CHAPTER 5. SCIENTIFIC RESEARCH INSTITUTE OF CONTROL COMPUTERS

5.1 Introduction

In this chapter we examine high performance computing developments within the Problem-oriented Computing Systems Division of the Scientfic-Research Institute of Control Computers (NIIUVM) of the Impul's Scientific Production Association in Severodonetsk, Ukraine. This division was created to develop high-performance systems oriented towards problems that are computationally intensive and have a high degree of data parallelism. The reconfigurable structure (PS-) parallel processors are among the most successful in Soviet high performance computing. While not the most powerful machines, the PS-2x00 claimed more installations than any other Soviet high-performance machines and have had a signficiant impact in a host of applications domains from geophysics to atomic energy to space research.

These machines were developed in an industrial setting and historical tradition which strongly influenced their design, construction, and use. After examining some of the history of computer development at NIIUVM, we look in detail at the PS-2x00 computers, highlighting their technical and performance characteristics. With the third generation of the PS-2x00 machines currently under development, we have the opportunity to trace their evolution, identifying the environmental, technological, and organizational factors and design principles and strategies which shaped the machines' development. To a lesser degree, we examine their contribution to the Soviet economy and to developments in computing as a whole.

The Soviet Union witnessed many changes between the start of *perestroika* and its dissolution in 1991 but not all of them had a significant impact on HPC research and development at NIIUVM. We examine the most significant changes in the environment of the Problem-oriented Computer Division, including changes in legislation, the relationship with suppliers, factories, the ministry, and with the sponsors of the HPC work. We also track the changing demand (market) for these machines.

We look at the impact of change in two areas: the evolution of HPC technology and organizational change in the Problem-oriented Computing Systems division. We conclude by discussing the prospects for HPC at NIIUVM.

5.2 History of NIIUVM Research

During the early 1930s, the Soviet government founded the Lisichansk Chemical Combinat in a virtually uninhabited portion of the Donetsk region in Eastern Ukraine. The village that grew up around the Combinat was renamed from Liskhimstroy to Severodonetsk in 1950, and attained the status of city in 1958.

In 1956, the Lisichansk Subsidiary of the Institute of Automation (LFIA) was created in Severodonetsk to provide research and development facilities for the growing chemicals industry. The first director, V. Yu .Tolkachev, attracted many bright young scientists and engineers to Severodonetsk to work on the design and implementation of electronic equipment for industry [Raki91, 1-9]. Two of these, V. V. Rezanov and I. I. Itenberg, were to play critical roles in the development of both general-purpose control computers and the parallel processing computers discussed below. During the early 1960s, the institute was renamed the Scientific Research Institute of Control Computing Machines (NIIUVM).

Built at approximately the same time as the institute, the Severodonetsk Instrument-Building Factory (SPZ) in 1959 began series production of its first devices, temperature and pressure regulators for the chemical industry. Since then, it has focused primarily on electronic instrumentation for the chemical and other industries, and only secondarily on computers [Pert86, 22]. The first series production of digital computers here began during the mid-1960s

In 1971, the Impul’s Scientific-Production Association (NPO Impul’s) was formed on the basis of a number of Severodonetsk organizations as part of a national reform effort to improve the linkages between science and industry through the creation of the scientific production associations [Pert86, 21]. Impul’s consisted of three structural units and three independent enterprises. The head organization, NIIUVM, was responsible for all Impul’s research and development activities in the area of computing systems for data processing and process control systems. The Experimental Factory of Computer Technology was closely associated with NIIUVM and built prototype models of computers designed at NIIUVM. The third structural unit was the Impul’s Foreign Trade Firm. The three independent Impul’s enterprises were the Severodonetsk Instrument-building Plant (SPZ) which was responsible for series production of computers developed at NIIUVM, the Severodonetsk Training-Computing Center which trained thousands of users of Impul’s computers through regular seminars, and the Volgograd Design-Engineering Bureau of Control Computing Systems which developed computer technology for building control computer systems [Impu89c]. By the end of the 1980s, approximately 12-15,000 people worked at NPO Impul’s.

NIIUVM became a pioneer in Soviet computing. Numerous innovations in the areas of control systems and data processing appeared for the first time in the Soviet Union in NIIUVM machines. As described below, these include the first commercial multi-machine and single-instruction multiple data (SIMD) parallel processors. Thousands of NIIUVM control systems were installed throughout Soviet industry, in atomic and other power stations, chemical plants, metallurgy enterprises, engine manufacturing plants, seismic data processing centers, hospitals, etc. NIIUVM computers were used in the first

Soviet airline reservation system SIRENA [Zhoz85; Pert86; Raki91] and in an information system supporting the 1980 Summer Olympics [Vdov81; Pert86; Impu91, 5; Raki91, 15]. NIIUVM machines reportedly have provided climate control in Lenin’s Mausoleum [Impu91, 5].

Initially, NIIUVM research focused on developing small-scale control systems called ‘‘auto-dispatchers,’’ punched-card readers, and other small process control devices, but during the 1960s researchers developed a number of indigenous computers. Among their accomplishments were the MPPI-1 (1963) for gathering information in technical processes, the UM-1 (1965) for real-time control, and the KVM-1 (1965) coordinating computer to manage installations consisting of a number of subordinate control computers. These systems together formed a hierarchical System of Operational Control of Production (SOU) which was one of the earliest Soviet attempts to create an upwardly compatible family of computers [Rudi70, 29; Apok74, 275; Raki91, 7]. Also developed at NIIUVM were the Severodonetsk, the first stack-based architecture with zero addressing built in the USSR and the M-2000 and M-3000, the first Soviet systems based on the architecture and instruction set of the IBM System/360 computers, even though they were built using transistors instead of integrated circuits [Prsu70; Yers80b, 10; Raki91, 6].

Most of the computers developed at NIIUVM from the mid-1960s on fell into the ASVT (Aggregate System of Computer Technology) classification. ASVT computers fall into four primary categories: ASVT-D, ASVT-M, ASVT-SM, and ASVT-PS. The principal machines are shown in Table 5-1. Development of ASVT machines, particularly the ASVT-M, was spread out over a number of institutes, but the machines reflected a consistent underlying design philosophy: the use of modular designs, standardized internal and external interfaces, and upward compatibility of software would support the creation

	ASVT-D		ASVT-M		
	M-2000	**M-3000**	**M-6000**	**M-7000**	
Word length (bits)	32	32	16	16	
Peak Performance (add) (KOPS)	40	100	200	400	
Main memory, maximum (bytes)	48K	96K	32/64K	128/256K	
Start of series production	-	1971	1972	1975	
Factory	-	SPZ	VUM/SPZ/ELVA	SPZ	

	ASVT-SM				ASVT-PS
	SM-1	**SM-1M**	**SM-2**	**SM-2M**	**PS-1001**
Word length (bits)	16	16	16	16	16
Peak Performance (add) (KOPS)	200	200	400	450	1000
Main memory, maximum (bytes)	32/64K	128K	128/256K	128/256K	512K-4M
Start of series production	1977	1982	1978	1983	1989
Factory	SPZ	OZVM/SPZ	SPZ	OZVM/SPZ	SPZ

Key:	
VUM	- Computer and Electronic Control Machines Plant (Kiev)
SPZ	- Severodonetsk Instrument-building Plant (Severodonetsk)
OZVM	- Orel’ Computing Machine Plant (Orel’)
ELVA	- Elva Scientific-Production Association (Tblisi)

Table 5-1 NIIUVM Control Computers
Sources: [Rudi70; Shel73; Apok74; Kovt75; Naza75; Pevn76; Aman77,405; Eg78; Uprs77; Grub77; Polu78; Naum79; Zhim79; Zamo85; Grub89; Impu89b,5]

of a wide array of problem-oriented, upgradable configurations with clear migration paths for users without a corresponding proliferation of hardware and software nomenclature.

The M-1000, M-1010, M-2000, and M-3000 were called ASVT-D computers since they were built using discrete components. These machines were targeted chiefly towards data processing applications. The M-3000 entered series production around 1971

[Prsu70; Reza71; Aman77, 405]. It had a reported average performance of 50 KOPS [Rudi70, 31; Doly82, 13-16].

The ASVT-M class consisted of the M-40, M-400, M-4000, M-4030, M-5000, M-5010, M-5100, M-6000, and M-7000 models, and all were manufactured using integrated circuits [Shel73; Grub89, 89]. The development of these machines was divided between NIIUVM and other institutes, including the Institute of Electronic Control Computers (INEUM) in Moscow which was chosen to be the lead institute in a large-scale CMEA program initiated in 1974 to develop the SM- series of ‘‘small systems’’ (*sistema malykh*) minicomputers. NIIUVM was responsible for the M-6000 and M-7000. Together, the ASVT machines represent a considerable range of underlying architectures. While the ASVT-D systems, as well as the M-4000 and M-4030, were based on IBM’s mainframe architectures, the M-6000 and M-7000 were based on the architecture of Hewlett-Packard’s HP-2116. The M-400, on the other hand, designed at INEUM, was based on the Digital Equipment Corporations PDP family of minicomputers [Boya77, 7; Grub80, 19-20; Hamm84; Sini87]. They were not copies of their Western counterparts, but incorporated modifications which facilitated interoperability in keeping with the ASVT philosophy. Thousands of M-6000 and M-7000 were manufactured and employed in process control applications in all branches of industry.

The SM-1 and SM-2 both were upward compatibile with the M-6000 and M-7000. The SM-1 was designed for use in simple, single-level control systems as a control unit for instrumentation and a real-time processor for data generated by these instruments. The SM-2 was designed for complex, multi-level, multi-machine process control systems with stringent performance and reliability requirements. It could be configured as a dual-processor systems with shared memory to enhance reliability.

The SM-1M and SM-2M differed from the SM-1 and SM-2 chiefly in that they were manufactured using smaller boards with an improved component base consisting of LSI chips and semiconductor memory, making it possible to reduce their size significantly, decrease the number of internal connections between the processor and memory, and improve reliability overall [Grub89, 126-127].

The SM-2M reflected a number of further developments in NIIUVM systems design. First, the SM-2M was designed to have much higher levels of reliability than its predecessors. In a dual-processor configuration, the SM-2M had full redundancy such that the failure of any one component would not cause the failure of the system as a whole [Bara82; Grub89, 127]. Second, the systems software reflected the application in systems software of the modular principles used in the hardware of earlier machines. In machines like the M-6000 the systems software was quite inflexible and closed; there was little consistency in the interfaces of the composite modules and no easy way of adding modules to existing software. The essence of the Aggregate System of Software (ASPO) was a uniform, or compatible set of interfaces between systems modules, a macro-language for describing how modules could be combined, and a program generator which would translate the high-level configuration description into a full program which could then be compiled in the conventional manner. Using libraries of macro specifications and relocatable code, even the operating system could be tailored to the requirements of a specific installation [Klio89, 127-129].

The PS-1001, the primary system in the ASVT-PS category, was a considerable advancement in all respects over the SM-2M. Introduced into series production at least eight years after the SM-2M, the PS-1001 consists of much greater amounts of memory, higher processing rates, and additional functionality for the most demanding process control applications. These included networking capabilities, and the ability to configure the machine into a system with triple-redundancy and backups of I/O and communications channels. It also can be equipped with external solid-state memory [Prsu87, 50]. This system was developed primarily for use in nuclear power plants [Astr88].

Each of the machines of the ASVT series developed at NIIUVM was developed for industry. Close relationships with users and real-world requirements forced developers to design entire systems–not just isolated computational engines–which could be manufactured in a reasonable manner and at a reasonable cost and be of high utility to users. The industrial nature of the R&D environment served as a filter for design ideas. Practicality was favored over theory. Building machines which could be used was more important than demonstrating innovative architectural ideas. This tradition had a profound effect on the development of a new line of high-performance data processing machines initiated in the mid-1970s.

5.3 The PS-2x00 Parallel Processors

Two early Western high-performance computer prototypes, the ILLIAC IV and the STARAN, inspired Soviet designers to work on a related architecture. Although the Soviet work cannot be considered an effort to duplicate the Western machines, many design features of the latter were closely studied and, to varying degrees, adopted by the Soviets. In some cases, the Soviet work was an improvement over Western approaches. We highlight the most significant similarities and differences in this section.

5.3.1 Western Antecedents

The ILLIAC IV project was initiated in 1966 when the Advanced Research Projects Agency (ARPA) of the U.S. Department of Defense awarded the University of Illinois a contract to build a computer based on a design called the SOLOMON (Simultaneous Operation Linked Ordinal MOdular Network) proposed in 1962 by D. L. Slotnick and others. The 1962 design described a two-dimensional array of 1024 processing elements performing bit-serial arithmetic on a single instruction stream originating from a single control unit. The ILLIAC IV was not an exact implementation of the SOLOMON but was a pioneering effort into the realm of highly parallel systems.

Designed for solving partial-differential equations, the original ILLIAC IV plans called for a system consisting of four quadrants of 64 floating-point processing elements. Each quadrant was to have a single control unit interpreting a single stream of instructions to be executed simultaneously by all the processing elements in the quadrant. The quadrants were to be connected with one another by a high-speed parallel I/O bus. Each processing element was to have a local store of 2K 64-bit words, and within each quadrant the processing elements were to be arranged in an 8x8 array using a NEWS (North East West South) nearest-neighbor interconnect topology, connecting each processing element with four others. The main memory of the system was implemented in disk with a capacity of 16 megawords and a transfer rate of 500 Megabits per second. The individual processing elements operated in two modes, enabled or disabled. The mode could be set independently under program control for each of the 64 processing elements. The ILLIAC IV was attached to a conventional computer, the PDP-10, which managed I/O, the memory subsystem, and interaction with users [Hord90, 21].

The Burroughs Corporation was selected as the systems contractor and the first unit, consisting of only one quadrant, was delivered to the NASA Ames Research Center in 1972, although usable service was not offered until 1975 [Hock88, 24-26; Hord90, 28]. Although its development path was long and arduous, the ILLIAC IV had a profound impact on computer development. An internal memo at the Institute for Advanced Computation by G. Feierbach and D. Stevenson in August 1976 listed sixteen distinct advances attributed to the ILLIAC IV. These included the first large-scale use of ECL integrated circuits which served as a major impetus to the development of this technology, the first significant use of semiconductor memory, the first successful implementation of large multilayer laminated boards, a definitive demonstration of the array approach to computation, the development of new algorithms, and others [Hord90, 79-84]. The ILLIAC IV was to be a strong source of inspiration for Soviet computer developers as well.

The STARAN associative processor, built by Goodyear Aerospace, was conceived in 1962 and completed in 1972. It incorporated some associative, or content-addressable ideas proposed in 1960 by W. Shooman. The STARAN consisted of four array modules, each incorporating 256 single-bit processing elements. In contrast to the SOLOMON in which each processing element had a local memory store, the STARAN processing elements shared a single store of between 64K and 64M bits through a flexible ‘FLIP’ network [Hock88, 29-30]. Memory was accessed in terms of 256-bit ‘‘slices’’ according to a 256-bit code, or pattern. Data elements were passed between processing elements through the FLIP network which provided a highly flexible interconnect system. The STARAN was particularly suitable for image processing applications.

5.3.2 The PS-2000

5.3.2.1 History

In the years following the 1972 introduction of the ILLIAC IV and STARAN, literature about these machines was widely circulated in the Soviet Union. A group of computer scientists at the Institute of Control Problems (IPU) in Moscow began considering how they might make use of some of the ideas presented in these two machines to achieve high performance computing capability in the Soviet Union. The two Western machines were attractive in large part because they presented the possibility of achieving relatively high computation rates using a rather slow component base, albeit in a rather narrow application domain. Under the direction of Izrail Medvedev, this group of enthusiasts developed a proposal for a machine consisting of a single sequential, scalar processor controlling a field of processors executing identical instructions simultaneously [Mnw84; Smyk85].

The design grew out of an analysis of Western developments as well as an examination of the computational models for a set of core algorithms found in many HPC applications. These included regular algorithms like FFT, recursive sum, sorting, difference equations, multi-dimensional array transformation, systolic algorithms, and others [Medv92]. Medvedev's team developed a methodology for transforming these computational models to computer architectures which could execute them. The PS-2000 was designed using this methodology to identify an architecture which would best support the computational models (or parallel-sequential transformations of them). The identified three categories of models. The first consists of identical parallel processing streams with no program branches. The second have multiple segments of type-1 processes with intersecting sections which are of the same type. Problems like matrix algebra, FFT, and others fall into this category. The third are arbitrarily structured sets of arbitrary operators. At any given time there are groups of similar operators to be executed. Their readiness to execute depends on the availability of input data and the truth of an associated predicate function. Data searching and sorting, and a number of iterative procedures have this quality. Medvedev determined that the third type of problem can be run of SIMD computers if each processing element is provided with a) an address processor for independent memory access and activation, and b) an activation processor for predicate computation. The PS-2000 is a machine designed for type-3 problems [Medv92, 50-52; Medv92b, 19-21].

In 1976 the group from IPU met with representatives from NIIUVM who had worked on the M-6000 and M-7000 under I. I. Itenberg. NIIUVM was a logical partner in this effort. IPU had dual subordination to the USSR Academy of Sciences and the Ministry of Instrument-Building, Means of Automation, and Control Systems (Minpribor), for which NIIUVM was a leading computer R&D facility. NIIUVM had extensive experience in industrial computer development and was closely associated with prototype development facilities. Furthermore, IPU and NIIUVM had already worked together on a number of projects. From 1976-1978 groups from the two institutes worked closely together to develop a design which could realistically be implemented. They were able to put together a rough design which called for 64 processing elements housed in five standard computer racks. Although the earliest design proposals called for an associative processing field inspired by the STARAN architecture, this idea was rejected in the draft design completed in 1978 because the components necessary to build it at that time were not available from the Ministry of the Electronics Industry (Minelektronprom), and NIIUVM did not have sufficient capability to design and manufacture the chips itself.

To build the machine, dubbed the PS-2000, financial and material backing was necessary. For several years the designers had been seeking to find potential customers who would be willing to support the project. The design, laying out in rough terms how one might go about developing a machine with a performance rate of 200 million simple addition operations per second was greeted with considerable skepticism, both at NIIUVM and elsewhere. Until that point, the most powerful machines developed at NIIUVM were only capable of 400 thousand operations per second. According to Aleksandr Nabatov, one of the principals from NIIUVM on the project, for the first few years they ‘‘had to work under conditions in which everyone doubted in what we were doing.’’ Thanks to the strong support and involvement of V. V. Rezanov and I. I. Itenberg who had gained considerable authority through earlier, successful, NIIUVM projects, work on designing the parallel processor continued. Foreign opinion of the machine's design also played an important role. Around 1974, Control Data Corporation had examined premiliary designs at IPU and expressed an interested in some joint work on such a project. The U.S. government vetoed any collaborative work in this area, but CDC's enthusiasm was a big help in convincing Soviet authorities that the project should be supported [Smyk85; Gure85]. Minpribor agreed to fund the project, with the Ministry of Geology and the Ministry of the Oil and Gas Industry serving as principal target users.

