XiangShan: Industrial-Competitive RISC-V CPUs in the Era of Open-Source Chips

Yinan Xu 徐易难

Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS)

October 2025

Outline

- The Era of Open-Source Chips
- Open-Source Industrial-Competitive RISC-V Chips
- Open-Source Chip Design Tools
- Open-Source Development & Business Models
- Conclusions

Dozens of ISAs in the past 50 years

- *Over the past half-century*, **dozens** of instruction sets have appeared, **all owned by private companies**.
  - Most have *disappeared* as their companies merged or shut down.

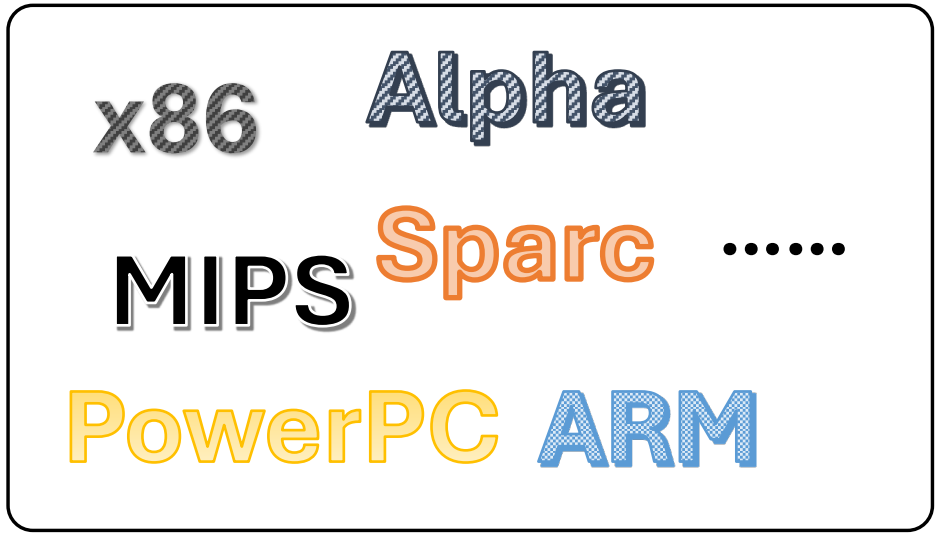

Today: CPU Market Dominated by x86/ARM

- *Over the past half-century*, **dozens** of instruction sets have appeared, **all owned by private companies**.
  - Most have *disappeared* as their companies merged or shut down.
- *Today*, **only x86 and ARM** remain as *mainstream* instruction sets, each with its own ecosystem.

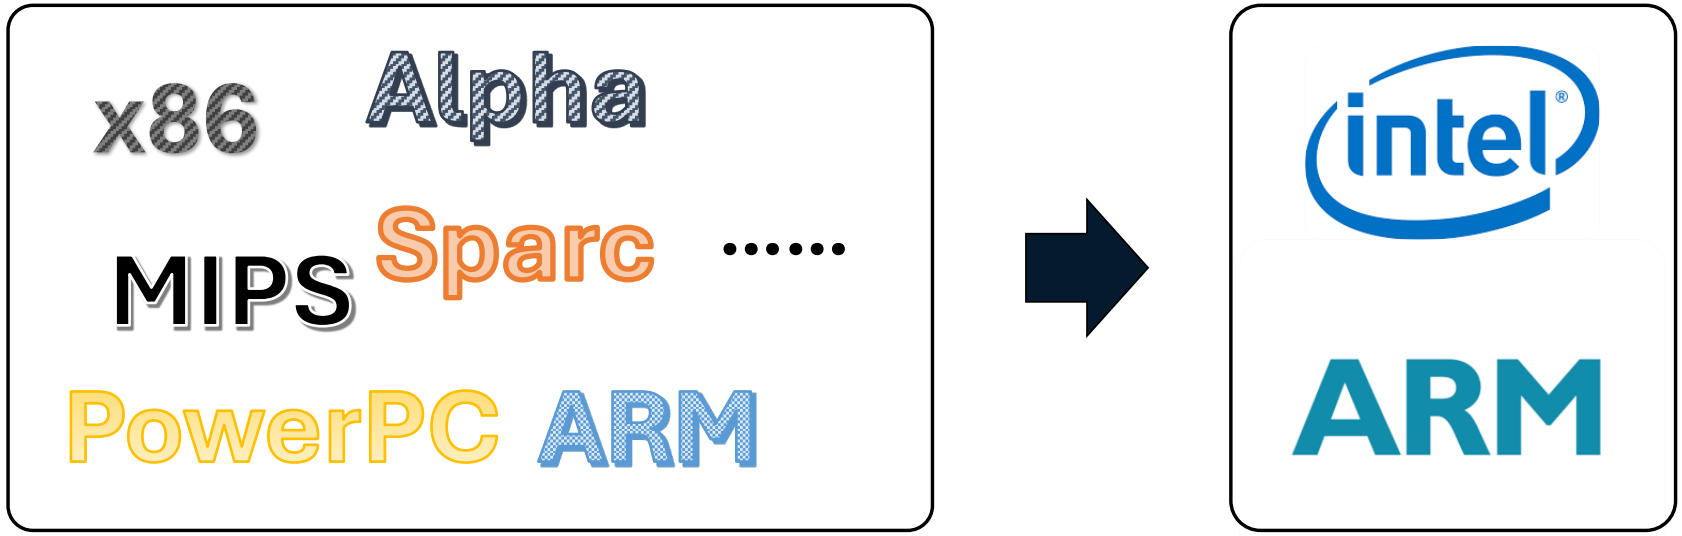

Open-Source Software Ecosystem

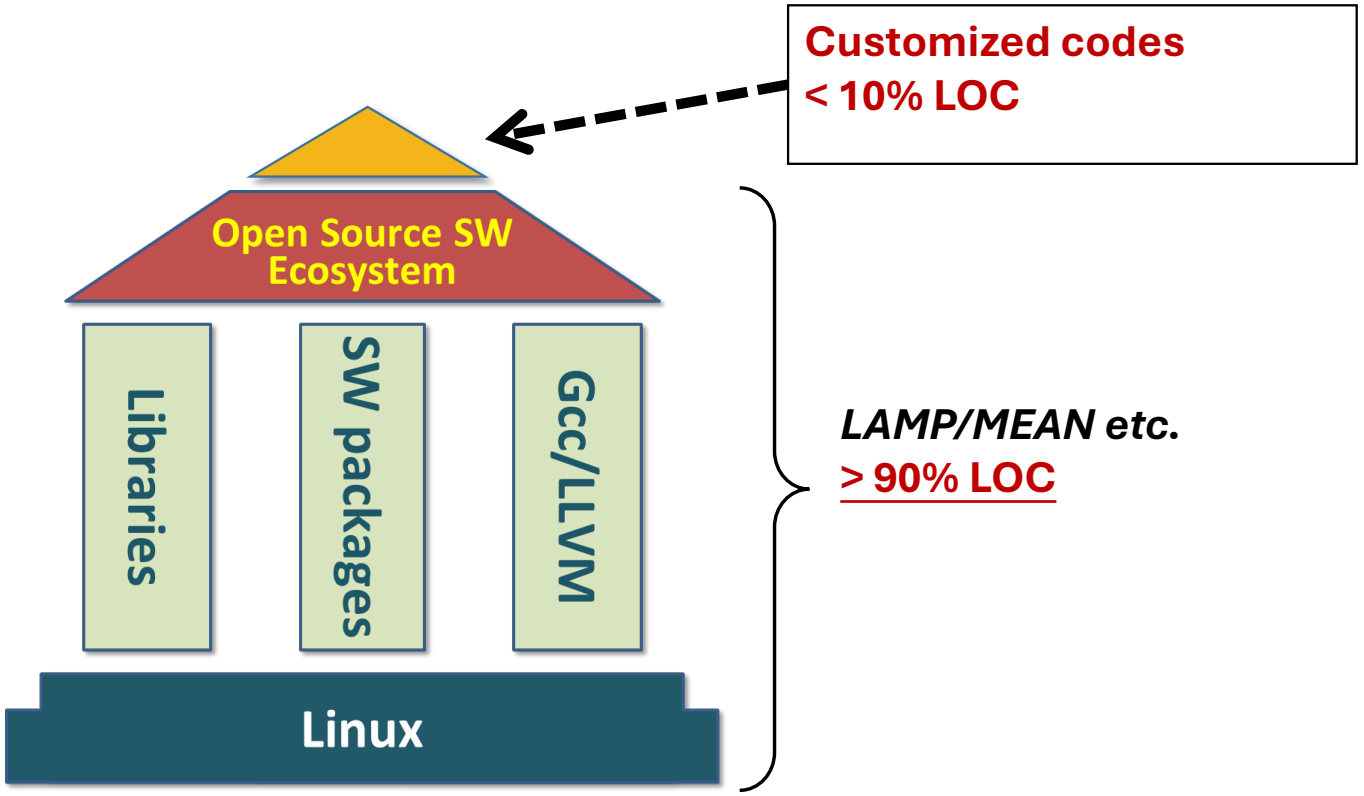

Mirror the success of the open-source software ecosystem?

Open-Source Chip Ecosystem (OSCE)

To Lower the barrier of chip development

By saving the cost of IPs, EDA tools and engineers in chip design

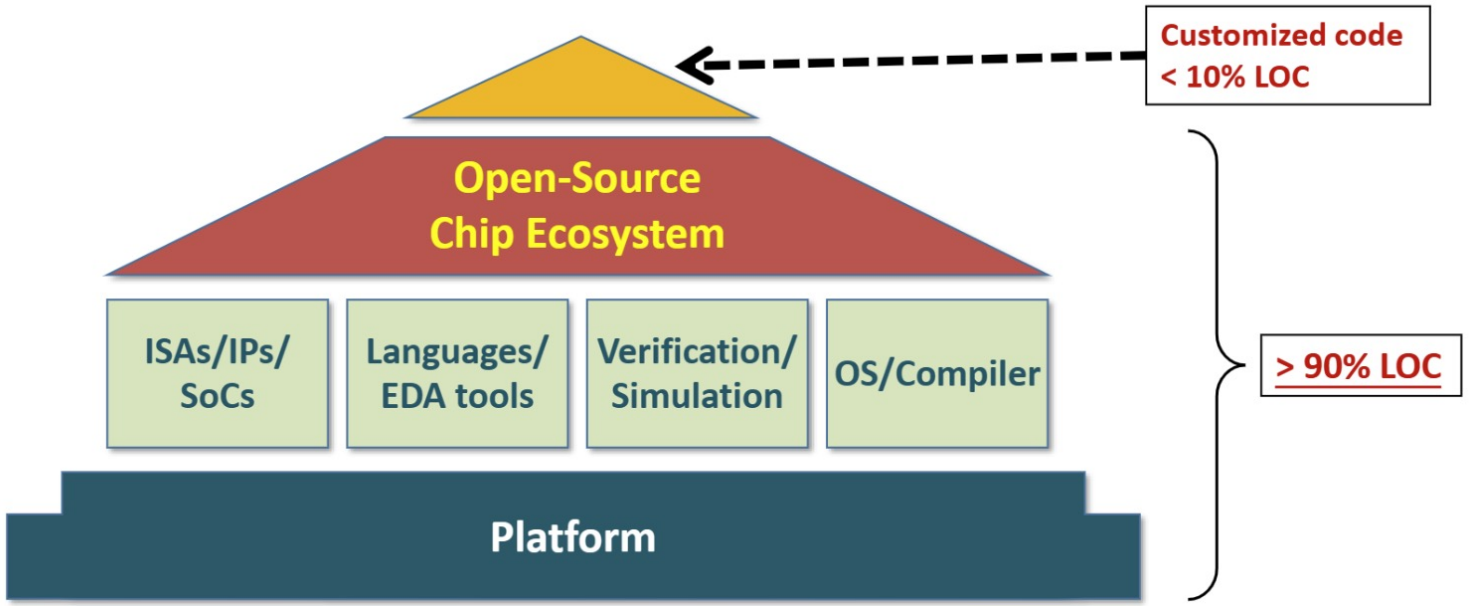

* Yungang Bao, The Four Steps to An Open-Source Chip Design Ecosystem, ACM SIGARCH Visioning Workshop, June 2019

Three levels of OSCE

- **L1: OPEN** ISA
- **L2: OPEN** Design & Implementation
- **L3: OPEN** Tools & Infrastructure

3 **Open Tools & Infrastructure**

Microarch. → Engineering → EDA Tools

ISA Spec.

RISC-V®

1 **Open ISA**

Docs

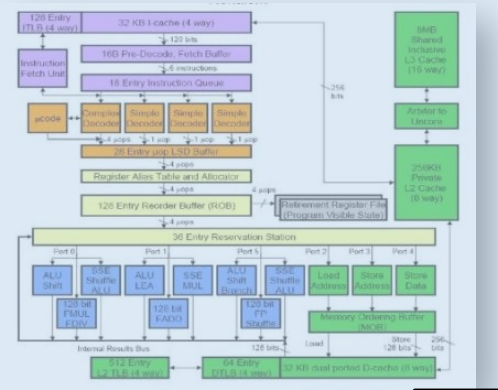

RTL codes

```
component DebugCoreTop is
  port (
    -- Trigger and Data
    cu_Clk      : in    std_logic_vector(2 downto 0) := (others => '0');
    cu0_Trig    : in    t_trig_0 := (others => (others => '0'));
    cu1_Trig    : in    t_trig_1 := (others => (others => '0'));
    cu2_Trig    : in    t_trig_2 := (others => (others => '0'));
    cu0_Data    : in    t_data_0 := (others => (others => '0'));
    cu1_Data    : in    t_data_1 := (others => (others => '0'));
    cu2_Data    : in    t_data_2 := (others => (others => '0'));

    -- Downstream I2C
    SCL         : in    std_logic := '0';
    SDA         : inout std_logic := '0';

    -- Upstream
    gt_RefClk_p : in    std_logic := '0';
    gt_RefClk_n : in    std_logic := '0';
    gt_RX_p     : in    std_logic_vector(2 downto 0) := (others => '0');
    gt_RX_n     : in    std_logic_vector(2 downto 0) := (others => '0');
    gt_TX_p     : out   std_logic_vector(2 downto 0);
    gt_TX_n     : out   std_logic_vector(2 downto 0)
  );
end component;
```

2 **Open Design & Implt**

Layout

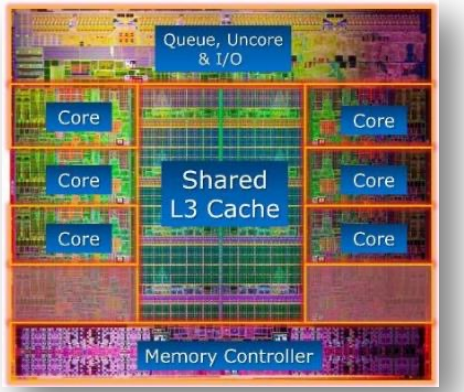

Open Standards: Instruction Sets Want to be Free!

- In 2010, UC Berkeley launched the RISC-V free and open ISA.

Field	*Standard*	*Free, Open Impl.*	*Proprietary Impl.*
Networking	Ethernet, TCP/IP	Many	Many
OS	Posix	Linux, FreeBSD	M/S Windows
Compilers	C	gcc, LLVM	Intel icc, ARMcc
Databases	SQL	MySQL, PostgresSQL	Oracle 12C, M/S DB2
Graphics	OpenGL	Mesa3D	M/S DirectX
ISA	??????	----------	x86, ARM, IBM360

Open-Source: Democratization and Building Consensus

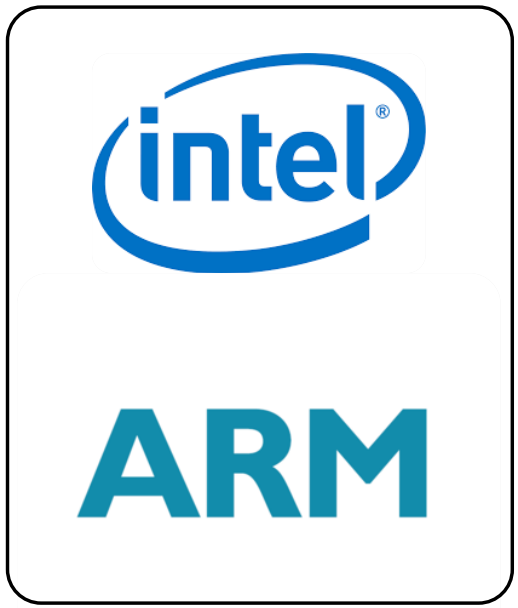

V.S.

- Private owned
- Monopoly

- Open development
- Open and shared ecosystem

A chip design that changes everything

- **10 Breakthrough Technologies 2023**

Ever wonder how your smartphone connects to your Bluetooth speaker, given they were made by different companies? Well, Bluetooth is an open standard, meaning its design specifications, such as the required frequency and its data encoding protocols, are publicly available. Software and hardware based on open standards—Ethernet, Wi-Fi, PDF—have become household names.

Now an open standard known as RISC-V (pronounced "risk five") could change how companies create computer chips.

