Keynote at the 23rd IEEE International Symposium on Network Computing and Applications (NCA 2025)
November 5-7th, 2025, Instituto Superior Técnico, Lisbon, Portugal

Evolution of Resource Allocation Algorithms and Protocols with Network Slicing

Stefano Secci
Professor
Cnam, Paris, France
stefano.secci@cnam.fr

Credit: Francesca Fossati, Sahar Hoteit, Stéphane Rovedakis, Mario Patetta, Zeidan Braik
https://roc.cnam.fr

Funded by:

Context: towards coordinated end to end resource allocation

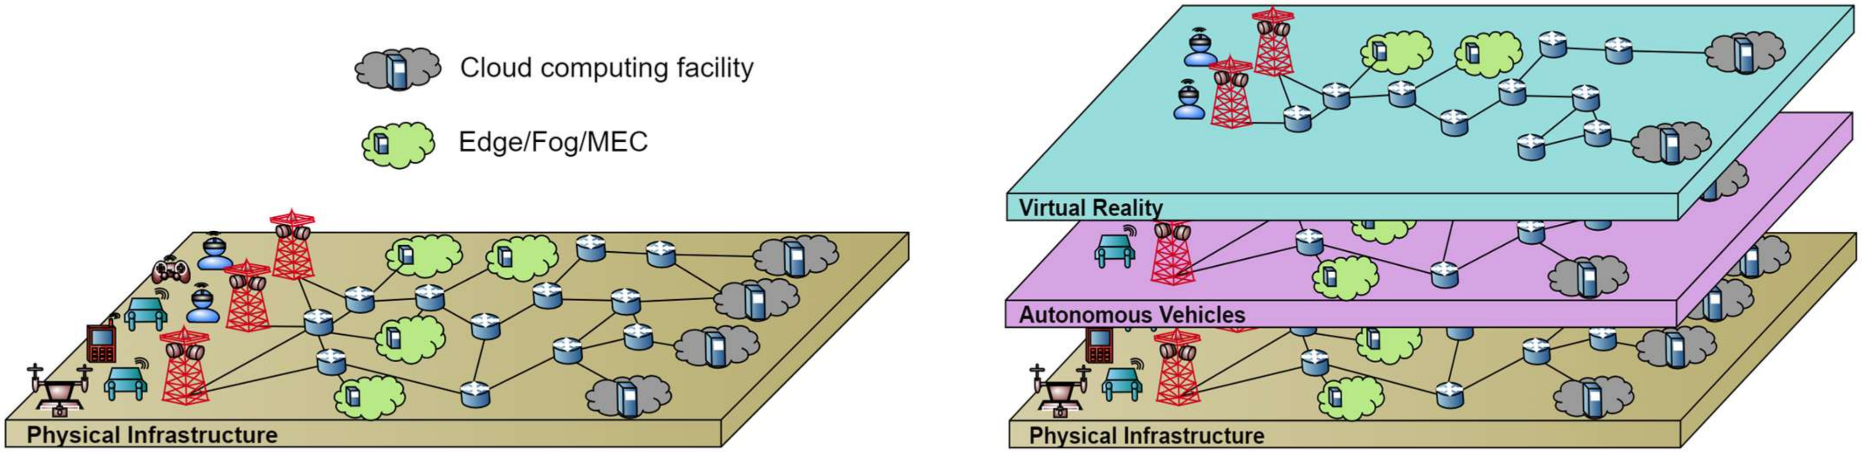

Credit: Mandar Datar, et al

- Services supported by modern networks imply:
  - Different resources (radio spectrum, bandwidth, computing power, ...)
  - Tailored allocation → network slicing
  - Shared physical infrastructure, managed by several operators

Outline

- Legacy Networks
- Verticalized Networks
- Single Resource Allocation
- Multi Resource Allocation
  - Centralized
  - Distributed
  - Congestion notification

Horizontal networks?

