



The system software I'm using:

RECONFIGURABLE COMPUTING

The Art of Change: Technologies for designing our future

A look into QuickSilver's ACM architecture

By Paul Master

[EE Times](#)

September 12, 2002 (4:39 p.m. EST)



QuickSilver Technology's Adaptive Computing Machine (ACM) is a new class of digital integrated circuit designed to bring adaptive computing into commercial use.

The ACM is designed to "adapt" — change its on-chip hardware configuration in response to software — many tens or hundreds of thousands of times per second, while consuming very little power. Such a capability allows software algorithms to be directly converted into dynamic hardware that is specific for that instant in time, during run-time, which results in optimum performance, low power consumption and the most efficient use of silicon real estate, which, in turn, translates to lower cost. The ACM's flexible architecture enables longer battery life and multifunctionality — ideal attributes for next-generation wireless devices.

Algorithmic evaluation

Mainstream research has focused on FPGA-based approaches to reconfigurable computing (RC). But, such conventional RC technology approaches the problem at too macro a level. That is, RC tends to work at the level of entire applications or algorithms. In reality, it is critical to consider the problem at the micro level of algorithmic elements.

A flexible architecture

Consider, for example, the number of elements used in word-oriented algorithms, such as the compute-intensive time-division multiple access (TDMA) algorithm employed in digital wireless transmission. Variants such as Sirius, XM Radio, EDGE and so forth, form a subset of this algorithmic class. A flexible architecture that can handle high-end TDMA should also be able to handle its less sophisticated cousins.

The Algorithm Space diagram maps word-oriented algorithms and their bit-orientated counterparts into a composite landscape. This includes algorithms for wideband code division multiple access (W-CDMA), which is used for wideband digital radio communications of Internet, multimedia, video, and other



READER CONTEST
STARTS SEPT. 16
What technologies will reshape the world?
What vendors will lead?
Enter and win a chance at up to \$1000 in cash prizes.



search go

THE ART OF CHANGE

- ▶ Editorial
- ▶ Introduction
- ▶ Indium Phosphide
- ▶ Nanostructures
- ▶ Ultrawideband Technology
- ▶ Platform-based design
- ▶ Reconfigurable Computing
- ▶ MEMS
- ▶ Optical Devices
- ▶ Mobile Robotics
- ▶ Companies to Watch

ANNIVERSARY LINKS

- ▶ NetSeminars
- ▶ Reader Contest
- ▶ Sponsors
- ▶ About this site

DEPARTMENTS

- SEMICONDUCTORS
- SYSTEMS & SOFTWARE
- EE DESIGN
- TECHNOLOGY
- THE WORK CIRCUIT
- COMMSDESIGN
- PLANET ANALOG
- EMBEDDED.COM
- IAPPLIANCWEB

HOT LINKS

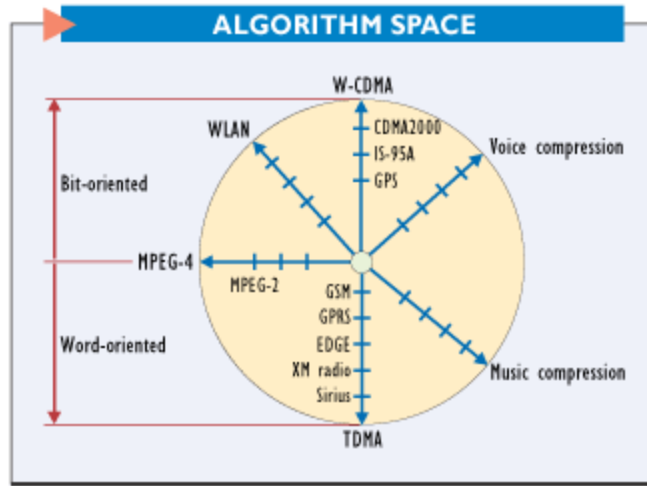
- ▶ Ask the Headhunter
- ▶ Books
- ▶ Calendar/Events
- ▶ Conference Coverage
- ▶ EET Info/Reader Service
- ▶ In Focus Reports
- ▶ Mentors Board - Students
- ▶ Online Seminars
- ▶ Opinion/Columns
- ▶ ProductWeek
- ▶ Salary Survey 2001
- ▶ Special Reports
- ▶ TWC on Campus

