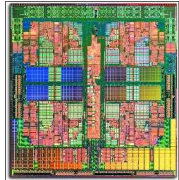



1

Caching Issues in Multicore Performance

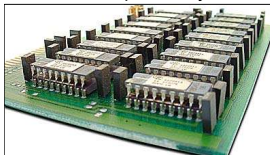
CPU Chip







Oregon State University
Mike Bailey
mjb@cs.oregonstate.edu


Off-chip Memory







This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/)



Oregon State University
Computer Graphics

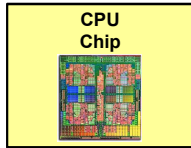
cache.pptx

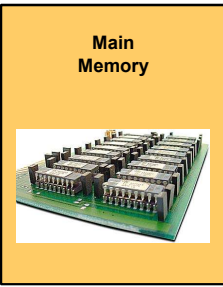
mjb - March 9, 2023

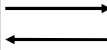
2

Problem: The Path Between a CPU Chip and Off-chip Memory is Slow

CPU Chip








This path is relatively slow, forcing the CPU to wait for up to 200 clock cycles just to do a store to, or a load from, memory.

Depending on your CPU's ability to process instructions out-of-order, it might go idle during this time.

This is a *huge* performance hit!



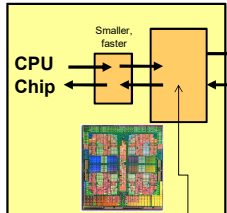
Oregon State University
Computer Graphics

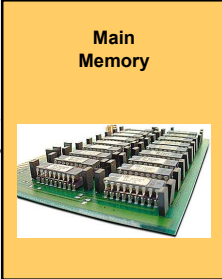
mjb - March 9, 2023


3

Solution: Hierarchical Memory Systems, or "Cache"

CPU Chip








The solution is to add intermediate memory systems. The one closest to the ALU (L1) is small and fast. The memory systems get slower and larger as they get farther away from the ALU.

L3 cache also exists on some high-end CPU chips



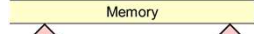
Oregon State University
Computer Graphics

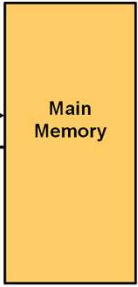
mjb - March 9, 2023

4


Cache and Memory are Named by "Distance Level" from the ALU

Memory

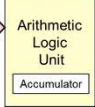




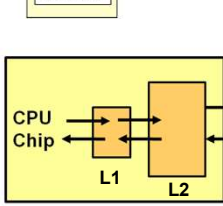
Control Unit

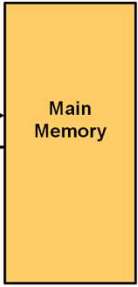


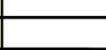
Arithmetic Logic Unit
Accumulator




CPU Chip







L3 cache also exists on some high-end CPU chips



Oregon State University
Computer Graphics

mjb - March 9, 2023

Storage Level Characteristics

	L1	L2	L3	Memory	Disk
Type of Storage	On-chip	On-chip	On-chip	Off-chip	Disk
Typical Size	100 KB	8 MB	32 MB	32 GB	Many GBs
Typical Access Time (ns)	.25	.50	10.8	50	5,000,000
Scaled Access Time	1 second	2 seconds	43 seconds	3.3 minutes	231 days
Managed by	Hardware	Hardware	Hardware	OS	OS

Adapted from: John Hennessy and David Patterson, *Computer Architecture: A Quantitative Approach*, Morgan-Kaufmann, 2007. (4th Edition)

Usually there are two L1 caches – one for Instructions and one for Data. You will often see this referred to in data sheets as: "L1 cache: 32KB + 32KB" or "I and D cache"



mjb - March 9, 2023

Cache Hits and Misses

When the CPU asks for a value from memory, and that value is already in the cache, it can get it quickly.

This is called a **cache hit**

When the CPU asks for a value from memory, and that value is not already in the cache, it will have to go off the chip to get it.

This is called a **cache miss**

While cache might be multiple kilo- or megabytes, the bytes are transferred in much smaller quantities, each called a **cache line**. The size of a cache line is typically just **64 bytes**.



Performance programming should strive to avoid as many cache misses as possible. That's why it is very helpful to know the cache structure of your CPU.



mjb - March 9, 2023

Spatial and Temporal Coherence

Successful use of the cache depends on **Spatial Coherence**:

"If you need one memory address's contents now, then you will probably also need the contents of some of the memory locations around it soon."

Successful use of the cache depends on **Temporal Coherence**:

"If you need one memory address's contents now, then you will probably also need its contents again soon."

If these assumptions are true, then you will generate a lot of cache hits.