While work on the draft design of the PS-2000 was in progress, the Soviet leadership initiated a massive campaign to develop the nation's oil and gas resources [Gust83]. During the 1960s the Soviet Union had enjoyed a steady supply of cheap hydrocarbons, but from the mid-1970s, the Soviet leadership grew increasingly concerned about the state of energy production. The 10th five-year plan, trumpeted at the 25th Party Congress in February 1976, placed great emphasis on coal reserves which were greater than those of oil or gas. When coal production during 1976 and 1977 fell far short of the annual targets, the Soviet leadership perceived a looming energy crisis.

In late 1977, general secretary Leonid Brezhnev announced a sharp change in policy and initiated a crash program to speed up the development of West Siberian oil [Gust83, 29-31]. Beginning in 1978 oil investment was increased sharply, growing by roughly 100% in absolute terms by 1980 [Gust83, 31]. In 1980, the Soviet leadership once again shifted the direction of their energy policy, launching a campaign to increase the output of natural gas by over 50% in five years [Gust83, 1]. In his address to the 26th Party Congress in February, 1981, Brezhnev stated: ‘‘I consider it necessary to single out the rapid development of Siberian gas output as a task of first-class economic and political importance’’ [Gust83, 34].

The NIIUVM/IPU efforts profited from the intense, high-level attention being focused on hydrocarbon exploitation. The oil, gas, and geological ministries had a desperate need for computing power to process seismic data. The SIMD architecture is very suitable for a large number of seismogram processing functions including filtering, tracing, amplitude control, static correction, seismic formation analysis, etc. [Akhm82; Neft91]. The PS-2000 offered a peak performance rate of 200 MIPS, far beyond the capabilities of any indigenous computer then available. The project quickly gained considerable high-level support and was included in the plans realizing the energy campaigns [Akhm82].

While the crisis atmosphere of the oil and gas campaigns generated material and political support for the new parallel computers, it provided a set of constraints which shaped the development of the system as well. In 1978 the PS-2000 was only at the draft design stage. Powerful computers were needed quickly so the PS-2000 had to be built in as short a time as possible. In particular, with the 26th Party Congress looming ahead in February, 1981, it was very desirable that working models be in existence by that time [Trap81, 31]. Furthermore, the systems had to be usable. There was little freedom to experiment with radical approaches or to test new theories.

The pressure to get machines into production in under three years forced a number of design decisions. First, components already in series production had to be used. From the start, PS-2000 designers had sought to create a machine using existing components, but the present time pressures confirmed this approach; there simply was no time to develop new, more suitable chips. The available component base determined the physical size and performance of subsystems and therefore constrained the functionality that could be built into the system. Second, since there was limited time in which to develop hardware, peripheral systems, and a full complement of systems software, the system would have to be built to make the greatest use possible of existing hardware and software systems. These decisions, discussed at greater length below, included the use of existing, general-purpose host computers, a simplified, linear interconnect system, and a relatively short word-length.

From the outset, IPU and NIIUVM engineers began working closely with the Ministry of Geology, in particular with scientists under Vladimir Kreysberg in the Department of Automated Computer Systems in the All-Union Scientific Research Institute of Geophysics (NIIGeofizika) in Moscow. The geophysicists provided valuable input into the nature of seismic applications, played a significant role in evaluating design decisions and defining the software library, and were responsible for much of the applications software development for the system [Trap79; Impu91, 6]. The first applied program package developed for the PS-2000 was a seismic data processing system which allowed one to obtain time profiles of geological structures. This package was also developed for use in the acceptance testing of the state commission in 1980. During these early years a geophysical job description language was developed to enable geophysicists to interact with the system using geophysics terminology [Trap81, 29-31].

Several PS-2000 prototypes were built between 1978 and 1980 when the machine passed state testing. One year later it entered series production at the Severodonetsk Instrument-building Plant (SPZ). Completed in time for the 26th Party Congress, the machine was touted by the Minister of Minpribor, M. S. Shkabardnya [Shka81].

5.3.2.2 Architecture and Construction

The PS-2000 is a SIMD machine which consists of a four basic parts: a monitor system, a parallel processor (called the PPS-2000) with between 8 and 64 processing elements (PE) (in 8 element blocks), a control unit for the processing elements, and an external memory system. These are shown in Figure 5-1. The PPS-2000 parallel processor can

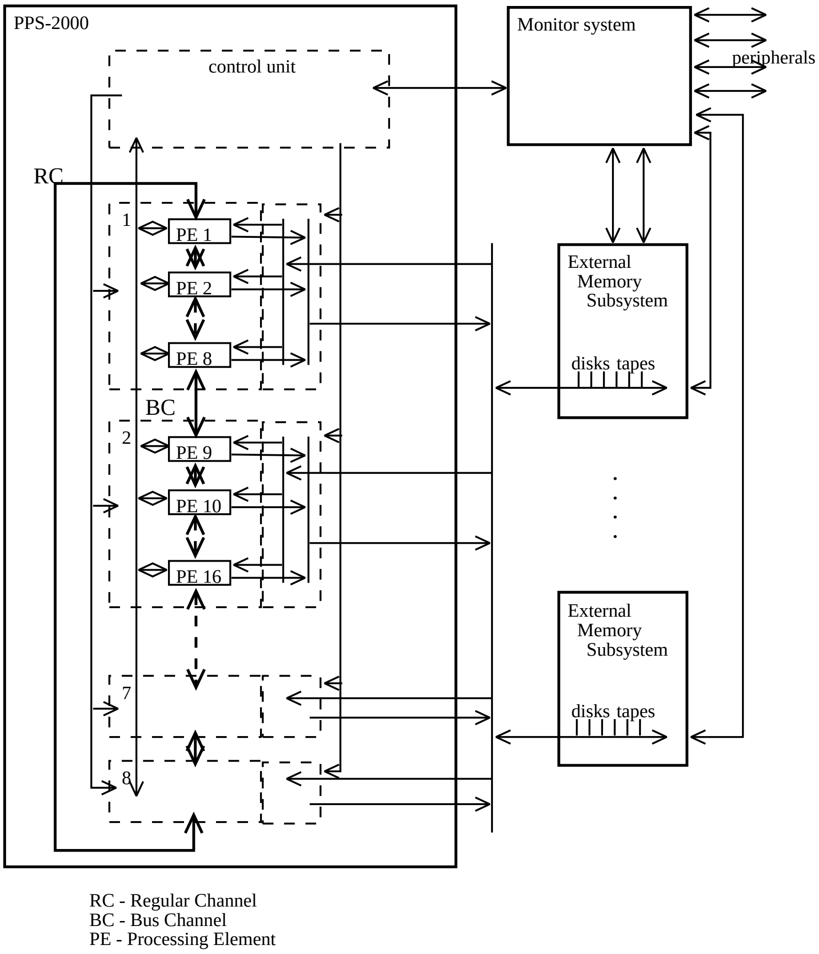

Figure 5-1 PS-2000 Multiprocessor
Source:[Pran83].

be considered a special processor which is attached to the monitor system, a standard control computer. A program runs on the monitor system, but uses the parallel processor to execute special functions such as fast Fourier transforms (FFT), linear algebra operations, etc. The PS-2000 was issued in six standard configurations. Three configurations had 8, 16, and 32 processing elements respectively. Two configurations had 64 processing elements, differing primarily in the amount of secondary storage provided. A dual configuration, containing 2x64 processing elements but only one monitor system was also installed in a number of locations [Grub89, 183]. Even larger configurations with up to 6x64 processing elements were reportedly installed [Medv92b, 30].

Monitor system. The PS-2000 was designed as a high-performance computing attachment to a standard-manufacture host, the SM-2, SM-2M, or SM-1210. The monitor system provides the primary user interface to the PS-2000. A standard production computer, the monitor system runs the customary operating system (with a few enhancements) and the systems utilities familiar to SM- users. In the PS-2000, the monitor system [Pran83, 114]:

- loads the control unit of the PPS-2000 with microprograms, programs, and constants;
- initiates and terminates execution on the PPS-2000;
- monitors execution on the PPS-2000 and performs diagnostics functions for the entire system;
- executes functions from the engineering console;
- supports the writing and compiling of PPS-2000 programs;
- manages time-sharing in a multiprocessing mode; and
- runs user programs which include routines to be run in the PPS-2000.

The user's application runs primarily on the host computer, the monitor system. Selected routines run on the PPS-2000. The monitor system loads the microprogram routines, initiates them at the appropriate time in the run of the application program, and waits for their completion. Once loaded into the PPS-2000, a routine –which would typically perform an operation like matrix generation, FFT, transformation of a tri-diagonal matrix, etc.–remains resident and is subsequently invoked simply by sending the PPS-2000 the correct parameters to initiate it. If there is insufficient memory in the control unit to accommodate a routine, systems routines are invoked to swap out one or more of the currently resident routines.

Rather than create specialized hardware and systems software for a host, the PS-2000 designers took as their base the SM-2 and SM-2M with the associated systems software, and merely had to modify these to support interaction with the parallel processor. The ‘‘Aggregate System of Software’’ (ASPO) which included the operating systems and a variety of systems utilities on the ASVT models was modified by including means of loading and initiating the execution of PS-2000 parallel code, a protocol for information and data exchange between the host and the parallel unit, a set of software development tools to facilitate the writing and debugging of parallel code, and a library of widely-used functions designed for the parallel unit [Pran83, 136]. In this manner, the user could operate in an environment with which s/he was already familiar, and the developers did not have to spend a large amount of time re-creating systems software which had been developed earlier. The PS-2000 operating system is a variant of the ASPO disk operating system of the SM-2x machines and could thus operate in real-time multitasking, timesharing, and batch modes [Pran83, 136].

Using a standard host also meant that existing applications written for the SM-2, -2M, and -1210 could be adapted to run on the PS-2000 by modifying only the portions of the application to be executed on the parallel processor.

A pragmatic solution, the idea of attaching the parallel unit to a conventional host computer was by no means a new idea. The ILLIAC IV, the STARAN and other early parallel machines such as the Parallel Element Processing Ensemble (PEPE) used conventional hosts to handle many systems functions.

The SM-2M and SM-1210 could be installed in dual configurations in which two machines shared a common memory space. Two monitor systems could therefore be incorporated into a PS-2000 system to improve reliability through redundancy; only one system could function as a monitor at a time, however [Pran83, 135]. The mean-time-to-failure rate for the PS-2000 could reportedly reach 3000 hours.

Parallel processor. A processing element contains:

- a 24-bit arithmetic-logic unit;
- a 24-bit main memory unit of between 4K and 16K words, depending on the capacity of the memory chips;
- a 14-bit index processor for indexed access to memory, incorporating ten index registers;
- an element activation processor for evaluating functions determining whether the processing element is to be activated or not;
- input and output registers which serve as buffers between main memory and the I/O channels.

The 24-bit word-length was chosen chiefly because of space considerations. Implementing a larger word-length would have required considerably more hardware per processing element, more cabinets, greater cost, etc. While technically feasible, larger word lengths were, at the time the PS-2000 was designed and built, not strictly necessary. For the target geophysics applications, 24-bit fixed- and floating-point formats were sufficient. In this respect the PS-2000 differs from the ILLIAC-IV which used 64-bit words to obtain the precision necessary for its intended use, solving partial differential equations [Hock88, 25]. Furthermore, the ILLIAC IV was intended as a computational vehicle for users with a variety of computationally intensive tasks such as linear programming, hydrodynamic simulations, ballistic missile defense analyses, and many more.

The PPS-2000 ALU includes an extension bit which was used to indicate double-precision, 48-bit values. This feature was not incorporated into the earliest units, apparently, for [Trap81, 20] claims that a 48-bit fixed-point processing format ‘‘is envisioned.’’ Being an add-on to the original design, operations on such values run very slowly, relative to the 24-bit operations. While each processing element executes one 24-bit register-register add per 320 nsec clock cycle and one 24-bit floating-point operation in three cycles, double-precision floating-point additions requires approximately 143 cycles [Impu89, 7].

The index processor of the memory module includes an ALU, a result register, a memory address register, and ten index registers which can be loaded in parallel. The use of multiple index registers gives the processing element a flexible, multi-level memory addressing scheme which allows many independent data exchange processes to occur simultaneously in different areas of memory without the need to reload index registers or use multiple operations [Pran83, 119].

The ILLIAC IV processing elements incorporated an enable/disable mode which could be set directly by the control unit or on the basis of a comparison of the contents of some registers, but the PS-2000 activation processor greatly expanded on this capability. The activation processor can enable/disable a processing element on the basis of logic operations performed on five internal registers or the results register of the ALU. It can also modify the links between adjacent registers, control data reception and transmission sequences on the bus channel, control the modification of the address data in the index-math processors, carry out associative attribute-based data search operations, and more [Pran83, 119-120]. The idea of building a parallel processor with associative memory as in the STARAN had been abandoned because of the lack of adequate components with which to build it. A remnant of this idea survived, however, in the form of associative capabilities in the activation processors.

The interconnection scheme for the processing elements is a novel features of the PS-2000 and it differs considerably from that of the ILLIAC IV. The ILLIAC IV used a nearest-neighbor interconnection scheme in which each processing element was directly connected to four others. Routing between processing elements took place via a ROUTE instruction which used a circular shift mechanism. This mechanism can be thought of as a circle with 64 slots in it. Each processing element can (optionally) place a data element in a slot. Then the entire circle is rotated some number of slots, and the data unloaded at the destination processing element. While this scheme can transfer up to 64 words at a time, it is restrictive in that all data items are moved the same fixed distance [Hord90, 31].

The PS-2000 uses two different interconnect channels to link processing elements, and a third, a data channel, to link the processing elements directly to external storage. A so-called ‘‘regular channel’’ provides a parallel link between each processing element and its two neighbors, forming a circular chain of processing elements. During *n* cycles, it can shift values which are located in identical registers, or in identical locations in memory in different processing elements, *n* processing elements to the right or left [Trap81, 29-31]. This mechanism is similar in concept to that of the ILLIAC IV, and suffers from the same restrictions. A so-called ‘‘bus channel’’ provides a serial link between each processing element and the control unit. The bus channel has a broadcast nature in which only one processing element (or the control unit) can send a data word at a time, but any number can receive it.

The PS-2000 approach differs principally from that of the ILLIAC IV in the reconfigurable nature of the links. The PS-2000 gets its ‘‘reconfigurable structure’’ (*perestraivayemaya struktura*) designation from the fact that these channels can be segmented and software reconfigured into 8, 4, or 2 segments consisting of 8, 16, or 32 processing elements, respectively [Berk82; Pran83, 119]. This serves two primary purposes. First, it becomes easier and faster to transfer data which is localized to small(er) sets of processing elements. Second, it makes it possible to broadcast different data elements to different processors over the bus channel.

When designers discussed how the processing elements should be interconnected, they naturally considered the four-way scheme used in the ILLIAC IV. Such a scheme was not feasible given the manufacturing technology available to them, however, since the number of interconnections was large, and would have required too much hardware to implement. They therefore considered linear interconnect structures and examined a scheme in which all 64 processing elements were linked in a single chain. They soon realized that considerably greater flexibility and performance could be achieved if instead of using one single chain they used a set of linear rings which could be combined to form larger rings.

Such an arrangement would have been worthless if it did not adequately support the information flows in the parallel algorithms for which the PS-2000 was designed. Designers asked the algorithm specialists to analyze the proposed configuration and determine to what degree their linear approach was worse than a four-way matrix interconnect. This interconnect did prove to be somewhat worse than the matrix scheme to no one’s surprise, but it did prove acceptable, especially considering the hardware constraints they were facing. Certain algorithms such as the FFT in fact map well to sets of processing elements with a ring interconnect.

Control unit. The control unit is linked through a serial channel to the monitor system. Its primary function is to control the execution of processes on the parallel processor. In addition, it receives commands from the monitor system, executes them, and manages the exchange of data between the PE memories and monitor system or external memory.

The control unit processes two types of instructions, each of which can be downloaded from the monitor system. One kind, 64-bit microinstructions, provides the control information that each PE needs to perform some action. The microinstructions, in some respects similar to horizontal architectures, simultaneously contain operators for the processing element ALU, the activation processors, the regular and bus channels, the memory unit, the index-math processor, and all microinstruction processor devices in the control unit. Thus the microinstructions can contain information about how data is to be moved among registers and memory within each PE, the arithmetic-logic operations to be performed, the activation status of a PE, transmission on the various channels, etc. The microinstructions allow complex operations to be represented in a single expression. For example, one such microinstruction can mean: ‘‘in device S add the contents of the R3 and B registers in each processing element, store the results in the C and M registers of the processing elements for which (T1 or T3 or T4) is true and jump to label M2 if the contents of the U register are less than or equal to the contents of register I2.’’ This control information is transmitted to each PE via a control signal bus. There are various formats of microinstructions which execute in 1-, 2-, or 3-cycles. Most microinstructions execute in one 320 nsec cycle, however [Pran83, 30-31].

The control unit also contains a data channel interface through which passes all data going to processing elements from the monitor system, or vice-versa. This interface makes the conversion between the 24-bit data used by the parallel processor, and the 16-bit used by the monitor system and external memory. This interface contains an input channel which has a throughput of 1.8 Mbytes/s and an output channel with a throughput of 1.4 Mbytes/s.

External memory. The ILLIAC IV used 16 Mbytes of disk memory, directly accessible from each processing element, as main memory. Access to additional external storage took place through the minicomputer host. The PS-2000 provides processing elements with direct access to external memory subsystems through an I/O bus. The notion of ‘‘main memory on disk’’ is not present. Each set of 8 processing elements has an input channel and an output channel onto an I/O bus which leads to one or more external memory subsystems. Input and output channels can be linked together to form sections of 16, 32, or 64 processing elements. Each external memory subsystem contains a maximum of two disk-drive controllers and four tape unit controllers for a total of four disks and eight tapes. Each complex can execute two transfers (one input and one output) simultaneously. PS-2000 configurations with 64 and 2x64 processing elements were equipped with 16 ES-5061 (29 Mbyte) removable-disk drives and 16 ES-5012 tapes [Grub89, 183].

Alternately, data exchange with peripheral storage can take place via the monitor system’s memory. The monitor executes I/O instructions storing data in main memory; this data is passed to or received from the parallel processor through standard channels. The monitor system can in this case be used as a sort of filter to selectively send data to the processing elements [Trap81].

