Technology Report

Chinese Research in Processor Designs for HPC

Bob Sorensen, Alex Norton, Earl Joseph, and Steve Conway
April 2018

HYPERION RESEARCH OPINION

China is a leading contender to stand up the first exascale systems in the world with a planned availability in the 2020-2021 time frame. To meet that aggressive goal, Chinese HPC developers are focusing on the use of indigenous processors to power these systems. There are a number of indigenous processor developments in China destined for the HPC sector that span a wide range of processor types: some are wholly unique to a single Chinese developer, while others are based on existing schemes that include ARM and even an x86 variant.

Many of the processors currently under development in China targeted for the exascale systems have a long history that can trace their origins back to the early 2000's. Despite this committed effort to Chinese processor development, for most of that time the use of Chinese processors has been sporadic. It is only in the last few years that all of the major Chinese leadership=class HPC developers have turned to some form of indigenous chips for their systems.

- This shift has emerged for two main reasons: Chinese processors have largely closed the computational gap with counterparts from overseas - at least for the specialized architectures that Chinese HPC designers currently favor - and US exports controls are seen as being increasingly rigid in limiting Chinese access to key US components, a situation that Chinese HPC planners and developers want to avoid going forward.

Hyperion believes the development and use of Chinese processors for their pre-exascale and exascale systems offers a number of key advantages for the Chinese. By having a number of competing Chinese processor development efforts, Chinese leadership-class HPC developers - the current class that includes Sugon, NUDT, and Sunway - enjoy the advantage of being able to design a matched processor/system architecture that is finely tuned for selected workloads or use cases.

- As such, these processors can incorporate specific hardware features that may not be widely available from other potential processor suppliers due to the realities of market forces for such specialized and narrowly focused capabilities.
- However, the use of a custom chips engenders a heavy reliance on largely untested designs that may not meet performance expectations and that, at a minimum, create the need for a wide range of custom operating systems, middleware and applications software to take advantage of the specific unique hardware features.
- As such these Chinese systems cannot take full advantage of the wide range of open source HPC software currently and instead must independently develop their own custom counterparts.

Note: This page is intentionally blank.

IN THIS REPORT

China is investing significant resources in developing indigenous leadership capability across different technology sectors. One of their most active - and promising - effort to date is the development of a complete end-to-end indigenous high performance computing system. A critical element of this strategy involves developing enabling technologies such as processors, interconnect technologies, and systems software.

This document provides insight into the indigenous processor technology developments established in China with a focus on HPC systems. These technological leaps have been enabled through sustained investments in technology research and development for more almost two decades.

The current leading computational processor families explored in this assessment include:

- The ShenWei Processor Family
- The Loongson/Godson/Higon Processor Family
- The FeiTeng Processor Family
- Note: the appendix includes a number of previous and non-HPC processor efforts in China

SITUATION OVERVIEW

Hyperion Research data shows that the high performance computing market is dominated by the x86 processor today, with more than 90% of the current HPC market using processors based on the x86 architecture. Other processor architectures in use include Power, SPARC and ARM, but these are generally found in only a small number of the world's most powerful HPCs. Bucking this trend, researchers and scientists in China are working hard to develop indigenous processors that advance the state of the art in processor design and that are uniquely configured to support critical Chinese HPC applications and workloads. The sections below cover a snapshot of the major HPC-centric processor families in China today.

The ShenWei Processor Family

ShenWei (or Sunway) is a premier indigenous Chinese designed microprocessor that has been an integral part of the HPC systems designed by China's National Parallel Computer Engineering Technology Research Center and used in a number of leadership-class Sunway HPCs.

Currently, the most notable example is the Sunway TaihuLight installed at the National Supercomputer Center in Wuxi that has a peak performance of over 126 Pflops, and a Linpack rating of over 93 Pflops.

- This system is number one the November 2017 Top 500 list, and it incorporates over 40,000 ShenWei SW26010 processors, the most recent version of the ShenWei processor line, which was introduced in 2016.

The first ShenWei chip, the SW-1, was introduced in 2006 with a RISC design loosely based on the US-built DEC Alpha 2116. The ShenWei SW-1 had a single core, ran at 900 MHZ, and was manufactured at SMIC - a Chinese semiconductor fab house - using a relatively conservative 130 nm process.

- The second generation SW2 microprocessor was released in 2008, had 2 cores running at 1400 MHz, used 70-100 W, and was manufactured using SMIC's 130 nm process. The third

generation of the ShenWei SW3/SW1600 was released in 2010, and consisted of 16 core, 64-bit RISC design running at 975 – 1200 MHz The SW3 was manufactured using a 65 nm fabrication process and was capable of 140.8 GFlops running at 1.1 GHz.

