Diversifying the Internet

Jon Turner and David Taylor
Washington University in Saint Louis

The Internet has fallen victim to its own stunning success. The interplay of the end-to-end design of IP and the vested interests of competing stakeholders has led to its growing ossification. Alterations to the Internet architecture that address its fundamental deficiencies or enable new services have been restricted to incremental changes. The slow pace of this process stifles innovation and the adoption of disruptive technology. A recent call to arms advances a research agenda to confront this impasse through virtualization [1]. In addition to describing a virtual testbed for the evaluation of new network architectures, it poses a question about the long-term role of virtualization in the Internet. The architectural "purist" views virtualization as a tool for architecture evaluation and the periodic deployment of successive, singular Internet architectures. In this paper, we advance the "pluralist" view that seeks to make virtualization an architectural attribute of the Internet. By enabling a plurality of diverse network architectures to coexist on a shared physical substrate, virtualization mitigates the ossifying forces at work in the current Internet and enables continual introduction of innovative network technologies. Such a diversified Internet would allow existing architectural deficiencies to be holistically addressed as well as enable the introduction of new architectures supporting new types of applications and services.

Network ossification

In a relatively short period of time, the Internet has become critical infrastructure for global commerce, media, and defense. Like many successful technologies, the Internet is suffering the adverse effects of inertia. The significant capital investment and competing interests of its major stakeholders creates a barrier to the introduction of disruptive technologies. Furthermore, the end-to-end design of IP requires global agreement and coordination to deploy changes. Working in concert, these two factors prevent existing problems from being holistically addressed and stifle innovation. In the current environment, the incremental deployment of the most basic and necessary changes such as IPv6 is painfully slow [2].

While the Internet's success speaks to the utility of its architecture for a wide range of applications, a best-effort packet delivery service is not best for all purposes. The ever-expanding scope and scale of the Internet's use has also exposed a number of fundamental deficiencies in the current architecture. Security, routing stability and control, and quality of service guarantees are a few of the most significant [3-5]. For over ten years, applications such as global video conferencing, telephony, and broadcast television have been touted as promising "next-generation" applications. It was widely believed that this class of applications justified the tremendous investment in dark fiber and would spur the next round of vigorous innovation and deployment [6]. This promise remains largely unrealized due to the inability of the current architecture to support these applications, the inability to change the architecture, and the prohibitively high cost of deploying a custom global network.

The ossification of the Internet architecture and its inability to support many types of applications impose a significant barrier to innovation. For the research community, the ability to perform large-scale experimentation that seeks to affect the Internet's operation is almost entirely limited to overlay networks. It can bring innovation to the network layer only through more efficient implementations of existing mechanisms. While there is opportunity to deploy new higher layer protocols and new link and physical layer technologies, the network layer has become untouchable.

Enabling diversity

Virtualization has been advanced as a means to break this impasse [1, 7, 8]. By enabling experimentation with new network architectures, a virtual testbed provides immediate value to the research community.

We discuss the development of such a testbed in the last section of this paper. While reinvigoration of applied architectural research has clear benefits, it is the introduction of new technologies into the Internet that can provide the more profound impact on society and which also presents the most significant challenge.

In the context of deployment, we characterize the architectural purist as viewing virtualization as a tool for experimentation and the periodic deployment of each in a possible succession of new Internet architectures. The theory here is that by supporting applications that attract a significant number of users, a successful network architecture operating on a virtual testbed can apply pressure on the current Internet stakeholders to adopt a new architecture. In this environment, experimentation may be continual, but adoption of a new network architecture is relatively rare. We seek to advance the pluralist view that virtualization should become a fundamental attribute of the Internet so that the problem of network ossification can be solved, once and for all.

In crafting a network, architects must invariably choose from among numerous parameters. One of the first considerations is often scope. Networks may be application-specific, general-purpose, or a hybrid that supports a certain class of applications. Networks may provide best-effort service or multiple flavors of guaranteed service. Other considerations include the type and level of security, control granularity and interfaces, and network scale. These architectural parameters then drive the design of the data and control planes. The architectural purist requires that a network architecture be general-purpose. It must provide a suitable platform for the set of all existing applications, while attempting to support the needs of future applications [9]. If the architecture is insufficient for a given application, then the application must wait until a new, more suitable architecture is adopted. The purist view also requires that each network node provide all of the shared services, regardless of the needs of the application traffic traversing the node. More importantly, we believe that the purist view does not sufficiently address the ossifying forces currently at work in the Internet. If the "narrow waist" of the Internet is a single end-to-end packet delivery service, then modification of that packet delivery service still requires universal agreement and coordination amongst the competing stakeholders. In this environment, virtualization (i.e. a virtual testbed) does not possess significantly more leverage than previous testbeds that have been unable to successfully impact the Internet.

