

# Parallel(la) Programming

## Concurrency, Deadlock and Interconnect




**James Pallister**

Research Engineer, Embecosm  
PhD Student, University of Bristol



# About me

- PhD studying *software-based* energy efficiency.
  - Currently looking at compiler technology
- Last summer:



# Outline

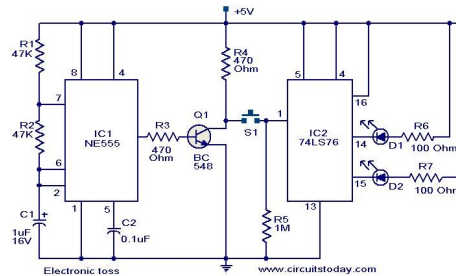
Parallelism

Deadlock & race conditions

Epiphany's interconnect

# Why Parallel?




The world is parallel.



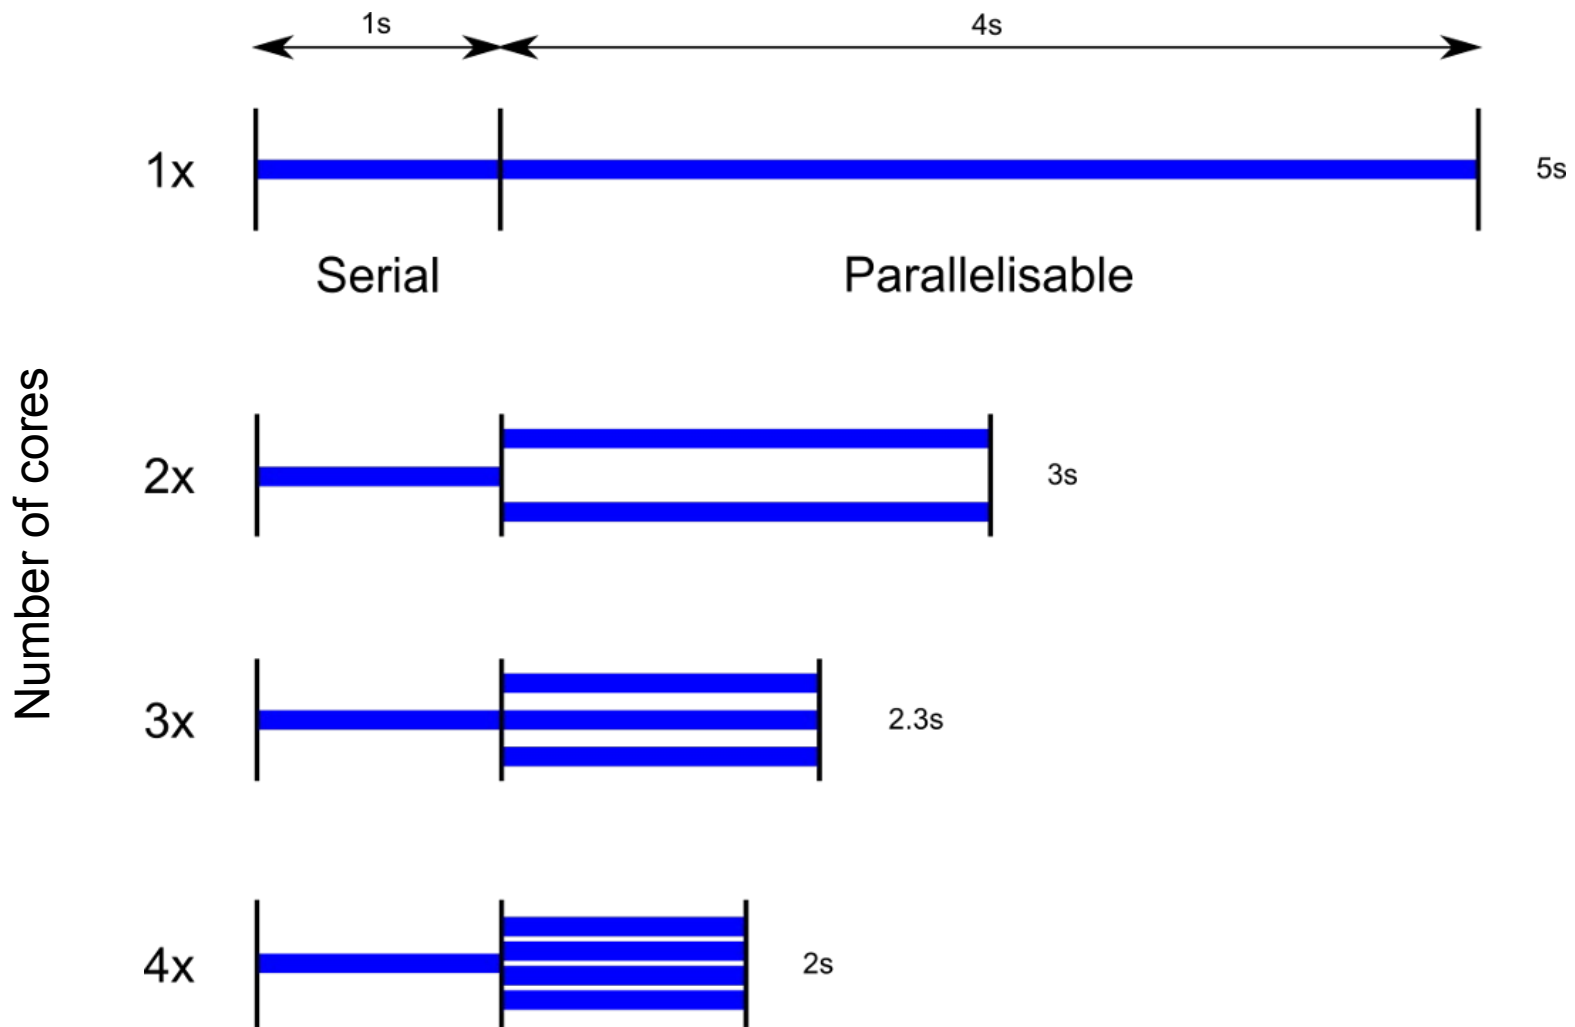
But when programming, we've been taught to think sequentially.

