

# FireDeX: a Prioritized IoT Data Exchange Middleware for Emergency Response

**Georgios Bouloukakis<sup>1,2</sup>**

Rennes, France, December 2018

2018 ACM/IFIP International Middleware Conference

Joint work with Kyle Benson<sup>1</sup>, Casey Grant<sup>3</sup>, Valérie Issarny<sup>2</sup>, Sharad Mehrotra<sup>1</sup>, Ioannis Moscholios<sup>4</sup>, Nalini Venkatasubramanian<sup>1</sup>

<sup>1</sup>Donald Bren School of Information and Computer Sciences, UC Irvine, USA

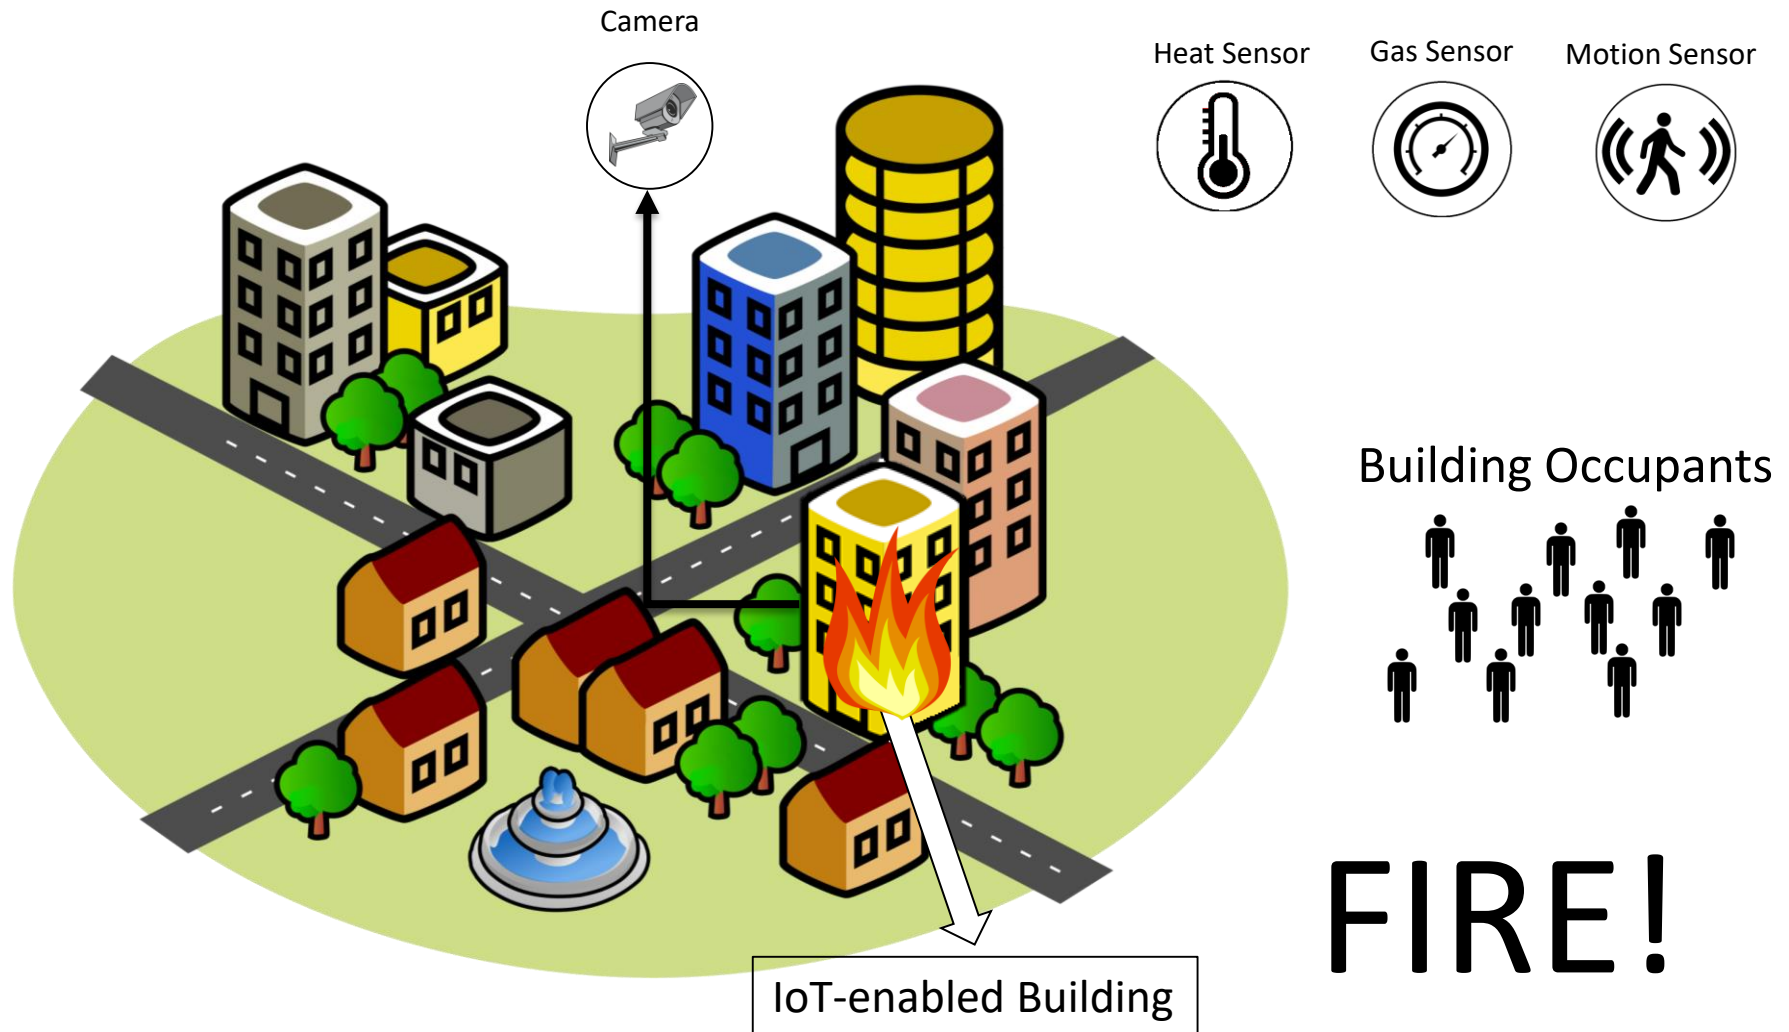
<sup>2</sup>MiMove team, Inria Paris, France

