

Optimal placement of Controllers in a resilient SDN Architecture

DRCN 2016

Nancy Perrot , Orange Labs,
with Thomas Reynaud (Centrale Paris, Orange)

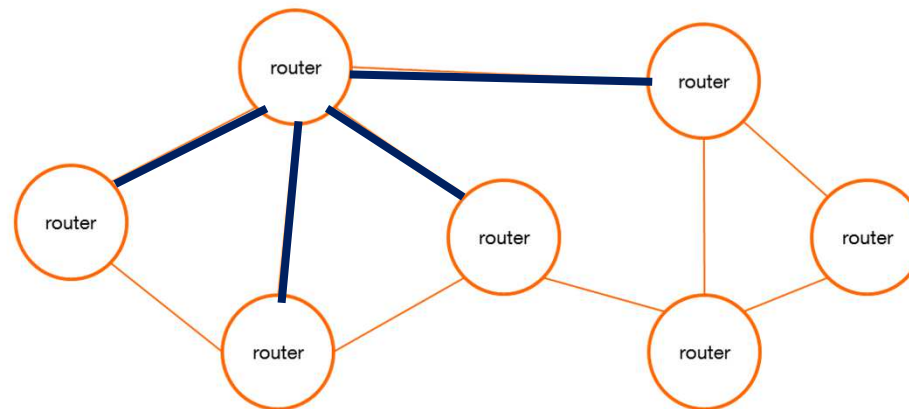


Outline

1. SDN in a few words
2. Motivation
3. Controller Placement Problem (CPP)
4. Resilient Controller Placement Problem (RCPP)
5. Conclusions and perspectives

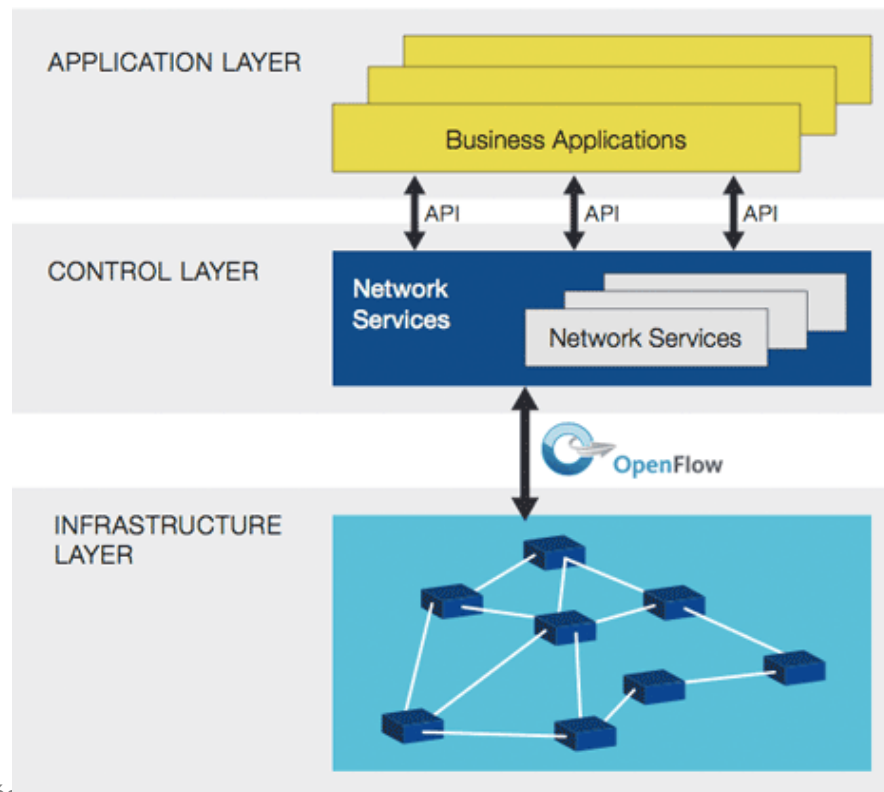
Introduction to SDN

- Data network devices (switch, router) have always been embedding three planes of operation :
 - **Forwarding Plane :**
 - Responsible for carrying user traffic, it moves packets from input to output
 - **Control Plane :**
 - Determines how packets should be forwarded
 - Responsible for signaling
 - **Management plane :**
 - Responsible for configuration of the control plane



