

**The Architecture of Problems and
Portable Parallel Software Systems**

Geoffrey Fox

**CRPC-TR91172
August, 1991**

Center for Research on Parallel Computation
Rice University
P.O. Box 1892
Houston, TX 77251-1892

**The Architecture of Problems and
Portable Parallel Software Systems***

Geoffrey C. Fox
Northeast Parallel Architectures Center
111 College Place
Syracuse, NY 13244-4100
315-443-1723
gcf@nova.npac.syr.edu

Abstract

We show how the structure or architecture of applications suggest the nature of parallel software systems that will run portably on a variety of parallel machines - both those available now and those expected during the coming decade. The discussion is illustrated by lessons learned from real applications implemented on current MIMD and SIMD machines. These are mainly academic problems and the extrapolation to complex industrial and government applications is unproven but we believe our methodology will still be applicable. We suggest that problems consist of a heterogeneous mixture of at least three distinct classes for which the parallel software support could be different. Further, we need approaches to manage the integration of these different software paradigms.

*This work was partially supported by ASAS Program Office and the National Science Foundation under Cooperative Agreement No. CCR-8809165-- the Government has certain rights in this material.

I: Introduction

Over the last decade, parallel computing has been explored in a research environment and there seems little doubt that

Parallel Computing Works.

What does one mean by this and what does it imply? Current commercial parallel computers are immature, but they offer better peak performance and better cost performance than conventional machines for many problems. A large number of parallel applications have been developed. There are only a tiny fraction of the computations running on sequential machines, but the initial parallel experimental applications cover a wide range of algorithms and perhaps the "essence" (computational kernels) of the majority of large scale scientific and engineering calculations. Clearly, the main limitation to the rapid spread of parallel computing is not hardware or algorithms but rather software. In this paper, we show how portable parallel software systems can be optimized around three broad application classes, or more pretentiously, three problem architectures. The current analysis is incomplete; we only have extensive experience with small academic codes - typically up to 10,000 lines in length. The extrapolation to the much larger and complex industrial and government applications is uncertain. For instance, we have little understanding of the parallelization of problems requiring real-time and extensive input/output support.

In Figure 1, we follow Kennedy (private communication, 1988) and divide software into five layers, where we are most interested here in the second layer which should be portable to a variety of hardware architectures, including both SIMD and MIMD machines. We propose that the process of theory, simulation and computation be viewed as a sequence of mappings between related systems as shown in Figure 2. The goal of the high level software in the second level of Figure 1 is to be general enough to represent a reasonable number of problems but not so general (such a sequential Fortran) that we have lost so much specific information about the problem that it cannot be parallelized. Both problems and machines have an architecture and the software system is dependent on both. However, we expect that the high level software systems discussed here will be

built around the problem and not the machine architecture. The latter is reflected in low level machine dependent support software which should be hidden from most users. We can probably persuade many users to produce parallel versions of their code; however, they will be loath to do this more than once. It is not reasonable that each new parallel architecture require a significantly different software implementation.

In Section II we review an architectural classification for problems, and in Section III we analyze current parallel software experience from this point of view. This leads to a tentative parallel software strategy for a general application which we present in Section IV. In Section V, we discuss software support for synchronous problems and in Section VI, the harder and still uncertain irregular loosely synchronous case is treated briefly.

II: Problem Architectures

We have introduced three broad classes of problem [Fox:88b, Denning:90, Angus:90a]. These were deduced from our experience at Caltech combined with a literature survey which was reasonably complete up to the middle of 1989. At Caltech, we developed some fifty applications on parallel machines of which twenty-five led to publications in the scientific literature describing the results of simulations performed on our parallel computers [Fox:88a, Fox:89n, Fox:87d, Fox:88oo]. Our work was mainly on the hypercube, but the total of three hundred references cover work on the Butterfly, transputers and the SIMD Connection Machine and DAP. We were interested in applications and algorithms where we could evaluate the scaling to very large parallel machines. Table 1B illustrates what we mean by an application -- "modelling the acoustic signature of a submarine using direct simulation of turbulence" would be another example and in Table 1A we divide eighty-four application areas into eight disciplines.

Our analysis of problem architecture or structure is based on a break up of each problem into spatial(data) and temporal (control) aspects. These key ideas are developed in [Fox:88b] and [Fox:88tt]. We introduce three broad classes of problem architectures which technically describe the temporal (time or synchronization) structure of the problem [Fox:88b]. The three general temporal

structures are called synchronous, loosely synchronous, and asynchronous; we sometimes shorten these to Classes I, II, and III, respectively. The temporal structure of a problem is analogous to the hardware classification into SIMD and MIMD. Further detail is contained in the spatial structure or computational graph describing the problem at a given instant of simulation time [Fox:88tt]. The spatial structure of a problem is analogous to the interconnect or topology of the hardware. The detailed spatial structure is important in determining the performance of an implementation [Fox:88a] but it does not affect the broad issues discussed here. In Table 1C, we only single out one special spatial structure, "embarrassingly parallel," where there is little or no connection between the individual parallel program components. For embarrassingly parallel problems, the synchronization (both software and hardware) issues are greatly simplified. As shown in Table 1C, asynchronous problems do not clearly scale to massively parallel systems unless they are embarrassingly parallel.

We originally introduced this classification to quantify an analysis of which problems ran best on which parallel machines. In this paper, we will extend these ideas and apply them to possible approaches to high level software paradigms.

Synchronous problems are data parallel in the language of Hillis [Hillis:87a] with the restriction that the time dependence of each data point is governed by the same algorithm. Both algorithmically and in the natural SIMD implementation, the problem is synchronized microscopically at each computer clock cycle. Such problems are particularly common in academia as they naturally arise in any description of some world in terms of identical fundamental units. This is illustrated by quantum chromodynamics (QCD) simulations of the fundamental elementary particles which involve a set of gluon and quark fields on a regular four dimensional lattice. These computations form the largest use of supercomputer time in academia [Baillie:89e, 90f].

Loosely synchronous problems are also typically data parallel but now we allow different data points to be evolved with distinct algorithms. Such problems appear whenever one describes the world macroscopically in terms of the interactions between irregular inhomogeneous objects evolved in a time

synchronized fashion. Typical examples are computer or biological circuit simulations where different components or neurons are linked irregularly and modelled differently. Time driven simulations and iterative procedures are not synchronized at each microscopic computer clock cycle but rather only macroscopically "every now and then" at the end of an iteration or a simulation time step.

Loosely synchronous problems are spatially irregular but temporally regular. The final *asynchronous* class is irregular in space and time. A good example is an event driven simulation which can be used to describe the irregular circuits we discussed above, but now the event paradigm replaces the regular time stepped simulation. Other examples include computer chess [Felten:88i] and transaction analysis. Asynchronous problems are hard to parallelize and some may not run well on massively parallel machines. They require sophisticated software and hardware support to properly synchronize the nodes of the parallel machine as is illustrated by time warp mechanism [Wieland:89a].

