WiOpt 2022

20th International Symposium on Modeling and Optimization
in Mobile, Ad hoc, and Wireless Networks
September 22, 2022

Economics of Multi-Operator Network Slicing

George Darzanos, Iordanis Koutsopoulos, Katia Papakonstantinopoulou, George D. Stamoulis

Department of Informatics, Athens University of Economics and Business (AUEB)

Athens, Greece

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

5G-VINNI

Context

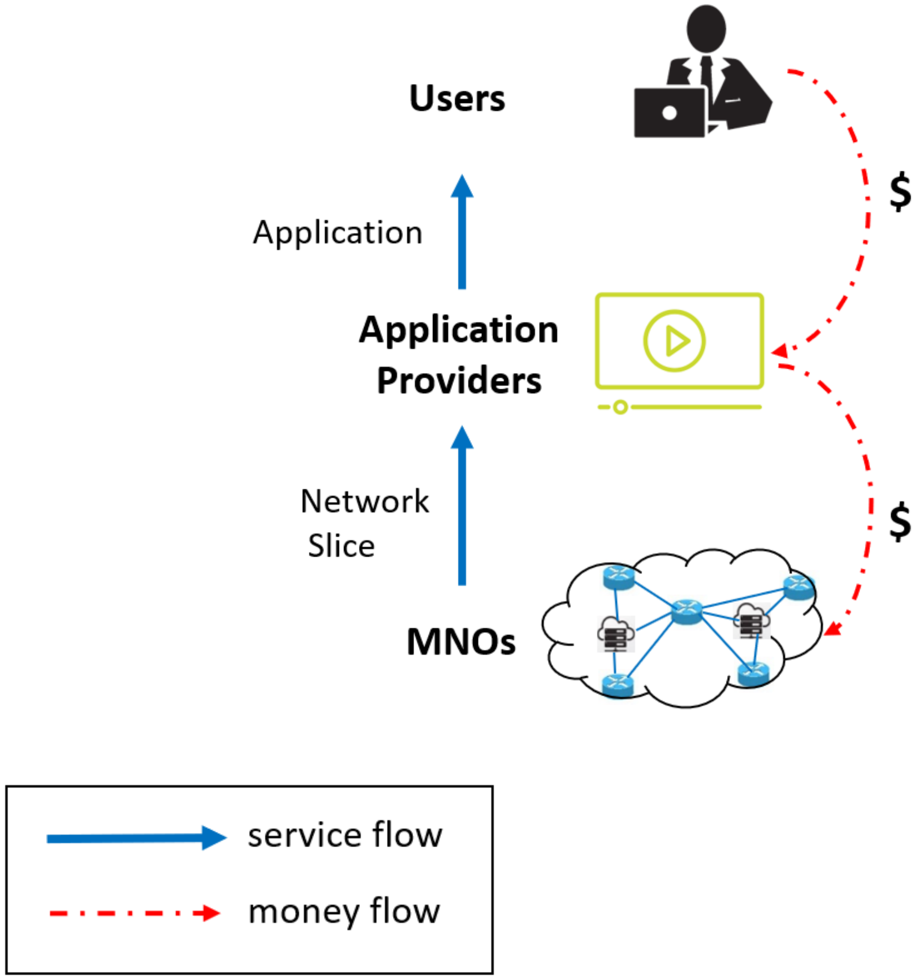

A **5G vertical application** is

- provided by an **Application Provider**
- over a **network slice** offered by one or more **MNOs**
- and consumed by the **Users**.

Network Slice: *A sequence of interconnected **Virtual Network Functions (VNFs)** with strict QoS requirements.*

✓ *enabled by **Virtual Machines (VMs)** and **Virtual Tunnels (VTs).***

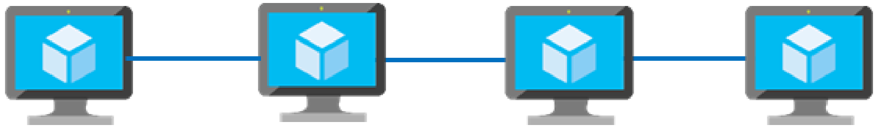

Use case: health sector

- **User:** A **hospital** whose doctors wish to perform surgeries remotely.

- **Application Provider:** Provides such an **e-Health application.**

- **MNOs:** Provide an end-to-end **network slice** from source (surgeon) to destination (hospital).
    - certain **VNFs** must be deployed in the source and destination regions

