



Designing Heatsinks and Thermal Solutions for Xilinx Devices

XAPP1377 (v1.0) June 6, 2022

Summary

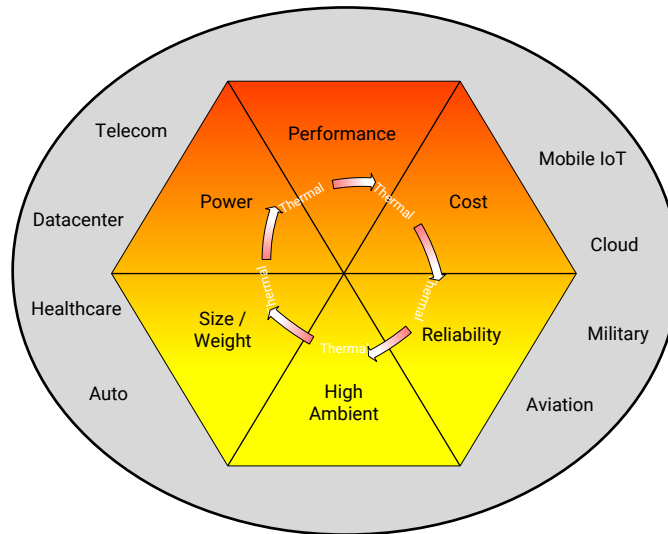
This application note describes guidelines and best practices for designing heatsinks and thermal solutions for Xilinx® devices. It includes upfront considerations for selection, design, and test to adequately plan and execute a successful thermal solution for all Xilinx devices.

Introduction

The overriding requirement for any thermal solution is to maintain the device temperature within operating specifications. However, for most applications, there are several other important considerations such as cost, area, weight, acoustical noise, reliability, ability to be manufactured, and many more. An optimal thermal solution balances all these requirements. However, in most applications these requirements are different, and some of these considerations could be more important than others. Other unique parameters of any given application are the operating environment and power consumption. Differences in local ambient temperature seen by the device, airflow variation, case and chassis size, orientation and configuration, and other parameters are often unique to a given design, as is the operational power depending on how the device is used. While devising a thermal solution, you should be mindful of all of these considerations and how they apply to the specific design to ensure the solution not only meets the thermal requirements, but also does so in a way that meets all of the goals and specifications for the system. The following figure demonstrates how the thermal design must balance many goals for any application.

Xilinx is creating an environment where employees, customers, and partners feel welcome and included. To that end, we're removing non-inclusive 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.

Figure 1: Thermal Design Balancing Example



X26621-050322

Steps to Designing a Thermal System

1. Identify the package lid type and associated documentation.
2. Determine the contact area for the package and height requirements.
3. Design the heatsink base.
4. Design fins and airflow (or fluid flow if liquid cooling).
5. Design the proper attachment.
6. Choose thermal interface material (TIM) and determine application.
 - a. TIM is the substance used between the device and thermal stack up, and is intended to improve thermal contact.
7. Determine thermal parameters for simulation and validate thermal solution.
8. Determine proper assembly, test, and debug (if necessary).

Choosing and Determining Package Lid Type

Xilinx offers different packaging types for different devices, and each have benefits and trade-offs. It is important to understand what lid type is offered for the target device during both the selection and the design process. Some devices offer different types of lids, so it is important to both select the proper lid type for the system goals, and to know that the lid type could dictate a different thermal design strategy.

Lidded vs. Lidless Devices

The first-order classification is whether or not the package has a lid. In general, lidless devices offer better thermal performance due to the fact they are not adding additional junction-to-case thermal resistance in the form of a TIM1 inside the package and the lid itself. TIM1 is the material applied by Xilinx within the lidded device package to thermally couple the die to the device lid. Lidless devices allow the thermal solution to get as close as possible to the source of the heat, which is the die. On the other hand, lidded devices can be easier to design, as they are the more traditional approach to IC packaging, and off-the-shelf thermal solutions can be easier to find.

From a reliability standpoint, lidded devices do offer more protection to the die while lidless devices with stiffener ring offer higher rigidity. This allows it to have higher board level reliability (BLR) characteristics. BLR refers to factors that can impact the device adhesion and signal contact to the board such as shock, drop, and vibration. It also can encompass factors like board bend, flex, and deflection that can result in partial or full separation of the device from the board.


There is also a common misconception that lidded devices offer protection from contaminants entering the package body and touching the die, inner package substrate, or packaging capacitors while lidless does not. While it is undeniable that lidded devices do offer more protection from this occurrence, the fact that lidded parts have vent holes to allow off gassing and moisture evaporation means that lidded packages cannot be considered sealed bodies, and similar protections against such contamination can occur. While there are certainly differences at the macro level for the reliability of lidded versus lidless, overall both are highly reliable package types with similar considerations for designing thermal systems.

Lidded Devices

Wire Bond Packages

Wire bond packages are generally more cost-effective and thus found on the low-cost Xilinx devices. The packages are durable, but from a thermal perspective, are not the highest performing and thus not recommended for high-power and/or high ambient environment operation. The IC is often encapsulated in a plastic mold which can impact the thermal resistance to the case. This is fine for lower power operation, but could pose challenges in higher power and higher operating temperature environments.

Figure 2: Cross Section of Wire Bond Package



X26618-050322


Lidded Flip Chip

The most common package used in Xilinx devices for over 10 years is lidded flip chip packages. Lidded flip chip packages offer very good power and signal integrity, and improved junction to board and junction-to-case thermal resistance as compared to wire bond packages. The lids are generally constructed from a nickel plated copper material. Xilinx uses a TIM1 material, which is thermal interface material applied by Xilinx within the lidded device package to thermally couple the die to the device lid. This material ensures good coverage as well as good thermal conductivity to further improve the junction to case thermal performances. However, the TIM 1 applied offers optimal thermal performance across the wide range of applications and industries. It must also be able to withstand the high temperatures of solder re-flow. This means while it is optimal, it might not perfectly meet a customer's requirements for High power or High ambient applications. Lidless packages allow users to select the TIM based on their requirements and it is also applied after the solder re-flow process. There are two types of lids used in Xilinx lidded flip chips, and it is good to become familiar with this for mechanical fit and to evaluate thermal properties.

Stamped Flip Chip Lid

Stamped flip chip lids are identified as a trapezoidal profile that is flat above the die surface and tapers towards the package substrate body edges. The heatsink surface generally only touches the flat surface at the top of the package. The following figure shows the cross section and top view of a stamped lid, also known as a type II lid.


Figure 3: Cross Section and Top View of a Stamped Lid



X26617-050322

Forged Flip Chip Lid

Forged flip chip lids are identified as having a rectangular profile where the top of the package extends to the full package body dimensions. From a thermal perspective, forged lids are generally superior to stamped lids, particularly in reduced conduction situations due to the fact the lid has more surface area for a heatsink to touch for a given package size. The lid itself is generally thicker leading to more heat spreading within the lid prior to heat sink extraction. The following figure shows the cross section and top view of a forged lid, also known as a type I lid.



X26616-050322

Lidless Devices


The most popular packaging option in more recent Xilinx devices is lidless packaging. There are two primary reasons for this. As previously mentioned, lidless devices generally have superior thermal performance as compared to lidded devices, so lidless devices are often preferred in addressing challenging thermal situations. Also, particularly for the larger package body devices, large package lids tend to deform over temperature, and overall reliability is challenging. On the other hand, lidless packages with stiffener ring have far fewer issues with overall coplanarity as we move to larger, multi-die packaging options.

As with lidded devices, we can further break down the lidless classification into three packaging types.

Bare Die Lidless

As the name implies, these have exposed die sitting on top of a package substrate without a lid or stiffener ring. This packaging option is generally found on smaller package body devices. It is found to be both cost effective and have very good thermal performance. This is a good choice for thermally demanding small-device applications. The following figure shows the cross section and top view of a bare die lidless package.

Figure 4: Bare Die Lidless Package




X26615-050322

Lidless with Stiffener Ring

As the name implies, an exposed die packaging option, which includes a stiffener ring around the periphery of the package substrate, provides additional rigidity and protection to the device. This packaging option is generally found in mid to large size packages and is a good choice for higher power and/or higher operating temperature environments. The following figure shows the cross section and top view of a lidless with stiffener ring package.


Figure 5: Lidless with Stiffener Ring Package



InFO

InFO is currently the smallest packaging option in the X-Y direction, as well as the Z direction, which refers to height. It is best suited for very area constrained applications that need the smallest packaging option. Thermally, this package option is very similar to bare die packages. However, due to how thin the package is, an additional step to edge bond the component to the board is required. InFO packages also require less asserted pressure by the thermal solution. This could increase cost and make assembly difficult. The following figure shows the cross section and top view of a ZU3P InFO package.


Figure 6: ZU3P InFO Package



Kria SOM with Heat Spreader

The Xilinx Kria™ system on a module (SOM) is a product containing multiple devices on a miniature printed circuit board (PCB) centering on a Xilinx device surrounded by other active and passive components. The SOM greatly simplifies design by including memory, voltage regulation, and other chip-level design requirements into a single, pre-designed module. Since the SOM consists of multiple chips, there are multiple points of cooling needed on a SOM. Xilinx simplifies this by including a pre-assembled heat spreader to the SOM. This means that the only requirement is to attach a heatsink (if necessary) to the heat spreader to control the temperature of the SOM.

