GS464V: A High-Performance Low-Power XPU with 512-Bit Vector Extension

Weiwu Hu, Yunji Chen
Institute of Computing Technology, Chinese Academy of Sciences
Loongson Technologies Corporation Limited
hww@ict.ac.cn, cyj@ict.ac.cn

Contents

- **Background**
- **The XPU architecture of GS464V**
- **Godson chips based on GS464V**

Godson is the academic name of Loongson

TOP10 HPCs in 2008.11

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz , Voltaire Infiniband / 2008 IBM	129600	1105.00	1456.70	2483.47
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT5 QC 2.3 GHz / 2008 Cray Inc.	150152	1059.00	1381.40	6950.60
3	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz / 2008 SGI	51200	487.01	608.83	2090.00
4	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60
5	Argonne National Laboratory United States	Blue Gene/P Solution / 2007 IBM	163840	450.30	557.06	1260.00
6	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband / 2008 Sun Microsystems	62976	433.20	579.38	2000.00
7	NERSC/LBNL United States	Franklin - Cray XT4 QuadCore 2.3 GHz / 2008 Cray Inc.	38642	266.30	355.51	1150.00
8	Oak Ridge National Laboratory United States	Jaguar - Cray XT4 QuadCore 2.1 GHz / 2008 Cray Inc.	30976	205.00	260.20	1580.71
9	NNSA/Sandia National Laboratories United States	Red Storm - Sandia/ Cray Red Storm, XT3/4, 2.4/2.2 GHz dual/quad core / 2008 Cray Inc.	38208	204.20	284.00	2506.00
10	Shanghai Supercomputer Center China	Dawning 5000A - Dawning 5000A, QC Opteron 1.9 Ghz, Infiniband, Windows HPC 2008 / 2008 Dawning	30720	180.60	233.47	

Top10 HPCs in 2009.11

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron Six Core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.60
2	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2009 IBM	122400	1042.00	1375.78	2345.50
3	National Institute for Computational Sciences/University of Tennessee United States	Kraken XT5 - Cray XT5-HE Opteron Six Core 2.6 GHz / 2009 Cray Inc.	98928	831.70	1028.85	
4	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution / 2009 IBM	294912	825.50	1002.70	2268.00
5	National SuperComputer Center in Tianjin/NUDT China	Tianhe-1 - NUDT TH-1 Cluster, Xeon E5540/E5450, ATI Radeon HD 4870 2, Infiniband / 2009 NUDT	71680	563.10	1206.19	
6	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0 GHz/Nehalem EP 2.93 Ghz / 2009 SGI	56320	544.30	673.26	2348.00
7	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60
8	Argonne National Laboratory United States	Blue Gene/P Solution / 2007 IBM	163840	458.61	557.06	1260.00
9	Texas Advanced Computing Center/Univ. of Texas United States	Ranger - SunBlade x6420, Opteron QC 2.3 Ghz, Infiniband / 2008 Sun Microsystems	62976	433.20	579.38	2000.00
10	Sandia National Laboratories / National Renewable Energy Laboratory United States	Red Sky - Sun Blade x6275, Xeon X55xx 2.93 Ghz, Infiniband / 2009 Sun Microsystems	41616	423.90	487.74	

Top10 HPCs in 2010.6

Rank	Site	Computer/Year Vendor	Cores	R_{max}	R_{peak}	Power
1	Oak Ridge National Laboratory United States	Jaguar - Cray XT5-HE Opteron Six Core 2.6 GHz / 2009 Cray Inc.	224162	1759.00	2331.00	6950.60
2	National Supercomputing Centre in Shenzhen (NSCS) China	Nebulae - Dawning TC3600 Blade, Intel X5650, NVidia Tesla C2050 GPU / 2010 Dawning	120640	1271.00	2984.30	
3	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2009 IBM	122400	1042.00	1375.78	2345.50
4	National Institute for Computational Sciences/University of Tennessee United States	Kraken XT5 - Cray XT5-HE Opteron Six Core 2.6 GHz / 2009 Cray Inc.	98928	831.70	1028.85	
5	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution / 2009 IBM	294912	825.50	1002.70	2268.00
6	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX/8400EX, Xeon HT QC 3.0/Xeon Westmere 2.93 Ghz, Infiniband / 2010 SGI	81920	772.70	973.29	3096.00
7	National SuperComputer Center in Tianjin/NUDT China	Tianhe-1 - NUDT TH-1 Cluster, Xeon E5540/E5450, ATI Radeon HD 4870 2, Infiniband / 2009 NUDT	71680	563.10	1206.19	
8	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60
9	Argonne National Laboratory United States	Intrepid - Blue Gene/P Solution / 2007 IBM	163840	458.61	557.06	1260.00
10	Sandia National Laboratories / National Renewable Energy Laboratory United States	Red Sky - Sun Blade x6275, Xeon X55xx 2.93 Ghz, Infiniband / 2010 Sun Microsystems	42440	433.50	497.40	

What Next ?

No. 1 ?

Build HPCs with domestic designed CPU !

CPU Plan of China

- **Design CPU for HPC is part of the National S&T Major Project for CPU**
  - **the National Mid-Term S&T Plan from 2006-2020**
- **16 Major Projects, each fund USD 5-10B from 2006-2020**
  - **CPU and OS**
  - **VLSI process technology**
  - **Next-generation (4G) wireless network**
  - **High-end digital machine tool**
  - **Advanced nuclear fission power plant**
  - **Water pollution control and treatment**
  - **Large aircraft**
  - **High-resolution earth-observation system**
  - **Manned space flight and lunar exploration**
  - **.....**

CPU Development Plan

- **10th Five Year Plan (2001-2005):**
  - **Startup and key technology research**
  - **Four-Issue OOO Architecture, 1.0GHz**
- **11th Five Year Plan (2006-2010):**
  - **From emulation to innovation, low- to high-end, research to product**
  - **Multi-core CPU with leading performance, CPU company setup**
  - **Desktop, servers, and HPC products based on domestic designed CPU**
- **12th and 13th Five Year Plan (2011-2020):**
  - **Build a new ecosystem to support the IT industry in China**
  - **Start from National Security, education, e-government, ...**

Godson CPU Briefs

- **Research Stage: started in 2001.**
  - **The 32-bit Godson-1 in 2002 is the first CPU in China**
  - **The 64-bit Godson-2B in 2003.10**
  - **The 64-bit Godson-2C in 2004.12**
  - **The 64-bit Godson-2E in 2006.03**
  - **Each Triple the performance of its previous one**
  - **SPEC int2000 and SPEC fp2000 of Godson-2E> 500**
- **Product Stage: started in 2008**
  - **Low end: SOCs for low-cost PC and consumer electronics**
  - **High end: multi-core CPUs for server and HPCs**

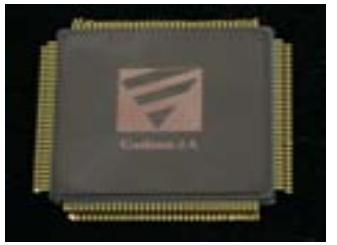

High end Roadmap: More cores on a chip

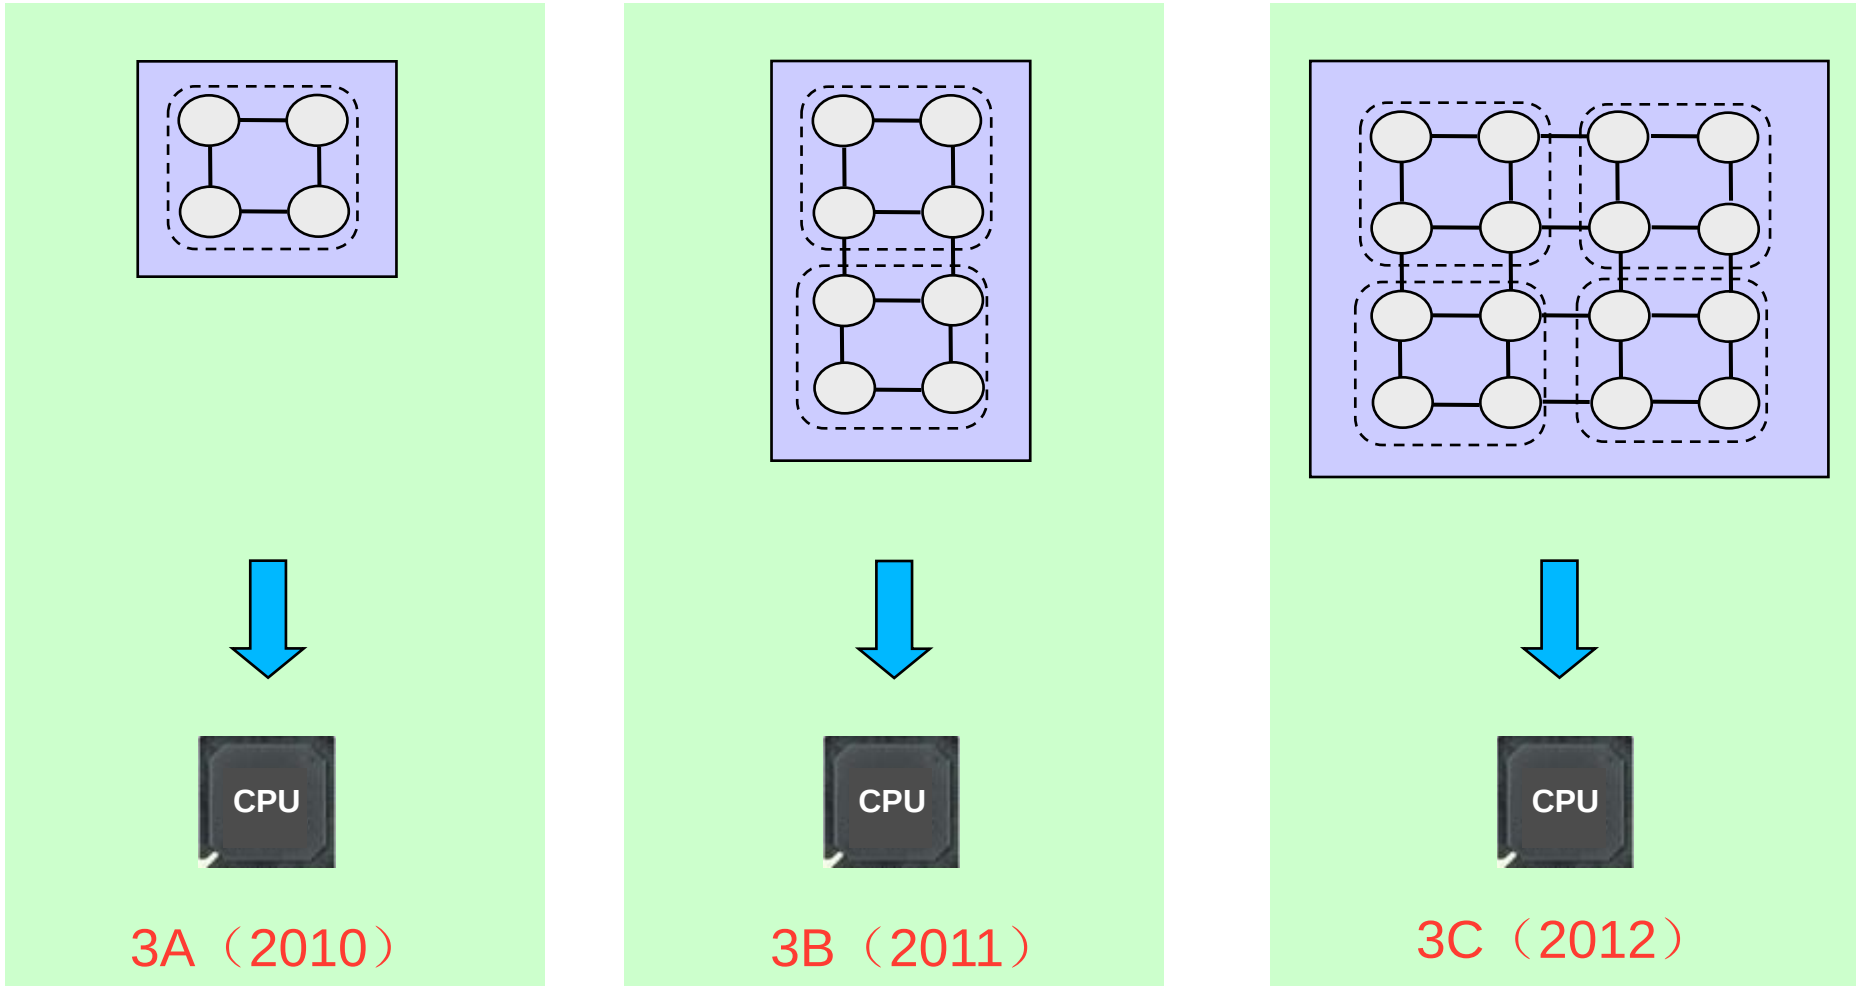

Godson-3A

- 1.0GHz@65nm CMOS, 10W
- 425M transistors, area 174.5mm^2
- Four GS464 cores
  - 64-bit MIPS64 Compatible
  - HW support for X86 binary translation
  - Four-issue, OOO
  - 64KB+64KB L1 (four-way)
- 4MB L2
- Two on-chip DDR2/3 controller.
- Two 16-bit HT
- PCI/PCIX, LPC, GPIO, etc.
- SPEC int2000 rate and SPEC fp2000 rate 25
- The architecture presented in HOTCHIPS2008

