CURTISS-WRIGHT

Cores and Threads: Hybrid Processors for Today's Multitasking World Part-2

Introduction

In Part 1 of this White Paper series, we learned about the evolution of modern processors. We presented the inner workings of single-core, multi-core, and hyper-threading processors. We also introduced the concept of a processor with different classes of cores, with Arm® processors adopting the big.LITTLE architecture and Intel® adopting a similar architecture called hybrid core processors which incorporates Performance cores (P-cores) and Efficient cores (E-cores). Finally, we explored how today's operating systems (OS) use hybrid core processors, and algorithmically select which type of tasks to assign to the different core type.

In Part 2 of this series, we present performance testing results on Intel hybrid core processors, exploring the performance and efficiency of P-cores vs. E-cores and single-threaded cores vs. hyper-threading cores. We also explore the incredible performance boost provided by using processor Turbo modes and we highlight the extreme power/efficiency penalties their use incurs. Finally, we'll consider how these technologies can benefit the system and software developer of embedded systems and summarize the specific core and thread configurations of today's popular embedded processors.

Contents

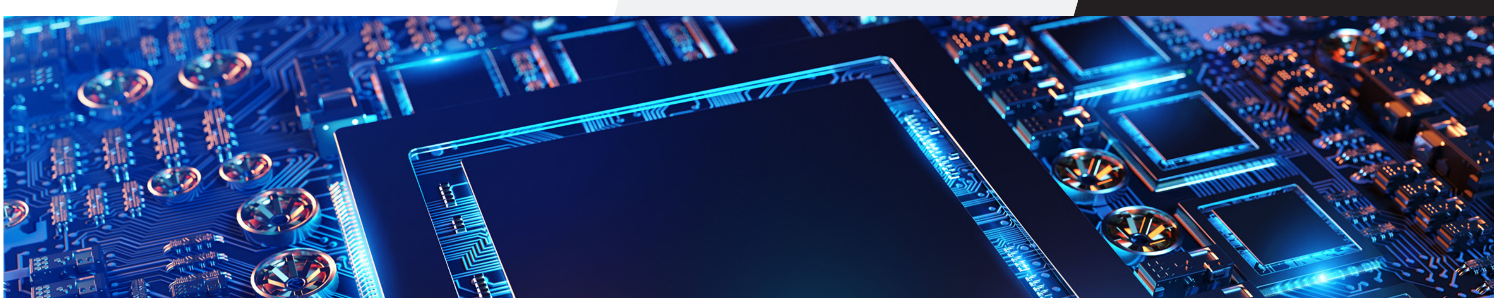

CURTISS-WRIGHT

Performance Exploration of the Intel Hybrid Processor

To evaluate the comparative performance of P-cores and E-cores, benchmarking tests were conducted on an Intel Alder Lake NUC computer[1] built around the 12th Gen i7-12700H processor featuring 6x hyper-threading P-cores and 8x E-cores (6P + 8E) for a total of 20 threads. In addition to the Intel Alder Lake NUC computer, we used a Dell Precision 7670[2] laptop built around the 12th Gen i7-12850HX processor (8P + 8E = 24 threads). While our testing did not include any Intel 13th Gen Raptor Lake processors, we can expect the 12th Gen Alder Lake and the 13th Gen Raptor Lake technologies to perform approximately equal, assuming similar clock speeds. Benchmark testing was performed under Windows 10 using PassMark Performance Test 10 software, which allowed specific testing of integer, floating point, and SSE/AVX instructions, as well as measuring performance in several real-world composite application algorithms. Tests were performed under controlled core usage configurations, allowing for a directed understanding of P-core vs. E-core performance, as well as comparing primary core vs. hyper-threading performance. Processor power was monitored using the CPUID HWMonitor Pro utility to determine processor efficiency.