Figure 7: Kria SOM with Pre-Assembled Heat Spreader



Early Evaluation of Device/Package Thermal Properties

For device families prior to Versal® ACAPs, the JEDEC junction-to-board (θ_{JB}) and junction-to-case (θ_{JC}) resistance values are published in the appropriate packaging and pinouts user guide. They are generated by using the JEDEC specification boundary conditions, a small two signal, two power plane (2S2P) boards with mounted cold plates to represent close to infinite cooling. This can be useful at the very early stages of comparing packages within a product family or from multiple vendors, as it allows you to easily compare thermal parameters using common boundary conditions. However, it is not suggested to use these values for evaluating thermal performance in a given system due to the fact that the boundary conditions are very different from most realized environments. For a rough first pass evaluation of a device thermal performance in a given environment, it is suggested to obtain the modified θ_{JC} and θ_{JB} based on different h (W/m² K) reflecting the system thermal design targets. This is documented in the appropriate thermal characterization document for your device/package, and can be obtained by request. This should

give a more realistic high-level understanding of the amount of power dissipation that can be achieved with a given thermal solution or vice versa, the thermal solution that should be designed for a given power dissipation at a given ambient with minimal calculation. After this is understood and a more precise assessment is desired, it is suggested to use the Xilinx thermal models to get a more accurate view of the end thermal operation of the device/package.

Sizing and Designing the Heatsink

Most solutions have unique sizing requirements both to properly interface the thermal solution to the device, and to fit into the overall system or chassis. Because of this, upfront design work must be done to ensure proper fitting. This section discusses how to properly size the heatsink to the device and share some considerations for sizing it to the system.

Device, Lid Type, and Associated Documentation Identification

The package description and all mechanical specifications/dimensions can be found in publicly available packaging and pinouts documents. These documents are organized by device family, so after a family member is known, this document should be downloaded to understand the details of the packages available for the selected device. Below you will find links to each of these documents:

- *7 Series FPGAs Packaging and Pinout Product Specification (UG475)*. See the 7 Series FPGAs Package Specifications table to understand the package type and body size. Consult the Mechanical Drawings chapter for precise dimensions for all packages for these devices. The Thermal Specifications chapter provides additional thermal information for these packages.
- *Zynq-7000 SoC Packaging and Pinout Product Specifications (UG865)*. See the Zynq-7000 SoC Package Specifications table to understand the package type and body size. Consult the Mechanical Drawings chapter for precise dimensions for all packages for these devices. The Thermal Specifications chapter provides additional thermal information for these packages.
- *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*. See the Package Specifications table to understand the package type and body size. Consult the Mechanical Drawings chapter for precise dimensions for all packages for these devices. The Thermal Specifications and Thermal Management Strategy chapters provide additional thermal information for these packages.
- *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*. See the Package Specifications table to understand the package type and body size. Consult the Mechanical Drawings chapter for precise dimensions for all packages for these devices. The Thermal Specifications and Thermal Management Strategy chapters provide additional thermal information for these packages.
- *Versal ACAP Packaging and Pinouts Architecture Manual (AM013)*. See the Package Specifications table to understand the package type and body size. Consult the Mechanical Drawings chapter for precise dimensions for all packages for these devices. For additional information on thermal packages, see the Thermal Specifications and Thermal Management Strategy chapters.

- *Kria K26 SOM Thermal Design Guide (UG1090)*. This guide covers many details of thermal design specific for Kria SOMs. It is suggested to read this entire document for Kria SOM thermal design. The K26 SOM Thermal Specifications table shows thermal targets and the Recommended Contact Area and 7X Thermocouple Locations on Top Heat Spreader for Characterization and Debug figure shows recommended contact area.

Determining Height and Contact Area for the Heatsink


Understanding the height and contact area of the heatsink to the device is important for any thermal solution. For lidless with stiffener ring, it is critical in determining proper heat sink fit. In all cases, this information is needed to understand the vertical tolerance and for calculations to ensure the proper amount of pressure is applied to the device so that no inference to neighboring components can exist. Using the Mechanical Drawings chapter from the appropriate packaging and pinouts document, locate the package outline drawing for the selected device/package.

The heatsink should be sized so that it has full contact to the top surface plane of the device to avoid potential hot spots or mechanical failure on the device. In many cases, it is oversized beyond the surface area without issues, however, in some cases there are limits to this sizing. The contact area and height specifications depend on the package and lid type.

Lidded


The contact area depends on the geometry of the lid. For forged lid and wire-bond chip-scale devices, the lid area generally mimics the package footprint area. Because of this, it is easy to understand and offers the most thermal contact possible for the device. For stamped lid flip chip and many standard wire bond packages, the lid contact area can be different from the package dimensions and can be calculated using the top view perspective of the package. The following figure is a top-view of an example stamped lid. The values indicated in red should be used for calculating surface area.

Figure 8: Example Stamped Lid Top View



The height or Z direction is generally indicated by the A parameter, which measures from the seating plane to the height of the device. The associated table in the appropriate user guide specifies the MIN, NOM, and MAX specification to allow for proper tolerance calculation for the neighboring component clearance, as well as when the heatsink is shared between multiple devices and tolerances needs to be calculated between devices. The following figure shows an example of a cross section from Xilinx mechanical drawings. The A parameter is used to determine package height.

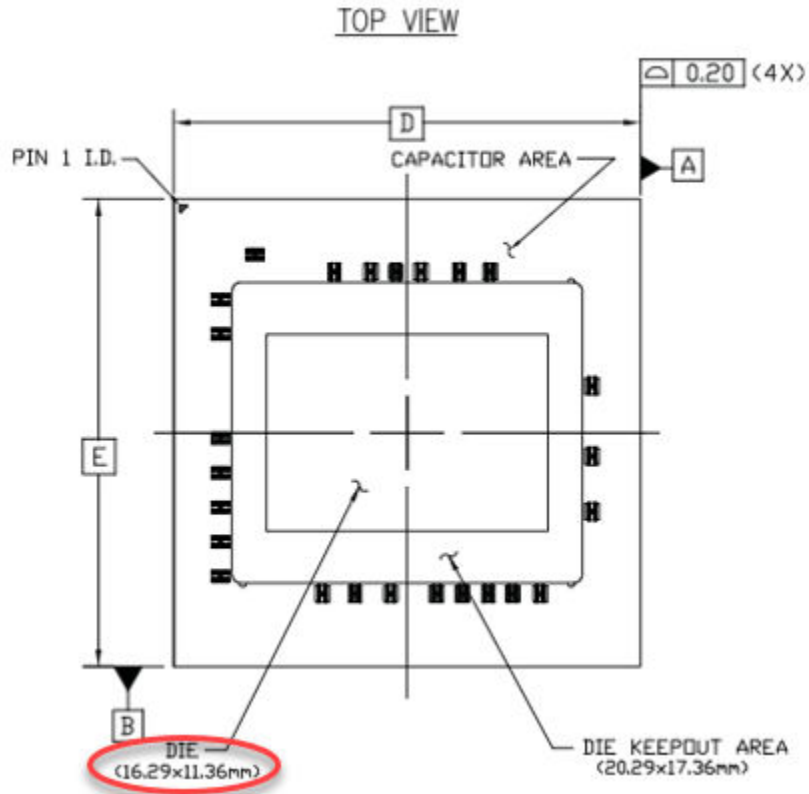
Figure 9: Example Forged Lid Cross Section



Bare Die


For bare die packages, the contact area is the area of the die. This is specified in the top-view perspective of the mechanical drawings. In this packaging type, the die is the highest point of the package and thus additional clearance to other points on the package like decoupling capacitors is not necessary. The following figure shows a top view of an example bare die. The value in red should be used to calculate surface area. Use the appropriate packaging and pinouts guide to find the values for specific devices.

Figure 10: Example Bare Die Top View



The package height is denoted by the A dimension in the mechanical drawings. This represents the height from the seating plane to the top of the die. The associated table in the appropriate user guide specifies the min, max, and nom parameters. The following figure shows an example diagram denoting the A parameter used for height.

Figure 11: Example Diagram Denoting A Parameter Used for Height




Lidless with Stiffener Ring

For lidless with stiffener ring packages, the contact surface area is indicated in the top view aspect of the mechanical drawing. Many of these package types use multiple die in the package, and the thermal solution should be constructed to fully contact all of them. Also, the stiffener ring is higher than the die surface, and thus an island or pedestal needs to be constructed to reach into the stiffener ring to allow proper contact to the die. The thermal solution should not touch the stiffener ring in any direction. Thus, the minimum island size should be the overall area

of all die in the package, and the maximum area should be sized so that it does not interfere with the stiffener ring. Some stiffener rings have a non-rectangular shape at the inner corners, so the island must be sized appropriately. The following figure is a top-view example of a lidless device with stiffener ring for reference. Use the appropriate packaging and pinouts guide to find the values for specific devices.

Figure 12: Example Top View of Lidless Device with Stiffener Ring




Using the previous figure as an example, the two numbers in red indicate the minimum island size. The values in blue can be used to calculate the maximum area. For this example, the minimum island size is 29.16 × 29.23 mm, and the maximum island size assuming a purely rectangular island is (47.5 – (2 × 6.00)) × (45 – (2 × 6.00)) or 35.5 × 33 mm. It is suggested to size this somewhere in between these boundaries. These values can be different for your selected package, as this is only an example of how to calculate island size. The following figure shows the pedestal extruding from the heatsink base calculated from the device package mechanical drawing to allow proper contact to die surface.

