Network Simulation and Testing

Polly Huang
Department of Electrical Engineering
National Taiwan University
http://cc.ee.ntu.edu.tw/~phuang
phuang@cc.ee.ntu.edu.tw

Topology Papers

- E. W. Zegura, K. Calvert and M. J. Donahoo. A Quantitative Comparison of Graph-based Models for Internet Topology. IEEE/ACM Transactions on Networking, December 1997.
- M. Faloutsos, P. Faloutsos and C. Faloutsos. On power-law relationships of the Internet topology. Proceedings of Sigcomm 1999.
- H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, W. Willinger. Network Topology Generators: Degree-Based vs. Structural. Proceedings of Sigcomm 2002.
- D. Vukadinovic, P. Huang, T. Erlebach. On the Spectrum and Structure of Internet Topology Graphs. In the proceedings of I2CS 2002.

Identifying Internet Topology

Random Graphs

Power law

Practical Model

The Problem

- What does the Internet look like?

- Routers as vertices
- Cables as edges
- Internet topologies as graphs
  - Which is this part of the Internet…

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

The Network Core

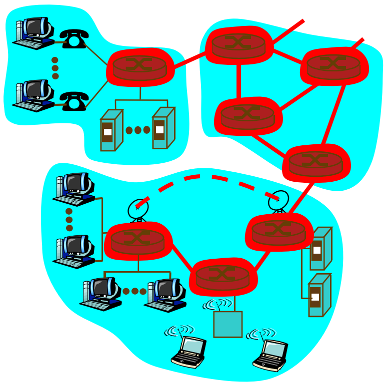

The Inter-connected Routers and Cables
(The Red Stuff)

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

For Example

The Internet, Circa 1969

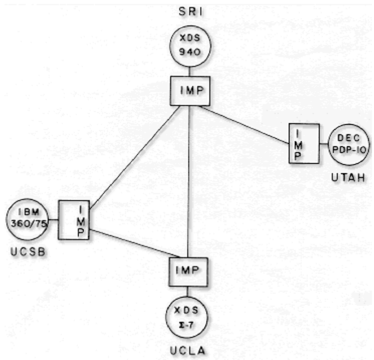

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

A 1999 Internet ISP Map

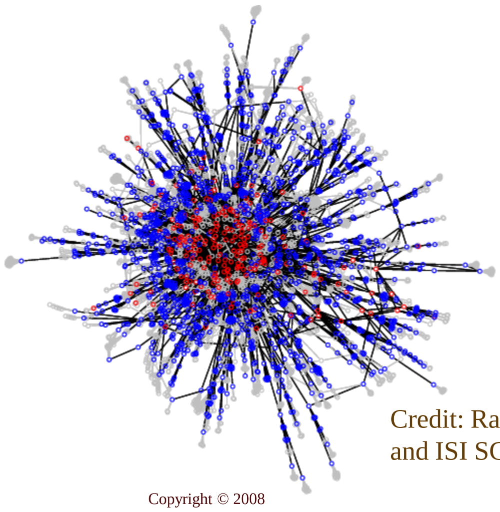

Credit: Ramesh Govindan
and ISI SCAN project

So?

- Tell me what this is
- Well.... Perhaps just give me a few of these so I can run my experiments

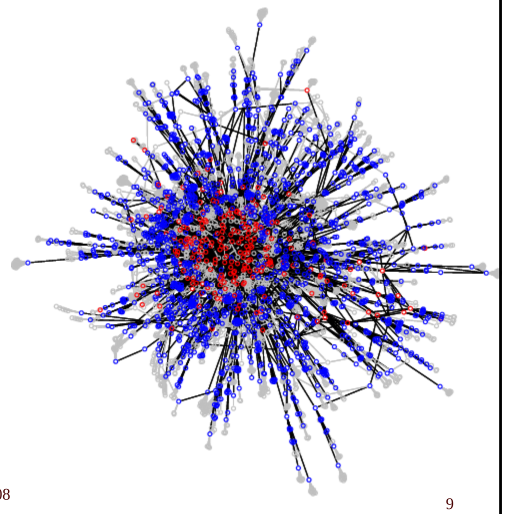

Back To The Problem

- What does the Internet look like?

- Equivalent of
  - Can we describe the graphs
    - String, mesh, tree?
    - Or something in the middle?
  - Can we generate similar graphs
    - To predict the future
    - To design for the future
- Not a new problem, but

Becoming Urgent

- Packet filter placement for DDoS
  - Equivalent of the vertex cover problem, NP complete
  - Exist a fast and optimal solution if the graphs are of certain type
  - How can the algorithm **be improved** with Internet-like topologies?
- VPN provisioning
  - Equivalent of the fluid allocation problem, NP complete
  - Exist heuristics and greedy algorithms performing differently depending on the graph types
  - How will the algorithm **perform** with Internet-like topologies?

