

Urgent Computing: Requirements from a case study

Daniel Balouek

STACK team, Inria, IMT Atlantique, Nantes, France



Inria



A Motivating Use case: Wildfire in California...

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park

By Jason Hanna, [Rebekah Riess](#), [Sara Smart](#) and Andy Rose, CNN

🕒 4 minute read · Updated 2:45 AM EDT, Mon July 25, 2022



A Motivating Use case: Wildfire in California leads to Air pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park

By Jason Hanna, Rebekah Riess, Sara Smart and Andy Rose, CNN
Updated 0645 GMT (1445 HKT) July 25, 2022



Detection → Emission → Transport

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2}, A R Ravishankara^{1,2}, Steven J Brey², Emily V Fischer², John Volckens³ and Sonia Kreidenweis²

Published 22 February 2021 • © 2021 The Author(s). Published by IOP Publishing Ltd

Air Pollution Health Effects

Respiratory

Coughing, wheezing, reduced lung function

Exacerbation of asthma, COPD

Lung cancer

Respiratory mortality

Reproductive

Low birth weight

Preterm births; intrauterine growth retardation

Birth defects

Central Nervous

Stroke (?)

Cognitive effects(?)

Cardiovascular


HRV reduction, dysrhythmias

Systemic inflammation

Atherosclerosis

Myocardial infarctions (Heart Attacks)

CV mortality






A Motivating Use case: Wildfire in California leads to Air Pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park


By Jason Hanna, Rebekah Riess, [Sara Smart](#) and Andy Rose, CNN

🕒 Updated 0645 GMT (1445 HKT) July 25, 2022

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2} , A R Ravishankara^{1,2} , Steven J Brey² , Emily V Fischer² , John Volckens³ 
and Sonia Kreidenweis² 

Published 22 February 2021 · © 2021 The Author(s). Published by IOP Publishing Ltd








A Motivating Use case: Wildfire in California leads to Air Pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park

By Jason Hanna, Rebekah Riess, [Sara Smart](#) and Andy Rose, CNN

Updated 0645 GMT (1445 HKT) July 25, 2022

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2} , A R Ravishankara^{1,2} , Steven J Brey² , Emily V Fischer² , John Volckens³ 
and Sonia Kreidenweis² 

Published 22 February 2021 · © 2021 The Author(s). Published by IOP Publishing Ltd



2013 Rim Fire Smoke (California) to Reno (Nevada) - 190 km (118 miles)






A Motivating Use case: Wildfire in California leads to Air Pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park

By Jason Hanna, Rebekah Riess, [Sara Smart](#) and Andy Rose, CNN

⌚ Updated 0645 GMT (1445 HKT) July 25, 2022

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2} , A R Ravishankara^{1,2} , Steven J Brey² , Emily V Fischer² , John Volckens³ 
and Sonia Kreidenweis² 

Published 22 February 2021 · © 2021 The Author(s). Published by IOP Publishing Ltd



We seem to have everything to deal with such situation, but...

- ✓ Sensors
- ✓ Advanced application models
- ✓ Large, powerful compute resources





A Motivating Use case: Wildfire in California leads to Air Pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park


By Jason Hanna, Rebekah Riess, [Sara Smart](#) and Andy Rose, CNN

🕒 Updated 0645 GMT (1445 HKT) July 25, 2022

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2} , A R Ravishankara^{1,2} , Steven J Brey² , Emily V Fischer² , John Volckens³ 
and Sonia Kreidenweis² 

Published 22 February 2021 · © 2021 The Author(s). Published by IOP Publishing Ltd



We seem to have everything to deal with such situation, but...

- ✓ Sensors
- ✓ Advanced application models
- ✓ Large, powerful compute resources
- ✗ No agreed-upon infrastructure and application models
- ✗ End-to-end performance and guarantees