The most recent generation of the Sunway line is the SW26010, a 64 bit RISC chip, introduced in 2016. The processor has 260 cores laid out as four powerful core groups (CGs) - connected via a network on chip (NoC) - each of which consists of a manage processing element (MPE) and 64 compute processing elements (CPEs) arranged in an eight by eight mesh.

- The MPE and CPEs are both 64-bit RISC, single-threaded cores working at 1.45 GHz and supporting 256-bit vector instructions with 32 vector registers.
- The MPE is designed to manage communications and scheduling tasks across the CPE computer cluster and to support interrupt functions and out-of-order execution.
- The MPE has two floating-point pipelines each of which performs 8 32-bit/64-bit flops per cycle with a peak performance of 23.2 GFlops per core. Each MPE has a 32 KB L1 data cache, a 32 KB L1 instruction, and a 256 KB L2 unified cache
- Each CPE has a dual pipeline, but only one of them issues floating-point instructions that performs 8 32-bit/64-bit flops per cycle, reaching a peak performance of 11.6 GFlops per CPE, providing a total of 2.969 TFlops for the combined 256 CPEs in a single processor.

FIGURE 1

ShenWei Chip Architecture

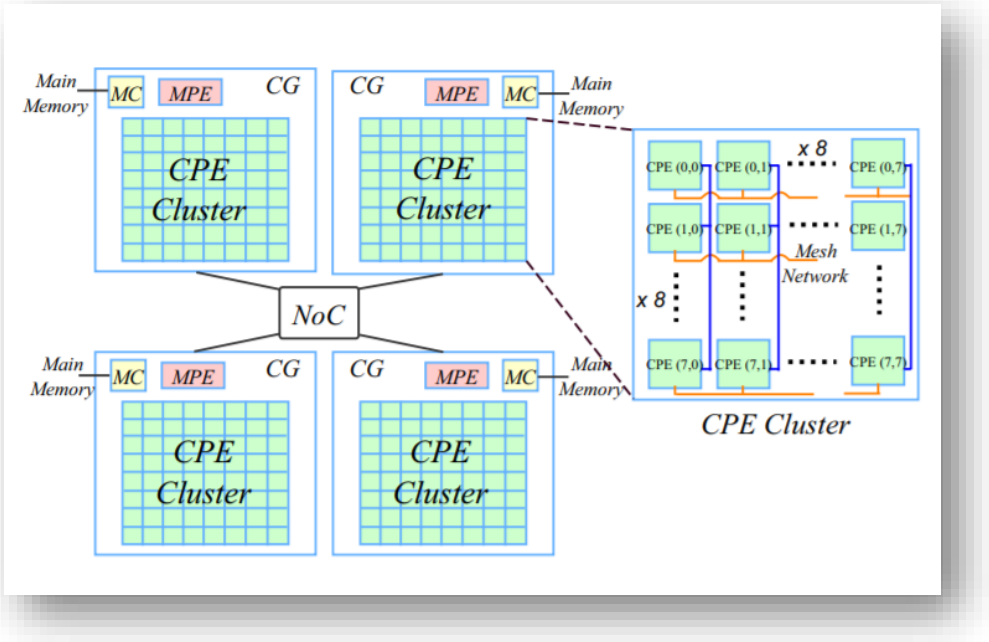

Source: James Lin et al, 2017 46th International Conference on Parallel Processing (ICPP), Hyperion Research, 2018

The Loongson/Godson/Higon Processor Family

The Loongson/Godson processors have been a staple of Chinese processor development with roots going back almost 20 years. However, recent events by Sugon to introduce a new processor into their pre-exascale system due out in the last half of 2018 signals a significant shift in processor design philosophy within this family.

The Loongson family of processors is a general-purpose MIPS64 processor developed at the Institute of Computing Technology (ICT) within the Chinese Academy of Sciences. Loongson development arises from a novel public-private partnership model where BLX IC Design Corporation designs and develops 64-bit Loongson processors while ST Microelectronics fabricates and markets the processor.

- Although the Loongson instruction set is MIPS64 based, the internal microarchitecture was independently developed by ICT. Since 2007, STMicroelectronics and Institute of Computing Technology have licensed technologies from MIPS Technologies.

The Loongson 1 (also called the Godson-232 core), introduced in 2002, was a 32 bit processor with a 266MHz clock speed featuring 8kB instruction cache, 8 kB data cache, and a 64 bit floating point unit. The processor was fabricated with 0.18 micron CMOS process and was capable of 200 double precision MFlops. The Loongson 1 was mainly targeted for embedded systems including point-of-sale systems.

