

Using the Cloud

- the Newcastle experience

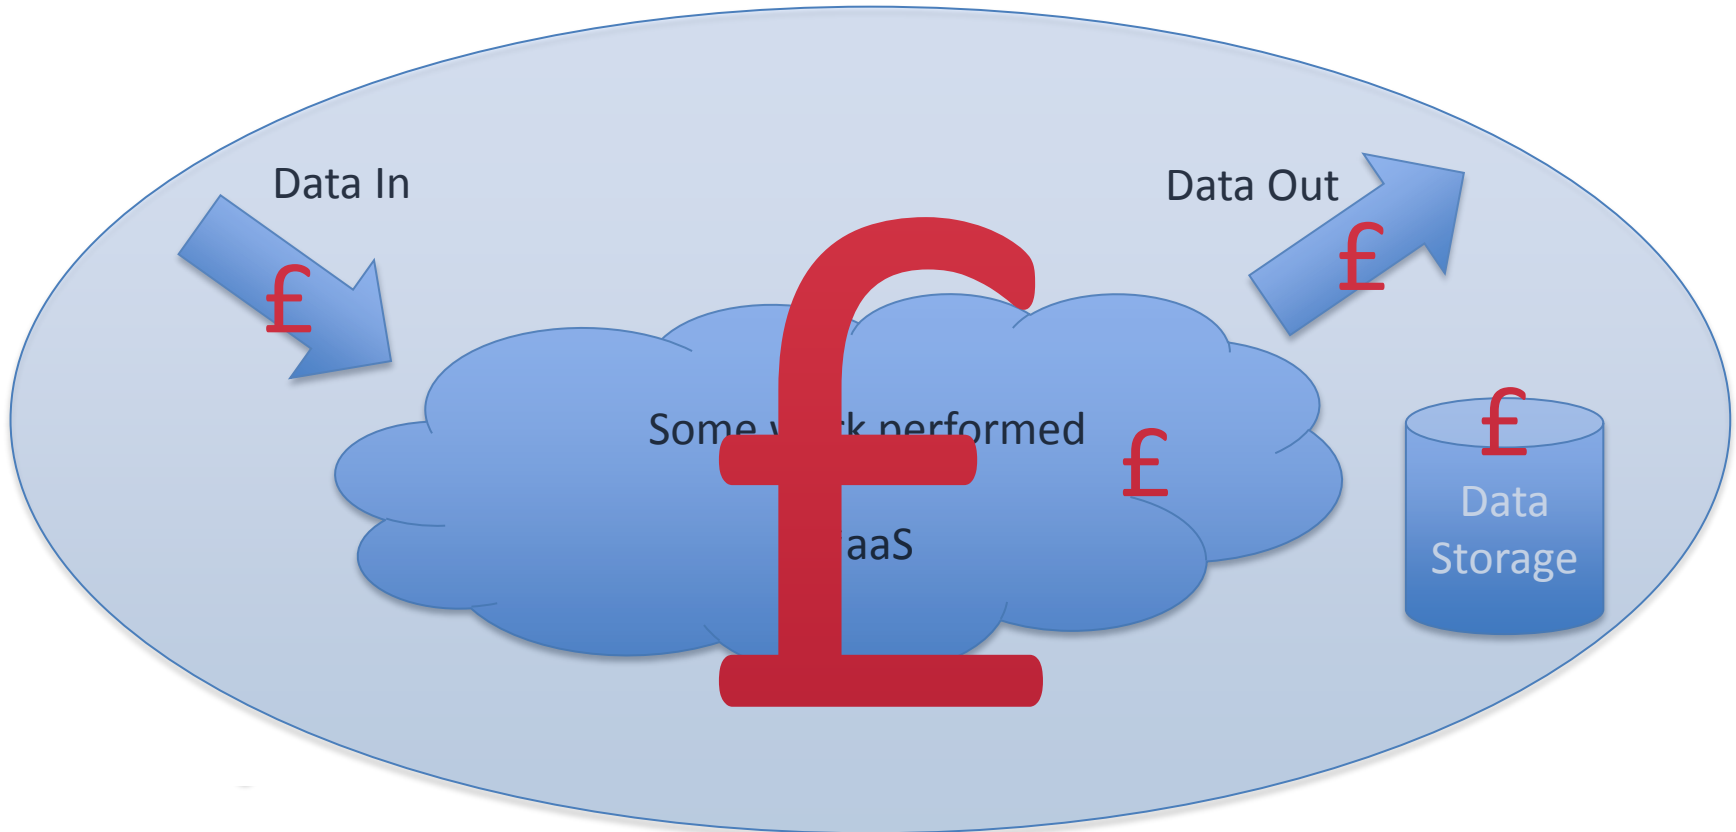
Stephen McGough
Vassilis Glenis, Chris Kilsby,
Vedrana Kutija, Simon Wodman
University of Sussex : HPC Workshop 2013
January 16th 2013