Synchronous or loosely synchronous problems parallelize on systems with many nodes. The algorithm naturally synchronizes the parallel components of the problem without any of the complex software or hardware synchronization mentioned above for event driven simulations. As shown in Table 1C, 90% of the surveyed applications fell into the classes which parallelize well. This also includes the embarrassingly parallel I, II, III-EP classes. It is interesting that massively parallel distributed memory MIMD machines which have an asynchronous hardware architecture are perhaps most relevant for loosely synchronous scientific problems.

In Table 2, we give details behind some of the applications in Table 1C by listing a few of the recent (end of 1989) Caltech applications with their problem architectures and an estimate of the appropriateness of SIMD or MIMD hardware. The software portability issue is illustrated by noting that all these codes were originally developed for the MIMD Hypercube, but few have been re-implemented for the SIMD machines with their currently distinct software model -- even though, as shown in Table 2, some 50% of these applications would naturally use a SIMD architecture. This is one of our motivations to

develop software systems designed for particular problem architectures and not for particular machines. Thus, we are developing extensions to Fortran to obtain an approach, termed FortranD, for synchronous problems which can be mapped to both SIMD and MIMD machines.

We have looked at many more applications since the detailed survey in [Fox:88b] and the general picture described above remains valid! We will emphasize later that many complicated problems are mixtures of the basic classifications. An important case is illustrated by a battle management simulation implemented by my collaborators at JPL [Meier:89a]. This is formally asynchronous with temporally and spatially irregular interconnections between various modules, such as sensors for control platforms and input/output tasks. However, each module uses a loosely synchronous algorithm such as the multi-target Kalman filter [Gottschalk:90b] or the target-weapon pairing system. Thus, we had a few (~ 10-50) large grain asynchronous (Class III) objects, each of which was a data parallel Class I or II algorithm. This type of asynchronous problem can be implemented in a scaling fashion on massively parallel machines. We will denote this IIICG-IIFG to indicate the Coarse Grain asynchronous controlling of Fine Grain loosely synchronous subproblems. A similar example of this problem class is machine vision and signal processing, where one finds an asynchronous collection of data parallel modules to perform various image processing tasks, such as stereo matching and edge detection. A somewhat different example is a project of Dennis from Rice in the NSF center CRPC [J. Dennis, Private Communication, 1991] to study optimal well placement in an oil reservoir. Here, the reservoir simulation for a given well placement is loosely synchronous, whereas the overall optimization is a naturally asynchronous Class III algorithm. Figure 3 illustrates another example where we outline an approach to designing a new airframe which involves aerodynamics, structures, radar signature and the optimization discussed above for the oil reservoir case. This Figure 3 also points the interesting analogy between heterogeneous problems, class III CG-IIFG and a heterogeneous computer network.

In the above cases, the asynchronous components of the problems were large grain modules with modest parallelism. This can be contrasted with Otto and Felten's MIMD computer chess algorithm, where the asynchronous evaluation of the pruned tree is "massively parallel" [Felten:88i]. Here, one can

break the problem up into many loosely coupled but asynchronous parallel components which give excellent and scalable parallel performance. Each asynchronous task is now a Class I or II modestly parallel evaluation of a given chess position.

III: Current Software Scenario and Lessons

The dominant software environment on SIMD machines has been Fortran 90, which on the CM-2 has replaced *LISP and C* as the primary language for scientific codes because of the quality of the compiler and the familiarity of scientific users with Fortran [TMC:89a].

On MIMD machines, the major environment has been Fortran or C plus explicit message passing. This has been adequate for synchronous and loosely synchronous problems, which dominated both our work at Caltech, and the applications surveyed in Tables 1 and 2. OCCAM has been used extensively on transputer systems but this has not gained general acceptance, and it is also a system with explicit message passing. The success of the early applications on parallel machines is exciting -- it certainly shows that "parallel computing works." But what do we understand by this? It means that nearly all large scale problems parallelize, but not that we have the best software methodology. Further, most of the current implementations are small academic or research codes; for instance, the fifty Caltech codes were nearly all between 1000 and 10,000 lines long. Longer industrial codes will require better software approaches. These need to address several issues. Explicit message passing, which we have used up to now, is formally portable among MIMD machines, as you can parameterize the number of nodes so that a given program will run on any size machine. However, this is deceptive as performance optimization does make the message passing approach machine dependent. One must consider issues such as the overlap of communication and calculation, decomposition choices, and message location trade-offs for latency and bandwidth. These introduce machine dependence, especially for the irregular Class II problems. To be truly portable, the user must implement an arbitrary decomposition, and this is impractical. The problem is clear; Fortran (C) plus message passing or OCCAM are software models built around the machine and not the problem architecture.

We should stress that what we did was "correct"; namely, we used the available software that allowed us to quickly explore the initial rounds of parallel machines. We must use these lessons to design better software that can address the key large industrial applications. These applications have buried in them essentially the same algorithms that we have shown to work with current software environments.

One lasting lesson is that parallelization of an application requires an understanding of the problem architecture. We can see how this is manifested in three different approaches to parallel software. In the Fortran 90 approach used on SIMD machines, such as the CM-2 and Maspar, the data parallel objects are manipulated explicitly, and parallelization is technically hard but well defined. On MIMD machines, we have currently required the user to know the structure of their problems and use this to parallelize "by hand." A traditional parallelizing computer uses a dependency graph to extract the problem architecture and so parallelize "sequential languages" such as Fortran 77. We believe that users are good at understanding their own problem structure, and this explains the success of Fortran 90 and Fortran plus message passing. Compilers cannot reliably extract problem structure from existing programs without user help. Thus, automatic compiler parallelization of Fortran 77 for distributed memory machines has not been very successful so far. As described in the next section, Kennedy's group at Rice is implementing an extended Fortran 77D, which has additional user decomposition commands which essentially specify the problem architecture for the compiler [Fox:91e].

We will use these lessons and the problem architecture discussion of Section II to analyze what software systems could be appropriate for our different problem architectures.

Class I Synchronous Problems

These problems are tightly coupled synchronous problems which are regular in space and time. Their data or geometric parallelism can be naturally expressed in Fortran 90D (appropriately extended Fortran 90) or similar languages such as CM Fortran [TMC:89a], Crystal [Chen:88b], C* [Quinn:90a, 90b], or even APL. This allows the user to specify the problem structure in a natural

high level fashion using the vector and matrix constructs of Fortran 90. The compiler can take care of mapping this onto different machines including those of SIMD and MIMD architecture [Wu:91a, Fox:91b].