- Loongson 2, which was released in early 2008, added 64-bit capability to the design. This Loongson processor featured a 4-way superscalar, out of order execution, 64 bit MIPS architecture processor core. The processor used a little-endian MIPS III compatible ISA comprising 5 execution units including 2 ALUs, 2 FPUs, and 1 address generation unit.
- Loongson 3, released in April 2010, was manufactured using 65 nm process and initially featured 4 cores operating at 1GHz consuming 15 W of power. Later versions of the Loongson 3 had 8 cores and utilized 40 W of power.
- More recent generations (Loongson 3B) featured an updated 8-core processor capable of running at 1.05 GHz with a peak performance of 128 GFlops double precision and 256 GFlops single precision. This was accomplished by adding two 256-bit vector-processing units in each core. These enhancements enable a peak performance of 8 double precision floating point fused Multiply-Add results per cycle or 16 GFlops per core at 1GHz.

Recently, the development arc of the Loongson chip shifted significantly when Chinese HPC maker Sugon (previously the Dawning Corporation that was spun out of the ICT as a commercial venture), established with the US chip maker AMD a joint venture called THATIC (Tianjin Haiguang Advanced Technology Investment Co., Ltd). to co-develop new chips based on the AMD x86 chip set

Essentially, AMD licensed its x86 processor technology and the intellectual property related to designing an SoC (system on chip) to Sugon support their development of a new Higon x86 CPU that Sugon plans to use in its pre-exascale prototype scheduled for operation the second half of 2018.

- Figure 2 is a recent series of technical specifications for a Sugon/AMD chip that may have features that will be in the processors bound for the Sugon pre-exascale system.
- Not much is known about the performance parameters of this new chip, but many experts indicate that it will be based on AMD's Zen microarchitecture, which is marketed under the brand name EPYC.

- The high-end 32 core EPYC option offers CPU, memory, I/O, Server, controller hub, and security in one package. As such, the EPYC architecture is geared primarily towards high-end cloud computing environments centered on a system on a chip architecture.
- It remains to be seen what kinds of modifications Sugon will make to the chip in order to tune its use for HPC workload environments.

FIGURE 2

Technical Specification of Recent Higon Chip

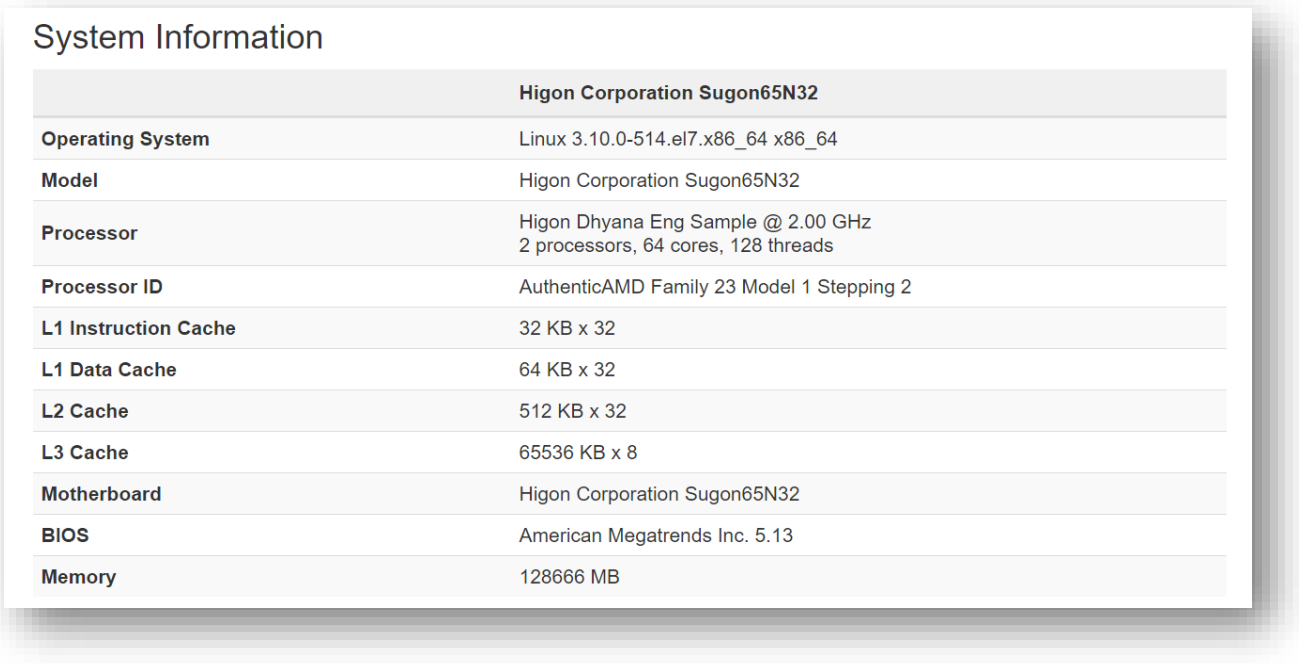

System Information

	Higon Corporation Sugon65N32
Operating System	Linux 3.10.0-514.el7.x86_64 x86_64
Model	Higon Corporation Sugon65N32
Processor	Higon Dhyana Eng Sample @ 2.00 GHz 2 processors, 64 cores, 128 threads
Processor ID	AuthenticAMD Family 23 Model 1 Stepping 2
L1 Instruction Cache	32 KB x 32
L1 Data Cache	64 KB x 32
L2 Cache	512 KB x 32
L3 Cache	65536 KB x 8
Motherboard	Higon Corporation Sugon65N32
BIOS	American Megatrends Inc. 5.13
Memory	128666 MB