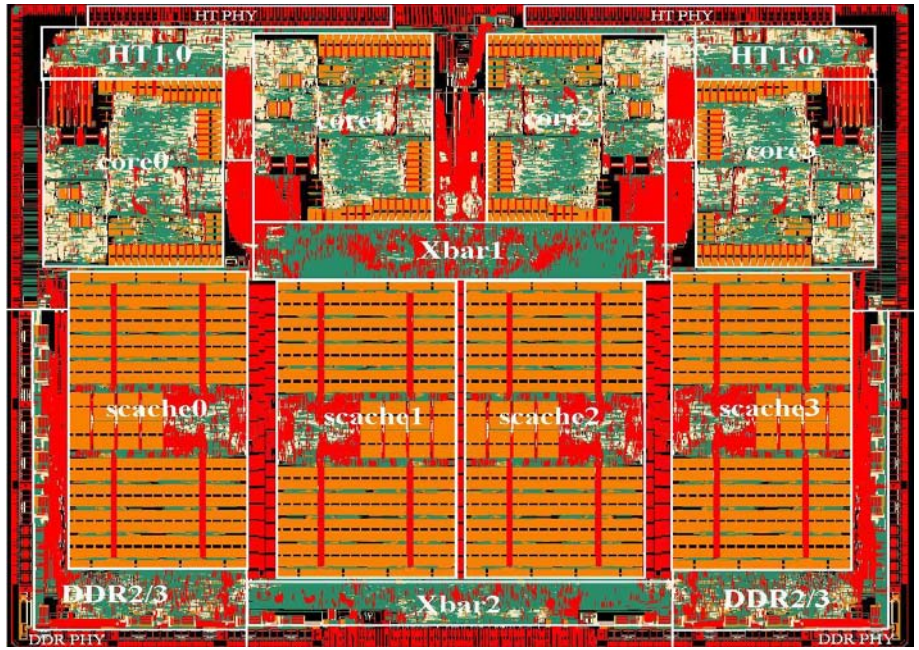

Low end Roadmap: From CPU to SOC

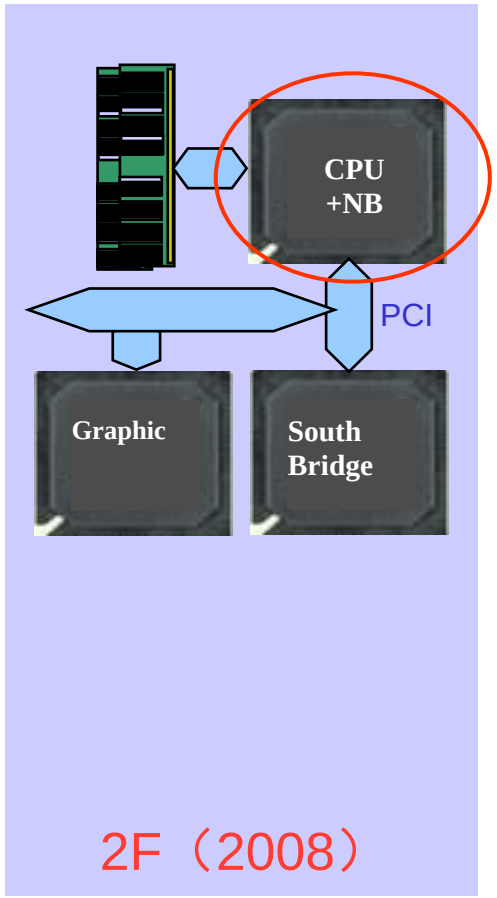

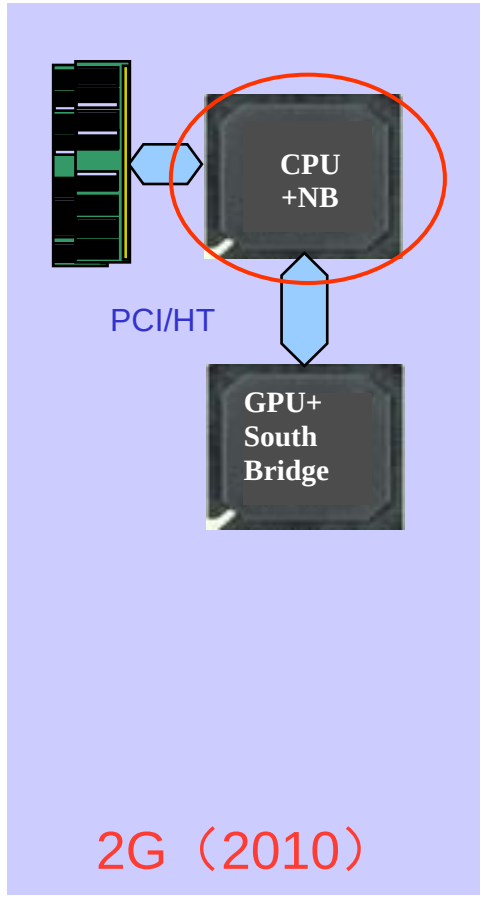

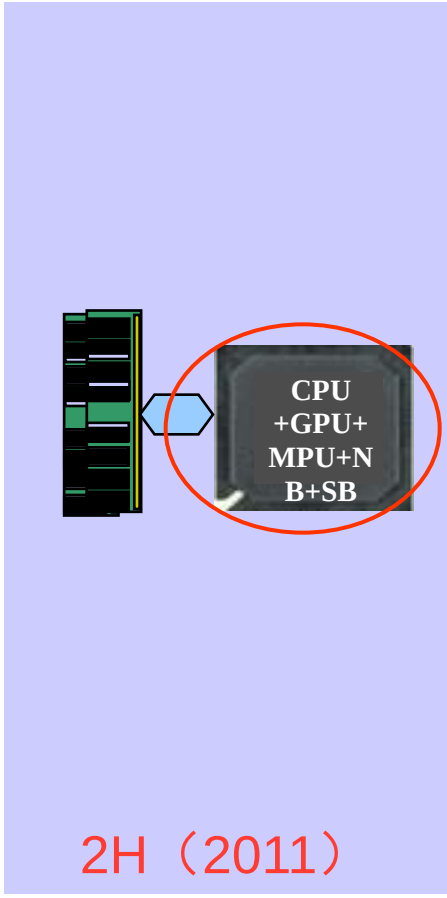

Godson-2G

- 1.0GHz@65nm CMOS, 3W
- 100M transistors, area 60mm^2
- Single GS464 cores
  - ◆64-bit MIPS64 Compatible
  - ◆HW support X86 binary translation
  - ◆Four-issue, OOO
  - ◆64KB+64KB L1 (four-way)
- 1MB L2
- On-chip DDR2/3 controller.
- 16-bit HT
- PCI/PCIX, LPC, GPIO, etc.

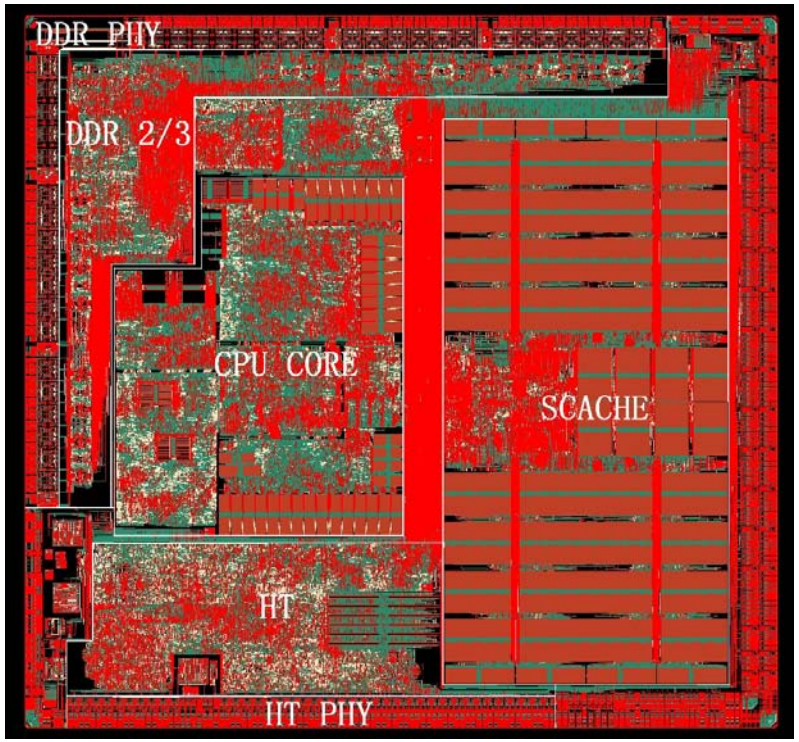

GS464: A “typical” 4-issue 64-bit CPU Core

- **MIPS64 compatible, 200+ instructions for X86 emulation**
- **Four-issue 64-bit superscalar OOO pipeline**
- **Two fix, two FP, one memory units**
- **64KB icache and 64KB dcache, 4-way**
- **64-entry TLB, 16-entry ITLB**
- **Directory-based cache-coherence**
- **Parity check for icache, ECC for dcache**
- **EJTAG for debugging**
- **Adopted by Godson-2G and Godson-3A**