Intel Performance Classification

While Intel does not explicitly declare the performance levels of cores vs. threads, a look into their thread classification algorithm in publicly available OS scheduler software finds Intel classifying the performance levels on hybrid processors in the following order (from highest performance to lowest performance):

- One thread per core on a P-core
- One thread on an E-core
- SMT (simultaneous multi-threading or "hyper-threaded") threads on P-cores

These performance levels mean that a system will load up one thread per P-core, then load up all the E-cores, before next moving to the hyper-threads on the P-cores. Using hyper-threads on P-cores may reduce performance to below that of an E-core, however, Intel does not publish its rationale behind these performance considerations. It may be based on absolute performance or due to non-obvious concerns such as shared bus contention, thermal hotspots on hyper-threading P-cores vs. distributed thermal heat when using P and E cores separately, etc. These performance classifications serve as guidelines for conventional OSs when they are tasked with assigning processes to processing cores.

Testing Scenarios

We evaluated the following testing scenarios, where "P" represents a P-core, "E" represents an E-core, and "T" represents threads.

1. Intel "Serpent Canyon" model NUC12SNKi72, https://www.intel.ca/content/www/ca/en/products/sku/231480/intel-nuc-12-enthusiast-mini-pc-nuc12snki72va/specifications.html
2. Precision 7670 Workstation | Dell Canada

P-Core Only

1P1T This scenario provides a baseline best-case for a single-threaded P-core.

1P2T This scenario provides an indication of the performance of a hyper-threading core, and is useful when comparing against 1P1T to understand if two hyper-threading cores offer double the performance of a single-threading core or perhaps something less.

2P1T This scenario provides a comparison of two P-cores, useful when compared to the 1P2T hyper-threading scenario.

4P2T This scenario provides a relative benchmark for the large installed base of "quad-core" hyper-threading processors (i.e., 4th Gen Haswell i7-4700EQ or 5th Gen Broadwell i7-5850EQ).

6P2T This scenario provides maximum performance when using only P-cores of the processor and is useful for comparison against the installed base of 6-core hyper-threading processors such as the 9th Gen Coffee Lake Xeon E-2276ME.

E-Core Only

1E This scenario provides a baseline best-case for a single E-core.

8E This scenario provides maximum performance when using only E-cores of the processor.

Combined P-Core + E-Core

6P2T+8E This scenario provides maximum processor performance using all cores and threads.

Base Clock vs. Turbo Modes

Base Clock (aka non-Turbo) Tests performed with base clock and non-Turbo settings are generally more deterministic and power efficient.

Turbo Modes Tests performed under Turbo conditions can offer significant performance boosts but sacrifice determinism and power efficiency.

Performance Tests

Tests were conducted on two separate processor systems:

1. **Intel NUC12SNKi72 "NUC" computer with a 12th Gen Alder Lake i7-12700H processor (6P+8E):** With the processor Turbo mode disabled, the computer operates with a P-core base clock frequency of 2.3 GHz and an E-core base clock frequency of 1.7 GHz. Under Turbo conditions, P-clock frequency increases up to 4.7 GHz and E-clock frequency increases up to 3.5 GHz.

2. **Dell Precision 7670 laptop computer with a 12th Gen Alder Lake i7-12850HX processor (8P+8E):** With the processor Turbo mode disabled, the computer operates with a P-clock frequency of 2.1 GHz and E-core clock frequency of 1.5 GHz. Under Turbo conditions, P-clock frequency increases up to 4.8 GHz and E-clock frequency increases up to 3.4 GHz.

The tests were performed using PassMark Performance Test 10 software on systems running Windows 10, focusing on CPU benchmarks. As with most benchmark software, scores should not be used as absolute values, but rather as comparative values with the software running under different processor configurations and Turbo settings.

Core configurations were changed in the UEFI BIOS prior to booting. Scripts were developed to assign specific processor cores and threads to the benchmarking software. These scripts ensured that the benchmark software processes ran on the designated test cores (aka: core affinity), while all other Windows processes were directed to non-benchmark software processor cores.

Comparing P-cores vs. E-cores

Single Core P vs. Hyper-Threading P: Hypothesis and Expectations
Based on architecture analysis and normalizing all results to show performance relative to the 1P1T (best-case) configuration, we can expect the following generalized results:

- 1P1T is our baseline performance using a single thread on a single processor core.
- 1P2T is not expected to perform fully twice that of a 1P1T test. Dual-thread cores typically share L1 cache and main memory, and this cache contention reduces the overall performance of both threads. While the total performance is significantly higher than a single thread, each thread is expected to be less than 100%, potentially lowering the performance for both threads.
- 6P1T would similarly perform below 6x 1P1T, due to core contention to shared L3 cache and main memory.

