Communications Network Design

lecture 10

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

March 2, 2009

Concave costs

When costs are concave, the network design problem has properties like single path routing. A common example is linear costs. Also we present a simple heursitic approach.

Notation recap

Mostly as before (lecture 6)

- A **network** is a graph $G(N,E)$, with **nodes** $N = \{1,2,\ldots n\}$ and **links** $E \subseteq N \times N$
- Offered traffic between O-D pair (p,q) is t_{pq}
- The set of all **paths** in $G(N,E)$ is $P = \cup_{[p,q]\in K} P_{pq}$
- Each link $e \in E$ has
  - a **capacity**, denoted by $r_e (\geq 0)$
  - a **distance** $d_e (\geq 0)$
  - a **load** $f_e (\geq 0)$
- The vector $\mathbf{x} = (x_\mu : \mu \in P)$ is called the **routing**

$$f_e = \sum_{\mu \in P : e \in \mu} x_\mu$$

Multicommodity flow problem

Completely general case

- objective: minimize some cost function
  - construction costs based on capacities r_e
  - performance costs (e.g. delays, reliability, ...) based on r_e and f_e
- input:
  - a set of nodes N
  - forecast traffic demands t_{pq}
- constraints are flow based (as before)
  - loads on links are implied by routing of traffic
  - link loads $\leq$ capacities

Call it the **multicommodity flow problem**

A simplified problem

- in general costs depend on r_e and f_e
  - lets us start a little simpler
  - only include construction costs
    - not performance costs
- assume we choose $r_e = f_e$
  - choose capacities to carry required loads
    - could include some overhead, e.g. $r_e = \gamma f_e$ for some $\gamma > 1$
- problem simplifies to choosing which links we need in our network
  - it becomes an integer programming problem
  - it has a direct relationship to least-cost routing on a complete graph

Formal problem specification

Formal problem specification:

$$
\text{(P)} \quad \text{min.} \quad C(\mathbf{f}) = \sum_{e \in E} c_e(f_e)
$$

$$
\begin{aligned}
\text{s.t.} \quad f_e &= \sum_{\mu \in P : e \in \mu} x_\mu && \forall e \in E. \\
x_\mu &\geq 0 && \forall \mu \in P \\
\sum_{\mu \in P_{pq}} x_\mu &= t_{pq} && \forall [p,q] \in K.
\end{aligned}
$$

Where we then take $r_e = \gamma f_e, \; \forall e \in E$

This looks the same as for routing, but the set E is the set of **all possible** links, rather than a given set, and the cost function C will be different (though still separable).

Typical cost function

- assume the cost function is continuous on $[0, \infty)$ and differentiable on $(0, \infty)$
- assume the cost function nondecreasing
- assume the cost function is separable

$$C(\mathbf{f}) = \sum_e c_e(f_e)$$

- assume the cost function is concave
  $C(\mathbf{f})$ is concave over Ω if for all $\lambda \in [0, 1]$, and all feasible loads $\mathbf{f}_1, \mathbf{f}_2 \in \Omega$,

$$C(\lambda \mathbf{f}_1 + (1-\lambda)\mathbf{f}_2) \geq \lambda C(\mathbf{f}_1) + (1-\lambda)C(\mathbf{f}_2)$$

