CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next-Generation Services

Wei Koong Chai, University College London
Ning Wang, University of Surrey
Ioannis Psaras and George Pavlou, University College London
Chaojiong Wang, University of Surrey
Gerardo García de Blas and Francisco Javier Ramon Salguero, Telefónica Investigación y Desarrollo S.A.U.
Lei Liang, University of Surrey
Spiros Spirou, Intracom SA Telecom Solutions
Andrzej Beben, Warsaw University of Technology
Eleftheria Hadjioannou, PrimeTel PLC

ABSTRACT

CURLING, a Content-Ubiquitous Resolution and Delivery Infrastructure for Next Generation Services, aims to enable a future content-centric Internet that will overcome the current intrinsic constraints by efficiently diffusing media content of massive scale. It entails a holistic approach, supporting content manipulation capabilities that encompass the entire content life cycle, from content publication to content resolution and, finally, to content delivery. CURLING provides to both content providers and customers high flexibility in expressing their location preferences when publishing and requesting content, respectively, thanks to the proposed *scoping* and *filtering* functions. Content manipulation operations can be driven by a variety of factors, including business relationships between ISPs, local ISP policies, and specific content provider and customer preferences. Content resolution is also natively coupled with optimized content routing techniques that enable efficient unicast and multicast-based content delivery across the global Internet.

INTRODUCTION

The original Internet model focused mainly on *connecting machines*, whereby addresses point to physical end hosts, and routing protocols compute routes to specific destination endpoints. Nowadays the Internet is primarily used for *transporting content/media*, where a high volume of both user-generated and professional digital content (webpages, movies/songs, live video streams, etc.) is delivered to users who are usually only interested in the content itself rather than the location of the content sources. Human needs along with the nature of communication technologies have transformed the Internet into a new content marketplace, generating revenue for various stakeholders. In fact, the Internet is rapidly becoming a superhighway for massive digital content dissemination.

In this context, many researchers have advocated a transition of the Internet model from *host-centric* to *content-centric*, with various architectural approaches proposed [1–7]. Many of these proposals support the key feature of *location independence*, where content consumers do not obtain explicit location information (e.g., the IP address) of the targeted content source a priori, before issuing the consumption request [1–3, 5, 7]. Nevertheless, location requirements are sometimes still demanded by both content consumers and providers. On one hand, content providers may want their content accessed only by content consumers from a specific region (known as *geo-blocking*); for example, BBC iPlayer, Amazon Video-on-Demand, Apple iTunes Store, and Sina video services. On the other hand, content consumers may prefer content originated from specific regions in the Internet; for instance, a U.S.-based shopper might only like to check the price of an item sold in Amazon online stores in North America rather than anywhere else in the world. Today, this is typically achieved through the user's explicit input in the URL (e.g., Amazon.com

 0163-6804/11/$25.00 © 2011 IEEE

and Amazon.ca) and supported by name resolution through the standard Domain Name System (DNS) [8], with the relevant requests directed toward the specific regional web server. Similar practice can be observed in multimedia-based content access (e.g., video-on-demand services), where consumers have specific requirements regarding the location/area of content sources.

In this article, we introduce a new Internet-based content manipulation infrastructure — CURLING: Content-Ubiquitous Resolution and Delivery Infrastructure for Next Generation Services. The objective is to both accurately and efficiently *hit* (or *not hit*) content objects in *specific regions/areas of the Internet*, based on user requirements and preferences. Such an approach, deployed by Internet service providers (ISPs), allows both content providers and consumers to express their location requirements when publishing/requesting content, thanks to the supported content *scoping/filtering* functions. In particular, instead of following the conventional DNS-like approach, where a content URL is translated into an explicit IP address pointing to the targeted content server, the proposed content resolution scheme is based on hop-by-hop *gossip*-like communication between dedicated content resolution server (CRS) entities residing in individual ISP networks. Content resolution operations can be driven by a variety of factors, including the business relationships among ISPs (provider/customer/peer), content consumer preferences and local ISP policies. This resolution approach is natively coupled with content delivery processes (e.g., path setup), supporting both unicast and multicast functions. Specifically, a content consumer simply issues a single *content consumption request* message (capable of carrying his/her location preferences on the content source candidate(s)), and then individual CRS entities collaboratively resolve the content identifier in the request, in a *hop-by-hop manner*, toward the desired source. Upon receiving the request, the selected content source starts transmitting the requested content to the consumer. During this content resolution operation, *multicast-like* content states are installed along the resolution path so that the content flows back immediately upon completion of the resolution process. By exploiting multicast delivery techniques, we increase the sustainability of the system in view of the expected explosion of content in the Internet.