CONTEXT: **Harnessing The Computing Continuum for Urgent Science**

CASE STUDY: **Earthquake Early Warning**

ON-GOING WORK: **Building data-driven pipelines at middleware-level**

CONCLUSION: **The Open-Road Ahead**

Urgent Computing: Managing uncertainty under constraints

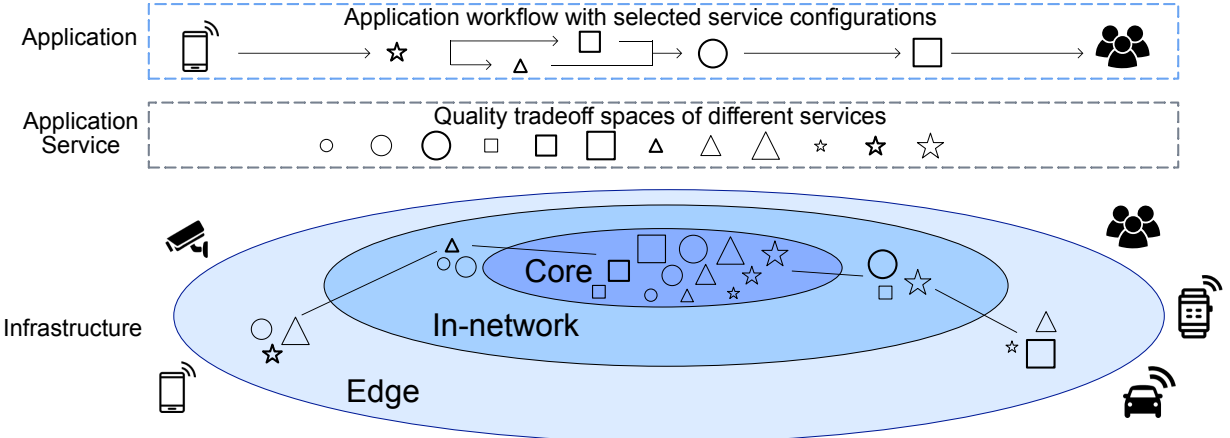
DEFINITION: Computing under **strict time** and **quality constraints** to support decision making with the desired confidence within a defined time interval



Urgent applications need to react to unforeseen events and to manage complex cost/benefit tradeoffs to meet constraints

- Steer online/on-demand computations based on data/content
- Balance costs of computations versus efficiency and availability of resources
- Support urgency constraints for timely decision-making
- ...

The Computing Continuum as a system of systems



FOG/EDGE COMPUTING

- Closer to the data
- Lower latencies
- Limited & expensive computation/storage
- Unreliable connectivity

IN-TRANSIT COMPUTING

- Distributed along the data path
- Limited, but can be effective
- Fewer guarantees

CORE DATA CENTERS/CLOUDS


- Far from data
- Expensive data access
- Relatively inexpensive
- Seemingly infinite

Combining the right compute resources at optimal processing points requires novel programming models


Case study: EEW

Earthquake Early Warning

Goal: Increase the confidence and response time of seismic event detection



Approach: Combine multiple data sources to cover the whole spectrum of events; leverage the CI continuum for timely processing



Earthquake Early Warning (EEW)



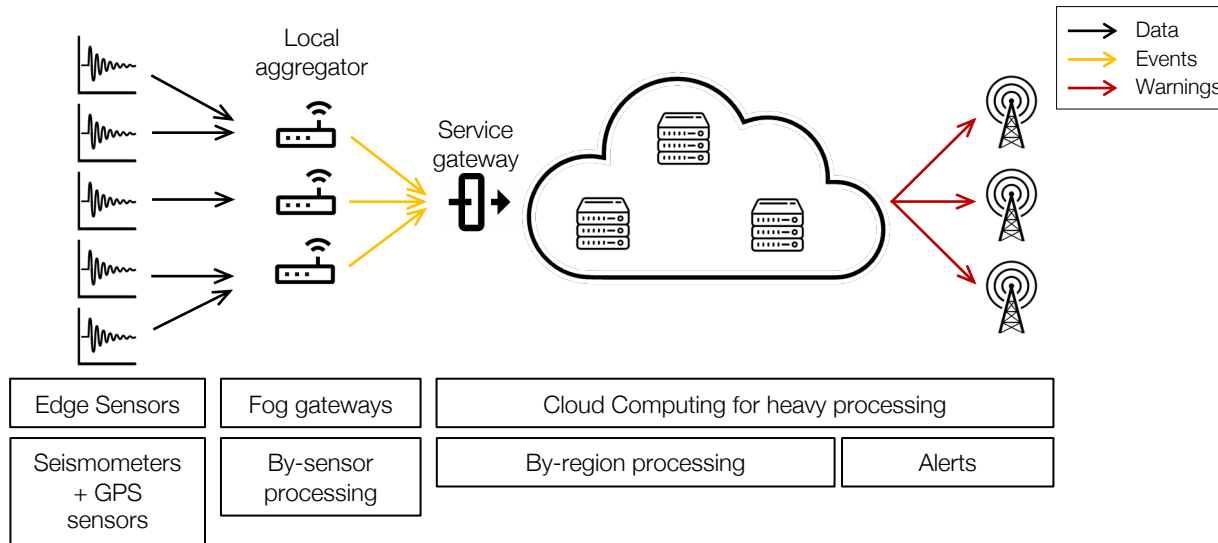
PROBLEM: How to program and distribute a complex ML model while preserving latency and resiliency across the Edge-to-Cloud Continuum?