Data exchange with external memory has proved to be a serious barrier to high performance. In 1983 the system used 29 Mbyte disks with data transmission rates of 312 KBytes per second. To achieve the highest I/O rate, the I/O channels should be segmented so that each group of eight or sixteen processing elements has independent input and output channels. In a maximum configuration up to four disks can feed data to the processing elements at an aggregate rate of 1.2 Mbytes/sec [Pran83, 134-135]. The input and output channels from each set of eight PEs have a combined throughput of 3.2 Mbytes/s (1.8 Mbytes/sec input, 1.4 Mbytes/sec output), or 25.6 Mbytes/s total for eight sets of eight processors. If the aggregate input transmission rate is 1.8x8 = 14.4 Mbytes/sec, then the bottleneck is clearly the data transmission rates of the magnetic disks. Later installations such as that at the Tsentrprogrammsistem software distribution center incorporated ES-5067 (100 Mbyte) disks with transfer rates of 806 KBytes/sec. Even these would be able to feed the processing elements at a rate of only 3.1 Mbytes/sec, a fraction of what the I/O busses can accommodate.

5.3.2.3 PS-2000 Success

The PS-2000 proved to be a highly successful computer, by Soviet standards. Only four years passed from the start of prototype development in 1976 to state testing. Series production was achieved only a year later. Between 1981 and 1989, approximately 200 units were manufactured at SPZ and used not only for geophysical applications, but also weapons design, mission control for the space program, metrology, medical diagnosis, atomic energy plant operation, etc. [Pipk83]. In 1988 the machine was nominated for (but did not receive) a State Prize in Science and Technology for its contributions [Izv880329]. The PS-2000 was a popular machine. It filled a high-performance computing niche which desperately needed computing resources. Other Soviet computers in series production during the early 1980s, mainframes, offered general-purpose processing capability, but had a peak-performance of less than 10 MIPS. It reportedly offered 5-10 times better performance on many applications than systems with attached array processors (see section 7.12) [Medv92b, 21]. The El'brus-1 was, for all but a small number of users, non-existent, and the El'brus-2 had not yet gone into series production. The PS-2000 was available, usable, reasonably reliable and could be purchased for several hundred thousands of rubles, a modest sum compared with high-end mainframes and the nascent El'brus computers which cost several millions of rubles. [Trap79; Pran83, 455; Mnw84; Smyk85].

The PS-2000 was successful for a number of reasons. First, the developers were experienced, capable computer engineers who had a good understanding of computer use in industry. Having designed computers since the late 1950s, NIIUVM researchers had much experience in developing computers geared towards use in industrial applications. They worked together with the IPU computer scientists to translate the ideas of the IPU ‘‘ideologists’’ into a viable design. The project was very much a joint effort. Every couple of months a team from Severodonetsk would travel to Moscow or vice versa to review alternatives, devise compromise solutions, make design decisions.

Since NIIUVM was part of the same association as the factory, SPZ, it had ready access to the same tools, materials, and components available to the factory. The prototype could therefore be built with an orientation towards production, shortening the time needed to put the prototype into series production.

Second, by using components already in existence, the length and complexity of the development cycle could be reduced considerably. Projects which push many technological boundaries simultaneously consume greater developmental resources, contend with greater uncertainties in supplies, are difficult to construct and debug because of the levels of new technology, and are very likely to experience significant delays.

Third, for its performance, the PS-2000 was not a terribly complicated machine to produce. The SIMD architecture required the manufacture of up to 64 identical processors per unit. With this degree of replication the cost per processor of setting up production was reduced. Since production quality for a given manufacturing line generally improves over time as the problems are ironed out, the high degree of replication of boards enabled stable quality levels to be achieved more quickly than would have been the case for a machine with little duplication of boards [Akhm82].

Fourth, the project had very focused goals. Developers had to create a machine which could be produced and used, and had to do so in a short period of time. These constraints served to focus the development effort and further keep the design within the bounds of what was feasible.

Fifth, the project enjoyed considerable high-level support. This ensured that the necessary links with Minelektronprom would be established expeditiously, that requisitions for development equipment and other necessary materials would be provided, and that the necessary monetary funds were allocated. The strong demand for the machines and political pressure also served to ensure that once prototyped the machine would be manufactured.

A sixth factor, difficult to quantify but shaping both Soviet policy towards computing and the market for the PS-2000, was the effect of Western export control policies. During the mid-1970s Control Data Corporation was selling a number of Cyber 17x computers to the Soviet oil and gas industry. During the last year of the Ford Administration, CoCom relaxed some restrictions on the export of computers to the Soviet Union and the U.S. administration approved the export of computers for use in Soviet air traffic control systems and production planning in the Kama River Truck Plant (KamAZ). Amendments passed in 1977 to the Export Administration Act of 1969 served to lessen export control restrictions by, among other things, calling for more expeditious handling of export applications, shifting emphasis away from shipments to ‘‘Communist’’ countries to emphasis on the commodity to be exported, and limiting the grounds under which the Secretary of Defense could recommend against export for national security reasons [Rich80, 165-167].

When Jimmy Carter became president, relations between the superpowers chilled considerably. Growing Soviet activism in developing countries and a military buildup which had given the USSR rough military parity with the U.S. had tarnished relations between the two countries. With his emphasis on human rights, and faced with the need to counter Soviet initiatives throughout the world, President Carter adopted a more confrontational approach.

Beginning in 1977, a series of executive actions and policy statements by both the Carter and Reagan administrations made it more difficult for the Soviets to acquire high-technology. In 1977 the Carter administration canceled the sale of a Cyber 76 to the Soviet Union’s weather research center on grounds that it was ‘‘a scientifically oriented computer that has wide uses in the United States for military research, development and support’’ [Nyt770624; Rich80]. In July, 1978, the Administration canceled the sale of a Sperry Univac computer to the Tass press agency for use during the 1980 Olympics. The reasons behind the cancellation were not only a claim that the Soviets might use the system for military purposes, but also a protest of the trials of two dissidents, Anatoly Shcharanskiy and Aleksandr Ginzburg, in Moscow [Burt78b; Wsj790328]. The export of a Cyber 173 for seismic research was turned down when the Department of Defense withdrew its support in 1979 [Rich80, 179]. When the Soviet Union invaded Afghanistan, U.S. export control policy shifted significantly. While the earlier measures were rulings on specific export licenses for specific machines, Carter’s response to the Soviet invasion, announced on January 4, 1980, included a cut-off of sales of high technology such as advanced computers and oil-drilling equipment until further notice [Cart80]. The tightened export controls on computing remained in place when Reagan assumed the presidency and were further strengthened after the military crackdown in Poland, for which Reagan held the Soviet Union responsible. ‘‘The issuance or renewal of licenses for the export to the USSR of electronic equipment, computers, and other high-technology materials is being suspended’’ [Reag81].

The PS-2000 was a source of great pride in the Soviet Union for it provided the Soviets with a system which could be used as an example of technological achievement at a time when the Soviet government felt considerable pressure to demonstrate its ability to develop all the technology it needed for its own uses. Numerous reports in the domestic and foreign press proclaimed the Soviet Union’s ability to proceed with world-class technological developments in spite of a moratorium on technology transfer from the West. R. Akhmetov’s comments in *Sotsialisticheskaya Industriya* about the use of the PS-2000 in the oil and gas campaign were typical of the tone which characterized these reports: ‘‘Incidentally, the Reagan administration, which has blocked the exchanges with the USSR of any information on computer technology, has done this to the detriment of science in its own country. More and more American scientists and specialists justly believe that USSR achievements...cannot continue to be ignored’’ [Akhm82]. The Moscow World Service in 1984 claimed that the PS-2000 was ‘‘quickly developed’’ after the American administration imposed trade restrictions on the sale of computers to the Soviet Union [Mosc84]. The PS-2000 project began before Carter’s more comprehensive restrictions in 1980, but a recognition of the writing on the wall following the Cyber 76 cancellation certainly played a role in advancing the project [Suga77].

The U.S. export controls also served to bolster demand for the PS-2000. When sales of computers and spare parts by Control Data Corporation–many to the Ministry of Geology–‘‘dried up’’ following the tightening of export control in 1980, the PS- computers lost a major competitor [Bosg811230]. ‘‘If the flow of these [CDC] machines had been without restrictions, then perhaps we would not have had any customers. To some degree [the export controls] helped us’’ stated A. S. Nabatov.

5.3.3 PS-2100

5.3.3.1 History

During 1981-1982, NIIUVM engineers worked primarily on installing and popularizing the PS-2000. During the first year of production the production process was still being perfected and customers needed considerable hand-holding to bring up their systems, keep them working, and learn to program and use them effectively.

In 1982, however, the NIIUVM engineers began to consider building a PS-2000 successor called the PS-2100. There were a considerable number of orders for the PS-2000 and the initial field results confirmed the system's usefulness. There was a demand, and there was reason to believe that there would be a demand for a successor as well. This work was carried out solely at NIIUVM without the participation of IPU researchers.

The early phases of development involved a careful assessment of both the requirements for a new machine and the current and future capabilities of NIIUVM, the production factory, and the supporting industries. The principal sponsor of the PS-2100 remained the Ministry of Geology, through NIIGeofizika, and its requirements continued to dominate the formulation of the technical statement of work. But during the design phase NIIUVM engineers made a concerted effort to broaden the market through an analysis of the nature of the computational problems faced not only by the geophysicists, but also by many of the other Impul's customers in the atomic energy, chemical, and other industries. For some months representatives from NIIUVM traveled around the country and held seminars in Severodonetsk to meet with actual and potential users of the PS-2x00 machines to analyze their applications and the associated hardware and software requirements. These requirements were packaged into the ‘‘*tekhnicheskoye zadaniye,*’’ or technical statement of work, which had to be approved by the primary sponsors, the geophysicists.

At the same time NIIUVM engineers made a careful evaluation of what technology–components, cabling, production tools, design tools, power supplies, etc.–was available. They determined which components and subsystems already in production could be used and which would have to be developed. In the latter case, they determined which of the missing components and subsystems could be developed in-house, and which would have to be developed by others.

In reviewing their work on the PS-2000 they concluded that the PS-2000 architecture used components effectively enough that a new machine built using the same MSI component base would yield only marginally better performance. The new machine should be built with a new generation of components.

In 1982, however, large-scale integrated circuits were not yet available. The component base was manufactured in another ministry, Minelektronprom, which, according to developers, was overloaded with orders from all quarters. Although the centralized planning mechanisms could specify that new types of production were to be initiated, if a factory had more total orders than it could fulfill, the director had considerable *de facto* authority to decide which part of the plan was to be met and which was to be disrupted. It was more in factories’ interests to manufacture chips with small- and medium-scale integration. Production for these was well established and they were being ordered in sufficient quantities to keep factories producing near capacity. By minimizing disruption in a factory’s production, higher volumes of production could be sustained, and the political fallout for factory directors for missed production targets could be minimized. To be sure, Minelektronprom was manufacturing components with large-scale integration for powerful military, or military-related customers. Insufficient numbers of such components were available for Minpribor’s non-military projects, however.

During 1983-1984, Itenberg’s team searched for an appropriate component base for the PS-2100. Finding nothing suitable from Minelektronprom, they formed a branch at NIIUVM to work on the development of integrated circuits. Lacking the technology to manufacture chips themselves, they experimented with hybrid technology, since this would allow them to reduce the amount of space occupied by components without having to use higher levels of integration in the circuits themselves. Hybrid technologies were strongly emphasized throughout Minpribor during the 11th Five Year Plan (1981-1985) [Gavr86, 42]. A hybrid integrated circuit is similar in concept to the printed circuit board and consists of multiple components interconnected on a single ceramic substrate. Unlike the components used on a printed circuit board which are each encapsulated in rather large packaging (relative to the size of the circuit itself), components in a hybrid circuit are devoid of packaging and can therefore be combined in a smaller physical space. Hybrid technology, now often refered to as multi-chip modules, has been used for many forms of analog circuitry, but until relatively recently has not been very popular for digital circuitry [Bras83, 773; Muku93]. One package developed at NIIUVM could contain up to 10 MSI chips. The engineers at NIIUVM developed in particular some mock-ups of hybrid integrated circuits oriented towards floating-point operations, because they had no such hardware at all at that time or for many years following.

In 1984 NIIUVM researchers learned that gate arrays were to become available from Minelektronprom in the near future. Gate arrays are a regular configuration of logic gates whose interconnects can be relatively easily customized to produce special-purpose chips. Their advantage is that without changing the underlying manufacturing technology, a wide array of chips can be manufactured simply by using a lithographic mask unique to each different kind of chip. Although their levels of integration are not necessarily as high as custom-made chips, they provide a very cost-effective solution to system development. Gate arrays have been popular as building blocks for many, if not most, Western computers.

In 1985 NIIUVM researchers began designing specialized integrated circuits and over a two year period developed seven types of gate array based chips which would be incorporated into the PS-2100. In the West at this time, chips based on gate arrays with on the order of 400 gates could be developed in a matter of months. There were a number of reasons why the NIIUVM efforts took more time. First, the technology was new to the engineers although processor design was not; there was a learning curve that needed to be climbed. Second, the base technology itself had not been perfected by Minelektronprom. At the time NIIUVM engineers started designing their chips, Minelektronprom’s Elektronika Scientific-Production Association in Voronezh’ was manufacturing experimental units. The electronics industry understood the technology, but had not yet mastered production. Third, development was slowed by a lack of sophisticated design tools. Such tools were developed both at NIIUVM and by a Lithuanian firm, but this took time as well.

The first prototype gate arrays for the PS-2100 were obtained at the end of 1987. They were correct, so at the beginning of 1988 they were manufactured in sufficient quantities to construct a complete base module. Several units were constructed during 1988, and the PS-2100 passed state testing in December, 1988 [Bere88]. By September, 1990, 15 PS-2100 base modules had been constructed and full series production began in late 1990 or early 1991.

5.3.3.2 Requirements

From their analysis of current and prospective applications, a number of architectural requirements emerged as dominant. Not surprisingly, higher performance was key. All the applications being considered, both the traditional seismic data processing as well as real-time image processing and process control, etc. needed higher processing rates, more memory, improved I/O, faster communications channels, and so forth. Designers realized they needed to concentrate on particular bottlenecks in the PS-2000. These included the inter-processing element transmission time, the efficiency of the object code, the time required to configure the system for a particular problem, and the access time to secondary storage. Of these, the time required to transmit data between secondary storage and the processing elements was the most severe bottleneck.

A second dominant requirement was reliability. Reliability had always been a primary concern in Impul’s computers since most of them had been oriented towards real-time processing control applications. Although the primary application of the PS-2000 was seismic data processing, it too had real-time capabilities and could be configured–such as in the dual configuration–to give it a mean-time-between-failure (MTBF) of reportedly 3000-5000 hours. For many customers whom NIIUVM was courting, reliability was even more important than high performance and MTBF rates approaching half a million hours became development goals. The emphasis on reliability was especially acute in the atomic energy industry [Impu91, 7]. According to Itenberg, the actual MTBF for the PS-2100 is several tens of thousands of hours.

The use of 32/64-bit data formats became a third major requirement. While 24-bit operands had been acceptable for geophysical applications during the late 1970s, they did not have the precision needed for modeling, solving systems of linear equations, etc. The shift to a 32/64-bit format was also attractive because these word-lengths had become standard in high-performance computing systems throughout the world.

These requirements demanded that changes be made to the system, but these changes were to be made based on the existing SIMD-oriented approaches. This was necessary for two primary reasons. First, a real demand for such machines existed, and would continue to exist for the next generation machines. In particular, the geophysicists who sponsored development had found the basic architecture to be a sound one for their purposes. Second, the existing installed base of PS-2000s made it necessary to retain a large measure of compatibility between the the PS-2000 and PS-2100. Radical changes to the basic architecture were, early on, ruled out. The SIMD architecture had proven its usefulness in many applications; the fundamental nature of these applications had not changed, so a SIMD approach would continue to be useful. As A. S. Nabatov put it,

> ...in our collective we understood that we had to continue with the SIMD architecture as long as the market for it exists. There were lots of discussions and ideas about how to upgrade this architecture. The fight was not regarding the architecture, but how to implement [it] [Naba91, 5].

As long as there were no fundamental problems with the basic SIMD approach and a demand for machines of this sort existed, work in this direction could be justified. Reinforcing this perspective was the fact that lacking a compelling reason to change, engineers are likely not to spend a large amount of time and effort considering radically different alternatives. They are likely to continue introducing incremental modifications to familiar approaches. When asked what they would change if they had the liberty to build the system from scratch with no compatibility requirements, Nabatov and Itenberg answered that

> ...it’s a very complex question. You understand, we have already become accustomed to this type of architecture. We think that it is rather promising for the future under our conditions, and in the West, too, a number of machines of this type of architecture–mostly experimental–are being developed. With regard to programmability, we would change part of the mechanism, but this is perhaps a question of small-scale changes in the architecture. Perhaps...I haven’t thought about this...we would move in a very different direction, but still [stay] within the area of parallel architectures [Iten90b, 42].

The installed base created pressures for preserving compatibility between generations of systems. Considerable volumes of software had been written for the PS-2000 and, naturally, users strongly resisted any efforts to alter the PS-2000 structure in a way that would require that they rewrite their software. To accommodate the peripheral devices which were to be attached to it, the PS-2100 naturally had to accommodate established peripheral interfaces.

5.3.3.3 Architecture and Construction

The PS-2100 is compatible with the PS-2000, but only at the level of ASPS, not at the level of executable code [Iten90d, 12]. ASPS is a hybrid language, having low-level features typical of assembly languages such as register access and bit-wise operations and high-level constructs such as DO-loops, complex data types, IF-THEN-ELSE clauses. It also contains instructions for managing the flow of data between processing elements [Ryad88]. The chief sources of incompatibility were changes in some of the register structures and data formats of the new machine. The decision to maintain compatibility only at the level of ASPS reflects the tension between the need to improve the performance of the machine, expand its capabilities, and preserve the existing software base. It also reflects certain characteristics of the PS-2000 and its software development which made a lack of compatibility at the lowest levels more tolerable.

Moving software from the PS-2000 to the PS-2100 requires recompiling existing software; but only software written in the lowest-level language, MNEMOKOD, would have to be rewritten into MIKROKOD PS2100, the microprogramming language for the PS-2100. While some users did write their own routines in MNEMOKOD, the bulk of MNEMOKOD programming, an arduous task, was done at NIIUVM. Programmers here were responsible for the library of approximately 200 routines and systems software which execute on the PPS-2000. The applications programs were written chiefly by users, but these typically were written at least at the ASPS level and relied heavily on the parallel routine library for code which actually ran on the PPS-2000. To be sure, some users did develop some code using microcode [Ivan85], but the number of programmers with sufficient knowledge of PS-2000 to write code at this level was probably not great. Finally, much of the systems software ran not on the PPS-2000, but on the host computer. In sum, the amount of code that needed to be rewritten was relatively limited and was chiefly the responsibility of NIIUVM programmers. The cost to applications developers and users, apart from recompilation, would be bearable. The PS-2100 engineers felt that the cost to their own programmers was tolerable, and worth the gains in performance that could be achieved.

