A flexible hydrodynamic modelling framework for GPUs and CPUs: Application to urban flood events

Luke S. Smith
Newcastle University

RainGain Intl. Workshop on Urban Pluvial Flood Modelling, Met Office, 6th October 2014

With thanks to Qiuhua Liang, Paul Quinn, Jingming Hou, Dave Alderson, Chris Kilsby, Jeff Neal, Environment Agency and Met Office.

Part-funded by grants from EPSRC (EP/K031678/1) and NERC (NE/K00896X/1).

Newcastle University

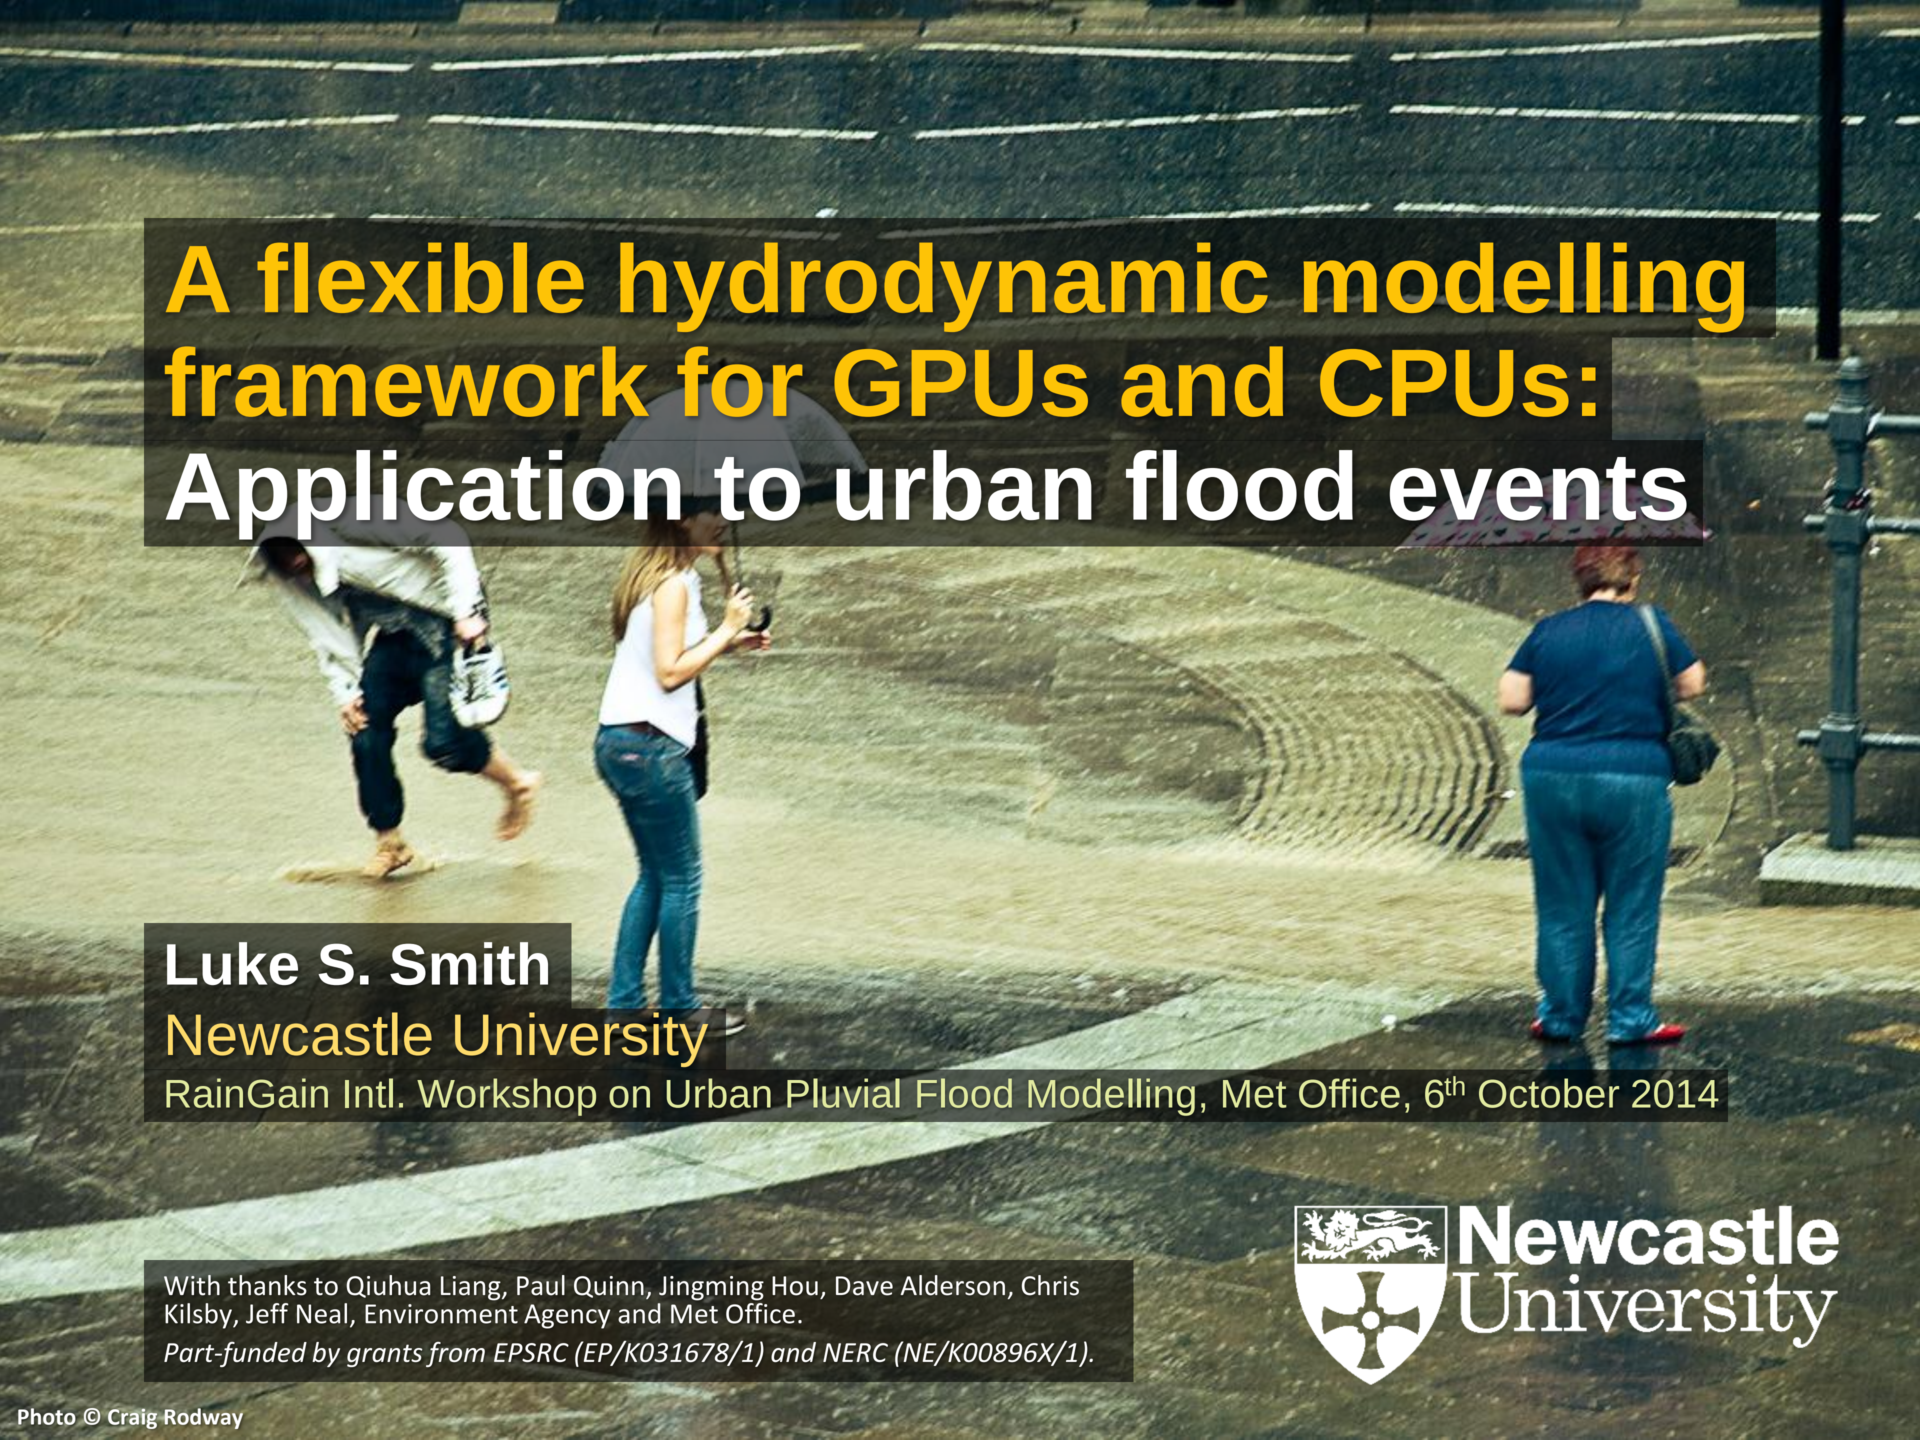

Rationale

- **Detailed catchment-scale modelling**
  Belford, Northumberland
  - Natural flood management
  - Barriers, pipes, hydrology…
- **Tsunami propagation**
  Japan/China
  - Multi-scale real-time warnings
- **Broad-scale flood risk analyses**
  London
  - Complex urban topography
  - Multi-source flood potential

Shallow flow models

- **Godunov-type finite-volume scheme**
  (Smith *et al.* 2014 in Urban Water Journal)
  - Shock-capturing
  - Stencil operations
  - Directionally unsplit
- **Second-order accuracy**
  (Smith and Liang 2013 in Computers & Fluids)
  - MUSCL-Hancock
- **Partial-inertial simplification**
  (Bates *et al.* 2010 in J. Hydrology)
  - No Riemann solver required

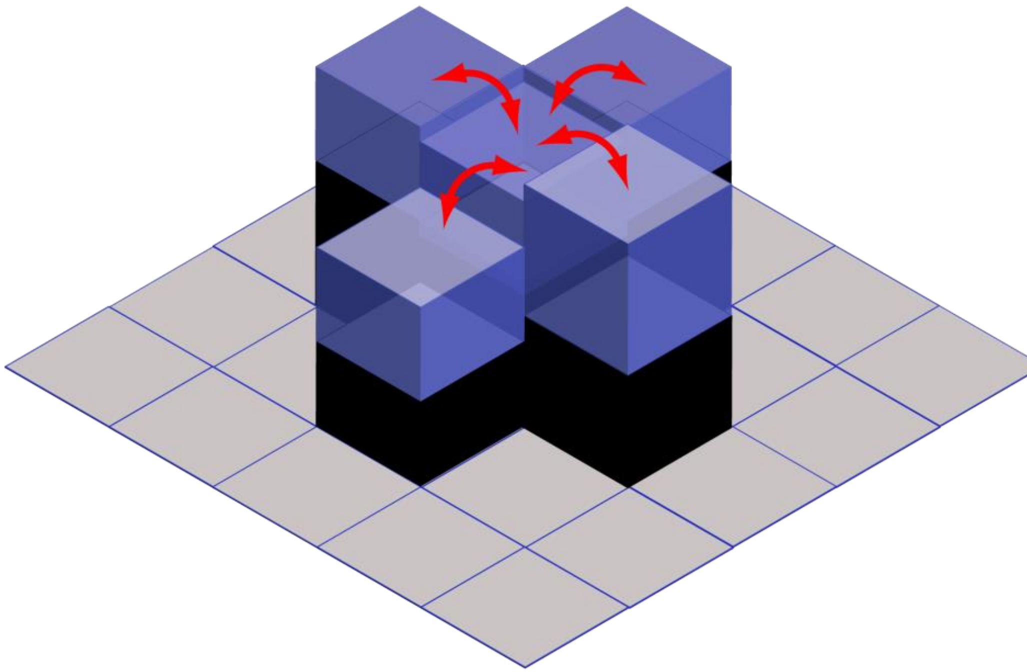

History and Moore's 'law'

The CPU in 1997

- 8.8 million transistors
- $469 at launch

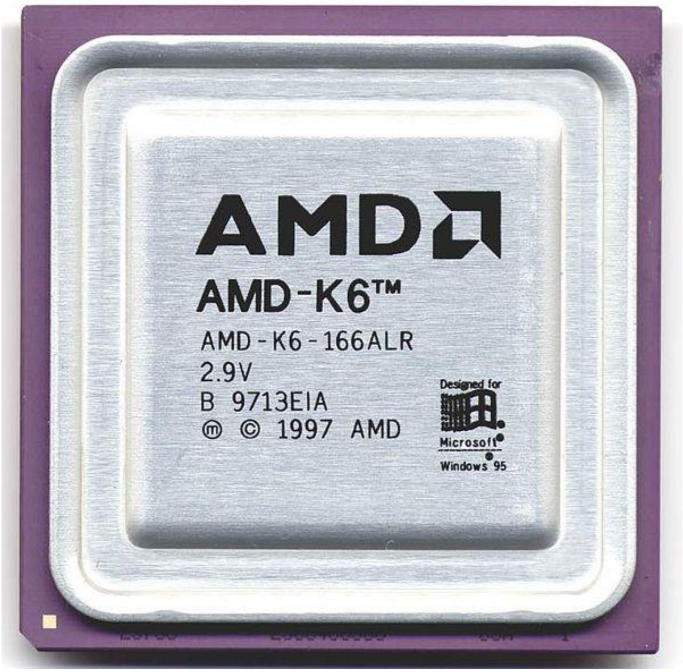

History and Moore’s ‘law’

The CPU in 1997

- 8.8 million transistors
- $469 at launch

Exponential increase in power

- Gordon Moore’s ‘Law’
- Electronics magazine, 1965

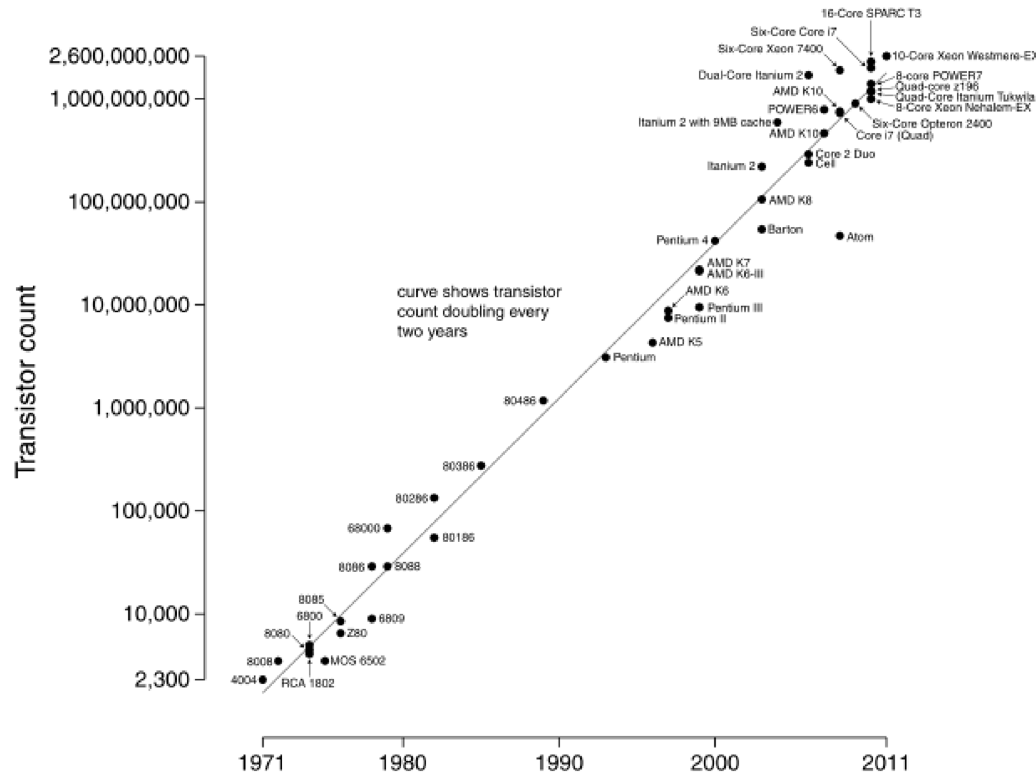

History and Moore's 'law'

The CPU in 1997

- 8.8 million transistors
- $469 at launch

Exponential increase in power

- Gordon Moore's 'Law'
- Electronics magazine, 1965

Multi-core transition

- Stagnant clock speeds

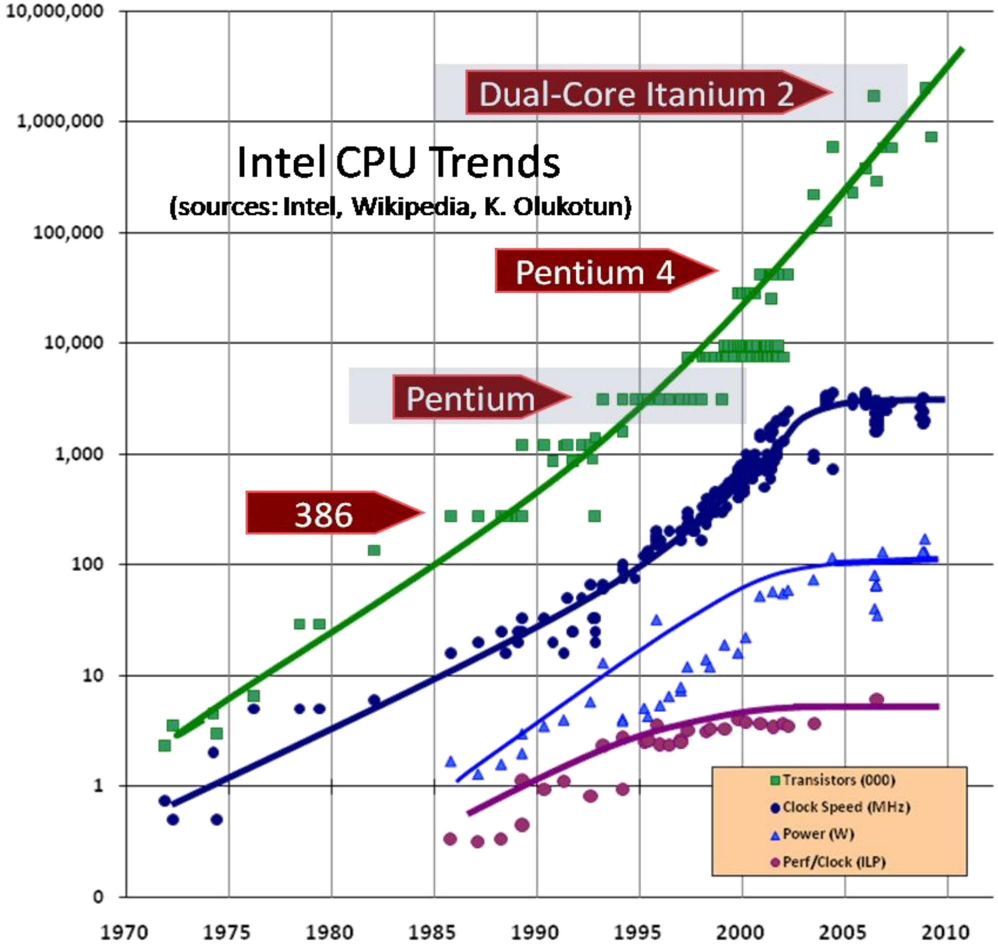

Heterogeneous architectures

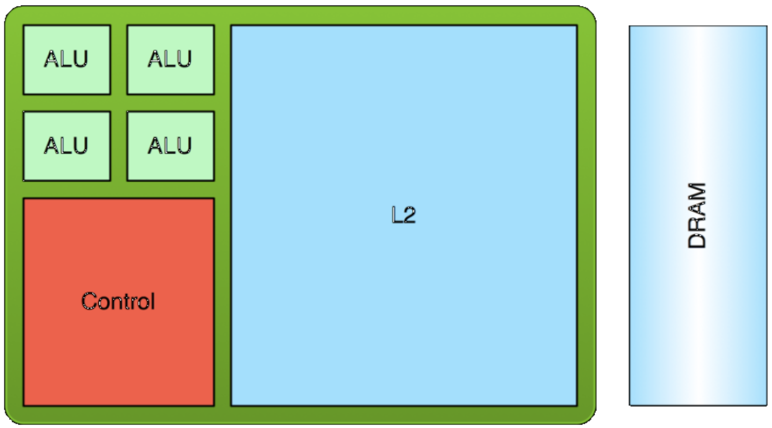

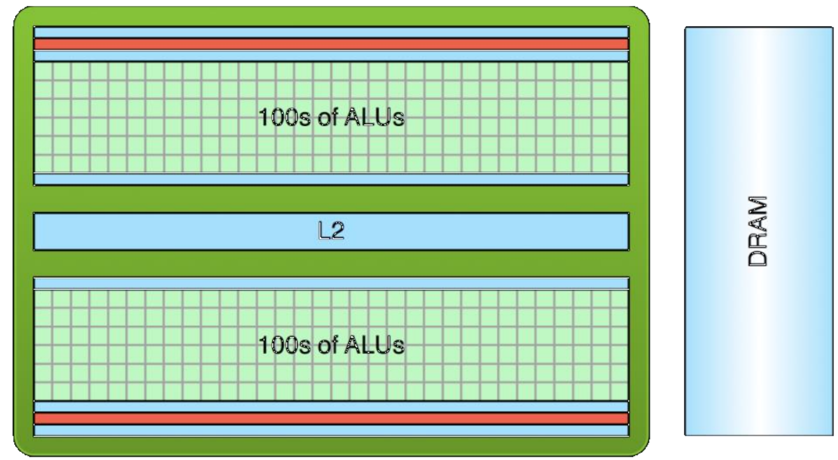

CPU

- **Low-latency** processor
- Low ratio of ALUs/control units
- **Fast response** times
- High RAM capacities
- General purpose computation

GPU… (APUs, Coprocessors, etc.)

- High-throughput processor
- High ratio of ALUs/control units
- Significant **host-bus latency**
- **Limited DRAM** capacities
- Challenges for optimisation

Graphics processing unit

	CPU	High-grade GPUs	Compute Server
			
Model:	Intel Xeon E5-2609	AMD FirePro V7800 (x2)	NVIDIA Tesla M2075 (x4)
Cost:	£311	£1,200	£8,400
SP GFLOPS:	*Unknown*	4032	5152
DP GFLOPS:	77	806	2064
Memory:	2GB (in price)	4GB (2GB x 2)	24GB (6GB x 4)
Cost/GFLOP:	£4.04	£1.49	£4.07
DP multiplier/device:	1.0x	5.2x	6.7x

HiPIMS

- **Research platform**
- **Plug-in based software**
  - Sediment transport
  - Discrete element modelling
- **Visualisation in real-time**
- **Dynamic code generation**
  - Cross-platform, cross-architecture
  - Intel CPU, NVIDIA/AMD GPU, IBM Cell, Parallella, etc.

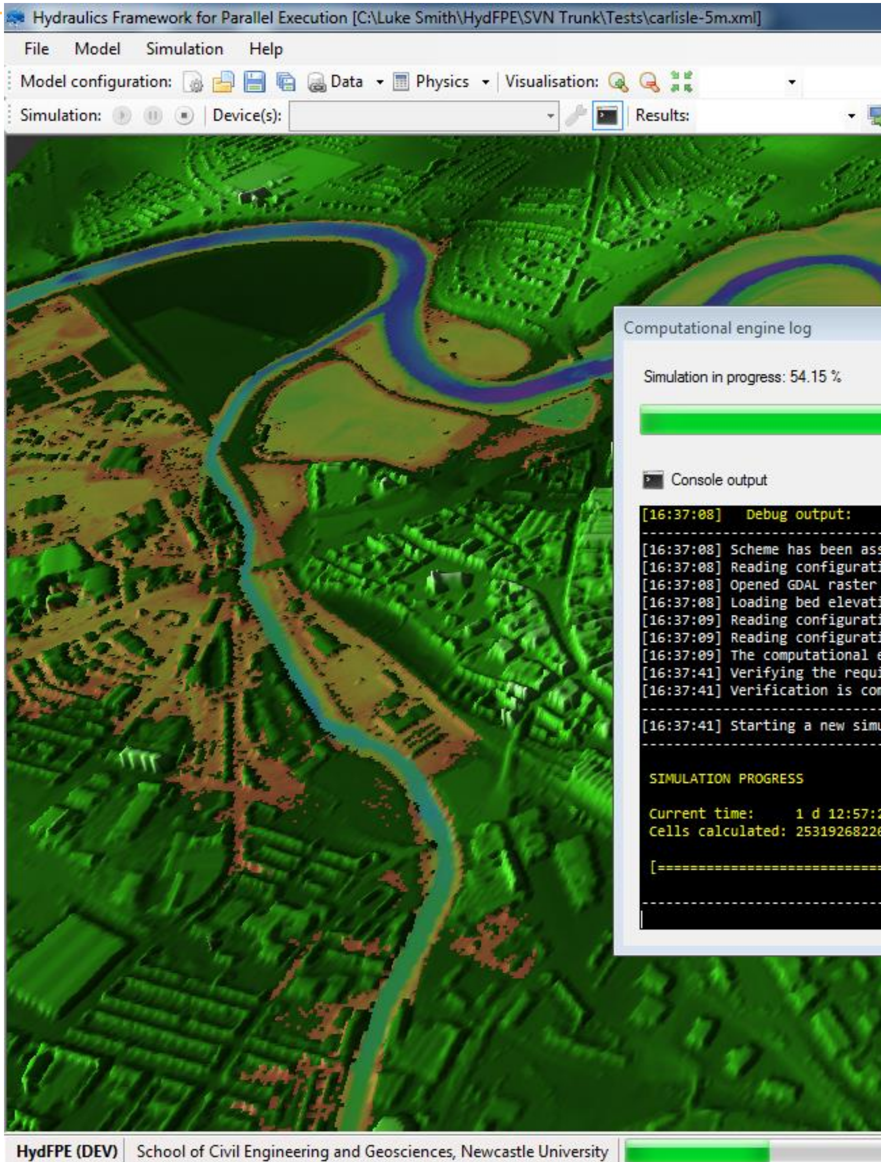

HiPIMS

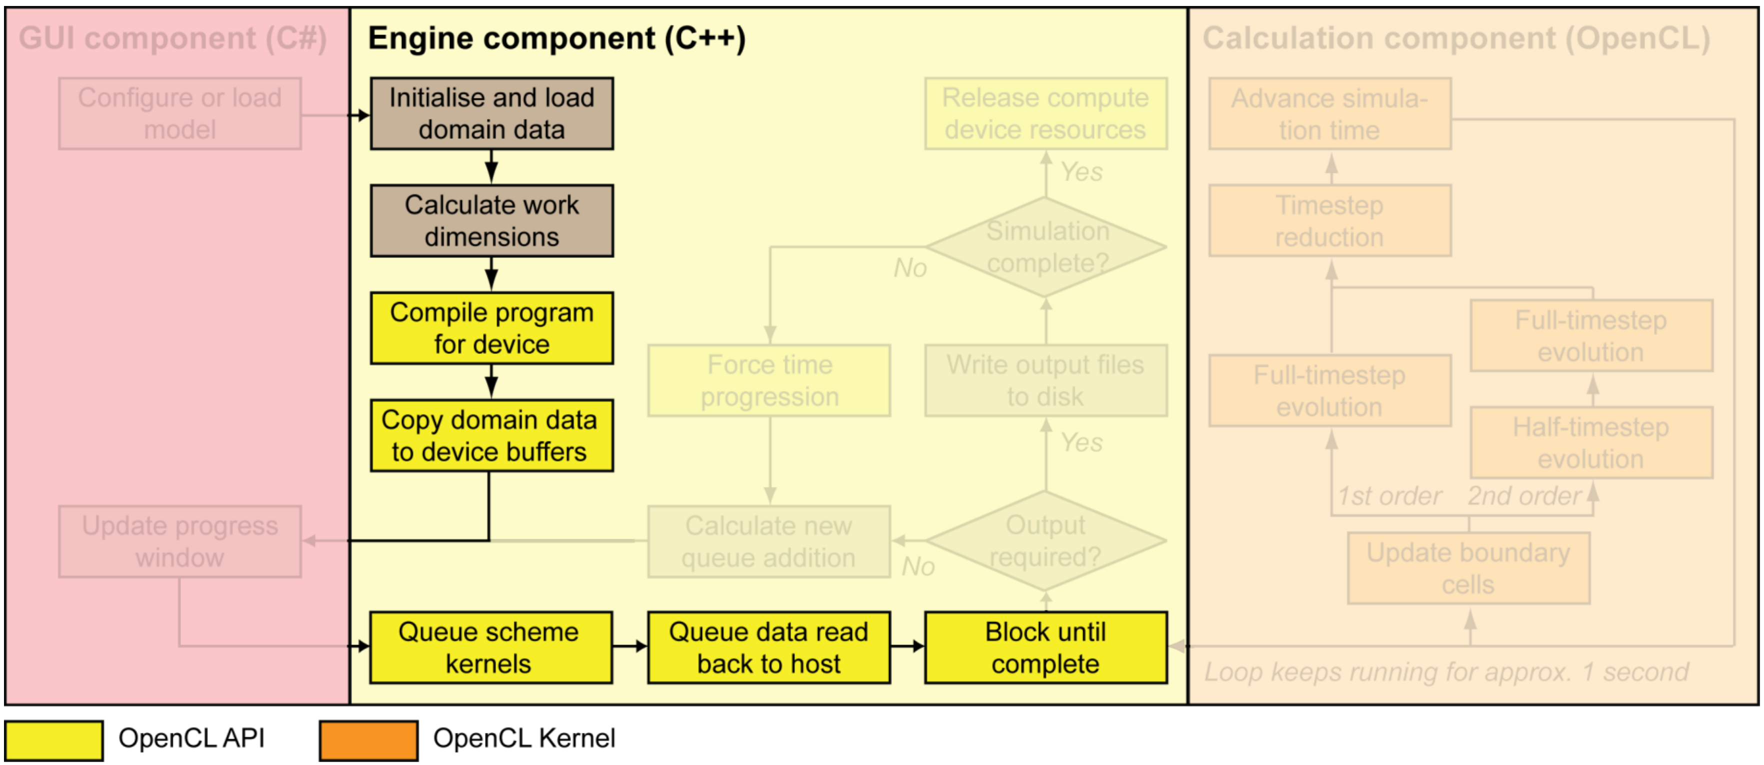

Smith LS, Liang Q (2013) Towards a generalised GPU/CPU shallow-flow modelling tool, *Computers & Fluids*, 88:334-343.

