

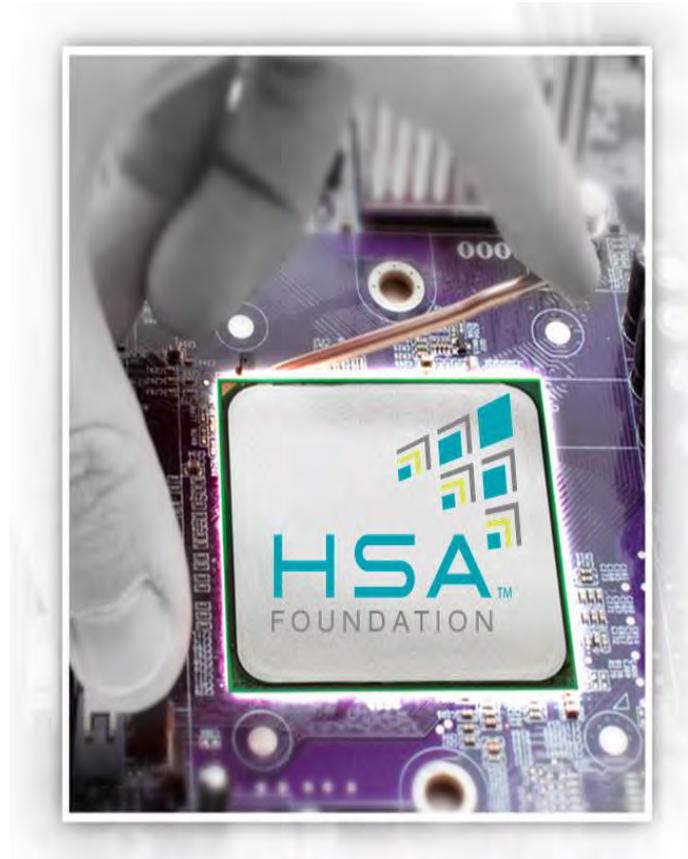
The background of the slide features a high-magnification, colorful image of a microchip die. The die is oriented diagonally, with the top-left corner cut off by a white diagonal line. The chip's surface is a complex grid of various colored regions: a large orange area at the top, a yellow area in the middle, and a green area at the bottom. Numerous small, dark blue and black rectangular structures are scattered across the surface, representing individual components or circuitry. A large, solid purple shape is overlaid on the bottom-left portion of the slide, partially obscuring the chip image. A smaller, solid green arrow-shaped graphic points towards the top-left corner of the chip image.

AMD POWER MANAGEMENT

GREG RODGERS
11/17/2014

APU: ACCELERATED PROCESSING UNIT

- ◆ The APU has arrived and it is a great advance over previous platforms
- ◆ Combines scalar processing on CPU with parallel processing on the GPU and high-bandwidth access to memory
- ◆ Advantages over other forms of compute offload:
 - ◆ Easier to program
 - ◆ Easier to optimize
 - ◆ Easier to load-balance
 - ◆ Higher performance
 - ◆ Lower power



WHAT IS HSA?



Joins CPUs, GPUs, and accelerators into a unified computing framework

Single address space accessible to avoid the overhead of data copying

Use-space queuing to minimize communication overhead

Pre-emptive context switching for better quality of service

Simplified programming

Single, standard computing environments

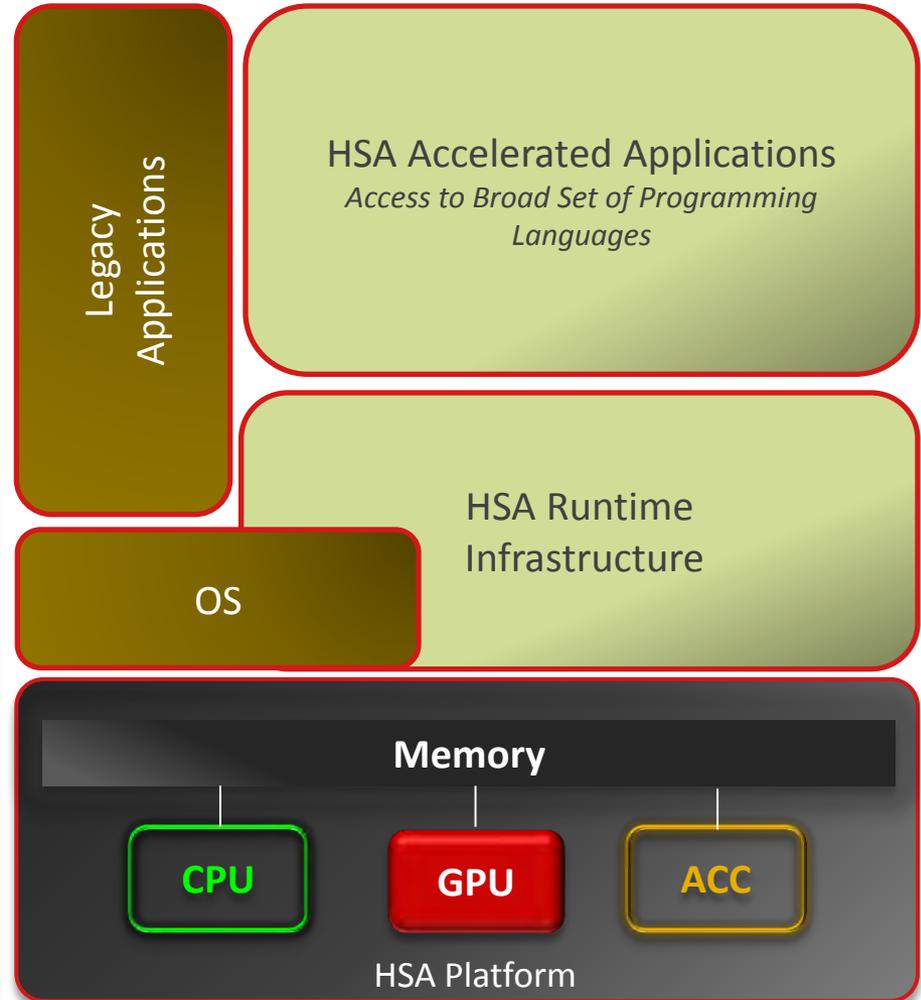
*Support for mainstream languages—
C, C++, Fortran, Java, .NET*

