

Public Internet - The Challenges Ahead

or

Is it Possible to Extend the Internet
Architecture in an Open Way?

José Legatheaux Martins
Departamento de Informática da
FCT/UNL

Warning

This is a kind of provocative question since it assumes that the Internet architecture is already being extended in proprietary ways 😊

This presentation only elaborates on results and proposals put forward by others, which are given proper credit.

Internet Principles (Established in Simpler Times)

- IP: A common best-effort connection-less network level protocol
- A network of autonomous networks and an inter-AS routing protocol - BGP
- A global and uniform interfaces addressing scheme (32/128 bits)
- And a common naming system - DNS

These principles were defined and become well understood and mature during the 90's when the Internet and the web exploded

Today's Internet

User and Device Numbers and Link Bit rates	Applications	Required Computational Support
<p>≈ 5.000.000.000 users</p> <p>≈ 20.000.000.000 devices connected</p> <p>Access bit rates: [10 Mbps, 1 Gbps [</p> <p>Core bit rates: [10 Gbps, 400 Gbps [</p>	<ul style="list-style-type: none">+ Web Access and Search,+ All kind of Client / Server Applications,+ Electronic Commerce+ Social Networks,+ Many Million Interactive User Applications,+ Person-2-Person Interactive Applications,+ Streaming,+ Data Acquisition,+ IoT,+ Remote Control,+ ...	<ul style="list-style-type: none">+ Massive Data Centres+ Private Data Centres with Millions of Cores and Specialized Hardware for AI+ Massive Parallel Architectures and Applications+ Massive use of Caching and Machine Learning

Dealing With New Requirements

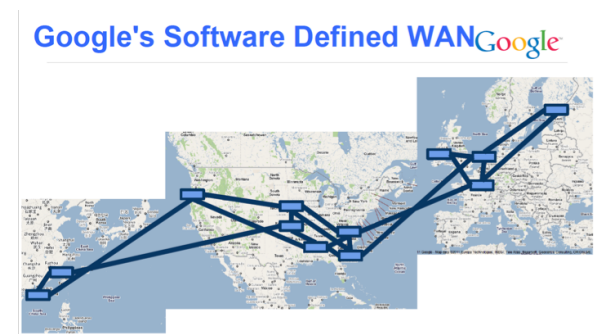

- Besides best-effort unicast IP, domains (ASs), by relying on IP and BGP, cannot provide new standardized services dependent of other domains
- Clean slate approaches failed
 - Information centric networking and variants
 - Inter-domain UDP Multicast
 - Generalized Inter-domain Publish / Subscribe
 - and several other ...

Because new global standardized Internet features require

- A smooth transition path, and
- Clear financial incentives for early adopters

But Network Innovation Made Huge Progress

- In data centres
 - “Software Defined Networking”
 - Network virtualization
 - On-demand closed subnetworks
 - Flexible adaptation to application needs
- In worldwide private data centres networks
 - Flexible management and optimization
- Inside autonomous domains
 - Data link and radio innovations (5G)
 - Traffic Management
 - Quality of service
 - Network slicing
- Higher levels
 - TLS, QUIC, BBR, ..
- However, nothing specially relevant at the IP or BGP levels



Addressing Application Needs (Outside) Public Internet Tiers

- First, servers clusters
- then, massive parallelism through cloud centres
- and reverse proxies and caches near the end users
- and, finally, private networks of worldwide data centres with Points of Presence (PoPs) inside ISPs or very close to them

