

Trading off bandwidth for memory in a future, information-centric Internet

Jim Roberts (IRT-SystemX)

2 October 2014



Projet porté par

Campus Paris Saclay
FONDATION DE COOPERATION SCIENTIFIQUE

Labellisation principale



Labellisations secondaires




Soutien de collectivités territoriales




Engineering the Internet

- the Internet is a “network of networks”, a set of routers interconnected by links allowing communication between hosts
- the traffic engineer's task: how much capacity (eg, link bandwidth) and how it should be shared to satisfy demand



Engineering the Internet

- the Internet is a “network of networks”, a set of routers interconnected by links allowing communication between hosts
- the traffic engineer's task: how much capacity (eg, link bandwidth) and how it should be shared to satisfy demand
- understanding a three-way relation between
 - **demand**: a succession of packet flows generated by diverse applications, typically modelled as a stochastic process
 - **capacity**: how much, but also how it is shared by different flows
 - **performance**: time to download, communication quality,... expressed in terms of probabilities



Engineering the Internet


- the Internet is a “network of networks”, a set of routers interconnected by links allowing communication between hosts
- the traffic engineer's task: how much capacity (eg, link bandwidth) and how it should be shared to satisfy demand
- understanding a three-way relation between
 - **demand**: a succession of packet flows generated by diverse applications, typically modelled as a stochastic process
 - **capacity**: how much, but also how it is shared by different flows
 - **performance**: time to download, communication quality,... expressed in terms of probabilities
- this talk is about a traffic engineering issue in a future **information-centric network**

Outline

1. engineering an information-centric Internet
2. NDN: a proposed information-centric network architecture
3. modelling cache performance
4. evaluating the memory-bandwidth tradeoff


An information-centric network

- more than 90% of Internet traffic is content retrieval
 - web pages, documents,... and videos
- network performance and costs are highly dependent on where the content is stored
 - eg, in remote servers or local cache memories




An information-centric network

- more than 90% of Internet traffic is content retrieval
 - web pages, documents,... and videos
- network performance and costs are highly dependent on where the content is stored
 - eg, in remote servers or local cache memories




An information-centric network

- more than 90% of Internet traffic is content retrieval
 - web pages, documents,... and videos
- network performance and costs are highly dependent on where the content is stored
 - eg, in remote servers or local cache memories




An information-centric network

- more than 90% of Internet traffic is content delivery
 - web pages, documents,... and videos
- network performance and costs are highly dependent on where the content is stored
 - eg, in remote servers or local cache memories



An information-centric network

- more than 90% of Internet traffic is content retrieval
 - web pages, documents, ... and videos
- network performance and costs are highly dependent on where the content is stored
 - eg, in remote servers or local cache memories
- what is the optimal memory-bandwidth tradeoff?
 - bigger caches means more hits and less network traffic
- can this tradeoff be achieved incrementally or do we need a new Internet architecture ?



Towards a new Internet architecture

- IP, the **Internet Protocol**, was designed some 40 years ago
 - some brilliant design decisions
 - connectionless packet switching, the "end-to-end principle", a layered architecture,...
 - but continuing success may be due more to Moore's law
 - increasing processing power, middle boxes, overlays, over-provisioning,...
- the Internet was not designed for present needs...
 - from 200 hosts in 1980 to more than 10^9 in 2014
 - from messaging, Telnet, FTP... to Web, social networks, video,...
- ... leading to some serious problems
 - viruses, attacks, phishing, identity theft, cyber crime,...
 - unreliable performance, difficult mobility management,...
 - a network hard to engineer, operate and troubleshoot
 - an improvised business model

Incremental change or a “clean slate” design

- since 2005, much world-wide research on new Internet architectures, deliberately ignoring the existing network



Incremental change or a “clean slate” design


- since 2005, much world-wide research on new Internet architectures, deliberately ignoring the existing network
- proposals include:
 - new architectural principles (eg, other than layering)
 - network virtualization
 - improved network management
 - ...
 - a network focused on content retrieval
- two major trends emerge from this effort
 - **software-defined networks** (virtualization, network management)
 - **information-centric networking** (content retrieval)
- the first is not clean slate, some argue the second can also be realized by incremental changes to IP

Outline

1. engineering an information-centric Internet
2. NDN: a proposed information-centric network architecture
3. modelling cache performance
4. evaluating the memory-bandwidth tradeoff


Named data networking (NDN)

