



300

White Paper | *ADVANCED POWER MANAGEMENT HELPS
BRING IMPROVED PERFORMANCE TO HIGHLY
INTEGRATED X86 PROCESSORS*






TABLE OF CONTENTS

THE IMPORTANCE OF POWER MANAGEMENT	3
THE X86 EXAMPLES	3
ESTABLISH A REALISTIC WORST-CASE FOR POWER	4
POWER LIMITS CAN TRANSLATE TO PERFORMANCE LIMITS	4
AMD TACKLES THE UNDERUSED TDP HEADROOM ISSUE	5
GOING ABOVE TDP	6
INTELLIGENT BOOST	7
CONFIGURABLE TDP	8
SUMMARY	9

Complex heterogeneous processors have the potential to leave a large amount of performance headroom untapped when workloads don't utilize all cores. Advanced power management techniques for x86 processors are designed to reduce the power of underutilized cores while also allowing for dynamic allocation of the thermal budget between cores for improved performance.

THE IMPORTANCE OF POWER MANAGEMENT

Those with experience implementing microprocessors know the importance of proper power management. Whether for simple applications processors or high-end server processors, the ability to down-clock, clock-gate, power-off, or in some manner disable unused or underused hardware blocks is crucial in limiting power consumption.

Better power management benefits range from energy savings within the data center to improved battery life in mobile devices. But don't underestimate the value of reducing power and increasing efficiency. In fact, power reduction and increased efficiency is even more important today, as processors integrate more and varied functional blocks.

THE X86 EXAMPLE

Typical x86 processors widely used in both consumer and embedded applications are a perfect example: Integration of network and security engines, memory controllers, graphics processing units (GPUs), and video encode/decode engines has effectively turned them into heterogeneous compute units that excel at a wide variety of workloads.

The notable thing about traditional reduction-based power management is that a particular functional block is only turned off when unused, or down-clocked when higher performance is not needed by the application. What about applications that desire more performance? Shouldn't saving power in one area allow you to utilize it in another?

Specifying power usage is complex, particularly with highly integrated processors. If the worst-case power for each individual hardware block in a heterogeneous processor were added together, the resulting total could be several times the achievable worst-case power for the device. The fact that it is nearly impossible to write software that will simultaneously utilize all functional blocks to their fullest extent is one reason. Simply feeding the various compute engines and I/O ports with enough data to keep them all 100% utilized would likely exceed the available bandwidth of internal buses. Central processing unit (CPU) cores manage data movement, and time spent there is less time spent executing higher-power instructions.

Another issue is that different instruction sequences can incur vastly different power usage, which can further complicate specifying processor power. For instance, complex floating-point instructions burn much more power than a simple I/O data read due to the significant difference in transistor logic they activate during execution. The combination of varying instruction types and utilized hardware blocks makes the actual power usage of the processor highly workload-dependent, and explains why it is rare to see a “typical” power specification for this device type. Still, implementers expect a maximum power specification on which to base their design.


ESTABLISH A REALISTIC WORST-CASE FOR POWER

The pragmatic approach for silicon providers is to survey real-world application software to establish a more realistic worst-case power and add some guard-band for safety. Both AMD and Intel use this type of methodology and specify it as thermal design power (TDP). TDP is essentially the maximum sustained power a processor can draw with “real world” software while operating under defined temperature and voltage limits.

POWER LIMITS CAN TRANSLATE TO PERFORMANCE LIMITS

Most embedded x86-based systems are power-constrained in some way. Designers will look for the best performance they can get in a given power envelope, at a price they can afford. The worst-case power limit can translate directly into a performance limit for a given processor product by effectively defining the maximum operating frequency.

Using TDP as a worst-case power specification instead of the cumulative per-block maximum power helps to increase that operating frequency, but it's also based on an assumption of the software workload. Applications using fewer hardware blocks, or using them to a lesser extent, use less power and effectively leave performance headroom on the table.



