Micron Automata Processor

A New Approach to Information Analysis

Paul Dlugosch
Director – Automata Processor Technology
Micron Technology, Inc.

Micron at a Glance

Founded: October 1978, Boise, Idaho

FY2013 Net Sales: $9.0 billion

NASDAQ Symbol: MU

Employees: ~30,000 worldwide

Products: We offer one of the world's broadest memory portfolios, including: DRAM components and modules, SSDs, NAND, and NOR, as well as other innovative memory technologies, packaging solutions and semiconductor systems

Markets We Serve: Micron's products are designed to meet the diverse needs of computing, networking, server, consumer, mobile, automotive, and industrial applications

Patents: ~26,000

Expansive Product Offering

DRAM Families

SDRAM
DDR
DDR2
DDR3
DDR4
RLDRAM®
Mobile LPDRAM
PSRAM/
CellularRAM
Hybrid Memory Cube

DRAM Modules

FBDIMM
RDIMM
VLP RDIMM
VLP UDIMM
UDIMM
SODIMM
SORDIMM
Mini-DIMM
VLP Mini-DIMM
LRDIMM
NVDIMM

Bare Die

Multiple
Technologies

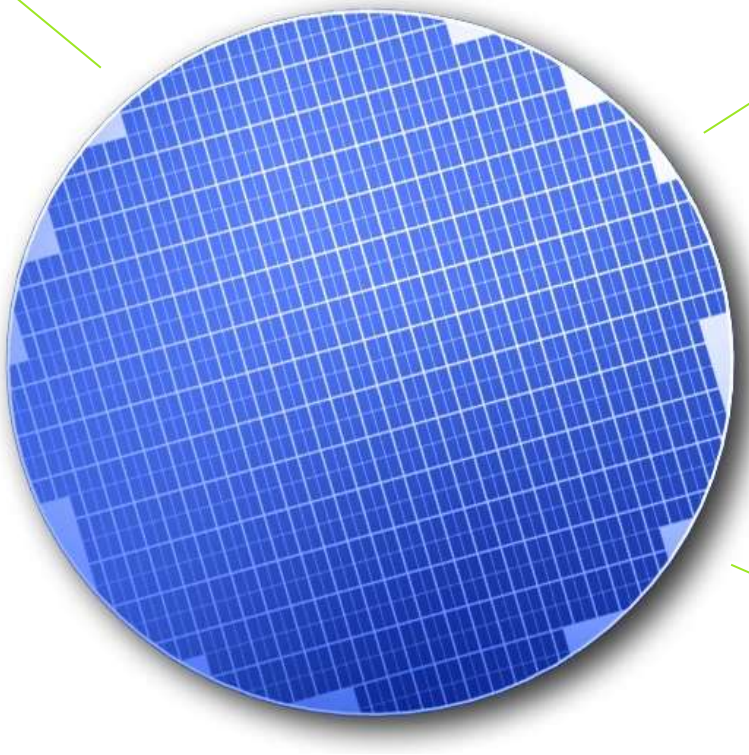

NAND Flash

TLC, MLC, SLC
Serial NAND
Enterprise NAND

Solid State Drives

Client SSD
Enterprise SATA
Enterprise SAS
Enterprise PCIe

Managed NAND

MCP
eMMC™
ClearNAND Flash
Embedded USB

NOR Flash

Parallel NOR
Serial NOR

Five Big Technology Trends

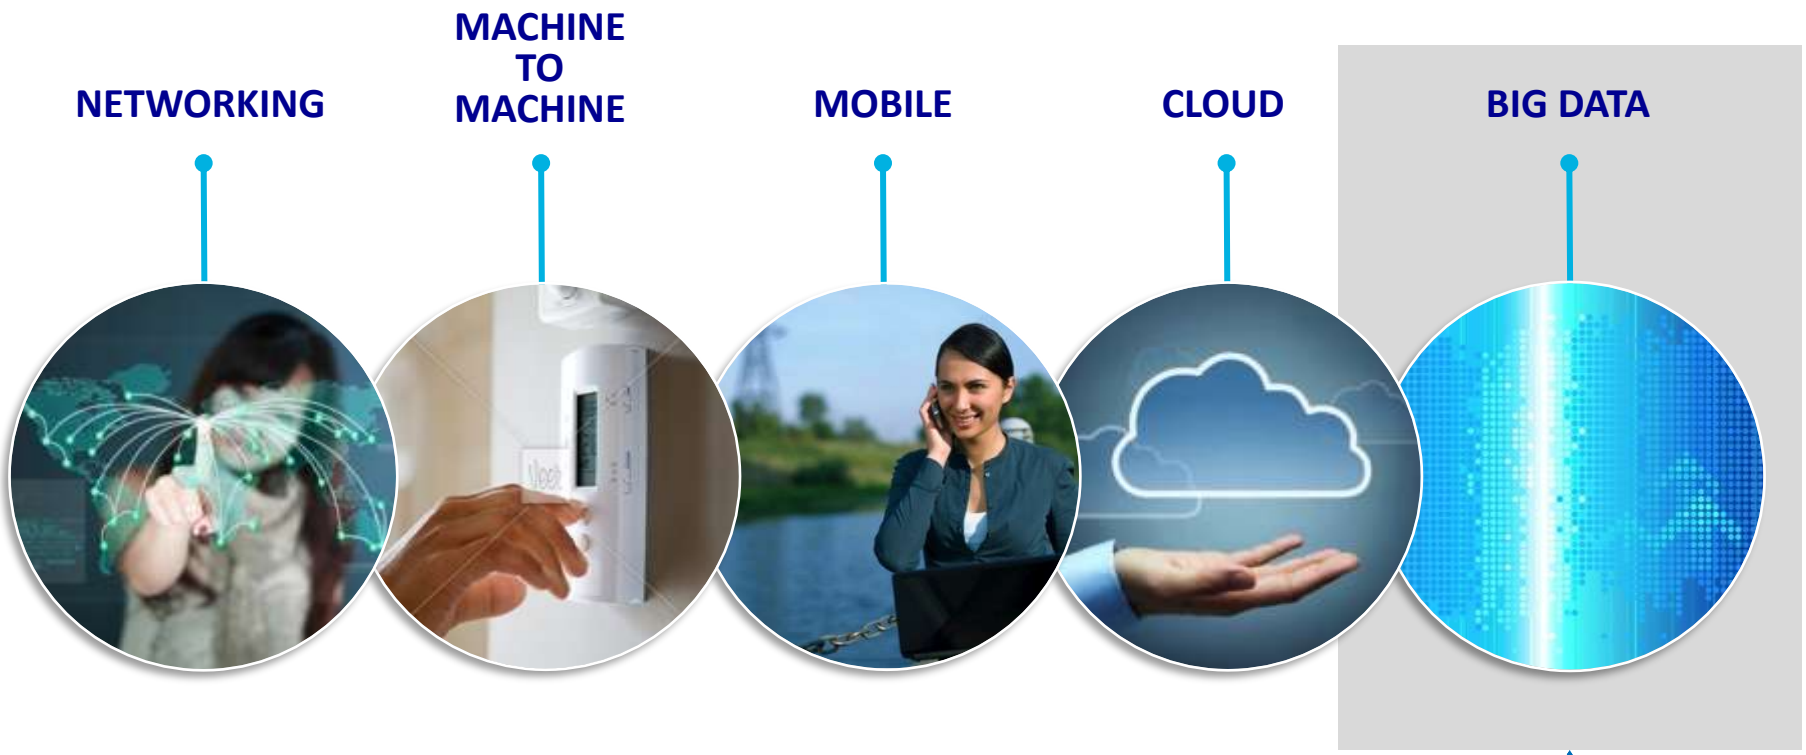

Big Data Presents A Unique Challenge for Memory Systems

- Customers demand high performance for analytics.
- Increasing levels of parallelism drive complexity in system architectures.
- Massive scale requires aggressive power targets.

