Resilient distributed routing in dynamical flow networks

Giacomo Como
Department of Automatic Control
Lund University

LCCC workshop on Information and Control in Networks
Lund, October 19, 2012

joint work with: K. Savla, D. Acemoglu, M. Dahleh, and E. Frazzoli at MIT, E. Lovisari at Lund Automatic Control

Resilient infrastructure networks

- good in business as usual, prone to disruptions
- cascade effects

$\Longrightarrow$ network vulnerability $>> \sum$ component vulnerabilities

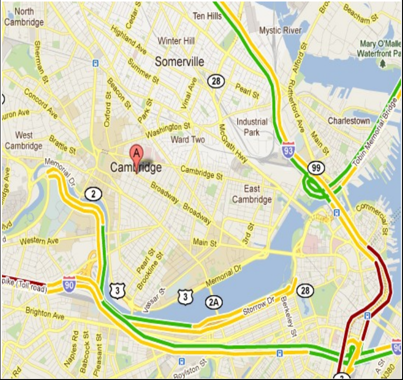

Typical Monday at 18:30

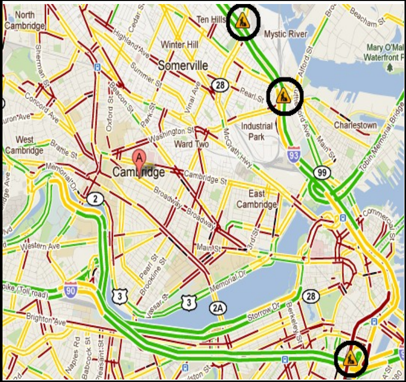

Monday, July 11, 2011, at 18:30

Static flow networks

di-graph $(\mathcal{V}, \mathcal{E})$

origin O

destination D

capacities $C_e > 0$

$\mathcal{E}_v^+ = \{\text{out-links of } v\}$

$\mathcal{E}_v^- = \{\text{in-links of } v\}$

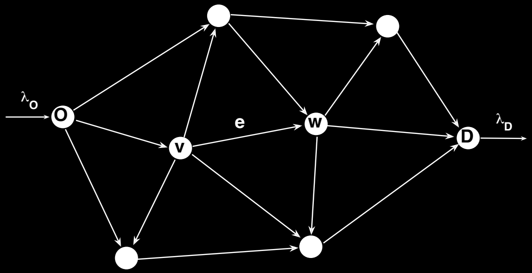

- equilibrium flow: $f = \{f_e\}$ such that $0 \le f_e < C_e$ and

$$\sum_{e \in \mathcal{E}_v^-} f_e = \sum_{e \in \mathcal{E}_v^+} f_e \qquad \forall v \neq O, D$$

Max-flow min-cut theorem

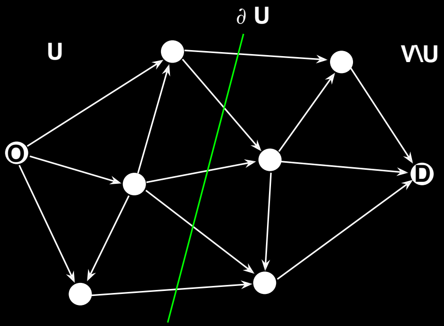

min-cut capacity: $$C := \min_{\text{O-D cut } \mathcal{U}} \sum_{e \in \mathcal{U}^+} C_e$$

$$\exists \text{ equilibrium flow } f : \sum_{e \in \mathcal{E}_O^+} f_e = \lambda_O \quad \Longleftrightarrow \quad C - \lambda_O > 0$$

Max-flow min-cut theorem

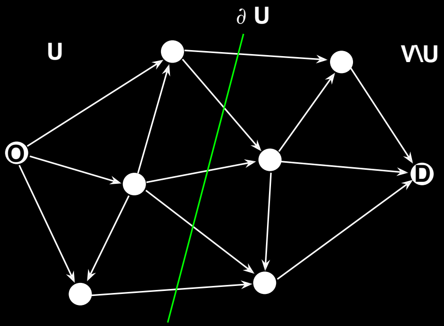

min-cut capacity: $C := \min_{\text{O-D cut } \mathcal{U}} \sum_{e \in \mathcal{U}^+} C_e$

$\exists$ equilibrium flow f : $\sum_{e \in \mathcal{E}_O^+} f_e = \lambda_O \iff C - \lambda_O > 0$

- static, centralized, global information

Dynamical flow networks

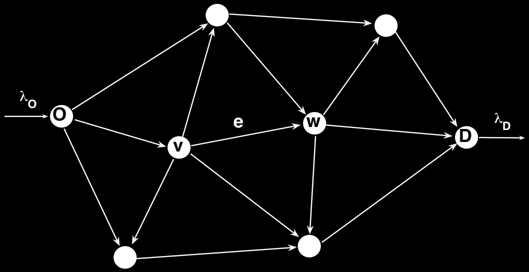

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_e = \lambda_v \; G_e^v - f_e$$

