

# **Movellus Aeonix: An Intelligent Clock Network**

By Linley Gwennap Principal Analyst and Aakash Jani  
Senior Analyst

December 2021



The Linley Group

[www.linleygroup.com](http://www.linleygroup.com)

# Movellus Aeonic: An Intelligent Clock Network

By Linley Gwennap, Principal Analyst, The Linley Group and Aakash Jani, Senior Analyst at The Linley Group

*Movellus is increasing silicon clock performance, power, and reliability through its Aeonic technology, which uses optimized IP to mitigate on-chip variation, skew, and voltage droop. It's well suited to high-performance systolic arrays, low-power edge IoT devices, and similar designs. Movellus sponsored this white paper, but the opinions and analysis are the authors'.*

Twenty years ago, many people were skittish about mail-order shopping. Orders placed over the phone or (courageously) online could take weeks or even months to arrive. Since then, Amazon completely changed the shopping paradigm by moving to distributed warehouses and optimizing supply chains. Shipping times dropped from months to a week or less, and the company even provides same-day delivery in some cases. It accomplished this feat by adding intelligence to logistics.

Many aspects of hardware design still lack intelligence, creating opportunities for newcomers. For decades, SoC architects were stuck with two options, clock trees and meshes that limited them in power or performance. Movellus' clocking solution follows Amazon's path by adding intelligence to clock networks. Intelligent clock networks reclaim performance and power by facilitating communication between the network's drop-off points.

Intelligent clock networks do much more, however: They create larger regions of synchronization. They reduce data bottlenecks between heterogeneous intellectual-property (IP) blocks. And they enable deterministic execution, necessary for autonomous driving and robotics. In short, intelligent clock networks will accelerate the production of emerging hardware designs.

## Introduction

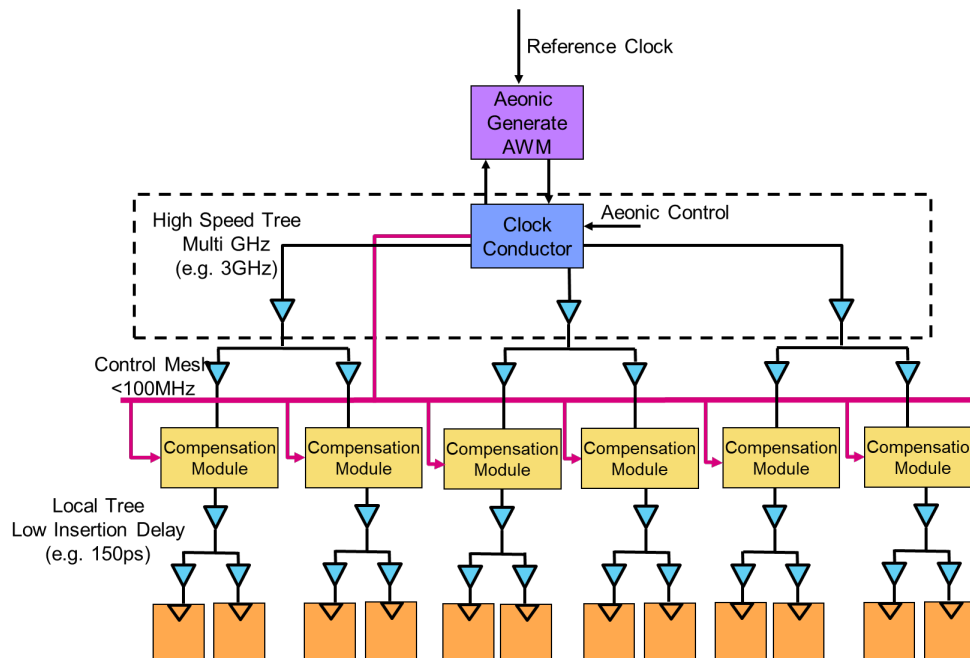
The clock signal is essential to any processor, yet it often receives little attention compared with more-glamorous features such as the CPU microarchitecture and memory subsystem. Leading-edge chips must carefully route the clock signal to every nook and cranny, consuming critical layout resources and as much as 8W of the chip's total power, even while running at 40% utilization. Clock-signal imperfections detract from the timing budget, forcing the entire chip to run slower. Rather than ignore the clock, Movellus has embraced it, delivering new technology that simplifies clock design and ultimately enables faster, lower-power chips.

