

ex uno pluria:
**The Service-Infrastructure Cycle, Ossification,
and the Fragmentation of the Internet**

Mostafa H. Ammar
School of Computer Science
Georgia Institute of Technology
Atlanta, GA
ammar@cc.gatech.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

In this article I will first argue that a *Service-Infrastructure Cycle* is fundamental to networking evolution. Networks are built to accommodate certain services at an expected scale. New applications and/or a significant increase in scale require a rethinking of network mechanisms which results in new deployments. Four decades-worth of iterations of this process have yielded the Internet as we know it today, a common and shared global networking infrastructure that delivers almost all services. I will further argue, using brief historical case studies, that success of network mechanism deployments often hinges on whether or not mechanism evolution follows the iterations of this Cycle. Many have observed that this network, the Internet, has become ossified and unable to change in response to new demands. In other words, after decades of operation, the Service-Infrastructure Cycle has become stuck. However, novel service requirements and scale increases continue to exert significant pressure on this ossified infrastructure. The result, I will conjecture, will be a fragmentation, the beginnings of which are evident today, that will ultimately fundamentally change the character of the network infrastructure. By ushering in a ManyNets world, this fragmentation will lubricate the Service-Infrastructure Cycle so that it can continue to govern the evolution of networking. I conclude this article with a brief discussion of the possible implications of this emerging ManyNets world on networking research.

CCS Concepts

•Networks → Network design principles; Network services; Routing protocols;

Keywords

network evolution, network ossification, Internet fragmentation, bypass networks, routing protocol evolution

1. INTRODUCTION

Necessity is the mother of invention. This has indeed been the case in the development of the global data networking infrastructure. In networking parlance one can restate the proverb as “Service requirements motivate infrastructure deployment.” With the “necessity” that motivates networking

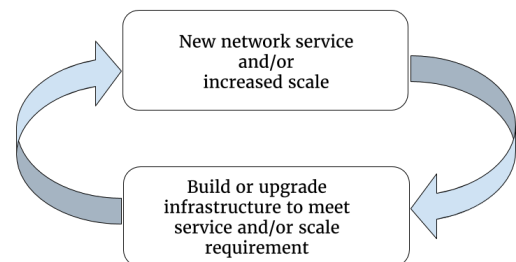


Figure 1: The Service-Infrastructure Cycle

continuously evolving, the “invention” has never been a one-time innovation. Rather, a *Service-Infrastructure Cycle* (or simply *the Cycle*) as shown in Figure 1 has been the framework that governed networking evolution.

In this framework, progress in *deployed* (as opposed to experimental or proposed) networking infrastructure is a continuous cycling. At any given point in time the networking infrastructure is designed and works well to support the popular uses and applications of the time. Network mechanisms and resources are also adapted to the current expected scale: the number of connected users, their network workload and their performance expectations.

At any point in time, the network infrastructure, typically, has some flexibility to support new applications and its resources are provisioned with an anticipation of growth. As time goes on, however, new applications and/or a significant increase in scale of network usage stresses the network infrastructure in ways that require a rethinking of current mechanisms or a significant step increase in network resources. That, in addition to technological advances, results in new deployments that satisfy the changed service requirements.

Consider, for example, that in the early days data networking was built to enable remote access to mainframe computers [36]. The system's design never anticipated email, which quickly emerged as the dominant use of the network.

The evolution of the network continued to respond to new

applications as file transfer, the web, content distribution, electronic commerce and video streaming challenged the design and deployment of the networking infrastructure.

Another force driving changes in networks is scale. As the uses of the network multiplied the number of attached end points increased and the networking technology had to evolve to accommodate the scale. Examples of this over the last four decades, to name just a few, include the development of the Domain Name Systems (DNS) to provide large scale address resolution service, the hierarchical inter-domain structure of the current Internet, the adoption of classless IPv4 addressing, and the development and increasing deployment of IPv6.

The Service-Infrastructure Cycle worked well to give us the Internet until it got stuck in the early 2000s. In 2005 Anderson et al. [9] complained about “the inability [of the Internet architecture] to adapt to new pressures and requirements.” In related work around the same time Turner et al. [52] argued that the Internet infrastructure suffers from *ossification*. This was defined as a situation where “alterations to the Internet architecture that address its fundamental deficiencies or enable new services have been restricted to incremental changes.” It was further stated that this ossification “stifles innovation and the adoption of disruptive technology.” Despite new service demands or increased scale, the Cycle has become stuck and the network infrastructure is no longer amenable to modification in meaningful ways because it has become ossified. Ossification and the reasons for it have received significant attention in the literature [52, 9]. To summarize, the scale of the Internet and the multiple authorities that govern its operation, make it difficult to agree on required changes and to deploy new features. It should be noted that, despite the technical feasibility of certain modifications to the Internet infrastructure, ossification results from the inability of network operators to make a business case for changes to their networks.

