Switch Chip Evolution

Uri Cummings
Co-founder, CTO

FULCRUM microsystems

Fulcrum's Next Generation

- First to 1 BBPS and 300nS, no corner cases
- First programmable and deterministic pipeline

Fulcrum Innovation Track Record

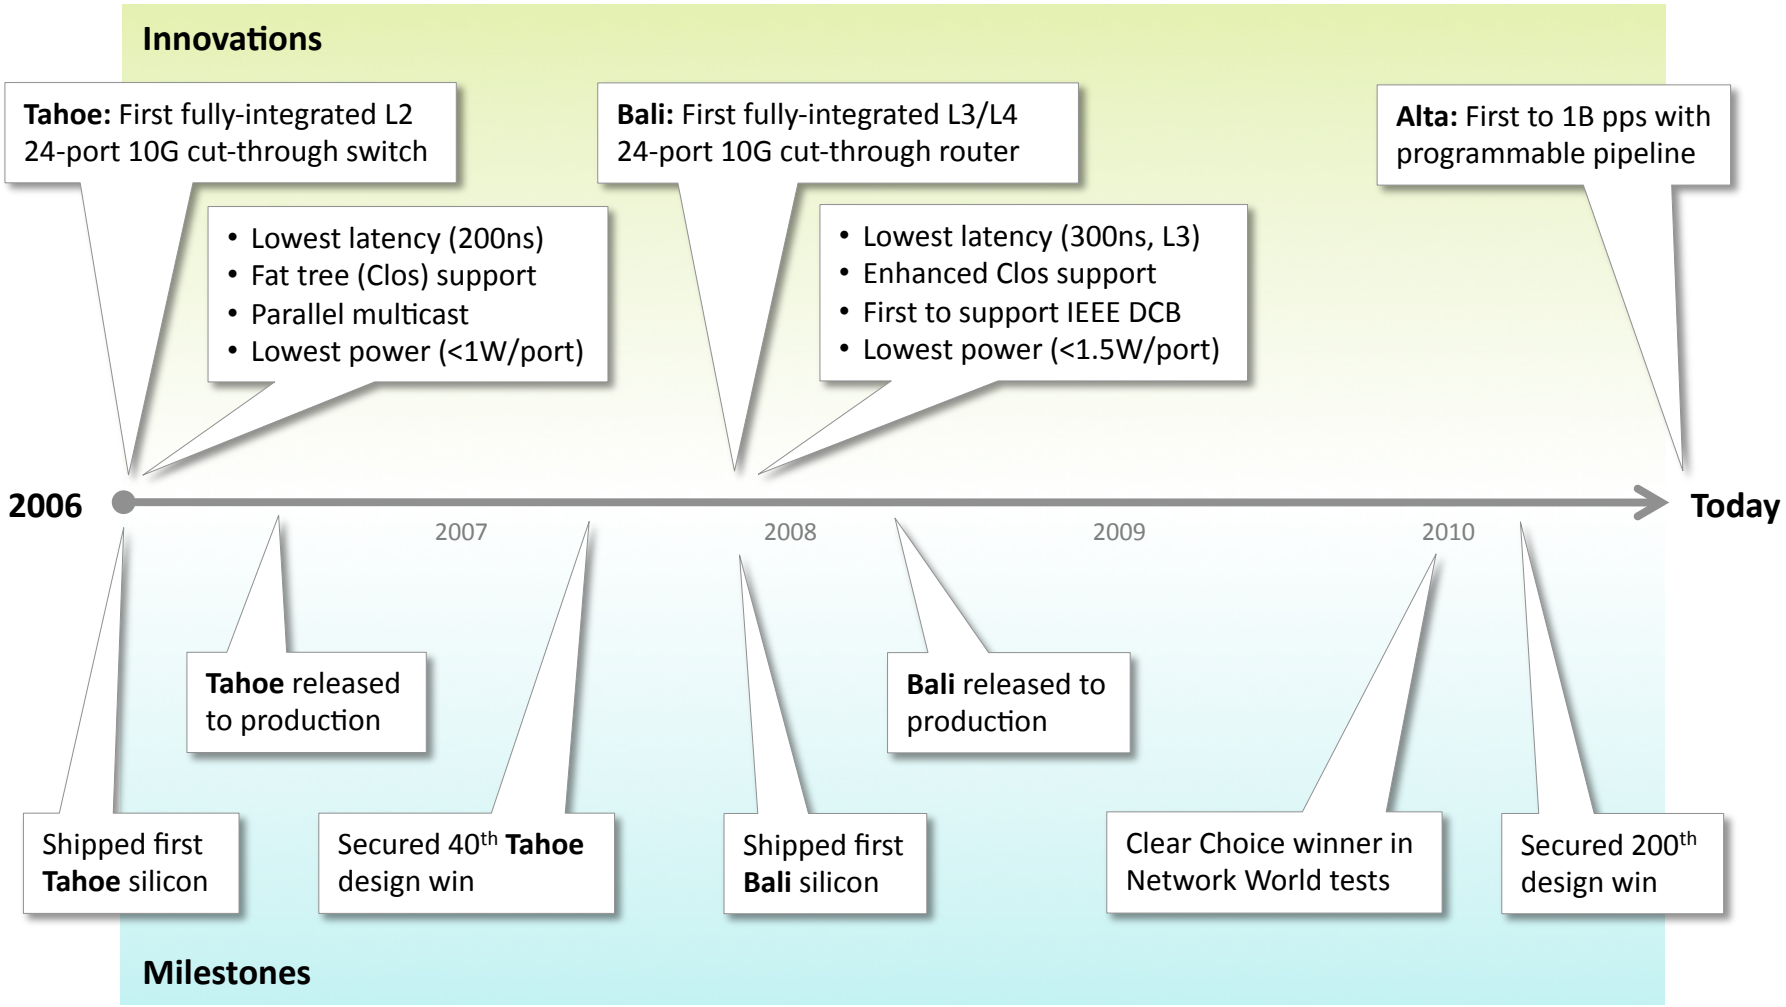

FULCRUM microsystems