P-Cores vs. E-Cores: Hypothesis and Expectations
With the emphasis on E-cores being power efficient, we expect to see significant power reductions for workloads operating on E-cores compared to P-cores. We also expect lower performance levels as the compromise penalty in exchange for the resulting power savings[3].

Test Results
Testing was performed on the NUC computer (i7-12700H 6P+8E @ 2.3+1.7GHz) to show the effect of thread operation on P-cores, hyper-threading P-cores, and E-cores. Figure 1 summarizes these tests, with bars indicating Integer Math and Composite test scores. The orange line and numbers represent performance relative to the 1P1T test score.

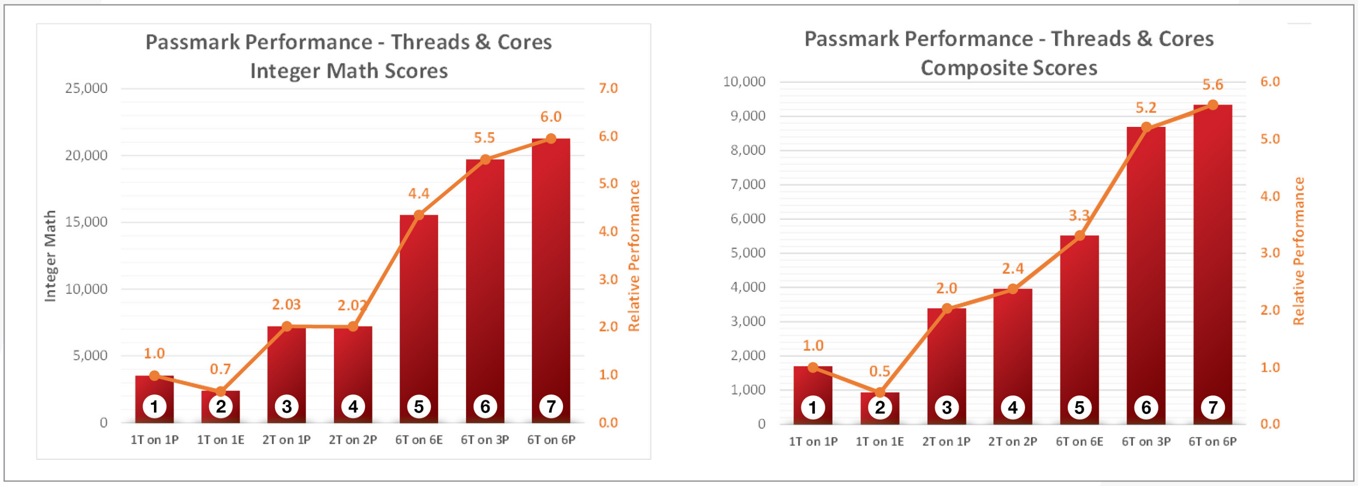

Figure 1: Performance of Threads and Cores

The following tests were performed using a single thread operating on a P-core as our baseline (❶):

❶ 1 thread operating on a P-core
❷ 1 thread operating on an E-core
❸ 2 threads operating on a single hyper-threading P-core
❹ 2 threads operating on separate P-cores
❺ 6 threads operating on 6 E-cores
❻ 6 threads operating on 3 hyper-threading P-cores
❼ 6 threads operating on 6 P-cores

3. Intel does not publicly quantify the power or performance differences between P-cores and E-cores, and upon request, they did not offer any further insight.

Test results confirm our expectations with the following observations:

- An E-core performs at ~70% of a P-core (❷) for Integer Math operations and ~50% for overall Composite workloads.
- While there may be a small difference between two threads operating on a single hyper-threading P-core vs. two separate P-cores (❸ vs.❹), the difference was negligible in Integer Math tests and more pronounced for Composite workloads.
- Multiple thread scenarios may show a very slight improvement because of shared cache efficiencies (❸ and ❹ are slightly higher than double ❶).
- Six threads running on independent P-cores perform at approximately 5.6x to 6x that of a single P-core (❼).
- Six threads running on three hyper-threading P-cores performed less than when running on independent P-cores (❻ vs. ❼).
- Six threads running on E-cores performed at 73% (Integer Math) and 59% (Composite) compared with independent P-cores (❺ vs. ❼) and performed at 80% and 64% compared to hyper-threading P-cores (❺ vs. ❻).