Lower development costs

Optimized compute density

*Radical performance improvement for
HPC, big data, and multimedia
workloads*

Low power to maximize performance per watt



A NEW ERA OF PROCESSOR PERFORMANCE

Single-core Era

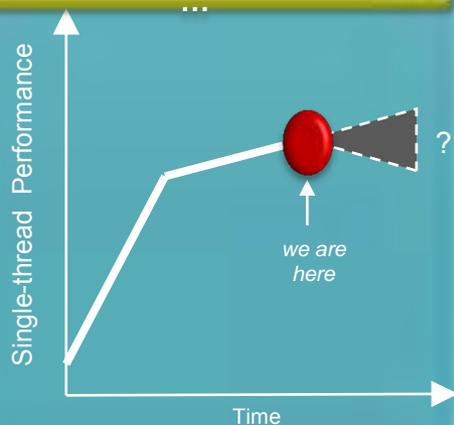
Enabled by:

- ✓ Moore's Law
- ✓ Voltage Scaling

Constrained by:

- ✗ Power
- ✗ Complexity

Assembly → C/C++ → Java



Multi-core Era

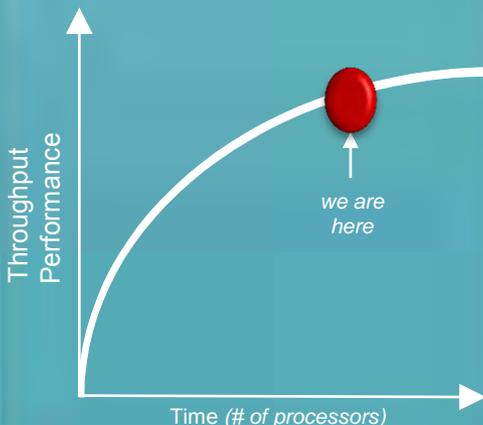
Enabled by:

- ✓ Moore's Law
- ✓ SMP architecture

Constrained by:

- ✗ Power
- ✗ Parallel SW
- ✗ Scalability

pthread → OpenMP / TBB ...



Heterogeneous Systems Era

Enabled by:

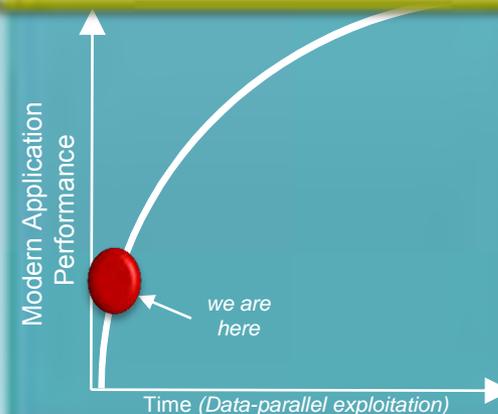
- ✓ Abundant data parallelism
- ✓ Power efficient GPUs

Temporarily

Constrained by:

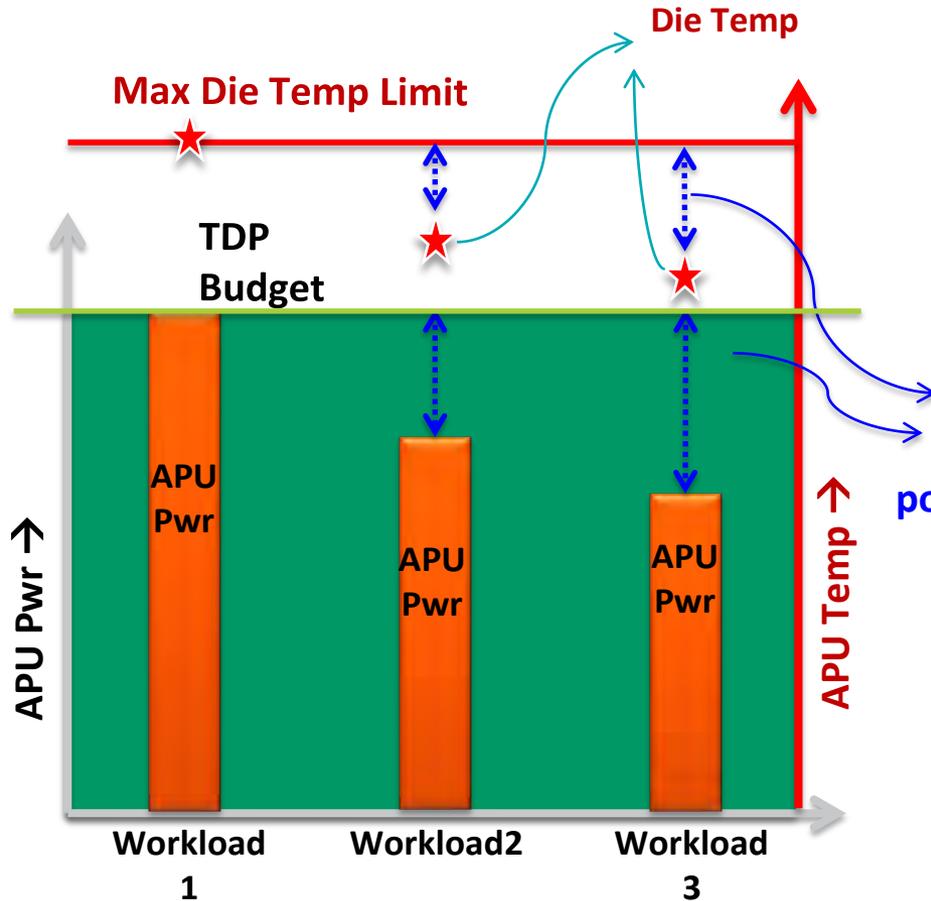
- ✗ Programming models
- ✗ Comm.overhead

Shader → CUDA → OpenCL
OMP4 → C++ → Java



AMD TURBO CORE TECHNOLOGY

MOTIVATION

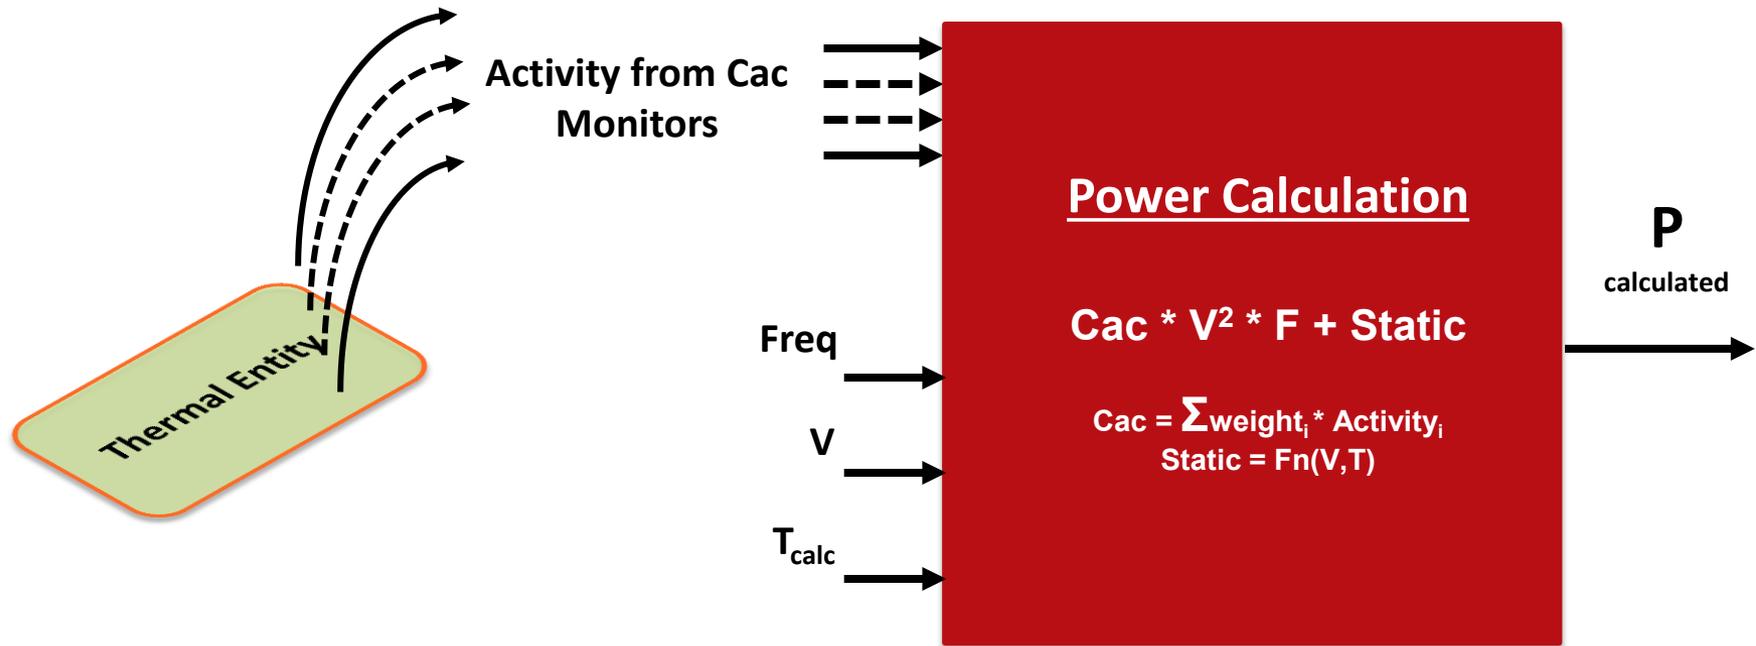


Without AMD Turbo CORE technology,
power and temperature headroom may be left
untapped
in many workload scenarios.