Besides some low-level changes, the basic architecture of the PS-2000 remained intact. The PS-2100 architecture is best understood as a set of extensions to the PS-2000 architecture to improve performance and reliability. Table 5-2 compares the PS-2000 and the PS-2100.

The PS-2100 consists of the following basic parts [Impu89, 4-5]:

- 1–10 so-called ‘‘base modules’’ (BM), each containing 64 processing elements;
- switching system for intra-system data exchange;
- external semiconductor memory (SSM);
- external magnetic disk memory subsystem;
- external magnetic tape memory subsystem;
- monitor subsystem consisting of one or two PS-1001s;
- simplified configuration subsystem for dual systems.

The PS-2100 is shown in Figure 5-2.

	PS-2100	PS-2000
Data Formats (bits)		
fixed-point	16;32	12;16;24
floating-point	32;64	24;48
Number of base modules in configuration	1-10	1-2
Ability to join base modules in configuration into a single processor field.	yes	no
Number of processing elements in base module	64	8-64
Maximum main memory per processing element (bytes)	128/512K	48K
Theoretical peak performance of one processing element (base module):		
fixed-point addition (MIPS)	2.38 (150)	3.125 (200)
floating-point addition (single precision) (Mflops)	1.02 (65)	1.04 (67)
floating-point addition (double precision) (Mflops)	.12 (7.6)	.02 (1.39)
fixed-point multiply (single precision) (Mflops)	.40 (25)	.45 (29)
floating-point mult. (single precision) (Mflops)	.51 (33)	.35 (22)
floating-point mult. (double precision) (Mflops)	.07 (4.2)	.04 (2.3)
Throughput of one data channel (Mbytes/sec)	6	1
Total throughput of all data subchannels (Mbytes/sec)		
in one base module	48	25
in maximum configuration	480	50

Table 5-2 Comparison of the PS-2100 and PS-2000
Source: [Impu89]

Base modules. The architecture of each of the base modules is virtually the same as the PPS-2000 [Impu89, 6]. Each base module operates under SIMD control and contains the same types of reconfigurable channels interconnecting processing elements as the PPS-2000. During the early design phase of the PS-2100 a series of discussions were held about how these transmission channels might be changed, since during the development of the PS-2000 it was clear that the linear, reconfigurable rings were not optimal. As with

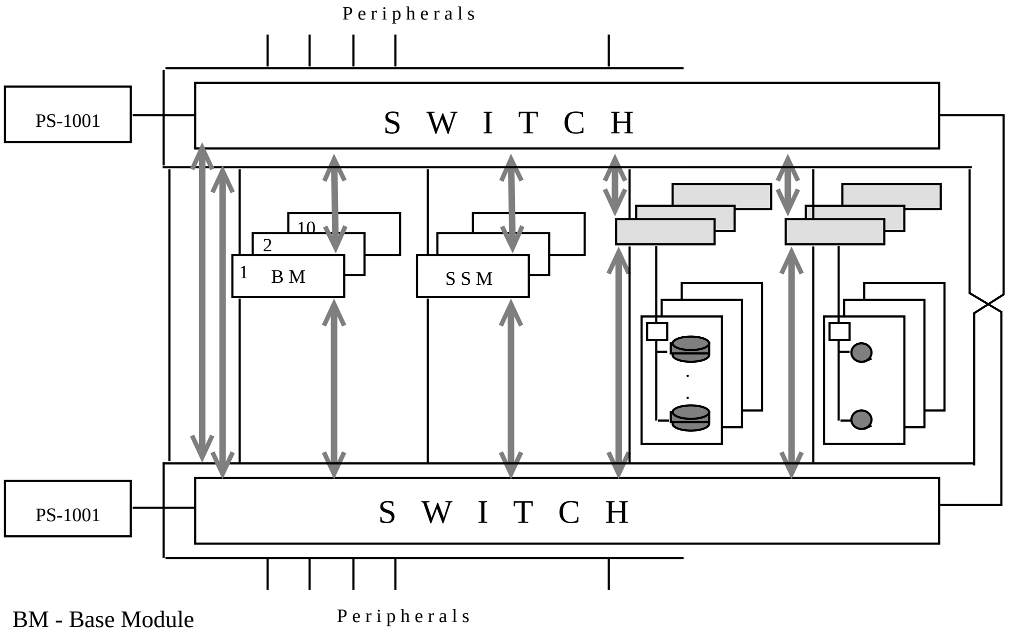

BM - Base Module
SSM - Semiconductor External Memory System

Figure 5-2 PS-2100 Multiprocessor
Source: [Impu89]

the PS-2000, the constraints of construction played a determining role. As Nabatov put it [Naba91, 6]:

> Each time we re-examined the possibilities [for changing the transmission channels] at the initial stages of each development....But the closer we got to the concrete implementation, the less willing we were to make a step forward in the area of lines....[I]t is a question of the construction and technology.

While improvements in technology made it possible to reduce the amount of hardware needed for a processing element, it also reduced the number of external connections to the processing element. The transmission channels have a fixed bus-bit width, so the fewer the external connections the more difficult it becomes to implement large numbers of transmission channels. In the PS-2000, eight boards were needed for one processing element. Together these offered on the order of 1200 external connections. The PS-2100 processing elements reside on one board, so the number of connections is reduced by a factor of eight. For this reason, improvements in the component base have not made it easier to increase the number of transmission channels between processing elements, in spite of a desire, in principle, to do so. Similar constraints are encountered in the PS-2300 (discussed below) in which each processing element will reside in one chip with 132 (128 data) outputs. Such constraints limit the degrees of freedom designers have to develop creative interconnection schemes.

In contrast to the PS-2000 in which only one instruction is sent from the control unit to the processing elements, a setup mechanism is used in the PS-2100 base module to increase the rate at which instructions can be sent to each processing element. In addition, the base modules are equipped with enhanced features for real-time operation: an interrupt system and a real-time timer.

Up to ten base modules can be configured in a single system. Each base module can execute an independent instruction stream, making the PS-2100 as a whole a MIMD/SIMD architecture. The ten base modules can execute either separate programs, different parts of the same program, or the same part of a single program on different data sets.

The ability to configure 1–10 based modules together serves a number of purposes. First, it makes it possible to increase the theoretical peak performance of a system by a factor of ten. Each base module has a theoretical peak performance of 150 MIPS; a full configuration, 1.5 GIPS [Impu89, 7; Iten89, 3]. The base module approach makes it possible to tailor a configuration to the performance requirements of individual users, thus making it possible to serve a broader market. In this respect the PS-2100 continues the principle incorporated into other Impul’s machines.

The base module approach also makes it possible to increase significantly the reliability of a configuration through redundancy. Individual subsystems, including the base modules, can be switched off and the configuration automatically reconfigured via a sliding backup mechanism [Impu89, 15]. In addition, when a processing element fails within a base module, the half of the base module containing the failed processing element is switched off and the job completed on the remaining half or on another base module [Impu89, 15]. To reduce the amount of down-time, each PS-2100 base module comes equipped with extra boards of each type. The user does not therefore have to wait for a replacement board to arrive from Severodonetsk but can replace a defective board immediately.

Internal exchange switch. To provide the flexibility and high data throughput needed in configurations with multiple base modules, engineers developed a so-called ‘‘internal exchange switch.’’ This switch, organized in a matrix configuration and based on the K1509KP1 switch developed at the Scientific Research Institute of Multiprocessor Computing Systems in Taganrog [Iten89, 2], permits data exchange between arbitrary systems components which are attached to it. It can support the simultaneous transmission of up to 31 data streams. The internal exchange switch comes in two configurations. A single-section switch contains 11 ports; a dual-section switch, 31. A synchronous interface at each port permits full duplex transmission. At any given moment, devices at any two ports can exchange data at a maximum rate of six Mbytes/sec [Impu89, 8-10]. The switch is tuned via instructions issued by the monitor subsystem.

I/O and external storage. Relieving the severe I/O bottleneck required multiple solutions at different levels of the architecture. A fundamental constraint was the capabilities of the secondary storage available to PS-2100 manufacturers. The most powerful Eastern Bloc disk drives in series production were the Bulgarian ES-5063 317.5 Mbyte drives. Produced since the mid-1980s, these drives have a data transmission rate of 1.25 Mbytes/sec [Impu89].[1]

Magnetic disk subsystems consisting of up to six ES-5063 disks attached to a single ES-5563 disk controller and each controller can be attached to one port of the internal exchange switch. This transfer rate per disk is approximately four times faster than that of the ES-5061 disks used in the PS-2000, although throughput from a disk subsystem is limited by the capabilities of the controller. Each controller can probably sustain read operations from not more than one or two disks simultaneously, putting an upper limit of 2.5 Mbytes/sec transfer rate per disk subsystem. The number of such subsystems is variable, theoretically limited only by the number of available ports in the internal exchange switch. In practice, the maximum configuration was equipped with four subsystems, each containing six disks for a total of 7.6 Gbytes [Iten89, 2-3].

The PS-2100 has fewer, but faster, I/O channels per set of 64 processing elements than the PS-2000. Each base module has four pairs of I/O channels, each accessible by 16 processing elements. Each channel has a throughput of six Mbytes/sec [Impu89, 6]. The aggregate throughput of the I/O channels for one base module is therefore 48 Mbytes/sec (relative to 25.6 Mbytes per second for a 64-processor PS-2000 configuration) so the gap between the capability of the disks and the I/O channels was reduced, but remained significant [Impu89, 8]. In contrast to the PS-2000 I/O channels, those in the PS-2100 have multiplexed subchannels, allowing greater flexibility in data exchange with external memory [Iten89, 1]. The throughput of the magnetic disks could not be increased, since larger capacity disks such as the 635 Mbyte ES-5065 were not being produced in large enough quantities to be accessible to Impul’s [Vtp89, 19; Gors89, 4].

---

[1] Other sources such as [Dani84b; Dani86] put the transfer rate at 1.198 Mbytes/sec, comparable to the 317.5 Mbyte IBM-3350 disk drives which were introduced in 1976 [Rdd82, 5; Dpro82, 40].

To compensate for the slow disk storage, designers greatly increased the amount of semiconductor memory. Most significantly, they added a semiconductor external memory unit consisting of blocks of 16 Mbytes each. Over a dozen such blocks could be installed, if desired, but as a rule 1–6 were [Impu89, 12]. External solid state memory devices have been used widely in Western supercomputers to provide data transfer rates high enough to keep the central processing units loaded. The PS-2100 is the first Soviet high-performance computer to use such a device, however. The semiconductor memory has a peak transfer rate for read operations of five Mbytes/sec through each of four ports which interface to the internal exchange switch for an aggregate rate of 20 Mbytes/sec [Impu89, 12]. This nearly fully loads the four input channels internal to the base module.

A second means of relieving the I/O bottleneck involved increasing the amount of local storage in each processing element. While each PS-2000 processing element contained up to 48 KBytes of local memory, each PS-2100 processing element contains 128-512 KBytes, depending on whether 64 Kbit or 256 Kbit memory chips are used.

The ability to expand the PS-2000 architecture by multiplying subsystems was largely facilitated by improvements in construction which permitted greater functionality to be packaged in a smaller space. In the PS-2000, each processing element occupies eight boards. A 64-processing element configuration occupied five cabinets [Trap79; Impu89, 8]. In the PS-2100, 6-8 layer boards were used in place of the PS-2000 2-layer boards, and gate arrays with 400 gates per chip made it possible to place an entire processing element on a single board. As a result, an entire 64-processing element base unit could be housed within one cabinet. This in turn made it feasible to incorporate multiple base modules in a single configuration. The need to manage a growing number of possible data transmission paths between the base modules and the peripheral units necessitated the creation of a new form of intra-system data exchange, the internal exchange switch.

5.3.3.4 Performance

The gains in processing speed came primarily from these extensions to the architecture rather than from increased processing rates in each processing element. The size of each processor was reduced making it possible to reduce the clock time from 320 nsec to 140 nsec. The number of cycles to perform basic operations was increased, however, largely to accommodate the increased word-length. While a single-precision fixed-point addition in the PS-2000 required one 320 nsec cycle, a similar operation in the PS-2100 required three 140 nsec cycles. As a result, the aggregate peak performance of the PS-2100 base module was 150 MIPS (32-bit operands) vs. the 200 MIPS (24-bit operands) of the PS-2000 [Impu89, 7].

Using gate arrays, NIIUVM engineers were able to create their own floating-point unit. While in the PS-2000 floating-point operations were something of an afterthought, in the PS-2100 they were incorporated into the design from the beginning; a double-precision floating-point addition takes only 20 times longer than a fixed-point addition, rather than 143 times longer as in the PS-2000 [Impu89, 7].

Between September, 1990 and January, 1991, LINPACK, Livermore FORTRAN Kernels, and Los Alamos National Laboratory Vector Operations (VECOP) benchmarks were executed on a single PS-2100 base module with 64 processing elements. The effort was not as comprehensive or rigorous as comparable Western efforts and in some cases basic conditions of the benchmarks were unavoidably violated.[2] While the results cannot be considered conclusive and their generalizability is questionable[3], they do provide a

---

[2] For example, LINPACK policy is that the code be compiled and run without modification–even to the comments– except to incorporate a timing routine local to the system [Dong90]. However, by very nature of the PS-2100, additional routines to initiate and manage the execution of parallel routines on the parallel unit must be added to the source code.

[3] Jack Dongarra, the author of the LINPACK tests states, "Benchmarking, whether with the LINPACK Benchmark or some other program, must not be used indiscriminantly to judge the overall performance of a computer system" [Dong88, 14].

rough orientation of system performance on various types of computational algorithms. A complete account of these tests can be found in [Wolc91].

LINPACK. The LINPACK Benchmark is a collection of FORTRAN subroutines for solving certain systems of linear equations [Dong88]. The benchmark was originally developed by Jack J. Dongarra and others to give users of the LINPACK software package data with which to estimate execution times on their own machine. Since then it has become one of the most widely used (and abused) benchmarks. The basic LINPACK Benchmarks operate on a 100x100 matrix. This was originally felt to represent a ‘‘large enough’’ problem [Dong88]. In recent years the LINPACK has been modified to operate on a 1000x1000 matrix. Called ‘‘Toward Peak Performance (TPP),’’ these benchmarks present a problem which is very large, giving computers the opportunity to use the hardware as efficiently as possible by reducing the relative amount of time spent on overhead. LINPACK policy allows users to make changes to the benchmark code so that this test would represent a user’s ‘‘best effort.’’

Table 5-3 shows the results obtained on the PS-2100 for single and double precision LINPACK tests (unoptimized), and TPP, double precision (optimized):

Test	Performance (Mflops)
Single precision (32-bit), unoptimized	1.73
Single precision (32-bit), optimized	5.3
Double precision (64-bit), unoptimized	.568
TPP double precision (64-bit), optimized	1.65

Table 5-3 PS-2100 LINPACK Results

Table 5-4 gives a comparison of these results with those of Western computers.

Computer	LINPACK (Mflops)	TPP (Mflops)	Theoretical Peak Performance (Mflops)
Cray-1S	27	110	160
Cyber 205 (2 pipe)	17	113	200
IBM RS/6000-320H	12	37	50
PS-2100	.57 (unopt)	1.7	7.6
DEC VAX 6000/410	1.2	1.5	2.6
CDC 6600	.48		
VAX 11/780 FPA	.14		

Table 5-4 PS-2100, Western Machines LINPACK Performance
Source: [Dong92; Wolc91]

As will be discussed below, the relationship of the order of the problem to the number of processing elements plays an important role in determining overall performance. The LINPACK Benchmarks use problems of order 100, 300, 1000, i.e. not even multiples of 64. When the LINPACK algorithms were run on a 1024x1024 matrix (optimized), the performance was 10.877 Mflops.

Livermore FORTRAN Kernels. The Livermore FORTRAN Kernels (LFK) are a set of 24 small algorithms drawn from applications used at Lawrence Livermore National Laboratory. The degree to which the Kernels accurately characterize the LLNL workload has been a point of discussion [Lube88; Mcma86].

The LFK were executed on 32-bit data. Most published figures for Western machines report performance on 64-bit data. Table 5-5 compares PS-2100 performance with that of a Cray-1 running the CFT ’84 compiler. An analysis of PS-2100 performance on these algorithms is given in [Wolc91, 21-27]. Although there are some surprises, such as the low performance on Kernel 21 which involves a matrix multiply–something that should parallelize very nicely–, the variation in performance on the various kernels can be ex-

Kernel	PS-2100 (Mflops)	Cray-1 (Mflops)
1	2.315	100.0
2	0.081	41.7
3	0.643	33.3
4	0.066	24.3
5	0.113	7.7
6	3.638	7.0
7	4.787	120.0
8	10.500	55.4
9	6.671	68.0
10	2.364	36.0
11	0.211	2.9
12	0.829	25.0
13	0.778	4.0
14	0.103	5.6
15	1.305	
16	0.189	
17	0.033	
18	5.786	
19	0.163	
20	0.094	
21	0.040	
22	1.101	
23	0.281	
24	0.037	
Arithmetic mean	1.757	37.92
Harmonic mean		
Loops 1-14	0.263	11.1
Loops 1-24	0.148	
Maximum	10.5	120.0
Minimum	0.033	2.9
Ratio (Max/Min)	318.375	41.1

Table 5-5 PS-2100, Cray-1 Performance on Livermore FORTRAN Kernels
Source:[Worl84; Wolc91]

plained roughly by two basic factors: the degree to which iterations of the loop can be executed on multiple processors, and the relationship between the number of PS-2100 processors and the number of loop iterations in the kernel.

VECOP Tests. Researchers at Los Alamos National Laboratory developed a simple model to help understand the performance of supercomputers [Buch84]. The vector operation (VECOP) tests are a simple set of routines to measure the performance on 64-bit floating-point operations as a function of vector length.

