Robustness of Complex Networks

Piet Van Mieghem
Network Architecture and Services (NAS)

TU Delft

Outline

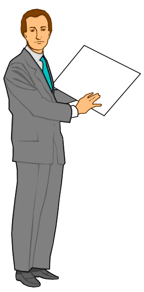

- Introduction
- Graph metrics
- Spectrum
- Network models
- Attacks & failures
- Framework for robustness

TU Delft

Birth of graph theory: the Königsberg bridge problem (Euler, 1736)

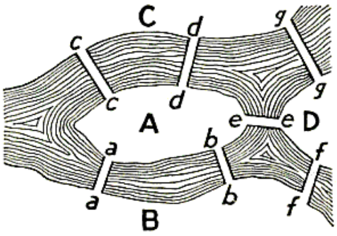

FIGURE 98. *Geographic Map: The Königsberg Bridges.*

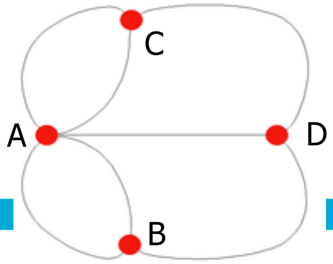

Can one walk across the seven bridges and never traverse the same bridge twice?

Leonhard Euler

Eulerian walk: zero or two nodes with odd degree

TU Delft

What is a network?

A graph G(N, L) specifies how items, called nodes, are interconnected or related to other nodes by links.

Trees

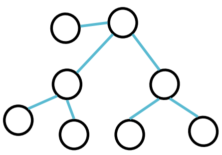

ring

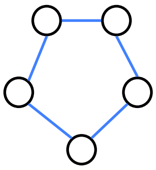

Complete graph

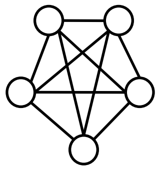

$L = N - 1$ $\quad$ $L = N$ $\quad$ $L = N(N-1)/2$

Fractal Nature of the Internet

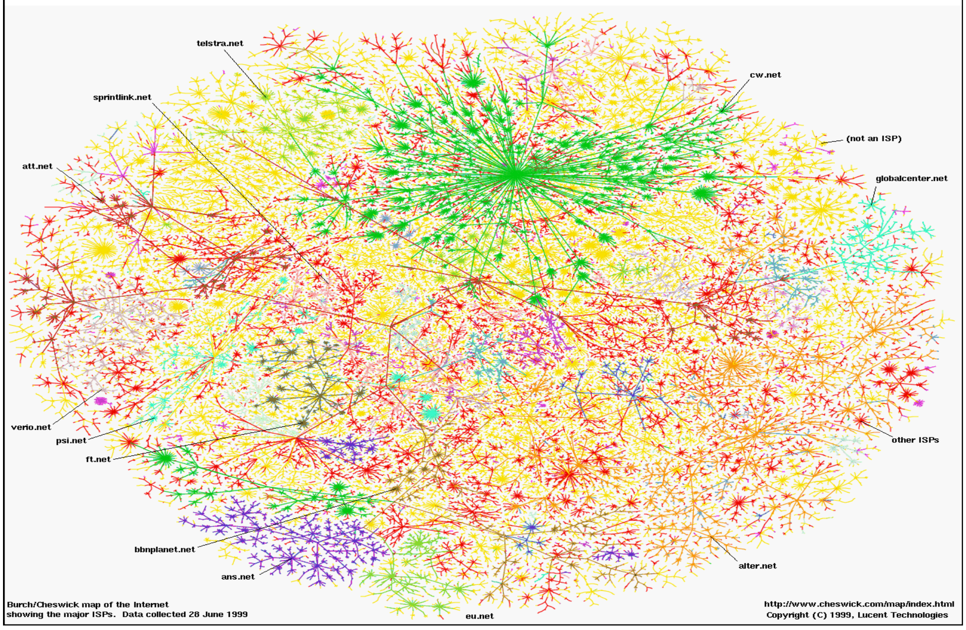

Our Brain network

The average human cerebral cortex contains

$N = 10^{11}$ neurons

$L = 10^{14}$ connections

500,000 km of wiring

moon-earth: 384 405 km

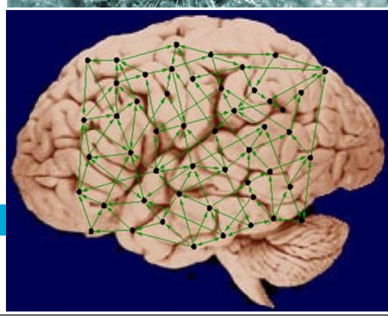

Network Science

- Are there properties common to all complex networks?
- if so, why?
- Can we formulate a general theory of the structure (topology), evolution and dynamics of complex networks?
- How do complex networks give rise to "adaptive", "living", "intelligent" behavior?
- How can we learn from nature to design robust, efficient, self-adaptive "man-made" networks ?

TU Delft

Outline

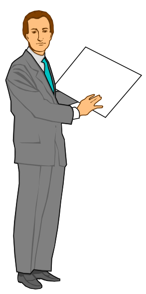

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

TU Delft

graph metric: degree

degree d_j of node j: number of neighbors of j

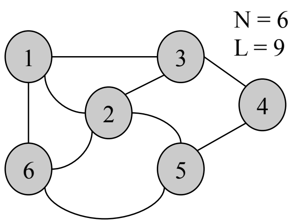

$d_1 = 3$
$d_2 = 4$
$d_3 = 3$
$d_4 = 2$
$d_5 = 3$
$d_6 = 3$

$d_1+d_2+d_3+d_4+d_5+d_6 = 18$

$$\sum_{j=1}^{N} d_j = 2L$$

average degree in G equals

$$E[D] = \frac{1}{N}\sum_{j=1}^{N} d_j = \frac{2L}{N}$$

bounds : $2 - \frac{2}{N} \leq E[D] \leq N - 1$

TUDelft

Adjacency matrix A

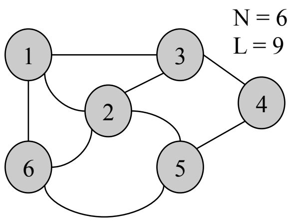

$$A_{N\times N} = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix}$$

For an undirected graph: $A = A^T$ is symmetric

Number of neighbors of node i is the degree: $d_i = \sum_{k=1}^{N} a_{ik}$

TUDelft

Ajacency matrix A

Degree vector d: $A\,u = d$ and $u = (1,1,...,1)$

Walk of length k from node i to j: succession of k links(arcs)

$$(n_0 \to n_1)(n_1 \to n_2)\ldots(n_{k-1} \to n_k) \text{ where } n_0 = i \text{ and } n_k = j$$

Path: a walk in which all nodes/vertices are different

Number of *k*-hop walks between node *i* and *j*: $(A^k)_{ij}$

***Open*: Number of *k*-hop paths** between node *i* and *j*
in terms of adjacency elements

TU Delft

graph metric: degree

Airline transportation network

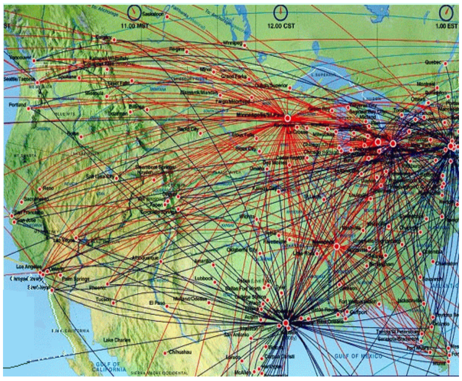

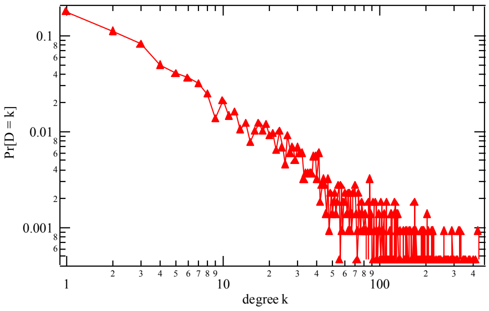

$$\Pr[D_{\text{Air}} = k] \sim k^{-1.21}$$

Internet: $\Pr[D_{\text{Internet}} = k] \sim k^{-\tau}, \quad \tau \in (2.2, 2.5)$

TU Delft

Graph Metric: Clustering coefficient

The clustering coefficient of node *v* is $c_G(v) = \frac{2y}{d_v(d_v - 1)}$

where y is the number of links between neighbors.
If $d_v = 1$, $c_G(v) = 0$.

$$c_G(A) = \frac{2}{4(4-1)} = \frac{1}{6}, \quad c_G(C) = 1$$

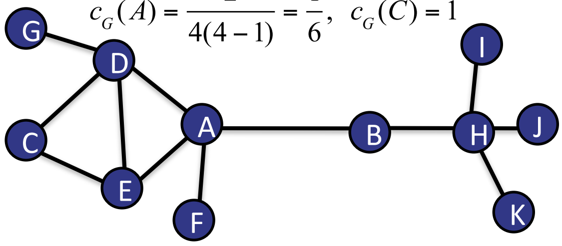

The clustering coefficient of a graph G: $c_G = \frac{1}{N}\sum_{v=1}^{N} c_G(v)$

TU Delft

Graph Metric: Hopcount

hopcount *H*: number of links in a shortest path in *G*

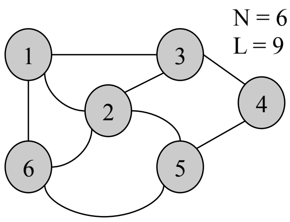

$H_{14} = 2$

diameter of G : hopcount of the longest shortest path in G

average hopcount E[H] reflects "efficiency" of transport in G

14

TU Delft

Graph Metric: Betweenness

The betweenness B_l of a link l is the number of shortest paths between all possible node pairs in G that traverse the link.

$$H_G = \sum_{i=1}^{N} \sum_{j=i+1}^{N} H_{i-j} = \sum_{l=1}^{L} B_l \quad \text{and} \quad E[B] = \frac{\binom{N}{2}}{L} E[H_N] \geq E[H_N]$$

H_{i-j} is the hopcount of the shortest path between i and j.

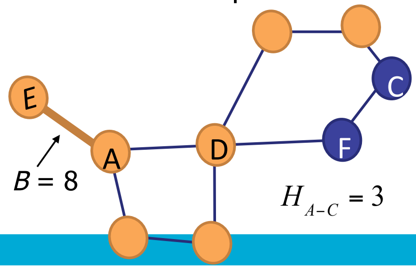

TU Delft

Assortativity

How are D_i and D_j (cor)related?

