White Paper
Intel® Xeon® Processor
Processor Architecture
Analysis

Measuring Processor Power

TDP vs. ACP

Specifications for the power a microprocessor can consume and dissipate can be complicated and may vary from one manufacture to the other. When comparing power specifications from one manufacture to the other it is important the comparisons are made with specifications which are most similar to each other.

This paper specifically compares how Intel and AMD* specify processor power for their server processors and how you can use those specifications to more accurately compare processors from each manufacturer.

April 2011

Revision: 1.1

Introduction

Intel has traditionally specified processor power with a specification called Thermal Design Power (TDP). AMD also has a TDP specification. Recently, AMD has introduced an additional power value called ACP (Average CPU Power). As of this writing, ACP is only specified on AMD's server processors. This white paper examines the power specifications from both manufactures to help customers understand the best way to compare them and to clarify the differences.

Let's start with a definition of the two key power specifications, and then look at how these specifications are typically used.

TDP (Thermal Design Power)

Intel defines TDP as follows: *The upper point of the thermal profile consists of the Thermal Design Power (TDP) and the associated Tcase value. Thermal Design Power (TDP) should be used for processor thermal solution design targets. TDP is not the maximum power that the processor can dissipate. TDP is measured at maximum TCASE.*[1]. The thermal profile must be adhered to to ensure Intel's reliability requirements are met. Note: Different processors SKU's have different TDP's. At the time of this writing, Intel® Xeon® processors for 2 socket servers (5600 series) are available with a TDP specification from 40W up to 130W depending on the particular SKU[1].

AMD Opteron* processors also have a TDP specification. According to AMD documentation[2] their TDP specification is as follows*:*

*"**TDP**. Thermal Design Power. The thermal design power is the maximum power a processor can draw for a thermally significant period while running commercially useful software. The constraining conditions for TDP are specified in the notes in the thermal and power tables."*
Notes:
- TDP is measured under the conditions of all cores operating at CPU COF, Tcase Max, and VDD at the voltage requested by the processor. TDP includes all power dissipated on-die from VDD, VDDNB, VDDIO, VLDT, VTT and VDDA.
- The processor thermal solution should be designed to accommodate thermal design power (TDP) at Tcase,max.TDP is not the maximum power of the processor.

And from the same document, current AMD Opteron* processors have TDP specifications which range from 35W up to 140W depending on the particular SKU.

A couple of additional comments about TDP specifications that applies to Intel® Xeon® processors, and most should also apply to AMD's Opteron* processors:

- Due to normal manufacturing variations, the exact thermal characteristics of each individual processor are unique. Within the specified parameters of the part, some processors may operate at a slightly higher or lower voltage, some may dissipate slightly higher or lower power and some may draw slightly higher or lower current. As such, no two parts have identical power and thermal characteristics. However the TDP specifications represent a "will not exceed" value.

- Because TDP is a worst case value when running a "worst case" application, most processors, when running a more "typical" workload, will dissipate power that is less than the rated TDP value; how much less will depend on the application and the specific part being tested. This is true not only for the processor, but for every electrical component in a server.

- It is important to note that thermal design power is the maximum thermal power the processor will dissipate, but not the same as the maximum power the processor can consume. It is possible for the processor to consume more than the TDP power for a short period of time that isn't "thermally significant". For example, a processor might consume slightly more power than the rated TDP value for say one microsecond…but then consume less power than the rated TDP value for a long period of time.

Such operation is considered normal. Because the processor temperature does not exceed the specified limits during such a short excursion, the processor will continue to operate correctly.

- It is possible to write "virus-like" code that toggles transistors in the processor on and off, but doesn't do any real work. Such "virus-like" code could cause the processor to exceed the rated TDP value for a much longer, "thermally significant" period of time. To ensure the processor stays within the thermal specification under such "virus-like" type conditions, Intel processors have a built-in "thermal control circuit" which reduces processor power by reducing the processor voltage and/or modulates the clock frequency.

What is AMD's ACP?

