

# *A Disruptive Solution to Parallel Processing*

William C. Cave†, Robert E. Wassmer†, Kenneth T. Irvine†, Henry F. Ledgard‡

† Prediction Systems, Inc.; ‡ University of Toledo

December 6, 2017 - [www.VisiSoft.com](http://www.VisiSoft.com)

## GENERIC PROBLEM OVERVIEW

For a nation to be a top competitor on a global economic scale, one of the most important capabilities it must sustain is its computer technology. Computers affect a nation's personal environment today as well as the competitive position of its industries. People now carry more computer power in their pocket than that contained in a huge room of equipment 50 years ago. Instead of using megawatts of power with backup generators, they use small batteries.

Yet application requirements continue to grow, sustaining the need for warehouses containing many racks of computers running in parallel. The energy used by these facilities has grown so great it requires special agreements with large power companies. Managers of these facilities must plan together for runs that would otherwise drain power from local towns. The growth in demand to solve large scale application problems has led to major international competition across the globe to support requirements for High Performance Computing (HPC).

At the same time that needs of HPC (also known as Exascale) have been rising, internal clock rates - the major driver of computer speed since the early 1960s - came to a halt in 2006. Loss of this "automatic" rise in speed forced hardware manufacturers to put multiple processors on a chip. A decade later, they have advanced to many full-up processors (18+) - each with its own co-processor - on a single chip, while putting 2 to 4 chips on a PC board.

To make software run faster, programmers are told to use more processors, as if 72 processors running in parallel can make an application run 72 times faster. But with today's approaches to building software, one finds Speed Multipliers (over a single processor) as low as 2 to 7 when using 100 processors in parallel. This translates to low Processor Utilization Efficiency (PUE), a measure of the speed multiplier divided by number of processors used. Since the 1980s, many studies have shown that current software approaches are slower than older approaches when compared on single processors. Separately, software development productivity has been declining, [1], [2], [3], [7], [11], [14]. The industry has been told by many highly experienced computer designers that there is no way for current software approaches to take advantage of parallel processors, [9], [12]. So why do they persist?

This situation is not restricted to the computer field. It applies to all growing technologies. History has shown that, as technologies grow, they get more complex. The basic problem here is dealing with increasing complexity. This requires breaking the current mold of thinking so that the problem is seen from a totally new perspective.

Once the solution becomes obvious to a few out-of-the-box thinkers, a whole new approach can be taken. And therein lies the real problem - resistance to what is known as a *Disruptive Technology*. Those who have invested large sums of time or money in the old approach do not want to hear that their investments are no longer useful or valuable - until they finally realize they are being left behind, and start to think about the need to change.

## MEASURABLE GOALS

To fairly compare solution approaches, one must observe a wide spectrum of applications that are critical to the future of U.S. technology and are known to be difficult to address effectively. In the case of simulation, application experts must understand the models as implemented in software to ensure that a system is correctly represented. One must then focus on fair measurement of results relative to achieving application needs. Only application experts can judge whether those needs are met. This implies the following measurable goals:

- Provide many orders of magnitude improvement in application run-time speed;
- Provide an order of magnitude reduction in the time and cost to develop software;
- Allow application experts to design, build, and test software directly;
- Allow newcomers to a project to quickly learn and understand complex software.

## CLASSIFYING PROPERTIES OF COMPLEX APPLICATION REQUIREMENTS

The defining property of applications requiring parallel processors is the need for large numbers of processors that are tightly coupled under a single Operating System (OS). These applications typically contain substantial inherent parallelism, i.e., independent elements that share information while operating concurrently. Designing systems to meet these requirements is substantially different from independent tasks on a server. *Parallel processor systems that minimize cost while maximizing speed require a new approach to software design.*

### Representative Applications

Those who have worked many parallel processor applications know that each class of applications appears different. Yet many share common properties. The applications listed below represent an immediate market, and can be used to characterize common properties that directly affect the measurable goals stated above.

1. Real-Time Control of Large Groups of Autonomous Moving Platforms
2. Human Body Organ simulation
3. Global Climate prediction
4. Fluid Flow simulation
5. Biological Particle simulation
6. Chemical - Molecular structure simulation
7. Scanning, sorting, and correlating massive databases (Big Data)
8. Weather prediction in mountainous terrain
9. Power distribution simulation
10. Electro-magnetic wave simulation
11. Global HF power transmission
12. Global Military Planning - Multiple moving platform simulation

To meet the economic requirements of each application, one must minimize run-time cost (number of processors, energy consumption) while meeting their performance constraints. Tailoring a hardware design for each is economically unacceptable. The alternative is to develop the equivalent of von Neumann's Instruction Set Architecture (ISA) to meet all of the applications to be addressed. The parallel processor technology described here provides the equivalent of von Neumann's single processor ISA. Known as the *Application Space Architecture* (ASA), it supports all of the above parallel processor applications.

Many of the above applications have requirements that fall into similar categories. Many require solutions to various types of mathematical formulations. Some applications can use heterogeneous data spaces to save time, using much larger cell sizes where the accuracy of representation of the physical space can suffice with a single cell instead of 100 cells. A number of the more pressing technology problems require extremely complex event-oriented simulations. Some of these are real-time and embedded applications that must be tied to the real-time clock. Most applications require server facilities for some type of I/O - while the main part of the application can run on large numbers of tightly coupled parallel processors. Such applications typically require database inputs where creation, modification, and use of the data is critical to initialization. Most applications require multiple outputs for analysis, often creating large output databases. These requirements are supported using tightly coupled parallel processors connected to a large server environment, see Figure 1.

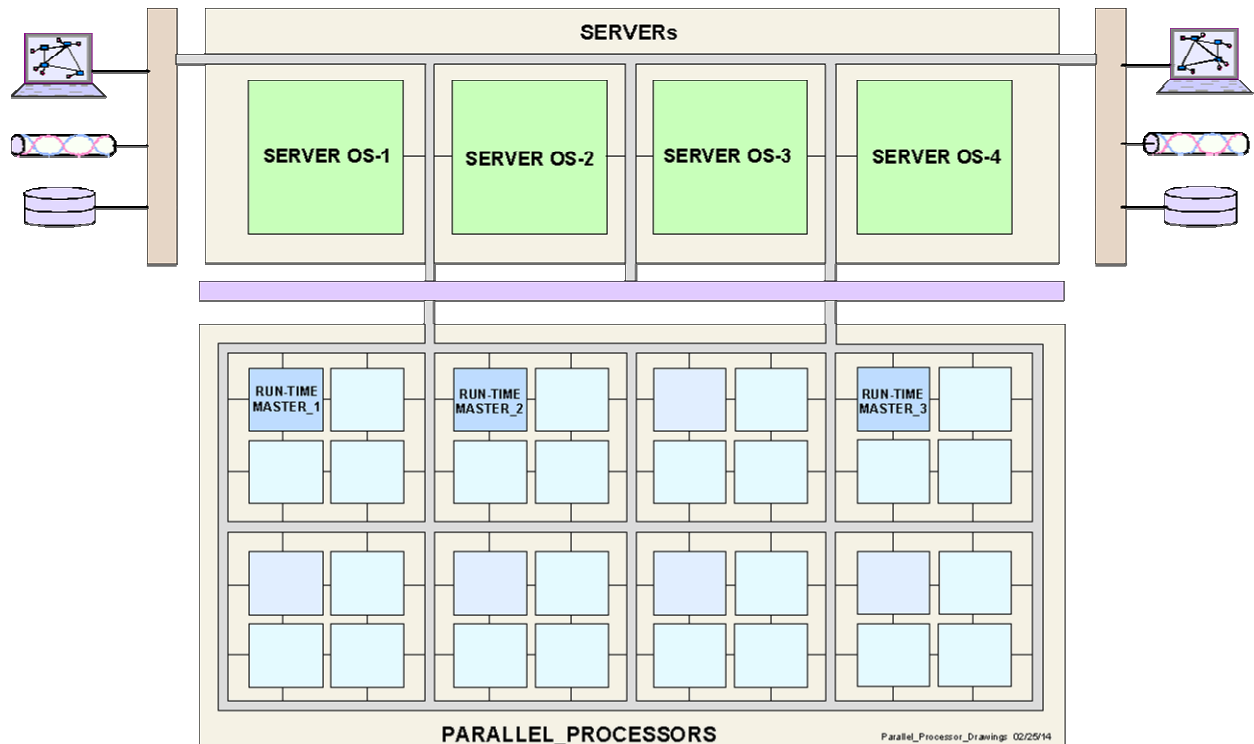


Figure 1. Parallel Processors tied to Servers.

## A SEVEN LAYER MODEL FOR COMPUTER TECHNOLOGY

Much of the discussion in this paper is derived from communication system design. In the mid 1970s, users and developers faced severe difficulties specifying requirements for interconnecting digital communication systems. This problem was most prevalent when trying to gain agreement on the application of various standards. The solution became apparent with agreement among nations on what has become known as the OSI Seven Layer Protocol Model, illustrated in Figure 2.

### Open Systems Interconnection (OSI) - Seven Layer Model for Communication Technology

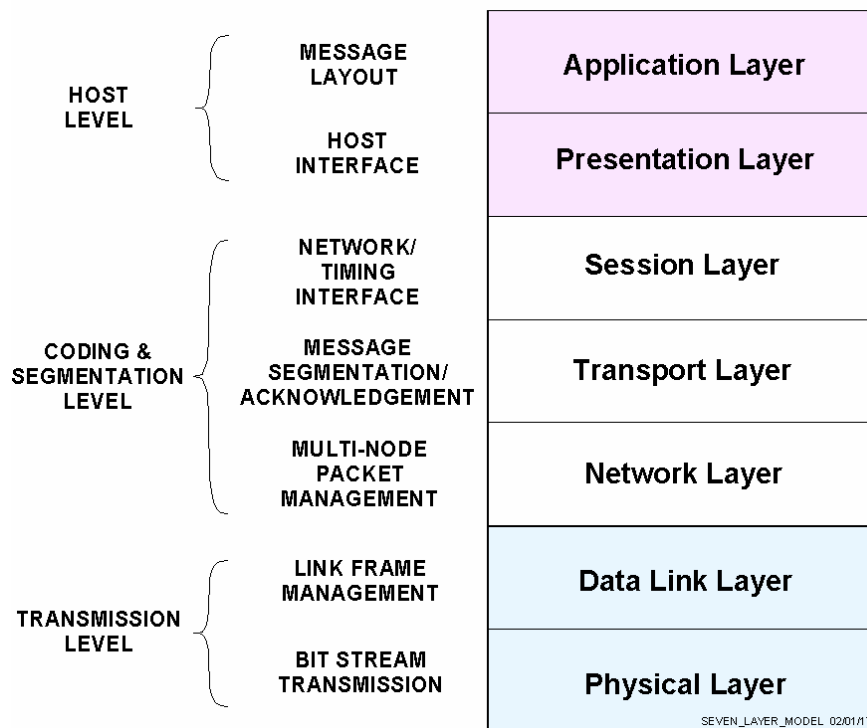


Figure 2. The Seven Layer Model for Communications Technology.