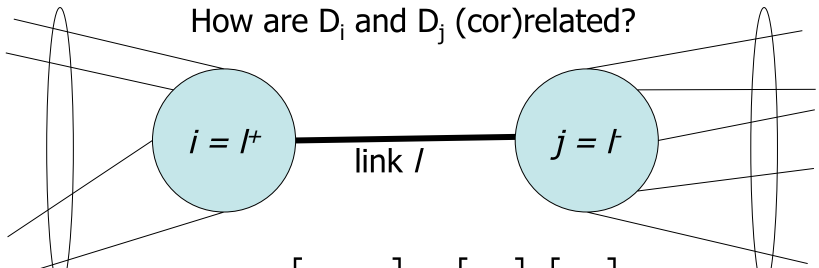

$$D_{l^+} - 1 \qquad \rho_D = \frac{E\left[D_{l^+} D_{l^-}\right] - E\left[D_{l^+}\right] E\left[D_{l^-}\right]}{\sqrt{Var\left[D_{l^+}\right] Var\left[D_{l^+}\right]}} \qquad D_{l^-} - 1$$

A network is (degree) *assortative* if $\rho_D > 0$

A network is (degree) *disassortative* if $\rho_D < 0$

16

TU Delft

(dis)assortativity

Reformulation of Newman's definition into algebraic graph theory

$$\rho_D = \frac{E\left[D_{l^+}D_{l^-}\right] - E\left[D_{l^+}\right]E\left[D_{l^-}\right]}{\sqrt{Var\left[D_{l^+}\right]Var\left[D_{l^+}\right]}} = \frac{N_1N_3 - N_2^2}{N_1\sum_{j=1}^{N} d_j^3 - N_2^2}$$

where $N_k = u^T A^k u$ is the total number of walks with k hops:

$$N_0 = \sum_{j=1}^{N} d_j^0 = N \qquad N_1 = \sum_{j=1}^{N} d_j^1 = 2L \qquad N_2 = \sum_{j=1}^{N} d_j^2 = d^T d$$

$$N_k \leq \sum_{j=1}^{N} d_j^k$$

17

P. Van Mieghem, H. Wang, X. Ge, S. Tang and F. A. Kuipers, 2010, "Influence of Assortativity and Degree-preserving Rewiring on the Spectra of Networks", The European Physical Journal B, vol. 76, No. 4, pp. 643-652

TU Delft

Degree-preserving rewiring

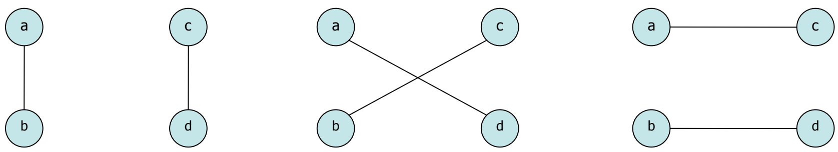

$$\rho_D = 1 - \frac{\sum_{i\sim j}\left(d_i - d_j\right)^2}{\sum_{j=1}^{N} d_j^3 - \frac{1}{2L}\left(\sum_{j=1}^{N} d_j^2\right)^2}$$

→ only two terms change

↓

degree-preserving rewiring algorithm

18

How many graphs do there exist with given N and d?

Open question (see B. McKay)

TU Delft

Connectivity of a Graph

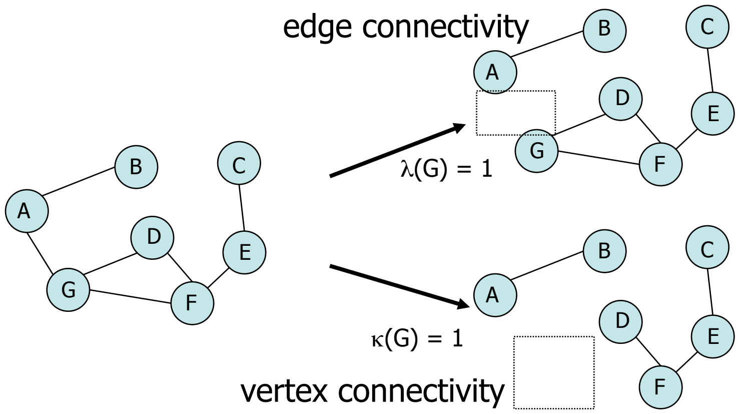

$\lambda(G)$ (or $\kappa(G)$) : the minimum number of links (nodes) whose removal disconnects G

Menger's Theorem

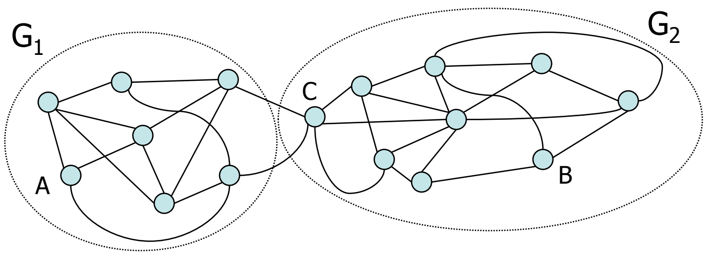

$d_a = 4$
$d_{min}(G) = 3$
$\lambda(G) = 2$
$\kappa(G) = 1$

Important inequality: $\kappa(G) \le \lambda(G) \le d_{\min}(G) \le \frac{2L}{N}$

***Menger's Theorem*:**

The maximum number of link(node)-disjoint paths between A and B is equal to the minimum number of links(nodes) seperating A and B.

There are at least $\lambda(G)$ link-disjoint and at least $\kappa(G)$ node-disjoint paths between any pair of nodes in G

List of topological metrics (undirected, unweighted graphs)

- hopcount
- closeness
- eccentricity
- diameter
- radius
- girth
- expansion
- distortion
- degree
- entropy
- joint degree
- assortativity
- modularity
- coreness
- clique number
- clustering coefficient
- rich club coefficient
- size giant component
- (node/link) connectivity
- coloring
- effective graph resistance
- *and many more*

21

TUDelft

Outline

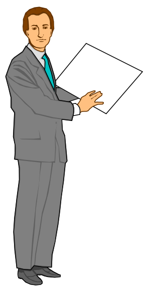

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

TUDelft

Eigenvalues and eigenvectors

$$Ax = \lambda x$$

$$A\begin{bmatrix} x_1 & x_2 & \cdots & x_N \end{bmatrix} = \begin{bmatrix} x_1 & x_2 & \cdots & x_N \end{bmatrix} \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_N \end{bmatrix}$$

$$AX = X\Lambda \quad \Rightarrow \quad A = X\Lambda X^{-1}$$

$$A = A^T = X\Lambda X^T = \sum_{k=1}^{N} \lambda_k x_k x_k^T$$

23

TU Delft

Basic theorem for symmetric matrices

Any **real symmetric matrix *S*** can be written as $S = X\Lambda X^T$, where X is the orthogonal matrix with real eigenvectors in the columns and $\Lambda = diag(\lambda_1,\ldots,\lambda_N)$, where λ_j is the j-th real eigenvalue.

The real eigenvalues can be ordered as

$$\lambda_N \le \lambda_{N-1} \le \cdots \le \lambda_2 \le \lambda_1$$

The eigenvalues are the zeros of the characteristic polynomial

$$\det(A - \lambda I) = 0$$

24

TU Delft

Algebraic graph theory

Any graph *G* with *N* nodes and *L* links can be represented by an adjacency matrix *A* and an incidence matrix *B*, and a Laplacian *Q*

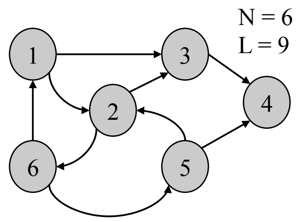

$$A_{N\times N} = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 \end{bmatrix} = A^T$$

$$B_{N\times L} = \begin{bmatrix} 1 & 1 & -1 & \dots & 0 \\ -1 & 0 & 0 & & 0 \\ 0 & -1 & 0 & & 0 \\ 0 & 0 & 0 & & 0 \\ 0 & 0 & 0 & & -1 \\ 0 & 0 & 1 & & 1 \end{bmatrix}$$

$$Q = BB^T = \Delta - A$$

$$\Delta = diag(d_1 \quad d_2 \quad \dots \quad d_N)$$

TU Delft

Basic properties of graph spectra

Spectrum of *A*:1) all eigenvalues lie in the interval $(-d_{max}, d_{max}]$

2) $\sum_{j=1}^{N} \lambda_j = 0 \qquad \sum_{j=1}^{N} \lambda_j^2 = 2L \qquad \sum_{j=1}^{N} \lambda_j^k = Trace(A^k) = \sum_{j=1}^{N} (A^k)_{jj}$

3) Perron-Frobenius Theorem: λ_1 non-negative and components eigenvector are non-negative. (irreducible = connected: positive)

Spectrum of *Q*: 1) any eigenvalue μ_k is non-negative and the smallest $\mu_N = 0$

2) complexity (number of spanning trees) is $\xi(G) = \frac{1}{N}\prod_{k=1}^{N-1} \mu_k$

3) the second smallest eigenvalue,

algebraic connectivity $a(G) = \mu_{N-1}$,

is related to how strongly a graph is connected

There exists a wealth of properties of graph spectra: see e.g.

P. Van Mieghem, *Graph Spectra of Complex Networks*, Cambridge University Press, 2011

TU Delft

Refreshing your knowledge (1/2)

- Adjacency matrix

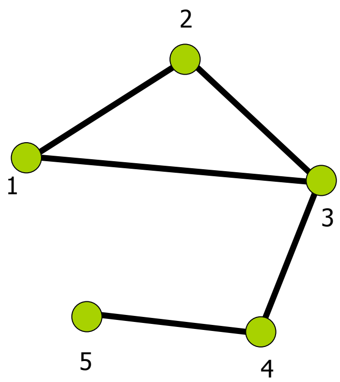

$$A = \begin{pmatrix} 0 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$

- λ_1 = spectral radius = largest eigenvalue of A

Refreshing your knowledge (2/2)

- Laplacian matrix

$$Q = \Delta - A$$

$$\Delta = diag(\ d_1 \quad d_2 \quad \dots \quad d_N)$$

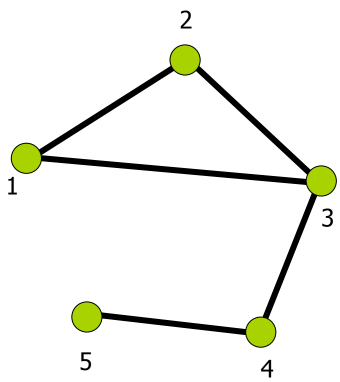

$$Q = \begin{pmatrix} 2 & -1 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 \\ -1 & -1 & 3 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & -1 & 1 \end{pmatrix}$$

- a(G) = *algebraic connectivity* = second smallest eigenvalue of Q

Connectivity of graphs

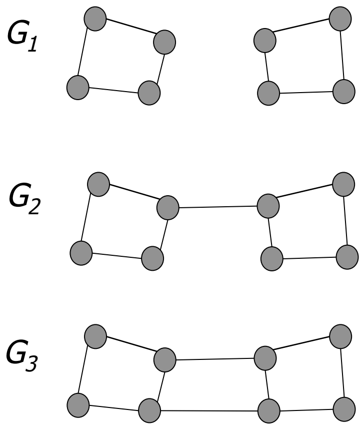

$a(G_1) = 0$

$a(G_2) = 0.29$

$a(G_3) = 0.59$

Difficulty to disconnect

TUDelft

Largest eigenvalue of a symmetric matrix

if $Ax = \lambda x$ then $A^k x = \lambda^k x$ for nonnegative integers k

Power method: $A^k w = \alpha_1 \lambda_1^k x_1 \left(1 + O\left(\left| \frac{\lambda_2}{\lambda_1} \right|^k \right) \right)$

Gerschgorin's theorem: $\lambda_1 \leq d_{\max}$

Rayleigh principle: $\lambda_1 \geq \frac{w^T A w}{w^T w}$ with equality only if $w = x_1$

There are many variations possible on the Rayleigh principle:

1) find suitable vector w
2) apply to powers of A recalling that $N_k = u^T A^k u$ is the total number of walks with k hops

P. Van Mieghem, Graph Spectra for Complex Networks, Cambridge University Press, 2011

TUDelft

Bounds largest eigenvalue adjacency matrix

Classical bounds: $d_{max} \geq \lambda_1(A) \geq \frac{2L}{N} = E[D]$

Walks based: $\lambda_1(A) \geq \left(\frac{N_{2k}}{N}\right)^{1/(2k)} \geq \left(\frac{N_k}{N}\right)^{1/k}$

$$k = 2 : \lambda_1(A) \geq \left(\frac{d^T d}{N}\right)^{1/2} = \frac{2L}{N}\sqrt{1 + \frac{Var[D]}{(E[D])^2}}$$

$$k = 3 : \lambda_1^3(A) \geq \frac{N_3}{N} = \frac{1}{N}\left(\rho_D\left(\sum_{j=1}^{N} d_j^3 - \frac{N_2^2}{N_1}\right) + \frac{N_2^2}{N_1}\right)$$

Optimized: $\lambda_1(A) \geq \frac{NN_3 - N_1N_2 + \sqrt{(NN_3)^2 - 6NN_1N_2N_3 + others}}{2(NN_2 - N_1^2)}$

31

P. Van Mieghem, Graph Spectra for Complex Networks, Cambridge University Press, 2011

TU Delft

Optimization

- Remove *m* nodes in *G* such that each removal decreases λ_1(A) maximally.
- Remove *l* links in *G* such that each link removal decreases λ_1(A) maximally.
  - What are the optimal strategies?
- Unfortunately, these problems are NP-complete...

