Illuminating the hidden challenges of data-driven CDNs

Theophilus A. Benson

Carnegie Mellon University

Digital Services and their Importance

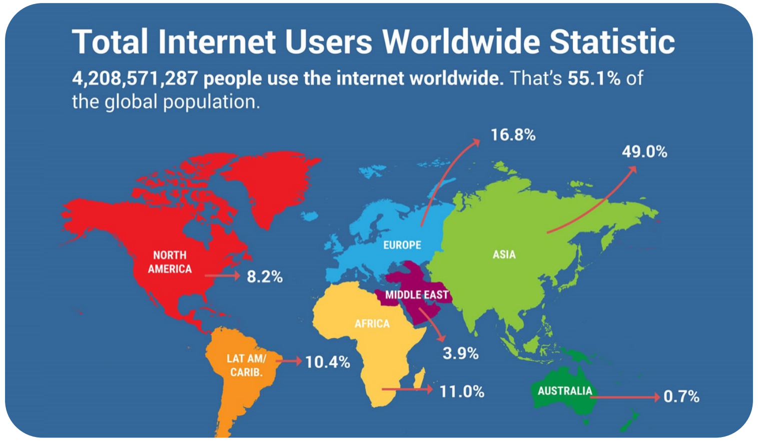

Impacts productivity

"89% of remote workers lost 6.3% productivity due to poor connectivity."

[Zen Internet – broadband from U.K]

Impacts quality-of-life

"19% elevation in stress due to a 6s delay in video loading"

[Ericsson Mobility Report]

Impacts business revenue

"10% users lost for 1s PLT inflation" [BBC study]

"$530K gain for 100ms PLT improvement."

[Mobify report]

"$300K-100M loss for an hour downtime."

[Atlassian report, Fastly study]