HiPIMS

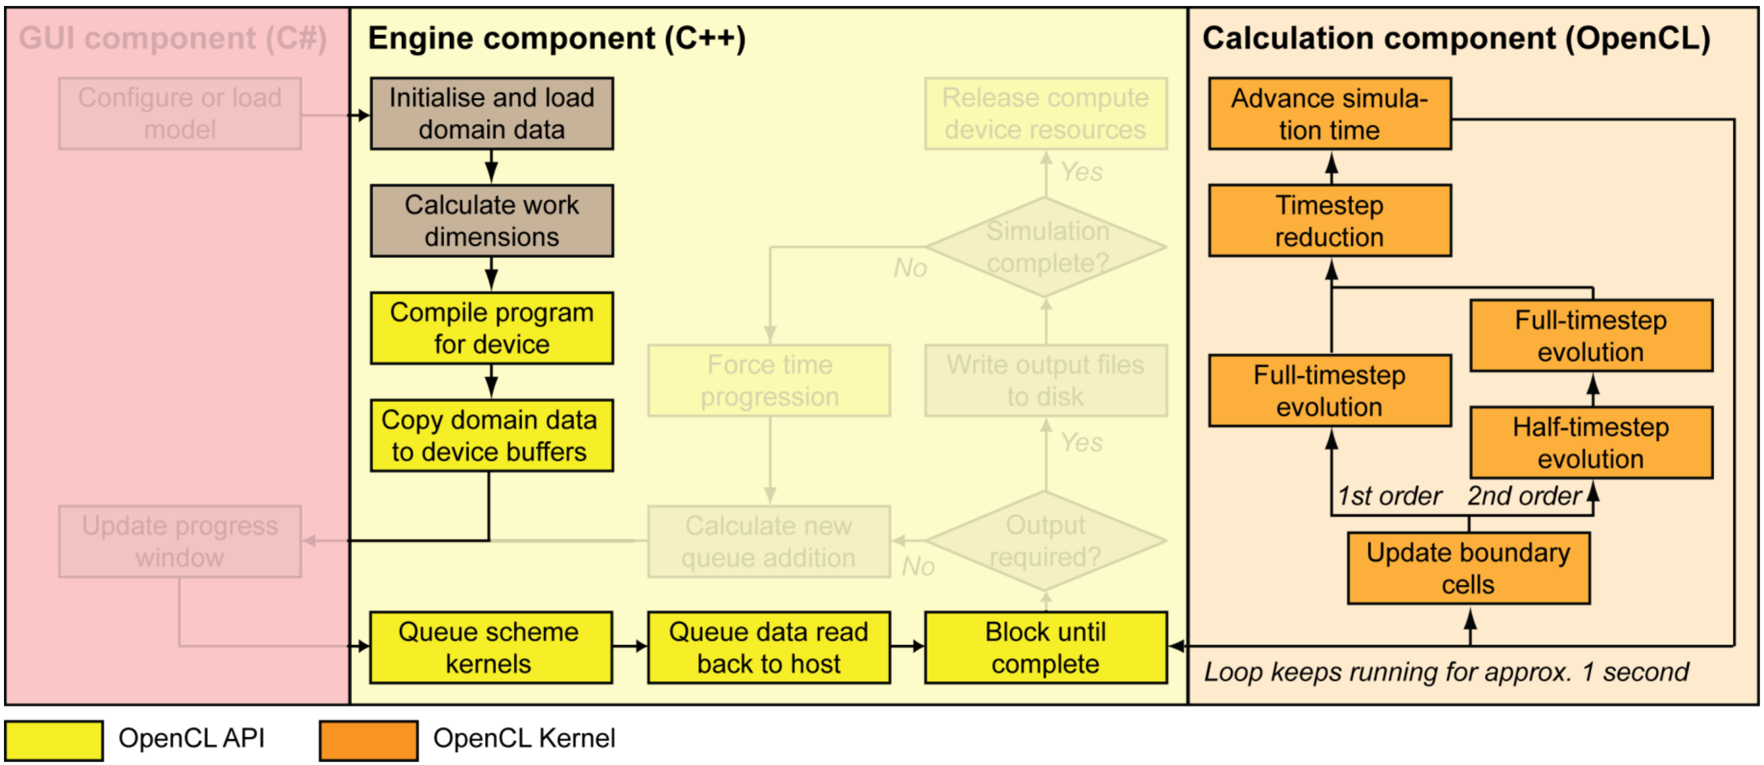

Smith LS, Liang Q (2013) Towards a generalised GPU/CPU shallow-flow modelling tool, *Computers & Fluids*, 88:334-343.

HiPIMS

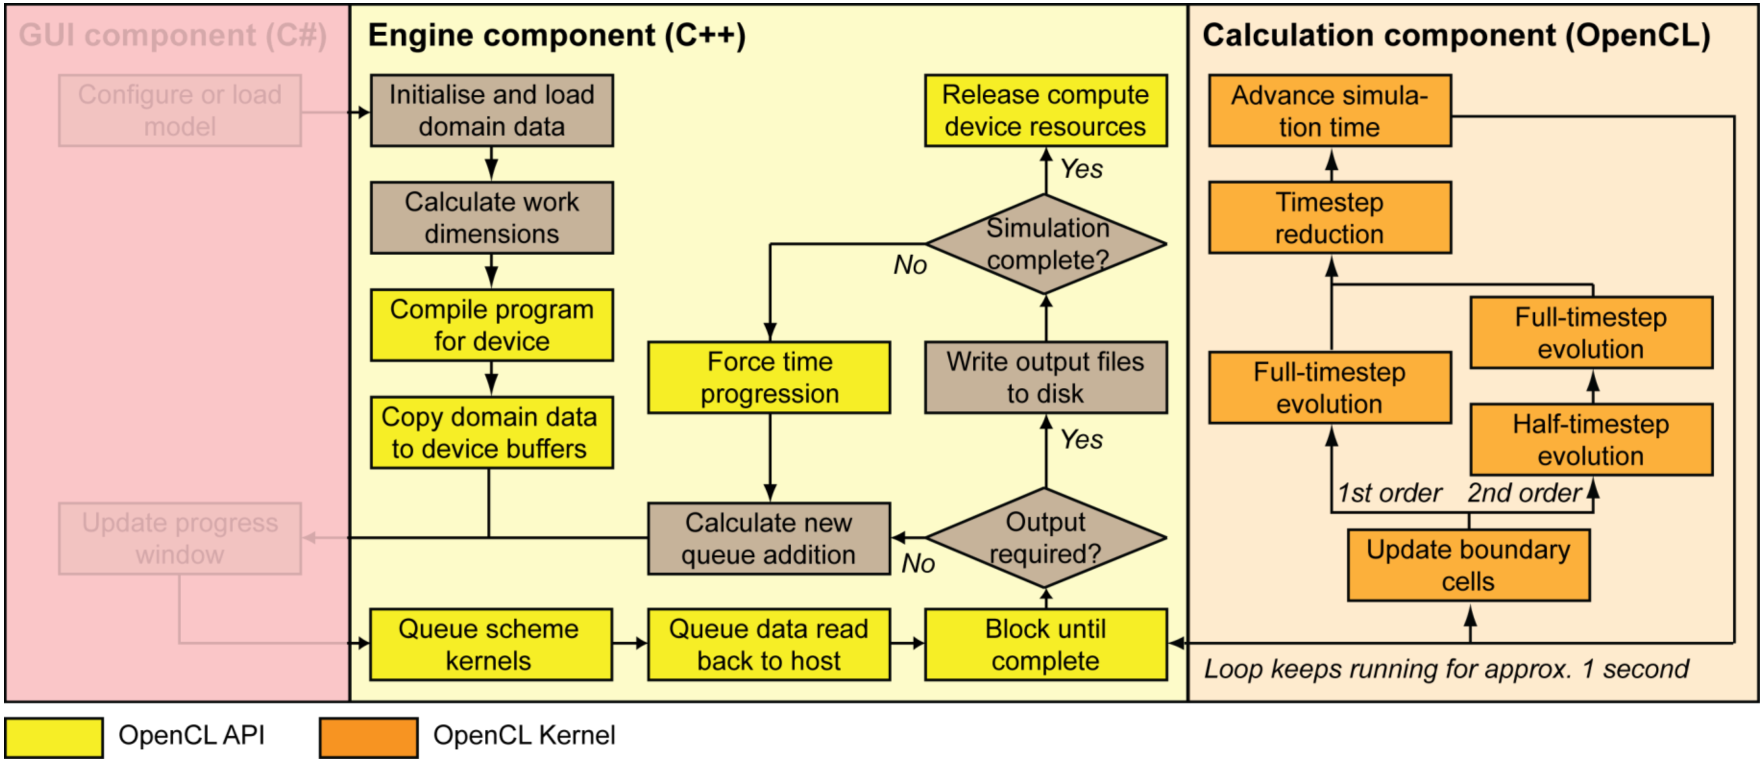

Smith LS, Liang Q (2013) Towards a generalised GPU/CPU shallow-flow modelling tool, *Computers & Fluids*, 88:334-343.

HiPIMS

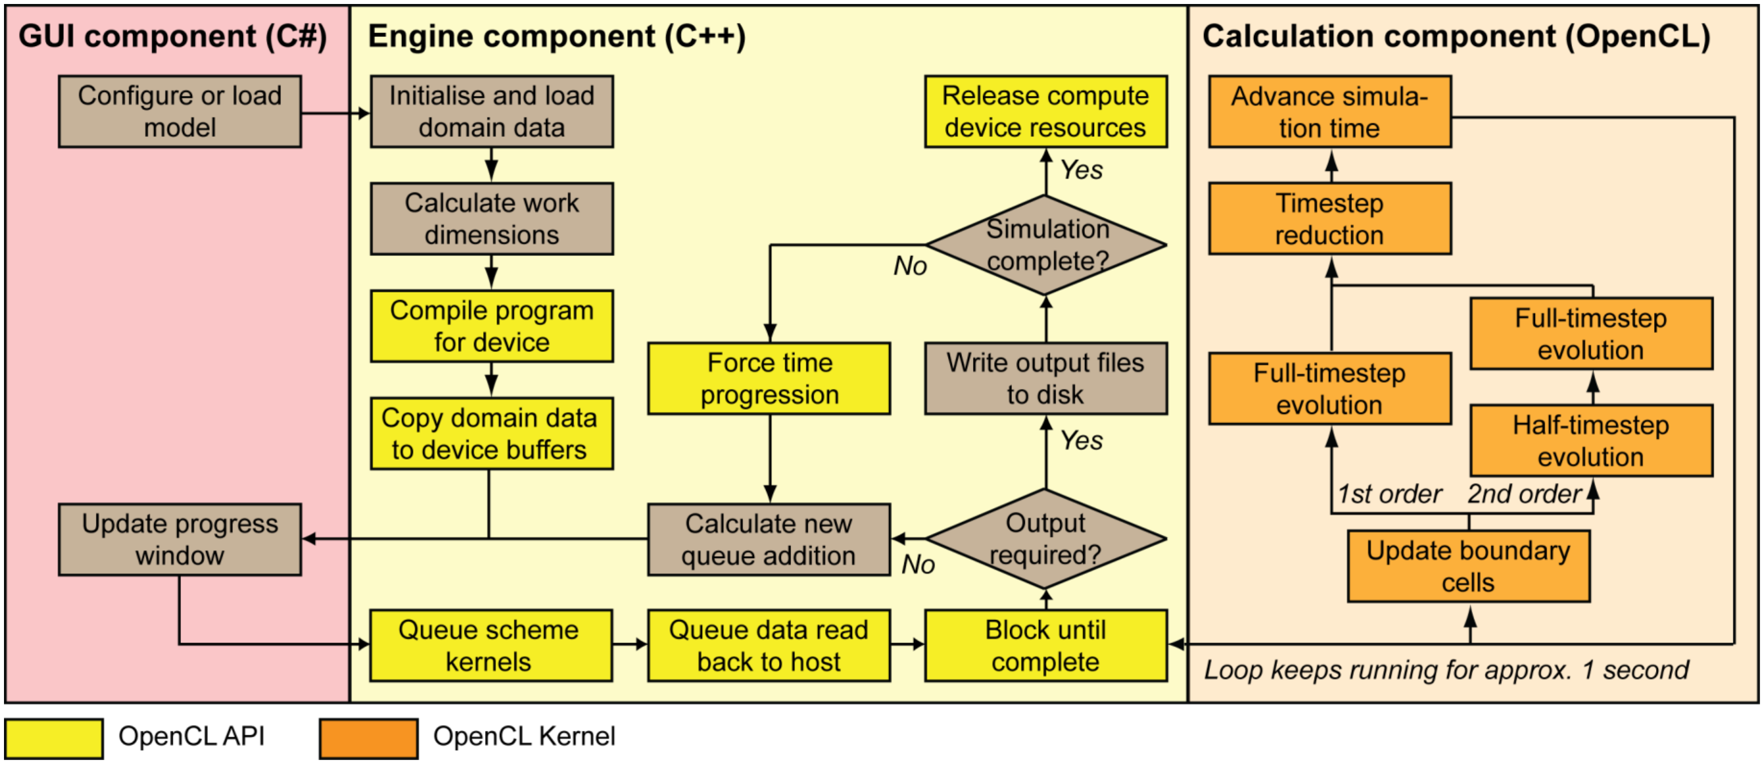

Smith LS, Liang Q (2013) Towards a generalised GPU/CPU shallow-flow modelling tool, *Computers & Fluids*, 88:334-343.

Why bother?

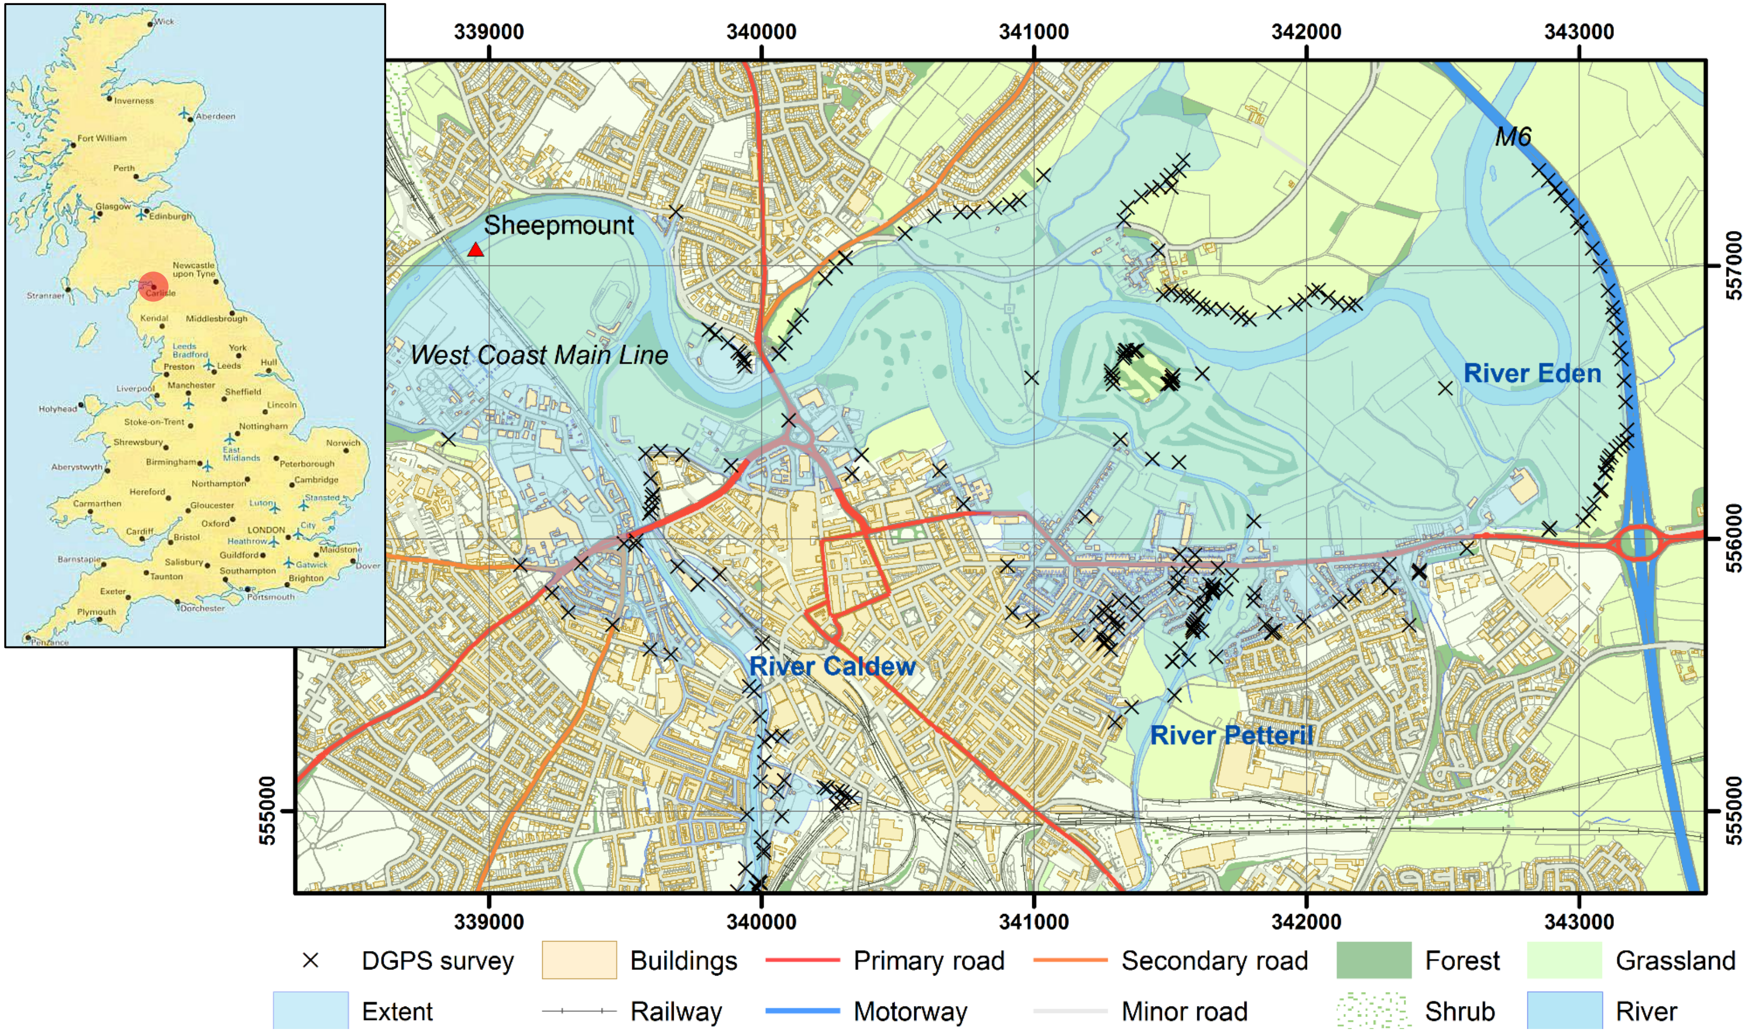

339000
340000
341000
342000
343000
557000
556000
555000
Sheepmount
West Coast Main Line
M6
River Eden
River Caldew
River Petteril
DGPS survey
Buildings
Primary road
Secondary road
Forest
Grassland
Extent
Railway
Motorway
Minor road
Shrub
River

Why bother?

...grid resolution

[25m]

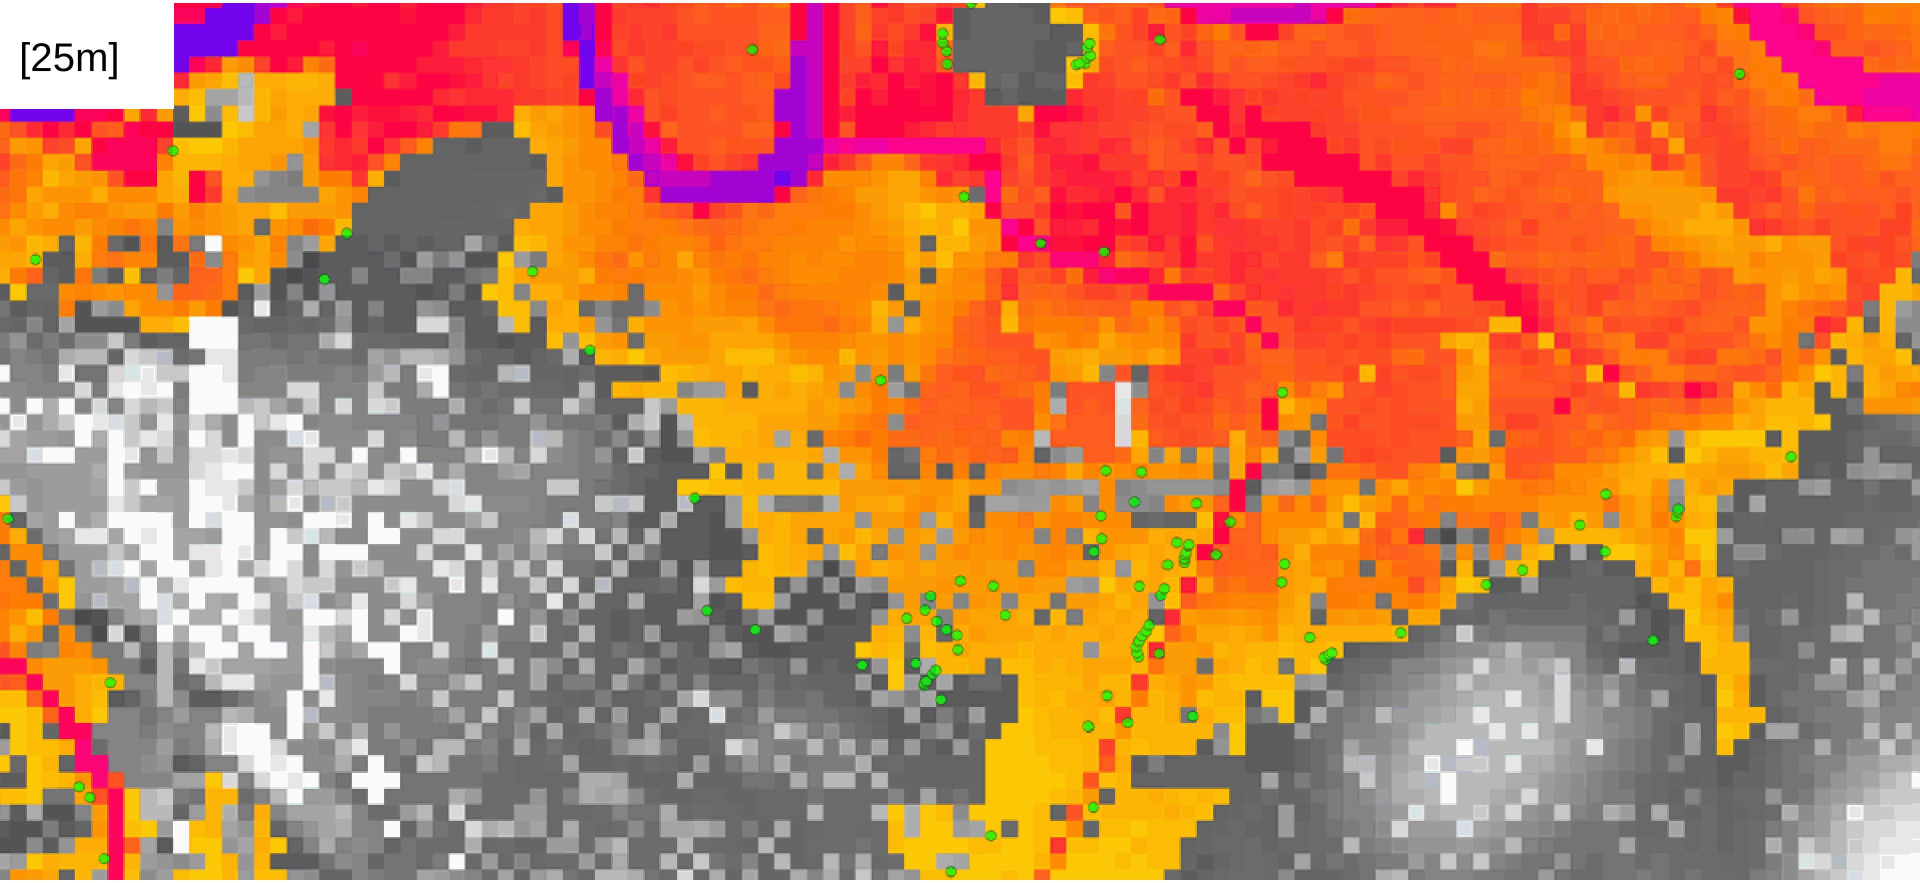

Simulation: 50 hours with FV SWEs 1st order

Why bother?

