

## **Introduction to the MMAGIX Multithreading Supercomputer**

A supercomputer is defined as a computer that can run at over a billion instructions per second (BIPS) sustained while executing over a billion floating point operations per second (GFLOPS) peak.

This is not the same as running at 1 GHz or higher. Intel and AMD chips run at very high clock speeds, but they take multiple clock cycles to execute an instruction even with long pipelines, superscalar execution and speculative execution, due to their very difficult to decode Complex Instruction Set Computer (CISC) architecture. These methods to gain a little performance cost a lot of silicon area and consume a lot of power, generating a lot of heat, while still wasting many clock cycles due to memory latency, falling well short of being true supercomputers.

The main competition to Intel's Pentium and AMD's Athlon and Opteron has been Reduced Instruction Set Computer (RISC) processors such as IBM PowerPC G4/G5, SUN UltraSPARC, ARM or MIPS and Digital Signal Processors (DSPs) such as TI's and Sony's.

RISC architectures execute at peak one instruction or more per clock cycle at lower power consumption. However, they also suffer from memory latency and fall well short of being supercomputers. DSPs are very hard to program and are not for general purpose use.

The major competitive advantage that the Intel/AMD CISC chips have had over the RISC chips has been the large range of Microsoft Windows software. They are not much used where Windows software compatibility is not the prime requirement.

As the industry moves from 32 bits to 64 bits and from Windows to portable multithreading operating systems such as Linux, this competitive advantage is lost.

All of these prior art uniprocessor architectures do not, can not and will not execute multiple simultaneous program threads effectively and efficiently. Therefore they can never deliver the high speed real-time simultaneous multithreading supercomputer performance and the very high bandwidth superconnectivity that are essential for true multimedia digital convergence that is the future in the computer, entertainment, telecommunications, military and aerospace industries.

These industries have realized that the limits of uniprocessor performance have been reached and that the only way forward is multithreading multiprocessor technology.

All the manufacturers mentioned have developed multicore versions of their prior art uniprocessor architectures; essentially just 2 or more of their existing processors on a chip. This is a very expensive and high power consumption solution that still does not deliver true supercomputing.

Our MMAGIX Technology R&D has taken a different approach, researching and developing from first principles a highly scalable and highly silicon area, memory, power and heat efficient, economical 64 bit symmetric multithreading (SMT) multiprocessor architecture with our many groundbreaking proprietary patent pending innovations, so that the family of MMAGIX multithreading multiprocessor chips will have supercomputer performance and superconnectivity at low fabrication cost, low power consumption and low heat generation.

FPGA technology has recently advanced to the point that MMAGIX can now be implemented in FPGA, for the first time making multithreading supercomputer performance possible in FPGA.

*"We have the Technology that others can only dream of."*

## **The MMAGIX Multithreading Supercomputer Advantage**

A MMAGIX supercomputer has one or more multithreading processors known as MMAGIX IP Cores. The number of IP Cores depends on the product version and scales well to many cores. Each MMAGIX IP Core combines 6 multithread register and instruction decode pipelines with a set of specialized instruction execution pipelines that can be shared between up to 6 instructions simultaneously. On average each MMAGIX IP Core can execute 5-6 instructions per clock cycle, each from a different thread. MMAGIX has an innovative compact RISC/CISC instruction set, rich in instructions closely targeted to the needs of optimizing compilers for modern high level languages such as C, C++, Java, Ada, Fortran and Cobol, almost doubling code density over RISC. The instruction set includes atomic instructions for interprocessor coordination and dynamic library linkage, which dramatically improve the real-time performance of multithreading operating systems such as Linux and BSD Unix.

The MMAGIX architecture gains its performance advantage not by pushing the limits of clock frequency at any cost, but by parallel processing and a number of silicon area efficiency and functional sharing innovations that result in very low power consumption and heat generation.

