

A Survey on the Design of Multiprocessing Systems for Artificial Intelligence Applications

BENJAMIN W. WAH, SENIOR MEMBER, IEEE, AND GUO JIE LI, MEMBER, IEEE

Abstract—Some important issues in designing computers for artificial intelligence (AI) processing are discussed. The issues discussed are divided into three levels: the representation level, the control level, and the processor level. The representation level deals with the knowledge and methods used to solve the problem and the means to represent it. The control level is concerned with the detection of dependencies and parallelism in the algorithmic and program representations of the problem, and with the synchronization and scheduling of concurrent tasks. The processor level addresses the hardware and architectural components needed to evaluate the algorithmic and program representations. Solutions in each level are illustrated by a number of representative systems. Design decisions in existing projects on AI computers are classified into the top-down, bottom-up, and middle-out approaches.

I. INTRODUCTION

IN RECENT YEARS, artificial intelligence (AI) techniques have been widely used in various applications, such as natural-language understanding, computer vision, and robotics. As AI applications move from the laboratories to the real world and as AI software grows in complexity, the computational throughput and cost are increasingly important concerns. The conventional von Neumann computers are not suitable for AI applications because they were designed mainly for sequential and deterministic numeric computations. Extensive efforts have been devoted to investigate and develop efficient AI architectures [188]. This paper provides a state-of-the-art assessment of AI-oriented systems and discusses the major issues involved in such designs.

A. Characteristics of AI Computations

To develop a special-purpose computer to support AI applications, the requirements of these applications must be fully understood. Many conventional numeric algorithms are well analyzed, and bounds on their computa-

tional performance have been established. In contrast, many AI applications are characterized by symbolic processing, nondeterministic computations, dynamic execution, large potential for parallel and distributed processing, management of extensive knowledge, and an open system.

Symbolic Processing: Data are generally processed in symbolic form in AI applications. Primitive symbolic operations, such as comparison, selection, sorting, matching, logic set operations (union, intersection, and negation), contexts and partitions, transitive closure, and pattern retrieval and recognition, are frequently used. At a higher level, symbolic operations on patterns such as sentences, speech, graphics, and images may be needed.

Nondeterministic Computations: Many AI algorithms are nondeterministic, that is, planning in advance the procedures to execute and to terminate with the available information is impossible. This is attributed to a lack of knowledge and a complete understanding of the problem; it may result in exhaustively enumerating all possibilities when the problem is solved or in a controlled search through a solution space.

Dynamic Execution: With a lack of complete knowledge and anticipation of the solution process, the capabilities and features of existing data structures and functions may be defined and new data structures and functions created while the problem is actually being solved. Further, the maximum size for a given structure may be so large that it is impossible to allocate the necessary memory space ahead of time. As a result, when the problem is solved, memory space and other resources may have to be dynamically allocated and deallocated, tasks may be dynamically created, and the communication topology may be dynamically changing.

Large Potential for Parallel and Distributed Processing: In parallel processing of deterministic algorithms, a set of necessary and independent tasks must be identified and processed concurrently. This class of parallelism is called AND parallelism. In AI processing, the large degree of nondeterminism offers an additional source of parallel processing. Tasks at a nondeterministic decision point can be processed in parallel. This latter class is called OR parallelism.

Manuscript received June 24, 1987; revised September 8, 1987 and January 29, 1989. This work was supported by the National Aeronautics and Space Administration under Contract NCC2-481.

B. W. Wah is with the Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801.

G. J. Li was with the University of Illinois, Urbana, IL. He is now with the Institute of Computing Technology, Academia Sinica, P.O. Box 2701, Beijing, People's Republic of China.

IEEE Log Number 8927232.

Knowledge Management: Knowledge is an important component in reducing the complexity of solving a given problem: more useful knowledge means less exhaustive searching. However, many AI problems may have very high inherent complexity, hence the amount of useful knowledge may also be exceedingly large. Further, the knowledge acquired may be fuzzy, heuristic, and uncertain in nature. The representation, management, manipulation, and learning of knowledge are, therefore, important problems to be addressed.

Open System: In many AI applications, the knowledge needed to solve the problem may be incomplete because the source of the knowledge is unknown at the time the solution is devised, or the environment may be changing and cannot be anticipated at design time. AI systems should be designed with an open concept and allow continuous refinement and acquisition of new knowledge.

In general, two basic approaches are available for improving the computational efficiency of processing AI tasks: having heuristic knowledge to guide searches and using faster computers. In the following sections, these approaches will be discussed.

B. Heuristic Searches

The key performance-related feature of AI computations is their nondeterminism, which results from a lack of complete understanding of the solution process. In other words, when a problem becomes well understood and can be solved by a deterministic algorithm, we usually cease to consider it "intelligent," although the problem may still be symbolic [155].

The starting point of conventional computations is deterministic algorithms, whereas efficient deterministic algorithms to solve a given AI problem is result from the knowledge accumulated and the gradual refinement of the computations. This involves the succinct choice of an appropriate knowledge-representation scheme, learning mechanisms to acquire the related knowledge, and a suitable architecture to support the computations. Good heuristics designed from previous experience may allow a complex problem to be solved efficiently, even on a serial processor.

Since the mid-1960's, the AI community has realized that inference alone was often inadequate to solve real-life problems. To enhance the performance of AI algorithms, they must be augmented with knowledge and metaknowledge of the problem domain in addition to formal reasoning methods. Metaknowledge refers to the control information to guide the search. This realization gave birth to *knowledge engineering* and *knowledge-based systems*, the field of applied AI [54]. Since the knowledge stored in any knowledge-based system may be incomplete and inaccurate, combinatorial searches are still needed.

C. Faster Technologies and Parallel Processing

An AI computer system must support both knowledge-base management and heuristic searches. Faster technolo-

gies and parallel processing are means to improve the computational efficiency. For many applications, such as natural-language understanding and computer vision, the current achievable performance is much lower than that needed. For example, according to DARPA's Strategic Computing proposal, it was estimated that an equivalent of one trillion von Neumann computer operations per second were required to perform the vehicle-vision task at a level that would satisfy the long-range objective of the Autonomous Vehicle Project [1]. At best, current sequential computers of reasonable cost achieve processing rates below 100 million operations per second, which implies at least 10^4 times improvement in performance are required.

Newer technologies can help in designing faster computers. For example, using GaAs high-electron-mobility transistors (HEMT's), it was estimated for a computer with over 500000 gates operating at 77 K and 15 levels per pipeline stage, the cycles times were predicted to be 2.7 ns with 5 W and 3200 gates per chip, and 2.0 ns with 20 W and 5200 gates per chip, respectively [10]. In contrast, a liquid-cooled Cray 2 supercomputer built using ECL technologies has eight levels per pipeline stage, more than 500000 gates, and operate at 300 K and 4.1 ns/cycle. The delay of one ECL gate level is approximately translated into 1.5 GaAs HEMT gate levels; hence correcting the cycle time of the Cray 2 supercomputer into HEMT technologies and 15 levels results in 5.1-ns cycle time for the Cray 2 computer. In short, there is a factor of two in using the newer technologies available today.

Another way to reduce the cycle time is to reduce the interconnect delay. It was estimated that with GaAs HEMT operating at 2-ns cycle time, the switching, fan-out, and interconnect delays were approximately 2, 10.5, and 87.5 percent of the cycle time, respectively [10]. Although superconductivity can be used to reduce the interconnect delays, it is less desirable with GaAs technologies due to the high impedance in the gates, and more desirable with ECL technologies. When combined, these newer technologies available today may allow improvement in the cycle times of one to two orders of magnitude.

The trend in design AI computers has been toward applying faster technologies and parallelism to process computation-intensive AI tasks. Examples of parallel AI systems currently available or under research/development include Alice, Aquarius, Butterfly, Concurrent Lisp machine, Connection Machine, Dado, Faim-1, FFP, iPSC, Japanese Fifth Generation Computer System (FGCS), NETL, Non-Von, Rediflow, Soar, Spur, and ZMOB [188]. Some of these computers, such as the Aquarius, Butterfly, iPSC, and ZMOB, were designed for both numeric and symbolic processing.

Recently, there is another trend to design small-grain massively parallel architectures for AI applications. These architectures are sometimes called *connectionist systems*; they are composed of a very large number of simple processing elements. Knowledge of a given entity in such systems are distributed on a number of processing elements and links, and each processor or link may be shared

TABLE I
RELATIVE PROBLEM SIZES SOLVABLE IN A FIXED AMOUNT OF TIME ASSUMING
LINEAR SPEEDUP*

Complexity to find Optimal Solution	Number of Processors					
	1	N	N ²	N ³	N ^k	2 ^N
N	N	N ²	N ³	N ⁴	N ^{k+1}	N ^{2N}
N ²	N	N ^{1.5}	N ²	N ^{2.5}	N ^{k/2+1}	N ^{2N/2}
N ³	N	N ^{1.33}	N ^{1.67}	N ²	N ^{k/3+1}	N ^{2N/3}
N ^k	N	N ^{1+1/k}	N ^{1+2/k}	N ^{1+3/k}	N ²	N ^{2N/k}
2 ^N	N	N + log N	N + 2 log N	N + 3 log N	N + k log N	2N

* Problem size, when sequential processing is used, is N.

by multiple entities. The use of connections rather than memory cells as the principal means to store information leads to the name "connectionism" [53]. The resemblance to neurons in a brain also results in the name "neural networks." Many computers can simulate connectionist systems. An example is the Connection Machine developed by Thinking Machines Inc., which can perform neural-network simulations two to three orders of magnitude faster than serial machines of comparable cost [87], [189].

The high performance in many parallel AI computers is achieved through associative processing and "data-level parallelism." This approach is suitable for operations on large databases, such as sorting, set operations, statistical analysis, and associative pattern matching. Yet data-level parallelism alone is not enough. For general AI applications involving heuristic searches, control-level parallelism should be involved. Unfortunately, early experience with multiprocessor architectures for Hearsay-II [55], Eurisko [111], OPS-5 [60], and others have led to a belief that parallel AI programs will not have a speedup of more than one order of magnitude. A possibly revolutionary approach to designing parallel languages and systems for AI processing may be needed.

One misconception in parallel processing is to use the total computing power of a parallel system to characterize the rate at which a given AI application is processed. Due to the nondeterminism in AI computations, a high computing power does not always imply a shorter completion time. Since most AI applications involve heuristic searches, resources may be devoted to fruitless searches, which use more computing power but do not help to decrease the time to find a solution. In fact anomalies may happen such that increasing the degree of parallelism may even increase the completion time in nondeterministic searches [116], [187]. What is important is how to allocate resources so only useful tasks are performed. The question of solving an AI problem in a parallel processing environment is still largely unanswered.

Another misconception about parallel processing is that it can be used to extend the solvable problem size of AI problems. Due to the high complexity of AI problems, parallel processing is useful in *improving the computational efficiency*, but not in *extending the solvable problem size* [187]. For example, a problem of size N and complexity N^k can be solved in N^k time units by a sequential processor. Assuming that N processors are used, the new prob-

lem size X that can be solved in the same amount of time satisfies

$$N^k N^k = X^k.$$

The left side of the equation represents the total computing power in N^k units of time with N processors, and the right side represents the number of operations to be performed in solving a problem of size X. Solving the previous equation yields

$$X = N^{1+1/k}.$$

Table I summarizes the results for other cases. It is assumed that the size of the problem solved by a sequential processor is N, that the number of parallel processors ranges from 1 to 2^N, that linear speedup is achievable, and that the same amount of time is allocated to both sequential and parallel processing. The first column shows the complexities of solving the problem optimally, and the other columns show the corresponding sizes of the same problem that can be evaluated in the same amount of time for various number of processors. The extension in problem size is minimal when the problem involved is complex. This is evident in the last row in which the problem solved has exponential complexity. In this case only a logarithmic increase in problem size is achieved when a polynomial number of processors are used, and a linear increase is resulted with exponential number of processors.

In essence, parallel processing alone cannot circumvent the difficulty of combinatorial explosion. The power of multiprocessing should not be overemphasized and must be combined with heuristic information to solve complex AI problems. Currently methods for combining heuristic information and massive parallelism are still largely unknown. The publication in 1985 of the Sixth Generation Computing System development proposal shows a serious intention in Japan to go beyond the current FGCS activities and address the AI aspects of computations [4].

D. Design Issues of Parallel AI Architectures

The essential issues in designing a computer system to support a given AI application can be classified into the representation level, the control level, and the processor level. The *representation* level deals with the knowledge and methods used to solve a given AI problem and the means to represent it. Design issues related to the representation level are discussed in Section II. The *control*

level is concerned with the detection of dependencies and parallelism in the algorithmic and program representations of the problem. Design issues related to the control level are presented in Section III. The processor level addresses the hardware and architectural components needed to evaluate the algorithmic and program representations. Issues related to the processor level are discussed in Section IV. Examples of issues in each level are shown in Table II.

Developing an AI architecture requires solutions to many issues in each level. Yet some of these issues are still open at this time. In this paper, we do not provide an exhaustive survey of all reported projects and their relevant issues. Instead, we discuss some important issues in the three levels and illustrate the solutions by a number of representative systems.

II. REPRESENTATION LEVEL

Since 1950, knowledge-representation schemes have been widely discussed in the literature [20], [48]. The representation level is an important element in the design process and dictates whether or not the given problem can be solved in a reasonable amount of time. Although various paradigms have been developed, most existing knowledge-representation methods and AI languages were designed for sequential computations, and the requirements of parallel processing were either not taken into account or were only secondary considerations. Moreover, many designers of AI computers start with a given language or knowledge-representation scheme; hence the representation level is already fixed. Research in designing AI computers has focused on automatic methods to detect parallelism and providing hardware support for time-consuming operations in a given representation but has not provided much to aid users in collecting and organizing knowledge or in designing efficient algorithms.

A. Domain Knowledge Representations

Domain knowledge refers to objects, events, and actions. From an implementation point of view, the criteria to evaluate a representation scheme for a multiprocessing system are its declarative power, the degree of knowledge distribution, and its structuralization.

Declarative versus Procedural Representations: The major knowledge-representation paradigms used today can roughly be classified into declarative and procedural ones, although most practical representation schemes combine features from both. Declarative representations specify static knowledge, while procedural ones specify static knowledge as well as the control information that operates on this static knowledge. Horn clauses (or even first-order logic), semantic networks, and rule-based production systems are examples of declarative representations, while Lisp programs are procedural representations. Frames combine both declarative and procedural information to represent structured knowledge. Attached to each frame is

TABLE II
EXAMPLES OF ISSUES IN DESIGNING AI COMPUTERS

Representation Level
Choosing an appropriate knowledge representation
Representing meta-knowledge
Acquiring and learning domain knowledge and meta-knowledge
Representing knowledge in a distributed fashion
Declaring parallelism in AI languages
Control Level
Analyzing data-dependencies
Synchronization
Maintaining consistency
Partitioning AI problems
Deciding granularity of parallelism
Dynamic scheduling and load balancing
Efficient search strategies
Trade-offs on using heuristic information
Predicting performance and linear scaling
Processor Level
Defining computational models
Developing methods to pass information
Designing hardware for overhead-intensive operations
Designing interconnection structure for load balancing and communication of guiding and pruning information
Managing large memory space

various heuristic information, such as a procedure on using the information in the frame.

A declarative approach allows the hiding of procedural control-flow information, thereby resulting in an easily created, modified, and understood knowledge representation. Declarative representations are referentially transparent; that is, the meaning of a whole can be derived solely from the meaning of its parts and is independent of its historical behavior. This may significantly increase program productivity because of its user orientation and user friendliness.

Declarative representations offer higher potential for parallelism than procedural ones for the same problem, because a declarative representation specifies tasks as a set, while a procedural representation may overconstrain the order of execution by the implicit order of statements. Parallel versions of procedural representations, such as parallel Lisp programs, achieve a limited amount of concurrency, while relying on programmers to specify the parallel tasks [76], [80]. However, parallelism in a declarative representation may be restricted by the implementation of the language translators. For example, interpreters for rule-based production systems can be viewed as pattern-directed procedure invocations. Although pattern matching may provide a rich source of parallelism, the match-select-act cycle is a bottleneck and restricts the potential parallelism. Less restrictions are seen in the implementation of logic programming and semantic networks. This is the key reason for the Japanese FGCS project to choose logic as the basic representation. It has also been reported that if 256000 processing units were used, the Connection Machine, using a semantic network representation, can execute four orders of magnitude faster than a sequential Lisp machine with respect to a number of object-recognition problems [59].

A disadvantage of declarative representations is that their nondeterminism is usually associated with a large search space that may partly counteract the gains of parallel processing; whereas procedural schemes allow the specification and direct interaction of facts and heuristic information, hence eliminating wasteful searches. A trade-off between the degree of parallelism and the size of the search must be made in designing a representation scheme.

Distributed Knowledge Representations: A second criterion to evaluate a representation scheme is its degree of distribution. In a local representation, each concept is stored in a distinct physical device, and each device may be shared among multiple concepts. Although this simplifies their management, the knowledge will be lost if the device fails. Most current AI systems adopt the local representations.