A Repetitive Cycle…

The Consistent Message

"Memory is the bottleneck!" — CPU Vendor

"We need faster memory!" — System OEM

The Response

"Sure, we can do that!" — Memory Industry

Innovations in memory interfaces…

… have been critical to improving performance.

Hybrid Memory Cube: A New Level of Performance

Revolutionary Approach to Break Through the "Memory Wall"

- Evolutionary DRAM roadmaps hit limitations of bandwidth and power efficiency
- Micron introduces a new class of memory: Hybrid Memory Cube
- Unique combination of DRAMs on Logic smashes through the memory wall

Unparalleled Performance

- Provides 15X the bandwidth of a DDR3 module
- Uses 70% less energy per bit than existing memory technologies
- Reduces the memory footprint by nearly 90% compared to today's RDIMMs

Key Applications

- Data packet processing, data packet buffering, and storage applications
- Enterprise and computing applications

How did we do it?

- Micron-designed logic controller
- High speed link to CPU
- Massively parallel "Through Silicon Via" connection to DRAM

…Breaking the Cycle

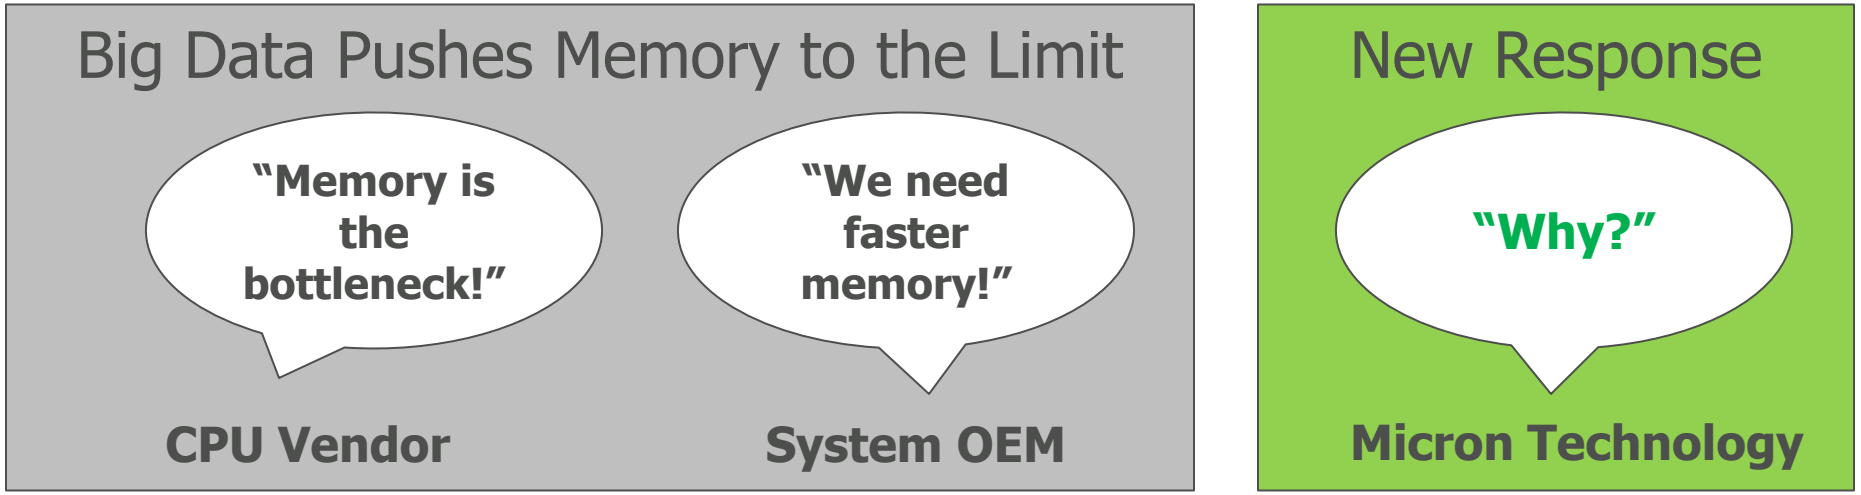

- The modern relationship between processor and memory was conceived to avoid complications associated with **physical reconfiguration of ENIAC**.
- Since the mid 1940's, most computer systems have been built on this basic architectural concept. The **role of memory** in systems was firmly cast.
- Micron concluded that important advancements can be made if we **challenge** this deeply rooted historical concept.

Micron Innovation: Automata Processor

Advanced Architecture Unleashes Massive Parallelism and new Era in Data Processing

- Unique solution to 'Big Data' problems that require data analysis
- Massively parallel memory based design (~50,000 vertex processors).
- Scalable for consumer through super computer applications.
- Allows processing performance to scale in capacity.
- Deployable in single chip, module and multi-module applications.
- Programmable architecture via Micron developed Software Dev Kit
- Custom developed Automata development language (ANML) allows full exploitation of chip parallelism and data processing capabilities.

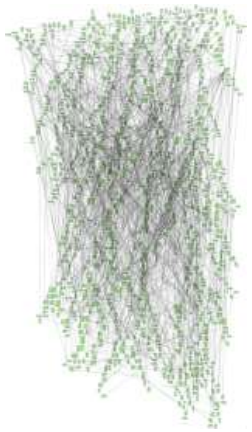

AUT•MATA PROCESSING

↖ Cyber-security Graph Pattern

Key Characteristics & Program Status

Physical Interface:	DDR3
Bus Width:	×8
State Cache:	512 vectors on chip
Symbol input rate:	128M Symbols/sec
Power:	4 W Max (all blocks active)
Sample Availability:	Mid - 2014
SDK Status:	V1.4 available now
Market Stage:	Evaluating early market initiatives

Application Examples

- Network Security*
  - Deep Packet Inspection
  - QoS
- Data Analytics
- Bioinformatics*
- Medical Diagnostics
- Video/Image analytics*

