


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

Polly @ NTU Copyright © 2008 1



Today

General System Analysis
The Internet
Evaluating the Internet

Polly @ NTU Copyright © 2008 2

The Engineering Cycle

```

graph TD
    WC[Workload Characterization] --> PP[Performance Prediction]
    PP --> AS[Alternatives Selection]
    AS --> SI[System Instrumentation]
    SI --> UM[Usage Monitoring]
    UM --> WC
    subgraph Cycle
    WC
    PP
    AS
    SI
    UM
    end
    GS((General System))
  
```

- For a running system
 - Monitor the usage
 - Characterize the workload
 - Predict for the future
 - Revise original design
 - Instrument the changes
 - And back to the top

Polly @ NTU Copyright © 2008 3


Chicken or Egg

```

graph TD
    WC[Workload Characterization] --> PP[Performance Prediction]
    PP --> AS[Alternatives Selection]
    AS --> SI[System Instrumentation]
    SI --> UM[Usage Monitoring]
    UM --> WC
    subgraph Cycle
    WC
    PP
    AS
    SI
    UM
    end
    GS((General System))
  
```

- But where did it all start?
 - It depends
 - The Internet case?
 - The alternatives selection
 - But it's really just experts' intuition


Polly @ NTU Copyright © 2008 4



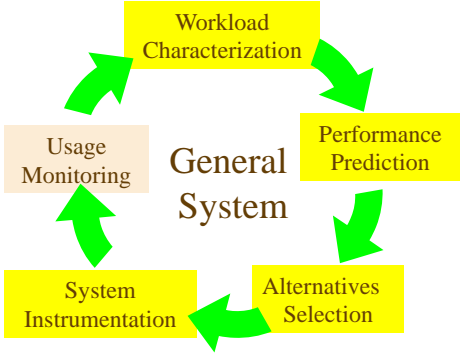
Sounds Easy...

Yah, easy to talk about it...
All Sorts of Problems in Practice...

Polly @ NTU Copyright © 2008 5



Monitoring the Usage



- Monitor the usage
 - Measurement methodology
 - Measurement tools
- Characterize the workload
- Predict for the future
- Revise original design
- Instrument the changes

Polly @ NTU Copyright © 2008 6

Characterizing the Workload

```

graph TD
    WC[Workload Characterization] --> PP[Performance Prediction]
    PP --> AS[Alternatives Selection]
    AS --> SI[System Instrumentation]
    SI --> UM[Usage Monitoring]
    UM --> WC
    subgraph General_System [General System]
        WC
        PP
        AS
        SI
        UM
    end
  
```

- Monitor the usage
- Characterize the workload
 - Modeling the measured data
 - The model needs to remain valid for data from taken at different time/location
- Predict for the future
- Revise original design
- Instrument the changes

Polly @ NTU Copyright © 2008 7

Predicting the Performance

```

graph TD
    WC[Workload Characterization] --> PP[Performance Prediction]
    PP --> AS[Alternatives Selection]
    AS --> SI[System Instrumentation]
    SI --> UM[Usage Monitoring]
    UM --> WC
    subgraph General_System [General System]
        WC
        PP
        AS
        SI
        UM
    end
  
```

- Monitor the usage
- Characterize the workload
- Predict for the future
 - Anticipate the user access pattern, demand increase
 - Evaluate the existing's capacity
- Revise original design
- Instrument the changes

Polly @ NTU Copyright © 2008 8

Designing the Alternatives


- Monitor the usage
- Characterize the workload
- Predict for the future
- Revise original design
 - If the current system won't live up to the challenge, how can it be changed...
 - Effective solutions
- Instrument the changes

Polly @ NTU Copyright © 2008 9

Instrumenting the Changes

- Monitor the usage
- Characterize the workload
- Predict for the future
- Revise original design
- Instrument the changes
 - Politics


Polly @ NTU Copyright © 2008 10



If This is the Telephone Network

- Monitor the usage
 - The big players place monitors all over the places in their own networks to collect data
- Characterize the workload
 - Fit the collected data to the well-known models
 - Human voice is Poisson
 - For some reason, the nature works this way.


Polly @ NTU Copyright © 2008 11



Telephone Network (cont)

- Predict for the future
 - Queuing theory:
 - Safe to supply λ_1 bandwidth for a call of average rate λ_1
 - $\lambda_1 + \lambda_2$ bandwidth for calls of average rate λ_1 and λ_2
 - Linear programming:
 - Given the max tolerable blocking rate, max the profit
- Revise original design
 - Mostly infrastructure-ral
 - I.e., rearranging or adding switches&cables
- Instrument the changes
 - Have full authority to change

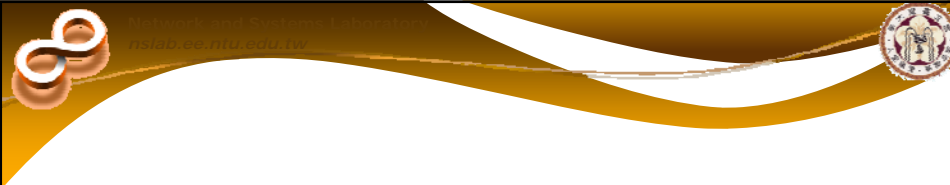
Polly @ NTU Copyright © 2008 12



No Real Difficulties

Quite a Profitable Business

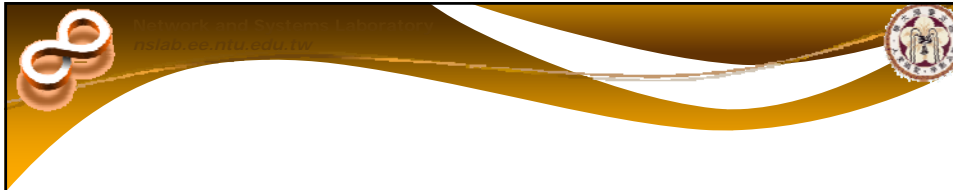
Polly @ NTU Copyright © 2008 13



Is the data network as profitable?

Let's review the data network.