If these assumptions are not true, then you will generate a lot of cache misses, and you end up re-loading the cache a lot.



mjb - March 9, 2023

How Bad Is It? -- Demonstrating the Cache-Miss Problem

C and C++ store 2D arrays a row-at-a-time, like this, $A[i][j]$:



For large arrays, would it be better to add the elements by row, or by column? Which will avoid the most cache misses?

```
sum = 0;
for( int i = 0; i < NUM; i++ )
{
    for( int j = 0; j < NUM; j++ )
    {
        float f = ???
        sum += f;
    }
}
```

Sequential memory order → float f = Array[i][j] ;

Jump-around-in-memory order → float f = Array[j][i] ;



mjb - March 9, 2023

Demonstrating the Cache-Miss Problem – Across Rows

9

```
#define NUM 10000
float Array[NUM][NUM];
double MyTimer( );

int
main( int argc, char *argv[ ] )
{
    float sum = 0.;
    double start = MyTimer( );
    for( int i = 0; i < NUM; i++ )
    {
        for( int j = 0; j < NUM; j++ )
        {
            sum += Array[ i ][ j ];    // access across a row
        }
    }
    double finish = MyTimer( );
    double row_secs = finish - start;
```



mjb - March 9, 2023

Demonstrating the Cache-Miss Problem – Down Columns

10

```
float sum = 0.;
double start = MyTimer( );
for( int i = 0; i < NUM; i++ )
{
    for( int j = 0; j < NUM; j++ )
    {
        sum += Array[ j ][ i ];    // access down a column
    }
}
double finish = MyTimer( );
double col_secs = finish - start;
```





mjb - March 9, 2023

Demonstrating the Cache-Miss Problem

11

Time, in seconds, to compute the array sums, based on by-row versus by-column order:



mjb - March 9, 2023

Array-of-Structures vs. Structure-of-Arrays:

12

```
struct xyz
{
    float x, y, z;
} Array[N];

float X[N], Y[N], Z[N];
```

X0
Y0
Z0
X1
Y1
Z1
X2
Y2
Z2
X3
Y3
Z3

X0
X1
X2
X3
...
Y0
Y1
Y2
Y3
...
Z0
Z1
Z2
Z3
...

1. Which is a better use of the cache if we are going to be using X-Y-Z triples a lot?
2. Which is a better use of the cache if we are going to be looking at all X's, then all Y's, then all Z's?

I've seen some programs use a "Shadow Data Structure" to get the advantages of both AOS and SOA



mjb - March 9, 2023

Computer Graphics is often a Good Use for Array-of-Structures:

X0	struct xyz { float x, y, z; } Array[N]; ...
Y0	
Z0	
X1	glBegin(GL_LINE_STRIP); for(int i = 0; i < N; i++) { glVertex3f(Array[i].x, Array[i].y, Array[i].z); } glEnd();
Y1	
Z1	
X2	
Y2	
Z2	
X3	
Y3	
Z3	



mjb - March 9, 2023

A Good Use for Structure-of-Arrays:

X0	float X[N], Y[N], Z[N]; float Dx[N], Dy[N], Dz[N]; ... for(int i = 0; i < N; i++) { Dx[i] = X[i] - Xnow; Dy[i] = Y[i] - Ynow; Dz[i] = Z[i] - Znow; }
X1	
X2	
X3	
...	
Y0	
Y1	
Y2	
Y3	
...	
Z0	
Z1	
Z2	
Z3	
...	




mjb - March 9, 2023

Good Object-Oriented Programming Style can sometimes be Inconsistent with Good Cache Use:

```
class xyz
{
public:
    float x, y, z;
    xyz *next;
    xyz( );
    static xyz *Head = NULL;
};

xyz::xyz( )
{
    xyz * n = new xyz;
    n->next = Head;
    Head = n;
};
```

This is good OO style – it encapsulates and isolates the data for this class. Once you have created a linked list whose elements are all over memory, is it the best use of the cache?




mjb - March 9, 2023

But, Here Is a Compromise:

It might be better to create a large array of xyz structures and then have the constructor method pull new ones from that list. That would keep many of the elements close together while preserving the flexibility of the linked list.

When you need more, allocate another large array and link to it.



mjb - March 9, 2023

But, Here Is a Compromise:

```

#include <stdio>
#define NUMALLOC 1024

struct node
{
    float data;
    bool canBeDeleted;
    struct node *next;
};
struct node *Head = NULL;



struct node *
GetNewNode()
{
    if( Head == NULL )
    {
        struct node *array = new struct node[NUMALLOC];
        Head = &array[0];
        for( int i = 0; i < NUMALLOC - 1; i++ )
        {
            array[i].canBeDeleted = false;
            array[i].next = &array[i+1];
        }
        array[NUMALLOC-1].next = NULL;
    }
    struct node *p = Head;
    Head = Head->next;
    return p;
}

void
DeleteNode( struct node *n )
{
    n->canBeDeleted = true;
}

```

Remember: in this scheme, you cannot delete an individual node because it was allocated as part of an array. The best you can do is track which nodes can be deleted and then when all of an array's nodes are flagged, delete the whole array.