STRATEGY: Decompose end-to-end services into manageable units
 Integrate the resulting units back into a functioning system at scale
 Adapt the system according to execution context.



SOLUTIONS: **Model for distributed Stream Processing**
Programmable reactive behaviors



Lessons from the deployment




Fig. Latency resulting from processing one seismic event on WE and WE-CO scenarios for different network configurations.

- 1,000 sensors emulated by virtual machines deployed on 25 host machines.
- Each set of 20 sensors composes an Edge layer and it is connected to one of the 50 Fog gateways.
- Each sensor produces 100 messages per second for 2 minute (Total of 12M messages)

Observations from this evaluation

The approach allowed us to model and implement the Earthquake Early Warning application

- Edge can be an interesting option for network constrained scenarios, but centralizing the processing is a better alternative in terms of latency
- Impact of the Libraries and Application Logic

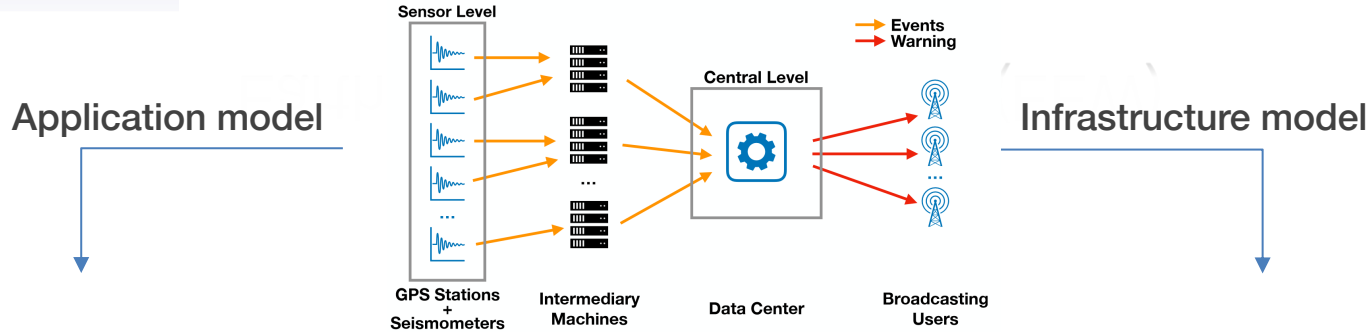
Still...

How to detect changes, and move from one configuration to another?

How to program data reactions to unforeseen events?

Contribution: Application Level

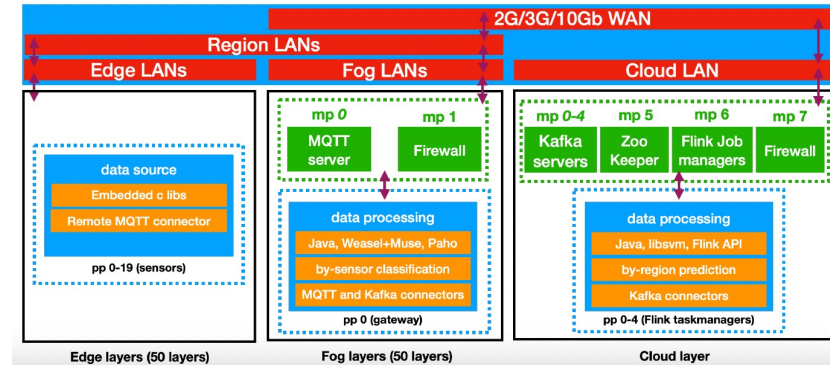
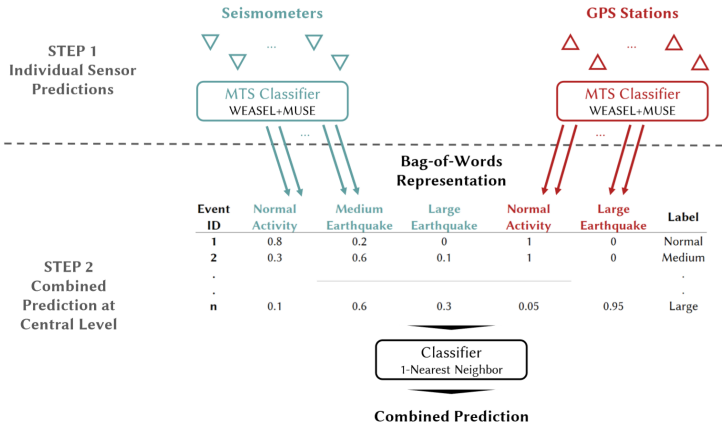
Earthquake Early Warning (EEW)