…grid resolution

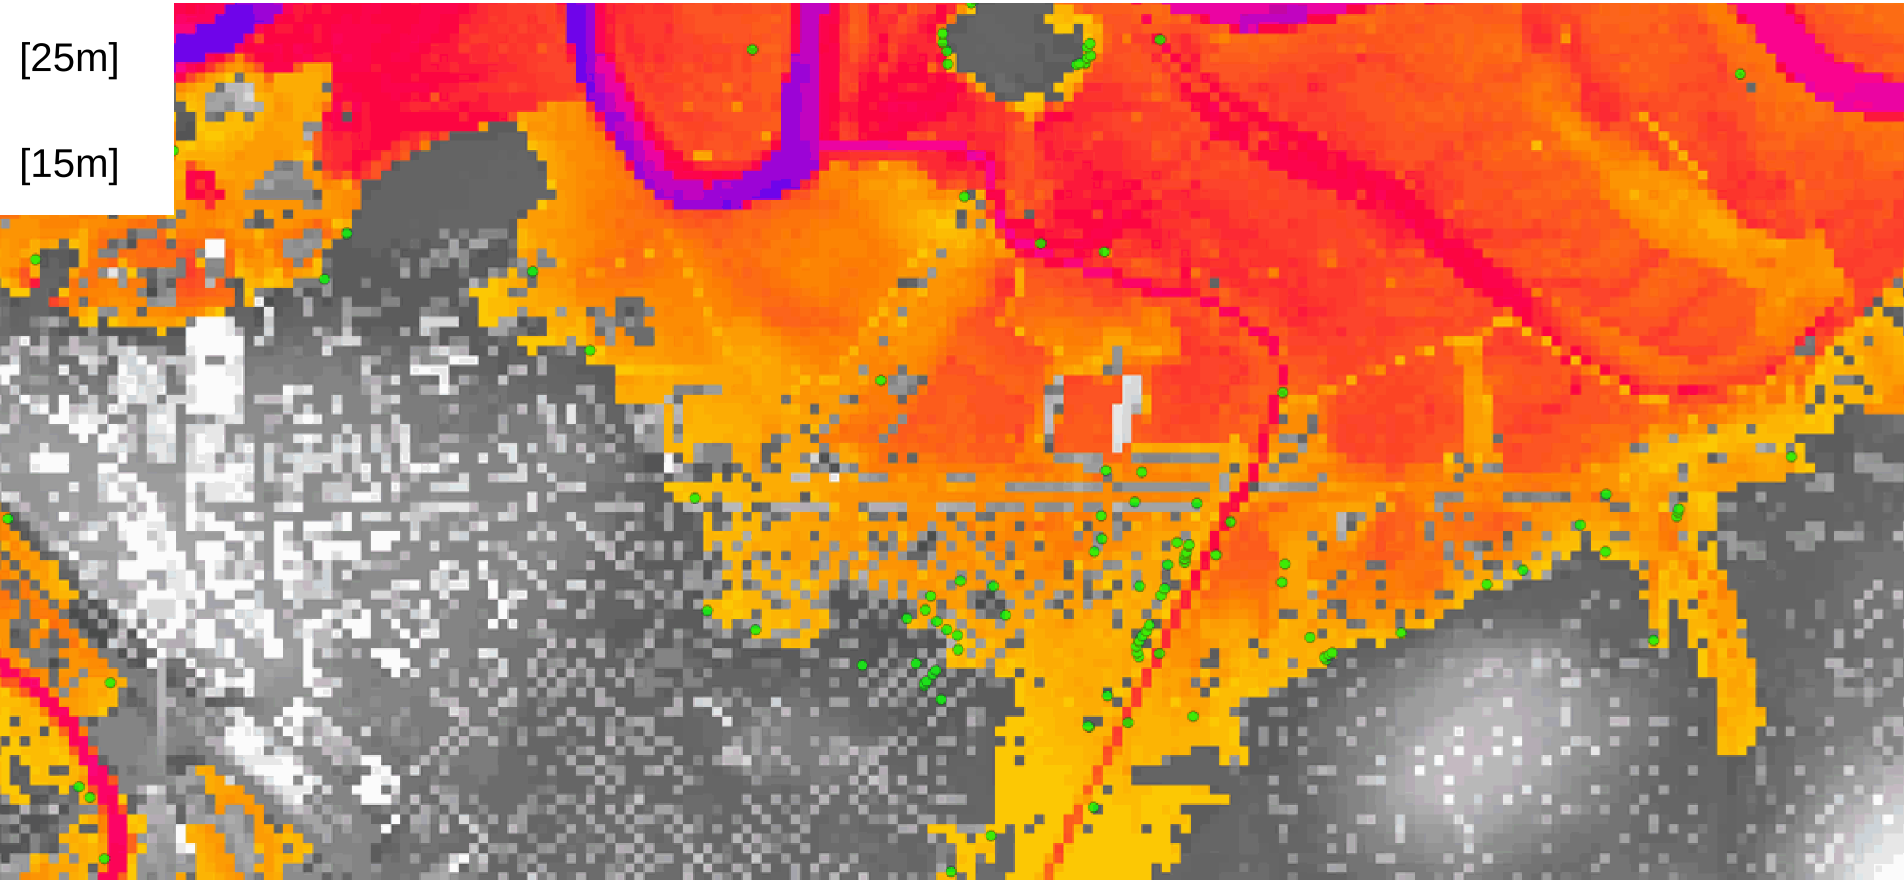

Simulation: 50 hours with FV SWEs 1st order

Why bother?

…grid resolution

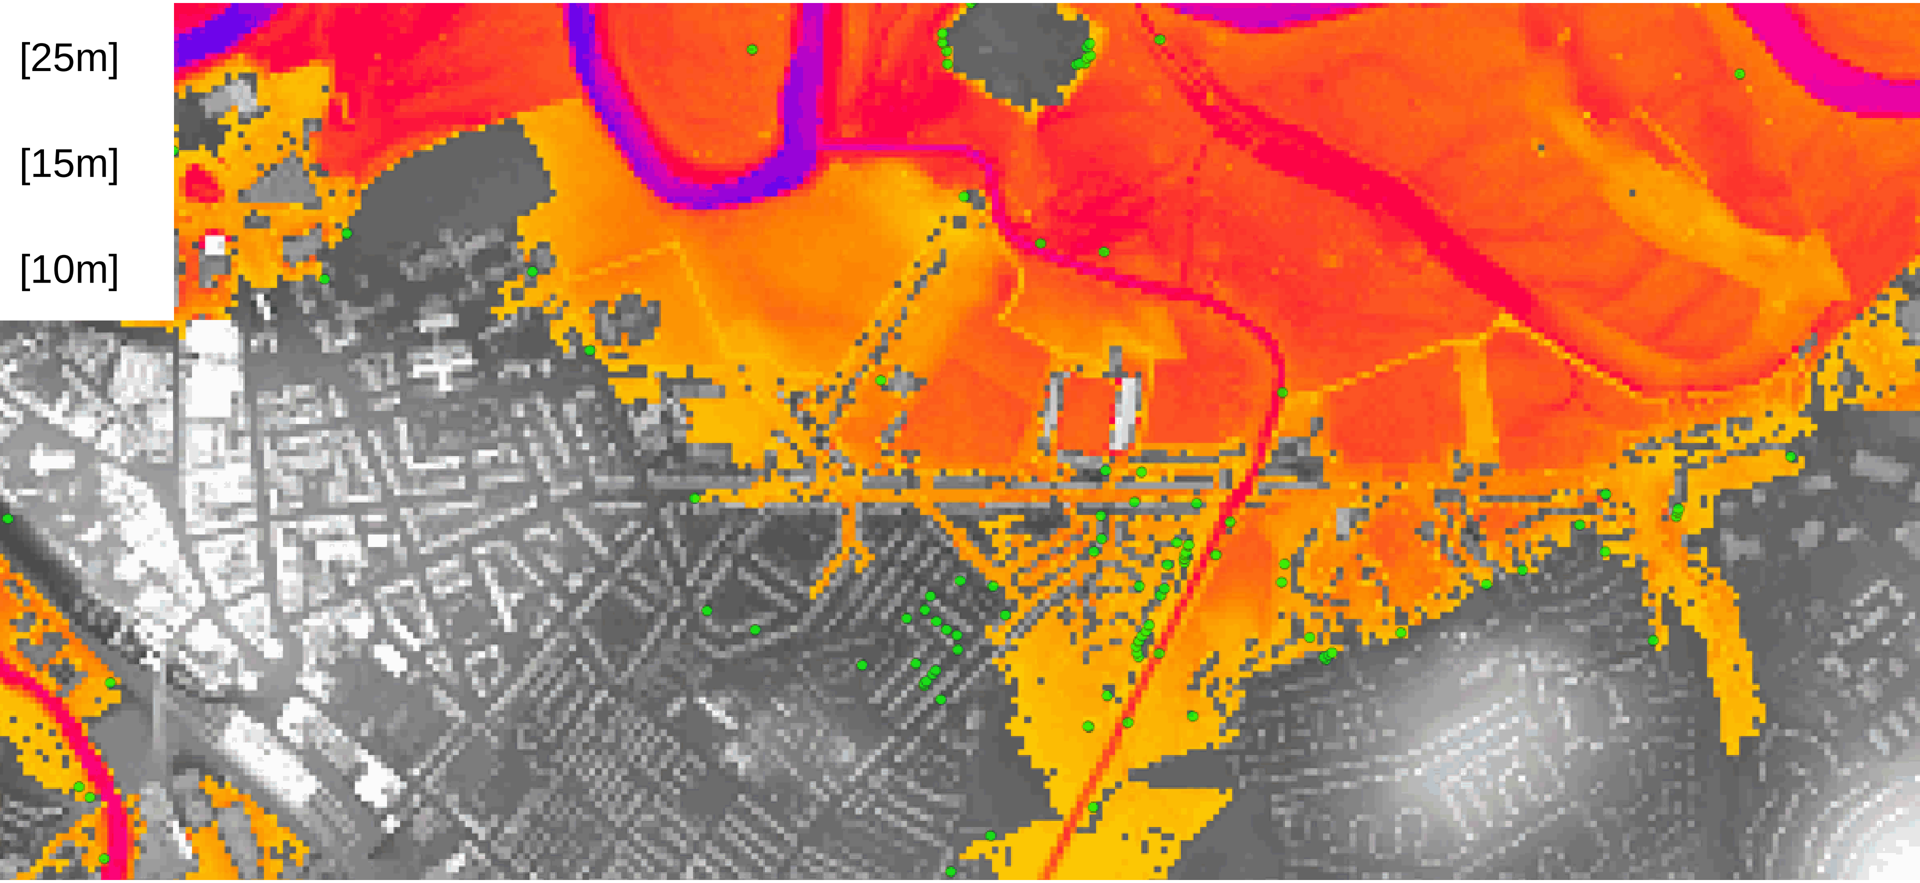

Simulation: 50 hours with FV SWEs 1st order

Why bother?

…grid resolution

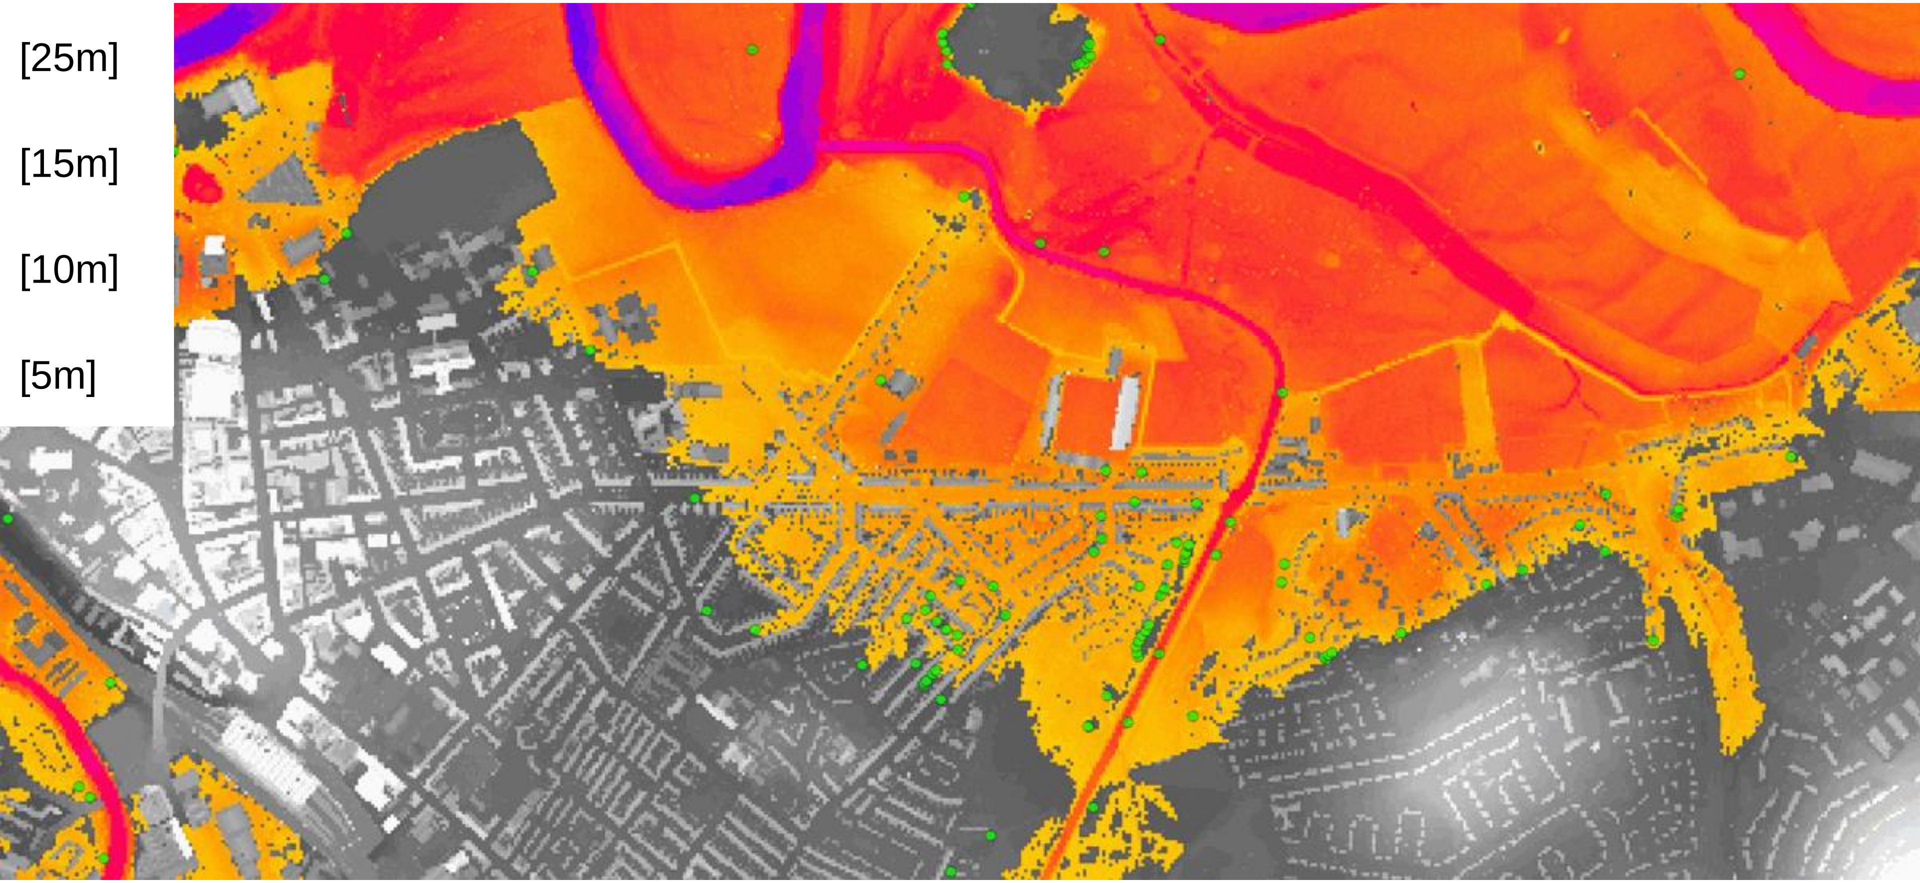

Simulation: 50 hours with FV SWEs 1st order

Why bother?

…grid resolution

[25m]

[15m]

[10m]

[5m]

[2m]

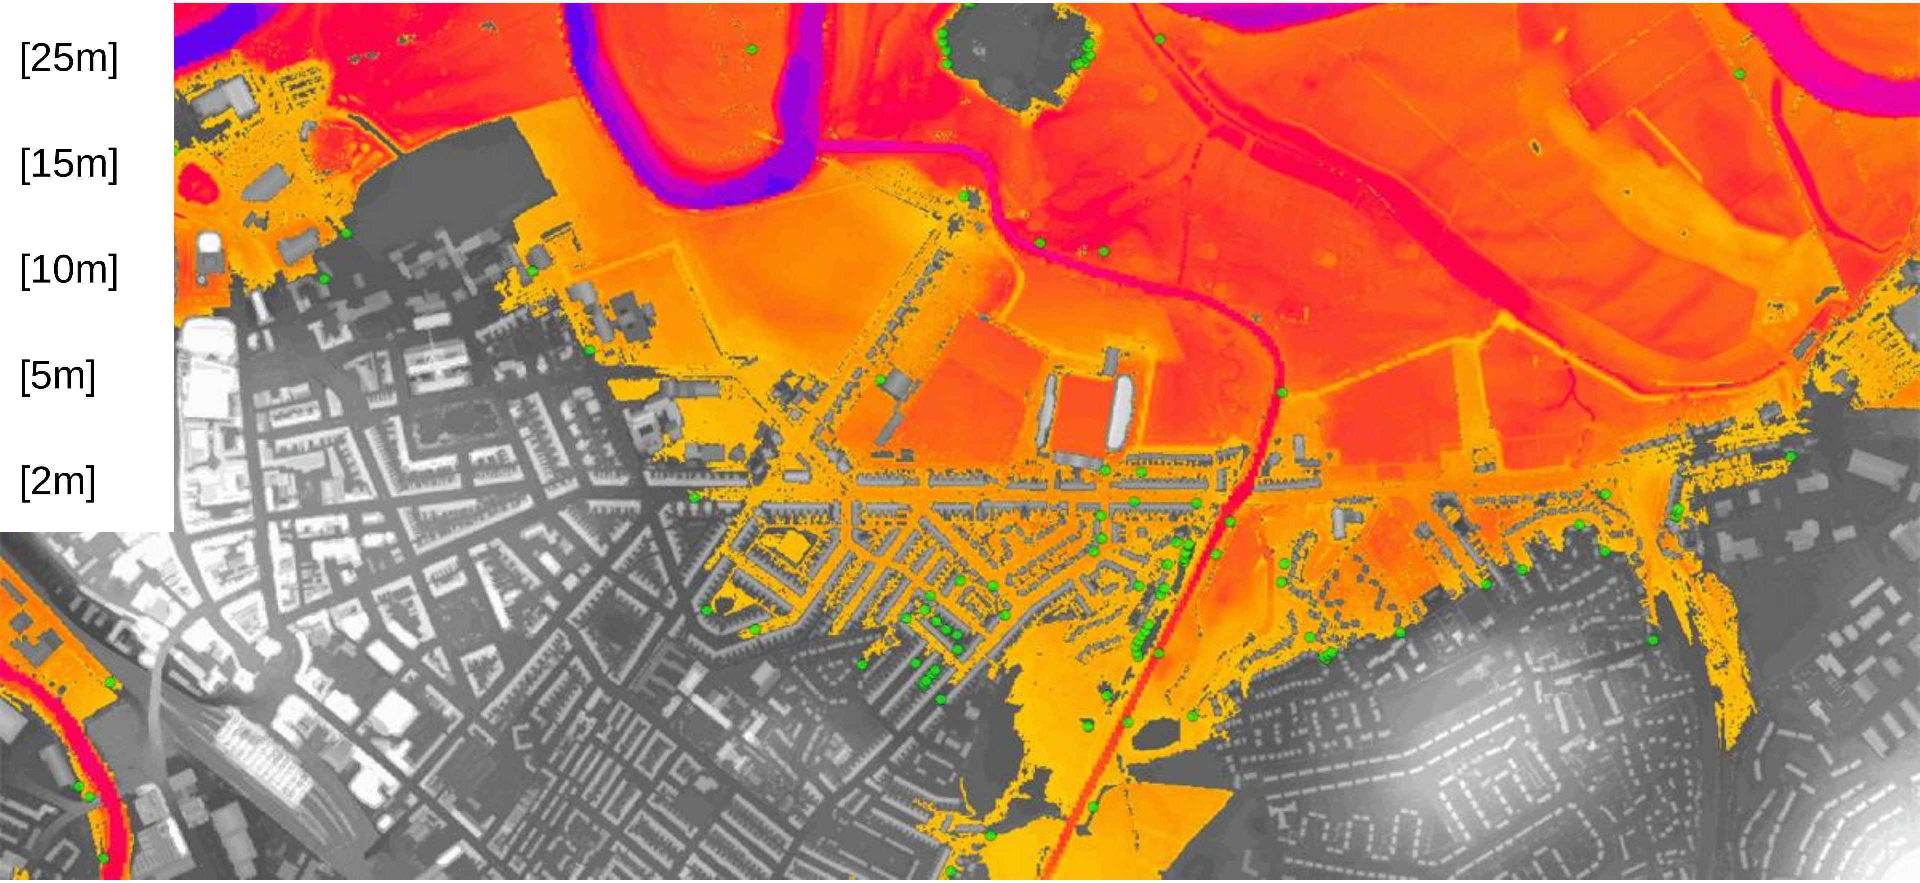

Simulation: 50 hours with FV SWEs 1st order

Why bother?

...sensitivity

- Good match attainable at all resolutions tested
- Calibration only possible against known extent and river hydrometry
- Sensitivity varies significantly
    - Low velocity flows, flood defences overtopped
    - Expect high resolutions to give low floodplain sensitivity

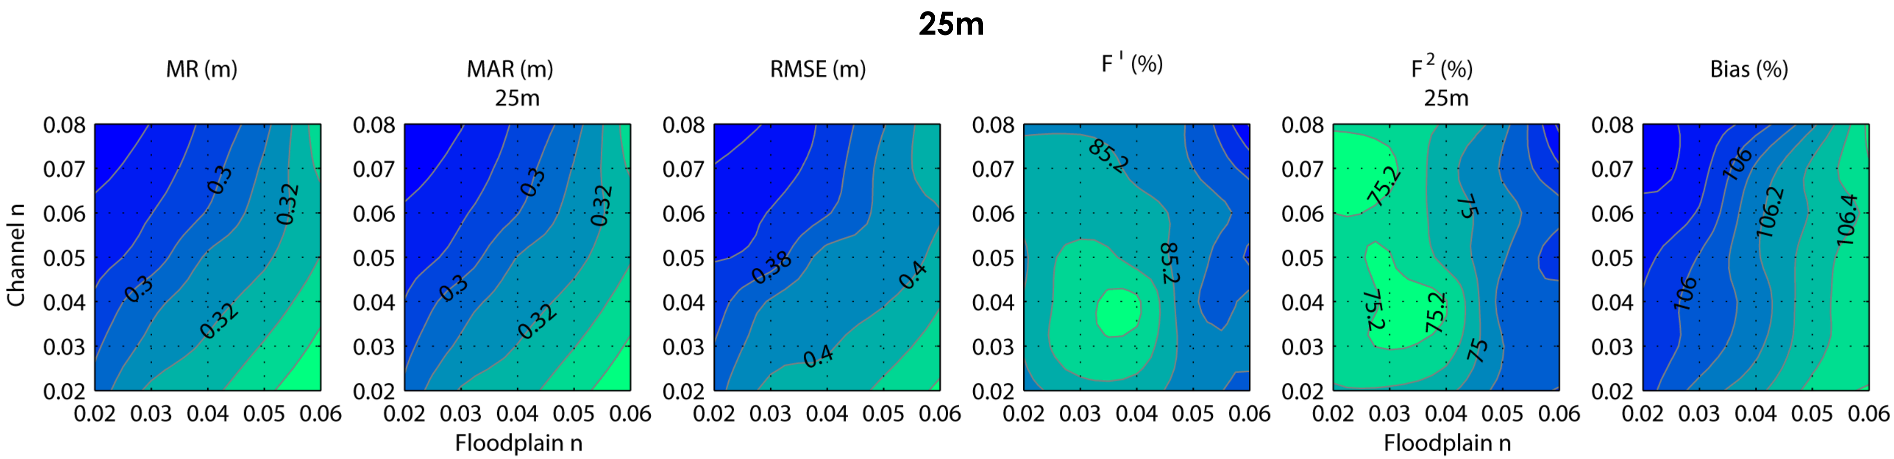

Smith LS, Liang Q, Quinn PF (2014) Towards a hydrodynamic modelling framework appropriate for applications in urban flood assessment and mitigation using heterogeneous computing, *Urban Water Journal*, [in press].

Why bother?

...sensitivity

- Good match attainable at all resolutions tested
- Calibration only possible against known extent and river hydrometry
- Sensitivity varies significantly
  - Low velocity flows, flood defences overtopped
  - Expect high resolutions to give low floodplain sensitivity

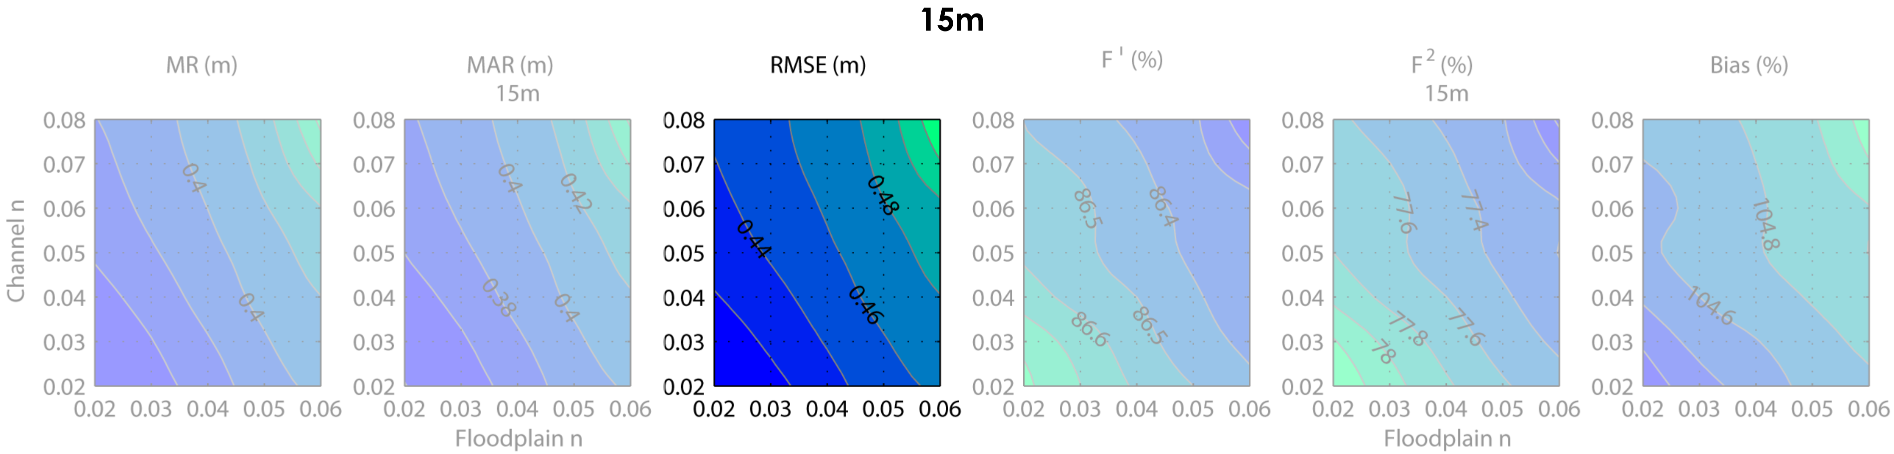

Smith LS, Liang Q, Quinn PF (2014) Towards a hydrodynamic modelling framework appropriate for applications in urban flood assessment and mitigation using heterogeneous computing, *Urban Water Journal*, [in press].