GS464 Architecture

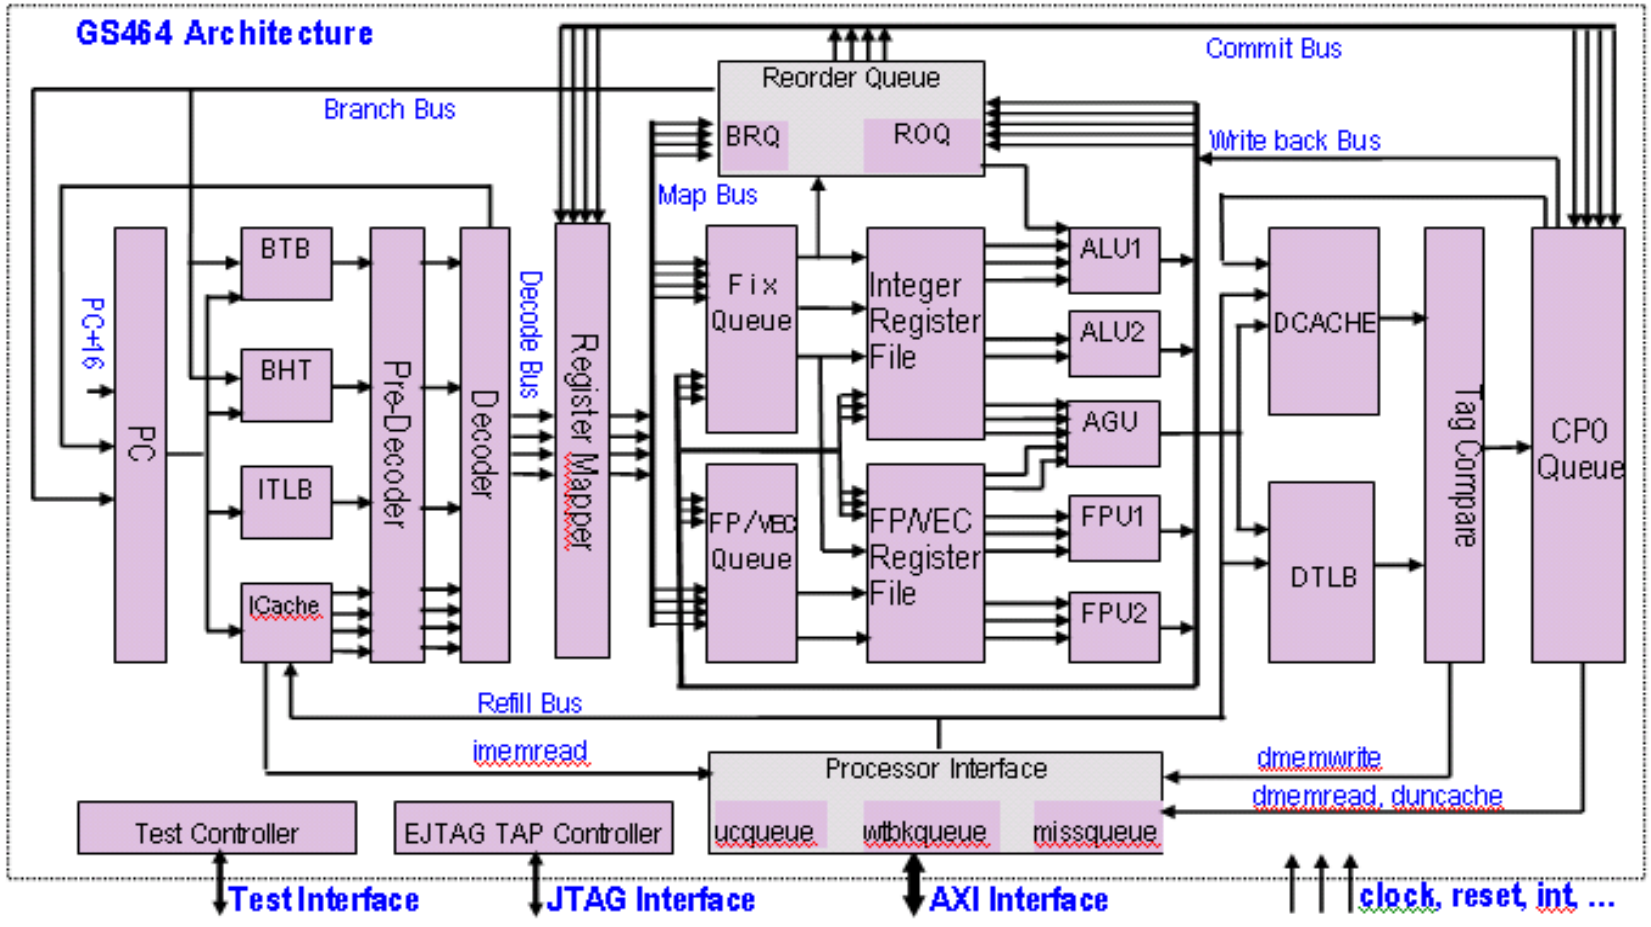

AGU: Address Generation Unit	EJTAG: Enhanced JTAG
ALU: Arithmetic Logic Unit	FPU: Floating-Point Unit
BTB: Branch Target Buffer	ITLB: Instruction Translation Look-aside Buffer
BHT: Branch Target Buffer	ROQ: ReOrder Queue
BRQ: Branch Reorder Queue	TAP: Test Access Port
DTLB: Data Translation Look-aside Buffer	UCQ: Uncache Access Queue

Contents

- **Background**
- **The XPU architecture of GS464V**
- **Godson chips based on GS464V**

XPU Idea

- **Obstacles of Moore's Law**
  - **Memory Wall: memory hierarchy**
  - **Power consumption Wall: multi-core architecture**
  - **Bandwidth Wall: XPU architecture**
- **Why XPU?**
  - **CPU: separate data and computation, flexible & low efficiency**
  - **GPU: tightly coupled data and computation: not flexible & high efficiency**
  - **XPU: Reconfigurable data for computation, flexible & high efficiency**
- **XPU=CPU+DSP+MPU+...**

XPU Idea

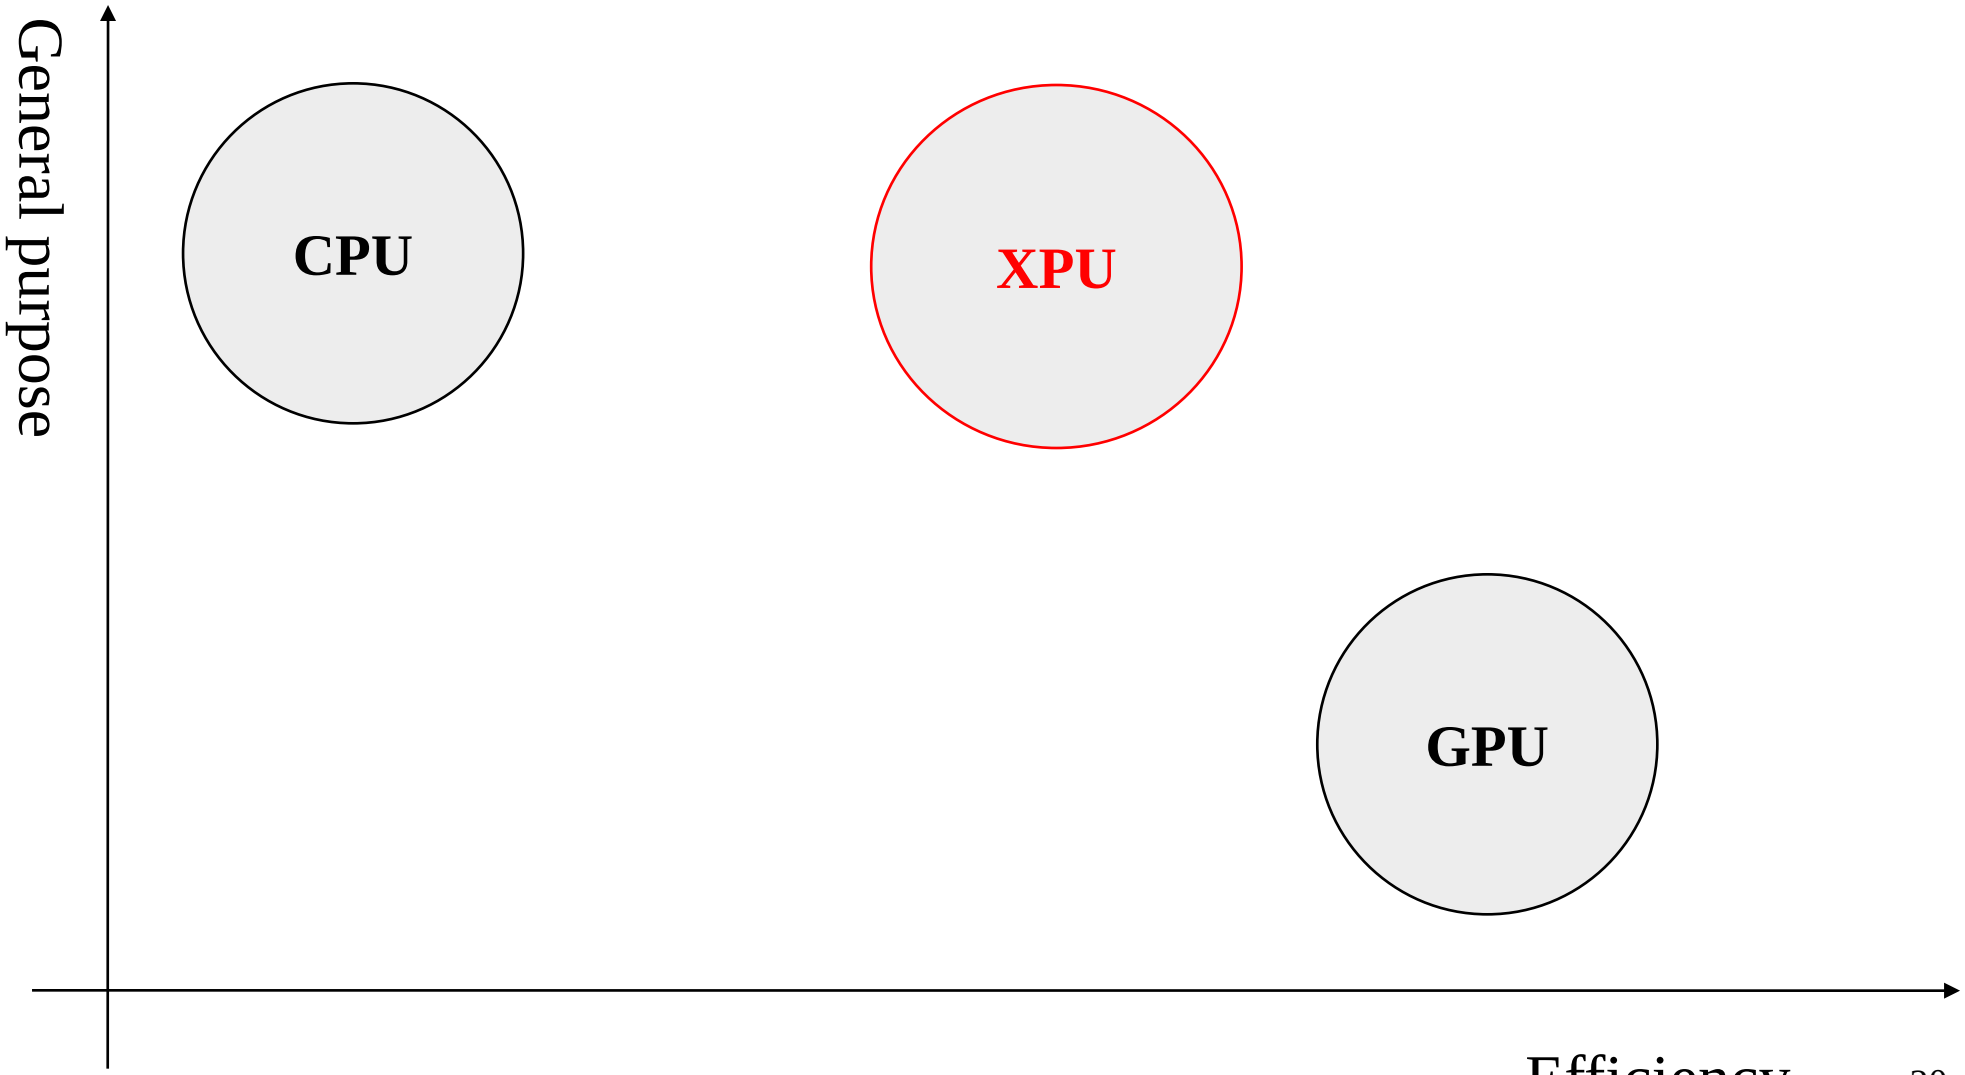

GS464V Architecture Features

- **Keeps all GS464 features, and extended by:**
- **Extend each 64-bit FP unit to 256-bit SIMD vector unit**
  - **Two vector units (dual issue 256-bit SIMD)**
  - **Each core has eight 64-bit MACs**
  - **Keep MIPS64 compatible**
- **128-entry 256-bit register file**
- **300+ SIMD instructions (Linpack, FFT, filter, media……)**
- **The XPU architecture**
- **Adopted by Godson-2H/Godson-3B**

