

Designing an Internet for Content Delivery and Not Communication

Krishna Gummadi

Max-Planck Institute for Software Systems
Saarbruecken, Germany

Outline

- A challenge for the future Internet
- Why it requires rethinking the Internet's design
- What infrastructure is needed to evaluate new designs

Challenge:

Efficient and cheap bulk content transfers

- Is the Internet the preferred medium for all digital content?
 - not for bulk content beyond a certain size
 - e.g., movies, home videos, software downloads, data backups, and scientific data
- Postal networks are more efficient & cheaper for bulk data
 - netflix mails DVDs, even Google uses sneaker-nets
- However, transfers over the Internet are more convenient
 - more demand for bulk content than Internet can handle

Why rethink Internet design for bulk transfers?

- Motivating case study: My relocation from U.S. to Germany
- Moving bulky house items via postal network was a breeze
 - different mailing rates for different delivery dates and reliability
- Moving 1.2 TB of bulk data via Internet was a nightmare
 - MPI's happy hours (12 - 6 AM) were UW's unhappy hours
 - TCP b.w. limited by path RTT, tweaking led to firewall alarms
 - after a week, used physical disks to move data

The real problem

- The Internet is designed for interactive communication, not bulk content delivery
 - well-suited for short bursty traffic, not for long duration flows
 - postal network is well designed for content delivery
- Communication traffic vs. content delivery traffic
 - online vs. offline, instantaneous delivery vs. delayed delivery
- Bulk content can be delayed at intermediate storage locations
 - the network can optimize routes and schedules for efficiency & cost

How would we design an Internet for content delivery and not communication?

- We would design it as a cargo network, and not as a telecommunication network
 - storage and accumulation centers at intermediate nodes in the network
 - coordinate route selection and transfer schedules across different content transfers
 - maximize network utilization, while meeting delivery deadlines

Designing such a network poses several research challenges

- Intermediate storage centers:
 - how large? how many? who manages resource contention? what are good policies? is store an infinite queue?
- Reliability of data transfers:
 - recovery end-to-end or between intermediate stores? how does the end host track status of transfer?
- Routing and scheduling:
 - who controls it? to what extent? how? do end hosts have a role?
- Other optimizations: multicast and caching
 - leverage content caches for other transfers? if so, how to name, address, and route content?

Requirements from experimental infrastructure

- Equip every router with large and high-performance storage systems
 - store at line speeds(?)
- Given trends in storage capacities and costs, the additional costs may not be substantial(?)
- Would benefit many other on-going research projects
 - delay-tolerant-networking, logistical networking, data-oriented transfers, and traffic analysis / monitoring

Conclusion

- Communication and content delivery are very different beasts
 - very different requirements and priorities
- The Internet today is a great communication medium but a poor content delivery medium
- Lets design an Internet for content delivery
 - it can co-exist with the interactive Internet
 - I can't wait to have Wikipedia, Google, YouTube on my desktop

Thanks