AMD FPGA Advantages over Competing Legacy LUT4 Architectures

WP558 (v1.0) September 24, 2024 **White Paper**

Abstract

Optimizing FPGA performance and power consumption while minimizing costs remains a key challenge for FPGA designers. AMD cost-optimized FPGAs provide a balanced solution that delivers high performance and low power consumption. The LUT6 architecture enables utilization advantage that helps customers choose smaller and lower cost devices. The utilization advantage combined with thermally efficient packaging enables both static and dynamic power savings.

AMD Adaptive Computing is creating an environment where employees, customers, and partners feel welcome and included. To that end, we're removing noninclusive language from our products and related collateral. We've launched an internal initiative to remove language that could exclude people or reinforce historical biases, including terms embedded in our software and IPs. You may still find examples of non-inclusive language in our older products as we work to make these changes and align with evolving industry standards. Follow this link for more information.

AMD Cost-Optimized Proven Track Record

AMD has a strong proven track record of offering cost-optimized devices, going back across multiple generations. Various vendors are trying to enter this market, and it is important for designers to understand the trade-offs between architectures. This document provides a data-driven comparison so customers can make a well-informed decision.

Low-Density Applications Priorities: Cost, Power, Size

Many applications that require sub-200 kLC, referred to as "low-density offerings" in this context, usually prioritize three main factors: cost, power, and size. Despite the perceived plethora of options, most competing devices in this space incorporate outdated legacy architecture, with legacy interfaces, to provide a low-cost offering.

Other Considerations

Deciding which FPGA to choose for a given application typically involves consideration of many factors including silicon capabilities, software design environment, tools, and field support.

AMD provides access to the latest technology for:

- Programmable logic
- Interfaces
- Packaging
- Design Tools and IP

6-Input LUTs: Optimized for Performance

Since the introduction of FPGAs, a variety of programmable logic architectures have been assessed, starting with the foundation of LUT architectures. Generally:

- 4-input look-up tables (LUTs) are optimized for area utilization
- 6-input LUTs are optimized for performance

This white paper highlights the advantages of leveraging innovations from the high-end, the utilization advantages of 6-input LUT architecture, as well as performance and power advantages for low-end applications.

LUT6 Silicon Utilization Advantage

Xilinx (now AMD) created the first FPGA in 1984, containing two 3-input LUTs and one register. Shortly after, it was succeeded by the 4-input LUT. While various FPGA vendors have experimented with 5-, 6-, and 7-input LUTs to improve performance, the majority of competitive low-density FPGAs continue to be based on LUT4 architecture.

All AMD FPGAs Introduced in the Past Decade Use LUT6 Architecture

LUT implementations on AMD and Lattice devices, respectively, are shown below.

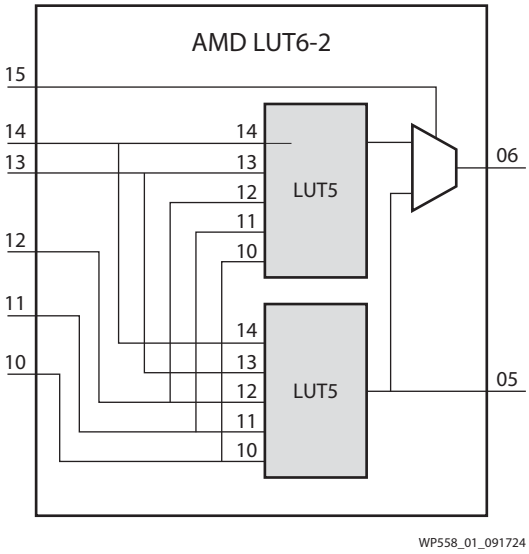

Figure 1: **One AMD LUT6 FPGA Architecture Block Used in UltraScale+ Devices**

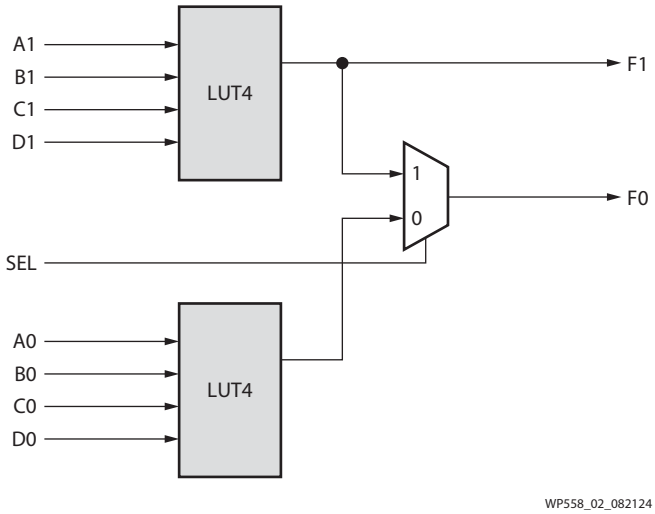

Figure 2: **Two Lattice LUT4 FPGA Architecture Blocks**

Devices, Speed Grades, and Software Tools Used for Analysis

When evaluating the capabilities of a given LUT architecture, one key parameter is utilization, which measures how well a given design can be mapped to these resources.