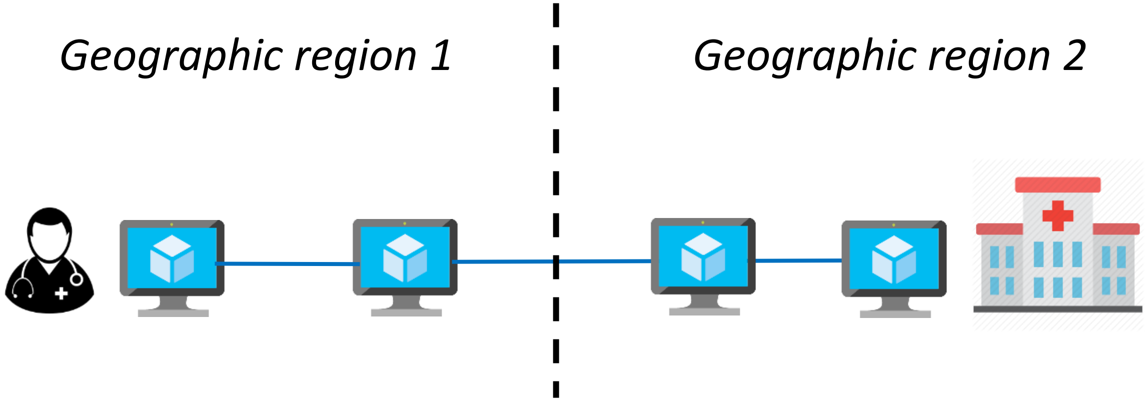

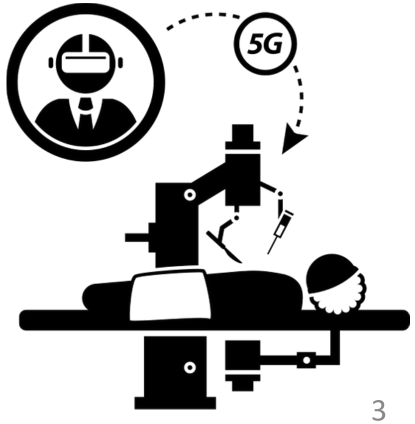

Motivation and Challenges

Motivation. The involvement of **multiple MNOs** is *necessary* for certain applications

- User Equipment (UE) in remote geographic regions

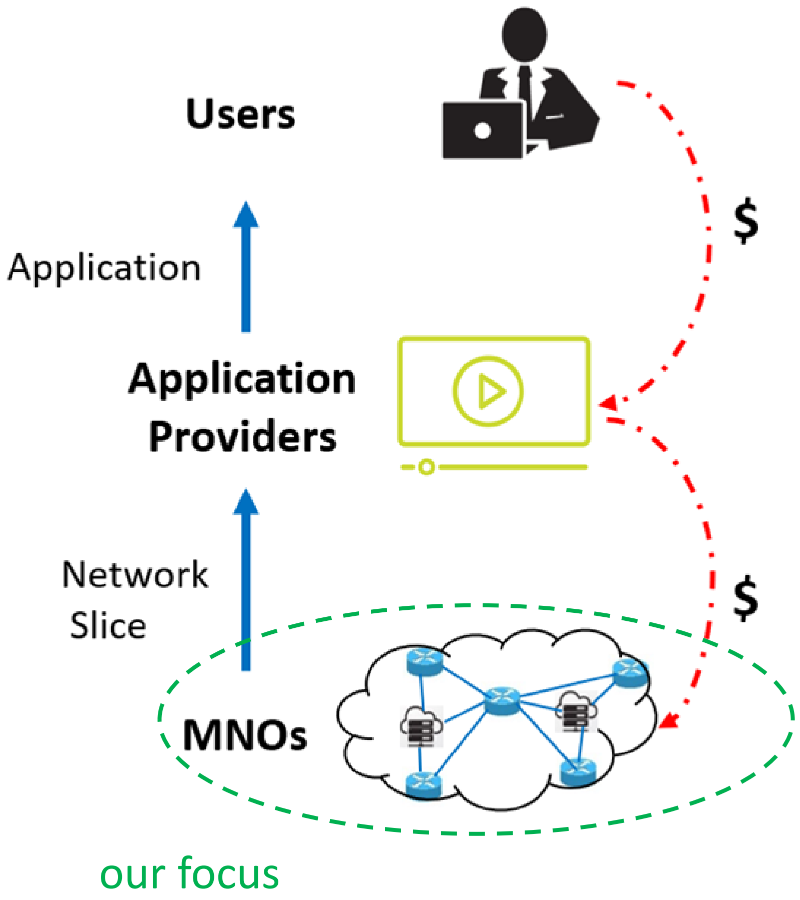

Challenges

- **Competition** among MNOs
- **Lack of information** (e.g., on topology, route availability etc.)

