Poolability

Damon Wischik, UCL

trilogy

Question

We have built a multipath version of TCP. For it to be useful, end systems need to have access to multiple paths. What mechanisms / protocols / algorithms should the Internet have for providing multiple paths?

For example:
Can we rely on multihomed users? Or on peer-to-peer applications? Or do we really do need support in the core of the Internet, e.g. through a change to BGP?

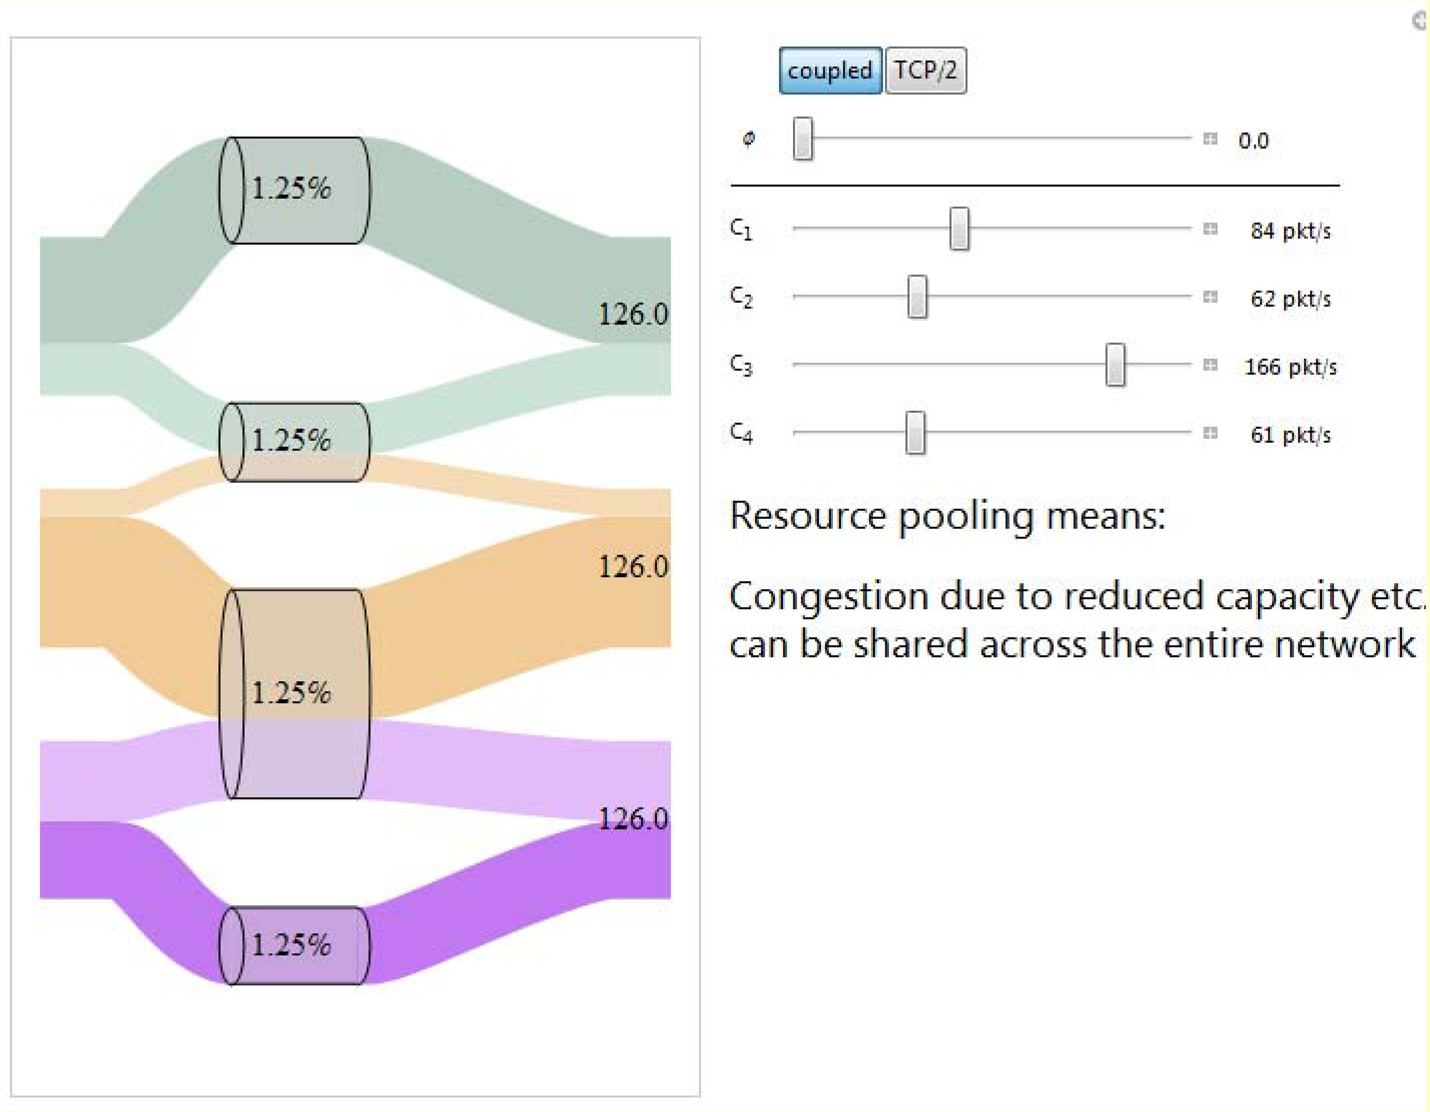

Note: these figures show the outcome from simulations of multipath TCP.

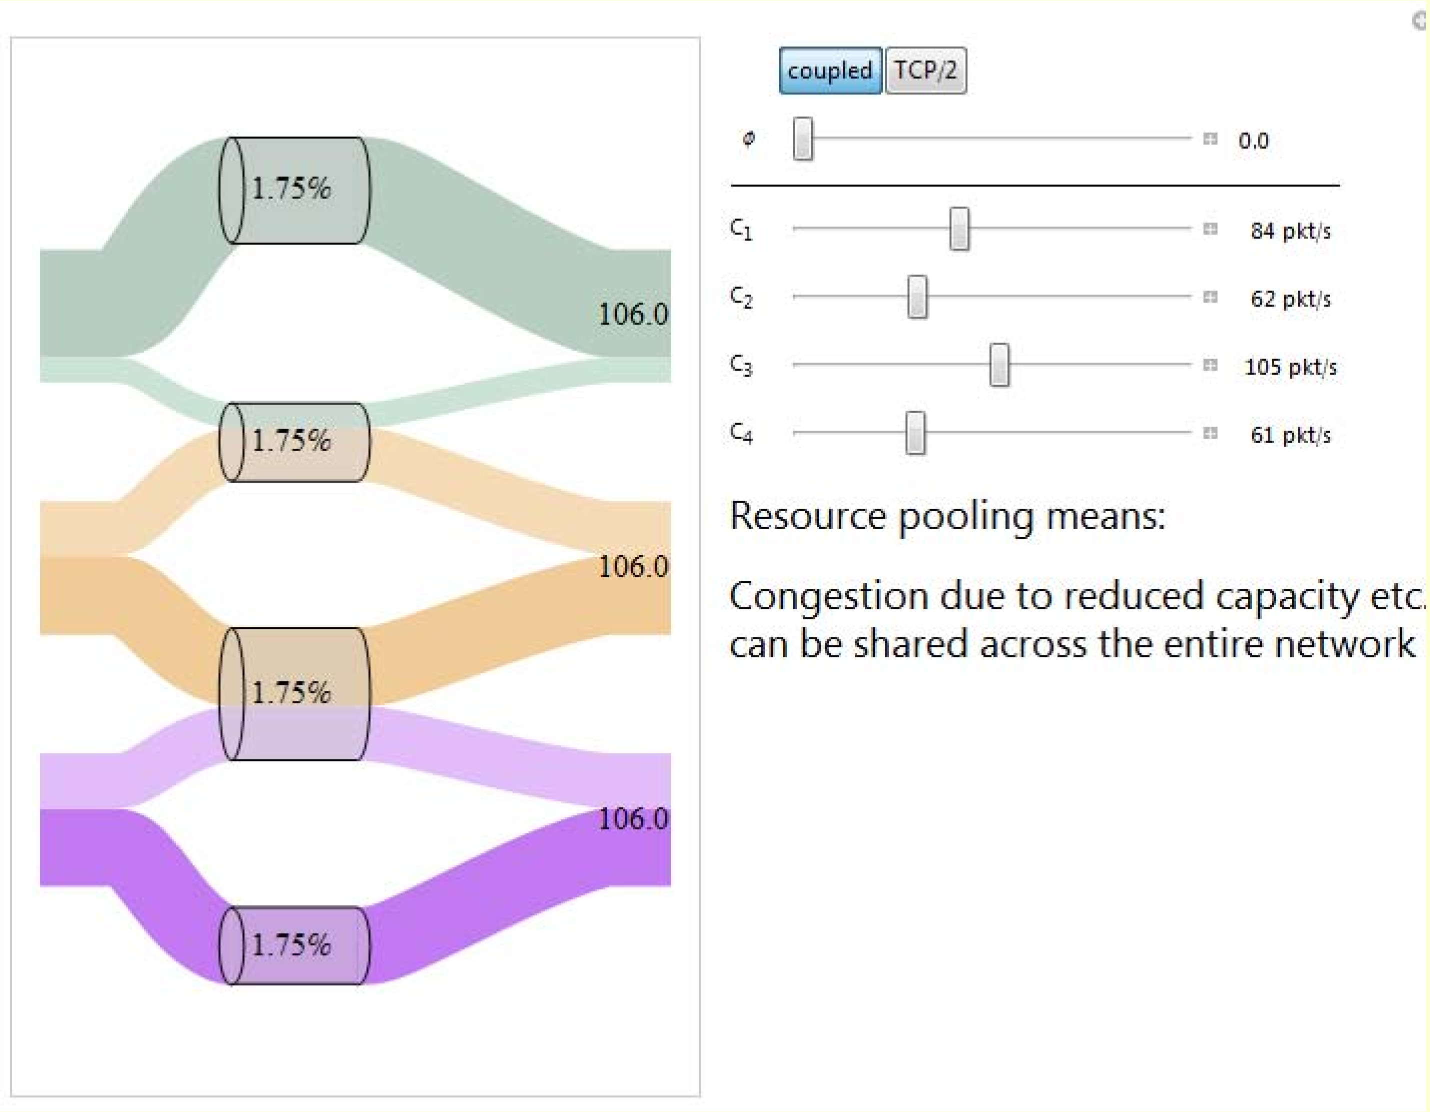

Note: these figures show the outcome from simulations of multipath TCP.