The Interlacing Theorem

For a real symmetric n x n matrix A and any principal m x m submatrix B of A obtained by deleting n-m same rows and columns in A, the eigenvalues of B interlace with those of A as $\lambda_{n-m+i}(A) \leq \lambda_i(B) \leq \lambda_i(A)$ for any $1 \leq i \leq m$

32

P. Van Mieghem, D. Stevanovic, F. A. Kuipers, C. Li, R. van de Bovenkamp, D. Liu and H. Wang, 2011, "Decreasing the spectral radius of a graph by link removals", Physical Review E, Vol. 84, No. 1, July, p. 016101

TU Delft

USA air transportation network
N = 2179 and L = 31326

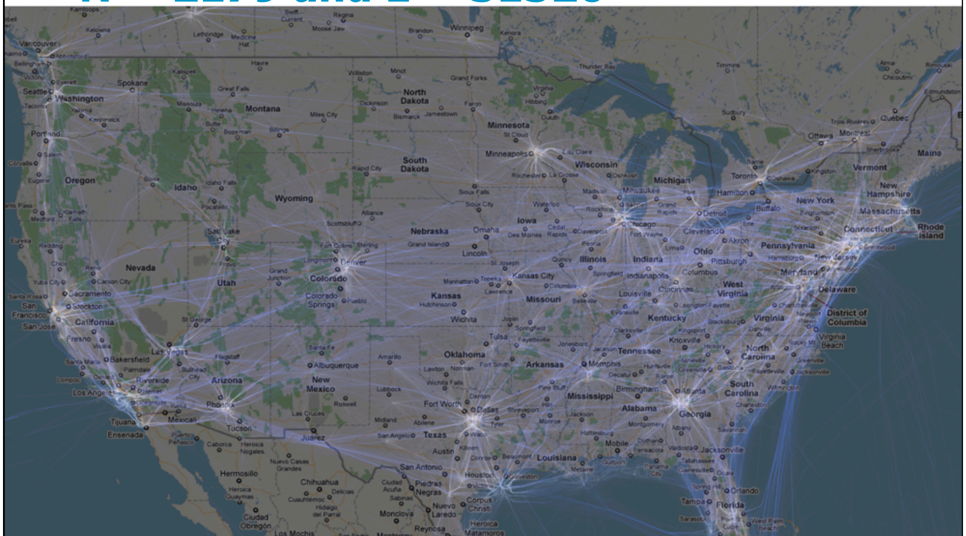

Van Mieghem, P., H. Wang, X. Ge, S. Tang and F. A. Kuipers, 2010, "Influence of Assortativity and Degree-preserving Rewiring on the Spectra of Networks", The European Physical Journal B, vol. 76, No. 4, pp. 643-652.

Degree-preserving rewiring USA air transport network: adjacency eig.

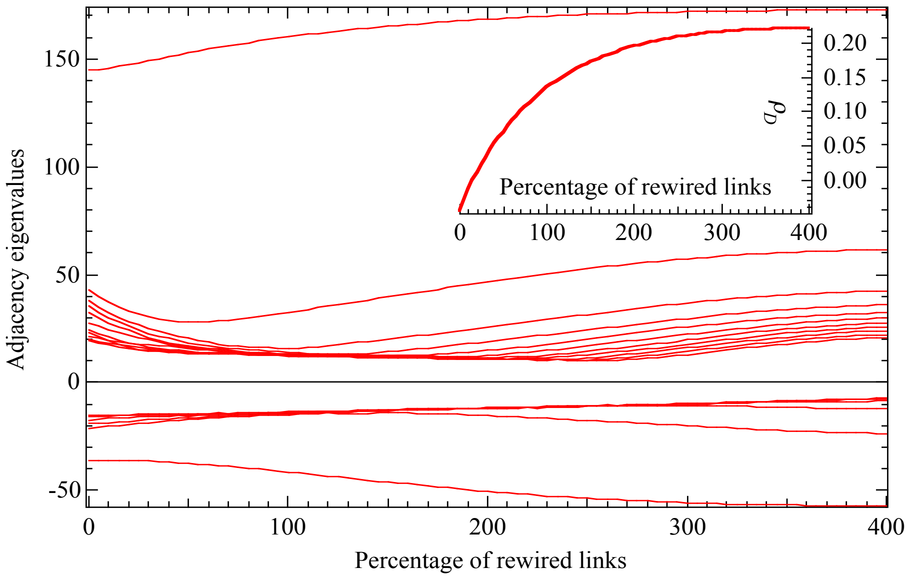

Degree-preserving rewiring USA air transport network: Laplacian eig.

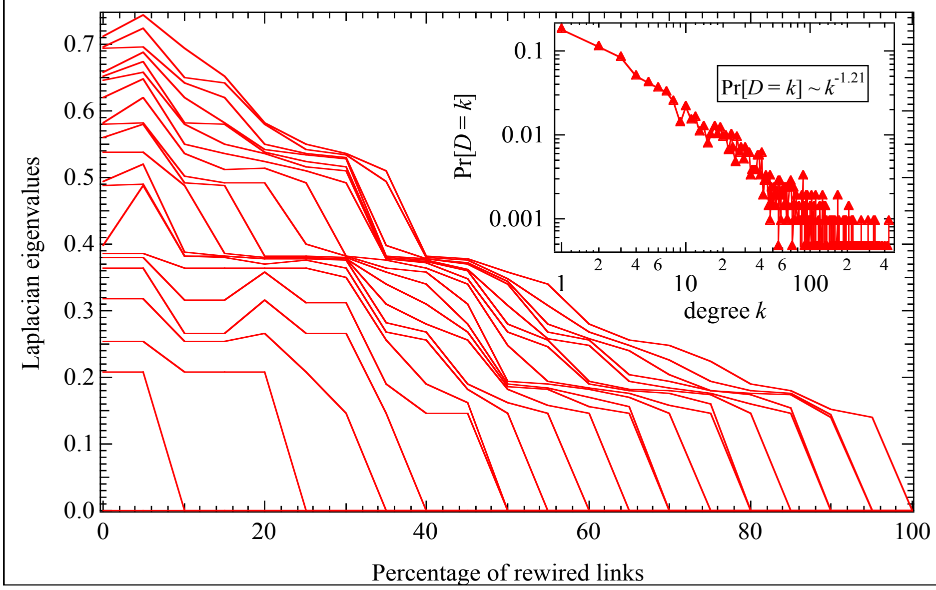

Outline

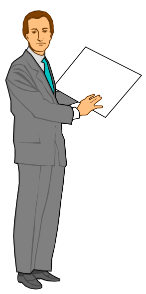

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

TU Delft

Complex network models

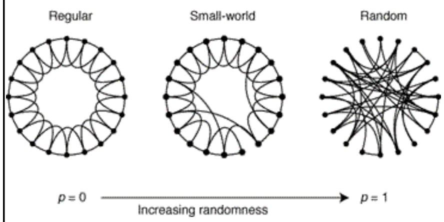

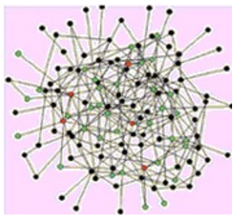

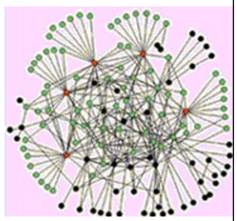

Small-World Graph of Watts-Strogatz

Erdös-Rényi Random Graph

Scale-Free Graph of Barabási-Albert

TU Delft

Network Model: Erdös-Rényi random graph

It is a class of graphs with *N* nodes and each node pair is connected independently with probability *p*.

$$E[L] = \frac{N(N-1)}{2} p$$

The average clustering coefficient follows

$$E[c_{G_p(N)}] = p$$

Network Model: Erdös-Rényi random graph

Degree distribution: Binomial distribution

$$\Pr[D_{rg} = k] = \binom{N-1}{k} p^k (1-p)^{N-1-k}$$

For large N and p = λ/N approaches a Poisson distribution

$$\Pr[D_{rg} = k] \simeq \frac{(pN)^k}{k!} e^{-Np}$$

Simplest proof via pgf:

$$E\left[z^{D_{rg}}\right] = (1-p+pz)^{N-1} = \left(1 + \frac{\lambda(z-1)}{N}\right)^{N-1} \to e^{\lambda(z-1)}$$

P. Van Mieghem, Performance Analysis of Communications Networks and Systems, Cambridge University Press, 2011

TUDelft

Network Model: Erdös-Rényi random graph

The critical link density p_c

$\Pr[G_p(N) \text{ is connected}] = \Pr[D_{min} \geq 1]$

$$\Pr[G_p(N) \text{ is connected}] = \begin{cases} 0 & \text{if } p < \frac{\log N}{N} \\ 1 & \text{if } p > \frac{\log N}{N} \end{cases}$$

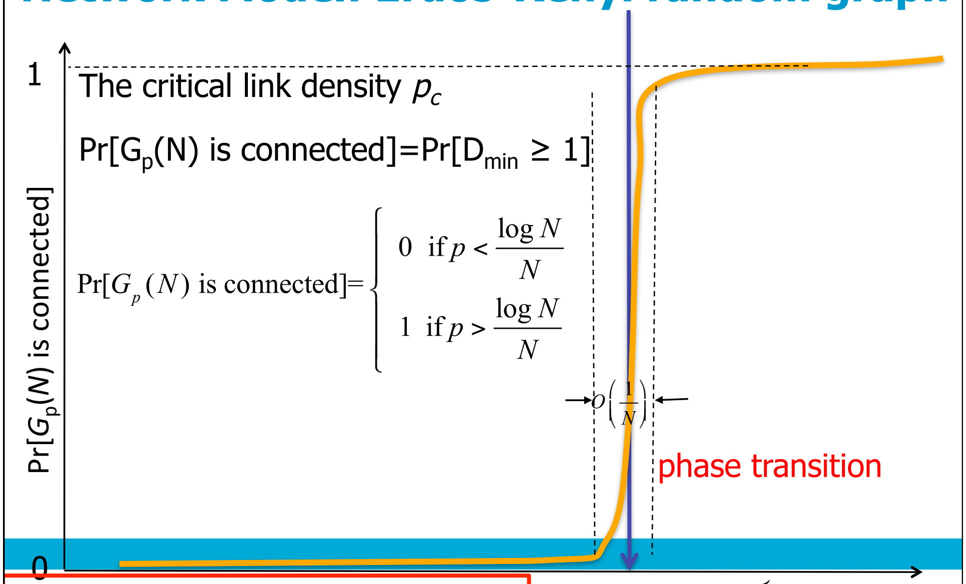

$$\Pr\left[G_p(N) \text{ is connected}\right] \approx \exp\left(-Ne^{-E[D]}\right)$$

