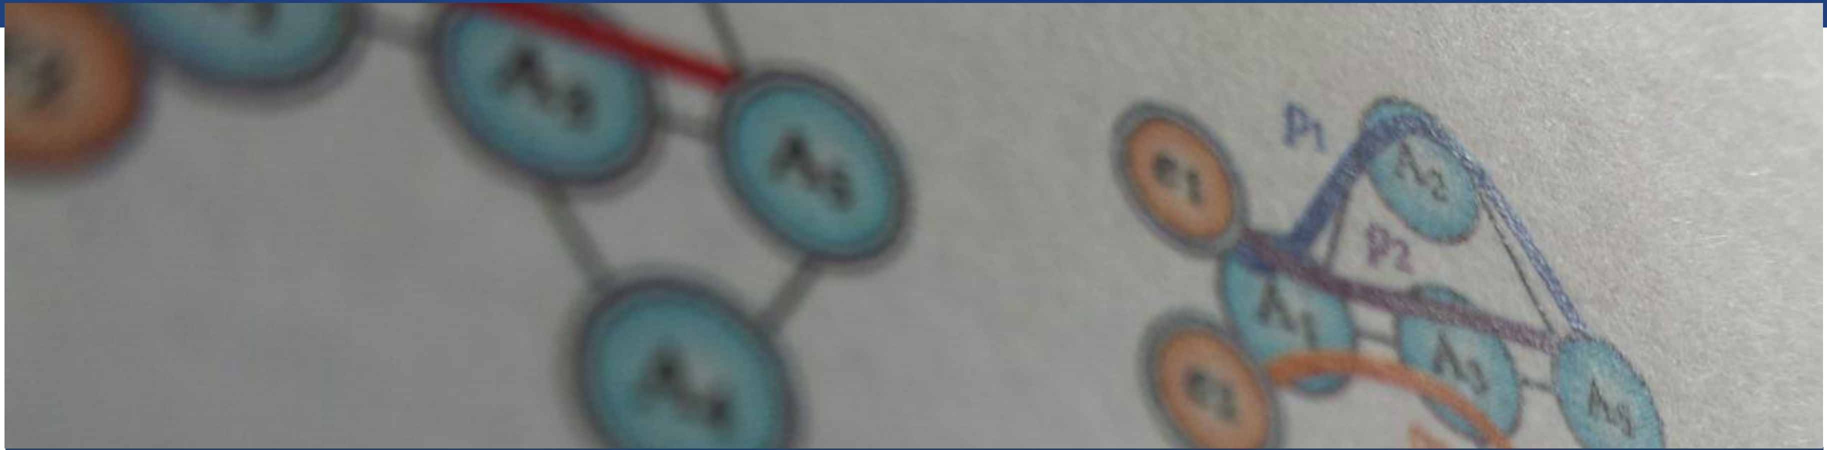

The Value of Information in Selfish Routing

Simon Scherrer, Adrian Perrig, Stefan Schmid

27th International Colloquium on Structural Information and Communication Complexity (SIROCCO)
June 29 - July 1, 2020

