Multi core Computer Architecture – Storage and Interconnects

Dr. John Jose
Department of Computer Science & Engineering
Indian Institute of Technology Guwahati, Assam

Tutorial - 06
Lecture – 29
Concepts in TCMP systems

So, welcome to the last tutorial session. Already we have completed the prescribed syllabus of Multi Core Computer Architecture Storage and the Interconnects. And this tutorial is the winding up session, where concepts of a distributed cache memories spreading over a tiled chip multi core processor is discussed. A couple of problems will give you better intuition regarding, how cache misses are handled in a Tiled Chip Multi Core Processor. We will move on to the first problem.

(Refer Slide Time: 00:57)

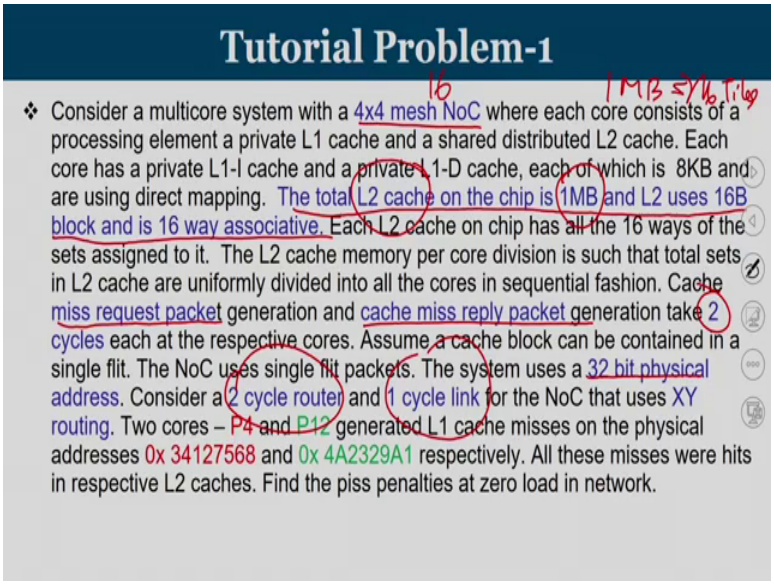

So, consider this problem, consider a multi core system with a 4 by 4 mesh NoC, where each core consists of a processing element, a private L1 cache and a shared distributed L2 cache. Each core has a private L1-I cache and a private L1-D cache, each of which is 8 k b and are using direct mapping. L2 cache on chip is 1MB, and L2 uses 16-byte block and it is 16 way associative.

Each L2 cache on chip has all 16 ways of the sets assigned to it. The L2 cache memory per core division is such that total sets in L2 cache are uniformly divided into all the

cores in sequential. Fashion cache miss request packet generation and caches miss reply generation take 2 cycles each at the respective cores.

Assume a cache block can be contained in a single flit. The NoC uses single flit packets the system uses a 32 bit physical addresses, consider a 2 cycle router and one cycle link for the NoC that uses x y routing. 2 cores P4 and P12 generated L1 cache misses on the physical addresses given. All these misses were hits in respective L2 caches, find the miss penalties at 0 load in the network.

So, let us try to summarize what this problem is all about. In this particular problem, they are given a 16 core tiled TCMP; where each of the tile has a process or a private L1 caches and shared L2 cache. Very time whenever the processor generates the address it goes to the private L1 cache let us I or D; whatever be the cache may be and if it is is a miss then it will go to L2. L2 is sheared and disturbed; that means, if there are 160 sets in the L2 cache for example, you have 16 tiles, so this one 16 is divided by 16.

So, each of the tile will hold the 10 sets each. So, you have set number 0 to 9 in tile 0, set number 10 to 19 in tile 1. Set number 20 to 29 in tile 3 like that. So, the entire sets in the L2 cache, rather than locating in 1 place, it is distributed at the same time it is shared.