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

More Specific

- Insights to **design**
  - What are the characteristics?
- Confidence in **evaluation**
  - Can we generate random topologies with the characteristics?
    - Why not use current Internet topologies?
    - Want the algorithm continue to work
    - Can't really predict the future
    - Thus, try with a few highly probably futures

In Another Sense

- Need to **analyze**
  - dig into the details of Internet topologies
  - hopefully to find invariants
- Need to **model**
  - formulate the understanding
  - hopefully in a compact way

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Background

- As said, the problem is not new!
- Three generations of network topology analysis and modeling already:
  - 80s - No clue, not Internet specific
  - 90s - Common sense
  - 00s - Some analysis on BGP Tables
- To describe: basic idea and example

Early Models

- A Quantitative Comparison of Graph-based Models for Internet Topology
- E. W. Zegura, K. Calvert and M. J. Donahoo..
- IEEE/ACM Transactions on Networking, December 1997

The No-clue Era

- Heuristic
- Waxman
  - Define a plane; e.g., [0,100] X [0,100]
  - Place points uniformly at random
  - Connect two points probabilistically
    - p(u, v) ~ 1 / e^{d} ; d: distance between u, v
- The **farther apart** the two nodes are, the **less likely** they will be connected

Waxman Example

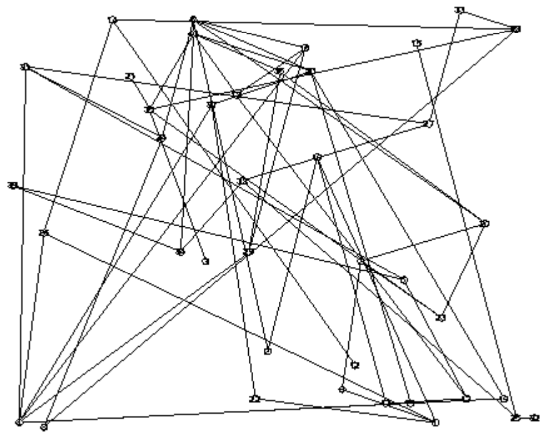

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

More Heuristics

- Pure Random
  - $p(u, v) \sim C$
- Exponential
  - $p(u, v) \sim 1 / e^{d/(L-d)}$
    - d: distance between u, v
    - L: ? → Restrict the max distance
- Locality
  - $p(u, v) \sim \begin{cases} \alpha & \text{if } d < r \\ \beta & \text{if } d \geq r \end{cases}$
    - D: distance between u, v
    - r: ? → Boundary of connectivity change

Copyright © 2008

These are also referred to as the:

Flat random graph models

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

The question is…

Is Internet flat?

Copyright © 2008

Remember This?

Inter-AS border (exterior gateway) routers

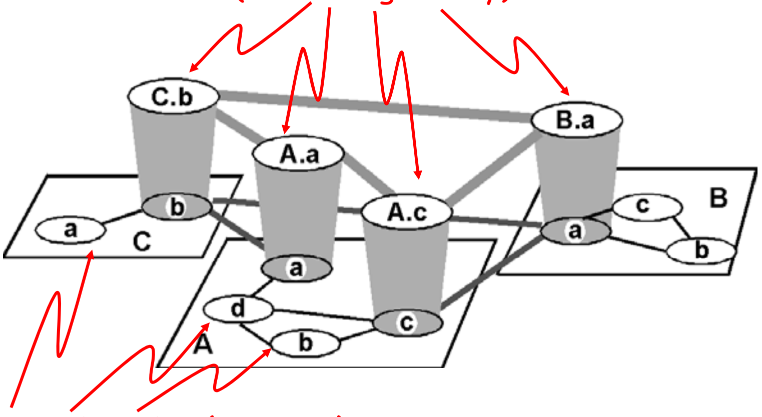

Intra-AS interior (gateway) routers

Internet: The Network

- The Global Internet consists of Autonomous Systems (AS) interconnected with each other:
  - **Stub AS**: small corporation: one connection to other AS's
  - **Multihomed AS**: large corporation (no transit): multiple connections to other AS's
  - **Transit AS**: provider, hooking many AS's together
- Two-level routing:
  - **Intra-AS:** administrator responsible for choice of routing algorithm within network
  - **Inter-AS:** unique standard for inter-AS routing: BGP

Therefore…

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

The Common-sense Era

- Hierarchy
- Tier
  - In a geographical sense
  - WAN, MAN, LAN
- GT-ITM
  - In a routing sense
  - Transit (inter-domain), stub (intra-domain)

Tier

- One big plane
  - Divide to random # of **WAN** partitions
  - Pick a random point in a partition
- One WAN
  - # of **MAN** partitions
  - point in a partition
- One MAN
  - # of **LAN** partitions
  - point in a partition

