bar{\mu}_{t-\Delta t})$ is the proportion of population of type s' that migrates to type s in Δt time. In particular, if $\bar{\mu}_t$ denotes the mean field limit of the normalised process at time $t + \log_\rho(\epsilon N)$, then, the number of infections observed in a smaller model with population N is approximately $N * \bar{\mu}_t$ and that of a larger model is approximately $kN * \bar{\mu}_t$. Thus, the larger model infection process can be approximated by the smaller model infection process by scaling it by k.

3. REFERENCES

[1] E. Hunter, B. Mac Namee, and J. D. Kelleher, "A taxonomy for agent-based models in human infectious disease epidemiology," *JASSS*, 2017.

[2] Ferguson et. al., "Report 9: Impact of non pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," 2020.

[3] Agrawal et. al., "City-scale agent-based simulators for the study of non-pharmaceutical interventions in the context of the covid-19 epidemic," *JIISC*, 2020.

[4] S. Juneja and D. Mittal, "Modelling the Second Covid-19 Wave in Mumbai," *arxiv*, May 2021.

[5] IISc-TIFR Covid-19 City-Scale Simulation Team, "COVID-19 Epidemic Study II: Phased Emergence From the Lockdown in Mumbai," *arxiv*, June 2020.

[6] K. Athreya and P. Ney, *Branching processes*, 1972.

[7] T. Britton and E. Pardoux, *Stochastic Epidemic Models with Inference*, 2019.