Routing on Multiple Optimality Criteria

João L. Sobrinho

Miguel A. Ferreira

Instituto de Telecomunicações,
Instituto Superior Técnico, Universidade de Lisboa

October 2020

Optimal Path Routing

- Performance metrics
  - ➔ Examples: any combination of cost, delay, capacity or available bandwidth. In wireless communications, loss probability and co-channel interference may be taken into account.
- Optimality criterion
  - ➔ Relative valuation of a performance metric.
- Optimal path
  - ➔ Most valuable performance metric among those of all paths from source to destination.
  - ➔ Examples: quickest paths [Chen & Chin 1990]; widest paths with co-channel interference [Draves 2004];

New Results

- Optimal path routing for an arbitrary optimality criterion
  - ➔ Standard routing protocols (EIGRP, BGP, DSDV, and OSPF) do not route on optimal paths for all but a very restricted class of optimality criteria.
  - ➔ Examples: Standard routing protocols do not route neither on quickest paths, nor on widest-paths with co-channel interference.

- Optimal path routing for multiple optimality criteria
  - ➔ Goal: for each flow of data-packets, route data-packets on an optimal path for the most appropriate optimality criterion for that particular flow.

Table of Contents

Table of Contents

Optimal Paths

- Set *S* of attributes
  - ➔ Represents a performance metric.
- Binary extension operation on *S*, $\oplus$
- Total order on *S*, $\leqslant_i$
  - ➔ Defines an optimality criterion *i*.

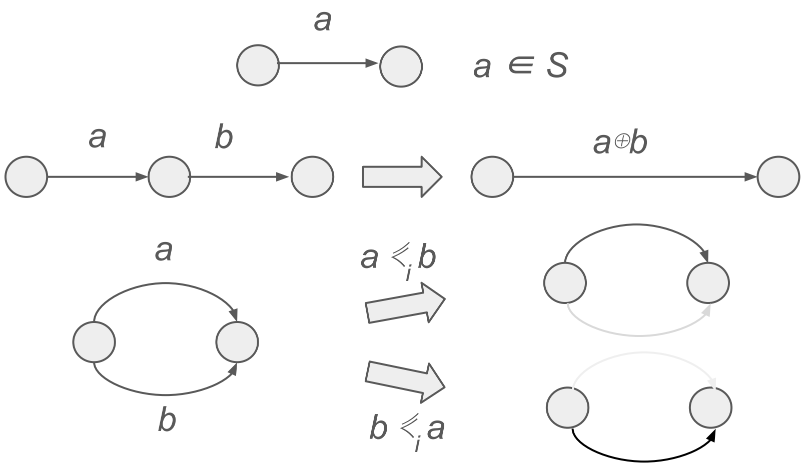

- Optimal attribute of set *A*:
  - ➔ Attribute that is preferred to every other attribute of *A*.

- Optimal attribute from source *s* to destination *t*:
  - ➔ Optimal attribute of the set of path attributes from *s* to *t*.

Quickest Paths

- Optimality criterion: minimize the transfer time of a file of size *K* along a path with capacity *c* and delay *d*, *K/c + d.*

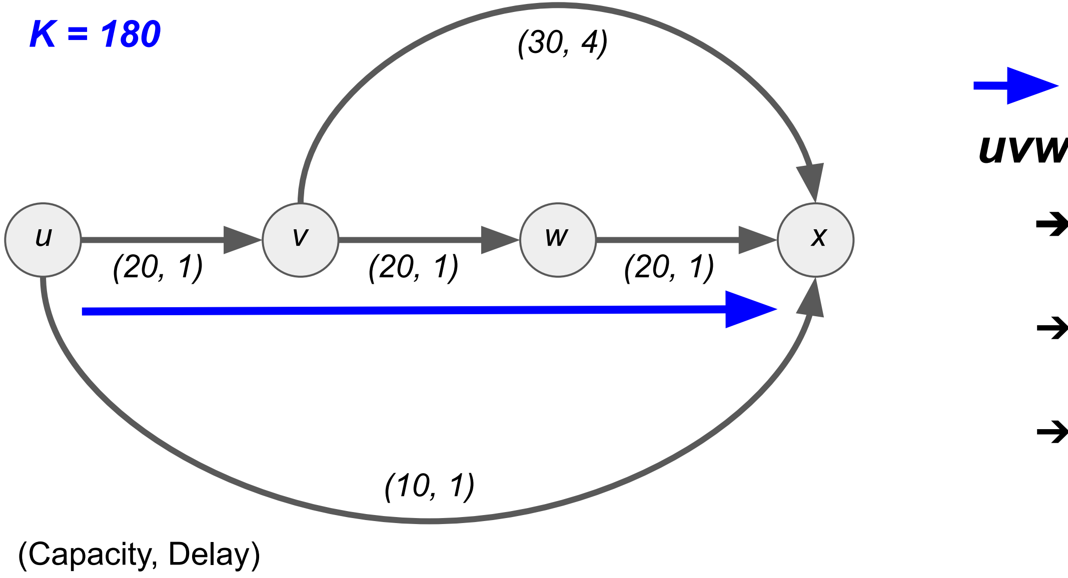

K-quickest path from *u* to *x* is ***uvwx***:

- **Transfer time along *uvwx, (20, 3)* is *12 = 180/20 + 3*.**
- Transfer time along *uvx, (20, 5)*, is *14 = 180/20 + 5*
- *Transfer time* along *ux, (10, 1)*, is *19 = 180/10 + 1*.

Quickest Order

- Cartesian product between capacities and delay:
  - ➔ $\{(c, d) \mid c \in \mathbb{R}^+ \cup \{\infty\}, d \in \mathbb{R}_0^+\} \cup \{(0, \infty)\}$.
- Product extension:
  - ➔ *(c, d)* extended with *(c', d')* is *(min{c, c'}, d+d')*.
- *K*-quickest order:
  - ➔ *(c, d)* is prefered to *(c', d')* if:
    $K/c + d < K/c' + d$, or $K/c + d = K/c' + d$ and $c > c'$.

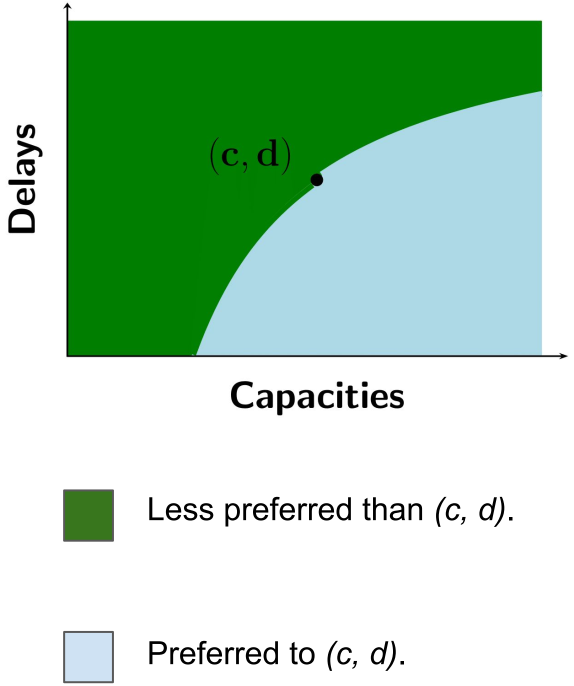

Standard Vectoring Protocols (EIGRP, BGP, DSDV)

- Separate path attribute computation per destination.

