GS464V: A High-performance Low-Power Processor Core with 512-bit Vector Extension

Weiwu Hu, Yunji Chen
Institute of Computing Technology, Chinese Academy of Sciences

***Abstract*—As the processor core of the next generation Godson microprocessors, GS464V targets high performance computation with low power consumption. It integrates two 256-bit vector processing units (VPUs) so as to complete 8 double-precision floating-point MADD, or 16 single-precision floating-point MADD per clock cycle. To fully take advantage of the computational power of GS464V in many important applications (e.g., scientific computation, signal processing, multimedia, et al.), a dedicated vector instruction set is also proposed. Furthermore, in additional to the normal memory functional unit, GS464V employs a programmable memory access coprocessor to both tackle the performance bottleneck of memory access and reduce power consumption.**

The first microprocessor adopting GS464V, which is named Godson-3B, has been taped out with 65nm process. Godson-3B integrates 8 GS464V cores, thus its peak performance achieves 256/128 GFlops (single/double precision) at 1GHz frequency. The power consumption of Godson-3B is about 20 Watt depending on applications. Godson-3C, which integrates 16 GS464V cores to achieve 768/384 GFlops (single/double precision) peak performance at 1.5GHz frequency, is also in the physical design phase.

I. Introduction

Godson project is a national project of China undertaken by Loongson Technologies Corporation Limited and Institute of Computing Technology, Chinese Academy of Sciences. As the high-end production of Godson project, Godson-3 is a series of high-performance Chip Multi-Processor (CMP) targeting scientific computing, high-end embedded applications, desktop applications and so on. In this paper, we will introduce the microarchitecture of GS464V, which is the processor core of the next generation Godson-3 micro-processors (including Godson-3B and Godson-3C).

II. Microarchitecture of GS464V

GS464V is a high-performance low-power processor core upgraded from GS464 (the processor core in Godson-3A). Similar with GS464, GS464V is MIPS64-compatible, and provides x86-emulation hardware support. As shown in Fig. 1, GS464V adopts a nine-stage dynamical pipeline, where the nine stages are responsible for fetching, predecoding, decoding, register mapping, issuing, reading register, execution, writing back and committing respectively. As a four-issue superscalar processor core, the widthes of fetching, predecoding, decoding, register mapping and committing pipeline stages are 4, while the widthes of other stages are 5 to support the 5 functional units (two fix-point functional units ALU1 and ALU2, two floating-point/vector functional units VPU1 and VPU2, and one memory functional units).

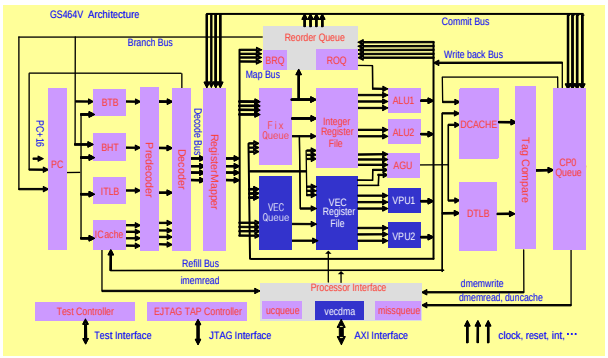

Fig. 1. The microarchitecture of GS464V. GS464V adopts a four-issue nine-stage dynamical pipeline. The blue blocks (vector queue, vector register file, vpu1, vpu2, and vecdma) are designed for vector extension.

The out-of-order execution mechanism of GS464V is a combination of the register renaming, dynamic scheduling, and branch prediction techniques. It has a 64-entry fix-point physical register file, a 16-entry fix-point reservation station reservation station issuing instructions out-of-order, and a 64-entry ROQ (reorder queue) committing out-of-order executed instructions in the program order. The branch prediction support of GS464V include a 16-entry BTB (branch target buffer), an 8K-entry BHT (branch history table), a 9-bit GHR (global history register), and a 4-entry RAS (return address stack).

The memory functional unit of GS464V is similar with GS464. It supports one 128-bit memory access per clock cycle. It has a two-way associative 32-KB L1 data cache, a fully associative 64-entry TLB, and a 16-entry load/store queue which supports out-of-order memory accesses, non-blocking cache and load speculation. When a load/store instruction misses in L1 data cache, GS464V can fetch data from L2 cache or memory controller through the 128-bit AXI interface.

III. Vector Extension of GS464V

The most attractive feature of GS464V may be its 512-bit vector extension. The blue blocks in Fig. 1 are the modules designed for vector extension. In each clock cycle, a 28-entry vector queue can out-of-orderly dispatch two vector instructions to the two 256-bit vector processing units (VPU1 and VPU2). Each VPU can perform 4 double-precision floating-point multiply and add (MADD) operations, or 8 single-precision floating-point MADD operations simultaneously,

or at most 32 fix-point operations. Furthermore, to keep the compatibility with MIPS64, each VPU can also perform one MIPS-64 floating-point instructions per clock cycle.