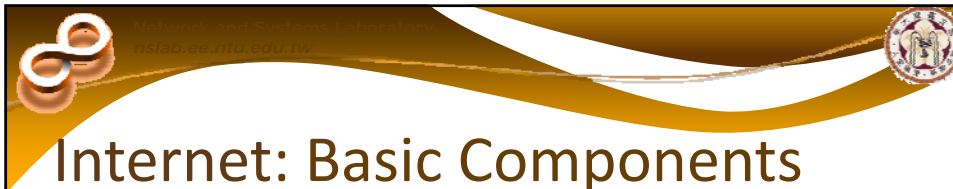
Polly @ NTU Copyright © 2008 14



Today

General System Analysis
The Internet
Evaluating the Internet


Polly @ NTU Copyright © 2008 15



Internet: Basic Components

- Think the **postal system**
- Nodes
 - End hosts and less number of routers
 - **Homes** and local/remote **post offices**
- Links
 - Connecting nodes (Access net, Ethernet, T1, T3, OC3, OC12, etc)
 - **Roads/streets** between homes and post offices


Polly @ NTU Copyright © 2008 16



Internet: Basic Constructions

- Packets
 - Destined to IP addresses (129.132.66.28)
 - Destined to **postal addresses** (1, Sec. 4 Roosevelt Rd.)
- Protocols
 - Packets sent with TCP (reliable)
 - Packets sent with **registered mail with confirmation**
 - But **no congestion control**
 - Other protocols...

Polly @ NTU Copyright © 2008 17



Internet Protocol Stack

- **Application:** supporting network applications
 - FTP, SMTP, HTTP
- **Transport:** host-host data transfer
 - TCP, UDP
- **Network:** routing of datagrams from source to destination
 - IP (addressing, routing, forwarding)
- **Link:** data transfer between neighboring network elements
 - Error Checking, MAC, Ethernet
- **Physical:** bits "on the wire"

application
transport
network
link
physical

Polly @ NTU Copyright © 2008 18

Data Network Research

(Side-Bar)

- **Application**
 - HTTP evolution
 - Web caching
- **Transport**
 - TCP evolution
- **Network**
 - Unicast routing
 - Multicast routing

→ • **P2P**

- Distributed directory for effective searching

• **Mobile and Wireless**

- Isolating drops due to bit error for TCP

• **QoS**

- = Problem Detection (loss congestion)
- Handling the hidden terminal problem (end-to-end, router-assisted)

Polly @ NTU Copyright © 2008 19

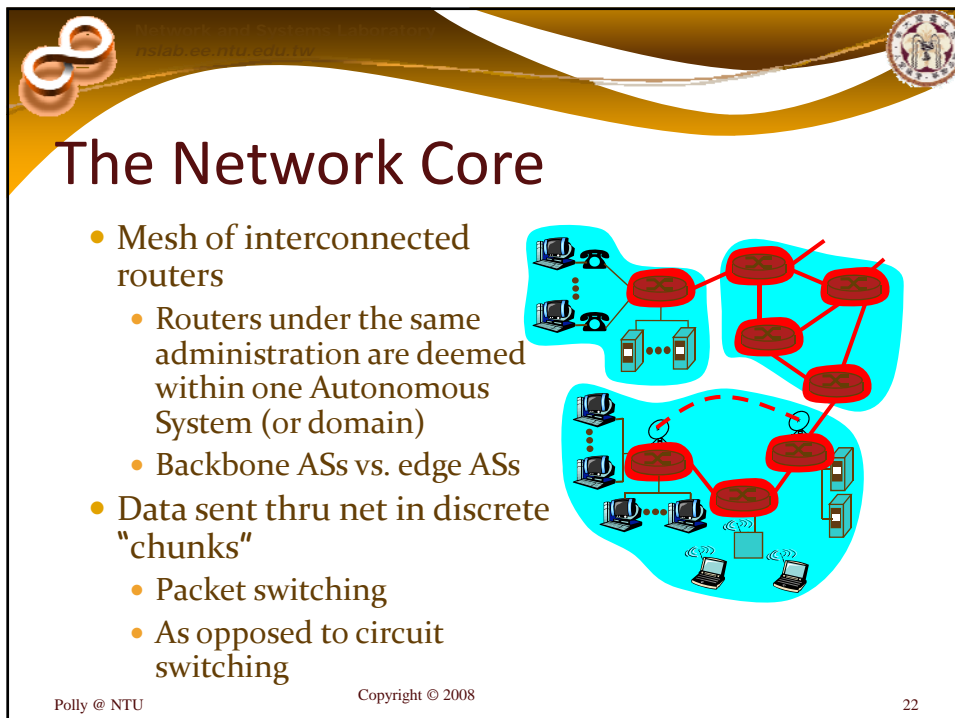
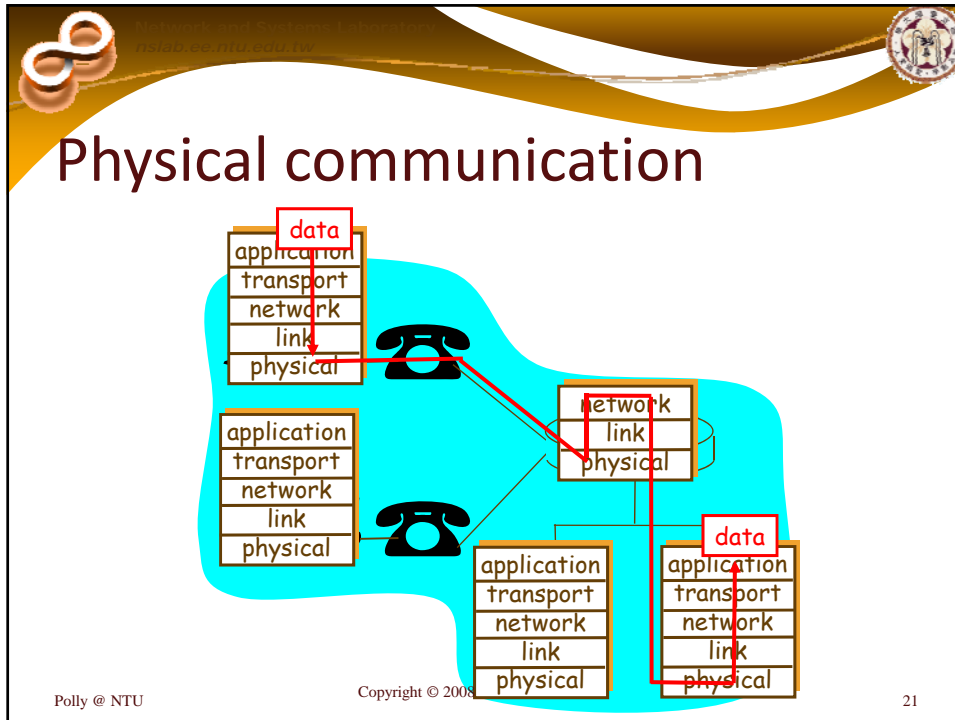
Internet Protocol Stack

