



MEDIATEK

CorePilot 3.0

*Max.Mid.Min (Tri-Cluster) technology
to maximize power efficiency with
extreme computing performance*

WHITE PAPER

Table of Contents

Introduction to Tri-Cluster: Max.Mid.Min.	4
Analysis of Recent Mobile Use Cases	5
Details of Max.Mid.Min. Technology	8
Mediatek Coherent System Interconnect (MCSI)	11
Thermal Management	12
Power Management (iCore)	12
Power Management (Fast DVFS)	13
Scheduler	14
Benefits of Max.Mid.Min. Technology	16
Energy Savings for Mobile Use Cases	17
Performance Uplift for Single Thread Applications	18
Sustained Performance Uplift for Multi-Thread Applications	19
Summary	19

List of Tables

Table 1 Trend of Mobile Use Cases	6
Table 2 Summary of Mobile Use Cases	8

Lists of Figures

Figure 1 Bad Fuel Efficiency When Using Wrong Gear for Loading	3
Figure 2 Good Fuel Efficiency When Using Right Gear for Loading	4
Figure 3 Tri-Cluster Architecture Briefing (Max.Mid.Min)	4
Figure 4 Tri-Cluster Key Points Briefing (Max.Mid.Min)	5
Figure 5 Categories of Mobile Use Cases (Source: Flurry Analytics)	5
Figure 6 Mobile Gaming Profiling	6
Figure 7 Social Messaging App Profiling	7
Figure 8 Utilities App Profiling	7
Figure 9 Browser Profiling	8
Figure 10 Performance Range of B.L. Architecture	9
Figure 11 Advantage of Min CPU	9
Figure 12 Advantage of Mid CPU	10
Figure 13 Advantage of Max CPU	10
Figure 14 CorePilot v.3 Technology Block Diagram	11
Figure 15 Tri-Cluster+MCSI Block Diagram	12
Figure 16 Traditional Way of Enabling Cluster1	14
Figure 17 Lower DVFS OPP Leads Higher Energy Efficiency.	15
Figure 18 Advantages of Fast DVFS: (1) Performance Gain (2) Power Saving	15
Figure 19 Intelligent Core Activation Technology Example	16
Figure 20 System Runs at Higher Frequency of CPUs of Min & Mid Clusters	17
Figure 21 Instant Max CPUs Activation Example	17
Figure 22 B.L. and Optimized Performance-Power Curve	18
Figure 23 Power Comparison Between B.L. and Max.Mid.Min.	19
Figure 24 Power Comparison Over Varying Mobile Use Cases	19
Figure 25 Peak & Sustainable Performance Comparison for Single-Thread App	20

MediaTek CorePilot v.3

In 2013, MediaTek delivered CorePilot v.1 which enabled the industry's first heterogeneous multiple processing (HMP) of big.LITTLE architecture. Since then, big.LITTLE architecture has been adopted widely as one of the industry standards. This Dual-Cluster architecture combines high-performance CPUs with energy-efficient CPUs on the same SoC to reduce energy consumption (and hence preserve battery power), while still delivering peak performance. Since both CPUs are architecturally compatible, workloads can be assigned to each CPU, on demand, to suit performance needs. High intensity tasks like games are assigned to the high-performance CPUs, for example, while less demanding tasks like email and audio playback are assigned to the energy-efficient CPUs. The load balancing triggers quickly enough to be completely transparent to the user and can greatly reduce energy consumption for common tasks.

From 2013 to 2015, high-end mobile devices' computing power has increased dramatically while the available power has remained at a consistent level. Power efficiency has now become the dominant issue for mobile devices. In addition, mobile use cases for ordinary people have grown more diverse and varied. For high-end smart phones, the existing big.LITTLE architecture is unable to reduce power consumption to mitigate thermal issues. For instance, medium workload may be generated when heavy mobile users play games. big.LITTLE lacks the intermediate level of CPU core that could provide sustainable performance and power efficiency. Thus, the system powered by big.LITTLE is very likely to allocate the big CPU to process a medium workload which causes an unnecessary waste of energy. Another case is medium-to-light workload. Most of the time, social messaging apps require low computing power on average, but a medium-to-light workload is still needed for some specific operations like scrolling the page to browse messages. The system tends to allocate the LITTLE CPU to process medium-to-light workloads which may degrade user experience. As shown in Figure 1, taking the car as an example, fuel efficiency is degraded when using second gear for medium loading and first gear for heavy loading.




Figure 1 Bad Fuel Efficiency When Using Wrong Gear for Loading



Figure 2 Good Fuel Efficiency When Using Right Gear for Loading

MediaTek's CorePilot v.3 introduces a scheduling algorithm for Tri-Cluster architecture. In comparison with conventional approaches to lowering the voltage and frequency of big CPU to act as the intermediate level of CPU, the Tri-Cluster architecture could lead to additional power reduction. In terms of peak performance, clustered Max CPUs deliver the highest performance, followed by clustered Mid CPUs and Min CPUs. On the contrary, the Mid and Min CPUs are much more power-efficient than the Max CPU. Figure 2 shows that three gears provide for better fuel efficiency across different road conditions. Using the right gear for the right load is the key to fuel efficiency. In sum, Tri-Cluster enables mobile devices to provide optimized performance and an extended range of power efficiency.

Introduction to Tri-Cluster: Max.Mid.Min.

The industry's first Tri-Cluster CPU architecture proposed by MediaTek introduces a brand new Mid cluster. The block diagram of Tri-Cluster is depicted in Figure 3.

