北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

香山服务器IP当前状态和路线图

产品经理：张健

2024-07-18

Contents

目录 /

北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

香山开源IP现状

香山开源IP的发展

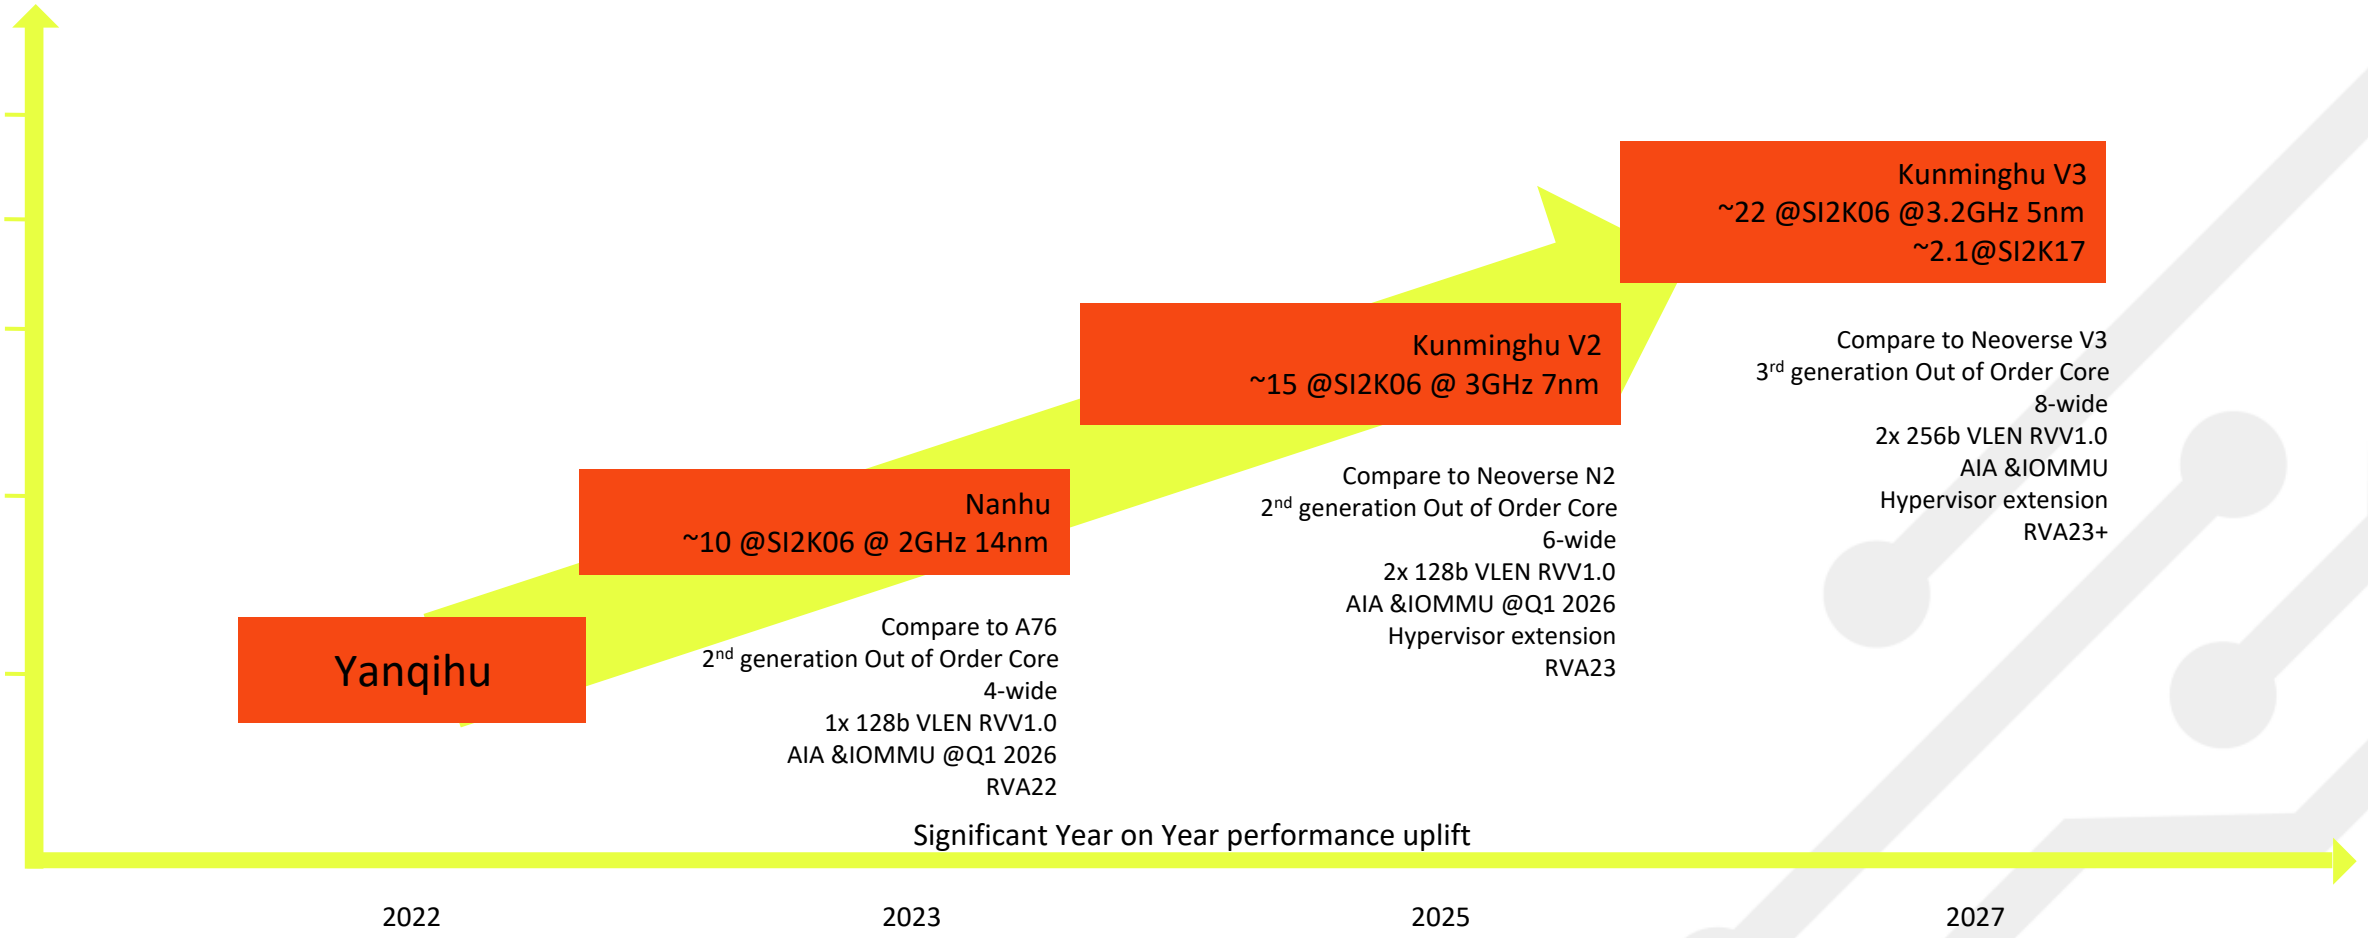
Yanqihu
Nanhu
~10 @SI2K06 @ 2GHz 14nm
Compare to A76
2nd generation Out of Order Core
4-wide
1x 128b VLEN RVV1.0
AIA &IOMMU @Q1 2026
RVA22
Kunminghu V2
~15 @SI2K06 @ 3GHz 7nm
Compare to Neoverse N2
2nd generation Out of Order Core
6-wide
2x 128b VLEN RVV1.0
AIA &IOMMU @Q1 2026
Hypervisor extension
RVA23
Kunminghu V3
~22 @SI2K06 @3.2GHz 5nm
~2.1@SI2K17
Compare to Neoverse V3
3rd generation Out of Order Core
8-wide
2x 256b VLEN RVV1.0
AIA &IOMMU
Hypervisor extension
RVA23+
Significant Year on Year performance uplift
2022
2023
2025
2027

