

STANFORD ARTIFICIAL INTELLIGENCE PROJECT
MEMO AIM-146

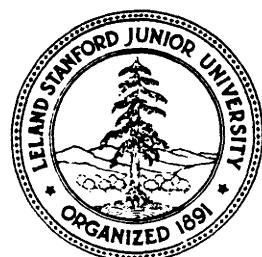
COMPUTER SCIENCE DEPARTMENT
REPORT NO. CS-224

PARALLEL PROGRAMMING

BY
ANDREI P. ERSHOV

JULY 1971

COMPUTER SCIENCE DEPARTMENT
STANFORD UNIVERSITY



Parallel Programming

by

Andrei P. Ershov

ABSTRACT: This report is based on lectures given at Stanford University by Dr. Ershov in November, 1970,

Introduction

The author is deeply convinced that programming problems are decisive in the overall design and utilization of parallel computing systems. Of course, the engineer using his imagination can fairly rapidly propose one or another design for such a computing system and even bring it into being; he does this through analogy or, conversely, through the method of selecting a contrasting design. This approach, however, necessarily leads to a situation in which programming problems become secondary, and the degree of freedom available in solving them is automatically limited by a prescribed structure that is frequently not responsive to essential points.

The problem of programming is of essentially secondary importance when one speaks of a particular task that must be solved on a given machine. However, the problem of programming is of primary importance in any thoughtful approach to the design of a computing system, in the sense that the developer must by the strength of his own penetration into the essence of the matter mentally include in the design the complete set of all possible programs and algorithms. On the basis of an analysis of their internal properties, structures, and external characteristics, he formulates design principles for the computing system. These principles must be adequate for the examined properties of the programs, and--most important--must be appropriate for reducing the programming task to a formal procedure and for automating it.

Automation of programming for Parallel computing systems is a basic necessity for a number of reasons.

First, the description of a parallel program is not characteristic of the nature of man's algorithmic thought processes. In fact, this is true only to the extent that we are considering the implementation on a computing system of a task specified in algorithmic form. Today this is the most common way of stating a problem; however, in the future this could change, since parallelism in a problem is frequently implicit in its initial formulating and is artificially driven out by algorithmization. This will be discussed in more detail below.

Second, Independent of whether man thinks "sequentially" or "in parallel," he tends to think laconically--to seek compactness in problem formulation. At the same time, studies on the design of parallel programs indicate that in many cases complete parallelism can be achieved only under conditions of potentially unlimited "multiplication" of some initial program constructions, where the multiplication is associated with systematic modification of the constructions. Thus, some portion of the programming process becomes an organic part of the computational process and therefore must necessarily be transferred to the machine.

In this work an attempt will be made to demonstrate the pre-eminence of programming over hardware design of a computing system. The concepts of parallel programming established at the present time will be examined, and on this basis an evaluation will be made of several approaches to the structure of computing systems. Due to the general lack of development of the problem, this analysis will be superficial and in places even speculative. However, the fog that obscures our ultimate goal can be dissipated only if we boldly step out into it. On the other hand, an attempt will be made to put forth some recommendations from a realistic position, taking into consideration contemporary technical capabilities, historically established frameworks, and other current limitations. Finally, the paper is limited in the sense that it is not a balanced, inclusive survey, but only reflects the personal viewpoint of the author, nurtured however over a long period of time.

Basic Concepts Of Parallel Programming

In any theory of parallel programming aspiring to completeness, there are three components that must be present:

- a way of describing a parallel program (parallel programming facilities);
- a way of deriving a parallel program from ordinary algorithmic notation (desequencing of sequential algorithms);
- a way of executing a parallel program (specification of the computing system, assignment of processors to jobs, control of the parallel computational process),

The recommended structure for a computing system must be derived from basic principles reflecting the most essential aspects of each of these components.