Digital Institute

CLOUD COMPUTING

The Cloud model



SaaS – Software as a Service, you use programs provided by someone else

PaaS – Platform as a Service, you provide the program and above


IaaS – Infrastructure as a Service, you provide the OS and above

Clouds are built on Data Centers

(acknowledgements to Dennis Gannon, Microsoft)

- Economies of scale
 - Approximate costs for a small size center (1000 servers) and a larger, 100K server center.

Technology	Cost in small-sized Data Center	Cost in Large Data Center	Ratio
Network	\$95 per Mbps/month	\$13 per Mbps/month	7.1
Storage	\$2.20 per GB/month	\$0.40 per GB/month	5.7
Administration	~140 servers/Administrator	>1000 Servers/Administrator	7.1



Each data center is
11.5 times
the size of a football field

Advances in DC deployment

- Conquering complexity.
 - Building racks of servers & complex cooling systems all separately is not efficient.
 - Package and deploy into bigger units:



(acknowledgements to Dennis Gannon, Microsoft)

Microsoft

Containers: Separating Concerns



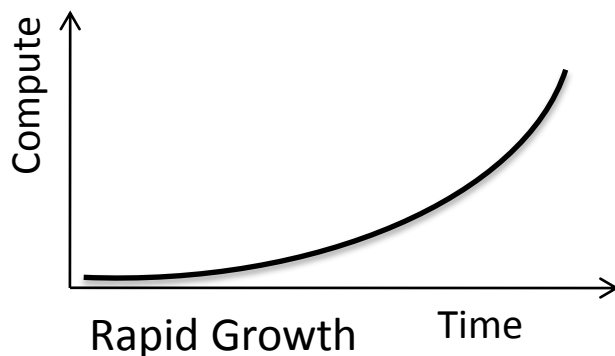
(acknowledgements to Dennis Gannon, Microsoft)



What is the Cloud Good for?

Digital Institute

- Highly democratic access to resources
 - Anyone with a credit card can use
- Rapid provisioning - Minutes as opposed to days / weeks / months
- Only pay for the time you use
 - Operational cost rather than capital cost
- Very good for small players
 - Can access more resources than they could buy
 - No worry for scale-up : Fast Growth
- Short-term goals
 - Turn-around time more important than long-term goals
 - Bursty usage



- General problems
 - Data (program) security / legality
 - Can I send my data (program) there?
 - Interaction with program
 - GUI interfaces
 - Interactive use
 - Data volumes
 - How long to transfer data?
 - data to program vs program to data
- Cloud specific problems
 - Data volumes
 - Cost of transferring
 - Execution time
 - Execution time (cost) vs ownership cost




Digital Institute

THE PROBLEM

Flood Modelling

- Landscape topology
- Add buildings
- Simulate rain storms
- Flow rainfall over landscape
- Determine which areas are prone to flooding



- Available resources
 - Simulate a city suburb (small area low resolution) requires ~1.5GB memory
 - Really want to simulate an entire city (large area) and/or at higher resolution
 - Requiring much more memory (up to 48GB)
- Running multiple ensembles
 - Different rain storms
 - Once every 2/10/20/50/100/200 year storm event
 - Different durations
 - 15 minutes, 30 minutes, 1 hour, 2 hours, 3 hours and 6 hours
 - Use high throughput computing
- Run time
 - Simulations last 1 – 13 days