# What is Parallelism?

- Multiple things happening at the same time to progress towards a common goal.




# Amdahl's Law




# Amdahl's Law

Amdahl's Speed-up



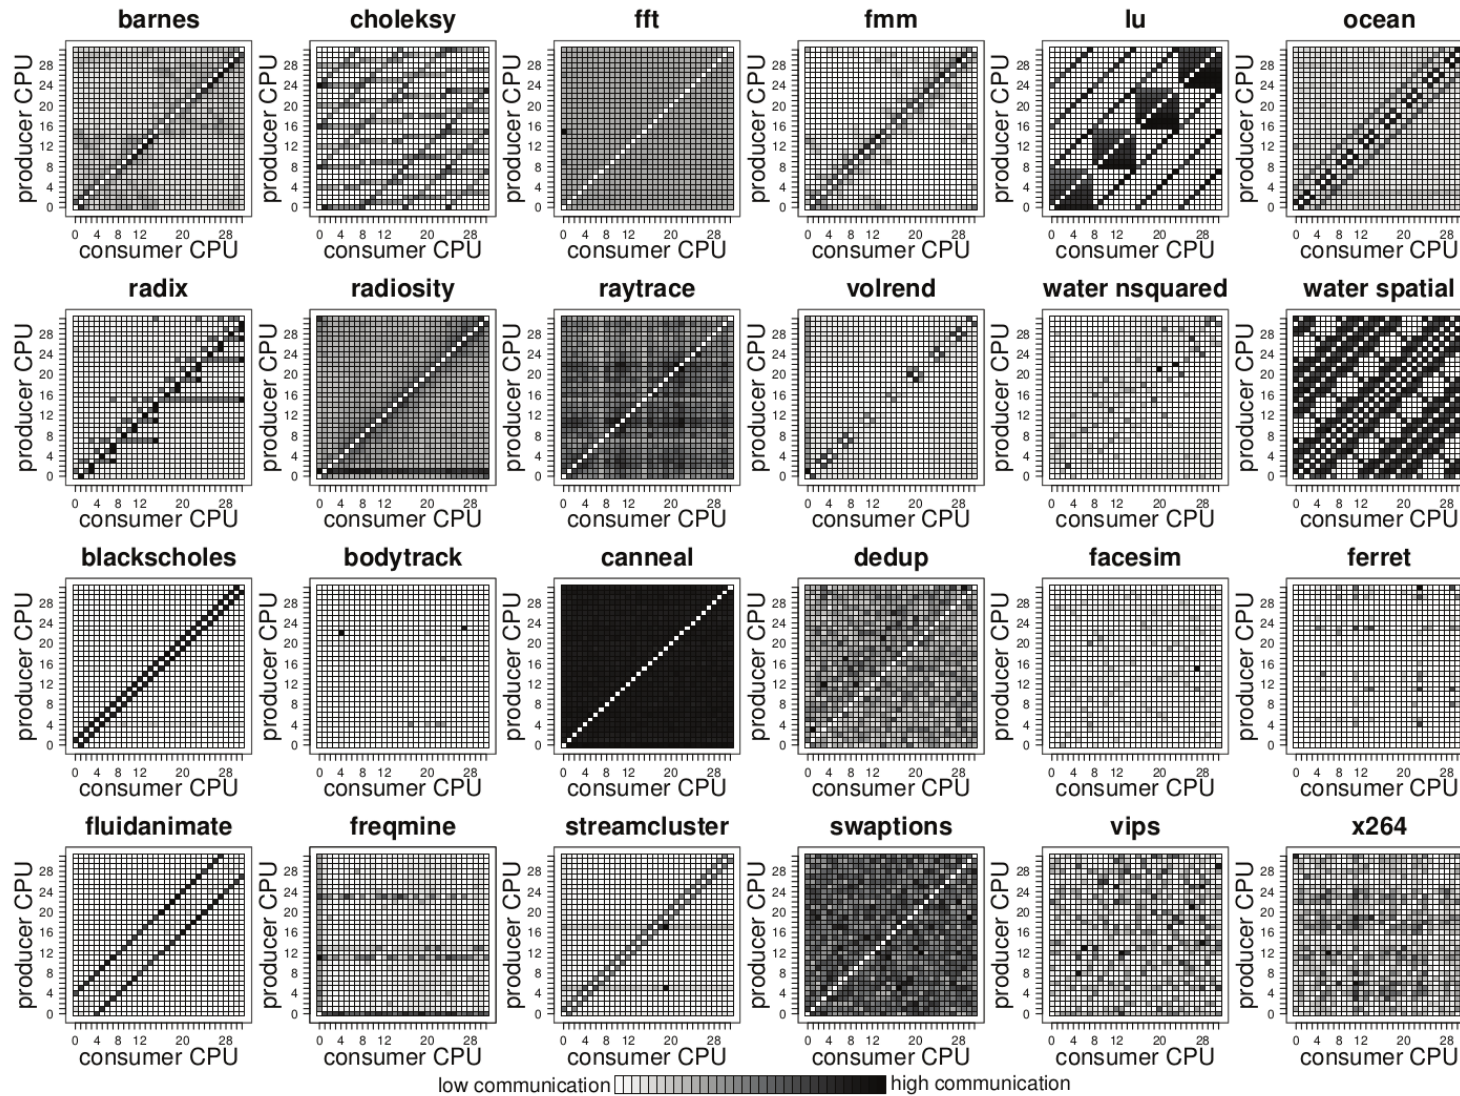
# Real life Amdahl's Law

Real life Amdahl's Law










# Communication Patterns




# Flynn's Taxonomy




# The 7 