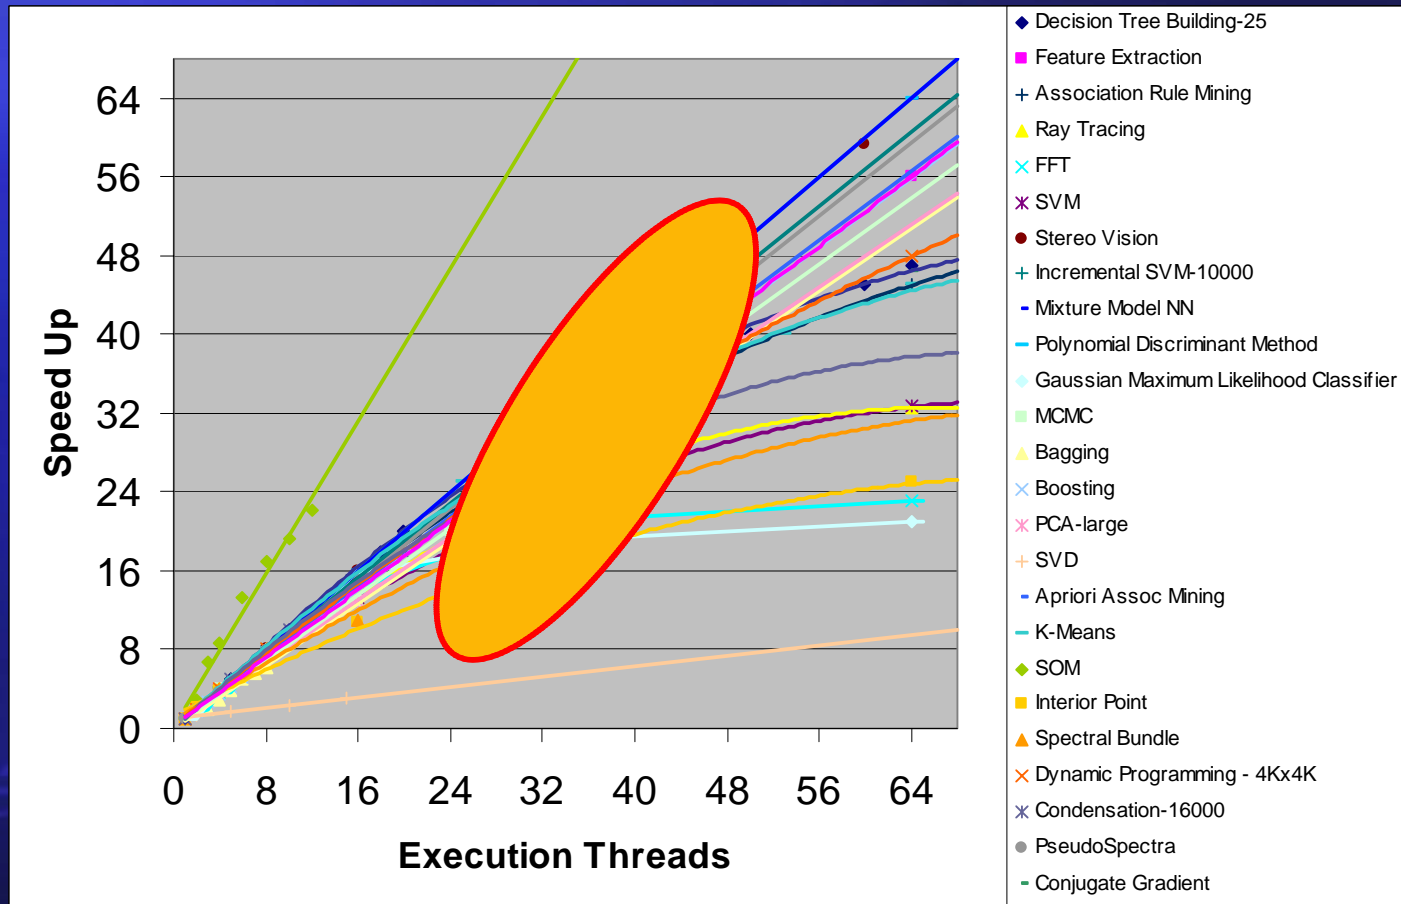


The RMS View



RMS: Emerging general purpose workloads that will drive the need high performance platform architectures