Utilizes estimated and measured metrics to optimize for performance under power and thermal constraints

BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY

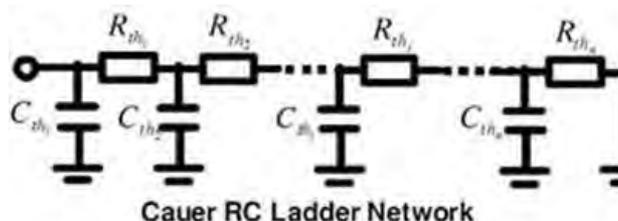
POWER CALCULATION



BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY

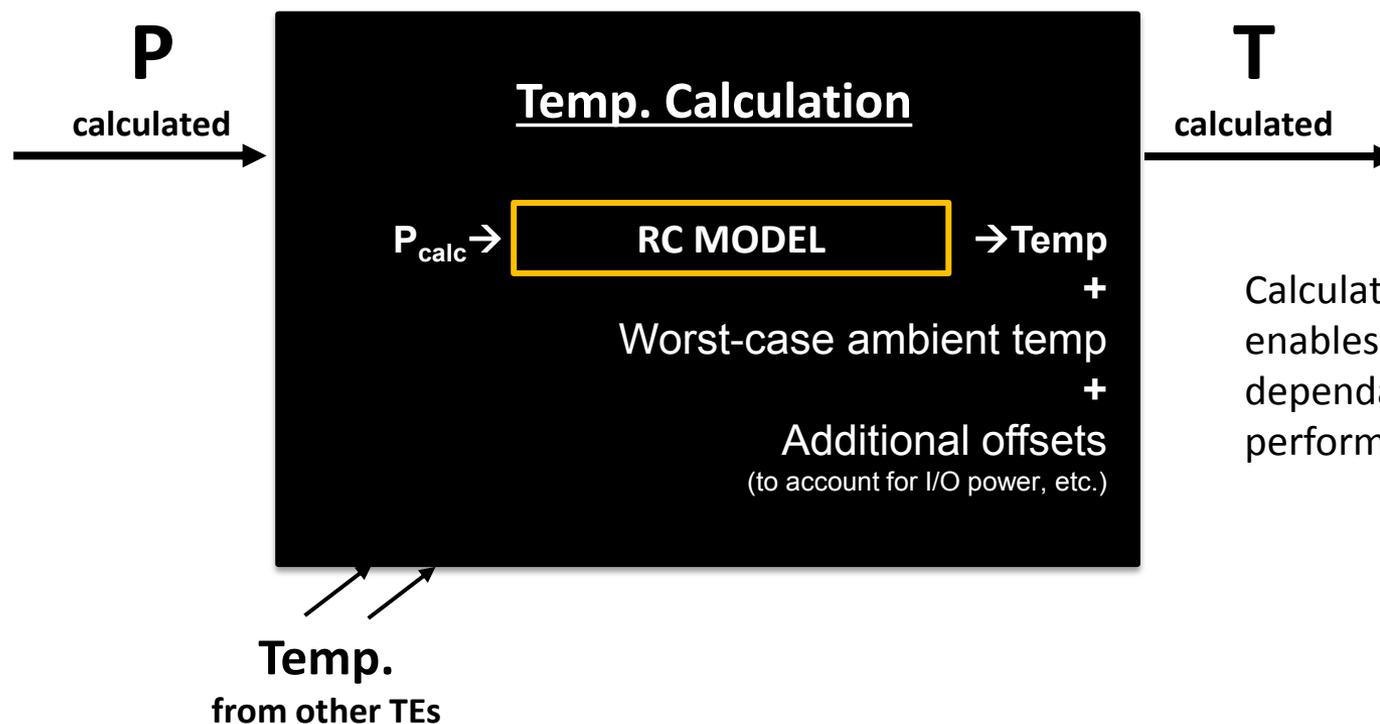


TEMPERATURE CALCULATION



RC MODEL represents

- Cooling solution (APU heat sink, fan, etc.)
- Temperature influence of other TEs