dwarves of parallel computation




Structured Grid




Monte Carlo




Spectral




Dense Linear Algebra



Unstructured Grid



N-Body

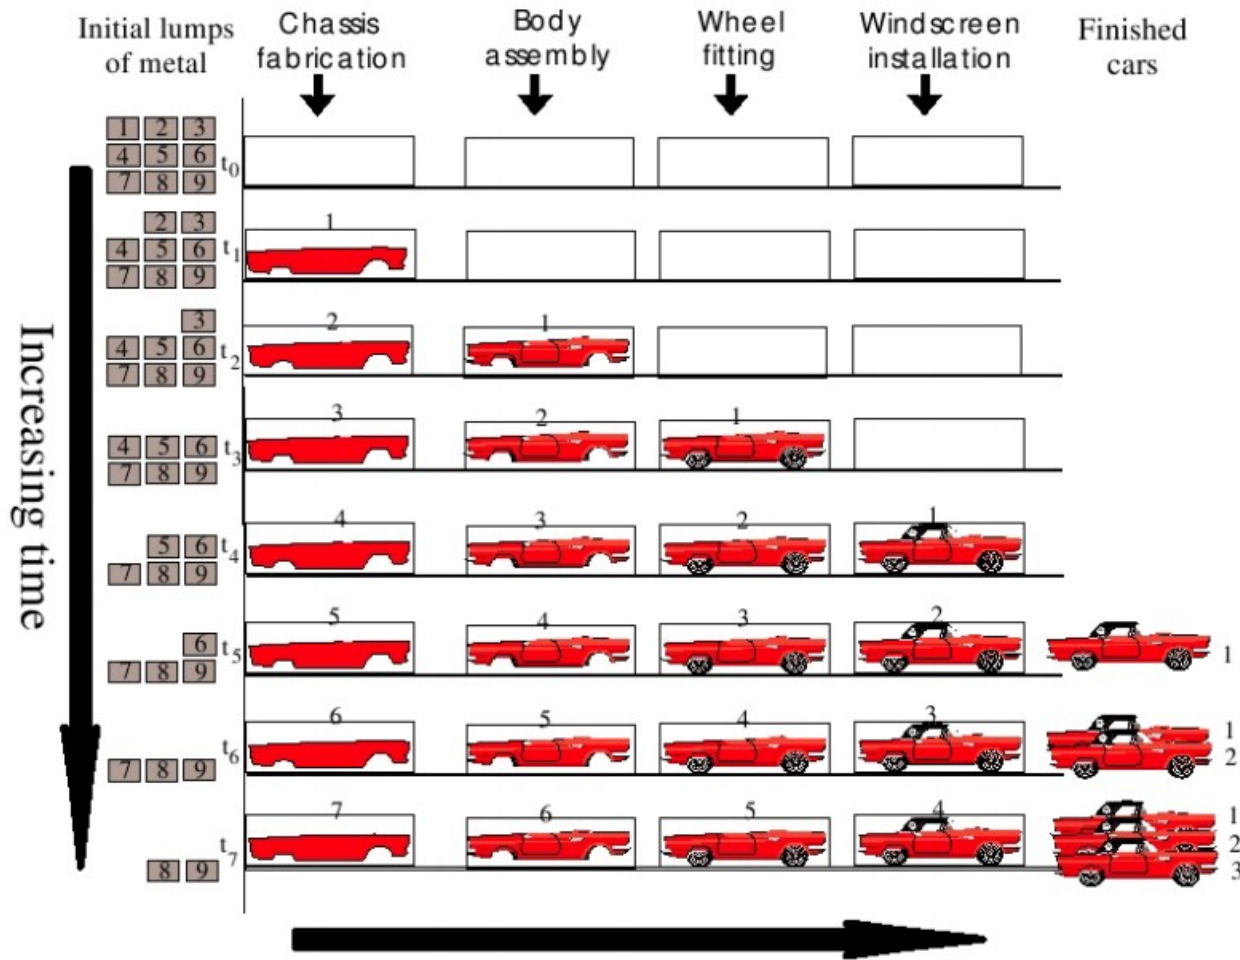


Sparse Linear Algebra

# Parallelism - Paradigms

- Pipelines
- Task farms
  - Server - client
- Geometric
  - Matrix multiply
  - Structured grid