* Application development in progress

Micron Automata Processor: Silicon

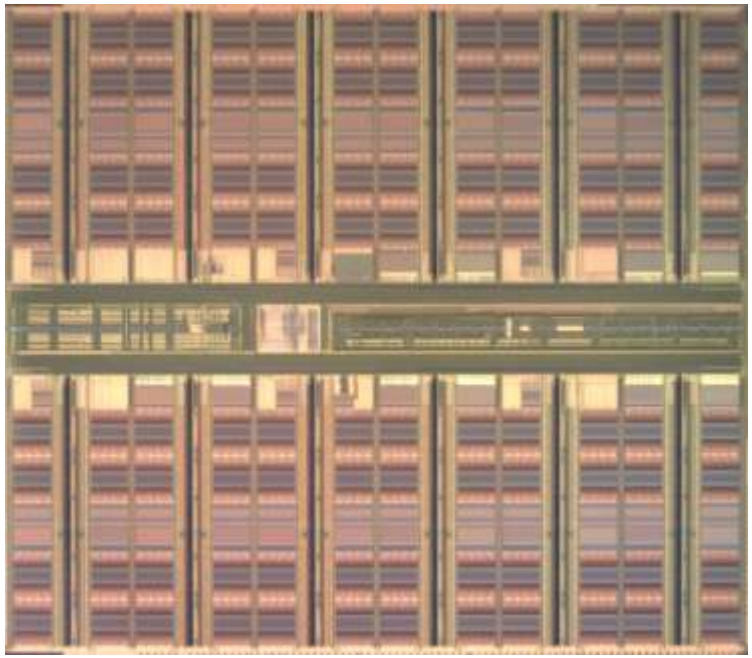

Key Device Parameters:

- 129.3 mm2 (12.15 × 10.64)
- 128M Symbols/Second
- 49,152 State Transition Elements
- 24,576 STE Max Automata Size
- 4W TPD (Estimated)
- 512 Entry State Cache
- 6,144 STE Match Capacity (Max)
- 6 Independent Result Regions
- Event Vector Division (2,4,8,16)
- 15.4mm × 12.0 mm 144b FPBGA

Sample availability targeted for Q3, 2014.

Market Opportunity

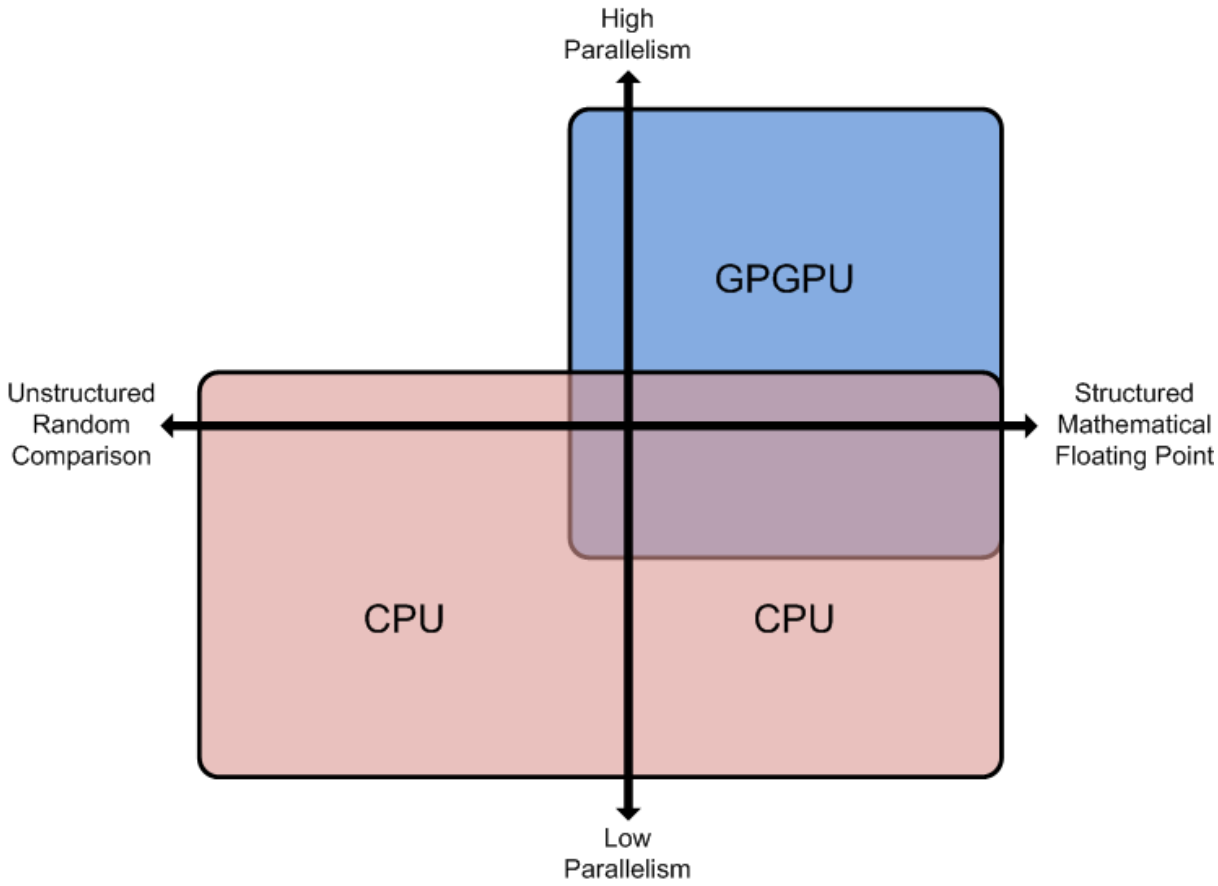

Problems Aligned with the Automata Processor

Applications requiring **deep analysis** of **data streams** containing **spatial** and **temporal** information are often impacted by the **memory wall** and will benefit from the **processing efficiency** and **parallelism** of the Automata Processor.

Network Security:

- Millions of patterns
- Real-time results
- Unstructured data

Bioinformatics:

- Large operands
- Complex patterns
- Unstructured data

Video Analytics:

- Highly parallel operation
- Real-time results
- Unstructured data

Data Analytics:

- Highly parallel operation
- Real-time results
- Unstructured data

Automata Processor: The Fabric

Match Elements:

- State Transition Element (STE)
- Determine match of input symbol
- Can support high in/out degree

Aggregation Elements:

- Counter Element
- Combinatorial Element

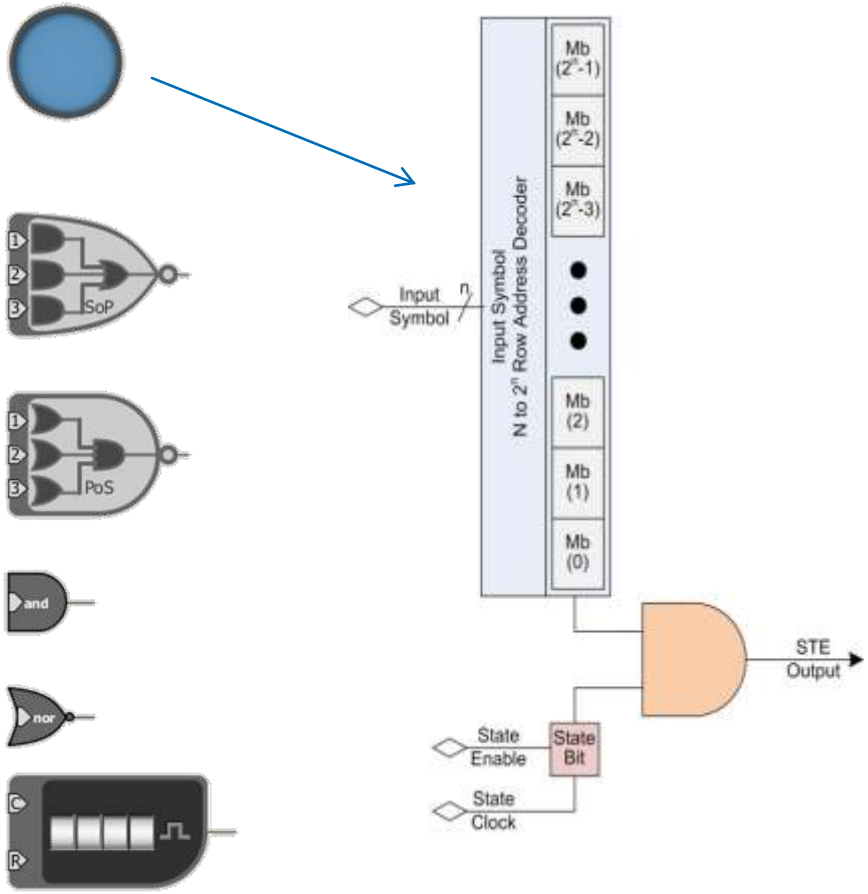

STE Activation (Routing) Parallelism

STE Array (Block Architecture)