Hypergiants Off-Nets Footprint

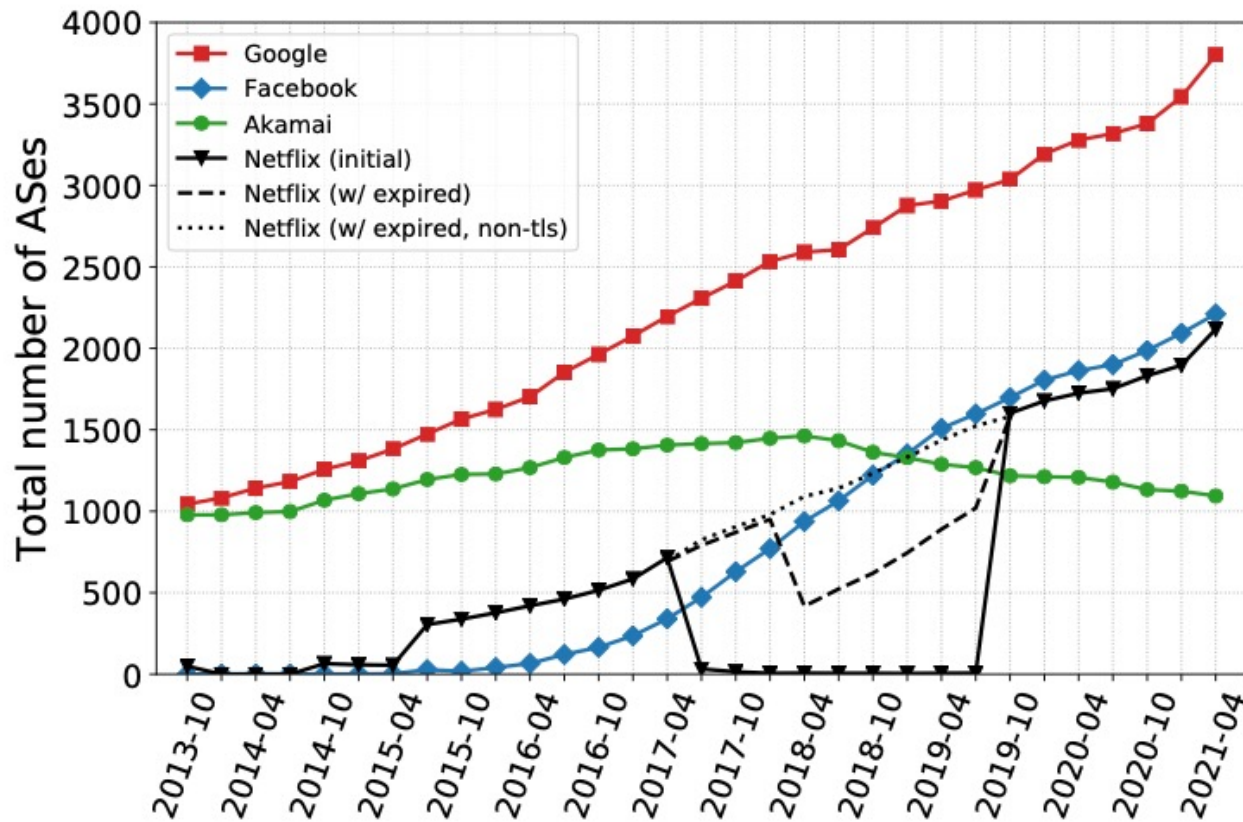


Figure 3: Off-net footprint growth for top-4 HGs over time.

Source: Petro Gigis et al., "Seven Years in the Life of Hypergiants' Off-Nets," SIGCOMM'21, August 2021, USA virtual event

Reaching Users Example



Figure 8: Percentage of a country's Internet users within the customer cones of ASes hosting Google off-net servers (April 2021).

