

WHY PROCESSOR PERFORMANCE IS MORE THAN FREQUENCY AND CORE COUNT

TODAY'S EVOLVING PROCESSOR ARCHITECTURES REQUIRE NEW SPECIFICATIONS AND PERFORMANCE BENCHMARKS

EXECUTIVE SUMMARY

The IT industry is undergoing significant workload shifts that challenge the status quo. Previously, the industry generally accepted that improving central processing unit (CPU) core count, clock speed, and instructions per cycle (IPC) would yield the best performance and power consumption results. In the last few years, the world's leading processor manufacturers introduced performance and efficient processor cores that challenge this status quo.

Processors are becoming more complex, with multiple types of cores with variable frequencies – and variable power and thermals to match those frequencies - designed into a single chip. Traditional methods of assuming performance based on processor base frequencies and thermal design power (TDP) are no longer accurate.

When deciding which processor is most suitable for their workloads, we at Moor Insights & Strategy encourage organizations to look at benchmarks and actual usages, along with application performance, rather than base frequencies and TDP. Furthermore, we encourage PC OEMs and processor vendors to change how they communicate processor specifications to better translate real-world performance through their specification sheets and channel marketing.

Intel's new 12th Gen Core Alder Lake processors adopt a more comprehensive approach to communicating processor performance through core counts and max frequency clock speeds for both types of cores inside of the processor and a range of operating power levels. IT buyers should pay attention to the changing industry landscape – which is moving toward mixed-core configurations and dynamic frequency and power levels – to find the right products for their organizations.

SHIFTING INDUSTRY TRENDS

The IT industry is entering a new era: multicore CPUs utilizing different architectures that deliver adaptive performance and power profiles.

This trend began with performance and efficient cores in the mobile phone world and has now transitioned to the PC with Intel's performance hybrid technology. Another new trend in processor capabilities is the dynamic frequency range of the modern system on a chip (SoC), which is more dependent on - and aware of - the workload and thermal conditions of the system and chip.

This trend of increased complexity is not isolated to CPUs; graphics processing units (GPUs) now feature many types of cores that perform various specialized tasks. Applications may use artificial intelligence (AI) instructions on the GPU in conjunction with the CPU to accelerate performance further. Today, some GPUs add matrix multiplication and ray-tracing cores with different frequencies than the familiar GPU shader cores.

Performance CPU cores differ from efficient cores in power and clock speed characteristics. These new processor designs have the flexibility to shift tasks between the different types of cores to maximize the balance between performance and power consumption.

Intel's 12th Gen performance hybrid technology uses multiple types of cores in mobile devices and PCs with differing approaches. Intel's 12th Gen Core, Alder Lake, incorporates these processors across the product line for the first time. Intel's first hybrid processor, Lakefield, was a limited-release product, and Alder Lake represents many of Intel's learnings from that first approach.

Balancing workloads between these different cores and dynamically clocking these diverse core architectures optimizes total processor performance and power consumption. Operating systems like Windows 10 and 11 have schedulers that dictate where workloads go on the processor, taking advantage of and recognizing the different cores to maximize performance and power savings. This makes assessing processor performance and power purely on the number of cores or clock speeds increasingly inaccurate. The combined complexity makes it much harder to determine the best suited SoC without benchmarking performance in the real world with real workloads.

These new paradigms in PC performance and power operations render the traditional way of defining processor specs outdated. Core counts must now consider performance and efficient cores (where applicable). Clock speed must now consider how long a processor stays at its peak clock and the average speed, rather than only its base clock.

AI performance is increasingly important across the IT industry and cannot be simply quantified or compared with core counts or clock speeds. Architectures can optimize for different AI workloads, so comparing AI performance across processors becomes even more challenging without running industry-accepted benchmarks like MLperf.

While most AI workloads on PCs are likely to be inference workloads, not all AI algorithms perform the same on all hardware. Since AI algorithms perform different tasks – like object detection, natural language processing, denoising, and super sampling – they all perform differently on different architectures and memory configurations. This disparity in AI performance doesn’t translate well across heterogenous core types like Intel’s performance and efficient cores.

TABLE 1: INTEL’S UPDATED TERMINOLOGY

| Old Specification Terminology (11th Gen or prior) | Updated Specification Terminology (12 th Gen) | Recommended New Specification Replacement |
|--|---|--|
| TDP | Processor Base Power | Processor Base Power |
| CTDP Down Power | Minimum Assured Power | Minimum Assured Power |
| CTDP Down Frequency | N/A | N/A |
| CTDP Up Frequency | N/A | Performance-core Max Turbo Frequency Efficient-core Max Turbo Frequency |
| CTDP Up Power | N/A | Maximum Turbo Power |
| Base Frequency | N/A | Performance-core Max Turbo Frequency Efficient-core Max Turbo Frequency |
| Processor Cores | Performance-cores Efficient-cores | Total Cores (P-cores plus E-cores) |

Source: Intel

These differences make it incredibly difficult to compare processor vendors to one another, primarily when one utilizes a heterogenous multi-architecture multicore design and the other does not.

