

Point/Counterpoint

Future Internet Architecture: Clean-Slate Versus Evolutionary Research

Should researchers focus on designing new network architectures or improving the current Internet?

OVER THE PAST several years, the networking research community has engaged in an ongoing conversation about how to move the field—and the Internet itself—forward. These discussions take place in the context of the tremendous success of the Internet, begging the question of whether researchers should focus on understanding and improving today's Internet or on designing new network architectures that are unconstrained by the current system. Ultimately, individual researchers have their own styles, often a unique blending of both approaches. In this Point/Counterpoint, Jennifer Rexford and Constantine Dovrolis debate the pros and cons of “clean slate” and “evolutionary” approaches to networking research, reflecting on the larger discussion taking place in the networking research community.


Point: Jennifer Rexford

The Internet is an undeniable success—a research experiment that escaped from the lab to become a major part of the global communications infrastructure. The seeds of the Internet's success lie in its “underspecified” design—a minimalist network providing a simple best-effort packet-delivery service coupled with programmable computers at the end points. These early design decisions were so important because they lowered the barriers to in-

novation in new applications (created by anyone who wants to program these computers) and link technologies (that can be easily adopted if they support the basic packet-delivery model). This has led to innovation far beyond what any of the early designers of the Internet could have ever imagined.

Given the Internet is so successful, and apparently so accommodating of innovation, “clean slate” networking

research may seem strange, even superfluous. Yet, nothing could be further from the truth. In fact, clean-slate design is important for enabling the networking field to mature into a true discipline, and to have a future Internet that is worthy of society's trust. Contrary to the very premise of our debate, I do not believe that evolutionary and clean-slate research are at odds. Insights from clean-slate research can



Nodal representation of the Internet.

(and should) help guide the ongoing evolution of the Internet, and a clean-slate redesign may be necessary for the Internet's continued evolution into a secure, reliable, and cost-effective infrastructure. Most importantly, as a research community, we should plant the seeds that will enable future research experiments to "escape from the lab."

Toward a Networking Discipline

The success of the Internet does not mean the field of networking is mature. Far from it. The Internet has grown and changed much faster than our own understanding of how to design, build, and operate large, federated networks. This is a common phenomenon in engineering. The great medieval cathedrals were built long before the field of civil engineering was in place. As a result, many of these early cathedrals collapsed under their own weight after decades of construction. Even the collapsed cathedrals were an invaluable learning experience along the long road toward a more rigorous approach to designing and building large structures. They were a step in the journey, not the destination itself. The way we design large buildings today reflects more than incremental improvements in engineering techniques, but a fundamentally more principled approach to the problem.

Whenever the Internet faces new challenges, from the fears of congestion collapse in the late 1980s to the pressing cybersecurity concerns of today, new patches are introduced to (at least partially) address the problems. Yet, we do not yet have anything approaching a discipline for creating, analyzing, and operating network protocols, let alone the combinations of protocols and mechanisms seen in real networks. Networking is not yet a true scholarly discipline, grounded in rigorous models and tried-and-true techniques to guide designers and operators. Witness any networking class or textbook, riddled as they are with descriptions of existing protocols rather than a top-down treatment of the "laws" or even "rules of thumb" governing the design, analysis, and operation of these protocols. Given the critical importance of communication networks, we need the field to mature into a discipline we can apply confidently in practice and teach

As a research community, we should plant the seeds that will enable future research experiments to "escape from the lab."

effectively to our students.

While studying today's Internet is clearly an important part of maturing the field, it is not enough; we also need exploration that is unfettered by today's artifacts. To be clear, ignoring today's artifacts does not mean ignoring reality. Any new designs must still grapple with practical constraints (such as the speed of light, or limitations on computation, memory, and bandwidth resources) and design requirements (for goals like efficiency, security, privacy, reliability, performance, ease of management, and so on). Yet, a clean-slate design process could remain free of the considerable minutiae of today's protocols and operational practices, and the challenges of incremental deployment.

A clean-slate design process can topple the underlying assumptions of today's architecture, such as asking whether we can achieve scalability without relying on hierarchical addressing, route traffic directly on the name of a service rather than the address of a machine, or have notions of identity that cannot be spoofed. This clean-slate exploration can lead to valuable new designs that fill out the large design space, expanding our knowledge and experience. This exploration can, perhaps more importantly, lead to new methodologies for designing networks and protocols. Whether and how to deploy these new ideas in today's Internet, while certainly a worthy topic in its own right, should sometimes be secondary to the broader goal of deepening our understanding of the field. The measure of successful

research should be the greater depth of our understanding, not just the breadth of deployment.