Three approaches to parallel programming can be examined, and will be conditionally designated as follows:

- natural Parallelism,
- concurrent computations in algorithmic languages,
- asynchronous programming-,

Natural parallelism is far from what might be somehow considered a comprehensive theory. However, some of its fundamental considerations are quite clear and in many respects highly attractive and distinctive, giving us the right to regard it as a basic approach. This theory does not have a specific advocate, although it is carried on the wind and appears tacitly in many works. Therefore, it is difficult to clearly specify sources for it,

The basic thesis of this theory is that in nature everything in fact takes place in parallel, and the more literally some actual process is described, the closer such a description is to a parallel program. The general problem is that one must be able to successfully and simply capture this natural parallelism and describe

it in some universal System of concepts, based on such principles as simultaneity of many events, proximity of actions, etc., i.e., a system of concepts in which a phenomenon is represented as a set of many elementary processes occurring in a space simultaneously and interacting only with neighboring processes. The principle of determinism (Indeed understood in a broad sense) is the basis of this approach; it says in the given case that if a phenomenon at some moment is described as a set of initial states and mechanism of interaction, then the entire subsequent history is predetermined and can be observed, measured, and computed. The space referred to does not necessarily have to be physical and continuous—it can be discrete, finite, etc.; it is only important that the concept of the immediate neighborhood be determined.

This principle is quite well known and in fact we apply it in many situations (in particular, all methods of describing phenomena in differential form are based on it). The real and far from trivial problem is to constructively formulate this principle in a universal form representing all the essentials that must be embodied in parallel programs, in view of the diverse computational universes in which informational models of phenomena of interest to us may be embedded,

On the whole, programming under such an approach disappears as a separate task and becomes an inseparable constituent and maybe even the whole of the Problem of describing the initial situation for the given phenomenon.

Under a pure formulation of this question, the problem of translating "ordinary" notation into a Parallel one also does not arise, since natural parallelism is primary and the concept of "ordinary" notation loses its meaning. However, a pure formulation of the question is not possible, since the entire computational aspect concerned with the design of the arithmetic model of the phenomenon bears imprints of the inherently sequential processes of logical reasoning which are the foundation of any algorithm. Therefore, any systematic attempt to construct a direct bridge from natural Parallelism to its implementation in a computing system requires an essential revision of all numerical analysis and the apparatus of mathematical physics. An additional and no less vast problem is the design of such a system of concepts in which any interaction occurring in nature is described as a process of information exchange, and in which the interacting element of a space is represented as a miniature, universal computer with many input and output lines doing the required information Processing,

If we turn to the third component of natural parallelism, then it is immediately necessary to make a hypothesis about an appropriate computing system. It is perfectly obvious that a computer environment [1] is the first candidate for such a system. A consideration of such an environment properly is deferred until the conclusion of this paper, but it should be mentioned here that if there is enough computing power, we can associate a separate unit of the environment with each element of the space, all control and

implementation of the computational process reduces to the construction and start-up of the environment. The environment reproduces the described phenomenon almost "one-for-one", and thereby frees us of concerns about organizing the computational process.

However, the problem becomes substantially more complex if the parallel process consists of just one element more than the available number of units of the environment. Very complex problems of accounting for boundary effects and of rearranging the environment during implementation of the computational process then arise.

Concluding this consideration of natural parallelism, we note that we began with a statement as to the attractiveness of the approach and then turned attention to the fact that with its some problems disappear. However, the entire remaining discussion was directed towards discrediting the approach. The discreditation lies in the fact that for development of the computational environment as a universal approach it is too revolutionary, breaking with many past concepts and leading to many problems on the level of fantasy; therefore, it cannot be recommended as a guide for immediate action.

At the same time, the author hopes that this discreditation does not result in any weakening of the search for a general-purpose scheme for introducing natural parallelism into programming. In particular, this must not occur because successful synthesis of natural parallelism and a computational environment in efficiently operating technical devices would have a decisive result on research into the structure and operating mechanisms of the brain--the most perfect computational environment created by nature.

Concurrent computations in algorithmic languages, in comparison to the approach just examined, lower us from the skies to the earth. This approach represents a somewhat minimal attempt to introduce parallelism into algorithmic languages, fully taking into account existing technology, contemporary machine architecture, and the organization of operating systems. The approach puts aside many questions of parallel programming, but in return it permits us to use, even if only partially, those capabilities that today's technology represents.