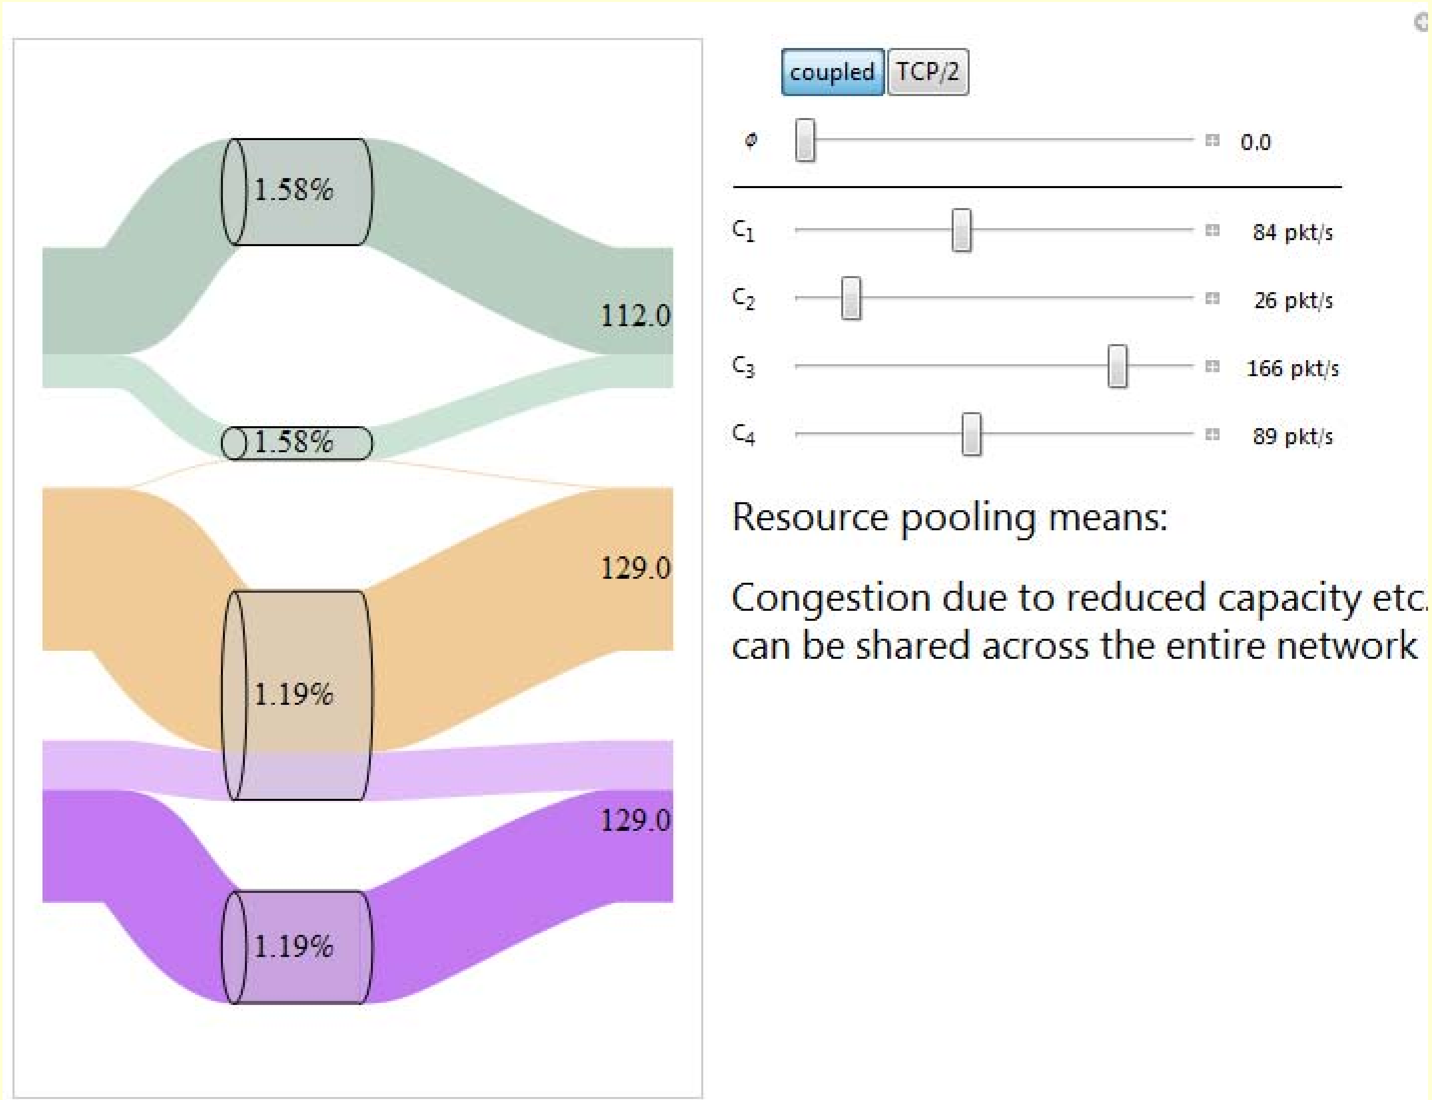

Note: these figures show the outcome from simulations of multipath TCP.

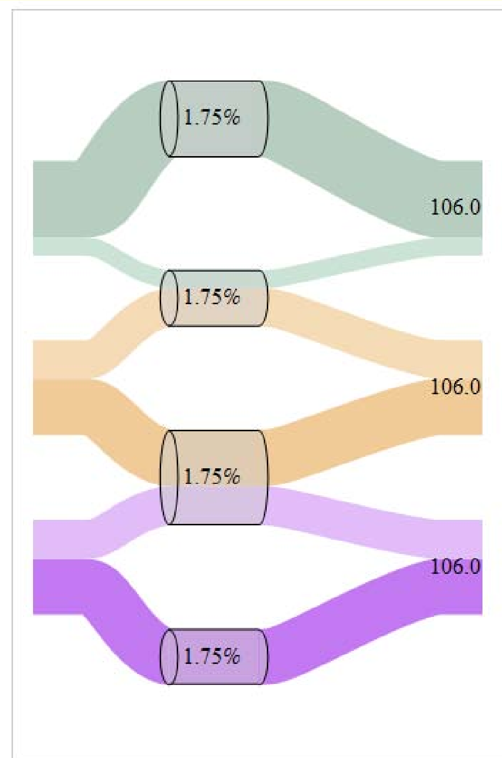

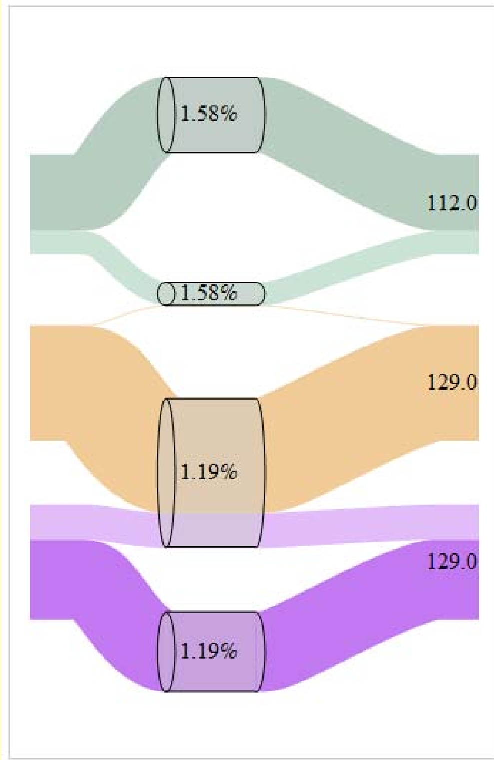

Sometimes the network behaves like a single "pooled resource". This means it can easily adapt to changing conditions, and it shares total capacity as fairly as possible.

But sometimes the network behaves like several separate "pooled resources".

How can we predict which?

