Internet Routing Games

João P. Hespanha

Center for Control Dynamical Systems and Computation

University of California
Santa Barbara

UCSB

In collaboration with: S. Bohacek (Univ. Delaware), K. Obraczka (UC Santa Cruz)
J. Lee (Postdoc, UC Santa Barbara), C. Lim (PhD candidate, USC)

Network Security vs. Fault-Tolerance

UCSB

The basic principle behind the design of the Internet was to utilize massive redundancy to achieve fault-tolerance

but this does not necessarily result in *security against malicious attacks*

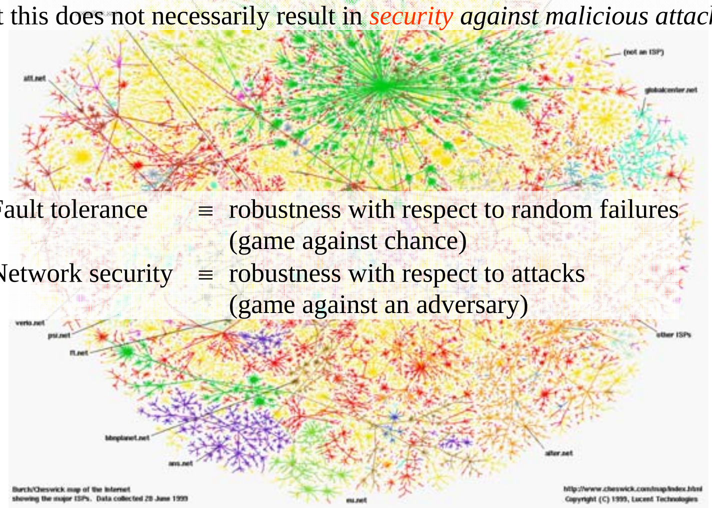

Fault tolerance ≡ robustness with respect to random failures
(game against chance)

Network security ≡ robustness with respect to attacks
(game against an adversary)

An adversary can explore weaknesses that chance will not easily find

Security vs. Fault-Tolerance in Routing

single-path routing — stochastic multi-path routing

Suppose all links are equally likely to fail, and one of them does fail…

Which routing strategy results in higher probability that a packet will reach destination?

link labels refer to probability of forwarding a packet

Both routing schemes result in exactly the same probability (50%)…

Security vs. Fault-Tolerance in Routing

single-path routing — stochastic multi-path routing

~~Suppose all links are equally likely to fail~~, and one of them does fail…

Assume that fail was caused by an attacker that selects the link

Which routing strategy results in higher probability that a packet will reach destination?

Attacker can learn routing policy and prevent all communication by compromising a single link

Compromising a single link, probability of intercepting packet is only 50% (assuming stochastic multi-path)

later we will find other reasons why multi-path may be advantageous…

Outline

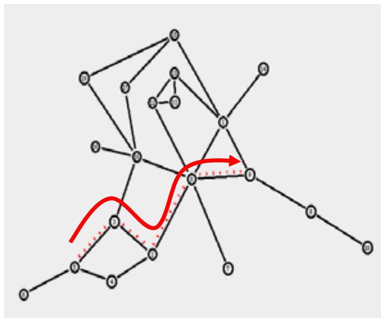

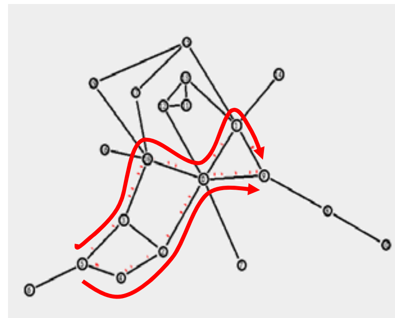

1. How to compute stochastic multi-path routing tables for general networks?
   Noncooperative game—explore redundancy in an adversarial context
2. Multi-path routing for multi-agent & networked control systems

