Abstract Submitted
for the DPP19 Meeting of
The American Physical Society

A machine-learning checkpoint/restart algorithm for particle-in-cell simulations.[1] LUIS CHACON, GUANGYE CHEN, Los Alamos National Laboratory — With ever-increasing computing power and memory capacity, particle check-pointing for fault recovery of particle-in-cell simulations is stressing I/O subsystems, and becoming prohibitive. Given that future exascale computers are expected to be significantly more vulnerable to hard faults than current HPC systems, the availability of a fast and accurate recovery strategy is absolutely essential. In this study, we consider compression of the particle distribution function (PDF) by unsupervised machine-learning techniques.[2] Specifically, we approximate the PDF with a Gaussian mixture.[3] The Gaussian mixture is found by employing maximum likelihood principle with an information criterion, the minimum-message-length principle, for determining an optimal density estimation of the PDF.[2] Restart is conducted by moment-matching sampling of the Gaussian mixture, which strictly conserves charge/mass, momentum, and energy. We demonstrate the effectiveness of the method with various electrostatic and electromagnetic particle-in-cell simulations in 1D and 2D.

[1]This work was performed under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory, managed by Triad National Security, LLC under contract 89233218CNA000001.
[2]G. Chen and L. Chacón, "A machine-learning checkpoint/restart algorithm for particle-in-cell simulations", in preparation
[3]Geoffrey McLachlan and David Peel. Finite Mixture Models. John Wiley & Sons, 2004.

Luis Chacon
Los Alamos National Laboratory

Date submitted: 27 Jun 2019 Electronic form version 1.4