Class III Asynchronous Problems

This class is irregular in space and time and often exhibits functional or process parallelism. Considering the battle management problem discussed in Section II, there is a natural class of parallel components formed by the different sensors and control platforms, and these objects communicate with messages even in the real world! Thus, in this architecture, we see a natural break-up into processes and message passing at the problem level, and software engineering approaches, such as object oriented programming, ADA, C++, Strand, [Foster:90] PCN [Chandy:90a], ISIS [Birman:91a] or Linda [Gelernter:89a] are possibilities. In many cases we do not need to use carefully optimized decompositions but rather, use statistical load balancing and decomposition methods. This problem class includes distributed computing and the software such as ISIS designed to support it. As well as this loosely coupled category, we also see the event driven simulations with their specialized software, which we discussed in Section II. These can be tightly coupled but the effectiveness of large scale parallelism is unclear.

We use "loosely-coupled" or "tightly-coupled" to mean a low or high volume of message traffic between the individual components; this can be measured quantitatively as a ratio of communication to the computational complexity. One could use this to quantify my rather vague statements above and we hope to follow with a more detailed analysis in a later paper.

Class II Loosely Synchronous Problems

These problems are irregular in space but regular in time. Often their spatial structure changes dynamically, and adaptive algorithms are needed. This class is hard because the tightly coupled spatial structure demands the same kind of detailed optimizations provided by the Fortran 90D or Fortran 77D compiler for Class I. However, the irregularities make this hard to implement.

We know that Fortran (or C) plus message passing works for this problem class, but we need a more portable user friendly approach. This could involve new data structures to extend languages like Fortran 90. It needs sophisticated run time support, such as that provided by the PARTI system from ICASE [Saltz:90a]. In particular, we need dynamic load balancing modules for which the basic research has been done, but no general implementations are yet available [Fox:88mm]. We will expand this brief discussion of loosely synchronous problems in Section V.

IV: A Strategy for Portable Parallel Programming

The field of parallel computing is advancing so rapidly that there is no time to develop a major new software environment. Further, there seem to be no compelling new ideas that would warrant this. Thus, we expect that the realistic strategy is to build on existing sequential languages - ADA, C, Fortran. We can expect that reasonable node compilers will exist for these languages and all the necessary support software, such as (node) libraries. Extending well known languages will also aid in the migration of existing codes and make good use of users' experience.

We see two important ways in which we can extend existing languages. The first is illustrated by Chandy's PCN, C++ and Birman's ISIS, which essentially allow one to build and manipulate tasks written in languages such as sequential ADA, C, or Fortran 77. This is a reasonable software environment for Class III problems. The second way of extending C and Fortran 77 adds the parallel data structures to obtain C* and Fortran 90 appropriate for Class I and possible Class II problems. We emphasize that these are complementary and not competing approaches. Indeed, we have shown in Section II the importance of the mixed Class IIICG-IIIFG, which is naturally supported by, for example PCN, where each module could be a data parallel Fortran 90D or Fortran 77D code. Future software system development should be coordinated so that such mixed systems are possible by integrating the development of groups concentrating on these different extensions.

This will give us a hybrid software system, which appears suitable for most problems. There is an overall software environment oriented towards

asynchronous applications with full functionality for creating and controlling objects communicating with a sophisticated message passing environment. Often at this level, we will only see modest parallelism. Each of the "asynchronous objects" is potentially massively data parallel synchronous or loosely synchronous module supported by languages optimized for these classes. We will expand on the latter in Sections V and VI.

We can also note that languages like ADA can support Class III problems but probably need extension for data parallelism. This could be obtained either by adding high level data structures to ADA or by allowing mixed languages with data parallel (say, C*) modules integrated into ADA. Clearly, this is a sophisticated software environment which has to map mixed problem architectures onto heterogeneous distributed computer systems with networks of SIMD and MIMD parallel machines. I believe that we know "in principle" how to tackle these issues but we have a lot of technology development for the separate components of the system before we can hope to implement the full sophisticated mixed environment.

V: Fortran D as a Parallel Software Environment

The success of CM Fortran as the programming environment for the CM-2 suggests that it is a good approach for our synchronous Class I applications. As discussed earlier, we view Fortran 90D (CM Fortran, C*) as programming systems for "SIMD" (synchronous) problems and not as languages for SIMD machines. Compilers can map Fortran 90D effectively into all parallel architectures suitable for this problem class including MIMD, SIMD parallel machines, systolic arrays and heterogeneous networks. Fortran 90 was not originally designed as a massively parallel programming system but it has one key attribute that makes it effective. It uses high level data structures explicitly (as vectors and matrices) and so the problem architecture is clear and not hidden in values of pointers and DO loop indices. It is portable, as high level constructs such as $A=B*C$ with A, B, and C matrices, can be optimized by the compiler for each new machine. Our experience has been that in many cases, users prefer Fortran 90 to Fortran 77, even for sequential applications, as it expresses applications naturally with much shorter code. Often one finds a factor of 2 to 3 reduction for Fortran 90 compared to Fortran 77.

However, Fortran 90 needs to be extended in significant ways for parallelism. We have designed one modest set of extensions to handle decomposition and "embarrassingly parallel" *forall* statements. As described in the next section, we are investigating further enhancements, especially in the area of new data structures. We term the resultant system Fortran 90D [Fox:91c].

Table 3 displays one example that helped us evaluate Fortran 90 as a portable environment [Keppenne:89a, 90a]. We isolated a 1500 line computational kernel from climate code using spectral methods. Extensive use of pointers made this code perform poorly on vector machines, such as the Cray YMP and made it essentially impossible for either a compiler or an outside person to improve or parallelize code. However, the code was rewritten by the original developer in Fortran 90, reducing the code size to 600 lines. This new code had an order of magnitude better performance on the Cray YMP while an outside "computer scientist" was able to convert it into Fortran 77 and Fortran 77 plus message passing without difficulty. We believe that this last step can be performed by a compiler using lessons from this and other manual conversions. In this sense, the new Fortran 90 code is portable and scalable to new machines.

We are developing with Kennedy's group at Rice and the Parasoft Corporation an integrated Fortran environment, illustrated in Figure 4, which emphasizes that the language extensions supply equally well to Fortran 77 and Fortran 90 [Fox:91c]. In this picture, we can view Fortran 90D as a "permanent annotation language" for user assisted parallelization of Fortran 77. It will require more experimentation with real application codes to compare the relative merits of parallelizing Fortran 77D versus Fortran 90D.

VI: Loosely Synchronous Extensions of Fortran 90D

In Table 4, we illustrate how increasingly complex problem architectures require extensions to a Fortran 90D environment. We see a progression of extensions to Fortran 90 including:

- a) Decomposition directives

- b) *forall* commands to control aspects of problems involving asynchronous but uncoupled calculations
- c) Run-time support for decomposition of irregular scientific computations such as those found in molecular dynamics and unstructured finite element calculations [Johnsson:89b, 90a]. This area has been pioneered by Saltz with the PARTI system [Saltz:90a, Saltz:87a].
- d) The above extensions of Fortran 90 handle problems in which the data structure is an array — including arrays of pointers needed in c). However, there are important cases where more general data structures are needed to naturally capture the architecture of the problem. This area has received little attention in the computer science community. I see it as a critical motivation for new parallel computing environments and languages. Thus, Fortran 90 handles simple array data structures quite well; one may prefer comparable array extensions of C (i.e., C*), ADA, or functional languages such as Crystal [Chen:88b]. I believe our study of Fortran 90D will naturally extend to comparable parallel versions of other languages. However, the key uncertainties are in the support of the difficult Class II and III problems. One data structure of importance is that of a tree which occurs in any recursive algorithm, such as sorting [Fox:88a] or most importantly, in scientific simulations using one of the various multiscale approaches. These are of growing interest in vision, partial differential equations, and particle dynamics.