Our Contribution

We introduce policies for the effective **multi-MNO network slice provisioning** under different:

- Ecosystem structures – **Centralized vs Peer-to-Peer**
- Degrees of Information availability – **Information Sharing Mechanism**
- Level of trust and MNOs' collaboration – **Cooperative vs Coopetitive**

System Model

MNOs' Topology

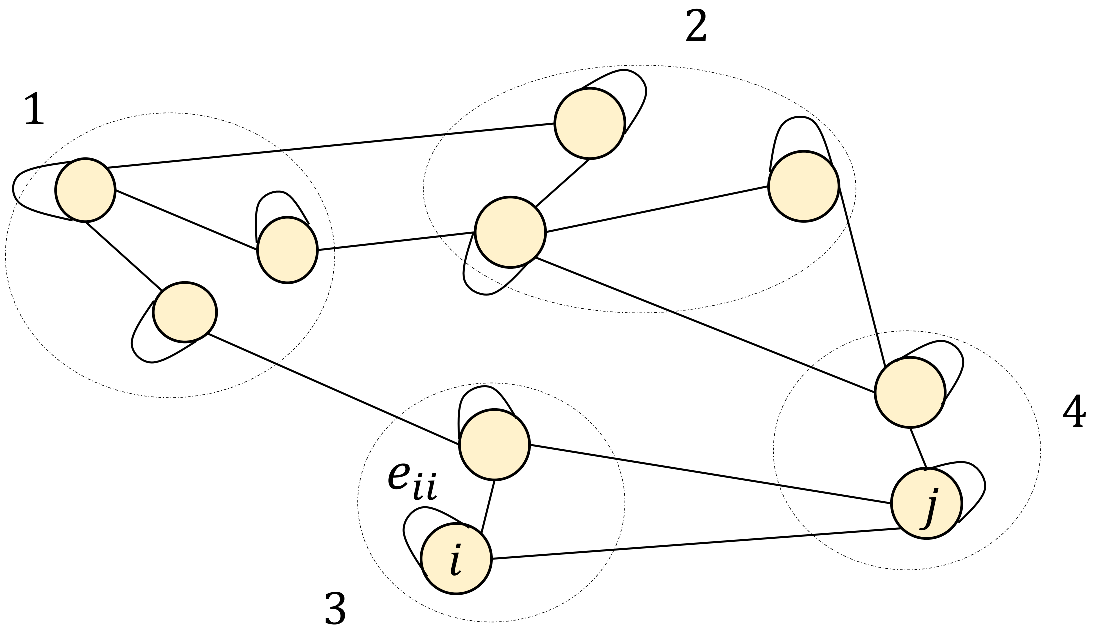

$\mathcal{I}$: set of MNOs

E: set of physical links

$\mathcal{L}$: set of geographic locations

Each MNO (node) $i \in \mathcal{I}$

C_i : CPU cores

$L_i \in \mathcal{L}$: location of MNO's i presence

Each physical link (edge) $e_{ij} \in E$

B_{ij} : Bandwidth (bits)

D_{ij} : Latency (secs)

Network Slice as a Service Model

Network Slice

- **Service-wise** → a chain of interconnected Virtual Network Functions (VNFs) with strict QoS requirements.
- **Resource-wise** → a set of virtualized network, computational and storage resources.

Application requirements → Network slice service requirements → Virtualized resources

Service graph (sequence of interconnected VNFs)

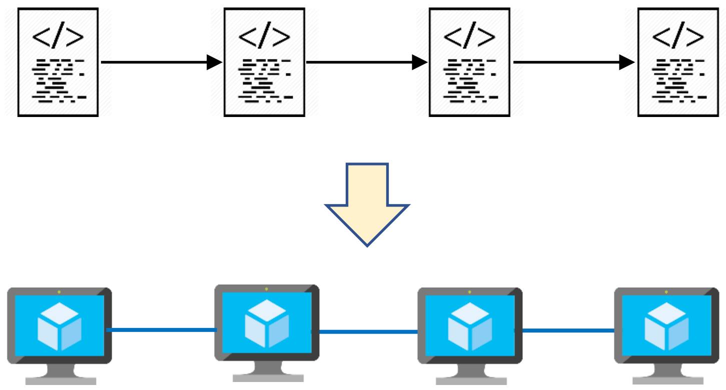

Virtualized resource (VMs and virtual tunnels)

Resource provisioning

✓ Cloud resources as **Virtual Machines (VMs)**

- c_v → CPU cores allocated to VM v

✓ Network resources as **Virtual Tunnels (VT)**

- b_τ → guaranteed bandwidth of VT τ
- d_τ → guaranteed latency of VT τ

Network Slice Request

Network slice request r → sets values on the parameters of a ***service template***

- **Network slice type** (e.g., uRLLC, eMBB, etc.) – type $t \in T$ determines:
  - Set and sequence of VNFs to be deployed, $\mathcal{F}_t$
  - Data packet size K_t
- **Quality class** (e.g., standard, premium) – class $q \in Q$ determines:
  - Target throughput: $B(t, q)$
  - Target latency: $D(t, q)$
- **Region of source and destination**
- **VM placement restrictions**
- **Traffic volume**
- **Price**