Detection of **all large earthquakes** without false alerts [AAA120]

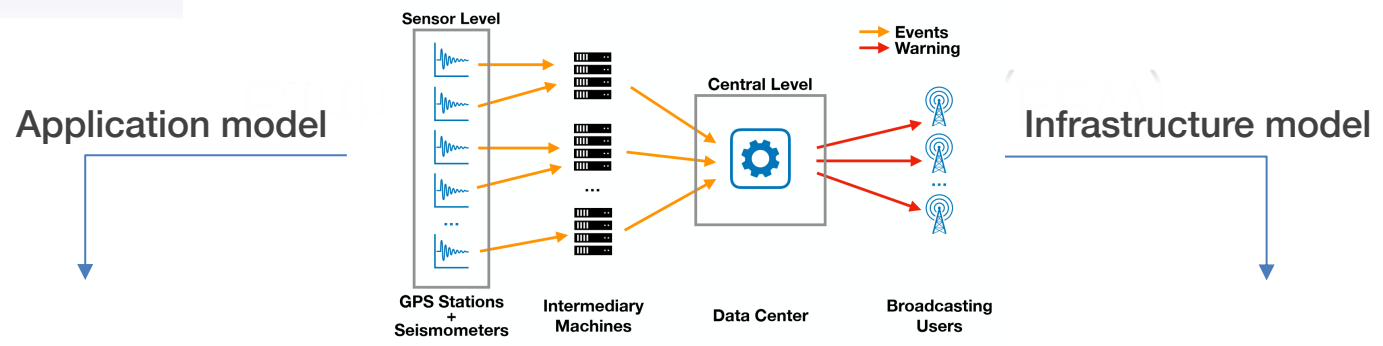
Geographically distributed infrastructure [UrgentHPC21, DML-ICC21]

- Reduce the volume of data transmitted in the network
- Increase the robustness of the EEW system


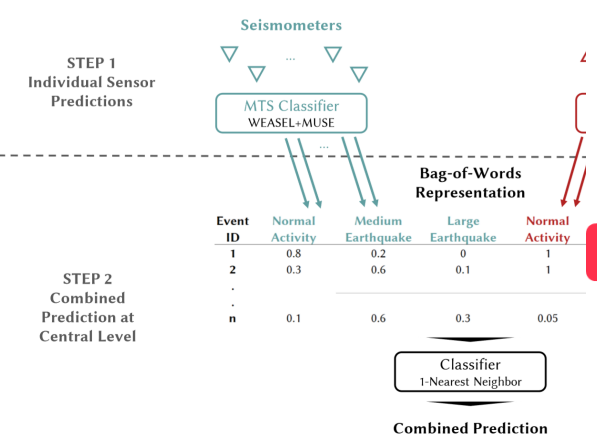


Contribution: Application Level

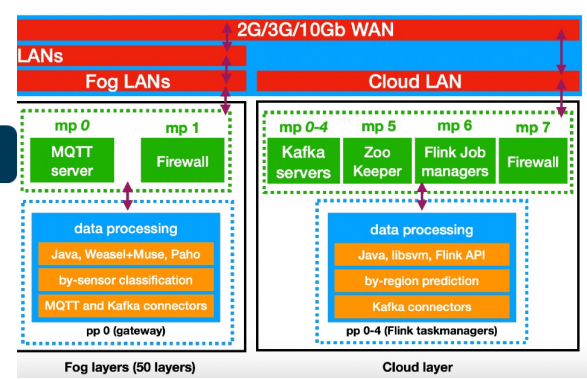
Earthquake Early Warning (EEW)