We are collecting as many examples as we can of these "skeletons in the parallel language closet" which can help motivate extensions to Fortran 90D and other languages.

One particularly good example is the Barnes-Hut [Barnes :86a] clustering algorithm [Appel:85a], which Salmon and Warren have implemented on the hypercube [Fox:89n, Fox:89t, Salmon:90a]. This is not only an interesting algorithm, but has been used for large scale astrophysics simulations on the NCUBE-1, NCUBE-2 and JPL Mark IIIfp hypercubes. These simulations have

studied the collision of galaxies and the evolution of fluctuations in the early universe [Salmon:89b] [Warren:91a]. Consider the evolution of a collection of N stars where the long range force between stars gives a complexity of $O(N^2)$ for the direct calculation. As illustrated in Figure 5, this can be reduced by noting that for widely separated systems, one can approximate the effect of the M stars by their centroid, or generally, their multipole expansion [Greengard:88a]. Applied recursively, this approach reduces the complexity to $O(N)$ or $O(N \log N)$, depending on the details of the implementation [Zhao:87a, 89a]. This gives the tree like data structure exemplified in Figure 6, where we form (in two dimensions) successive quad-trees until there is, at most, one star in each final "leaf" of the tree. As shown in Figure 7, this method parallelizes well with efficiencies of 80% on large realistic three dimensional problems on the 512 node NCUBE-1. However, this careful user decomposition and parallelism is not easy to capture in current languages. The "natural" (in C) data structure is a linked list to represent the dynamic tree. Presumably, no static compiler analysis can decode the pointer values to uncover this data structure. Subtleties used by Salmon -- including replication of the top of the tree among all nodes to avoid a hot spot there -- are hard to automate with current approaches. We have given a more detailed analysis of this problem and also a multiscale statistical physics example [Baillie:91a] in [Fox:90p]. Features that could have general significance include

- The concept of *locally essential data* which is "all the information" needed to evolve forward the stars inside the subdomain contained in a given node. This concept implies a replication of the top of the tree in all nodes.
- A natural hypercube (recursive) algorithm to build the parallel tree which is related to other global problems such as parallel quicksort [Fox:88a].
- A loosely synchronous implementation with a communication phase followed by a calculation phase. In spite of the complex global method needed to build the tree, the total communication overhead is, at worst, a few percent.

We are currently investigating this and other difficult examples, such as high level image analysis and other multiscale algorithms, to see if they can be

supported by additional data structures (e.g., a tree) and a new run-time library to manipulate these structures and relate them to existing Fortran (C*) constructs. As shown in Table 4, we expect to enhance Fortran D with these new concepts and extend this language to a possible portable approach to all synchronous and loosely synchronous problems.

VII: Conclusions

We have surveyed many applications and shown how study of problem architectures allows one to clarify which software approaches are appropriate for which problems and will give scalable portable code.

Our current research is concentrating on implementing and understanding these lessons as described in Sections V and VI. We are also hoping to broaden our application survey and, in particular, study the large codes seen in government and industrial problems.

Acknowledgements

I would like to thank Ken Kennedy, Adam Kolawa, Joel Saltz, John Salmon and Min-You Wu for helping me understand these issues.

References

- [Angus:90a] Angus, I. G. Fox, G. C., Kim, J. S., and Walker, D. W. *Solving Problems on Concurrent Processors: Software for Concurrent Processors*, volume 2. Prentice-Hall, Inc., Englewood Cliffs, NJ 07632, 1989.
- [Appel:85a] Appel, A. W. "An efficient program for many-body simulation," *Sci. Stat. Comput.*, 6:85, 1985.
- [Baillie:91a] Baillie, C. F., and Coddington, P. D. "Cluster identification algorithms for spin models," *Concurrency: Practice and Experience*, 3(2):129-144, April 1991. Caltech Report C3P-855.
- [Barnes:86a] Barnes, J., and Hut, P. "A hierarchical $O(N \log N)$ force calculation algorithm," *Nature*, 324:446, 1986.
- [Birman:91a] Birman, K. "Recent Developments with ISIS," Talk given at DARPA Workshop, Providence, Rhode Island, February 28, 1991.

- [Chandy:90a] Chandy, K., and Taylor, S. "A primer for program composition notation." Technical Report CRPC-TR90056, California Institute of Technology, June 1990.
- [Chen:88b] Chen, M., Li, J., and Choo, Y. "Compiling parallel programs by optimizing performance," *Journal of Supercomputing*, 2:171-207, 1988.
- [Denning:90] Denning P. J. and Tichy W. F. "Highly Parallel Computation," *Science* 250, 1217-1222 (1990).
- [Felten:88i] Felten, E. W., and Otto, S. W. "A highly parallel chess program," in *Proceedings of International Conference on Fifth Generation Computer Systems 1988*, pages 1001-1009. ICOT, November 1988. Tokyo, Japan, November 28-December 2. Caltech Report C3P-579c.
- [Foster:90] Foster, I., Taylor, S. "Strand™: New Concepts in Parallel Programming," Prentice-Hall, Englewood Cliffs, NJ 07632, 1990.
- [Fox:87d] Fox, G. C. "Questions and unexpected answers in concurrent computation," in J. J. Dongarra, editor, *Experimental Parallel Computing Architectures*, pages 97-121. Elsevier Science Publishers B. V., North-Holland, 1987. Caltech Report C3P-288.
- [Fox:88a] Fox, G. C., Johnson, M. A., Lyzenga, G. A., Otto, S. W., Salmon, J. K., and Walker, D. W. *Solving Problems on Concurrent Processors*, Volume 1. Prentice-Hall, Inc., Englewood Cliffs, NJ 07632, 1988.
- [Fox:88b] Fox, G. C. "What have we learned from using real parallel machines to solve real problems?," in G. C. Fox, editor, *The Third Conference on Hypercube Concurrent Computers and Applications*, Volume 2, pages 897-955. ACM Press, 11 West 42nd Street, New York, NY 10036, January 1988. Caltech Report C3P-522.
- [Fox:88mm] Fox, G. C. "A review of automatic load balancing and decomposition methods for the hypercube," in M. Schultz, editor, *Numerical Algorithms for Modern Parallel Computer Architectures*, pages 63-76. Springer-Verlag, 1988. Caltech Report C3P-385.
- [Fox:88oo] Fox, G. C. "The hypercube and the Caltech Concurrent Computation Program: A microcosm of parallel computing," in B. J. Alder, editor, *Special Purpose Computers*, pages 1-40. Academic Press, Inc., 1988. Caltech Report C3P-422.
- [Fox:88tt] Fox, G. C., and Furmanski, W. "The physical structure of concurrent problems and concurrent computers," *Phil. Trans. R. Soc. Lond. A*, 326:411-444, 1988. Caltech Report C3P-493.

