TEADAL

Controlling Data Gravity and Data Friction: From Metrics to Multidimensional Elasticity Strategies

TU Wien: Boris Sedlak, Victor Casamayor Pujol, Praveen Kumar Donta, Schahram Dustdar

04/07/2023

WWW.TEADAL.EU

Illustrative Scenario

- Medical data from IoT sensors **accumulated** locally, before transferring to cloud storage, resulting in a personal data lake [1] and high **data gravity**.
- Data must be **transformed** to respect privacy policies before ingesting it into the cloud storage, thus creating **data friction**.

[1] Alrehamy et al., Personal Data Lake With Data Gravity Pull, IEEE Big Data 2015

Illustrative Scenario

- Medical data from IoT sensors **accumulated** locally, before transferring to cloud storage, resulting in a personal data lake [1] and high **data gravity**.
- Data must be **transformed** to respect privacy policies before ingesting it into the cloud storage, thus creating **data friction**.

- Collect metrics at **cluster leader**
- Evaluate **device state** due to metrics
- Plan **corrective measures**
  → Scale down resolution of data
  → Move processing to nearby device

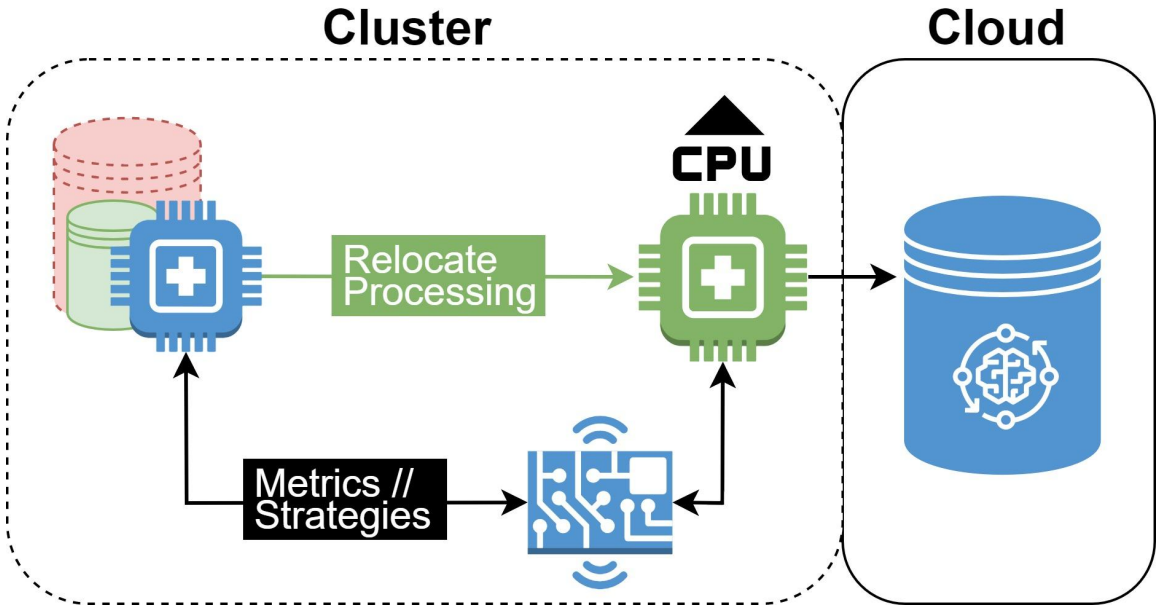

[1] Alrehamy et al., Personal Data Lake With Data Gravity Pull, IEEE Big Data 2015

Markov blankets (1/3)

- **Which situation triggers which adaptation mechanism?**
  → Highly dependable on the given context

- Require a modelling framework for evaluating the system according to requirements (i.e., **SLOs**) and act **accordingly**
- **Markov blankets** as a model for mapping sensory state (i.e., metrics) to action state (i.e., elasticity strategies)

Markov blankets (2/3)

Suppose we have an **SLO** that is concerned with **limiting** data gravity and data friction, we can create a **Bayesian network of variables** that reflect the sensory and active state

Include for each SLO all metrics that determine its state (**parents**), elasticity strategies (**children**), and contextual factors influencing strategies (**children's parents**)