Flat Networks Need High Radix Switches

Clos in the cloud

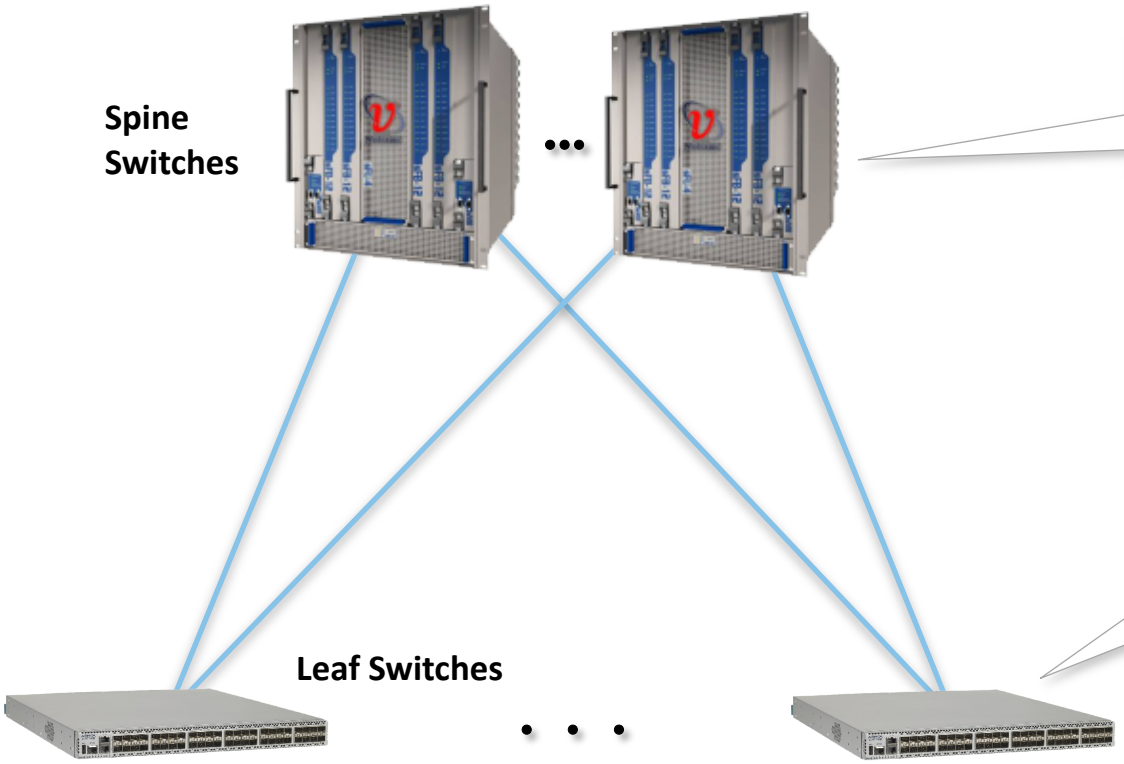

The Big Switch
- Significant investment
- Doesn't change that often

TOR Layer
- Create a Clos with the spine
- Port density determines scale
- Full bandwidth (no aggregation)