<sup>3</sup>National Fire Protection Association, USA

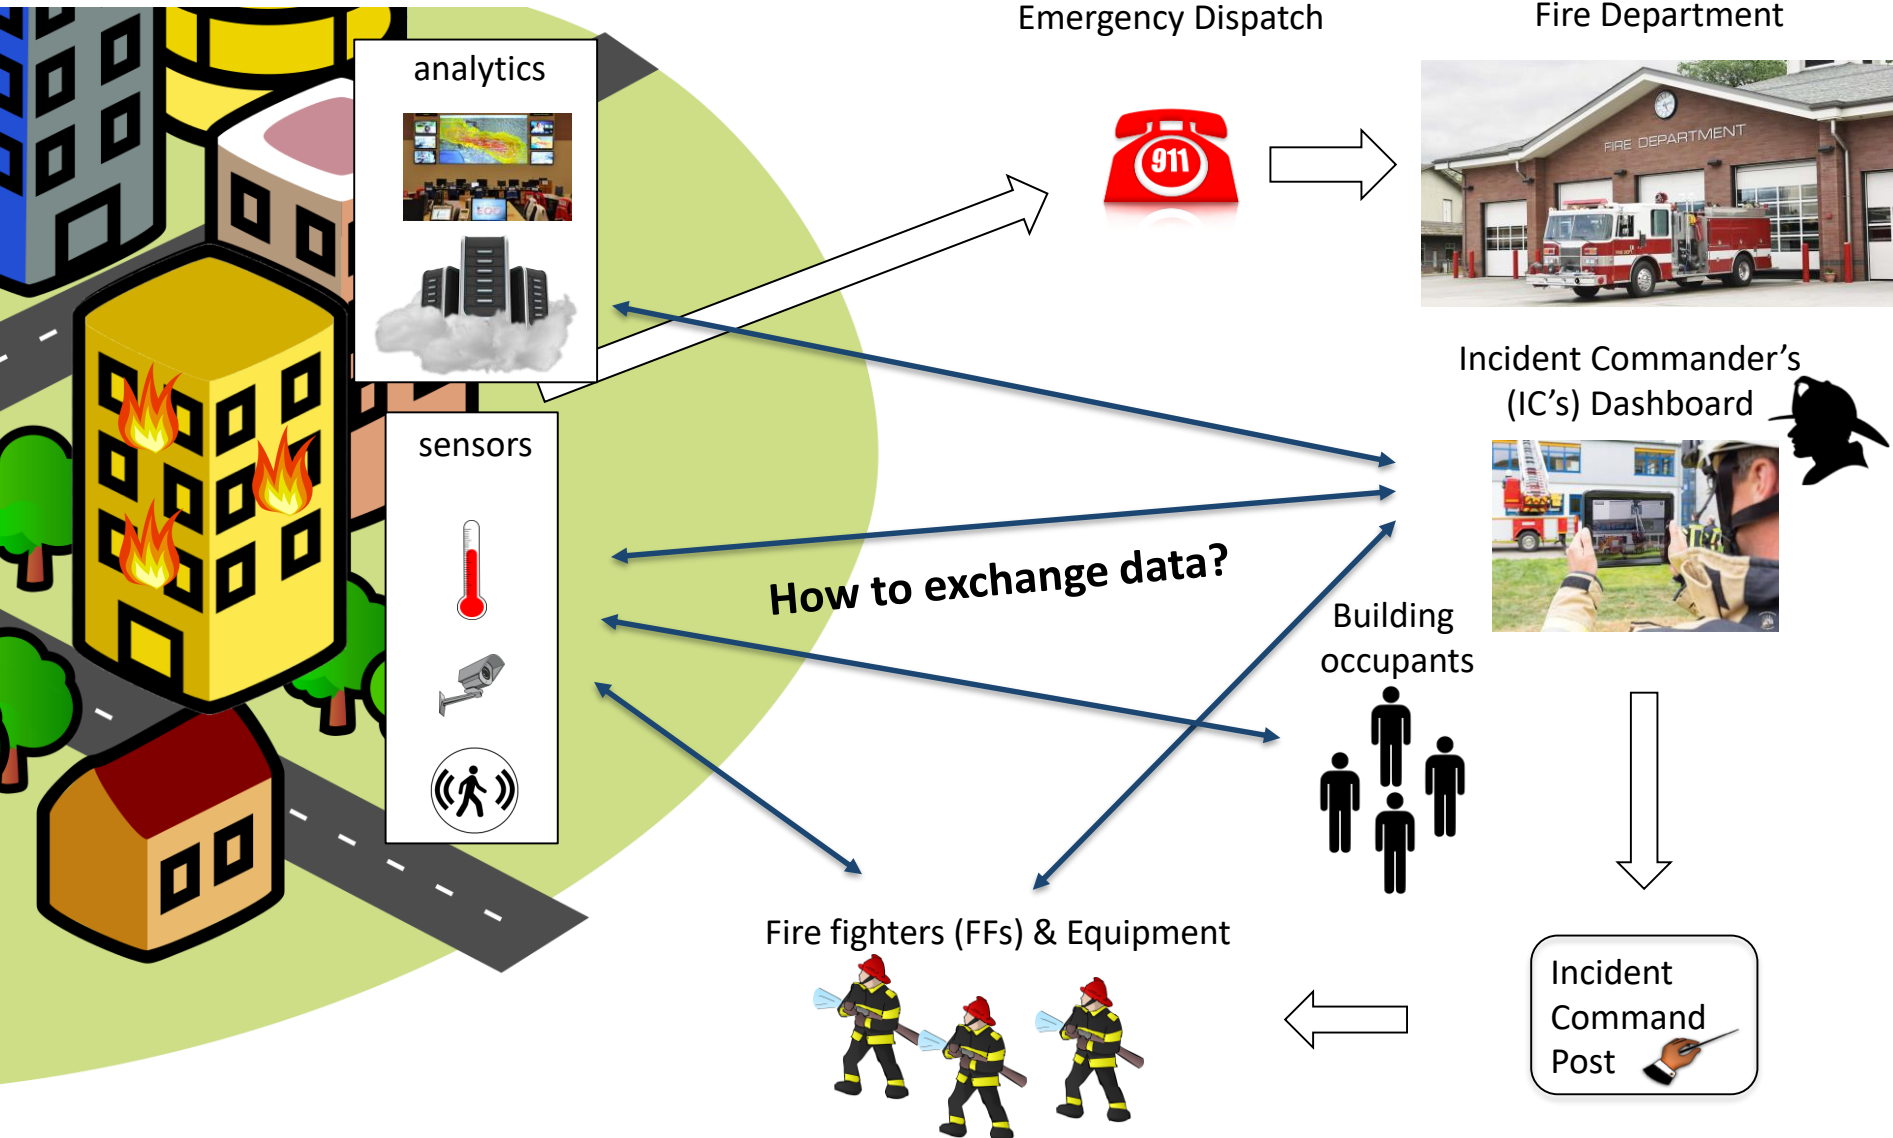
<sup>4</sup>Dept. of Informatics & Telecommunications, Univ. of Peloponnese, Greece