- When a node receives an **attribute** from an out-neighbor:
  - ➔ Extension of the link attribute to the out-neighbor with the received attribute originates a candidate attribute. ($\oplus$)
  - ➔ Election of the **optimal attribute** of the set of all candidate attributes ($\preccurlyeq_i$); advertisement of the elected attribute to in-neighbors.

- Destination-based forwarding:
  - ➔ Node forwards data-packets to out-neighbor from which elected attribute was learned.

Optimal path routing for arbitrary optimality criterion?

Routing on Quickest Paths

***K* = 180**

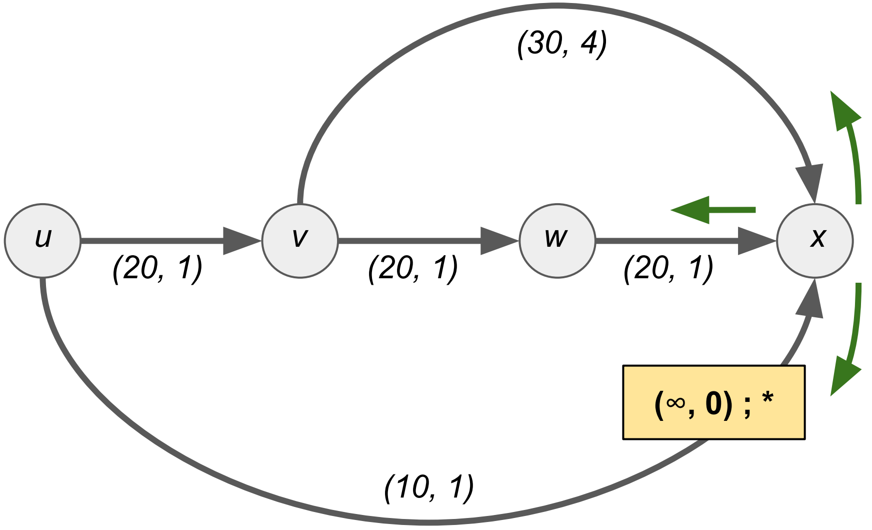

Elected pair ; out-neighbor
Candidate pair ; out-neighbor

Routing on Quickest Paths

_K = 180_

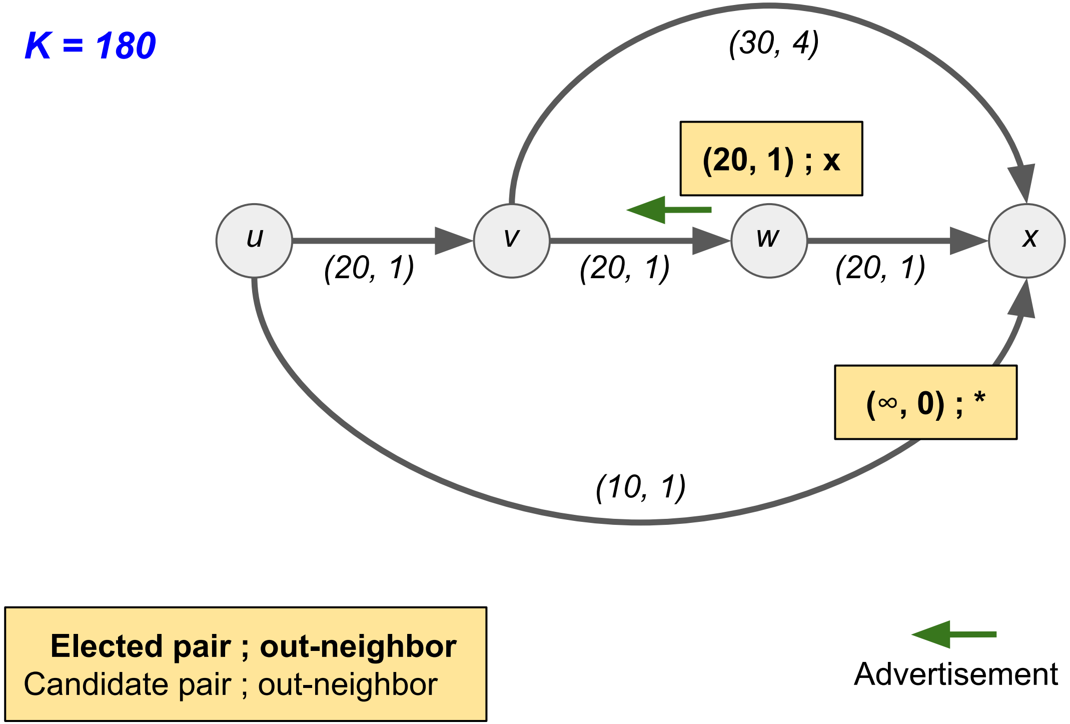

- Node *w*:

➔ *(20, 1)* is the extension of *(20, 1)* with *(∞, 0)*.

Routing on Quickest Paths

K = 180

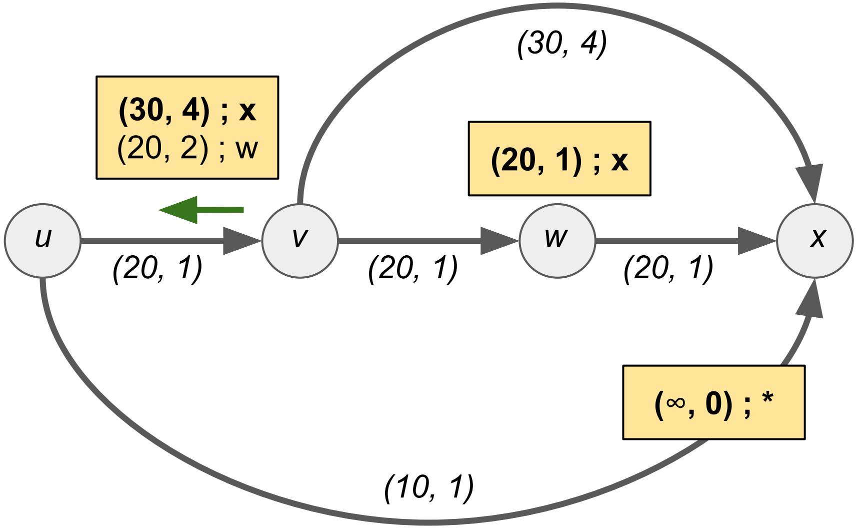

- Node *v*:

➔ *(30, 4)* is the extension of *(30, 4)* with *(∞, 0)*.