Detection of **all large earthquakes** without



Infrastructure [UrgentHPC21, DML-ICC21]
Volume of data transmitted in the network
Performance of the EEW system



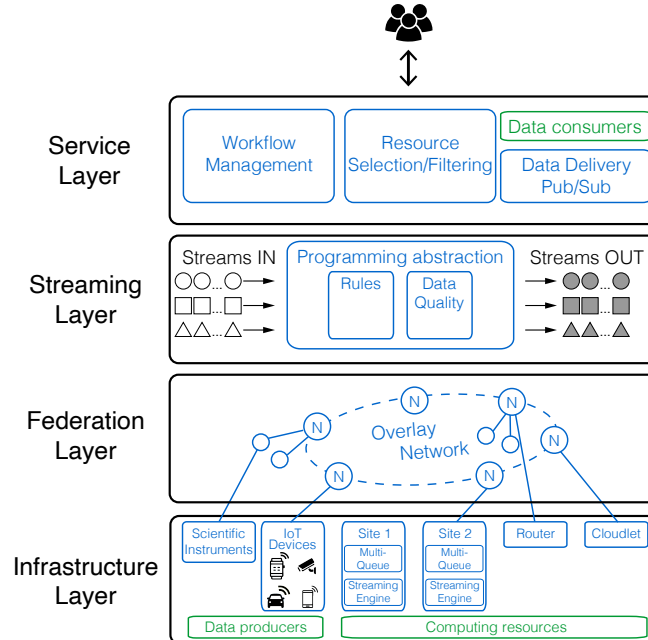
R-Pulsar: Enabling Data-driven Workflows across the Continuum (rpulsar.org)



PROBLEM: What, where and when to compute unbound data streams generated by multiple, distributed sources?



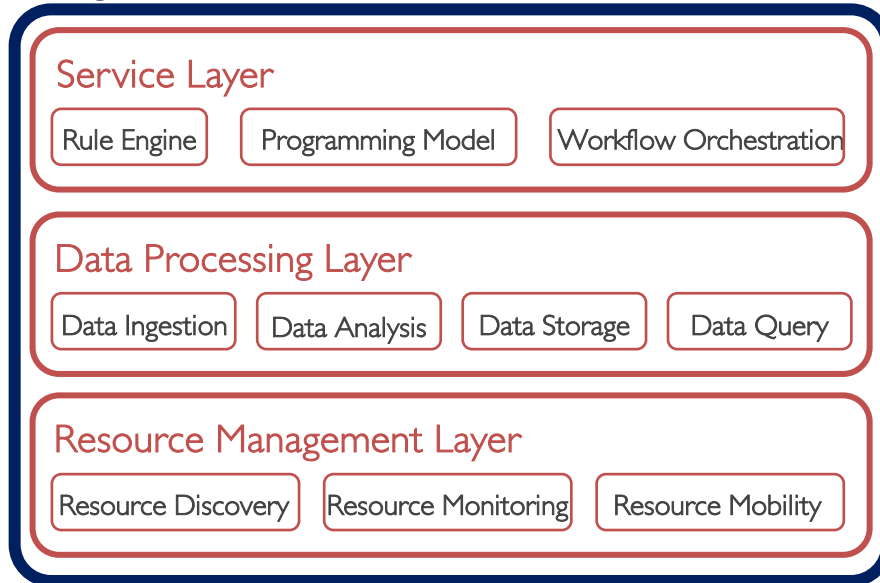
STRATEGY: Flexible programming model that enables data driven, location aware and resource aware mechanisms between Edge and Cloud.



Survey of Edge and IoT Platforms

- 100+ edge-based middleware platforms in the market
 Study of contributions and systems through a simplified architecture
- Components and design goals
 - Academic and commercial systems