Discussion and Take-Aways

Comparing P-core vs. E-core performance: these tests demonstrate that E-cores operate at 50-73% compared to a P-core. However, when considering that P-cores operate with a clock frequency of 2.3 GHz and E-cores operate with a clock frequency of 1.7 GHz, the clock speeds of the E-cores are 26% lower than the P-core clock speed and would thus be expected to provide lower performance consistent with this 26% lower clock speed. P-cores and E-cores, if operated at the same frequency, may provide similar performance, at least for integer math operations. Under composite workloads, the E-cores produced scores of approximately 50% compared to the P-cores, perhaps highlighting other performance compromises made by Intel in the pursuit of higher efficiency.

When considering hyper-threads, there is a slight decrease in performance compared to separate cores. Comparing hyper-threads vs. independent cores (❻ vs. ❼), an 8% performance reduction was measured.

Comparing Base Clock vs. Turbo Operation

Hypothesis and Expectations

Processor performance is directly proportional to clock speed. Without a doubt, we expect far better performance with Turbo clock operation when compared to base clock operations due to its increased clock speeds. However, we must also be prepared to accept a power penalty for this performance boost. When the processor is operating at the highest turbo speeds, there is simply no comparison: we expect Intel processor performance with Turbo clocks to surpass any non-Turbo performance by a factor of 2:1, assuming the Turbo clock speeds are approximately double that of base clock speeds.

Test Results

Testing was performed on the Dell 7670 computer (i7-12850HX 8P+8E @ 2.1+1.5GHz, Turbo to 4.8+3.4GHz).

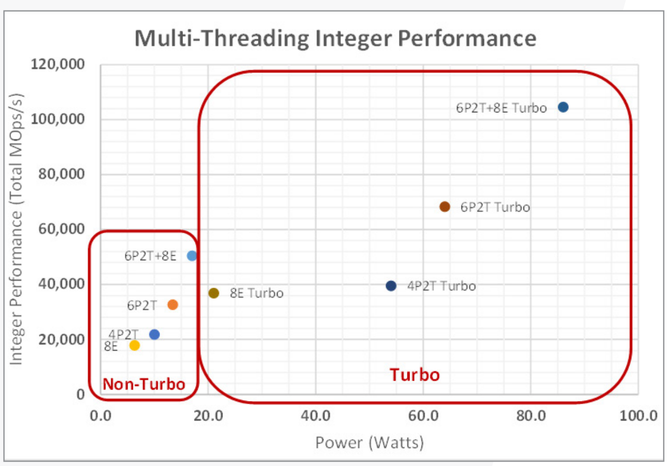

Figure 2: Base vs. Turbo Integer Performance for Different Core Configurations

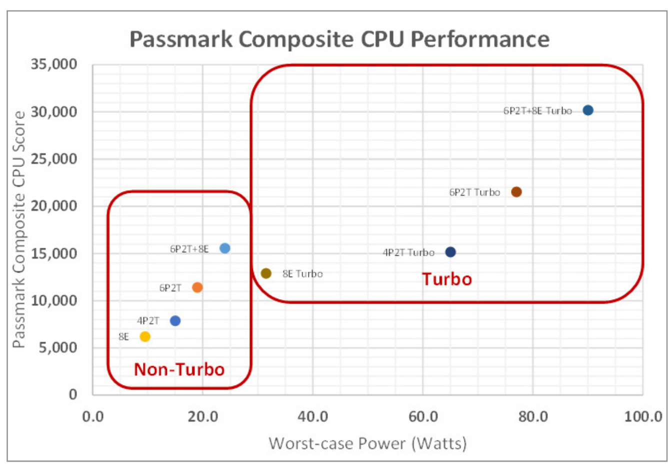

Figure 3: Base vs. Turbo Composite Performance for Different Core Configurations

Discussion and Take-Aways

Figure 2 and Figure 3 make starkly evident the massive performance boost gained by Turbo operations. In all test configurations, the Turbo performance roughly doubled the non-Turbo performance. However, while the non-Turbo processor power consumption was generally within a narrow band, the processor power consumption when operating under Turbo configurations was considerably higher, consuming and dissipating as much as 5.3 times more power for a doubling in performance. Clearly, power efficiency is sacrificed with Turbo clock configurations.