➔ *(20, 2)* is the extension of *(20, 1)* with *(20, 1)*.

➔ *(30, 4)* is preferred to *(20, 2)*: *180/30 + 4 = 10 < 11 = 180/20 + 2*.

Elected pair ; out-neighbor
Candidate pair ; out-neighbor

Advertisement

Routing on Quickest Paths

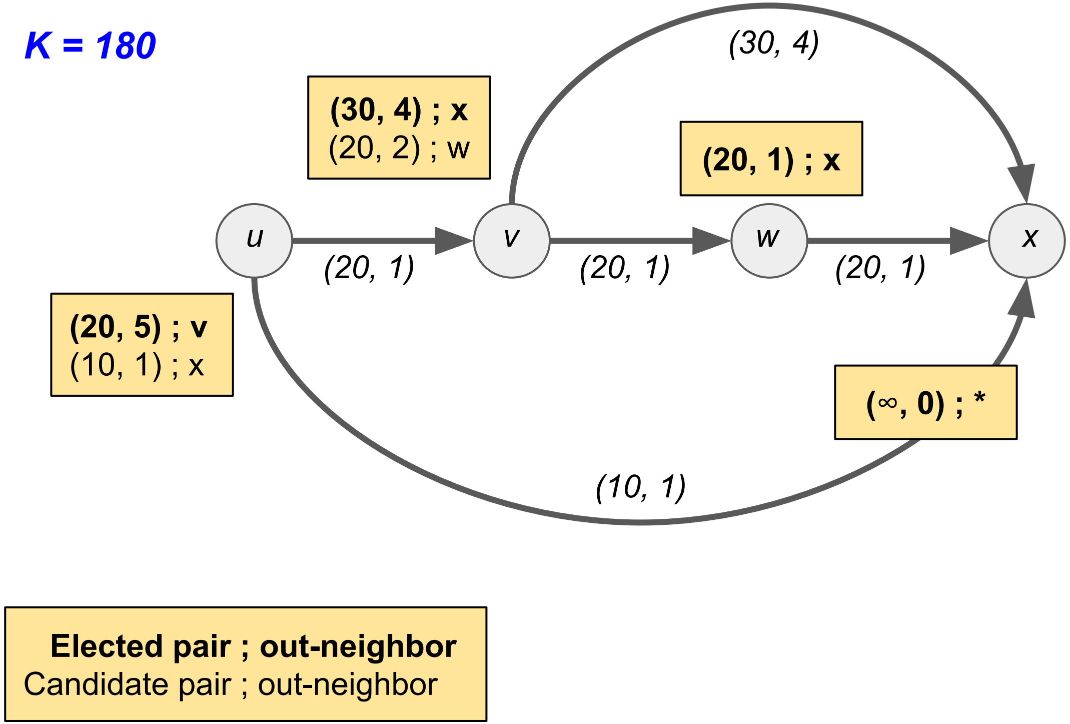

- Node *u*:

➔ *(20, 5)* is the extension of *(20, 1)* with *(30, 4)*.

➔ *(10, 1)* is the extension of *(10, 1)* with *(∞, 0)*.

➔ *(20, 5)* is preferred to *(10, 1)*: $180/20 + 5 = 14 < 19 = 180/10 + 1$

Routing on Quickest Paths

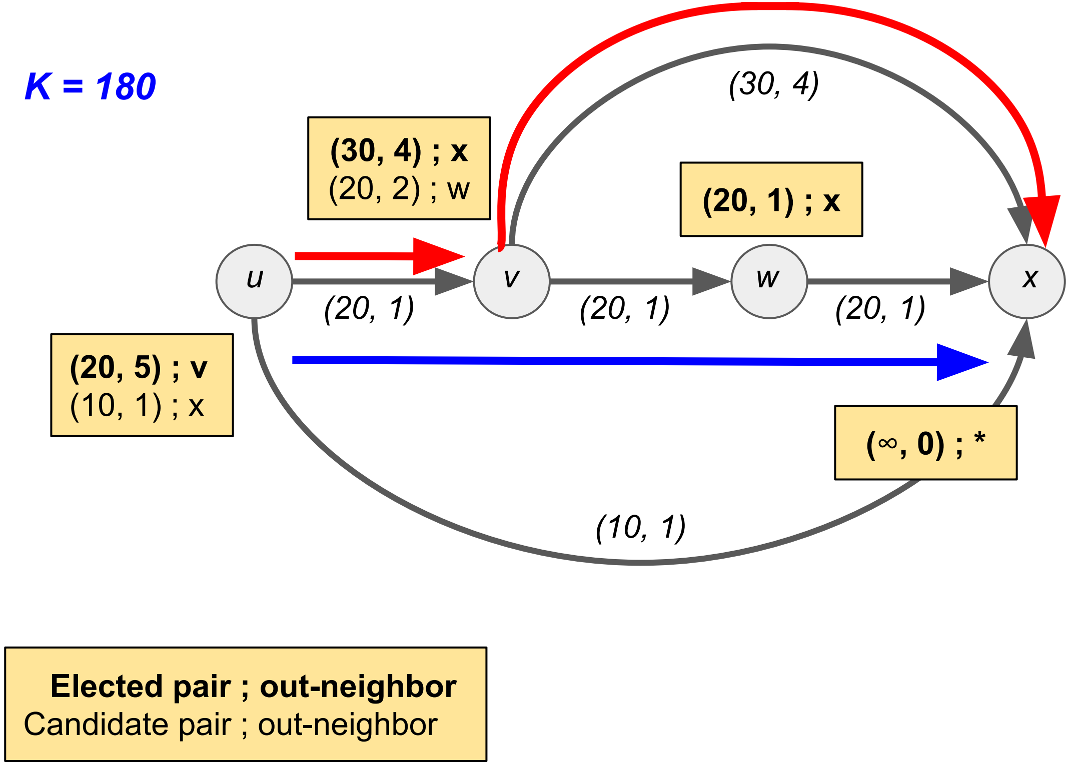

Data-packets traverse path ***uvx***.

K-quickest path from *u* to *x* is ***uvwx***.

Isotonicity and Optimality

Extension preserves the relative preference among attributes.

$\oplus$ is **isotone** for $\preceq_i$ if:

$b \preceq_i c$ implies $a \oplus b \preceq_i a \oplus c$, for all attributes a, b, c.

Isotonicity $\Rightarrow$ Optimal path routing for arbitrary optimality criterion

Table of Contents

Partial Orders and Isotonic Reductions

- Partial order on S, ⩿
  - ➔ Pairs of attributes may be comparable or **incomparable**.