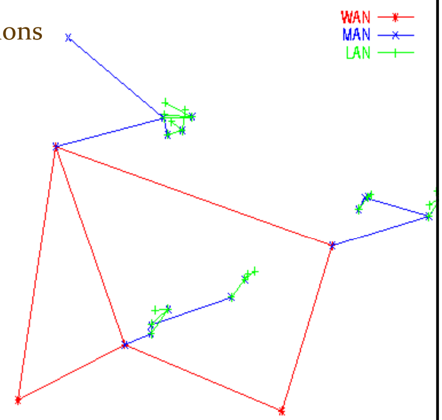

GT-ITM

- **Transit**
  - Number
  - Connectivity
- Stub
  – Number
  – Connectivity
- Transit-stub
  – Connectivity

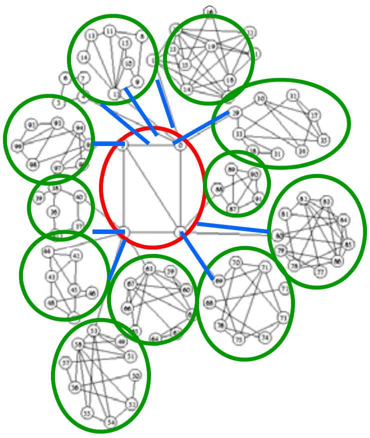

Now the question is…

Does it matter which model I use?

A Quantitative Comparison

- Compare these **models**
  - Flat: Waxman, pure, exponential, locality
  - Hierarchical: Tier (N-level), TS
- With these **metrics**
  - Number of links
  - Diameter
    - For all pairs of nodes, the longest distance of all shortest paths
  - Number of biconnected components
    - Biconnected component: max set of a sub-graph that any 2 links are on the same cycle

Methodology

- Fixed the number of nodes and links
- Find the parameters for each model
  - that will in result generate the number of nodes and links
- Reverse engineering
  - Some with only 1 combination
  - Some with multiple combinations
    - TS usually

Comprehensible Results

- Amongst the flat random models
  - **Pure** random longer in **length diameter**
- Amongst the hierarchical random models
  - **TS** higher in **# of bicomponents**
- Between the flat and hierarchical models
  - **Flat** lower on # of **bicomponents**
  - **Flat** lower in **hop diameter**

Statistical Comparison

- KS test for hypothesis
  - For any pair of models
    - X; Y
  - Generate N number of graphs
    - $X_1,\ldots,X_N$; $Y_1,\ldots,Y_N$
  - Find the metric value M for each graphs
    - $M(X_1, X_2, \ldots X_N)$; $M(Y_1, Y_2, \ldots Y_N)$
  - Find if the 2 samples are from the same population
    - Confidence level 95%
  - Yes meant X and Y are 95% the same

Quantify the Similarity

- Home-bred test for degree of similarity
  - For any pair of models
    - X; Y
  - Generate N number of graphs
    - $X_1,\ldots,X_N$; $Y_1,\ldots,Y_N$
  - Find the metric value M for each graphs
    - $M(X_1, X_2, \ldots X_N)$; $M(Y_1, Y_2, \ldots Y_N)$
  - For i = 1,…N, compute the probability of
    - $M(X_i) < M(Y_i)$
  - 0.5 meant X and Y are similar relative to M
  - All black or all white → very different

Harder to Grasp Results

- Confirm the simple metric comparison results
- Results of different sizes and degrees being Consistent
- Length-based and hop-based results are quite different
- Significant diff between N-level and TS

Making Another Statement

- The use of graph model is application dependent
- Show in multicast experiments
  - Delay and hop counts of the multicast trees
  - Different graph models give different results

Nice Story, But is This Real?

- What is TS
  - Composition of flat random graphs
  - Which random really?
- Measurement infrastructure is maturing
- Repository of real Internet graphs

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Identifying Internet Topology

Random Graphs

Power law

Practical Model

Break-through

- On power-law relationships of the Internet topology
- M. Faloutsos, P. Faloutsos and C. Faloutsos
- Proceedings of Sigcomm 1999.

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

A Study of BGP Data

- Analyze BGP routing tables
  - November 1997 to December 1998
  - Autonomous System level graphs (AS graphs)
- Find power-law properties in AS graphs
  - 3.5 of these power-law relationships