A useful first step:
remind ourselves of the classic multi-commodity flow problem

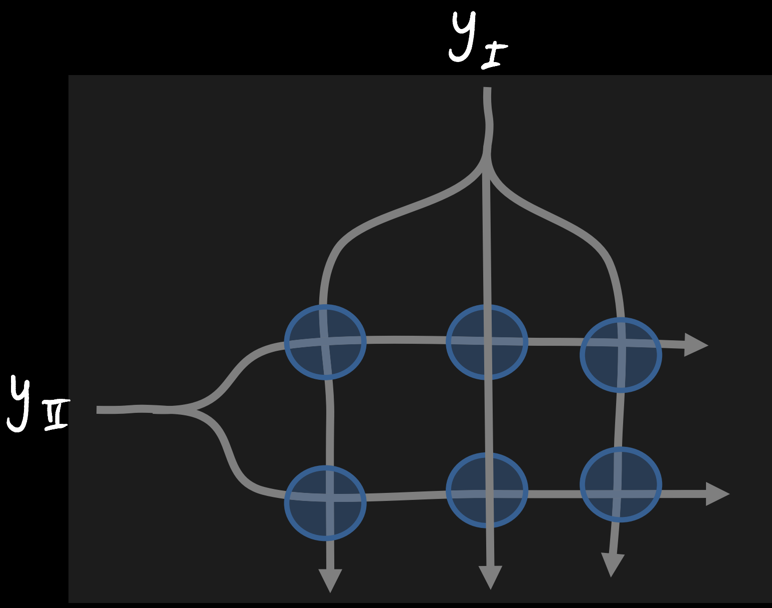

Given the capacities of the six resources in this diagram, what flow rates y_I and y_{II} are feasible?

The capacity constraints in a multicommodity flow problem can be written as *generalized cut constraints.*

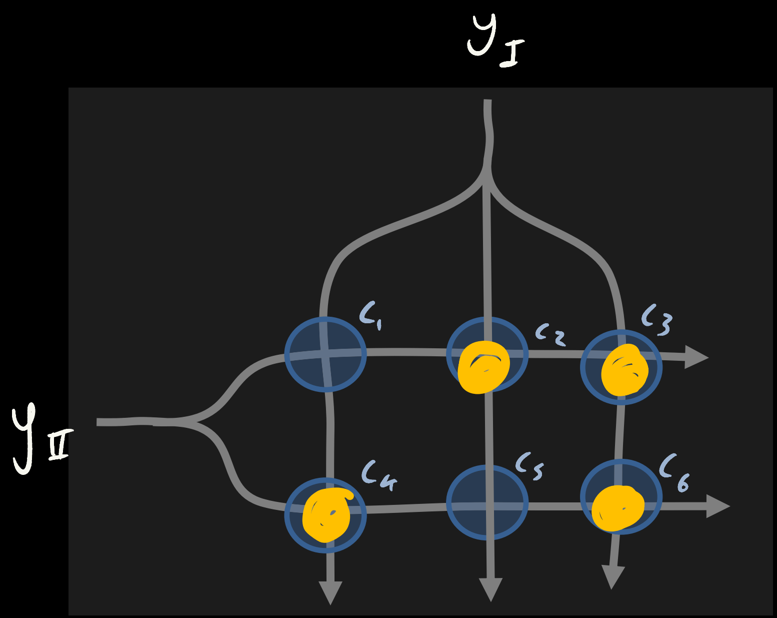

For (y_I, y_{II}) to be feasible, it is necessary that

$$y_I \le c_1 + c_2 + c_3$$
$$y_{II} \le c_1 + c_4$$
$$y_{II} \le c_1 + c_5$$

$$2y_I + 3y_{II} \le 2c_2 + 2c_4 + c_3 + c_6$$

and 25 other inequalities.

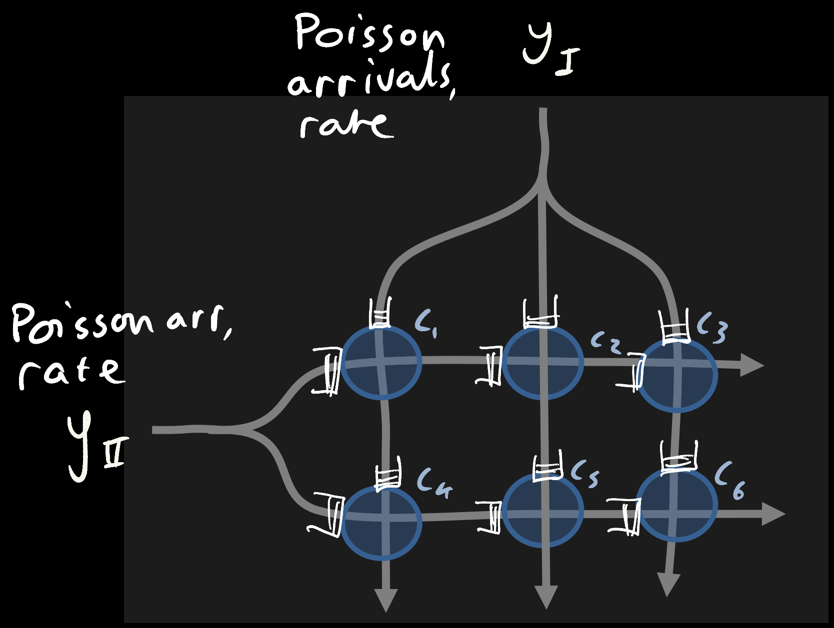

These generalized cut constraints also tell us something useful about queueing networks in heavy traffic.

[Laws, 1992, “Resource pooling in queueing networks”.]

The generalized cut constraints show up in the multicommodity flow problem, and in the heavy traffic queueing problem.

Do they also show up multipath TCP?

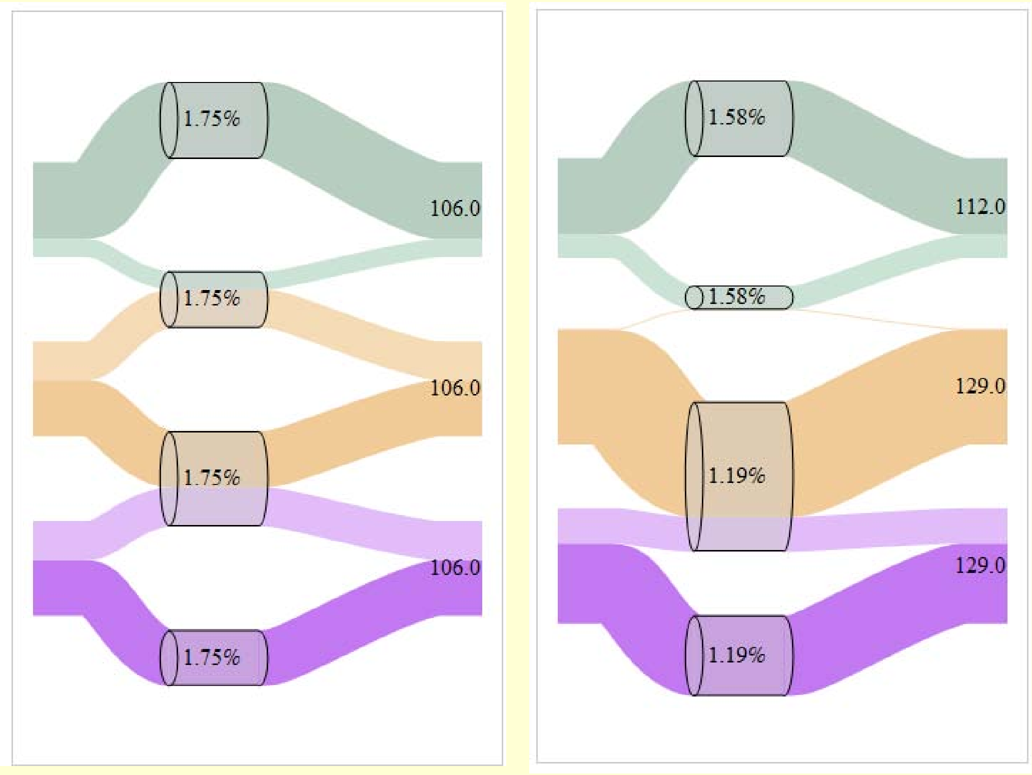

In this experiment, I asked:

"What is the effect on flow rates and on drop probabilities, of changing the capacity at one of the links?"

Let's turn this into a maths problem.

We'll know we've got the right answer, if the generalized cut constraints drop out in the analysis.

