CSx55: Distributed Systems [Architectures/Topology]

Decentralized topologies

Nodes without a weave
  Like wings without flight

Connect them near and far
  And watch it soar
Imbuing each with a nifty quirk
  Traits that make them tick

This you probably knew,
  Your networks tell a lot about you

Shrideep Pallickara
Computer Science
Colorado State University

Frequently asked questions from the previous class survey

- If a site pre-loads a page, is that stateful?
- Are there times where threads-per-connection is worse than threads-per-request?
- Is a client stateless when it fully opts out of any “cookies”?
- How do we design a system to safely isolate slow processes?
- Was there ever an IPv5?

Topics covered in this lecture

- Decentralized architectures
- Topologies
  - Regular graphs
  - Random graphs
  - Small world graphs
  - Power law networks

Ryan Stern and Shrideep Pallickara. On the Role of Topology in Autonomously Coping with Failures in Content Dissemination Systems. Proceedings of the *ACM Cloud and Autonomic Computing Conference*. Miami, USA. 2013.

Decentralized Architectures

Decentralized architectures

- Server may be split up into logically equivalent parts
  - Each part operates on its share of the dataset
  - Balance the load
- Interaction between processes is **symmetric**
  - Each **peer** acts as a client and a server

Structured Peer to Peer Architectures: Distributed hash tables

- **Data items** are assigned an identifier from a large random space
  - 128-bit UUIDs (2^{128} or 10^{38}) or 160-bit SHA-1 digests {2^{160} or 10^{48}}
- **Nodes** are also assigned a number from the ***same*** identifier space

Crux of the DHT problem

- Implement an efficient, **deterministic** scheme to *map data item to node*
- When you look up a data item?
  - Network address of node holding the data is returned

Unstructured P2P networks typically rely on random graphs

- Maintain connections to randomly chosen live nodes
- To locate a data item
  - **Flood** the network