Network-based path selection

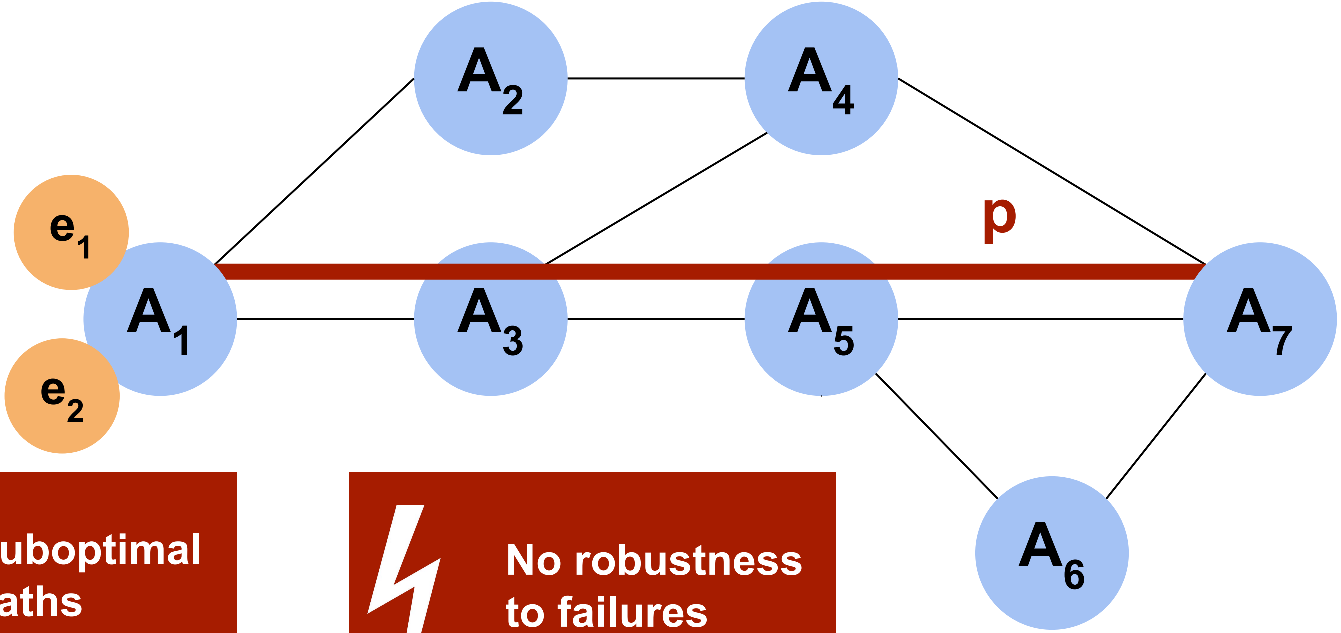

Suboptimal paths

No robustness to failures

Source-based path selection

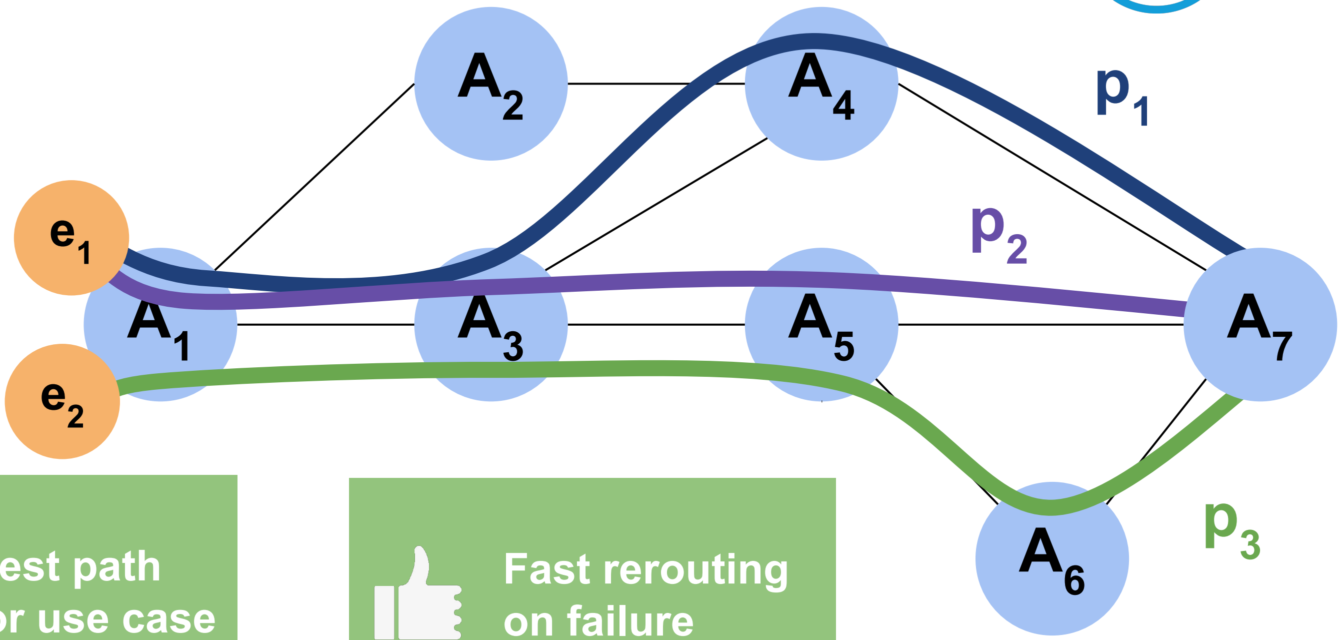

Best path for use case

Fast rerouting on failure

Network-based path selection: Network operator view

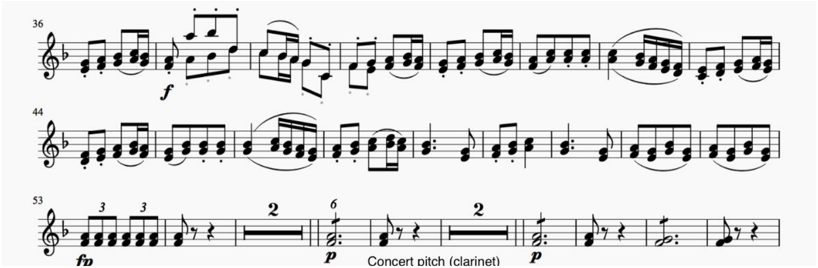

Source-based path selection: Network operator view

Goals of our work

Revisit selfish-routing concepts to investigate two issues arising in emerging path-aware Internet architectures:

- **Impact of information:** What network state information should be shared with end-hosts?

- **Impact on network operators:** What is the impact of selfish routing on the cost of network operators?

Price of Anarchy: Three components

C Social cost function

F^{opt} Social optimum

F^{eq} Equilibrium

$$PoA = \frac{C(F^{eq})}{C(F^{opt})}$$

Adapted Wardrop model of source-based path selection

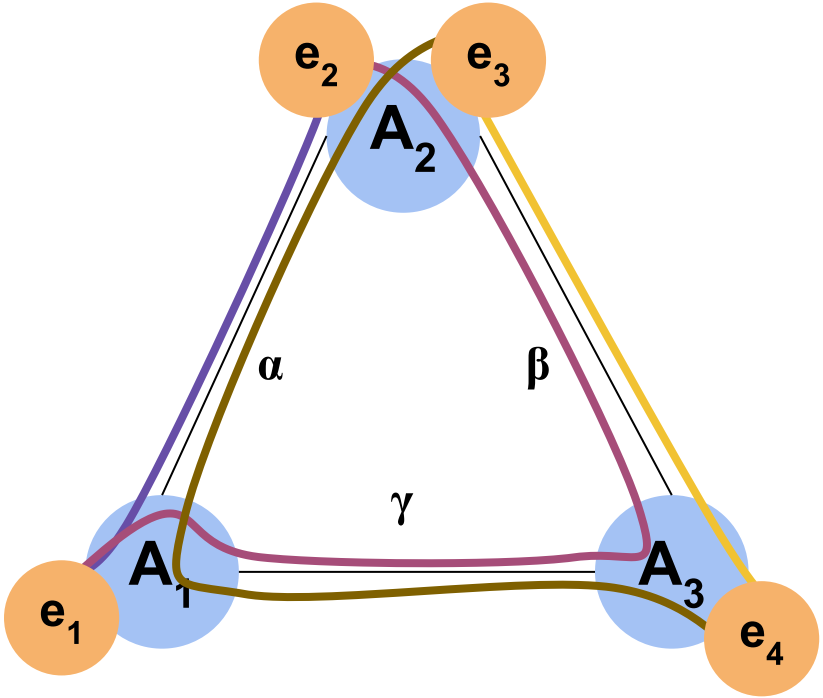

$$\mathbf{d} = (d_{1,2}, d_{3,4}) = (1, 1)$$

$$\mathbf{F} = (F_{\alpha}, F_{\gamma\beta}, F_{\beta}, F_{\alpha\gamma})$$

$$\mathbf{f} = (f_{\alpha}, f_{\beta}, f_{\gamma})$$

$$c_{\alpha}(f_{\alpha}) = 1$$
$$c_{\beta}(f_{\beta}) = f_{\beta}^{2}$$
$$c_{\gamma}(f_{\gamma}) = f_{\gamma}$$

$$C_{\pi}(\mathbf{F}) = \Sigma_{\ell \in \pi} c_{\ell}$$

Total cost functions and social optima

End-host cost function: $C^* = \Sigma_{\text{end-hosts}} \Sigma_{\text{paths}}$ flow on path $\cdot$ path cost

