NEXTSILICON

THE NEXT REVOLUTION OF HARDWARE IS SOFTWARE

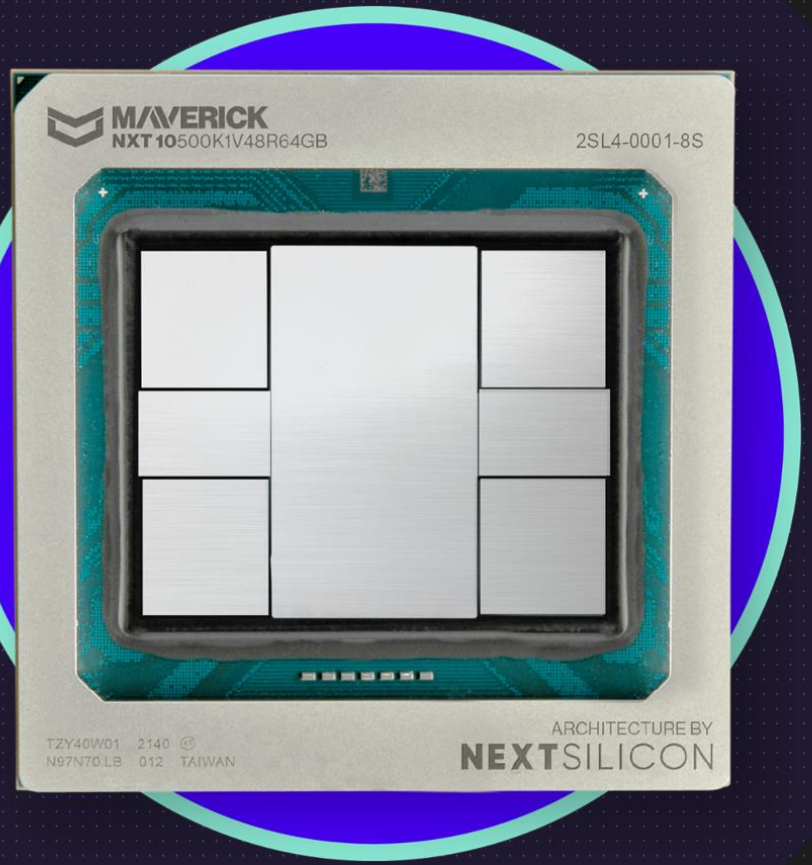

INTRODUCING:

DYNAMIC COMPUTE

A platform that reconfigures itself in run time.