NETWORK RESOURCES

JOB SEARCH go

capacity-demanding applications, and subvariants such as CDMA2000, IS-95A and so forth. The diagram also shows various mixes of word-oriented and bit-oriented components, such as MPEG, and voice and music compression. The ACM architecture is able to cover this very large problem space and all the points in between.

Algorithms are heterogeneous in nature, which means that within complex algorithms, constituent elements are substantively different. In turn, this indicates that the homogeneous architectures associated with traditional FPGA-based RC approaches — which have the same lookup table replicated tens of thousands of times — are not appropriate for most algorithmic tasks. Even newer, more advanced FPGAs with complex elements like 18 x 18 multipliers do not satisfy the requirements of adaptive computing.



How, then, is it possible to create a flexible architecture that rapidly adapts to algorithmic input without compromising the ASIC gold standard of speed and low power consumption?

The solution is to create a heterogeneous architecture that fully addresses the heterogeneous nature of the algorithms (see figure). Start with five types of nodes: arithmetic, bit manipulation, finite state machine, scala, and configurable input/output used to connect to the outside world.

Memory cache

Each node consists of computational gates and its own local memory cache (approximately 75 percent of a node is in the form of memory). Additionally, each node includes configuration memory, but unlike the serial configuration bit stream used in an FPGA, an ACM has a dedicated 128-bit or 256-bit bus to carry the data used to adapt the device.

It's important to realize that each node performs tasks at the level of complete algorithmic elements. For example, a single arithmetic node can be used to implement different, variable-width, linear arithmetic functions such as a FIR filter, a discrete cosign transform, a fast Fourier Transform (FFT) and so forth. Such a node can also be used to implement variable-width nonlinear arithmetic functions such as $((1/\sin A) \times (1/x))$ to the 13th power.

Similarly, a bit-manipulation node can be used to implement different, variable-width bit-manipulation functions, such as a Linear Feedback Shift Register, Walsh code generator, GOLD code generator, TCP/IP packet discriminator and others.


[eencyclopedia](#) [go](#)
[elibrary](#) [go](#)
NetSeminar Services [go](#)
EE TIMES NETWORK
Online Editions
[EE TIMES](#)
[EE TIMES ASIA](#)
[EE TIMES CHINA](#)
[EE TIMES KOREA](#)
[EE TIMES TAIWAN](#)
[EE TIMES UK](#)
Web Sites
[CommsDesign](#)
[GaAsNET.com](#)
[iApplianceWeb.com](#)
[Microwave_Engineering](#)
[EEdesign](#)
[Deepchip.com](#)
[Design & Reuse](#)
[Embedded.com](#)
[Elektronik i Norden](#)
[Planet Analog](#)
[Semiconductor Business News](#)
[The Work Circuit](#)
[TWC on Campus](#)
ELECTRONICS GROUP SITES
[ChipCenter](#)
[EBN](#)
[EBN China](#)
[Electronics Express](#)
[NetSeminar Services](#)
[QuestLink](#)
THIS WEEK'S ISSUE

[Click here for "EE TIMES DIGITAL," the weekly issue in PDF format](#)
SIGN UP FOR NEWSLETTERS
[CLICK HERE](#)

A finite state machine node can be used to implement any class of finite state machine (FSM). In the case of a really large or complex FSM, the machine can be spread across multiple FSM nodes, or different portions of the state machine can be time-sliced across a single node. This means that the node can be adapted to execute different portions of the state machine on the fly.

A scalar node can be used to execute legacy code, while a configurable input/output node (not shown in the figure) can be used to implement I/O in the form of a UART or bus interfaces such as PCI, USB, Firewire and other I/O-intensive actions.