Recently, distributed representations have been proposed. In this scheme, a piece of knowledge is represented by a large number of units and distributed among multiple physical devices, and each device is shared among multiple knowledge entities. The resulting system is more robust because the failure of one physical device may cause some but not all information to be lost in multiple knowledge entities. Neural networks [90] and the Boltzmann Machine [89] are examples in this class. The proposed Boltzmann Machine consists of a very large network of binary-valued elements that are connected to one another by bidirectional links with real-value weights. The weight on a link represents a weak pairwise constraint between two hypotheses. A positive weight indicates that the two hypotheses tend to support one another, while a negative weight suggests that the two hypotheses should not both be accepted. The quality of a solution is then determined by the total cost of all constraints it violates.

Another interesting distributed knowledge-representation scheme, called Sparse Distributed Memory (SDM), has been proposed by Kanerva [98]. The SDM has a 1000-bit address to model a random sample of 2^{20} physical locations. Given a 1000-bit read/write address, the locations in the SDM that are within 450 bits of this address are selected associatively. Statistically, nearly 1000 memory locations will be selected. The word read is a statistical reconstruction by a majority rule. The SDM model was designed with an analogy to the human brain and can perform pattern computations such as looking up patterns similar to a given pattern and generating a pattern that is an abstraction of a given set of similar patterns [42]. Although it is much simplified with respect to the human brain, its concept may lead to a new class of computers suitable for pattern computations.

Distributed representations are generally fault-tolerant in that, within a large parallel network with a few faulty units, the remaining pattern is still usable. This property is very attractive for wafer-scale integration. The disadvantage of distributed representations is that they are hard for an outside observer to understand and modify, so automatic learning schemes must be employed. An open problem at this time is to combine local and distributed repre-

sentations by decomposing a large knowledge base into partitions and using a local representation for each.

Structuralization of Knowledge: A third criterion to evaluate knowledge-representation schemes is their structuralization; this is related to the inference time and the amount of memory space required to store the knowledge. In general, the more structured a knowledge representation is, the less inference time and the more memory space are needed. An experimental comparison of efficiency has been reported for four kinds of knowledge-representation schemes for a pilot expert system, namely, a simple production system, structured production system, frame, and logic [132]. It was found that the volume of knowledge bases for the four schemes were different. In one case, both production systems have 263 rules and 15000 characters, the frame system has 213 frames and 29000 characters, and the logic system has 348 clauses and 17000 characters. The memory space required by the frame system is the largest because some related pieces of knowledge have to be replicated in different frames. Since at most one conclusion is allowed in each Horn clause, the space of the logic system is larger than that of the production systems. The experimental results also show that, with respect to forward and backward reasoning, the frame system is the fastest, while the logic system is the slowest. The efficiency of the frame system is relatively insensitive to the size of the knowledge base because related pieces of knowledge are connected to one another by pointers, thereby limiting searches. The inference time of the simple production system is moderately sensitive to changes in the size of the knowledge base, while that of the logic system is markedly sensitive to changes in size.

Structured knowledge representations are usually desirable as long as the memory space needed is reasonable. To achieve this end, metaknowledge may be included in the knowledge base to reduce the search overhead needed. There are two problems in using metaknowledge. First, it consumes more memory space and may increase the overheads in memory management and communication. Second, metaknowledge in a poorly understood domain may be fallible and may lead the search in the wrong direction, thereby increasing the total search time. Theoretical studies and experimental comparisons are urgently needed to address this space-time trade-off.

B. Metaknowledge Representations

Metaknowledge includes the extent and origin of domain knowledge of a particular object, the reliability of certain information, the possibility that an event will occur, and the precedence constraints. In other words, metaknowledge is knowledge about domain knowledge. Metaknowledge can be considered to exist in a single level or in a hierarchy [19]. In a hierarchical form, metaknowledge is used to decide which domain-dependent actions to perform, while *meta-metaknowledge* is the control knowledge about metaknowledge. Higher level metaknowledge is common-sense knowledge known to humans.

lge base into
r each.
erion to eval-
eir structural-
d the amount
nowledge. In
resentation is,
ry space are
fficiency has
epresentation
a simple pro-
n, frame, and
of knowledge
In one case,
15000 charac-
29000 charac-
s and 17000
he frame sys-
ces of knowl-
mes. Since at
rn clause, the
of the produc-
ow that, with
ng, the frame
is the slowest.
ely insensitive
ated pieces of
by pointers,
of the simple
changes in the
logic system is

usually desir-
reasonable. To
cluded in the
head needed.
edge. First, it
ease the over-
ication. Sec-
domain may
ong direction,
oretical stud-
ly needed to

in of domain
ty of certain
l occur, and
taknowledge
taknowledge
a hierarchy
is used to
perform, while
about meta-
mmon-sense

The use of metaknowledge allows one to express the partial specification of program behavior in a declarative language, hence making programs more aesthetic, simpler to build, and easier to modify. It facilitates incremental system development; that is, one can start from a search-intensive program and incrementally add control information until a possibly search-free program is obtained. Lastly, many knowledge-representation schemes and programming paradigms, such as logic, frames, semantic networks, and object-oriented programming, can be integrated with the aid of metaknowledge [19], [64]. Metaknowledge can be classified as deterministic and statistical according to the correctness and efficiency considerations.

Deterministic Metaknowledge: Deterministic metaknowledge is related to the correct execution of the algorithm. Metaknowledge about precedence relationships results from a better understanding of the problem; this helps reduce the resource and time complexities. For instance, to solve the problem of sorting a list, it is necessary to analyze the problem, find the appropriate representation, and evaluate the necessary tasks. A list of n elements can be sorted by searching in parallel in $O(\log n!)$ average time ($= O(n \cdot \log n)$) one of the $n!$ permutations that contain the sorted elements; however, an algorithm such as Quicksort contains functionally dependent subtasks and can sort the list in $O(n \cdot \log n)$ average time using one processor. In general, the deeper we understand the problem to be solved, the larger is the set of necessary precedence constraints and the more efficient is the solution to the problem.

Many AI languages allow programmers to specify the sequence of executions in a serial computer, but the metaknowledge to specify the correct execution in a multiprocessing environment is incomplete or missing. In programs written in pure declarative languages, the static aspects of the represented knowledge are stressed, while the controls are left to the compiler/interpreter. For instance, in a logic program, a clause $a: - a_1, a_2, a_3$, means that a is implied by a_1, a_2 , and a_3 , but nothing about their functional dependencies is represented. The sequence of executions in a serial computer is correct because a definite search order is imposed, but the precedence relationships among subgoals are unknown to the scheduler in a multiprocessor.

In a number of AI languages such as Prolog, the type and meaning of variables and functions are dynamic and query dependent and cannot be completely specified at compile time. To use metaknowledge in this regard, the semantic meaning of subgoals and operations can be specified, which can be interpreted as precedence relationships by the scheduler at run time. In logic programming, the method to represent semantic information in a general and efficient way is still open.

The metarules used must be sufficient and precise such that all precedence relationships can be derived unambiguously and easily. An important consideration is the scope within which metarules can be applied. Common-sense metarules should be included to operate on more specific metarules specified by the programmers. Using the

metarules, the interpreter/compiler generates the necessary synchronization primitives.

Several researchers have addressed this problem. Gallaire and Lasserre used metaknowledge expressed as a general or special control strategy in a Prologlike interpreter [63]. In their approach, metaknowledge is made explicit through metarules, each of which describes an action to be undertaken by the interpreter whenever the interpreter focuses its attention on an object involved in the metarule. In LP, a Prolog equation-solver learning system [154], control information is expressed in a declarative representation, and inference is performed at the meta-level. Search at the object level is replaced by search at the "meta" level. Research is necessary to provide a practical method to specify unambiguously the needed synchronization through metaknowledge.

Statistical Metaknowledge: Statistical metaknowledge can be used to enhance the computational efficiency of an AI program. Warren used a simple heuristic and reordered only the goals of compound queries written in pure Prolog [190]; even so, he typically obtained query speedups of an order of magnitude. The probability of success of a subgoal and the associated search cost have been found to be useful in guiding the search of logic programs [69], [114]. In general, clauses in Prolog with the same head should be ordered such that those likely to succeed with a smaller expected search cost are searched first. In contrast, subgoals within a clause should be ordered such that those likely to fail with a smaller expected search cost are searched first.

In many expert systems, the *belief* and other measurements of accuracy of the information have been widely used. For example, in MYCIN, the *confidence factor* (CF) is used to decide among alternatives during a consultation session [21]. The CF of a rule is a measurement of the association between premises and actions. A positive CF indicates that the evidence confirms the hypothesis, while a negative CF indicates disconfirming evidence.

The representation of metaknowledge about uncertainty is an active topic in AI research. Several methods, such as fuzzy logic and Dempster-Shafer theory, are being studied currently. The proper choice is still unclear.

C. AI Languages and Programming

Conventional imperative languages are inefficient and complex to program for symbolic and pattern processing; hence the design of AI programming languages has had a central role in the history of AI research. Frequently, new ideas in AI were accompanied by a new language that was natural for expressing the ideas.

To enhance programmer productivity and take full advantage of parallel processing, declarative languages have been designed for AI programming. Function-, logic-, and object-oriented languages are the major programming paradigms today. Lisp is an early and widely used functional language; it is characterized by symbolic computations, representation of information by lists, and recursion

as the only control mechanism. Numerous imperative features have been incorporated into different dialects of Lisp, so most Lisp programs are not actually declarative, but a large enough subset allows declarative programming to be done.

Hybrids of programming paradigms have been developed. One simple approach to combining features from two languages is to provide an interface between the two. Examples include Loglisp [142], Funlog [39], and Oil [36]. The provision of features from multiple languages within a single unified framework, such as Lambda Prolog, has also been proposed. A different approach called *narrowing* involves replacing pattern matching in functional languages by unification [140]. Logic programs can then be expressed as functions. Recently, three commercial programming tools Kee, Art, and Loops have been introduced, which provide a mechanism to allow multiple paradigms to be used in a program.

New AI languages feature large declarative power, symbolic processing constructs, representing information by lists, and using recursion as the only control mechanism. These languages differ in their expressive power, their ease of implementation, their ability to specify parallelism, and their ability to include heuristic knowledge. A language-oriented AI computer will inherit all the features and limitations of the language it implements. Note that no single paradigm is appropriate for all problems, as one language may be more "natural" than another, depending on the requirements and the personal view. Hence intelligent systems should allow multiple styles, including function-, object-, and logic-oriented paradigms.

Expressive Power versus Ease of Implementation: Functional languages, such as pure Lisp [122], Backus' FP [13], Hope [14], and Val [123], share many features with logic languages, including the declarative nature, reliance on recursion, and potential for execution parallelism. Yet they have vital individual features as well. First, in functional programs, input and output variables are fixed, while in logic programs, the modes of variables are query dependent. For example, the statement $z = \text{plus}(x, y)$ in a functional program implies that x and y are inputs and z is output. In contrast, in a logic program, the goal sum (X, Y, Z) has eight possible combinations of modes of variables X , Y , and Z . For instance $(\text{in}, \text{out}, \text{in})$ means that $Y = Z - X$. Second, in a functional program, only constant and constructor functions can appear in the output; while in a logic program, logic variables can be used as output. Third, pure functional programs are deterministic, and no search is needed, while logic programs are inherently nondeterministic and require searches. Finally, functional programming provides the ability to write high-order functions; that is, a function can be passed as an argument. In contrast, Prolog is a first-order language, although some logic programming languages are not.

The first three properties, especially the nondirectionality, make logic languages more expressive in the sense that a single logic program corresponds to several functional programs. Moreover logic and functional programs are

executed using resolution and reduction (or term rewriting), respectively. Note that resolution can use input information implicit in the patterns to cut down the size of the set to be examined. For example, to solve the append subgoal, $\text{append}([P], [Q, R], [1, 2, 3])$, resolution makes no distinction between inputs and outputs and uses the input information (length of the lists) to select the appropriate clauses and produce bindings for the variables involved. However, in the corresponding functional formulation $([P], [Q, R]) = \text{split}([1, 2, 3])$, all possible splits of $[1, 2, 3]$ are produced, and the one that splits the list into $[P]$ and $[Q, R]$ will be selected. The previous example illustrates that reduction can lead to overcomputation as compared to resolution.

The crucial disadvantage of functional programming lies in the difficulty to represent the inherent nondeterminism in AI problems. Although the recursive formulation and the leftmost-outermost reduction of functional programs enable depth-first searches naturally, it is difficult to write a heuristic search program by a pure functional language since heuristic searches are inherently history-sensitive. In fact, best-first-search programs written in Lisp include a lot of "setq" and "prog" statements, which are not pure functional primitives [195]. Due to their less expressive power for representing nondeterminism and their inefficiency in dealing with large data structures, pure functional languages are unsuitable for general AI applications.

Although logic languages are more expressive, their implementations, especially in a parallel processing environment, are more difficult due to the nondirectionality of variables. The dynamic nature of modes requires run-time analysis. In contrast, the run-time behavior of functional programs is much simpler to control than that of logic programs, particular in a parallel context. Techniques such as graph reduction and data flow have been developed for the parallel evaluation of functional languages. Further, Lisp has only a few primitive operators and provides unique list structures to compound data objects. These features simplify the implementation of Lisp compilers/interpreters. In fact, Scheme, a dialect of Lisp, has been implemented in a single chip [166]. The implementation, however, may be complicated by the dynamic nature and primitives with side-effects introduced in practical functional languages. Dynamic features, such as random accesses to linked lists, garbage collection, frequent function calls, and dynamic binding of functions, incur extensive run-time overheads.

Obviously, it would be advantageous if the simple controls of functional languages could be implemented in the more expressive logic languages. Considerable efforts have been devoted to combining functional and logic programming [39]. One approach to simplifying logic languages is to introduce directionality of modes of variables [140]. This method degrades its expressive power to that of first-order functional languages. Others attempt to extend functional languages to achieve the expressive power of logic languages but retain most of the underlying functional simplicity. An example is Hope with unification [34]. Unfortunately, up to now, no language exists that has

rewriting),
informa-
of the set
d subgoal,
no distinc-
input infor-
ate clauses
However,
P], [Q, R])
produced,
R] will be
reduction
solution.
mming lies
eterminism
lation and
l programs
ult to write
al language
ensitive. In
o include a
re not pure
expressive
their ineffi-
pure func-
pplications.
ve, their im-
ng environ-
nality of
res run-time
f functional
hat of logic
niques such
developed for
es. Further,
d provides
ects. These
mpilers/in-
o, has been
ementation,
nature and
ctical func-
random acc-
nt function
ir extensive
simple con-
nted in the
efforts have
e program-
anguages is
bles [140].
to that of
t to extend
e power of
ylying func-
unification
sts that has

good expressive power while being flexible enough for parallel execution. Efforts are needed in this direction.

Specification of Parallelism: Since parallel processing was not a consideration when most existing AI languages, such as Lisp and Prolog, were designed, the precedence restrictions implicit in a sequential execution order cannot be detected easily in a parallel execution. To extend these languages in a parallel processing environment, explicit primitives may have to be included.

In a pure functional language (data-flow language), the meaning of an expression is independent of the history of computations performed prior to the evaluation of this expression. Precedence restrictions occur as a result of function application. Notions such as side effects do not exist, hence all arguments and distinct elements in a dynamically created structure can be evaluated concurrently. For example, to compute the average of numbers in a list s , $(1.(2.(3.nil)))$, using the function $average(s) = div(sum(s), count(s))$, the computations of $sum(1.(2.(3.nil)))$ and $count(1.(2.(3.nil)))$ can proceed concurrently. It has been reported that implementations of functional languages on parallel computers seems easier than that on sequential ones [33].

Note that Lisp and many of its dialects are not pure functional languages. Referential transparency is lost in most Lisp languages due to side effects. The precedence restrictions are represented not only in function calls but also in procedures.

Several parallel Lisp languages have been proposed and implemented. Multilisp, developed by Halstead, has been implemented on a 128-processor Butterfly computer. Concurrency in Multilisp can be specified by means of the *pcall* and *future* constructs [77]. *pcall* embodies an implicit fork join. For example, $(pcall\ ABC)$ results in the concurrent evaluation of expressions A , B , and C . The form $(future\ X)$ immediately returns *future* (a pseudo value) for X and creates a task to evaluate X concurrently, hence allowing concurrency between the computation and the use of X . When the evaluation of X yields a value, it replaces the *future*. The *future* construct is good in expressing mandatory parallelism but is quite expensive in the current Multilisp implementation.

Another parallel Lisp language, Concurrent Lisp [163], is extended from Lisp 1.5 and has three additional primitive functions to specify concurrency: STARTEVAL for process activation, and CR (critical region function), and CCR (conditional critical region function) for mutual exclusion. A multiprocessing program written in Concurrent Lisp is a set of cooperating sequential processes, each of which evaluates its given form. Similar to P/V primitives, CR and CCR have enough power to express process interactions.

In Parlog, a parallel logic programming language [27], every argument has a mode declaration that states whether the argument is input (?) or output (^). For example, in the following statements, mode $merge(??)$.

$$merge([U|X], Y, [U|Z]) \leftarrow merge(X, Y, Z).$$

the first two lists are merged to form the result. In Concurrent Prolog [150], a read-only annotation (?) is used. For example,

$$merge([U|X], Y, [U|Z]) \leftarrow merge(X?, Y, Z).$$

indicates that X must have a value before $merge(X?, Y, Z)$ can be invoked. Another way to specify the concurrency is to use different symbols to distinguish between "parallel AND" and "sequential AND" such as "||" and "&" in Parlog. Guarded clauses are used in Parlog and Concurrent Prolog partly to specify parallelism. A guarded clause has a format, $h:-g|b..$, where g is the *guard* of the clause and b is its *body*. Subgoals in the body can only be evaluated when all subgoals in the guard have succeeded, and values bound have been committed to the body.

