МЦСТ ЭЛЬБРУС

Russian Microprocessors of the Elbrus Architecture Series for Servers and Supercomputers

Alexander Kim, Ignat Bychkov, Vladimir Volkonskiy, Feodor Gruzdov, Sergey Semenikhin, Vladimir Tikhorsky, Vladimir Feldman

JSC “MCST”

Russian Supercomputing Days, September 2015

Products and Technologies

- Microprocessors (CPU)
  “Elbrus” and “MCST-R” Lines
- Controllers (south bridge)
- Computers, computer modules
- System Software
  - –Operating System
  - –Software Development Kit
    - 3-way Parallelism Support by Compiler
  - –Binary Compatibility Technology
  - –Secure Program Execution Technology

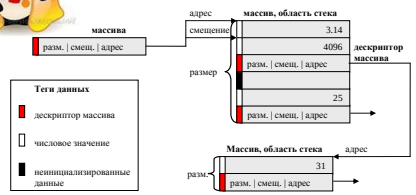

МЦСТ ЭЛЬБРУС

Outline

- Elbrus technologies
- Elbrus products
- Elbrus future

МЦСТ ЭЛЬБРУС

Deep Hardware & Software Integration in the Elbrus Architecture

- *HW architecture provides*
  - Parallel resources by wide instruction (VLIW-like)
    - Up to 25 scalar operations per cycle per core
    - Up to 12 Flops DP (24 packed SP) per cycle per core
      - Doubling in the Elbrus-8CV
    - Multicore
    - Multiprocessor support (ccNUMA)
  - Large-scale register file
  - Optimization supporting features
  - Binary compatibility supporting features
  - Secure program execution supporting features
- *Compilers and OS provides*
  - Program parallelization by optimizing compiler
    - Instruction level parallelism (many operations per cycle)
    - Packed (vector) operation parallelism
    - multicore, multithreading parallelism
  - Viable binary compatibility with Intel x86, x86-64 on the basis of transparent dynamic binary compilation technology
  - Programming languages implementation for secure program execution

General purpose MP architecture – wide range of use
High MP resource utilization
Power efficient on HPC applications

МЦСТ ЭЛЬБРУС

Key Elbrus CPU Technologies

VLIW-like architecture (25+ ops per cycle)

- Supported by optimizing compiler

Binary compatibility with Intel x-86, x86-64 via BT

- Functionality
  - Direct execution of 20+ operating systems, including: MSDOS, Windows XP, 7, Linux, QNX, PS/2
  - Direct execution of 1000+ popular applications
  - Execution of applications under operating system Elbrus (Linux Distributive)
- Performance – up to 80% from native
  - By transparent optimizing binary translation system
  - Based on strong and powerful hardware support
- Independent from Intel license

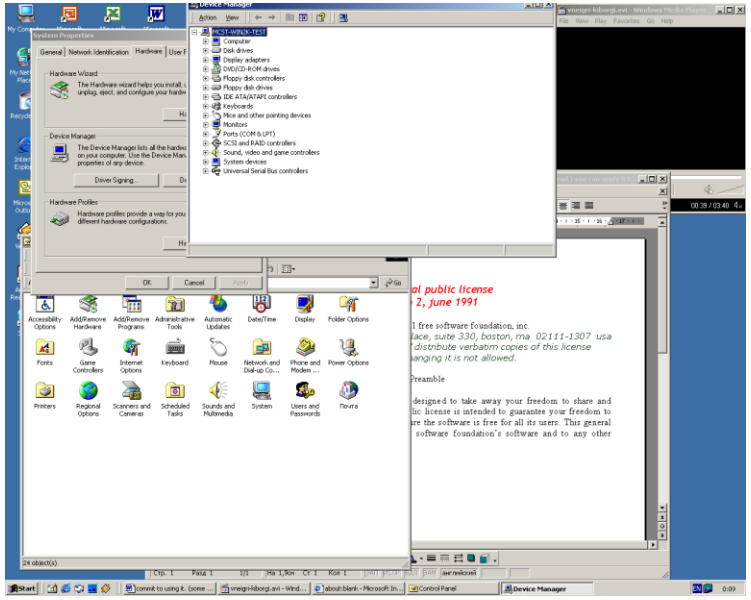

Secure program execution technology

- Memory and data protection
  - Structured memory
  - Object access by descriptors
  - Scopes access supported
- Critical vulnerabilities detection
  - Buffer overflow
  - Uninitialized data access
  - Dangling pointer access

