NISMOD: National Infrastructure Systems Model Improving Performance and Scalability

[1]Environmental Change Institute, School of Geography and the Environment, University of Oxford
[2]School of Engineering, University of Newcastle
[3]Data and Analytics Facility for National Infrastructure, Science and Technology Facilities Council

R. Schoenmakers[1], M. Ives[1], C. Robson[2], S. Nagella[3], R. Fowler[3], S. Chorlton[3], E. Yang[3], J. W. Hall[1]

Introduction

NISMOD is the first national infrastructure system-of-systems long-term planning tool. The system was initially designed to be used solely by ITRC consortium researchers. As NISMOD reached a more mature state, stakeholders such as policy makers developed an interest in using the system for their own analyses.

With this growing user group came an increased demand for model runtime, and a need to understand the workings of each model, through sensitivity analyses of model parameters.

The NISMOD software was not designed to multiple users and such large-scale analyses. The ITRC Mistral development team was therefore set with the challenge to improve the performance and scalability of NISMOD in order to meet the increasing and changing demand for analyses.
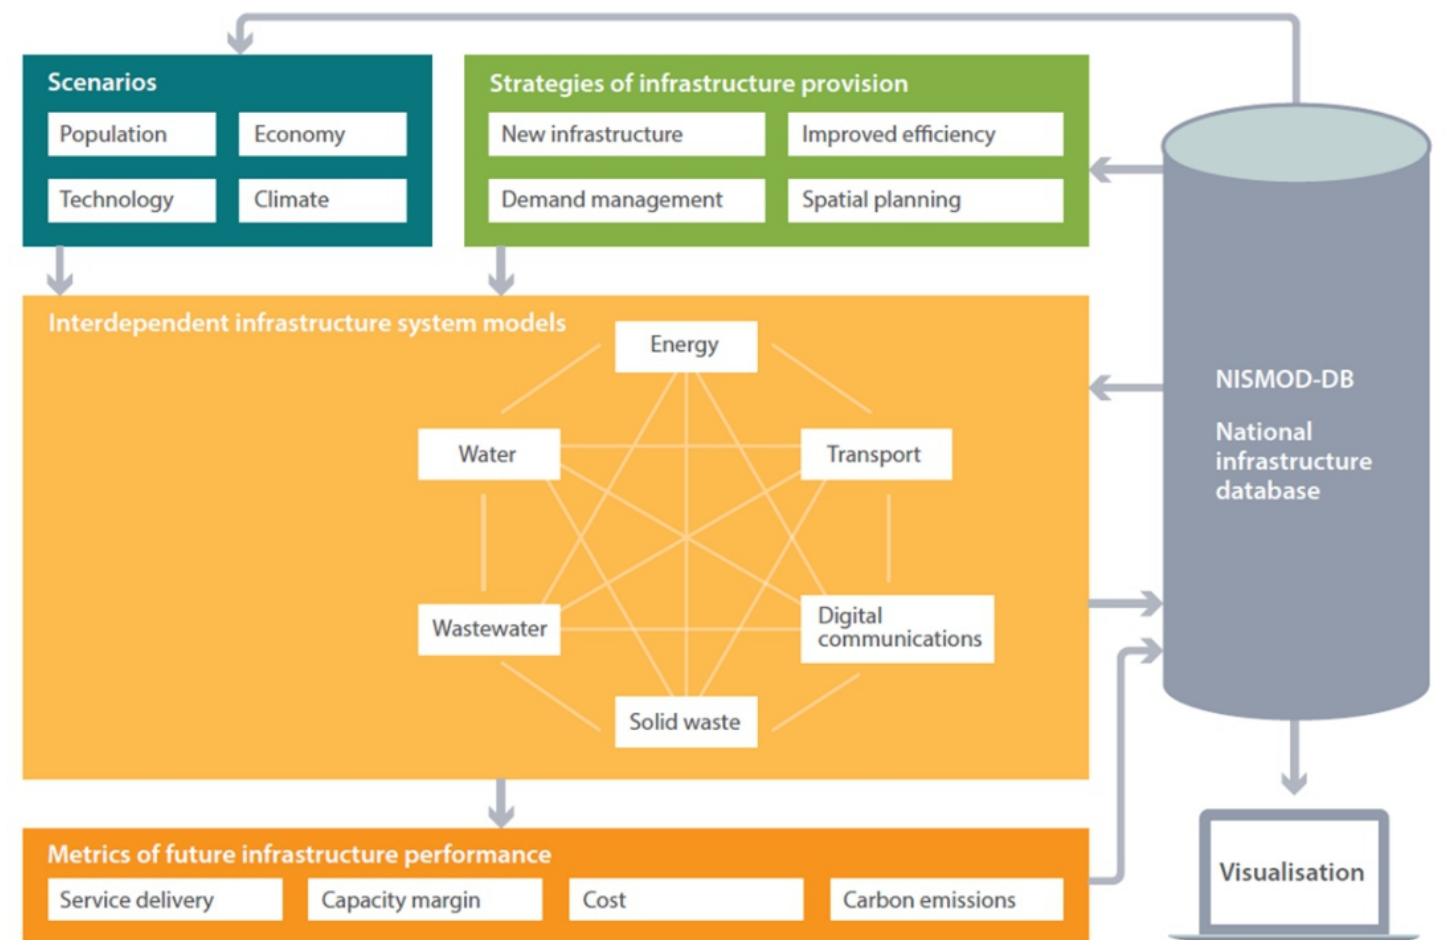
Aims and Objectives
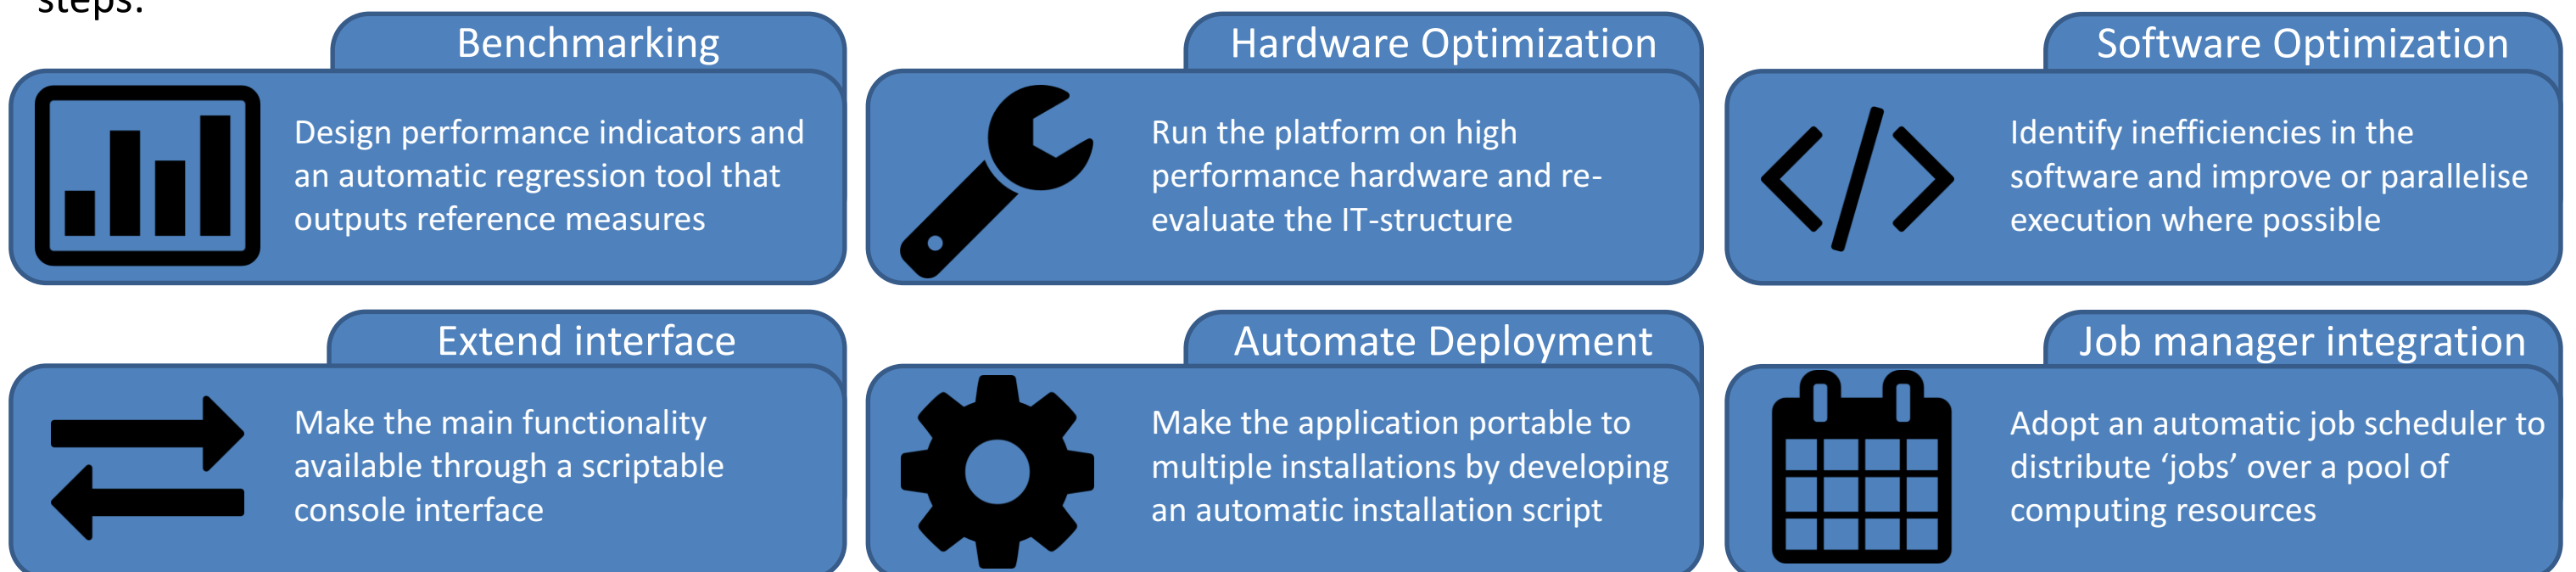
Goal 1: Improve model runtime performance