↑ density on v

↑ flow through v

↑ fraction routed to e

↑ outflow of e

Dynamical flow networks

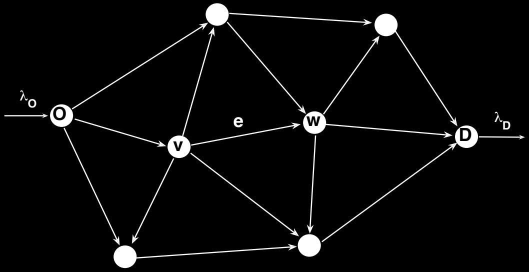

$$\frac{\mathrm{d}}{\mathrm{d}t} \rho_e = \lambda_v \quad G_e^v - f_e$$

(cf. with $\frac{\partial}{\partial t}\rho = -\nabla_x \cdot f$)

Flow function

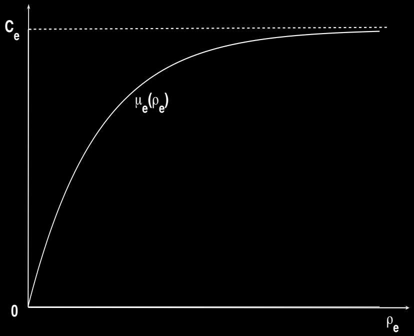

$$\forall e \in \mathcal{E}$$

$$f_e = \mu_e(\rho_e)$$

$$\mu_e(0) = 0 \qquad \sup_{\rho_e \geq 0} \mu_e(\rho_e) = C_e$$

$$\frac{\mathrm{d}}{\mathrm{d}\rho_e}\mu_e > 0$$

Distributed routing

$\lambda_v = \sum_{e \in \mathcal{E}_v^-} f_e$ total flow through v

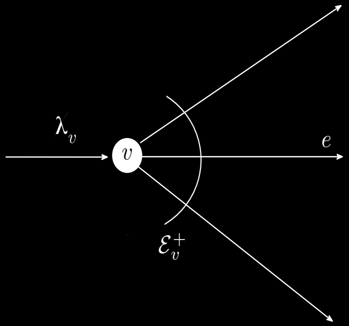

$G_e^v =$ fraction of λ_v routed to e

local information: $\rho^v = \{\rho_e : e \in \mathcal{E}_v^+\}$

$$G^v : \mathbb{R}_+^{\mathcal{E}_v^+} \longrightarrow \mathcal{P}(\mathcal{E}_v^+)$$

↗ distributed routing policy

↑ local info space

↑ probability simplex on $\mathcal{E}_v^+$

Locally responsive distributed policies

$$G^v : \mathbb{R}_+^{\mathcal{E}_v^+} \to \mathcal{P}(\mathcal{E}_v^+)$$

(i) $\frac{\partial}{\partial \rho_j} G_e^v \geq 0 \qquad \forall e \neq j \in \mathcal{E}_v^+$

(ii) $\rho_e \to \infty \quad \Rightarrow \quad G_e^v(\rho^v) \to 0$

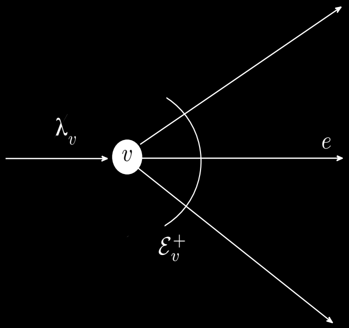

Locally responsive distributed policies

$$G^v : \mathbb{R}_+^{\mathcal{E}_v^+} \to \mathcal{P}(\mathcal{E}_v^+)$$

(i) $\frac{\partial}{\partial \rho_j} G_e^v \geq 0 \qquad \forall e \neq j \in \mathcal{E}_v^+$

(ii) $\rho_e \to \infty \quad \Rightarrow \quad G_e^v(\rho^v) \to 0$

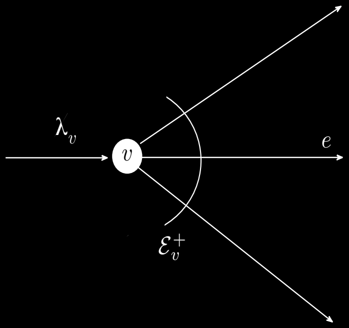

Ex.: i-logit

$$G_e^v(\rho^v) \propto \alpha_e \exp(-\beta \rho_e) \qquad \beta > 0,\ \alpha \in \mathbb{R}_+^{\mathcal{E}_v^+}$$

Locally responsive distributed policies

$$G^v : \mathbb{R}_+^{\mathcal{E}_v^+} \to \mathcal{P}(\mathcal{E}_v^+)$$

(i) $\frac{\partial}{\partial \rho_j} G_e^v \geq 0 \qquad \forall e \neq j \in \mathcal{E}_v^+$