香山开源IP性能对比

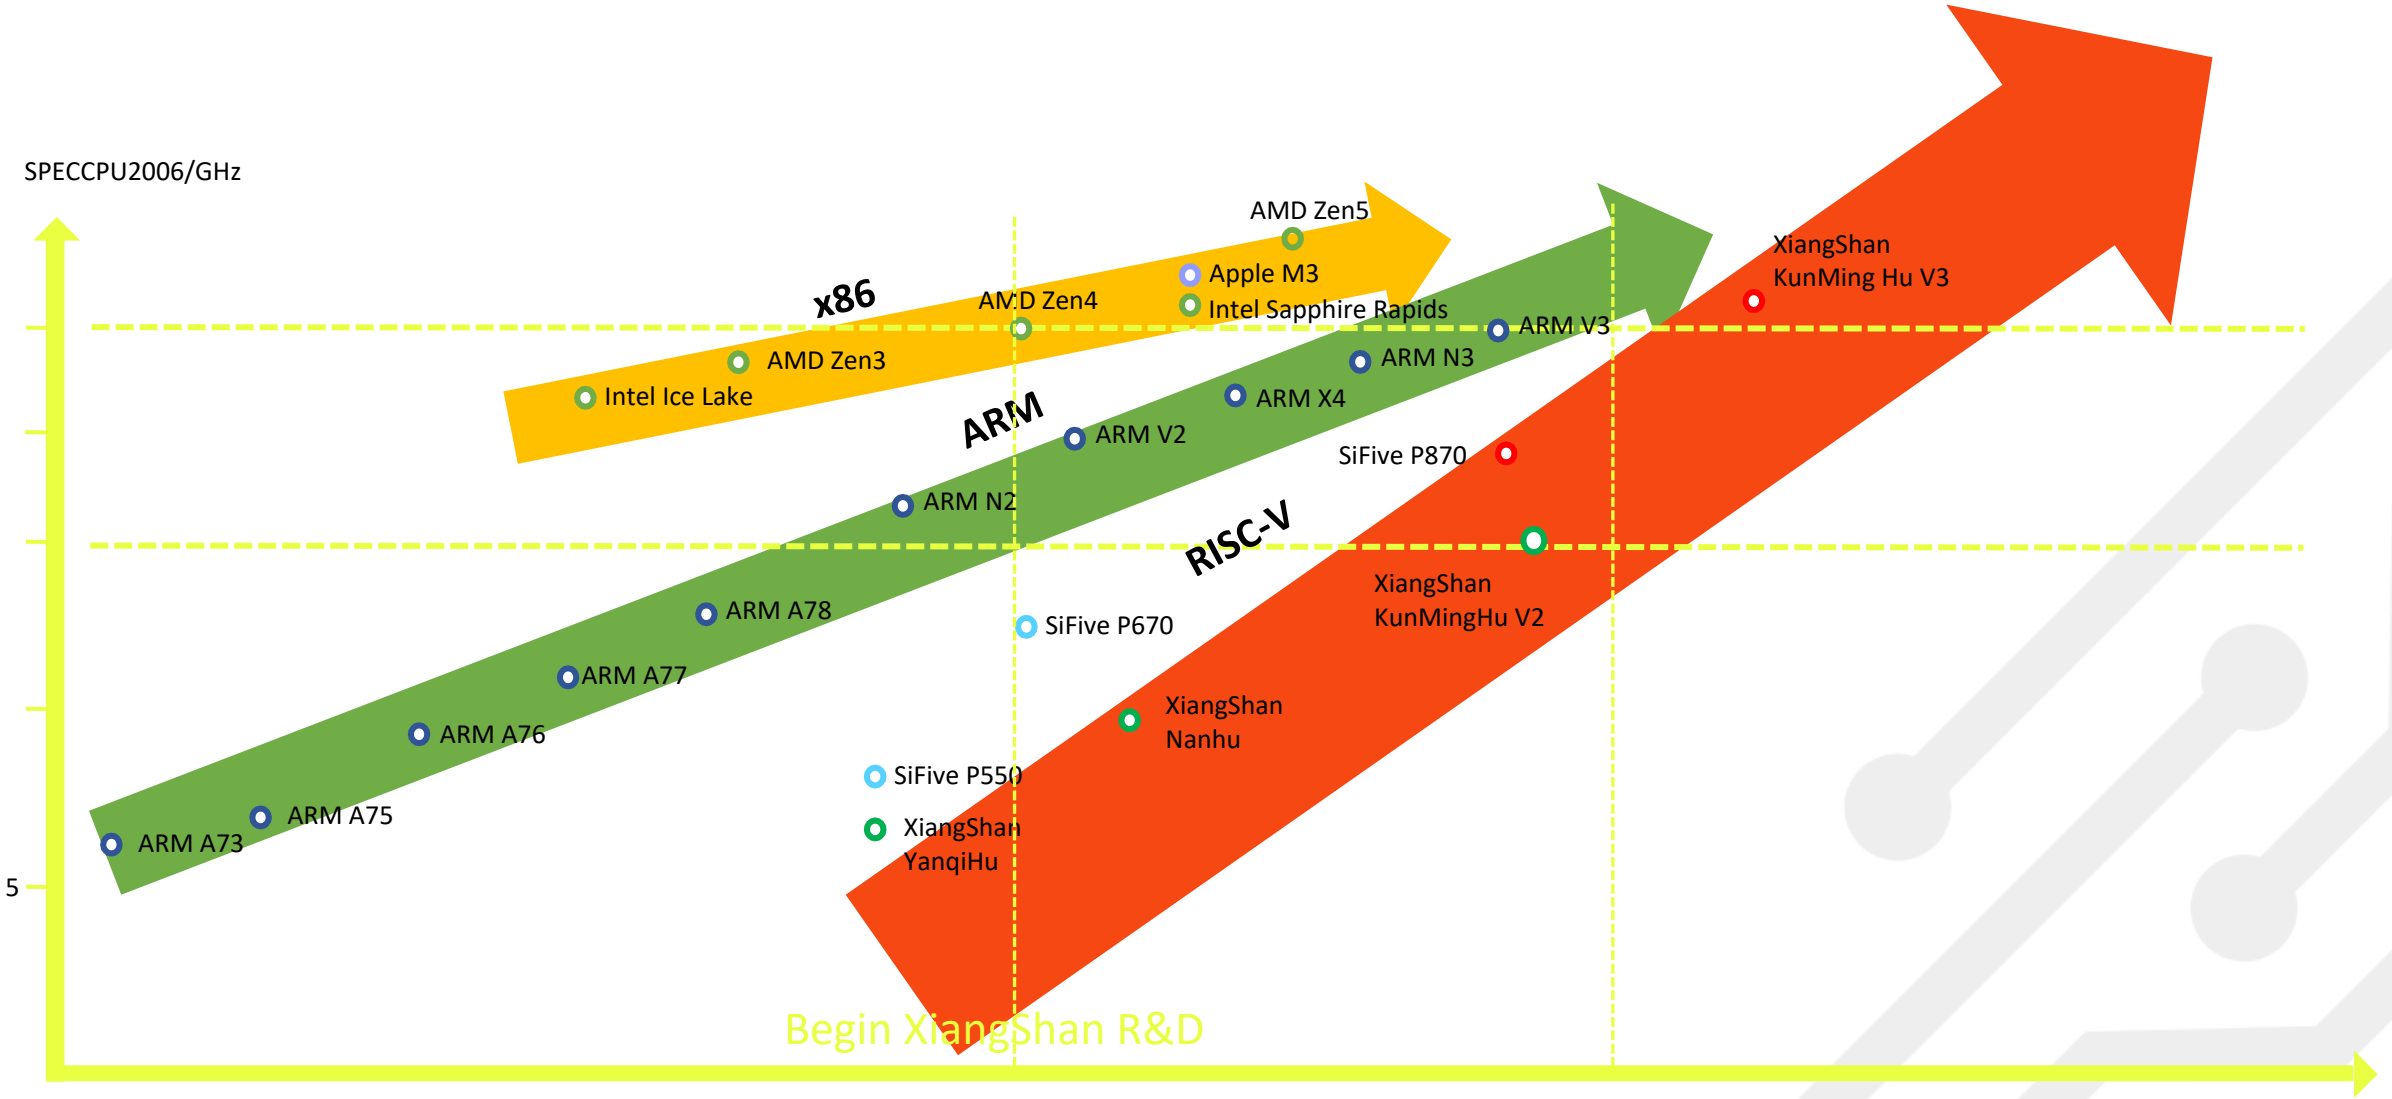
SPECCPU2006/GHz
5
x86
ARM
RISC-V
AMD Zen5
Apple M3
Intel Sapphire Rapids
AMD Zen4
AMD Zen3
Intel Ice Lake
ARM V3
ARM N3
ARM X4
ARM V2
ARM N2
ARM A78
ARM A77
ARM A76
ARM A75
ARM A73
XiangShan
KunMing Hu V3
SiFive P870
XiangShan
KunMingHu V2
SiFive P670
XiangShan
Nanhu
SiFive P550
XiangShan
YanqiHu
Begin XiangShan R&D

香山开源IP规格对比

【最新版】昆明湖V2R2与N2_CSS对比

	A	B	C	D	E	F
1	feature	Neoverse N2 CSS	RISC-V spec	XS-KMH V2R1(2024/06)	Q3	Q4
2	spec2006(GHz)	15	**6-15**	**15**		
3	Address space	VA 48bit, PA 48bit	**SV48, PA48**	**VA39bit, PA 36bit**	**SV48, PA48**	
4	memory region	find-grained	**coarse-grained(PMA&PMP) & Svpbmt**	**coarse-grained(PMA&PMP)**	**coarse-grained(PMA&PMP) & Svpbmt**	
5	outstanding transaction	memory, uncache memory, MMIO	same	memory	Yes(lack MMIO OT)	**Yes(IMMIO OT)**
6	Time keeping	private timer(with global timestamp)	**CLINT, sstc**	**CLINT**	**CLINT, sstc**	
7	NMI	Yes	**ss1p13**	**No**	**Yes(ss1p13)**	
8	vector	NEON/SVE/SVE2	**V**	**V**		
9	Performance monitor	Yes	**Shcounterenw,Sscofpmf,Sscounterenw,Zihpm**	**Shcounterenw,Sscofpmf,Sscounterenw,Zihpm**		
10	Debug	coresight debug	**Debug spec(0.13, 1.0)**	**Debug 0.13**		
11	Trace	Yes	**E-trace**	**No**	**E-trace**	
12	CPU virtualization	CPU, memory	**H,Shgatpa. Shvsatpa. Shvstvecd**	**H,Shgatpa. Shvsatpa. Shvstvecd（性能持续优化）**		
13	Interrupt virtualization	GIC	**AIA**	**AIA csr and IMSIC**		
14	misalignment access	Yes	**Zicclsm**	**No**	Zicclsm(wo vector)	**Zicclsm(w vector)**
15	cache maintenance	Yes	**zicbom,zicboz,zicbop**	**No**	**zicbom,zicboz,zicbop**	
16	Atomic	L1$, L2$, SLC with cas	**Zacas, Ziccamoc**	**Only D$ lrsc**	*(Ziccamoa)*	**Zacas, Ziccamoc**
17	CPU interface	CHI Issue E.b	**Tilelink, Tilelink2, CHI**	**CHI Issue B(subset)**	**CHI Issue B&E.b(subset)**	**Yes(with ~~AMO~~, CMO)**
18	power	retention, power down	retention, power down	**gating**	**retention**	**Yes(gating;core retention; power down)**
19	HW page table A/D update	Yes	**Yes(svadu)**	**No**		**~~Yes(svadu)~~(Delayed)**
20	nr of cores	up to 128	***Yes***	***up to 2***		Yes(up to 16)
21	CPU subsystem	Yes	**Maybe**	**NO(Yes for tilelink)**		**Yes**
22	RAS	I$ parity, D$ ECC, L2$ ECC	**Yes**	**I$ parity, D$ ECC(Default OFF)**	**(I$ parity + D$ ECC on)**	**(I$ parity + D$ ECC + L2 $ ECC)**
23	IOMMU	Yes(SMMU)	**IOMMU**	**open source(note1)**		Yes(wo PCIe)
24	DFT&MBIST	Yes	Yes	No	**Yes**	
25	Security	Realm	**Security Model(?)**	**No**		**机密虚拟机(CPU+interrupt; wo IO)**
26	multi-channel ddr	Yes	Maybe	**Yes(Tilelink)**	**Yes(Tilelink&CHI)**	
27	memory interleave	Yes	**Maybe**	**Yes(no work for CPU subsys)**		
28	multi-channel PCIe	Yes	**Maybe**	**Yes(in verification)**		
29	Die2Die support	Yes	**Maybe**	**No**		**Yes(IMISC, DebugModule)**
30	memory tag	Yes(MTE, PAC)	**Maybe**	**No**		**Yes(Ssnpm&Supm)**
31	cache partition and monitor	Yes(MPAM)	**Maybe**	**No**		

香山开源IP RVA23开发状态

	A	B	C	D	E	F	G	H	I	J	K
1		Finished before 20240630					2024Q3		2024Q4	Not supported by KMH-V2R2&NH5	
2	RVA23 U mandatory	A	C	D	F	M	Zicclsm(scalar only)	Zvfhmin	Zicclsm		
3		V	Zicntr	Zihpm	Ziccif	Ziccamoa	Zicbom	Zfhmin	Zawrs		
4		Zba	Zbb	Zbs	Zic64b	Zicsr	Zicbop		Zihintpause		
5		Zicboz	Zkt	Zicond	Za64rs		Zfa		Supm		
6		Ziccrse	Zcmop	Zvbb	Zvkt		Zimop		Zihintntl		
7	RVA24 U mandatory								Zacas		
8									Ziccamoc		
9	RVA23 S mandatory	Zifencei	Svbare	Sv39	Svade	Ssccptr	Svpbmt		Ssnpm		
10		Svinval	H	Ss1p13	Sstvecd	Ssu64xl	Svnapot				
11		Sstvala	Shtvala	Shgatpa	Ssstateen	Sscofpmf	Sstc				
12		Shvstvecd	Shvsatpa	Shvstvala	Sscounterenw	Shcounterenw					
13	Server mandatory (RVA23 optional)	Sdtrig			Sscsrind	Ssaia	Sv48	sv48x4	Ssstrict	Svadu	Zkr
14	RVA23 U optional	zbc	Zvfh				Zama16b		Zfh	Zvksg	Zvkng
15										Zfbfmin	Zvfbfmin
16										Zvbc	Zvfbfwma
17	RVA23 S optional	Sdext								Sv57	Svvptc
18											
19	南湖(NH)V5可选										