So, L1 cache misses what is given in this question is going to generate packets in NoC, and this packets round trip travel from the source tile; which create a packet, travel all the way to the distance comeback and that is going to be the miss penalty. So, some properties of the network in terms of latency of the routers and many other parameters are given.

So, from this big paragraph of the question that is given, even though the story is big, let us try to break it down in to easily understandable format. So, if you look at this slide, I have made ample changes such that it is easy for you to grasp it. So, it is a 4 by 4 mesh NoC; that means, there are 16 tiles. And we are not really bothered about what is the specification of L1 cache, because whenever you get a miss in L1 cache the network comes into picture. So, the total L2 cache on chip is 1MB, and L2 uses 16-byte block and is 16 way associative.

So, L2 cache whatever you have it is total of 1MB. So, this 1MB is scattered across 16 tiles. So, that we have to further sub divide, and the L2 cache in general is 16-byte block

capacity and it is 16 way associative. Now this is big uniformly divided. Cache miss request that is time a packet is created, and the miss reply packet we will take 2 cycles each. And you assume that one cache block can be contained in a single flits, so 16 bytes will travel together. Now system use 32-bit physical address.

So, whatever is the physical address it is 32 bit. You have a 2 cycle router. So, it takes 2 cycle in each of the router and one cycle in the link. So, totally in order to move from one router to another, it is 2 plus 1 3 cycles it will take x y routing. And these are the addresses that is been given.

(Refer Slide Time: 05:09)

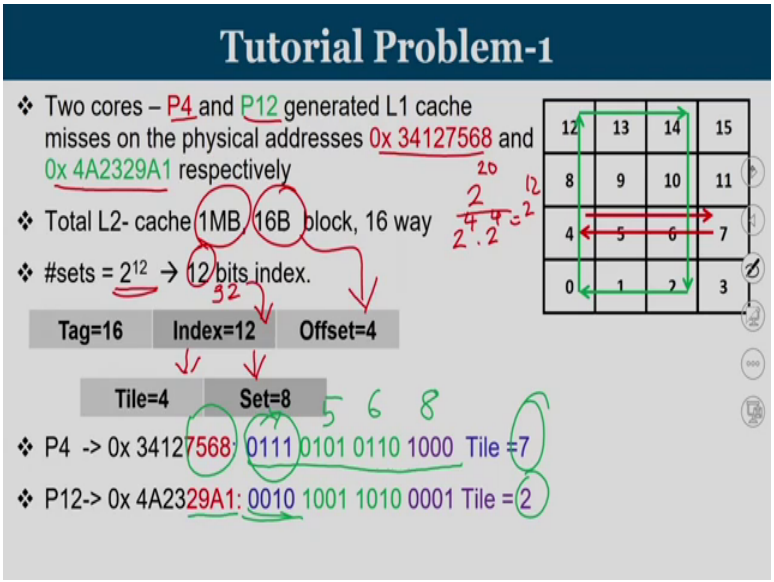

Let us try to summarize. There are 2 cores P4 is going to generate an address, and P12 is going to generate the another address.

So, this is the structure of the tile. So, whenever in these kind of problems, we have to use the standard numbering. Numbering starts from 0 at the bottom left, and 15 at the top right. That is already been discussed in our last tutorial, that is the standard numbering that we have to follow. Now for this particular case, let us try to find out what is the split up of the address. So, there are total L2 cache on the chip. It is 1MB, this 1MB is scattered across all the 16 tiles.

So, without considering the distributed nature, let us try to find how many sets are there in L2 cache. So, number sets in L2 cache is 2 power 12, from where will you get? It is

basically 2 power 20 is the cache size 1MB, divided by 16 by block. So, block size is 2 power 4, and associativity is also 2 power 4, and that will give u 2 power 12, so it is a 12-bit index.

Now, the address it is totally 32 bit is the physical address. This 32 bits is divided into 16-bit tag, 12-bit index and offset of 4 bit. So, you get this 12 bits directly you are going to use the index. Since I am using 16 byte as the offset these are the cache memory concepts we learned.