The layers above represent the sequence of requirements, where each layer determines the requirements for the next layer. At first there was resistance from those building the hardware. In the end, definition of the problem became so obvious that it was quickly solved.

The requirement for two modules to communicate directly while running concurrently on parallel processors is critical. Understanding the analogous problem in communications has been simplified by the *Seven Layer Protocol Model*. Depicted in Figure 3 is the computer technology equivalent. It describes the sequential layers of design information that must be produced to create a complete working system, starting with the application requirements and ending with those for hardware. It is aimed specifically at parallel processing applications.

## Seven Layer Model for Computer Technology

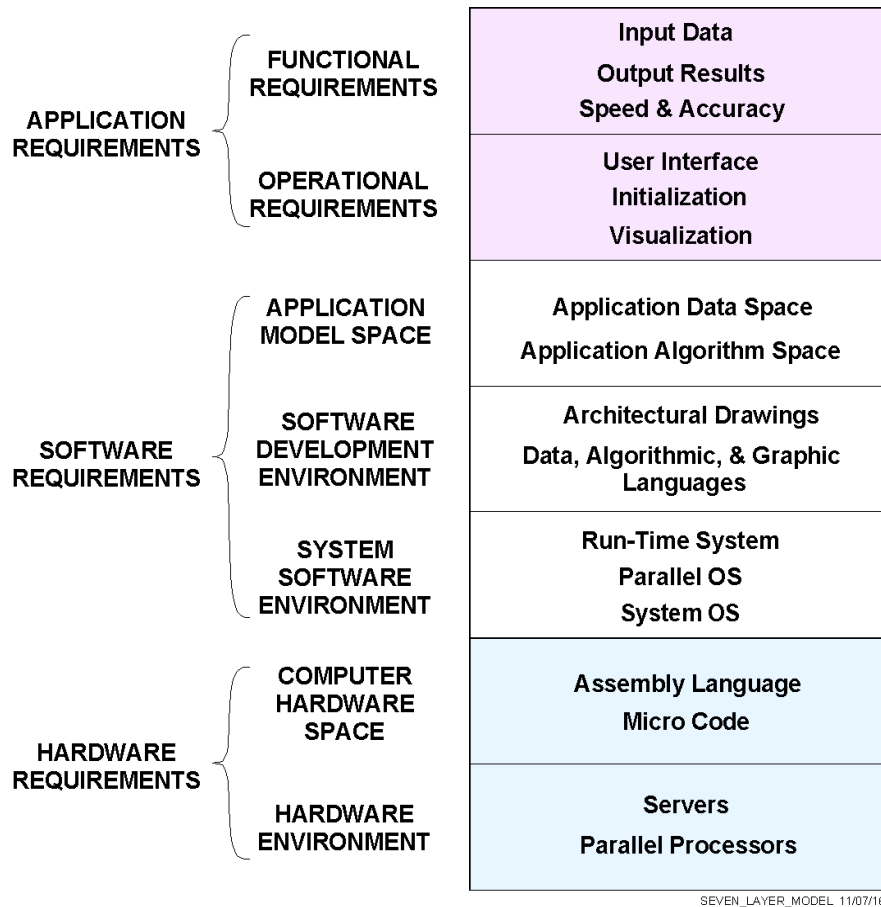


Figure 3. The Seven Layer Model for Computer Technology.

Figure 4 illustrates the interconnection of stacks at three different communication nodes. The ovals connecting these stacks represent data messages passed between stacks. The goal of a communication system is to pass messages - or packets forming a message - between nodes. This must be done reliably, and within fractions of a second. The speed and reliability of message transmission is critical. The coding and segmentation layers between the application requirements and hardware requirements were the critical layers in solving the problems of algorithmic development that dramatically enhanced speed and reliability.

The seven layer model for computer technology applies directly to a large scale simulation or control system running on a parallel processor. Here, the ovals in Figure 4 represent memory shared directly between processors. When processes are running concurrently on multiple processors, one is concerned about minimizing the time to exchange information between processors as well as minimizing unproductive time across all processors. This requires maximizing the concurrent processing time on each processor, measured by Processor Utilization Efficiency (PUE), the underlying speed multiplier.




Figure 4. Multiple Seven Layer Model nodes.

In the case of independent server tasks, time to run may be important. But in that case, memory is not shared directly between tasks. It is left to the Operating System (OS) to assign tasks to processors and move instructions and data to the assigned processor. As illustrated in Figure 1, servers are typically tied to external memory devices, communication channels, and interactive terminals. In the case of a parallel processor where speed is critical, a subset of parallel processors may interface with the server processors, but need not interface directly with I/O devices due to the waiting periods required to service those devices.

## SOFTWARE LANGUAGES

Figure 3 illustrates the importance of software to implement data spaces and algorithms. Experienced computer technologists have pointed to software languages as the cause of lack of speed and other problems. So what is the purpose of a language? Languages have evolved as the principle means for people to communicate. Thousands of languages have come and gone over the course of human evolution. Language is a major factor in determining how people in one society are able to compete with those in another. When comparing languages, the measure used is the ability for people to communicate with speed and reliability. How fast can they get the information across without misunderstanding. This is what Information Theory, defined by Claude Shannon and others, is all about, see [15].

Fast and reliable transfer of information is the goal of information theory. Understanding this theory is critical to the design of information systems that require speed. The underlying concept is that: *Reliability of information transfers is increased by adding redundancy* (i.e., additional data). This may be as simple as sending the same message twice, or by using additional words, such as articles, adjectives, or adverbs.

Redundancy is used when writing and reading computer memory. Bits are added to data transfers for error detection and correction. In wireless communications, it is not unusual to double or triple the size of the original data stream (redundancy) to ensure reliable transfers. English is considered to have a high degree of redundancy compared to most other languages, implying it is more likely that information is transferred reliably - and key to the survival of its users.

The following is quoted from Bjarne Stroustrup, inventor of C++: “English is arguably the largest and most complex language in the world (measured in number of words and idioms), but also one of the most successful,” see [16]. It dominates the world of free trade. Considering the small size of the islands where it originated, its survival is attributed to its reliability. To understand this, consider the military motto, “information is power” - the more information one has to make a decision, the more likely a good outcome. If the information is misunderstood, the outcome is likely to be less than expected. So what are the rules that ensure the reliable transfer of information and understanding when dealing with computer languages?

## **Measurement - The Foundation Of Science**

In his ACM article in 2000, *The Emperor with No Clothes*, [11], Henry Ledgard quotes W. Edwards Deming who stated “If you can’t measure it, you can’t improve it.” The message carried the same point as that made by David Parnas, [13], 10 years earlier: “Without measures from repeatable experiments, software is not a science.” Both Ledgard and Parnas are in the top of the list of those knowledgeable in computer languages.

Studies comparing interactive languages have shown that errors increase as statements move from good English to a more terse form, see Ledgard, [10]. Comparisons of COBOL, FORTRAN and C-based languages will typically derive the following programmer reactions: COBOL is verbose; FORTRAN is fair; C-based languages are terse.

At the top of the list of language designers is Grace Hopper, who wrote the first compiler while working at Univac in 1952. In 1959, after the CODASYL conference started the formal development of COBOL, Hopper’s programming group at Univac spearheaded the language design based upon her own FLOW-MATIC language, see Wikipedia, and Beyer, [4]. Hopper's experience that programs are best written in a language close to English - rather than in machine code or languages close to machine code - was captured in the new language, and COBOL went on to be the most ubiquitous data system language to date. Hopper then developed CMS-2, a language for the U.S. Navy that added math and scientific facilities to COBOL. CMS-2 provided the same hierarchical data and hierarchical instruction syntax that contributes huge productivity gains and applies directly to parallel processing. As shown below, software productivity and parallel processing are intertwined.

The typical misperception is that a verbose language correlates to slower run-times. In fact, these properties are totally unrelated. With some terse languages, the reverse is true. Since source language is translated to machine language, speed depends on both the language design and the translator. Understandability and speed can be improved simultaneously. If we are after reliability, verbose is best. If we are after speed, COBOL is clearly the fastest at handling data. FORTRAN inverts large matrices faster than C-based languages. But it takes much more than a language to solve the parallel processing problem for the applications listed above. It requires the ability to deal with substantial complexity. This becomes obvious with measurements.

## **Using A Single Processor To Implement The Required Decision Processes**

When using single processors to implement desired application outcomes, some form of differential representation (e.g.,  $\Delta T$ ) is typically used to represent evolving actions as a function of time. On a single processor, this must be implemented using a set of sequential instructions.

This implies selecting a  $\Delta T$  that is sufficiently small so that the result of evolving outcomes is represented with sufficient accuracy at the end of each  $\Delta T$  increment. In the case of large complex systems, these decision processes may take substantial amounts of time. If the application time constraints cannot be met with a single processor, then one may use a parallel processor. However, the development of data spaces and corresponding algorithms will likely be substantially different to take advantage of the potential parallel processing speeds.

In the case of nonlinear or nonstationary systems, looping through repetitive sets of calculations and decisions may be required to ensure convergence of solutions to meet predetermined or evolving error criteria. In the case of large complex systems, these repetitive processes may take substantial amounts of time. Again, one may meet the required run-time constraints using parallel processing, and system development will likely be substantially different to achieve the potential parallel processing speeds.

### **Using Parallel Processors To Implement The Required Decision Processes**

The obvious difference when using parallel (as opposed to single) processor systems is the ability to perform instructions concurrently on separate processors. The obvious implication is that the application contains reasonable inherent parallelism. If not, then the instructions must be performed sequentially. Clearly applications that can be sped up using many parallel processors have substantial inherent parallelism.

One must beware of measures that are based on “Embarrassingly Parallel” applications. Embarrassingly parallel applications can be split easily into totally independent parts that may run concurrently with little if any intra-communications, except at the start and end of a run.

For example, in the case of parametric analyses, one may run many simulations using different random number seeds to observe how parameter variations that fall randomly within a distribution affect the outcome of a sequence of events. These types of problems may be run on servers or a cluster of small machines (e.g., PCs) since the individual simulations are virtually independent (they share no data while running). This property underlies the definition of problems known as “Embarrassingly Parallel”. It implies they may be designed to run independently - in serial or parallel - to obtain the same answers. When one is concerned about overall run time, the set of sequential tasks may be run - independently - in parallel.

