

Slide 1

Administrivia

- Reminder: Homework 2 deadline extended to Monday.

Slide 2

Example Applications

- Before starting on *Finding Concurrency* patterns — two example applications to be used as running examples.

Slide 3

Example — Molecular Dynamics

- Goal is to simulate what happens to large molecule. Of interest, e.g., in modeling how a drug interacts with a protein.
- Approach is to treat molecule as a collection of balls (atoms) connected by springs (chemical bonds). Then do “standard time-stepping” — divide time into discrete steps, and at each step use classical mechanics to figure out new positions for atoms based on current positions and forces among them.
In more details . . .

Slide 4

Molecular Dynamics — Computation

- At each time step:
 - Compute forces (vibrational and rotational) on atoms caused by chemical bonds between them. Short-range interaction, so not too much computation here.
 - Compute forces on atoms caused by their electrical charges. Potentially must consider all pairs of atoms, so lots of computation here.
 - Use forces to update atoms' positions and velocities.
 - Compute other physical properties of the system — e.g., energies.
- To reduce the computational load, can limit computation of electrical-charge-induced forces to atoms that are “close”. To do this, calculate for each atom a list of “neighbors”. If time steps are short, atoms don't move much, and we don't have to do this every step.

Molecular Dynamics Pseudocode

Slide 5

```

Int const N // number of atoms
Array of Real :: atoms (3,N) //3D coordinates
Array of Real :: velocities (3,N) //velocity vector
Array of Real :: forces (3,N) //force in each dimension
Array of List :: neighbors(N) //atoms in cutoff volume

loop over time steps
  vibrational_forces (N, atoms, forces)
  rotational_forces (N, atoms, forces)
  neighbor_list (N, atoms, neighbors)
  non_bonded_forces (N, atoms, neighbors, forces)
  update_atom_positions_and_velocities
    (N, atoms, velocities, forces)
  physical_properties ( ... Lots of stuff ... )
end loop

```

Pseudocode for Non-Bonded Force Computation

Slide 6

```

function non_bonded_forces (N, Atoms, neighbors, Forces)
  Int const N // number of atoms
  Array of Real :: atoms (3,N) //3D coordinates
  Array of Real :: forces (3,N) //force in each dimension
  Array of List :: neighbors(N) //atoms in cutoff volume
  Real :: forceX, forceY, forceZ

  loop [i] over atoms
    loop [j] over neighbors(i)
      forceX = non_bond_force(atoms(1,i), atoms(1,j))
      forceY = non_bond_force(atoms(2,i), atoms(2,j))
      forceZ = non_bond_force(atoms(3,i), atoms(3,j))
      force(1,i) += forceX;   force(1,j) -= forceX;
      force(2,i) += forceY;   force(2,j) -= forceY;
      force(3,i) += forceZ;   force(3,j) -= forceZ;
    end loop [j]
  end loop [i]
end function non_bonded_forces

```

Example — Heat Diffusion

Slide 7

- A simple example, representative of a big class of scientific-computing applications — “heat distribution problem”.
- Goal is to simulate what happens when two ends of a pipe are put in contact with things at different (constant) temperatures — pipe conducts heat, its temperature changes over time, eventually converging on a smooth gradient.
- Can model mathematically how temperature in pipe changes over time using partial differential equations.
- Can approximate solution by “discretizing” — spatially and with regard to time.

Heat Diffusion Code

Slide 8

```
double *uk = malloc(sizeof(double) * NX);
double *ukpl = malloc(sizeof(double) * NX);
double *temp;
double dx = 1.0/NX; double dt = 0.5*dx*dx;
double maxdiff, diff;

initialize(uk, ukpl);

for (int k = 0; (k < NSTEPS) && (maxdiff >= threshold); ++k) {

    /* compute new values */
    for (int i = 1; i < NX-1; ++i) {
        ukpl[i]=uk[i]+ (dt/(dx*dx))*(uk[i+1]-2*uk[i]+uk[i-1]);
    }

    /* check for convergence */
    maxdiff = 0.0;
    for (int i = 1; i < NX-1; ++i) {
        diff = fabs(uk[i] - ukpl[i]);
        if (diff > maxdiff) maxdiff = diff;
    }

    /* "copy" ukpl to uk by swapping pointers */
    temp = ukpl; ukpl = uk; uk = temp;

    printValues(uk, k);
}
```

Finding Concurrency Design Space

Slide 9

- Starting point in our grand strategy for developing parallel applications.
Overall idea — capture how experienced parallel programmers think about initial design of parallel applications. Might not be necessary if clear match between application and an *Algorithm Structure* pattern.
- Idea is to work through three groups of patterns in sequence (possibly with backtracking):
 - Decomposition patterns (*Task Decomposition, Data Decomposition*):
Break problem into tasks that maybe can execute concurrently.
 - Dependency analysis patterns (*Group Tasks, Order Tasks, Data Sharing*):
Organize tasks into groups, analyze dependencies among them.
 - *Design Evaluation*: Review what you have so far, possibly backtrack.
- Keep in mind — best to focus attention on computationally intensive parts of problem.