(Back to the Topic)

- **Application:** supporting network applications
 - FTP, SMTP, HTTP
- **Transport:** host-host data transfer
 - TCP, UDP
- **Network:** routing of datagrams from source to destination
 - IP (addressing, routing, forwarding)
- **Link:** data transfer between neighboring network elements
 - Error Checking, MAC, Ethernet
- **Physical:** bits "on the wire"

application
transport
network
link
physical

Polly @ NTU Copyright © 2008 20



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


Polly @ NTU

Copyright © 2008

23

Internet AS Hierarchy

Inter-AS border (exterior gateway) routers

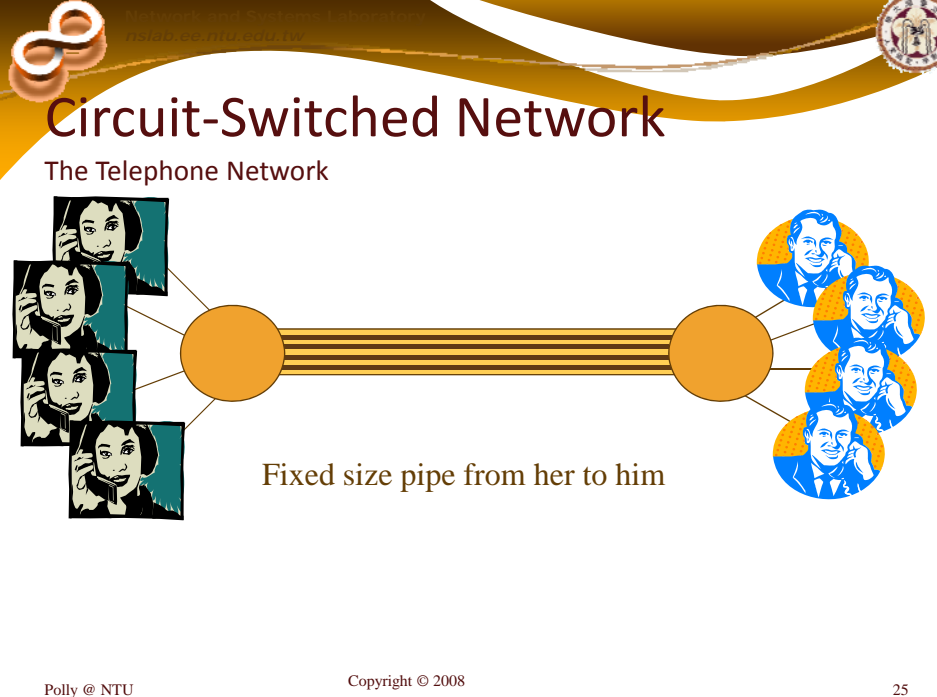


Intra-AS interior (gateway) routers

Polly @ NTU

Copyright © 2008

24

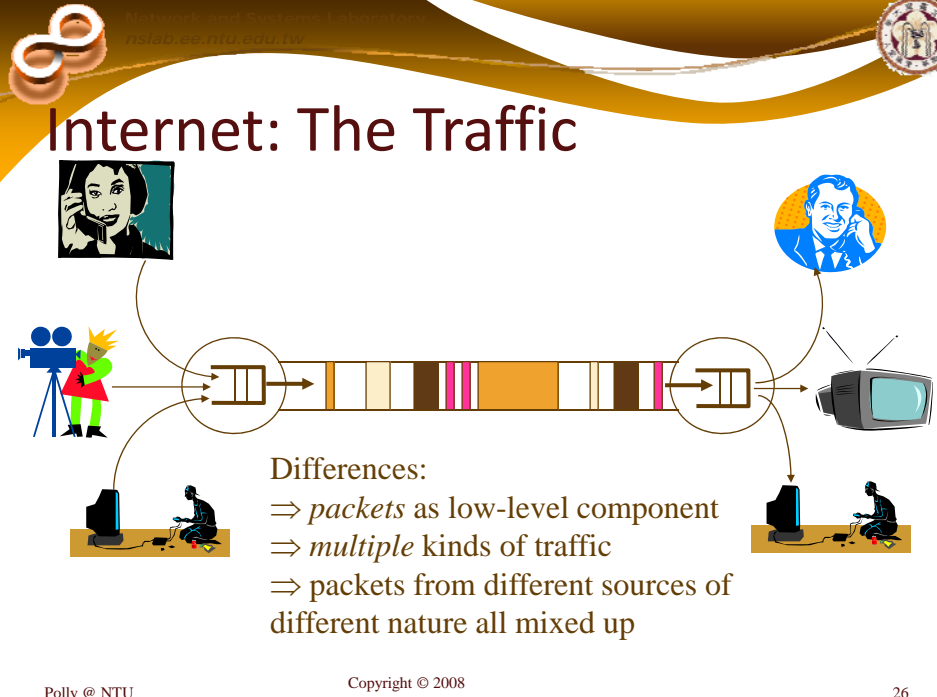


Circuit-Switched Network
The Telephone Network

Fixed size pipe from her to him

Polly @ NTU Copyright © 2008 25

The diagram illustrates a circuit-switched network. On the left, four female icons representing callers are connected to a central orange circular node. On the right, four male icons representing recipients are connected to another central orange circular node. These two nodes are connected by a thick, multi-line orange pipe, representing a dedicated circuit. The text 'Fixed size pipe from her to him' is placed below the pipe.




Internet: The Traffic

Differences:
⇒ packets as low-level component
⇒ multiple kinds of traffic
⇒ packets from different sources of different nature all mixed up

Polly @ NTU Copyright © 2008 26

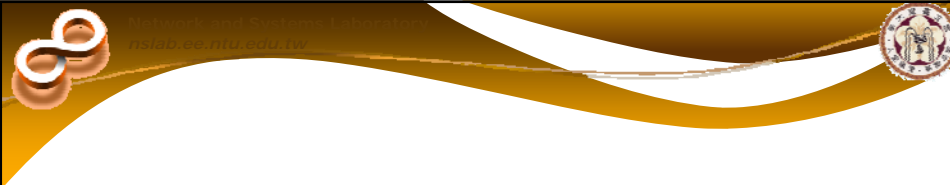
The diagram illustrates packet-switched traffic on the Internet. On the left, three icons (a woman on a phone, a person with a video camera, and a person at a computer) are connected to a circular node. On the right, three icons (a man on a phone, a television, and a person at a computer) are connected to another circular node. Between these nodes is a long horizontal pipe filled with a sequence of multi-colored rectangular blocks representing packets. The text 'Differences:' is followed by three bullet points describing packet-switched traffic.