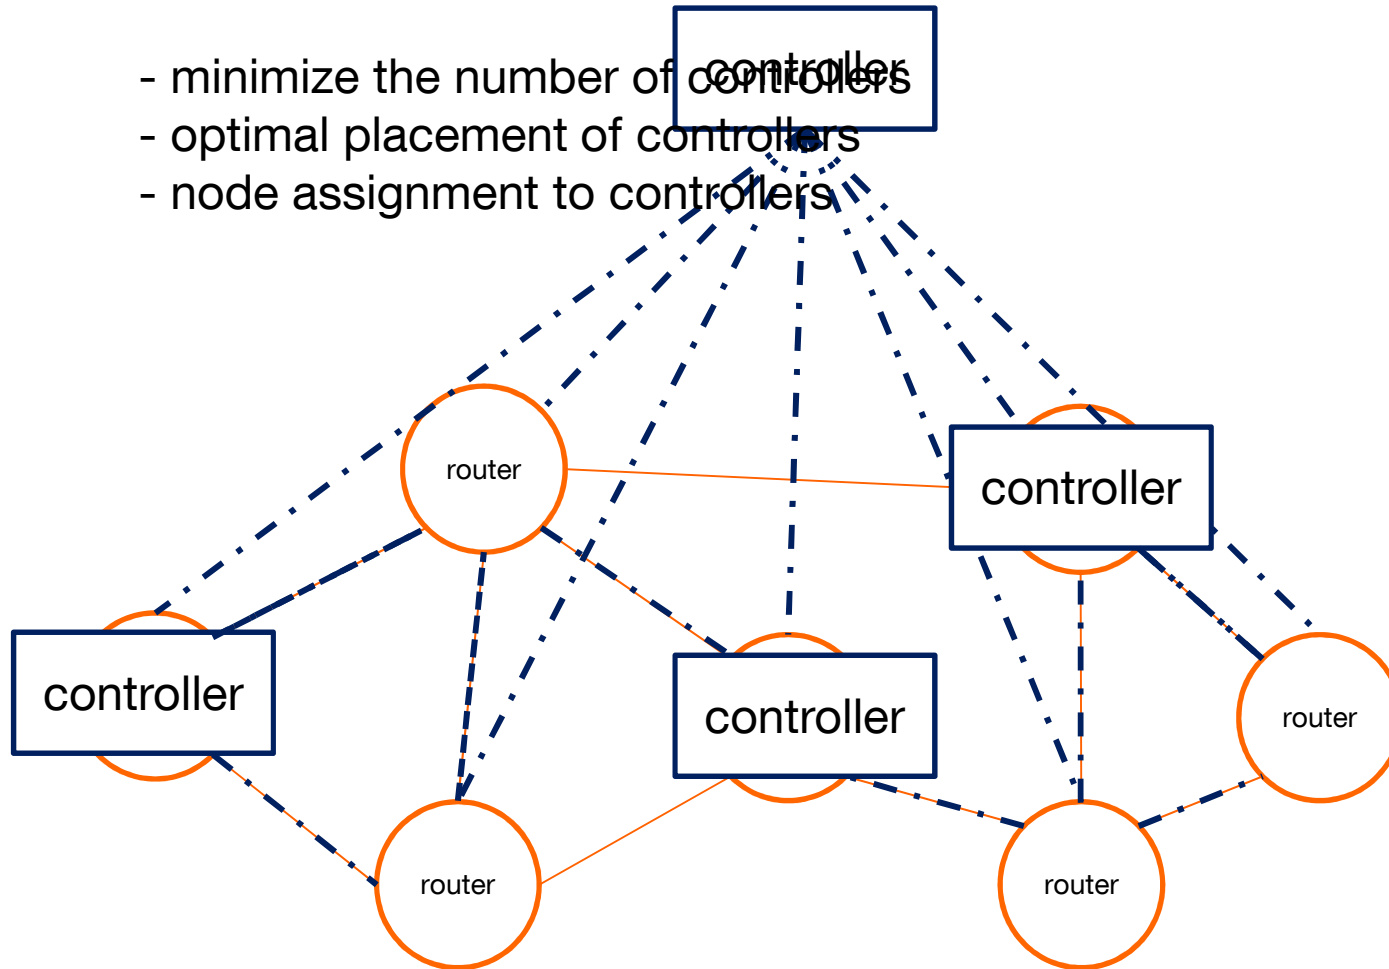
The SDN Paradigm

- SDN is the physical separation of the network control plane from the forwarding plane. The control plane controls several devices.
 - The network control becomes directly programmable
 - The underlying infrastructure is abstracted for applications and network services.