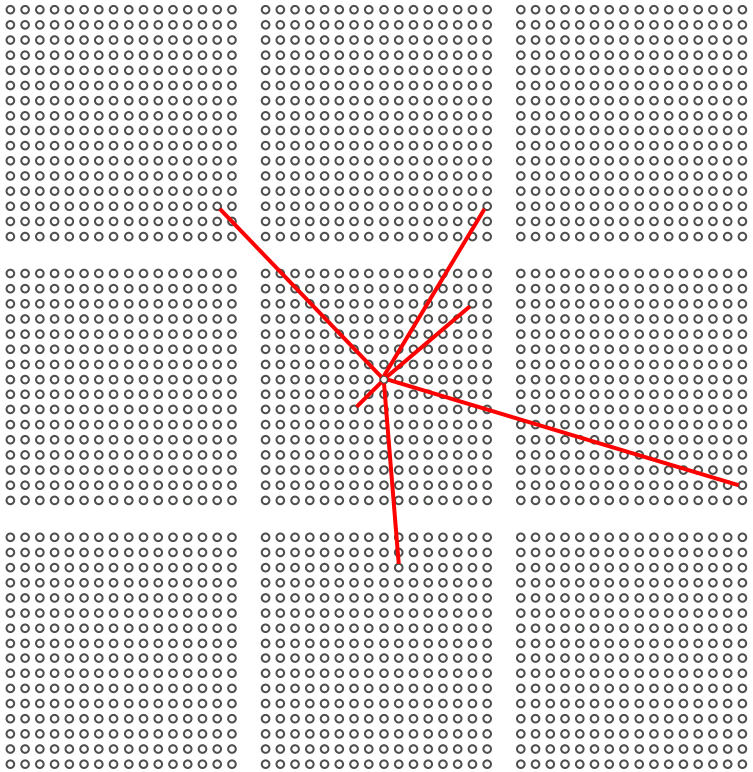

- STE's may route to any other STE within it's Block.
- STE's may route to other STE's in adjacent blocks.
- STE's may route (activate) multiple STE's simultaneously.

Routing lines effectively replace the overhead associated with a traditional random access to system memory.

Automata Processor – Basic Operation

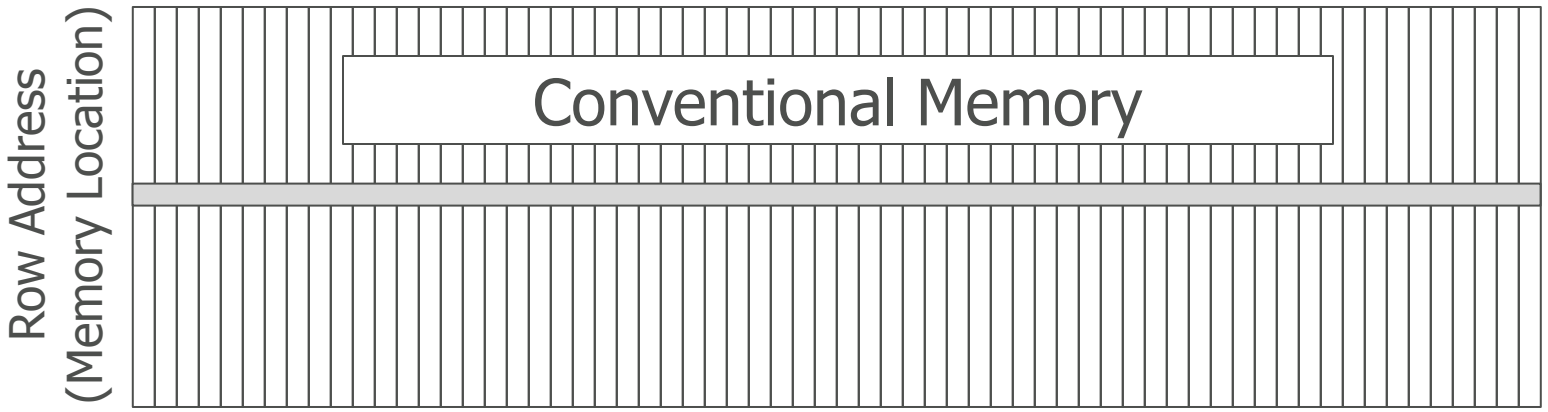

Row Access results in **one** word being retrieved from memory.

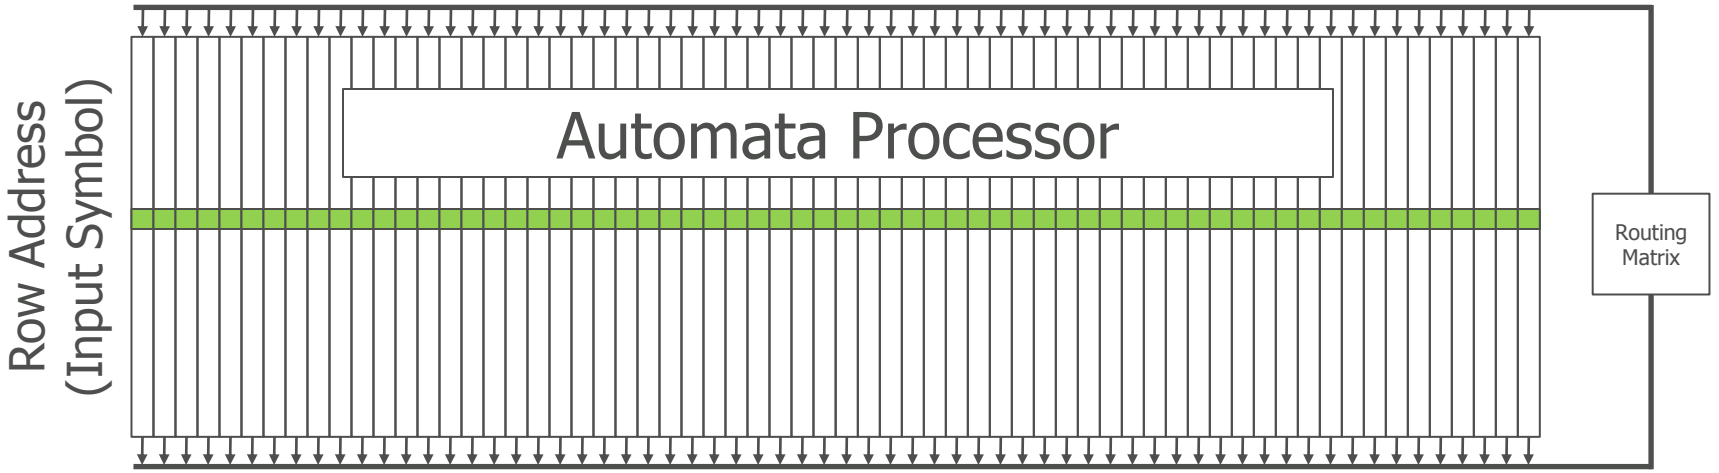

Row Access results in **49,152** match & route operations.