- [Fox:89n] Fox, G. C. "Parallel computing comes of age: Supercomputer level parallel computations at Caltech," *Concurrency: Practice and Experience*, 1(1):63-103, September 1989. Caltech Report C3P-795.
- [Fox:89t] Fox, G. C., Hipes, P., and Salmon, J. "Practical parallel supercomputing: Examples from chemistry and physics," in *Proceedings of Supercomputing '89*, pages 58-70. ACM Press, November 1989. IEEE Computer Society and ACM SIGARCH, Reno, Nevada, Caltech Report C3P-818.
- [Fox:90p] Fox, G. D. "Hardware and software architectures for irregular problem architectures," in *ICASE Workshop on Unstructured Scientific Computation on Scalable Microprocessors*, October 1990. Held in Nags Head, North Carolina. SCCS-111.
- [Fox:91b] Fox, G. C. "Achievements and Prospects for Parallel Computing", Invited Talk at *International Conference on Parallel Computing: Achievements, Problems and Prospects*, Anacapri, Italy, June 3-9, 1990. C3P-927.
- [Fox:91c] Fox, G. C. "Parallel problem architectures and their implications for portable parallel software systems." Technical Report C3P-967, Northeast Parallel Architectures Center, May 1991. CRPC-TR91120, SCCS-78, Presentation at DARPA Workshop, Providence, Rhode Island, February 28, 1991.
- [Fox:91e] Fox, G. C., Hiranadani, S., Kennedy, K., Koebel, C., Kremer, U., Tseng, C.-W., Wu, M.-Y., "Fortran D language specification." Technical Report SCCS-42c, Rice Center for Research in Parallel Computation; CRPC-TR90079, April 1991.
- [Fox:91p] Fox, G. C. "Parallel problem architectures and their implications for portable parallel software systems." Technical Report C3P-967, Northeast Parallel Architectures Center, May 1991. CRPC-TR91120, SCCS-78, Presentation at DARPA Workshop, Providence, Rhode Island, February 28, 1991.
- [Gelernter:89a] Gelernter, D. "Multiple tuple spaces in Linda," in *Proceedings of Parallel Architectures and Languages Europe*, volume 2, page 366. Springer-Verlag, LNCS, June 1989.
- [Gottschalk:90b] Gottschalk, T. D. "Concurrent multi-target tracking," Technical Report C3P-908, California Institute of Technology, April 1990. Published in *Proceedings of the Fifth Distributed Memory Computing Conference*, April 9-12, Charleston, South Carolina.

- [Greengard:88a] Greengard, L. "The rapid evaluation of potential fields in particle systems," in *ACM Distinguished Dissertation Series, Vol. IV*. MIT Press, Cambridge, Mass., 1988. Yale research report YALEU/DCS/RR-533, April 1987.
- [Hillis:87a] Hillis, W. D. "The Connection Machine," *Scientific American*, page 108, June 1987.
- [Johnsson:89b] Johnsson, S. L., and Mathur, K. "Data structures and algorithms for the finite element method on a data parallel supercomputer." Technical Report CS89-1, Thinking Machines Corporation, 1989.
- [Johnsson:90a] Johnsson, S. L., and Mathur, K. K. "Experience with the conjugate gradient method for stress analysis on a data parallel supercomputer." Technical report, Thinking Machines Corporation, 1990.
- [Keppenne:89a] Keppenne, C. L. *Bifurcations, Strange Attractors and Low-Frequency Atmospheric Dynamics*. PhD thesis, Universite Catholique de Louvain, 1989.
- [Keppenne:90a] L., K. G., Ghil, M., Fox, G. C., Flower, J. W., Kolawa, A., Papaccio, P. N., Rosati, J. J., Shepanski, J. F., Spadaro, F. G., and Dickey, J. O. "Parallel processing applied to climate modeling." Technical Report SCCS-22, Syracuse University, November 1990.
- [Meier:89a] Meier, D. L., Cloud, K. C., Horvath, J. C., Allan, L. D., Hammond, W. H., and Maxfield, H. A. "A general framework for complex time-driven simulations on hypercubes," Technical Report C3P-761, California Institute of Technology, March 1989. Published in the *Proceedings of the Fourth Conference on Hypercubes, Concurrent Computers and Applications*.
- [Quinn:90a] Quinn, M. J., and Hatcher, P. J. "Data-parallel programming on multicomputers," *IEEE Software*, pages 69-76, September 1990.
- [Quinn:90b] Quinn, M. J. "Compiling SIMD Programs for MIMD Architectures," *Proceedings of the IEEE Computer Society 1990 International Conference on Computer Languages*, March 1990.
- [Salmon:90a] Salmon, J. *Parallel Hierarchical N-Body Methods*. PhD thesis, California Institute of Technology, December 1990.
- [Saltz:87a] Saltz, J., Mirchandaney, R., Smith, R., Nicol, D., and Crowley, K. "The PARTY parallel runtime system," in *Proceedings of the SIAM Conference on Parallel Processing for Scientific Computing*. Society for Industrial and Applied Mathematics, 1987. held in Los Angeles, CA.

- [Saltz:90a] Saltz, J., Berryman, H., and Wu, J. "Multiprocessor and runtime compilation," *Concurrency: Practice and Experience*, 1991. To be published.
- [TMC:89a] Thinking Machines Corporation, Cambridge, MA. *CM Fortran Reference Manual*, version 5.2-0.6 edition, September 1989.
- [Warren:91a] Warren, M. S., Zurek, W. H., Quinn, P. J., and Salmon, J. K. "The shape of the invisible halo: N-Body simulations on parallel supercomputers." Technical Report C3P-961, California Institute of Technology, 1991. Submitted to Proceedings of After the First Three Minutes, ed. S. Holt, V. Trimble, and C. Bennetti, AIP, 1991; Los Alamos Technical Report LA-UR-90-3915.
- [Wieland:89a] Wieland, F., Hawley, L., Feinberg, A., Diloroto, M., Blume, L., Ruffles, J., Reiher, P., Beckman, B., Hontalas, P., Bellenot, S., and Jefferson, D. "The performance of a distributed combat simulation with the time warp operating system," *Concurrency: Practice and Experience*, 1(1):35-50, 1989. Caltech Report C3P-798.
- [Wu:91a] Wu, Min-You and Fox, Geoffrey C. "Compiling Fortran90 programs for distributed memory MIMD parallel computers." Technical Report SCCS-88, Syracuse Center for Computational Science, April 1991. CRPC-TR91126.
- [Zhao:87a] Zhao, F. *An $O(N)$ Algorithm for Three-dimensional N-body Simulations*. PhD thesis, Massachusetts Institute of Technology, 1987. October.
- [Zhao:89a] Zhao, F., and Johnsson, S. L. "The parallel multipole method on the Connection Machine." Technical Report CS89-6, Thinking Machines Corporation, October 1989.