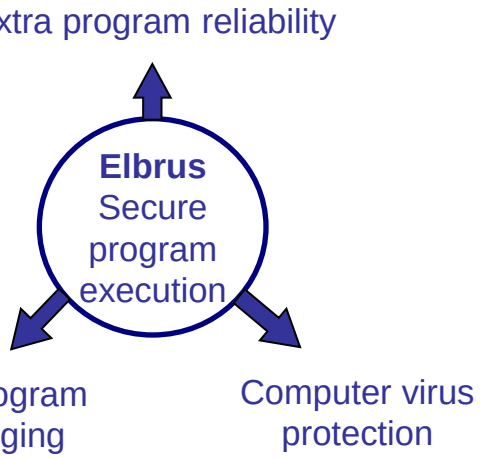

x86, x86-64 compatibility on the basis of binary compilation system

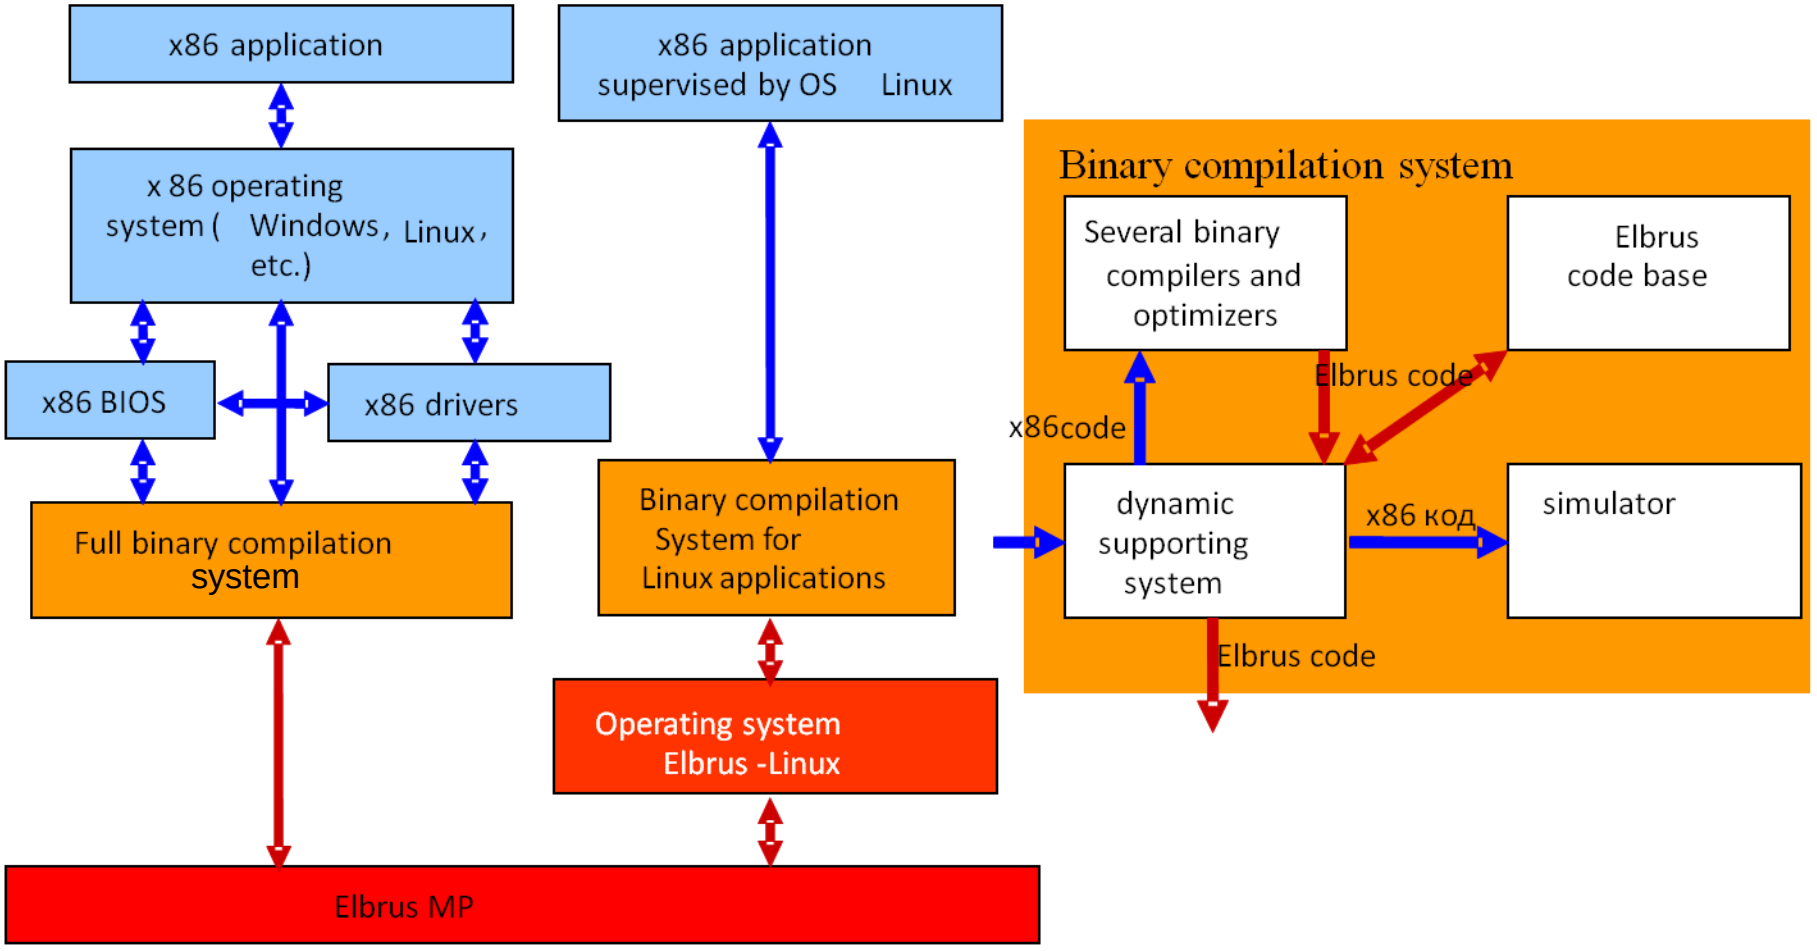

Viable binary compatibility on the basis of ILP and special HW support

МЦСТ ЭЛЬБРУС

Binary compatibility system details

- Several optimization levels
  - Simplest, template based, fast compilation, poor code
  - intermediate, region based, rather fast, viable code
  - Highest, utilizes all parallelism of the Elbrus MP architecture
    - Efficient execution of multithreading applications and OSs
    - Efficient implementation of precise and asynchronous exceptions and interrupts
- Free cores are used for parallel dynamic compilation and optimization
- Well optimized regions saved in special Elbrus code base
  - Used in repeating execution
- Feedback control for performance tuning
  - Regions with negative impact on performance recompiled

МЦСТ ЭЛЬБРУС

Secure execution technology details

Security in the Elbrus

- All pointers protected by tags
  - Impossible to construct or to fake pointer
- Objet bounds are controlled by descriptor
- High level language scopes supported
  - Access to visible data in scope and through pointers passed from other modules (methods)

Antivirus protection

High program reliability

Traditional architectures

- Arithmetic data and pointers undistinguished
