

More Patterns for Parallel Application Programs*

Berna L. Massingill[†] Timothy G. Mattson[‡]
Beverly A. Sanders[§]

Abstract

We are involved in an effort to develop a pattern language for parallel application programs. The pattern language consists of a set of patterns that guide the programmer through the entire process of developing a parallel program, including patterns that help find the concurrency in the problem, patterns that help find the appropriate algorithm structure to exploit the concurrency in parallel execution, and patterns describing lower-level implementation issues. The current version of the pattern language can be seen at <http://www.cise.ufl.edu/research/ParallelPatterns>.

In this paper, we present three patterns from our pattern language, selected from the set of patterns that are used after the problem has been analyzed to identify the exploitable concurrency. *ChooseStructure* addresses the question of how to select an appropriate pattern from the others in this set. *DivideAndConquer* is used when the problem can be solved by recursively dividing it into subproblems, solving each subproblem independently, and then recombining the subsolutions into a solution to the original problem. *PipelineProcessing* is used when the problem can be decomposed into ordered groups of tasks connected by data dependencies.

1 Introduction

1.1 Overview

We are involved in an effort to design a pattern language for parallel application programs. The goal of the pattern language is to lower the barrier to parallel programming by guiding the programmer through the entire process of developing a parallel program. In our vision of parallel program development, the programmer brings into the process a good understanding of the actual problem to be solved, then works through

*Copyright © 2001, Berna L. Massingill. Permission is granted to copy for the PLoP 2001 conference. All other rights reserved.

[†]Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL; blm@cise.ufl.edu (current address: Department of Computer Science, Trinity University, San Antonio, TX; bmassing@trinity.edu).

[‡]Parallel Algorithms Laboratory, Intel Corporation; timothy.g.mattson@intel.com.

[§]Department of Computer and Information Science and Engineering, University of Florida, Gainesville, FL; sanders@cise.ufl.edu.

the pattern language, eventually obtaining a detailed design or even working code. The pattern language is organized into four design spaces, which are visited in order.

- The *FindingConcurrency* design space includes high-level patterns that help find the concurrency in a problem and decompose it into a collection of tasks. (The patterns in this design space are presented in [7].)
- The *AlgorithmStructure* design space contains patterns that help find an appropriate algorithm structure to exploit the concurrency that has been identified. (Three patterns from this design space are presented in [6].)
- The *SupportingStructures* design space includes patterns that describe useful abstract data types and other supporting structures.
- The *ImplementationMechanisms* design space contains patterns that describe lower-level implementation issues.

The latter two design spaces (slightly stretching the typical notion of a pattern) might even include reusable code libraries or frameworks. We use a pattern format for all four levels so that we can address a variety of issues in a unified way. The current, incomplete, version of the pattern language can be seen at <http://www.cise.ufl.edu/research/ParallelPatterns>. It consists of a collection of extensively hyperlinked documents, such that the designer can begin at the top level and work through the pattern language by following links.

In this paper, rather than describing the pattern language as a whole, we present the complete text of three selected patterns from the *AlgorithmStructure* design space. The chosen patterns are relatively mature, and significant enough to stand alone. The patterns in the *AlgorithmStructure* design space help the designer find an appropriate algorithm structure suitable for parallel implementation and are applicable after the concurrency in a problem has been identified. Thus, before attempting to apply these patterns the designer should have determined (i) how to decompose the problem into tasks that can execute concurrently, (ii) which data is local to the tasks and which is shared among tasks, and (iii) what ordering and data dependencies exist among tasks.

The concurrency in parallel programs introduces potentially nondeterministic behavior and the possibility of race conditions. Correctness concerns thus play a large role in parallel programming and are addressed by describing constraints on the problem and implementation. The goal is to provide rules that, if followed, will preclude concurrency errors. These constraints are typically described first in informal but precise language; in some cases this informal discussion is followed by a more formal and detailed discussion including references to supporting theory.

1.2 Example

As mentioned above, before attempting to apply patterns from the *AlgorithmStructure* design space, the designer should have analyzed the problem to identify potential concurrency. As an example of this analysis, consider the following problem taken from the field of medical imaging. (This example is presented in more detail in [7].) We can decompose this problem in two ways — in terms of tasks and in terms of data.

An important diagnostic tool is to give a patient a radioactive substance and then watch how that substance propagates through the body by looking at the distribution of emitted radiation. Unfortunately, the images are of low resolution, due in part to the scattering of the radiation as it passes through the body. It is also difficult to reason from the absolute radiation intensities, since different pathways through the body attenuate the radiation differently.

To solve this problem, medical imaging specialists build models of how radiation propagates through the body and use these models to correct the images. A common approach is to build a Monte Carlo model. Randomly selected points within the body are assumed to emit radiation (usually a gamma ray), and the trajectory of each ray is followed. As a particle (ray) passes through the body, it is attenuated by the different organs it traverses, continuing until the particle leaves the body and hits a camera model, thereby defining a full trajectory. To create a statistically significant simulation, thousands if not millions of trajectories are followed.

The problem can be parallelized in two ways. Since each trajectory is independent, it would be possible to parallelize the application by associating each trajectory with a task. Another approach would be to partition the body into sections and assign different sections to different processing elements.

As in many ray-tracing codes, there are no dependencies between trajectories, making the task-based decomposition the natural choice. By eliminating the need to manage dependencies, the task-based algorithm also gives the programmer plenty of flexibility later in the design process, when how to schedule the work on different processing elements becomes important.

The data decomposition, however, is much more effective at managing memory utilization. This is frequently the case with a data decomposition as compared to a task decomposition. Since memory is decomposed, data-decomposition algorithms also tend to be more scalable. These issues are important and point to the need to at least consider the types of platforms that will be supported by the final program. The need for portability drives one to make decisions about target platforms as late as possible. There are times, however, when delaying consideration of platform-dependent issues can lead one to choose a poor algorithm.

