ExAMPLER

An ExCALIBUR Project

WiFi
FM Events
#1739FDM

Making High-End Computing Easier for Social Scientists

18-20 March 2025

The Foundling Museum, London

Sponsor: ExCALIBUR-ECP Exchange Programme

Urban Big Data Centre

casa

The James Hutton Institute

Agenda (Tuesday PM)

- 14:00 Welcome and introduction to ExAMPLER
- 14:30 Agent-based modelling on GPUs with FLAME GPU
- 15:30 Coffee
- 16:00 Future directions for NetLogo:
  Lower threshold, higher ceiling
- 17:00 General discussion
- 17:30 Close
- 18:30 Dinner:
  Antalya Restaurant, 103-105 Southampton Row

ExAMPLER

An ExCALIBUR Project

Introduction to ExAMPLER

Gary Polhill, with

Alison Heppenstall, Mike Batty, Matt Hare, Ricardo Colasanti, Richart Milton, Doug Salt

Urban Big Data Centre

casa

The James Hutton Institute

Briefly…

- ExAMPLER is an 18 month project funded by the EPSRC under the ExCALIBUR programme
  - How to get ABM to exascale computing?
- Systematic Literature Review
  - State-of-the-Art in HPC use by ABM community
- Visions workshops
  - What could people do with exascale ABM?
- Roadmaps workshops
  - How do we get there?
- Engage with other ExCALIBUR projects

And…?

- 18 months became 22 months…
- SLR was a pain in the neck
  - Scopus and Web of Science are ‘unclean’
  - Read 220 papers: hardly any used HPC
    - Hardly any even said what they did use…
  - Not even sure ‘agent-based model’ means anything!

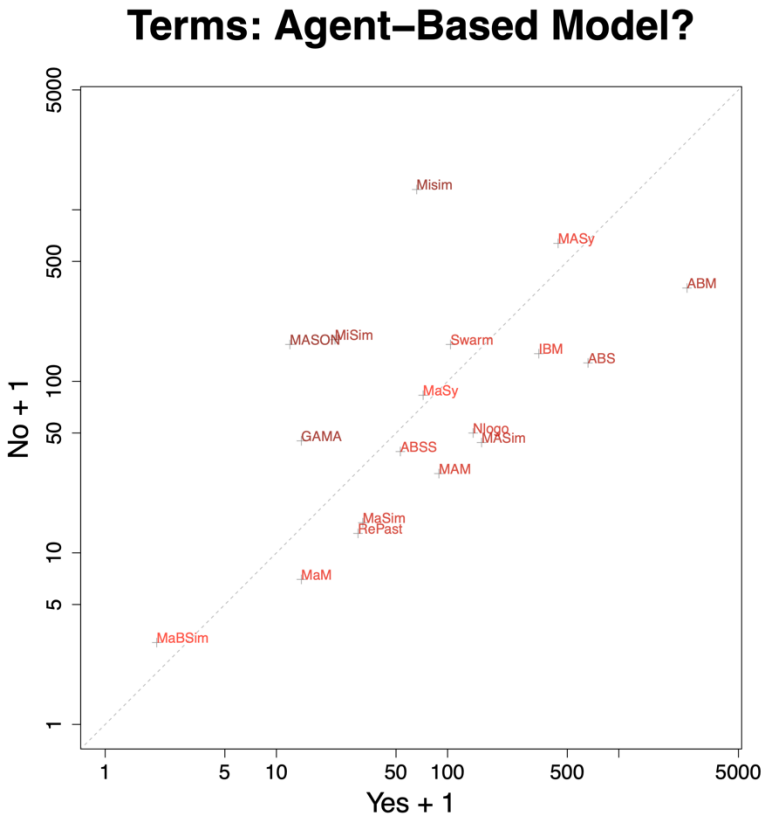

Visions workshops; Glasgow & London ’23

- Explored potential of exascale for ABM
- Participatory approach
  - Used group model building to co-create causal loop models
- Glasgow (Social Simulation Conference)
  - 2h (1h introduction; 1h group modelling)
- London
  - 1.5 days (expert presentations; 1.5h group modelling; refinement discussions)

