Time, Clocks and Efficiency of Population Protocols

Leszek Gąsieniec[1] ✉ 🏠 ⓘ
University of Liverpool, UK

Grzegorz Stachowiak ✉ ⓘ
University of Wrocław, Poland

─── **Abstract** ───

The model of population protocols is used to study distributed processes based on pairwise interactions between simple anonymous agents drawn from a large population of size n. The order in which agents meet in pairs is determined by the random scheduler, s.t., each consecutive pair is chosen uniformly at random. After each interaction the state of the relevant agents are amended according to the predefined transition function (the actual code of the algorithm) which governs the considered process. The state space of agents is often fixed and the size n is not known in advance, i.e., not hard-coded in the transition function. We assume that a population protocol starts in the predefined initial configuration of agents' states representing the input. And if successful, the protocol stabilises in a final configuration of states forming the output representing the solution to the considered problem.

The time complexity of a population protocol refers to the number of interactions required to stabilise this protocol in a final configuration. We also define *parallel time* as the time complexity divided by n. Note that the parallel time of the system and the expected local time of each agent, i.e., the number of interactions observed by each agent, are correlated. Several mechanisms, known as *phase clocks*, have been developed to measure parallel time more accurately than counting local interactions. Most of the clocks target counting $\Theta(\log n)$ parallel time required to fully synchronise all agents in the population. There are leader (and junta) based phase clocks which utilise a fixed number of states [2, 4]. This type of clocks allows also counting any poly-logarithmic time while preserving fix state utilisation. The other type refers to leaderless clocks utilising $\Theta(\log n)$ states [1, 5]. This type allows approximate counting of parallel time as fixed resolution clocks [5] or oscillators [1]. Another clock type introduced recently in [3] enables counting $\Theta(n \log n)$ parallel time utilising a fixed number of states and either leaders or connections in the network constructor model.

We also discuss parallel efficiency of population protocols referring to protocols operating in $\Theta(\text{poly} \log n)$ parallel time, we propose extensions of the population protocol model leading to further improvement in state and time utilisation, and we state some open problems.

─── **References** ───

1 D. Alistarh, J. Aspnes, and R. Gelashvili. Space-optimal majority in population protocols. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2221–2239. SIAM, 2018.

2 D. Angluin, J. Aspnes, and D. Eisenstat. Fast computation by population protocols with a leader. *Distributed Comput.*, 21(3):183–199, 2008.

---

[1] Corresponding author

3   L. Gąsieniec, P.G. Spirakis, and G. Stachowiak. New clocks, fast line formation and self-replication population protocols. *CoRR*, abs/2111.10822, 2021. `arXiv:2111.10822`.

4   L. Gąsieniec and G. Stachowiak. Enhanced phase clocks, population protocols, and fast space optimal leader election. *J. ACM*, 68(1):2:1–2:21, 2021.

5   D. Doty, M. Eftekhari, L. Gąsieniec, E.E. Severson, P. Uznanski, and G. Stachowiak. A time and space optimal stable population protocol solving exact majority. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021*, pages 1044–1055. IEEE, 2021.