Yet, clean-slate networking research cannot stop at pencil-and-paper designs. In addition to new ideas, and rigorous theoretical models and analysis, we need to push our ideas further into real implementations and (ideally) deployments. The "Eureka" moments that lead to real progress happen when we encounter surprises, when something happens that we could never have planned or predicted. Building, evaluating, and deploying real systems—on experimental facilities such as the proposed GENI and Federica platforms (in the U.S. and Europe, respectively)—exposes our nascent ideas to the harsh light of day, and gives us the feedback necessary to help our ideas grow sharper and stronger as we address the unexpected setbacks and limitations, and embrace the practical constraints and design requirements we were unwittingly ignoring.

Building and deploying our designs is more than just the last step in evaluating an idea—it is part of a continuous cycle of research, constantly refining the problem, the models, and the solutions until a more complete understanding emerges. This approach to networking research should sound familiar—it is exactly how the early ARPAnet was designed and built, leading to the amazing advances we have seen in the 40 years since the first message was delivered over the network we would come to call "the Internet." At the time, the notion that the ARPAnet would eventually overtake the established telecommunication networks of its day was inconceivable to most people. But, we know now how that story turned out.

Toward an Internet Worthy of Our Trust

The Internet is showing signs of age. Pervasive security problems—spam, denial-of-service attacks, phishing, and so on—are only the most visible symptoms. The Internet also does not handle mobile hosts, whether users on the move or virtual machines migrating from one computer to another, all that well. The Internet's best-effort service model is a poor match for many real-time applications, such as IPTV and videoconferencing. The Internet is not

reliable enough, due to equipment failures, software bugs, and configuration mistakes. Managing a large network is too expensive—often costing more than the underlying equipment—and tremendously error prone. The Internet consumes too much energy, in an era of serious concern about global warming. The Internet does not seem ready to handle the coming onslaught of countless small sensor devices that have the potential to revolutionize our world. The list goes on and on.

Many of these pressing challenges are deeply rooted in early design decisions underlying the Internet, and may not be solvable without fundamental architectural change. For example, many security problems relate to the Internet's weak notions of identity, and particularly the ease of spoofing everything from IP addresses to domain names, from email addresses to routing information. Stronger notions of identity are not easily retrofitted on today's architecture. Mobility is difficult to handle because IP addresses are hierarchical and tightly coupled with the scalability of the routing protocols. Breaking this coupling may require a new relationship between naming, addressing, and routing. Network management is difficult because of the current "division of labor" between the distributed protocols running on the network elements and the management systems that can only indirectly tune the many knobs these protocols expose. Solving these problems may require us to revisit some of the most basic principles underlying the Internet of today.

Clean-slate research allows us to explore radically new designs, to see if

they are viable alternatives to the solution we have now. Some of these clean-slate solutions may very well have an incremental path to deployment. But, as the American baseball legend Yogi Berra famously said, "You've got to be very careful if you don't know where you're going, because you might not get there." Clean-slate research can help us determine where we should be going. Clean-slate design may also help us decide what parts of the Internet should *not* change. Perhaps, despite the challenges facing today's Internet, we fundamentally cannot do much better along some dimensions (say, security) without paying too high a price along some other dimension. Clean-slate research can help us understand those trade-offs, to guide decisions about whether and what to change.

Finally, perhaps wholesale change is both necessary and possible. Despite enabling innovation in applications and link technologies, the Internet architecture itself is remarkably resistant to change. In redesigning the Internet, we can direct much-needed attention to this problem. Making the inside of the network more programmable, and allowing multiple independent designs to coexist in parallel, are a promising start in this direction. Perhaps the future Internet could have the seeds for its own constant reinvention lying within it. We are already seeing the early fruits of this kind of clean-slate thinking, in software-defined networking infrastructures like OpenFlow (<http://www.openflowswitch.org/>) that are being deployed in several enterprise, datacenter, and backbone networks. Even experimental infra-

structures like GENI and Federica, designed as they are to enable multiple simultaneous experiments with new network architectures, are themselves examples of this kind of change.

Fundamental change like this is, indeed, possible and it is already starting to happen, due to the early clean-slate research efforts over the past several years. Further, more substantive change can happen in the years ahead. Given the Internet largely supplanted the circuit-switched telephone networks, is it so farfetched to think that something else might supplant the Internet, or so significantly alter the Internet that we no longer recognize it from the descriptions we see in today's networking textbooks?

Conclusion

Networking is still a young field. While the Internet's success is something we should admire and celebrate, we should not be content with our current understanding of the field or view the Internet architecture as set in stone. Perhaps a new generation of researchers and practitioners will turn the future Internet into something that only vaguely resembles its predecessor. Perhaps this future network will accommodate change more broadly and deeply than even today's Internet has. A willingness to step back, and design from scratch, is an important part of the research repertoire that can enable these advances in the field, and of the Internet itself. □

Jennifer Rexford (jrex@cs.princeton.edu) is a professor in the computer science department at Princeton University in New Jersey.

Copyright held by author.