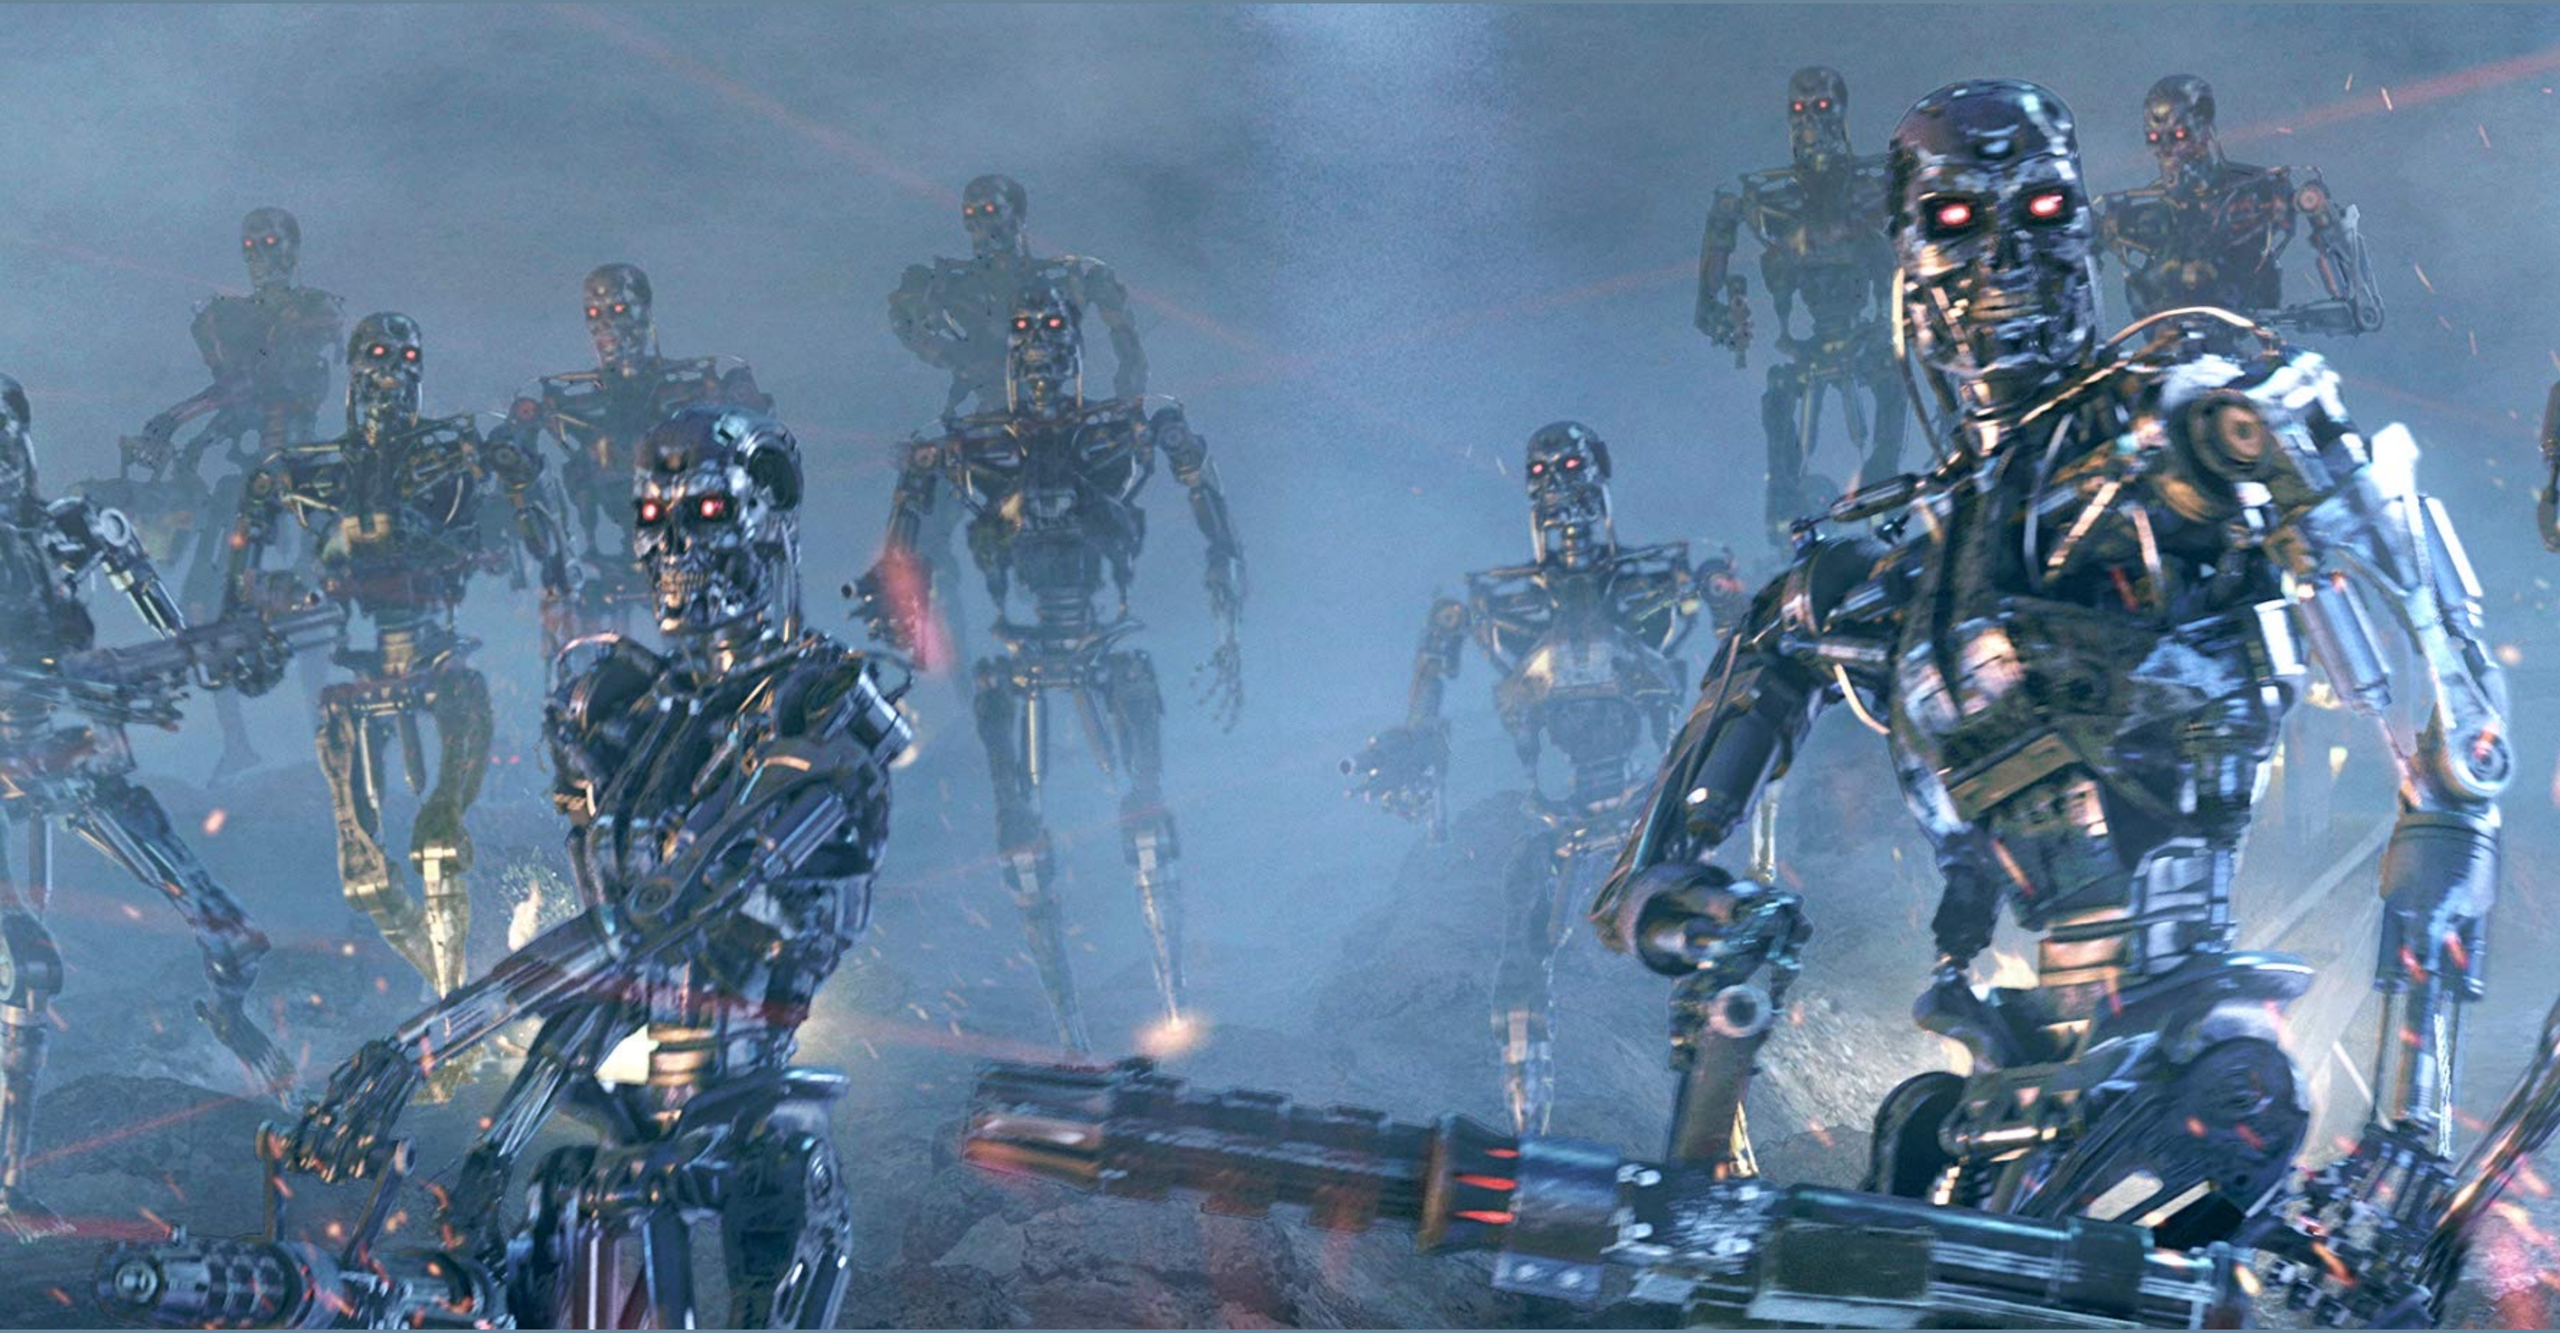

Use of CDNs to Maintain QoE

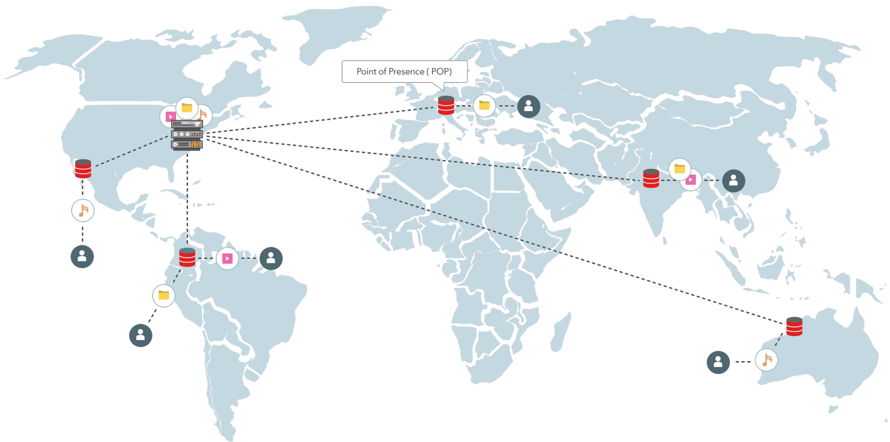

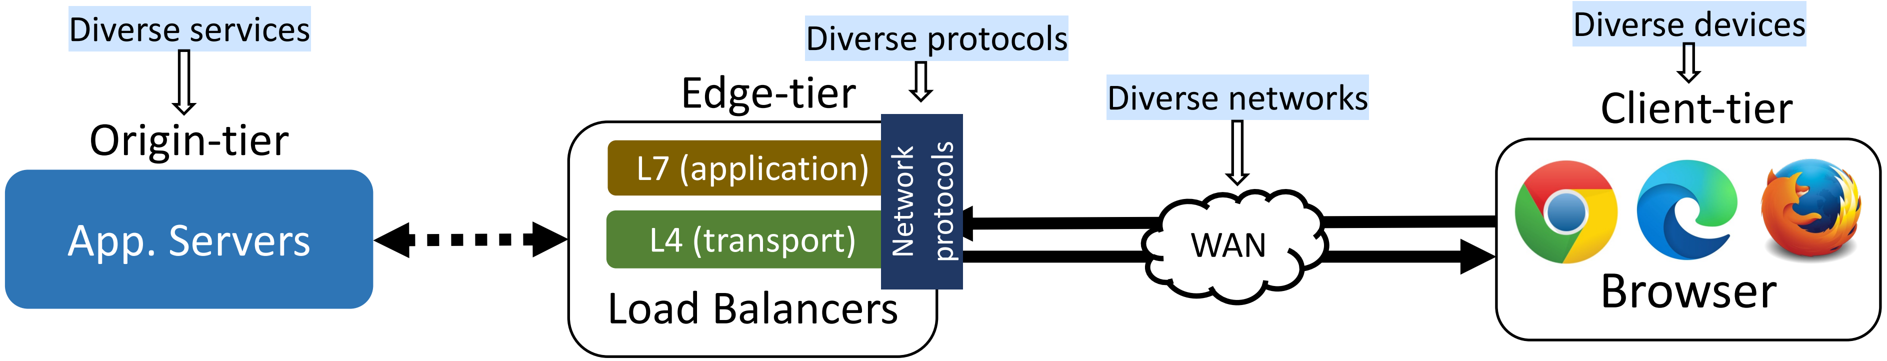

Data-driven == Systematic Approach to Tackle Diversity