Why bother?

…sensitivity

- Good match attainable at all resolutions tested
- Calibration only possible against known extent and river hydrometry
- Sensitivity varies significantly
    - Low velocity flows, flood defences overtopped
    - Expect high resolutions to give low floodplain sensitivity

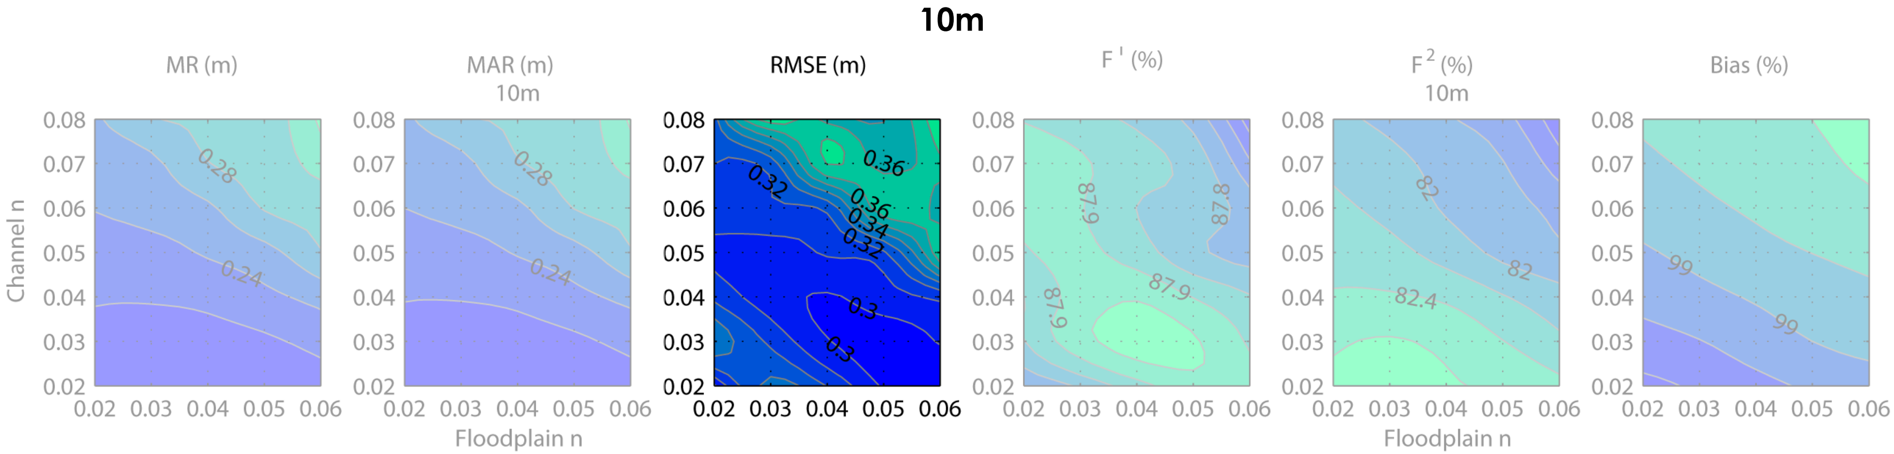

Smith LS, Liang Q, Quinn PF (2014) Towards a hydrodynamic modelling framework appropriate for applications in urban flood assessment and mitigation using heterogeneous computing, *Urban Water Journal*, [in press].

Why bother?

...sensitivity

- Good match attainable at all resolutions tested
- Calibration only possible against known extent and river hydrometry
- Sensitivity varies significantly
  - Low velocity flows, flood defences overtopped
  - Expect high resolutions to give low floodplain sensitivity

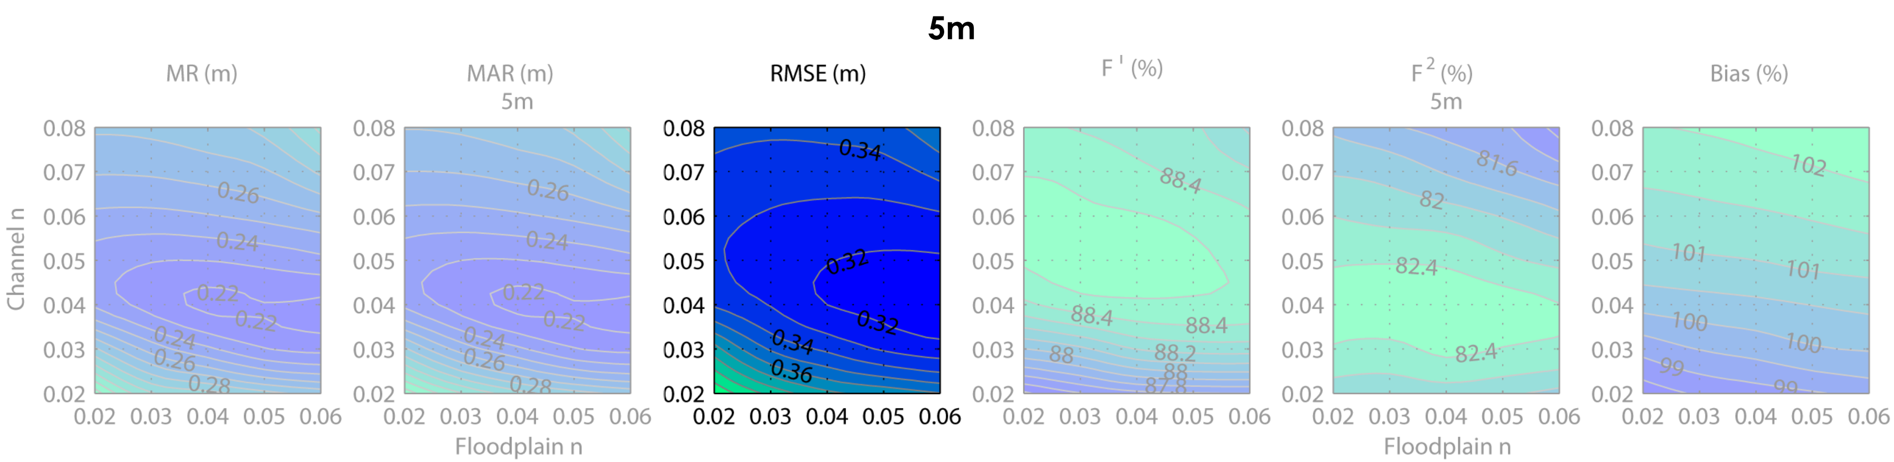

Smith LS, Liang Q, Quinn PF (2014) Towards a hydrodynamic modelling framework appropriate for applications in urban flood assessment and mitigation using heterogeneous computing, *Urban Water Journal*, [in press].

...sensitivity

- Good match attainable at all resolutions tested
- Calibration only possible against known extent and river hydrometry
- Sensitivity varies significantly
  - Low velocity flows, flood defences overtopped
  - Expect high resolutions to give low floodplain sensitivity

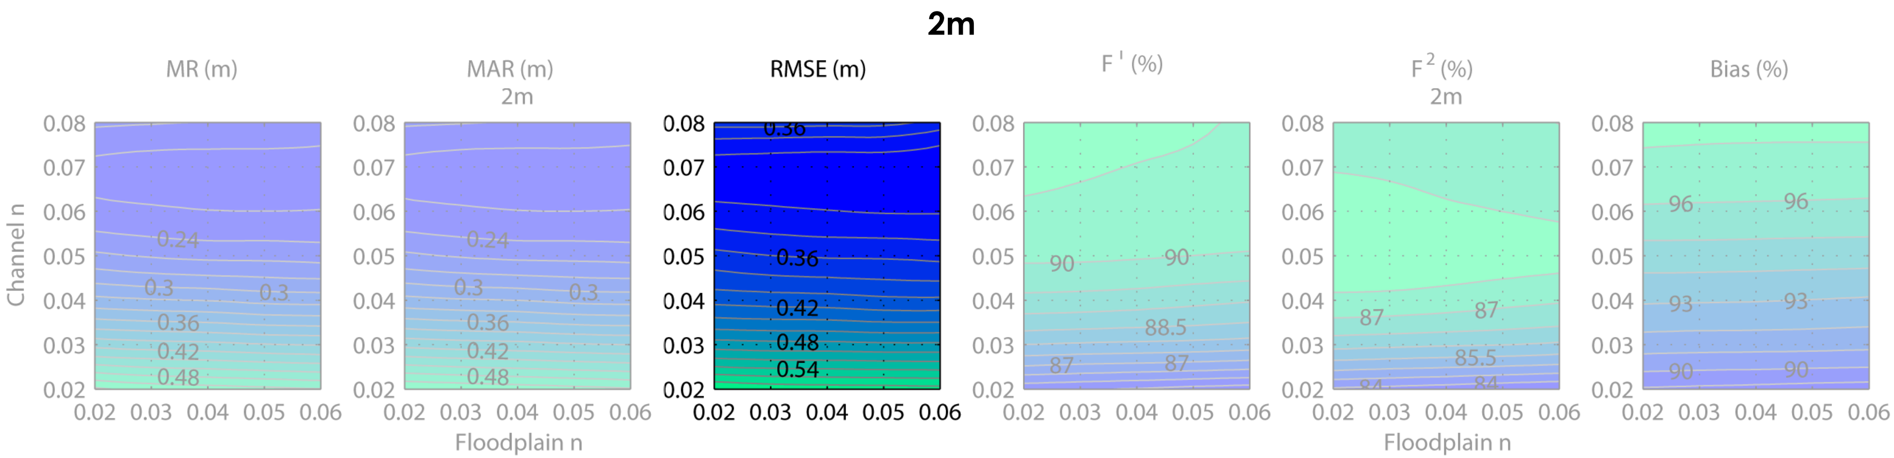

Smith LS, Liang Q, Quinn PF (2014) Towards a hydrodynamic modelling framework appropriate for applications in urban flood assessment and mitigation using heterogeneous computing, *Urban Water Journal*, [in press].

Why bother?

…sensitivity

- Hypothetical breach in defences in Thamesmead, London
- High flow velocities and high sensitivity

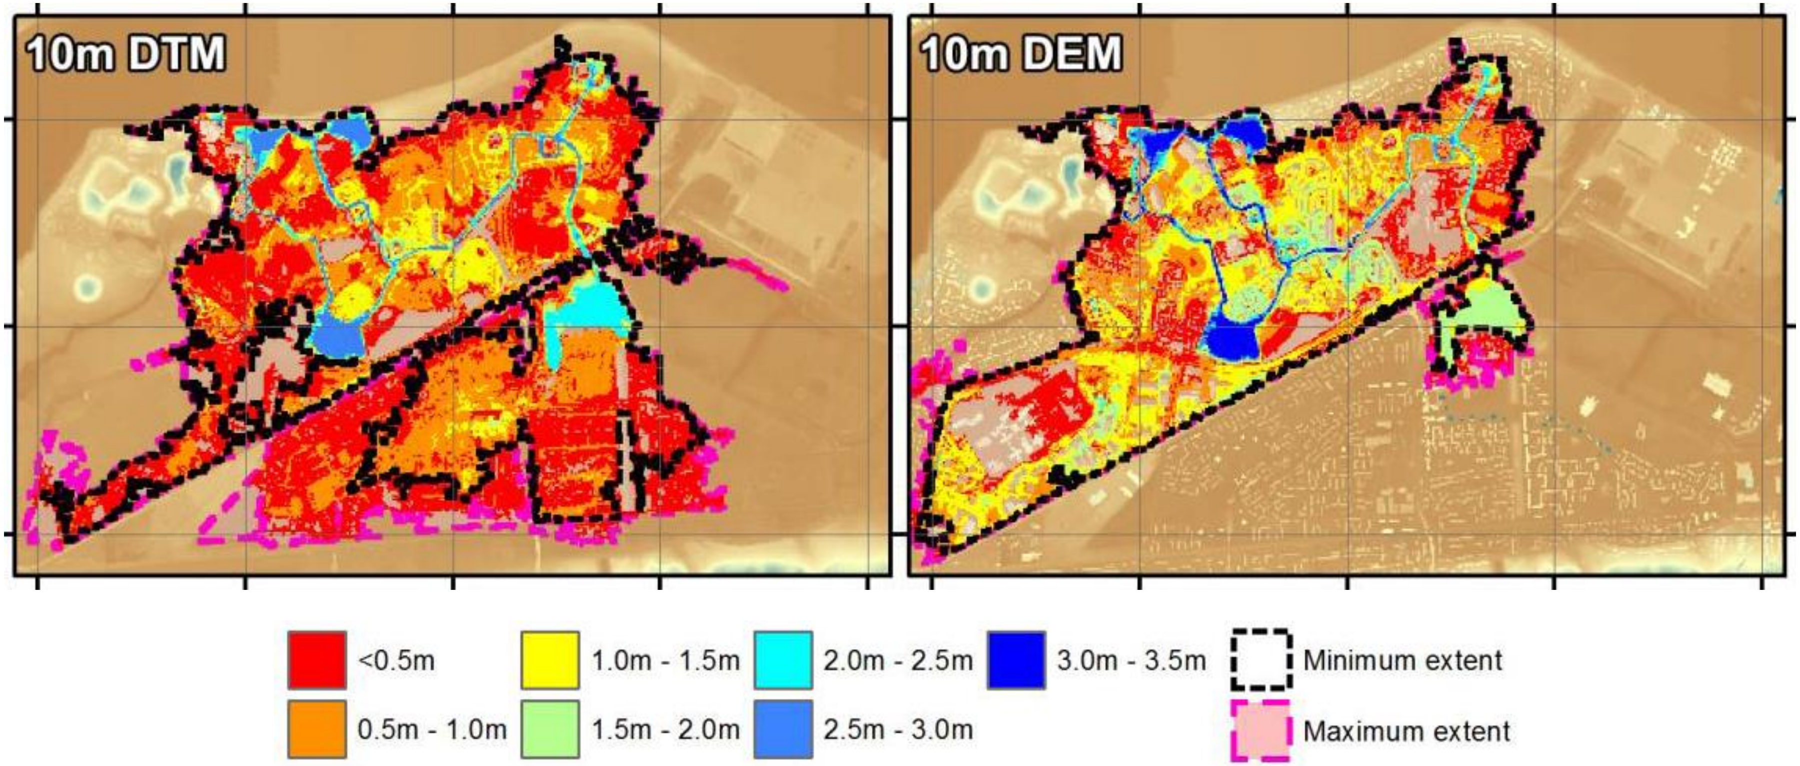

Liang Q, Smith LS (2014) A High-Performance Integrated Hydrodynamic Modelling System for Urban Flood Simulations, *Journal of Hydroinformatics*, [under review].

Why bother?

…sensitivity

- Hypothetical breach in defences in Thamesmead, London
- High flow velocities and high sensitivity

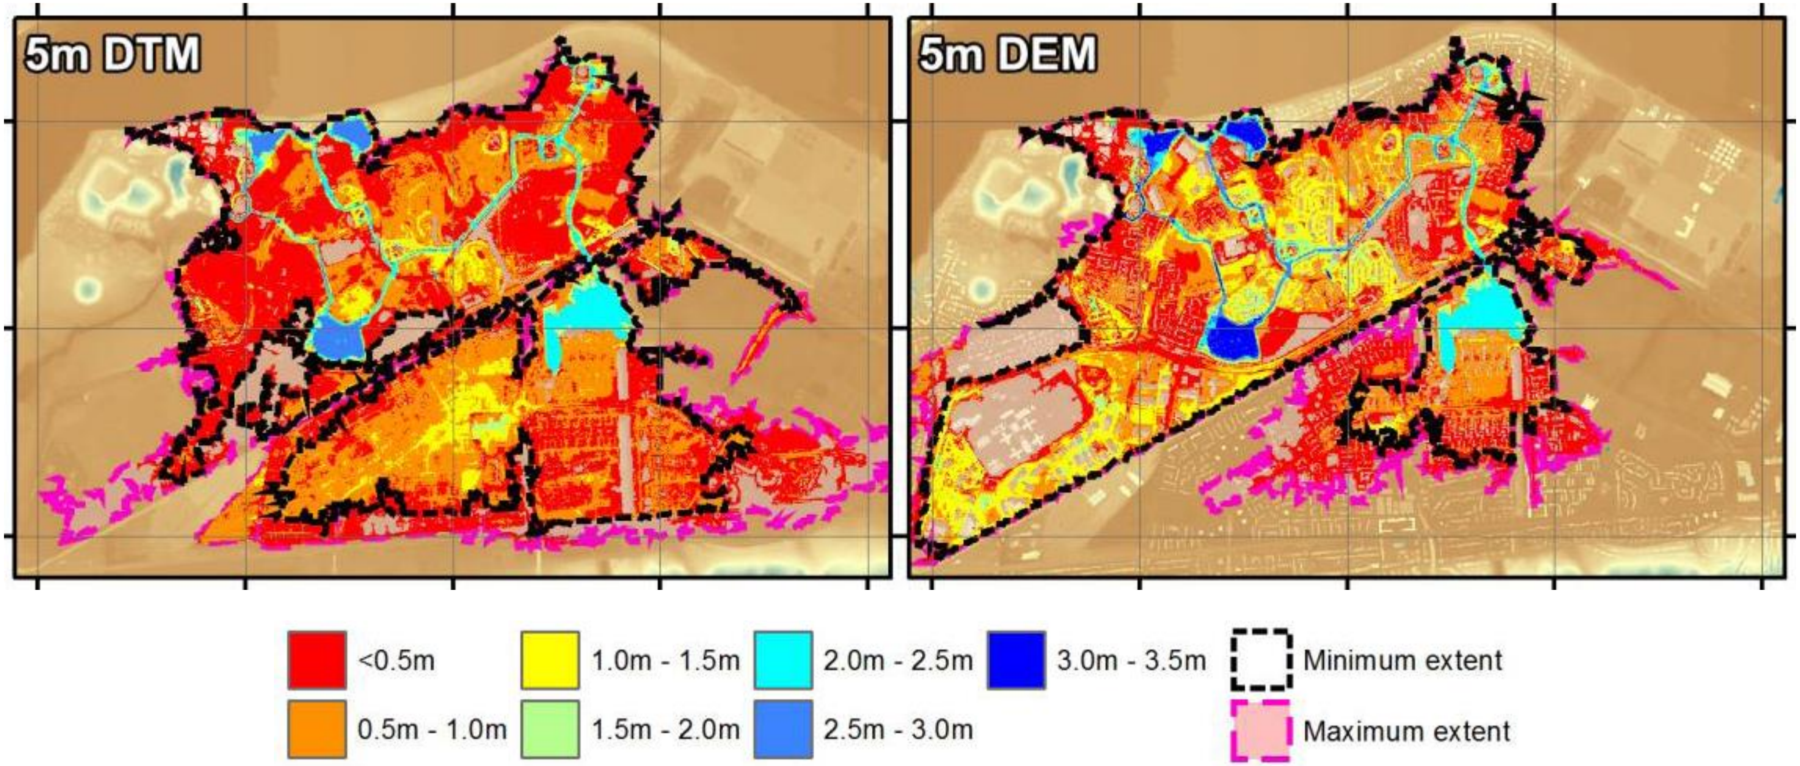

Liang Q, Smith LS (2014) A High-Performance Integrated Hydrodynamic Modelling System for Urban Flood Simulations, *Journal of Hydroinformatics*, [under review].

Why bother?

…sensitivity

- Hypothetical breach in defences in Thamesmead, London
- High flow velocities and high sensitivity

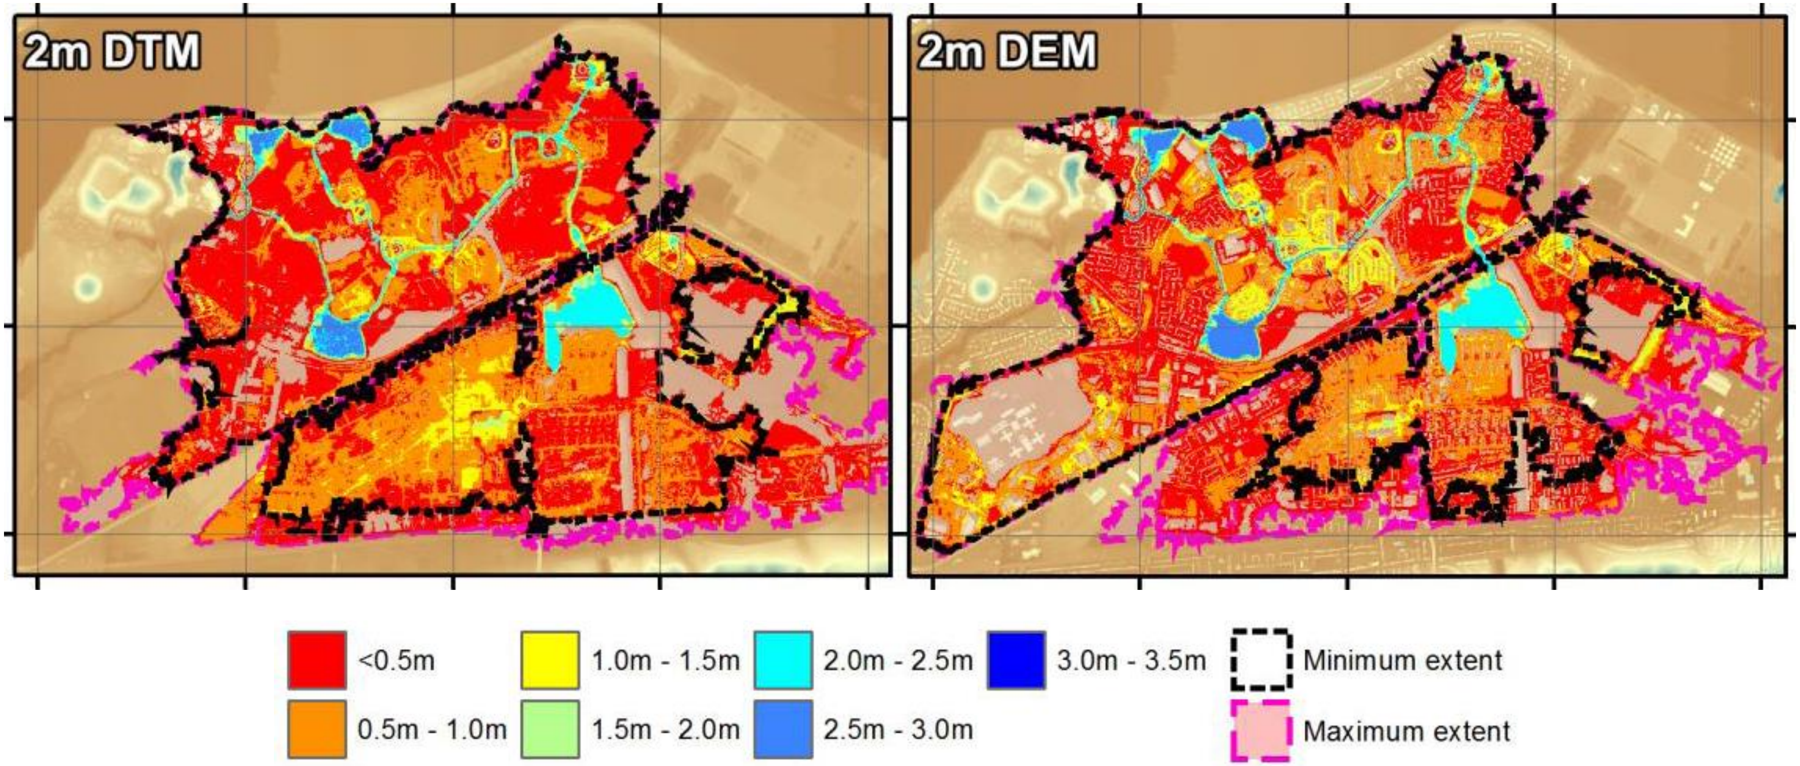

Liang Q, Smith LS (2014) A High-Performance Integrated Hydrodynamic Modelling System for Urban Flood Simulations, *Journal of Hydroinformatics*, [under review].

