SummerSoc
Service Oriented Computing

18th Symposium and Summer School On Service-Oriented Computing
(June 24th - 29th, 2024)

Computing Continuum – Session Chair: Schahram Dustdar

Distributed Intelligence in the Computing Continuum using Active Inference

Prof. Schahram Dustdar & Dr. Praveen Kumar Donta

Distributed Systems Group, TU Wien, Vienna 1040, Austria.

(26th June 2024)

TEADAL
https://www.teadal.eu/

TU WIEN dsg DISTRIBUTED SYSTEMS GROUP

Outline

- Distributed Systems and State-of-the-Art
- Distributed Intelligence
- Human Eco-systems vs Distributed Computing Continuum Systems
- Computing Continuum Monitoring and Self-healing
- Data Protocols for Computing Continuum
- Clustered Edge Intelligence through Active Inference

Evolution in Distributed Systems

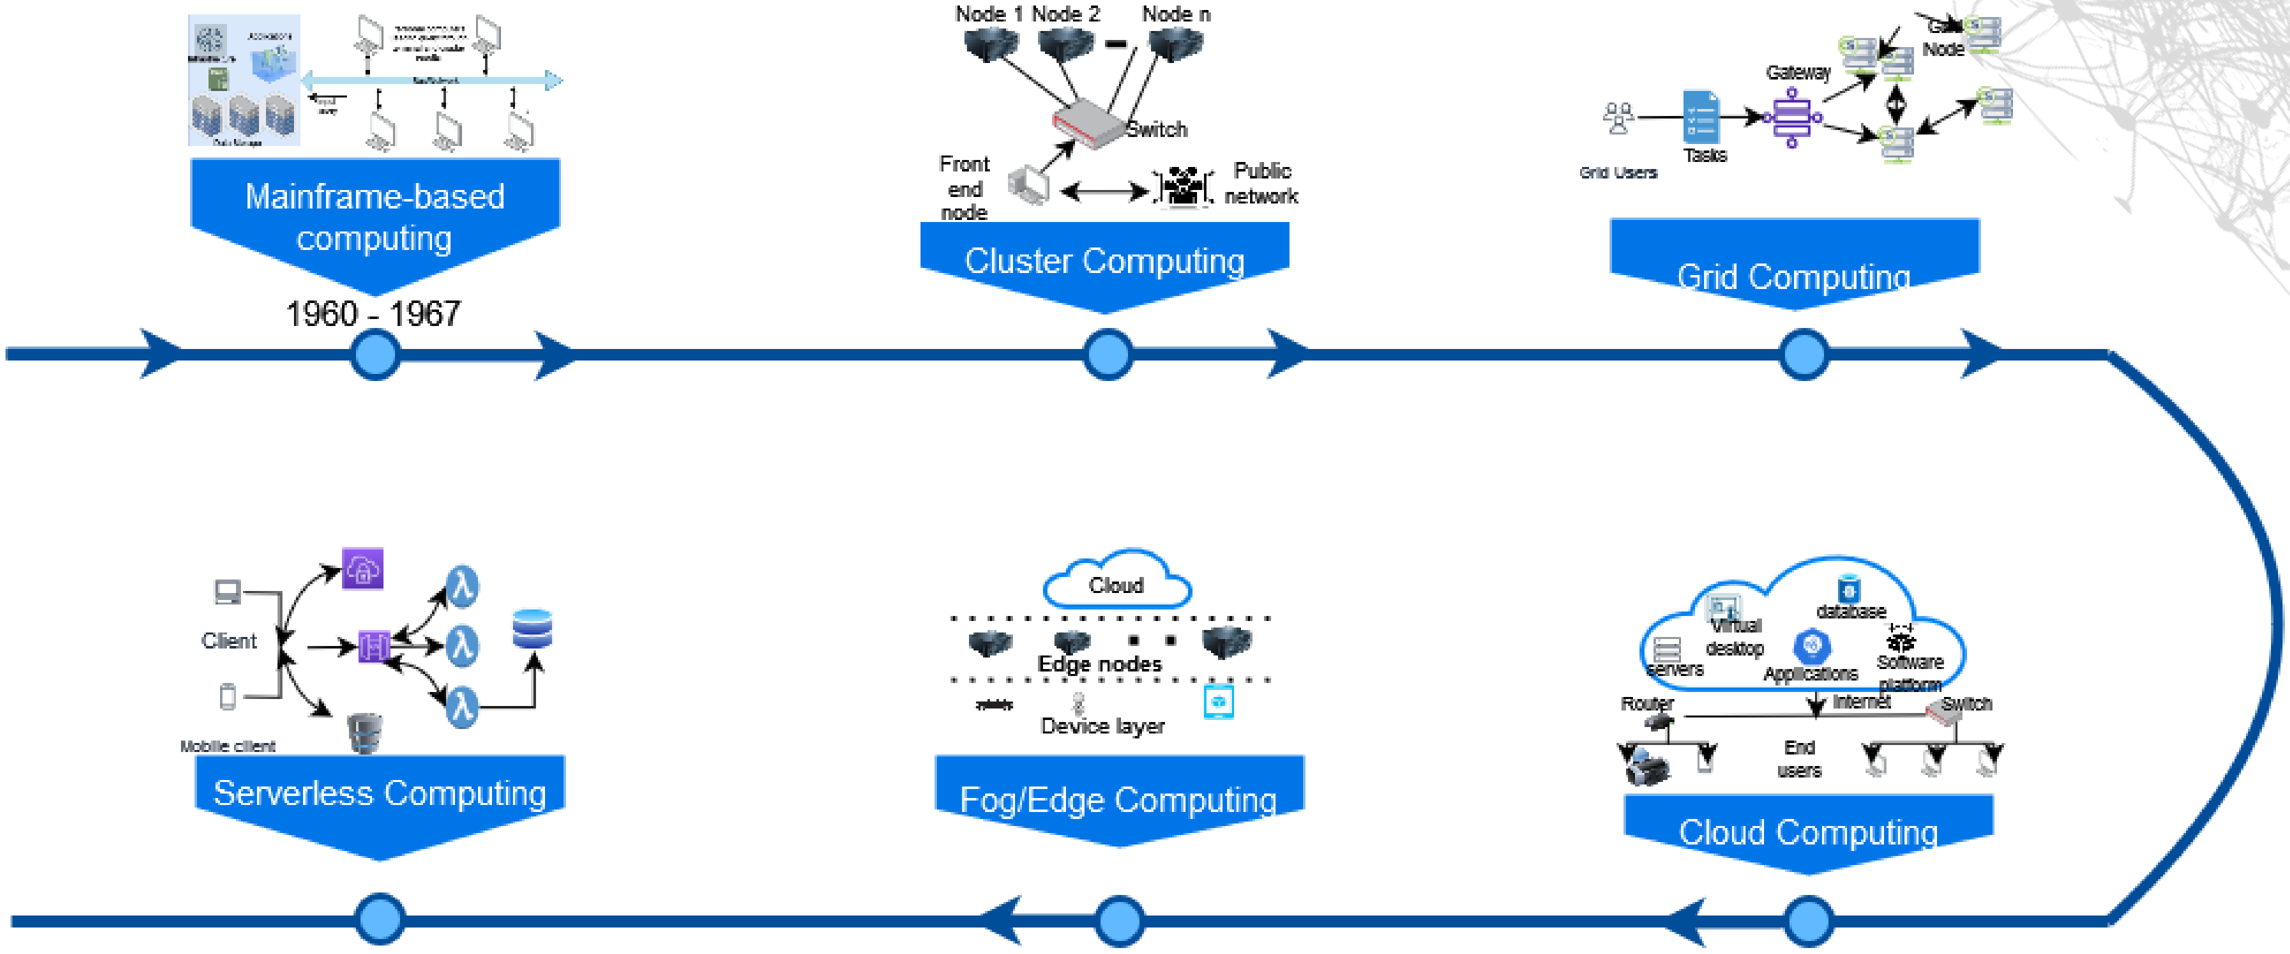

Donta, P.K., et al. "Exploring the potential of distributed computing continuum systems." *Computers* 12.10 (2023): 198.

Current State

- Distributed Systems are key to our society
- Underlie our critical infrastructures and applications (Smart cities, Healthcare, Autonomous vehicles,…)
- Interconnectedness (fabric) of components (HW, SW, People) induces complexity
- We increasingly see fundamental issues we need to address

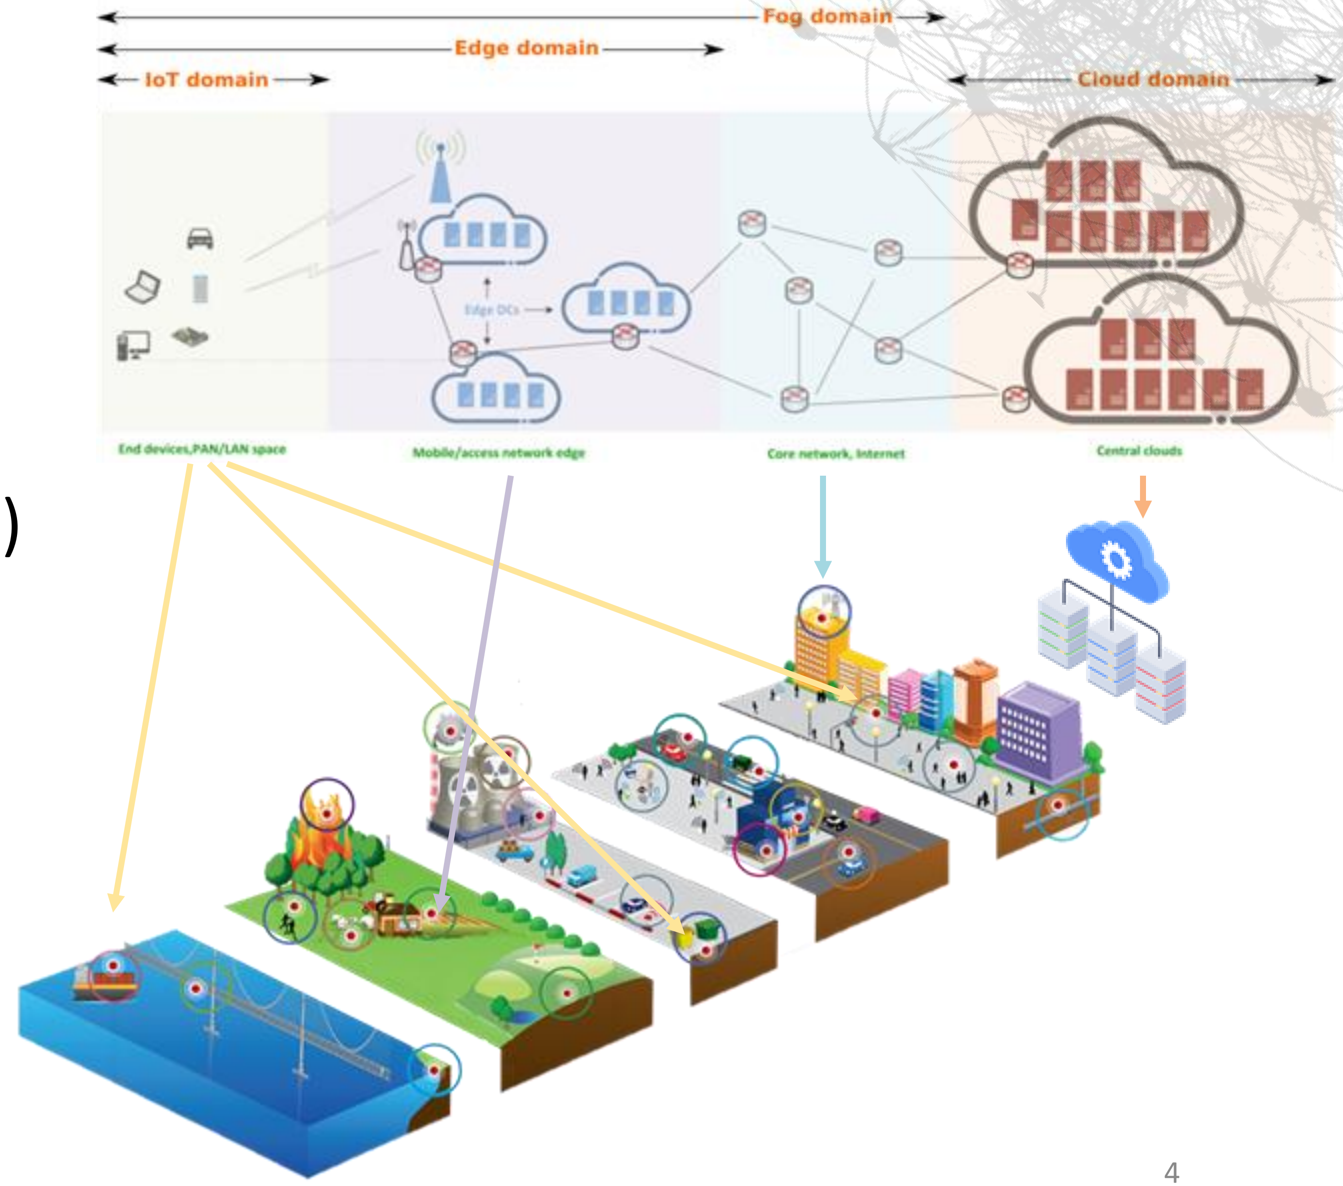

Infrastructure

➢ Computing continuum

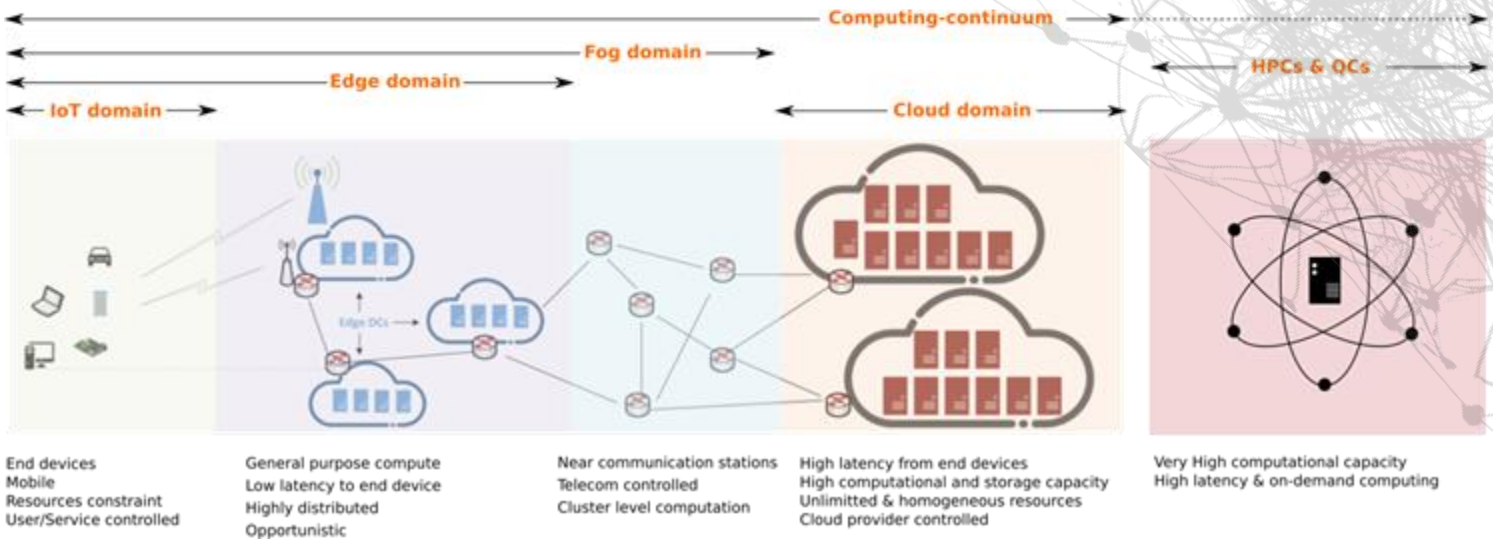

➢ Application performance highly dependent on the underlying infrastructure
  - Heterogeneity of resources & heterogeneous distribution
  - Resources diverse interconnection