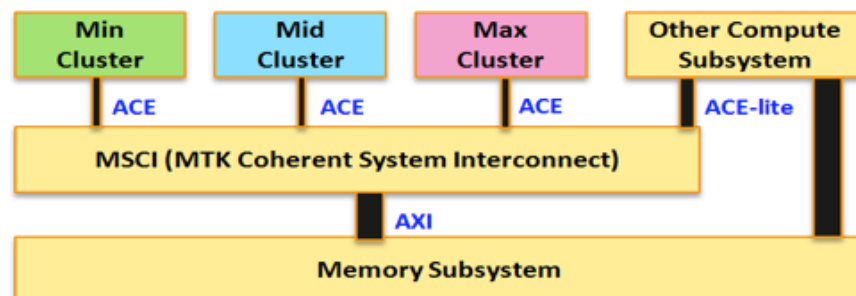



Figure 3 Tri-Cluster Architecture Briefing (Max.Mid.Min)

Following the block diagram of Tri-Cluster is the power performance curve. The left side of Figure 4 compares the differences in power and performance between Dual-Cluster and Tri-Cluster. The orange and blue blocks represent Dual-Cluster and Tri-Cluster,

respectively. Tri-Cluster provides more gears for different road conditions. On the right side of Figure 4, the key points of Tri-Cluster are listed. Point number two mentions two constraints of Dual-Cluster. Constraint one is the limited range of performance and power. The power of the LITTLE cluster is not low enough for light workload, whereas the power surges in the big cluster when pushed for peak performance. Constraint two is the lack of granularity that forces the system to use the big cluster when only medium performance is required. These two constraints can be largely mitigated by incorporating a Mid, or third, cluster in CPU architecture. With the help of a Mid cluster, the Min cluster is able to provide a lower minimum power and Max cluster can aim for higher peak performance.



- 1 A brand new Mid cluster is introduced
- 2 Min cluster goes lower in minimum power; Max cluster aims for higher peak performance
- 3 Reduced power consumption vs. conventional dual cluster across all performance demand conditions

Figure 4 Tri-Cluster Key Points Briefing (Max.Mid.Min)

Through analysis of various applications on the mobile platform, this whitepaper will illustrate how the Tri-Cluster technologies can enhance the user experience with regard to performance or power efficiency.

Analysis of Recent Mobile Use Cases

Flurry, an App analytics provider which gathers data from its network of more than 450,000 mobile applications installed on more than 1.3 billion devices worldwide, indicates that daily mobile use scenarios falls into four basic categories: (1) Gaming; (2) Social Messaging; (3) Utilities; and, (4) Web Browsing. Based on the data released by Flurry in 2013 and 2014, there was a strong increase in social network and utilities activities.

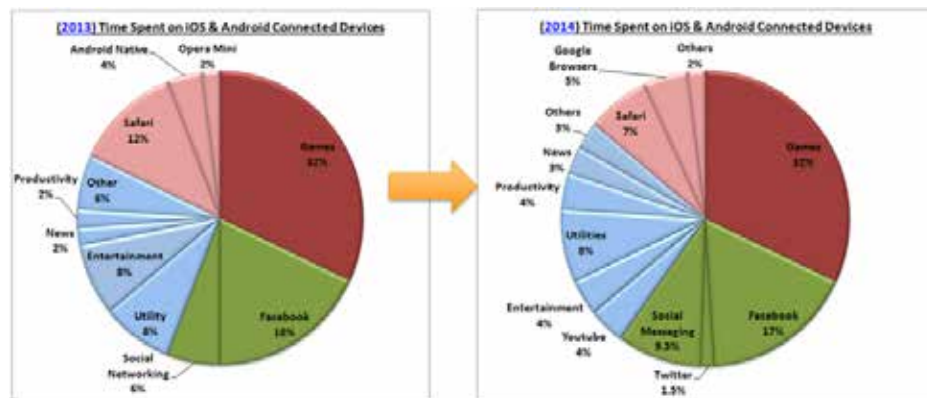


Figure 5 Categories of Mobile Use Cases (Source: Flurry Analytics)

Scenarios	Example Application	Time Spent % (2013) Per Day	Time Spent % (2014) Per Day
1. Gaming	Temple Run 2	32%	32%
2. Social Messaging	Facebook	24%	28%
3. Utilities & Others	Mail, YouTube	24%	26%
4. Web Browsing	Chrome Browser	20%	14%

Table 1 Trend of Mobile Use Cases

MediaTek's Tri-Cluster technology aims to provide the greatest flexibility to fit in with these use scenarios, but also targets social media and utilities scenarios. The following sections would explain how Tri-Cluster architecture can meet these requirements.

Scenario 1 – Gaming

Several experiments were conducted to profile mobile gaming. In terms of time distribution, gaming continuously generate a large amount of loading. Seeing it from an energy perspective, the heavy workload of gaming occupies almost the total power consumption and often requires high and sustainable performance. A good user experience requires a system to provide stable and sufficient FPS (frame-per-second). Accordingly, the most suitable CPU for mobile gaming would be a high performance CPU, the Tri-Cluster's Max CPU or Mid CPU.

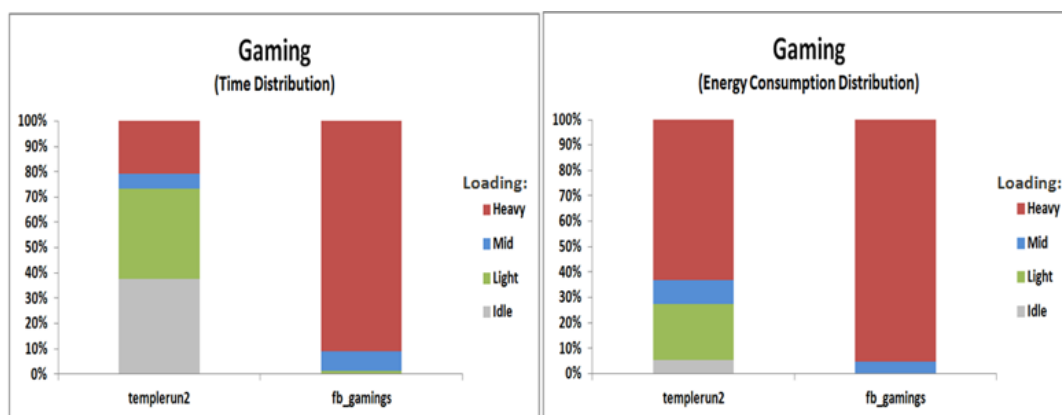


Figure 6 Mobile Gaming Profiling

Scenario 2 – Social Messaging

Most social and messaging apps only require average-low computing power. The low performance CPU could be suitable for social and messaging, Tri-Cluster's Mid CPU or Min CPU.