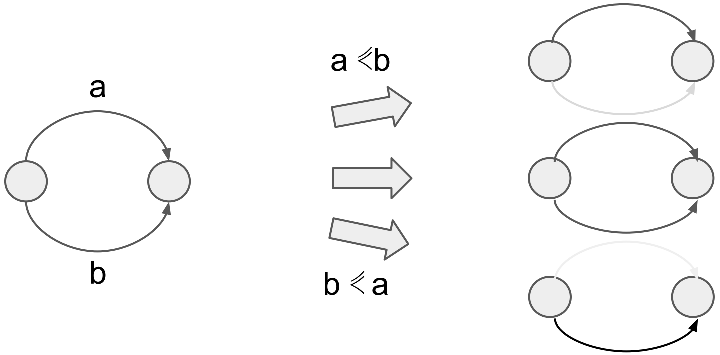

⩿ is an **isotonic reduction** of $⩿_i$ for ⊕ if [Lengauer & Theune 1991]:

(1) ⩿ is contained in $⩿_i$;
(2) ⊕ is isotone for ⩿.

Reducing an order declaring some pairs of attributes incomparable.

Dominant Paths

- Dominant attribute of set *A*:
  - ➔ Attribute that is not less preferred than any attribute of A.

- Set of dominant attributes from source *s* to destination *t*:
  - ➔ Set of dominant attributes of the set of path attributes from *s* to *t*.

If $\preccurlyeq$ is an isotonic reduction of $\preccurlyeq_i$ for $\oplus$, then:

set of dominant attributes from *s* to *t* **includes** optimal attribute from *s* to *t*.

Optimal path routing for non-isotone optimality criterion *becomes* Dominant path routing for isotone partial order

Product Order: Isotonic Reduction of Quickest Order

- Product order:

→ *(c, d)* is preferred to *(c', d')* if:
$c \geq c'$ and $d \leq d'$.

- *K*-quickest order

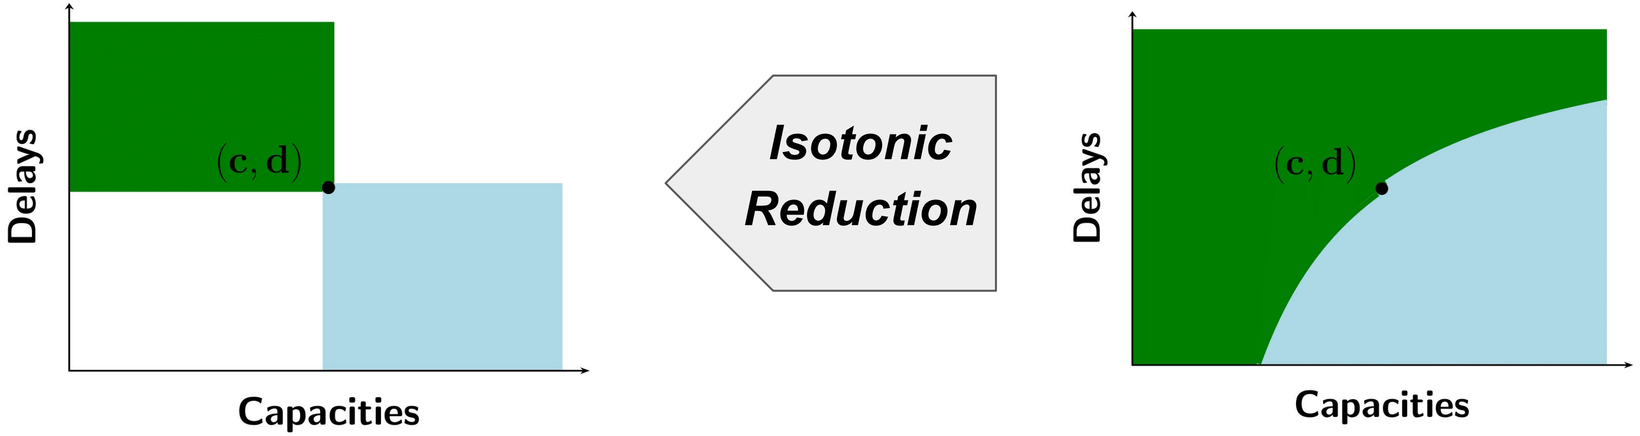

Less preferred than (c, d).

Preferred to (c, d).

Incomparable to (c, d).

Largest Isotonic Reductions

Largest isotonic reductions promote more efficient computations.

Largest isotonic reduction $\leqslant_R$ of $\leqslant$ for $\oplus$:

$a \leqslant_R b$ if $x \oplus a \leqslant x \oplus b$ for all attributes x, for all attributes a, b.