mjb - March 9, 2023



An Example of Where Cache Coherence Really Matters: Matrix Multiply

The usual approach is multiplying the entire A row * entire B column
This is equivalent to computing a single dot product

Row i of A * Column j of B → Element (i,j) of C

```

for( i = 0; i < SIZE; i++ )
    for( j = 0; j < SIZE; j++ )
        for( k = 0; k < SIZE; k++ )
            Sum and store

```

$\sum A[i][k]$
*
 $B[k][j]$
→
 $C[i][j]$

Problem: Column j of the B matrix is not doing a unit stride

mjb - March 9, 2023

An Example of Where Cache Coherence Really Matters: Matrix Multiply

Scalable Universal Matrix Multiply Algorithm (SUMMA)
Entire A row * one element of B row
Equivalent to computing one item in many separate dot products

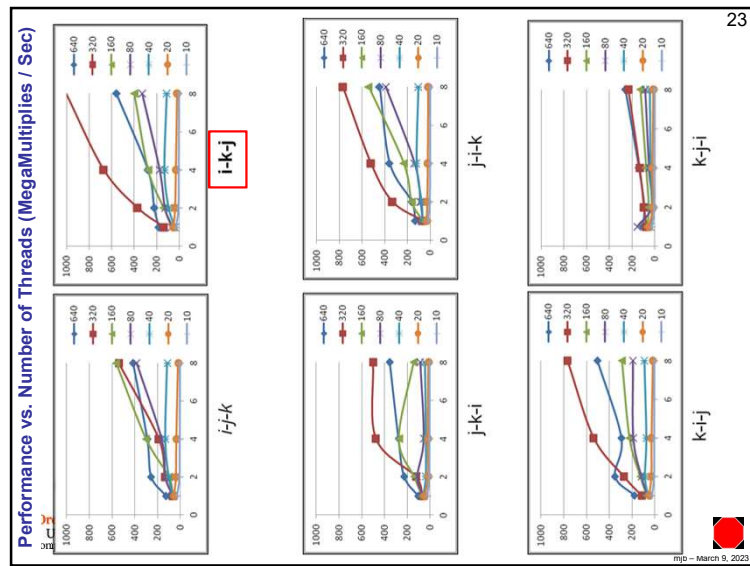
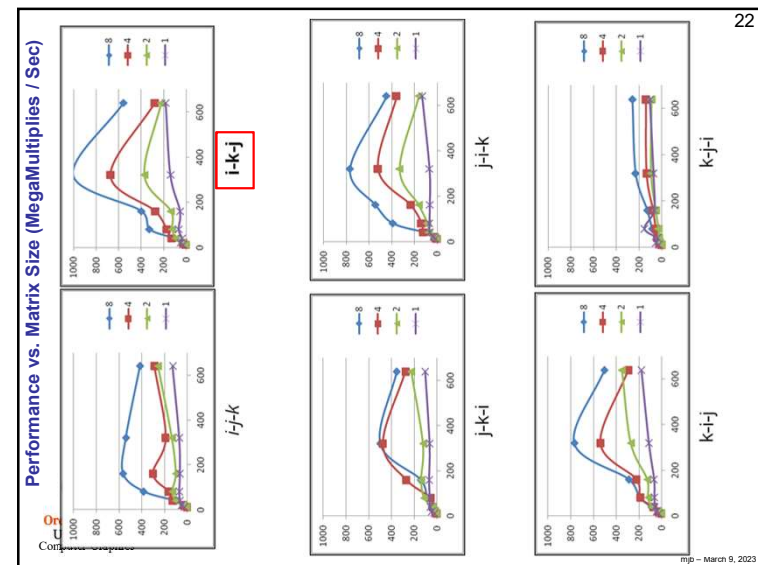
Row i of A * Row k of B = Element (i,j) of C

```

for( i = 0; i < SIZE; i++ )
  for( k = 0; k < SIZE; k++ )
    for( j = 0; j < SIZE; j++ )
      A[i][k] * B[k][j]
  
```

Add to C[i][j]

mjb - March 9, 2023



Cache Architectures

N-way Set Associative – a cache line from a particular block of memory can appear in a limited number of places in cache.
Each "limited place" is called a **set** of cache lines. A set contains N cache lines.

The memory block can appear in any cache line in its set.