--- MIT Technology Review

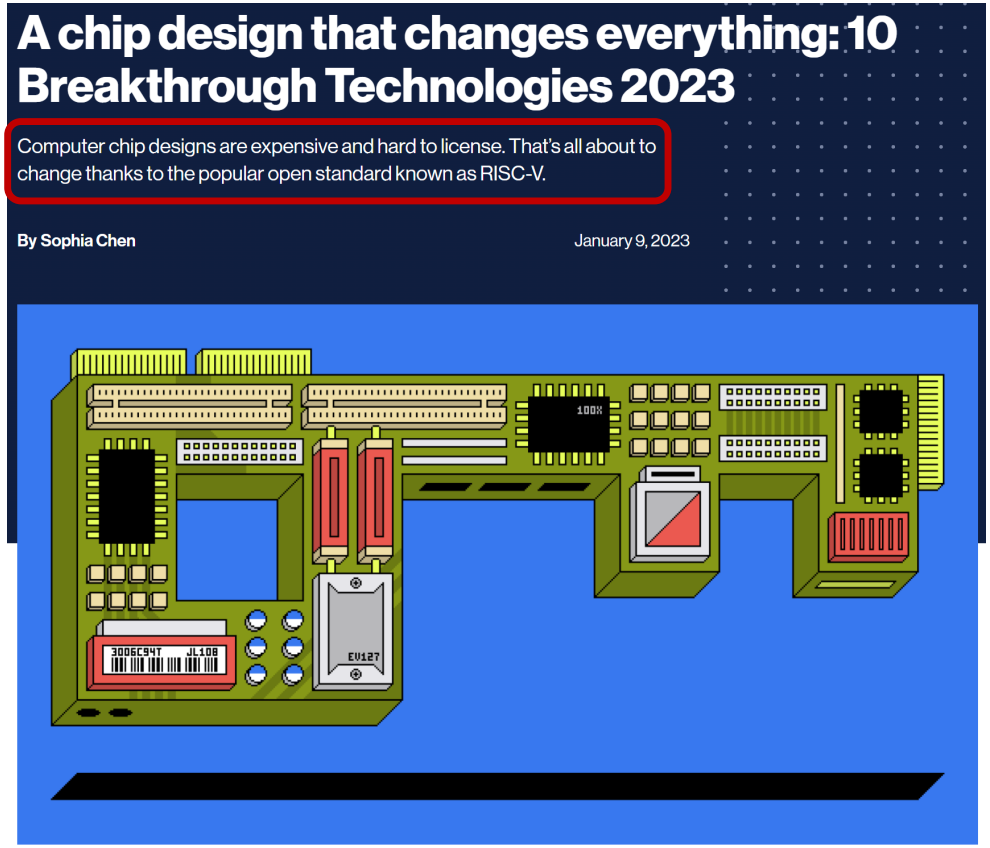

RISC-V: A mainstream market (> $90 billion) by 2030

- The global market size is projected to reach **$92.7 billion** by 2030, with a compound annual growth rate of **47.4%.**

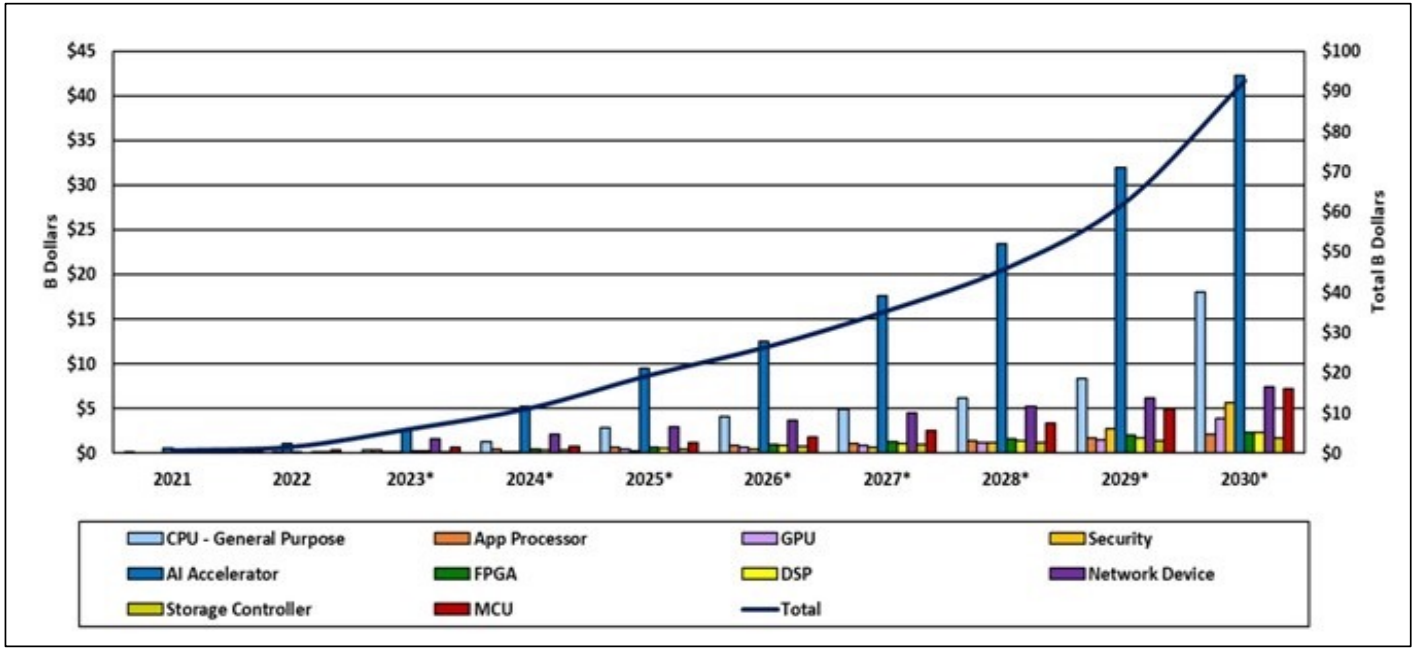

SHDgroup

Global Revenue Forecast for RISC-V SoC Chips

The XiangShan Project

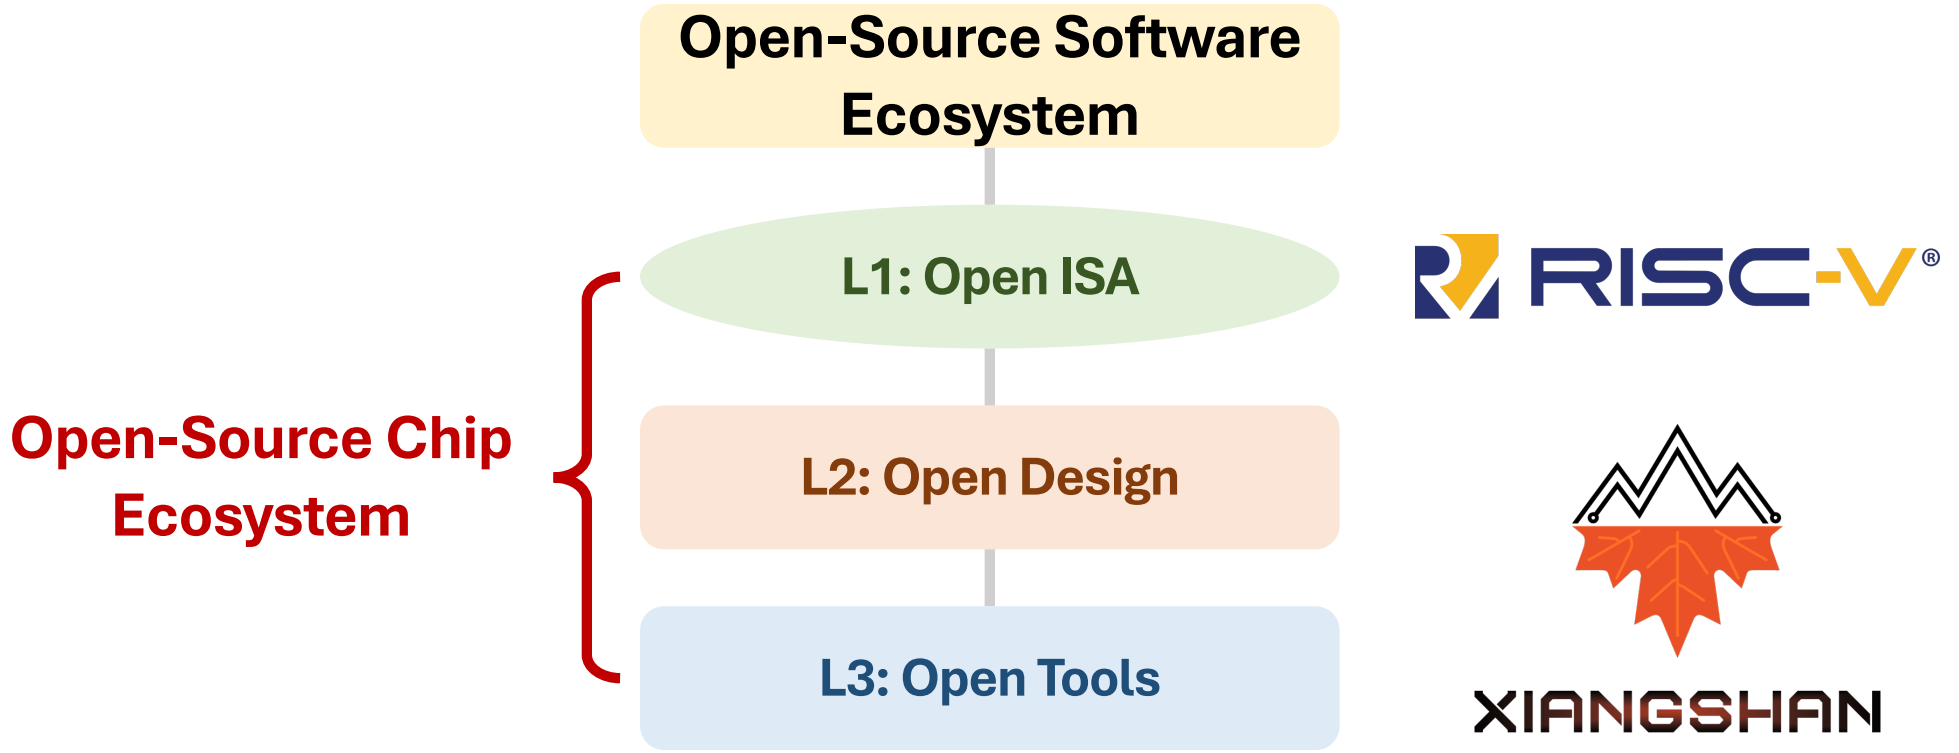

Outline

- The Era of Open-Source Chips
- Open-Source Industrial-Competitive RISC-V Chips
- Open-Source Chip Design Tools
- Open-Source Development & Business Models
- Conclusions

XiangShan: Open-Source High Performance Processors

L1: OPEN ISA (RISC-V)

L2: OPEN Design/Implementation

L3: OPEN Framework/Tools

Fragrant Hill in Beijing

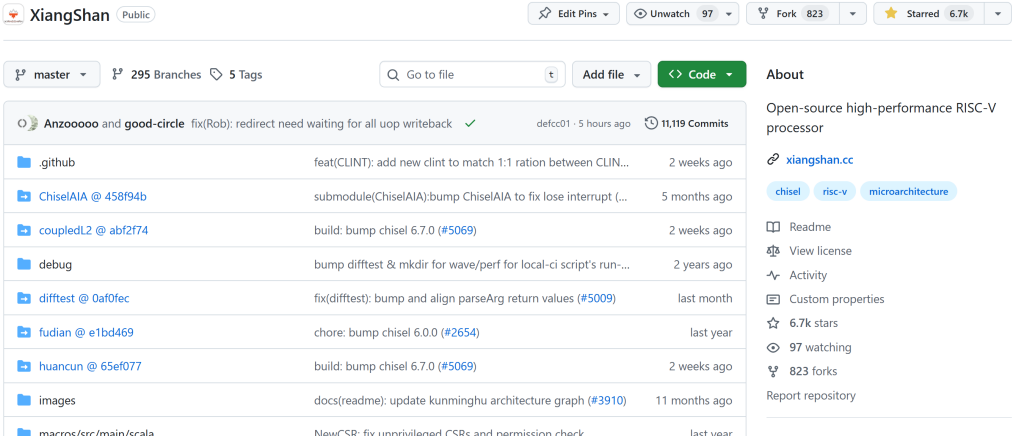

> 6.7K stars, > 820 forks on GitHub

The Open-Source RISC-V CPU

- **Vision**: To establish an open-source RISC-V core backbone like **Linux**, which can be widely used in *industry* and support innovative ideas from *academia*.

Linux

V.S.

XIANGSHAN

XiangShan: Open-Source High Performance Processors

https://github.com/OpenXiangShan/XiangShan

- 1st generation: Yanqihu (YQH)
  - RV64GC, single-core, superscalar OoO
  - **28nm tape-out, 1.3GHz, July 2021**
  - **SPEC CPU2006 7.01@1GHz, DDR4-1600**
- 2nd generation: Nanhu (NH)
  - RV64GCBK, dual-core, superscalar OoO
  - **14nm GDSII delivery, 2GHz, 2023 Q3**
  - **Estimated** SPECint 2006 19.10@2GHz**
- 3rd generation: Kunminghu (KMH)
  - RV64GCBKHV, quad-core, superscalar OoO
  - **Advanced-node, 3GHz, 1.5x IPC of NH**
  - **Close collaboration with industrial partners**

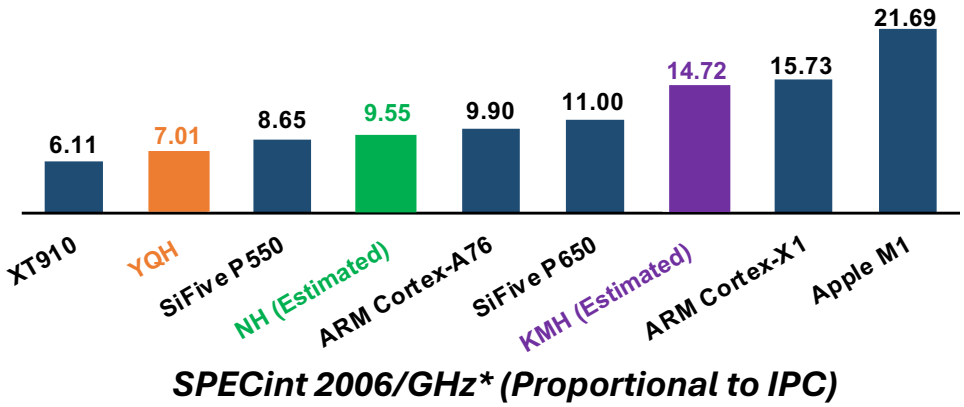

SPECint 2006/GHz (Proportional to IPC)*

* Source: XT910@ISCA'20, SiFive, AnandTech
** Updated January 5, 2023

XiangShan Gen 3: Kunminghu

- **Target ARM Neoverse N2**
  - SPECCPU2006: 45@3GHz (15/GHz)
  - Vector/Hypervisor extension supported
- **A Joint Dev Team** (coordinated by BOSC)

Kunming Lake in Beijing

Highlights in XiangShan Gen 3 Kunminghu

- **Functional Enhancement**
  - Support RISC-V **Vector/Hypervisor** extension
  - Support **RVA23 profile**
  - Support interconnection based on **CHI protocol**
- **Performance Exploration**
  - Performance boost in frontend, backend, load-stor
  - Performance model calibrated with RTL
  - Workflow: **DSE on perf model => Impl. & fine tunning on RTL**
- **Functional Verification**
  - **Hierarchical verification flow** spanning system/integration/unit level + FPGA prototyping
  - Industrial-grade verification process
- **Physical Design**
  - Experienced physical design team
  - Simultaneous iteration of RTL coding based on timing evaluation

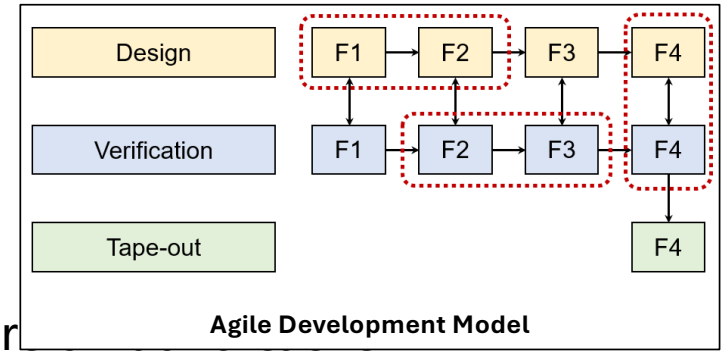

Performance Evaluation of Gen 3 Kunminghu

- Method: SPEC CPU checkpoints selected by Simpoint
  - Base: GCC 12 –O3, RV64GCB, jemalloc
  - 1MB L2 and 16MB L3
  - Simulated@3GHz with DRAMsim3 DDR4-3200

SPECint 2006 est.@ 3GHz		SPECfp 2006 est.@ 3GHz	
400.perlbench	35.88	410.bwaves	66.89
401.bzip2	25.50	416.gamess	40.89
403.gcc	46.72	433.milc	45.25
429.mcf	58.13	434.zeusmp	52.10
445.gobmk	30.26	435.gromacs	33.65
456.hmmer	41.60	436.cactusADM	46.16
458.sjeng	30.53	437.leslie3d	46.01
462.libquantum	122.50	444.namd	28.88
464.h264ref	56.57	447.dealII	73.43
471.omnetpp	39.37	450.soplex	51.99
473.astar	29.23	453.povray	53.44
483.xalancbmk	72.03	454.Calculix	16.38
GEOMEAN	**44.15**	459.GemsFDTD	37.18
		465.tonto	36.67
		470.lbm	91.24
		481.wrf	40.62
		482.sphinx3	48.57
		GEOMEAN	**44.60**