Micro-architecture of GS464V

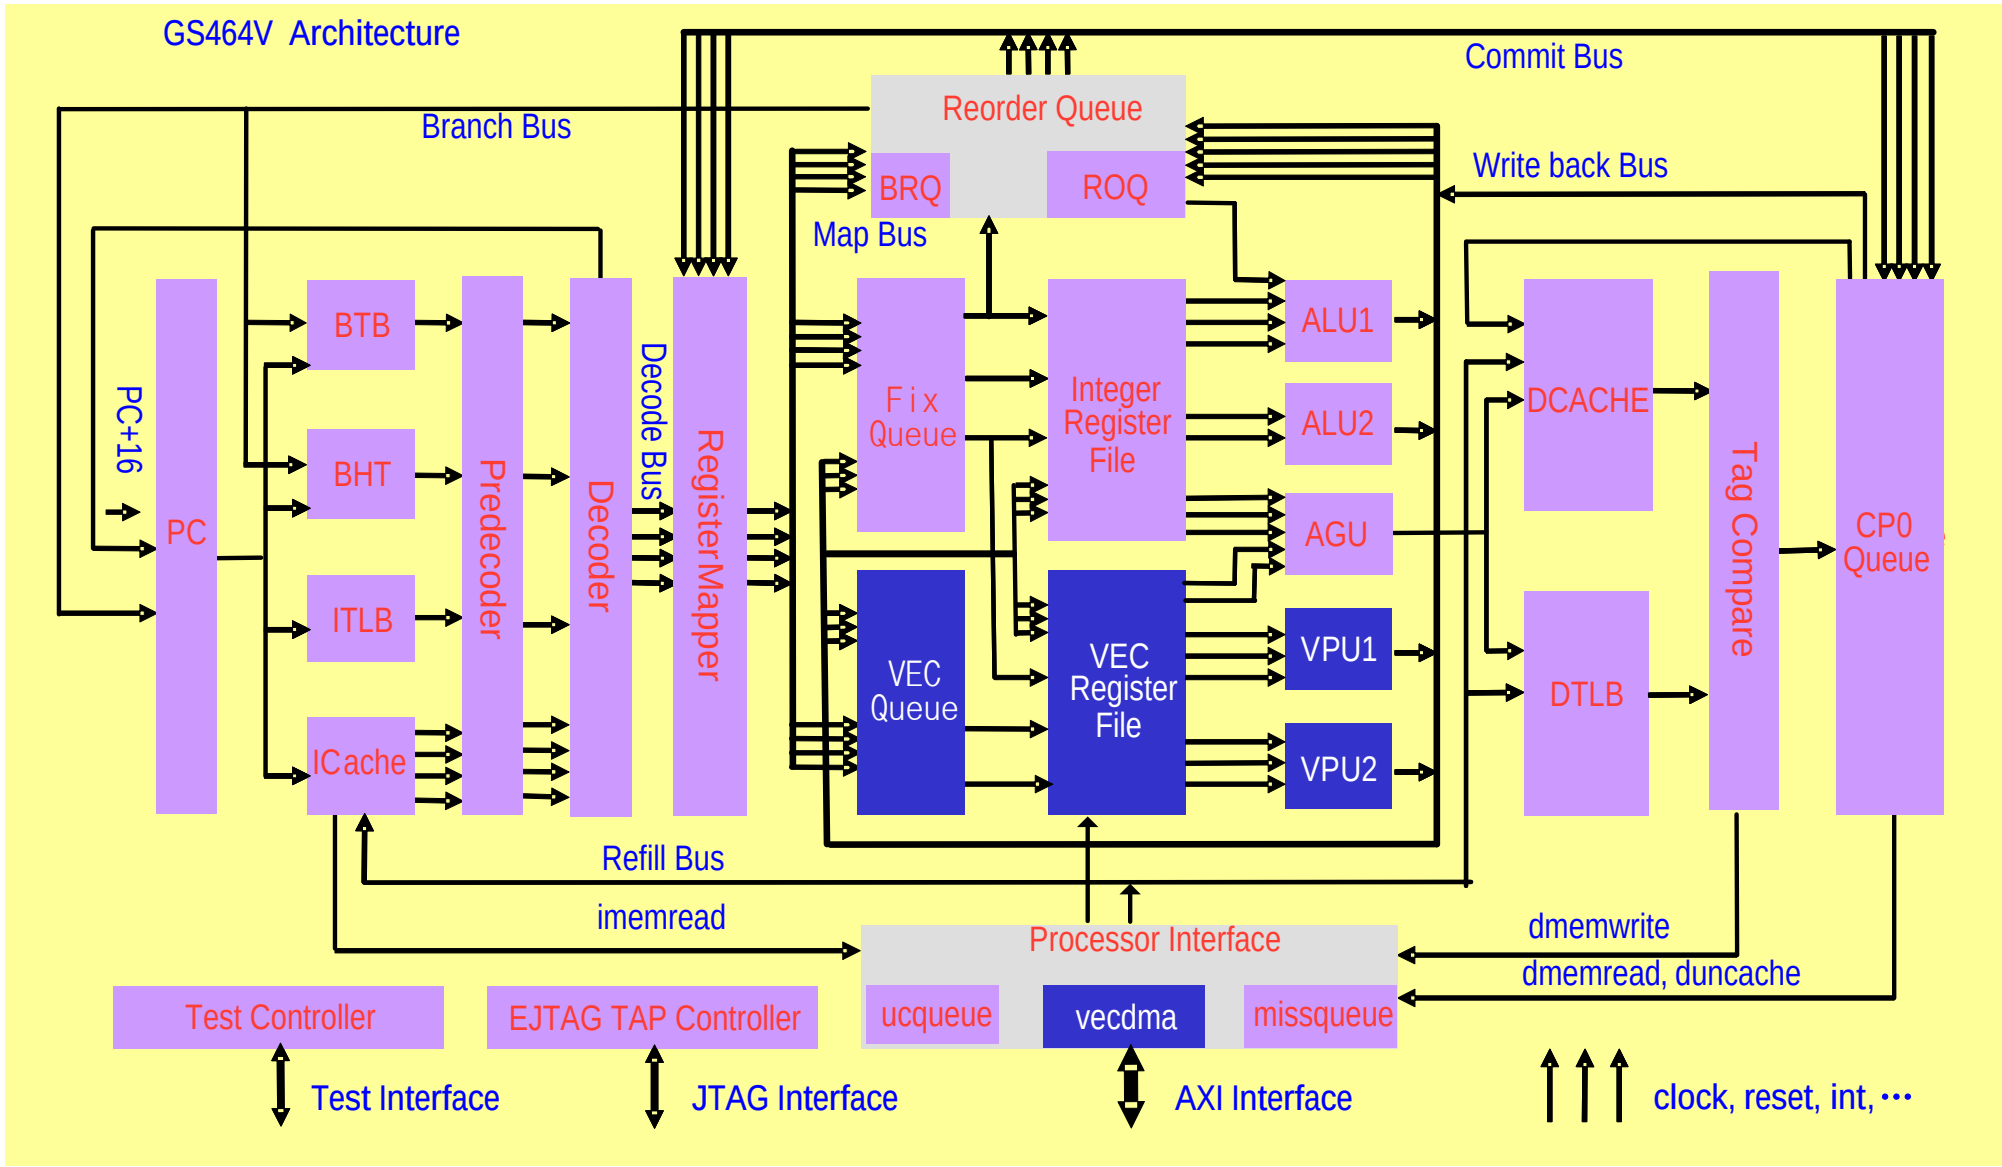

Feeding the Starving Vector Unit

- **With limited bandwidth, how to provide enough data in required format to feed the vector unit ?**
  - **Two 256-bit vector unit, limited bandwidth**
  - **Use the bandwidth more efficiently**
- **Special data link for vector unit**
  - **Data moves in parallel with computation**
  - **Data are move continuously in large blocks**
  - **Reorganizing data in the way from memory/cache to VR, as required by applications.**

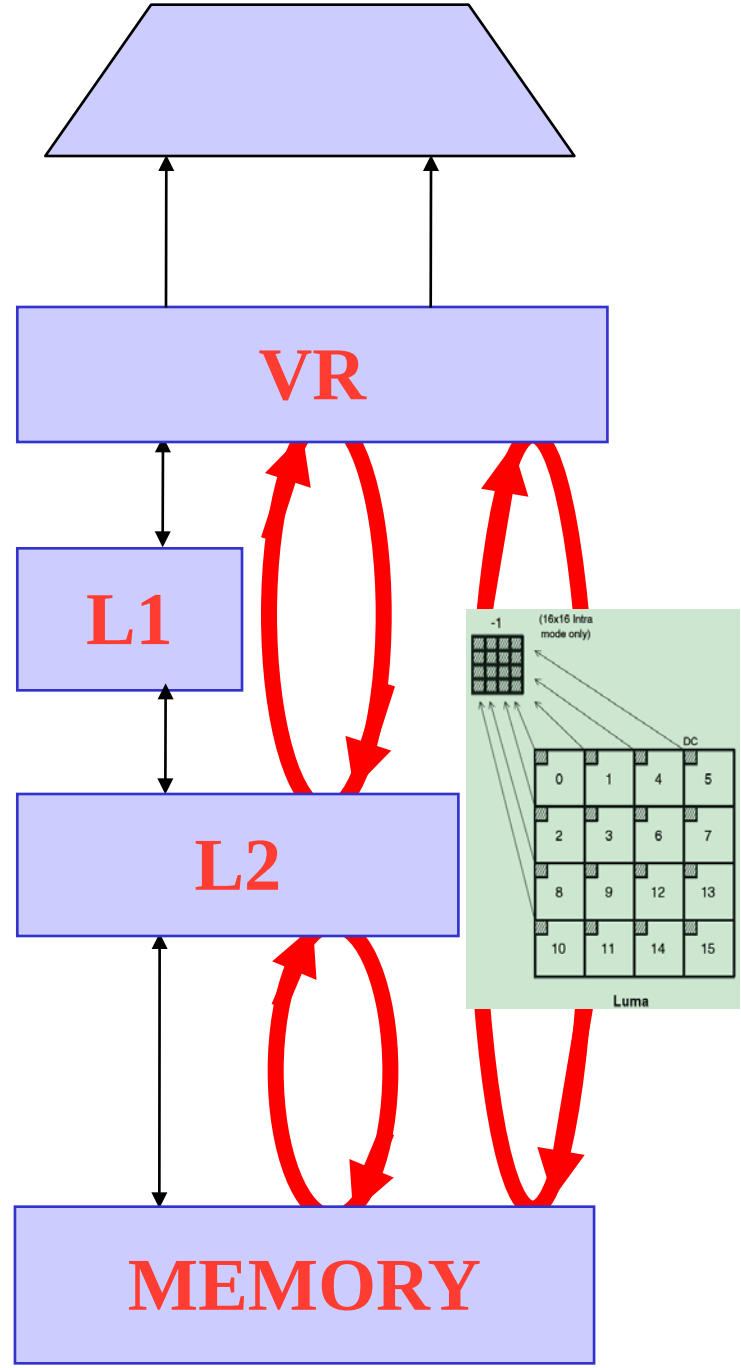

Godson Super Link (GSL)

- **Direct and reconfigurable data transfer between vector register and cache/memory**
  - **Matrix transposing**
  - **Bit revert**
  - **Entropy decoding**
  - **……**
- **Memory access coprocessor**
  - **Multiple Godson Super-Link**
  - **Flow control among GSLs**
  - **Synchronize with GS464V**

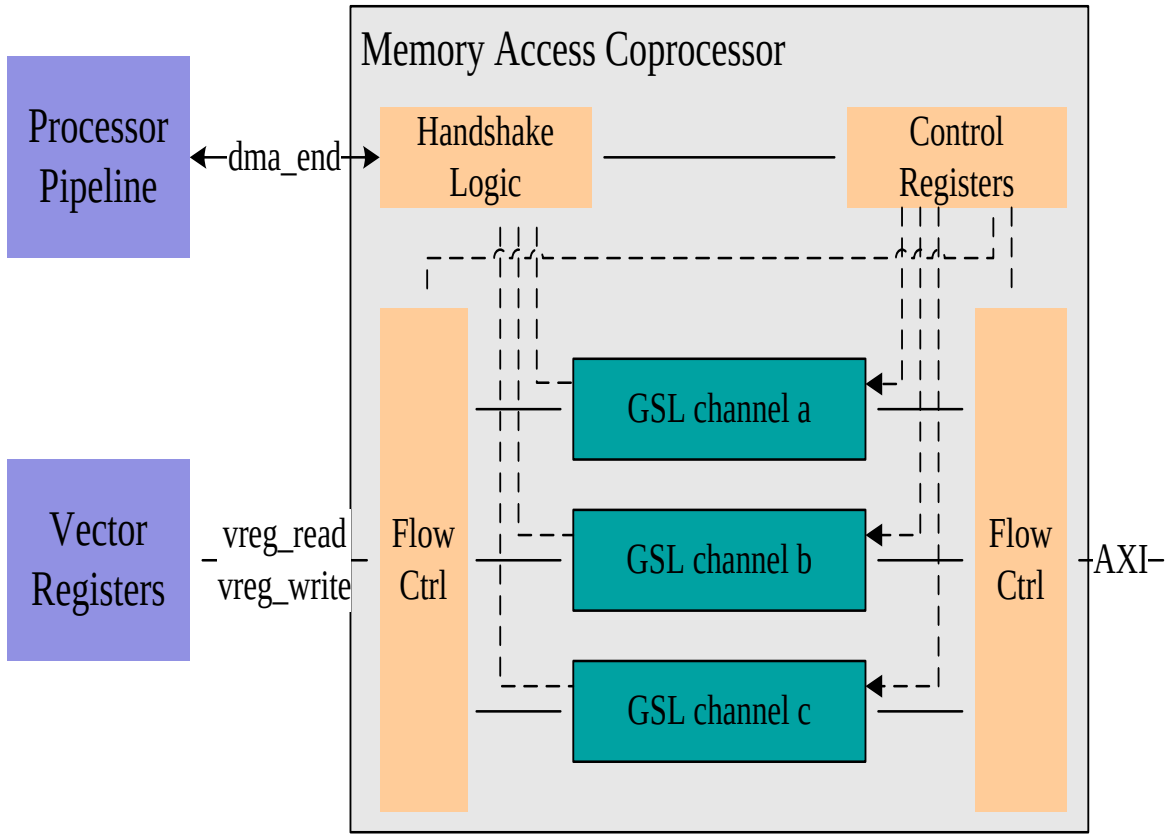

Programing Model

```
{
  ......
  GSL_transfer_data(from, to, format);
  ......                  // independent codes
  GSL_wait_transfer_complete;
  ......                  // codes depend on the transferred data
}
```