Most Caches today are N-way Set Associative
N is typically 4 for L1 and 8 or 16 for L2

Cache:

Sets of Cache Lines

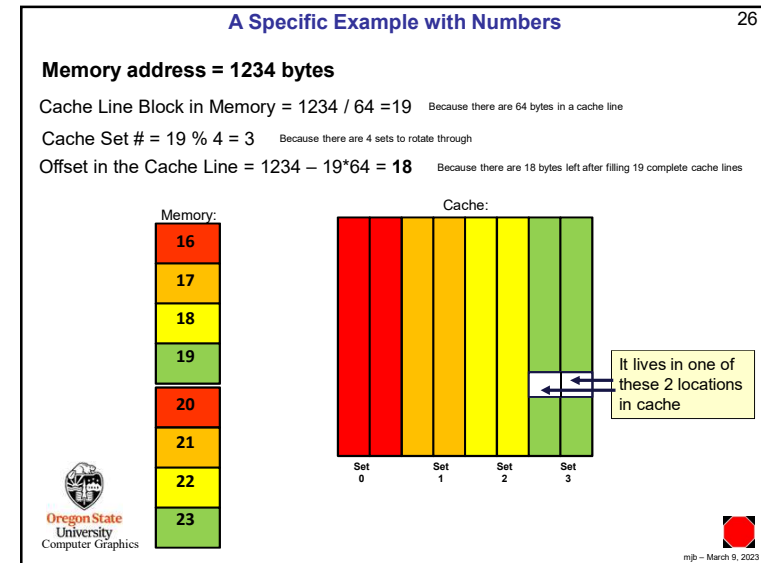
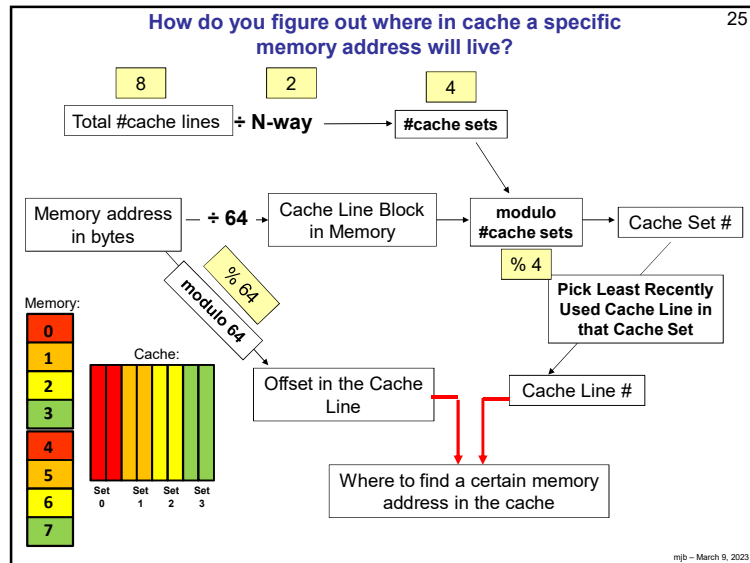
64 bytes

Memory:

Cache line blocks in memory (the numbers) and what cache line set they map to (the colors)

This would be called "2-way"

mjb - March 9, 2023



How Different Cores' Cache Lines Keep Track of Each Other

27

Each core has its own separate L2 cache, but a write by one can impact the state of the others.

For example, if one core writes a value into one of its own cache lines, any other core using a copy of that same cache line can no longer count on its values being up-to-date. In order to regain that confidence, the core that wrote must flush that cache line back to memory and the other core must then reload its copy of that cache line.

To maintain this organization, each core's L2 cache has 4 states (**MESI**):

1. Modified
2. Exclusive
3. Shared
4. Invalid

Oregon State University Computer Graphics

mjb - March 9, 2023

A Simplified View of How MESI Works

28

1. Core A reads a value. Those values are brought into its cache. That cache line is now tagged **Exclusive**.
2. Core B reads a value from the same area of memory. Those values are brought into its cache, and now both cache lines are re-tagged **Shared**.
3. If Core B writes into that value. Its cache line is re-tagged **Modified** and Core A's cache line is re-tagged **Invalid**.
4. Core A tries to read a value from that same part of memory. But its cache line is tagged **Invalid**. So, Core B's cache line is flushed back to memory and then Core A's cache line is re-loaded from memory. Both cache lines are now tagged **Shared**.