Combined Causal Loop Model

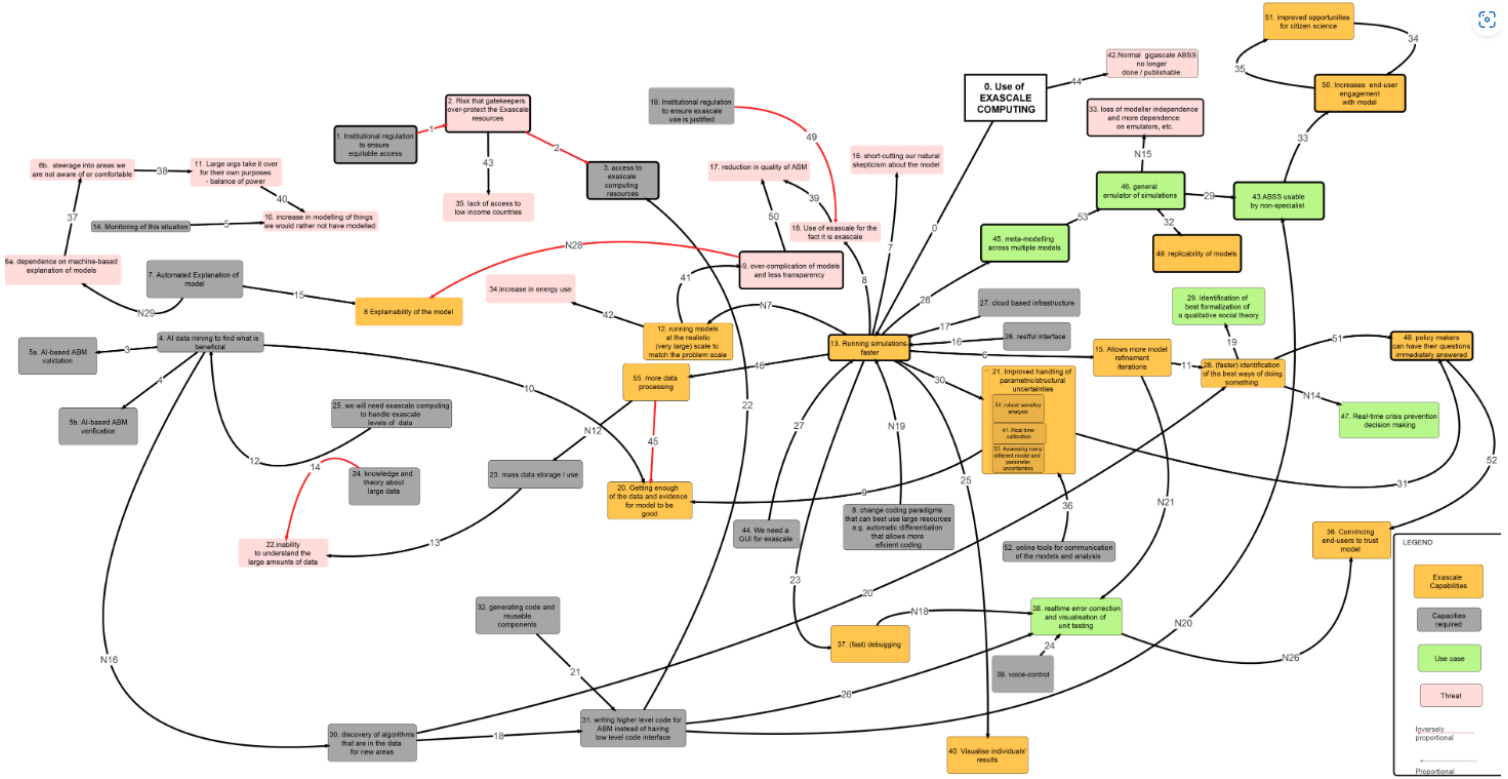

Exascale Capabilities

Capacities Required

Use Cases

Threats

Inversely Proportional

Proportional

Visions model: Use Cases

- Individual-based social science as a new field
- Identifying best formalizations for social theories
- Real-time decision-making support
- Meta-modelling and general simulation emulators
- Real-time visualization and debugging tools