Unstructured Data – Unstructured Processor

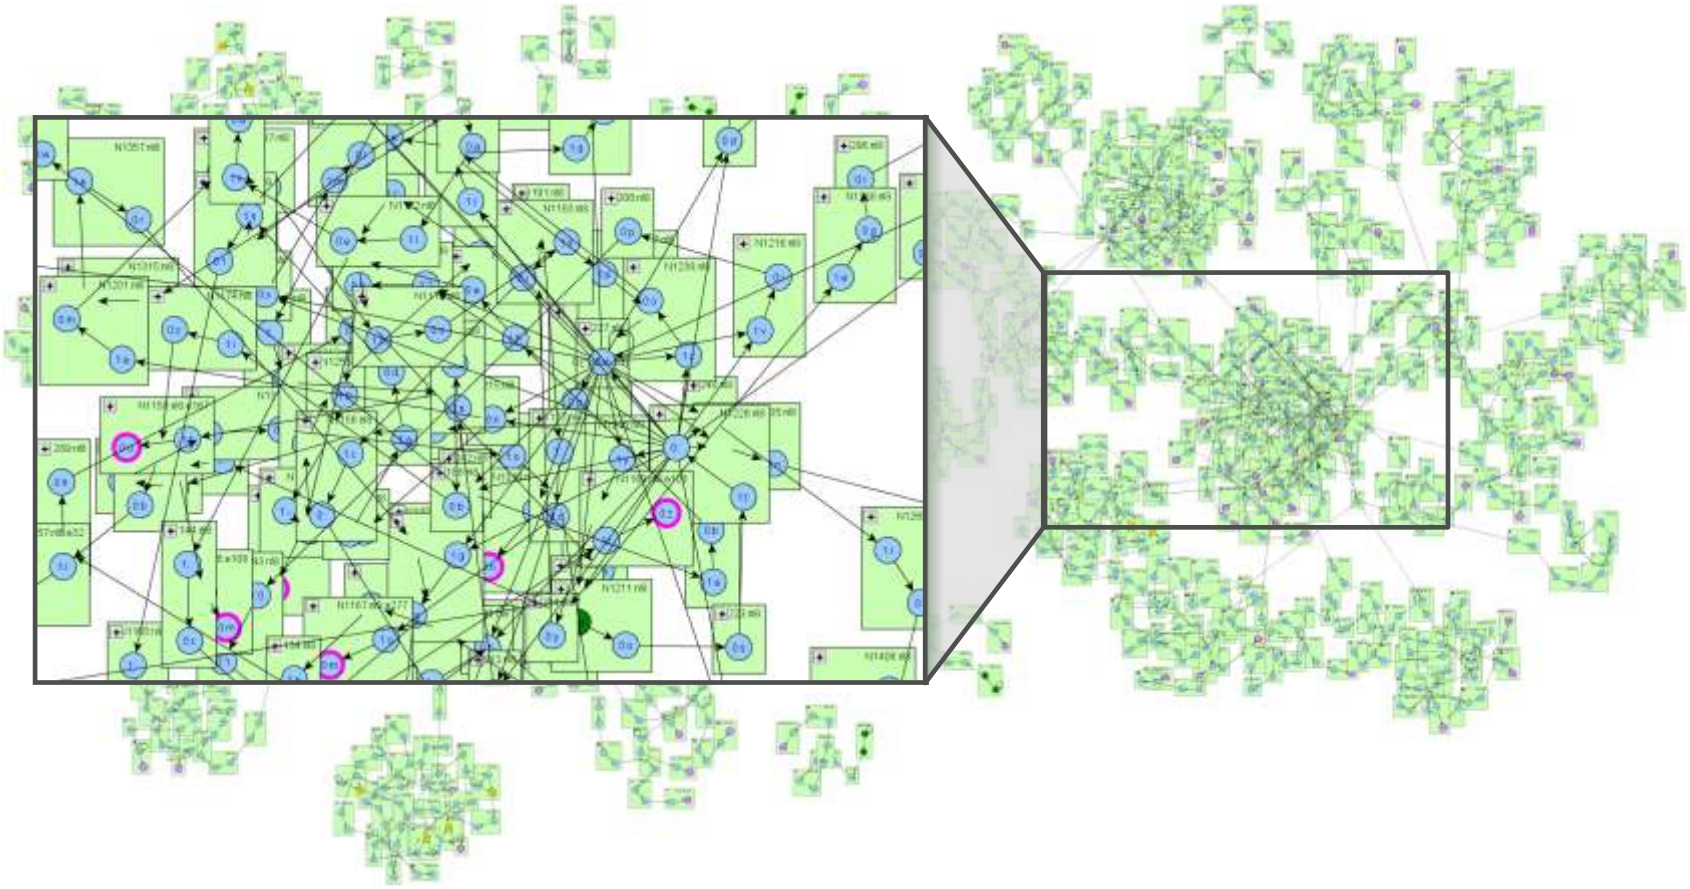

ROW Access is Fully Utilized on Every Cycle

The Linley Group

MICROPROCESSOR *report*

Insightful Analysis of Processor Technology

"Micron's 48-chip evaluation board scales this bandwidth to a ridiculous 38TB/s, which enables Automata to solve problems that traditional processors cannot."

