

The Lenovo logo is displayed in white text on a black rectangular background.

Balanced Memory Configurations with 2nd Gen and 3rd Gen AMD EPYC Processors

Demonstrates balanced memory guidelines for second and third generation AMD EPYC processors

Compares the performance of balanced and unbalanced memory configurations

Explains memory interleaving and its importance

Provides recommendations on how to balance memory and maximize performance

Nathan Pham
Joseph Jakubowski
Jamal Ayoubi



Abstract

Configuring a server with balanced memory is important for maximizing its memory bandwidth and overall performance. Lenovo ThinkSystem servers running AMD EPYC processors have eight memory channels per processor and up to two DIMMs per channel, so it is important to understand what is considered a balanced configuration and what is not.

This paper defines three balanced memory guidelines that will guide you to select a balanced memory configuration. Balanced and unbalanced memory configurations are presented along with their relative measured memory bandwidths to show the effect of unbalanced memory. Suggestions are also provided on how to produce balanced memory configurations.

This paper is for customers and for business partners and sellers wishing to understand how to maximize the performance of Lenovo ThinkSystem servers with second-generation and third-generation AMD EPYC processors.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

<http://lenovopress.com>

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction	3
Memory topology	3
Memory interleaving	4
Balanced memory configurations	4
About the tests	5
Applying the balanced memory configuration guidelines	5
Summary of the performance results	17
Summary	19
Change history	20
Authors	20
Notices	21
Trademarks	22

Introduction

The memory subsystem is a key component of AMD EPYC server architecture which can greatly affect overall server performance. When properly configured, the memory subsystem can deliver extremely high memory bandwidth and low memory access latency. When the memory subsystem is incorrectly configured, memory bandwidth available to the server can become limited and overall server performance can be reduced.

This paper explains the concept of balanced memory configurations that yield the highest possible memory bandwidth from the AMD EPYC architecture. By increasing the number of populated memory DIMMs from one to sixteen, examples of balanced and unbalanced memory configurations are shown to illustrate their effect on memory subsystem performance.

This paper specifically covers the second-generation AMD EPYC processor family (EPYC 7002), formerly code named *Rome*, and the third-generation AMD EPYC processor family (EPYC 7003), formerly code named *Milan*.

Memory topology

Access to the information stored on memory DIMMs is controlled by memory controllers within the processor. The AMD EPYC family processors have eight Unified Memory Controllers (UMC). Each UMC has one memory channel, and each memory channel supports up to two memory DIMM slots.

Figure 1 illustrates how an AMD EPYC family processor's memory controllers are connected to memory DIMM slots.

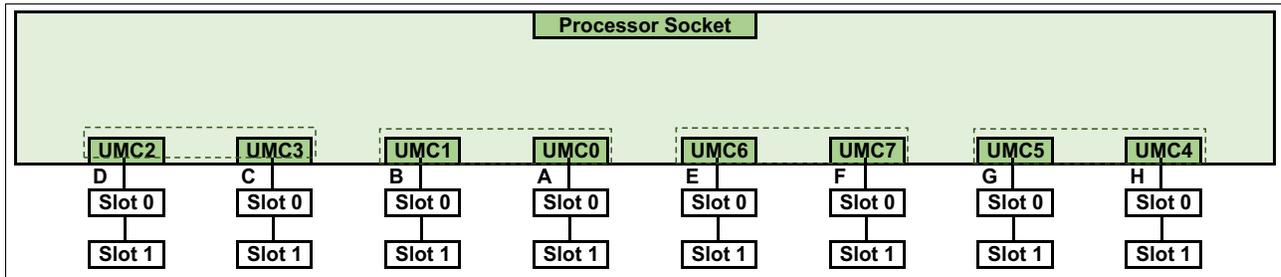


Figure 1 AMD EPYC family processors with eight Unified Memory Controllers (UMC), eight memory channels, and sixteen memory DIMM slots

Memory channels are grouped in pairs. As shown in Figure 1 above, the channel pairs are AB, CD, EF, and GH. To illustrate various memory topologies for a processor, different memory configurations will be designated as D:C:B:A:E:F:G:H where each letter indicates the number of memory DIMMs populated on each memory channel.

- ▶ A refers to Memory Channel 0 on UMC0
- ▶ B refers to Memory Channel 1 on UMC1
- ▶ C refers to Memory Channel 3 on UMC3
- ▶ D refers to Memory Channel 2 on UMC2
- ▶ E refers to Memory Channel 6 on UMC6
- ▶ F refers to Memory Channel 7 on UMC7
- ▶ G refers to Memory Channel 5 on UMC5
- ▶ H refers to Memory Channel 4 on UMC4

As an example, a 2:2:2:2:1:1:1:1 memory configuration has 2 memory DIMMs on Memory Channels A, B, C, D and 1 memory DIMM on Memory Channels E, F, G, H.

Memory interleaving

The AMD EPYC processor family optimizes memory accesses by creating interleave sets across the memory controllers and memory channels. For example, if two memory channels on the same channel pair have the same total memory capacity and number of ranks, a 2-channel interleave set is created across the two memory channels.

Interleaving enables higher memory bandwidth by spreading contiguous memory accesses across both memory channels rather than sending all memory accesses to one memory channel. In order to form an interleave set between two channels on the same pair, the two channels are required to have the same memory type and total capacity.

For EPYC 7002 processors, the number of memory channels in an interleave set has to be 1, 2, 4 or 8. For EPYC 7003 processors, the number of memory channels in an interleave set can be 1, 2, 4, 6 or 8.

If one interleave set cannot be formed for a particular memory configuration, it is possible to have multiple interleave sets. When this happens, performance of a specific memory access depends on which memory region is being accessed and how many memory DIMMs comprise the interleave set. For this reason, memory bandwidth performance on memory configurations with multiple interleave sets can be inconsistent.

Contiguous memory accesses to a memory region with fewer channels in the interleave set will have lower performance compared to accesses to a memory region with more channels in the interleave set.

Figure 2 illustrates a 4-channel interleave set which results from populating identical memory DIMMs on two memory channel pairs. This 4-channel interleave set interleaves across memory controllers and between memory channels. Consecutive addresses alternate between memory controllers with every fourth address going to each memory channel.

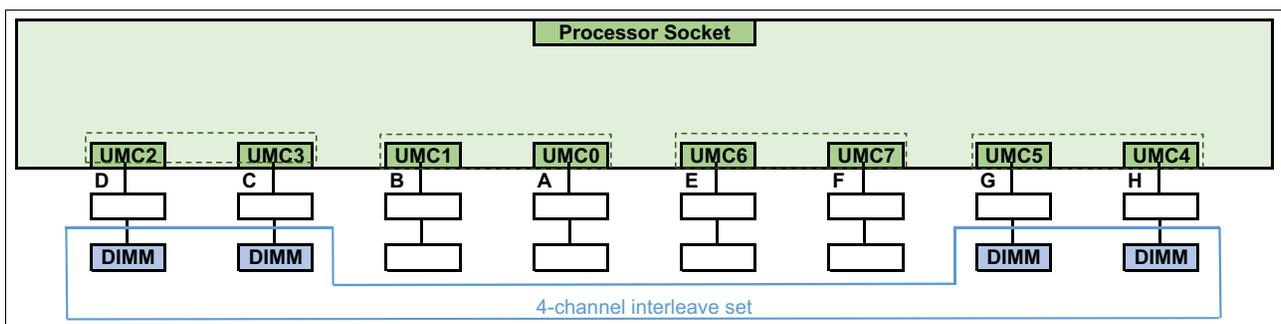


Figure 2 4-channel interleave set across memory controllers and between memory channels

Balanced memory configurations

Balanced memory configurations enable optimal interleaving which maximizes memory bandwidth. Memory interleaving starts with a channel pair, so memory populated in two channels of a channel pair need to have the same total memory capacity and total number of ranks to form an interleave set.

Optimal socket memory bandwidth is achieved when all channel pairs have the same configuration. System level memory bandwidth is optimal when each physical processor socket has the same physical memory capacity.