Problem Size

Problem	X Small	Small	Medium	Large	Thames
Simulation size	2km by 2km	2km by 2km	2km by 2km	10km by 12km	200km by 55km
Sim resolution	2m by 2m	1m by 1m	0.5m by 0.5m	4m by 4m	15m by 15m
Cells	1 million	4 million	16 million	7.5 million	5 million
Memory	1.5GB	6GB	40GB	24GB	24 GB
Disk Space	100GB	100GB	200GB	200GB	200GB
Runtime	5 – 24 hours	1 – 3 days	7 – 10 days	7 – 11 days	8 – 13 days



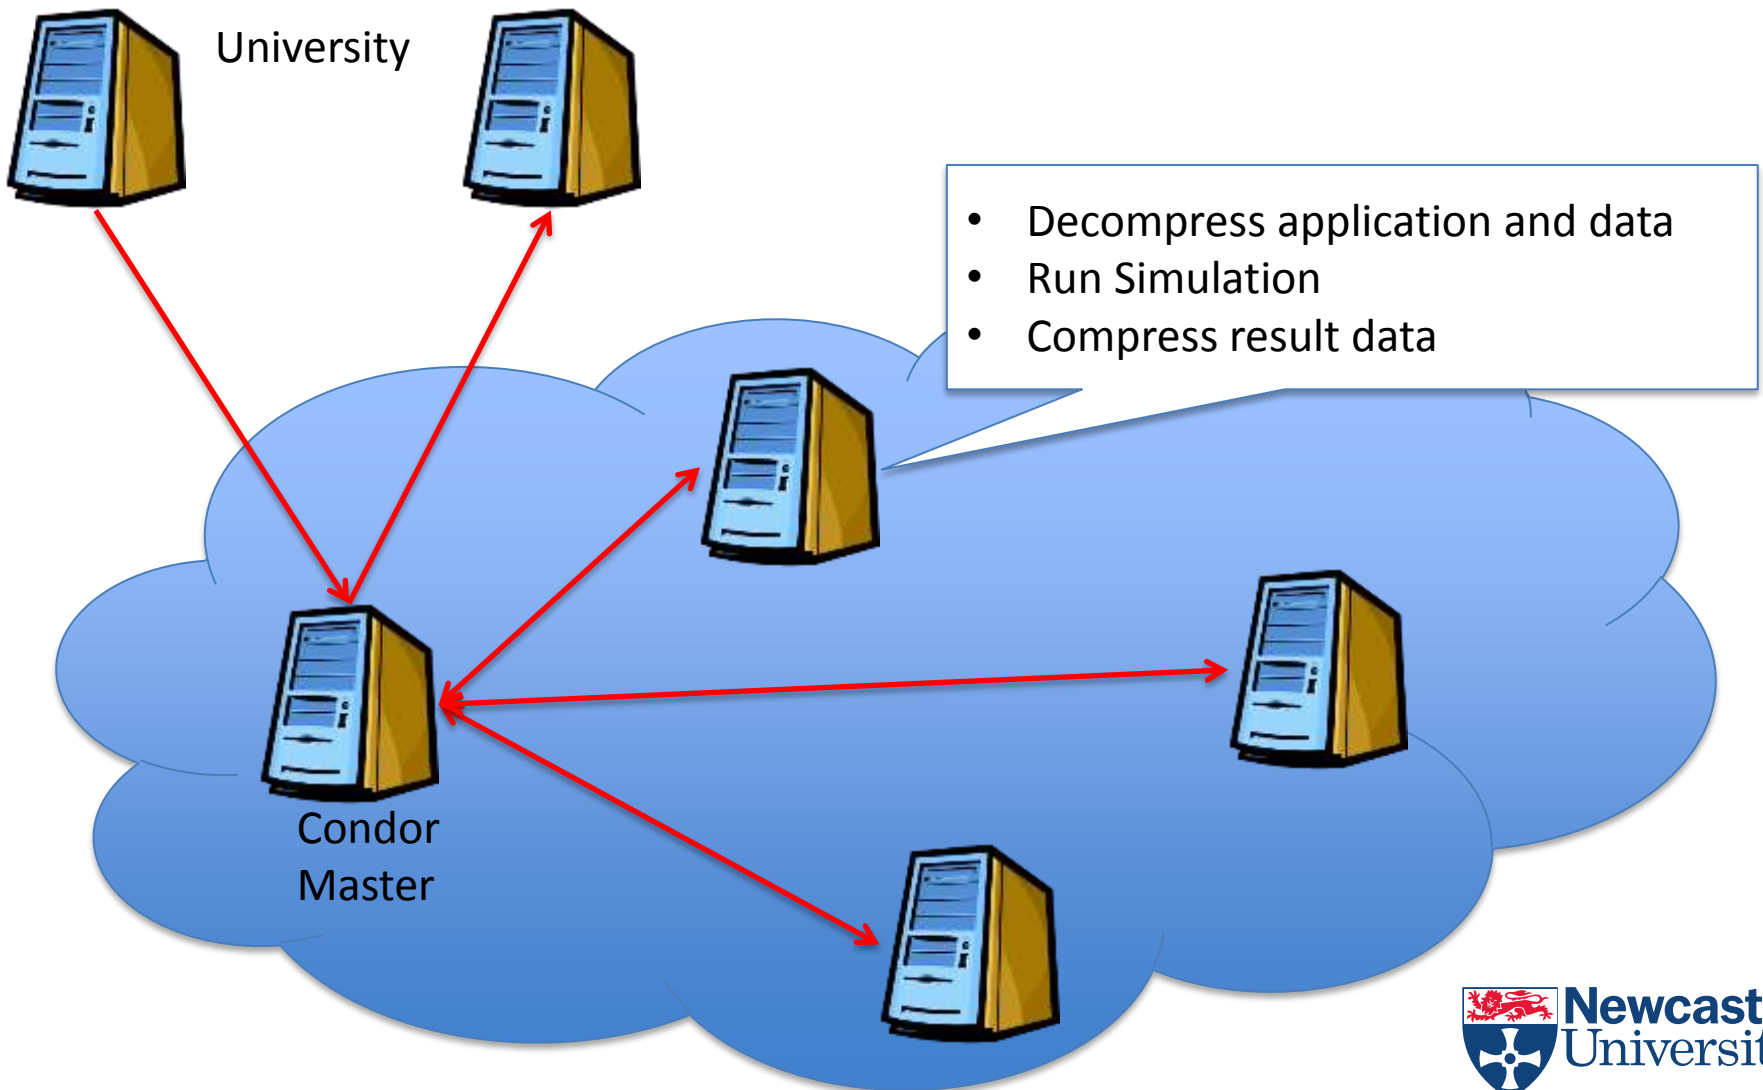
Digital Institute

THE PLAN

The Plan

- Allocated ~£20,000 to run work on the Cloud
- Using IaaS
- Deploy Linux Based Virtual Machines
- Run Condor over virtual Cluster
 - Cluster completely in the Cloud
 - Well understood system
- Instances can be added and removed as required
- Data Staged back to computer in the university on completion of work

