



NUMA Topology for AMD EPYC™ Naples Family Processors

Publication #	56308	Revision:	0.70
Issue Date:	May 2018		

© 2018 Advanced Micro Devices, Inc. All rights reserved.

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

Trademarks

AMD, the AMD Arrow logo, AMD EPYC, 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.

Contents

Revision History	4
NUMA Topology in Naples SoC.....	5
Channel-Interleaving: Four NUMA Nodes Per Socket	6
Channel-Pair Interleaving: One NUMA Node Per Socket	7
Socket Interleaving (2P system)	7

List of Figures

Figure 1. Illustration of a Naples Die and Socket.....	5
Figure 2. Different Interleaving Options Available on AMD EPYC™	6
Figure 3. Channel Interleaving in an AMD EPYC™ 2P System	6
Figure 4. Channel-Pair Interleaving in AMD EPYC™ 2P System	7



Revision History

Date	Revision	Description
May 2018	0.70	Initial public release.

NUMA Topology in Naples SoC

The NUMA topology in the SoC, codenamed “Naples”, is comprised of the following:

- 4 Dies per package.
- 2 Core-Complexes (CCXs) per Die.
- Up to 4 Cores per CCX sharing an L3 cache. All CCXs configured equally.
- 2 Threads per Core (SMT) sharing an L2 cache.
- 2 Memory channels per die
- 8 memory channels per package with up-to 2 DIMMs per channel.
- Platform support for one or two SoCs (1P or 2P).

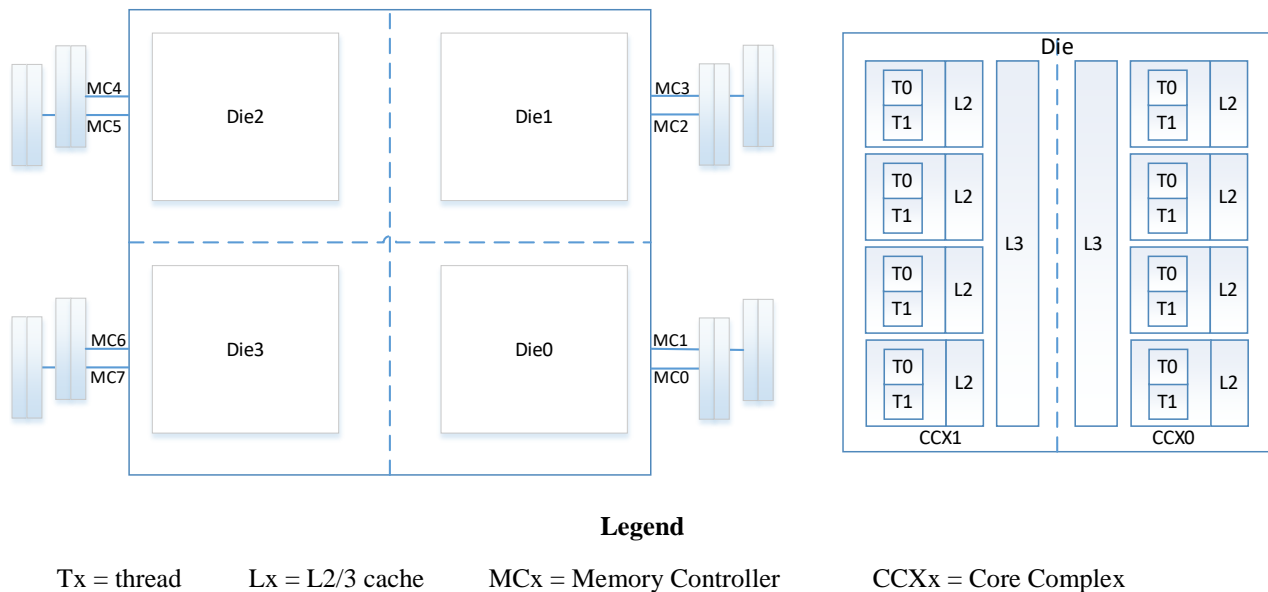


Figure 1. Illustration of a Naples Die and Socket

Interleaving Options

The default memory interleaving setting for Naples is “Channel Interleaving.” This configuration provides the highest throughput and lowest average latency for memory transactions. As the granularity of interleaving increases, the average memory latency also increases, with the highest average latency resulting from the “Socket Interleaving” setting.




Figure 2. Different Interleaving Options Available on AMD EPYC™

Channel-Interleaving: Four NUMA Nodes Per Socket

This NUMA setting represents the interleaving of the two memory channels on each die, with each die configured as a NUMA node. This configuration requires that a DIMM is populated on each channel, otherwise any non-symmetrical DIMM population will not be interleaved and will instead be stacked on top as a general memory pool. Furthermore, any die that does not have one of the two channels populated will not be interleaved.