1.3 In this paper

The remainder of this paper consists of the complete text of three patterns from the *AlgorithmStructure* design space:

- *ChooseStructure* addresses the question of how to select an appropriate pattern from the others in this set.
- *DivideAndConquer* is used when the problem can be solved by recursively dividing it into subproblems, solving each subproblem independently, and then recombining the subsolutions into a solution to the original problem.
- *PipelineProcessing* is used when the problem can be decomposed into ordered groups of tasks connected by data dependencies.

Each numbered major section represents one document in the collection of hyperlinked documents making up our pattern language; each document represents one pattern. To make the paper self-contained, we replace hyperlinks with text formatted like this and footnotes or citations. To make it easier to identify patterns and pattern sections, we format pattern names as *SomePattern* and pattern section names as **SomeSection**.

2 The ChooseStructure Pattern

Problem

After you have analyzed your problem to identify exploitable concurrency, how do you use the results to choose a structure for the parallel algorithm?

Context

The first phase of designing a parallel algorithm usually consists of analyzing the problem to identify exploitable concurrency, usually by using the patterns of the *FindingConcurrency*¹ design space. After performing this analysis, you should have (1) a way of decomposing the problem into a number of tasks, (2) an understanding of how the problem's data is decomposed onto and shared among the tasks, and (3) an ordering of task groups to express temporal or other constraints among the tasks. To refine the design further and move it closer to a program that can execute these tasks concurrently, you need to map the concurrency onto the multiple units of execution (UEs)² that run on a parallel computer.

Of the countless ways to define an algorithm structure, most follow one of nine basic design patterns. The key issue is to decide which pattern is most appropriate for your problem.

Forces

There are competing forces to keep in mind in deciding which overall structure fits your problem best:

- Different aspects of the analysis may pull the design in the direction of different structures.
- A good algorithm design must strike a balance between (1) abstraction and portability and (2) suitability for a particular target architecture. The challenge faced by the designer, especially at this early phase of the algorithm design, is to leave the parallel algorithm design abstract enough to support portability while ensuring that it can eventually be implemented effectively for the parallel systems on which it will be executed.

This pattern describes a way of balancing these forces and choosing an overall structure for the algorithm.

¹A set of patterns in our pattern language; see [7].

²Generic term for a collection of concurrently-executing entities, usually either processes or threads.

Solution

Overview

Our approach to using the results of the preliminary analysis to choose an overall structure for the algorithm has four major steps:

- **Target platform:** What constraints are placed on the parallel algorithm by the target machine or programming environment?
- **Major organizing principle:** When you consider the concurrency in your problem, is there a particular way of looking at it that stands out and provides a high-level mechanism for organizing this concurrency? Notice that in some situations, a good design may make use of multiple algorithm structures (combined hierarchically, compositionally, or in sequence), and this is the point at which to consider whether such a design makes sense for your problem.
- **The *AlgorithmStructure* decision tree:** For each subset of tasks, how do you select an *AlgorithmStructure* design pattern that most effectively defines how to map the tasks onto UEs?
- **Re-evaluation:** Is this chosen *AlgorithmStructure* pattern (or patterns) suitable for your target platform?

Steps

- **Consider the target platform.** What constraints are placed on your design by the target computer and its supporting programming environments? In an ideal world, it would not be necessary to consider such questions at this stage of the design, and doing so works against keeping the program portable (which is also desirable). This is not an ideal world, however, and if you do not consider the major features of your target platform, you risk coming up with a design that is difficult to implement efficiently.

The primary issue is how many UEs (processes or threads) your system will effectively support, since an algorithm that works well for ten UEs may not work well at all for hundreds of UEs. You do not necessarily need to decide on a specific number (in fact to do so would overly constrain the applicability of your design), but you do need to decide on an order-of-magnitude number of UEs.

Another issue is how expensive it is to share information among UEs. If there is hardware support for shared memory, information exchange takes place through shared access to common memory, and frequent data sharing makes sense. If the target is a collection of nodes connected by a slow network, however, sharing information is very expensive, and your parallel algorithm must avoid communication wherever possible.

When thinking about both of these issues — the number of UEs and the cost of sharing information — avoid the tendency to over-constrain your design. Software usually outlives hardware, so over the course of a program's life, you may

need to support a tremendous range of target platforms. You want your algorithm to meet the needs of your target platform, but at the same time, you want to keep it flexible so it can adapt to different classes of hardware.

Also, remember that in addition to multiple UEs and some way to share information among them, a parallel computer has one or more programming environments that can be used to implement parallel algorithms. Do you know which parallel programming environment you will use for coding your algorithm? If so, what does this imply about how tasks are created and how information is shared among UEs?

- **Identify the major organizing principle.** Consider the concurrency you found using the patterns of the *FindingConcurrency* design space. It consists of tasks and groups of tasks, data (both shared and task-local), and ordering constraints among task groups. Your next step is to find an algorithm structure that represents how this concurrency maps onto the UEs. The first step is to find the *major organizing principle* implied by the concurrency. This usually falls into one of three camps: *organization by orderings*, *organization by tasks*, or *organization by data*. Algorithms in the first two groups are task-parallel, since the design is guided by how the computation is decomposed into tasks. Algorithms in the third group are data-parallel, because how the data is decomposed guides the algorithm. We now consider each of these in more detail.

For some problems, the major feature of the concurrency is the presence of well-defined interacting groups of tasks, and the key issue is how these groups are ordered with respect to each other. For example, a GUI-driven program might be parallelized by decomposing it into a task that accepts user input, a task that displays output, and one or more background computational tasks, with the tasks interacting via “events” (e.g., the user does something, or a part of the background computation completes). Here the major feature of the concurrency is the way in which these distinct task groups interact.

For other problems, there is really only one group of tasks active at one time, and the way the tasks within this group interact is the major feature of the concurrency. Examples include so-called “embarrassingly parallel” programs in which the tasks are completely independent, as well as programs in which the tasks in a single group cooperate to compute a result.