A New Class of Chip for Networking

Programmable vs. deterministic pipeline

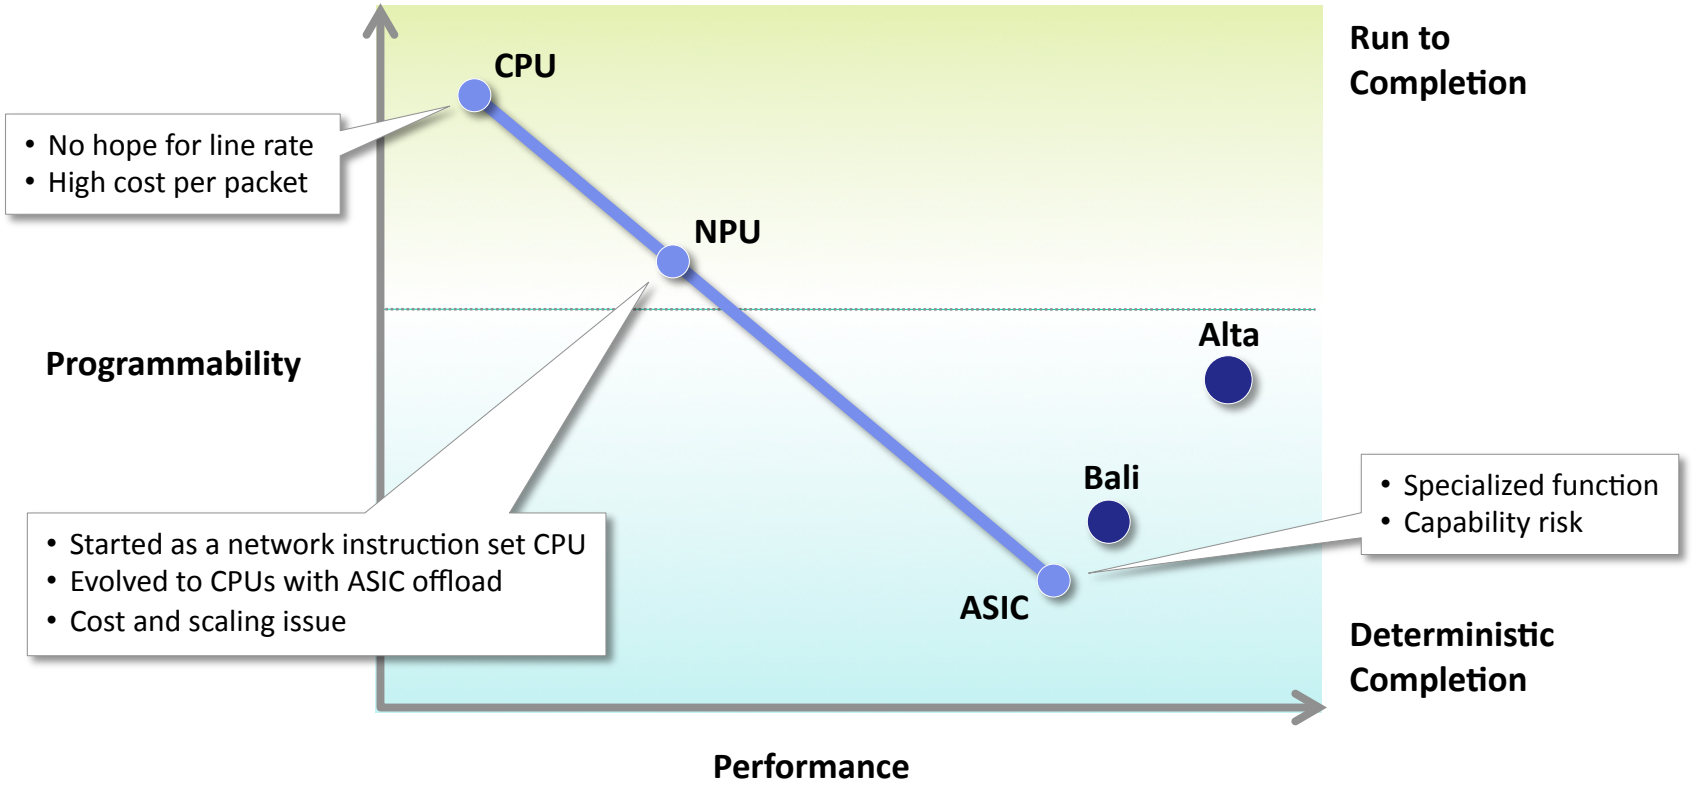

FULCRUM microsystems

Switch Architecture Best Practices

Holy grail: Fully-provisioned, output-queued, shared memory

- **Deterministic frame processing**
  - Enhances performance
  - Avoids corner cases
- **Pure shared memory**
  - Efficient buffering
  - Low latency
  - Full multicast
- **Compact, flexible port logic**
  - Drive wire cost and density
- **Fully provisioned**
  - No corner cases