Legacy Networks

What we teach to students

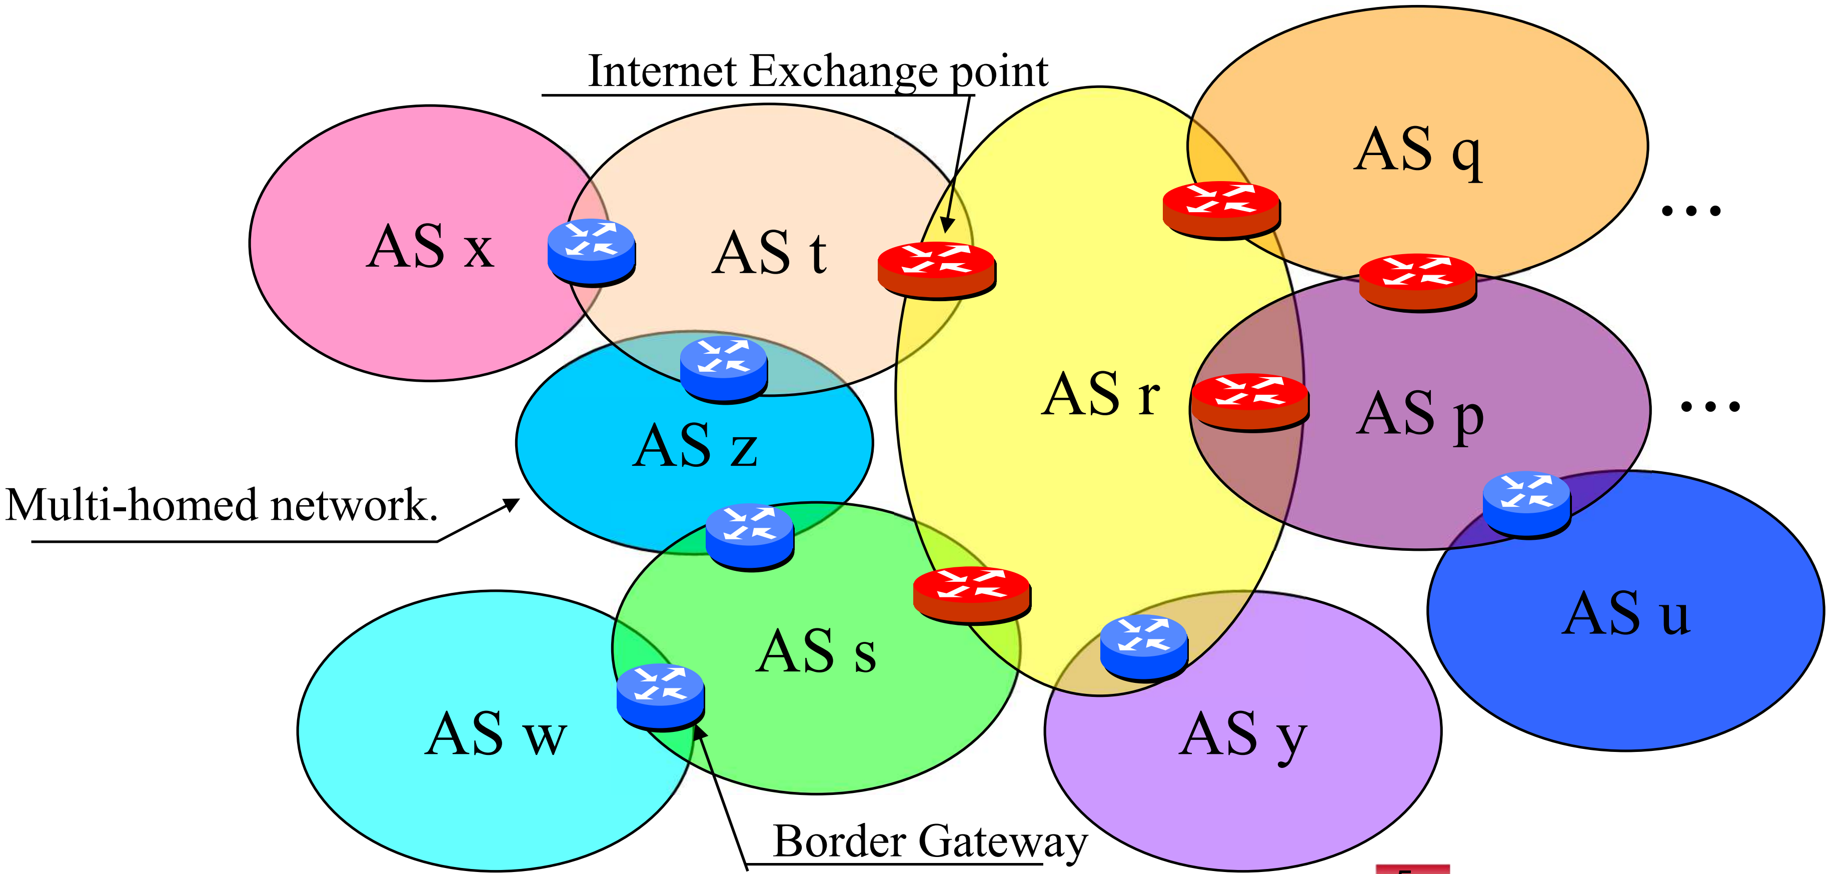

What we teach students

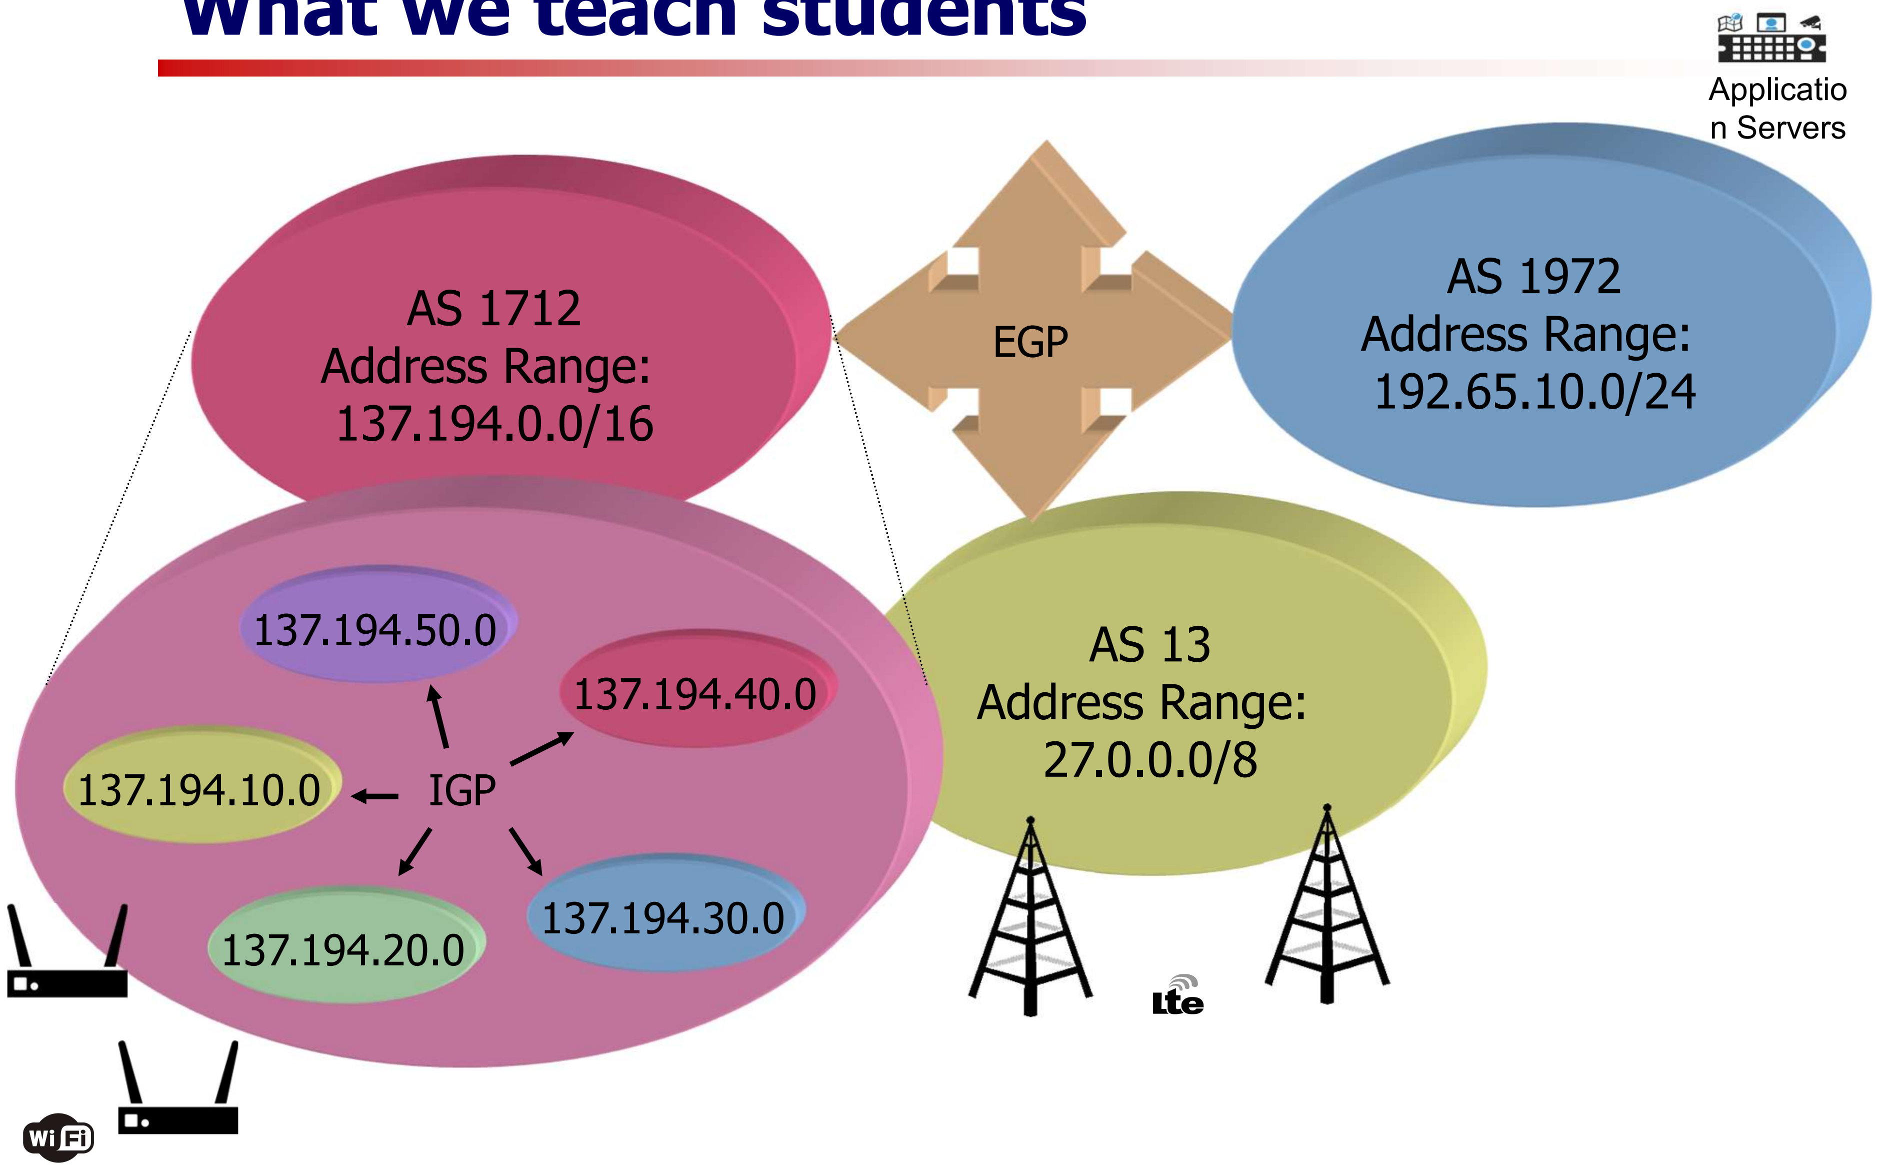

Cloud Networking evolution

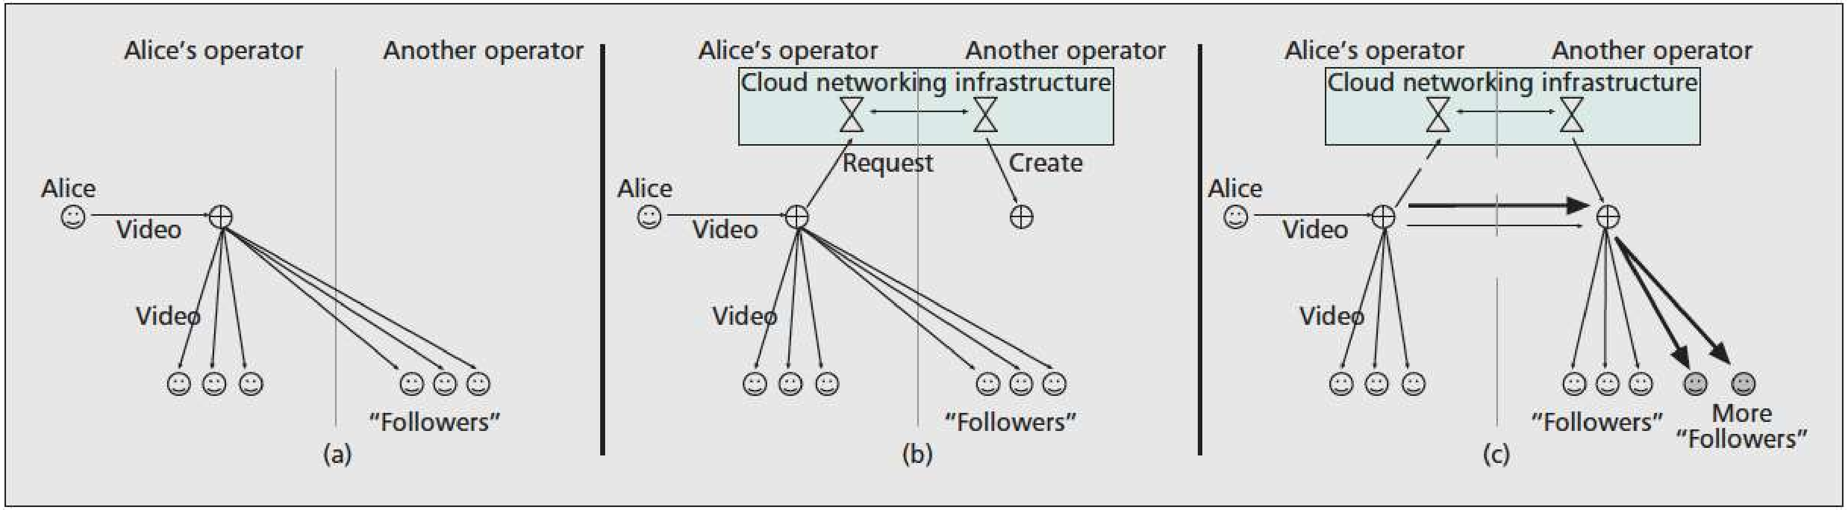

Source: B. Ahlgren et al, "Content, Connectivity, and Cloud: Ingredients for the Network of the Future", IEEE Communications Magazine, 2011

From localized cloud (a) to multi-site cloud (b) and multi-tenancy networks (c)

Cloud resources (shown in green) are spread much finer and deeper into the network, close to the actual point of usage.

The challenge of decentralization

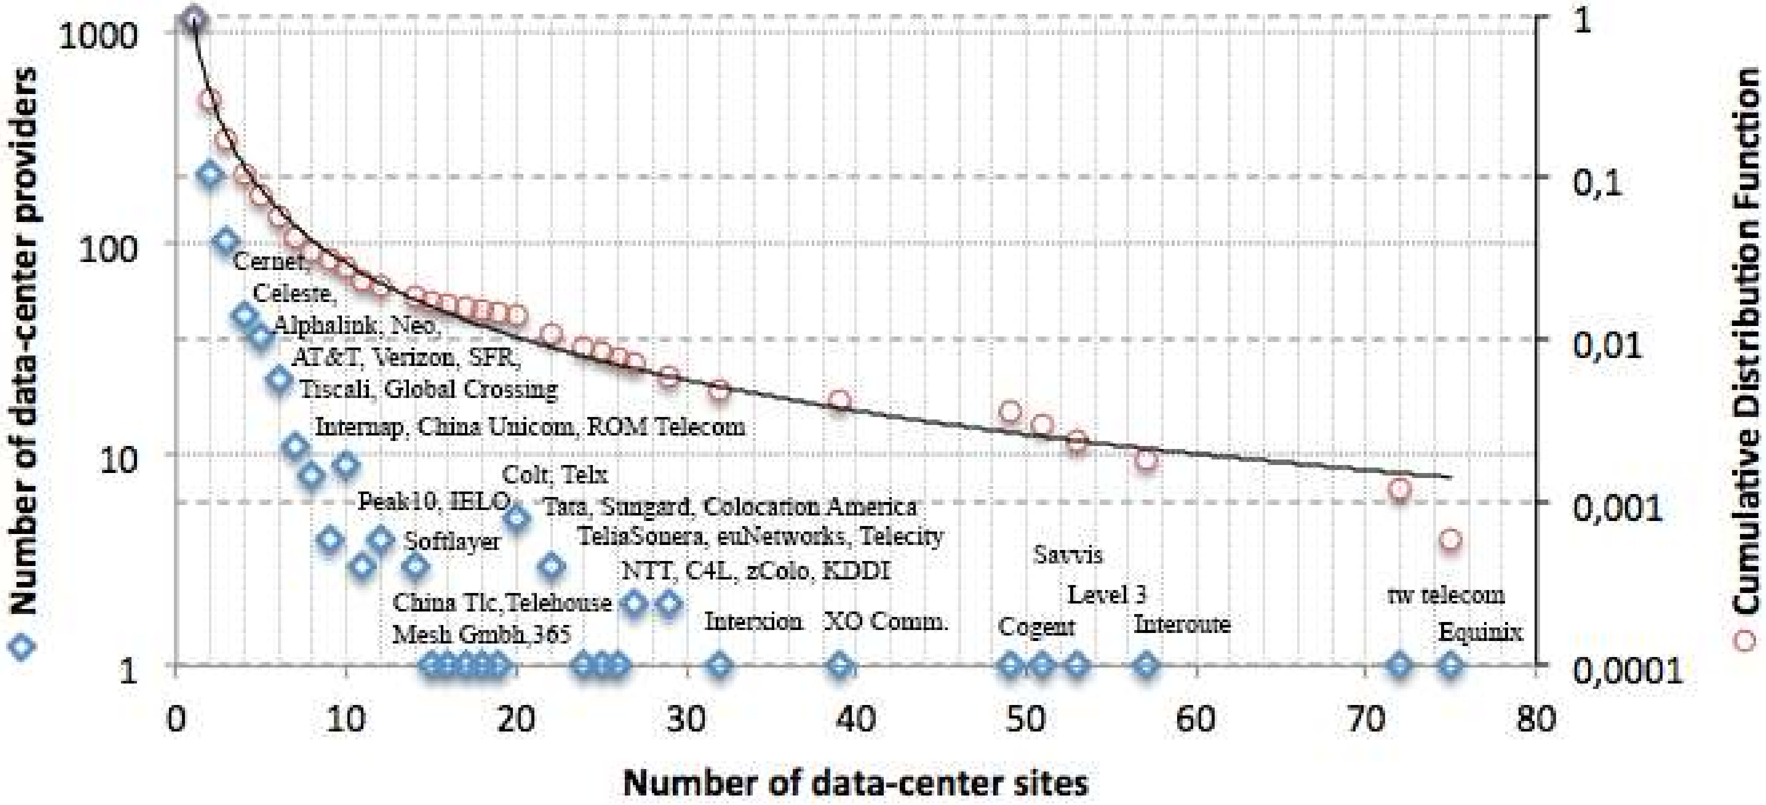

S. Secci, S. Murugesan, Cloud Networks: Enhancing Performance and Resiliency", IEEE Computer Magazine, Oct. 2014.

In-network / edge Clouds

a.k.a. CloudLets, Mobile Edge Computing

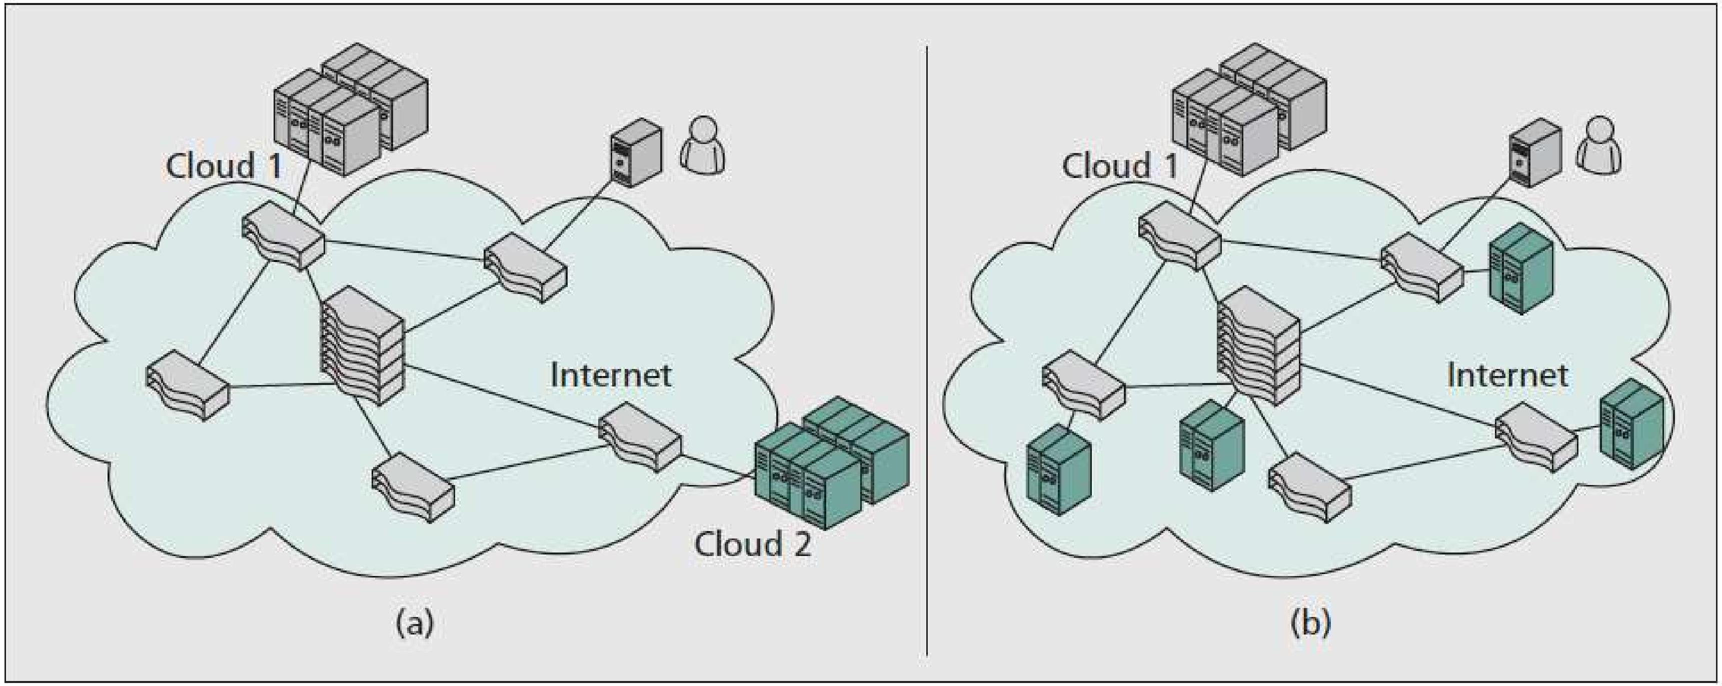

Source: B. Ahlgren et al, "Content, Connectivity, and Cloud: Ingredients for the Network of the Future", IEEE Communications Magazine, 2011

