

CS244 Lecture 3

Architecture and Principles

**End-to-end Arguments In System Design
(1981)**

Chang Kim



Logistics

- I need to know you better
 - Please join the Slack org and send me and the Tas your name and picture
 - When you ask questions or make comments during the lecture in the classroom, please tell me your name (just a few times starting today)

Previous Paper's Conclusion

- “Datagram” good for most important goals, but poor for the rest of the goals.
- Processing packets in isolation, resource management, accountability all hard.
- Anticipates flows and “soft-state” for the future.

End-to-End Arguments in System Design

[Saltzer, Reed, Clark 1981]

End-to-end in a nutshell

“The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the end points of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)”

Case studies

- Error handling in file transfer
- Encryption and authentication
- The partition between TCP, IP, and the link layer of error handling, flow control and congestion control.
- Fairness in resource allocation
- Addressing

Packet error handling across layers

- Ethernet FCS
 - CRC32
- IP header checksum
 - One's complement sum over IP header
- TCP checksum
 - One's complement sum over TCP header, payload, and IP pseudo header
- TLS (SSL)
 - Cryptographic MAC for integrity
- Why all these seemingly redundant checks?

Some consequences

- In layered design, the E2E principle provides guidance on where functions belong.
- “Dumb, minimal” network and “intelligent” end-points.

Many argue that:

E2E principle allowed the Internet to grow rapidly because innovation took place at the edge, in applications and services.

Ex. WWW, Skype, BitTorrent, Bitcoin

On the other hand...

E2E principle appears to have become diluted: NATs, firewalls, VPN tunnel endpoints, ...

- Perhaps not surprising: E2E principle grew in an era of trust among users. Now network must protect *itself*.

The network is no longer “dumb, minimal”

- Now over 7,000 RFCs.
- Router OS's based on 100M lines of source code.

Q: Is this a problem?

What you said

Hannes:

“What might be an example of a layered system that is or is not implemented with the end-to-end argument in mind?”

Jared:

“What are examples of top systems today that violate the end-to-end principle? Would you say that Software Defined Networking, Caching and CDNs, etc violate this principle to some extent?”

Agata:

“Does the end-to-end principle mean programmers need to know less or more about the underlying network and possible errors?”

What you said

Purvi:

“Is there a way to design a system in which the application or the user can adjust how “strongly” the end-to-end argument is applied?”

Maggie:

“While I appreciate the usefulness and benefits of end-to-end argument, I think performing the function at the lower level has another benefit: it can reduce the costs of implementing the function by individual application. Although it could also be an extra cost if the application does not need the function.”

Pratyusha:

“Is there a unifying principle(s), guidelines, or a litmus test(s) for when a design decision may be becoming out of the scope of the end-to-end argument?”

What you said

Sarah:

“What is a communication use where accuracy is more important than speed? What is a use case (other than audio) where correct information is not more important than speed?”

Michelle:

“It seems that this end-to-end argument was necessary because, contrary to historical communication systems like the postal service or telegram or phone, computer networks were meant to serve a much broader variety of purposes, so the tendency to perform low-level optimization was actually potentially counterproductive.”

What belongs in, what out?

Questions:

- Does routing belong in the “dumb, minimal” network?
- How about multicast, mobility, QoS...?
- Are NATs necessary, good, or evil?
- Is the E2E principle constraining innovation of the infrastructure?

Additional references

[rfc3724] “The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture” - [Kempf et al.](#)

[Blumenthal] “Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world”, [ACM Transactions on Internet Technology](#), Vol. 1, No. 1, August 2001, pp 70-109.