Hierarchical organization of nodes

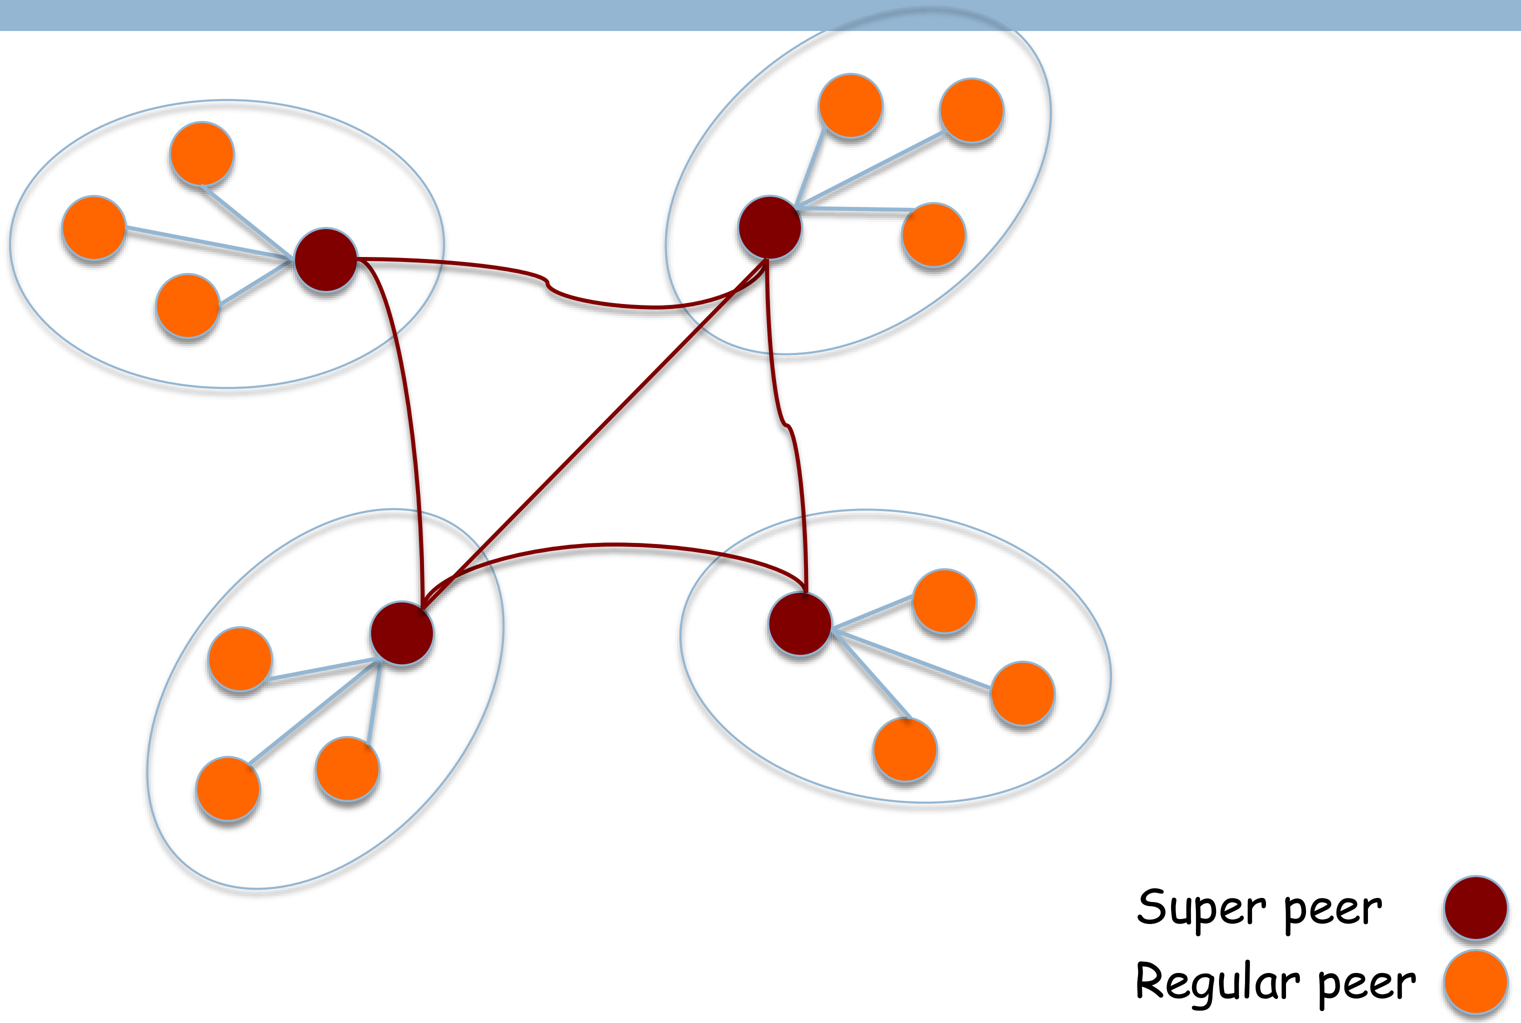

Superpeer networks

- The client-superpeer relationship is **fixed**
  - When a peer joins, it *attaches* itself to the superpeer and stays attached till it leaves
- Superpeers are expected to be *long-lived* processes with *high-availability*
- Selecting nodes that are eligible to be superpeers?
  - Closely related to the leader **election** problem

Some declare their lives are lived
as true profundity,
and others claim they really live
the real reality.

...

In minor ways we differ,
in major we're the same.

I note the obvious differences
between each sort and type,
but we are more alike, my friends,
than we are unalike.

Human Family, Maya Angelou

SMALL WORLDS

Stanley Milgram's experiment on social networks

- In 1967 he mailed 160 letters
- People were randomly chosen from Omaha, Nebraska
- Objective was to pass their letter
  - TARGET: Stock broker in Boston, MA
  - CONSTRAINT: Use *intermediary* known to them on a **first-name** basis

Results: It's a small, small world

- 42 letters made it through
  - Median was just 5.5 intermediaries
  - US Population in 1967: 200 million
    - 2025: 347 million

- First demonstration of what is known as the **small world** effect

Intuitively it seems that the pathlengths should have been much higher

- People's social circle is cliquish or **clustered**
- People you know, know each other

The key is the distribution of links within social networks

- Some acquaintances are relatively isolated
- Some have wide ranging connections
  - Play a critical role in bringing network **closer** together
- Milgram experiment
  - ¼ of the successful chains passed through a local storekeeper

The Hollywood Network:

- Here we organize all actors in a graph
- If they have co-starred with someone in a movie
  - They have a direct link to them (1 hop)
- Some actors have more links than others because they have acted in so many movies
  - E.g., Kevin Bacon

The Hollywood Network: 6 degrees of Kevin Bacon

- John Carradine: 4000 links
- Robert Mitchum: 2905 links
- But acting in the most movies does not always translate into shortest hops to a random node in the network
- Rankings:
  - Rod Steiger: 2.53
  - Donald Pleasence: 2.54
  - Martin Sheen, Christopher Lee, Robert Mitchum, Charlton Heston
  - Kevin Bacon? 2.79 pathlength and ranked 876th

Turns out even a small number of bridges can dramatically reduce pathlengths

Duncan Watts and Steven Strogatz (1998), "*Collective Dynamics of 'Small-World' Networks*," Nature 393, p 440.

Regular Graphs

- Ring of `n` vertices
- Each of the nodes are connected to its **nearest** `k` neighbors