To evaluate utilization differences between LUT architectures, AMD conducted a comprehensive analysis. 30 Open Core (OC) designs were implemented targeting AMD and equivalent Lattice devices. The following table describes the devices, speed grades, and the software tools used for testing.

Table 1: **AMD and Lattice Devices, Speed Grades, and Tools Used for Testing**

Vendor	AMD		Lattice	
Node	**28 nm**	**16 nm**	**28 nm**	**16 nm**
Family	Artix™ 7 FPGA	Artix UltraScale+™ FPGA	MachXO5™-NX FPGA	Avant™-E FPGA
Device	7A100T	AU7P	LFMXO5-25	LAV-AT-E70
Speed Grade	-1L	-2, -1, -1LV	9HP, 7LP, 9LP, 7HP	-1, -3
Software Tool	AMD Vivado™ Design Suite		Lattice Radiant Design Software	
Tool Version	Vivado 2024.1		Radiant 2024.1	

Result: LUT6 Uses 40% Fewer LUTs on Average than LUT4 Architecture

The LUT utilization was recorded for each OC design on both the AMD and Lattice tools. Geomeans of the ratios of LUT utilization were taken, resulting in ~40% on average fewer LUTs utilized for AMD devices. As a result, LUT6 architecture can be inferred to use 40% fewer LUTs on average than LUT4 FPGA architecture.

This 40% reduction leads to several advantages. There is a direct impact to design footprint, enabling you to choose a lower density device from AMD, compared to the size of device from Lattice, to perform the same functionality. Choosing a smaller logic density device allows for static power savings, which we cover in a further section. Refer to Figure 3 for a visual representation of the 40% design utilization reduction enabled by LUT6 architecture. In this example, a design that fits into an Avant E 50 could fit into an AMD Spartan™ UltraScale+ SU200P FPGA due to the 40% LUT6 reduction.

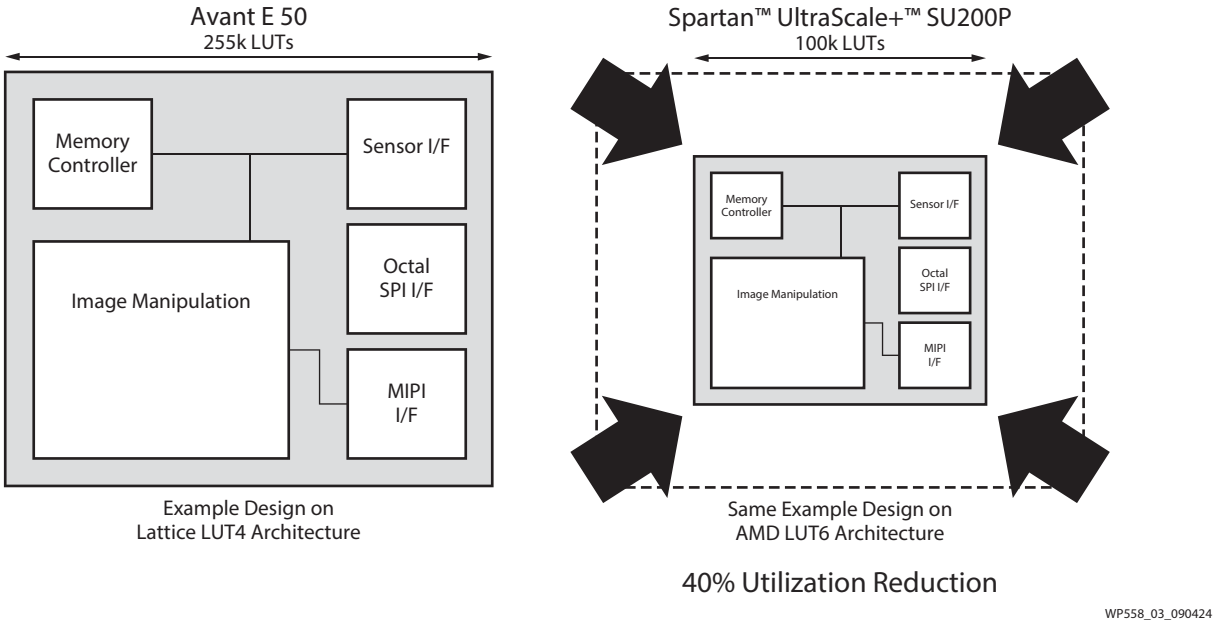

Figure 3: **Visual Representation of the 40% Design Utilization Reduction Enabled by LUT6 Architecture**

Figure 4 highlights the AMD to Lattice device mappings, after accounting for the 40% utilization advantage.