As a result, the basic guidelines for a balanced memory subsystem are as follows:

1. Memory should be populated in channel pairs. Both channels in a channel pair should have the same total memory capacity and the same number of ranks.
2. All populated channel pairs should have the same memory configuration. The number of populated channels has to be 2, 4, or 8 for the EPYC 7002 processors, or 2, 4, 6, or 8 for the EPYC 7003 processors.
3. All processor sockets on the same physical server should have the same configuration of memory DIMMs

We will refer to the above guidelines as *balanced memory guidelines 1, 2 and 3* throughout this paper.

Tip: For best performance, choose balanced memory populations with 2, 4, 8 or 16 DIMMs per processor. For Gen 3 processors (EPYC 7003) only, 6-DIMM and 12-DIMM population can form a 6-channel interleave for best performance. All other memory populations are unbalanced and memory performance can be severely degraded.

About the tests

STREAM Triad is a simple, synthetic benchmark designed to measure sustainable memory bandwidth. Its intent is to measure the highest memory bandwidth available. STREAM Triad will be used to measure the sustained memory bandwidth of various memory configurations to see the effect of suboptimal memory configurations on memory bandwidth. Unless otherwise stated, all test configurations were done at the same memory speed, 2933 MHz.

For more information about STREAM Triad, see the following web page:

<http://www.cs.virginia.edu/stream/>

As described in “Memory interleaving” on page 4, it is possible to get inconsistent memory bandwidth performance when there are multiple interleave sets in a memory configuration. To account for this behavior, we configure each test run to make sure the STREAM benchmark accesses the entire memory address map. The reported test result reflects the average memory bandwidth performance for these configurations. In reality, memory bandwidth performance in configurations with multiple interleave sets can be unpredictable, depending on which interleave set is being accessed. For this reason, we also provide memory bandwidth measurement for the worst case. This happens when memory in the interleave set with the lowest number of memory channels is being accessed.

Applying the balanced memory configuration guidelines

We will start with the assumption that balanced memory guideline 3 (described in “Balanced memory configurations” on page 4) is followed: all processor sockets on the same physical server have the same configuration of memory DIMMs. Therefore, we only have to look at one processor socket to describe each memory configuration.

In our lab measurements, all memory DIMMs used were 32 GB dual-rank (2R) RDIMMs. The number of memory DIMMs installed will be increased from 1 to 16 to see the effect on memory bandwidth. For each memory configuration, it will be determined which balanced memory guidelines are followed, and the number and type of interleave sets will be shown. Any recommendations for improving the performance of the memory configuration will also be pointed out.

When installing memory DIMMs into your server, follow the DIMM installation sequence for your particular model. The examples in this brief follows the recommended memory population sequence as shown in Table 1 on page 6. The table also includes the recommended 6-DIMM and 12-DIMM configuration that are only for EPYC 7003 processors.

Table 1 Memory population sequence for the AMD EPYC processor (7002 and 7003, except where noted)

Number of DIMMs	UMC		UMC2		UMC3		UMC1		UMC0		UMC6		UMC7		UMC5		UMC4	
	Channel	D1	D0	C1	C0	B1	B0	A1	A0	E0	E1	F0	F1	G0	G1	H0	H1	
	DIMM 1	DIMM 2	DIMM 3	DIMM 4	DIMM 5	DIMM 6	DIMM 7	DIMM 8	DIMM 9	DIMM 10	DIMM 11	DIMM 12	DIMM 13	DIMM 14	DIMM 15	DIMM 16		
1			X															
2	X		X															
3	X		X											X				
4	X		X											X		X		
5	X		X				X							X		X		
6 (EPYC 7002)	X		X		X		X							X		X		
6 (EPYC 7003)	X		X				X			X				X		X		
7	X		X		X		X			X				X		X		
8	X		X		X		X			X		X		X		X		
9	X		X	X	X		X			X		X		X		X		
10	X	X	X	X	X		X			X		X		X		X		
11	X	X	X	X	X		X			X		X	X	X		X		
12 (EPYC 7002)	X	X	X	X	X		X			X		X	X	X	X	X		
12 (EPYC 7003)	X	X	X	X			X	X	X	X			X	X	X	X		
13	X	X	X	X	X		X	X		X		X	X	X	X	X		
14	X	X	X	X	X	X	X	X		X		X	X	X	X	X		
15	X	X	X	X	X	X	X	X	X	X		X	X	X	X	X		
16	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X		

Configuration of 1 DIMM - unbalanced

We will start with one memory DIMM which yields the 0:1:0:0:0:0:0:0 memory configuration shown in Figure 3.

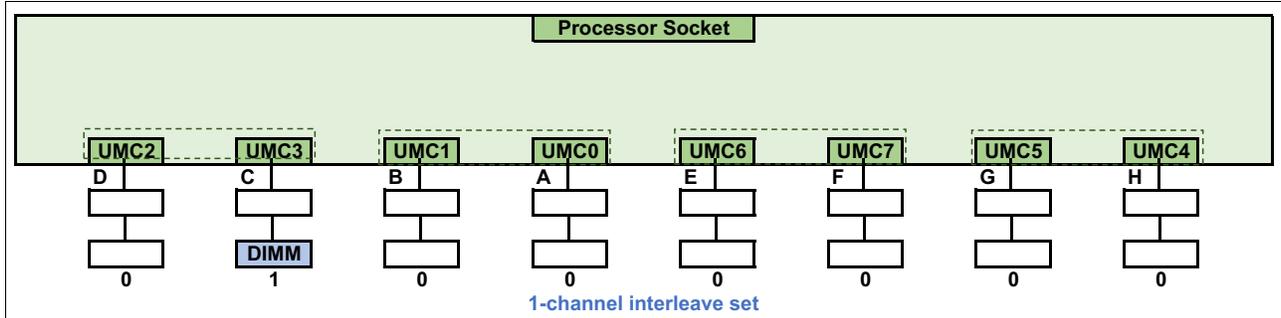


Figure 3 0:1:0:0:0:0:0:0 memory configuration relative memory bandwidth: 14%

Balanced memory guideline 1 is not followed with only one populated memory channel. Guideline 2 is not applicable since only one memory channel is populated. This is not a balanced memory configuration.

A single 1-channel interleave set is formed. Having only one memory channel populated with memory greatly reduces the memory bandwidth of this configuration which was measured at 14% or about one eighth of the full potential memory bandwidth.

The best way to increase the memory bandwidth of this configuration is by using more memory DIMMs. Two 16 GB RDIMMs populated on two channels of a channel pair would provide the same memory capacity while nearly doubling the memory bandwidth.

Configuration of with 2 DIMMs - balanced if installed correctly

The recommended memory configuration with 2 memory DIMMs is the 1:1:0:0:0:0:0:0 memory configuration shown in Figure 4.

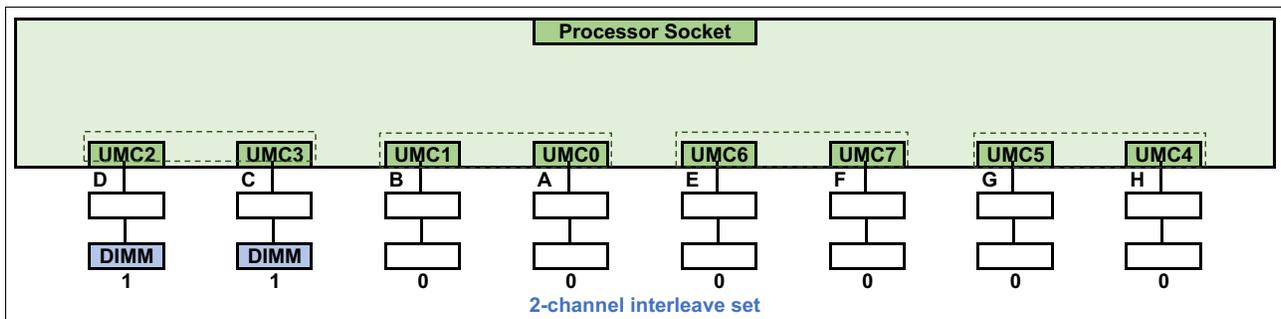


Figure 4 1:1:0:0:0:0:0:0 memory configuration relative memory bandwidth: 28%

This memory configuration follows balanced memory guideline 1 and 2 as both memory channels in a channel pair are populated with the same DIMM. This is a balanced memory configuration.