According to AMD documentation[3], ACP (Average CPU Power) is the average (Geometric Mean) power a processor was measured to dissipate while running a collection of 5 different benchmarks (Transaction Processing Performance Council (TPC Benchmark*-C), SPECcpu*2006, SPECjbb*2005, and STREAM.) It is also noted in the AMD documentation these measurements to determine the ACP value are not done on every processor, but only on a particular processor that was selected from the "hotter side" of their manufacturing distribution.

Similar to the normal manufacturing variances that may result in a processor consuming less than the rate TDP value, it is expected that each AMD processor will have a different actual ACP value. As of this writing, there are AMD Opteron* processors (6100 and 4100 series) with ACP values that range from 32W up to 105W.

Although AMD's thermal specification document[2] does not mention ACP, based on earlier AMD documentation[4], we believe the tables below accurately reflects the TDP value of a processor with a given ACP value. We believe the table on the left represents previous generation "Istanbul" processors, and the table on the right represents current generation "Magny Cours" processors.

Previous Generation AMD Opteron Processors (Istanbul)[4]	
ACP	**TDP**
105 W	137 W
75 W	115 W
55 W	79 W
40 W	60 W

Table 1: "Istanbul" TDP & ACP Values

Current Generation AMD Opteron Processors (Magny Cours)[2,3]		
ACP	**TDP**	**CPU Series**
105 W	140 W	6100 (G34)
80 W	115 W	6100 (G34)
65 W	85 W	6100 (G34)
75 W	95 W	4100 (C32)
50 W	65 W	4100 (C32)
32 W	35 W	4100 (C32)

Table 2: "Magny Cours" TDP & ACP Values

From AMD's description, we conclude the ACP value is not a guaranteed specification and some parts may have ACP values that are higher or lower than the ACP value quoted by AMD for a given SKU. Also, AMD only quotes an ACP value for the Opteron* class of processors and does not quote ACP values for any of their other processor lines.

What thermal specifications are most useful?

If you are a thermal engineer, the processor TDP specification is very important because your thermal solution (fans, heat sink, etc.) must be able to be able to dissipate the rated TDP value. Intel and AMD both agree on this point. If a thermal engineer designs to ACP, it is likely the design will be undersize and not keep the processor within its thermal specifications.

If you are a server customer, you are probably not concerned about how much power any given component dissipates, but instead the power of the entire server when running your workload. To measure actual server power consumption, you can easily connect a power meter to the input cord(s) of you server while running a variety of workloads.

ACP, while an interesting concept, does not help the thermal engineer design heat sinks; nor is it the best method to determine the power of their server while running a given workload. But you may ask, doesn't ACP represent the power dissipated by a more "typical" workload, instead of a worst case workload? Yes it does, but there are still flaws in using ACP to estimate "real world" processor and server power:

1. Since ACP is not measured on every processor, the actual power a given AMD processor will dissipate may be higher or lower than the ACP value. TDP values, on the other hand, are a specified value and are measured on every device that ships from the factory.
2. The workloads used to generate the ACP vale may not match the workload(s) of your specific use.
3. The processor isn't the only component in the server. Even if the ACP value is "close enough" to estimate the power of the processor, no other component in the server has an equivalent ACP value. And even if other components did have "ACP like" values it would be tedious to collect those values and add them up to estimate the system power under "real world" conditions.
4. It is very easy to measure the actual server power while running your specific workload. Just plug your server into a power meter while running your specific application(s). So there is no need to use a processor ACP value. You can simply measure system power which is more accurate and less effort than item #3 above.

ACP does not equal TDP

A common misuse of specs is to assume that TDP and ACP are similar specifications. As shown above, TPD and ACP have very different definitions and different uses, and cannot be compared.

As an example, below are the SPECpower_ssj2008* results of two similar 2-socket servers from Hewlett Packard, one with Intel processors and the other with AMD processors. Note the ACP value for the AMD processors is 80W, while the TDP specification for the Intel processors is 95W. The discrepancy in ACP vs. TDP power values might suggest the AMD platform should dissipate ~30W lower system power. Yet according to the SPECpower results, the Intel platform is actually 49W lower in overall system power at 100% target load. Thus the lower ACP value does not translate to a lower system power.