The benchmark measures the performance of 1,000,000 executions of such simple vector operations as adding a scalar to a vector (V+S), adding two vectors (V+V), adding a scalar to the product of two vectors (V*V+S), etc. The results are dependent on both the vector length and the manner in which operands are stored in memory. In particular, the benchmarks consider three storage cases: operands and results are stored in contiguous memory locations, in non-contiguous locations but with constant stride, and in a random fashion (scatter/gather). Some VECOP results are shown in Figures 5-3 and 5-4. A more complete set of data can be found in [Wolc91]. Not surprisingly, the VECOP results rather clearly reflect the nature of the PS-2100. Some of the conclusions to be drawn from the tests are [Wolc91, 28-30]:

- Performance is very sensitive to the way data elements are stored in memory and the length of the vector. If the number of processors is a multiple of the stride (such as when stride = {2,4,8}, vector elements will be doubled up on some processing elements but not on others. When stride=2, as shown in figure 5-3, data are stored only at the even (or odd) processing elements. A maximum of 32 processing elements execute, leading to the observed decrease in performance, when the vector length is greater than 32. If the vector is short enough that all elements can be distributed to all active processing elements, then there is no degradation of performance.
- Performance is very sensitive to the relationship between the number of processing elements and the number of elements in the vector. In the graph in Fig-

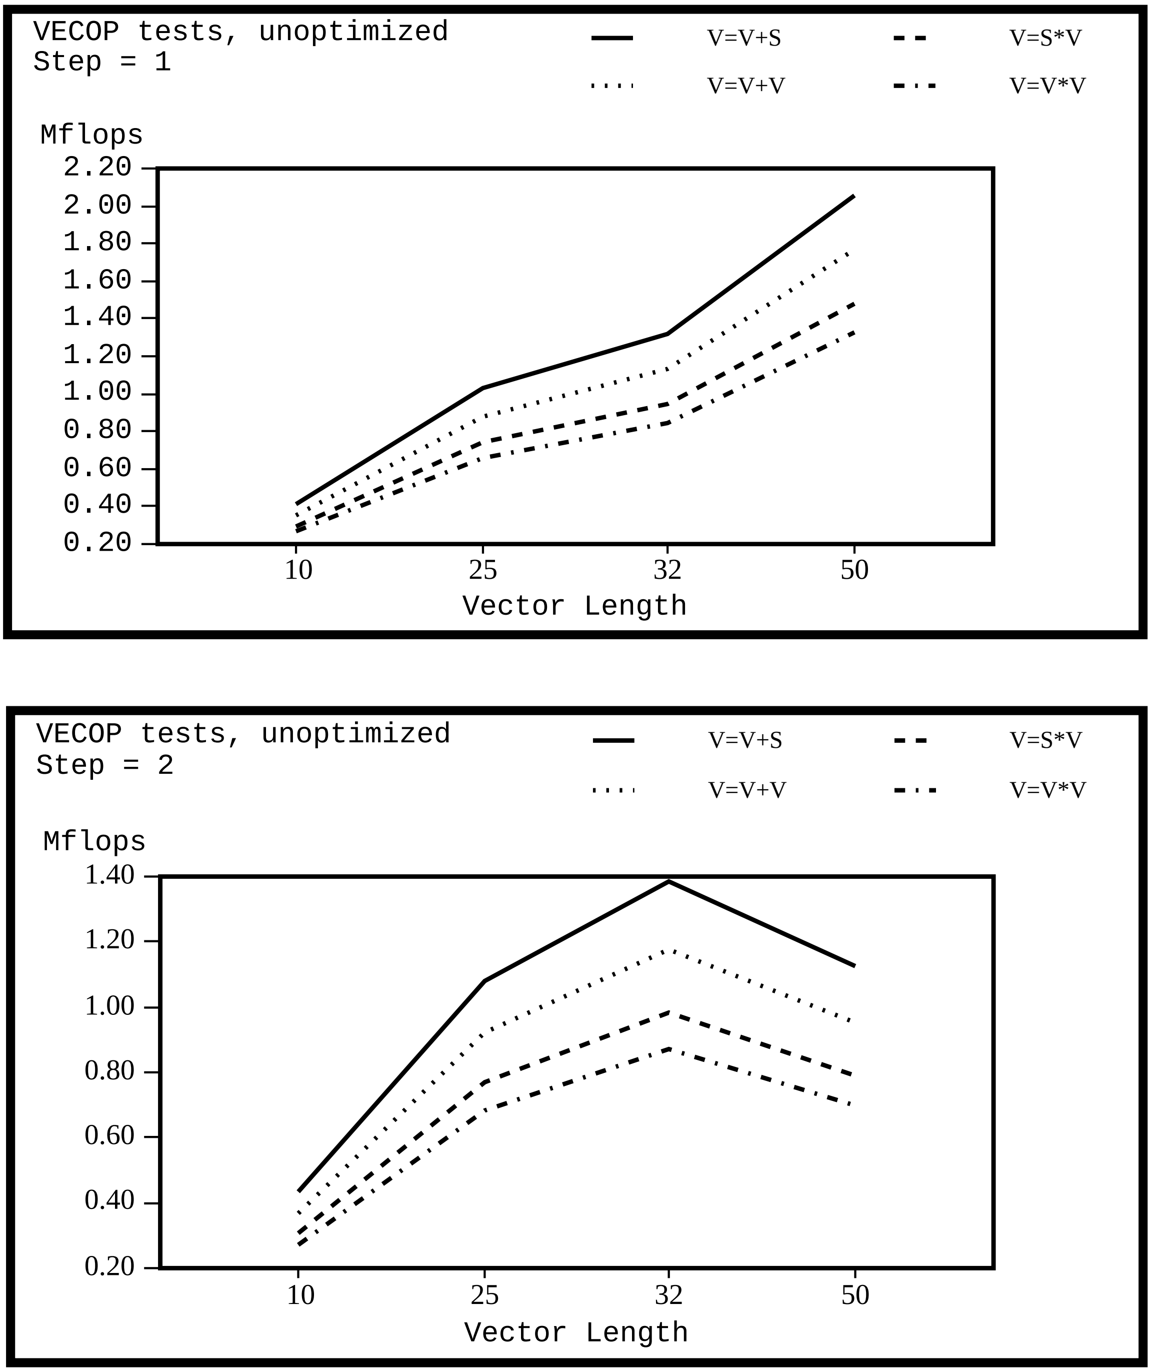

Figure 5-3 VECOP Results, Unoptimized, for Step = 1,2

ure 5-4, when the vector length <= 64, vector elements can be fully distributed among the processing elements. Increasing the vector length allows more

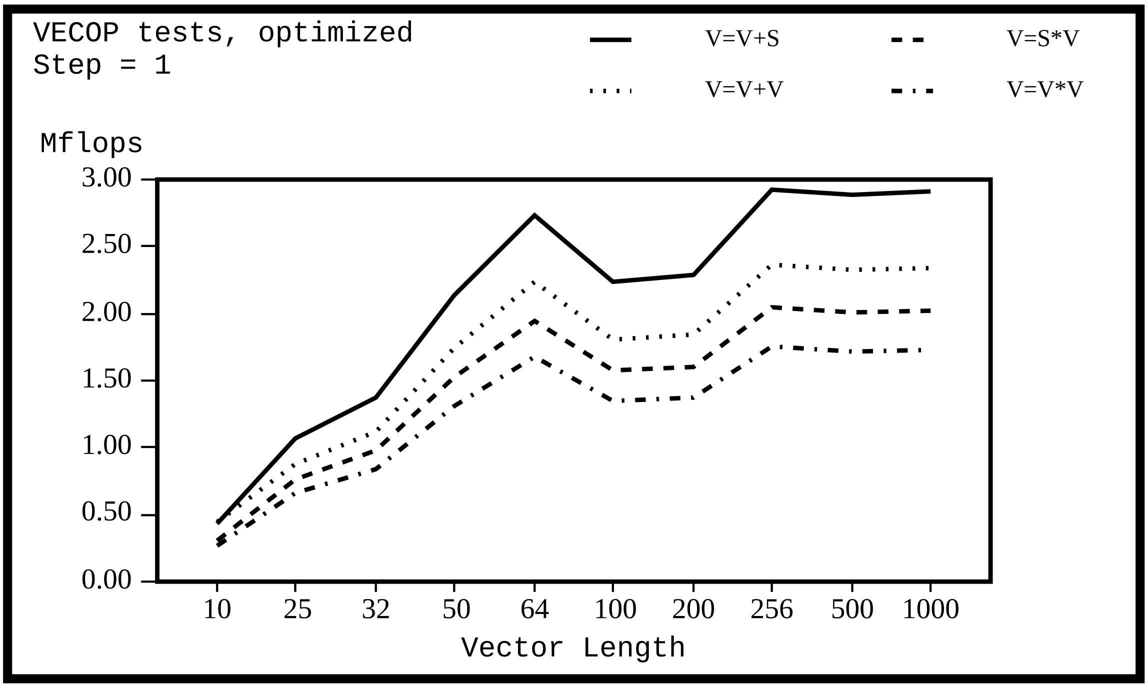

Figure 5-4 VECOP Results, Optimized, for Step = 1

processing elements to execute during a given time period, leading to higher performance rates. The performance peaks when the vector length (n) is a multiple of 64 since all processing elements are busy all the time. When n>64, the multiprocessor must ‘‘strip mine’’, or process data in sets of 64. The performance decreases if during some time period not all processing elements have data to process. This occurs whenever the vector length is not a multiple of 64.

- The PS-2100 has very weak performance on scatter-gather operations. These operations involve indirect addressing of memory locations. A vector contains addresses which are accessed to load (gather) or store (scatter) data values.

Weaknesses of the PS-2100. While a considerable improvement over the PS-2000, the PS-2100 still suffers from a number of weaknesses [Iten90]:

- Scalar processing. A SIMD machine, the PS-2100 naturally exhibits higher performance rates when multiple processing elements operate simultaneously. Not only do scalar operations leave many processing elements idle, but the base performance of each processor is not fast compared with world standards. The LINPACK benchmarks involve significant amounts of scalar computation. The lack of powerful scalar capabilities makes the PS-2100 particularly susceptible to Amdahl's Law.

  Amdahl's Law states that the achievable performance of a computer that has a parallel component and a scalar component is a function of the speed of the scalar unit and the fraction of the computation performed on it, and that the degree of speedup is dominated by the scalar portion. Specifically, Amdahl's Law states that if R is the ratio of the parallel processing rate to scalar processing rate, then the final speedup, S, of a process that is P percent vectorized is

$$\frac{1}{\left(\left[1-P\right]+P/R\right)}$$

In the PS-2100, scalar computation can be performed in one of two ways. The host computer can execute the source program, with only the parallel routines being off-loaded onto the parallel processor. Thus the host computer can execute the non-parallel portions of the code. On a PS-1001 host, single-precision fixed-point addition requires 980 nsec (1.02 MIPS) [Impu89b, 5]. The execution time on a single PS-2100 processing unit is 420 nsec (2.38 MIPS). Thus the ratio R of parallel execution rates to scalar is (64 x 2.38)/1.02 = 149. If a process is 90% parallelized, the total speed-up would be only 9.4, or 6.3% of the theoretical maximum speed-up of 149 which would occur if the process were 100% parallelizable.

Alternatively, scalar computation can be performed on a single processing element. In this case, the ratio R=64, and the total speed-up for the example above would be 8.8, or 13.8% of the maximum speed-up of 64.

- Data transmission system. Although the PS-2100 is an improvement over the PS-2000, the system still lacks a truly high-speed communications system for transmitting data between processing elements.
- Memory bandwidth. In each memory cycle, only one 16-bit word can be read from the processing element.
- Floating-point and double-precision computation. The PS-2100 does not incorporate high-performance hardware floating-point units. Floating-point and double-precision computations require a significantly greater number of machine cycles than fixed-point and single-precision operations.
- I/O remains one of the greatest bottlenecks. According to Itenberg, because of this bottleneck, one is able to attain little more than 10% of the theoretical peak performance of the system. While the system was designed for use with 317.5 Mbyte drives, in practice these drives have not been available. PS-2100 configurations are often equipped with imported Winchester drives made for personal computers [Impu91, 9].

Benchmark data must be handled carefully. Although they are often designed to give results which are representative of performance on a particular class of data, the degree to which they succeed in this has been a popular point of discussion and dispute among computer users and designers for years. While results for the PS-2100 may be considered low in comparison with many commercially available computers in the West, there is at least one example of how such an architecture (in this case a dual-configuration PS-2000) can be used to obtain results which are surprisingly competitive.

In September, 1990, a group from Sandia National Laboratories and Lawrence Livermore National Laboratory visited the Applied Physics Institute in Novosibirsk, which was solving high-speed projectile impact problems on a dual PS-2000 configuration. During the visit, the US group was asked to define a problem to be run on the PS-2000. One of the members of the delegation, Jim Asay, set up a two-dimensional hyper-velocity impact problem which involved the impact of an aluminum pellet into a double-walled shield followed by an aluminum target. The impact of a high-speed object on a satellite would be a similar problem. The description was entered into the computer through a reportedly user-friendly interface. The computation was performed on a coarse grid and was solved in about 12 minutes.

After returning to the US, the scientists reproduced the results on a Cray Y-MP using one processor. They tried two different approaches, giving processing times of 4.1 minutes and 14.25 minutes respectively. The latter used a method that was not vectorized.

The level of resolution was not nearly as fine as on the Cray, and the application was undoubtedly highly optimized (coded in either ASPS or, perhaps, in microcode) for this particular problem. Nevertheless, the machine gave useful results for a computationally intensive problem.

5.4 A Period of Change

In the years following Gorbachev’s appointment, changes in legislation and policy unleashed forces which quickly snowballed, causing dramatic changes in the economic, social, and political environments within which organizations like NPO Impul’s functioned. Change occured in the legislative environment, in the PS-2x00 market, in relationships with R&D sponsors, suppliers, customers, the ministry, factories, and foreign organizations. The changes placed extraordinary pressures on NPO Impul's, created unprecedented opportunities and degrees of freedom to adapt to these changes. We discuss here those changes which appear to have had the strongest impact on the Problem-oriented Computing Systems division in terms of its structure and activities, and the PS-2x00 computers.

5.4.1 Relationships with Suppliers

The development and manufacture of the PS-2x00 line is clearly dependent on the ability to acquire the necessary inputs. These include everything from subsystems to components to cabling to packing crates. During the first few years after the introduction of the Law on State Enterprises (Associations) in 1987 NPO Impul's had relatively few problems acquiring the necessary component base. Here the nature of the component base, the structure of the market for NPO Impul's machines, the existing relationships with components manufacturers, and the timing of development played helpful roles.

The PS-2000 had been built almost entirely with chips in series production at the time the system was developed. Even acquiring such chips during the late 1970s was not easy, for Minelektronprom was overloaded with its own tasks and problems at the time. Significant assistance had to be given to NPO Impul's by the USSR Council of Ministers to establish the necessary relationship with Minelektronprom. Once the link to Minelektronprom factories had been established, however, subsequent interactions could take place largely without the involvement of the higher levels of government. NPO Elektronika in Voronezh' had not developed gate arrays specifically because of needs at NPO Impul's; semi-custom gate arrays can easily serve as building blocks for a large number of different computers. Given existing production of such chips and the establishment of relations between NPO Impul's and NPO Elektronika, however, low-level working contacts between Itenberg’s engineers and their counterparts in Voronezh’ were, for the most part, sufficient to work out the preparation of gate arrays tailored to the PS-2100.

When the PS-2100 was completed, The Elektronika Plant contracted with Impul’s to move forward and develop the next generation of chips. It is unlikely that NPO Impul’s was the only customer for such chips, although the funding for development came primarily from the financing provided to NIIUVM for development of the PS-2300 (discussed below). Although NPO Elektronika was reportedly operating under 100% state orders (*goszakaz*) in 1989, there was at this time much talk about the level of state orders dropping significantly in the coming year or two. NPO Elektronika therefore felt that it was in its interests to work on the next generation of gate arrays, a chip with wide potential applicability and a proven market.

We speculate that timing factors also played a role in NIIUVM’s ability to acquire the chips necessary for the PS-2100 and also the PS-2300. The chips for the PS-2100 had already reached the production stage by the time the new economic mechanisms introduced by the Law on State Enterprises (Associations) began to take effect in Minelektronprom, in 1988-1989. When the contract between NIIUVM and NPO Elektronika to work together on chips for the next generation PS-2x00 was agreed upon, probably in 1988, the prospects for the machines looked good. What would have happened if the PS-2100 prototype had not been developed until 1991 or 1992? A drop-off in demand (described below) for these and other machines using gate arrays would surely have been felt and NPO Elektronika would by this time have been exercising much greater autonomy in determining its product mix. Would it have elected to proceed with development of the new gate arrays? We cannot be certain, but there is a non-trivial probability that they would not have, leaving NIIUVM stranded.

In 1990, the principal shortages were not in components, but in other, less exotic materials. Regarding components, S. N. Krashenko, deputy chief-engineer for new technology at SPZ, commented that, ‘‘[i]t is now possible to obtain everything, but the question is how much it will cost’’ [Kras90, 2]. In shorter supply than components were machine-tool steel, quality wire, good insulation, glass-cloth-base laminate for multilayer boards, and lumber for packaging. ‘‘The component base is right now, I’d say, tolerable’’ said Krashenko. According to Itenberg in September, 1991, the first gate arrays with 17,000 gates were scheduled to be delivered towards the end of that month.

5.4.2 Demand for PS-2x00 Computers

Under the reforms introduced by the Law on State Enterprises (Associations), two changes in particular had a significant impact on factories and, in turn, on the relationship between those factories and their customers and the R&D organizations which developed new technology. First, the reforms greatly increased the ability and need for factories to expand the so-called horizontal links, the links with their customers and suppliers. Second, the system of state planning which dictated what items were to be manufactured, who would manufacture them, where the inputs were to come from and where the outputs were to go, began to recede during the 12th (and final) Five Year Plan, giving way to the system of state orders (*goszakaz)*. While in many respects the state orders served the same sort of coordinating and supply functions as the annual and Five Year Plans (FYP) the total volume of state orders was decreased during the years following their introduction. This was both because of a conscious effort to move away from a plan-based system to a system based more on direct contracts between suppliers and customers, and because of a decrease in military orders stemming in part from efforts to convert military production to civilian.

While state orders were decreasing, direct contracts were not fully replacing orders that earlier had been incorporated into the Plan. As PS-2x00 customers were given greater responsibility for their own financial status and the uncertainty of the economy increased, they became more cautious in their expenditures, reducing the number of orders which they were submitting. While under the former state budget allocations money was allocated specifically for the purchase of computers and could be spent only on computers, under the new economic mechanisms such purchases had to come out of an organization's income. At the same time, the money was no longer ear-marked specifically for the purchase of one type of equipment or another. Computers had to compete with building apartments, providing social services for workers, the purchase of production inputs, etc.

As a result, NPO Impul's experienced a considerable financial crisis during 1988 and 1989 [Shka88, 78]. The market for the high-performance computers was weakening for two additional reasons. First, the PS-2000 was being phased out of production. Organizations planning to buy such a machine would do better to wait for series production of the PS-2100. Second, the accounting practices of Soviet organizations served to reduce demand among organizations that were simply planning to upgrade from a PS-2000 to a PS-2100. The planned amortization period for the PS-2000 was initially set at ten years. An enterprise which, for example, purchased a unit in 1985 would be required to use the system through 1995. It could sell the system, but not simply throw it away. While not, strictly speaking, limiting the purchase of a new machine, it was still legally difficult to replace a PS-2000 with a newer model.

