the network (TAPs) as well as end-points. Also, in order to provide rich end-to-end semantics, our transfer service is implemented *below* an end-to-end session whereas DOT works on top of existing transport layer protocols.

Recently, Popa et al [33] propose the use of HTTP as the narrow waist of the Internet, so it can be viewed as a transfer service. While the use of HTTP is certainly easier to use in the short-term, it is difficult to exploit many content centric and multi-path optimizations due to the inherent limitations of HTTP (i.e., naming, rigid semantics, etc), Moreover, we also define the roles of layers below and above the transfer service i.e., segment and session layers, which play a key role in the Tapa architecture.

Future Internet Architectures: Tapa's pull mode is inspired by data oriented architectures (e.g., CCN [27]), but there are important differences as we discuss below. First, unlike CCN, Tapa is not a "pure" content-centric architecture, so host and destination identifers are present in the packets even if we are retrieving an ADU based on its identifer. This results in different per-hop router operations as well as different failure modes in Tapa compared to CCN. Second, CCN operates at a per-packet granularity while Tapa uses ADUs. This difference mandates the need for different mechanisms to support traditional transport functions like congestion control, reliability, and data reassembly.

A better positioning of Tapa compared to CCN and other new network architectures is to view it as a transport architecture that can leverage these proposals as part of its segment protocol. For example, if Tapa is implemented over content centric network architectures [27, 26, 29], or service centric architectures [32, 26] then Tapa's transfer service can leverage the inherent features provided by these architectures for content/service routing. This will simplify Tapa's transfer layer and can also potentially improve performance due to late binding, native support for content/service discovery, and intrinsic security [26].

7. FINAL THOUGHTS

We presented the design, implementation, and evaluation of Tapa, a transport architecture that accommodates network heterogeneity and rich in-network services. Tapa unbundles today's transport and makes explicit use of in-network services that operate on ADUs. Our practical experience, as well as the case studies in this paper, confirms that Tapa offers great flexibility at multiple levels: customized solutions as segment protocols; diverse data oriented optimizations at the transfer level; and services with new semantics at the session level.

8. ACKNOWLEDGMENTS

We thank the anonymous reviewers, XIA team members, Ratul Mahajan, Ihsan Qazi, and Dina Papagiannaki for providing useful feedback on this work. This research was funded in part by NSF under award number CNS-1040801.

9. REFERENCES

[1] CMU Wireless Emulator. www.cs.cmu.edu/ emulator/.
[2] Dynamic site acceleration. http://www.akamai.com/html/solutions/dsa_curriculum.html.
[3] Service Location Protocol. RFC 2608.
[4] Vanlan. research.microsoft.com/en-us/projects/vanlan/.
[5] A. Akella, et al. Self-management in chaotic wireless deployments. In *MobiCom '05*, pp. 185–199. 2005.
[6] M. Alizadeh, et al. Data center tcp (dctcp). In *Proceedings of the ACM SIGCOMM 2010*, pp. 63–74. 2010. ISBN 978-1-4503-0201-2.
[7] A. Anand, et al. Redundancy in network traffic: findings and implications. In *Proc. of SIGMETRICS*, pp. 37–48. 2009.
[8] D. G. Andersen, et al. Accountable Internet Protocol (AIP). In *SIGCOMM*. 2008.
[9] A. V. Bakre, B. Badrinath. Implementation and performance evaluation of indirect tcp. *IEEE Transactions on Computers*, 46(3):260–278, 1997.
[10] S. Biswas, R. Morris. Exor: opportunistic multi-hop routing for wireless networks. *SIGCOMM CCR*, 35(4):133–144, 2005.
[11] J. Border, et al. Performance enhancing proxies intended to mitigate link-related degradations, 2001.
[12] D. Clark. The design philosophy of the darpa internet protocols. In *SIGCOMM '88*, pp. 106–114. 1988. ISBN 0-89791-279-9.
[13] D. D. Clark, D. L. Tennenhouse. Architectural considerations for a new generation of protocols. In *SIGCOMM*. 1990.
[14] D. D. Clark, et al. Addressing reality: an architectural response to real-world demands on the evolving internet. *SIGCOMM Comput. Commun. Rev.*, 33(4):247–257, 2003.
[15] D. D. Clark, et al. Making the world (of communications) a different place. *SIGCOMM CCR.*, 35(3):91–96, 2005.
[16] J. Crowcroft, et al. Plutarch: an argument for network pluralism. *SIGCOMM Comput. Commun. Rev.*, 33:258–266, August 2003. ISSN 0146-4833.
[17] F. R. Dogar. Architecting for diversity at the edge: Supporting rich network services over an unbundled transport. PhD thesis. 2012.
[18] F. R. Dogar, P. Steenkiste. M2: Using Visible Middleboxes to Serve Pro-active Mobile-Hosts. In *ACM SIGCOMM MobiArch '08*, pp. 85–90. 2008.
[19] F. R. Dogar, P. Steenkiste. Segment based internetworking to accommodate diversity at the edge. *Technical Report - CMU-CS-10-104*, 2010.
[20] F. R. Dogar, P. Steenkiste, K. Papagiannaki. Catnap: Exploiting high bandwidth wireless interfaces to save energy for mobile devices. In *ACM MobiSys*, pp. 107–122. 2010.
[21] F. R. Dogar, et al. Ditto: a system for opportunistic caching in multi-hop wireless networks. In *ACM MobiCom*. 2008.
[22] J. Eriksson, H. Balakrishnan, S. Madden. Cabernet: vehicular content delivery using wifi. In *MobiCom '08*, pp. 199–210. 2008.
[23] K. Fall. A delay-tolerant network architecture for challenged internets. In *SIGCOMM '03*, pp. 27–34. 2003.
[24] B. Ford, J. Iyengar. Breaking up the transport logjam. In *ACM Hotnets*. 2008.
[25] S. Guha, P. Francis. An end-middle-end approach to connection establishment. In *SIGCOMM*. 2007.
[26] D. Han, et al. XIA: Efficient support for evolvable internetworking. In *Proc. 9th USENIX NSDI*. Apr. 2012.
[27] V. Jacobson, et al. Networking named content. In *CoNEXT '09*, pp. 1–12. 2009.
[28] S. Kandula, et al. FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In *NSDI*. April 2008.
[29] T. Koponen, et al. A data-oriented (and beyond) network architecture. In *SIGCOMM '07*, pp. 181–192. 2007.
[30] M. Li, et al. Block-switched networks: a new paradigm for wireless transport. In *NSDI'09*, pp. 423–436. 2009.
[31] X. Liu, X. Yang, Y. Xia. Netfence: preventing internet denial of service from inside out. In *ACM SIGCOMM 2010*.
[32] E. Nordstrom, et al. Serval: An end-host stack for service-centric networking. In *Proc. 9th USENIX NSDI*. April 2012.
[33] L. Popa, A. Ghodsi, I. Stoica. HTTP as the narrow waist of the future Internet. In *Hotnets 2010*.
[34] S. Roy, et al. Application level hand-off support for mobile media transcoding sessions. In *NOSSDAV '02*, pp. 95–104. 2002.
[35] J. H. Saltzer, D. P. Reed, D. D. Clark. End-to-end arguments in system design. *ACM Trans. Comput. Syst.*, 1984.
[36] I. Stoica, et al. Internet indirection infrastructure. *SIGCOMM Comput. Commun. Rev.*, 32(4):73–86, 2002.
[37] N. Tolia, et al. An architecture for internet data transfer. In *NSDI '06*.
[38] M. Walfish, et al. Middleboxes no longer considered harmful. *OSDI*, pp. 215–230, 2004.