(classic) $= \Sigma_{\pi \in \Pi} F_\pi \cdot C_\pi(\mathbf{F}) = \Sigma_{\ell \in L} f_\ell \cdot c_\ell(f_\ell)$

End-host optimum: $\mathbf{F}^* = \text{argmin}_{\mathbf{F}} C^*(\mathbf{F})$

Network-operator cost function: $C^\# = \Sigma_{\text{links}}$ link cost $= \Sigma_\ell c_\ell(f_\ell)$

Network-operator optimum: $\mathbf{F}^\# = \text{argmin}_{\mathbf{F}} C^\#(\mathbf{F})$

Characterizing social optima: Suboptimal path flow pattern

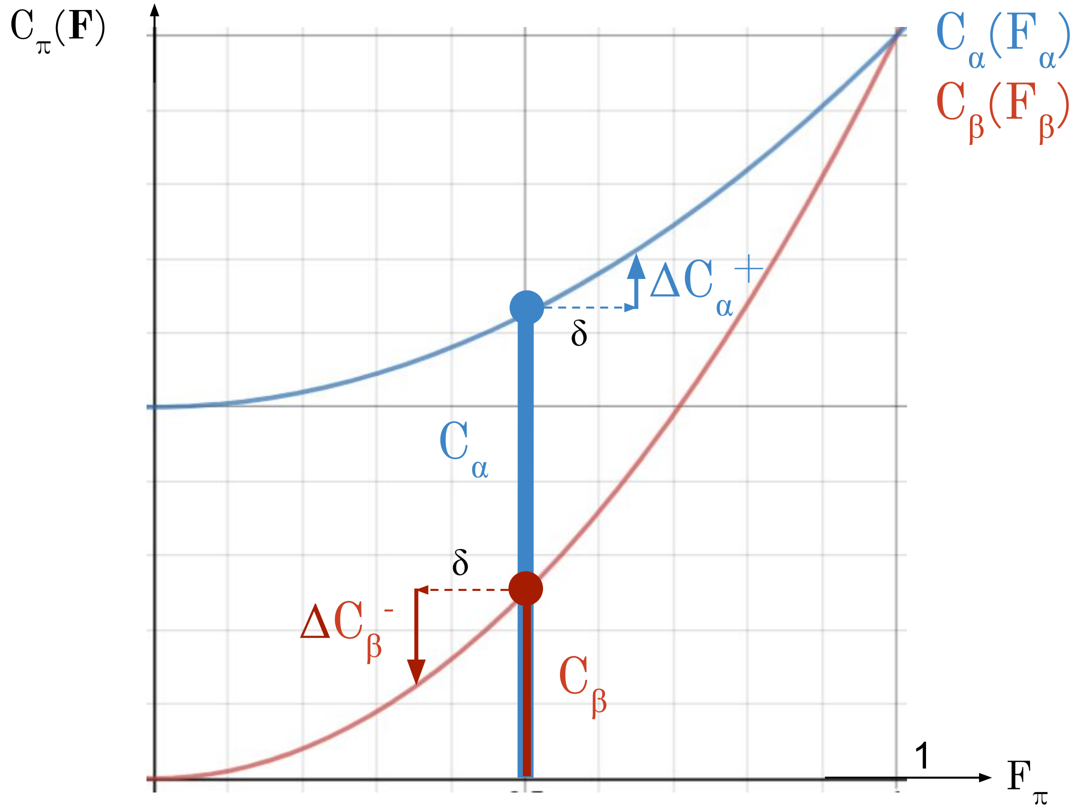

$$\mathbf{d} = (d_{1,2}) = (1)$$

$$\mathbf{F} = (F_\alpha, F_\beta)$$

$$C(\mathbf{F}) = C_\alpha(F_\alpha) + C_\beta(F_\beta)$$

$$\exists\ \delta.\ |\Delta C_\alpha^+| < |\Delta C_\beta^-|$$

$\Rightarrow$ C can be reduced

Characterizing social optima: Optimal path flow pattern

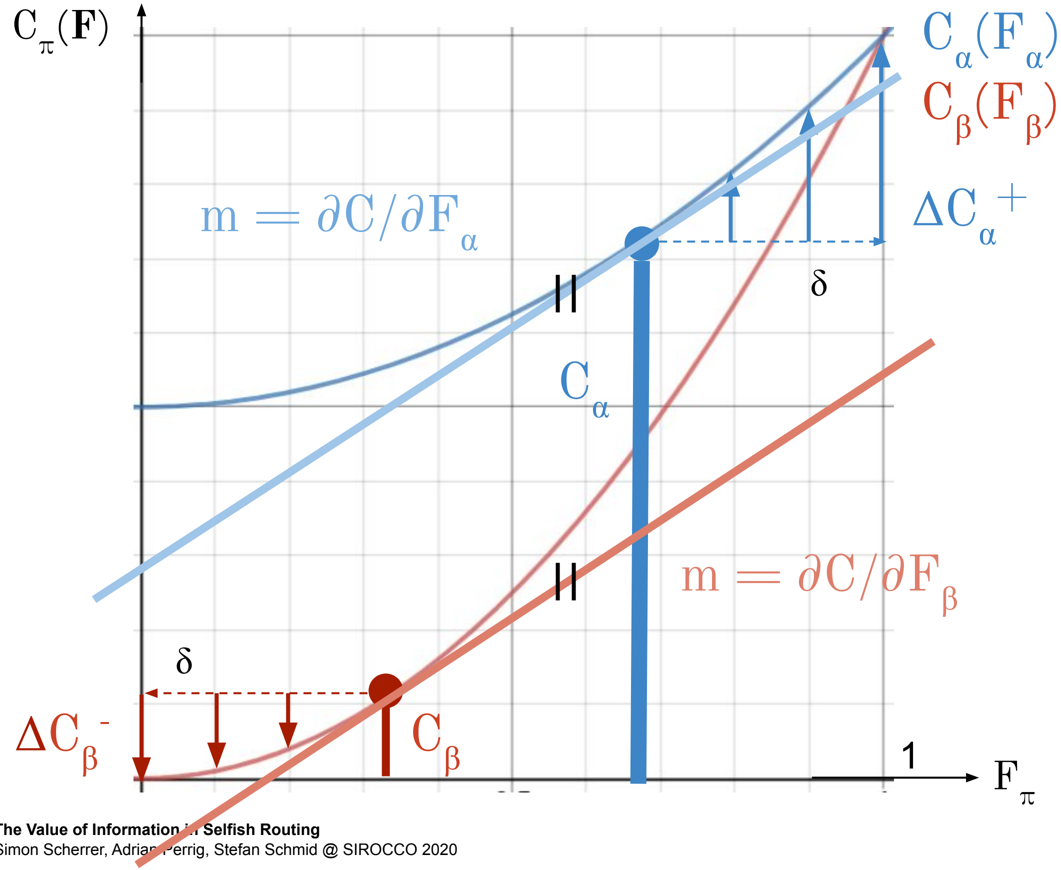

$\forall \delta. |\Delta C_\alpha^+| > |\Delta C_\beta^-|$

$\Rightarrow$ C **cannot** be reduced

$\Rightarrow$ C is optimal

$\partial C/\partial F_\alpha = \partial C/\partial F_\beta$

$\Rightarrow \forall \delta. |\Delta C_\alpha^+| > |\Delta C_\beta^-|$

$\Rightarrow$ C is optimal

Socially optimal marginal costs

$\partial C(\mathbf{F})/\partial F_\pi$ is the *marginal cost* of path π