Example regular graph with k = 4

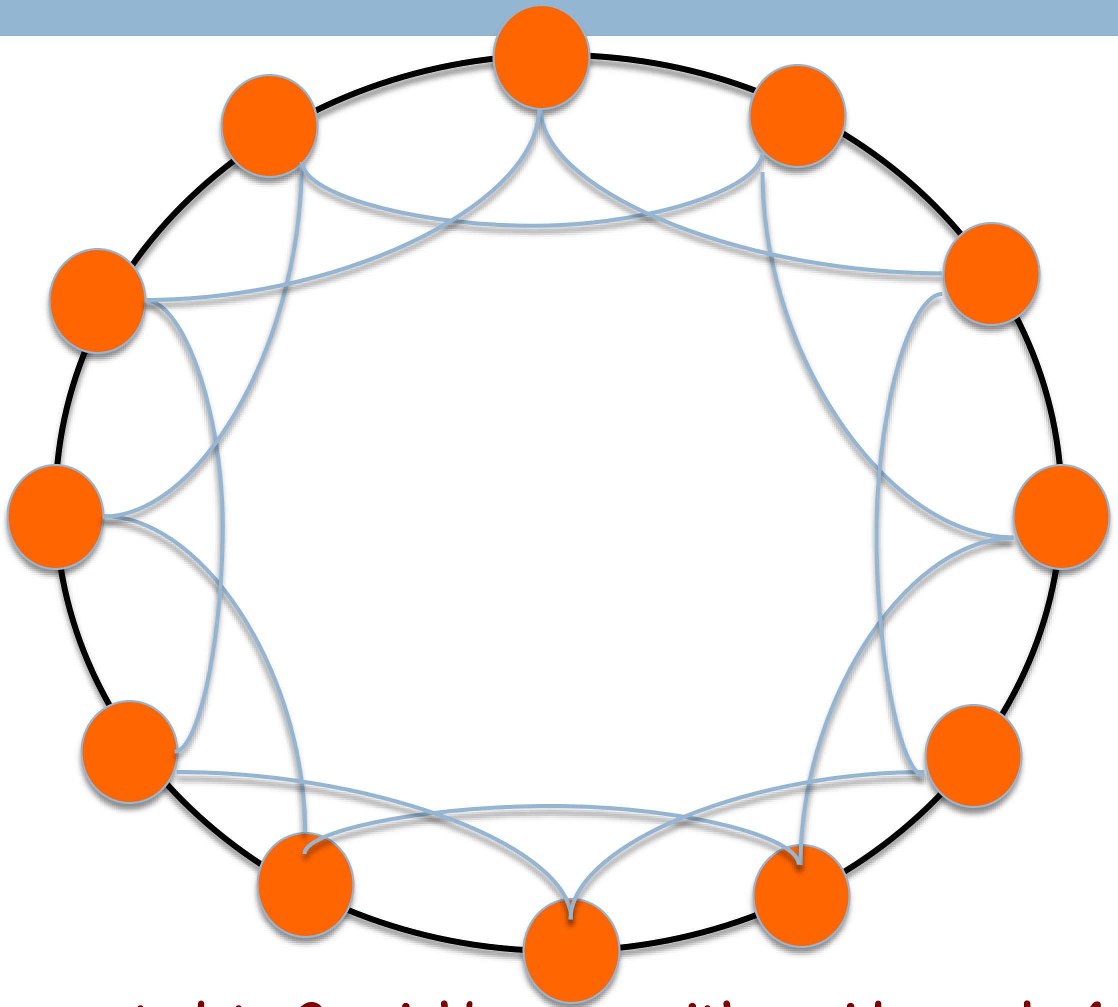

Each node is connected to 2 neighbors on either side; so k=4

Pathlength in a graph

- Average number of **hops** to reach any node in the system
  - For **each pair** of vertices, compute shortest path
  - Take the average over all pairs
- Gives a sense of **how far apart** points are in the network

Clustering coefficients are a measure of the level of clustering

- For **k** neighbors of a vertex, the number of possible connections between them is

$$C_2^k = \frac{k(k-1)}{2}$$

  - **Clustering coefficient** of a vertex
    - Proportion (0 ~ 1) of possible links actually present in graph

Pathlength in Regular graphs

- Approximately **n/2k**
- If **n**=4096 and **k**=8
- Pathlength = **n/2k** = 256
  - Very large!