Finally, for some problems, the way data is decomposed and shared among tasks stands out as the major way to organize the concurrency. For example, many problems focus on the update of a few large data structures, and the most productive way to think about the concurrency is in terms of how this structure is decomposed and distributed among UEs. Programs to solve differential equations or perform linear algebra often fall into this category, since they are frequently based on updating large data structures.

As you think about your problem and search for the most productive way to begin organizing your concurrency and mapping it onto UEs, remember that the most effective parallel algorithm design may be hierarchical or compositional: It often happens that the very top level of the design is a sequential composition

of one or more *AlgorithmStructure* patterns (for example, a loop whose body is an instance of an *AlgorithmStructure* pattern). Other designs may be organized hierarchically, with one pattern used to organize the interaction of the major task groups and other patterns used to organize tasks within the groups.

- **Identify the pattern(s) to use.** Having considered the questions raised in the preceding sections, you are now ready to select an algorithm structure, guided by an understanding of constraints imposed by your target platform, an appreciation of the role of hierarchy and composition, and a major organizing principle for your problem. You make the selection by working through the decision tree presented in Figure 1 Starting at the top of the tree, consider your concurrency

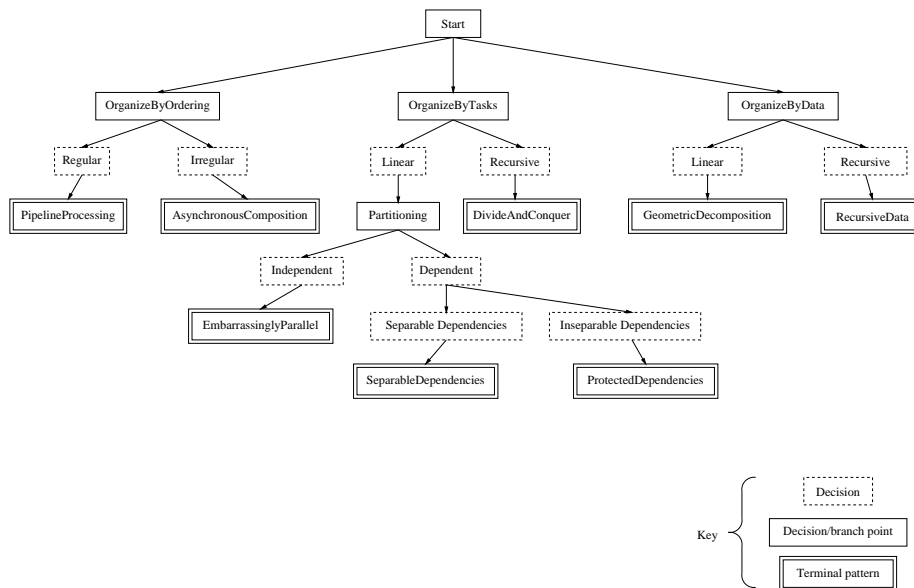


Figure 1: Decision tree for the *AlgorithmStructure* design space.

and the major organizing principle, and use this information to select one of the three branches of the tree; then follow the discussion below for the appropriate subtree. Notice again that for some problems the final design may combine more than one algorithm structure; if no one of these structures seems suitable for your problem, it may be necessary to divide the tasks making up your problem into two or more groups, work through this procedure separately for each group, and then determine how to combine the resulting algorithm structures.

- *Organize By Ordering.* Select the *OrganizeByOrdering* subtree when the major organizing principle is how the groups of tasks are ordered with respect to each other. This pattern group has two members, reflecting two ways task groups can be ordered. One choice represents “regular” order-

ings that do not change during the algorithm; the other represents “irregular” orderings that are more dynamic and unpredictable.

- * *PipelineProcessing*³: The problem is decomposed into ordered groups of tasks connected by data dependencies.

- * *AsynchronousComposition*⁴: The problem is decomposed into groups of tasks that interact through asynchronous events.

- *Organize By Tasks*. Select the *OrganizeByTasks* subtree when the execution of the tasks themselves is the best organizing principle. There are many ways to work with such “task-parallel” problems, making this the largest group of patterns. The first decision to make is how the tasks are enumerated. If they can be gathered into a set linear in any number of dimensions, take the *Partitioning* branch. If the tasks are enumerated by a recursive procedure, take the *Tree* branch.

If the *Partitioning* branch is selected, the next question is to consider the dependencies among the tasks. If there are no dependencies among the tasks (i.e., the tasks do not need to exchange information), then you can use the following pattern for your algorithm structure:

- * *EmbarrassinglyParallel*⁵: The problem is decomposed into a set of independent tasks. Most algorithms based on task queues and random sampling are instances of this pattern.

If there are dependencies among the tasks, you need to decide how they can be resolved. For a large class of problems, the dependencies are expressed by write-once updates or associative accumulation into shared data structures. In these cases, the dependencies are separable and you can use the following algorithm structure:

- * *SeparableDependencies*⁶: The parallelism is expressed by splitting up tasks among units of execution. Any dependencies among tasks can be pulled outside the concurrent execution by replicating the data prior to the concurrent execution and then reducing the replicated data after the concurrent execution. That is, once the data has been replicated, the problem is greatly simplified and looks similar to the embarrassingly parallel case.

If the dependencies involve true information-sharing among concurrent tasks, however, you cannot use a trick to make the concurrency look like the simple embarrassingly parallel case. There is no way to get around explicitly managing the shared information among tasks, and you will probably need to use the following pattern:

- * *ProtectedDependencies*⁷: The parallelism is expressed by splitting up tasks among units of execution. In this case, however, variables in-

³Section 4 of this paper.

⁴A pattern in the *AlgorithmStructure* design space.

⁵A pattern in the *AlgorithmStructure* design space; see [6].

⁶A pattern in the *AlgorithmStructure* design space; see [6].