A single 2-channel interleave set is formed across the two memory channels. Only two memory channels are populated with memory which greatly reduces the memory bandwidth of this memory configuration to about one fourth of the full potential memory bandwidth. It was measured at 28% of the full potential memory bandwidth.

Once again, the best way to increase the memory bandwidth of this configuration is by using more memory DIMMs. Four 16 GB RDIMMs would provide the same memory capacity while nearly doubling memory bandwidth.

To illustrate the effect of an unbalanced memory configuration, another way to arrange two memory DIMMs is to install them on two memory channels that are not on the same channel pair as in the 0:1:1:0:0:0:0 memory configuration shown in Figure 5.

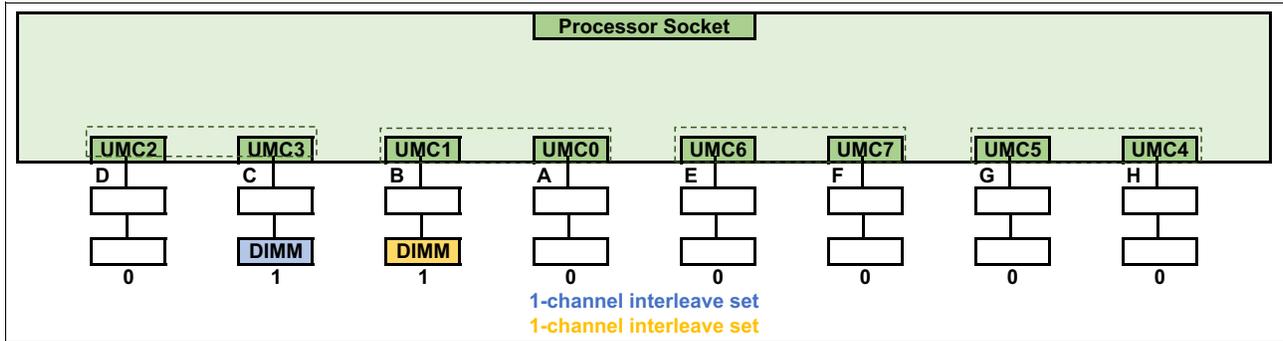


Figure 5 0:1:1:0:0:0:0 memory configuration relative memory bandwidth: 14%

This configuration does not follow balanced memory guideline 1 having memory populated on channels of different channel pair. Two 1-channel interleave sets are formed and this greatly reduced memory bandwidth performance. It was measured at 14%. This shows the importance of populating both memory DIMMs on two channels of the same channel pair for better performance.

The best way to improve memory bandwidth performance for this configuration even further is to increase the number of DIMMs and spread them across more channel pairs.

Configuration of 3 DIMMs - unbalanced

The best configuration for three memory DIMMs is to populate two DIMMs on the same channel pair. The third DIMM can be populated in any other channel. An example of this configuration is shown in Figure 6.

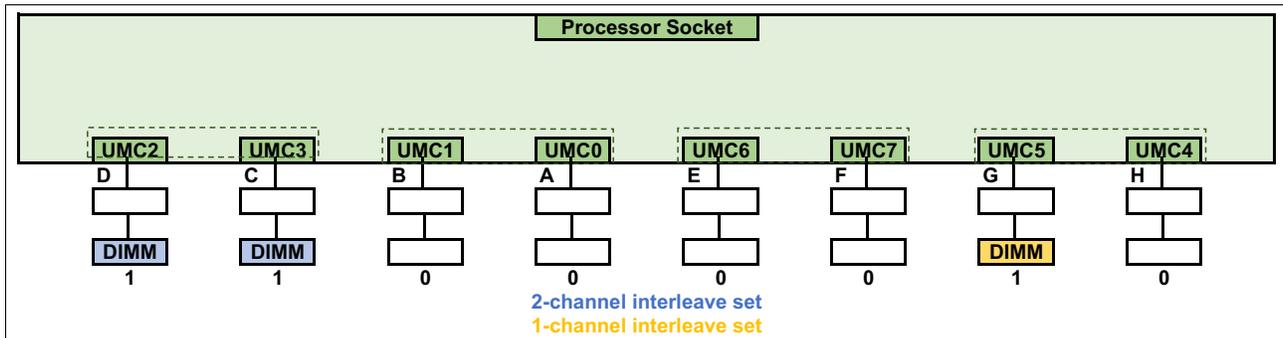


Figure 6 1:1:0:0:0:0:1:0 memory configuration relative average memory bandwidth: 24%, min memory bandwidth: 14%

This configuration results in one 2-channel interleave set and one 1-channel interleave set. Memory bandwidth performance for this configuration is not consistent depending on which interleave set is being accessed. Average bandwidth was measured at 24%. Worst case bandwidth was measured at 14%.

Another way to populate three memory DIMMs is to spread them on three memory channels on three different channel pair. An example of this configuration is 1:0:1:0:1:0:0 and is shown in Figure 7. This configuration does not follow guideline 1 and guideline 2 and results in non-optimal memory configuration. Three 1-channel interleave sets are formed, and bandwidth was measured at 14%.

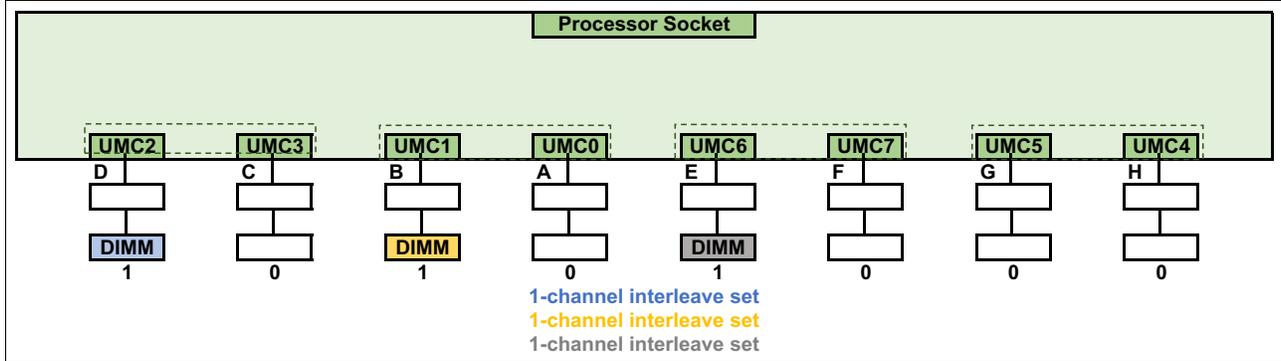


Figure 7 1:0:1:0:1:0:0 memory configuration relative memory bandwidth: 14%

Configuration of 4 DIMMs - balanced if installed correctly

Four memory DIMMs can be populated in the 1:1:0:0:0:0:1:1 memory configuration shown in Figure 8.