Performance – Each of the integrated engineering-based simulation models within NISMOD require a significant amount of processing time, with a full system-of-systems analysis run taking up to 12 hours to complete. As well as limiting the throughput of analysis, this affected the efficiency and quality of model/software development as debugging and testing became very time consuming. These average model run times must be improved.

Goal 2: Enable sensitivity analyses on an HPC Cluster

Scalability – The initial requirements of the system-of-systems model resulted in a single-node architecture design. This required each computing resource to be setup and started manually, which was inefficient, limited the number of concurrent analyses and made it impossible to perform large-scale sensitivity analyses within a reasonable time.

Process

The DAFNI Pilot-1 project started in August 2017 to increase the performance and scalability of NISMOD v1 and to demonstrate the value of state-of-the-art computing facilities, modelling and visualisation capabilities to the UK infrastructure innovation and research community. The ITRC MISTRAL and Science and Technology Facilities Council (STFC) teams undertook these challenges through the following process steps:

Benchmarking	Hardware Optimization	Software Optimization
Design performance indicators and an automatic regression tool that outputs reference measures	Run the platform on high performance hardware and re-evaluate the IT-structure	Identify inefficiencies in the software and improve or parallelise execution where possible

Extend interface	Automate Deployment	Job manager integration
Make the main functionality available through a scriptable console interface	Make the application portable to multiple installations by developing an automatic installation script	Adopt an automatic job scheduler to distribute 'jobs' over a pool of computing resources

Results

Performance – By improving the hardware setup, IT configuration and software bottlenecks we have been able to achieve ***30 fold reductions in the average runtime of simulation models***.
Scalability – At present we have a HTCondor pool running which enables us to run up to 7 models concurrently distributed over 3 nodes and the freedom to scale this up whenever necessary.

Lessons Learned

Options for improved performance and scalability are available even for systems not designed for high performance computing.

Further work

Use the lessons learned from this Pilot-1 project to design and build systems that are capable of performing within the parallelisation of a High Performance Computing environment.