Figure 7 Social Messaging App Profiling

Scenario 3 – Utilities

Like social and messaging apps, most utilities apps only need average-low computing power. Only a few cases require burst performance. Therefore, the most suitable CPU for social and messaging should be high efficiency CPU (mostly low performance), Tri-Cluster's Min CPU.

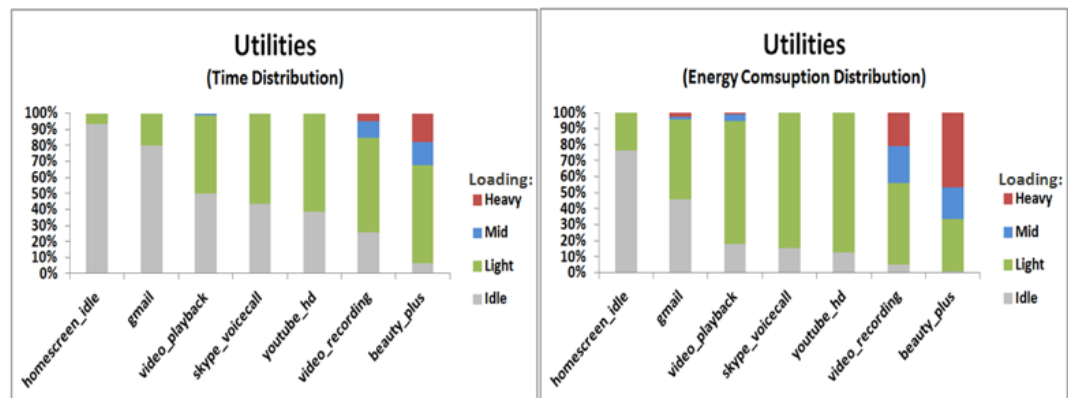


Figure 8 Utilities App Profiling

Scenario 4 – Browsing

Web browsing demands burst computing performance while loading web pages. The most suitable CPU would be the high performance CPU, Tri-Cluster's Max CPU.




Figure 9 Browser Profiling

Scenario Summary

MediaTek's Tri-Cluster is optimized for different user scenarios. The Min CPU minimizes power for light loading but long-lasting activities. The Mid CPU is "just enough" performance for medium loading jobs to avoid wasting energy. And, the Max CPU maximizes performance for high demand activities.

Application Scenarios	CPU Load Requirement	Dual-Cluster CPU Configuration	Tri-Cluster CPU Configuration
1. Gaming	Heavy and Mid mixed	Max CPU only (waste power)	Smart switch btw Max/Mid CPUs
2. Social Messaging	Light and Mid mixed	Min CPU only (slow)	Smart switch btw Mid/Min CPUs
3. Utilities & Others	Mostly Light	Min CPU	Min CPU
4. Web Browsing	Mostly Heavy	Max CPU	Max CPU

Table 2 Summary of Mobile Use Cases

Details of Max.Mid.Min. Technology

According to the analysis of mobile use cases, Tri-Cluster can optimally meet all the use-case scenarios. But, can Dual-Cluster architecture achieve the same performance as Tri-Cluster architecture? The paragraph below explains the differences between Dual-Cluster and Tri-Cluster.

Dual-Cluster offers a good range of power efficiency, but not enough.

- The big CPU is used for middle-performance application.
- The LITTLE CPU is not energy-efficient enough for the extremely light-load tasks of daily operation.

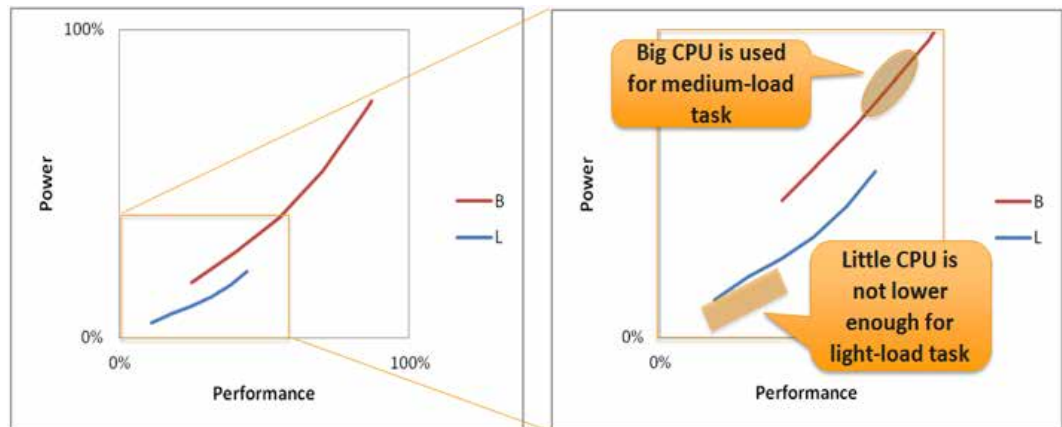


Figure 10 Performance Range of big.LITTLE Architecture.

Tri-Cluster architecture offers a fully-optimized range of power and performance.

- The Min CPU is much more energy-efficient for daily operations.