Figure 13: Pedestal Extruding from Heatsink Base




There are two vertical (height) parameters of interest, one from the seating plane to the top of the stiffener ring and one from the stiffener ring to the die surface. The first parameter is denoted by the letter A in the mechanical drawing. That is to be used to understand the clearance under the heatsink for nearby components like capacitors as well as to properly size the attachment for the heatsink. The height difference between the die surface area and stiffener ring is denoted by the A3 parameter and needed to properly size the minimum height of the pedestal or island for the heatsink. The following figure is a diagram denoting the A parameter used for height for clearance and attachment considerations.

Figure 14: Example Diagram Denoting the A Parameter used for Height Clearance



The following figure is a diagram denoting the A3 parameter. It indicates the minimum height requirements for the pedestal or island constructed into the heatsink.

Figure 15: Example Diagram Denoting the A3 Parameter for Pedestal or Island



InFO

For InFO packages, the contact area is the package body area. This is because the periphery of the package body is filled with a mold material that is planar to the die. This measurement is specified in the top view perspective of the mechanical drawings. The height is derived from the A parameter.

Figure 16: Example Top View of InFO Device

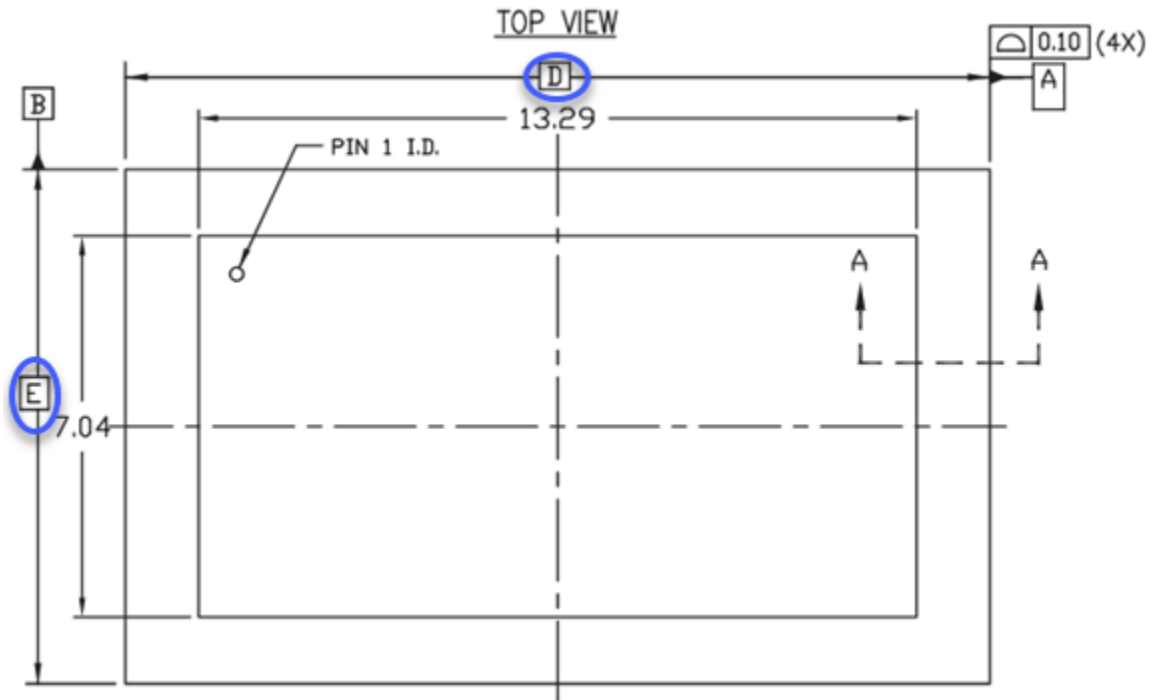




Figure 17: Example Diagram Denoting the A Parameter used for Clearance and Attachment Considerations



Kria SOM

Because the Kria devices have a heat spreader, consult the appropriate thermal design guide to find the recommended contact area. Mounting is done directly to the M3 mounting holes and thus there are no specific height tolerances to adhere to.

Figure 18: Recommended Contact Area for K26 SOM



Mechanical CAD Models and Samples

You can request CAD (STP) models and/or order mechanical samples if you need additional clarity on the mechanical specifications. The STP file can be used in many popular CAD tools to help understand and model the thermal apparatus, and to see how it fits before committing to prototyping or manufacturing. For further assurance in forming and fitting the device package, and to further inspect and test prototype systems, mechanical samples are available by order.

Coplanarity and Etching of the Contact Surface Area


It is important that the contact surface of the heatsink be parallel to the mating surface to ensure the best thermal performance, improve overall reliability, and to minimize the possibility of damage. When the surface of the heatsink is not parallel, the best case is that the TIM thickness is suboptimal over portions of the device. This compromises thermal performance and causes uneven cooling. If the heatsink is askew and there is enough uneven pressure applied to a small portion of the device contact area, it can cause lid deformity. In the case of lidless devices it can damage the die. For lidded or lidless devices, this should be avoided by ensuring in the design that the heatsink contact surface is parallel to the device plane. Check with the heatsink manufacturer to make sure the manufacturing process maintains this.

Traditionally, the heatsink contact surface is generally manufactured as a smooth plane. Xilinx recommends the use of an etched pattern to improve the overall thermal performance. The etched pattern should be controlled to have a coplanarity of 0.050 mm to achieve the best contact. This is especially beneficial for lidless devices but applicable to any Xilinx device. Using an etched pattern compared to a smooth plane offers several benefits:

- It reduces the overall contact resistance to the device by allowing micro-bubbles in the TIM1.5 or TIM2 material a place to recede rather than pooling and creating larger voids in the TIM.
- The TIM materials with metallic particles can allow for a thinner bond line thickness (BLT) by allowing larger particles to settle in the etching grooves creating a thinner and more uniform TIM thickness.
- The etching pattern adds more surface area for the TIM material to touch further reducing contact resistance.
- As the contact surfaces deflect due to thermal expansion over temperature, the etching region serves as a reserve for additional TIM allowing the TIM material to transfer between surfaces rather than create larger voids or leakage.

The amount of improvement varies depending on application. Xilinx has observed anywhere from 1°C to 5°C improvement using this technique, however, more improvement is possible in the case of larger TIM voids.

Figure 19: Xilinx Recommended Etching Pattern for Contact Surface of Heatsink




Selecting a Heatsink Base

After you know the contact area for the device, you can think about the design of the heatsink base. Decisions made here can significantly impact the area, performance, and cost of the thermal solution and is one that should not be overlooked. One consideration that can have significant impact on the base design is how many devices the heatsink must cool. Having a heatsink dedicated to a single device allows sizing and height tolerances to be optimized for that device often leading to both easier design and higher performance. On the other hand, a single heatsink that must touch multiple devices often must have compromises in size and height tolerance optimization that leads to less than ideal results. There are ways to compensate for this. However, for the basis of this application note, we assume the heatsink is dedicated to the Xilinx device or in cases that it is not, the Xilinx device is the most thermally critical and thus the heatsink is optimized towards the Xilinx device.

Size

The size of the heatsink base must be the package body plus the area needed for the attachment at a minimum. See [Heatsink Attachment and Mounting](#) for more information on heatsink attachment. As noted in [Coplanarity and Etching of the Contact Surface Area](#), the package body can be larger than the surface area. Refer again to the mechanical drawing and the parameters labeled D and E that pertain to the package body size. As indicated in the following figure, use the D and E parameters from the top view of the mechanical drawing to determine package body size.

Figure 20: D and E Parameters from the Top View of the Mechanical Drawing




It can be sized larger if more cooling is necessary. The thickness of the heatsink base depends on several factors including the material used, necessary rigidity of the heatsink, the amount of heat spreading necessary, the type and attachment of fins, the use of 2-phase cooling, and other factors.

One-Piece Heatsink vs Floating Lid

When using a lidless with stiffener ring package, there are two choices of base type that should be considered: a one-piece heatsink and a floating lid. A one-piece heatsink typically consists of the heatsink base with one side containing a properly sized island or pedestal, and the other side consisting of the fins or another means to conduct the heat to where it will be dissipated. Another option is a floating lid which is more-or-less a dedicated heat spreader that is used to transport heat from the die to a flat surface where a separate heatsink that transports the heat can be attached. Floating lids are commonly used during the prototype and bring-up phase of a product where a socket is used, the production heatsink has not been manufactured, or it must

be operated as an open-box where the proper thermal solution is possibly not fully attached or operational. In some cases, a floating lid can be used in a production environment. However, in general, one-piece heatsinks perform better, are less prone to reliability issues such as high-G drops and vibration, and can be more cost effective for higher volume procurement. Due to this, most engineers start with a floating lid and then transition into a one-piece heatsink as they go to a higher volume production product.