Business Model

We first present a basic business model that involves relevant stakeholders and their business interactions. The following top-level roles can be envisaged:

- Content providers (CPs): the entities that offer content to be accessed and consumed across the Internet. These include both commercial CPs and end users who publish their content in the Internet.
- Content consumers (CCs): the entities that consume content as receivers.

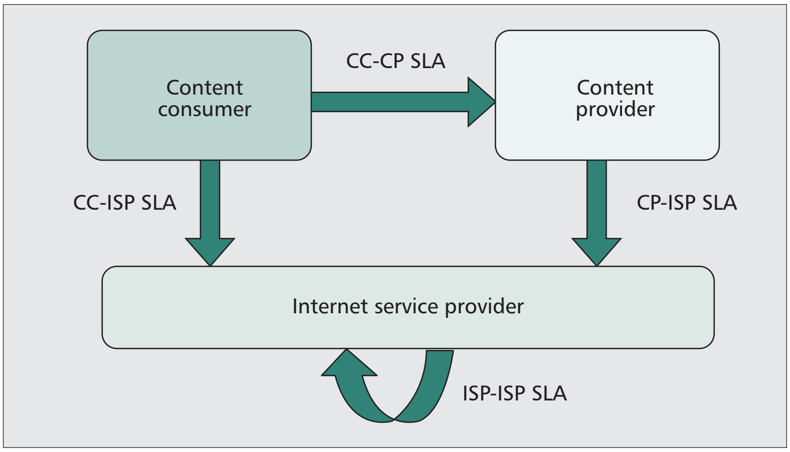

Figure 1. *Business model.*

- ISPs: Equipped with the CURLING content-aware infrastructure, ISPs are responsible for dealing with the content publication requests from CPs, and content consumption requests from consumers, and for the actual delivery of the content, possibly with quality of service (QoS) awareness.

Figure 1 shows the business interactions between individual roles. Since CPs rely on the underlying content-aware infrastructure owned by ISPs , they are expected to establish a service level agreement (SLA) involving relevant payment to the ISP for content publication services (CP-ISP SLA). In addition, since ISPs offer content searching/location and delivery services to CCs, a CC-ISP SLA can be established. Sometimes, CCs may need to pay CPs for consuming charged content (e.g., pay-per-view). This can be covered by the CC-CP SLA between the two. Finally, business contracts are also established between ISPs (ISP-ISP SLA), given a provider-customer or peering relationship between them. A low-tier ISP needs to pay its provider ISP not only for content traffic delivery, but also for *delegated* content publication/resolution services on behalf of its own customers, including directly attached CPs and consumers.

The CURLING Architecture

Our solution requires a form of aggregatable labels capable of being sequentially ordered to which we refer to as *content identifiers* (IDs). A content to be published and accessed is allocated a globally unique content ID. Multiple copies of the same content that are physically stored at different sites in the Internet share one exclusive ID.

Content manipulation operations rely on two distinct entities in the CURLING architecture:

- The content resolution server (CRS), which handles content publication requests, discovers the requested content, and supports content delivery
- The content-aware router (CaR), which collaborates with its local CRS(s) to enforce receiver-driven content delivery paths

At least one CRS entity is present in every domain for handling *local* publication requests and content consumption requests, and interact-