- Power-law by definition
  - Linear relationship in log-log plot

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

2 Important Power-laws

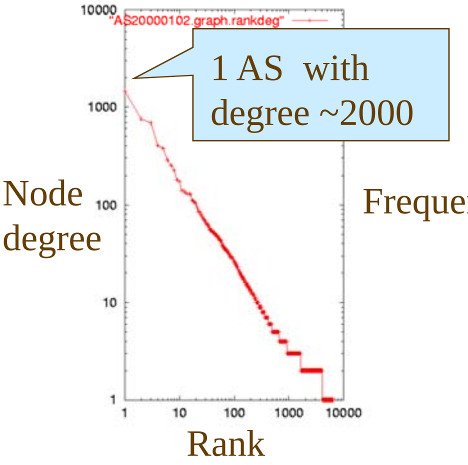

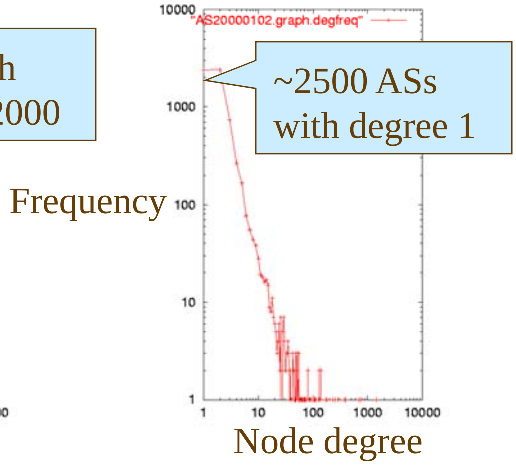

Polly @ NTU

39

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

1.5 More Power-Laws

- Number of h-hop away node pairs to h
  - Actually, this one, not quite…
- Eigenvalues λ_i to i
  - A graph is an adjacency matrix
  - λ_i, eigenvalues of that matrix

Polly @ NTU

40

The Power-law Era

- Models of the 80s and 90s
  - Fail to capture power-law properties
- BRITE
  - Barabasi's incremental model
- Inet
  - Fit the node degree power-laws specifically
- Won't show examples
  - Too big to make sense

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

BRITE

- Create a random core
- Incrementally add nodes and links
- Connect new link to existing nodes probabilistically
  - Waxman or preferential

- Node degrees of these graphs will magically have the power-law properties

Inet

- Generate node degrees with power-laws
- Connecting links preferentially to node degree at random

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Are They Better?

- Network Topology Generators: Degree-Based vs. Structural
- H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, W. Willinger..
- Proceedings of Sigcomm 2002

A Newer Comparison

Paper #1 vs. Paper #3

- Methodology the same
  - Given the random graph models
  - And a set of metrics
  - Find differences and similarities

Relevance Enriched

Paper #1 vs. Paper #3

- Up-to-date models
  - Adding the power-law specific models into the comparison
- Network-relevant metrics
  - Expansion, resilience, distortion, link value
- Concrete reference data
  - BGP table derived AS graphs
  - Can say more or less realistic

Structural vs. Degree-based

- Structural
  - Tier and TS
- Degree-based
  - Inet, BRITE, and etc.

Metrics for Local Property

- Expansion
  - Size of neighborhood per node
  - Control message overhead
- Resilience
  - Number of disjoint path per node pair
  - Probability of finding alternative routes
- Distortion
  - Min cost of spanning tree per graph
  - Cost of building multicast tree

Measure of Hierarchy

- Link Value
  - Home-bred
  - Degree of traversal per link
    - Each link maintains a counter initialized to 0
    - For all pair of nodes
    - Walk the shortest path
    - For each link walked, increment the link's counter
    - Looking at the distribution of the counter values
  - Location and degree of congestion

Polly @ NTU

49

Result in a Sentence

- Current degree-based generators **DO** work better than Tier and TS.
- This doesn't mean structure isn't important!

Polly @ NTU

50

There's yet another question…

BRITE or Inet?

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Which is better?

- Compare **AS**, **Inet**, and **BRITE** graphs
- Take the AS graph history
  - From NLANR
  - 1 AS graph per 3-month period
  - 1998, January - 2001, March

Methodology

- For each **AS** graph
  - Find number of nodes, average degree
  - Generate an **Inet** graph with the same number of nodes and average degree
  - Generate a BRITE graph with the same number of nodes and average degree
  - Compare with addition metrics
    - Number of links
    - Cardinality of matching