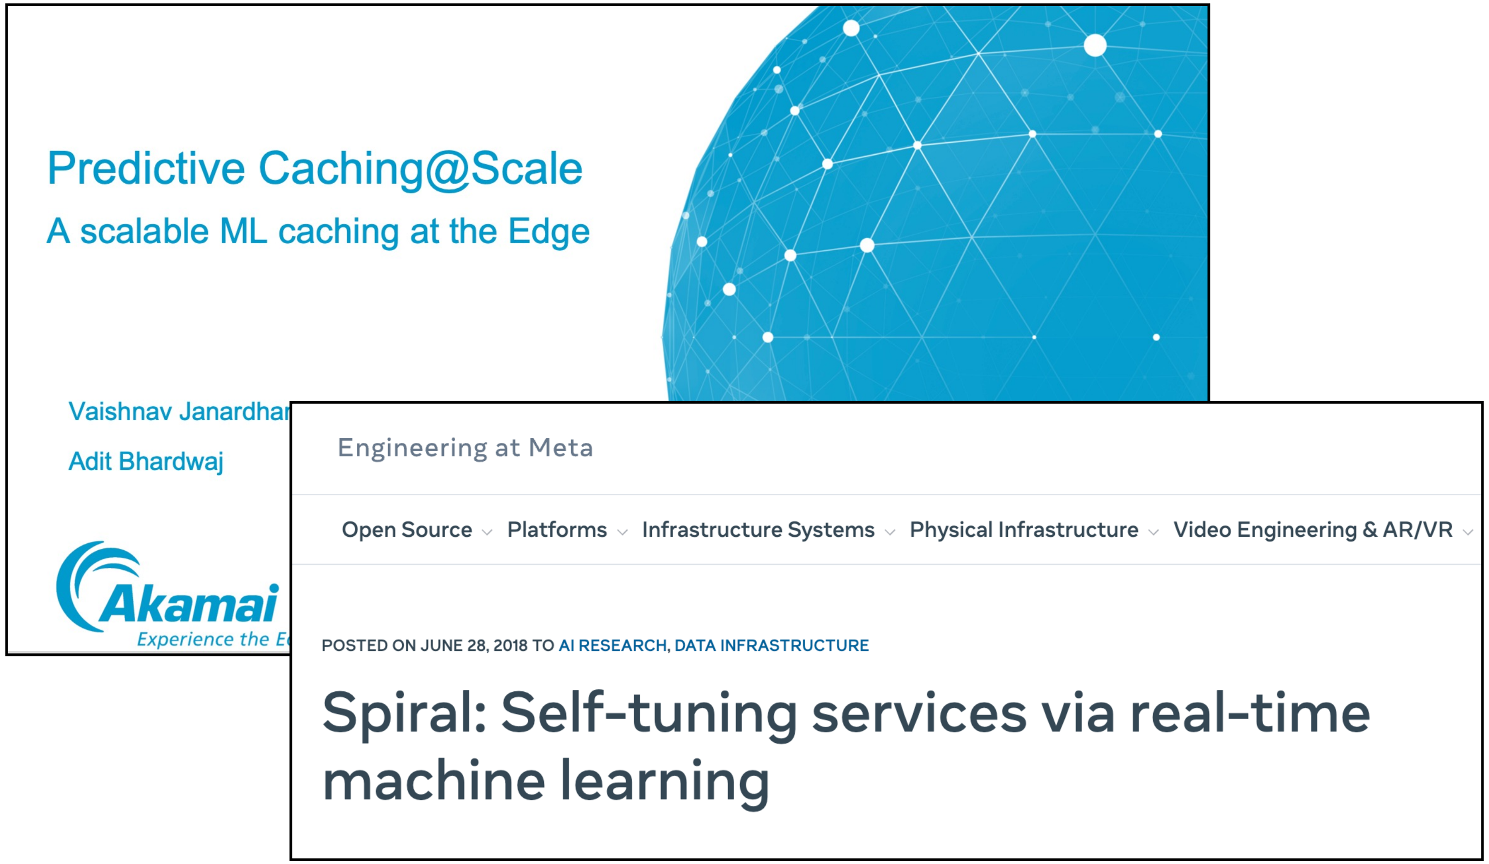

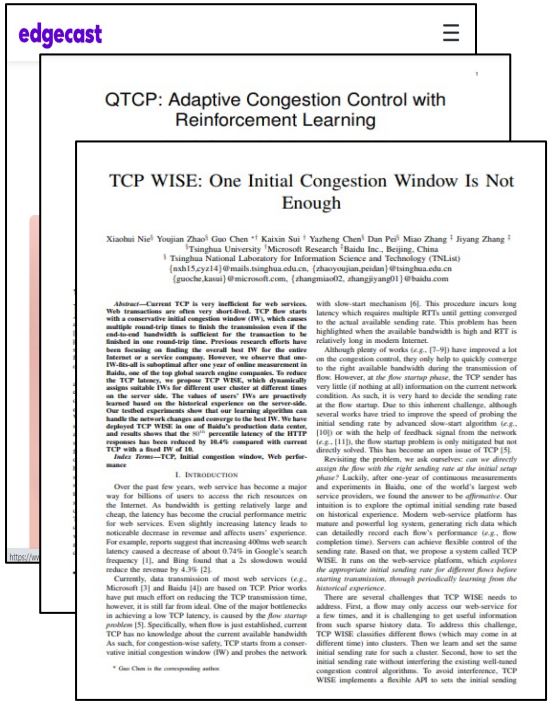

Data-driven == Systematic Approach to Tackle Diversity

Lacking systematic and principled investigations of the deployment implications of data-driven CDNs