Figure 21: Floating Lid Showing Island, Pre-Applied TIM, and Etch Pattern



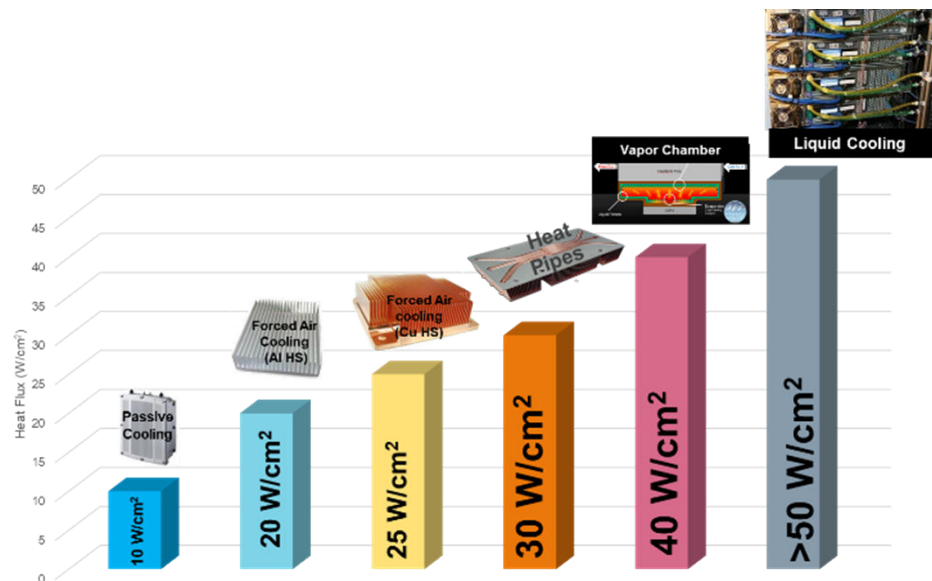
Sizing a Floating Lid

A floating lid still needs an island and should be sized the same whether using a one-piece heatsink or floating lid. The primary difference is in the sizing and construction of the base or heat spreader portion for the floating lid. While the base of the floating lid can be sized larger than the package body, it generally is sized to match the body. The stiffener ring also has guide holes at the four corners of the device that with appropriately placed pins on the floating lid, guides it for proper alignment to the device package without the need for additional attachments. The top surface is flat to allow interface to any additional thermal management component like an off-the-shelf heatsink. The following figure depicts an example floating lid mechanical drawing.

Airflow and Determining Heatsink Fins

With the exception of liquid cooled systems and those deployed in a vacuum like space, most systems will use air to transport the heat away from the device either by forced or natural convection. In most cases this necessitates fins, pins, or other extrusions being used either directly attached to the heatsink base or indirectly through conduction from the device to where the heat can be dissipated to the air. To get a suitable exchange of heat within the physical constraints of the system to the air, the size, number, attachment, position, and orientation of the fins is dependent on many system-specific parameters. These include ambient temperature, power dissipation, airflow speed, air pressure, turbulence, manufacturability, and even gravity. Thus, the specific design of the fins is particular to the application and open to you to find the right combination for your conditions and constraints. Most Xilinx evaluation boards use fan sinks, as airflow is not guaranteed by other means. Fansinks provide an effective and controlled means to ensure a minimum airflow to the heatsink. More typical to most systems, forced air is provided by fans in the case or chassis for higher power systems and natural convection for lower power systems. In any of these cases, analysis of the maximum power dissipation with maximum ambient conditions can be used as a first order indication of fin effectiveness. But as covered in a later section, thermal simulation tools should be used to evaluate and refine the fin effectiveness for a given design.

Figure 23: Example Thermal Solutions for Differing Demands



TIM Selection

The best designed heatsinks can significantly under-perform if there is too much contact resistance between them and the device. The purpose of thermal interface material is to minimize contact resistance and maximize overall thermal conductivity from the device to the thermal solution. TIM2 is used for lidded devices, and TIM1.5 for lidless devices.

The purpose of thermal interface material (TIM2 for lidded devices and TIM1.5 for lidless devices) is to minimize contact resistance and maximize overall thermal conductivity from the device to the thermal solution. It is a very important aspect of any thermal design, and one that deserves some time and thought so as to not compromise the overall thermal integrity of the solution.

Considerations when Selecting a TIM

Often, thermal conductivity is one of the first and primary parameters used when selecting a TIM. However, often what is far more important is choosing a TIM that can both minimize contact resistance (voids) and BLT. Thinking about it, even a TIM with relatively poor conductivity could provide superior performance if coverage and thickness is far better than a higher conductivity material. For this reason, Xilinx suggests focusing on the characteristics and application of TIM that ensures best coverage, contact, and BLT first, and the overall conductivity as a secondary parameter for TIM choice. Also, while the TIM material has to be thermally conductive, it should not be electrically conductive to minimize the chances of electrical shorts both inside and outside packages. Finally, long term stability and reliability should also be an obvious consideration with TIM choice.

Types of TIM Material

Grease, Putty, and Paste

The use of unsolidified form of TIM-like grease, putty, or paste is commonly used to thermally couple the heatsink to the device. These often can be easily dispensed and controlled and have good thermal properties (conductivity and low thermal contact resistance). Xilinx has found some phase change material (PCM) forms of paste to be a very effective TIM due to further reducing contact resistance. PCM is a type of TIM that employs a matter phase transition usually to/from a solid and liquid form as the material heats and cools, allowing for easy application while delivering superior contact resistance compared to single-phase materials. Xilinx is currently using Laird 780 SP PCM in its Alveo™ products, as it is a good balance of minimum BLT, applicability, and contact resistance. It also provides high reliability due to resistance to leaching, dry out, and other longer-term issues. This and similar materials have been proven to perform well with Xilinx lidded and lidless products.

Thermal Pads

Thermal pads are generally used when the thermal solution cannot be in direct contact with the device such as in the event multiple devices are connected to the same heatsink. Thermal pads compensate for greater tolerances in height than greases and putty. Xilinx does not recommend the use of thermal pads in direct contact with lidless devices because large BLT compensation compromises the performance improvements that lidless devices offer. Similarly, for lidded devices in thermally challenging situations, Xilinx recommends optimization of BLT that can allow for either a very thin pad, or move to an unsolidified TIM-like grease or paste. When using a pad, PCM-based materials tend to perform better than single-phase materials further optimizing contact resistance.

Thermal Epoxy, Tape, and Adhesive

Xilinx does not suggest the use of thermal epoxy, thermal tape, or other adhesives for use of a TIM or heatsink attachment in a production environment. The following are possible problems that can occur when using adhesives.

- Increased contact resistance
- Uneven thermal contact
- Reliability issues, including the possibility of lid separation
- Voids in TIM1
- Other damaging effects under shock, drop, and vibration
- Difficulty in debug and rework if deemed necessary

Because of this, Xilinx suggests using other materials for TIM rather than adhesives.

TIM Quantity and Application

The purpose of dispensing TIM by machine or by hand is to get coverage in the contact area as close to 100% as possible. Other goals include limiting voids, minimizing BLT, and controlling over dispensing to avoid TIM spillage and runoff. How this is done depends on the material and specific application details. This is being mentioned here as something to be discussed with manufacturing partners to ensure the goals are understood and well managed. It is expected to have some TIM runoff. As long as the TIM material is not electrically conductive or corrosive, no damage is expected in lidded or lidless applications. Being non-electrically conductive or corrosive is common for most TIM materials, however not all TIM materials. As previously mentioned, the use of etching on the contact surface also helps with achieving this goal.

Heatsink Attachment and Mounting

The last and arguably one of the more important parts of heatsink design is the heatsink attachment. The primary goal for most heatsink attachments is to apply an even 20 to 50 PSI of pressure for all lidded and lidless devices. InFO differs in that it requires a reduced 3 to 5 PSI due to the thin non-organic package substrate. SOM's target pressure to the heat spreader is 20 to 30 PSI. Pressure outside these ranges can cause the following problems.

- Degraded thermal performance
- Decreased overall reliability of the device
- Stressed corners or edges of the device, causing potential damage to the device

Because of those reasons, and as stated before, Xilinx does not recommend the use of adhesives for heatsink attachment to our devices. Furthermore, most clips that attach to our package substrate, are not recommended due to the following.

- Proper and even pressure to the device is difficult to control.
- Reliability to vibration, drops, and other G-forces is lessened.

- Amount of force is not enough with Z-clips to be considered a reliable option for heatsink attachment.

Therefore, Xilinx recommends using a dynamic mounting (pressure compensating) attachment that can apply enough force to at least four evenly spaced points around the device. This will allow proper design to these specifications. A common means to do this is with the use of spring screws at the four corners of the device. This generally requires through-holes in the board and thus should be considered prior to PCB layout so that the proper keep out and routing considerations are made for this attachment.

Figure 24: Example Heatsink with Four Spring Screws for Attachment



Calculating Spring Constant for Attachment

To properly size the spring constants for applying the 20 to 50 PSI of pressure to one of our lidded or lidless devices, we first need to calculate how much force is needed to determine which springs can provide that force. The equation for force is relatively simple: $\text{force} = \text{pressure} \times \text{area}$. The following steps show you how to find the force in pounds for our example.

1. Convert the units from mm into pounds by dividing the calculated mm^2 surface area calculated earlier by 645.16 to convert to inches^2 .
2. Multiply that area by the desired pressure (generally in the mid-range of the device pressure specification) to find the target force.
3. Divide that number by the number of attachment screws to get the force per screw. This force can be used to determine the spring constant needed based on the desired compression of the spring.