Use Cases

Visions model: Threats

- Increased energy consumption and cost
- Inefficient use of exascale resources
- Loss of research community control over agent-based modelling
- Overly complex models and gatekeeper control
- Over-reliance on AI for development and interpretation

Threats

Visions workshops: Challenges

- High cost of hardware might limit funding for research and development
- Parallelization of agent-based models not trivial
  - Risk that this is ignored, and we ‘oversimplify’ ABMs to enable full benefits of parallelism
    - Synchronous behaviour, matrix/network-based formalisms, simplified interactions, simplified decision-making
  - Introduce artefacts ‘for the sake of simplicity’
    - Exactly what ABM was supposed to avoid!

Roadmap: steps to exascale

- Exascale ABM is sensitive to interaction mode
  - In some cases, within-run parallelism is too constrained for exascale compute to be beneficial
- We need to work up to exascale ABM
  - Laptop (gigascale) -> HPC (terascale) -> petascale
- We need a competitive science case
  - …that beats simulating galaxies, metabolisms and protein molecules
- We need to scale up using proxy apps
  - Use smaller number of agents, fewer time steps, smaller spatial extents
  - But provide meaningful indicators of performance
    - On GPU and CPU and/or hybrid computing environments…
    - Can these be 'reusable' from one model to another…?
- Training metamodels of ABMs could make their results more immediately accessible

How could we…?

Access GPU
programming skills

How could we…?

Access GPU programming skills

Buy non-standard IT equipment

How could we…?

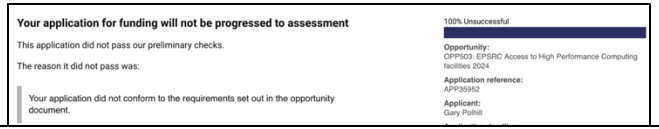
Your application for funding will not be progressed to assessment

This application did not pass our preliminary checks.

The reason it did not pass was:

> Your application did not conform to the requirements set out in the opportunity document.

100% Unsuccessful

Opportunity:
OPP503: EPSRC Access to High Performance Computing facilities 2024

Application reference:
APP35952

Applicant:
Gary Polhill

> Your application did not conform to the requirements set out in the opportunity document.

Access GPU
ogramming skills

Buy non-standard
IT equipment

Make a successful
HPC request

How could we…?

Access GPU
ogramming skills

Buy non-standard
IT equipment

Make a successful
HPC request

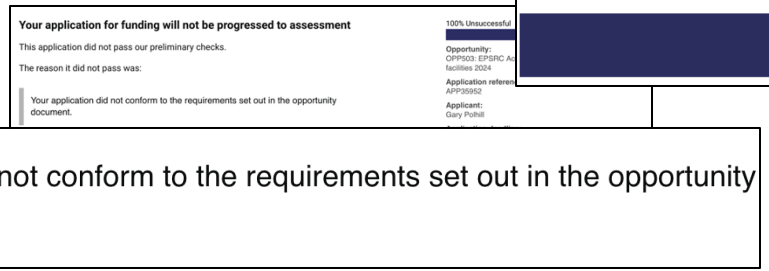

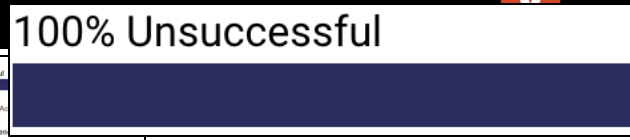

> Your application did not conform to the requirements set out in the opportunity document.