given path-flow pattern $\mathbf{F}$

A path-flow pattern $\mathbf{F}$ is optimal w.r.t. a cost function $C \in \{C^*, C^\#\}$

if for every origin-destination pair:

$$F_\alpha, \ldots, F_\rho > 0 \qquad F_\sigma, \ldots, F_\omega = 0$$

$$\frac{\partial C(\mathbf{F})}{\partial F_\alpha} = \ldots = \frac{\partial C(\mathbf{F})}{\partial F_\rho} \leq \frac{\partial C(\mathbf{F})}{\partial F_\sigma} \leq \ldots \leq \frac{\partial C(\mathbf{F})}{\partial F_\omega}$$

Social optimum: Comparison (Example)

$$F^{\#} = (F_\alpha, F_\beta, F_\gamma) = (\tfrac{1}{2}, 0, \tfrac{1}{2})$$

$$F^{*} = (F_\alpha, F_\beta, F_\gamma) = (\tfrac{2}{3}, \tfrac{1}{3}, 0)$$

$\alpha \quad C_\alpha(F_\alpha) = F_\alpha^{\,2}$

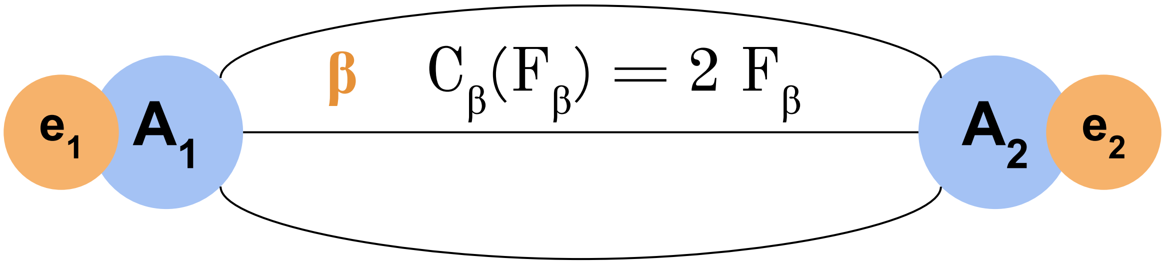

$\beta \quad C_\beta(F_\beta) = 2\,F_\beta$

$\gamma \quad C_\gamma(F_\gamma) = F_\gamma + 2$

Different optima! Network operators prefer usage of links with little variable cost (here: γ)

Price of Anarchy: Where are we?

C Total cost function ✓

F^{opt} Social optimum ✓

F^{eq} Equilibrium

$$PoA = \frac{C(F^{eq})}{C(F^{opt})}$$

Equilibrium with latency-only information (LI equilibrium)