$p_c \sim \log N/N$

TUDelft

Random Graph $G_{0.002}(300)$

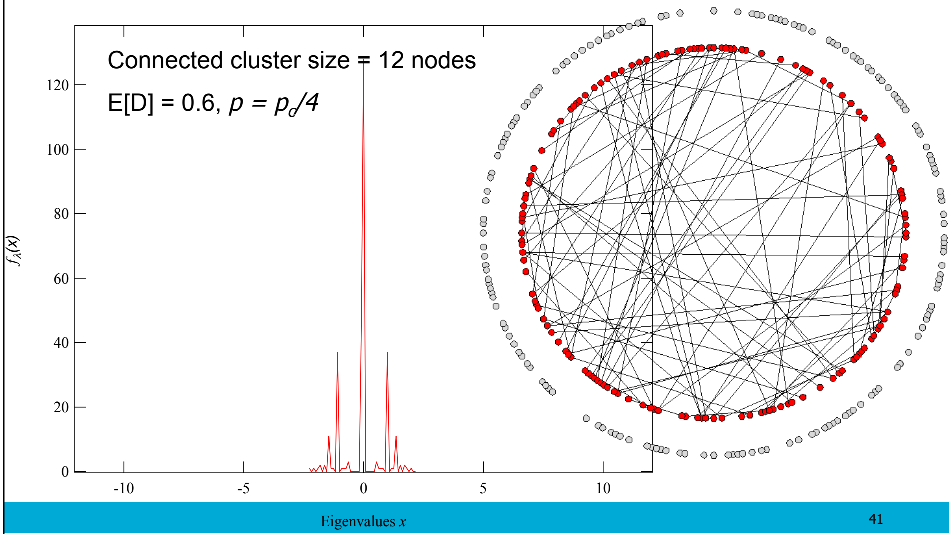

41

Isolated node

Node with degree ≤ 9

Node with degree > 9

from Javier Martin Hernandez

TUDelft

Random Graph $G_{0.004}(300)$

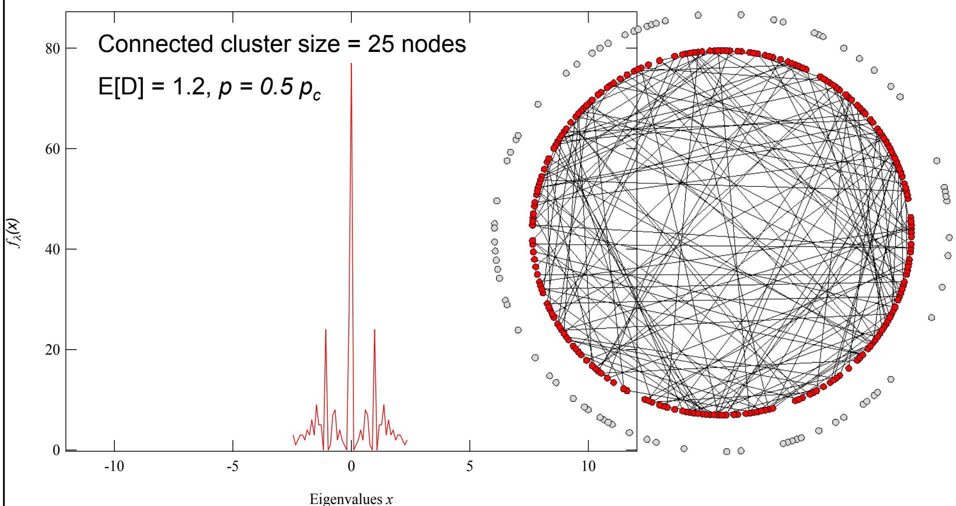

42

Isolated node

Node with degree ≤ 9

Node with degree > 9

TUDelft

Random Graph $G_{0.008}(300)$

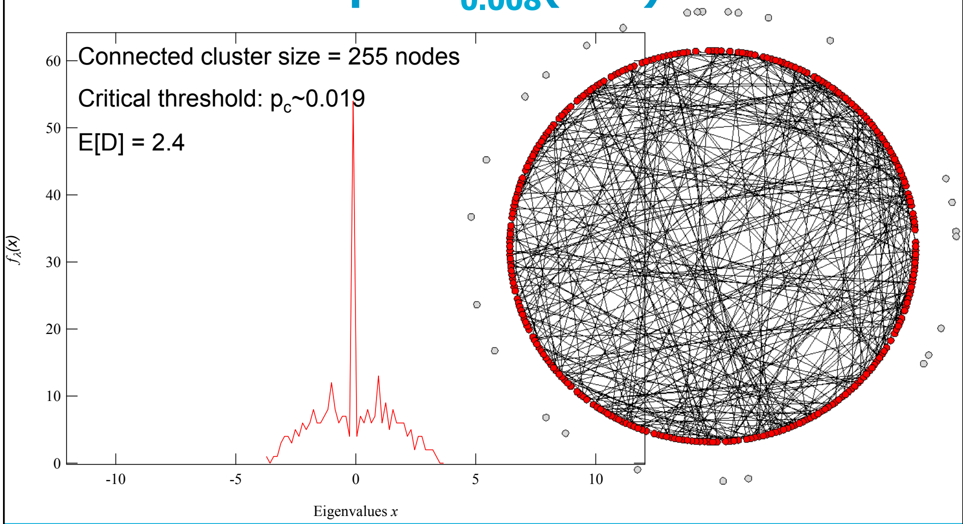

43

Isolated node

Node with degree ≤ 9

Node with degree > 9

TU Delft

Random Graph $G_{0.016}(300)$

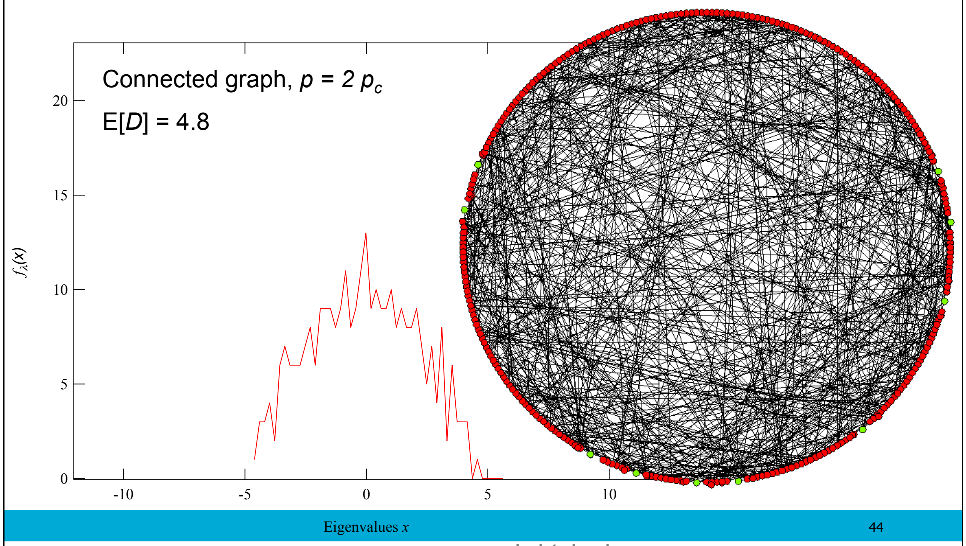

44

Isolated node

Node with degree ≤ 9

Node with degree > 9

TU Delft

Random Graph $G_{0.032}(300)$

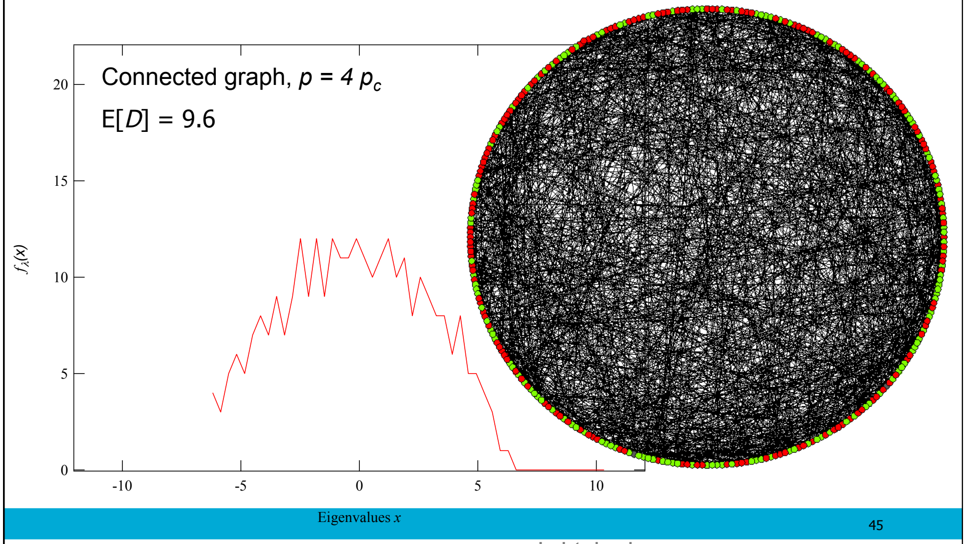

TU Delft

Wigner's semicircle law for $G_p(N)$

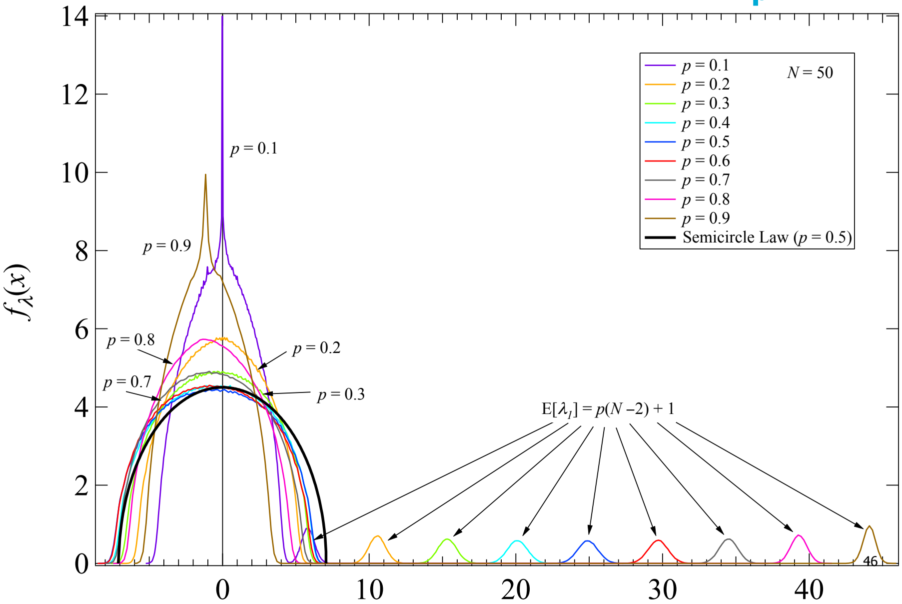

P. Van Mieghem, *Graph Spectra of Complex Networks*, Cambridge University Press, 2011

Network Model: Small-world graph

Collective dynamics of 'small-world' networks

letters to nature

Duncan J. Watts* & Steven H. Strogatz

Department of Theoretical and Applied Mechanics, Kimball Hall, Cornell University, Ithaca, New York 14853, USA

