

Analyzing the performance of the StreamEngine® 5000 Processor Family versus traditional RISC architectures in communications & media applications

Prepared by Ubicom Inc.

June 2006

Introduction

Historically, performance has been the key differentiator for processors. Most commonly, the clock frequency in megahertz (MHz) has been the scale used in measuring processors' performance. However, next generation of processors with more efficient architectures are changing this old and familiar notion. An excellent example is the Ubicom **StreamEngine** family of multithreading processors. Benefiting from a memory to memory instruction set architecture, packets can stream directly to on-chip memory and be processed to completion very efficiently by any of the 10 independent threads, and a real time operating system optimized for the environment. This means that their performance characteristics are quite different from traditional processors such as ARM or MIPS running an embedded operating system or Linux.

The objective of this paper is to provide guidance on how to predict the performance of these different architectures in communication & media applications.

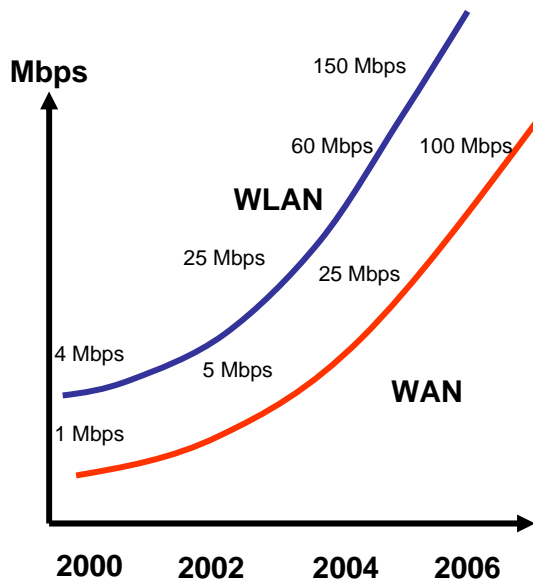
Comparing and Measuring Performance

Every network link has a peak wire-speed (physical link) performance. However, the actual usable throughput is reduced from this peak by network and protocol overheads. These overheads vary from link to link and protocol to protocol; therefore, comparing wire speeds is not a

good way to predict end-to-end performance.

Some protocols such as Ethernet are full duplex (simultaneous transmit and receive), and others like IEEE802.11g are half duplex (either transmit or receive). Obviously, this makes a two times difference in throughput. In this paper all network throughput numbers are end-to-end, measured by the industry standard Chariot measurement tool, doing multiple TCP file transfers, one for each link and direction. These streams have two 1500-byte data packets for each 64-byte TCP ACK packet, so the average packet size is about 1000-bytes. For example, Fast Ethernet is a 100 Mbps wire speed technology, but it is full duplex and has low protocol overhead. Chariot measures a peak throughput of 185 Mbps. 802.11g has a peak radio rate of 54 Mbps, but it has high protocol overhead, and Chariot measures only 24 Mbps of actual throughput. Comparing physical link performance, we might think that 802.11g is 54% of Fast Ethernet speed, but in fact it's only 13% (24 Mbps vs. 185 Mbps) as fast as Ethernet.

Ubicom provides advanced tools for detail measurement of CPU performance characteristics. Another reference source is "Computer Architecture, a Quantitative Approach", by Hennessy and Patterson (referenced as H&P below) which has been used throughout of this paper.



Home Networking Performance Requirements

Multimedia Applications are Fueling the Need for More Performance

The performance requirement in home routers has increased much faster than Moore's law during the past five years. The WAN interface sets the required routing performance. It increased from 1 Mbps for ADSL to 5 Mbps for ADSL2, to 25 Mbps for triple play routers that will be deployed this year. In Japan and Korea there are deployments above 50 Mbps, and fiber to the home requires 100 Mbps or higher. Wireless radio throughput increased from 4 Mbps for 802.11b, to 25 Mbps for 802.11g, to 60 Mbps for turbo-g, and will be over 150 Mbps in 2006 with 802.11n. The total WAN plus WLAN throughput has increased from 5 Mbps in late 2000, to 250 Mbps in mid 2006, a 50x increase in required performance. This is a 100% per year growth rate. An ARM7 had sufficient performance in 2000, but in 2006 a 500 MHz MIPS core or StrongARM is not fast enough.

Qualitative Performance Characteristics

The Ubicom, ARM, and MIPS processors all execute 1 instruction per clock peak, so the perception would be that their performance is completely

determined by the clock frequency. However, this is not the case. There are additional factors to consider:

1. Architectural efficiency
2. Memory time
3. Hazards

The architectural efficiency is the number of instructions it takes to complete a typical task or benchmark. If one instruction set uses fewer instructions, it will complete the task sooner, even if the clock frequency is the same. The other factors are lost time per instruction when the processor has to wait for instructions or data from off-chip memory, and lost time due to hazards, where one instruction has to wait for the results of a previous instruction.

