The Future of Massively Distributed Services on the Internet

Mark Hapner
Sun Microsystems

Massively Distributed

- Work done at edge with minimal central coordination
- Work done with distributed infrastructure
- Work done by large communities
- The Web itself ...
  - Content
  - Function

Examples

Computing	Community
Virus	eBay
Botnet	Wikipedia contributors
Seti@home	Blog-osphere
BitTorrent	Amazon reviewers
eDonkey	Newspapers
Skype	Bulletin Boards
Network News	Podcast-osphere
DNS	
eMail	

Non-Examples

IM	
Google	
eBay	

Commercial Models

- Legal
  - Selling words
  - Selling eyes
  - Selling communities
  - Selling computing resources?
  - Selling information?
- Illegal
  - File sharing
  - Stealing identities
  - Stealing money
  - Extortion

The Long Tail

- The long tail of demand
- Democratization of production
- Democratization of delivery

Around 80 per cent of all traffic in the Internet is already P2P. This traffic will increase 1,000-fold in the next five years and most of it will be encrypted P2P, according to a study by Staselog and researchers at Finnish Universities.

Distributed Hash Tables

Distributed hash tables (DHTs) are a class of decentralized distributed systems which partition ownership of a set of keys among participating nodes, and can efficiently route messages to the unique owner of any given key. Each node is analagous to a bucket in a hash table. DHTs are typically designed to scale to large numbers of nodes and to handle continual node arrivals and failures. This infrastructure can be used to build more complex services, such as distributed file systems, peer-to-peer file sharing systems, cooperative web caching, multicast, anycast, and domain name services.

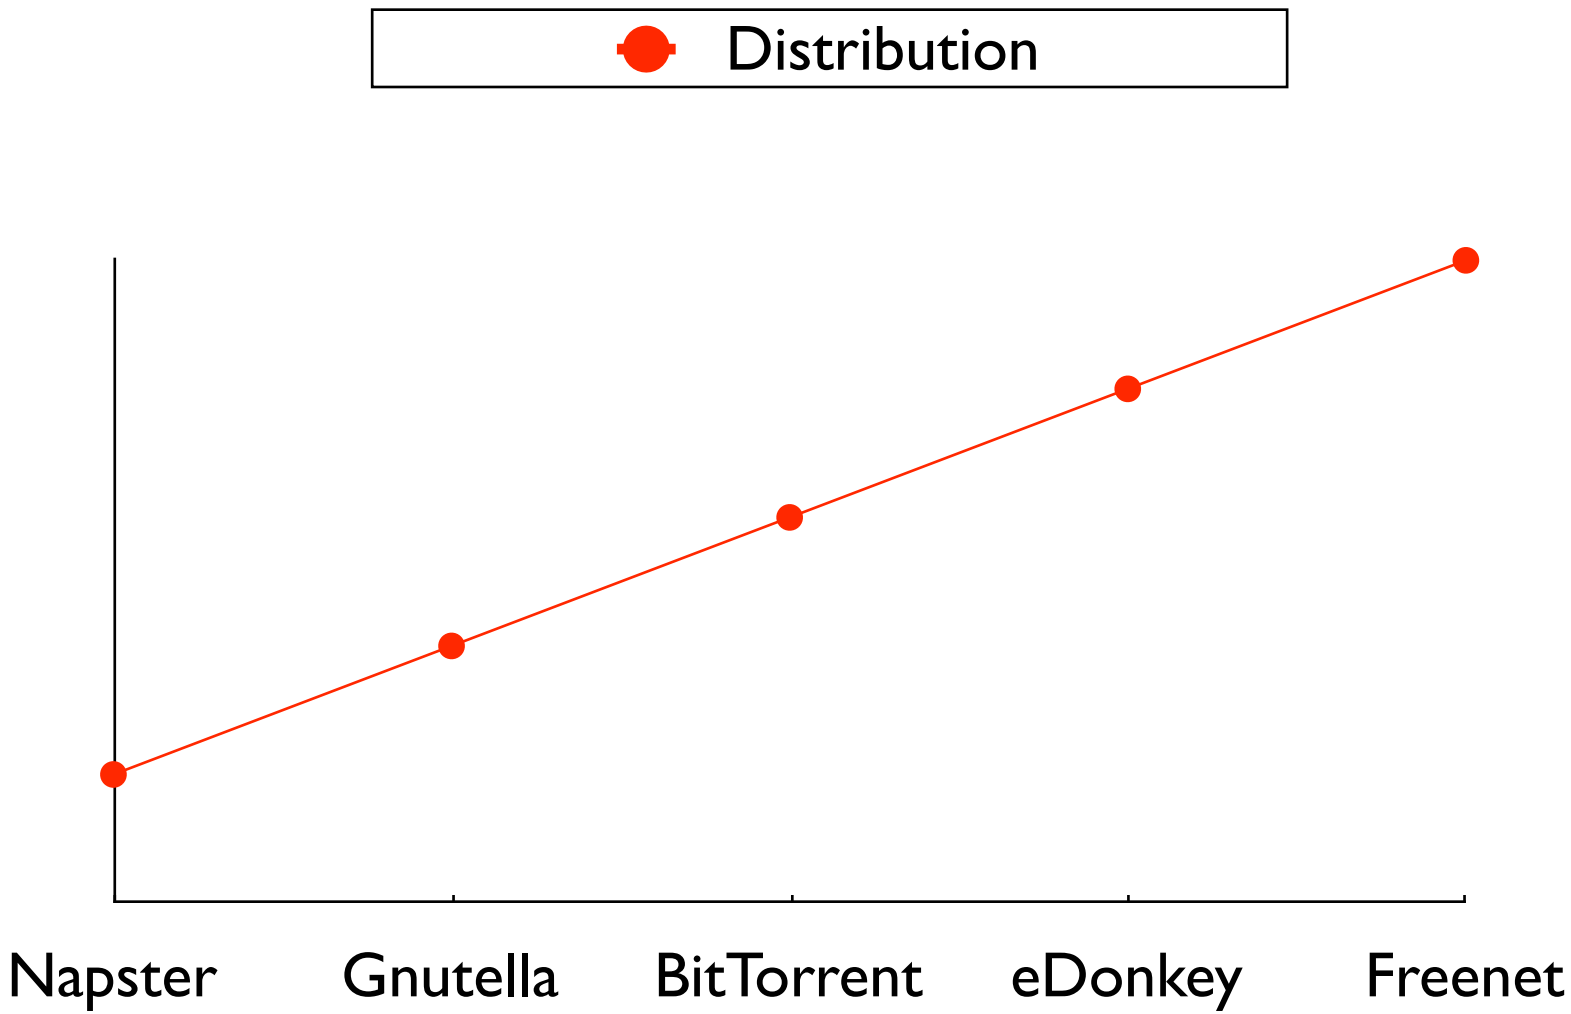
Distribution
Napster
Gnutella
BitTorrent
eDonkey
Freenet