Source: Geekbench.com, Hyperion Research, 2018

The FeiTeng Processor Family

The National University of Defense Technology(NUD) in Changsha has been developing processors for its internally-designed HPCs for almost 20 years. Its first processor, the Yin He FeiTeng series of high performance general purpose CPUs, was derived from a SPARC-based design but with an instruction set fully compatible with Intel's Itanium 2. Early iterations of the FeiTeng processors were implemented on a 130 nm process technology and delivered a peak performance of 16 GFlops while running at 500 MHz.

The second-generation FT64 chip used system-on-chip architecture comprised a general purpose CPU and a stream processor on the same chip. The processor was deployed as an accelerator in the Yin He High Performance Computing system.

- The usage model of FT 64 was similar to that of a GPGPU accelerator, with the FT64 running as a co-processor driven by a host CPU.

- Each HPC system board comprised a driver host FT64 processor, driving eight FT64 processors interconnected by an on chip host interface.

In contrast, the third generation FeiTeng1000 was designed as a traditional standalone general purpose processor. The 350 million-gate FeiTeng processor was manufactured with 65 nm process and capable of a clock frequency of 0.8 – 1 GHz.

- Although NUDT's Tianhe 1A supercomputer used 2048 FeiTeng 1000 processors, they were not part of the computational complement and instead were used in the system's front-end administrative server.

Going forward, NUDT now has two separate processor development lines: the FT 2000+, which is a general-purpose 256 core processors chip targeted for use in the soon to be available TianHe-3 and the Matrix 2000 accelerator used in the Tianhe-2A.

FT2000+

According to early reports, the FT2000+ is an ARM V8 processor with 64 cores capable of 615 GFlops running at 2.2-2.4 Ghz and using about 100 watts of power. Figure 3 outlines some of its key features.

FIGURE 3

FeiTeng 2000+ Processor Features

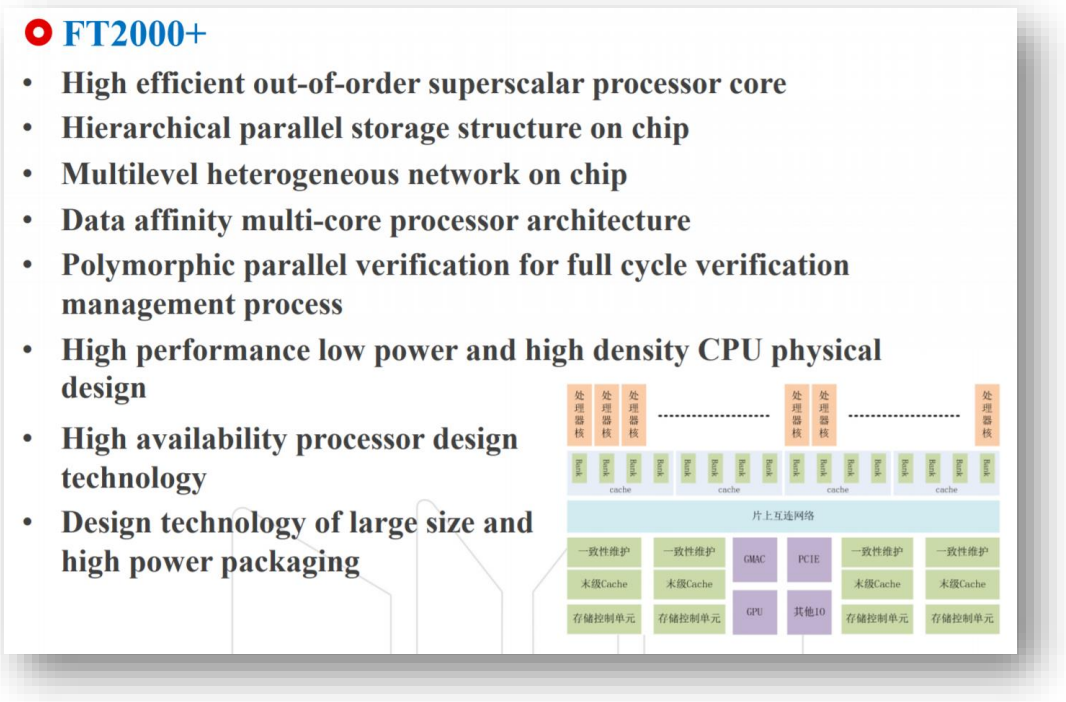

Source: NUDT, Hyperion Research, 2018

Matrix 2000 Accelerator

The Matrix-2000 is a 64-bit many-core processor designed by NUDT and introduced in 2017. This chip was designed exclusively as an accelerator for China's Tianhe-2 supercomputer in order to upgrade and replace the aging Intel's Knights Corner accelerators after the United States banned the sale of high-performance accelerators to China.