北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

香山核内特性进展

单核编译器性优化/服务器特性功能验证/开源软件

(only for唐丹）编译器性能摸高性能

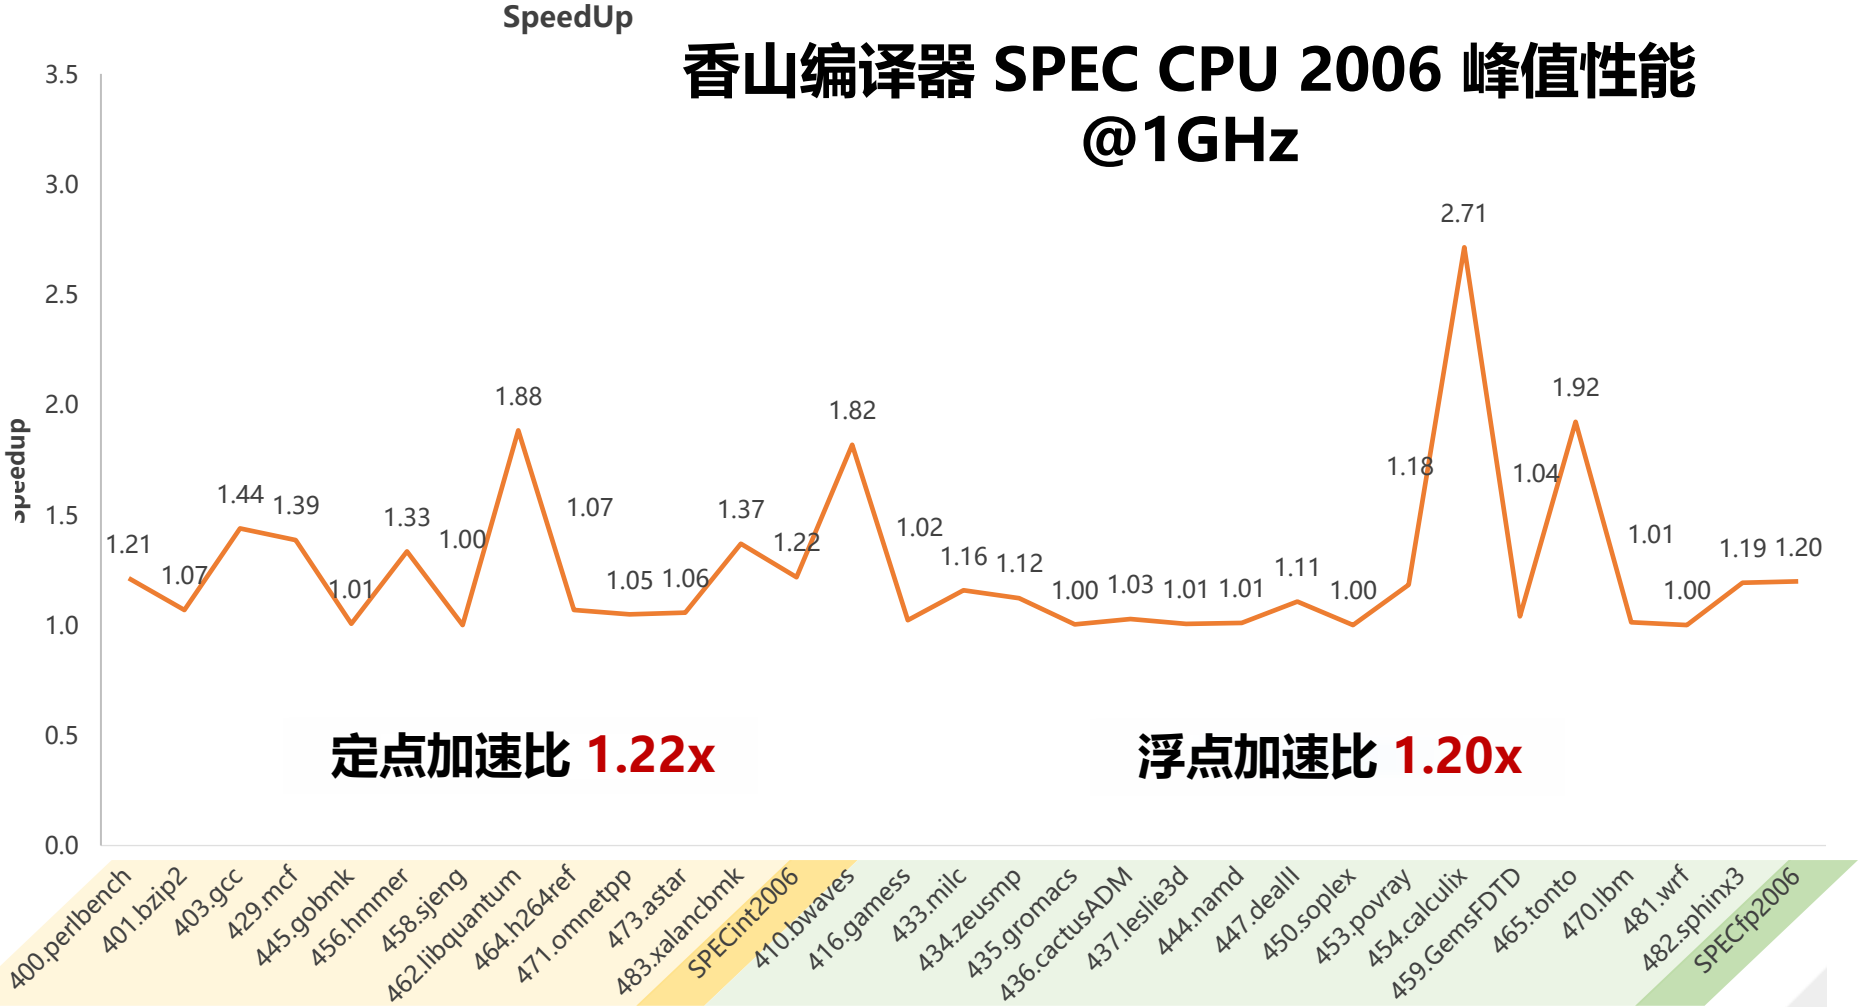

17.45

SPECint2006

18.00

SPECfp2006

- 相比于 LLVM 19.1.0　　编译选项：-O3 -ffast-math -flto -jemalloc　测试平台：XiangShan-Kunminghu RTL 仿真测试 (d084f29c)

香山开源编译器的性能提升

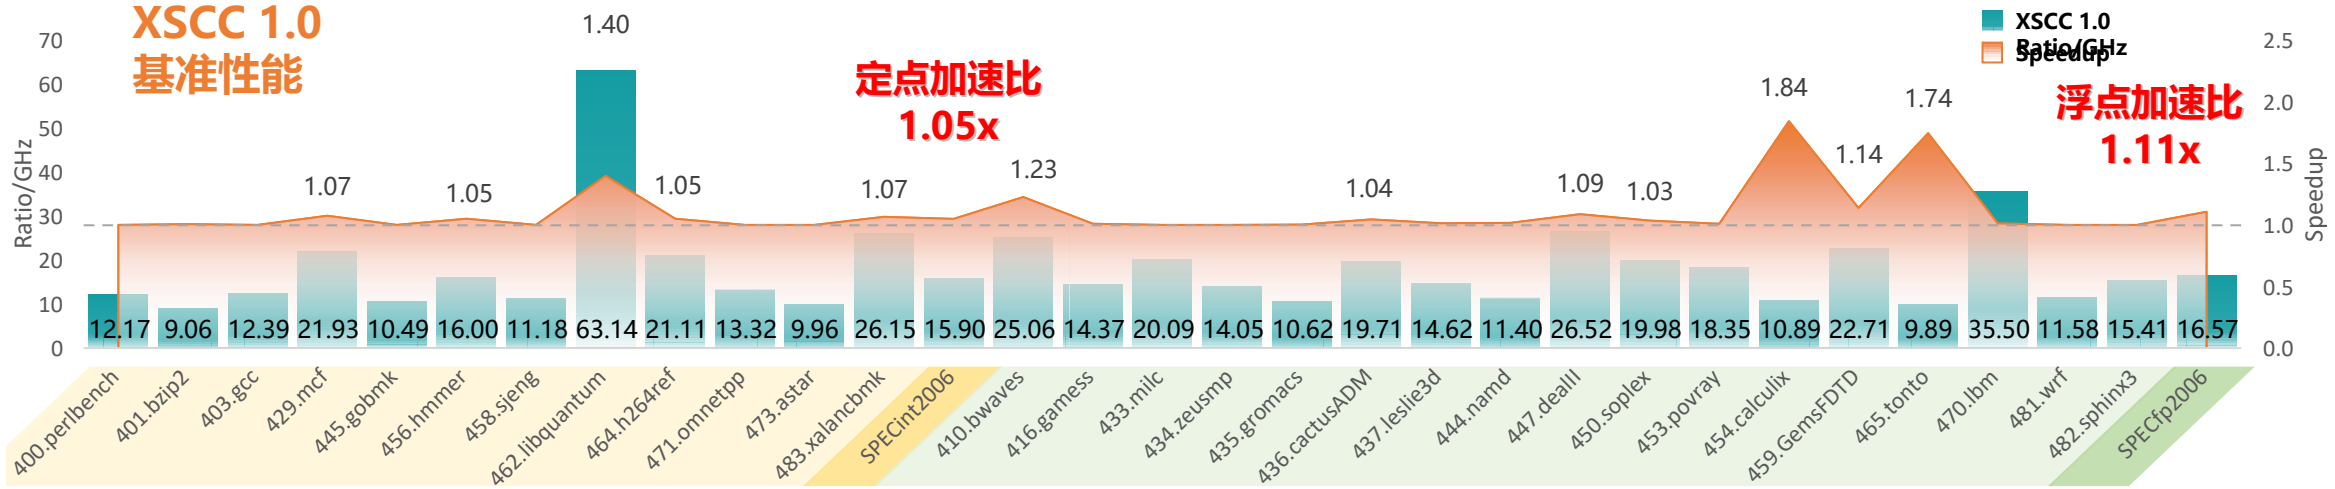

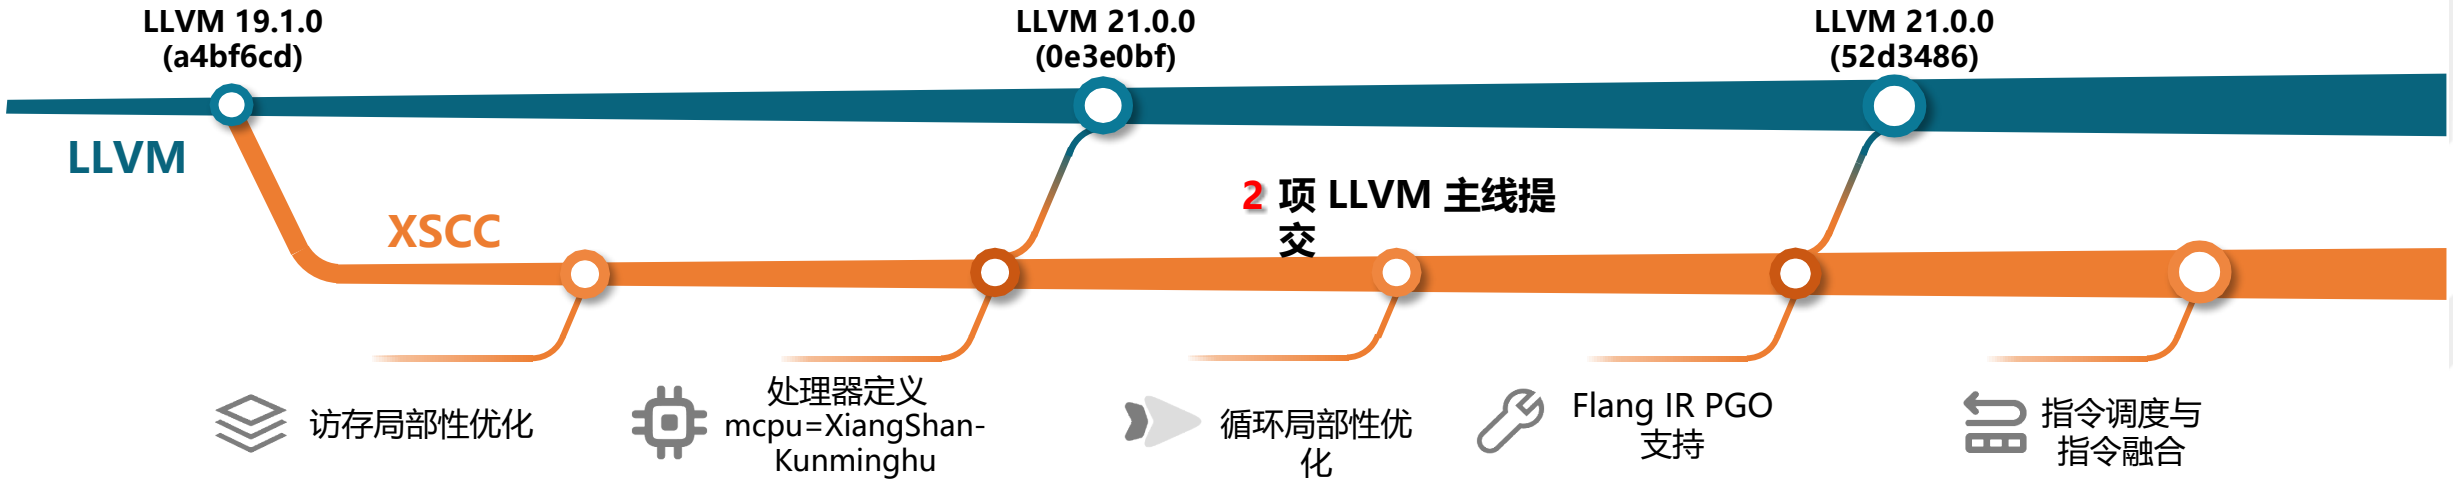

- 基于编译器性能摸高结果（specint 17.45），开芯院编译器团队逐步开源编译
- XSCC 编译选项: -O3 –ffast-math –flto –jemalloc　测试平台: XiangShan-Kunminghu GEM5 仿真测试 (7eb7cce4) 相比于 LLVM 19.1.0 编译选项: -O3 –ffast-math –flto –jemalloc