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Number of Links

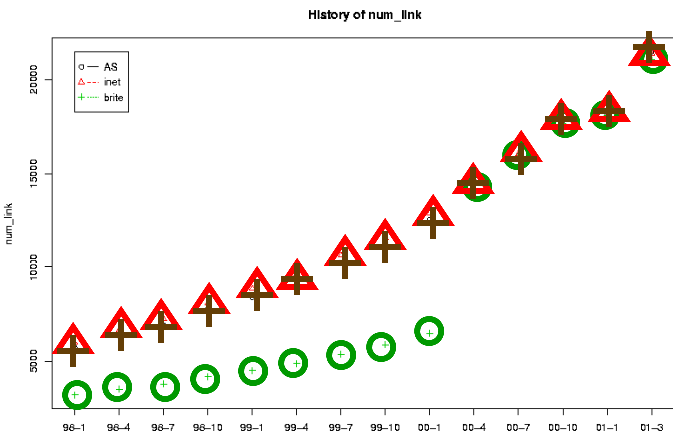

Date

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Matching Cardinality

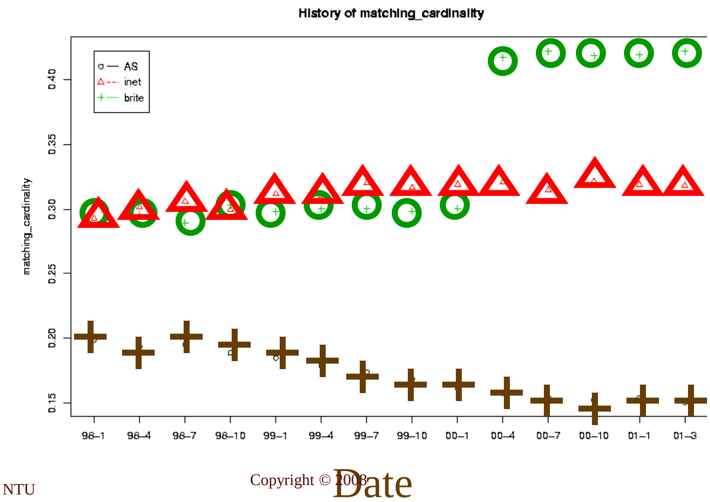

Polly @ NTU

55

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Matching Cardinality What?

- G = (V, E)
- M
  - A subset of E
  - No 2 edges share the same end nodes
- Matching Cardinality
  - Maximum Cardinality of Matching (MCM)
  - |Largest possible M| / |E|

Polly @ NTU

56

Summary of Background

- Forget about the heuristic one
- Structural ones
  - Miss power-law features
- Power-law ones
  - Miss other features
  - But what features?

No Idea!

- Try to look into individual metrics
  - Doesn't help much
  - A bit information here, a bit there
- Tons of metrics to compare graphs!
- Will never end this way!!

Identifying Internet Topology

Random Graphs
Power law
Practical Model

Spectral Analysis

- On the Spectrum and Structure of Internet Topology Graphs
- D. Vukadinovic, P. Huang, T. Erlebach
- In the proceedings of I2CS 2002.

Our Rationale

- So power-laws on node degree
  - Good
  - But not enough
- Take a step back
  - Need to know more
  - Try the extreme
  - Full details of the inter-connectivity
  - Adjacency matrix

The Research Statement

- Objective
  - Identify missing features
  - Hopefully the invariants
- Approach
  - Analysis on the **adjacency matrix**
    - can re-produce the complete graph from it
  - To begin with, look at its **eigenvalues**
    - Condensed info about the matrix

No Structural Difference

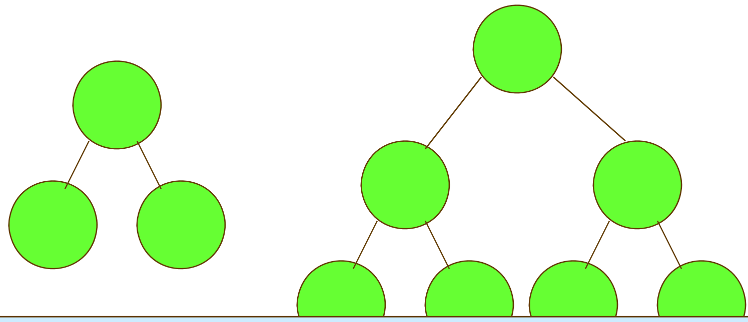

Eigenvalues are proportionally larger.
# of Eigenvalues is proportionally larger.

Normalization

- Normalized adjacency matrix
  - Normalized Laplacian
  - Eigenvalues always in [0,2]
- Normalized eigenvalue index
  - Eigenvalue index always in [0,1]
- Sorted in an increasing order
- N

Looking at a whole spectrum
Thus referred to as spectral analysis

Features of nls

- Independent of
  - size, permutation, mirror
- Similar structure <-> same nls
- Usually true but…

- Good candidate as the signature or fingerprint of graphs

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Tree vs. Grid

Trees

Eigen-values [0,2]

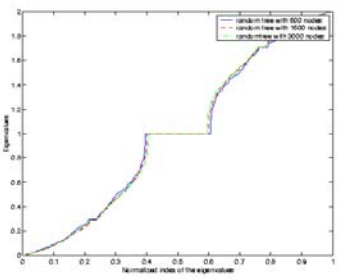

Eigenvalue Index [0,1]

Grids

Eigen-values [0,2]

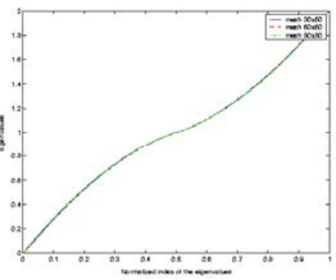

Eigenvalue Index [0,1]