  - chords lie below the function

Concave

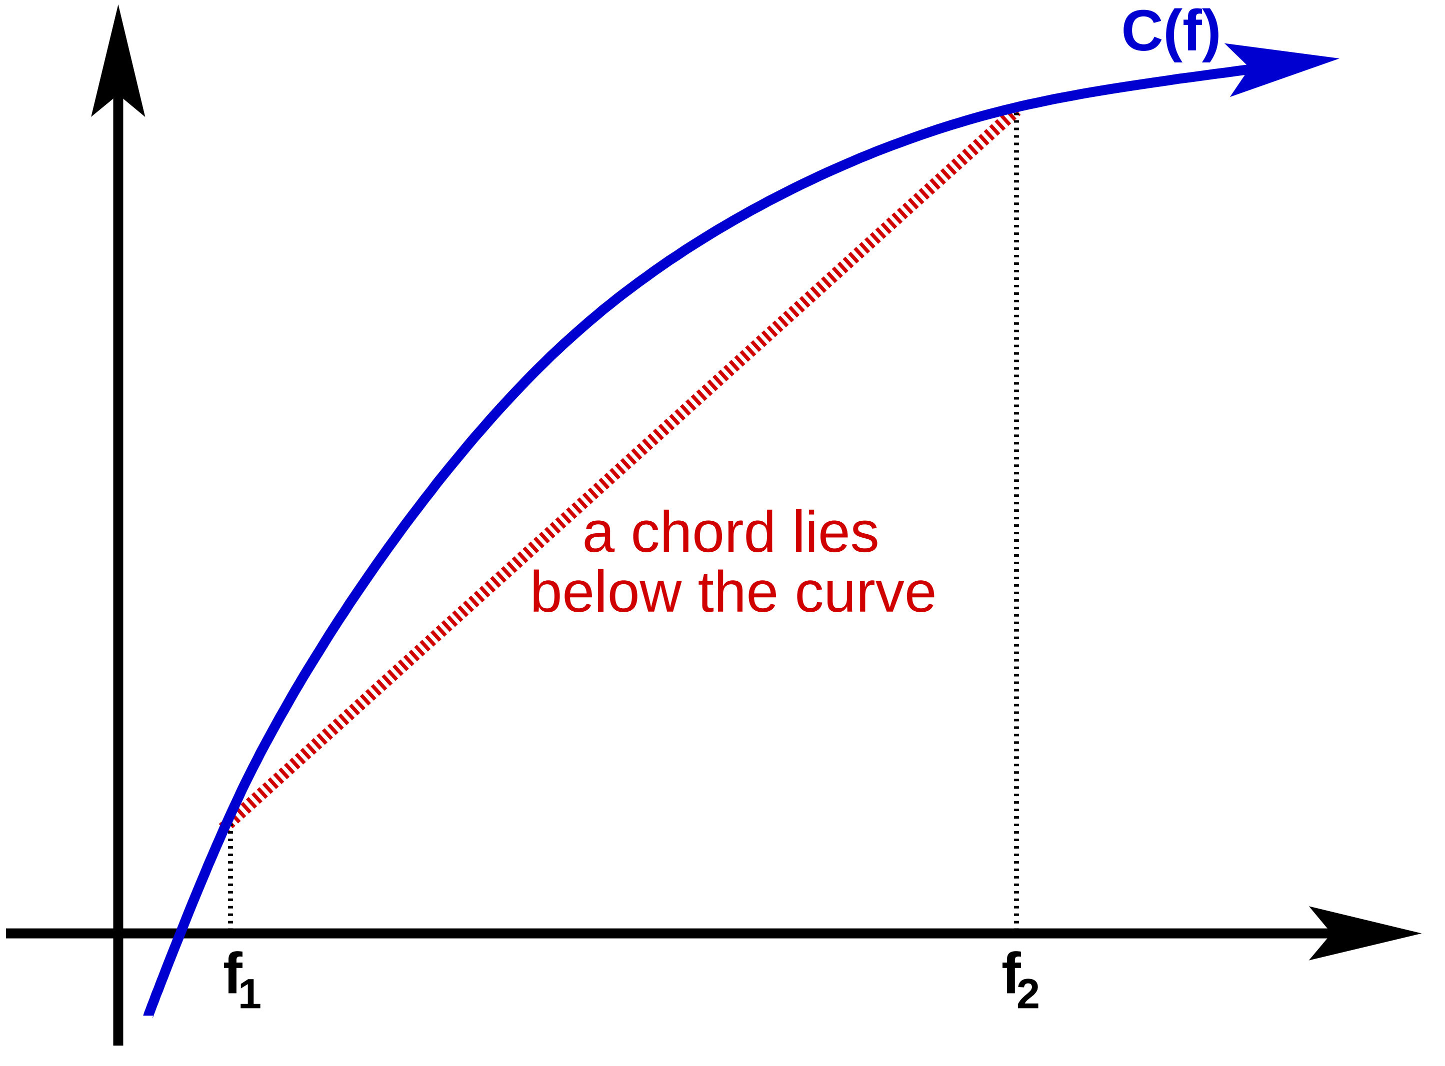

Concave costs

- concave costs represent "economy of scales"
  - operations at a larger scale have a smaller marginal cost, e.g. $\frac{\partial c_e}{\partial f_e}$ is decreasing
  - operations at a larger scale have a smaller average cost, e.g. $\frac{c_e(f_e)}{f_e}$ is decreasing
- alternative view "multiplexing gain"
  - multiplexed (grouped) traffic has a lower relative variance, and so is less "bursty"
  - less overhead is required for smoother traffic
- Example

$$c_e(f_e) = k_e f_e^{\alpha}, \quad k_e = \text{constant}, \, \alpha \in (0.4, 0.6)$$

Concave costs and routing

- we have so far (today) ignored routing
- nice result that shows for concave costs, we only need to consider single path routing (no load sharing)

Proposition: If $C(\mathbf{f})$ is a concave cost function of load $\mathbf{f}$, then the minimum is attained by routing t_{pq} on a single path $\hat{\mu}_{pq}$ for all O-D pairs $[p,q] \in K$.