In this article I consider the question of whether and how the data networking infrastructure can continue its evolution to respond to new demands despite the ossification barrier. To this end, I will first argue that the Service-Infrastructure Cycle is fundamental to networking evolution. Historically, success of network mechanism deployments often hinges on whether or not mechanism evolution follows the Cycle. In fact, anticipating the Cycle, i.e., trying to do too much too soon, can contribute to deployment failure. In support of this argument, I present in Section 2 two case studies on the evolution of unicast and multicast routing. The aim is to showcase deployment success when the mechanism evolution was allowed to follow the Service-Infrastructure Cycle, and failure when developments do not follow the iterative nature of the Cycle.

Next, in Section 3, I will start with the observation that four decades-worth of iterations of the Service-Infrastructure Cycle have yielded a *common and shared* global network that has become ossified. I will argue, however, that novel service requirements and scale increases continue to exert significant pressure on this ossified infrastructure, seriously hindering the operation of the Cycle. The result, I will conjecture, will be a fragmentation, the beginnings of which are evident today, that will ultimately fundamentally change the character of the network infrastructure. This fragmentation will “lubricate” the Service-Infrastructure Cycle and help it get unstuck so that it can continue to govern the evolution of

networking.

A few caveats: First, it is not my intention to (re)write networking or Internet history. There has been many excellent efforts in that regard (see for example [36] and [28]). Rather, my goal is to add structure to this history in an effort to understand it better. Second, my attempt to add structure and recognize patterns will leave large uncovered gaps in time and technology. I aim to get this type of discussion started and look forward as others contribute to improve on this humble start.

2. A TALE OF TWO ROUTINGS

To understand the workings of the Service-Infrastructure Cycle let us briefly trace the contrasting histories of unicast and multicast routing. I will argue that successful unicast routing deployment progressed until its Cycle was hindered by ossification. Multicast routing, on the other hand, did not see wide-scale deployment because, in large part, it *anticipated* the Cycle by setting early goals that targeted future applications and future network scale.

2.1 Unicast routing follows the Cycle

Unicast routing protocols are used to build forwarding tables that guide network packets to one uniquely identified destination. Unicast routing is a fundamental function of networks. As such, it was the focus of research from the very early days of networking.

Table 1 shows some (by no means all) important milestones in the evolution of deployed unicast routing protocols. The work by Prosser [45] from 1962 represents work that laid the conceptual foundation for network routing. The table also shows how, over the years, unicast routing has gone through many iterations of the Cycle. Each time a particular routing protocol is deployed experience is gained. This experience as well as the need to address increased network scale, cause the next iteration and a new version is deployed.

Consider, for example, the deployment of the Exterior Gateway Protocol (EGP) which enabled unicast routing in a novel larger-scale Internet multi-domain structure. As the Cycle iterations continued, the Border Gateway Protocol (BGP) was developed a few years later and was, according to RFC 1105 [38], “built on experience gained with EGP as defined in RFC 904 and EGP usage in the NSFNET Backbone as described in RFC 1092 and RFC 1093”; clear evidence that the Cycle was in operation.

The table also shows a vibrant landscape for unicast routing from the early 1960s. Things slowed down considerably by the early 2000s as ossification set in. BGP went from version 1 to version 4 in a span of 6 years, but nothing after that. Efforts to address security vulnerabilities in BGP have had mixed success because of the “[lack of] a single centralized authority that can mandate the deployment” [27].

To summarize, the operation of the Cycle is evident in the history of innovation and deployment of unicast routing. The effect of ossification is also evident as further innovation in deployment stalls.

2.2 Multicast routing anticipates the Cycle

Now let’s contrast the evolution of unicast routing with that of multicast routing. Multicasting is the act of sending data from a source to multiple destinations using a *single* transmit operation. Because of this, multicast can be highly efficient when applications call for it. Multicasting can occur

Year	Milestone	Comments
1962	Prosser [45]	Conceptual Framework for Packet Routing
1969	Distance Vector Routing [30]	First operational deployment of distributed routing
1979	Link State Routing [40]	Better Scalability than Distance Vector Routing
1982	Exterior Gateway Protocol [47]	Enables simple multi-domain hierarchy
1989-1995	Border Gateway Protocol v1-v4 [38, 46]	Multi-domain policy Routing
Early 2000s	Secure BGP [34, 39]	Wide deployment lagging [27]

Table 1: Unicast routing incrementally deployed before ossification

at several layers. Although sometimes it is not called out explicitly, it is used extensively at the application layer in group messaging (e.g., email, text) and live video streaming. For the most part, however, multicasting at the application layer is supported by multiple unicasts at the network layer, because multicasting at the network layer (i.e., multicast routing) is not widely deployed.