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 12:45 UTC

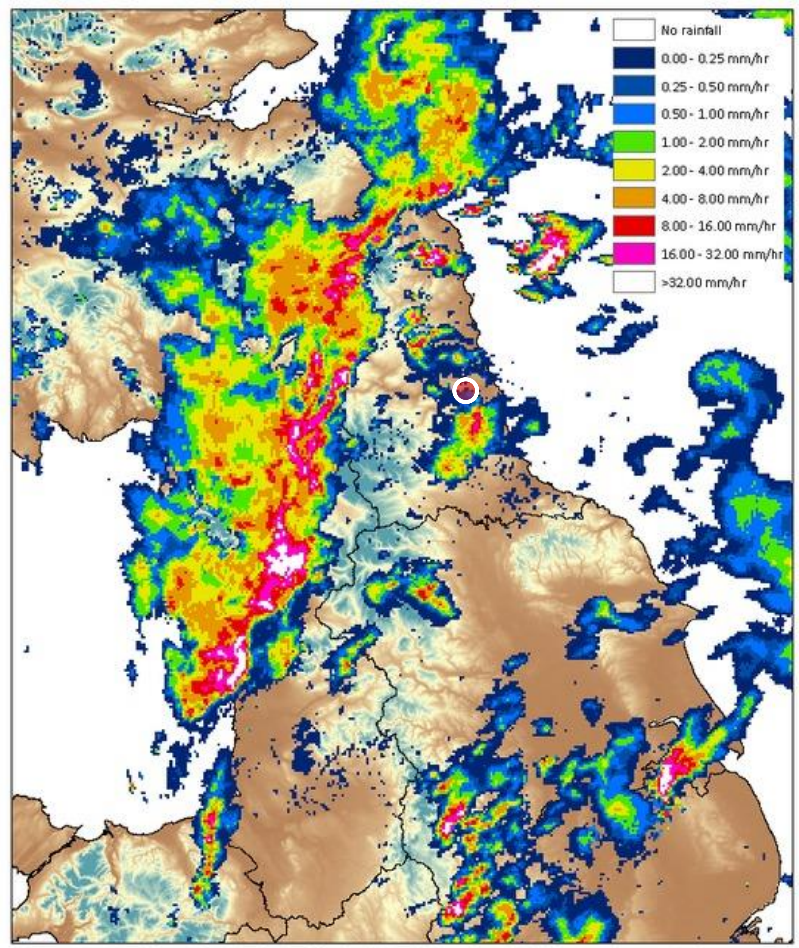

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 13:35 UTC

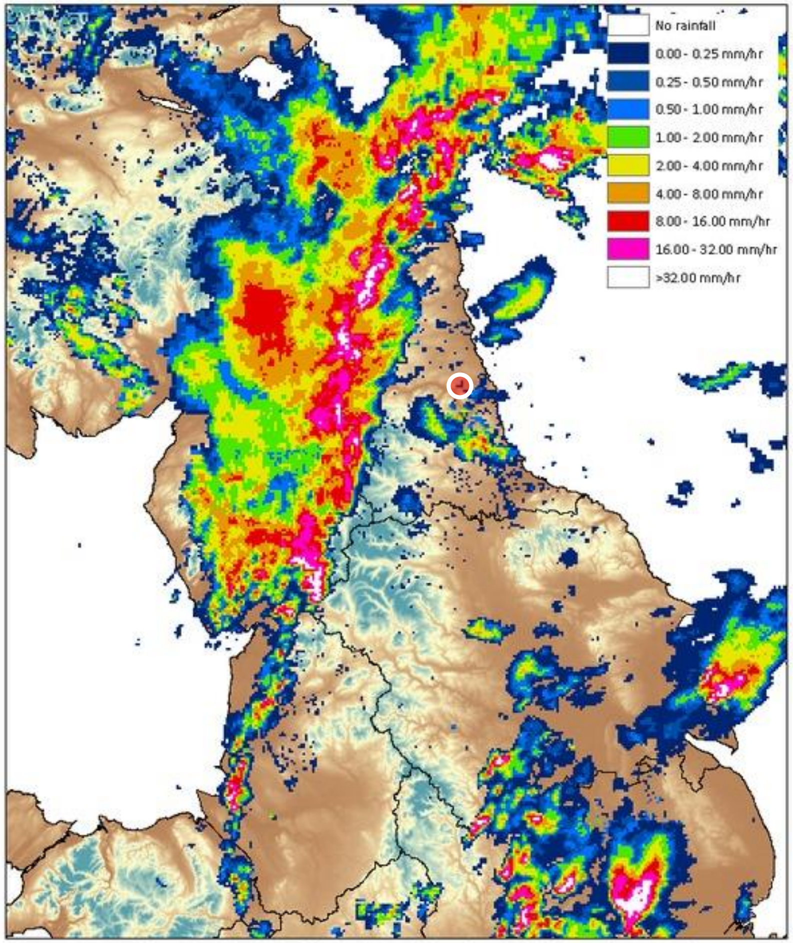

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

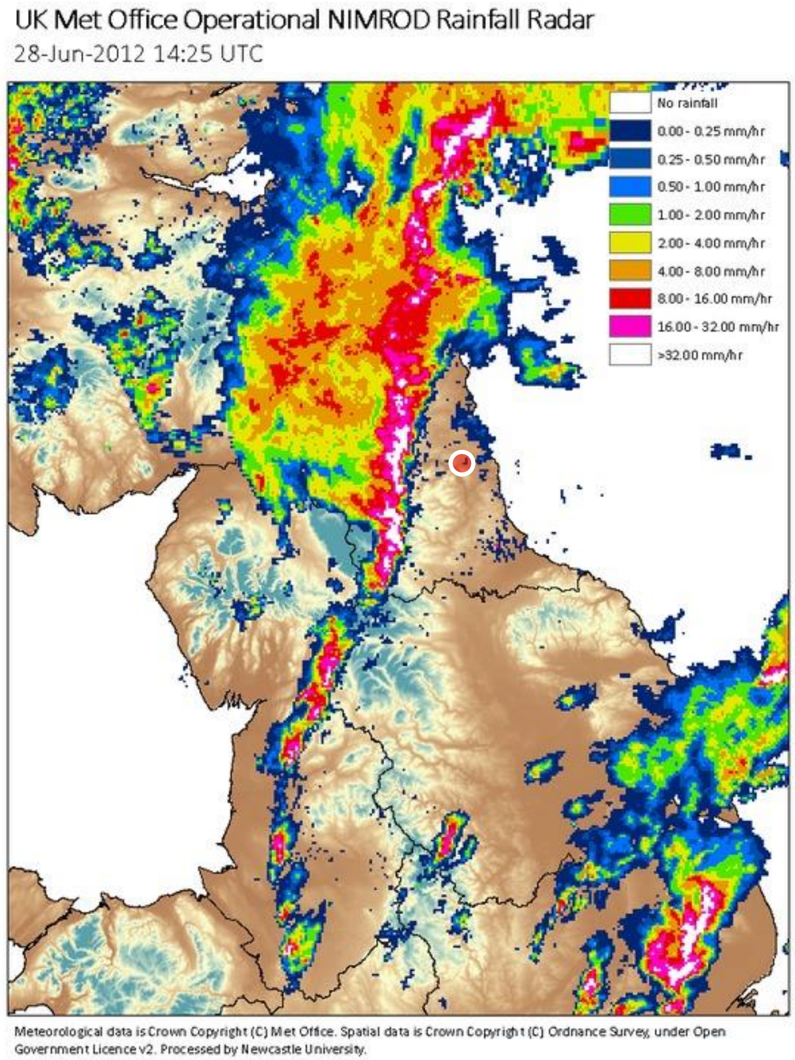

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

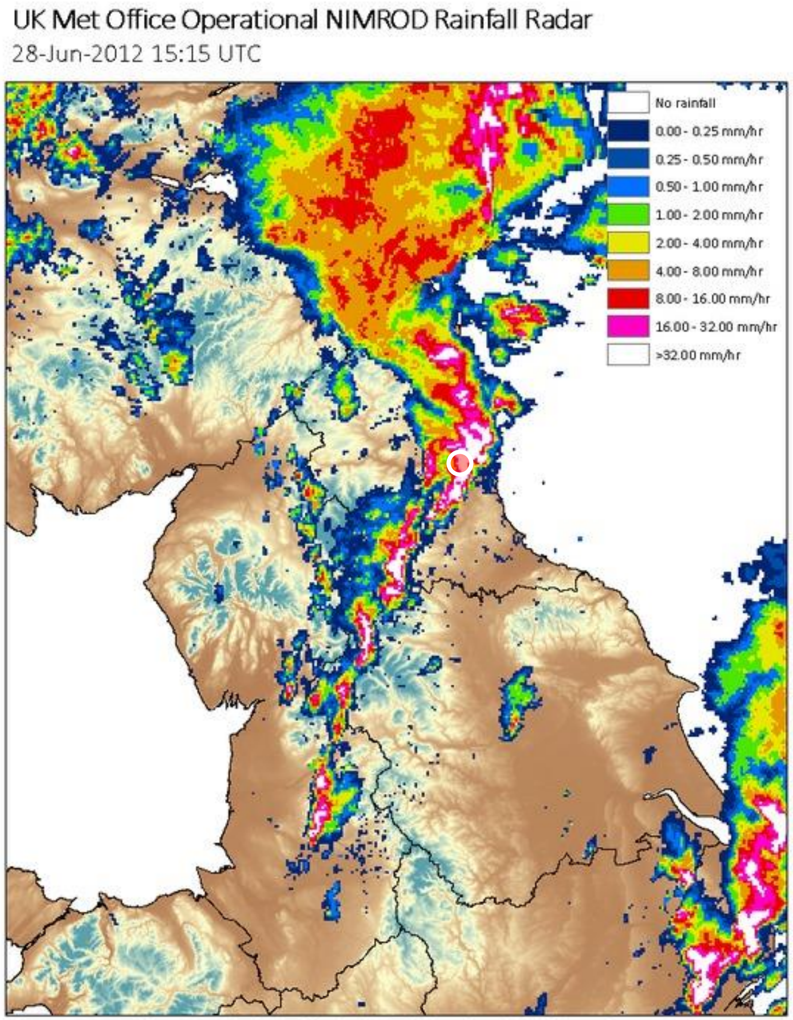

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
  - Strategic transport routes
  - Blue light routes
  - Light rail services
  - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 16:05 UTC

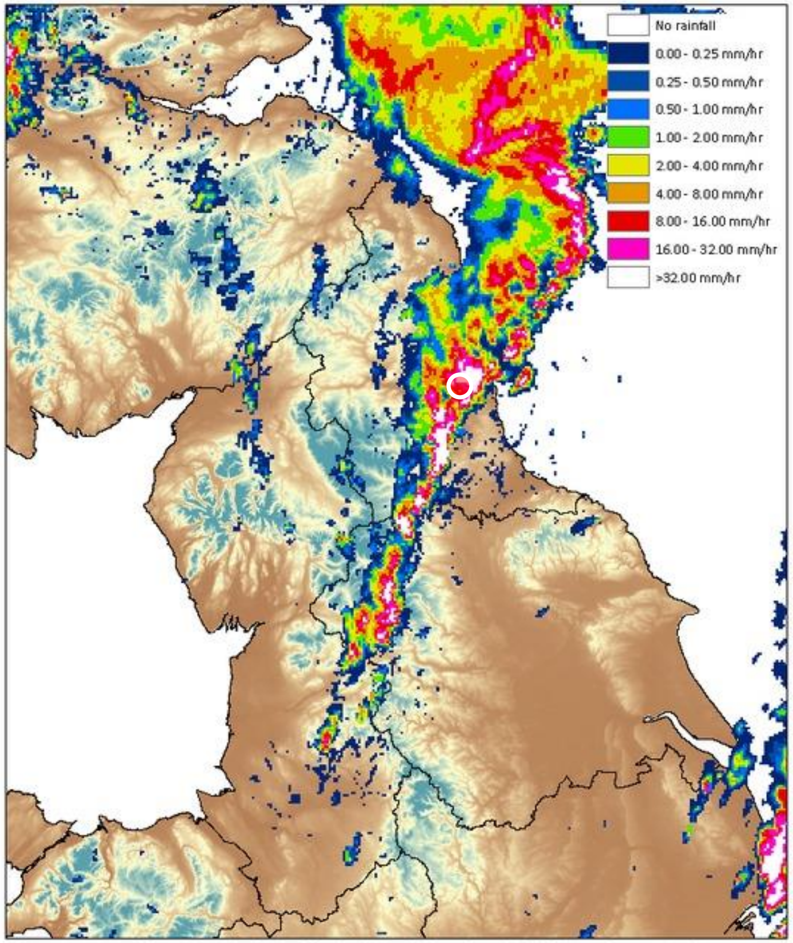

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 16:55 UTC

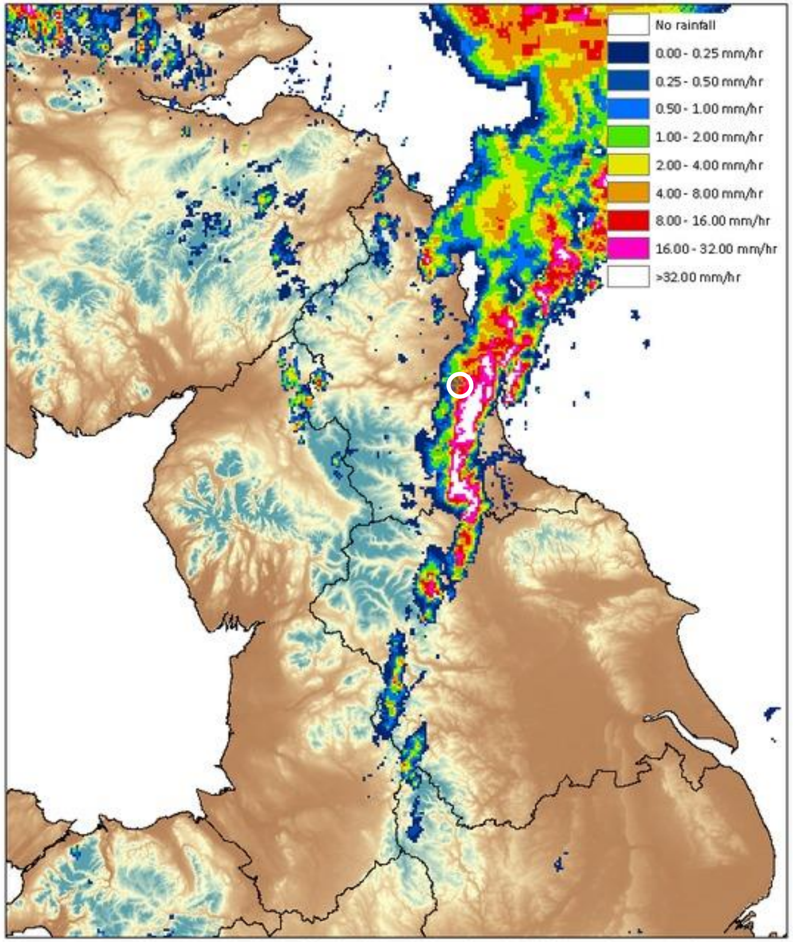

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 17:45 UTC

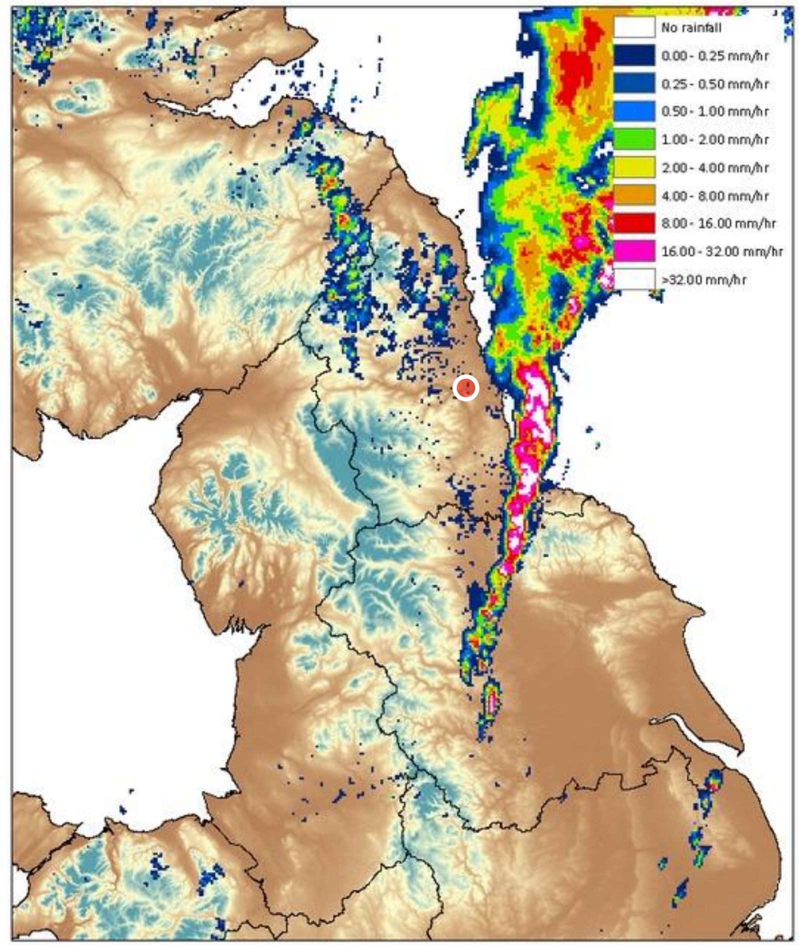

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Pluvial flooding in Newcastle

- 15:55 on 28 June 2012
- **50mm** widely in 2 hours
- Severe transport disruption
    - Strategic transport routes
    - Blue light routes
    - Light rail services
    - National rail services

UK Met Office Operational NIMROD Rainfall Radar
28-Jun-2012 18:00 UTC

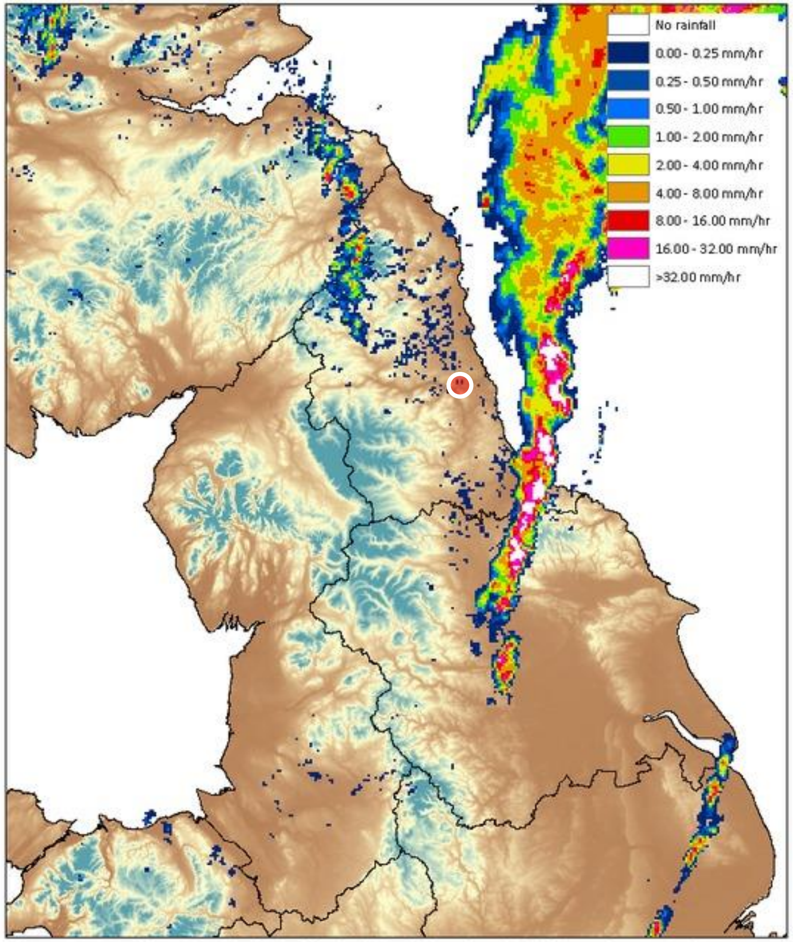

Meteorological data is Crown Copyright (C) Met Office. Spatial data is Crown Copyright (C) Ordnance Survey, under Open Government Licence v2. Processed by Newcastle University.

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 15:15 UTC

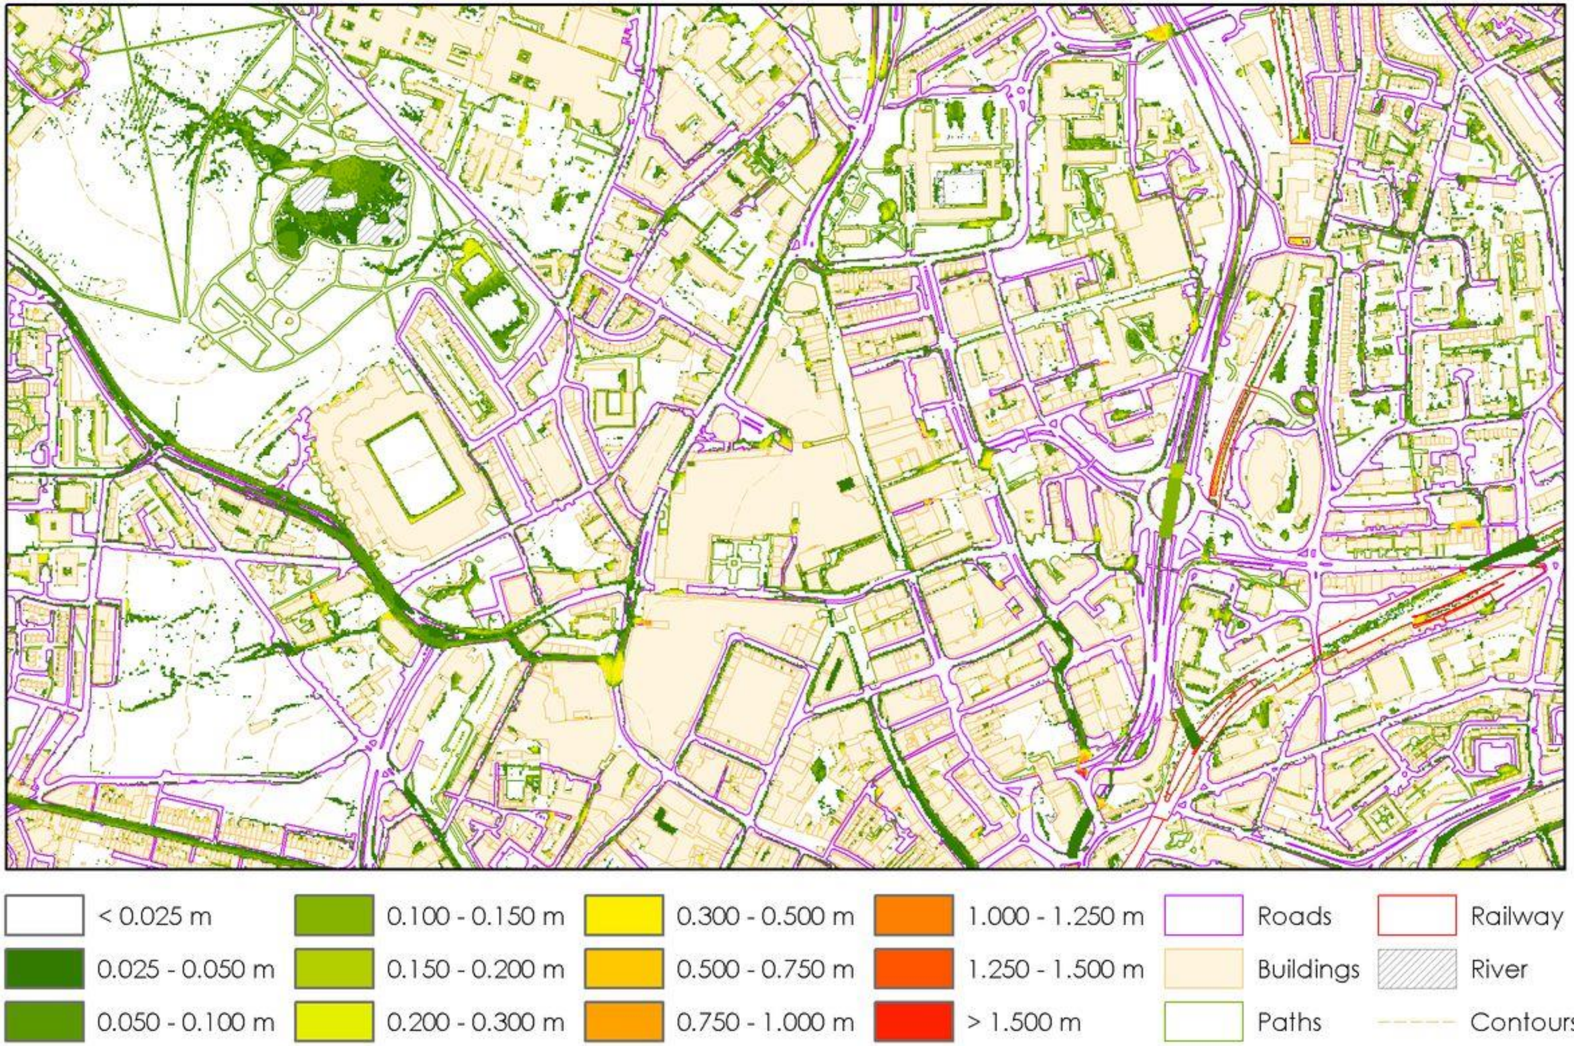

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 15:30 UTC

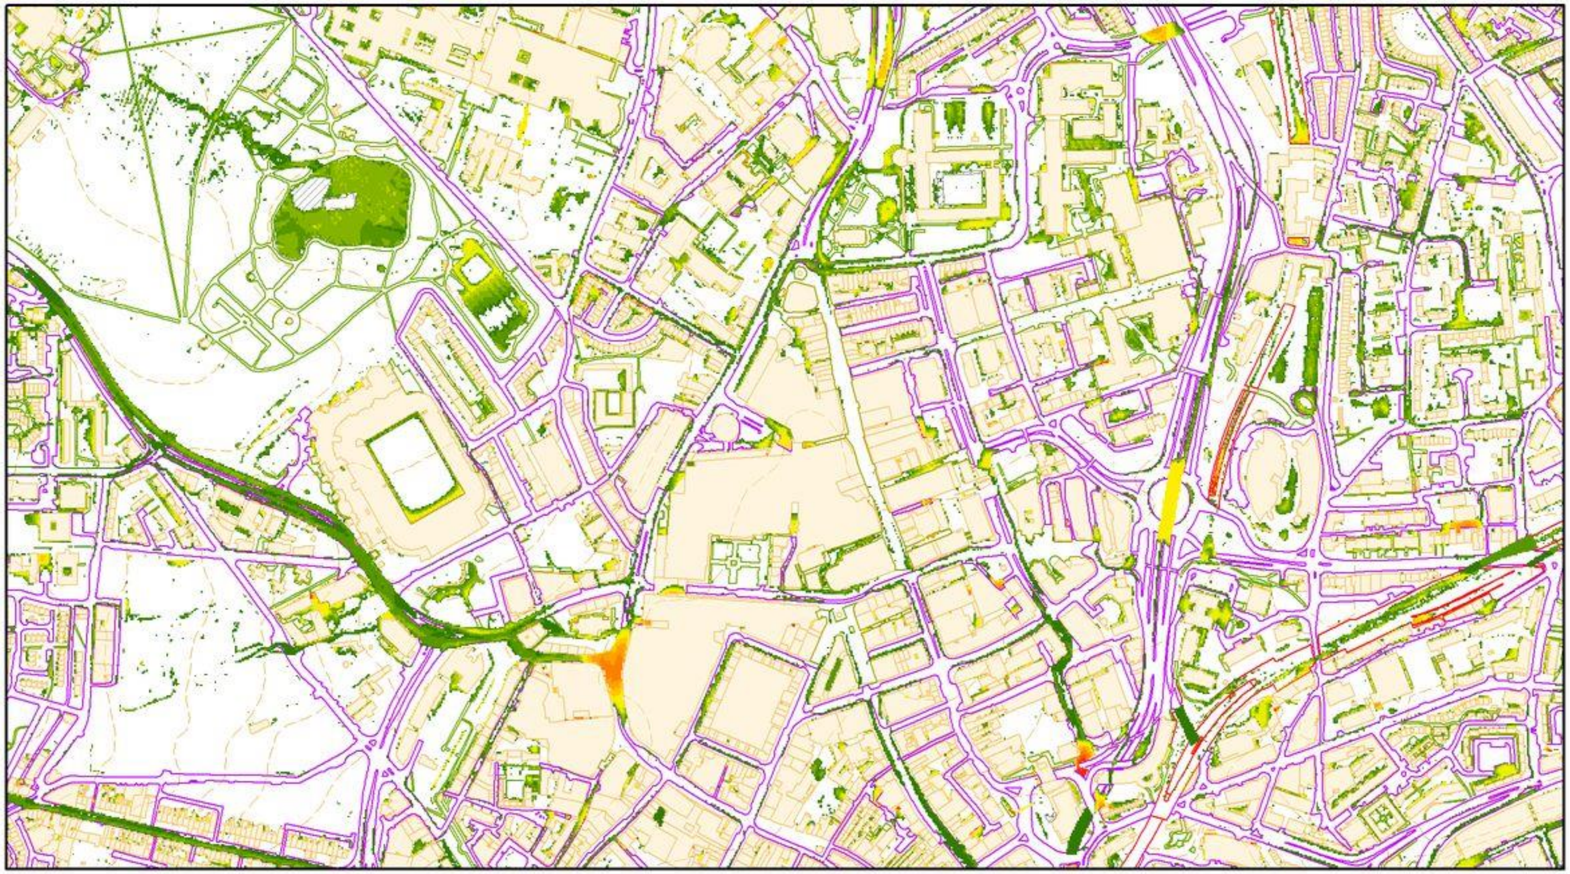

< 0.025 m	0.100 - 0.150 m	0.300 - 0.500 m	1.000 - 1.250 m	Roads	Railway
0.025 - 0.050 m	0.150 - 0.200 m	0.500 - 0.750 m	1.250 - 1.500 m	Buildings	River
0.050 - 0.100 m	0.200 - 0.300 m	0.750 - 1.000 m	> 1.500 m	Paths	Contours