In 1990, many of the traditional customers began placing orders. By September, 1990, orders for about 150 PS-2100 base modules reportedly had been received. Itenberg interpreted this as an indication that customers who, a year or two earlier, had been postponing purchases of new technology were recognizing the benefit of the computers to their operations and deciding to make purchases in spite of the economic turmoil.

In spite of an apparent genuine desire to purchase the machines, most of these customers were not able to carry through their orders. The increasingly chaotic economic conditions proved too strong. Even in September, 1990, Itenberg realized the difference between an order on paper and a purchase. He stated that an actual purchase was dependent on three things: a) the customer actually having the money to purchase a machine, b) the customer having the desire to purchase, and c) the factory's being able to acquire the components necessary to construct the machine. During 1988 and 1989, a significant number of customers who had planned to purchase a computer changed their plans, deciding to purchase the computer not immediately, but a year or two later.

Itenberg's fears about the orders were realized towards the end of 1990. The Soviet financial year coincides with the calendar year, and many of those who had placed orders had done so in anticipation of funds being available in the 1991 budget for their purchase. In a large number of cases, this proved not to be the case. Money was not available, contract prices for a many products were rising, and the potential customers once again became circumspect in their purchasing, waiting to see what the future would bring. Impul's' market was being dragged down by the general economic crisis. Production levels throughout the country were falling precipitously and the level of debt between enterprises was growing rapidly. Few organizations had money available to purchase the new computers, or chose to purchase the fashionable Western personal computers.

A final reason for the drop in demand, not easily measured, is the influx of Western technology. According to A. S. Nabatov, given the opportunity to travel abroad to sign a deal to acquire a piece of Western technology, many potential Impul's customers would purchase the foreign technology rather than buy the domestic product. Of course, in industries where hard currency is not available, this factor plays a smaller role.

By September, 1991, approximately forty PS-2100 base modules had been manufactured but only approximately 20-30 had been sold [Tcha92, 33]. The number of customers willing and able to pay for real machines had shrunk from several tens, to single digits. Particularly serious was the fact that a couple of customers who had ordered fifteen base modules between them had canceled their orders after the machines had entered production.

Although the Law on State Enterprises (Associations) gave factories greater freedom in establishing their own contractual prices, NPO Impul's has kept prices for the PS-2100 at a constant level, in effect reducing their relative price as inflation causes price increases throughout the economy. The cost of a PS-2100 base module plus peripherals was 500,000 rubles in 1989, of which approximately 2-300,000 rubles was for the base module itself [Iten89, 3]. This figure was the same in 1990. In 1991, a base module alone was still being sold for 200,000 rubles. The sale price did not exceed the official ‘‘limit price,’’ set under the state price formation formulas while the machine was being designed, primarily because raising the price would hurt an already very weak market.

At the same time as demand for the PS-2100 was dropping off, financing for the next generation, the PS-2300, remained firm, at least for the time being. The Ministry of Geology remained the primary sponsor, supplying roughly five million rubles per year. According to Itenberg, through 1991 funding was sufficient for R&D work and the manufacture of the number of machines needed by the sponsor. Unlike many other divisions at NIIUVM, Itenberg's division was receiving by some accounts 95% of that which was needed to carry on development of the PS-2300. For this reason, Problem-oriented Computing Systems division did not, as of 1990, have to draw on credit to support the work although NPO Impul's as a whole had had to. Many of the other NIIUVM divisions which had received financing through the Ministry of Electronic Equipment and Instrument-building (Minelektrotekhpribor), the successor to Minpribor, were in catastrophic shape as the ministry failed to obtain the usual funds; some of these were obtaining only 10-15% of what they needed to carry out their research. Thus Itenberg's division was in better shape than most.

The prospects for development were highly uncertain, however. The demand for the current generation and research support for the next are inter-related. If branch organizations decided, for example, to purchase IBM workstations rather than PS-2100s, or otherwise refuse to purchase the PS-2100 computers, the ministry is not likely to continue funding development of such machines.

5.4.3 Relationship with the Factories

NPO Impul's includes two factories, a prototype development facility and the Severodonetsk Instrument-building Plant for series production. NIIUVM has had a close working relationship with both of these units for many years, greatly facilitating the introduction of new machines into production.

During the scientific-research (NIR) and prototype-development (OKR) phases, both the series production and prototype development factories–but particularly the latter–are actively involved in the project. Technicians from the prototype development factory participate in the actual construction of prototypes, but representatives from the series production plant are also involved to make sure that the capabilities of the factory are taken into account. Decisions about the types of cabinets to use, the number of layers in the boards, the types of interconnects, and the precision class of the boards are examples of decisions which depended on the factory. The prototype development team uses the same components, materials, and construction tools as the series production plant. The documentation for series production is written during the prototype construction phase at the prototype development facility.

When a few prototype units have been completed and the production documentation is finished, the documentation is delivered to SPZ where the process of assimilation into production begins. In the case of the PS-2000 this process took a year. For the PS-2100, it took closer to two years. During the assimilation period the series production facilities are set up, modifications are made to the design to make it better suited to series production or to correct weaknesses discovered during the construction of the prototype.

The relationships between NIIUVM and SPZ have been strong, given that they were physically close together and both subordinate to the NPO Impul’s General Director and scientific-council. Production orders for SPZ were worked out not only with NPO Impul’s planners but also with the higher level ministerial and state planning organizations. Getting the factory to manufacture new machines was not always a smooth process, but through the NPO Impul’s ties, NIIUVM was able to work more closely with the factory and have greater influence over its production schedule than could, for example, research institutes in the Academy of Sciences which had to rely on ministerial factory which were, administratively, far removed.

The R&D efforts at NIIUVM were not funded directly from the sales of the products manufactured by SPZ. One or more primary sponsors–in this case the Ministry of Geology–who had a strong need for a given technology funded the entire research and development process. In turn, these sponsors had the right to dictate major systems requirements and receive the first, prototype, units. The cost of development was determined through a series of state normatives regarding the wages to be paid, contributions to other funds, the costs of materials and other inputs, etc.

When the prototype development phase ended, the production documentation was ‘‘thrown over the wall,’’ gratis, to the series production plant which assimilated production. The cost of the series produced units was calculated according to state norms for the cost of materials, equipment, wages, various factory funds including a technological retooling fund, and profit, etc. The profit is used for various expenses such as construction of housing for workers, construction of new buildings, equipment acquisition, etc. These figures are worked out in the Laboratory of Price Formation which establishes a so-called ‘‘limit price’’ for each product. Before the reforms, a system could not be priced above the limit price.

The cost of production did not include R&D costs, however. The principal sponsors of R&D carried this cost themselves, gaining the right to obtain the first units produced. Subsequent customers pay only for the hardware and software; no further payment for the R&D costs of that machine were made. As a result, primary sponsors paid considerably more than subsequent customers.

The significance is that the R&D facility does not profit monetarily from the sale of systems in series production. Although profits from the sale of computers were absorbed into the NPO Impul’s and Impul’s in turn provided money for the budgets for its constituent parts, NIIUVM would receive not an additional kopek if 100 systems were manufactured rather than 50. While NPO Impul’s has a certain fund available to sponsor in-house R&D, it is not enough to support the development of systems like the PS-2x00 machines. Funding for the research and development of the next generation system would again have to take the form of specific, ear-marked funding from a sponsor rather than from proceeds of the sale of previous generation systems. In this respect Impul’s differs considerably from Western corporations which rely heavily on income from sales to fund their R&D programs. Dependent on R&D funding from sponsors or the Ministry, NIIUVM has little access to hard currency, further limiting its ability to acquire foreign technology.

In 1990 serious discussions were held about the possibilities of altering the financial relationship between NIIUVM and SZP such that documentation developed by NIIUVM was sold to SPZ. No concrete changes were made as a result of these discussions. A stated reason was that the accountants at Impul's had not learned to properly evaluate the value of intellectual labor such as prototype development. A more likely reason is that the factory's commitment to many of the new NIIUVM systems was marginal enough that requiring the factory to purchase production documentation could have jeopardized entire classes of machines, such as the PS-2x00.

As the constituent parts of NPO Impul's gained greater autonomy and responsibility for their own affairs, the SPZ naturally was increasingly concerned that there be a market for whatever it was manufacturing and that the manufacture be profitable. No longer was the production schedule determined exclusively by the Plan. Except for orders which were part of state orders, SPZ had to find customers itself. In general, the greatest profits were gained from machines which had been in production for more than three years. The PS-2000 and SM-2M were both in this category. In 1990 it was too early to determine the profitability of the PS-1001 and the PS-2100. Production of the PS-1001 peaked in 1991 at 100 units, declining to 70 in 1992.

The PS-2100 was not the main SPZ product, accounting for only about 2-3% of SPZ's annual gross income. Nevertheless, such orders were not to be dismissed lightly, especially in times when the overall number of orders for the factory for all products was decreasing. But the future of the PS-2100 was not assured and there was a significant amount of pressure on developers to prove its worth to SPZ, to prove that there were customers for the system. When the orders for 15 base modules were cancelled, Itenberg felt obligated to purchase these units from the factory to prevent it from shutting down production completely. The total cost of these base modules, at 200,000 rubles each, was three million rubles, a heavy debt for Impul's and, in particular, for the Problem-oriented Computing Systems division.

As profit grew as a motivating factor for SPZ, it was less inclined to engage in marginally profitable activities. One example of this was the factory's efforts to provide systems maintenance and support for installations in the field. Providing these services required a highly-trained, highly qualified staff. Itenberg's team and engineers at the prototype development factory wrote the production documentation, so SPZ engineers did not develop the expertise necessary to tune and service the machines. The factory did not prove willing to train individuals, since this would require a significant investment of time and resources. Several of the SPZ workers who knew the system best reportedly left the factory to start their own small enterprise.

In short, the relationship between NIIUVM and SPZ through 1991 was stable, but uncertainty about its future was growing. SPZ appeared willing to continue production of NIIUVM machines already in production, but lacked the internal expertise to develop new machines itself.

In 1992, because of a lack of orders for the PS-2100, SPZ ceased production and reoriented much of its capacity to consumer goods for which there was a market, such as telephones, watches and washing machines. While in the past computers constituted 80% of SPZ's production, but the end of 1992 the percentage had dropped to 20%. SPZ became completely independent of NPO Impul's and had no administration in common with NIIUVM.

5.4.4 Relationship With The Ministry

Prior to the perestroika reforms, the relationship between NPO Impul's and its ministry, Minpribor, was similar to that elsewhere in the economy. The ministry owned all Impul's facilities, established the production and development Plans and defined the indicators and forms associated with the planning process, provided funding for operations, and absorbed income generated through the sale of Impul's products. Changes to Impul's structure or activities had to be approved at higher levels in the ministry.

Following a poor performance during the 11th Five Year Plan, Minpribor in 1986 underwent a reorganization [Pano85; Prav860709; Shka86]. In an effort to tighten the link between R&D and production facilities, nearly all research facilities not already part of a production (PO) or scientific-production association (NPO) were incorporated into one [Shka86; Shka86b]. Since NIIUVM was already a part of NPO Impul's, these changes had minimal impact on it.

One consequence of the *perestroika* reforms was that Minpribor began to play a smaller and smaller role in the life of Impul's, particularly in the day-to-day activities. On January 1, 1987, Minpribor became one of five ministries to implement full *khozraschet* in its enterprises and associations, anticipating the passage of the Law on State Enterprises (Associations) which went into full effect January 1, 1988 [Sukh87, 9; Rumy88]. The Ministry had made a transition to *khozraschet* on the sectorial level in 1970, but this was the first time this accounting method had been applied in its full form to individual associations and enterprises [Shka87].

In July 1989, Minpribor was dissolved in a major governmental reshuffling, at least in part to reduce administrative overhead [Gorb89, 1]. Minpribor organizations were predominantly absorbed into the Ministry of the Electronic Equipment Industry (Minelektrotekhprom) to form the Ministry of Electronic Equipment and Instrument-building Industry (Minelektrotekhpribor), although some institutes and factories were absorbed into the Ministry of the Radio Industry (Minradioprom). Under these circumstances, the new ministry leadership played an even smaller role than that of Minpribor and was, according to Itenberg, ‘‘barely noticed’’ by those working in the high-performance systems division. The high-performance development was financed chiefly by the Ministry of Geology, but other NIIUVM research had been financed by Minpribor. Following the absorption of Minpribor into Minelektrotekhpribor the level of financing dropped however, and with it, the involvement of the ministry in the activities of NIIUVM.

5.5 The Response to Change

The changes described above created both significant crises and new degrees of freedom to shape the mission, structure, and technology at NPO Impul's.

As the economic crisis deepened and demand for NPO Impul's products dropped, survival became a basic goal. But the goal was more than just survival at any cost; survival meant preserving the capability of staying in the same principal niche–control and high-performance computing systems–that NPO Impul's had occupied for decades. During the plenary session of the 35th Anniversary Jubilee Conference at NIIUVM, general-director V. G. Rakitin spelled out his goals in an address to Impul's employees and users of Impul's computers [Impu91, 1]:

> We are joined by one work. You are and will remain our users and technicians. We'll retain our orientation, because our work is needed in our economy. Without electricity, chemistry, automation, the economy can't function. Our task in these difficult times is to preserve the collective, preserve the direction. We will work together to raise the quality of work to a more competitive level.

In this section we examine efforts to preserve the collective, to preserve the direction of development with the Problem-oriented Computing Systems Division by focusing on two areas of change: modifications in organizational structure, and changes in the high-performance computing systems.

5.5.1 Changes in Structure

5.5.1.1 Traditional NIIUVM Structure

Traditionally, scientific-production associations drew together under one administrative roof organizations which covered the entire product life cycle: R&D facilities such as a scientific-research institute (NII), design bureaus, prototype development facilities, and manufacturing plants. Additionally, as in the case of Impul’s, the association could incorporate training centers, associated laboratories, and other organizations. Production associations in general were created for a number of reasons, including administrative economy, achievement of economies of scale, promotion of the status and power of managers of very large enterprises, and the consolidation of a given sector. Scientific production associations in particular were designed to forge a closer link between research and production facilities [Nove86, 69-70].

Throughout the Soviet Union the structure of industrial research institutes has been cast from the same mold: an NII consists of multiple divisions, each consisting of multiple laboratories. Each division is responsible for its own line of work. At NIIUVM, I. I. Itenberg is the head of the division of problem-oriented high-performance computer technology which consists of approximately 75 people, including systems engineers, software engineers, lay-out engineers, operators, and others.

The divisions are subordinate to a scientific-technical council, which consists of the division and sub-division heads and deputy-directors of the institute and is chaired by the general-director. The general director and the scientific-technical council form the common administrative point between the NII and the other component organizations of the NPO.

Deputy-directors, such as V. V. Rezanov, were responsible for coordinating the activities of the various divisions and for allocating resources to each. Rezanov was responsible for ensuring that the various divisions all adhered to the same interface standards so that their devices would be interoperable, in keeping with the ASVT philosophy. If Itenberg and his engineers wished to depart from such a standard, they would have to get permission from Rezanov and the scientific-technical council. He also distributed financing to each. Requests were, of course, submitted by the divisions themselves and refined by the institute's economists who saw that all figures adhered to the myriad of economic norms established by the government and ministerial planners. But the link between the request and what was allocated passed through Rezanov, who was a critical link in the financial chain between the sponsor and the R&D departments. Rezanov and the general-director were also key players whenever new ties were being established with a supplier organization, such as with Minelektronprom factories, but once the ties were established, most contact took place between technicians in the respective organizations.

Itenberg's division was subordinate to Rezanov, but from the outset it had greater independence than did many of the other divisions at NIIUVM. A principal reason was that the PS-2x00 computers represented a qualitatively new line of development from the traditional, control system research. As such, the research was more self-contained than in other divisions.

Prior to the *perestroika* reforms, this structure was virtually invariant. It was created according to the so-called state schedule (*struktura shtatnogo raspisaniya*). The creation of new divisions and laboratories had to be approved at ministerial levels, sometimes even at the level of the USSR Council of Ministers. Similarly, the appointment of individuals as directors, deputy-directors, heads of divisions and laboratories had to be approved by the ministry and, especially in the case of director and deputy-directors, by the Council of Ministers. The job titles and their associated pay levels were all fixed by the government.

The laboratory structure was virtually invariant. Once created, laboratories were seldom broken up. Once assigned to a laboratory, a researcher moved out of a laboratory only through promotion, dismissal, or voluntary departure; rarely did someone transfer laterally to another laboratory.

The head of each division was also the chief engineer for the work carried out. He had to adhere to technical policy set by the scientific-technical council and the general-director and deputy-directors.

The chief engineer had full responsibility for work carried out in his division, and was the final arbiter for technical and procedural questions which did not need approval at higher levels.

5.5.1.2 Changes to the Structure of the Problem-oriented Computing Systems Division

The reforms placed nearly complete authority over the structure of the association and the institute into the hands of the institute’s directorate. No longer were approval or resolutions from the ministry or the Council of Ministers needed to create new laboratories or divisions, or alter the structure of the existing ones. With this freedom, this degree of authority, came the need to address the questions of what would be the best form of organization for NPO Impul’s, NIIUVM, and the constituent divisions. While the basic laws allowing changes to the structure were passed in 1987 as part of the Law on State Enterprises (Associations), the existing system had considerable inertia, and it wasn’t until 1990 that new organizational forms began to appear at NIIUVM.

The transition towards a market economy placed increasing financial pressure on NPO Impul's. All recognized that work would have to be carried out more effectively, and many viewed structural change as key to accomplishing this. Specifically, many felt that greater effectiveness could be achieved by increasing the independence of individual divisions and expanding the authority of the division heads over the scientific-technical and financial activities. This would, it was argued, increase the effectiveness of decision-making, increase personal responsibility for these activities, and reduce overhead.

Gradually, authority was decentralized. The first major change to NIIUVM following the reforms was the creation of so-called temporary scientific-technical collectives (VNTK). In contrast to laboratories which were long-term structural entities which carried out a variety of activities for many, many years at a time, the VNTK were task or program oriented and much more flexible. A group drawn from different laboratories, different divisions, or even different institutes would be created to address a given project, and that project alone. When the project was completed, the collective could disband and its members join new and different groups being formed to work on other projects. The VNTK organizational form dates back to an August, 1983 Council of Ministers resolution, but had been used only infrequently, for relatively large-scale programs such as the START new generation project discuessed in chapter 6 [Fort90, 112]; Itenberg sought to apply this form on a much more modest scale, to individual tasks.

Itenberg's division was the first at NPO Impul's to experiment with the VNTK, creating the initial one in November, 1990. The idea arose when tasks came up for which Itenberg's division did not have the necessary personnel and expertise. At the same time, some of the other divisions had the needed expertise and less than a full workload. A key enabling factor was the fact that responsibility for budgets was being pushed down into the divisions themselves as the *khozraschet* mechanisms took hold more fully. While in the past laboratories and divisions had been very protective of their workers, they were now more willing to let them work on someone else’s project and get paid through that project. As Itenberg stated, ‘‘This allowed us to obtain a certain degree of independence. Now we ourselves set the rate of pay, we ourselves determine the incentives, we ourselves hire people, we ourselves can attract people from other quarters. It’s more flexible.’’