So, offset is going to be 4 bits and whatever is remaining the balance is tag, so you have 16 bits of tag. Now this 2 power 12 sets of the cache memory is distributed across 16 cores. So, out of the 2 power 12 sets, it is been represented by a 12 bit set number, this 12 bit set number has to be divided into most significant 4 bits is for the tiles, and the next 8 bits will tell within the tile you have 256 sets.

So, that is going to be the split up. So, the 32-bit physical address is divided into 16-bit tag, 12-bit index and 4 bit offset. This 12-bit index which represents the set number is again further sub divided into 4-bit tile number, the most significant 4-bit tile number, and the next 8 bit. So, when you have 1MB of cache, and 2 power 12 cache sets are there, this is divided across 16 tile; meaning, one tile can be accommodate 256 sets.

Now, our question is this is the address, P4 is going to miss on the address, it is hexadecimal address. 0 x 3412 7568 this is a 32-bit value; we know that whatever has be shown in this black color that represents the tag. The red portion is going to indicate the remaining.

So, the 7568 is our point of interest. So, 7568 and that consists of your index bit and the offset. So, the index bit and the offset bit 7568 alone is written here. 7 stands for 0111, 5 means 0101, this is 7, this is 5 this is 6 and this is 8.

So, you can see that the color difference this first blue indicates these are indicating the tile number. Then next 8 bits indicate within the tile what is the set number and the last 4 bits indicate what is the offset. So, from this this value is is representing as 9. So, this address corresponds to tile 7; meaning, this physical address 0 x 3412 7568 is mapped in tile number 7.

So, a packet which starts from tile number 4 will travel through 5 6 and reach 7, and then the reply packet travel from 7 6 5 and it is reaching back to 4. So, this is the way the packet is going to travel. Similarly, if you take the second address that is the miss that is generated from P12 for this address 4A23 29A4 take the index and the offset portion 29A1 correspond to 0010100110100001; that means, it is mapped to tile 2, so this value indicates tile 2.

So, the miss from P4 is meant to tile 7, and miss from P12 is meant for tile 2. So, P12 is in the top left corner and tile 2 is in the bottom side. So, as per x y routing packet will move from 12 to 13 to 14, and then to 10 6 and 2. And from 2 the packet is returning back 2 10 4 8 and reaching back to 12. So, this is we are trying to understand where is the miss starting, the misses are staring from core number 4, and it is destination is on core number 7. So, a packet has to travel from 4 to 7 and it has to come back.

Similarly, the second miss start from 12, it goes all the way to 2 by x y routing, from 2 it comes to back 12.

(Refer Slide Time: 10:24)

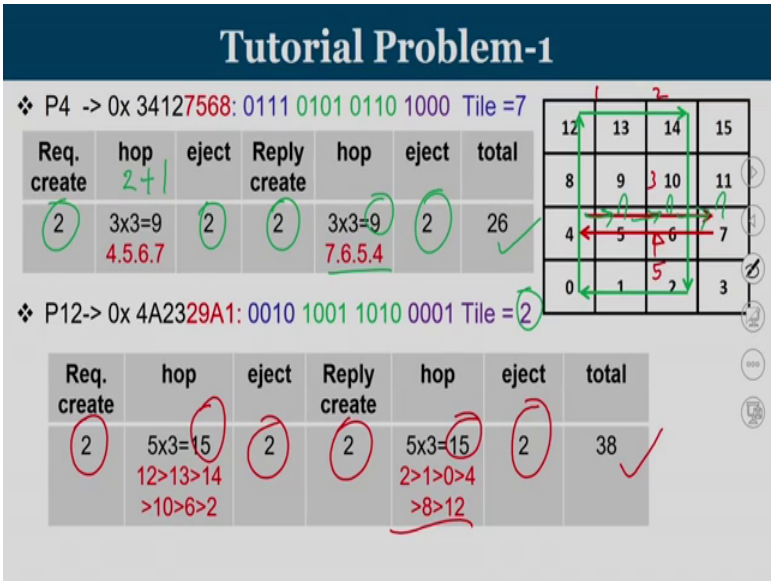

Now, consider the first miss, we have seen that it is on tile number 7. In the question it is mentioned that, packet creation, so a packet will be created at tile number 4 that takes 2 cycles, and this packet is going travel across this.