SDN in Wide Area Networks

- minimize the number of controllers
- optimal placement of controllers
- node assignment to controllers



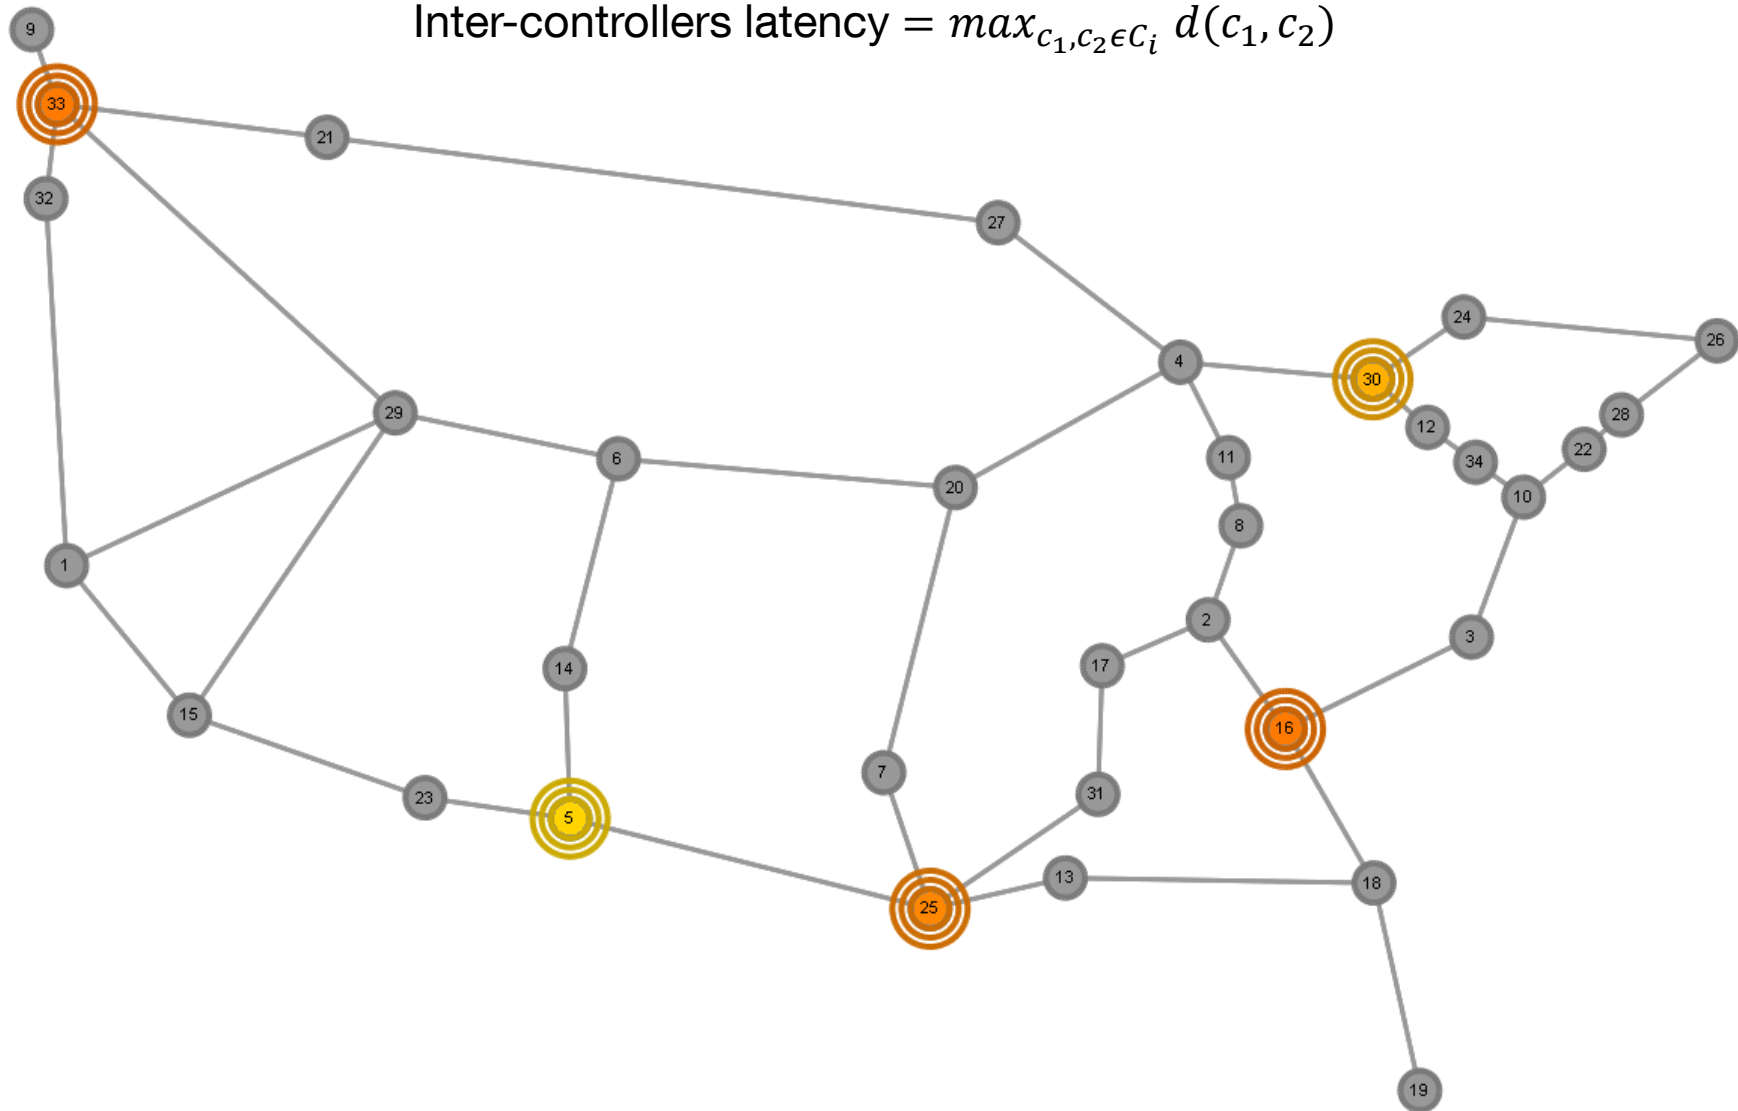
What is a “good” placement ?