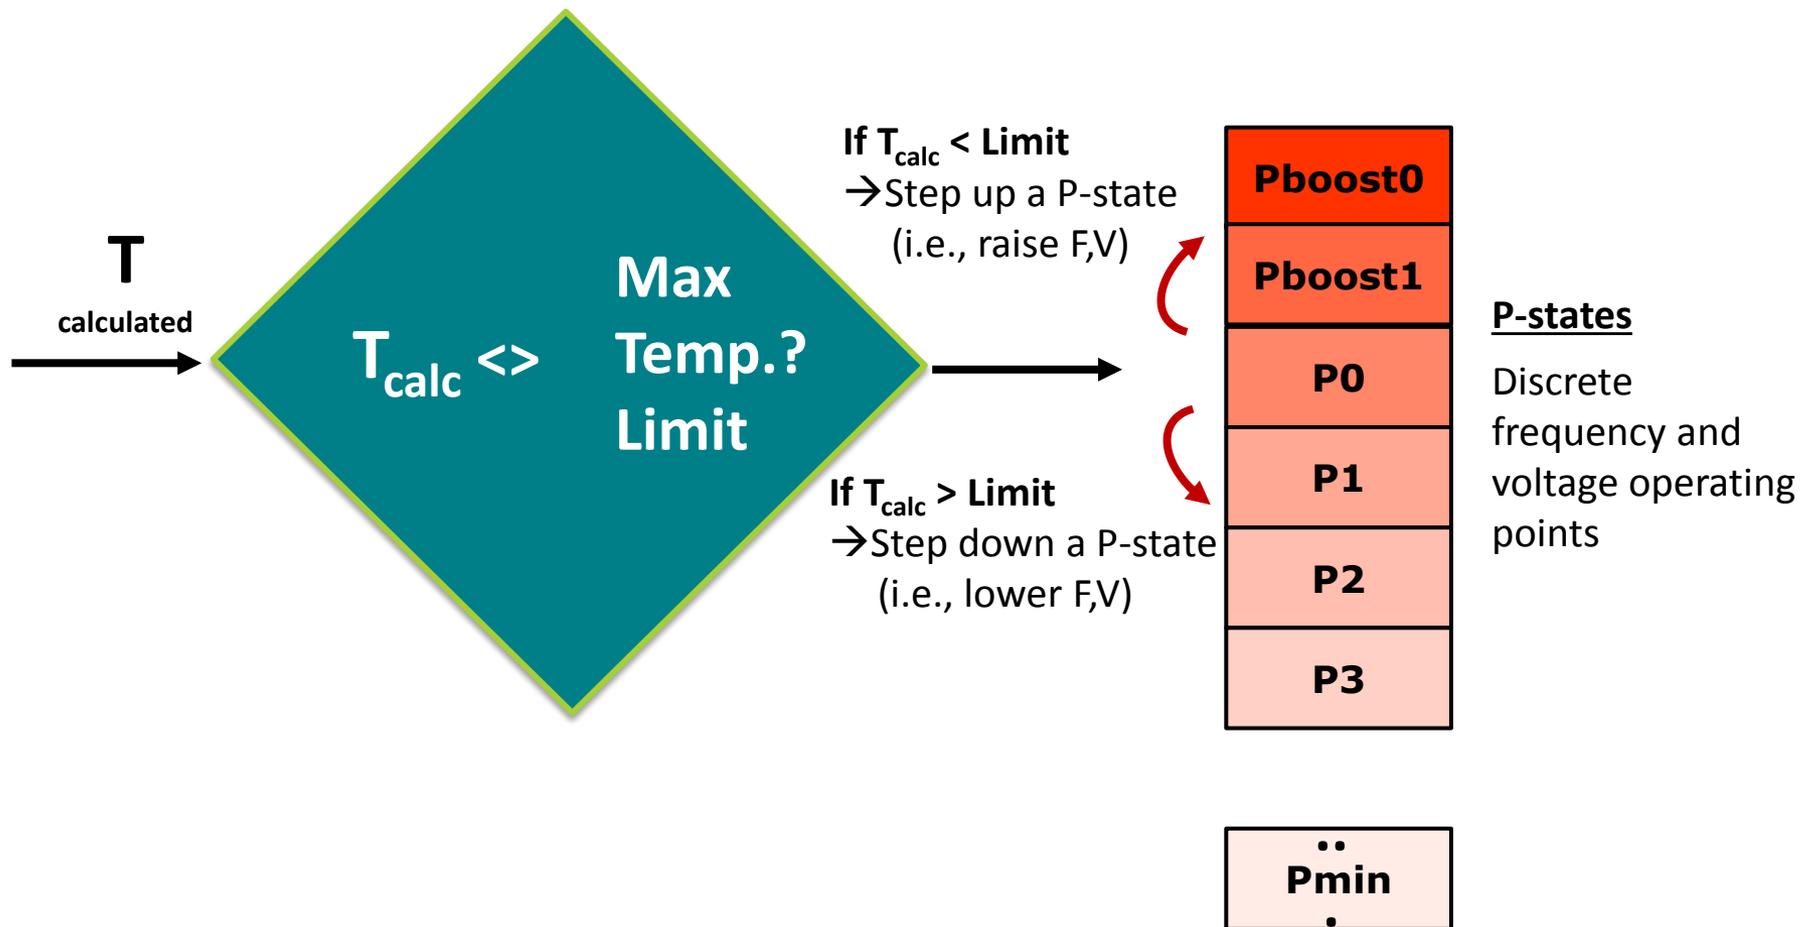


Calculated temperature enables AMD to deliver robust, dependable, and repeatable performance