A simple flow allocation problem

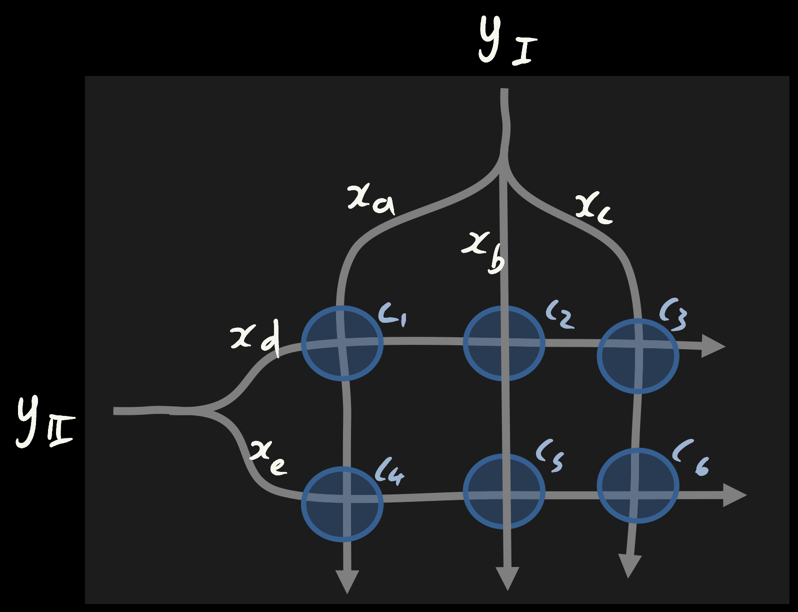

Suppose we are given concave utility functions. Consider the problem:

maximize $U_I(y_I) + U_{II}(y_{II})$

over $x \geq 0,\ y \geq 0$

such that

$$y_I = x_a + x_b + x_c$$

$$y_{II} = x_d + x_e$$

$$x_a + x_d \leq c_1$$

$$x_a + x_e \leq c_4$$

$$\vdots$$

A simple flow allocation problem

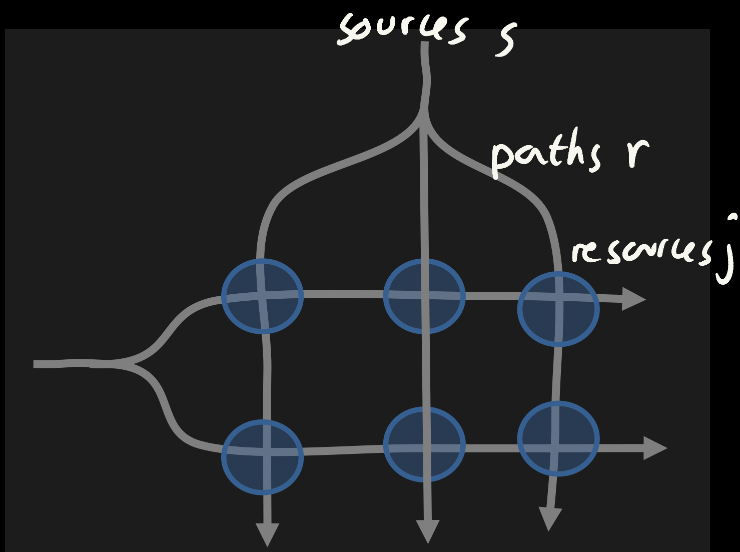

maximize $\sum_s U_s(y_s)$

over $x \geq 0, \quad y \geq 0, \quad z \geq 0$

such that

$$y = Hx$$

$$z = Ax, \quad z \leq C$$

$$H = \begin{array}{c} \\ I \\ II \end{array} \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{bmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{array}$$

$$A = \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{array} \begin{array}{c} \begin{array}{ccccc} a & b & c & d & e \end{array} \\ \begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \end{bmatrix} \end{array}$$

A simple flow allocation problem (relaxed version)

It can be shown that a suitable multipath TCP solves this problem.

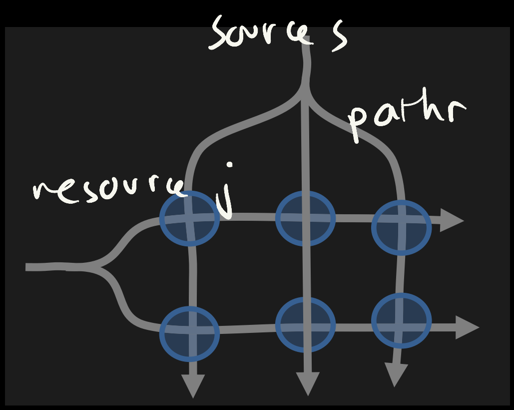

maximize $\sum_s U_s(y_s) - \sum_j c_j L_j(\rho_j)$

over $x \geq 0,\ y \geq 0,\ z \geq 0$

such that

$$y = Hx$$
$$z = Ax$$
$$\rho_j = z_j / c_j$$

where $L_j(\rho) = \int_0^\rho p_j(\rho)\, d\rho$

and $p_j(\rho)$ = packet drop probability at link j, when the load is ρ

We want to know how the solution changes when capacities change.
I shall take *y* to be fixed, and only look at how *x* changes.

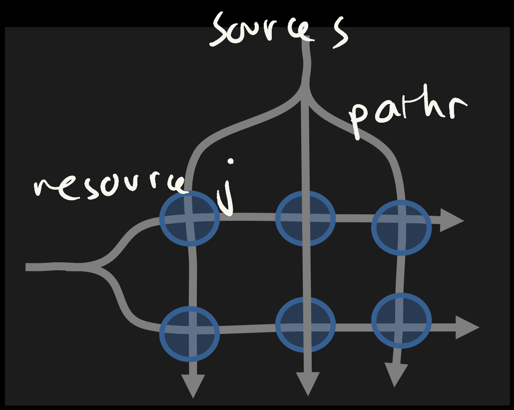

maximize $\sum_s U_s(y_s) - \sum_j c_j L_j(\rho_j)$

over $x \geq 0,\ y \geq 0,\ z \geq 0$

such that
$$y = Hx$$
$$z = Ax$$
$$\rho_j = z_j / c_j$$

where $L_j(\rho) = \int_0^\rho p_j(\rho)\, d\rho$

and $p_j(\rho)$ = packet drop probability at link j, when the load is ρ

We want to know how the solution changes when capacities change. I shall take y to be fixed, and only look at how x changes.

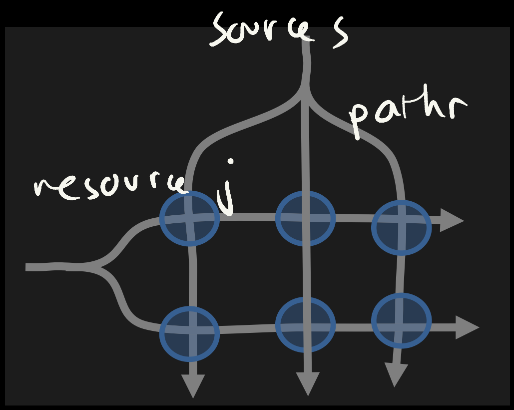

$$\text{minimize} \quad \sum_j C_j L_j(\rho_j)$$

$$\text{over} \quad x \geq 0, \quad z \geq 0$$

$$\text{such that} \quad y = Hx$$

$$z = Ax$$

$$\rho_j = z_j / C_j$$

Write out the complementary slackness conditions

Take the total derivative with respect to C_j for some j

Solve for dz_i/dC_j using linear algebra

Theorem

At an isolated link, $\dfrac{dp}{dC} = \dfrac{p}{\tilde{C}}$ where $\tilde{C} = \dfrac{C}{L''(p)}$

In a network with idealized multipath congestion control

$$\frac{dp_j}{dC_j} = \frac{p_j}{\left(\dfrac{\tilde{C}_j}{1-\Psi_{jj}}\right)} \quad \text{where} \quad \tilde{C}_j = \frac{C_j}{L''(p_j)}$$

I call Ψ_{jj} the "poolability score", and $C_j/(1-\Psi_{jj})$ the "effective pooled capacity". If $\Psi_{jj} \approx 1$ then the link sheds load easily. If $\Psi_{jj} \approx 0$ then the link is "solitary".

Here, $$\begin{bmatrix} \Psi \\ \Phi \end{bmatrix} = \begin{bmatrix} \bar{A} & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} A^T \tilde{C}^{-1} \bar{A} & -\bar{H}^T \\ H & 0 \end{bmatrix} \begin{bmatrix} A^T \tilde{C}^{-1} \\ 0 \end{bmatrix}$$

and $$\tilde{C} = \begin{bmatrix} c_1/L_1''(p_1) & & 0 \\ & \ddots & \\ 0 & & c_n/L_n''(p_n) \end{bmatrix}$$

and $\bar{A}, \bar{H}$ are the adjacency matrix and the source/path matrix, restricted to paths with non-zero traffic.

My goal was to devise a metric for resource pooling. What is resource pooling? It means

> *"making a collection of resources behave like a single pooled resource".*

My goal was to devise a metric for resource pooling. What is resource pooling? It means

> *"making a collection of resources behave like a single pooled resource".*

"behave":

respond to changes in capacity

My goal was to devise a metric for resource pooling. What is resource pooling? It means

"making a collection of resources behave like a single pooled resource".

"behave":

respond to changes in capacity

"like a single resource":

At a single resource, $\frac{dp}{dc} = \frac{\rho}{\tilde{c}}$.

In a multipath network, $\frac{dp_j}{dc_j} = \frac{\rho_j}{\tilde{c}_j / 1-\Psi_{jj}}$.

So it's like a single link with capacity $\frac{\tilde{c}_j}{1-\Psi_{jj}}$. Call this the "effective pooled capacity"

If the poolability score is $\Psi_{jj} \approx 1$ then the link sheds load easily.
If the poolability score is $\Psi_{jj} \approx 0$ then the link is "solitary".