Automata Processor Workbench

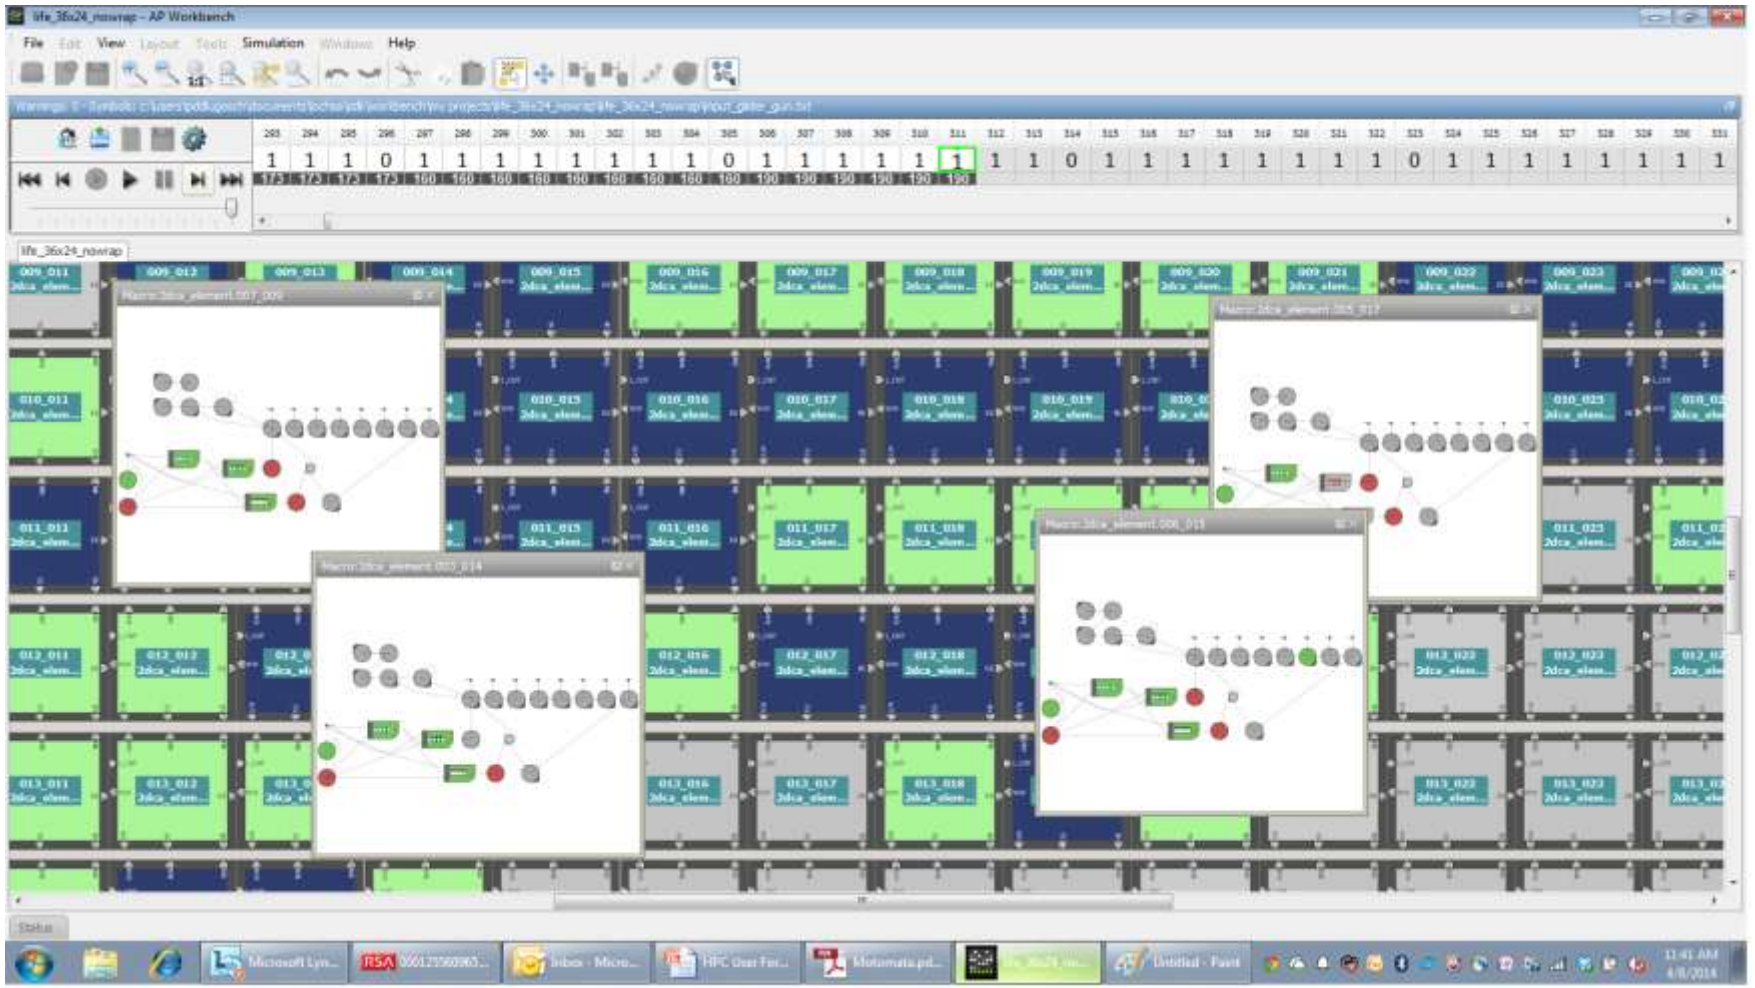

AP Workbench: Hierarchical Design & Macro Support

Programmer Productivity

Pattern #1 →

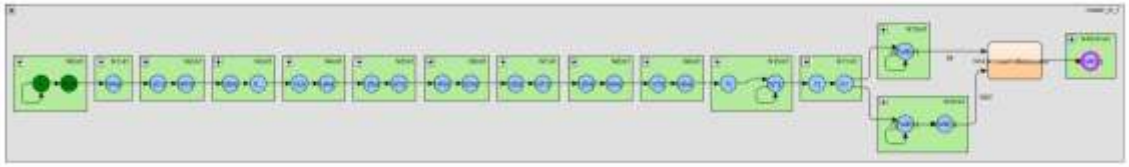

Pattern #2 →

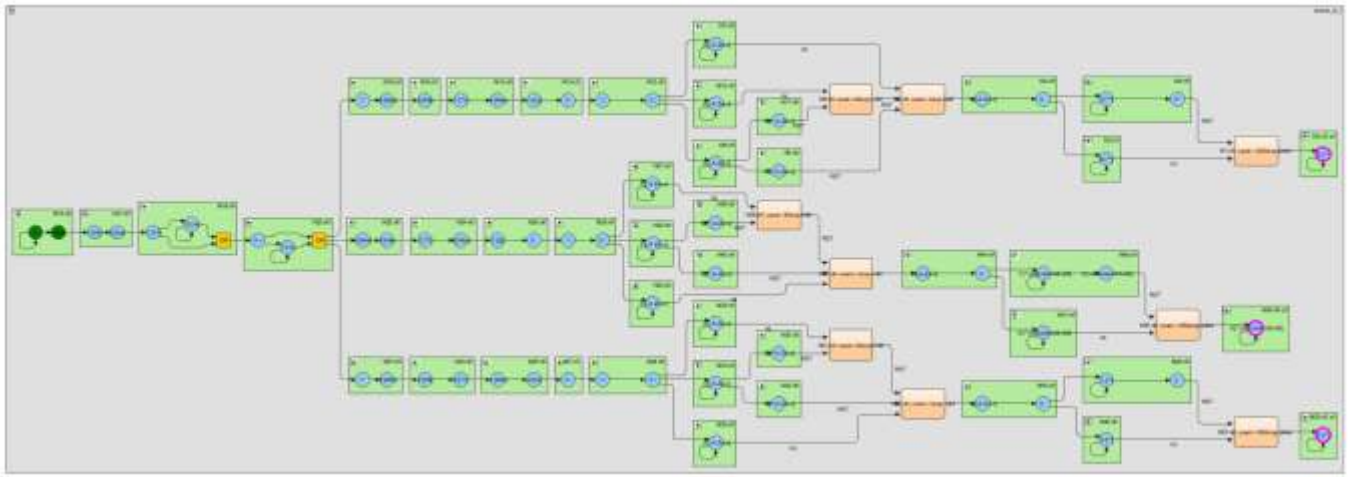

Pattern #3 →

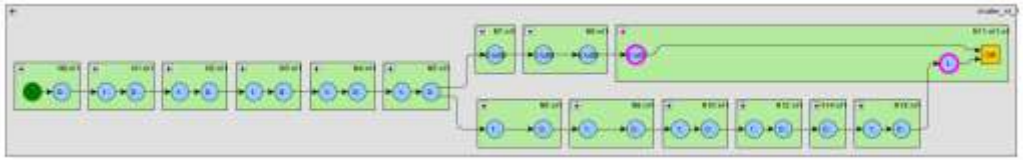

Parallelization of automatons requires no special consideration by the user. Each automaton operates independently upon the input data stream.