Multicast routing is typically deployed today for niche applications that run on privately managed IP networks (see for example [41, 10]). There have been many discussions of why multicast is not deployed widely today in public networks (see for example [22]).

Milestones in the development of multicast routing are shown in Table 2. A good starting point is Dalal and Metcalfe’s 1978 paper that, while aiming to propose the reverse path forwarding technique, summarizes most other options for deploying multicast routing.

Probably the first proposal for deploying multicast in the Internet protocol stack as we know it today, was the 1984 paper by Aguilar in which he proposed using a form of multi-destination addressing, where the addresses of members of the group would be listed in each multicast packet. The technique had scalability limitations since one could only list up to 9 addresses in a single IP packet. Additionally, multicasting to this many destinations could easily be accomplished with separate unicast transmissions. For these reasons Aguilar’s proposal was not deployed.

I argue that this was an *anticipation* of the Cycle. This occurs when protocols are not deployed because they fail to satisfy some future requirement, even though they can be helpful in the present. In the case of Aguilar’s multicast proposal, the requirement for highly scalable multicast routing was not borne out of existing or even near term reality. The Internet around that time had around 1000 hosts¹. Aguilar’s mechanism would provide one order of magnitude scalability which would have been significant at the time. There was no urgent need for scalability beyond this potential order of magnitude savings. Furthermore, a multicast group exceeding 9 hosts could have easily been split into multiple smaller groups; something that was ultimately explored later [33, 8]. A modestly scalable multicast routing deployment would have been extremely useful to get experience and provide a starting point for the Cycle’s iterations.

Compare this with what happened in unicast routing, where the first distance vector protocol deployed had issues, including scalability limitations. Yet it served the purpose of getting things started. The main difference, of course, is the network *had* to have some form of unicast routing from the beginning, so we did not have the luxury of delaying deployment.

¹<https://www.zakon.org/robert/internet/timeline/>

As shown in Table 2, the late 1980s and early 1990s saw a progression of multicast routing protocols², each addressing the shortcomings of the one before or working to address a different set of constraints and goals. This may look like the Service-Infrastructure Cycle in operation, except it is important to note that none of these protocols saw contemporaneous *operational* deployment in support of applications or services beyond testbed environments such as the Mbone [23]. Around 2005 there seemed to be a consensus around a protocol that enabled so-called *Single Source Multicast (SSM)* as one that met scalability and other requirements. The problem is that by then the Internet had ossified and it was nearly impossible to deploy such a protocol, even though there were compelling arguments for its deployment.

To summarize, multicast routing development never had a chance of “riding” the Cycle. Early non-scalable ideas were rejected when scalability was not really needed. Experience with several generations of multicast routing protocols was gained only in testbed environments. When, finally, there appeared to be consensus on a workable multicast routing approach, it was no longer feasible to deploy it widely as ossification had already set in.

2.3 Takeaways

The takeaway message from this historical analysis is that before ossification, deployment success was achieved when functions had an opportunity to ride the Service-Infrastructure Cycle. Over multiple iterations, experience gained from deployment combined with motivation from new service or increased scale provided impetus for new and improved function deployment.

The two routing case studies illustrate how progress in support for novel services and in meeting demands imposed by increased scale requires the ability to ride the Service-Infrastructure Cycle. Delaying early deployment, for reasons such as the lack of compelling applications or a business case, and waiting for mechanisms that satisfy long-term predicted goals had a detrimental impact on deployment success.

The case studies also show how ossification of infrastructure, when it occurs, can have a profound effect on the ability of the infrastructure to support novel services. In the unicast routing case, ossification made it difficult to widely deploy enhancements to BGP4. Multicast routing deployment was hindered because by the time there was agreement on deployable and practical techniques, ossification had set in.

So what happens when the network infrastructure is unable to respond to new service or scale demands? We discuss this question in the next section.

²The table only highlights routing developments and omits many important support mechanisms and protocols that were also needed to completely deploy multicast.