Slicing, shared infrastructure

Verticalized Networks

From 4G to 5G, then 6G

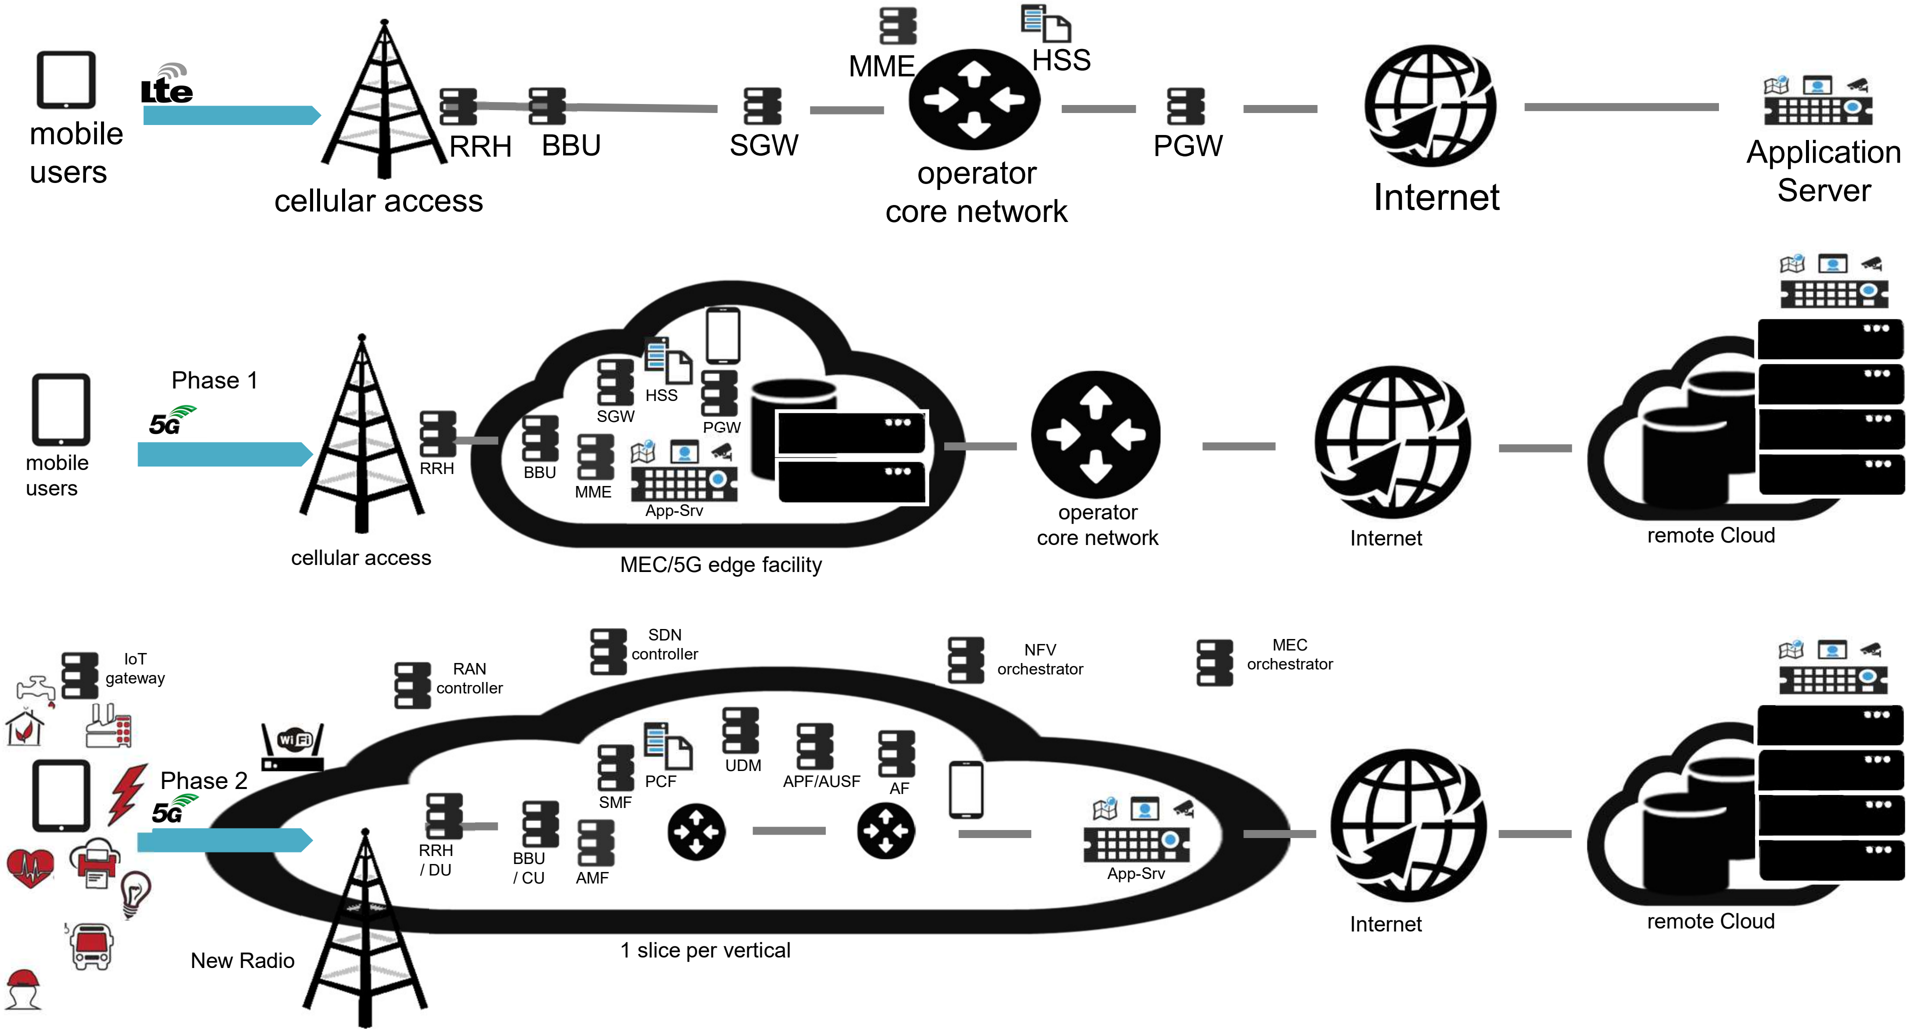

6G ? extreme low latency (<0.1 ms) & high reliability (7 9s), NTN, submarine

Network Verticalization

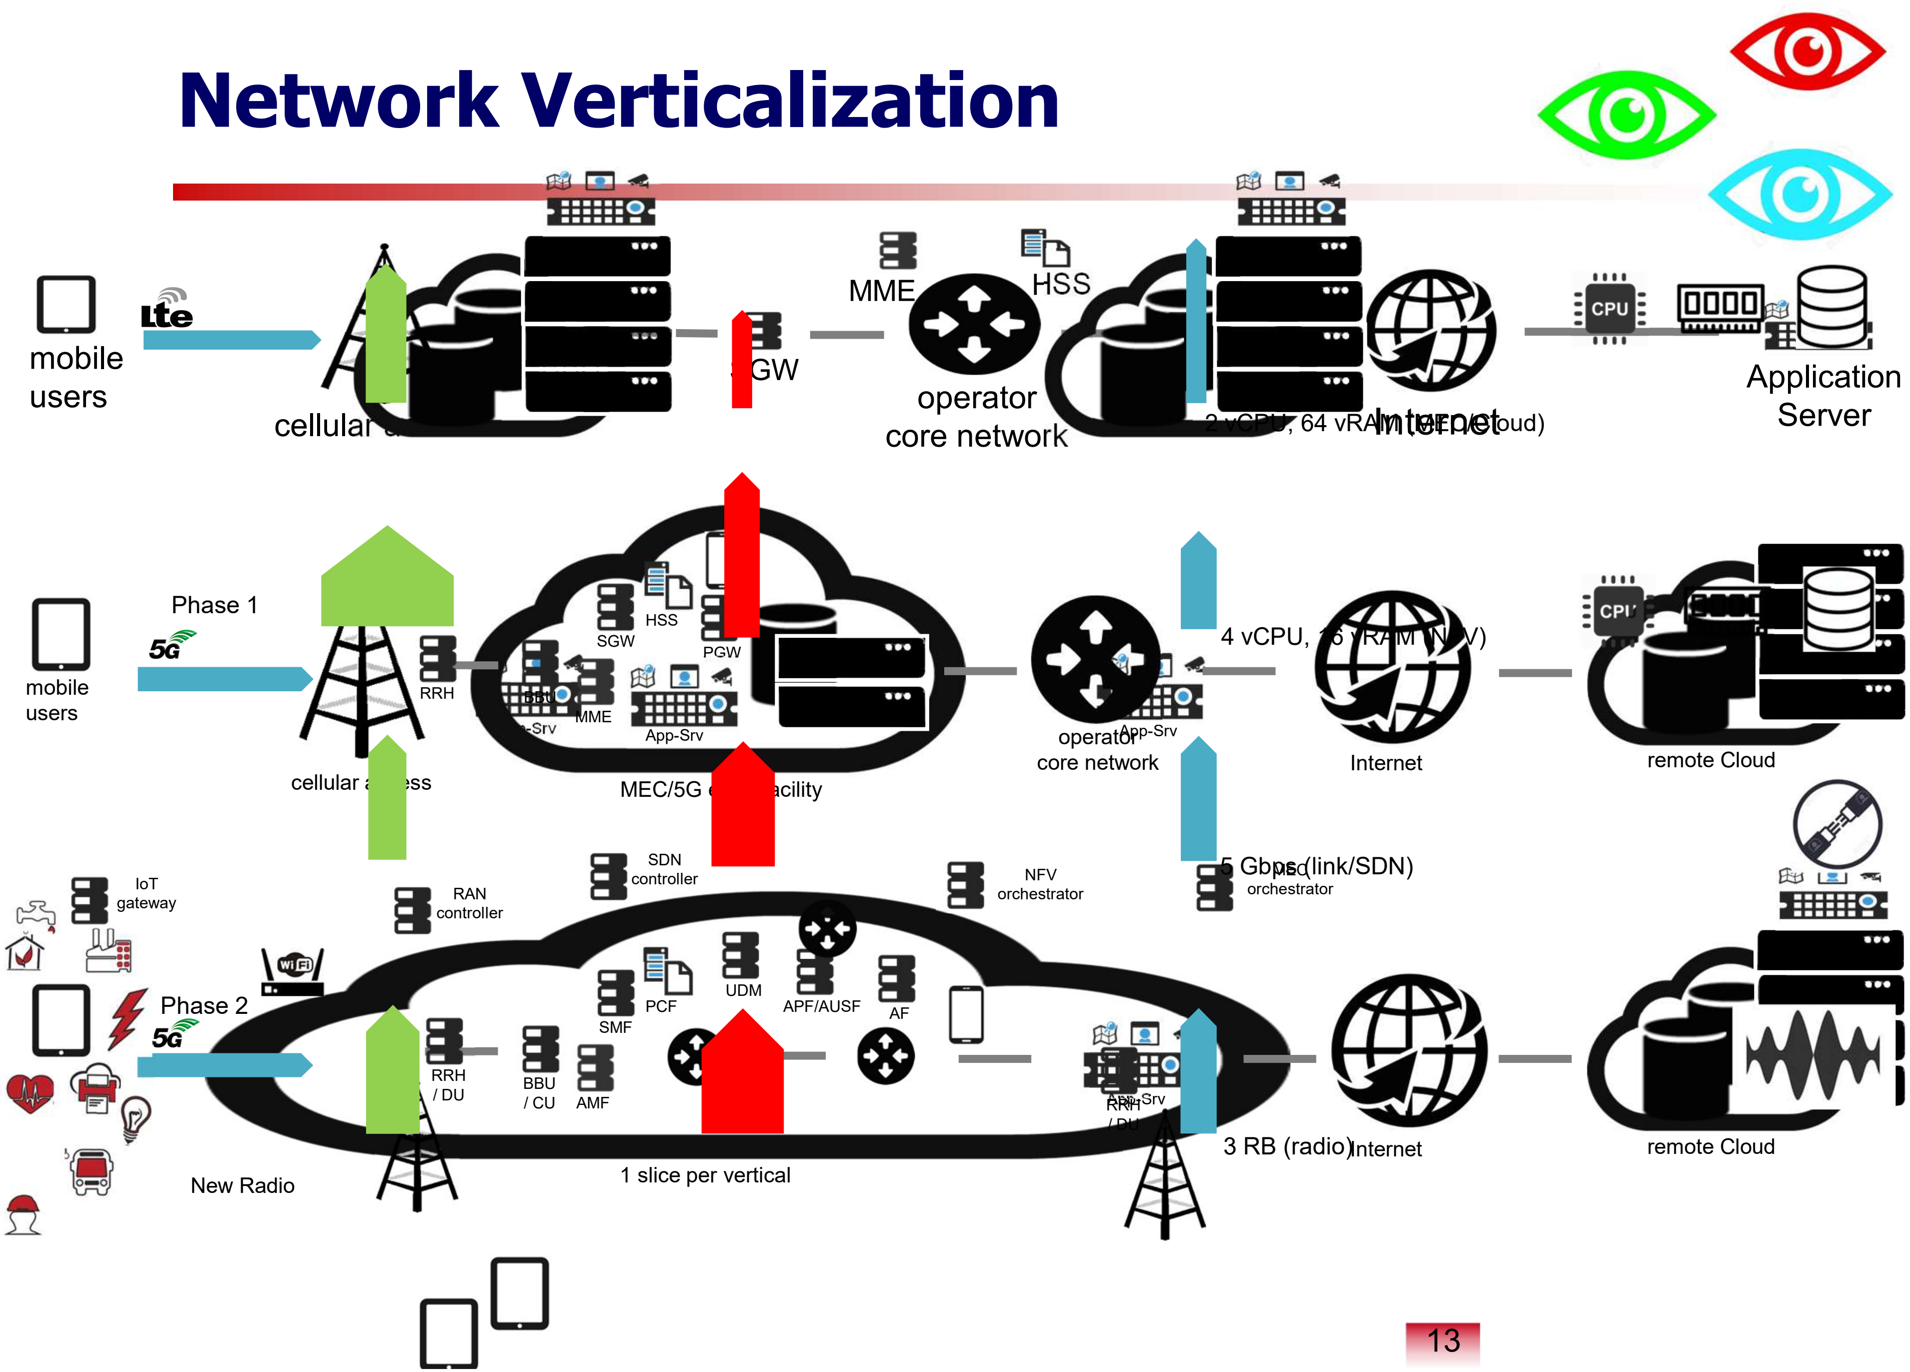

Infrastructure sharing

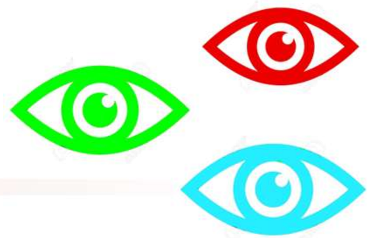

- An Infrastructure Network Provider (INP) has to allocate (multiple) resource to verticals (tenants)

- Competition among INPs
  - Verticals (tenants) may be able to choose one among multiple INPs, for each resource
  - The most performant/reliable INP attracts tenants
  - If an INP delivers service to two concurrent verticals, trust/fairness problem

    → tenants would ask the infrastructure to be auditable

- Monopolist or INP that cannot be dynamically changed
  - If for one resource there is no INP choice to take, the INP has to guarantee the resource is fairly shared