# Pipelines




Motion of cars through pipeline

from "Alan Chalmers, Practical Parallel Processing, 1996"

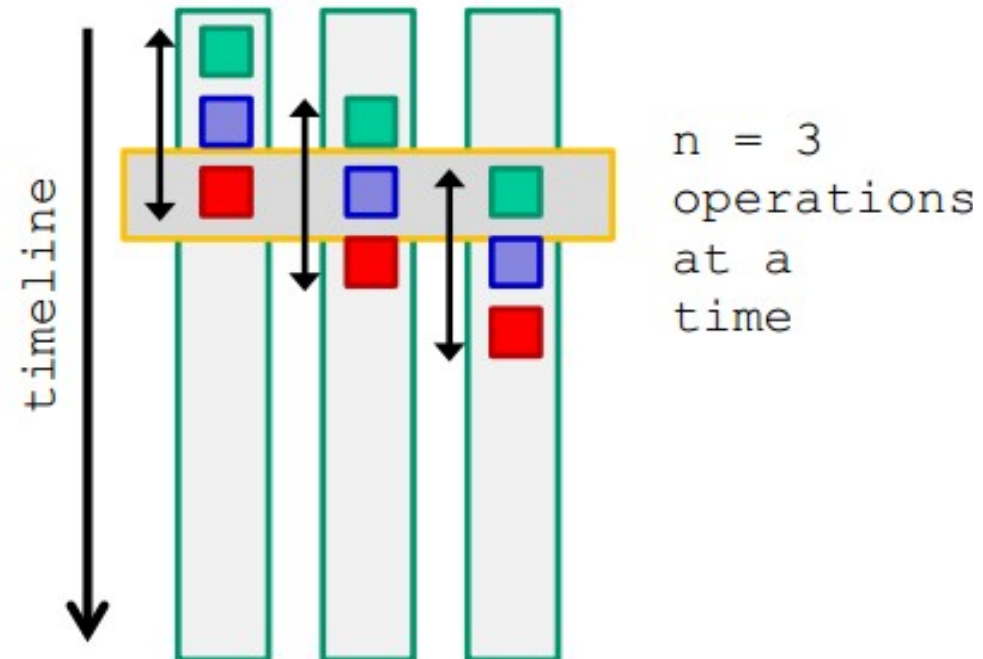


# Pipelining

## Sequential Loop:

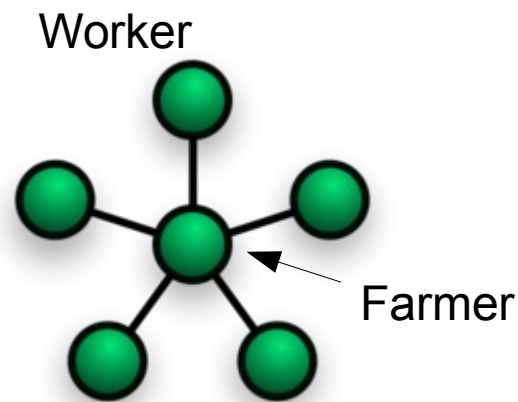


## Pipeline:



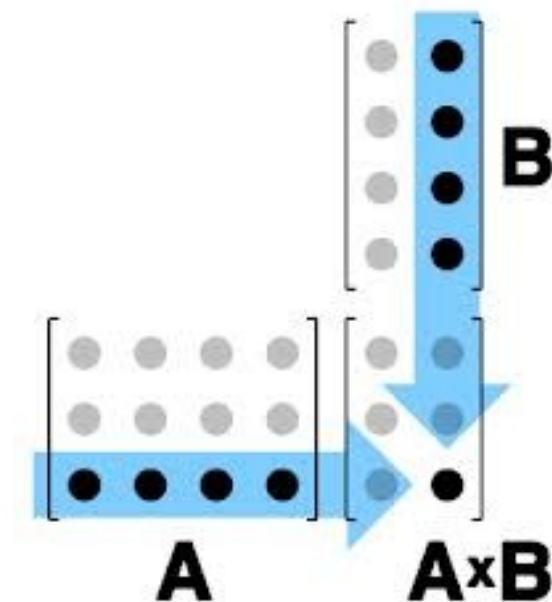

# Task Farms

- The 'farmer' distributes work
- The workers do the work, and when finished ask for the next item.
- The farmer does a lot of communication – can become communication bound