Source: Petro Gigis et al., "Seven Years in the Life of Hypergiants' Off-Nets," SIGCOMM'21, August 2021, USA virtual event

According to TeleGeography.com(*) Reports

- Long-haul bandwidth and IP transit prices are falling 15% to 30% each year.
- 60% of inter country bandwidth is serving private data centres interconnections and this trend is increasing. 90% of the trans-Atlantic bandwidth is used by content providers
- According to some observers, only 30% of end user traffic comes from the Public Internet, i.e. from the traditional transit core.
- And a significant part of this "outside" traffic comes from close data centres of other nearby big players

(*) TeleGeography.com provides "Telecom intelligence" tells their site

A Possible Proprietary Internet Extension

- As browser and mobile applications can easily be extended by importing software frameworks
- making this application code resort to DNS over TLS redirected to a proprietary DNS-like search service
- would allow the implementation of Information Centric Network Internet extensions
- Terminals would be at low latency (1 - 5 ms) from the POPs of this new network architecture, provided by a world-wide private network
- These private "Internet" backbones may be the first to explore customers 5G speeds

Unified Internet is Being Replaced by Successful Private Networks

- This is a worrisome development since it can hamper innovation and competition
- However, it also shows that the unicast host-to-host IP packet delivery can be complemented and augmented with a set of new functionalities provided using processing in an interconnected set of Hypergiants POPs:
 - Caching,
 - Flow termination and reduced latency,
 - Load balancing,
 - Massive content diffusion,
 - Improved security,
 - Improved name to resource binding, ...

Back to an Open, Unified and Enhanced Internet Architecture

- Some recent papers propose an enhanced Internet architecture:
 - James McCauley et al, "Enabling a Permanent Revolution in Internet Architecture," SIGCOMM'19, August 2019
 - Hari Balakrishnan et al, "Revitalizing the Public Internet By Making it Extensible," ACM SIGCOMM Computer Communications Review, Volume 51, Issue 2, April 2021
- The goal is to allow the *"public Internet to offer similar capabilities - of the private networks - (and eventually much more) in a open process (in terms of process and code) and extensible manner."*

Proposed Principles

- Services should be layered on top of packet delivery (the so called 3.5 layer of interconnection of hosts and service nodes - SN)
- Services need only be implemented at network edges
- Services should be initially deployed via open-source (curated) modules designed to run on commodity servers (SN)
- Services should focus on simple processing, not general computation

Proposed Preliminary Design

- Service nodes are commodity servers deployed inside networks and data centres (the Service Node Providers) and running an open-source runtime
- Hosts are connected to service nodes by a discovery protocol (initially deployed as open source in browsers and mobile applications)
- Hosts / SN and SN / SN communications use tunnels and should be orthogonal to BGP
- Each service is instantiated as an open-source module and runs in VMs or containers executing on service nodes
- Once approved, a service open-source module can be world wide deployed using software deployment techniques

Thank You for Your Time

Slides in Stock

(Hyper) Giant Market Presence

Aplicação	Região	Empresas dominantes (faturação)
Search (*)	World except China	Google – 92% , Microsoft – 3%, ...
Search (*)	China	Alibaba + Tencent – 91% , Microsoft – 4%, ...
Browsers (*)	World	Google – 65 % , Apple, Microsoft, Firefox, ...
Devices software (*)	World	Google 41%, Microsoft 32%, Apple 23% , ...
Social Networks (*)	World except China	Facebook 76% , Twiter, Pintrest, Google, ...
Online advertising (**)	World	Google 28%, Facebook 25% , Alibab 10%, Amazon 7%, Tencent, ...
Cloud (***)	World except China	Amazon – 32%, Microsoft – 20% , Google,
Mail	World except China	Google – 20%, Microsoft – 11% , ...
CDNs, Security (##)	World	Akamai – 35%, Cloudflare 19% , Fastly, Verizon, Amazon, ...