Parallel data movement and computation: Linpack as an example

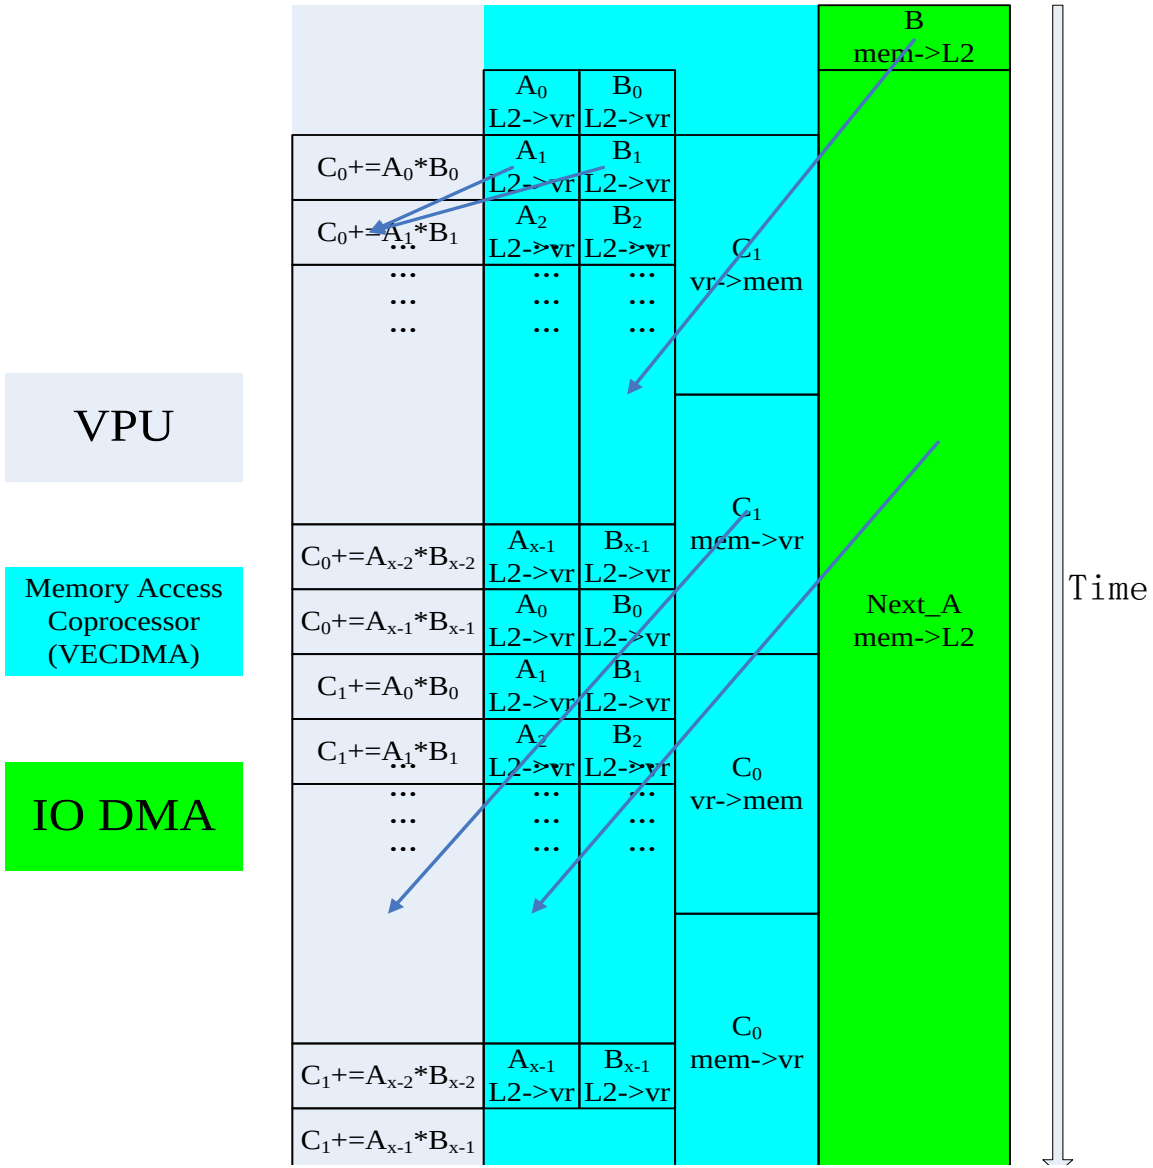

- **Computation and shuffle in parallel**
- **Computation and data transferring in parallel**
  - **VR↔L2**
  - **VR↔mem**

