



Sony Computer Entertainment Europe Research & Development Division

Pitfalls of Object Oriented Programming

Tony Albrecht – Technical Consultant
Developer Services



PS3
PlayStation 3

What I will be covering

- A quick look at Object Oriented (OO) programming
- A common example
- Optimisation of that example
- Summary



Object Oriented (OO) Programming

- What is OO programming?
 - a programming paradigm that uses "objects" – data structures consisting of datafields and methods together with their interactions – to design applications and computer programs.
(Wikipedia)
- Includes features such as
 - Data abstraction
 - Encapsulation
 - Polymorphism
 - Inheritance



What's OOP for?

- OO programming allows you to think about problems in terms of objects and their interactions.
- Each object is (ideally) self contained
 - Contains its own code and data.
 - Defines an interface to its code and data.
- Each object can be perceived as a 'black box'.



Objects

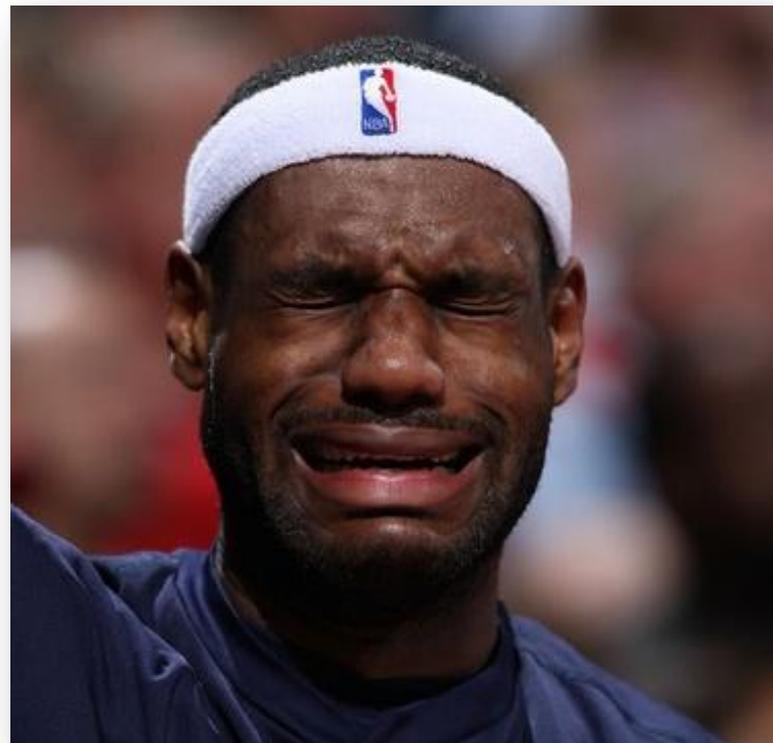
- If objects are self contained then they can be
 - Reused.
 - Maintained without side effects.
 - Used without understanding internal implementation/representation.

- This is good, yes?

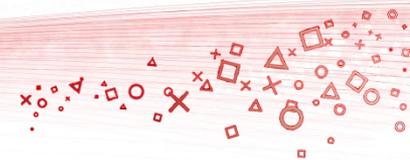


Are Objects Good?

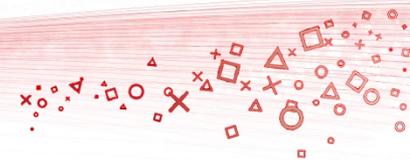
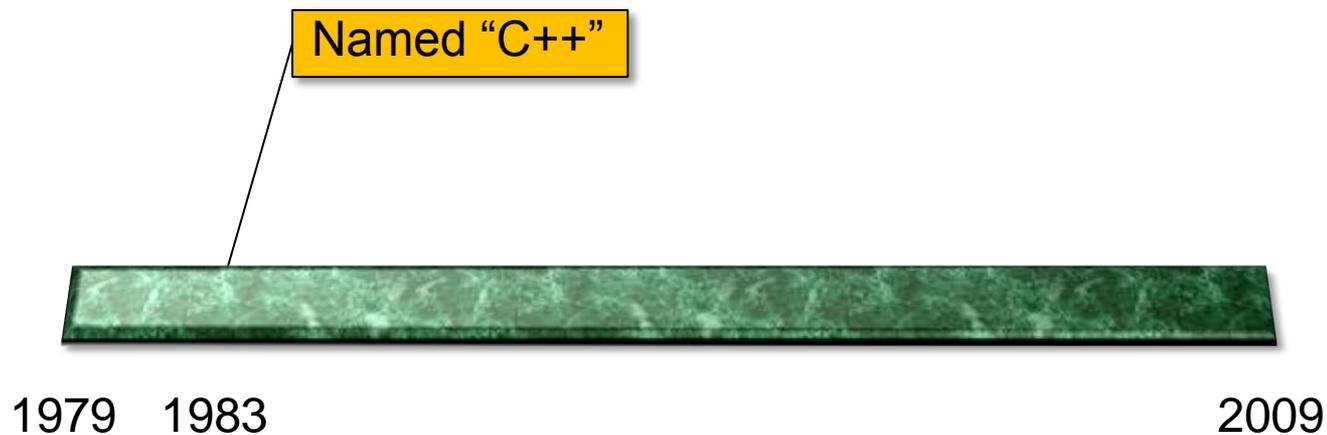
- Well, yes
- And no.
- First some history.



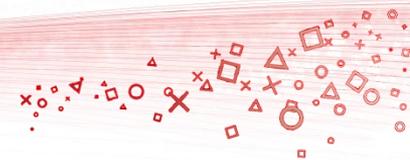
A Brief History of C++



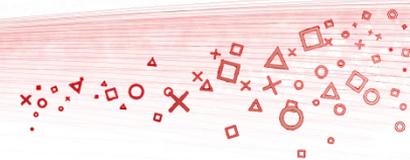
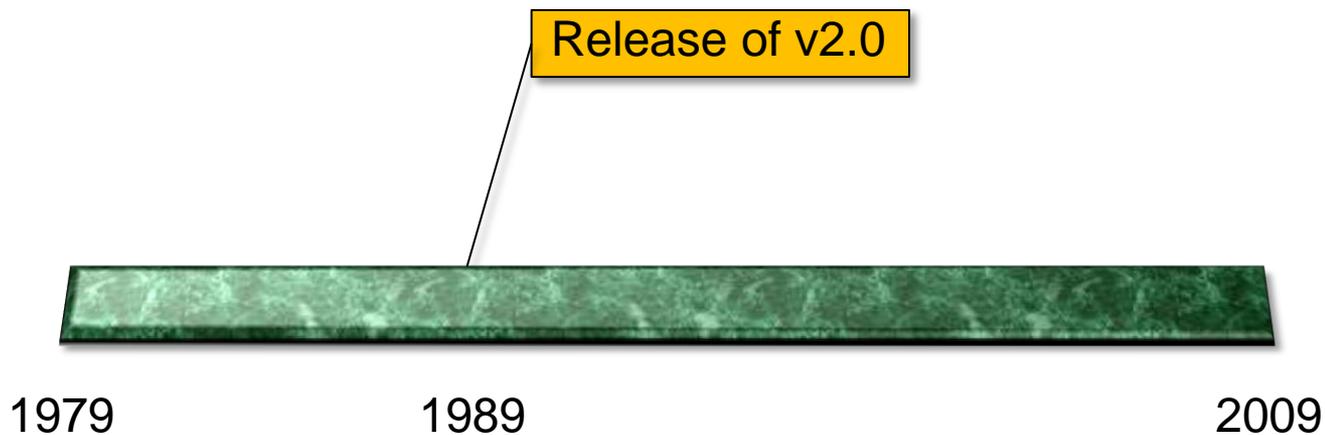
A Brief History of C++



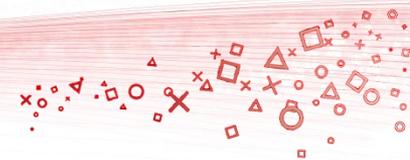
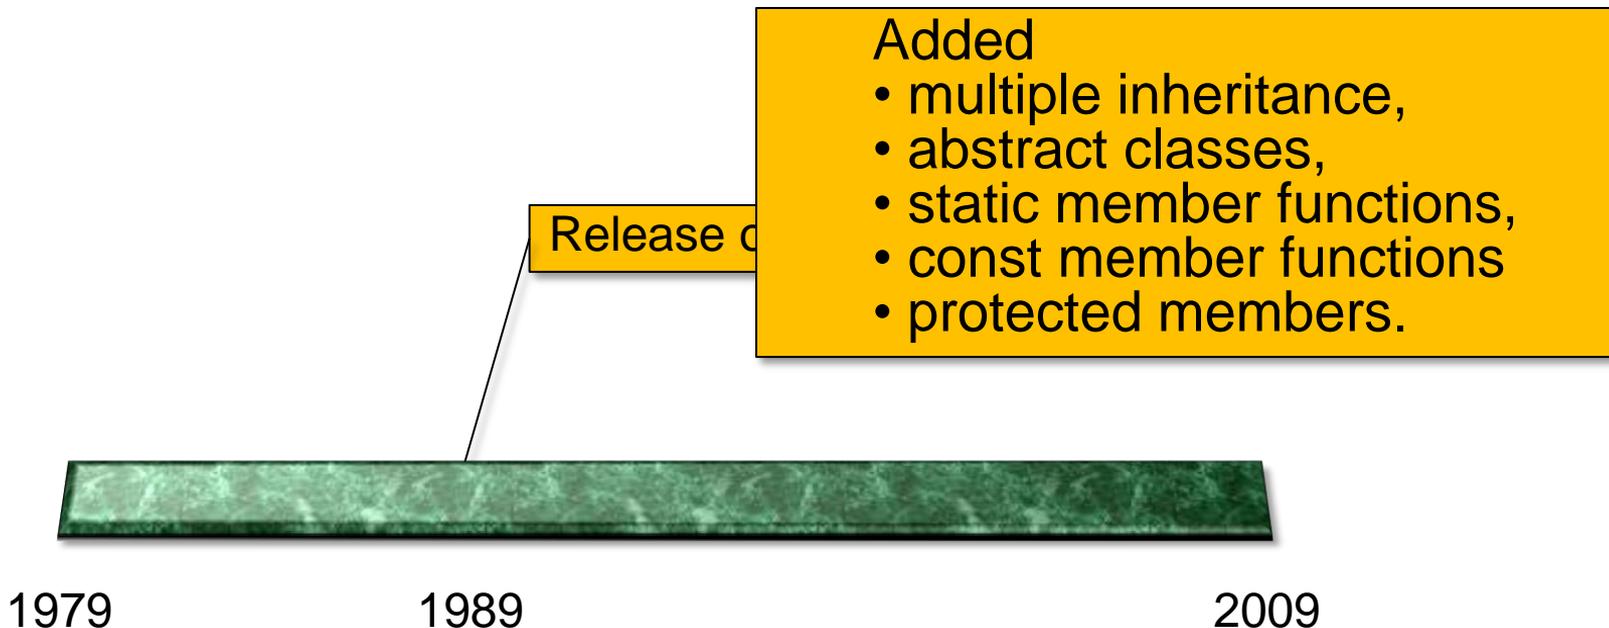
A Brief History of C++



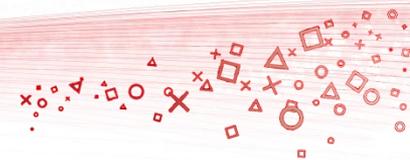
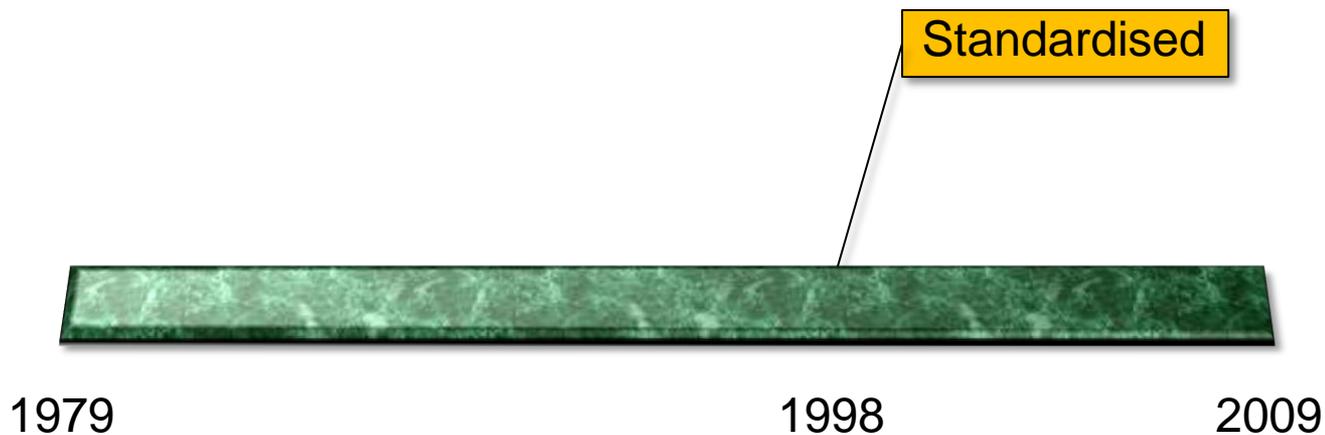
A Brief History of C++



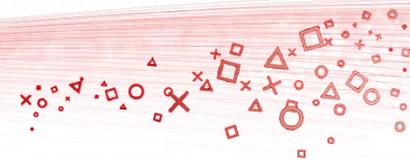
A Brief History of C++



A Brief History of C++



A Brief History of C++



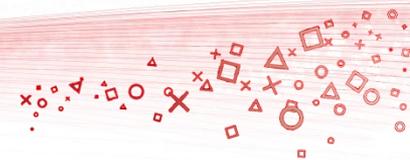
A Brief History of C++



1979

2009

?



So what has changed since 1979?

- Many more features have been added to C++
- CPUs have become much faster.
- Transition to multiple cores
- Memory has become faster.



• 16-bit microprocessor
 • 16K RAM
 • 13" color monitor
 (24 lines of 32 chrs.)
 • 26K ROM operating system
 (includes 14K BASIC)
 • Sound - 3 tones, 5 octaves
 • 16 colors: 192 x 256 res.
 • Large TI library of ROM
 programs available.

only \$1150
Includes 13" Color Monitor!

FINALLY TEXAS INSTRUMENTS TI-99/4 Home Computer

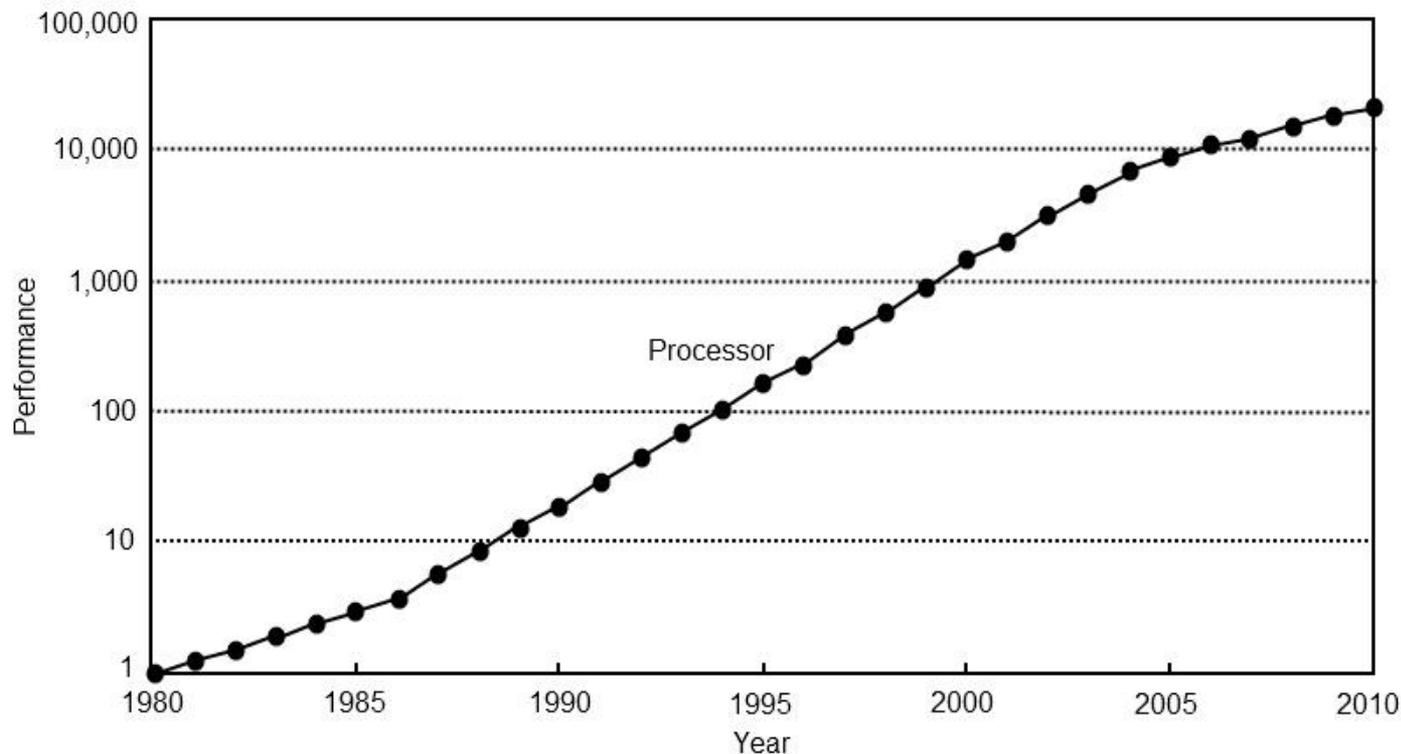
Many Peripherals. Coming soon!

Over 1000 software tapes, books, disks on display. Come in and browse.