Example of single path routing

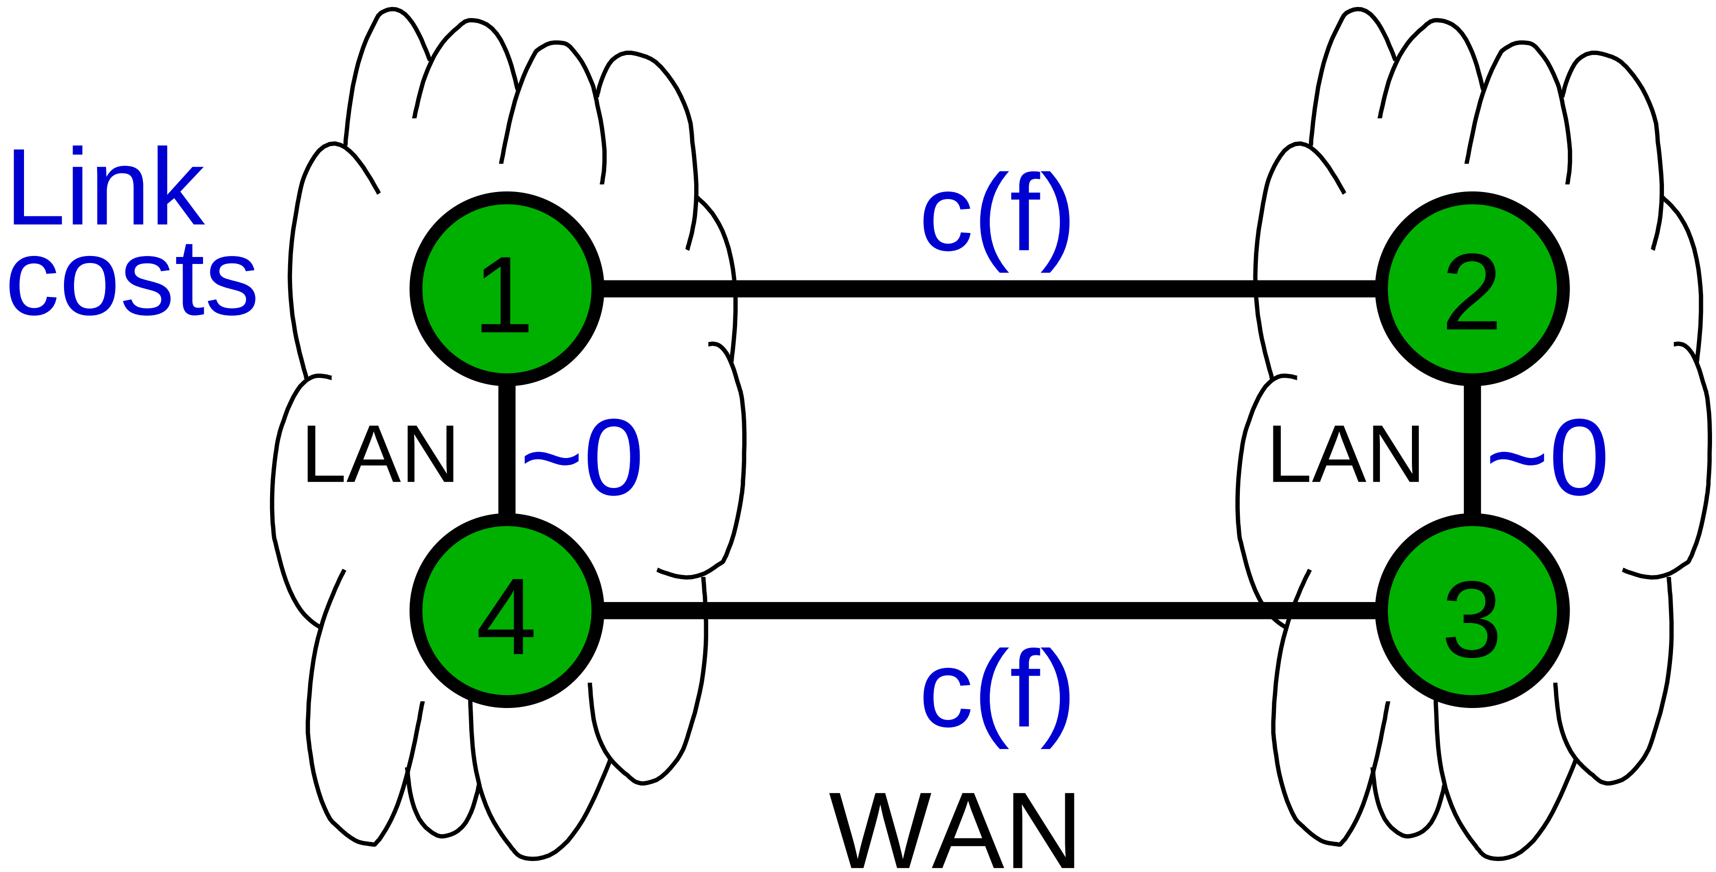

Example of single path routing

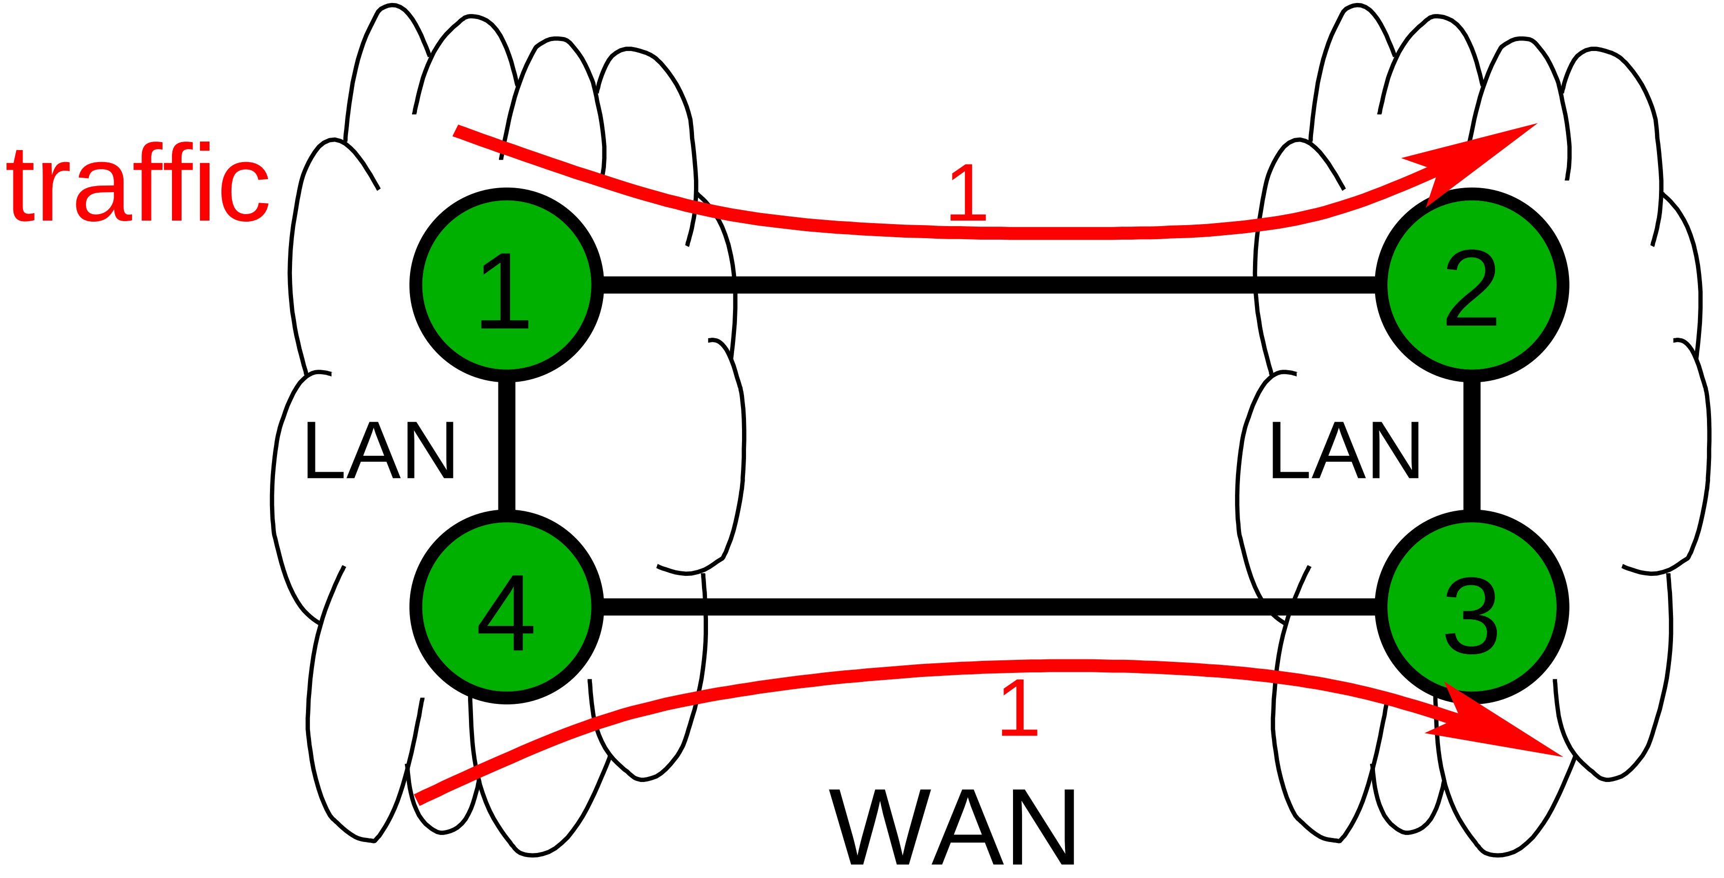

Example of single path routing

