Parallel Programming with MPI

N-Body Codes and Collective Communication

Paul Gray, University of Northern Iowa
**David Joiner, Shodor Education Foundation**
Tom Murphy, Contra Costa College
Henry Neeman, University of Oklahoma
**Charlie Peck, Earlham College**

N Bodies

N-Body Problems

An N-body problem is a problem involving N "bodies" – that is, particles (e.g., stars, atoms) – each of which applies a force to all of the others.

For example, if you have N stars, then each of the N stars exerts a force (gravity) on all of the other $N–1$ stars.

Likewise, if you have N atoms, then every atom exerts a force (nuclear) on all of the other $N–1$ atoms.

1-Body Problem

When *N* is 1, you have a simple 1-Body Problem: a single particle, with no forces acting on it.

Given the particle's position P and velocity V at some time t_0, you can trivially calculate the particle's position at time $t_0+\Delta t$:

$$P(t_0+\Delta t) = P(t_0) + V\Delta t$$

$$V(t_0+\Delta t) = V(t_0)$$

2-Body Problem

When *N* is 2, you have – surprise! – a 2-Body Problem: exactly two particles, each exerting a force that acts on the other.

The relationship between the 2 particles can be expressed as a differential equation that can be solved analytically, producing a closed-form solution.

So, given the particles' initial positions and velocities, you can immediately calculate their positions and velocities at any later time.

N-Body Problems

For N of 3 or more, no one knows how to solve the equations to get a closed form solution.

So, numerical simulation is pretty much the only way to study groups of 3 or more bodies.

Popular applications of N-body codes include astronomy and chemistry.

Note that, for N bodies, there are on the order of N^2 forces, denoted O(N^2).

N Bodies

N-Body Problems

Given N bodies, each body exerts a force on all of the other $N-1$ bodies.

Therefore, there are $N \cdot (N-1)$ forces in total.

You can also think of this as $(N \cdot (N-1))/2$ forces, in the sense that the force from particle A to particle B is the same (except in the opposite direction) as the force from particle B to particle A.

Aside: Big-O Notation

Let's say that you have some task to perform on a certain number of things, and that the task takes a certain amount of time to complete.

Let's say that the amount of time can be expressed as a polynomial on the number of things to perform the task on.

For example, the amount of time it takes to read a book might be proportional to the number of words, plus the amount of time it takes to sit in your favorite easy chair.

$$C_1 \cdot N + C_2$$

# Big-O: Dropping the Low Term

$$C_1 \cdot N + C_2$$

When N is very large, the time spent settling into your easy chair becomes such a small proportion of the total time that it's virtually zero.

So from a practical perspective, for large N, the polynomial reduces to:

$$C_1 \cdot N$$

In fact, for any polynomial, all of the terms except the highest-order term are irrelevant, for large N.

Big-O: Dropping the Constant

$$C_1 \cdot N$$

Computers get faster and faster all the time. And there are many different flavors of computers, having many different speeds.

So, computer scientists don't care about the constant, only about the order of the highest-order term of the polynomial.

They indicate this with Big-O notation:

$$O(N)$$

This is often said as: "of order N."

N-Body Problems

Given *N* bodies, each body exerts a force on all of the other *N*–1 bodies.

Therefore, there are $N \cdot (N-1)$ forces in total.

In Big-O notation, that's $O(N^2)$ forces to calculate.

So, calculating the forces takes $O(N^2)$ time to execute.

But, there are only *N* particles, each taking up the same amount of memory, so we say that N-body codes are of:

- $O(N)$ spatial complexity (memory)
- $O(N^2)$ time complexity

O(N^2) Forces

Note that this picture shows only the forces between A and everyone else.

How to Calculate?

Whatever your physics is, you have some function, F(A,B), that expresses the force between two bodies A and B.

For example,

$$F(A,B) = G \cdot \text{dist}(A,B)^2 \cdot m_A \cdot m_B$$

where G is the gravitational constant and m is the mass of the particle in question.