Clock trees were the first widely used approach, back when processors were simpler and frequencies lower. They begin with a single clocking source, such as a ring oscillator, and

distribute that signal through a branching set of buffers that eventually reach all end points (clock consumers) in the chip. The time difference between the end points (called *skew*) varies on the basis of branch length and number of buffers. This architecture minimizes clock power, but as the tree grows, so does skew. Some manual tricks can reduce the skew, but as designs become more complex, it eventually grows too large. The higher the target clock speed, the smaller the allowable skew.

To solve this problem, designers turned to clock meshes. This architecture drives the clock source into a mesh of wires that equalize the skew, and it makes balancing synchronous clock networks easier. But the extra wires result in power-hungry meshes and consume lots of routing area. Additionally, each chip requires a custom mesh design, which necessitates considerable engineering resources and expertise.

An evolution of both clock technologies, Movellus' Aeonix combines the benefits of a high-speed clock tree and clock mesh, as Figure 1 shows. It relies on the company's all-digital IP portfolio, which comprises an adaptive-workload-management (AWM) module, a clock conductor, and a smart clock module (SCM). Aeonix helps chip architects reduce power and increase performance by controlling process variation and combating variation in current and voltage noise.



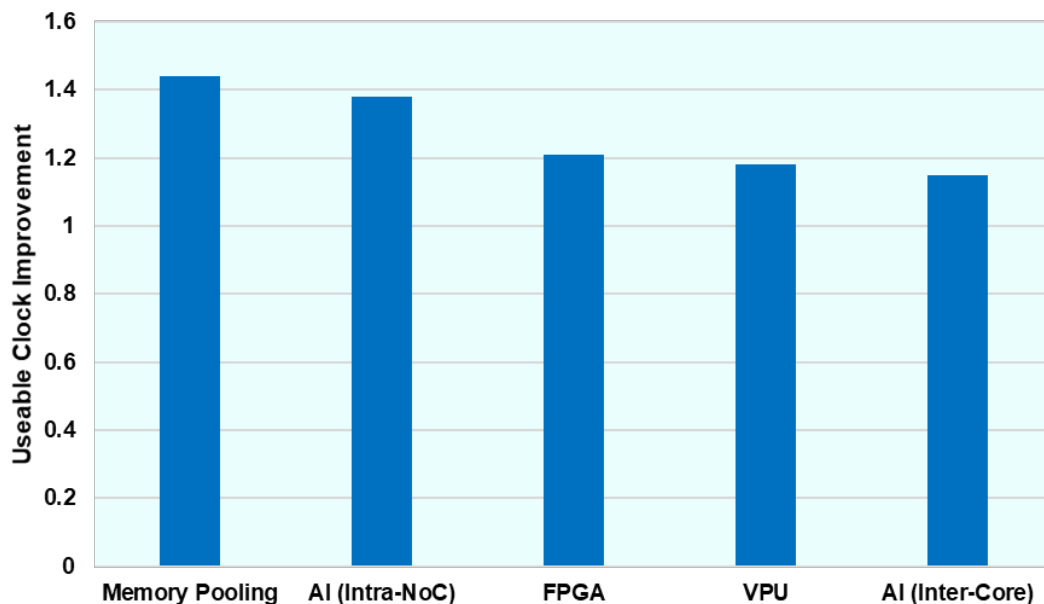
**Figure 1. Aeonix Diagram.** Movellus' Aeonix combines a clock tree with a mesh network to deliver large reductions in skew and on-chip variation while also reducing power compared to a traditional mesh.

## Increasing Clock Frequency

Chip designers can extract linear performance improvements from pipelined architectures by increasing the peak frequency ( $F_{MAX}$ ). But this task sounds easier than it is. In advanced nodes such as 7nm and 5nm, random process variation alters the physical properties of nets, vias, and transistors. This variation means that some clock paths may be faster than others and that each die could have a different set of fast and slow paths. During the clock-design (CTS) stage, designers hedge their bets against this variation by adding margin to the timing budget, thus limiting  $F_{MAX}$ .

Aeonic helps recapture lost frequency by optimizing the timing margins. Using the control mesh, the SCMs automatically detect timing changes caused by on-chip variation (OCV) and Voltage/ Temperature fluctuations. If a path is too fast or too slow, the SCM compensates by providing the exact delay necessary to balance the path. This approach can accommodate any combination of fast and slow paths. By neutralizing the effect of OCV as well as VT changes, the Movellus technology allows clock designers to reallocate the timing margin to increase the clock frequency.