Use cases

Smart City

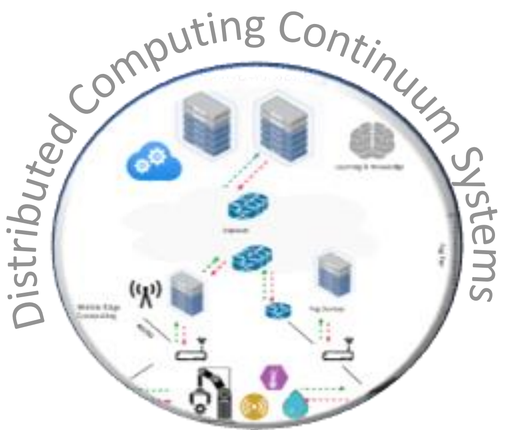

Healthcare

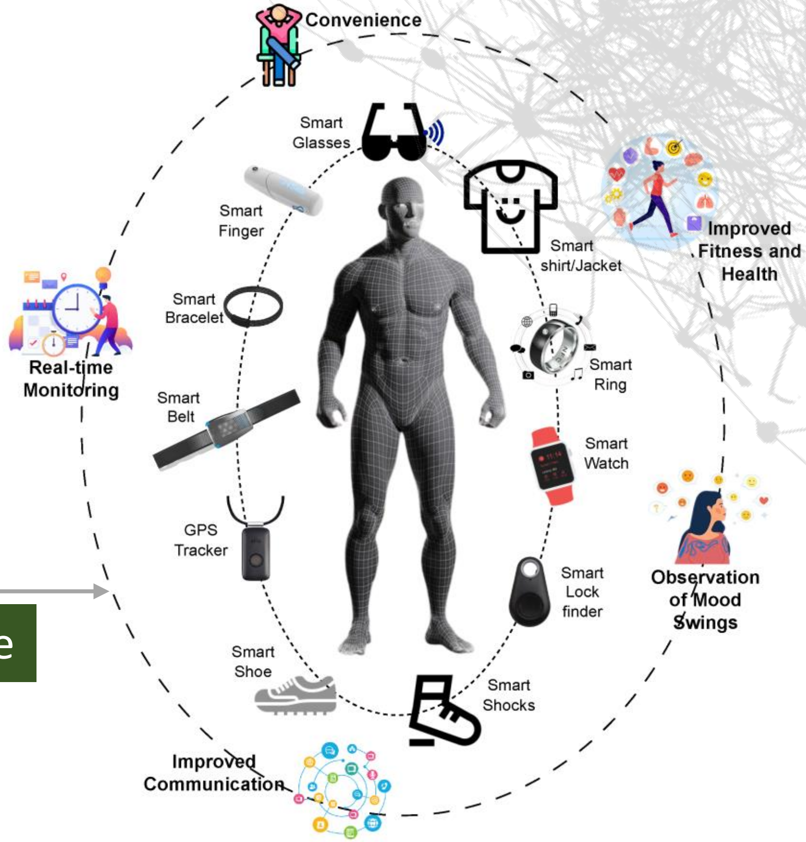

Industry Automation

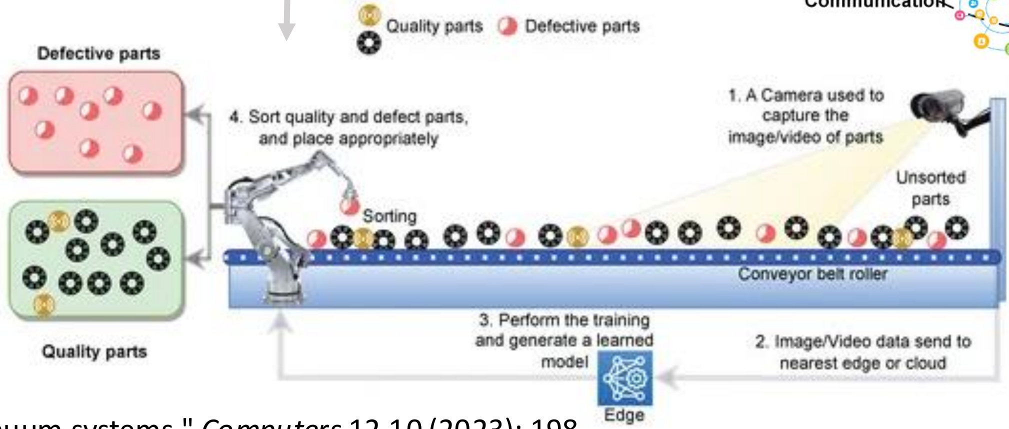

Donta, P.K., et al. "Exploring the potential of distributed computing continuum systems." *Computers* 12.10 (2023): 198.

Distributed Computing Continuum Systems

Benefits	**Challenges**
• Bandwidth Optimization • Scalability & Availability • Low latency • Resource optimization & Load balancing • Reliability or Flexibility	• Interoperability • Data Synchronization • Governing resources • Privacy

Distributed Intelligence in Computing Continuum

Casamayor Pujol, V., Donta, P. K., Morichetta, A., Murturi, I., & Dustdar, S. (2022, November). Distributed computing continuum systems–opportunities and research challenges. In *International Conference on Service-Oriented Computing* (pp. 405-407). Cham: Springer Nature Switzerland.

Traditional Approaches

Techniques used in Previous and Majority of Ongoing works

Artificial Intelligence

Machine Learning

Requirement of huge input/training data

Computational hungry

Relatively Expensive (Cost vs Benefits)

Adversarial Attacks

Difficulty in Real-time Accuracy verification

Resource constraint

Human Eco-system vs. Computing Continuum

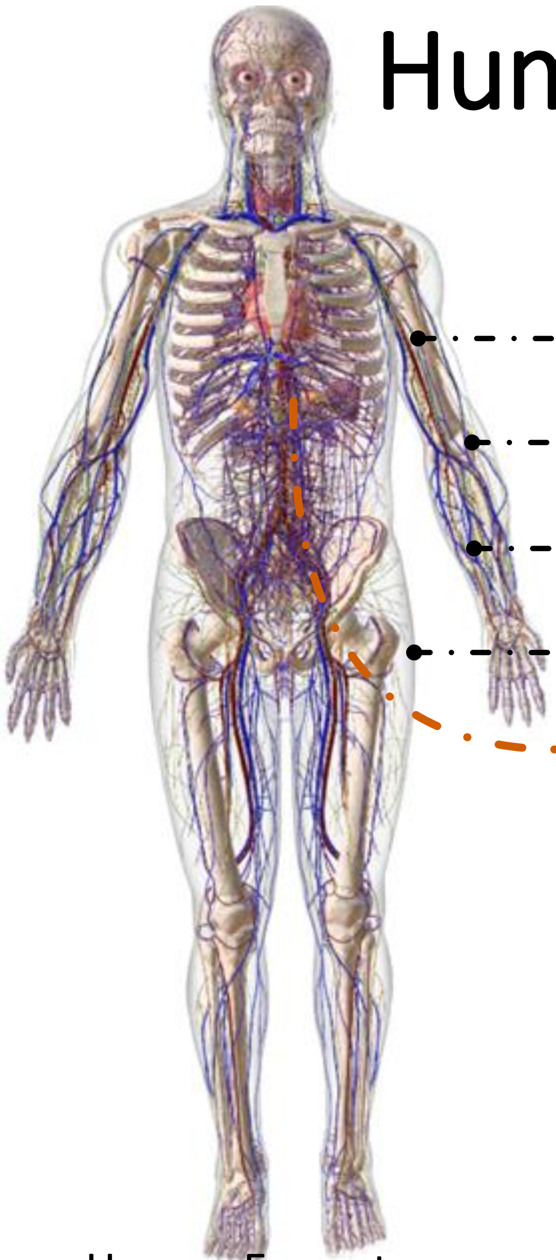

The human body is comprised of a series of complex systems, including:

- Skeletal System
- Nervous System
- Cardiovascular System
- Lymphatic System
- Endocrine System

Nervous System, Cardiovascular System → Infrastructure Systems

Lymphatic System, Endocrine System → Regulation Systems

- Brain
- Spinal Cord
- Cranial Nerves
- Spinal Nerves

- Oxygen
- White Blood Cells
- Hormones
- Nutrients

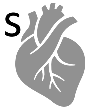

Helping the body meet the demands (**40k neurons)**

Control Internal Environment, Memory and Learning (**86 billion neurons**)

Alshami, A. M. (2019). Pain: is it all in the brain or the heart?. *Current Pain and Headache Reports*, *23*, 1-4
Pujol, V., *Donta, P.K.*, et. Al.., 2022, November. Distributed computing continuum systems–opportunities and research challenges. In *International Conference on Service-Oriented Computing* (pp. 405-407). Cham: Springer Nature..

Human Eco-system vs. Computing Continuum

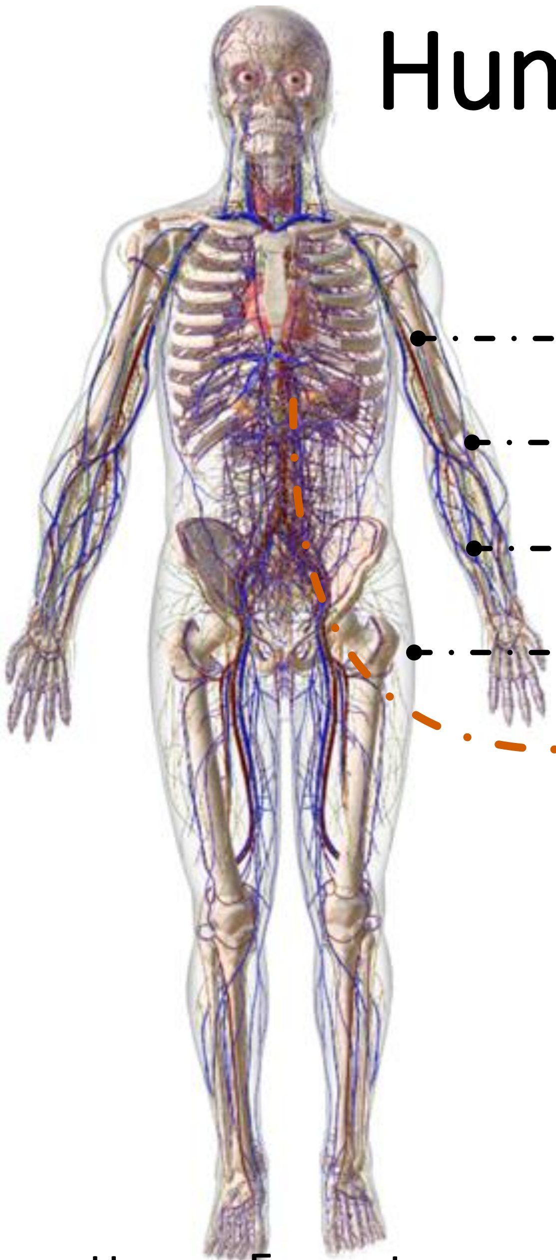
Human Ecosystem

The human body is comprised of a series of complex systems, including:

- Skeletal System
- Nervous System
- Cardiovascular System
- Lymphatic System
- Endocrine System

→ Infrastructure Systems

→ Regulation Systems

The human body is comprised of a series of complex systems, including:

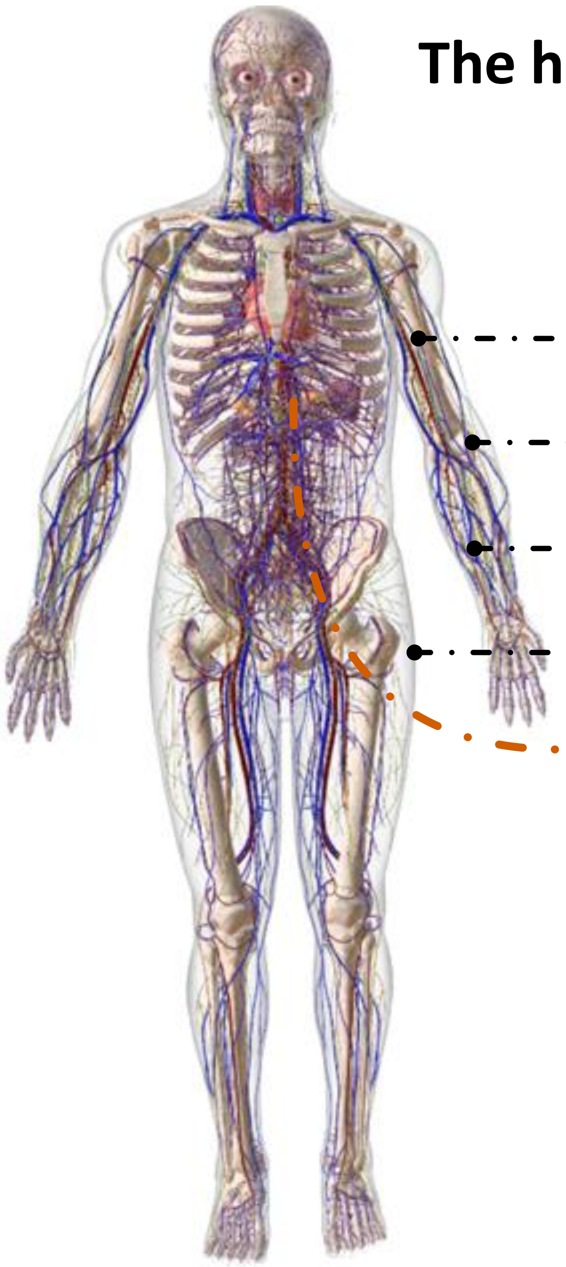

Skeletal System

Nervous System

Cardiovascular System

Lymphatic System

Endocrine System

→ Infrastructure Systems

→ Regulation Systems

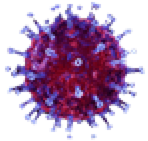

- Part of the immune system
- Protects your body against foreign invaders

Distributed Computing Continuum Systems are composite complex systems

Infrastructure Systems

Skeletal System
Nervous System
Cardiovascular System
⋮

- Devices & Sensors
- Connection & Communication
- Data Flow
- Learning & Knowledge

Regulation Systems

- Elasticity Systems
- Self-adaptive Systems
- Fault-tolerance Systems
- Privacy & Security

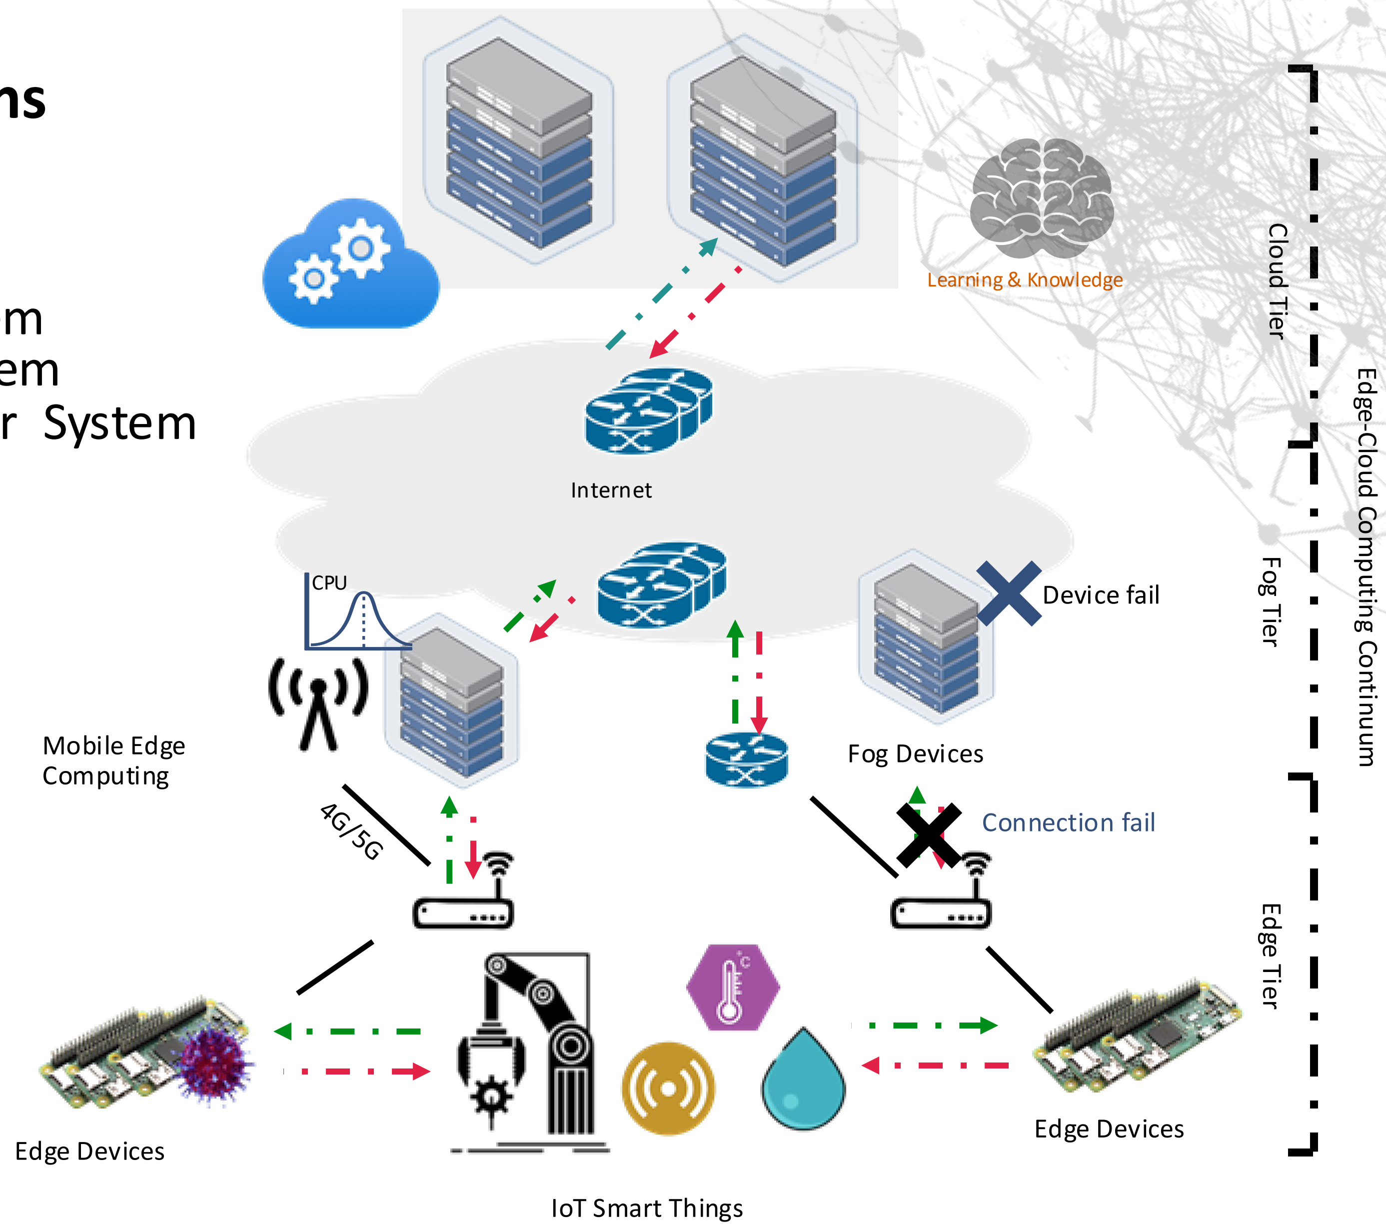

Human Ecosystem vs. Distributed Continuum Systems

Human Ecosystem

Feedback Loops and adaptability

Communication and coordination with other parts

Decision-making

Consciousness

Self-healing mechanism

Distributed Computing Continuum Systems

QoE & QoS

Handle *unpredictable* issues

Synchronization and Coordination

Design intelligent data protocols[2]

Efficient and Adaptable Resource management

Explainability through Causal relations

Monitoring and self-healing devices[1]

[1]P.K., Donta, et al. "Governance and sustainability of distributed continuum systems: A big data approach." *Journal of Big Data* 10.1 (2023): 53.
[2]P.K., Donta and S. Dustdar "Intelligent Data protocols at the Edge", IEEE SERVICES 2023.