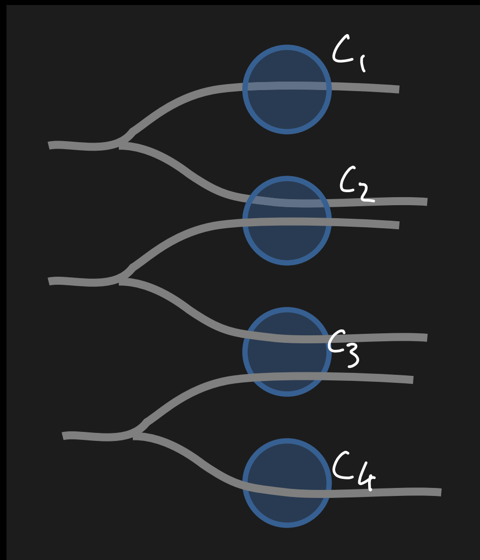

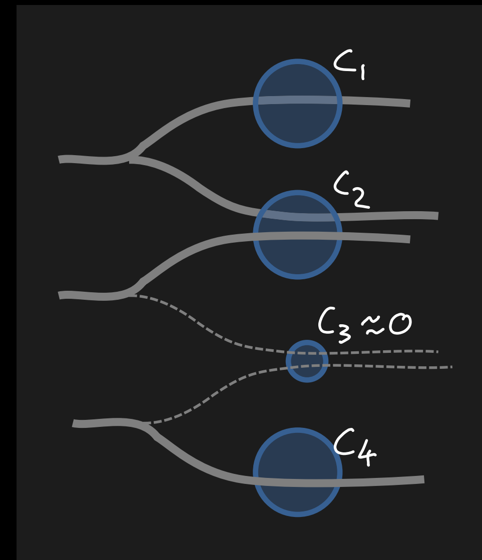

$$\Psi_{11} = \frac{\hat{c}_2 + \hat{c}_3 + \hat{c}_4}{\hat{c}_1 + \hat{c}_2 + \hat{c}_3 + \hat{c}_4}$$

$$\Psi_{22} = \frac{\tilde{c}_1 + \tilde{c}_3 + \tilde{c}_4}{\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3 + \tilde{c}_4}$$

$$\Psi_{33} = \frac{\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_4}{\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3 + \tilde{c}_4}$$

$$\Psi_{44} = \frac{\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3}{\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3 + \tilde{c}_4}$$

$$\Psi_{11} = \frac{\tilde{c}_2}{\tilde{c}_1 + \tilde{c}_2}$$

$$\Psi_{22} = \frac{\tilde{c}_1}{\tilde{c}_1 + \tilde{c}_2}$$

$$\Psi_{44} = 0$$

If the poolability score is $\Psi_{jj} \approx 1$ then the link sheds load easily.
If the poolability score is $\Psi_{jj} \approx 0$ then the link is "solitary".

For one particular set of link capacities,

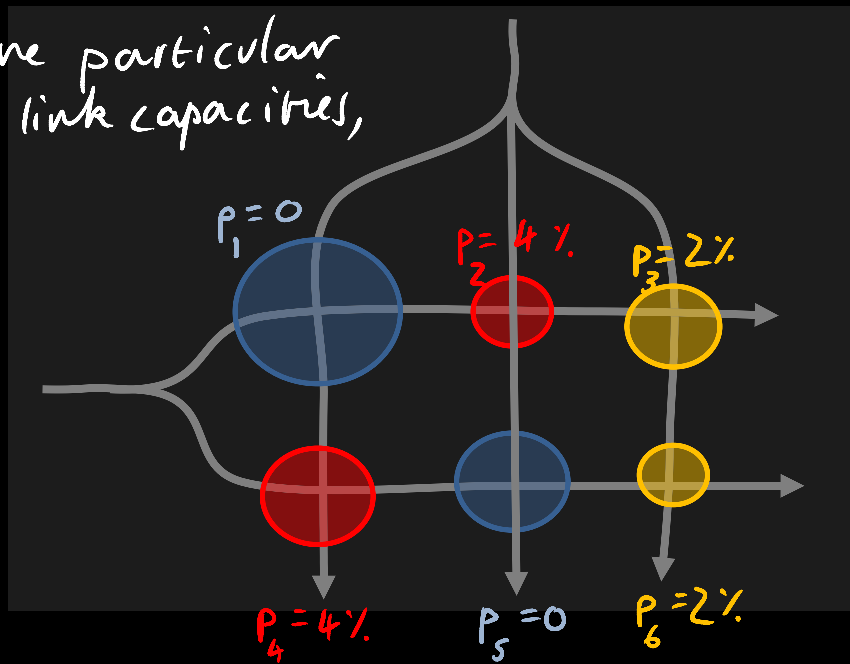

$$\Psi_{11} = 0$$

$$\Psi_{22} = \frac{4c_4 + c_3 + c_6}{4c_2 + 4c_4 + c_3 + c_6}$$

$$\Psi_{33} = \frac{4c_2 + 4c_4 + c_6}{4c_2 + 4c_4 + c_3 + c_6}$$

$$\Psi_{44} = \frac{4c_2 + c_3 + c_6}{4c_2 + 4c_4 + c_3 + c_6}$$

$$\Psi_{55} = 0$$

$$\Psi_{66} = \frac{4c_2 + 4c_4 + c_3}{4c_2 + 4c_4 + c_3 + c_6}$$

If the poolability score is $\Psi_{jj} \approx 1$ then the link sheds load easily.
If the poolability score is $\Psi_{jj} \approx 0$ then the link is “solitary”.

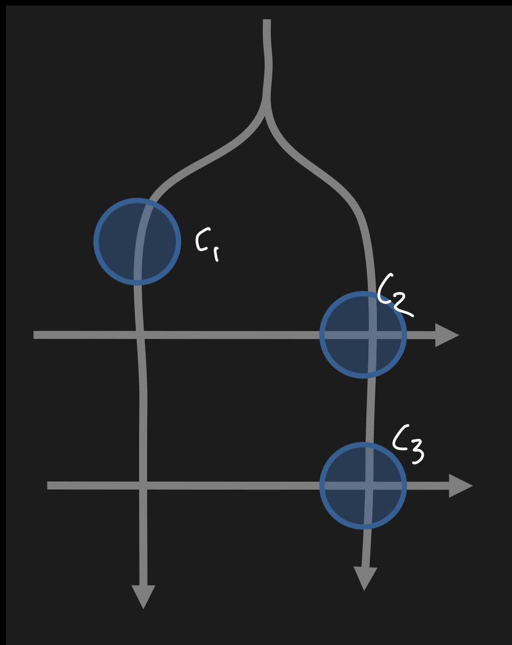

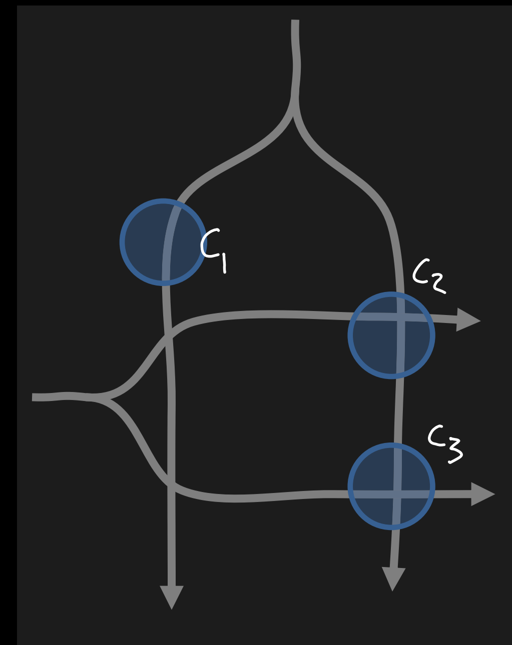

$$\Psi_{11} = \frac{1/\tilde{c}_1}{1/\tilde{c}_1 + 1/\tilde{c}_2 + 1/\tilde{c}_3}$$

$$\Psi_{22} = \frac{1/\tilde{c}_2}{1/\tilde{c}_1 + 1/\tilde{c}_2 + 1/\tilde{c}_3}$$

$$\Psi_{33} = \frac{1/\tilde{c}_3}{1/\tilde{c}_1 + 1/\tilde{c}_2 + 1/\tilde{c}_3}$$

$$\Psi_{11} = \frac{\tilde{c}_2 + \tilde{c}_3}{4\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3} = \frac{1/\tilde{c}_1}{1/\tilde{c}_1 + 2/\frac{1}{2}(\tilde{c}_2 + \tilde{c}_3)}$$

$$\Psi_{22} = \frac{4\tilde{c}_1 + \tilde{c}_3}{4\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3}$$

$$\Psi_{33} = \frac{4\tilde{c}_1 + \tilde{c}_2}{4\tilde{c}_1 + \tilde{c}_2 + \tilde{c}_3}$$

The same "magic coefficients" turn up in the multicommodity flow problem, in the heavy traffic queueing model.

Now we know they turn up in the rate allocation problem, and we know how to interpret them.

Can we use them to decide on a good protocol for multipath routing? Can they help in working out which parts of the network would benefit from extra capacity, or which flows would benefit from extra paths?

Is there an inductive way to calculate the "magic coefficients", based on the formulae for series and parallel resistor networks?