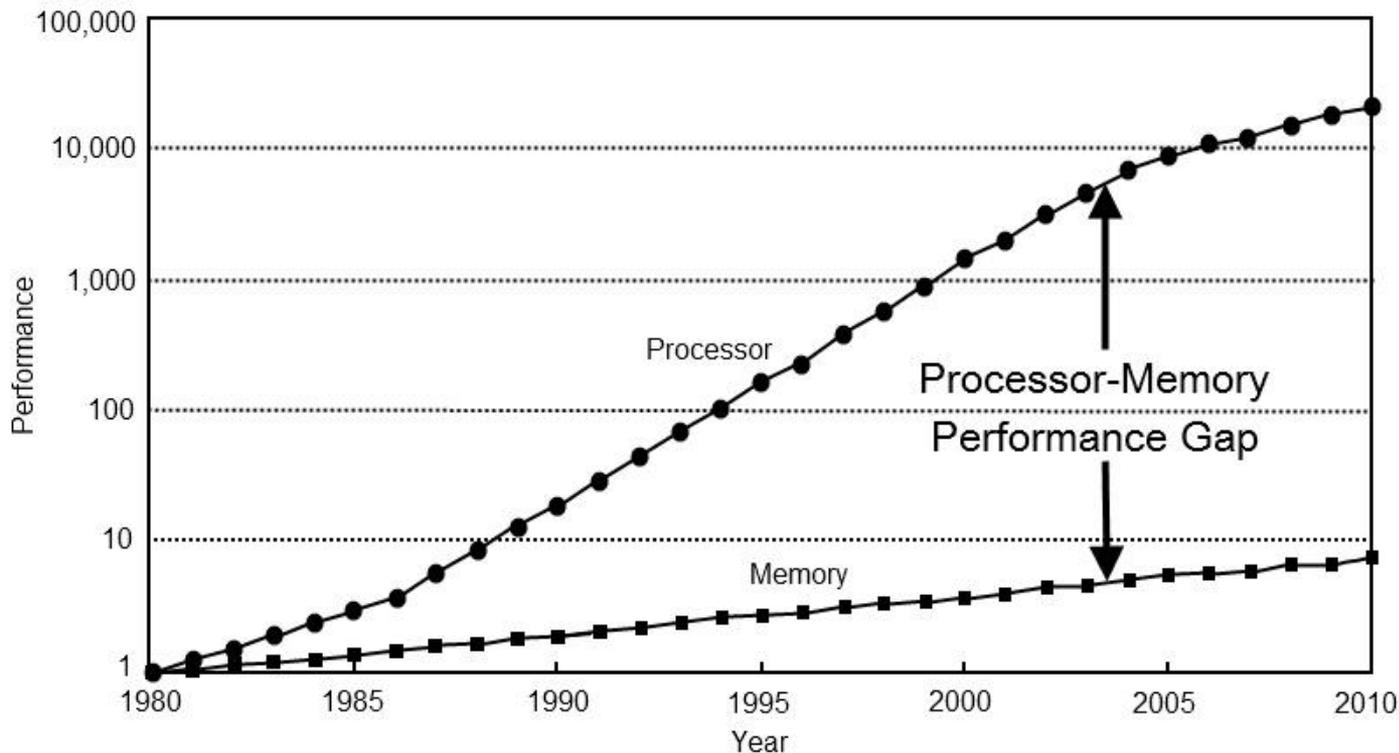
<http://www.vintagecomputing.com>



CPU performance



CPU/Memory performance



What has changed since 1979?

- One of the biggest changes is that memory access speeds are far slower (relatively)
 - 1980: RAM latency ~ 1 cycle
 - 2009: RAM latency ~ 400+ cycles
- What can *you* do in 400 cycles?



What has this to do with OO?

- OO classes encapsulate code and data.
- So, an instantiated object will generally contain all data associated with it.



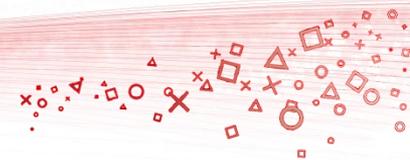
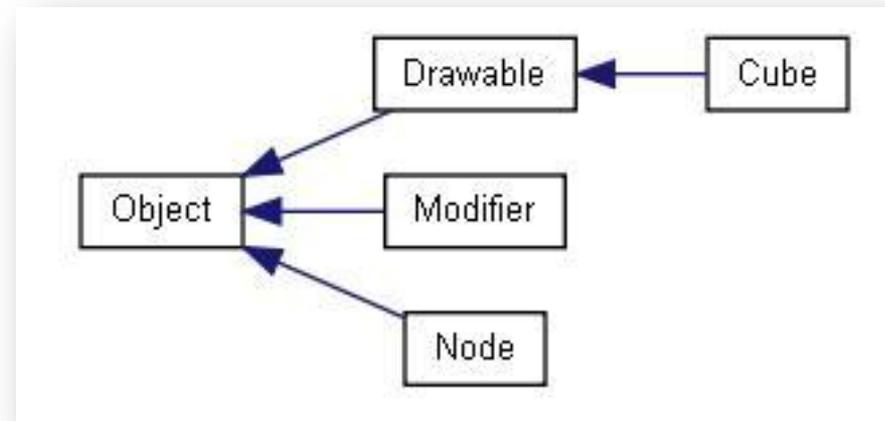
My Claim

- With modern HW (particularly consoles), excessive encapsulation is BAD.
- Data flow should be fundamental to your design (Data Oriented Design)



Consider a simple OO Scene Tree

- Base Object class
 - Contains general data
- Node
 - Container class
- Modifier
 - Updates transforms
- Drawable/Cube
 - Renders objects



Object

- Each object
 - Maintains bounding sphere for culling
 - Has transform (local and world)
 - Dirty flag (optimisation)
 - Pointer to Parent

```
class object
{
    // <methods removed for clarity>
    Matrix4 m_Transform;
    Matrix4 m_worldTransform;
    BoundingSphere m_BoundingSphere;
    BoundingSphere m_worldBoundingSphere;
    char* m_Name;
    bool m_Dirty;
    object* m_Parent;
};
```



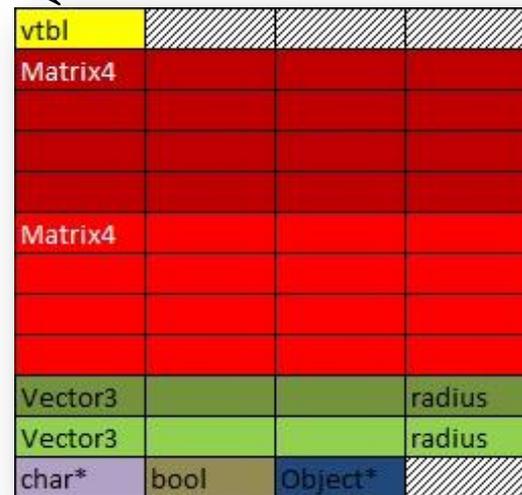
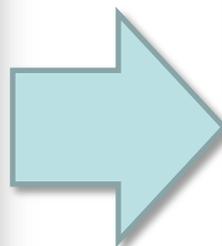
Objects

Class Definition

```
class object
{
    // <methods removed for clarity>
    Matrix4 m_Transform;
    Matrix4 m_worldTransform;
    BoundingBox m_BoundingBox;
    BoundingBox m_worldBoundingBox;
    char* m_Name;
    bool m_Dirty;
    object* m_Parent;
};
```

Each square is 4 bytes

Memory Layout



Nodes

- Each Node is an object, plus
 - Has a container of other objects
 - Has a visibility flag.

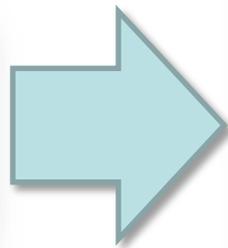
```
class Node : public Object
{
    // methods removed for clarity
    std::vector<Object*> m_Objects;
    bool m_Isvisible;
};
```



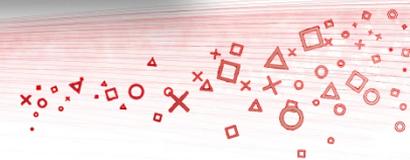
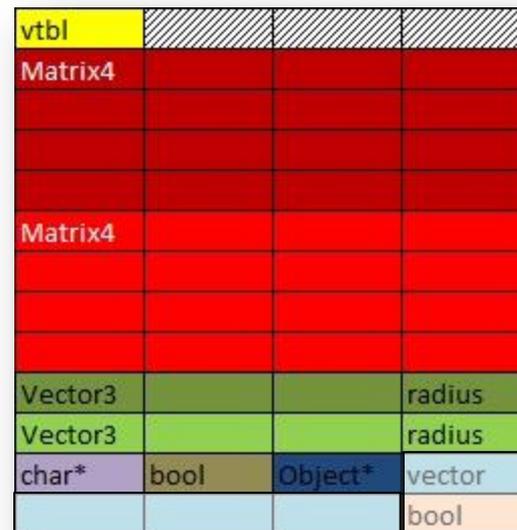
Nodes

Class Definition

```
class Node : public Object
{
    // methods removed for clarity
    std::vector<Object*> m_Objects;
    bool m_IsVisible;
};
```



Memory Layout



Consider the following code...

- Update the world transform and world space bounding sphere for each object.

```
const BoundingBox& Node::GetWorldBoundingBox(const Matrix4& parentTransform)
{
    m_worldTransform = parentTransform*m_Transform;

    for(std::vector<Object*>::const_iterator itr = m_Objects.begin();
        itr != m_Objects.end();
        ++itr)
    {
        m_worldBoundingBox.ExpandBy((*itr)->GetWorldBoundingBox(m_worldTransform));
    }
    return m_worldBoundingBox;
}
```



Consider the following code...

- Leaf nodes (objects) return transformed bounding spheres

```
virtual const BoundingBox& GetworldBoundingBox(const Matrix4& parentTransform)
{
    if(m_Dirty)
        m_worldBoundingBox = m_BoundingBox.Transform(parentTransform);
    return m_worldBoundingBox;
}
```



Consider the following code...

- Leaf node (with children) return transformed bounding

What's wrong with this code?

```
virtual const BoundingBox& GetworldBoundingBox(const Matrix4& parentTransform)
{
    if(m_Dirty)
        m_worldBoundingBox = m_BoundingBox.Transform(parentTransform);
    return m_worldBoundingBox;
}
```



Consider the following code...

- Leaf node (transform) transformed bounding

If m_Dirty=false then we get branch misprediction which costs 23 or 24 cycles.

```
virtual const BoundingSphere& GetworldBoundingBox(const Matrix4& parentTransform)
{
    if(m_Dirty)
        m_worldBoundingBox = m_BoundingBox.Transform(parentTransform);
    return m_worldBoundingBox;
}
```



Consider the following code...

- Leaf nodes (objects) return transformed bounding

Calculation of the world bounding sphere takes only 12 cycles.

```
virtual const BoundingBox& GetWorldBoundingBox(const Matrix4& parentTransform)
{
    if(m_Dirty)
        m_worldBoundingBox = m_BoundingBox.Transform(parentTransform);
    return m_worldBoundingBox;
}
```



Consider the following code...

- Leaf nodes (objects) return transformed bounding

So using a dirty flag here is actually slower than not using one (in the case where it is false)

```
virtual const BoundingBoxSphere& GetworldBoundingBoxSphere(const Matrix4& parentTransform)
{
    if(m_Dirty)
        m_worldBoundingBoxSphere = m_BoundingBoxSphere.Transform(parentTransform);
    return m_worldBoundingBoxSphere;
}
```



Cache usage

Main Memory



parentTransform is already in the cache (somewhere)

L2 Cache



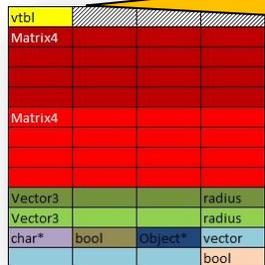
```
m_worldTransform = parentTransform*m_Transform;
```



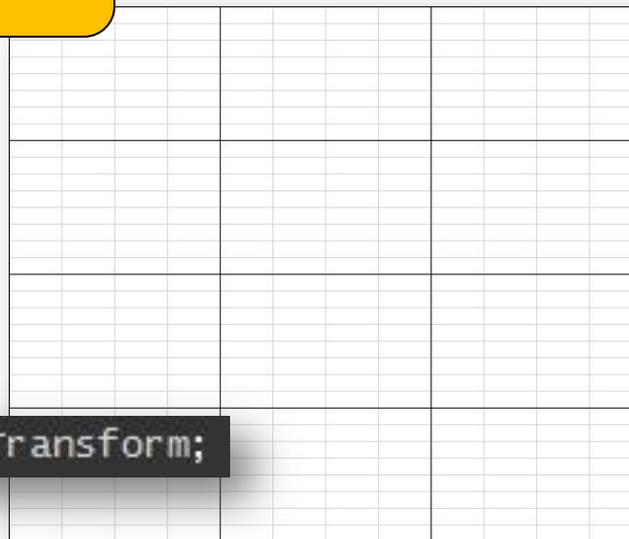
Cache usage

Assume this is a 128byte boundary (start of cacheline)

Main Memory



L2 Cache



```
m_worldTransform = parentTransform*m_Transform;
```

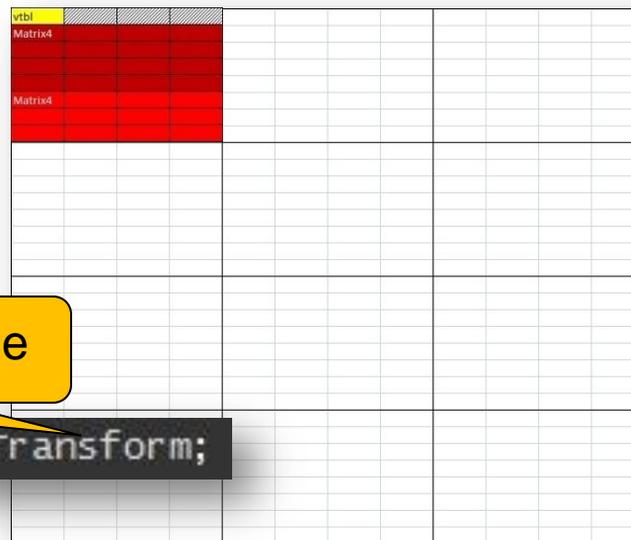


Cache usage

Main Memory



L2 Cache



Load m_Transform into cache

```
m_worldTransform = parentTransform*m_Transform;
```

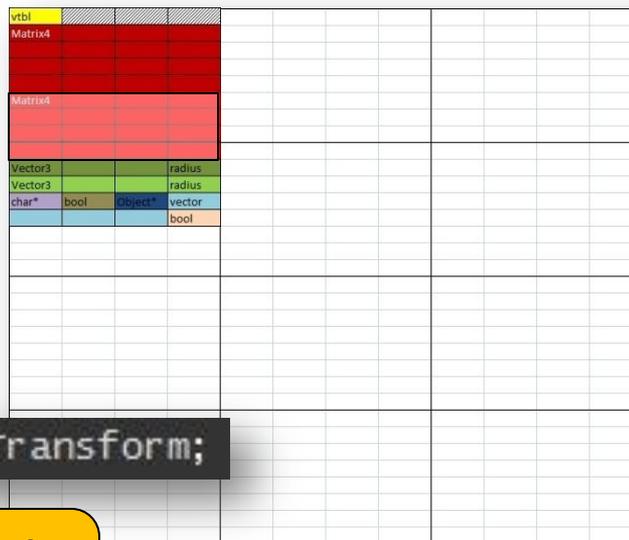


Cache usage

Main Memory



L2 Cache



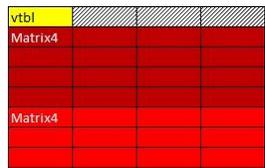
```
m_worldTransform = parentTransform*m_Transform;
```

m_WorldTransform is stored via cache (write-back)



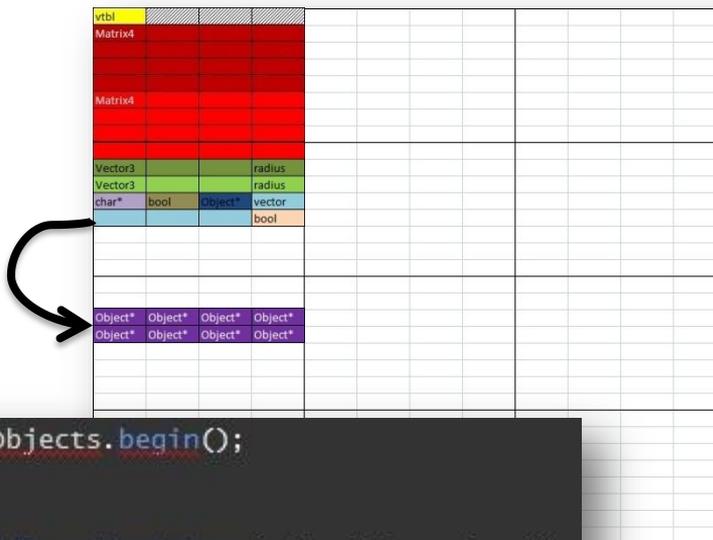
Cache usage

Main Memory



Then a pointer is pulled from somewhere else (Memory managed by std::vector)

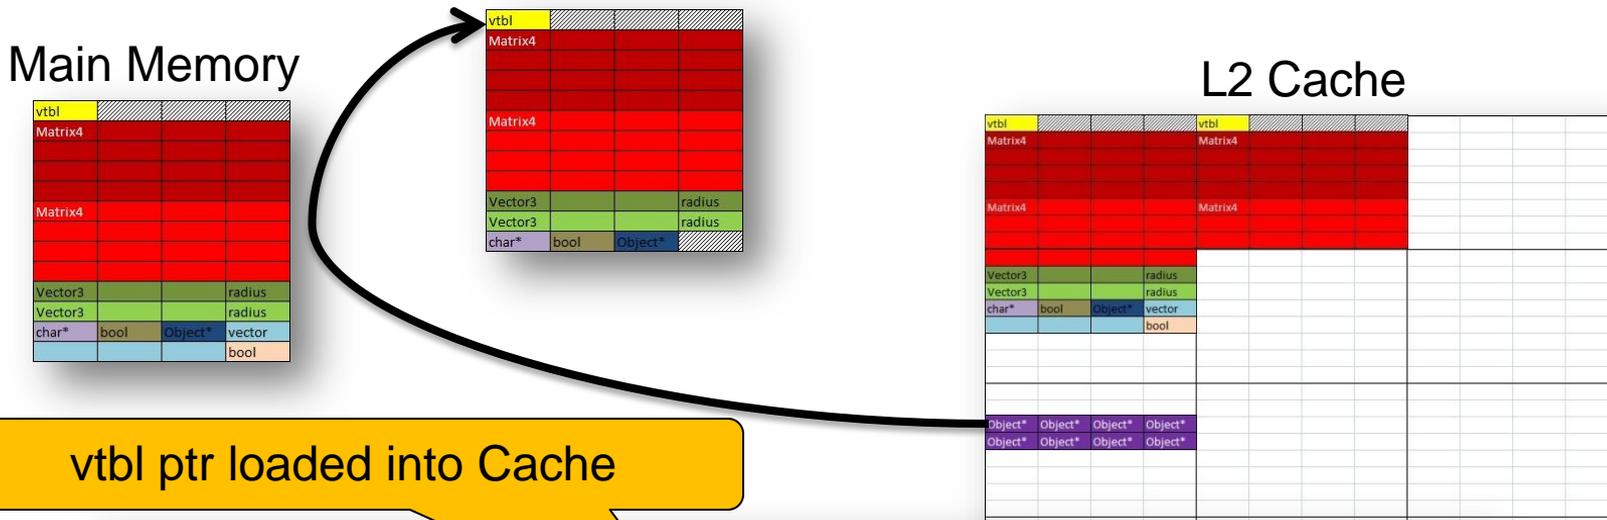
L2 Cache



```
for(std::vector<Object*>::const_iterator itr = m_Objects.begin();
    itr != m_Objects.end();
    ++itr)
{
    m_worldBoundingSphere.ExpandBy((*itr)->GetWorldBoundingSphere(m_worldTransform));
}
```