*“The MMAGIX Multiprocessor is the Supercomputer Technology for the 21<sup>st</sup> Century.”*

## **The MMAGIX Multithreading Supercomputer in FPGA on AMC**

Field Programmable Gate Array (FPGA) technology has recently advanced to the point that a MMAGIX symmetric multithreading (SMT) multiprocessor IP core can now be implemented in FPGA without the fabrication cost of a custom chip and can support superconnectivity interfacing to high bandwidth encryption/decryption/compression/decompression and telecommunications switching circuitry on the same FPGA chip.

The telecommunications industry has standardized on the Advanced Telecommunications Computer Architecture (ATCA) and Micro Telecommunications Computer Architecture (MTCA), which are telecommunications rack, backplane and carrier card standards for 10+ Gb/s data switching with a free market in standards compliant equipment. The ATCA and MTCA standards are respectively for larger and smaller telecommunications systems and both standards have the actual telecommunications circuitry in the form of small plug in Advanced Mezzanine Cards (AMCs) that can contain control and switch fabric circuitry and 10+ Gb/s data interfaces. As predicted, the US military and aerospace industries have also recently adopted the MTCA standard, so the future of ATCA, MTCA and AMC is assured.

Because of the need to upgrade telecommunications equipment when data transmission standards change, the industry has become the major user of high performance FPGAs in volume. The actual circuitry in an FPGA can be field upgraded, unlike a custom chip.

Consequently several telecommunications equipment manufacturers are producing AMC products with an FPGA, a DDR2 SODIMM for memory and 1 or 2 10 Gb/s SPI-4.2 interfaces. These products need 1+ BIPS of multithreading supercomputer power on FPGA and only MMAGIX can deliver a licensable IP core of this performance that will fit on the Xilinx Virtex-5 FPGAs in use. We are in close contact with this market and its immense opportunities.

Professional telecommunications equipment is high profit margin and the FPGAs cost over \$1000 each and depend totally on licensed IP for their circuitry, making it very viable to charge substantially for each MMAGIX IP Core on FPGA.

MMAGIX supercomputer technology in FPGA is perfectly suited to the fast growing AMC market and AMC is currently the biggest market opportunity for MMAGIX.

# The MMAGIX Quantum Leap

*“Both multithreading and multi-core approaches exploit the concurrency in a computational workload. The cost, in silicon, energy and complexity, of making a CPU run a single instruction stream ever faster goes up nonlinearly and eventually hits a wall imposed by the physical limitations of circuit technology. ... multithreading is processor-level optimization to improve area and energy efficiency. Multithreaded architecture is driven to a large degree by the realization that single-threaded, high-performance processors spend a surprising amount of time doing nothing. When the results of a memory access are required for a program to advance, and that access must reference RAM whose cycle time is tens of times slower than that of the processor, a single-threaded processor can do nothing but stall until the data is returned.” EETIMES 26 September 2007: “Demystifying multithreading and multi-core.”*

MMAGIX IP comprises Pending Patents and novel Architecture, Instruction Set & Designs for:

- High performance 64 bit wordsize Symmetric Multithread IP Cores optimized for Multicore.
- Each superpipelined Symmetric Multithread IP Core simultaneously fetches, decodes and executes up to 6 instructions per clock cycle, drawn from 6 Thread Register and Instruction Decode pipelines of 16 threads each, absorbing instruction and memory access latencies.
- With multiport L1 I&D-Caches and symmetric execution pipelines, on average over 5 of the 6 instruction execution slots per core will be filled per clock cycle.
- As all instructions have some latency, sequential instructions are executed from different threads, eliminating the need for complex and congested pipeline bypass circuitry.
- By absorbing instruction and memory access latencies, the actual instruction execution rate is multiplied compared to a unithread processor core.
- Each thread has a bank of 32 64 bit registers, encoded as 16 data registers and 16 address registers, plus an extra bank of 32 registers for operating system traps and interrupts.
- The 16/32 bit RISC/CISC instruction set is optimized for C/C++/Java/Ada/Fortran/Cobol.
- By optimizing for the instructions, condition codes and traps used by high level language compilers such as GCC, common instruction sequences need less executable MMAGIX instructions and so take less I-Cache memory and less clock cycles to execute.
- SIMD Vector & String instructions for operations on multiple 8, 16, 32 or 64 bit variables per clock cycle multiply the speed of operations on vectors & strings of smaller variables.
- Zero cycle looping speeds up all loops and is optimized for scalar & SIMD vector & string.
- Load and Store instructions support all C/C++ data types of 1, 8, 16, 32, 64 & 128 bits and do not have to be at aligned addresses.
- Most function parameters are passed in registers and leaf functions do not need to save context on the stack. When needed, register save and restore is doubleword at a time.
- Branch, Call and Return instructions have their latency absorbed by multithreading and so do not need branch prediction hardware for further silicon area savings.
- These features combined make for exceptionally efficient scalar & SIMD vector & string loops, either inline or as functions, aligned or unaligned, challenging Altivec-like SIMD designs for high performance while being implemented as simple compiler optimization of ordinary C/C++ software loops, not nonportable hand crafted assembler software.
- Prior art SpinLocks obsoleted by our patent pending L1 & L2 Cache and Memory Controller support for Atomic Memory instructions such as SleepLock, Unlock, Semaphore and Signal, greatly enhancing the efficiency of real time operating systems and device drivers in the critical regions where threads interact, unblocking a major performance bottleneck.
- Our patent pending Shared Coprocessors save much silicon area and power by sharing complex circuitry such as encryption/decryption between cores, cycle by cycle.
- High bandwidth Multistream Data Flow Controllers can be interfaced to multiple MMAGIX IP Cores as Shared Coprocessors so multiple data streams can be simultaneously serviced by very low latency threads. This is the major application for Multithread IP Cores.
- MMAGIX in FPGA delivers 1+ BIPS per core @ 200 MHz and superconnectivity @ 20+ Gb/s.

## Feature Comparison between leading Multicore Architectures

*Any inaccuracies are due to difficulty in obtaining some technical information and will be corrected if necessary. Suggestions for features to compare are welcome.*

*No multithreading multicore architecture other than MMAGIX is available in FPGA. At 1+ BIPS per MMAGIX IP Core in FPGA, MMAGIX stands alone without any competition.*

<b>Feature</b>	<b>MMAGIX</b>	<b>Cortex-A9</b>	<b>Cavium</b>	<b>Raza XLR</b>	<b>IBM Cell</b>
Processor core	MMAGIX	ARM	MIPS	MIPS	PowerPC
Data/Address bits	64	32	64	64	64
16/32 bit instructions	16 and 32	16 or 32	16 or 32	16 or 32	32
ASIC frequency MHz	666-2000+	1000+	300-1000	600-1600	3000+
FPGA frequency MHz	200	-	-	-	-
Multithread/Superscalar	SMT	SS	SS	MT	SMT
Pipeline stages	11	8	5	10	21
Cores/chip	1-64+	1-4	1-16	2-16	1+8 DSP
Cores/FPGA	1+	-	-	-	-
Instructions/cycle/core	~5.5	<0.20	<0.20	~0.70	~0.35+DSP
Threads/core	96	1	1	4	2
Registers/thread	32	16	32	32	32
SIMD multimedia ops	Yes	Optional	No	No	Yes
SIMD multiply/divide ops	Yes	Optional	No	No	Yes
SIMD binary FP ops	Optional	Optional	No	No	Yes
SIMD decimal FP ops	Optional	No	No	No	No
Load/store unaligned	Yes	Difficult	Difficult	Difficult	Difficult
Full word literals	Yes	No	No	No	No
Zero cycle looping	Yes	No	No	No	No
Register subscripts	Yes	Yes	No	No	Yes
Scaled subscripts	Yes	Yes	No	No	Yes
Trap/interrupt registers	Yes	Yes	No	No	No
Threads for interrupts	Yes	No	No	No	No
L2 atomic memory ops	Yes	No	No	No	No
L2 sleeplock/unlock ops	Yes	No	No	No	No
L2 memory mapping	Yes	No	No	No	No
L2 dual port I/O buffers	Yes	No	No	No	No