- The Matrix-2000 features 128 RISC cores operating at 1.2 GHz achieving 2.46 / 4.92 TFlops (DP/SP) with a peak power dissipation of 240 W.
- As seen in Figure 4, the Matrix-2000 consists 128 cores, eight DDR4 memory channels, and x16 PCIe lanes with four supernodes (SN) consisting of 32 cores each.

FIGURE 4

The Matrix 2000 Accelerator

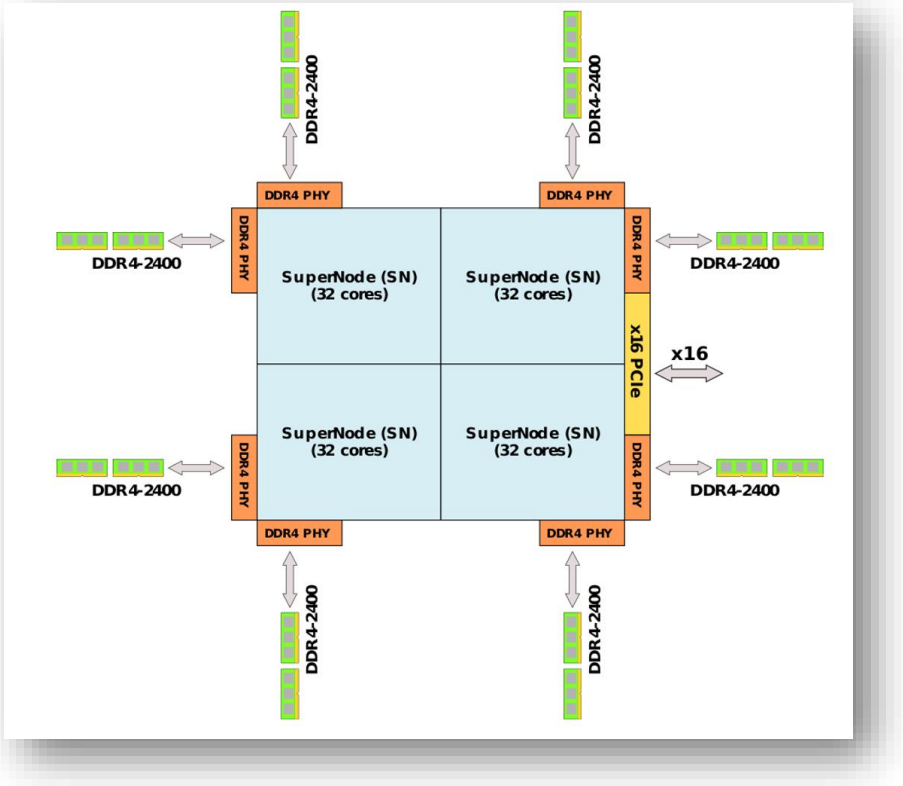

Source: NUDT, Hyperion Research, 2018

Other Potential HPC-Relevant Processors

There are two additional HPC-related CPU design efforts on-going in China that merit monitoring, although to date they have made little technical details public, and they have not yet been used in any HPC systems.

HiSilicon is a Chinese fabless semiconductor company based in Shenzhen, Guangdong and fully owned by Huawei. The firm produces processors based on ARM64 V8 that are used in Huawei servers.

- HiSilicon holds licenses from ARM Holdings for a broad range of ARM cores including the ARM Cortex-A9 MPCore, ARM Cortex-M3, ARM Cortex-A7 MPCore, ARM Cortex-A15 MPCore ARM Cortex-A53, and ARM Cortex-A57.
- HiSilicon works closely with Huawei to provide advanced processor solutions particularly for surveillance, set-top-dox digital TV and home connectivity, including CPUs, GPUs and DSPs.
- For example, HiSilicon has released world first 64bit CPU smart TV SOC and world first high performance cortex A73 CPU smart TV SOC.

Guizhou Huaxintong (HXT) is a joint venture between the People's Government of Guizhou Province and US semiconductor maker, Qualcomm, registered at Gui'An New Area, Guizhou Province, with operations and R&D centers in Beijing and Shanghai.

- HXT Semiconductor specializes in designing, developing, and selling sophisticated chips targeted for the server sector.
- The company plans to start shipping China-customized ARM server chips sometime 2018.