$$\mathbf{d} = (d_{1,2}) = (1)$$

$$\mathbf{F} = (F_\alpha, F_\beta) = (1, 0)$$

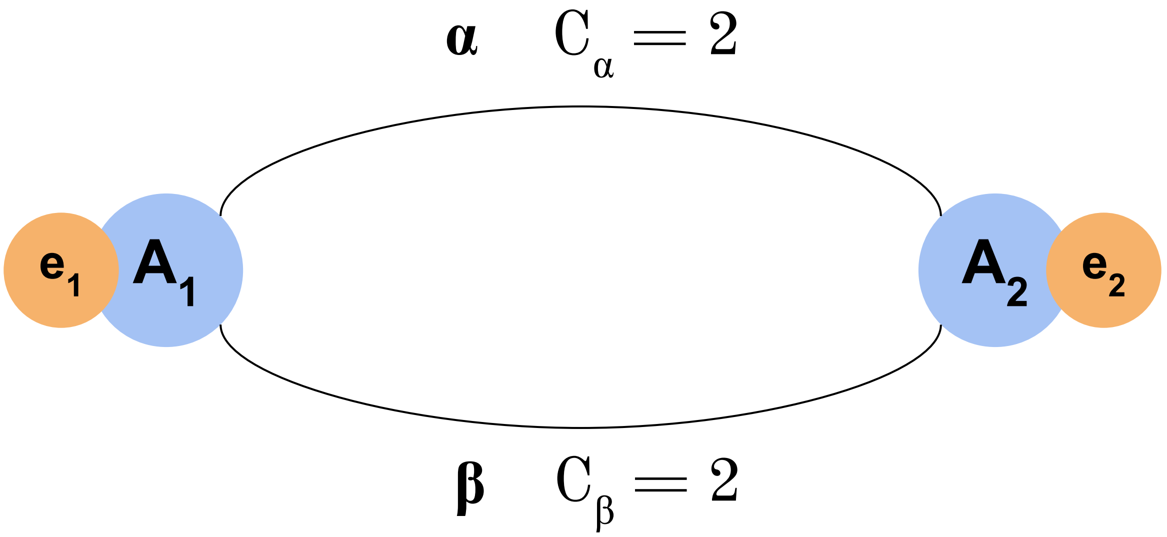

$C_\alpha = C_\beta \Rightarrow \mathbf{F} = (1,0)$ is an LI equilibrium

Characterizing the LI equilibrium

A path flow pattern $\mathbf{F}$ is an LI equilibrium

if for every origin-destination pair:

$$F_\alpha, ..., F_\rho > 0 \qquad F_\sigma, ..., F_\omega = 0$$

$$C_\alpha(\mathbf{F}) = ... = C_\rho(\mathbf{F}) \leq C_\sigma(\mathbf{F}) \leq ... \leq C_\omega(\mathbf{F})$$

Equilibrium with perfect information (PI equilibrium)

$\mathbf{d}_{(1)} = (d_{1,2}) = (1)$

$\mathbf{F}_{(1)} = (F_\alpha, F_\beta)$

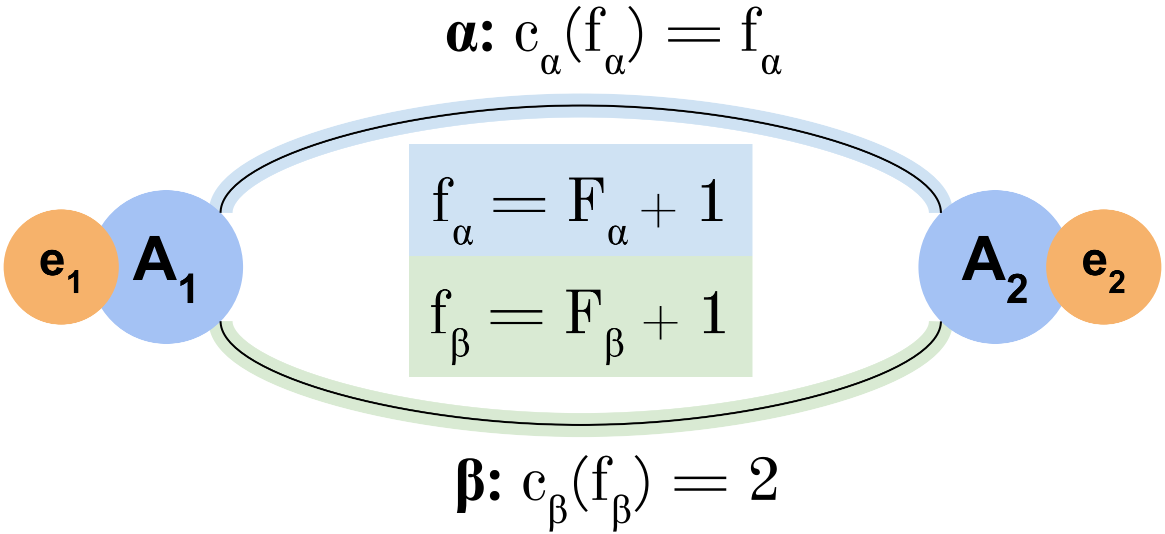

Minimize selfish cost $C_{(1)}(\mathbf{F}_{(1)}) = F_\alpha \cdot (F_\alpha + 1) + F_\beta \cdot 2$

$\Rightarrow (F_\alpha, F_\beta) = (2/3, 1/3)$ is a PI equilibrium

Characterizing the PI equilibrium

A path flow pattern $\mathbf{F}$ is a PI equilibrium

if for every origin-destination pair of any end-host e:

$$F_\alpha, \ldots, F_\rho > 0 \qquad F_\sigma, \ldots, F_\omega = 0$$

$$\frac{\partial C_{(e)}(\mathbf{F})}{\partial F_\alpha} = \ldots = \frac{\partial C_{(e)}(\mathbf{F})}{\partial F_\rho} \leq \frac{\partial C_{(e)}(\mathbf{F})}{\partial F_\sigma} \leq \ldots \leq \frac{\partial C_{(e)}(\mathbf{F})}{\partial F_\omega}$$

Capturing the value of information

Information assumption	Latency-only Information (LI)	Perfect Information (PI)
Equilibrium	$\mathbf{F}^0$	$\mathbf{F}^+$
Price of Anarchy	$\mathrm{PoA}^0 = \frac{\mathrm{C}(\mathbf{F}^0)}{\mathrm{C}(\mathbf{F}^{\mathrm{opt}})}$	$\mathrm{PoA}^+ = \frac{\mathrm{C}(\mathbf{F}^+)}{\mathrm{C}(\mathbf{F}^{\mathrm{opt}})}$

$\Delta =$ Value of Information (VoI)

The benefits of information

$$\mathrm{VoI} > 0$$

The benefits of information: Network of parallel links

(cf. Roughgarden 2003)