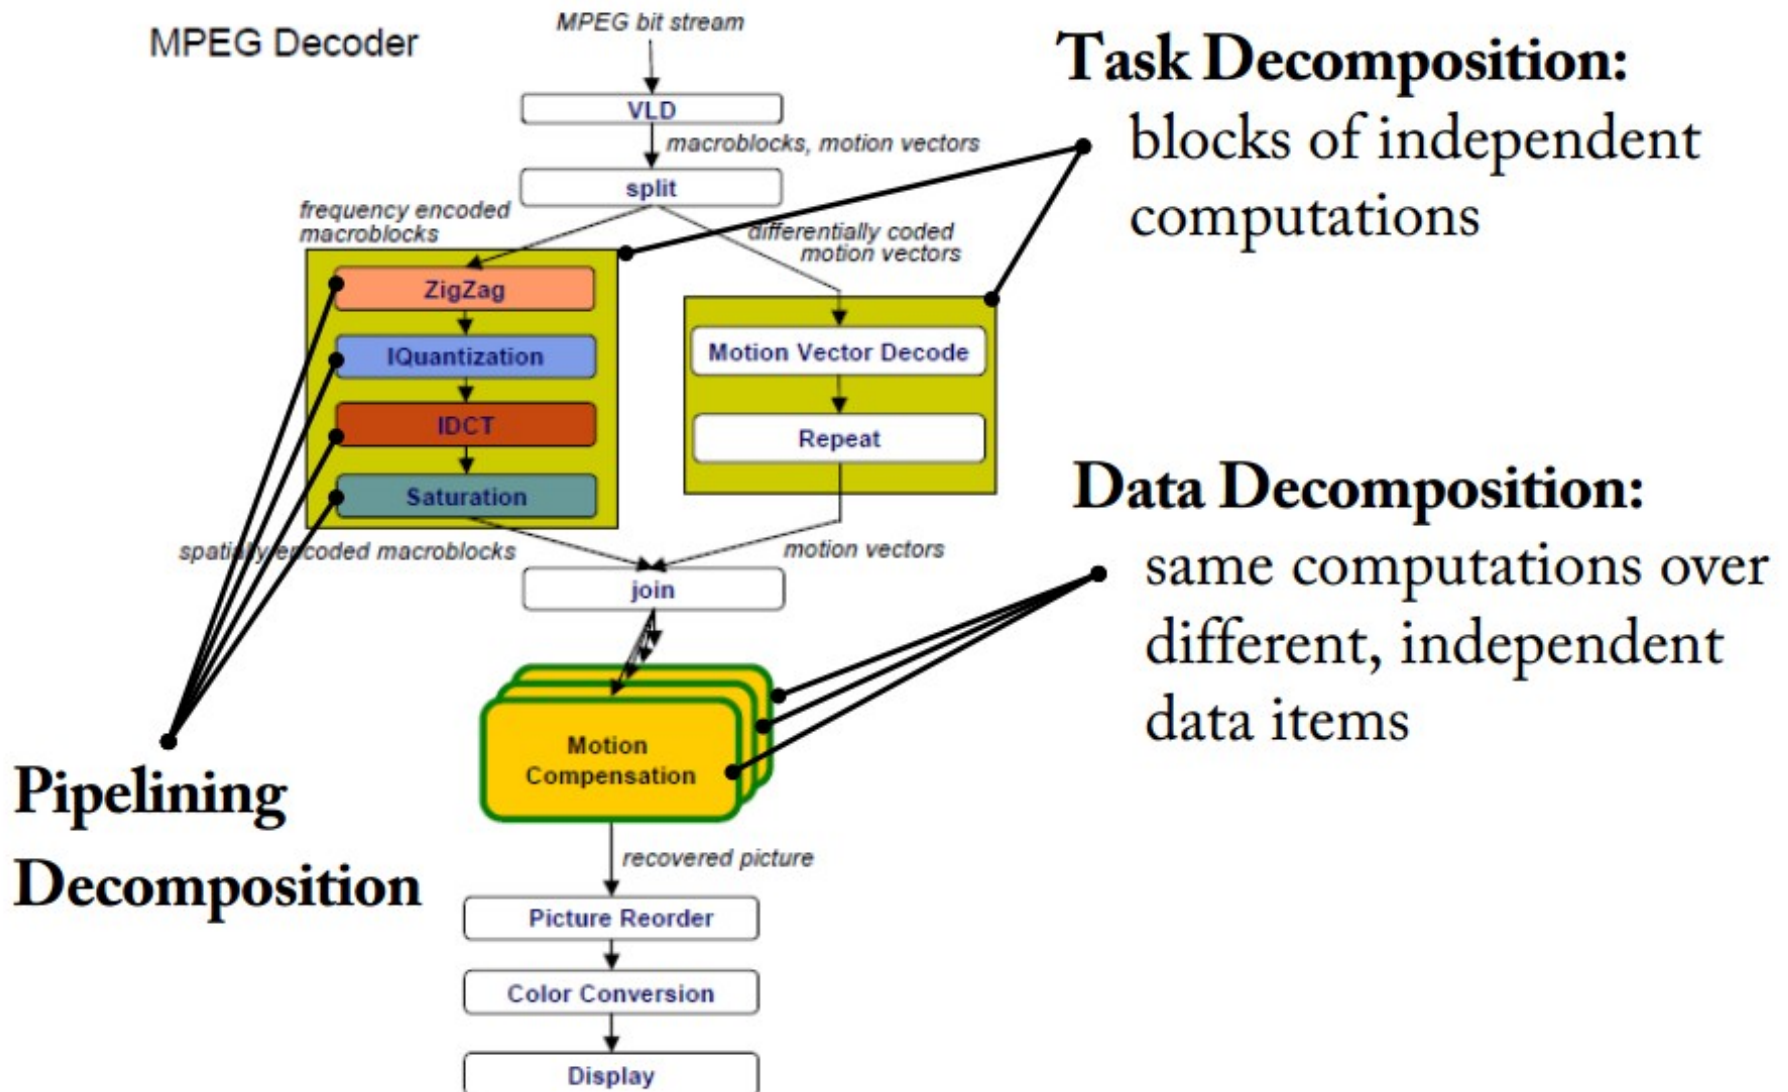
# Geometric Parallelism

- Often based on the layout of the task
- Matrix multiply
  - Grid-based
- Physical simulation
  - Often grid-based