Configanator Workflow

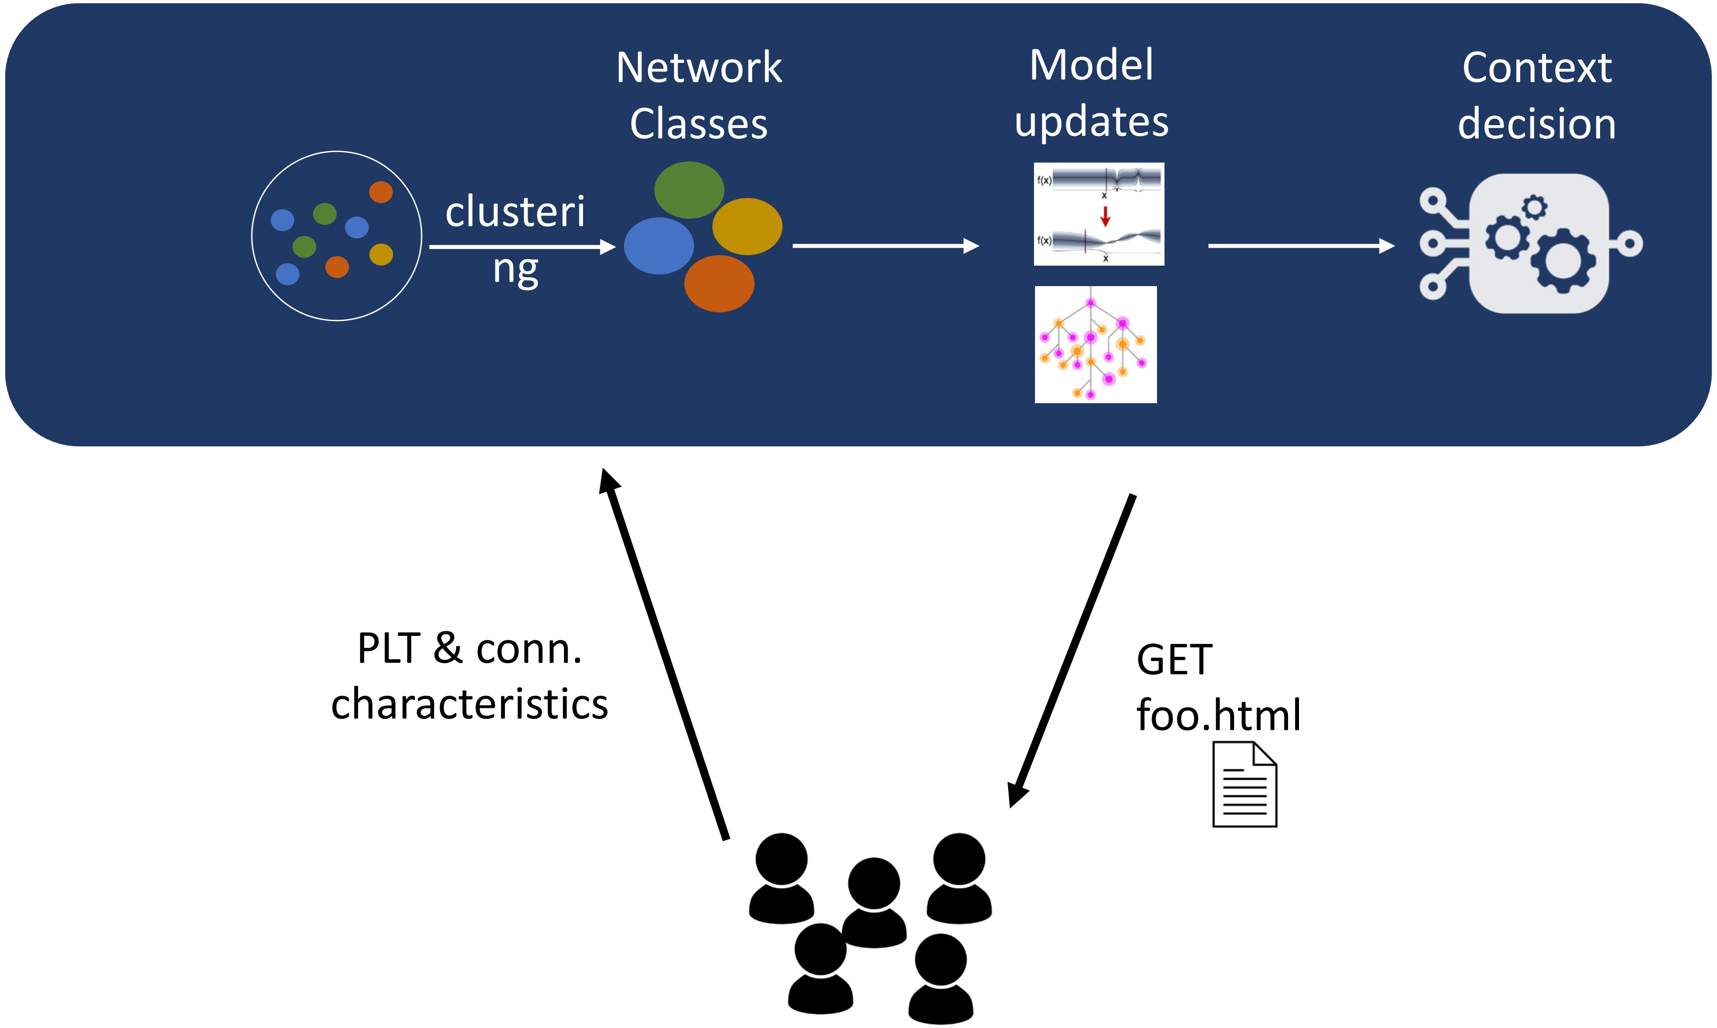

Evaluation Setup

Trace-driven simulation

Traces from multiple regions.

- Global (CDN trace)

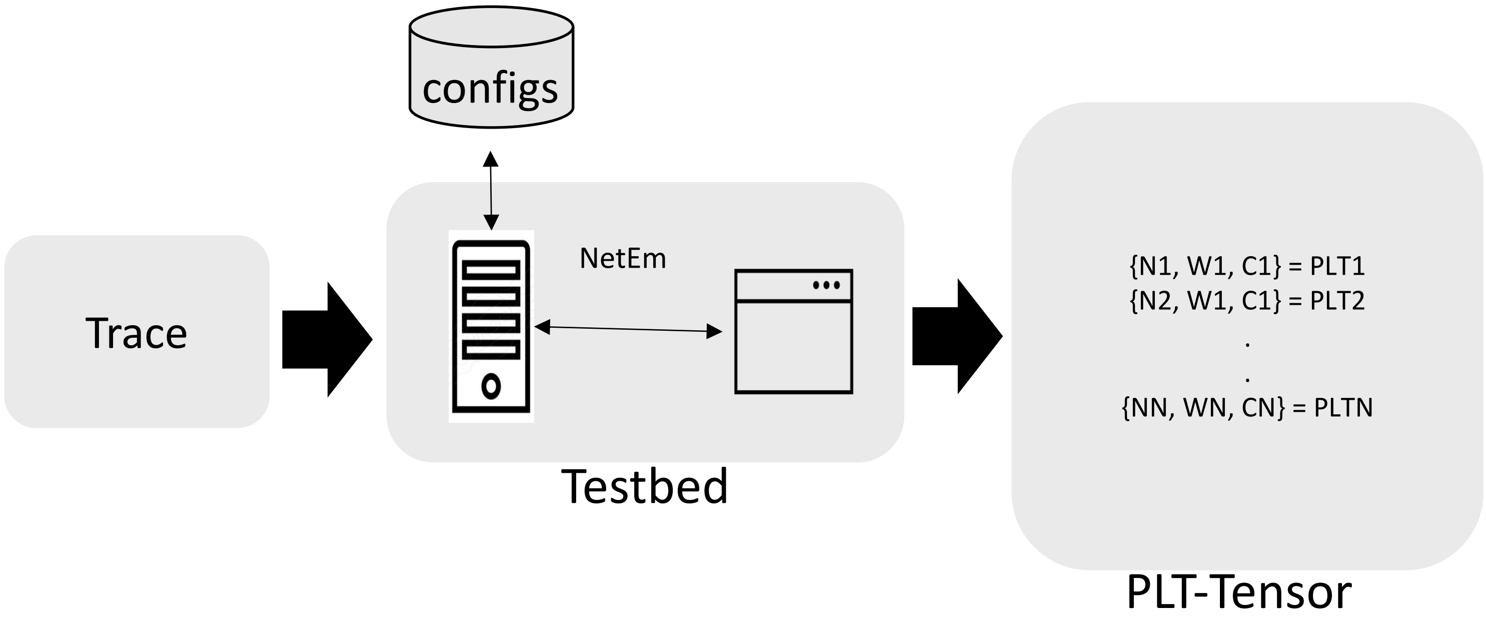

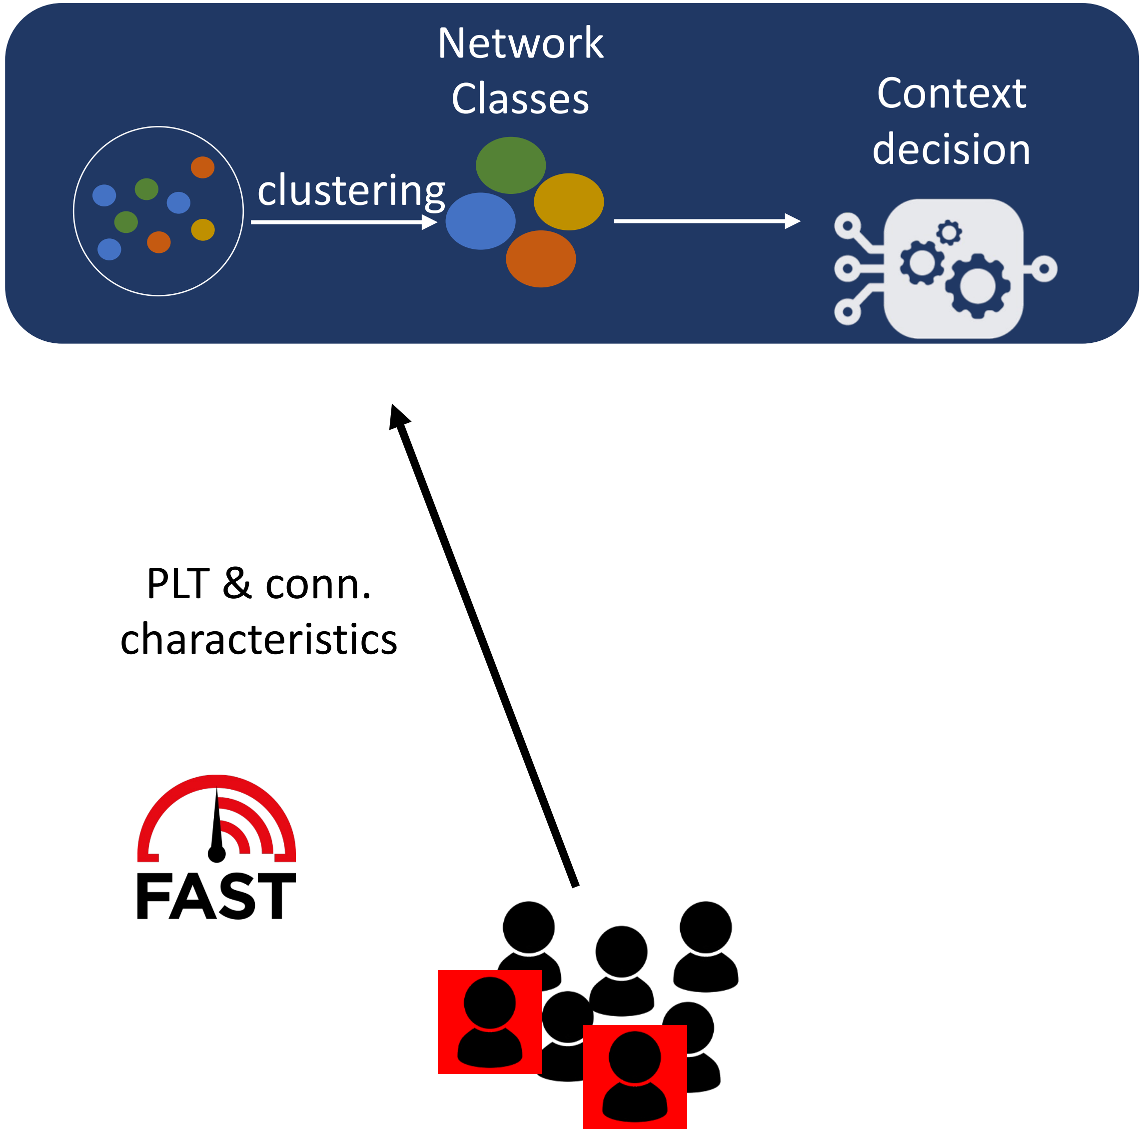
Network
Classes
clustering
Context
decision
PLT & conn.
characteristics
FAST