None of the Representative Applications listed above are embarrassingly parallel. These applications generally require a totally different approach to their software (and hardware) design, one that maximizes Processor Utilization Efficiency (PUE), see [6].

## **DESIGNING COMPLEX SYSTEMS**

Development of complex systems that do not need the speed of parallel processors has been a major problem in the software field. It has been well described in the literature concerned with the development of large systems since the 1980s, see [1], [2], [3], [7], [9], [11], [12], [13], [14]. When comparing development of these systems to that of large architectural structures, e.g., buildings, aircraft, power plants, communication networks, etc., these software application systems should be easier to build. So what is the problem?



## **Engineering Drawings**

If one looks at the current methods for building software, one finds critical faults in the approach. These faults are similar to those observed in the early centuries as engineering fields evolved. For example, architectural engineering did not exist until more recent times. Artists provided renditions of what a building should look like. There were no engineering drawings and no measurements. The art was turned over to builders who worked to figure out how to do the construction, often redoing their work during development. Today, skyscrapers require hundreds of engineering drawings and books of specifications to ensure that construction proceeds in a well defined, orderly and high-quality manner. Masons, carpenters, plumbers, electricians, etc. are expected to follow the drawings and specifications. They do not do the designs as they go. Architects inspect their work as it progresses.

## **Architecture**

Although the word architecture is well defined when used in hardware engineering disciplines, it remains unspecified when used in software literature. Similarly, the word module is well defined in hardware engineering fields. Its use in software appears to imply similarity to hardware but remains undefined. The word module is well understood in computer chip and board design, but clearly misunderstood when used in relation to software. One does not build complex hardware systems without first creating an architecture. The architecture is defined using engineering drawings.

## **Modules**

Complex hardware systems define modules using engineering drawings. Their use is critical when it comes to maintenance. Modules must be easy to replace when they fail. Therefore, they are designed to be independent, implying they can be pulled and replaced easily when they fail. The property of independence provides the ability to change the interior parts of a module without affecting the use of its external connections or use by other modules. Software modules are independent when they share no data and can run concurrently on parallel processors.

## **A NEW APPROACH TO BUILDING COMPLEX SOFTWARE**

The answer to the general problem of expanding technology is to linearize growing complexity. Engineers, scientists, and mathematicians study linear systems in school. They are taught how the property of independence can be used to linearize a complex nonlinear system. This is accomplished by using spaces that fit the problem, so that each coordinate is independent of the others. More generally, breaking a problem into a set of independent parts can effectively linearize an otherwise nonlinear problem, making solutions more understandable.

When working as a team in the development of complex systems, it becomes apparent that complexity is reduced by breaking such systems into a hierarchy of spaces. The selection of spaces determines the level of simplification of the design. This concept is no different from the solution of complex mathematical problems where the selection of the “best” space leads to maximum simplicity of the algorithms. This occurs when the subspaces are maximally independent. This results in simplification of the algorithms and reduction of the time to achieve solutions. Selecting the best space requires deep understanding of the dynamics of a problem.

The applications listed above require designing the best spaces in which to map the inherent parallelism of each into software. As shown in Figure 5, this requires mapping the application requirements into an overall software/hardware space defined by experts using a software environment that aids in the design effort. This environment must support ease of translation of the application requirements into a software space that maximizes understanding as well as speed. This is accomplished in various engineering fields using Computer-Aided Design (CAD) systems. The complex translation into a hardware language is now shifted to the computer - where the burden should be, see [17].




Figure 5. Spaces for translation of application requirements into software and hardware.

### Instruction Set Architecture (ISA)

When Von Neumann defined requirements for the first stored program digital computer in the early 1950s, he based it on a variety of applications, from weather prediction, to solving ballistic equations, and solving the compression equations to produce enough energy to set off a nuclear reaction. He derived the *Instruction Set Architecture* (known as the ISA), the set of instructions needed to write the computer programs to implement a wide range of applications. While he did not design the first stored program digital computer (the MANIAC), he provided the programming interface requirements (the ISA) to those who did the logical design to implement his stored language instructions. Memory size quickly became the key to speed.

### Application Space Architecture (ASA)

The representative applications enumerated above all have different software design requirements. As in most difficult applications, the mathematical spaces selected to minimize the complexity of their solutions are all different. Representing these problems requires hierarchical data spaces and discrete event spaces - as well as continuous and discrete time spaces - all in the same application. To support this translation requires an extension of mathematics that helps one to conceive the design of these spaces and corresponding decision processes. The application software must also invoke synchronization facilities for sharing temporally independent data spaces (independent at specific times).

Only application experts can fully understand these spaces and their independence properties. Such software development facilities must be easy for application experts to use. Given languages that efficiently implement these facilities, one is on the road to determining the requirements for parallel processor hardware design.

Simplifying the software design process for parallel processing requires more than just new programming languages or hardware facilities. It requires the ability to map the inherent parallelism of an application system into software modules that can run concurrently. This is met when application experts can easily visualize the architecture of an application system and the corresponding module designs. The key properties required of application software architecture are the *understandability* and *independence* of modules. These requirements are met in engineering fields where CAD systems use visualization to create architectures that represent higher levels of the design process that are most critical.

Design of software for an HPC application requires decomposition into a hierarchy of independent modules that may run concurrently. Given a sufficient degree of inherent parallelism, one must map that parallelism into a *software architecture* tailored to the specific application. To develop optimal decompositions, one must start with the top-level requirements, working down to the bottom of the algorithms. This process requires application expertise.

What is *software architecture*? It is the same as other engineering architectures. Skyscrapers, airplanes, ships, electronic circuits and chips could not be produced without the visualization of architecture using engineering drawings. Architecture defines many critical system properties. First, it is *time invariant* - it does not represent flow of control. Second, it provides a *visualization of the connectivity (independence)* of modules designed to implement the system. Independence is critical to both construction and maintenance of complex systems. Visual depiction of the independence of modules is critical to their design and their ability to run concurrently, i.e., in parallel.

### **Visual Software Engineering**

To meet parallel processor speed constraints, software architectures must be decomposed into independent modules that can share data while running concurrently. To accomplish this, application experts must be able to easily understand the detailed design - down to the code. These requirements led to *VisiSoft*®, a visual software engineering approach, using a CAD system for designing parallel processor software described in [6], and illustrated in Figure 6.

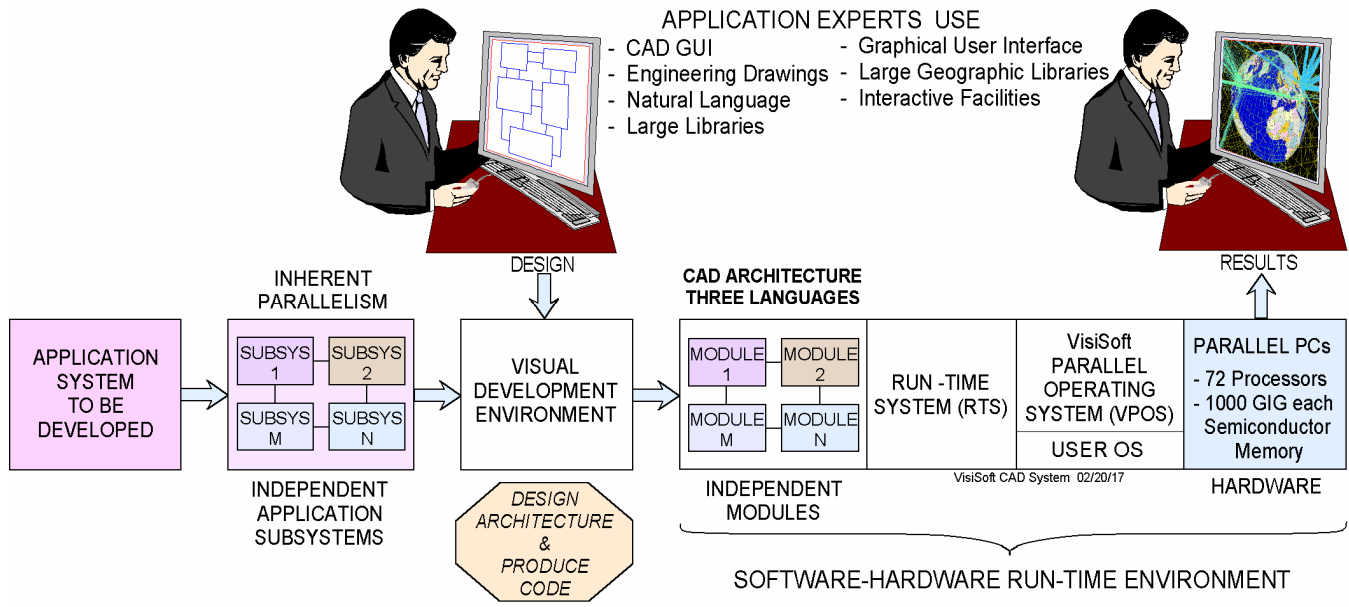


Figure 6. Visual software engineering using a CAD system.

The context of architecture as used here is no different from the hierarchy of drawings required to support the design of skyscrapers. This is unnecessary when building dog houses or small software applications. Similarly, architectural drawings must be accompanied by a set of easily understood specifications. This raises the question: What does it take to produce engineering drawings of software, and languages that are easily understood by application experts?

### VisiSoft LANGUAGES

The VisiSoft languages have been designed to ease the understanding of detailed algorithms required to implement extremely complex systems, see [5]. As understood by Grace Hopper, [4], such organizations are best supported by deep hierarchies of both data and instructions. This is particularly true when dealing with large modules that must be designed to run independently.

### The Separation Principle

The underlying principle supporting the visual characterization of module independence is the ability to determine which modules share what memory resources (data). This is accomplished by separating data from instructions at the language level. Defined in 1982 by Cave while designing the General Simulation System (GSS), this has become known as the *Separation Principle*, [8]. It maps directly into today's hardware designs where data memory is separated from instruction memory. As indicated above, GSS evolved into the VisiSoft CAD system for software, [5]. Using what is defined as the Generalized State Space framework, the Separation Principle is achieved by storing all data in *Resources*. A memory Resource is illustrated in Figure 7.