Cache usage



vtbl ptr loaded into Cache

```

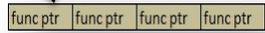
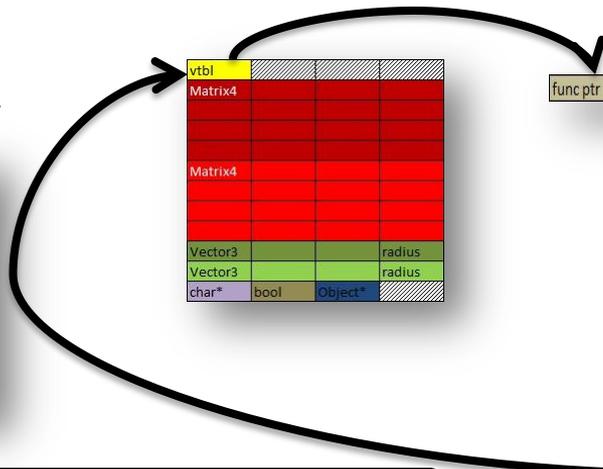
for(std::vector<Object*>::iterator itr = m_Objects.begin();
    itr != m_Objects.end();
    ++itr)
{
    m_worldBoundingSphere.ExpandBy((*itr)->GetWorldBoundingSphere(m_worldTransform));
}

```

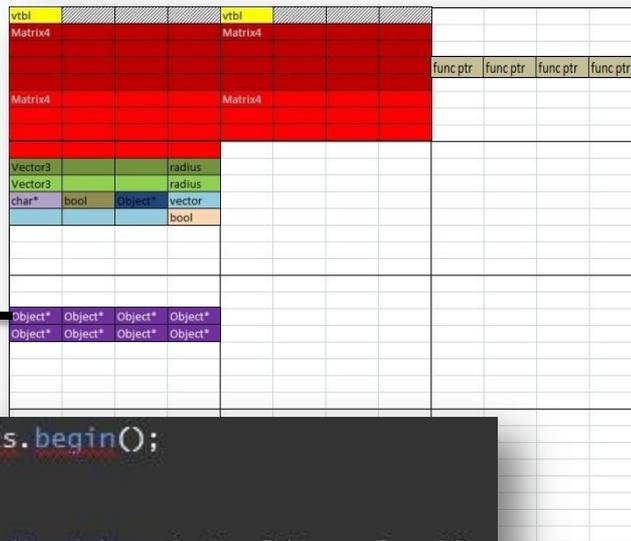


Cache usage

Main Memory



L2 Cache



Look up virtual function

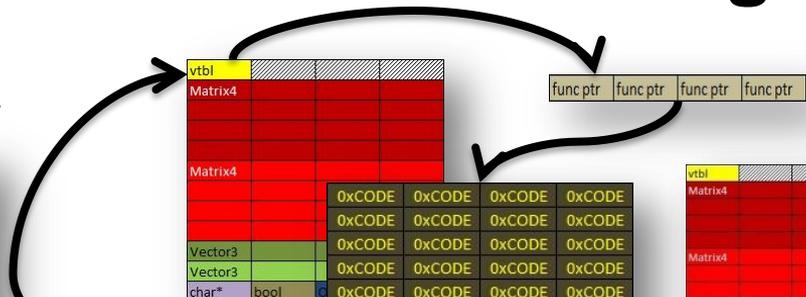
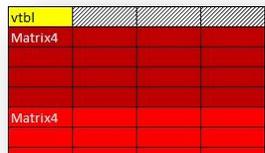
```

for(std::vector<Object*>::iterator itr = m_Objects.begin();
    itr != m_Objects.end();
    ++itr)
{
    m_worldBoundingSphere.ExpandBy((*itr)->GetworldBoundingSphere(m_worldTransform));
}
    
```

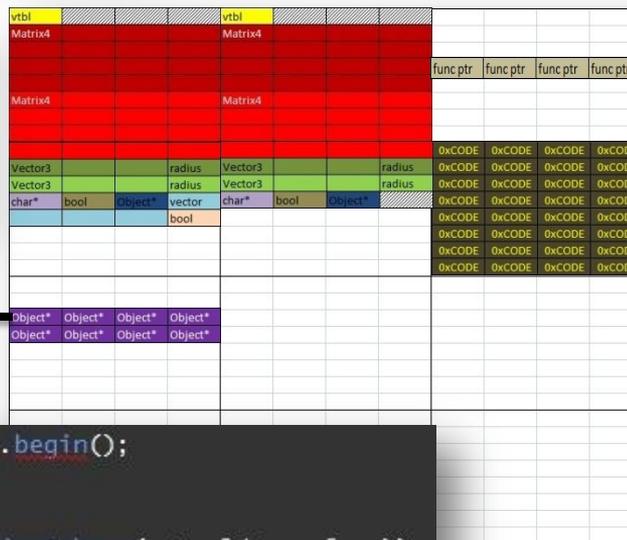


Cache usage

Main Memory



L2 Cache



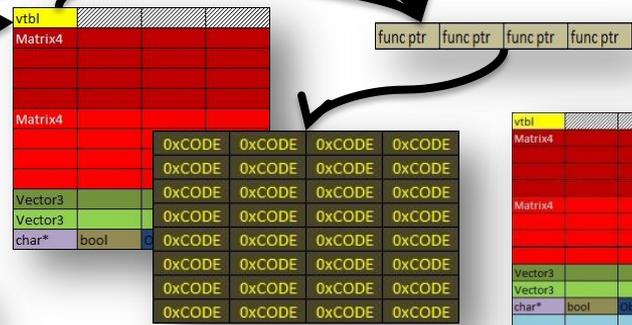
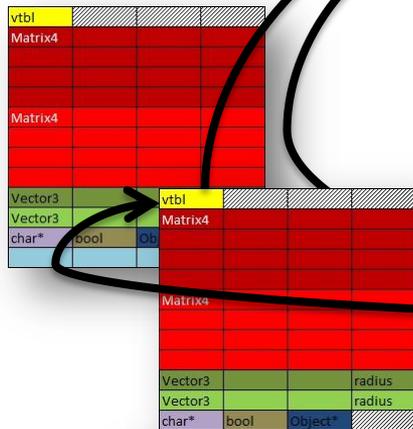
Node's World Bounding Sphere is then Expanded

```
for (std::vector<Object*>::const_iterator itr = m_objects.begin();
     itr != m_objects.end();
     ++itr)
{
    m_worldBoudingsphere.ExpandBy((*itr)->GetworldBoudingsSphere(m_worldTransform));
}
```

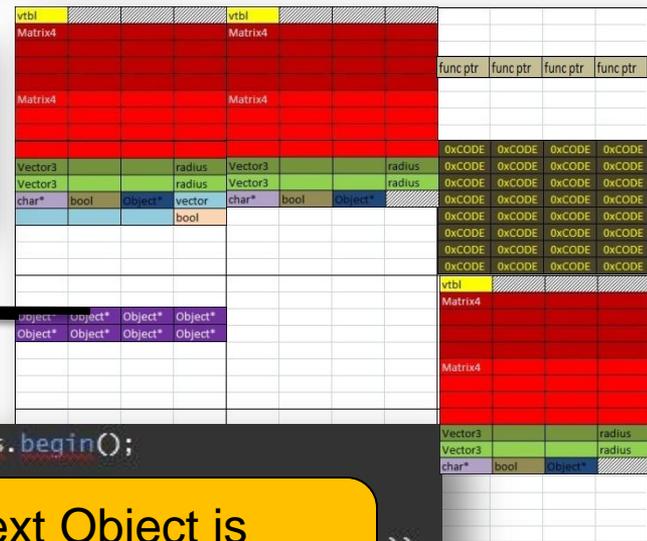


Cache usage

Main Memory



L2 Cache



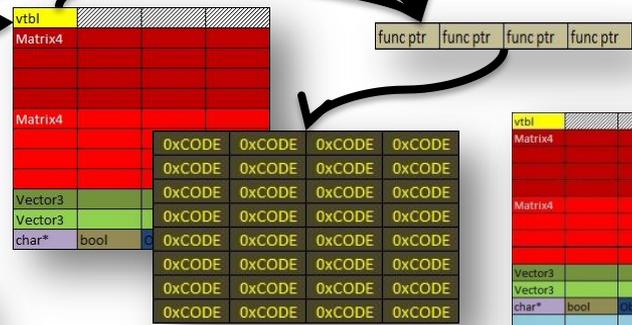
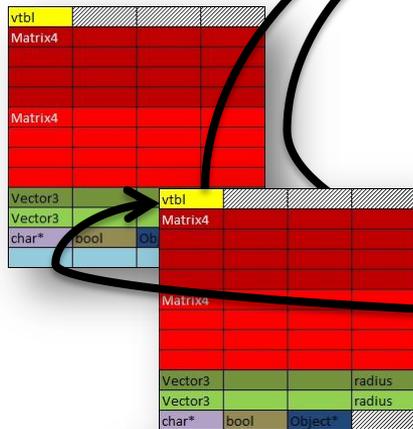
```
for(std::vector<Object*>::const_iterator itr = m_objects.begin();
    itr != m_objects.end();
    ++itr)
{
    m_worldBoundingsphere.Expand(itr);
}
```

Then the next Object is processed



Cache usage

Main Memory



L2 Cache



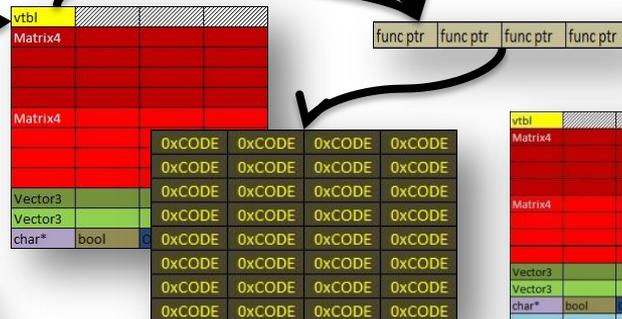
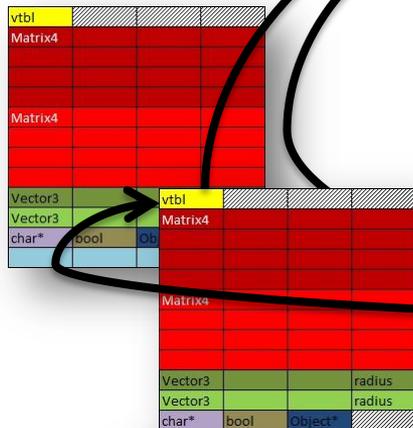
```
for(std::vector<Object*>::const_iterator itr = m_objects.begin();
    itr != m_objects.end();
    ++itr)
{
    m_worldBoundingsphere.Expand(itr);
}
```

First object costs at least 7 cache misses



Cache usage

Main Memory



L2 Cache



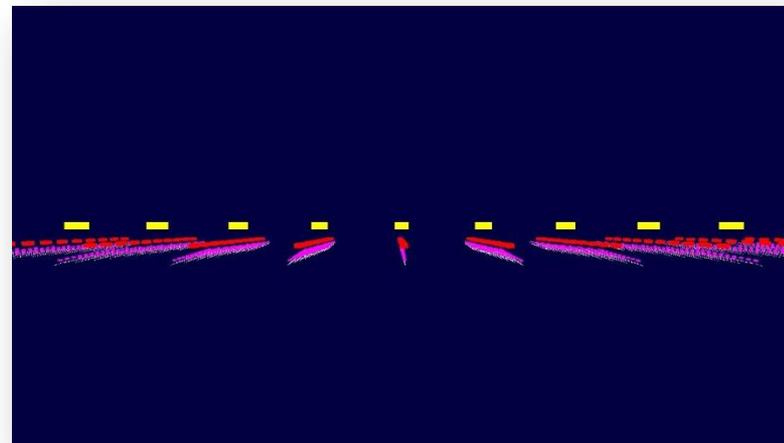
Subsequent objects cost at least 2 cache misses each

```
for (std::vector<Object*>::
itr!= m_objects.end();
++itr)
{
m_worldBoudingsphere.ExpandBy((*itr)->GetworldBoudingsSphere(m_worldTransform));
}
```

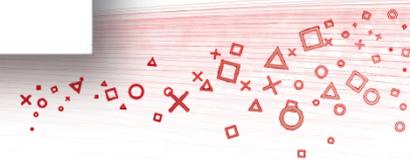
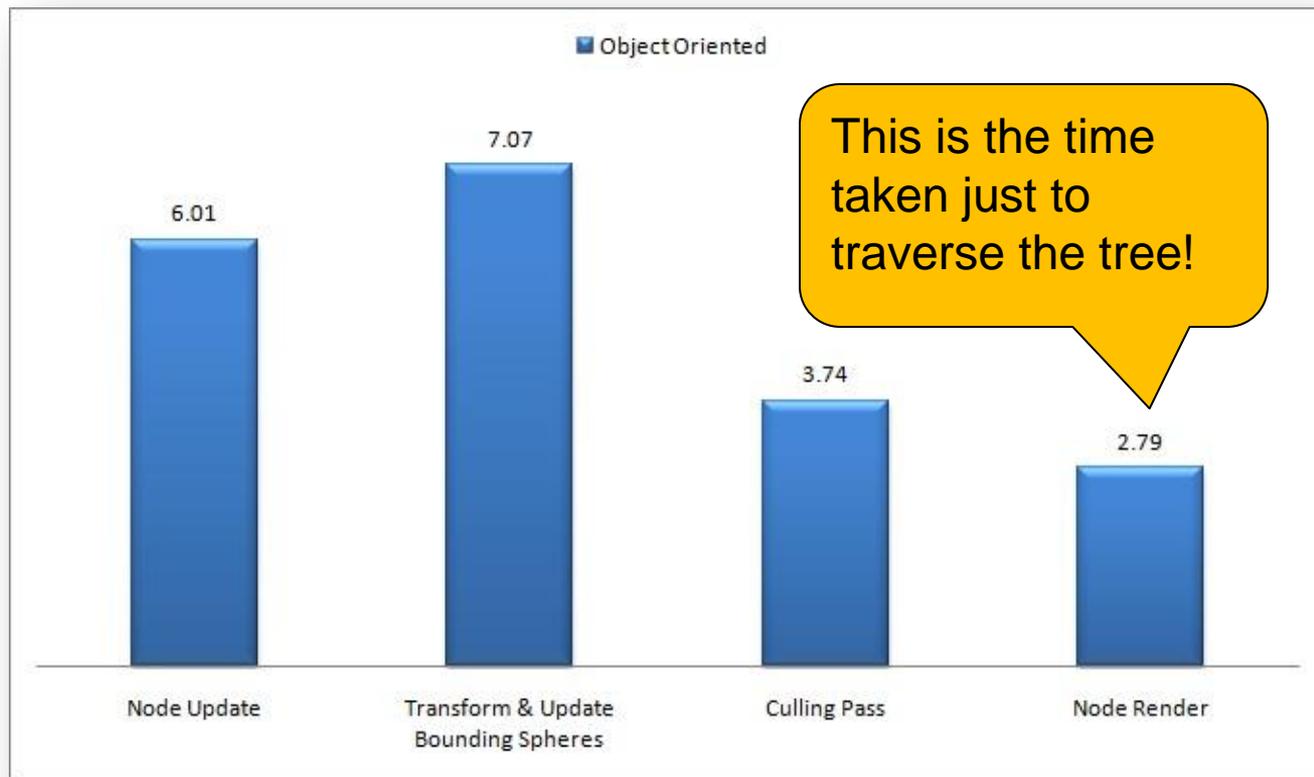


The Test

- 11,111 nodes/objects in a tree 5 levels deep
- Every node being transformed
- Hierarchical culling of tree
- Render method is empty

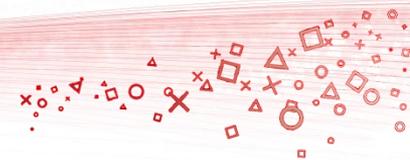


Performance

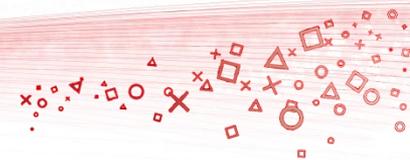
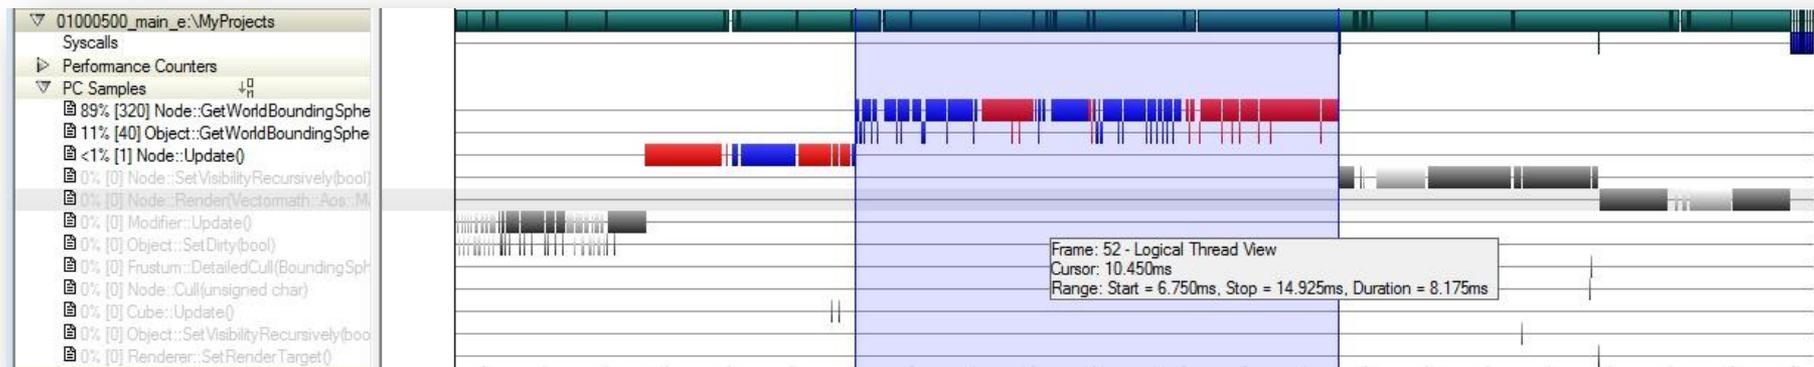


Why is it so slow?