Largest Isotonic Reduction of Quickest Order

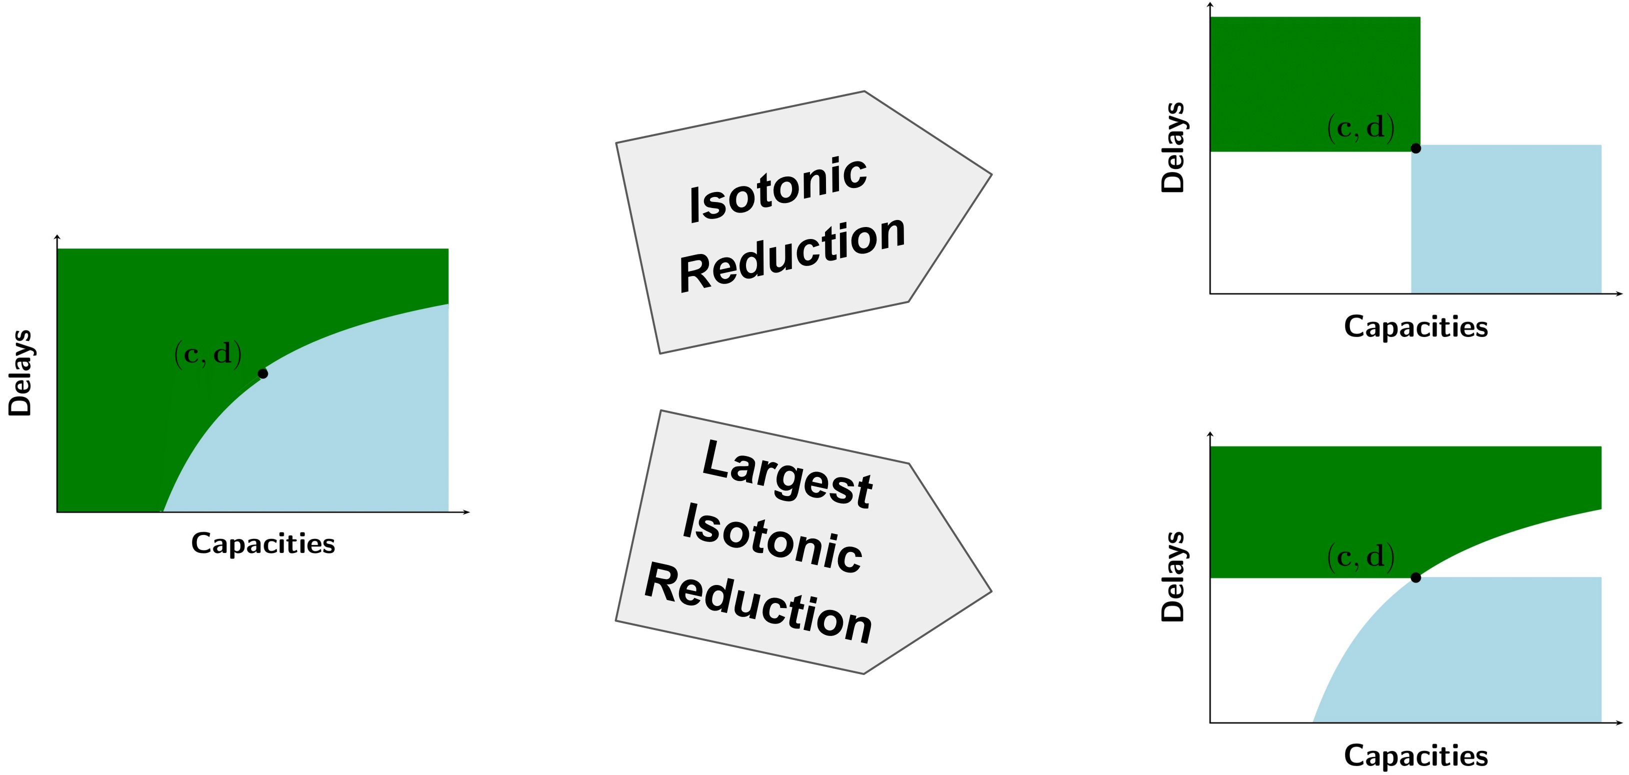

Dominant Paths for the Product Order

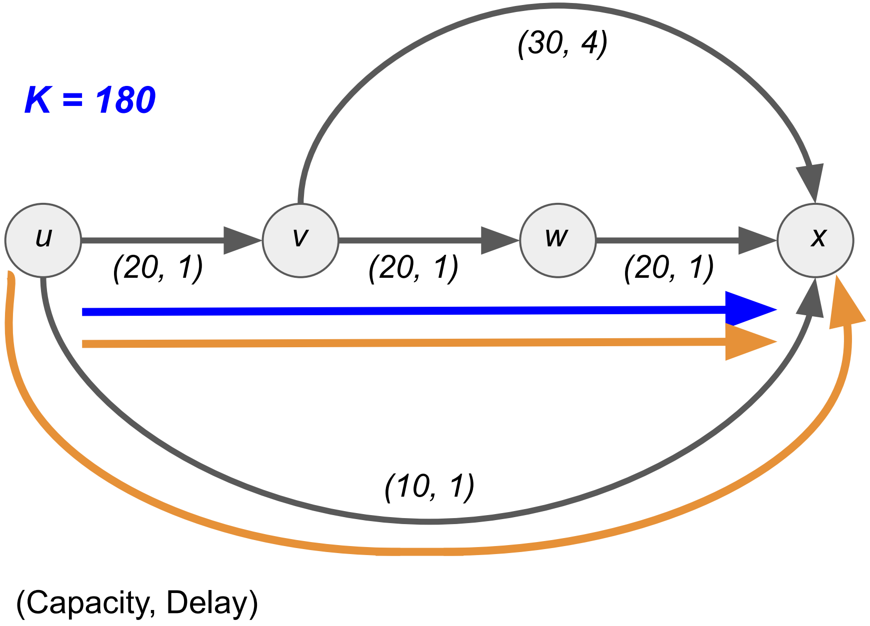

K-quickest path from *u* to *x* is ***uvwx***.

Dominant paths from *u* to *x* are ***uvwx*** and ***ux***:

- ➔ *(20, 3)* is incomparable to *(10, 1)*: *20 > 10, but 3 > 1;*
- ➔ *(20, 3)* is preferred to *(20, 5)*: *20 ≥ 20* and *3 < 5.*

Partial-Order Vectoring Protocols

- Separate path attribute computation per destination.
- When a node receives a **set of attributes** from an out-neighbor:
  - ➔ Extension of the link attribute to the out-neighbor with each of the received attributes originates a set of candidate attributes via the out-neighbor. ($\oplus$)
  - ➔ Election of the **set of dominant attributes** of the set of all candidate attributes ($\lessdot$); advertisement of the set of elected attributes to in-neighbors.
- Label-switching per destination:
  - ➔ Node assigns distinct labels to each elected attribute, advertised alongside the attributes [Chandranmenon & Varghese 1996].

Isotonicity ⇒ Dominant path routing

Routing on Quickest Paths (Revisited)

K = 180

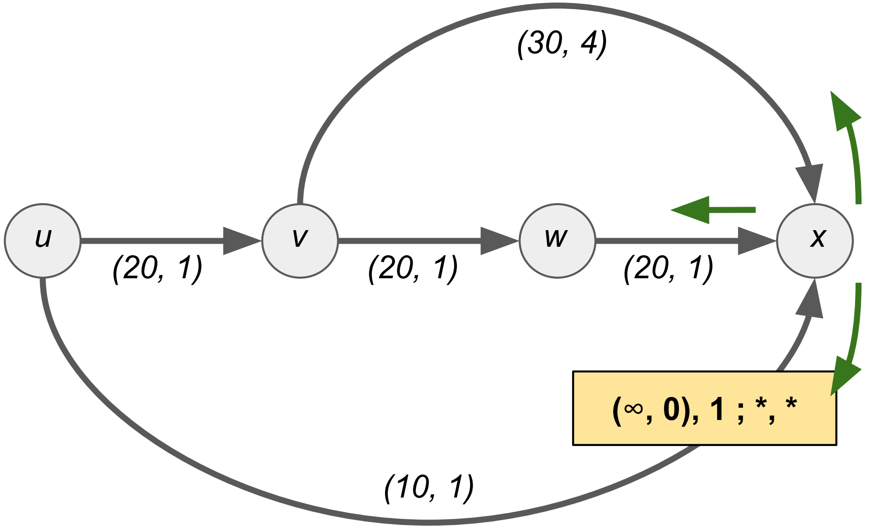

Elected pair, in-label ; out-neighbor, out-label
Candidate pair, - ; out-neighbor, out-label

Advertisement

Routing on Quickest Paths (Revisited)

K = 180

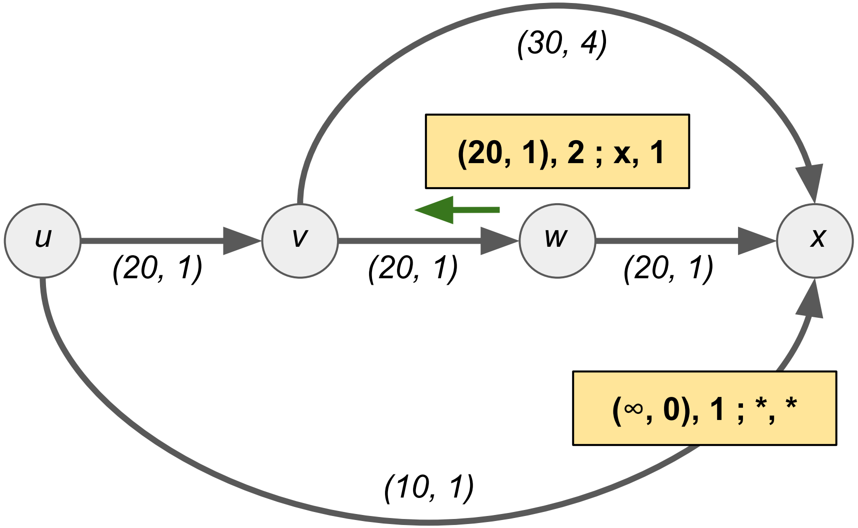

- Node *w*:

→ *(20, 1)* is the extension of *(20, 1)* with *(∞, 0)*.