RESOURCE: TRANSCEIVER		INSTANCES: TRANSMITTER RECEIVER	
<b>GENERAL PARAMETERS</b>			
1	TRANSMITTER_POWER	REAL	INITIAL_VALUE 100
1	RECEIVER_THRESHOLD	REAL	INITIAL_VALUE 120
<b>RADIO</b>			
1	TRANSCEIVER	<b>STATUS</b>	TRANSMITTING RECEIVING IDLE OFF
1	LOCATION		
2	X_POSITION	REAL	
2	Y_POSITION	REAL	
2	ELEVATION	REAL	
1	ANTENNA_HEIGHT	REAL	
1	ANTENNA_GAIN	REAL	
<b>RECEIVER CONNECTIVITY VECTOR</b>			
1	POWER_AT_RECEIVER	REAL	
1	TOTAL_NOISE_POWER	REAL	
1	CONNECTIVITY_MATRIX		
2	PROPAGATION_LOSSES		
3	TERRAIN_LOSS	REAL	
3	FOLIAGE_LOSS	REAL	
3	TOTAL_LOSS	REAL	
2	SIGNAL_POWER	REAL	
2	SIGNAL_TO_NOISE_RATIO	REAL	
2	LINK_DELAY	REAL	
2	LINK	<b>STATUS</b>	GOOD FAIR POOR
<b>TRANSCEIVER RULES</b>			
1	TRANSCEIVER_PROCESS	<b>RULES</b>	GOOD_RECEPTION CONFLICTING_RECEPTION CONFLICTING_BROADCAST

Figure 7. Example of a hierarchically structured state vector (RESOURCE).

Deep hierarchies are required to design complex spaces. They must also allow large complex data structures to be moved in a single instruction fetch, with all of the individual fields directly available to instruction hierarchies in *Processes* as illustrated in Figure 8. This provides ease of understanding as well as orders of magnitude improvement in single processor speeds. Experimental results of speed comparisons are described in [6], Chapter 18.

Note that subscripts are not used in Figures 7 or 8. This is because the resource and process shown are part of an instanced module, where instance pointers (TRANSMITTER & RECEIVER) are automatically handled at the module level. These are set when a process within an instanced module is CALLED or SCHEDULED. Moving instance implementation to the module level substantially enhances understanding of the code.

```

PROCESS: RECEPTION
INSTANCES: TRANSMITTER
            RECEIVER
RESOURCES: TRANSCEIVER
            MESSAGE_FORMATS
            TRANSMITTER_OUTPUT

START_RECEPTION
  IF TRANSCEIVER IS IDLE
    EXECUTE GOOD_RECEPTION
  ELSE IF TRANSCEIVER IS RECEIVING
    EXECUTE CONFLICTING_RECEPTION
  ELSE IF TRANSCEIVER IS TRANSMITTING
    EXECUTE CONFLICTING_BROADCAST .

GOOD_RECEPTION
  IF SIGNAL_TO_NOISE_RATIO IS GREATER_THAN RECEIVER_THRESHOLD
    SET TRANSCEIVER TO RECEIVING
    ADD SIGNAL_POWER TO TOTAL_POWER_AT_RECEIVER .
    CALL DECODE_MESSAGE .

  IF MESSAGE_TYPE IS FORMAT_A
  AND SYNC_CODE IS VALID
  AND LAST_SYMBOL IS A_TERMINATOR
    EXECUTE SEND_ACKNOWLEDGEMENT .

CONFLICTING_RECEPTION
  IF POWER_AT_RECEIVER IS GREATER_THAN SIGNAL_POWER
    SCHEDULE ABORT_RECEIVE NOW .

CONFLICTING_BROADCAST
  CANCEL END_RECEIVE NOW
  SCHEDULE START_RECEIVE IN EXPON(0.83) MILLISECONDS
  WITH PRIORITY 80

SEND_ACKNOWLEDGEMENT
  MOVE ACKNOWLEDGEMENT TO TRANSMIT_MESSAGE_BUFFER
  IF DESTINATION IS BROADCAST
    SEARCH LINK_CONNECTIVITY_VECTOR OVER RECEIVER
    EXECUTING_TRANSMISSION
    WHEN LINK IS GOOD
  ELSE EXECUTE TRANSMISSION .

TRANSMISSION
  SCHEDULE LINK_RECEPTION
  IN LINK_DELAY MICROSECONDS
  USING TRANSMITTER, RECEIVER

```

Figure 8. Example of a hierarchically structured transformation (Process).

More importantly, VisiSoft allows users to assign selective sets of instanced modules to specified processors using the third language, the Control Specification, see [6]. The assignment process is aided by a graphical interface so applications experts can take advantage of their knowledge of the physical properties of the application relative to physical processor assignments. This language provides many facilities, e.g., sections for defining files, libraries, initialization and multiple reruns for simulation and optimization, as well as eliminating scripts.

Most important, these languages support design of software spaces that simplify human translation of inherently parallel physical entities into an organization of independent workloads or modules. Design of the resource and process languages were driven in part by factors somewhat akin to those motivating the use of *tiling* in parallel versions of FORTRAN. However, they minimize memory management overhead of swapping instructions and paging data.

This is accomplished by maximizing the work done on each processor while processes run concurrently with those on the other processors, thus maximizing the PUE. Figures 7 and 8 provide examples of hierarchical resources and processes that help simplify such designs.

The instruction language supports hierarchies of rule structures, where looping and complex IF ... THEN ... ELSE statements are flattened - no nesting. What is known as *Waterfall* or *Fall through* code is gone - without GOTOs. These properties dramatically simplify design of complex algorithms. They lead to substantial increases in both understanding and run time speed - on single as well as parallel processors. We note that there is no global data in this system, and all data is automatically referenced by pointer.

When building complex software, human translation is simplified if a language supports obvious representation of physical behavior. Redundancy in a language, e.g., English, supports ease of understanding, and the likelihood that information will be communicated correctly to another person, see [10] and [15]. The examples in Figures 7 and 8 are taken directly from large detailed simulations of Packet Radio networks. With hierarchical data structures like those shown, one can represent complex algorithms associated with physical systems with ease. Actual systems may entail more complex resources and processes than those shown, but are easily understood by subject area experts.

The Separation Principle also underlies visualization of software architectures using engineering drawings. Resources are depicted as ovals in architectural drawings as illustrated in Figure 9. *Processes* containing instructions that implement transformations are depicted as rectangles. The lines connecting them determine which processes have access to what resources. In this figure, each process has a dedicated resource and shared resources. Transformation 1 has state vector A as input, state vector B for dedicated use, and shares state vector C with transformation 2. Therefore, Transformations 1 and 2 are not *independent*.

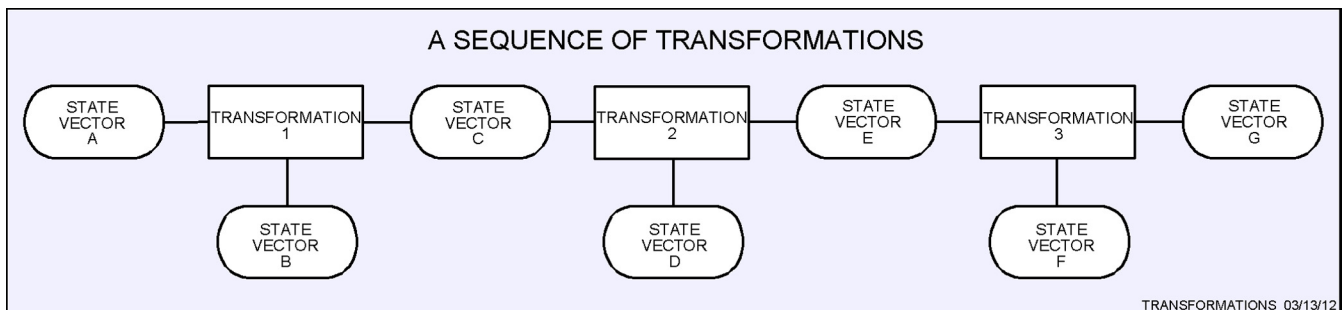


Figure 9. State vectors and transformations.

### Properties Of Independence

As used here, the property of *independence* ensures that processes running on a parallel processor produce *complete and consistent* results for a given set of initial conditions. Consider that state vectors C, D, and E have initial values  $C_i$ ,  $D_i$ , and  $E_i$ . When run on a single processor (sequential machine), Transformation 2 will produce the same outputs:  $C_o$ ,  $D_o$ , and  $E_o$  for a given set of inputs every time it runs; i.e., the results will be *complete and consistent*. If while it is running, one of the resources is changed from the outside, the results may not be complete and consistent. This is because the data being accessed is not *consistent* relative to Transformation 2.

If Transformations 1 and 2 run concurrently, shared state vector C could be changed by either, rendering the data as recognized by the other as *potentially inconsistent*. Therefore, in general, they cannot operate concurrently.

Similarly, Transformation 2 is directly coupled to Transformation 3 by shared state vector E, is not independent of it, and thus cannot run concurrently with it. However, Transformations 1 and 3 can operate concurrently since they share no state vector directly and are therefore *spatially independent*. Transformation 2 can operate only when Transformations 1 and 3 are both idle; in that case they are *temporally independent*.

## SOFTWARE ARCHITECTURE

As illustrated in Figure 10, software architects can decompose a system into modules by grouping resources and processes into an *elementary module*. *Hierarchical modules* are created by grouping modules into higher level modules. Figure 10 shows a library module, PROPAGATION\_PREDICTION, that is sufficiently complex to warrant its own drawing. In general, modules are independent if they share no resources (i.e., they are not connected). Having designed an architecture, developers can implement the data structures and rules using the *resource* and *process* languages. Using this CAD system, resources and processes may be edited directly on the drawing as illustrated in Figure 11. The languages do not permit the declaration of scope rules. It is the architecture that determines how data is shared, and the corresponding independence of modules. Most important, the languages are designed to provide for deep hierarchies in both data structures and rule structures to support important software architecture requirements. These language properties are critical to simplifying the understandability of large complex software systems.

### Parallelism, Architecture, and Decomposition

When striving to take advantage of the inherent parallelism in a system, one must determine the architecture of software modules that maximizes concurrency on a parallel processor. Picking the best set of state spaces is key to solving this problem. Again, best translates to simplicity of transformations and run-time speed.

Having defined *Generalized State Space* as the framework, the mathematical analogy becomes one of selecting the best set of information vectors (Resources) to represent the system attributes. Depending upon how the resources are designed and structured, the rules (Processes) are much easier to understand, build, and modify. This is also determined by the *independence properties of the architecture*, i.e. the interconnection of resources and processes. It is the architectural drawings that determine which processes share what resources, eliminating scope rules in the language. This system also eliminates C++ type pointers and global data types. All of these unnecessary complexities are in conflict with the design of software for parallel processing.