    → the infrastructure should be auditable,

    by tenant request or by regulation

Infrastructure sharing

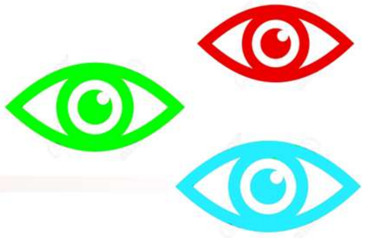

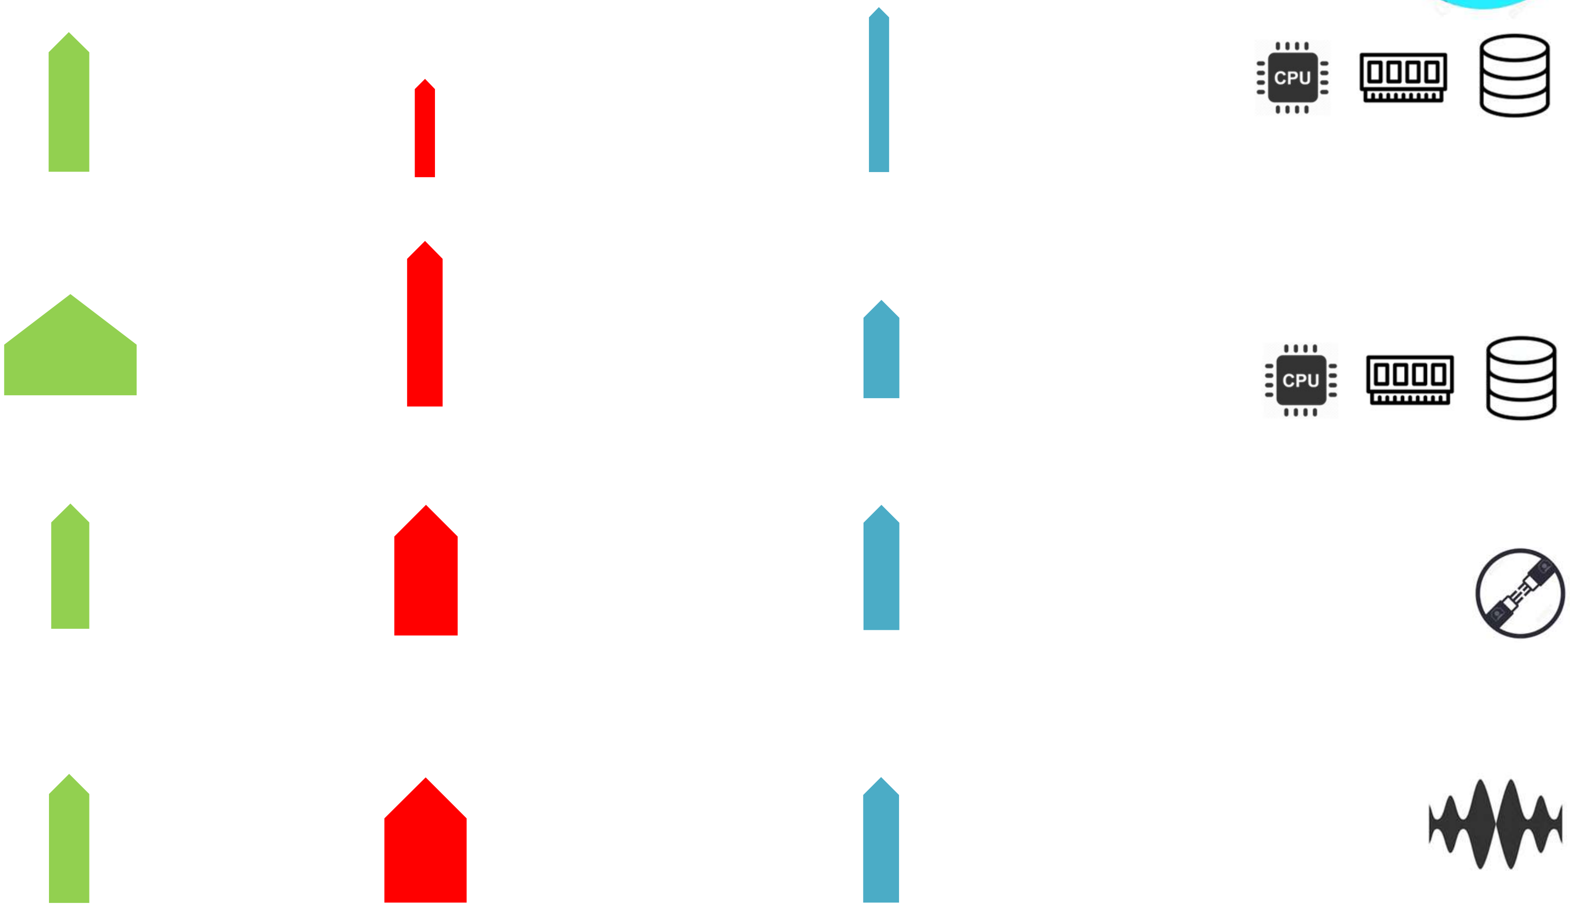

Francesca FOSSATI, Stefano MORETTI, Stefano SECCI, "A Mood Value for Fair Resource Allocations", Proc. of IFIP Networking 2017, June 12-16, 2017, Stockholm, Sweden.
Francesca FOSSATI, Sahar HOTEIT, Stefano MORETTI, Stefano SECCI, "Fair Resource Allocation in Systems with Complete Information Sharing", IEEE/ACM Transactions on Networking, Vol. 26, No. 6, pp: 2801-2814, Dec. 2018.
Francesca FOSSATI, Deep MEDHI, Stefano MORETTI, Stefano SECCI, "Error Estimate and Fairness in Resource Allocation with Inaccurate Information Sharing", IEEE Networking Letters, Vol. 1, No. 4, pp: 173-177, Dec. 2019.

Fairness, MMF, proportional allocation, game theoretic allocations

Single Resource Allocation

Resource allocation problem

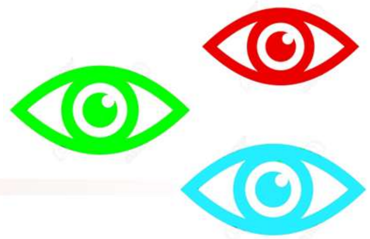

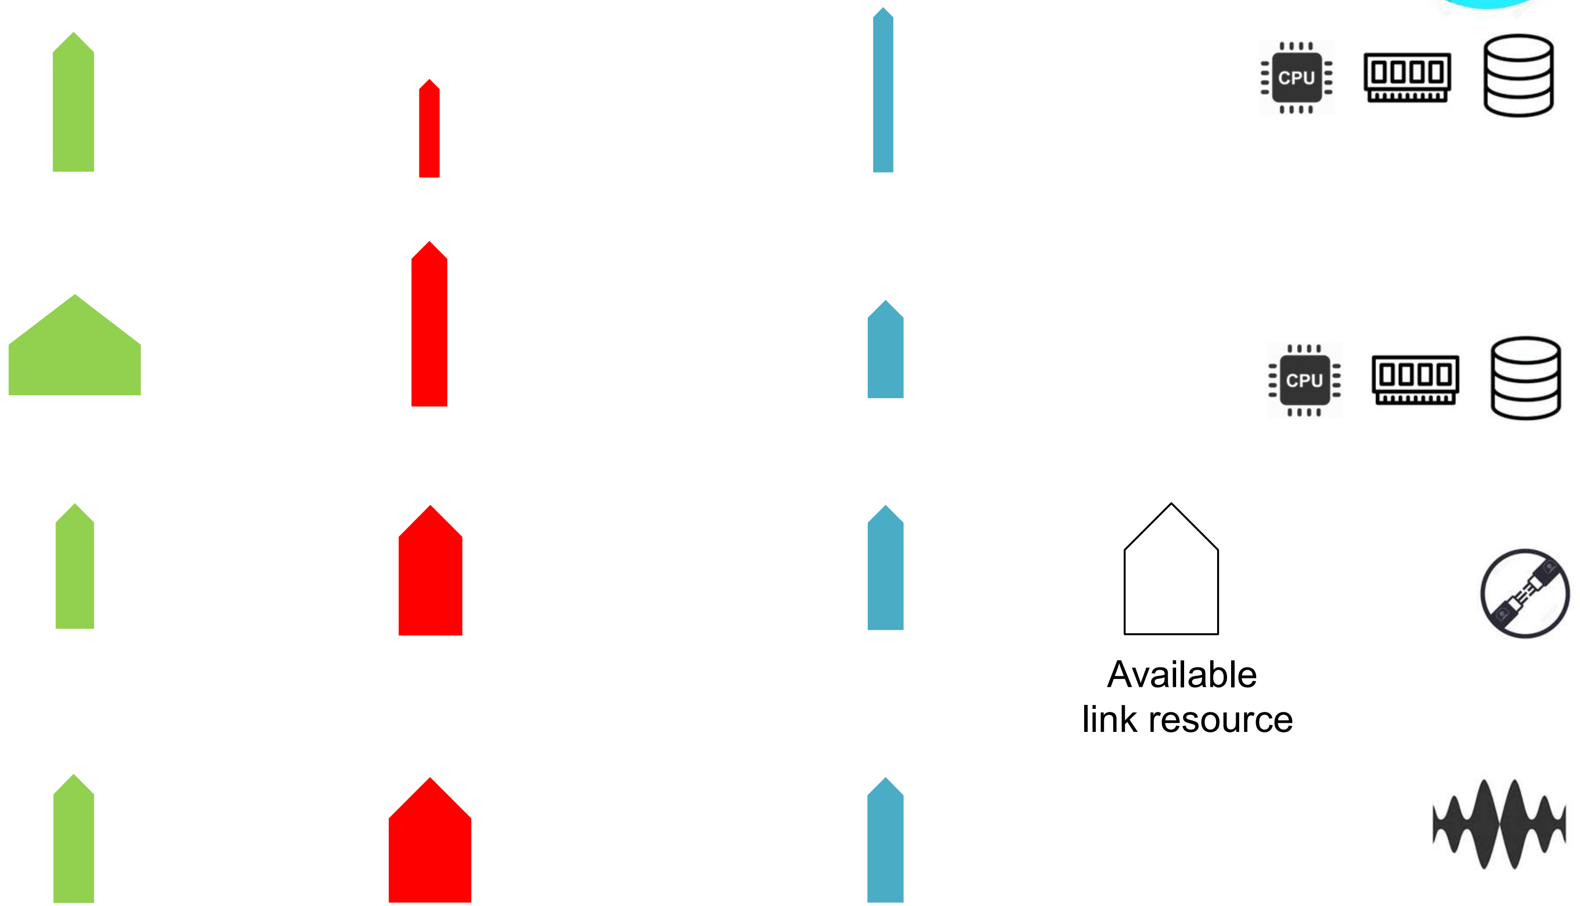

Single resource allocation problem
legacy systems

- Resource scarcity - How to allocate?

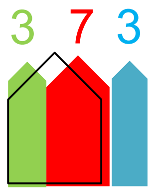

Available resource: 10

- Legacy systems : <u>you know only your demand</u>
  - Proportional fairness

    10 * 3 / (3+7+3) = 3/13 = 2.3

    10 * 7 / (3+7+3) = 7/13= 5.4

    10 * 3 / (3+7+3) = 2.3

  - max-min fairness (MMF)

    3

    (10-3-3) = 4

    3

Single resource allocation problem

complete information systems

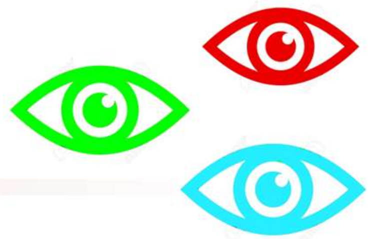

- Resource scarcity - How to allocate?

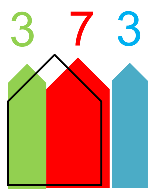

Available resource: 10

- Auditable systems :

  <u>you know your demand, others' demand, available resource</u>

  - If  gets 7, how much should it be satisfied from 0 to 10 ?
  - If  gets 4, how much should it be satisfied from 0 to 10 ?
  - If  gets 6, how much should it be satisfied from 0 to 10 ?

Single resource allocation problem
complete information systems

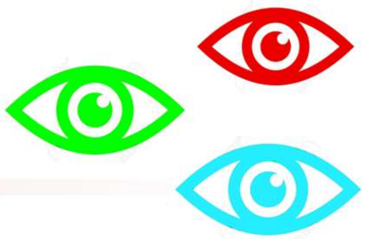

- Resource scarcity - How to allocate?

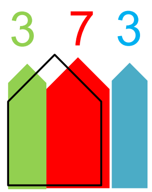

Available resource: 10

- Auditable systems :

  <u>you know your demand, others’ demand, available resource</u>

  - If  gets 7, how much should it be satisfied from 0 to 10 ? **10**

    because it’s the max it could get

  - If  gets 4, how much should it be satisfied from 0 to 10 ? **0**

    Because it’s the min it could get

  - If  gets 6, how much should it be satisfied from 0 to 10 ? **6.6**

    ????

    It’s a measure of your mood

New notion of satisfaction rate

- Player Satisfaction (PS) rate

$$PS_i = \frac{x_i - min_i}{max_i - min_i}$$

c_i *i th demand*

$$min_i = v(i)$$

$$max_i = v(N) - v(N \setminus i)$$

$$v(S) = \max\{E - \sum_{i \in N \setminus S} c_i, 0\}$$

What's left for S if all those not in S take all they ask

Available resource: 10

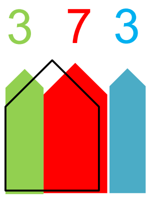

If  gets 6, then (6-4)/(7-4)=2/3

- If $\sum_{i=1}^{n} c_i = E$ then $PS_i = 1, \forall i \in N$

Francesca FOSSATI, Sahar HOTEIT, Stefano MORETTI, Stefano SECCI, "Fair Resource Allocation in Systems with Complete Information Sharing", IEEE/ACM Transactions on Networking, Vol. 26, No. 6, pp: 2801-2814, Dec. 2018.

New notion of fairness

- Players' fairness index
  - I.e., Jain's index using the PS rate instead of the demand fraction satisfaction

$$J_p = \left[\sum_{i=1}^{n}(PS_i)\right]^2 \Big/ n\sum_{i=1}^{n}(PS_i)^2$$

Francesca FOSSATI, Sahar HOTEIT, Stefano MORETTI, Stefano SECCI, "Fair Resource Allocation in Systems with Complete Information Sharing", IEEE/ACM Transactions on Networking, Vol. 26, No. 6, pp: 2801-2814, Dec. 2018.