Task-Based Versus Data-Based Decomposition

Slide 10

- Two basic approaches to decomposing a problem — task-based and data-based. Usually one will seem more logical than the other, but may need to think through both.
- Either way, you'll look at both tasks and data; difference is in which you look at first, and then the other follows.

Slide 11

Task Decomposition

- Goal here is to break up (some of) computation into “tasks” — logical elements of overall computation that might be independent enough to do concurrently.
- At this stage, try to stay abstract and portable; also try to identify lots of tasks (can always recombine them later if too many), as independent of each other as possible.
- Places to look for tasks include groups of function calls (e.g., in divide-and-conquer strategy), loop iterations (e.g., many examples we’ve discussed).
- Simple example — matrix multiplication.
- Once you have this, consider data related to each task (*Data Decomposition*).

Slide 12

Data Decomposition

- Goal here is to break up (some of) problem data into parts (“chunks”) that can be operated on concurrently. Good choice if most computation consists of updates to big data structure(s).
- Again, try to stay abstract and portable; also try to “parameterize” decomposition so you can easily try various choices at runtime.
- Data structures to look at include arrays, recursive structures such as trees.
- Simple example — matrix multiplication.
- Once you have this, consider computation related to each chunk of data (*Task Decomposition*).

Decomposition — Examples

- Next slides will show working through our two examples. For purposes of illustration, we'll do one starting with a *Task Decomposition* and inferring a *Data Decomposition*, the other one the other way around.

Slide 13

Molecular Dynamics Example — Task Decomposition

- Tasks that find the vibrational forces on an atom.
- Tasks that find the rotational forces on an atom.
(Together, these are tasks to compute “bonded forces” — those due to chemical bonds.)
- Tasks that find the non-bonded forces on an atom (the ones due to electrical charges).
- Tasks that update the position and velocity of an atom.
- Tasks that update the neighbor list for an atom. (Or we could consider updating all the neighbor lists as one task, as in the book, if we think it won't be done very often and therefore is not worthwhile to parallelize.)

Slide 14

Slide 15

Molecular Dynamics Example — Data Decomposition

- Key data structures:
 - An array of atom coordinates, one element per atom.
 - An array of atom velocities, one element per atom.
 - An array of lists, one per atom, each defining the neighborhood of atoms considered to be “close”.
 - An array of forces on atoms, one element per atom.
- Decompose each of these to match task decomposition — into elements corresponding to individual atoms.

Slide 16

Heat Diffusion Example — Data Decomposition

- Key data structures:
 - Array for “old values” (time step k).
 - Array for “new values” (time step $k + 1$).
- Most computation involves updating or otherwise operating on these two arrays. Think of partitioning into “chunks”.

Heat Diffusion Example — Task Decomposition

- Tasks to compute new values from old values, one per chunk.
- Tasks to compute maximum difference between new and old values, one per chunk.
- Task to swap pointers (to fake copying new values to old values).

Slide 17

Group Tasks

- Once you've broken down problem into tasks / data chunks, need to put it back together as design for parallel algorithm.
- First step — look for “groups of tasks” — logically related, or interdependent, or all with same constraints, etc. Often just one group.

Slide 18

Order Tasks

- Next step — identify constraints on groups of tasks. Possibilities:
 - “First this, then that.”
 - “All of these together.”

Slide 19

Molecular Dynamics Example — Group Tasks, Order Groups

- Task groups based on list of a few slides back — each type of task (e.g., compute rotational forces) defines a task group.
- Ordering constraints, for each timestep:
 - Task group to compute neighbor list must run before task group to compute non-bonded forces.
 - Task groups to compute bonded and non-bonded forces must run before task group to update positions and velocities.
 - Task group to update positions and velocities must run before next timestep.
- (Also see Figure 3.4 in the book.)

Slide 20

Slide 21

Heat Diffusion Example — Group Tasks, Order Groups

- Task groups based on list of a few slides back — each type of task (e.g., compute new values from old values) defines a task group.
- Ordering constraints, for each timestep:
 - Three tasks groups must run in sequence.
 - All task groups must run before next timestep.

Slide 22

Data Sharing

- Sometimes tasks are totally independent, each executes on totally separate data, etc. Usually not, though. Point here is to think through dependencies.
- Useful to think in terms of:
 - “Task-local” data — variables used only/mainly by single task, particularly the ones being updated. Example — chunks in heat diffusion problem.
 - Globally shared data — variables not associated with any particular task(s). Example — sum in numerical integration problem.
 - Data shared among smaller groups of tasks. Example — “boundary” points in heat diffusion problem.

Data Sharing, Continued

Slide 23

- Potential problems different in different environments; goal is to ensure correctness without adding too much overhead:
 - With shared memory, all UEs (can) have access to all data, but must use synchronization to prevent “race conditions”.
 - With distributed memory, each UE has its own data, so race conditions not possible, but must use communication to (in effect) share data.
- Basic approach — first identify what data is shared, then figure out how it's used.