Clearly, the previous approach of specifying parallelism by users detracts from the objective of declarative programming, which separates logic from control, or "what" from "how." Both mode declarations in Parlog and read-only annotations in Concurrent Prolog impose a fixed execution order on subgoals, which may be inefficient in parallel processing. On the other hand, distinguishing the guard from the body cannot completely specify the precedence relationships because subgoals in the guard and body may be dependent. The use of guards is also complicated by a lack of general methodology to select subgoals in the guard. Moreover, precedence relationships are a partial order, so the distinction between "sequential AND" and "parallel AND," which are linear orders, is insufficient to specify all precedence relationships. Lastly, owing to the nondeterministic behavior of AI programs, users cannot always specify the parallelism perfectly. A desirable parallel AI language should allow its compiler to detect the parallelism and schedule parallel executions as efficiently as possible.

Object-Oriented Languages: Object-oriented programming holds promise as a framework for concurrent programming that can be extended to data bases and knowledge bases. It is expected that "object-oriented programming will be in the 1980's what structured programming was in the 1970's" [141]. A variety of object-oriented languages include Smalltalk [68], Loops [161], Actor [6], CommonObjects [158], and many others [191]. Recently, CommonLoops was suggested as a standard for object-oriented extensions to Common Lisp by the Lisp community [16].

Object-oriented programming has been used to express different concepts, but the concept of an object is the common feature in these languages. Objects are entities that combine the properties of procedures and data. Object-oriented programming replaces the conventional operator-operand concept by messages and objects. All actions in an object-oriented program result from sending messages among objects. A selector in the message specifies the operation to be performed. An object responds to messages using its own procedures (called *methods*) for performing operations. Message sending supports data abstractions, a concept that is necessary but not sufficient for

the language to be object-oriented. Object-oriented languages must additionally support the management of data abstractions using abstract data types and the composition of abstract data types through *inheritance*. Inheritance is used to define objects that are almost like other objects. In fact, object-oriented programming should be characterized by the nature of its type mechanisms rather than the nature of its communication mechanisms; that is, object-oriented programming can be defined as

object-oriented = data abstraction + data types
+ type inheritance.

Object-oriented programming is a paradigm for organizing knowledge domains while allowing communications. Concurrent models, operating systems, and coordination tools are built from low-level objects, such as processes, queues, and semaphores. Hewitt's Actor model is a formalization of the ideas of object-oriented languages: an actor in his model is the analogue of a class or type instance but considers the added effects of parallelism [83]. Computations in the Actor model are partial orders of inherently parallel events having no assignment commands. The language Act3, based on the Actor model, combines the advantages of both object-oriented and functional programming [5]. To support object-oriented programming, appropriate objects representing data structures should exist at the hardware level as objects of "machine data-structure type." This gives birth to the data-type architecture [67]. The Apiary network architecture is based on the Actor model [84], [85].

D. Summary

A major problem in the representation level lies in the large amount of knowledge needed to define a good representation and the imprecise nature of this knowledge. Efforts have been directed toward the automatic acquisition of domain knowledge and metaknowledge to lead to a good representation and the design of a language that is more expressive and yet easy to implement in a parallel processing environment. The design of a systematic method to generate alternate representations is particularly desirable. The methodology should start with the problem specification, use automated tools to transform the problem specifications into problem representations, compare alternate representations, and use metaknowledge to guide the generation of different representations.

III. CONTROL LEVEL

There are four basic issues in the control level of computer-system design. Maintaining consistency of knowledge is important, as incomplete and inconsistent knowledge is often dealt with in AI computations. As multiprocessing is widely used in AI computations, related issues include the decomposition of a problem (or program) into subproblems, the synchronization of cooperating processes, and the scheduling of processes for efficient execution. Although the design issues in the control level are similar to

those in traditional multiprocessing systems, AI problems often start with different representations, hence their solutions in the control level may be very different from traditional ones.

A. Consistency Maintenance

Traditional logic is *monotonic* because new axioms are only added to the list of provable theorems and never cause any to be withdrawn. However, knowledge-based systems on changing real-world domains have to cope with the maintenance of consistent deduction. Classical symbolic logic lacks the tools to deal with inconsistencies caused by new information. *Nonmonotonic reasoning* has been developed to deal with this problem [194].

Early attempts at consistency maintenance evolved around explicit manipulation of statements. The major system developed was Strips, which dealt with the manipulation of blocks of various sizes, shapers, colors, and locations by a robot [56]. In Strips, the entire data base is searched for inconsistencies when the robot moves a block. System applied inference refers to a system in which the architecture provides a mechanism to automatically maintain the consistency of the data base. The widely publicized system of this nature was Microplanner [165]. In Microplanner, the operators of Strips are replaced by "theorems." There is no automatic inference mechanism, and the programmer is required to encode all possible implications of a theorem. An improvement to Strips is Doyle's Truth Maintenance System (TMS) in which the reasons for beliefs are recorded and maintained, and these beliefs can be revised when discoveries contradict assumptions [47]. To attach a justification to a fact, a TMS is designed with a goal that efficiently links consequences and their underlying assumptions. In TMS, each relation has an associated IN and OUT nodes. The statement at this node is true if the statements in the IN list are known to be true and the statements in the OUT list are not true.

A different approach to consistency maintenance was adopted in designing the IBM Yes/MVS expert system that operates on a System 370 computer under the MVS operating system [149]. This expert system is used to schedule a real-time system in which contradiction occurs between the changed facts and the previous consequences. The system removes inconsistent deductions and computes new consequences in accordance with the changed facts. The consistency maintenance mechanism has three parts: recognition of inconsistencies, modification of the resultant state to remove inconsistencies and deduce consistent consequence, and hidden control to ensure that all inconsistencies are detected and corrected properly.

Experience on the design of Yes/MVS shows a pitfall in which correcting an inconsistency may cause another inconsistency, which in the process of being corrected reintroduces the first inconsistency. It was also found that knowledge represented in a style for consistency maintenance turned out to be quite modular, and maintaining it has been easier than initially expected.

stems, AI problems
ns, hence their solu-
very different from

Nonmonotonic logic has been demonstrated to be feasible but inefficient to implement in a large system. To allow the system to be used in real time, hardware support has to be provided on the time-consuming operations. Fundamental operations such as standard data base functions may have to be implemented in hardware. The management of a virtual memory system to support frequent additions and deletions in a TMS is an important design issue. The maintenance of the appropriate storage organization such that locality is maintained among relations affecting each other is a nontrivial problem. Finally, parallel processing may introduce additional problems of consistency; efficient parallel architectures to process concurrent queries have to be investigated.

B. Partitioning

In parallel computations, determining the granularity or the minimum size of a subproblem that should be computed by a single processor depends on the inherent parallelism in the problem to be solved. Partitioning can be implemented in different levels. In the higher levels, a complex AI problem is partitioned into several functional tasks, each of which is processed by a functionally distributed computer system. In the lower levels, the control graph of the program is partitioned into atomic operations, each of which can be processed independently.

Partitioning can be performed by users at design time or compilers at compile time or schedulers at run time. In the first method, programmers use a parallel language to specify and partition problems. These languages can define parallel tasks and the associated data communications. Design issues of parallel languages were discussed in Section II-C. In this section, we discuss static and dynamic partitioning.

Inherent Parallelism and Granularity: The proper granularity of parallelism should be determined from the inherent parallelism in the problem and the communication overheads involved in synchronization and scheduling. In general, finding the optimal granularity is difficult; however, the degree of parallelism inherent in the problem may provide useful information to guide the design of the architecture.

An example to illustrate the choice of the proper granularity is shown in the design of parallel rule-based systems. Forgy *et al.* observed that each OPS-5 production, when it fires, manipulates a few (usually two or three) working memory elements and affects only a small number (20-30) of productions [60]. According to this analysis, it appears that only limited speedups are available and that massive parallelism may not be needed. To improve the degree of parallelism, further efforts should be devoted to a) investigating parallel match algorithms, b) designing efficient partitioning strategies, and c) developing techniques to rewrite sequential OPS5 programs into versions more suitable for parallel processing.

Gupta estimated that the hardware utilization will be around two percent if the Rete match algorithm is mapped directly onto the Dado architecture [74]. He recommended

partitioning OPS5 production rules into 32 subsets to exploit the modest amount of production-level parallelism.

Based on Gupta's algorithm, Hillyer and Shaw studied the execution of production systems on the Non-Von computer, a heterogeneous system with 32 large processor elements (LPE's) and 161000 small processor elements (SPE's) [88]. Each SPE has 64 bytes of Ram to store a condition-element term. The large number of SPE's, which can be viewed as an active memory of LPE's, perform intraproduction tests in a massively associative fashion. The performance is predicted at a rate of more than 850 productions fired per second using hardware comparable in cost to a VAX 11/780. This shows that two orders of magnitude of speedup is achievable by properly partitioning production systems.

The partitioning algorithm used may have significant effects on performance. If a majority of node activations occur within a single partition, then the performance will not be good. Some researchers have reported heuristics for partitioning production systems, such as assigning productions that are sensitive to the same context, goal, or task to different processors in a round-robin fashion. However, preliminary results have shown that these strategies do not bring significant improvement as compared to random partitioning [134]. Intelligent partitioning strategies, using knowledge previously known, remain to be developed.

In a multiprocessing system, it is hoped that equal-sized tasks are distributed evenly to all processing units. The above example, however, has shown that this may be impractical because the problems to be solved may have irregularly structured control- and data-flow graphs and data-dependent workloads. In practice, efficient heuristic methods may have to be used to partition the task graph into granules that can be executed in parallel. Important related issues to be studied in this case are the design of heterogeneous architectures and the dynamic distribution of workload.

Compiler Detection of Parallelism: Based on the data dependencies in a program, a compiler may be able to detect the parallel modules in it and partition the program at compile time. An example is the post-compiler of Faim-1, called an *allocator*, that performs data-flow analysis on the procedural code and inference connectivity analysis on the logic behavior to statically distribute the fragments to the processing elements [11]. Similar work has been done on partitioning programs for numeric applications [108].

Detection of parallelism in logic programs has centered on detecting AND parallelism and OR parallelism. AND parallelism in logic programs involves the simultaneous execution of subgoals in a clause. Due to shared variables, concurrent execution of two or more subgoals in a clause may result in binding conflicts. The detection of AND parallelism is based on the analysis of input-output modes of arguments in a subgoal. The input and output variables in a logic program denote the direction of binding transfers during unification, in a way similar to the input and output arguments in procedure calls. However, an argument in a logic program can be in the input mode in one instance and in the output mode in another, or may

remain unbound. This dynamic behavior prohibits a complete static analysis. Previous research, therefore, developed methods either to provide primitives for users to specify the modes or to assign modes automatically to arguments that can be analyzed at compile time and leave the rest to be resolved at run time. Automatic detection of AND parallelism at compile time can be classified into two types.

a) *Detection of restricted AND-parallelism:* DeGroot proposed a typing algorithm to detect restricted AND parallelism [38]. The essential concept is to monitor all potentially executable subgoals and ensure that no two subgoals will share one or more unbound variables if they execute in parallel. A term in a clause can be in one of three types: 1) grounded (or constant), 2) nongrounded nonvariable (an input variable), or 3) variable (an uninstantiated variable). To lower the run-time overhead of checking the contents of terms, a partial check is made at compile time; only terms of type 1 and that of type 3 with different variable names are detected to be independent. All other possibilities remain to be detected at run time. A consequence of this partial check is that a term may occasionally be typed too strongly.

b) *Detection of coupled data-dependencies:* Chang, Despain, and DeGroot updated the above typing algorithm by testing for coupled data dependencies at compile time to reduce the run-time overhead [24], [40]. In this scheme, variables in a clause are classified into three groups: grounded, coupled, and independent. (An independent variable is neither grounded nor coupled with other variables.) Two terms are said to be coupled if they share at least a common unbound variable. Two variables are in the same coupled group if the compiler detects that there is a chance for them to be coupled. To find the group a variable belongs to, a programmer has to supply the activation mode of the query and the entry points of the program. The compiler then classifies the variables in the subgoals from left to right and derives the execution graphs and backtracking based on the worst case activation mode of each variable. Multiple execution graphs may be generated at compile time and the appropriate one selected at run time. This scheme has been adopted in the PLM of the Aquarius project [43].

Other heuristic methods of checking types at compile time are also possible. Tung and Moldovan have also investigated a number of heuristics to infer the modes of a given variable and mark all possible input-output modes of arguments in the clauses [178].

Compiler detection of parallelism has the advantages of reduced run-time overhead and programming efforts. Its disadvantage is that it may not be able to detect all the inherent parallelism in a highly expressive AI language and may have to be combined with user declaration and dynamic detection. The restrictions of compiler detection are briefly summarized below.

Special cases: The extraction of parallelism from data-dependency analysis is based on the assumption that if two subgoals do not share any unbound variable, then they can be executed concurrently. This assumption is not

true in some special features of the language, such as outputs in Prolog. A solution to this problem is proposed by DeGroot [41].

Procedural dependencies: A procedural dependency exists between two subgoals if their execution order is fixed by their semantics. For example, in the following clause,

$$a(X) : - \text{test_for_ok}(X), \text{work_on}(X)$$

the subgoal "test_for_ok(X)" must be executed first. Note that the subgoals in this example cannot be executed concurrently even if X is grounded, because the second subgoal may contain meaningless, inaccurate, or unbound work unless the first subgoal is true. In declarative languages such as Prolog, it is difficult to specify the semantics of subgoals without specifying its explicit control for parallelism. A solution to this problem is proposed by DeGroot [41].

Exponential complexity: It may be difficult to define all possible combinations of modes at compile time as they grow exponentially with the number of potential output variables.

Dynamic Detection of Parallelism: Many data dependencies in a highly expressive AI language cannot be resolved until run time. For example, a subgoal $p(X, Y)$ in a logic program may be called as $p(X, X)$, which is a coupling dependency on a query with coupled terms introduced at run time. This dependency cannot be detected at compile time. Due to the dynamic nature of AI computations, an AI computer should provide a mechanism to map the program and data onto hardware dynamically.

In general, the computational model can be represented as a token-flow graph with four kinds of nodes: and-decomposition, or-decomposition, and-join, and or-join. The tokens passed along the edges can be demand tokens, data tokens, or control tokens. Conery and Kibler described an AND/OR process system based on a producer-consumer model that dynamically monitors variables and continually develops data-dependency networks to control the order of execution of subgoals, never allowing two potential procedures with the same variable to be executed in parallel [31]. An ordering algorithm, called a connection rule, is used dynamically to determine a generator for each unbound variable. When a subgoal is completed, it is checked to ensure that it did produce all variable bindings it was supposed to; otherwise, the ordering algorithm is evaluated again. Improvements were made to the above scheme to reduce further the run-time overhead and extract more parallelism [103], [118].

Since dynamic partitioning must be repeatedly executed at run time, it may reduce the performance gains and could even produce negative gains. The trade-off between static partitioning by an intelligent compiler and dynamic partitioning by a sophisticated operating system is an important issue to be addressed in parallel AI processing. Dynamic partitioning is closely related to dynamic scheduling, and related issues will be discussed in a subsequent section.

Bottleneck Analysis: An important issue in partitioning is to decompose the problem evenly, so bottlenecks in performance do not exist. It is easy to see that if a bottleneck requires a fraction of the total computations, then the speedup cannot be more than the reciprocal of this fraction, regardless of how the rest is partitioned. It is well-known that the performance bottleneck of an application executing on a vector computer is its scalar code. Similarly, the performance bottleneck of a parallel AI computation is its sequential part (sequential inference or I/O). An important problem is to find the bottleneck in the problem to be solved.

Experience with designing the Fido vision system at Carnegie-Mellon University has shown that an unbalanced partitioning algorithm can substantially degrade the performance [104]. Adding Warp, a systolic system with peak processing rate of 100 mflops, to a host (a Sun computer and three "standalone processors") seems only to double or triple the speed of the Fido loop. This means that Warp is definitely underutilized; functions on the standalone processors, either in preprocessing or post-processing in using the Warp array, take up a substantial amount of time. It is expected that proper partitioning of vision algorithms will improve its performance significantly.

C. Synchronization

Synchronization refers to the control of deterministic aspects of computations, while scheduling handles mainly the nondeterministic aspects. The objective of synchronization is to guarantee the correctness of parallel computations such that the results of execution in parallel are the same as those of a sequential execution. That is, the parallel execution is serializable. In some nondeterministic problems, the generation of the same set of results as a sequential execution may not be necessary. For example, a user may wish to obtain a small subset of answers from a large set; the particular answers obtained do not have to be the same in the serial and parallel cases. In this case requirements on synchronization can be relaxed in parallel processing.

Many synchronization primitives used in AI processing are the same as those used in conventional computers. Examples include semaphores, test-and-set, full/empty bits, fetch-and-add, and synchronization-keys. Additionally, new or extended concepts related to synchronization have been introduced by AI researchers, such as the blackboard and actors. In this section, we will survey the synchronization of AI computations in the control and data levels and mechanisms using shared memory and message passing.

Two Levels of Synchronization: In procedural languages, if a statement precedes another statement in the program, the implication is that this statement should be executed before the second statement if the two statements share common variables; that is, control-level synchronization is implicit when data-level synchronization is needed. This

implicit execution order may overspecify the necessary precedence constraints in the problem.