Let's come back to this question:

Is the data network as profitable?


Polly @ NTU Copyright © 2008 27



Today

General System Analysis
The Internet
Evaluating the Internet


Polly @ NTU Copyright © 2008 28



What If This is the Internet

- Monitor the usage
 - The big players place monitors all over the places in their own networks to collect data
 - Would this give you representative data?
- Characterize the workload
 - Fit the collected data to the well-known models
 - Human voice is Poisson
 - Are Web browsing, Email, P2P, etc traffic Poisson?


Polly @ NTU Copyright © 2008 29



Internet (cont)

- Predict for the future
 - Queuing theory:
 - Save to supply λ_1 bandwidth for a call of average rate λ_1
 - $\lambda_1 + \lambda_2$ bandwidth for calls of average rate λ_1 and λ_2
 - Average, a good measure? Does traffic add up?
- Revise original design
 - Mostly infrastructure-ral
 - Still infrastructure-ral?
- Instrument the changes
 - Have full authority to change
 - Can the big players dictate?

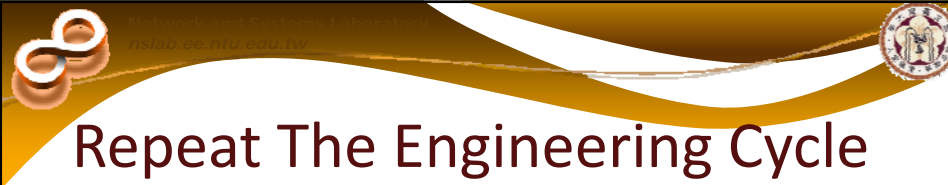
Polly @ NTU Copyright © 2008 30



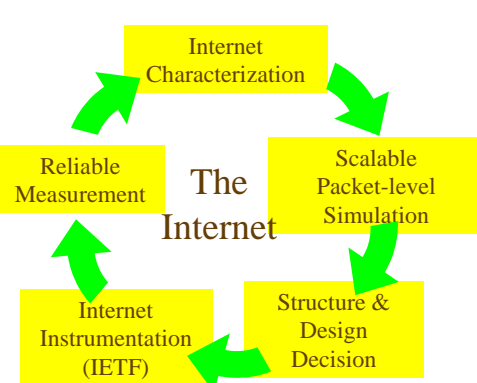
For the Most of These Questions

The Answer is NO

Polly @ NTU Copyright © 2008 31



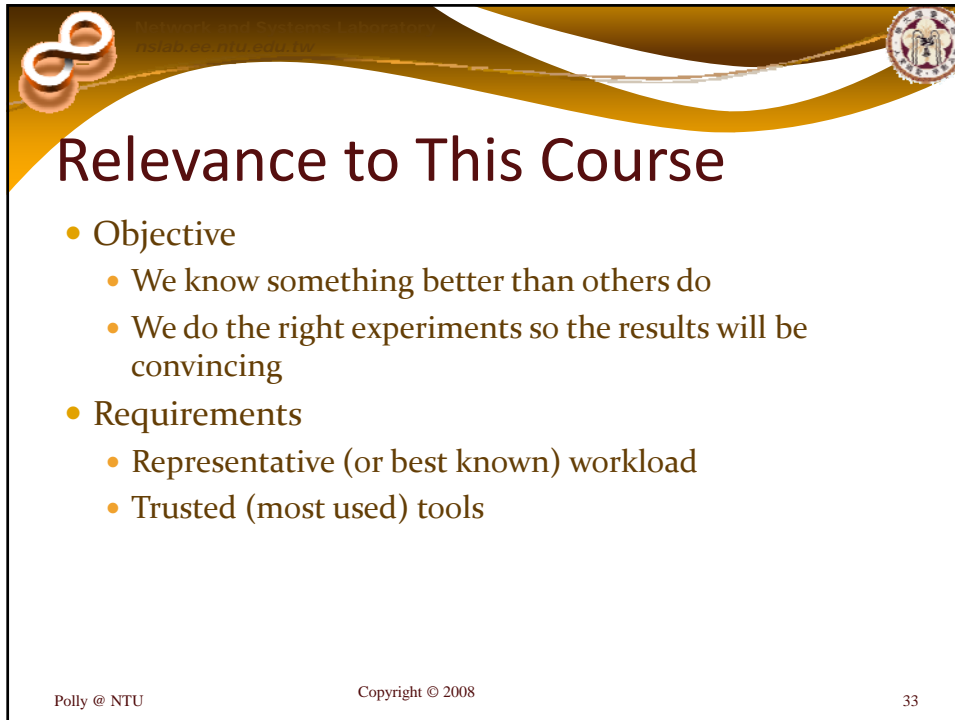
Repeat The Engineering Cycle



The Internet

- For the Internet
 - Monitor the usage
 - Passive and active measurement
 - Characterize the workload
 - Traffic, topology, routing errors, access pattern modeling
 - Predict for the future
 - Scalable simulation & testing tools
 - Revise original design
 - Protocol and Infrastructure
 - Instrument the changes
 - IETF