$$\text{maximal latency} = \max_{v \in V} \min_{c \in C_i} d(v, c)$$



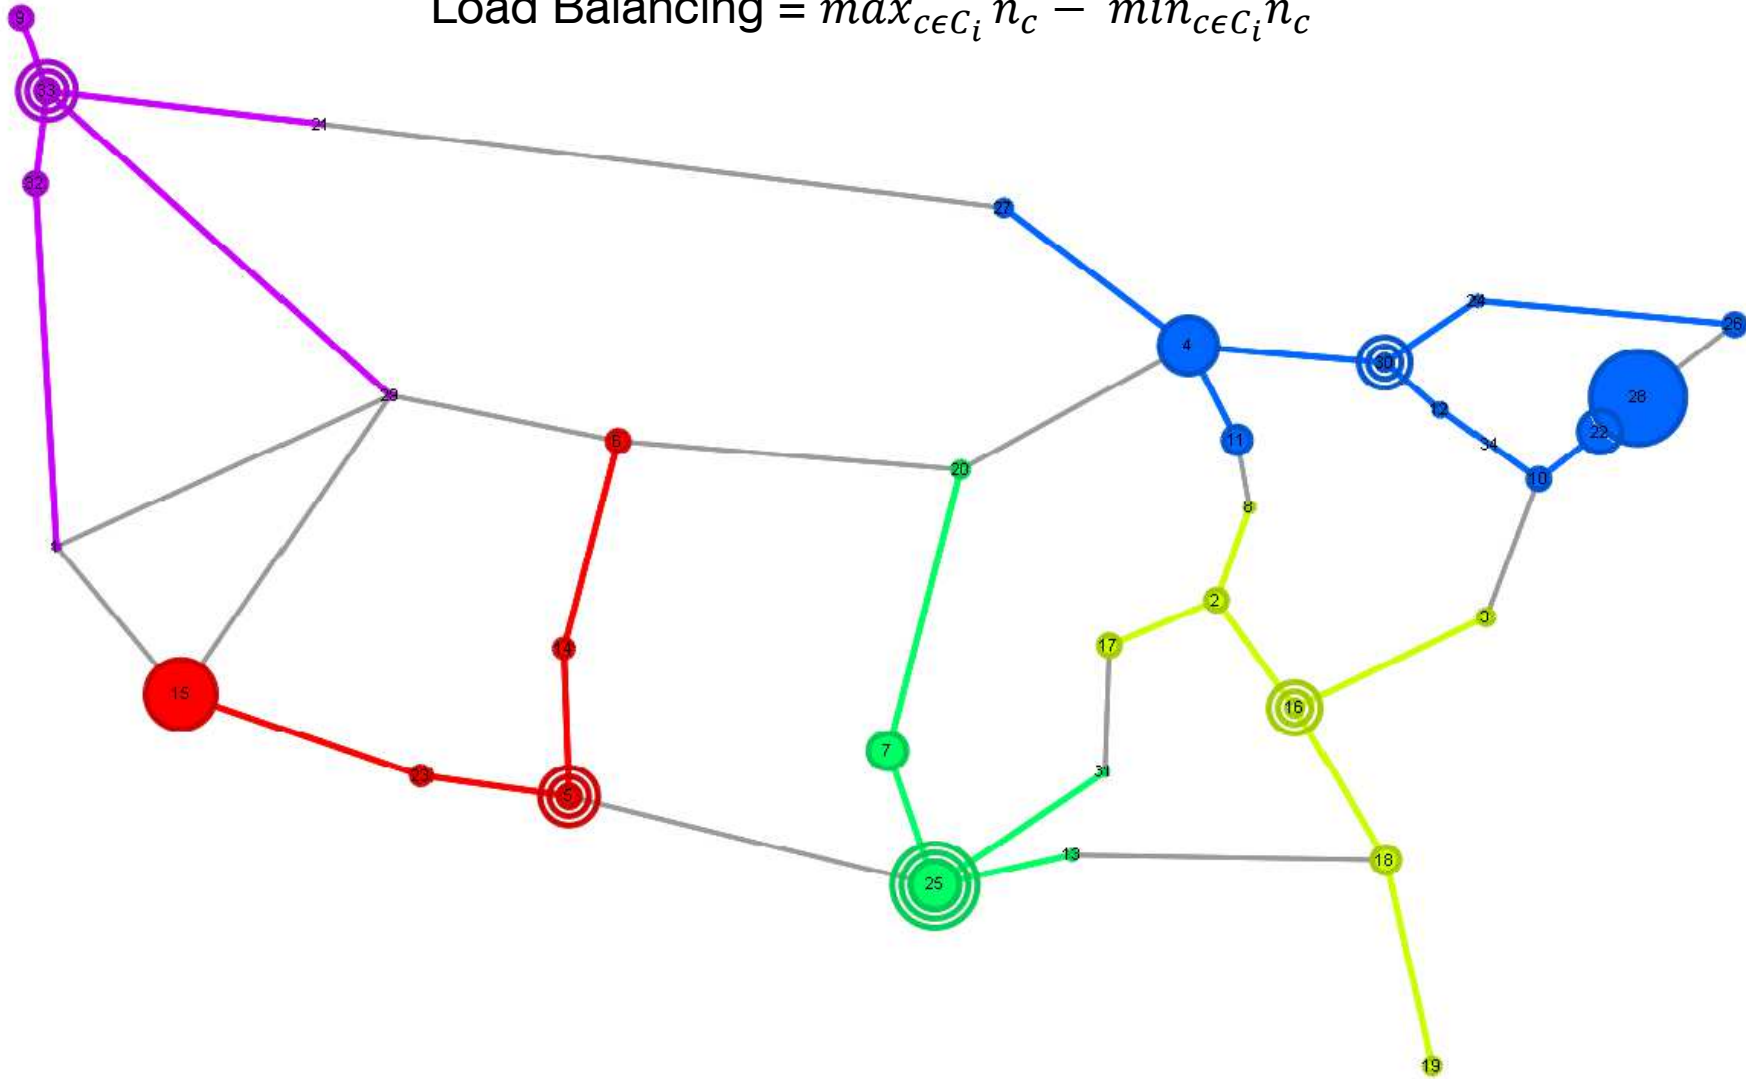
What is a “good” placement ?

