

# **In-transit Analytics on Distributed Clouds**

## *Applications and Architecture*

Omer Rana  
*Cardiff University, U.K.*

**Abstract:** The increasing deployment of sensor network infrastructures (in a variety of applications, ranging from environmental monitoring, "Smart Cities", energy demand forecasting, social media analysis to emergency response) has led to large volumes of data becoming available, leading to new challenges in storing, processing, analysing and transferring such data. This is especially true when data from multiple sensors is pre-processed prior to delivery to users. Where such data is processed in-transit (i.e. from data capture to delivery to a user) over a shared distributed computing infrastructure, and due to the increasing availability of software defined networks, it is necessary to provide some Quality of Service (QoS) guarantees to each user. This talk provides: (i) scenarios of applications that have these types of characteristics; (ii) a computational architecture for supporting QoS for multiple concurrent scientific workflow data streams being processed (prior to delivery to a user) over a shared infrastructure. The architecture is used to demonstrate how a streaming pipeline, with intermediate data size variation (inflation/deflation), can be supported and managed using a dynamic control strategy at each node. Such a strategy supports end-to-end QoS with variations in data size between the various nodes involved in the workflow enactment process.

### **BRIEF BIOGRAPHY**

Omer Rana is Professor of Performance Engineering at Cardiff School of Computer Science & Informatics. He was formerly the deputy director of the Welsh eScience Centre at Cardiff University -- where he had an opportunity to collaborate with a number of scientists working in computational science and engineering. He holds a PhD in "Neural Computing and Parallel Architectures" from Imperial College (University of London, UK). His research interests are in the areas of high performance distributed computing, data mining and analysis and multi-agent systems. Prior to joining Cardiff University he worked as a software developer with Marshall BioTechnology Limited in London, working on projects with a number of international biotech companies, such as Merck, Hybaid and Amersham International. He has been involved in the Distributed Programming Abstractions and the 3DPAS themes at the UK National eScience Institute. He is an associate editor of the ACM Transactions on Autonomous and Adaptive Systems, IEEE Transactions on Cloud Computing, series co-editor of the book series on "Autonomic Systems" by Birkhauser publishers, and on the editorial boards of "Concurrency and Computation: Practice and Experience" (John

Wiley), the Journal of Computational Science (Elsevier) and the recently launched IEEE Cloud Computing magazine. Along with his co-researchers, he was recipient of the best paper award at CLOSER 2013 (Aachen, Germany).