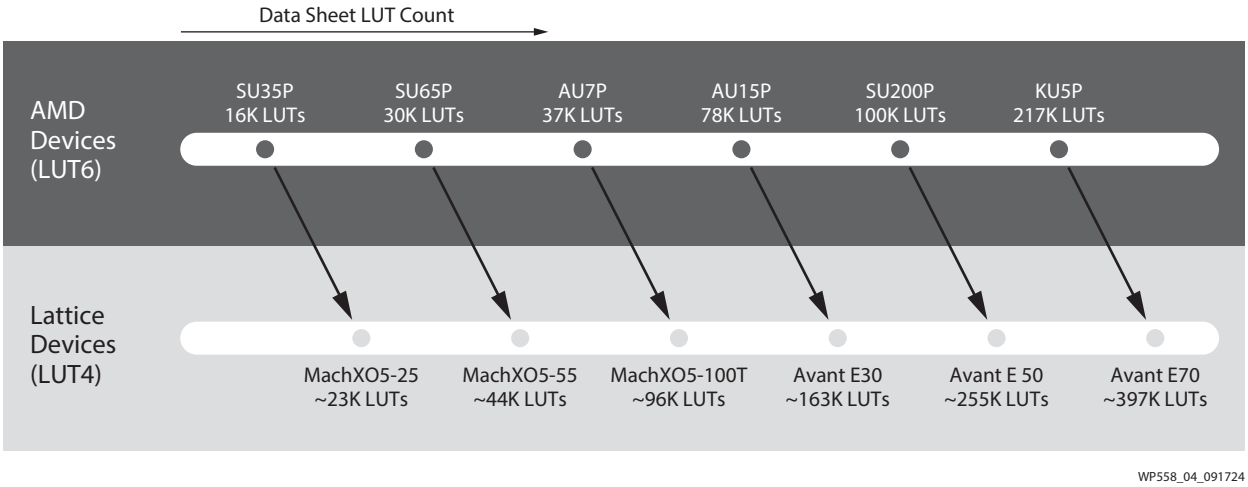

Figure 4: **Competing Lattice Devices Mapped to Smaller AMD Devices**

Superior Performance with LUT6 Architecture

In addition to the LUT6 utilization advantage mentioned above, our benchmarks revealed that AMD UltraScale+ FPGAs offer a 1.8X higher F_{MAX} advantage over Lattice Avant and 2.7X higher F_{MAX} than Lattice Nexus.

The same 30 open core designs were run on the devices outlined above. From there, the F_{MAX} was determined for each device by iteratively running the designs until the device failed to meet timing constraints. The results and geomeans were recorded.

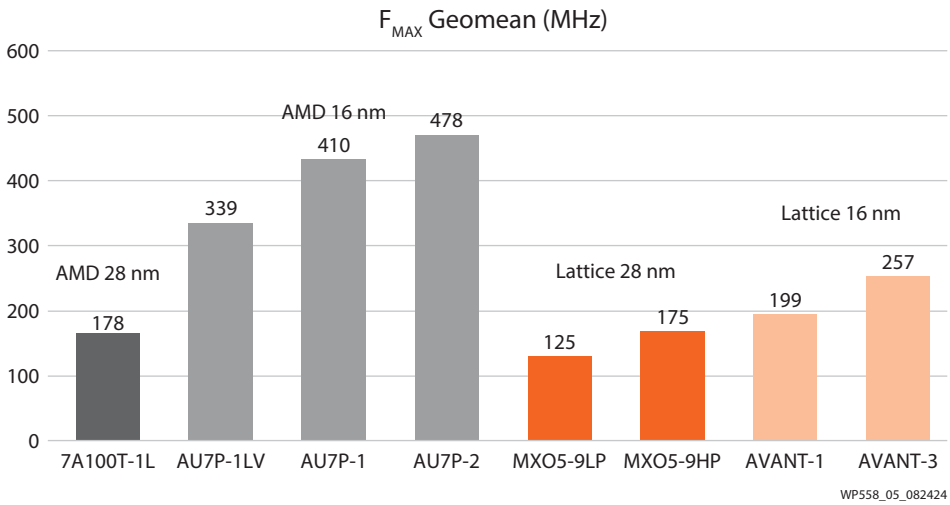

Figure 5: **Geomean F_{MAX} for AMD and Lattice Devices**

Result: AMD is Approximately 2 Speed Grades Faster

Summary of data in Figure 5:

- AMD UltraScale+ -1LV device is 27% faster than Avant -3, making it roughly "2 speed grades faster"
- Lattice Avant fastest (-3) speed grade is 29% faster than their slowest (-1)
  - Suggests that a ~29% F_{MAX} delta roughly represents "2 speed grades apart" for Lattice Avant

In Figure 5, taking a ratio of the AMD 16 nm highest speed grade offering (AU7P -2, 507 MHz) vs. the Lattice 16 nm highest speed grade offering (Avant E70 -3, 278 MHz) yields a 1.8X higher F_{MAX}.