INTEL'S ALDER LAKE 12TH GEN CORE FOR MOBILE

With the introduction of Intel's new 12th Gen Alder Lake processors, the company has changed how it and its partners communicate processor specifications to better demonstrate performance expectations. Intel's ARK page for specifications has adopted a more nuanced approach to cores, clock speed, and power.¹ This also comes with a need to update terminologies to match the changes in types of CPUs and power.

Table 2 below demonstrates Intel's current methodology with the new Intel Core i9-12900H processor compared to the old methodology with the last generation's Intel Core i9-11900H processor.

TABLE 2: INTEL'S CPU SPECIFICATION SIDE-BY-SIDE COMPARISON

| ARK Compare Intel® Products | | |
|---|----------------------------------|----------------------------------|
| | Intel® Core™ i9-12900H Processor | Intel® Core™ i9-11900H Processor |
| CPU Specifications | | |
| Total Cores | 14 | 8 |
| # of Performance-cores | 6 | |
| # of Efficient-cores | 8 | |
| Total Threads | 20 | 16 |
| Intel® Turbo Boost Max Technology 3.0 Frequency ‡ | 5.00 GHz | 4.90 GHz |
| Performance-core Max Turbo Frequency | 5.00 GHz | |
| Efficient-core Max Turbo Frequency | 3.80 GHz | |
| Cache | <u>24MB</u> Intel® Smart Cache | <u>24MB</u> Intel® Smart Cache |
| Processor Base Power | 45 W | |
| Maximum Turbo Power | 115 W | |
| Minimum Assured Power | 35 W | |
| Max Turbo Frequency | | 4.90 GHz |
| Configurable TDP-up Frequency | | 2.50 GHz |
| Configurable TDP-up | | 45 W |
| Configurable TDP-down Frequency | | 2.10 GHz |
| Configurable TDP-down | | 35 W |

Source: Intel

¹ Intel ARK – Product Specifications

Let's look at some of Intel's new specifications (as illustrated above) more closely.

- **Total Cores:** Intel's CPU specifications now include a new total core count, which isn't an entirely new specification but incorporates both performance and efficient core counts.
- **Total Threads:** Thread count, which also isn't a new specification, requires updating because it now includes two different types of cores that are threaded differently, with 12th Gen P-cores having two threads per core and E-cores having one thread per core. Therefore, it's crucial to understand the difference between old and new thread counting.
- **Max Turbo Frequency:** Similarly, Intel has added Max Turbo frequency specifications for both the performance and efficient cores; this only required one number for one type of core previously.
- **Cache:** This represents the total of L2 and L3 cache, which is a methodology comparable to what the industry counts when it compares CPU cache. Intel calls this Smart Cache, terminology that has existed for over a decade.
- **Processor Base Power:** Intel's replacement for TDP, processor base power more accurately represents the long-term sustained power consumption of the processor at its default power management configuration.
- **Max Turbo Power (MTP):** MTP is also a more representative measure of the processor's power state when it clocks itself to the absolute maximum, if even for a millisecond. It may not represent real-world power consumption during day-to-day use.
- **Minimum Assured Power:** Minimum assured power is a potential state the processor can be powered down to but could represent the processor's guaranteed base performance level (PL1).

Intel's approach to frequency specifications aligns with the changes happening in the industry, and better accounts for multicore CPUs with heterogeneous architectures with different performance and power profiles. Its new approach to power specifications also better matches how processors behave with dynamic power, based on the dynamic nature of the processors' clock speeds based on workload. Intel's new mobile CPU specifications show today's complex market and the nuance necessary when listing specifications.

While these new architectures may introduce new challenges to the IT hardware space, they also emphasize benchmarks and testing real-world workloads to test performance more critically.

CALL TO ACTION

Processors are becoming increasingly more complex, and with that added complexity comes more nuance in the types of cores that reside within them. Traditional methods of evaluating processors based purely on core counts, clock speeds, and TDPs are obsolete and require more context, including pertinent benchmarks.

Intel's new approach to core count, clock speeds, and power does a much better job of reflecting real-world performance. It should help set buyers' expectations for mobile platforms' total performance and battery life. Good benchmarks based on actual user workflows and applications are still the most crucial metric for fully understanding actual head-to-head performance.² With these learnings and changes in definitions, we hope the industry can work together to enable easier decision-making that ultimately delivers better end-user experiences.

For more information on mobile and desktop platform benchmarks, go [here](#).

² [What to Consider When Benchmarking Workstations – Moor Insights & Strategy - 2021](#)

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

[Anshel Sag](#), Vice President and Principal Analyst at [Moor Insights & Strategy](#)

PUBLISHER

[Patrick Moorhead](#), CEO, Founder & Chief Analyst at [Moor Insights & Strategy](#)

INQUIRIES

[Contact us](#) if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

DISCLOSURES

Intel commissioned this paper. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2022 Moor Insights & Strategy. Company and product names are used for informational purposes only and may be trademarks of their respective owners.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.