Multi-MNO Network Slicing Example

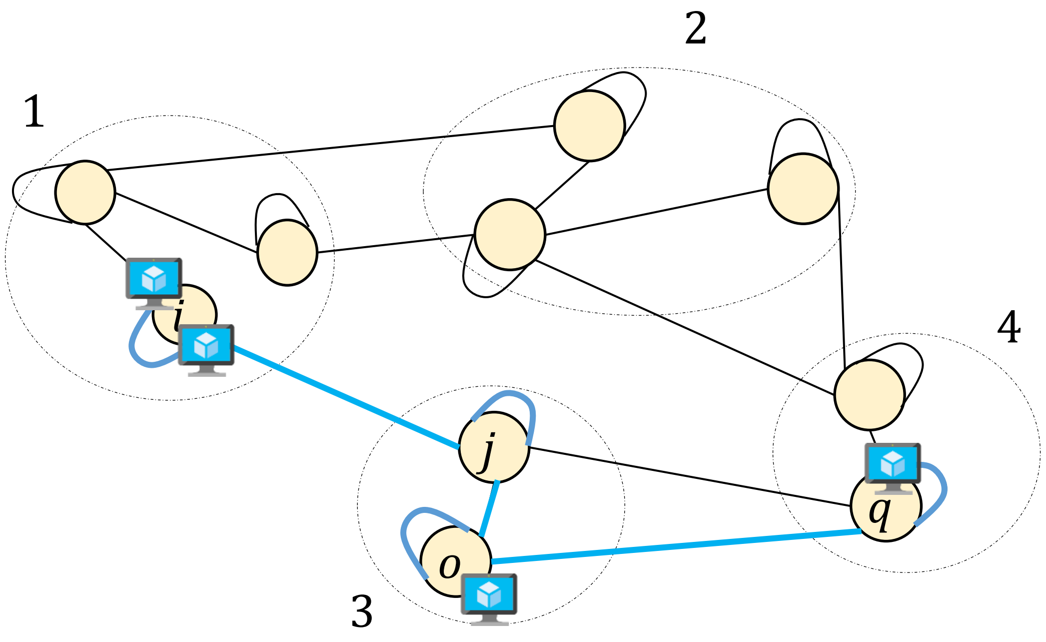

Assuming that there is a **request** r for a network slice

- with certain **QoS requirements**
- with **location 1** as source and **location 4** as destination

Decisions to be made:

1. **Dimensioning process** – *(agnostic to the topology)*
   - **Resources** that should be allocated in each **VM** v (i.e., c_v) and **VT** τ (i.e., b_τ)
   - Output →

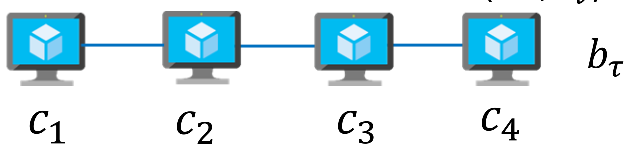

2. **Embedding process** – *(considering the topology)*

Decision Variables

   - **Placement** of **VMs** and **VT** over the topology - set of paths $\mathcal{P}$
   - $x_{v,i} \in \{0,1\}$ → determines if VM v (with capacity c_v) is placed in MNO i
   - $y_{\tau,\pi} \in \{0,1\}$ → determines if VT τ (with bandwidth b_τ) is placed over path $\pi \in \mathcal{P}$

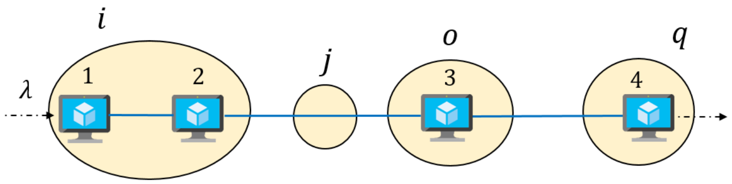

Network Slice Dimensioning

- **Dimensioning process** for network slice request r based on the *template inputs*.
  - **The number of VMs** to be deployed is determined by the service type t_r.
  - **Bandwidth** to be allocated across the links that the **virtual tunnel** traverses

$$b_{\tau_r} = B(t_r, q_r)\, K_{t_r}$$

Target throughput

Packet size

  - **Computational resources** to be allocated to all **VMs** of r, $\boldsymbol{c}_r = \{c_v\}_{v \in \mathcal{V}_r}$
    - ✓ i.e., the number of **CPU cores** for achieving throughput $B(t_r, q_r)$ in each VM v

$$\mu_v(c_v, f) = B(t_r, q_r)$$

VM service rate: $\mu_v(c_v, f) = \sigma_f c_v$

Embedding Process

Approaches

Centralized

- **A Broker** determines the embedding of all network slice requests $\mathcal{R}$
  - The **Broker** has full information
  - Is the contact point for **Application Providers**

