Live Lecture Chat Window
Monday, May 23, 2022

Announcements:

- Reminder: P6 is due Sunday, May 29
- No class (on-campus or Live Lecture) Monday May 30 (Memorial Day!)
- But I will have Office Hours Monday May 30 from 1:00-3:30 PDT
- Last Live Lecture will be Wednesday June 1 (sniffle)

14:58:49 The slower computers do add speed to the cluster though, right? So, they are better than nothing?

Yes, as long as you don't run afoul of the *Compute : Communicate* issues. MPI clusters have relatively slow communication paths since they involve a physical network, not just shared memory.

14:59:56 What is managing the flow of data from the network level to the other CPUs?

It is socket connections between all the processors in the communicator.

15:36:32 For broadcast, why would you send the same data to all the compute nodes? It seems like a waste of compute vs scatter/gather.

It *would* be a waste if you were sending the *entire* dataset to each processor, knowing that each processor will only be working on a subset of the data. Typically, an MPI Broadcast is used to send small things that each processor needs to know about, such as constants (like in the heat transfer example).

15:59:14 Quick question for Project 7A - for the 2nd bumper position and velocity, we're free to customize those right?

Yes, although your life would be lots easier if you used a second sphere. In CG, we *really* like spheres. It is just one total surface and the normal vector (the perpendicular, which you need for bouncing) is simple to compute. It's just a vector from the sphere center to the point where the particle is.