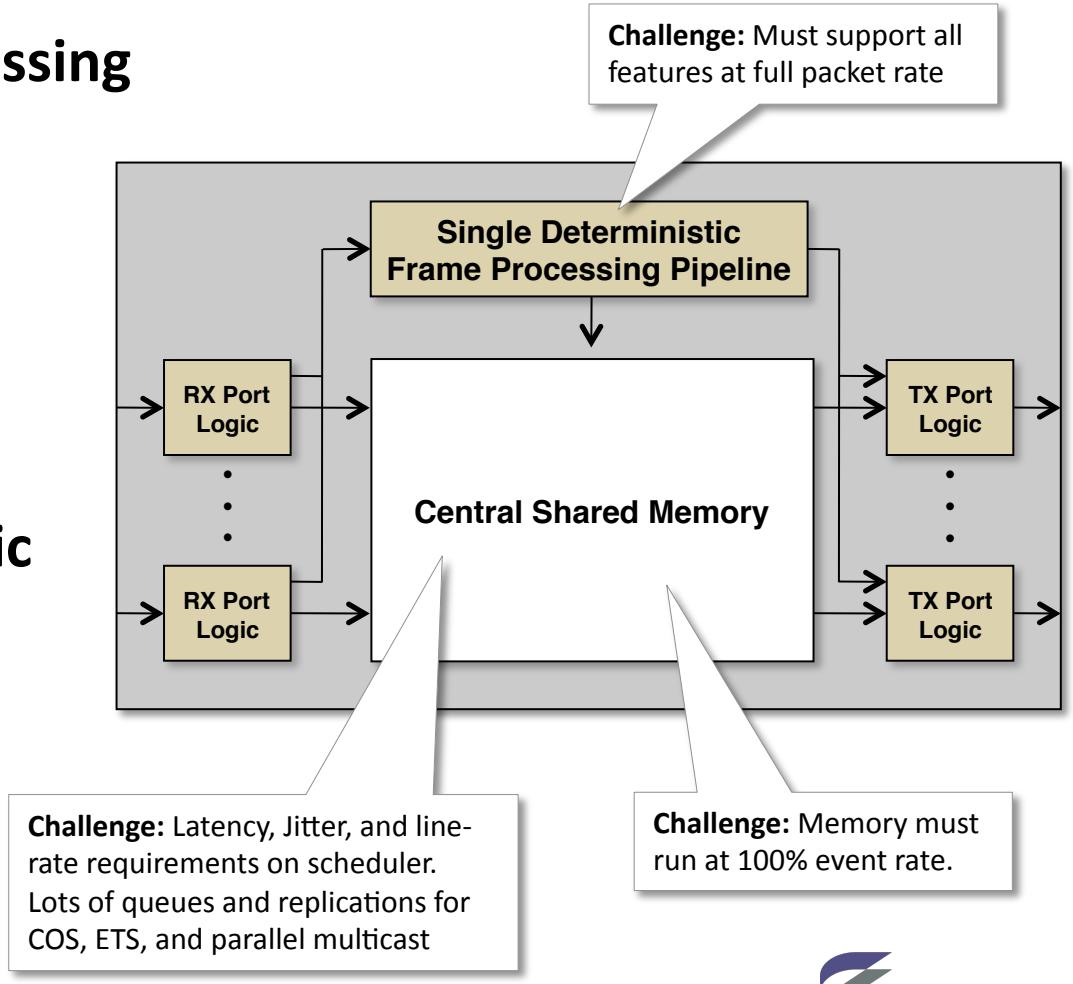

Common Compromises in the Market

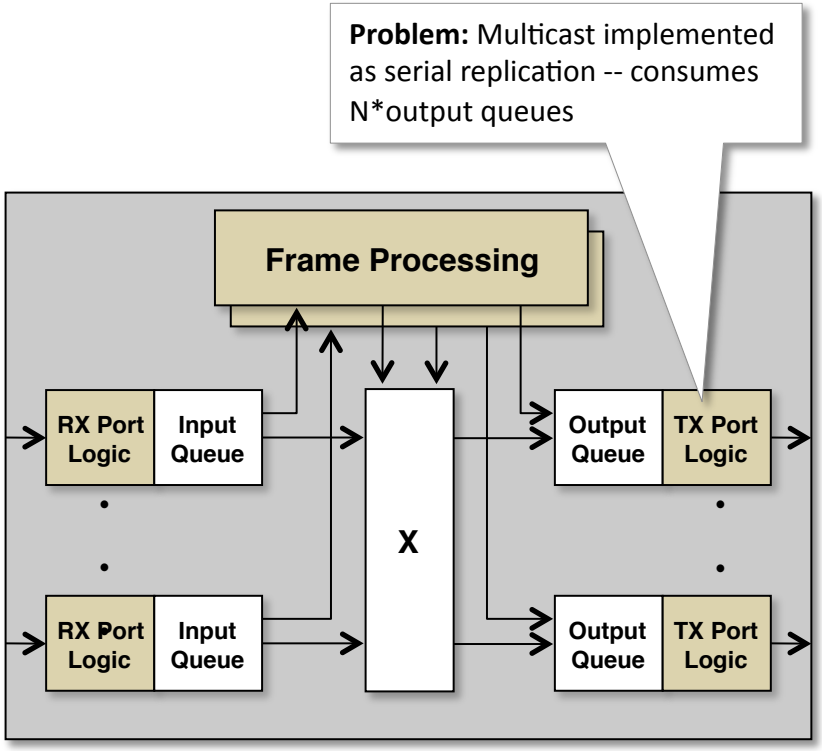

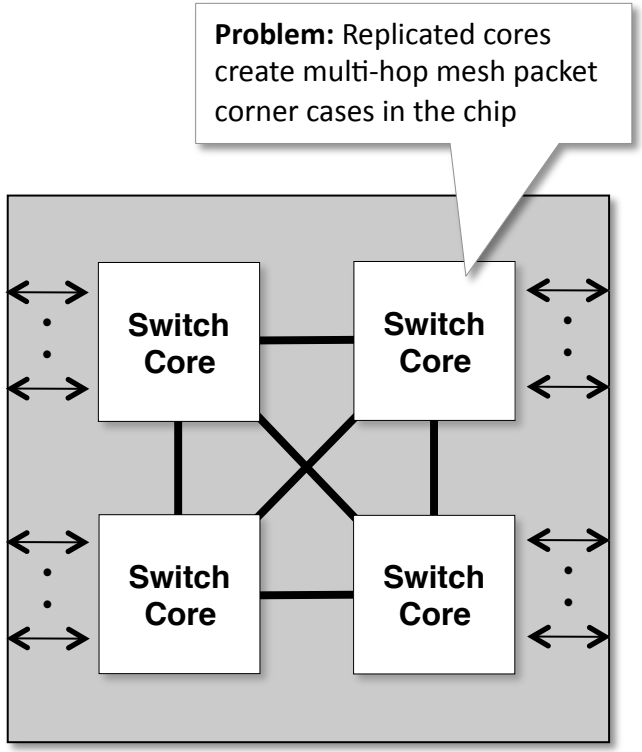

- **Even some other shared memory switches have significant limitations**
  - Must have a pseudo-dual ported memory that is greater than the packet rate (1BPPS)
  - Frame processing pipeline must also be > 1 GHz for all features
    - Hashing and other cache concepts don't work for layer 3

The Alta Architecture

Output queued, shared memory. No compromises – ever!