Edge Platform Architecture



System	Data Processing Layer				Resource Management Layer			Service Layer		
	Data Ingestion	Data Analysis	Storage	Data Query	Resource Discovery	Resource Monitoring	Resource Mobility	Workflow Orchestration	Rule Engine	Programming Model
AWS Greengrass [4]	✓	✓				✓			✓	✓
Azure IoT Edge [58]	✓	✓	✓	✓		✓			✓	✓
Eaas [94]	✓	✓							✓	✓
Google Cloud IoT [33]	✓	✓			✓	✓			✓	✓
Everyware IoT [29]	✓	✓			✓	✓			✓	✓
Predix [26]	✓	✓			✓	✓			✓	✓
Bosch IoT [12]	✓	✓	✓		✓	✓			✓	✓
Yanzi et al. [95]	✓	✓			✓	✓			✓	✓
R-Pulsar [74,77]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
FogHorn [31]	✓	✓			✓	✓			✓	✓
Geelytics [16]	✓	✓							✓	✓
Fogflow [17]	✓	✓	✓	✓	✓	✓		✓	✓	✓
OpenMTC [64]	✓	✓			✓	✓			✓	✓
SiteWhere [85]	✓	✓	✓	✓	✓	✓			✓	✓
SmartThings [86]	✓	✓							✓	✓
Kaa [42]	✓	✓	✓	✓	✓	✓	✓		✓	✓
Samsung Artik [80]	✓	✓							✓	✓
Ayla Network [62]	✓	✓							✓	✓
Altair SmartWorks [87]	✓	✓							✓	✓
EdgeX [24]	✓	✓			✓	✓			✓	✓
PlCasso [47]	✓	✓			✓	✓		✓	✓	✓
Hua-Jung Hog et al. [36]	✓	✓	✓	✓	✓	✓			✓	✓
Cloud4IoT [69]	✓	✓			✓			✓	✓	✓
Nebulae [61]	✓	✓							✓	✓
FogGIS [8]	✓	✓	✓	✓					✓	✓
FOG-engine [57]	✓	✓	✓	✓				✓	✓	✓
CEFIoT [40]	✓	✓	✓	✓				✓	✓	✓
SAVI-Iot [43]	✓	✓	✓	✓				✓	✓	✓
Foggy [97]	✓	✓						✓	✓	✓
ISYMPHONY [96]	✓	✓						✓	✓	✓
Macchina.io [51]	✓	✓	✓	✓					✓	✓
Clearblade [21]	✓	✓	✓	✓					✓	✓
IBM Watson IoT [38]	✓	✓			✓	✓			✓	✓

Survey of Edge and IoT Platforms: [IC321] Balouek-Thomert, D., & Parashar, M. (2021, August). Platforms for Edge Computing and Internet of Things applications: A survey. In 2021 Thirteenth International Conference on Contemporary Computing (IC3-2021) (pp. 140-149).

R-Pulsar: Enabling Data-driven Workflows across the Continuum (rpulsar.org)



PROBLEM: What, where and when to compute unbound data streams generated by multiple, distributed sources?



STRATEGY: Flexible programming model that enables data driven, location aware and resource aware mechanisms between Edge and Cloud.

Producers

Traffic



Smoke



Any physical/virtual sensors


DATA COLLECTION

Consumers / Filters

Individual values



Windows



FILTERING / RULE MATCHING

Actions

Send a message

Approximate computations

Change system parameters
or execution mode

Launching computation
(local, cloud, hybrid, etc.)

ACTION TRIGGERS

Policy-driven adaptation using R-Pulsar

Event-driven programming model based on control-flow constructs

- **Resources and data profiles** based on *attributes and contents/description*
 - Supports content-, location-aware, resource-aware operators
- **Workflow execution** driven by *availability/content of data and availability/state of resources*
- **Autonomic runtime** allows management of *tradeoffs* and/or *service configurations*
 - Cost/Quality/Deadline-oriented runtime policies