Instead of attempting to describe architectures at the language level, the problem is separated into the natural hierarchy of drawings and language. To follow this concept further, there are specific module types, designated at the drawing level using different colors. These are described below. The simplifications that these visual hierarchies provide become obvious after using them.



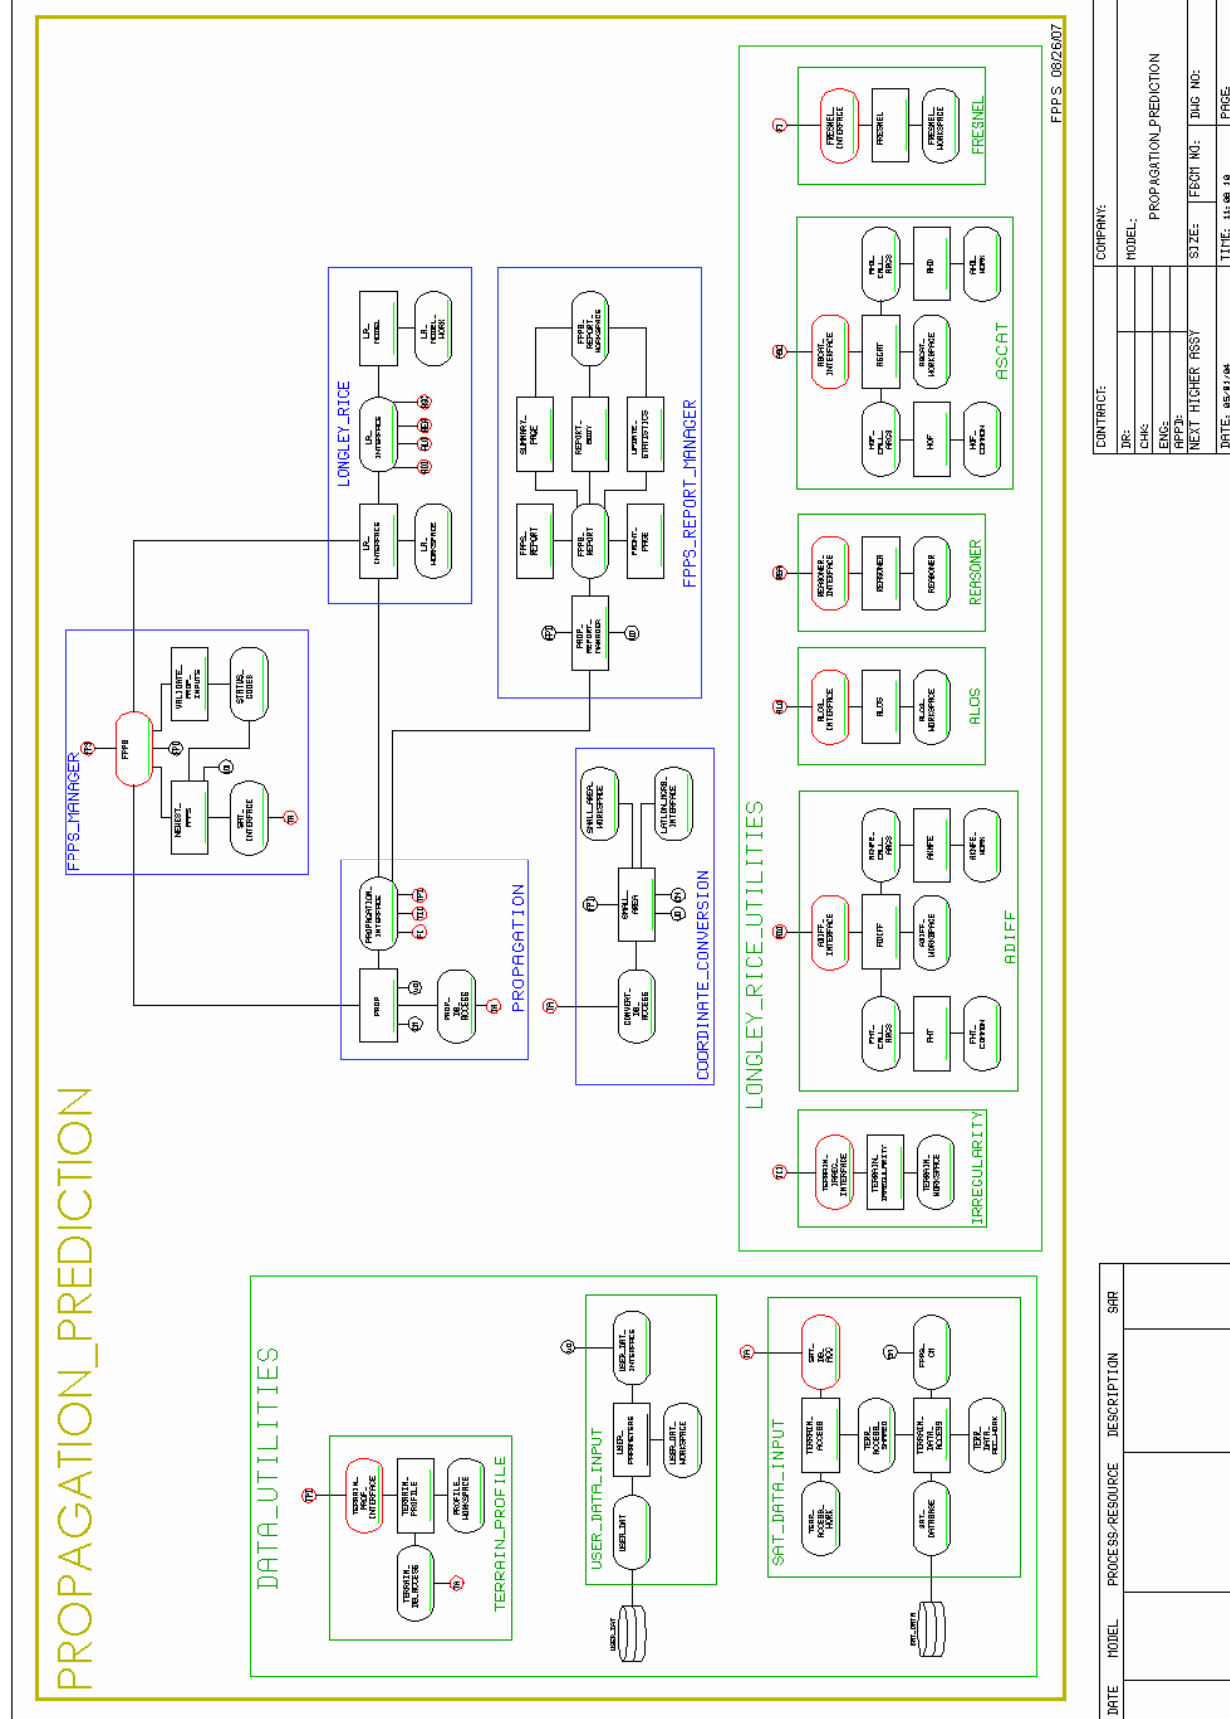


Figure 10. Engineering drawing of a library module.



## Types Of Modules

The types of modules that make up the layers of a software design hierarchy are shown in Figure 12. These types provide different levels of protection with regard to their reuse in different hierarchies. Both elementary and hierarchical modules can reside within each type.

The rules for these types are described below with examples that follow.

- **Modules** - have a blue border. These are the basic building blocks in a task. In the CAD system described here, modules may be decomposed hierarchically, i.e., they may contain submodules and sub-submodules, etc.
  - **Elementary Modules** - contain resources and processes.
  - **Hierarchical Modules** - contain elementary modules and hierarchical modules.
- **IND Modules** - have a brown border. IND Modules only share Inter-Processor (IP) Resources externally - and only with other IND Modules. When using parallel processors, IND Modules must be the highest level modules on a processor. IND Modules may reside on the same or different processors.
- **Utility Modules** - have a green border. These are modules that are reused by processes in the same directory, and can appear in more than one hierarchy in different drawings. They are typically used to manage separate databases or perform utility type functions. The green color distinguishes them for change protection. They can only be changed in their own drawing. If they are changed to accommodate a different requirement, that change must be compatible with all processes that use them. Separate copies automatically reside on each processor that uses them.
- **Library Modules** - have a gold border. These are highly protected utility modules that can be shared from different directories and computers, being stored as object modules in special object library files. The source only appears in the directory where they are maintained. Library module Processes are called from an application using their process name, module name, and library name. Since each of these names must be unique within the next level of hierarchy, there can be no duplicate names when linking to library modules in the CAD environment described here. Separate copies reside on each processor that uses them.

The functions of a library module may be upgraded while at the same time preserving the original module in the library for prior users. Users can call the new function using the same process name within the same library by using a new module name. The existing CAD system has a large set of libraries that support various applications, including 3D graphics. These are shared easily.

## Instanced Modules

Modules and IND Modules can be instanced (number of instances appear in red). This allows assignment of groups of instances to different processors. This facility is supported by the Control Specification language, providing the ability to assign specific sequences of instances to different processors to maximize PUE.

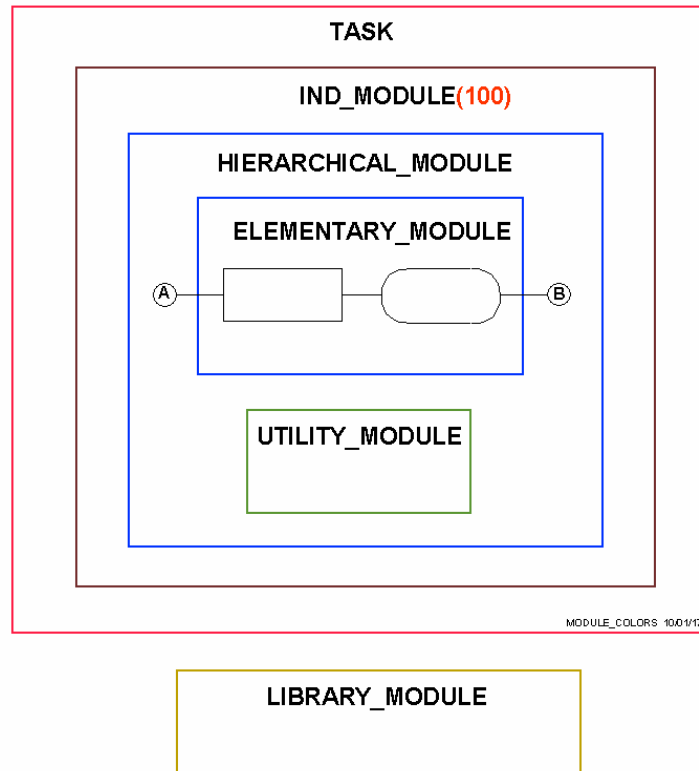


Figure 12. VisiSoft hierarchy of modules.