香山开源编译器的社区提交工作

香山昆明湖处理器定义 Merged

Commit 0e3e0bf

4 people authored on Apr 21 · ✕ 21 / 23 · Verified

```
[RISCV] Add processor definition for XiangShan-KunMingHu-V2R2 (#123193)

XiangShan-KunMingHu is the third generation of Open-source
high-performance RISC-V processor developed by Beijing Institute of Open
Source Chip (BOSC) , and its latest version is V2R2.

Co-authored-by: Shenglin Tang <tangshenglin@ict.ac.cn>
Co-authored-by: Xu, Zefan <ceba_robot@outlook.com>
Co-authored-by: Tang Haojin <tanghaojin@outlook.com>
```

main (#123193)

Fortran PGO 功能支持 Merged

Fix and reapply IR PGO support for Flang #142892

Merged tarunprabhu merged 5 commits into llvm:main from OpenXiangShan:OpenXiangShan/reapply-flang-pgo last month

SPEC2006	Runtime without PGO (Sec)	Runtime with PGO (Sec)	Speedup
410.bwaves	101.0	97.6	**1.03x**
416.gamess	259.0	244.0	**1.06x**
437.leslie3d	94.7	94.1	1.01x
454.calculix	182.0	180.0	1.01x
481.wrf	93.4	91.4	1.02x

IR结构化简 Open

[IndVarSimplify] Eliminated Pure LoopCounter #146845

Open buggfg wants to merge 1 commit into llvm:main from OpenXiangShan:OpenXiangShan/elimininting-pure-loop-counter

- 放宽 LoopCounter 步幅限制，激发多循环优化潜力
- 合并PHI节点，增加循环展开与向量化机会

计算结构化简 Open

[Flang] Canonicalize divdc3 calls into arithmetic-based complex division #146017

Open Hanwen-ece wants to merge 17 commits into llvm:main from OpenXiangShan:OpenXiangShan/flang-new-generalizes-complex-division

SPEC2017	Runtime with __divdc3 function calls (Sec)	Runtime with arithmetic-based complex division (Sec)	Speedup
627.cam4_s	1688.0	1529.0	**1.10x**

（only for唐丹）香山编译器在其它RISC-V架构同样可获得性能提升

SpacemiT X60

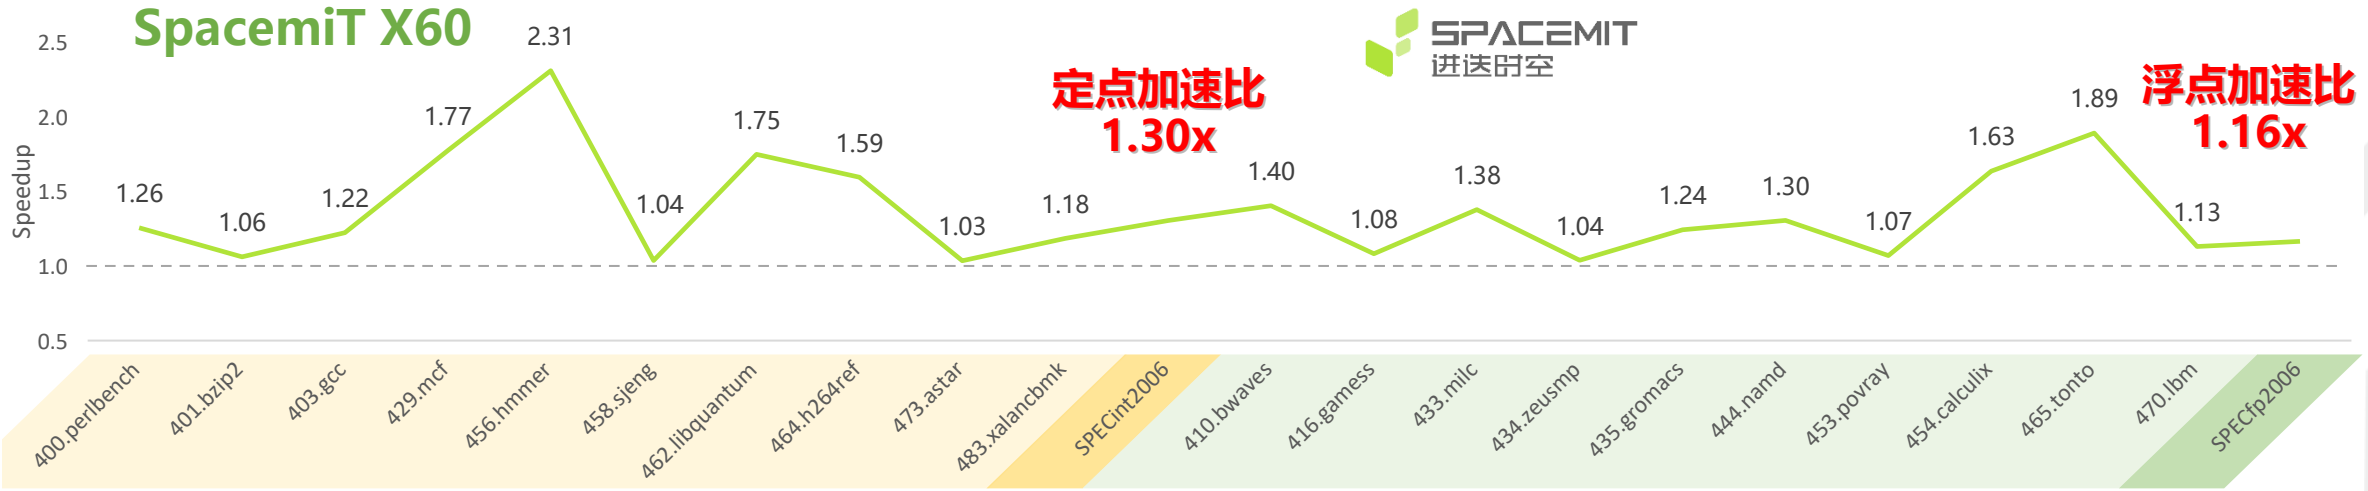

XuanTie C920

砥砺不殆
期待更多

* 相比于 LLVM 19.1.0　编译选项: -O3 –ffast-math –flto –jemalloc

RVA23功能验证

EDA验证

UT

BT

IT

面向业务的ST用例开发（GOS）

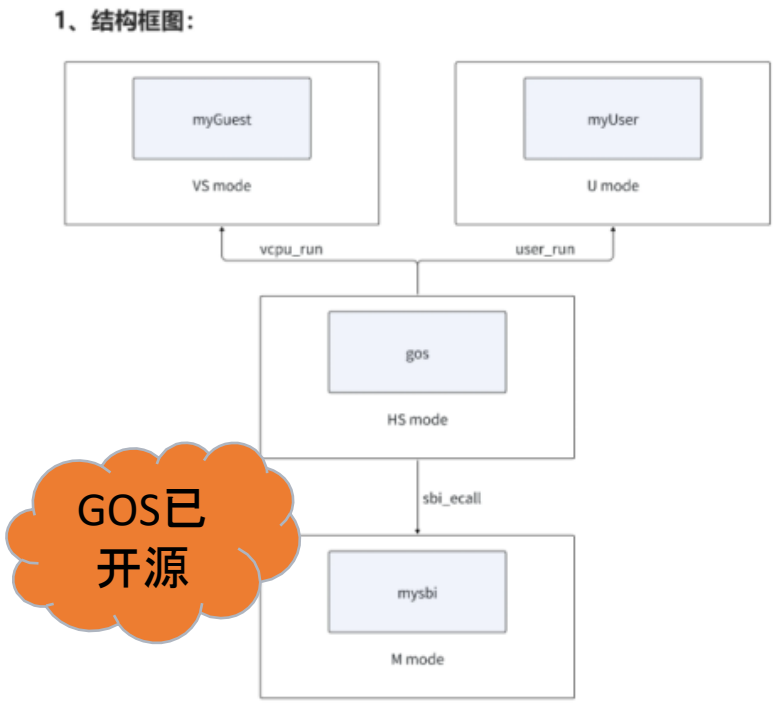

openEuler RVA23用例

用例名称	验证状态
Kernel build	进行中
C++/protobuf	单核完成
Python3	单核完成
golang	单核完成
openBLAS	单核完成
mysql client	单核完成
memcached	单核完成
Java/hadoop	单核完成
mysql	单核完成

感谢openEuler RISC-V SIG提供用例和执行验证

北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

从CPU到计算子系统

昆明湖+温榆河
南湖+珠江

单核性能提升（SPECCPU base)