Table 1 Empirical examples of small-world networks

	L_{actual}	L_{random}	C_{actual}	C_{random}
Film actors	3.65	2.99	0.79	0.00027
Power grid	18.7	12.4	0.080	0.005
C. elegans	2.65	2.25	0.28	0.05

Network Model: Small-world graph

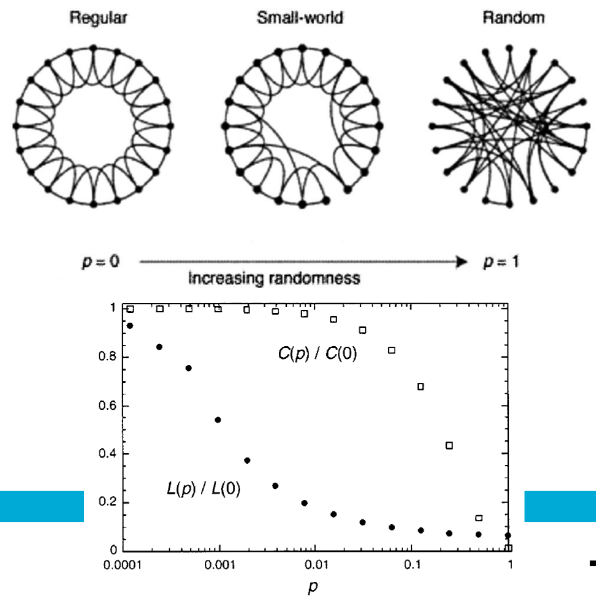

Spectrum small world graph

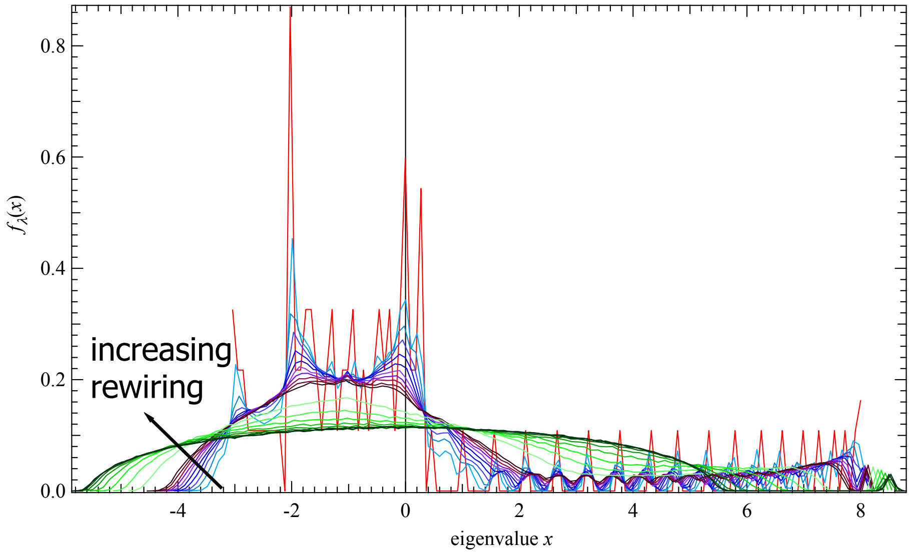

P. Van Mieghem, *Graph Spectra of Complex Networks*, Cambridge University Press, 2011

TU Delft

Network Model: BA power law graph

Metabolic network: nodes(chemicals) and links(bio-chemical reactions)

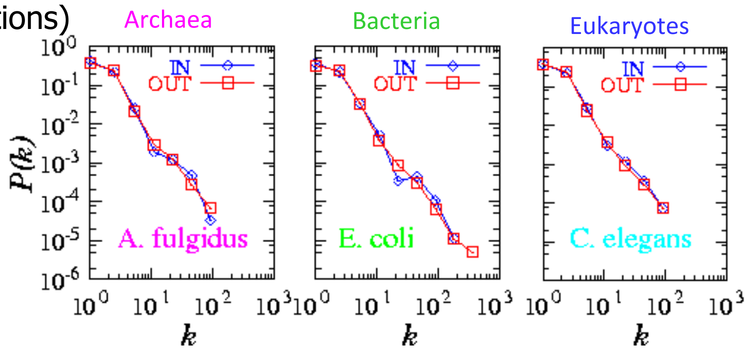

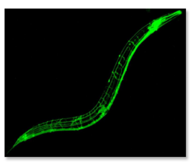

H. Jeong, B. Tombor, R. Albert, Z.N. Oltvai, and A.L. Barabasi, *Nature*, 407 651 (2000)

TU Delft

“Power laws” in complex networks

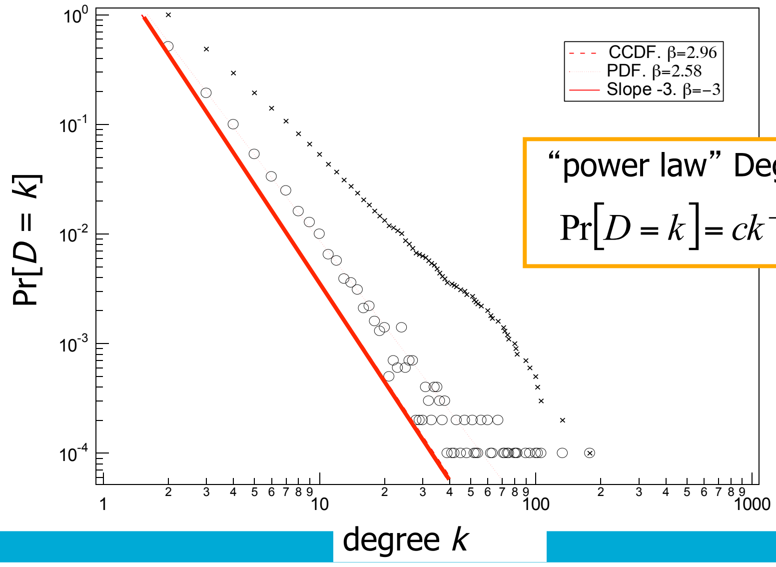

“power law” Degree:

$$\Pr[D=k] = ck^{-\beta}$$

TU Delft

Power law graphs

Measurements of the degree of nodes in (subgraphs of) the Internet topology indicate that

$$\Pr[D=k] \approx ck^{-\tau}$$

A power law degree distribution is also called scale-free:

$$\Pr[D=ak] = ca^{-\tau}k^{-\tau} = a^{-\tau}\Pr[D=k]$$

Any number a just multiplies the probability density; there is no characteristic length

Moreover,

$$E[D] = \frac{\zeta(\tau-1)}{\zeta(\tau)} \text{ provided } \tau > 2, \text{ where } \zeta(\tau) = \frac{1}{\tau-1} + \gamma + O(\tau-1)$$

$$E[D^k] = \frac{\zeta(\tau-k)}{\zeta(\tau)} \text{ provided } \tau > k+1$$

The simplest family of “power law” graphs have been proposed by Barbasi-Albert:

1) start with n nodes
2) attach a new node with m links to a node proportionally to its degree
3) repeat 2) until size N is reached

This construction of “preferential attachment”, “rich get richer”, is observed in many large complex networks (webgraph, proteins, social relations, etc...)

TU Delft

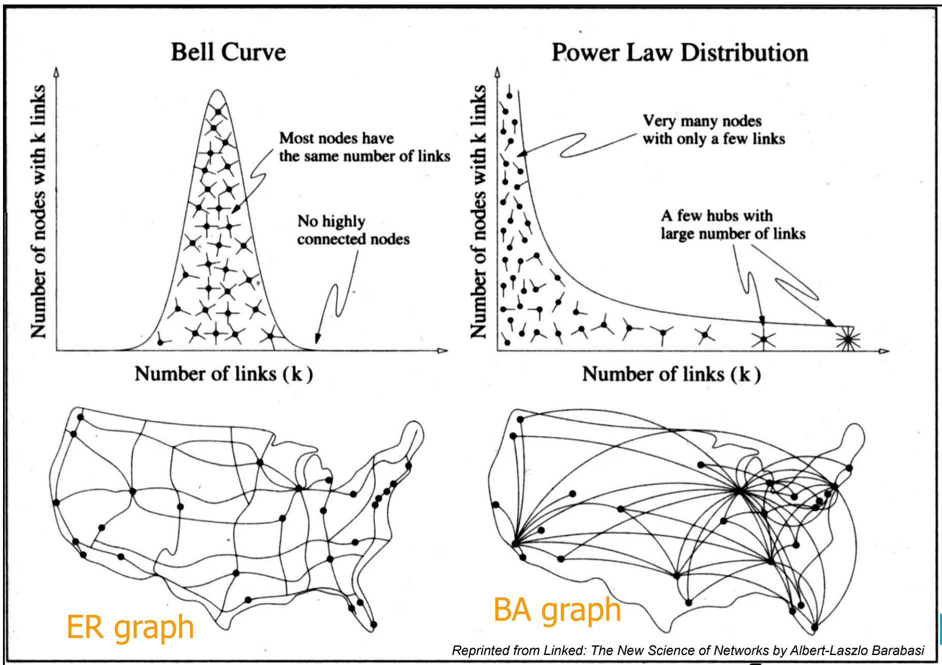

Mystery of "power laws"

Power law of a "property" appear if the "system" grows **exponentially**:

- if *X* grows exponentially with *Y* and *Y* has an exponential distribution, then *X* will have a power-law distribution *(Proof PA, p. 324)*

The exponential function *f(t)* has a linear differential equation

$$\frac{df(t)}{dt} = af(t)$$

which essentially means "growing proportional to its size"

At phase transitions, quantities of interest also change in a "power law" fashion

54

TU Delft

Observed common properties

- *small-world property*
  - average length of a path is short compared to the size N of the network ($E[H] = O(\log N)$)
- *scale-free degree distribution*
  - heavy tails (non-Gaussian behavior)
- *clustering and community structure*
  - network of networks
- *robustness to random node failure*
- *vulnerability to targeted hub attacks and cascading failures*