BUILDING BLOCKS OF AMD TURBO CORE TECHNOLOGY

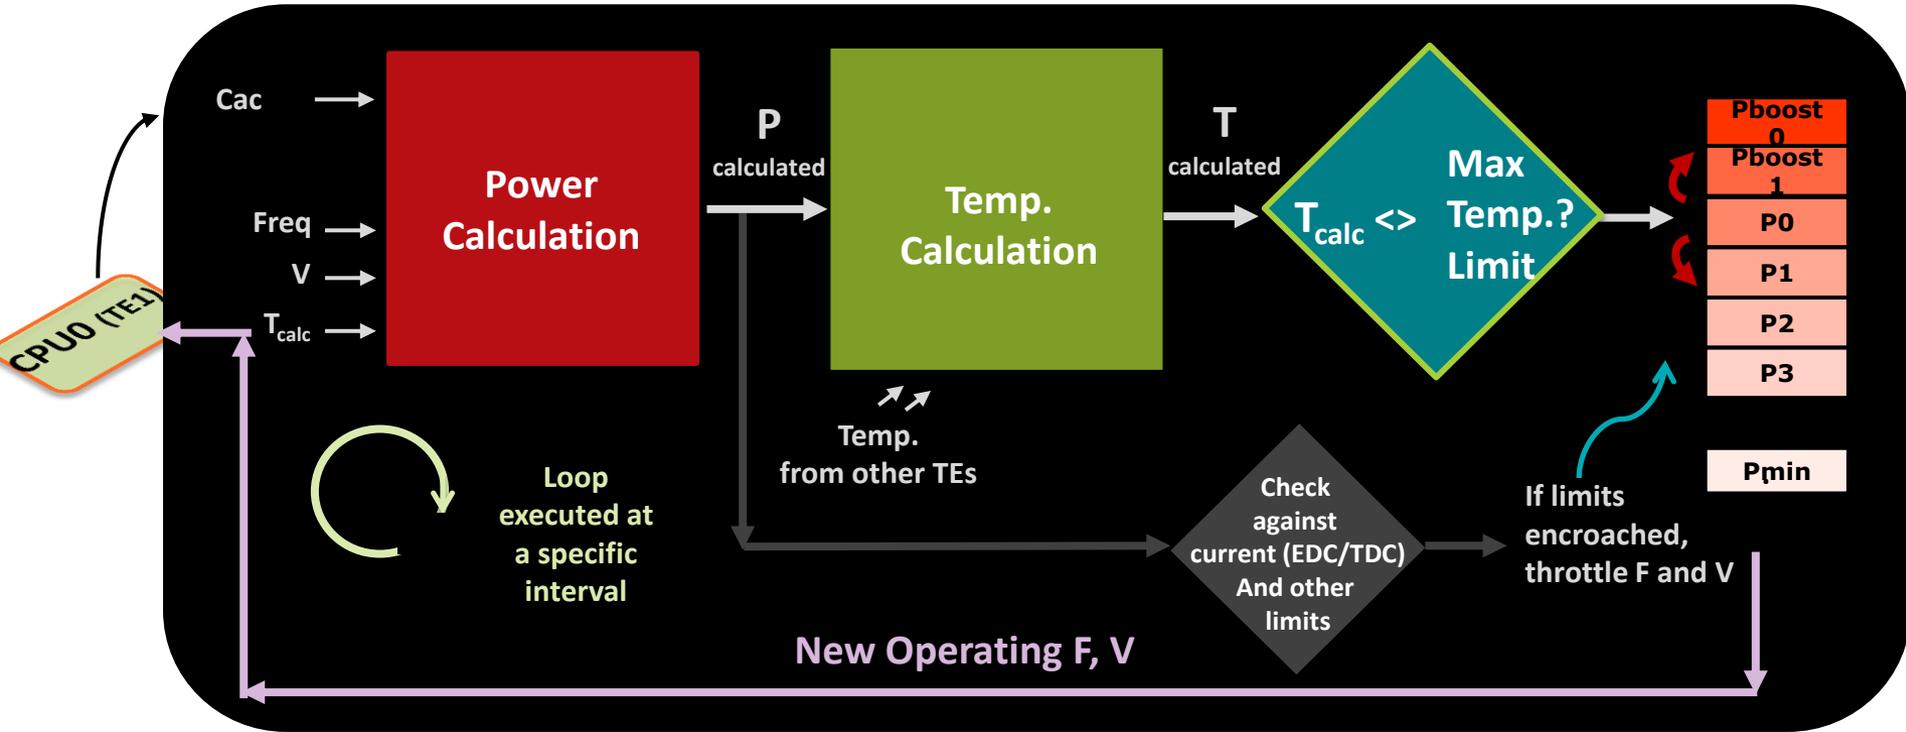


P-STATE SELECTION



AMD TURBO CORE CONTROL LOOPS

PUTTING THE PIECES TOGETHER