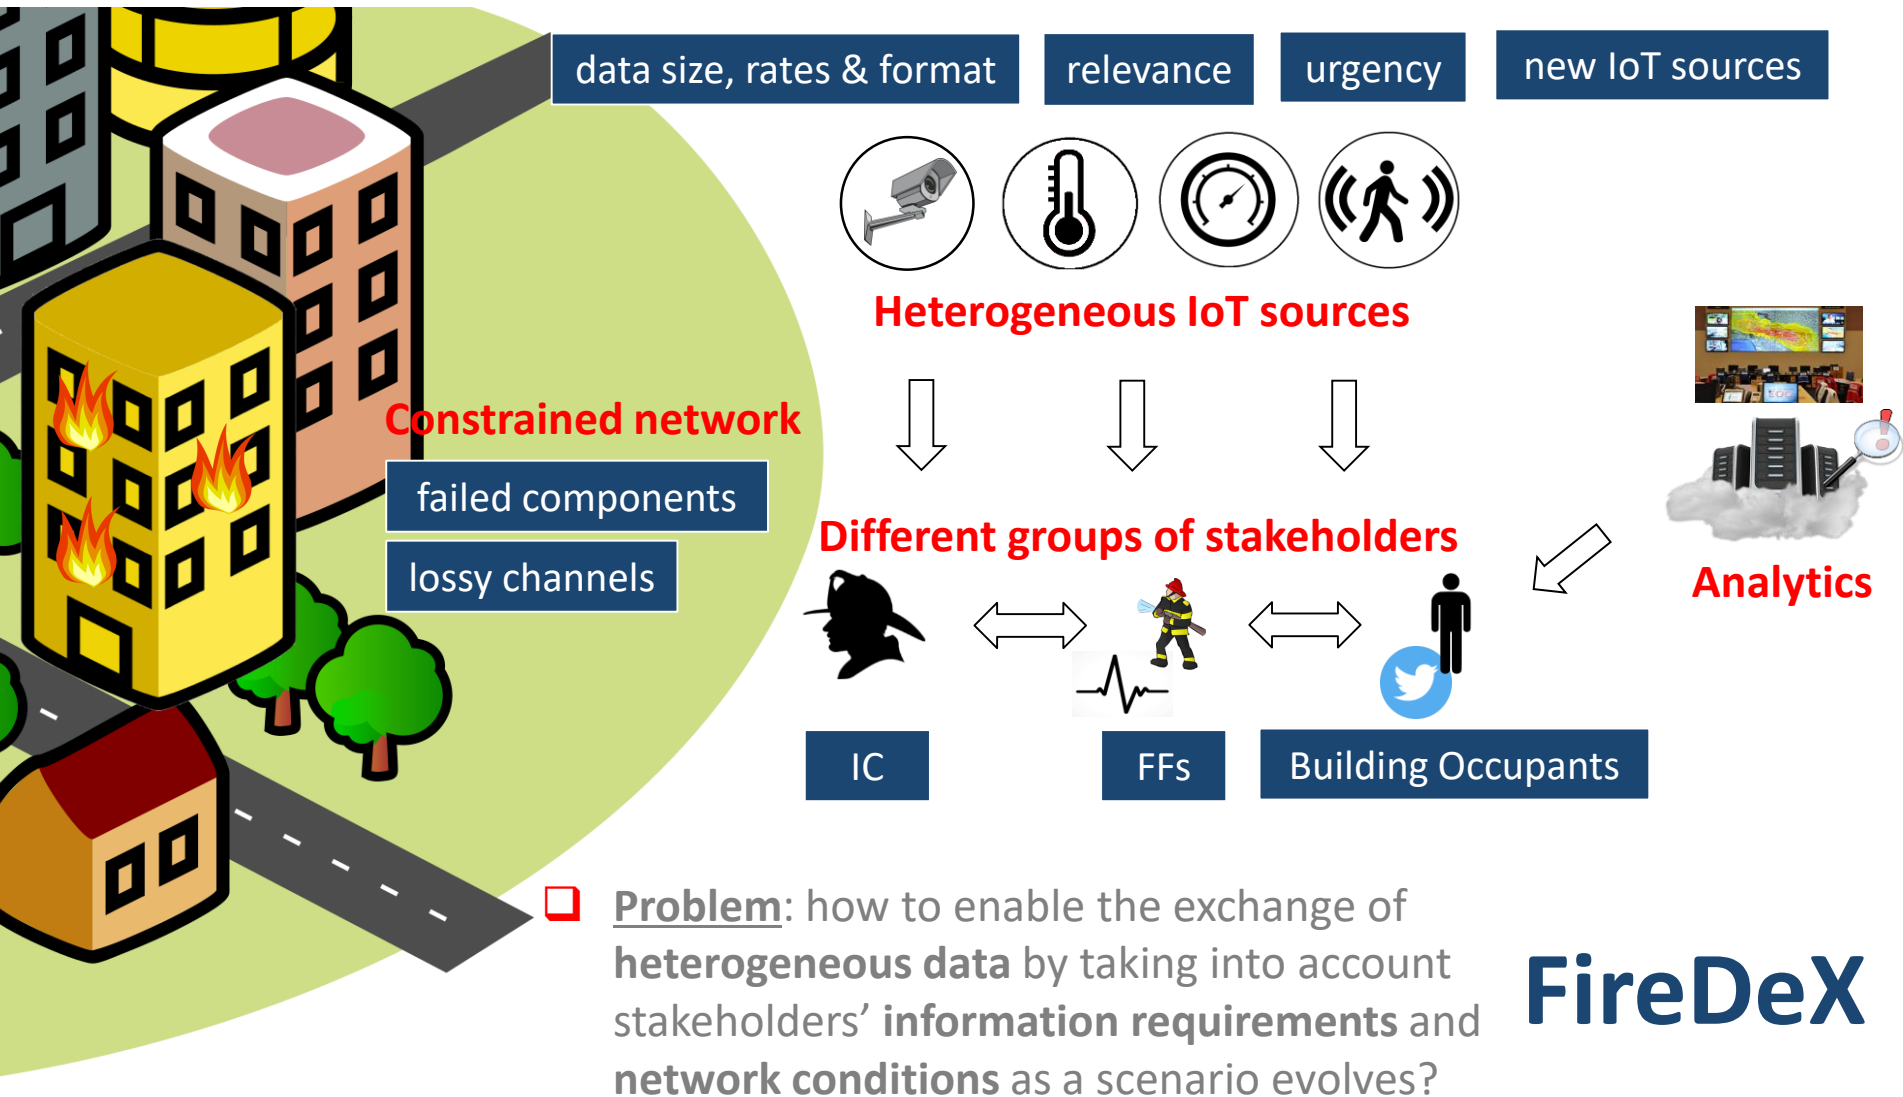
# Motivation: IoT-enhanced structural fire response