To fully take advantage of the vector computational ability, GS464V defines a vector instruction set considering the requirement of important applications (e.g., scientific computation, signal processing, multimedia, et al.) as well as vector compiler. It is worth noting that the vector instruction set of GS464V includes a series of shuffle-computation mixed instructions to reduce shuffle instructions. Traditionally, vector programs may include many shuffle instructions to reorganize the data residing in vector registers. Thus, many previous processors with vector unit adopt dedicated shuffle unit, which consumes the width of superscalar and incurs additional power consumption. The shuffle-computation mixed instructions of GS464V eliminate the necessity of dedicated shuffle unit (although the VPU of GS464V can still perform shuffle operation). For example, there can be no shuffle instruction in matrix multiplication and fast fourier transform (FFT) implemented on GS464V. Hence both the number of instructions and power consumptions are reduced.

Another important issue about processor with vector unit is the memory bandwidth. The computation ability of vector extension can not be fully explored without a corresponding promotion of memory bandwidth. On one hand, to relieve the burden of memory bandwidth through exploiting the spatial locality, GS464V employs a 128-entry 256-bit vector register file with 4 write and 8 read ports, which can hold 4KB data; on the other hard, GS464V employs an additional programmable memory access coprocessor (the vecdma block in Fig. 1) which is independent with the normal memory unit. The non-reusable data can be automatically and directly transferred between vector register file and L2 cache/memory controller by the memory access coprocessor, while reusable data can reside in L1 data cache and be accessed by normal load/store instructions. Thus, the memory bandwidth is enlarged. Concretely, as shown in Fig. 2, the memory access coprocessor integrates three channels named Godson Super-Links (GSLs), to flexibly support memory access flow. Each GSL can automatically move a data flow with arbitrary size between the vector register file and exterior memory based on the configured direction, buffer addresses, buffer size, unit size, stride, length of GSL. In addition, GSL can automatically calculate its accessing memory address with some address mapping patterns (e.g., matrix transposing, bit reverting, et al.) which are common in scientific computation and signal processing. Furthermore, the three GSLs can work in parallel, and the flow controllers can arbitrate between them based on their priorities, latencies, bandwidths given by control registers. To keep the synchronization between GS464V and the memory access coprocessor, a group of cyclic dma_end signals are introduced as the handshake mechanism: The coprocessor may set a dma_end signal to inform GS464V when it has transferred a given size of data; GS464V can branch based on the value of a dma_end signal, and clear a dma_end signal to inform coprocessor that a data block has been processed.

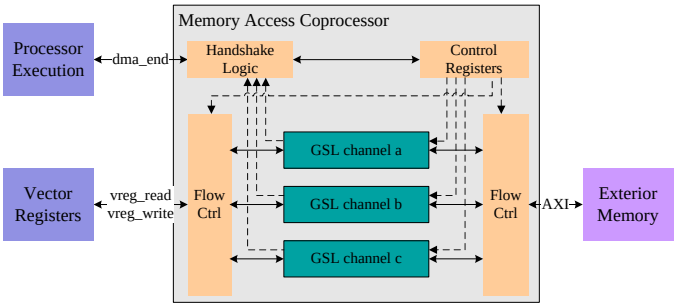

Fig. 2. Flow graph of the memory access coprocessor. The three Godson Super-Links (GSLs) can transfer three flows with arbitrary size between vector register file and exterior memory.

Benefited by the programmable memory access coprocessor, the vector computational ability of GS464V can be fully utilized in many important applications. For example, GS464 can achieve $\sim 93\%$ efficiency for matrix multiplication, and $\sim 87\%$ efficiency for FFT.

IV. GS464V-BASED CMP

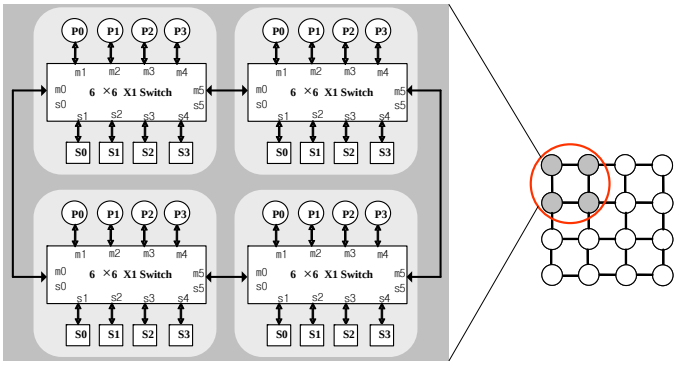

Fig. 3. Overall Architecture of Godson-3C. There are four nodes in Godson-3C. The latency of one hop on each mesh node is 4 cycle. Each node contains four GS464V cores, four L2 cache banks, one HyperTransport controller, one DDR2/3 memory controller, and crossbars which connect these components together.

Several CMP chips with different process technologies are in progress based on GS464V. Godson-3B, which integrates 8 GS464V cores, has been taped out with 65nm process. Its peak performance achieves 256/128 GFlops (single/double precision) at 1GHz frequency. The power consumption of Godson-3B is about 20 Watt depending on applications. Thus, the performance/power ratio of Godson-3B is about 6.4GFlops/Watt, which is higher than most stat-of-art high-performance microprocessor. The overall area of Godson-3B is 299.8 square millimeter, while a GS464V core is 15.5 square millimeter. The single-core version of Godson-3B, which is named Godson-2H, will be taped out in 2010 with 65nm process. Furthermore, Godson-3C, which integrates 16 GS464V cores (as shown in Fig. 3), is in the physical design phase, and will be taped out with 32nm process. At 1.5GHz frequency, Godson-3C can achieve the peak performance of 768/384 GFlops (single/double precision).