Stochastic routing policies

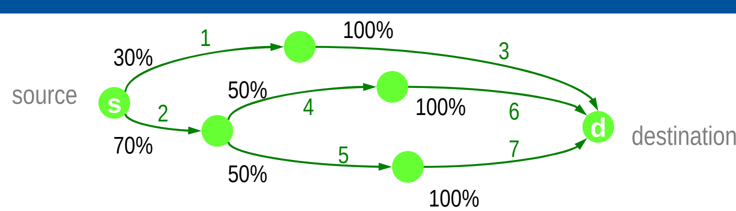

probability that a packet arriving at the node where ℓ starts will be routed though link ℓ

set of (unidirectional) links

stochastic routing policy $\equiv R := \{ r_\ell \geq 0 : \ell \in \mathcal{L} \}$

for every node n $\sum_{\ell \in \mathcal{L}[n]} r_\ell = 1$

summation over links that exit node n

e.g., $R := \{ .3, .7, 1, .5, .5, 1, 1\}$

add to 1 add to 1

$\mathcal{R}_{\text{stoch}}$ $\equiv$ set of all routing policies

Stochastic routing policies

UCSB

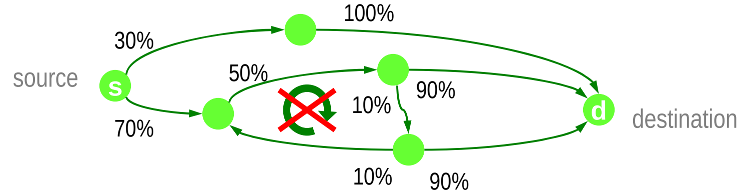

probability that a packet arriving at the node where ℓ starts will be routed though link ℓ

set of (unidirectional) links

stochastic routing policy $\equiv R := \{ r_\ell \geq 0 : \ell \in \mathcal{L} \}$

for every node n $\displaystyle\sum_{\ell \in \mathcal{L}[n]} r_\ell = 1$

summation over links that exit node n

$\mathcal{R}_{\text{stoch}}$ $\equiv$ set of all routing policies

$\mathcal{R}_{\text{no-cycle}}$ $\equiv$ set of all cycle-free policies, i.e., for which there is no closed sequence of links all with positive routing probability

Attack space

UCSB

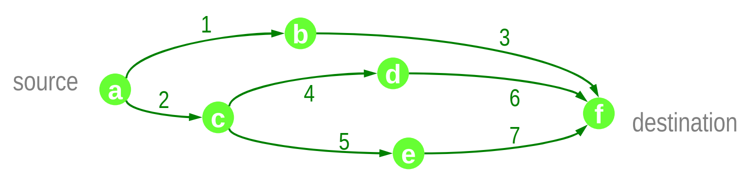

probability that packets in link ℓ are compromised

set of links

pure attack $\equiv$ $P := \{ p_\ell : \ell \in \mathcal{L} \}$

attacker has available a pool of "pure attacks" and will select the one that is more likely to prevent communication

e.g., pure attack at link 3 with 10% probability of success:

$P_3 := \{ 0, 0, .1, 0, 0, 0, 0 \}$

pure attack at node f with 20% probability of success:

$P_f := \{ 0, 0, .2, 0, 0, .2, .2 \}$

$\mathcal{P} \equiv$ set of all (pure) attacks available to attacker

Mixed attacks

UCSB

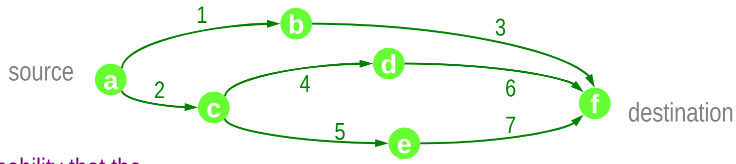

probability that the attacker will intercept a packet traveling in link ℓ

set of links

attacker is allowed to randomize between pure attacks with appropriate probabilities

(pure) attack $\equiv \quad P := \{ \, p_\ell : \ell \in \mathcal{L} \, \}$

$\mathcal{P} \equiv$ set of all pure attacks available to attacker

mixed attack policy $\equiv M := \{ \, m_P : P \in \mathcal{P} \, \} \in [0,1]^{\mathcal{P}}$

$$\sum_{P \in \mathcal{P}} m_P = 1$$

probability that the attacker select the pure attack P

Example

UCSB

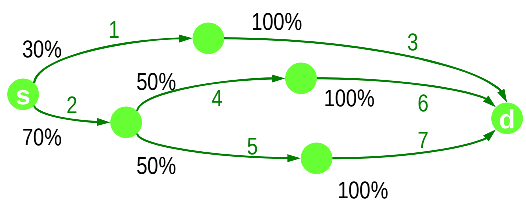

stochastic routing policy $\equiv R := \{ \, .3, \, .7, \, 1, \, .5, \, .5, \, 1, \, 1 \}$

pure attacks available to attacker $\equiv \mathcal{P} := \{ \, \{.1,0,0,0,0,0,0\},\{0,.1,0,0,0,0,0\}, \ldots$
$\ldots,\{0,0,0,0,0,.1,0\}, \, \{0,0,0,0,0,0,.1\} \, \}$