- ***Solution Performance***
  - 1 BPPS, fully provisioned
  - 300 nS latency
- ***RapidArray***
  shared memory
  - Fastest crossbars and memories
  - Lowest latency scheduler
- ***FlexPipe***
  frame processing pipeline
  - Supports full event rate
  - Deterministic & programmable
- **Compact, flexible port logic**
  - Every feature on any lane
  - Integrated PHY
- **Fully provisioned**

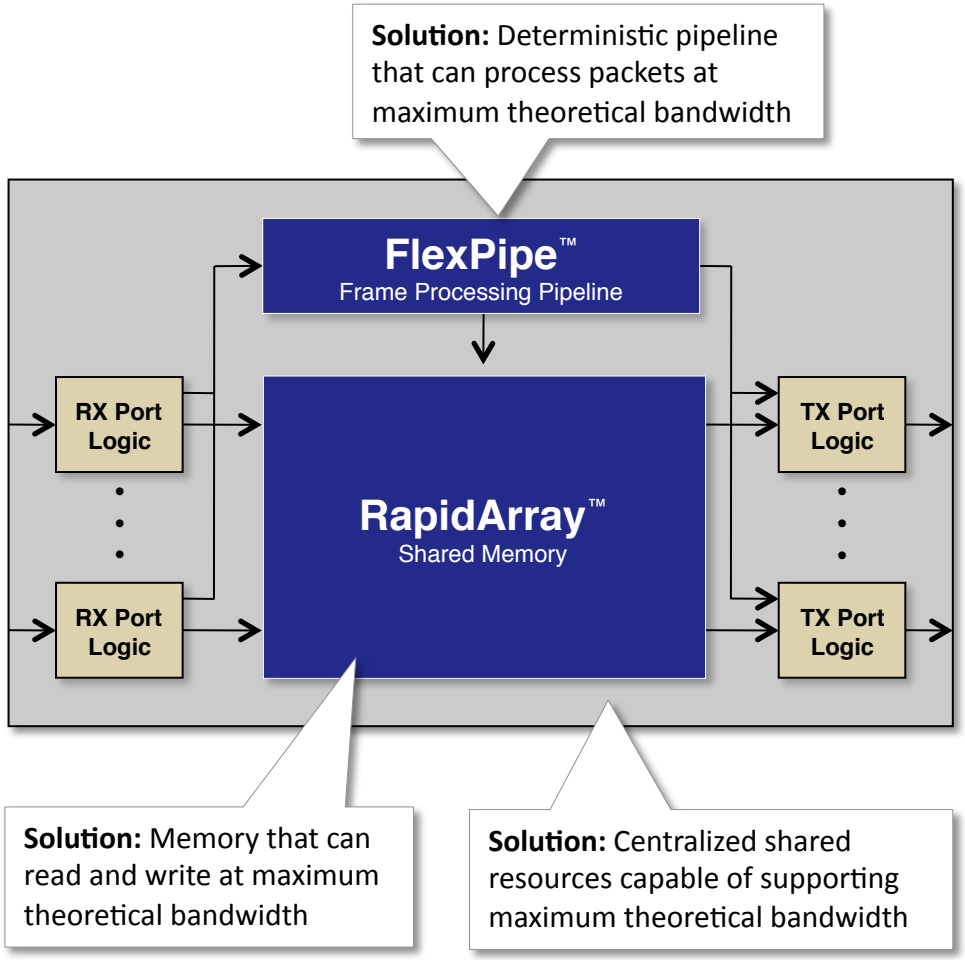

RapidArray™ Shared Memory Performance

- **Flexible reads and writes**
  - Any port can write to any location
  - Every port can read from any location
  - A packet can be any size, 64B-10kB (including 1B larger than you want)
  - All packet rewrite on Egress at line rate
- **High-throughput scheduling**
  - Separate link list for RX and TX queues
  - In segment cut-through
  - Parallel no-drop multicast
  - Transmission Selection
  - Segment usage reported per segment time to parallel watermark calculations
- **Unparalleled performance**
  - 1BPPS
  - About 1000 queues
  - 67 nS scheduling ring
  - Fine-grained segment size