FUTURE OUTLOOK

Several diverse and well-funded efforts in China are working towards developing innovative and indigenous processor capabilities. As the global race towards exascale computing heats up, these diverse capabilities will almost certainly enable Chinese HPC suppliers to move closer to their goal of developing an end-to-end Chinese built HPC system.

- Processor approaches in China are diversified and incorporate a mix of older open processor architectures like MIPS, Alpha, SPARC, etc. while blending in the newer architectures like ARM and evolving legacy designs such as x86. To support these chips, Chinese scientists and researchers in collaboration with companies are also developing critical supporting middleware software including compilers and more.
- Chinese organizations will likely soon turn their attention to a number of key supporting technologies for interconnects, memory, and shortage, technologies that still are primarily based outside of China.

Hyperion believes the development and use of Chinese processors for their pre-exascale and exascale systems offers the Chinese a number of key advantages. By having a many competing Chinese processor development efforts, Chinese leadership-class HPC developers enjoy the advantage of being able to design a matched processor/system architecture that is finely tuned for selected workloads or use cases.

- These processors can include specific hardware features that may not be widely available from other potential processor suppliers due to the realities of market forces for such a specialized chip.
- However, the use of a custom chips engenders a reliance on largely untested designs that may not meet performance expectations and that, at a minimum, create the need for a wide range of custom operating systems, middleware, and applications software to take advantage of the specific unique hardware features.
- As such these Chinese systems cannot take full advantage of the wide range of open source HPC software and instead must independently develop their own custom counterparts adding, perhaps additional development complexity, time, and cost.

APPENDIX: PREVIOUS CHINESE PROCESSOR ACTIVITIES

The ShenWei Processor Family

ShenWei is a series of microprocessors developed by Jiangnan Computing Research lab. ShenWei microprocessors were primarily developed for military use of the PRC, and more recent versions of the ShenWei Family have been utilized in the indigenously developed Sunway Blue Light supercomputer. Initially, these microprocessors are based on the DEC Alpha design. The first generation ShenWei SW-1 microprocessor was released in 2006. The ShenWei SW-1 comprised a single core, running at 900 MHZ, and it was manufactured at SMIC's 130 nm process.

- The second generation SW2 microprocessor was released in 2008, comprised of 2 cores running at 1400 MHz, using 70-100 W, and manufactured using SMIC's 130 nm process.
- The third generation of the ShenWei SW3 was released in 2010, and consisted of 16 core, 64-bit RISC running at 975 – 1200 MHz. The SW3 was manufactured using a 65 nm fabrication process and is capable of 140.8 GFLOPS @1.1 GHz, and it was designed to support a quad channel 128 bit DDR3 bus, capable of addressing a maximum memory capacity of 16 GB.

The ShenWei SW1600 was the third generation microprocessor that was capable of achieving 140 GFLOPS using 16 cores based on RISC architecture, operating at 1.1 GHz. The SW1600 microprocessor featured a four-issue superscalar CPU architecture, with two integer and two floating point execution units, and a 7 stage integer and 10 stage floating point pipeline.

- The SW1600 was capable of addressing up to 8 terabytes of virtual memory and nearly 1 terabyte of physical memory. The processor features an L1 cache of nearly 8KB instruction cache and 8KB data cache, in addition to 96KB of L2 Cache.

The Loongson/Godson Processor Family

The Loongson family of processors is a general-purpose MIPS64 processor developed at the Institute Of Computing Technology, Chinese Academy of Sciences (CAS) in PRC. The Loongson was developed using a novel public-private partnership model where the BLX IC Design Corporation designs and develops 64-bit Loongson processors, while ST Microelectronics fabricates and markets the processor. Although the Loongson instruction set is a MIPS64 based, the internal microarchitecture was independently developed by ICT. Since 2007, STMicroelectronics and Institute of Computing Technology have licensed technologies from MIPS Technologies.

The Loongson 1 (Godson-232 core) was a 32 bit processor capable of 266MHz clock speed featuring an 8KB instruction cache, 8 KB data cache, and a 64 bit floating point unit. The processor was fabricated with 0.18 micron CMOS process, and it was capable of 200 double precision Mflops. The Loongson 1 was targeted for embedded systems including point-of-sale systems.

The Loongson 2, which was released in early 2008, added 64-bit capability to the Loongson 1. This Loongson processor featured 4-way superscalar, out of order execution, and a 64 bit MIPS architecture processor core.

- The processor utilized a little-endian MIPS III compatible ISA comprising 5 execution units including 2 ALUs, 2 FPUs, and 1 address generation unit (AGU). One of the 2 FPUs featured an integrated SIMD unit.

