

ICS 362 Distributed Systems

Distributed Systems: Part 5

Lecturer: Toby Daniel

Distributed Systems Topologies

- The peer-to-peer explosion has reminded people of the power of decentralized systems.
- The promise of robustness, open-endedness, and infinite scalability have made many people excited about decentralization.
- But in reality, most systems we build on the Internet are largely centralized.

Decentralised Systems

- The Internet itself is the largest decentralized computer system in the world. But ironically in the 1990s many systems built on the Internet were completely centralized.
 - The growth of the Web meant most systems were single web servers running in fabulously expensive facilities.
- Now with peer-to-peer, the pendulum has swung the other way to radically decentralized architectures such as Gnutella and Kazaa.
- In practice, extreme architectural choices in either direction are seldom the way to build a usable system.

Topology


- The debate between centralized and decentralized systems is fundamentally about topology
- in other words, how the nodes in the system are connected.
- Topology can be considered at many different levels:
 - physical,
 - logical,
 - connection,
 - organizational.

Network Diagrams

- For this analysis, topology is considered in terms of the information flow.
- Nodes in the graph are individual computers or programs.
- links between nodes indicate that those nodes are sharing information regularly in the system.
- Typically, an edge implies that the two nodes are directly sharing bits across a network link.


Centralised

- Centralised systems are the most familiar form of topology, typically seen as the client/server pattern used by databases, web servers, and other simple distributed systems. All function and information is centralized into one server with many clients connecting directly to the server to send and receive information. Many applications called "peer-to-peer" also have a centralized component. SETI@Home is a fully centralized architecture with the job dispatcher as the server. And the original Napster's search architecture was centralized, although the file sharing was not.




Ring

- A single centralized server cannot handle high client load, so a common solution is to use a cluster of machines arranged in a ring to act as a distributed server. Communication between the nodes coordinates state-sharing, producing a group of nodes that provide identical function but have failover and load-balancing capabilities. Unlike the other topologies here, ring systems are generally built assuming the machines are all nearby on the network and owned by a single organization.




Hierarchical

- Hierarchical systems have a long history on the Internet, but in practice are often overlooked as a distinct distributed systems topology. The best-known hierarchical system on the Internet is the Domain Name Service, where authority flows from the root name-servers to the server for the registered name and often down to third-level servers. NTP, the Network Time Protocol, creates another hierarchical system.



Decentralised

- Decentralised systems, where all peers communicate symmetrically and have equal roles. Gnutella is probably the most "pure" decentralized system used in practice today, with only a small centralized function to bootstrap a new host. Many other file-sharing systems are also designed to be decentralized, such as Freenet or OceanStore. Decentralized systems are not new; the Internet routing architecture itself is largely decentralized, with the Border Gateway Protocol used to coordinate the peering links between various autonomous systems.




Hybrid Topologies

- Distributed systems often have a more complex organization than any one simple topology. Real-world systems often combine several topologies into one system, making a hybrid topology. Nodes typically play multiple roles in such a system. For example, a node might have a centralized interaction with one part of the system, while being part of a hierarchy in another part.


Centralised + Ring

- Serious web server applications often have a ring of servers for load balancing and failover. The server system itself is a ring, but the system as a whole (including the clients) is a hybrid: a centralized system where the server is itself a ring. The result is the simplicity of a centralized system (from the client's point of view) with the robustness of a ring.



Centralised + Centralised

- The server in a centralised system is itself often a client of one or more other servers. Stacking multiple centralized systems is the core of n-tier application frameworks. For example, when a web browser contacts a server, the software on that server may just be formatting results into HTML for presentation, and itself calling to servers hosting business logic or data. Web services intermediaries such as Grand Central Networks also create several layers of centralized system. Centralized systems are often stacked as a way to compose function.



Centralised + Decentralised

- A new wave of peer-to-peer systems is advancing an architecture of centralized systems embedded in decentralized systems. This hybrid topology is realized with hundreds of thousands of peers in the FastTrack file-sharing system used in KaZaA and Morpheus. Most peers have a centralized relationship to a "supernode," forwarding all file queries to this server (much like a Napster client sends queries to the Napster server). But instead of supernodes being standalone servers, they band themselves together in a Gnutella-like decentralized network, propagating queries. Internet email also shows this kind of hybrid topology. Mail clients have a centralized relationship with a specific mail server, but mail servers themselves share email in a decentralized fashion.