So, a packet from 4 we will take 2 cycle in 4 and 1 cycle to travel from 4 to 5. It takes 2 cycle in 5 and take one cycle to reach from 5 to 6. Again 2 cycle in 6, we will take one cycle to reach from 7. It takes 2 cycle in 7, that is routing and we see a location those things will be over, and then only the packet is going to be getting ejected. So, if you look it take 2 cycle for creation of the request, and then the packet will move from 4 to 5 from 5 to 6 and 6 to 7, that is 3 hope, each of the hope will take 3 cycles, that is 2 cycle in the router, plus 1 cycle in the link.

So, total 9 cycle it is spend in travelling from 4 to 7. Now in 7 it takes 2 cycles that is for the ejection process routing, and other thing will be done. It will come to know it has to be ejected. So, 2 cycles for eject, and then reply packet is created, and new packet is created at 7; which will take 2 cycles, the packet will be travelling in the x y direction from 7 to 6 to 5 to 4, again 3 hopes 9 cycles and then the packet has to be ejected.

So, you have total of 18 hopes 9 plus 9 18 hopes plus another 8, 26 cycle is the time that is needed for a packet to start from tile 4, go all the way to tile 7 create a new packet and come back. So, let us us try to summarize what we have done. You have a miss; we will first identify where this particular address is mapped on to. We came to now that a miss that occurred in 4 this particular address was mapped in tile number 7.

So, upon encountering a miss, the packet has to be created at 4, 2 cycles packet has to travel from 4 to 7, that is 3 hopes, each hope will take 3 cycles each; that means, 2 cycles in the router and one cycle in the link. So, 3 hopes 3 cycle each 9 hopes the packet will reach from 4 to 7. At 7, it will taking the 2 cycle delay of the router to get there is no more link travel only operation in the router. So, let us call the ejection it takes 2 cycles. Now the reply packet is created take another 2 more cycles, it has to travel 3 hopes again 9 cycles. It has to be getting ejected at 4, 2 cycle. So, that is makes the total 26.

So, the miss penalty means it is 26 cycle is the miss penalty for this. Now consider the second one we have processor P12, it is given an address and we know that it is mapped to the this particular address is mapped to tile number 2. Now let us see what happens in tile number 2, packets travels from 12 to 2. So, it takes 2 cycle for request creation, and it is having 5 hopes, 12 to 3 is 1, 3 to 4 is 2, 4 to 10 is 3, then 10 to 6 is 4, and 6 to 2 is 5. So, it takes 5 hopes to reach to 12 to 13 to 14 to 10 to 6 to 2. So, 5 into 3, 15 cycles 2 cycle like in the previous case 2 cycle to eject the packet, 2 cycle to create the reply

packet, again 5 hopes travelled that is 15. And this is the path in which the packet is travelling, 2 cycle for ejection, that makes it total of 38 cycles.

So, the cache miss from tile 12 which is mapped on to tile 2, we will take 38 clock cycles to reach the destination. So, in this way miss penalty in a TCMP system, because the cache memory L2 cache memory is set and distributed, miss penalty will also varies. Certain misses we will go to nearby cores and come back miss penalty is less. Certain misses will go to further core, so miss penalty is going to be relatively a bit higher.

