

Statement: Transmit/Receive and Request/Response Models for Communication

Patrick Crowley Washington University in St. Louis St. Louis, MO, USA pcrowley@wustl.edu

ABSTRACT

Communication models are important because in most systems, and certainly in large or distributed ones, it is often difficult to distinguish between cause and effect, and to know if solutions match with problems. We proceed with the premise that challenges in today's networked systems may be due to their underlying communication models.

CCS CONCEPTS

Networks → Network architectures → Network design principles

KEYWORDS

Communication models, Internet protocols, computer networking

ACM Reference format:

Patrick Crowley. 2023. Statement: Transmit/Receive and Request/Response Models for Communication. In *Proceedings of 10th ACM conference on Information-Centric Networking (ACM ICN '23), October 9-10, 2023, Reykjavik, Iceland. ACM, New York, NY, USA, 3 pages.* https://doi.org/10.1145/3623565.3623760

1 Transmit/receive & IP/QUIC

By communications model, we mean the basic abstractions used to reason about communication; in other words, the mental model and vocabulary that describe the process of communication in simple terms. Claude Shannon articulated [SH48] the landmark model for communications in 1948, which we refer to as the transmit/receive model, consisting of a source, transmitter, channel, receiver, and destination. Information is conveyed through a channel in the form of a signal. Shannon's use of this model enabled him to develop a mathematical characterization of its components and their interactions, which in turn led to his discovery of bits (binary digits as information theory, and digital information representation, and these inventions in turn birthed the digital systems and digital communications we enjoy today. Shannon's model for transmitting



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License. *ACM ICN '23*, October 9–10, 2023, Reykjavik, Iceland © 2023 Copyright is held by the owner/author(s). ACM ISBN 979-8-4007-0403-1/23/10.

ACM ISBN 979-8-4007-0405-1725/10. https://doi.org/10.1145/3623565.3623760 discrete symbols across a noisy channel was not invented tabula rasa, rather, it was a simplified representation of the telegraph, telephone, and radio communication systems of the day. These real-world systems came first, and Shannon's model followed.

The transmit/receive model can be logically extended and made digital, to include digital representations for information, a bidirectional channel, source and destination addresses, and packets as logical signals. With these logical extensions, the transmit/receive model can be seen as the communications model that forms the basis of the Internet Protocol (IP). IP operates over a switched packet network, in which messages are sent between devices with IP addresses. While Shannon's ideas did not include an analogue to packet switching, which has been instrumental to the Internet's success, IP is otherwise a straightforward extension. The impact of this model for communication has been profound and is difficult to overstate. In the minds of many, this model is synonymous with the idea of communication itself, rather than just one model amongst others.

The most notable recent networking transmit/receive protocol to be developed and deployed is QUIC [LA17], which was originated by Google and subsequently refined and adopted by the IETF as the protocol to run above IP as part of HTTP/3 [BI20]. Notably, the QUIC handshake includes the TLS 1.3 [RE18] handshake, and hence offers private sessions. QUIC addresses problems with long lived HTTP connections, head-of-line blocking, the handshake penalty for security, and client mobility via connection IDs. It is based on the transmit/receive model and bears the characteristics of channel-oriented protocols.

Research efforts in data-, content- and information-centric networking [KO07, JA09, ZH14] have pointed out that channeloriented protocols have become limiting factors in today's internet, causing problems of security & privacy (you can only "armor the pipe" not the content), scalability & efficiency (often, network paths are carrying multiple copies of the same data in parallel), and application simplicity (logic must be added and maintain to apps to map from the data-oriented app logic to the connection-oriented communication model). While used universally, the transmit/receive model has been rivaled in usage and popularity by the request/response model.

2 Request/response & HTTP

Beginning with the invention of the world wide web [BL94] by Tim Berners-Lee in 1989, a global model for linking data (HTML) was connected to a global network (IP) via an application-layer transport protocol (HTTP.) Rather than linking between data objects in one machine, the way HyperCard [GO88] and prior hypertext systems did, the web and HTTP allowed for linking data objects between machines across an IP network. HTTP is a request/response protocol in which a client sends a request (e.g., to get a URL) to a server, and the server replies with a response to a client (e.g., the web page corresponding to the URL). As much of human activity has come to be mediated by the Internet, this request/response model of communication and HTTP have come to dominate global communications.

3 Strengths and Weaknesses

Both models are asymmetric with respect to the communication participants: transmitters transmit, receivers receive, clients request, and servers respond. In order to permit symmetric, bidirectional communication, participants must operate a second instance in the opposite direction. Here, we see an important difference between the models and their leading exemplars. IP deployments are nearly always bidirectional with both sides capable of communicating in the same way; HTTP deployments rarely do so. (IP exceptions include networks sometimes used in sensitive environments which permit one-way communication, generally outbound, through the use of an IP-layer data diode).