# An Example




# More Examples




SIMD Parallelism



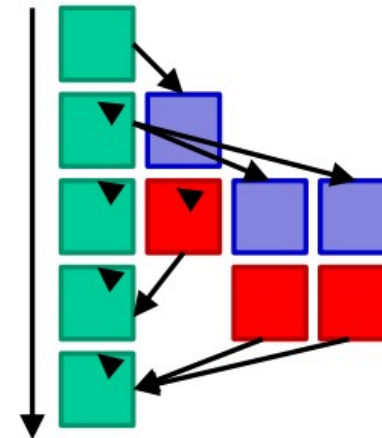
Pipelining (Chain Process)



Replication




Hierarchical Fork




Farming in Star Topology

# Deadlock and Race Conditions




# The Dining Philosophers




# Conditions for deadlock

- Mutual exclusion
- Resource holding
- No pre-emption
- Circular waiting


# The Dining Philosophers



Epicurus




Plato



Paul of  
Tarsus



Descartes



Aristotle



# Race Conditions



# Race Conditions

Thread 1

```
printf("hello ");  
printf("world ");
```

Thread 2

```
printf("good ");  
printf("bye ");
```

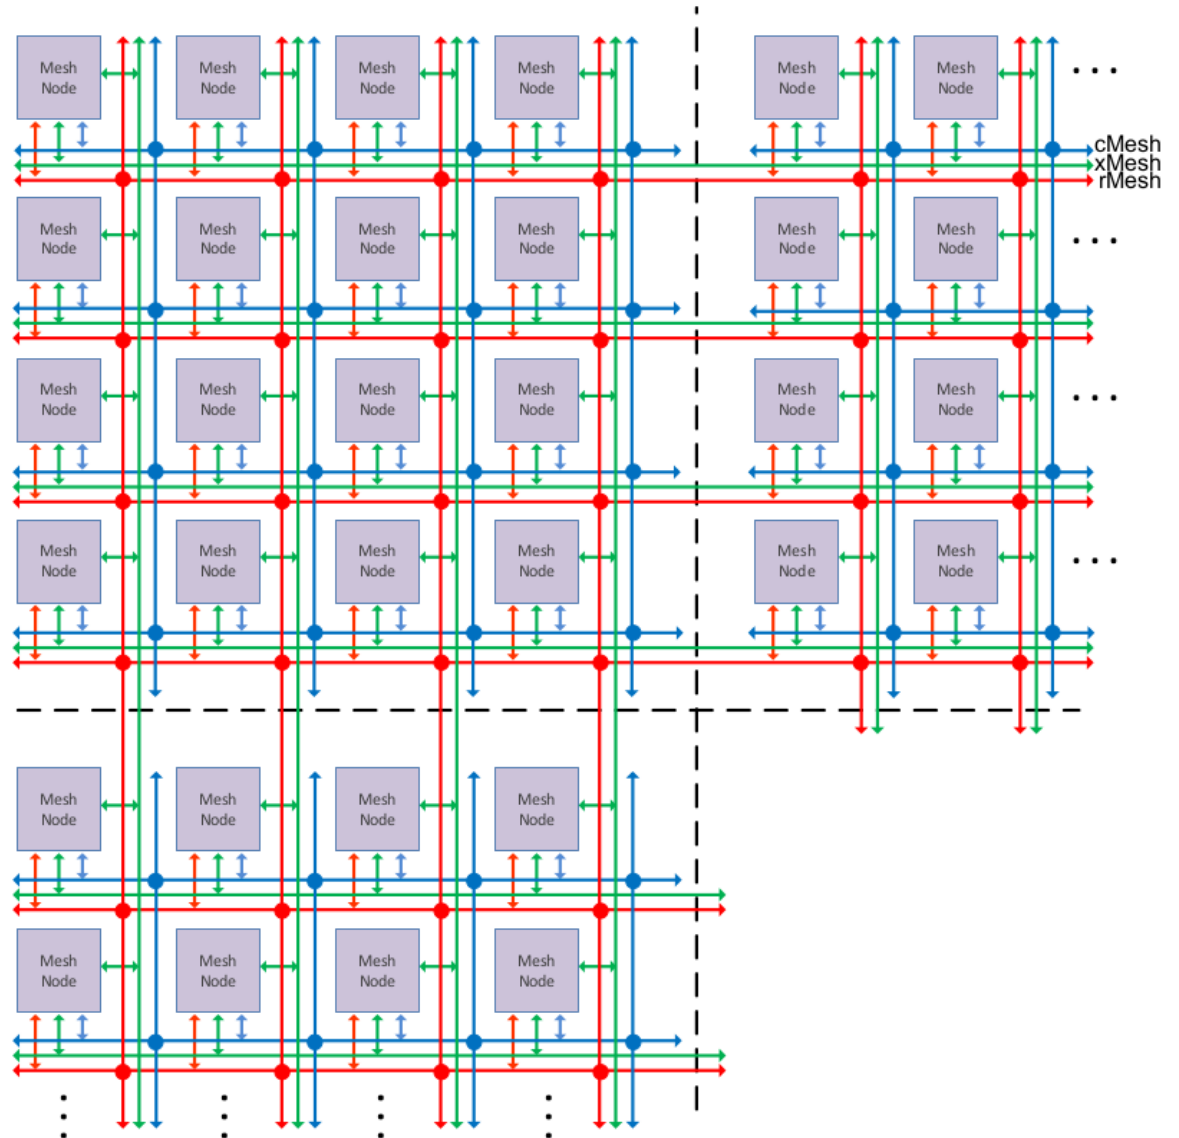
hello world good bye  
hello good bye world  
hello good world bye