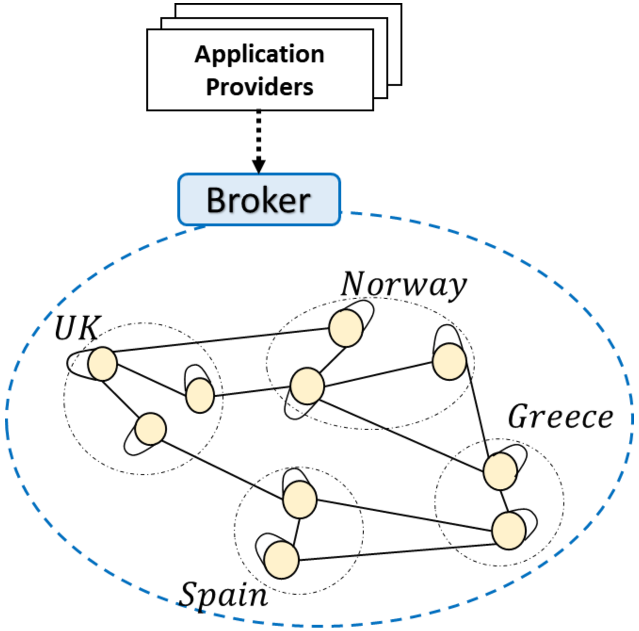

Peer-to-peer

- **Each MNO** determines the embedding of **requests from his own customers** $\mathcal{R}_i \subseteq \mathcal{R}$
  - Distributed **information sharing mechanism**
    - ✓ *Each MNO may have incomplete information*
    - ✓ *Tunable level of information availability*

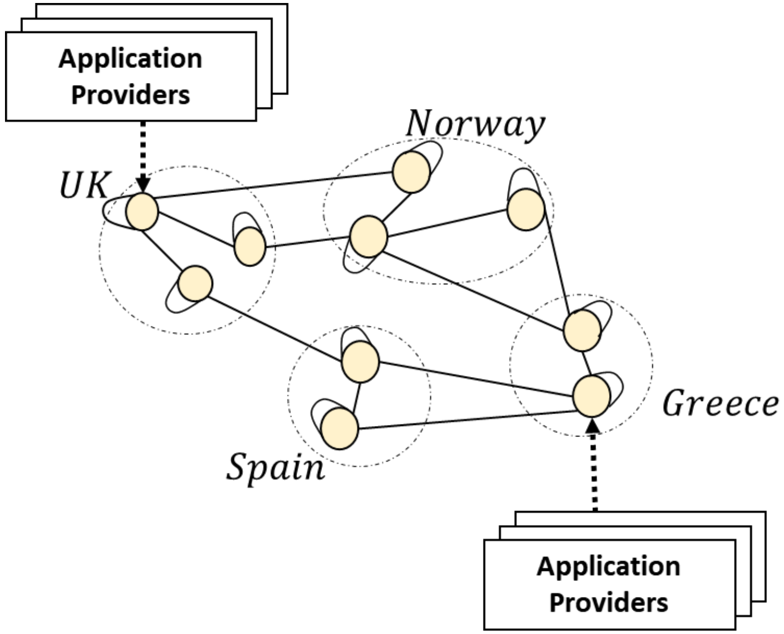

Centralized Approach (I)

Cooperative Policy: Broker solves a **global total Profit maximization** problem → Mixed Integer Program

$$\max_{\mathbf{X},\mathbf{Y}} \sum_{r\in\mathcal{R}} \Bigg[\overbrace{\sum_{\pi\in\mathcal{P}} y_{\tau_r,\pi}\, \hat{p}_r}^{\text{APs' payments}} \overbrace{- \sum_{i\in\mathcal{I}} \sum_{\nu\in\mathcal{V}_r} x_{\nu,i}\kappa_i(c_\nu) - \sum_{\pi\in\mathcal{P}} y_{\tau_r,\pi} \sum_{e_{ij}\in\pi} \kappa_{ij}(b_{\tau_r})}^{\text{Cost of resources}} \Bigg]$$

Decision variables

$x_{\nu,i} \in \{0,1\}$ → **VMs** placement

$y_{\tau_r,\pi} \in \{0,1\}$ → **VTs** placement

Constraints

- Infrastructure **capacity**
- VMs and VTs **unique placement**
- VM placement **location** restrictions
- VMs and VTs **alignment**
- End-to-end **latency**
- **Prices** that the Application Providers are willing to pay

Revenue sharing (per request) → $\hat{p}_{i,r}(\mathbf{X}^*,\mathbf{Y}^*) = \underbrace{K_{i,r}(\mathbf{X}^*,\mathbf{Y}^*)}_{\text{Cost of } i} + \frac{K_{i,r}(\mathbf{X}^*,\mathbf{Y}^*)}{\sum_{i\in\mathcal{I}} K_{i,r}(\mathbf{X}^*,\mathbf{Y}^*)} \underbrace{S_r(\mathbf{X}^*,\mathbf{Y}^*)}_{\text{Profit from } r}$

Centralized Approach (II)

Coopetitive Policy: Broker solves a **local Profit maximization** per **MNO** – for $\mathcal{R}_i \subseteq \mathcal{R}$ and $\mathcal{P}_i \subseteq \mathcal{P}$

- Local problems are solved sequentially in a **Round-Robin** approach
- **MNOs** publish **prices** instead of costs
- Assumption: All MNOs follow a common pricing scheme -- *Infrastructure utilization-driven pricing*

$$p_i(c_\nu) = c_\nu \; p_{i,co} \left[1 + \frac{\log(C_i - \tilde{C}_i + c_\nu)}{\log(C_i)} \right] \longrightarrow \text{CPUs}$$

$$p_{ji}(b_\tau) = b_\tau \; p_{ji,bw} \left[1 + \frac{\log(B_{ji} - \tilde{B}_{ji} + b_{\tau_r})}{\log(B_{ji})} \right] \longrightarrow \text{Bandwidth}$$

Peer-to-peer Approach

Coopetitive Policy: Each **MNO** i solves a **Local Profit maximization** problem for the requests of his own customers, $\mathcal{R}_i \subseteq \mathcal{R}$.