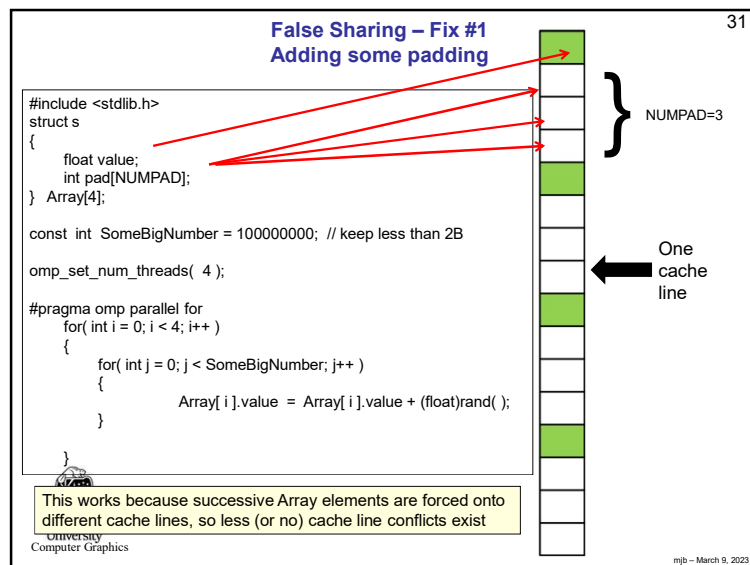
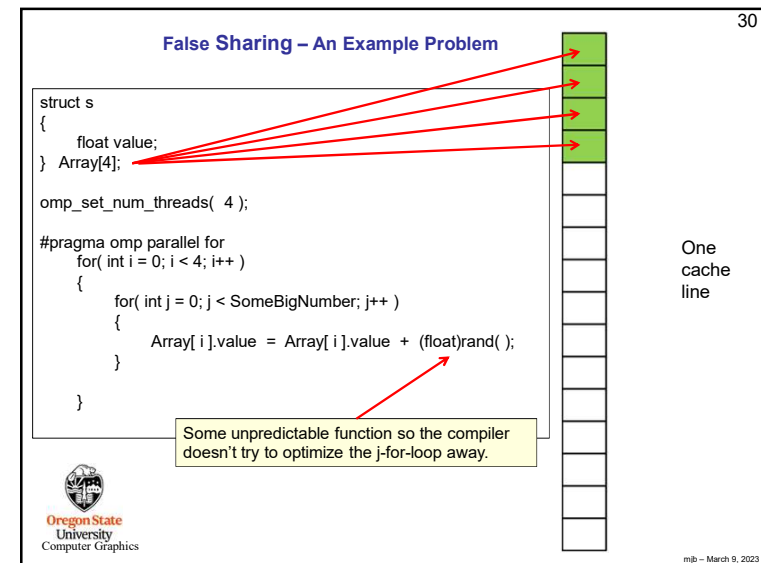
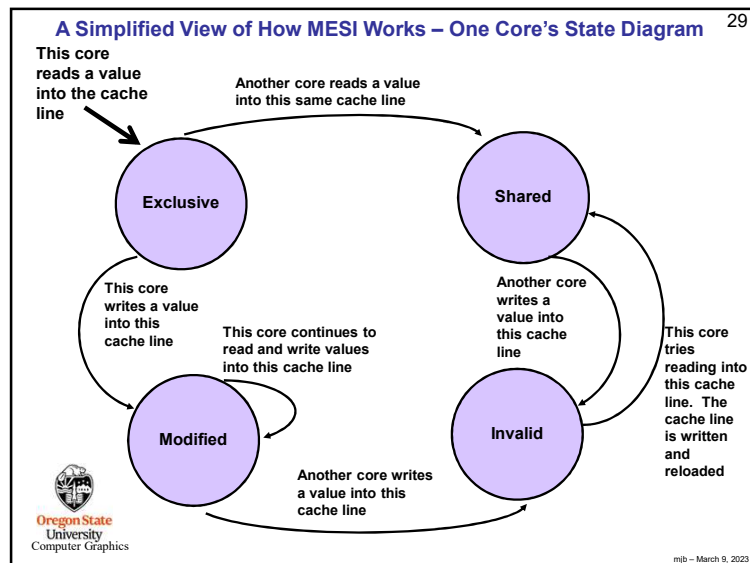
Step	Cache Line A	Cache Line B
1	Exclusive	-----
2	Shared	Shared
3	Invalid	Modified
4	Shared	Shared

This is a huge performance hit, and is referred to as **False Sharing**

Note that False Sharing doesn't create incorrect results – it just creates a performance hit. If anything, False Sharing prevents getting incorrect results.