GEANT data provided by UCL Belgium
multipath routes, link capacities, and traffic matrices

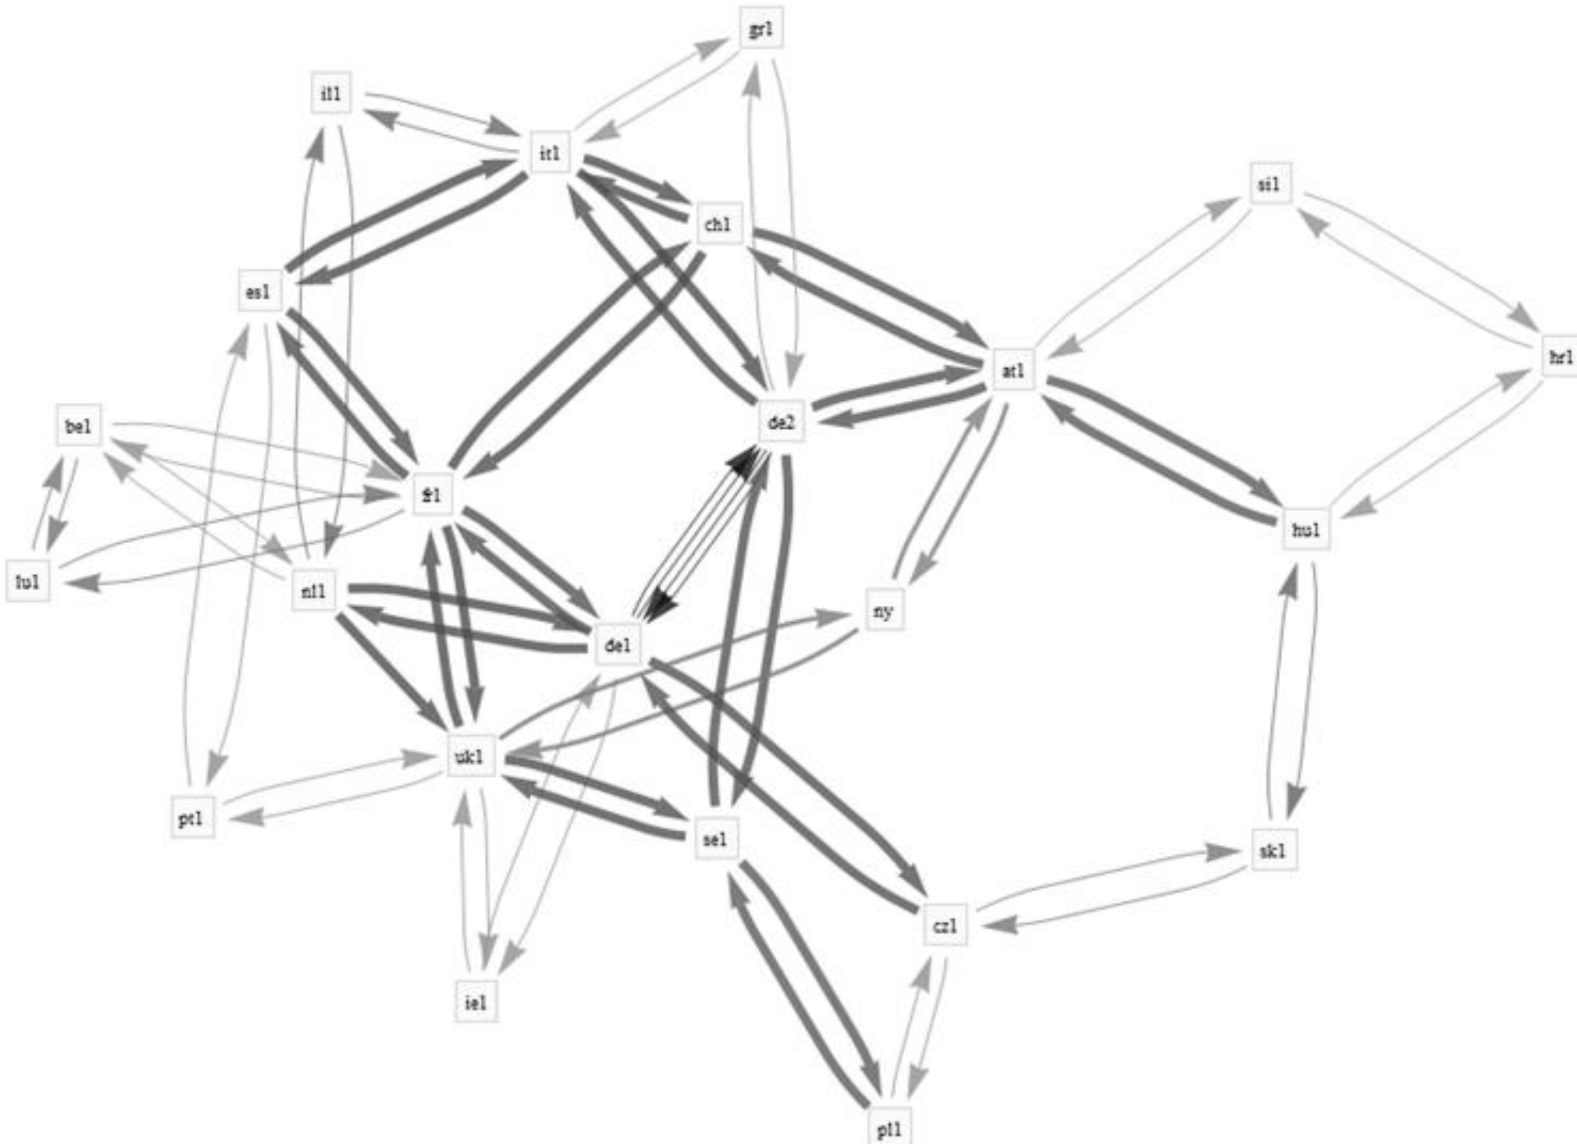

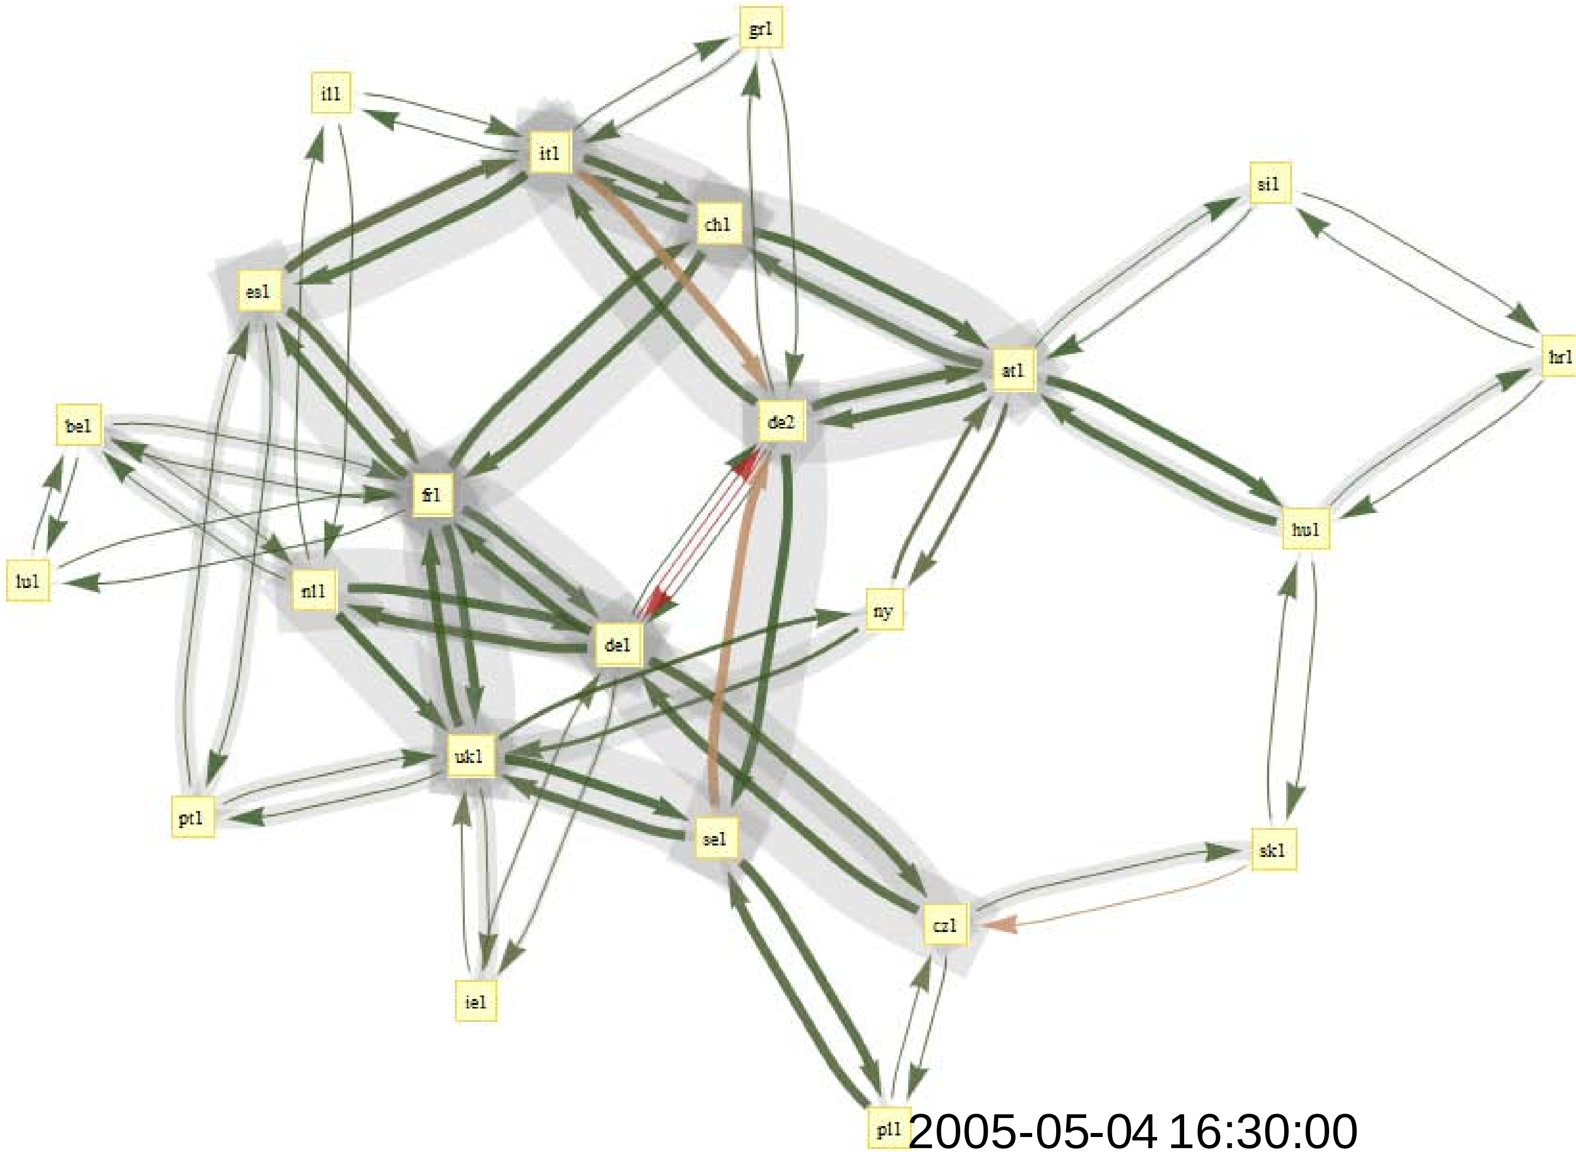

2005-05-04 16:30:00
Colours show utilization
Grey shows effective pooled capacity