The descriptive facilities for denoting parallelism, about which we speak in this section, are found in such languages as those for simulation (SIMSCRIPT [2], SIMULA [3]) and the "new generation" of algorithmic languages--PL/1 [4], ALGOL-68 [5]. The existence of these parallel, or concurrent, computational branches are clearly explicated in an algorithm. Branches parallel to one another usually have a common beginning and end. Each branch, in turn, can consist of sub-branches. Thus, all program parallelism is explicitly stated, and the points of branching and joining can be determined.

*Translator's note: "environment" is used as a translation of the Russian word "sreda" which is evidently used as a technical term in the paper referred to.

syntactically. Branches can use common variables and can include means for interruption or suspension of branches depending on the value of the variable quantity. Thus, synchronization of concurrent computations is accomplished, the explicit programming of which is a task for programmers,

The described approach still does not offer in its general notation any way of dividing the program into parallel branches, since use of the descriptive facilities of parallelness is not formalized and is completely up to the programmer,

In the majority of cases, the introduction of parallel branches permits parallel execution but does not prescribe it. In the ease of a deviation from full parallelism, a well-written program permits the branches to be executed in arbitrary order, which eases the work of the operating system. However, in truth the existing languages do not give formal rules for defining a "well-formed" program,

The organization of computations for a program with parallel branches takes place according to the Principles of multiprogramming, where the program for solving a large problem is represented to the "eyes" of the operating system supervisors as a stream of related small jobs. One parallel branch is represented as one job. A kind of passport is set up for this job, and stored in the operating system. When the control reaches the branching point, a call (established there by the translator) to the supervisor is executed, which activates the passports of the branches beginning at that point and makes them candidates for processing. The supervisor carries out some planning at the assignment of branches for processing. Its purpose is to maximally load units (or processors) that are able to operate together, and it attempts to achieve completion of all parallel branches as rapidly as possible, since only a concurrent output at the convergence point will permit an advance to the next point in the program.

This method is practical only when points of branching and convergence occur relatively infrequently. This at once imposes a serious limitation on the degree of parallelism and on the number of permissible branches. Another limitation is that the parallel structure of the program is static and in fact does not permit, for example, expansion of loops.

Asynchronous programming, from the author's viewpoint, has the most right to be called a theory of Parallel programming, and therefore will be considered more thoroughly. The first results of the six-year work of the young mathematicians V.E. Kotov and A.S. Narin'yaniform the fundamentals of this theory [6-9]. The work of the American mathematicians Karp and Miller [10] is related to this direction. Asynchronous programming from the very beginning has developed as a mathematical theory, based on certain axiomatics. And in this lies its merit and its imperfections. No computing systems have yet been designed or translators created according to the principles of asynchronous programming. But the theory of

asynchronous programming promises to:

- provide a precise concept of a parallel (asynchronous) program;
- describe a class of computing systems;
- rigorously define program equivalence;
- formulate criteria of the correctness of parallel programs for a given system;
- define degrees of parallelism and maximum parallelism;
- constructively prove the possibility of formal translation of an ordinary program into a parallel program possessing maximum parallelism;
- study certain internal problems of parallel programming.

An additional characteristic of asynchronous programming is that in its theoretical aspect it is linked with the already rather well-developed theory of Sequential Programs--Program Schemes [11]. This association is still not fully developed, but the prospects of doing so are clear,

we shall examine the fundamental concepts of asynchronous programming. The initial concept is that of a quasiprogram, which is represented as an arbitrary set of statements operating on a memory which consists of variables. A statement is a multi-pole "black box" which assigns output values to its output lines as functions of its inputs. Associated with each statement input and output pole are variables. The input variables supply the statement with values from memory, and the output variables accept the outputs of the statement. Each statement can be in one of three states: dormant, ready, or functioning.