(ii) $\rho_e \to \infty \quad \Rightarrow \quad G_e^v(\rho^v) \to 0$

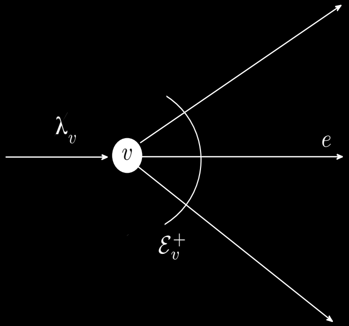

- (i) $\Rightarrow$ network is monotone (i.e., Jacobian is Metzler) $\Rightarrow \ldots$
  $\Rightarrow \ldots || \ ||_1$ contraction $\Rightarrow \lim_{t\to\infty} \lambda_D(t)$ independent of $\rho(0)$

Dynamical flow networks

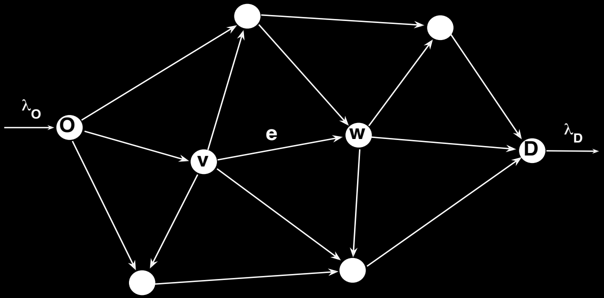

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_e = \lambda_v G_e^v(\rho^v) - \mu_e(\rho_e) \qquad e \in \mathcal{E}_v^+,\ v \in \mathcal{V} \setminus \{D\}$$

transferring if $\lim_{t\to\infty} \lambda_D(t) = \lambda_O$

Perturbed dynamical flow networks

Perturbation : $\forall e \in \mathcal{E}$

$$\mu_e(\rho_e)$$

$$\downarrow$$

$$\tilde{\mu}_e(\rho_e) \leq \mu_e(\rho_e)$$

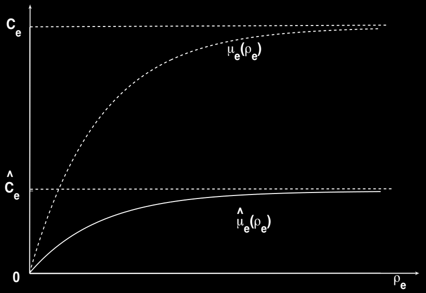

Magnitude $\delta := \sum_e \delta_e = \sum_e ||\mu_e(\cdot) - \tilde{\mu}_e(\cdot)||_\infty$

Margin of resilience

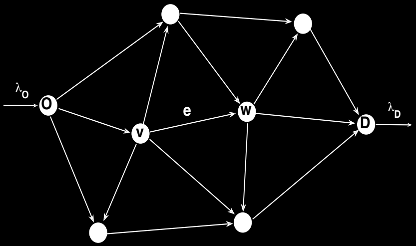

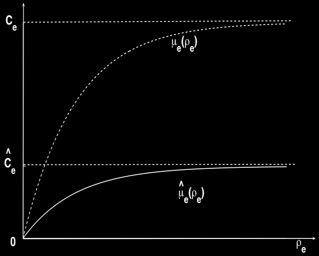

perturbed network: $$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{\rho}_e = \tilde{\lambda}_v G_e^v(\tilde{\rho}^v) - \tilde{\mu}_e(\tilde{\rho}_e) \qquad e \in \mathcal{E}_v^+, \; v \neq D$$

$$\gamma := \text{ inf magnitude of disruption s.t. } \lim_{t \to +\infty} \tilde{\lambda}_D(t) < \lambda_O$$

Optimal resilience

residual node capacity

for limit flow f^*

$$R(f^*) := \min_{v \neq D} \sum_{e \in \mathcal{E}_v^+} C_e - f_e^*$$

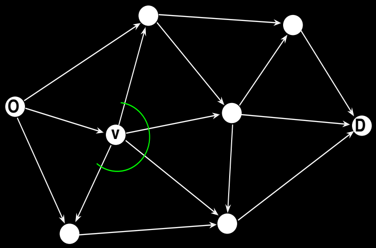

Optimal resilience

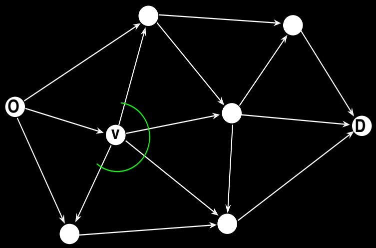

$$R(f^*) := \min_{v \neq D} \sum_{e \in \mathcal{E}_v^+} C_e - f_e^*$$

Theorem: acyclic network, limit flow f^*

1. distributed routing

$$\gamma \leq R(f^*)$$

2. locally responsive routing $\Rightarrow$ f^* globally attractive

$$\gamma = R(f^*)$$

Optimal resilience (cont'd)