(*) Source: <https://statcounter.com>

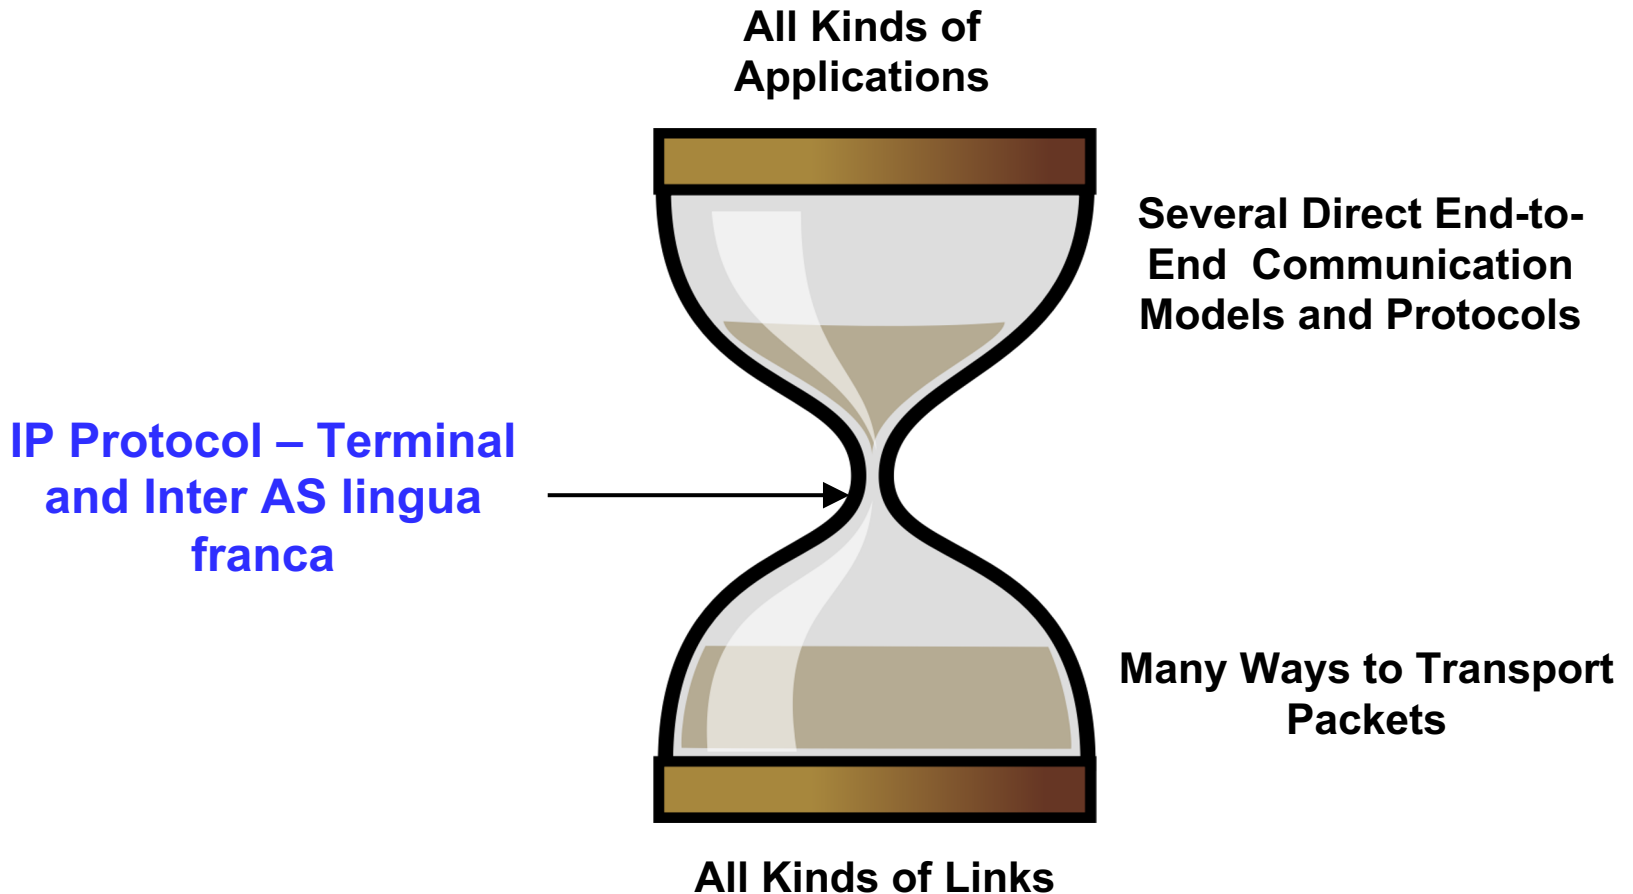
(**) Source: <https://www.emarketer.com>