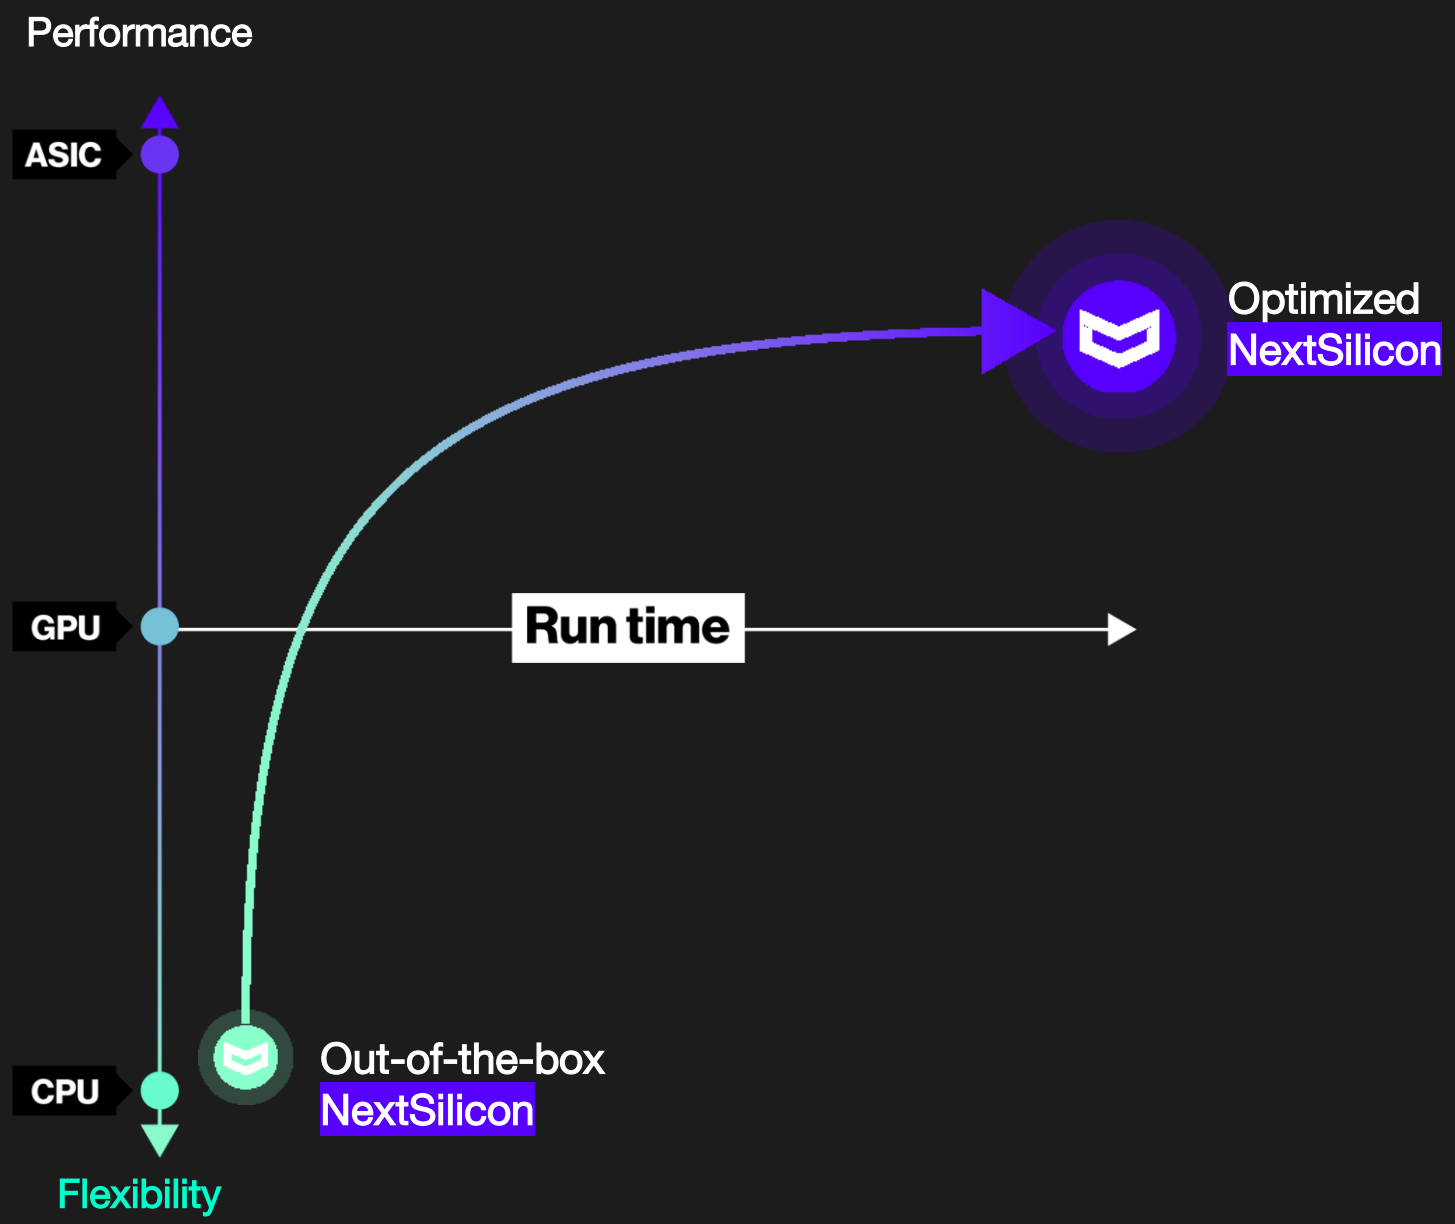

Workload-Aware, Software Defined Hardware

WHY

- **Novel architecture** - Non Von Neumann - presenting new opportunities
- **Asymmetrically** advantage run-time workloads
- **Minimal optimization** required for your multi-threaded CPU code
- **Massive parallel** execution
- **DESIGNED for HPC** workloads