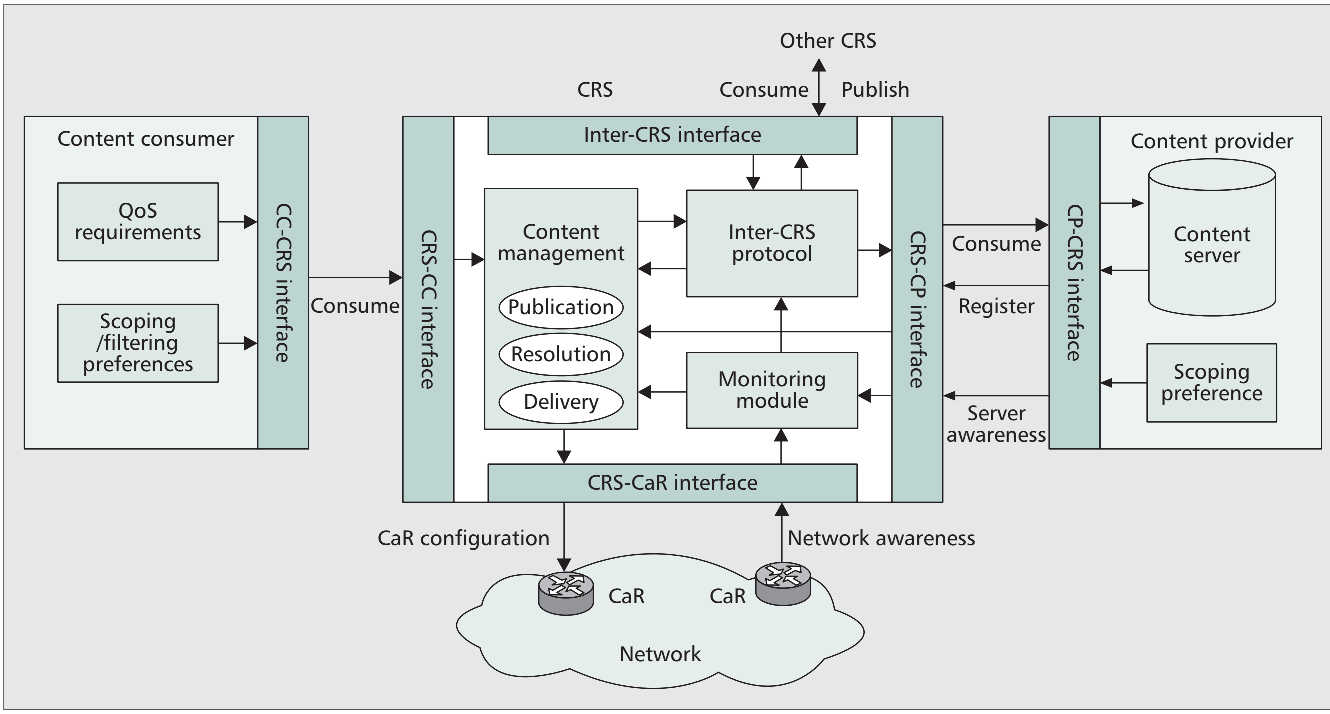

Figure 2. *High-level architecture of the hop-by-hop hierarchical content resolution approach.*

ing with other neighboring CRS entities for content publication/resolution across domains. Both CPs and consumers are configured to know their local CRS. The number of CRSs in each domain depends on performance and resilience considerations. Figure 2 depicts the functional view of the CURLING architecture. The internal structure of the CRS entity consists of three logical components. The *content management* block is responsible for dealing with requests from both CPs and CCs (via CRS-CP and CRS-CC interfaces, respectively), including content ID allocation and entry creation upon new content registrations, and also content ID lookup upon each content consumption request from a CC. A dedicated content record repository is also maintained, including not only content ID lookup information, but also ingress and egress(es) CaR(s) within the local domain for each active content session being delivered in the network. The *inter-CRS protocol component* enables the communication between neighboring CRSs for handling interdomain content publication/consumption requests. Finally, the *monitoring module* gathers necessary *near-real-time* information on content server and underlying network conditions for supporting optimized content resolution and delivery configuration operations.

CRSs communicate with other entities via specialized interfaces:

- *Inter-CRS interface* enables interaction among CRSs in neighboring domains, especially when they cooperate in content publication and searching for a requested content across domains.
- *CRS-CP interface* connects content servers owned by CPs with CRSs, and allows CPs to publish content, optionally with scoping requirements on potential CCs. This interface is also responsible for passing information on server load conditions to a CRS for enabling optimized content resolution operations.
- *CRS-CC interface* connects CC devices with the CRSs and allows consumers requesting and receiving content with scoping/filtering preferences on candidate content sources.
- *CRS-CaR interface* allows a CRS to actively configure relevant CaRs for each content session (e.g., content state maintenance). It also gathers necessary information from the underlying network that will be used for optimized content resolution processes.