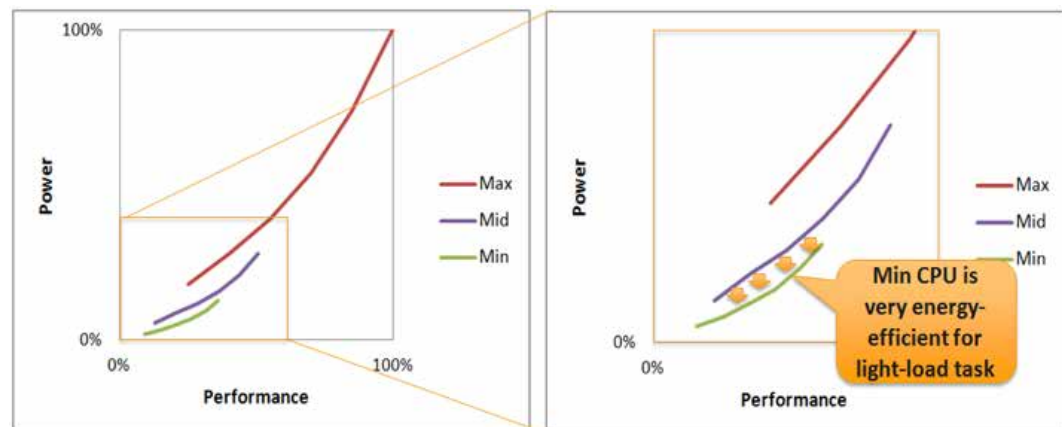


Figure 11 Advantage of Min CPU

- Mid CPU mitigates the power consumption issue of middle-performance requirements. The conventional approach lowers the voltage and frequency of a performance-oriented CPU, such as the big CPU of a Dual-Cluster architecture, when performing medium-load tasks. The Mid CPU saves more power, with enough performance for medium-load tasks as shown in Figure 12.

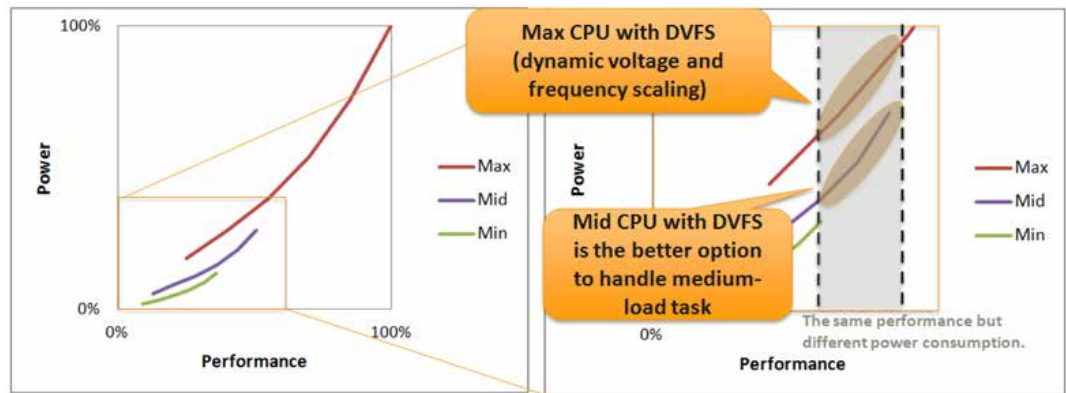


Figure 12 Advantage of Mid CPU

- The Max CPU can aim for higher performance due to the existence of the Mid CPU.

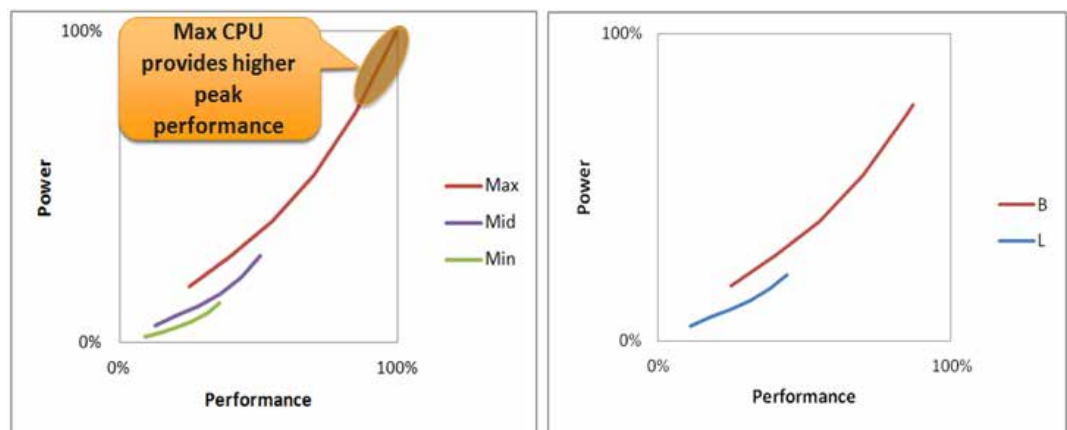


Figure 13 Advantage of the Max CPU

It is clear that Tri-Cluster architecture can not only provide higher peak performance for heavy-load tasks, but also reduce the power consumption for light-load or medium-load tasks.

MediaTek's CorePilot v.3 (Tri-Cluster: Max.Mid.Min.) maximizes the performance and power-saving potential of heterogeneous computing with a host of technologies, including "MediaTek Coherent System Interconnect (MCSI)," "Thermal Management," "Power Management (iCore and Fast DVFS)," and "Scheduler."

The relationships of all of the CorePilot v.3 components are listed below.