New allocation rule

- The Mood Value
  - The allocations x_i such that all users have the same mood, i.e., the same Ps_i
  - That is, the ones that maximize the PS fairness (=1)

- What is the mood ?
  - A satisfaction rate equal for everybody

$$m = \frac{E - min}{max - min}$$

$$min = \sum_{i=1}^{n} min_i$$
$$max = \sum_{i=1}^{n} max_i$$

- What is the Mood Value then ?
  - An allocation that is higher than the minimum one by a quantity given by the product of the common mood times what's left to be fully satisfied

$$x_i^m = min_i + m(max_i - min_i)$$

Francesca FOSSATI, Sahar HOTEIT, Stefano MORETTI, Stefano SECCI, "Fair Resource Allocation in Systems with Complete Information Sharing", IEEE/ACM Transactions on Networking, Vol. 26, No. 6, pp: 2801-2814, Dec. 2018.

Francesca FOSSATI, Stefano MORETTI, Patrice PERNY, Stefano SECCI, "Multi-Resource Allocation for Network Slicing", *IEEE/ACM Transactions on Networking*, Vol. 28, No. 3, pp: 1311-1324, June 2020.

MURANES

Multi-Resource Allocation

Multi-resource allocation problem

Resource dependency

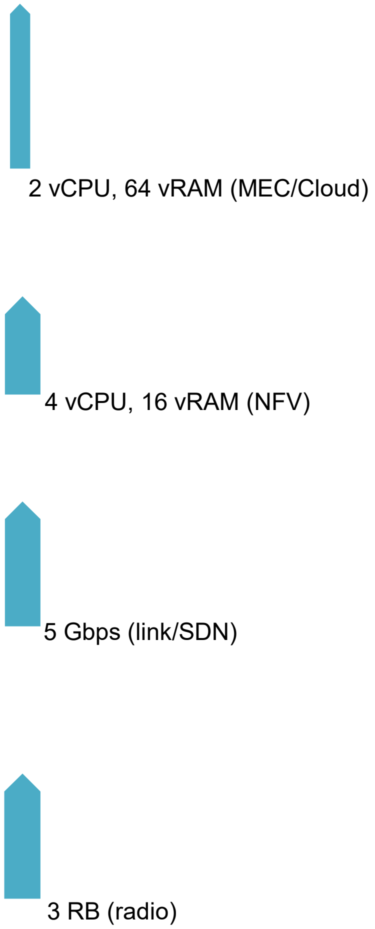

- If I get only 2 Gbps instead of 5 Gbps?
  - I no longer really need 4 vCPU and 16 vRAN for my VNFs
  - Because I'm ok with a lesser number of VNF instances
- If I get only 1 Gbps throughput at the radio access ?
  - I no longer need 5 Gbps virtual link
  - I no longer need that much vCPU and vRAM for VNFs and application servers

Need to take into account inter-resource relationship to avoid resource waste

Multi-resource vs single-resource allocation

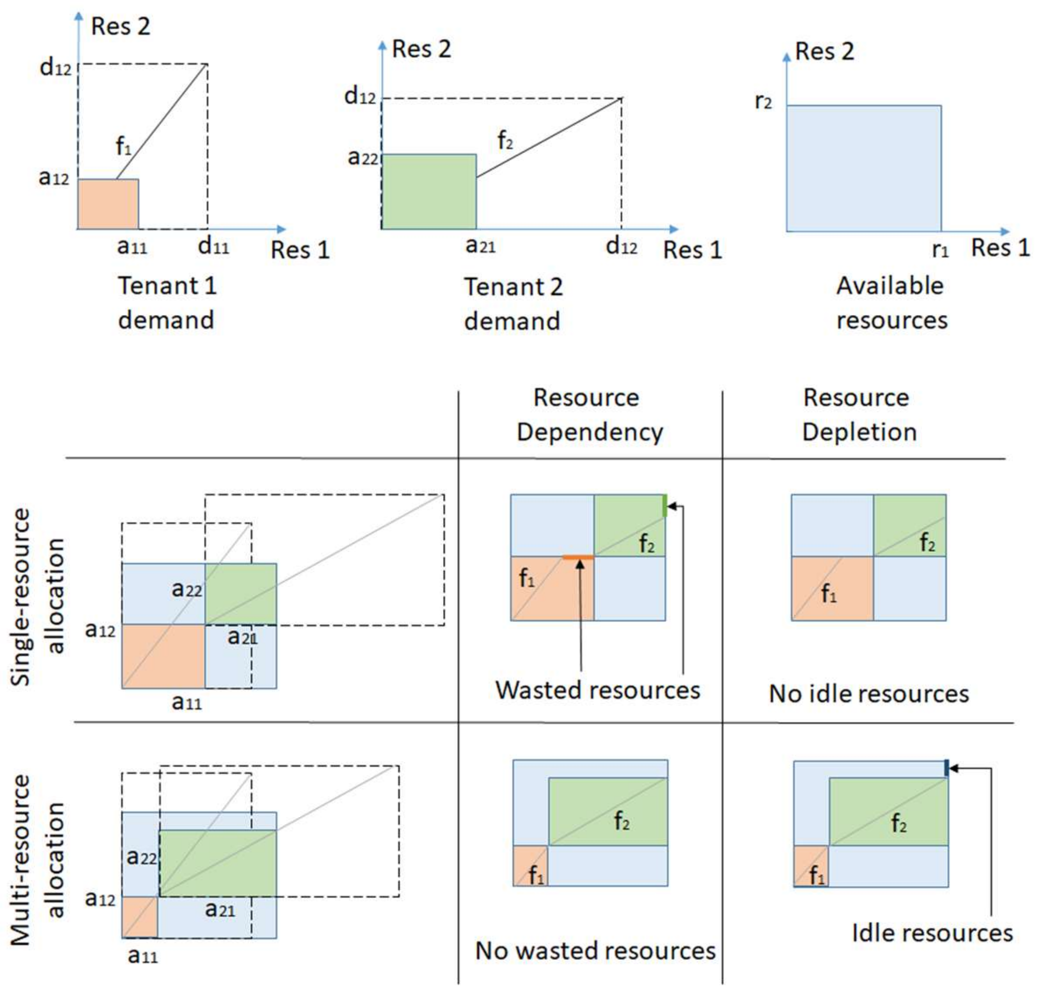

MURANES framework

Multi Resource Allocation for Network Slicing

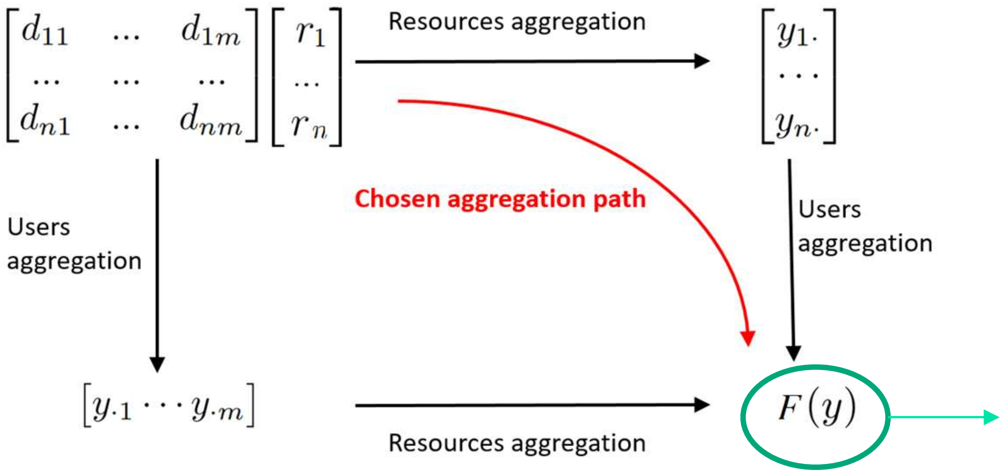

An Ordered Weighted Average (OWA) is:
a scalarizing function F: $\mathbb{R}^n \to \mathbb{R}$
using a weighting vector $w \in \mathbb{R}^n_+$
of the form $F(v_1, \dots v_n) = \sum_{j=1}^{n} w_j v_{(j)}$, where $v_{(j)}$
is the *j*-th smallest element of $(v_1, \dots v_n)$.

OWA used with a decreasing weighting vector, i.e., $w_i \geq w_{i+1}, \forall\, i < n$ allows to model a wide range of "fair" attitudes going from the egalitarianism to the utilitarianism.

MURANES framework

Multi Resource Allocation for Network Slicing

Definition 1. An *OWA* is a scalarizing function $F : \mathbb{R}^n \to \mathbb{R}$ parametrized by a weighting vector $w \in {\mathbb{R}_+}^n$ of the form $F(v_1, ...v_n) = \sum_{j=1}^{n} w_j v_{(j)}$, where $v_{(j)}$ is the j-th smallest element of $(v_1, ..., v_n)$.

$$\begin{array}{ll} \text{maximize} & OWA(v) \\ \text{subject to} & x \in \mathscr{F} \\ & 0 \leq x_i \leq 1, \forall i \in N \end{array}$$

$\mathscr{F}$ s.t. $\sum_{i \in N} a_{ij} \leq r_j, \forall j \in M$

Dominant share

$$ds_i = \max_j \{\frac{d_{ij}}{r_j}\}$$

where v can be equal to:

x	classical satisfaction
$ds \cdot x$	weighted classical satisfaction
ps	player satisfaction
$ds \cdot ps$	weighted classical satisfaction

		System		
		w=(1,0,⋯,0)	⋯	w=(1,1,⋯,1)
Individual	x	$\max\min x_i$	⋯	$\max \sum_{i=1}^{n} x_i$
	$ds \cdot x$	$\max\min ds_i x_i$	⋯	$\max \sum_{i=1}^{n} ds_i x_i$
	ps	$\max\min ps_i$	⋯	$\max \sum_{i=1}^{n} ps_i$
	$ds \cdot ps$	$\max\min ds_i ps_i$	⋯	$\max \sum_{i=1}^{n} ds_i ps_i$

Fossati F., Moretti S., Perny P., Secci S. *Multi-Resource Allocation for Network Slicing*. IEEE/ACM Transactions on Networking, Vol. 28, No. 3, pp: 1311-1324, June 2020.

MURANES framework: some results

Multi Resource Allocation for Network Slicing

Test of

- classical single-resource rules
- new multi-resource allocations.

100 allocation problems

- 3 resources (Memory, vCPU, link capacity),
- 10 slices.

Random generation of the slice demands using a subset of Amazon EC2 instances.

All the resources are congested. Different congestion levels (fraction of the global demand not allocated due to resource scarcity).

g-prop: x	g-mood: ps
g-drf: ds·x	gm-drf: ds·ps

<u>*Congestion level*</u> (cl_j) of a resource j : $cl_j = \sum_{i=1}^{n} d_{ij} \,/\, R_j$

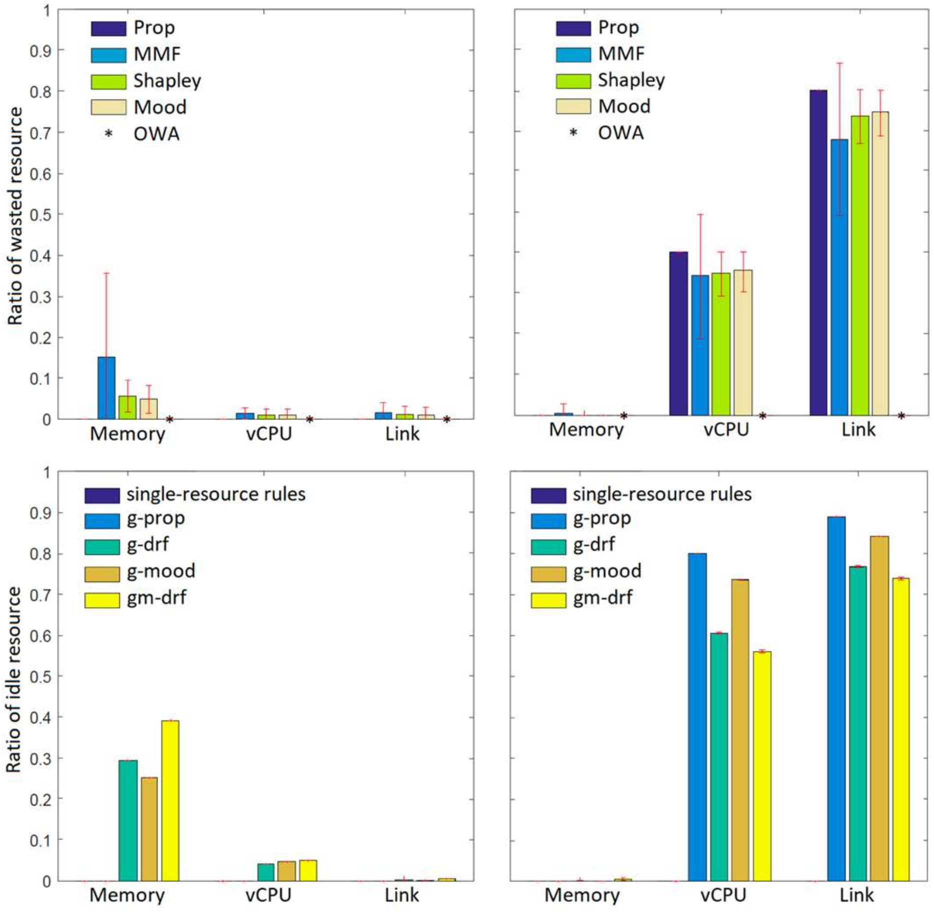

cl = (10,10,10)

cl = (10,2,1.1)

Fossati F., Moretti S., Perny P., Secci S. *Multi-Resource Allocation for Network Slicing*. IEEE/ACM Transactions on Networking, Vol. 28, No. 3, pp: 1311-1324, June 2020.