How could we…?

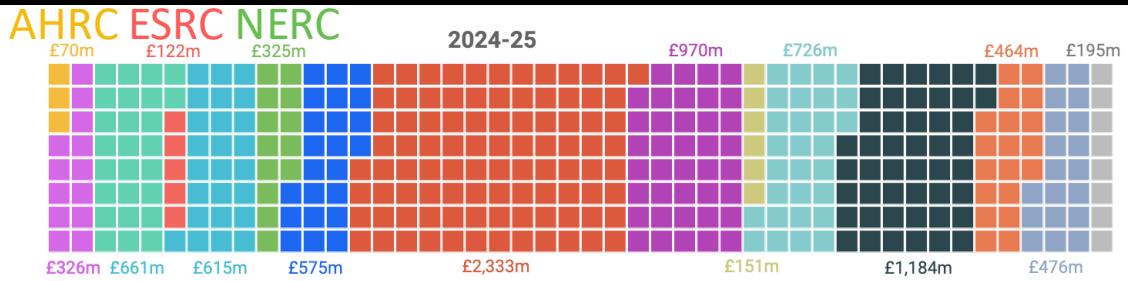

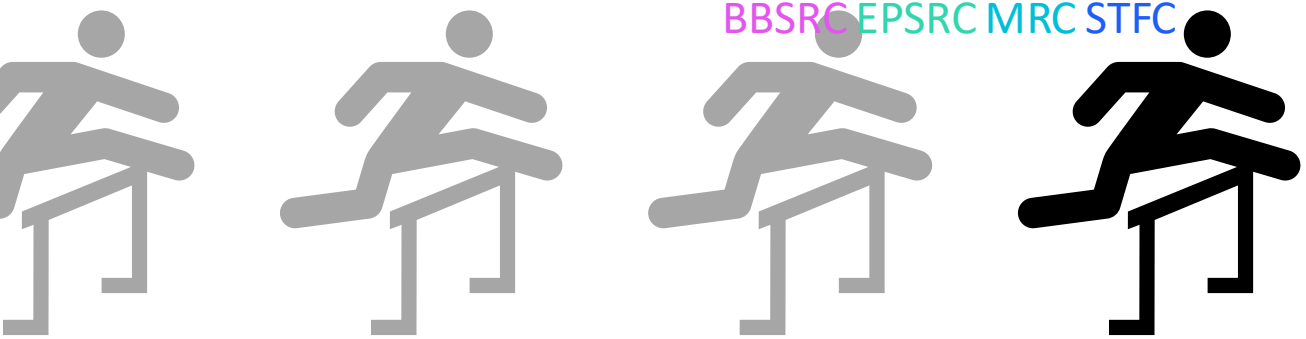

cess GPU
mming skills

Buy non-standard IT equipment

Make a successful HPC request

Justify investment and opportunity costs

Central Funding
Existing Strategic Prog.
New Strategic Prog.
Infrastructure
Collective Talent
Existing Commitments
Core IUK
Core RE

How could we…?

s GPU
ning skills

Buy non-standard IT equipment

Make a successful HPC request

Justify investment and opportunity costs

Design, build and maintain reusable software

How could we…?

Why do you want to model the whole of the UK?

