

A Prototype Computer with Non-von Neumann Architecture Based on Strategic Domestic J7 Microprocessor

A. S. Molyakov

Peter the Great St. Petersburg Polytechnic University, St. Petersburg, 195251 Russia

e-mail: andrei_molyakov@mail.ru

Received June 16, 2016

Abstract—We consider a prototype of a computer with non-von Neumann architecture based on the domestic J7 microprocessor and nonclassical massively parallel program organization with globally addressable memory and new data transmission technology, new design of 3D electronics packaging to enhance the security level.

Keywords: virtualization, information security, non-von Neumann architecture, 3D package

DOI: 10.3103/S0146411616080137

INTRODUCTION

Rapidly developing hardware tools are usually a driving force for fundamental changes in systems and applications. Software which develops more slowly seldom creates possibilities for crucial improvements. Software, however, does develop, and its improvement makes possible and the necessary reconsideration of obsolete approaches.

The basic specific feature of modern supercomputers is the possibility of efficient handling of large volumes of accumulated and dynamically varied information, as well as operation with many high intensity input data flows.

Large data volumes usually represent graph databases, operation with them is characterized by high intensity nonregular memory access, which is not easy for available commercial high performance microprocessors.

New approaches, economically efficient and more easily accessible, can be of practical importance to Russia because of the lack of industrial background for the creation of her own component base and specialized supercomputers based on it for the highest quality solution of this problem.

NEW CONCEPT FOR DEVELOPMENT OF SUPERCOMPUTER COMPLEXES

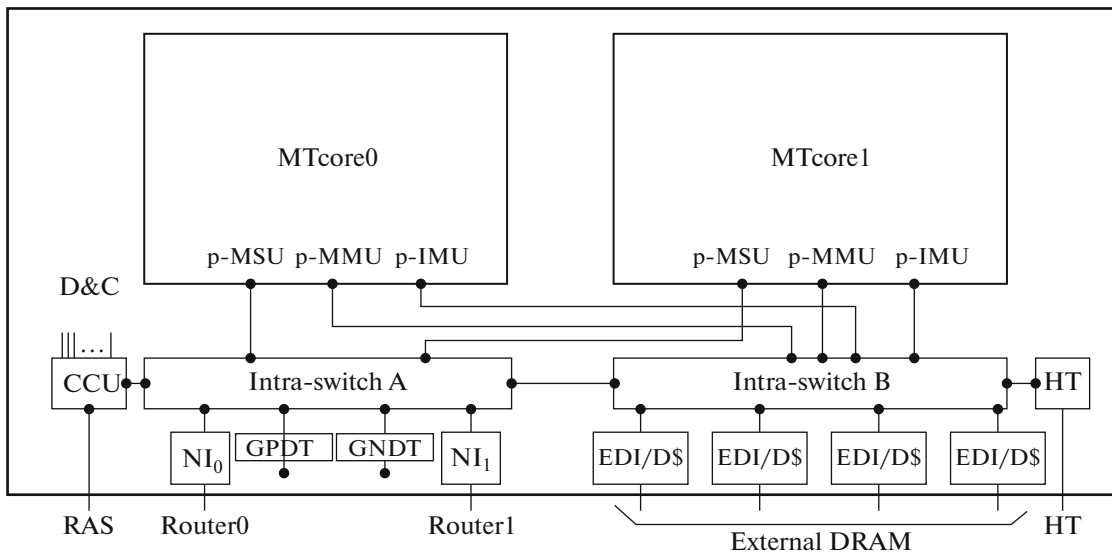
The new concept for development of supercomputer complexes includes the following:

- nonclassical massively parallel software models with global addressable memory and highly asynchronous parallel processes;
- new data transfer technologies, new 3D packaging of electronic units, four privilege levels (“traditional” Unix and Windows NT OS use two protection rings);
- domain data, command, and thread devices protection, and development of own OS and fault-tolerance subsystem, transparent technologies of packaging on chips and boards (domestic component base).

The J7 microprocessor (figure) has two multithread cores (MTcore0 and MTcore1); it was developed taking into account specific features of the 90-nm ASIC technology (Fujitsu, series 101), and is designed for operation at 1 GHz [1, 2].

One multithread core contains four instruction pipelines, each operating with 16 thread devices capable of executing one thread process each.

Each thread device contains control registers for instruction streams of the corresponding thread (status word with instruction counter, exception flags and masks, and other registers), and rather large sets of 64-bit architectural registers for storing fixed point and floating point numbers, addressable 1-bit flag registers, and control transfer address registers. Switching from execution of one thread device instructions



Microarchitecture of the dual-core J7 microprocessor.

to execution of another thread device instructions takes place without register reloading during one processor cycle [3].

Moreover, several thread devices can simultaneously in one cycle issue instructions for execution in functional devices.

The information dependence of executed instructions from one thread device is hardware-monitored via register occupancy flag tables.