A quasiprogram is executed in a computing system in the following manner. At the moment of start-up, the initial state of memory is assigned. All statements are dormant. The operation of the system consists of a succession of changes in the states of the computation at separate, discrete moments of time. At each such moment, the computing system can transfer statements from one state to another, with only a ready statement able to become a functioning one,

If a statement becomes functioning, it receives at that moment from memory values of arguments. If a statement stops functioning, it transmits to memory at that moment the values of computed results,

The progress of the computations can be depicted graphically in the form of a computational process, represented as a time diagram of the switching statements on and off, whereby "on" is understood as transfer to the functioning state, and by "off" is understood exit from the functioning state. An interval of the time diagram spanning moments of on and off switchings reflects one statement action (Fig. 1).₁

From this computational process it is possible to

unambiguous, plot its important characteristics--to construct a n information-logic graph. This oriented graph is without cycles, and its vertices are the statement actions (Fig. 2). From vertex A there is an arc to vertex B if at action B some of its arguments receive information assigned by some result of action A, or if action A in a direct way has an effect on action B. The concept of a direct effect of one statement on another is quite complexly determined, but the example of Fig. 3 gives some impression of its nature. Here, statement S_1 directly affects statement S_2 , whereas S_1 has no effect on S_3 . Statement S_0 affects S_2 only indirectly, and S_1 and S_3 directly,

The information-logic graph, on the one hand, retains some necessary minimum of information that makes it possible to re-establish the way of processing the input data of the computational process into the final results. On the other hand, it ignores less essential details of the progress of the computational process,

The fundamental characteristic of asynchronous programming is the assumption of non-uniqueness of the execution of the computing system to a quasiprogram. Generally speaking, it is assumed that the system has some parameters or degrees of freedom that are not fixed by the quasiprogram. This means that for a given quasiprogram for some initial memory state, the system can implement some set, possibly infinite, of computational processes.

A quasiprogram is called a program for a given computing system if for any assigned initial memory state all computational processes implemented for it by the system have the same information-logic graph. Thus, an information-logic graph is invariant, guaranteeing accuracy of information processing under any behavior of the computing system executing the given program,

The computing system's greater degree of freedom is due to the following. It is assumed that each statement in a program is equipped with a predicate that has some inputs from memory; this predicate is called a trigger function. The statement and its trigger function are called a block. In executing the program, the system continuously computes the values of all trigger functions of the program blocks. Block9 with trigger functions equal to 1 are considered to be ready, while the others are dormant. At any moment in time, the system can switch any ready blocks on or switch any functioning ones off. Computing systems of this type are called asynchronous systems,

Thus, we see that asynchronous programs have a very important quality: protection from arbitrariness that the system may exhibit with respect to the moments of switching of ready operators, the time of their execution, or the amount of computing facilities available for appointment of the ready operators for operation. An asynchronous program does not impose any specific requirements for the computing system regarding its time characteristics or computing power (number of processors). Moreover, system characteristics can

be altered dynamically, without any loss of program run validity.

The concepts of the degree of parallelness, or the degree of asynchronousness, are introduced in asynchronous programming in an interesting manner. Quantitative measures of asynchronousness are not used; however, it is possible to determine that one program is more asynchronous than another, and also to determine the most asynchronous program among a set of programs being compared for asynchronousness.

We shall consider two programs, P_1 and P_2 , equivalent in the sense that for identical initial memory states they generate computational processes with identical information-logic graphs. For each information-logic graph IL there are, correspondingly, sets M_1 and M_2 of the computational processes possessed by such graph IL . Thus, if for any IL , $M_1 \supset M_2$, then P_1 possesses greater asynchronousness. In this way, the program that allows the computing system greater flexibility in the range of computational processes resulting in the given information-graph is recognized as being more asynchronous. There are no explicit statements about more or less parallelism, and this is correct since the degree of parallelness of a program is actually determined not only inherently, but also by the capabilities of the system. However, if program P_1 can be executed by all the same ways as can P_2 , and in addition by still other ways, then there is some chance that it has greater parallelness, and anyway not less.

The degree of asynchronousness also has a more constructive form of determination. The degree of diversity in executing programs in the final analysis rests on the set of binary relations among program statements for determining whether or not restrictions exist on the relative order of statement execution. The sets of these binary relations form a dependence graph. The fewer edges the dependence graph has, the more asynchronous is the program. If the program has a dependence graph such that removal of any edge from it clearly destroys the information-logic connections of some execution of the program, then such a program has maximum asynchronousness.