- The Loongson 2 processor family featured separate 64KB instruction and data caches (L1). Each chip also hosted a 512KB 4-way set associative L2 Cache and an integrated DDR memory controller. The Loongson 2E was capable of 1GHz at 7W of power.
- The Loongson 2F additionally featured an integrated simple video accelerator capable of delivering 1GHz while consuming 4W of power. Other processors built on the Loongson 2 platform included the Godson 2G and Godson 2H each of which is a 1 GHz processor, fabricated using a 65nm CMOS process, HW supported x86 binary translation, and on chip DDR2/DDR3 controller.

The Loongson3, released in April 2010, was manufactured using 65 nm process and initially featured 4 cores operating at 1GHz consuming 15 W of power. Later versions of Loongson3 featured 8 cores and utilized 40 W of power. The Loongsoon3 also featured DDR2/DDR3 DRAM support.

- More recent generations (Loongson 3B) featured an updated 8-core processor capable of running at 1.05 GHz with a peak performance of 128 Gflops double precision and 256 Gflops single precision. This was accomplished by adding two 256-bit vector-processing units in each core.
- These enhancements enabled a peak performance of 8 double precision floating point fused Multiply-Add results per cycle or 16 Gflops per core operating at 1GHz while delivering remarkable energy efficiency of 128 Gflops using 40 W.

The FeiTeng Processor Family

The Yin He FeiTeng series of high performance general purpose CPUs, was designed and developed by National University of Defense Technology in the Hunan province. The first generation of the YHFT processor adapted the Explicitly Parallel Instruction Computing (EPIC) architecture with an Instruction Set Architecture (ISA) fully compatible with Intel Itanium2. Early iterations of the FeiTeng processors were implemented on the 130 nm process technology and delivered a peak performance of 16 Gflops while running at 500 MHz.

The second-generation system-on-chip architecture comprised a general purpose CPU and a stream processor on the same chip. The second-generation processor was implemented under the code name FT64 (FeiTeng 64), and it was purported to be the world's first 64-bit stream processor dedicated for high performance computing.

- The processor was deployed as an accelerator in the Yin He High Performance Computing system. The usage model of FT 64 is very similar to that of a GPGPU accelerator, in that FT64 was designed to run as a co-processor, driven by a host CPU. Each HPC system board comprised a driver host FT64 processor, driving eight FT64 processors interconnected by an on chip host interface.
- The system architecture also consisted of separate processors and dedicated FT64 memory. The host FT64 also had its own dedicated host memory.
- The systems developers of the FT64 also developed a stream programming language called SF95 which added 10 compiler directives to FORTRAN95.

The third generation FeiTeng1000 is a SPARCv9 compatible processor that was designed and produced in Hunan province in China. Unlike its first and second generation cousins, the third generation FT1000 was designed as a traditional standalone general purpose processor. The 350 million-gate FeiTeng processor was manufactured with 65 nm process and was capable of a clock frequency of 0.8 – 1 GHz.

- The 8-core SPARC based processor was capable of executing 64 threads and the processor hosted 3 Hyper Transport channels, 4 DDR3 memory controllers capable of a memory bandwidth of 32 GBps. The FT 1000 was designed to support 8 lane PCIe 2.0 interfaces and 3 CPU direct interfaces.
- The Tianhe 1A supercomputer deployed 2048 FeiTeng 1000 processors.

The Unity/PKUnity Processor Family

The indigenously developed Unity line of processors utilized a novel RISC based instruction set architecture designed by the Microprocessor Research and Development Center (MPRC) of Peking University in China. The Unity line of processors was built on a fully functional SoC utilizing architecture similar to ARM but leveraging a different instruction set named Unicore.

- The Unicore instruction set is identical to that of a standard ARM instruction set except that conditional execution have been removed and the register specifier bits have been reassigned.

The Unity Line of processors evolved over multiple generations of research and development, and it is also known is some literature as PKUnity. The first generation Unicore 16 microprocessor was developed in 1999. Unicore 16 microprocessor was designed and developed for embedded processors scenarios. The second generation UniCore32 was developed in December 2000. The Unicore 32 was a RISC based microprocessor which featured a 32/16 bit mixed instruction set.

- In addition to the Unicore32 microprocessor, MPRC also developed the supporting software system including BIOS, operating system, simulator, debugger and full compilation tool-chain.
- In 2001 MPRC developed a 64bit floating point co-processor Unicore-F64, along with an optimized Unicore 32 and other indigenously developed controllers including PCI, SDRAM, 10/100 Ethernet MAC and DMA controllers.

MPRC designed the Unity-863 SoC (PKUnity 1) in 2002 largely building upon the prior engineering efforts in Unicore32 and Unicore46. The Unity-863 SoC was developed with 250nm techniques contained nearly 8,000,000 transistors. The Unity-863 SoC processor runs at a frequency of 200MHz, with less than 800mW power dissipation.