Taking a deeper look, Figure 6 illustrates individual F_{MAX} results for each device tested for each OC design run. AMD offers superior performance across almost all designs.

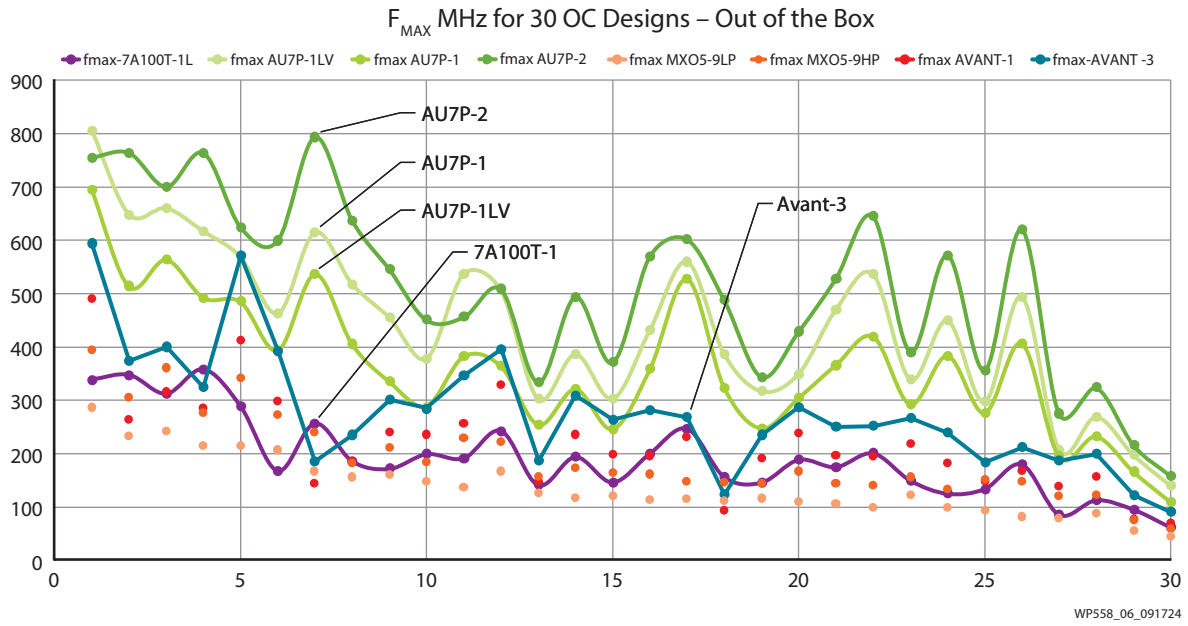

Figure 6: **F_{MAX} Plotted across All Designs Run for AMD and Lattice Devices**

Result: AMD UltraScale+ FPGAs Offer 1.8X Higher F_{MAX} over Lattice Avant and 2.7X Higher F_{MAX} than Lattice Nexus

AMD Multiple Speed and Core Voltage Options

AMD Spartan and Artix UltraScale+ devices are offered in multiple speed and core voltage options. Table 2 provides a comparison of clock network F_{MAX} based on the publicly available data sheets.

Table 2: **AMD vs. Lattice FPGAs Clock F_{MAX} as Available in the Public Data Sheets**

Core Voltage (V)	NX Speed Grades (Mach, Certus, CertusPro)	NX F_{MAX} (MHz)	Avant Speed Grades	Avant F_{MAX} (MHz)	Artix 7 Speed Grades	Artix 7 F_{MAX} (MHz)	UltraScale+ Speed Grades	UltraScale+ F_{MAX} (MHz)
0.72							-1LV -2LV	667 725
0.82			-1 -2 -3	375 500 625				
0.85							-1 -2	667 775
0.9					-2L	394	-3	891
0.95					-1L	464		
1.0	-7 -8 -9	276 325 400			-1 -2 -3	628		

Result: With a Two Speed-Grade Advantage, You Can Do More in a Lower Density Part

Summary of Table 2:

- The AMD UltraScale+ -1LV (V_{LOW}, 0.72V) devices can match the performance of the Lattice Avant -3 speed grade devices running at a higher core voltage of 0.82V
- This lower-voltage operation enables static & dynamic power reductions
  - This is often overlooked by our competitors with respect to power-performance analysis, often positioning AMD V_{NOM} (0.85V) devices against theirs—not a true comparison since AMD offers a lower voltage option that can match or beat the Lattice V_{NOM} performance. This demonstrates that AMD -1LV parts can compete on performance with the Lattice portfolio.

Do More with Less Logic

The AMD Cost Optimized FPGA Portfolio offers flexible FPGA architecture that enables your designs to be run between 150 to 500+ MHz to fit your application's needs. If a slower Lattice device does not meet timing requirements, you are forced to migrate to a faster Lattice device, negating any potential power savings associated with choosing the slower device. This is not the case with AMD. With a two speed-grade advantage, you can do more with less logic. Choose a smaller LUT6-enabled device to perform the same functionality as a larger device built on LUT4.