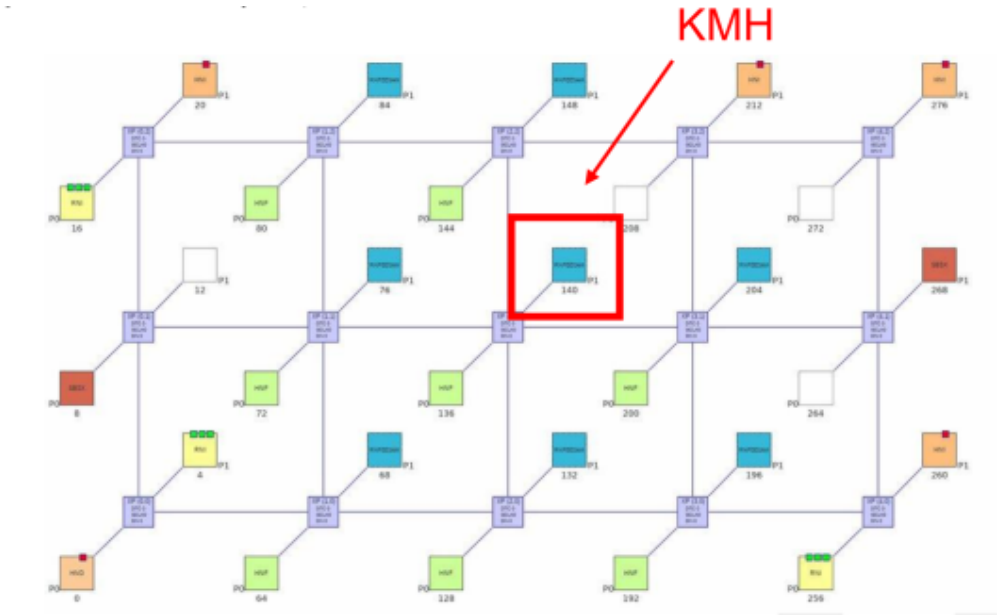

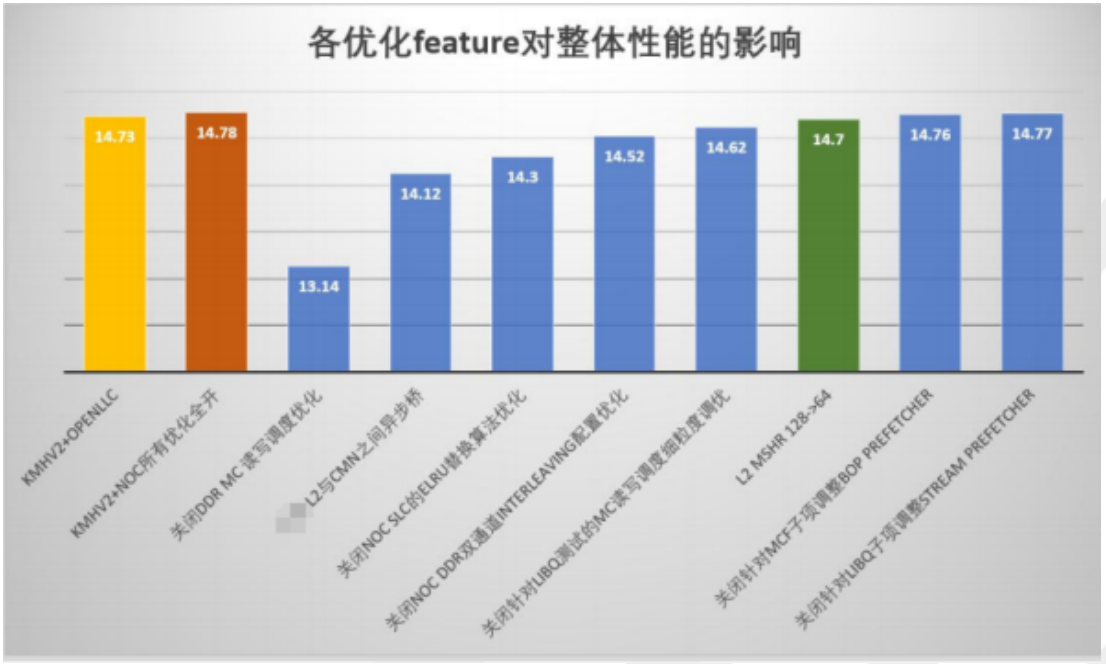

- 通过包括CPU,NoC和MC累计8项硬件性能优化，在给定延迟和带宽下获得24.9%性能提升（11.83->14.78）；
- 调优展示了香山开源项目对于不同mesh尺寸（2x2 vs 3x5），不同latency（186cycles->219cycles），有很好的适配能力；
- 昆明湖支持通过预取提前量调整适配mesh NoC的延迟。

备注：1. 3x5的mesh最大可以支持16核（通过CAL）；2. 系统中只有单核昆明湖运行，在次系统下，昆明湖实测可用的最大带宽为39.9GB/s；3. 使用ddr4 3200

香山开源计算子系统

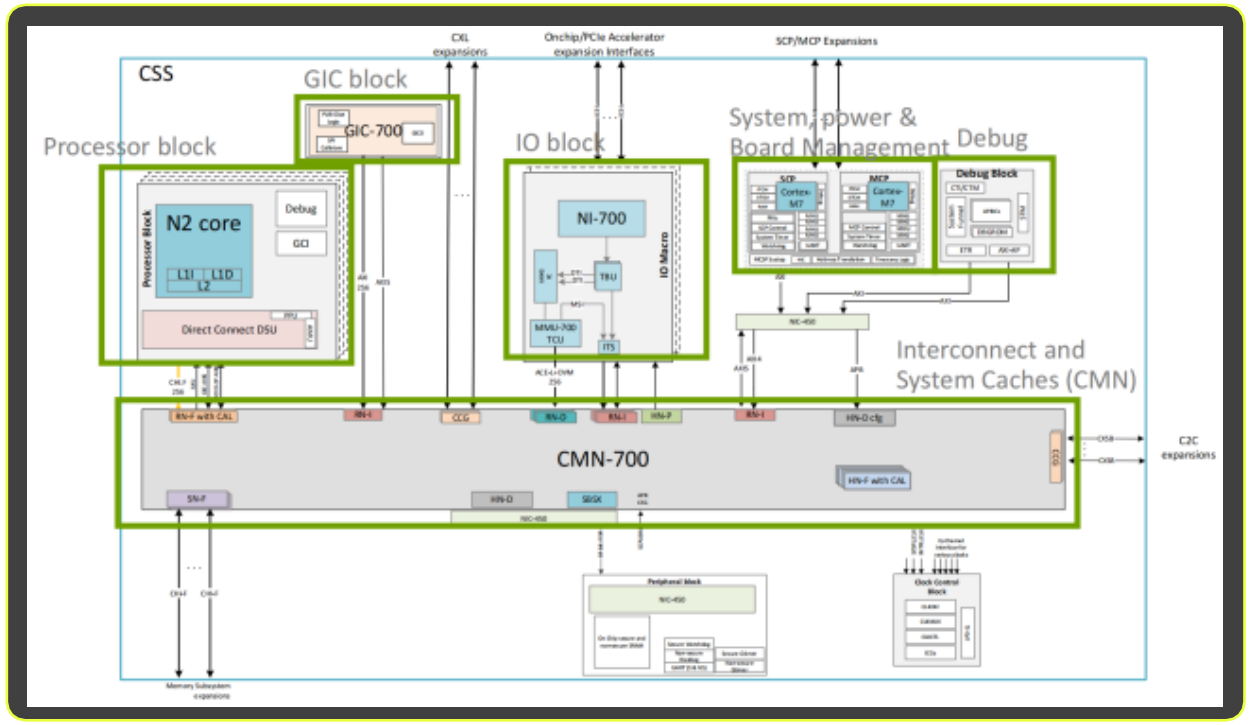

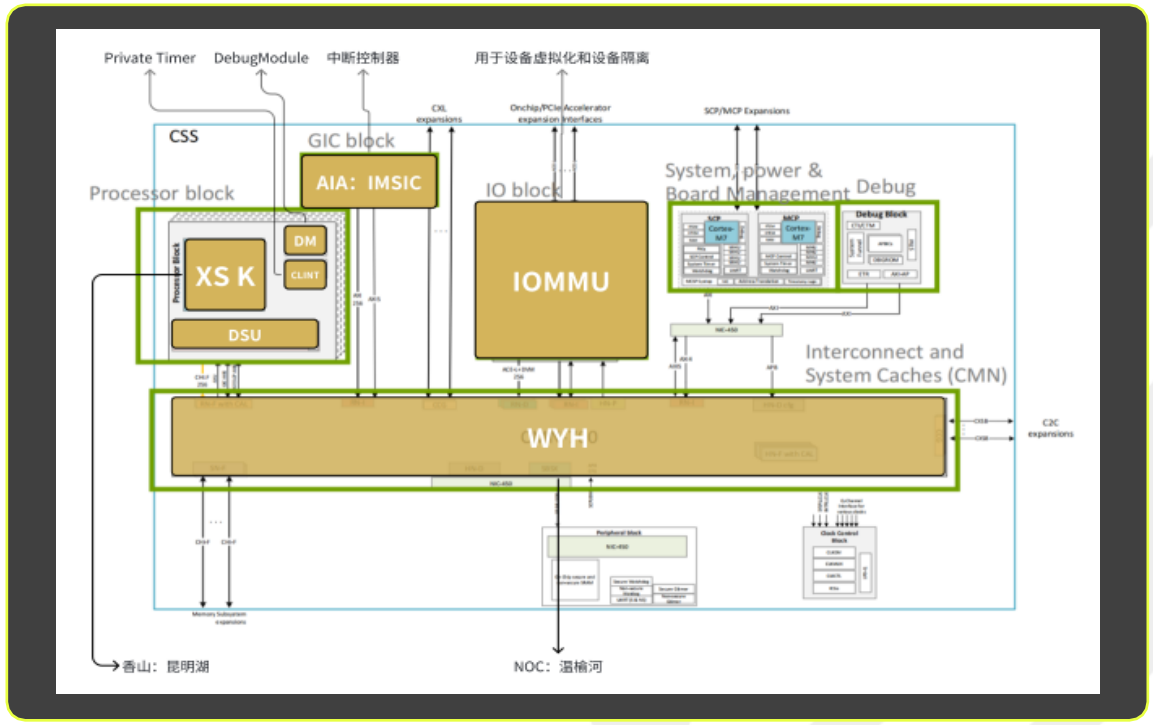

相关上游规范/认证

- SBSA(Server Base System Architecture)
- SBBR(Server Base Boot Requirements)
- ARM Server Ready

相关上游规范

- RISC-V server platform specification
  - RVA profile
  - RISC-V server SOC specification
  - Boot and Runtime Services specification
  - RISC-V platform security model
- Certification(Certification Steering Commitee)

昆明湖V2：满足服务器CPU规格的CPU子系统

■ CPU子系统可配置选项

类别	可配置项名称	选项
CPU Core	L2 Cache size(KB)	128, 256, 512, 1024, 2048
中断(AIA)	IMSIC MSI接口	MSI写接口，AX4总线
中断(AIA)	IMSIC vs interrupe file	5-63
中断(AIA)	IMSIC 中断数量	255, 511, 1023, 2047
CLINT	是否采用全局时间戳	Y/N
Debug	是否使用私有mmio通路	Y/N
Debug	Debug rom地址	48bit PA
Debug	external trigger	0-16

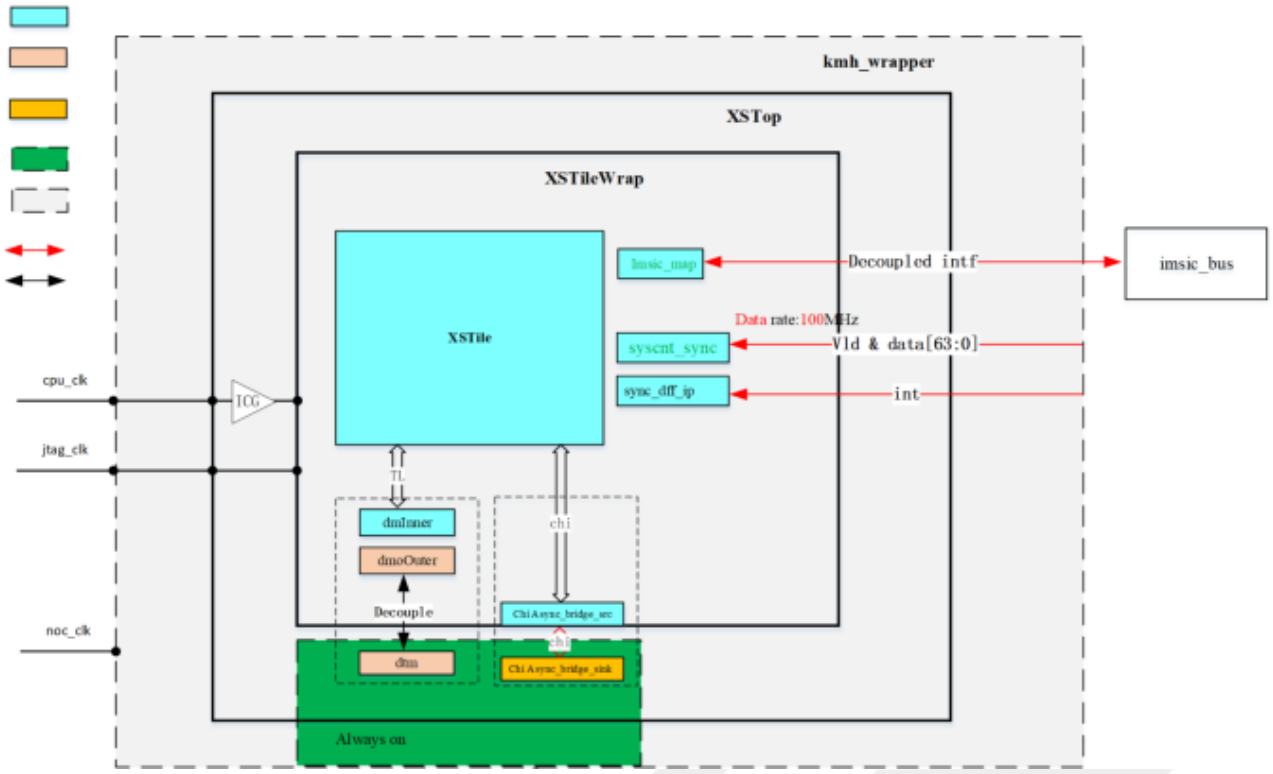

昆明湖对Server SOC IP的支持

AIA	
IIC_010	支持AIA规范（CSR和IMSIC）
IIC_020	外部中断支持MSI
IIC_030	IMSIC支持S mode
IIC_040	IMSIC支持5个VS interrupt file
IIC_050	IMSIC S mode interrupt file至少支持255个中断
IIC_060	IMSIC VS mode interrupt file至少支持255个中断
IIC_070	IMSIC支持MMIO访问
IIC_080	支持APLIC