Computing Continuum Monitoring, Diagnosis and Cure

Diagnosis and Cure

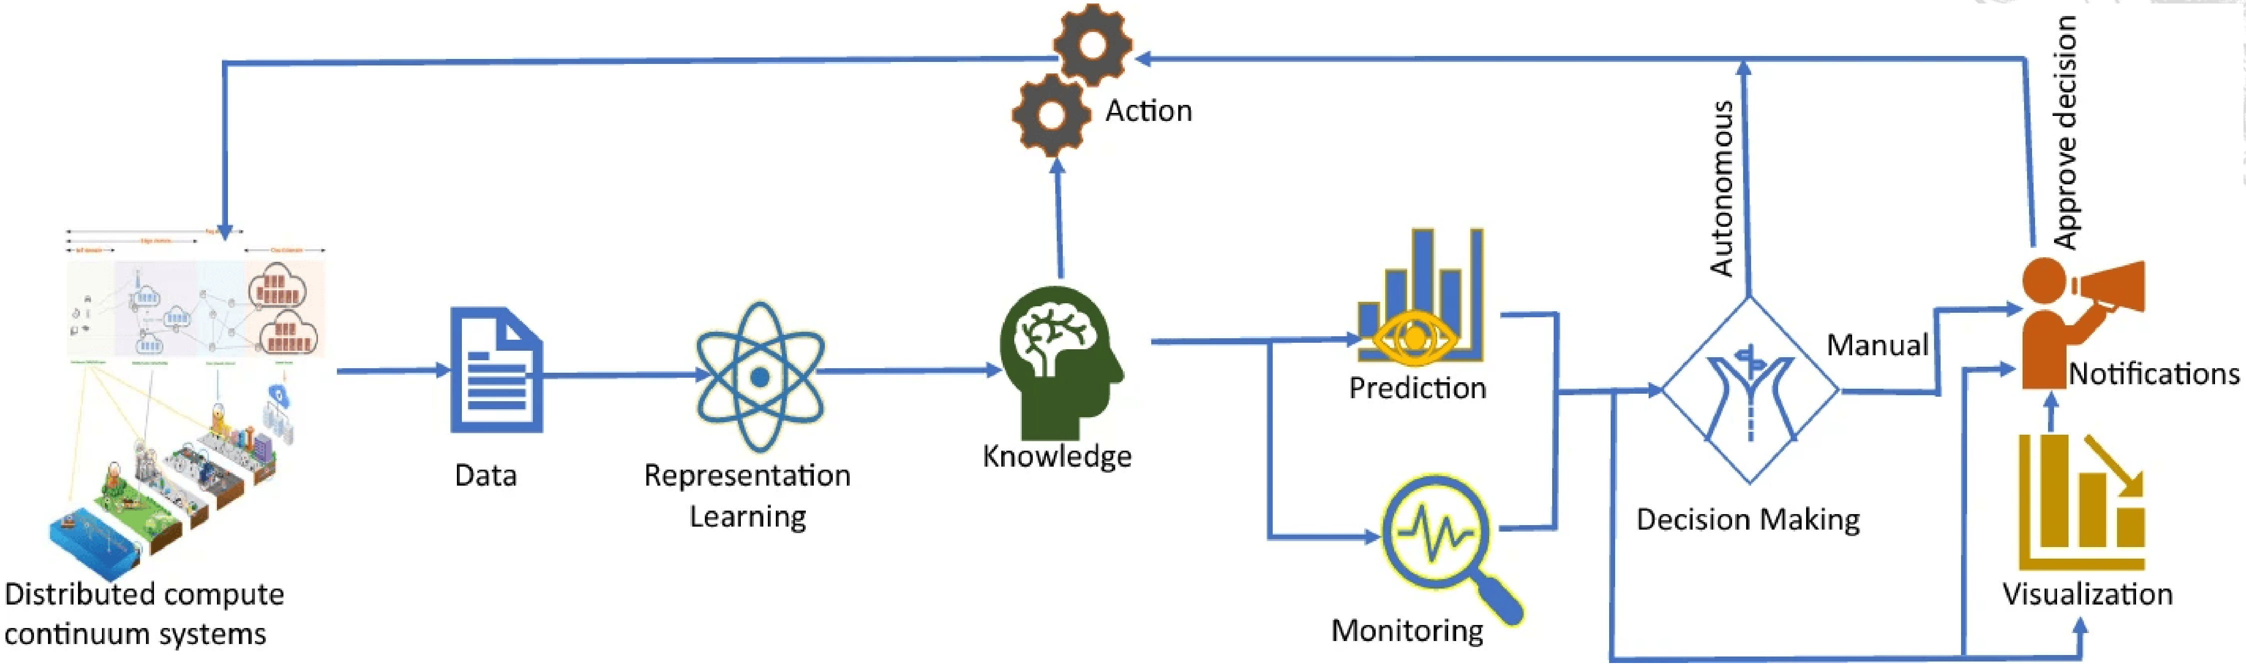

P.K., Donta, et al. "Governance and sustainability of distributed continuum systems: A big data approach." *Journal of Big Data* 10.1 (2023): 53.

Self-healing Agent

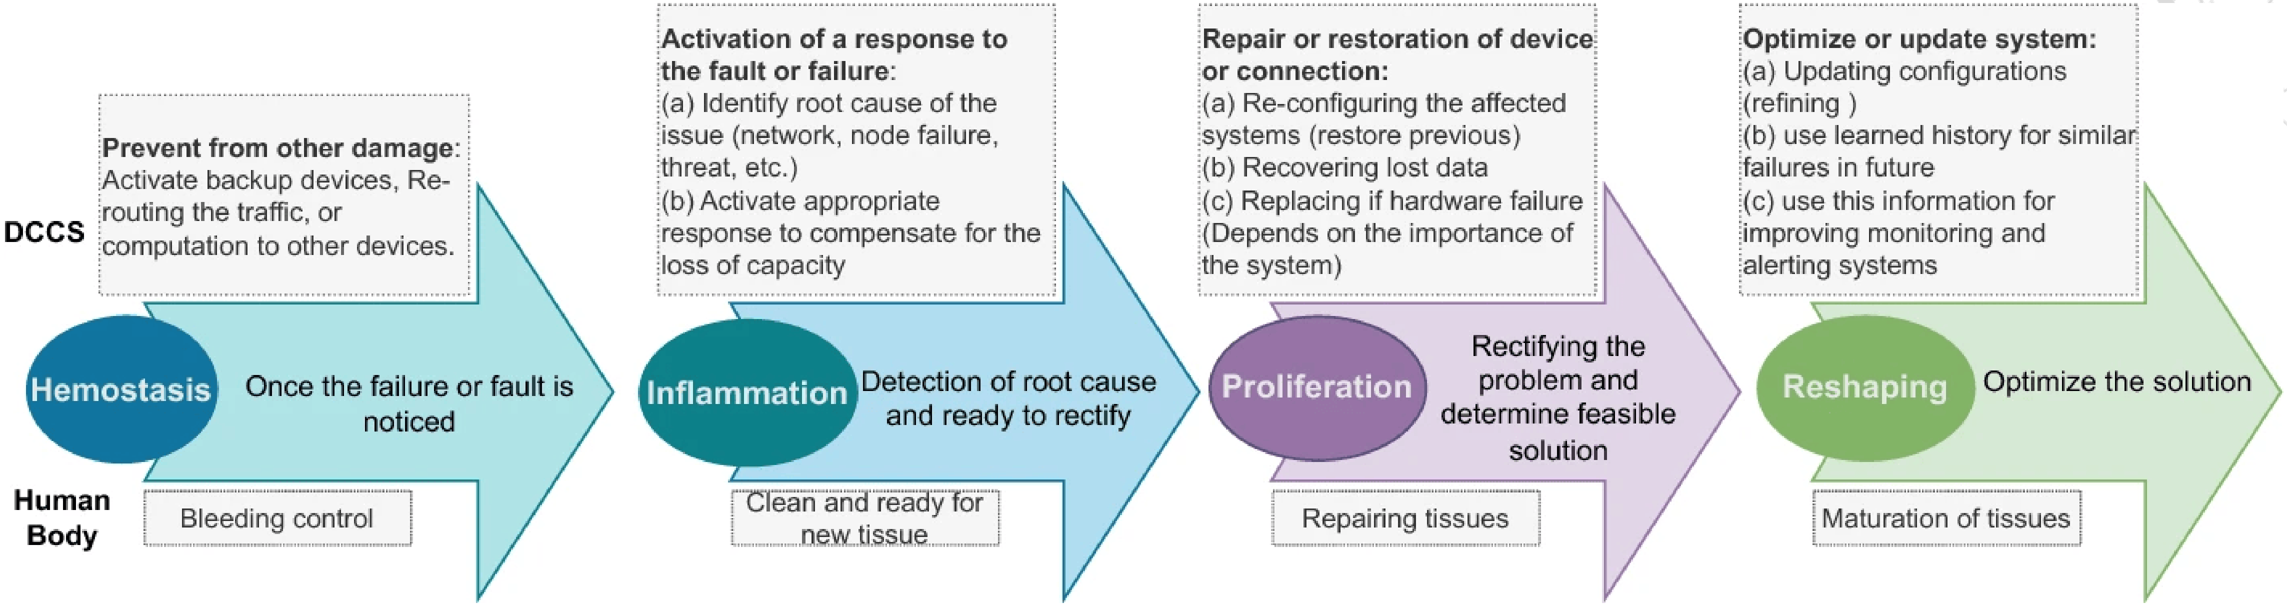

P.K., Donta, et al. "Governance and sustainability of distributed continuum systems: A big data approach." *Journal of Big Data* 10.1 (2023): 53.

Data Protocols and Intelligence

- P. K. Donta and S. Dustdar (2023) Towards Intelligent Data Protocols for the Edge, IEEE intelligent EDGE 2023 (IEEE SERVICES 2023)
- Donta, P. K., Srirama, S. N., Amgoth, T., & Annavarapu, C. S. R. (2022). Survey on recent advances in IoT application layer protocols and machine learning scope for research directions. *Digital Communications and Networks*, *8*(5), 727-744.

Data Protocols

CoAP (Constrained Application Protocol)

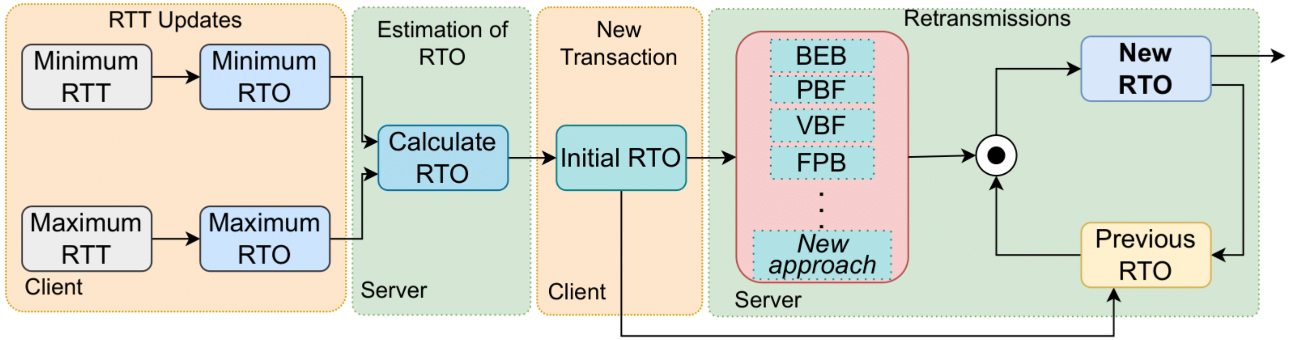

MQTT (Message Queue Telemetry Transport)

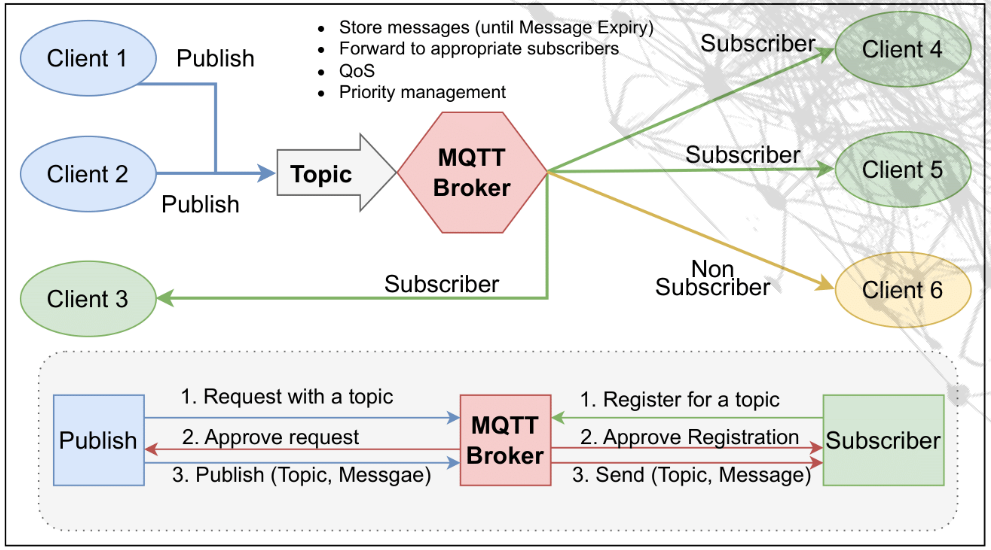

AMQP (Advanced Message Queuing Protocol)

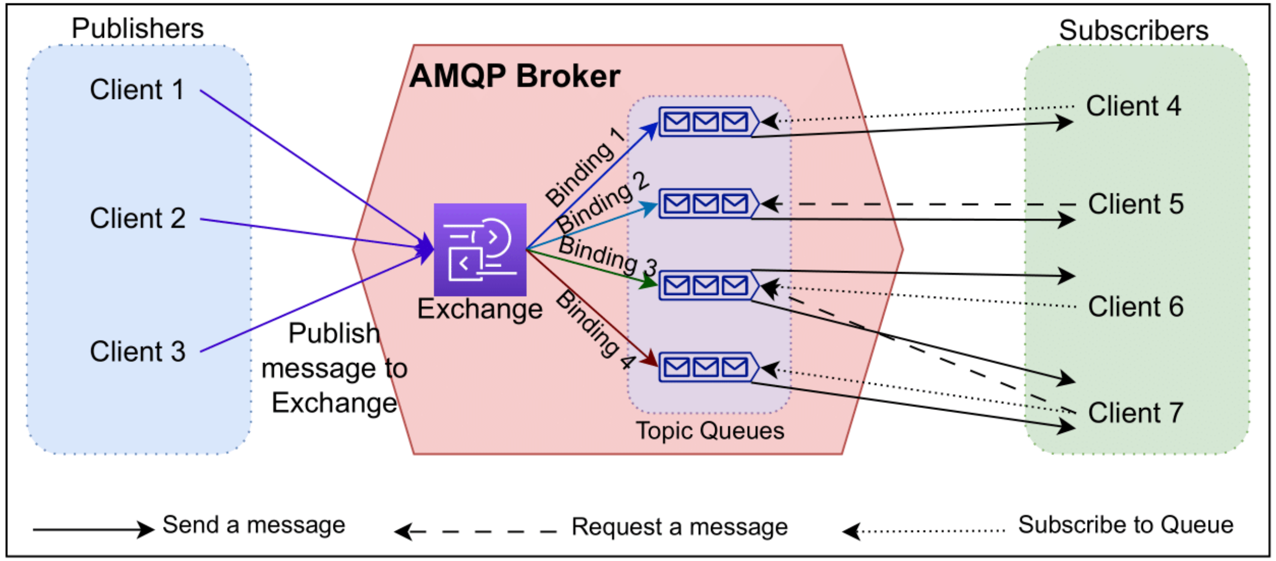

DDS (Data Distribution Service)