Outline

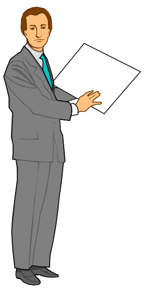

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

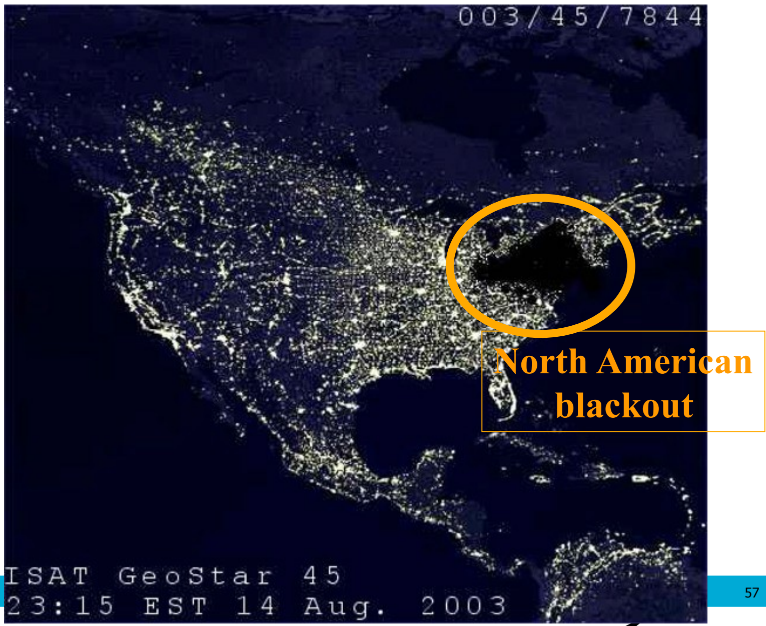

57

TU Delft

Cause: somewhere in Ohio

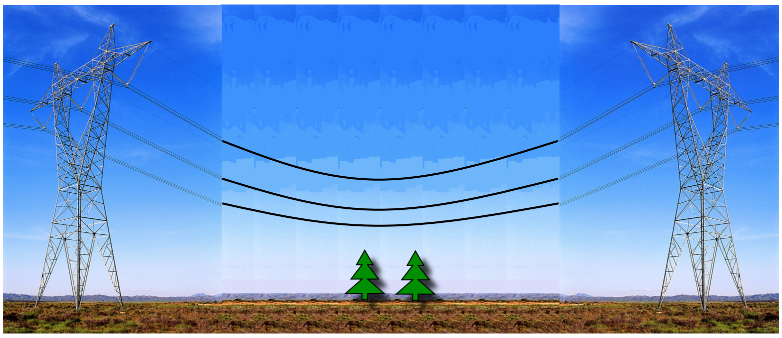

58

TU Delft

Cause: somewhere in Ohio

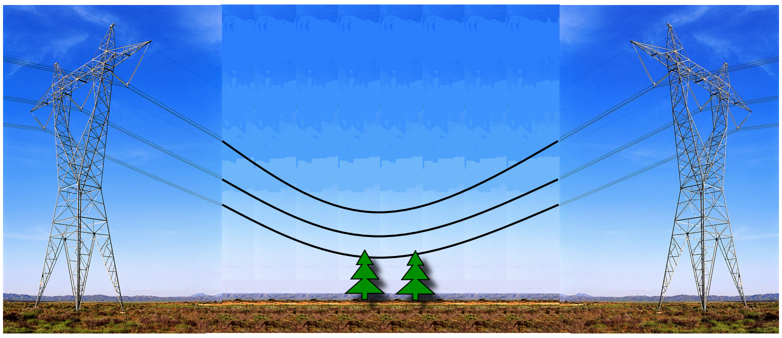

TU Delft

Cause: somewhere in Ohio

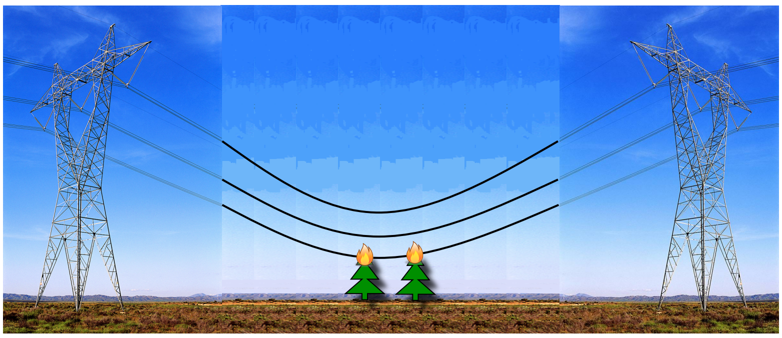

TU Delft

A few days earlier…

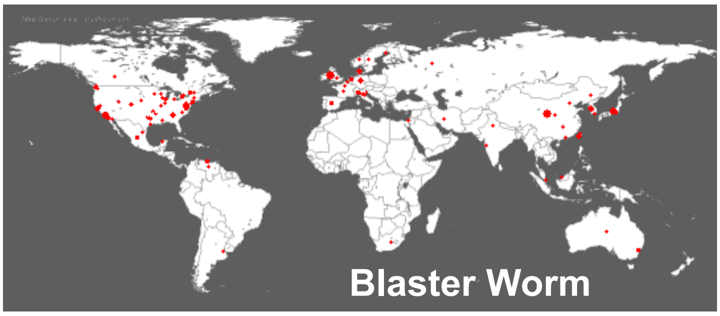

- “Alarm systems failed due to infection with Blaster Worm”

Introduction (1/2)

- Society is critically depending on complex networks
  - Internet
  - Transportation networks
  - Energy networks
  - Communication networks
- Severe consequences if networks are disrupted
- Robustness is defined as the extent to which the complex network is able to cope with perturbations imposed on it

Introduction (2/2)

- Examples of perturbations
  - Failures
    - Broken fibre cables
    - Malfunctioning switches
    - …
  - Attacks
    - Denial-of-Service
    - Physical attack on Internet Exchange
    - …
- Aim of this lecture: discuss examples of how graph-theoretic metrics can be used to quanity robustness

Motivation for virus spread in networks

- Computer viruses
  - security threat to Internet
  - annoyance
  - very costly
    - Code Red worm: several billion $$ in damage
- Why do we care?
  - Understanding the spread of a virus is the first step in preventing it
  - How fast do we need to disinfect nodes so that the virus dies quickly? Which nodes?

Applications of virus spread models

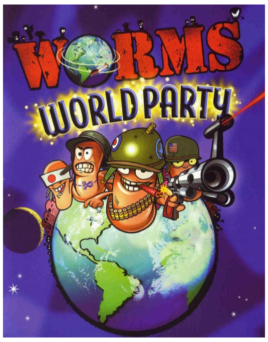

- Computer virus and worms modelling
- Epidemic algorithms
- Error propagation in networks
- Any self-replicating object on a dynamic network
- Emotions as infectious diseases in social networks

Simple SIS model (1)

- Homogeneous birth (infection) rate β on all edges between infected and susceptible nodes
- Homogeneous death (curing) rate δ for infected nodes

$\tau = \beta / \delta$: effective spreading rate

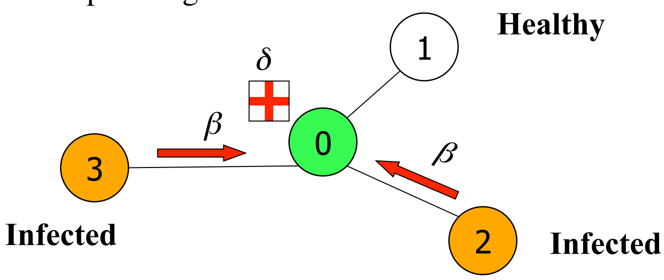

Simple SIS model (2)

- Each node *j* can be in either of the two states:
  - "0": healthy
  - "1": infected
- Markov continuous time:
  - infection rate β
  - curing rate δ
- Mathematically:
  - X_j is the state of node j
  - infinitesimal generator $Q_j(t) = \begin{bmatrix} -q_{1j} & q_{1j} \\ q_{2j} & -q_{2j} \end{bmatrix} = \begin{bmatrix} -q_{1j} & q_{1j} \\ \delta & -\delta \end{bmatrix}$

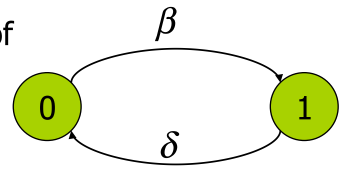

Simple SIS model (3)

- Nodes are interconnected in graph:

$$Q_j(t) = \begin{bmatrix} -q_{1j} & q_{1j} \\ \delta & -\delta \end{bmatrix}$$

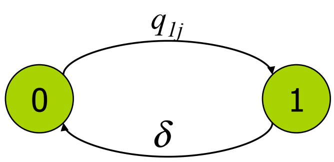

where the infection rate is due all infected neighbors of node *j:*

$$q_{1j}(t) = \beta \sum_{k=1}^{N} a_{jk} 1_{\{X_k(t)=1\}}$$

and where the adjacency matrix of the graph is

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1N} \\ a_{21} & a_{22} & \cdots & a_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ a_{N1} & a_{N2} & \cdots & a_{NN} \end{bmatrix}$$

Simple SIS model (4)

- Markov theory requires that the infinitesimal generator is a matrix whose elements are NOT random variables
- However, this is not the case in our simple model:

$$q_{1j}(t) = \beta \sum_{k=1}^{N} a_{jk} 1_{\{X_k(t)=1\}}$$

- By conditioning to each possible combination of infected states, we finally arrive to the exact Markov continuous SIS model
- *Drawback*: this exact model has 2^N states, where N is the number of nodes in the network.

69

TU Delft

Exact SIS model
N = 4 nodes

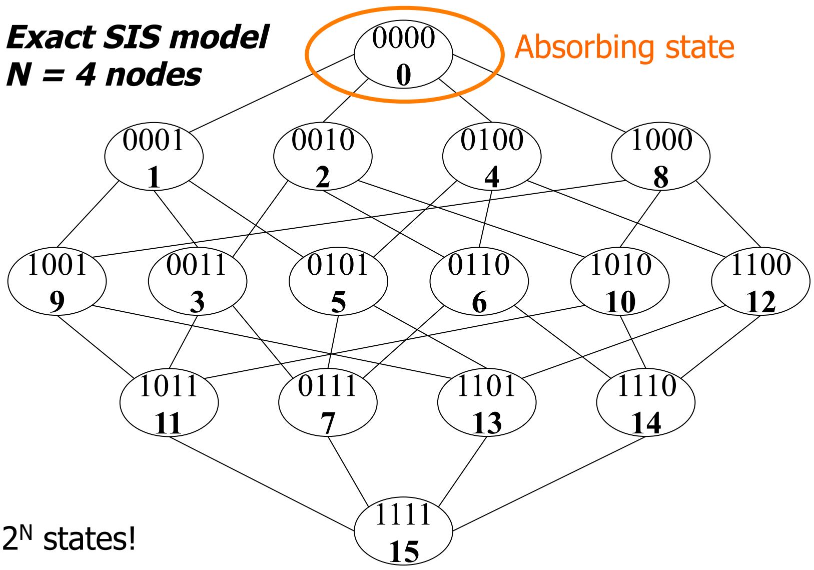

2^N states!

TU Delft

Simple SIS model (5): mean field

- The infinitesimal generator

$$Q_j(t) = \begin{bmatrix} -q_{1j} & q_{1j} \\ \delta & -\delta \end{bmatrix}$$

$$q_{1j}(t) = \beta \sum_{k=1}^{N} a_{jk} 1_{\{X_k(t)=1\}}$$

is replaced by its mean (***the only approximation***!)

$$Q_j = \begin{bmatrix} -E[q_{1j}] & E[q_{1j}] \\ \delta & -\delta \end{bmatrix}$$

$$E[q_{1j}(t)] = \beta \sum_{k=1}^{N} a_{jk} \Pr[\{X_k(t)=1\}]$$

- Being able now to apply ordinary Markov theory, we arrive at our **_N_-intertwined model for virus spread**

$$\begin{cases} \dfrac{dv_1}{dt} = (1-v_1)\beta \sum_{k=1}^{N} a_{1k} v_k - \delta v_1 \\ \dfrac{dv_2}{dt} = (1-v_2)\beta \sum_{k=1}^{N} a_{2k} v_k - \delta v_2 \\ \vdots \\ \dfrac{dv_N}{dt} = (1-v_N)\beta \sum_{k=1}^{N} a_{Nk} v_k - \delta v_N \end{cases}$$

where $v_k(t) = \Pr[X_k(t)=1]$

N-intertwined virus spread model

- Non-linear matrix equation:

$$\frac{dV(t)}{dt} = \beta A.V(t) - diag(v_i(t))(\beta A.V(t) + \delta u)$$

where the vector u^T =[1 1 … 1] and V^T = $[v_1\ v_2 \ldots v_N]$

- *Results:*
  - Probability of infection v_k for each node *k* separately
  - Number of infected nodes in the steady state
  - Phase transition phenomena for any network (largest eigenvalue of the adjacency matrix *A*)
  - Analytic computations feasible:
    - expansions of v_k as a function of the effective infection rate around the epidemic threshold and around infinity

P. Van Mieghem, J. Omic, R. E. Kooij, "Virus Spread in Networks", IEEE/ACM Transaction on Networking, Vol. 17, No. 1, pp. 1-14, (2009).