Elected pair, in-label ; out-neighbor, out-label
Candidate pair, - ; out-neighbor, out-label

Advertisement

Routing on Quickest Paths (Revisited)

K = 180

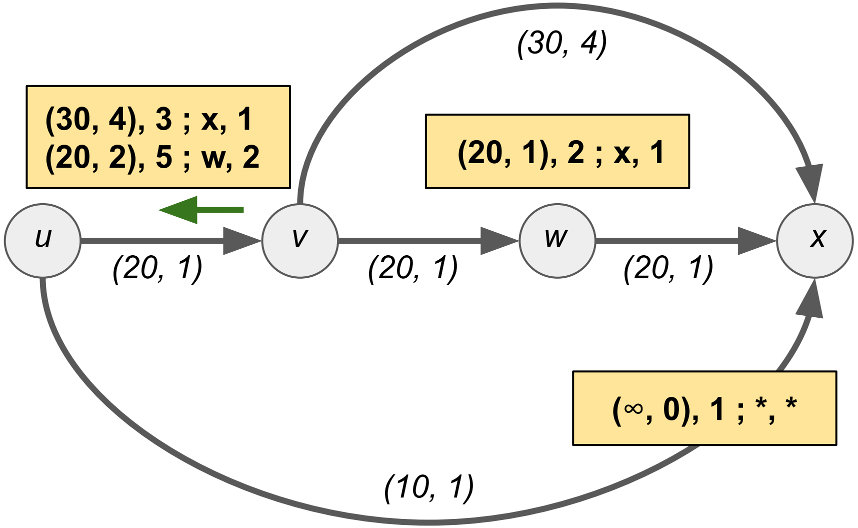

- Node *v*:

➔ *(30, 4)* is the extension of *(30, 4)* with *(∞, 0)*.

➔ *(20, 2)* is the extension of *(20, 1)* with *(20, 1)*.

➔ *(30, 4)* is incomparable to *(20, 2): 30 > 20, but 4 > 2.*

Elected pair, in-label ; out-neighbor, out-label
Candidate pair, - ; out-neighbor, out-label


Advertisement

Routing on Quickest Paths (Revisited)

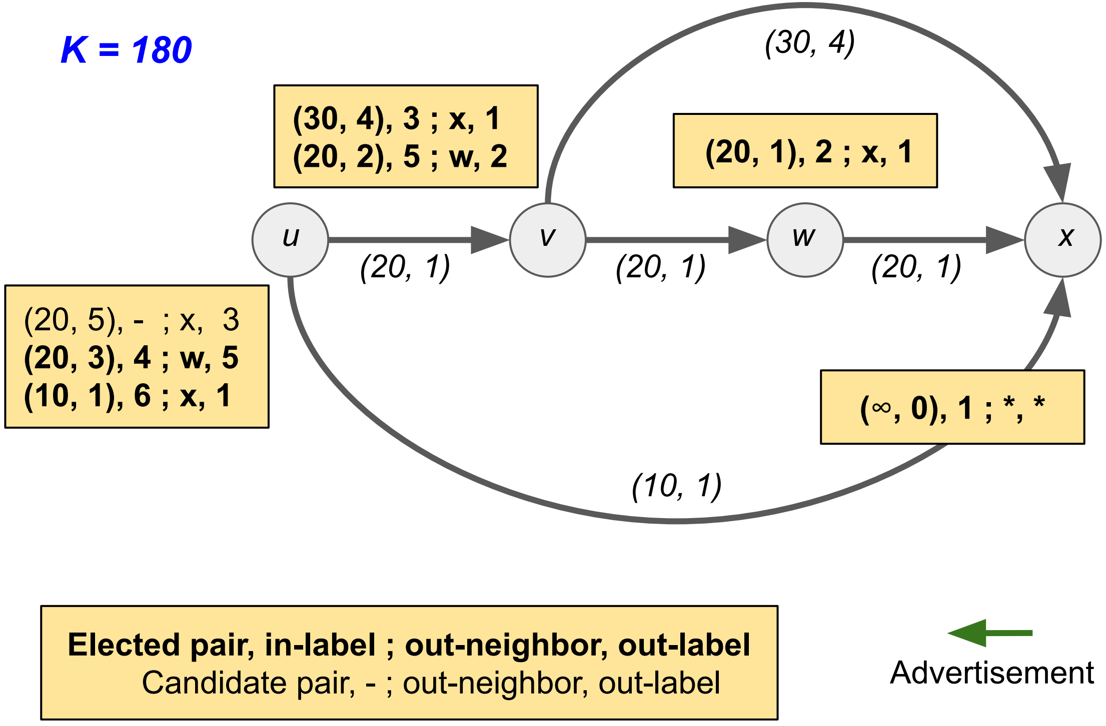

- Node *u*:

➔ *(20, 5)* is the extension of *(20, 1)* with *(30, 4)*.

➔ *(10, 1)* is the extension of *(10, 1)* with *(∞, 0)*.

➔ *(20, 3)* is incomparable to *(10, 1)*: *20 > 10, but 3 > 1.*

➔ *(20, 3)* is preferred to *(20, 5)*: *20 ≥ 20* and *3 < 5*.