STEP 01

IDENTIFY THE LIKELY FLOW

- In massively parallel applications, a small portion of code (the **“likely” flows**) takes up most of the run time
- NextSilicon **automatically** identifies likely and unlikely flows and performance bottlenecks
- NextSilicon reconfigures the hardware **asymmetrically** to favor the likely flows, so applications run orders-of-magnitude faster

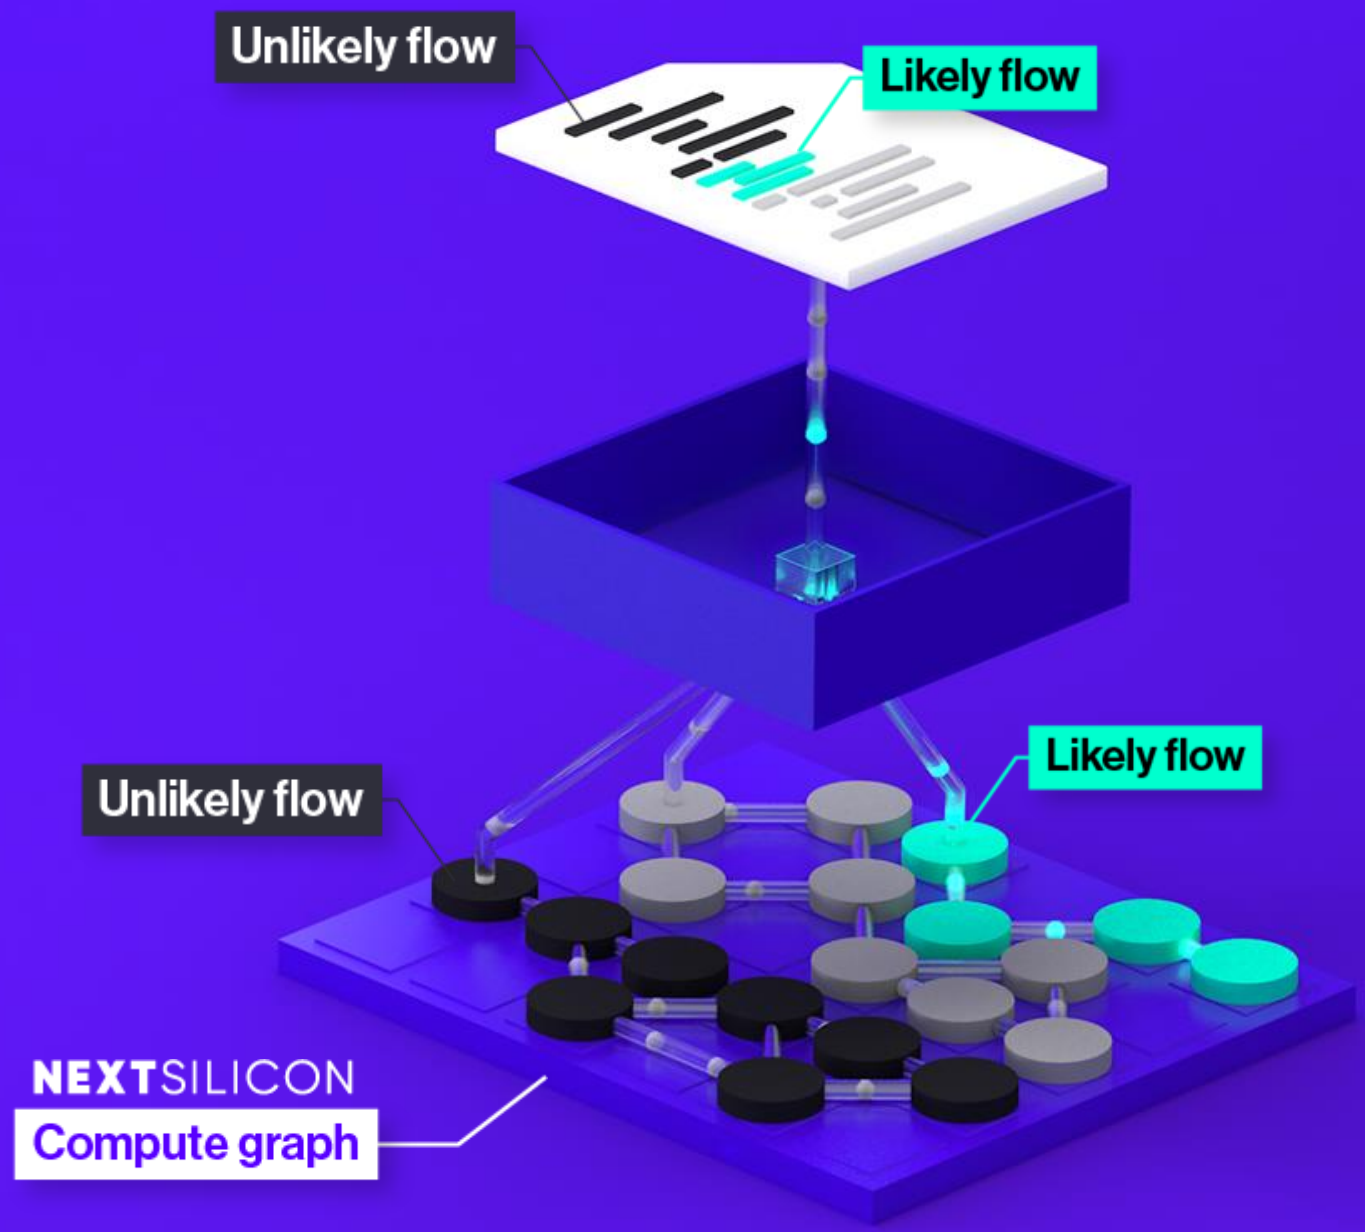

STEP 02

OPTIMIZE THE LIKELY FLOW

- We use an **iterative** optimization process at **run time** to **reconfigure** likely flows onto our dataflow grid
- On-chip **telemetry** constantly relays runtime information to the optimization process
- Each iteration further **optimizes** the compute graph for likely flows, until the majority of resources are devoted to them
- We call a fully optimized grid a **mill core**