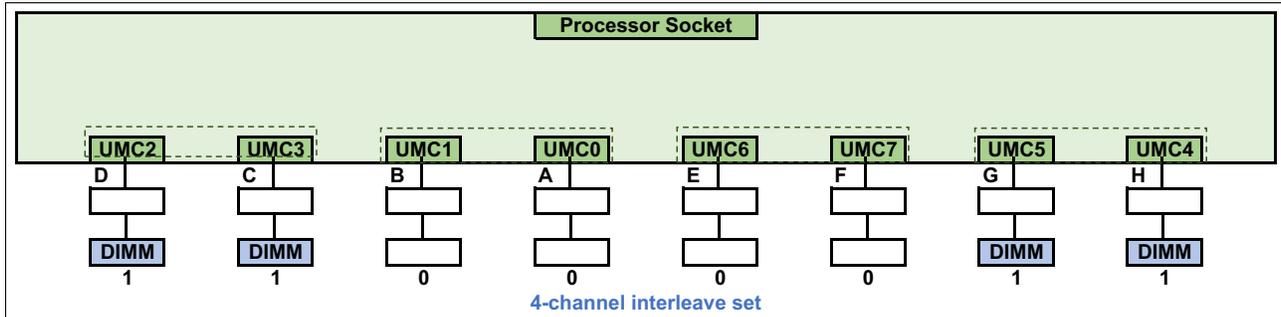


Figure 8 1:1:0:0:0:0:1:1 memory configuration relative memory bandwidth: 54%

This memory configuration follows balanced memory guideline 1 as each channel on the four populated channels have the same memory capacity and ranks. It also follows balanced memory guideline 2 as two populated channel pairs have identical DIMM configuration. This is a balanced memory configuration. A single 4-channel interleave set is formed across the four populated channels. Bandwidth was measured at 54%.

The best way to increase memory bandwidth for this configuration is to use more DIMMs and spread them across all memory channels.

Four memory DIMMs can also be populated in the 1:0:1:0:1:0:1:0 memory configuration as shown in Figure 9. This memory configuration follows balanced memory guideline 2, but not 1. It is an unbalanced memory configuration.

Four 1-channel interleave sets are formed. As seen before, more than one interleave set is detrimental to memory bandwidth. This memory configuration bandwidth was measured at 14% which is about one fourth the bandwidth of the balanced configuration shown above.

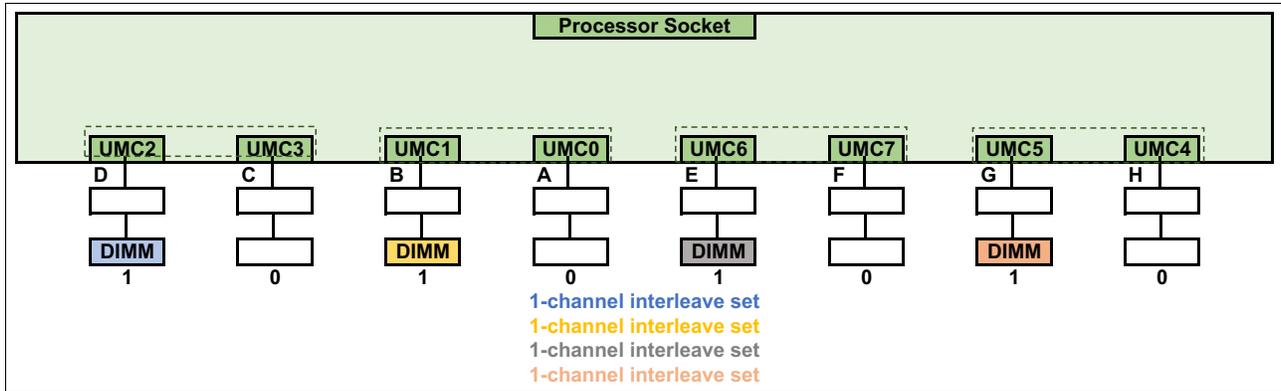


Figure 9 1:0:1:0:1:0:1:0 memory configuration relative memory bandwidth: 14%

Configuration of 5 DIMMs - unbalanced

Five memory DIMMs can be populated in a 1:1:0:1:0:0:1:1 configuration as shown in Figure 10. This configuration is not a balanced configuration, and it results in one 4-channel interleave set and one 1-channel interleave set. Average bandwidth for this configuration was measured at 30%. Worst case memory bandwidth was measured at 14%.

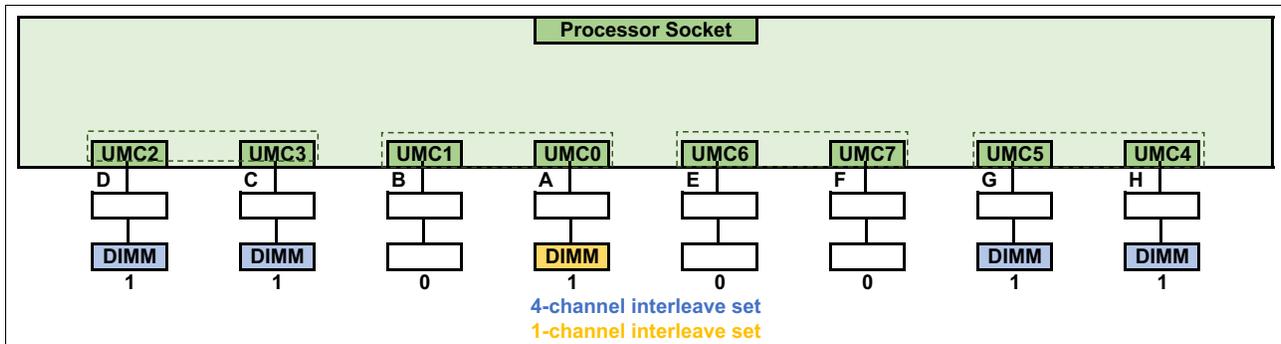


Figure 10 1:1:0:1:0:0:1:1 memory configuration relative average memory bandwidth: 30%, min memory bandwidth: 14%

Configuration of 6 DIMMs - unbalanced

Six memory DIMMs can be populated in a 1:1:1:1:0:1:1 configuration as shown in Figure 11. This configuration follows balanced memory guideline 1, but not guideline 2. This results in an unbalanced configuration. It results in one 4-channel interleave set and one 2-channel interleave set. Average bandwidth for this configuration was measured at 40%. Worst case memory bandwidth was measured at 28%.

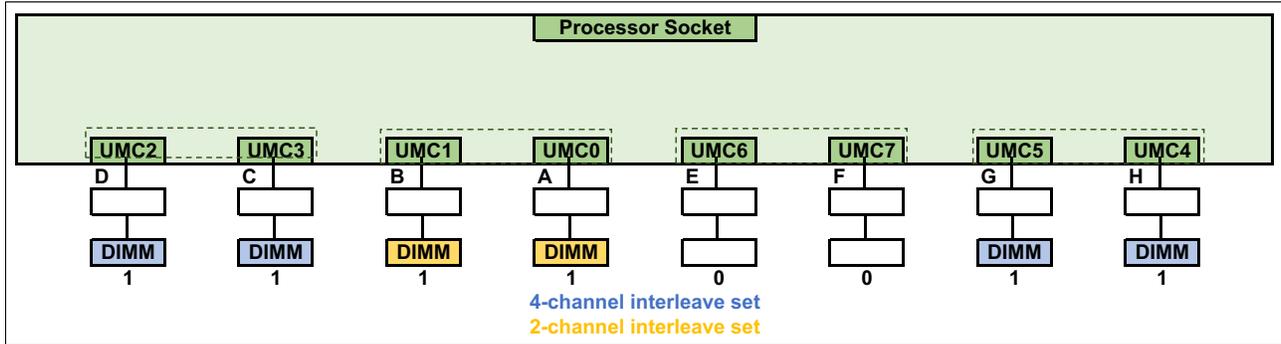


Figure 11 1:1:1:1:0:1:1 memory configuration relative average memory bandwidth: 40%, min memory bandwidth: 28%

EPYC Gen 3 added support to 6-channel memory interleaving when memory is populated in the 1:1:0:1:1:0:1:1 configuration as shown in Figure 12. While technically unbalanced when memory is populated in this specific configuration, a 6-channel interleave set can be formed, and memory bandwidth was measured at 71%.