Table 1: Summary of Problem Architectures

A. Data Sample from 300 Papers [Angus 90a, Fox 88b]

<u>84</u>	<u>Total Applications</u>
9	Biology
4	Chemistry and Chemical Engineering
14	Engineering
10	Geology & Earth Science
13	Physics
5	Astronomy and Astrophysics
11	Computer Science
18	Numerical Algorithms

B. Typical Applications

Calculate Proton Mass	Evolution of the Universe
Seismic Modelling	Optimization of Oil Well
Dynamics of H+HO	Placement
Image Processing	Analyse Voyager Data from
Multiple Target Tracking	Neptune
	Computer Chess

Table 1C. Conclusions of Survey of Applications

About 50% of applications clearly run well on SIMD machines.

About 90% of applications scale to large SIMD/MIMD machines.

Category	Number	Fraction		Natural Support Hardware
I: Synchronous	34	0.4	Total Class I and II	SIMD
II: Loosely Synchronous (not Synchronous)	30	0.36	Spatially Connected 0.76	MIMD Distributed Memory
I: Embarrassingly Parallel	6	0.07		SIMD
II or III: Embarrassingly Parallel but asynchronous and needs MIMD	6	0.07		MIMD Distributed Memory
III: Truly Asynchronous (Spatially connected)	8	0.1	Unclear Scaling	Unclear Maybe MIMD Maybe Shared Memory

Table 2: Problem Architecture of 14 Selected Caltech Parallel Applications
[Fox:89t]

Application	Problem Architecture	Does SIMD Perform Well
QCD	I - Regular	Yes
Continuous Spin (High T_c)	I - Regular	Yes
Ising/Potts Models	I - Regular	Yes
Strings	III - Embarrassingly Parallel (forall)	No
Particle Dynamics $O(N \log N)$ $O(N^2)$	II - Irregular I - Regular	Maybe Yes
Astronomical Data Analysis	IIICG-IIFG	Unknown
Chemical Reactions $H + H_2$ Scattering $e^- + CO$ Scattering	I - Regular + forall I - Regular + forall	Probably Probably
Grain Dynamics	I - Regular	Yes
Plasma Physics	II - Can Be Irregular	Probably
Neural Networks	II - Typically Irregular	Sometimes
Computer Chess	III Asynchronous	No
Multi-target Tracking	II - Irregular	Maybe

Table 3: Performance of a Climate Modelling Computational Kernel

Code	Machine	Performance Mflops	
Original C	CRAY Y-MP (1 head)	1.5	Old Code
Fortran 90 (CM Fortran)	8K CM-2	66 (problem too small)	New Portable code
Fortran 77 Generated from Fortran 90	CRAY Y-MP	20	
Fortran 77 + Message Passing Generated from Fortran 90	NCUBE-1 (16 node) hypercube	3.3	
	NCUBE-2 (16 node) hypercube	20	
	Intel i860 (16 node) hypercube	80	

In each case only minor [obviously needed] optimizations were performed.

Table 4.
Fortran 90D for Synchronous (SIMD) and Loosely Synchronous (MIMD)
Data Parallel Programming (about 90% of Scientific and Engineering Computations)

Program Class	Language and Environment Features
a) I-Regular Geometry e.g., full matrix e.g., finite difference e.g., Monte carlo	"pure" Fortran 90 with arrays of values Need decomposition directives in Fortran D
b) I-Regular+III-EP e.g., chemical potential and dynamics problems Calculate matrix elements (needs <i>forall</i>); full matrix algebra (Class I) for energies and cross sections	Add <i>forall</i> to Fortran 90
c) I/II-Regular Topology but irregular geometry e.g., finite element	Add arrays of pointers to arrays of values. Need new run-time library as in PARTL.
d) "True" Loosely Synchronous (II) Irregular Problems e.g., High-level image processing e.g., Multiscale simulations Problem architectures are more general than that of array.	New data structures in Fortran90D
e) IIICG-1, IIFG Complex System Simulations (See Section II and Figure 3)	Fortran90D modules controlled by object oriented systems.