Note: Take special care to ensure all units are converted to compatible formats to ensure proper calculation.

After a spring is selected, we suggest that you back-calculate the pressure the springs will apply on the device to ensure the minimum or maximum pressure specification is not violated. The simple equation to calculate the necessary spring constant (k) is as follows.

Equation 1: Necessary Spring Constant (k)

$$k = \frac{P}{A \times x}$$

Where:

- k is the spring constant we are solving for in pound-force
- P is target pressure in PSI
- A is calculated surface area to the device in inches²
- x is the spring compression in inches

Evaluate Board Bracing

The amount of force needed to provide adequate pressure needs to be evaluated in terms of board reliability, particularly for larger devices where more force needs to be delivered and is spaced across greater distance potentially causing board deflection. In some cases, it can require additional board bracing or provision of a heatsink backing plate or brace to ensure the reliability of the Xilinx device and neighboring components is not compromised. This should be considered early because such bracing often requires board keep out areas, as well as modification of component placement like decoupling capacitors or memory components. The following figure is an example board back brace to be used to better distribute the force to a larger board area to prevent board deflection.

Figure 25: Example Board Back Brace



Designing a Heatsink for Multiple Devices and Packages

Many systems have multiple devices that need thermal management. Often it is desirable to simplify the heatsink design by having a single heatsink touch all these devices. There are many unique considerations for these designs, and this section is not intended to cover many of these aspects. It is, however, intended to highlight some focus areas common to this design.

The primary trade-off of this approach is that often thermal performance is sacrificed due to accommodating the different height tolerances of all the devices the heatsink must contact by increasing TIM thickness, and thus thermal resistance to the heatsink. The general suggestion here is to optimize TIM thickness for the most thermally critical component to get the best thermal transfer to the heatsink for that device, often at the cost of slightly increased thickness of the less critical devices. This can be a tricky balance, particularly if multiple devices have tight thermal margin. However, if the devices are expected to be close to or exceed the temperature limits, this is a necessary step to get the best system performance.

When it comes to lidless devices, managing the contact resistance to the heatsink is more critical, which generally lead to one of two choices. The first and most thermally efficient choice is to construct the proper island in the heatsink and optimize the BLT to the lidless device. The second option is to use a floating lid or an attached intermediate heat spreader connected to the top of the lidless device as a means to ensure good thermal contact to the die of the lidless device. It is not suggested to use a thick TIM or gap filler on lidless devices, as it is more susceptible to voids and larger thermal resistance resulting in poor thermal performance and potential reliability issues.

Due to the larger size of the heatsink base for multiple device thermal attachments, another consideration is that managing heat spreading within the base becomes more important. This generally leads to a higher importance to consider two phase cooling in the heatsink base as a means to maximize cooling and minimize the impact of one device heating another.

Lastly, it is often an even greater consideration where the attachments for the heatsinks exist on multiple device thermal management solutions. Xilinx devices still require even application of 20 to 50 PSI for lidded and lidless packages of the thermal management solution, and thus, having attachments not located very close to the device can make it difficult to manage this. If it is uncertain or impossible that the minimum pressure can be applied to the device, you should consider installing an attached heat spreader that can ensure the proper pressure on your Xilinx device.

Thermal Simulation

As thermal margin reduces, the need for accurate thermal analysis increases. The best way to make truly informed thermal decisions prior to building and characterizing a system is using thermal simulation. Thermal simulation can also be a useful tool to understand and correlate measured values, particularly when the results do not match anticipated outcomes. Because more and more systems have eroded thermal margin, the value of thermal simulation has increased dramatically in recent years.

When to Simulate

First simulation often occurs prior to device selection. Simulation can be used to evaluate different device and packaging options to ensure they can operate within its specifications for the given environmental constraints. Omitting this step can lead to difficulties, unperceived costs, and delays later in the design process. First simulation can often be simplified by just modeling the devices in a simplified boundary condition with the expected power to do a first pass evaluation. As a specific device is selected and more is finalized about the system, the thermal modeling should evolve to more accurately represent the board, chassis, and airflow characteristics to bring greater confidence and greater thermal margin to the thermal design. As power is refined and as board and system-level modifications and assumptions change, simulation should be reevaluated to ensure adequate thermal margin remains. This process helps to ensure that any thermal or mechanical issues are identified as early as possible to allow the most time to resolve. Omitting this can lead to identification of such issues towards the end of the project, potentially leading to compromises that can be undesirable or missing schedules.

Obtaining Thermal and Power Targets to use with Thermal Simulation

Prior to running a thermal sim, the power dissipation and temperature targets for the device need to be understood. As for thermal targets, it is best to first understand the minimum and maximum device operating temperatures for the device. This can be obtained from the appropriate Xilinx product selection guide or data sheet. In general, most devices fall into one of four classifications.

Table 1: Operating Temperatures for Xilinx Devices

Ordering Code	Temp Grade	Operating Range
C	Commercial	0 to 85°C
E	Extended	0 to 100°C ¹
I	Industrial	-40 to 100°C ¹
M	Military	-55 to 125°C
Q	Q-grade	-40 to 125°C

Notes:

1. These devices can offer an excursion temperature to 110°C depending on device.

For devices offering excursion temperature operation, the junction temperature might be able to operate at higher temperatures for short periods of time. To use this thermally beneficial feature, a mission profile needs to be generated and evaluated for the design. For more information on generating a mission profile and using an excursion temperature, refer to *Extending the Thermal Solution by Utilizing Excursion Temperatures* (WP517).

For devices using HBM memories, the additional step of setting temperature goals for HBM operation must also be taken. The maximum operating temperature for HBM memories is 95°C for constant operation. For the Virtex® UltraScale+™ and Versal® devices with -2 LE speed grade, excursion temperature operation can be up to 105°C for a limited time. Higher operating temperatures can impact refresh rates that affect the operational bandwidth, so that should be taken into consideration when setting HBM memory target temperatures.

Determining the target temperature for simulation depends on design goals. Often the target is set to the maximum operating temperature of the device or the calculated excursion temperature if allowed by the device and design. However, sometimes a lower target temperature is desired. For instance, if minimum operating power is desired, further reduction of the junction temperature results in less power dissipation in the devices. Another example is in the case of a hand held unit. Sometimes lower operating temperatures are necessary so as to not burn the user or keep battery or other component temperatures within reasonable limits. In any case, one of the first steps prior to thermal simulation is determining the appropriate junction temperature target for the application.

Xilinx devices often can have a wide range of power depending on use, and thus Xilinx provides the Xilinx® Power Estimator (XPE) tool to understand operational power that can be used during thermal analysis and regulator selection. XPE allows a user to provide high-level parameters for how they intend to operate the device and can provide a power calculation for thermal simulation. The tool can be obtained from the [Power](#) page of the Xilinx website. For Kria SOM users, the Power Design Manager (PDM) tool is used to determine power dissipation and apply predicted power to the SOM thermal model. The tool can be obtained from the [Kria K26 System-on-Module](#) page of the Xilinx website. After the tool is properly filled out, you can find the power to be applied to the simulation in the Total On-Chip Power section of the tool.

Figure 26: Example Total On-Chip Power Results from XPE

Total On-Chip Power	9.8 W	
Junction Temperature	100.0 °C	
Thermal Margin	0.0°C	0.0W

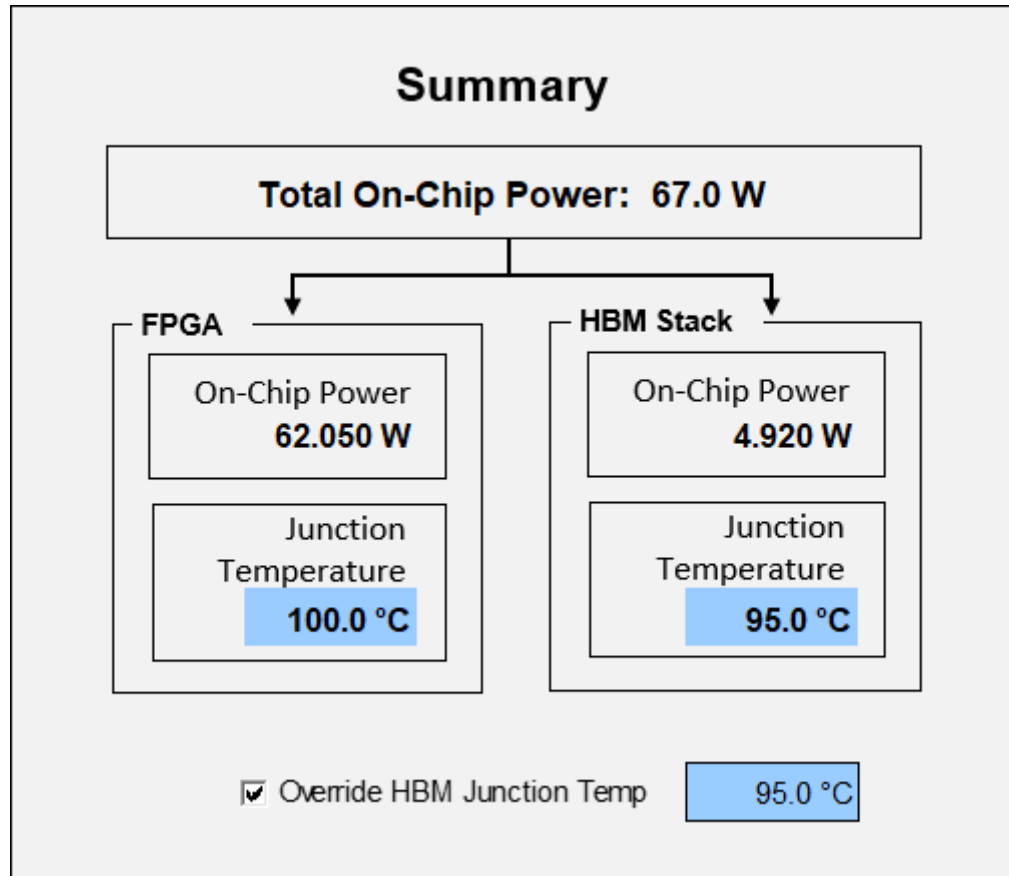
For an accurate estimation, use the determined thermal target for the device by forcing the junction temperature to that value. This can be done on the Summary page in the Environment section.