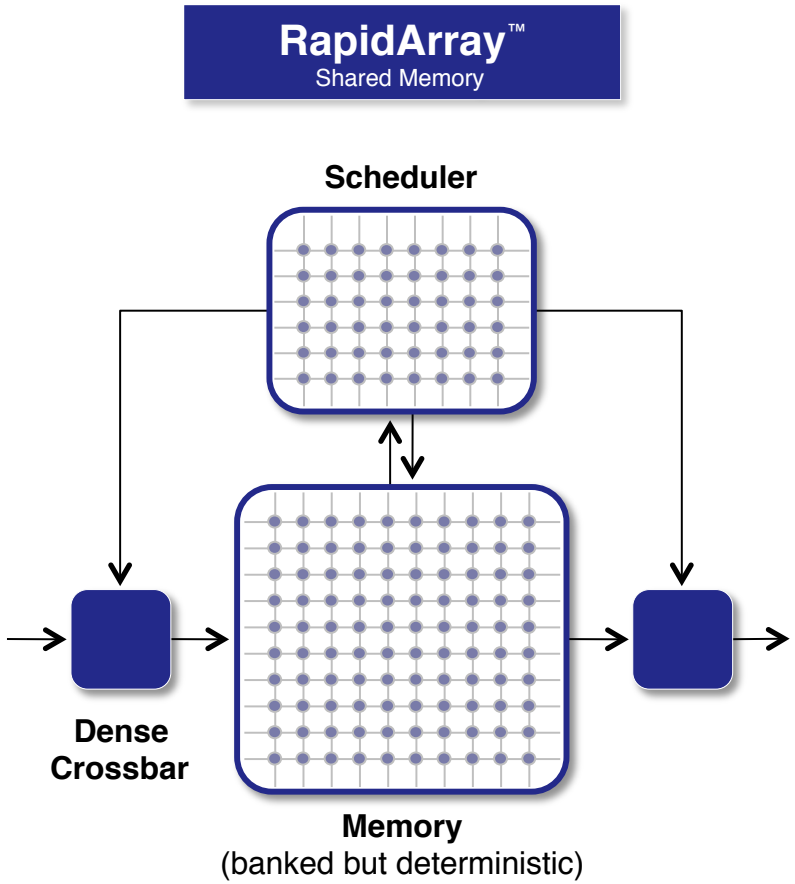

Alta's FlexPipe™ Packet Processing

Programmable and deterministic pipeline

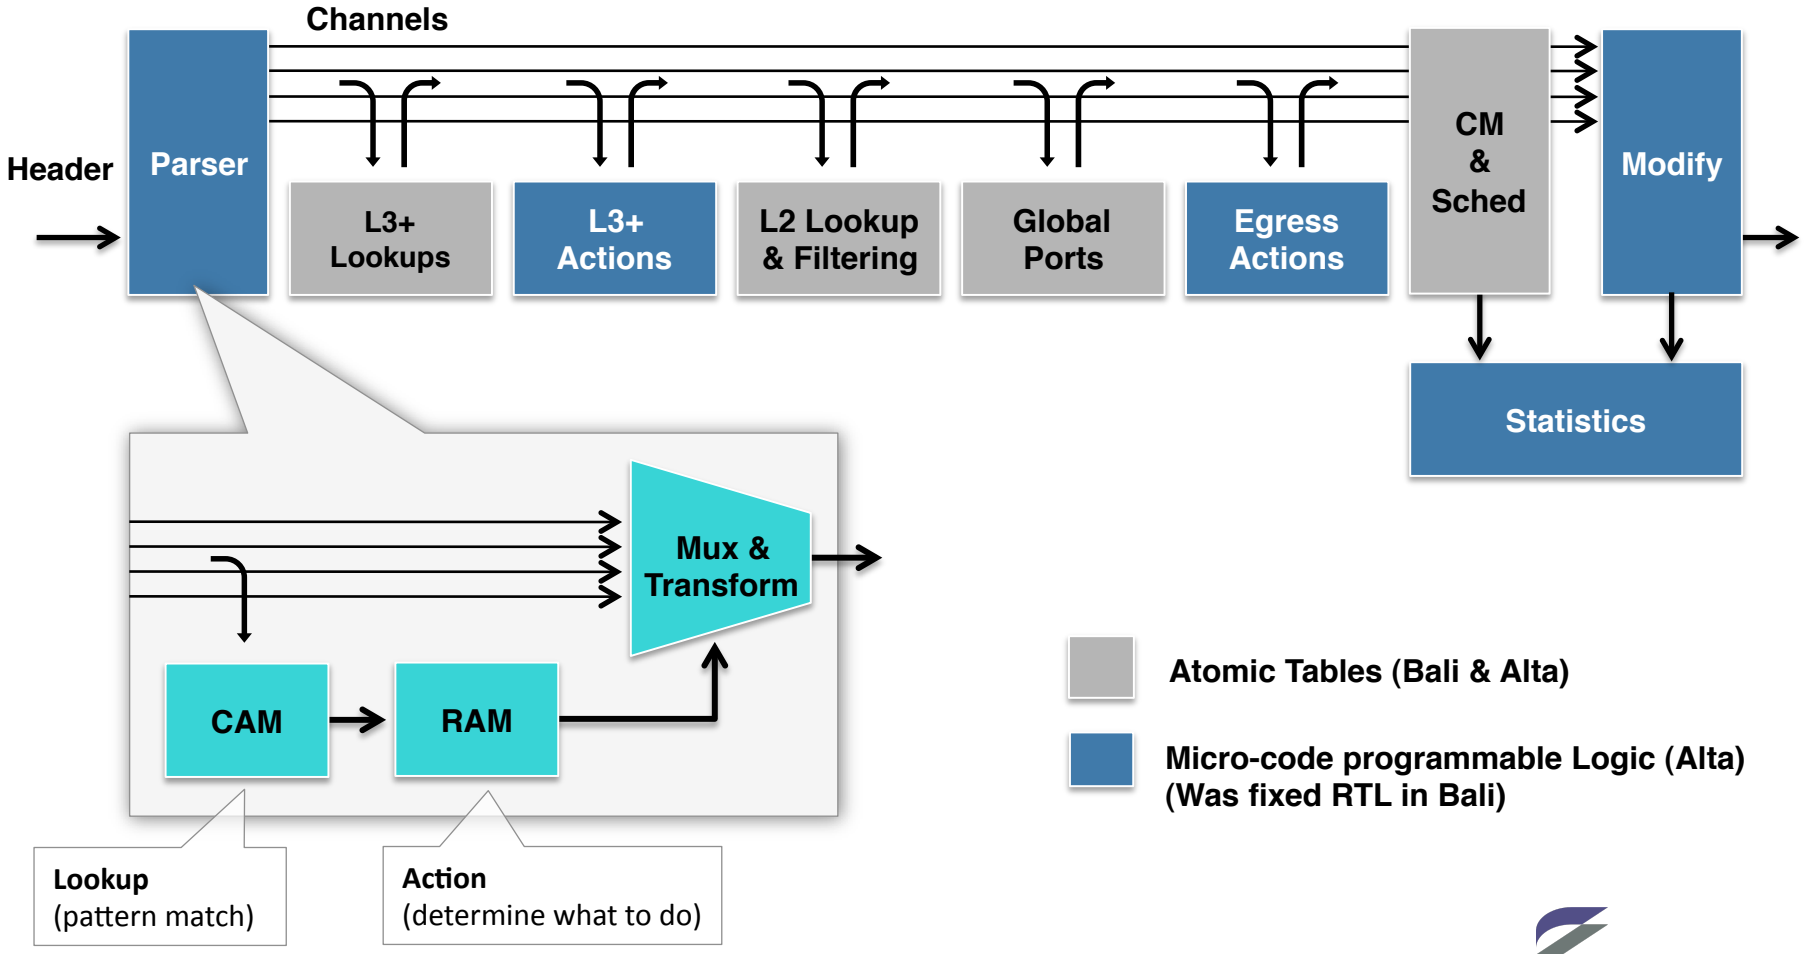

The FlexPipe™ FFU Architecture

Efficient filtering and forwarding using CAM and BST