The pluralist approach provides a means to permanently guard against these ossifying forces while providing additional design freedom for network architects [10]. In the remainder of this paper, we advance this view by describing a diversified Internet that supports a plurality of network architectures coexisting on a shared physical substrate network. We refer to the diverse networks comprising this diversified Internet as interlay networks to emphasize the system's interlaid structure, to distinguish the concept from overlay [8] and underlay [7] approaches, and to avoid the much overloaded term, "virtual network". A diversified Internet that supports multiple interlay networks allows new network architectures to be introduced alongside incumbent architectures and given the opportunity to succeed (or fail) on their own merits. By allowing end users to opt-in to a variety of interlay networks, network architectures are exposed to the same market forces as other technologies. We argue that exposing network technologies to such competitive forces is the only sure way to enable continuing innovation and change.

Interlay networks

Interlay networks enable fundamentally different end-to-end packet delivery mechanisms by allowing each network to specify packet formats, addressing methods, packet forwarding, routing protocols, etc. Interlay networks are implemented on top of a *substrate* comprising the physical resources used by the various interlays. In this model, the "narrow waist" of the Internet architecture is a thin network provisioning layer that provides automated mechanisms for allocating the substrate's resources to different interlays. The role of the provisioning layer is to allow network providers to automatically deploy, configure and operate interlay networks. In general, the provisioning layer does not restrict the kind of new services that may be deployed.

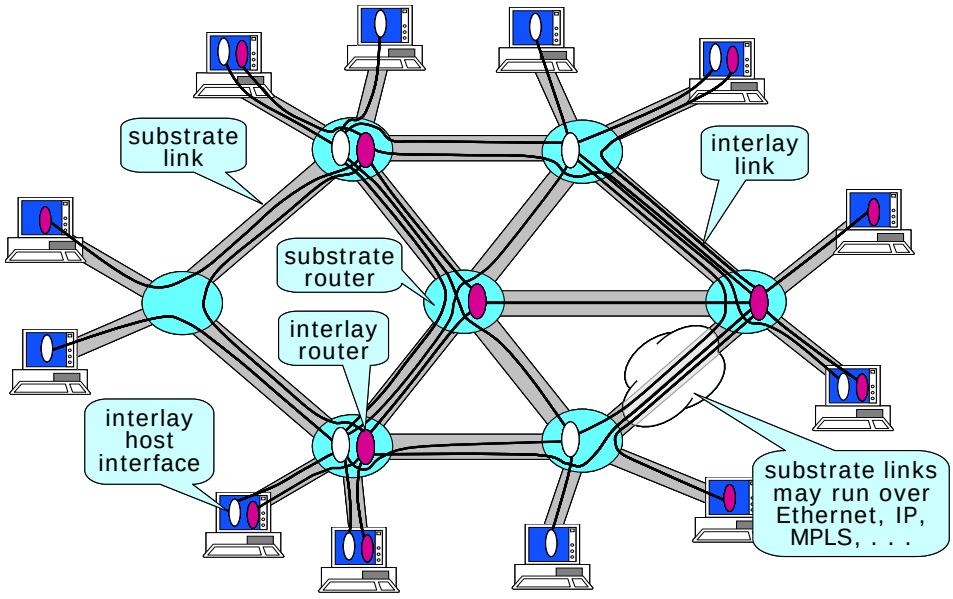

Figure 1: Diversified Internet supporting multiple interlay networks on a shared substrate

We define an interlay network to be a network of interlay links joining interlay routers to each other and to the host interfaces of end users. Once provisioned, the interlay links and routers "look" just like physical network resources. It is up to the interlay networks (not the substrate) to implement end-to-end packet delivery by selecting packet formats, forwarding mechanisms, control and management planes, etc. This freedom allows network architects to construct networks that best suit the intended application or range of applications that will operate over the interlay network. The substrate is concerned only with providing the resources needed by distinct interlay networks and ensuring that the different interlays can co-exist without interference. Our bias is to include as little as possible in the substrate. While this may mean that different interlay networks duplicate certain common services, it also means that interlay networks with simple needs are not burdened with the cost of supporting some shared service that they don't use. While this potentially increases the effort required to develop new interlay networks, we believe that this can be offset through the use of design tools and development of modular components that can be used in a variety of different interlays.