Counterpoint: Constantine Dovrolis

LET US FIRST identify the major difference between the two approaches. Evolutionary Internet research aims to understand the behavior of the current Internet, identify existing or emerging problems, and resolve them under two major constraints: first, backward compatibility (interop-

erate smoothly with the legacy Internet architecture), and second, incremental deployment (a new protocol or technology should be beneficial to its early adopters even if it is not globally deployed).

On the other hand, clean-slate research aims to design a new "Future Internet" architecture that is significantly better (in terms of performance, security, resilience, and other properties) than the current Internet without

being constrained by the current Internet architecture.

Clean-Slate Research and Its Real-World Impact

Clean-slate Internet research is not something new. In fact, there is a long history of such efforts and we can learn something by analyzing whether earlier clean-slate protocols and architectures have been adopted or not. To name a few examples, consider active

networks, per-flow QoS guarantees and admission control, the connection-less network protocol CLNP, transport protocols such as XCP, or interdomain routing architectures such as Nimrod. There is also a large number of protocols that are more or less backward compatible but not truly incrementally deployable, such as IPv6, interdomain IP multicast, RSVP, and IntServ, IPsec, or S-BGP. Arguably, these protocols have not seen large-scale deployment, at least so far. The “real world” adopted instead evolutionary approaches such as NATs, caching and content distribution networks, DiffServ, adaptive applications, and various security mechanisms (such as end-host security, intrusion detection systems, and routing filters) that work well with the legacy architecture. Why does clean-slate architectural research, or even protocols and designs that attempt to be backward compatible, often fail to be adopted in practice?^a

In industrial economics, it is well known that an emerging technology that is subject to network externalities will probably not be able to replace a widely deployed but inferior technology, as long as there are costs involved in switching from the incumbent to the emerging technology (see Arthur¹ and related papers). Instead, the more relevant question is whether the emerging technology offers a valuable new service the current technology cannot

a I do not claim that the research on those earlier clean-slate protocols was mediocre or that it did not have academic impact—I am strictly focusing on their deployment and real-world impact.

How does the additional value of a new technology, relative to the incumbent technology, compare to the transition cost?

The ARPANET architecture was only one of several competing architectures and it was through a long evolutionary process that it prevailed.

provide directly or indirectly. In other words, how does the additional value of a new technology, relative to the incumbent technology, compare to the transition cost?

It is not enough for a clean-slate architecture to be “better” than the current Internet architecture. For the former to have real impact it should be able to replace the latter—otherwise it will remain an intellectual exercise. It is the question of real-world impact that differentiates clean-slate from evolutionary research and design. And at least so far, the proponents of clean-slate research have not shown instances of such new applications or services that cannot be directly or indirectly constructed for the current Internet. Incidentally, the promise of a “secure and trustworthy Future Internet” is appealing but not convincing: there is no way to provide security guarantees with an open-ended threat model. Further, it is very likely that a brand-new internetworking architecture will have more design and implementation bugs and security holes than the current Internet architecture (which is being “debugged” for more than 30 years now).

The proponents of clean-slate design emphasize they will not stay with “paper designs”—they will build and experiment with the proposed architectures in testbeds such as GENI. But what would that prove? Several previous clean-slate protocols were also implemented and tested 10 or 20 years ago. The issue was not the lack of implementation or experimentation, but the fact that those protocols could not compete with incumbent

technologies, considering the actual benefits they provide to users and the costs involved in the technological transition. These are issues of mostly economic nature that GENI or other testbeds cannot help us study. Further, these testbeds are not used by real applications and people and they do not operate under the economic and policy constraints of the real world. The early ARPANET succeeded because it was not just a testbed: it was also used as a production network, connecting some universities and research labs, while at the same time networking researchers could experiment with new protocols and technologies.

Another popular claim is that the current Internet architecture is the result of clean-slate thinking back in the 1960s or 1970s. However, we should not ignore that packet switching or TCP/IP were not inventions that “came out of nowhere”—they resulted from an evolutionary process that started from synchronous multiplexing in circuit-switched networks, moving to asynchronous multiplexing and then to datagram forwarding. Further, the ARPANET architecture was only one of several competing architectures (such as IBM SNA, DECnet, ITU X.25, Xerox Pup, SITA HLN, or CYCLADES), and it was through a long evolutionary process that the former eventually prevailed.

Is the Internet Architecture Really “Ossified”?

One of the primary arguments for clean-slate research has been that the current Internet architecture is ossified, especially at the central layers of the protocol stack (IP and TCP), and that ISPs have no incentive to adopt any architectural innovations. This is a rather negative view of what happens. The Internet architecture maps an ever-increasing diversity of link-layer technologies to a rapidly increasing range of applications and services. To support this innovation at the lowest and highest layers of the architecture, the central protocols of the architecture must evolve very slowly so that they form a stable background on which diversity and complexity can emerge.