Design





Digital Institute

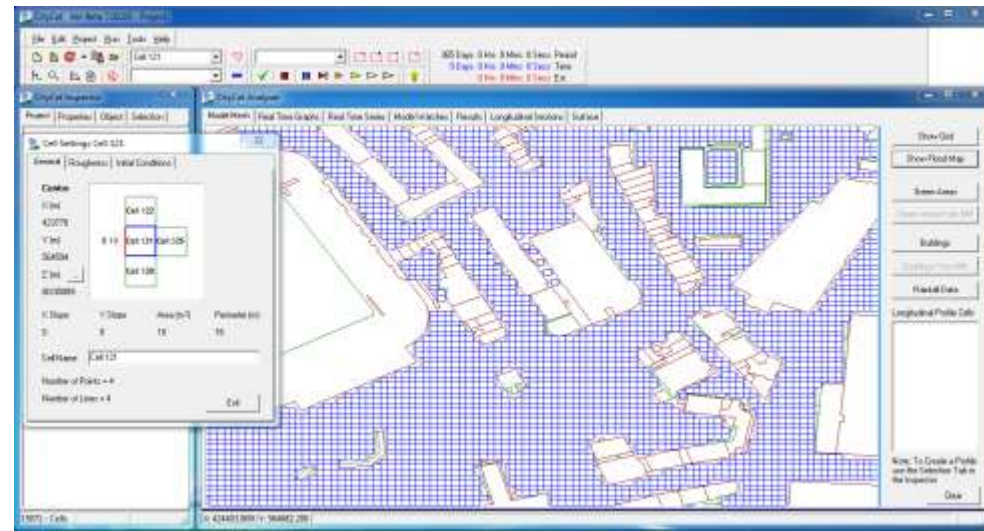
ISSUES



Technical Issues: Software

Digital Institute

- Software
 - Driven through GUI interface
 - Originally command line but evolved into GUI application
 - Written in Delphi
 - Limits use to Windows based Instances
 - Higher cost for using Windows
- Solution:
 - Removal of GUI interface code – large effort
 - Recompiled with open source Delphi compiler for Linux (Lazarus, Linux free Pascal compiler)



Legal Issues: Licensing

- Landscape topology data free to use under any circumstance
- Building data under license from external company
 - Licensed to use for non commercial use within the University
 - Cloud not in the University
- Solution:
 - Needed to draft up new collaboration agreement with company

- £20,000
 - Can't do this by stealth
- University decided that this should be spent in-house not go to external company
 - Can't we do this in-house?
 - No appropriate resources
 - Can't we buy the resources?
 - Not enough money, £20,000 would buy 2-3 servers
 - ~ 3 years work
 - European Tendering laws
 - Producing a price comparison report
 - Can't we do this in-house – we have resources?
 - Resources are running key university services
- No infrastructure in place to deal with Cloud

Pricing Comparison

- Two metrics defined
 - Cost per simulation hour
 - Cost per cloud hour / simulations per cloud instance
 - Hours in one month
 - The number of hours in one month we could purchase
- Used to compare between cloud providers and purchased resources
- Look at two problem cases
 - Small jobs (3GB), large jobs (40GB)

- Server bought November 2011
 - 12 Core, 128GB RAM £3,182 ~ \$5,142
 - Could buy 6 of these

Table 4 Cost-time analysis for locally provisioned resources

Simulation Memory	Cost per simulation hour	Max hours
3GB	\$0.576	53,568
40GB	\$2.304	13,392

Cloud : Azure

- No 40GB instances available
- All resources are priced the same

Table 7 Cost-time analysis for 3GB simulations on Azure

Name	Cores	Memory	Concurrent	Unit Cost	Hours	Cost per simulation hour
Medium	2	3.5	1	\$0.24	134,583	\$0.24
Large	4	7	2	\$0.48	134,583	\$0.24
Extra Large	8	14	4	\$0.96	134,583	\$0.24

- No 40GB instances available
- Prices the same for all

Table 8 Cost-time analysis for 3GB simulations on GoGrid

Name	Cores	Memory	Conc- urrent	Unit Cost	Hours	Cost per simulation hour
Server 4/4/200	4	4	1	\$0.76	42500	\$0.76
Server 8/8/400	8	8	2	\$1.52	42500	\$0.76
Server 16/16/800	16	16	5	\$3.04	53125	\$0.608

- Only one 3GB option available

Table 9 Cost-time analysis for 3GB simulations on RackSpace

Name	Cores	Memory	Concurrent	Unit Cost	Hours	Cost per simulation hour
4096/160	1	4	1	\$0.252	128,174	\$0.252

Table 5 Cost-time analysis for 3GB simulations on EC2

Name	Cores	Memory	Conc- urrent	Unit Cost	Hours	Cost per simulation hour
Large	4	7.5	2	\$0.34	190,000	\$0.17
Extra Large	8	15	4	\$0.68	190,000	\$0.17
High Memory XL	6.5	17.1	5	\$0.50	323,000	\$0.10
Double XL	13	34.2	11	\$1.00	355,300	\$0.091
Quad XL	26	68.4	22	\$2.00	355,300	\$0.091
High CPU XL	20	7	2	\$0.68	95,000	\$0.34
Cluster Compute	33.5	23	7	\$1.60	141,312	\$0.229