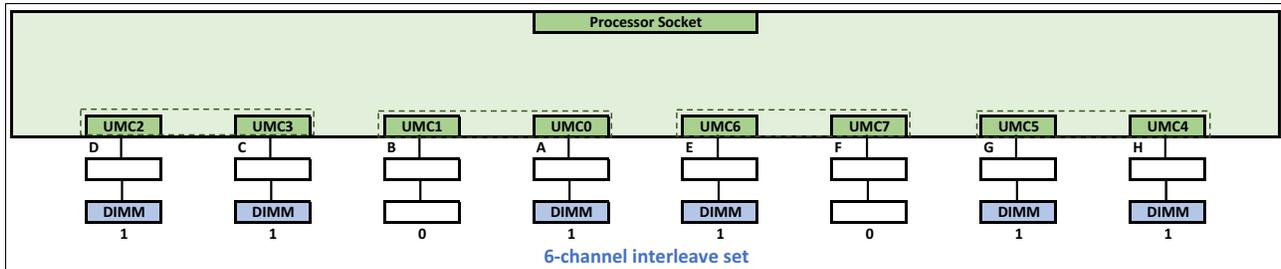


Figure 12 1:1:0:1:1:0:1:1 memory configuration relative memory bandwidth: 71% with EPYC Gen 3 processors

Configuration of 7 DIMMs - unbalanced

Seven memory DIMMs can be populated in the 1:1:1:1:1:0:1:1 memory configuration shown in Figure 13. This memory configuration does not follow balanced memory guideline 1 as only one channel was populated in channel pair EF. It also does not follow balanced memory guideline 2 as not all populated channel pairs have the same memory configuration. This is an unbalanced memory configuration.

This configuration results in one 4-channel interleave set, one 2-channel interleave set, and one 1-channel interleave set. Average bandwidth for this configuration was measured at 35%. Worst case bandwidth was measured at 14%.

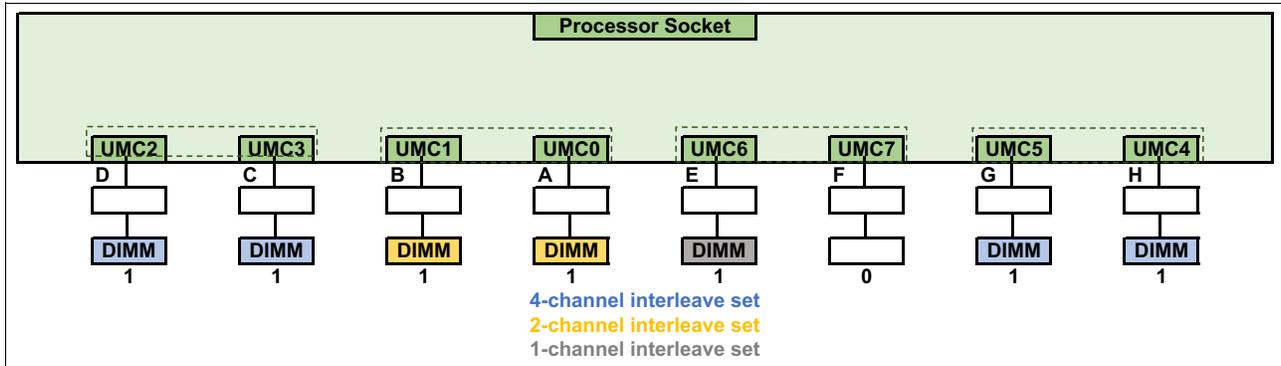


Figure 13 1:1:1:1:1:0:1:1 memory configuration relative average memory bandwidth: 35%, min memory bandwidth: 14%

Configuration of 8 DIMMs - balanced

Eight memory DIMMs can be populated in the 1:1:1:1:1:1:1:1 memory configuration shown in Figure 14. This memory configuration is a balanced configuration since all memory channels were populated with the same capacity and ranks. A single 8-channel interleave set is formed across all DDR channels and that results in a max memory bandwidth. Memory bandwidth was measured at 100%.

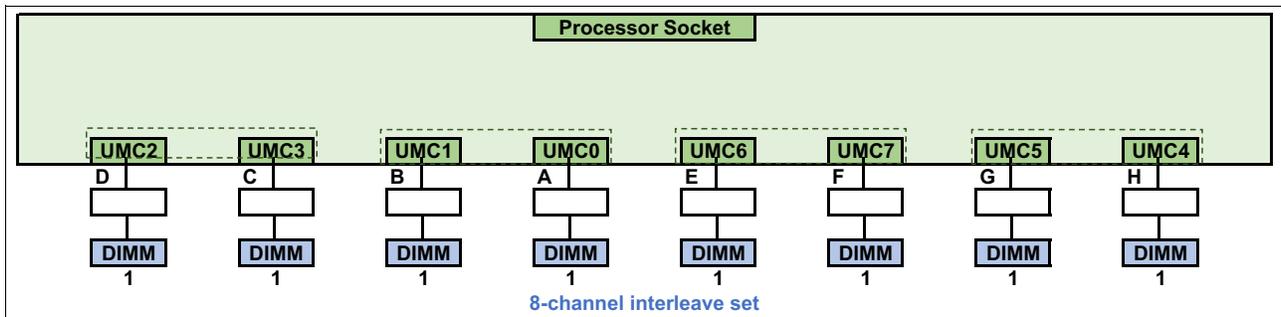


Figure 14 1:1:1:1:1:1:1:1 memory configuration relative memory bandwidth: 100%

Configuration of 9 DIMMs - unbalanced

Nine DIMMs can be populated in a 1:2:1:1:1:1:1:1 configuration as shown in Figure 15. Channel pairs AB, EF, and GH meet guideline 1, while channel pair CD does not. The configuration also does not follow guideline 2. This is an unbalanced configuration and multiple interleave sets are formed. Average memory bandwidth performance for this configuration was measured at 33% of the full bandwidth potential. Worst case bandwidth was measured at 14%.

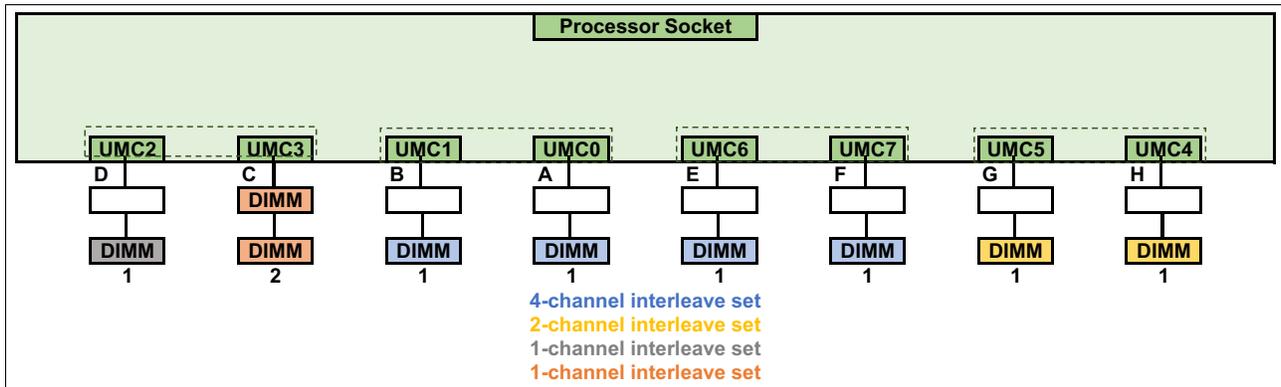


Figure 15 1:2:1:1:1:1:1:1 memory configuration relative average memory bandwidth: 33%, min memory bandwidth:14%

Configuration of 10 DIMMs - unbalanced

Ten DIMMs can be populated in a configuration as shown in Figure 16. All channel pairs have followed guideline 1 since both channels on the same pair have the same memory capacity. However, this configuration does not follow guideline 2 since channel pair CD has different memory configuration compared to AB, EF, and GH. As a result, one 4-channel interleave set and two 2-channel interleave sets are formed. Average memory bandwidth was measured at 38% of the max bandwidth potential. Worst case memory bandwidth was measured at 28%.