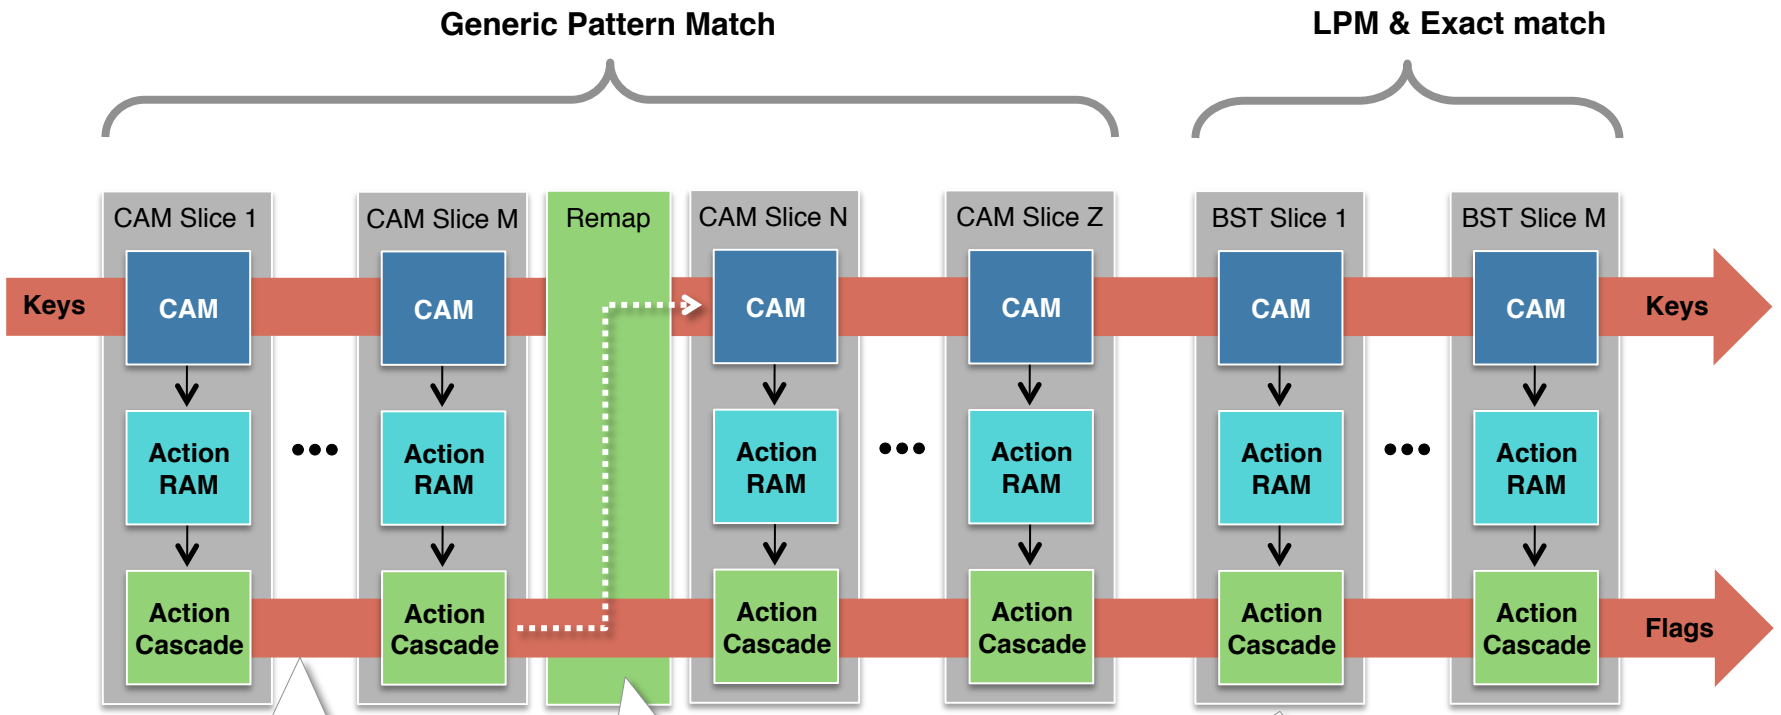

FULCRUM microsystems

Companion devices with FocalPoint

Application Accelerators with FocalPoint

- **FlexPipe programs "application tags" exposing switch resources to the application**
- **Further processing done in application accelerator**

