

Internet Technology History — The Case of My PhD Thesis on End-to-End Arguments

Matthias Bärwolff

March 5, 2011

1 Introduction

A history of the Internet is always a history of at least two things – the conspicuous surface of things, and the hidden workings beneath. While sometimes it is already hard enough to make sense of *what* goes on at the surface, the more challenging part is, of course, the latter: looking at *why* those things go on in the first place, what allows them to happen, and how those workings compare to alternative (possibly hypothetical) arrangements in the overall system. In my PhD thesis (Bärwolff 2010) I have looked at one of the design principles widely considered to be constitutional to the Internet: the ‘end-to-end principle’, which argues that, to the maximum extent possible, functions of computer networking ought to sit with the application end points, not random intermediary nodes. In this essay I shall briefly discuss the approach taken in the thesis, and some of the more interesting conclusions to be drawn from that approach.

2 Computer Science and History

Some consider it adventurous to even call computer science a science, for unlike most other sciences it typically chooses to build (with little ‘natural’ bounds) what it may then describe, measure, and reason about. Yet there is an increasing feeling among many in the field that science is an entirely appropriate term to name the host of practices and principles found there. If we accept this position as a reasonable premise, then it is as reasonable to look at the history of such practices and principles, and thus advance the science of computing, or computer science.

My thesis has done just that for one particular – some would say particularly important – principle of computer networking, namely the end-to-end principle and the practices to be associated with it. I have ploughed through the history of how the balance of functions between network intermediaries and end points has evolved over time, how it was modeled, and how

the dynamics between theorizing and engineering played out through the decades from the early 1960s to today. In the process I have not only come to appreciate the difficulty of articulating the principle, but have also added fairly substantial elaboration to the proper articulation of the principle, its foundation in historical and empirical fact, and its proper scope and limitations. By taking a history approach to developing the end-to-end principle I have uncovered elaborations that would be hard to arrive at by mere blackboard reasonings or limited laboratory experiments.

The value of my exercise is hard to judge fully at this point. At least, I have added to the record of properly researched historical fact with regard to the Internet and its immediate predecessor Arpanet. However, despite extensive triangulation and the conscious waiving of prior hypotheses that could have biased my research priorities, mine remains a subjective take on various, sometimes highly contentious matters. It is absolutely possible that I have overlooked certain historical threads others might deem vital to elaborating the end-to-end principle. Plus, of course, I have brought a subjective disposition to the table shaped by my prior research and opinions on certain points. Yet, even though one may question plausibility of some of my elaborations about the end-to-end principle, I believe that the historical part of my thesis is particularly valuable in its own right.

3 Sources, Opinions, and Interpretations

The motivation for my thesis was twofold: for one, it was the lack of proper articulation and elaboration of the end-to-end principle; and, for another, it was the conspicuous lack of thorough history of the notion that could serve to inform any such elaboration. The former is not to say that previously existing elaborations of the end-to-end arguments were false or futile; however, over time they increasingly added to the confusion of what exactly the initially very simple end-to-end principle means once you add complications about real world developments that have taken place since. Our goal was thus, to put it with von Hayek (1973, p. 60) “to search for a correct statement of what before was known implicitly.” No doubt this is an ambitious feat, hence my reservations above.

The more surprising find, very early in my endeavor, was the poor state of historical research about the history of the Internet that could aid such novel interpretations of the original end-to-end principle. To be sure, the primary sources are all out there, but sensible compilations beyond more or less random collections of ‘interesting things that happened in the course of the development of the Internet’ are hard to get by. Some of the most widely cited secondary sources on the Internet’s evolution are very lax in their proper referencing and discussing primary sources and interviews; and some of the more scholarly approaches are tilted by certain inclinations

about normative questions that come to mix with technical ones in the history of the Internet.

So how did I go about compiling a secondary account that could serve as a sound footing for elaborating the end-to-end principle? My first strategic decision was a pragmatic one about structure and the necessary sequentiality of my text: in the first part of my thesis I traced the history of the end-to-end arguments as a subject of academic discourse in the open literature about the Internet and its design principles. This turned out to be a fairly mechanical exercise, albeit one that turned up a number of original insights about the meaning of the end-to-end principle. Also, it is the most complete literature review of its kind, going back to well before the well-known 1984 articulation by Saltzer, Reed, and Clark.

