Geophysical Research Abstracts
Vol. 21, EGU2019-10456, 2019
EGU General Assembly 2019

Building the Urban Observatory: Engineering the largest set of publicly available real-time environmental urban data in the UK

Luke Smith (1) and Mark Turner (2)

(1) Urban Observatory, Newcastle University, Newcastle-upon-Tyne, United Kingdom (luke.smith@newcastle.ac.uk), (2) Digital Institute, Newcastle University, Newcastle-upon-Tyne, United Kingdom (mark.turner@newcastle.ac.uk)

The Urban Observatory (UO) at Newcastle represents a multi-million-pound investment in ubiquitous urban sensing, providing one of the largest real-time environmental datasets in the world. This research facility, majority staffed by research software engineers, provides open data used by community groups, local and national government, and research projects ranging from cyber-security to flood forecasting. Requirements and demands have developed over the last three years, and as new data sources were incorporated, RSEs were introduced with expertise in machine learning, high-performance computing, and hybrid database architecture.

The UO's journey has not been plain sailing. Early efforts were plagued by communication problems and inconsistent clocks, rapid growth and attempts to maintain flexibility resulted in severe performance degradation, and progression into providing statistics about crowd density and people movements has introduced cyber-security and data protection challenges. This mixture of skills has broadened the remit of the RSEs involved, emphasised the need to communicate the lessons learned by means other than traditional academic publishing, and created opportunities for staff development in the form of new skills and supervision.

The UO benefited from simultaneous improvements made by RSEs across the university; staff in similar roles were beginning to meet and discuss software engineering in their own research contexts, the need for formal recognition of their roles, and career progression. This began as links between software engineers within computing and engineering but has grown to incorporate diverse backgrounds from across the university. From initial interactions, these have become a series of meetups where staff in different disciplines can come together to discuss best practices when developing software for research. Progress is being made gaining recognition for the RSE role within the university, and the visibility of RSEs in the UO is benefiting the cause.

Today, the UO receives and archives more than one hundred metrics per second, in addition to still images, video, radar and laser-scan matrices. Machine learning methods are applied in real-time across CCTV streams, and significant changes are identified using edge processing to reduce network loading. The introduction of data queues, progressively more sophisticated data sharding, thirty edge processors, and hourly replication has drastically reduced the frequency of problems. This would not have been possible without recognition for RSEs and their skills, and leveraging knowledge from the whole university.

Incorporating robust software engineering principles into the UO played a key role in allowing it to grow from a research project in a single school to one of the university's flagship research facilities, and a world-leading exemplar for running city-scale experiments on IoT networks. Designing systems that scale, are reliable, and well-maintained is the central aim of an RSE. Examples such as this will help both in Newcastle and in the wider national community for a long term aim of getting RSEs recognised as a career path within academic institutions. The UO in Newcastle is the first of six observatories in the UK, and as these other observatories begin provisioning their networks, we hope it will facilitate RSE collaboration nationally and internationally.