When building complex software systems on parallel processors, it is important to select the proper module types at the engineering drawing level. Module types are used to create hierarchies, and to differentiate how they are used and accessed. Depending upon the module type and what is required architecturally, modules must share specified resource types as dictated by design for concurrent operations. Equally important are selection of specific types of resources, also designated at the drawing level using different colors as described below.

### Types Of Resources

VisiSoft provides a hierarchy of types of data memory resources to support different architectural requirements. It is the architecture - augmented by different resource types - that clearly simplifies the design of complex software systems, particularly those implemented using parallel processors. Resources also facilitate direct mapping of complex mathematical spaces into well organized data descriptions using an easily understood English-like language. The language itself provides for complex hierarchical data descriptions within a resource. Designers can create or modify resource types using buttons and panels in the Visual Development Environment (VDE) in Figure 6. The hierarchy of different resource types is shown in Figure 13.

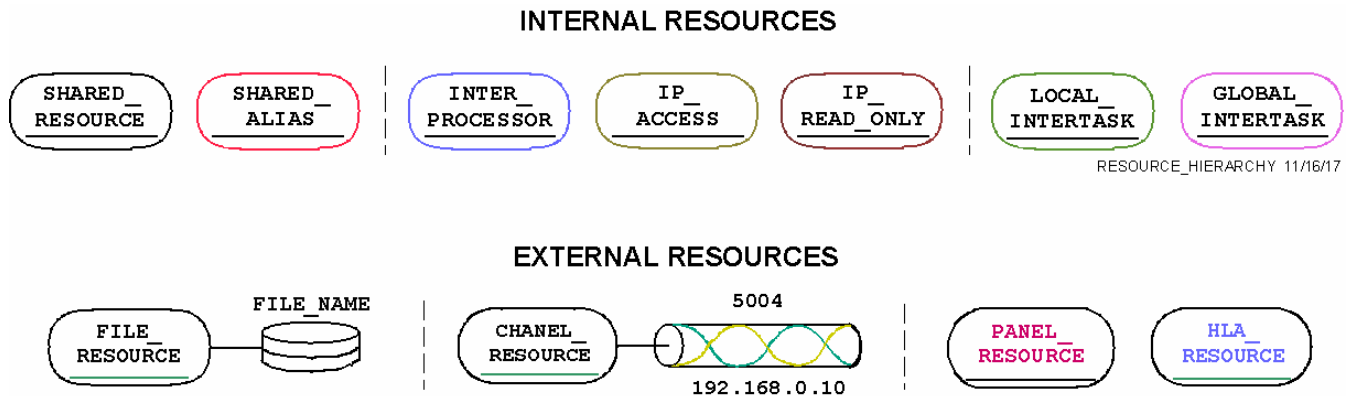


Figure 13. Different types of memory resources.

### Internal Resource Types

There are three types of Internal Resources. These are defined below, along with their sub-types. Refer to Figure 14 below for some examples of the first two types.

1. Shared Directly Within A Task
  - Shared Resource - Directly between processes
  - Shared Alias - By pointer to a resource shared with the Calling process
2. Shared Across Parallel Processors on a SIMPLEX Basis Within A Task
  - Inter-Processor - Written only by Attached Processes in same IND module
  - IP Access - Accessed by Attached Processes (can be Modified)
  - IP Read-Only - Accessed Read-Only by Attached Processes
3. Shared Between Tasks
  - Local Inter-Task - Can only be Accessed by processes in the Family
  - Global Inter-Task - Can be Accessed by processes in Any Task

In Figure 14, the IND\_PLATFORM module is instanced 8 times with all of the IP resources connected to those on other processors. The interior module MODULE\_M\_2 is instanced 100 times. In this module, the IP Access resource HR\_C2 is very small, does not eat up much memory, and is used by process HP\_AC2 to make changes in it before moving these changes to HR\_Z2. However, HR2\_RO is very large and does not have to be changed (written to). The 100 instances of this resource are multiplied by all 8 of the IND\_PLATFORM module instances yielding 800 total instances of this large resource. Using the READ ONLY resource, these are eliminated and replaced by a single resource that may be accessed by all processes that reside on that processor. This saves a huge amount of memory.

Even if some of the IND Module instances reside on a different processor, only one additional copy of HR2\_RO is needed for that processor. This holds for different IND Modules that also access HR2\_RO. Only one Read-Only copy is needed on a processor.

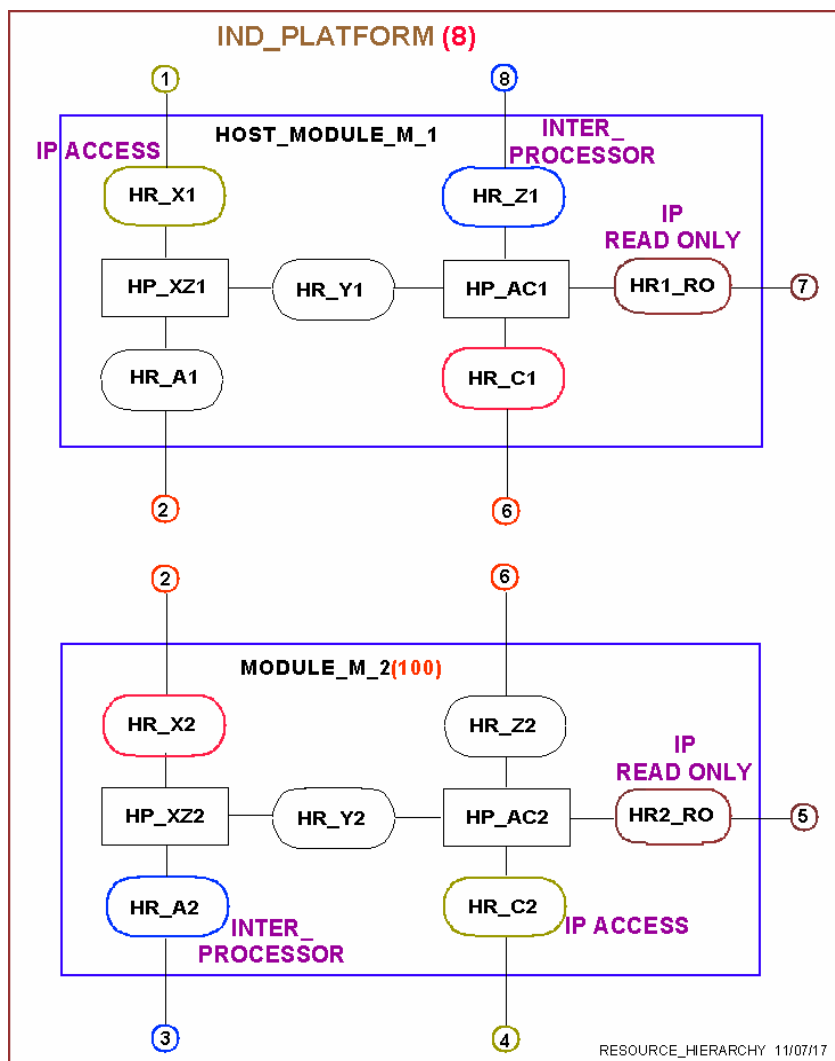


Figure 14 - The Architecture of Memory Resources (Spaces).

IP/IP Access Resource affinity is similar to that between a SHARED AS resource and an ALIAS resource, with the following exceptions:

- ALIASed resources do not have their own memory - just a template. For example, looking at Figure 14, when process HP\_XZ1 CALLs HP\_XZ2 - it automatically points to the SHARED\_AS resource HR\_A1 to be used as a data map of the ALIASed resource HR\_X2. Conversely, in Figure 14, the IP\_ACCESS resource HR\_X1 has its own memory. An ACCESS statement in HR\_XZ1, e.g., ACCESS HR\_X1OP (on another processor), will cause the most recently RELEASEd copy of the IP resource HR\_X1OP to be copied to HR\_X1. Note that HR-X1OP need not reside on a different processor.
- The SHARED AS/ALIAS architecture establishes an automatic link based on the architecture - without the need for any special statements. IP resources are also automatically linked based on the architecture and supported by the RELEASE and ACCESS statements.

## External Resource Types

There are four types of External Resources shown in Figure 13. These are defined below.

1. File Input / Output Resource

This resource is used to describe the records to be read from - or written to - a file.

2. Channel Input / Output Resource

This resource is used to describe the packets to be read from - or written to - a TCP/IP Channel given the Port and Channel numbers.

3. PANEL Resource

This resource is used to interface with interactive Panels created using the VisiSoft Panel Library Manager to observe, enter, and modify information interactively during a task.

4. HLA Resource

This resource is used to describe messages to be taken from - or sent to - another member of an HLA network using IP protocols.

The Resource language has been designed to support hierarchical data descriptions that map directly into complex application spaces that are easily understood by application experts. This is particularly apparent when defining complex file, record, and packet structures containing multiple layers of hierarchy.

## Module Independence

To assess parallelism, one must understand the property of module independence. Two modules are *spatially independent* (independent of time) if they share no resources. From an application standpoint, these modules have no direct influence on each other.

Two modules are *temporally independent* if they only influence each other at specified points in times. A typical example is when using differential equations to represent dynamic systems. Computer solutions to differential equations typically assign a  $\Delta T$  time step that is small enough to ensure sufficient accuracy of solutions when compared to actual tests ( $\Delta T$  may vary for nonlinear systems). This implies that equations are solved at successive time steps,  $(T + \Delta T)$ , so that changes in one part of a system cause sufficiently accurate effects in another part. Information on changes can be exchanged at the beginning or ending of each time step, affecting calculations within the desired time step.

More generally, modules are *temporally independent* if they share information on a synchronized basis. This requires a facility where only one module can copy the information to be shared into a resource held for other modules. Other modules can copy that information into their own resources making the information available for follow-on calculations. System level states may be defined by the designer to ensure synchronization. These states are shared by pairs of IND Modules. The module changing the resource can set a state to indicate it has updated the resource shared by the pair. The module reading that resource then resets it showing that the latest copy has been read. Just as the independence of the real physical systems they may represent, neither module need stop to wait for the other, thus maximizing PUE. An example of such a physical system is a group of people exchanging different emails.



## **Use Of Facilities That Support The Requirements**

Designing systems that use the above facilities to run modules concurrently on parallel processors requires an in-depth knowledge of the particular application being developed. Only application experts will know how to use the above facilities to deal with the timing and synchronization necessary to ensure proper operation of such systems. Therefore, it is imperative that these facilities be easy to use by those experts.