Integration of large GPU cores, as done in AMD R-Series APUs, increases the potential for unused power budget.

AMD's recent move to integrate discrete-class GPUs with x86 processor cores in accelerated processing units (APUs) underscores this power management challenge. Some APUs contain a GPU that accounts for more than half of the silicon die and a proportional amount of the power budget. A much larger potential for under-utilization of the APU's power envelope exists in this scenario if the software workload is highly CPU-centric or GPU-centric. The trend toward integration of these complex, heterogeneous cores is likely to continue and necessitates a means of harnessing the excess thermal headroom.

AMD TACKLES THE UNDERUSED TDP HEADROOM ISSUE


AMD Turbo CORE technology¹ was launched several years ago to address underutilized TDP headroom. AMD Turbo CORE began with a simple core-counting mechanism that allowed some CPU cores to use higher-frequency "boost" states while other CPU cores were idle. This approach only affected the CPU cores, and was primarily targeted at accelerating single-threaded applications that didn't leverage a multi-core architecture.

Generational improvements have increased the granularity and effectiveness of the technology by adding more boost states for CPU and GPU cores, real-time power and temperature monitors, and enabling dynamic power budget allocation between cores.

Increasing performance by boosting to higher frequencies is relatively simple, since the use of multiple performance states (voltage and frequency operating points) has been around for a while. However, the complexity lies in determining when and which cores to boost. For AMD Embedded R-Series APUs, the process starts by dividing the processor into separate thermal entities: one for each CPU core-pair and one for the GPU. I/O power is small by comparison, so it is defined as a fixed value based on characterization to reduce complexity.

An integrated microcontroller manages AMD Turbo CORE calculations, allowing a more complex and therefore more effective algorithm. In deciding whether boosting a given core is possible, the

power usage of each thermal entity must be determined. On-die analog power measurement at many amps is not practical in a 32nm silicon on insulator (SOI) process, and external measurement is not possible because the various cores share power rails.



Applications with a low C_{AC} can leave unused TDP and temperature headroom. New power management techniques can exploit both for improved performance.

Alternatively, proprietary activity monitors that are integrated throughout the processor architecture model current logic activity as an AC capacitance (C_{AC}). The C_{AC} monitors effectively profile the running application to determine if it is one of those “worst-case” workloads that defines TDP or something less laborious. Static power of the core is determined by transistor leakage at a given voltage and temperature which can be characterized for the device and hard-coded into the algorithm as a function of temperature. A calculated temperature value from a previous iteration is used for reasons that will


be explained later. Total instantaneous power of the thermal entity can then be calculated by $P = C_{AC} * V^2 * f + P_{static}$, and total power for the APU equals the summation of the power for each thermal entity and the I/O power offset. The instantaneous power calculation result is compared to an allocated power budget for the thermal entity, as well as the device’s thermal design current specification to ensure that current demand does not exceed what the voltage regulator can provide. If either value is too close to the limit, firmware can impose throttling by reducing the core’s performance state. The ability to boost the performance state is maintained when headroom exists on both parameters.

GOING ABOVE TDP

Even if an application with a high C_{AC} drives the APU to consume the full TDP, operation at this level may occur in bursts or be preceded by idle time such that the die temperature at the start of the high C_{AC} period is far below the maximum specification. The latest version of AMD Turbo CORE also takes the opportunity to boost in this scenario by allowing brief excursions above TDP when there is adequate temperature headroom. After all, the purpose of a TDP limit is only to ensure die temperature stays in check.

Real-time temperature values from around the thermal entity provide a scaling factor to the power calculation so they influence the boost decision without controlling it directly. Derivation of a calculated temperature comes from application of the calculated power to a reference thermal solution model. Reducing the influence of actual die

temperature on the boost algorithm is an intentional tradeoff to increase deterministic performance of the device. The calculated temperature is then combined with temperature data from other thermal entities to determine if thermal headroom exists.