- Difference with Centralized Coopetitive → Lack of full information with respect to paths
- **Information sharing mechanism** (inspired by Border Gateway Protocol - BGP)
  - Each MNO maintains a **table of preferable paths** **–** *e.g., **cheapest** and **feasible** paths*
    - ✓ *α paths for each combination of (destination location, slice type, quality class)*
  - Each **path** is characterized by **attributes** that capture
    - Estimated end-to-end latency
    - Estimate throughput (bottleneck)
    - Average resource unit price (computational, network)

Main MNO processes

- Path augmentation
- Update rule for preferable paths
- Path forwarding

Peer-to-peer Approach (II)

Information sharing mechanism - Processes

- **Path Augmentation process.** Assuming that MNO i received a path $\pi^{'}$ from "neighbor" MNO j, he generates the path $\pi = \pi^{'} \cup \{e_{ji}, e_{ii}\}$
  - Estimated end-to-end latency**:** Additive to the one receive by j
  - Estimated throughput: Re-evaluates the minimum throughput (bottleneck) across the path
  - Average unit prices of computational and network resources across the path

- **Update rule for preferable path →** Cheapest **and** feasible path
  - Feasible path → A path that **satisfies the throughput and latency target values**
  - Cheapest path → The path with the **minimum estimated price**

- **Path forwarding →** After a path update, an MNO pushes the new preferable path to its neighbors.