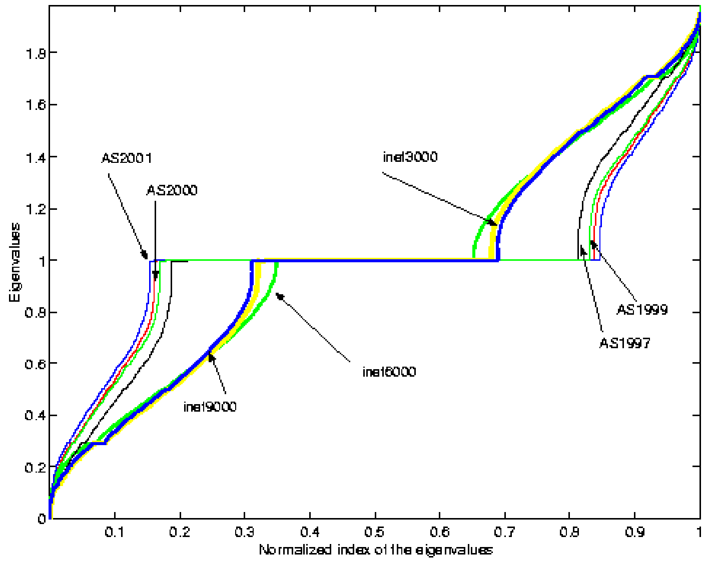

nls as Graph Fingerprint

- Unique for an entire class of graphs
  - Same structure ~ same nls
- Distinctive among different classes of graphs
  - Different structure ~ different nls
- Do have exception but rare

Spectral Analysis

- Qualitatively useful
  - nls as fingerprint
- Quantitatively?
  - Width of horizontal bar at value 1

Width of horizontal bar at 1

- Different in quantity for types of graphs
  - AS, Inet, tree, grid
  - Wider to narrower
- Polly: What is this?
- Theory colleague: Multiplicity 1, $m_G(1)$

Tight Lower Bound

- Polly: Any insight about this $m_G(1)$?
- Theory colleague: $m_G(1) \geq |P| - |Q| + |I|$

- Polly: P, Q, and I???
- Theory colleague: Components of the original graph…

For a Graph G

- P: subgraph containing **p**endant nodes
- Q: subgraph containing **q**uasi-pendant nodes

- Inner: G - P - Q
- I: **i**solated nodes in Inner
- R: Inner - I (R for the **r**est)

Enough Theory!

- Not really helping!
- P, Q, R, I in networking terms

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Physical Interpretation

- Q: high-connectivity domains, **core**
- R: regional alliances, partial **core**
- I: multi-homed leaf domains, **edge**
- P: single-homed leaf domains, **edge**

- Core vs. edge classification
  - A bit fuzzy
  - For the sake of simplicity

Validation by Examples

- Q
  - UUNET, Sprint, Cable & Wireless, AT&T
- R
  - RIPE, SWITCH, Qwest Sweden
- I
  - DEC, Cisco, HP, Nortel
- P
  - (trivial)

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Revisit the Theory

- $m_G(1) \geq |P| + |I| - |Q|$
- Correlation
  - Ratio of the edge components ->
  - Width of horizontal bar at value 1
- Grid, tree, Inet, AS graphs
  - Increasingly larger $m_G(1)$
  - Likely proportionally larger edge components

Evolution of Edge

Ratio of nodes in P

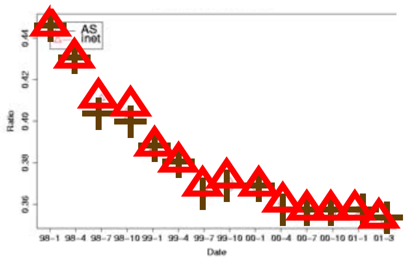

Ratio of nodes in I

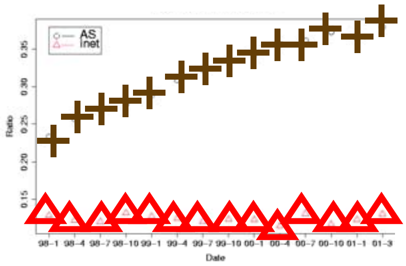

The edge components are indeed

Strong growth of I component ~ increasing number of multi-homed domains

Evolution of Core

Ratio of nodes in Q

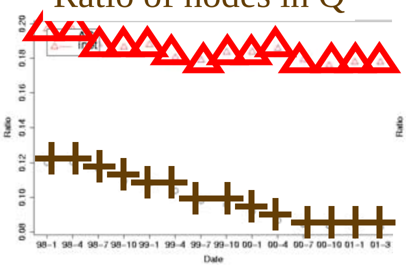

Ratio of links in Q

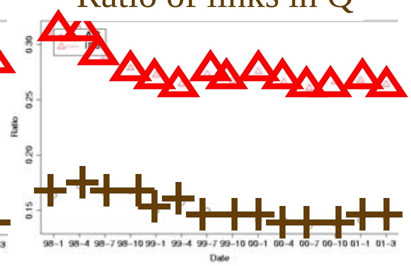

The core components get more links than nodes.