$\Sigma_k d_{k,T} = 1$

$\mathbf{F} = (F_{1\alpha}, F_{1\beta}, \ldots F_{K\alpha}, F_{K\beta})$

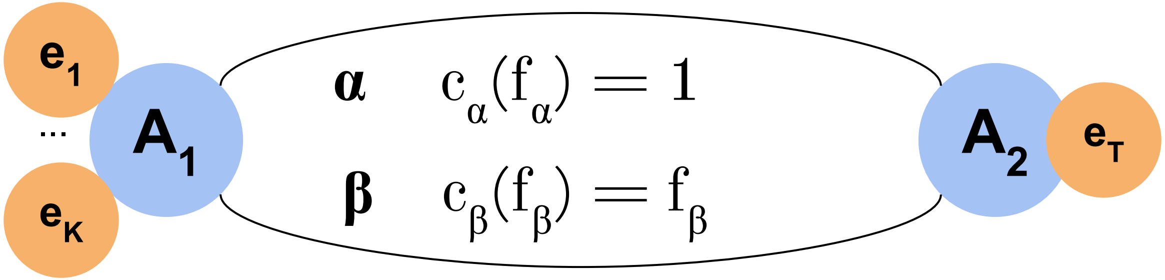

EH Opt: $\mathbf{F}^*$ s.t. $f_\beta = 1/2$

LI Eq: $\mathbf{F}^0$ s.t. $f_\beta = 1$

NO Opt: $\mathbf{F}^\#$ s.t. $f_\beta = 0$

PI Eq: $\mathbf{F}^+$ s.t. $f_\beta = K/(K+1)$

The benefits of information: Network of parallel links

	LI equilibrium	PI equilibrium
End-host perspective	$\mathrm{PoA}^{*0} = 4/3$	$\mathrm{PoA}^{*+} = (K^2+K+1)/(K^2+2K+1) \cdot 4/3 \leq \mathrm{PoA}^{*0}$
Network-operator perspective	$\mathrm{PoA}^{\#0} = 2$	$\mathrm{PoA}^{\#+} = 1 + K/(K+1) \leq 2 = \mathrm{PoA}^{\#0}$

The benefits of information: Network of parallel links

	LI equilibrium	PI equilibrium
End-hos perspecti		$\text{PoA}^{*+} = (K^2+K+1)/(K^2+2K+1)\cdot 4/3 \le \text{PoA}^{*0}$
Network operato perspecti		$\text{PoA}^{\#+} = 1 + K/(K+1) \le 2 = \text{PoA}^{\#0}$

PI equilibrium cheaper than LI equilibrium

$\text{VoI} > 0$

The drawbacks of information

$$\mathrm{VoI} < 0$$

The drawbacks of information: Ladder network

$$\mathbf{d} = (d_{11,12}, d_{21,22}) = (1, 1)$$

$$\mathbf{F}^{\rightarrow} = (1, 0, 1, 0)$$

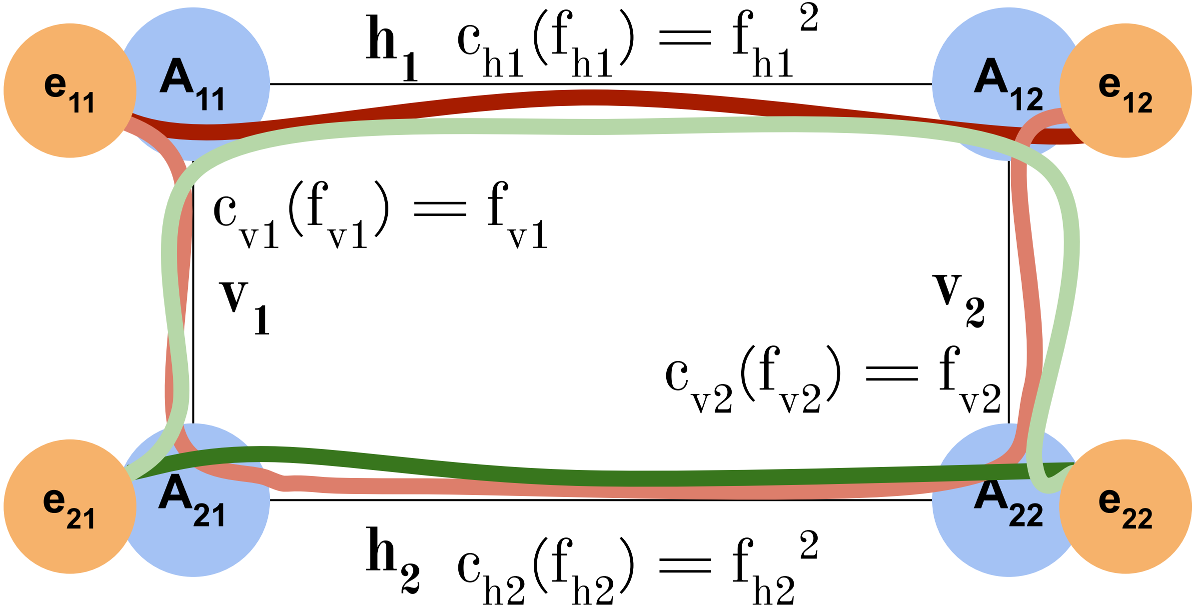

Direct-only $\mathbf{F}^{\rightarrow}$ is universally optimal: $\mathbf{F}^{\rightarrow} = \mathbf{F}^{*} = \mathbf{F}^{\#}$