Figure 4 plots Integer performance on a per-core basis. Again, we see approximately double the performance of each processor configuration when operating with Turbo clocks and the high-power consumption of the Turbo configurations. We can also see that the per-core performance of the 6P2T+8E configuration falls between the "E-core only" and the "P-core only" test cases. This performance drop is due to shared resource contention within the processor when both the P-cores and E-cores are operating.

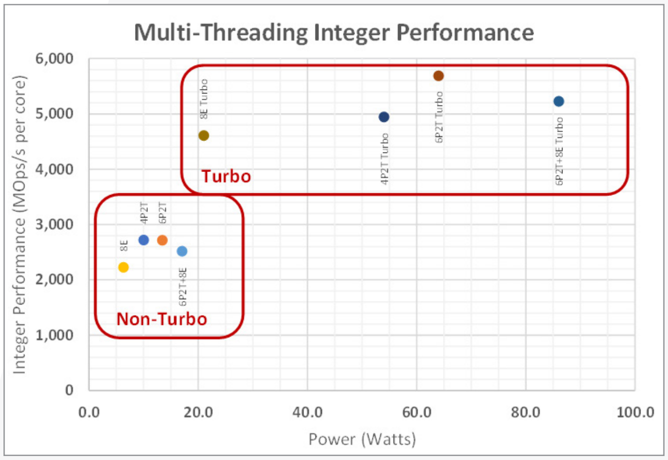

Figure 4: Base vs. Turbo Performance, Multi-Threading Integer Performance, Plotted Per-Core

Comparing P-cores vs. E-cores

Power Efficiency of P vs. E Cores: Hypothesis and Expectations

With the considerable efforts that Intel has put into developing its Efficient E-cores, we expect to see significantly higher power efficiency compared to the standard P-cores used in most mainstream computing platforms today.

Power Efficiency of P vs. E Cores: Test Results

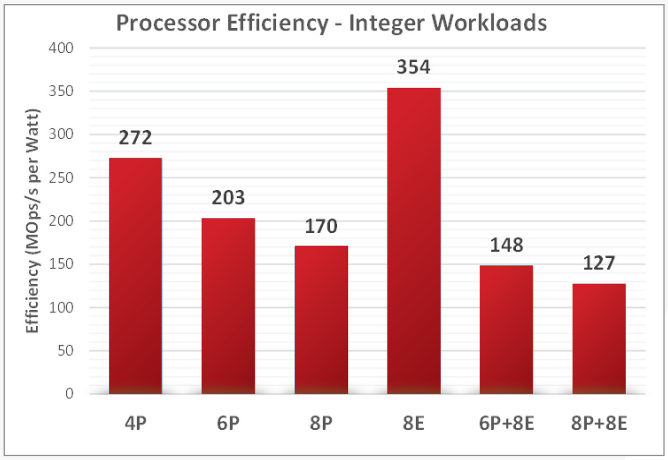

Figure 5: Power Efficiency of P vs. E Cores, Integer Workloads

Power efficiency was calculated by simply dividing the test performance score by the processor power consumption to determine an efficiency measurement in MOps per watt. Reviewing Figure 5, we observe:

- As the number of cores increases, efficiency decreases. This would be expected as we know overall performance per core will decrease due to shared resource contention.
- Looking at the 8-thread test case, the efficiency of eight E-cores (354 MOps/W) is more than double that of eight P-cores (170 MOps/W).

Perhaps the most surprising result of this testing was the truly massive increase in power efficiency of Intel's E-cores. While supporting 100% compatibility with P-cores, the E-cores offer more than twice the power efficiency over P-cores. As we learned earlier, we only lose approximately 30% in absolute performance for integer workloads.

Power Efficiency of Base vs. Turbo Clocks: Hypothesis and Expectations

While operation with Turbo clocks offered roughly double the performance, it came at a power penalty of more than five times for some test cases. As Turbo power consumption does not track performance linearly, we can expect that processors operating with Turbo clocks will also suffer from power efficiency losses.