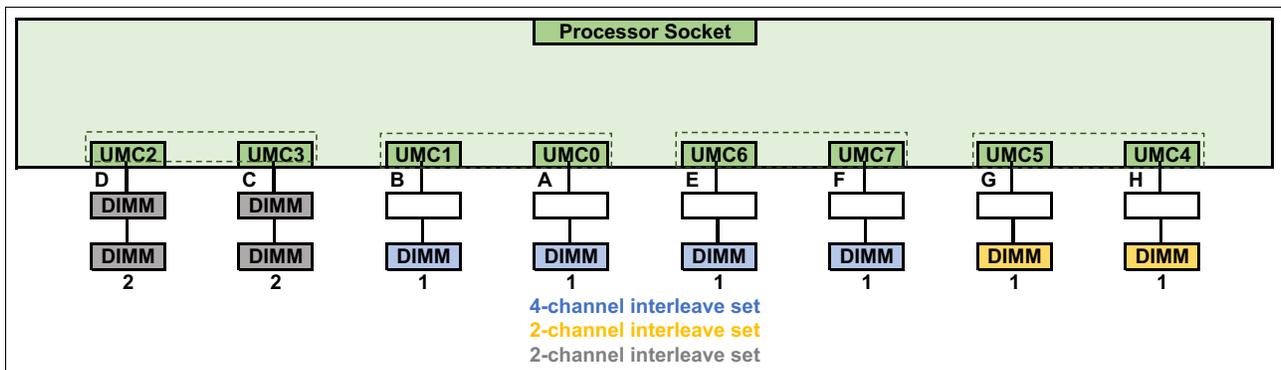


Figure 16 2:2:1:1:1:1:1:1 memory configuration relative average memory bandwidth: 38%, min memory bandwidth: 28%

Configuration of 11 DIMMs - unbalanced

Eleven DIMMs can be populated in a configuration as shown in Figure 17. Channel pairs AB, CD, and EF follow guideline 1 while GH does not. This is an unbalanced configuration, and as a result, multiple interleave sets are formed as shown in the figure below. Average memory bandwidth was measured at 31% of the max bandwidth potential. Worst case memory bandwidth was measured at 14%.

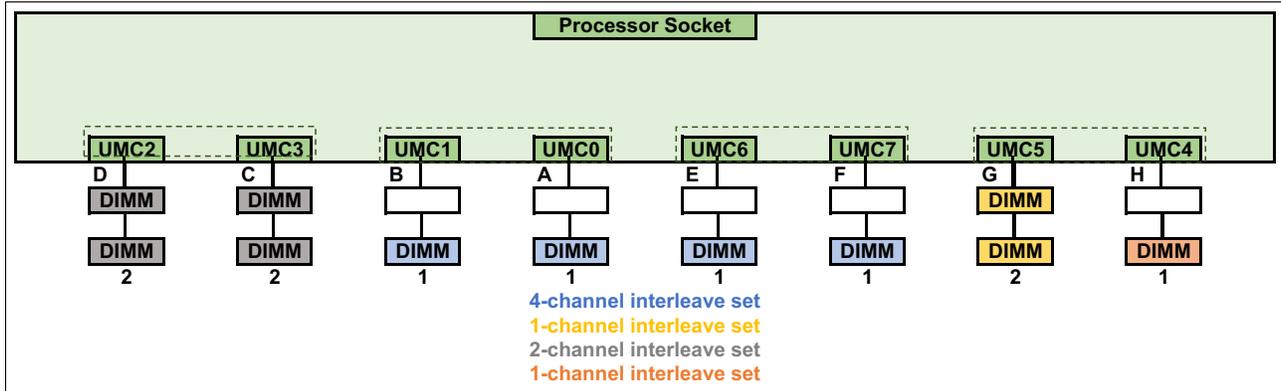


Figure 17 2:2:1:1:1:1:2:1 memory configuration relative average memory bandwidth: 31%, min memory bandwidth: 14%

Configuration of 12 DIMMs

Twelve DIMMs can be populated in a configuration as shown in Figure 18. All channel pairs follow guideline 1. However, guideline 2 is not followed as channel pairs do not have identical memory configuration. As a result, two 4-channel interleave sets are formed. Average memory bandwidth was measured at 35% of the max bandwidth potential. Worst case memory bandwidth was measured at 28%.

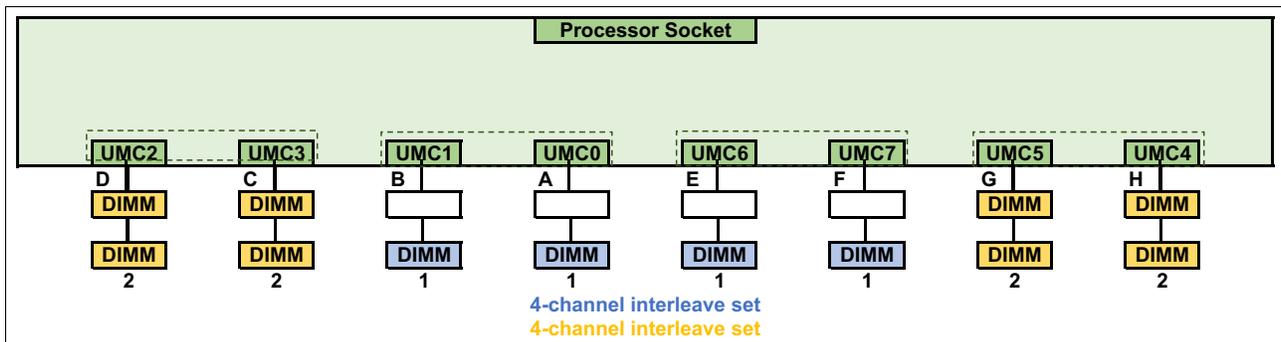


Figure 18 2:2:1:1:1:1:2:2 memory configuration relative average memory bandwidth: 35%, min memory bandwidth: 28%

With EPYC Gen 3 processors, the 2:2:0:2:2:0:2:2 memory configuration results in a single 6-channel interleave as shown in Figure 19 even though this configuration is technically unbalanced since channels AE are not a channel pair. Memory bandwidth was measured at 71%.

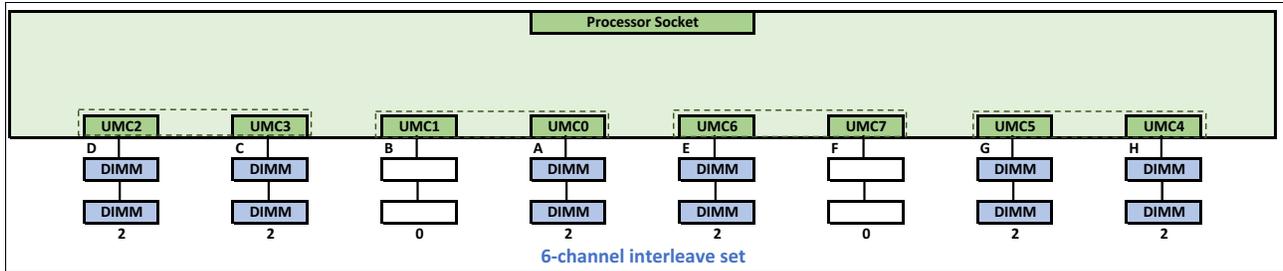


Figure 19 2:2:0:2:2:0:2:2 memory configuration relative memory bandwidth: 71% with EPYC Gen 3 processors

Configuration of 13 DIMMs - unbalanced

Thirteen DIMMs can be populated in a configuration as shown in Figure 20. Channel pairs CD, EF, and GH follow guideline 1 while AB does not. This is an unbalanced configuration, and as a result, multiple interleave sets are formed. Average memory bandwidth was measured at 34% of the max bandwidth potential. Worst case memory bandwidth was measured at 14%.