Oregon State University Computer Graphics

mjb - March 9, 2023



False Sharing – the Effect of Spreading Your Data to Multiple Cache Lines 33



NUMPAD = 0



mjb - March 9, 2023

False Sharing – the Effect of Spreading Your Data to Multiple Cache Lines 34



NUMPAD = 1



mjb - March 9, 2023

False Sharing – the Effect of Spreading Your Data to Multiple Cache Lines 35


NUMPAD = 2





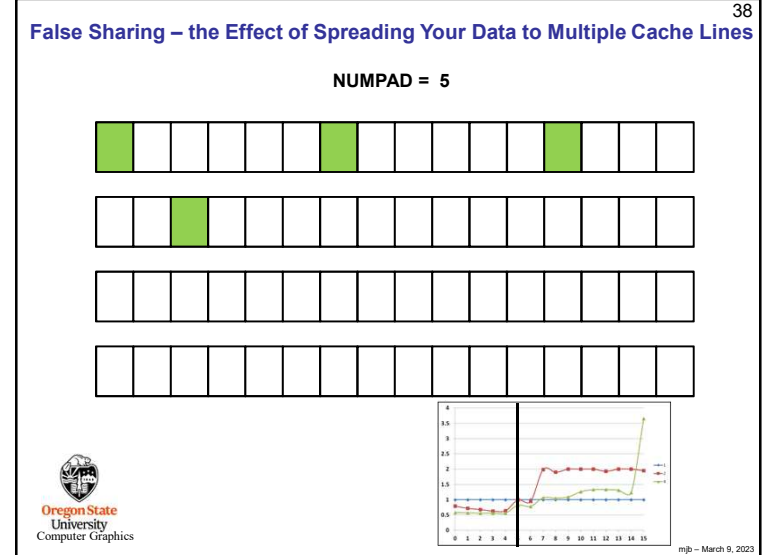
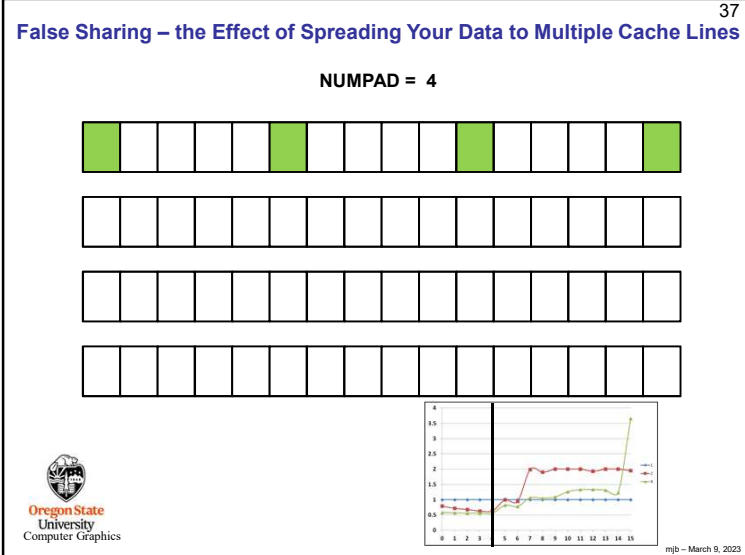
mjb - March 9, 2023

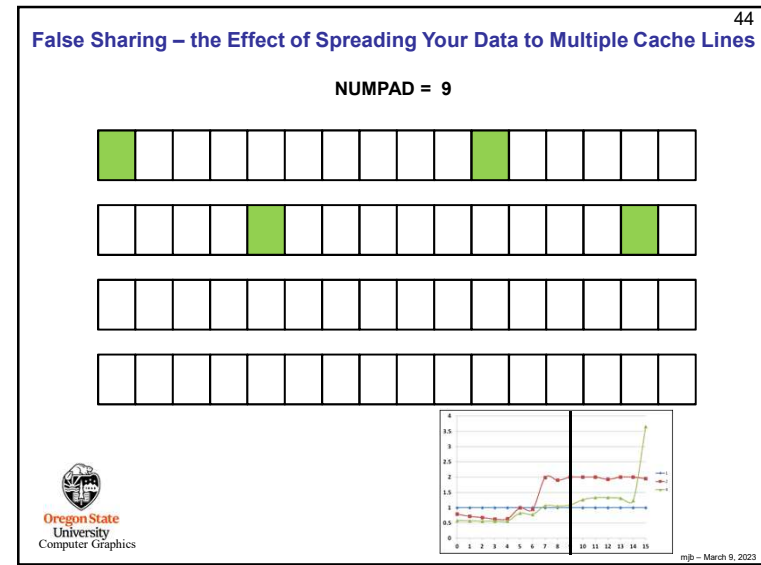


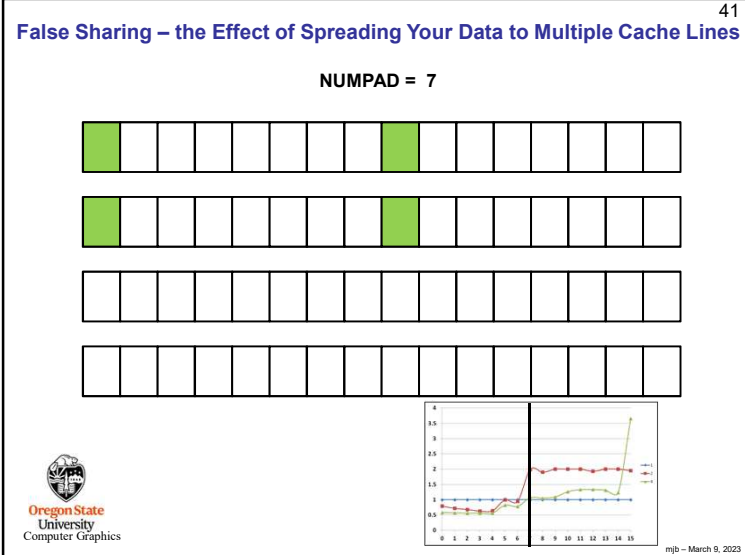
False Sharing – the Effect of Spreading Your Data to Multiple Cache Lines 36