(Refer Slide Time: 14:40)

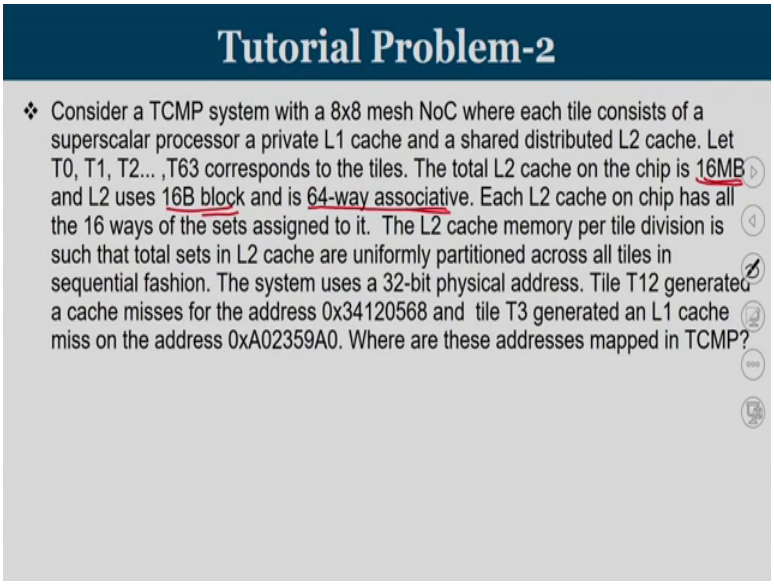

I hope you have clearly understood the concepts of distributed cache structure in TCMP try to work on with more problems that is given. Now we have a another tutorial problem. Consider a TCMP system with a 8 by 8 mesh, this is a larger mesh 64 core where each tile consists of a superscalar processor a private L1 cache and a shared distributed L2 cache. Tiles are numbered as T 0 to T 63.

Now, total L2 cache here is 16MB here is the difference. So, 16MB L2 cache, 16-byte block and 64 way associative. So, there is a small change here. It is 16-byte block and 64 way associative.

(Refer Slide Time: 15:26)

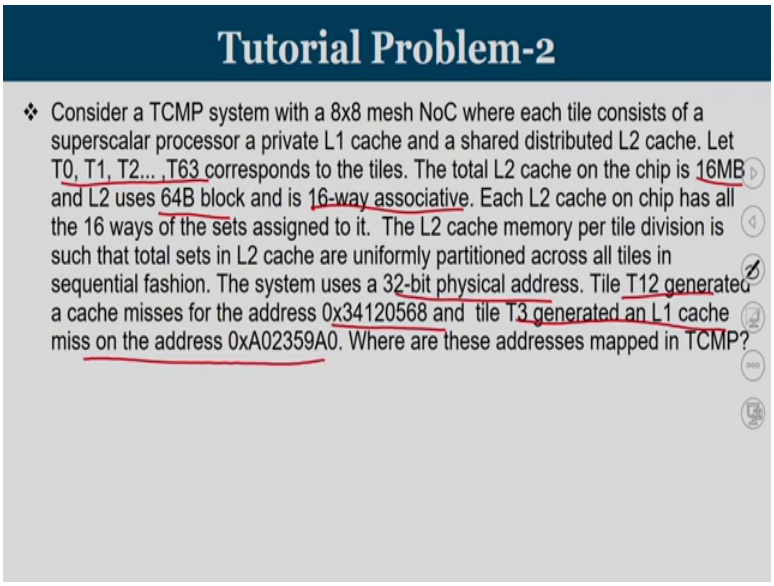

So, come to the second problem, problem number 2. Consider a TCMP system with an 8 by 8 mesh NoC where each tile consist of a superscalar processor a private L1 cache and a shared distributed L2 cache. The tiles are numbered as T 0 to T 63, tiles are numbered as T 0 to T 63.

The tiles are numbered as T 0 to T 63, and the L2 cache is totally 16MB 64-byte block and it is 16 way associative. Similar to the previous problem, the L2 sets are been uniformly distributed in a sequential fashion across all the tiles. System uses 32 bit of physical address, and tile T1 generated a cache miss for the address this, and tile 3 generated an L1 cache miss on this address. Where are these addresses mapped in the TCMP?

