



DOI:10.1145/3110531

Vinton G. Cerf

# In Praise of Under-Specification?

Eric Schmidt, executive chairman of Alphabet (Google's parent company), recently drew my attention to the notion of "under-specification." He reminded me that the Internet had benefited

strongly from this concept. Several specific examples came to mind. The Internet Protocol (IP) specification does not contain any information about routing. It specifies what packets look like as they emerge from or arrive at the hosts at the edge of the Internet, but routing is entirely outside of that specification<sup>a</sup> partly because it was not entirely clear what procedures would be used for Internet routing at the time the specification was developed and, indeed, a number of them have been developed over time. There is nothing in the specification that describes the underlying transmission technology nor is there anything in the specification that speaks to how the packet's payload (a string of bits) is to be interpreted. These matters are open to instantiation independent of the specification of packet formats.

Some of the under-specification can be a manifestation of *layering* that figured strongly in the ARPANET host-host protocols and was carried over in the Internet Protocol suite. The idea is that while there is a well-defined interface between the layers that specifies how information crosses the layer boundary, the details of the layer above or below are hidden. This feature allows for changes in the implementation of and even the characteristics of the upper or lower layer. For example, above the IP layer, one finds a number of different

protocols such as User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) or Real-Time Protocol (RTP) that all send and receive Internet packets but they use and interpret the IP packet payloads in different ways. Below the IP layer one finds a variety of different transmission technologies including Ethernet, Multi-protocol Label Switching, Frame Relay, Asynchronous Transfer Mode, Dense Wave-Length Division Multiplexing, and many others. The IP layer doesn't really care how the packets are transported.

What is interesting to contemplate is whether the notion of under-specification that induces flexibility and anticipates new but unknown developments can be codified in a concrete way beyond the purely conceptual. Is there a way to measure the degree of specification in the way that Claude Shannon found to specify information as entropy independent of semantics? Can something be fully specified, partly specified, or completely unspecified and how would these be described or measured more precisely? In circuit design, for instance, there is the notion of "don't care" for some values in a Boolean representation. Can this notion be applied to program specification as well as to protocol specifications? Are there design principles that one can derive from this notion of under-specification?

I am reminded of an anecdote told about doing business with Chinese manufacturers. American companies produced very detailed specifications of what was to be fabricated down to the last detail and the Chinese compa-

nies produced exactly what was asked, at a price. But a Chinese company produced a less specific specification, leaving room for the manufacturer to innovate, leading to a design that was less expensive, easier to manufacture, and to maintain.

One of my oldest friends, Jonathan Postel, was the Internet Assigned Numbers Authority for many years and was often quoted: "Be liberal in what you accept and conservative in what you send," in reference to the implementation of protocols. His dictum was aimed at improving interoperability. Of course, people who are particularly concerned about security might take issue with this particular nostrum (and some have!).

As may be apparent to readers who have gotten this far, I am not yet sure there is a *there there*, but I am fascinated by the possibility that it might be possible to extract some design principles from this notion that would lead to potentially more robust and adaptable designs. Think about what makes a chair a *chair*. It's a thing to sit on, has legs and usually a back and maybe some arms. But there are *so many* things we recognize as chairs that are quite varied in their specifics. Flexible design suggests to me that under-specification has something to do with *essence* or core concepts. I hope interested readers will take a moment to share their thoughts, particularly if they see more deeply into this idea than I have at the present. 

Vinton G. Cerf is vice president and Chief Internet Evangelist at Google. He served as ACM president from 2012–2014.

Copyright held by owner/author.

<sup>a</sup> This is not precisely correct since the notion of "source routing" is part of the specification and allows a host to force packets to flow along a path specified by intermediate IP addresses, but the general route generation and selection process is independent of the IP specification itself.