In the second part of my thesis I elaborated the history of the Internet as a real-world artifact and with particular regard to the notion of the end-to-end principle – beginning with the Arpanet, featuring various digressions about other ‘local area’ networks in the 1960s and 1970s, and discussing in some detail various issues such as packet fragmentation and congestion management in the Internet. This part of the thesis turned out to be the far more challenging one, for it involved working through a world of primary sources of which to make proper sense. In what may be described as a happy recursion, the Internet afforded me with a wealth of sources and communication options to complete this enterprise. As for written sources, I have amassed some 2000 papers, taking up more than 3 GB of hard disk space – all largely through various free online repositories plus resources accessible via university libraries (TU Berlin, HU Berlin, MIT). Also, I have done a dozen or so mostly informal personal interviews, and have exchanged hundreds of correspondences with other knowledgeable people via mailing-lists (particularly the Internet History List at www.postel.org) and personal email. The result of my work is a very well-researched account of a number of key episodes in the broader technological history of the Internet. While there have been other works to comparable effect (Hafner and Lyon 1998; Abbate 1999) there had (to the best of my knowledge) been no account yet which synthesizes from a large enough set of *explicitly referenced* primary sources. To a certain extent this is of course due to my luxury of being able to add some 250 pages of endnotes to a relatively slim main body of text without an editor stopping me. The completeness of the record thus obtained, however, is probably unparalleled.

A more general methodological note: the huge number of sources accessible with relative ease and convenience is more than just ‘a larger set of sources’, it amounts to an explosion in our ability to triangulate facts. It also allows a much improved ‘iteration facility’ for fine-tuning one’s premises, thus allowing one to quickly elevate through increasingly informed stages of understanding. While none of those options necessarily suffices to obtain consistent, valid, or instructive historical findings, the task of doing ‘tech-

no-history’ has become a lot easier – as I hope I have demonstrated in my thesis.

4 Getting the Word Out

The most amazing thing about writing these days may not be the incredible wealth of prior works one may access as a ‘shoulder’ to stand on. It may rather be the ease with which one can produce a book in arbitrarily mint typesetting quality (using tools like L^AT_EX), plus the availability of print-on-demand publishers that drop the upfront cost for any would-be private publisher next to zero (see for example www.createspace.com). Of course, neither availability nor affordability make for automatic awareness and acceptability; and even the availability of a free PDF (with fancy internal hyperlinking and search facility) may not help much when the average attention span of a growing number of people is quickly dropping to Twitter message length.

Anyway, to me the option of getting my thesis out as a nicely printed book for an extremely competitive price (cheaper than self-printing, let alone binding) was a most welcome one, especially since all the typesetting had already been done for submission of the thesis, anyway. Getting from the initial PDF to a print ready book version took some additional effort (scaling the size, removing some of the margin jazz that only makes sense in the electronic version), but I think it was well worth it. Self-publishing works such as mine may not revolutionize the traditional publishing industry business with its often useful separation of various roles in the value chain, but it sure is a nice complement to these practices.

References

- Abbate, Janet (1999). *Inventing the Internet*. MIT Press.
- Bärwolff, Matthias (2010). “End-to-End Arguments in the Internet: Principles, Practices, and Theory”. PhD thesis. Technische Universität Berlin. URL: <http://baerwolff.de/publications/2010-10-PhD-thesis.html>. Also self-published at CreateSpace.com.
- Hafner, Katie and Matthew Lyon (1998). *Where Wizards Stay up Late: The Origins of the Internet*. New York, NY: Touchstone.
- Saltzer, Jerome Howard, David Peter Reed, and David Dana Clark (1984). “End-to-End Arguments in System Design”. In: *ACM Transactions in Computer Systems* 2.4, pp. 277–288. URL: <http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf>. Edited version of a conference paper published in 1981.

Von Hayek, Friedrich August (1973). "Rules and Order". In: *Law, Legislation and Liberty: A New Statement of the Liberal Principles of Justice and Political Economy*. Vol. 1. 3 vols. University of Chicago Press.