Here is the equation to calculate performance:

$$CPU_Performance = \frac{clock_rate \times architecture_efficiency}{1 + memory_time + hazard_time}$$

Software Efficiency

On the same processor, there might be two different implementations of the same application, with very different performance. A more efficient software implementation uses fewer instructions to do the same amount of work. Measured router performance indicates that Ubicom ipOS uses fewer instructions than Linux in router applications. For any application, it is possible to tune the software to increase its efficiency and improve performance, but this adds time to market. IpOS was designed from the start for fast packet processing, and includes patented software techniques to process packets as efficiently as possible. Linux, however, was designed to be a general purpose operating system, so it is not surprising that it is slower at routing.

$$\text{System_Performance} = \frac{\text{CPU_Performance}}{\text{instructions_executed}}$$

The full performance equation is shown at the foot of this page

Product	Min	Max
MIPS 4K	170	230
MIPS* 24K	300	333
ARM9E**	150	210
Ubicom IP3023	250	325
Ubicom IP51xx	275	350

CPU Frequency Comparison (0.13u Technology)

* TSMC web site

** ARM web site

Clock Rate

The Ubicom IP3023 has a deep 9 stage pipeline, which gives it a higher clock rate than most MIPS or ARM implementations. The IP3023 in 0.13u is available in two frequencies, 248 MHz and 325 MHz. The IP51xx processors in 0.13u are available at 275 MHz and 350 MHz. The MIPS web site gives 210 to 255 MHz for an 0.13u MIPS 4K implementation, but actual implementations range from 170 MHz to 230 MHz. The ARM9E core in 0.13u is 150 to 210 MHz.

The MIPS 24k core has a deeper pipeline than the MIPS 4K, and can reach 330 MHz in 0.13u. This is a new core and is not in production as of late 2005.

More Efficient Instruction Set

MIPS, ARM, Power and other RISC processors have a load/store instruction set that means data must be moved to registers before being used. Moving data from one location to another takes two instructions (a load and a store). This is a good

choice for servers, but not optimized for packet processing (a packet is a sequence of bytes in memory). For processing, the CPU scans the packet, makes minor changes, and perhaps copies it. For packet processing a memory to memory instruction set is the most efficient architecture.

The Ubicom instruction set operates directly on data in memory, and can read and write memory in the same instruction. That means it can move data using one instruction per word, twice the rate of a RISC CPU.

In applications running on RISC CPUs, 35% to 45% of the instructions are LOAD or STORE (H&P). Routing applications are at the high end of that range. Ubicom CPUs running routing applications have only 30% of the instructions moving data to/from CPU registers. 35% of the instructions combine some other operation with a memory access.

MIPS CPUs do not have an auto-increment addressing mode. This means that scanning a packet requires a LOAD to get the data, and an ADD to increment the address. Ubicom's instruction set can perform an address increment with most memory accesses. This is why typical MIPS code has 15% to 25% ADD instructions, and Ubicom code has only 6% ADDs.

Overall, the Ubicom CPU executes only 75% to 80% of the number of instructions of a MIPS CPU, and is 1.25 to 1.3 times more efficient.

On-Chip Memory

All Ubicom routing code executes at full speed from on-chip memory. ARM and MIPS applications run from off-chip DRAM, through a cache, and spend 30 to 60 percent of their time waiting for cache misses. IP30xx applications keep all packet

$$\text{System_Performance} = \frac{\text{clock_rate} \times \text{architecture_efficiency}}{\text{instructions_executed} \times (1 + \text{memory_time} + \text{hazard_time})}$$

data in on-chip memory, so there is never a cache miss. IP51xx applications keep most packets on-chip, but move some packets off-chip when required.

For example, say a 170 MHz MIPS core is connected to a 133 MHz SDRAM through a cache. On average each instruction makes 1.4 memory references (one for the instruction, and the rest for data). Each memory reference to DRAM takes 20 clocks. If the cache has a 3% miss rate, the average instruction will spend 0.8 clocks waiting for memory, slowing performance from one clock per instruction to 1.8 clocks per instruction.

Off-chip memory used by ARM or MIPS reduces the performance, increases the cost, and consumes additional power, compared to IP3023. As the ARM or MIPS CPU core goes to higher frequency, it takes more clocks to access the DRAM, and the memory penalty goes up.