On the other hand, if the tasks are specified as a set using a declarative language, then control-level synchronization is absent, and they can be processed concurrently if they do not share common variables. If they have common variables but are semantically independent, then they can be processed sequentially in an arbitrary order to maintain data-level synchronization.

The difficulty of specifying control-level synchronization when tasks are semantically dependent is a major problem in declarative languages such as Prolog. For example, the decomposition of a set into two subsets in Quicksort must be performed before the subsets are sorted. Hence the tasks for decomposition and for sorting are both semantically and data dependent. To overcome this problem, programmers are provided with additional tools, such as specifying the input/output modes of variables in a Prolog program, to specify control-level synchronization. These primitives may have side effects and may not be able to specify completely all control-level synchronization in all situations. These problems may have to be dealt with at run time until sufficient information is available.

In general, process activations and deactivations can be considered as control-level synchronization, while passing arguments in procedure calls can be considered as data-level synchronization. Both methods can be implemented through a shared memory or by message transfers.

Shared Memory: In tightly coupled multiprocessor systems, synchronization is done through a shared memory. Examples of such existing and proposed AI computers include Aquarius [43], Concurrent Lisp machine [164], Concert Multilisp machine [79], and Parallel Inference Engine [70]. In what follows, we will discuss synchronization using blackboards and show methods using shared variables in logic programs.

a) Blackboard: Historically, the blackboard model was developed for abstracting features of the Hearsay-II speech-understanding system [50]. The model is usually viewed as a problem-solving framework; however, we discuss only its control aspect here. The model consists of three major components: a knowledge source, a blackboard data structure, and control. The knowledge to solve the problem is partitioned into knowledge sources that are kept independently. The data needed to solve the problem concerned include input data, partial solutions, alternatives, and final solutions, which are kept in a global data base, the blackboard. The blackboard can be divided into multiple blackboard panels that correspond to the hierarchy of solution space. Knowledge sources result in changes in the blackboard, which lead to a solution to the problem. Communications and interactions among knowledge sources take place solely through the blackboard. A monitor is needed to ensure that no more than one knowledge source can change the blackboard at one time. There are a set of control modules that monitor changes in the blackboard and decide the appropriate action to take next. The sequence of knowledge-source invocations is dynamic.

The blackboard model provides a useful framework for diverse types of knowledge to cooperate in solving a problem and has been used to many AI applications. Its implementation is similar to that of a critical section in operating systems. In the pure model, the solution is built one step at a time. Currently, extensive research on concurrent access to blackboards is conducted.

Hayes-Roth has proposed a more powerful *blackboard control architecture* in which control information (meta-knowledge) is also stored and updated on a separated control blackboard [82]. This approach adapts to complex control plans as a whole. Operational strategies, heuristic, and scheduling rules can change repeatedly in the course of problem-solving.

b) *Synchronization via shared memory variables*: Although Lisp contains a "pure function" subset, it also supports many functions with side effects, such as *rplaca*, *rplacd*, *set*, and input/output functions. These side effects result from procedural dependencies and global (or free) variables and resemble problems in conventional parallel languages. In fact, some shared-memory multiprocessors, such as Concert and Butterfly, support both Multilisp, Simultaneous Pascal, and other parallel languages [79]. Multilisp provides a simple method to wait for values generated in the future. However, as in other languages, procedure activations in Multilisp may not be well nested, and an activation can terminate before another activation it contains. This exception-handling problem has to be addressed in programming the system [78].

Pure Prolog is a single-assignment language. Under this restriction, the distinction between a shared-memory variable and a communication channel vanishes. Since a logic variable is not allowed to be rewritten through side effects, conventional hardware-synchronization mechanisms, such as test-and-set, full/empty-bit method and fetch-and-add, are no longer needed in multiprocessing of pure logic programs [119]. The popular strategy taken now is to provide the programmer with a mechanism to delay process reduction until enough information is available so that a correct decision can be made. Currently, the Concurrent Prolog group is concentrating their efforts on Flat Concurrent Prolog, a subset of Concurrent Prolog. In Guarded Horn Clause (GHC) [181], ICOT's current choice for Kernel Language 1, OR parallelism was eliminated from Concurrent Prolog, and a strict synchronization rule that suspends a subgoal if it tries to write in the parent environment is adhered. This rule made the read-only annotation somewhat superfluous. Although it simplifies the implementation of GHC, some expressive power is lost due to a weaker notion of unification [168].

c) *Joins*: As similar to conventional *fork-join* primitives, static joins can be used for synchronization in parallel AI processing. For example, in multiprocessing of logic programs, a parent node can activate its children in parallel, and each child begins producing all possible answers. The parent waits for each child to complete, collects their answers, computes the "join" of their answers, and passes the entire set of results as its answer. This approach

uncovers the greatest AND parallelism in a logic program but is efficient only if the program consists mostly of deterministic procedures and clauses; that is, most variables have only a single binding. For nondeterministic AI problems, joins are impractical because the nondeterminism increases the uncertainty whether a given AND node should be evaluated. Note that if joins are computed dynamically, that is, a parent node collects separate answers from each child as they are produced, then the data-level synchronization employed forms a pipelined computation called dynamic joins. This scheme will be discussed later with respect to synchronization in semantic networks.

Message Passing: In passing messages, a communication channel between the sender and receiver processes is required. Synchronization via messages can be achieved through software protocols or specialized hardware. Many existing and proposed AI computers pass around messages of arbitrary complexity and perform complex operations on them. The computing elements are complex, and the communication costs are high. Alternatives to passing messages are discussed in this section.

a) *Message passing in production systems*: Reasoning using forward chaining in production systems has different behavior from reasoning using backward chaining. The behavior in forward chaining is illustrated in OPSS, whose interpreter repeatedly executes a match-select-act cycle. In the match phase, all rules whose conditions are satisfied by the current content of the working memory are selected. This set is called the conflict set. In the select phase, conflict resolution is performed to select one of the productions in the conflict set. In the act phase, the working memory is modified according to the action part of the selected rule. Although the three phases can overlap in a multiprocessing environment, synchronization must be performed to ensure that the result is consistent with that of a sequential execution; that is, all changes in the conflict set must be known prior to the completion of conflict resolution in the next cycle.

Synchronization in the efficient Rete interpreter for OPS5 is based on a data-flow graph, which can be viewed as a collection of tests that progressively determine the productions ready to fire. Inputs to the graph consist of changes to the working memory encoded in tokens. Output tokens specify changes that must be made to the conflict set. Tokens are sent via messages in a multiprocessing system.

b) *Marker passing and value passing*: Marker passing has been studied as an alternative to message passing. In such systems, communications among processors are in the form of single-bit markers. An important characteristic is that there is never any contention: if many copies of the same marker arrive at a node at once, they are simply OR'ed together. The order of markers to be passed is determined by an external host.

Marker passing is suitable for systems implementing semantic networks. Nodes in the semantic network are mapped to processors in the system. An example of such a

system is NETL [51]. A basic inference operation in semantic networks is set intersection. Analogous to dynamic joins in data bases, set intersections are implemented using data-level synchronization. If an object with n properties is searched, then n commands are sequentially broadcast to all corresponding links, the associated nodes are marked, and the node with n markers reports its identity to the controller. Marker passing is adequate for many recognition problems; however, it may not be sufficient to handle general AI problems. The Connection Machine was originally developed to implement marker passing to retrieve data from semantic networks, but its current version has more powerful processing units that can manipulate address pointers and send arbitrary messages.

In value passing, continuous quantities or numbers are passed around the system, and simple arithmetic operations are performed on these values. Like marker-passing systems, there is no contention in value passing: if several values arrive at a node via different links, they are combined arithmetically, and only one combined value is received. In this sense, value passing systems can be considered as an analog computer. Examples of value-passing system are the Boltzmann machine [52] and other "neural" computation systems [91].

Marker-passing systems do not gracefully handle recognition problems in which the incoming features may be noisy. These problems can be better handled by value-passing system in which each connection has an associated scalar weight that represents the confidence on the incoming values. Many iterative relaxation algorithms that have been proposed for solving low-level vision and speech-understanding problems are ideally suited to value-passing architectures.

c) Object-Oriented and Actor Approaches: In the object-oriented approach, and in particular, the Actor model, an actor is a virtual computing unit defined by its behavior when messages are received. Actors communicate via point-to-point messages that are buffered by a mail system. The behavior of an actor consists of three kinds of actions: 1) communicate with specific actors of known mail addresses; 2) create new actors; and 3) specify a replacement that will accept the next message. Actor languages avoid the assignment command but allow actors to specify a replacement. Replacements can capture history-sensitive information, while allowing concurrent evaluation of data-independent expressions [6]. Message passing in actors, which can be viewed as a parameter-passing mechanism, differs from both call-by-value and call-by-reference.

D. Scheduling

Scheduling is the assignment of ready tasks to available processors. It is especially important when there is nondeterminism in the algorithm. Scheduling can be static or dynamic. Static scheduling is performed before the tasks are executed, while dynamic scheduling is carried out as the tasks are executed. The actions to be performed in scheduling include 1) determination of dependent tasks, 2)

static reordering of tasks at compile time, 3) dynamic selection of tasks at run time when free processors are available, and 4) determination of the number of processors to solve a given class of problems cost-effectively. All schedules can be considered as a search strategy based on a search tree or search graph [136].

Identifying Dependencies: Parallel scheduling of AI programs is complicated by their dynamic functional and shared-variable dependencies and the high expressive power of many AI languages. Due to high expressive power, the same program can be used to represent many different dependencies, each of which may be scheduled differently. Identifying dependencies at compile time is also difficult due to the dynamic and nondeterministic nature of executions.

If functional dependencies exist among tasks, then the scheduler must find these dependencies dynamically; if there are no functional dependencies but only shared-variable dependencies, then the scheduler has to compare the merits of all possible schedules. Both cases are not practical because of the high dynamic overhead. As discussed earlier, solutions to detect dependencies are not satisfactory at this time.

A viable approach is to identify the possible dependencies at compile time, statically order all sibling nodes in a search tree for each case, and schedule them according to a parallel depth-first strategy. A simple method was proposed by Warren [190], which orders the subgoals in a clause according to the number of possible solutions generated under the given subgoal. Our experimental simulations indicated that the worst case evaluation time resulting from this method can be worse than the case without reordering, but the best case time can be 2-30 times better. Warren's method does not consider the effects of backtracking, the possible dependencies among subgoals and clauses, and the overhead of finding the solutions. We have proposed a method to represent the effects of backtracking as an absorbing Markov chain [117]. By assuming that sibling nodes are independent, they are reordered to minimize the total expected search cost of the program. Heuristics have been developed to reorder subgoals when they are dependent and have side effects. Our preliminary simulations indicated that the performance is substantially better than that of Warren's method.

Selection Strategies: Suppose in the course of evaluating an AI program n active tasks and m processors are available, $1 \leq m < n$. The ideal scheduling algorithm should select m active tasks such that this decision will minimize the expected computational time. It is difficult to design such an optimal selection algorithm because 1) the metrics to guide the search are estimated heuristically and may be fallible, 2) the metrics may be dynamically changing during the search, and 3) problem-dependent precedence restrictions may exist that cannot be detected at compile time. As a result, unexpected anomalies may occur when parallel processing is applied.

The potential parallelism in an AI computation can be classified into two types: deterministic parallelism and

TABLE III
SELECTING THE m SMALLEST NUMBERS FROM n NUMBERS

Approach	Time Complexity in Each Iteration	Space/Hardware Complexity for Selection	Accuracy of Selection
Multistage selection network	$O(\log m \cdot \log n)$	$O(n \cdot \log^2 m)$	1.0
Single-stage network	$O(m)$	$O(n)$	1.0
No-wait policy	$O(1)$	$O(m)$	0.63

nondeterministic parallelism. Deterministic parallelism refers to the concurrent execution of two or more units of computations, all of which are necessary for the completion of the given job. The computational units can be tasks, processes, and/or instructions. Since all units of computation, which are performed concurrently, have AND relations, this kind of parallelism is traditionally called AND-parallelism. Nondeterministic parallelism refers to the search of multiple potential solutions in parallel. Since all potential solutions have OR relations, this kind of parallelism is traditionally called OR parallelism.

Although AND parallelism is treated as deterministic and OR parallelism as nondeterministic in conventional studies, the selection of descendants of an AND task to evaluate is also nondeterministic, as the aim is to select one that fails as soon as possible. Hence scheduling is important for tasks that are nondeterministic but may not be specific with respect to AND or OR parallelism.

In nondeterministic searches, heuristic information to guide the scheduler in selecting nondeterministic tasks is more important than the design of parallel processors, as the number of processors is almost always smaller than the number of processable tasks.

As an example, in selecting nodes to evaluate in a branch-and-bound search tree, which is an OR tree with lower bound values to guide the search, the problem is reduced to finding the m smallest numbers from n numbers. Table III shows the results obtained by three architectural approaches. In the first approach, a multistage selection network was designed to perform the selection exactly [185]. In the second approach, a single-stage ring network was used to shuffle the nodes until a complete selection was obtained [186]. In the third approach, a no-wait policy was applied. It was recognized that the heuristic information to guide the search might not always be accurate. Hence the "most promising" task in local memory was always evaluated in each cycle, while the fetch of the "more promising" tasks from other processors was initiated. It was found that, on the average, a minimum of 63 percent of the desirable tasks to be selected were selected by the no-wait policy without any additional overhead on selection, assuming that the m most promising tasks were randomly distributed among the processors [186], [187].

The management of the large memory space to store the heuristic information and the large number of intermediate

nodes in the search tree is another difficult problem to solve. A trade-off must be made to decide for a given amount of heuristic information and a given architectural model whether the amount of heuristic information should be increased or decreased, and how effective should the new heuristic information be.

The memory space required to store enough heuristic information to avoid backtracking is often prohibitive. For example, assume that all solution trees of a complete binary AND/OR tree with n levels are equally likely. The leaves are assumed to be OR nodes and are at level 0, while the root is an AND node and is a level n . We have that $f(n)$, the total number of solution trees, satisfy the following recurrence.

$$f(n) = \begin{cases} 1, & n = 0 \text{ or } n = 1 \\ 4f^2(n-2), & n \geq 2 \end{cases}$$

$$= 2^{2(2^{n/2}-1)}.$$

For $n=0$, there is only one node, hence there is one solution tree. For $n=1$, the root is an AND node with two descendants (see Fig. 1(a)). Again, this represents one solution tree. For the general case, each node in level $n-2$ has $f(n-2)$ solution trees (see Fig. 1(b)). A solution tree for the root at level n consists of picking two nodes in level $n-2$, a total of four combinations. Each pair of nodes selected in level $n-2$ represent two solution trees, all possible combinations of which will yield a new solution tree. This case is depicted in Fig. 1(b).

Since all solution trees are equally likely, the entropy of the heuristic information to guide the search at the root such that a correct decision is always made without backtracking is

$$I = \sum_{j=1}^{f(n)} \frac{1}{f(n)} \log_2 f(n) = 2(2^{n/2}-1),$$

which is exponential with respect to the height of the tree.

To manage the large memory space incurred by the storage of intermediate subproblems that may lead to solutions, we have investigated three alternatives to support branch-and-bound algorithms with a best-first search, the results of which are displayed in Table IV. In a direct implementation, the best-first search was implemented on an existing virtual-memory system, a VAX 11/780 computer running 4.2 BSD Unix. In the second approach, a modified virtual memory with specialized fetch and replacement policies was designed to adapt to the characteristics of the search algorithm. In the third approach, the no-wait policy discussed earlier was used to select subproblems in the main memory without waiting for the "most promising" subproblem to be accessed from the secondary memory. Again the no-wait policy is superior in performance [199].

The nondeterministic nature of computations and the fallibility of heuristic guiding may lead to anomalies of

problem to for a given architectural ation should : should the

gh heuristic hibitive. For a complete y likely. The level 0, while We have that fy the follow-

$t = 1$

there is one node with two represents one e in level $n - 2$ A solution tree two nodes in Each pair of solution trees, ld a new solu-

the entropy of ch at the root without back-

-1).

ght of the tree. curred by the may lead to natives to sup- est-first search. IV. In a direct mplemented on 11/780 com- d approach, a fetch and re- the character- approach, the select subprob- for the "most the secondary rior in perfor-

ations and the anomalies of




Fig. 1. Binary AND/OR trees. (a) With two levels. (b) With n levels. (Circles represent AND nodes; boxes represent OR nodes.)

TABLE IV
RELATIVE TIMES TO COMPLETE A BRANCH-AND-BOUND ALGORITHM FOR VARIOUS MEMORY-MANAGEMENT TECHNIQUES

Approach	0/1 Integer Programming Problems	0/1 Knapsack Problems
Direct implementation	1	1
Modified virtual memory	0.6	0.1
No-wait policy	0.1	0.001

parallelism. When n processors are used to solve the problem, the resulting speedup as compared to a single processor may be less than one, greater than n , or between one and n . The reasons for this anomalous behavior are due to 1) ambiguity in the heuristic information, 2) more than one solution node, and 3) approximation and dominance tests [113]. As a result, subtrees searched under serial processing will be terminated, and the search will be misled into a different part of the search tree.