$$\text{Inter-controllers latency} = \max_{c_1, c_2 \in C_i} d(c_1, c_2)$$



What is a “good” placement ?

$$\text{Load Balancing} = \max_{c \in C_i} n_c - \min_{c \in C_i} n_c$$



The Controller Placement Problem

- A Facility Location Problem with
 - maximal latency between the controller and its assigned nodes l_{max}
 - maximal latency between two controllers l_{cc-max}
 - load balancing constraints
- The binary variables are
 - assignment variables : $x_{ij} \in \{0,1\}$
 - active controller variables : $y_i \in \{0,1\}$
 - linearization variables : $t_{ii'} = y_i y_{i'} \in \{0,1\}$
- $d_{(i,j)}$ shortest path between router j and controller i
- covering matrix
 - $a_{ij} = \begin{cases} 1 & \text{iff } d_{ij} \leq l_{max} \\ 0 & \text{otherwise} \end{cases}$
- objective : $\min \sum_{i \in C} y_i$

The CPP - explicit formulation

- each router j must be covered by at least one controller within the latency bound :

$$\sum_{i \in C} a_{ij} y_i \geq 1 \quad \forall j \in R$$

- each router j must be assigned to the nearest active controller i :

$$\left\{ \begin{array}{l} \sum_{i \in C} x_{ij} = 1 \quad \forall j \in R \\ x_{ij} \leq y_i \quad \forall i \in C, \forall j \in R \\ y_{\sigma_{jq}} \leq \sum_{m=1}^q x_{j\sigma_{jm}} \quad \forall j \in R, \forall q \in [1, |C| - 1] \end{array} \right.$$

- all pairs of controllers must respect the allowed inter-controllers latency

$$t_{ii'} d_{ii'} \leq l_{cc-max} \quad \forall i, i' \in C$$

The CPP - explicit formulation

- the difference of load between all pairs of controllers must be at most δ :

$$-\delta - (|R| - \delta)(1 - t_{ii'}) \leq \sum_{j \in R} (x_{ij} - x_{i'j}) \leq \delta + (|R| - \delta)(1 - t_{ii'}) \quad \forall i, i' \in C$$


