


# Complexity and Service Provider Networks in the 21st Century



SANOG V  
Dhaka, Bangladesh  
February 2005

David Meyer  
dmm@1-4-5.net  
<http://www.1-4-5.net/~dmm/talks/SANOGV>

# Agenda

---

- Goals and Objectives
- What is Complexity?
- Why do we Care?
- A Few Examples From Everyday Life
- Summary/Key Takeaways
- Q & A

# Goals and Objectives

---

- Goal: Socialize the understanding of complexity and related design principles in the Internet engineering community
  - Focus on higher level principles, many of which aren't well understood in the Internet engineering community
  - Point out where theory meets application
  - List some key takeaways for everyday use
- These slides:
  - <http://www.1-4-5.net/~dmm/talks/SANOGV>

# Before We Dive In To All Of This...

---

If you take one thing away from this talk, it should be this:

Keep your network designs and implementations as simple and understandable as possible

As we shall see, physics (and human nature) will be pushing you in the other direction

# So What is Complexity?

---

- Many ways to study complexity, including
  - Classical Computational Complexity
    - ▶ Deals with upper/lower bounds on time, memory, etc of a given algorithm
    - ▶ For example, a full mesh is said to have complexity  $O(n^2)$ , since you need  $n*(n-1)/2$  unique connections, and  $n*(n-1)/2 \sim c*n^2 = O(n^2)$
  - Many others
- In this talk we'll look at "Dynamical Systems Theory"
  - Focuses on a system's large scale, macroscopic properties
    - ▶ That is, those properties that emerge at scale
  - Applies to the Internet as a system
    - ▶ Which is what we are interested in from an operational perspective
  - What kind of problems can we analyze this way?
    - ▶ The convergence behavior of BGP
    - ▶ BGP UPDATE propagation
    - ▶ The behavior of IGPs under congestion
    - ▶ PWE3 Alarm "multiplication"
    - ▶ ...

# Why Do We Care?

---

Because, as we will see, this type of complexity and the dynamic behavior it gives rise to is a major factor in the cost of both buying and operating our networks

This idea was eloquently stated by Mike O'Dell and come to be known as "The Simplicity Principle"

# The Simplicity Principle

---

- The Simplicity Principle states that "Complexity is the primary mechanism which impedes efficient scaling, and as a result is the primary driver of increases in both capital expenditures (CAPEX) and operational expenditures (OPEX)."
  - Corollaries
    - ▶ KISS Principle
    - ▶ Law of Diminishing Returns
    - ▶ Related: Occam's Razor
- That is, if you don't keep designs and implementations as simple as possible, it is going to wind up costing you both more to operate your network (OPEX), and more to grow as your business grows (CAPEX)
  - But how simple is "simple as possible"?

# Well watch out

---

- Everything should be made as simple as possible, but not simpler.
  - Albert Einstein
- As we shall see, there is a close relationship between complexity and robustness
  - Complexity is an important source of robustness in your networks
    - Example: Multihoming
- Takeaway: Complexity is not inherently "bad"

# Where is this complexity coming from?

---

- Lots of sources. I'll focus on the non-linearity inherent in large systems such as the Internet
- Relatively recent work in the dynamical and non-linear systems field suggests that highly designed systems (such as the Internet) are "Robust yet Fragile"
- Sounds like a contradiction of sorts, doesn't it?


# Robust yet Fragile Systems?

---

- Basically, a system needs some degree of complexity (design) to achieve robustness, but it does so at the cost of increasing sensitivity to (hopefully) rare events
- However, while a system's robustness is bounded, its complexity is not.
  - This gives rise to a robustness-complexity distribution is heavy-tailed (we will see a picture of this on the next slide).
  - For example, if you build your network such that when you add a router, that router is more likely to be directly connected to a well connected router, the distribution of the number of circuits per router will be heavy-tailed [LUN2004].
  - Current POP design practices produce just this effect

# Well, what does this all of this mean?

---



Robustness v. Complexity

See "Complexity and Robustness", J.M Carlson and John Doyle, Proc. Natl. Acad. Sci. USA, Vol. 99, Suppl. 1, 2538-2545, February 19, 2002.

<http://www.pnas.org/cgi/doi/10.1073/pnas.012582499> for a nice introduction.

# Sources of Complexity

---

- There are two well-known principles from non-linear systems theory that apply here, namely:
  - Amplification
  - Coupling

# Amplification Principle

---

- The Amplification Principle states that there are non-linearities which occur at large scale (which are not observed at smaller scales).
- COROLLARY: In many large networks, even small things can and do cause huge events.
  - In system-theoretic terms, in large systems such as these, even small perturbations on the input to a process can destabilize the system's output.
- Takeaway: As your networks grow, even things that once looked small can turn out to be big!