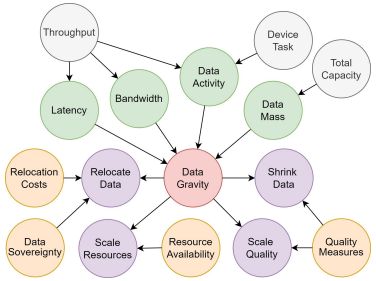

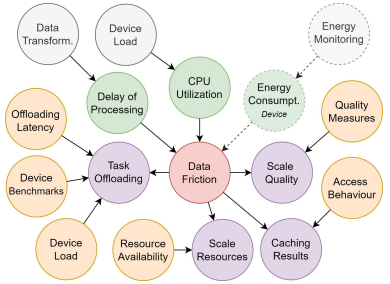

Markov blankets (2/3)

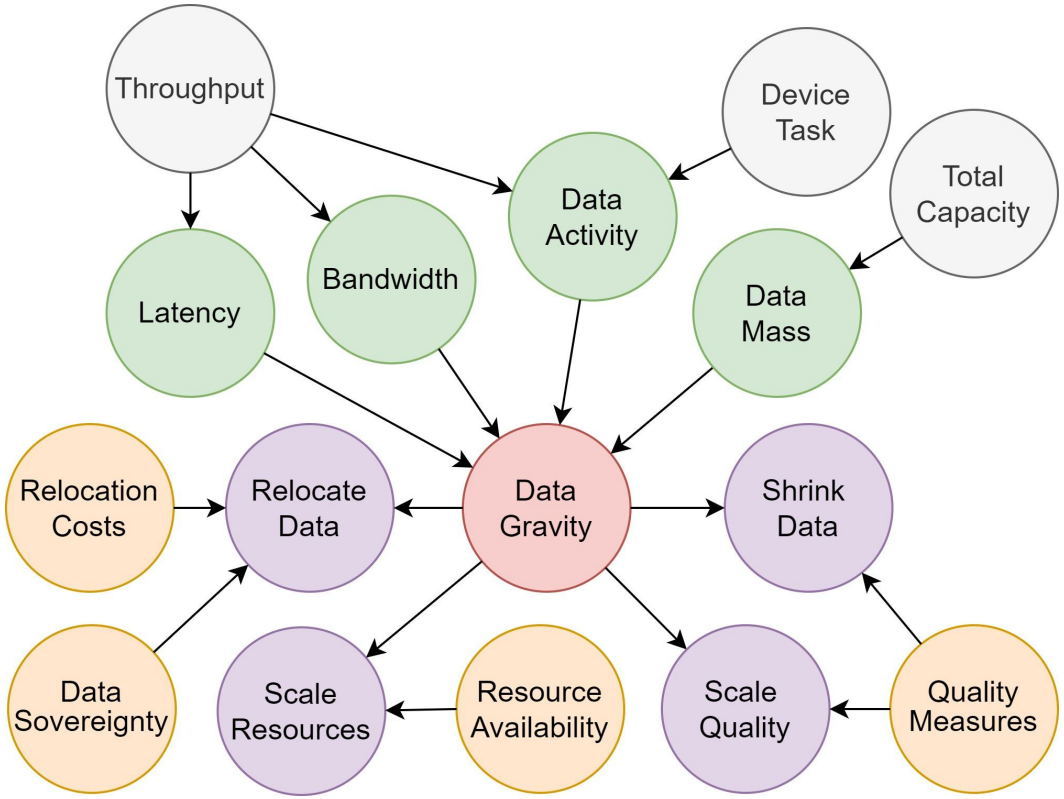

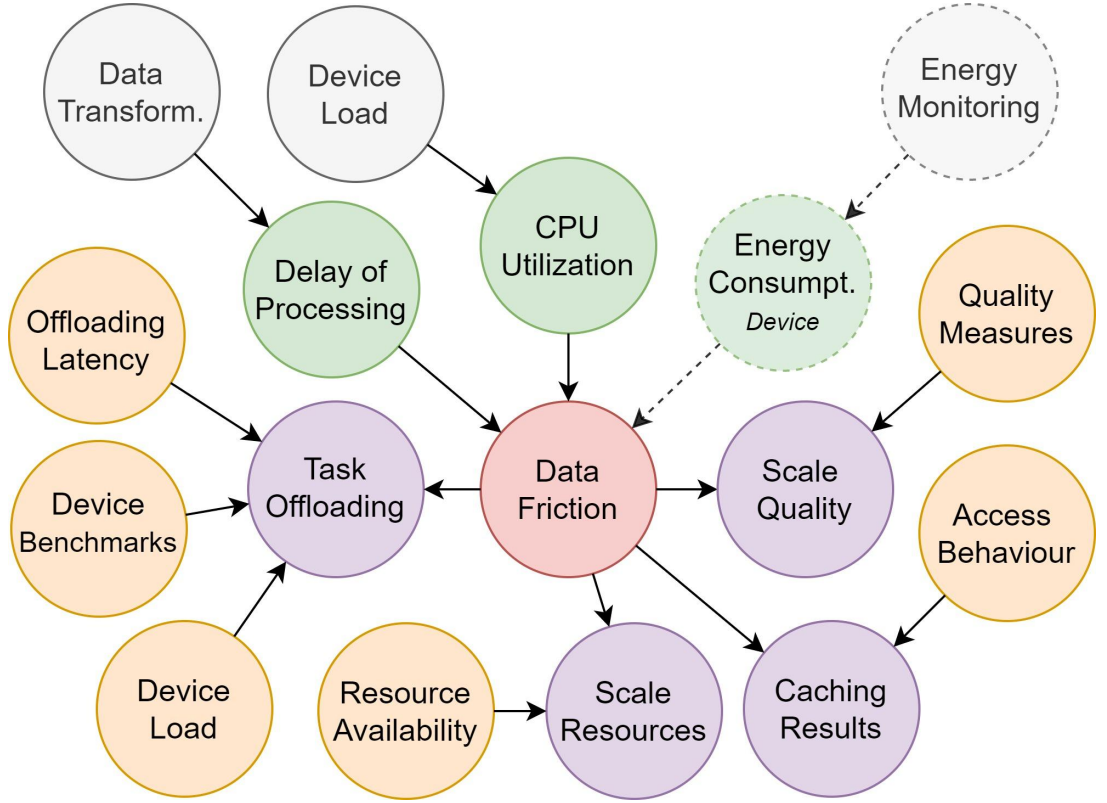

Include for each SLO all metrics that determine its state (**parents**), elasticity strategies (**children**), and contextual factors influencing strategies (**children's parents**)

Markov blankets (3/3)

- Contains metrics, elasticity strategies, and contextual factors, which we all include within a **Markov SLO Configurations (MSC)**
- Contextual information allows to switch between multiple dimensions of elasticity strategies, e.g., by **scaling resources** or **QoS**
- MSC can change over time, including the perception of metrics and the mapping to respective elasticity strategies

Open architectural questions

- Summarized under 4 main challenges, which follow the **MAPE+K** cycle
- Conceptually build upon the Polaris framework [2] to solve them

[2] Nastic et al., Polaris Scheduler: Edge Sensitive and SLO Aware Workload Scheduling in Cloud-Edge-IoT Clusters, IEEE Cloud 2021

Open architectural questions

- Summarized under 4 main challenges, which follow the **MAPE+K** cycle
- Conceptually build upon the Polaris framework [2] to solve them
- **SLO controller** for evaluating metrics and mapping to single strategy

SLO
Controller

[2] Nastic et al., Polaris Scheduler: Edge Sensitive and SLO Aware Workload Scheduling in Cloud-Edge-IoT Clusters, IEEE Cloud 2021

MAPE+K - Collect Metrics in the Compute Continuum

- Collecting device metrics at the cloud consumes **bandwidth**
- Accumulate metrics locally or at a cluster leader (e.g., a fog router)
- Store metrics in a time-series DB until consumed by the SLO controller