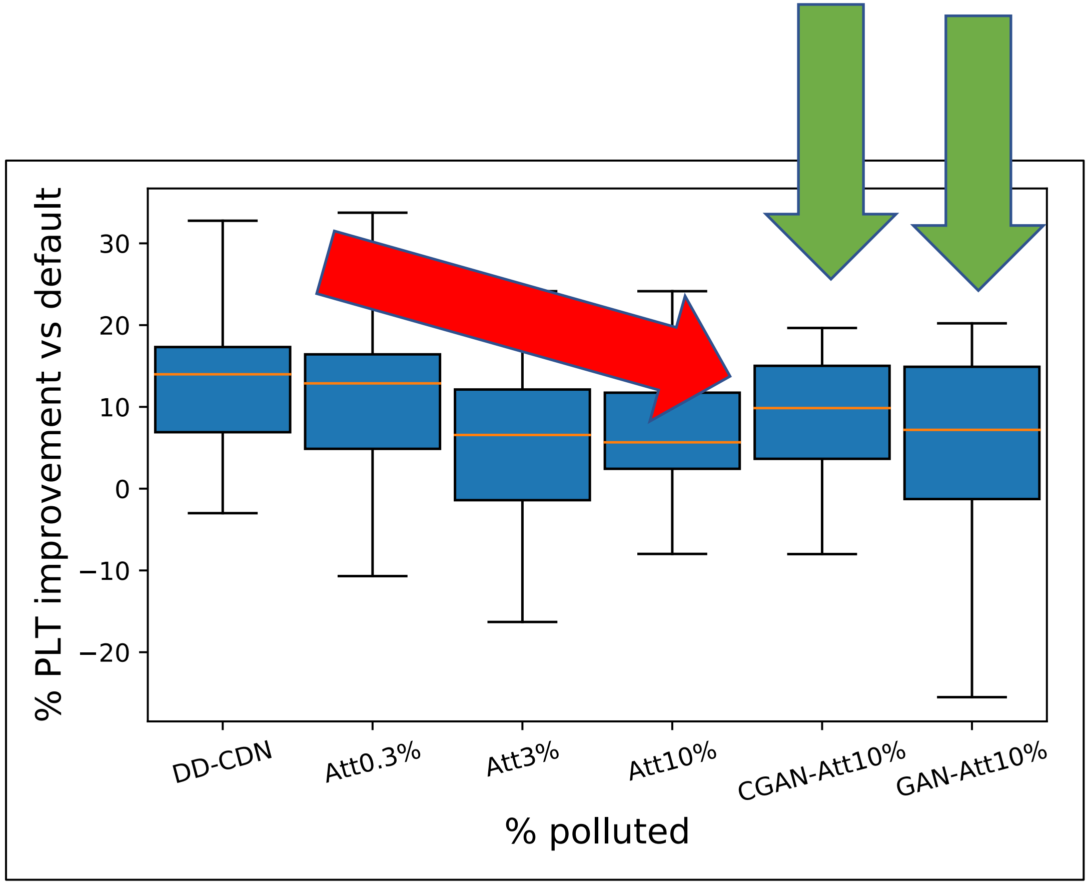
% PLT improvement vs default
30
20
10
0
−10
−20
DD-CDN
Att0.3%
Att3%
Att10%
CGAN-Att10%
GAN-Att10%
% polluted