Design Methodologies: ANML, REGEX, Scripting, Full Custom

Application: Bioinformatics
Function: Fuzzy Matcher

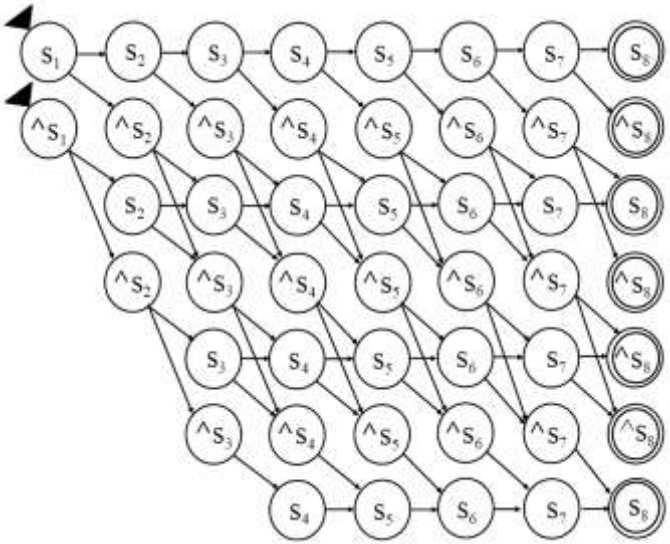

Compiler Input: Scripted ANML

Application: Cyber security
Function: DoS Attack Apache

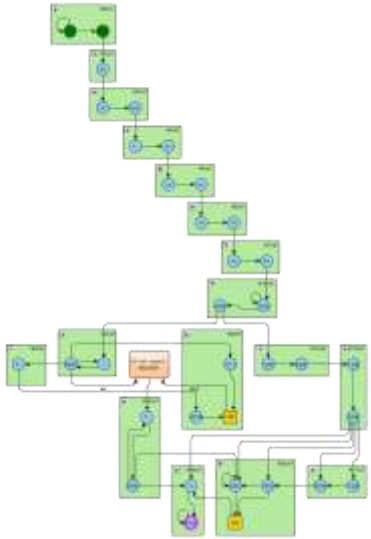

Compiler Input: REGEX

Automata Processor SDK supports a variety of input methods.

Automata Processor: Scalability

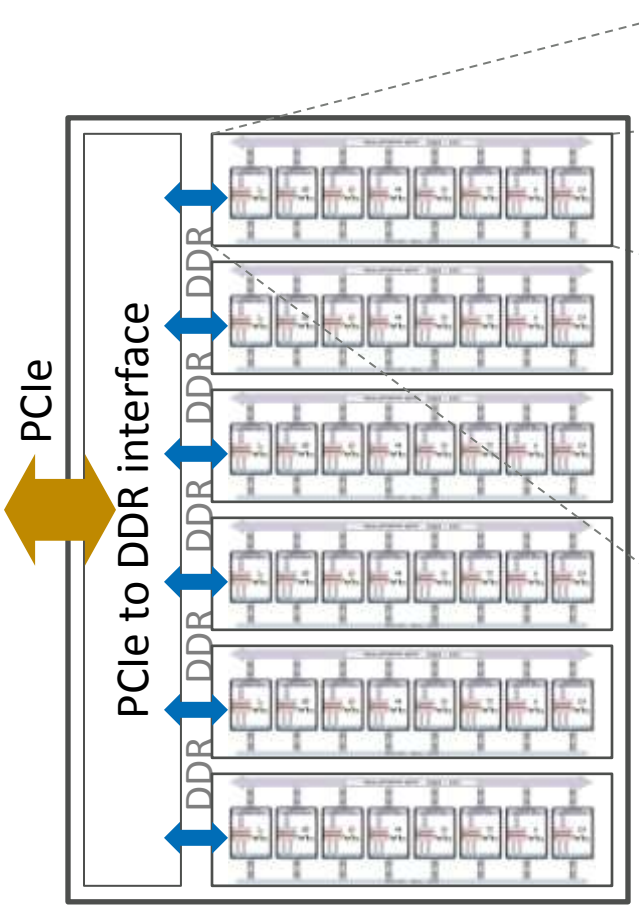

Automata Processor board

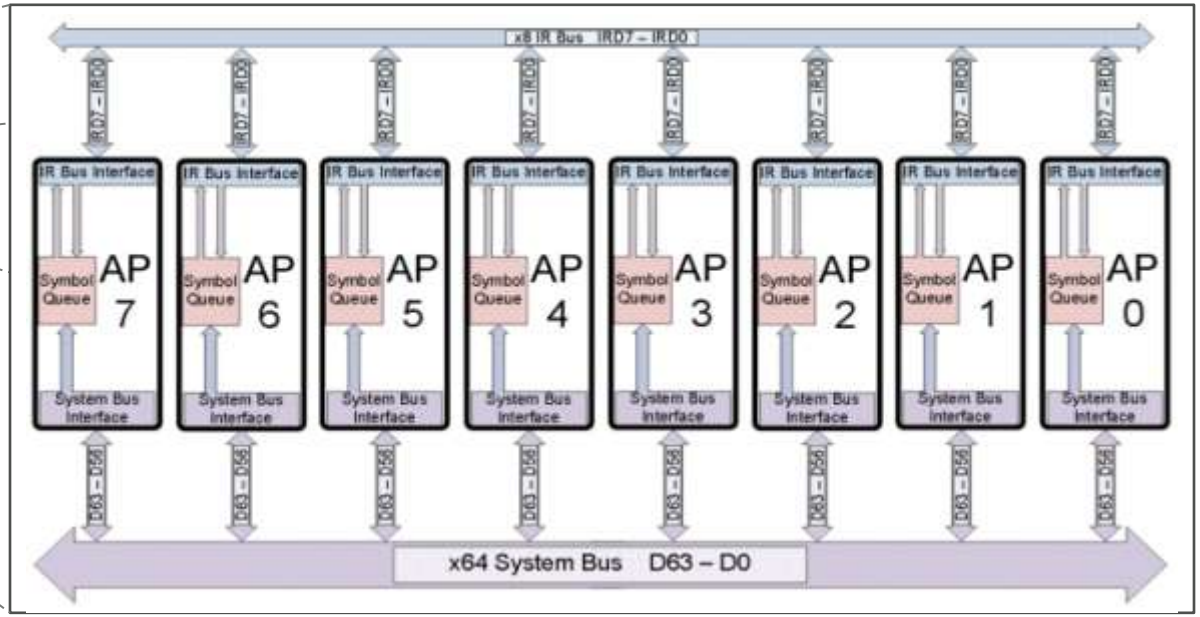

Rank of Automata Processor chips

Scaling:

- Configurable for speed vs. capacity.
- Host controller must manage streams.