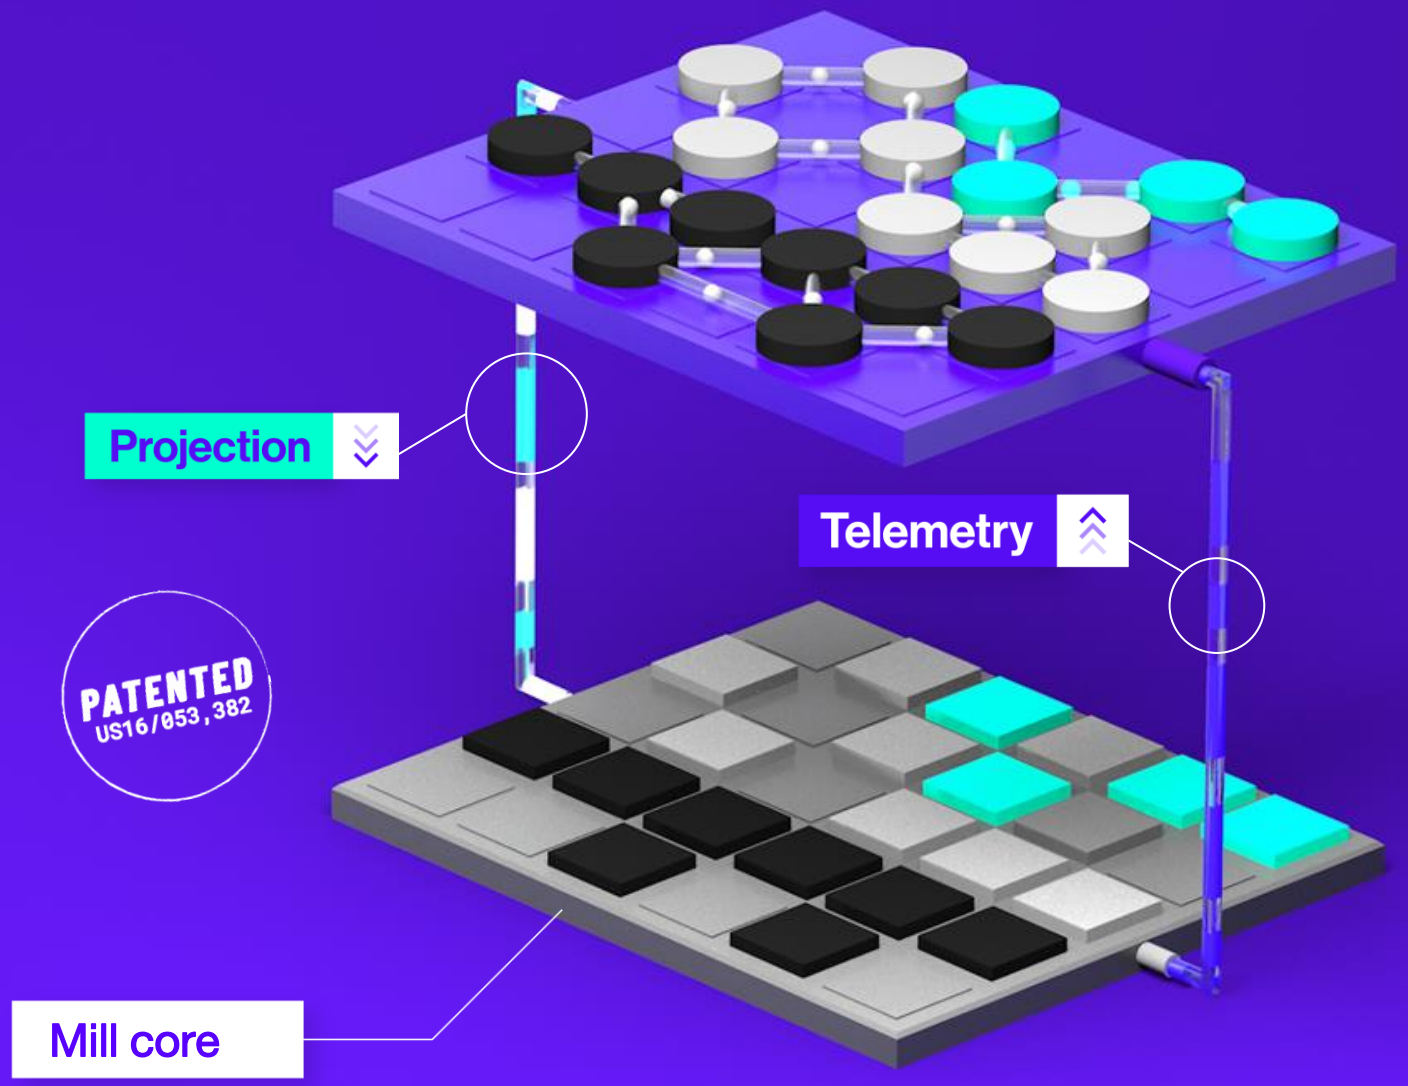

STEP 02

OPTIMIZE THE LIKELY FLOW

- A mill core is a **software-defined** processing core, optimized for a particular application by the runtime optimization process
- Mill cores are **optimized for throughput** over latency, with high **power efficiency**
- An optimized mill core can run **hundreds** or **thousands** data streams in parallel