Core Connectivity

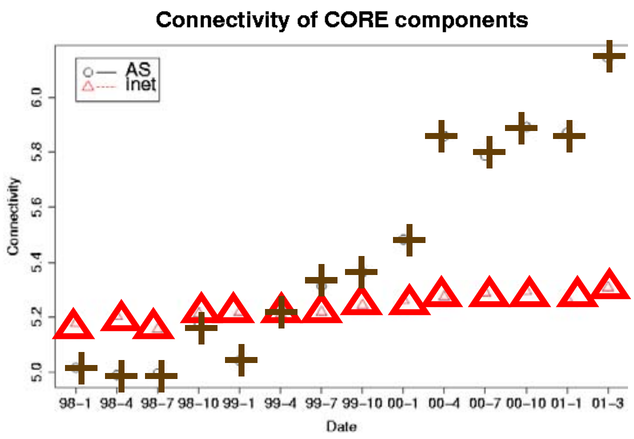

Polly @ NTU

79

What can we conclude here?

Edge and core behave differently. Structure is important!

Polly @ NTU

80

But is this going to change?

I.e., is this the invariant that we're looking for?

Search of Invariants

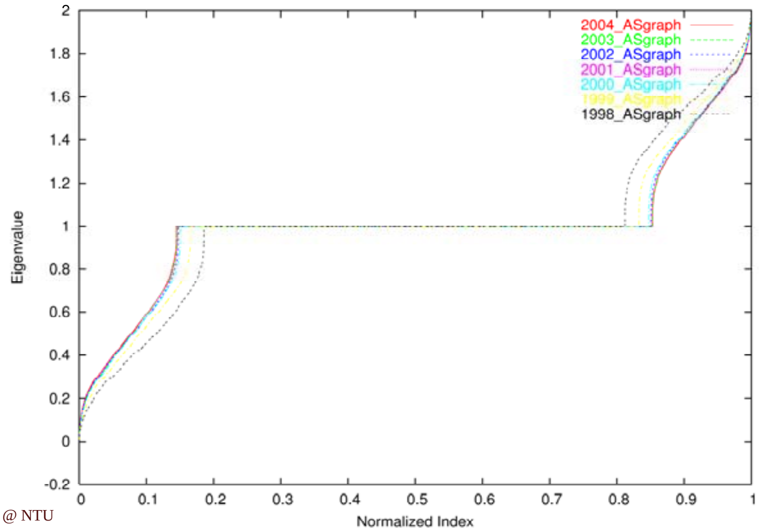

MG(1)