Actual measurements of a 170 MHz MIPS CPU/133 MHz DRAM router implementation measured 11 million memory accesses per second (0.6 clock memory penalty). A 180 MHz MIPS CPU/90 MHz DRAM low speed router does 5.5 million memory accesses per second (1.0 clock memory penalty). These CPUs had 32 KB of cache (the maximum for this core). Processors with smaller caches would have larger penalties.

The memory penalty depends on the cache size and application.

Memory time per instruction

IP30xx: 0
 IP51xx 0 to 0.1
 MIPS: 0.5 to 0.75
 ARM9: 0.5 to 0.75

Pipeline Hazards and Multithreading

Any pipelined CPU has pipeline hazards, where a later instruction requires data that is not yet avail-

able from an earlier instruction. When this happens the pipeline must stall, reducing performance. Taken and mis-predicted branches are the largest hazards. RISC CPUs have a hazard when data from a LOAD instruction is not available before it is used. The Uvicom CPUs do not have a LOAD-use hazard, since the operand memory fetch is part of the instruction.

The Uvicom processors have a deep pipeline, which causes a large penalty for mis-predicted branches. The instruction set includes static branch prediction to reduce mis-predictions. The compiler and profiler enable OVER 75% accurate static branch prediction.

Multithreading in the Uvicom processors hides these penalties, and when four or more threads are simultaneously active, the pipeline penalties are negligible. The initial ipOS operating system uses two threads, so the measured average hazard is 0.3 clocks per instruction. Future OS releases will use more threads and dramatically reduce this penalty.

Hazard Time per Instruction

IP3023: 0.3
 IP51xx 0.1 (with ipOS multithreading)
 MIPS: 0.3 to 0.5 (H&P, page A-65)
 ARM9: 0.3 to 0.5

Overall Performance

Applying this information to the performance equation derived earlier, the following comparison can be constructed:

Product	Clock (MHz)	Archt. Efficny.	Mem. Time	Hazard Time	Relative Perf.
ARM9	180	1.15	0.75	0.5	92
MIPS 4K	230	1.0	0.5	0.5	115
IP3023	250	1.25	0	0.3	240
IP5160	275	1.25	0.1	0.1	286
IP5170	350	1.25	0.1	0.1	364

The Ubicom CPU core is almost three times the performance of ARM or MIPS in the same process technology.

Software I/O

The IP3023 and IP51xx do many I/O functions in software that other communications processors do in hardware. This typically uses 10 to 20 percent of the CPU performance in exchange for a large reduction in die area and cost.

Software I/O MIPS used:

IP30xx PCI: 1 MIPS per 2 Mbps of TCP WLAN throughput. A turbo-g router with 55 Mbps throughput uses 27 MIPS for PCI (11%).

IP51xx PCI: 1 MIPS per 8 Mbps of TCP WLAN throughput. An 802.11n router with 150 Mbps throughput uses 19 MIPS for PCI (7%).

IP30xx 10/100 Ethernet: The Ethernet HRT uses 32 MIPS for full duplex full speed operation (185 Mbps). A 55 Mbps stream from a turbo-g radio uses 9.5 MIPS (4%).

IP51xx 10/100 Ethernet (WAN port): The Ethernet HRT uses 20 MIPS for full duplex full speed operation (185 Mbps). A 60 Mbps fiber broadband link uses 6.5 MIPS (2%).

IP5160/70 Gbit Ethernet (LAN port): The IP5160/70 Gbit Ethernet HRT has a peak throughput of 400 Mbps using 34 MIPS. A 150 Mbps 802.11n stream uses 13 MIPS (5%).

Outbound CRC generation: Packets that are routed or generated by the chip need CRC generation. Inbound WLAN packets also need CRC generation. The IP30xx family does CRC generation in software, and the IP51xx family does it in hardware.

IP30xx CRC Generation: 1 MIPS per 3.2 Mbps WAN or 6.4 Mbps WLAN. A turbo-g router with 55 Mbps WLAN throughput and 6 Mbps WAN uses 10 MIPS (7%)

An IP30xx SOHO router with a 6 Mbps WAN port and 55 Mbps Turbo-g WLAN uses:

PCI:	27 MIPS
WAN:	1 MIPS
LAN:	10 MIPS
CRC8:	10 MIPS
Total:	48 MIPS
	<u>19% of total CPU</u>

An IP51xx SOHO router with a 60 Mbps WAN port and 150 Mbps 802.11n WLAN uses:

PCI:	19 MIPS
WAN:	6 MIPS
LAN:	13 MIPS
Total:	38 MIPS
	<u>14% of total CPU</u>

Packet Buffers in Code Memory

The IP30xx only has 64-KB of on-chip data memory, so packets are stored in on-chip code memory which introduces some overhead. The IP51xx has a more flexible on-chip memory system, so packets are stored in data memory, increasing performance about 10 percent.

ipOS Efficiency

The ipOS operating system is designed explicitly for fast packet processing, and it is much more efficient than VxWorks or Linux. It uses a run-to-completion model, which avoids the performance loss due to context switching. It has a set of packet access APIs that make packet processing code more efficient and avoids copying packets.

ipOS is 1.5 to 2.5 times as efficient as Linux, depending on the Linux version.