Year	Milestone	Comments
1978	Dalal and Metcalfe [20]	Conceptual Framework for multicast routing
1984	Aguilar [6]	First IP-based routing using multidestination addressing
1988	Distance Vector Multicast Routing (DVMRP) [54]	Deployed in the MBone experimental testbed [16]
1994	Protocol Independent Multicast (PIM) [21]	Support for sparse multicast groups
2000	Small-Group Multicast [14]	A modernized version of Aguilar's 1984 proposal
2003	Source Specific Multicast [13]	Multicast support that finally addresses a real need

Table 2: Multicast routing does not get the benefit of deployment iterations before ossification

3. FROM MANYNETS TO ONENET (AND BACK AGAIN?)

I now consider what happens when the “unstoppable force” of continuously changing service demands and increasing scale meet the “immovable object” that is the ossified Internet. I will first dwell on something that many take for granted, namely, the fact that today a shared network infrastructure is being used to deliver communication services. I will then briefly discuss some of the new demands that the network infrastructure will need to address in the near future. Finally, I will conjecture that we are beginning to see evidence that the “immovable object” is showing signs of yielding to the “unstoppable force”.

3.1 The OneNet and its Consequences

One of the most under-appreciated milestones in networking is the transition from a many networks (ManyNets) world to one where one common network (OneNet), the Internet, provides global connectivity for almost all services. The history of how this transition occurred is well-documented [36] and I will not repeat it here. Suffice it to say that economy of scale and ubiquitous connectivity, both of which were primary goals of *initial* networking efforts, are well-served with a common global network.

The transition to OneNet was clearly in response to scale, global reach, and other demands on the network. As such, it fits within the general evolutionary paradigm of Figure 1. The transition itself took time and many iterations of the Cycle. First all data communication, email, file download and some content services were consolidated. Networks and network technology that were so commonplace to be textbook material (see for example [49]) like CCITT's X.25/X.75, IBM's SNA, Tymnet, and DECNET became of interest only as historical curiosities. Then, over time, voice and video services began to also be incorporated into the OneNet.

The OneNet transition had important consequences. The multiple-administrative-domains structure of the Internet today was conceived because it is not desirable to have a single entity manage a common network with global reach. Also a common network requires global agreement on network protocol standards which led to the OneNet agreement battles [48] ultimately leading to the adoption of the ARPANet architecture (a.k.a. the TCP/IP protocol suite) as the standard for global connectivity.

The OneNet world also has had strong implications for networking researchers. For a number of years, there was little acceptance for work that could not be ultimately deployed in the Internet. The OneNet world also made it very difficult for researchers to validate their proposals at scale through experimentation. This led to the many efforts

to build experimental networks: MBone [23], QBone [51], 6Bone [24], PlanetLab [19], and most recently GENI [12].

For our purposes, however, the most important OneNet consequence was its inevitable ossification. This would not have happened had we stayed with ManyNets. *By its very nature, a ManyNets world cannot ossify*, since new demands can always be satisfied with new networks.

3.2 New demands on network infrastructure

Not surprisingly, novel demands on the network infrastructure continue unabated. These demands need to be satisfied for the network infrastructure to maintain its usability. In the near term, the network infrastructure will need to handle:

- A significant increase in the number of connected devices with the advent of the Internet of Things (IoT). Estimates vary widely with a rough consensus around a total of 14 billion additional connected devices on the Internet by 2022 [43].
- The need for content providers to have increased control and accountability for the characteristics of network paths to their customers.
- Unique service demands for emerging applications - low latency (gaming, real-time control, high-frequency trading), prioritization and preemption (critical communication), and high-bandwidth (high-quality video streaming).

3.3 Fragmentation of network infrastructure

So how is the networking infrastructure responding to these demands? Since the ossified OneNet cannot, it appears that we are slowly being ushered back into a fragmented ManyNets world. In this world, some (most?) services and content providers will use physically and/or virtually separate networks to reach their users. A shared infrastructure will continue to be available to reach services that don't find it feasible to provide a separated infrastructure or as backup. Individual users will possibly not be aware of this fragmentation. I discuss below some evidence of the emergence of this ManyNets world.

The Flattening of the Internet.

The earliest symptoms of Internet fragmentation were detected in the work by Gill et al. [26] in the form of a *flattening* of the Internet. The paper observes that “large content providers are assembling their own wide-area networks.” This was followed by a study by Labovitz et al. [35] that discovered that “the majority of inter-domain traffic by volume now flows directly between large content providers, data center / CDNs and consumer networks.” More recently

Chiu et al. [18] observe that “Google connects directly to networks hosting more than 60% of end-user prefixes, and that other large content providers have similar connectivity.”

From a user’s perspective, a flat Internet connectivity is as depicted in Figure 2a. Most of the Internet traffic travels over connections between a user’s access network (access ISP) and a content provider (CP) network. The structure in Figure 2a represents, in principle, a common ubiquitous shared network providing global connectivity. In effect, however, for the majority of traffic this shared network exists to provide access to independently operated content provider networks.