If you have all of the forces for every pair of particles, then you can calculate their sum, obtaining the force on every particle.

How to Parallelize?

Okay, so let's say you have a nice serial (single-CPU) code that does an N-body calculation.

How are you going to parallelize it?

You could:

- have a master feed particles to processes;

- have a master feed interactions to processes;

- have each process decide on its own subset of the particles, and then share around the forces;

- have each process decide its own subset of the interactions, and then share around the forces.

Do You Need a Master?

Let's say that you have *N* bodies, and therefore you have ½*N*(*N*-1) interactions (every particle interacts with all of the others, but you don't need to calculate both A → B and B → A).

Do you need a master?

Well, can each processor determine on its own either (a) which of the bodies to process, or (b) which of the interactions?

If the answer is yes, then you don't need a master.

Parallelize How?

Suppose you have P processors.

Should you parallelize:

- by assigning a subset of N/P of the bodies to each processor, or

- by assigning a subset of $\frac{1}{2}N(N-1)/P$ of the interactions to each processor?

Data vs. Task Parallelism

- <u>Data Parallelism</u> means parallelizing by giving a subset of the data to each processor.

- <u>Task Parallelism</u> means parallelizing by giving a subset of the tasks to each processor.

OU Supercomputing Center for Education & Research

Data Parallelism for N-Body?

- If you parallelize an N-body code by data, then each processor gets N/P pieces of data.

- For example, if you have 8 bodies and 2 processors, then:
  - P_0 gets the first 4 bodies;
  - P_1 gets the second 4 bodies.

- But, every piece of data (i.e., every body) has to interact with every other piece of data.

- So, every processor will send all of its data to all of the other processors, for every single interaction that it calculates.

Task Parallelism for N-body?

- If you parallelize an N-body code by task, then each processor gets all of the pieces of data that describe the particles (e.g., positions, velocities).

- Then, each processor can calculate its subset of the interaction forces on its own, without talking to any of the other processors.

- But, at the end of the force calculations, everyone must share all of the forces that have been calculated, so that each particle ends up with the total force that acts on it.

- These is called a *global reduction*.

MPI_Reduce

Here's the syntax for **MPI_Reduce**:

```
 MPI_Reduce(sendbuffer, recvbuffer, count,
   datatype, operation, root,
   communicator);
```

For example, to do a sum over all of the particle forces:

```
mpi_error_code =
     MPI_Reduce(
           local_sum_of_particle_forces,
           global_sum_of_particle_forces,
           number_of_particles, MPI_DOUBLE,
           MPI_SUM, master_process,
           MPI_COMM_WORLD);
```

Sharing the Result

In the N-body case, we don't want just one processor to know the result of the sum, we want everyone to know.

So, we could do a reduce followed immediately by a broadcast.

But, MPI gives us a routine that packages all of that for us: `MPI_Allreduce`.

`MPI_Allreduce` is just like `MPI_Reduce` except that every process gets the result (so we drop the `master_process` argument).

MPI_Allreduce

Here's the syntax for **MPI_Allreduce**:

```
mpi_error_code =
    MPI_Allreduce(sendbuffer, recvbuffer,
        count, datatype, operation,
        communicator);
```

For example, to do a sum over all of the particle forces:

```
mpi_error_code =
    MPI_Allreduce(
        local_sum_of_particle_forces,
        global_sum_of_particle_forces,
        number_of_particles, MPI_DOUBLE,
        MPI_SUM, MPI_COMM_WORLD);
```

Collective Communications

A <u>collective communication</u> is a communication that is shared among many processes, not just a sender and a receiver.

MPI_Reduce and **MPI_Allreduce** are collective communications.

Others include: broadcast, gather/scatter, all-to-all.

Collectives Are Expensive

Collective communications are very expensive relative to point-to-point communications, because so much more communication has to happen.

But, they can be much cheaper than doing zillions of point-to-point communications, if that's the alternative.