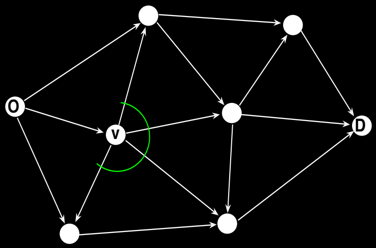

$$R(f^*) := \min_{v \neq D} \sum_{e \in \mathcal{E}_v^+} C_e - f_e^*$$

- $1 \leq \dfrac{C - \lambda_O}{R(f^*)}$ can be arbitrarily large

local information constraint $\implies$ resilience loss

Optimal resilience (cont'd)

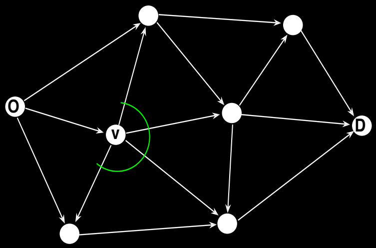

$$R(f^*) := \min_{v \neq D} \sum_{e \in \mathcal{E}_v^+} C_e - f_e^*$$

- perturbations and distributed routing effect only downstream
- locally responsive policies

$\Downarrow$

locally optimal load rebalance

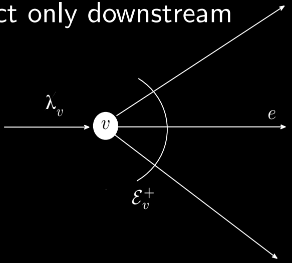

Bounded density capacities

$$0 < \rho_e^{\max} < +\infty$$

$$\mu_e(0) = 0 \qquad \mu_e(\rho_e^-) = C_e$$

$$\mu_e'(\rho_e) > 0 \qquad \forall \rho_e \in (0, \rho_e^{\max})$$

$$\mu_e(\rho_e) = 0 \qquad \forall \rho_e \geq C_e$$

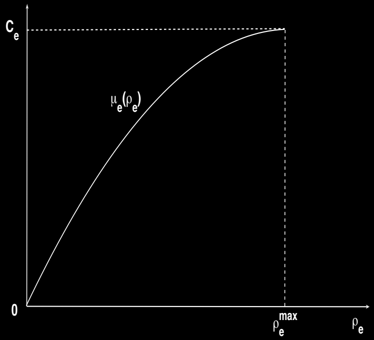

Bounded density capacities

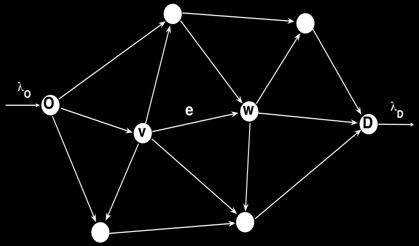

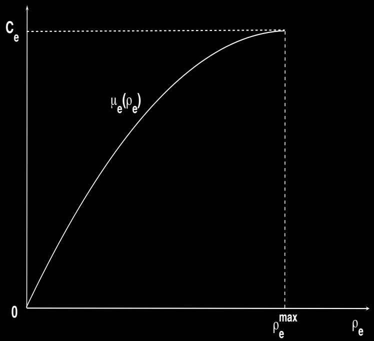

$$e = (v, w) \Rightarrow \qquad \frac{\mathrm{d}}{\mathrm{d}t}\rho_e = \chi_v(t)\lambda_v G_e^v(\rho^v) - \chi_w(t)\mu_e(\rho_e)$$

$$\chi_v(t) := 1 - \prod\nolimits_{e \in \mathcal{E}_v^+}(1 - \xi_e(t)) \qquad \xi_e(t) := \mathbb{1}_{[0, \rho_e^{\max})}(\rho_e)$$

Cascade propagation mechanism

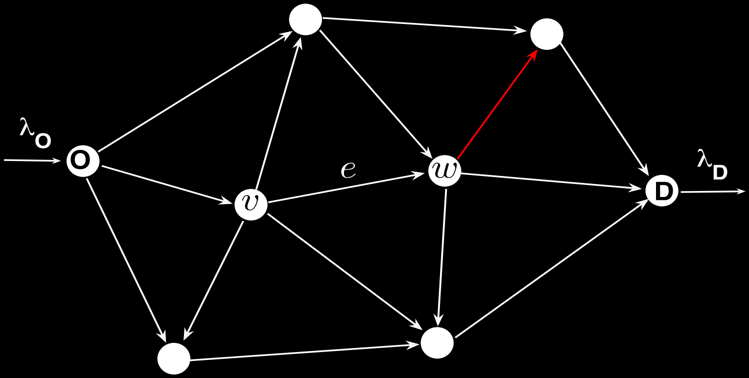

Perturbation localized on $\mathcal{E}_w^+$, such that $\sum_{\mathcal{E}_w^+} \tilde{C}_e < \lambda_w^*$