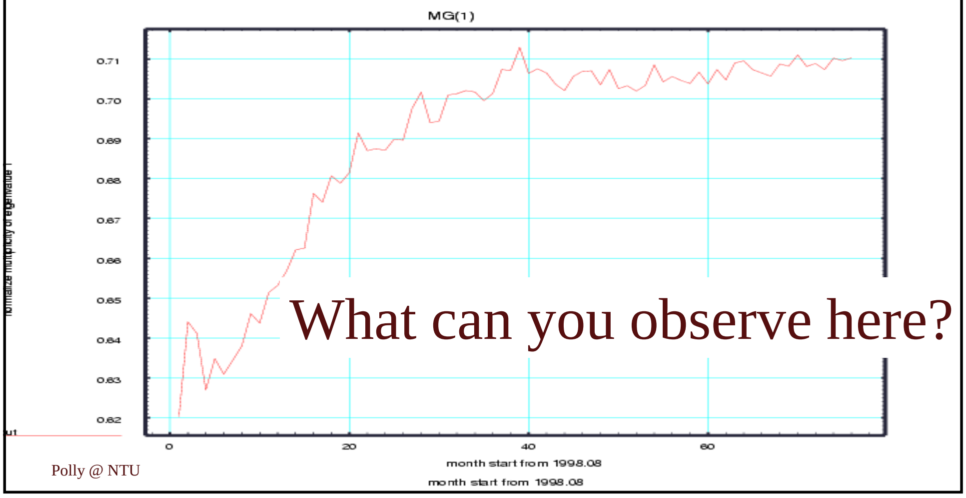

Polly @ NTU

Internet Economics Lesson #1

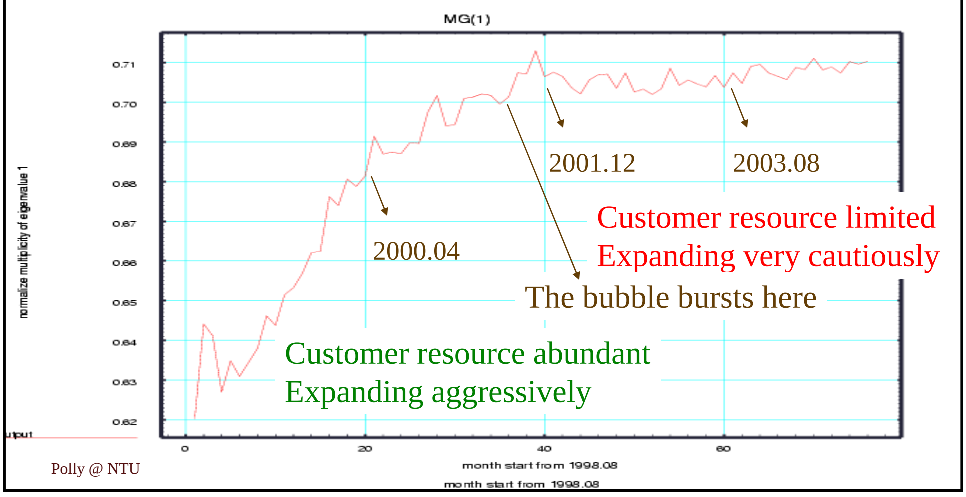

Polly @ NTU

Internet Economics Lesson #2

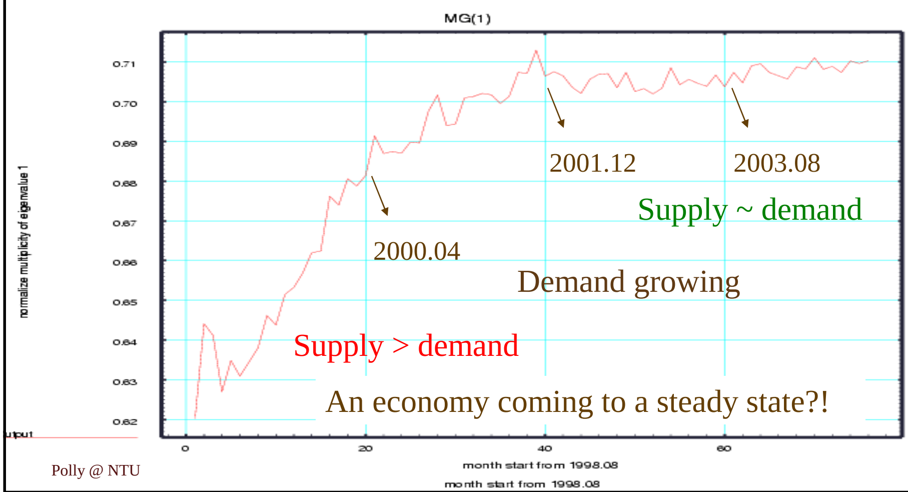

Oh my god, I can completely see the Internet economy here!

But is MG(1) the topology invariant?

Since there is no better invariant, we will take it for now.

Economists can probably confirm whether MG(1) will be the 'invariant' we are looking for…

Towards a Hybrid Model

- Form Q, R, I, P components
  - Average degree -> nodes, links
  - Radio of nodes, links in Q, R, I, P
- Randomly linking P-Q, I-Q, R-Q, R-R, Q-Q
  - With the preferential function identified connecting nodes from different components

Illustrated

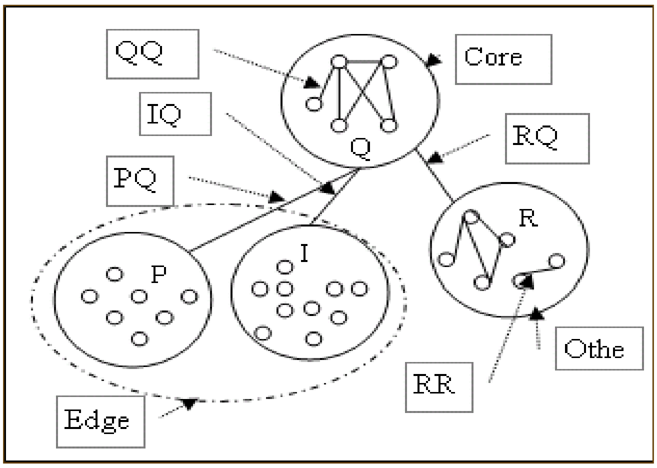

Our Premise

- Encompass both **statistical** and **structural** properties
- No explicit degree fitting

- Not quite there yet, but… do see an end
- … no real practical model at the moment (>”<)

Conclusion

- Firm theoretical ground
  - nls as graph fingerprint
  - Ratio of graph edge -> multiplicity 1
- Plausible physical interpretation
  - Validation by actual AS names and analysis
  - Explanation for AS graph evolution
- Framework for a hybrid model

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Observed Features

- Internet graphs have relatively **larger** **edge** components
- Although ratio of **core** components decreases, average degree of **connectivity increases**

Research Statement

- Objective
  - Identify missing features
  - Hopefully the invariants
- Approach
  - Analysis on the **adjacency matrix**
    - can re-produce the complete graph from it
  - To begin with, look at its **eigenvalues**
    - Condensed info about the matrix

Immediate Impact

- DDoS Attack Prevention
  - Efficient algorithm for optimal solution
  - Applicable only to graphs with **large edges**
  - Internet graphs!!!
  - ~50% faster
  - solution slightly better than the algorithm in SIGCOMM 2001

What Should You Do?

- Large-scale network required
  - Inet 3.0
- Hierarchical network required
  - GT-ITM
- Network not really important
  - Dumbbell

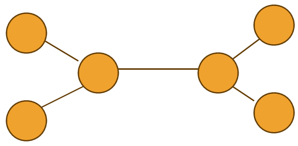

Network and Systems Laboratory
nslab.ee.ntu.edu.tw

Questions?