Routing on Quickest Paths (Revisited)

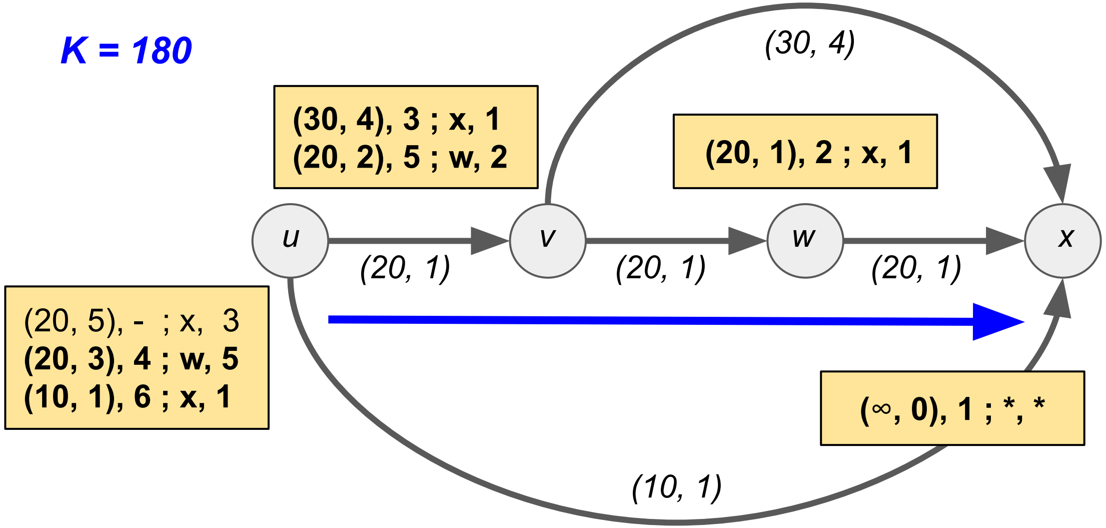

K-quickest path from *u* to *x*:

K = *180*: Data-packets traverse path **uvwx**.

Multiple Optimal Paths and Intersection

- Collection of optimality criteria, *I*.
  - ➔ Each optimality criterion $i \in I$ defined by a total order $\leqslant_i$ on *S*.

Intersection $\leqslant$ of all total orders $\leqslant_i$, $i \in I$:

$a \leqslant b$ if $a \leqslant_i b$ for all *I*, for all attributes *a*, *b*.

Partial order $\leqslant$ *is a reduction of each of the total orders* $\leqslant_i$, $i \in I$.

Product Order: Intersection of the Quickest Orders

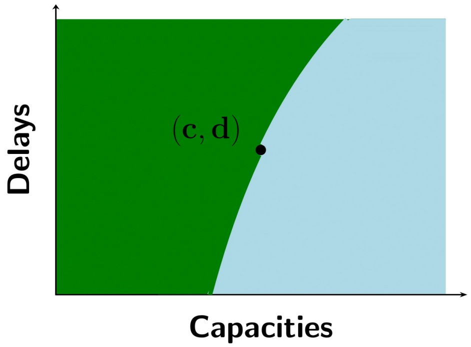

⋮

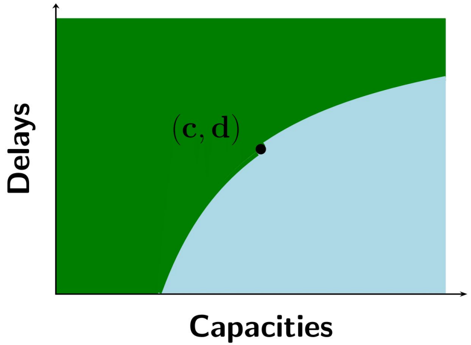

Intersection

(all K-quickest orders)

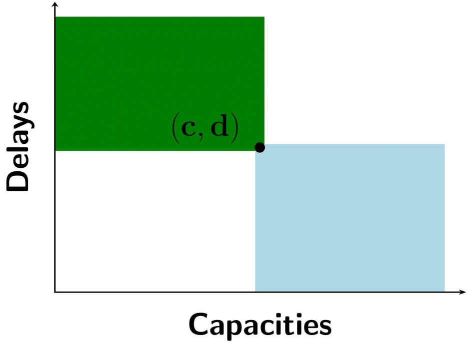

Routing on Multiple Quickest Paths

K = 20

K = 180

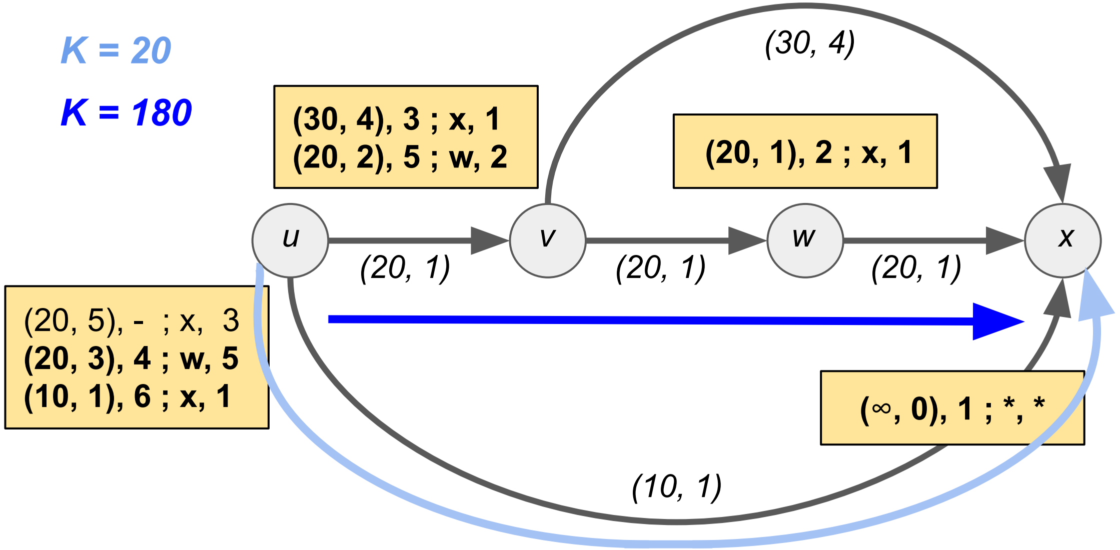

K-quickest path from *u* to *x*:

K = *20*: Data-packets traverse path ***ux***.

K = *180*: Data-packets traverse path ***uvwx***.

Solution Overview

Collection of Optimality Criteria → **Intersection** → Partial Order → **Isotonic Reduction** → Isotone Partial Order

Partial-Order Vectoring Protocols → *allow* → Optimal Path Routing for Multiple Optimality Criteria

Non-Restarting and Restarting Vectoring Protocols

	Non-Restarting	Restarting
Routing State	➔ Set of candidate attributes per out-neighbor. ➔ Set of elected attributes; improves and worsens.	➔ *Only* set of elected attributes; can solely improve.
Stable State	➔ With isotonicity, delivery on dominant paths;	
	➔ With or without isotonicity, delivery of data-packets.	➔ Without isotonicity, permanent black-holes.
Transient Behaviour	➔ Count-to-infinity. ➔ Forwarding-loops.	➔ Black-holes.

Table of Contents

Main Questions

- Input networks:
  - ➔ Rocketfuel networks [Spring 2004]: topologies of ISPs with every link annotated with a capacity and a delay.
- Computational toolbox:
  - ➔ Dijkstra’s-like algorithm designed to operate on isotone partial order;
  - ➔ Homegrown simulator of partial-order vectoring protocols.
- Output answers:
  - ➔ What are the sizes of dominant attributes sets?;
  - ➔ What are the termination times of partial-order vectoring protocols?