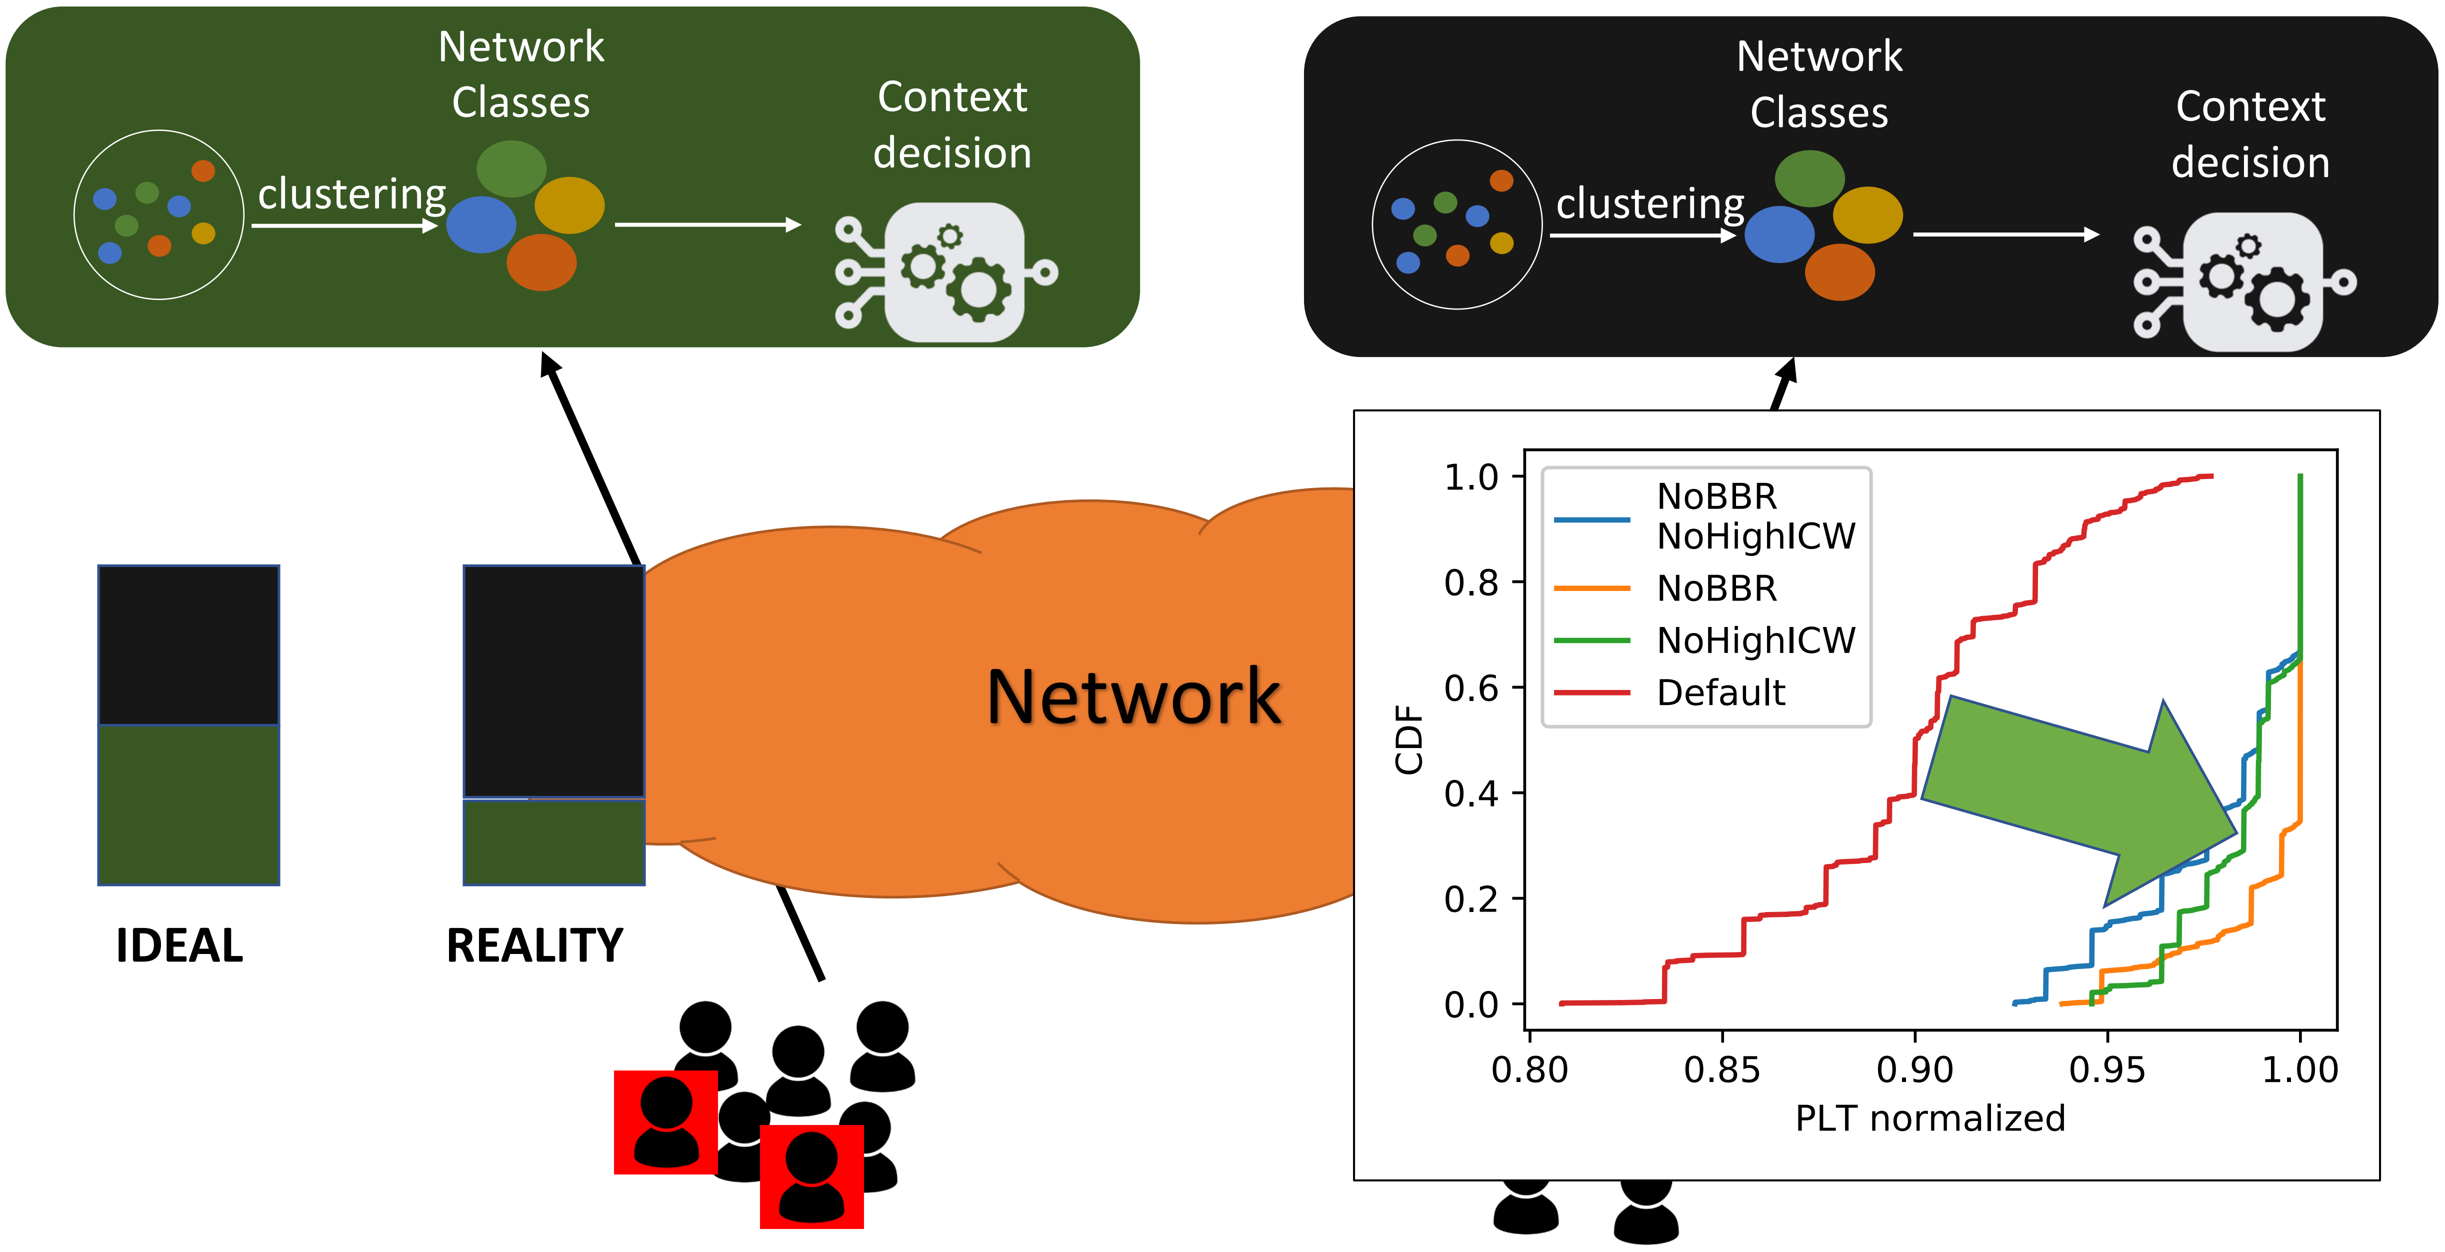

Network Classes
Context decision
clustering
Network Classes
Context decision
clustering
Network
IDEAL
REALITY
1.0
0.8
0.6
0.4
0.2
0.0
CDF
NoBBR NoHighICW
NoBBR
NoHighICW
Default
0.80
0.85
0.90
0.95
1.00
PLT normalized