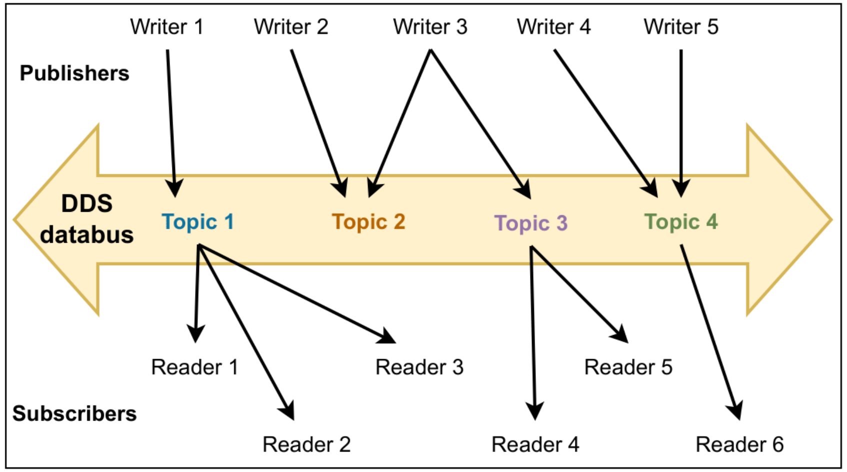

P. K. Donta and S. Dustdar (2023) Towards Intelligent Data Protocols for the Edge, IEEE intelligent EDGE 2023 (IEEE SERVICES 2023)

Deterministic Rules – Data Protocols

- Protocols follow deterministic set of rules
- **Benefits**:
  - Low overhead
  - Limited computing or memory resources
  - Easy to track the misconducts
  - Predictability
- **Limitations**
  - Flexibility and Adaptability
  - Handling uncertainty
  - Scalability
  - Dynamic decision making systems

Intelligence in Data Protocols

Benefits

- Scalability
- Adaptability
- Handling uncertainty
- Dynamic decision making, such as
  - Message filtering
  - Message expiry prediction
  - Resource usage (memory and CPU)
  - Congestion control
- Reliability

Implications

- Need of additional resources (to store history and perform analysis)
- Decide where learning to be performed
- Deciding run time accuracy
- Deciding suitable AI/ML for the protocols properties

Develop *Novel* Intelligent Data Protocol

- We intended to develop an lightweight intelligent protocol for Computing continuum systems by mitigating the following issues:
  - Message filtering
  - Message expiry prediction
  - Resource usage (memory and CPU)
  - Congestion control
- Adaptable and
- Interoperable features

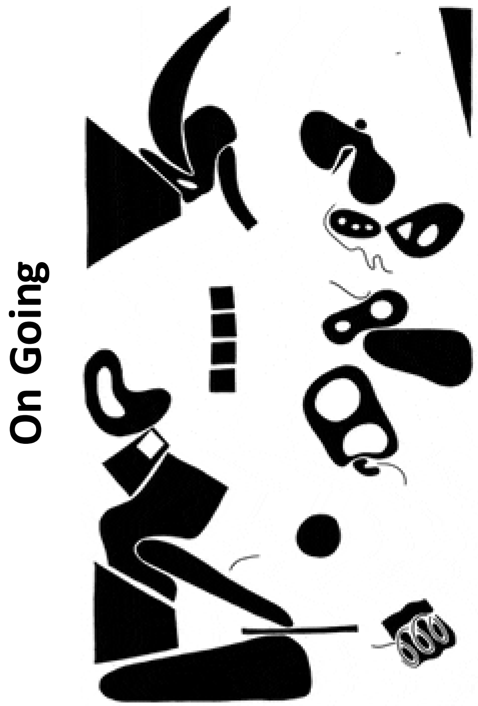

Image source: https://giphy.com/gifs/traceloops-ongoing-Oj7pPn80t5Js3pFyfM

Equilibrium in Computing Continuum through Active Inference

Sedlak, B., Pujol, V. C., Donta, P. K., & Dustdar, S. (2024). Equilibrium in the Computing Continuum through Active Inference. *Future Generation Computer Systems*.

Bayesian Network Learning

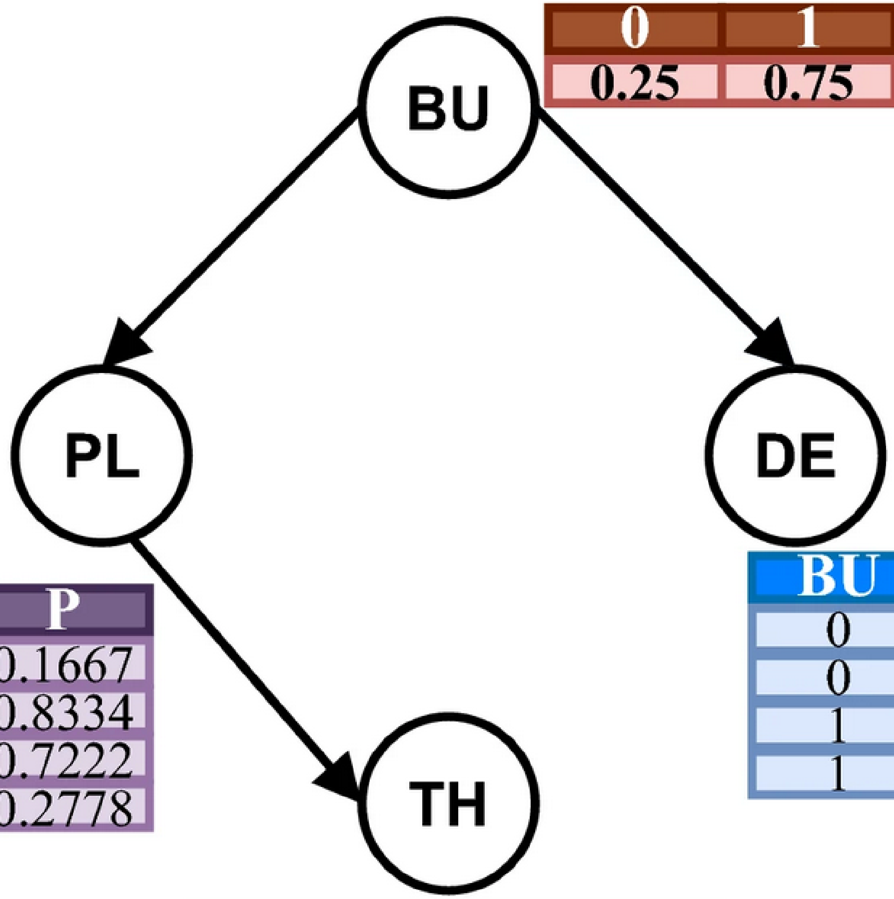

0	1
0.25	0.75

BU	PL	P
0	0	0.1667
0	1	0.8334
1	0	0.7222
1	1	0.2778

BU	De	P
0	0	0.1667
0	1	0.8334
1	0	0.7222
1	1	0.2778

PL	Th	P
0	0	0.0714
0	1	0.9286
1	0	0.5
1	1	0.5

Causality

Causality is the study of how things influence one other, how causes lead to effects.

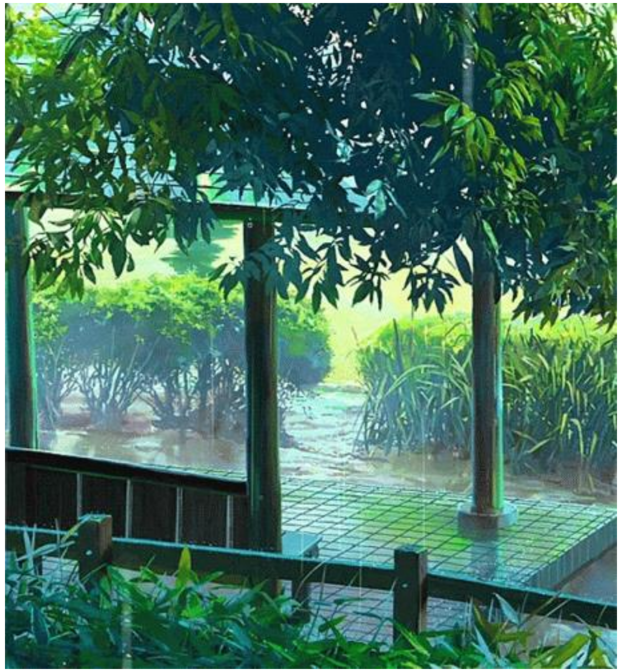

Surprises are generally created by low-probability outcomes

Image sources: https://www.btelligent.com/en/blog/127801ea523cbf313725fcb60289fbf8-4/
https://ru.pinterest.com/pin/489555421981720364/
https://imgflip.com/gif/6keaol

Free Energy

The discrepancy (or uncertainty) between the agent's understanding of the process and reality is called *Free Energy* (FE)

Lower the *Free Energy* (FE), higher the accuracy.

This is commonly known as **Bayesian inference** and allows agents to use existing beliefs (widely known as priors) to calculate the probability of related events.

Image source: https://tenor.com/view/rain-umbrella-thor-gif-17980481

Markov Blanket

The Markov Blanket of a random variable is the subset of nodes that provide enough information to statistically infer its value. Concept from Judea Pearl [1].

In a Bayesian Network, the Markov Blanket of a node (N) is composed of the parents (P), the children (C) and the co-parents of the children (S).

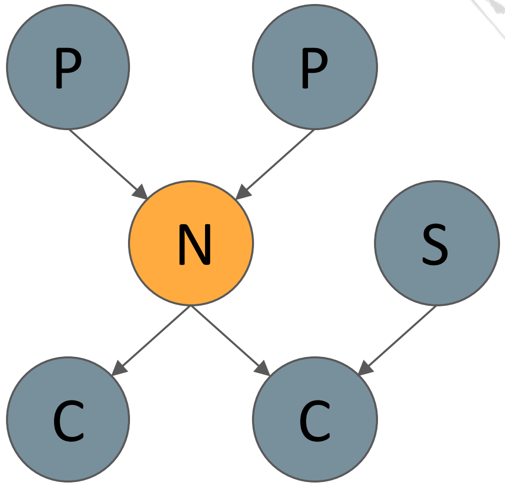

A tool for *causal* filtering.

[1] Pearl, J. (1988). Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers Inc.

Causal Inference

Causal inference is a technique aiming at learning ***causal relationships*** between ***variables***, and then being able to *model and predict* the system causal behavior.

- 3 Rungs on the ladder of causation. [2]
  - Observational
  - Interventional
  - Counterfactual
- Explainability capacity

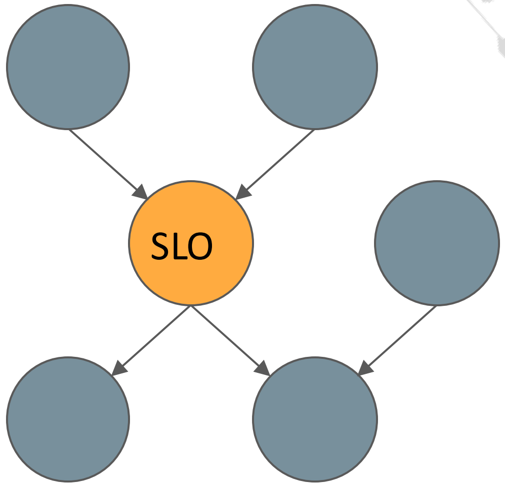

[2] J. Pearl and D. Mackenzie (2018), The Book of Why: The New Science of Cause and Effect. USA: Basic Books, Inc.,

Active Inference

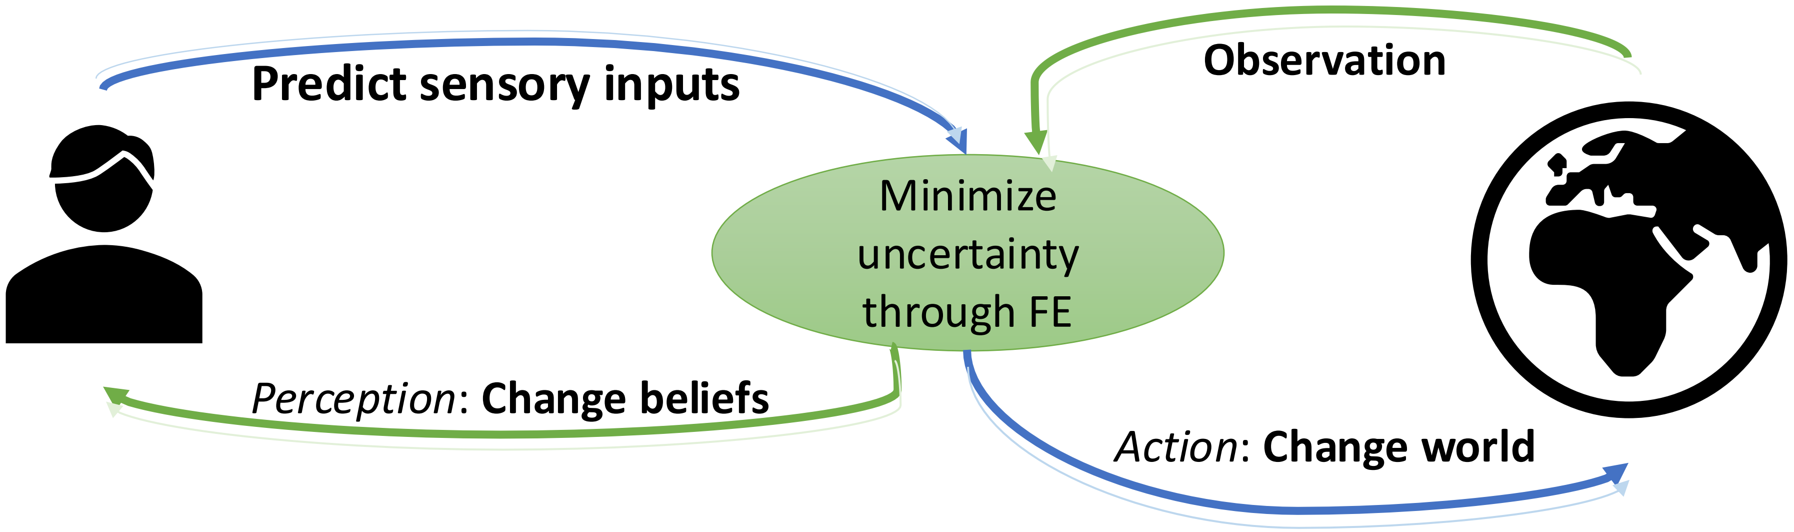

Figure: Action – Perception Cycle

Service Level Objective - Type

whether a system is able to fulfill its intended function at a point in time

Three types of SLOs

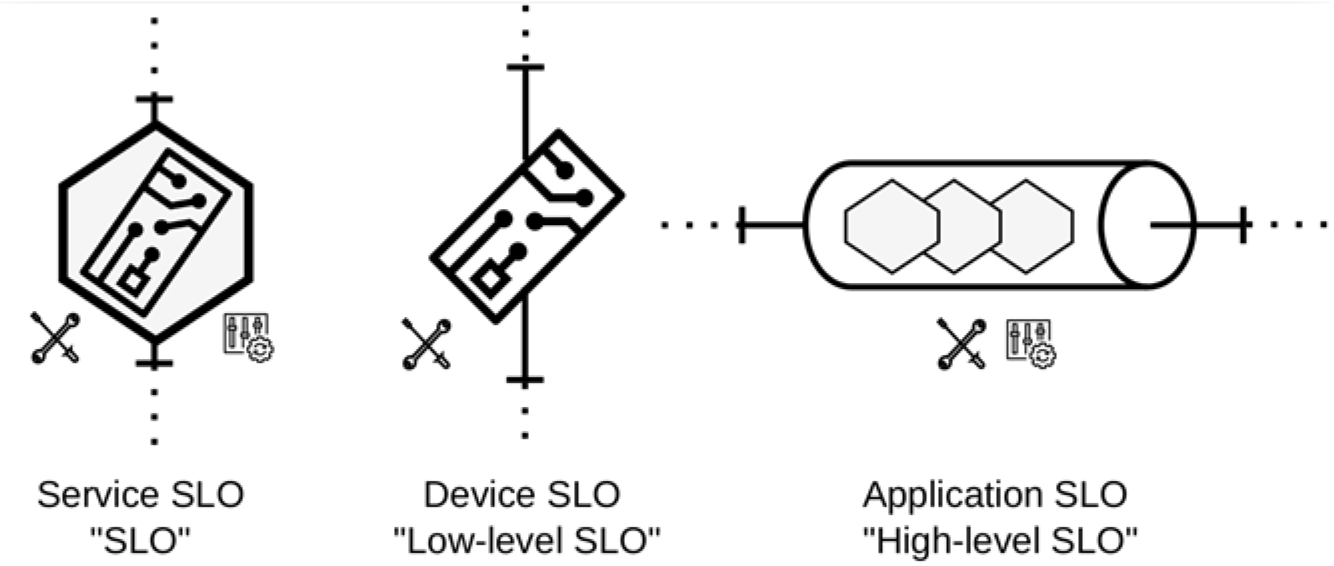

Service Level Objective as system requirements

Components: Service + host + SLO + elasticity strategies
Tailored adaptations to service and host