Going beyond linear dependencies

- Resource 1
- Resource 2

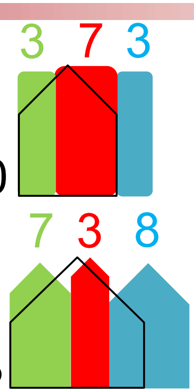

- Dependencies:

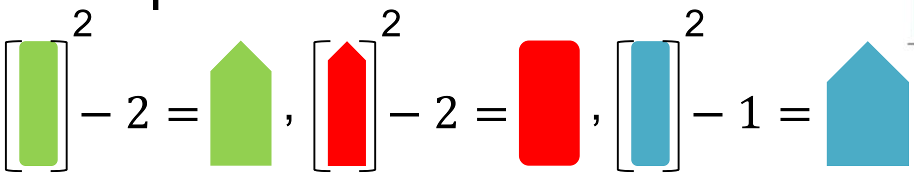

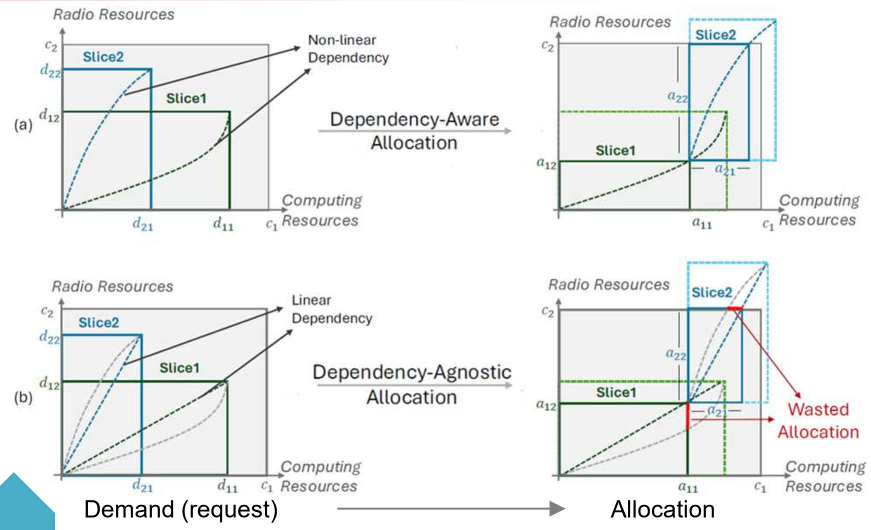

- DRF (assuming linear proportional dependencies):

Allocation: 2.4 ~~5.6~~ 3.7 ~~4.2~~ 1.24 1.8 2.1 ~~5.6~~ 3.4

Satisfaction: 0.8 ~~0.8~~ 0.53 ~~0.6~~ 0.177 0.6 0.7 ~~0.7~~ 0.426

- DDRF:

Allocation: 2.692 5.25 4.25 2.5 2.5 5.25

Satisfaction: 0.897 0.75 0.607 0.833 0.833 0.6562

Total Resource Capacity Partitioning

- *How to compute the actual satisfaction users can achieve accounting for resource dependencies, which may reduce the allocated satisfaction ?*
- Effective Satisfaction rate:
  - A measure for user's actual satisfaction on each resource, for a given allocation
  - $\mathrm{X} \in [0,1]^{\mathrm{N}\times\mathrm{M}}$ is the allocated satisfaction matrix, users $\mathcal{N} = [1,..,\mathrm{N}]$, resources $\mathcal{M} = [1,..,\mathrm{M}]$, $\mathcal{D}$ is the set of all dependencies between resources for each user:

$$X^{\text{effective}} = \arg \max_{e \in \mathcal{S}(X,\mathcal{D})} \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{M}} e_{ij}$$

$$\mathcal{S}(X,\mathcal{D}) = \left\{ e \in [0,1]^{N \times M} \mid 0 \leq e_{ij} \leq X_{ij}, \forall i,j, \ \ e \in \mathcal{D} \right\}$$

- How is available capacity partitioned?
  - Used Capacity: $\sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{M}} X_{ij}^{\text{effective}} d_{ij}$
  - Wasted Capacity: $\sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{M}} (X_{ij} - X_{ij}^{\text{effective}}) d_{ij}$
  - Idle Capacity: $\sum_{j\in\mathcal{M}} \left(c_j - \sum_{i\in\mathcal{N}} X_{ij} d_{ij}\right)$

DDRF framework

Dependency-aware Dominant Resource Fairness

$$\max \sum_{i=1}^{N}\sum_{j=1}^{M} x_{ij}$$

(Maximize overall satisfaction)

$$\text{s.t.} \quad \sum_{i=1}^{N} d_{ij}x_{ij} \le C_j, \qquad \forall j \in \mathcal{M}$$

$$X \in \mathcal{D},$$

(Satisfaction respect the dependencies)

$$\hat{y}_{ij}\hat{y}_{kj}\hat{\mu}_{ij}\hat{x}_{ij} = \hat{y}_{ij}\hat{y}_{kj}\hat{\mu}_{kj}\hat{x}_{kj}, \quad \forall i \neq k \in \mathcal{N},\ j \in \mathcal{M}$$

(Satisfaction of the users' congested resource with dominant share equalizes their share)

$$0 \le x_{ij} \le 1, \qquad \forall i \in \mathcal{N},\ j \in \mathcal{M}$$

Advantages:

- Fair treatment of low-demanding (weak) tenants like MMF in single-resource.
- Considers the generic dependencies in the system.
- Anchor the allocation on the system's most limiting (bottlenecks) resources.
- Removes allocation waste, improves effective satisfaction.
- Outperforms DRF satisfaction rate when weak tenants exists in the linear dependency case.

Two questions left...

- Who is going to compute the allocation, given the extreme latency requirements ?
- How to allow line-rate adaptations ?

Who is going to compute the allocation?

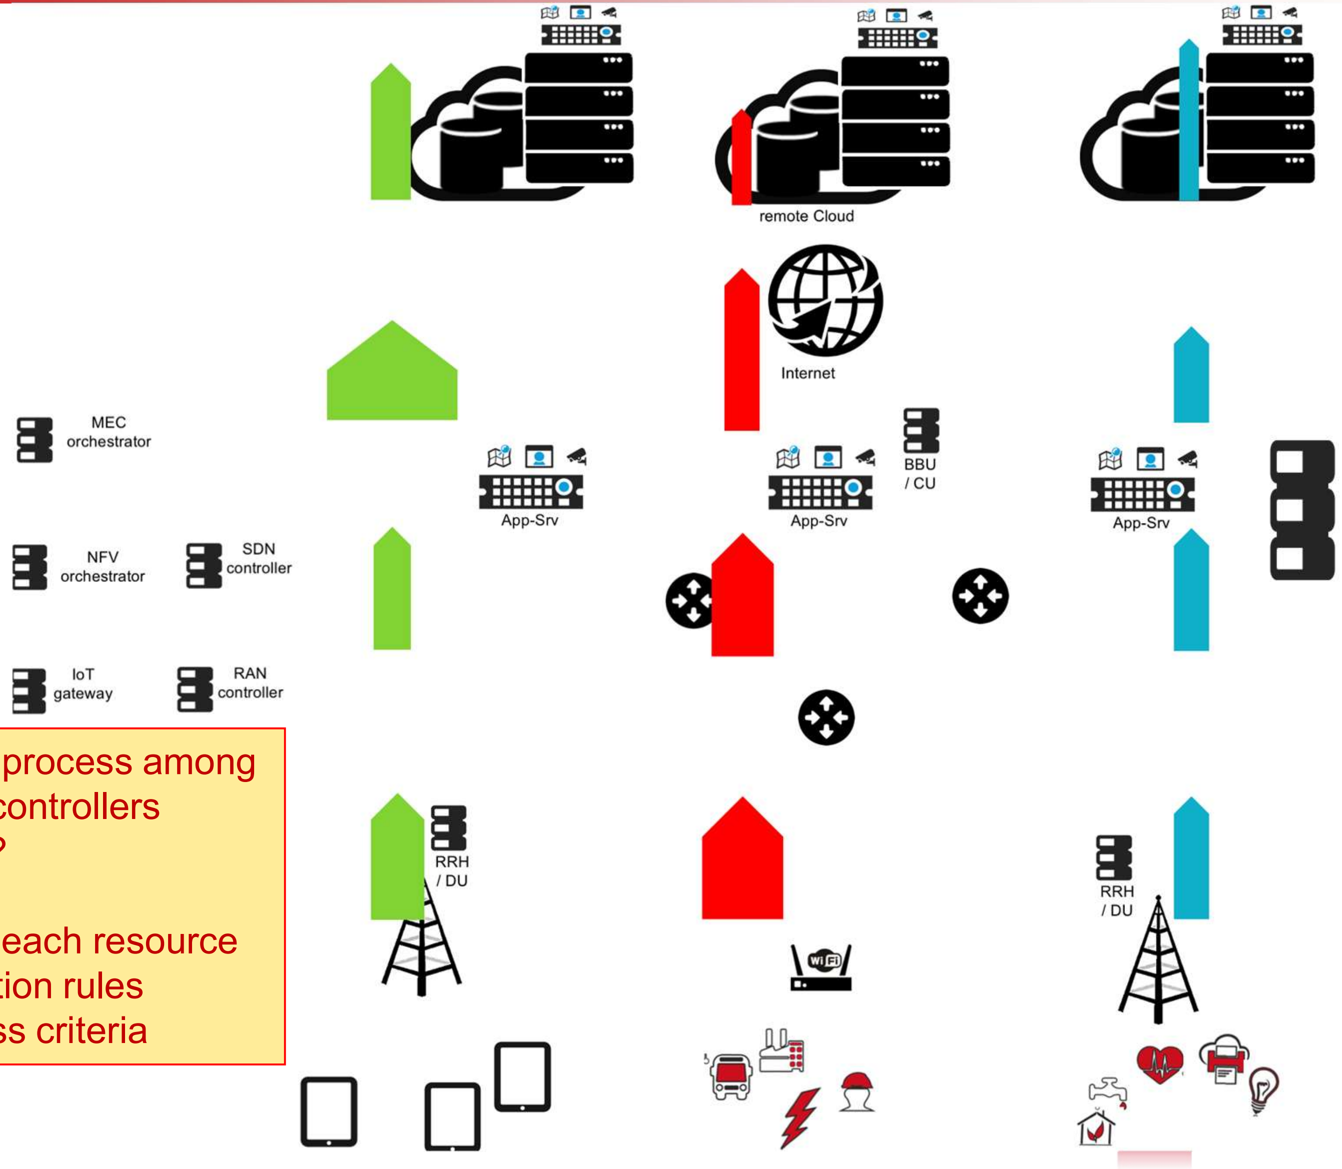

A super-orchestrator
?

Or a distributed process among resource controllers
?

Advantage – for each resource
- different allocation rules
- different fairness criteria

Francesca FOSSATI, Stefano MORETTI, Stéphane ROVEDAKIS, Stefano SECCI, "Decentralization of 5G slice resource allocation", Proc. of 2020 IFIP/IEEE Int. Symposium on Network Operations and Management Systems (NOMS), Apr. 20-24, 2020, Budapest, Hungary.
Francesca FOSSATI, Stephane ROVEDAKIS, Stefano SECCI, "Distributed algorithms for multi-resource allocation ",*IEEE Transactions on Parallel and Distributed Computing*, Vol. 33, No. 10, pp:2524-2539, Oct. 2022.

Distributed Multi-Resource Allocation

MAESTRO_5G HEIDIS

Possible distributed MRA approaches

MRA: multi-resource allocation

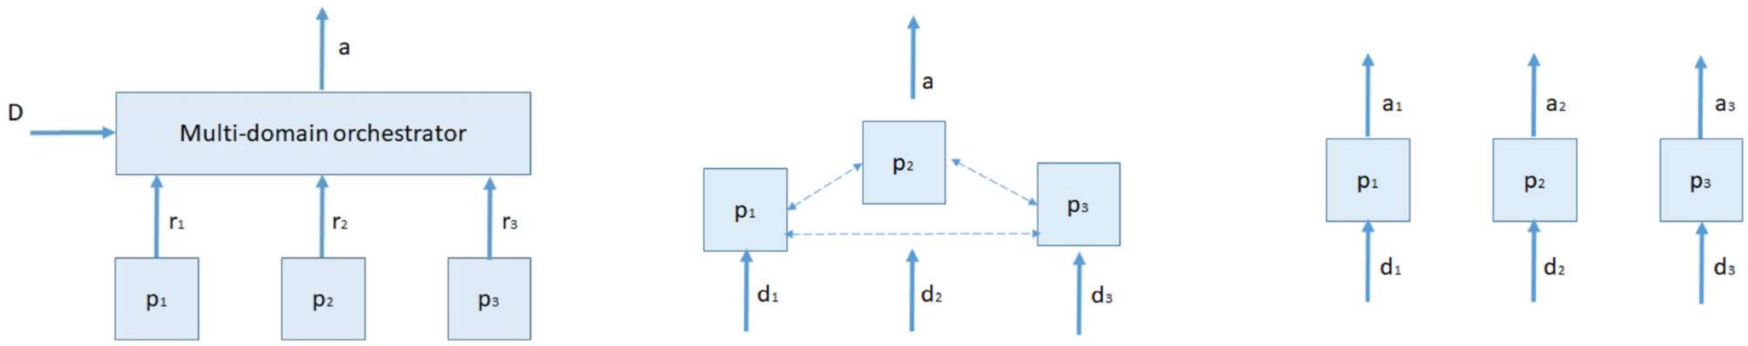

(a) Centralized algorithm (b) Distributed algorithm (c) Decentralized algorithm

Fig. 2: Diagrams of multi-resource allocation approaches for an example with three providers p_i providing one resource r_i, with $i = 1, 2, 3$. D is a $n \times 3$ demand matrix, where n is the tenant number and d_i, with $i = 1, 2, 3$ is a demand vector for the resource i, a is the $n \times 3$ allocation matrix and a_i, with $i = 1, 2, 3$ is an allocation vector for the resource i.