- linking variables constraints :

$$\begin{aligned} t_{ii'} &\geq y_i + y_{i'} - 1 && \forall i, i' \in C \\ t_{ii'} &\leq y_i && \forall i, i' \in C \\ t_{ii'} &\leq y_{i'} && \forall i, i' \in C \end{aligned}$$


- $x_{ij}, y_i, z_j, t_{ii'} \in \{0,1\}$

The Controller Placement Problem

Effect of load balancing constraints on COST topology




optimal solution with $\delta=3$




relaxation of load balancing constraints

The Controller Placement Problem

Effect of the maximal delay on the COST topology




delay max is 15% of graph diameter




delay max is 39% of graph diameter

The Controller Placement Problem

Effect of the maximal delay on Zib topology



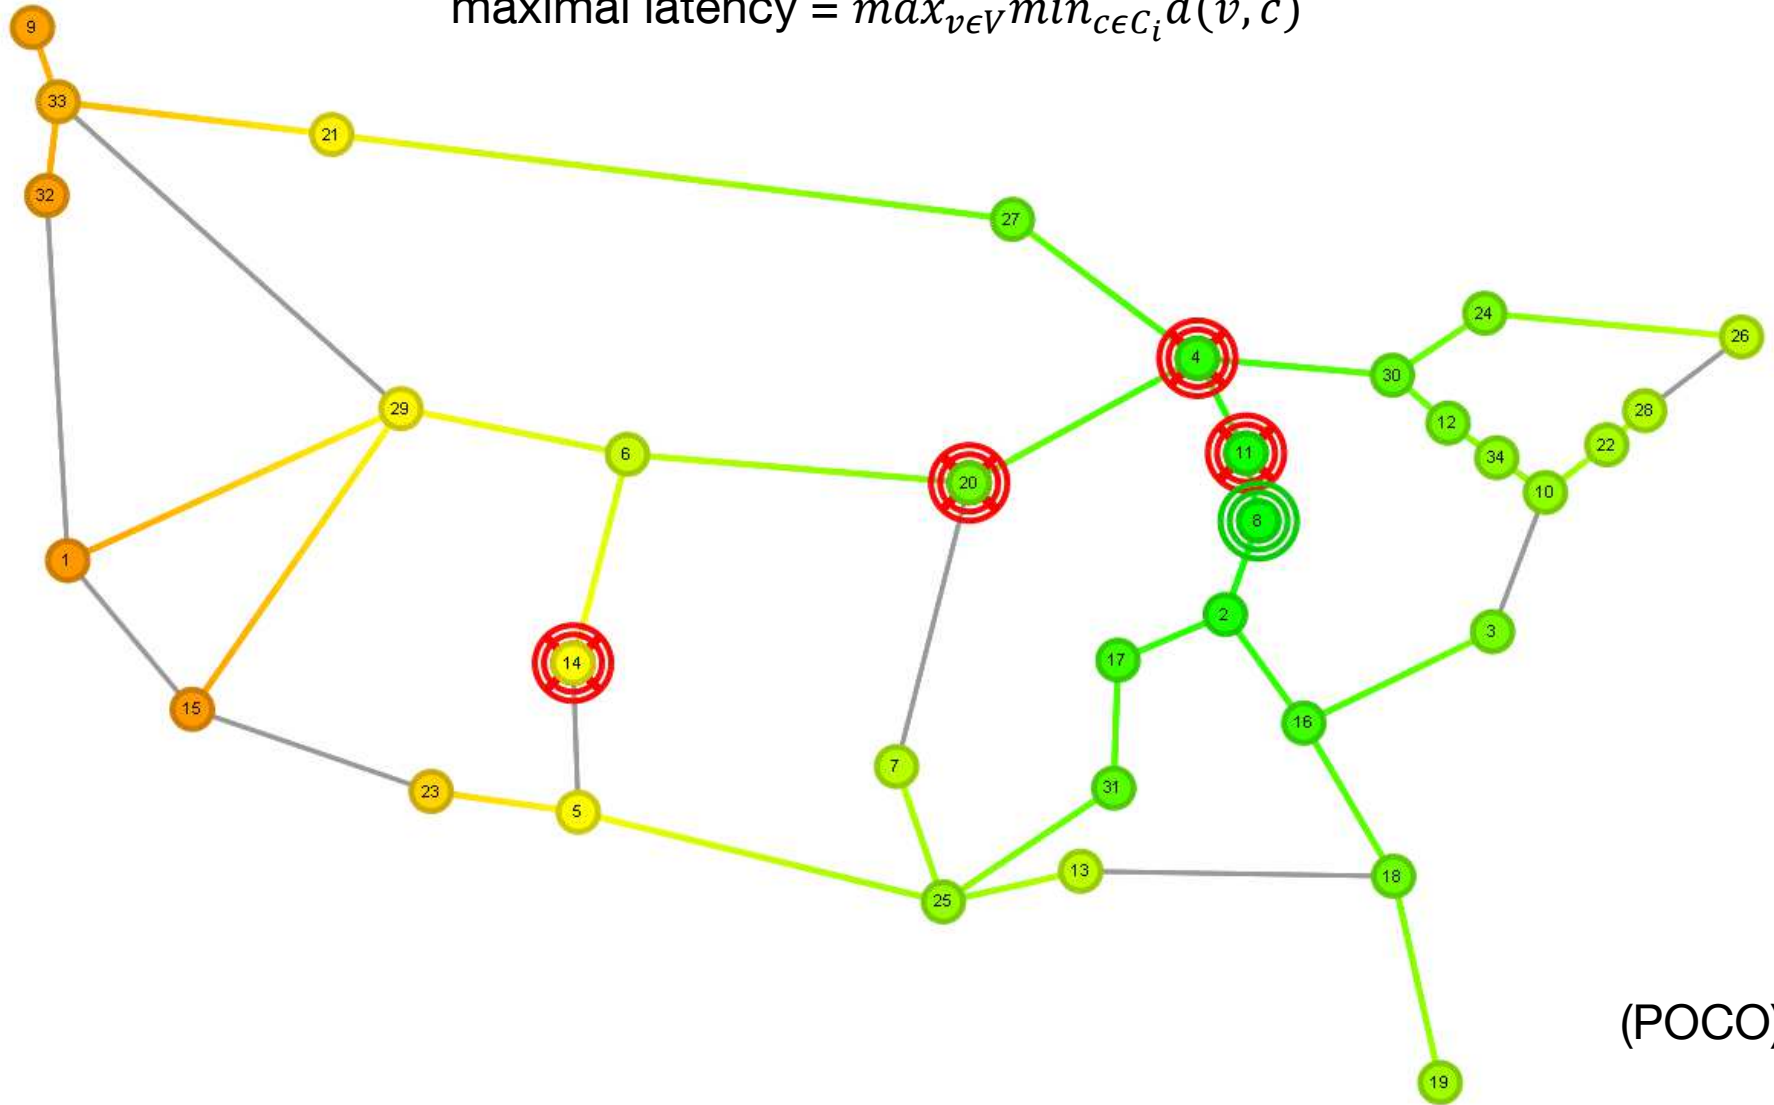
max latency is 30%



max latency is 50%

What is a “good” placement ?

$$\text{maximal latency} = \max_{v \in V} \min_{c \in C_i} d(v, c)$$



The Resilient Controller Placement Problem

- If a controller fails ? the nodes are assigned to another one
 - increases the latency between routers and controller
 - unbalanced domains (especially if the secondary controller takes the management of all the routers of the failed controller)
- we consider **simultaneously k levels** of controller failures.
- let p be the failure probability of a controller
- x_{ij}^k are re-assignment variables for each level of failure : 1 if controller i is the k^{th} backup controller of router j
- $z_j^k = 1$ if j has a $(k - 1)^{th}$ backup controller but not a k^{th} backup controller