Applications run on top of interlay networks. Users may opt-in to one or many interlay networks, thus freeing application developers of the constraints imposed by any single network architecture. Developers can target their applications for the interlay best-suited to their application's requirements. While traditional applications, such as email and the web might well use an IP interlay, other applications could be designed to target interlays with different characteristics. Indeed, an application developer might even choose to develop an application in concert with a new interlay architecture, integrating application and network in ways that have not been possible with conventional network architectures. As an illustration of this, consider an interlay network designed to support distance learning. It would likely include QoS support for audio and video, but might also include specialized multicast mechanisms, with distributed audio bridging (to enable natural, high quality interaction) and audio-triggered video switching, so that the focus of video would track the current speaker. These mechanisms could be incorporated into the interlay's routers and coordinated using protocols that operate over an omni-directional multicast tree. Such an interlay could also include built-in format translators, to enable seamless participation among users with incompatible equipment. While integrating application and network in this way is something that the networking community has made a point to avoid in recent decades, there is no need to avoid it in a diversified network environment.

What's in it for stakeholders?

One lesson we draw from the current ossification of the Internet is that enabling change will require convincing stakeholders that change is in their interest. Fortunately, there is much for current stakeholders to be unhappy about in the Internet as it stands. The best-effort Internet of today is a commodity service that affords network providers limited opportunities to distinguish themselves from their competitors. A diversified Internet can provide a much richer environment for innovative organizations. This, in turn, can stimulate the development of a wide range of new services and will drive investment in the core infrastructure components needed to deliver those services to end users.

In a diversified internet, the current ISPs become Substrate Network Providers (SNPs), implementing the automated provisioning layer used to configure resources for interlay networks. SNPs can distinguish themselves from their peers through the quality of the resources they provide, the tools they provide to facilitate the development of new interlay networks and the services they provide to operating interlay networks (for example, providing hardware fault tolerance, or integrated end-user billing services).

Interlay networks would be deployed and operated by Interlay Network Providers (INPs). Interlay networks could span multiple substrate networks, and thus would acquire network resources from multiple SNPs. In this model, INPs are not required to purchase, deploy, or maintain physical infrastructure. This significantly lowers the barrier to entry and significantly accelerates deployment of new, possibly disruptive, network architectures. INPs distinguish themselves by the service they provide to applications, the platform they provide for application developers, and the value of the supported applications to the end-users.

A diversified Internet also creates new opportunities for network equipment vendors. The systems needed to implement a diversity of interlay routers within a common physical platform will afford vendors with new opportunities to develop unique products. These systems will provide configurable resources that can be used by different interlay routers to deliver their individual services. The emergence of high performance network processors [11] and configurable logic devices [12] provide the core elements needed to build such systems, but the competitiveness of equipment products will rest in how easy vendors make it for interlay networks to design and configure new types of interlay routers using those resources.

Application developers have a much richer set of choices in a diversified Internet than they have today. They may choose the most suitable interlay network from among those available, or they may design a new overlay network, tailored to their application's specific needs.

Technical challenges

To turn the vision of a diversified Internet into reality, the research community will need to address some key technical challenges. We conclude by outlining a few of these challenges, and inviting the research community to develop options for meeting them.

One of the primary challenges centers on the design of the protocols and associated mechanisms needed to enable the automated provisioning of diverse interlays in a global, multi-domain substrate network. The problem to be solved is very different from that addressed by conventional network resource reservation mechanisms. Some of these differences make it easier, others make it more complex. Unlike flow-based reservation where end-to-end resources must be reserved for individual communication sessions, the interlay provisioning system will allocate resources on a relatively coarse-grained and long-term basis. This is because interlay networks will be set up and operated to serve large numbers of users and resources will be reserved to handle the aggregate traffic from these many users over an extended period of time. These differences dramatically reduce the scaling and performance challenges associated with resource provisioning, relative to resource reservation in conventional networks. On the other hand, the nature of the interaction between interlay networks and the multi-domain substrate is intrinsically more complex, as the interlay networks must be able get information

about capabilities available through different SNPs and evaluate available options, in order to determine which SNPs can best serve its needs. They must then obtain the required resources from multiple SNPs and combine them to create their interlay.