Three proposed distributed MRA schemes (2)

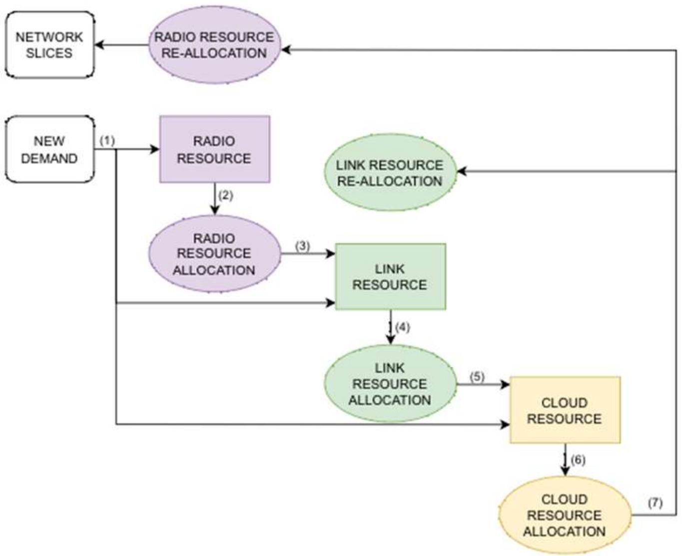

(a) Cascading Resource Allocation (CRA)

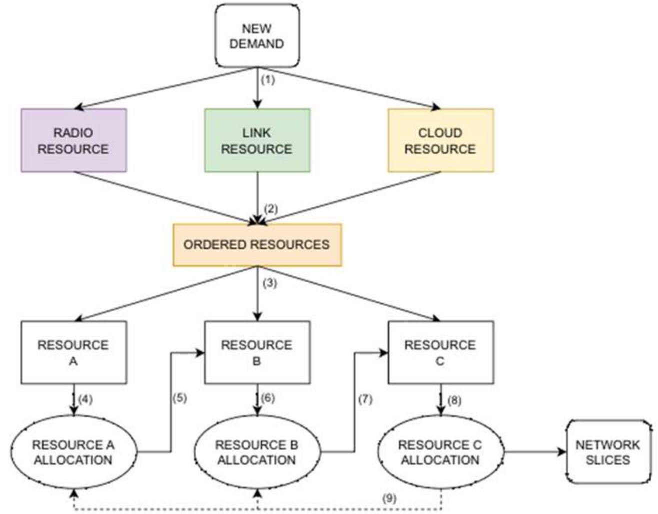

(b) Ordered Cascading Resource Allocation (OCRA)

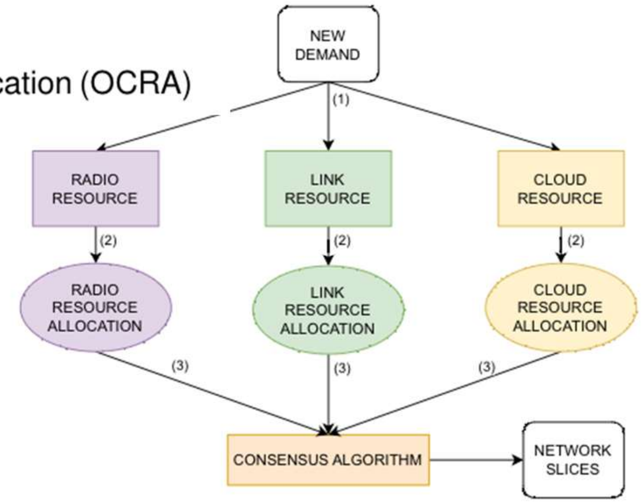

(c) Parallel Resource Allocation (PRA)

Delay-msg complexity tradeoffs

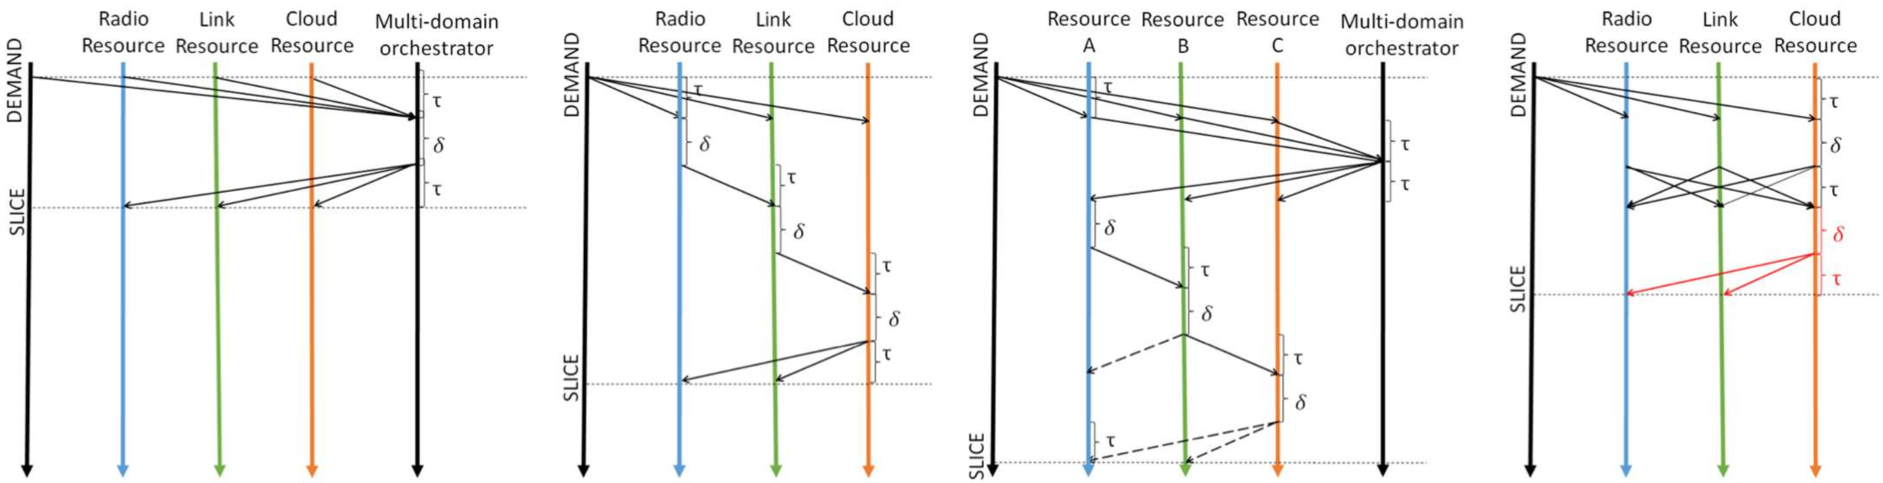

(a) Centralized (b) CRA (c) OCRA (d) PRA

Fig. 6: Involved signaling for the centralized algorithm and the proposed distributed ones, for the radio-link-cloud scenario, as a function of time, under the hypothesis of upper bound on transfer times (τ) and equal allocation computing times (δ). The dashed arrows indicate not necessary steps, and the red arrows correspond to extra steps of the 2-phase consensus algorithm.

Algorithm	Best case	Worst case	Message complexity
Centralized	$2\tau+\delta$	$2\tau+\delta$	$2p+1$
CRA	$(p+1)\tau+\delta$	$(p+1)\tau+p\delta$	$3p-2$
OCRA	$(p+2)\tau+\delta$	$(p+3)\tau+p\delta$	$[4p-1, \frac{p^2+7p}{2}-1]$
PRA-1	$2\tau+\delta$	$2\tau+\delta$	p^2
PRA-2	$3\tau+2\delta$	$3\tau+2\delta$	p^2+p-1

TABLE I: Delay budget and message complexity - General case with p resource providers.

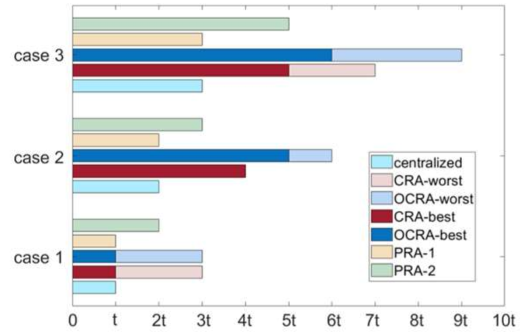

Fig. 7: Comparison of delay budgets with $p=3$. Case 1: $\tau \ll \delta$, $t=\delta$. Case 2: $\delta \ll \tau$, $t=\tau$. Case 3: $\tau=\delta=t$.

Summary

Algorithm	Advantages	Disadvantages
Centr.	Low delay budget	Multi-domain orchestrator High confidentiality disclosure
CRA	No multi-domain orcherstrator	Re-allocation
OCRA	Rarely re-allocation	High delay budget Multi-domain orchestrator
PRA-1	No multi-domain orchestrator Low delay budget Independent radio allocation	Pareto efficient solution not guaranteed High message complexity
PRA-2	No multi-domain orchestrator Low delay budget Independent radio allocation	High message complexity Low confidentiality disclosure

TABLE II: Pros vs cons of studied algorithms.

Mario Patetta, Stefano Secci, Sami Taktak, Multi-Resource Explicit Congestion Notification, submitted to IEEE Transactions on Networking.
Mario Patetta, Sami Taktak, Stefano Secci “Congestion Notification for Multi-Resource Allocation Control”, draft-ietf-iccrg-mrecn-00 (work in progress), presented at IETF 124, Montreal, Canada, Nov 4, 2025.

Multi-Resource Explicit Congestion Notification

Toward real-time signaling

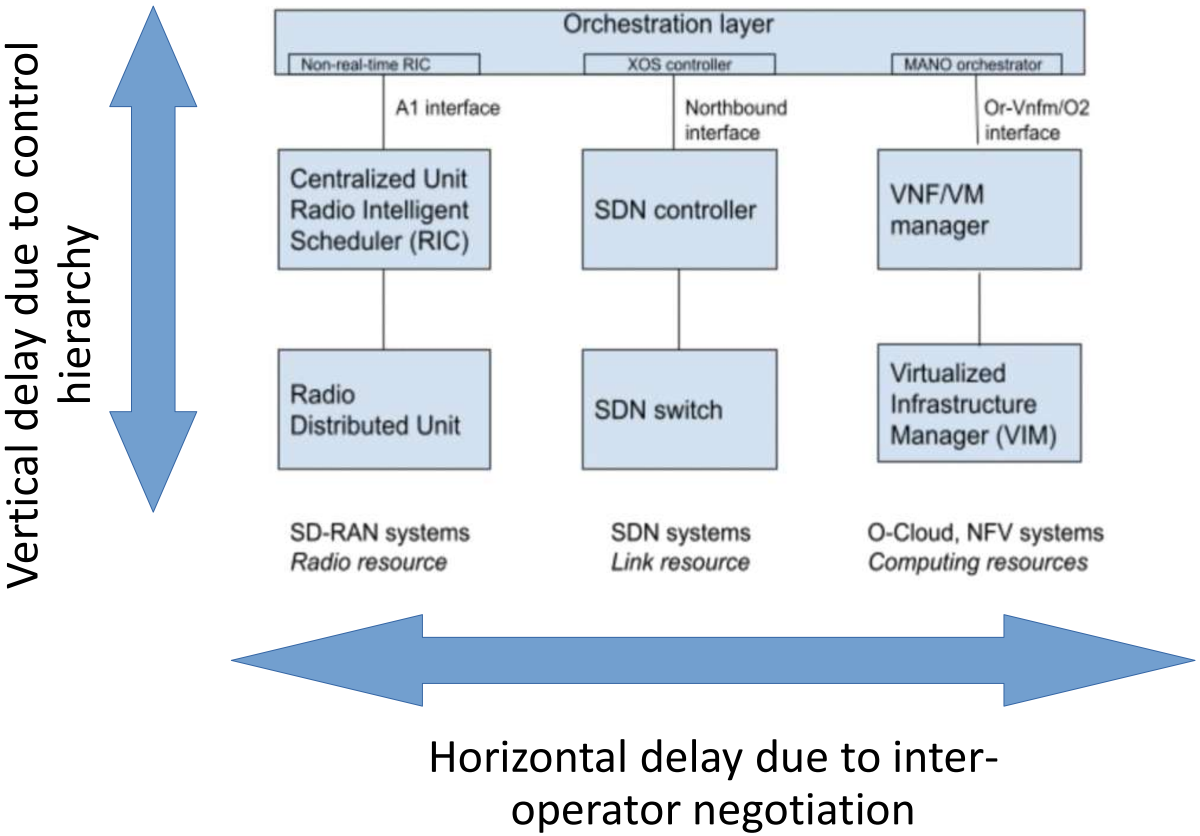

Multi-Resource Explicit Congestion Notification (MR-ECN)

- Notification injected into affected packets (both directions)
- Recipients → equipment traversed by the affected 'slice' (instead of end devices)
- Granularity → slice (instead of flow)

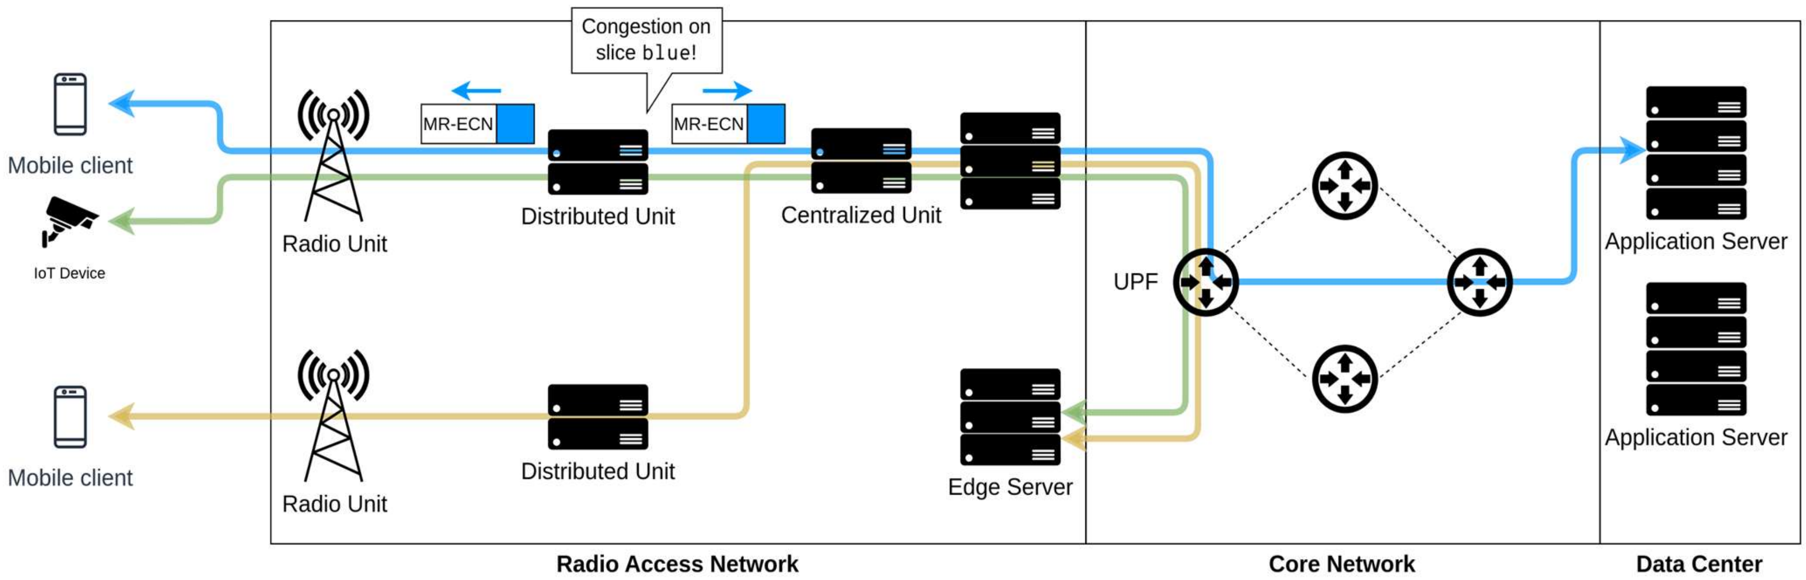

ECN → MR-ECN

from flow congestion to resource congestion notification

	ECN	MR-ECN
Purpose	Bypass pkt loss TCP's congestion avoidance	Bypass controller hierarchies in multi-resource allocation adjustment
Resource	Link capacity	Any
Granularity	Flow or MPLS trunk (RCF 5129)	Slice resource sw/hw stack
Recipients	End devices	data-plane node (IP router, vswitch, gNB)

0	1	2	3	4	5	6	7
MRCE	Resource ID		Congestion Severity			ECN	

- Possible use of DSCP bits (next to ECN) to be discussed @IETF
  - Multi-Resource Congestion Experienced (1 bit)
  - Resource ID (2 bits) → radio, computing (RAN), link bandwidth, computing (cloud/edge)
  - Congestion Severity (3 bits)
  - ECN: untouched (11 if congestion experienced)

Mario Patetta, Sami Taktak, Stefano Secci "Congestion Notification for Multi-Resource Allocation Control", draft-ietf-iccrg-mrecn-00 (work in progress), presented at IETF 124, Montreal, Canada, Nov 4, 2025.