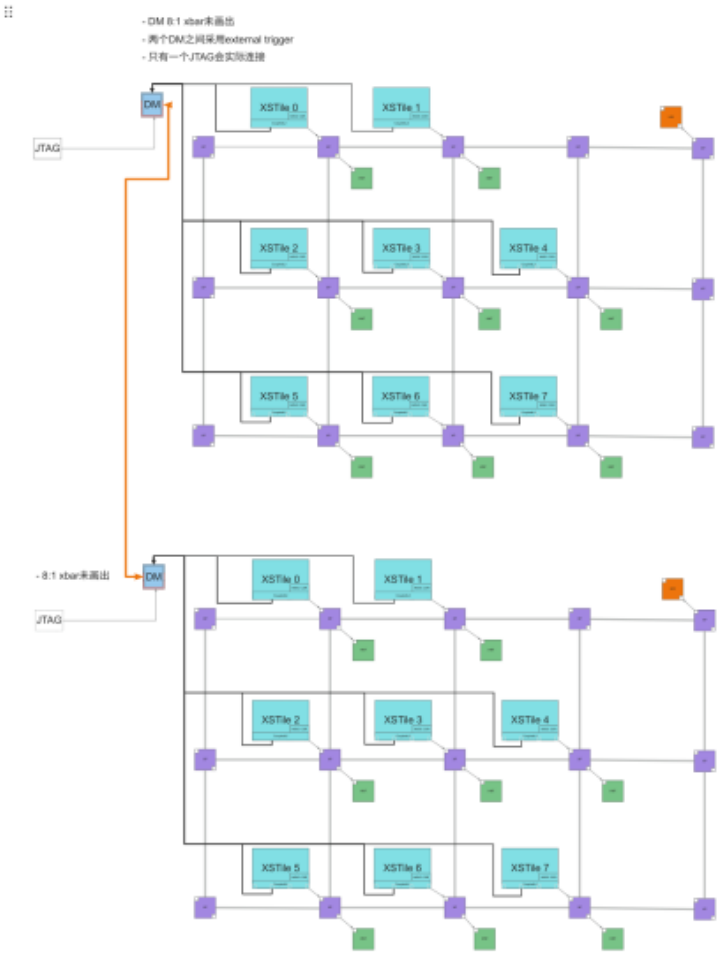

iommu开发现状
IOM_010
IOM_020
IOM_030
IOM_040
IOM_050
IOM_060
IOM_070
IOM_080
IOM_090
IOM_100
IOM_110
IOM_120
IOM_130
IOM_140
IOM_150
IOM_160
IOM_170
IOM_180
IOM_190
IOM_200
IOM_210
IOM_220
IOM_230
IOM_240
IOM_250
IOM_260
IOM_270
IOM_280
IOM_290
IOM_300
IOM_310
IOM_320

机密虚拟机方案

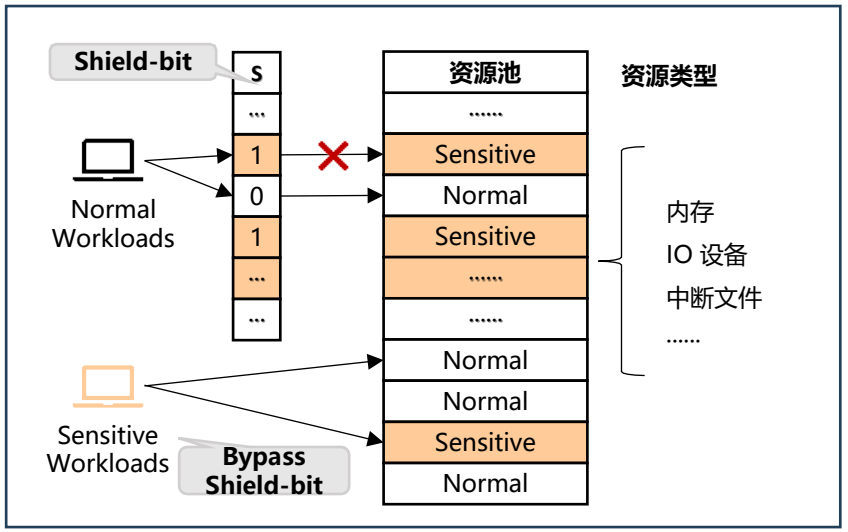

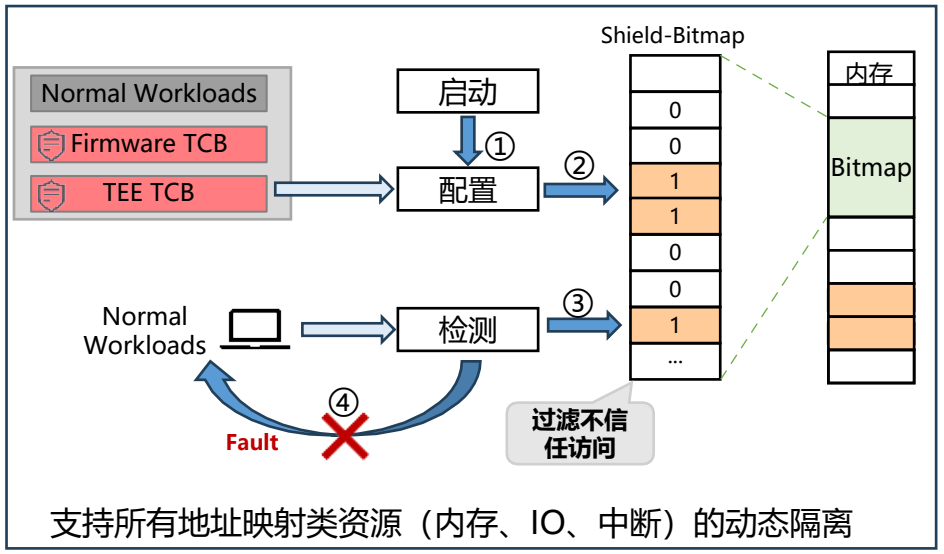

支持所有地址映射类资源（内存、IO、中断）的动态隔离

- **Configuration**
  - ✓ Configure Shield-Bitmap CSR
    - Enable Shield-XS Isolation
    - Set Shield-Bitmap Base Address
  - ✓ Shield-Bitmap Cache Size = **128 x 8 Bytes**
  - ✓ KunminghuV2 Configuration (TileLink Prototype)
    - L1 Instruction/Data Cache Size = 64KB
    - L1 Instruction/Data TLB = **48-Full Association**
    - L2 Cache Size = 1MB
    - L3 Cache Size = 16MB

- **Performance**
  - ✓ SPECInt2006 Simpoint est.@**3GHz**
    - **GEOMEAN 44.62 → 44.29 (↓0.72%)**

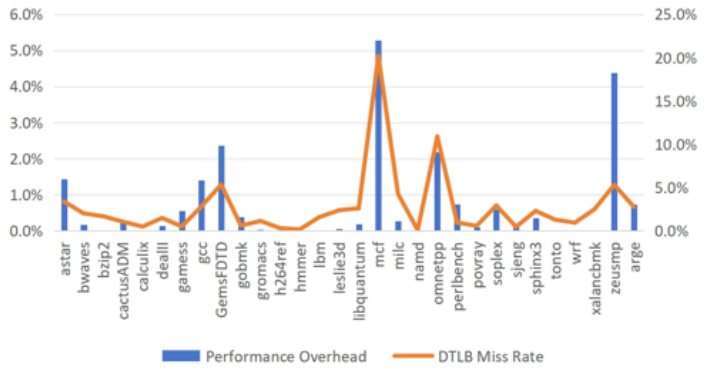

The performance overhead increases with the DTLB miss rate.

- **Area**
  - ✓ Process Design Kit (PDK) 7nm Evaluation
    - Hardware Overhead ≈ **0.3%**

Module	Cell Area (μm^2)
Shield-Bitmap Cache	5,075
Shield-Bitmap Checker	1,088
MMU Area	50,843
KMH V2 Core Area	2,000,000
Area Percentage	**0.3%**

北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

SOC参考设计的验证

IT/ST

多核验证环境

tl-test-new：支持多核的L2验证环境

```
2140  [CLog.B] closed file: clog.l3top.b
2141  saving memdb to chiseldb.db ...
2142  [7200000] [tl-test-new-FINAL] void TLSequencer::Finalize(): called at 7200000
2143  [7200000] [tl-test-new-FINAL] "                                                "
2144  [7200000] [tl-test-new-FINAL] "  ▄▄▄▄▄▄ ▄▄ ▄▄  ▄▄ ▄▄ ▄▄▄▄▄▄ ▄▄  ▄▄ ▄▄▄▄▄▄ ▄▄▄▄▄  ▄▄  "
2145  [7200000] [tl-test-new-FINAL] "  ▄▄  ▄▄ ▄▄▄▄  ▄▄ ▄▄ ▄▄  ▄▄ ▄▄  ▄▄ ▄▄  ▄▄ ▄▄  ▄▄ ▄▄  "
2146  [7200000] [tl-test-new-FINAL] "  ▄▄▄▄ ▄▄ ▄▄ ▄▄ ▄▄ ▄▄▄▄▄▄ ▄▄▄▄▄▄ ▄▄▄▄▄▄ ▄▄  ▄▄ ▄▄  "
2147  [7200000] [tl-test-new-FINAL] "  ▄▄   ▄▄ ▄▄ ▄▄▄▄ ▄▄    ▄▄ ▄▄ ▄▄ ▄▄  ▄▄ ▄▄  ▄▄   "
2148  [7200000] [tl-test-new-FINAL] "  ▄▄   ▄▄ ▄▄  ▄▄▄▄ ▄▄▄▄▄▄ ▄▄ ▄▄ ▄▄▄▄▄▄ ▄▄▄▄▄  ▄▄  "
2149  [7200000] [tl-test-new-FINAL] "                                                "
2150  [7200000] [tl-test-new-FINAL] void TLSequencer::Finalize(): tl-test subsystem finalized
```

TXRSP		RXRSP		TXDAT	
Transaction	Count	Transaction	Count	Transaction	Count
RespLCrdReturn	0	RespLCrdReturn	0	DataLCrdReturn	0
SnpResp	2351705	SnpResp	0	SnpRespData	2337024
CompAck	29067959	CompAck	0	CopyBackWrData	7347689
RetryAck	0	RetryAck	23889421	NonCopyBackWrData	1652662
Comp	0	Comp	15343294	CompData	2436240
CompDBIDResp	0	CompDBIDResp	3673848	SnpRespDataPtl	0
DBIDResp[1]	0	DBIDResp[1]	0	SnpRespDataFwded	1363540
PCrdGrant	0	PCrdGrant	23888886	WriteDataCancel	0
ReadReceipt	0	ReadReceipt	1653580	DataSepResp	0
SnpRespFwded	536350	SnpRespFwded	0	NCBWrDataCompAck	0