Running a design faster can enable one of two potential scenarios:

1. Leveraging that performance to use fewer resources, thereby shrinking the overall design and choosing a smaller device, resulting in total power and BOM cost savings

—OR—

2. Keeping the device size as-is and opting to leverage the performance for future scaling, enabling further longevity

As seen from the data presented in this section, leveraging architecture innovations from the high-end portfolio can be incredibly beneficial at the low-end. Naturally, this leads to the question on power consumption and thermals, which will be analyzed in the following sections.

Packaging Thermal Advantages

When considering the power consumption of a device, you need to evaluate two areas:

- Dynamic power: Directly derived based on the resources used, clock rates, and toggle rates
- Static power: Independent of resource utilization and instead scales with the junction temperature (T_j) of the device

To minimize system cost, cost-sensitive products are often deployed with a passive thermal solution. This makes the thermal performance of the package critical in determining static power consumption.

Evaluating Thermal Performance

To evaluate the thermal performance of a package, designers must consider three primary thermal dissipation paths:

- θja - Coefficient of Thermal Resistance from junction to ambient
- θjb - Coefficient of Thermal Resistance from junction to the board
- θjc - Coefficient of Thermal Resistance from junction to case or heat sink (if applicable)

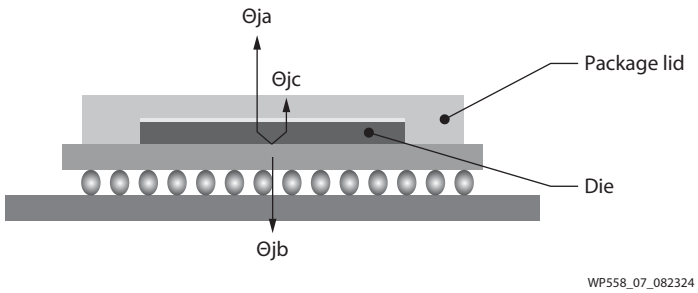

Figure 7: **Thermal Dissipation Paths on a Lidded FPGA Device**

A higher coefficient of thermal resistance means less heat transfer through the material, which translates to a higher silicon temperature for the same load and thermal solution. This ultimately leads to higher

static power as described above. For low-cost applications without a heat sink, designers should consider θja 0 LFM, which measures thermal dissipation to ambient with no air flow.

Based on analysis of public data sheets, user guides, and the Radiant tool, it was found that AMD packages offer up to 34% lower junction to ambient (θja) dissipation on average for still air conditions.[1] This allows for a simpler thermal solution, and lower static power, assuming similar dynamic power is dissipated. This implies that, depending on the environmental conditions and the dynamic power of the design, no thermal solution may be required. This simplified thermal solution then enables the potential for lower total system cost.

Result: AMD Offers Up to 34% Lower Junction to Ambient Dissipation

For the total power analysis, a fixed junction temperature of 100°C (Tj) was used for evaluation. This replicates the standard customer comparison.

Then the "maximum safe ambient" temperature listed in the Radiant tool was noted for the JEDEC θja specification for still air. This maximum safe ambient was applied to the AMD devices along with their θja to the AMD estimation.

Result: Lower θja Means Lower T_j and Lower Static Power

Due to lower θja values in AMD devices for equivalent size, the result was lower Tj and lower static power. These results are shown as the "thermal advantage" in a later analysis. It should also be noted as the dynamic power of a device increases, this AMD packaging results in significantly lower junction temperatures. In some scenarios, the Lattice device's "maximum safe ambient" is listed as a subzero temperature, indicating that a thermal solution or dedicated airflow to the package is required, increasing cost and complexity of the solution.

Enabling Lower Total Power

This section provides a detailed comparison of power consumption for AMD LUT6 vs. Lattice LUT4 devices. These power estimations were performed using each vendors' power estimation tools: Lattice Radiant Software, and Xilinx Power Estimator (XPE) or Power Design Manager (PDM), depending on the AMD device targeted.

Power Consumption Comparison: The Set Up

For device utilization, we assume that the Lattice device will be filled to 80% of available fabric resources. The next step is to leverage the 40% reduction for LUT6 to LUT4, described in the previous section. This 40% reduction is then applied to get an equivalent AMD device. In both cases, these resources are entered into the respective power estimation tools. Note, the 40% reduction was only applied to LUTs, all other resources were kept as-is for comparison.

These assumptions were used for the dynamic power comparison. Often, when performing a comparison between vendors, the same junction temperature is assumed. However, this does not take into account the

1. This analysis was based on JESD51 specifications for packages; actual results will vary.

thermal benefits of the available packages. To consider static power in the comparison, the ambient temperature (TA) of the estimation is set so that the Lattice device junction temperature (T_j) is 100°C, assuming worst case scenario and the typical cost-constrained deployment.

That same ambient temperature is applied to XPE or PDM depending on the device targeted, along with the package θja to estimate the junction temperature (T_j). The junction temperature is noted as that leads to lower static power.