Numerical Results

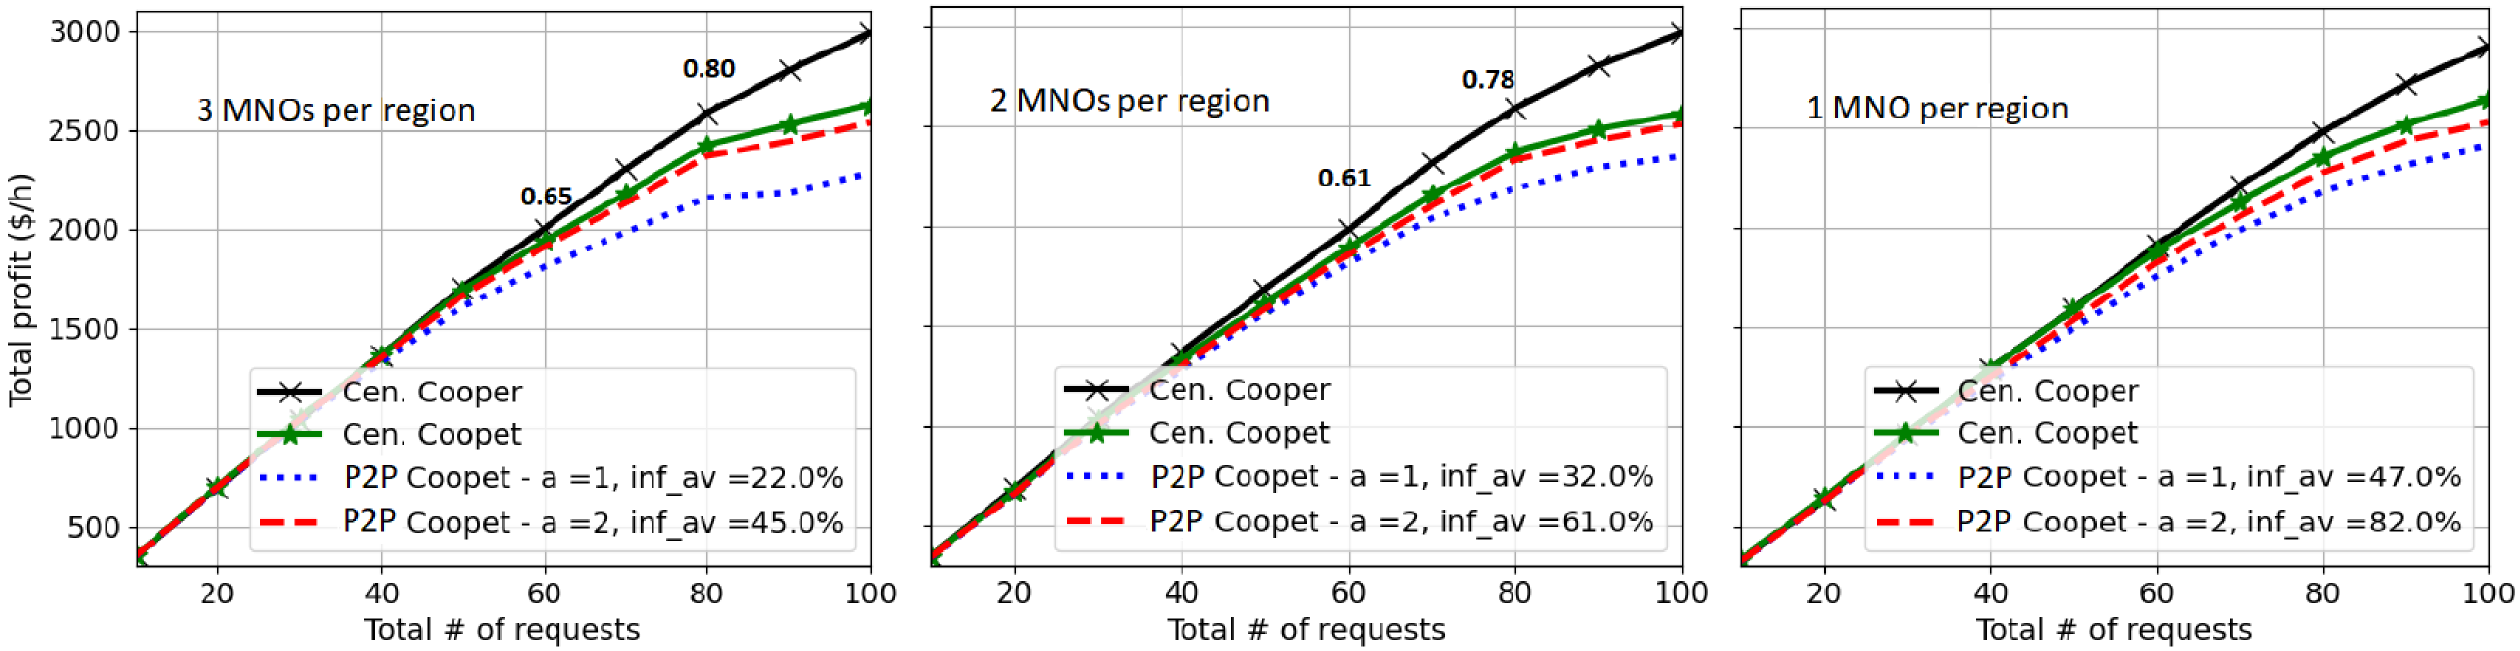

(a) 10 MNOs, 3 Geographic regions (b) 10 MNOs, 5 Geographic regions (c) 10 MNOs, 10 Geographic regions

Total Profits

- **Centralized.** For **utilization** < **~ 0.65**, the difference between *Cooperative* and *Coopetitive* is small (up to 5%).
  - increases significantly for utilization > **~ 0.8** – reaches **15%** at utilization 1.
  - Increases when the **# MNOs per geographic location** is greater than 1

- **Peer-to-Peer.** The performance of *Peer-to-Peer Coopetitive* improves with the size of the “table of preferable paths”
  - close to *Centralized Cooperative (upper bound) for* ***information availability > ~45 %.***
  - as the **# of MNOs per region** decreases, a **higher** value of α is needed for achieving an *adequate performance.*