# Motivation: IoT-enhanced structural fire response

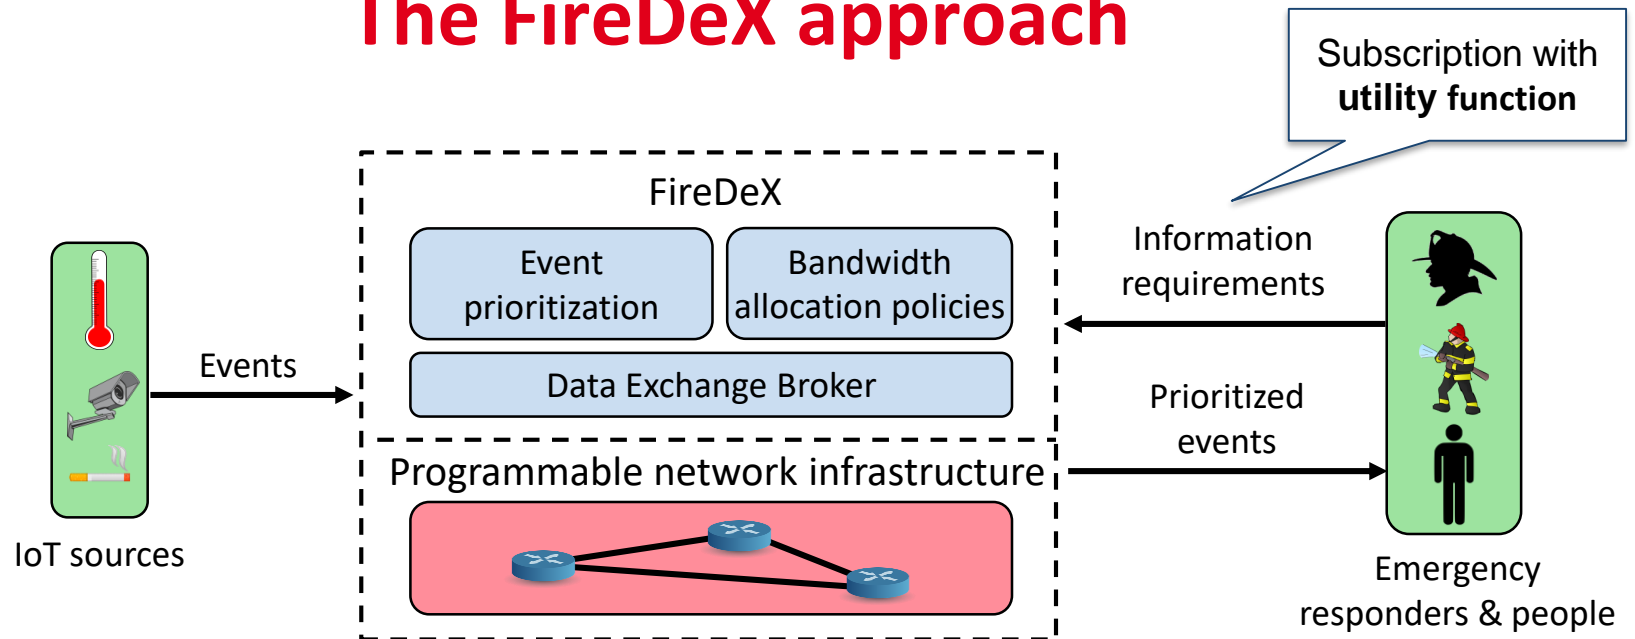


# Motivation: IoT-enhanced structural fire response



FireDeX


# The FireDeX approach



- FireDeX middleware configures the data exchange & network with prioritization and bandwidth allocation policies based on:
  - information requirements
  - network resource constraints
- Rely on SDN to bridge critical information requirements with network flows.
- Model the performance of FireDeX across multiple layers using Queueing Theory.
- Use the underpinning formal model for deriving novel algorithms that prioritize IoT events and tune notification delivery/response times.

Goal: timely and reliable delivery of the most critical data to relevant subscribers despite challenging network conditions.