In summary, scheduling is important when there is non-determinism in the problem. Good heuristic metrics to guide the search are usually difficult to design and depend on statistics such as success probabilities, search costs, and problem-dependent parameters. Trade-offs must be made among the dynamic overhead incurred in communicating the heuristic-guiding information, the benefits that would be gained if this information led the search in the right direction, and the granularity of tasks. In practice, the merits of heuristic guiding are not clear, since the heuristic information may be fallible. As a result, some AI architects do not schedule nondeterministic tasks in parallel. The excessive overhead coupled with the fallibility of heuristic information also leads some designers to apply only static scheduling to AI programs.

Pruning: Pruning can be considered as a negative form of heuristic guiding which guides the search to avoid subproblems that will never lead to better or feasible solutions. Pruning is useful in both backward and forward chaining. In backward reasoning, problems are decomposed into smaller subproblems and evaluated independently. There are usually redundant evaluations of the same task in different parts of the search tree when the

search trees are recursive. Likewise, in forward reasoning, the more primitive facts are reduced to form more general facts until the query is satisfied. Unnecessary results are generated because it is not clear which reduction will lead to a solution of the problem.

Pruning in search problems can be carried out by dominance relations. When a node P_i dominates another node P_j , it implies that the subtree rooted at P_i contains a solution with a value no more (or no less) than the minimum (or maximum) solution value of the subtree rooted at P_j .

As an example, consider two assignments P_1 and P_2 on the same subset of objects to be packed into a knapsack in the 0/1 knapsack problem. If the total profit of the objects assigned to the knapsack for P_1 exceeds that of P_2 and the total weight of the objects assigned in P_1 is less than that of P_2 , then the best solution expanded from P_1 dominates P_2 .

When parallel processing is used, it is necessary to keep the set of current dominating nodes (denoted by N_d) in memory [187]. These are nodes that have been generated but not yet dominated. In general, N_d can be larger than the set of active nodes. A newly generated node, P_i has to be compared with all nodes in N_d to see whether P_i or any nodes in N_d are dominated.

If N_d is small, it can be stored in a bank of global data registers. However, centralized comparisons are inefficient when N_d is large. A large N_d should then be partitioned into m subsets, N_d^0, \dots, N_d^{m-1} , and distributed among the local memories of the m processors. A subproblem P_{ij} generated in processor i , is first compared with N_d^i ; any subproblems in N_d^i dominated by P_{ij} are removed. If P_{ij}

is not dominated by a subproblem in N_d^i , it is sent to a neighboring processor and the process repeats. If it has not been dominated by any node in N_d, P_i , eventually returns to processor i and is inserted into N_d^i .

There are several problems associated with the use of dominance tests in AI applications. First, dominance relations are very problem-dependent and cannot be derived by a general methodology. Most of the dominance relations have been developed for dynamic programming problems. To derive a dominance relation in a search process, a dominance relation is hypothesized, and a proof is developed to show that the dominance relation is correct. Some progress has been made on using learning-by-experimentation to derive dominance relations for dynamic programming problems [200]. However, automatic proof techniques are largely missing. Moreover, learning-by-experimentation is applicable if there are a very small number of dominance relations that are used frequently in the problem. In many AI applications coded in Prolog, there is a large number of dominance relations, each of which is used infrequently in the program. Some special cases can be solved, such as finding redundant computations in recurrences [25]. For the general case, it is sometimes difficult to find these dominance relations without human ingenuity. Second, many dominance relations are related to the semantics of the applications. A good language to represent semantics is missing at this time. Lastly, the overhead of applying dominance relations is usually very high, and sequential and parallel implementations will incur prohibitive overhead.

Granularity of Parallelism: When a parallel computer system with a large number of processors is available, it is necessary to determine the granularity of parallelism, that is, the size of tasks that will be executed as an indivisible unit in a processor. Since many AI problems can be represented by AND/OR trees, some processors have to be idle when nodes close to the root are evaluated. The proper number of processors should be chosen to match the inherent parallelism in the problem to be solved.

The proper granularity is a function of the problem complexity, the shape of the AND/OR tree, and the distribution of processing times of tasks. Many of these parameters are dynamically changing and data-dependent, and only special cases can be analyzed [115]. An important functional requirement for parallel processing of AI programs is the ability to dynamically distribute the workload. For a system with a small granularity, an efficient interconnection network is required to transfer data and control information. In a loosely coupled system with a coarse grain, an effective load balancing mechanism is also needed.

IV. PROCESSOR LEVEL

The VLSI technology that has flourished in the past ten years has resulted in the development of many special-purpose computers for AI processing. Architectures for AI processing can be classified into the micro-, macro-, and

system-level architectures. Microlevel and macrolevel architectures are discussed in the next two sections. Sections IV-C-IV-G briefly discuss the system-level architectures. A taxonomy of architectures implementing AI systems have also been discussed by Hwang *et al.* [93].

A. Microlevel Architectures

The microlevel architectures consist of architectural designs that are fundamental to applications in AI. In the design of massively parallel AI machines [52], some of the basic computational problems recognized are set intersection, transitive closure, contexts and partitions, best-match recognition, Gestalt recognition, and recognition under transformation. These operations may not be unique to AI, and many exist in other applications as well. Due to the simplicity of some of these operations, they are usually implemented directly in hardware, especially in systolic arrays. Many other basic operations can also be implemented in VLSI. Examples include sorting and selection, computing transitive closure, string and pattern matching, selection from secondary memories, dynamic programming evaluations, proximity searches, and unification.

Some AI languages such as Lisp differ from traditional machine languages in that the program/data storage is conceptually an unordered set of linked record structures of various sizes, rather than an ordered indexable vector of numbers or bit fields of a fixed size. The instruction set must be designed according to the storage structure [160]. Additional concepts that are well-suited for list processing are the tagged memory [127] and stack architectures.

B. Macrolevel Architectures

The macrolevel is an intermediate level between the microlevel and the system level. In contrast to the microlevel architectures, macrolevel architectures are (possibly) made up of a variety of microlevel architectures and perform more complex operations. However, they are not considered as a complete AI system but can be taken as more complex supporting mechanisms for the system level. The architectures can be classified into those that manage data, such as dictionary machines, data base machines and structures for garbage collection, and those for searching.

A dictionary machine is an architecture that supports the insertion, deletion, and searching for membership, extremum, and proximity of keys in a data base [148]. Most designs are based on binary-tree architectures; however, designs using radix trees and a small number of processors have been found to be preferable when keys are long and clustered [57].

Data base machines depend on an architectural approach that distributes the search intelligence into the secondary and mass storage and relieves the workload of the central processor. Extensive research has been carried out in the past decade on optical and mass storage, back-end storage systems, and data base machines. Earlier data base machines developed were mainly directed toward

general-purpose relational data base management systems. Examples include the DBC, Direct, Rap, CASSM, associative array processors, text retrieval systems, and CAFS [12], [92], [109]. Nearly all current research on data base machines to support knowledge data bases assume that the knowledge data base is relational, hence research is directed toward solving the disk paradox [17] and enhancing previous relational data base machines by extensive parallelism [128], [143], [153], [173]. Commercially available data base and backend machines have also been applied in knowledge management [102], [131].

Searching is essential to many applications, although unnecessary combinatorial searches should be avoided. The suitability of parallel processing to searching depends on the problem complexity, the problem representation, and the corresponding search algorithms. Parallel algorithms and architectures to support divide-and-conquer, branch-and-bound, and AND/OR-graph search have been developed [187].

Extensive research has been carried out in supporting dynamic data structures in a computer with a limited memory space. *Garbage collection* is an algorithm that periodically reclaims memory space no longer needed by the users [30]. This is usually transparent to the users and could be implemented in hardware, software, or a combination of both. For efficiency reasons, additional hardware such as stacks and reference counters are usually provided.

C. Functional-Programming-Oriented System-Level Architectures

The objective of writing a functional program is to define a set of (possibly recursive) equations for each function [33]. Data structures are handled by introducing a special class of functions called constructor functions. This view allows functional languages to deal directly with structures that would be termed "abstract" in more conventional languages. Moreover, functions themselves can be passed around as data objects. The design of the necessary computer architecture to support functional languages thus centers around the parallel evaluation of functional programs (function-oriented architectures) and the mechanisms of efficient manipulation of data structures (list-oriented architectures).

In *function-oriented architectures*, the design issues center on the physical interconnection of processors, the method used to "drive" the computation, the representation of programs and data, the method to invoke and control parallelism, and the optimization techniques [184]. Desirable features of such architectures should include a multiprocessor system with a rich interconnection structure, the representation of list structures by balanced trees, and hardware supports for demand-driven execution, low-overhead process creation, and storage management.

Architectures to support functional-programming languages can be classified as uniprocessor architectures, tree-structured machines, data-driven machines, and demand-driven machines. In a uniprocessor architecture, be-

sides the mechanisms to handle lists, additional stacks to handle function calls and optimization for redundant calls and array operations may be implemented [23], [159], [179]. Tree-structured machines usually employ lazy evaluations, but suffer from the bottleneck at the root of the tree [35], [120], [133]. Data-flow machines are also natural candidates for executing functional programs and have tremendous potential for parallelism. However, the issue of controlling parallelism remains unresolved. A lot of the recent work has concentrated on demand-driven machines which are based on reduction machines on a set of load-balanced (possibly virtual) processors [28], [32], [100], [101], [105], [175], [176].

List-oriented architectures are architectures designed to support the manipulation of data structures and objects efficiently. Lisp, a mnemonic for list processing language, is a well-known language to support symbolic processing. There are several reasons why Lisp and list-oriented computers are really needed. First, to relieve the burden on the programmers, Lisp was designed as an untyped language. The computer must be able to identify the types of data, which involves an enormous amount of data-type checking and the use of long strings of instructions at compile and run times. Conventional computers cannot do these efficiently in software. Second, the system must periodically perform garbage collection and reclaim unused memory at run time. This amounts to around 10-30 percent of the total processing time in a conventional computer. Hardware implementation of garbage collection is thus essential. Third, due to the nature of recursion, a stack-oriented architecture is more suitable for list processing. Lastly, list processing usually requires an enormous amount of space, and the data structures are so dynamic that the compiler cannot predict how much space to allocate at compile time. Special hardware to manage the data structures and the large memory space would make the system more efficient [37], [58].

The earliest implementation of Lisp machines were the PDP-6 computer and its successors the PDP-10 and PDP-20 made by the Digital Equipment Corporation (DEC) [122]. The half-word instructions and the stack instructions of these machines were developed with Lisp's requirements in mind. Extensive work has been done for the DEC-system 10's and 20's on garbage collection to manage and reclaim the memory space used.

The design of Lisp machines was started at MIT's AI Laboratory in 1974. Cons, designed in 1976 [106], was superseded in 1978 by a second generation Lisp machine, the CADR. This machine was a model for the first commercially available Lisp machines, including the Symbolics LM2, the Xerox 1100 Interlisp workstation, and the Lisp Machine Inc. Series III CADR, all of them delivered in 1981. The third-generation machines were based on additional hardware to support data tagging and garbage collection. They are characterized by the Lisp Machines Inc. Lambda supporting Zetalisp and LMLisp, the Symbolics 3600 supporting Zetalisp, Flavors, and Fortran 77, the Xerox 1108 and 1132 supporting Interlisp-D and Smalltalk,

and the Fujitsu FACOM Alpha Machine, a backend Lisp processor supporting Maclisp. Most of the Lisp machines support networking using Ethernet. The LMI Lambda has a NuBus developed at MIT to produce a modular, expandable Lisp machine with multiprocessor architecture.

A single-chip processor to support Lisp has been implemented in the MIT Scheme-79 chip [166]. Other experimental computers to support Lisp and list-oriented processing have been reported [44], [71]–[73], [130], [139], [144]–[146], [170]. These machines usually have additional hardware tables, hashing hardware, tag mechanisms, and list processing hardware, or are microprogrammed to provide macroinstructions for list processing. A Lisp chip built by Texas Instruments implements over half a million transistors on a 1-cm² chip for 60 percent of the functions in a TI Explorer. The implementation on a single chip results in five times improvement in performance [121]. Experimental multiprocessing systems have been proposed to execute Lisp programs concurrently [75], [86], [124], [125], [163], [164], [193]. Data-flow processing is suitable for Lisp as these programs are generally data driven [7], [8], [196], [197]. Other multiprocessing architectures to support list processing have been proposed and developed [29], [45], [66], [84], [176].

Architectures have also been developed to support object-oriented programming languages. Smalltalk, first developed in 1972 by the Xerox Corporation, is recognized as a simple but powerful way of communicating with computers. At MIT, the concept was extended to become the Flavors system. Special hardware and multiprocessors have been proposed to directly support the processing of object-oriented languages [96], [138], [167], [183].

Owing to the different motivations and objectives of various functional-programming-oriented architectures, each machine has its own distinct features. For example, the Symbolics 3600 [127] was designed for an interactive program development environment where compilation is very frequent and ought to appear instantaneous to the user. This requirement simplified the design of the compiler and results in only a single-address instruction format, no indexed and indirect addressing modes, and other mechanisms to minimize the number of nontrivial choices to be made. On the other hand, the aim in developing Soar [183] was to demonstrate that a reduced instruction set computer (RISC) could provide high performance in an exploratory programming environment. Instead of microcode, Soar relied on software to provide complicated operations. As a result, more sophisticated software techniques were used.

D. Logic- and Production-Oriented System-Level Architectures

Substantial research has been carried out on parallel computational models of utilizing AND parallelism, OR parallelism, and stream parallelism in logical inference systems, production systems and others. The basic prob-

lem on their exponential complexity remains open at this time.

Sequential Prolog machines using software interpretation, emulation, and additional hardware support such as hardware unification and backtracking [174] have been reported. Single-processor systems for production systems using additional data memories [110] and a RISC architecture [60] have been studied.

New logic programming languages suitable for parallel processing have been investigated. In particular, the use of predicate logic [49], extensions of Prolog to become Concurrent Prolog [150], Parlog [26], and Delta-Prolog [137], and parallel production systems [182] have been developed. One interesting parallel language is that of systolic programming, which is useful as an algorithm design and programming methodology for high-level-language parallel computers [151].

Several prototype multiprocessor systems for processing inference programs and Prolog have been proposed, some of which are currently under construction. These systems include multiprocessors with a shared memory [18], ZMOB, a multiprocessor of Z80's connected by a ring network [192], Aquarius, a heterogeneous multiprocessor with a crossbar switch [43], and Mago, a cellular machine implementing a Prolog compiler that translates a Prolog program into a formal functional program [107]. Techniques for analyzing Prolog programs such that they can be processed on a data-flow architecture have been derived [9], [15], [81], [95], [97]. An associative processor has been proposed to carry out propositional and first-order predicate calculus [46].

Dado is a multiprocessor system with a binary-tree interconnection network that implements parallel production systems [162]. Non-Von is another tree architecture used to evaluate production systems at a lower level of granularity [152].

E. Distributed Problem-Solving Systems

Knowledge in an AI system can sometimes be represented in terms of semantic nets. Several proposed and experimental architectures have been developed. NETL [51] and its generalization to Thistle [52] consist of an array of simple cells with marker-passing capability to perform searches, set-intersections, inheritance of properties and descriptions, and multiple-context operations on semantic nets. Thinking Machine Inc.'s Connection Machine is a cellular machine with 65 536 processing elements. It implements marker passing and virtually reconfigures the processing elements to match the topology of the application semantic nets [87]. Associative processors for processing semantic nets have also been proposed [126].

Some AI architectures are based on frame representations and may be called object-oriented architectures. For example, the *Apiary* developed at MIT is a multiprocessor actor system [84]. An efficient AI architecture may also depend on the problem-solving strategy. A general form of

architectures called connectionist architectures evolve from implementing neurons in brains [53]. The basic idea of the *Boltzmann machine* is the application of statistical mechanics to constrained searches in a parallel network [89]. The most interesting aspect of this machine lies in its domain-independent learning algorithm [3].

With the inclusion of control into stored knowledge, the resulting system becomes a distributed problem-solving system. These systems are characterized by the relative autonomy of the problem-solving nodes, a direct consequence of the limited communication capability. With the proposed formalism of the contract net, contracts are used to express the control of problem solving in a distributed processor architecture [157]. Related work in this area include Petri-net modeling [135], distributed vehicle-monitoring testbed [112], distributed air-traffic control system [22], and modeling the brain as a distributed system [61], [65].

F. Hybrid Systems

It has been suggested that a combination of Lisp, Prolog, and an object-oriented language such as Smalltalk may be a better language for AI applications [169]. This approach can be carried out in two ways. First, multiple AI languages can be implemented using microprogramming on the same computer, so programs written in these languages can be executed independently. For example, Prolog is available as a secondary language on some Lisp machines. A version of a Prolog interpreter with a speed of 4.5 KLIPS (kilo lines of interpreted statements) has been developed for Lisp Machine's Lambda. A second approach is to design a language that combines the desirable features from several AI languages into a new language. Some of the prototype multiprocessors, such as ZMOB and Mago, were developed with a flexible architecture that can implement object-oriented, functional, and logic languages. FAIM-1, a multiprocessor connected in the form of a twisted hex-topology, was designed to implement the features of object-oriented, functional, and logic programming in the Oil programming language [11]. Currently, a parallel version of Scheme similar to MultiLisp is being implemented. Hope, a hybrid functional and logic language, is currently being implemented on Alice [156].