基于定向随机的覆盖率收集（局部）

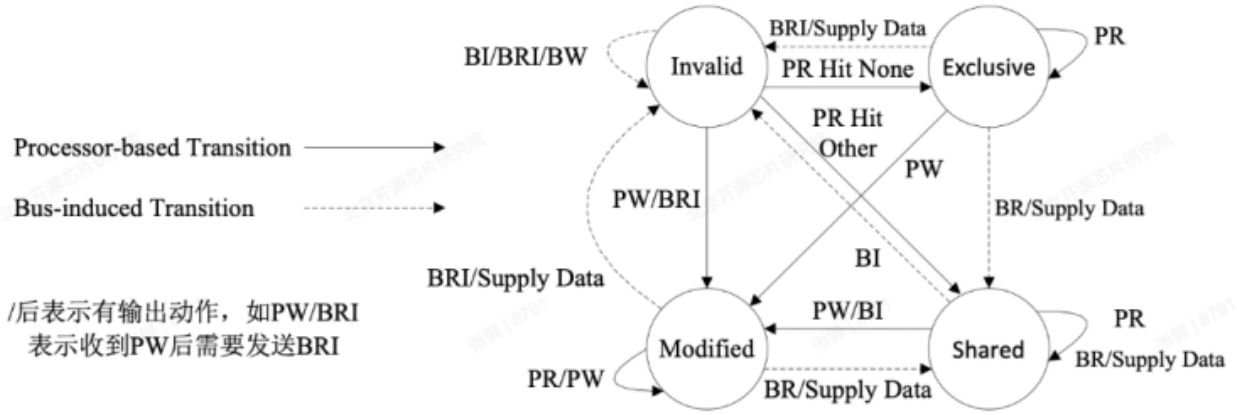


```
i == 0x0 overlap_producer_stress
i == 0x0 overlap_consumer_stress
overlap_producer_stress    data check pass!
overlap_consumer_stress    data check pass!
i == 0x0 rw_consumer_stress
i == 0x0 rw_producer_stress
rw_producer_stress    data check pass!
rw_consumer_stress    data check pass!
i == 0x0 rw_producer_stress
i == 0x0 rw_consumer_stress
rw_consumer_stress    data check pass!
rw_producer_stress    data check pass!
i == 0x0  ww_consumer_stress
i == 0x0  ww_producer_stress
ww_consumer_stress    data check pass!
ww_producer_stress    data check pass!
i == 0x0  ww_producer_stress
i == 0x0  ww_consumer_stress
ww_producer_stress    data check pass!
ww_consumer_stress    data check pass!
Child with PID 82 exited:0x0.
Child with PID 83 exited:0x0.
```

验证内容

cache test

验证结果（局部）

DiffTest-H: 基于硬件加速器的协同仿真框架

- **特点：对验证数据包的通信优化**
  - **Batch**：降低通信频次
  - **Squash**：减少通信数据量
  - **Replay**：保持指令级别的调试粒度
- 仿真验证加速比达 **57~109** 倍
- DiffTest-H 帮助香山发现了 **151** 个硬件 bug，涉及到 **780** 行和 **19** 个 PR 的代码修改

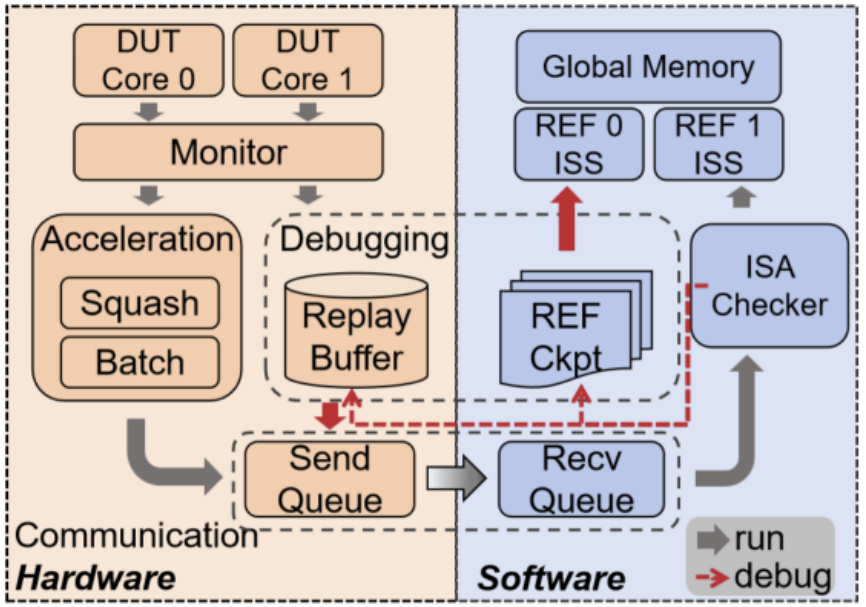

CPU	NutShell	香山昆明湖	香山昆明湖
环境	帕拉丁	帕拉丁	FPGA
配置	默认配置	默认配置	mini config
无优化	14KHz	4KHz	0.1MHz
+Batch	102 KHz (7×)	24 KHz (6×)	1.3 MHz (13×)
+NonBlock	398 KHz (28×)	71 KHz (12×)	2.2 MHz (22×)
+Squash	1080 KHz (77×)	478 KHz (109×)	5.7 MHz (57×)

IO通路的验证：定制化EP

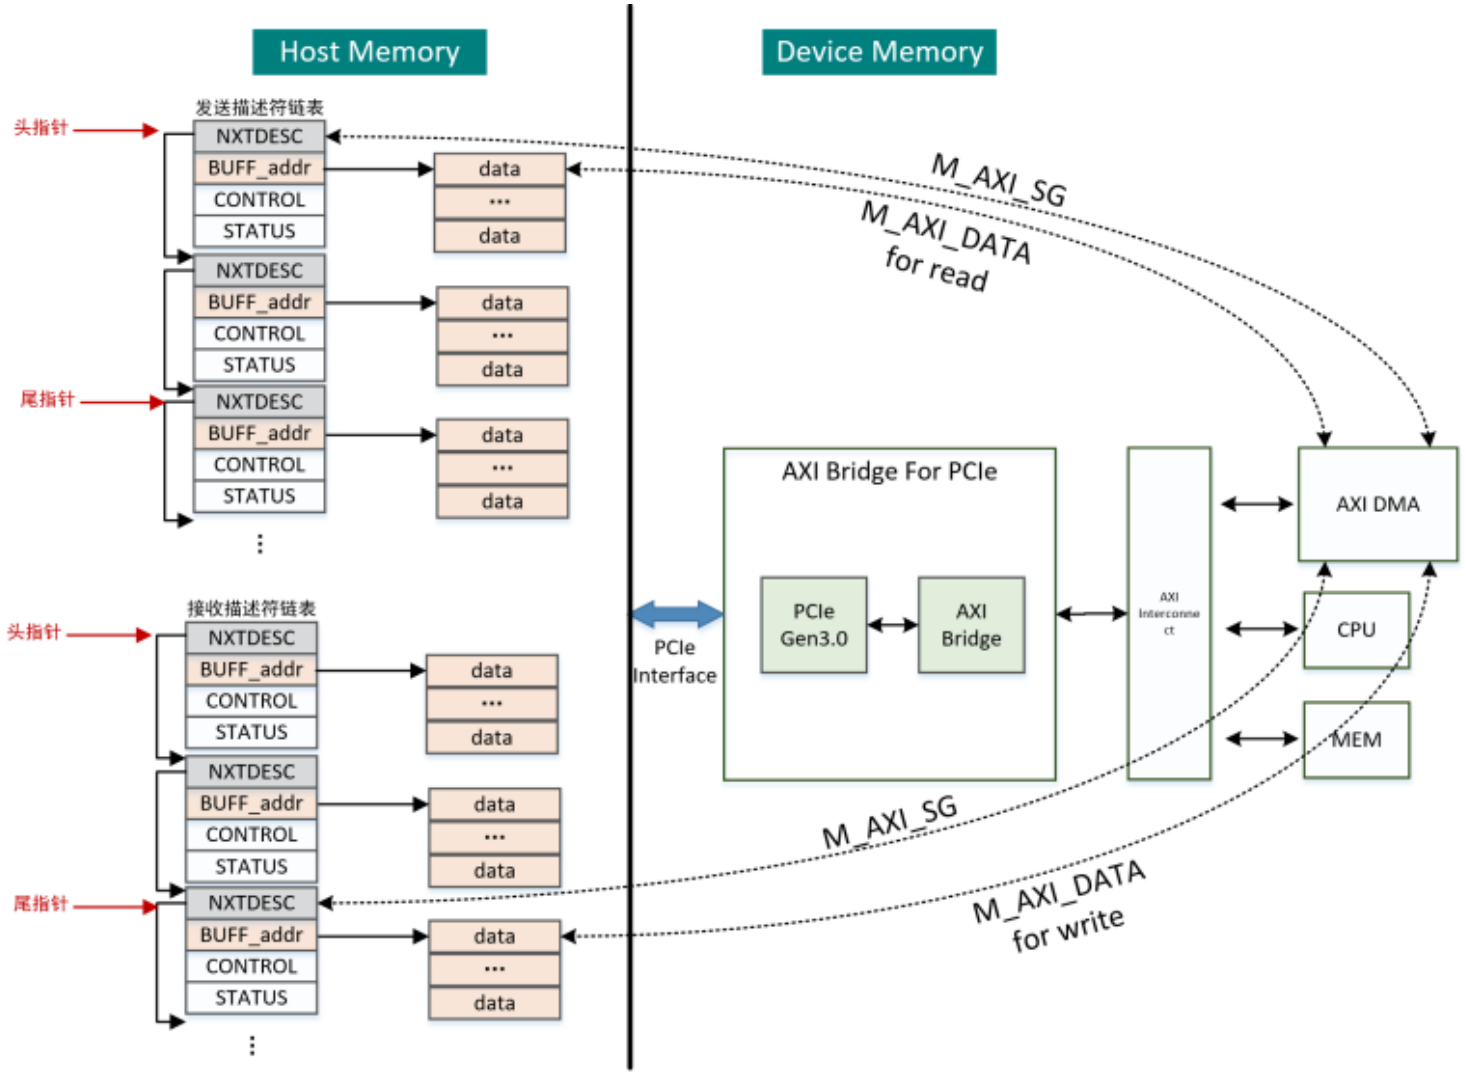

- 基于生产者消费者模型的多核+ dma cache一致性测试
- 可灵活配置dma搬运地址长度，core访问的地址和长度，支持数据正确性检测
- dma 支持multi-channel， sg dma和block dma模式
- dma支持D2M, M2M, M2D 数据搬运和正确性检测
- 中断测试模式：支持自主发送中断