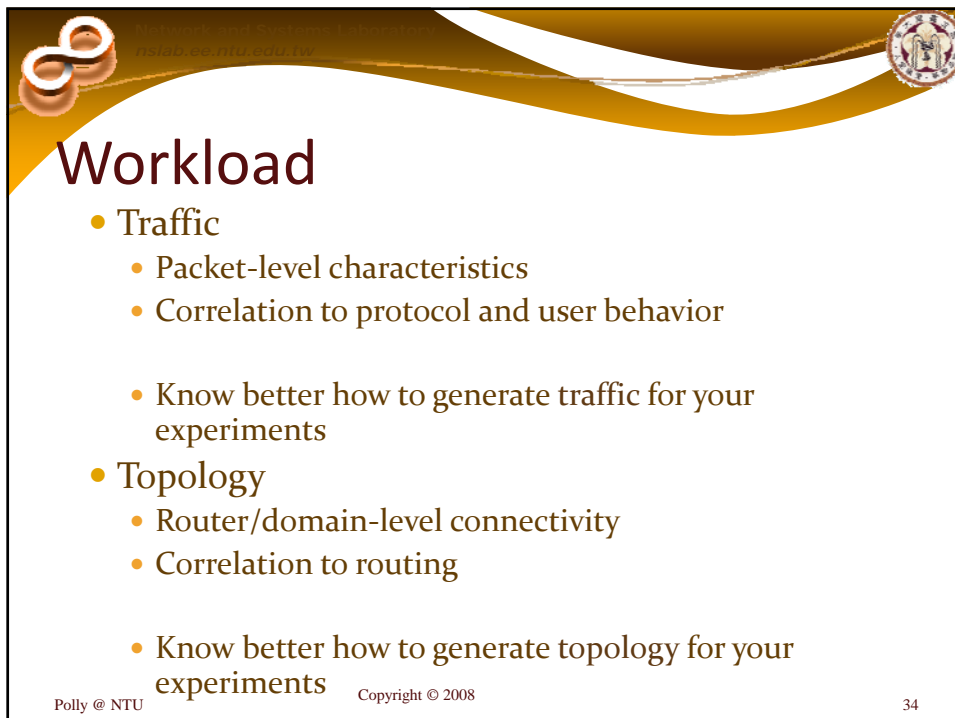
Polly @ NTU Copyright © 2008 32



Relevance to This Course

- Objective
 - We know something better than others do
 - We do the right experiments so the results will be convincing
- Requirements
 - Representative (or best known) workload
 - Trusted (most used) tools

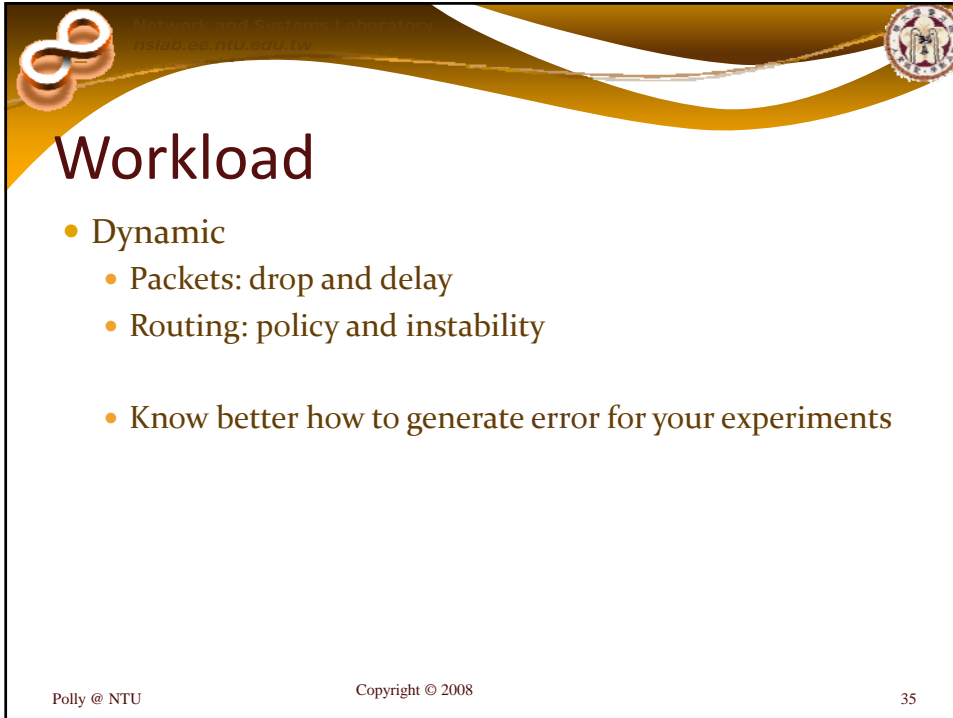
Polly @ NTU Copyright © 2008 33



Workload

- Traffic
 - Packet-level characteristics
 - Correlation to protocol and user behavior
 - Know better how to generate traffic for your experiments
- Topology
 - Router/domain-level connectivity
 - Correlation to routing
 - Know better how to generate topology for your experiments

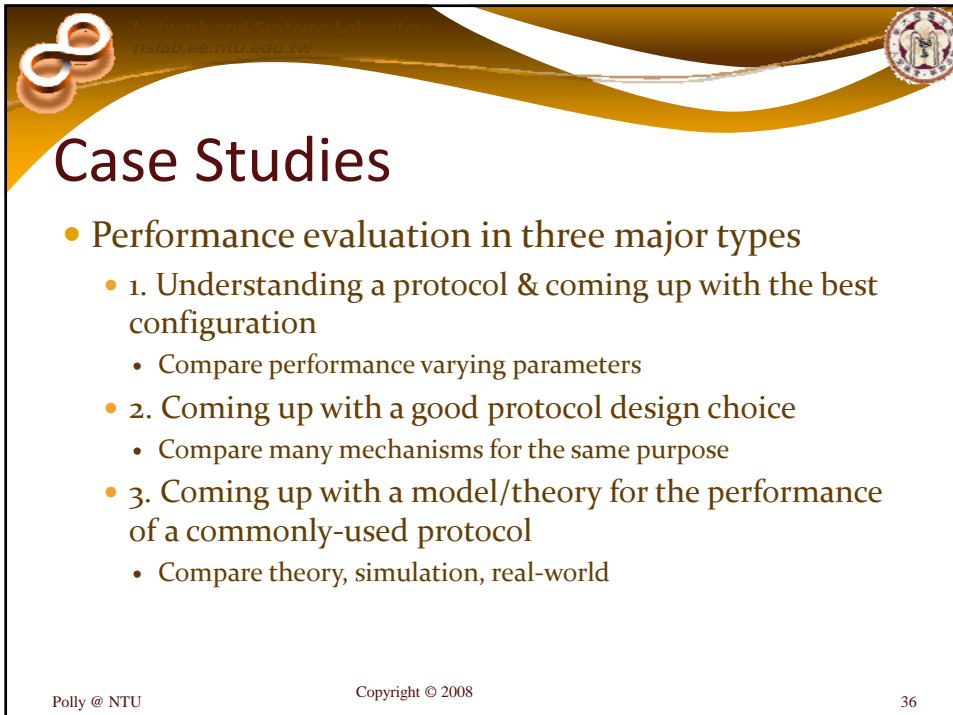
Polly @ NTU Copyright © 2008 34



Workload

- Dynamic
 - Packets: drop and delay
 - Routing: policy and instability
- Know better how to generate error for your experiments