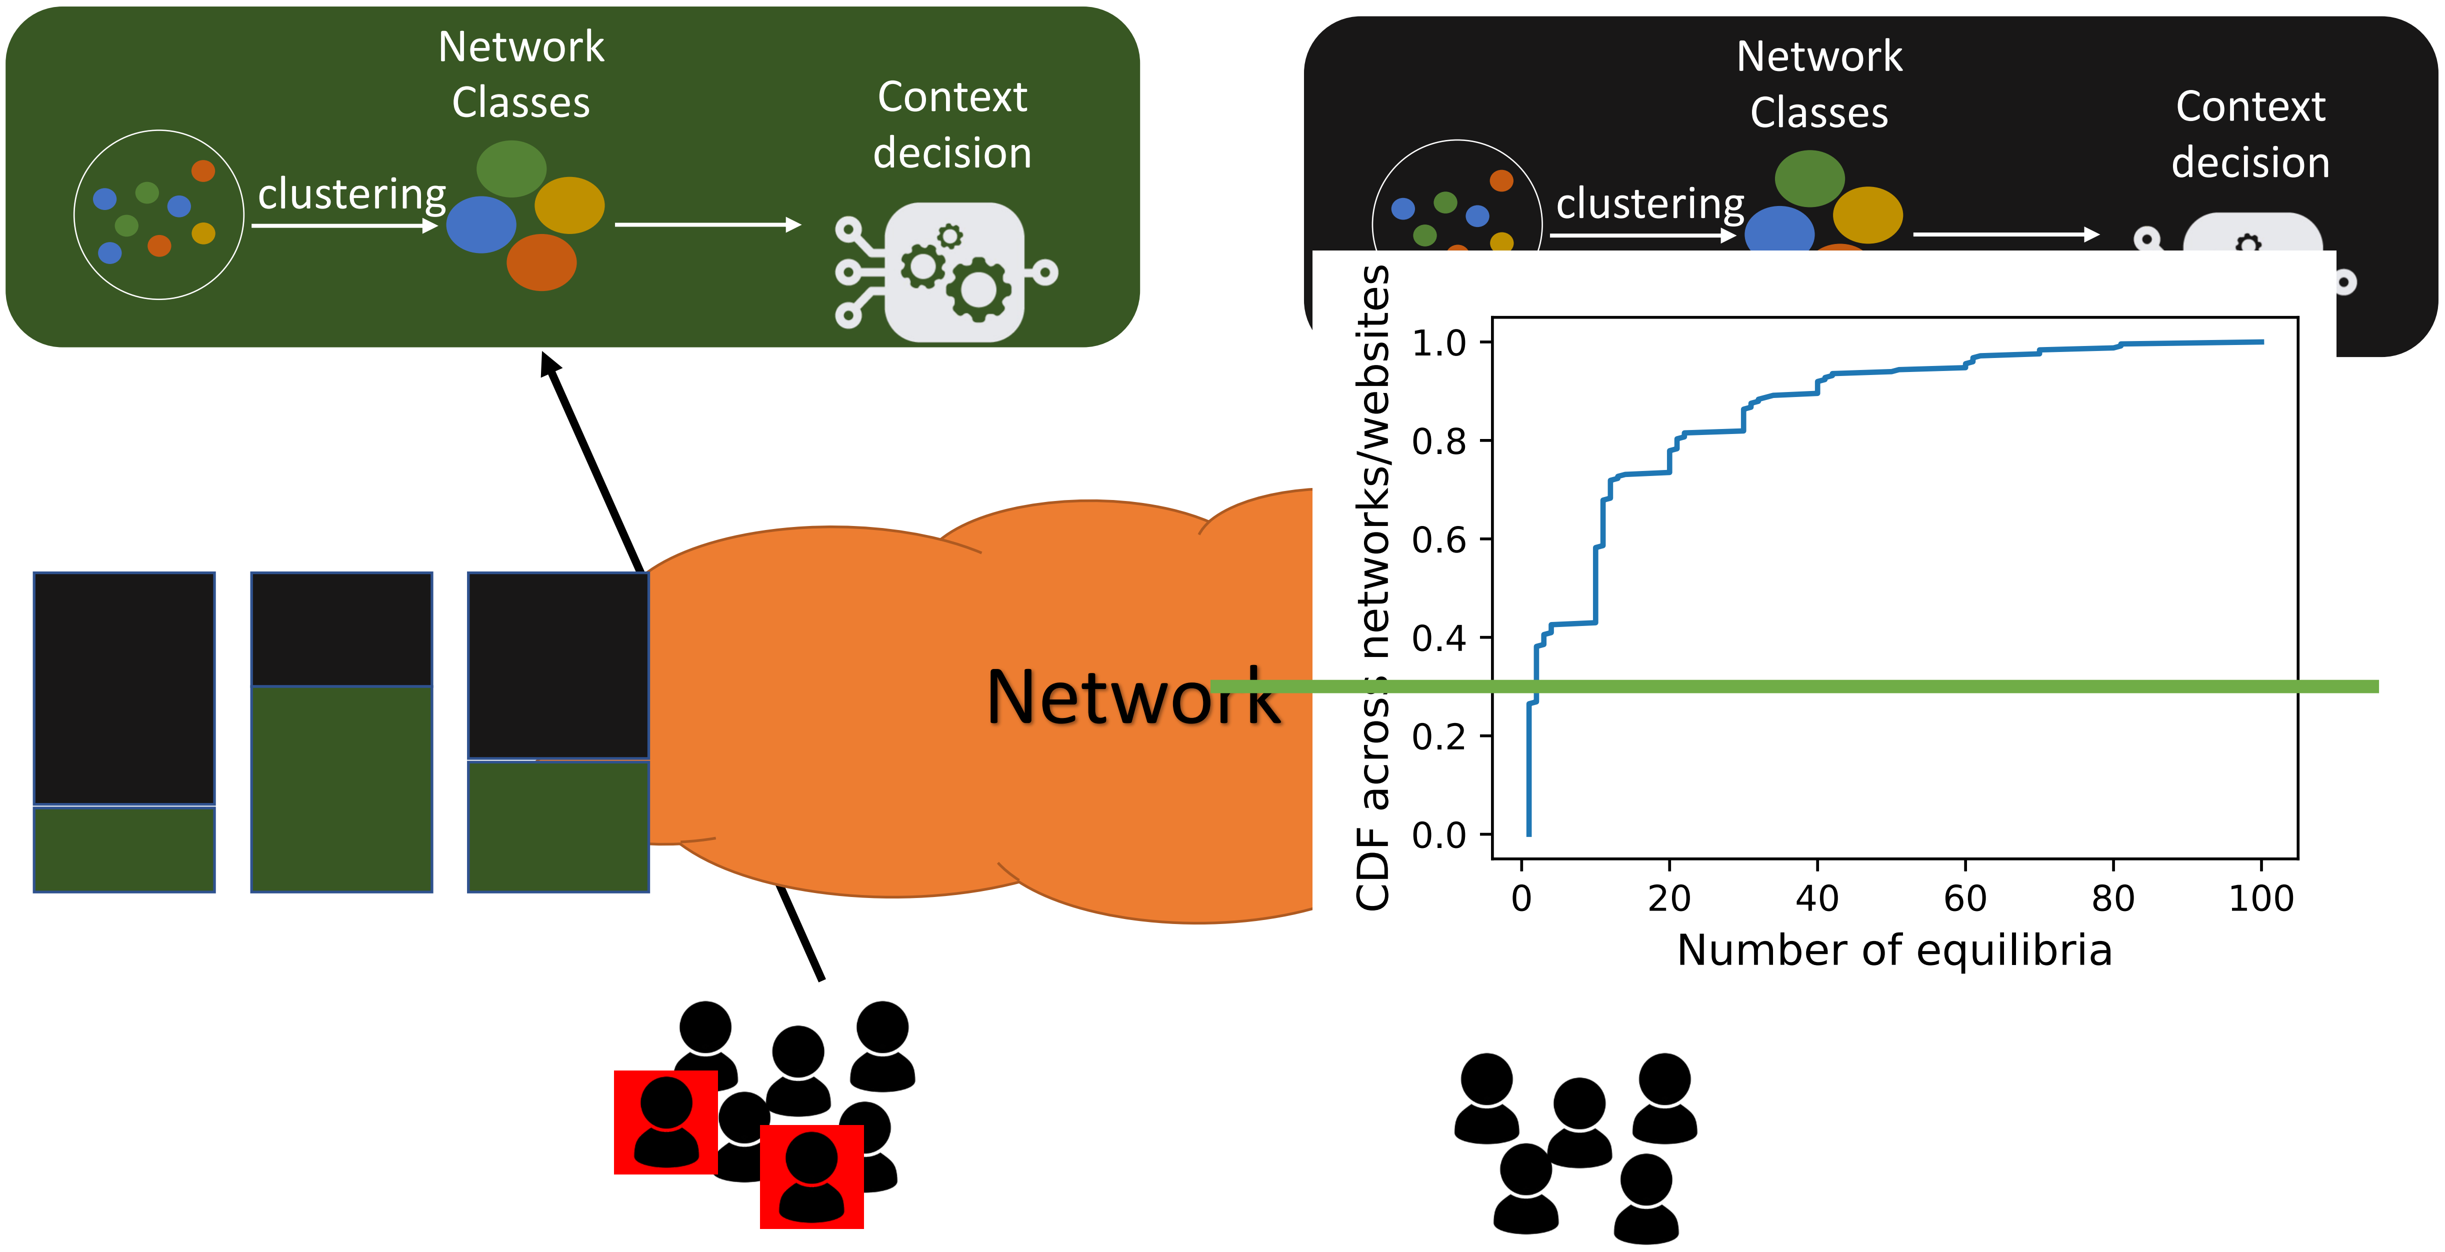

Network
Classes
Context
decision
clustering
Network
Classes
Context
decision
clustering
Network
CDF across networks/websites
1.0
0.8
0.6
0.4
0.2
0.0
0
20
40
60
80
100
Number of equilibria