Power Efficiency of Base vs. Turbo Clocks: Test Results

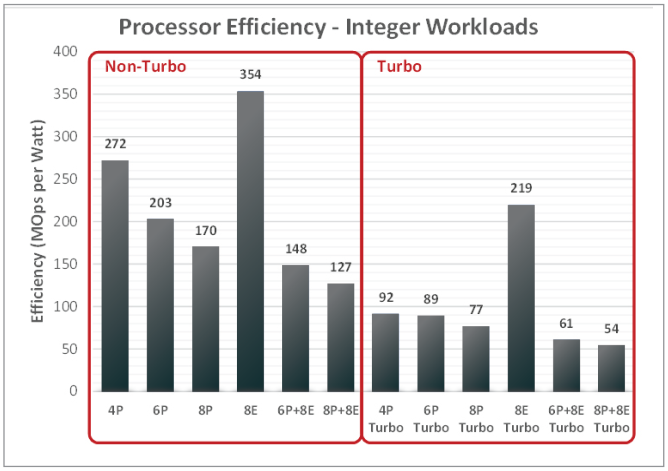

Figure 6: Power Efficiency of Non-Turbo vs. Turbo Clock Settings

Examining Figure 6 we compare the performance efficiency under non-Turbo base clock vs. Turbo clock settings. While Turbo clock modes offer significantly higher overall performance, they come at a significant power penalty, considerably reducing power efficiency.

Important Considerations for Turbo Operations

Curtiss-Wright strongly recommends against operating with Turbo clock modes for several reasons. First and foremost, Intel processor reliability analysis is only performed under base clock (non-Turbo) conditions[4]. Operating with Turbo clocks may reduce the operational life and reliability of the processor.

Operating with the Turbo clocks enabled also produces non-deterministic performance, as the processor clocks (and thus performance) are not predictable. Intel does not publish the details of Turbo clock algorithms, and we cannot determine when a processor clock will go into Turbo modes or how variable the clock speeds will be.

Finally, Curtiss-Wright has evaluated actual processor die thermal performance under Turbo operations. We have measured die "hot-spots" which can measure 15-20°C higher than other parts of the same silicon die. As a processor die is generally cooled equally across the entire die, these hot-spots will quickly bring the processor into thermal throttling, responding to the over-temperature condition by lowering processor clock rates – effectively canceling the benefits of Turbo operation. It has been demonstrated that better overall performance can be consistently achieved at high operating temperatures when the processor is configured for non-Turbo operations.

Core Configurations of Popular Intel Processors

Intel processors are not normally known by their P-core or E-core designations. Instead, we generally refer to the processors by their generational trade names and SKU part numbers. A list of commonly used Intel processors for embedded applications in the Aerospace and Defense industry is summarized in in the Table below.

Intel Core Generation	Family Name	Processor SKU	Core Config.	Threads
4th Gen Core	Haswell	i7-4700EQ	4P	8
5th Gen Core	Broadwell	i7-5850EQ	4P	8
5th Gen Xeon D	Broadwell DE	Xeon D-1559	12P	24
9th Gen Core	Coffee Lake	E-2276ME	6P	12
10th Gen Ice Lake Xeon D	Ice Lake D	D-1700 (LCC) D-2700 (HCC)	10P 20P	20 40
11th Gen Core	Tiger Lake	W-11865MRE	8P	16
13th Gen Core	Raptor Lake	i7-13800HRE	6P+8E	20

We often compare processors with very high-level parameters, such as the number of cores or the overall clock speed. With today's OSs, there is little consideration for how cores and threads are used. As discussed in Part 1 of this White Paper series, a contemporary OS carries a complexity tax of many hundreds of background processes and daemons that slowly steal precious processing cycles.

With the introduction of hybrid processing cores, we can now assign processes to different types of cores, making more optimal use of the performance levels needed by each class of usage. Under real-world conditions, a 20-core processor with lower overall performance may, in fact, operate more efficiently and responsively than a 16-core processor with higher overall performance.

4. Per Intel: "Industrial use conditions are base frequency operation only – no core or graphics turbo operation enabled or active."

CURTISS-WRIGHT

Benefits, Drawbacks and Recommendations

Do Background Applications Really Matter?