Polly @ NTU Copyright © 2008 35



Case Studies

- Performance evaluation in three major types
 - 1. Understanding a protocol & coming up with the best configuration
 - Compare performance varying parameters
 - 2. Coming up with a good protocol design choice
 - Compare many mechanisms for the same purpose
 - 3. Coming up with a model/theory for the performance of a commonly-used protocol
 - Compare theory, simulation, real-world

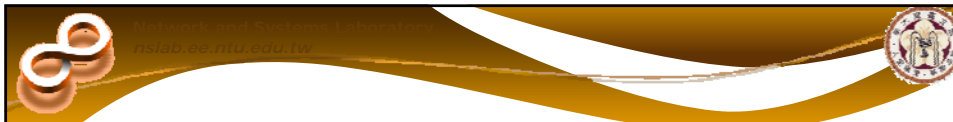
Polly @ NTU Copyright © 2008 36



Tools

- ns-2
 - About the most popular in the research community
 - Platform for cross-examination
- tcpdump
 - Not the only one but the most efficient one
 - Also the most popular one in the research community
- dummynet
 - Not the only one
 - But an easy and thus often-used one
- PlanetLab
 - The only one
 - Real Internet (A small subset)


Polly @ NTU Copyright © 2008 37



The Useful Theory

Statistics
Evaluation Methodology


Polly @ NTU Copyright © 2008 38



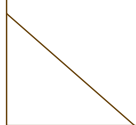
Keyword: Heavy-Tailed

- It turned out computer processes tend to be heavy-tailed or power-law distributed!
 - CPU time consumed by Unix processes
 - Size of Unix files
 - Size of compressed video frames
 - Size of FTP bursts
 - Telnet packet interarrivals
 - Size of Web items
 - Ethernet bursts

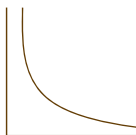
Polly @ NTU Copyright © 2008 39



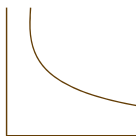
Illustrated



$y = a - b x$ linear



$y = a e^{-bx}, b > 0$ exponential




$y = a / x^b, b > 0$ heavy-tailed

Polly @ NTU Copyright © 2008 40



How to tell when you see one?


Polly @ NTU Copyright © 2008 41



Review Some Statistics

- Density vs. Distribution
- Poisson
- Exponential
- Pareto
- Self-similarity...

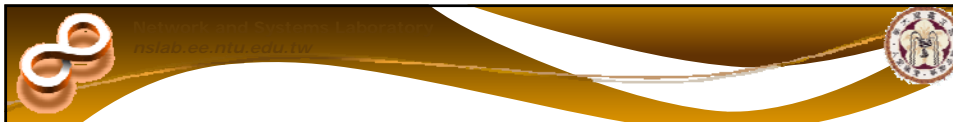
Polly @ NTU Copyright © 2008 42



Density vs. Distribution

- Density is the probability of certain events to happen
 - $f(x)$
- Distribution is usually referred to as the accumulative density
 - $f(0)+f(dz)+f(2*dz)+...+f(x)$
 - $F(x) = \int_{0 \rightarrow x} f(z) dz$

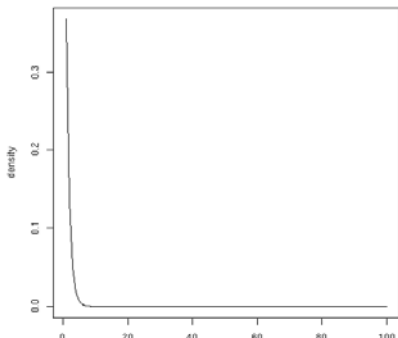
Polly @ NTU Copyright © 2008 43



Exponential

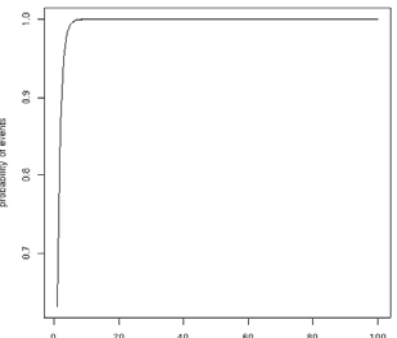
- # of time units between events
- $f(x) = ae^{-bx}$