Figure 27: Example of Overriding Junction Temperature to Excursion 110°C for Power Estimation in XPE

Environment		
Junction Temperature	<input checked="" type="checkbox"/> User Override	110.0 °C
Ambient Temp		107.5 °C
Effective Θ_{JA}	<input type="checkbox"/> User Override	

For devices using HBM memories, the HBM target temperature and power summaries can be found in the following summary table on the Summary page. Two power entries should to be provided to the simulation model: one for the FPGA and one for the HBM stacks. Both can be obtained from the summary table in XPE. Even for devices that contain multiple die or multiple HBM stacks, only a single power value for each should be provided to the model comprising the total power of all die or stacks.

Figure 28: Example HBM Power and Thermal Information from XPE



Supported Simulators and Simulation Models for Xilinx Devices

To get accurate thermal results from a simulator, Xilinx suggests using a CFD simulator that is designed specifically for electronics. For this reason, Xilinx directly supports the use of Siemens (formerly Mentor) FloTHERM and Ansys IcePak simulators, and provides natively compiled models for those tools. While it is suggested to use one of the simulators supported by Xilinx, if another CFD is the only tool available, it is better to use it than to not perform thermal simulation. If the simulator supports the generic ECXML format, Xilinx can provide that model by request. If it is not supported, thermal models of Xilinx devices must be hand built by the user.

There are two methods to build such models: resistor-based (DELPHI) and material-based (Detailed). It is not suggested to build simple 2-resistor (2-R) models based on JEDEC values, because while those are developed using a well-defined method, its primary intention is used to compare the thermal performance of different packages but not necessarily under the operating conditions most users face. A more accurate method is to build a multi-resistor DELPHI based model. For our UltraScale+™ and prior families, the details on the dimensions and resistor values can be found in the appropriate thermal characterization documentation that is either included in the models or available by request. DELPHI models provide a good balance of accuracy and computational speed. If greater accuracy is required, a material detailed model can be built. This will require more computation time. This can be requested from Xilinx. For Versal and later devices, only material detailed models are supported.

For devices that support DELPHI, it is suggested to use the DELPHI models early in the thermal design process for several quick iterations of the design to allow quicker convergence of the thermal system. As the variations of the simulations narrow, a detailed model can be substituted to ensure adequate thermal margin exists when using a more accurate model.

Obtaining Xilinx Thermal Models and Information

Xilinx thermal models are available on the [Power Efficiency](#) page of the Xilinx website. For UltraScale+ and prior device families, the models on the website are DELPHI and often include the thermal characterization report. This report shares details of the use of the model as well as how to reproduce the model on other simulators. The document is available for most of the recent device families, but if the document does not exist, you can request it.

Detailed thermal models are also available by request for most devices. For Versal devices and later, Xilinx provides two forms of detailed models: simplified and full detail. The simplified model is a material model that allows the thermal results to be extracted without including the stiffener ring and other details in the model. These additional package details often impact meshing and convergence of the models, so it is suggested to use the simplified models for most thermal needs. The full detail models include all package details and better represent the physical form factor of the device.

Note: These models can have excessive runtime and other undesirable characteristics, so it is not suggested for most thermal simulation needs. These models are better suited for mechanical use.

3-D CAD models of the device packages are also available by request for most Xilinx devices. It is suggested to use these to ensure proper fitting of the heatsink prior to commitment for manufacturing.

Guidance for Modeling the Etching Pattern for Simulation

As mentioned earlier, Xilinx recommends the use of an etching pattern on the contact surface of the heatsink, and it is even more important for bare die and lidless devices as a means to allow minimal BLT to the die. The thermal modeling of the etching is very difficult and can make meshing and computation time become excessive. For these reasons, when using etching, Xilinx recommends modeling a flat surface with very low BLT and/or slightly higher thermal conductivity than the TIM material provides by modeling it implicitly. For example, through experimental testing, it was determined for the Laird SP780 material, a modeled thickness of 70 um and conductivity of 20 W/mK would model the effectiveness of the etching. This is a simple way to understand the in-system benefits of the etching without the time and effort needed to properly model the etching.

Evaluating Simulation Results

Determining Device Maximum Temperature from Simulation

A successful simulation usually results in the temperature monitor point (generally at the center of the die) being within the temperature targets determined for the simulation. For devices using HBM, there are two monitor points: one for the FPGA die and one for the HBM stacks. Even for devices that is composed of multiple die or multiple HBM stacks, the model simplifies the analysis to a single point for a grouping of die. Temperature gradients and hot-spot mitigation is generally not required in this analysis as long as proper contact of the designed thermal solution exists. The programmable nature of the design means that power density maps can shift and change over time, and thus Xilinx only requires a uniform power map for simulation and measurement to the device thermal diode system monitor (SYSMON) to remain in specification. Xilinx takes care of any additional higher temperature gradient through device testing. The end goal of the thermal designer is to have thermal simulations meet targets as specified by the monitoring point within adequate thermal margin, and that the device operates within that same target temperature as measured by the thermal diode within the measurement tolerances of the circuitry.

Determining Amount of Thermal Margin

The amount of thermal margin you should apply to the simulation results is unique to each design and designer. However, the one common thing is that all thermal designs should incorporate some thermal margin to ensure that uncertainties in the simulation do not lead to the device missing thermal targets in actual operation. Uncertainties in thermal design can occur in many places including:

- Power used during simulation
- Thermal model uncertainties
- Device measurement system monitor (SYSMON) uncertainty
- TIM contact and thickness variations
- Heatsink manufacturing tolerances
- Airflow deviations
- Other device power and thermal exhaust disparities

Designers often incorporate margin by using worst-case or sometimes worse than worst-case parameters in simulations. For example, they might incorporate highest possible power, upper limits of ambient, impeded airflow, improbably high TIM thickness and contact resistance, and more. Sometimes, they will apply additional margin on top of that. In general, Xilinx does not recommend this practice, as it can lead to worse than worst-case conditions that result in thermal over-design, or sometimes near impossible scenarios that will never be seen. The likelihood of all worst-case parameters occurring at once is very low so proper judgment must be made when entering simulation constraints/boundaries. In addition, proper judgment is needed when interpreting the results as a means of getting relative assurance that the end design will perform as needed but not be over designed resulting in greater cost, area, weight, and other undesirable characteristics.

In the end, it is up to the designer to determine the amount of margin that should be used to gain enough confidence in the thermal performance of the design. In general, it should be based on the confidence in the parameters in the design actually occurring in the final system collectively. If there is low confidence that most simulation parameters can be exceeded, greater margin might be more necessary than a system that is not expected to exceed any of the parameters. In the end, the quality and certainty of the simulation parameters and the simulation itself is the primary guide to how much margin is enough.

Relating Thermal Simulation Results Back to the Power Estimations

For the initial power estimations, an assumed fixed junction temperature is set to the target for the thermal design. After the thermal design becomes more final, it is suggested to back annotate simulation results to the power estimations to dynamically calculate the device junction temperature as a means to allow more accurate power estimations and monitor thermal margins as the internal design evolves. This can serve as a constraint for the logical design allowing better understanding and management of the design power to ensure that the thermal design is not over-designed later. To do this, derive a local ambient and effective θ_{JA} . This will serve as a simplified representation of the thermal performance of the designed thermal system as it relates to the junction temperature of the Xilinx device. The local ambient is generally the ambient temperature in the simulation as seen by the device. It can differ from the system-level ambient depending on whether the device is exposed to exhaust heat of other devices. However, local ambient should be set to the value as seen within the simulation. The effective θ_{JA} is a simplified thermal resistance value from device junction to ambient. This is used in power calculations to allow simple calculation of junction temperature based on estimated power. It also is used to assess the impact of increased/decreased power due to the effectiveness of the thermal system. It is calculated from simulation results by taking the power applied to the device in simulation divided by the calculated junction temperature as seen by the monitor point in the simulation model minus the local ambient temperature used.

Equation 2: Effective θ_{JA} Calculation for Power Estimation Based on Simulation Results

$$\text{Effective } \theta_{JA} = \frac{(\text{Simulated } T_J - \text{Simulated } T_A)}{\text{Simulated Power}}$$

The local ambient and effective θ_{JA} from the simulation results can be entered into XPE in the environment settings after the thermal design has been established as shown in the following figure.