The women's session is in the other room

skills

Buy non-standard IT equipment

Make a successful Tier-2 HPC request

Justify investment and opportunity costs

Design, build and maintain reusable software

Overcome imposter syndrome & discrimination

Gains without exascale…

The Economist

Try for free | Log in

Weekly edition | The world in brief | War in the Middle East

Britain | Exaflop

Britain's government pulls the plug on a superfast computer

The decision is likely to hurt researchers and AI companies

Exaflop PHOTOGRAPH: PAUL ZANRE

Aug 22nd 2024

- Better exploration of behavioural dynamics and parameter spaces
- Better understanding of relationship between dynamics and computational demand
  - More insights into social complexity?
- More routine use of HPC by social scientists
  - More HPC environments configured to suit their needs and software
- GPUs use less energy per executed instruction

Agenda (Tuesday PM)

- 14:00 Welcome and introduction to ExAMPLER
- 14:30 Agent-based modelling on GPUs with FLAME GPU
- 15:30 Coffee
- 16:00 Future directions for NetLogo:
  Lower threshold, higher ceiling
- 17:00 General discussion
- 17:30 Close
- 18:30 Dinner:
  Antalya Restaurant, 103-105 Southampton Row

Agenda (Wednesday AM)

- 09:00 Arrival and coffee

- 09:30 Hands-on cloud-based FLAME GPU Tutorial
- 10:30 Running NetLogo on ARCHER-2
- 11:00 Coffee

- 11:30 Requirements discussion
- 12:30 Lunch (and some people leave)

Running NetLogo on ARCHER-2

The James Hutton Institute

Gary Polhill

Command-line access

Wolf-Sheep experiment (16 runs)

- Estimate resource use using the ‘`mgr`’ extension
  - Provides access to the Java Virtual Machine’s runtime environment data
    - `max-mem ; go: max (list max-mem mgr:mem)`
    - **`mgr:cpu-time`** `- start-cpu-time`
      `; setup: start-cpu-time = mgr-cpu-time`
- Run the wolf-sheep model until
  - No sheep and/or no wolves
  - Number of sheep or wolves > 10000

ARCHER-2 Setup

- You have a home directory on login.archer2.ac.uk
  - The ‘head’ node
  - e.g. /home/x999/x999/fblogs
- And a work directory accessible from ARCHER-2’s other machines
  - e.g. /work/x999/x999/fblogs
- Your home directory is *only* accessible on the head node!
  - Even though $HOME is /home/x999/x999/fblogs
  - All your runtime stuff needs to be in /work/…
  - And software should not use your $HOME!

NetLogo

- Uses the Java Virtual Machine
  - Designed for laptop use, not HPC
    - Assumes your computer isn't being used by others
  - Expects your $HOME to be accessible!
- Look at netlogo-headless.sh and JVM_OPTS
  - -XX:MaxRAMPercentage=50
    - Antisocial!

What I had to do in netlogo-headless.sh

- Add this to stop userPrefs error messages:
  - prefs=/var/tmp/javaPrefs.$$
  - mkdir $prefs
  - JVM_OPTS: -Djava.util.prefs.userRoot=$prefs
  - rm -rf $prefs
- Add this to tell Java where my $HOME should be:
  - JVM_OPTS: -Duser.home=/work/x999/x999/fblogs
- Add this to stop Java hogging all the cores on the node with garbage collection threads:
  - JVM_OPTS: -XX:+UseParallelGC -XX:ParallelGCThreads=3
- Add this to tell Java not to use half the RAM on the node!
  - JVM_OPTS: -Xmx1024M

Almost there… ‘Partition’ and ‘QoS’

- ARCHER-2 has a complicated way of allocating jobs to queues according to how ‘big’ they are
  - Number of nodes needed (128 cores each)
  - How much RAM per node
  - How long (‘wallclock time’) to run your job
- You need to know all this (hence the mgr extension)
  - BUT: *In the general case* it is provably impossible to know it!
    - Rice (1953)
  - AND: We should expect ABMs to be unpredictable in this way!
    - Polhill (2022)

CONTENT

ANTISOCIAL SIMULATION: USING SHARED HIGH-PERFORMANCE COMPUTING CLUSTERS TO RUN AGENT-BASED MODELS

DECEMBER 14, 2022 THESUBMISSIONAUTHOR LEAVE A COMMENT

By Gary Polhill

Information and Computational Sciences Department, The James Hutton Institute, Aberdeen AB15 8QH, UK.