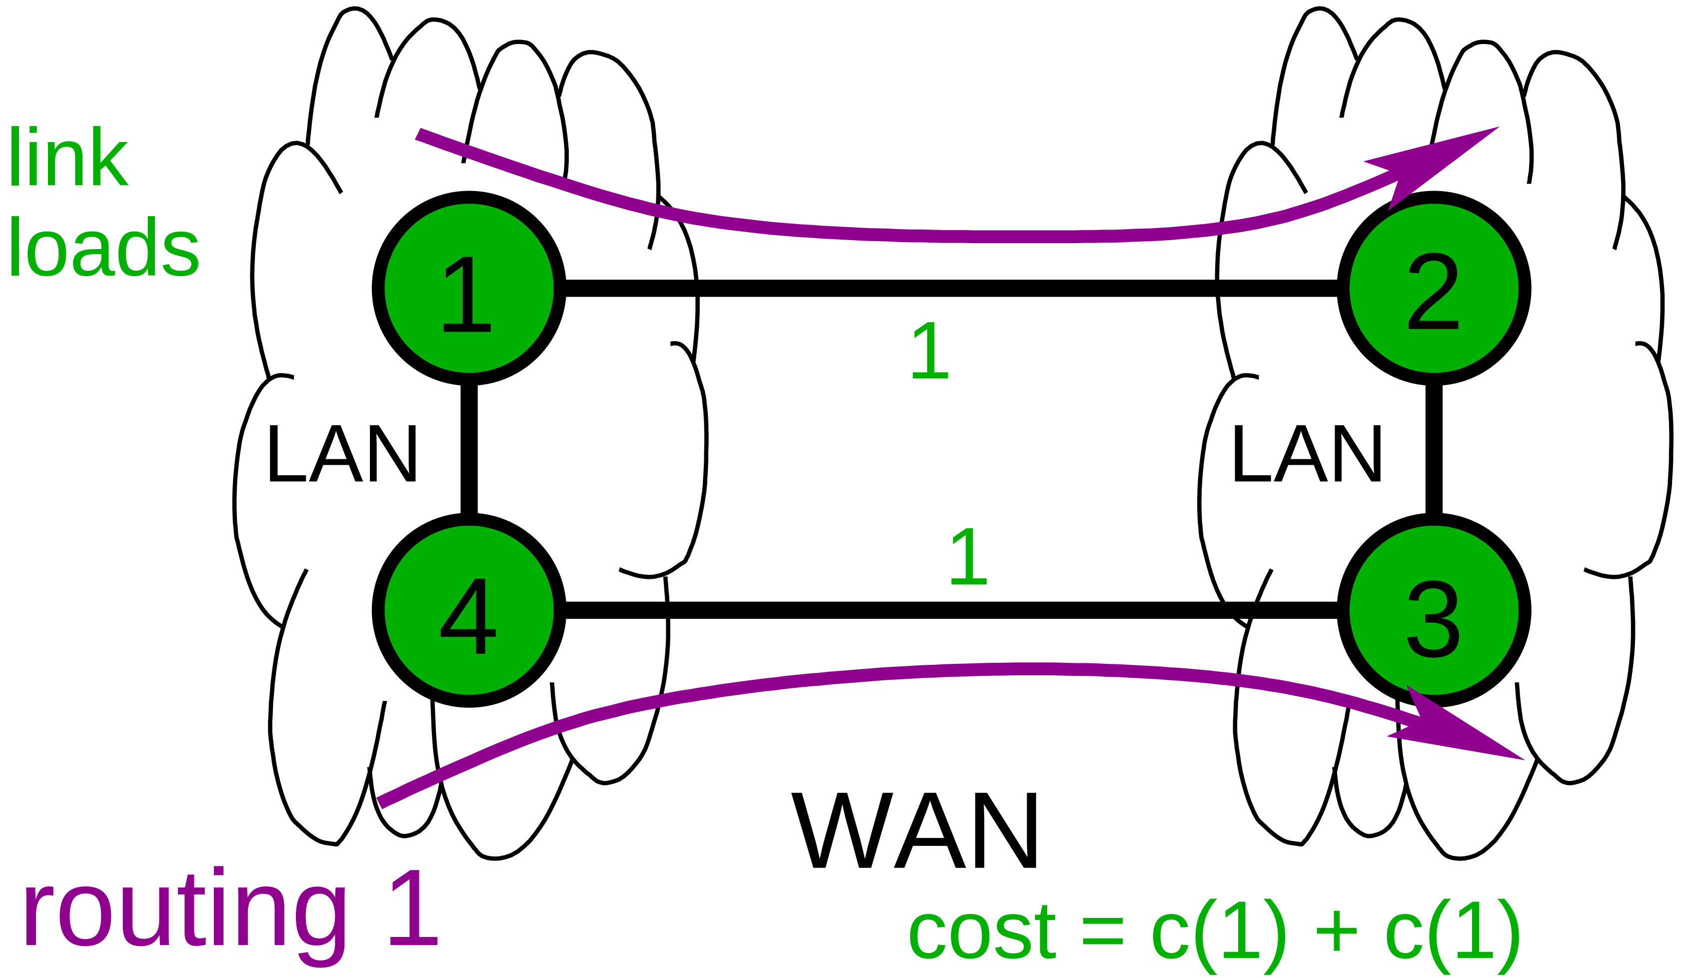

Example of single path routing

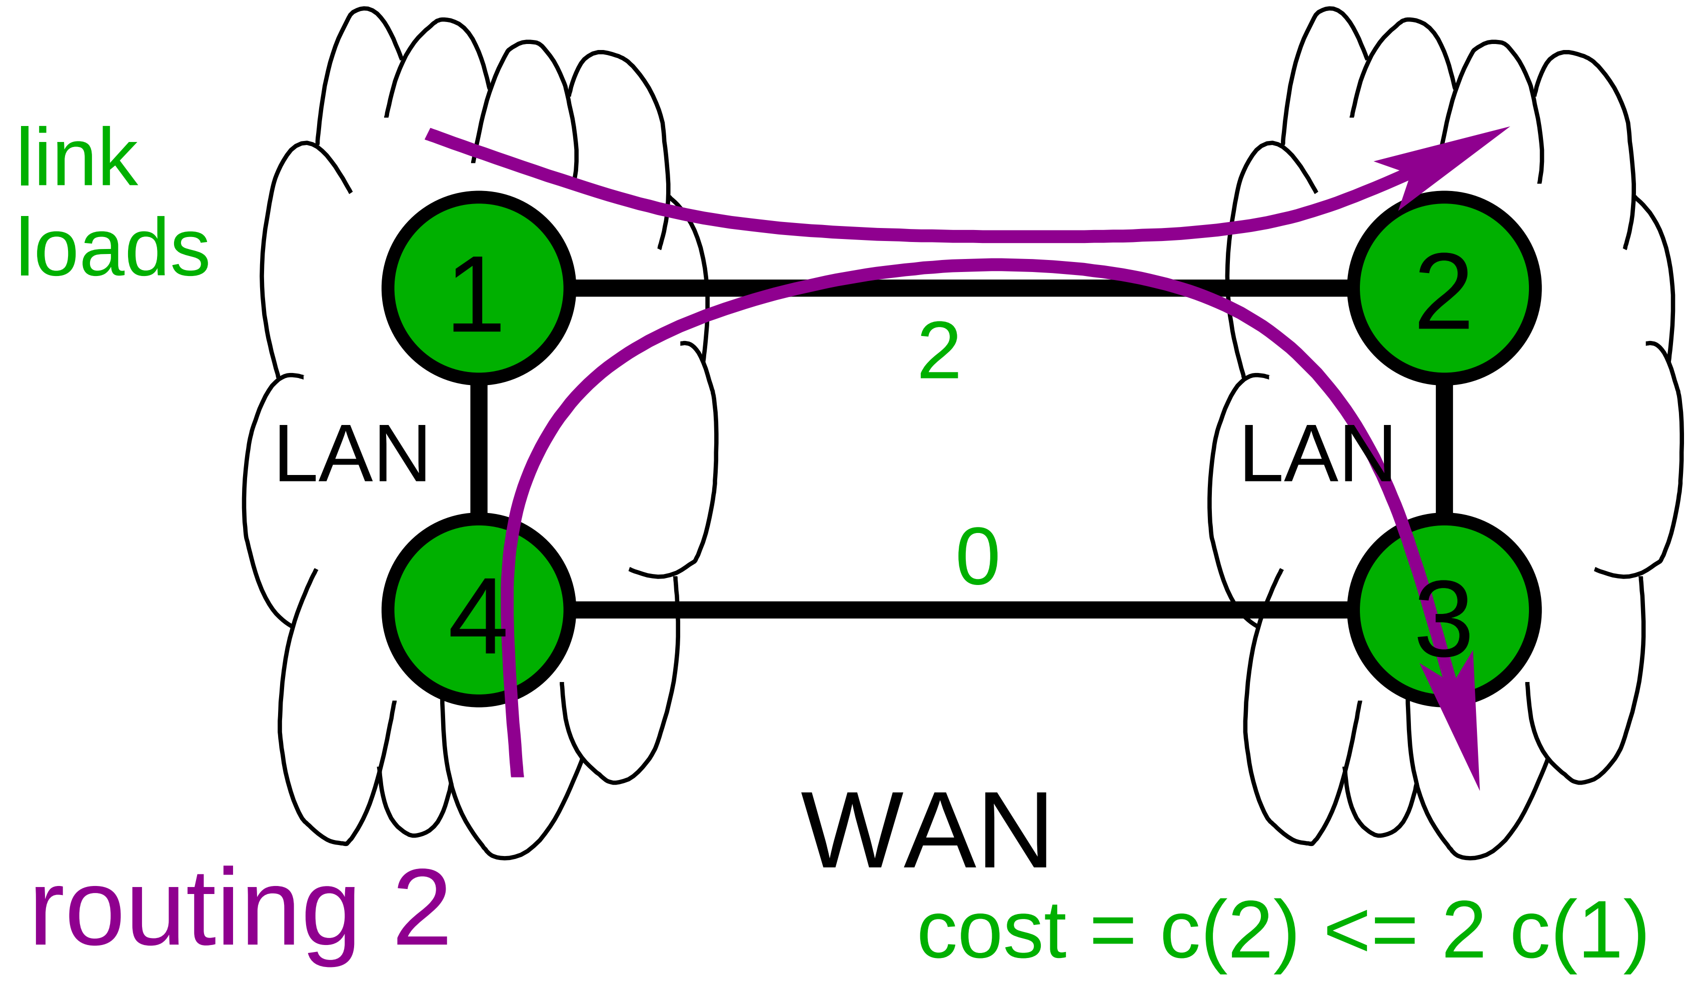

Concave costs and routing: proof

Proof: Let us take two paths $\mu_1, \mu_2 \in P_{pq}$. Suppose there is a routing $\mathbf{x} = (x_\mu : \mu \in P)$ such that the traffic between the O-D pair $[p,q]$ is routed across both μ_1 and μ_2, i.e. $x_{\mu_1} > 0$ and $x_{\mu_2} > 0$.

