

Slide 1

Administrivia

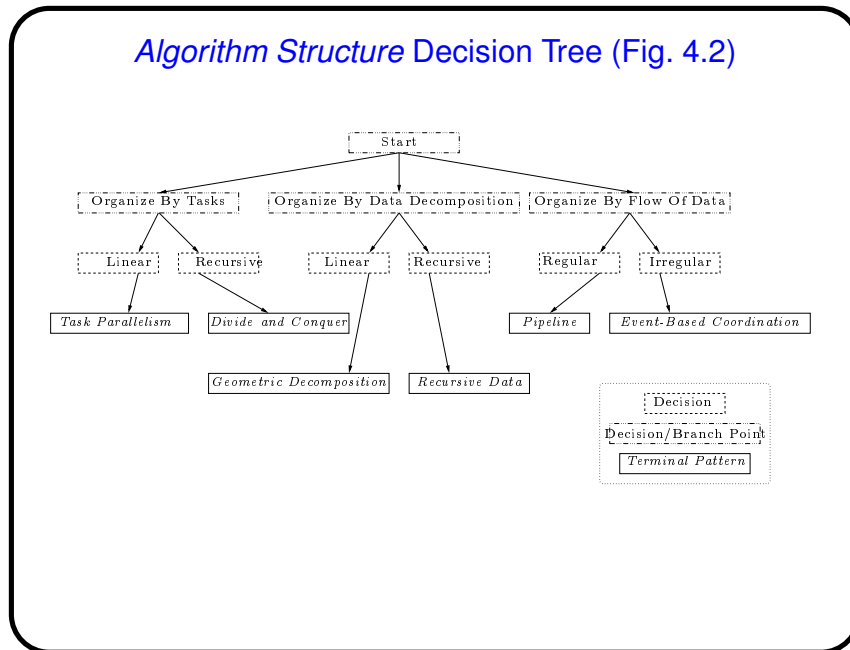
- (None? other than to nag you about Homework 3?)

Slide 2

Algorithm Structure Design Space

- Historical note: These are the patterns with the longest history. We started out trying to identify commonly-used overall structures for parallel programs (these patterns), and then at some point added the other “design spaces”.
- After much thought, writing, revision, and arguing, we came up with . . .

Slide 3



Slide 4

Algorithm Structure Patterns

- If decomposition/analysis reveals organization in terms of tasks: *Task Parallelism* (probably most common strategy) or *Divide and Conquer*.
- If decomposition/analysis reveals organization in terms of data: *Geometric Decomposition* (second most common strategy) or *Recursive Data*.
- If organization is in terms of flow of data: *Pipeline* or *Event-Based Coordination*.
- Note, by the way, use of a "pattern format" — consistent list of sections ("Problem", "Context", etc.). Supposedly this is one reason design patterns are effective ways of presenting information.

Slide 5

Task Parallelism

- Problem statement:
When the problem is best decomposed into a collection of tasks that can execute concurrently, how can this concurrency be exploited efficiently?
- Key ideas in solution — managing tasks (getting them all scheduled), detecting termination, managing any data dependencies.
- Many, many examples, including:
 - Numerical integration example (next slide).
 - Molecular dynamics example (after that).
 - Mandelbrot set computation.
 - Branch-and-bound computations: Maintain list of “solution spaces”. At each step, pick item from list, examine it, and either declare it a solution, discard it, or divide it into smaller spaces and put them back on list. Tasks consist of processing items from list.

Slide 6

Numerical Integration and Finding Concurrency Patterns

- A task decomposition probably makes sense here, with the tasks being the iterations of the main loop.
- There's only one group of tasks, and the tasks in the group can execute concurrently.
- Data shared among tasks includes a read-only variable (`step`), a variable that could be made task-local (`x`), and an “accumulate data” variable (`sum`).

Numerical Integration and *Task Parallelism*

Slide 7

- How to define tasks so we get “enough but not too many”?
One task per loop iteration is really too many, since each task is so small, but we can get away with it if we keep the overhead of managing the tasks small — as all our solutions do.
- How to manage data dependencies (if any)?
Dependency involving `x` can be managed by just giving each UE its own copy.
Dependency involving `sum` can be managed by giving each UE a local copy and combining all copies at end.
- How to assign tasks to UEs? statically (at compile time) or dynamically (at runtime)?
All tasks are the same size, so static assignment will work and probably be most efficient.

Molecular Dynamics and *Task Parallelism*

Slide 8

- How to define tasks so we get “enough but not too many”?
One task per atom pair is too many; one task per atom is probably right.
- How to manage data dependencies (if any)?
Dependency involving `forces` array — potentially any UE can write to any element, if we exploit symmetry resulting from Newton’s third law. But computation is accumulation/reduction, so just give each UE a local copy and combine all copies at end.
- How to assign tasks to UEs? statically (at compile time) or dynamically (at runtime)?
Work per task can vary, since how many atoms are “close” varies. Decide at next level.

Geometric Decomposition

Slide 9

- Problem statement:
How can an algorithm be organized around a data structure that has been decomposed into concurrently updatable “chunks”?
- Key ideas in solution — distributing data, arranging for needed communication.
- Probably second most common pattern. Examples include:
 - Heat-diffusion problem previously discussed (next slide).
 - Matrix multiplication using blocks.

Heat Diffusion and Geometric Decomposition

Slide 10

- How to distribute data?
One chunk per UE will probably work well. (Note that for other problems it might not.) Might be nice to include in data structure a place to store values from neighboring chunks. More in *Distributed Array*, next chapter.
- How to synchronize/communicate?
With shared memory, just need barrier synchronization.
With distributed memory, need to exchange values with neighbor UEs, also perform reduction.

Divide and Conquer

Slide 11

- Problem statement:
Suppose the problem is formulated using the sequential divide and conquer strategy. How can the potential concurrency be exploited?
- Key idea in solution — create new task(s) every time we split (sub)problem, recombine when we merge.
- Examples include mergesort and some non-naive algorithms for N -body problem.
- Straightforward if you already have a sequential divide-and-conquer solution, but scalability is somewhat limited.

Recursive Data

Slide 12

- Problem statement:
Suppose the problem involves an operation on a recursive data structure (such as a list, tree, or graph) that appears to require sequential processing. How can operations on these data structures be performed in parallel?
- Key idea in solution — “out of the box” thinking to expose concurrency.
- Probably least-used structure at the time of writing the book (because it doesn’t map well to then-current architectures); included for completeness and because examples are interesting — e.g. “roots in forest” example.

Remaining Patterns

- Next time: *Pipeline and Event-Based Coordination*.

Slide 13

Sidebar: gnuplot

- Tool I like for both quick interactive plots and nice-looking ones to use in papers: `gnuplot`. Available on most UNIX-like systems and (I think!) optionally for other operating systems. Home page at `gnuplot.sourceforge.net`. Can do 2D and 3D plots, the former with Cartesian or polar coordinates.
(Interestingly(?) enough, the name has nothing to do with the GNU project!)
- To start it, `gnuplot`. Brings up a command-line interface. Online help available with `help`.

Slide 14

gnuplot, Continued

Slide 15

- Useful commands include `plot` to plot function(s) or data from file(s), `set` to set various things (e.g., x and y ranges).
- Default output to terminal, but with `set terminal` and `set output` you can instead store to a file in various formats.
- Can also put commands (`plot` etc.) in a file and execute batch-style, or with `load`. Useful if you want to regenerate plots when data changes.
- (Examples.)

Minute Essay

Slide 16

- None — I'll just record attendance.