Figure 29: Environment Settings

Environment		
Junction Temperature	<input type="checkbox"/> User Override	
Ambient Temp		44.6 °C
Effective θ_{JA}	<input checked="" type="checkbox"/> User Override	3.4 °C/W

This should also be used in the Xilinx Vitis™ or Vivado® software as XDC constraints to allow for proper understanding of the thermal design capabilities during FPGA design development:

```
set_operating_conditions -ambient_temp <temp_value> -thetaja  
    <thetaja_value>
```

The units for the command arguments *-ambient_temp* is °C and for *-thetaja* is °C/W. To convey the same values as shown in [Figure 22: Example Floating Lid Diagram](#), the following should be added to the Vivado XDC constraint file:

```
set_operating_conditions -ambient_temp 44.6 -thetaja 3.4
```

See [XPE for Early Thermal Analysis](#) for a video with more details on using thermal simulation results for power estimation.

Assembly, Test, and Debug

This section provides some guidance to the mechanical assembly of the thermal solution as well as best practices for thermal testing, characterization, and debug of the Xilinx device in-system.

Managing Transient Pressure

During the attachment of the thermal solution, short durations of higher pressure can be applied to the device as the force from the attachment increases and so as the TIM material compresses. When left unmanaged, it is possible the amount of pressure, even for a short time, can cause damage to the device.

One such method to do so is using proper installation equipment that can control the torque asserted on the screw attachments as to not exert excessive pressure on the device. An example tool is a smart torque screwdriver that can adjust angle, speed, and torque. This reduces stress to the device during assembly and allows the proper pressure profile to be applied to the device. This then allows safe assembly of the thermal system.

Measuring Heatsink Pressure

Use of a pressure sensor or gauge placed in between the thermal system and the device is a prudent method to measure and fine tune both transient and steady-state pressure applied to the device. This should only be used early in the prototyping stages, but not during operation. This will adjust and refine the assembly process to ensure transient pressure does not exceed the maximum 50 PSI exerted on the device. Further, it will ensure that the final steady-state pressure is even and within device pressure targets after the settling of the TIM and any pressure compensating features of the attachment.

Board Strain Measurement

It is also suggested to characterize the maximum board flex or strain made during construction, assembly, test, and use to ensure excessive board stress does not contribute to potential BLR issues.

The way to measure this is through the use of a strain gauge. Xilinx suggests keeping the stress on the board within + or - 500 μ strain. Dye and pry analysis can also be used in conjunction to strain measurement to ensure proper device contact to the pads exists after final assembly and test as further evidence of a reliable process.

Testing TIM Coverage

As mentioned in the [TIM Selection](#), poor TIM coverage or the formation of voids can seriously compromise the thermal integrity of the system. The target should be greater than 95% coverage to ensure proper heat conduction from the device without causing hot spots or other undesirable effects. The TIM coverage should periodically be verified after assembly to ensure application methods adequately meet this coverage. Additional qualification testing after several thermal cycles is also recommended to help ensure the longevity of the TIM performance.

Heatsink Removal Considerations

During the initial bring-up and debug stages of the design, removal of the heatsink is not an uncommon practice. Care should be taken so that device damage does not occur and that the device and heatsink are properly cleansed to allow proper reassembly without compromising the thermal integrity of the system. The thermal attachment should be disassembled so as to not create excessive uneven pressure to the device. This should only be used early in the prototyping and assembly stages, but not during operation. Often this is opposite to the steps that should be taken during assembly. The TIM often forms a bond that remains after the attachment is removed that can cause device damage if the heatsink is directly lifted. It is suggested to remove the thermal solution with a twisting motion as a means to break this bond. Sometimes it becomes necessary to heat the heatsink to further soften the TIM material so that excessive extracting force is not exerted on the device.

The TIM material should be cleaned from the device and heatsink and new TIM applied prior to reapplying the thermal solution. Follow the guidelines provided by the TIM provider for proper methods of TIM removal. Some consideration to the solvents used during the cleaning process should be taken to ensure corrosion or damage does not occur to the device or thermal solution. Some solvents that Xilinx has had success with include:

- Best – Toluene
- Better – Acetone
- Better – Isoparaffinic hydrocarbon (Isopar and Soltrol)
- Okay – Isopropyl alcohol

Creating a Thermal Test Design

The power within a programmable device is variable depending on use, so an upfront consideration should be made towards using the device for thermal testing and characterization. Of course it is best to use the final device programming image to do overall thermal characterization as it most accurately represents the power in the product. That is not always possible because often the board is designed and completed prior to the final image. The programmable nature of the FPGA allows you to create a power consuming design for the

purpose of thermal testing if a more accurate image is not available. Such a design generally is not functional but dissipates power to whatever levels desired for testing. There are several methods to design such a test image. If such a test image is required, it is suggested that the image is requested from Xilinx. Depending on the target device, there can be specific considerations to understand.

Using SYSMON for Operational Junction Temperature Measurement

Evaluating the in-system thermal performance should always be gauged by reading the system monitor circuitry (sometimes called SYSMON or XADC) in the device. This is an integrated ADC which can read an on-chip thermal diode as a means to get an accurate temperature reading. The diode is capable of providing the junction temperature of the device during operation. The temperature specifications for a device are always relative to the SYSMON readings and not to case or other measurement points. For the device to be considered to be operating within operating limits, the SYSMON temperature reading must take into account any uncertainties in the temperature measurement which can be obtained from the device datasheets. For example, an UltraScale+ device has a plus or minus accuracy of 3°C when using an external voltage reference. This means for the device to be considered to be operating in temperature specification, SYSMON should read at or below 97°C (100°C max operating temperature minus 3°C uncertainty). If the device supports temperature excursions, a SYSMON measurement reading is at or below 107° (110°C minus 3°C uncertainty). For Versal and later devices, the uncertainty into the device testing has been eliminated, thus simplifying the measurement to be interpreted without any uncertainty considerations. For Versal E and I-grade devices, a SYSMON reading of 100°C can be considered within specifications (110°C when considering excursion). SYSMON data can be accessed either within the design circuitry or using the JTAG debug interface and the Vivado software.

Power Measurement

As mentioned earlier, power dissipation is variable in the device, and power estimation is not always accurate and aligned to the actual operation. For this reason, a very important aspect to in-system thermal analysis is accurate measurement of the power being consumed in the device during testing and operation. The device can accurately measure the voltage via the SYSMON circuit. The current, however, cannot be measured directly from the device. This requires upfront considerations during board design to incorporate a means to measure the current being delivered to the power rails of the device so that actual power can be calculated. Many voltage regulators include some telemetry circuitry. Some do not, and some can have compromised accuracy. If the current measurement is either not possible or trustworthy from the regulator, it is suggested to add such telemetry to the board as a means to properly characterize power dissipation to the device. Without this key information, it is difficult to understand if the thermal management is under-performing or the system is simply consuming more power than anticipated. While device current measurement is generally good to have at any point of a

product lifetime, it is particularly important at the prototyping, characterization, and bring-up stages. In order to balance system cost and debug visibility, some customers choose to incorporate current measurement capability to the board which can be removed after production runs occur as a cost-saving measure. If the ability to accurately determine power does not exist, it can severely impede the ability to characterize and debug thermal issues.

Using Thermocouples for External Debug

A common thermal debug and characterization technique is the use of one or more thermocouples or thermistors placed in the system to get real-time temperature measurements. This can give much more insight to the thermal performance and limitations in the system. There are some precautions for this to ensure the results gathered are accurate and properly interpret thermal operation.

Xilinx does not recommend placing a thermocouple between the heatsink and the device to perform case measurements. This can often provide erroneous readings, compromise the performance of the system, and in some cases cause damage to the device. It is suggested instead to first measure at multiple points on the external base and/or fins of the heatsink to gauge thermal conduction to the heatsink. If enough temperature variance is seen between the SYSMON reading and heatsink measurement, investigate the heatsink contact to the device and/or heatsink construction as a general first course of action. If no issues are found there and a device case measurement is still desired, the best method to do so is to drill a hole through the heatsink to the area of the device that is intended to be probed.

Ensure there is very good contact between the thermocouple and the device through the given hole. This can be difficult, and poor contact to the device often leads to erroneous readings. Due to this reason as well as the fact that the thermal system must be modified/damaged to do this measurement, it is suggested to only be done as a last resort as in most cases this is not a necessary step to understand thermal issues in the system. When adhering the thermocouple to the measurement point, as repeatably mentioned, good contact to the measuring surface must be maintained. Improper application of thermal epoxy, such as when the thermocouple does not have direct contact to the case or excessive epoxy is applied, is shown to cause several degrees of difference to actual surface measurement.

Using thermal tape is a more common method to fix the thermocouple to the measurement surface because it is less permanent. However, the common issue here again is contact. If the tape is not taut enough and/or has a good surface to affix to, the thermocouple can lift, causing reduced readings. Coupling these common issues with the inherent uncertainties of the thermocouple measurement itself can often lead to misinterpretation of results. Thus, these and other possible sources of error should be eliminated before making conclusions from the data.