High-performance computing (HPC) clusters are increasingly being used for agent-based modelling (ABM) studies. There are reasons why HPC provides a significant benefit for ABM work, and to expect a growth in HPC/ABM applications:

NetLogo experiments – embarallelism

- Suppose I want to sample parameter space multiple times
- Or to find out how long the wolf-sheep model runs to extinction…
- These are embarrassingly parallel problems
  - Tens, or hundreds of thousands of separate runs of the model.
- ARCHER-2 isn't really set up for that…

QoS on ARCHER-2

QoS	Nodes per Job	Wallclock Max	Tasks queued	Tasks running
standard	1024	24h	64	16
highmem	256	24h	16	16
taskfarm	16	24h	128	32
short	32	1200s	16	4
long	64	96h	16	16
largescale	5860	12h	8	1
lowpriority	2048	24h	16	16

https://docs.archer2.ac.uk/user-guide/scheduler/

So how to set up an experiment for ARCHER-2?

- 128 cores per node (standard partition)
  - One big NetLogo BehaviorSpace run per node?
    - 32 threads for runs, 96 for garbage collection?
      - 8GB per run max
  - 1h per run; 24h max = 768 runs per node
    - 16 jobs max at once = 12,288 maximum runs
  - 10 min per run: 73,728 maximum runs
- But we could have up to 1024 nodes…
  - Would need a NetLogo instance that occupied multiple machines simultaneously!
    - Then: 786,432 maximum runs at 1h per run

My wolf-sheep experiment SLURM script


```
#!/bin/sh
#SBATCH --time=00:15:00
#SBATCH --partition=standard
#SBATCH --qos=lowpriority
…
#SBATCH --array=1-16
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=16
…
WORK=/work/x999/x999/fblogs
XML=estimate-resource-use-properly.xml
printf -v EXPT_ID "estimate-resource-use-properly-%04d" $SLURM_ARRAY_TASK_ID
DIR=$WORK/wolf
OUT="$DIR/$EXPT_ID.out"
CSV="$DIR/$EXPT_ID-table.csv"
srun  env HOME=$WORK $WORK/NetLogo-6.4.0-64/netlogo-headless.sh \
     --model "$DIR/wolf.nlogo" --setup-file "$DIR/$XML" --experiment "$EXPT_ID" \
     --threads 1 --table "$CSV" > "$OUT" 2>&1
```

Results (wolf-sheep)

Memory (bytes)	Time (seconds)
356958392	0.69658919
356600192	0.68873805
356606016	0.69343138
357083360	0.80104185
356818600	0.71493286
356886120	0.93375915
356797544	1.00320204
356397464	**1.03023495**
356893344	0.87355938
356838240	0.70507674
356155576	0.76428515
356973656	0.75747006
338193464	**0.35546979**
356175992	0.64679629
357138160	0.65883272
356228080	0.94630621