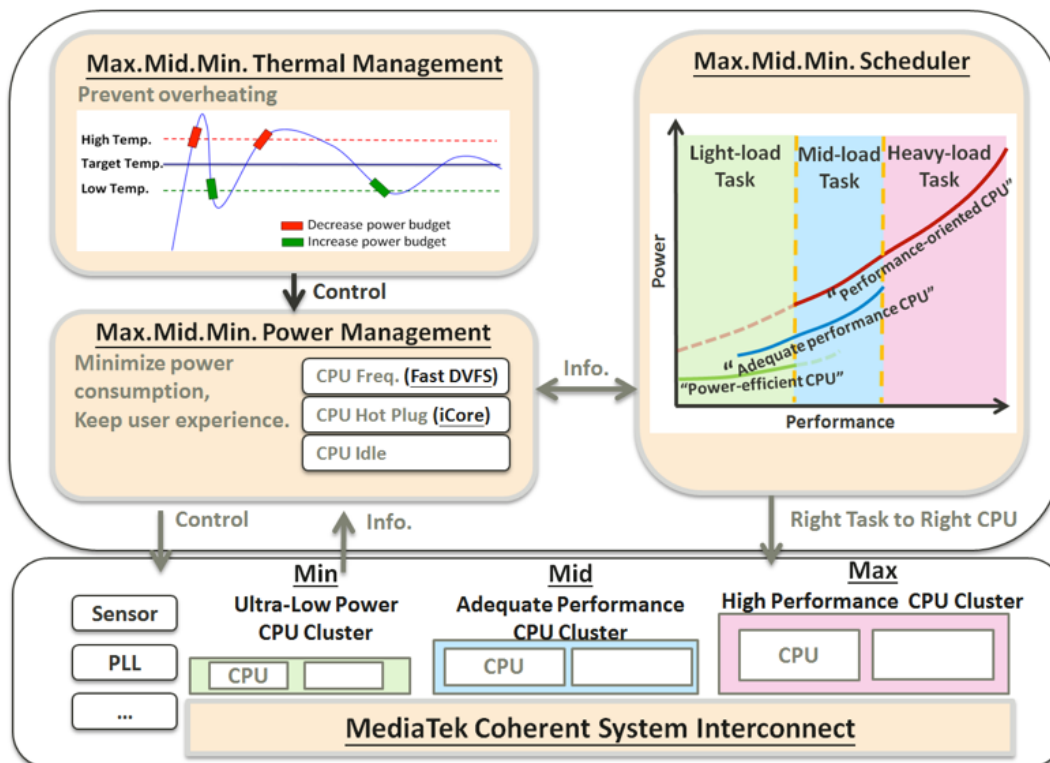


Figure 14 CorePilot v.3 Technology Block Diagram

- **MediaTek Coherent System Interconnect (MCSI)** – MCSI is the fundamental component that makes Tri-Cluster available. With its low latency and high throughput, it enables high performance, fully-hardware-coherent operations and IO coherence maintenance for the architecture.
- **Thermal Management** – This technology applies several mechanisms to calculate allowable power budget for SoC components based on computing requirements, temperature changes, and form factor heat up conditions, and aims to maintain an optimal user experience.
- **Power Management (iCore)** – With the iCore technology, physical CPU0 and cluster0 can be adaptively powered down as needed. This allows the system to respond more accurately to system loadings which will lead to a better user experience and improved power savings.
- **Power Management (Fast DVFS)** – Fast DVFS is a new technique which improves performance and reduces power consumption by increasing the sampling rate (40 times) of DVFS operation with an optimized OPP control algorithm. It also collaborates with scheduler, thermal management, and dynamic power management to optimize the system performance with maximum energy efficiency.

- **Scheduling** – CPUs are powered on or off according to MAX.MID.MIN. Power Management's policy. Once the CPUs are turned on, the scheduler gives software simultaneous access to all the active CPUs in the Tri-Cluster, and tasks are assigned selectively to strategically manage resources and increase power efficiency.

Mediatek Coherent System Interconnect (MCSI)

Mobile device users always demand better user experience and longer battery life. A better user experience requires ever-increasing performance, and longer battery life is only possible when lower power is required to run systems. Heterogeneous multi-processing and computing become crucial to meet both requirements - a high performance processing unit to address short-burst requirement to enhance user experience and an energy-efficient processing unit to address low-power requirement to facilitate longer battery life.

Among heterogeneous processing units and clusters, more and more shared data and memory should be coherent and consistent. System and cache coherence mechanisms are critical to ensure coherency of shared data and memory. In order to eliminate software overhead for cache maintenance (flushing or invalidating cache) by a software-based coherence mechanism, a hardware-based coherence mechanism is developed to reach higher performance and lower system power. Hardware-based coherence mechanisms allow for high performance by enabling data coherence and facilitating fast software thread migration across processing units and across processing clusters. Hardware-based coherence mechanisms also bring further power reduction by reducing external memory access and putting processing units into a low-power state.

Tri-Cluster is an innovative heterogeneous multi-processing and computing architecture which tangibly enhances user experience and prolongs battery life. In order to enable Tri-Cluster architecture, MediaTek developed MediaTek Coherent System Interconnect (MCSI) to support hardware coherency among heterogeneous processing clusters. MCSI uses more than two fully coherent ports, based on AMBA 4 AXI Coherency Extensions (ACE), to enable Tri-Cluster in a mobile platform.

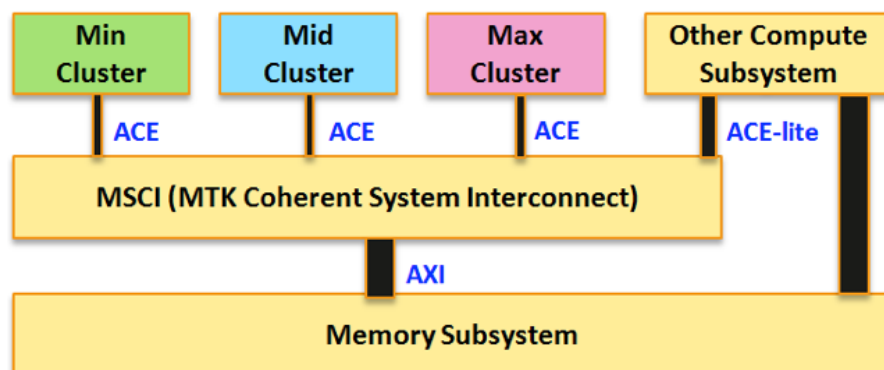


Figure 15 Tri-Cluster+MCSI Block Diagram

MCSI is scalable and supports more than two ACE ports for full coherence among heterogeneous processing clusters. It also supports several ACE-lite ports for IO-coherence between heterogeneous processing clusters and hardware-accelerated engines. With low latency and high throughput, MCSI can fully enable the high performance Tri-Cluster coherence operations and IO coherence maintenance for architecture.