The flattening of the network infrastructure has been going on for some time. In itself, it does not necessarily represent the beginning of Internet fragmentation. Recent trends, however, take this flattening a step further and are allowing content providers to *extend their control* into the shared access network itself. These include:

- The placement of content provider servers (also known as appliances) within access networks [5, 15]. These require access networks to provision specific bandwidth for access to these appliances, dedicating resources within access networks to specific content providers.
- Specific frameworks and interfaces made available to content providers to enable collaboration and coordination between them and access networks (see for example [3, 4, 25]).

These recent trends represent, in effect, a fragmentation of the *access* network, the only part of the flat Internet shared among the CPs in Figure 2a. They are leading us into a Zero-Hop Networking architecture [7] as shown in Figure 2b. This further confirms that the initial flattening was, indeed, a first step to a more complete Internet fragmentation.

Emergence of Bypass Networks.

I define a *bypass network* as a network that operates and is administered separately from the OneNet global communication infrastructure. Such networks carry traffic with special requirements that cannot be accommodated by the global infrastructure for performance and/or capacity (or scale) reasons. By definition, such networks are a phenomenon of the OneNet world. Bypass networks can be *private*, built and operated for use by a specific enterprise to carry the enterprise’s own traffic or *public*, meaning anyone can subscribe, attach to and use the network.

Private bypass network examples include Google’s own “purpose-built network,” [32, 53] which was built because Google “equipment vendors [were] slow in terms of delivering the capabilities requested by the company.” This is another way of saying that the vendors (or is it the network standards the vendors were following?) were slow or incapable of following the Service-Infrastructure Cycle with Google’s WAN demands being the driving force.

Another example of private bypass networks are the IPTV networks that cable television providers use to deliver digital content to users [10, 41]. These implement various technologies that are not deployed in the global infrastructure including IP Multicast to efficiently deliver digital content to cable modems.

One can argue that private bypass networks do not necessarily represent a fragmentation of the global Internet infrastructure. Google’s WAN and IPTV networks have limited

accessibility and as such are not in competition with the OneNet.

Motivated, however, by demands that the ossified OneNet cannot satisfy and, perhaps, encouraged by the feasibility of building separate networks demonstrated by private bypass networks, recent efforts have aimed to build global, publicly accessible, “purpose-built” networks. These networks are true bypass networks because: 1) their traffic could have been carried on the global Internet, if only it were capable of delivering adequate service, and 2) are publicly available for subscription and use.

One of the earliest public bypass networks was the one built by Spread Networks [2] in 2010 to support the low-latency requirements of high-frequency trading. The network boasts sub 13 ms round trip latency between New York and Chicago.

Another type of bypass network, that has seen recent significant development work and some global deployment, centers around the requirements of the Internet of Things (IoT). IoT has unique traffic requirements: traffic characterized often by the need for the sporadic exchange of short messages, the need to be power-efficient, and the expectation of very large scale deployments. The belief that the current Internet infrastructure cannot in its current form handle these requirements and cannot be changed to meet them has led to the development of specialized network technologies and architectures, known collectively as Low-Power Wide-Area Networks (LPWANs)[17]. LPWANs start with innovations in the access technologies, including the physical layer. Commercial deployments, however, such as SIGFOX [55] have included full networking protocol stacks and provide complete end-to-end service that does not include the use of Internet infrastructure.

Other examples of public bypass networks include *Haste* [29] and *FirstNet* [1]. *Haste* is a commercial bypass network that aims to provide very low latency for interactive gaming applications. It is a complete end-to-end solution that combines dedicated hardware as well as sufficient provisioning of resources to achieve its aims. *FirstNet* [1] is a network built specifically with disaster and emergency first responder communication needs. Its functions include prioritization and preemption capabilities.

Internet fragmentation is also evident in efforts to define 5G networks, a set of emerging conceptual architectures that aim to satisfy the demands of the future mobile and connected society [50]. Viewed within our Service-Infrastructure Cycle framework, 5G networks aim to provide the flexibility that will ensure the Cycle continues to operate well into the future. A key enabler of this “lubrication” of the Cycle is the idea of 5G slicing. In the 5G slicing proposal [31] the end-to-end network is sliced into networks with three different sets of capabilities: 1) A high throughput slice, 2) an ultra low latency slice, and 3) a large scale connectivity slice (e.g., for IoT devices). While all of this is still on the drawing board for now, it clearly points to a mindset where the future networking needs are no longer satisfied by a common global infrastructure.