~22ms



Look at GetWorldBoundingSphere()



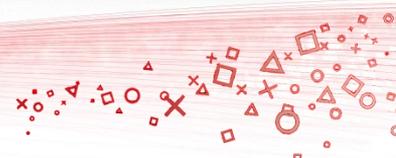
Samples can be a little misleading at the source code level

```
372 9 57 const BoundingSphere& Node::GetWorldBoundingSphere(const Matrix4& parentTransform)
- - 58
2998 49 59     if(!m_Dirty)
- - 60         return m_WorldBoundingSphere;
- - 61
- - 62     // if it was dirty, then we need to update the bounding volumes and transforms
2212 32 63     if(m_Parent)
- - 64         m_WorldTransform = parentTransform*m_Transform;
- - 65     else
- - 66         m_WorldTransform=m_Transform;
- - 67     // was dirty, so we need to recalculate the bounds of the children
42 1 68     m_WorldBoundingSphere=BoundingSphere(); // zero it
- - 69
2179 46 70     for(std::vector<Object*>::const_iterator itr = m_Objects.begin();
- - 71         itr!= m_Objects.end();
- - 72         ++itr)
- - 73     {
9085 174 74         m_WorldBoundingSphere.ExpandBy((*itr)->GetWorldBoundingSphere(m_WorldTransform));
- - 75     }
123 3 76     m_Dirty=false;
18 1 77     return m_WorldBoundingSphere;
314 6 78 }
- - 79
```



<pre> 372 9 57 const BoundingSphere& - - 58 { 2998 49 59 if(!m_Dirty) - - 60 return m_World - - 61 - - 62 // if it was dirty 2212 32 63 if(m_Parent) - - 64 - - 69 2179 46 70 for(std::vector<Ob - - 71 itr!= m_Object - - 72 ++itr) - - 73 { 9085 174 74 m_WorldBoundin - - 75 } 123 3 76 m_Dirty=false; 18 1 77 return m_WorldBou 314 6 78 } - - 79 </pre>	<pre> 9 - 00012724 F8010100 std r0,0x100(r1) PIPE 8 - 00012728 DBE100E8 stfd f31,0xE8(r1) - - 0001272C 386300B0 addi r3,r3,0xB0 8 - 00012730 FB8100C0 std r28,0xC0(r1) 12 2 00012734 FBA100C8 std r29,0xC8(r1) PIPE 9 1 00012738 881B00D4 lbz r0,0xD4(r27) 23 1 0001273C FBC100D0 std r30,0xD0(r1) PIPE ----- 2989 48 00012740 2F800000 cmpwi cr7,r0,0x0 ----- - - 00012744 FBE100D8 std r31,0xD8(r1) PIPE - - 00012748 419E026C beq cr7,0x000129B4 - - 0001274C 801B00D8 lwz r0,0xD8(r27) 2 00012750 2F800000 cmpwi cr7,r0,0x0 01 (0001274C) REG - - 00012754 419E02A0 beq cr7,0x000129F4 01 (00012750) REG - - 00012758 397B0010 addi r11,r27,0x10 - - 0001275C 7D2020CE lvx v9,0,r4 - - 00012760 38E00010 li r7,0x10 1 00012764 79690020 clrldi r9,r11,32 PIPE - - 00012768 39000020 li r8,0x20 - - 0001276C 38000030 li r0,0x30 PIPE 2 00012770 788A0020 clrldi r10,r4,32 19 - 00012774 7D8058CE lvx v12,0,r11 3 00012778 100004C4 vxor v0,v0,v0 - - 0001277C 7D6938CE lvx v11,r9,r7 3 00012780 3B9B0050 addi r28,r27,0x50 - - 00012784 7CA940CE lvx v5,r9,r8 18 - 00012788 7DA900CE lvx v13,r9,r0 - - 0001278C 7B890020 clrldi r9,r28,32 01 (00012780) REG 1344 19 00012790 10202A8C vspltw v1,v5,0 - - 00012794 11012A8C vspltw v8,v5,1 PIPE 9 1 00012798 1040628C vspltw v2,v12,0 </pre>
---	--

if(!m_Dirty) comparison



Stalls due to the load 2 instructions earlier

```

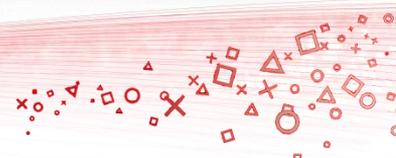
372 9 57 const BoundingSphere&
- - 58 {
2998 49 59     if(!m_Dirty)
- - 60         return m_World
- - 61
- - 62     // if it was dirty
2212 32 63     if(m_Parent)
- - 64         m_WorldTransfo
- - 65     else
- - 66         m_WorldTransfo
- - 67     // was dirty, so v
42 1 68     m_WorldBoundingSpl
- - 69
2179 46 70     for(std::vector<Ob
- - 71         itr!= m_Object
- - 72         ++itr)
- - 73     {
9085 174 74         m_WorldBoundin
- - 75     }
123 3 76     m_Dirty=false;
18 1 77     return m_WorldBou
314 6 78 }
- - 79

```

```

9 - 00012724 F8010100 std r0,0x10
8 - 00012728 DBE100E8 stfd f31,0xE
- - 0001272C 386300B0 addi r3,r3,0
8 - 00012730 FB8100C0 std r28,0xC0(r1)
12 2 00012734 FBA100C8 std r29,0xC8(r1)
9 1 00012738 881B00D4 lbz r0,0xD4(r27) PIPE
23 1 0001273C FBC100D0 std r30,0xD0(r1) PIPE
2989 48 00012740 2F800000 cmpwi cr7,r0,0x0
- - 00012744 FBE100D8 std r31,0xD8(r1) PIPE
- - 00012748 419E026C beq cr7,0x000129B4
- - 0001274C 801B00D8 lwz r0,0xD8(r27)
32 2 00012750 2F800000 cmpwi cr7,r0,0x0 01 (0001274C) REG
30 - 00012754 419E02A0 beq cr7,0x000129F4 01 (00012750) REG
1 - 00012758 397B0010 addi r11,r27,0x10
3 - 0001275C 7D2020CE lvx v9,0,r4
5 - 00012760 38E00010 li r7,0x10
1 - 00012764 79690020 clrdi r9,r11,32 PIPE
- - 00012768 39000020 li r8,0x20
22 - 0001276C 38000030 li r0,0x30 PIPE
2 - 00012770 788A0020 clrdi r10,r4,32
19 - 00012774 7D8058CE lvx v12,0,r11
3 - 00012778 100004C4 vxor v0,v0,v0
- - 0001277C 7D6938CE lvx v11,r9,r7
3 - 00012780 3B9B0050 addi r28,r27,0x50
- - 00012784 7CA940CE lvx v5,r9,r8
18 - 00012788 7DA900CE lvx v13,r9,r0
- - 0001278C 7B890020 clrdi r9,r28,32 01 (00012780) REG
1344 19 00012790 10202A8C vspltw v1,v5,0
- - 00012794 11012A8C vspltw v8,v5,1 PIPE
9 1 00012798 1040628C vspltw v2,v12,0

```



```

372 9 57 const BoundingSphere&
- - 58 {
2998 49 59     if(!m_Dirty)
- - 60         return m_World
- - 61
- - 62     // if it was dirty
2212 32 63     if(m_Parent)
- - 64         m_WorldTransfo
- - 65     else
- - 66         m_WorldTransfo
- - 67     // was dirty, so v
42 1 68     m_WorldBoundingSpl
- - 69
2179 46 70     for(std::vector<Ob
- - 71         itr!= m_Object
- - 72         ++itr)
- - 73     {
9085 174 74         m_WorldBoundin
- - 75     }
123 3 76     m_Dirty=false;
18 1 77     return m_WorldBou
314 6 78 }
- - 79

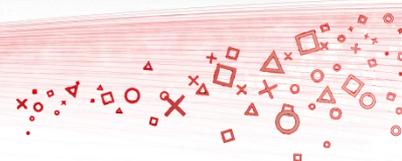
```

```

9 - 00012724 F8010100 std r0,0x100(r1) PIPE
8 - 00012728 DBE100E8 stfd f31,0xE8(r1)
- - 0001272C 386300B0 addi r3,r3,0xB0
8 - 00012730 FB8100C0 std r28,0xC0(r1)
12 2 00012734 FBA100C8 std r29,0xC8(r1) PIPE
9 1 00012738 881B00D4 lbz r0,0xD4(r27)
23 1 0001273C FBC100D0 std r30,0xD0(r1) PIPE
-----
2989 48 00012740 2F800000 cnpwi cr7,r0,0x0
- - 00012744 FBE100D8 std r31,0xD8(r1)
- - 00012748 419E026C beq cr7,0x0
- - 0001274C 801B00D8 lwz r0,0xD8(r1)
32 2 00012750 2F800000 cnpwi cr7,r0,0x0
30 - 00012754 419E02A0 beq cr7,0x0
1 - 00012758 397B0010 addi r11,r27
3 - 0001275C 7D2020CE lvx v9,0,r4
5 - 00012760 38E00010 li r7,0x10
1 - 00012764 79690020 clrdi r9,r11,32 PIPE
- - 00012768 39000020 li r8,0x20
- - 0001276C 38000030 li r0,0x30 PIPE
2 - 00012770 788A0020 clrdi r10,r4,32
19 - 00012774 7D8058CE lvx v12,0,r11
3 - 00012778 100004C4 vxor v0,v0,v0
- - 0001277C 7D6938CE lvx v11,r9,r7
3 - 00012780 3B9B0050 addi r28,r27,0x
- - 00012784 7CA940CE lvx v5,r9,r8
18 - 00012788 7DA900CE lvx v13,r9,r0
- - 0001278C 7B890020 clrdi r9,r28,32
1344 19 00012790 10202A8C vspltw v1,v5,0
- - 00012794 11612A8C vspltw v8,v5,1 PIPE
9 1 00012798 1040628C vspltw v2,v12,0

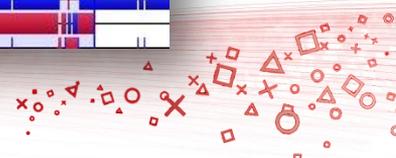
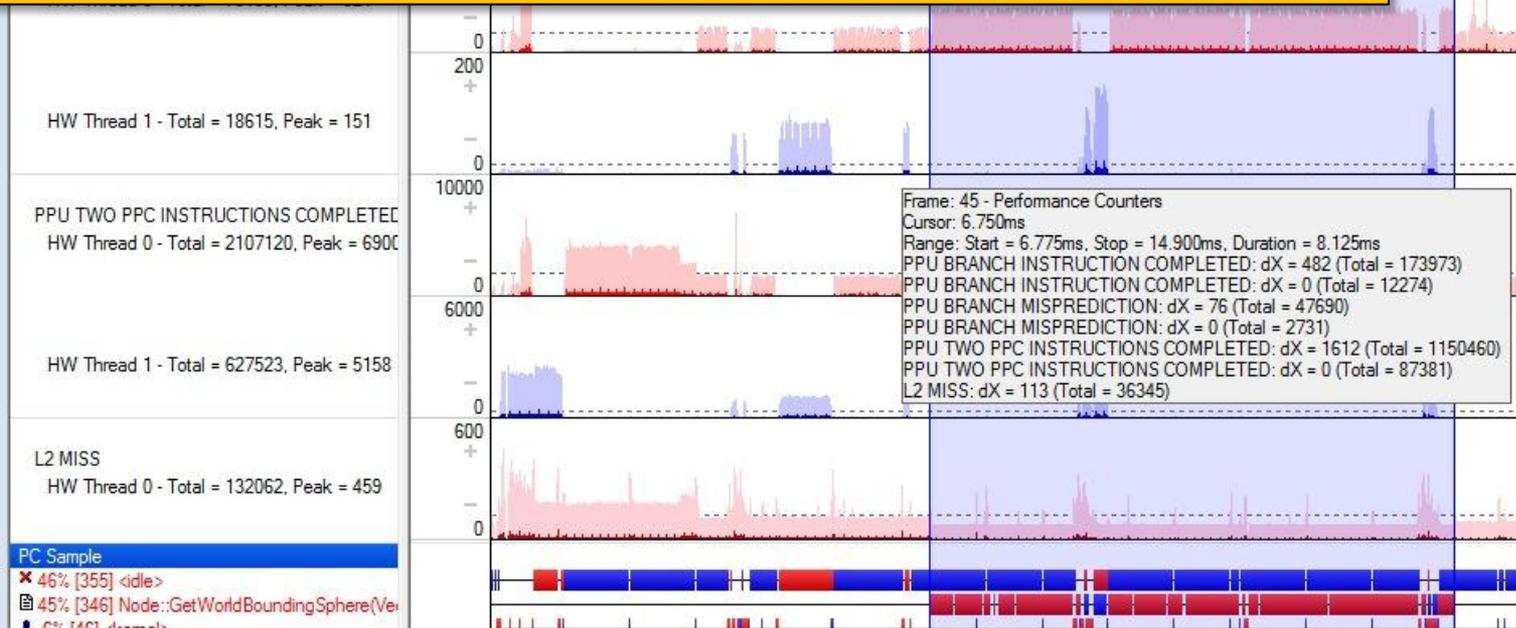
```

Similarly with the matrix multiply



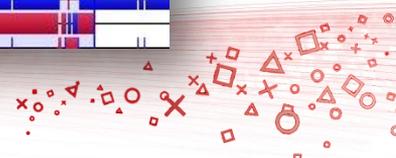
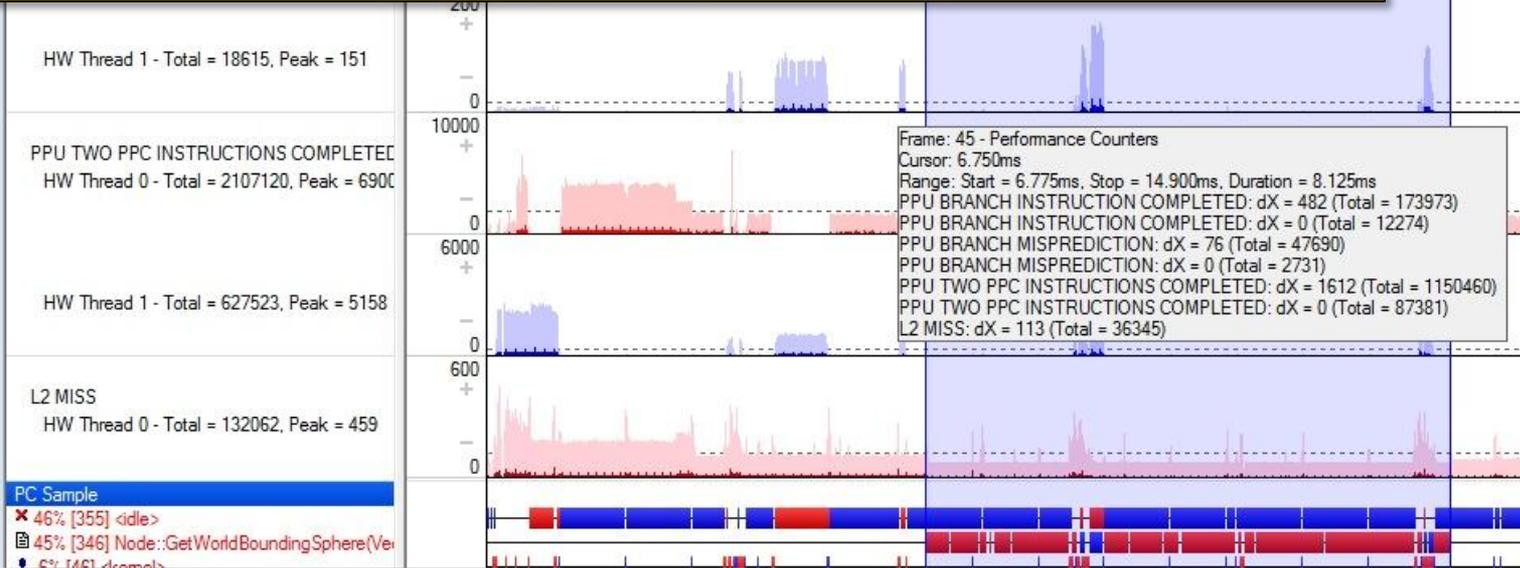
Some rough calculations

Branch Mispredictions: 50,421 @ 23 cycles each $\approx 0.36\text{ms}$



Some rough calculations

Branch Mispredictions: 50,421 @ 23 cycles each $\approx 0.36\text{ms}$
 L2 Cache misses: 36,345 @ 400 cycles each $\approx 4.54\text{ms}$



- From Tuner, ~ 3 L2 cache misses per object
 - These cache misses are mostly sequential (more than 1 fetch from main memory can happen at once)
 - Code/cache miss/code/cache miss/code...



Slow memory is the problem here

- How can we fix it?
- And still keep the same functionality and interface?



The first step

- Use homogenous, sequential sets of data

```
class Object
{
    // <methods removed for clarity>
    Matrix4 *m_Transform;
    Matrix4 *m_worldTransform;
    BoundingSphere *m_BoundingSphere;
    BoundingSphere *m_worldBoundingSphere;
    char* m_Name;
    bool m_Dirty;
    Object* m_Parent;
};
```

vtbl	Matrix4*	Matrix4*	Bsphere*
Bsphere*	char*	bool	Object*
vector			

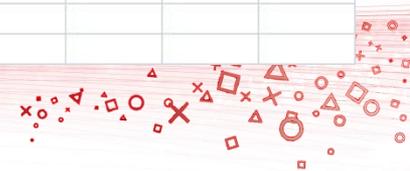
```
class Node : public Object
{
    // methods removed for clarity
    std::vector<Object*> m_Objects;
    bool m_IsVisible;
};
```



Homogeneous Sequential Data

vtbl	Matrix4*	Matrix4*	Bsphere*
Bsphere*	char*	bool	Object*
vector			

World transforms				Local Transforms				World Bounding Spheres				Local Bounding Spheres			
Matrix4				Matrix4				Vector3				Vector3			
								Vector3				Vector3			
								Vector3				Vector3			
Matrix4				Matrix4				Vector3				Vector3			
								Vector3				Vector3			
								Vector3				Vector3			
Matrix4				Matrix4				Vector3				Vector3			
								Vector3				Vector3			
								Vector3				Vector3			
Matrix4				Matrix4				Vector3				Vector3			
								Vector3				Vector3			
								Vector3				Vector3			
								Vector3				Vector3			

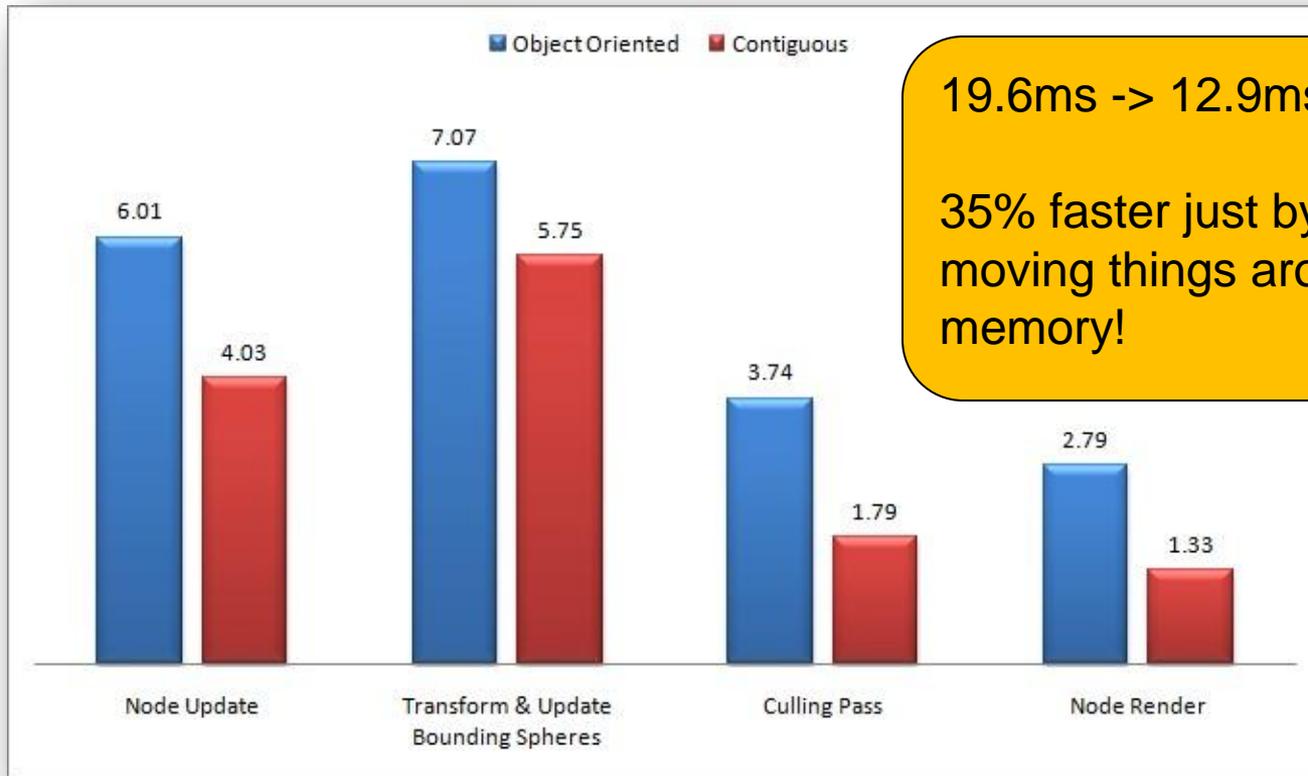


Generating Contiguous Data

- Use custom allocators
 - Minimal impact on existing code
- Allocate contiguous
 - Nodes
 - Matrices
 - Bounding spheres



Performance



What next?