good bye hello world  
good hello bye world  
good hello world bye



# Interconnect

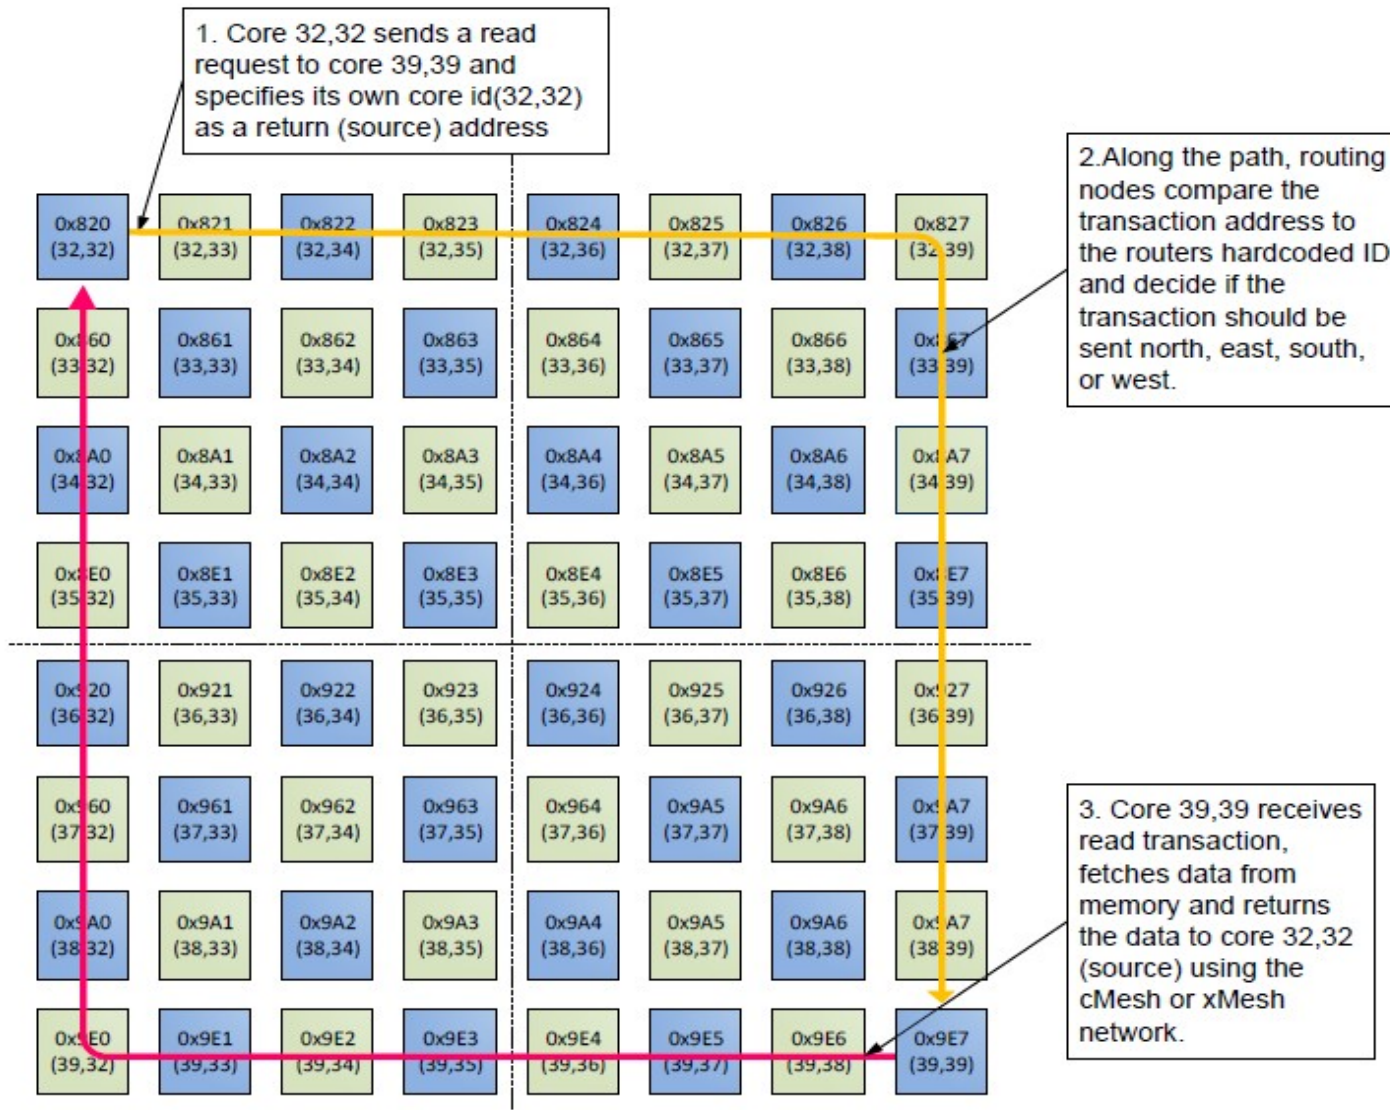
The eMesh



# The eMesh

- 3 meshes
  - Read, write, off-chip
  - Fast write mesh
- Read/write to any core
- Everything handled transparently
  - But it is useful to know some of the details for performance reasons


# Routing



# Mesh Determinism

- Relaxed consistency
- Be careful when writing to other cores!
- Order of writing to 2 different cores is not deterministic – race conditions and deadlock
- **Writing to a remote core, then attempting to read back the value is not deterministic**

# Safe Example




send to core 1

receive from core 0  
print "hello"  
send to core 2

receive from core 1  
print "world"

hello world

# Unsafe example



send to core 1  
send to core 2


receive from core 0  
print "hello"

receive from core 1  
print "world"

hello world

Lucky

# Unsafe example



send to core 1  
send to core 2

receive from core 0  
print "hello"


receive from core 1  
print "world"

world hello

Unlucky

# Make it safe

- How do we make it safe?
  - Synchronisation



send to core 1  
receive from core 1  
send to core 2

receive from core 0  
print "hello"  
send to core 0

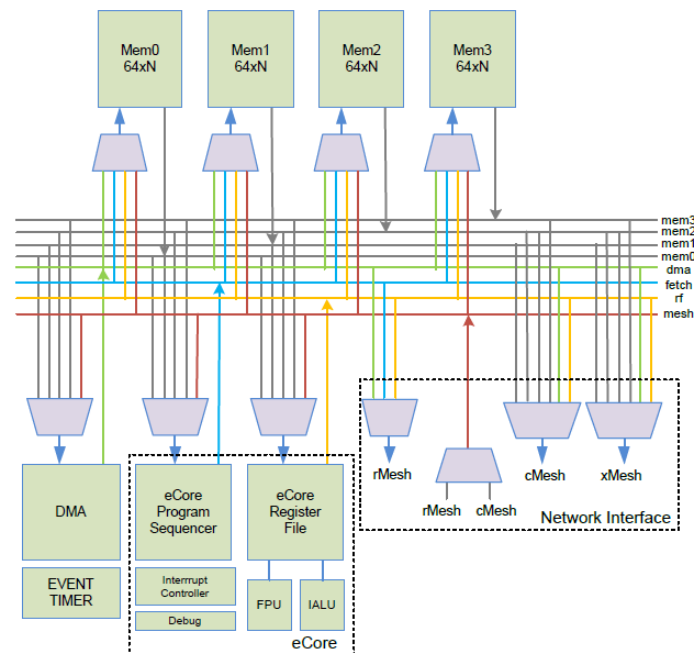
receive from core 1  
print "world"

hello world



# Make it fast

- Write mesh is non-blocking
- Try to structure programs and algorithms to only write to *non-local* memory