Clustering Coefficient: Regular graph k=4

$$\text{Clustering coefficient} = \frac{3(k-2)}{4(k-1)}$$

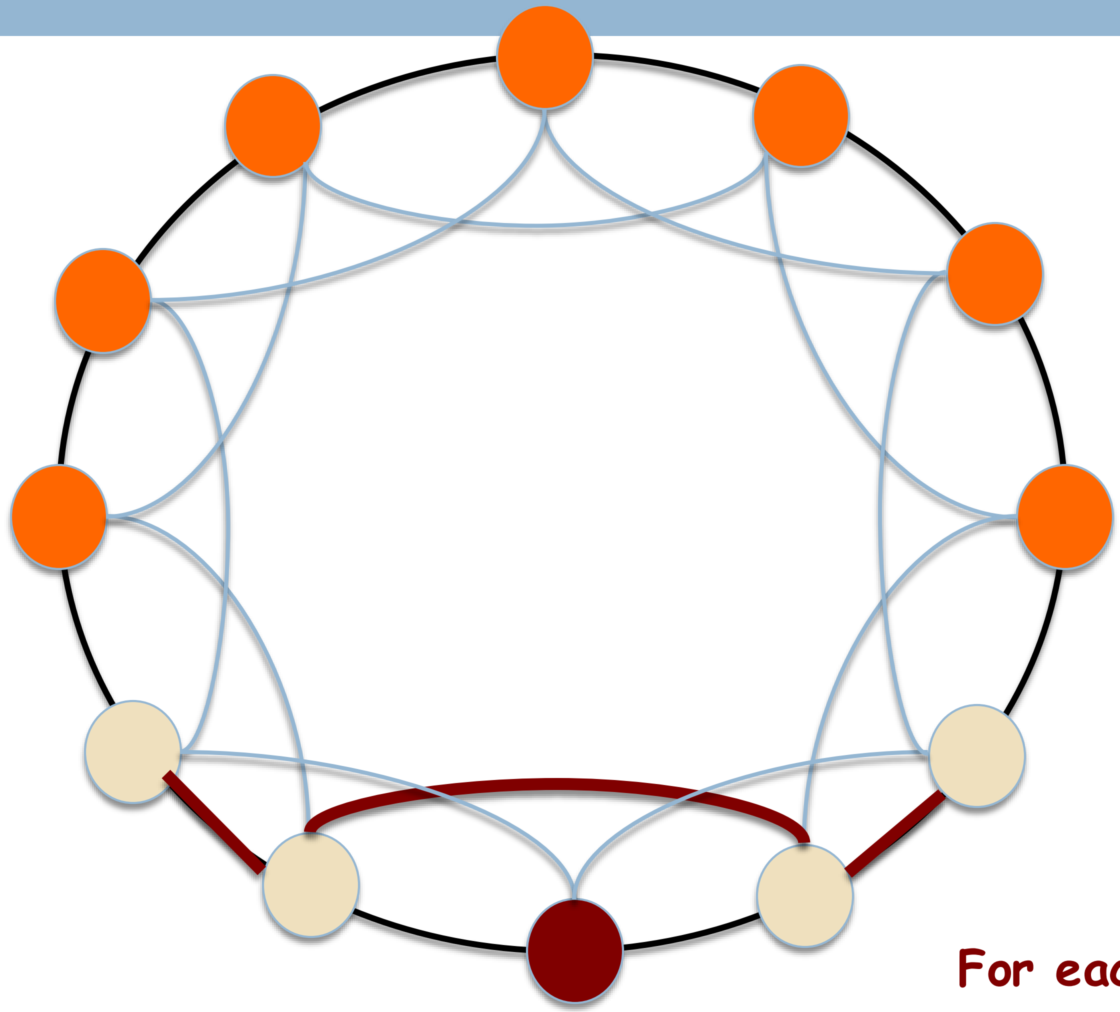

For each vertex = 3/6

Random Graphs

- Opposite of regular graphs
- Vertices are connected to each other at **random**

Random Graphs

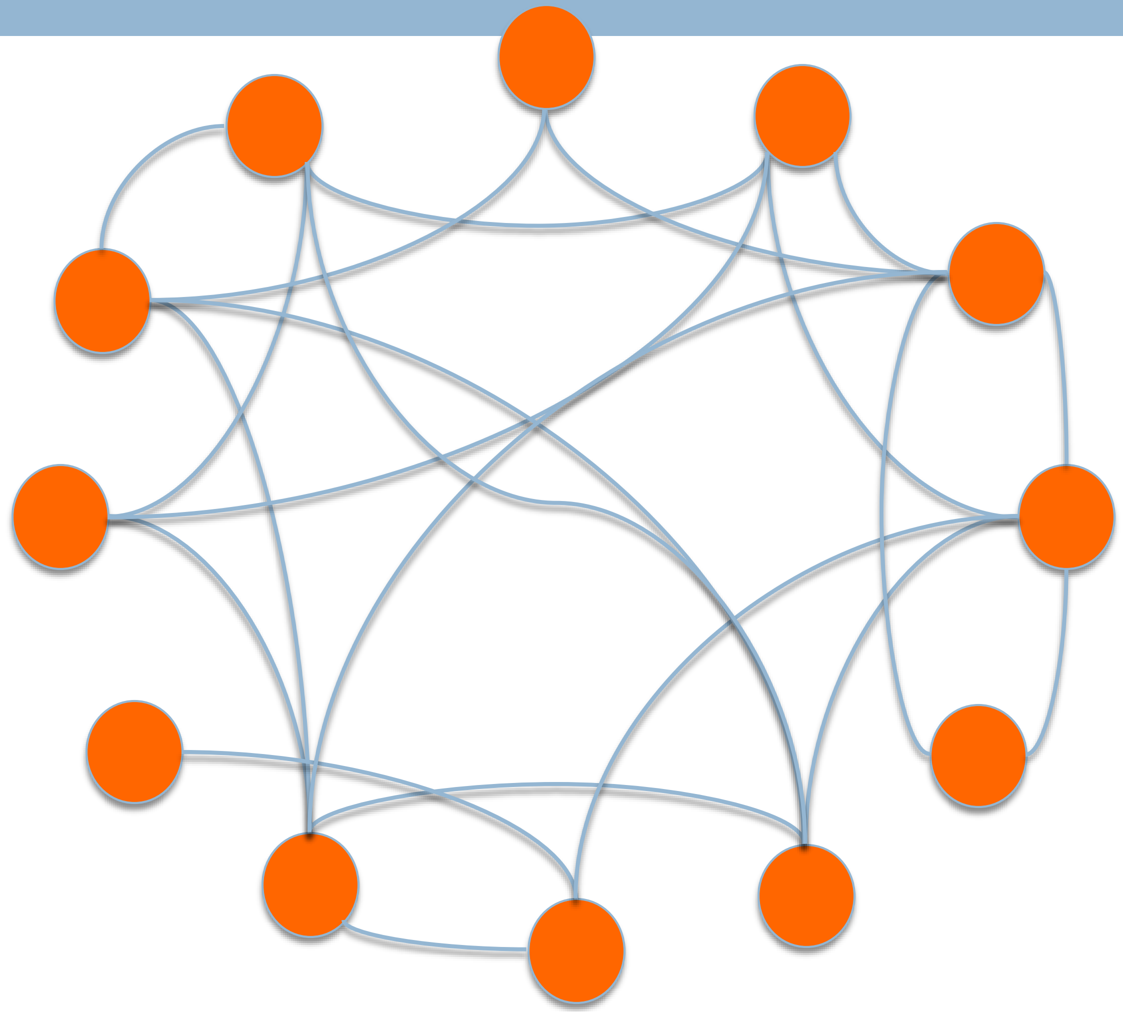

Pathlength and clustering coefficients in Random Graphs

- Pathlength is approximately `log n/log k`
- Clustering coefficient is approximately: `k/n`