⁷A pattern in the *AlgorithmStructure* design space.

volved in data dependencies are both read and written during the concurrent execution and thus cannot be pulled outside the concurrent execution but must be managed during the concurrent execution of the tasks. Notice that the exchange of messages is logically equivalent to sharing a region of memory, so this case covers more than just traditional shared-memory programs.

This completes the *Partitioning* branch; now consider the patterns in the *Tree* branch. Here we have two cases. These cases are very similar, differing in how the subproblems are solved once the tasks have been recursively generated:

- * *DivideAndConquer*⁸: The problem is solved by recursively dividing it into subproblems, solving each subproblem independently, and then recombining the subsolutions into a solution to the original problem.
- *Organize By Data*. Select the *OrganizeByData* subtree when the decomposition of the data is the major organizing principle in understanding the concurrency. There are two patterns in this group, differing in how the decomposition is structured — linearly in each dimension or recursively.
 - * *GeometricDecomposition*⁹: The problem space is decomposed into discrete subspaces; the problem is then solved by computing solutions for the subspaces, with the solution for each subspace typically requiring data from a small number of other subspaces. Many instances of this pattern can be found in scientific computing, where it is useful in parallelizing grid-based computations, for example.
 - * *RecursiveData*¹⁰: The problem is defined in terms of following links through a recursive data structure.
- **Re-evaluate**. After choosing one or more *AlgorithmStructure* patterns to be used in your design, skim through their descriptions to be sure they are reasonable suitable for your target platform. (For example, if your target platform consists of a large number of workstations connected by a slow network, and one of your chosen *AlgorithmStructure* patterns requires frequent communication between tasks, you are likely to have trouble implementing your design efficiently.) If the chosen patterns seem wildly unsuitable for your target platform, try identifying a secondary organizing principle and working through the preceding step again.

Examples

Medical imaging

For example, consider the medical imaging problem described in *DecompositionStrategy*¹¹. This application simulates a large number of gamma

⁸Section 3 of this paper.

⁹A pattern in the *AlgorithmStructure* design space; see [6].

¹⁰A pattern in the *AlgorithmStructure* design space.

¹¹A pattern in the *FindingConcurrency* design space; see [7]. This example is also summarized in Section 1 of this paper.

rays as they move through a body and out to a camera. One way to describe the concurrency is to define the simulation of each ray as a task. Since they are all logically equivalent, we put them into a single task group. The only data shared among the tasks are read-only accesses to a large data structure representing the body. Hence the tasks do not depend on each other.

Because for this problem there are many independent tasks, it is less necessary than usual to consider the target platform: The large number of tasks should mean that we can make effective use of any (reasonable) number of UEs; the independence of the tasks should mean that the cost of sharing information among UEs will not have much effect on performance.

Thus, we should be able to choose a suitable structure by working through the decision tree in Figure 1. Given that in this problem the tasks are independent, the only issue we really need to worry about as we select an algorithm structure is how to map these tasks onto UEs. That is, for this problem the major organizing principle seems to be the way the tasks are organized, so we start by following the *OrganizeByTasks* branch.

We now consider the nature of our set of tasks. Are the tasks arranged hierarchically, or do they reside in an unstructured or flat set? For this problem, the tasks are in an unstructured set with no obvious hierarchical structure among them, so we follow the *Partitioning* branch of the decision tree.

The next decision takes into account the dependencies among the tasks. In this example, the tasks are independent. This implies that the algorithm structure to use for this problem is described in *EmbarrassinglyParallel*. (Notice, by the way, that an embarrassingly parallel algorithm should not be viewed as trivial to implement — much careful design work is still needed to come up with a correct program that efficiently schedules the set of independent tasks for execution on the UEs.)

Finally, we review this decision in light of possible target-platform considerations. As we observed earlier, the key features of this problem (the large number of tasks and their independence) make it unlikely that we will need to reconsider because the chosen structure will be difficult to implement on the target platform. Nevertheless, to be careful we also review *EmbarrassinglyParallel*; fortunately, it appears to be suitable for a variety of target platforms.

3 The *DivideAndConquer* pattern

Problem

How can you exploit the potential concurrency in a problem that can be solved using the divide-and-conquer strategy?

Context

Consider the divide-and-conquer strategy employed in many sequential algorithms. With this strategy, a problem is solved by splitting it into subproblems, solving them independently, and merging their solutions into a solution for the whole problem. The

subproblems can be solved directly, or they can in turn be solved using the same divide-and-conquer strategy, leading to an overall recursive program structure. The potential concurrency in this strategy is not hard to see: Since the subproblems are solved independently, their solutions can be computed concurrently. Figure 2 illustrates the strategy and the potential concurrency.




Figure 2: The divide-and-conquer strategy.

The divide-and-conquer strategy can be more formally described in terms of the following functions (where N is a constant):

- `Solution solve(Problem P)`: Solve a problem (returns its solution).
- `Problem[] split(Problem P)`: Split a problem into N subproblems, each strictly smaller than the original problem (returns the subproblems).
- `Solution merge(Solution[] subS)`: Merge N subsolutions into solution (returns the merged solution).
- `boolean baseCase(Problem P)`: Decide whether a problem is a “base case” that can be solved without further splitting (returns `true` if base case, `false` if not).
- `Solution baseSolve(Problem P)`: Solve a base-case problem (returns its solution).

The strategy then leads to the top-level program structure shown in Figure 3.

```
Solution solve(Problem P) {
    if (baseCase(P))
        return baseSolve(P);
    else {
        Problem subProblems[N];
        Solution subSolutions[N];
        subProblems = split(P);
        for (int i = 0; i < N; i++)
            subSolutions[i] = solve(subProblems[i]);
        return merge(subSolutions);
    }
}
```

Figure 3: Sequential pseudocode for the divide-and-conquer strategy.

Indications

Use *DivideAndConquer* when:

- The problem can be solved using the divide-and-conquer strategy, with subproblems being solved independently.