In 2003 an improved version of Unity-863 (PKUnity 1) was delivered, the key enhancement being utilization of 180nm fabrication process. The Unity-863 (PKUnity 1) CPU SoC features a Unicore32 processor utilizing RISC architecture with a 5 stage pipeline, capable of 300 MHz. The unicore32 processor part of the Unity863 features improved DSP processing capabilities.

The PKUnity SoC series include PKUnity1 SoC (released 2001), Unity 805-1 (released 2002), PKUnity-2 (released 2003), Unity 805-2(released 2004), PKUnity-3(released 2008), PKUnity-4/4D (released 2010).

- The Unity 805-1/2 SoCs integratedultra low power versions of UniCore-1 CPUs.
- PKUnity1 utilized a 250nm TSMC fabrication process to develop a SoC processor that was 6.5x6.5 mm2 in size. The PKUnity1 utilized 4 W and contained nearly 8M transistors with a pin count of 376.
- PKUnity2 utilized 180nm TSMC fabrication process to develop a SoC processor capable of running at 300 MHz while consuming only 3.5 W of power.
- The PKUnity-2 SoC was developed on a die size 4.8x4.8 mm2 consisting over 12 M transistors and 376 pins.

- PKUnity3 utilized 130nm TSMC fabrication process to develop a SoC processor capable of running at 700 MHz while consuming 3.0 W of power.
- The PKUnity-3 SoC was developed on a die size 9.2x7.3 mm2 consisting over 55 M with 976 pins.
- Unity 805-1 utilized a 250nm TSMC fabrication process to develop a SoC processor capable of running at 100MHz. The processor had a die size of 4x4 mm2 with a 4M-transistor count, and 208 pin count.
- Unity 805-2 utilized 180nm TSMC fabrication process to develop a SoC processor capable of running at 200 MHz. The processor had a die size of 4.5x4.5 mm2 with an 8M-transistor count, and 256 pin count.

The ARCA Processor Family

ARCA Technologies started out in 2001 through investment by a Chinese Government funded venture. ARCA (Advanced RISC Computer Architecture) Technologies indigenously designed and developed the Arca series of processors starting with Arca1. Arca Technologies initial approach towards processors was unique in that Arca sought to design both the processors and develop an indigenous Instruction Set Architecture known as the ArcaISA.

- The ArcaISA is predominantly RISC based comprising of over 80 instructions with each instruction consisting of 3 operands and contains 32 general purpose registers.

The Arca-1 processor was designed using a 0.25 micron fabrication process. The Arca-1 processor had a 5-stage pipeline and used 1.2 W of power at the peak clock speed of 166MHz.

- The Arca1 also contained separate 32 way associative 8k instruction cache and 8k data cache. Arca1 also hosted a DSP unit that utilized a pair of multiply-add units in addition to SIMD support for media acceleration.

The second generation of Arca processors, the Arca2, was released in December 2002. The Arca2 processor was manufactured using a 0.18 micron fabrication process on a die size of 5 mm2. Arca2 was able to operate at 400 MHz while utilizing 400 mW power, which made it ideal for embedded devices. The size of instruction caches and data caches in Arca2 remained the same at 8KB.

In 2004 Arca began development of a faster processor codenamed Arca-3. Arca-3 was designed using the same 0.18 micron fabrication process and could clock up to 600 MHz. In order to achieve this level of performance, the chip utilized 7 stage pipelines instead of 5 and utilized branch prediction to offset branch penalties from longer pipelines. The instruction cache was doubled to 16k as well as the data cache which was also increased to 16k.

However, the development of this processor was impacted by non-technical factors including scandals and corruption. In 2006 Arca purchased ARM 926EJ-S license in a remarkable shift from designing indigenous ISAs and architectures towards licensing and building on existing architectures.

Other Chinese Processor Efforts

There were several other indigenous processor developments beyond the ones listed above:

- The Hangzhou-based NationalChip licensed ARM in 2008 and could deliver processors for consumer grade devices including display devices and digital entertainment.
- NationalChip's processor offering included the GX 1100, 1200, 1500, and 3000 families. The GX series feature SoC based architecture that were optimized for entertainment devices.

- The Fuzhou-based RockChip pioneered an ARM Cortex A8-based custom CPUs and SoC chips for consumer grade devices.
  - One RockChip offering, the RK29xx, was the first chip to decode Google's WebM VP8 in hardware. The CPU was capable of 1.2 GHz with 512 KB L2 cache. The processor also featured an integrated 60 million polygons/s GPU as well as DSP-accelerated 1080p playback and encoding in most formats. RK29xx supports tablets and smartphones with up to 1280×800 displays.
- Other processor companies developing ARM based processors include: LeadCore Technologies, AllWinner, Brite Semiconductor Corp. and Shanghai InfoTM microelectronics.

ABOUT HYPERION RESEARCH, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user and vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com