* Updated in March 2025

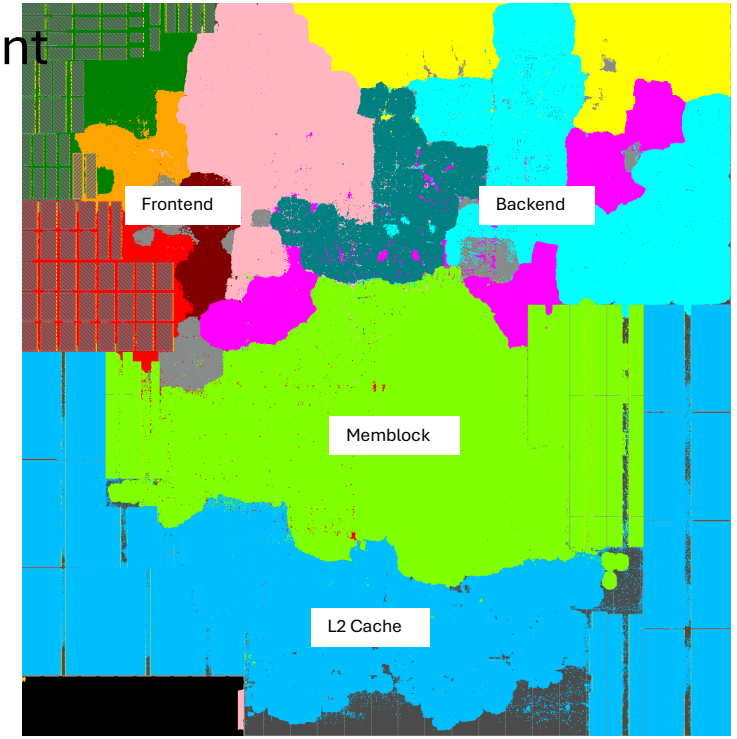

Floorplan of KMH V2R2 (single core)

XiangShan: Open-Source High Performance Processors

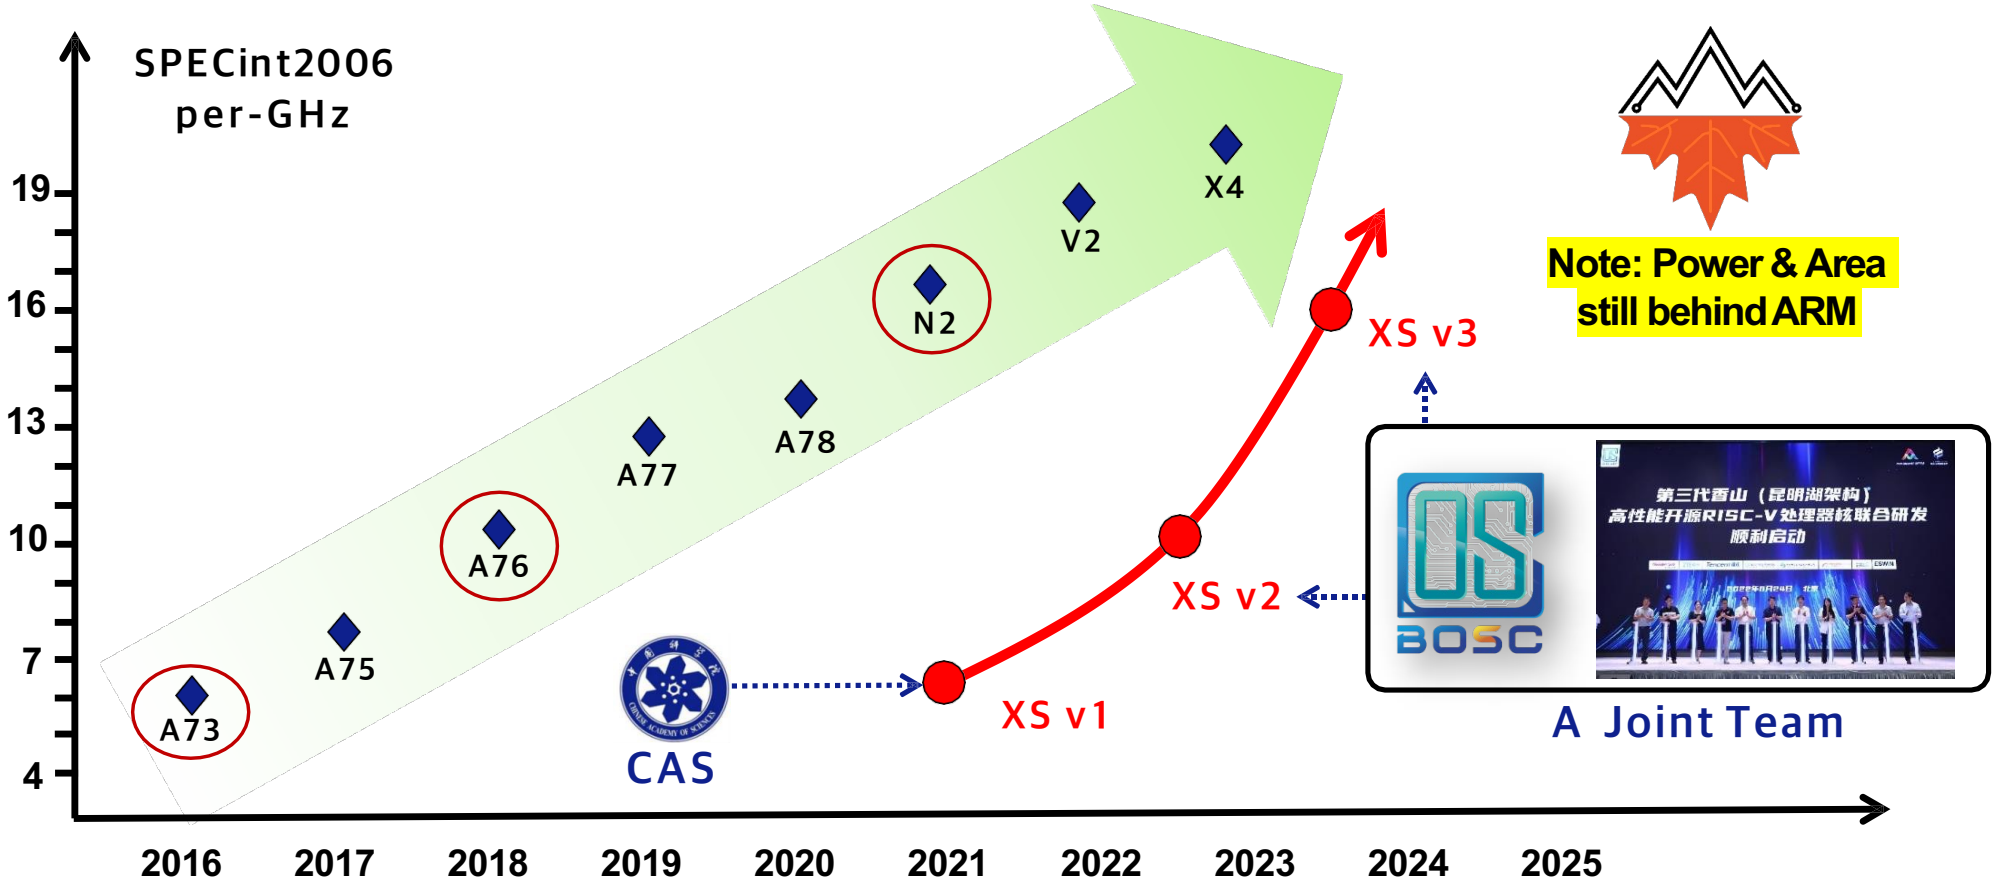

Compiler Support for Nanhu and Kunminghu

- LLVM & GCC support XiangShan optimizations (BOSC, ISCAS, ICT)

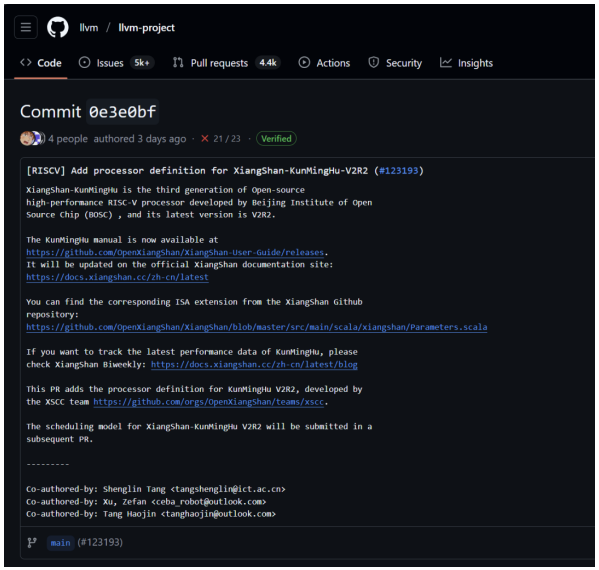

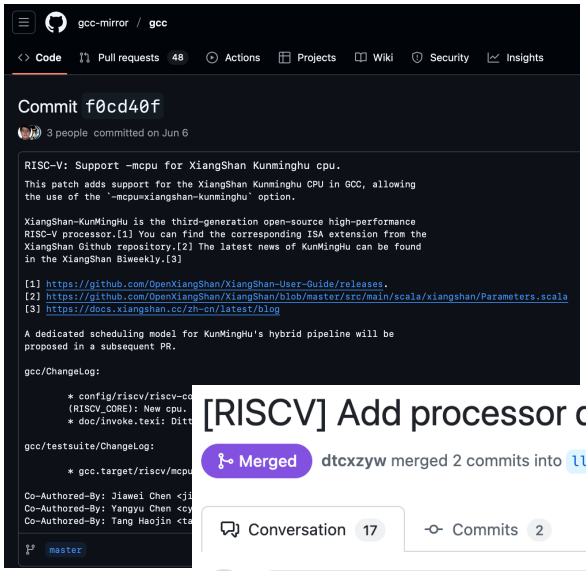

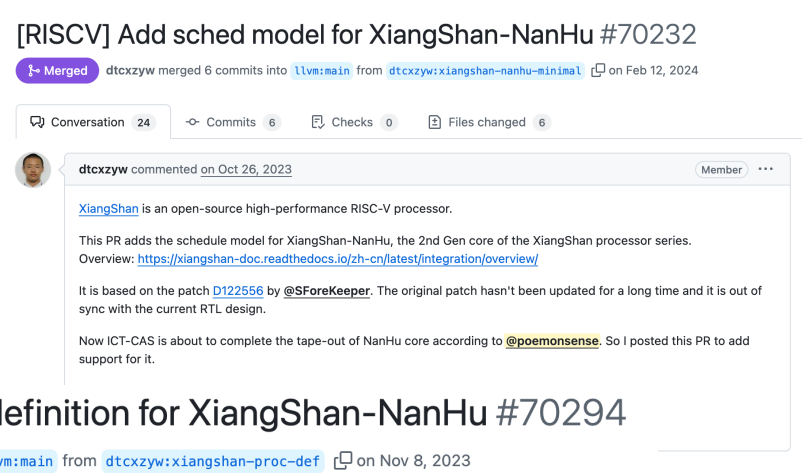

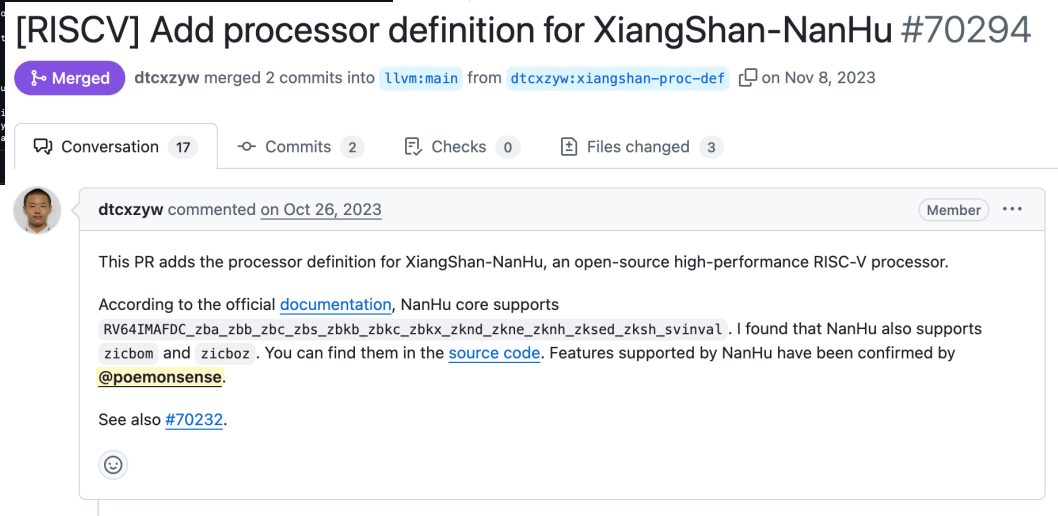

OpenXiangShan: Empowering Architecture Research

- **An open-source, continuously developing research platform**

Your research field	XiangShan provides …
Microarchitecture	**Performance**: An industrial-competitive, high-performance superscalar OoO microarchitecture
	Functionality: RVA23-compatible RISC-V design
	Development: mature, user-friendly design flows
	Tapeouts by the XiangShan team and leading industrial partners with real-world deployment
Chip development tool	Realistic and challenging **research problems**

An Effective Infrastructure for Research

- ***Topic: Computer Architecture***
  - XiangShan: a realistic out-of-order RISC-V implementation with industry-competitive performance and an active open-source community
  - MinJie provides the toolchains
- *Microarchitecture, accelerators, novel architectures, profiling, systems, benchmarking, security, compilers, …*

midgard EPFL

Imprecise Store Exceptions, ISCA'23
Single Address Space Faas with Jord, ISCA'25

- ***Topic: Agile Chip Development***
  - XiangShan is a progressive, configurable, complicated, challenging benchmark
  - MinJie provides a good startpoint
- *HDLs, verification, performance, power, area, prototyping, DFT, synthesis, placement, routing, ECO, …*

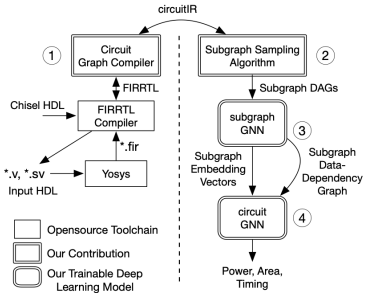

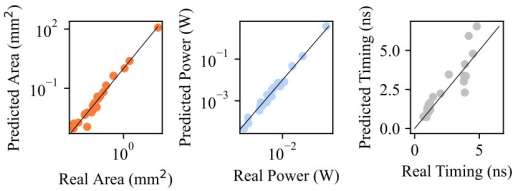

SNS v2, MICRO'23 (Duke University)

Outline

- The Era of Open-Source Chips
- Open-Source Industrial-Competitive RISC-V Chips
- Open-Source Chip Design Tools
- Open-Source Development & Business Models
- Conclusions

Open-Source Infrastructures (Tools) for Chips

XiangShan Chips

- Complex ISAs
- Microarchitecture
- PPA improvements
- ……
- Lower thresholds
- Agile development
- Parameterization
- ……

61

Chisel: agile hardware description language

- 2018: quantitatively comparing Chisel and Verilog, reducing code size by 80%
- 2020: completed the 1st generation of XiangShan, booting Linux within 3 months
- 2022: 67,000 lines of design code and 31,000 lines of verification code
- 2024: 214,000 lines of code in all code repositories

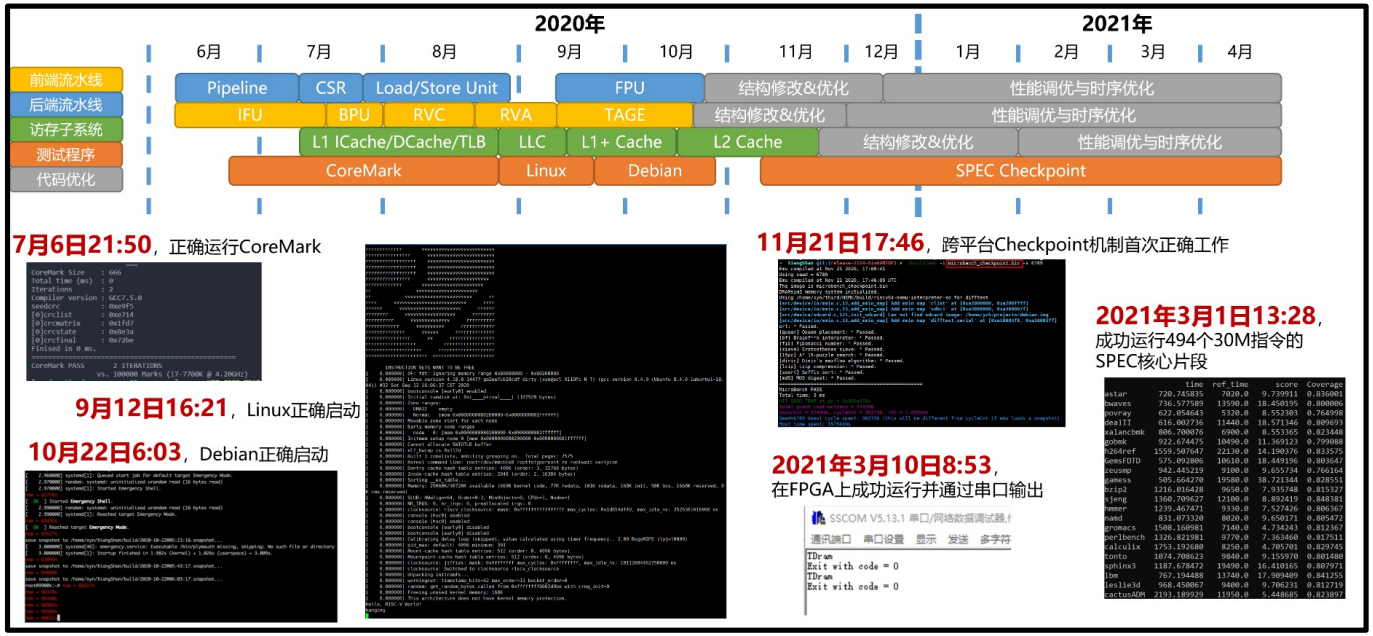

Better design flexibilities

Kunminghu has **212** configurable parameters, and its L2/L3 cache has **65** configurable parameters

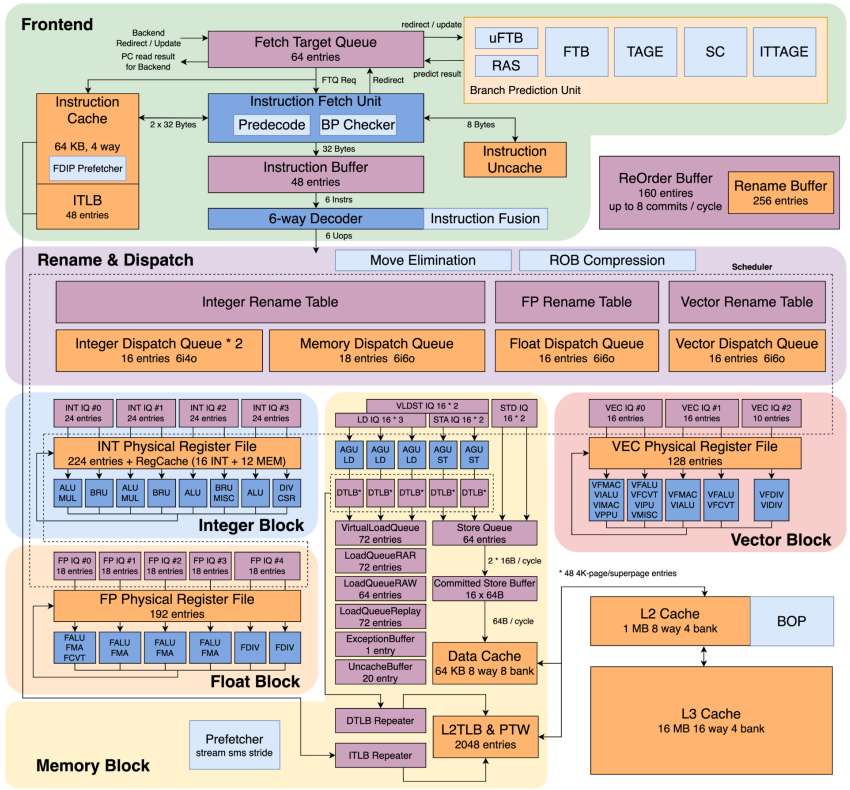

ARM A76 has **8** configurable parameters, and the DSU (L3) has **25** configurable parameters.