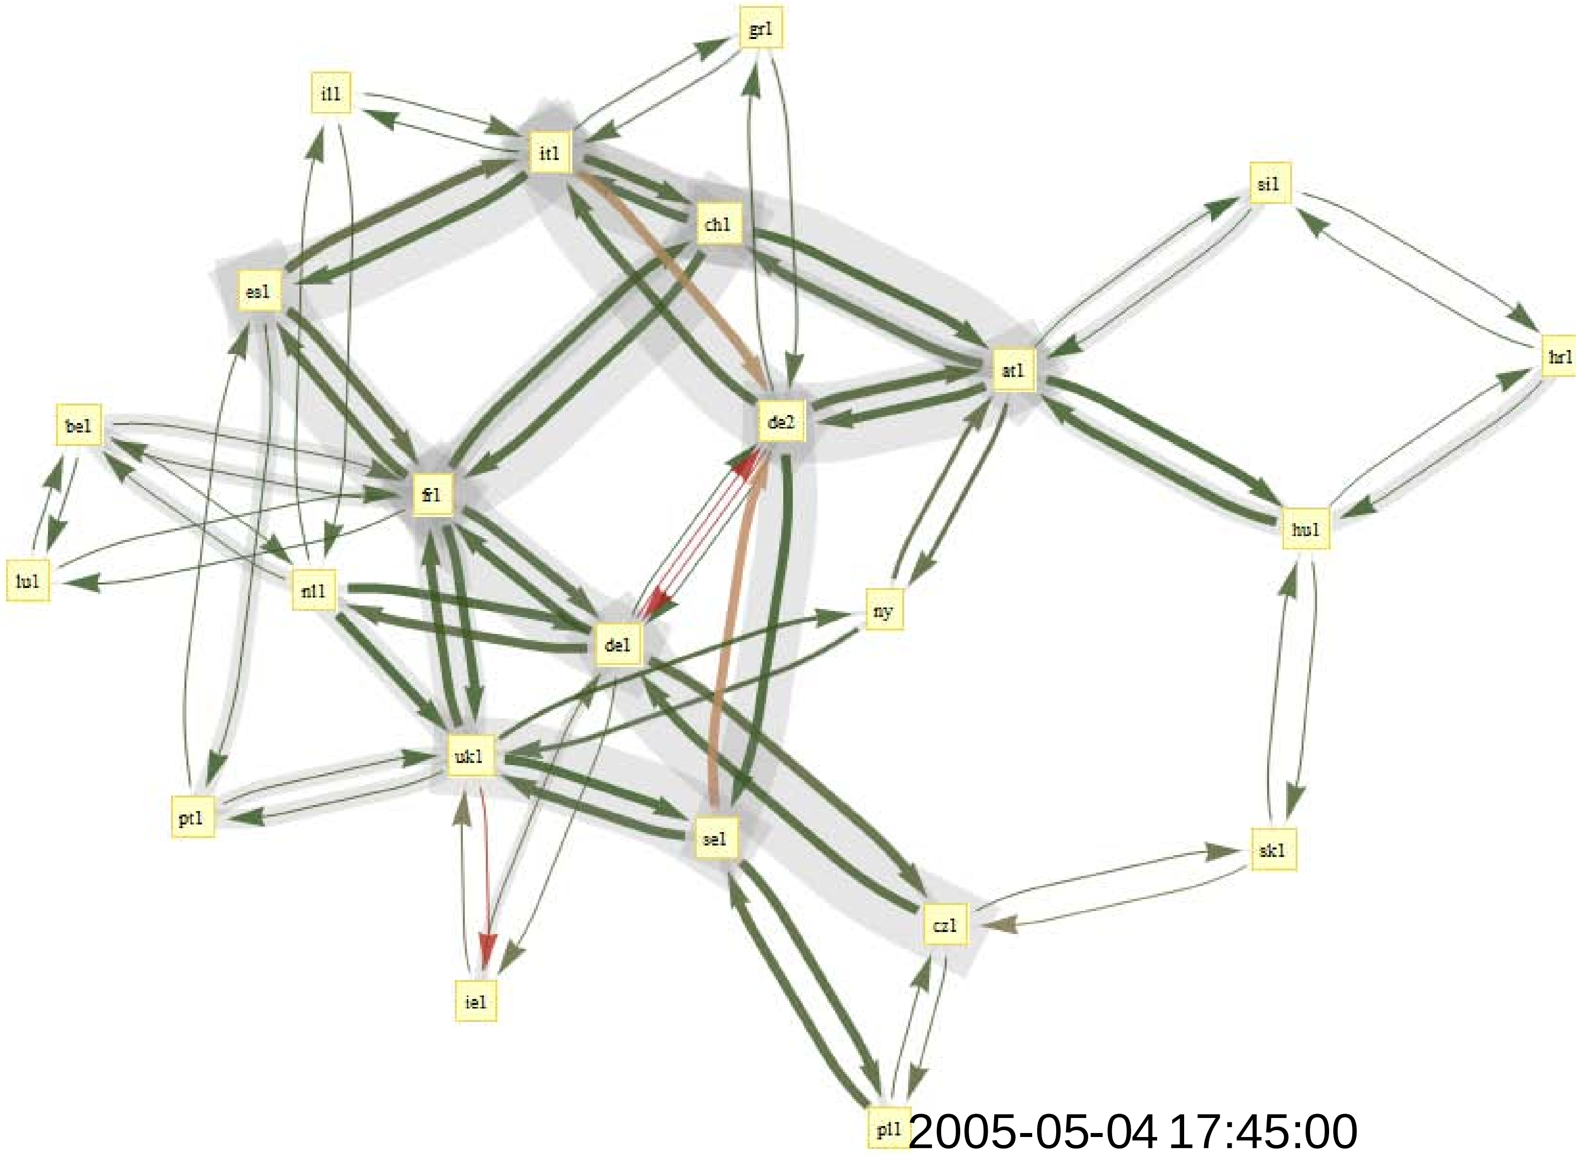

2005-05-04 17:45:00
Colours show utilization
Grey shows effective pooled capacity