Example Application

- **market data message switch**

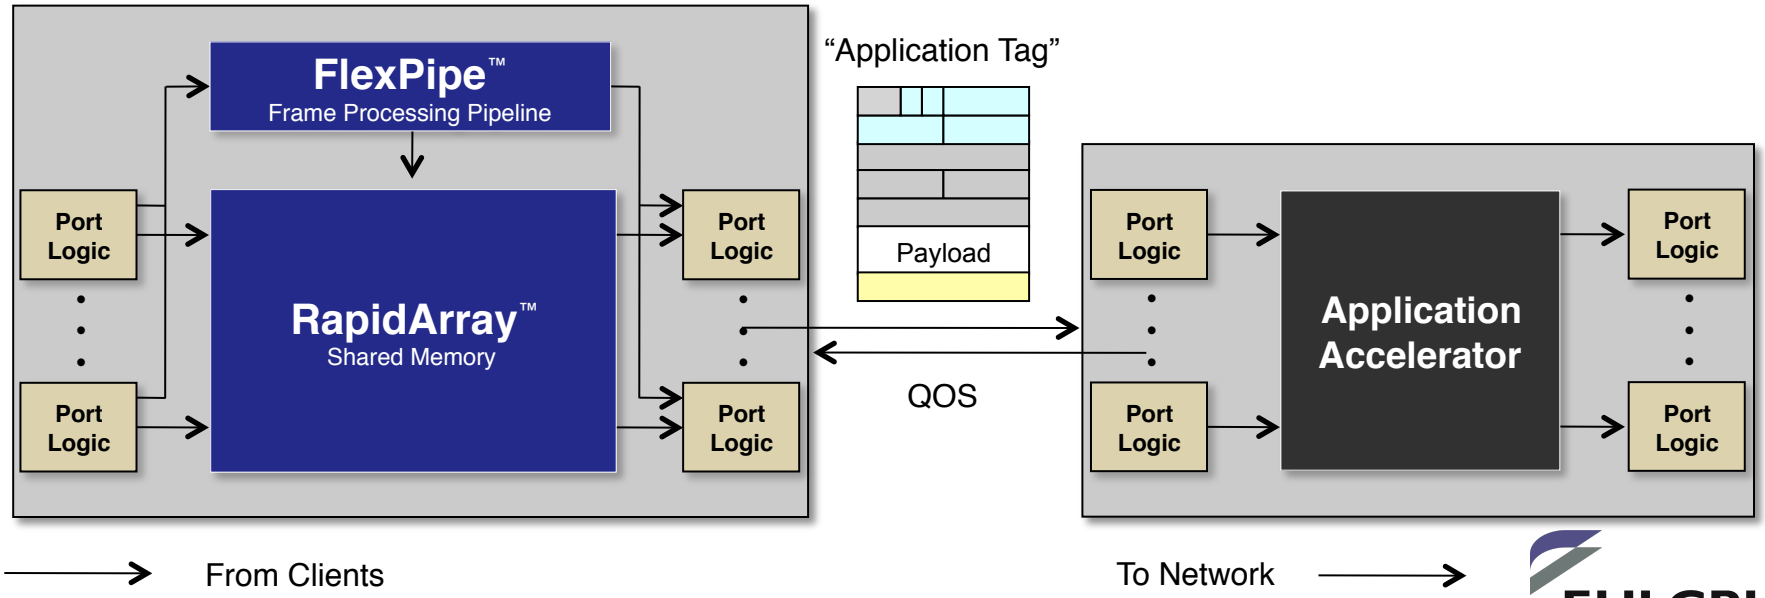

Thank You!

Uri Cummings

Founder, CTO

uri@**fulcrum**micro.com

FULCRUM microsystems

818.871.8100

www.**fulcrum**micro.com

26630 Agoura Road

Calabasas, CA 91302