As already mentioned, the theory of asynchronous programming offers a constructive method for formal transfer from sequential programs, considered in the form of program schemata, to asynchronous programs of maximum parallelness. This transfer takes place in two stages. In the first stage, several equivalent transformations of the scheme itself occur, and then a single transformation of the scheme into an asynchronous program is performed; the latter consists of "splitting" the scheme into the individual statements and of assigning for each statement its trigger function.

It appears that a barrier in the path of converting from a program scheme to a maximally asynchronous program lies in the fact that the internal asynchronousness of a scheme can be proven to be actually less than its potential asynchronousness. This difference results in an unsuccessful allocation of memory, imposing artificial connections between statements, and in imperfect multiplication of

some statements into several instances which, when executed simultaneously, can extend the range of program execution. Two pairs of statements, depicted in Fig. 4, serve as an example of the difference between internal and potential asynchrony. Since the same value y is used in each pair to transfer information, these two pairs of operators cannot be executed concurrently, which would be permissible if in each pair its own variables were used.

In the work of V.E. Kotov, a methodology is offered for transforming any program scheme into an equivalent scheme, the internal asynchrony of which reaches its potential asynchrony.

In the development of a theory of parallel programming, it was unclear earlier whether or not, without sacrificing generality of results, quantitative indicators of the passage of time could be included, threading on the temporal axis only the fact of the occurrence of events, and not their duration. The assumption that all changes in state occur instantaneously, including the actions of statements, is the extreme expression of this abstraction. It has been shown by A.S. Narin'yan that for any information-logic graph there exists an executable computational process in which all statements activate instantaneously (the so-called reduced process). Moreover, it has been shown by him that if a quasiprogram is a program for a computing system with instantaneous action of statements, then it will also be the program for a computing system with any prolonged action of statements.

Continuous computation of trigger functions occupies a significant place in the execution of asynchronous programs. If such a process is to be implemented in a real system, maximum simplification and acceleration of this computation becomes very important. From this viewpoint, the result obtained by Z. Zvinogradskij* is very interesting. It shows that for any asynchronous program it is possible to constructively find its equivalent program such that all trigger functions in that program have the form of logical functions of elementary conditions expressing only the fact of the action of the operator. This makes it possible without loss of generality to attribute to the computation of trigger functions a specific appearance that permits realization using interrupt registers and other high-speed hardware facilities.

* Private communication

Discussion of the Structure of Computing Systems

The use of binary switching and the realization of Boolean functions remain for the foreseeable future the firm foundation of electronics, based in its most direct form on the principle of discrete actions. However, the "embedding" of elementary binary structures within larger structures reflecting an a priori

algorithmic basic of programming must be one of the fundamental principles for achieving pre-eminence of programming over hardware. The author has already mentioned that the established algorithmic basis (arithmetic actions and relations on numbers, associative and addressed information search) has an historical nature and, possibly, is subject to revision from the tenets of natural parallelism. Nevertheless, we do not presently see the possibility of stepping beyond the limits of this algorithmic base, which, to a point, is taken fully into account in contemporary algorithmic languages and the axiomatics of the theory of programming,

Therefore, the author specifically does not recommend homogeneous binary computational environments with dynamic restructuring as the primary candidate for a computational universe. Because of its universality, any such complex will lose out in productivity to a computing system in which the structure of the algorithmic base (accumulators, multipliers, index registers, address matrices, etc.) is introduced a priori at the moment of design and is implemented with the assistance of a rich arsenal of special devices.

In particular, the author considers advisable the preservation of the established separation between the organization of information processing in active Processors and its Storage in a passive store. A large volume of a common memory with random access is the most suitable unit for programming, "absorbing" and decoupling all difficulties concerned with boundedness, with the necessity for fast commutation of processors for some operation in a sequence that was not predictable earlier.

Distribution of memory among processors leads to the necessity of establishing more rigid synchronization. In fact, if processor A transfers results of its operation to another processor B, then this means that the results at A consist of only part of that information which is needed for the operation of B (if this were not so, then A could continue the operation itself, without transferring it to B). But consistent receipt of information by processor B from various parts of the system can lead to a conflict with the desire to utilize as fully as possible the productivity of the processor and to not create downtimes, since generally speaking it is possible to guarantee such productivity only under "equi-accessibility" of any operation to any processor, which contradicts the principle of memory distribution. Another threat to productivity from memory distribution is that in order to provide remote data transmission connections, too many processors must operate only as transit points. Thus, a homogeneous system of many processors with a distributed memory becomes economically justified, in the opinion of the author, only when the utilization factor for the processors may be comparatively low and is measured by roughly the same values as the utilization factor for the core memory locations, determined by the number of location accesses relative to the duration of the storage of information in the location.