Cascade propagation mechanism

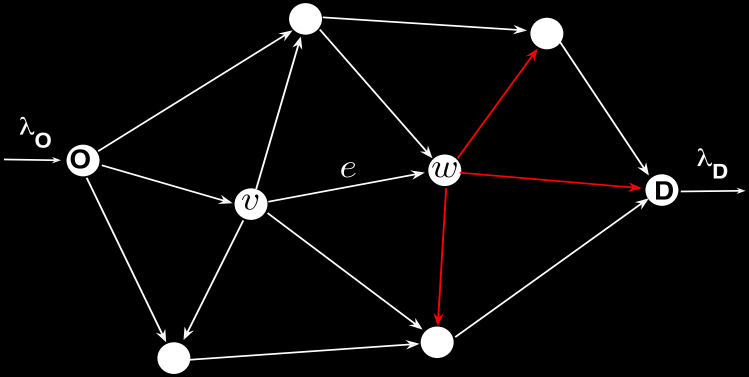

Node w overloaded $\Rightarrow$ every $e \in \mathcal{E}_w^+$ drops

Cascade propagation mechanism

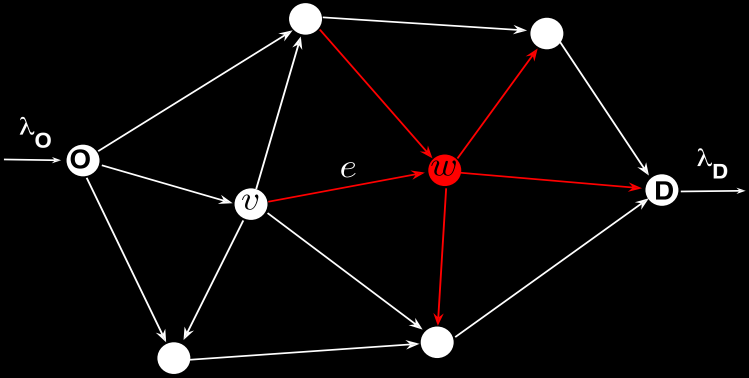

Hence w drops together with all $e \in \mathcal{E}_w^-$

Cascade propagation mechanism

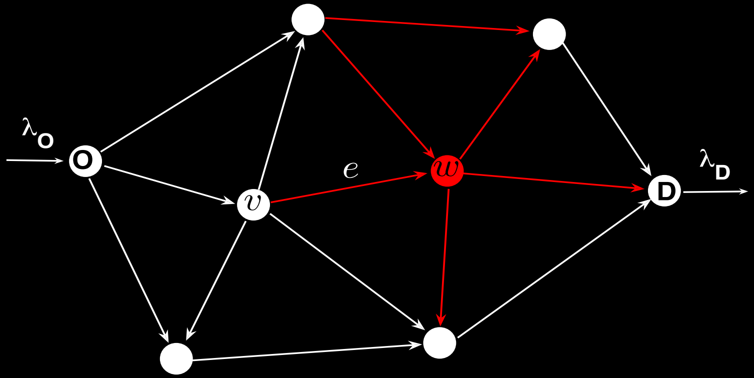

Possibly other nodes downstream become overloaded...