Figure 8 outlines the total power measured at 100 MHz for MachXO5-NX LFMXO5-25, Spartan UltraScale+ SU35P, and Artix 7 XA25T—for various speed grades. Figure 9 includes total power at 150 MHz for Artix UltraScale+ AU7P, CertusPro-NX-50, MachXO5-100, Artix 7 50T, and Artix 7 75T.

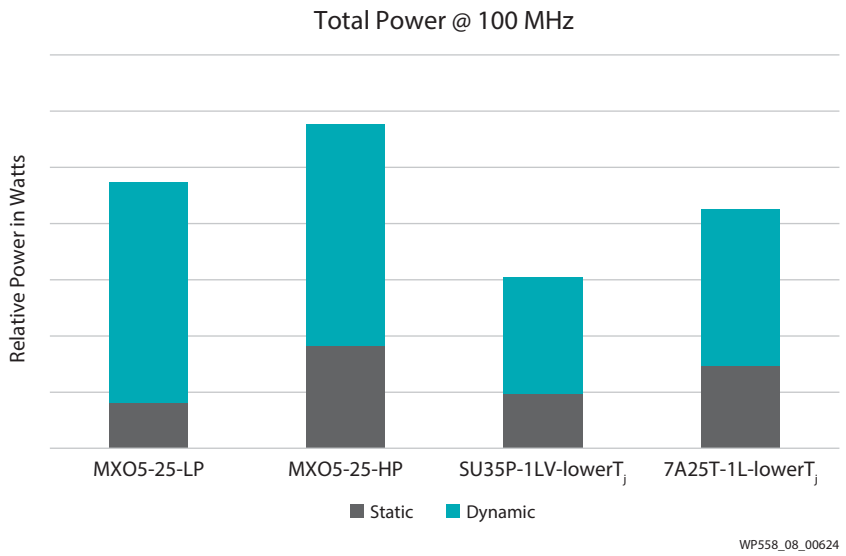

Figure 8: **Total Power Consumption in Watts: MachXO5-25, Spartan UltraScale+ SU35P, and Artix 7 25T**

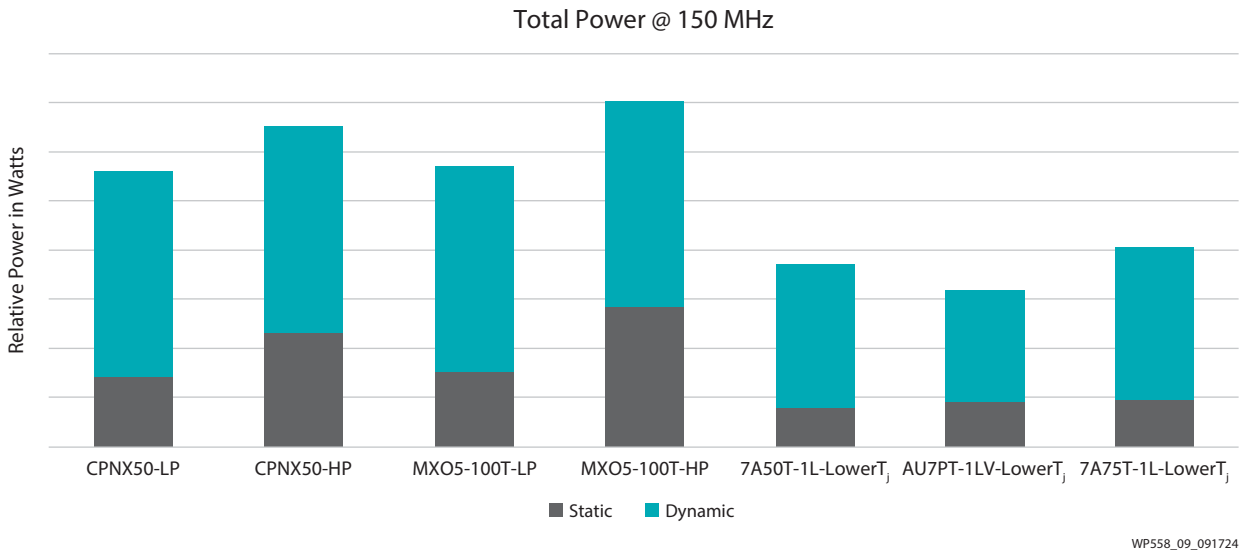

Figure 9: **Total Power Consumption, in Watts: Artix 7 7A50T, Artix 7 7A75T, Artix UltraScale+ AU7P, CertusPro-NX-50, and MachXO5-100**

Result: AMD UltraScale+ FPGAs Beat Lattice on Total Power across Nexus and Avant Platforms

Spartan UltraScale+ FPGA power numbers are an estimate based on UltraScale+ architecture, and preliminary package data. Additionally, the 40% LUT6 advantage is included for LUT only, all other resources are the same (block RAM, DSP, HDIO, etc.).