- So, with **n**=4096 and **k** = 8
  - Average pathlength = `log 4096/log 8` = 4
  - Much better than regular graphs
- Clustering coefficient = 8/4096 = 0.002
  - Much lower than regular graphs

Comparing regular and random graphs

- Regular graph
  - **High clustering**
  - High pathlength

- Random graph
  - Low clustering
  - **Low pathlength**

Small world graphs: Add a few random links to the regular graph

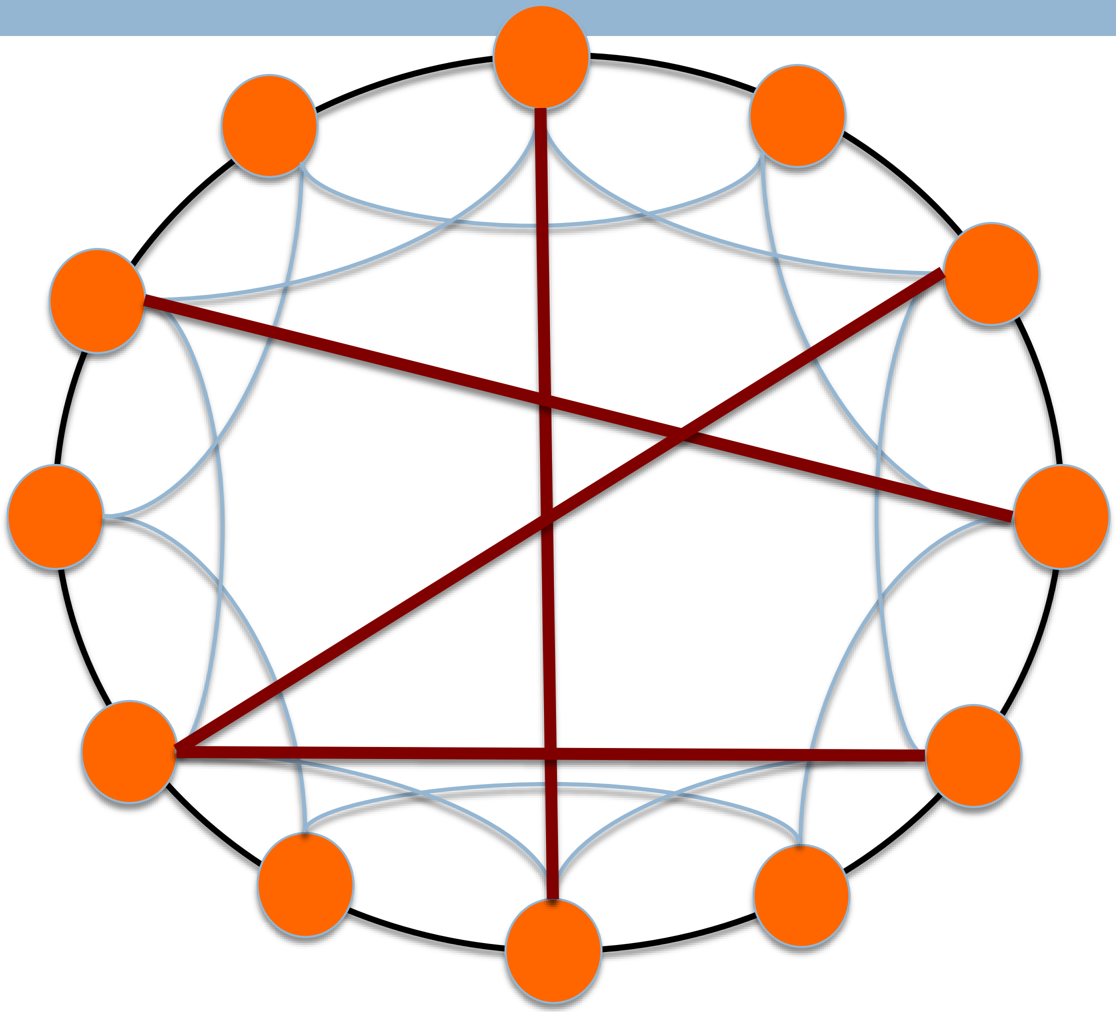

Small world graphs

- High local clustering
- Short global pathlengths

- Implications:
  - Small amount of **rewiring** needed to promote the transition
  - Transition is *barely noticeable* at the local level

As restless as we are
We feel the pull
In the land of a thousand guilts
And poured cement
Lamented and assured
To the lights and towns below
Faster than the speed of sound
Faster than we thought we'd go
Beneath the sound of hope