In addition to enabling high performance for Tri-Cluster architecture, MCSI is designed with advanced low-power technology. Compared to the traditional solution, MCSI saves more than half of coherence-interconnect power in most mobile operations. The advanced low-power design of MCSI will extend battery life for normal operations on handheld device by reducing coherent interconnect power. This optimal low-power technology also increases sustainable performance, which is seriously constrained under the thermal envelope of handheld device by reserving more power for processing clusters.

Thermal Management

As mentioned, thermal has become a critical issue for mobile devices and can directly impact the total SoC power budget. There is a strong relationship between thermal and performance because SoC performance is limited by the power budget and the heat dissipation capability is limited by form factor and cost. The question becomes how to maximize performance under thermal constraints, phone skin temperature, and silicon junction temperature (also known as die temperature.)

Two categories of challenges were met during thermal management development:

- For general challenges, thermal management needs to cover silicon process variation and be able to adapt itself to different environments.
- For Tri-cluster challenges, better thermal management is needed to achieve higher performance because the higher performance requirement usually results in more power consumption and higher temperature.

Thermal management technology was developed to deal with these challenges. Thermal management periodically calculates the allowable power budget for SoC components based on computing requirements, temperature changes, and form factor heat-up conditions. With more accurate temperature control, performance is typically increased by 10% for compute intensive applications. In addition, for some specific cases like larger power chip corners, the thermal management prevents temperature from spiking by proactively predicting temperature-rise bursts and limiting the temperature-rise speed.

User experience is an important factor that thermal management takes into consideration. The thermal management smoothly adjusts system performance in order to maintain good user experience when the system starts thermal throttling.

Power Management (iCore)

To be able to reduce system power consumption and at the same time provide the required computing power, the number of CPU cores activated in the system at any time should match the requirement of the target workload. When the system workload rises, the system should be able to quickly turn on additional CPU cores to meet the required computing demand. When the system workload is reduced, the system should be able to quickly turn off unneeded CPU cores and CPU clusters to save power.

In the default Linux kernel, the boot CPU (physical CPU0) of the system cannot be powered down due to its special roles for handling tasks such as interrupt handling and secure world functions. This means that CPU0 and cluster0 must be powered on at all times, which in some cases can lead to suboptimal power utilization and degraded overall system performance.




Figure 16 Traditional Way of Enabling Cluster1

The scenario illustrated in Figure 16 is an example of the workload calling for enabling the 4 CPU cores in cluster1. Ideally, CPU0 and cluster0 should be shutdown to save power. However, if the responsibilities of the physical CPU0 cannot be transferred to the CPUs in cluster1, then it will cause the following drawback: because CPU0 and cluster0 must remain in the powered on state, the system will consume more power than needed.

The performance of the CPUs in cluster1 will be affected due to the cache snooping latency incurred on the cache coherent interconnect by cluster0. To support powering down the physical CPU0 and cluster0, the tasks that are statically assigned to the physical CPU0 must be able to be transferred to the other CPUs in the system.

With the iCore technology, physical CPU0 and cluster0 can be adaptively powered down as needed. The end result is improved system power use and better computation performance.

Power Management (Fast DVFS)

DVFS (Dynamic Voltage and Frequency Scaling) is an effective power management technique where the clock frequency of a process is decreased to allow a corresponding reduction in the supply voltage. Hence, DVFS is widely used on modern mobile devices to help reduce power waste and avoid providing excess voltage and frequency. The conventional DVFS mechanism is controlled by a

software DVFS governor that adjusts DVFS according to system performance, battery condition, and the thermal situation for a better user experience and longer battery life. Figure 17 shows that the low DVFS OPP (Operating Performance Point) can achieve higher energy efficiency.




Figure 17 Lower DVFS OPP Leads Higher Energy Efficiency.

Hence, rapidly following current CPU activity with appropriate DVFS OPP can optimize system performance and power consumption. However, a conventional solution makes the most of DVFS operation by quick following because the penalty becomes considerable when the sampling rate increases for checking CPU activity and then switching CPU OPPs. Fast DVFS technology solves this issue by not only increasing the sampling rate (up to 40 times) but also decreasing the interruption on current tasks for better CPU utilization. There are two major benefits of increasing the sampling rate with Fast DVFS Technology: (1) Rapidly raise DVFS OPP to follow instant performance responsiveness for better user experience and (2) Swiftly lower DVFS OPP for power saving when the current activity passes into light loading.




Figure 18 Advantages of Fast DVFS: (1) Performance Gain (2) Power Saving.

Fast DVFS Technology increases the sampling rate without a considerable switching penalty, which reduces the switching period. Although the rapid sampling rate benefits performance and power saving, inappropriate operation might cause a downgrade of power saving, such as overshooting or switching too often. Hence, Fast DVFS Technology not only increases the sampling rate but also optimizes and precisely controls OPP operation among Tri-Clusters. In addition to maximizing power savings, optimized DVFS OPP switching also helps increase the margin of power budget and uplifts the performance before hitting the thermal wall.

The analysis of power reduction on Fast DVFS Technology shows that more than 10% of the CPU's power can be saved for some daily applications - such as web browsing, video recording, and gaming - compared with conventional DVFS operation. Fast DVFS Technology will become more advantageous if the activity of the CPU dramatically varies often in the application.