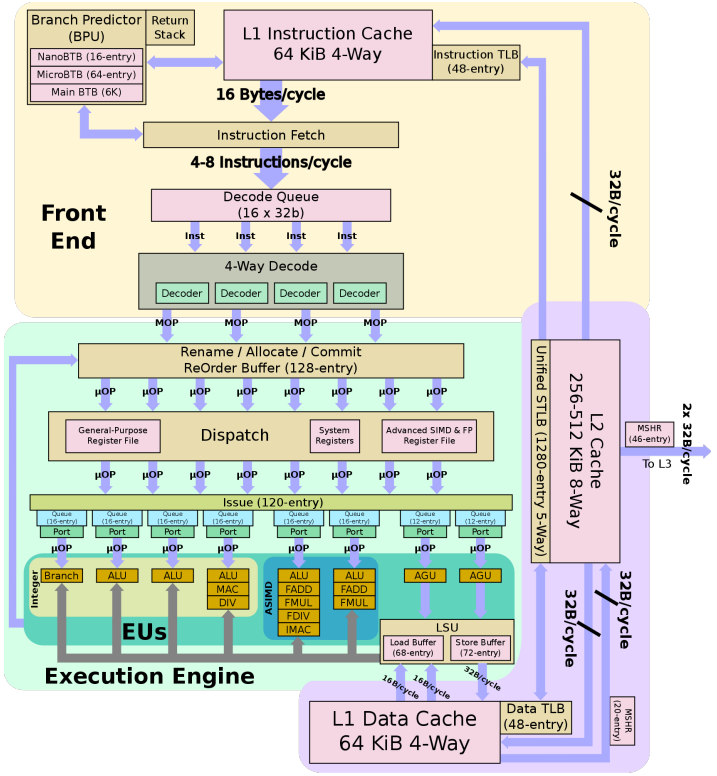

63

DiffTest: a co-simulation CPU verification framework

- **Co-simulation workflow**
  - Instructions commit/other states update
  - The simulator executes the same instructions
  - Compare the architectural states
  - Abort or continue

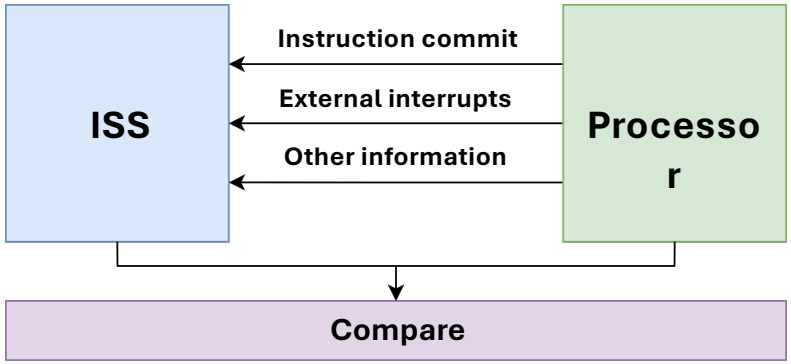

Basic architecture

- **Verification infrastructures for CPUs**
  - APIs for HDLs such as Chisel/Verilog
  - RTL simulators such as Verilator, VCS, Palladium
  - RISC-V ISS such as Spike, NEMU

```
while (1) {
    icnt = cpu_step();
    ref_step(icnt);
    r1s = cpu_getregs();
    r2s = ref_getregs();
    if (r1s != r2s) { abort(); }
}
```

Online checking

https://github.com/OpenXiangShan/difftest

Standard interfaces for RISC-V CPU verification

- Key idea: information probes support flexible combination for different scenarios

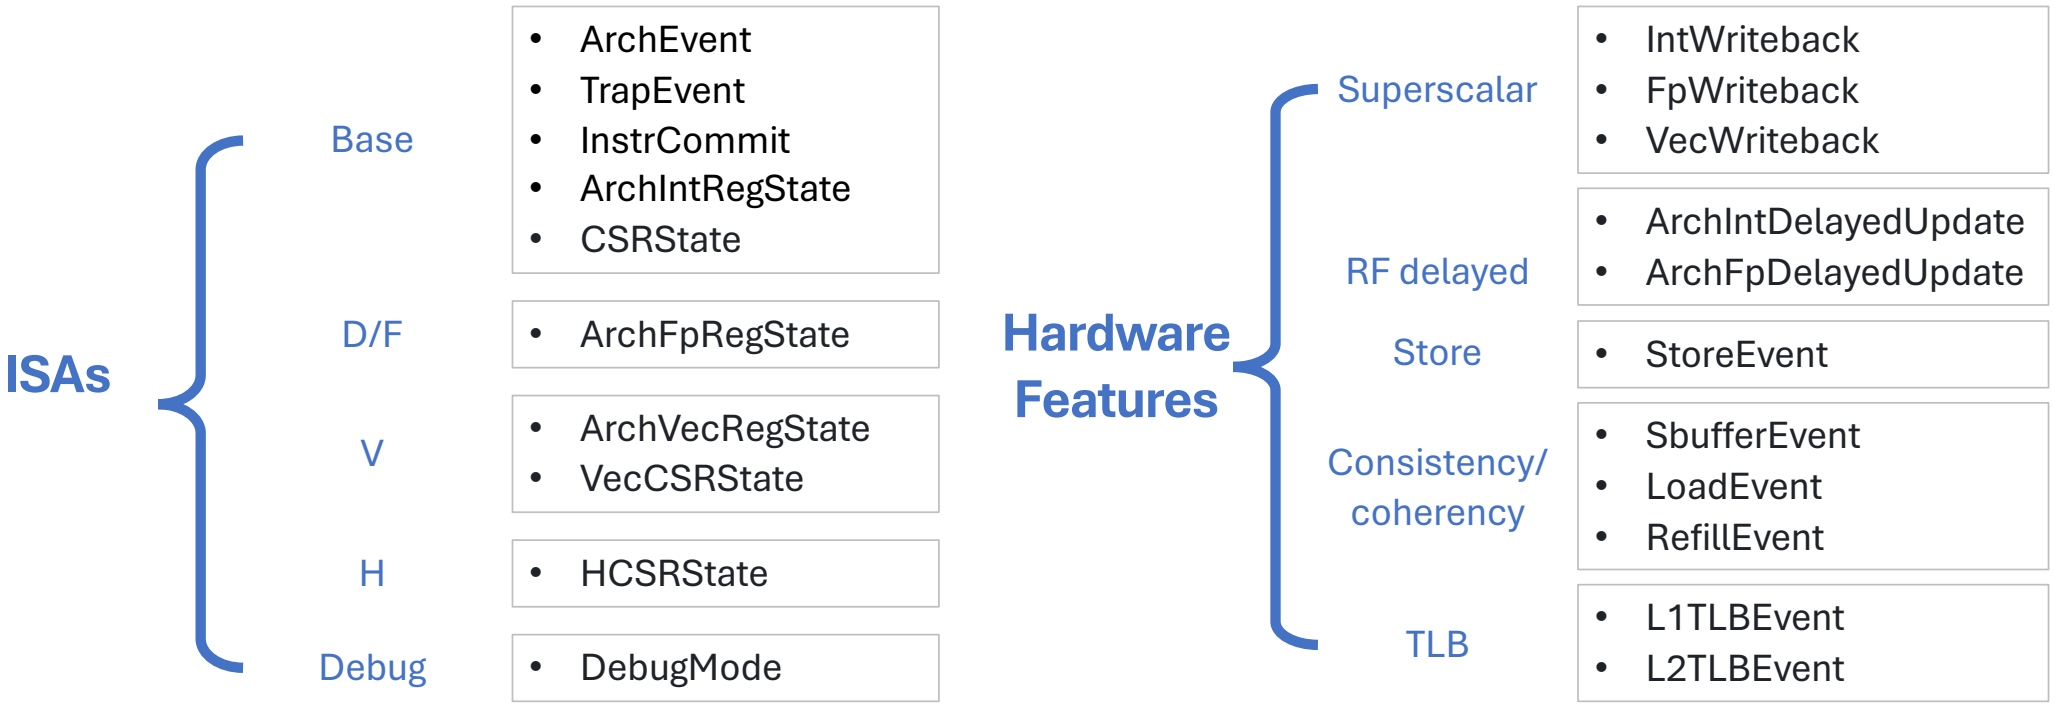

Accelerated co-simulation on Emulator/FPGA

- **DiffTest now supports hardware-accelerated co-simulation**

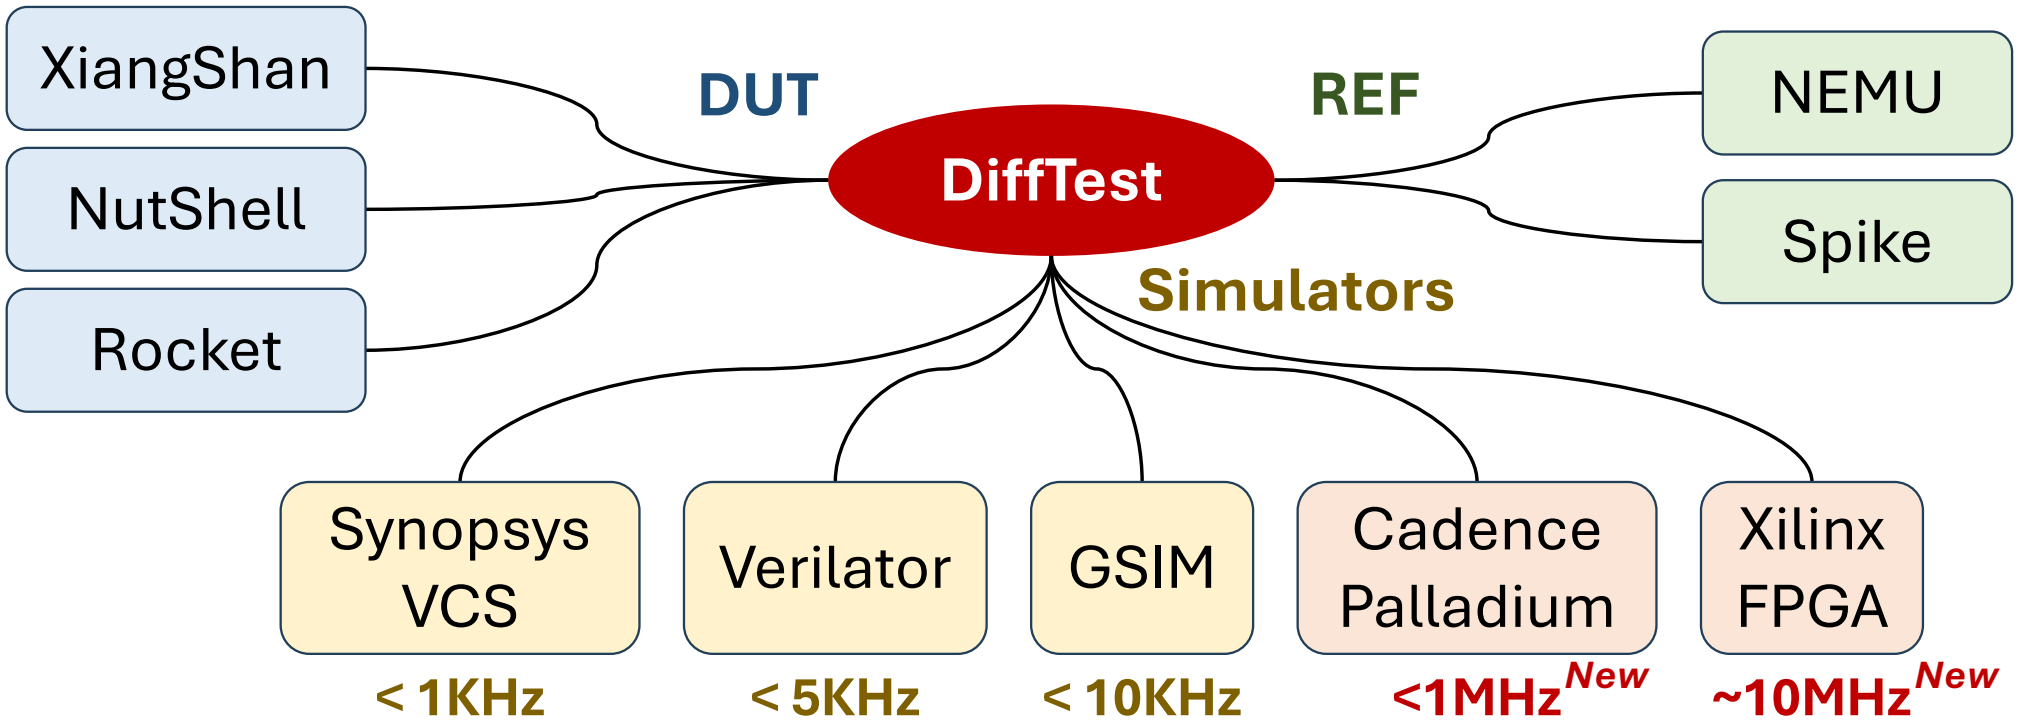

DiffTest-H: semantic-aware co-simulation acceleration

- Optimizes communication overhead for verification data packets
  - **Batch**: Reduces communication frequency
  - **Squash**: Reduces communication data volume
  - **Replay**: Maintains debugging granularity

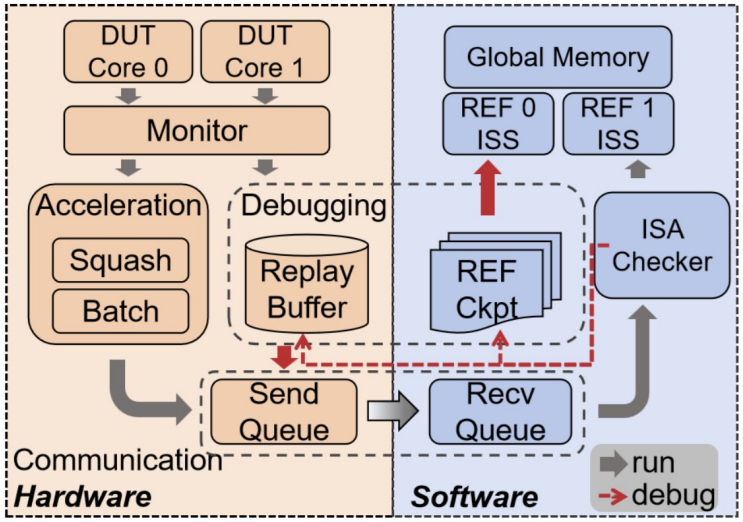

Figure: workflow

Category	Types	Representative Examples
Control Flow	5	Exceptions and interrupts, Instruction commits, Traps, ...
Register Updates	9	CSRs, General-purpose registers, Floating-point registers, ...
Memory Access	3	Load/store operations, Atomic memory operations, ...
Memory Hierarchy	6	Cache refill operations, L1/L2 TLB operations, ...
RISC-V Extensions	9	Vector/Hypervisor CSRs, Vector registers, ...