G. Fifth Generation Computer Projects

The fifth generation computer system (FGCS) project was started in Japan in 1982 to further the research and development of the next generation of computers. It was conjectured that computers of the next decade will be used increasingly for nonnumeric data processing such as symbolic manipulation and applied AI. The goals of the FGCS project are

- a) to implement basic mechanisms for inference, association, and learning in hardware;
- b) to prepare basic AI software to utilize the full power of the basic mechanisms implemented;
- c) to implement the basic mechanisms for retrieving and managing a knowledge base in hardware and software;
- d) to use pattern recognition and AI research achievements in developing user-oriented man-machine interfaces; and
- e) to realize supporting environments for resolving the "software crisis" and enhancing software production.

The FGCS project is a marriage between the implementation of a computer system and the requirements specified by applications in AI, such as natural-language understanding and speech recognition. Specific issues studied include the choice of logic programming over functional programming, the design of the basic software systems to support knowledge acquisition, management, learning, and the intelligent interface to users, the design of highly parallel architectures to support inferencing operations, and the design of distributed-function architectures that integrates VLSI technology to support knowledge data bases [99], [177], [180].

A first effort in the FGCS project is to implement a sequential inference machine, or Sim [198]. Its first implementation are two medium-performance machines for software development known as personal sequential inference (PSI) machine and cooperative high-speed inference (CHI) machine [171]. The PSI and CHI machines have further been implemented in custom LSI's into PSI-II and CHI-II. The PSI-II has been found to have a performance that ranges from 100 to 333 KLIPS for various benchmark programs. Another architectural development is on the knowledge-base machine, Delta [129].

The current efforts in the intermediate stage are on the parallel inference machine, or PIM, and the multi-PSI computers [129]. As an intermediate target, PIM-1 is being built now. It consists of about 100 processing elements, with a total speed of 10-20 MLIPS (mega lines of interpreted statements) including overhead caused by Pimos. Eight processing elements with private caches in a cluster are connected through a shared memory, and a switching network is used to connect the clusters. Each processing element will be implemented in standard-cell VLSI chips. The machine language is KL1-B based on GHC [147]. Lastly, the development of the basic software system acts as a bridge to fill the gap between a highly parallel computer architecture and knowledge information processing [62]. The Pimos was designed as a single unified operating system to control the parallel hardware [172]. It was built on the multi-PSI (version 2) system. Each PE consists of a PSI-II, 16-MW main memory, and interfaces to the mesh interconnection network. The KL1-B interpreter is implemented in firmware and attains a speed of 100-150 KLIPS [94].

In the final stage, a parallel computer with about 1000 processing elements and attaining 100 mlips to 1 GLIPS (giga lines of interpreted statements) is expected to be built. Although the projects are progressing well, there is a

recognition that more research is needed on exploiting intelligence rather than brute-force parallelism. The proposal of the sixth generation computer system project is an indication of efforts in this direction [4].

The Japanese FGCS project has stirred intensive responses from other countries. The British project is a five-year \$550 million cooperative program between government and industry that concentrates on software engineering, intelligent knowledge-based systems, VLSI circuitry, and man-machine interfaces. Hardware development has focused on Alice, a Parlog machine using data-flow architectures and implementing both Hope, Prolog, and Lisp [156]. The European Commission has started the \$1.5 billion five-year European Strategic Program for Research in Information Technologies (Esprit) in 1984 [2]. The program focuses on microelectronics, software technology, advanced information processing, computer-integrated manufacturing, and office automation. In the U.S., the most direct response to the Japanese FGCS project was the establishment of the Microelectronics and Computer Technology Corporation in 1983 [1]. The project has an annual budget of \$50-\$80 million per year. It has a more evolutionary approach than the revolutionary approach of the Japanese and would yield technology that the corporate sponsors can build into advanced products in the next 10-12 years. Meanwhile, other research organizations have formed to develop future computer technologies of the U.S. in a broader sense. These include Darpa's Strategic Computing and Survivability, the semiconductor industry's Semiconductor Research Corporation, and the Microelectronics Center of North Carolina [1].

V. DESIGN DECISIONS OF AI-ORIENTED COMPUTERS

The appropriate methodology to design an AI computer should utilize a top-down design approach: functional requirements should be developed from the problem requirements, which are mapped into hardware based on technological constraints. Similar to the design of conventional computers, a bottom-up design approach is not adequate since special requirements of the applications may not be satisfied. Before a design is made, it is important to understand the applicability of the system to a class of problems and then to strive for high performance in a prototype implementation. Thus knowing that an m -processor system gives a k -fold increase in performance over a single processor is more important than knowing the maximum instruction rate of a prototype. Proper understanding and analysis of the problem is probably more important than applying brute-force parallelism randomly in the design.

The issues classified in Table II provide a view to the sequence of design decisions made in developing a special-purpose computer to support AI processing. The various approaches can be classified as top-down, bottom-up, and middle-out.

Top-Down Design Decisions: This approach starts by defining, specifying, refining, and validating the require-

ments of the application, devising methods to collect the necessary knowledge and metaknowledge, choosing an appropriate representation for the knowledge and metaknowledge, studying problems related to the control of correct and efficient execution with the given representation scheme, identifying functional requirements of components, and mapping these components into software and microlevel, macrolevel and system-level architectures subject to technological and cost constraints. The process is iterative. For example, the representation of knowledge and the language features may be changed or restricted when it is discovered that the functional requirements found cannot be mapped into a desirable and realizable system with the given technology and cost. In some projects, the requirements may be very loose and span across many different applications. As a result, the languages and knowledge-representation schemes used may be oriented towards general-purpose usage. The Japanese FGCS project is an attempt to use a top-down approach to design an integrated user-oriented intelligent system for a wide spectrum of applications.

Bottom-Up Design Decisions: In this approach, the designers first design the computer system based on a computational model, such as data flow, reduction, and control flow, and the technological and cost limitations. Possible extensions of existing knowledge-representation schemes and languages developed for AI applications are implemented. Finally, AI applications are coded using the representation schemes and languages provided. This is probably the most popular approach to apply a general-purpose or existing system for AI processing. However, it may result in inefficient processing, and the available representation schemes and languages may not satisfy the application requirements completely, ZMOB and Butterfly Multiprocessor are examples in this class.

Middle-Out Design Decisions: This approach is a short cut to the top-down design approach. It starts from a proven and well-established knowledge-representation scheme or AI language (most likely developed for sequential processing) and develops the architecture and the necessary modifications to the language and representation scheme to adapt to the application requirements and the architecture. This is the approach taken by many designers in designing special-purpose computers for AI processing. It may be subdivided into top-first and bottom-first, although both may be iterative. In a top-first middle-out approach, studies are first performed to modify the language and representation scheme to make it more adaptable to the architecture and computational model. Primitives may be added to the language to facilitate parallel processing. Nice features from several languages may be combined. The design of the architecture follows. Alice and FAIM-1 are examples of architectures designed using this approach. In the bottom-first middle-out approach, the chosen language or representation scheme is mapped directly into architecture by providing hardware support for the overhead-intensive operations. Applications are implemented using the language and representation scheme

ds to collect the
choosing an ap-
edge and meta-
the control of
given representa-
ments of com-
into software and
architectures sub-
s. The process is
on of knowledge
ged or restricted
al requirements
le and realizable
ost. In some pro-
and span across
he languages and
may be oriented
these FGCS pro-
oach to design an
for a wide spec-

approach, the de-
based on a com-
ction, and control
ditions. Possible
entation schemes
ations are imple-
d using the repre-
ed. This is proba-
a general-purpose
However, it may
available represen-
atisfy the applica-
l Butterfly Multi-

approach is a short
It starts from a
ge-representation
oped for sequen-
itecture and the
nd representation
irements and the
y many designers
or AI processing,
bottom-first, al-
-first middle-out
modify the lan-
e it more adapt-
model. Primi-
facilitate parallel
anguages may be
e follows. Alice
s designed using
e-out approach,
heme is mapped
hardware support
applications are
entation scheme

provided. Lisp computers are examples designed with this approach.

VI. CONCLUSION

Although many AI computers have been proposed or built, Lisp computers are probably the only architecture that have had widespread use for solving real AI problems. This is probably due to the large investment in software for many applications coded in Lisp. At present, there is no comprehensive methodology for designing parallel AI computers. Research on AI in the past three decades and the recent experience in building AI computers have led to a view that the key issue of an AI system lies in the understanding of the problem rather than efficient software and hardware. In fact, most underlying concepts in AI computers are not new and have been used in conventional systems. For example, hardware stack and tagged memory were proposed before they were used in Lisp computers. However, the above argument does not imply that research on hardware and architectures is not necessary.

To support efficient processing of AI applications, research must be done in developing better AI algorithms, better AI software management methods, and better AI architectures. The development of better algorithms can lead to significant improvements in performance. Many AI algorithms are heuristic in nature, and upper bounds on performance to solve these problems have not been established as in traditional combinatorial problems. As a consequence, the use of better heuristic information, based on common-sense or high-level metaknowledge and better representation of the knowledge, can have far greater improvement in performance than improved computer architecture. Automatic learning methods to aid designers in systematically acquiring and managing new knowledge to be available in the future are very important.

Better AI software management methods are essential in developing more efficient and reliable software for AI processing. AI systems are usually open and cannot be defined based on a closed-world model. The language must be able to support the acquisition of new knowledge and the validation of existing knowledge. Probabilistic reasoning, fuzzy knowledge, and nonmonotonic logic may have to be supported. The verification of the correctness of an AI program is especially difficult due to the imprecise knowledge involved and the disorganized way of managing knowledge in a number of declarative languages and representation schemes. Traditional software engineering design methodologies must be extended to become knowledge engineering to accommodate the characteristics of knowledge in AI applications. Automatic programming is important to aid designers to generate the AI software from specifications.

The role of parallel processing and innovative computer architectures lies in improving the processing time of solving a given AI problem. It is important to realize that parallel processing and better computer architectures can-

not be used to overcome the exponential complexity of exhaustive enumeration (unless an exponential amount of hardware is used) and are not very useful to extend the solvable problem space. For a problem with a size that is too large to be solved today by a sequential computer in a reasonable amount of time, it is unlikely that it can be solved by parallel processing alone, even if a linear speedup can be achieved. The decision to implement a given algorithm in hardware depends on the complexity of the problem it solves and its frequency of occurrence. Problems of low complexity can be solved by sequential processing or in hardware if they are frequently encountered; problems of moderate complexity should be solved by parallel processing; and problems of high complexity should be solved by a combination of heuristics and parallel processing. In many AI systems developed today, tasks and operations implemented in hardware are those that are frequently executed and have polynomial complexity. These tasks or operations are identified from the languages or the knowledge-representation schemes supported. The architectural concepts and parallel processing schemes applied may be either well-known conventional concepts or new concepts for nondeterministic and dynamic processing. The role of the computer architects lies in choosing a good representation, recognizing overhead-intensive tasks to maintain and learn metaknowledge, identifying primitive operations in the languages and knowledge-representation schemes, and supporting these tasks in hardware and software.

REFERENCES

- [1] Special Issue on Tomorrow's Computers. *IEEE Spectrum*, vol. 20, no. 11, pp. 51-58, 69, Nov. 1983.
- [2] "ESPRIT: Europe challenges U.S. and Japanese competitors." *Future Generation Computer Syst.*, vol. 1, no. 1, pp. 61-69, 1984.
- [3] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for Boltzmann machines." *Cognitive Sci.*, vol. 9, no. 1, pp. 147-169, 1985.
- [4] Science and Technology Agency. *Promotion of Research and Development on Electronic and Information Systems that May Complement or Substitute for Human Intelligence*. Tokyo, Japan: Science and Technol. Agency, Tokyo, 1985.
- [5] G. Agha and C. Hewitt, "Concurrent programming using actors: Exploiting large-scale parallelism." *Lecture Notes in Comput. Sci.*, no. 206, pp. 19-41, Dec. 1985.
- [6] G. Agha, *Actor: A Model of Concurrent Computation in Distributed Systems*. Cambridge, MA: MIT Press, 1986.
- [7] M. Amamiya, R. Hasegawa, O. Nakamura, and H. Mikami, "A list-processing-oriented data flow machine architecture." in *Proc. Nat. Computer Conf.*, pp. 144-151, AFIPS Press, 1982.
- [8] M. Amamiya and R. Hasegawa, "Dataflow computing and eager and lazy evaluations." *New Generation Computing*, vol. 2, no. 2, pp. 105-129, 1984.
- [9] M. Amamiya, M. Takesue, R. Hasegawa, and H. Mikami, "Implementation and evaluation of a list-processing-oriented data flow machine." in *Proc. 13th Annu. Int. Symp. Computer Architecture*, Tokyo, Japan, June 1986, pp. 10-19.
- [10] G. M. Amdahl, "Tempered expectations in massively parallel processing and semiconductor industry," presented at the 2nd Int. Conf. Supercomputing, Santa Clara, CA, May 1987.
- [11] J. M. Anderson et al., "The architecture of FAIM-1." *IEEE Computer*, vol. 20, no. 1, pp. 55-65, Jan. 1987.
- [12] E. Babb, "Joined normal form: A storage encoding for relational databases." *Assoc. Comput. Mach. Trans. Database Syst.*, vol. 7, no. 4, pp. 588-614, Dec. 1982.