Cascade propagation mechanism

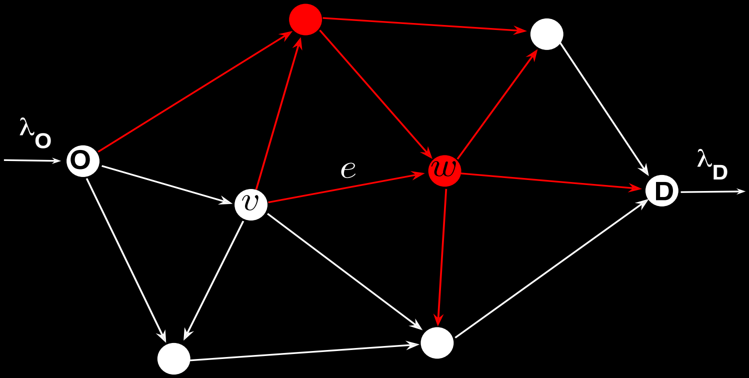

...and eventually drop

Cascade propagation mechanism

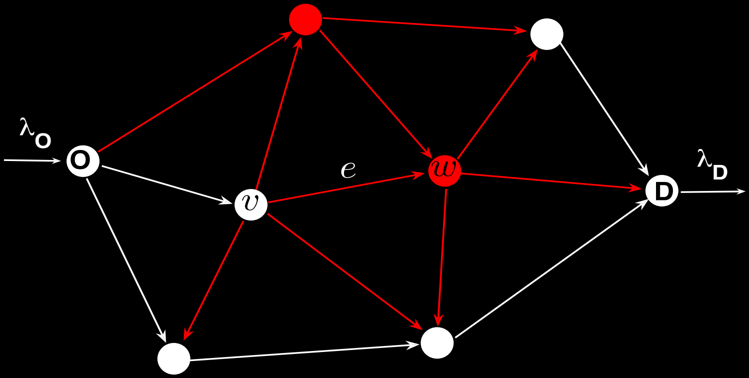

Either a new equilibrium flow is eventually achieved

Cascade propagation mechanism

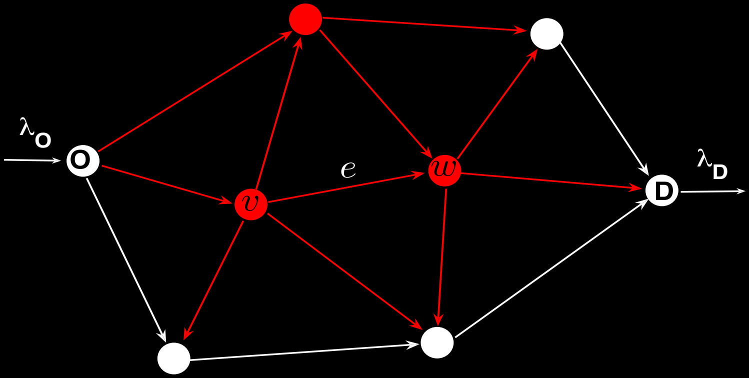

...or not.

Cascade propagation mechanism

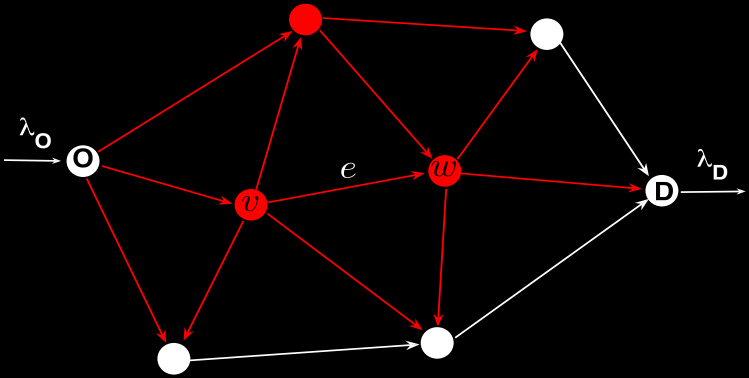

If the perturbation overloads the origin...

Cascade propagation mechanism

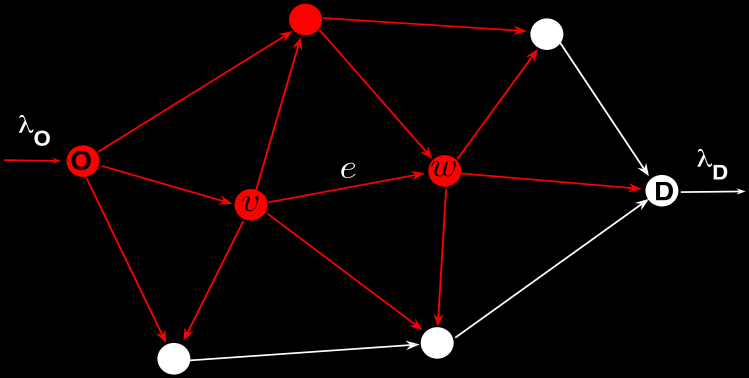

… no more flow passes through the network.

Dynamical flow networks with bounded density

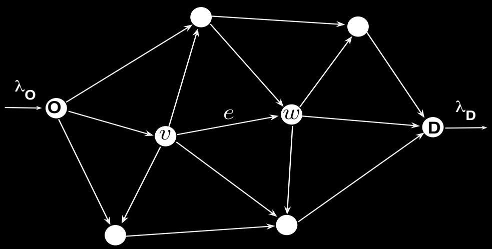

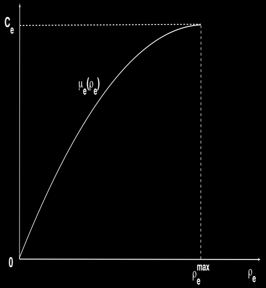

Theorem: $\rho_e^{\max} < +\infty \quad \forall e \quad \Longrightarrow \quad \lim_{t\to\infty} \lambda_D(t) \in \{0, \lambda_O\}$

1 acyclic network, any distributed routing

$$\gamma \leq \Gamma(f^*) \qquad \Gamma(f^*) := \dots \text{ backward induction}$$

2 locally responsive routing, $d_v^+ \leq 2 \quad \Longrightarrow \quad \gamma = \Gamma(f^*)$