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 15:45 UTC

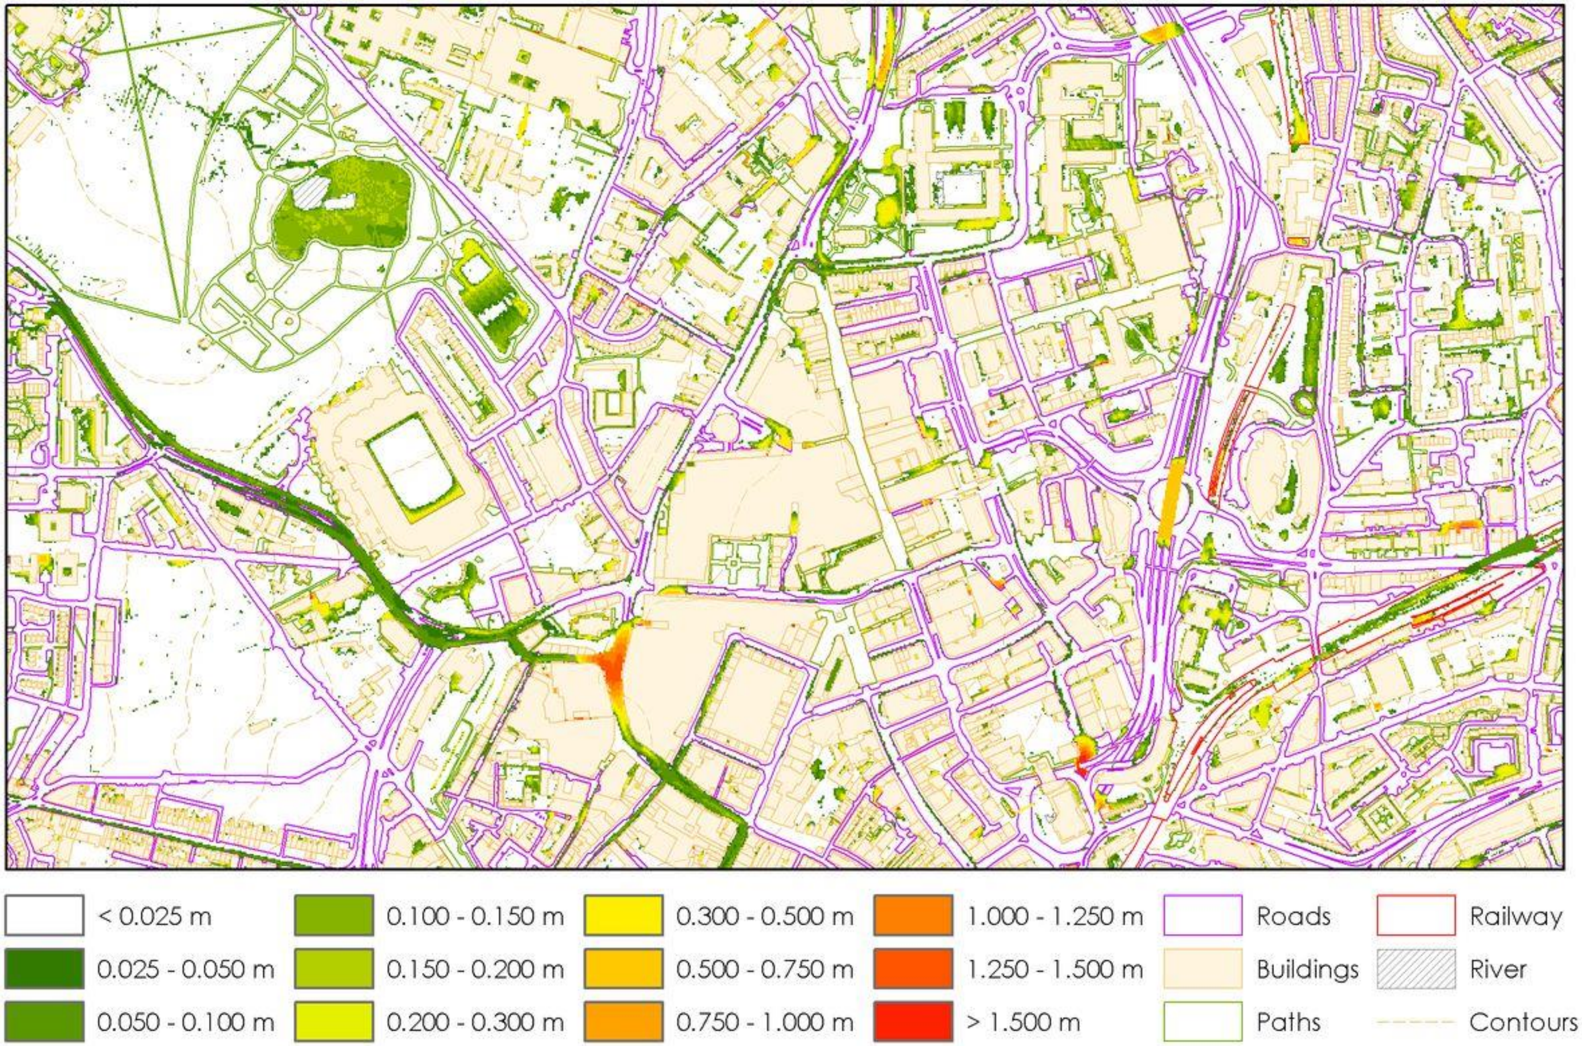

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 16:00 UTC

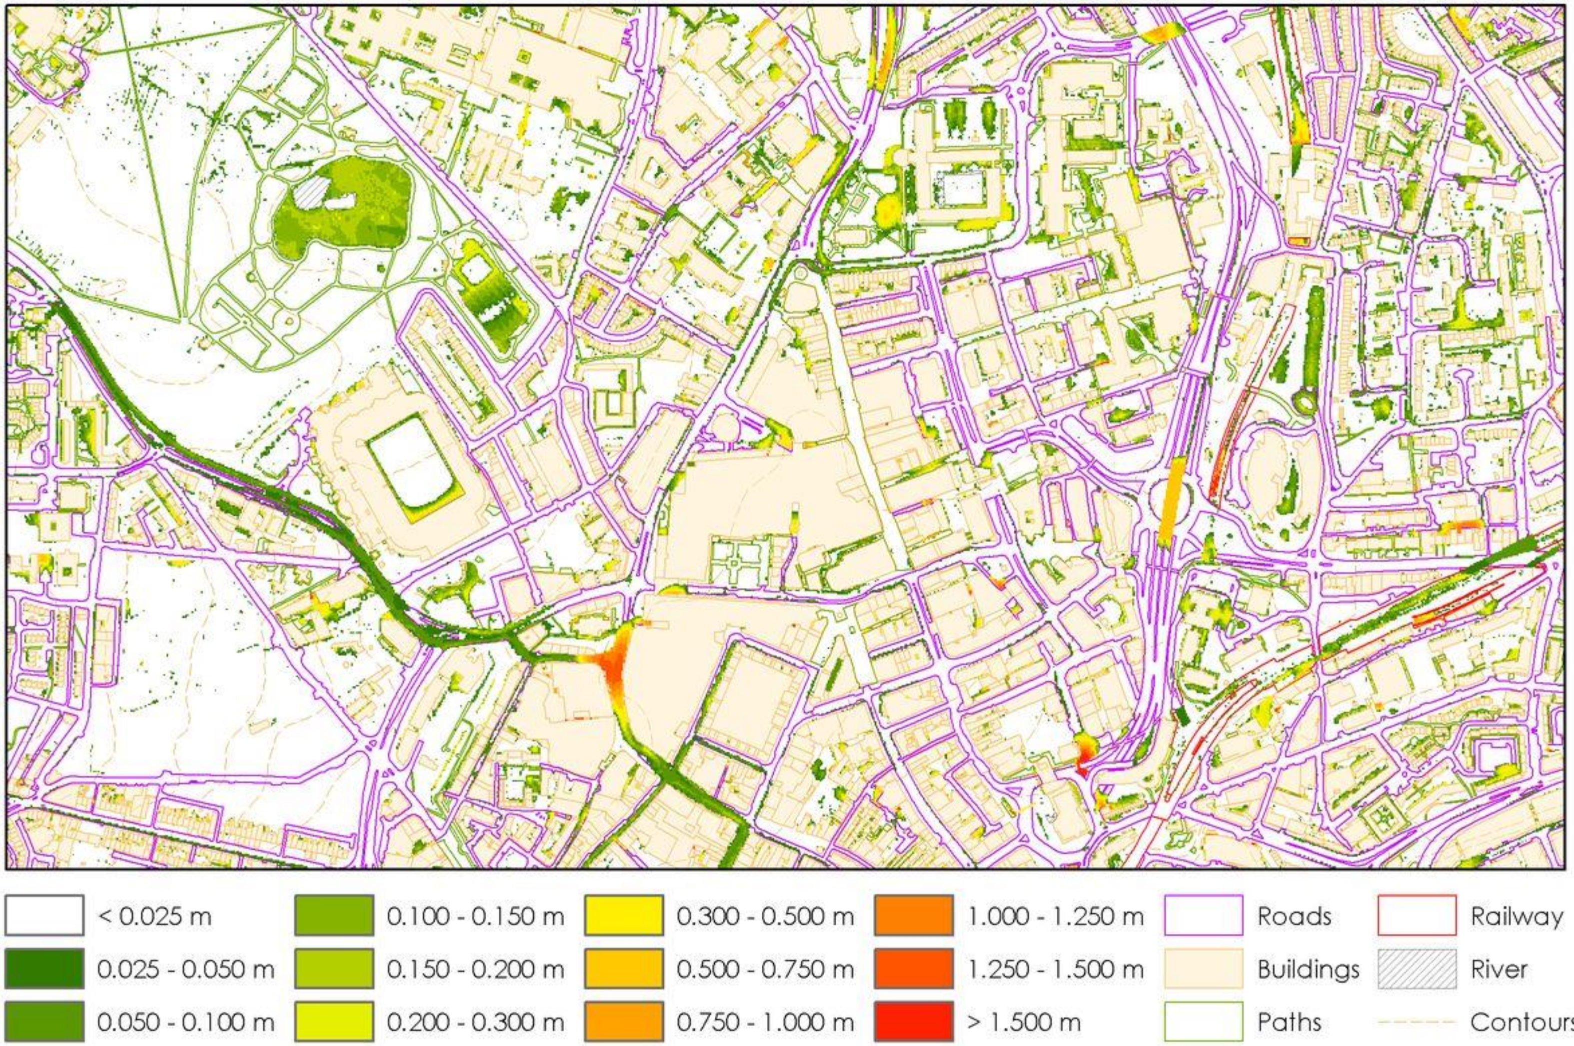

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 16:15 UTC

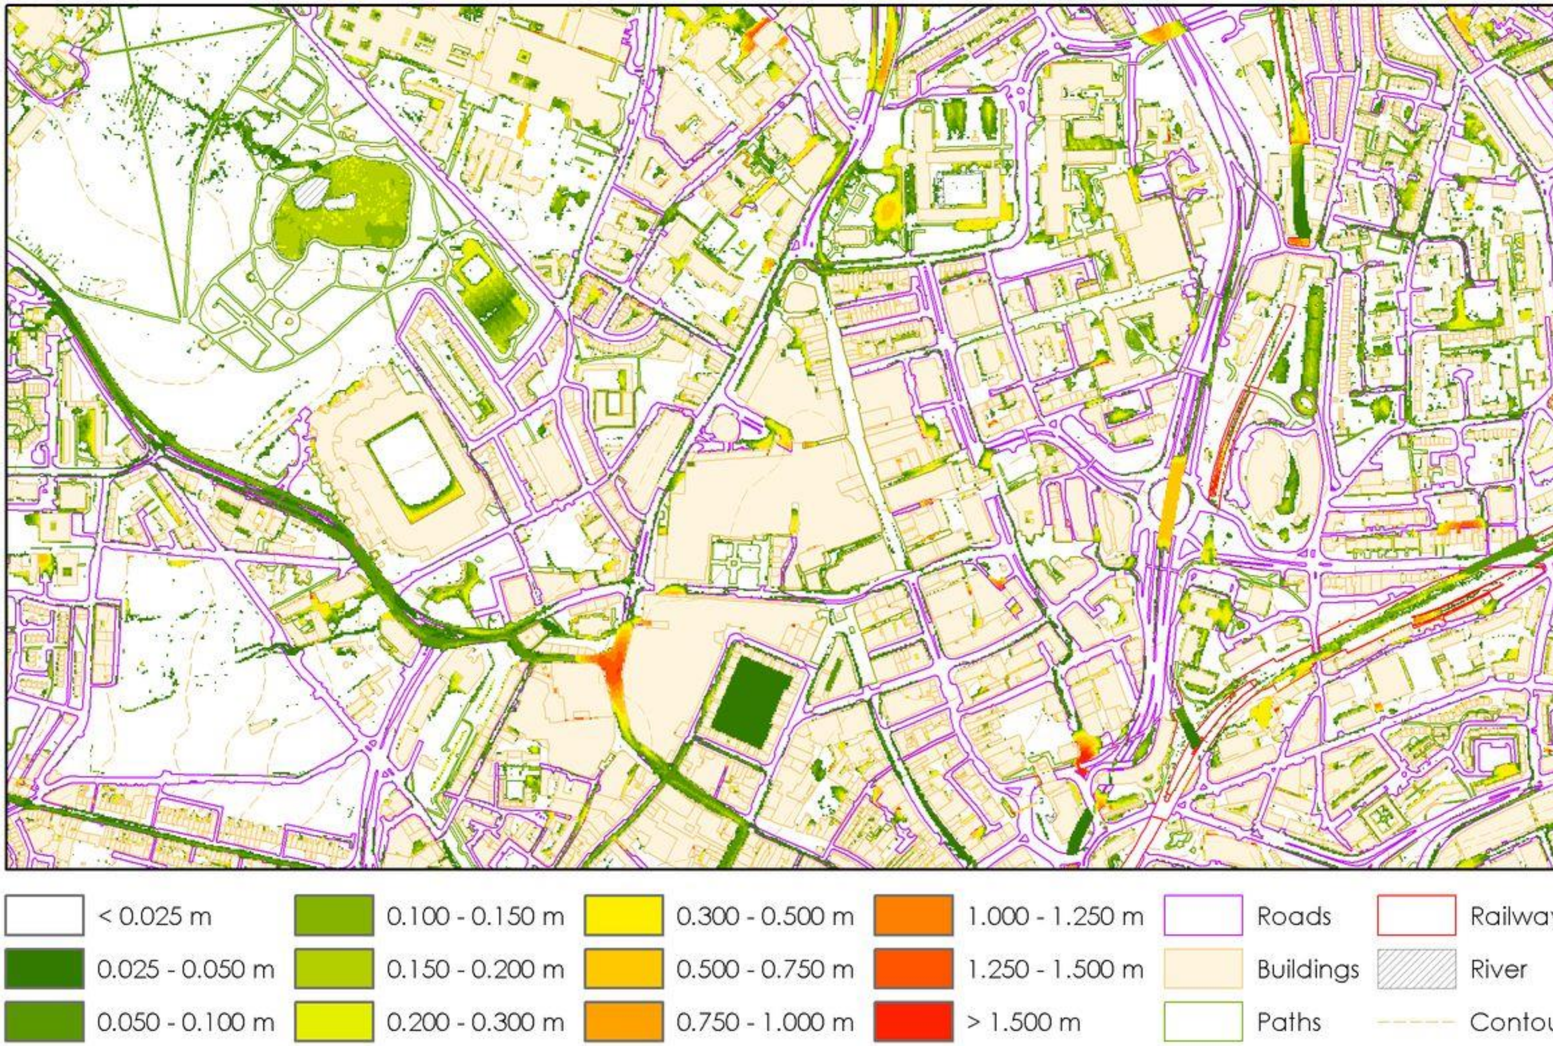

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 16:30 UTC

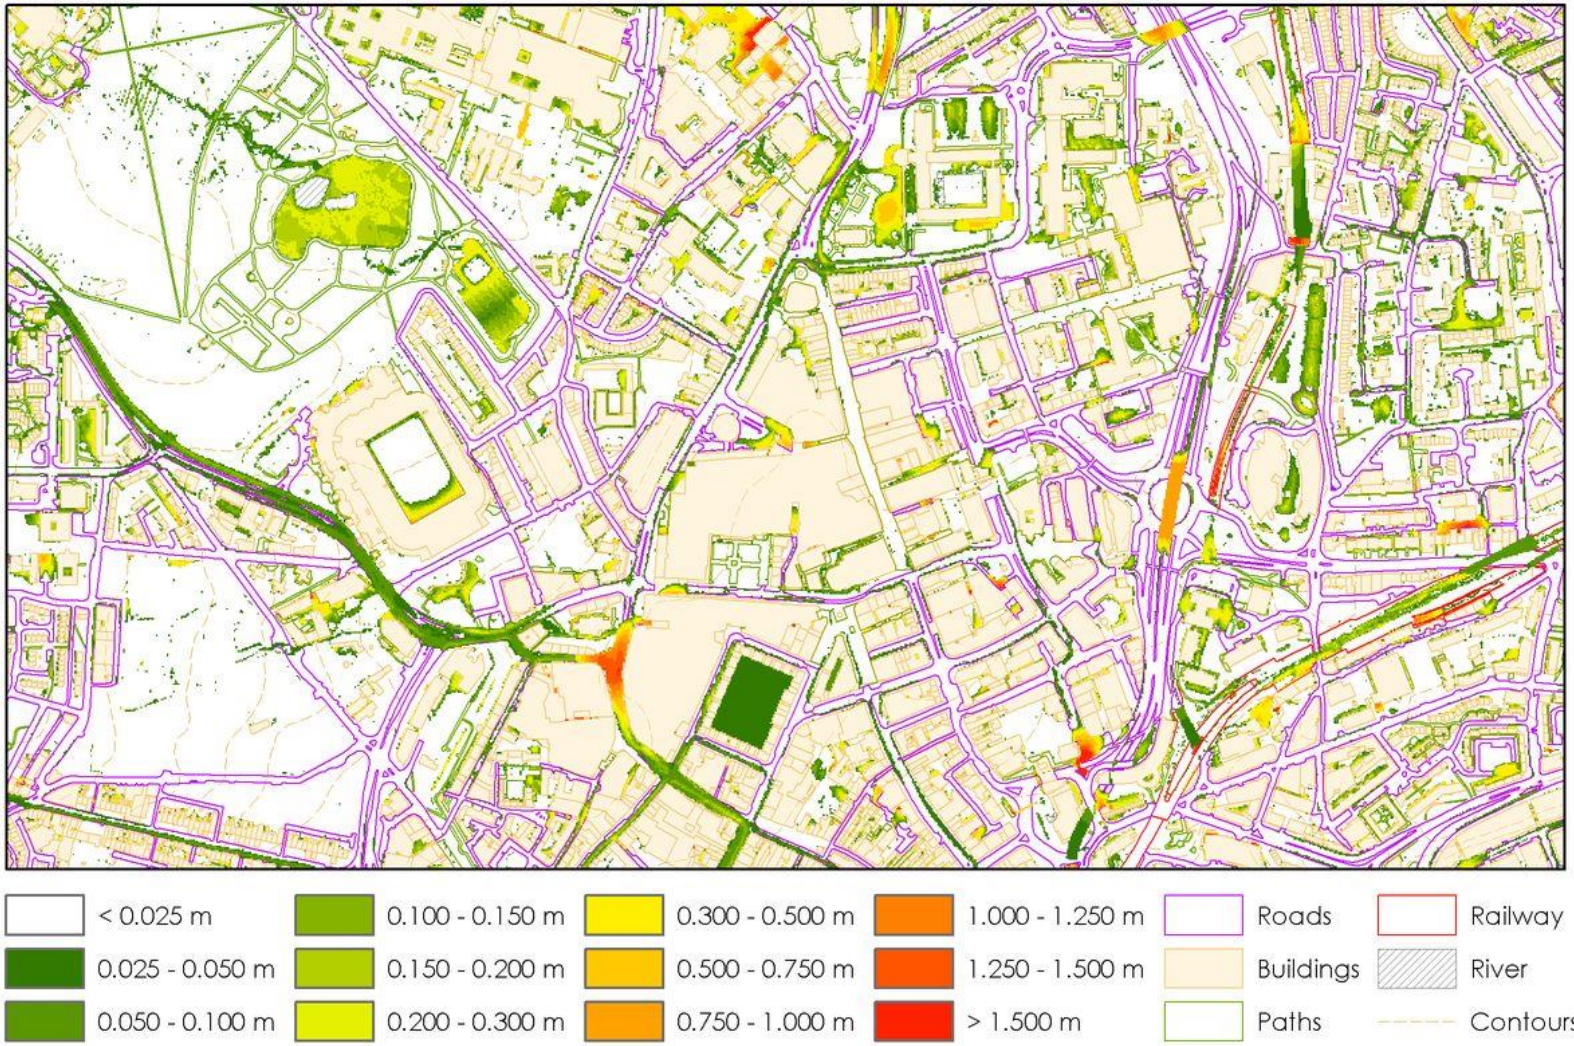

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 16:45 UTC

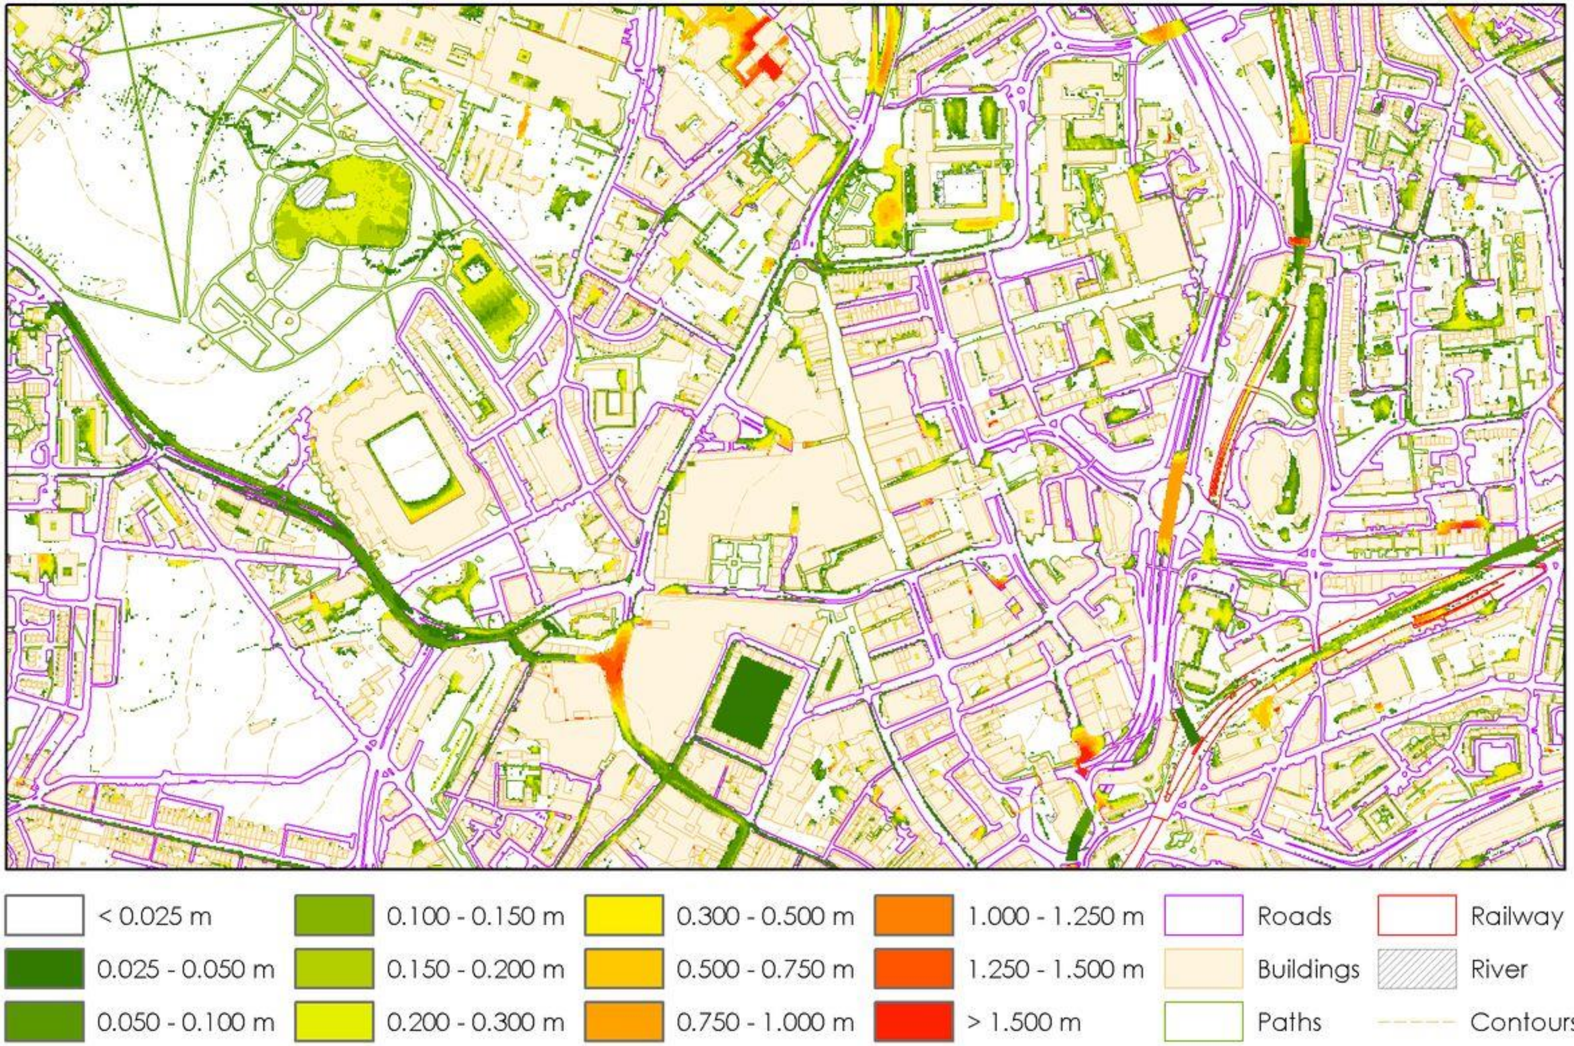

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:00 UTC

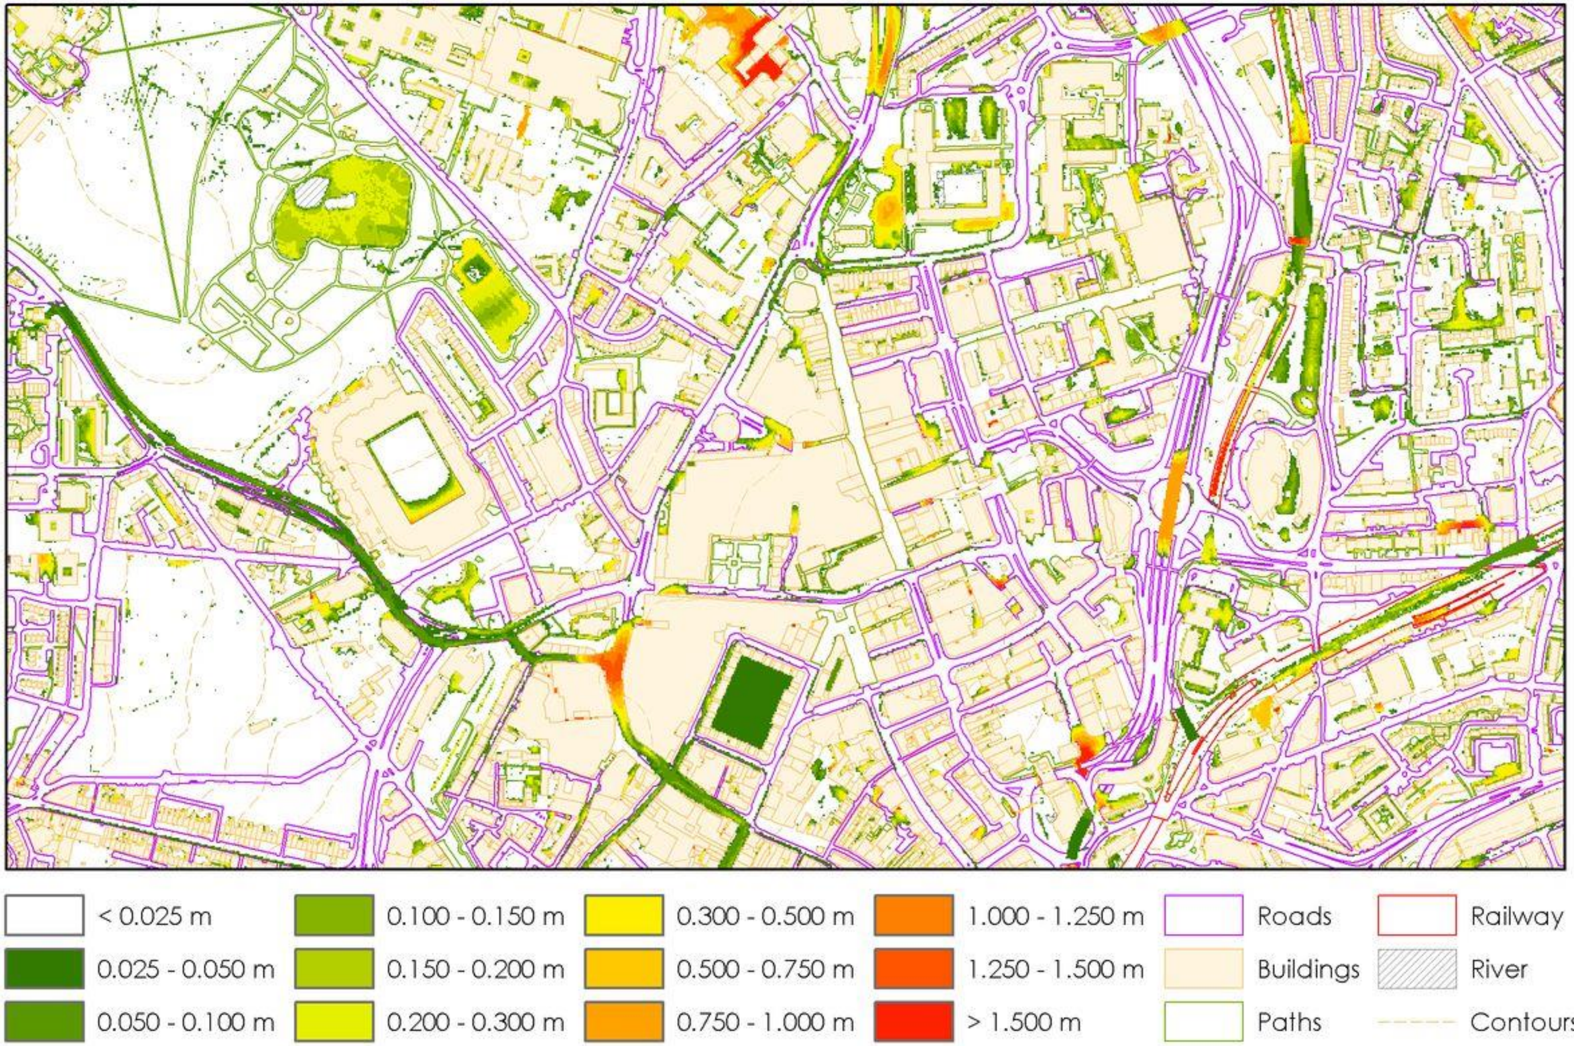

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:15 UTC