- Process data in order
- Use implicit structure for hierarchy
 - Minimise to and fro from nodes.
- Group logic to optimally use what is already in cache.
- Remove regularly called virtuals.



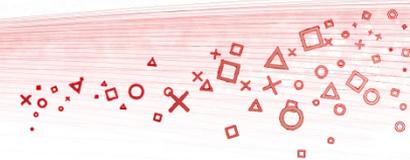
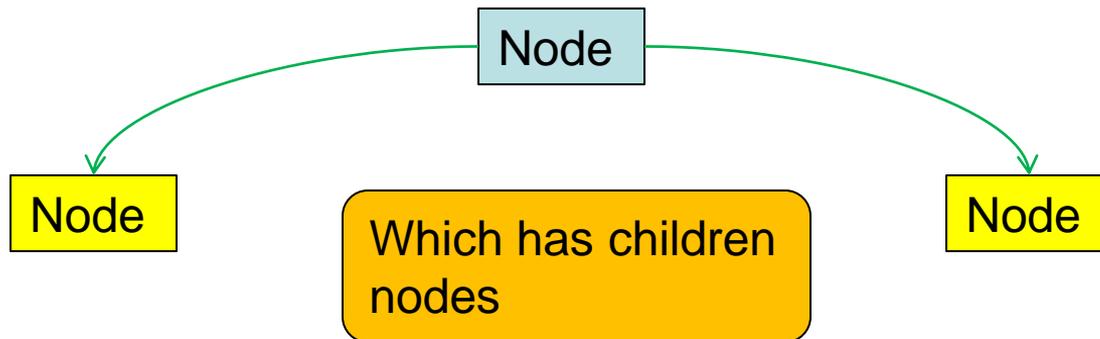
Hierarchy

We start with
a parent Node

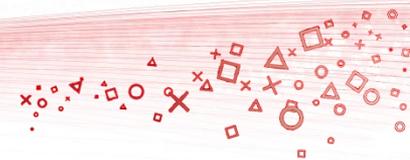
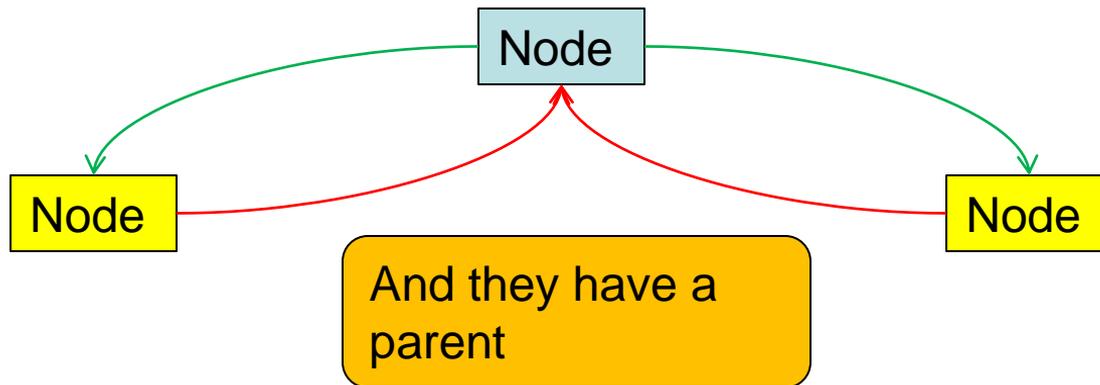
Node



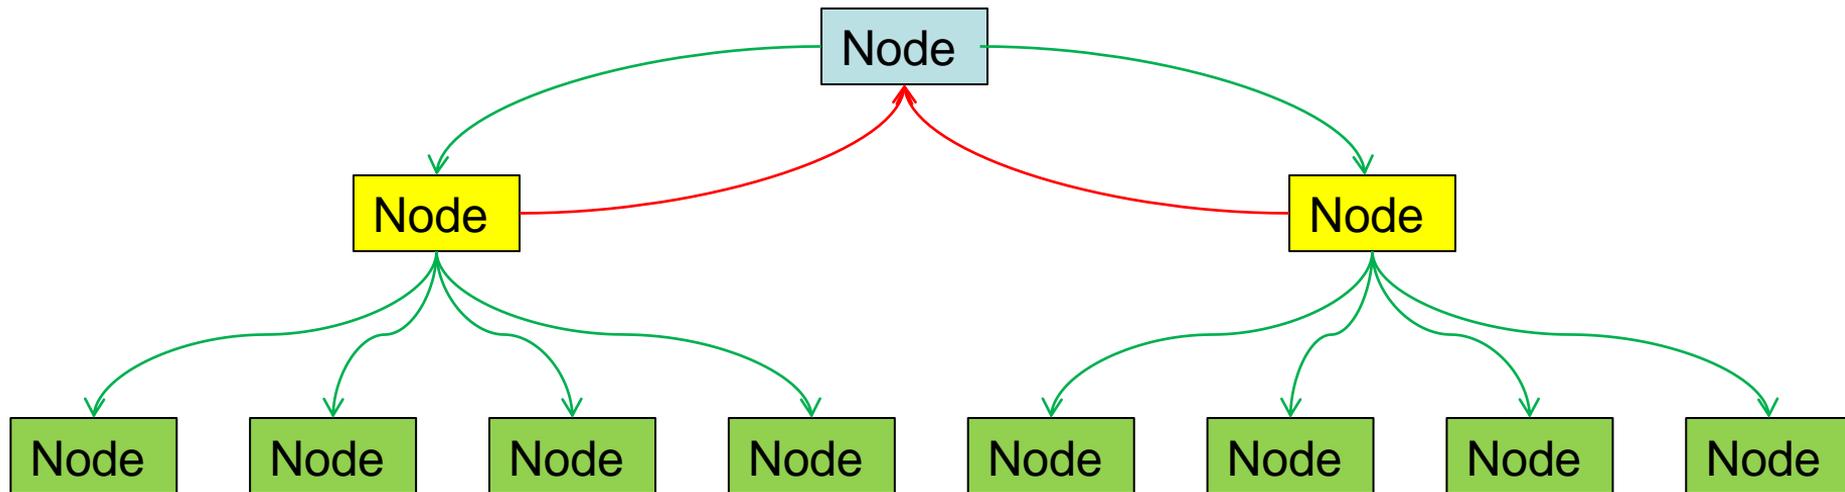
Hierarchy



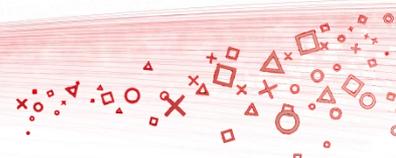
Hierarchy



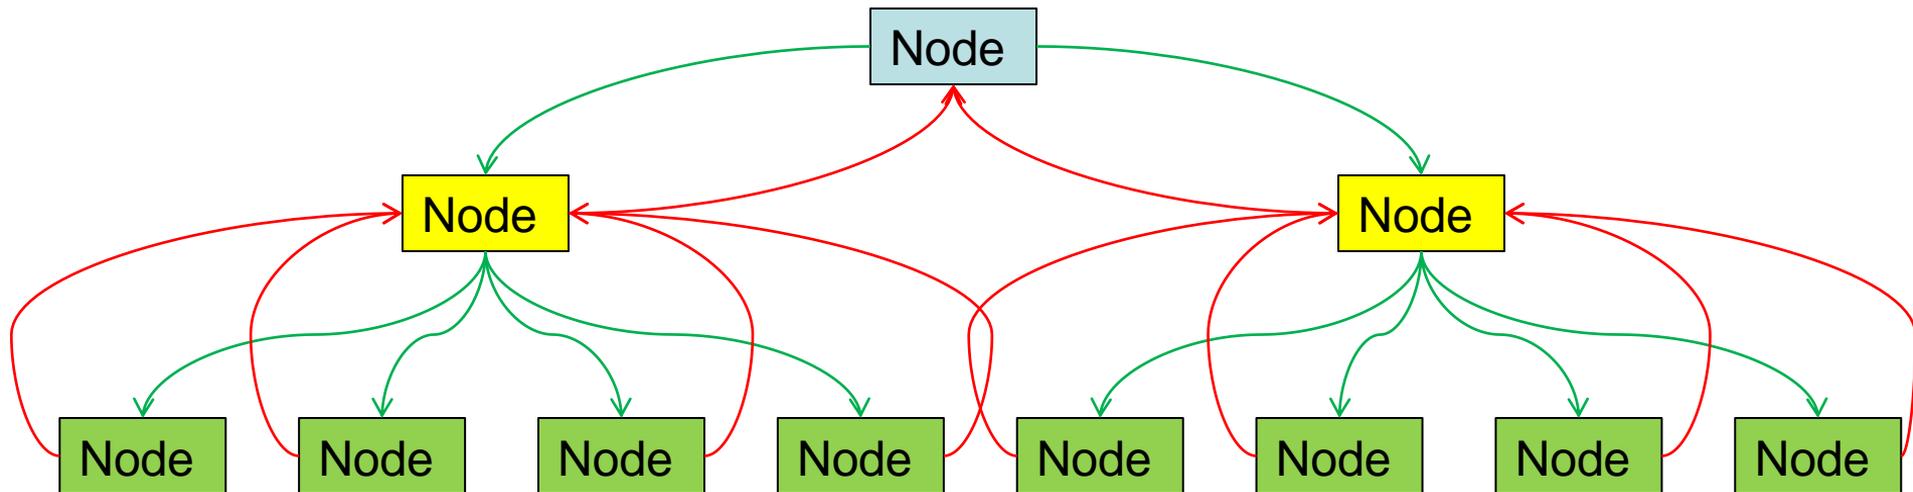
Hierarchy



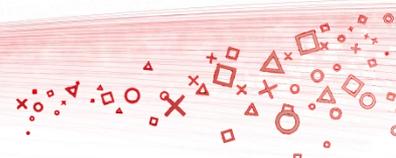
And they have children



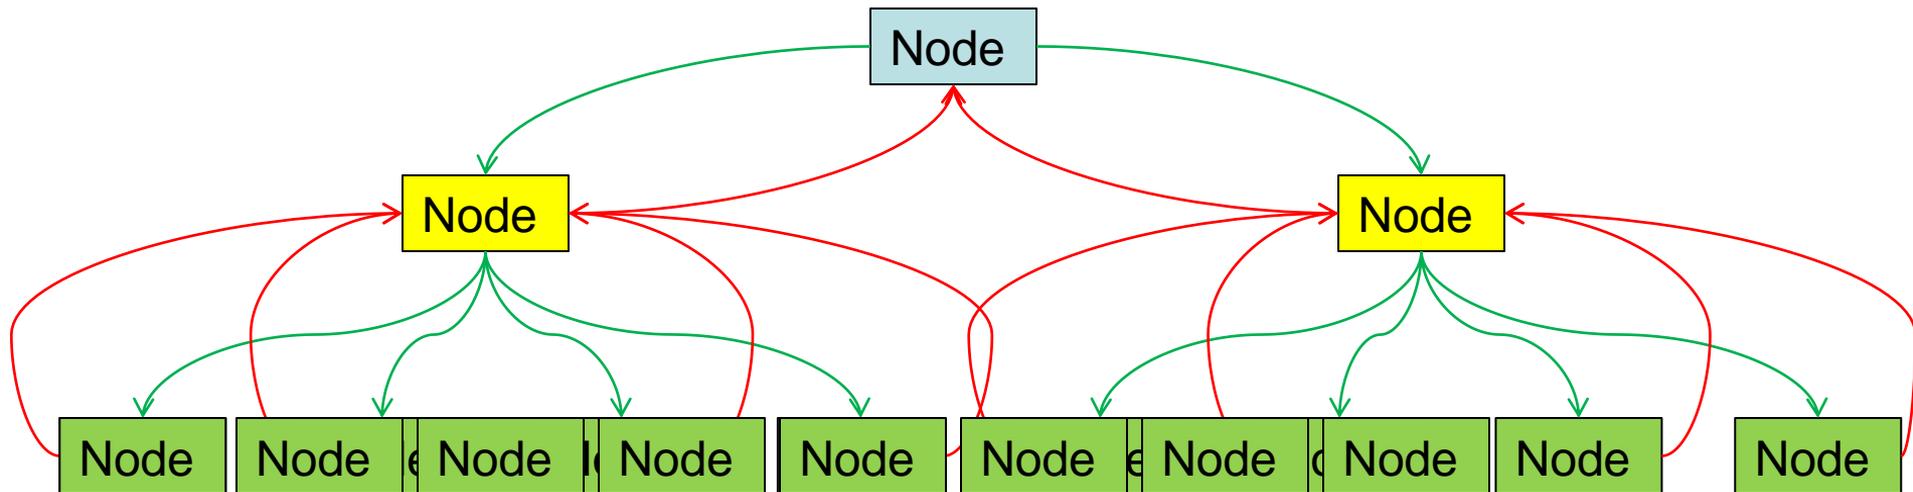
Hierarchy



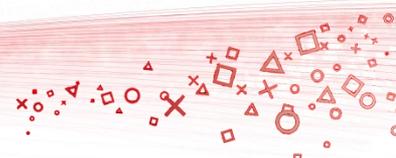
And they all have parents



Hierarchy



A lot of this information can be inferred



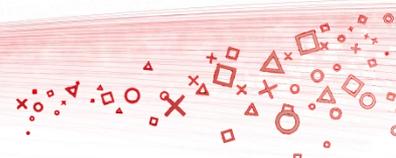
Hierarchy

Node

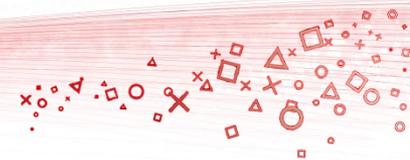
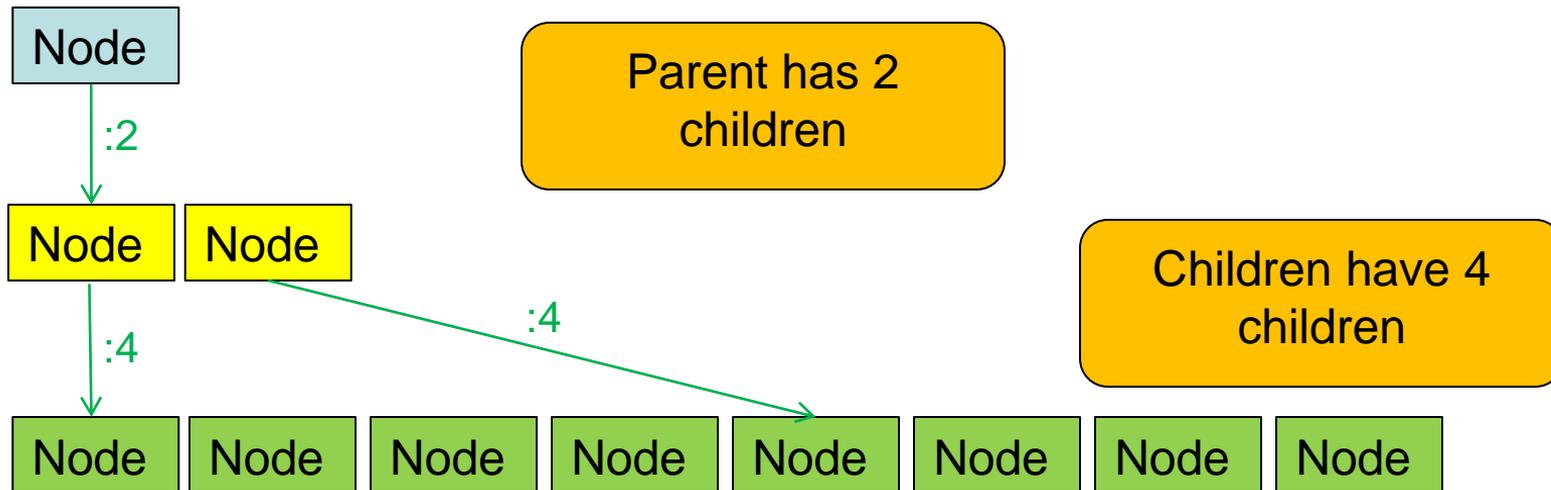
Use a set of arrays,
one per hierarchy
level

Node Node

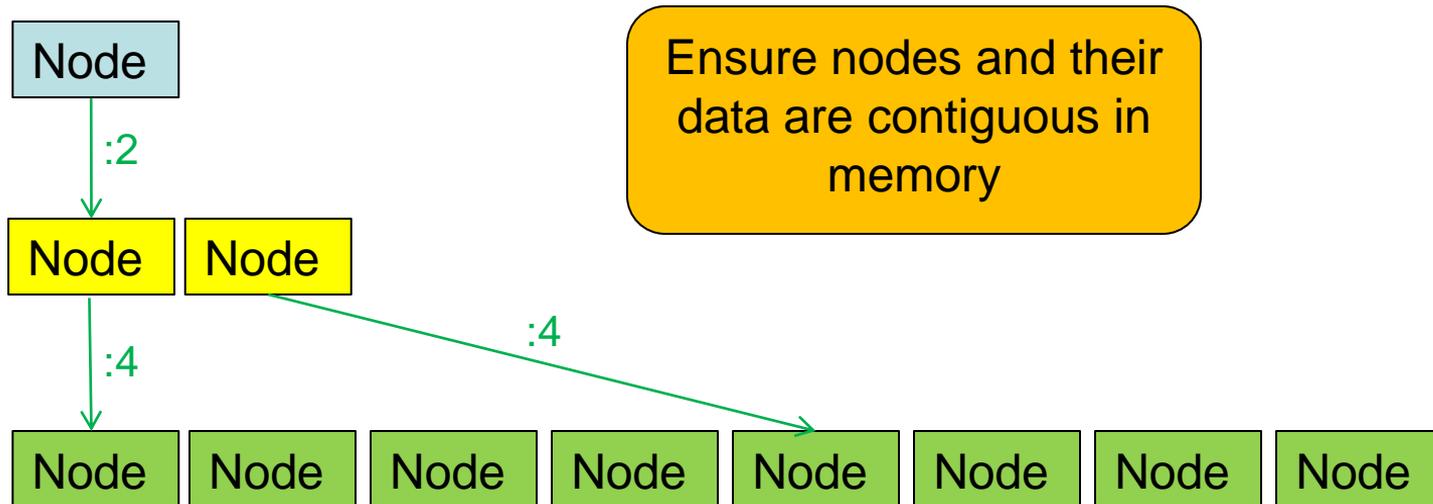
Node Node Node Node Node Node Node Node



Hierarchy



Hierarchy



Ensure nodes and their data are contiguous in memory



- Make the processing global rather than local
 - Pull the updates out of the objects.
 - No more virtuals
 - Easier to understand too – all code in one place.

