Applications of GPU-Euterpe to W7-X turbulence

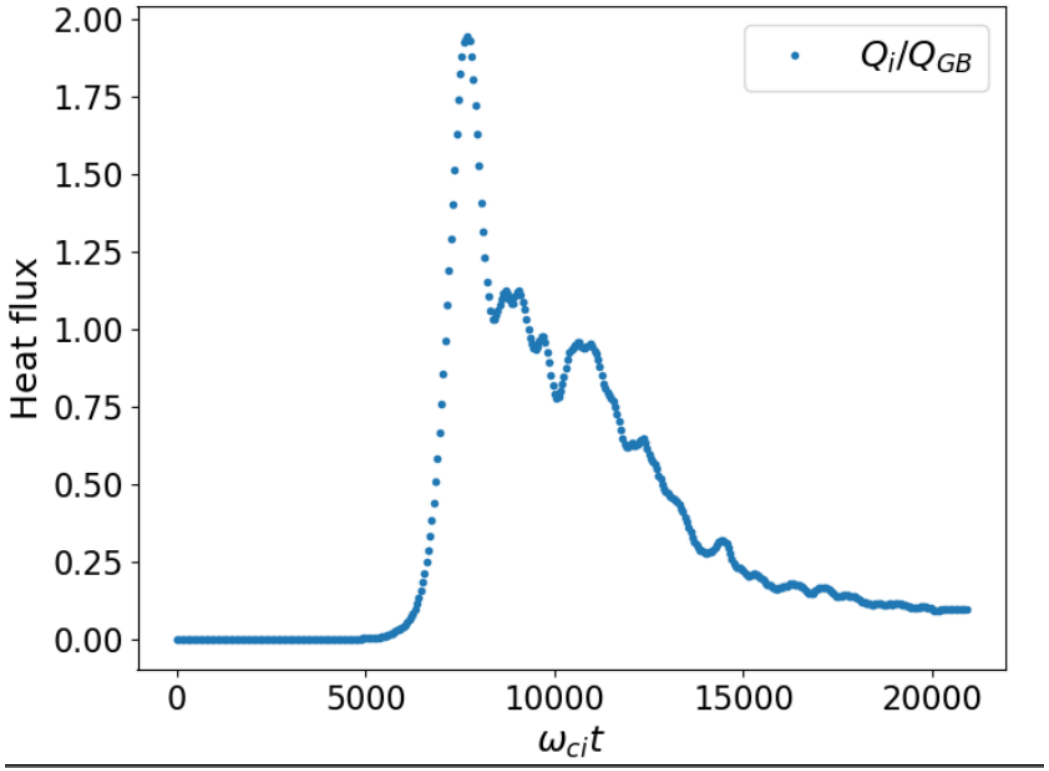

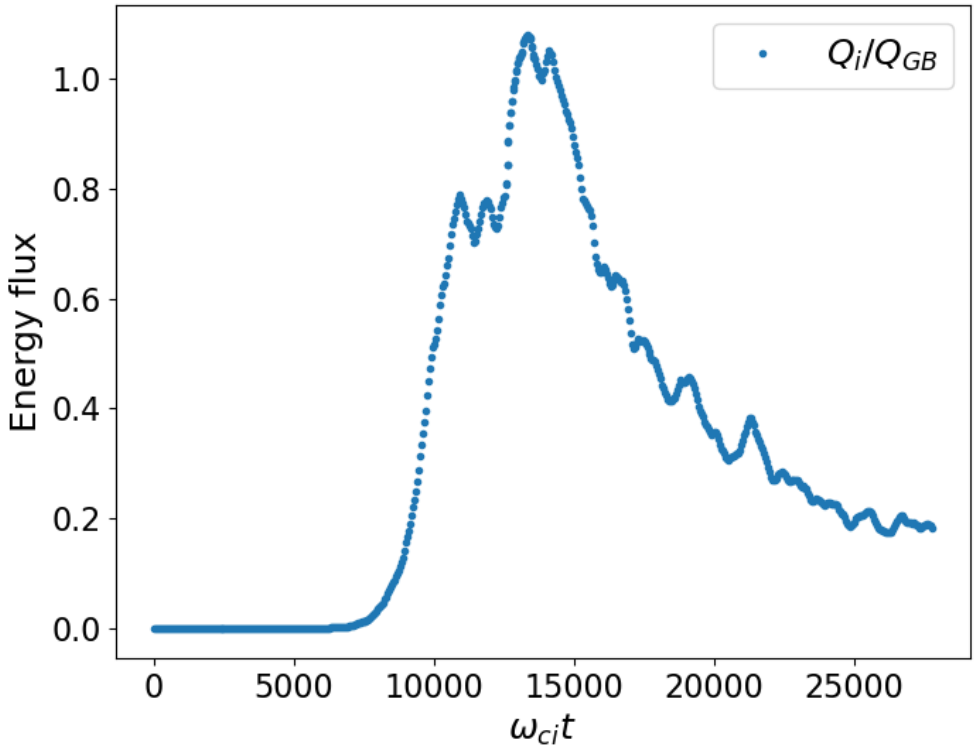

Electrostatic turbulence in W7-X, Lx = 200,
adiabatic electrons; 32 Marconi100 nodes; 20 hours,
450 000 000 ions, 128(rad) x 256(pol) x 64(tor) grid,
dt = 5 (cyclotron units)

Electrostatic turbulence in W7-X, Lx = 400,
adiabatic electrons, 32 Marconi100 nodes, 48 hours,
450 000 000 ion markers, 128(rad) x 512(pol) x 128(tor) grid,
dt = 5 (cyclotron units): can be larger for larger Lx

Applications of GPU-Euterpe to W7-X turbulence

- Results
  - GPU-Euterpe works (proof-of-principle) on Marconi100
  - Computing time comparable to large Marconi jobs
  - Reasonable stellarator turbulence (incl. EM) simulations in reasonable times
- Problems
  - GPU-memory insufficient for larger machine size (e.g. Lx>600)
  - More nodes (> 32) cannot be used on current Marconi100 allocation
- To-do
  - EM linear runs (stellarator KBMs, AEs, kink), hybrid CKA-Euterpe runs
  - Formal scalings for GPU-Euterpe: grid size; several species on GPUs
  - Improvements of algorithm („larmor“ data, omp): ACH request for 2023