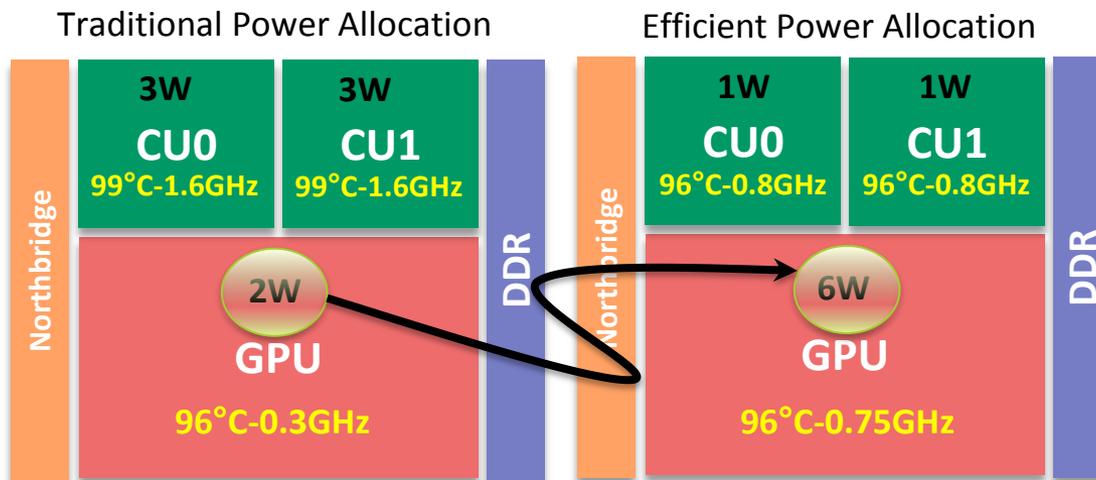


CPU1 (TE2) Control Loop (similar to above)

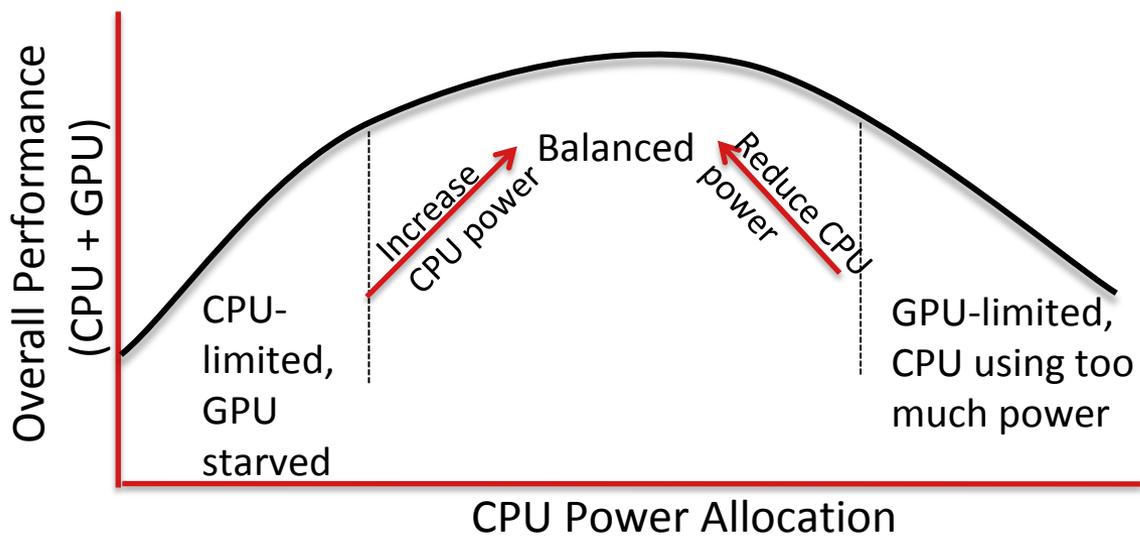
GPU (TE3) Control Loop (similar to above)