In fact, if you look at the difference in processor TDP between these two platforms (180W for the Intel-based platform and 230W for the AMD-based platform) the 50W difference in TDP values for these platforms is nearly identical to the 49W system power difference measured at 100% target load. This would suggest comparing processor TDP's is more accurate at predicting relative platform power.

	Vendor/Server	Processor	Processor TDP	Processor ACP
Server A	2 socket HP* ProLiant* DL380 G7	2 x Intel® Xeon® processor X5675 3.07GHz	2 x 95W = 180W	N/A
Server B	2 socket HP ProLiant DL385 G7	2 x AMD Opteron processor 6174 2.2GHz	2 x 115W = 230W	2 x 80W = 160W

Table 3: System Configuration Comparison for SPECpower_ssj2008

	Performance[1] (higher is better)	System Power[1] (Lower is better)	Performance/Watt[2] (Higher is better)
Server A	894,314	222W	4020
Server B	888,819	271W	3281

Table 4: SPECpower_ssj2008* Results

1. Performance and power measured in ssj_ops and watts @ 100% target load while running SPECpower_ssj2008
2. Performance to Power Ratio at 100% target load per SPECpower_ssj2008 results

Server A used 222W@100% target load compared to Server B's 271W@100% target load. Server A has a SPECpower_ssj2008 result of 3,197 overall ssj_ops/Watt and Server B has a SPECpower_ssj2008 result of 2,355 overall ssj_ops/Watt.

Data Source for Server A: (as of March 30, 2011)
http://www.spec.org/power_ssj2008/results/res2011q1/power_ssj2008-20110209-00353.html
Data Source for Server B: (as of March 30, 2011)
http://www.spec.org/power_ssj2008/results/res2010q2/power_ssj2008-20100323-00242.html

Bottom line: TDP and ACP are not equal. If you want to compare "apples to apples" it is recommend you compare Intel's TDP specification with AMD's TDP specification. AMD's ACP value has no equal from Intel. In addition, it is expected that both Intel and AMD's processor will typically dissipate less than the rated TDP value when running real world applications.

Common misuses of processor power specifications

Max Voltage * Max Current = Max Power

As noted above, the power, thermal and current specifications are used by specific engineers for specific purposes. A common misuse of the specifications is to multiply the maximum current specification, by the maximum voltage specification and determine the maximum power the processor will dissipate. While the below equation is true:

Voltage x Current = Power

Such a calculation does not mean a processor can or will dissipate this much power. The actual voltage and current values of a processor are managed together to ensure the processor power remains below the rated TDP value. If a given processor is running near its maximum current value, for example, the voltage for this processor is likely well below the maximum voltage value.

Lower Power is Better

There are many press articles which assume a lower power server is also more energy efficient. This is actually far from the truth. Power by itself is not a measurement of overall server efficiency. Performance of the server, in conjunction with the power consumed is what defines energy efficiency. A system which is lower power, but is also lower performance will take longer to perform a task, and

may ultimately consume more energy. The most efficient server is one which has the best performance per watt. As an example, consider these two servers with the following performance and power specifications:

	Vendor/Server	Processor	Processor TDP
Server A	2 socket HP ProLiant DL380 G7	2 x Intel® Xeon® processor X5675 3.07GHz	2 x 95W = 180W
Server B	2 socket ZT Systems 1253Ra Datacener Server*	2 x AMD Opteron* processor 4164EE 1.8GHz	2 x 35W = 70W

Table 5: System Configuration Comparison for SPECpower_ssj2008

At first glance you might think Server B will be more efficient because it has lower total processor TDP (70W vs. 180W). However, when you look at the overall server performance per watt, you see a different story. Server A has 218% higher performance, yet server power is only 176% greater than Server B. Even though Server A does consume more power, it is 23.6% more energy efficient as shown by its better performance/watt.

