5. Zhang R., Son O.-S. A Measurement Model for Visual Complexity in HCI: Focusing on Visual Elements in Mobile GUI Design. Electronics. 2025. Т. 14, вип. 5. С. 942. URL: https://doi.org/10.3390/electronics14050942.

ПРИСКОРЕНЕ МОДЕЛЮВАННЯ ЕПІДЕМІЙ У ВЕЛИКИХ МЕРЕЖАХ: ОПТИМІЗАЦІЯ МЕТОДУ ГІЛЛЕСПІ ТА ДВОШАРОВИЙ ПІДХІД / ACCELERATED EPIDEMIC SIMULATION IN LARGE-SCALE NETWORKS: OPTIMIZATION OF THE GILLESPIE ALGORITHM AND A TWO-LAYER APPROACH

Куриляк Ю.А.[1], Еммеріх М.Т.М, [2]/ Kuryliak Y.A. [1], Emmerich M.T.M.[2]
[1] Національний університет «Львівська політехніка» / Lviv Polytechnic National University
вулиця Степана Бандери, 12, Львів, Львівська область, 79000, Україна
[2] Університет Юваскуля / University of Jyväskylä
Agora, Mattilanniemi 2, 40100 Jyväskylä, Jyväskylä, 40014, Finland
E-mail: [1]yulian.a.kuryliak@lpnu.ua, [2]michael.t.m.emmerich@jyu.fi

Abstract: This study addresses the challenge of accelerating epidemic simulations in large-scale complex networks through algorithmic and structural optimization. Traditional Gillespie-based stochastic simulations accurately reproduce epidemic dynamics but become computationally prohibitive for networks exceeding tens of thousands of nodes. To overcome this limitation, we build upon such efficiency techniques as local rate updates and ordered event-selection structures, which reduce the computational complexity of each simulation step from O(n) to O(log(n)). Building on these principles, we propose a two-layer (micro–macro) modeling framework: the micro layer simulates intra-community dynamics, while the macro layer captures inter-community infections using hazard-integral rates derived from mobility data and the epidemic states of metanodes. This hierarchical approach enables scalable and parallelizable simulations that preserve stochastic accuracy while substantially reducing computational cost, allowing realistic modeling of epidemic spread across millions of agents and multiple cities.

Epidemics have accompanied humanity throughout history. Each time a new wave of disease emerges, societies seek to understand how it spreads and what can be done to contain it. Historically, epidemic dynamics were modeled using differential equations based on the mean-field assumption — treating the population as a homogeneous system without accounting for its contact structure. Today, with increasing computational power and advances in simulation technology, it has become possible to reproduce infection dynamics in realistic, irregular contact networks that reflect individual-level interactions.

Contact networks are sparse, heterogeneous in degree, and exhibit strong local clustering and a pronounced community structure — groups of nodes densely connected within communities but weakly linked to others. Therefore, accurate modeling requires capturing explicit node-to-node connections and their temporal state transitions. Agent-based modeling achieves this by representing each node as an individual whose contacts define potential infection paths.

Such processes can be formulated as Continuous-Time Markov Chains (CTMC). To simulate them efficiently, the Gillespie algorithm provides an exact stochastic representation of system dynamics. Its key principles are: (I) only feasible transitions from the current state are considered; (II) only one event occurs at a time; and (III) the time to the next event follows an exponential distribution, after which the event is selected proportionally to its rate. This ensures accuracy and avoids state-space explosion, but computational complexity still limits the model size. As shown in [1], the classical Gillespie algorithm can simulate up to several tens of thousands of nodes, which remains insufficient for realistic population scales. For comparison, modern agent-based frameworks such as Covasim [2] implement stochastic transitions in discrete time with a daily step, which simplifies computations but sacrifices temporal precision.

The first optimization involves local rate updates. In the standard algorithm, every state change triggers recalculation of all possible transitions. In the improved implementation, only the node that changes state and its immediate neighbors are updated, dramatically reducing the number of operations while preserving accuracy.

The second optimization employs ordered data structures for rapid event selection. Whereas the classical approach scans all transition rates, the proposed sorted-list method achieves logarithmic-time selection [1]. An alternative solution is the binary-tree approach described in [3].

Despite these optimizations, simulating very large networks remains computationally demanding. To further enhance scalability, we introduce a two-layer model that hierarchically partitions the network into communities. At the micro layer, each community is simulated using the optimized Gillespie algorithm; at the macro layer, rarer interactions between communities are modeled using CTMC dynamics. This hierarchical approach, detailed in [4], drastically reduces the number of processed events while maintaining node-level granularity.

In the full-scale simulation, the underlying contact network is assumed to be static, and communities are connected through a limited number of bridge nodes that mediate inter-community infections. Once such a bridge node becomes infected, additional external infections cannot enter the same community until recovery occurs, effectively constraining cross-community transmission. In contrast, the two-layer abstraction treats each community as a metanode at the macro layer, where new infections are assigned to a random node within the community rather than fixed bridges. This eliminates topological bias, better reflects short-term human mobility, and makes the model behavior effectively more dynamic while preserving consistency with the underlying network structure.