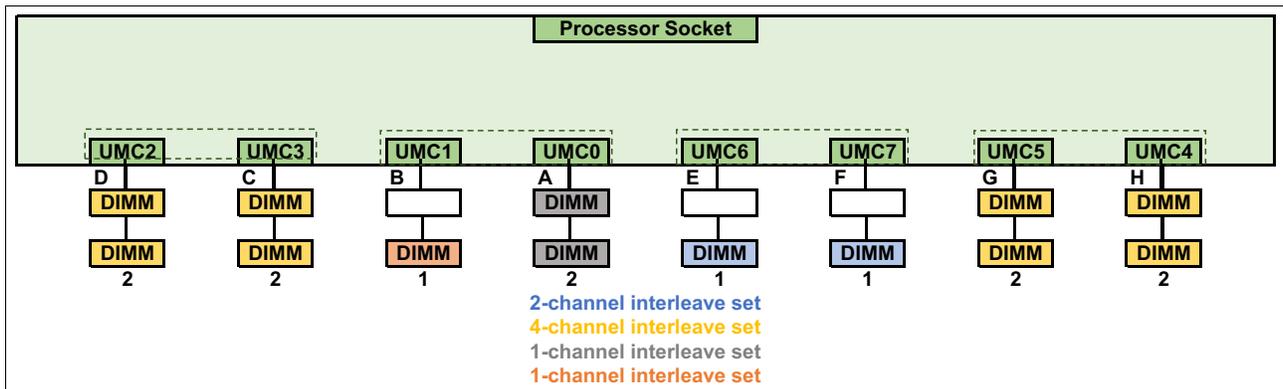


Figure 20 2:2:1:2:1:1:2:2 memory configuration relative average memory bandwidth: 34%, min memory bandwidth: 14%

Configuration of 14 DIMMs - unbalanced

Fourteen DIMMs can be populated in a configuration as shown in Figure 21. All channel pairs follow guideline 1. However, guideline 2 is not followed as channel pairs do not have identical memory configuration. As a result, multiple interleave sets are formed. Average memory bandwidth was measured at 35% of the max bandwidth potential. Worst case memory bandwidth was measured at 28%.

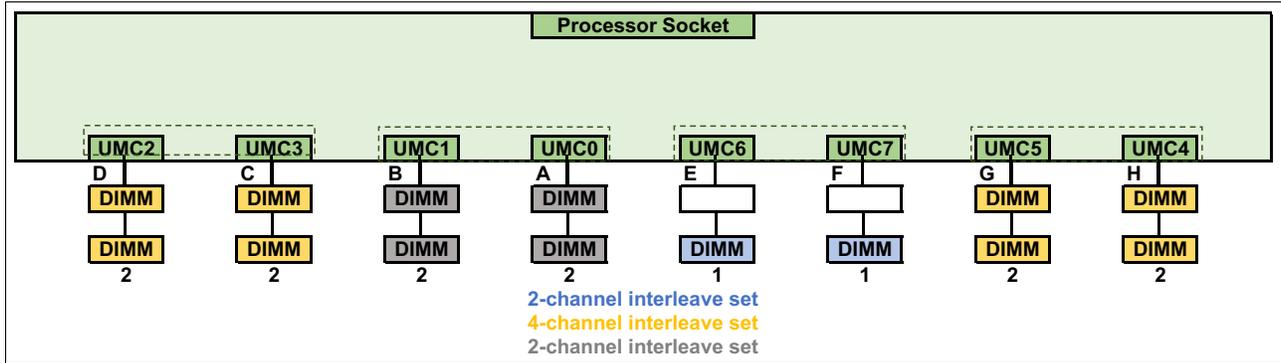


Figure 21 2:2:2:2:1:1:2:2 memory configuration relative average memory bandwidth: 35%, min memory bandwidth: 28%

Configuration of 15 DIMMs - unbalanced

Fifteen DIMMs can be populated in a configuration as shown in Figure 22. Channel pairs AB, CD, and GH follow guideline 1 while EF does not. This is an unbalanced configuration, and as a result, multiple interleave sets are formed. Average memory bandwidth was measured at 34% of the max bandwidth potential. Worst case memory bandwidth was measured at 14%.

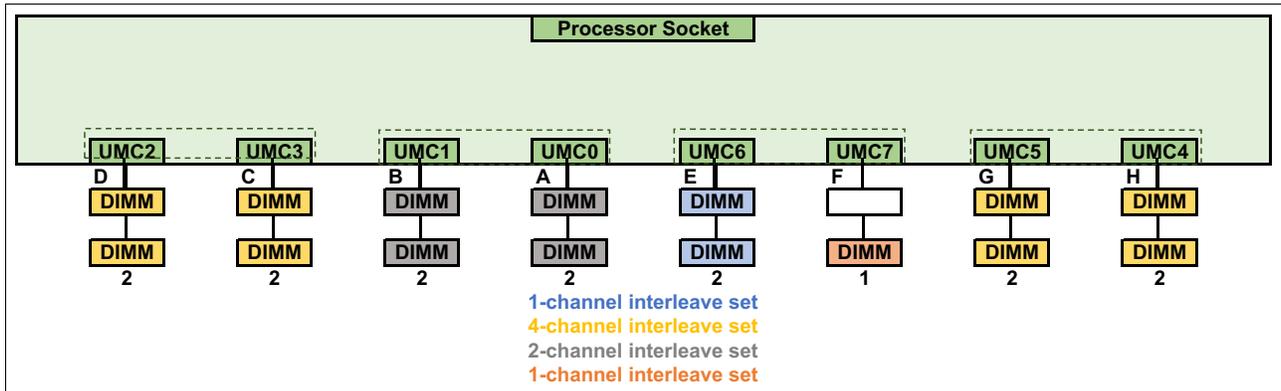


Figure 22 2:2:2:2:2:1:2:2 memory configuration relative average memory bandwidth: 34%, min memory bandwidth: 14%

Configuration of 16 DIMMs - balanced

Sixteen DIMMs can be populated in a configuration as shown in Figure 23. All channels have the same memory capacity so both guideline 1 and 2 are followed. One 8-channel interleave set is formed. This is a balanced memory configuration, and it yields the highest possible memory bandwidth. It was measured at 100% of the max bandwidth potential.

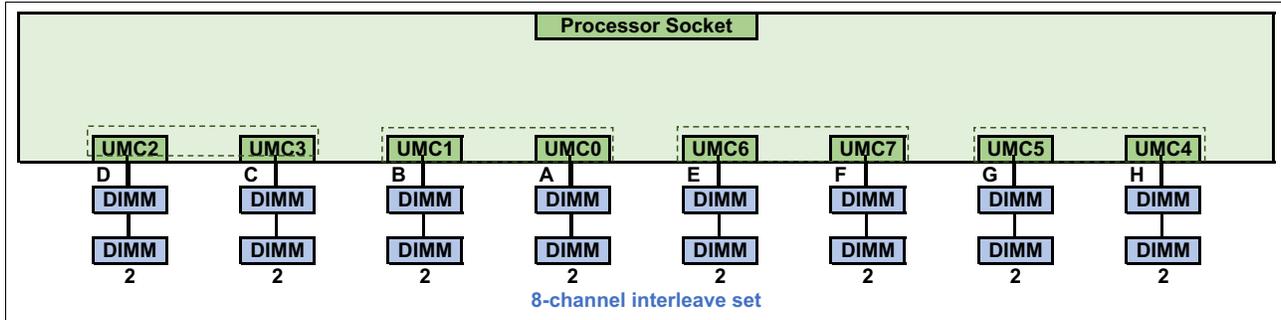


Figure 23 2:2:2:2:2:2:2:2 memory configuration relative memory bandwidth: 100%

Summary of the performance results

Table 2 shows a summary of the relative memory bandwidth of all the memory configurations that were evaluated. It also shows the number of interleave sets formed for each and whether it is a balanced or unbalanced memory configuration.