Automata Processor: Support & Tools

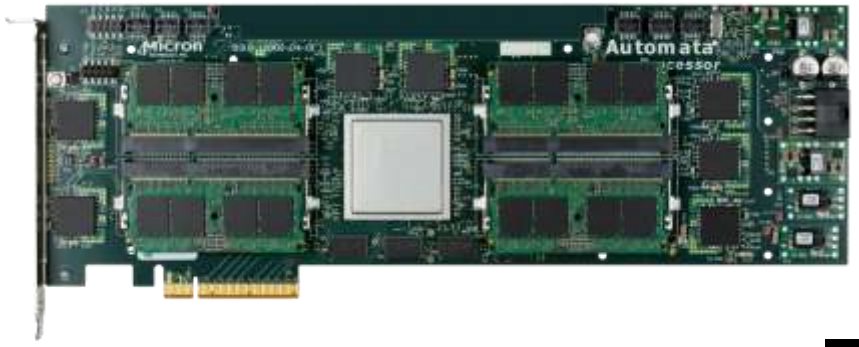

PCIe Development Board

- Industry Standard PCIe bus interface
- Capacity for up to 48 AP's
- Large FPGA capacity
- DDR3 for local storage

Software Development Kit

AP Optimization, loading & debugging tools & compiler.

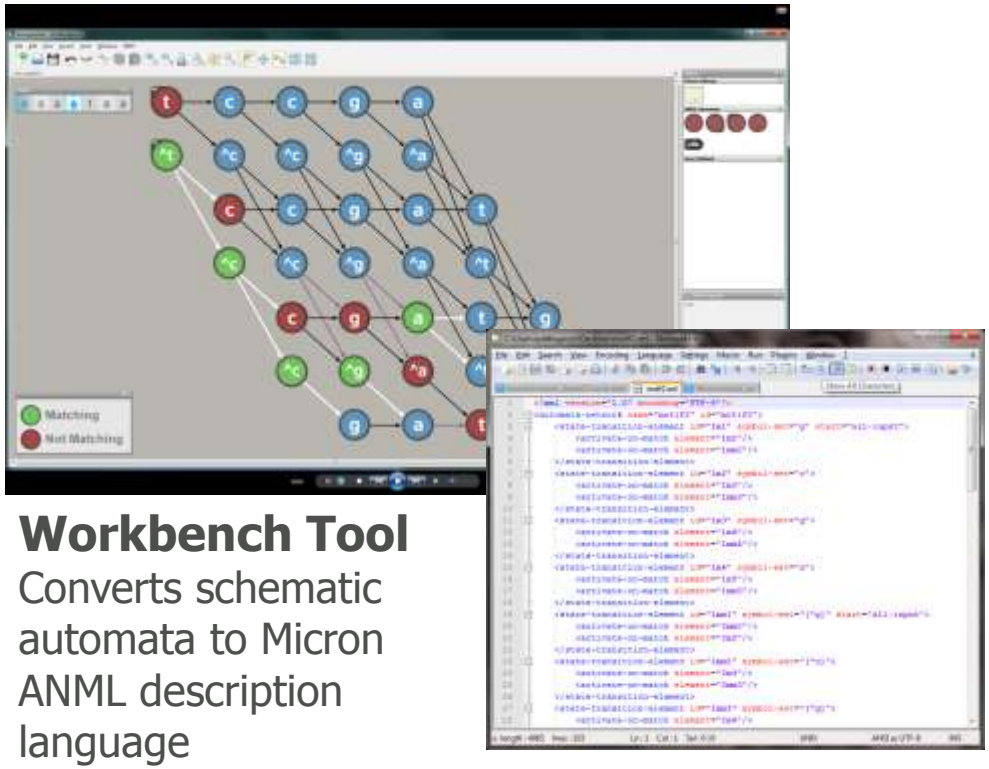

Workbench Tool

Converts schematic automata to Micron ANML description language

A Fundamental Breakthrough

CPU Block Diagram

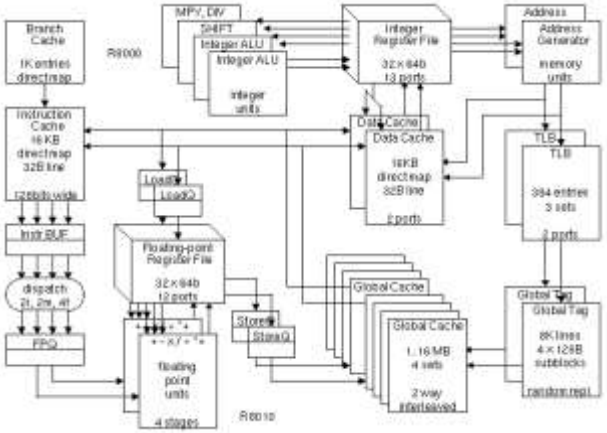

Human Neocortex

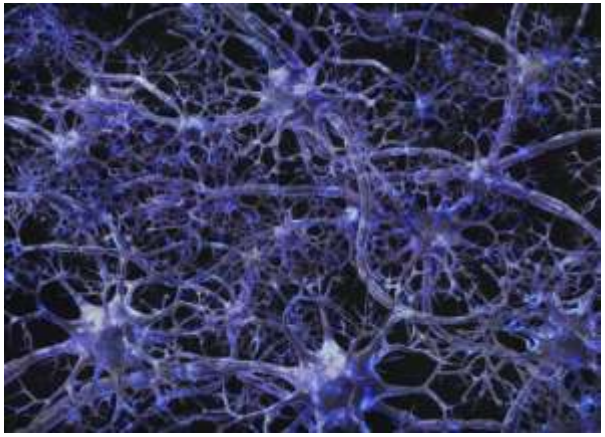

CPU Program

```
ORG   0H        ;start (origin) at location 0
MOV   R5,#25H   ;load 25H into R5
MOV   R7,#34H   ;load 34H into R7
MOV   A,#0      ;load 0 into A
ADD   A,R5      ;add contents of R5 to A
                ;now A = A + R5
ADD   A,R7      ;add contents of R7 to A
                ;now A = A + R7
ADD   A,#12H    ;add to A value 12H
                ;now A = A + 12H
HERE:SJMP HERE  ;stay in this loop
     END        ;end of asm source file
```

Automata Processor Program

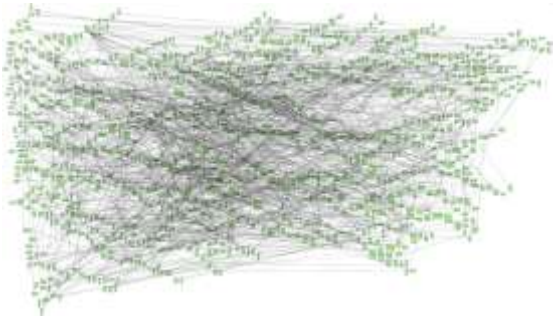


Micron®
Focused on Memory | Engineered for your Innovation