This pattern is particularly effective when:

- The amount of work required to solve the base case is large compared to the amount of work required for the recursive splits and merges.
- The split produces subproblems of roughly equal size.

Forces


- The traditional divide-and-conquer strategy is a widely useful approach to algorithm design. Algorithms based on this strategy are almost trivial to parallelize based on the obvious exploitable concurrency.
- As Figure 2 suggests, however, the amount of exploitable concurrency varies over the life of the program; at the outermost level of the recursion (initial split and final merge) there is no exploitable concurrency, while at the innermost level (base-case solves) the number of concurrently-executable tasks is the number of base-case problems (which is often the same as the problem size). Ideally, you would like to always have at least as many concurrently-executable tasks as processors, and clearly this pattern falls short in that respect, and the problem only gets worse as you increase the number of processors, so in general this pattern does not scale well.
- This pattern is more efficient when the subproblems into which each problem is split are roughly equal in size / computational complexity.

Solution

Overview

If the subproblems of a given problem can be solved independently, then you can solve them in any order you like, including concurrently. This means that you can produce a parallel application by replacing the `for` loop of Figure 3 with a parallel-`for` construct, so that the subproblems will be solved concurrently rather than in sequence. It is worth noting at this point that program correctness is independent of whether the subproblems are solved sequentially or concurrently, so you can even design a hybrid program that sometimes solves them sequentially and sometimes concurrently, based on which approach is likely to be more efficient (more about this later).

This strategy can be mapped onto a design in terms of tasks by defining one task for each invocation of the `solve` function, as illustrated in Figure 4 (rectangular boxes correspond to tasks).



(boxes correspond to tasks)

Figure 4: Parallelizing the divide-and-conquer strategy.

Note the recursive nature of the design, with each task in effect generating and then absorbing a subtask for each subproblem.

Note also that either the `split` or the `merge` phase can be essentially absent:

- No `split` phase is needed if all the base-case problems can be derived directly from the whole problem (without recursive splitting). In this case, the overall design will look like the bottom half of Figures 2 and 4.

- No merge phase is needed if the problem can be considered solve when all of the base-case problems have been identified and solved. In this case, the overall design will look like the top half of Figures 2 and 4.

Designs based on this pattern include the following key elements:

- Definitions of the functions described in the **Context** section above (`solve`, `split`, `merge`, `baseCase`, and `baseSolve`).
- A way of scheduling the tasks that efficiently exploits the available concurrency (subproblems can be solved concurrently).

Key elements

- **Definitions of functions.** It is usually straightforward to produce a program structure that defines the required functions: What is required is almost the same as the equivalent sequential program, except for code to schedule tasks, as described in the next section.
- **Scheduling the tasks.** Where a parallel divide-and-conquer program differs from its sequential counterpart is that the parallel version is also responsible for scheduling the tasks in a way that exploits the potential concurrency (subproblems can be solved concurrently) efficiently.

The simplest approach is to simply replace the sequential `for` loop over subproblems with a `parallel-for` construct, allowing the corresponding tasks to execute concurrently. (Thus, in Figure 4, the two lower-level splits execute concurrently, the four base-case solves execute concurrently, and the two lower-level merges execute concurrently.) To improve efficiency (as discussed later in this section), you can also use a combination of `parallel-for` constructs and sequential `for` loops, typically using `parallel-for` at the top levels of the recursion and sequential `for` at the more deeply nested levels. In effect, this approach combines parallel divide-and-conquer with sequential divide-and-conquer.

Correctness issues

Most of the correctness issues in implementing this pattern are the same ones involved in sequential divide-and-conquer, plus a few additional restrictions to make the concurrency work properly.

- Considering first the sequential divide-and-conquer strategy expressed in the pseudocode of Figure 3, you can guarantee that `solve(P)` returns a correct solution of `P` if the other functions meet the following specifications, expressed in terms of preconditions and postconditions. (As before, `N` is an integer constant.)

```
- Solution baseSolve(Problem P):  
  Precondition: baseCase(P) = true.  
  Postcondition: returned value is a solution of P.
```

- `Problem[] split(Problem P)`:
Precondition: `baseCase(P) = false`.
Postcondition: returned value is an array of N subproblems, each strictly smaller than P , whose solutions can be combined to give a solution of P . Here, “strictly smaller” means smaller with respect to some integer measure that, when small enough, indicates a base-case problem. (This ensures that the recursion is finite.)
 - `Solution merge(Solution[] subS)`:
Precondition: `subS` is an array of N solutions such that for some problem P and array of subproblems `subP = split(P)`, `subS[i]` is a solution of `subP[i]`, for all i from 1 through N .
Postcondition: returned value is a solution of P .
- To make the concurrency work, it is sufficient for the solutions of subproblems to be computed independently. That is, for two distinct subproblems `subP[i]` and `subP[j]` of P , `solve(subP[i])` and `solve(subP[j])` must be computed independently. This will be true if neither call to `solve` modifies variables shared with the other call. If the solutions of subproblems cannot be computed independently, then any access to shared variables must be protected with appropriate synchronization. This, of course, tends to reduce the efficiency of the calculation.

Efficiency issues

Effective use of this pattern depends on reducing the fraction of the program’s lifespan during which there are fewer concurrently-executable tasks than processors, and there are several factors that contribute to this goal:

- Having a problem whose split and merge operations are computationally trivial compared to its base-case solve.
- Having a problem size that is large compared to the maximum number of processors available on the target environment.
- Reducing the number of levels of recursion required to arrive at the base-case solve by splitting each problem into more subproblems. This generally requires some algorithmic cleverness but can be quite effective, especially in the limiting case of “one-deep divide-and-conquer”, in which the initial split is into P subproblems, where P is the number of available processors. See the **Related Patterns** section for more discussion of this strategy.
- If problem size is large compared to the number of available processors, at some point in the computation the number of concurrently-executable tasks will exceed the number of processors. If you take the simple approach of always using the `parallel-For` construct to schedule the tasks corresponding to subproblems, this approach produces a situation in which at some point the number of units of

execution exceeds the number of available processors. If such a situation would be inefficient in the target environment (i.e., if context-switching among UEs¹² is expensive), or if there is significant overhead associated with the `parallel-for` construct, it will probably be more efficient to use the `parallel-for` construct only for the outer levels of the recursion, switching to a sequential loop when the total number of subproblems (number of subproblems per split multiplied by recursion level) exceeds number of available processors.