# SDN-background



# Mapping info. reqs. to network state

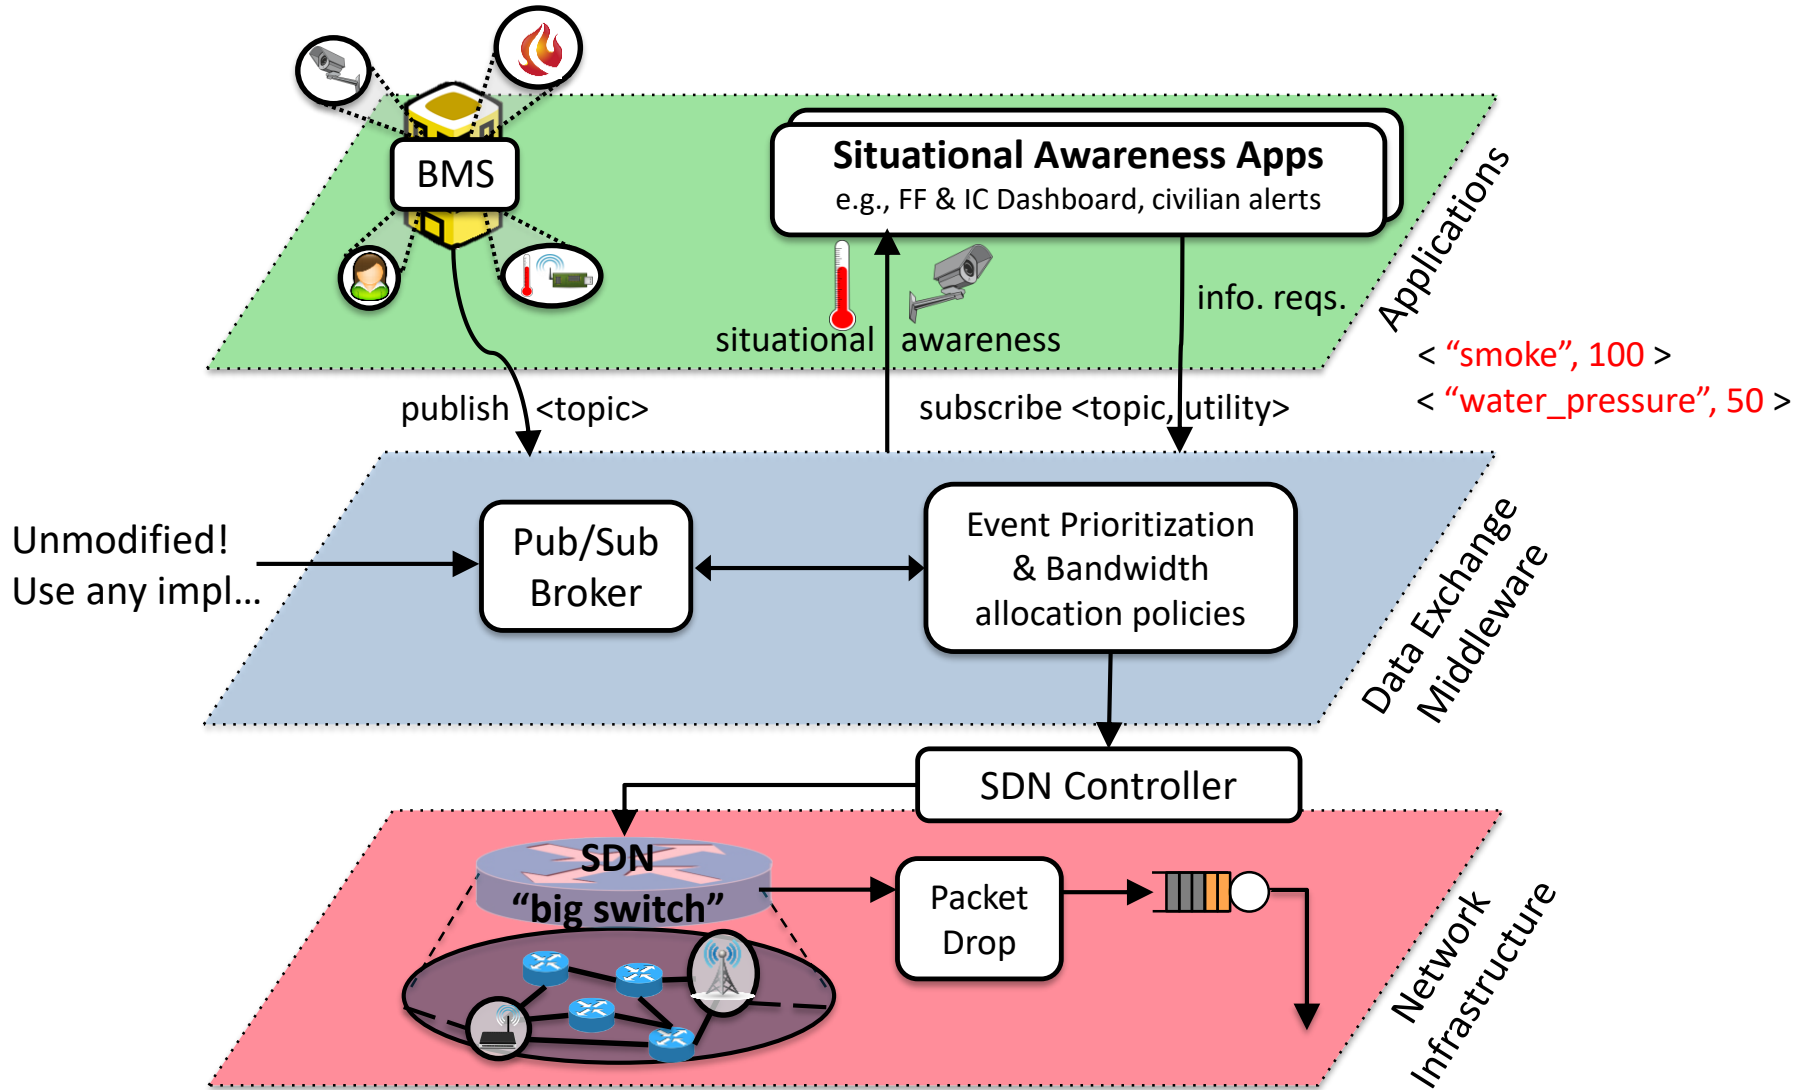
App / DeX concept

FireDeX configurations



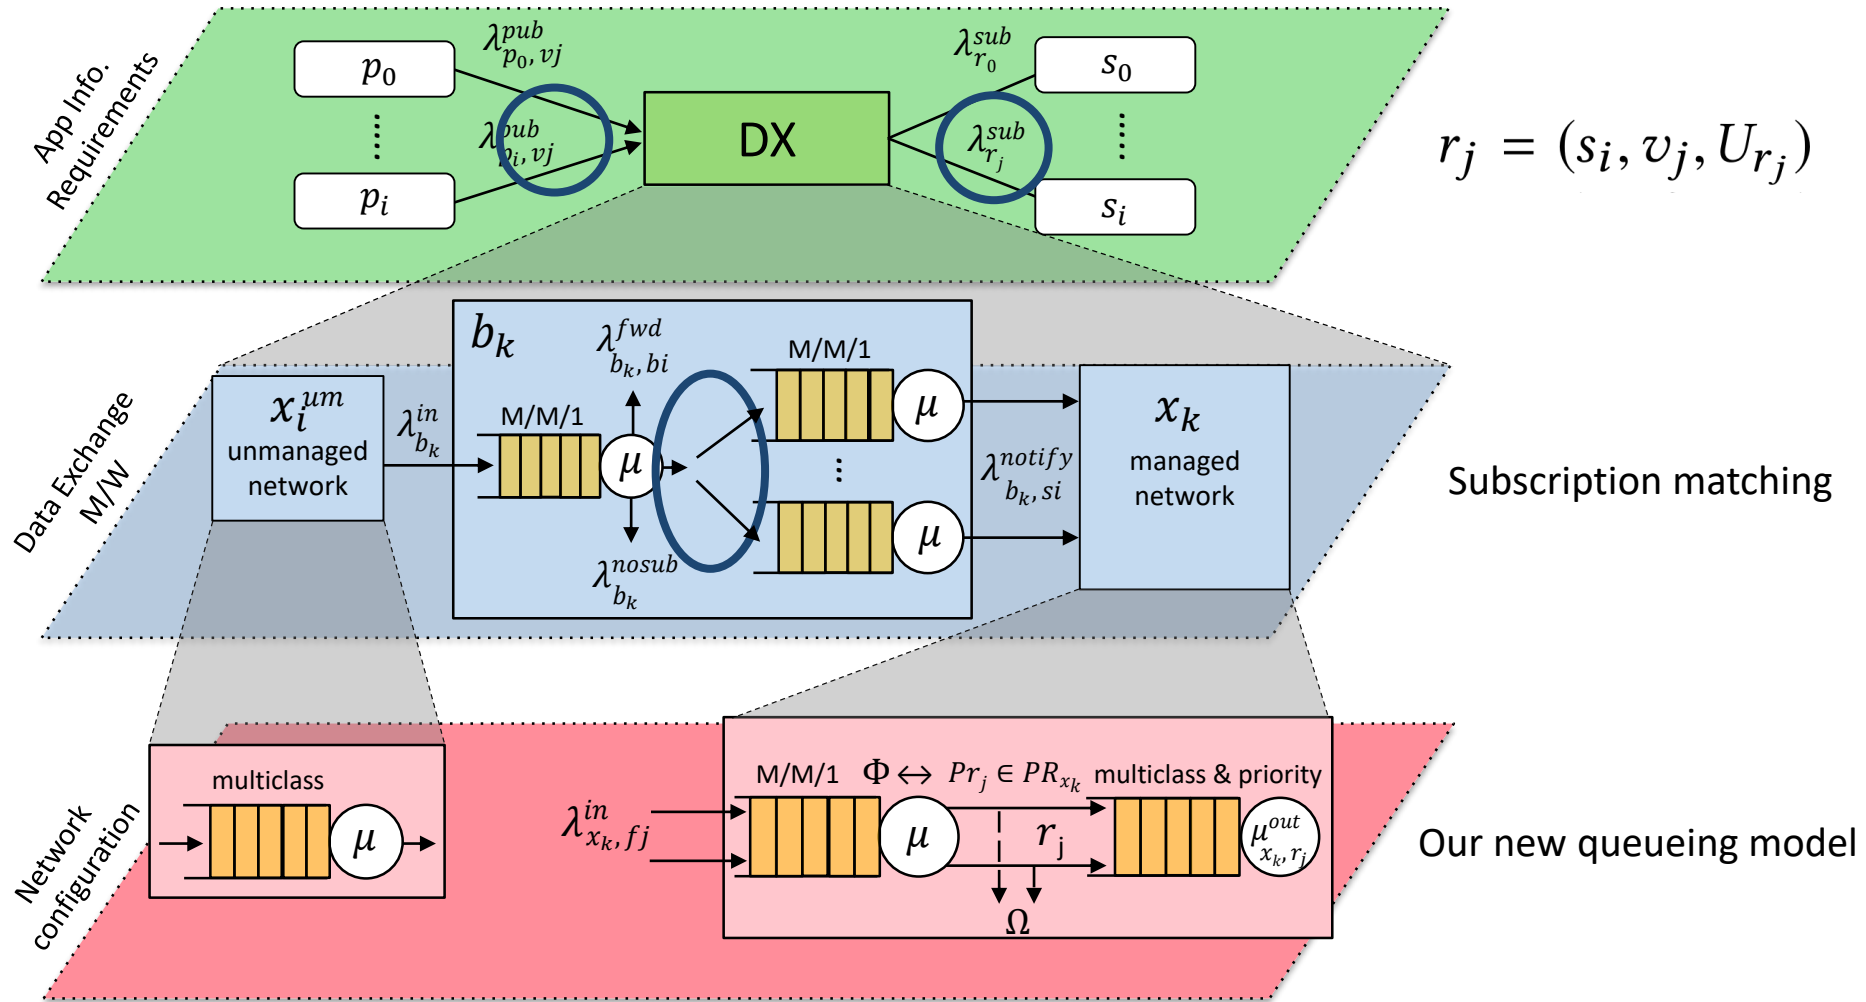
- Network flows enable SDN infrastructure to differentiate subscriptions (e.g. by UDP/TCP port number + IP addr).

# FireDeX across layers





# Modeling FireDeX using queuing theory



# Prioritization algorithm


1. Estimate the adjusted utility function per subscription: information value per unit of *bandwidth*.
2. Sort subscriptions.
3. Group them into approximately equal-sized network flows.
4. Priorities assigned to approximately equal-sized groups of network flows.

$$A_{r_j} = \frac{\hat{U}_{r_j}}{G_{v_j} \lambda_{b_k, r_j}^{notify}}$$

Maximum utility achievable for subscription  $r_j$

Serialized packet size for topic  $v_j$  notifications

Rate of notifications (publications) for subscription  $r_j$



# Drop rate algorithms

Flat

Drop rates for each network flow

$$\Omega(f_j) = \beta$$

Linear

$$\Omega(f_j) = \beta \Phi(f_j)$$

Mapped priority to network flow

Exponential

$$\Omega(f_j) = 1 - \beta^{-\Phi(f_j)}$$

Optimized

1. Formulated as a convex optimization problem.
  - Maximize overall utility as sum of all subscriptions' utilities.
  - Enabled by choice of logarithm for utility function.
2. 2nd constraint: queue stability condition.
  - Ensures allocated bandwidth within that available.

$$\begin{aligned}
 &\text{maximize} && \sum_{r_j \in R} \alpha_{r_j} \log(1 + \Xi_{r_j}) \\
 &\text{subject to} && \Omega(f_j) \in [0, 1], \forall f_j \in F \\
 &&& \sum_{r_j \in R_{x_k}} \frac{\lambda_{x_k, r_j}^{thru}}{\mu_{x_k, r_j}^{out}} \leq 1 + \tilde{\rho}, \forall x_k \in X
 \end{aligned}$$

"Rho tolerance" enables keeping a buffer within the bandwidth (~0.1)