The drawbacks of information: Ladder network

$$\mathbf{F}^{\rightarrow} = (1, 0, 1, 0)$$

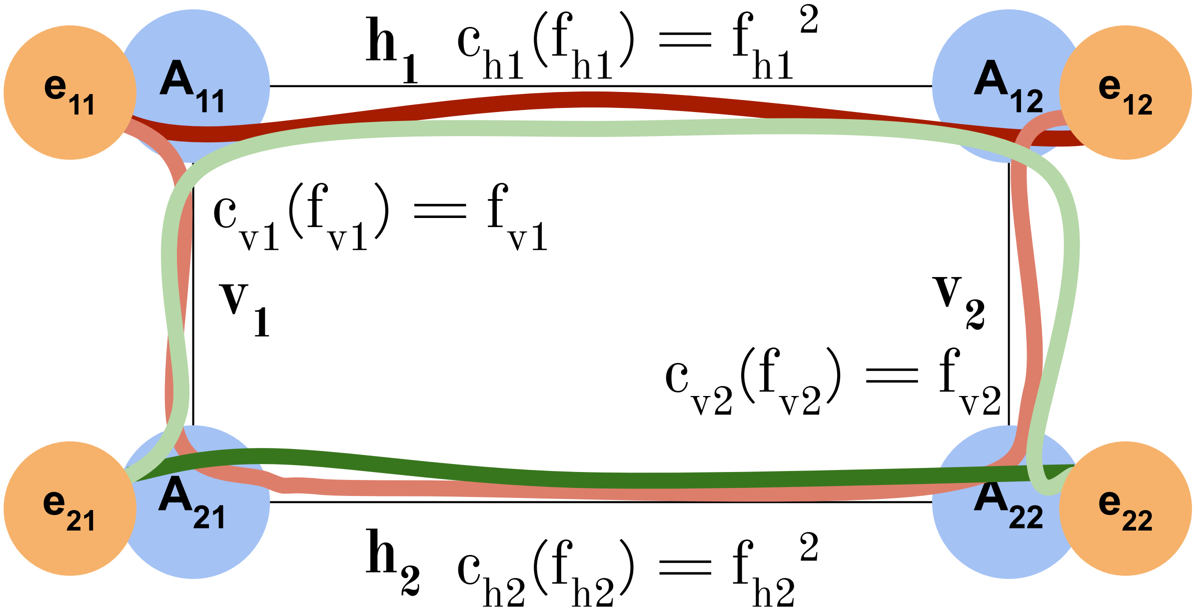

$$C_{1H}(\mathbf{F}^{\rightarrow}) = 1 = C_{1V}(\mathbf{F}^{\rightarrow}) \quad \Rightarrow \quad \mathbf{F}^{\rightarrow} = \mathbf{F}^{0}$$

(LI equilibrium is optimal)

The drawbacks of information: Ladder network

$$\mathbf{F}^{\rightarrow} = (1,\ 0,\ 1,\ 0)$$

$$\mathbf{F}^{\sim} = (0.9,\ 0.1,\ 1,\ 0)$$

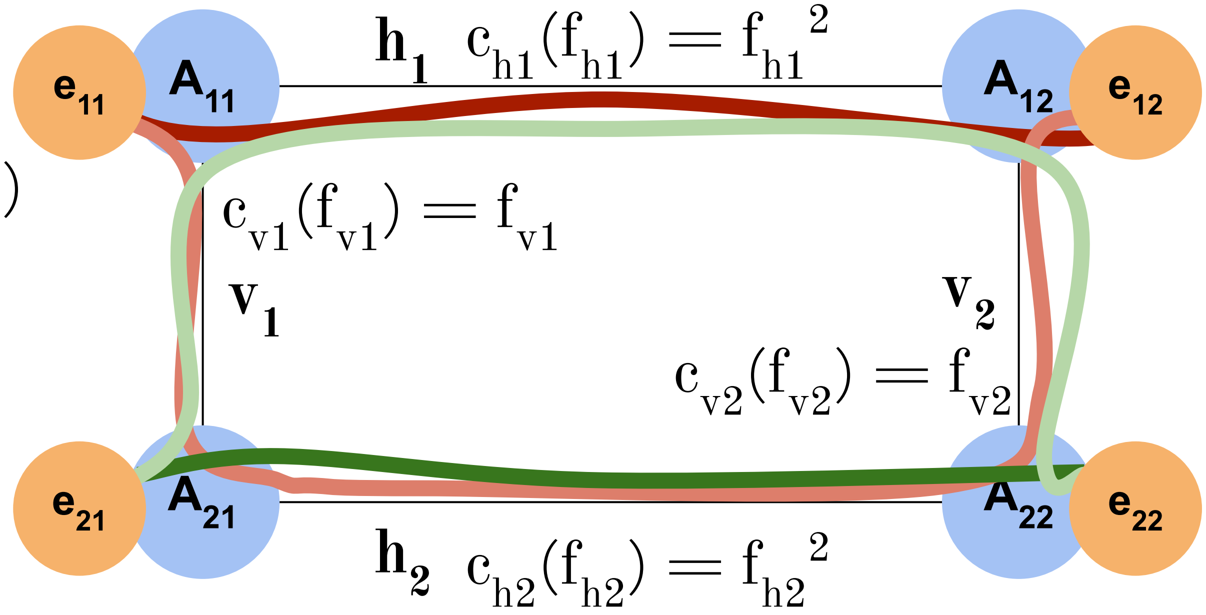

$$C_{(1)}(\mathbf{F}^{\rightarrow}) = 1 > C_{(1)}(\mathbf{F}^{\sim}) = 0.87 \Rightarrow \mathbf{F}^{\rightarrow} \neq \mathbf{F}^{+}$$

(PI equilibrium is suboptimal)