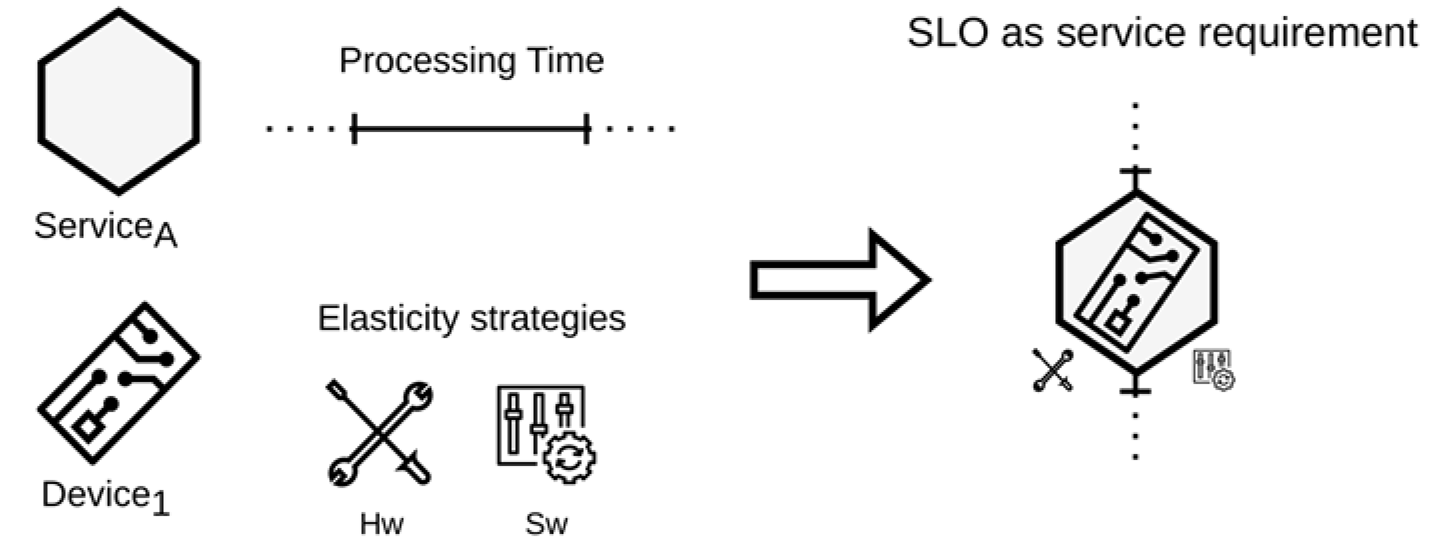

DeepSLOs

- A construct we envision relating SLOs
- Provides a complete view of DCC system
- Allows aggregation towards higher abstractions

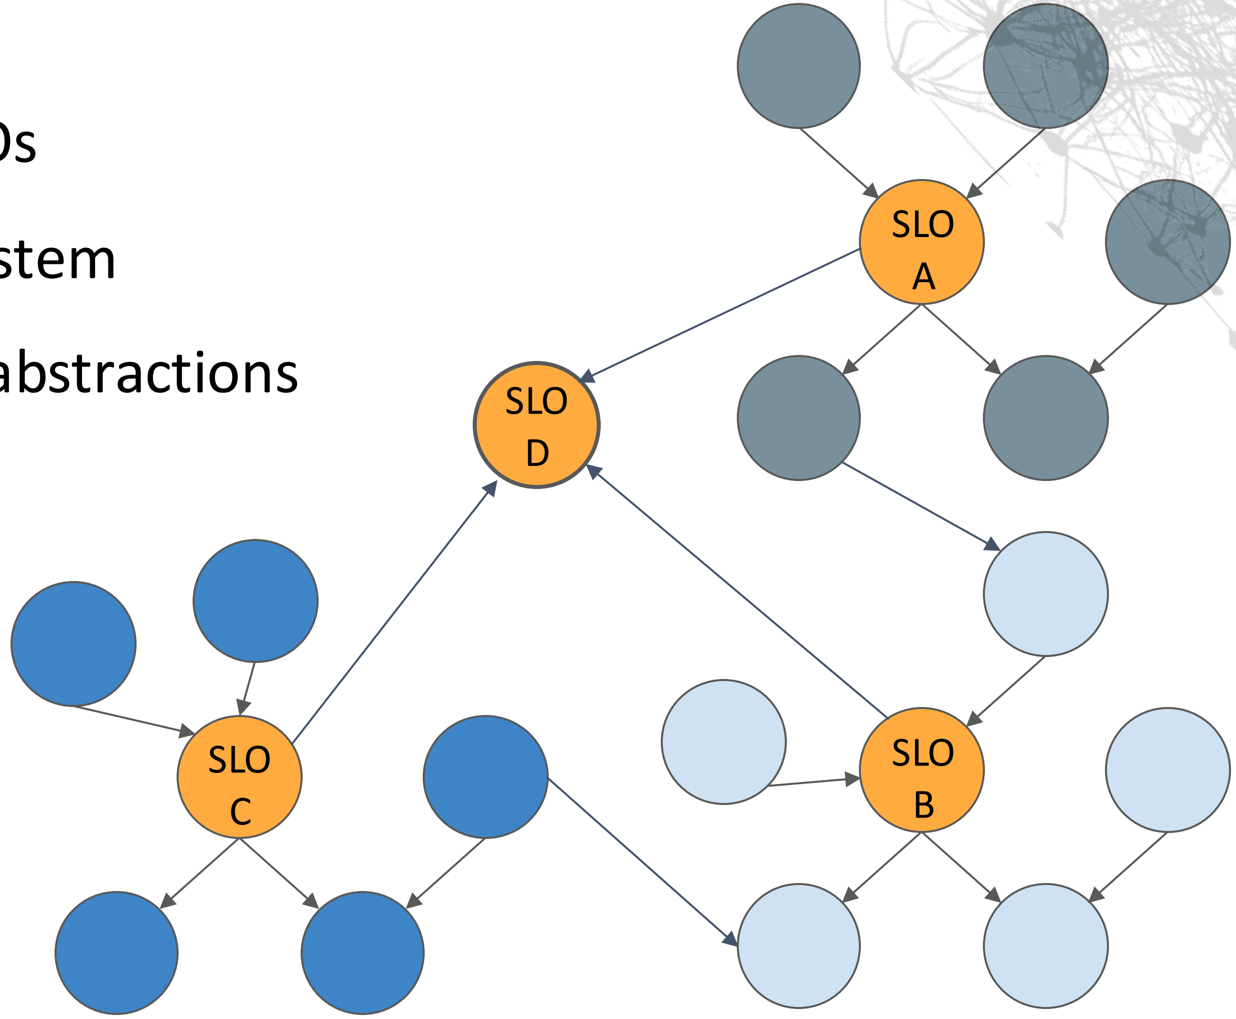

Running Example

- Reflected in most of the architecture

- Use Case
  Distributed video processing architecture where IoT streams are transformed on **edge devices** to preserve individual's privacy. After privacy enforcement, distribute streams over **cloud**.

- Hierarchical network structure
  IoT devices provide streams to edge devices; streams processed locally at edge devices; video stream properties are **configurable**

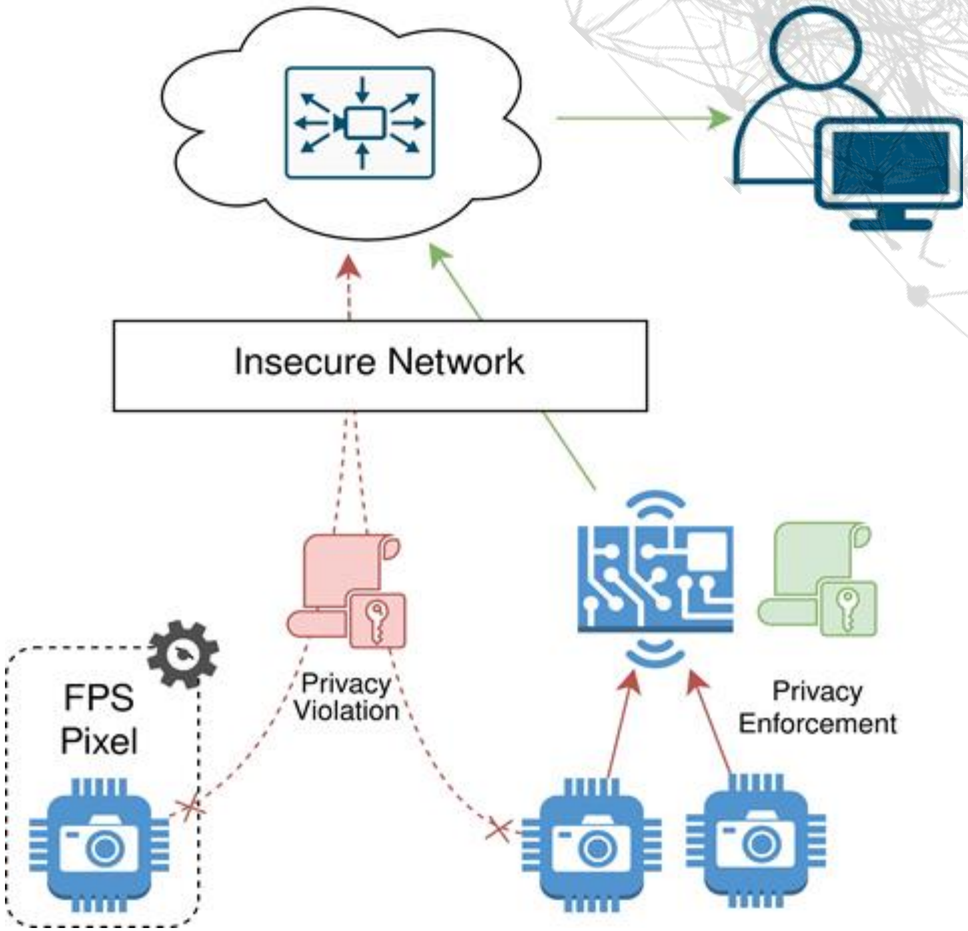

Collaborative Edge Intelligence Framework

3 major contributions in interplay:

1. Continuous model accuracy and local SLO fulfilment
    a. Static BNL and Inference
    b. Continuous BNL and Inference (**AIF**)

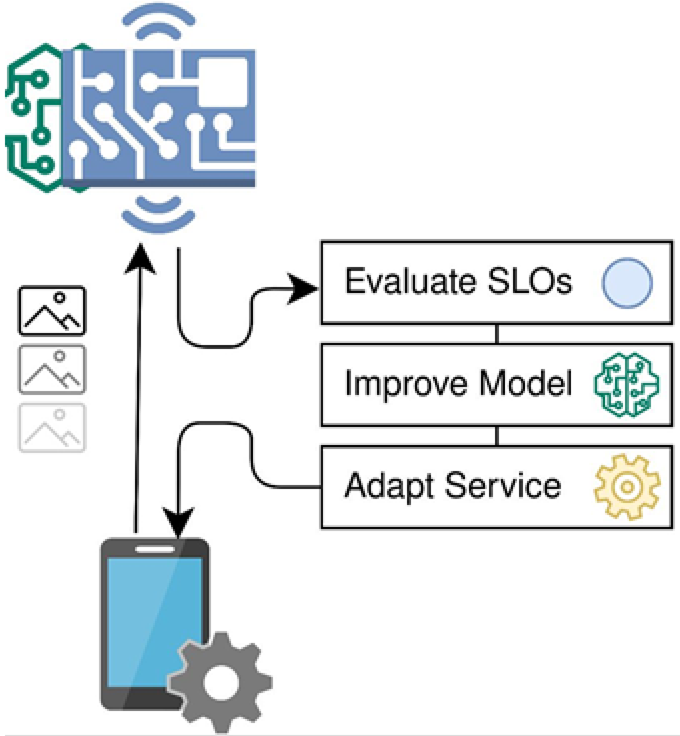

Collaborative Edge Intelligence Framework

3 major contributions in interplay:

1. Continuous model accuracy and local SLO fulfilment
    a. Static BNL and Inference
    b. Continuous BNL and Inference (**AIF**)
2. Federation and combination of models

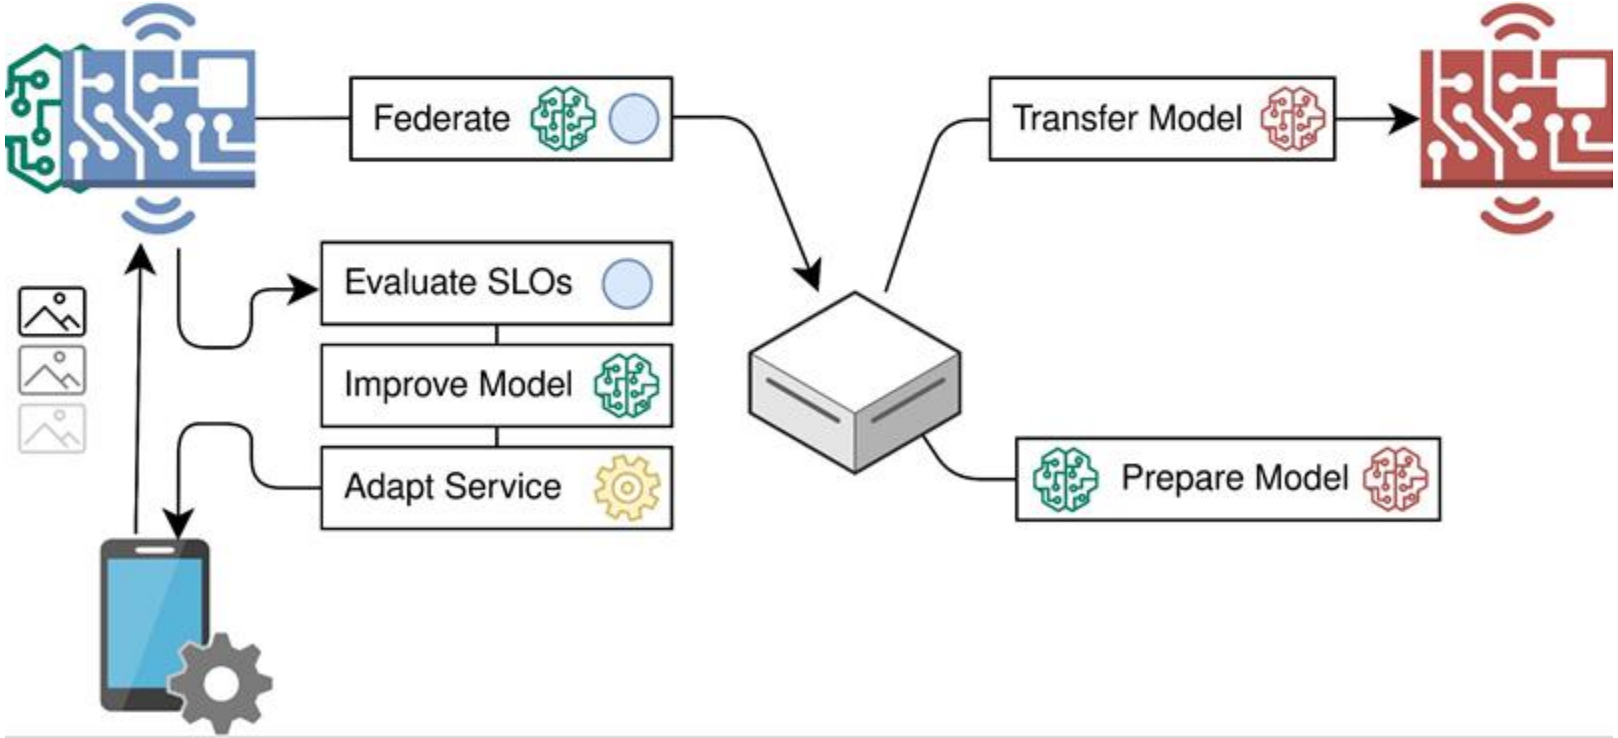

Collaborative Edge Intelligence Framework

3 major contributions in interplay:

1. Continuous model accuracy and local SLO fulfilment
    a. Static BNL and Inference
    b. Continuous BNL and Inference (**AIF**)
2. Federation and combination of models
3. Collaboration between cellular structures