Thus, as a primary candidate for the computing system of the immediate future, the author advocates a homogeneous computing system

on a common memory with sufficiently well developed processors structured on an established algorithmic base. This reasoning, however, is in need of some closer examination.

First, the programming for such a system must be based on the principles of **asynchronousness**, since the architecture of this type of system is most adequate to an abstract model of the computing system used in asynchronous programming theory. In addition, program asynchronousness makes very simple the problem of assigning processors to operations, including dynamic assignment, if only the problem of fast scanning of the program's trigger functions is solved.

Second, the system must include the principle of separating strictly computational processes (execution of program statements) from control processes (scanning trigger functions or external interrupt signals and assigning processors to an operation).

It is very tempting to concentrate the control process in a special processor, called here the monitor. Processor general registers must be accessible, on the one hand, to the monitor, while on the other hand the monitor must have in its own memory the complete "logical scheme" of any program presented to it. In its capacity as a generator of a flow of requests for a job to be executed, processing these requests by the method of assigning processors to various jobs constitutes the essential aspect of control.

Such an approach is of interest also because, it seems to the author, it permits the union within the framework of a single system architecture of what would appear to be contrasting types of operations, such as the organization of multiaccess use (i.e., allocation of the computational resources among a flow of weakly connected service requests), and parallel, multiprocessor operation (i.e., concentration of the computational resources for the solution of one large problem). This unity is achieved by the fact that if we examine the "anatomical structure" of the process of controlling the solution of a large problem written in the form of an asynchronous program, then this structure would appear as a limit case of the flow of requests for operations in a time-sharing system.

Separating control out as an individual process, serviced by the monitor, has its own weaknesses. One of these is that critical requirements are imposed on the speed of scanning trigger functions, assigning processors to an operation, and loading processor register stores. Another factor is that centralization of control reduces system reliability. A different scenario is one in which each processor itself searches out its own new job as soon as it finishes its existing operations. This requires the creation in the common memory of a special "labor exchange", access to which is open to every processor. This exchange can be duplicated or can have "branches" in segments of the store, preferably those associated with given processors.

Acknowledging that it is premature to make final judgments on the US8 of centralized or distributed control, the author nevertheless considers that it is necessary to fully carry out the development of methods for centralized control. A consideration of control as a self-contained function, suitable also for special hardware implementation, has particular significance in any objective solution of the problem of the relationship between programming and hardware facilities for exercising control over a computing system. Operational experience with well-developed operating systems for batch batch processing and multiaccess use has indicated that the control processes have their own algorithmic base, the structure of which can and must be embedded in special hardware. This does not exclude, however, the fact that this currently specific base will eventually merge with the "general algorithmic" base. However, this should not be taken for granted; it would have to come about as the result of solving a clearly stated problem.

Conclusion

What, from the viewpoint of the author, must be the fundamental directions for work in the area of programming for parallel computing systems?

First, it is necessary to implement in contemporary algorithmic languages such parallel programming facilities as parallel branches of concurrent computations.

On the basis of the methodology of asynchronous programming, it seems that one could supplement "manual" methods by facilities for automatic determination of parallel branches for reduction of the number of parallel branches, making remaining ones longer, and the checking of a parallel program's asynchronousness. These techniques, by the mid-seventies, could be fully implemented on even operating systems for third-generation multiprocessor computer configurations. Some fundamentals for such techniques can be found already in our existing operating systems [12].

The second direction, which the author believes can provide the richest results by the end of the seventies, must be for a comprehensive design of asynchronous programming as a completed theory, linked with the formal theory of sequential programs in that it would become a working instrument of parallel programming.