MR-ECN integration in Smart-NICs

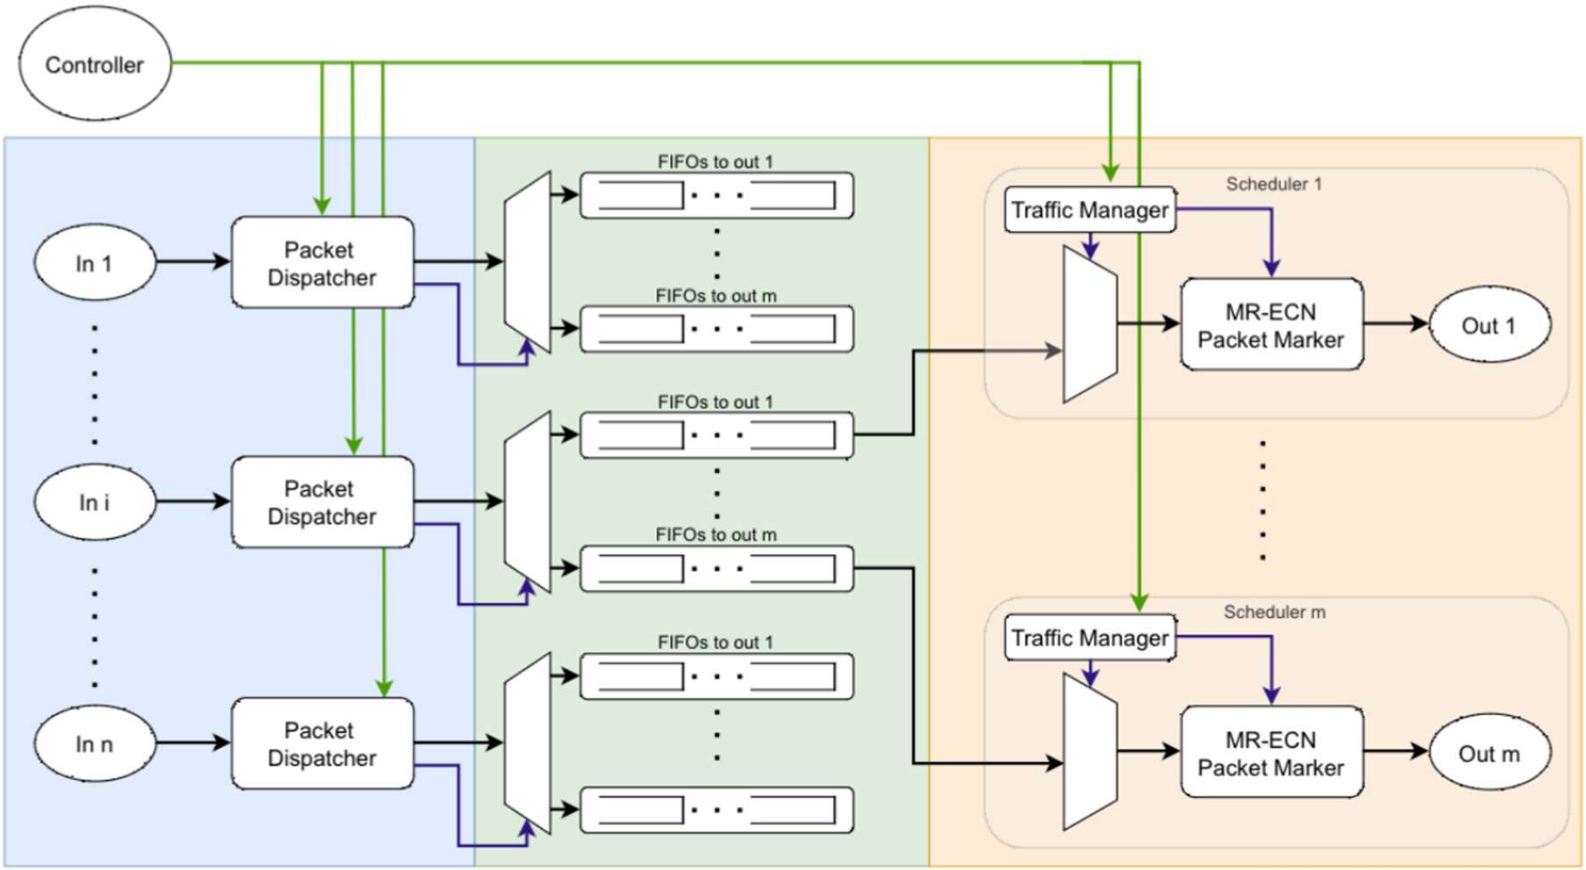

MR-ECN programmable scheduler block diagram.

Example of MRECN outcome

- Congestion of resource *j* allocated to 'slice' *i*:

$$congestion_{ij} = \left(1 - \frac{available_{ij}}{allocated_{ij}}\right)$$

- Congested resources constitute a **bottleneck**
  - → should adjust other allocations

$$waste_i = \max_j \left(\sum_{t=0}^{T} \left| \frac{adjusted_{ij}(t)}{allocated_{ij}} - \frac{available^c(t)}{allocated^c} \right| \Delta t \right)$$

- If no adjustment (adjusted = allocated)
  - → accumulated waste is the cumsum of congestion over time

Example:
Slice *i* is assigned resources:
R1=20 R2= 480 R3=8

R3 availability suddenly drops to 6 (congestion=0.25) for T=10 Δt

→ 25% of R1 and R2 is wasted*

After 3 Δt, R1 is adjusted to 20*0.75=15
After 4 Δt, R2 is adjusted to 480*0.7=336 (over adjusted)

→ accumulated waste over time
wasted R1 = 0.25*3 Δt = 0.75 Δt
wasted R2 = 0.25*4 Δt + 0.05*6 Δt = 1.3 Δt

Waste = max(0.75 Δt, 1.3 Δt) = 1.3 Δt

* assuming proportional dependency, aka Leontief utility

Sample of MR-ECN behavior

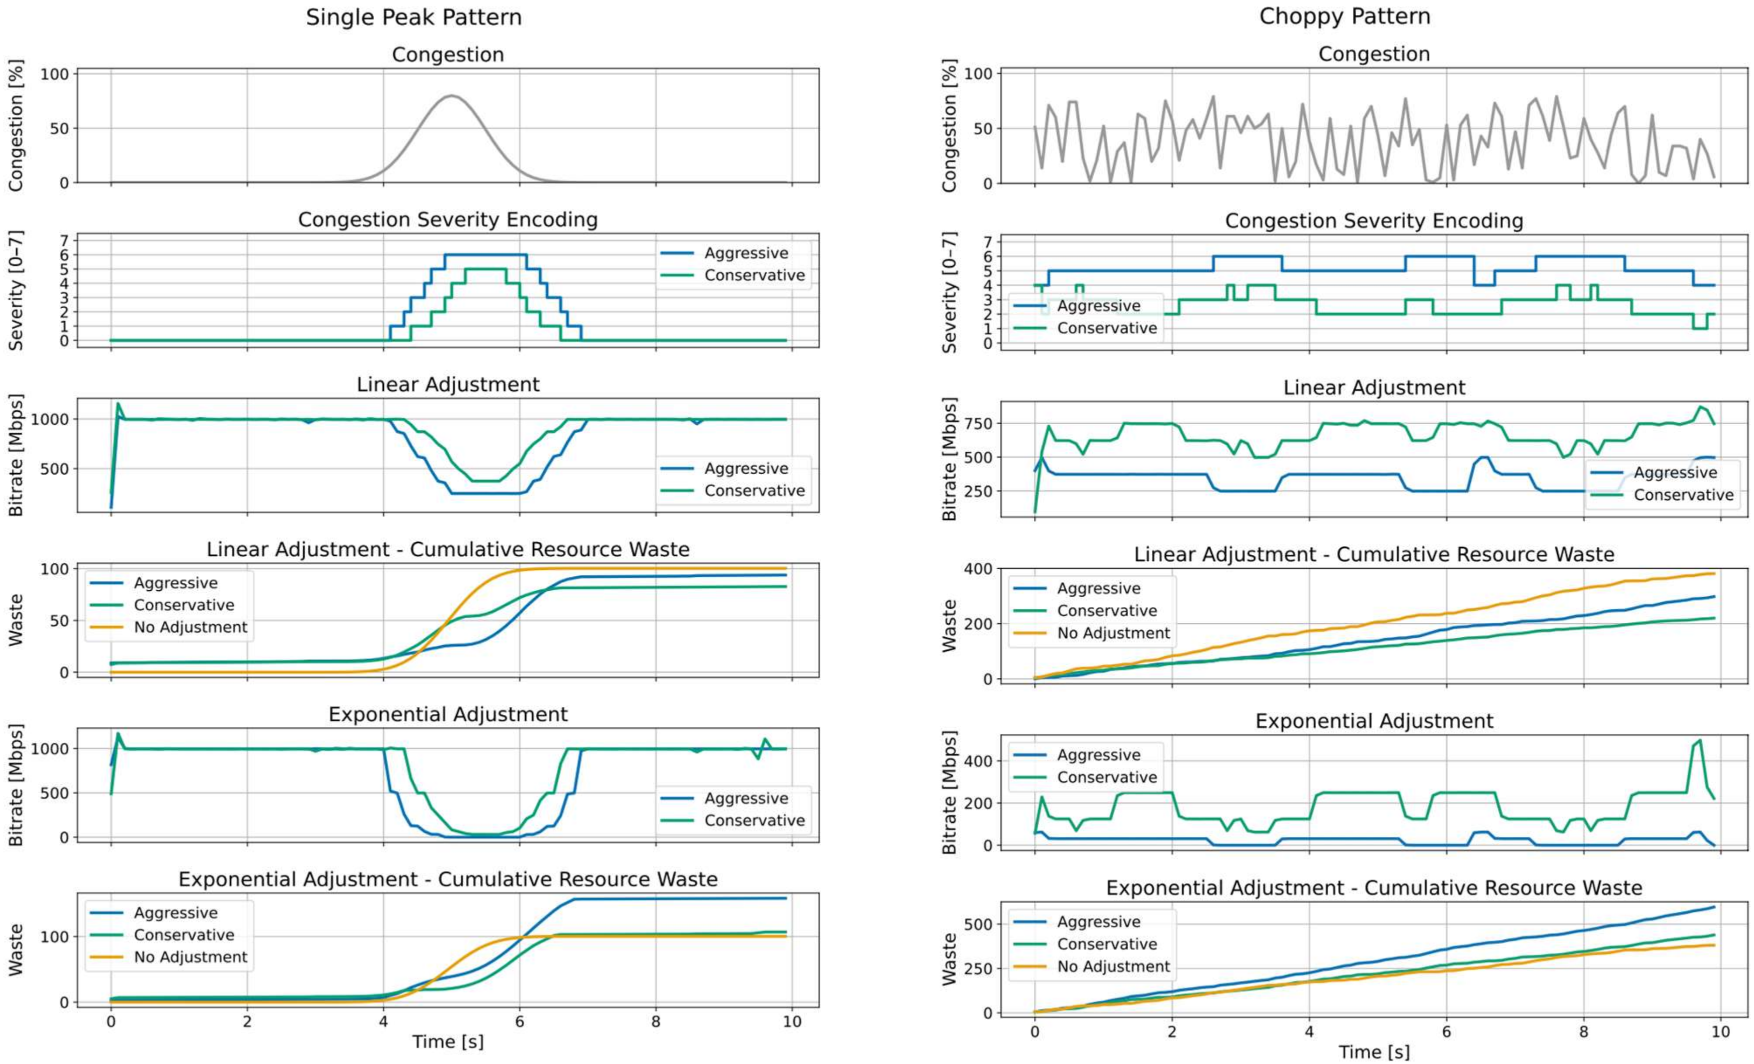

Questions ?

References

- Francesca FOSSATI, Stefano MORETTI, Stefano SECCI, "A Mood Value for Fair Resource Allocations", Proc. of IFIP Networking 2017, June 12-16, 2017, Stockholm, Sweden.
- Francesca FOSSATI, Sahar HOTEIT, Stefano MORETTI, Stefano SECCI, "Fair Resource Allocation in Systems with Complete Information Sharing", IEEE/ACM Transactions on Networking, Vol. 26, No. 6, pp: 2801-2814, Dec. 2018.
- Francesca FOSSATI, Deep MEDHI, Stefano MORETTI, Stefano SECCI, "Error Estimate and Fairness in Resource Allocation with Inaccurate Information Sharing", IEEE Networking Letters, Vol. 1, No. 4, pp: 173-177, Dec. 2019.
- Francesca FOSSATI, Stefano MORETTI, Stéphane ROVEDAKIS, Stefano SECCI, "Decentralization of 5G slice resource allocation", Proc. of 2020 IFIP/IEEE Int. Symposium on Network Operations and Management Systems (NOMS), Apr. 20-24, 2020, Budapest, Hungary.
- Francesca FOSSATI, Stefano MORETTI, Patrice PERNY, Stefano SECCI, "Multi-Resource Allocation for Network Slicing", *IEEE/ACM Transactions on Networking*, Vol. 28, No. 3, pp: 1311-1324, June 2020.
- Francesca FOSSATI, Stephane ROVEDAKIS, Stefano Secci, "Distributed algorithms for multi-resource allocation ",*IEEE Transactions on Parallel and Distributed Computing*, Vol. 33, No. 10, pp:2524-2539, Oct. 2022.

- Mario Patetta, Stefano Secci, Sami Taktak, Multi-Resource Explicit Congestion Notification, submitted to IEEE Transactions on Networking.
- Zeidan Braik, Sahar Hoteit, Francesca Fossati, Stefano Secci, DDRF: Dependency-aware Dominant Resource Fairness, submitted to IEEE Transactions on Networking.
- Mario Patetta, Stefano Secci, Sami Taktak, “Congestion Notification for Multi-Resource Allocation Control”, draft-ietf-iccrg-mrecn-00 (work in progress), presented at IETF 124, Montreal, Canada, Nov 4, 2025.