- [13] J. Backus, "Function-level computing," *IEEE Spectrum*, vol. 19, no. 8, pp. 22-27, Aug. 1982.
- [14] R. Bailey, "A Hope tutorial," *Byte*, vol. 10, no. 8, pp. 235-258, Aug. 1985.
- [15] L. Bic, "Execution of logic programs on a dataflow architecture," in *Proc. 11th IEEE/ACM Annu. Int. Symp. Computer Architecture*, June 1984, pp. 290-296.
- [16] D. G. Bobrow, et al., "CommonLoops: Merging common Lisp and object-oriented programming," Xerox Palo Alto Research Center, Tech. Rep. ISL-85-8, Aug. 1985.
- [17] H. Boral and D. DeWitt, "Database machine: An idea whose time has passed?" *Database Machines*, pp. 166-167, 1983.
- [18] P. Borgwardt, "Parallel Prolog using stack segments on shared-memory multiprocessors," in *Proc. IEEE Int. Symp. Logic Programming*, pp. 2-11, Feb. 1984.
- [19] K. Bowen, "Meta-level programming and knowledge representation," *New Generation Computing*, vol. 3, no. 4, pp. 359-383, 1985.
- [20] R. Brachman and H. Levesque, Eds., *Readings in Knowledge Representation*. Los Altos, CA: Morgan Kaufmann, 1985.
- [21] B. G. Buchanan and E. H. Shortliffe, *Rule-Based Expert Programs: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley, 1984.
- [22] S. Cammarata, D. McArthur, and R. Steeb, "Strategies of cooperation in distributed problem solving," in *Proc. 8th Int. Joint Conf. Artificial Intelligence*, Aug. 1983. Los Altos, CA: Kaufmann, pp. 767-770.
- [23] M. Castan and E. I. Organick, "M3L: An HLL-RISC Processor for Parallel Execution of FP-Language Programs," in *Proc. 9th Annu. IEEE/ACM Symp. Computer Architecture*, 1982, pp. 239-247.
- [24] J. H. Chang, A. M. Despain, and D. DeGroot, "AND-parallelism of logic programs based on a static data dependency analysis," in *Proc. IEEE COMPCON Spring*, 1985, pp. 218-225.
- [25] H.-Y. Chen and B. W. Wah, "The RID-REDUNDANT procedure in C-Prolog," in *Proc. Int. Symp. Methodologies for Intelligent Systems*, Charlotte, NC, Oct. 1987, pp. 71-75.
- [26] K. Clark and S. Gregory, "PARLOG: Parallel programming in logic," Imperial College, London, England, Res. Rep. DOC 84/4, 1984.
- [27] K. Clark and S. Gregory, "Note on system programming in PARLOG," in *Proc. Int. Conf. 5th Generation Computer System*, 1984, pp. 299-306.
- [28] T. Clarke, P. Gladstone, C. Maclean, and A. Norman, "SKIM—The S, K, I reduction machine," in *Conf. Rec. Lisp Conf.*, Stanford Univ., Menlo Park, CA, 1980.
- [29] G. Coghill and K. Hanna, "PLEIADES: A multimicroprocessor interactive knowledge base," *Microprocessors Microsyst.*, vol. 3, no. 2, pp. 77-82, Mar. 1979.
- [30] J. Cohen, "Garbage collection of linked data structures," *Computing Surveys*, vol. 13, no. 3, pp. 341-367, Sept. 1981.
- [31] J. S. Conery and D. F. Kibler, "AND parallelism and nondeterminism in logic programs," *New Generation Computing*, vol. 3, no. 1, pp. 43-70, 1985.
- [32] J. Darlington and M. Reeve, "ALICE and the Parallel Evaluation of Logic Programs," Dept. Computing, Imperial College of Science and Technology, London, England, Preliminary Draft, June 1983.
- [33] J. Darlington, "Functional programming," in *Distributed Computing*, F. B. Chambers, D. A. Duce, and G. P. Jones, Eds. London: Academic, 1984, ch. 5.
- [34] J. Darlington, A. J. Field, and H. Pull, "The unification of functional and logic languages," Imperial College, London, England, Tech. Rep., Feb. 1985.
- [35] A. L. Davis, "A data flow evaluation system based on the concept of recursive locality," in *Proc. Nat. Computer Conf. AFIPS Press*, 1979, pp. 1079-1086.
- [36] A. L. Davis and S. V. Robison, "The FAIM-1 symbolic multiprocessing system," in *Proc. IEEE COMPCON Spring*, 1985, pp. 370-375.
- [37] M. F. Deering, "Architectures for AI," *Byte*, pp. 193-206, Apr. 1985.
- [38] D. DeGroot, "Restricted AND-Parallelism," in *Proc. Int. Conf. 5th Generation Computers*, Nov. 1984, pp. 471-478.
- [39] D. DeGroot and G. Lindstrom, Eds., *Logic Programming*. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [40] D. DeGroot and J. H. Chang, "A comparison of two AND-parallel execution models," in *Proc. Hardware and Software Components and Architectures for the 5th Generation*, AFCET Informatique, Paris, France, Mar. 1985, pp. 271-280.
- [41] D. DeGroot, "Restricted AND-parallelism and side-effects in logic programming," in *Supercomputers and AI Machines*, K. Hwang and D. DeGroot, Eds. New York: McGraw-Hill, 1988.
- [42] P. Denning, "A view of Kanerva's sparse distributed memory," NASA Ames Research Center, Moffett Field, CA, RIACS Tech. Rep. TR-86.14, June 1986.
- [43] A. M. Despain and Y. N. Patt, "Aquarius—A high performance computing system for symbolic/numeric application," in *Proc. IEEE COMPCON Spring*, Feb. 1985, pp. 376-382.
- [44] P. Deutsch, "Experience with a microprogrammed interlisp systems," in *Proc. MICRO*, vol. 11, Nov. 1978.
- [45] H. Diel, "Concurrent data access architecture," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 373-388.
- [46] W. Dilger and J. Muller, "An associative processor for theorem proving," in *Proc. IFAC Symp. Artificial Intell.*, 1983, pp. 489-497.
- [47] J. Doyle, "A truth maintenance system," *Artificial Intell.*, vol. 12, no. 3, pp. 231-272, 1979.
- [48] H. Dreyfus and S. Dreyfus, "Why expert systems do not exhibit expertise," *IEEE Expert*, vol. 1, no. 2, Summer 1986.
- [49] M. H. van Emden and G. J. de Lucena-Filho, "Predicate logic as a language for parallel programming," in *Logic Programming*, S. A. Tarnlund and K. Clark, Eds. New York: Academic, 1982, pp. 189-198.
- [50] L. D. Erman, F. Hayes-Roth, V. R. Lesser, and D. R. Reddy, "The Hearsay-II speech-understanding system: Integrating knowledge to resolve uncertainty," *Assoc. Comput. Mach. Computing Surveys*, vol. 12, no. 2, pp. 213-253, June 1980.
- [51] S. E. Fahlman, *NETL: A System for Representing and Using Real-World Knowledge*. Series on Artificial Intelligence, Cambridge, MA: MIT Press, 1979.
- [52] S. E. Fahlman and G. E. Hinton, "Massively parallel Architectures for AI: NETL, THISTLE, and Boltzmann Machines," in *Proc. AAAI Nat. Conf. Artificial Intell.*, 1983, pp. 109-113.
- [53] S. E. Fahlman and G. E. Hinton, "Connectionist architecture for artificial intelligence," *IEEE Computer*, vol. 20, no. 1, pp. 100-109, Jan. 1987.
- [54] E. A. Feigenbaum, "Knowledge engineering: The applied side," in *Intelligent Systems: The Unprecedented Opportunity*, J. E. Hayes and D. Michie, Eds. Chichester, England: Ellis Horwood Ltd., 1983, pp. 37-55.
- [55] R. D. Fennell and V. R. Lesser, "Parallelism in artificial intelligence problem solving: A case study of Hearsay-II," *IEEE Trans. Computers*, vol. C-26, no. 2, pp. 98-111, Feb. 1977.
- [56] R. E. Fikes and N. J. Nilsson, "Strips: A new approach to the application of theorem proving to problem solving," *Artificial Intell.*, vol. 2, no. 3 & 4, pp. 189-208, 1971.
- [57] A. L. Fisher, "Dictionary machines with a small number of processors," in *Proc. 11th Annu. IEEE/ACM Int. Symp. Computer Architecture*, June 1984, pp. 151-156.
- [58] J. Fitch, "Do we really want a Lisp machine?" presented at the ACM SEAS/SMC Annu. Meeting, Jan. 1980.
- [59] A. M. Flynn and J. G. Harris, "Recognition algorithms for the connection machine," in *Proc. Int. Joint Conf. Artificial Intell.*, 1985, pp. 57-60.
- [60] C. L. Forgy, A. Gupta, A. Newell, and R. Wedig, "Initial assessment of architectures for production systems," in *Proc. AAAI Nat. Conf. Artificial Intell.* Aug. 1984, pp. 116-120.
- [61] W. Fritz, "The Intelligent System," *SIGART Newsletter*, no. 90, pp. 34-38, Oct. 1984.
- [62] K. Furukawa and T. Yokoi, "Basic software system," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 37-57.
- [63] H. Gallaire and C. Lasserre, "Metalevel control for logic programs," in *Logic Programming*, K. L. Clark and S. A. Tarnlund, Eds. New York: Academic, 1982, pp. 173-185.
- [64] M. R. Genesereth, "An overview of meta-level architecture," in *Proc. AAAI Nat. Conf. Artificial Intell.*, 1983, pp. 119-124.
- [65] A. S. Gevins, "Overview of the human brain as a distributed computing network," in *Proc. IEEE Int. Conf. Computer Design: VLSI in Computers*, 1983, pp. 13-16.
- [66] W. K. Giloi and R. Gueth, "Concepts and realization of a high-performance data type architecture," *Int. J. Comput. Inform. Sci.*, vol. 11, no. 1, pp. 25-54, 1982.
- [67] W. K. Giloi, "Advanced object oriented architecture," *Future Generation Comput. Sys.*, vol. 2, no. 2, pp. 169-175, 1985.

- tion, *AFCET Informa-*
 30.
- m and side-effects in
and AI Machines, K.
 McGraw-Hill, 1988.
- "distributed memory,"
 field, CA, RIACS Tech.
- A high performance
 application," in *Proc.*
 376-382.
- grammed interisp sys-
 78," in *Proc. Int. Conf.*
 373-388.
- processor for theorem
Intell., 1983, pp.
- Artificial Intell.*, vol. 12,
- systems do not exhibit
 1986.
- no. "Predicate logic as a
Logic Programming, S. A.
 Academic, 1982, pp.
- ser, and D. R. Reddy,
 tem: Integrating knowl-
Comput. Mach. Computing
 1980.
- Representing and Using*
Artificial Intelligence, Cam-
- sively parallel Architec-
 tzmann Machines," in
 1983, pp. 109-113.
- ectionist architecture for
 1. 20, no. 1, pp. 100-109.
- ng: The applied side," in
 pportunity, J. E. Hayes
 d. Ellis Horwood Ltd.,
- lism in artificial intelli-
 arsary-II," *IEEE Trans.*
 eb. 1977.
- A new approach to the
 lem solving," *Artificial*
- h a small number of
 M Int. Symp. Computer
- ine?" presented at the
 1980.
- ion algorithms for the
 Conf. Artificial Intell.,
- Wedig, "Initial assess-
 ems," in *Proc. AAAI*
 116-120.
- RT Newsletter, no. 90,
- e system," in *Proc. Int.*
 84, pp. 37-57.
- ontrol for logic pro-
 k and S. A. Tarnlund,
 -185.
- level architecture," in
 83, pp. 119-124.
- rain as a distributed
 Conf. Computer Design:
- and realization of a
 ut. J. Comput. Inform.
- architecture," *Future*
 69-175, 1985.
- [68] A. J. Goldberg and D. Robson, *Smalltalk-80: The Language and Its Implementation*. Reading, MA: Addison-Wesley, 1983.
- [69] M. M. Gooley and B. W. Wah, "Efficient reordering of Prolog programs," in *Proc. 4th IEEE Int. Conf. Data Engineering*, Los Angeles, CA, Feb. 1988, pp. 71-75.
- [70] A. Goto, H. Tanaka, and T. Moto-oka, "Highly parallel inference engine PIE—Goal rewriting model and machine architecture," *New Generation Computing*, vol. 2, no. 1, pp. 37-58, 1984.
- [71] E. Goto, T. Ida, K. Hiraki, M. Suzuki, and N. Inada, "FLATS, A machine for numerical, symbolic and associative computing," in *Proc. 6th Int. Joint Conf. Artificial Intell.*, Los Altos, CA: William Kaufman, Aug. 1979, pp. 1058-1066.
- [72] N. Greenfield and A. Jericho, "A professional's personal computer system," in *Proc. IEEE/ACM 8th Int. Symp. Computer Architecture*, 1981, pp. 217-226.
- [73] M. Griss and M. Swanson, "MBALM/1700: A microprogrammed Lisp machine for the Burroughs B1726," in *Proc. ACM/IEEE MICRO-10*, 1977.
- [74] A. Gupta, "Implementing OPS5 production systems on DADO," in *Proc. IEEE Int. Conf. Parallel Processing*, 1984, pp. 83-91.
- [75] A. Guzman, "A heterarchical multi-microprocessor Lisp machine," in *Proc. IEEE Workshop on Computer Architecture for Pattern Analysis and Image Database Management*, Nov. 1981, pp. 309-317.
- [76] R. H. Halstead, Jr., "Implementation of MULTILISP: LISP on multiprocessor," in *Proc. ACM Symp. LISP and Functional Programming*, 1984.
- [77] R. Halstead, "Parallel symbolic computing," *IEEE Computer*, vol. 19, no. 8, pp. 35-43, Aug. 1986.
- [78] R. Halstead, Jr., and J. Loaiza, "Exception handling in multilisp," in *Proc. Int. Conf. Parallel Processing*, Aug. 1985, pp. 822-830.
- [79] R. Halstead, Jr., T. Anderson, R. Osborne, and T. Sterlig, "Concert: Design of a multiprocessor development system," in *Proc. IEEE/ACM Int. Symp. Computer Architecture*, June 1986, pp. 40-48.
- [80] R. Halstead, Jr., "Design requirements of concurrent Lisp machines," in *Supercomputers and AI Machines*, K. Hwang and D. DeGroot, Eds. New York: McGraw-Hill, 1988.
- [81] R. Hasegawa and M. Amamiya, "Parallel execution of logic programs based on dataflow concept," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 507-516.
- [82] B. Hayes-Roth, "A blackboard architecture for control," *Artificial Intell.*, vol. 26, no. 3, pp. 251-321, July 1985.
- [83] C. Hewitt, "Viewing control structure as patterns of passing messages," *Artificial Intell.*, vol. 8, no. 3, pp. 323-364, 1977.
- [84] C. Hewitt, "The apary network architecture for knowledgeable systems," in *Conf. Rec. Lisp Conf.*, Stanford Univ., Menlo Park, CA, 1980, pp. 107-117.
- [85] C. Hewitt and H. Lieberman, "Design issues in parallel architectures for artificial intelligence," in *Proc. IEEE COMPCON Spring*, Feb. 1984, pp. 418-423.
- [86] M. Hill et al., "Design decisions in SPUR," *IEEE Computer*, vol. 19, no. 11, pp. 8-22, Nov. 1986.
- [87] W. D. Hillis, *The Connection Machine*. Cambridge, MA: MIT Press, 1985.
- [88] B. K. Hillyer and D. E. Shaw, "Execution of OPS5 production systems on a massively parallel machine," *J. Parallel Distributed Computing*, vol. 3, no. 2, pp. 236-268, 1986.
- [89] G. E. Hinton, T. J. Sejnowski, and D. H. Ackley, "Boltzmann machine: Constraint satisfaction network that learns," *Carnegie-Mellon Univ., Pittsburgh, PA*, 1984.
- [90] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Nat. Acad. Sci.*, vol. 79, pp. 2554-2558, 1982.
- [91] J. J. Hopfield and D. W. Tank, "Neural computation of decisions in optimization problems," *Biol. Cybern.*, vol. 52, no. 3, pp. 1-25, July, 1985.
- [92] D. K. Hsiao, Ed., Special Issue on Database Machines, *IEEE Computer*, vol. 12, no. 3, Mar. 1979.
- [93] K. Hwang, R. Chowkanyun, and J. Ghosh, "Computer architectures for implementing AI systems," in *Supercomputers and AI Machines*, K. Hwang and D. DeGroot, Eds. New York: McGraw-Hill, 1988.
- [94] N. Ichijoshi, T. Miyazaki, and K. Taki, "A distributed implementation of flat G1IC on the multi-PSI," presented at the Int. Conf. Logic Programming, 1987.
- [95] K. B. Irani and Y. F. Shih, "Implementation of very large prolog-based knowledge bases on data flow architectures," in *Proc. 1st IEEE Conf. Artificial Intell. Applications*, Dec. 1984, pp. 454-459.
- [96] Y. Ishikawa and M. Tokoro, "The design of an object-oriented architecture," in *Proc. 11th IEEE/ACM Int. Symp. Computer Architecture*, 1984, pp. 178-187.
- [97] N. Ito, H. Shimizu, M. Kishi, E. Kuno, and K. Rokusawa, "Data-flow based execution mechanisms of parallel and concurrent prolog," *New Generation Computing*, vol. 3, pp. 15-41, 1985.
- [98] P. Kanerva, "Parallel structures in human and computer memory," NASA Ames Research Center, Moffett Field, CA, RIACS Tech. Rep. TR-86.2, Jan. 1986.
- [99] K. Kawanobe, "Current status and future plans of the fifth generation computer system project," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 3-36.
- [100] R. M. Keller and F. C. H. Lin, "Simulated performance of a reduction-based multiprocessor," *IEEE Computer*, vol. 17, no. 7, pp. 70-82, July 1984.
- [101] R. M. Keller, F. C. H. Lin, and J. Tanaka, "Rediflow multiprocessing," in *Proc. IEEE COMPCON Spring*, 1984, pp. 410-417.
- [102] C. Kellogg, "Intelligent assistants for knowledge and information resources management," in *Proc. 8th Int. Joint Conf. Artificial Intell.*, Los Altos, CA: William Kaufman, 1983, pp. 170-172.
- [103] S. Kim, S. Maeng, and J. W. Cho, "A parallel execution model of logic program based on dependency relationship graph," in *Proc. Int. Conf. Parallel Processing*, Aug. 1986, pp. 976-983.
- [104] G. Klinker, E. Clune, J. Crisman, and J. Webb, "The implementation of a complex vision system on systolic array machine," *Dep. Comput. Sci., Carnegie-Mellon Univ., Pittsburgh, PA*, Tech. Rep., May 1986.
- [105] W. E. Kluge, "Cooperating reduction machines," *IEEE Trans. Computers*, vol. C-32, no. 11, pp. 1002-1012, Nov. 1983.
- [106] T. Knight, "The CONS Microprocessor," *Mass. Inst. Technol., Cambridge, AI Working Paper 80*, Nov. 1974.
- [107] A. Koster, "Compiling Prolog programs for parallel execution on a cellular Machine," in *Proc. ACM'84 Annu. Conf.*, Oct. 1984, pp. 167-178.
- [108] D. J. Kuck, E. S. Davidson, D. H. Lawrie, and A. H. Sameh, "Parallel supercomputing today and the cedar approach," *Science*, pp. 967-974, Feb. 1986.
- [109] G. G. Langdon Jr., Ed., Special issue on database machines, *IEEE Trans. Comput.*, vol. C-28, no. 6, June 1979.
- [110] D. B. Lenat and J. McDermott, "Less than general production system architectures," in *Proc. 5th Int. Joint Conf. Artificial Intell.*, Los Altos, CA: William Kaufman, 1977, pp. 923-932.
- [111] D. B. Lenat, "Computer software for intelligent systems," *Sci. Amer.*, vol. 251, no. 3, pp. 204-213, Sept. 1984.
- [112] V. R. Lesser and D. D. Corkill, "The distributed vehicle monitoring testbed: A tool for investigating distributed problem solving networks," *AI Mag.*, pp. 15-33, Fall 1983.
- [113] G. J. Li and B. W. Wah, "Computational efficiency of parallel approximate branch-and-bound algorithms," in *Proc. IEEE Int. Conf. Parallel Processing*, Aug. 1984, pp. 473-480.
- [114] ———, "MANIP-2: A multicomputer architecture for evaluating logic programs," in *Proc. IEEE Int. Conf. Parallel Processing*, Aug. 1985, pp. 123-130. (Also in *Tutorial: Computers for Artificial Intell. Applications*, B. W. Wah, Ed., New York: IEEE Computer Society, 1986, pp. 392-399.)
- [115] ———, "Optimal granularity of parallel evaluation of AND-trees," in *Proc. Fall ACM/IEEE Joint Computer Conf.*, Nov. 1986, pp. 297-306.
- [116] ———, "Coping with anomalies in parallel branch-and-bound algorithms," *IEEE Trans. Computer*, vol. C-34, no. 6, pp. 568-573, June 1986.
- [117] ———, "How good are parallel and ordered depth-first searches?" in *Proc. IEEE Int. Conf. Parallel Processing*, Aug. 1986, pp. 992-999.
- [118] Y. J. Lin and V. Kumar, "A parallel execution scheme for exploiting AND-parallelism of logic programs," *Artificial Intell.*, pp. 972-975, Aug. 1986.
- [119] G. Lindstrom and P. Panangaden, "Stream-based execution of logic programs," in *Proc. IEEE Int. Symp. Logic Programming*, Feb. 1984, pp. 168-176.
- [120] G. Mago, "Making parallel computation simple: The FFP machine," in *Proc. IEEE COMPCON Spring*, 1985, pp. 424-428.