Simulations

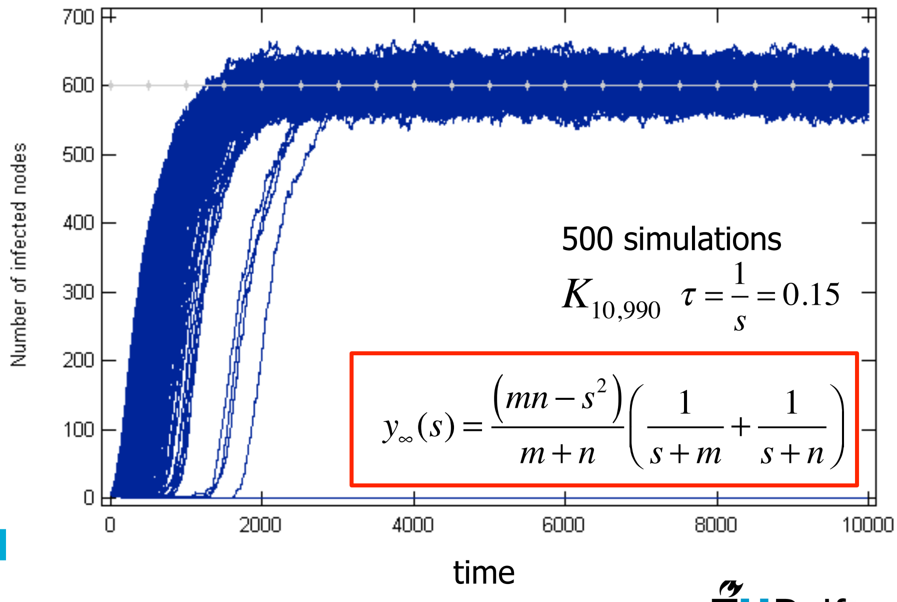

$K_{10,990} \quad \tau = \frac{1}{s} = 0.15$

$$y_\infty(s) = \frac{\left(mn - s^2\right)}{m+n}\left(\frac{1}{s+m} + \frac{1}{s+n}\right)$$

TU Delft

Kephart-White model

Assume perfect homogeneity & symmetry: a graph of degree *r*

$$\begin{cases} \frac{dv_1}{dt} = (1-v_1)\beta\sum_{k=1}^{N} a_{1k}v_k - \delta v_1 \\ \frac{dv_2}{dt} = (1-v_2)\beta\sum_{k=1}^{N} a_{2k}v_k - \delta v_2 \\ \vdots \\ \frac{dv_N}{dt} = (1-v_N)\beta\sum_{k=1}^{N} a_{Nk}v_k - \delta v_N \end{cases}$$

$$\frac{dv}{dt} = (1-v)\beta r v - \delta v$$

steady-state

$$v = 1 - \frac{1}{r\tau} \quad \text{where} \quad \tau = \frac{\beta}{\delta}$$

threshold

$$\tau \geq \tau_c = \frac{1}{r} \quad \text{then} \quad v \geq 0$$

74

J. O. Kephart and S. R. White, "Directed-graph epidemiological models of computer viruses," *Proc. IEEE Comput. Soc. Symp. Research in Security and Privacy*, May 1991, pp. 343–359.

TU Delft

What is so interesting about epidemics?

β : infection rate per link
δ : curing rate per node
τ= β/ δ : effective spreading rate

- **Final epidemic state**
- **Rate of propagation**
- **Epidemic threshold**

$$\tau_c = \frac{1}{\lambda_1(A)}$$

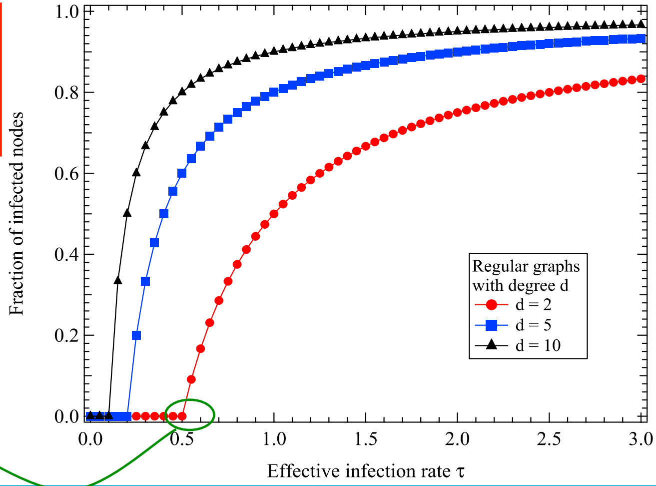

$$E[D] = \frac{2L}{N} \leq \lambda_1(A) \leq d_{\max}$$

TU Delft

Affecting the epidemic threshold

- Degree-preserving rewiring
  - Changing the assortativity of the graph
  - Van Mieghem, P., H. Wang, X. Ge, S. Tang and F. A. Kuipers, 2010, "Influence of Assortativity and Degree-preserving Rewiring on the Spectra of Networks", The European Physical Journal B, vol. 76, No. 4, pp. 643-652.
- Removing links/nodes (optimal way is NP-complete)
  - Van Mieghem, P., D. Stevanovic, F. A. Kuipers, C. Li, R. van de Bovenkamp, D. Liu and H. Wang, 2011, "Decreasing the spectral radius of a graph by link removals", Physical Review E, Vol. 84, No. 1, July, p. 016101.
- Quarantining: Removing inter-module links
  - Omic, J., J. Martin Hernandez and P. Van Mieghem, 2010, "Network protection against worms and cascading failures using modularity partitioning", 22nd International Teletraffic Congress (ITC 22), September 7-9, Amsterdam, Netherlands.

E.R. van Dam, R.E. Kooij, The minimal spectral radius of graphs with a given diameter, *Linear Algebra and its Applications,* 423, 2007, pp. 408-419.

TU Delft

Outline

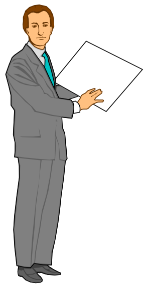

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

TUDelft

ResumeNet

ResumeNet

A framework for network robustness: the R-model

Define robustness Classes or “how good is the network?”

To which class does a given network belong?

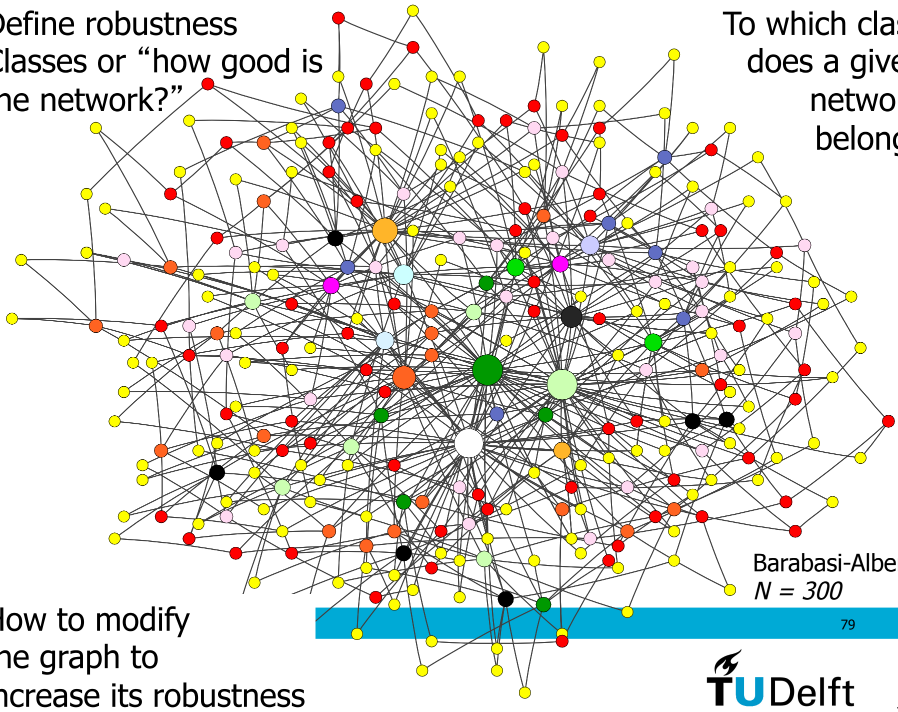

How to modify the graph to increase its robustness

79

TU Delft

Network: topology + service(s)

- Topology (or network infrastructure):
  - graph *G* with *N* nodes and *L* links
  - link weights
  - “hardware”
- Service:
  - more abstract and less clearly defined
  - uses the network infrastructure to transport items (e.g. email service, telephony, video, cars on roads, neurons in brain, etc.)
  - “software”
- Topology and service
  - own specifications
  - service is often designed independently of the topology
  - often more than 1 service on a same topology

80

TU Delft

High-Level Goal: Express Network Resilience in a Number *R*

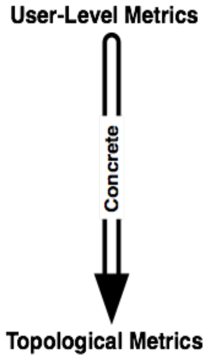

$R \in [0,1]$

R = 0: absence of resilience
R = 1: ideally resilient

TU Delft

Simple framework

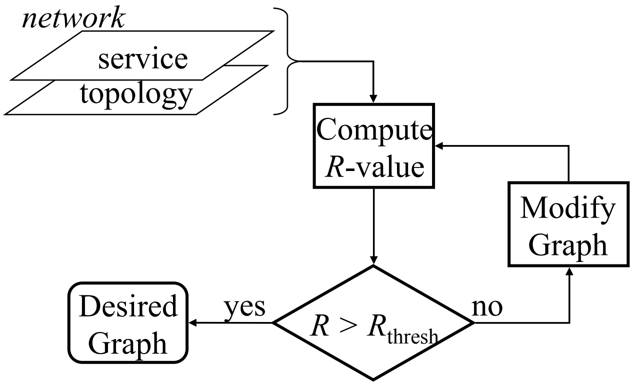

Goals:
1. define "*R*-value" that characterizes the level of robustness in *any* network
2. compute the *R*-value
3. robustness classes (understanding): which *R* is desirable and what is $R_{threshold}$

82

TU Delft

R-model

$$R = \sum_{k=1}^{m} s_k t_k = s.t \qquad (0 \le R \le 1)$$

s: the service vector with m components (interpreted as weights)
t: the topology vector where each component is a metric
(e.g. average degree, clustering coefficient, algebraic connectivity, minimum degree, diameter/hopcount, betweenness, etc...)

- Normalization: R = 0 (absense of network robustness)
  R = 1 (perfect robustness)
- **Linear**:
  - simplest m-dim expression & geometric interpretation (vector)
  - expectation E[R] easy
- no constraints on component values (else linear programming model)

83

P. Van Mieghem, C. Doerr, H. Wang, J. Martin Hernandez, D. Hutchison, M. Karaliopoulos and R. E. Kooij, 2010, "A Framework for Computing Topological Network Robustness", Delft University of Technology, report20101218.

TU Delft

Issues with R-model

- dimension m: trade-off between accuracy and computational complexity (*not problematic, consensus*)
- orthogonality of the metrics (*fundamental problem*)
  - each metric should ideally be a basis vector in m-dim space
  - almost all topology metrics are dependent
  - degree of dependence depends on the graph
  - *Solution*: no metrics, but matrices (adjacency A, incidence B, Laplacian Q) or spectra (graph theory)?
- Normalization of graphs: how to compare graphs with different number of nodes and links?
- unclear how to map a service onto a service vector (recall s_k is projection of s on k-th metric)

84

TU Delft

Which metrics to choose?