Density: Exponential

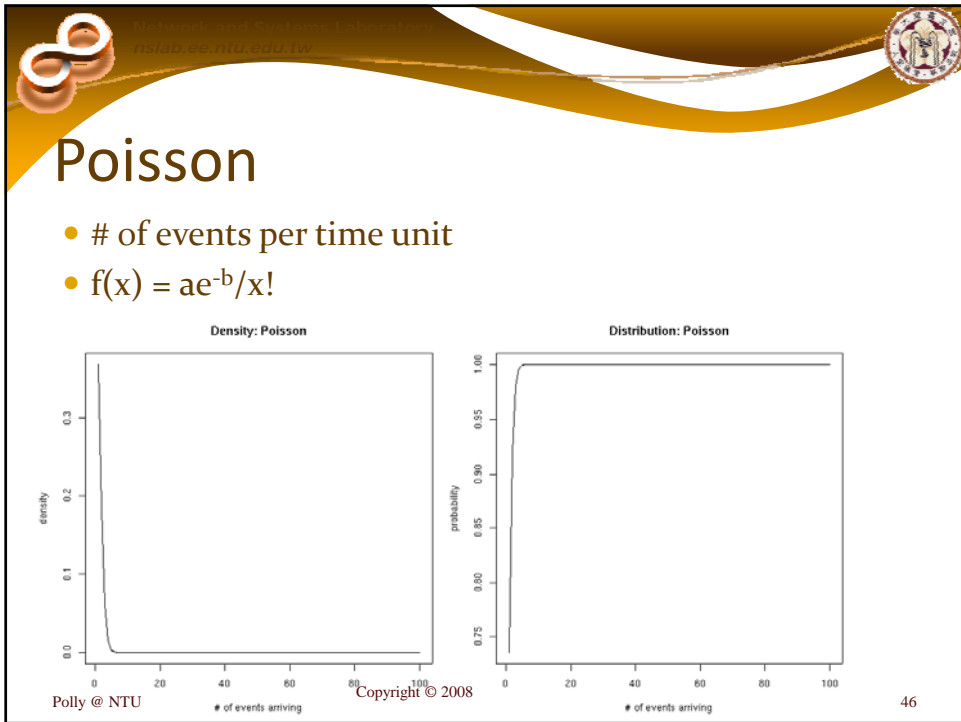



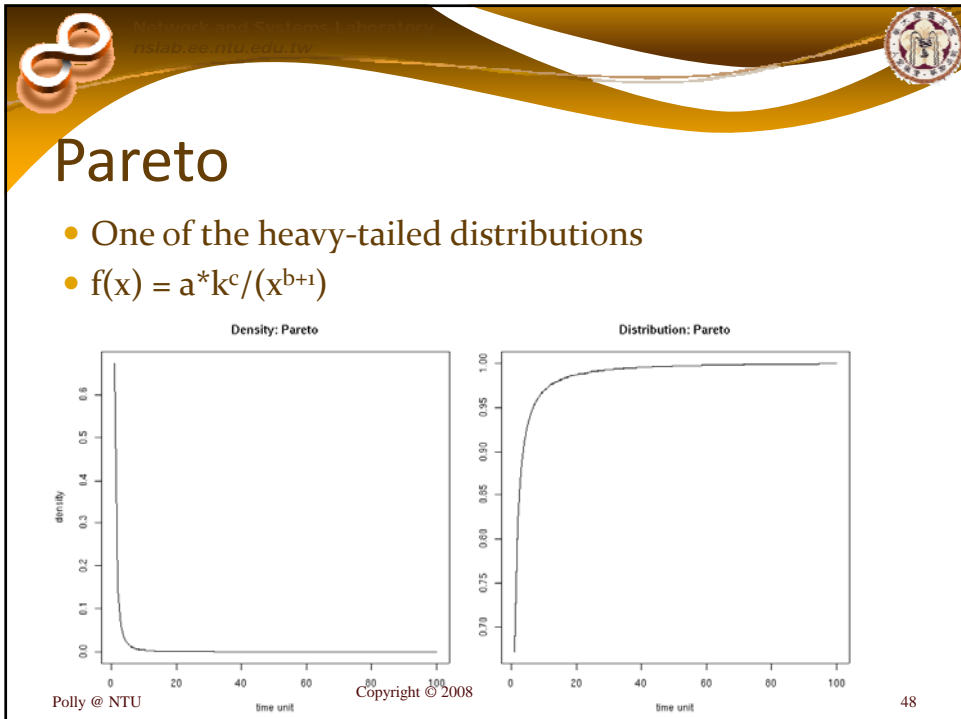

Polly @ NTU Copyright © 2008

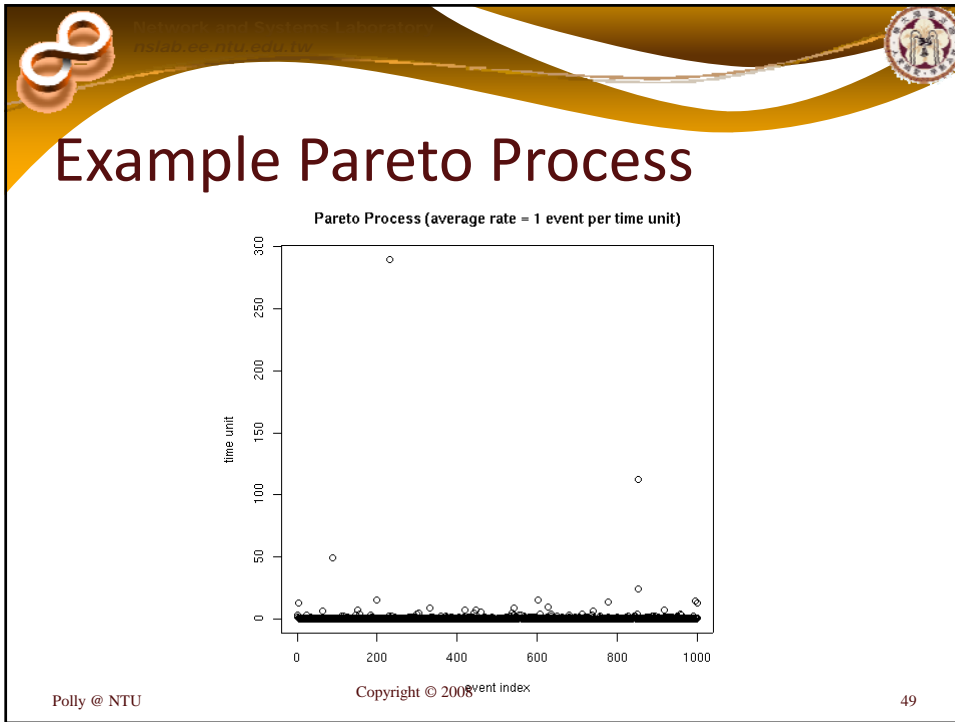
Distribution: Exponential



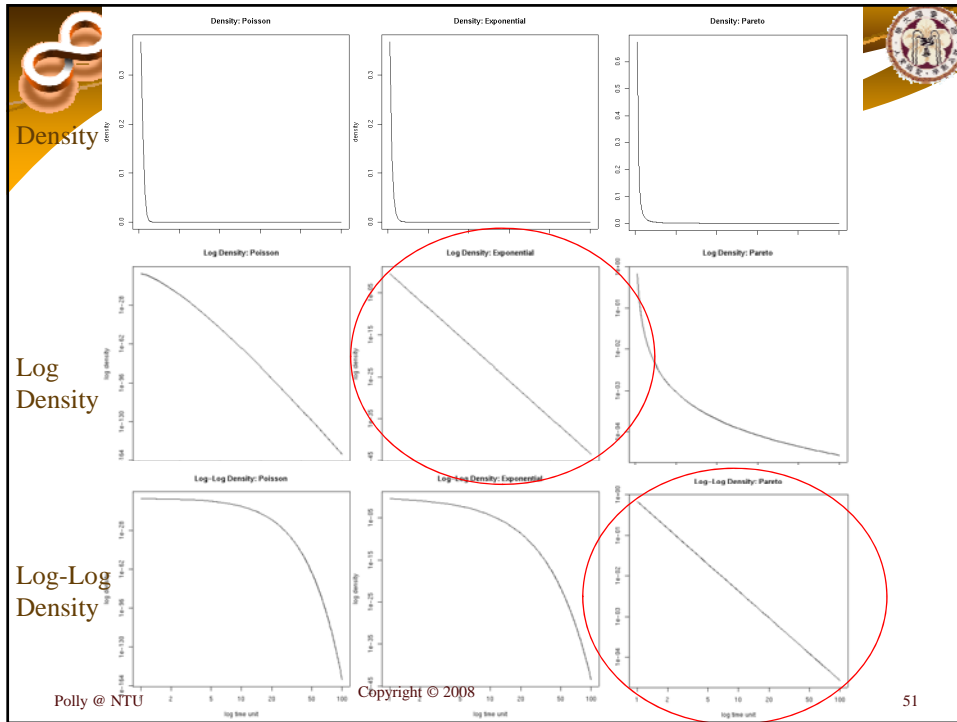
44







- ## Distinguishing Them
- Density
 - Log density
 - Log-log density
- Polly @ NTU Copyright © 2008 50



Now, what's with self-similarity...