(##) Source: <https://www.t4.ai>

(***) Source: <https://www.statista.com>

(#) Source: <https://w3techs.com>

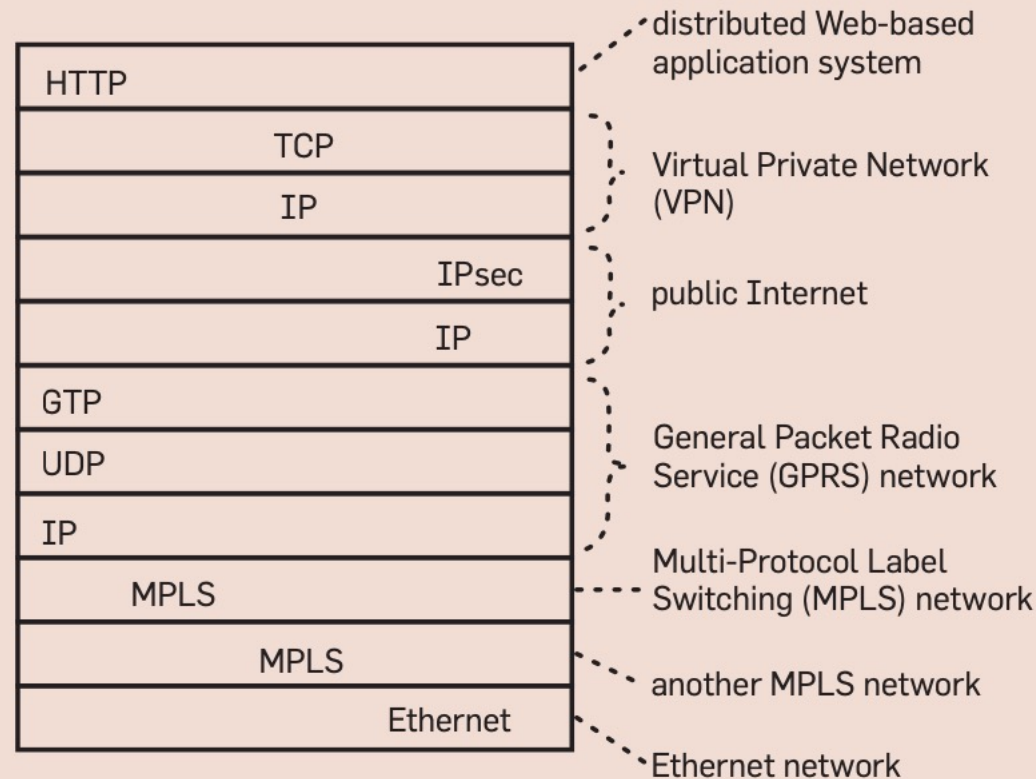
A Clear Success of "Under Specification"



Domains and Layering Hide Complexity

Figure 1. Headers of a typical packet in the AT&T backbone network.

Headers lower in the diagram are outermost in the actual packet.



Source: Pamela Zave and J. Rexford, "The Compositional Architecture of the Internet," CACM, 2019, Vol 62, N. 3, pp.s 78-87

"Wrestling with Alligators" (*)

- Some of these "alligators" are:
 - Privacy and Identity
 - Security
 - Intellectual Property
 - Uncontrolled use of AI
 - Gig-economy
 - Cyber-war
 -
- Governing and answering to "who knows, who decides, and who decides who decides,"
- Transit and access networks and publicly access networks are subject to regulation, but applications and private networks are not



(*) Source: Vint Cerf

(**) Source: Shohsana Zuboff