10% effective link attacks (7 attacks)
(attacker can target any link, it will succeed in compromising packet delivery with 10% probability)

Example

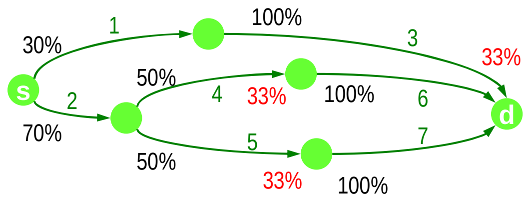

stochastic routing policy $\equiv R \coloneqq \{ .3, .7, 1, .5, .5, 1, 1\}$

pure attacks available to attacker $\equiv \mathcal{P} \coloneqq \{ \{.1,0,0,0,0,0,0\},\{0,.1,0,0,0,0,0\}, \ldots$
$\ldots,\{0,0,0,0,0,.1,0\}, \{0,0,0,0,0,0,.1\} \}$

mixed attack policy $\equiv M \coloneqq \{ 0, 0, .33, .33, .33, 0, 0 \}$

$$\mathrm{P}_{R,M}(\text{capture}) = \big(30\% \times 33\% + 2 \times 35\% \times 33\%\big) \times 10\% = 3.3\%$$

probability that packet is captured for routing policy R and mixed attack policy M

but not really rational…

Example

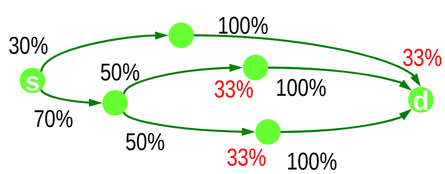

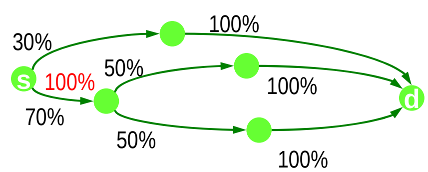

mixed attack policy
$M \coloneqq \{ 0, 0, .33, .33, .33, 0, 0 \}$

stochastic routing policy
$R \coloneqq \{ .3, .7, 1, .5, .5, 1, 1\}$

$$\mathrm{P}_{R,M}(\text{capture}) = \big(30\% \times 33\% + 2 \times 35\% \times 33\%\big) \times 10\% = 3.3\%$$

for this attack policy M
router cannot do better

$$\mathrm{P}_{R,M}(\text{capture}) = 3.3\% \quad \forall R$$

but attacker could do better against R with
$M \coloneqq \{ 0, 1, 0, 0, 0, 0, 0 \}$

$$\mathrm{P}_{R,M}(\text{capture}) = \big(70\% \times 100\%\big) \times 10\% = 7\%$$

but then …

neither of the above policies is an "equilibrium" since
at least one player can improve its outcome by changing its policy

Routing game

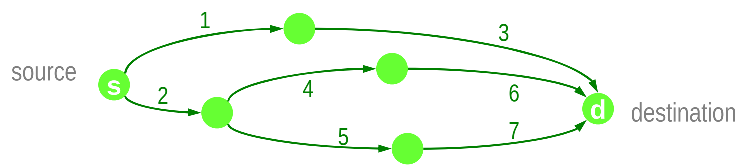

Compute saddle-point equilibrium policies:

$R^* \in \mathcal{R}_{\text{no-cycle}}$ (cycle-free stochastic routing policy)

$M^* \in [0,1]^{\mathcal{P}}$ (mixed attack policy)

for which

$$\mathrm{P}_{R^*,M^*}(\text{capture}) = \min_{R\in\mathcal{R}_{\text{no-cycle}}} \max_{M\in[0,1]^{\mathcal{P}}} \mathrm{P}_{R,M}(\text{capture})$$
$$= \max_{M\in[0,1]^{\mathcal{P}}} \min_{R\in\mathcal{R}_{\text{no-cycle}}} \mathrm{P}_{R,M}(\text{capture})$$