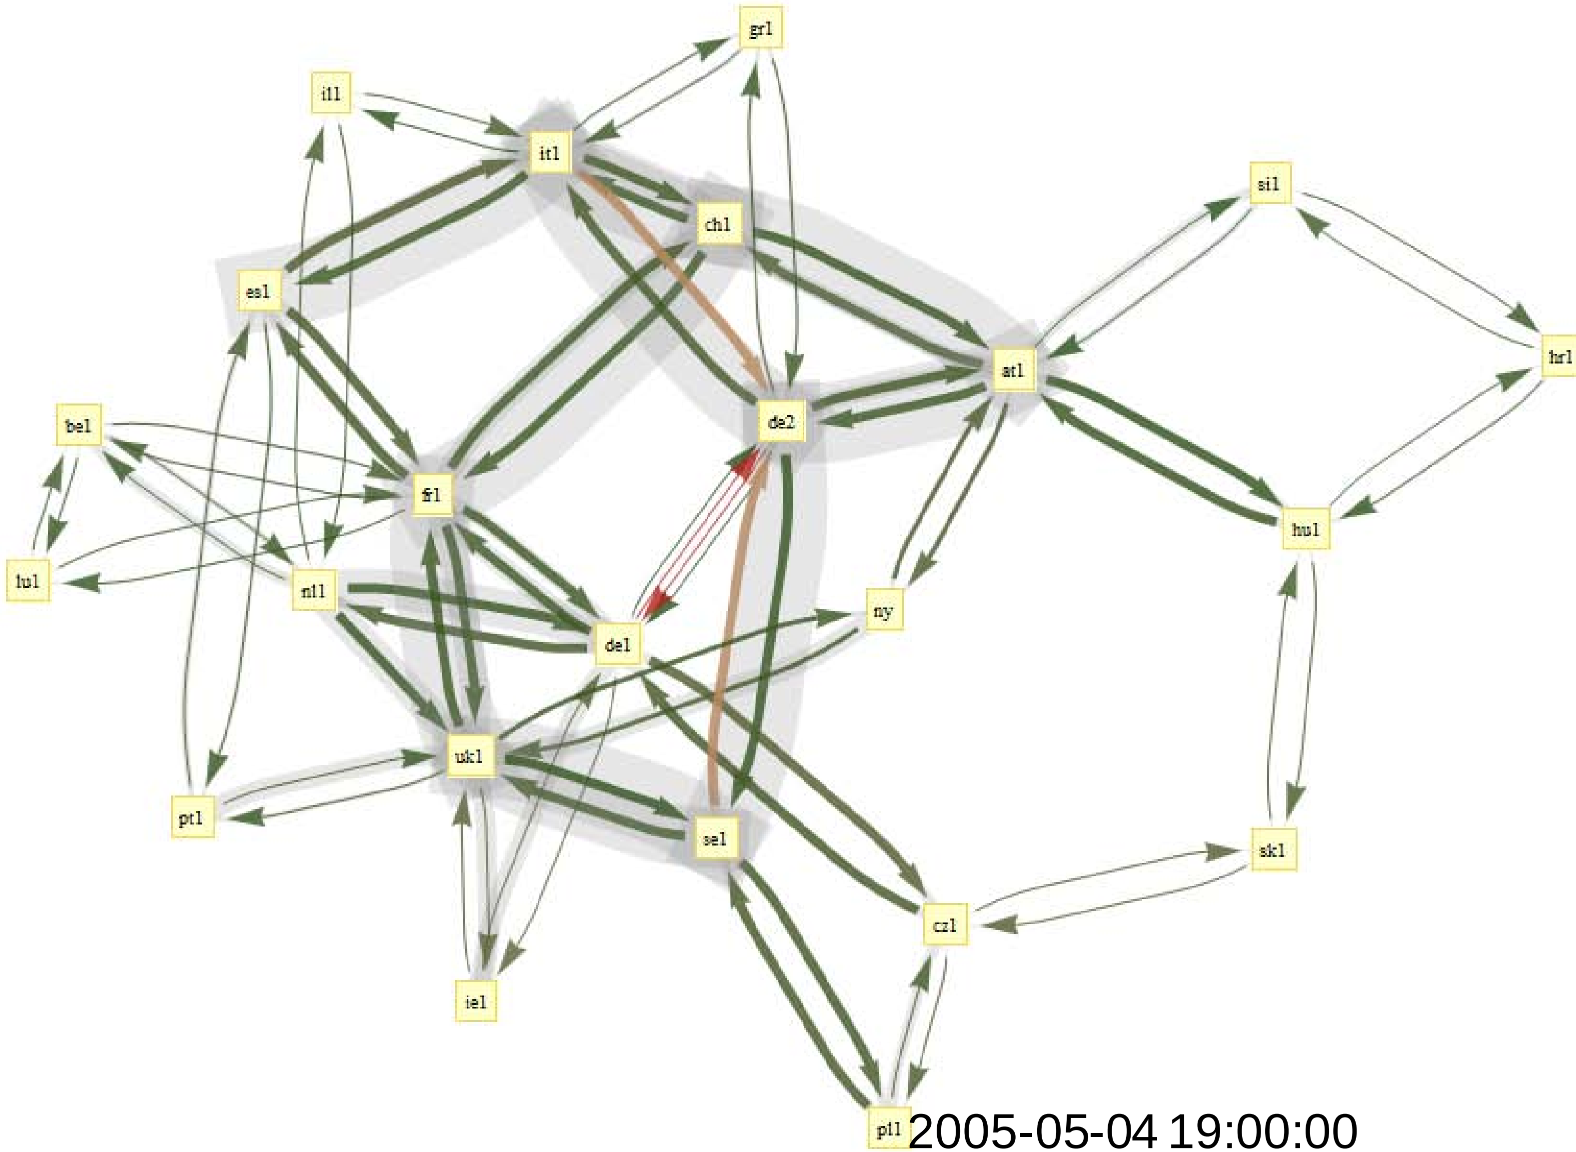

2005-05-04 19:00:00
Colours show utilization
Grey shows effective pooled capacity

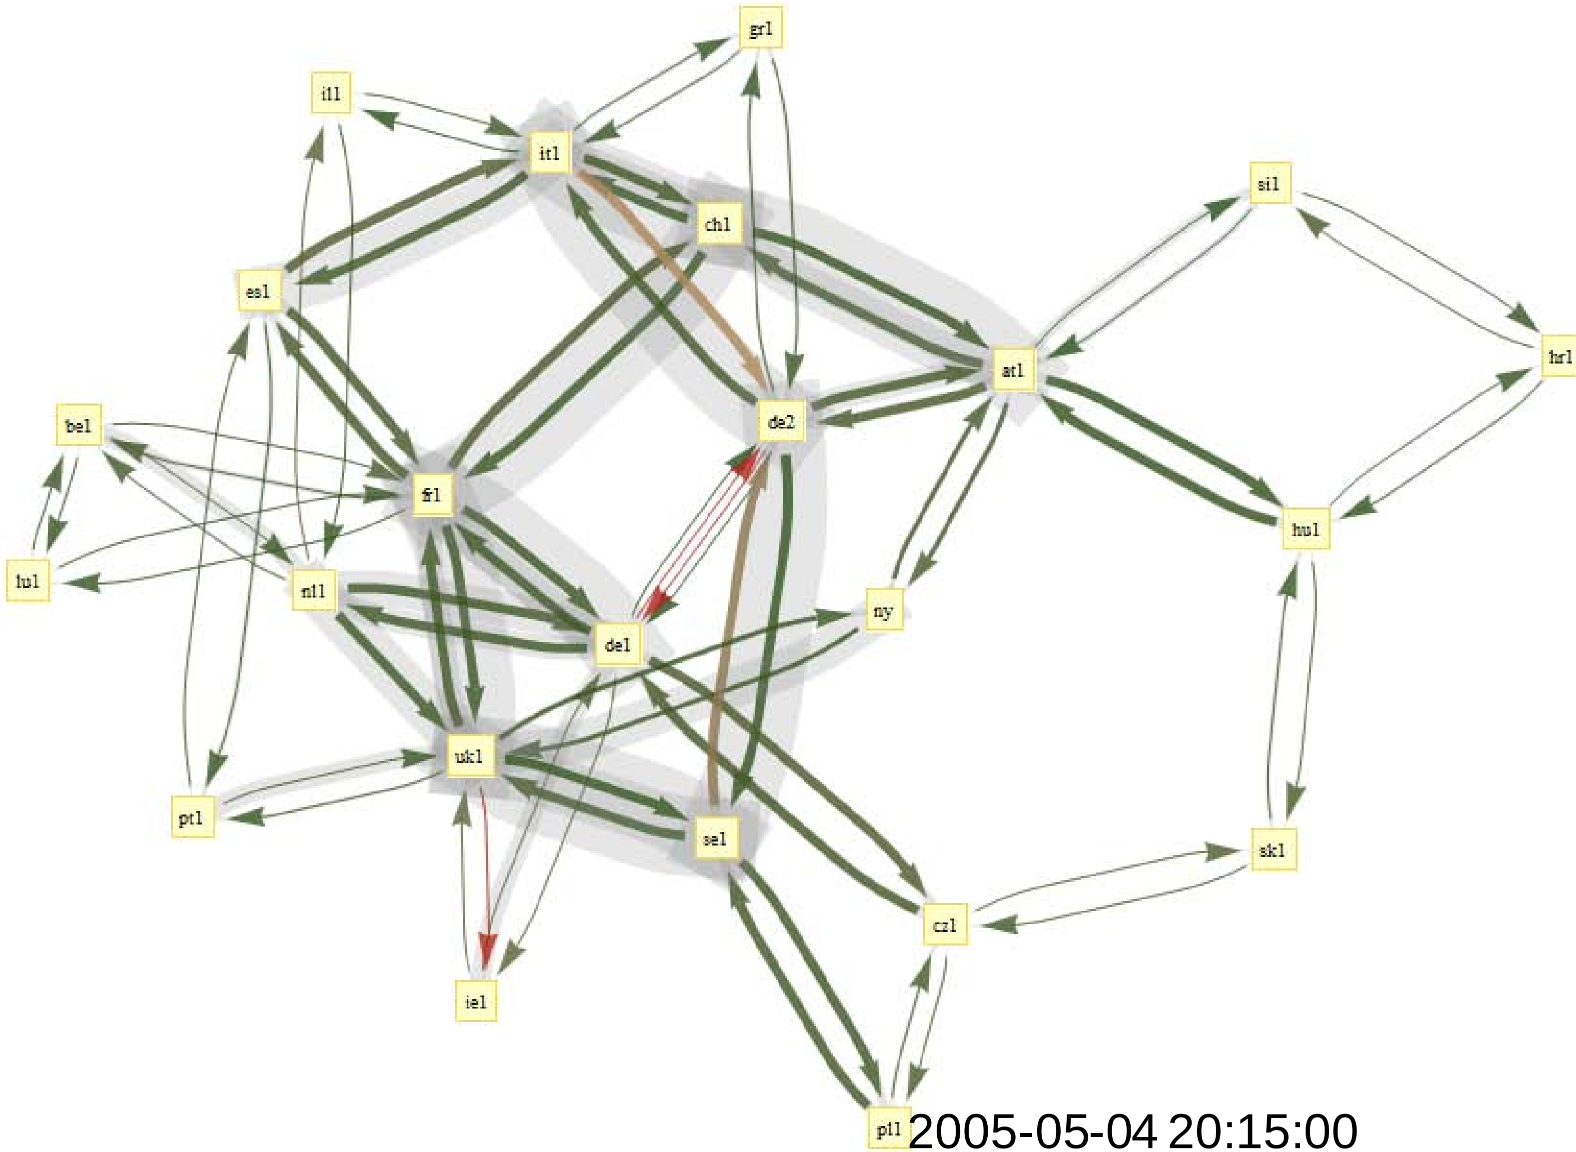

2005-05-04 20:15:00
Colours show utilization
Grey shows effective pooled capacity