Correlate to Simulation and Update Thermal Power Estimation Data

An important but often overlooked step for the debug and finalization stages of the design is to relate any measurement data to the simulation and thermal power estimation constraints. Often towards the end of the design, proper time is not allocated to take the measurement data and compare with the initial assumptions and constraints as a means to both refine them for the current project as well as improve the initial assumptions for the next project. This allows for more confident updates and debug for the project if the need arises. Also, for future designs, it allows for a tightening of assumptions to allow for greater optimization and accuracy during the iteration.

The first step to this is to compare the measurements, particularly the junction temperature, with those seen during the simulation. If there are any noticeable deviations, explore how the thermal modeling can be improved to get a more accurate result. As for the power estimations, the design continues to develop and evolve after this thermal characterization takes place, and many times future updates can be applied. By reassessing and updating the thermal parameters to the XPE and Vivado power estimations, any subsequent updates can more accurately assess the impact against the fixed thermal solution. To do so, simply update the prior Effective θ_{JA} to the measured data.

Equation 3: Effective θ_{JA} Calculation for Power Estimation Based on Measurement Results

$$\text{Effective } \theta_{JA} = \frac{(\text{Measured } T_J - \text{Measured } T_A)}{\text{Measured Power}}$$

This new data should be applied to XPE and Vivado similarly to what was referenced in the [Evaluating Simulation Results](#) section.

Example Heatsink Design

Using the information discussed up to this point, this section briefly describes an example heatsink design for the VC1902-VSVA2197 device fitting the needs of the application in terms of device resources, I/O, GT, and package size.

1. Perform an initial thermal evaluation of the package.

Get an initial power estimate from XPE for the desired application. Obtain the thermal simulation model from the [Power Efficiency](#) page of the Xilinx website. Perform a quick simulation using a simple representative boundary condition and ensure there is some margin in the operating junction temperature. If adequate margin is not found, iterate with different acceptable device, package, power and environmental conditions until a workable solution is found. For the purposes of this example, it is found that there is adequate thermal margin.

2. Identify the associated documentation and package lid type, contact area, and height requirements.


Because this is a Versal device, we used the *Versal ACAP Packaging and Pinouts Architecture Manual (AM013)* to determine the target package is a 45 x 45 mm lidless package.

Table 2: Example Table from the Packaging and Pinouts Guide Specifying High Level Information of a Target Package

Packages	Description	Package Specifications			
		Package Type	Pitch (mm)	Size (mm)	LSC Ball Grid Size (balls)
SFVB625	Super-fine pitch with forged lid	BGA	0.8	21 x 21	-
VBVA1024	Fine pitch bare-die		0.92	31 x 31	4 x 4
VFVB1024	Fine pitch with forged lid			31 x 31	4 x 4
VFVB1369	Fine pitch with forged lid			35 x 35	5 x 5
VSVE1369	Fine pitch lidless with stiffener ring			35 x 35	5 x 5
VSVA1596	Fine pitch lidless with stiffener ring			37.5 x 37.5	5 x 5
VIVA1596	Overhang fine pitch lidless with stiffener ring			40 x 40	6 x 6
VFVA1760	Fine pitch with forged lid			40 x 40	6 x 6
VFVC1760	Fine pitch with forged lid			40 x 40	6 x 6
VSVD1760	Fine pitch lidless with stiffener ring			40 x 40	6 x 6
VSVA2197	Fine pitch lidless with stiffener ring			45 x 45	7 x 7
VSVA2785	Fine pitch lidless with stiffener ring			50 x 50	9 x 9

The package physical dimensions are located within the mechanical drawings section of the document.

Figure 30: Package Physical Dimensions for Example Target Package from the Packaging and Pinouts Guide



VC1902-VSVA2197 is a lidless design. As previously defined, the contact area is defined by the area of the die surface. This is represented in the drawing with the dimensions of 25.75 × 17.78 mm, and is the minimum dimension of the island on the heatsink base. The maximum dimension in the X direction is 45 – (2 × 4.00) - 0.20 = 36.8 mm and in the Y direction 45 – (2 × 6) - 0.20 = 32.8 mm. However, you cannot have it overhang this much in both directions due to the protrusions in the inner corners of the stiffener ring. Thus, it is

NOTES:

1. ALL DIMENSIONS AND TOLERANCES CONFORM TO ASME Y14.5M-1994
2. SYMBOL "M" IS THE PIN MATRIX SIZE.

suggested to slightly oversize it, but not to the point that it touches the corners of the stiffener ring. Anticipating the uncertainties in the manufacturing of the island as well as the placement of the heatsink, we allocate a little more than 2 mm to each dimension resulting in a 30 × 22 mm contact surface for the island. This should ensure that the entire die surface area is covered under all circumstances without interference from the stiffener ring.

Next, the height of the island is determined by the offset of the die plane to that of the stiffener ring. This would be the MAX of the A3 dimension of 1.15 mm. The height of the island must be sized so that when assembled, the heatsink base does not touch the stiffener ring so some additional margin can be added to ensure this. Finally, the A dimensions is used for the minimum height of the heatsink base from the PCB.

A 3D STP file is then requested from Xilinx so that these sizing assumptions can be tested with the model to ensure good fit without interference in your CAD program. Further fine tuning could persist. It is also a good time to order a mechanical sample of this package so that it is available in time for receiving the first heatsink prototype for proper fitting analysis.

3. Design heatsink base, fins, airflow, and attachment.

The heatsink base for this example design is simple aluminum because adequate airflow and not excessively high ambient conditions are expected. If the expectation is to operate in a more challenging environment, other materials and/or 2-phase cooling should be evaluated. We plan to allocate 100 × 70 mm to the heatsink base area which should allow enough area to a fansink, and fins to allow for sufficient cooling. These parameters can be refined after thermal simulation. Using the parameters for the island gathered previously, the bottom center of the heatsink base is designed. An etched surface is specified into the island surface area in contact to the die to further improve effectiveness.

Several fin designs were discussed and evaluated and eventually modeled to determine an optimal arrangement. A 2 1/5 inch diameter fan is recessed into the fins to accommodate height requirements for the overall design.

To allow for even and proper pressure compensation for the life of the product, four equally spaced spring screws are designed into the corners of the heatsink as well as the proper keepouts specified for the PCB design.

The spring constants for the springs are calculated. Initial assumptions are target pressure of 30 PSI for this device/package, area is the die area of 25.75 × 17.78 mm = 457.835 mm² and the amount of spring compression is 0.5 inches. The first step is to convert the die area to inches: 457.835 / 645.16 = 0.7096 in².

$$k = \frac{30}{0.7096 \times 0.5} = 84.5 \frac{\text{pounds}}{\text{inch}} \text{ or } 114.57 \frac{\text{newton}}{\text{meters}}$$

4. Choose TIM and determine application.

- VC1902-VSVA2197 is a lidless design. A TIM 1.5 is suitable for this application. In order to get the best contact resistance we will select a phase change material (PCM) composite. Taking into consideration availability, contract manufacturer capabilities, and prior experience, we have selected the Laird 780SP material. The maximum possible coplanarity deviation for the die is 0.1 mm and thus enough material should be dispensed so that it covers the volume of the maximum deviation of this mating surface or $25.75 \times 17.78 \times 0.1 \text{ mm} = 45.8 \text{ mm}^3$. Additional material is specified to ensure proper coverage and margin. Details of the dispensing of material is worked out later with the contact manufacturer.
5. Determine thermal parameters for simulation and validate thermal solution.

Using the thermal model obtained in the initial evaluation, now elaborate the thermal model to more accurately represent the thermal system. Perform, evaluate, and iterate as necessary to gain as much thermal margin as possible. If adequate thermal margin is not obtained, re-evaluate prior decisions to obtain adequate thermal margin before proceeding.
 6. Finalize the design.

After a final design is achieved, the heatsink design is then sent to manufacturing. More details need to be worked out with the contact manufacturer for assembly and test. The details of this is dependent on the manufacturer. After the heatsink is received, it should be tested for mechanical fit as well to ensure that pressure and board strain meet the design goals. Then thermal characterization should proceed to validate that the thermal simulation and design are still within design margins.

Conclusion

As stated in the introduction, every thermal design is unique. This application note and the examples included are intended to point out some best practices, but they are not necessarily a precise procedure to how every thermal design should proceed. Following these guidelines can provide better thermal management for your products.

References

These documents provide supplemental material useful with this guide:

1. Xilinx [Power Efficiency](#)
2. *7 Series FPGAs Packaging and Pinout Product Specification (UG475)*
3. *Zynq-7000 SoC Packaging and Pinout Product Specifications (UG865)*
4. *UltraScale and UltraScale+ FPGAs Packaging and Pinouts Product Specification (UG575)*
5. *Zynq UltraScale+ Device Packaging and Pinouts Product Specification User Guide (UG1075)*
6. *Versal ACAP Packaging and Pinouts Architecture Manual (AM013)*
7. *Kria K26 SOM Thermal Design Guide (UG1090)*
8. [Thermal Models Download](#)
9. [XPE Download](#)
10. [XPE for Thermal Analysis Video](#)

Revision History

The following table shows the revision history for this document.

Section	Revision Summary
06/06/2022 Version 1.0	
Initial release.	N/A

Please Read: Important Legal Notices

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby **DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE**; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>.

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

© Copyright 2022 Xilinx, Inc. Xilinx, the Xilinx logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.