INTELLIGENT BOOST

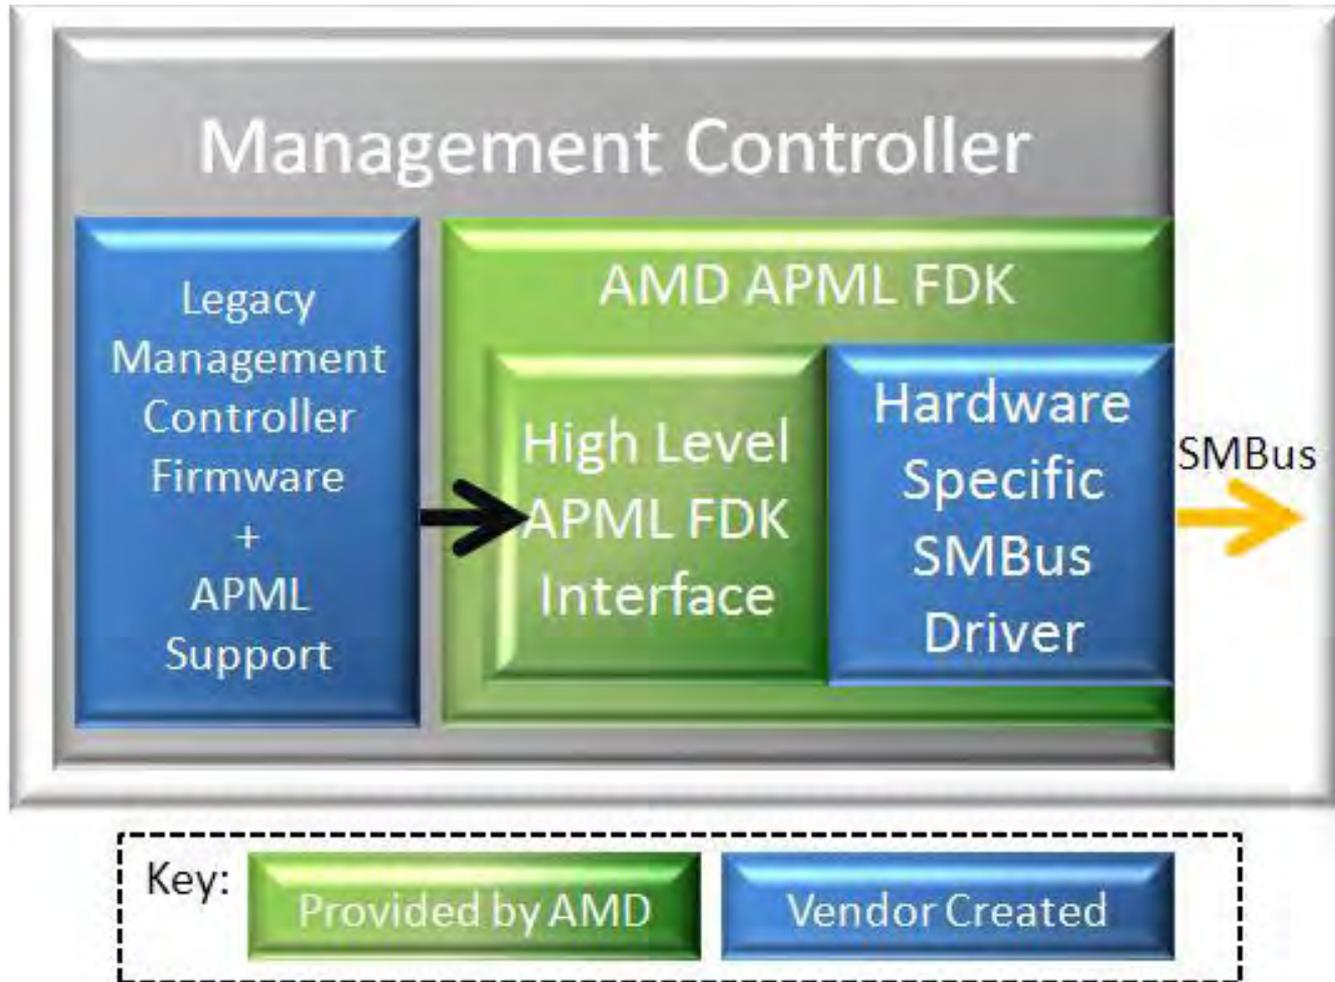
EFFICIENT ALLOCATION OF POWER



- CPU power budget at minimum level to keep GPU fully utilized
- Reduced CPU temperature designed to allow GPU to sustain higher power level
- Total system performance increases



APML: ADVANCED POWER MANAGEMENT LINK



▲ Examples:

- Processor power-capping
- Processor asset identification (including CPUID)
- Machine check register access with optional alerts to management subsystems

▲ Specifics

- 100 kHz standard, and 400 kHz fast mode
- SB-TSI: Sideband temperature sensor interface
 - Access internal temperature sensor and specify temperature thresholds
- SB-RMI: Sideband remote management interface
 - Monitor the current P-state
 - Set and read current maximum P-state

SUMMARY



OPTIMAL UTILIZATION UNDER GIVEN PHYSICAL CONSTRAINTS

▲ HSA framework for effective utilization of available compute resources

- Unified, coherent address space
- Flexible and powerful programming interface
- GPU as first-class compute engine

▲ Power management solutions for effective operation under physical constraints

- Extensive monitoring of internal events and activity
- Intelligent algorithms for translating thermal headroom into performance
- Management of power limits (cTDP)
- Deterministic results using estimated values
- Better margining using measured temperature values

▲ APML

- ▲ Provides a way to characterize AMD Processors and associated BKDG
- ▲ Motivation to expose features via BKDG and higher level APIs.
- ▲ Set requirements for co-design of future processors

Measurements: Nodes

(Info) A node-level measurement shall consist of a combined measurement of all components that make up a node in the architecture. For example, these components may include the CPU, memory, and the network interface. If the node contains other components such as spinning or solid state disks, they shall also be included in this combined measurement. The utility of the node-level measurement is to facilitate measurement of the power or energy profile of a single application. The *node* may be part of the network or storage equipment, such as network switches, disk shelves, and disk controllers.

(important) The ability to measure the current and voltage power and energy of any and all nodes must be provided.

The measurements shall provide a readout capability of:

- (mandatory)** ≥ 100 per second
- (important)** ≥ 1000 per second
- (enhancing)** ≥ 10000 per second

(mandatory) The power and energy data must be real electrical measurements, not based on heuristic models.

DISCLAIMER & ATTRIBUTION



The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

© 2014 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.