Fast DVFS Technology is a new technique that improves performance and reduces power consumption. This technology increases the sampling rate (40 times) of DVFS operation with an optimized OPP control algorithm. It also collaborates well with scheduler, thermal management, and dynamic power management to optimize the system performance with maximum energy efficiency. It provides a better user experience and longer battery life on mobile devices.

Scheduler

Intelligent Core Activation Technology

To fulfill daily-used or sustainable operations, the CPU management and frequency governors will continue to monitor the system workloads and power budgets to dynamically activate/deactivate CPU cores or up-lift and down-lift the frequency of Min & Mid CPU clusters. Meanwhile, the scheduler will arrange the workloads according to the capacity of Min and Mid clusters at the moment. Typical mobile applications such as social or messaging, email, basic gaming, and web browsing, etc., can be satisfied by better performance and energy-efficient trade-offs to prolong battery life.

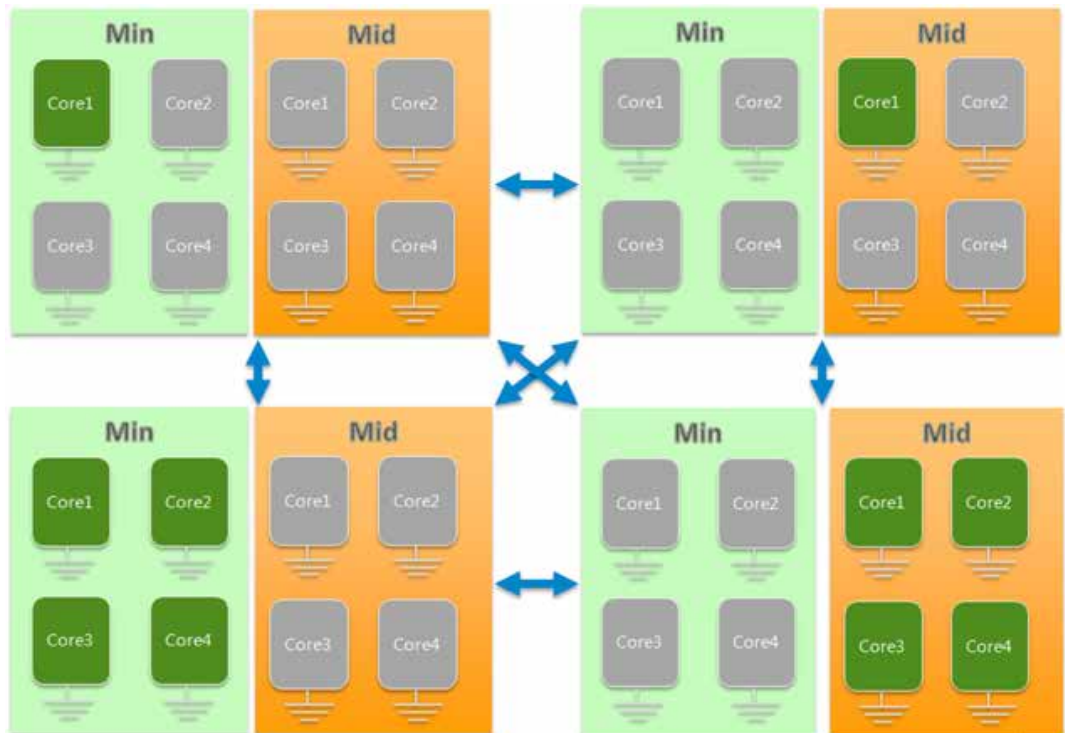


Figure 19 Intelligent Core Activation Technology Example

To meet the high degree of multi-thread or multi-tasking workloads, the system can run at a higher frequency of Octa-cores of Min & Mid clusters.

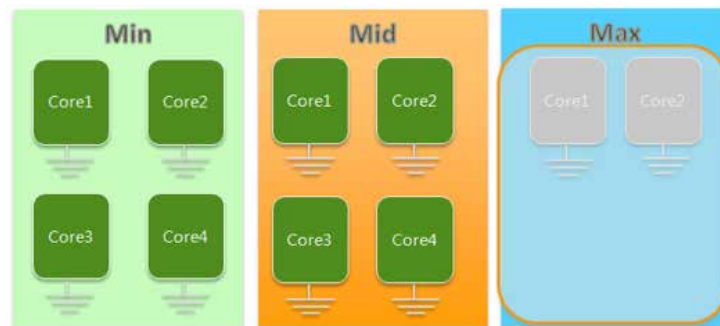


Figure 20 System Runs at Higher Frequency of CPUs of Min&Mid Clusters.

To deliver the instantaneous peak performance demanded by the computing-intensive or time-responsive, urgent tasks, the CPU management governor cooperates with the scheduler to activate the Max cores and assign these tasks to the Max cluster immediately, so that the computing demands can be satisfied to improve overall system responsiveness.

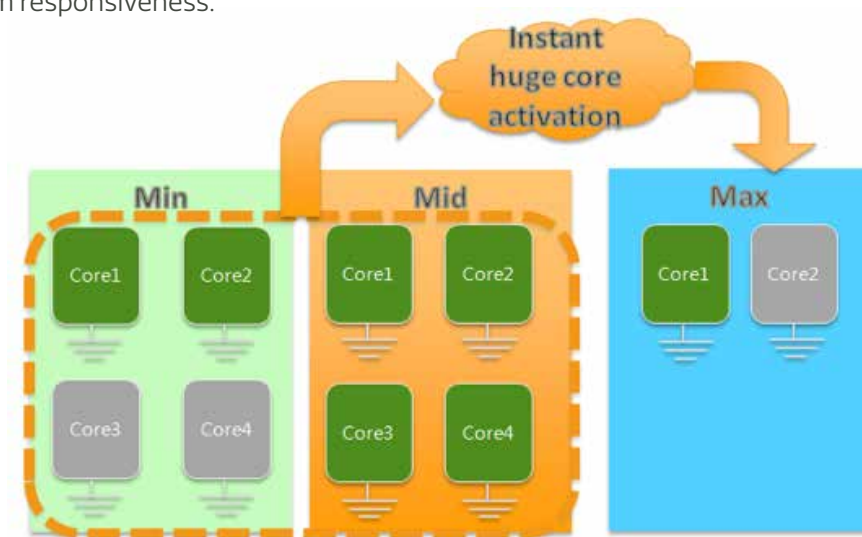


Figure 21 Instant Max CPUs Activation Example

Advanced Global Task Scheduling (GTS)