Let $\mathbf{f}$ be the link loads induced by $\mathbf{x}$; so

$$f_e = \sum_{\mu : e \in \mu} x_\mu$$

Consider two cases:

- the traffic x_{μ_2} on μ_2 is moved to μ_1 inducing loads $\mathbf{f}^{(1)}$
- the traffic x_{μ_1} on μ_1 is moved to μ_2 inducing loads $\mathbf{f}^{(2)}$

Concave costs and routing: proof

The net result is:

$$\mathbf{f} = \frac{x_{\mu_1}\mathbf{f}^{(1)} + x_{\mu_2}\mathbf{f}^{(2)}}{x_{\mu_1} + x_{\mu_2}} = \frac{x_{\mu_1}}{x_{\mu_1} + x_{\mu_2}}\mathbf{f}^{(1)} + \frac{x_{\mu_2}}{x_{\mu_1} + x_{\mu_2}}\mathbf{f}^{(2)} \qquad (1)$$

and therefore, for all $e \in E$,

$$f_e = \frac{x_{\mu_1} f_e^{(1)} + x_{\mu_2} f_e^{(2)}}{x_{\mu_1} + x_{\mu_2}} \qquad (2)$$

- In both cases links $e \notin \mu_1, \mu_2$ and links $e \in \mu_1, \mu_2$ have load unaltered, e.g. $f_e^{(1)} = f_e^{(2)} = f_e$.
- Only those links on precisely one of the paths μ_1, μ_2 have loads altered by this process.

Concave costs and routing: proof

In more detail: check this out for links $e \in E$:

- if $e \in \mu_1$ and $e \in \mu_2$ then $f_e^{(1)} = f_e^{(2)} = f_e$ so equation (2) correctly gives the load as f_e.
- if $e \notin \mu_1$ and $e \notin \mu_2$ then $f_e^{(1)} = f_e^{(2)} = f_e$, and (2) is OK.
- if $e \in \mu_1$ but $\notin \mu_2$ then $f_e^{(1)} = f_e + x_{\mu_2}$ and $f_e^{(2)} = f_e - x_{\mu_1}$. So RHS of (2) above gives
$$\frac{x_{\mu_1} f_e^{(1)} + x_{\mu_2} f_e^{(2)}}{x_{\mu_1} + x_{\mu_2}} = \frac{x_{\mu_1}(f_e + x_{\mu_2}) + x_{\mu_2}(f_e - x_{\mu_1})}{x_{\mu_1} + x_{\mu_2}} = f_e$$
- if $e \notin \mu_1$ but $\in \mu_2$ then $f_e^{(1)} = f_e - x_{\mu_2}$ and $f_e^{(2)} = f_e + x_{\mu_1}$. So RHS of (2) above gives
$$\frac{x_{\mu_1} f_e^{(1)} + x_{\mu_2} f_e^{(2)}}{x_{\mu_1} + x_{\mu_2}} = \frac{x_{\mu_1}(f_e - x_{\mu_2}) + x_{\mu_2}(f_e + x_{\mu_1})}{x_{\mu_1} + x_{\mu_2}} = f_e$$

Concave costs and routing: proof

Take $\lambda = \frac{x_{\mu_1}}{x_{\mu_1} + x_{\mu_2}} \in (0,1)$ and $1 - \lambda = \frac{x_{\mu_2}}{x_{\mu_1} + x_{\mu_2}} \in (0,1)$.
When C is concave. By definition, for all $\lambda \in [0,1]$,

$$C(\lambda \mathbf{f}_1 + (1-\lambda)\mathbf{f}_2) \geq \lambda C(\mathbf{f}_1) + (1-\lambda)C(\mathbf{f}_2)$$

Given that $\mathbf{f} = \lambda \mathbf{f}^{(1)} + (1-\lambda)\mathbf{f}^{(2)}$ we get

$$C(\mathbf{f}) \geq \lambda C(\mathbf{f}^{(1)}) + (1-\lambda)C(\mathbf{f}^{(2)})$$

If $C(\mathbf{f}^{(1)}) \leq C(\mathbf{f}^{(2)})$, then $\lambda C(\mathbf{f}^{(1)}) + (1-\lambda)C(\mathbf{f}^{(2)}) \geq C(\mathbf{f}^{(1)})$ and therefore, $C(\mathbf{f}) \geq C(\mathbf{f}^{(1)})$. This means the traffic can t_{pq} can all be re-routed onto μ_1 with less cost.

If $C(\mathbf{f}^{(1)}) \geq C(\mathbf{f}^{(2)})$ then, re-route traffic t_{pq} onto μ_2. $\square$

Concave costs and routing

- The result above means that with concave costs
  - we can assume that single paths are used for end-to-end demands.
- Heuristic for network design
  - adapt the Frank-Wolfe method
    - remember this was used for routing with **convex** costs
  - assumptions
    - we start with a single path routing $\mathbf{x}$
    - the corresponding induced load is $\mathbf{f}$
    - the routing is **not** a shortest path routing

Heuristic Method

- If all traffic is allocated to a shortest path, STOP.
- Else, select for all $k \in K$, a shortest length path $\hat{\mu}_k$ of length $l_{\hat{\mu}_k}$.
- Allocate t_k to its shortest path $\hat{\mu}_k$ for all $k \in K$.
- Call this routing $\mathbf{z}$.
- Re-calculate shortest paths; go to first step.

Note we have concave cost, so there is no guarantee that the shortest path routing we find will be the minimal cost routing!