Table: Packets

DiffTest-H: Hardware-Accelerated Co-Simulation Verification

- **13.8 MHz** on FPGA, **instruction-level** debugging
- Deployed on **XiangShan**, with **151 bugs uncovered**
- **Open-Source** at *github/OpenXiangShan/difftest*

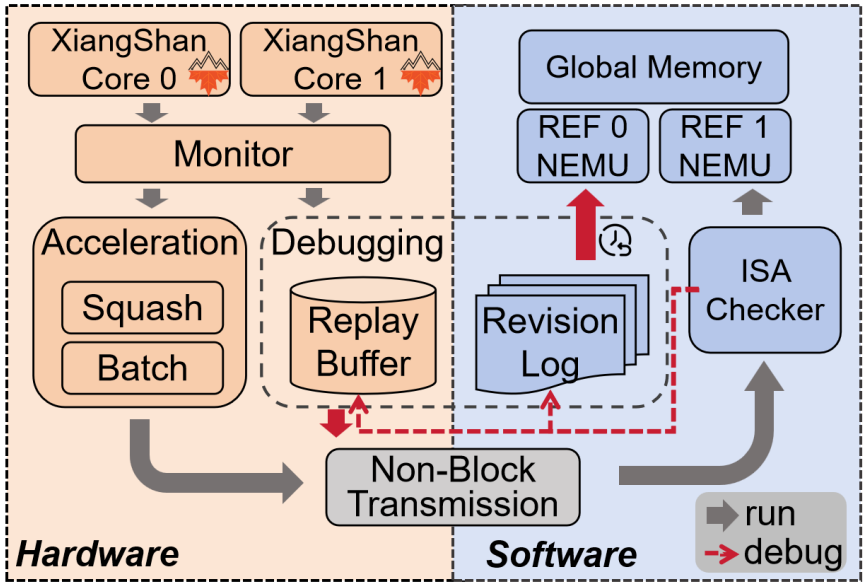

Verilua: an easy-to-use unit-testing framework

New option for hardware verification

	Verilua	Cocotb	Fault	PyMTL/PyHGL/Chisel	UVM/SV
核心技术	VPIML	VPI	IR + 多后端	单一语言垂直整合	--
执行模型	**在线仿真+离线分析**	在线仿真	--（元编程）	在线仿真	在线仿真
验证资产复用	**跨语言、跨场景**	**否**	跨语言	**否**	--
HDL	--	--	--	Y	Y
HVL	Y	Y	Y	Y	Y
HSE	**Y**	--	--	--	--
WAL	**Y**	--	--	--	--
学习门槛	**低**	中	中	**高**	**高**

https://github.com/cyril0124/verilua

mux.sv

```
module mux(
    output wire out,
    input wire sel,
    input wire a,
    input wire b
);

assign out = sel ? b : a;

endmodule
```

main.lua

```
fork {
    main_task = function()
        sim.bypass_initial()

        dut.a:set(0)
        dut.b:set(1)

        dut.sel:set(1)
        await_time(1)
        dut.out:expect(1)

        dut.sel:set(0)
        await_time(1)
        dut.out:expect(0)

        print("TEST PASS!")

        sim.finish()
    end
}
```

Test result:

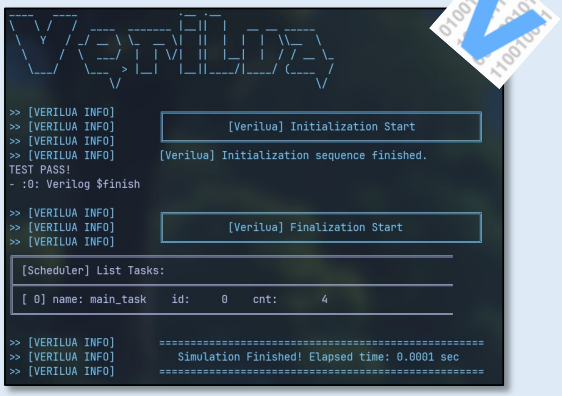

XS-GEM5: Calibrated Gem5 Simulator

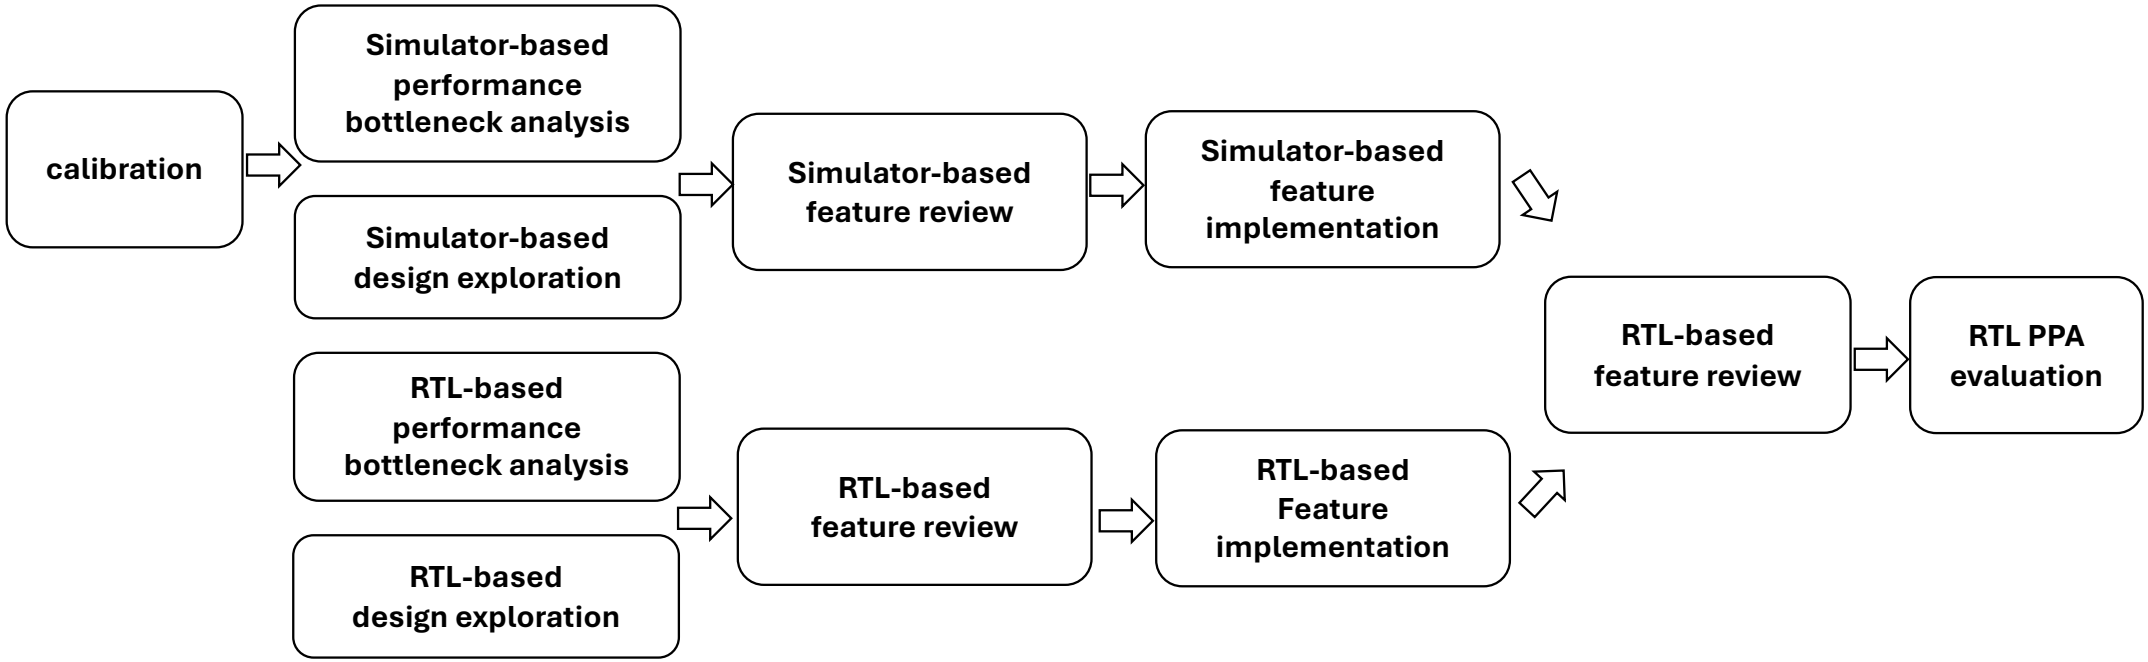

https://github.com/OpenXiangShan/GEM5

XS-GEM5: Calibrated Gem5 Simulator

- Based on GEM5, calibration was completed for XiangShan
- The overall error with RTL (KunminghuV2) in SEPC06 is less than 3%
  - The error of a single benchmark is less than 5%
  - Currently being used for architecture exploration of the KunminghuV3

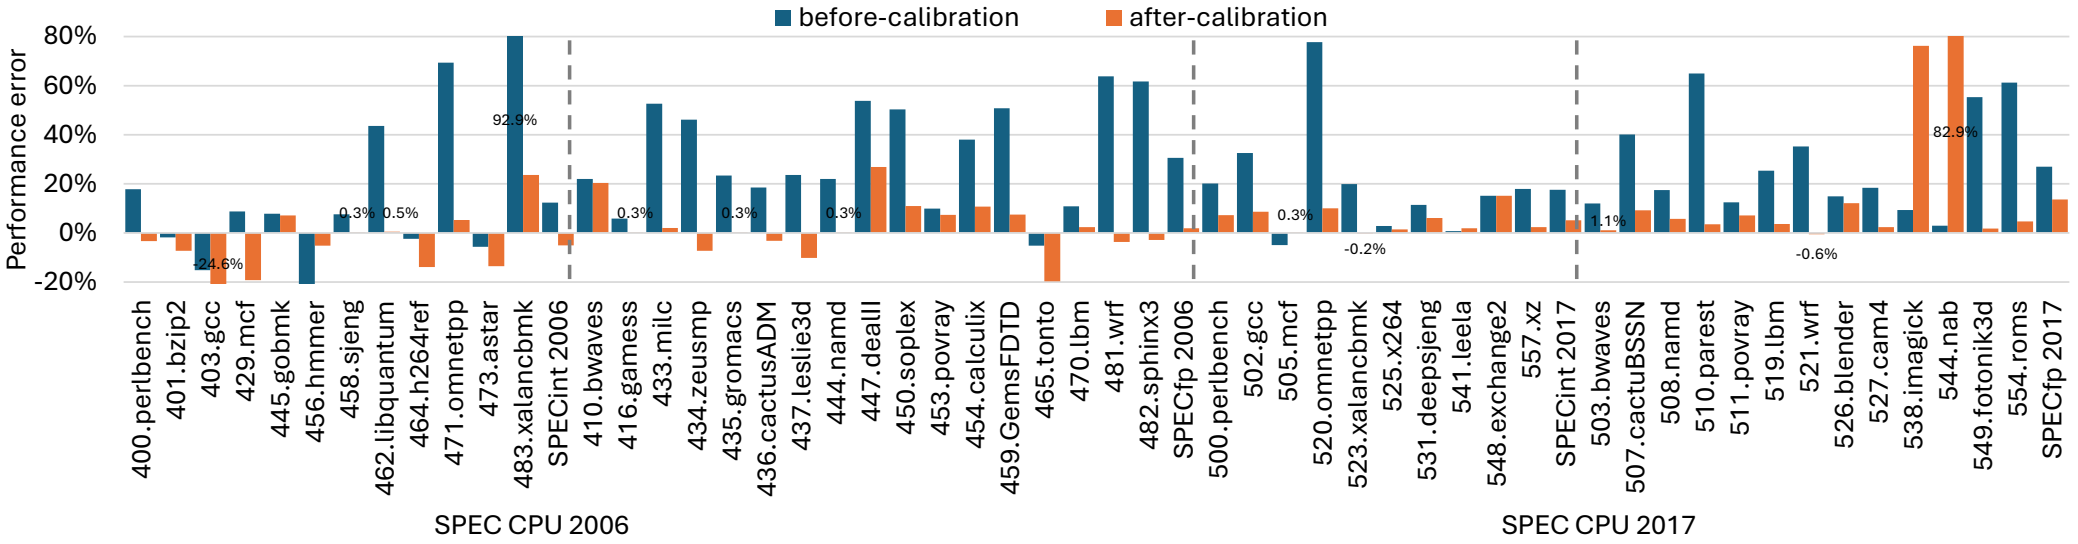

TraceRTL: Trace-driven RTL CPU Model

- Applying mature architecture exploration techniques *from simulators to RTL*
- **Insight**: Architecture exploration **only requires essential performance** components; complete functional correctness is *unnecessary*
  - Performance evaluation is performed before ensuring CPU functional correctness

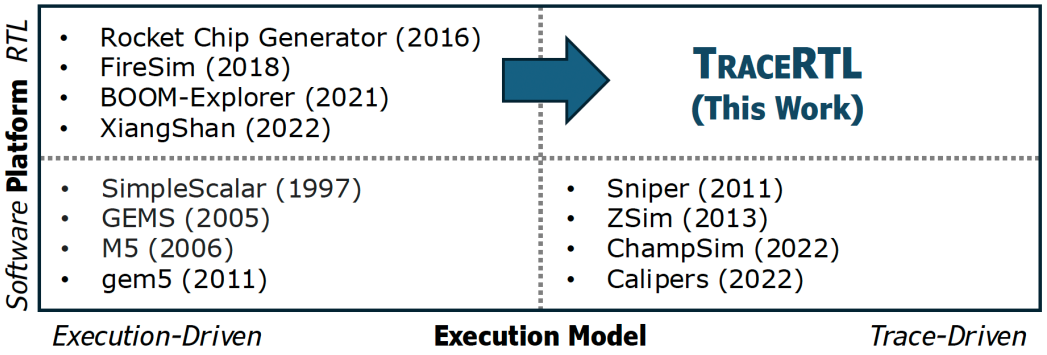

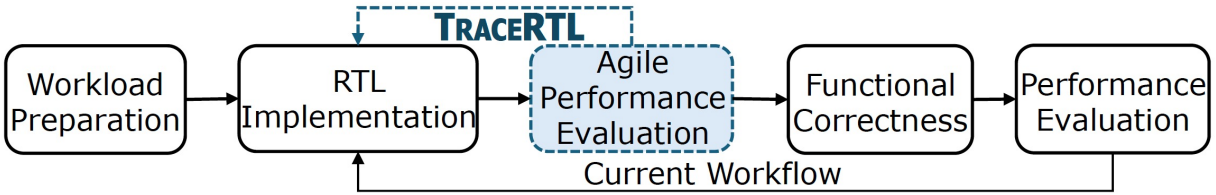

TraceRTL: Trace-driven RTL CPU Model

- Automated, low-intrusion modification of **a trace-driven XiangShan**
  - **Eliminating functional dependencies**: Chisel, circuitry, architecture, functional abstraction, performance sensitivity, etc.
  - **Eliminating performance errors**: The impact of lost information in the trace, including addresses, data, etc.

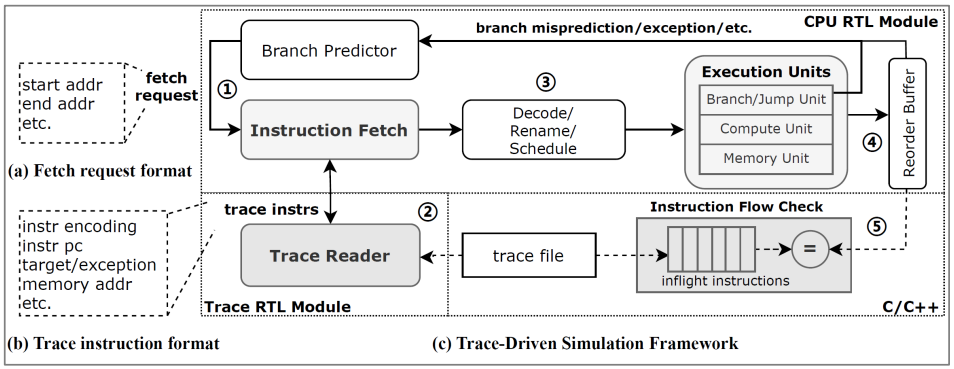

Trace-driven CPU from XiangShan

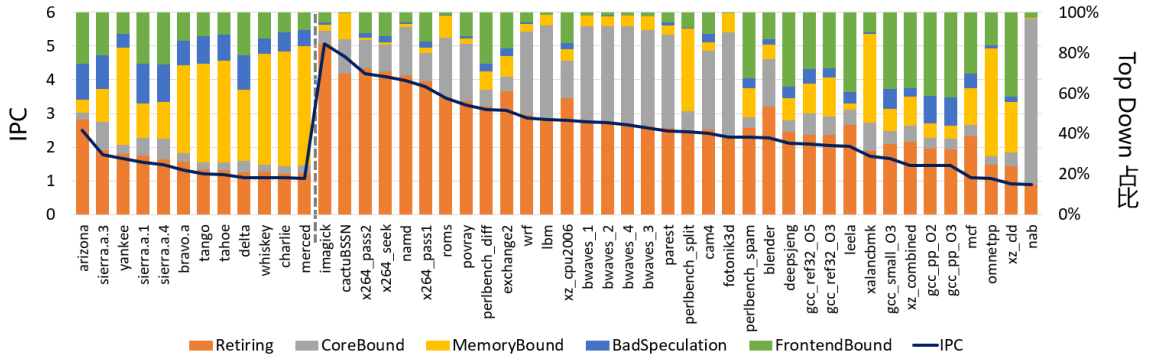

Top-Down differences between Google Datacenter Workloads and SPEC CPU

FPGA accelerated design space exploration

- **Design Space Exploration (DSE)**: Searching the CPU design parameter space to find the PPA balance point
- Due to several *limitations*, all current work is performed on a simulator.
  - **Speed Limitation**: Single simulation speed should be fast enough.
  - **Parameter Limitation**: Simulators offer flexible parameter tuning, while RTL is challenging

FastDSE: FPGA + logic/physical parameter decoupling

- FastDSE: FPGA Acceleration + Logic/Physical Parameter Decoupling
  - Stable operation at 50MHz
  - accelerating the DSE process by 68.7 times.

