
ESTD1845
QUEEN'S
UNIVERSITY
BELFAST

DISTRIBUTED STREAM PROCESSING ON EDGE AND FOG COMPUTING VIA TRANSPRECISE COMPUTING

HANS VANDIERENDONCK

DiPET COORDINATOR

13 APRIL 2021

SCIENTIFIC BACKGROUND (1) - MOTIVATION

Edge and fog computing
Distributed computing paradigms focused on performing services and computations closer to end users. SDCDN considers the question of distributing computations across networks.

Stream processing
Stream processing encompasses a class of applications that operate on periodic events or data streams in a dataflow-like fashion.

Transprecise computing
Traditionally, computing has been designed to be exact, often resulting in excessive computation and data footprints. Transprecision aims to adapt precision as requirements and compute power change.

Quality of Service / Experience
Edge/fog performance quantified through QoS/QoE metrics; align closely to precision (detail in computation) and accuracy (value of computation).

QUEEN'S UNIVERSITY BELFAST

SCIENTIFIC BACKGROUND (2) - SCENARIOS

Malleable applications
Adapt resource usage and QoE of streaming applications at runtime based on availability of resources and/or characteristics of the data.

Deploying applications on insufficient resources
Where applications require more compute power are memory than a device can provide, it may be preferable to execute it at reduced QoE rather than not at all.

Migration of applications
Availability of resources can be a trigger for application migration, as well as a contributing factor to decide which applications to migrate and where to migrate them to.

Oversubscription of resources
Quantifiable loss of QoE in the presence of temporary oversubscription of resources. Instead of applications competing for resources, QoE-tolerant applications can be scaled back.

QUEEN'S UNIVERSITY BELFAST

KEY CHALLENGES AND POTENTIAL IMPACTS

CHALLENGES

Explore and characterize the application domains likely to benefit from transprecision

- Network monitoring (anomaly detection, intrusion detection)
- Video analytics (object detection, scene understanding)
- Streaming analytics

Expressing transprecision and adapting the execution of streaming applications

Efficient distribution of streaming applications

POTENTIAL IMPACTS

Cost-effective applications, infrastructures

Enable unprecedented response latency, throughput, energy consumption and system utilization

Integrate the project outcomes in reproducible, deployable technology

PARTNERS

Coordinator, Lead WP5 Management
Transprecise computing (Lee, Vandierendonck)
Edge computing and middleware (Varghese)

Lead WP2 Programming Model and Runtime Implementation
Programming languages and runtimes, analytics (Pratikakis)

Lead WP3 Dynamic Resource Mapping and Middleware
Fog computing (Mondal, Pierre)
Geo-distributed data stream processing (Mondal, Pierre)

Lead WP4 Deployment and Evaluation, community network testbed
Networked distributed systems (Vilchez, Freitag, Navarro)
Wireless Community Networks (Cerdà-Alabern, Baig)

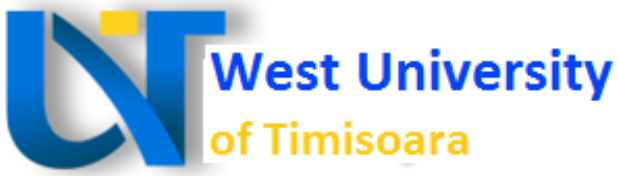

Lead WP1 Streaming Algorithms with Tunable Precision
AI/ML, Distributed Computing (Iuhasz, Petcu)
HPC, Cloud Computing, Service Oriented Computing (Panica)

MAIN SCIENTIFIC RESULTS

A. Developing a new technology based on transprecise computing

- Identified plausible transprecise streaming applications
- Dissemination and adoption of the transprecise computing idea
- Performance monitoring and modelling scalability of transprecise applications w/ ±2% error

B. Transprecise Video Analytics

Using Deep Neural Networks (DNN)

- Malleable application: Dynamically select DNN depending on video content in order to minimize computation, energy consumption
- Reduced precision arithmetic: Experimental evaluation of hybrid floating-point (HFP8) in DNN training

C. Testbed Development

- Deploy network monitoring in the Guifi.net community network
- Collect and anonymise netflows; comprehensive analysis of algorithms for anomaly detection using the event detection engine (EDE)

SUSTAINABILITY AND VALORISATION

INDUSTRY ENGAGEMENT

- QUB through the ECIT engineering team (knowledge transfer support mechanism) and its industry partnerships
- UPC will engage with eXo to apply transprecise network monitoring to support network management

FOLLOW-ON PROJECTS

- UR1 will leverage the autoscaling mechanisms to minimise reconfigurations (TERRA FORMA, France)
- UVT will continue the development of the event detection engine (EDE), previously designed for exascale systems, in follow-on projects
- QUB will continue the exploration of transprecise computing in future projects (EU MSCA IF Fellowship SoftNum)
- Additionally, the concepts are disseminated through undergraduate and graduate courses and research projects, as well as PhD projects, by all partners