3 locally responsive routing $\quad \Longrightarrow \quad \gamma \geq R(f^*)$

Dynamical flow networks with bounded density

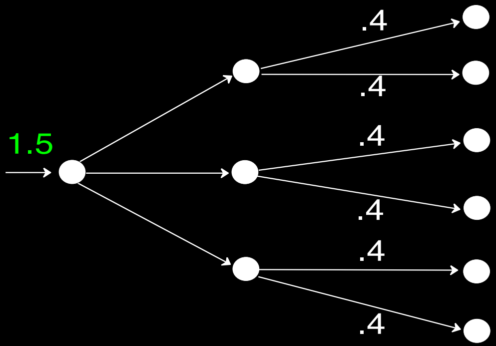

Dynamical flow networks with bounded density

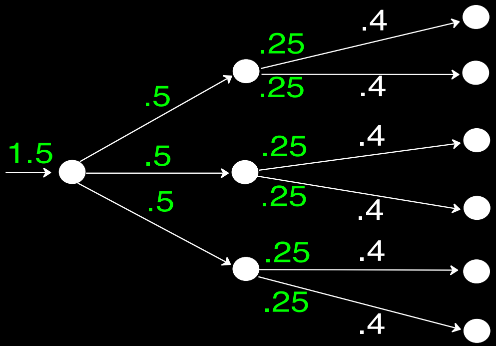

Dynamical flow networks with bounded density

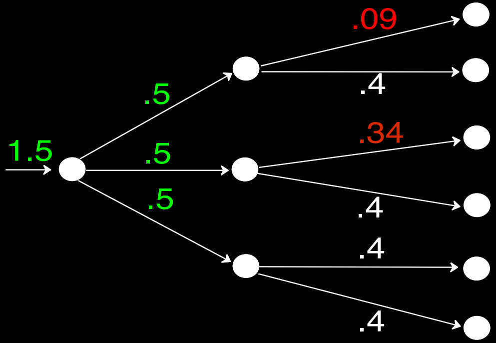

Dynamical flow networks with bounded density

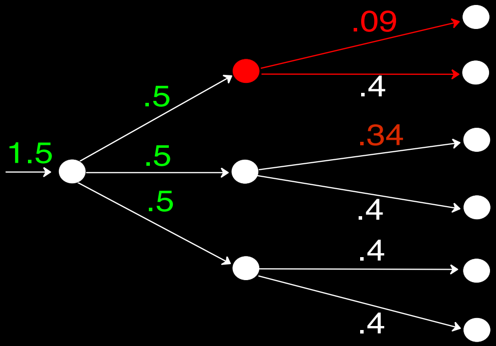

Dynamical flow networks with bounded density

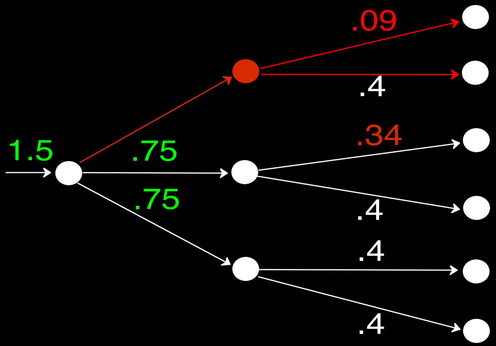

Dynamical flow networks with bounded density

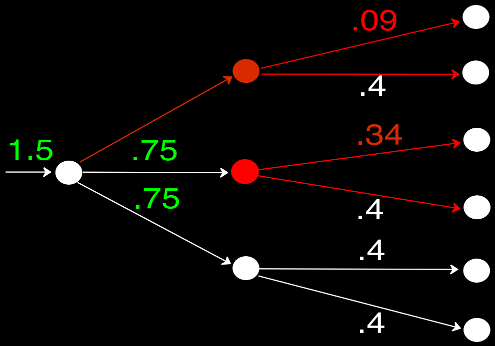

Dynamical flow networks with bounded density

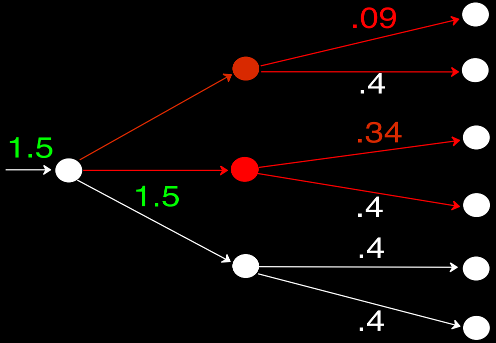

Dynamical flow networks with bounded density

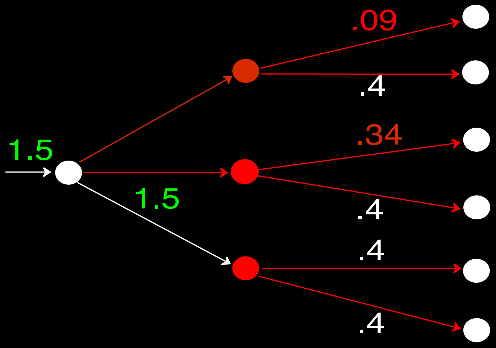

Dynamical flow networks with bounded density

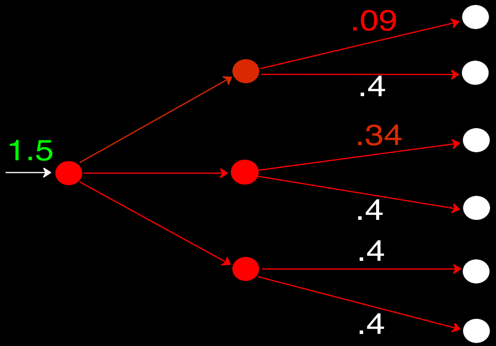

Slowing down local flow

can reduce flow through v

$j \in \mathcal{E}_v^-$

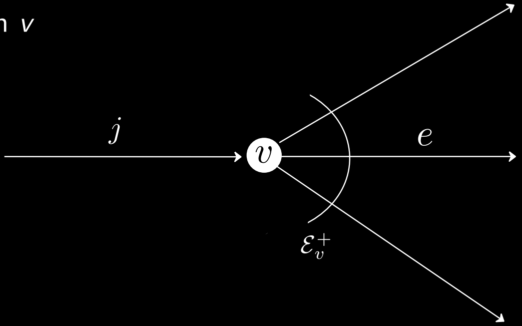

$$G^j \quad : \quad \mathbb{R}_+^{\{j\} \cup \mathcal{E}_v^+} \quad \longrightarrow \quad \mathcal{P}(\{j\} \cup \mathcal{E}_v^+)$$