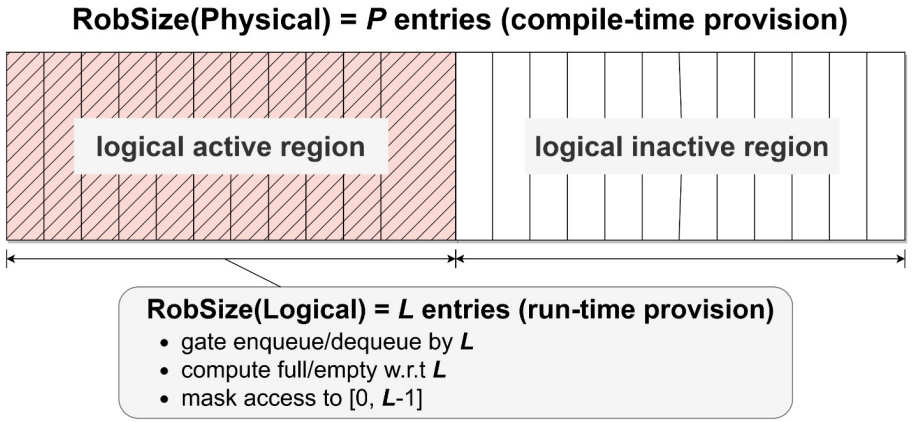

Logic/Physical Parameter Decoupling

- One physical param
- Multiple logical parameters
- Adjust dynamically

MinJie: Open & Agile Development Toolchain

- It supports a new model of collaborative chip development based on open source, continuously building a team of over 600 people (the largest in the world).
  - Selected as one of the 12 IEEE MICRO Top Picks

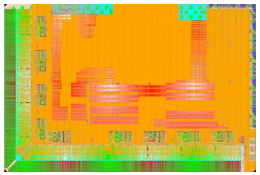

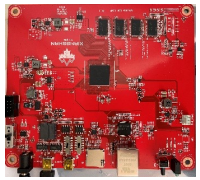

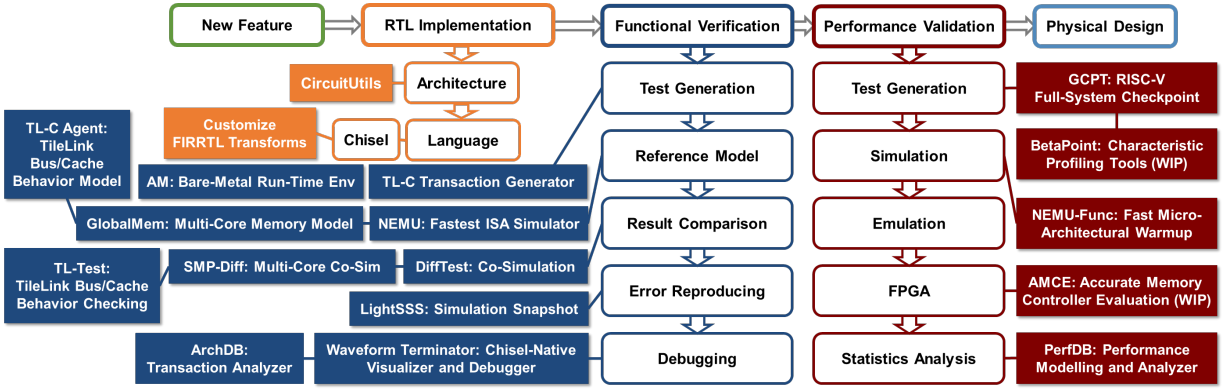

Outline

- The Era of Open-Source Chips
- Open-Source Industrial-Competitive RISC-V Chips
- Open-Source Chip Design Tools
- Open-Source Development & Business Models
- Conclusions

Development Model for Open-Source Chips

- With the support of the Beijing Municipal Government and the Chinese Academy of Sciences, **16 companies** jointly launched the **Beijing Institute of Open Source Chip (BOSC)** to accelerate the industrial development of Xiangshan.
- BOSC has assembled a team of **>500 people**, becoming one of the largest RISC-V CPU core R&D teams in the world.

16 Founding Members

Dec. 6, 2021

Tiered Rolling Open-Source Model

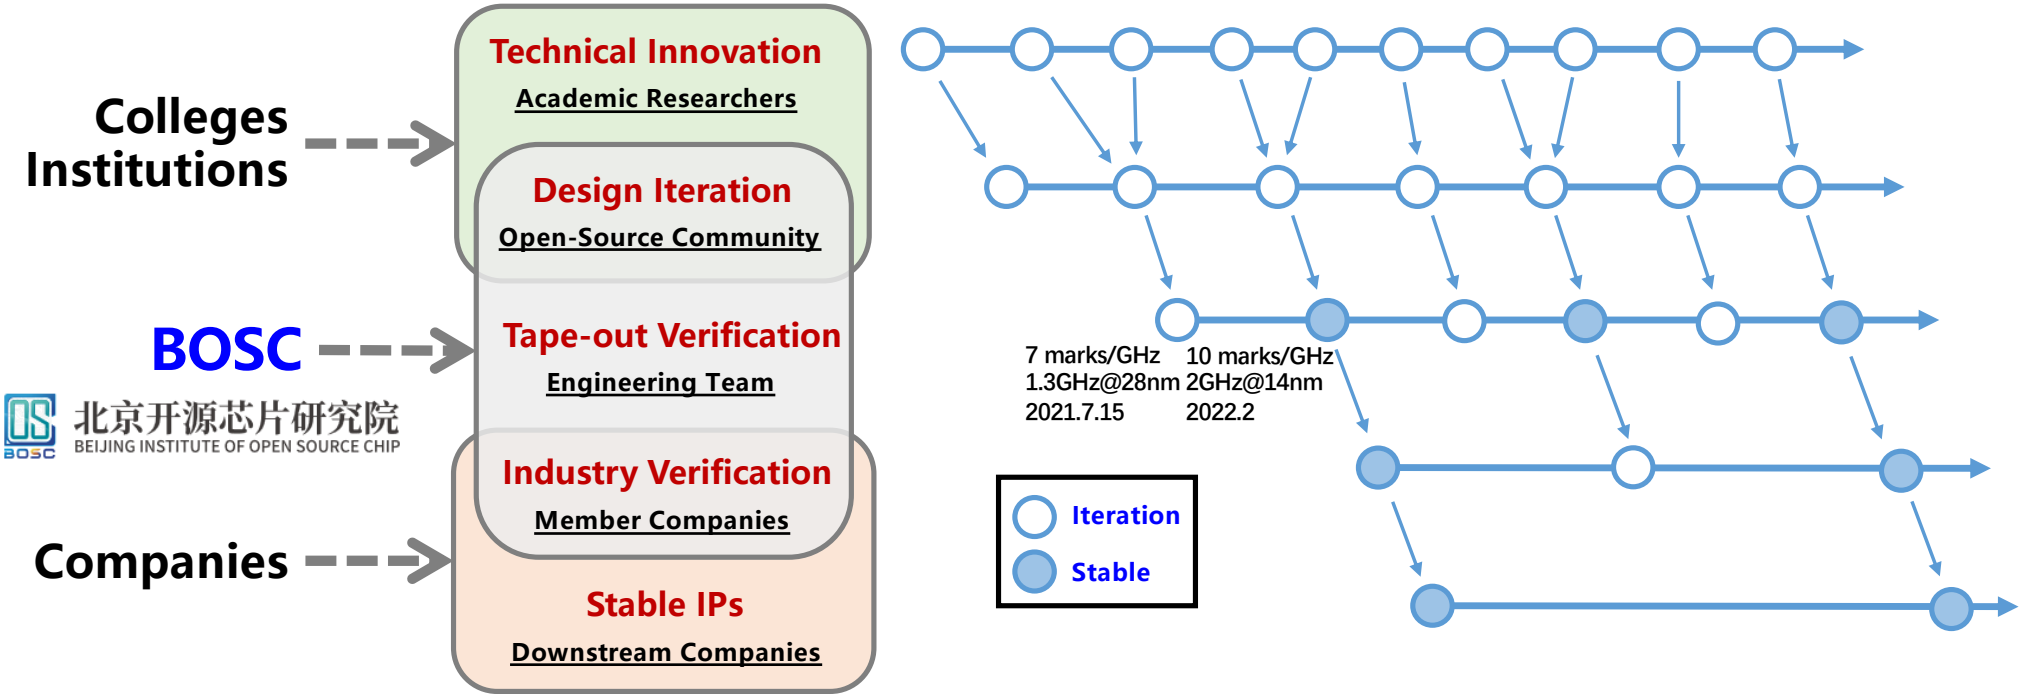

Chip products based on XiangShan

Kunminghu V2

User	Product	Tapeout Time
A	**8-core video codec**	**2025/9**
B	**64-core server**	**2025/9**
C	**128-core server**	**2025/10月**
D	128-core server	2026/3
B	8-core client SoC	2026/3
E	16-core	2026/3
F	64-core server	2026/3
G	4-core client SoC	2026/6
H	128-core server	2026/12

Nanhu

User	Product	Tapeout Time
I	**GPGPU**	**2024/production**
J	**GPGPU**	**2024/production(>10k)**
K	4-core FPGA ctrl.	2025/12
L	4-core sec. ctrl.	2025/12
M	4-core	2025/12
N	4-core router ctrl.	2026/3

XiangShan Nanhu Chips

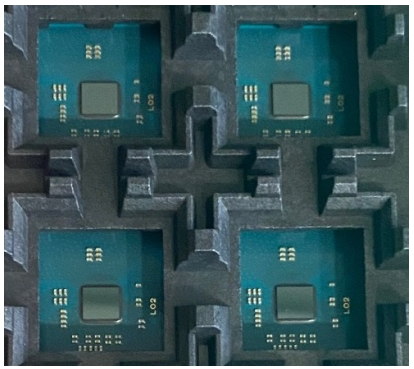

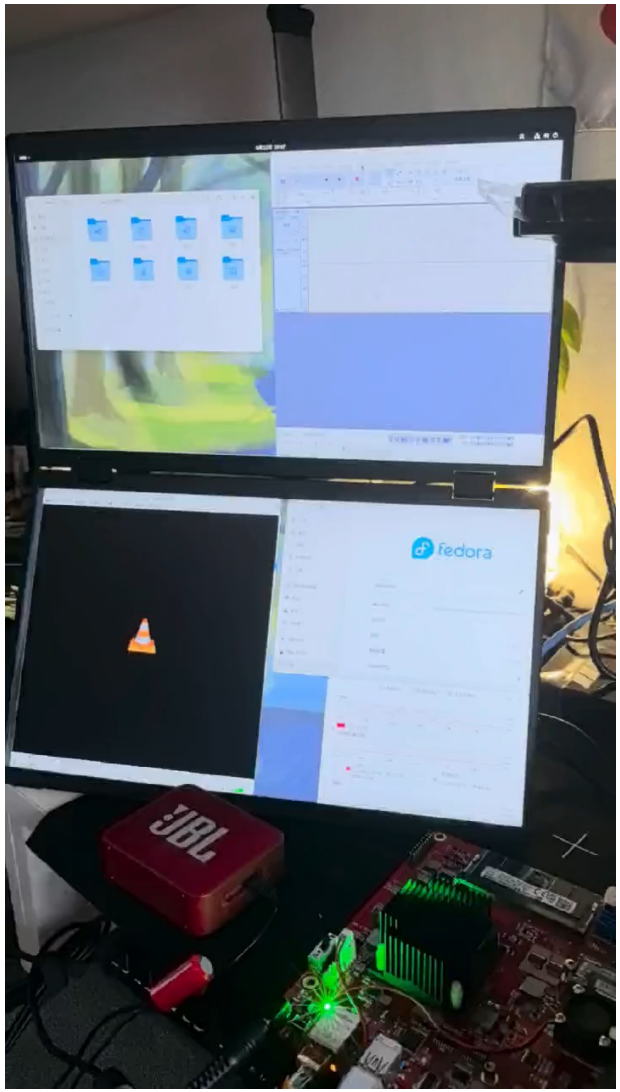

FPGA Prototype in Only Two Weeks

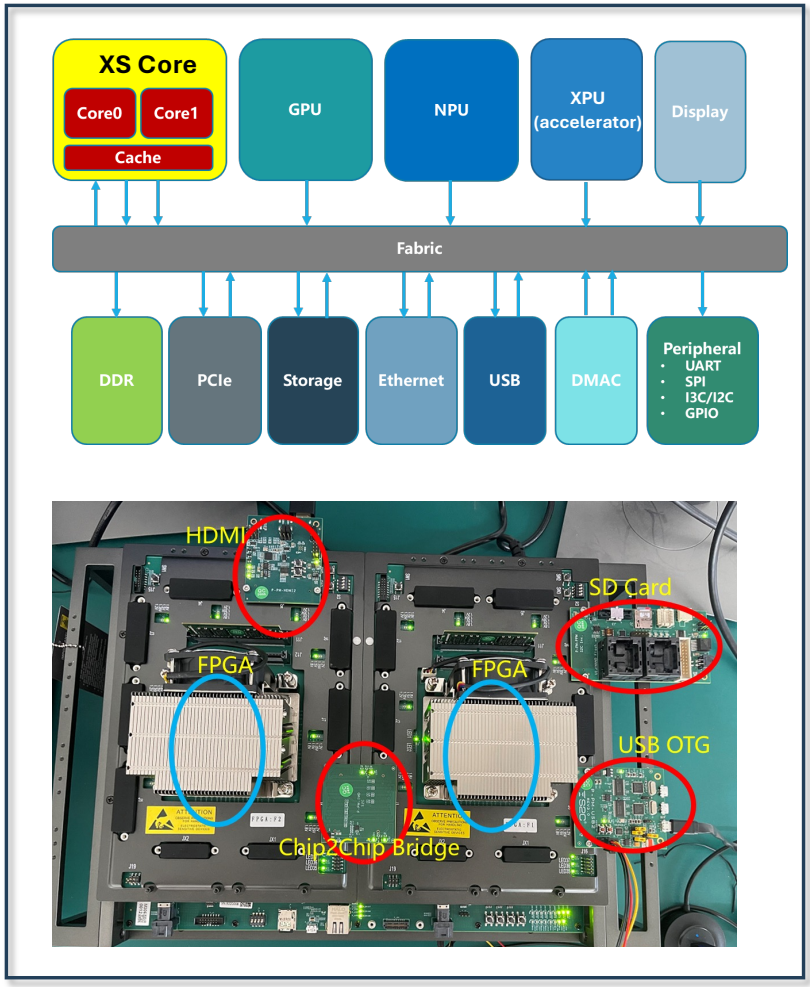

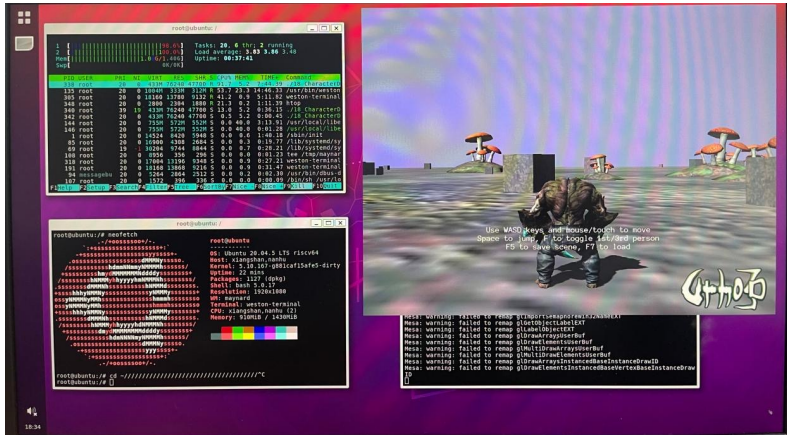

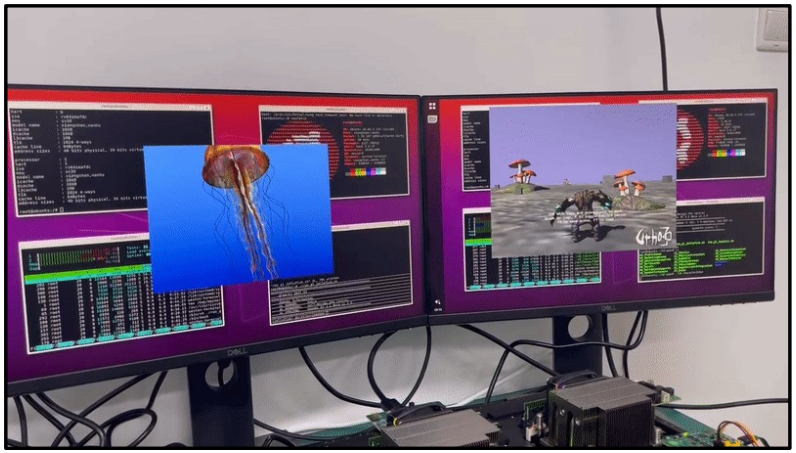

Source: Xinchen Technology

Real Chip of XiangShan Nanhu

- **Test chip was back in October 2023**
  - Successfully brought up Linux and working with PCIe device (GPU, Ethernet, USB..)