The point

- Routing and capacity are intricately linked
- We can solve the capacity problem (for the cases above) by solving the routing problem on a **complete graph**.
- Any link with zero traffic is eliminated
- other links have capacities designed to carry traffic plus some overhead.
- Different types of cost
  - routing $\Rightarrow$ convex costs $\Rightarrow$ SPF
  - construction $\Rightarrow$ concave costs $\Rightarrow$ unique routing
- special case: **linear costs**
  - best of both cases: unique SPF routing

Example with linear costs

Link costs α_e

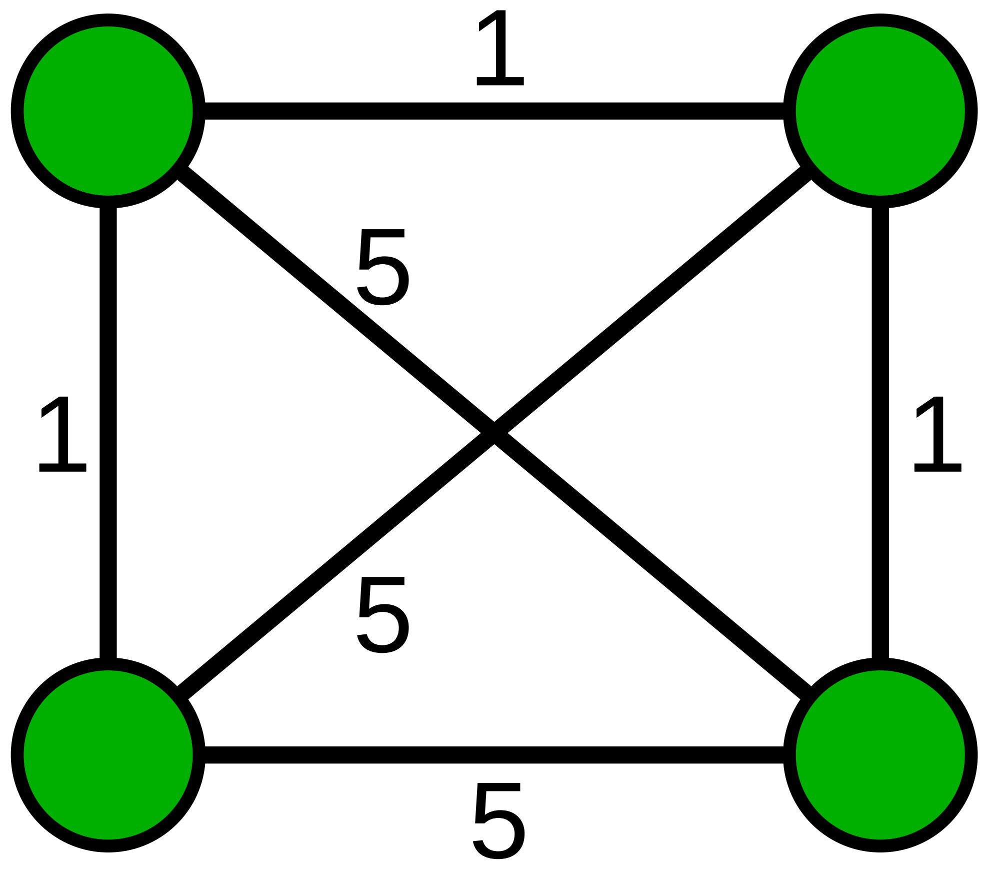

Example with linear costs

SPF tree

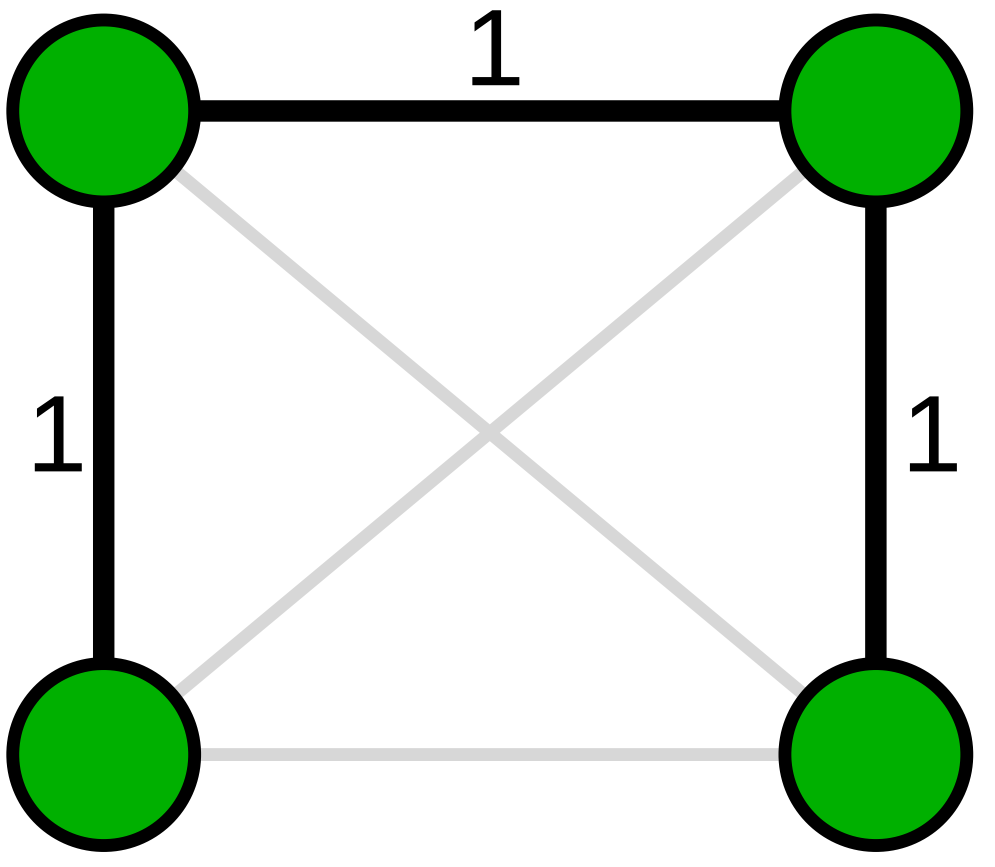

Example with linear costs

O–D	paths
1–2	1–2
1–3	1–3
1–4	1–2–4
2–3	2–1–3
2–4	2–4
3–4	3–1–2–4

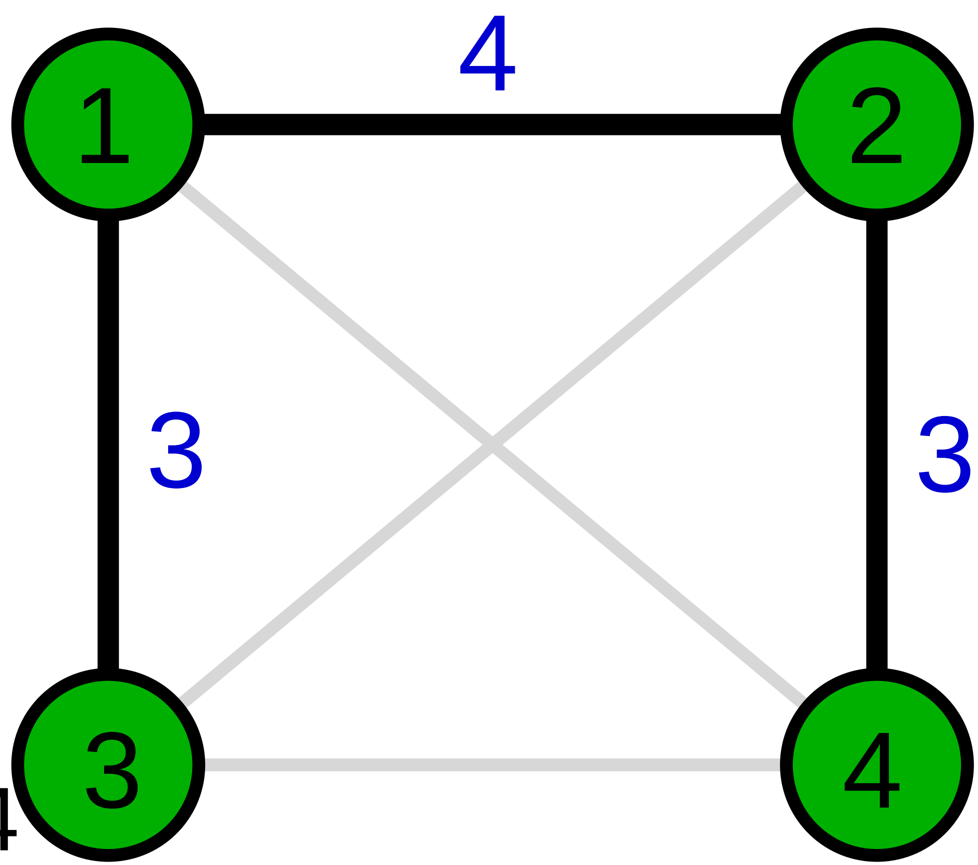

Example with linear costs

Link loads f_e

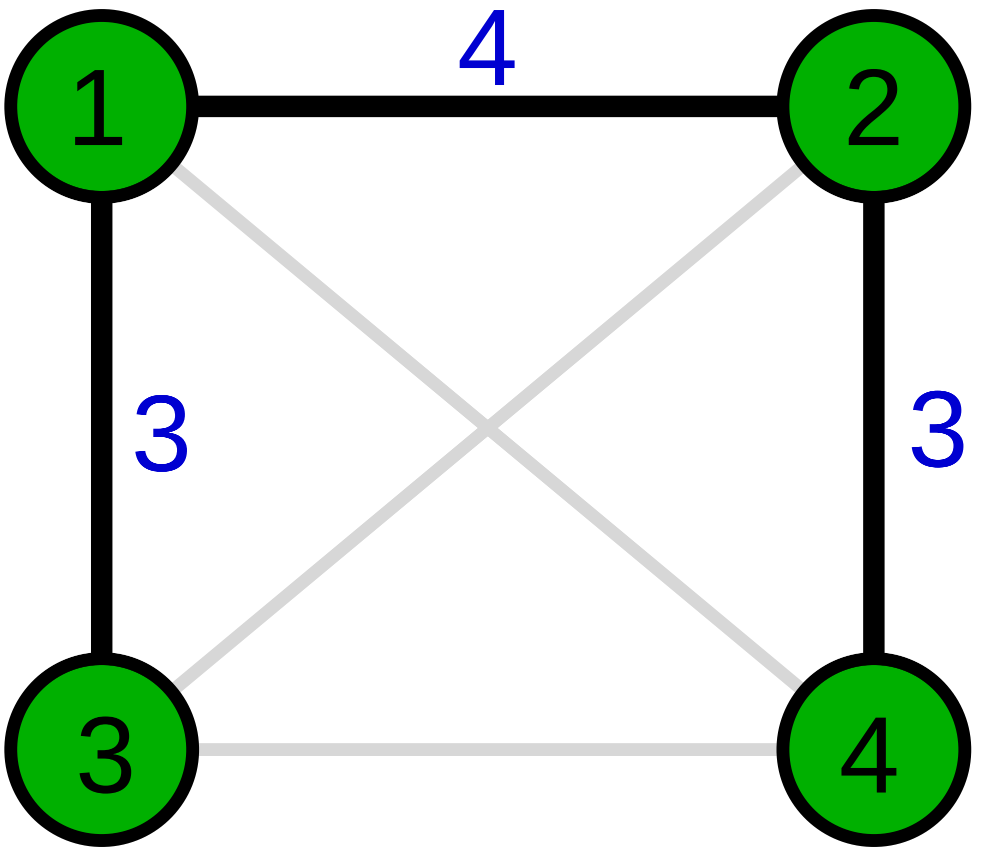

References