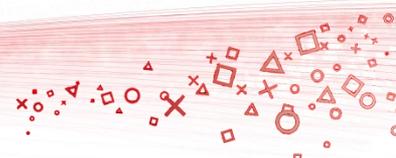
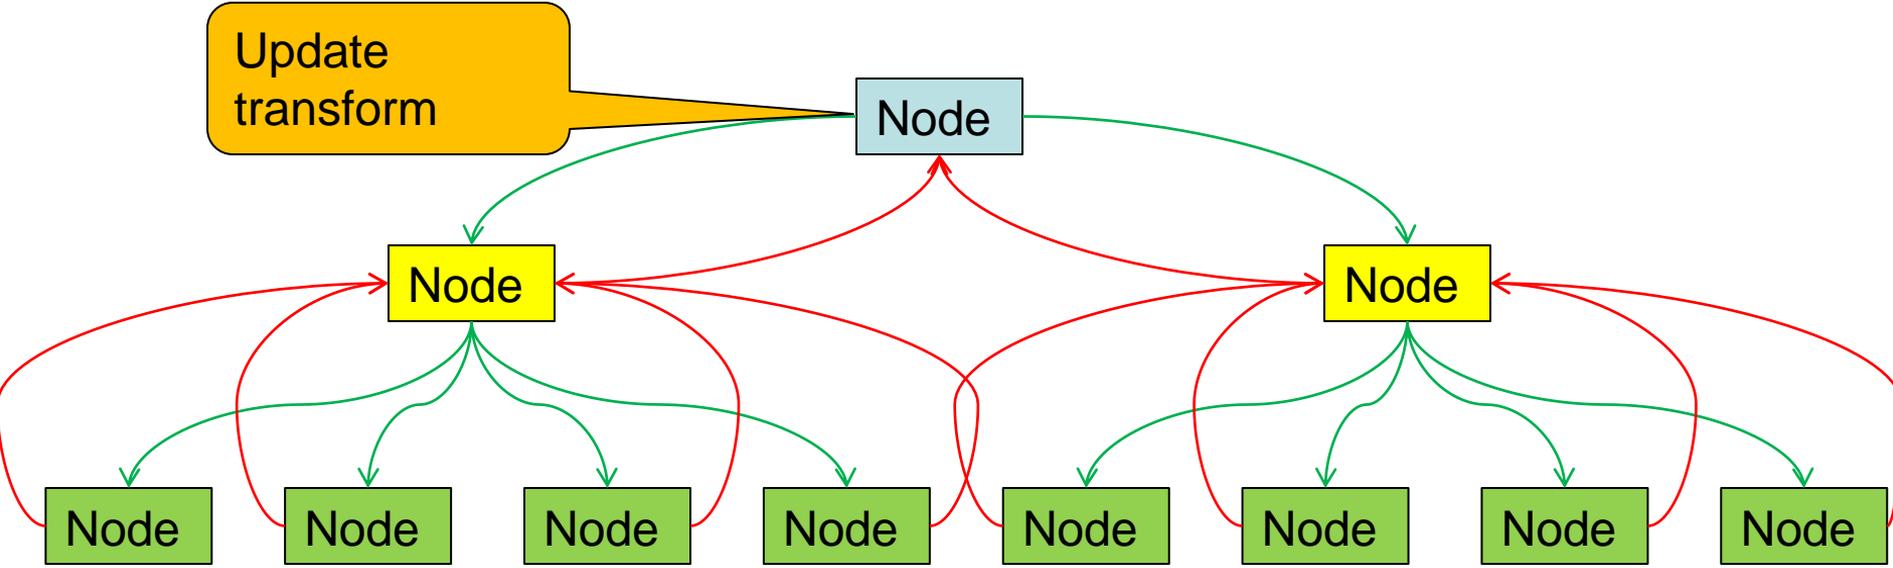


Need to change some things...

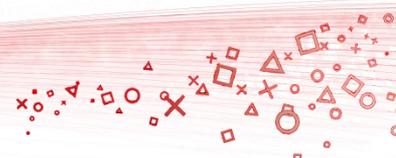
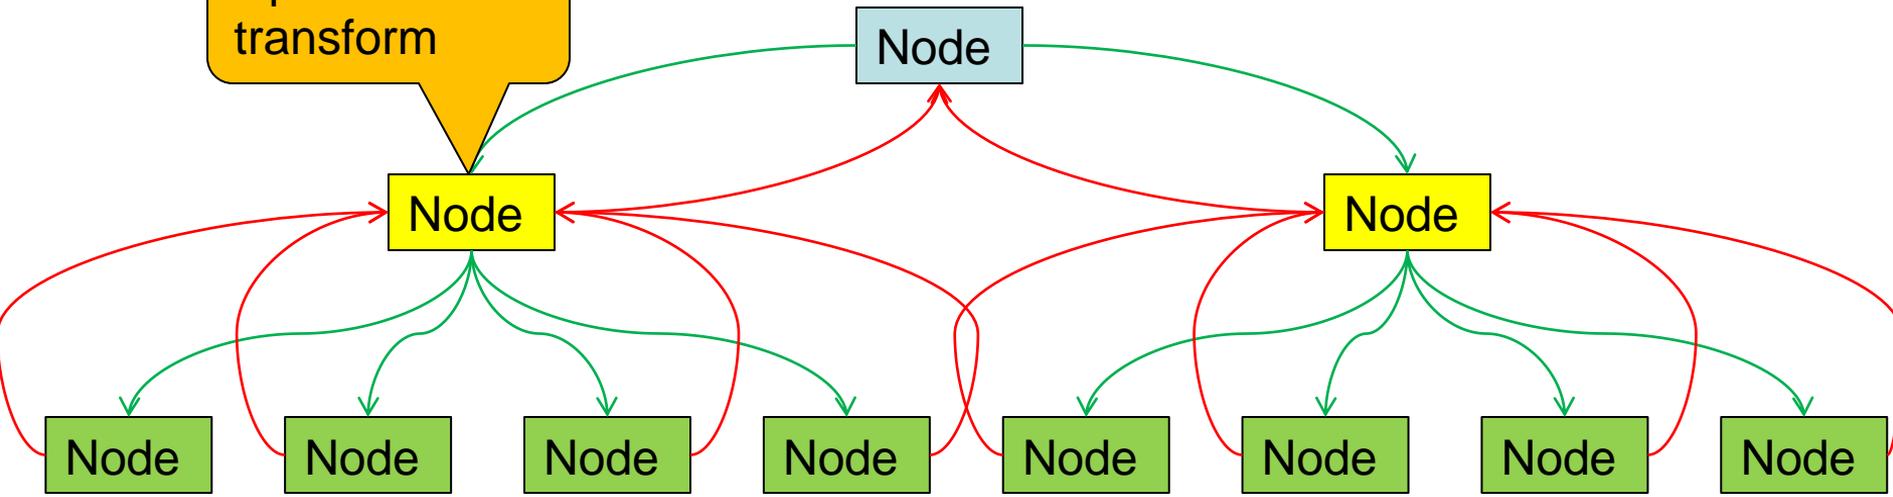
- OO version
 - Update transform top down and expand WBS bottom up

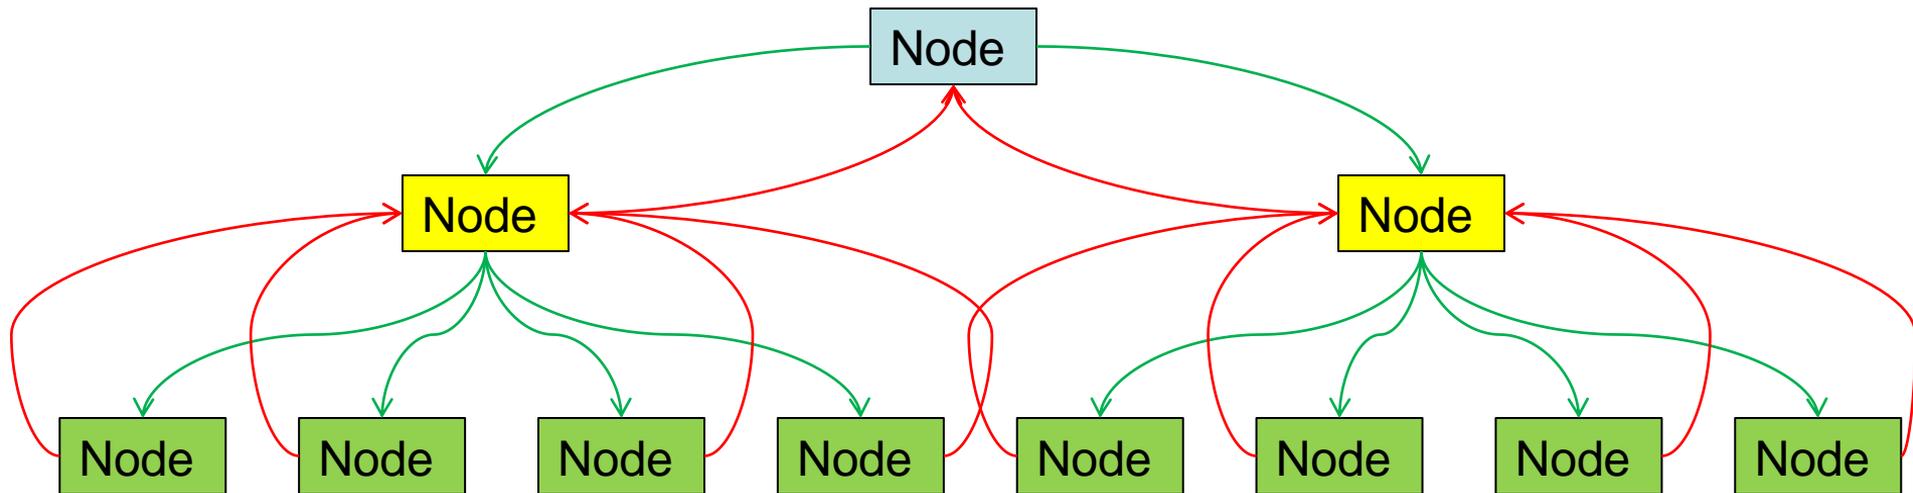


Update transform

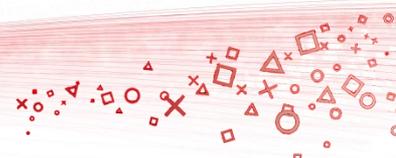


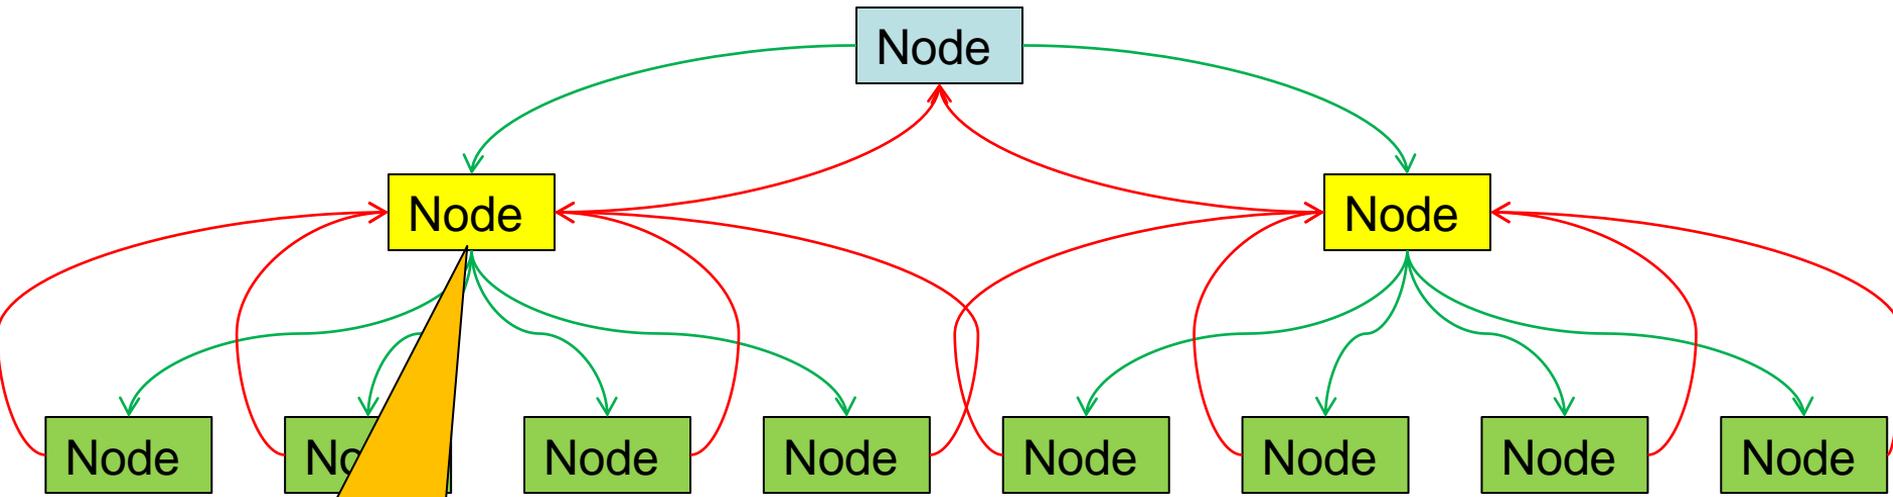
Update transform



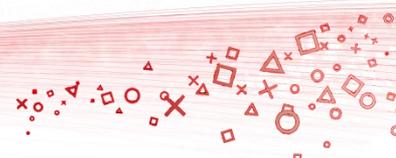


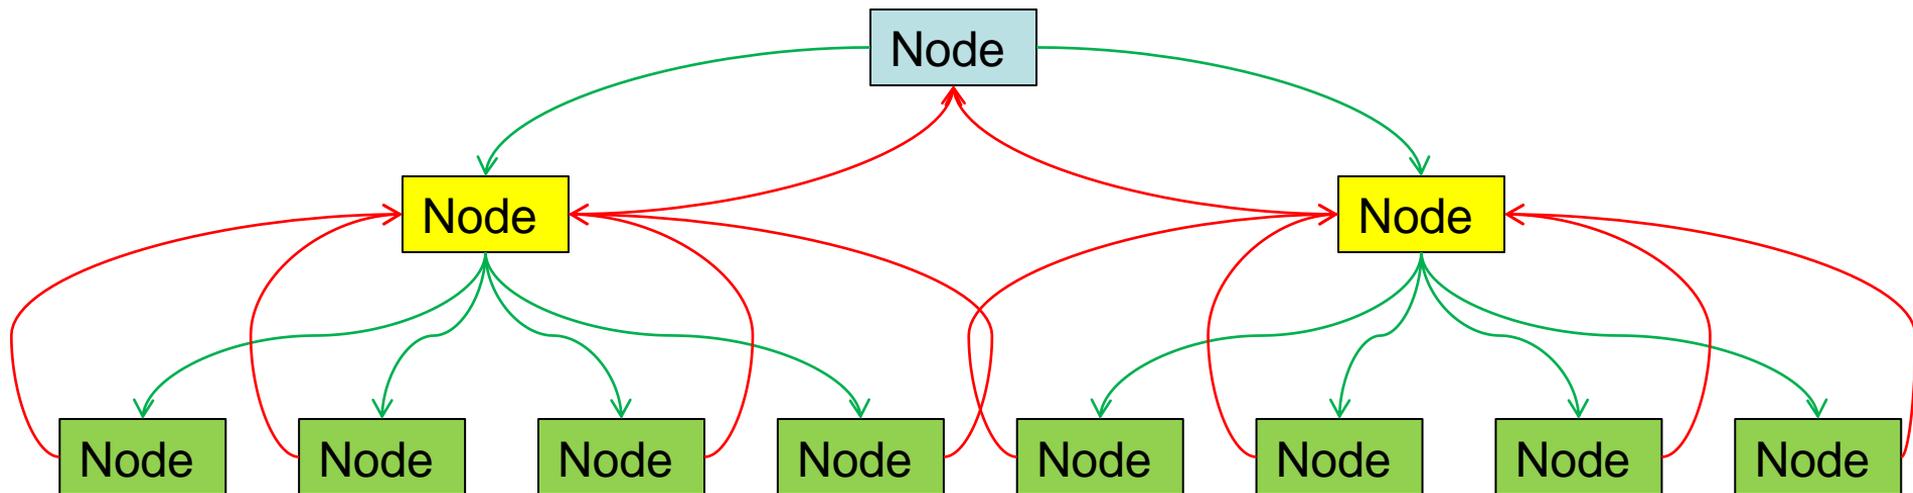
Update transform and world bounding sphere