# Amplification Examples

---

- Multiplication of BGP messages across the peering mesh of the Internet
  - For example, Tim Griffin's work [GRIFFIN2002] showing that a small amount of inter-connectivity causes the output of a routing mesh to be significantly more complex than its input.
- Alarm multiplication on pseudo-wires
  - Consider what happens when an OC192 (or so) full of PWs fails
- Takeaway: Look closely at aggregated and mesh structures in your network. They can be sources of unanticipated amplification.
- BTW, you see amplification in nature all the time

Think  $O(n!)$  convergence time for BGP is bad?

Well, it gets worse...



Tacoma Narrows Bridge

# Coupling Principle

---

- The Coupling Principle states that as things get larger, they often exhibit increased interdependence between components
- COROLLARY: The more events that simultaneously occur, the larger the likelihood that two or more will interact. This phenomenon has also been termed "unforeseen feature interaction" [Doyle and Willinger]
- Coupling can be either horizontal (within a protocol layer) or vertical (between layers)

# Coupling Principle Examples

---

- Canonical examples include Van Jacobson and Sally Floyd's work on the synchronization of various control loops, including routing updates and TCP (slow start synchronization)
- An important result of these observations is that coupling is closely related to synchronization
  - Injecting randomness into these systems is one way to reduce coupling -- This is the basis of RED/WRED
  - Another important method for reducing a system's coupling is ensure that local changes have only local effect
- Takeaway: Localize the impact of any design or configuration change