  - Pointer is a number
- Objects allocated in plain memory, object bounds aren't checked
- HW don’t understand scopes
  - Reliable module can be destroyed

No antivirus HW support

Low program reliability

МЦСТ ЭЛЬБРУС

Outline

- Elbrus technologies
- Elbrus products
- Elbrus future

MP Elbrus-2C+, Elbrus-4C, south bridge KPI-1

Elbrus-2C+:

- TSMC 90 nm process, 10 metal layers
- 0.5 GHz clock frequency
- Power - 25 W
- Chip structure
  2 Elbrus architecture cores,
  4 DSP Multicore architecture cores
- Total performance -
  28/8 Gflops sp/dp:
  2 Elbrus cores – 16/8 Gflops sp/dp,
  4 DSP cores – 12 Gflops sp
- Die size - 17,2x16,8 mm
- Sampling 2011
- Production 2012H1

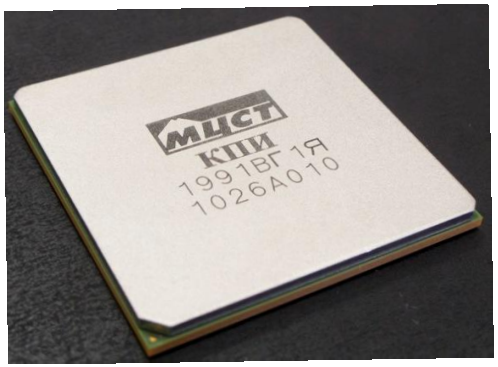

South bridge KPI-1:

- TSMC130 nm process, 9 metal layers
- 250 MHz clock frequency
- Power – 5 W
- 14 interfaces provides:
  system, PCI Express, PCI, Ethernet (10/100/1000), SATA 2.0, USB 2.0, RS 232/485, etc.
- Die size – 10,6x10,6 mm
- Sampling - 2010
- Production – 2011H1

Elbrus-4C:

- 4 Elbrus architecture cores
- 8 MB L2 cache (2 MB per core)
- TSMC 65 nm process
- Die size 380 mm^2
- Тактовая частота 0.8 GHz
- Power – 45 W
- Performance 50/25 Gflops sp/dp
- Memory throughput 38,4 GB/sec (3 DDR3 channels)
- 3x16 GB/sec inter CPU channels for 4 CPU ccNUMA 16 GB/sec
- 2 IO links
- Sampling 2013
- Production 2014H1