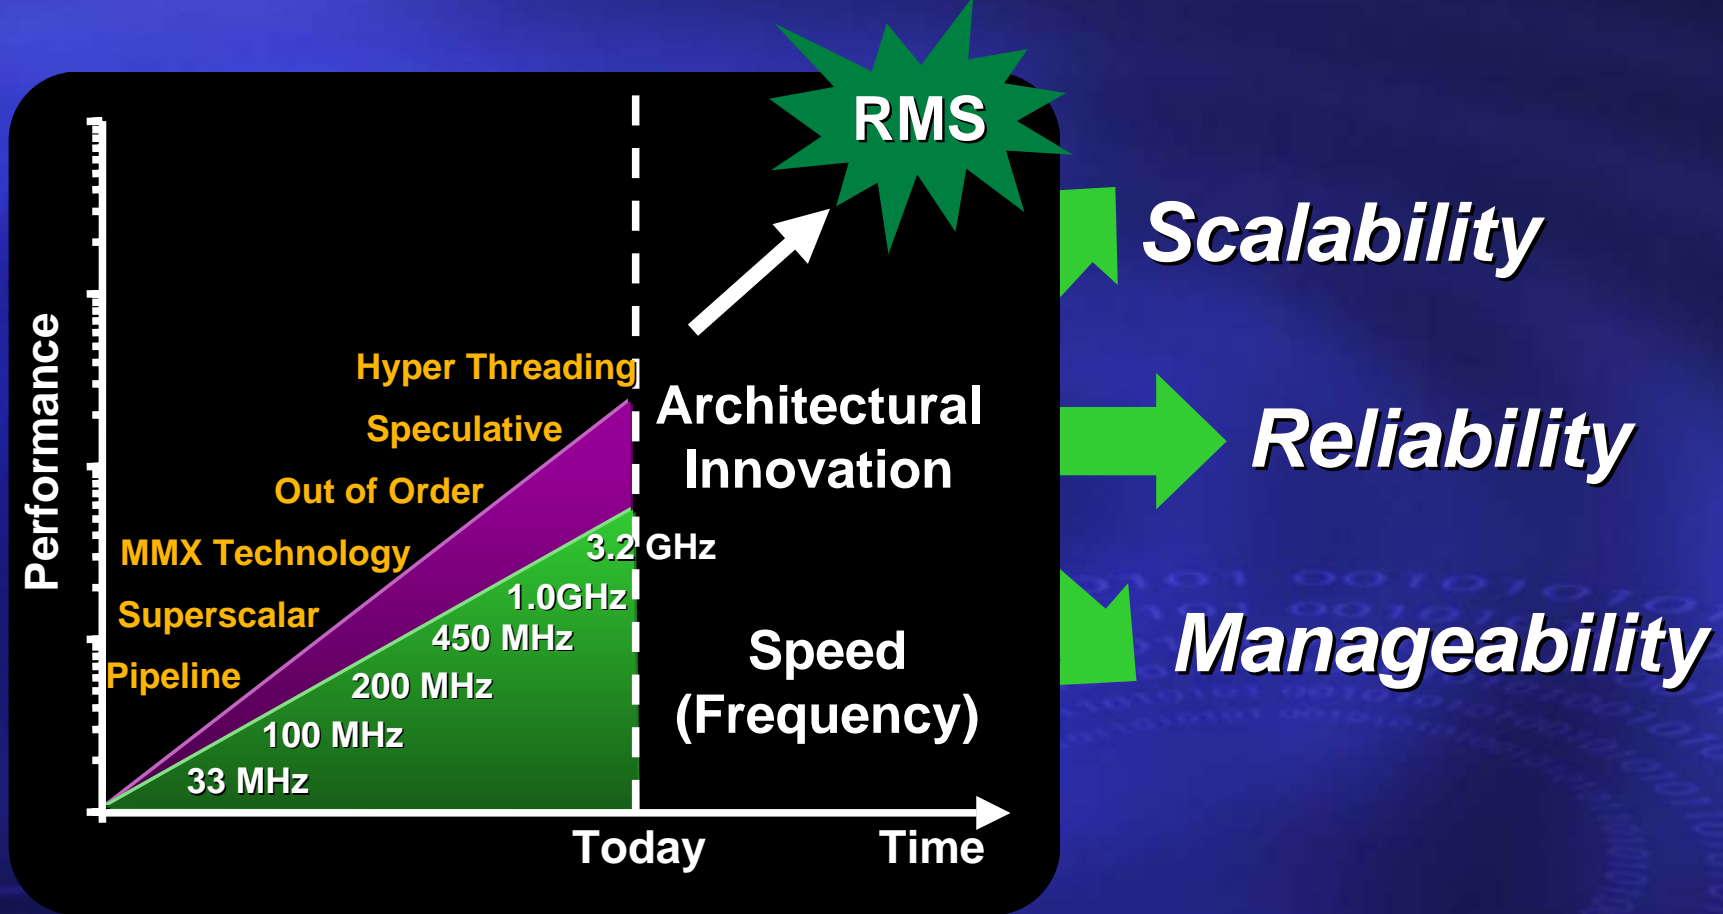
Thread level parallelism in RMS



Curves derived from published sources based on CMP, GPU, SMP, and clusters

RMS workloads should scale well to 100s of execution threads

New Architectural Approach




Value = Performance + Architectural Innovation

Scalability

Scalar plus many core for highly threaded workloads




Large, Scalar cores for high single-thread performance



Many-core array

- CMP with 10s-100s low power cores
- Capable of TFLOPS+
- Full System-on-Chip
- Servers, workstations, embedded...



Multi-core array

- CMP with ~10 cores

Dual core

- Symmetric multithreading