Existence?
Computation?

policies chosen by intelligent opponents to minimize their worst-case losses
(no player will improve its outcome by deviating from equilibrium)

Probability of capture

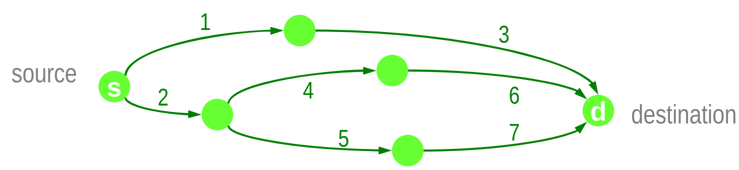

Given

$R \in \mathcal{R}_{\text{no-cycle}}$ (cycle-free stochastic routing policy)

$M := \{ m_P : P \in \mathcal{P} \} \in [0,1]^{\mathcal{P}}$ (mixed attack policy)

$$\mathrm{P}_{R,M}(\text{capture}) = \sum_{P\in\mathcal{P}} \Big(m_P \, \text{row}[P] \, x_P \Big)$$

diagonal matrix with all the elements of R

row vector with all the pure policies p_ℓ

unique solution to

$$x_P = \text{diag}[R]A(I - \text{diag}[P])x_P + \text{diag}[R]c$$

(matrix A and vector c only depend on the graph)

Linear (thus concave) in M (maximizer)
but not convex with respect to the routing policy R (minimizer)
so mini-max existence theorems do not apply…

Probability of capture

UCSB

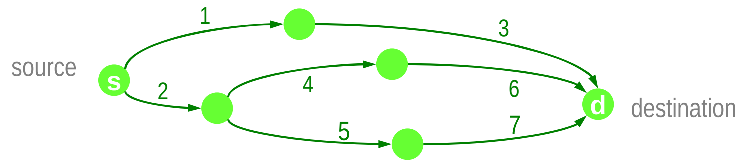

Under mild assumptions (*) on pure attacks

Given $R \in \mathcal{R}_{\text{no-cycle}}$, $M := \{ m_P : P \in \mathcal{P} \} \in [0,1]^{\mathcal{P}}$

$$\mathrm{P}_{R,M}(\text{capture}) = \Big(\sum_{P \in \mathcal{P}} m_P \, \mathrm{row}[P] \Big) x$$

row vector with all the pure policies p_ℓ

flow vector ≡ unique solution to

$$x = \mathrm{diag}[R] A x + \mathrm{diag}[R] c$$

(matrix A and vector c only depend on the graph)

(*) the same pure attack does not simultaneously targets two links in the same path (true for every single-link or single-node attacks)

Probability of capture

UCSB

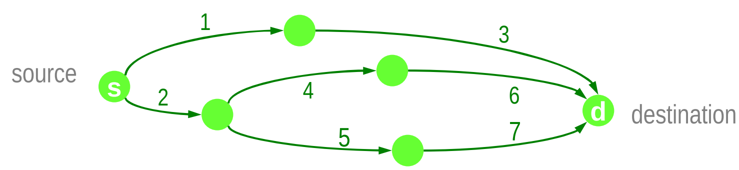

Under mild assumptions (*) on pure attacks

Given $R \in \mathcal{R}_{\text{no-cycle}}$, $M := \{ m_P : P \in \mathcal{P} \} \in [0,1]^{\mathcal{P}}$

$$\mathrm{P}_{R,M}(\text{capture}) = \Big(\sum_{P \in \mathcal{P}} m_P \, \mathrm{row}[P] \Big) x$$

row vector with all the pure policies p_ℓ

flow vector ≡ unique solution to

$$x = \mathrm{diag}[R] A x + \mathrm{diag}[R] c$$

(matrix A and vector c only depend on the graph)

Not convex with respect to the routing policy R but linear (convex!) with respect to the vector x…
Key idea: solve game for x & then compute R

Routing policies & Flow vectors

Theorem: i) There is a one-to-one correspondence between routing policies R in $\mathcal{R}_{\text{stoch}}$ & flow vectors x in a convex set $\mathcal{X} \subset \mathbb{R}^{\mathcal{L}}$

ii) For cycle-free $R \in \mathcal{R}_{\text{no-cycle}}$, the corresponding flow vector x satisfies

$$x = \text{diag}[R]Ax + \text{diag}[R]c$$

Therefore

$$\mathrm{P}_{R,M}(\text{capture}) = \sum_{P \in \mathcal{P}} m_P \,\text{row}[P]x$$