1979; William Patrick Corgan; The Smashing Pumpkins

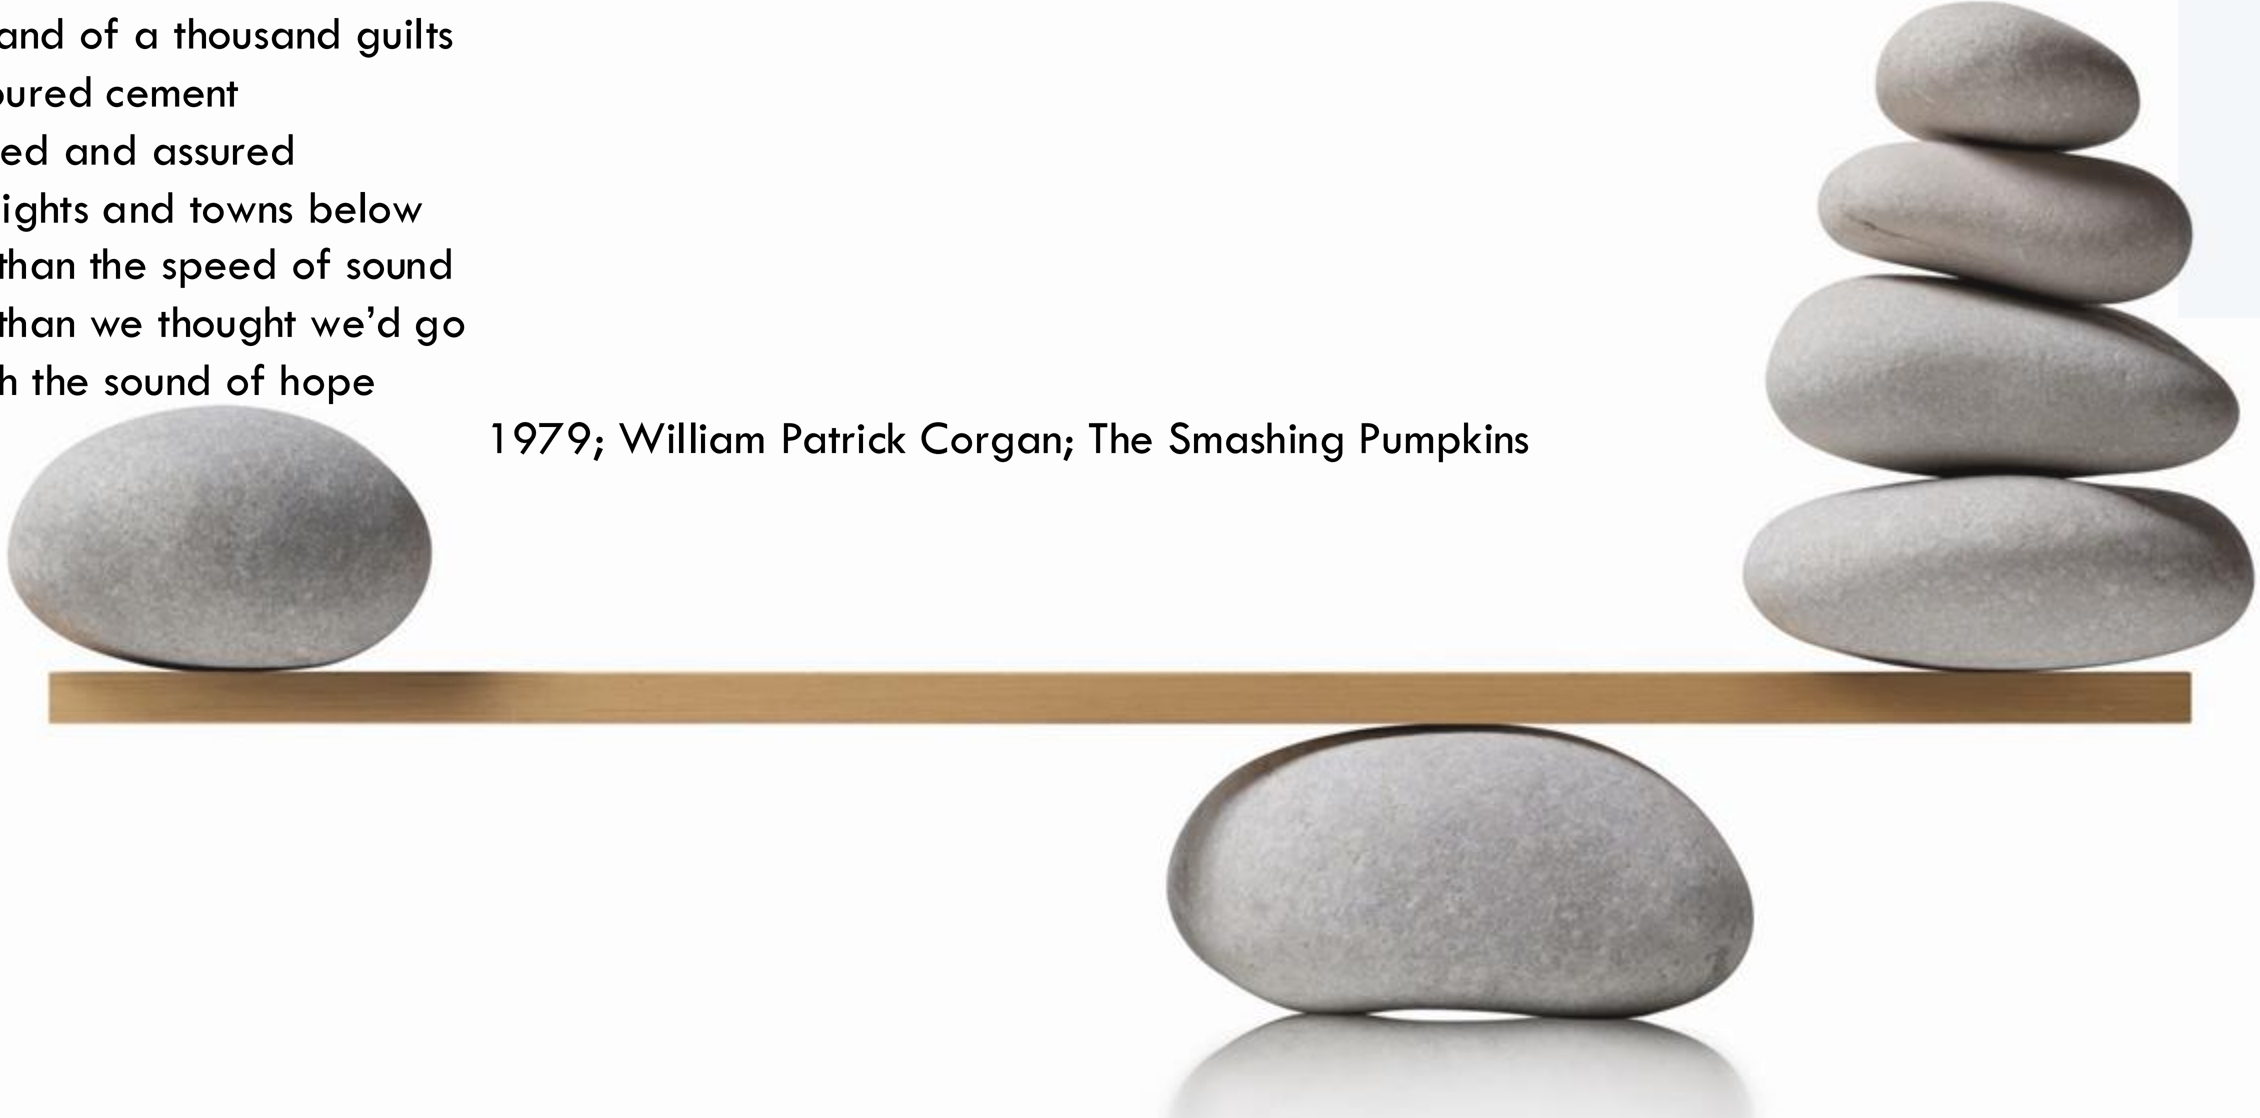

SCALE FREE NETWORKS

Power law is a special relationship between two quantities

- The number or frequency of the object
  - Varies as a **power**
- Of some attribute (size) of the object

- Earthquakes
  - The frequency of earthquakes varies as a power of the size of the earthquake

Power law and Random Networks: Real World examples

- Random networks
  - Eisenhower National Highway System
  - Nodes=Cities, Links=Highways connecting them
  - Most cities served by roughly the same number of highways
- Scale-free networks
  - Airport system
  - Large number of small airports served by a few major **hubs**