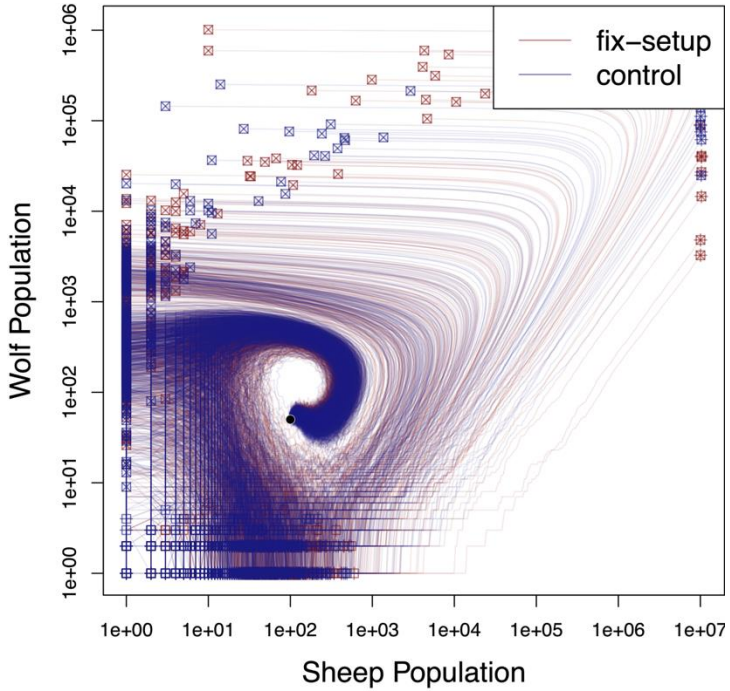

Polhill et al. (under *n*th revision) *JASSS*
(2000 runs on my own cluster)

Summary

- Had to hack netlogo-headless.sh
- Had to tell JVM *and* shell my ‘home’ (work) dir
  - Separately!
- Could not submit my 1,000-run experiment as planned
- Something went wrong with I/O

```
"BehaviorSpace results (NetLogo 6.4.0)","Table version
 2.0"
"/work/ x999/x999/fblogs /wolf/wolf.nlogo"
"estimate-resource-use-properly-0001"
"02/17/2025 02:58:06:432 +0000"
"min-pxcor","max-pxcor","min-pycor","max-pycor"
"-25","25","-25","25"
"[run number]","wolf-gain-from-food","show-energy?","i
nitial-number-wolves","wolf-reproduce","initial-number
-sheep","model-version","sheep-gain-from-food","grass-
regrowth-time","sheep-reproduce","same-setup?","[step]
","max-mem","mgr:cpu-time - start-cpu-time"
"1","20","false","50","5","100","sheep-wolves","4","30
","4","false","382","357138160","0.658832715"
99995"
```

what happened here?

I/O on ARCHER-2

- File Per Process
  - “One of the first parallel strategies people use for IO”
  - “Having thousands of files open at once can overload the system”
- File Per Node
  - Better… and achievable with one massive NetLogo instance per node
- Single Shared File with collective writes
  - Coordinated writing to avoid locking

https://docs.archer2.ac.uk/user-guide/io/

Agenda (Wednesday AM)

- 09:00 Arrival and coffee

- 09:30 Hands-on cloud-based FLAME GPU Tutorial
- 10:30 Running NetLogo on ARCHER-2
- 11:00 Coffee

- 11:30 Requirements discussion
- 12:30 Lunch (and some people leave)

Agenda (Wednesday PM)

- 14:00 Hackathon: decide on options
- 14:30 Hackathon: break out into groups and hack
- 15:30 Coffee
- 16:00 Hackathon: more hacking
- 17:00 Hackathon: report back
- 17:30 Close
- 18:30 Dinner:
  Tas Bloomsbury, 22 Bloomsbury Street

Agenda (Thursday AM)

- 09:00 Arrival and coffee

- 09:30 Hackathon: recap and plan
- 10:00 Hackathon: more hacking
- 11:00 Coffee

- 11:30 Hackathon: report back
- 11:50 Closing Remarks
- 12:00 Lunch

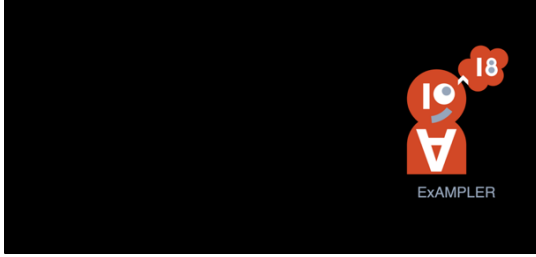

Funded by the Engineering and Physical Sciences Research Council, Grant Number EP/Y008839/1, with support from the Scottish Government Rural and Environment Science and Analytical Services Division (project reference JHI-C5-1)