Routing policies & Flow vectors

Theorem: i) There is a one-to-one correspondence between routing policies R in $\mathcal{R}_{\text{stoch}}$ & flow vectors x in a convex set $\mathcal{X} \subset \mathbb{R}^{\mathcal{L}}$

ii) For cycle-free $R \in \mathcal{R}_{\text{no-cycle}}$, the corresponding flow vector x satisfies

$$x = \text{diag}[R]Ax + \text{diag}[R]c$$

stochastic routing policy — flow vector

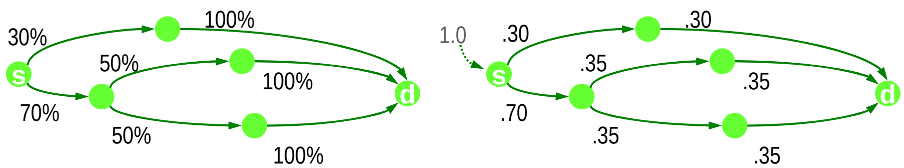

stochastic routing policy
R := { .3, .7, 1, .5, .5, 1, 1}

flow vector
x := { .3, .7, .3, .35, .35, .35, .35}

the vectors $x \in \mathcal{X}$ obey a "flow conservation law" at every node, with total unit flow exiting the source node

Flow game

flow vector

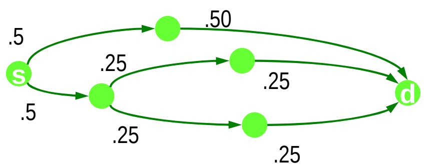

Compute saddle-point:

$x^* \in \mathcal{X}$ (flow vector)

$M^* \in [0,1]^{\mathcal{P}}$ (mixed attack policy)

for which

$$\sum_{P\in\mathcal{P}} m_P^* \operatorname{row}[P]x^* = \min_{x\in\mathcal{X}} \max_{M\in[0,1]^{\mathcal{P}}} \sum_{P\in\mathcal{P}} m_P \operatorname{row}[P]x$$

$$= \max_{M\in[0,1]^{\mathcal{P}}} \min_{x\in\mathcal{X}} \sum_{P\in\mathcal{P}} m_P \operatorname{row}[P]x.$$

Theorem: Every flow game has a saddle-point (x^*, M^*) with x^* cycle-free

by bilinearity of the criterion and
convexity and (almost) compactness of $\mathcal{X}$ & $[0,1]^{\mathcal{P}}$

Back to routing game…

flow vector

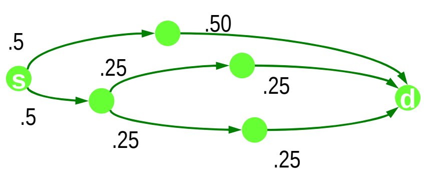

stochastic routing policy

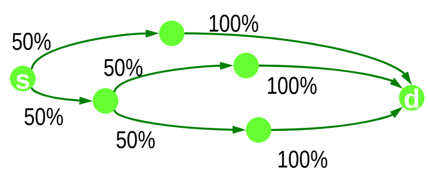

Theorem: The routing game has saddle-point policies.

Moreover, for every saddle-point (x^*, M^*) of the flow game with x^* cycle-free, the pair (R^*, M^*) is a saddle-point of the routing game, with R^* constructed from x^*:

$$r_\ell^* := \frac{x_\ell^*}{\sum_{\ell'\in\mathcal{L}[\ell]} x_{\ell'}^*} \qquad \forall \ell \in \mathcal{L}$$

summation over all links that exit
from the same node as ℓ

Solving the flow game actually solves the routing game…

Solution to the flow & routing games

flow vector

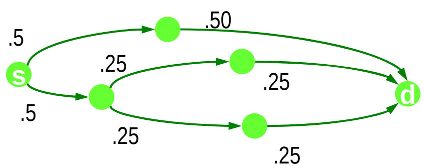

stochastic routing policy

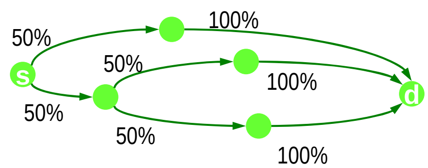

Theorem: The value V^* of the flow game is given by

$$V^* = \min_{\substack{x\in\mathcal{X} \\ \mathrm{row}[P]x\le\mu,\ \forall P}} \mu$$

max-flow problem solvable by linear programming

and the saddle-point x^* is any x at which the minimum is attained.