(Refer Slide Time: 16:21)

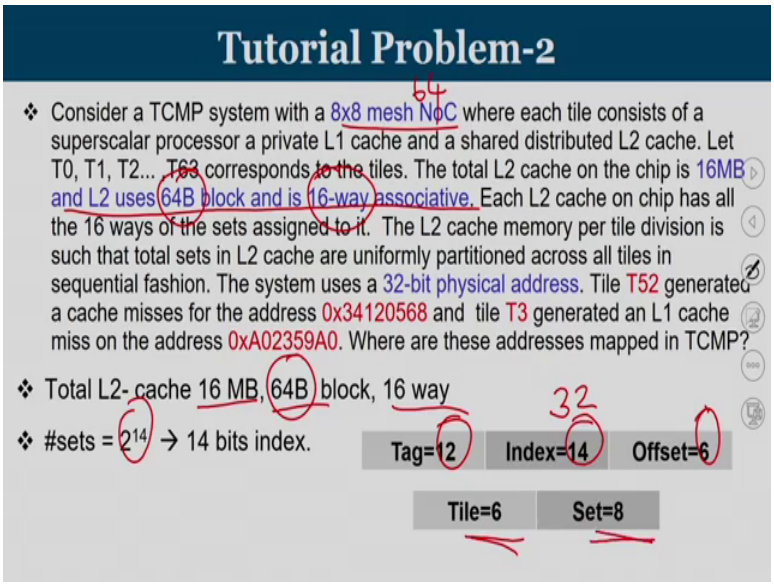

So, these are the important phrases in the question. It is a 64 core setup and your L2 cache totally 16MB, and this 16MB is shared and distributed across that. So, you have a 16-way associative cache that is there.

So, let us try to rephrase the important concepts given in this questions. So, I have change the colors, 8 by 8 mesh; that means, it is a 64 core and you have an L1 L2 cache size is L2 cache size is given it is 64-byte block, and 16 way associative, and you have 2 addresses that is there.

So, I have total of 16MB of L2 cache with a block size of 64 bytes, and 16 way associative. That will tell you total of 2 power 14 sets are there the same division cache size divided by block size divided by associativity. So, for a 14-bit index, the entire 32 bit in the physical address is divided into 14 bits for index, since I have 64-byte block offset is going to be 6 bit and the remaining is tag. So, it is a 12-bit tag 14-bit index and 6 bit offset.

Now, you have 2 power 14 sets are there, these 2 power 14 sets has to be distributed across 64 tiles. So, most significant 6 bits of the 14-bit index will be telling the tile number. So, 6-bit tile number and remaining 8 bits will tell you how many sets are there within a tile.

(Refer Slide Time: 17:52)