Personal Computers with Elbrus CPUs

Monoblock

- Display 21” 1920*1080
- Video card 2D/3D*
- HD: SATA 3.5” + DVD
- Interfaces (USB 2.0, WiFi, Bluetooth, DVI, Gigabit Ethernet, camera, microphone)
- Size 535x415x55(mm)

Compact computer

- CPU Elbrus-2C+

Desktop

- CPU Elbrus-4C
  - 4 cores
  - L2 cache – 2 MB per core
  - 800 MHz
  - 45 W
- 2D/3D video card
- interfaces
  - PCI Express 1.0 8 lines
  - Gigabit Ethernet
  - SATA 2.0

Servers and Clusters with Elbrus CPUs

Server Elbrus-4.4 (based on Elbrus-4C CPU)

- 4 CPUs Elbrus-4C (4 cores, 800 МГц), total of 200 GFLOPs , 2 southbridge controllers
- RAM: 96 GB, 12x DIMM DDR3-1600
- Interfaces: SATA 2.0 – 8 channels, Gigabit Ethernet – 2 channels, PCI Express 1.0 x8 – 2 slots, PCI – 2 slots, USB – 6 slots
- Case height: 2U,1U

Cluster based on Elbrus-4C CPUs

- Cabinet 47U – 1;
- 4-processor servers – up to 64
- CPUs – up to 256 (1024 cores)
- RAM – 6-12 TB
- HD – 32-64 TB
- FPGA-based interconnect (design by MCST)
- Air Cooling system
- Power – up to 20 KW
- Peak performance – up to 13,8 TFLOPs

МЦСТ ЭЛЬБРУС

Operating System Elbrus

- Elbrus OS kernel based on OS Linux kernel
  - Real time mode support
  - Elbrus technologies support
    - Binary compatibility for Linux applications in Intel x86 codes
    - Efficient secure execution of programs
- Software development kit
  - Optimizing compilers (C, C++, Fortran, Java), linker, debugger, profiler, math libraries
    - Program parallelization
      - MPI, OpenMP, automatic parallelization for ILP, vectorization, multithreading
    - Performance libraries
  - Open source software stack
    - Compatibility with GCC features

- Operating system user package
  - Utilities, services, general purpose libraries
  - Graphics subsystem, network, databases, office package
  - Cluster resource management
    - slurm, irqbalance, torque, ganglia, nfs-server, iscsi-target
  - Drivers from open-source Linux world

МЦСТ ЭЛЬБРУС

Outline

- Elbrus technologies
- Elbrus products
- Elbrus future

Next generation CPUs and controllers

Elbrus-8C

- 8 Elbrus cores
  - 30+ ops per cycle
- 1,3 GHz clock frequency
- Peak performance 125/250 Gflops dp/sp
- TSMC 28 nm process
- Die area 321,4 mm^2
- L2 Cache – 512 KB per core
- L3 Cache – 16 MB, shared
- sampling – 2015Q4
- production – 2016H1

South bridge KPI-2

- TSMC 65 nm process
- CPU channel - 16 GB/sec
- interfaces
  PCI Express 8+8+4 lines,
  Gigabit Ethernet – 3 ports,
  SATA 3.0 – 8 ports,
  USB 2.0 – 8 ports
- SPMC controller
- interrupt controller
- sampling – 2015Q4
- production – 2016H1

Elbrus-8CV

- 8 Elbrus cores
  - 50+ ops per cycle
- 1,5 GHz clock frequency
- Peak performance 512+/256+ Gflops sp/dp
- Die area 435 mm^2
- L2 Cache – 512 KB per core
- L3 Cache – 16 MB, shared
- TSMC 28 nm process

Sampling – 2018Q2

Production – 2018Q4

Performance increase of the Elbrus MP series

Elbrus-2C+
0.5 GHz, 2+4 C
2*DDR2-800
16+12 Gflops sp
25 W
90 nm

2 years

3x

Elbrus-4C
0.8 GHz, 4 C
3*DDR3-1600
50-60 Gflops sp
45…60 W
65 nm

2 years

4-5x

Elbrus-8C
1.3 GHz, 8 C
4*DDR3-1600
250 Gflops sp
~60…90 W
28 nm

2 years

2x+

Elbrus-8CV
1.5 GHz, 8 C
4*DDR4-2400
512+ Gflops sp
~60…90 W
28 nm

2011

2013

2015

2018

We are developing next generations of MP, computers, and system software

МЦСТ
ЭЛЬБРУС

Thank you