By September, 1991, Itenberg presided over nearly half a dozen VNTK. The principal ones were for the development of personal geophysics computing systems (PGVK), the development of software for the PGVK, and the installation and tuning of a PS-2100 system with ten base-modules. Additional VNTK were devoted to integrating a single base module with a personal computer and to developing systems software. A VNTK for installing and tuning a PS-2100 with six base modules had already completed its work and had been disbanded.

The organizational structure and the nature of some of the tasks carried out by Itenberg’s workers changed in tandem. Tasks which could increase the income of the division were desirable. One change, for example, involved the services provided to users. Traditionally, once the prototype machines had been built and installed at the principal sponsors, the Itenberg’s division turned over all responsibility for installation and maintenance to the manufacturer, SPZ. During 1990, two basic changes occurred which caused the Problem-oriented Computing Systems Division to become more heavily involved. First, several of the individuals at SPZ who had been primarily responsible for this work left the factory to start their own small enterprise. As a result, the factory was less capable of installing and maintaining the systems it was manufacturing. Second, Itenberg recognized that such work could be done by them on a contractual basis, bringing more income directly into his own division. Members of the division began working with the small enterprise on installation, debugging, and system verification projects.

After the directorate had given Itenberg permission to form a VNTK, others followed suite. By September, 1991, VNTK had spread throughout NIIUVM, to the extent that the division-laboratory structure had been replaced in practice by the VNTK as the dominant organizational form. The former divisions and laboratories still existed administratively, but the work was being carried out in VNTK. To be sure, some divisions like that working on the PS-3x00 made the transition by transforming an entire division into a single VNTK. But the principles of flexibility and fluidity increasingly dominated thinking about organizational structure and even these were considered temporary in nature. The only two divisions which continued to operate under the old structure were the scientific-technical information and standards divisions where the need for flexible working groups was less pronounced.

The issue of which organizational form would be the most appropriate was still under discussion in September, 1991. Since the creation of the VNTK, in August, 1990, a law was passed allowing the creation of a new organizational form, the small enterprise. In the months following its passage, a growing body of discussion and experience had been disseminated through the mass media. Soviet economists also had been promoting such organizations. In additional, several small enterprises were created in Severodonetsk by people who formerly worked at Impul's, with good results. Inspired by these examples, individuals in Itenberg's division–and throughout NIIUVM–grew more convinced that the old system had to be replaced by a more progressive one in which wages were more closely tied to the amount of work and the profitability of the organization. The risk that such efforts could fail was acknowledged, but many felt that the old system with its dependence on centralized authorities would not serve them well in the future. The possibility of creating small enterprises at NIIUVM was first raised in the beginning of 1991.

An arrangement in which independent collectives existed under the roof of the old organization had benefits for both. The independent collectives could use the parent organization's name recognition to acquire contracts more easily, would have greater access to production facilities than if they functioned in complete isolation, had a chance of getting internal or government credits, and could use the same material base as before. The parent firm would profit because the collectives, fighting for survival, would bring in profitable orders, they would pay rent to the parent organization, and the directorate of the parent organization would not have to concern itself with management or personnel issues within the collectives [Pivo92].

Many questions remained to be solved, however. The precise nature of the small enterprise had to be established. At least two types of small enterprise could be used. One variety could employ 50 people or fewer. Another, called a scientific small enterprise, could employ up to 100 people, or up to 250 if it were involved in production [Ezh9008]. Second, the precise nature of the relationship between the small enterprise and Impul's was still unclear. For example, if the small enterprise were a separate organization with no possessions of its own, renting facilities and equipment from Impul's, who would provide such social services as a kindergarten for employees' children? If an employee of Impul's had been waiting ten years for an apartment, would his or her status be jeopardized by joining the small enterprise? Many such questions had to be addressed.

A more fundamental issue had not been resolved, however. While decentralizing the structure of NIIUVM could lead to more effective work, it also threatened to destroy the integration between divisions necessary to carry out the core activities of the institute: the development of complete process control and high-performance data processing systems.

Allowing each division to drift along its own path threatened to undercut the core capability of the institute. For this reason, NIIUVM leadership was resistant to decentralization, but felt the financial pressures to do so accutely.

During 1991 and 1992 recommendations were made to the Directorate to implement a large measure of decentralization in the form of rental collectives (*arendnyy kolektiv*) without granting complete independence to individual divisions [Cher92]. Still a part of NIIUVM, such collectives enjoy increased financial independence, and have a formal agreement with the institute to rent facilities and equipment. They have their own bank accounts so that income received from customers would pass directly to the collective rather than through the institute accounts. They have control over wages, but must earn the money to pay them. They also have the right to fire workers. In non-financial respects, however, the workers continue to operate as employees of NIIUVM. Heads of the collectives are still subordinate to the director, and decisions of the scientific council were binding on them. In August, 1992, an order was issued to implement a number of rental collectives [Cher91]. By the end of 1992 there were five: high-performance computing, general-purpose systems on the basis of the PS-1001, controllers, publishing, and CAD systems. These constituted the core of the institute. Small enterprises did exist, but for peripheral tasks such as engaging in commercial ventures to supply workers with foodstuffs. The institute had a dual structure.

However, there was sufficient discontent with this order that a commission was created to resolve the issue of institute structure. The commission could not resolve differences of opinion among its members and instead of one compromise recommendation, several were offered. These ranged from a centralized structure in which the director had unified control over property, finances, and management to a loosely-coupled collection of units, each with full legal status and control over financial and management issues.

Separately, Rezanov argued passionately that the only way NPO Impul's could survive was if the association achieved a level of self-financing in which the sales of products to customers would generate a profit sufficient to cover the costs of R&D and the operations of the firm. The products with the greatest earning potential were the integrated process control and data processing systems and associated services. The requirements of these systems should determine the structure of the association with production, not R&D, as the core activity. Given the decentralization that had already taken place, he proposed a hybrid structure in which each of the rental collectives worked both on tasks related to the integrated systems under the coordination of the association's leadership, and cultivated individual market niches in which they had complete control [Reza92].

The issue still not been resolved. At the time data collection for this study ended in December, 1992, the issue had yet to be put before the expanded NIIUVM scientific council for a final decision [Cher92].

The issues of decentralization and flexibility versus integration and maintenance of capability were pressing issues at all levels withing NPO Impul's. As greater autonomy and flexibility were given to the divisions and collectives within NIIUVM, Itenberg's primary goals remained close to those stated by the general-director: to preserve the collective, to keep the PS-2x00 development team intact so that the line of development itself could be preserved. He used a three-prong strategy: incorporate flexible teams to address the division's tasks, retain leadership over the work of the division as a whole, and pursue opportunities for earning money outside the division's traditional sphere of activities.

Facing weakening demand for his machines, low profit margins on those few units which could be sold and growing uncertainties about future research funding, Itenberg had to try to find ways to keep his development team intact. Key to keeping his principal engineers was finding enough money to pay them a wage which could support them and their families.

A possible solution was to try to put non-core individuals to work on projects which would earn sufficient money for the division to support the engineers working on high-performance systems. In particular, if they could start up the manufacture of low-cost, low-overhead consumer products which enjoyed a large market, the twin goals of providing work for all and supporting the key lines of work could be achieved. The list of possible products is broad. They would be willing to manufacture anything from medical equipment to electronic games to coffee grinders to door-handles for cars. Any product would be considered, provided it would enjoy an extensive market and could be manufactured at NIIUVM in conjunction with the prototype development factory.

To cultivate the market for the division’s products and explore possibilities for new products Itenberg therefore chose to do more of his own marketing. The chief changes in ‘‘marketing’’ involved advertisement and pricing. Promotional literature (prospectuses) had always been printed, but they were now being printed earlier in the product development cycle. Itenberg gained much greater flexibility in setting prices. He sought to preserve and expand the customer base by 1) maintaining close contact with current and past customers, providing new equipment and maintenance, 2) identifying new customers through communications with old customers, 3) participating in conferences where he could both publicize the machines and learn of new applications to run on them. With new applications come new potential customers.

The plan favored by Itenberg and others in 1991 was to create a small enterprise within the framework of NPO Impul’s specifically for the production of consumer goods. Ideally, this enterprise would establish a joint operation with a foreign firm which would provide investment funds, or simply technical licenses or know-how to expedite the production of specific goods. Because the prefered organizational form at NIIUVM during 1992 was the rental collective and many issues surrounding the role of small enterprises had not been adequately resolved, Itenberg's division was transformed into the former.

The VNTK structures within the division changed little during this year, although their tasks changed as funding for PS-2300 development ended. The team doing hardware development began working on a new generation of control computers, based on Western microprocessors. The software development team worked on applications for existing PS-2100 users.

5.5.2 Changes in Technology

In spite of the many changes precipitated by the *perestroika* reforms, the core trajectory of the high-performance machines has remained very constant. Besides basic improvements in the traditional performance, reliability, and functionality, changes to the PS-2100 and its successor have been relatively small-scale, focusing on making the machines more suitable to a broader circle of users, while remaining within the framework dictated by the nature of their relationship with the principal sponsor and their traditional customer base.

5.5.2.1 The PS-2100

Modifications were made to the PS-2100 configurations to improve the user interface and adapt the machine for use by a wider circle of potential customers. Using a PC AT bus–IUS (‘‘universal systems interface’’) adapter, a personal computer could be used as the primary user console. One such configuration, called a personal geophysics computing system, consisted of an IBM PC-compatible personal computer used as a host machine attached via the bus adapter to a PS-2100 base module, a PS-1001 monitor subsystem, and a plotter [Iten91b, 10]. The same adapter could be used in larger configurations, attaching directly to multiple PS-2100 base modules, the external exchange switch, magnetic disks and tapes, and a number of additional personal computer user consoles [Iten91b, 11]. Other developments under way at the time of this study include the incorporation of standard Token Ring and Ethernet networking protocols.

The PS-2100 was being adapted to other host computers as well. The geophysicists at NIIGeofizika had talked to representatives from IBM who wanted to expand IBM's operations in the Russian oil industries. During the course of such discussions, the question arose about whether or not it was possible to use the RS/6000 workstation as a host to the PS-2100 multiprocessor. It was felt that this would be a way of both selling IBM workstations and the PS-2100 in Russia [Tcha92, 33]. Initial discussions took place during the summer of 1991 and in September A. S. Nabatov visited IBM's Norwegian subsidiary to explore the issue further. These discussions ultimately proved unfruitful.

To increase the range of machines available, a small, desk-side version with 32 processing elements was developed.

5.5.2.2 The PS-2300

Shortly after the PS-2100 prototypes were completed, work began on successor models. Designers explored the possibility of overcoming some technological limitations by building a machine consisting solely of imported components. When it became clear that such a project would not be feasible because of the difficulty of getting Western components and information about them, work began on an indigenous successor to the PS-2100. The aborted project was called the PS-2200 even though it never passed the design phase; efforts were redirected towards the creation of the PS-2300.

As had been the case with the PS-2100, the PS-2300 project began with an analysis of the weaknesses of the PS-2100, the requirements of users–principally of the geophysicists

	PS-2300/2	PS-2300	PS-2300/1	PS-2300 full configuration
Processing elements	16	64	256	2560
Theoretical Peak Performance (Mflops)	60	250	1000	10,000
Main Memory (Mbytes)	2	40	96	1000
Housing	desk top	desk side	cabinet	10 cabinets

Table 5-6 PS-2300 Configurations

who remained the primary sponsors–and the components and other technologies projected to be in production by the time the PS-2300 was to be completed.

Since the principal sponsors and applications remained the same as for the PS-2100, the basic requirements for the PS-2300 were simply to improve the basic operating parameters: increase the processing rate, improve reliability, increase the amount of internal and external memory, provide faster I/O, reduce the physical space, etc. Although not required by the principal sponsors, several measures were being taken to increase the potential market. These included development of a family of PS-2300 configurations, shown in table 5-6,ranging from a desk-top model to a full-scale multicabinet model, the incorporation of non-indigenous computers such as MS-DOS based personal computers and Unix based workstations as the host computers, the incorporation of standard network protocols such as the Token Ring and Ethernet and the use of the IEEE 754 floating-point format for data compatibility with other computers. According to Itenberg, ‘‘There was no such requirement [for using the IEEE 754 standard]. We simply are aware of the state of affairs in the development of computer technology and don’t want to be left without in our country, although the implementation of this is a bit more complicated..." [Iten91, 8].

The practice of providing compatibility at the level of ASPS (but not at the binary level) was continued. As had been the case with the PS-2100, this would involve recoding the software libraries and some of the systems software, but relatively little additional effort on the part of user applications developers.

A key development making such configurations possible was the development of an improved gate array component base. Gate arrays had been a boon to the PS-2100 project for they enabled designers to customize chips on a base of series production chips with far less effort–both technical and political–than would have been required to develop fully-custom chips. The CMOS gate arrays projected to be available in series production around 1993, the anticipated year of completion for the PS-2300 prototype, have 17,000 gates, over 40 times more than those used in the PS-2100 [Impu91, 9; Impu91b]. With such chips, an entire processing element could be put on a single chip, and four chips could be put on one board. Development of gate arrays funded through NIIUVM took place at the NPO Elektronika in Voronezh', Russia. The first 17,000 gate chips were to be delivered during the fall, 1991.

In designing chips based on these gate arrays, NIIUVM developers were able to integrate 64-bit operations more fully into the processing elements, such that 32- and 64-bit floating-point operations both execute in the same number of cycles [Impu91b, 1]:

	PS-2100 PE	PS-2300 PE
Number of PEs per board	1	4
Peak processing rate (Mflops) on reg.-reg.		
32-bit floating-point addition	1.0	4.0
32-bit floating-point multiply	0.5	4.0
64-bit floating-point addition	0.12	4.0
64-bit floating-point multiply	0.07	0.5

In addition, the PS-2300 will incorporate 128K-512K bytes (depending on whether 256 Kbit or 1 Mbit memory chips are used) of buffer memory per chip to facilitate I/O operations.

As with the PS-2100, designers reconsidered the question of modifying the processing element interconnect schema. Once again, as the design progressed, the decision was made to continue with the original linear-ring interconnect. The new gate arrays provided 128 data outputs per chip, too few to incorporate four-way interconnects easily.

The shift to incorporating the MS-DOS and Unix operating systems is the result of a strategy to use systems more common than the PS-1001 as the host computer. According to V. A. Largin, it is not clear what the future of the PS-1001 will be. Personal computers are widely available and the ability to attach a PS-2x00 parallel processor to IBM PC compatible machines would increase the potential user base. The MS-DOS machines are only single-user, single-task computers, however. Some sort of multi-user, multi-task operating system should be used. The natural choice is a widely-used operating system which runs on commonly available workstations. Unix is the natural choice.

The PS-2300 is an incremental evolution of the PS-2100. The basic requirements and target applications remained intact, and, at least through 1991, funding through NIIGeofizika, the principal sponsor, remained stable at levels consistent with prior years. The basic architectural decisions remained unchanged, while the operating and performance characteristics were improved, chiefly through the incorporation of an improved component base and construction technologies.

As has been noted, several measures were being taken to improve the ability to adapt a configuration to the needs and technology of the widest selection of users possible. The need to improve the marketability of the system was becoming more urgent at the end of the 1980s and early 1990s, but the basic philosophy was a continuation of that developed in the ASVT and earlier PS programs.

At the end of 1991, work on the PS-2300 came to a hault. The geophysicists who had been funding research did not have the money to continue support. Consequently, NIIUVM did not have the funding necessary to support the development of the new generation of gate arrays at NPO Elektronika. At the time of this writing, efforts were being made to secure some support for the PS-2300 through the Ukrainian government, but it was not clear whether these would be successful. Even if they were, there was a strong possibility that the engineers at NPO Elektronika would have found other work and would no longer be interested in working on PS-2300 chips.

A complicating factor was that Voronezh' is in Russia, now a different state. Thanks to difficult relations between Ukraine and Russia, movement of capital between the two countries is currently difficult and re-established working relations with NPO Elektronika could be problematic. There are no firms in Ukraine which can manufacture chips with the necessary levels of integration, and the PS-2300 depends on being able to place an entire processor on a single chip. Consequently, Itenberg's engineers were considering using general-purpose Western chips instead, should the PS-2300 project be restarted. Because of the specifics of the PS-2x00 processing elements a Western microprocessor would have to be augmented with a set of customized chips, but the latter could have a lower level of integration and, in principle, be manufactured domestically.

5.6 Discussion

The Problem-oriented Computer Systems Divisions within the Scientific Research Institute of Control Computers (NIIUVM) has developed high-performance computing systems with original designs for over a decade and a half. Although not the most powerful Soviet machines, the PS-2x00 computers throughout the 1980s constituted perhaps the most successful line of Soviet high-performance computers, enjoying high production rates, relatively short R&D cycles, considerable popularity among users, and generally good operating characteristics. The PS-2x00 are an exception within Soviet HPC, having R&D cycles considerably shorter than a decade. The PS-2000, prototyped in four years and assimilated into series production a year later, is unique among Soviet high-performance computing efforts in the speed with which it was brought to fruition. A comparison between this machine and others is instructive in identifying factors which impact the nature and speed of development within the Soviet context.

The third generation of PS-2x00 systems was under development through the end of 1992. Over the years, this family has grown steadily in performance and functionality, but has maintained a high degree of continuity and focus in underlying design and construction from one generation to the next, even during the turbulent *perestroika* era. In this section we summarize the many factors which have shaped development of these computers.

During the reform period, the structures of the Problem-oriented Computer Systems Division, NIIUVM, and NPO Impul's as a whole have undergone considerable change. Organizational structures became significantly more fluid, changing frequently, and the ability to make decisions about the structure, operations, and finances to a considerable degree was decentralized, placing much more authority and responsibility in the hands of the division heads. However, the issue of organizational structure has been a point of considerable debate and had, at the time of this writing, not been settle conclusively. In this section we summarize the factors affecting technological advance and organizational structure within the Problem-oriented Computing Systems division and, to the extent that it affects HPC, within NIIUVM as a whole from the perspective of the conceptual framework introduced in chapter 2. We will analyze the impact of change since the start of the *perestroika* reforms and discuss options for the future of HPC at NIIUVM.