Polly @ NTU Copyright © 2008 52

Teletraffic vs. Data traffic

- Teletraffic
 - Exp Exponential
- Data traffic
 - Heavy tailed

Polly @ NTU Copyright © 2008 53

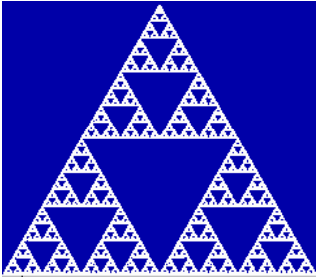
Measured Internet Traffic vs. Poisson Process

- Performance problem every 1-2 hours (network lags!)
- Profitable business?
- High variability (bursty)
- Long-range dependency
- Self-similarity

Polly @ NTU Copyright © 2008 54

Self-similarity?

- Distributions of #packets/unit look alike in **different time scale**



Sierpinski Triangles

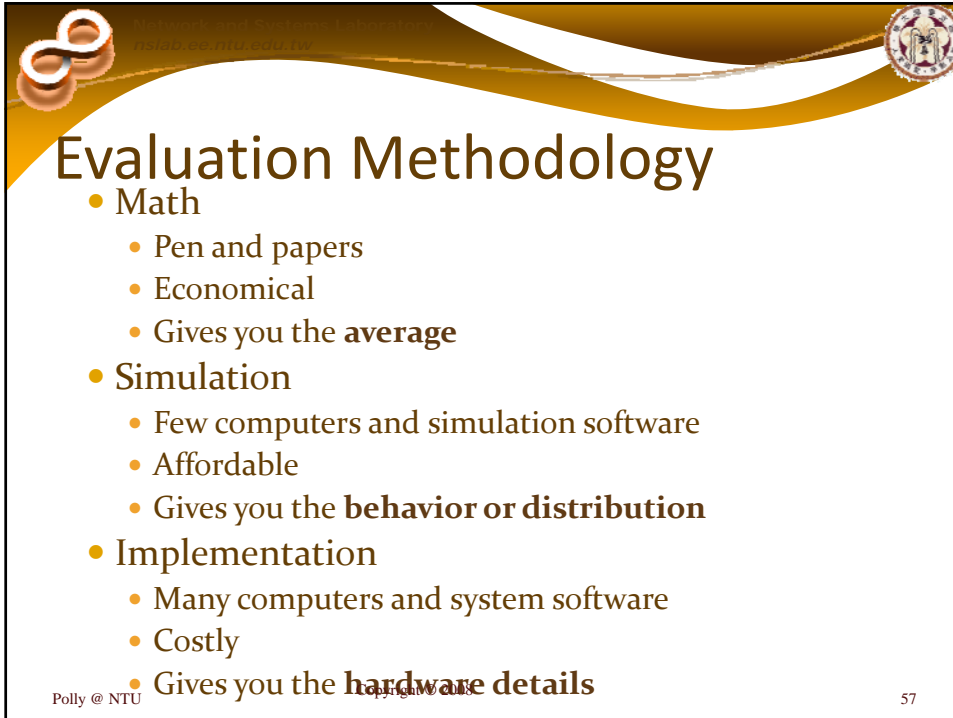
Polly @ NTU Copyright © 2008 55

The slide features a decorative header with a gold and white wavy background. On the left, there are three interlocking rings. On the right, there is a circular logo. The main content is centered on a white background. The title 'Self-similarity?' is in a dark red serif font. Below it is a bullet point with the text 'Distributions of #packets/unit look alike in different time scale', where 'different time scale' is in red. The central image is a blue square containing a white Sierpinski triangle fractal. Below the image is the text 'Sierpinski Triangles' in a dark red serif font. At the bottom, there are three small text elements: 'Polly @ NTU' on the left, 'Copyright © 2008' in the center, and '55' on the right.

End of Statistics Review

Polly @ NTU Copyright © 2008 56

The slide features a decorative header with a gold and white wavy background. On the left, there are three interlocking rings. On the right, there is a circular logo. The main content is centered on a white background. The title 'End of Statistics Review' is in a dark red serif font. At the bottom, there are three small text elements: 'Polly @ NTU' on the left, 'Copyright © 2008' in the center, and '56' on the right.



Evaluation Methodology

- Math
 - Pen and papers
 - Economical
 - Gives you the **average**
- Simulation
 - Few computers and simulation software
 - Affordable
 - Gives you the **behavior or distribution**
- Implementation
 - Many computers and system software
 - Costly
 - Gives you the **hardware details**

Polly @ NTU Copyright © 2008 57



Which should you use?

Depends on what you care for the problem
in hand!

Polly @ NTU Copyright © 2008 58



Assumptions


- It's OK to leave out details
- But
 - You need to be clear what details you leave out.
 - You need to argue it is OK to leave those details out for now.
 - And you are working on including those details and the results will be available in the future.

Polly @ NTU Copyright © 2008 59



Questions?


Polly @ NTU Copyright © 2008 60



Traffic Papers

- V. Paxson, and S. Floyd, Wide-Area Traffic: The Failure of Poisson Modeling. IEEE/ACM Transactions on Networking, Vol. 3 No. 3, pp. 226-244, June 1995
- M. E. Crovella and A. Bestavros, Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes. IEEE/ACM Transactions on Networking, Vol 5, No. 6, pp. 835-846, December 1997
- Anja Feldmann; Anna C. Gilbert; Polly Huang; Walter Willinger, Dynamics of IP traffic: A study of the role of variability and the impact of control. In the Proceeding of SIGCOMM '99, Cambridge, Massachusetts, September 1999
- Vern Paxson. End-to-end internet packet dynamics. ACM/IEEE Transactions on Networking, 7(3):277-292, June 1999

Polly @ NTU Copyright © 2008 61



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.

Polly @ NTU Copyright © 2008 62



Questions?

Polly @ NTU

Copyright © 2008

63