Efforts must be concentrated on solution of the following problems:

- the mechanism of computing trigger functions;
- derivation of the "logical scheme" of a parallel program during the translation process and its dynamic augmentation or alteration during program running;
- algorithms for dynamic and static assignment of processors to an operation;
- methodologies for fast commutation of processors during the

expansion of sequential cyclic computations;
--organization of buffering and dynamic loading of a faster store from a slower one;
- study of the capabilities of distributed and centralized controls;
--finding a special algorithmic base for control processors in computing systems.

Successful progress in these directions, in the author's opinion, will help us by the end of the seventies obtain fourth-generation computing systems with productivity ratings greater than 100-million operations per second,

As the third direction, it is necessary to support research on a broad front into computational environments with fundamental study of the following questions:

--the search for the most appropriate universal computational cell and determination of the degree of its connectivity with the environment;

--study of boundary effects in bounded environments and the problem of dynamic restructuring of the environment;

--research into the capabilities and feasibility of building "multiple", mutually penetrating environments with various functional purposes (for example, control and operating, computing and transporting environments);

--development of concrete special-application devices for problems whose structure is in itself sufficient for the selected environment structure;

--development of the principles of natural parallelism and then on that basis a reconsideration of our algorithmic base,

In the opinion of the author, in the eighties computational environments will become a competitor for conquering the problem of increasing the ceiling on the power of computational facilities and the creation of an artificial intelligence,

Acknowledgement

The author deeply thanks Mr. Wad8-B, Holland of RAND Corporation, who kindly agreed to translate this text into English.

References

1. E.V. Evreinov, "On the Microstructure of the Elementary Machines of a Computing System," Computing Systems [Vychislitel'nye sistemy], No. 4, Institute of Mathematics, Siberian Department, USSR Academy of Sciences, Novosibirsk, 1962 [in Russian].
2. H. Markowitz, et al., SIMSCRIPT--An Algorithmic Language for Simulation [SIMSKRIPT--algoritmicheskij jazyk dlya modelirovaniya], Soviet Radio Publishing House, Moscow, 1966 [in Russian translation].
3. O.I. Dahl and K. Nygaard, "SIMULA--A Language for the Programming and Description of Systems with Discrete Events,, Algorithms and Algorithmic Languages [Algoritmy i algoritmicheskie yazyki], No. 2, Computing Center, USSR Academy of Sciences, Moscow, 1967 [in Russian translation].
4. PL/1 General-Purpose Programming Language [Universal'nyj jazyk programmirovaniya PL/1], Mir Publishing House, Moscow, 1968 [in Russian translation].
5. A. Van Wijngaarden, et al., "Report on the ALGOL-68 Algorithmic Language,, Cybernetics [Kibernetika], No. 6, 1969; No. 1, 1970.
6. V.E. Kotov and A.S. Narin'yani, "Asynchronous Computing Processes on a Memory", Cybernetics [Kibernetika], No. 3, 1966 [in Russian].
7. V.E. Kotov and A.S. Narin'yani, "On Transformation of Sequential Programs Into Asynchronous Parallel Programs", Information Processing 68, North-Holland, Amsterdam, 1969,
8. A.S. Narin'yani, "Asynchronous Computing Processes on a Common Memory", Candidate Dissertation, Computing Center, Siberian Department, USSR Academy of Sciences, Novosibirsk, 1970 [in Russian].
9. V.E. Kotov, "Transformation of Operator Schemes Into Asynchronous Programs", Candidate Dissertation, Computing Center, Siberian Department, USSR Academy of Sciences, Novosibirsk, 1970 [in Russian].
10. R.M. Karp and R.X. Miller, "Parallel Program Schemata", J. Comp. and Syst. Sci., Vol. 3, 1969, pp. 147-195.
11. A.P. Ershov and A.A. Lyapunov, "Formalization of the Concept of a Program", Cybernetics [Kibernetika], No. 5, 1967 [in Russian].
12. I.B. Zadykhajlo, et al., "Operating System of the Institute of Applied Mathematics of the USSR Academy of Sciences for the BESM-6 (CS IPM)," Transactions of the Second All-Union Conference on Programming [Trudy 2-j Vsesoyuznoj konferentsii po programmirovaniyu], Session C, Novosibirsk, February 3-6, 1970 [in Russian].