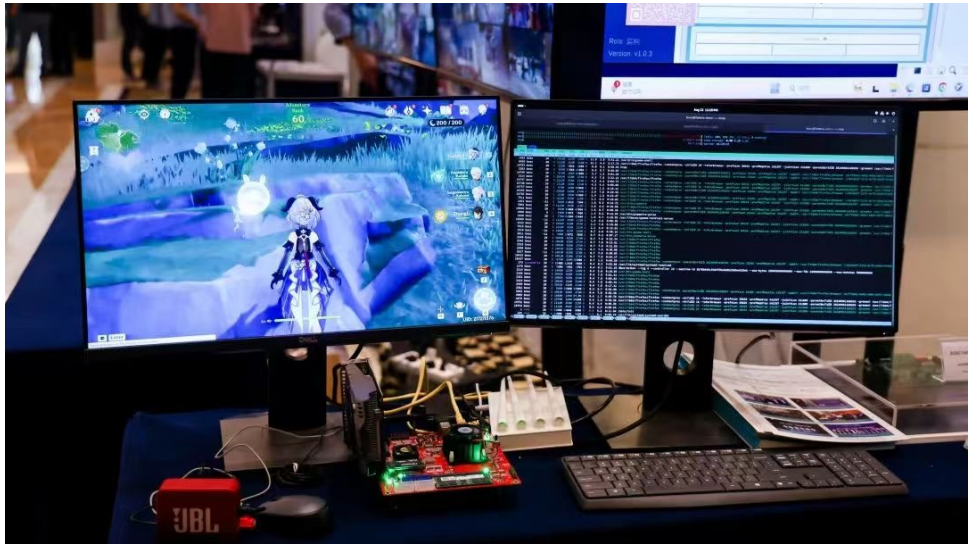

Nanhu test board @ RVSC'24

Running Desktop Linux & Cloud Game*

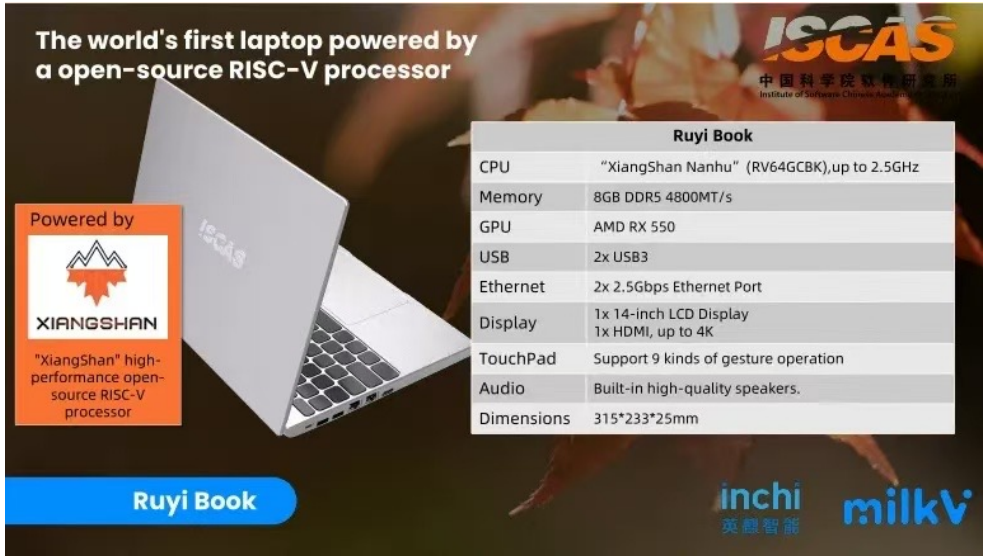

Ruyi XiangShan Book

Design by ISCAS, Inchi, MilkV

* Genshin Impact · Cloud, only ~ 1 fps, just for fun

Linux Boot on 8-Core XiangShan Kunminghu V2

- Lanxin Computing has successfully launched Linux on an 8-core SoC built on the Kunminghu V2

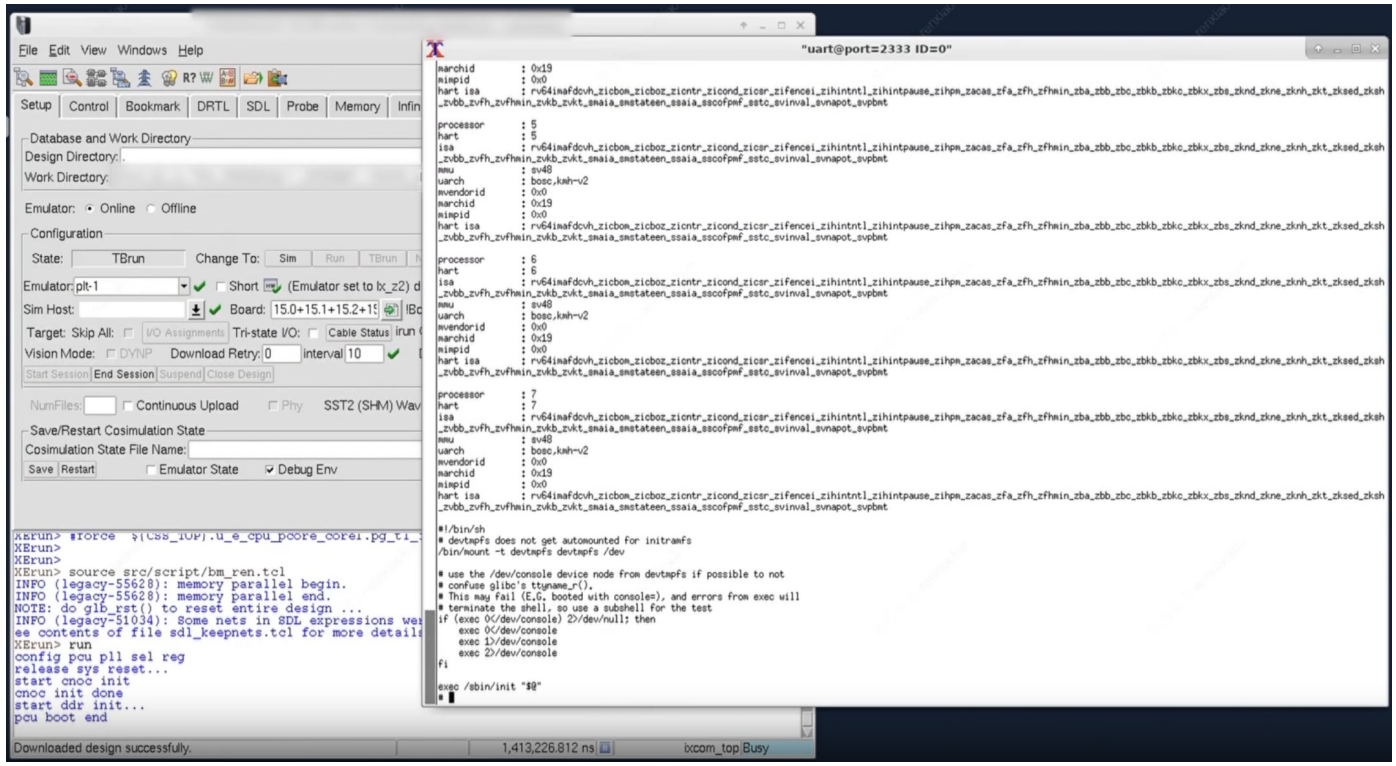

Roadmap 2025: 3 CPU Compute Systems

CPU

Nanhu V5

- SPEC06 10/GHz
- 2GHz@12nm
- Target ARM A76
- 2025/10 delivery

Kunminghu V2

- SPEC06 15/GHz
- 3GHz@7nm
- 64-core
- 2025/4 delivery

Kunminghu V3

- SPEC06 22/GHz
- 3GHz@7nm
- 128-core
- 2025/12 delivery

NoC

Zhujiang V1

- Ring Bus
- Max 16-core
- Cache Coherency
- 2025/10 delivery

Wenyuhe V2

- Target CMN-700
- 128-core
- Chiplet support
- 2025/12 delivery

Advantages of open-source collaborative development

- Traditionally, **test cases are *valuable private assets*** of chip companies.
- With open source, it can be directly deployed within different companies.
  - A total of 1,467 bugs have been fixed, of which **492 (33.4%) were submitted by companies.**

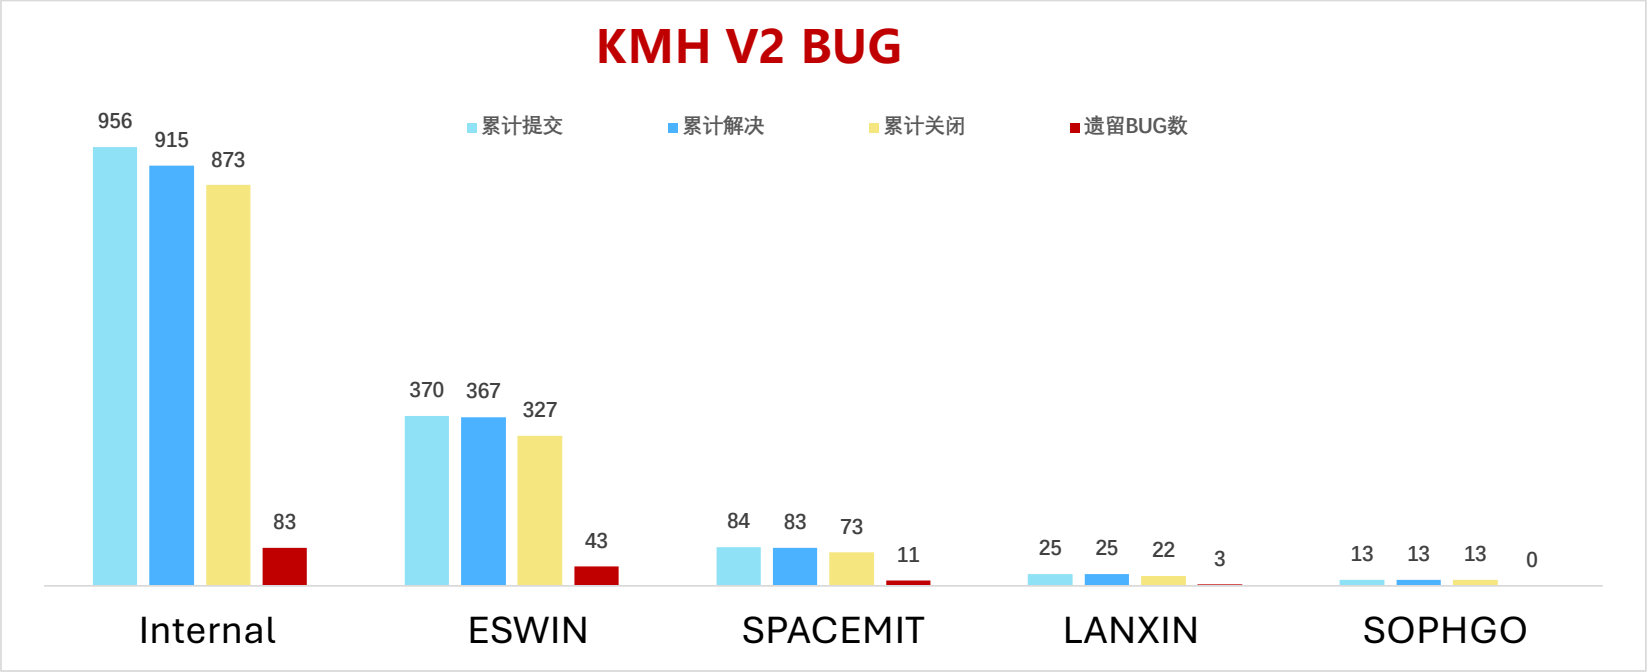

85

UnityChip Verification: Crowd-Sourcing Verification

- Chip verification using an open-source crowdsourcing model
- Bringing together software and hardware engineers
- A verification campaign launched on XiangShan
- Let 10,000 people participate in verification!
- https://open-verify.cc/en/

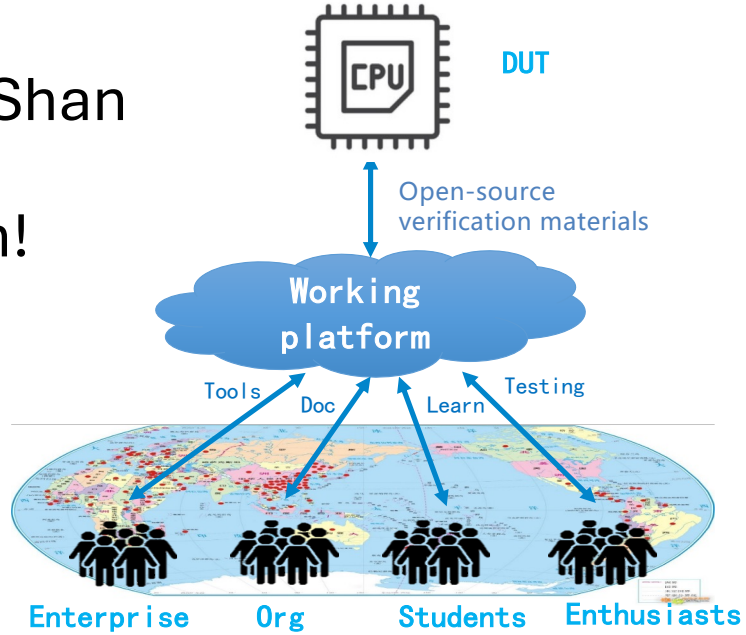

UnityChip: A Croudsourcing Platform for Chip Verification

- **Features: Involve software engineers**
  - Enable crowdsourcing
  - Support multi-languages
  - Be compatible with UVM
- **Effectiveness**
  - 5 **undergraduate** students found **10 bugs** **in 2 months**
- **Easy-to-Use**
  - Students not familiar with Linux and Python
  - Learn to use tools in **5 days**
  - Start adding test cases after another **10 days**
- **UnityChip Competition for XiangShan**
  - Task: Unit Test for IFU, BPU, ITLB, etc.
  - Scan the QR code for more information

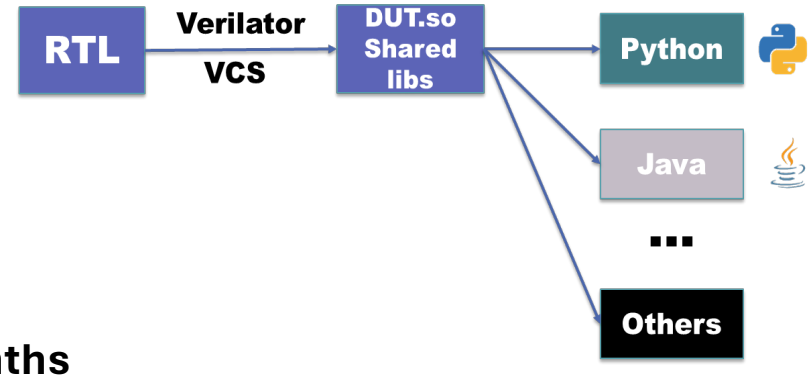

GitHub Link

Open Problems on Chip Development Infrastructures

- Sharing ***the real-world infrastructure challenges*** faced by the XiangShan project

- **Hardware Descriptions**
- **Functional Verification**
- **Performance Improvements**

Problems

- **Architecture**
- **Software Engineering**
- **EDA**
- **PL/Sys/…**

Directions

- **Discussions**
- **Interns**
- **Full-time**
- **Research/Industrial Collaborations**

Participation

Outline

- The Era of Open-Source Chips
- Open-Source Industrial-Competitive RISC-V Chips
- Open-Source Chip Design Tools
- Open-Source Development & Business Models
- Conclusions

XiangShan achieves L2.5

- **L1: OPEN** ISA
- **L2: OPEN** Design/Implementation
- **L3: OPEN** Framework/Tools

3 Open Framework/Tools

microarch. → coding → EDA Tools

ISA Spec.

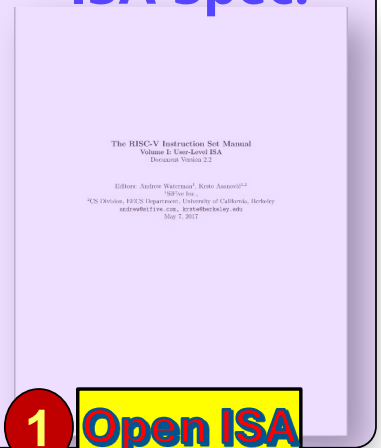

1 Open ISA

Docs

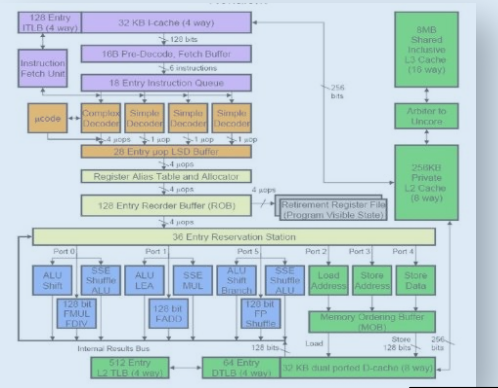

RTL

```
component DebugCoreTop is
  port (
    -- Trigger and Data
    cu_Clk      : in    std_logic_vector(2 downto 0) := (others => '0');
    cu0_Trig    : in    t_trig_0 := (others => (others => '0'));
    cu1_Trig    : in    t_trig_1 := (others => (others => '0'));
    cu2_Trig    : in    t_trig_2 := (others => (others => '0'));
    cu0_Data    : in    t_data_0 := (others => (others => '0'));
    cu1_Data    : in    t_data_1 := (others => (others => '0'));
    cu2_Data    : in    t_data_2 := (others => (others => '0'));

    -- Downstream I2C
    SCL         : in    std_logic := '0';
    SDA         : inout std_logic := '0';

    -- Upstream
    gt_RefClk_p : in    std_logic := '0';
    gt_RefClk_n : in    std_logic := '0';
    gt_RX_p     : in    std_logic_vector(2 downto 0) := (others => '0');
    gt_RX_n     : in    std_logic_vector(2 downto 0) := (others => '0');
    gt_TX_p     : out   std_logic_vector(2 downto 0);
    gt_TX_n     : out   std_logic_vector(2 downto 0)
  );
end component;
```