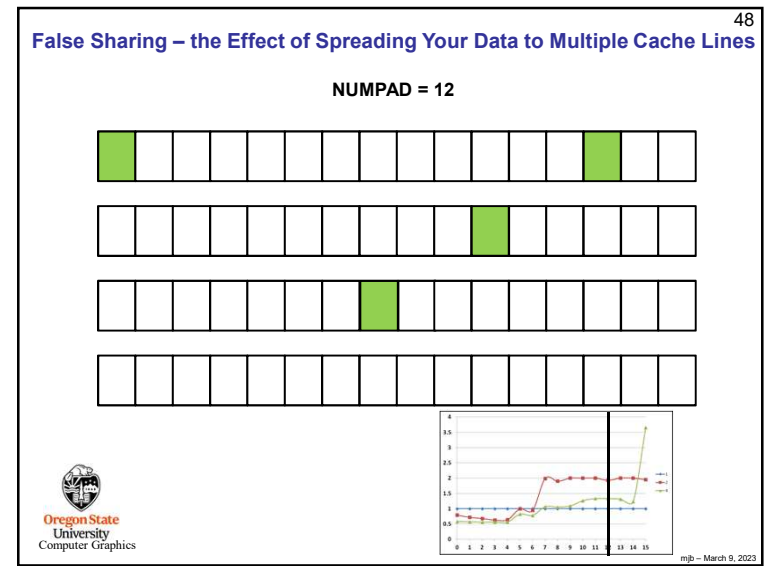
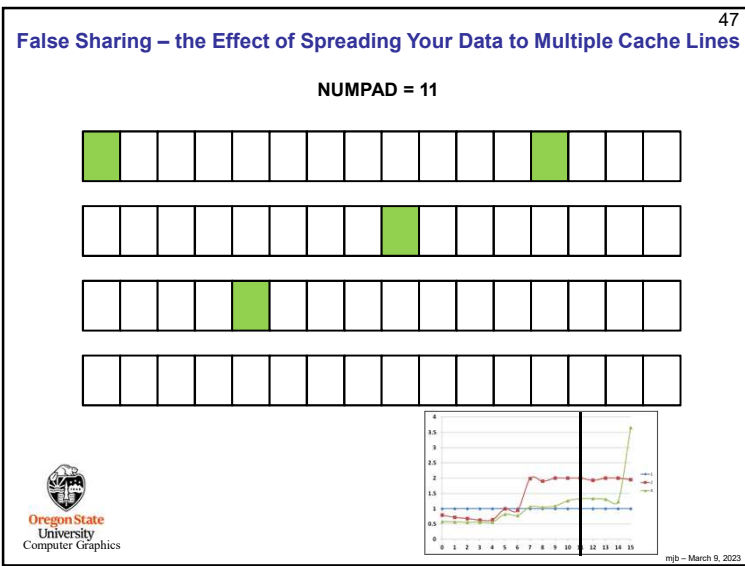


NUMPAD = 3



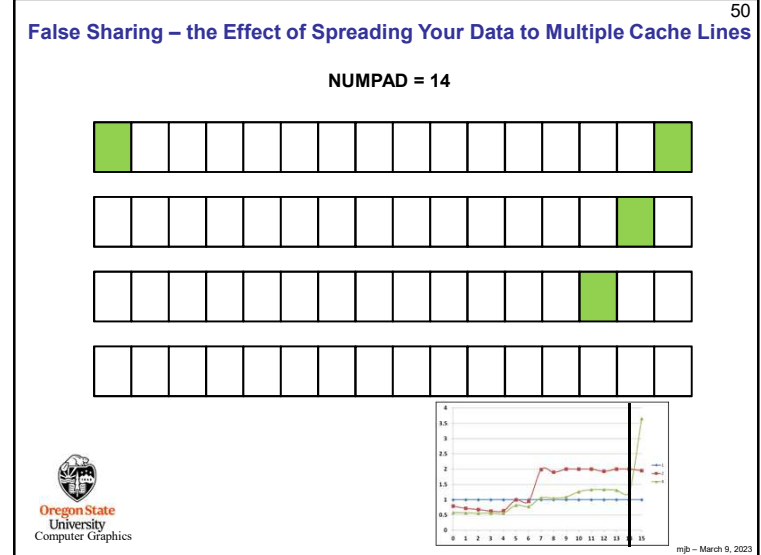
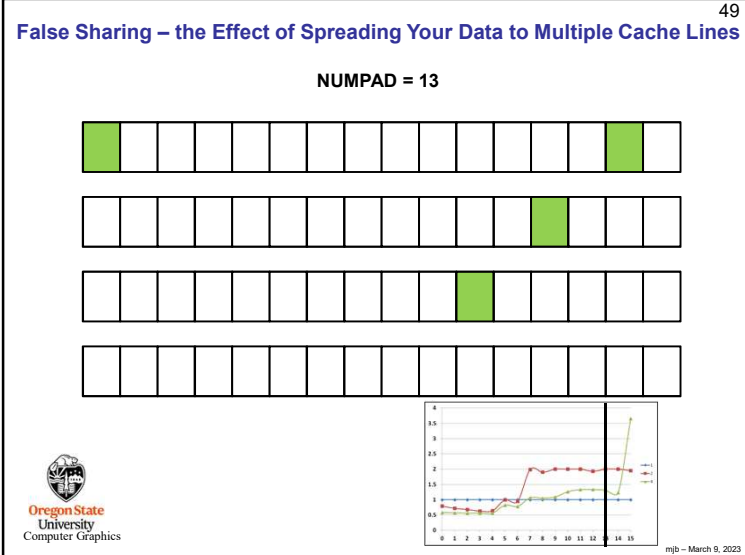