Computation and Shuffling in one Instruction: Vector Instruction for FFT

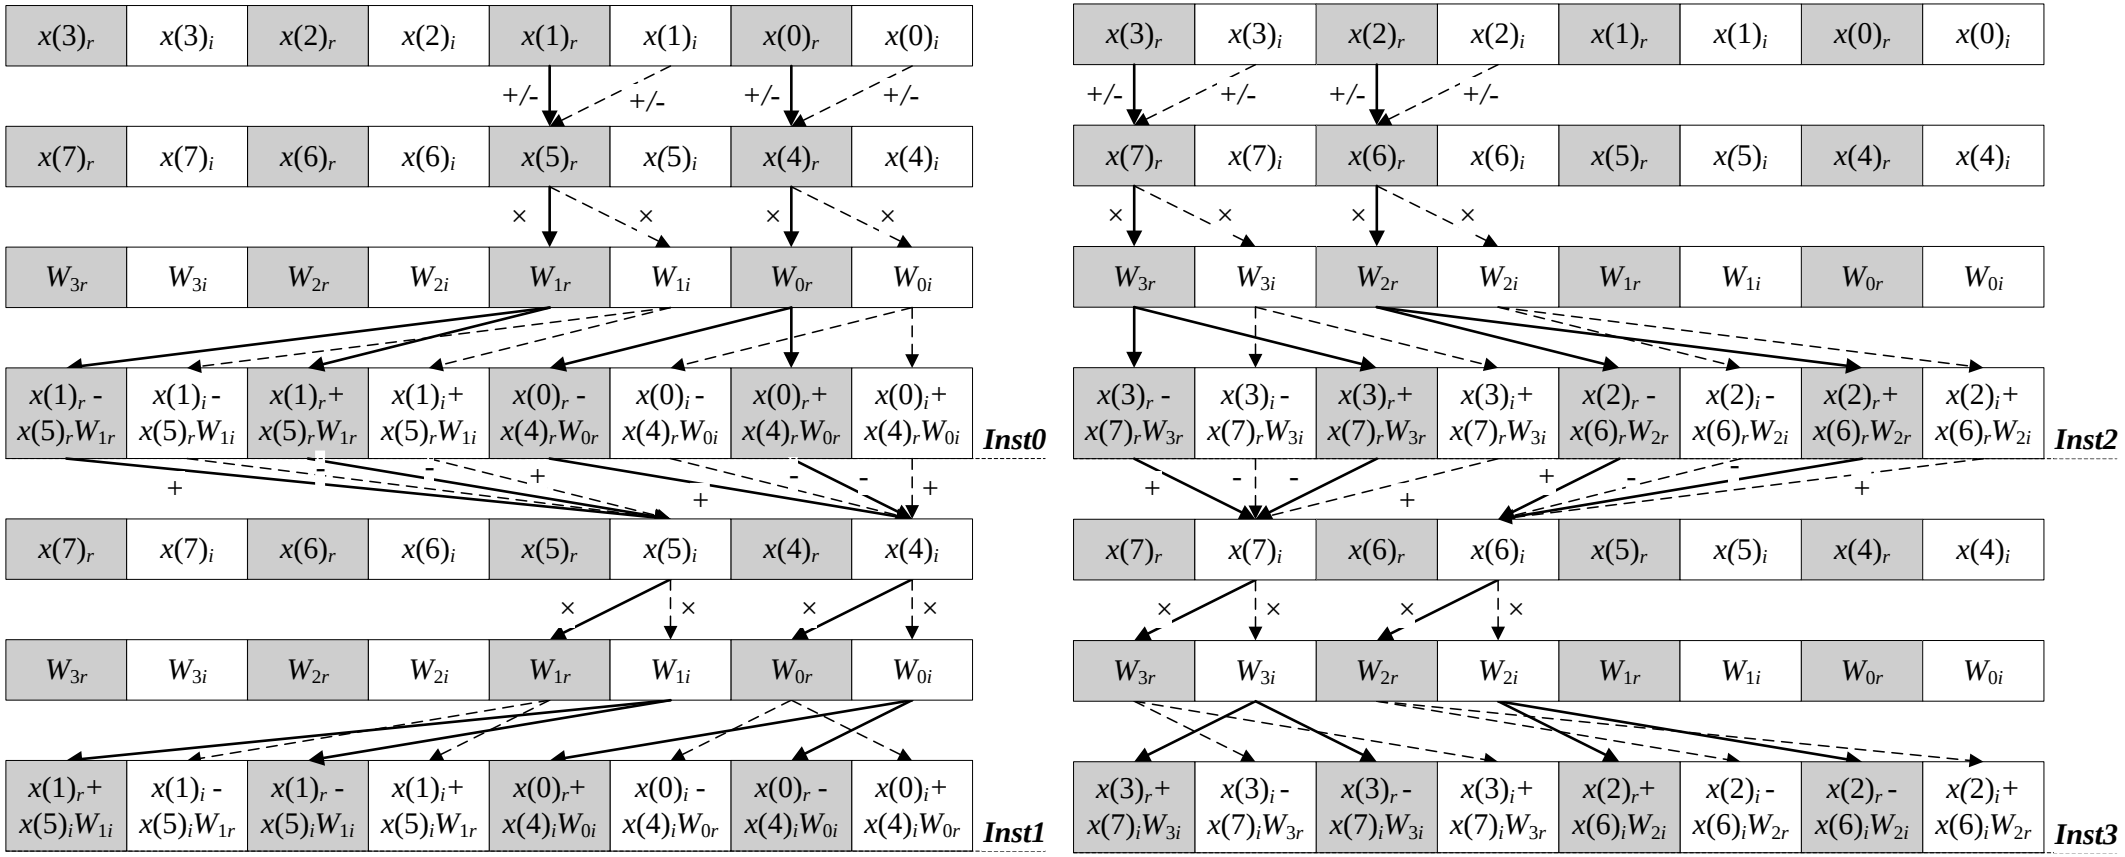

Computation and Shuffling in one Instruction: Vector Instruction for Linpack

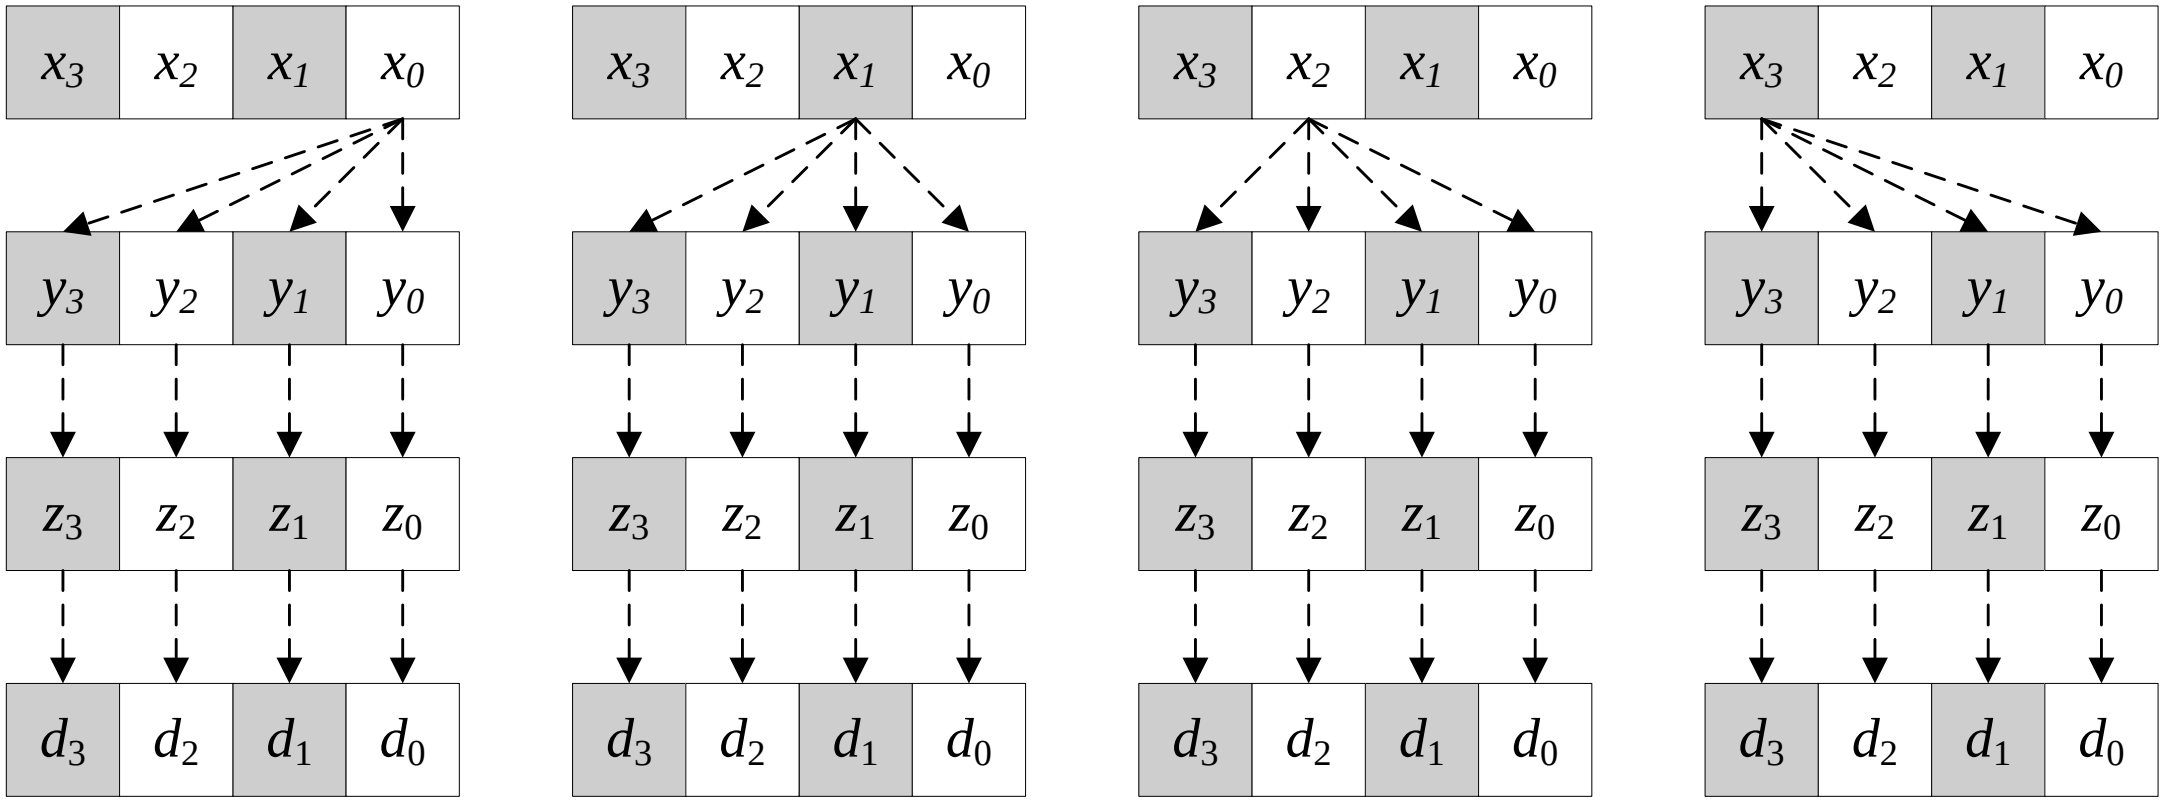

Computation and Shuffling in one Instruction: Vector Instruction for Media

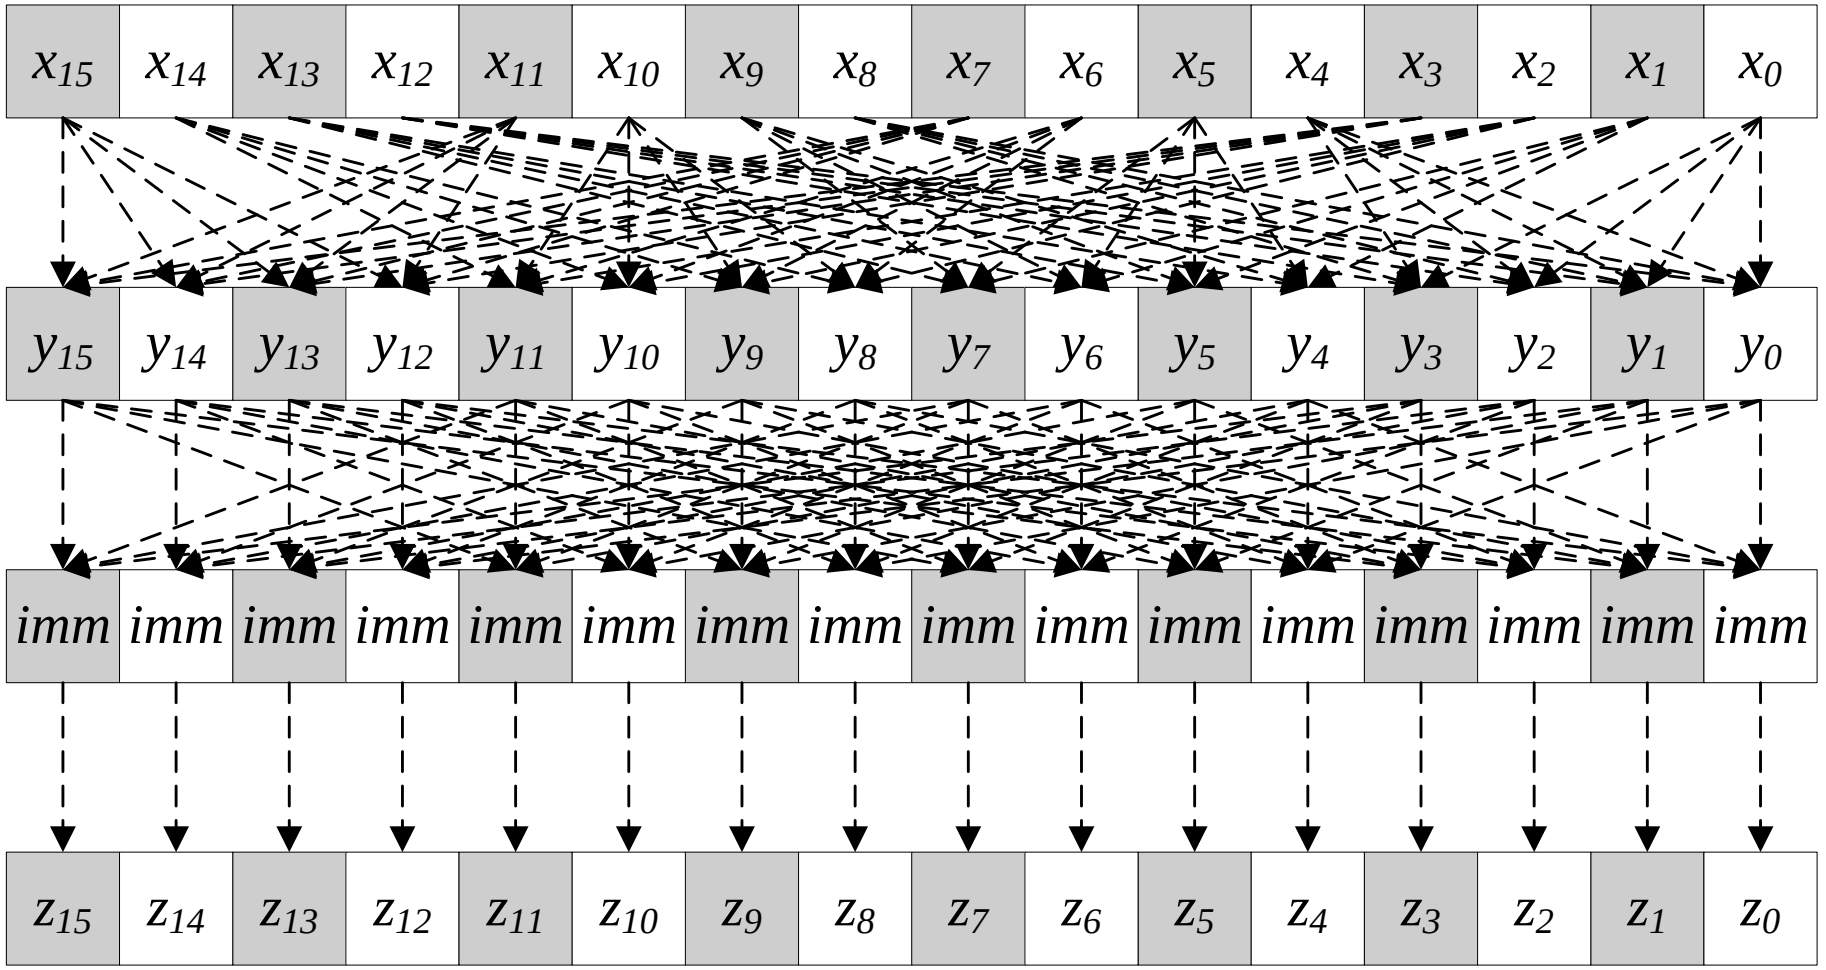

Performance Results

- **With FPGA prototyping and RTL simulation, GS464V achieves excellent performance**
- **Eight-core godson-3B (64 MACs, 2 DDR3) achieves**
  - **>93% of peak performance for matrix multiplication**
  - **>87% of peak performance for 1024 point complex FFT**
  - **0.37us for 1024 point floating point FFT at 1GHz**
- **1080p high definition H.264 decoding with single core at 1GHz**
  - **>100 frames per second**

GDSII of GS464V (65nm)

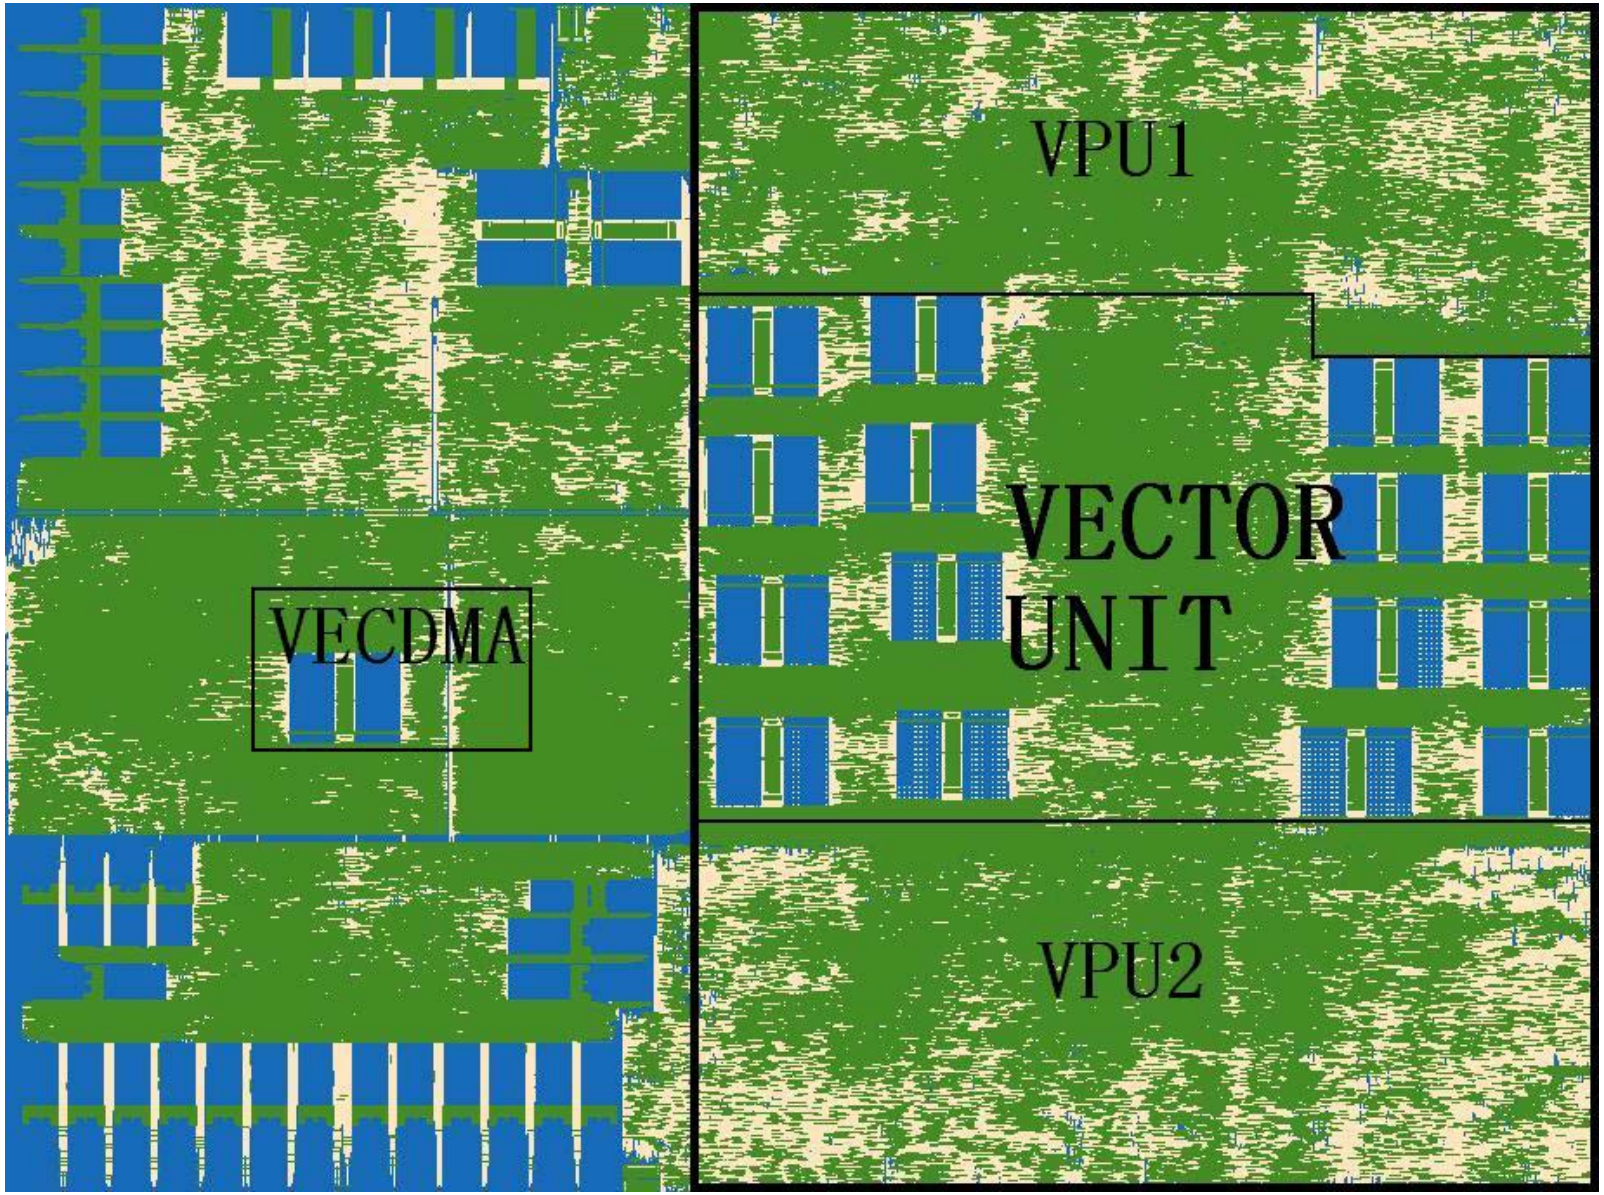

Contents

- **Background**
- **The XPU architecture of GS464V**
- **Godson chips based on GS464V**
  - **Multi-core chips for servers and HPCs**
  - **SOC chips for low-cost PC**