The advanced GTS provides superior responsiveness, performance, and power trade-offs through the following features:

- **Energy efficiency** – Dynamically adjusts workload migration thresholds between Min, Mid and Max clusters to react to diverse computation demands, so that each cluster can achieve energy-efficiency.
- **Inter-cluster cache coherence** - To effectively increase cache locality, the inter-cluster task assignment and migration obeys thread-group migration discipline, further reducing inter-cluster communication overhead.

- **Responsiveness** - Adequately separates light and heavy workloads between Min, Mid and Max clusters to improve overall system responsiveness.
- **Schedule inter-operations** – The default OS scheduler makes CFS-class and RT-class decision independently. In Tri-Cluster architecture, it incurs an unbalanced workload between CPU/Cluster and longer responsive time of tasks affecting the overall system efficiency. The impacts can be greatly reduced through enhanced inter-operations between CFS- and RT-class embedded in advanced GTS of Tri-Cluster scheduler.

Benefits of Max.Mid.Min. Technology

Combined with a state-of-the-art, energy-efficient scheduler, the left graph of Figure 22 depicts the optimized energy-efficiency curve that achieves the lowest power under specific performance requirement for conventional big.LITTLE architecture. The right side of Figure 22 shows the optimized energy-efficiency curve for advanced architecture.




Figure 22 big.LITTLE and Optimized Performance-Power Curve

Compared to the big.LITTLE optimized curve, the energy-efficiency-optimized curve could provide much lower power in a specific performance range. Figure 23 depicts the power ratio from big.LITTLE architecture to architecture over a wide-range of performance requirement. It demonstrates up to 51% CPU power savings for light-loading operations, and up to 28% CPU power reduction for medium-loading operations.




Figure 23 Power Comparison Between big.LITTLE and Max.Mid.Min.

Energy Savings for Mobile Use Cases

Comparing varying mobile use cases, Tri-Cluster architecture can provide substantial power savings. Compared to big.LITTLE architecture, Tri-Cluster technology could bring a 31% CPU energy reduction in mobile gaming use cases, and a 35%, 38% and 15% energy savings in social media and message, utilities and others, and browsing scenarios, respectively (under the same performance requirement). Considering the mobile user time-spent percentage over these four categories, the architecture contributes to an average of 30% CPU energy savings.




Figure 24 Power Comparison Over Varying Mobile Use Cases

Performance Uplift for Single Thread Applications

Along with Mediatek's advanced technology on Max CPU implementation, this technology could uplift peak CPU performance by 15% over conventional big CPU implementation for heavy loading and short-burst operations, enhancing user experience by offering faster responsiveness. The technology also enhances sustainable CPU performance by up to 12% for the medium loading and sustained scenarios, like gaming, when these use cases are constrained by the thermal envelope of handheld devices and CPU operations are confined by a limited power budget. Owing to much higher energy efficiency in lower CPU performance requirements, the technology could improve sustainable performance by up to 33% if CPU operations are severely constrained by a very limited CPU power budget.




Figure 25 Peak & Sustainable Performance Comparison for Single-Thread App

The state-of-the-art technology not only reduces CPU energy by 30% for varying mobile use cases which substantially extends battery life, but also uplifts CPU peak performance by 15% to tangibly shorten response time for interactive and burst operations. Plus, it enhances sustainable CPU performance by up to 12% for long-run operations.

Sustained Performance Uplift for Multi-Thread Applications

For highly parallel multi-threaded applications, the Tri-Cluster mobile system could offer higher multi-threaded computing energy efficiency and substantially uplift-sustained CPU performance. This is seriously constrained under the thermal envelope of handheld devices. Compared to the traditional 4+4 B.L. solution, the advanced 2+4+4 Max.Mid.Min. solution could increase the sustained CPU performance of highly parallel multi-threaded applications from +16% to +37%, which vary over the power

budget reserved for the CPU. For example, a +18% sustained performance increase can be seen if a 3.5W power budget is reserved for CPU, and a +32% sustained performance increase is possible if a 2W power budget is reserved for the CPU.

Summary

In July 2013, MediaTek delivered the industry's first mobile system on a Dual-Cluster chip with Heterogeneous Multi-Processing (CorePilot v.1 technology).

In 2013, CorePilot was improved to pilot GPU-favored tasks to the GPU. The technology was based on the CorePilot v.1, which can efficiently pilot the tasks to the suitable performance-oriented CPU or the energy-efficient CPU. Moreover, for some throughput-oriented tasks, CorePilot v.2 was capable of automatically executing tasks on CPUs and GPUs simultaneously.

In 2015, thermal has become the dominant issue for mobile devices because of the increased computing power. In addition, mobile use cases for ordinary people are diverse and varied. MediaTek has discovered these issues and further proposed CorePilot v.3 technology (Tri-Cluster "Max.Mid.Min") to maximize power efficiency with extreme computing performance.

- **Ultra-Low Power (Min)** - Applies advanced power technologies to achieve ultra-low power and extend battery life for daily operation.
- **Adequate Performance (Mid)** - Because of Min's ultra-low power, Mid can aim for an adequate performance target without sacrificing energy efficiency for sustainable operations.
- **High Peak Performance (Max)** - Because of Mid's adequate performance, Max can aim for a higher performance target without sacrificing energy efficiency to enhance short-burst user experience.