## **Timing & Synchronization Of Observations, Actions, And Outcomes.**

Software systems running on parallel processors typically represent Objects / Entities that:

- Operate Concurrently
- Act / Interact Based On factors such as:
  - Gravitational Forces
  - Electro-Magnetic Forces
  - Observation of states of other objects on different processors  
e.g., position, velocity, transmissions, other actions

Models of these systems must meet, support, and represent the required operational properties of the application with sufficient accuracy. This generally implies:

- Implementing decision processes that produce the desired behaviors of the application.
- These decision processes are typically implemented using complex algorithms operating on complex data spaces.
- These decision processes typically produce responses / actions that are observed by and affect the behavior of other entities in the system.
- The timing and synchronization of - the observations of and responses to - these actions are typically critical to affecting the correct outcomes and responses of the system.
- The representation and implementation of these concurrent observations and actions is critical to achieving the level of accuracy of outcomes required by the application.

This implies maintaining consistency and coherency of data exchanges while operating concurrently on parallel processors.

## **Capabilities Needed To Support The Above Requirements**

Models of entities must be able to:

- Obtain information describing the behavior of other models that are running concurrently.
- Provide information describing their own behavior that can be obtained by other models running concurrently.
- Observe the reactions of other models.
- Complete observations and reactions within specified time frames.



Physical systems can perform these functions while operating concurrently. With the right facilities, these systems can be represented by models operating concurrently on parallel processors.

### **Facilities To Support The Above Requirements**

Software facilities must exist that provide the ability to:

- Schedule events to occur in the future that represent time to complete actions.
- Release information to other models while running concurrently.
- Directly access information from other models that are running concurrently.
- If necessary, wait for information needed to make decisions.

### **TIME SYNCHRONIZATION AND SPEED**

Given the temporal independence condition, changes in one IND module that affects another on a different processor will be reflected with sufficient accuracy as long as the effects are resolved within a user specified  $\Delta T$  time period. This requires the ability to *synchronize information exchanges*. This implies that the times when: (1) a sending process updates its IP resource copy; and, (2) the receiving process reads it - both fall within an allowed  $\Delta T$ .

Synchronization is accomplished when the shared IP Resource is exchanged within successive  $\Delta T$  time periods.

In typical applications, speed is most important, determining whether application run-time constraints are met. When applications impose such a constraint, the optimal design problem is to minimize the number of processors required to meet that constraint. When time constraints can be met using a smaller number of processors, speed will rise exponentially just due to the smaller footprint, dramatically reducing power requirements as well as floor space.

Tests have shown that applications built in VisiSoft can increase speed by 2 orders of magnitude - compared to typical software approaches - *just on a single processor*. When using a parallel processor, reducing the number of processors by only a single order of magnitude can dramatically increase the speed further. Adding the PUE and other multipliers, *one may expect to achieve 3 to 4 orders of magnitude increase in speed using VisiSoft on a parallel processor*.

### **PARALLEL PROCESSOR SOFTWARE DESIGN**

To solve parallel processor software problems, one must take maximum advantage of the inherent parallelism in a particular application to minimize running time. Figure 15 provides an expanded depiction of the steps required to map parallel application requirements into a software space. This requires a software environment that simplifies the design effort for application experts. As described above, this environment must also support ease of translation of the software design into a hardware environment that maximizes understanding as well as speed. *Much of the burden of design and implementation is then shifted from the developer to highly sophisticated architecture and language translators* - using the computer. As described by Peter van der Linden, [17 - pg 64], this is just where the burden should be.

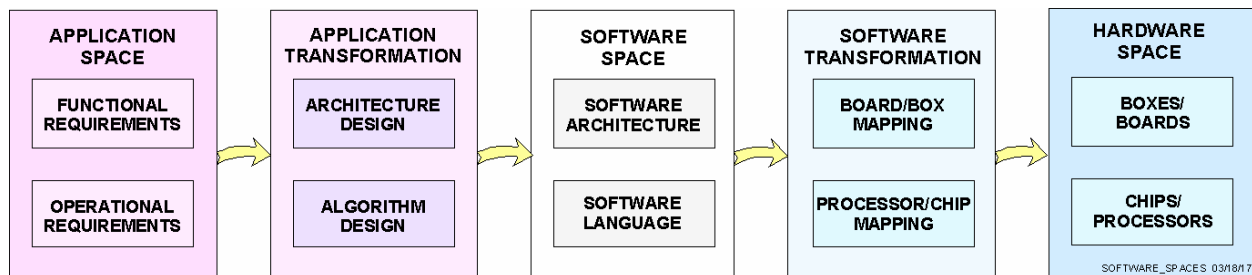


Figure 15. Spaces for translation of application requirements into software and hardware.

The illustration of transformations in Figure 15 implies the following facilities:

- Mapping the functional and operational requirements of an application into a software architecture and supporting algorithms requires: (1) Application experts with knowledge of the complex system and corresponding event spaces; and (2) Graphical facilities for visualization of the software architecture.
- Mapping application space requirements into a software design for a parallel processor hardware space requires special architectural and language facilities. Graphical visualization of the architecture is critical to this process.
- Complex software spaces require the use of deep hierarchical data structures, organized in accordance with the application space - not by data type. These spaces must be easily organized and understood by application experts.
- Complex software algorithms require the use of deep hierarchical rule structures. These must be organized in accordance with the application space - using one-in one-out control structures. The language must be easily understood by application experts.
- The need for application expertise is required to design Independent (IND) Modules. Special graphical visualization tools are critical to simplify understanding of both the module design and mapping process.
- Mapping Independent IND Modules over heterogeneous spaces can provide substantial improvements in speed. Tests must be run on each IND Module to determine their relative speeds - for processor placement - to maximize PUE.
- IND Module run-time speed measurements are provided automatically using the VisiSoft Parallel OS (VPOS) and may be viewed with bar graphs of IND Module run-times on each processor. Speed is improved by reorganizing module assignments.
- Assignment of IND Module Types across Boxes of chips based on physical system connectivity can substantially minimize delays over direct memory access channels.
- During initialization, the run-time system provides VPOS with architectural information derived from the development environment. This coupled with the expert user allocation of IND Modules eliminates most run-time memory management.
- Cross-processor event states save huge amounts of time by synchronizing the exchange of information between processes that continue to run concurrently.

The above bullets represent some of the many facilities that have been built, tested and proven to work on many projects since 1982. It is hard to see how one can obtain the many orders of magnitude of increased speed without the approach highlighted here.

### **The GLOBAL\_PLANNER Simulation**

Figure 16 illustrates the facilities described above in a highly nonstationary GLOBAL\_PLANNER simulation. This simulation contains 16 IND Modules, 9 of which are instantiated. Each module (instance) is mapped onto one of 14 processors to minimize PUE. Run-time results of the PUE mapping is shown graphically in Figure 17 for 3  $\Delta T$  time steps. This mapping shows the utilization of the 14 processors, where the only wasted time is at the end of each bar. This illustrates the typical result of 95% PUE for the 3  $\Delta T$  time steps, well above any other approach.

### **OPTIMIZING PHYSICAL INTERCONNECTIONS**

When considering the applications listed above, they all represent physical systems. These systems may be thought of in many ways, e.g., particles, molecular structures, moving platforms, etc., all with limited purview when embedded in huge groups. One can also think of the processing abilities of the brain, its speed and minimal energy dissipated compared to the current design of servers. Figure 18 illustrates the direct synapse interconnections. There is no separate communication system between synapses.

Server environments require a separate communication system, and the corresponding protocol layers built in software - so that one task can talk directly to another task on a different processor far away. When modeling physical systems, one can route shared information through adjacent processors. For the listed applications, just exchanging information with adjacent neighbors is generally sufficient. When information must be routed through a network to nodes farther away, the fastest approach is saturation signaling (also known as flooding) through sequences of nodes. This may be used to pass information to multiple processors or find the best (fastest) path to the desired processor, see [18]. This leads to a new hardware design approach for all of the parallel processor applications listed above.



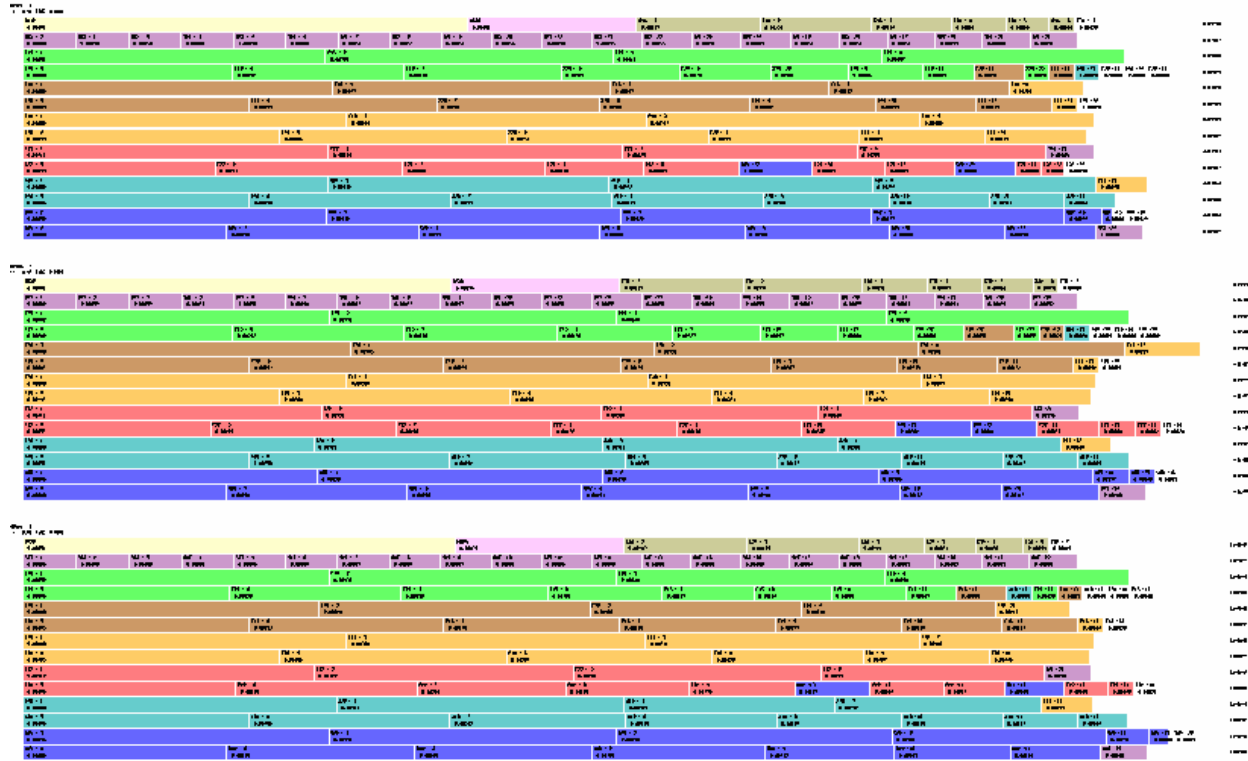


Figure 17. A graphical representation of the Processor Utilization Efficiency (PUE).