Individual Profits

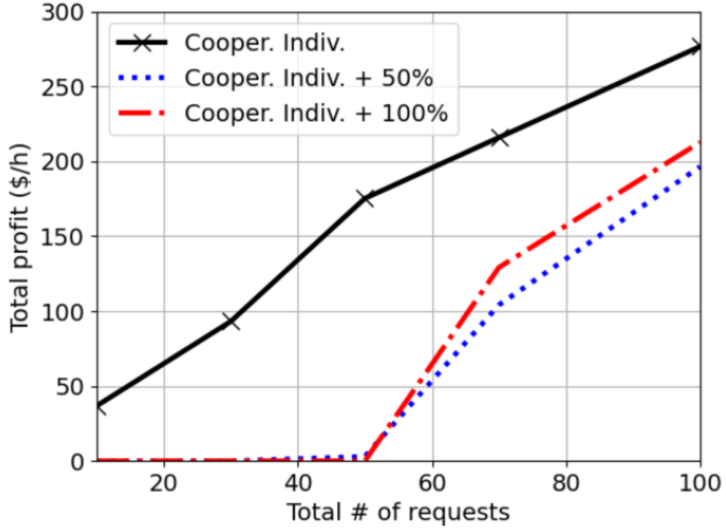

(a) 10 MNOs, 3 Geographic regions

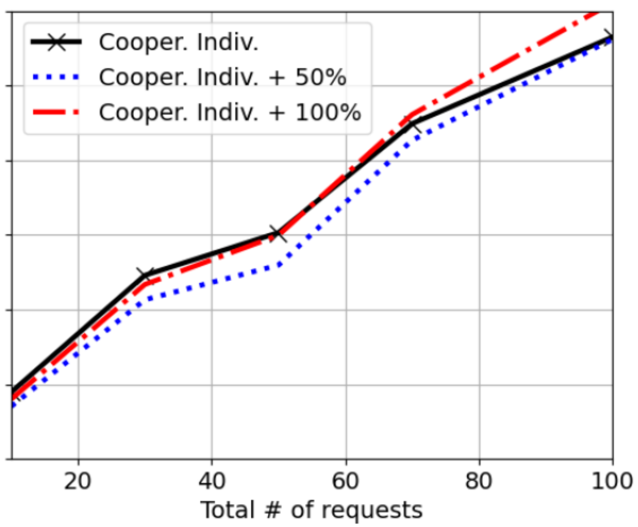

(c) 10 MNOs, 10 Geographic regions

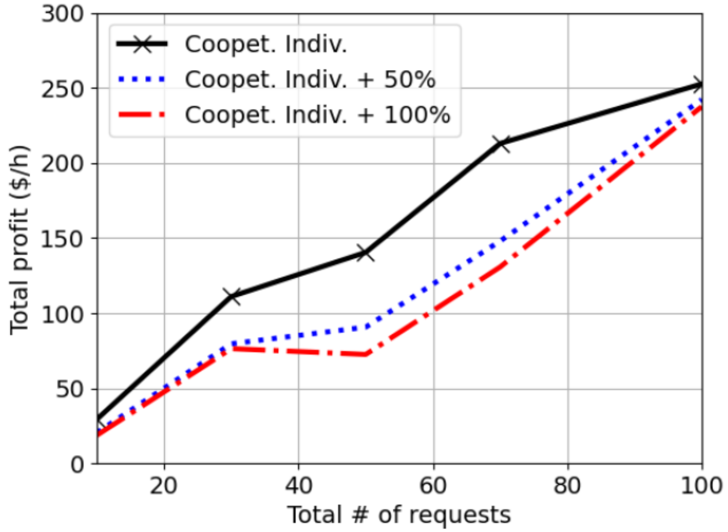

(a) 10 MNOs, 3 Geographic regions

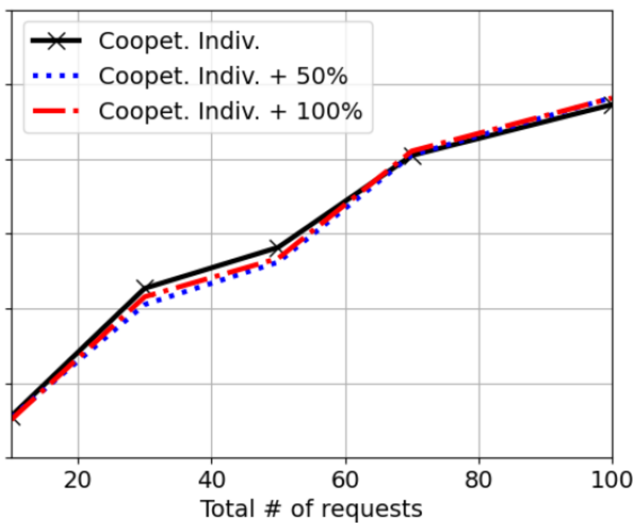

(c) 10 MNOs, 10 Geographic regions

Cooperative - Impact of untruthfulness.

- **More than 1** MNOs in each region
  - an untruthful MNO will always have **profit loss.** ✓
- **Only 1 MNO** in each region **AND high utilization**
  - an untruthful MNO **can** generate **higher profit** ✖
  - MNOs **do not have knowledge** about others' utilization ✓

The untruthful behavior is avoided

Coopetitive - Impact of strategic pricing.

- Similar observations to the cooperative mode.

Concluding Remarks and Future Directions

- The **profit of MNOs is maximized** when they comply with the policies and rules under all proposed approaches and modes.

- The MNOs' profit under the **peer-to-peer coopetitive** mode is comparable to those in the centralized coopetitive one, when applied in the appropriate network conditions.

- The **untruthful** or **strategic** behavior of MNOs is discouraged/avoided.

- **Future Work**
  - Extend the models to capture the provisioning of a *single network slice* over **multiple physical paths**.
  - Extend the *information sharing mechanism* to operate under **inaccurate information**.
  - Analyze the impact of *strategic* behavior of MNOs when **forwarding paths.**

Thank you for your attention!

Backup Slides

QoS Models

Throughput (packets/sec)

- **VT** τ: Depends on the allocated **bandwidth** b_τ and the **data packet** size K_t → $\boxed{b_\tau / K_t}$
- **VM** v: Depends on the VM **service rate** → $\boxed{\mu_v(c_v, f) = \sigma_f c_v}$
  - σ_f: packets/sec that VNF f can process over a unit of computational capacity

Latency (secs)

- **VT** τ: If deployed over physical path π, the **aggregate delay** along the path $\boxed{d_\tau(\pi) = \sum_{e_{ij}} D_{ij}}$
- **VM** v **:** Each VM is modeled as an **M/M/1 queueing system** with latency

$$\boxed{d_v(c_v, f) = \frac{1}{\mu_v(c_v, f) - \lambda_r}}$$