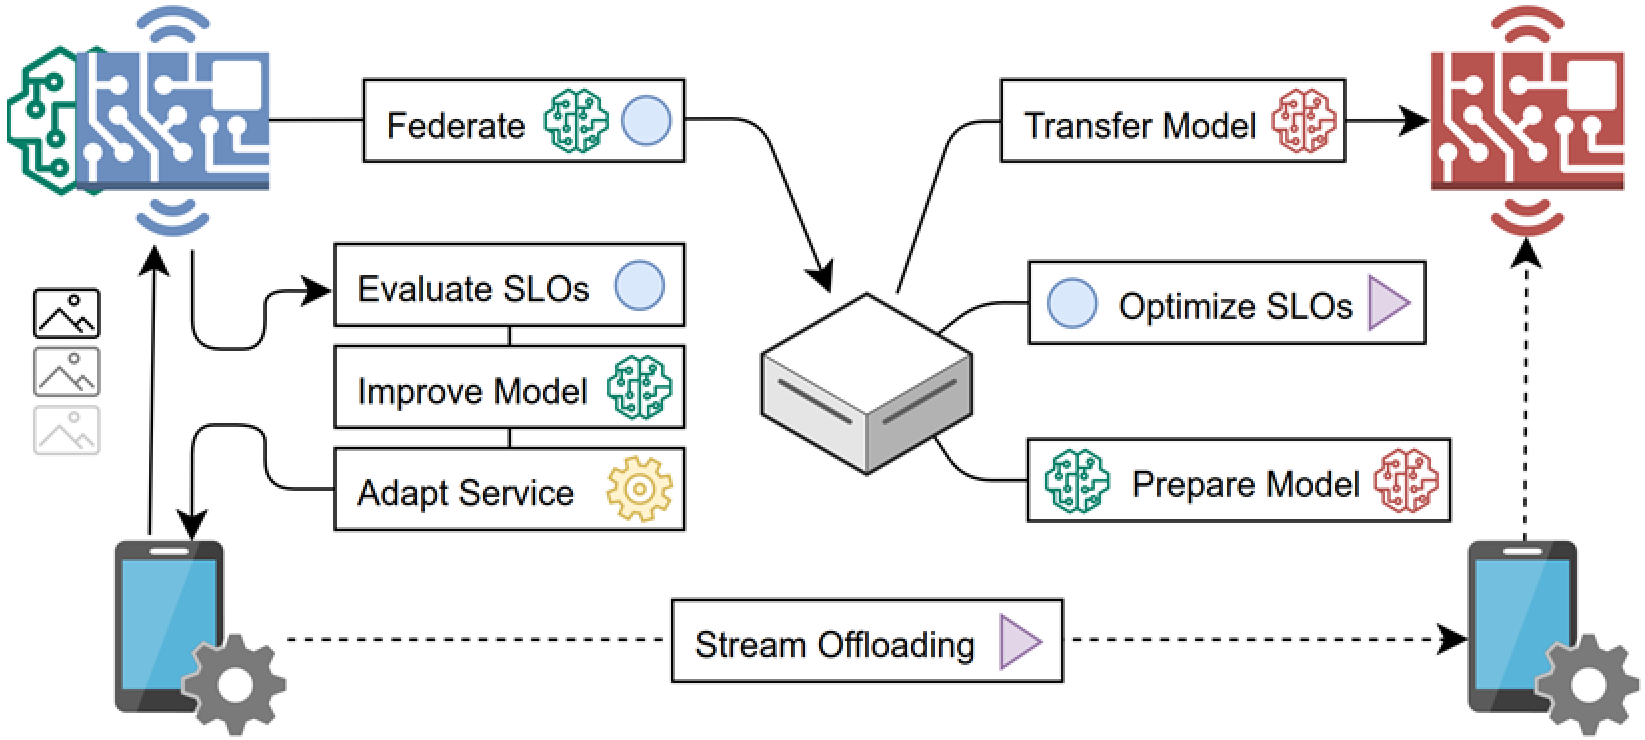

1a – Static BNL and Inference

- Basic mechanism for assuring SLOs at individual devices
- **Requires training data in upfront and is prone to data shifts**
- Evaluates possible configurations through a 3-step method

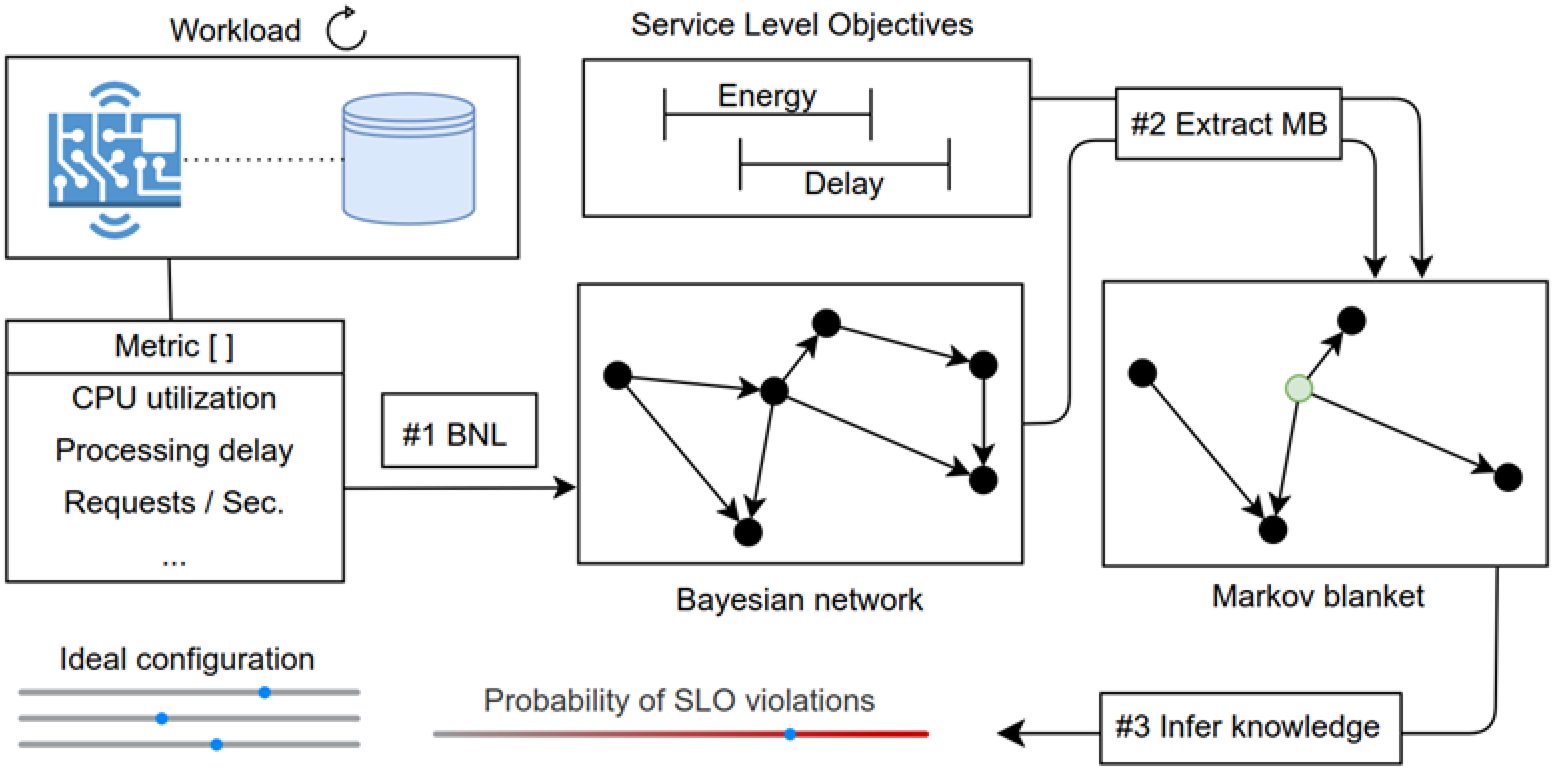

1a – Static BNL and Inference (2)

Bayesian Network Learning (BNL)

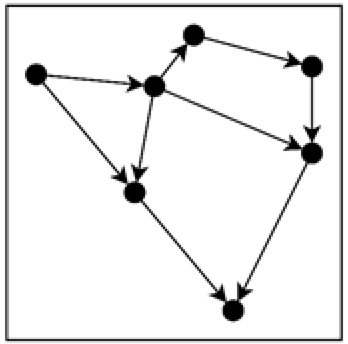
Bayesian network

- ❑ **Structure Learning**
  Hill-Climb Search (HCS)
  Dir. Acyclic Graph (DAG)
- ❑ **Parameter Learning**
  Max. Likelihood Estimation
  Conditional Prob. Table (CPT)

Markov Blanket (MB) Selection

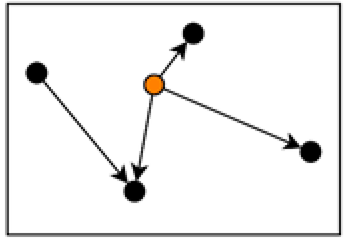
Markov blanket

- ❑ Causality filter [1,4]
- ❑ Identify variables that have an impact on **SLO fulfillment**

Knowledge Extraction

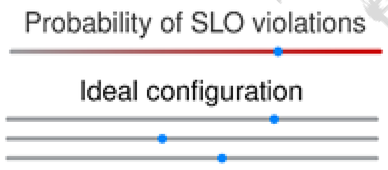

- ❑ P(SLO < x) for all variable combinations
- ❑ Find **Bayes-optimal** system configuration

1b – Active Inference Cycle

- **AIF agent** → Equilibrium-Oriented SLO Compliance (**EOSC**) model
- Agent uses SLOs as **preferences** during continuous adaptation
- BN trained incrementally from incoming observations
- Beliefs updated according to prediction errors

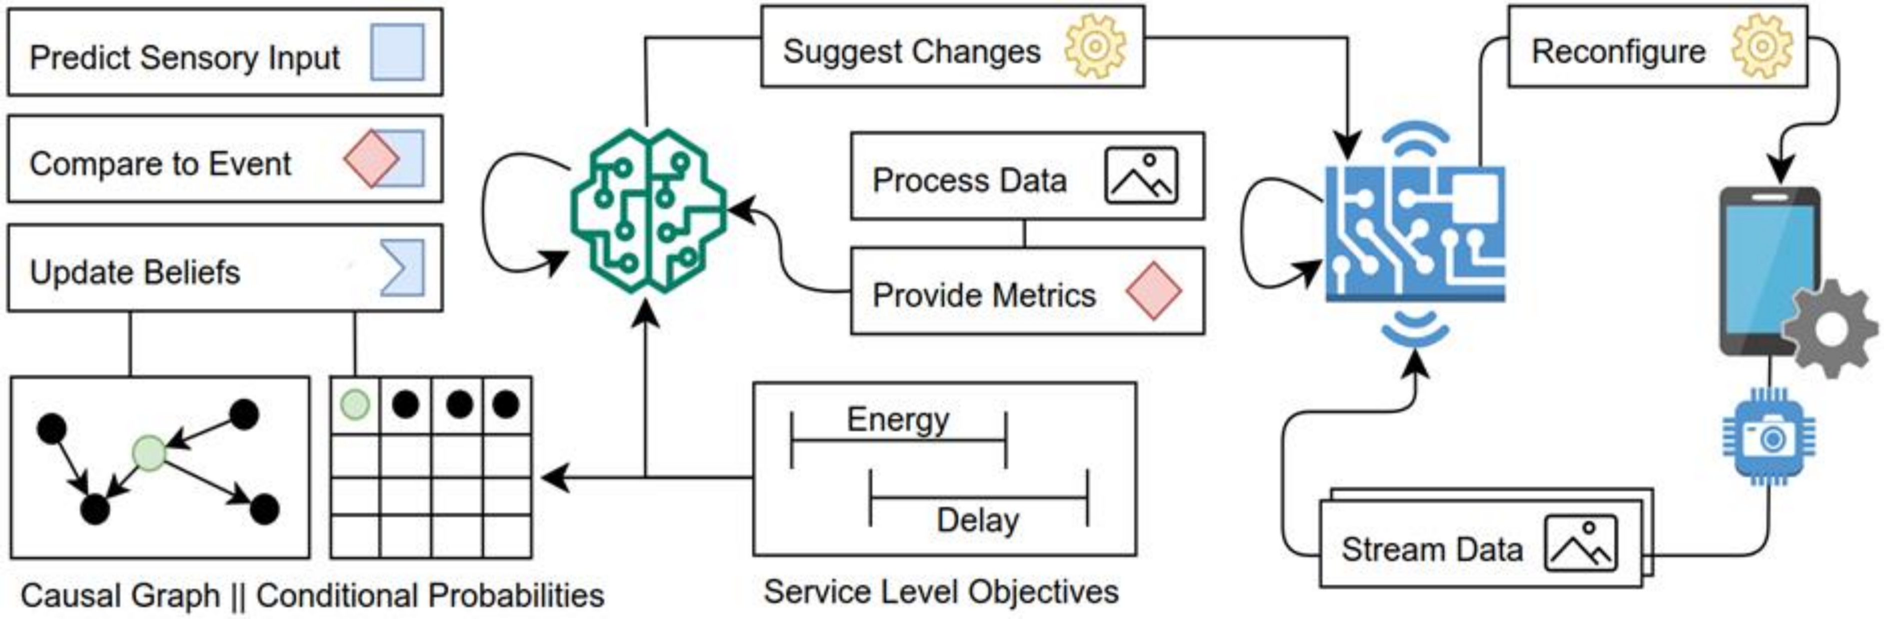

1b – AIF Agent Behaviour/Action

Determined by three factors:

- Pragmatic value (*pv*)
  Summarizes **QoE** SLOs (e.g., resolution)
- Risk assigned (*ra*)
  Summarizes **QoS** SLOs (e.g., network limit)

pv & *ra* calculated as **separate factors** from MBs; configurations rated according to SLO fulfilment; **interpolation** between known configurations

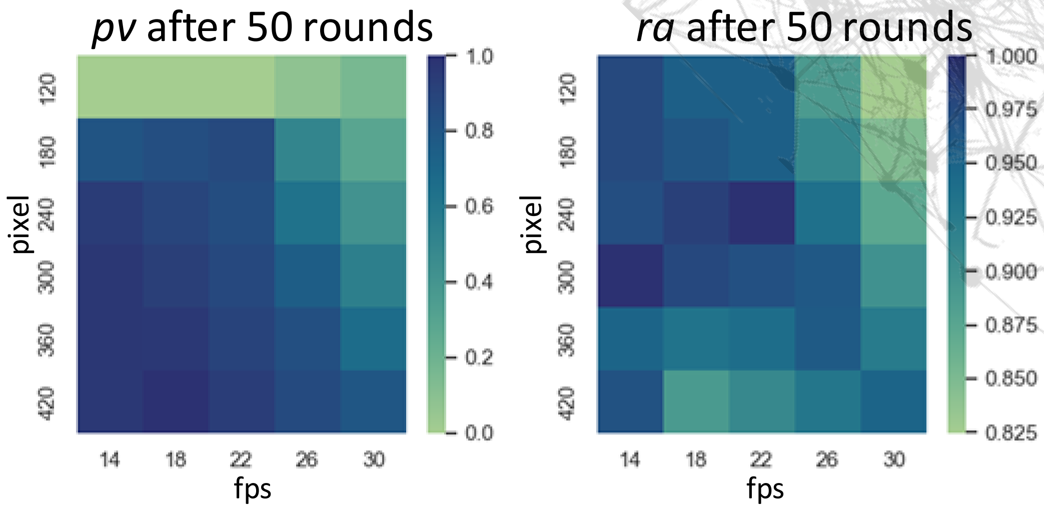

- Information gain (*ig*)
  Continued on the next slide

$$u_c = pv_c + ra_c + ig_c$$

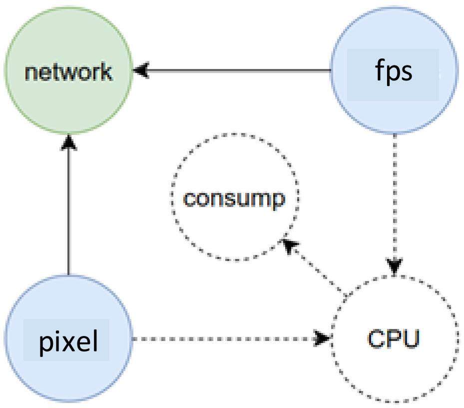

1b – AIF Agent Behaviour/Action (2)

- Information gain (*ig)*
  - Favours configurations that promise **model improvement**
  - Summarizes surprise for observations included in the **MB**
  - Hyperparameter (*e*) allows exploring designated areas

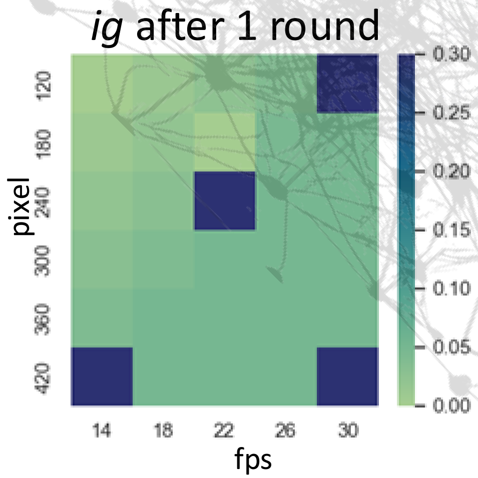

AIF *Action-Perception* cycle:

1. Calculate **surprise** for current batch of observations
2. Retrain structure (or parameters) depending on surprise
3. Calculate behavioural factor for **empirically evaluated** configs
4. **Interpolate** between known configurations in 2D (or 3D) space
5. Choose the highest-scoring (device) configuration

Agent gradually develops **understanding** how to ensure SLOs

$$ig(c) = e + \left(\frac{\tilde{\mathfrak{S}}_c}{\bar{\mathfrak{S}}}\right) \times 100$$

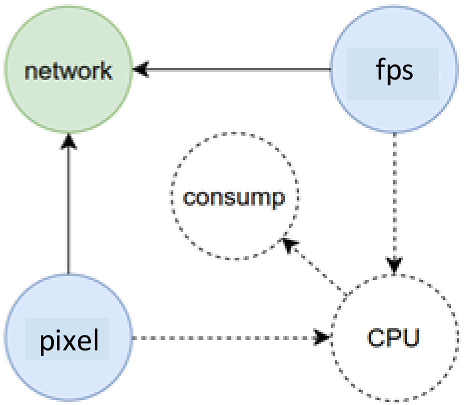

2 – Knowledge Exchange

Extend from single devices to the CC

Heterogeneity among the Edge

- Impedes simple transfer learning of models
- Low model accuracy → high surprise

- Requires a **cluster leader** (fog node or edge)
- EOSC models collected at a leader node
- Model selection according to hardware char.
- Merging models to provide tailor-fit one