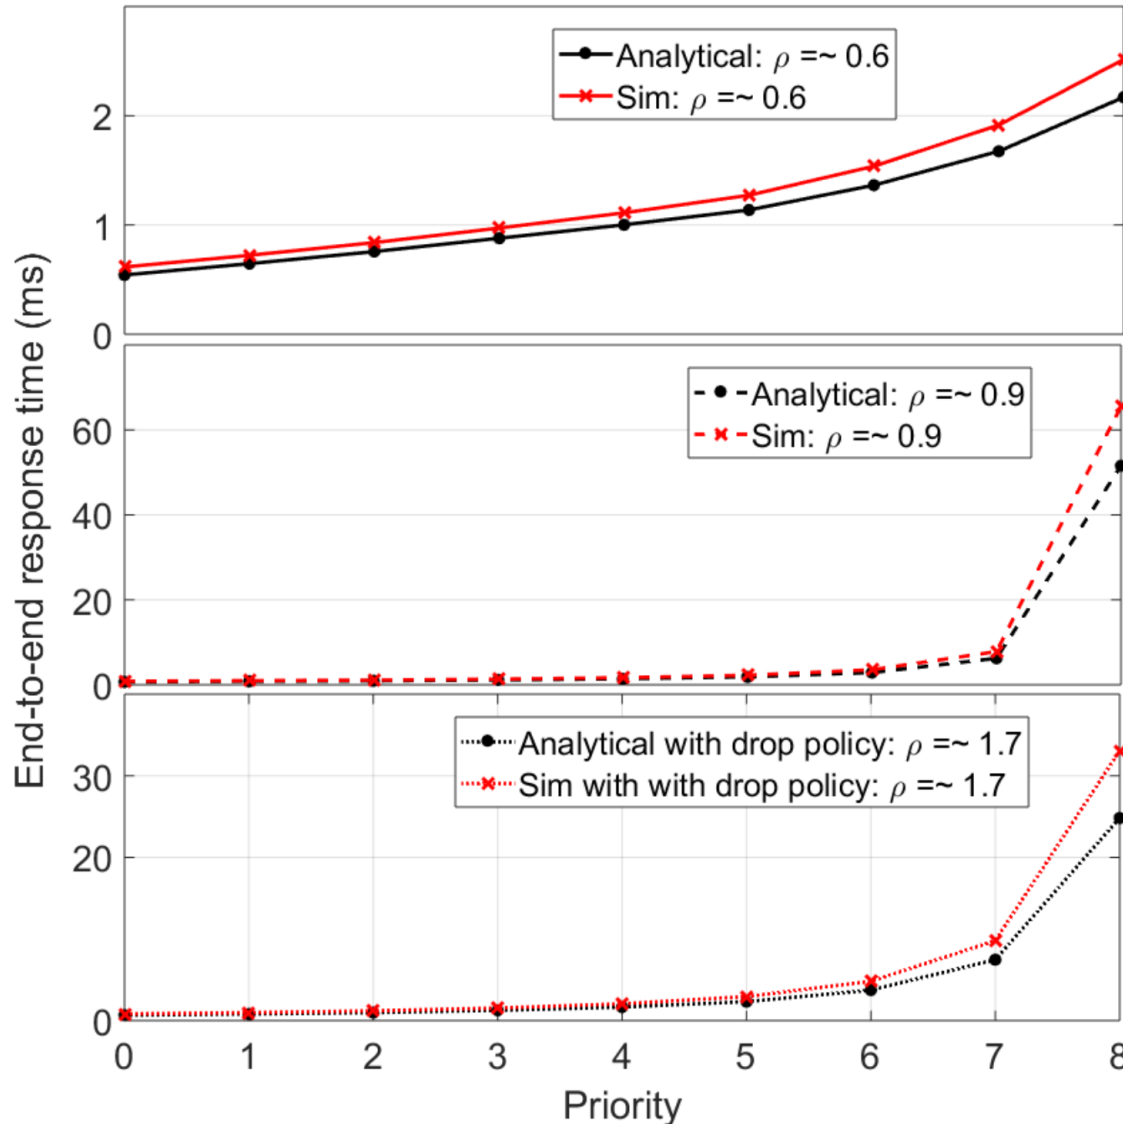
# Experimental setup

- We validate our theoretical model, evaluate the FireDeX approach and compare different prioritization and dropping algorithms.
- We use JINQS (Java Implementation of a Network-of-Queues Simulation) to build our queueing network: an open source simulator for building queueing networks.
  - We have extended JINQS to implement our new multiclass priority queueing model.

Data Exchange Parameters					
	#topics ( $ V $ )	pub rate ( $\lambda_{p,v_i}^{pub}$ )	event size ( $G_{v_i}$ )	#subscriptions ( $ R_{s_i} $ )	utility weight ( $\alpha_{r_i}$ )
Telemetry data	140	$\text{Exp}(\frac{1}{6}) \in [4,7]$	$\text{Exp}(\frac{1}{110}) \in [90,500]$	70	$\text{Exp}(\frac{1}{0.5}) \in [0.01,2]$
Async. events	60	$\text{Exp}(\frac{1}{4}) \in [3,5]$	$\text{Exp}(\frac{1}{800}) \in [500,1100]$	42	$\text{Exp}(1) \in [0.1,4]$

Network Parameters					
#subscribers ( $ S $ )	#publishers ( $ P $ )	#flows ( $ F_{s_i} $ )	#priorities ( $ Y $ )	bandwidth ( $w_{s_i}$ )	$\rho$ tolerance ( $\tilde{\rho}$ )
10	160	9	9	80 Mbps	0.1

# Model validation: varying traffic loads




→ Under-loaded

→ Saturated

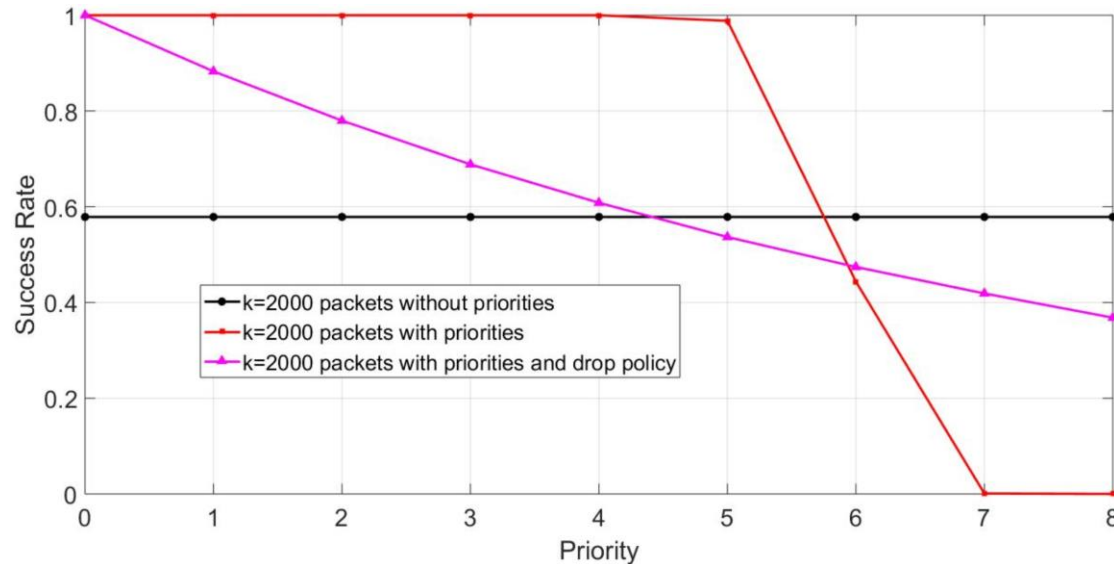
➤ Analytical model for lowest priorities is slightly less accurate.