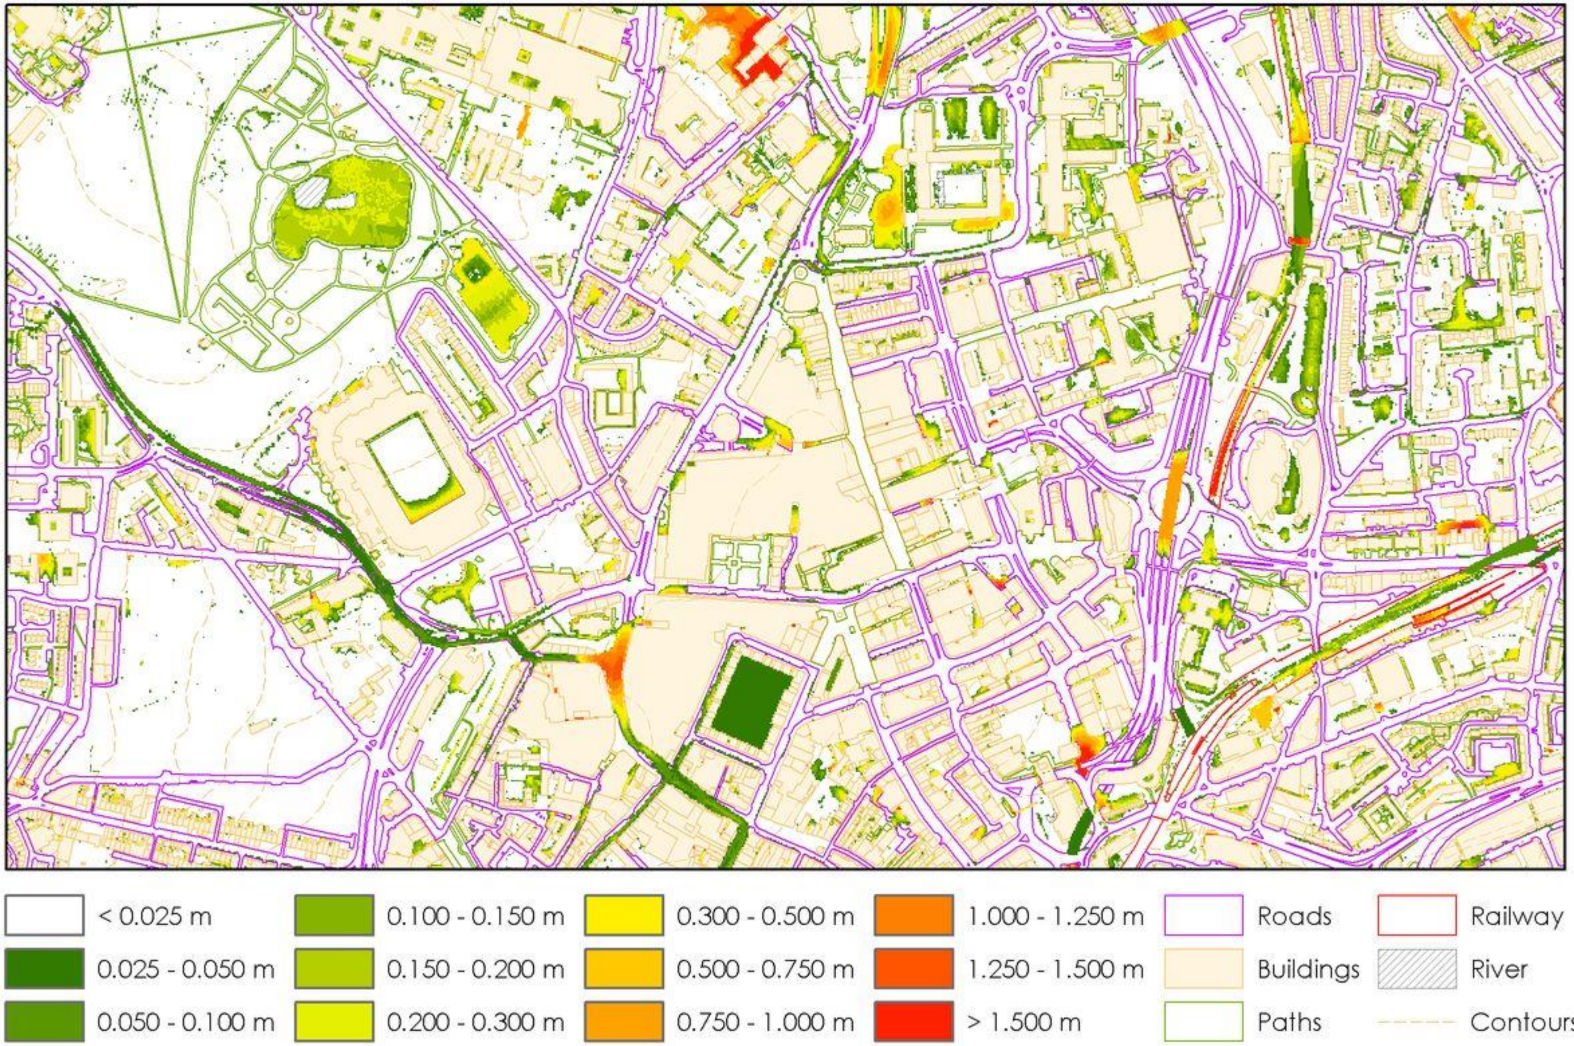

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:30 UTC

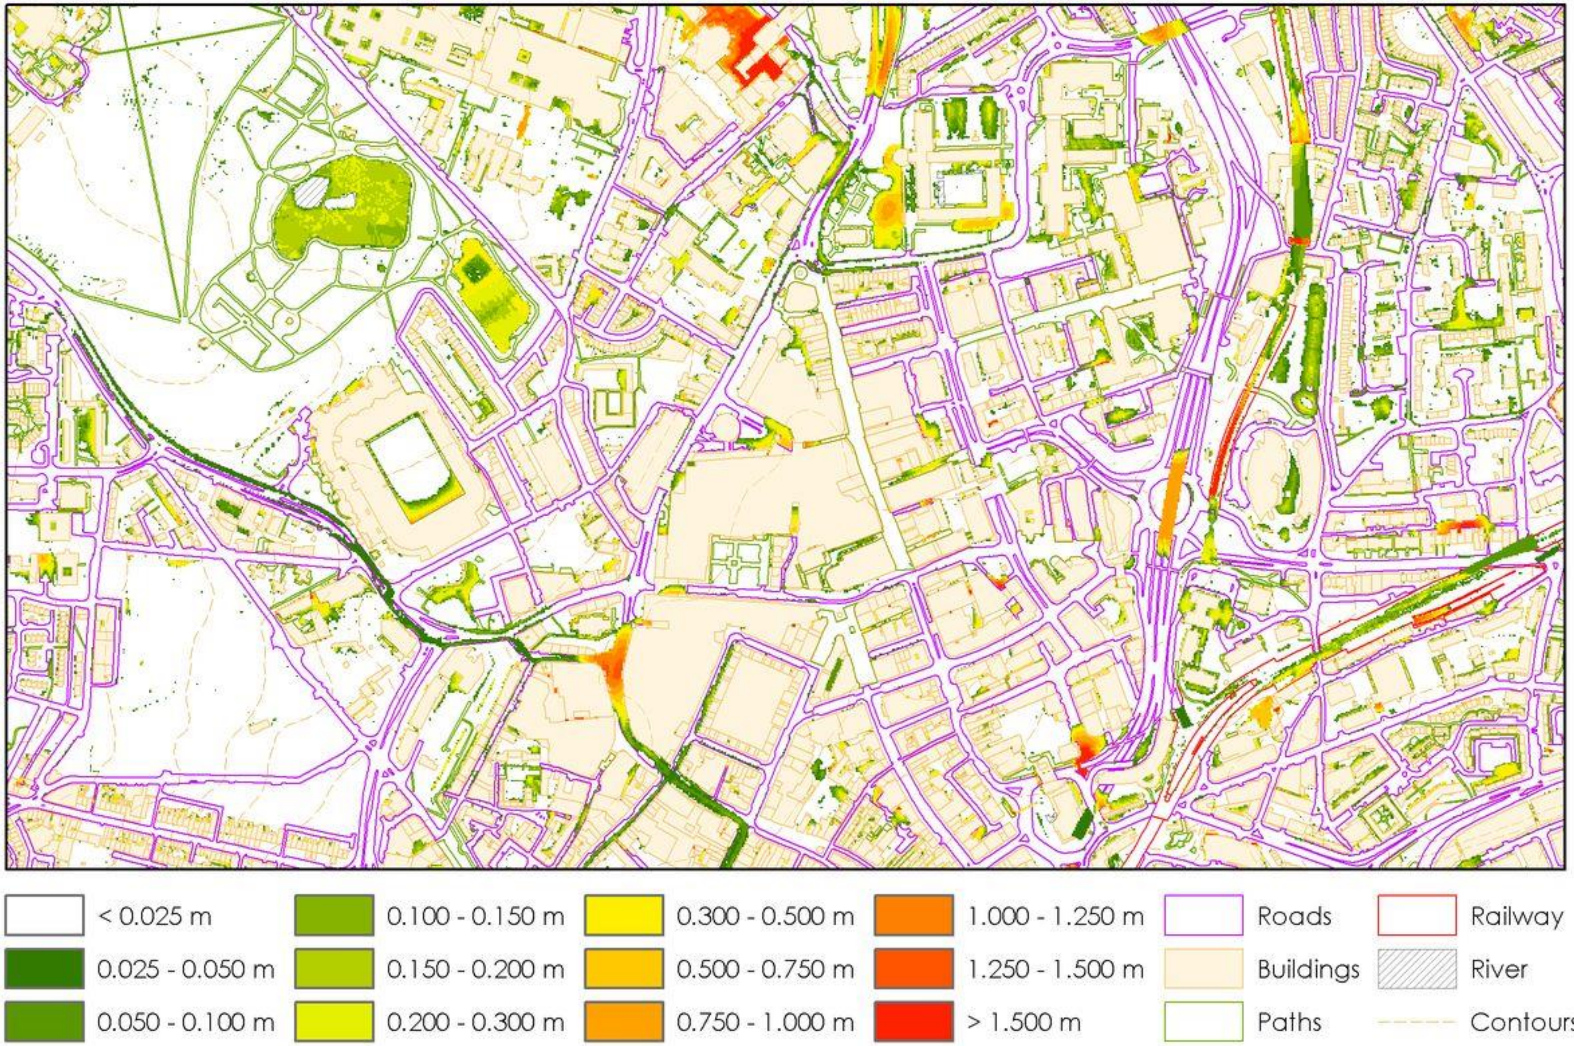

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:45 UTC

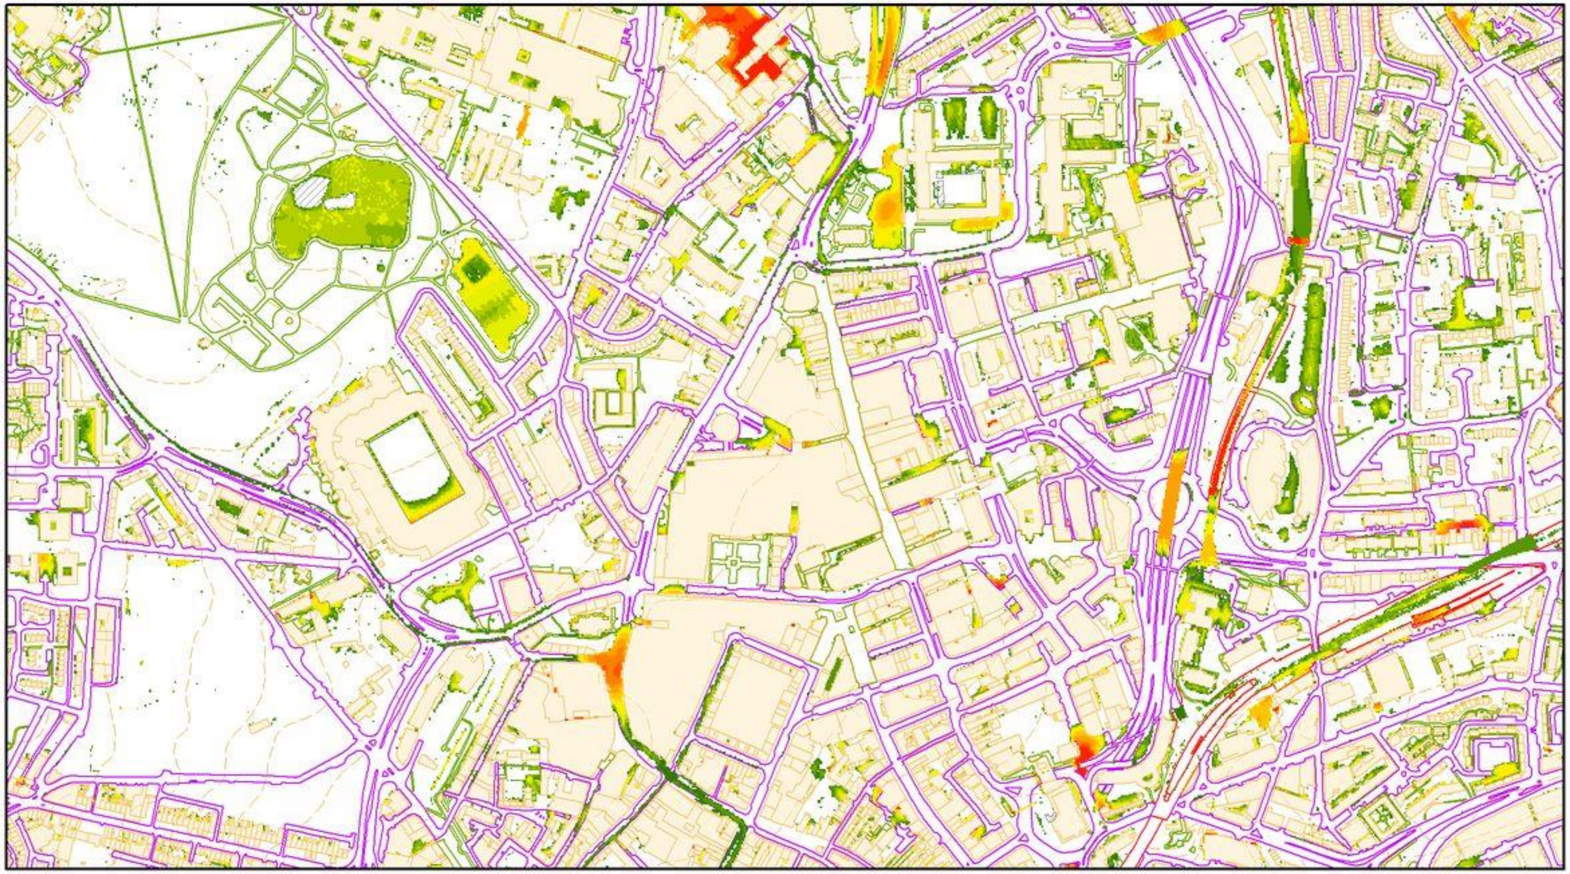

< 0.025 m	0.100 - 0.150 m	0.300 - 0.500 m	1.000 - 1.250 m	Roads	Railway
0.025 - 0.050 m	0.150 - 0.200 m	0.500 - 0.750 m	1.250 - 1.500 m	Buildings	River
0.050 - 0.100 m	0.200 - 0.300 m	0.750 - 1.000 m	> 1.500 m	Paths	Contours

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:50 UTC

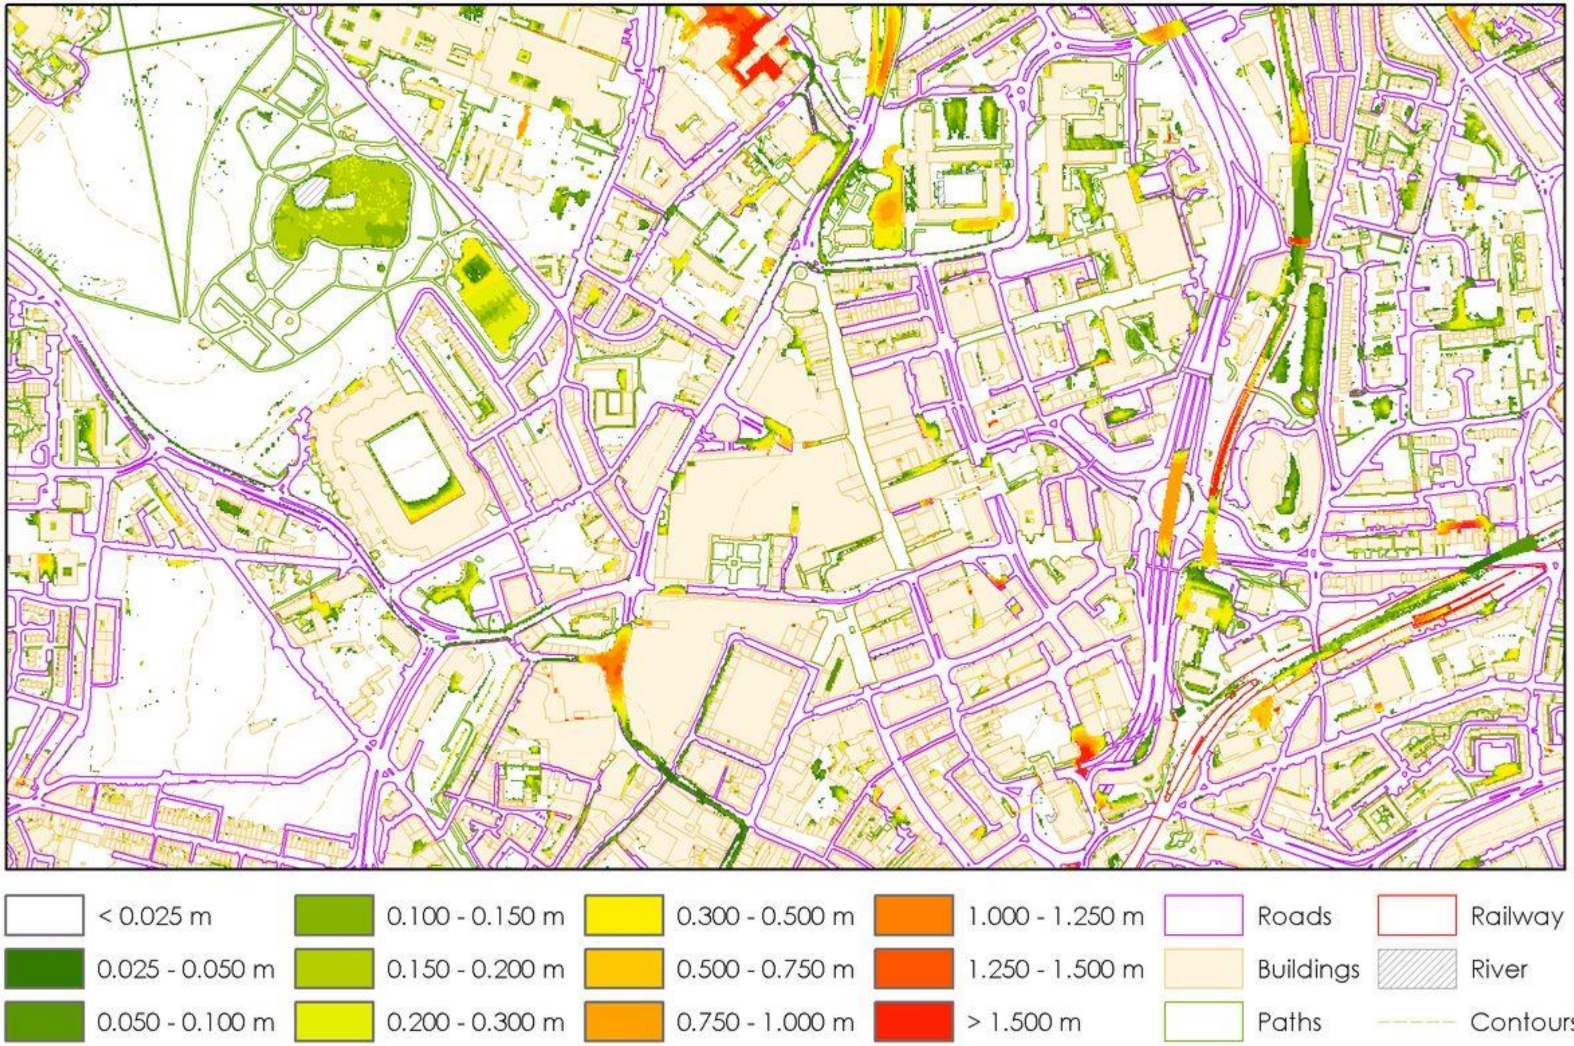

Newcastle upon Tyne Surface Water Model (Input from UKMO NIMROD)