- initially proposed by Jacobson et al. (Palo Alto Research Center) as content-centric networking (CCN), currently developed in NSF project called NDN
- instead of addresses, packets have names
 - users request a content chunk (eg, 4KB) with an **Interest packet**
 - network returns each requested chunk in a **Data packet**
- content chunks are cached in NDN routers




Forwarding in an NDN router

- a Forwarding Information Base (**FIB**) indicates useful output(s) for every **name** received in Interest packets
- a Pending Interest Table (**PIT**) records requesting input for each Interest until Data packet is received
- a Content Store (**CS**) temporarily caches returned Data



NDN saves bandwidth

- a content chunk cached in the CS can be downloaded directly
 - on receipt of Interest, first check content store
- simultaneous multicasting (eg, live video) realized using PIT entry
 - if content chunk name is already in PIT, do not send Interest, add new input for future Data transmission



A new ICN architecture or enhanced IP?


- Information-centric networking (eg, NDN) has compelling advantages over IP
 - saves bandwidth through deferred and simultaneous multicast
 - simplifies mobility management
 - facilitates network security by data encryption, as necessary
- on the other hand,
 - the Internet already saves bandwidth by caching: so-called Content distribution networks (CDNs) are already widely deployed
 - name-based forwarding brings severe scalability issues
 - currently there are no “killer applications” that make ICN (eg, NDN) an obvious winner
- the choice partly depends on cost-effectiveness
 - how should one engineer a network of caches?
 - what exactly is the memory-bandwidth tradeoff?

Outline

1. engineering an information-centric Internet
2. NDN: a proposed information-centric network architecture
3. modelling cache performance
4. evaluating the memory-bandwidth tradeoff

Engineering a content cache

- **demand**: a request process for content items of different popularities
- **capacity**: size and a replacement policy to keep most useful items in cache
- **performance**: the "hit rate" = proportion of requests served by cache




The Internet traffic mix

- in 2014, more than 96% of Internet traffic is content
- video counts for 60%, file sharing 20%, web, etc 20%
- video includes YouTube, Netflix, live video, webcams,...
- vast catalogues
 - $\sim 10^{11}$ web pages $\approx O(1 \text{ petabyte})$, ie, 10^{15} bytes
 - $\sim 10^6$ torrents $\approx O(1 \text{ petabyte})$
 - $\sim 10^8$ YouTubes $\approx O(1 \text{ petabyte})$
 - $\sim 10^4$ VoD items $\approx O(1 \text{ terabyte})$, ie, 10^{12} bytes
- but highly skewed demand
 - a relatively small number of highly popular items
 - popularity typically follows a generalized **Zipf law**...


Zipf popularity

- popularity is measured by request arrival rate
- measurements reveal popularity decreases as a power law:
 - request rate of n^{th} most popular object $\propto 1/n^\alpha$
 - typically, $\alpha \approx 0.8$




Zipf popularity

- popularity is measured by request arrival rate
- measurements reveal popularity decreases as a power law:
 - request rate of n^{th} most popular object $\propto 1/n^\alpha$
 - typically, $\alpha \approx 0.8$



Zipf popularity

- popularity is measured by request arrival rate
- measurements reveal popularity decreases as a power law:
 - request rate of n^{th} most popular object $\propto 1/n^\alpha$
 - typically, $\alpha \approx 0.8$



Modelling the request process: the independent reference model (IRM)

- assume a fixed catalogue of N objects
- requests arrive sequentially and any request is for object n with probability $\propto q(n)$ for $1 \leq n \leq N$
 - eg, $q(n) = 1/n^\alpha$
- ignores **time locality**, ie, assumes catalogue and popularities remain fixed over time
- a problem arises when trying to estimate $q(n)$
 - reliable statistics require long measurement periods
 - but catalogue and object popularities change
- the law for torrents avoids the **time locality** problem but ignores **space locality** since torrent trackers are global

Replacement policies

- when a cache is full, some objects must be removed to make room for new ones, eg,
 - least recently used (LRU): replace the object that has not been requested for the longest time
 - random: replace any object chosen at random
 - least frequently used (LFU): only cache the most popular objects
 - ...
- LRU is optimal among policies that cannot see into the future
- LRU appears as a reasonable compromise between complexity and performance
 - objects are indexed by a linked list that evolves at each request arrival



most recent




least recent

Calculating LRU hit rates (Che *et al.*, 2002)

