The Silicon Objects High-Performance Processor: An Overview

A New Approach to High Throughput Scalable and Reconfigurable Data plane Processor Design

MathStar's Silicon Object Array (SOA) products are massively parallel processor arrays with 1GHz internal clock speed, flexible I/O interconnect, and powerful high-level software tools. The first member of this product line – the Hydra – is aimed at multi-gigabit line-rate data and communications processing.

After a brief overview of the paper, the following areas are covered:

- *Hardware architecture*
- *Software architecture*
- *Hydra application space*

Low-level detail and programming examples are provided.

I Overview

Why Reconfigurable?

Transistor density is ever increasing at a tremendous rate as the industry marches toward sub-0.1µm technologies. Often, the processing potential this brings remains unlocked due to prohibitively high development complexity, time, and cost. Currently, full custom chip design in 0.13µm can reach as high as $5-10 M.

While traditional microprocessors and Field Programmable Gate Arrays (FPGA) based designs avoid these high Non-recurring Engineering (NRE) expenses, overall lack of performance (and efficiency) leads to large area and hence high per chip cost. In certain cases – such as multi-gigabit line rate communications processing – performance constraints render these solutions untenable.

To fill the void, vendors must move to reconfigurable silicon *platforms*. Through the use of application-targeted processor arrays, the trade-off between 1) design time and NRE cost versus 2) overall performance and per-chip cost can be mitigated.

Further, reconfigurable silicon arrays provide additional benefits of allowing changes to be made through software upgrades, allowing vendors to begin design before specifications (standards) are finalized. For many OEMs, these time-to-market (and hence time-in-market) advantages outweigh any advantages provided by custom silicon.

Due to the orders of magnitude in performance improvements that can be achieved, commercial reconfigurable platforms leveraging years of academic research are now poised to bring substantial value to the marketplace.

MathStar's Silicon Object array is a new approach to a reconfigurable platform that offers unheralded performance advantages. Further, it is the industry's first very high-performance multi-processor array targeted at data and communications processing.

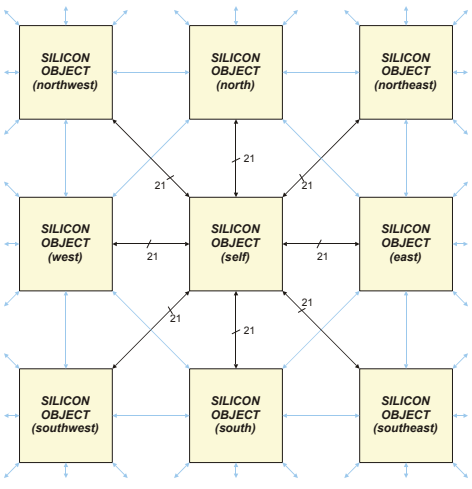

Figure 1-1: Silicon Object Array

The Silicon Objects Array: Architecture

The Silicon Objects array is a heterogeneous medium-grained array composed of hundreds of individual processing elements. Each element is called a Silicon Object.

Within the array the data path and control path are loosely coupled, yet independently configured. The data path is 16 bits wide while the control path is bit-wise granular.

Each Silicon Object has its own program and data memories. Further, each processor operates on its own without the aid of global control. Using Flynn's Taxonomy, this is best characterized as a Multiple Instruction Multiple Data (MIMD) machine.

Communication between Silicon Objects is primarily nearest neighbor. Party lines enable longer distance communication. Objects are allowed to change communication patterns on a per-clock basis (shown in Figure 1-1 above).

Each Silicon Object has a program memory of eight instructions that contain both operation and communication directions. The instructions are loaded at configuration time and cannot be changed. Intelligent compilation (scheduling and routing) tools deterministically allocate instructions to each object before run time.

The control path guides program execution while data is moved and operated upon via the 16-bit data path. From this view, instructions are the mechanisms that tie the independent control and data paths together within the array.

There are five data path Silicon Object types within the common Silicon Object array structure: the Arithmetic Logic Unit (ALU) Silicon Object, the Content Addressable Memory (CAM) Silicon Object, the Cyclic Redundancy Check (CRC) Silicon Object, the Integer/Real Multiplier Silicon Object, and the Galois Multiplier Silicon Object (for finite field arithmetic). The first three objects are incorporated within the Hydra array. The latter two are currently under development.

In addition, RAM memory resources are distributed in the array. The function and ratio of these different Silicon Objects were chosen based on detailed study of the attributes of communication processing algorithms.

To complement the flexibility of the digital processing core, the I/O ring allows user to define both LVDS and HSTL operation. Voltage levels are also selectable. Total bandwidth per pin reaches as high as 800MHz DDR (1.6Gb/s total). In addition, General Purpose I/O (GPIO) can be set to user-defined LVCMOS levels.

Future versions of the Silicon Object array are planned that incorporate MathStar's proven high-speed 3.125Gb/s SERDES technology.

The Silicon Objects Array: Applications

The Silicon Object Reconfigurable Array is a multi-processor array targeted at high-throughput data processing applications that exhibit high-levels of data flow determinism (i.e., regular dependencies) at a localized level. Irregular dependencies (e.g., interrupts, context switches) are handled as ordinary signals.

MathStar's initial Silicon Object arrays will be targeted at multi-gigabit communications processing, an application space that exhibits these algorithmic proprieties. Examples include data link layer (layer 2) processing, TCP/IP processing (layer 3+), and security processing.

At the functional level, these applications require frame/packet parsing and generation, finite state machines, CRC generation and detection, comma detection, statistics counters, hashing, and memory controllers.

This set of functionality allows the Silicon Object Array to operate as a highly powerful and flexible communication processor. Whereas most network processors only provide limited flexibility starting at Layer 2, the Hydra array allows processing beginning at Layer 1.

When combined with its flexible I/O ring, the Hydra can be used to implement various I/O standards such as 10G Ethernet, SPI-4.2, HyperTransport, and RapidIO.

Moving Forward

Finally, communication-processing problems are by no means the only algorithms that map efficiently to a processor array.

Over time, MathStar will release arrays targeted at signal processing applications such as image and video compression (e.g., JPEG2000 / MPEG), wireless LANs (802.11), and Forward Error Correction (G.709). These algorithms also exhibit favorable data flow dependencies.

In order to process these algorithms efficiently, both integer/real/complex multipliers and Galois multipliers are required. Development of high speed multipliers that fit within the common Silicon Objects framework is underway.

More Information?

Silicon Object hardware architecture, software tools, and application areas are covered in the following sections of this paper.

Detailed descriptions of communication resources, registers, instruction sets, and input/output buses are provided. Schematics and software programming examples are also given.

The information contained in this paper is presented in order to give a comprehensive picture of what the Silicon Object array is and of what it is capable. Therefore, complete details concerning the entire operation of the array are not provided.

If more detailed documents pertaining to the Hydra Silicon Object Array hardware, software, or applications are required, please contact MathStar directly.