28-Jun-2012 17:55 UTC

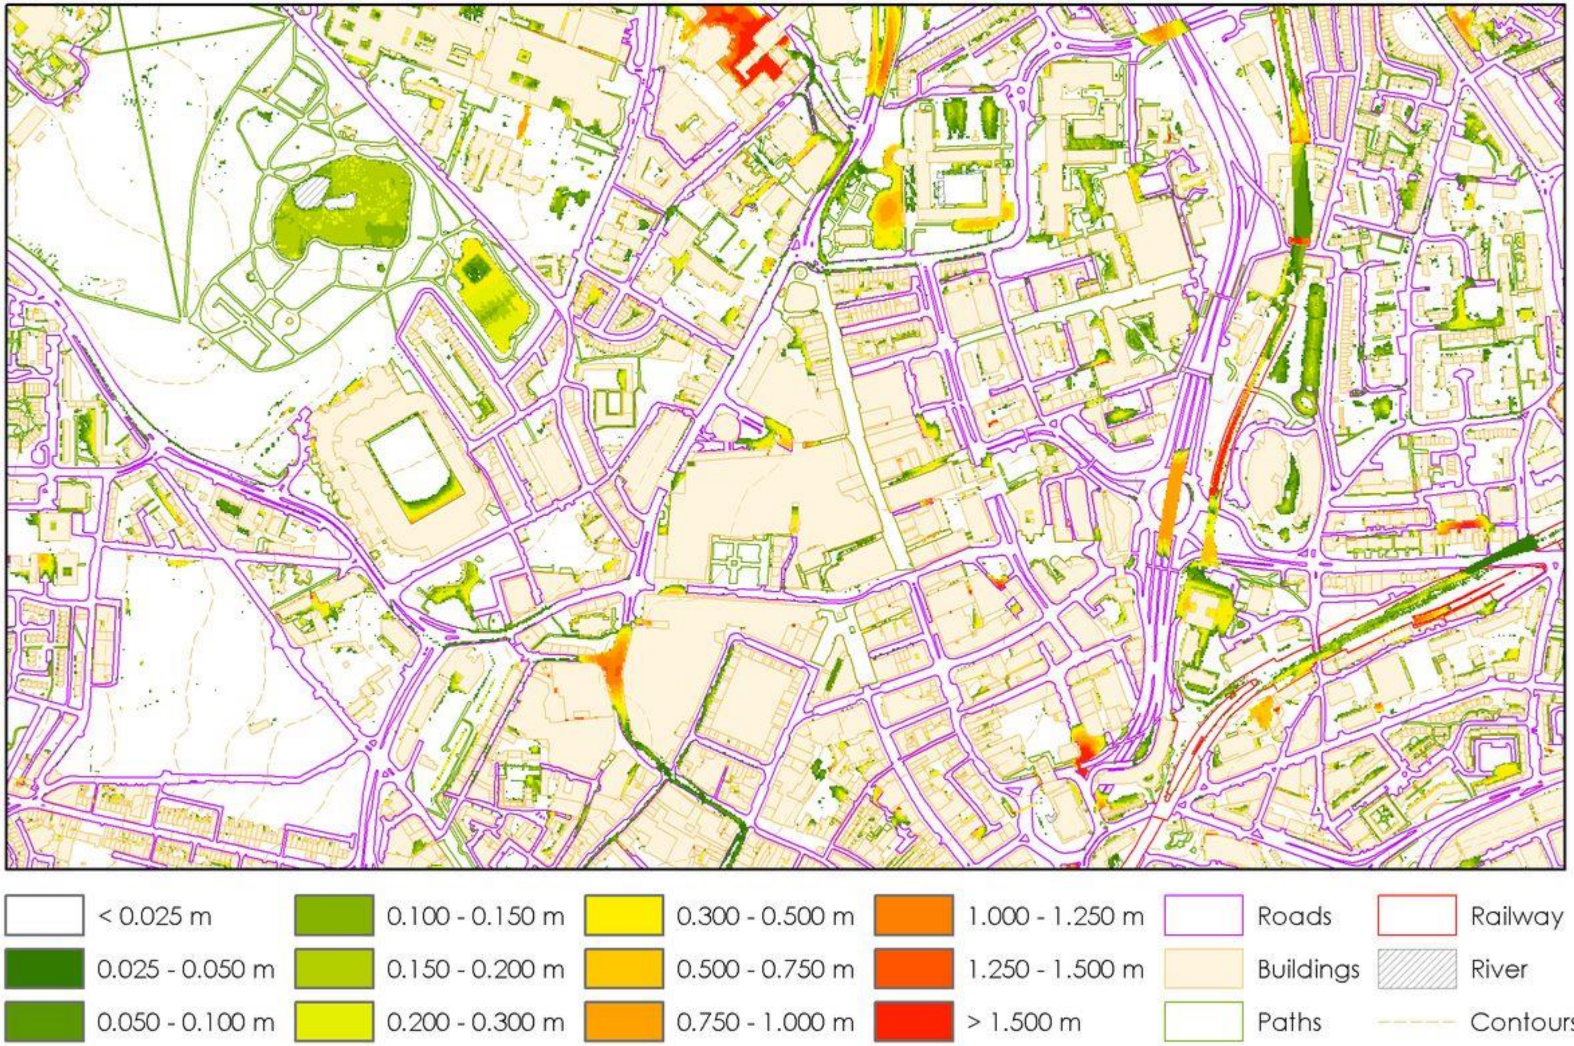

Modelling versus reality

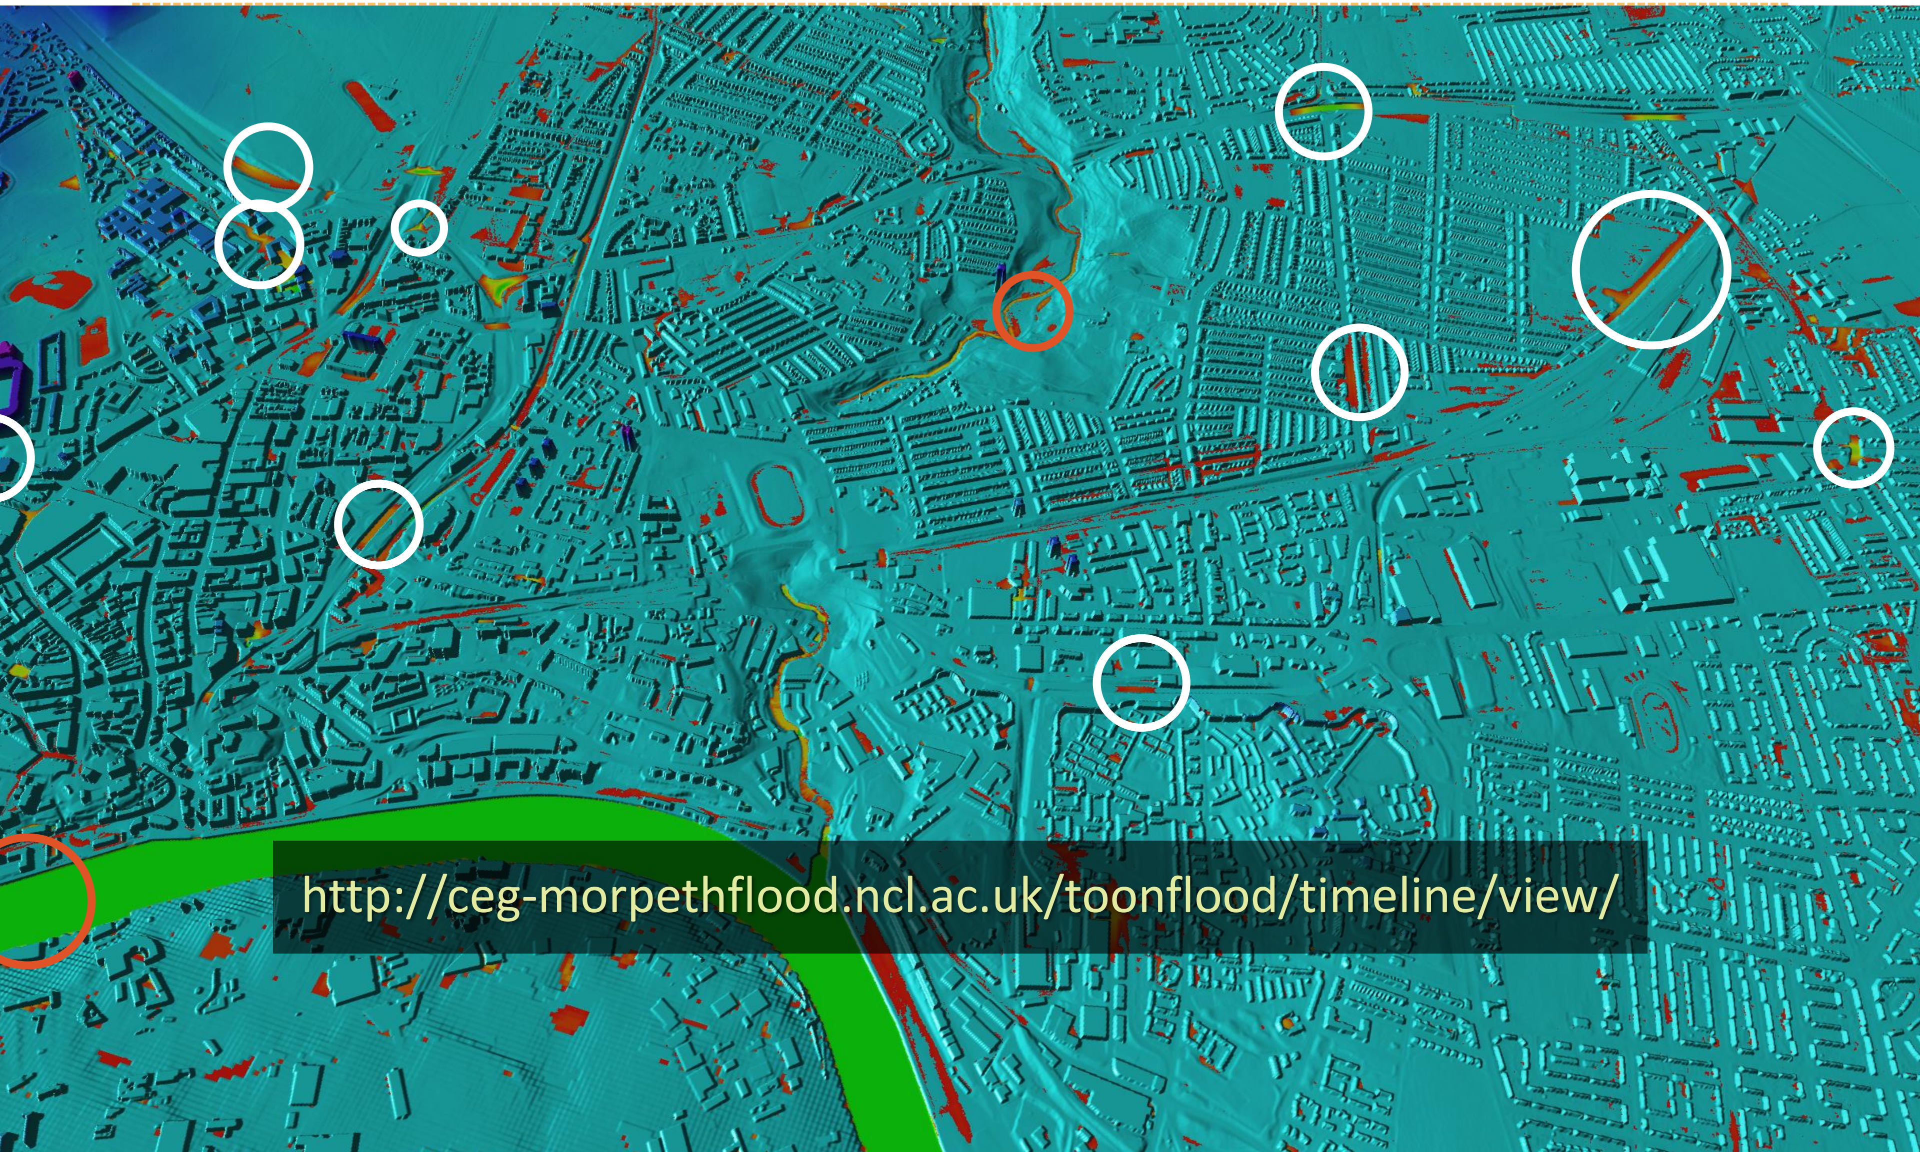

Computational performance

- Simulation times for **3 hours** from start of the event
- **31km²** area of central Newcastle
- Resolutions of **2m or better** are preferable

Resolution (cells)	1 x NVIDIA M2075 *Released 2010*		1 x NVIDIA K20 *Released 2012*	
4m (1,958,484)	00:45:21		00:31:18	0.69x
2m (7,833,169)	05:12:05	6.88x		

- Halve resolution, increase run-time approx. **eight-fold**
- Need to split work across **multiple devices**

Domain decomposition

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

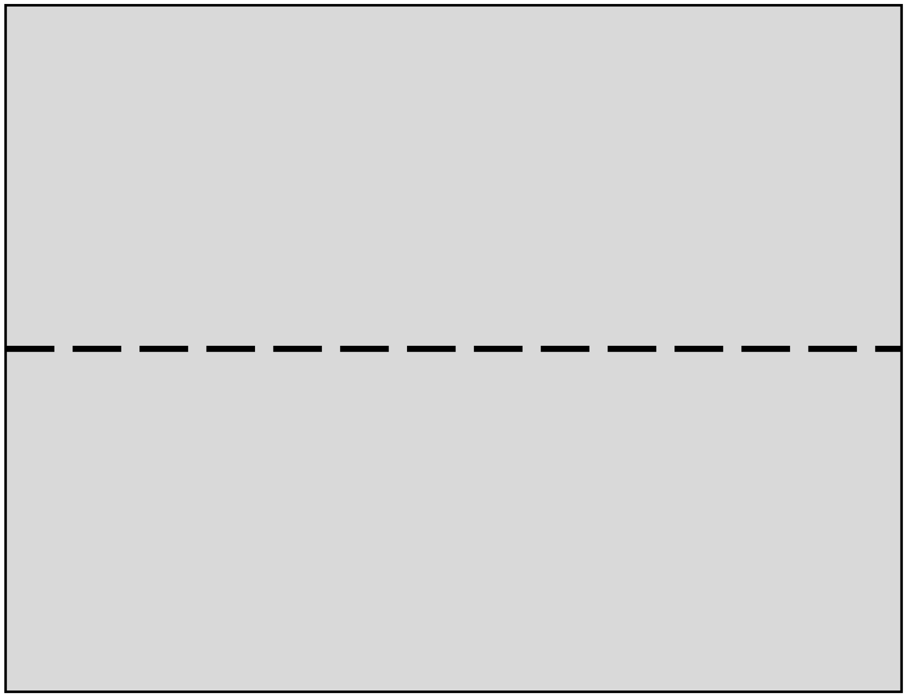

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

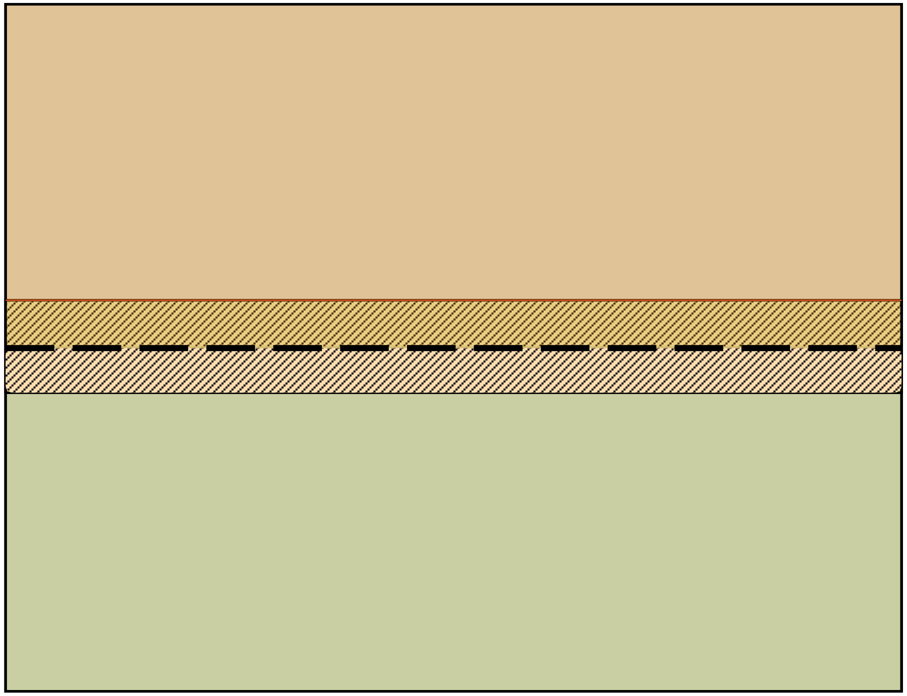

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

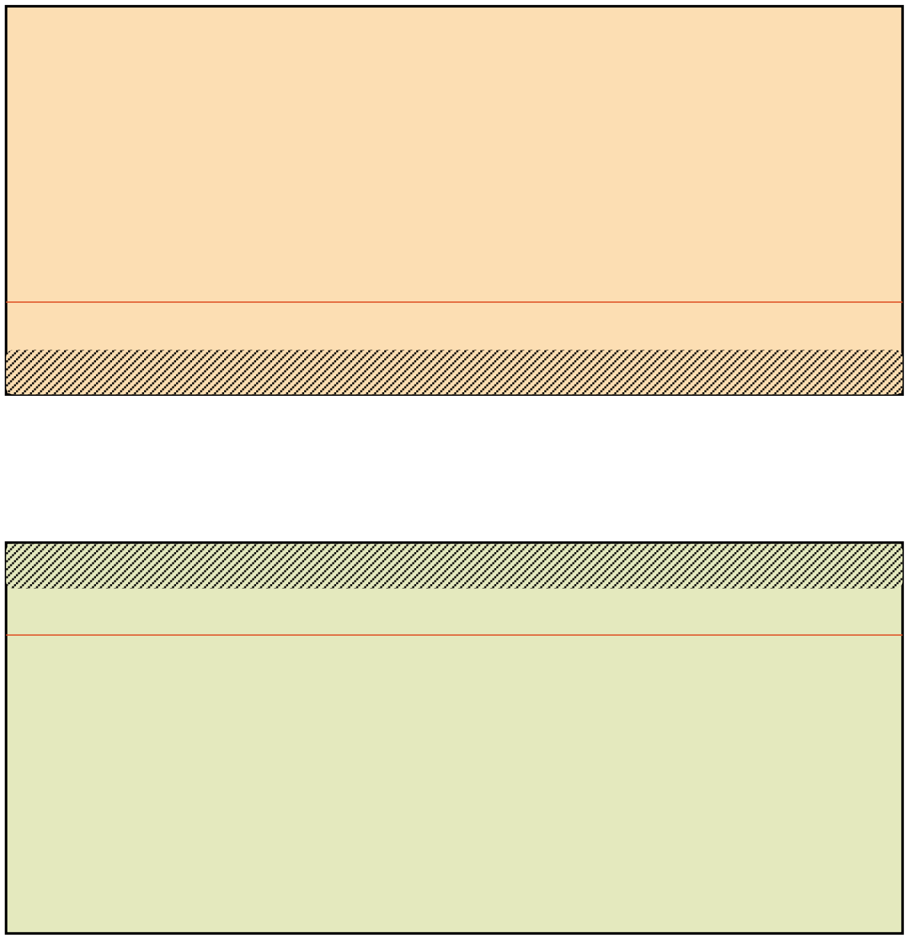

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

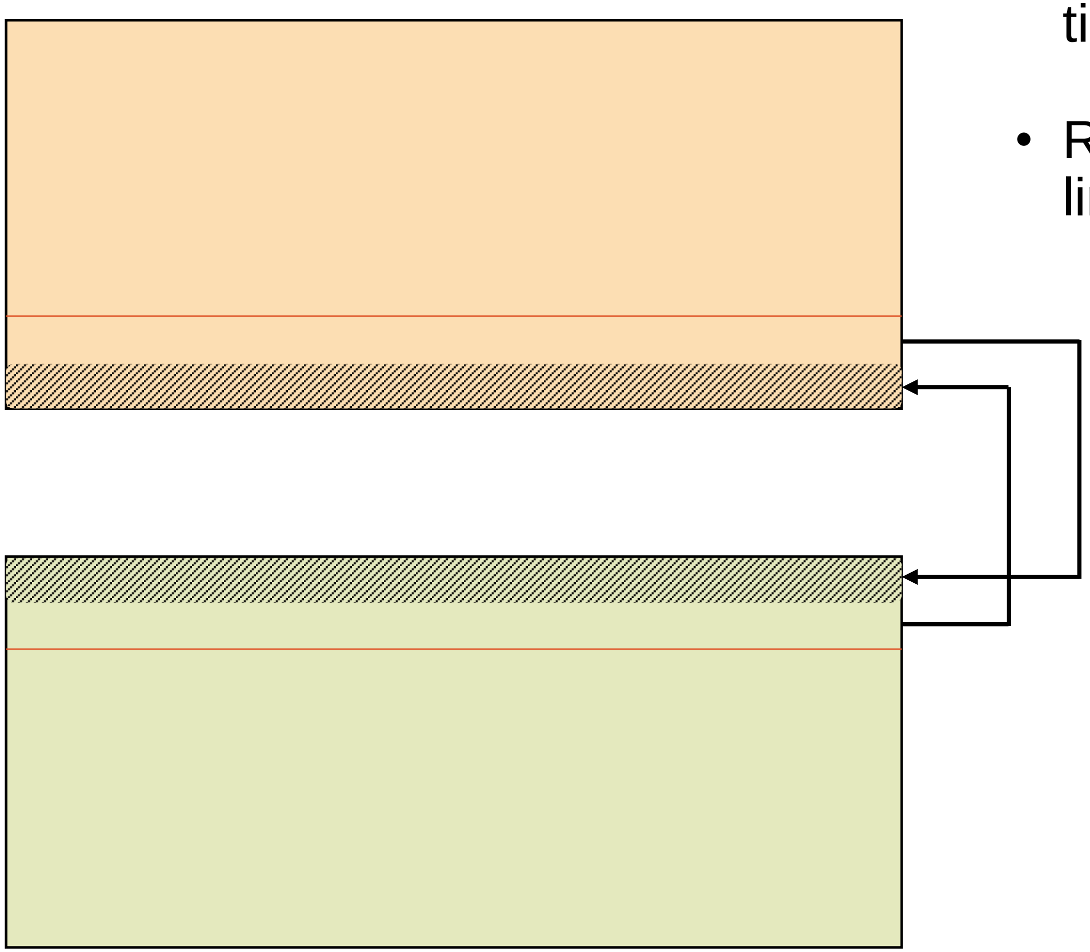

- Wave propagation 1 cell / timestep
- Row synchronisation limited by half overlap

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

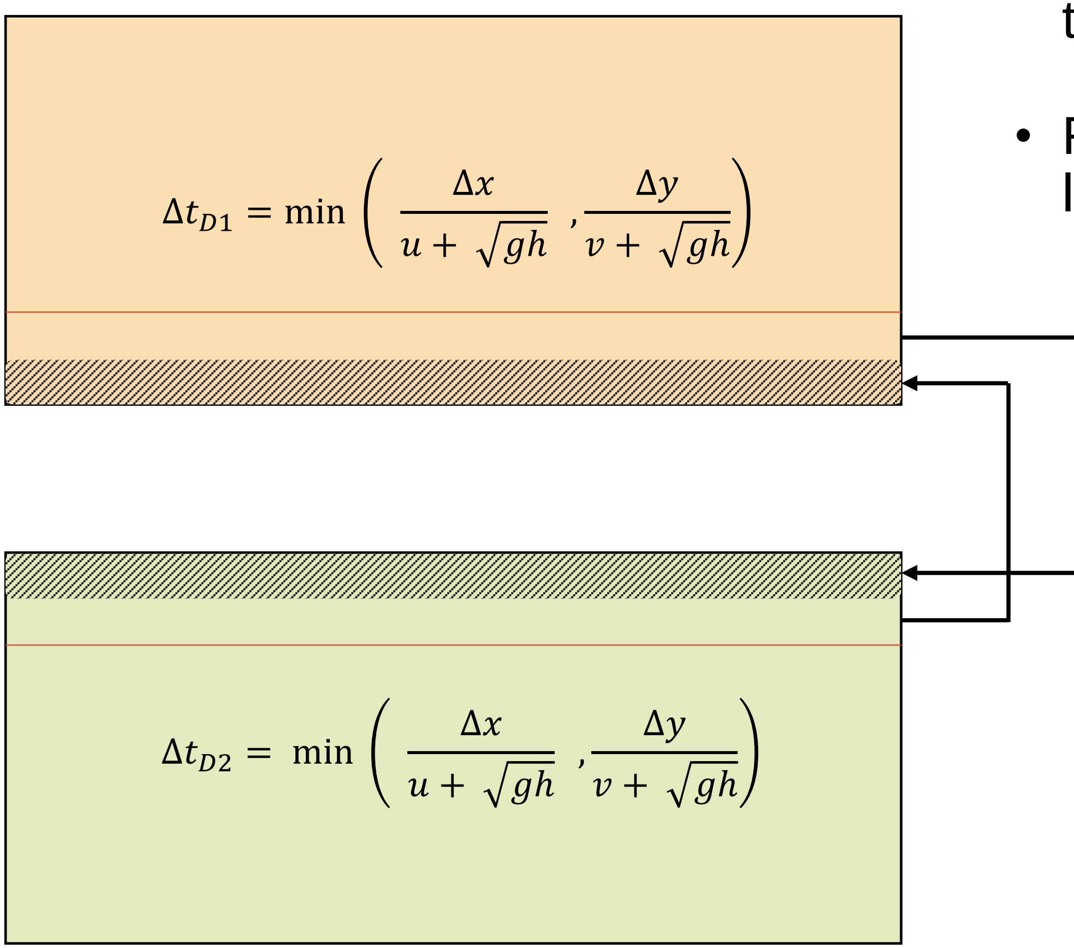

- Wave propagation 1 cell / timestep
- Row synchronisation limited by half overlap

$$\Delta t = \min\left(\Delta t_{D1}, \Delta t_{D2} \right)$$

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Domain decomposition

- Wave propagation 1 cell / timestep
- Row synchronisation limited by half overlap

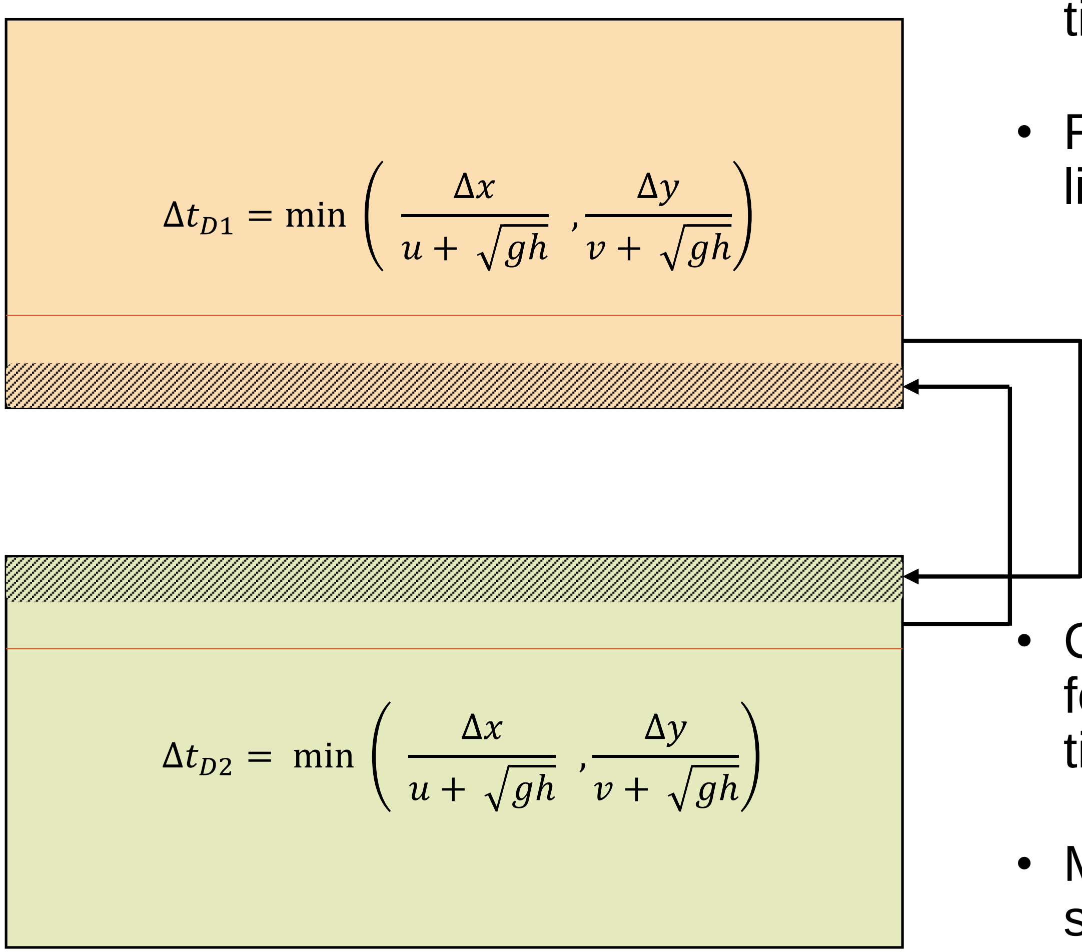

$$\Delta t = \min\left(\Delta t_{D1}, \Delta t_{D2} \right)$$

- Can be overcome by forecasting likely timesteps
- May require reversing a simulation

See works of Rostrup, Acuña, and Sætra for early work on domain decomposition.

Computational performance

- Domain decomposed across **four devices**
- Current servers are limited to **eight** PCI-e x16 slots
- Explicit timestep synchronisation or domain-independent timesteps
- Best performance scaling on **large domains** (ratio of sync:work)

Synchronisation	Overlap	Resolution (cells)	Devices	NVIDIA Tesla M2075	
None	**N/A**	**2m (7,833,169)**	1 x	05:12:05	*1.00x*
Every 1 iteration	**10 cells**	**2m (8,035,336)**	4 x	01:41:11	/3.08
Every 50 iterations	**50 cells**	**2m (8,805,496)**	4 x	01:28:31	/3.52

Computational performance

- Domain decomposed across **four devices**
- Current servers are limited to **eight** PCI-e x16 slots
- Explicit timestep synchronisation or domain-independent timesteps
- Best performance scaling on **large domains** (ratio of sync:work)

Synchronisation	Overlap	Resolution (cells)	Devices	NVIDIA Tesla M2075	
None	**N/A**	**2m (7,833,169)**	1 x	05:12:05	*1.00x*
Every 1 iteration	**10 cells**	**2m (8,035,336)**	4 x	01:41:11	/3.08
Every 50 iterations	**50 cells**	**2m (8,805,496)**	4 x	01:28:31	/3.52

- Independent domains only synchronise **every 3 seconds**

Conclusions

- **Supercomputer hydraulics feasible**
  - Potentially costly
  - Need software rewrite
- **Run-times will decrease for multi-core/heterogeneous softwares**
- **Work division is challenging**
  - Scaling vs. overheads
  - Minimising communication
  - Forecasting timesteps