Examples

Mergesort

Mergesort is a well-known sorting algorithm based on the divide-and-conquer strategy, applied as follows to sort an array of N elements:

- The base case is an array of size 1, which is already sorted and can thus be returned without further processing.
- In the split phase, the array is split by simply partitioning it into two contiguous subarrays, each of size $N/2$ (or $(N+1)/2$ and $(N-1)/2$, if N is odd).
- In the solve-subproblems phase, the two subarrays are sorted (by applying the mergesort procedure recursively).
- In the merge phase, the two (sorted) subarrays are recombined into a single sorted array in the obvious way.

This algorithm is readily parallelized by performing the two recursive mergesorts in parallel.

Matrix diagonalization

[3] describes a parallel algorithm for diagonalizing (computing the eigenvectors and eigenvalues of) a symmetric tridiagonal matrix T . The problem is to find a matrix Q such that $Q^T \cdot T \cdot Q$ is diagonal; the divide-and-conquer strategy goes as follows (omitting the mathematical details):

- The base case is a 1-by-1 matrix, which is already diagonal and can be returned without further processing.
- The split phase consists of finding matrix T' and vectors u, v , such that $T = T' + uv^T$, and T' has the form

$$\begin{bmatrix} T_1 & 0 \\ 0 & T_2 \end{bmatrix}$$

where T_1 and T_2 are symmetric tridiagonal matrices (which can be diagonalized by recursive calls to the same procedure).

¹²Units of execution — generic term for a collection of concurrently-executing entities, usually either processes or threads.

- The merge phase recombines the diagonalizations of T_1 and T_2 into a diagonalization of T .

Details can be found in [3] or in [4].

Other known uses

- Classical graph and other algorithms. Any introductory algorithms text will have many examples of algorithms based on the divide-and-conquer strategy, most of which can be parallelized with this pattern. (As noted in the **Consequences** section, however, such parallelizations are not always efficient.)
- Leslie Grignard’s Fast Multipole Algorithm.
- Floating Point Systems’ FFT (Fast Fourier Transform — an algorithm for computing discrete Fourier transforms).
- Tree-based reductions, particularly for the PRAM model, as described in [5].
- Certain well-known algorithms for solving the N -body problem, for example the Barnes-Hut algorithm and some algorithms of John Salmon.

Related Patterns

It is interesting to note that just because an algorithm is based on a (sequential) divide-and-conquer strategy does not mean that it must be parallelized with *DivideAndConquer*. A hallmark of this pattern is the recursive arrangement of the tasks, leading to a varying amount of concurrency. Since this can be inefficient, it is often better to rethink the problem such that it can be mapped onto some other pattern, such as *GeometricDecomposition* or *SeparableDependencies*¹³.

4 The PipelineProcessing Pattern

Problem

If your problem can be solved by an algorithm in which data flows through a sequence of tasks or stages (a *pipeline*), how can you exploit the potential concurrency in this approach?

Context

The basic idea of this pattern is much like the idea of an assembly line: To perform a sequence of essentially identical calculations, each of which can be broken down into the same sequence of steps, we set up a “pipeline”, one stage for each step, with all stages potentially executing concurrently. Each of the sequence of calculations is performed by having the first stage of the pipeline perform the first step, and then

¹³Patterns in the *AlgorithmStructure* design space.

the second stage the second step, and so on. As each stage completes a step of a calculation, it passes the calculation-in-progress to the next stage and begins work on the next calculation.

This may be easiest to understand by thinking in terms of the assembly-line analogy: For example, suppose the goal is to manufacture a number of cars, where the manufacture of each car can be separated into a sequence of smaller operations (e.g., installing a windshield). Then we can set up an assembly line (pipeline), with each operation assigned to a different worker. As the car-to-be moves down the assembly line, it is built up by performing the sequence of operations; each worker, however, performs the same operation over and over on a succession of cars.

Returning to a more abstract view, if we call the calculations to be performed C_1 , C_2 , and so forth, then we can describe operation of a *PipelineProcessing* program thus: Initially, the first stage of the pipeline is performing the first operation of C_1 . When that completes, the second stage of the pipeline performs the second operation on C_1 ; simultaneously, the first stage of the pipeline performs the first stage of C_2 . When both complete, the third stage of the pipeline performs the third operation on C_1 , the second stage performs the second operation on C_2 , and the first stage performs the first operation on C_3 . Figure 5 illustrates how this works for a pipeline consisting of four stages.

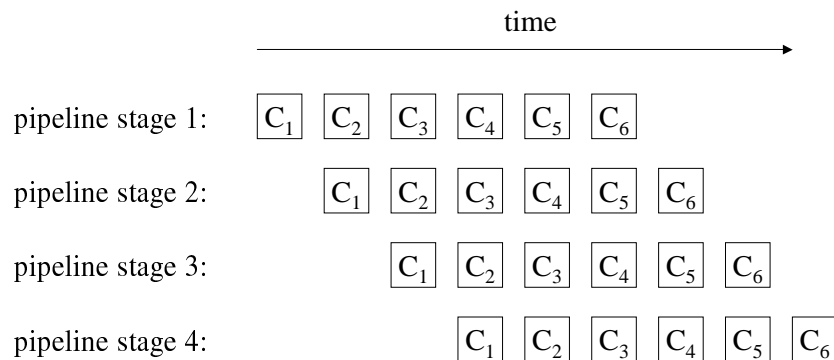


Figure 5: Pipeline stages.