Fast onboarding (= horizontal scaling) of devices

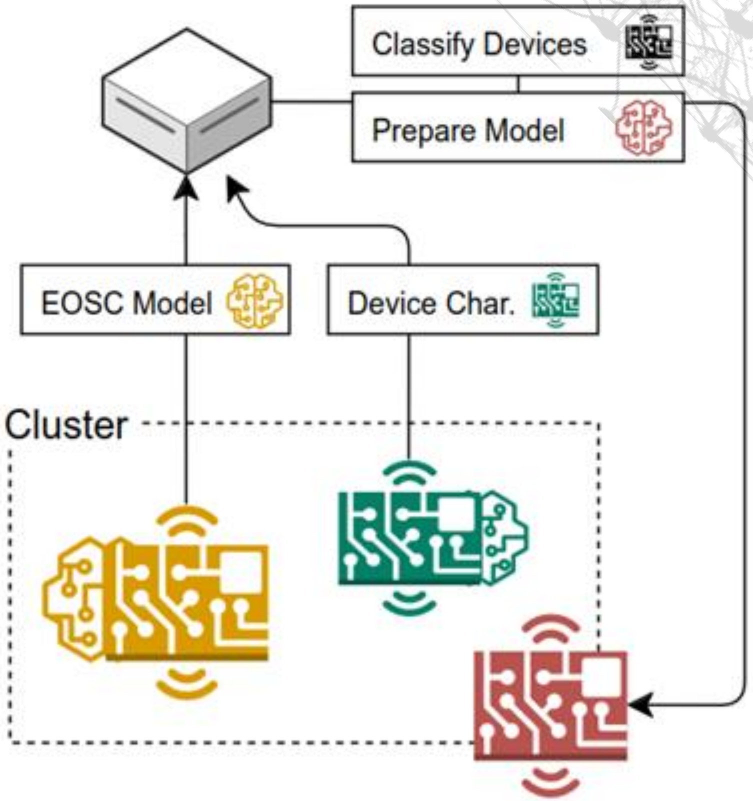

3 – Collaborative Scaling

Limited action scope of devices

- Individual devices restricted to local scope to resolve SLO violations

- Leader node collecting **environmental metrics** (e.g., network congestion)
- Incorporated to causal model, contrasted against local SLO fulfilment (**AIF**)
- Emerging structures allows optimizing cluster-wide SLO fulfilment
  - E.g., redistribute clients between impacted devices

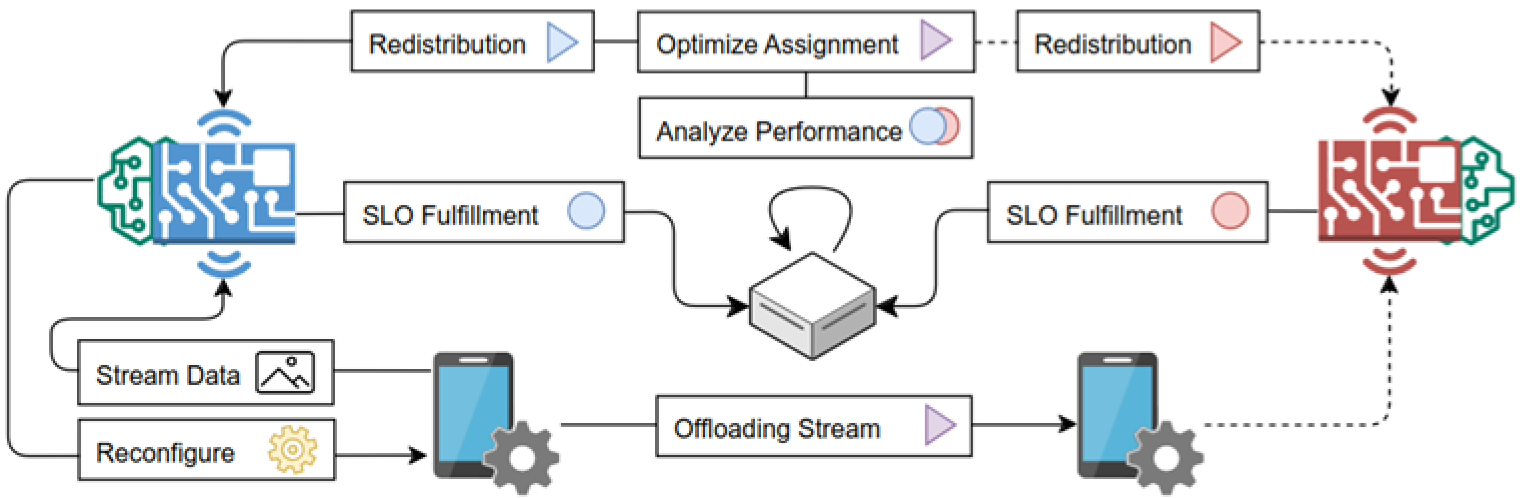

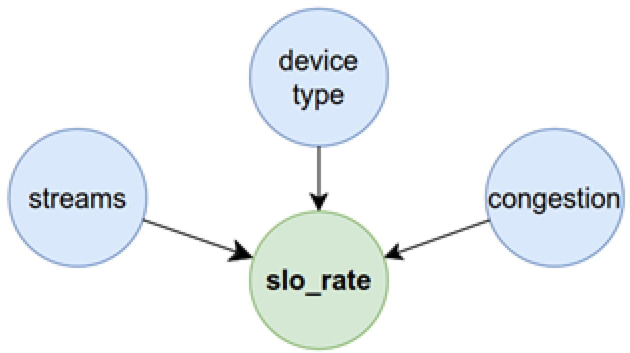

Evaluation - Overview

- Use Case
  Distributed video processing architecture where streams are transformed on **edge devices** to preserve privacy of individuals.

- Implementation
  Prototype including video transformations and the collaborative edge intelligence framework.

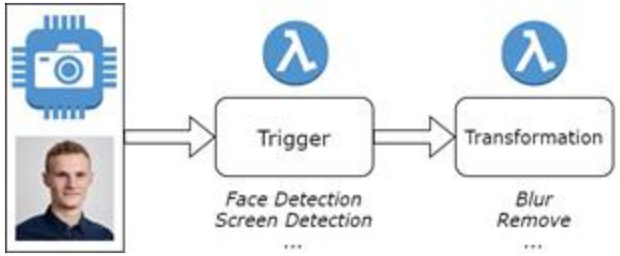

[5]

- Evaluation Scope
  Targeting each contribution with different aspects.

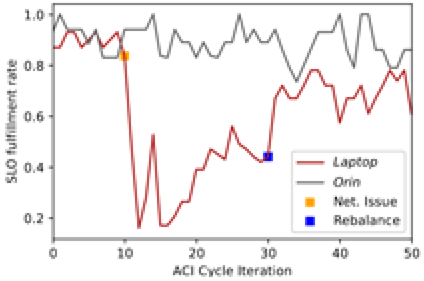

[5] Sedlak, B., Murturi, I., Donta, P.K., Dustdar, S.: A Privacy Enforcing Framework for Transforming Data Streams on the Edge. IEEE Emerging Topics in Computing (2023)

Evaluation - Use Case

BNL comprises metrics from various sources (e.g., IoT client or edge device);
Extended with target conditions (i.e., SLOs) to create the **EOSC** model:

Model training takes 11 (3) metrics

Table 1: List of metrics captured by the devices, which are turned into variables by ACI

Name	Origin	Unit	Description	Param
pixel	IoT	num	number of pixel contained in a frame	Edge
fps	IoT	num	number of frames received per second	Edge
bitrate	IoT	num	number of pixels transferred per second	No
cpu	Edge	%	utilization of the device CPU	No
memory	Edge	%	utilization of the system memory	No
streams	Edge	num	number of IoT devices providing data	Fog
consumption	Edge	W	energy pulled by the device	No
network	Edge	num	data transferred over network interface	No
delay	App.	ms	processing time per video frame	No
success	App.	T/F	if a pattern (i.e., face) was detected	No
distance	App.	num	relative object distance between frames	No
slo_rate	Edge	%	combined SLO Fulfillment rate ($pv \times ra$)	No
device_type	Edge	enum	physical device type	No
congestion	Edge	num	network congestion that increases latency	No

SLOs from model variables

Table 2: Extracted SLOs and their classification.

SLO	Condition	Tier	Type
network	$network < 1.6$ MB/s	Edge	QoS
in_time	$delay < 1/fps$	Edge	QoS
success	$success = True$	Edge	QoE
distance	$distance < 50$	Edge	QoE
slo_rate	$\max(slo\_rate)$	Fog	Both

Parameters allow **configuring** a component's environment

Evaluation - Implementation

Python prototype for which we provide:

- Github repository[1]
- Docker container[2]

Evaluation included a variety of edge devices:

bedded-systems/jetson-xavier-nx/

Table 3: List of devices used for implementing and evaluating the presented methodology

Full Device Name	ID	Price	CPU	RAM	GPU	p [1,4]	g [0,2]	Σ
ThinkPad X1 Gen 10	*Laptop*	1800 €	Intel i7-1260P (16 core)	32 GB	———	Very High (4)	None (0)	4
Nvidia Jetson Orin	*Orin*	500 €	ARM Cortex A78 (6 core)	8 GB	Volta (383 core)	High (3)	High (2)	5
Nvidia Jetson Nano	*Nano*	150 €	ARM Cortex A57 (4 core)	4 GB	———	Low (1)	None (0)	1
Nvidia Jetson Xavier	*Xavier*	300 €	ARM Carmel v8.2 (6 core)	8 GB	———	Medium (2)	None (0)	2
Jetson NX GPU	*NX*	300 €	ARM Carmel v8.2 (6 core)	8 GB	Amp (1024 core)	Medium (2)	Low (1)	3

Devices combined within a cluster and classified relatively to each other

[1]https://github.com/borissedlak/workload/tree/main/FGCS
[2]https://hub.docker.com/repository/docker/basta55/workload/

Evaluation - Aspects

We motivated, evaluated, and provided the results for **13** aspects:

A-1: Do MBs reduce the complexity of inference?
A-2: What is AIF's operational overhead?
A-3: How long require AIF agents to ensure SLOs?
A-4-1: Are the produced Bayesian networks interpretable?
A-4-2: Is the behaviour of AIF agents explainable?
A-5: What is the operational impact of including BNL in the AIF cycle?
A-6: Can changes in variable distribution be handled?
A-7: Can SLOs be modified during runtime?

K-1: What is the SLO fulfilment rate of transferred models?
K-2: Can knowledge transfer achieve any speedup?
K-3: Do tailored models have lover surprise compared to existing models?

S-1: How is the load distributed among resource-constrained devices?
S-2: Can intelligent CC structures optimize local SLO fulfilment?

Sedlak et al.. (2023). Equilibrium in the Computing Continuum through Active Inference. ***arXiv e-prints***, arXiv-2311.
https://arxiv.org/abs/2311.16769

Evaluation - Aspects (Filtered)

We motivated, evaluated, and provided the results for 13 aspects:

A-1: Do MBs reduce the complexity of inference?
A-2: What is AIF's operational overhead?
A-3: How long require AIF agents to ensure SLOs?
A-4-1: Are the produced Bayesian networks interpretable?
A-4-2: Is the behaviour of AIF agents explainable?
A-5: What is the operational impact of including BNL in the AIF cycle?
A-6: Can changes in variable distribution be handled?
A-7: Can SLOs be modified during runtime?

K-1: What is the SLO fulfilment rate of transferred models?
K-2: Can knowledge transfer achieve any speedup?
K-3: Do tailored models have lover surprise compared to existing models?

S-1: How is the load distributed among resource-constrained devices?
S-2: Can intelligent CC structures optimize local SLO fulfilment?

A-1: Do MBs reduce the complexity of inference?

- Setup
  Modify the AIF agent to calculate behavior factors (i.e., **surprise,** etc) for a reduced number of SLOs with or without MB

- Result
  Applying MBs reduced the median inference time of 4 SLOs from 197ms to 151ms

- Implication
  MB provided a decreased **system view**

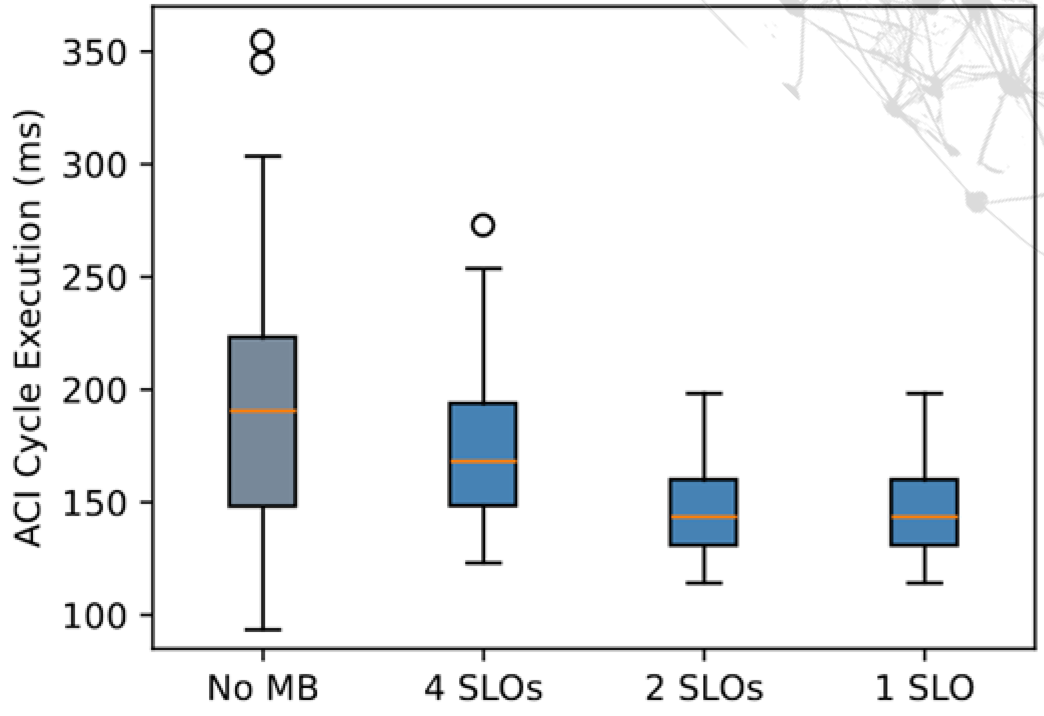

A-4-1: Are the produced Bayesian networks interpretable?

- Setup
  Train the EOSC model from scratch and extract the BN after X rounds

- Result
  Dependencies **gradually** revealed:

- Implication
  AIF can be used to identify **causal relations** according to current and upcoming observations. Results are intuitively comprehensible.

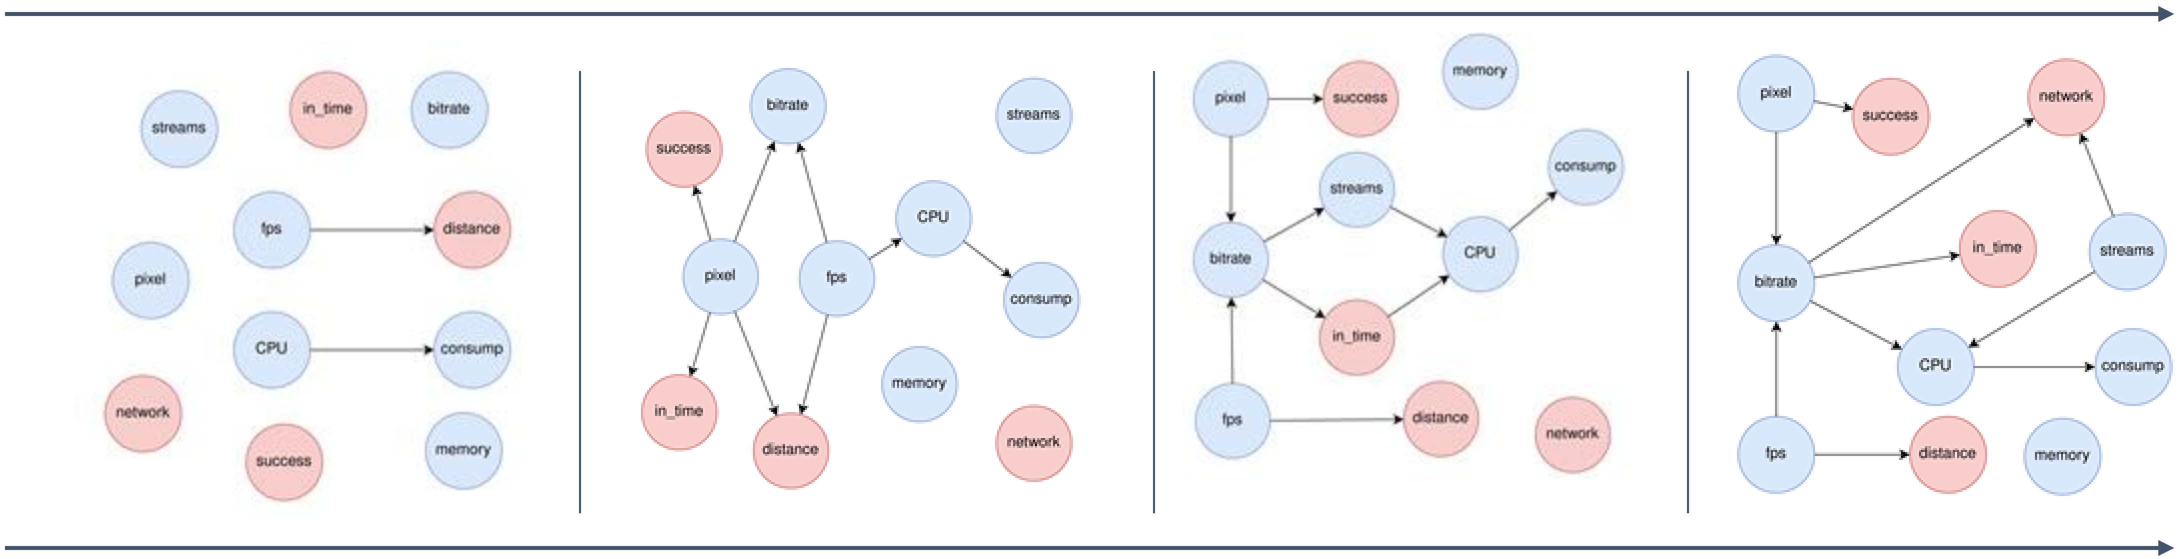

1 AIF round　3 AIF rounds　5 AIF rounds　10 AIF rounds

A-4-2: Is the behavior of AIF agents explainable?