- [121] G. Matthews, R. Hewes, and S. Krueger, "Single-chip processor runs Lisp environments," *Comput. Design*, pp. 69-76, May 1, 1987.
- [122] J. McCarthy, "History of Lisp," *SIGPLAN Notices*, vol. 13, no. 8, pp. 217-223, 1978.
- [123] J. R. McGraw, "Data flow computing: Software development," *IEEE Trans. Comput.*, vol. C-29, no. 12, pp. 1095-1103, 1980.
- [124] D. McKay and S. Shapiro, "MULTI—A Lisp based multiprocessing system," in *Conf. Rec. Lisp Conf.*, Stanford Univ., Menlo Park, CA, 1980.
- [125] M. Model, "Multiprocessing via intercommunicating Lisp systems," in *Conf. Rec. Lisp Conf.*, Stanford Univ., Menlo Park, CA, 1980.
- [126] D. I. Moldovan, "An associative array architecture intended for semantic network processing," in *Proc. ACM'84 Annu. Conf.*, Oct. 1984, pp. 212-221.
- [127] D. A. Moon, "Symbolics architecture," *IEEE Computer*, vol. 20, no. 1, pp. 43-52, Jan. 1987.
- [128] K. Murakami, T. Kakuta, and R. Onai, "Architectures and Hardware Systems: Parallel Inference Machine and Knowledge Base Machine," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 18-36.
- [129] K. Murakami, T. Kakuta, R. Onai, and N. Ito, "Research on parallel machine architecture for fifth-generation computer systems," *IEEE Computer*, vol. 18, no. 6, pp. 76-92, June 1985.
- [130] M. Nagao, J. I. Tsujii, K. Nakajima, K. Mitamura, and H. Ito, "Lisp machine NK3 and measurement of its performance," in *Proc. 6th Int. Joint Conf. on Artificial Intell.*, Los Altos, CA: William Kaufman, Aug. 1979, pp. 625-627.
- [131] P. M. Neches, "Hardware support for advanced data management systems," *IEEE Computer*, vol. 17, no. 11, pp. 29-40, Nov. 1984.
- [132] K. Niwa, K. Sasaki, and H. Ihara, "An experimental comparison of knowledge representation schemes," *AI Mag.*, pp. 29-36, Summer 1984.
- [133] J. T. O'Donnell, "A systolic associative Lisp computer architecture with incremental parallel storage management," Ph.D. dissertation, Univ. of Iowa, Iowa City, 1981.
- [134] K. Ofizer, "Partitioning in parallel processing of production systems," in *Proc. IEEE Int. Conf. Parallel Processing*, 1984, pp. 92-100.
- [135] J. Pavlin, "Predicting the performance of distributed knowledge-based systems: A modeling approach," in *Proc. Nat. Conf. Artificial Intelligence*, 1983, pp. 314-319.
- [136] J. Pearl, *Heuristics—Intelligent Search Strategies for Computer Problem Solving*. Reading, MA: Addison-Wesley, 1984.
- [137] L. M. Pereira and R. Nasr, "Delta-Prolog, A distributed logic programming language," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 283-291.
- [138] A. Plotkin and D. Tabak, "A tree structured architecture for semantic gap reduction," *Computer Architecture News*, vol. 11, no. 4, pp. 30-44, Sept. 1983.
- [139] E. von Puttkamer, "A microprogrammed Lisp machine," *Microprocessing Microprogramming*, vol. 11, no. 1, pp. 9-14, Jan. 1983.
- [140] U. S. Reddy, "On the relationship between logic and functional languages," in *Logic Programming*, ed. D. DeGroot and E. G. Lindstron, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1985.
- [141] T. Rentsch, "Object oriented programming," *SIGPLAN Notices*, vol. 17, no. 9, pp. 51-57, Sept. 1982.
- [142] J. Robinson and E. Sibert, "LOGLISP: Motivation, design, and implementation," in *Logic Programming*, K. Clark and S. Tarnlund, Eds. New York: Academic, 1982.
- [143] H. Sakai et al., "Design and implementation of relational database engine," in *Proc. 5th Generation Computer Systems*, 1984, pp. 419-426.
- [144] J. Sansonnet, D. Botella, and J. Perez, "Function distribution in a list-directed architecture," *Microprocessing Microprogramming*, vol. 9, no. 3, pp. 143-153, 1982.
- [145] J. P. Sansonnet, M. Castan, and C. Percebois, "M3L: A list-directed architecture," in *Proc. 7th Annu. IEEE/ACM Symp. Computer Architecture*, May 1980, pp. 105-112.
- [146] J. P. Sansonnet, M. Castan, C. Percebois, D. Botella, and J. Perez, "Direct execution of Lisp on a list-directed architecture," in *Proc. ACM Symp. Architectural Support for Programming Languages and Operating Systems*, Mar. 1982, pp. 132-139.
- [147] M. Sato, H. Shimizu, A. Matsumoto, K. Rokusawa, and A. Goto, "K11 execution model for PIM cluster with shared memory," presented at the Int. Conf. Logic Programming, 1987.
- [148] H. Schmeck and H. Schroder, "Dictionary machines for different models of VLSI," *IEEE Trans. Comput.*, vol. C-34, no. 5, pp. 472-475, May 1985.
- [149] M. Schor, "Declarative knowledge programming: Better than procedural," *IEEE Expert*, vol. 1, no. 1, pp. 36-43, Spring 1986.
- [150] E. Shapiro and A. Takeuchi, "Object oriented programming in concurrent Prolog," *New Generation Computing*, vol. 1, no. 1, pp. 25-48, 1983.
- [151] E. Shapiro, "Systolic programming: A paradigm of parallel processing," in *Proc. Int. 5th Generation Computer Systems*, 1984, pp. 458-470.
- [152] D. E. Shaw, "On the range of applicability of an artificial intelligence machine," *Artificial Intell.*, vol. 32, pp. 151-172, 1987.
- [153] S. Shibayama, T. Kakuta, N. Miyazaki, H. Yokota, and K. Murakami, "A relational database machine with large semiconductor disk and hardware relational algebra processor," *New Generation Computing*, vol. 2, no. 2, pp. 131-155, 1984.
- [154] B. Silver, *Meta-Level Inference: Representing and Learning Control Information in Artificial Intelligence*. Studies in CS and AI. Amsterdam, The Netherlands: North-Holland, 1986.
- [155] H. A. Simon, "Whether software engineering needs to be artificially intelligent," *IEEE Trans. Software Eng.*, vol. SE-12, no. 7, July 1986.
- [156] K. Smith, "New computer breed uses transputers for parallel processing," *Electronics*, pp. 67-68, Feb. 24, 1983.
- [157] R. G. Smith and R. Davis, "Frameworks for cooperation in distributed problem solving," *IEEE Trans. Syst. Man Cybern.*, vol. SMC-11, no. 1, pp. 61-70, Jan. 1981.
- [158] A. Snyder, "Object-oriented programming for common Lisp," Software Technology Lab., Hewlett-Packard Lab., Palo Alto, CA, Rep. ATC-85-1, 1985.
- [159] G. Steel and G. Sussman, "Design of Lisp-based processor, or SCHEME: A dielectric Lisp or finite memories considered harmful, or LAMBDA: The ultimate opcode," Mass. Inst. Technol., Cambridge, AI Memo 514, March 1979.
- [160] G. L. Steele, Jr., and G. J. Sussman, "Design of a Lisp-based microprocessor," *Comm. Assoc. Comput. Mach.*, vol. 23, no. 11, pp. 628-645, Nov. 1980.
- [161] M. Stefik and D. G. Bobrow, "Object-oriented programming: Themes and variations," *AI Mag.*, pp. 40-62, Spring 1986.
- [162] S. J. Stofo, "Initial performance of the DADO2 prototype," *IEEE Computer*, vol. 20, no. 1, pp. 75-84, Jan. 1987.
- [163] S. Sugimoto, K. Tabata, K. Agusa, and Y. Ohno, "Concurrent Lisp on a multi-micro-processor system," in *Proc. 7th Int. Joint Conf. Artificial Intell.*, Los Altos, CA: William Kaufman, Aug. 1981, pp. 949-954.
- [164] S. Sugimoto, K. Agusa, K. Tabata, and Y. Ohno, "A multi-micro-processor system for concurrent Lisp," in *Proc. IEEE Int. Conf. Parallel Processing*, 1983, pp. 135-143.
- [165] G. Sussman, T. Winograd, and E. Charniak, "Micro-planner reference manual," Mass. Inst. Technol., Cambridge, Tech. Rep. AIM-203, 1970.
- [166] G. J. Sussman, J. Holloway, G. L. Steel, Jr., and A. Bell, "Scheme-79—Lisp on a chip," *IEEE Computer*, vol. 14, no. 7, pp. 10-21, July 1981.
- [167] N. Suzuki, K. Kubota, and T. Aoki, "SWARD32: A bytecode emulating microprocessor for object-oriented languages," in *Proc. Int. Conf. 5th Generation Computer Systems*, pp. 389-397, 1984.
- [168] A. Takeuchi and K. Fukukawa, "Parallel logic programming languages," in *Proc. 3rd Int. Conf. Logic Programming*. New York: Springer-Verlag, 1986.
- [169] I. Takeuchi, H. Okuno, and N. Ohsato, "TAO—A harmonic mean of Lisp, prolog, and smalltalk," *SIGPLAN Notices*, vol. 18, no. 7, pp. 65-74, July 1983.
- [170] K. Taki, Y. Kaneda, and S. Maekawa, "The experimental Lisp machine," in *Proc. 6th Int. Joint Conf. Artificial Intell.*, Los Altos, CA: William Kaufman, Aug. 1979, pp. 865-867.
- [171] K. Taki et al., "Hardware design and implementation of the personal sequential inference machine (PSI)," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 398-409.
- [172] K. Taki, "The parallel software research and development tool: Multi-PSI system," in *Proc. France-Japan Artificial Intell. and Computer Science Symp.*, 1986, pp. 365-381.
- [173] Y. Tanaka, "MPIDC-massive parallel architecture for very large databases," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 113-137.

- [174] E. Tick and D. H. D. Warren, "Towards a pipelined Prolog processor," *New Generation Computing*, vol. 2, no. 4, pp. 323-345, 1984.
- [175] P. Treleaven and G. Mole, "A multi-processor reduction machine for user-defined reduction languages," in *Proc. 7th IEEE/ACM Int. Symp. Computer Architecture*, pp. 121-130, 1980.
- [176] P. C. Treleaven and R. P. Hopkins, "A recursive computer architecture for VLSI," in *Proc. 9th Annu. IEEE/ACM Symp. Computer Architecture*, Apr. 1982, pp. 229-238.
- [177] P. C. Treleaven and I. G. Lima, "Japan's fifth-generation computer systems," *IEEE Computer*, vol. 15, no. 8, pp. 79-88, Aug. 1982.
- [178] Y. W. Tung and D. Moldovan, "Detection of AND-parallelism in logic programming," in *Proc. Int. Conf. Parallel Processing*, pp. 984-991, Aug. 1986.
- [179] D. A. Turner, "A new implementation technique for applicative languages," *Software—Practice and Experience*, vol. 9, no. 1, pp. 31-49, 1979.
- [180] S. Uchida, "Inference machines in FGCS project," in *Proc. VLSI'87 Int. Conf.*, Aug. 1985, IFIP TC-10, WG 10.5.
- [181] K. Ueda, "Guarded horn clauses," ICOT, Tokyo, Japan, Tech. Rep. TR-103, 1985.
- [182] L. M. Uhr, "Parallel-serial production systems," in *Proc. 6th Int. Joint Conf. Artificial Intell.* Los Altos, CA: William Kaufman, Aug. 1979, pp. 911-916.
- [183] D. Ungar, R. Blau, P. Foley, D. Samples, and D. A. Patterson, "Architecture of SOAR: Smalltalk on RISC," in *Proc. 11th Annu. IEEE/ACM Int. Symp. Computer Architecture*, 1984, pp. 188-197.
- [184] S. R. Vegdahl, "A survey of proposed architectures for the execution of functional languages," *IEEE Trans. Comput.*, vol. C-33, no. 12, pp. 1050-1071, Dec. 1984.
- [185] B. W. Wah and K. L. Chen, "A partitioning approach to the design of selection networks," *IEEE Trans. Comput.*, vol. C-33, no. 3, pp. 261-268, March 1984.
- [186] B. W. Wah and Y. W. Ma, "MANIP—A multicompiler architecture for solving combinatorial extremum problems," *IEEE Trans. Comput.*, vol. C-33, no. 5, pp. 377-390, May 1984. (Also in *Tutorial: Computer Architecture*, D. D. Gajski, V. M. Milutinovic, H. J. Siegel, and B. P. Furth, Eds. IEEE Computer Soc., 1987, pp. 578-591 and *Tutorial: Parallel Architecture for Database Systems*, A. R. Hurson, L. L. Miller, and S. H. Pakzad, Eds. IEEE Computer Soc., 1988.)
- [187] B. W. Wah, G. J. Li, and C. F. Yu, "Multiprocessing of combinatorial search problems," *IEEE Computer*, vol. 18, no. 6, pp. 93-108, June 1985. (Also in *Tutorial: Computers for Artificial Intell. Appl.*, B. W. Wah Ed. IEEE Computer Soc., 1986, pp. 173-188.)
- [188] B. W. Wah and G. J. Li, Eds., *Tutorial on Computers for Artificial Intell. Appl.* New York: IEEE Computer Society Press, May 1986.
- [189] D. L. Waltz, "Applications of the connection machine," *IEEE Computer*, vol. 20, no. 1, Jan. 1987.
- [190] D. H. D. Warren, "Efficient processing of interactive relational database queries expressed in logic," in *Proc. 7th Int. Conf. Very Large Data Bases*, 1981, pp. 272-281.
- [191] P. Wegner and B. Shriver, Eds., "Special issue on object-oriented programming workshop," *SIGPLAN Notices*, vol. 21, no. 10, Oct. 1986.
- [192] M. Weiser et al., "Status and performance of the ZMOB parallel processing system," in *Proc. IEEE COMPCON Spring*, Feb. 1985, pp. 71-73.
- [193] R. Williams, "A multiprocessing system for the direct execution of Lisp," in *Proc. ACM 4th Workshop Computer Architecture for Non-Numeric Processing*, Aug. 1978.
- [194] T. Winograd, "Extended inference modes in reasoning by computer systems," in *Artificial Intell.*, vol. 13, pp. 5-26, 1980.
- [195] P. H. Winston and B. Horn, *Lisp*, 2nd ed. Reading, MA: Addison-Wesley, 1984.
- [196] Y. Yamaguchi, K. Toda, and T. Yuba, "A performance evaluation of a Lisp-based data-driven machine (EM-3)," in *Proc. 10th Annu. IEEE/ACM Int. Symp. Computer Architecture*, June 1983, pp. 363-369.
- [197] Y. Yamaguchi, K. Toda, J. Herath, and T. Yuba, "EM-3: A Lisp-based data-driven machine," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 524-532.
- [198] T. Yokoi, S. Uchida, "Sequential inference machine: SIM—Its programming and operating system," in *Proc. Int. Conf. 5th Generation Computer Systems*, 1984, pp. 70-81.
- [199] C. F. Yu and B. W. Wah, "Efficient branch-and-bound algorithms on a two-level memory system," *IEEE Trans. Software Eng.*, vol. SE-14, no. 9, Sept. 1988.
- [200] —, "Learning dominance relations in combinatorial search problems," *IEEE Trans. Software Eng.*, vol. SE-14, no. 8, Aug. 1988.



Benjamin W. Wah (S'74-M'79-SM'85) received the Ph.D. degree in computer science from the University of California, Berkeley, CA, in 1979.

He was on the faculty of the School of Electrical Engineering at Purdue University, West Lafayette, IN, between 1979 and 1985. He is now a Professor in the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory of the University of Illinois at Urbana-Champaign. Between 1988 and 1989, he was on leave at the National Science Foundation as a Program Director in the Microelectronic Information Processing Systems Division. His areas of research include computer architecture, parallel processing, artificial intelligence, distributed databases, and computer networks. For his contributions to research, he has been nominated as a University Scholar of the University of Illinois in 1989.

Dr. Wah is the Associate Editor-in-Chief of the IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, an area editor of the *Journal of Parallel and Distributed Computing*, and an editor of *Information Sciences*. He serves as a member of the Governing Board of the IEEE Computer Society and a program evaluator for ABET (computer engineering) and CSAC (computer science). Previously, he served as chairman and member of program committee of a number of international conferences, an Editor of the IEEE TRANSACTIONS ON SOFTWARE ENGINEERING, and a Distinguished Visitor of the IEEE Computer Society.

Guo-jie Li (S'83-M'86) graduated from Peking University, Beijing, China, in 1968. He received the M.S. degree in computer science and engineering from the University of Science and Technology of China and Institute of Computing Technology, Chinese Academy of Science, and the Ph.D. degree in electrical engineering from Purdue University in 1981 and 1985, respectively.

He was a Post-Doctoral Research Fellow at the University of Illinois, Urbana-Champaign, between 1985 and 1986. Currently, he is an Associate Researcher at the Academy of Science, Beijing, China. His research interests include parallel processing, computer architecture, and artificial intelligence.