This idea can be extended to include situations in which some operations can be performed concurrently. Figure 6 illustrates two pipelines, each with four stages. In the second pipeline, the third stage consists of two operations that can be performed concurrently.

These figures suggest that we can represent a pipeline as a directed graph, with vertices corresponding to elements of the calculation and edges indicating dataflow. To preserve the idea of the pipeline, it is necessary that this graph be acyclic, and it is

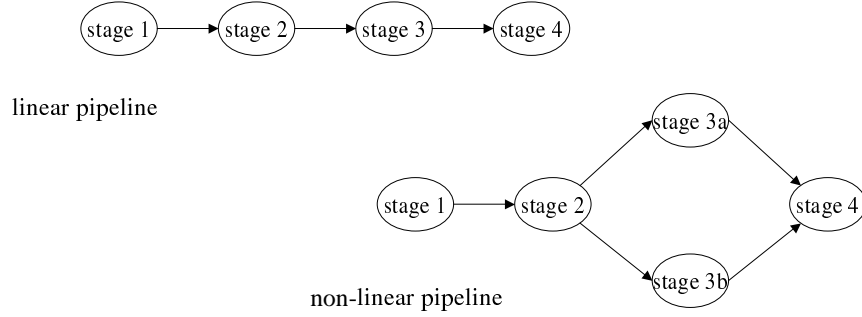


Figure 6: Example pipelines.

probably best if it does not depart too much from a basically linear structure, in which the elements can be divided into stages, with each stage communicating only with the previous and next stages.

We can describe the linear case more formally, as follows: This pattern describes computations whose goal is to take a sequence of inputs in_1, in_2 , etc. and compute a sequence of outputs out_1, out_2 , etc., where the following is true:

- out_i can be computed from in_i as a composition of N functions $f^{(1)}, f^{(2)}$, and so on, where (letting \circ denote function composition)

$$out_i = f^{(N)} \circ \dots \circ f^{(2)} \circ f^{(1)}(in_i)$$

- For i and j different, and m and n different, the computation of $out_i^{(m)}$ is independent of the computation $out_j^{(n)}$, where we have defined input and output sequences for the functions $f^{(k)}$ as follows:

- $out_i^{(m)} = f^{(m)}(in_i^{(m)})$
- $in_i^{(1)} = in_i$
- $out_i^{(N)} = out_i$
- $in_i^{(m+1)} = out_i^{(m)}$, for m between 1 and $N - 1$.

To restate this less formally: For different stages m and n of the pipeline, and different elements C_i and C_j of the sequence of calculations to be performed, stage m of the calculation for C_i can be done independently of stage n of the calculation for C_j . This is the key restriction in this pattern and is what makes the concurrency possible.

Indications

Use *PipelineProcessing* when:

- The problem consists of performing a sequence of calculations, each of which can be broken down into distinct stages, on a sequence of inputs, such that for each input the calculations must be done in order, but it is possible to overlap computation of different stages for different inputs as indicated in Figures 5 and 6.

The pattern is particularly effective when:

- The number of calculations is large compared to the number of stages.
- It is possible to dedicate a processor to each element, or at least each stage, of the pipeline.

This pattern can also be effective in combination with other patterns:

- As part of a hierarchical design, in which the tasks making up each stage of the pipeline are internally organized using another of the *AlgorithmStructure* patterns.

Forces

- This pattern can be straightforward to implement (as described in the **Solution** section below), particularly for message-passing platforms.
- However, in a pipeline algorithm, concurrency is limited until all the stages are occupied with useful work. This is referred to as “filling the pipeline”. At the tail end of the computation, again there is limited concurrency as the final item works its way through the pipeline. This is called “draining the pipeline”. In order for pipeline algorithms to be effective, the time spent filling or draining the pipeline must be small compared to the total time of the computation. This pattern therefore is most effective when the number of calculations is large compared to the number of operations/stages required for each one.
- Also, either the stages of the pipeline must be kept synchronized or there must be some way of buffering work between successive stages. The pattern therefore usually works better if the operations performed by the various stages of the pipeline are all about equally computationally intensive. If the stages in the pipeline vary widely in computational effort. The slowest stage defines a bottleneck for the algorithm’s aggregate throughput. Furthermore, a much slower stage in the middle of the pipeline will cause data items to back up on the input queue, potentially leading to buffer overflow problems.

Solution

Overview

Viewing the pattern in terms of tasks, define one task for each element of the pipeline (one element per stage in a linear pipeline, possibly more for a nonlinear pipeline). Each task can be thought of as having a predecessor (in the previous stage) and a successor (in the next stage), with obvious exceptions for tasks corresponding to the first and last stages of the pipeline and a straightforward generalization to nonlinear pipelines (where a task can have multiple predecessors or successors). Data dependencies are defined as follows. Each task requires as input a sequence of input items from its predecessor (or synchronized sequences from its predecessors); for each input item it generates an output item to be used by its successor (or synchronized groups of items, each group to be distributed among its successors).

Designs based on this pattern include the following key elements:

- A way of defining the elements of the pipeline, where each element corresponds to one of the functions that make up the computation. In a linear pipeline (such as the “linear pipeline” of Figure 6), these elements are the stages of the pipeline; in a more complex pipeline (such as the “non-linear pipeline” of Figure 6) there can be more than one element per pipeline stage. Each pipeline element will correspond to one task.
- A way of representing the dataflow among pipeline elements, i.e., how the functions are composed.
- A way of scheduling the tasks.