ETH zürich

The drawbacks of information: Ladder network

$\mathbf{F}^{\rightarrow} =$

$C_{(1)}(\mathbf{F}$

$F_{1H} \cdot$

$F_{1V} \cdot$

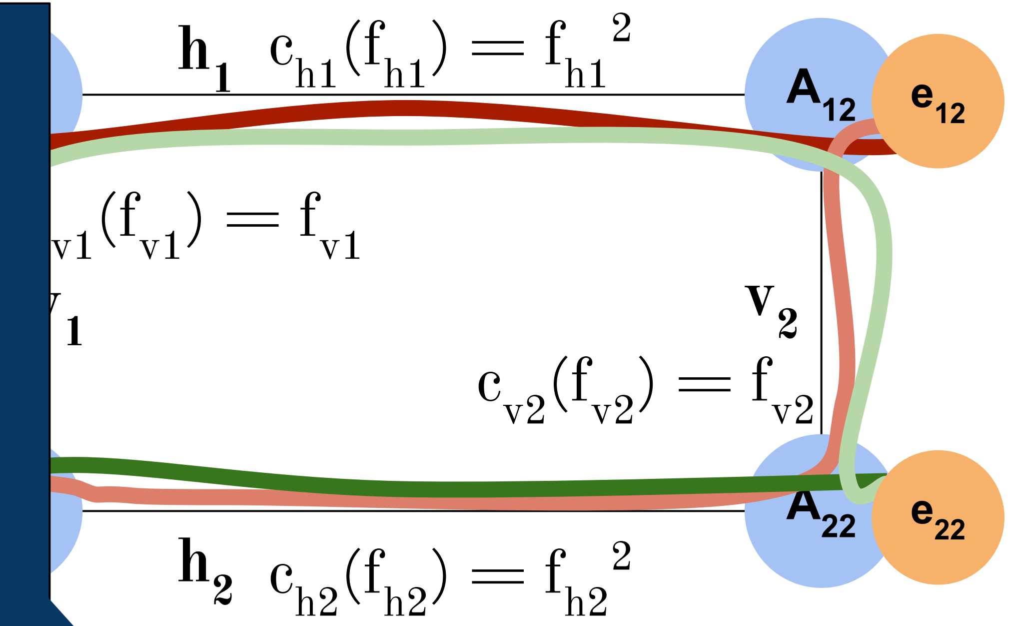

$\mathbf{F}^{\sim} =$

PI equilibrium

more costly than

LI equilibrium

$\mathrm{VoI} < 0$

$C_{(1)}(\mathbf{F}^{\rightarrow}) = 1 > C_{(1)}(\mathbf{F}^{\sim}) = 0.8 \Rightarrow \mathbf{F}^{\rightarrow} \neq \mathbf{F}^{+}$

(PI equilibrium is suboptimal)

The drawbacks of information: Generalized ladder network

Upper bound on PoA for network operators:

$$\mathrm{PoA}^{\#+} \leq 1 + \frac{2(H-1)}{3H}\, p$$

$$\leq 1 + \frac{2}{3}\, p$$

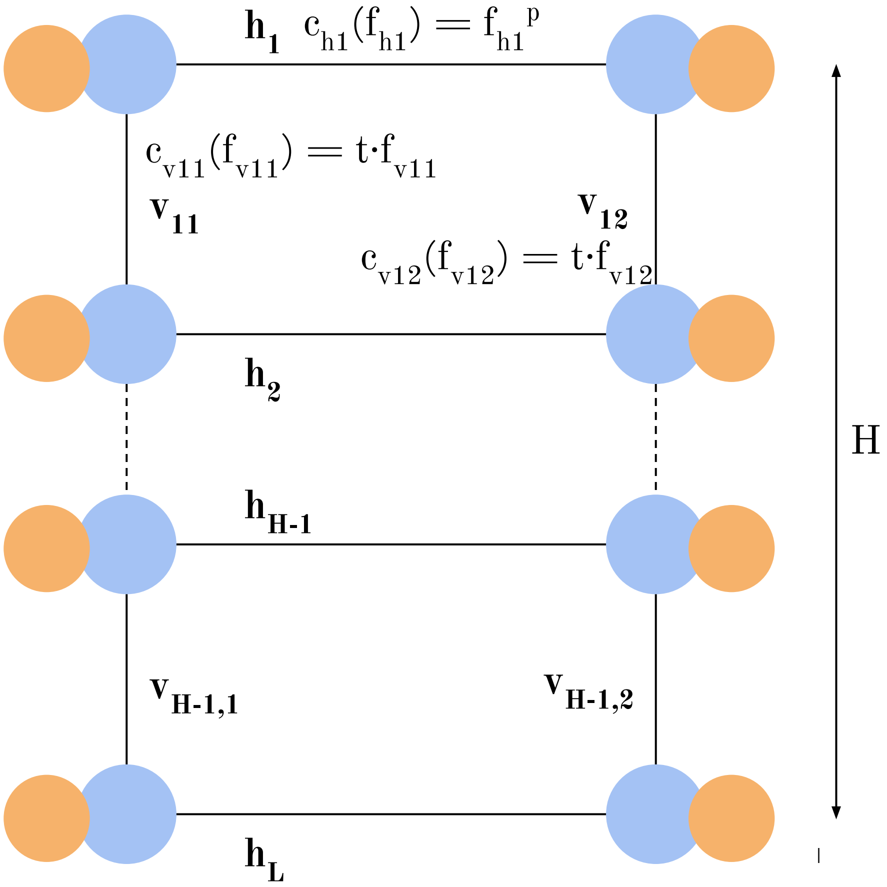

Drawback of information: Abilene Topology Case Study

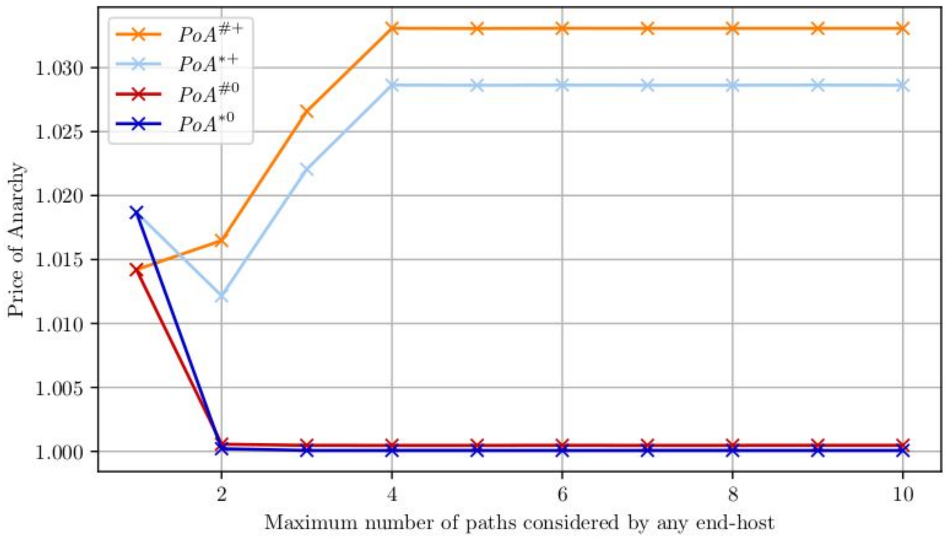

Questions

Thank you for your attention!

Happy to answer questions in the chat forum!

Or by email: Simon Scherrer
simon.scherrer@inf.ethz.ch