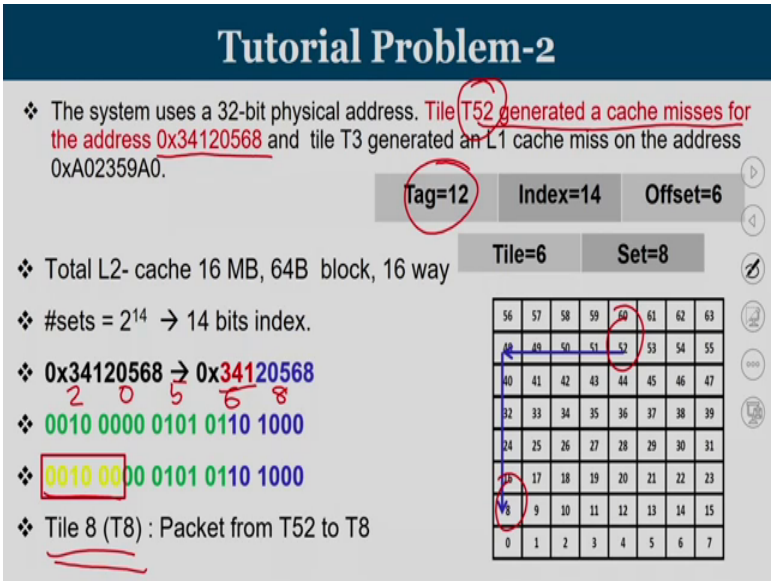

So, from the address now we have to split up system uses 32-bit address. So, T 52 is going to generate a cache miss for the address 0 x 3412 0568. This is what we have seen. So, 3412 0568 for the address it is 12 bits is for the tag. So, whatever you see in the red color that is tag, the blue portion indicates an index and the offset.

So, 2 0568 when I am going to write this this is 200102056 and 8, when you write that in binary, the green portion indicates the index. And the last 6 bit indicates the offset. Now out of the index my point of interest is only the first 6 bit.

So, those first 6 bit if you take, the value is 800100. Meaning, the miss that is generated from T 52 is 2 tile 8 the packet is moving from tile 52 to tile 8. This is what this tile number 52 and this is going to be tile number 8. So, the packet is travelling from T 52 to T 8 that is a mapping that is been given.

(Refer Slide Time: 19:05)

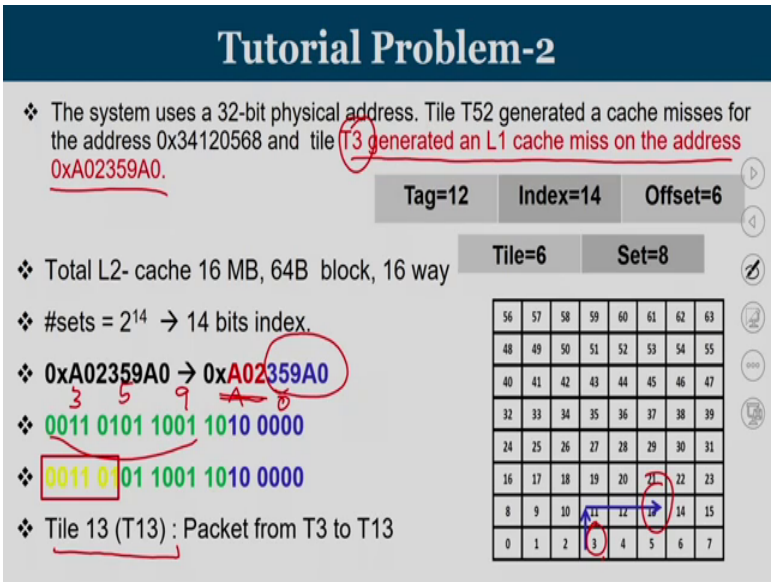

Now, the second portion of the question is request from T 3 is on this address 0 x A02359A0, when I write the address, the red portion indicates the tag bit the blue portion indicates the index and offset bits, I am going to write the expanded form. So, it is 3659A and 0. The green portion in this address indicates the index portion. This index is further divided into first 6 bit is going to be the tile number, and the last 8 bit is going to be the set number. Tile number indicates it is 13, 001101 is tile number 13. So, this request is from tile number T 3.

So, from 3 the request is going to tile number 13. So, this particular address is mapped to tile number 13, and the previous one was mapped to tile number 8. So, this is the way how things are been done. So, this question is all about tile mapping. So, today we have seen 2 different problem; one is about finding bout miss penalty and this is a relatively larger TCMP system 64 cores. And here we have to find out what is the address mapping to which tile it is been mapped.

So, this completes our todays tutorial. So, with this fair treatment of the topics on TCMP system which is given. So, I request all the learners in this course to kindly solve similar kind of problems so that you get a grip on how address mappings are done. This topics are very important. Kindly refer to similar questions before the final exams, and wishing you good luck for the exams.

Thank you.