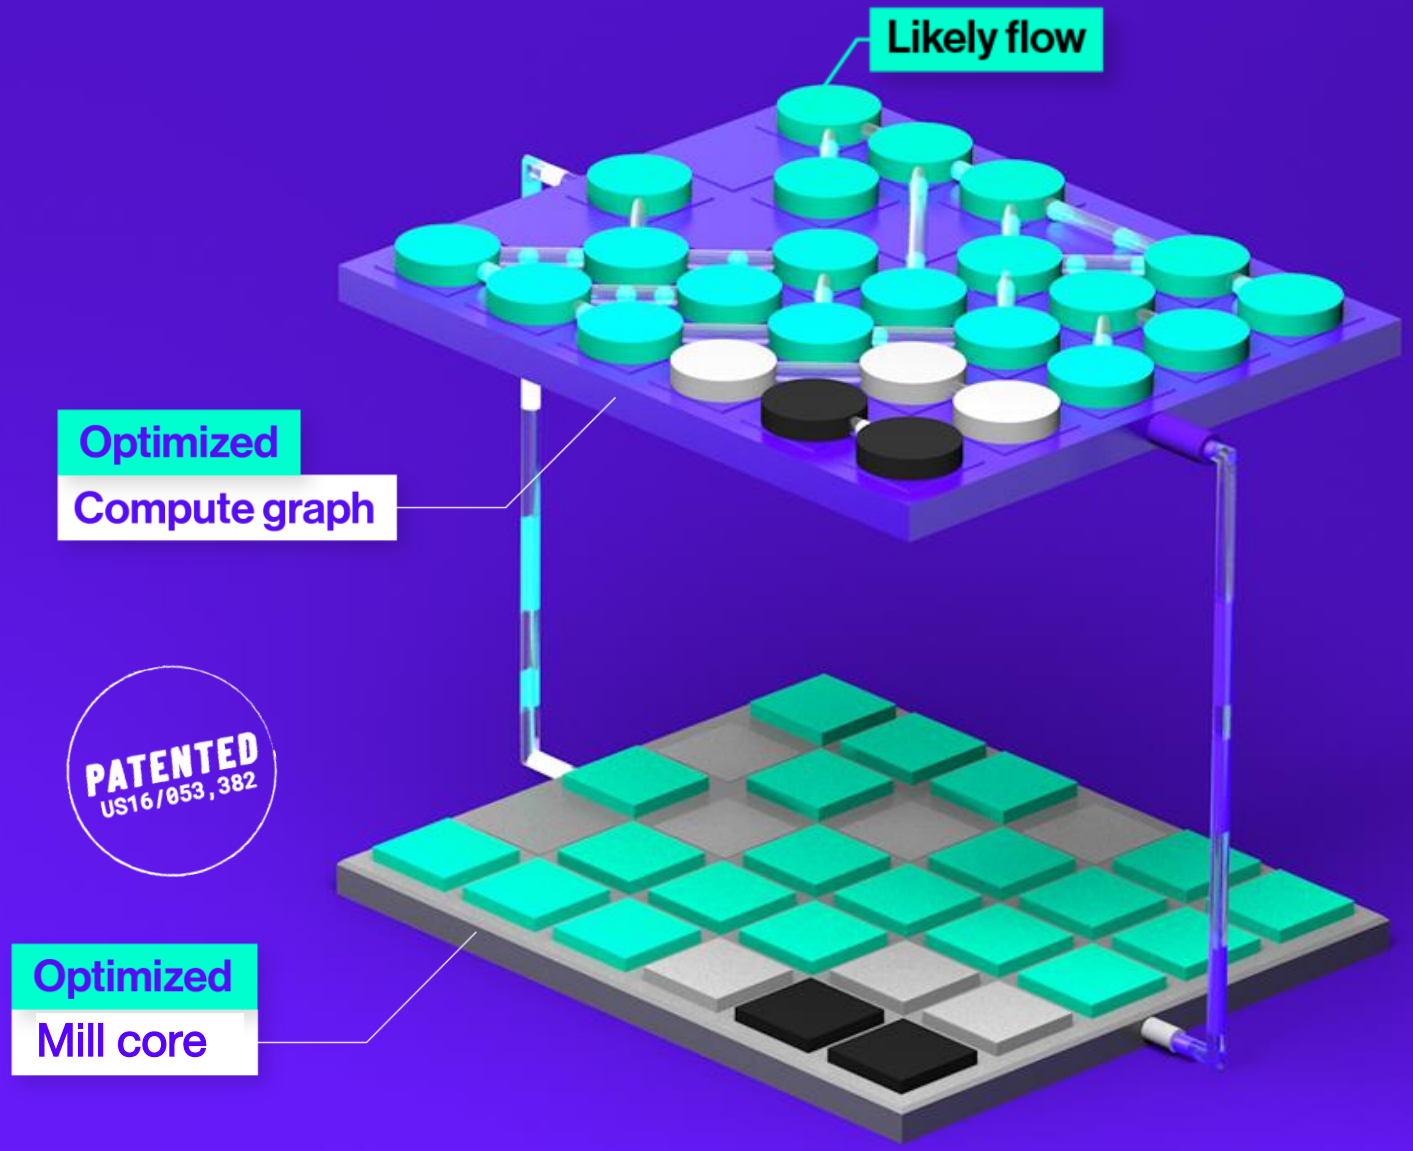

STEP 03

REPLICATE ACROSS THE CHIP

> Mill cores can be **replicated** hundreds of times for **massively parallel** processing