Yes, background applications matter. For example, the Windows 10 laptop computer used to create this White Paper had five foreground applications running, a whopping 136 background processes, and 111 active Windows processes running simultaneously. With many processes spawning multiple threads, the total thread count being managed is well over 500. It is hard to fathom all the things these background tasks are doing and monitoring, but they are all running, waiting for their turns on this author's older quad-core (Intel i5-7440HQ = 4P1T) processor. To run these background processes, the benefits of higher efficiency E-cores would definitely shine. It is highly unlikely that these background processes need the performance of a P-core, yet they are all taking their turns executing on the Performance cores and draining this laptop's battery. A few extra minutes of battery life to help reduce battery anxiety would certainly be welcomed.

For embedded systems where battery life is less critical, the extra power savings may seem insignificant when weighing the benefits of manually controlling core usage with a hybrid core processor against the complexity of this new undertaking. Yet, when considering the secondary benefits, one can realize with every Watt of power saved, the true benefits will multiply. Smaller power supplies and reduced cooling requirements lead to smaller systems with reduced weight for a smaller size, weight, and power (SWaP) footprint, and ultimately, these SWaP savings translate into extended mission duration and reduced operating costs.

Effective use of P-cores and E-cores for Embedded Systems Designers

In a desktop environment, we let the OS decide how to manage hardware and software resources. As users, we are content to accept the "mass-market" driven benefits that deliver good enough performance vs. power savings. However, in a highly engineered embedded system, where real-time performance and power savings matter, the implications of processor performance and efficiency must be more closely considered.

Embedded systems, which are highly engineered to perform a given set of tasks or functions, are typically well understood during systems design. Deep knowledge of an embedded system's processor power consumption, thermal dissipation, and software performance is crucial to reduce the system's SWaP footprint. These benefits are even more critical for airborne systems, where every additional pound of weight can have far-reaching implications.

Utilizing the latest Intel hybrid core processors can provide incredible benefits that have never been imagined before. Operating software efficiencies can be increased, reducing power consumption. System performance can be fine-tuned to take advantage of Performance P-cores when performance is needed, or applications can be assigned to Efficient E-cores to save power or reduce thermal loads when appropriate. Real-time systems can achieve better determinism and real-time response by utilizing the higher number of logical cores presented to the OS and carefully controlling thread affinity to meet real-time objectives.

Initially, letting an OS decide how to manage threads and cores automatically may be the easiest approach. After all, today's OSs are extremely competent and well designed to manage a wide range of application types. Over time, though, the software developer's ability to fine-tune these parameters will yield positive results that were not possible before the introduction of the Intel hybrid core processor.

Conclusion

As processing technologies evolve in the fast-paced commercial world, Curtiss-Wright constantly evaluates these new capabilities to understand how our Aerospace and Defense customers can best leverage these commercial developments. Some of the technologies described in this paper, such as multi-core and hyper-threading processors, are widely used today, and some, such as hybrid core processors with P-cores and E-cores, are just now beginning to attract the interest of industry architects. While the Aerospace and Defense industry frequently lags behind the commercial world, due to an appropriate abundance of caution and the criticality of the systems it relies on, the industry eventually embraces new technologies, taking advantage of their benefits while steering clear of their shortcomings. As this process is ongoing and as we start to realize these newfound benefits, other new emerging technologies will need to be analyzed and considered.

Related Products

Curtiss-Wright has a long history of providing COTS processing technology to the Aerospace and Defense industry. We offer the broadest range of performance vs. power and efficiency processing module technologies in many industry standard form factors. The list below summarizes a few of our most popular processing technologies in use today.

