



Orchestrating Applications in the Cloud-Edge Continuum

Giuseppe Bisicchia

Mauriana Pesaresi Seminar Series

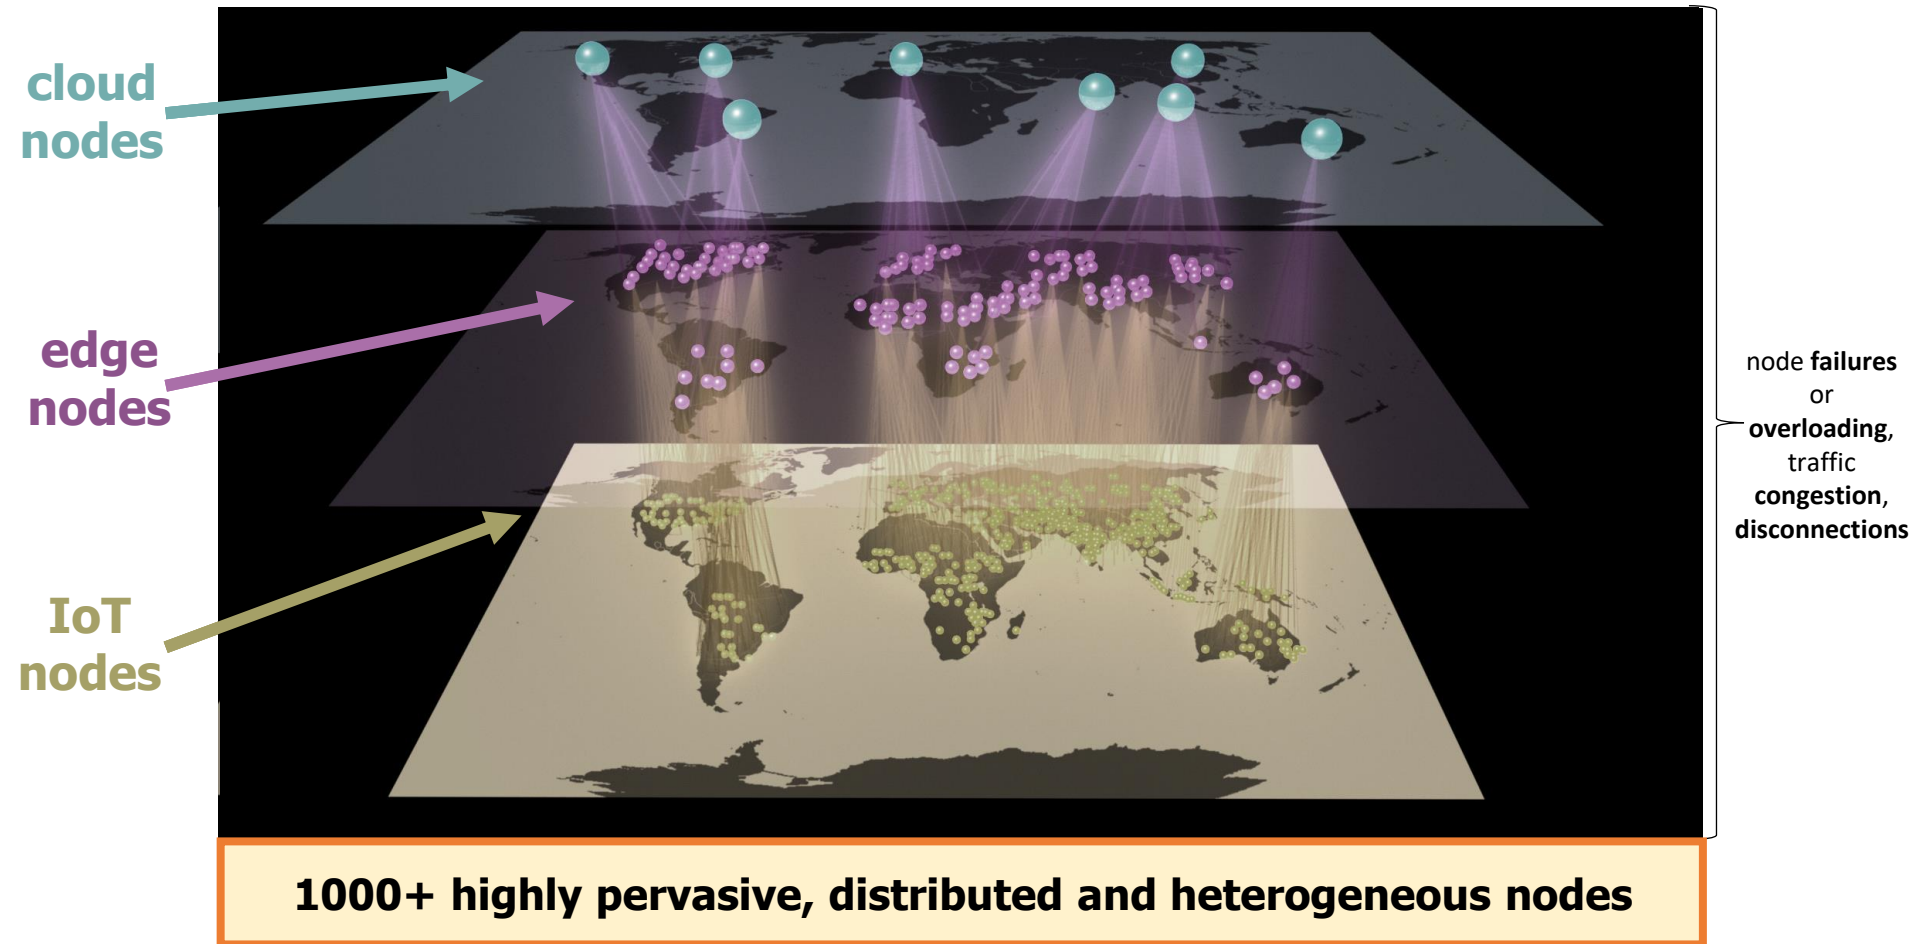
Context: Multi-Service Applications

strictly
hardware and
software
requirements,
QoS
...



100+ interacting (micro)services


Context: Cloud-Edge Infrastructures



Research Context

Cloud-IoT Infrastructures

100+ interacting (micro)services



strictly hardware and software requirements, QoS ...

1000+ highly pervasive, distributed and heterogeneous nodes

Multi-Service Applications

Research Context

Cloud-IoT Infrastructures

100+ interacting (micro)services

cloud nodes

edge nodes

IoT nodes

None of the existing tools supports a continuous, QoS- and context-aware management of microservices on Cloud-Edge infrastructures in continuity with the CI/CD pipeline


1000+ highly pervasive, distributed and heterogeneous nodes

Multi-Service Applications

strictly hardware and software requirements, QoS ...

node failures or overloading, traffic congestion, disconnections

Open Research Problem



How can we ensure the continuous satisfaction of application requirements in a scalable manner over the Cloud-Edge continuum?

Targets

- Reducing **decision-making times** and **unnecessary management operations**
- **Improve** the orchestration **quality** of next-gen **applications** by:
 - **continuously ensuring** their requirements and
 - **balancing** - possibly **contrasting** - **QoS objectives** (e.g., energy/carbon footprint vs operational costs).



Key Factors



Infrastructural
changes



Changes in services
requirements



Addition/removal of
services

Technical Issues



Required **new orchestration processes**, more **lightweight** and better suited for **Cloud-Edge** settings



Should consider **various application paradigms** (e.g. serverless, osmotic)



Support **data migrations** and **scaling** of services

Possible Research Lines



Orchestration
methodologies



Learning
functionalities




Formal Models



Cost models &
heuristics

How can we ensure the **continuous satisfaction** of **application requirements** in a **scalable** manner over the **Cloud-Edge continuum**?




<i>Application Requirements Data</i>	<i>Infrastructure Monitored Data</i>
---	---

Software Reqs	Software Caps
---------------	---------------

Hardware Reqs	Hardware Caps
---------------	---------------

IoT Reqs	IoT Caps
----------	----------

Service-to-Service Bandwidth/Latency	Featured Node-to-Node Bandwidth/Latency
--------------------------------------	---



Reducing **decision-making times** and **unnecessary management operations** while **improving the orchestration quality** of next-gen applications

Thank You