The new timing technology is application agnostic, since it only depends on the design's clock. Through signal analysis on live silicon, Movellus validated its efficacy across various designs including AI accelerators, FPGAs, and VPUs. As Figure 2 shows, the technology can extend the usable clock period by up to 44%, giving chip architects a healthy performance boost. Especially for AI devices, this boost translates directly to greater theoretical throughput.



**Figure 2. Useable clock period improvements.** Aeonic works across various architectures and compute domains. It can boost usable clock period by up to 44% and averages a 27% gain across the above designs.

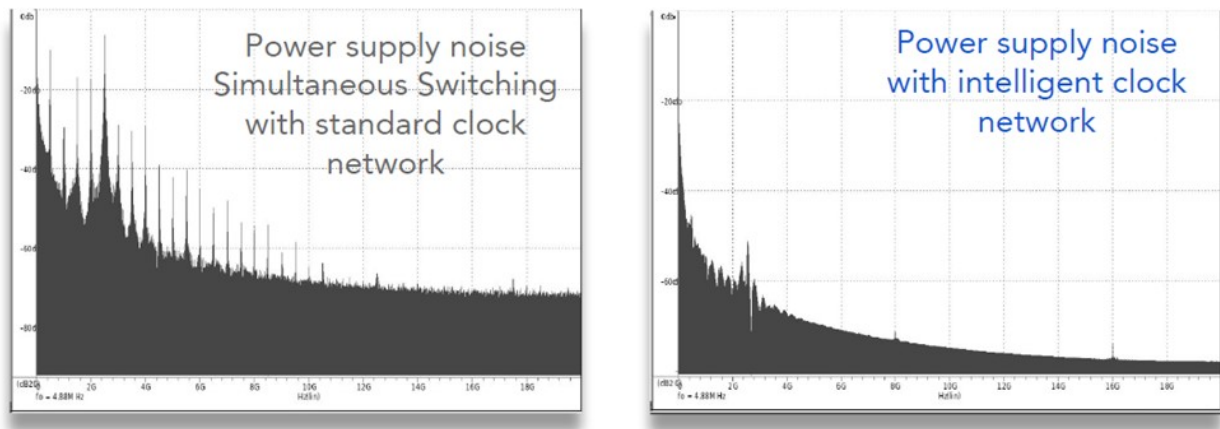
## ***Reducing Latency***

In manycore designs, such as GPUs and systolic arrays in AI accelerators, chip architects occasionally deactivate cores to conserve power. Consider a general-purpose AI accelerator. It will process neural networks of varying sizes; for smaller ones, it can switch off some cores, and for larger ones, it can switch on all cores. This process is software controlled during run time, but it isn't instantaneous. By constantly toggling cores, chip architects introduce latency caused by voltage droop (the immediate power decrease that results from increasing load).

Movellus' IP portfolio features an Adaptive Workload Module (AWM) to decrease this latency: AWM has a fast 10ns frequency-transition time. It helps reduce latency without creating system stalls. Because of its recovery time, AWM is well suited to chips with large synchronous clock domains—systolic arrays in AI devices, for instance.

Lower latency is critical for AI inference. According to AI-consortium MLCommons, edge inference times for object detection range from half a millisecond to several milliseconds. AWM can help reduce these inference times, providing an advantage in saturated AI markets such as automotive edge devices.

Voltage noise produces clock jitter, which causes timing errors by creating imprecise clock-capture windows. To accommodate these jitter effects, timing architects add margins or high-accuracy but expensive mica capacitors. By adapting to changes in power at run time, the adaptive phase-shifting modules and Aeonic help reduce voltage noise in clocking structures by up to 16x, as Figure 3 shows. To ensure each transistor switches at the same frequency in a synchronous domain, the module actively adapts to voltage changes. With a more reliable clock source through Adaptive Phase Shifting (APS) integration, timing architects can reduce their jitter budget to boost  $F_{MAX}$  by up to 30%.




**Figure 3. Reduction in peak current and voltage noise.** Aeonic is highly adaptable to changes in switching activity, which helps reduce the peak current by up to 24x.

### ***Cutting Power***

With the Internet of Things causing edge silicon to proliferate, small power envelopes have become paramount. Edge devices for factories and smart cities often run on just a few milliwatts. Instead of increasing performance, architects of these products focus on limiting power consumption, potentially making the difference between needing a wall outlet and needing just a battery.