- **Setup**
  Train the EOSC model from scratch and extract the agent's behavioral factors after X rounds

- **Result**
  Develops clear preferences

- **Implication**
  Allows to **empirically debug** the behavior and **fine-tune** agent by adjusting hyperparameters

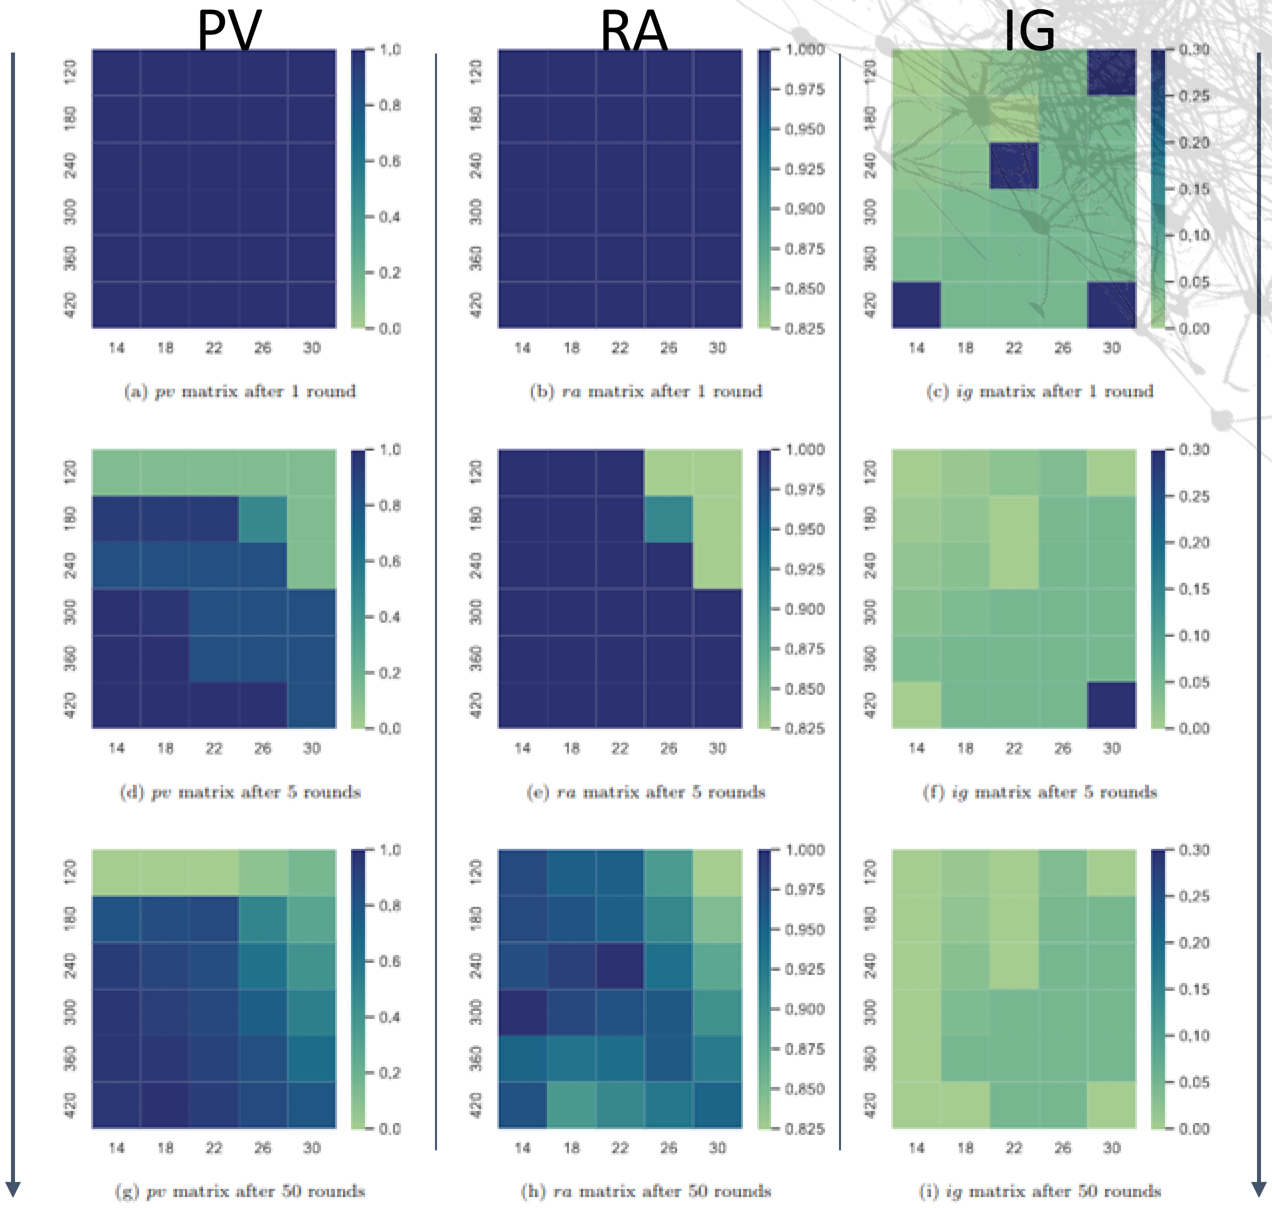

K-3: Do tailored models have lover surprise compared to existing models?

- Setup
  Federate EOSC models within the cluster, select and ***combine*** models for joining edge device; track retraining.

- Result
  Tailor-made model reported the lowest **surprise**, although remaining models improved through **retraining.**

- Implication
  Surprise can be decreased by choosing a (best-)fitting device model .

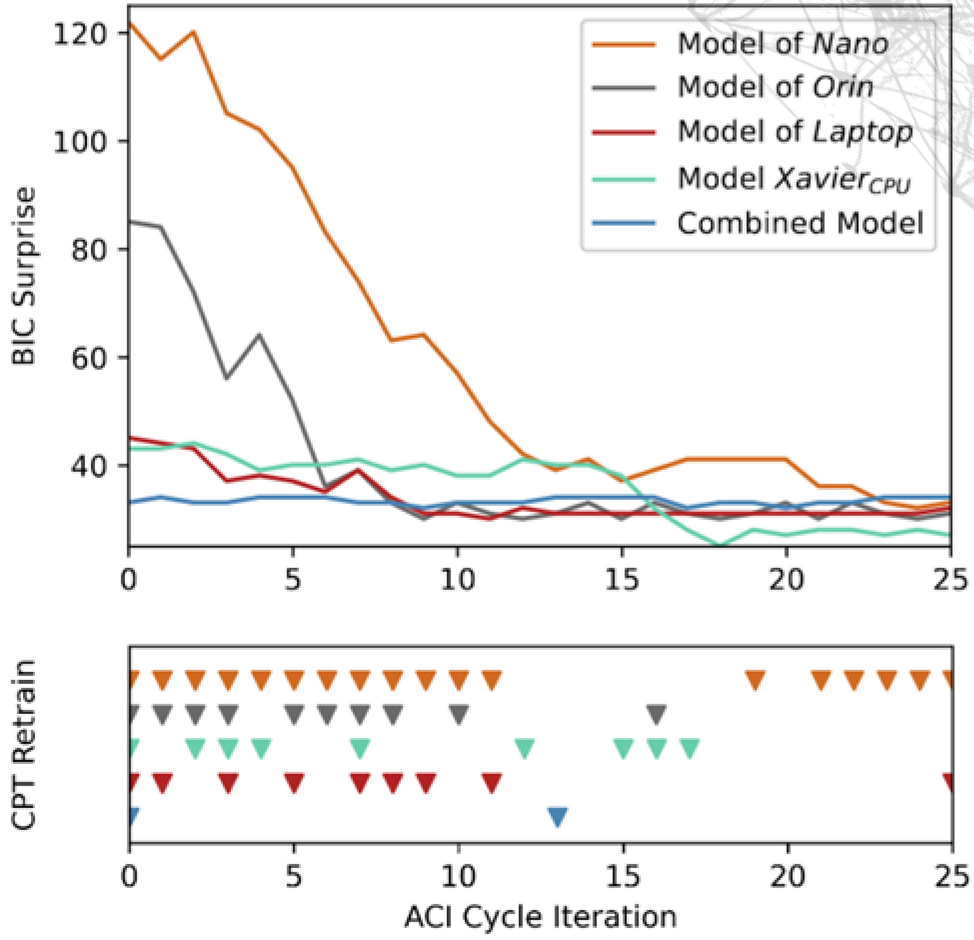

S-2: Can intelligent CC structures optimize local SLO fulfillment?

- Setup
  Clients distributed equally between **comparable** devices, introducing network *congestion* for one of them; rebalance load.

- Result
  Cluster-wide SLO fulfillment (Σ) improved from 1.03 to 1.53.

- Implication
  Was able to **raise the scope** of elasticity strategies, but requires sufficient data to model the relation of *congestion* → *slo_rate*.

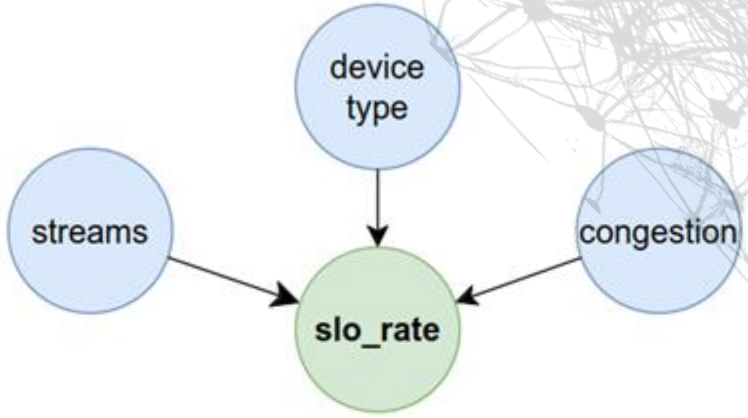

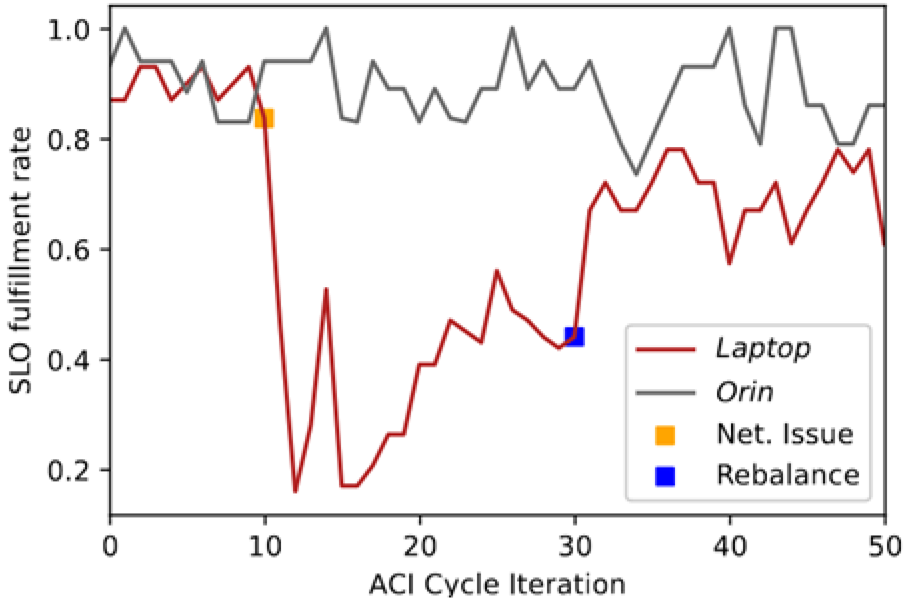

Summary

- Human Eco-systems' intelligent bring Distributed Continuum systems more intelligent and adaptable.
  - Specially Self-healing
  - Higher QoS and QoE
  - Self-adaption

- Active Inference as key method for **self-adaptation**
  - **Autonomous** EOSC model training and updating
  - Fulfill SLOs through **continuous** reconfiguration

Future plans

- Pending comparison with other (ML) approaches
  - Evaluation of more complex use cases
- Composition of MBs for larger structures (**DeepSLOs**)
  - Constrain one MB depending on another's SLOs

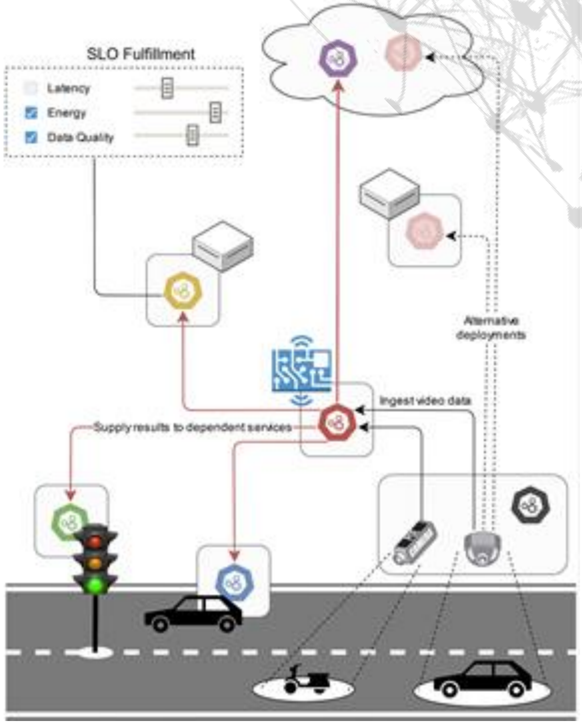

Thank you

Dr. Praveen Kumar Donta
Distributed Systems Group
TU Wien, Vienna, Austria
(pdonta@dsg.tuwien.ac.at)
https://dsg.tuwien.ac.at/team/pdonta/

References


Designing Reconfigurable Intelligent Systems with Markov Blankets


Active Inference on the Edge: A Design Study*


On Causality in Distributed Continuum Systems

Victor Casamayor Pujol, Boris Sedlak, Praveen Kumar Donta, and Schahram Dustdar, Vienna University of Technology, Vienna, 1040, Austria


On Distributed Computing Continuum Systems

Schahram Dustdar, Fellow, IEEE, Victor Casamayor Pujol, and Praveen Kumar Donta, Member, IEEE

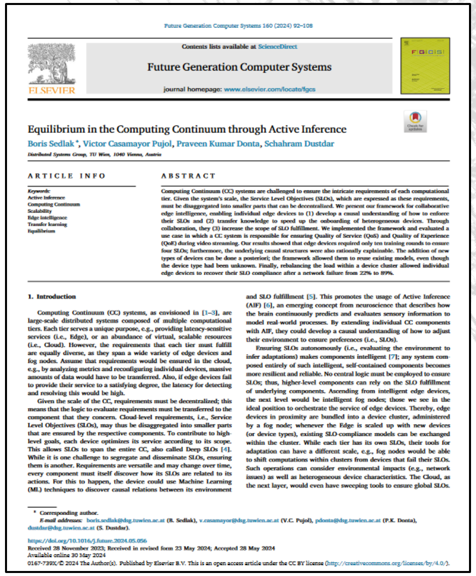
Future Generation Computer Systems

Equilibrium in the Computing Continuum through Active Inference

Boris Sedlak, Victor Casamayor Pujol, Praveen Kumar Donta, Schahram Dustdar

- Designing Reconfigurable Intelligent Systems with Markov Blankets, ICSOC 2023, https://doi.org/10.1007/978-3-031-48421-6_4
- Active Inference on the Edge: A Design Study, IEEE PerconAI 2024, https://doi.org/10.1109/PerComWorkshops59983.2024.10502828
- On Causality in Distributed Continuum Systems, IEEE Internet Computing 2024. https://doi.org/10.1109/MIC.2023.3344248
- On Distributed Computing Continuum Systems, IEEE Transactions on Knowledge and Data Engineering, 2023. https://doi.org/10.1109/TKDE.2022.3142856
- Equilibrium in the Computing Continuum through Active Inference, Future Generation Computer Systems, 2024. https://doi.org/10.1016/j.future.2024.05.056

Additional readings:

- Deep SLOs in the Computing Continuum, ACM AppLIED 2024. https://doi.org/10.1145/3663338.3663681
- Edge intelligence—research opportunities for distributed computing continuum systems, IEEE Internet Computing 2023. https://doi.org/10.1109/MIC.2023.3284693
- Governance and sustainability of distributed continuum systems: a big data approach, Journal of Big Data, 2023. https://doi.org/10.1186/s40537-023-00737-0
- Fundamental Research Challenges for Distributed Computing Continuum Systems , Information, 2023. https://doi.org/10.3390/info14030198
- Exploring the potential of distributed computing continuum systems. *Computers, 2023.* https://doi.org/10.3390/computers12100198