# Example: Achieving Five Nines

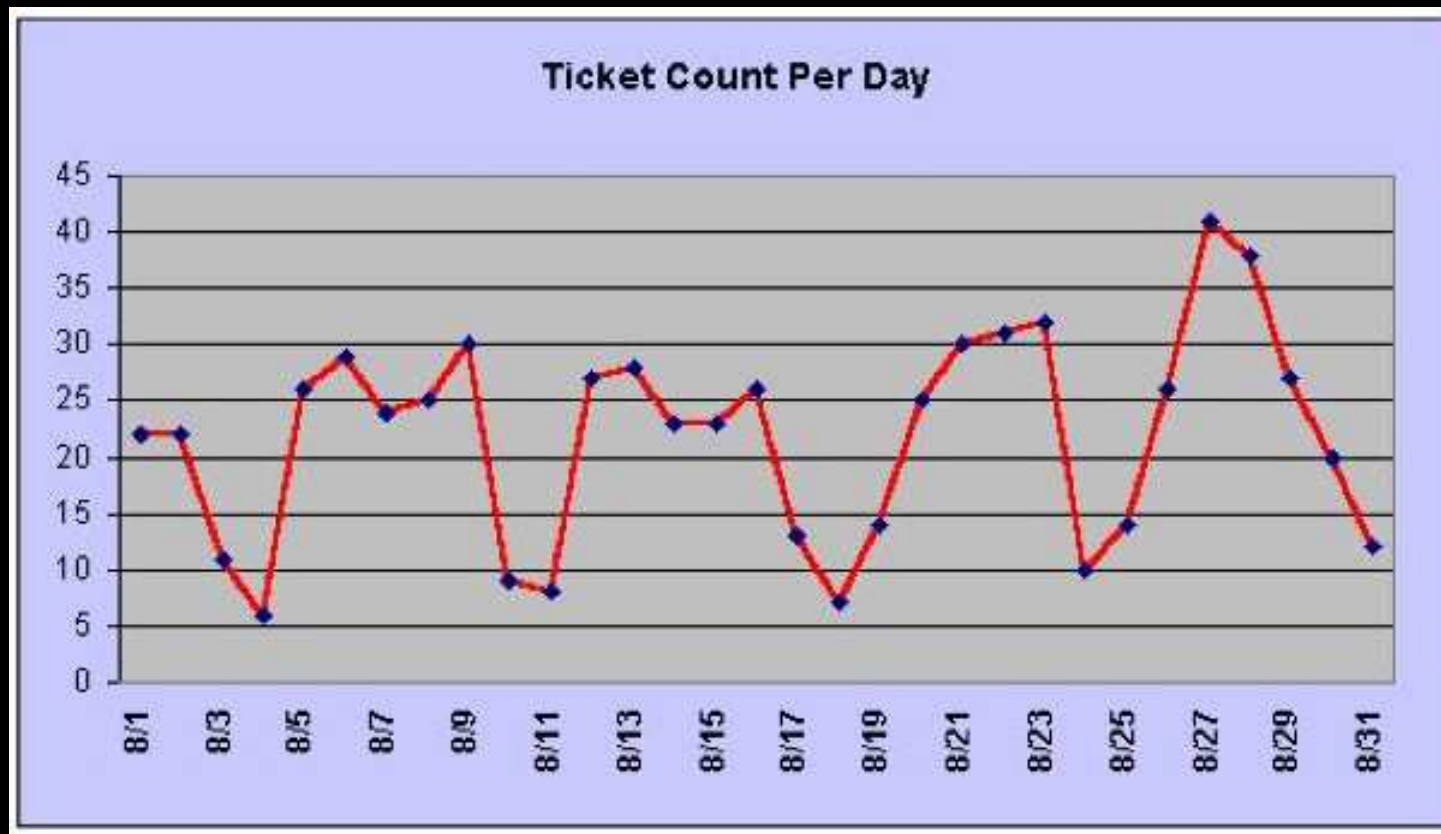
---

- OK, why is this so hard?
- First, note that 80% of unscheduled outages are caused by people or process errors [SCOTT, many others] --> There is only a 20% window in which to optimize component reliability.
  - Why? Because in order to increase component reliability, we typically add complexity (remember the relationship between robustness and complexity). This effectively narrows the 20% window.
- Bottom line: Optimization frequently increases complexity, which increases the likelihood of people and process failures

# Can it really be that...

---

80% of unscheduled outages are caused by people/process? Well...




Data and Image courtesy Vijay Gill/AOL, 2002

# Complexity/Robustness Spirals

---

The same phenomenon which makes five nines difficult to reach is also characterized as a "complexity/robustness" spiral [Doyle and Willinger], in which increases in complexity create further and more serious sensitivities, which then require additional robustness, which requires increases in complexity, which in turn requires additional robustness, and so on (hence the spiral).

# A "Well Known" C/R Spiral



The creature is called an "Ouroboros" (which means "devouring its tail"). See <http://www.dragon.org/chris/ouroboros.html> (URL courtesy Sean Doran)

# A Few Examples From Everyday Life

---

## ■ iBGP/BGP Topologies

- Amplification
  - Multiplication of messages through a mesh

## ■ Link-state routing protocols

- Amplification
  - Cascading failures in the face of congestion/hello timeouts
- Horizontal Coupling
  - Network congestion to routing stability

## ■ OAM Mapping

- Vertical Coupling

## ■ Pseudo-wire Alarm Aggregation

- Amplification/Vertical Coupling

# A Few Everyday Examples, cont

---

- L1/L2 Message Mapping/Interworking
  - Amplification/Coupling
  
- TDM over IP [TDMOIP]
  - Vertical Coupling (Synchronization)
  
- IMS (3GPP IP Multimedia Subsystem)
  - Vertical Coupling
    - Services to transport

# Summary

---

- First, keep things as simple as you can
  - Remembering that complexity isn't inherently bad
- Clearly understand (and document) the potential for amplification to occur
  - Look closely at aggregated and mesh structures in your network
    - And remember that as things scale up, what was small can become very large
- Clearly understand (and document) the where coupling occurs in your design and implementation
  - Localize the impact of any design or configuration change
  - Can also occur in operational procedures

# Finally

---

- Be Active in Standards Bodies
- There are a many SDOs (Standards Developing Organizations) which build technologies that directly effect you
- IETF, ITU-T, 3GPP, ATIS, ETSI, ...
- Make your requirements known

# Questions?

---

Thanks!

■ More information:

- [http://www.1-4-5.net/~dmm/complexity\\_and\\_the\\_internet](http://www.1-4-5.net/~dmm/complexity_and_the_internet)

# References

---

- [Doyle and Willinger] "Robustness and the Internet: Theoretical Foundations", John C. Doyle, et al. Draft, March, 2002.
- [Doyle and Willinger] "Robustness and the Internet: Design and evolution", Walter Willinger and John Doyle, 2002.
- [GRIFFIN2002] "What is the Sound of One Route Flapping", Timothy G. Griffin, IPAM Workshop on Large-Scale Communication Networks: Topology, Routing, Traffic, and Control, March, 2002.

# References

---

- [LUN2004] "A First-Principles Approach to Understanding the Internet's Router-level Topology", SIGCOMM 2004.
- [SCOTT] "Making Smart Investments to Reduce Unplanned Downtime", D. Scott, Tactical Guidelines, TG-07-4033, Gartner Group Research Note, March 1999.
- [TDMOIP] draft-ietf-pwe3-tdmoip-02.txt