Product	Form Factors	Architecture	Processor	CPU Configuration	DMIPS Performance[5]	DMIPS per Watt (efficiency)
VPX3-1703	3U VPX	Arm A53	NXP LS1043A, 9W	4P1T @ 1.6 GHz	19,584	2,176
XMC-120	XMC Mezzanine	Intel Atom "Bay Trail"	Atom E3845, 10W	4P1T @ 1.9 GHz	33,060	4,723
VPX3-133 VPX3-152 VME-196 VPX6-197	3U VPX 3U VPX VME 6U VPX	Power Architecture	NXP T2080, 25W	4P2T @ 1.8 GHz	38,880	1,555
XMC-121 VPX3-1220	XMC Mezzanine 3U VPX	Intel 7th Gen Core "Kaby Lake"	Xeon E3-1505Lv6, 25W	4P2T @ 2.2 GHz	235,664	9,944
VPX3-1260 VME-1910	3U VPX VME	Intel 9th Gen Core "Coffee Lake Refresh"	Xeon E-2276ME, 45W	6P2T @ 2.8 GHz	431,424	10,786
VPX6-1961	6U VPX	Intel 11th Gen Core "Tiger Lake"	Xeon E-11865MRE, 45W	8P2T @ 2.6 GHz	568,256	13,660
VPX3-1708 V3-1708	3U VPX 3U VPX Safety Certifiable	Arm A72	NXP LX2160A, 31W	16P1T @ 2.0 GHz	198,400	6,400
VPX3-482 (XD1)	3U VPX	Intel 5th Gen Broadwell DE	Xeon D-1539, 35W Xeon D-1559, 45W	8P2T @ 1.6 GHz 12P2T @ 1.5 GHz	306,176 (8c) 430,560 (12c)	8,748 (8c) 9,568 (12c)
VPX6-483 (XD2)	6U VPX	Intel 5th Gen Broadwell DE	Xeon D-1559, 45W Xeon D-1587, 65W	Dual processor, each: 12P2T @ 1.5 GHz 16P2T @ 1.7 GHz	Each CPU: 430,560 (12c) 650,624 (16c)	Each CPU: 9,568 (12c) 10,010 (16c)
VPX3-484 (XD3)	3U VPX	Intel 10th Gen Ice Lake D	Xeon D-1746TER, 67W	10P2T @ 2.0 GHz	594,800	8,878
VPX6-485 (XD4)	6U VPX	Intel 10th Gen Ice Lake D	Xeon D-2752TER, 77W Xeon D-2775TE, 100W Xeon D-2796TE, 118W	Dual processor, each: 12P2T @ 1.8 GHz 16P2T @ 2.0 GHz 20P2T @ 2.0 GHz	Each CPU: 642,384 (12c) 951,680 (16c) 1,189,600 (20c)	Each CPU: 8,343 (12c) 9,517 (16c) 10,081 (20c)

5. DMIPS Performance is calculated from manufacturer provided processor DMIPS/MHz specifications and are the theoretical maximum number of instructions that can execute per second. They do not take into consideration other processor capabilities, such as floating point or vector performance using AVX2/AVX512 instructions, etc.

Actual performance will be highly dependent on many other system considerations, as well as algorithm architectures such as the ability to utilize all cores and all threads simultaneously. These values are best used as general-purpose representative values when comparing one processor to another.

In 2023, Curtiss-Wright announced its first Intel hybrid core processing module, the VPX3-1262. Following the very successful VPX3-1260, which is based on Intel's 9th Gen Coffee Lake (6P) processor, the VPX3-1262 features an Intel 13th Gen "Raptor Lake" i7-13800HRE processor with a 6P8E core configuration and offers a total of 14-cores and 20-threads.

Product	Form Factors	Architecture	Processor	CPU Configuration	DMIPS Performance	DMIPS per Watt
VPX3-1262	3U VPX	Intel 13th Gen Core "Raptor Lake"	i7-13800HRE, 45W	6P2T + 8E = 14-core, 20-thread 6P2T @ 2.5 GHz + 8E1T @ 1.8 GHz	527,822	14,265

Author

Aaron Frank
Senior Product Manager
Curtiss-Wright Defense Solutions

Aaron Frank joined Curtiss-Wright in 2010 as a Senior Product Manager for the C5ISR group. He is responsible for a wide range of COTS products utilizing advanced processing, video graphics/GPU, and network switching technologies in many industry-standard module formats. Aaron has a long history of working with and promoting open standards, participating in IEEE and SMPTE standards since 1990, and was personally awarded an Emmy from the National Academy of Television Arts and Sciences for his work to create the openGear platform (www.opengear.tv), an open equipment standard used by over 80 of the top broadcasting corporations worldwide. Aaron holds a Bachelor of Science in Electrical Engineering from the University of Waterloo.