When Rules (When to move?): **Where Rules (Where to move it?):**

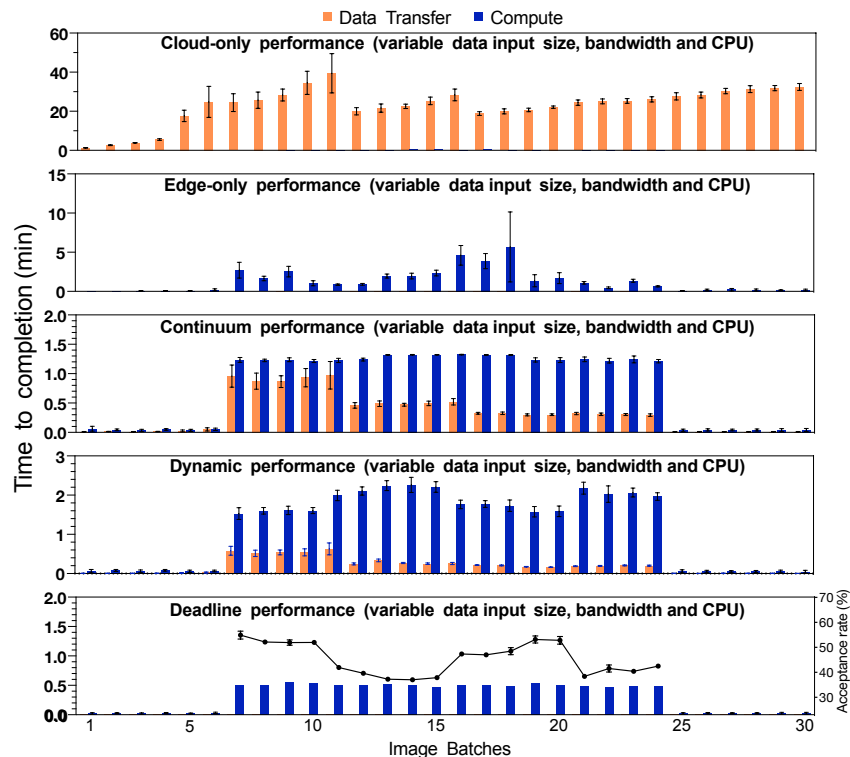




Fig. Time to completion resulting from processing 30 batches of images on each policy with a variable input data size, bandwidth capacity and CPU.







A Motivating Use case: Wildfire in California leads to Air Pollution in Utah

California's fast-moving Oak Fire burns 14,000 acres and forces thousands to evacuate outside Yosemite National Park

By Jason Hanna, Rebekah Riess, Sara Smart and Andy Rose, CNN

Updated 0645 GMT (1445 HKT) July 25, 2022

Could the exception become the rule? 'Uncontrollable' air pollution events in the US due to wildland fires

Liji M David^{1,2} , A R Ravishankara^{1,2} , Steven J Brey² , Emily V Fischer² , John Volckens³  and Sonia Kreidenweis² 

Published 22 February 2021 · © 2021 The Author(s). Published by IOP Publishing Ltd



Edge (Field cameras)

1. Smoke detection at the edge of the network using Camera Imagery




Cloud (SDSC)

2. Wildfire simulations to determine the severity and direction of fires



HPC (Utah)

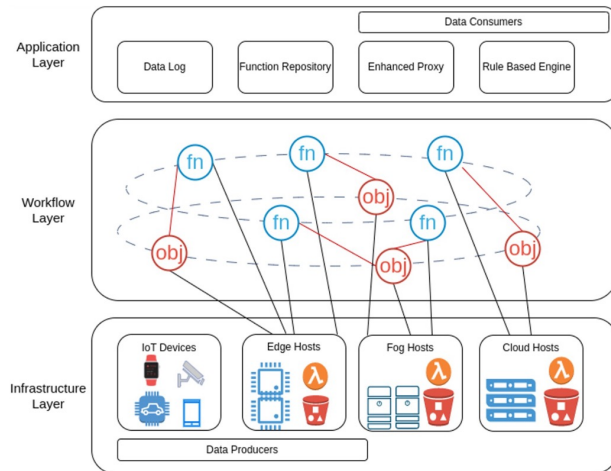
3. Pollution Concentration maps to support decision-making



GOAL

Realizing a fluid ecosystem where distributed computing resources and services ([Computing Continuum](#)) are composed on-demand to support delay-sensitive and data-driven workflows ([Urgent Applications](#))

RESEARCH DIRECTIONS: Taking processing to the data, not the other way around



How to orchestrate and where to run?

Requirements

Data collection and preparation

- Reflect **instant** changes in **real-time** systems
- Careful consideration is needed for **data quality** control

Computational resources

- **Requires** significant computational resources (include **HPC** resources, **GPUs**)
- **Sufficient** resources are needed to train and fine-tuning model

Faster development and deployment

- **Reduce** time and cost to validate new CI system
- Transfer knowledge from **existing** data to **new** applications

Integration and Interoperability

- **Integrate** with other existing system platform
- Enable broader system-level **optimization**

Thank you!

Acknowledgments



Daniel Balouek



Email: daniel.balouek@inria.fr

WWW: <http://daniel-balouek.com>

Project resources

R-Pulsar: <http://rpulsar.sci.utah.edu>

Collaborators



We are open to collaborations and internships!