Figure 18. A typical representation of the interconnection of synapses in the brain.

## HARDWARE DESIGN TO SUPPORT APPLICATION SPACE ARCHITECTURE (ASA)

Figure 19 illustrates the interconnection of 27 GreenGeneMachine (GGM) Boxes, sharing VisiSoft Resources directly via Direct Memory Access (DMA) channels between boxes. With 120 processors in each GGM Box (can go to 144), this configuration contains a total of 3240 processors. With only a 3 order of magnitude gain, this facility will beat any with 3,000,000 processors. But fair comparisons must use real applications - like those in the representative list provided above.

Virtually all of the physical systems of interest for parallel processing can be mapped into a generalized 3D space, sharing data with all of their adjacent neighbors. There is no need to waste time with cache coherency and communications protocols to route data between processors. Memory management is performed based on the architecture of Independent (IND) Modules - the basic component of a VisiSoft software application. Figure 19 illustrates the need to communicate directly with adjacent neighbors in all 26 directions. Figure 20 shows the special design of the GGM Box with direct connections using all faces of the Box to account for corners and edges as well as faces.

Without the VisiSoft architectural approach to software design, the special language facilities to implement this with huge speed gains, and the engineering facilities used to share memory using VisiSoft Resources (allowing the design of deep hierarchical data spaces), none of this would be possible. One can achieve 2 orders of magnitude improvement in speed using VisiSoft - just on a single processor. This translates to 4 orders of magnitude due to the reduced hardware footprint. With the parallel processor facilities built into the system, one can expect 4 to 6 orders of magnitude improvement over applications written using today's most popular software languages. With the use of heterogeneous data spaces, one may expect more than 4 to 6 orders of magnitude increase in speed over current approaches.

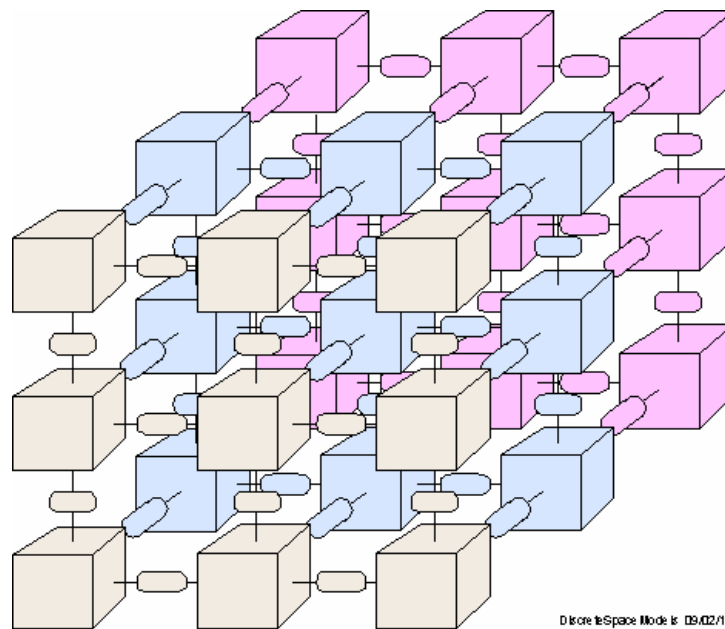


Figure 19. Processors in different GGM Boxes are directly connected by DMA channels.




Figure 20. Each Box has twenty-six adjacent Boxes, all directly interconnected.

From the above, one sees that software development environments affect the design of both the run-time environment and the operating system. Together, they both affect the design of the hardware. Given the VisiSoft CAD facilities described here, many hardware facilities can be eliminated from parallel processor chips, such as the following.

- Routers - replaced by DMA channels between chips, boards, and Boxes
- Cache coherency replaced by IND Modules
- Thread synchronization replaced by design
- Stack facilities

Note that these facilities are best replaced by more cache close to the processors or more processors on a chip, the major factors in achieving speed. We also note that the number of parallel processors in Figure 19 may be small compared to some actual environments. However, that configuration easily fits in a small rack that is sufficiently fast for most applications.

Unless one has participated directly in development of VisiSoft architectures, the power of these facilities may take time to comprehend. Once used, it is apparent that the architectural simplifications provided are just as critical to single processor software design. It also becomes obvious that the ability to design good architectures depends directly on the language. It is why productivity multipliers are very high when using this CAD environment, especially in the support mode when a new person has to understand what another has built. This requires a willingness to compare different approaches based on real measures of speed and productivity.

## SUMMARY

Design of complex application systems that require parallel processing to meet speed requirements is difficult using current approaches. This paper briefly skims aspects of system design to illustrate some of the top level concepts that must be applied to solve this problem. It is intended to demonstrate an architectural perspective that equates to similar engineering fields, e.g., aeronautical, architectural, electrical, etc. The skills required to design skyscrapers require higher level facilities for designing large complex structures. These skills and facilities are not needed to design simple buildings. Learning to program small problems using different languages is akin to learning the architectural design of small houses. Large complex systems require a hierarchy of architectural facilities that reside above the code.

This paper only covers top level concepts and facilities that simplify achieving high performance computing. They are part of a system that has evolved building hundreds of large simulations since 1982. This system ensures synchronization of modules running concurrently on parallel processors so they do not have to stop to exchange information. Explicit passing of control to processes on different processors is slow and unnecessary. Scheduling processes to run in future  $\Delta T$  time periods, or scheduling a process given that an event has occurred on another processor eliminates time wasted waiting for events or testing changes in states.

The facilities built into the VisiSoft Parallel OS (VPOS) support direct synchronization, and eliminate stack management, memory management, and especially coherency management, saving huge amounts of time. VPOS contains accurate clocking, providing measures of processor utilization using graphical representations and supporting optimization of independent module-processor assignments. The results of careful testing and design have produced performance improvements of four to six orders of magnitude for typical large scale parallel processor applications. Tests of these results can be reproduced independently.

The visual simplifications provided by these architectural facilities become more obvious as they are used. Without realization of the importance of the properties of independence and understandability, this architectural approach would have never been developed.

This paper shows that complex automation problems cannot be solved with simple programming languages. Software language must support numerous architectural requirements as well as ease of understanding. As in other engineering disciplines, it is the architectural technology that is required to design large complex systems and simulations. This is particularly true when designing systems to be run on parallel processors. Without the VisiSoft architectural facilities, it is hard to understand parallel processing problems let alone solve them.

We acknowledge the time it takes to evolve chip designs. However, our tests achieved huge order of magnitude increases in speed using existing chips with 18 processors each, and limited L1 and L2 cache sizes. These speed gains resulted from the software. New board designs can be produced rapidly. Using an incremental approach, this technology can evolve with additional speed gains every few years.

We also recognize that this represents a major paradigm shift, so that gaining acceptance by early adopters would likely be on new project development. But when one examines the rapid accumulation of savings on existing projects, and the ease with which existing software is translated into a much more understandable system, the conversion of existing projects becomes obvious. In any event, we believe this shift is inevitable.



## REFERENCES

1. Anselmo, Donald and Henry Ledgard, Measuring Productivity in The Software Industry, *Communications of the ACM*, Vol. 46. No.11, Nov 2003.
2. Anselmo, Donald, *Why Software Productivity Has Not Improved*, Software Summit, Washington, D.C., May 2004.
3. Armour, Phillip G., Software: Hard Data, *Communications of the ACM*, Vol. 49. No.9, Sept. 2006.
4. Beyer, Kurt W., *Grace Hopper and the Invention of the Information Age*. Cambridge, MA: The MIT Press, (2009). ISBN 978-0-262-01310-9.
5. Cave, W.C., et.al, *A CAD System For Developing Complex Software - An Engineering Approach*, Visual Software International, Spring Lake, NJ, September, 2014.  
[http://www.VisiSoft.com/PDF\\_Files/ACADSystemForSoftware.pdf](http://www.VisiSoft.com/PDF_Files/ACADSystemForSoftware.pdf)
6. Cave, W.C., et.al, *Software Theory For Parallel Processors*, Visual Software International, Spring Lake, NJ, March, 2016. [http://www.VisiSoft.com/PDF\\_Files/SoftwareTheoryBook.pdf](http://www.VisiSoft.com/PDF_Files/SoftwareTheoryBook.pdf)
7. Groth, R., *Is the Software Industry's Productivity Declining?*, *IEEE Software*, Nov/Dec 2004.
8. Kambayashi, Yasushi and Henry F. Ledgard, "The Separation Principle - A Programming Paradigm" *IEEE Software*, March/April 2004
9. Krishnadas, L.C., *EE Times Article*, Jan 17, 2008.
10. Ledgard, H., et al, "The Natural Language of Interactive Systems," *CACM* No. 10, October 1980, pp 556-563.
11. Ledgard, Henry F., *The Emperor with No Clothes*, *Communications of the ACM*, Oct 2000.
12. Merritt, R., *Multicore Puts Screws to Parallel Programming Modules*, *EE Times On Line*, Feb 15, 2008.
13. Parnas, D., "Education for Computer Professionals," *IEEE Computer*, January 1990, pp 17-22.
14. Poore, Jesse H., *A Tale of Three Disciplines and a Revolution*, *IEEE Computer Society*, Jan 2004.
15. Shannon, C.E., *A Mathematical Theory Of Communication*, *BSTJ*, Vol.27, pp 379 & 623, Jul & Oct 1948.
16. Stroustrup, B., "What is Object-Oriented Programming?" (1991 revised version). Proc. 1st European Software Festival. February, 1991. <http://www.public.research.att.com/~bs/whatis.pdf>
17. van der Linden, P, *Expert C Programming - Deep C Secrets*, SunSoft Press - Prentice Hall, Englewood Cliffs, NJ, 1994.
18. Cave. W.C. & Dunn, R.M., *Saturation Processing: An Optimized Approach to a Modular Computer System*, ECOM Technical Report-2636, Ft Monmouth, NJ, November 1965.  
[http://www.VisiSoft.com/PDF\\_Files/Saturation\\_Processing.pdf](http://www.VisiSoft.com/PDF_Files/Saturation_Processing.pdf)