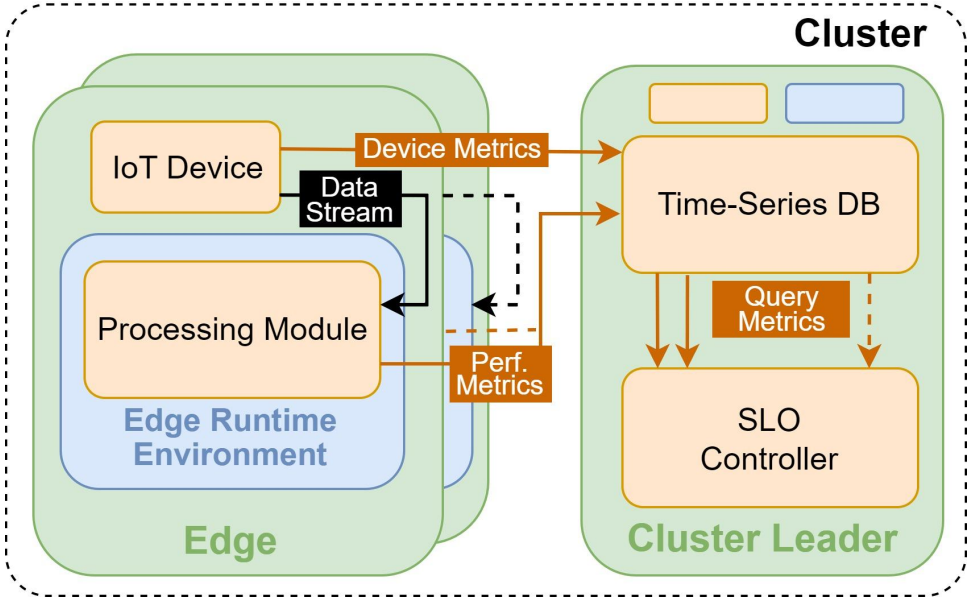

MAPE+K - SLO Specification & Analysis

- Analyzing metrics in the cloud introduces a considerable delay
- Evaluate SLOs **decentralized** to foster timely reactions
- Specify how the SLO result is determined by **metric formulas**

- For example, evaluate data friction according to **delay** and **CPU load**

$$processingDelay_{ms} \times f\left(cpuLoad_{\%}\right) \tag{1}$$

$$f(x) = \begin{cases} (x/t_x)^2, & \text{if } x \geq t_x \\ 1, & \text{otherwise} \end{cases} \tag{2}$$

MAPE+K - Context-Aware Planning of Strategies

- Planning predetermined strategy limited to one elasticity dimension
- Consider context for identifying the most beneficial dimension, e.g., QoS
- Analyzes dynamic and static configurations (e.g., quality measures)

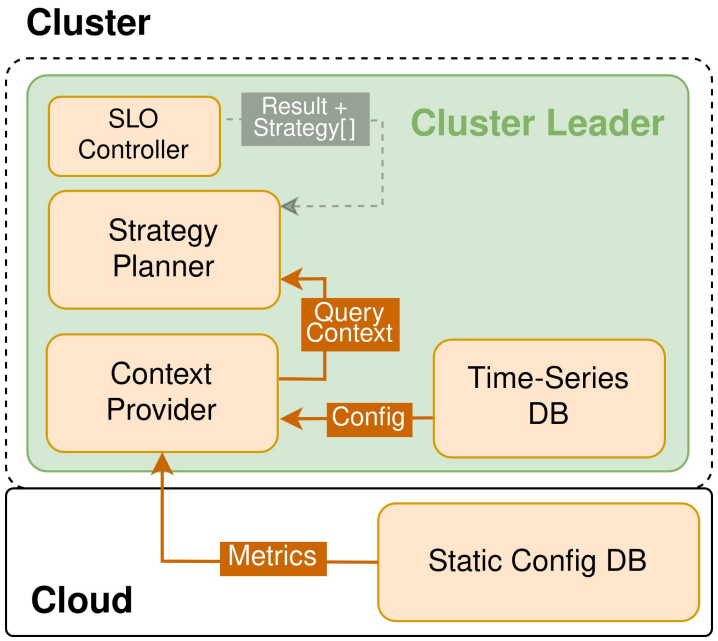

MAPE+K - Distributed Execution of Elasticity Strategies

- QoS (e.g., resolution) cannot be decreased at the consumer anymore
- Strategies that affect the QoS must be executed at the consumer directly