Legend for figure:

MTcore0, MTcore1 are the multithread cores, each containing 4 protection domains and 64 thread devices;

p-MSU is the multithread core port from the message I/O unit (MSU);

p-MMU is the port of the multithread core from the virtual address translation unit (memory management unit, MMU);

Intra-switch A is the intrachip switch for transmission of messages concerning network operations;

Intra-switch B is the intrachip switch for transmission of messages concerning execution of memory access operations;

CCU is the central control unit of the microprocessor;

D&C are the diagnostic and control lines;

RAS are the communication lines with the reliability, availability, and serviceability system (RAS system);

NI₀, NI₁ are the internode network adapters connected to the network;

GPDT is the global protection domain table which contains task unicodes for protection domains of multithread cores in which these tasks are solved;

GNDT is the global node descriptor table containing the mapping of virtual numbers of the task nodes onto logical ones;

EDI/D\$ are the off-chip memory interface units including a data cache, a unit for execution of atomic operations and handling tag bits, and memory controllers;

HT is the hypertransport unit (HyperTransport); p-IMU is the multithread core port for the unit for receiving and sending command cache lines.

One pipeline can issue one or two instructions per cycle. Two instructions can be issued if one of them is for global registers and the other one for floating point registers.

One thread of the pipeline can issue one (or two) instructions for execution once each four cycles. Thus, from 8 to 16 instructions can be issued in the J7 microprocessor per cycle. One core of the microprocessor can simultaneously execute several tasks. Each task corresponds to one protection domain, and one of the tasks in the core is necessarily OS. A user task executed in the microprocessor can be simulta-

neously executed in protection domains of different cores, and information on the task assignment to protection domains is stored in a special table of the microprocessor.

A PROTOTYPE COMPUTER WITH NON-VON NEUMANN ARCHITECTURE

The prototype computer with non-von Neumann architecture tested as a part of the “Alfa-monitor” software–hardware complex represents a “China pie,” i.e., has a multilayer structure (layers of memory modules, layers of I/O subsystem directly on the chip, layers of processor elements).

The modular packaging is performed using “dense packing” without contact wires and additional shunting lines. The main data exchange lines are assembled using the new technology implying “high temperature superconductivity,” namely, links with nanocoating manufactured from carbon nanotubes, carbon oxide structure, metal ligands.

Double-indexing register occupancy tables (along with occupancy counters) were created for imparting a pronounced asynchronism to the system and parallelization of operations at each thread device (multiple branching graphs, sophisticated case algorithms). In this technology, there are no delays for accessing register data. The chip modules have a multilayered structure (layers of memory modules, layers of I/O subsystem, layers of processor elements).

The concept of the “non-von Neumann architecture” is based on the following software architectural and structural technological principles:

MT – multithreading, multithread arrangement of the processor and executed programs;

DF – dataflow, dataflow control of computing using models with ordered data flows and static graphs (stream-based computing model) and unordered data flows and dynamic graphs (dataflow model). This affects both the processor and communication network arrangement, and that of the software;

DAE – decoupled access/execute model of the processor and programs separates the following parts: the computing part (heavy computing) with a large number of computing operations involving regular memory access with good space–time localization, and the non-computing part of memory handling with a large number of memory access instructions and address calculations providing data for the computing part;

PIM – processor in memory, the technologies of development of intelligent memory modules, i.e., memory modules with built-in processors;

CMT – chip multiprocessor technologies of development of very large scale integration circuits of “system on a chip” type.

Characteristics:

—12 multithread cores, each containing 256 thread devices and providing operation of 16 protection domains;

—each thread device can execute program threads differing with respect to the type of calculations and the distributed computing model (static and dynamic graphs of data flows, threading models);

—register memory of general purpose registers and registers for floating point data storage/processing is realized as large size two-level cached register files, which provides high efficiency (considerable gain) in an area at the chip assembly and optimal energy consumption;

—general purpose registers and floating point registers can be used both for temporary data storage and as I/O ports for superfast interaction and minimum read/write delays, moreover, additional CTX/CP modules for super-high performance memory and network handling are used;

—on the microarchitecture level the CT-2 processor has 300 built-in commands, it processes 32-bit and 64-bit signed and unsigned floating point numbers;

—the number of nodes on one chip is 63648;

—floating point units (FPU) and fixed point units (FXU) can operate in asynchronous and synchronous manners. Operating synchronously, they can synchronously execute SIMD operations over short vectors: sixteen operations with 32-bit numbers, eight operations with 64-bit numbers. Fast processing of “long vectors” is also implemented.

Feature 1.

The multithread core contains 256 threads. Along with usual functional devices, there exists a special unit for operations with 64×64 -bit matrices. There also exists a device for operations with 160-bit data. One hundred and sixty bit is the width of a special arithmetic–logic device that can execute operations over operands representing a composition of several operands. For example, 160-bit length can be represented as $128 + 32$ or $80 + 80$.

In the first case, this device can be used as a storage adder for 32-bit operands; and in the second case, the 3D pin-free packaging of alternating processor and memory chips is used for operations over “extended double precision” floating point numbers or group operations over 80 1-bit flag registers available in the thread device.