A key advantage of the ACM's architecture is that any node can be adapted to perform a new function, clock cycle by clock cycle. This means that any portion of the ACM — from just a few nodes and interconnects up to the entire chip — can be adapted in a single clock cycle. This results in a radical change in the way algorithms are implemented today. Rather than passing data from function to function, the data can remain resident in a node while the function of the node changes on a clock cycle-by-clock cycle basis. It also means that, unlike an ASIC implementation, the ACM can be adapted tens or hundreds of thousands of times a second, so that only those portions of an algorithm that are actually being executed need to be resident in the chip at any one time. This reuse of gates enables tremendous reductions in silicon area and power consumption.



Since the ACM's architecture is fractal in nature, it is totally scalable. A four-node cluster is formed from an arithmetic, bit-manipulation, FSM and scalar node, which is connected via a Matrix Interconnection Network (MIN). A 16-node cluster is formed from four 4-node clusters linked by their own MIN, while a 64-node cluster is formed from four 16-node clusters linked by their own MIN, and so forth. An ACM can contain from one to thousands of node clusters, as required.

Design advantages

ACM designs are represented in SilverWare, which is the C language augmented with temporal and spatial extensions. This means that, unlike an ASIC-based implementation in which algorithms are effectively frozen in silicon, the ACM can be quickly and easily adapted to accommodate the numerous evolving standards and protocols used in today's designs. In addition to accelerating time-to-market, this approach eases design reuse and reduces the risk of failure.

ACM technology also eliminates the very difficult problems associated with hardware/software co-design, because the entire system is initially represented as software. Having said this, it's important to understand that SilverWare is not executed by the ACM in the same way that machine code is processed by a DSP, i.e., executing a long stream of instructions. Instead, SilverWare is used to dynamically adapt the ACM on the fly to create the exact hardware needed to perform whichever algorithmic tasks are required at any particular time. Complex algorithmic elements can be thought of as the smallest operators, and many of these complex algorithmic elements are temporally or spatially combined to form an application. In essence, software becomes hardware.

Paul Master is the CTO of QuickSilver Technology.

THE RESOURCE CENTER

<p>INTERNATIONAL JOB SEARCH FIND A JOB</p> <input style="width: 100%;" type="text"/> <input type="button" value="GO"/> <p>NetSeminar FIND AN <u>ONLINE SEMINAR</u></p> <input style="width: 100%;" type="text"/> <input type="button" value="GO"/> <p><input type="button" value="PRINT THIS STORY"/> <input type="button" value="SEND AS EMAIL"/></p>	<p>search FIND A RELATED STORY</p> <input style="width: 100%;" type="text"/> <input type="button" value="GO"/> <p>cyclopedia LOOK UP A TECHNICAL TERM</p> <input style="width: 100%;" type="text"/> <input type="button" value="GO"/> <p>library FIND A WHITE PAPER</p> <input style="width: 100%;" type="text"/> <input type="button" value="GO"/>
---	--

Electronics Marketplace

- [We Ship Same Day -PCB protos/production](#)**
They offer lower price, but are you willing to burn pricey components? Choose wisely. 25% OFF Any Size order-Click here ...
 - [The Fastest Embedded Processor Ever - Xtensa V](#)**
Test drive Tensilica's Xtensa V processor, the fastest ever according to EEMBC Certification Labs.
 - [PCB123-the PCB Solution, Design to Order](#)**
1)Download FREE circuit board design software that converts schematic to layout. 2)Easy order process (quotes your board...)
 - [Prototype Circuit Boards from PCBexpress](#)**
Two-layer, 20-piece at \$11 each, ship in 24-hours. Four-layer, 10-piece at \$19 each, ship in 48 hours. No tooling charge...
 - [5 New FREE IC Technical Publications](#)**
Download these and view our entire library today!
- [Buy a link NOW:](#)

TSMC
[Check out TSMC's new SOC technology, Nexsys](#)

Verisity
[See an online demo of SureCov's code coverage technology!](#)