BitTorrent

With BitTorrent, files are broken into smaller fragments, typically a quarter of a megabyte each. As the fragments are distributed to the peers in a random order, they can be reassembled on a requesting machine. Each peer takes advantage of the best connections to the missing pieces while providing an upload connection to the pieces it already has.

A Torrent

To share a file using BitTorrent, a user creates a .torrent file, a small "pointer" file that contains:

- the filename, size, and the hash of each block in the file (which allows users to make sure they are downloading the real thing)
- the address of a "tracker" server
- and some other data (like client instructions).

eDonkey

- A decentralized peer-to-peer computer network that implements the Kademlia algorithm which is a form of DHT
- Content key: ed2k
- Number of concurrent Users: ~800,000
- Estimated to be 50% of web traffic prior to MPAA Cease and Desist

ed2k Link

A file is divided into 9.28Mb chunks and the hash is calculated for each one. The resulting hash table is hashed once again, and the final value is used as a part of the ed2k link. The typical ed2k link also includes the filename and the filesize.

ed2k://|file|The_Two_Towers-The_Purist_Edit-Trailer.avi|14997504965c013e991ee246d63d45ea71954c4d|/

Magnetic Link

A magnet link is like an ISBN book number in that it specifies only a specific content (or "resource"). No two different files will have the same magnet link, however the same file found at different locations will have the same magnet link. This is the same as with ISBN numbers: No 2 different books have the same numbers, but a specific book will have the same number regardless of where you put it.

magnet:?xt=urn:sha1:YNCKHTQCWBTRNJIV4WNAE52SJUQCZO5C

Freenet

Freenet is a decentralized censorship-resistant peer-to-peer distributed data store. Freenet works by pooling the contributed bandwidth and storage space of member computers to allow users to anonymously publish or retrieve various kinds of information. Freenet uses a kind of key based routing similar to a distributed hash table to locate peers' data.

The system has no central servers, is peer-to-peer, and is not subject to the control of any one individual or organization. Even the designers of Freenet do not have any control over the overall system.

Freenet

- Content key: SHA-1 160 bit digest
- Fully distributed search via dynamically constructed routing tables
  - Results are indeterminate
  - Nodes are buckets for close content
- Plausible deniability
  - Content encrypted with DSA or KSK
  - Nodes don't know what they store
  - File transits all look the same
  - v7 to create scalable Darknet

Today

- P2P
  - Lacks a business model
  - Motivation comes from community
  - Mass demand is entertainment
- Portals
  - Library vs yellow pages
  - Index value vs knowledge value
- Seeded vs crawled
- New models of organization emerging

The Future ...

- New ‘networks’ distribute
  - Content
  - Search
- Communities
  - Self organize
  - Provide their own resources
- Individuals
  - Host distributed communities
- Portals
  - Dissappear

Links

- http://www.thelongtail.com/
- http://www.edonkey2000.com/
- http://magnet-uri.sourceforge.net/
- http://freenetproject.org/