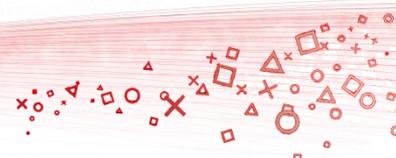


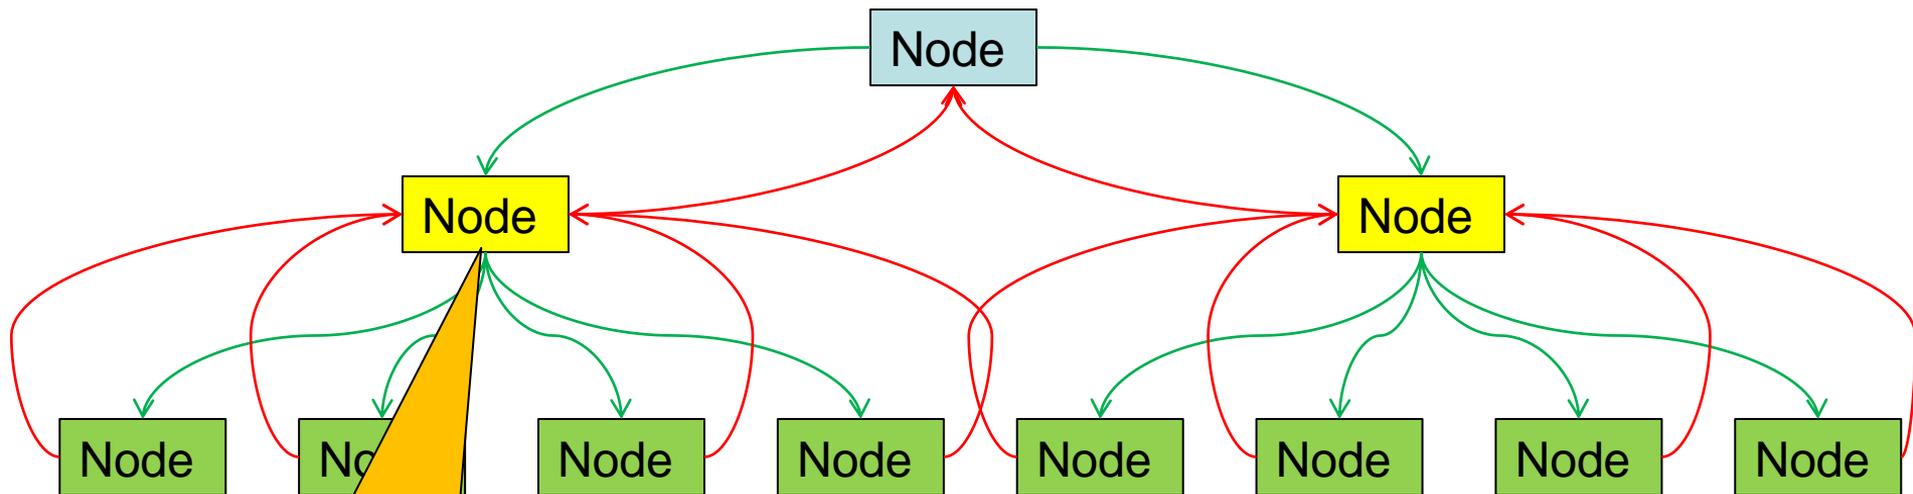
Add bounding sphere of child



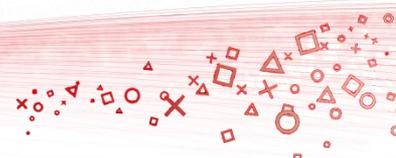


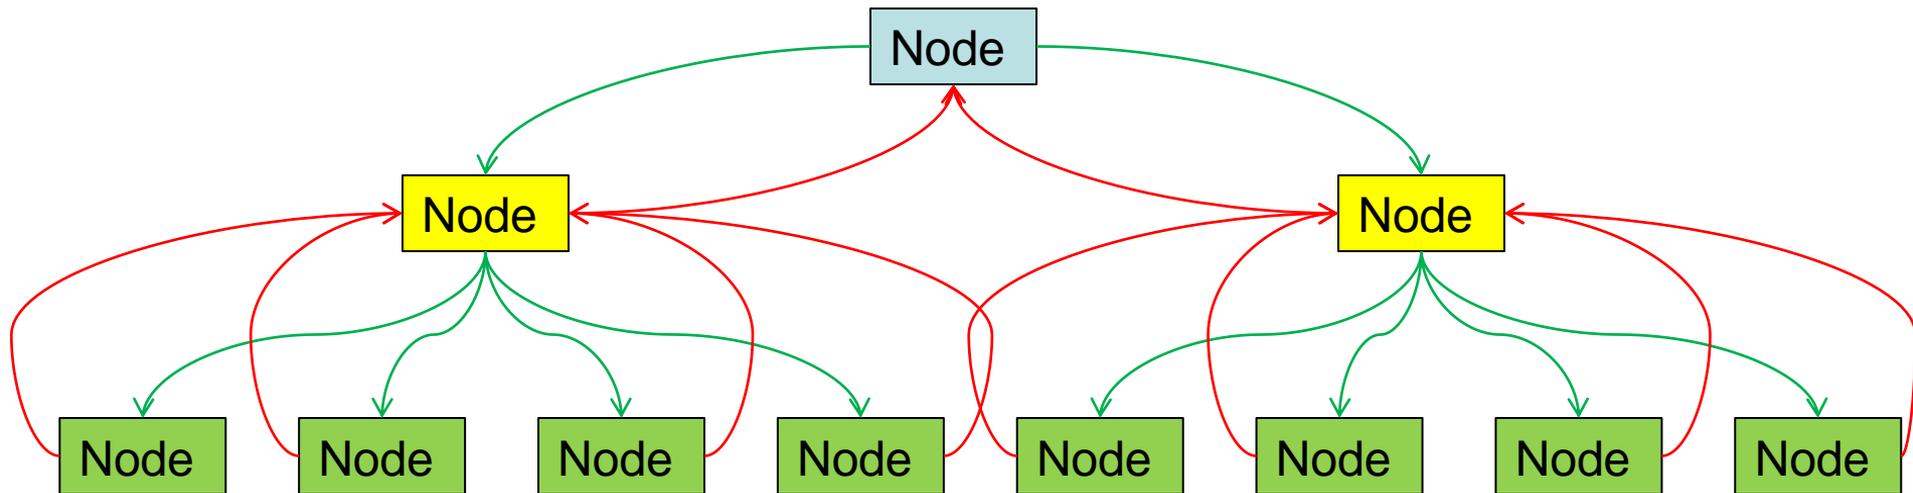
Update transform and world bounding sphere



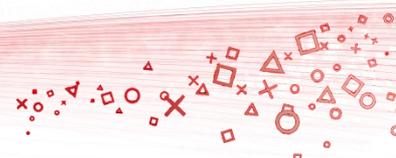


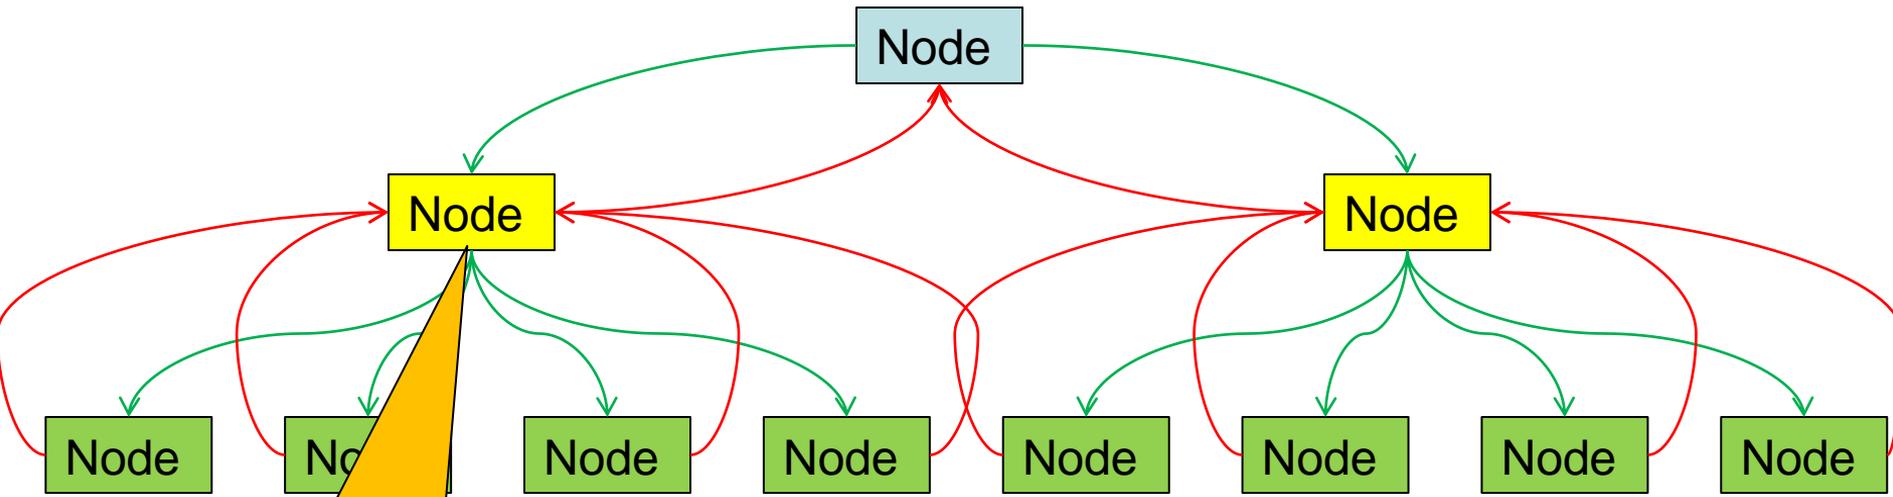
Add bounding sphere of child



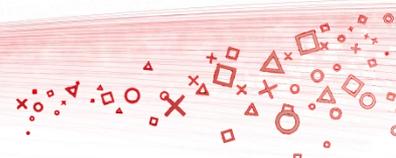


Update transform and world bounding sphere





Add bounding sphere of child



- Hierarchical bounding spheres pass info up
- Transforms cascade down
- Data use and code is ‘striped’.
 - Processing is alternating

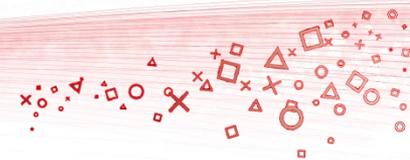
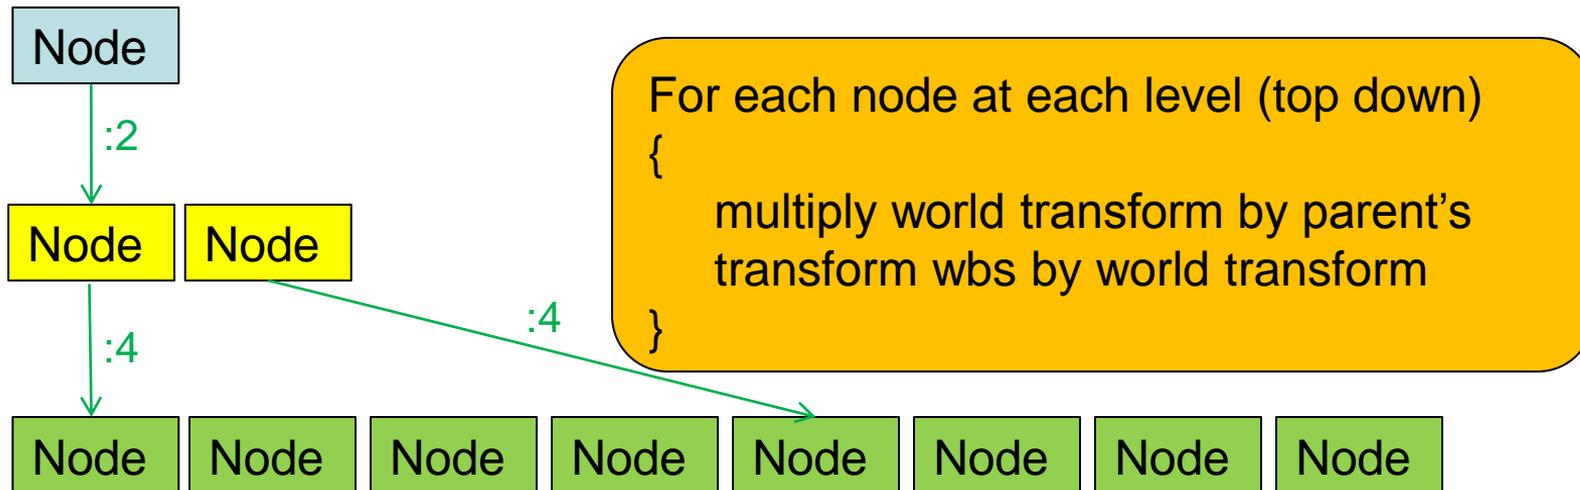


Conversion to linear

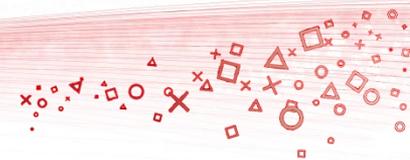
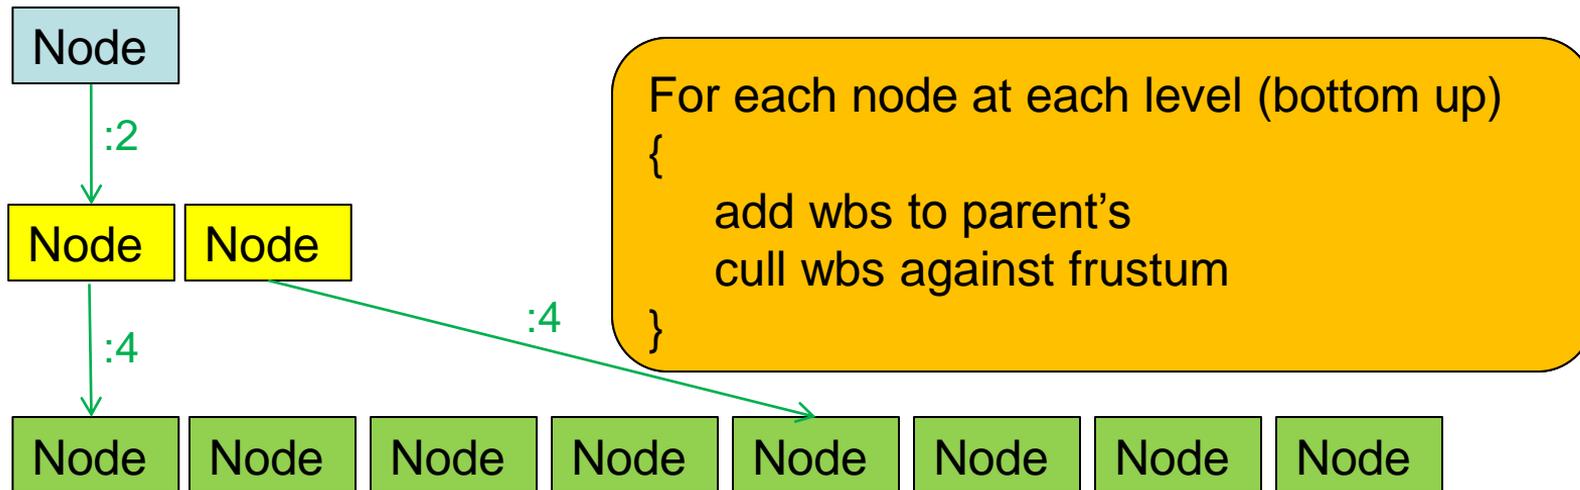
- To do this with a 'flat' hierarchy, break it into 2 passes
 - Update the transforms and bounding spheres (from top down)
 - Expand bounding spheres (bottom up)



Transform and BS updates



Update bounding sphere hierarchies



Update Transform and Bounding Sphere

How many children nodes
to process

```
for(int j=0;j<size;j++)
{
    const int innerSize = parent->m_Objects.size();
    const Matrix4 *parentTransform = parent->m_worldTransform;

    j+=innerSize;
    for(int k=0;k<innerSize;k++, wmat++, mat++, bs++, wbs++)
    {
        *wmat = (*parentTransform)*(*mat);
        *wbs = bs->Transform(wmat);
    }
    parent++;
}
```



Update Transform and Bounding Sphere

```
for(int j=0;j<size;j++)  
{  
    const int innerSize = parent->m_Objects.size();  
    const Matrix4 *parentTransform = parent->m_worldTransform;  
  
    j+=innerSize;  
    for(int k=0;k<innerSize;k++, wmat++, mat++, bs++, wbs++)  
    {  
        *wmat = (*parentTransform)*(*mat);  
        *wbs = bs->Transform(wmat);  
    }  
    parent++;  
}
```

For each child, update transform and bounding sphere



Update Transform and Bounding Sphere

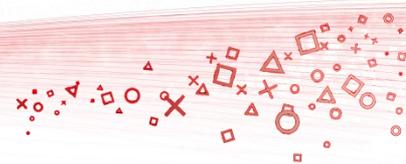
```

for(int j=0;j<size;j++)
{
    const int innerSize = parent->m_Objects.size();
    const Matrix4 *parentTransform = parent->m_worldTransform;

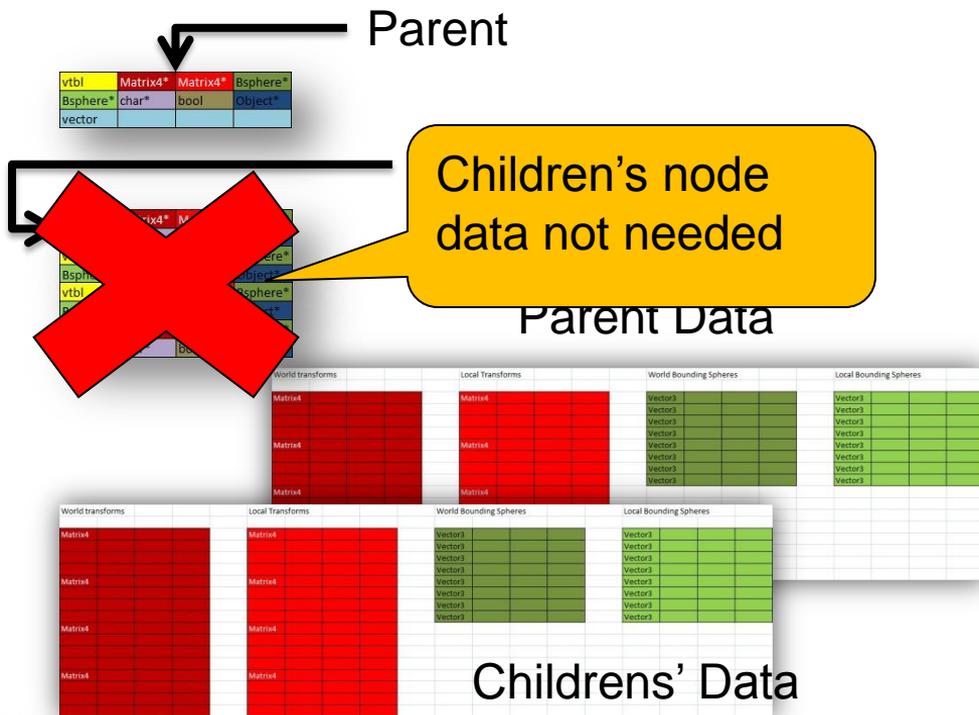
    j+=innerSize;
    for(int k=0;k<innerSize;k++, wmat++, mat++, bs++, wbs++)
    {
        *wmat = (*parentTransform)*(*mat);
        *wbs = bs->Transform(wmat);
    }
    parent++;
}

```

Note the contiguous arrays

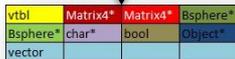


So, what's happening in the cache?