Key elements

- **Defining the elements of the pipeline.** What is needed here is a program structure to hold the computations associated with each stage. There are many ways to do this. A particularly effective approach is an SPMD¹⁴ program in which the ID associated with the UE¹⁵ selects options in a case statement, with each case corresponding to a stage of the pipeline.
- **Representing the dataflow among pipeline elements.** What is needed here is a mechanism that provides for the orderly flow of data between stages in the pipeline. This is relatively straightforward to accomplish in a message-passing environment by assigning one process to each function and implementing each function-to-function connection (between successive stages of the pipeline) via a sequence of messages between the corresponding tasks. Since the stages are hardly ever perfectly synchronized, and the amount of work carried out at different stages almost always varies, this flow of data between pipeline stages must

¹⁴“Single program, multiple data” — a parallel-programming style in which each thread or process runs the same program, but on different data.

¹⁵Unit of execution — generic term for one of a collection of concurrently-executing entities, usually either processes or threads.

usually be both buffered and ordered. Most message-passing environment (e.g., MPI) make this easy to do.

If a message-passing programming environment is not a good fit with the target platform, you will need to explicitly connect the stages in the pipeline with a buffered channel. Such a buffered channel can be implemented as a shared queue (using *SharedQueue*¹⁶).

- **Scheduling the tasks.** What is needed here is a way of scheduling the tasks that make up the design. Usually all tasks are scheduled to execute concurrently (e.g., as elements of an SPMD program), since this avoids bottlenecks and potential deadlock.

Correctness issues

- What makes concurrency possible here is the requirement that for different stages m and n of the pipeline, and different elements C_i and C_j of the sequence of calculations to be performed, stage m of the calculation for C_i can be done independently of stage n of the calculation for C_j . The tasks that make up the pipeline should therefore be independent, except for the interaction needed to pass data from one stage to the next. This happens naturally in a distributed-memory environment; in a shared-memory environment, it can be guaranteed by (i) making sure the mechanism used to pass data from one pipeline stage to the next is correctly implemented (e.g., by using a concurrency-safe shared data structure such as *SharedQueue*) and (ii) not allowing tasks to modify any shared variables except those used in this mechanism.

Examples

Fourier-transform computations

A type of calculation widely used in signal processing involves performing the following computation repeatedly on different sets of data:

- Perform a discrete Fourier transform (DFT) on a set of data.
- Manipulate the result of the transform elementwise.
- Perform an inverse DFT on the result of the manipulation.

Examples of such calculations include convolution, correlation, and filtering operations, as discussed in [8].

A calculation of this form can easily be performed by a three-stage pipeline:

- The first stage of the pipeline performs the initial Fourier transform; it repeatedly obtains one set of input data, performs the transform, and passes the result to the second stage of the pipeline.

¹⁶A pattern in the *SupportingStructures* design space.

- The second stage of the pipeline performs the desired elementwise manipulation; it repeatedly obtains a partial result (of applying the initial Fourier transform to an input set of data) from the first stage of the pipeline, performs its manipulation, and passes the result to the third stage of the pipeline.
- The third stage of the pipeline performs the final inverse Fourier transform; it repeatedly obtains a partial result (of applying the initial Fourier transform and then the elementwise manipulation to an input set of data) from the second stage of the pipeline, performs the inverse Fourier transform, and outputs the result.

Each stage of the pipeline processes one set of data at a time. However, except during the initial “filling” of the pipeline, all stages of the pipeline can operate concurrently; while the first stage is processing the N -th set of data, the second stage is processing the $(N - 1)$ -th set of data and the third stage is processing the $(N - 2)$ -th set of data.

Known uses

- Image processing applications.
- Problems in the CMU task-parallel Fortran test suite [1].
- Other pipelined computations, as described in [5].

Related Patterns

This pattern is very similar to the *Pipes and Filters* pattern of [2]; the key difference is that this pattern explicitly discusses concurrency. This pattern is similar to *AsynchronousComposition*¹⁷ in that both patterns apply to problems where it is natural to decompose the computation into a collection of semi-independent entities. The difference is that in *PipelineProcessing*, these entities interact in a more regular way, with all “stages” of the pipeline proceeding in a loosely synchronous way, whereas in the *AsynchronousComposition* pattern, there is no such requirement, and the entities can interact in very irregular and asynchronous ways.

Acknowledgments

We thank Intel Corporation, the National Science Foundation (grant #9704697), and the Air Force Office of Scientific Research (grant #4514209-12) for their financial support. We also thank Doug Lea and Alan O’Callaghan, who acted as shepherds for our two previous PLoP papers ([6] and [7]), and Bob Hanmer, our shepherd for this paper.

¹⁷Another pattern in the *AlgorithmStructure* design space.

References

- [1] The CMU task-parallel fortran test suite. Available at <http://www.cs.cmu.edu/afs/cs.cmu.edu/project/iwarp/member/fx/public/www/fx.html>.
- [2] Frank Buschmann, Regine Meunier, Hans Rohnert, Peter Sommerlad, and Michael Stal. *Pattern-Oriented Software Architecture, Volume 1: A System of Patterns*. John Wiley & Son Ltd, 1996.
- [3] J.J. Dongarra and D.C. Sorensen. A fully parallel algorithm for the symmetric eigenvalue problem. *SIAM J. Sci. and Stat. Comp.*, 8:S139–S154, 1987.
- [4] G.H. Golub and C.F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, 1989.
- [5] J. Jájá. *An Introduction to Parallel Algorithms*. Addison-Wesley, 1992.
- [6] Berna L. Massingill, Timothy G. Mattson, and Beverly A. Sanders. Patterns for parallel application programs. In *Proceedings of the Sixth Pattern Languages of Programs Workshop (PLoP99)*, 1999. See also our Web site at <http://www.cise.ufl.edu/research/ParallelPatterns>.
- [7] Berna L. Massingill, Timothy G. Mattson, and Beverly A. Sanders. Patterns for finding concurrency for parallel application programs. In *Proceedings of the Seventh Pattern Languages of Programs Workshop (PLoP'00)*, August 2000. See also our Web site at <http://www.cise.ufl.edu/research/ParallelPatterns>.
- [8] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 1988.