- cache size C objects, popularity of object $n \propto q(n)$
- assume "independent reference model" or, equivalently, Poisson request arrivals at rate $q(n)$ for object n
- "characteristic time" T_C is time for C different objects to be requested
- assume random variable T_C is approximately deterministic, $T_C \sim t_C$
- then, hit rate for object n is $h(n) = 1 - \exp\{-q(n)t_C\}$
- now, $C = \sum_n \mathbf{1}\{\text{object } n \text{ is in cache}\}$
- taking expectations, $C = \sum_n h(n) = \sum_n (1 - \exp\{-q(n)t_C\})$
- solving numerically for t_C yields $h(n)$
- the approximation is very accurate and we know why (Fricker, Robert & Roberts, 2012)


LRU hit rate for $q(n) = 1/n^\alpha$

- strong impact of Zipf parameter α




LRU cache performance

- strong impact of Zipf parameter α
- strong impact of object population N



LRU cache performance

- strong impact of Zipf parameter α
- strong impact of object population N




Outline

1. engineering an information-centric Internet
2. NDN: a proposed information-centric network architecture
3. modelling cache performance
4. evaluating the memory-bandwidth tradeoff

In-network caching or caches at the edge only?

- in NDN, caches are distributed through the network
 - given a “**cache budget**”, the optimal distribution depends on cooperative placement policy and Interest forwarding strategy
- recent research suggests the difference between the optimum and simply caching at the edge may only be slight
 - depending on policy, strategy and supposed demand model




In-network caching or caches at the edge only?

- in NDN, caches are distributed through the network
 - given a “**cache budget**”, the optimal distribution depends on cooperative placement policy and Interest forwarding strategy
- recent research suggests the difference between the optimum and simply caching at the edge may only be slight
 - depending on policy, strategy and supposed demand model
- the cache budget is a key parameter **but this is not a given**
 - eg, if memory is **very cheap**, edge caching is clearly preferable
 - the optimal cache budget in fact depends on the memory-bandwidth tradeoff


Tradeoff at the edge

- a simple symmetric network model
- cost of edge caching = $C \times S \times k_m$
 - C = cache size, S = number of sites, k_m = unit cost of memory
- cost of bandwidth = $T \times m(C) \times k_b$
 - T = total demand (bit/s), $m(C)$ = miss rate, k_b = unit cost of bandwidth




Tradeoff at the edge


- example data
 - empirical popularity law for torrents
 - $N = 1.6$ petabyte, $T = 1$ Tb/s, $S = 100$
 - $k_m = .15$ € per GB per month, $k_b = 15$ € per Mb/s per month




Costing the tradeoff



Costing the tradeoff




Costing the tradeoff




Observations on tradeoff

- key factor is $\Gamma = Tk_b / SNk_m$ where N is catalogue size
 - $\Gamma = \text{max bandwidth cost} / \text{max storage cost}$



Observations on tradeoff

- key factor is $\Gamma = Tk_b / SNk_m$ where N is catalogue size
 - Γ = max bandwidth cost/ max storage cost
- rarely advantageous to optimize cache size
 - if $\Gamma \ll 1$, no low level cache
 - if $\Gamma \gg 1$, cache (almost) all at low level




Observations on tradeoff

- key factor is $\Gamma = Tk_b / SNk_m$ where N is catalogue size
 - Γ = max bandwidth cost/ max storage cost
- rarely advantageous to optimize cache size
 - if $\Gamma \ll 1$, no low level cache
 - if $\Gamma \gg 1$, cache (almost) all at low level
- cost trends $\Rightarrow \Gamma$ is increasing with time
 - k_m decreases by 40% each year
 - k_b decreases by 20% each year
- there is scope for optimizing position of edge caches
 - S large for small catalogue (eg, VoD, $N = O(10^{12})$)
 - S small for large catalogue (eg, torrents, etc, $N = O(10^{15})$)

Is the future Internet a network of data centers?

- edge routers become data centers, storing huge amounts of content and performing multiple data processing functions, including routing
- still scope for caching in the access network, typically for limited size catalogues like VoD
- an active area of current research



Outline

1. engineering an information-centric Internet
 2. NDN: a proposed information-centric network architecture
 3. modelling cache performance
 4. evaluating the memory-bandwidth tradeoff
- conclusions

- traffic engineering and the Internet
 - relating demand, capacity and performance
 - for capacity planning and effective resource sharing
 - a new caching dimension for an information-centric network
- the Internet has become information-centric
 - new architectures are proposed (eg, Named Data Networking)
 - though IP might still be made to evolve
- cache performance is critical in information-centric networks
 - depends on content catalogue size and popularity distributions
 - Che: a method for calculating the hit rate for LRU replacement
- optimal structure depends on the memory-bandwidth tradeoff
 - caching at the edge appears better than in-network caching
 - with scope for specialized caches in the access network