→ Overloaded

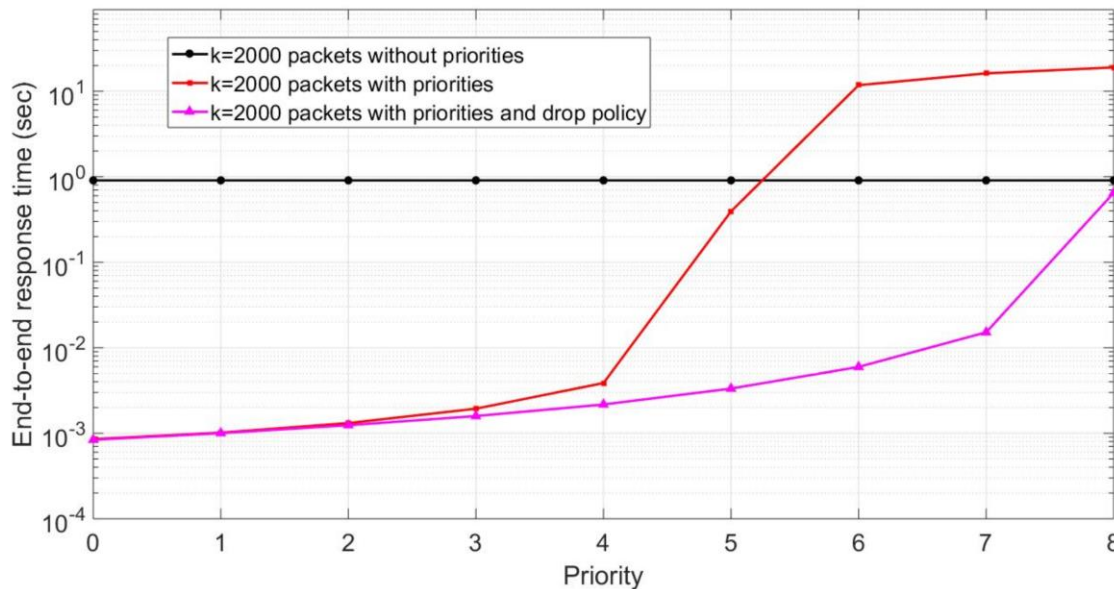
# Model validation: scalability



# FireDeX approach evaluation




- With an **overloaded** system, switch **buffers fill up** and cause high delay / packet drops.
- Our approach **delivers more** high priority events than finite buffers only.



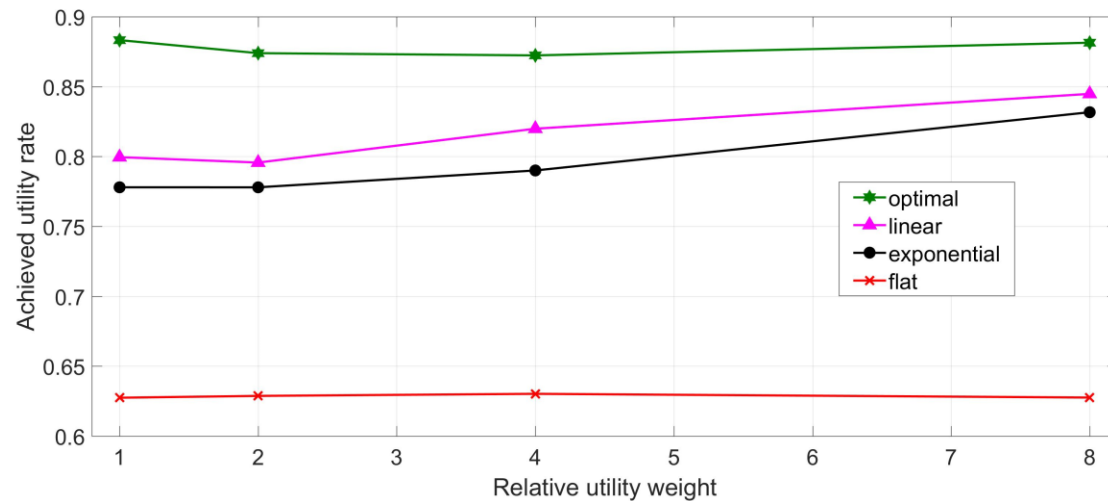
- High priority events also **delivered quicker**.
- Addition of **drop rate** policy smooths success rate while reducing end-to-end delay.

# Algorithms comparison



## Prioritization algorithms:

- Our **bandwidth-aware** greedy strategy performs better than bandwidth-unaware version.
- Both better than **no prioritization**.
- But **random** priorities are worst: need to set priorities correctly!



## Drop rate algorithms

- **Convex optimization** performs best in comparison to linear, exponential and flat policies (drop rates by **assigned priority**).
- Plot shows **varying utilities** of async events vs. data telemetry: simpler policies perform closer to optimal for larger differences.



# Conclusions & Next steps

## Conclusions

- We introduce a middleware that integrates application and network awareness.
- Our application-aware **prioritization algorithm** improves the value of exchanged information by 36% when compared with no prioritization.
- Network-aware **drop rate policies** improve this performance by 42% over **priorities only** and by 94% over **no prioritization**.

## Next steps

Queueing model:

- Consider non-Poisson arrival and service rates by using G/G/1 or G/D/1 queues.

System:

- Alternative utility functions.
- Tuning the entire broker network.
- Use our TIPPERS testbed and CFAST simulator to further evaluate the FireDeX approach.

# Thank you