To use a biological analogy, certain developmental Gene Regulatory Networks were established in the Early

Cambrian (about 510 million years ago) and they have not evolved significantly since then. These GRNs are referred to as evolutionary kernels, and it is now understood that they are largely responsible for major aspects of all animal body plans. For instance, the heart of a fruit fly and the heart of a human, despite distinct morphologies, develop using the same core cardiac GRN. Evolutionary kernels represent a stable basis on which diversity and complexity of higher-level processes can evolve.²

An Agenda for Evolutionary Internet Research

Instead of thinking about the Internet as an artifact that we designed in the past and we can now redesign, we can start thinking of the Internet as an evolving ecosystem that is affected by, and in turn is affecting, several disciplines and how we study them. Its evolution is controlled, not only by technology, but also by the global economy, creative ideas by millions of individuals, and a constantly changing set of “environmental pressures” and constraints. Our mission then, as Internet researchers, is to first measure and understand the current state of this ecosystem, predict where it is heading and the problems it will soon face, and create what could be referred to as intelligent mutations: innovations that can, first, avoid or resolve those challenges, and second, innovations that can be adopted by the current architecture in a way that is backward compatible and incrementally deployable. This is a pragmatic research agenda that can have real impact on millions of people.

Instead of testbeds, evolutionary research needs various experimental resources that will be integrated in the current Internet. First, we need a dense infrastructure of “Internet monitors” of various types that will allow us to accurately measure what is currently happening in this evolving ecosystem. It is embarrassing that (despite the tremendous value of the Route Views project) we still do not have an accurate way to measure the Internet interdomain topology. We also do not have an estimate of how much traffic flows between any two autonomous systems, even though that interdomain traffic matrix largely determines the economics of the global Internet. Plus, we have no way to know

We can start thinking of the Internet as an evolving ecosystem that is affected by, and in turn is affecting, several disciplines and how we study them.

how the Internet population uses the Internet and the Web across time and space. As this knowledge gap increases, I am concerned we will soon be unable to track our own creation, and much more to influence its future.

Together with an extensive monitoring infrastructure, evolutionary Internet research would greatly benefit if we could operate our own experimental ISP. This would be a real TCP/IP network, running all protocols of the current Internet architecture, present at many Internet Exchange Points, peering openly with other ISPs and content providers, and carrying traffic that belongs to real Internet users. One way to do so could be that universities use this experimental ISP to carry part of their traffic for free, with the understanding that this is a research network and so its traffic may be subject to experimental “mutations” of the Internet architecture. This is different than Internet2 or NLR, which are production networks, and certainly very different than isolated GENI-like testbeds.

Where is the Science, After All?

The proponents of clean-slate design claim their approach leads to a science of network design (sometimes referred to as “network science,” which is confusing because the same term is used in other disciplines to refer to the study of complex systems using dynamic graph models and network analysis techniques). It is also often claimed that evolutionary Internet research is not a science, but a collection of “hacks” and incremental improvements. This is a misleading position.

Several breakthroughs in networking research resulted from evolutionary research. For instance, major results in congestion control and active queue management resulted from attempts to understand and improve TCP, the discovery of fundamental properties of the Internet traffic and topology, the design of innovative peer-to-peer communication protocols, or the development of end-to-end network inference as well as network tomography methods.

A domain of knowledge does not become science because it is based on clean optimization frameworks or because it proves deep results about toy models. Good science requires relevance to the real world, measurements and experimental validation, testable hypotheses, and models with predictive power.

Epilogue

I often wonder, what is the main reason that well-respected Internet researchers have decided to pursue the clean-slate approach? It cannot be just the “funding carrot,” I am sure. Here is one possible answer from a science fiction TV series. In “Battlestar Galactica” (S4-E21),” Mr. Lampkin says to Commander Adama: “I have to say I’m shocked with how amenable everyone is to this notion of (...leaving everything behind and starting with nothing on the newly discovered planet Earth).” Commander Adama responds “Don’t underestimate the desire for a clean slate, Mr. Lampkin.” It may be that we find joy and pride in the idea that we can redesign the Internet from scratch, that we can avoid all previous mistakes and do it perfectly this time. If we do not want to sound like science fiction dialogue, however, it is important that we continue to foster the evolution of the current Internet, having positive impact on the way many millions of people live, work, and communicate. ☑

References

1. Arthur, W.B. Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99, 394 (1989), 116–131.
2. Dovrolis, C. and Streebman, T. Evolvable network architectures: What can we learn from biology? *ACM SIGCOMM Computer Communications Review (CCR)* 40, 2 (Apr. 2010).

Constantine Dovrolis (dovrolis@cc.gatech.edu) is an associate professor in the College of Computing at Georgia Tech in Atlanta, GA.

Copyright held by author.