AMD Turbo CORE algorithms use a variety of frequency, voltage, temperature and logic activity inputs to dynamically determine which cores need a performance boost and how much thermal headroom is available.

Other thermal entities can act as heat sources or sinks, depending on their temperature state, and therefore must be considered. Temperature offsets are also included to account for sensor tolerance and help make sure that the maximum junction temperature is never exceeded. The calculated temperature is compared to predefined thresholds to determine the amount of boost that is possible.

INTELLIGENT BOOST

The final stage of AMD Turbo CORE technology is called Intelligent Boost; it uses a proprietary algorithm that helps improve efficiency by only allowing a core to boost if it can translate that higher frequency to increased performance. If each thermal entity control loop operated independently

under a demanding workload, all cores would attempt to boost until they reach their maximum performance state or until the device thermal limit is reached. It is very unlikely that the application will be perfectly balanced, but rather limited by one core type (CPU or GPU) being saturated.

Intelligent Boost examines the workload at a very high frequency to give more thermal budget to the core that needs it the most by preventing the other cores from boosting more than necessary, maximizing efficiency without affecting overall processor performance.

With an understanding of how boost technologies work, designers should consider where it could affect their application or design practices. A common concern is that designers may

have become accustomed to the idea that the power draw of their software application doesn't come close to driving a processor near its TDP, leading them to design to a lower specification. Historically, this may have been safe, but boost technologies will tend to drive the processor closer to TDP than before by allowing the active cores to consume more power.

Operating closer to TDP may sound like a bad thing, but keep in mind that performance can be gained with the increase in power. An example could be a machine vision application achieving higher frame rates for faster recognition. Total processor power might increase in that scenario, but doesn't materially increase with applications such as media playback in a digital signage player. Fixed, periodic workloads may complete faster at a higher power level, but when the burst of activity is over, cores then spend more time in lower-power idle states, so the average power is approximately the same. Applications like this can still benefit from boosted performance through better responsiveness.

Some designers will dislike the idea of variable performance, especially in real-time applications, but the potential delta is actually quite small in the AMD case. The CAC value of the application has the biggest impact on how much a core will boost, and is very deterministic in behavior on a given processor model. Testing an application on the target processor is a simple way to determine actual performance. The temperature-based boost scalar is the only mechanism that provides variability and it is intentionally limited to minimize


its impact. While users will not see a significant performance delta across the operating temperature of the processor, any variation that does exist can be seen when operating the device near its maximum die temperature.

It is worth noting that the temperature-based boost scalar described only serves to increase boost but does not gradually scale down performance if the maximum die temperature is exceeded. A separate and less granular hardware thermal control mechanism can be used to drop cores to minimum performance states in the event of an over-temperature condition. For applications that are very sensitive to deterministic performance, the boost features can always be disabled.

CONFIGURABLE TDP

Beyond performance benefits, the ability of AMD Turbo CORE algorithms to control average power consumption of the processor also enables a new and interesting feature on the latest generation of AMD APUs, called configurable TDP. It essentially provides the system designer a knob to modify the processor TDP to better fit the needs of the application.

A useful example might be a system design with a thermal budget for a 20W processor but vendor offerings that only include 15W and 25W options. Configurable TDP enables flexibility so the designer isn't forced to choose a lower-performing 15W option in order to remain within the 20W power budget. Instead, the 25W processor might be used but configured for 20W.



SUMMARY

AMD Turbo CORE technology will help to dynamically provide the processor's best available performance while keeping thermal dissipation under the specified amount. Support for configurable TDP and the level of configurability varies by processor model, but it can be a very useful feature for those that support it.

System designers should keep these new concepts in mind when choosing and implementing embedded x86 processors. Power management isn't just about saving power anymore.

1. AMD Turbo CORE technology is available only with select AMD APUs and GPUs.

DISCLAIMER

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of non-infringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

AMD, the AMD Arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. PID 54696-A

© 2014 Advanced Micro Devices, Inc. All rights reserved.