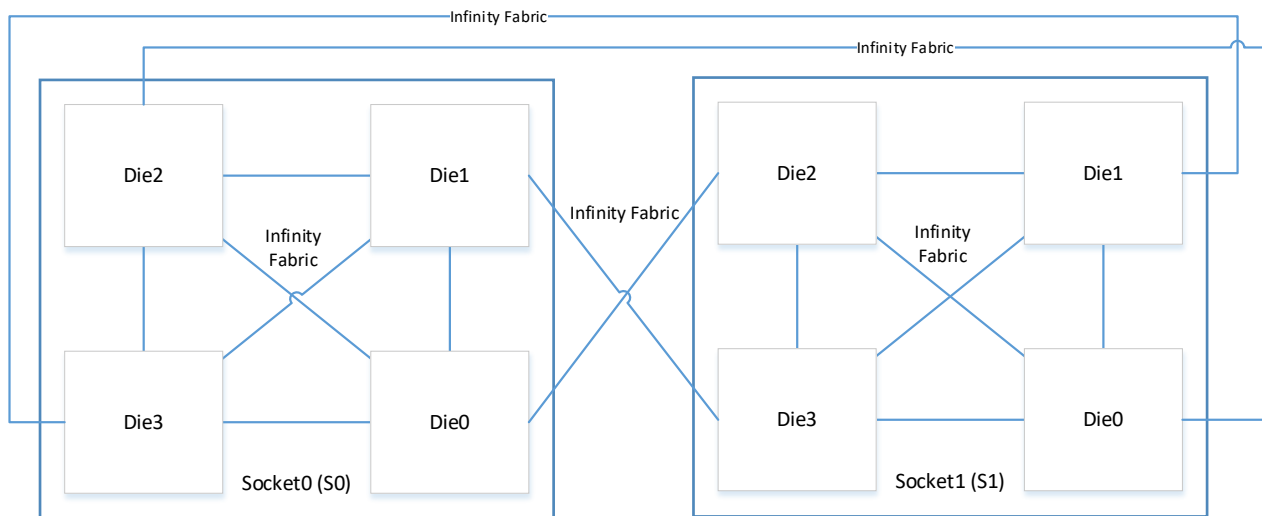


Figure 3. Channel Interleaving in an AMD EPYC™ 2P System

Channel-Pair Interleaving: One NUMA Node Per Socket

This NUMA setting represents the interleaving of all eight memory channels on each socket, with each socket configured as a NUMA node. This configuration will require that every channel is populated with equal size memory. In a 2P system both sockets must be in this mode or one of the sockets must have no memory populated.

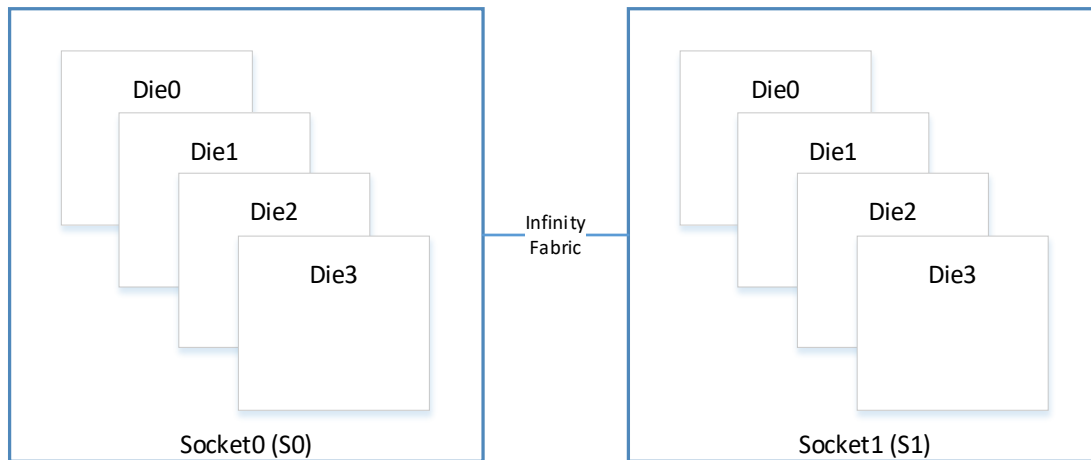


Figure 4. Channel-Pair Interleaving in AMD EPYC™ 2P System

Socket Interleaving (2P system)

On a 2P system, enabling socket interleaving effectively creates a single NUMA node for the whole system, in which case a Static Resource Affinity Table (SRAT) and System Locality Information Table (SLIT) are not necessary. This memory interleaving option is only available if all channels in both sockets have equal size memory.