$$\text{delay} = d_{ij} x_{ij}^k (1 - p)p^{k-1}$$

$$\text{penalty cost} = l_{max} z_j^k p^{k-1}$$

Resilient Controller Placement Problem

- bi-objective problem

- first objective

$$\min \sum_{i \in C} y_i$$

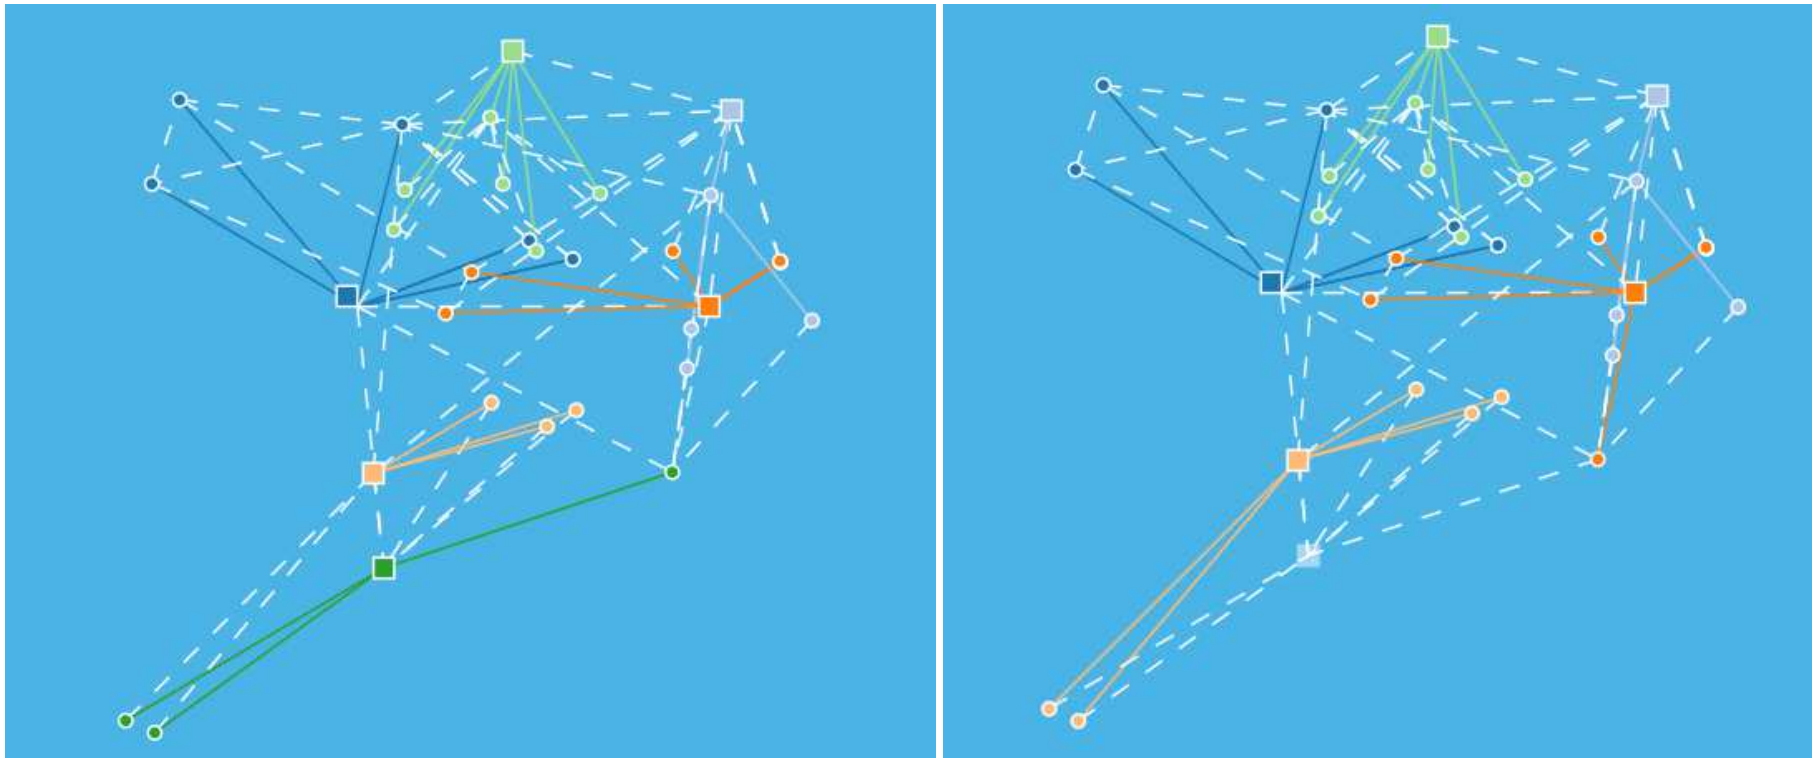
- second objective

$$\min \sum_{j \in R} \sum_{k=1}^{|C|} \sum_{i \in C} d_{ij} x_{ij}^k (1-p)p^{k-1} + l_{max} \sum_{j \in R} \sum_{k=1}^{|C|+1} p^{k-1} z_j^k$$

+ the same block of constraints than for CPP for all back-up levels k

- The solution consists of
 - the minimum number of controllers,
 - their placement among the candidate network nodes,
 - the assignment of network elements to controllers,
 - the re assignment in case of multiple failures of any controller with minimal degradation of QoS.


Resilient Controller Placement Problem



Simulation results

Parameters	Value
Number of nodes	[10, 80]
l_{max}	{3 000, 5 000, 7 000}
l_{cc-max}	7 000
δ	3
back-up levels	2


3 000 random graphs




■ CPU time on number of nodes

Simulation results

- Evolution of the number of controllers depending on the graph size/density.



Number of controllers on the number of nodes



Number of controllers on the number of arcs

Conclusions and Perspectives

- These formulations have been implemented in a decision-aid tool to simulate deployment scenario.
- Huge network instances : spectral clustering, relaxation-based heuristic.
- Study of dynamic resources re-assignment will depend on real use cases.