Optimal routing policy R^* can be computed using:

$$r^*_\ell := \frac{x^*_\ell}{\sum_{\ell'\in\mathcal{L}[\ell]} x^*_{\ell'}} \qquad \forall \ell \in \mathcal{L}$$

Max-flow interpretations

for pure attacks at individual links

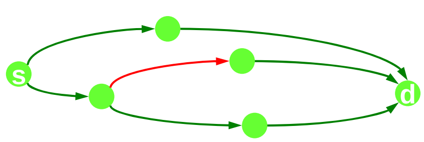

- Optimal routing *minimizes the maximum link flow* (subject to constraints that depend on the link reliability)
- In practice, maximizes throughput subject to link bandwidth constraints

for pure attacks at individual nodes

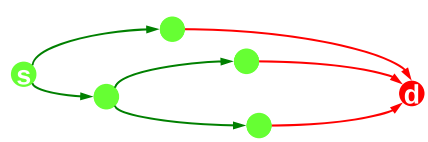

- Optimal routing *minimizes the maximum node load* (subject to constraints that depend on node reliability)
- In practice, balances the load between nodes (useful for energy-starved nodes)

Several reasons to use multi-path routing

UCSB

increase security

- Hespanha, Bohacek. Preliminary Results in Routing Games, 2001.
- Bohacek, Hespanha, Lee, Obraczka, Lim, Enhancing security via stochastic routing, 2002
- Papadimitratos, Haas, Secure message transmission in mobile ad hoc networks, 2003
- Lee, Misra, Rubenstein, Distributed Algorithms for Secure Multipath Routing, 2005

improve robustness

- Ganesan, Govindan, Shenker, Estrin, Highly Resilient, Energy Efficient Multipath Routing in Wireless Sensor Networks, 2002
- Wei, Zakhor, Robust Multipath Source Routing Protocol (RMPSR) for Video Communication over Wireless Ad Hoc Networks, 2004
- Tang, McKinley, A distributed multipath computation framework for overlay network applications, 2004

increase throughput

- Chen, Chan, Li, Multipath routing for video delivery over bandwidth-limited networks, 2004

maximize network utilization

- Elwalid, Jin, Low, Widjaja, MATE: MPLS adaptive traffic engineering, 2001
- Lee, Gerla, Split multipath routing with maximally disjoint paths in ad hoc networks, 2001
- Mirrokni, Thottan, Uzunalioglu, Paul, Simple polynomial time frameworks for reduced-path decomposition in multi-path routing, 2004

Estimation through network

UCSB

process

$$x_{k+1} = Ax_k + Bw_k$$

zero-mean stochastic disturbance

remote state-estimator

$$\hat{x}_{k+1} = A\hat{x}_k$$

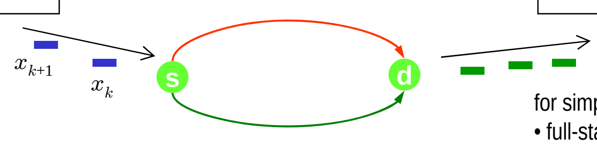

for simplicity:
- full-state available
- no measurement noise
- no quantization

Optimal remote state estimator:

$$\hat{x}_{k+1} = \begin{cases} Ax_k & \text{succ. transmission at time } k \\ A\hat{x}_k & \text{unsucc. transmission at time } k \end{cases}$$

Remote state estimation error: $e_k := x_k - \hat{x}_k$

$$e_{k+1} = \begin{cases} Bw_k & \text{succ. transmission at time } k \\ Ae_k + Bw_k & \text{unsucc. transmission at time } k \end{cases}$$

Estimation through network

UCSB

process

$x_{k+1} = Ax_k + Bw_k$

zero-mean stochastic disturbance

remote state-estimator

$\hat{x}_{k+1} = A\hat{x}_k$

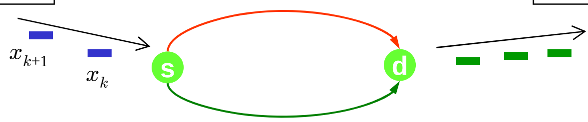

$$e_{k+1} = \begin{cases} Bw_k & \text{succ. transmission at time } k \\ Ae_k + Bw_k & \text{unsucc. transmission at time } k \end{cases}$$