> Overall run times improve by **orders of magnitude**

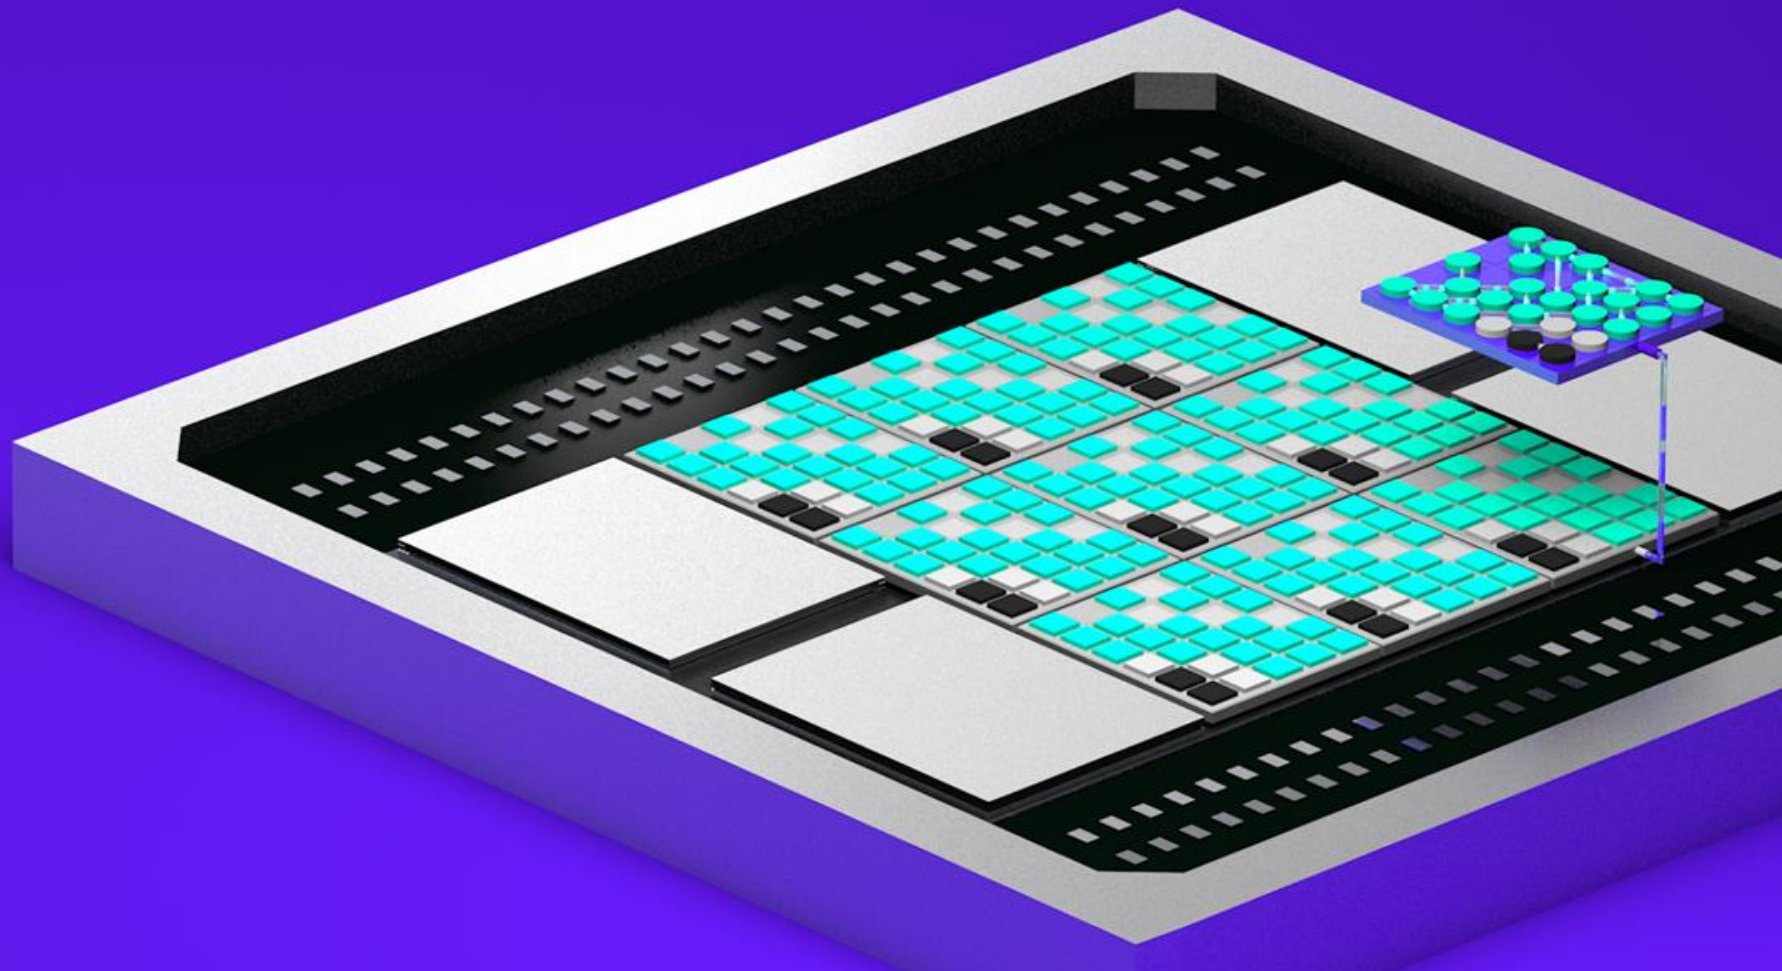

Novel Architecture = Novel Performance

The NextSilicon software suite intelligently adapts the Maverick processor into a workload specific ASIC, ***at runtime***. This novel architecture, with distributed HBM memory access, accelerates workloads the others can't

Some Unique characteristics:

- **Open languages** - no vendor specific lock-in
- No compromise in **FP64 performance**
- Highest **random access** memory bandwidth
- Automatic memory management
- **Atomic** operations without **penalty**
- **Extremely Sparse** workloads run **efficiently**
- Superior **branchy** code **performance**
- **Accelerated Integer performance** in hardware

NEXTSILICON

THANK YOU.