Security

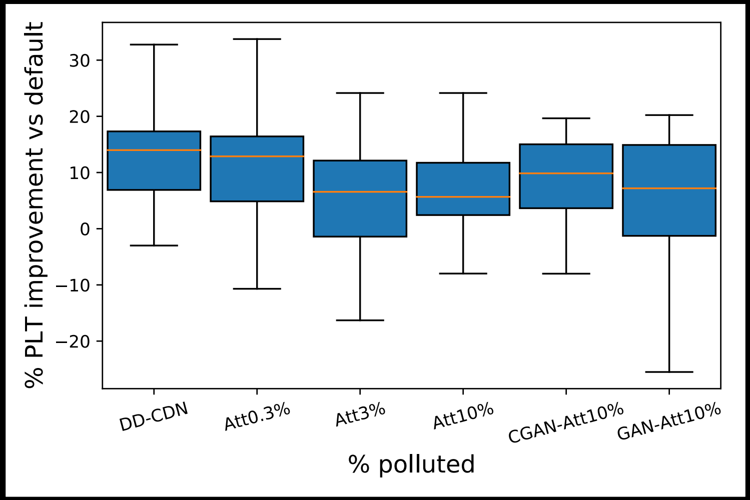

Use side-channel to detect anomalous inputs

Fairness

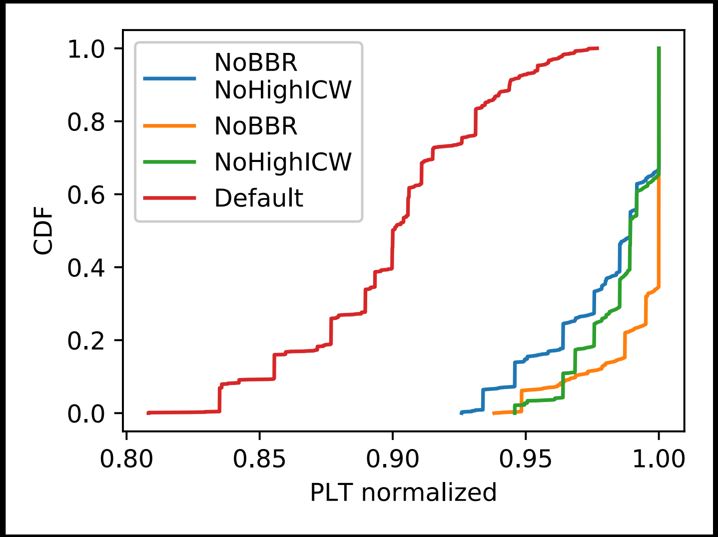

Pre-Filter unfair configs

Stability

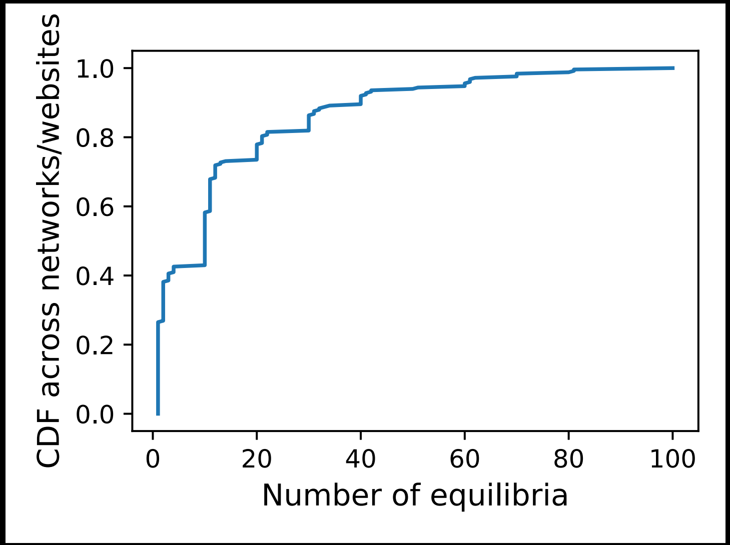

Pre-Filter configs that promote instability

GuardRails

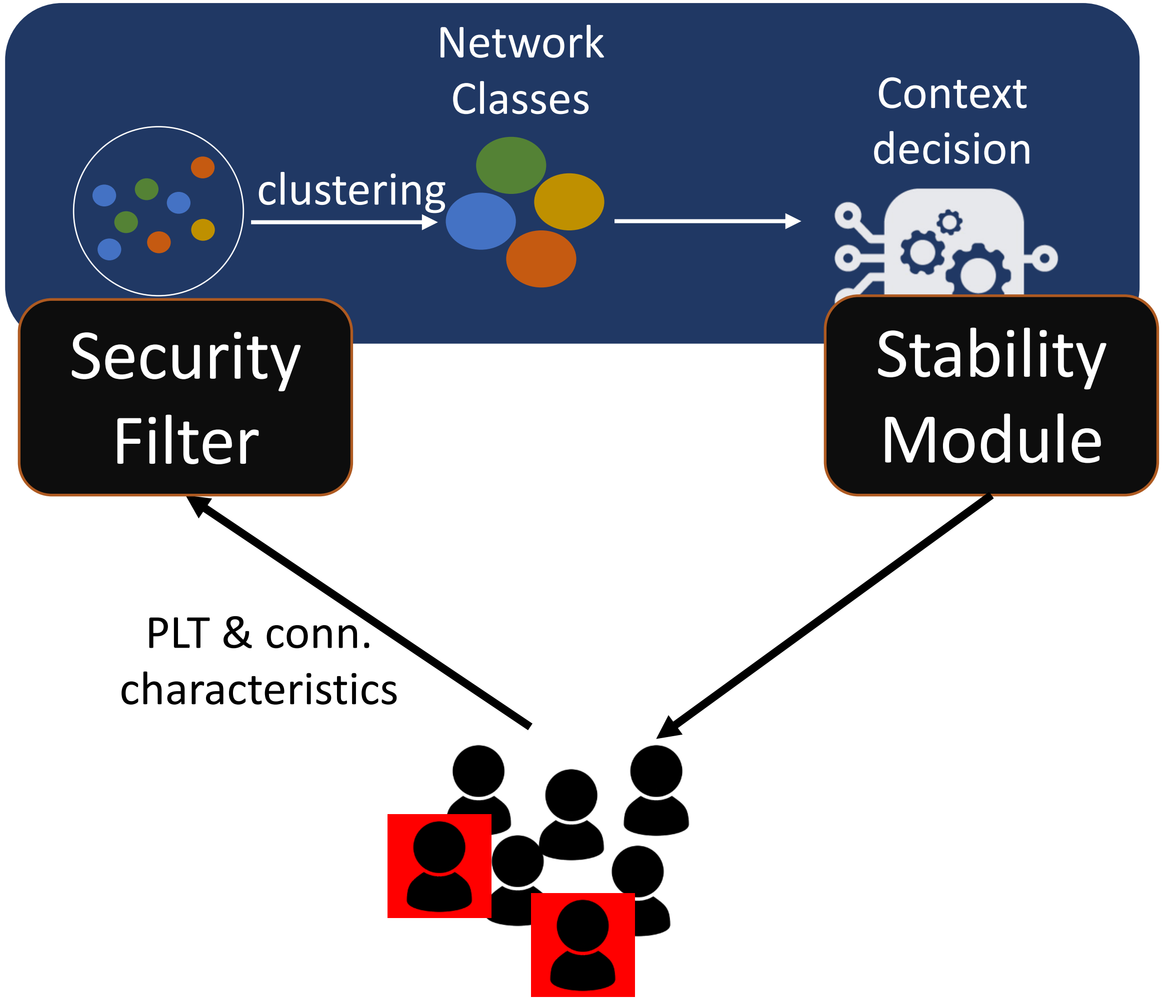

Network Classes
Context decision
clustering
Security Filter
Stability Module
PLT & conn. characteristics

Conclusion + Future Work

Growing deployment of Data-driven CDNs

Empirically investigated deployment challenges

Introduced a strawman solution: GuardRails

Plan to explore infrastructure support