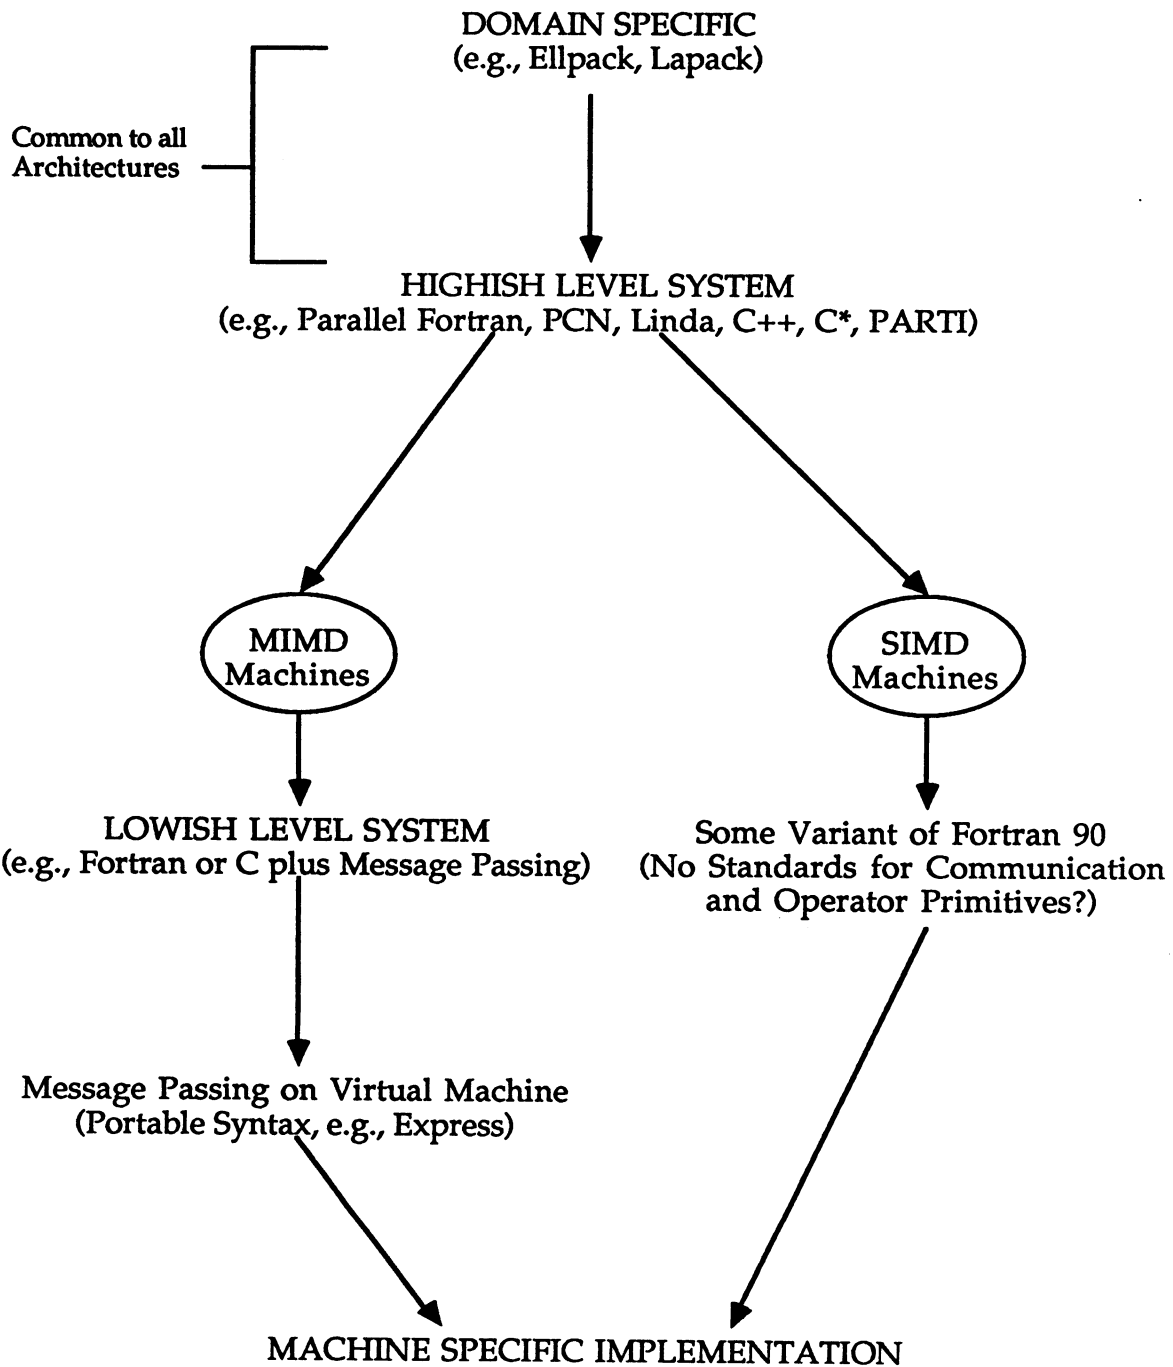


Figure 1. Software Layers for Parallel Machines

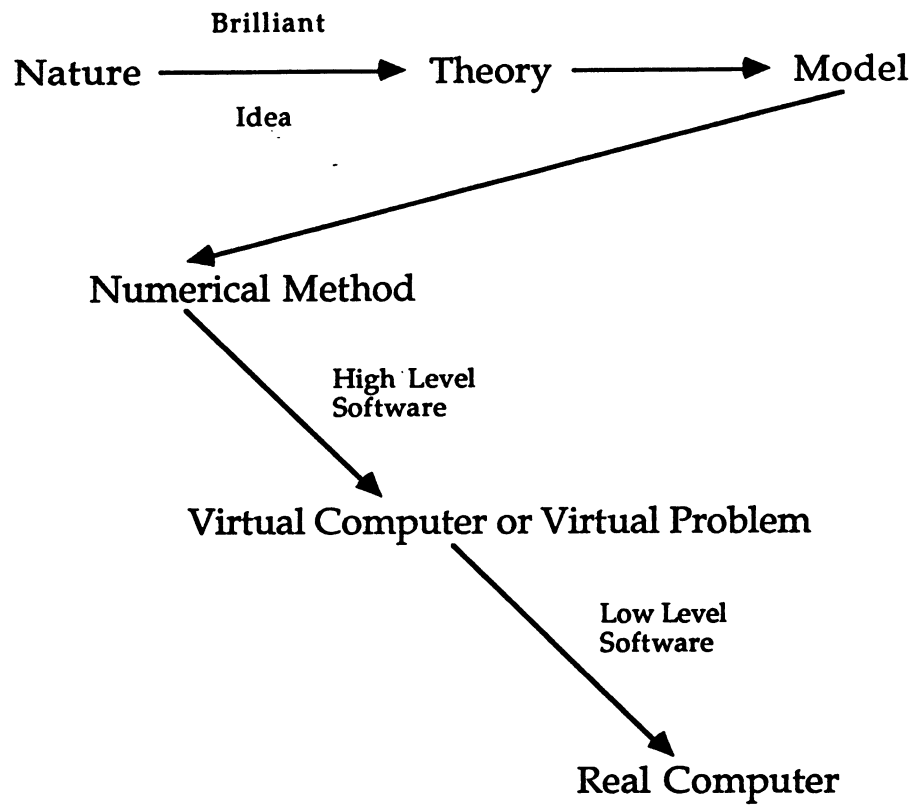


Figure 2. Computation and Simulation as a Series of Maps

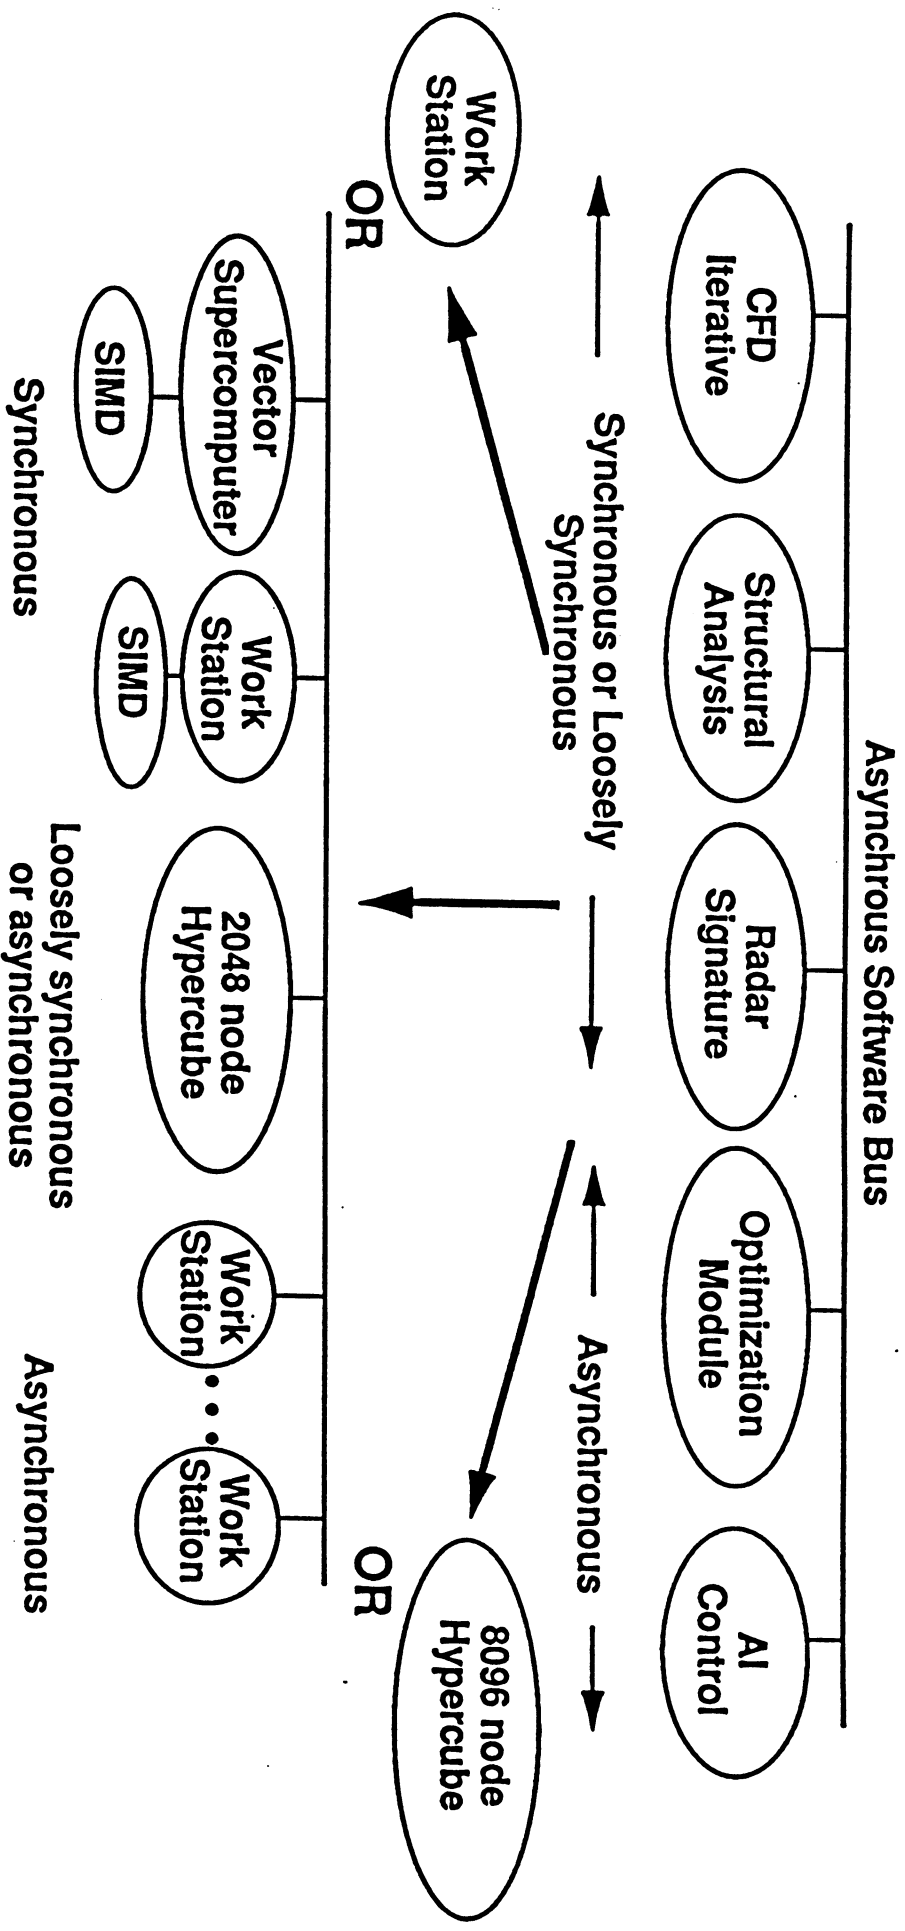


Figure 3: The Mapping of Heterogeneous Problems onto Heterogeneous Computer Systems.

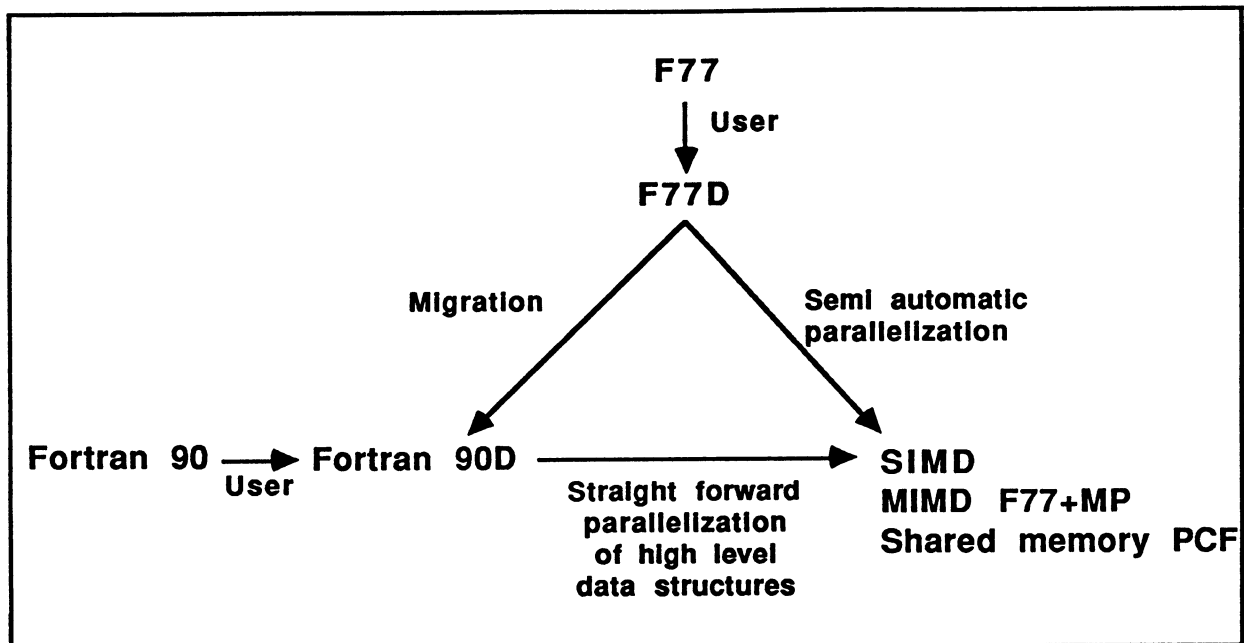


Figure 4: An integrated Fortran environment




Figure 5: A cluster of stars replaced by their centroid




Figure 6: A complex (tree-like) physical structure which must be captured by a data structure in the computation. It is currently represented by a linked list and it is not well expressed in sequential or parallel Fortran. "O" represents a star.




Figure 7: Speedup on the 1-512 node NCUBE-1 hypercube for three different astrophysical particle simulations [Salmon90:a].