mjb - March 9, 2023









False Sharing – Fix #2: Using local (private) variables

OK, wasting memory to put your data on different cache lines seems a little silly (even though it works). Can we do something else?

Remember our discussion in the OpenMP section about how stack space is allocated for different threads?

If we use local variables, instead of contiguous array locations, that will spread our writes out in memory, and to different cache lines.

Diagram illustrating memory layout for multiple threads. The memory is divided into sections: Stack (green), Common Program Executable (yellow), Common Globals (purple), and Common Heap (red). Red arrows indicate that each thread has its own private stack space, which is used to store local variables, avoiding false sharing.

Oregon State University
Computer Graphics

mjb – March 9, 2023

False Sharing – Fix #2

```
#include <stdlib.h>
struct s
{
    float value;
} Array[4];

omp_set_num_threads( 4 );

const int SomeBigNumber = 100000000;

#pragma omp parallel for
for( int i = 0; i < 4; i++ )
{
    float tmp = Array[ i ].value;
    for( int j = 0; j < SomeBigNumber; j++ )
    {
        tmp = tmp + (float)rand( );
    }
    Array[ i ].value = tmp;
}
```


Makes this a private variable that lives in each thread's individual stack

This works because a localized temporary variable is created in each core's stack area, so little or no cache line conflict exists

Diagram illustrating memory layout for multiple threads. The memory is divided into sections: Stack (green), Common Program Executable (yellow), Common Globals (purple), and Common Heap (red). Red arrows indicate that each thread has its own private stack space, which is used to store local variables, avoiding false sharing.

Oregon State University
Computer Graphics

mjb – March 9, 2023



malloc'ing on a cache line

What if you are malloc'ing, and want to be sure your data structure starts on a cache line boundary?

Knowing that cache lines start on fixed 64-byte boundaries lets you do this. Consider a memory address. The top N-6 bits tell you what cache line number this address is a part of. The bottom 6 bits tell you what offset that address has within that cache line. So, for example, on a 32-bit memory system:

32 - 6 = 26 bits	6 bits: 0-63
Cache line number	Offset in that cache line

So, if you see a memory address whose bottom 6 bits are 000000, then you know that that memory location begins on a cache line boundary.

Oregon State University
Computer Graphics

mjb – March 9, 2023

malloc'ing on a cache line

57

Let's say that you have a structure and you want to malloc an ARRAYSIZE array of them. Normally, you would do this:

```
struct xyzw *p = (struct xyzw *) malloc( (ARRAYSIZE)*sizeof(struct xyzw) );
struct xyzw *Array = &p[0];
...
Array[ i ].x = 10. ;
```

If you wanted to make sure that array of structures started on a cache line boundary, you would do this:

```
unsigned char *p = (unsigned char *) malloc( 64 + (ARRAYSIZE)*sizeof(struct xyzw) );
int offset = (long int)p & 0x3f; // 0x3f = bottom 6 bits are all 1's
struct xyzw *Array = (struct xyzw *) &p[64-offset];
...
Array[ i ].x = 10. ;
```

Remember that when you want to free this malloc'ed space, be sure to say:


not: ~~free(Array);~~

Oregon State University
Computer Graphics

mjb - March 9, 2023

Now, Consider This Type of Computation

58



Should you allocate the data as one large global-memory block (i.e., shared)? Or, should you allocate it as separate blocks, each local to its own core (i.e., private)? Does it matter? Yes!

If you allocate the data as one large global-memory block, there is a risk that you will get False Sharing at the individual-block boundaries. Solution: make sure that each individual-block starts and ends on a cache boundary, even if you have to pad it. **(Fix #1!)**

If you allocate the data as separate blocks, then you don't have to worry about False Sharing (**Fix #2!**), but you do have to worry about the logic of your program remembering where to find each Node #-1 and Node #i+1.

Oregon State University
Computer Graphics

mjb - March 9, 2023