- Read and write from *local* memory is fast



# A simple example - slow

Running on core 0

```
int a[10], b[10];      // On core 0
int c;                 // On core 1

for(i = 0, c = 0; i < 10; ++i)
    c += a[i] * b[i];
```

Note – without synchronisation this program is also non-deterministic

# A simple example - fast

Running on core 0

```
int a[10], b[10];           // On core 0
int c_temp;              // On core 0
int c;                       // On core 1

for(i = 0, c_temp = 0; i < 10; ++i)
    c_temp += a[i] * b[i];

c = c_temp;
```

Fortunately this program is deterministic

# A tricky example

Running on core 0

```
int a[10], b[10];    // On core 0
int c[10];          // On core 1

for(i = 0; i < n; ++i)
    c[i] += a[i] * b[i];
```

# Another Solution?

## Core 0

```
int a[10], b[10];
int *temp;

temp = &partial;

for(i = 0; i < 10; ++i)
{
    temp[i] = a[i] * b[i];
    sync();
}
```

## Core 1

```
int c[10];
int partial[10];

for(i = 0; i < 10; ++i)
{
    sync();
    c[i] += partial[i];
}
```

Is it faster?

# Conclusion/Tips

- Many different paradigms for parallelism
- Careful of resource sharing
- Race conditions
  - Unique challenges with accessing remote cores
- Writing is faster
  - But it's challenging to pick the best way of doing this

# More Info

Epiphany architecture reference manual

<http://www.adapteva.com>

<http://www.parallella.org>

<http://www.embecosm.com>

Questions?