2 Open Design/Implt

Layout

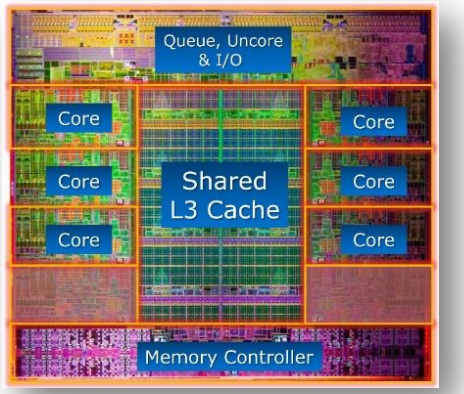

One of the Most Popular Open-Source Chip Projects

- **GitHub: >6700** Stars, > 800 Forks

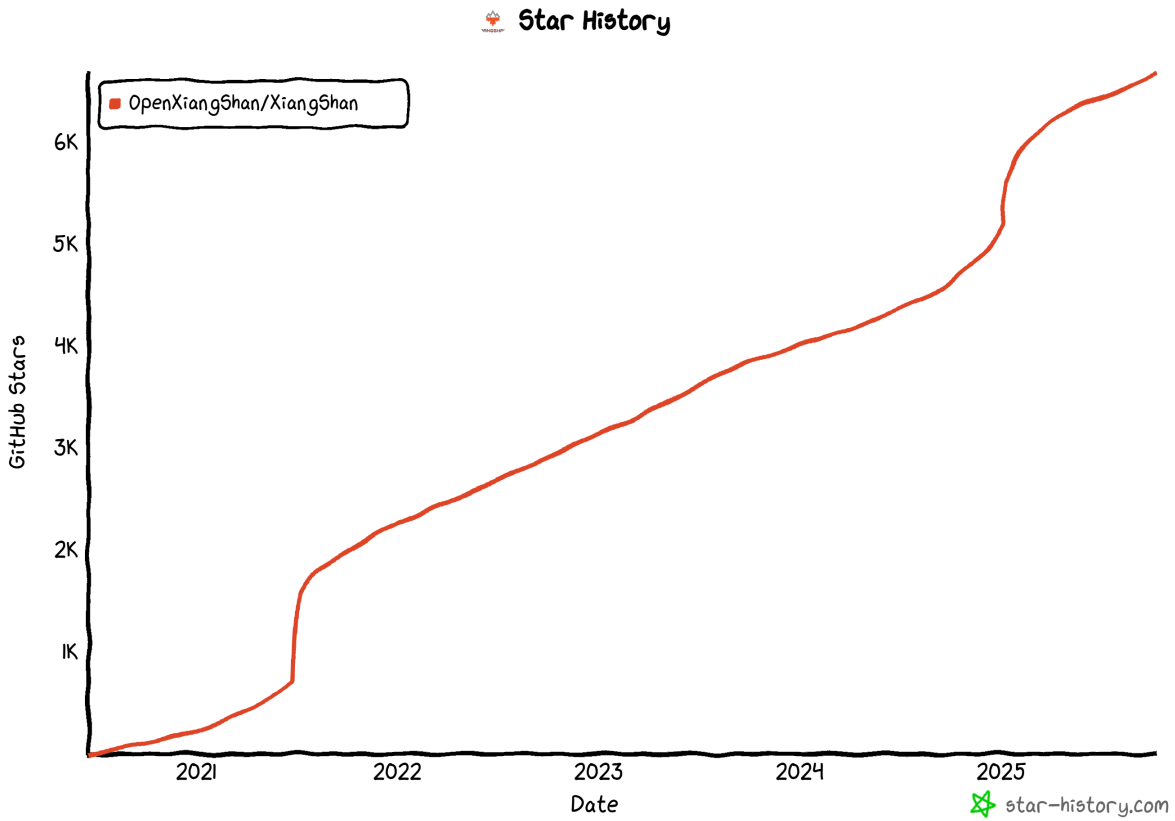

Data collected on Oct 10, 2025

15+ XiangShan Tutorials Around the World

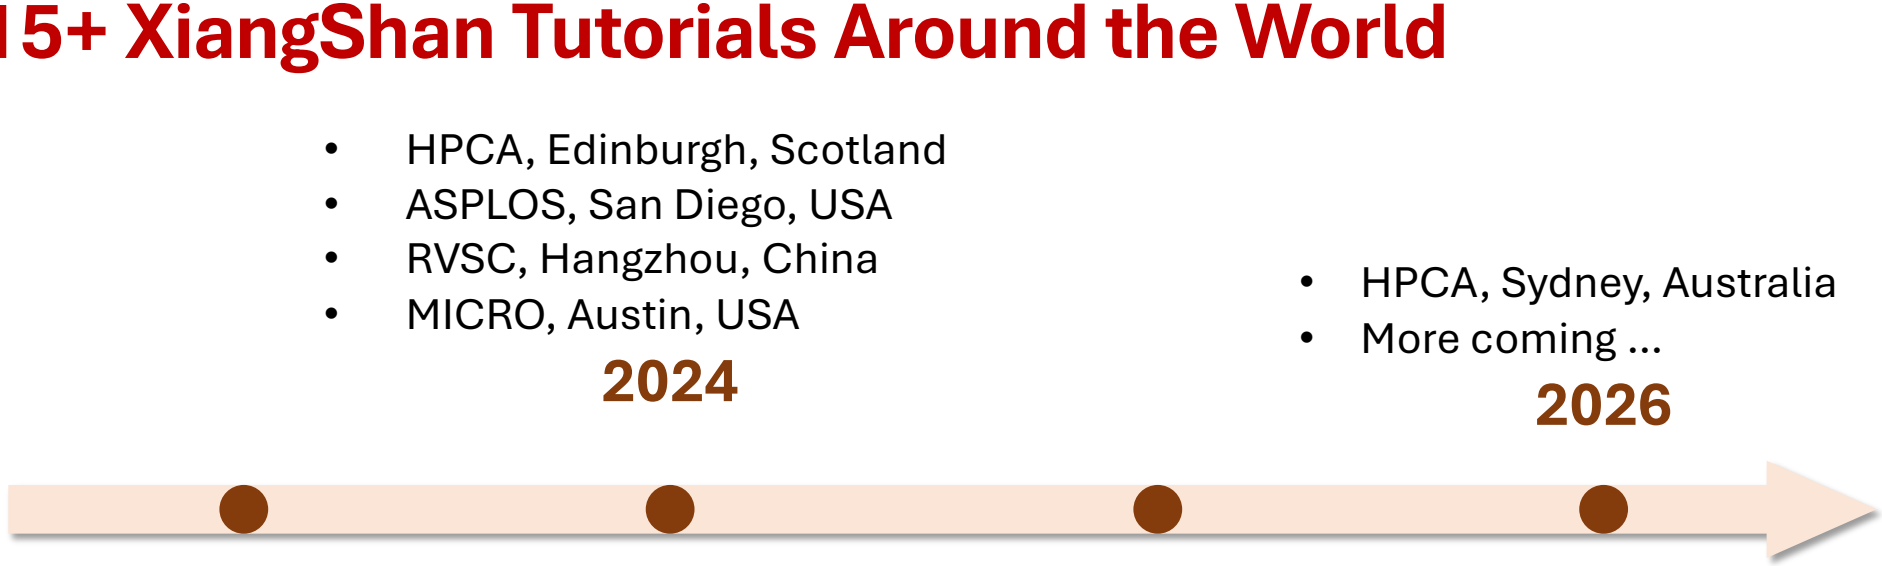

- ASPLOS, Vancouver, Canada
- RVSC, Beijing, China
- MICRO, Toronto, Canada

- HPCA, Edinburgh, Scotland
- ASPLOS, San Diego, USA
- RVSC, Hangzhou, China
- MICRO, Austin, USA

- HPCA, Las Vegas, USA
- ASPLOS, Rotterdam, Netherlands
- RVSE, Paris, France
- ISCA, Tokyo, Japan
- APPT, Athens, Greece
- RVSC, Shanghai, China
- ***MICRO, Seoul, Korea***

- HPCA, Sydney, Australia
- More coming …

Welcome old and new friends!

XiangShan Open-Source Community Conference

The 4th RISC-V Summit China 2024
Hangzhou, China

The 5th RISC-V Summit China 2025
Shanghai, China

93

GitHub Issues & Discussions

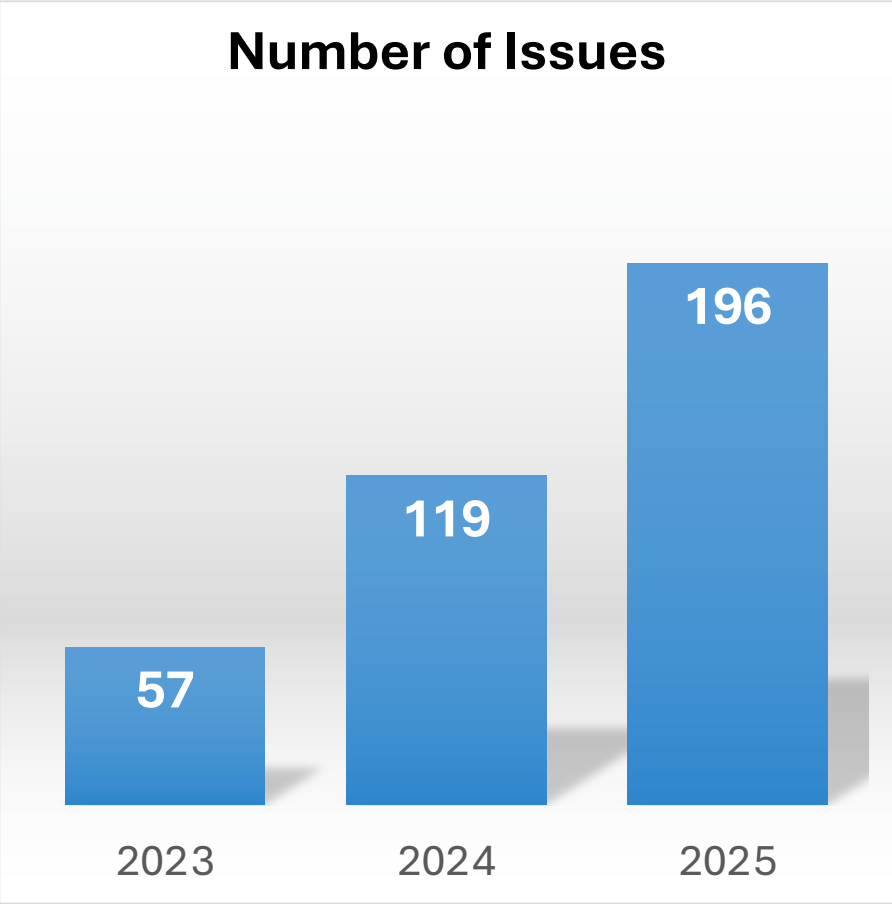

- Time to First Response for Issues
  - Average: 25 hours
  - Median: 13 hours
- Feel free to create an issue on GitHub

Acknowledgments in XiangShan

- A list outlines 32 techniques used in the XiangShan RTL codes
- **https://docs.xiangshan.cc/acknowledgments/**

Multi Port Data Cache

[1] Gurindar S. Sohi, and Manoj Franklin. "High-bandwidth data memory systems for superscalar processors." 4th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS). 1991. [RTL Codes]

Data Replacement

[1] Aamer Jaleel, Kevin B. Theobald, Simon C. Steely, and Joel Emer. "High performance cache replacement using re-reference interval prediction (RRIP)." 37th Annual International Symposium on Computer Architecture (ISCA). 2010. [RTL Codes]

Data Prefetch

[1] Jean-Loup Baer, and Tien-Fu Chen. "An effective on-chip preloading scheme to reduce data access penalty." ACM/IEEE Conference on Supercomputing. 1991. [RTL Codes]

[2] Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi and Andreas Moshovos. "Spatial memory streaming." 33rd International Symposium on Computer Architecture (ISCA). 2006. [RTL Codes]

[3] Santhosh Srinath, Onur Mutlu, Hyesoon Kim, and Yale N. Patt "Feedback directed prefetching: Improving the performance and bandwidth-efficiency of hardware prefetchers." IEEE 13th International Symposium on High Performance Computer Architecture (HPCA). 2007. [RTL Codes]

master XiangShan / src / main / scala / xiangshan / mem / prefetch / SMSPrefetcher.scala

Siudya feat(dfx): integerate dfx components (#4312)

Code Blame 1359 lines (1207 loc) · 58.6 KB

```
/***************************************************************************************
* Copyright (c) 2024 Beijing Institute of Open Source Chip (BOSC)
* Copyright (c) 2020-2024 Institute of Computing Technology, Chinese Academy of Sciences
* Copyright (c) 2020-2021 Peng Cheng Laboratory
*
* XiangShan is licensed under Mulan PSL v2.
* You can use this software according to the terms and conditions of the Mulan PSL v2.
* You may obtain a copy of Mulan PSL v2 at:
*          http://license.coscl.org.cn/MulanPSL2
*
* THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OF ANY KIND,
* EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO NON-INFRINGEMENT,
* MERCHANTABILITY OR FIT FOR A PARTICULAR PURPOSE.
*
* See the Mulan PSL v2 for more details.
*
*
* Acknowledgement
*
* This implementation is inspired by several key papers:
* [1] Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi and Andreas Moshovos. "[Spatial memory
* streaming.](https://doi.org/10.1109/ISCA.2006.38)" 33rd International Symposium on Computer Architecture (ISCA).
* 2006.
*

package xiangshan.mem.prefetch

import org.chipsalliance.cde.config.Parameters
import chisel3._
import chisel3.util._
import utils._
import utility._
```

Biweekly Report in English

- our recent progress and performance data
- https://docs.xiangshan.cc/zh-cn/latest/blog/category/biweekly-en/

Biweekly-en

2025年9月29日 · 分类于 Biweekly-en · 需要 5 分钟阅读时间

[XiangShan Biweekly 86] 20250929

Welcome to XiangShan biweekly column! Through this column, we will regularly share the latest development progress of XiangShan. We look forward to your contribution.

This is the 86th issue of the biweekly report.

We are very pleased to share two pieces of news with you.

On September 20, the XiangShan team won the first Open Source Contribution Award from the CCF Architecture Committee. This collective award holds special significance for the XiangShan team—it represents recognition and support from our academic peers for the open-source processor and the team itself, laying the foundation for XiangShan to have a broad impact. The XiangShan team will continue to move forward, step by step, striving to keep XiangShan alive for 30 years!

On September 22, Innosilicon released the "Fenghua 3" full-featured GPU. The "Fenghua 3" GPU successfully integrated the XiangShan "Nanhu" processor IP core, which is performance-competitive with the ARM Cortex-A76, as its high-performance on-chip main control CPU. This integration marks a new phase in the industrial application of open-source high-performance CPU IPs and signifies that RISC-V can carve out a path different from the traditional ARM model.

We believe that open-source chips do not equate to low performance or low quality. Open source will profoundly change the cost structure of chip development, providing a new paradigm for chip design in the industry.

User Guide

- For Software Developers and Hardware Integrators
- https://docs.xiangshan.cc/projects/user-guide/en/late

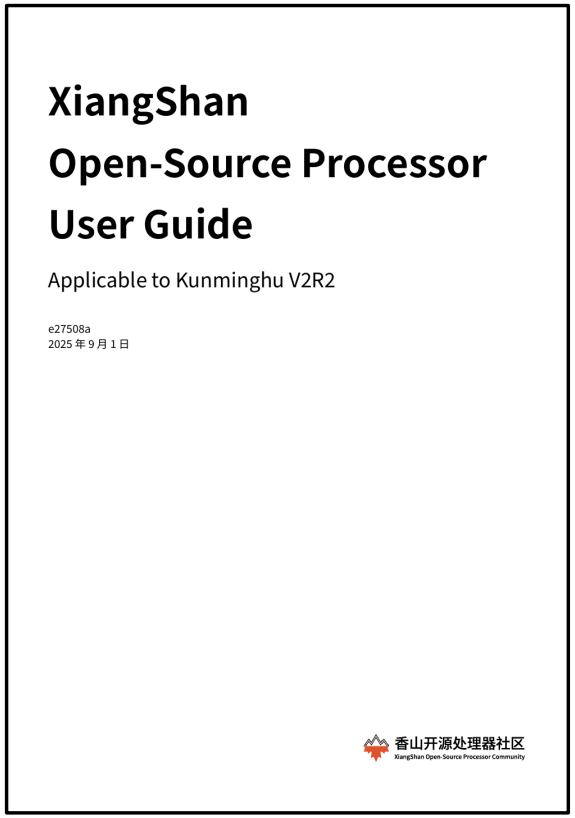

Design Document

- Detailed Docs on Microarchitecture and Modules
- https://docs.xiangshan.cc/projects/design/en/latest/

Open-Source: From Software to Hardware

- **96% software** codebases contain **open-source code** (overall 77%)
- *In the future, the proportion of **open-source IP in the chip industry will inevitably break through zero**, and will continue to increase.*

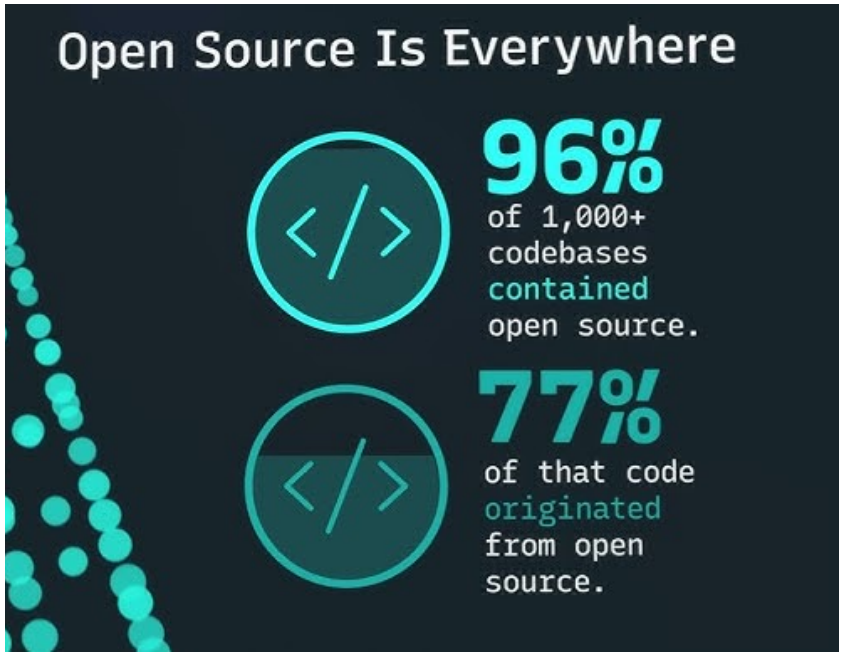

The 2024 Open Source Security and Risk Analysis (OSSRA) Report | Synopsys

OpenXiangShan: Together for a Shared Future

- Feel free to contact us through email or file issues on GitHub!
  - all@xiangshan.cc
  - https://github.com/OpenXiangShan/XiangShan

XiangShan Home page

XiangShan Document

XiangShan Biweekly Report

XiangShan User Guide

XiangShan Design Document