Number of Dominant Capacity-Delays

ASN	Nodes	Links	Distinct capacities	Dominant capacity-delays
1221	*50*	*194*	*8*	*1.9*
1239	*284*	*1882*	*19*	*2.5*
1755	*73*	*292*	*18*	*2.2*
3257	*113*	*558*	*21*	*3.5*
3967	*72*	*180*	*19*	*3.7*
6461	*129*	*726*	*19*	*2.8*

Average number of dominant capacity-delays *fall short of* number of distinct link capacities.

Number of Dominant Capacity-Delays

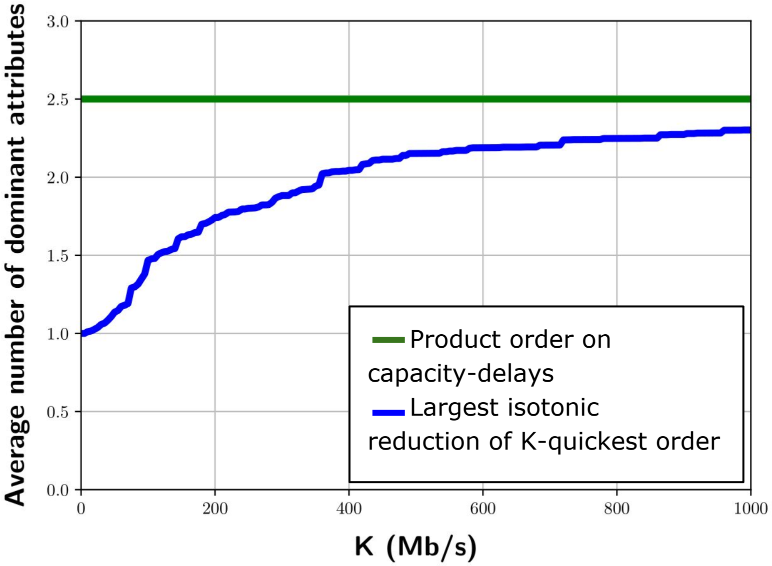

Largest isotonic reduction *improves* routing state of partial-order vectoring protocols.

Termination Times (Non-Restarting Protocols)

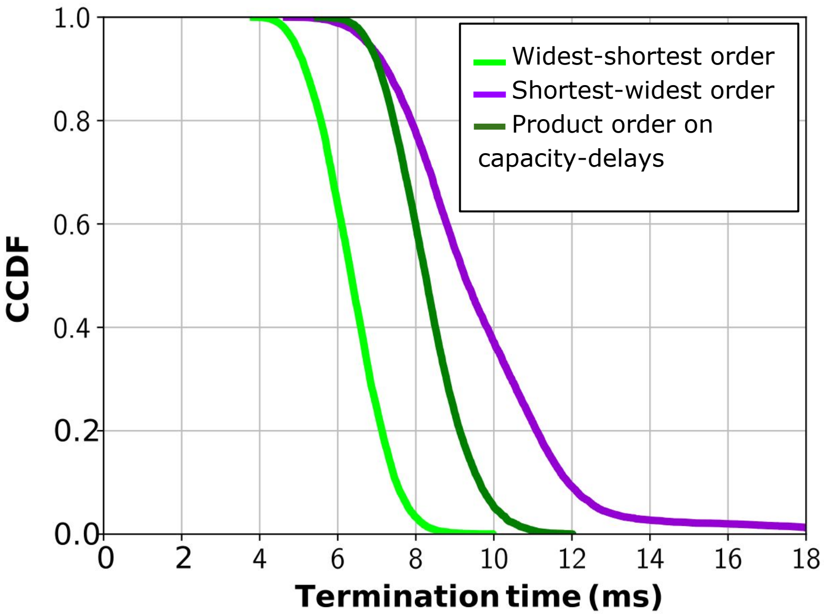

Network-wide announcement of a destination

Termination times for widest-shortest order and for product order have *similar behaviors*.

Termination times for product order are *better* than for shortest-widest order.

Termination Times (Non-Restarting Protocols)

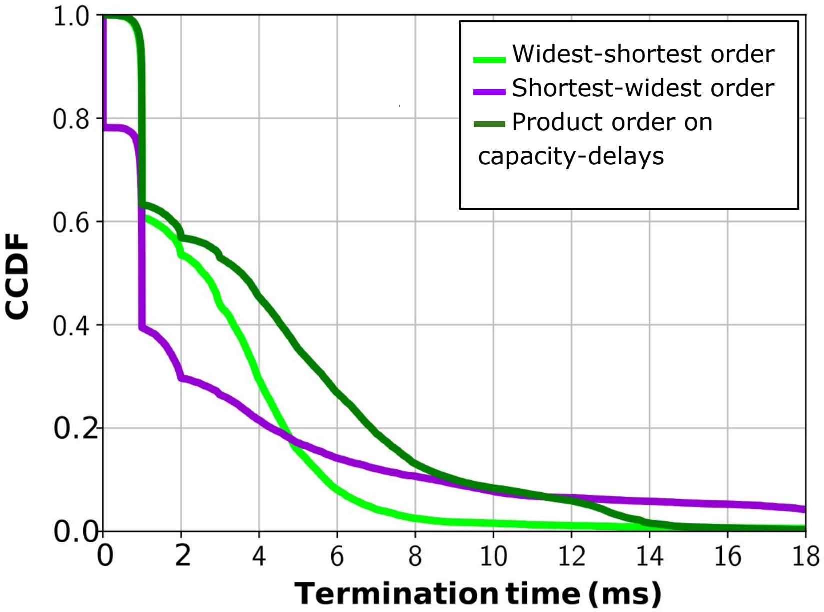

Failure of a link

Table of Contents

Summary of Results

- Formulation of two optimal path routing problems:
  - ➔ Optimal path routing for arbitrary optimality criterion.
  - ➔ Optimal path routing for multiple optimality criteria.
- Introduction of new routing concepts and protocols:
  - ➔ Intersection of collection of total orders and isotonic reduction of partial orders.
  - ➔ Non-restarting and restarting partial-order vectoring protocols.
- Approach is generic:
  - ➔ Other examples: wide-shortest paths; widest paths with co-channel interference [Draves 2004];
- Partial-order vectoring protocols termination fast with few elected attribute per destination

Ongoing Research

- Other routing problems:
  - ➔ Optimal path routing: routing on delay-constrained least-cost paths [Salama 1997]; routing with probabilistic delay guarantees [Brand 2008];
  - ➔ Regular-expression-constrained routing: routing on service-chaining-constrained paths [Choi 2001];
  - ➔ Interdomain routing.

- In-depth study of restarting partial-order vectoring protocols
  - ➔ Data plane.
  - ➔ Stability across computation instances.

Questions?

Routing on Multiple Optimality Criteria

João L. Sobrinho Miguel A. Ferreira

Instituto de Telecomunicações,
Instituto Superior Técnico, Universidade de Lisboa

October 2020