

Motivation

• Communication is a time and energy intensive operation compared to computation in large clusters



Courtesy: http://web.eecs.utk.edu/~huangj/hpc/hpc_intro.php

Main Idea

• Do local communication (halo exchange) in events only when value changes by some threshold, otherwise keep using last communicated values



- Due to lesser communication, numerical scheme may take more iterations to converge. However, some iterations do not involve communication, leading to overall savings in time
- Similar to asynchronous communication but saves energy as well
- Use tools from systems and control theory to show convergence asymptotically

Event-Triggered Communication in Parallel Computing

Soumyadip Ghosh¹, Kamal K. Saha¹, Vijay Gupta¹, and Gretar Tryggvason² ¹University of Notre Dame ²Johns Hopkins University

Simulations

- Event-triggered communication is compared with periodic communication (where communication happens with a period)
- Event-triggered communication saves simulation time compared to periodic communication



Event-triggered communication scales better





Periodic Communication

Implementation

• Events of communication only known by sender, not by receiver



Sender PE knows when to send



MPI Send

Sender PE knows when to send

Event-Triggered Communication



Choosing threshold for event-triggered communication is challenging



Adaptive Threshold

Circular domain in processor solved by any local solver

> Overlap region involving communication which can be event-triggered

- however hard to quantify!
- parallel machine learning

[1] S. Ghosh, K. K. Saha, V. Gupta, and G. Tryggvason, "Event-Triggered Communication in Parallel Computing", ScalA workshop, Supercomputing (SC) 2018. [2] S. Ghosh, K. K. Saha, V. Gupta, and G. Tryggvason, "Parallel Computation using Event-Triggered Communication", American Control Conference 2019. Email: sghosh2@nd.edu



Extensions

Adaptive Threshold with extrapolation

Can be generalized to Schwarz methods in a domain decomposition framework



Summary

 Showed faster convergence to same solution by reducing communication systematically Can save energy due to communication as well, Can be applied to parallel graph algorithms and

References