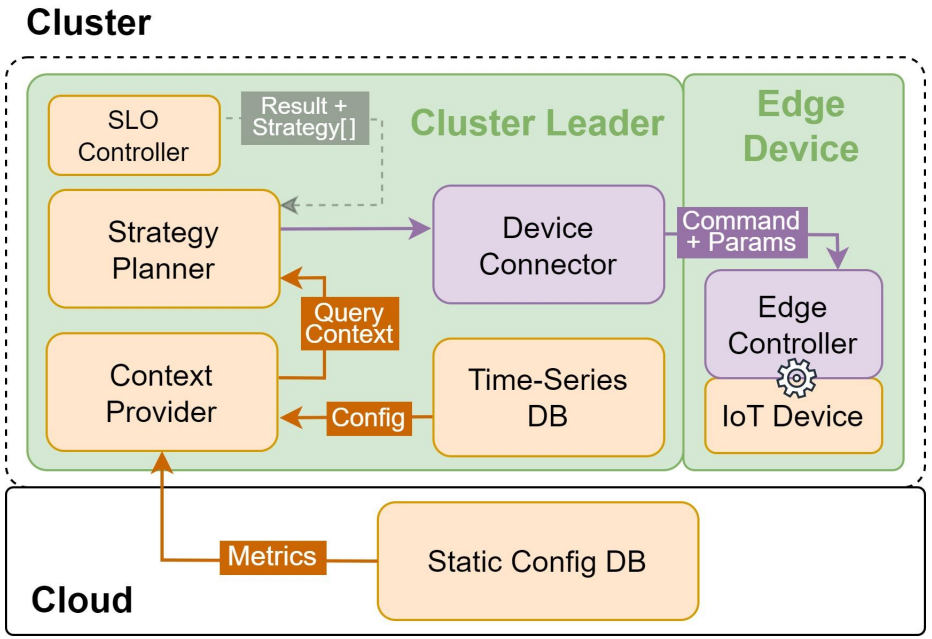

MAPE+K - Knowledge Transfer and MSC

- Metrics collected and evaluated at cluster leaders
- Cluster leaders plan elasticity strategies and orchestrate to devices
- Knowledge can be federated to create **hierarchical structures**

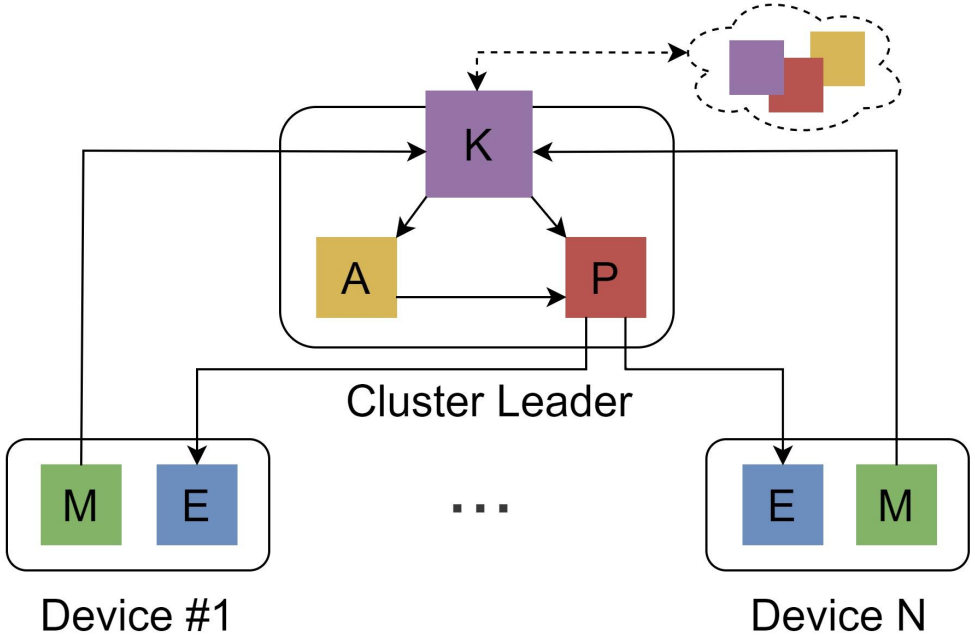

Summary

- MSCs as a method to **map** metrics to elasticity strategies
- Follows the MAPE-K cycle for continuous adaptation
- Context-based planning can consider multiple **elasticity dimensions**
- Computing Continuum as a **composition** of Markov blankets

TEADAL

TEADAL.EU

@TEADAL_eu

@TEADAL

TEADAL project is funded by the EU's Horizon Europe programme under Grant Agreement number 101070186