The difference from classical microprocessors is as follows:

- the unit of batch queries to other modules of the system containing physical memory addressable via a unified globally addressable virtual address space;

- direct interconnection of the microprocessor chips on three levels can be used for increasing the processor core number and increasing fault tolerance (resource backup upon the program execution and synchronization of results);

- atomic execution of integer arithmetic and logical operations directly over a memory cell, operation with access and status tag bits;

- direct network query transmission channel for queries to a physical address in the region not cached by the microprocessor in L2 cache. In this case, first, it is not necessary to additionally translate the virtual address in the microprocessor MMU using the second level of virtual addressing (R-segments and R-pages), and second, it is not necessary to conduct a search in cache memories. Thus, it is not necessary to transmit a message into the microprocessor via the NI unit, which means greatly increased efficiency;

- connection to a terabit network integrated in a communication line substrate using nanotechnologies and “high temperature superconductivity”;

- interface for packet transmission to intermodular network, superfast packet transmission with compression, which allows packet serialization/deserialization.

Feature 2.

The size of addressable data is increased to 32 PByte. The command address is extended to 48 bit. Thus, commands taken from another supercomputer node can be executed.

This solution considerably reduces the constraint on the size of executed programs and can be used for increasing reliability.

Feature 3.

The average number of r-type registers (registers for storing binary codes, data addresses, and fixed point numbers) and f-type registers (registers for storing floating point numbers) in the thread device is 32. The registers are 64-bit. The heterogeneous thread computing model is supported; the basic principles of the model are as follows:

- threading calculations executed at thread devices can be “heavy,” “intermediate,” and “light,” which results in the difference in the number of available register resources, addressable memory, complexity of the allowed state diagram of the calculation, service procedures of the thread device while selecting commands for execution in functional devices;

- “light” threads can operate with 16 r-registers and 16 f-registers, “intermediate” threads, with 32 registers of each type, and “heavy” threads, with 64 registers of each type.

Feature 4.

The microprocessor uses two-level register memory. There are many thread devices in the core, namely, 256, and each such device can contain, on average, 32 r- and f-registers.

Thus, very large register files are required, which is difficult to implement. The problem is solved by introducing two levels.

The first level of register memory is the real multi-input register memory, and the second level is the common static memory to which all register memory of the core is mapped. Register memory is divided into 8-word pages.

The exchange between the first and second levels is executed via register pages similar to an ordinary memory with page organization. Thus, register addressing is in reality two-level, the number of the register page and the 3-bit register number in the page.

Feature 5.

There exist 1-bit q-registers for storing flags of logical and arithmetic operations. These registers are used in logical calculations and for conditional command execution, which allows one to eliminate some conditional control transfer commands. The application of q-registers increases register resource and increases the capabilities of combined pipelined command execution.

The prototype computer contains 80 q-registers. They are 64 global and 16 local q-registers. This considerably increases the efficiency of checking complex logical expressions, which is typical for real-time systems.

Feature 6.

Registers can serve for temporary storage of values and as ports for data transmission between threads, which allows one to avoid using RAM. With a large number of threads in the cores and such a feature, it is possible to implement new computing models based on static graphs of data flows.

CONCLUSIONS

Important results were obtained in the last 2–3 years in the field of superconductor electronics and specialized quantum computers based on cryogenic superconducting electronics operating as analogue machines. These technologies were applied in prototypes and can be used for the development of specialized accelerators in the framework of exaflop supercomputers in the near future, which would allow one to increase the efficiency of solving engineering and numerical problems by 3–6 orders of magnitude and go beyond the physical performance constraints of common supercomputers (Landauer limit) of several tens of exaflops.

The promising post-Moore computer elements and new principles of supercomputer construction allow one to anticipate the creation of specialized military zettaflop supercomputers approximately in 2020, and yotaflop supercomputers, in 2024, with expected energy consumption of the latter on a level of 15 MW.

The complete solution of the problem of protection implies the development of supercomputers with hardware-based protection of programs and data. At present, four protection levels are implemented (user, executor OS, core OS, and initialization). A computer based on post-Moore elements and new construction principles of non-von Neumann architecture may possess a substantially increased fault-tolerance and security level.

The new concept of development of protected supercomputer complexes means the transition to non-classical massively parallel program models with globally addressable memory and highly asynchronous parallel processes, application of new data transmission technologies, and new designs of 3D electronic unit packaging.

REFERENCES

1. Slutskin, A.I. and Eisymont, L.K., The Russian supercomputer with globally addressable memory, *Otkrytye Sist.*, 2007, no. 9, pp. 42–51.
2. Mitrofanov, V.V. and Eisymont, L.K., The element base and architecture of high-performance multi-processor computing systems, perspective strategic and embedded supercomputers, in *Sb. Dinamika radioelektroniki* (Dynamics of Radioelectronics), Borisov, Yu.I., Ed., Moscow: Tekhnosfera, 2008, 2nd ed., pp. 70–76.
3. Zhirnov, V., et al., Limits to binary logic switch scaling, *A Gedanken Model – Proceedings of the IEEE*, 2003, vol. 91, no. 11, pp. 1934–1939.

Translated by E. Baldina