$G_j^j(\rho_j, \rho^v) :=$ fraction of flow kept on j

$1 - G_j^j(\rho_j, \rho^v) :=$ fraction of flow allowed through v

Slowing down achieves capacity

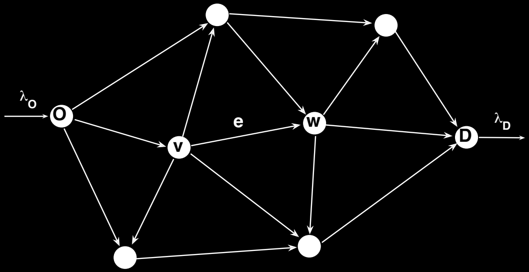

Theorem: Acyclic network, locally responsive policy

$$\gamma = C - \lambda_O$$

both with finite and infinite ρ_e^{max}

Dynamical flow networks with bounded density

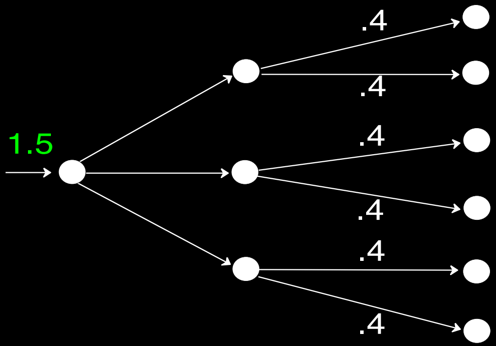

Dynamical flow networks with bounded density

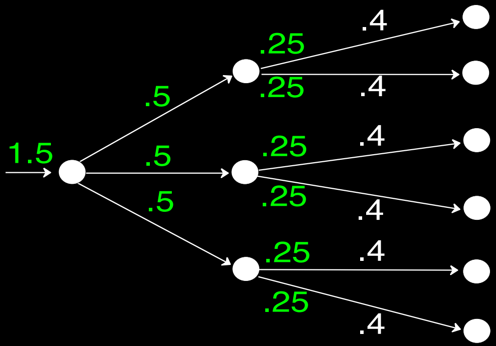

Dynamical flow networks with bounded density

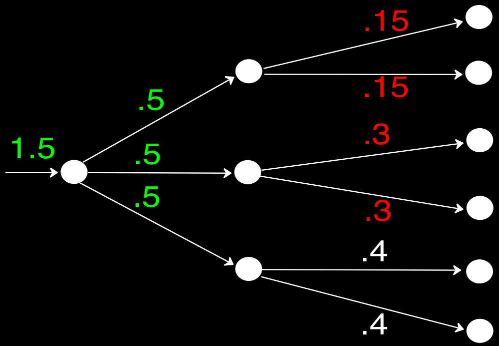

Dynamical flow networks with bounded density

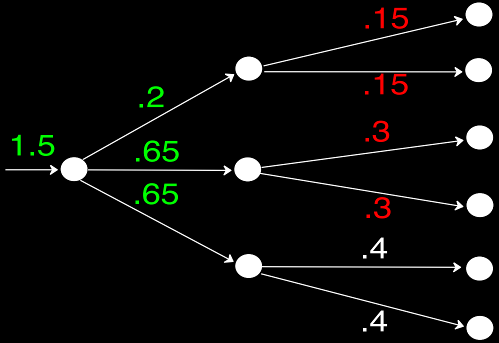

Dynamical flow networks with bounded density

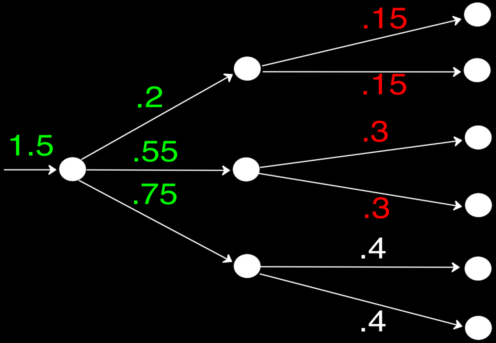

Conclusion

Summary

- robust distributed routing for dynamical flow networks
- margins of resilience
- role of cascades

Current/future directions

- cycles
- multicommodity flows
- value of communication
- resilience of networks with other dynamics