In Figure 10,[1] for high-performance designs, and in Figure 11,[1] for low-power designs, we can see the power advantage (shown as percentages) for each device: SU35P, SU50P, and SU100P.

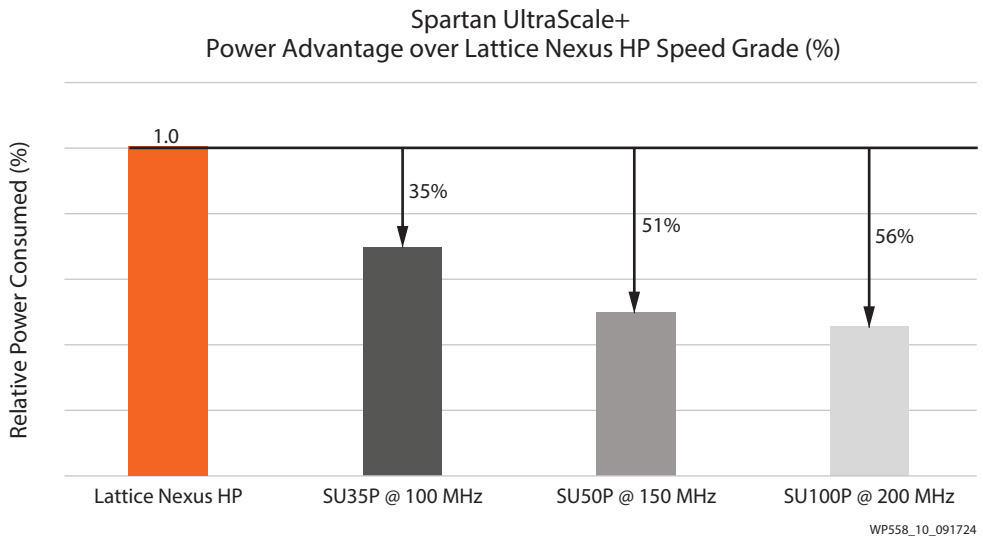

Figure 10: **AMD Spartan UltraScale+ FPGA Power Advantage over Lattice Nexus HP as a Percentage**

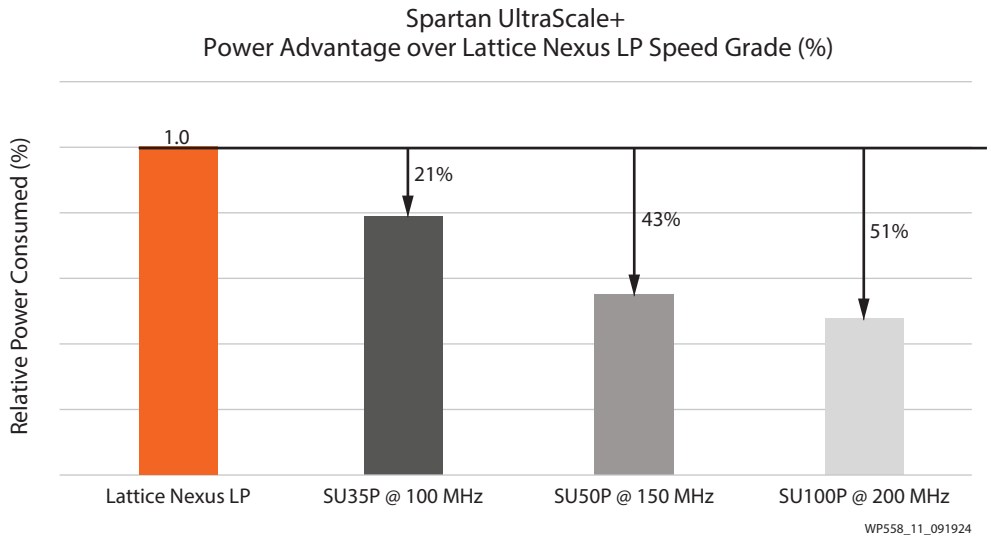

Figure 11: **AMD Spartan UltraScale+ FPGA Power Advantage over Lattice Nexus LP as a Percentage**

1. Lattice Nexus HP & LP include three devices: Mach XO5-NX 25, CertusPro-NX 100, and Mach XO5-NX 100T. All three are normalized to 100%.

The geomean of these results was then taken, resulting in Spartan UltraScale+ FPGAs offering up to 36% lower total power consumption than Lattice 28 nm Nexus platform. Spartan UltraScale+ FPGAs offer up to 46% lower total power consumption than Lattice Nexus FPGAs for high-performance designs.

In Figure 12, the AMD Kintex™ UltraScale+ KU5P FPGA static power is significantly lower than the Lattice Avant E70. This is due to the Radiant tool requiring a subzero ambient temperature (–18°C) to keep the Avant E70 under 100°C with still air dissipation. Consequently, operating the device above 200 MHz, the Lattice Avant E70 requires a heatsink at regular ambient temperature. Conversely, the AMD device (KU5P) offers lower dynamic power and superior package thermal characteristics, resulting in a much lower junction temperature and therefore lower static power.