8-core Godson-3B

- 8 four-issue 64-bit core
- 2*256-bit Vector Ext. per core
- 1.0GHz@65nm
- 128GFLOPS@40W
- 2 DDR3, 2 HT Controllers
- 583M xtors, 300mm^2
- Taped out 2010.5

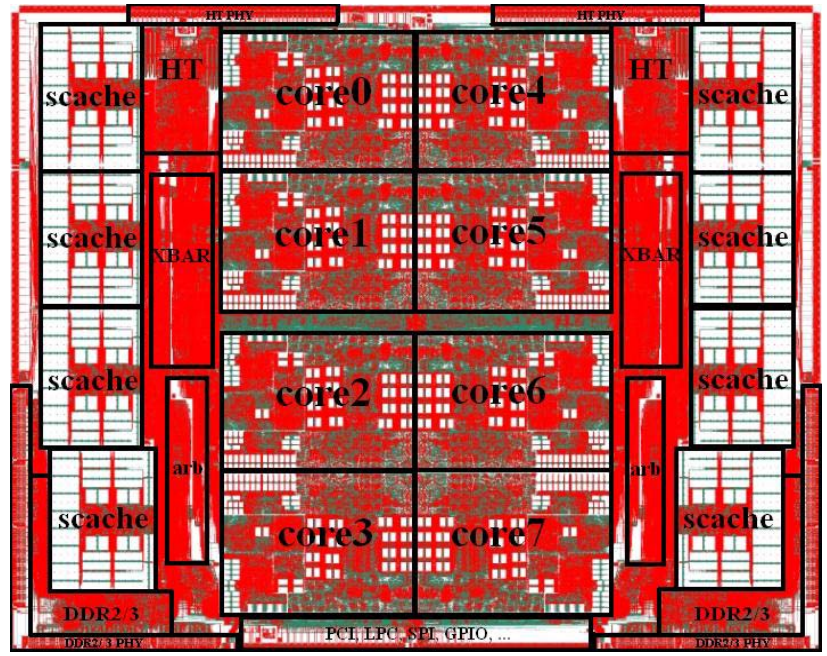

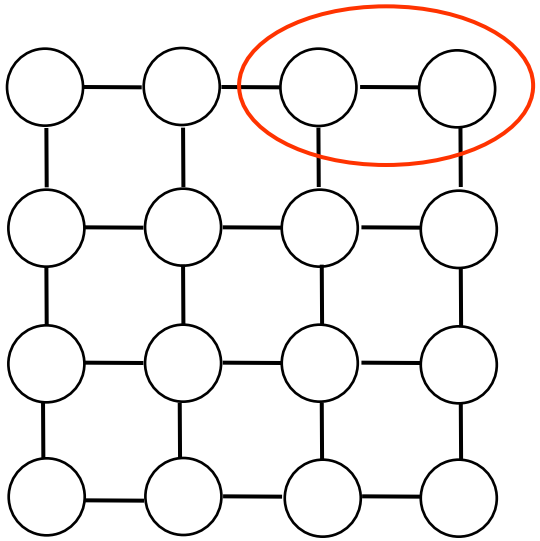

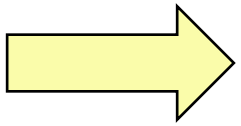

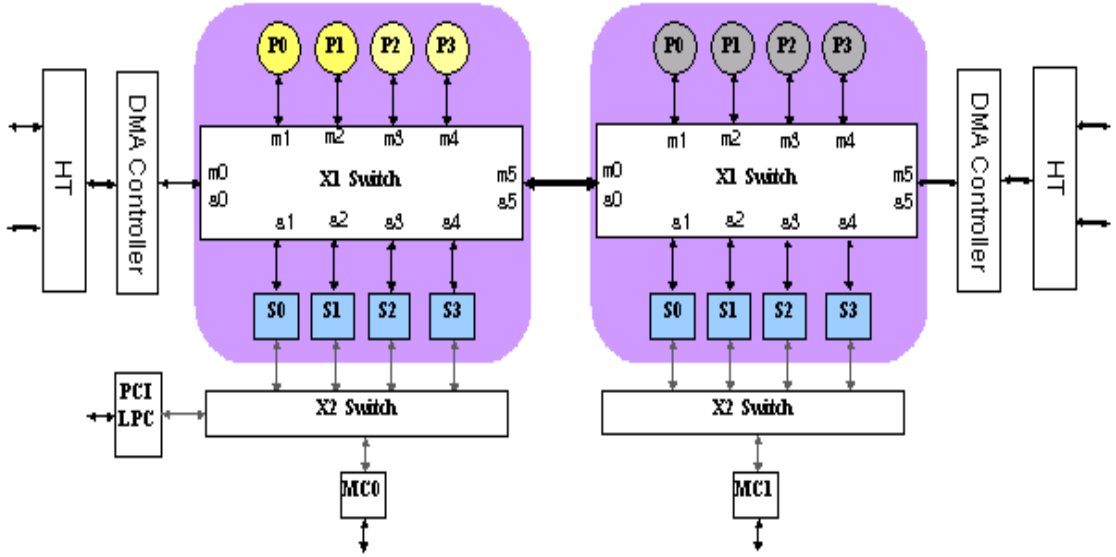

16-core Godson-3C

16 four-issue 64-bit Core

2*256-bit Vector Ext. per core

1.5-2.0GHz@28nm

384-512GFLOPS@20W

4 DDR3, 4 HT Controllers

To be taped out 2011

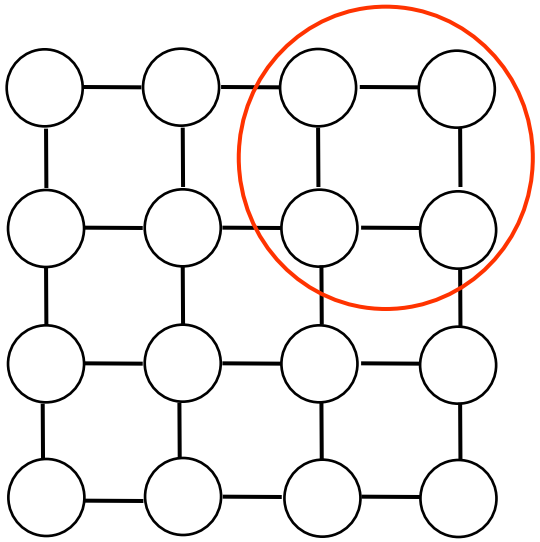

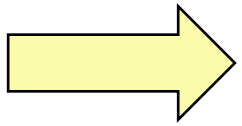

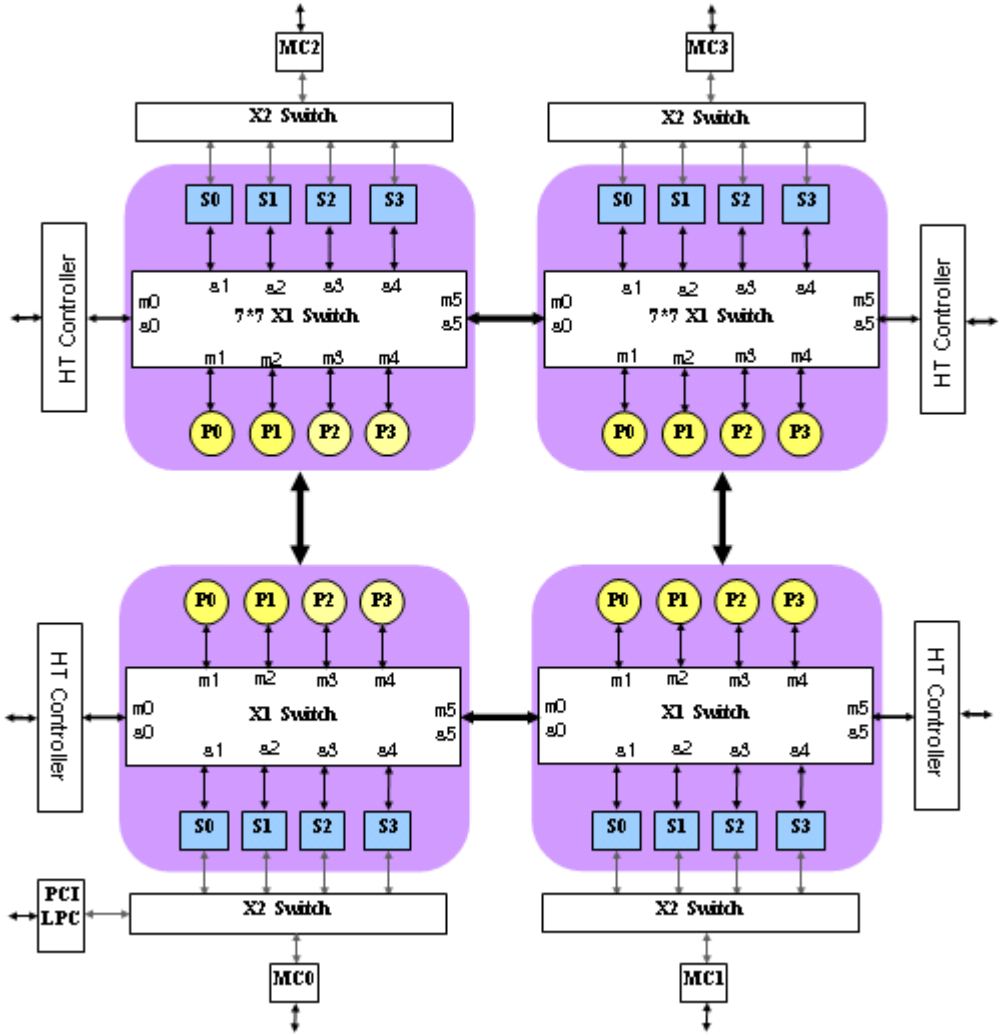

Godson-2H for Low-cost PC

- 1GHz@65nm
- GS464V (HD media decoding)
- 512KB L2
- 3D low power GPU
- DDR2/3 memory controller
- PCIE 2.0 controller
- SATA, USB, GMAC controller
- LPC, SPI, UART, etc.
- Single chip solution for low cost PC
- To be taped out 2010Q4