Failures caused by an attacker that succeeds with probability p_{att}

- single-path routing $\equiv$ probability of failed transmission = p_{att}

$$e_{k+1} = \begin{cases} Bw_k & \text{w.p. } 1 - p_{\text{att}} \\ Ae_k + Bw_k & \text{w.p. } p_{\text{att}} \end{cases}$$

mean-square stable iff $p_{\text{att}} < \frac{1}{|\lambda_i[A]|^2}$

- multi-path routing $\equiv$ probability of failed transmissions = $p_{\text{att}}/2$

$$e_{k+1} = \begin{cases} Bw_k & \text{w.p. } 1 - \frac{p_{\text{att}}}{2} \\ Ae_k + Bw_k & \text{w.p. } \frac{p_{\text{att}}}{2} \end{cases}$$

mean-square stable iff $p_{\text{att}} < \frac{2}{|\lambda_i[A]|^2}$

Estimation through network

UCSB

process

$x_{k+1} = Ax_k + Bw_k$

stochastic disturbance

remote state-estimator

$\hat{x}_{k+1} = A\hat{x}_k$

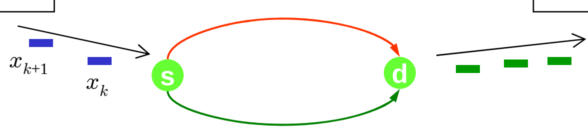

Consider random failures:

$p_{\text{fail}} \equiv$ probability that a link will fail
$T_{\text{ttr}} \equiv$ mean time-to-recover (exponentially distributed)

- single-path routing
- multi-path routing

At steady state:

$$\text{P(unsucc. transmission)} = \frac{p_{\text{fail}} T_{\text{ttr}}}{1 + p_{\text{fail}} T_{\text{ttr}}}$$

but drops are not i.i.d. …

Estimation through network

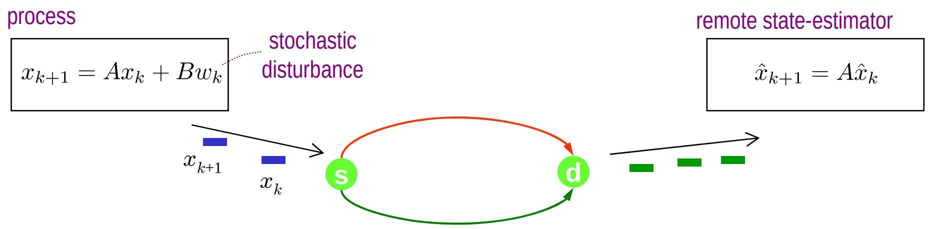

Consider random failures:

$p_{\text{fail}} \equiv$ probability that a link will fail
$T_{\text{ttr}} \equiv$ mean time-to-recover (exponentially distributed)

For: 1-dimensional quasi-stable process $A = 1 + \epsilon$, $\epsilon \ll 1$
low fail probability $p_{\text{fail}} \ll 1$

- single-path routing — mean-square stable iff $T_{\text{ttr}} \leq \frac{1}{2\epsilon}$
- multi-path routing — mean-square stable iff $T_{\text{ttr}} \leq \frac{1}{\epsilon}$

twice as large admissible mean time-to-recover

in this networked estimation problem, the maximum spread of packets is optimal even "against" random failures

Conclusions

- Communication networks are extremely vulnerable components to critical systems
  – multitude of individual components, spatially distributed, difficult to protect
  – especially true for wireless networks (jamming, eavesdropping, battery drainage due to overuse, etc.)
- Game theory is a natural framework to study robustness
  – redundancy, by itself, does not guarantee robustness
  – attacks are not random events: very unlikely events can be prompted by an attacker
- Determined routing polices that exploit multi-path routing
  – formulation as a zero-sum game between router and attacker
  – saddle-point solutions found by reducing problem to a flow-game
  – policies found also have applications to
    - throughput maximization
    - load balancing
    - improve robustness of NCSs (even against random failures)

[Observation: traditional measures of QoS such as probability of drop, expected delay are not sufficient to predict performance in NCSs]