Product	Routing Perf. (Mbps)	Frequency / Arch. (MHz)
Broadcom 5352	30	200 / MIPS
Atheros 2313	53	230 / MIPS 4K
RDC R3211	38	150 / (ARM 9)
Intel IXP-425	70	266 / (StrongARM with HW accl.)
Realtek* 8651B	170	200 / (Lexra with HW accl.)
Ubicom IP3023	150	250
Ubicom IP5160	185**	275

Routing Performance Table

* The Realtek chip has a hardware accelerator that offloads routing from the CPU, so this number doesn't show CPU performance.

** Routing performance would be over 240 Mbps if it were limited by the CPU.

Routing Performance

Router performance measured on several home gateways and on the IP3023 routing reference design. This is simple NAT routing of the Chariot TCP streams. Compared to routers that don't have hardware routing acceleration, the Ubicom router is 2.5 to 6 times faster. The CPU is about twice the performance of the other RISC CPUs. The remaining difference is due to ipOS efficiency.

Routers with hardware acceleration can approach Ubicom performance, but at a higher cost, and without the flexibility to implement changing protocols.

802.11g Bridge Performance

At 26 Mbps total chariot TCP throughput the wireless interface is the bottleneck on performance, however here is how many the processing breaks down as follows on an IP3023:

12 MIPS PCI HRT

5 MIPS 10/100 Ethernet HRT

18 MIPS lost due to hazards – only one thread is running

15 MIPS OS, stack, packet switching
60 MIPS total, or 24% total chip utilization, including hazards

30% of total time is spent in Hazards

802.11g MIMO Bridge Performance

At 56 Mbps total chariot TCP throughput the wireless interface is the bottleneck on performance, however here is how many the processing breaks down as follows on an IP51xx:

12 MIPS PCI HRT

5 MIPS 10/100 Ethernet HRT

31 MIPS OS, stack, bridging

25 MIPS lost due to hazards – only one thread is running

5 MIPS lost waiting for the caches

29% CPU utilization, 78 total MIPS

Security Performance

The IP3023 does not have security accelerator hardware.

RC4: 21 clocks per byte

DES: 120 clocks per byte (64 bit blocks)

3DES: 350 clocks per byte (64 bit blocks)

AES: 25 clocks per byte (128 bit blocks)

MD5: 12 clocks per byte (512 bit blocks)

SHA-1: 26 clocks per byte

Michael: 6 clocks per byte

Packet buffer processing overhead: 3 to 10 clocks per byte.

This results in approximately 5—10Mbps of VPN tunnel termination performance when using triple DES (3DES) for an IP3K software implementation.

RSA: About 90 million instructions for 1024 key ModExp (1 per SSL session)

The IP51xx includes hardware acceleration for DES, 3DES, AES, MD5, and SHA-1 to increase performance to over 100 Mbps. CPU instruction enhancements double the RSA performance.

DSP performance

The Ubicom instruction set includes 16-bit fixed point DSP instructions. DSP is typically used for

audio CODECs that need real-time performance. Most SOCs require a separate DSP core to meet the real time requirements, but Uvicom can run the CODEC on a dedicated hardware thread. The IP51xx adds additional DSP instructions which double the performance for audio processing.

G.729AB CODEC:

IP3023:	70 MIPS
IP51xx	35 MIPS

Summary

The StreamEngine 5000 family of processors coupled with the ipOS operating system builds on the IP3000 architecture to deliver unparalleled packet processing performance for embedded applications. In 0.13u technology, the architecture achieves higher clock frequencies and 2 to 3 times the packet processing efficiency per clock cycle versus MIPS or ARM architectures.

Combined with the flexibility afforded by Software I/O, the ability to process audio natively and the hard real time processing characteristics, the StreamEngine 5000 family offers a powerful design alternative for communications and media applications.



510 N Pastoria Avenue
Sunnyvale, CA 94085
Tel (408) 789-2200
Fax (408) 739-2427
Email sales@ubicom.com
Web www.ubicom.com

© 2006 Uvicom, Inc. All rights reserved. Features could be changed