Data Sharing — Categories of Shared Data

Slide 24

- Read-only: Easiest case. If shared memory, don't need to do anything. If distributed memory, consider giving each process a copy. Examples include global constants.
- Effectively-local (large data structure, but each element accessed by only one UE): Also easy. If distributed memory, give each process “its” data.

Slide 25

Data Sharing — Categories of Shared Data, Continued

- Read-write (accessed by more than one task, at least one changing it): Can be arbitrarily complicated, but some common cases aren't too bad:
 - “Accumulate” (variable(s) used to accumulate result — usually a reduction). Example — sum in numerical integration problem. Give each task (or each UE) a copy and combine at end.
 - “Multiple-read/single-write” (multiple tasks need initial value, one task computes new value). Example — points near boundaries of chunks in heat diffusion problem. Create at least two copies, one for task that computes new value, other(s) to hold initial value for other tasks.

Slide 26

Molecular Dynamics Example — Analyze Task/Data Dependencies

- Arrays of atom positions, velocities:
 - Read-only for most groups of tasks — but tasks may need access to many elements, so for distributed memory might want to duplicate.
 - Updated by one group of tasks, but each task updates its own element(s) — “effectively local”.
- Array of forces:
 - Read-only for group of tasks that update positions and velocities, and each task needs access only to “local” data.
 - Updated by several groups of tasks, but updates fit “accumulate data” model.

Molecular Dynamics Example — Task/Data Dependencies, Continued

Slide 27

- Array of neighbor lists:
 - Read-only for group of tasks that compute “non-bonded” forces, and each task needs access only to local data.
 - Updated by one group of tasks, but each task updates its own element(s).
- (Also see Figure 3.5 in book.)

Heat Diffusion Example — Analyze Task/Data Dependencies

Slide 28

- Arrays of old, new values:
 - Old values read-only for all groups of tasks, and each task needs access mostly to local data — plus “boundary values” for neighboring tasks.
 - New values updated by one group of tasks, and each task computes values only for “its” elements.

For distributed memory, could distribute among processes, with extra variable(s) to hold copy of boundary values.
- Maximum difference between old, new values is “accumulate data” in one group of tasks, read-only elsewhere.
- Pointers to old/new values — changed at end of time step by one task, read-only elsewhere. Could duplicate for distributed memory.

Design Evaluation

- Idea of this pattern — questions to ask yourself about design/analysis before going further, to reduce odds of costly mistakes.
- Ideal design is easy to implement/maintain and produces a fast program suitable for target architecture. (But keep in mind old saying from engineering: “Good, fast, cheap. Pick any two.”)

Slide 29

Design Evaluation — Suitability for Target Platform

- How many processing elements (PEs) are available? Need at least one task per PE, often want many more — unless we can easily get exactly one task per PE at runtime, with good load balance.
- How are data structures shared among PEs? If there's a lot of shared data, or sharing is very “fine-grained”, implementing for distributed memory will likely not be easy or fast.
- How many UEs are available and how do they share data? Similar to previous questions, but in terms of UEs — with some architectures, can have multiple UEs per PE, e.g., to hide latency. For this to work, “context switching” must be fast, and problem must be able to take advantage of it.
- How does time spent doing computation compare to overhead of synchronization/communication, on target platform? May be a function of problem size relative to number of PEs/UEs.

Slide 30

Slide 31

Design Evaluation — Design Quality

- Is it flexible? Will it adapt well to a range of platforms (if appropriate), differing numbers of UEs/PEs, different problem sizes? Does it deal gracefully with “boundary cases”?
- Is it efficient? Can you get good load balance? Is overhead minimal? consider UE creation and scheduling, communication, and synchronization.
- Is it (paraphrasing Einstein) “as simple as possible, but not simpler”? Is it reasonable to think mortals can produce working code relatively quickly? which can later be ported and/or enhanced?

Slide 32

Design Evaluation — Preparation for Next Phase

- How regular are tasks and their data dependencies?
- Are interactions between tasks (or groups of tasks) synchronous or asynchronous?
- Are tasks grouped in the best way?

Slide 33

Molecular Dynamics Example — Design Evaluation

- Major phases of computation seem to involve a lot of tasks, so we can take advantage of many processors.
- Data sharing seems more suited to shared memory than distributed memory, but the latter could work if we just duplicate data (have to think about how well that would “scale”).
- Tasks and data are fairly regular, with one exception: how many neighbors an atom has might vary a lot. Probably will affect how we split up work among UEs.
- Interaction among tasks is synchronous.

Slide 34

Heat Diffusion Example — Design Evaluation

- Major phases of computation seem to involve a lot of tasks, so we can take advantage of many processors.
- Data sharing seems suitable for either shared or distributed memory.
- Tasks and data are very regular, interaction is synchronous.

Minute Essay

- None — sign in.

Slide 35