FLORIAN PUCHHAMMER, University of Waterloo

Simulating Chemical Reaction Networks with Randomized Quasi-Monte Carlo

The τ-leap algorithm by Gillespie is used to simulate chemical reaction networks as a discrete time Markov chain, as an alternative to solving the typically high-dimensional chemical master equation. The objective is to estimate the expectation of a function of the molecule numbers at a given future time. We combine this approach with a powerful randomized quasi-Monte Carlo (RQMC) technique, Array-RQMC, which has been proven to outperform Monte Carlo and conventional RQMC in many applications. For chemical reaction networks, this combination can reduce the variance by factors in the thousands, compared to Monte Carlo. Even more than it was observed for ordinary RQMC by Beentjes and Baker (2018).

Array-RQMC efficiently simulates an array of realizations of the Markov chain in parallel, but requires to sort the chains at each step by their (multi-dimensional) states. In our simulations, Array-RQMC always performed better than ordinary MC, regardless of the sort. Yet, the choice of the sorting function can have a significant impact on both variance reduction and simulation speed. We demonstrate how different sorting algorithms can increase the efficiency of the method and discuss how they can be customized for specific applications and for different performance measures.