A second challenge concerns how to ensure that different interlay networks can co-exist on a common substrate without interference. While there are well-understood mechanisms for multiplexing different virtual links onto a single physical link so as to isolate the different traffic streams, there is no such direct precedent for substrate routers that can be shared by a variety of virtual routers. While virtual machines in computer systems provide a suitable conceptual framework for thinking about these issues, the specific mechanisms used by virtual machines have limited applicability, since they do not address performance isolation and are not directly applicable to systems comprising a large number of parallel subsystems.

Perhaps the most immediate challenge facing the concept of a diversified internet is deployment. While we believe that there are great potential benefits to a diversified Internet and that critical stakeholders can be motivated to grasp the opportunity that a diversified Internet affords, we need a strategy for getting from here to there. We argue that this strategy will have to embrace the approaches used in the early deployment of IP. Specifically, we expect it to start as an overlay, progress through a government-supported experimental backbone and shift to commercial operation as network operators recognize the opportunity and as the multi-domain issues become better understood. In some ways, this process has already begun, since the ideas we are advancing here are already partially reflected in the PlanetLab testbed [13], which supports virtual overlay networks on top of a shared global infrastructure. The diversified Internet concept carries these ideas further, pushing them into the core of the network, rather than limiting them to the overlay context. We believe that the next step in carrying these ideas forward is the creation of a national testbed that leverages key ideas developed in the PlanetLab context, while refining them to enable the large-scale deployment of diverse interlay networks that can support large numbers of users and serve as a vehicle for new applications [1]. A successful large-scale demonstration of a diversified Internet will stimulate the creation of innovative applications and network services. These, in turn, will attract users and they will attract commercial operators interested in bringing these services to the larger public. We expect this to be a long process (as was the development of the Internet itself) and we acknowledge that no one can predict with any confidence how it is likely to play out. However, the problems posed by the Internet's growing ossification make it essential that we find a way to overcome the current impasse and we argue that diversification of the Internet is the most promising approach to both addressing the current problem and avoiding its recurrence in the years to come.

References

1. Peterson, L., S. Shenker, and J. Turner. *Overcoming the Internet impasse through Virtualization*. in *ACM HotNets-III*. 2004.
2. Vaughan-Nichols, S.J., *We Love IPv6, We Love IPv6 Not*, in *Enterprise IT Planet*. 2004.
3. Crowcoft, J., et al. *QoS's Downfall: At the bottom, or not at all!* in *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*. 2003. Karlsruhe, Germany.
4. Feamster, N., H. Balakrishnan, and J. Rexford. *Some Foundational Problems in Interdomain Routing*. in *ACM HotNets-III*. 2004.
5. Sekar, V., et al. *Toward a Framework for Internet Forensic Analysis*. in *ACM HotNets-III*. 2004.
6. Grice, C., *Qwest expands amid concerns of fiber glut*, in *CNET News.com*. 2001.
7. Gold, R., P. Gunningberg, and C. Tschudin. *A Virtualized Link Layer with Support for Indirection*. in *ACM SIGCOMM Workshop on Future Directions in Network Architecture*. 2004. Portland, Oregon.
8. Touch, J.D., et al. *A Virtual Internet Architecture*. in *ACM SIGCOMM Workshop on Future Directions in Network Architecture*. 2003. Karlsruhe, Germany.

9. Clark, D.D., et al. *Addressing Reality: An Architectural Response to Real-World Demands on the Evolving Internet*. in *ACM SIGCOMM Workshop on Future Directions in Network Architecture*. 2003. Karlsruhe, Germany.
10. Crowcoft, J., et al. *Plutarch: An Argument for Network Pluralism*. in *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*. 2003. Karlsruhe, Germany.
11. Intel, *IXP2800 and IXP2850 Network Processors*. 2004, Intel Corporation: Datasheet.
12. Xilinx, *Virtex-II Pro Platform FPGAs: Introduction and Overview*. 2003: DS083-1 (v3.0).
13. Chun, B., et al., *PlanetLab: An Overlay Testbed for Broad-Coverage Services.* ACM Computer Communications Review, 2003. **33**(3).