Distribution of links in random networks

- Follows a **bell curve**
- Most nodes have the same number of links

Comparison of the distribution of links in random and scale-free networks

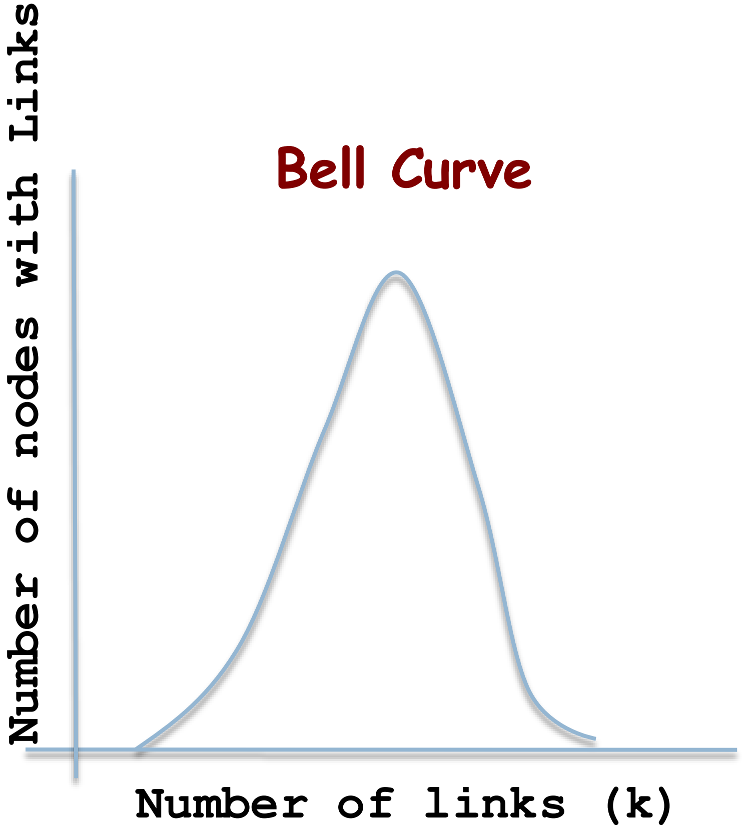

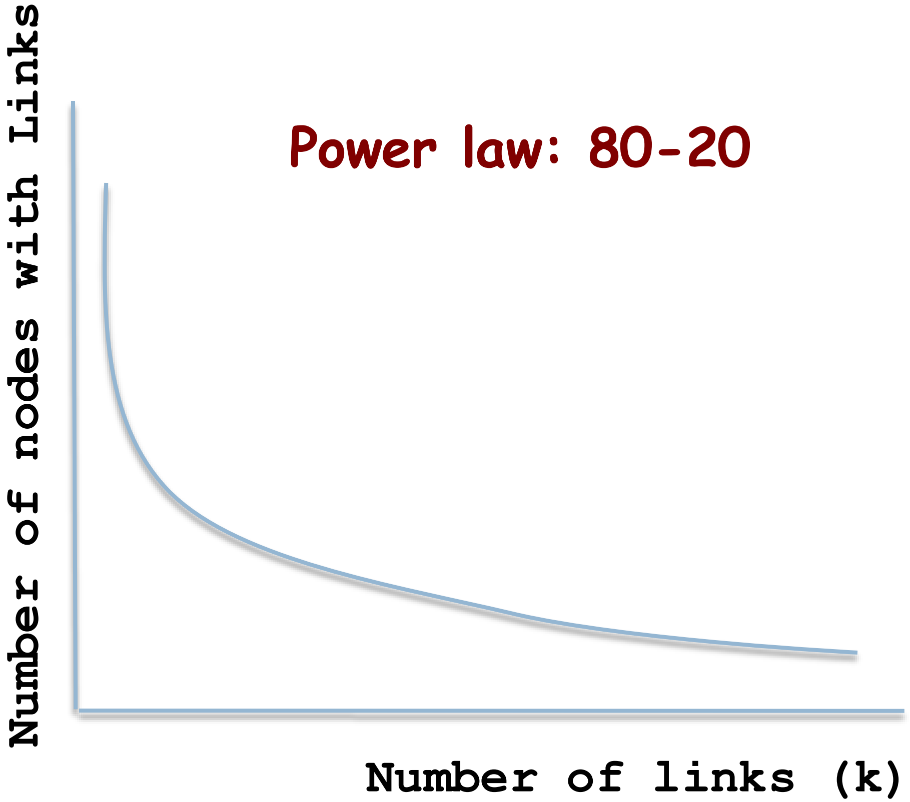

Growth of scale-free networks

- Addition of nodes
- **Preferential attachment**
  - Nodes prefer to attach to well-connected nodes

- **RESULT:** Highly connected nodes emerge

Power law distributions have no peak

- Continuously decreasing curve
- Many small events coexist with a few very large ones

- Imaginary planet:
  - Most people will be really short
  - Among 8 billion people, 1 person would be 8000 ft

Bell Curves vs Power Laws

- Bell Curves
  - Occur very often in nature
  - Exponentially decaying tail
    - Responsible for **absence** of hubs
- Power Laws
  - Emerge during phase transitions
    - Move from chaos to order: Self organization
  - Decay far more slowly
    - Allows for hubs

Why power law networks are called scale-free [1/2]

- In a random network vast majority of nodes have same number of links
  - Nodes deviating from average are rare
  - There is a characteristic **scale** in its connectivity
    - Embodied by the average node
    - Fixed by the peak of the degree distribution

Why power law networks are called scale-free [2/2]

- In a power law network
  - Absence of peak
- No such thing as a characteristic node
  - Continuous hierarchy of nodes spanning from rare hubs to numerous tiny nodes
- No intrinsic scale in power law networks
  - **Scale-free** networks

Achilles' heel in the power law network

- Power law networks are robust to random failures
- Vulnerable to a targeted attack on hubs

- Removal of hubs
  - Disintegrates these networks
  - Breaks them up into tiny non-communicating islands

Coexistence of robustness and vulnerability plays a role in complex systems

- Sea otters in California went nearly extinct because of excessive hunting for its pelts

- In 1911 federal regulators banned hunting them
  - Otters made a dramatic comeback

The case of the otter recovery [1/2]

- Because otters feed on *urchins*, increase in their numbers leads to a decrease in the number of urchins
- With fewer urchins around, the number of *kelps* went up dramatically
- Increased the supply of food for fish
  - Protected the coast from erosion

The case of the otter recovery [2/2]

- Protection of one species (a hub) altered economy and ecology of the coastline
- Finfish now dominate coastal fisheries
  - Once dedicated to shellfish

The contents of this slide set are based on the following references

- *Peer-to-Peer: Harnessing the Power of Disruptive Technologies. Edited by Andy Oram. O'Reilly Publishing. ISBN: 0-596-00110-X.*
  [Chapter 14 – Performance by Theodore Hong]
- *Linked: How Everything is Connected to Everything Else and What it Means for Business, Science, and Everyday Life*. Albert-László Barabási. Plume. ISBN: 0452284392/978-0452284395.
  [Chapters 4,5,6, and 7]