Another advantage of the two-layer model is its parallel execution capability. Since most events occur within communities, each community can be simulated on a separate processor core, yielding near-linear speedup as the number of cores increases — particularly when inter-community connectivity is low. In this regime, synchronization between micro-layer simulations is infrequent because cross-community infections are rare, allowing communities to evolve almost independently. As inter-community connections become denser, synchronization must occur more often, which slightly reduces efficiency; however, the overall runtime still remains significantly lower than in a monolithic model. Moreover, epidemic activity peaks in different communities typically occur at different times, enabling dynamic reallocation of computational resources to the most active regions. Under limited hardware resources, a subset of communities outside the primary focus of analysis can be replaced with precomputed infection curves, preserving the accuracy of global dynamics while further reducing computational cost.

The two-layer simulator can also be interpreted as a simplified social-mobility model, where macro-layer nodes represent cities and edge weights correspond to intercity travel flows. With appropriate mobility data, the model enables analysis of how changes in movement patterns or local restrictions affect the rate and scale of infection spread. Australia offers a convenient case study due to its geographic isolation; its intercity mobility network can be reconstructed from mobile-operator-derived datasets [5].

In conclusion, the optimized Gillespie algorithm combined with the two-layer modeling framework achieves a practical balance between stochastic accuracy and computational scalability. It enables realistic simulations of epidemic dynamics in large, structured populations and can be readily integrated with demographic and mobility datasets. This framework also supports direct comparison between simulated and real-world epidemic data, paving the way for quantitative evaluation of intervention and containment policies.

СПИСОК ВИКОРИСТАНИХ ДЖЕРЕЛ

1. Kuryliak, Y., Emmerich, M. T. M., Dosyn, D. Simulating epidemic peak dynamics on complex networks using efficient Gillespie algorithms. *Infection, Genetics and Evolution*, 2025, 105768. DOI: 10.1016/j.meegid.2025.105768

2. Kerr, C. C., Stuart, R. M., Mistry, D., et al. Covasim: An agent-based model of COVID-19 dynamics and interventions. *PLOS Computational Biology*, 2021, vol. 17, no. 7, e1009149. DOI: 10.1371/journal.pcbi.1009149.
3. Vo Hong Thanh, R. & Zunino, R. Tree-based search for Stochastic Simulation Algorithm. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing (SAC '12)*. Riva del Garda, Italy, 2012, pp. 144–151. DOI: 10.1145/2245276.2232001.
4. Kuryliak, Y. & Emmerich, M. T. M. Towards complexity reduction of large-scale epidemic simulation in two-scale networks. 2025, *CEUR Workshop Proceedings 4004: pp. 327-336.*
5. Du, Z., Yang, Y., Ertem, Z., et al. A review of human mobility: Linking data, models, and real-world applications. *Humanities and Social Sciences Communications*, 2025, vol. 12, no. 1, pp. 1–18. DOI: 10.1057/s41599-025-03160-7

ЗАСТОСУВАННЯ АЛГОРИТМІВ НАВЧАННЯ Q-МЕРЕЖ ДЛЯ ІТЕРАТИВНОЇ ЗАДАЧІ В'ЯЗНЯ

Ткач.Н.В.
Національний університет "Києво-Могилянська академія"
04655, м. Київ, вулиця Сковороди, 2, НаУКМА, Факультет інформатики. тел.(044) 426 60 64.
E-mail: n.tkach@ukma.edu.ua, факс (044) 426 60 64

This study is aimed at showcasing the performance of Deep Q-Networks (DQN) for the Iterated Prisoner's Dilemma (IPD) with a compact episodic state embedding. The agent compresses the interaction context into a fixed-size vector and is trained against deterministic Axelrod strategies. Evaluation of normalized payoff, pairwise cooperation rate of strategies, and the learned behavior of the agent suggests the possibility of efficiently clustering existing strategies by latent learnable features. This may lead to advancements in both game theory and reinforcement learning. The limitations are outlined for future research, including recurrent-based and transformer-based policy-learning networks, stochastic opponents, and comparative analysis to the baseline performance.

Ітераційна дилема в'язня (IPD) є класичною моделлю співпраці[1]. У даній роботі представлено результати агента, навченого на основі алгоритму навчання з підкріпленням глибинних Q-мереж ("deep Q-network")[3], здатного формувати контекстно-залежні рішення за рахунок стислого подання стану. Модель тренувалася проти 72 детерміністичних стратегій у 200-ходових епізодах, реалізованих згідно з турнірами Аксельрода [1] в однойменній бібліотеці.

DQN-агент приймає стислий стан,- 16-вимірний вектор поточної історії гри, що кодує ключові патерни взаємодії. Навчання здійснюється проти детерміністичної популяції у "сліпому" режимі - модель немає інформації про стратегію, проти якої навчається. Функція втрат L(θ) базується на TD-помилці функції оцінювання $V_{\psi}(s)$. Q-функцію[2] ми оцінюємо за допомогою багатошарового перцептрону, який приймає на вхід інформацію про стан s що містить контекст стратегії та останню дію a. Параметрами φ та θ інтерпретуємо як відповідні параметри нейронних мереж які надають оцінку поточного стану, та оцінку стратегії, яку агент має вибрати для мінімізації похибки часової різниці(TD).

$$\min_{\psi}\left(r+\gamma V_{\psi}\left(s'\right)-V_{\psi}(s)\right)^2 \tag{1}$$

$$y^{\mathrm{DQN}}=r+\gamma \max_{a'} Q_{\theta^-}\left(s',a'\right) \tag{2}$$

$$L(\theta)=E_{(s,a,r,s')D}\left[\left(y^{DQN}-Q_{\theta}(s,a)\right)^2\right] \tag{3}$$