Table 2 Summary of all the memory configurations

Number of DIMMs Populated	Configuration	Number of Interleave Sets	Relative Performance	Balanced or Unbalanced
1	0:1:0:0:0:0:0	1	14%	Unbalanced
2	1:1:0:0:0:0:0	1	28%	Balanced
2	0:1:0:0:0:0:1:0	2	14%	Unbalanced
3	1:1:0:0:0:0:1:0	2	24%	Unbalanced
3	1:0:1:0:1:0:0:0	3	14%	Unbalanced
4	1:1:0:0:0:0:1:1	1	54%	Balanced
4	1:0:1:0:1:0:1:0	4	14%	Unbalanced
5	1:1:0:1:0:0:1:1	3	30%	Unbalanced
6	1:1:0:1:1:0:1:1	1	71%	Unbalanced ^a
6	1:1:1:1:0:0:1:1	3	40%	Unbalanced
7	1:1:1:1:1:0:1:1	4	35%	Unbalanced
8	1:1:1:1:1:1:1:1	1	100%	Balanced
9	1:2:1:1:1:1:1:1	5	33%	Unbalanced
10	2:2:1:1:1:1:1:1	4	38%	Unbalanced
11	2:2:1:1:1:1:2:1	5	31%	Unbalanced

Number of DIMMs Populated	Configuration	Number of Interleave Sets	Relative Performance	Balanced or Unbalanced
12	2:2:0:2:2:0:2:2	1	71%	Unbalanced ^a
12	2:2:1:1:1:1:2:2	2	35%	Unbalanced
13	2:2:1:2:1:1:1:2:2	4	34%	Unbalanced
14	2:2:2:2:1:1:2:2	3	35%	Unbalanced
15	2:2:2:2:2:1:2:2	4	34%	Unbalanced
16	2:2:2:2:2:2:2:2	1	100%	Balanced

a. Exception to the balanced memory configuration guidelines 1 and 2 for EPYC Gen 3 resulting in an unbalanced memory configuration, however, a 6-channel interleave set can be formed for best performance.

When using the same memory DIMM, only memory configurations with eight or sixteen memory DIMMs provide the full potential memory bandwidth. These are the best memory configurations for performance. Balanced memory configurations can also be achieved with two and four memory DIMMs, but they do not populate all the memory channels which reduces their memory bandwidth and performance.

Balanced memory configurations are the only memory configurations that should be used if memory bandwidth and performance are important.

Near-balanced memory configurations

It is not required that only identical memory DIMMs be used to achieve good memory bandwidth and performance. A mix of two different capacity memory DIMMs may be the most cost-effective way to produce the needed memory capacity and performance. Memory configurations that use only identical memory DIMMs and follow all the balanced memory guidelines are referred to as *balanced*. Memory configurations that use two different memory DIMMs and follow all the balanced memory guidelines are referred to as *near-balanced*.

Measurements show about a 3% loss in memory bandwidth when using a near-balanced memory configuration as compared to a balanced memory configuration but only if there are an even number of total memory ranks on each memory channel.

Maximizing memory bandwidth

To maximize the memory bandwidth of a server, the following rules should be followed:

1. Balance the memory across the processor sockets - all processor sockets on the same physical server should have the same configuration of memory DIMMs
2. Balance the memory across the memory controllers – all memory controllers on a processor socket should have the same configuration of memory DIMMs
3. Balance the memory across the populated memory channels - all populated memory channels should have the same total memory capacity and the same total number of ranks
4. Peak memory performance is achieved with 8 or 16 DIMMs per socket.

For optimal memory bandwidth:

1. Determine your needed memory capacity per socket,
2. Divide this memory capacity by eight to determine the minimum memory capacity needed per DDR channel
3. Round this calculated memory capacity up to the closest capacity available with 1 DIMM Per Channel (DPC) or 2 DPC combination, and
4. Populate your server with eight identical DIMM combinations derived from Step 3.

For example:

- ▶ If 512GB of total memory capacity is needed per socket, you can populate each socket with 16x 32GB DIMMs.
- ▶ If 768GB of total memory capacity is needed per socket, each DDR channel needs to be populated with $(768/8)=96$ GB of memory, which can be achieved using one 64GB DIMM + one 32GB DIMM. You can populate each socket with 8x 64GB DIMMs and 8x 32GB DIMMs. In this example, the following 2933 MHz DIMMs can be populated on each DDR channel:
 - ThinkSystem 32GB TruDDR4 2933MHz (2Rx4 1.2V) RDIMM-A
 - ThinkSystem 64GB TruDDR4 2933MHz (2Rx4 1.2V) RDIMM-A

Memory population rules: Please refer to product guide for memory population rules:

- ▶ ThinkSystem SR635:
<https://lenovopress.com/lp1160-thinksystem-sr635-server#memory-options>
- ▶ ThinkSystem SR655:
<https://lenovopress.com/lp1161-thinksystem-sr655-server#memory-options>
- ▶ ThinkSystem SR645:
<https://lenovopress.com/lp1280-thinksystem-sr645-server#memory-options>
- ▶ ThinkSystem SR665:
<https://lenovopress.com/lp1269-thinksystem-sr665-server#memory-options>

Summary

Overall server performance is affected by the memory subsystem which can provide both high memory bandwidth and low memory access latency when properly configured. Balancing memory across the memory controllers and the memory channels produces memory configurations which can efficiently interleave memory references among its DIMMs producing the highest possible memory bandwidth. An unbalanced memory configuration can reduce the total memory bandwidth to as low as 14% of a balanced memory configuration with 8 or 16 identical DIMMs installed per processor.

Implementing all three of the balanced memory guidelines described in this paper results in balanced memory configurations producing the best possible memory bandwidth and overall performance.

Lenovo recommends installing balanced memory populations with 2, 4, 8 or 16 DIMMs per processor. Peak memory performance is achieved with 8 or 16 DIMMs per processor.

For EPYC Gen 3 (EPYC 7003) processors, unbalanced 6-DIMM and 12-DIMM configurations support forming a single 6-channel interleave set. The 6-channel interleave set greatly improves memory bandwidth compared to the multiple interleave sets formed with alternate 6-DIMM and 12-DIMM configurations.

Change history

June 2021:

- ▶ Added AMD EPYC 7003 processors

Authors

This paper was produced by the following team of performance specialists:

Nathan Pham is a Senior Engineer in the Lenovo Infrastructure Solutions Group Performance Laboratory in Morrisville, NC. He has spent 19 years between IBM and Lenovo working on system performance. His focus has been in system performance architecture, modeling, analysis, and optimization. In his current role, he is responsible for all aspects of future Lenovo DCG system architecture and design for optimized performance. Nathan holds a Bachelor of Science degree in Computer Engineering and a Master of Science degree in Computer Engineering, both from North Carolina State University.

Joe Jakubowski is the Principal Engineer and Performance Architect in the Lenovo Infrastructure Solutions Group Performance Laboratory in Morrisville, NC. Previously, he spent 30 years at IBM. He started his career in the IBM Networking Hardware Division test organization and worked on various token-ring adapters, switches and test tool development projects. For more than 25 years, he has worked in the x86 server performance organization focusing primarily on database, virtualization and new technology performance. His current role includes all aspects of Intel and AMD x86 server architecture and performance. Joe holds Bachelor of Science degrees with Distinction in Electrical Engineering and Engineering Operations from North Carolina State University and a Master of Science degree in Telecommunications from Pace University.

Jamal Ayoubi is a Systems Performance Verification Engineer in the Lenovo Infrastructure Solutions Group Performance Laboratory in Morrisville, NC. Before his current position, Jamal interned with the Lenovo Workstation team where he primarily worked on competitive benchmarking and technical solutions. His current role includes CPU, Memory, and PCIe subsystem analysis and performance validation against functional specifications and vendor targets. Jamal holds Bachelor of Science degrees in Electrical Engineering and Computer Engineering from North Carolina State University.

Thanks to the following people for their contributions to this project:

- ▶ Arunkumar Krishnamoorthy
- ▶ David Watts

This paper is based in part on the Lenovo Press paper, *Balanced Memory Configurations with Second-Generation Intel Xeon Scalable Processors*. Thanks to the authors:

- ▶ Dan Colglazier
- ▶ Joseph Jakubowski
- ▶ Jamal Ayoubi

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on June 28, 2021.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/lp1268>

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available from <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
Lenovo XClarity™

Lenovo(logo)®
ThinkAgile™

ThinkSystem™

The following terms are trademarks of other companies:

Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.