- Which metrics to choose is still an open question
  - Decomposition problem
  - Dependency problem
  - Normalization problem

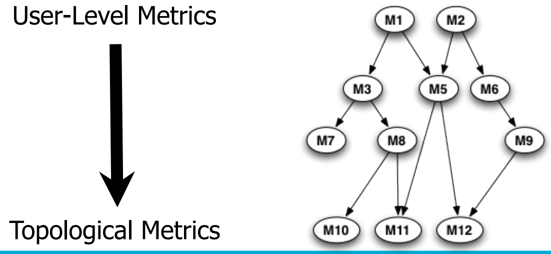

85

TU Delft

R as a function of "challenges"

- resilience is related to the network's capability to withstand perturbations from the outside during a given time interval

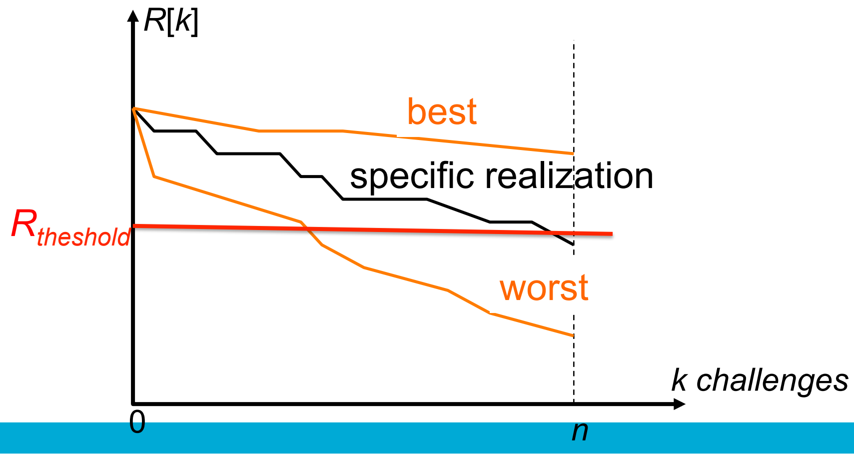

TU Delft

Understanding the Series of Events: Metric Envelopes

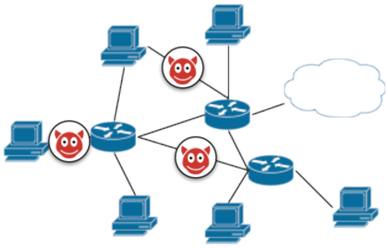

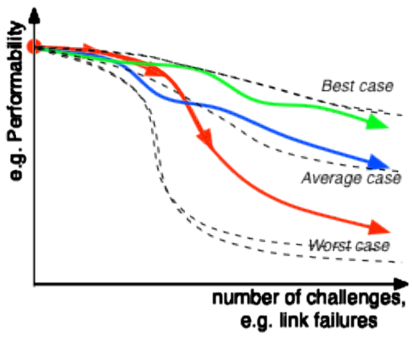

TU Delft

Understanding the Series of Events: Comparing Metric Envelopes

- Comparing resilience based on metric envelopes give a visual explanation of the network degradation process
- Depending on the application domain a more bounded envelope might be preferable
- The effect of various failure sources on the evaluated metric can be revealed

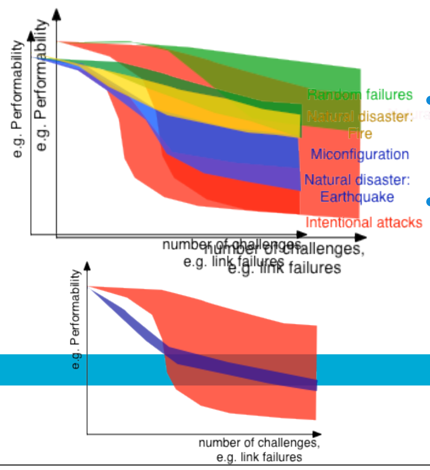

TU Delft

Computational Approach to a Measuring Resilience

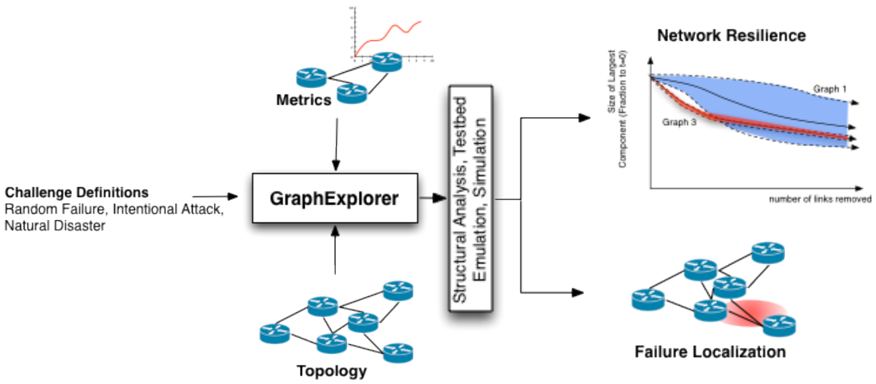

89

TUDelft

Case Study:
A Wireless GEANT2

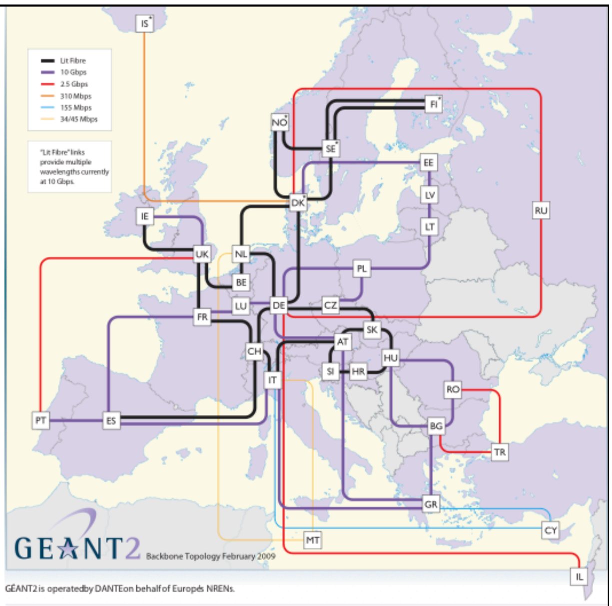

90

TUDelft

Show case for:
- Regional Challenges
- Fine-Grained, Intuitive Failures

W-GEANT2: Where are the weak points?

- Risk map indicates which areas are most vulnerable to challenges

Impact map visualizes the effect of a particular failure on the network as a whole

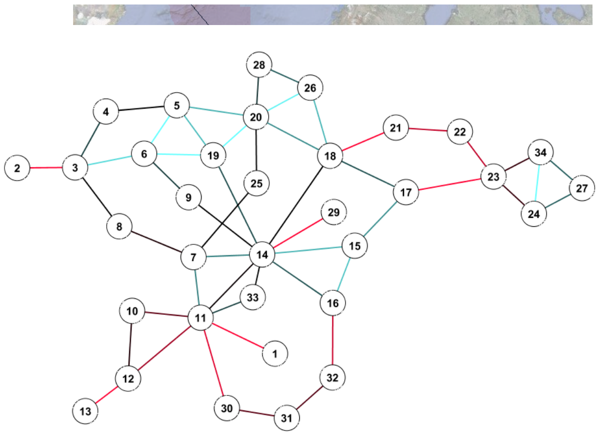

Let' s take a deeper look:
What concretely would happen?

TUDelft

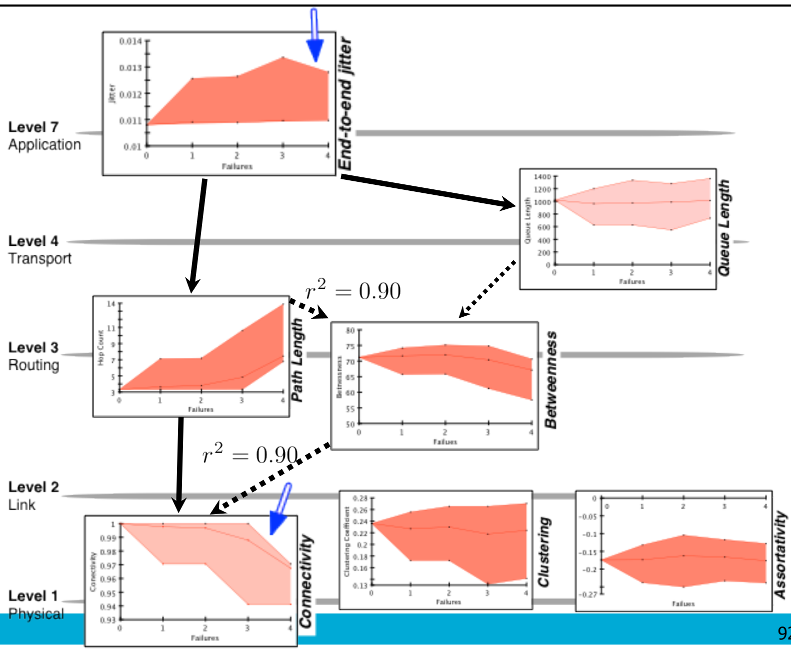

C. Doerr and J. Martin-Hernandez, "A Computational Approach to Multi-Level Analysis of Network Resilience," Proc. 3rd Int. Conf. Dependability, Venice, Italy, July 2010.

TUDelft

Outline

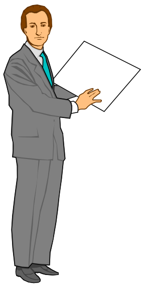

Introduction

Graph metrics

Spectrum

Network models

Attacks & failures

Framework for robustness

TUDelft

Books

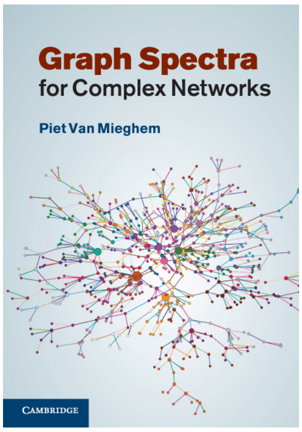

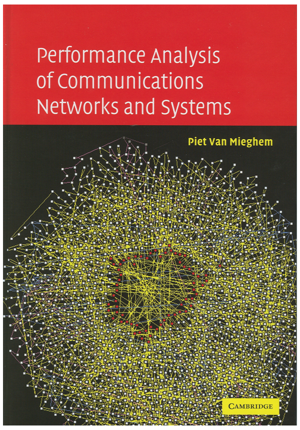

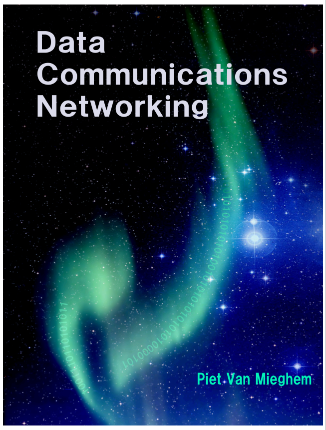

Articles: **http://www.nas.ewi.tudelft.nl**

94

TUDelft

IWSOS 2012 in Delft: March 15-16

supported by

euronf NETWORK OF EXCELLENCE

ifip

IWSOS 2012

Sixth International Workshop on Self-Organizing Systems
Delft, The Netherlands, March 15-16, 2012

TU Delft