4. CONCLUDING REMARKS

In the 1970s, a group of networking researchers and engineers set out to deploy a communication infrastructure that would act as “utilities [that] will offer a variety of information services and transactions, such as retrieval from

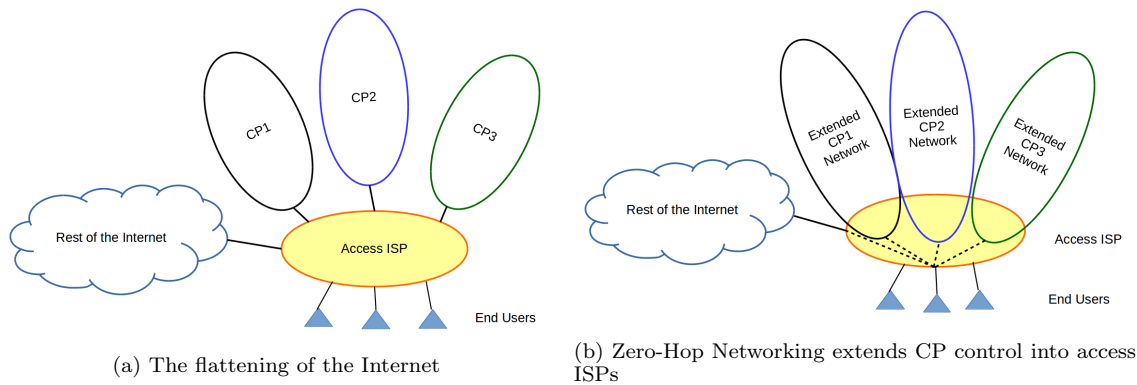


Figure 2: Flattening as a prelude to zero-hop fragmentation

multiple independent databases, messaging, electronic mail, conferencing, banking, tele-shopping and interest matching.” Meanwhile around the same time another group of communication engineers were beginning to deploy another communication network. That network, while motivated by a grand vision from the decade before of providing “Man-Machine Symbiosis” and of “The Computer as a Communication Device”, started with the simple goal of providing computer-to-computer connectivity and resource sharing.

If you guessed that the first communication infrastructure was the precursor of today’s Internet, you are mistaken. It was the vision of what were called then Videotex networks [11]. The second network – now you know – was indeed the Internet’s direct ancestor [36], delivering the beginnings of a vision laid out by J. C. R. Licklider in the 1960s[37].

Today the first network, Videotex, does not exist because it attempted to meet all its goals in one shot [42]. The second became the Internet and provides all the services envisioned by the first network and then some. I have argued here that the success of the Internet today in delivering services should be attributed to a modest start that was followed by iterations of the Service-Infrastructure Cycle. Similarly the failure of the Videotex network was not because of its vision, but rather because it anticipated the Cycle and tried to accomplish too much, too fast.

I have also argued and provided some evidence that the ossification of the Internet infrastructure along with continued new demands are requiring the fragmentation of the infrastructure so that the Service-Infrastructure Cycle can continue to operate – *from one, many*.

We are left with the question: *So what if the communication infrastructure is fragmenting?* Overall, it is not a bad thing and we may not have any choice in the matter. I do, however, have a couple of concerns.

My first concern is not about the fragmentation itself but that, despite the re-emergence of the ManyNets world, the networking research community might continue to expend its energy to further the illusion of a single global networking infrastructure. For a while after infrastructure ossification was observed, there was a search for either a fix that would de-ossify the Internet (through virtualization, for example [52]), or a single replacement for the Internet that would be immune from ossification by design [44]. While these efforts have had significant intellectual impact, they have yet to have a perceptible impact on the deployed infrastructure. The idea of a single unified networking infrastructure was

perhaps worth fighting for and saving. However, this battle is lost. We should, therefore, embrace the ManyNets world and the research agenda that emerges from it.

In fact, the emergence of the ManyNets world will be a boon to networking research. It is a fertile ground for innovation as increased scale, with no end in sight, and new services continue to drive network evolution.

The second concern I have is that, while it lasted, the OneNet has served well in providing cost-efficient networked services. Perhaps, more importantly, it has enabled significant innovation over the years by providing a low cost of entry to new services. The challenge for the networking research community is how to continue these benefits in the emerging ManyNets world.