Table 6 Cost-time analysis for 40GB simulations on EC2

Name	Cores	Memory	Conc- urrent	Unit Cost	Hours	Cost per simulation hour
Quad XL	26	68.4	1	\$2.00	16,150	\$2.00



Cost Summary

- Can't directly compare local with Cloud on cost as local resources are available after project
- But number of hours available in one month is far greater on the cloud
- EC2 offered the best price
- Not taken into account
 - The performance of instances
 - Differences between providers
 - Differences between instances from the same provider

Payment Issue

- Most cloud providers 'insist' on payment by credit card
- We wanted to do £20,000 worth of work in one month
 - Can't get a university credit card with this sort of limit
 - Special arrangements made to allow head of finance to have this limit
 - Doesn't work as he wants to pay for an item not have an open card payment
 - Raised existing card limit to £20,000 and split work over payment months
- Discovered too late that Cloud providers can take payment by invoice (not advertised)

- Upload files:
 - 3 – 25MB depending on simulation run
- Download files:
 - 1.4 – 9GB
- With compression we can bring this down to
 - 125 – 1024MB
- Uploads free downloads not
- Data transfer speeds
 - Dependent on time of day
 - Varied from 200-600KB/s (worst when US awake)
 - Keeping instances alive

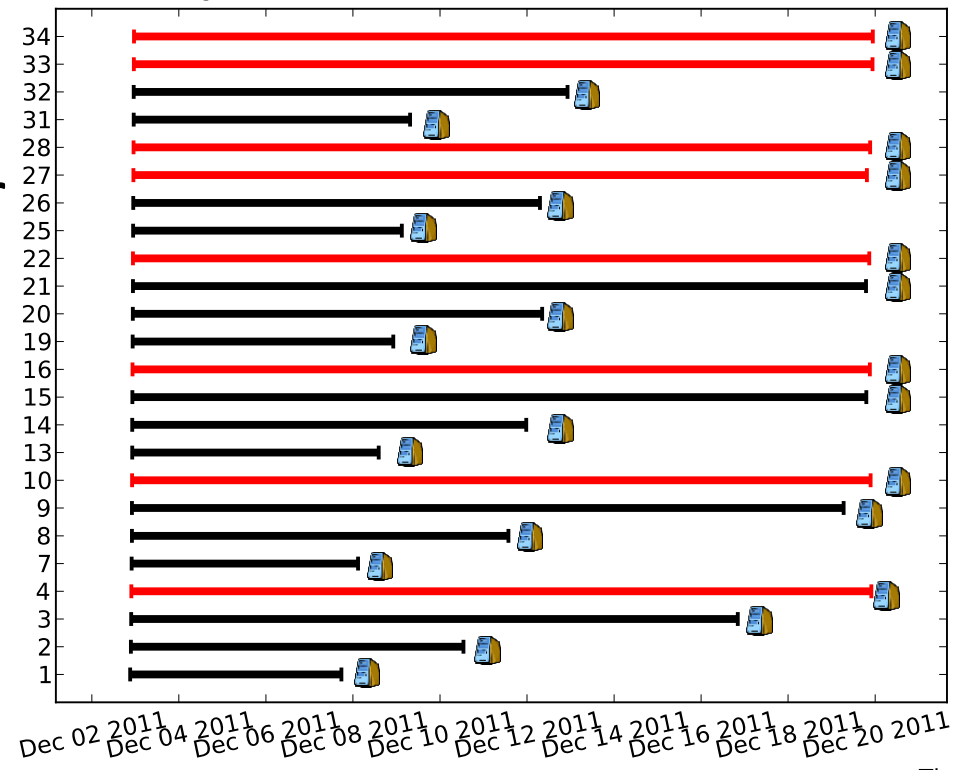


Digital Institute

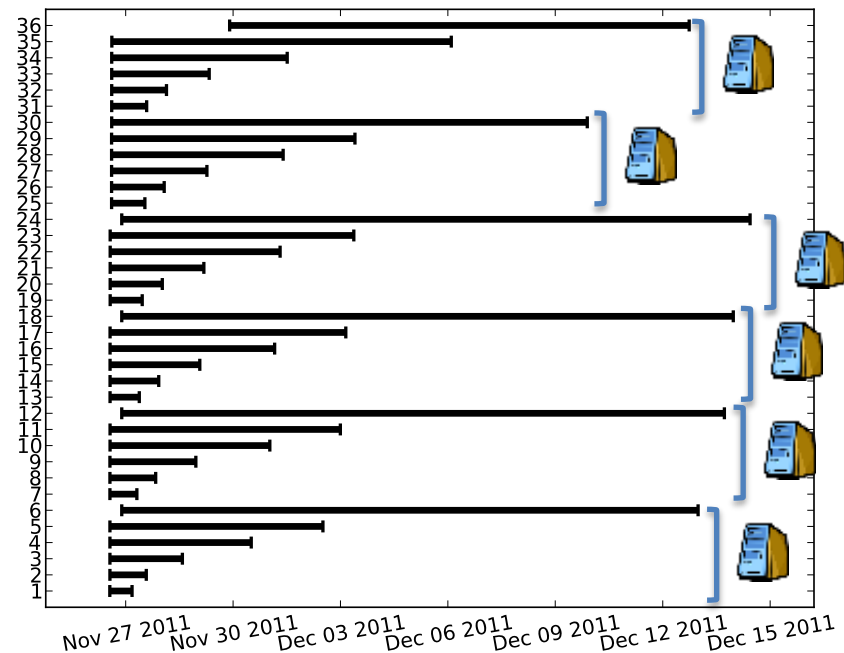
SOME PRELIMINARY RESULTS

Simulation Results

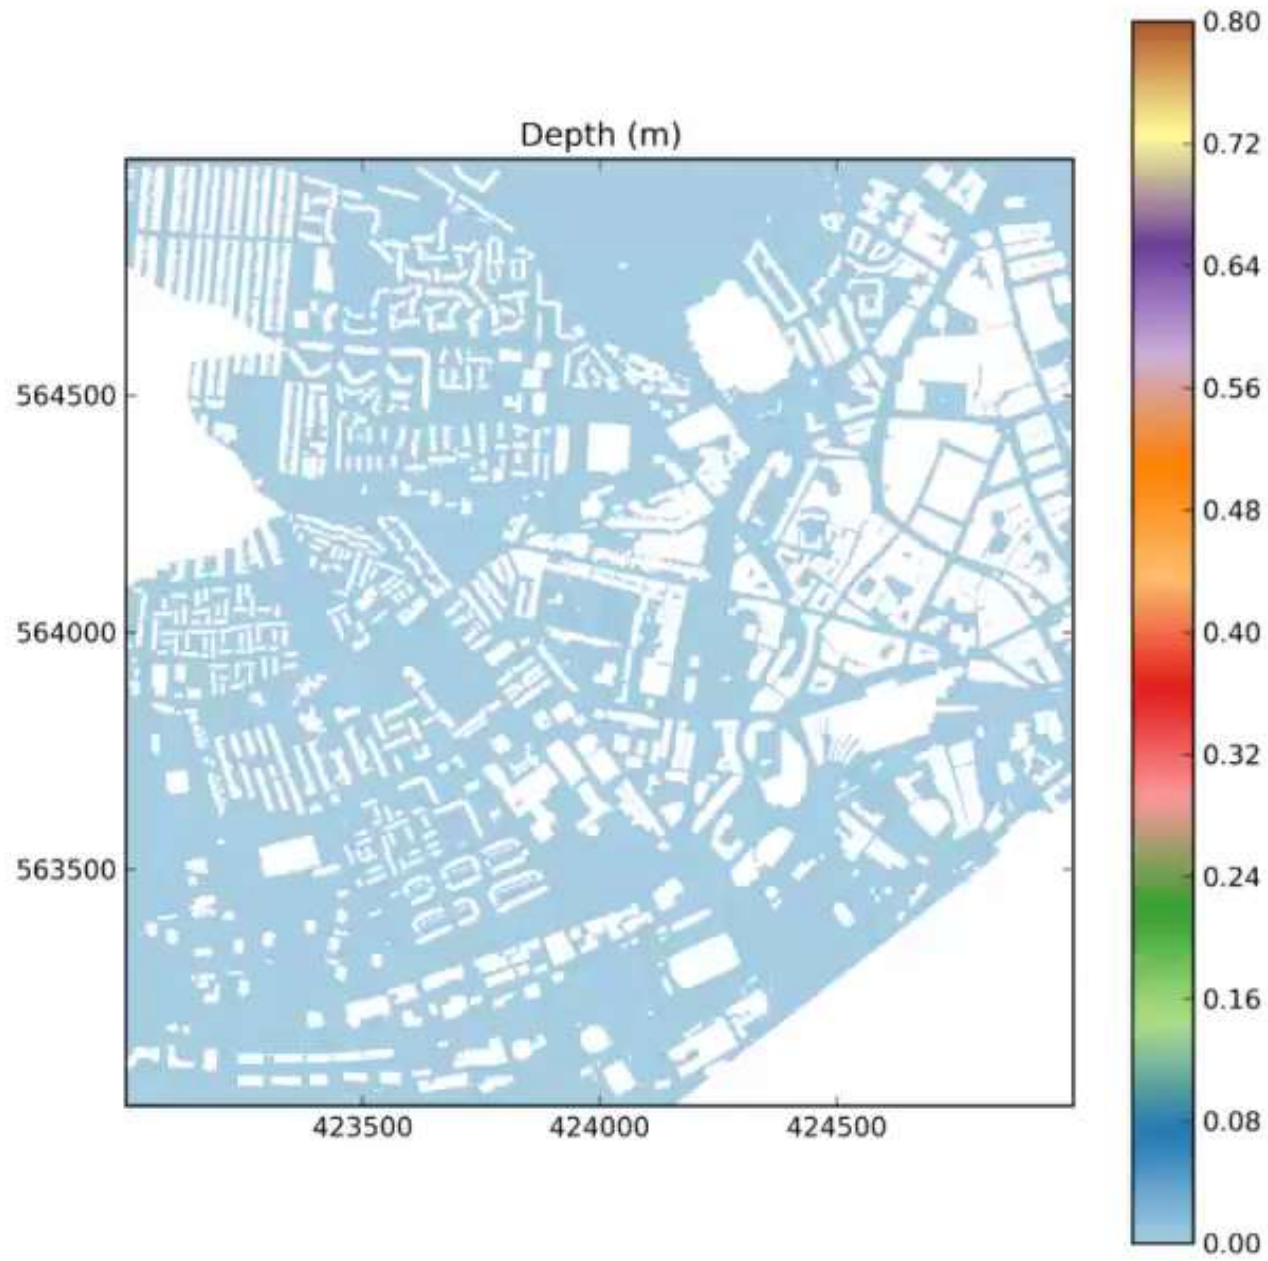
- Very large simulations – 1 simulation per instance
 - Instances shut down on completion
- Actual data costs
 - Compute \$2.02 / hour
 - Inc wasted time
 - Data \$2.16 for 18GB




- 11GB simulations
 - 6 jobs per extra-large instance
- Bad job alignment – jobs in order
 - Longest jobs keep boxes active
 - Effective cost \$1.03
 - If we grouped jobs more intelligently
 - Bring cost down to ~\$0.395
 - Close to \$0.333 optimal
 - Data transfer 10GB
 - \$1.08











Conclusions

- Technical issues on using the Cloud are relatively easy to solve
- Political issues are more difficult to solve
 - Breaks the model of how University does work
 - Lots of people to convince this is a valid way to do things
- Hopefully next time will be easier!
- Need more (Cloud) time to process results



Digital Institute

stephen.mcgough@ncl.ac.uk

Thank you

QUESTIONS