Relative to the transmit/receive model, the request/response model introduces asynchrony into communication. The request and receive operations are ordered but can happen at some delay from one another. This is a virtue when links are unreliable, especially with chunked data (i.e., when data objects are fragmented into smaller, individually deliverable pieces.) This also enables request pipelining, so that a group of requests can be sent concurrently.

HTTP is a client-server protocol and the request/response model, while general-purpose, has client-server affordances. This creates challenges when using the request/response model for non-clientserver applications. Consider a simple, point-to-point chat application. When one app has a message to send to the other, it must do so by making a request of some kind. One way to do this is for the sending application to create a message with name X, then make a request to the recipient of the form: "Ask me for message X", so that the other application can then turn around and make a request of the form: "Send me message X." It is possible to build this type of application, but it is not an application that is particularly natural to express with this communication model; this can lead to unnatural, hard-to-program application logic. The request/response model is more naturally suited to client-server application structures.

We can see this dynamic in protocols based on the request/response model, such as Named Data Networking (NDN) [ZH14], a research effort to develop a general-purpose network layer -- one with a narrow waist like IP -- based on the request/response model. To accommodate situations like the simple messaging application, HTTP permits data objects to be passed along with requests, something akin to parameters to a function call. With NDN, it is often awkward to create namespaces and application logic that work well with serverless or non-client-server

application structures, including the simple chat application discussed above. NDN's rigorous interest-data packet exchange reduces the effectiveness of including rich data in a request. NDN also exhibits complexity in namespaces with multiple, distributed producers of the same data [AB18]; our hypothesis is that NDN's multi-producer challenges arise because the request/response model is fundamentally a client-server pattern, and application architectures with no server or multiple servers run afoul of the communications model. Indeed, the author's work in using NDN in recent years to build and deploy mobile handheld systems and applications for real-world ad hoc data sharing [SH17] are a primary motivation of this perspective.

4 Discussion

Is there community consensus around the strengths and weaknesses of these models? Are researchers exploring others?

ACKNOWLEDGMENTS

This work was supported by NSF CNS Award 2234360.

REFERENCES

- [SH48] Shannon, Claude Elwood. "A mathematical theory of communication." The Bell system technical journal 27.3 (1948): 379-423.
- [LA17] Langley, Adam, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasic, Dan Zhang, Fan Yang et al. "The quic transport protocol: Design and internet-scale deployment." In Proceedings of the conference of the ACM special interest group on data communication, pp. 183-196. 2017.
- [BI20] Bishop, Mike. "Hypertext transfer protocol version 3 (HTTP/3)." Internet Engineering Task Force, Internet-Draft draft-ietf-quic-http-27 (2020).
- [RE18] Rescorla, Eric, and Tim Dierks. "The transport layer security (TLS) protocol version 1.3." (2018).
- [KO07] Koponen, Teemu, Mohit Chawla, Byung-Gon Chun, Andrey Ermolinskiy, Kye Hyun Kim, Scott Shenker, and Ion Stoica. "A data-oriented (and beyond) network architecture." In Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 181-192. 2007.
- [JA09] Jacobson, Van, Diana K. Smetters, James D. Thornton, Michael F. Plass, Nicholas H. Briggs, and Rebecca L. Braynard. "Networking named content." In Proceedings of the 5th international conference on Emerging networking experiments and technologies, pp. 1-12. 2009.
- [ZH14] Zhang, Lixia, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, K. C. Claffy, Patrick Crowley, Christos Papadopoulos, Lan Wang, and Beichuan

ACM ICN '23, October 9-10, 2023, Reykjavík, Iceland

Statement: Transmit/Receive and Request/Response Models for Communication

Zhang. "Named data networking." ACM SIGCOMM Computer Communication Review 44, no. 3 (2014): 66-73.

- [BL94] Berners-Lee, Tim, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. "The world-wide web." Communications of the ACM 37, no. 8 (1994): 76-82.
- [GO88] Goodman, Danny. The complete HyperCard handbook. Bantam Books, Inc., 1988.
- [AB18] Abraham, Hila Ben, Jyoti Parwatikar, John DeHart, Adam Drescher, and Patrick Crowley. "Decoupling information and connectivity via information-centric transport." In Proceedings of the 5th ACM Conference on Information-Centric Networking, pp. 54-66. 2018.
- [SH17] Secure Handhelds on Assured Resilient networks at the tactical Edge (SHARE). DARPA. https://www.darpa.mil/program/secure-handheldson-assured-resilient-networks-at-the-tactical-edge, 2017