At low voltages, silicon becomes more sensitive to fluctuations in voltage drop and voltage noise. To address this issue, architects must either increase the supply voltage (to reduce relative sensitivity) or buffer clock paths to resolve setup violations (where the data comes too slow). Either solution increases power—anathema in this segment.

Instead of using the lower OCV and jitter budgets to raise  $F_{MAX}$ , chip designers can slow the clock to reduce power by a factor of the square root of the frequency. Additionally, the Aeonic architecture boosts power efficiency at the transistor level by limiting the effects of simultaneous-switching noise (voltage drops). The timing architecture runs each subsystem at a separate phase, reducing IR spikes by 4x. It relaxes the lowest voltage for timing closure by 100mV. The technique produces a useful clocking cycle that's 50% larger, translating to an average 18% power savings, as Figure 4 shows.



**Figure 4. Power savings per architecture.** Aeonic limits the effect of simultaneous-switching noise, thereby relaxing the lowest voltage for timing closure by 100mV. This effect translates to an average 18% power decrease.

## ***Boosting Longevity***

Once deployed, silicon’s physical properties change over time owing to environmental stressors. For example, in high-voltage and high-temperature environments, critical timing paths will degrade because of natural electron drift. The chip’s performance will decline as a result, eventually leading to failure. In hard-to-reach deployments, such as cellular base stations and smart factories, replacement of a failing chip can cost the company more than the chip itself.

In specific applications, such as orbital systems, data centers, and nuclear-power plants, silicon experiences high particle and EMF radiation. High-energy radiation can create enough noise to disrupt transistors across the die, necessitating recalibration. Movellus’ all-digital IP portfolio includes radiation-hardened libraries that are process agnostic. They extend the lifetime of these difficult-to-replace parts and improve their functional safety for complex and dangerous tasks.

During the design process, engineers finish timing closure for low-level blocks (CPUs, GPUs, and other subsystems) and then send abstracted timing models to full-chip timing engineers. The process creates a bottleneck during closure that can hold up full-chip-timing (FCT) closure, which is necessary for tapeout. Through its intelligent clock network, Movellus helps companies parallelize the two efforts. The network automatically balances skew and OCV, allowing FCT engineers to focus on the high-speed tree while block engineers work on the intelligent mesh. For several of its customers, the company’s IP accelerated timing closure by a full quarter.

## Conclusion

By redefining architectures, Movellus is tracing a path analogous to that of network-on-a-chip (NoC) vendors. At one time, every complex chip had its own hand-designed NoC that was often inefficient and required extensive design resources. Today's clock solutions are just as exhaustive. Each company uses one of two predefined options and spends countless man-hours optimizing scripts for IP and clock integrations. Movellus eases clock-network design while reducing power and clock skew, directly improving chip-level performance.

To generate an Aeonic configuration for a customer, Movellus needs PPA requirements, the silicon platform, and workload characteristics. It can then deliver encrypted hard- and soft-IP blocks, abstracted timing data, and tool-flow scripts for customer integration, simplifying most CTS duties.

By using Movellus' silicon-proven technology, chip designers have produced working silicon with 38% greater performance or 10–30% lower energy consumption. Pair these gains with process-node advancements, and the potential result is more than 50% performance gains.

Across the clock-design and closure stages, the Aeonic architecture provides enhancements over existing architectures and creates new opportunities for greater clock synchronization across silicon. It improves chip-level performance and power, helping processor companies outdo the competition. System designers can reduce their time to market, increasing their chances of adoption. And finally, chip vendors can deliver more-resilient products, bolstering their reputations and their relationships with customers.

*Movellus is breathing new life into a once stagnant design component, boosting performance, power, and architectural freedom for its customers. For more information access <https://www.movellus.com/>*

*Linley Gwennap is a principal analyst at The Linley Group and a senior editor of Microprocessor Report. Aakash Jani is a senior analyst at The Linley Group and a senior editor of Microprocessor Report. The Linley Group, now a subsidiary of TechInsights, offers the most-comprehensive analysis of microprocessors and SoC design, going beyond business strategy to examine internal technology. Our in-depth articles cover topics that include embedded processors, mobile processors, server processors, AI accelerators, IoT processors, processor-IP cores, and Ethernet chips. For more information, access our website at [www.linleygroup.com](http://www.linleygroup.com).*