Load parent and its transform



```
const int innersize = parent->m_Objects.size();
const Matrix4 *parentTransform = parent->m_worldTransform;
```

Unified L2 Cache



Parent Data



Childrens' Data



Load child BS and set WBS

Next child is calculated with no extra cache misses !

```

    vtbl Matrix4* Matrix4* Bsphere*
    Bsphere* char* bool Object*
    vector

    Parent
    ↓
    for(int k=0;k<inners
    {
        *wmat = (*parentTransform)*(*mat);
        *wbs = bs->Transform(wmat);
    }
  
```

Parent Data



Childrens' Data

Unified L2 Cache



Load child BS and set WBS

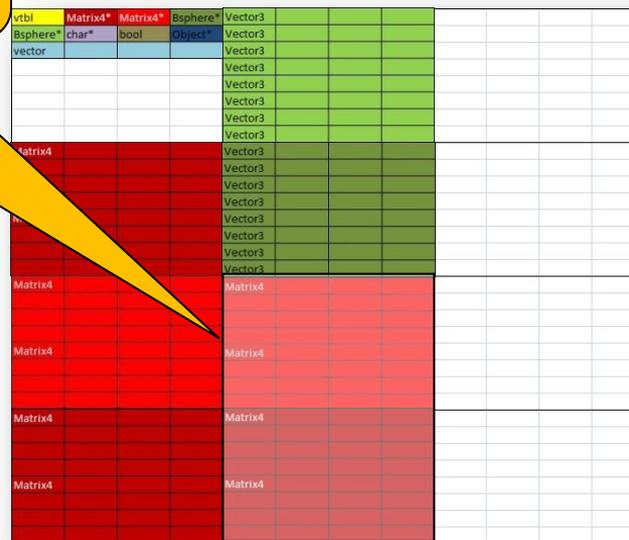
The next 2 children incur 2 cache misses in total

vtbl	Matrix4*	Matrix4*	Bsphere*
Bsphere*	char*	bool	Object*
vector			

Parent

```
for(int k=0;k<inners
{
    *wmat = (*parentTransform)*(*mat);
    *wbs = bs->Transform(wmat);
}
```

Unified L2 Cache

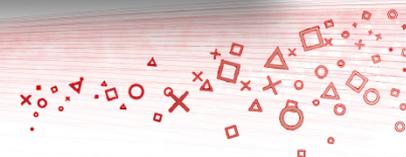


Parent Data

World transforms	Local Transforms	World Bounding Spheres	Local Bounding Spheres
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3

World transforms	Local Transforms	World Bounding Spheres	Local Bounding Spheres
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3
Matrix4	Matrix4	Vector3	Vector3

Childrens' Data



Prefetching

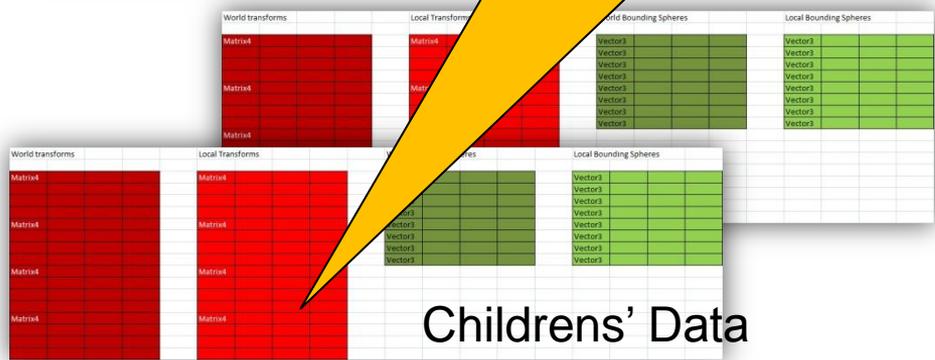
```

    Parent
    ↓
    Matrix4* Matrix4* Bsphere*
    char*    bool    Object*
    vector

    for(int k=0;k<inners
    {
        *wmat = (*parent
        *wbs = bs->Trans
    }
  
```

Because all data is linear, we can predict what memory will be needed in ~400 cycles and prefetch

Parent's Data



Children's Data

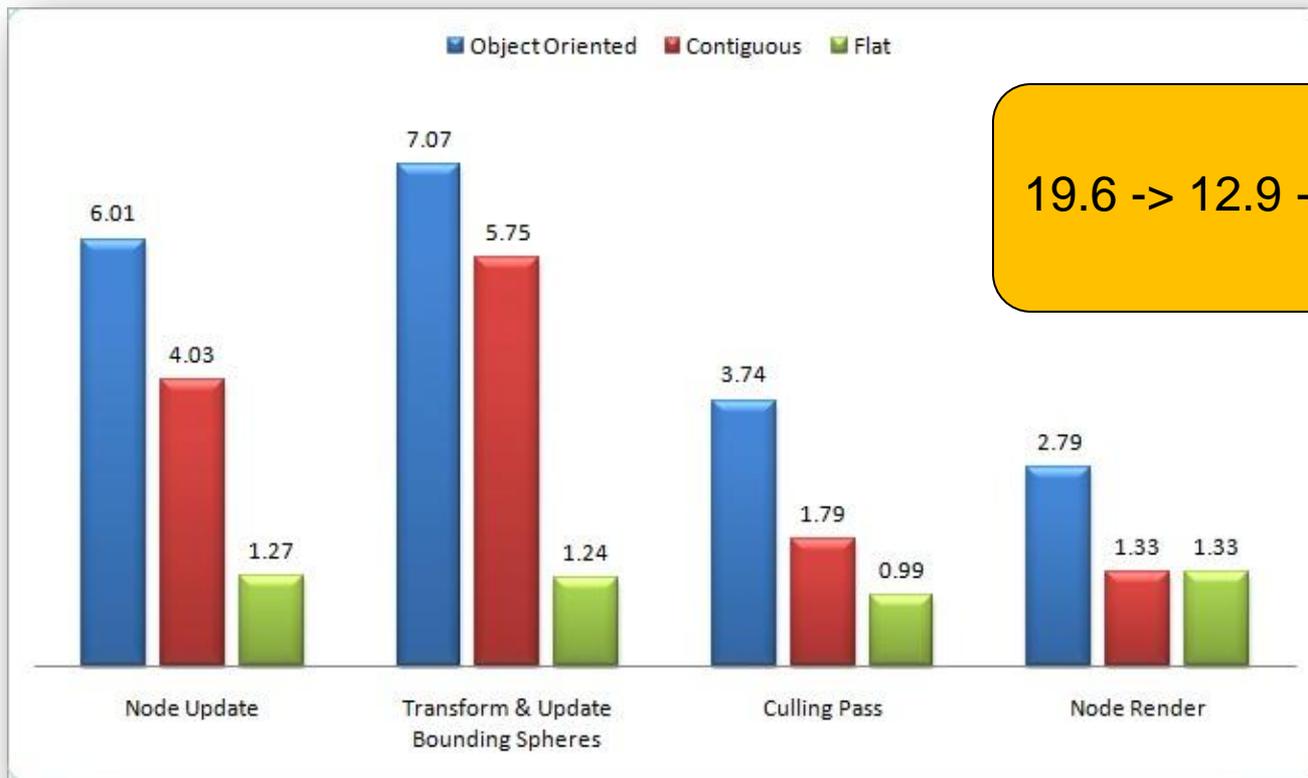
Unified L2 Cache



- Tuner scans show about 1.7 cache misses per node.
- But, these misses are much more frequent
 - Code/cache miss/cache miss/code
 - Less stalling



Performance



19.6 -> 12.9 -> 4.8ms



Prefetching

- Data accesses are now predictable
- Can use prefetch (dcbt) to warm the cache
 - Data streams can be tricky
 - Many reasons for stream termination
 - Easier to just use dcbt blindly
 - (look ahead x number of iterations)



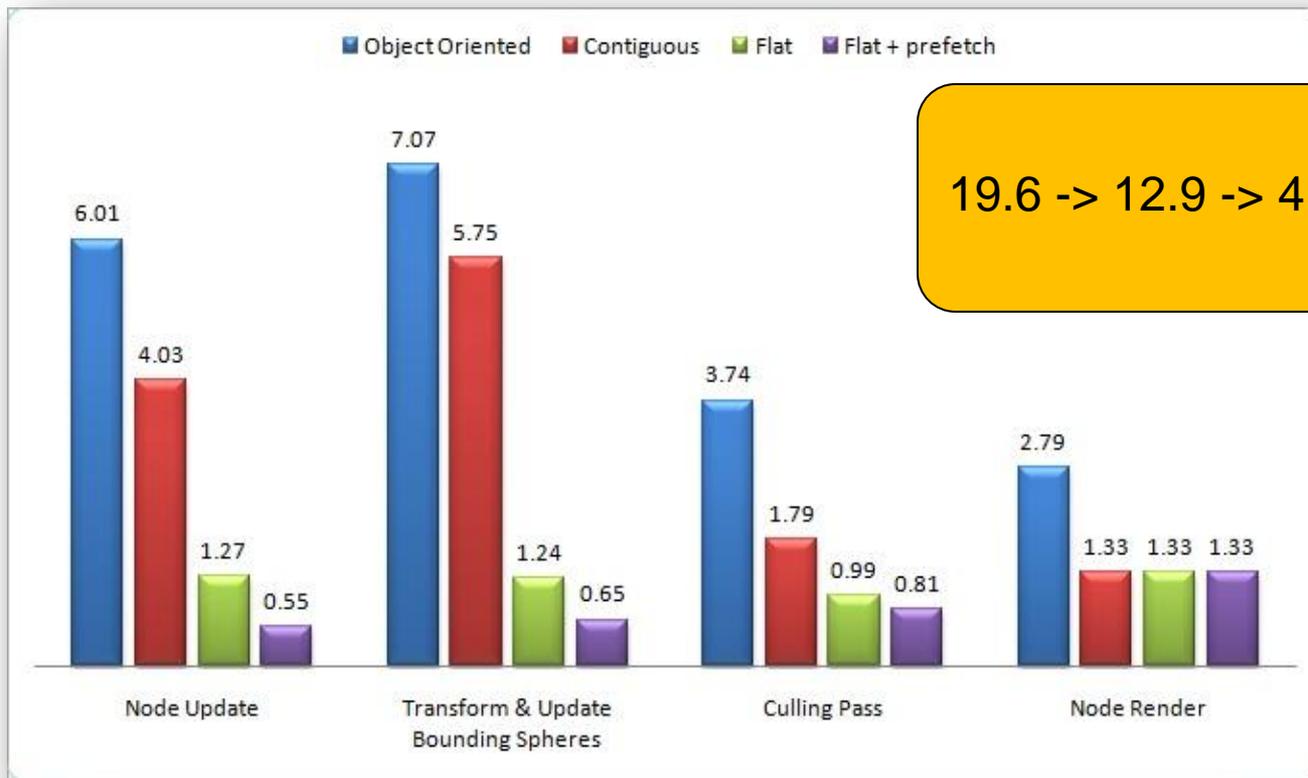
Prefetching example

- Prefetch a predetermined number of iterations ahead
- Ignore incorrect prefetches

```
for(int i=0;i<m_CurrentNumTransforms;i++)  
{  
    __dcbt(m_Transforms[i+20]);  
    *m_Transforms[i] = (*m_Transforms[i])*(*m_Transform);  
}
```



Performance



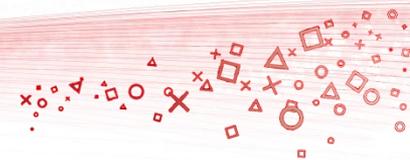
A Warning on Prefetching

- This example makes very heavy use of the cache
- This can affect other threads' use of the cache
 - Multiple threads with heavy cache use may thrash the cache



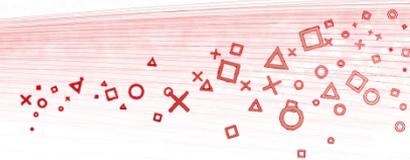
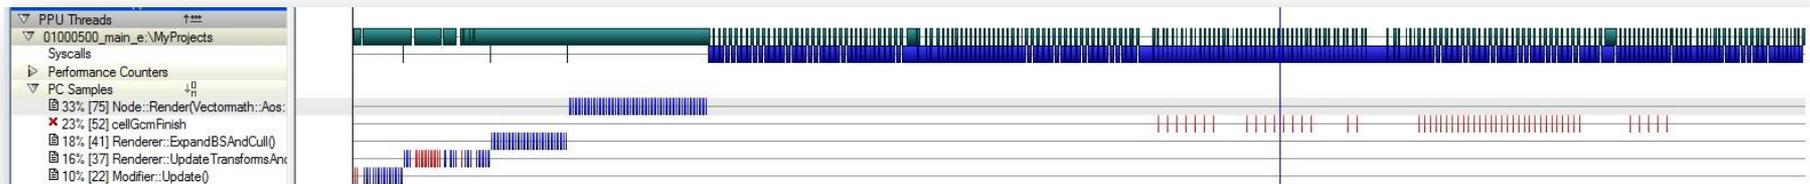
The old scan

~22ms

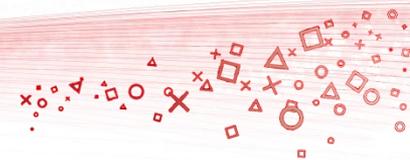
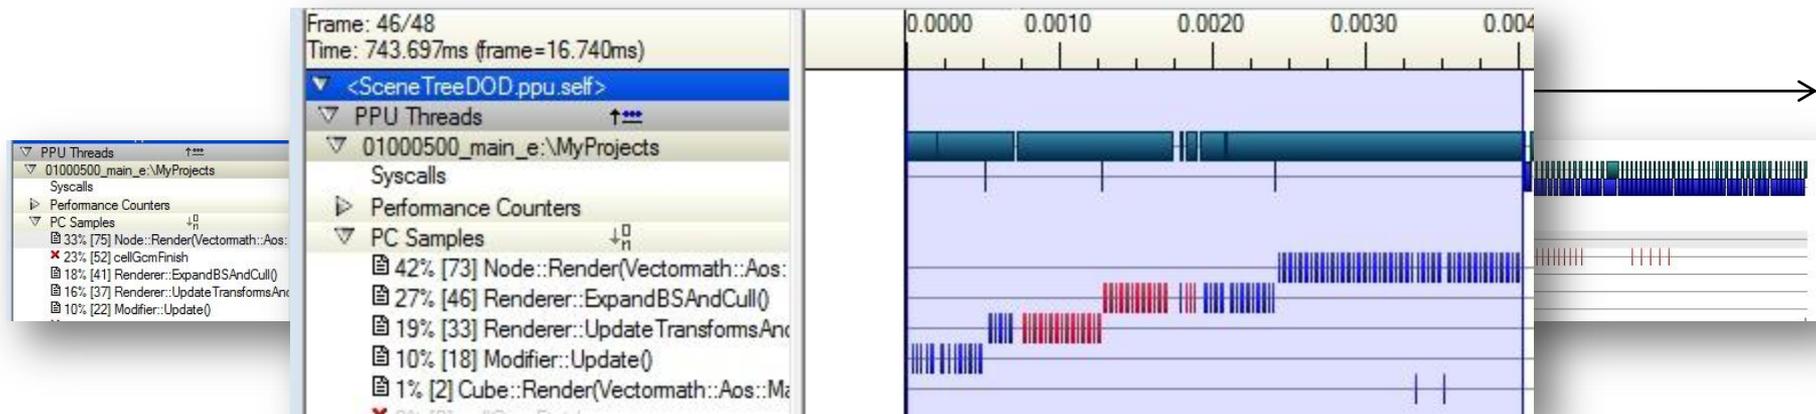


The new scan

~16.6ms



Up close



Looking at the code (samples)

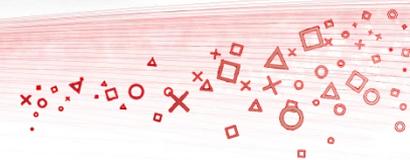
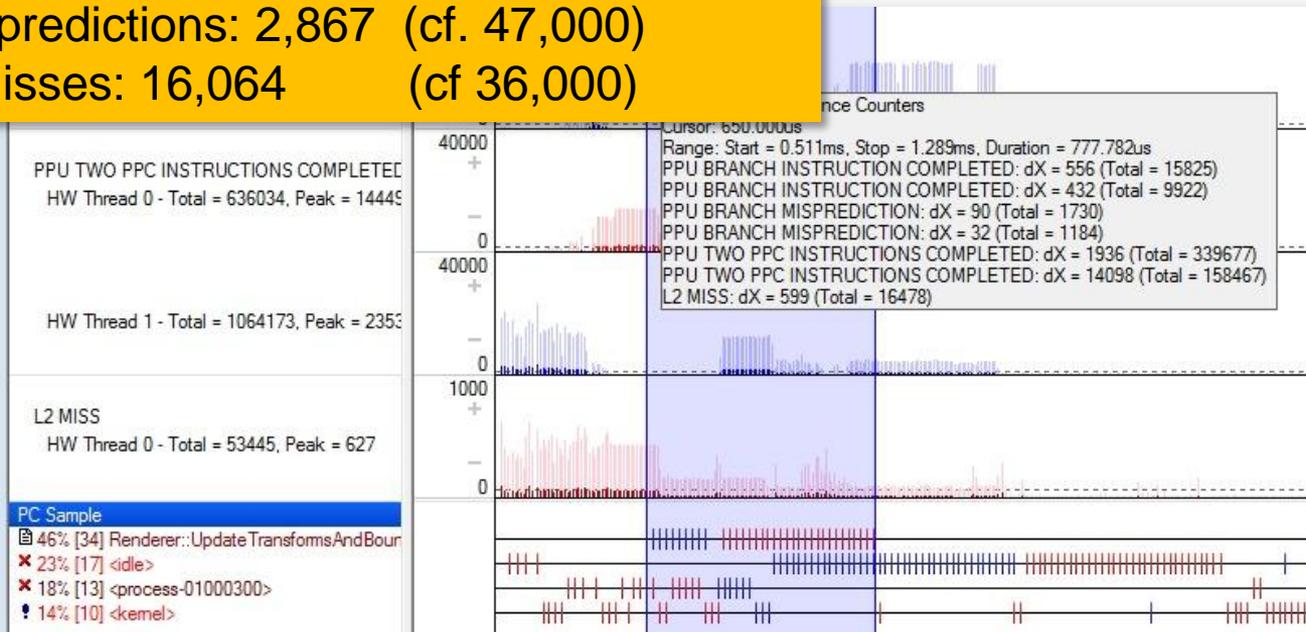
```
- - 449
151 3 450      const Node* parent = (Node*)node->m_Parent;
- - 451      // iterate through all the matrices at this level, multiplying them by their parent
8 - 452      for(int j=0;j<size;j++)//,node++, wmat++, mat++, wbs++, bs++)
- - 453      {
- - 454          const int innerSize = parent->m_Objects.size();
25 - 455          const Matrix4 *parentTransform = parent->m_WorldTransform;
- - 456
2 - 457          j+=innerSize;
372 9 458          for(int k=0;k<innerSize;k++, wmat++, mat++, bs++, wbs++, node++)
- - 459          {
233 2 460              __dcbt(wmat+lookAhead);
146 3 461              __dcbt(mat+lookAhead);
317 5 462              __dcbt(bs+lookAhead);
117 4 463              __dcbt(wbs+lookAhead);
- - 464              *wmat = (*parentTransform)*(*mat);
- - 465              *wmat = (*node->m_Parent->m_WorldTransform)*(*mat);
- - 466
32 1 467              *wbs = bs->Transform(wmat);
- - 468          }
- - 469          parent++;
```



Performance counters

Branch mispredictions: 2,867 (cf. 47,000)

L2 cache misses: 16,064 (cf 36,000)



In Summary

- Just reorganising data locations was a win
- Data + code reorganisation= dramatic improvement.
- + prefetching equals even more WIN.



OO is not necessarily EVIL

- Be careful not to design yourself into a corner
- Consider data in your design
 - Can you decouple data from objects?
 - ...code from objects?
- Be aware of what the compiler and HW are doing



Its all about the memory

- Optimise for data first, then code.
 - Memory access is probably going to be your biggest bottleneck
- Simplify systems
 - KISS
 - Easier to optimise, easier to parallelise



Homogeneity

- Keep code and data homogenous
 - Avoid introducing variations
 - Don't test for exceptions – sort by them.
- Not everything needs to be an object
 - If you must have a pattern, then consider using Managers



Remember

- You are writing a GAME
 - You have control over the input data
 - Don't be afraid to preformat it – drastically if need be.
- Design for specifics, not generics (generally).



Data Oriented Design Delivers

- Better performance
- Better realisation of code optimisations
- Often simpler code
- More parallelisable code



The END