5.6.1 The Technology

Table 5-7 gives an outline of some of the many factors shaping the evolution of the PS-2x00 series. The PS-2x00 family evolved within the context of a belief system–a set of guiding principles–which was in part shared by all R&D divisions of NIIUVM and in part unique to the Problem-oriented Computing Systems Division. A foundational belief was that the machines should be very strongly oriented towards industry, towards actual use in existing applications, rather than as a test-bed for theoretical ideas. The implication was that the machines should be developed as a system, with an integrated, complete set of hardware, software and peripherals, and should be developed and manufactured in a timely manner; simply developing an interesting and complex computational engine with good theoretical parameters would not be acceptable. A second belief characterizing NIIUVM development as a whole was that the systems should serve as broad a customer base as possible. In Western capitalist economic systems such beliefs are driven by basic profit motives: the more you sell the more you earn. Under the Soviet system, an R&D facility like NIIUVM would not earn more from increased sales of a machine, but would increase its stature and political and financial support base. Since users of NIIUVM machines had a wide variety of special-purpose control and computational applications, the strategy which evolved was to build modular systems which could be configured in a wide variety of ways, according to the specific needs of a given user. This strategy eventually became part of the philosophy of machine design and modularity, with its requirements for uniform internal and external hardware and software interfaces, is a strong characteristic of nearly all NIIUVM computers. As the centralized directive

Environment Directive form of economic management **Increasingly market driven** Monopolistic infrastructure High-level pressure to develop PS-2000 quickly Extensive set of upstream industries Indicator-based insentives for upstream enterprises **Incentives based on market signals** Stable user requirements: High performance for core set of data parallel tasks, real-time capability, high reliability, moderate computational precision, upward compatibility between generations Stable funding for R&D **(until 1992)** Close links with supporting industries, especially with production facilities. Links based on administrative ties, NPO structure **Weakening links with supporting industries, increasingly based on mutual interest and persuasion** Strong market **Very weak market** USSR integrated politically **Ukraine-Russia transactions problematic**	Technology SIMD-oriented architecture, Reconfigurable bus interconnect, Autonomous indexing, Predicate processor, Distributed memory, Share secondary storage, etc.
	Organizational structure (NIIUVM) Traditional NII division-oriented structure (HPC Division) Traditional structure based on laboratories **Flexible Structure based on VNTK, rental collectives, with shared organizational services** Largely functional division based on core systems, subsystems, tasks **Structure oriented towards new tasks**
Technological availability Examples and ideas from West: ILLIAC IV, STARAN Reliance on components and subsystems in series production. Generally available in sufficient quantities. Extensive computer development experience Some reliance on new technologies (i.e. gate arrays) **Expanded opportunities for use of Western technology. Still difficult to acquire.**	Beliefs (design principles) Build fastest machines possible Develop systems for real users in industry As much as possible, use availble technologies, and limit use of immature technologies (Don't seek to be an industry driver) Develop systems as quickly as possible Serve broad customer base Don't duplicate technology of other NIIUVM divisions
Organizational slack Steady funding through 1991 **Decreasing/terminated funding for HPC** Integrated funding stream for HPC **Greater reliance on small-scale/contract work**	Strategy Employ many moderately powerful processing elements in SIMD-based arrangement Build modular systems Incorporate existing hardware/software to degree possible Seek out a variety of customers

Table 5-7 Factors Influencing PS-2x00 Evolution

system gave way to a more market-oriented system, serving a broad user base still served the institute's interests.

A third basic belief, specific to the special-purpose computing division, was that high performance was a major design goal, and that high performance should be achieved through the use of parallelism, specifically through a single-instruction multiple-data (SIMD) approach.

Clearly there were some trade-offs to be made to accomodate the three pillars of the design philosophy. Achieving higher performance often comes at the cost of lengthened development cycles. But in the early days of development, several environmental factors strongly influenced the development strategy. The combination of an intense national energy campaign and Western restrictions on the export of computers to the Soviet Union increased the pressure to build a high-performance computer for seismic applications in as short a time as possible. Key strategic decisions about development arose out of the guiding principles and these environmental factors in particular. First, the PS-2000 would not serve as a driving force for support industries such as Minelektronprom; it would be built using components already in existence in order to speed up development and reduce the time, effort, and expense needed to develop or acquire customized chips and materials. Second, the machine would be designed as a problem-oriented parallel processor attached to a general-purpose host, thus reducing the hardware and systems software development effort, preserving much of the operating environment familiar to users of NIIUVM products, and making it possible to use the machine within the framework of the aggregate systems of computer technology and software (ASVT and ASPO).

A third environmental factor which strongly influenced the speed with which the PS-2000 could be brought into production was the relatively close relationship between NIIUVM, the prototyping factory, and series production facilities at SPZ. The ‘‘administrative gap’’ between these entities was small, they were geographically proximate, and the channels of communication and cooperation were in place long before the PS-2000 went into series production. The NIIUVM projects reflect the benefits and difficulties of conducting research and development within an important organizational form in ministerial branch science, the scientific-production association. They demonstrate the benefits of a relatively close, long-term association between R&D and production facilities and a strong industrial orientation, but also reveal some of the problems inherent in the ‘‘union of independent organizations,’’ particularly during a period of greater decentralization.

The primary sponsor for the machines, NIIGeofizika, was a strong source of environmental influence over the PS-2x00 machines. Its applications determined the basic set of requirements and the SIMD approach was verified as very applicable to many seismic applications. One of the key design decisions–the use of a reconfigurable bus interconnect–was agreed upon only after specialists had confirmed that it would support the primary applications.

NIIGeofizika remained the primary sponsor through the end of 1991 and its requirements and funding provided a stable, long-term foundation for development. Its on-going sponsorship was one of the important factors underlying a key characteristic of the PS-2x00 family: the continuity in architecture and design from one generation to the next. Improvements were largely incremental, directed at increasing processing speed, improving computational precision, providing greater amounts of primary and secondary storage, releaving bottlenecks such as I/O throughput, improving applications and systems software, etc. Even the shift to a MIMD/SIMD architecture in the PS-2100 simply reflects the aggregation of multiple basic SIMD modules to achieve higher total performance. The basic architecture proved satisfactory for the primary applications; there was little need to alter the fundamental approach. A related environmental factor was the need to maintain a level of software compatibility between one generation and the next because of the growing installed base of PS-2000 computers.

During the late 1980s and early 1990s when demand for the PS-2100 grew quite weak, developers were nevertheless focusing on developing yet another generation, more expensive, with much greater memory and numbers of processing elements, etc. Supporting this development was not investment by NPO Impul's with the prospect of a large market for such machines but the continuing support of the primary sponsor.

Seismic applications were not the only ones to influence development of the PS-2x00 line. A strategy stemming from the desire to cultivate as broad a market as possible was to incorporate features which would be useful to users outside of geophysics if they did not conflict with NIIGeofizika requirements. The atomic energy industry in particular had strong requirements for real-time features and reliability as did the space industry which used the systems for real-time processing of satellite data and control. Other potential users (as well as the geophysicists) desired greater computational precision than the 24/48-bit PS-2000 formats allowed.

The ability to develop all the PS-2x00 computers was strongly influenced by the availability of know-how, components and other supplies, and development tools. The Western experience with the ILLIAC IV and STARAN computers gave Soviet developers considerably inspiration. These machines pointed out a possible development path, and their existence encouraged developers to think about how machines of this nature could be developed in the Soviet Union. Much of the knowledge critical to building a usable, industrial machine had been built up during more than a decade of real-world computer development at NIIUVM.

The actual implementations were sharply constrained by the component technology available, however. Some design features, associative memory in particular, were judged not feasible given the state of Soviet microelectronics. The decisions to use existing components and production technology had a number of implications. The size of the components determined machine size, causing designers to limit the number of processing elements to 64 in the PS-2000 and the word-length to 24-bits. Similarly, a desire to keep the size of the machine to a reasonable level forced the use of a reconfigurable bus interconnect. Although it was clear that this interconnect was less desirable than, for example, a NEWS interconnect, the technology available for each generation of PS-2x00 computers convinced developers to continue using it. As more functionality was packaged into smaller components, it became feasible to increase the word-length, the number of processors, the amount of memory, etc.

The impact of component availability can be seen clearly when the PS-2100 and PS-2000 are compared. The PS-2100 development time was at least two years greater than that of the PS-2000 in spite of the fact that developers had accumulated a great amount of experience building the PS-2000 and the basic architecture of the PS-2100 base modules was very similar to that of the PS-2000. The principal delay factor was the lack of appropriate components (gate arrays) to construct the machine, together with a lack of experience in customizing them for the PS-2100, and only marginally adequate design tools.

The PS-2x00 themselves and the inherent know-how, architecture, and construction also played an important role. Because the PS-2000 user base was quite extensive, developers felt compelled to preserve the basic architecture and assembly-level instruction set. But beyond this, the basic SIMD-oriented approach was suitable enough for the core applications that developers did not feel the need to explore radically new architectures. Developers prefered to apply existing knowledge of this approach to a new generation of similar machines rather than explore significantly different approaches and face steep new learning curves and protracted development times. The next generation would look similar to the current one in large part because this was the technology that the developers were familiar with.

The PS-2000, -2100, and -2300 mark points on a ‘‘technological trajectory’’ which has been remarkably even. Each generation represented an incremental extension of its predecessor: key parameters were improved, functionality was increased in a cumulative manner, the underlying architecture was retained, and compatibility with previous generations was largely preserved. Until funding ended, the direction of the technological trajectory did not change in any essential manner.

We can postulate the existence of a technological paradigm embodying elements already discussed within which the technological trajectory progressed. Elements of the paradigm include achieving high performance and reliability through parallelism and modularity, the use of a SIMD-based architectural approach with independent memory indexing and activation processors within the processing elements, the incorporation of a linear, segmented-bus interconnect system, construction using standard components and subsystems, and other elements. In other words, the machines already built served as the models or patterns on which further developments were based.

Underlying this consistency were a number of factors which remained largely invarient during the development of these three machines. Table 5-7 shows little variation in user requirements, funding stability, design principles, and strategy. The existing approach met development goals adequately and was reinforced by the need for intergenerational compatibility, and technology was (or became) available to support a strategy of continuity. Thus the ‘‘selection environment’’ which gave developers signals about which developments were possible and beneficial strongly favored the continuation of the existing trajectory.

5.6.2 The Organization

The traditional structure of NIIUVM and the prevalence of similar structures throughout the USSR was discussed in section 5.5.1.1. The basic structure, with the various levels of job specifications, was established by the central authorities. The creation of new laboratories and divisions also had to be approved by higher-level authorities. As a result, while research projects naturally differed greatly from one organization to another, their basic structures were very similar throughout the country.

The specific laboratories and their research domains was determined largely by the nature of the technology being developed and the existing organizational structure. When a new task or development project arose, it was usually assigned to one or more existing laboratories/divisions; if the work was a significant enough departure in nature or scale from existing work, a new laboratory or division could conceivably be created, but only after multiple levels of approval from higher authorities.

During the reform period, the organizational structure of NIIUVM and the problem-oriented Computing Division changed in significant ways, leading to more flexible and autonomous structures. The factors influencing organizational structure most strongly are summarized in Table 5-8. These changes reflected the new opportunities created by changes in legislation, the growing financial crisis for NPO Impul's as a whole, and sometimes conflicting views on which organizational structures served best served the needs of the association, NIIUVM, and its constituent divisions.

Changes in the legal environment of state enterprises and associations had a significant enabling impact on organizational structure. The June, 1987 Law on State Enterprises (Associations) gave the leadership of the NPO the authority to alter the internal structure. Other laws defined the conditions for creating other organizational forms, including small enterprises, temporary scientific-technical collectives, rental collectives and cooperatives. While they did not explicitly mandate structural change, they a) made it possible to create a variety of new, legally recognized organizational forms, and b) gave the organization itself much greater authority to make such changes. As other organiza-

Environment Laws establishing norms for job titles, wage levels Involvement of higher-level officials in approving changes to organizational structures **Legislation allowing alternative organizational forms** **Legislation giving individual institutes the authority to determine their own structure** **Legislation implementing *khozraschet* principles at institute and sub-institute levels** Strong market for NPO Impul’s products **Rapidly declining market for NPO Impul’s computer products, especially HPC**	Technology Functional tasks of R&D programs **Some movement away from large-scale projects to smaller scale projects, contract work**
	Organizational Structure Traditional NPO, institute, division structures oriented towards basic R&D, production tasks. **Weakening link between R&D and production** **Appearance of more flexible and autonomous VNTK, rental collectives, small enterprises**
Technological Availability **Examples of organizational structure and structural change at other organizations**	Beliefs **Key to survival is maintenance of integrated structures** **Key to survival is giving organizational components greater responsibility to fend for themselves** **Key to survival is retaining core personnel**
Organizational Slack Stable funding for R&D at NIIUVM **Decrease in government funding for R&D** **Implementation of *khozraschet* principles**	Strategy **Maintain integrated divisional and institute structures** **Create flexible, autonomous organizational structures** **Seek foreign partners**

Table 5-8 Factors Influencing Organizational Structure within NIIUVM, Problem-oriented Computer Systems Division

tions throughout the country took advantage of the changes in legislation they became examples to decision-makers at NPO Impul’s.

The actual impetus to experiment with new organizational forms came from other quarters. The deepening financial crisis at NPO Impul’s forced management at all levels to search for was of improving the efficiency of work, find alternative funding sources, and increase the ability to apply existing resources in more effective ways. The VNTK, rental collectives, and small enterprises, taking responsibility for their own financial status helped accomplish these goals through their flexible organization, the strong incentives they had for finding contracts, and regulations which permitted greater freedom in setting wage levels, etc.

The creation of these new organizational forms required a shift in beliefs about the most appropriate organizational forms, a shift which has by no means been made willingly by all at NIIUVM. In spite of its inflexibility, the traditional NII structure had served the NPO Impul's quite well over the years, especially in terms of integrating the many development projects carried out through the institute. As the economic crisis grew in magnitude, the primary goal increasingly became survival, and the preservation of the collectives. The strategy of greater decentralization and autonomy of the low-level organizational structures was not adopted easily by high-level management which recognized that the cost of decentralization could very well be a resulting fragmentation of the research of the institute as a whole.

The choice was a difficult one. The nature of the market (or market prospects) for large-scale products which required inter-divisional coordination was not strong enough to alleviate the pressures on the individual organizational structures to find whatever sources of revenue they could. To some, the administrative burden of centralized coordination was an unwelcome overhead, reducing the individual structures' viability. To others, the ability to maintain the integration between NIIUVM divisions was the key to future survival. An important, but unknown variable for NIIUVM was how long it would be before the market for the integrated hardware/software process control and data processing systems recovered. The more appropriate organizational structure to no small degree depends on the answer to this question.

The Problem-oriented Computer Systems Division was very much affected by the forces just described at two levels. As a division, this organization was one of the constitutent parts of NIIUVM and felt the tradeoffs between operating as an autonomous unit finding and executing its own contracts, and remaining a contributing element of the broader NIIUVM research program. The same forces were operating at the sub-division level as well. As the need to find additional revenue sources grew and funding for the PS-2300 grew inadequate, the individual VNTK had to work on whatever projects they could find which would bring in revenue. However, the likelihood of fragmentation to the point were future work on an HPC project is hindered by intra-divisional organizational factors is less than for the institute as a whole. The division (the rental collective) not the constitutent VNTK has its own back account. Although diverging (the software collective continues to develop PS-2100 applications which the architecture collective is working on industrial computers based on the Intel family of microprocessors), the work of each VNTK is closely monitored and coordinated by Itenberg. The core set of PS-2x00 engineers, although working in different VNTK, is small and their experience in working together is great enough that it is likely that will continue to coordinate their efforts and seek projects which will benefit them as a team. Should funding for HPC reappear, it is still possible to reconstitute the core development team.

5.6.3 Prospects

What are the implications of these changes on the future of HPC activities at NIIUVM? The reforms and their consequences have brought about some improvements in the manner in which R&D is conducted, but these have been overshadowed by developments which have made it extreme difficult to continue development of the PS-2x00 line. Without a market and now without funding for HPC development, the prospects for HPC development are, at best, dismal.

Nevertheless, the reforms have brought some improvement in the manner in which R&D is conducted which are likely to be beneficial in the future. The shift away from a centralized planning and supply system and the very real possibility that traditional funding sources would dry up forced people like Itenberg to spend more effort trying to find and cultivate potential supporters and customers. These efforts include rudimentary marketing, more extensive and intensive meetings with those who have applications potentially suitable to PS- style high-performance computers, publicizing at conferences, etc. Similar activities have taken place for years as NIIUVM sought to expand its market, but they have taken on new urgency and developers are becoming each of these remains under the firm administrative control of Itenberg. Only the rental collective (i.e. the division) has its own bank account. much more sensitive to the real needs of actual and potential users. This, plus growing contact with the West has the potential for increasing the rate of idea generation.

Some aspects of ‘‘coalition-building’’–gathering the support necessary to obtain financing, materials, tools, and know-how–has become easier. It is no longer necessary to navigate through long chains of Ministry and government planning officials. Most series produced material inputs and components, either domestic or foreign, are available at a price. Negotiations can be carried out directly between suppliers and customers who now deal on the basis of self-interest rather than centralized directives.

In the short term, however, low organizational slack still make it very difficult to acquire inputs, and the poor state of the economy make finding contracts and other support a very time-consuming process. The geo-political distance from former customers and suppliers in Russia also severely hinder activities. Although NIIUVM does have a pilot production plant which can carry out low-volume production and generate income, it lacks the large-scale production facilities that used to be provided by SPZ and has never had discretionary use of the funds generated through the sale of SPZ products. Without the potential for generating income, NIIUVM will have a difficult time attracting investment and be forced to rely on piece-meal contract work.

Although NIIUVM in principle has creater opportunity to forge ties with the West and acquire foreign inputs, it is still difficult to do so. The information and financial flows between the West and Severodonetsk, located in a remote area of Eastern Ukraine, still move predominantly through middlemen in Russia. Consequently, acquiring information about Western products and companies and making the exchange from the Ukrainian kharbovantsy to the Russian rule to foreign currencies is time consuming, costly, and uncertain.

Some developments have had a positive impact on the product development cycle at NIIUVM. The introduction of temporary, flexible teams has enabled Itenberg to customize teams for the tasks at hand and in principle draw human resources from all parts of NPO Impul's. The concentration of the appropriate human resources is likely to help reduce R&D cycles. In addition, a stronger contract orientation is likely also to shorten development cycles by forcing engineers to meet the terms of contracts, or forfeit payment.

In the short term, however, these improvements will often be overshadowed by the amount of time engineers spend search for contracts or needed supplies, or simply caring for their families.

In sum, although seeds have been laid for successful R&D in the future, the prospects for HPC development are very poor. Successful development in this domain clearly depends on developing an HPC product for which there is a market. The development of HPC systems which are competitive with Western systems now available will require the use of components which, like the 17,000 gate CMOS chips, are currently not available from the domestic industry. To regain competitiveness, Itenberg may have to find specialized niches and increasingly incorporate Western components.