5. ACKNOWLEDGMENT

Thanks to Jorg Liebeherr for giving me the time and space at the University of Toronto to work on this article during summer of 2017. The many conversations we had on this topic and his feedback on an early draft helped shape some of the arguments. Marwan Fayed and Ellen Zegura also provided valuable feedback on early drafts. Ken Calvert, Raouf Boutaba and Yashar Ganjali provided valuable feedback through their contributions to the panel at the 2017 IEEE ICNP conference that debated the arguments in this article and considered its implications on the networking research community. Finally, thanks go to Olivier Bonaventure for his comments on the submitted version of this editorial.

6. REFERENCES

- [1] FirstNet. <https://www.firstnet.gov/>. [Online; accessed 29-September-2017].
- [2] Spread Networks. <https://SpreadNetworks.com/>. [Online; accessed 29-August-2017].
- [3] AT&T Domain 2.0 Vision White Paper. https://www.att.com/Common/about_us/pdf/AT&T%20Domain%202.0%20Vision%20White%20Paper.pdf, 2013. [Online; accessed 21-August-2017].
- [4] AT&T ECOMP (Enhanced Control, Orchestration, Management & Policy) Architecture White Paper. <https://about.att.com/content/dam/snrdocs/ecomp.pdf>, 2013. [Online; accessed 21-August-2017].

- [5] Netflix Open Connect. <https://openconnect.netflix.com/en/>, 2017. [Online; accessed 23-August-2017].
- [6] L. Aguilar. Datagram routing for Internet multicasting. In *ACM SIGCOMM Computer Communication Review*, volume 14, pages 58–63. ACM, 1984.
- [7] M. Ammar, E. Zegura, and Y. Zhao. A Vision for Zero-Hop Networking (ZeN). In *Distributed Computing Systems (ICDCS), 2017 IEEE 37th International Conference on*, pages 1765–1770. IEEE, 2017.
- [8] M. H. Ammar and L.-R. Wu. Improving the throughput of point-to-multipoint ARQ protocols through destination set splitting. In *INFOCOM'92. Eleventh Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE*, pages 262–271. IEEE, 1992.
- [9] T. Anderson, L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. *Computer*, 38(4):34–41, 2005.
- [10] P. Arberg, T. Cagenius, O. V. Tidblad, M. Ullerstig, and P. Winterbottom. Network infrastructure for IPTV. https://www.ericsson.com/res/thecompany/docs/publications/ericsson_review/2007/2_NetworkInfrastructure.pdf, 2007. [Online; accessed 29-August-2017].
- [11] A. Ball, G. Bochmann, and J. Gecsei. Videotex networks. *Computer*, (12):8–14, 1980.
- [12] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar. Geni: A federated testbed for innovative network experiments. *Computer Networks*, 61:5–23, 2014.
- [13] S. Bhattacharyya. An Overview of Source-Specific Multicast (SSM). Technical report, RFC 3569, 2003.
- [14] R. Boivie, N. Feldman, and C. Metz. Small group multicast: A new solution for multicasting on the Internet. *IEEE Internet Computing*, 4(3):75–79, 2000.
- [15] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the expansion of Google’s serving infrastructure. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 313–326. ACM, 2013.
- [16] S. Casner and S. Deering. First IETF Internet Audiocast. *SIGCOMM Comput. Commun. Rev.*, 22(3):92–97, July 1992.
- [17] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi. Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios. *IEEE Wireless Communications*, 23(5):60–67, 2016.
- [18] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are we one hop away from a better Internet? In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, pages 523–529. ACM, 2015.
- [19] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: an overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communication Review*, 33(3):3–12, 2003.
- [20] Y. K. Dalal and R. M. Metcalfe. Reverse path forwarding of broadcast packets. *Communications of the ACM*, 21(12):1040–1048, 1978.
- [21] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu, and L. Wei. An architecture for wide-area multicast routing. *ACM SIGCOMM Computer Communication Review*, 24(4):126–135, 1994.
- [22] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the IP multicast service and architecture. *IEEE network*, 14(1):78–88, 2000.
- [23] H. Eriksson. Mbone: The multicast backbone. *Communications of the ACM*, 37(8):54–61, 1994.
- [24] R. Fink and R. Hinden. 6bone (ipv6 testing address allocation) phaseout. Technical report, 2004.
- [25] B. Frank, I. Poese, Y. Lin, G. Smaragdakis, A. Feldmann, B. Maggs, J. Rake, S. Uhlig, and R. Weber. Pushing CDN-ISP collaboration to the limit. *ACM SIGCOMM Computer Communication Review*, 43(3):34–44, 2013.
- [26] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening Internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *International Conference on Passive and Active Network Measurement*, pages 1–10. Springer, 2008.
- [27] S. Goldberg. Why is it taking so long to secure Internet routing? *Communications of the ACM*, 57(10):56–63, 2014.
- [28] K. Hafner and M. Lyon. *Where wizards stay up late: The origins of the Internet*. Simon and Schuster, 1998.
- [29] haste. <https://haste.net/haste-technology/>. . [Online; accessed 29-August-2017].
- [30] F. Heart. Interface message processors for the ARPA computer network: Quarterly report number 1783, April 1969.
- [31] Huawei. 5G Network Architecture A High-Level Perspective. <http://www.huawei.com/minisite/hwmbbf16/insights/5G-Nework-Architecture-Whitepaper-en.pdf>, 2007. [Online; accessed 29-August-2017].
- [32] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al. B4: Experience with a globally-deployed software defined WAN. *ACM SIGCOMM Computer Communication Review*, 43(4):3–14, 2013.
- [33] S. K. Kasera, J. Kurose, and D. Towsley. Scalable reliable multicast using multiple multicast groups. *ACM Sigmetrics performance evaluation review*, 25(1):64–74, 1997.
- [34] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected areas in Communications*, 18(4):582–592, 2000.
- [35] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. In *ACM SIGCOMM Computer Communication Review*, volume 40, pages 75–86. ACM, 2010.
- [36] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff. A brief history of the Internet. *ACM SIGCOMM Computer Communication Review*, 39(5):22–31, 2009.
- [37] J. C. R. Licklider. Man-Machine Symbiosis and The Computer as a Communication Device. <http://memex.org/licklider.pd>. [Online; accessed