IO通路的验证：网络验证

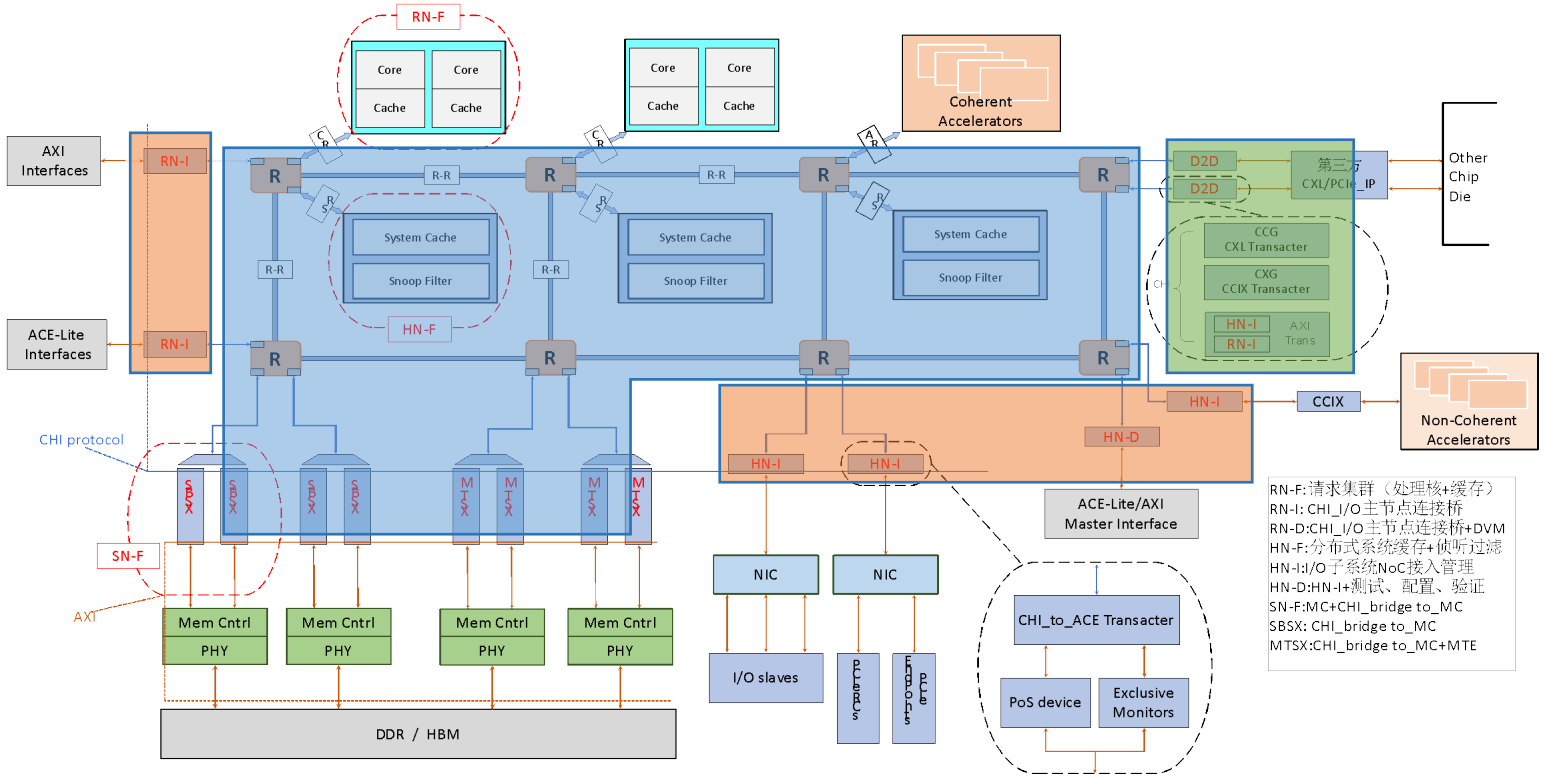


```
# while true;do iperf3 -c 192.168.1.20;sleep 1;done
[  892.831970] random: iperf3: uninitialized urandom read (37 bytes read)
Connecting to host 192.168.1.20, port 5201
[  892.901748] random: iperf3: uninitialized urandom read (131072 bytes read)
[  5] local 192.168.1.21 port 35368 connected to 192.168.1.20 port 5201
[ ID] Interval           Transfer     Bitrate         Retr  Cwnd
[  5]   0.00-1.96   sec  1.25 MBytes  5.35 Mbits/sec    0    127 KBytes
[  5]   1.96-3.89   sec  1.25 MBytes  5.43 Mbits/sec    0    134 KBytes
[  898.805286] random: crng init done
[  5]   3.90-5.74   sec  1.25 MBytes  5.71 Mbits/sec    0    134 KBytes
[  5]   5.76-7.66   sec  1.25 MBytes  5.50 Mbits/sec    0    134 KBytes
[  5]   7.67-9.51   sec  1.25 MBytes  5.70 Mbits/sec    0    134 KBytes
[  5]   9.53-11.42  sec  1.25 MBytes  5.52 Mbits/sec    0    134 KBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval           Transfer     Bitrate         Retr
[  5]   0.00-11.42  sec  7.50 MBytes  5.51 Mbits/sec    0             sender
[  5]   0.00-14.07  sec  7.50 MBytes  4.47 Mbits/sec                  receiver

iperf Done.
Connecting to host 192.168.1.20, port 5201
[  5] local 192.168.1.21 port 35372 connected to 192.168.1.20 port 5201
[ ID] Interval           Transfer     Bitrate         Retr  Cwnd
[  5]   0.00-1.86   sec  1.25 MBytes  5.65 Mbits/sec    0    134 KBytes
[  5]   1.86-3.83   sec  1.25 MBytes  5.32 Mbits/sec    0    134 KBytes
[  5]   3.84-5.89   sec  1.25 MBytes  5.12 Mbits/sec    0    134 KBytes
[  5]   5.90-7.91   sec  1.25 MBytes  5.21 Mbits/sec    0    134 KBytes
[  5]   7.92-9.93   sec  1.25 MBytes  5.24 Mbits/sec    0    134 KBytes
[  5]   9.94-11.90  sec  1.25 MBytes  5.36 Mbits/sec    0    134 KBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval           Transfer     Bitrate         Retr
[  5]   0.00-11.90  sec  7.50 MBytes  5.29 Mbits/sec    0             sender
[  5]   0.00-14.63  sec  7.50 MBytes  4.30 Mbits/sec                  receiver

iperf Done.
```

统一的ST验证环境

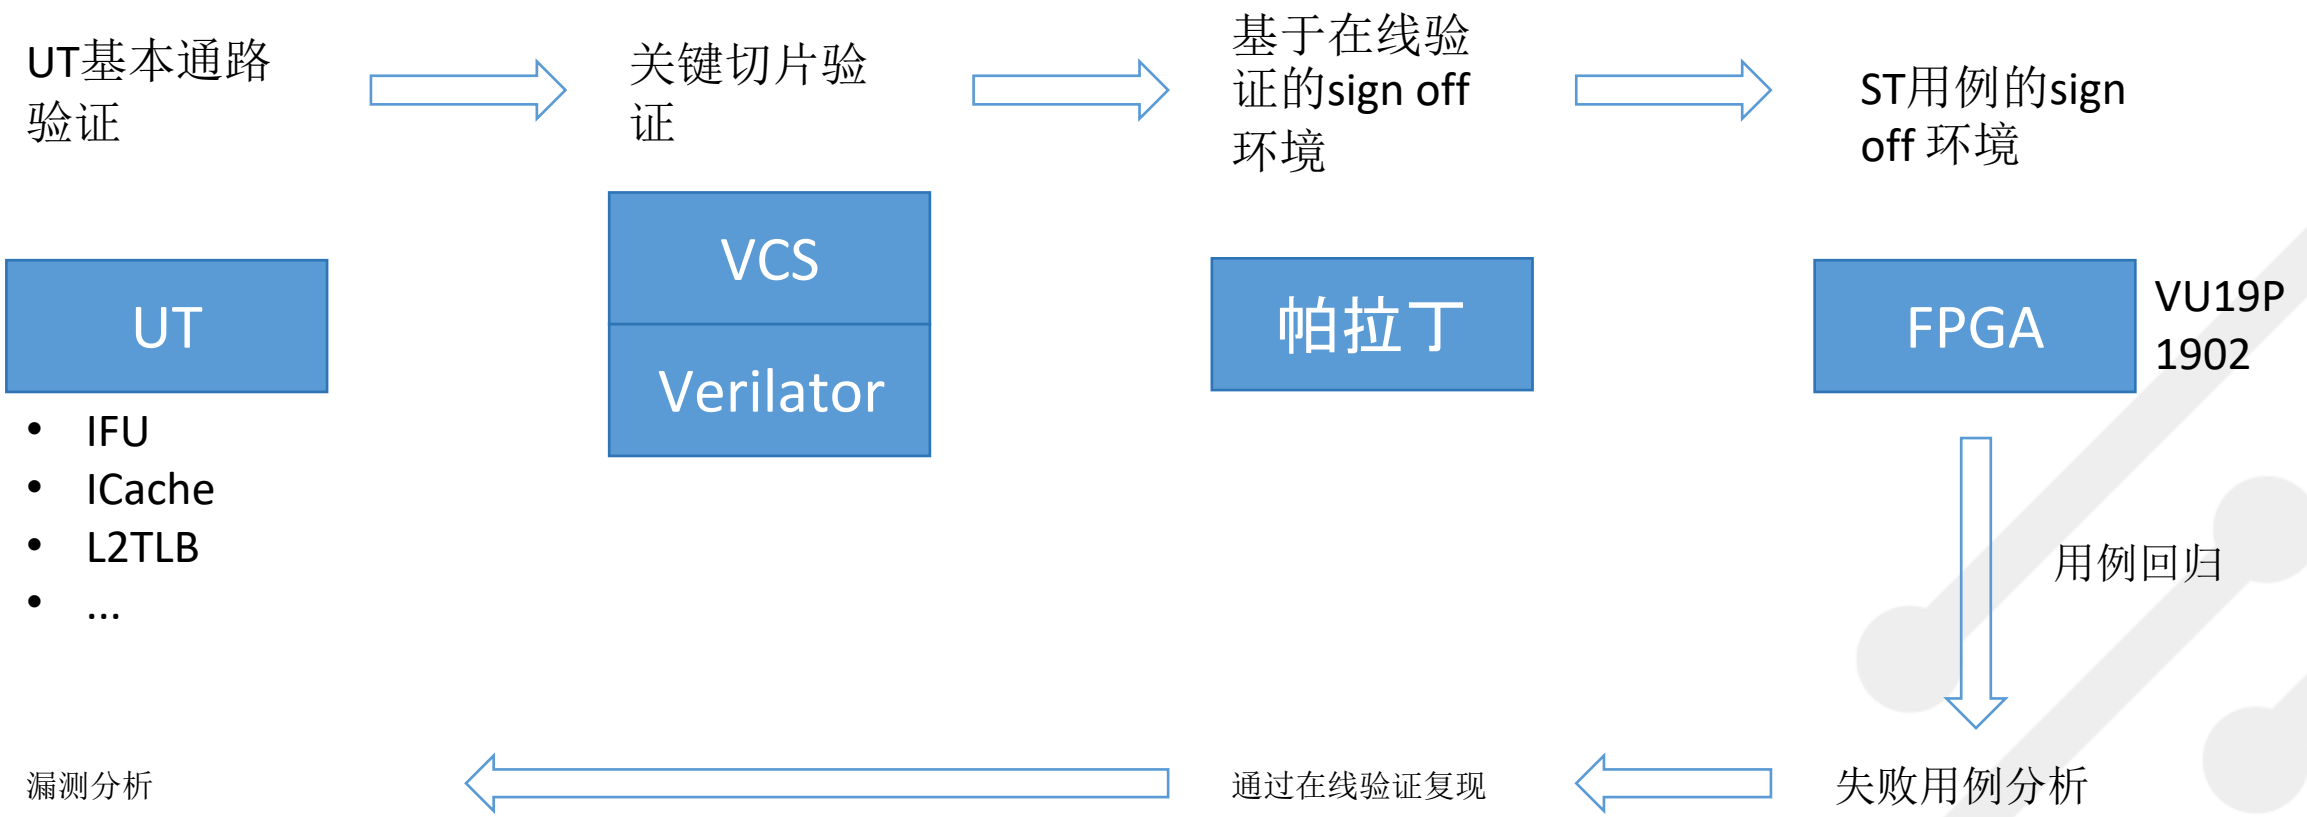

https://github.com/RV-BOSC/xs_build

北京开源芯片研究院
BEIJING INSTITUTE OF OPEN SOURCE CHIP

Thanks