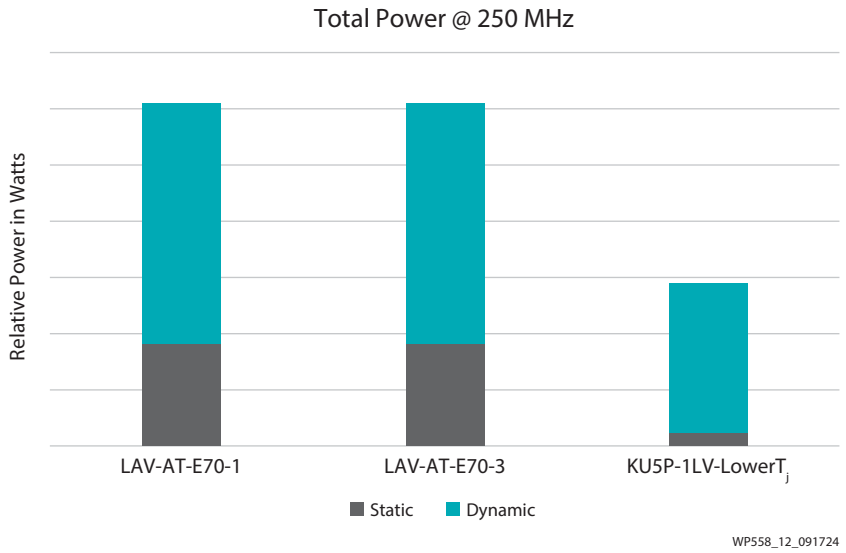

Figure 12: **Total Power Consumption in Watts: Lattice Avant E70, AMD Kintex UltraScale+ KU5P**

Power Consumption Comparison: The Summary

The above figures depicting total power consumption in watts outline how the AMD Cost-Optimized Portfolio can offer lower total power against the Lattice portfolio of Nexus and Avant. The AMD -1L parts can compete on both power and performance with the low power (LP) and high performance (HP) speed grades of the Lattice product line.

The dynamic power savings associated with choosing a smaller device, and the static power savings associated with more thermally efficient packages, result is lower total power than competitive offerings. AMD cost-optimized FPGAs offer higher performance within the same power envelope, or lower power at the same F_{MAX}.

By taking advantage of LUT6 enabled resource utilization reduction, you can select a smaller device, run it as fast as you need to, and run it cooler with thermally efficient packaging. All these factors combined result in a solution that meets your power envelope criteria. By choosing AMD, there is no need to compromise—on performance, or on power.

Conclusion

AMD cost-optimized devices provide a balanced solution that delivers high performance while minimizing power consumption. By using smaller devices with lower core voltages and thermally efficient packaging, you can achieve both dynamic and static power savings. This means you can enjoy higher performance within the same power envelope or lower power usage at the same maximum frequency—all without compromising on either performance or power efficiency. Additionally, these architecture-enabled benefits have the potential for lowering total BOM cost.

Lattice claims to "innovate" explicitly for the low end. This translates to leveraging legacy FPGA architecture, coupled with legacy interfaces, resulting in minimally viable products that are subject to becoming outdated. The AMD Cost-Optimized Portfolio offers industry leading silicon, with modern interfaces, resulting in future proof devices that maximize lifetime longevity.

If you are interested in running the numbers yourself, please check out the COP Competitive Lounge. For additional information, please reach out to AMD directly.

Finding Additional Documentation

Technical Information Portal

The AMD Technical Information Portal is an online tool that provides robust search and navigation for documentation using your web browser. To access the Technical Information Portal, go to https://docs.amd.com.

Documentation Navigator

Documentation Navigator (DocNav) is an installed tool that provides access to AMD Adaptive Computing documents, videos, and support resources, which you can filter and search to find information. To open DocNav:

- From the AMD Vivado™ IDE, select **Help > Documentation and Tutorials**.
- On Windows, click the **Start** button and select **Xilinx Design Tools > DocNav**.
- At the Linux command prompt, enter `docnav`.

Note: For more information on DocNav, refer to the *Documentation Navigator User Guide* (UG968).

Design Hubs

AMD Design Hubs provide links to documentation organized by design tasks and other topics, which you can use to learn key concepts and address frequently asked questions. To access the Design Hubs:

- In DocNav, click the **Design Hubs View** tab.
- Go to the Design Hubs web page.

Revision History

The following table shows the revision history for this document.

Section	Revision Summary
09/24/2024 Version 1.0	
General updates	Initial AMD release.

Please Read: Important Legal Notices

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AUTOMOTIVE APPLICATIONS DISCLAIMER

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.

© Copyright 2024 Advanced Micro Devices, Inc. AMD, the AMD Arrow logo, Artix 7, Kintex, Spartan, UltraScale+, Vivado, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.