- 29-August-2017].
- [38] K. Lougheed and Y. Rekhter. RFC 1105: A Border Gateway Protocol, 1989.
- [39] C. Lynn, J. Mikkelsen, and K. Seo. Secure BGP (s-bgp). *IETF Draft: draft-lynn-s-bgp-protocol-01.txt*, 2003.
- [40] J. M. McQuillan, I. Richer, and E. C. Rosen. An overview of the new routing algorithm for the ARPANET. In *Proceedings of the sixth symposium on Data communications*, pages 63–68. ACM, 1979.
- [41] C. I. M. D. Network. National market systems. 2017.
- [42] A. M. Noll. Videotex: Anatomy of a failure. *Information and Management*, 9(2):99 – 109, 1985.
- [43] A. Nordrum. Popular Internet of Things Forecast of 50 Billion Devices by 2020 Is Outdated. <http://spectrum.ieee.org/tech-talk/telecom/internet/popular-internet-of-things-forecast-of-50-...>, 2016. [Online; accessed 21-August-2017].
- [44] J. Pan, S. Paul, and R. Jain. A survey of the research on future Internet architectures. *IEEE Communications Magazine*, 49(7), 2011.
- [45] R. Prosser. Routing procedures in communications networks-part i: Random procedures. *IRE Transactions on Communications Systems*, 10(4):322–329, 1962.
- [46] Y. Rekhter and T. Li. RFC 1771. *A Border Gateway Protocol*, 4:1–54, 1995.
- [47] E. C. Rosen. RFC 827: Exterior gateway protocol (EGP). Retrieved September, 23:2008, 1982.
- [48] A. L. Russell. 'rough consensus and running code' and the Internet-OSI standards war. *IEEE Annals of the History of Computing*, 28(3):48–61, 2006.
- [49] M. Schwartz. *Telecommunication networks: protocols, modeling and analysis*, volume 7. Addison-Wesley Reading, MA, 1987.
- [50] M. Shafi, A. F. Molisch, P. J. Smith, T. Haustein, P. Zhu, P. De Silva, F. Tufvesson, A. Benjebbour, and G. Wunder. 5G: A Tutorial Overview of Standards, Trials, Challenges, Deployment, and Practice. *IEEE Journal on Selected Areas in Communications*, 35(6):1201–1221, 2017.
- [51] B. Teitelbaum, S. Hares, L. Dunn, R. Neilson, V. Narayan, and F. Reichmeyer. Internet2 QBone: building a testbed for differentiated services. *IEEE network*, 13(5):8–16, 1999.
- [52] J. S. Turner and D. E. Taylor. Diversifying the Internet. In *Global Telecommunications Conference, 2005. GLOBECOM'05. IEEE*, volume 2, pages 6–pp. IEEE, 2005.
- [53] A. Vahdat, D. Clark, and J. Rexford. A purpose-built global network: Google's move to SDN. *Queue*, 13(8):100, 2015.
- [54] D. Waitzman, S. Deering, and C. Partridge. Distance vector multicast routing protocol. 1988.
- [55] J. C. Zuniga and B. Ponsard. Sigfox system description. *LPWAN@ IETF97, Nov. 14th*, 2016.