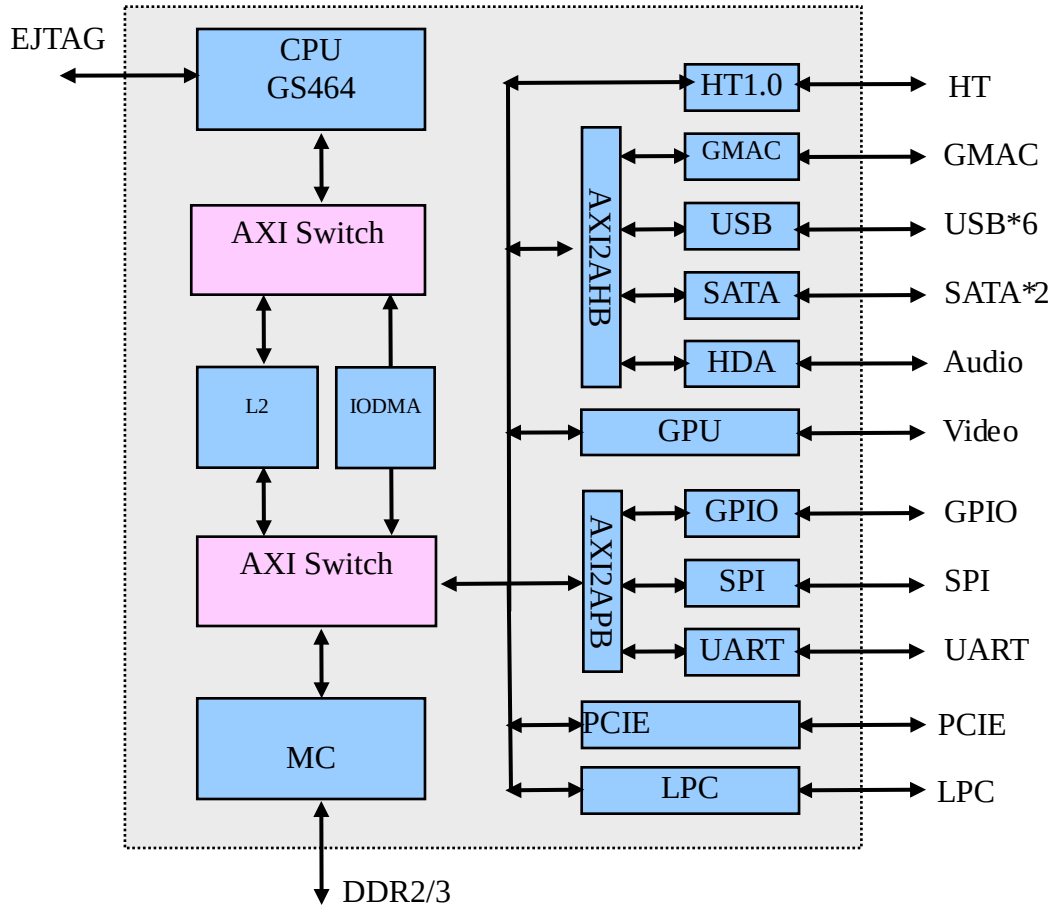

龙芯2H结构图

Dawning Blades with Godson-3A/3B

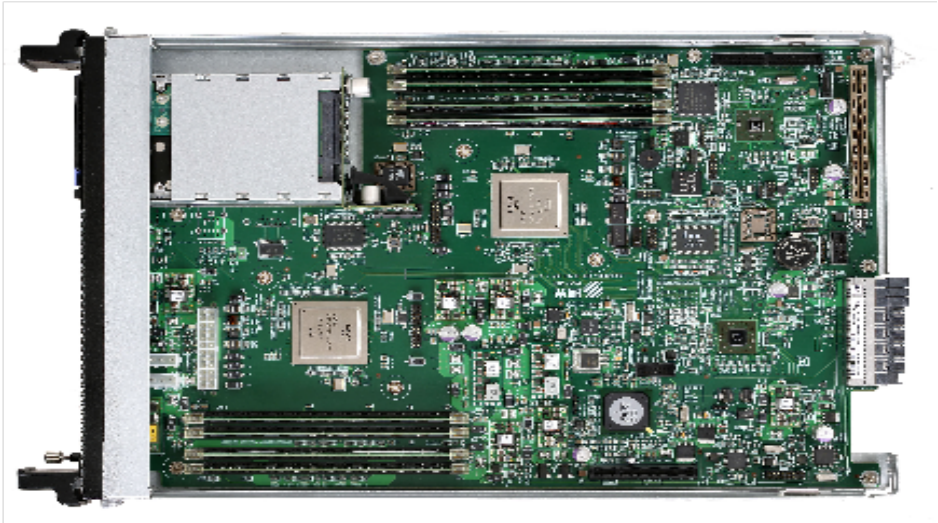

Personal HPC Node

Opteron Card

1

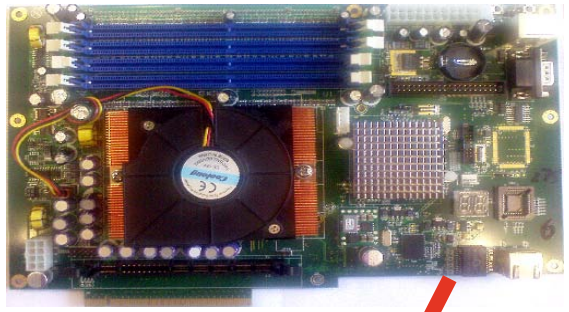

Dual Godson-3A/3B card

4

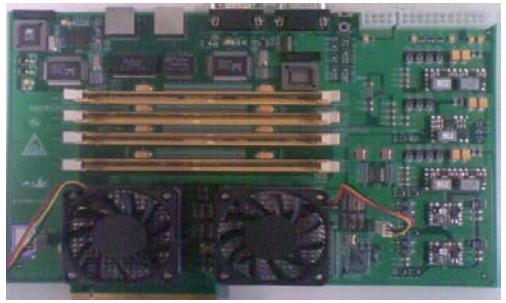

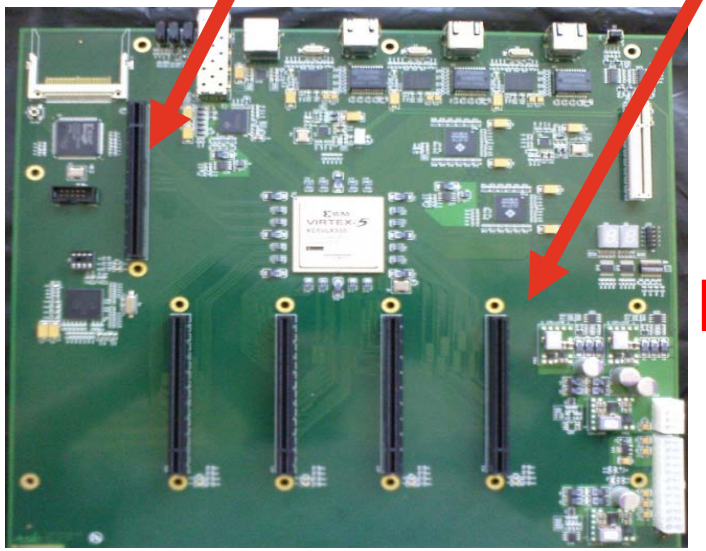

5HT+1 PCI-E x8 board

HPP node

1U16P Board for HPC

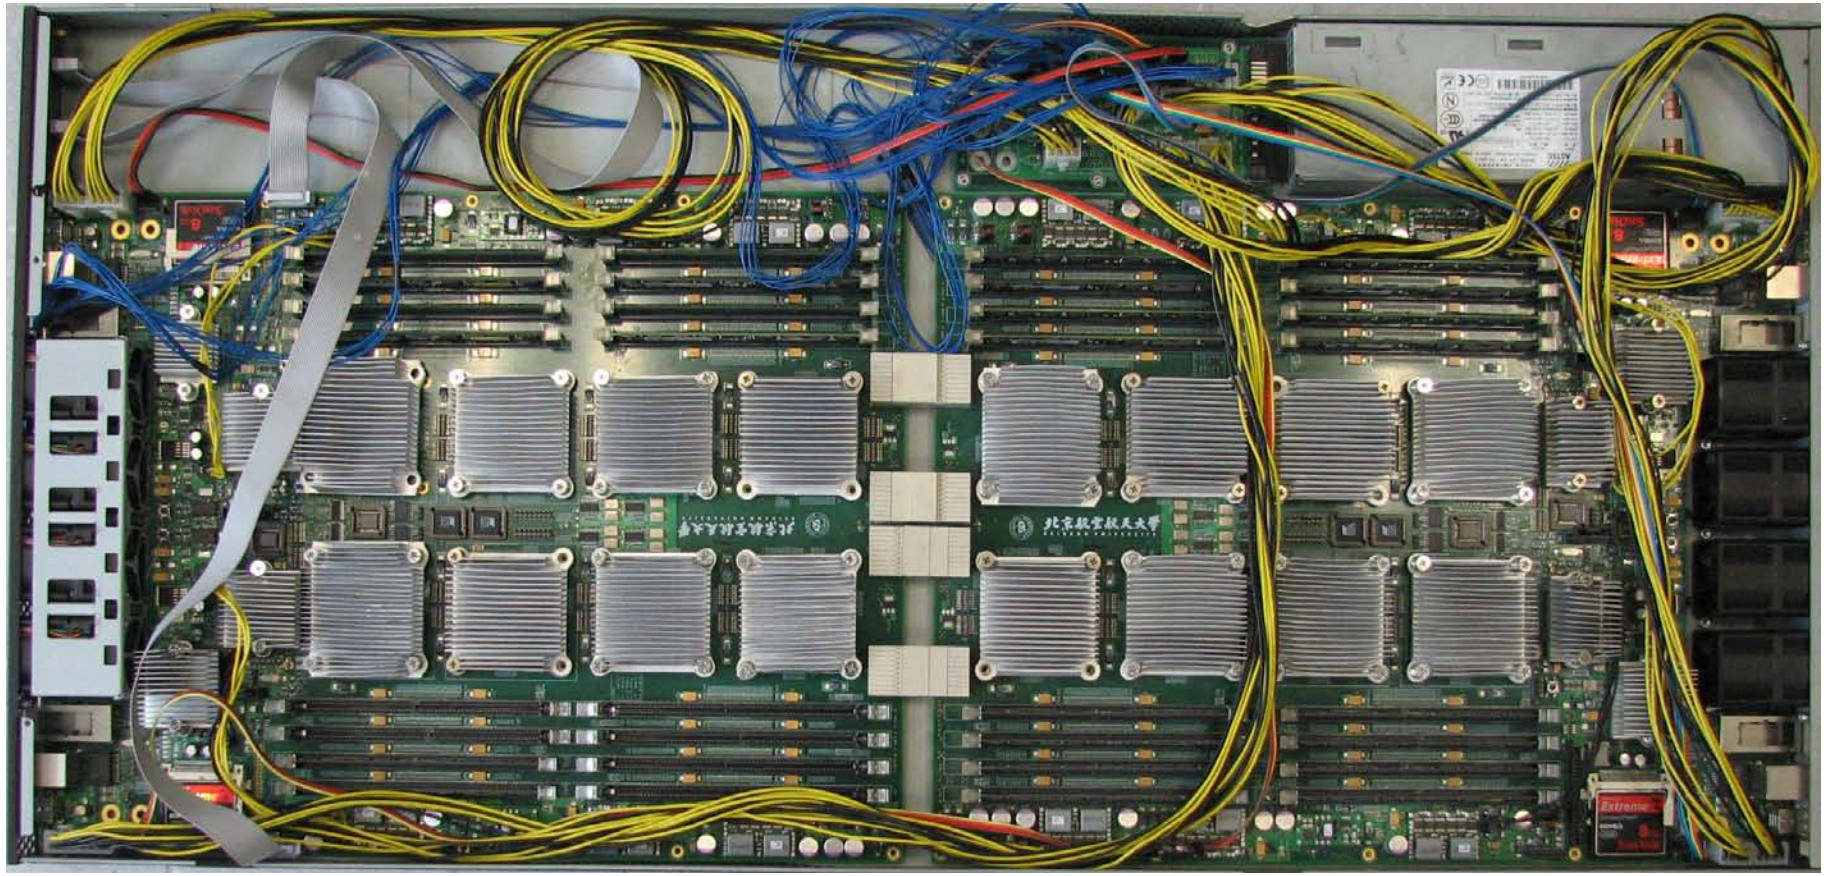

1U2T with Godson-3B

1U8T with Godson-3C

Conclusion

- **XPU: convergence of CPU/DSP/MPU/GPU**
  - **Vector Unit + Godson Super Link**
  - **Provides enough data with correct format to feed the starving CPU**
  - **Data is reorganized in the way from memory to register**
- **Achieves high performance**
  - **93% Matrix multiplication, 87% FFT for 8-core Godson-3B**
  - **>100 frames 1080p H.264 decoding for single core 1GHz Godson-2H**
- **Godson chips and applications**
  - **Godson-3B/3C: High-end computing**
  - **Godson-2H: Low-cost PC and media applications**

Thanks

Abbreviations

- **CPU** **Central Processing Unit**
- **MPU** **Media Processing Unit**
- **GPU** **Graphic Processing Unit**
- **XPU** **eXtra Processing Unit**
- **OOO** **Out-of-Order**
- **HT** **HyperTransport**
- **NB** **North Bridge**
- **SB** **South Bridge**
- **AXI** **An Open Standard On-chip Interconnect Specification of ARM**
- **GSL** **Godson Super-Link**
- **GDS** **Graphic Database System**
- **SPI** **Serial Peripheral Interface**
- **LPC** **Low Pin Count**