	Performance[1] (higher is better)	System Power[1] (Lower is better)	Performance/Watt[2] (Higher is better)
System A	894,314	222W	4020
System B	409,496	126W	3253

Table 6: SPECpower_ssj2008* Results

1. Performance and power measured in ssj_ops and watts @ 100% target load while running SPECpower_ssj2008
2. Performance to Power Ratio at 100% target load per SPECpower_ssj2008 results

Server A used 222W@100% target load compared to Server B's 126W@100% target load. Server A has a SPECpower_ssj2008 result of 3,197 overall ssj_ops/Watt and Server B has a result of 2,106 overall ssj_ops/Watt.

Data Source for Server A: (as of March 30, 2011)
http://www.spec.org/power_ssj2008/results/res2011q1/power_ssj2008-20110209-00353.html
Data Source for Server B: (as of March 30, 2011)
http://www.spec.org/power_ssj2008/results/res2010q2/power_ssj2008-20100601-00265.html

How should an end user measure/estimate "real world" server power?

It is reasonable end users would want to know the power their servers are expected to consume when running "real world" workloads. Estimating server power based on a worst case TDP specifications will result in over estimating server power. Intel and AMD both agree on this point: *"It is of little value to measure power consumption by only looking at the spec sheets for different components, adding the totals together, because these generally only report the maximum power consumption.[3]"* As such, an ACP like value seems reasonable at first glance. However ACP only gives you the power of the processors when running "real world" applications. It doesn't help estimating the power dissipated by the other components in the server such as memory, hard drives, I/O boards, disk controllers etc.

All is not lost, however, because there is a very accurate way to measure server power under a users "real world" conditions. And again, Intel and AMD both agree on this point: *"The best way to measure a server's power consumption is the power meter, an inexpensive tool that is plugged into the wall, and then your device, like a server, can be plugged into the power meter. The meter displays the wattage drawn "at the wall" and allows you to analyze the power consumption under a variety of different utilization levels."*[3] So if an accurate "real world" power value is needed, simply measure it with a power meter. Because of normal component power consumption tolerances, it is recommended that more

than one sever is measured. Power can vary slightly even between identically configured servers from the same vendor.

Why doesn't Intel provide an ACP value for their processors?

Intel sees no value adding another specification to our processors. As noted above, ACP is not useful for system or processor thermal engineers and end users can get more accurate power consumption values by simply measuring the actual power of their server while running their specific application. In addition, AMD does not specify exactly where in their silicon process distribution they measure ACP, so it would be impossible for Intel to create an identical ACP specification.

Conclusion

While you can come to your own conclusions on the value and purpose of ACP, we want to convey the following:

- ACP does not equal TDP. These specifications are quite different and are not comparable.
- When comparing processor power specifications, we believe comparing TDP to TDP is the most "apples-to-apples' comparison that can be made.
- If you want to determine the power of your server while running a "real world" workload, plug the best and most accurate is to plug your server into a power meter and measure system power at the wall.

Sources:

1. Intel® Xeon® Processor 5600 Series Datasheet Volume 1 http://www.intel.com/Assets/en_US/PDF/datasheet/323369.pdf (March 2010)
2. AMD Family 10h Server and Workstation Processor Power and Thermal Data Sheet. http://support.amd.com/us/Processor_TechDocs/43374.pdf (June 2010)
3. http://www.amd.com/us/Documents/43761D-ACP_PowerConsumption.pdf (2010)
4. http://www.amd.com/us/Documents/43761C_ACP_WP_EE.pdf

Author

Scott Huck is a Performance Architect in Intel's Data Center Group

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of servers available for purchase.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER. INCLUDING WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

SPEC and the benchmark name SPECpower_ssj are trademarks of the Standard Performance Evaluation Corporation. Benchmark results stated above reflect results published on http://www.spec.org as of November 22, 2010. For the latest SPECpower_ssj2008 benchmark results, visit http://www.spec.org/power_ssj2008/results/power_ssj2008.html.

Intel, the Intel logo, Intel®, Xeon®, and the Xeon® logos are trademarks or registered trademarks of Intel Corporation or its subsidiaries in other countries. * Other names and brands may be claimed as the property of others. All timeframes, dates and products are subject to change without further notification.