A CaR is the network element that natively processes content packets according to their IDs. Generally, it is not necessary for every router in the network to be a CaR, and typically CaRs are planted at the network boundary as ingress and egress points for content delivery across ISP networks. The function of CaRs will be specified later with the description of the content delivery process.

Hop-by-Hop Hierarchical Content-Based Operations

We envisage the following three-stage content operation life cycle: *publication*, *resolution*, and *delivery*. The task of content resolution is to:

- Identify the desired content source in the Internet according to the requested content ID and optionally CC preferences
- Trigger the content transmission by the selected content server

Once the content server starts the transmission of the content upon receiving the content consumption request, the content delivery function is responsible for enforcing the actual delivery path back to the consumer.

CONTENT PUBLICATION

Content publication is the process of making content available across the Internet. It consists of two stages.

Stage 1: Content Registration — It begins with the CP notifying the local CRS that a new content is now available via a `Register` message. In the case where multiple copies of the same content are available at different locations, the CP is responsible for informing the local CRS(s) of each content server hosting that specific content copy. Upon reception of the `Register` message, the CRS registers this content by creating a new record entry in its local content management repository containing a globally unique content ID assigned to that content, and the *explicit location* of the content (i.e., IP address of the content server).

Stage 2: Content Publication Dissemination — Once the content is registered to a CRS, this CRS is responsible for publishing it globally to ensure successful discovery by potential consumers. This is achieved through the dissemination of the `Publish` message across CRSs in individual domains according to their business relationships. A `Publish` message is created by the CRS where the content is actually registered by the CP. By default, each CRS disseminates a new `Publish` message towards its counterpart in the *provider domain(s)* until it reaches a tier 1 ISP network. Each CRS receiving a new `Publish` message updates its content management repository with a new record entry containing the content ID and the *implicit location* of the content (i.e., the IP prefix associated with the neighboring domain from where the `Publish` message has been forwarded). Following this rule, each CRS effectively knows the locations of all the content within its own domain (explicitly) and those under it (its customer domains, implicitly). Peer domains, however, will not know the content records of each other.

We introduce the concept of *scoped publication* to allow publication of content only to specific areas in the Internet as designated by the CP. This feature is able to natively support regionally-restricted content broadcasting services such as BBC iPlayer and Amazon VoD that are only available within the United Kingdom and United States, respectively. We achieve this through the `INCLUDE` option embedded in the `Register/Publish` messages where the CP specifies a scoped area in the Internet (e.g., only the IP prefix associated with the local ISP network where the content is registered). A special case of scoped publication is the *wildcard mode* (denoted by an asterisk symbolizing *all* domains) for which the CP has no restrictions on the geographical location of potential consumers in the Internet.

Figure 3 illustrates different scenarios in the publication process. It depicts the *domain-level* network topology with each circle logically representing a domain containing a CRS entity. We first assume that CP **S1** registers a content item (assigned with ID **X1** by the local CRS in the stub domain **A.A.A**) to the entire Internet by issuing a `Register` message with a wildcard. Each intermediate CRS along the publication path creates a content entry for **X1** associated with the IP prefix of its customer domain from where the `Publish` message has been forwarded. For clarity, the `Publish` messages are omitted in the figure for other scenarios. Our approach also allows local domain policies to influence the publication process (e.g., domain **B.A** has the policy of NOT propagating content **X2** originated from the multi-homed domain **A.B.B** to its own provider). **S3** illustrates the *scoped* registration by only registering content **X3** to tier-2 domain **A.A** from this CP. This effectively limits the access of content **X3** to domain **A.A** and its customer domain **A.A.A**. Finally, records for different copies of the same content can also be aggregated. For instance, both **S4** and **S5** host one copy of content **X4** respectively, but the two `Publish` messages from **B.B.A** and **B.B.B** are merged at **B.B**, in which case domain **B** only records aggregated location information (**X4** → **B.B**). A content consumption request for **X4** received at **B.B** can be forwarded to either **B.B.A** or **B.B.B** based on performance conditions such as content delivery path quality or server load.

CONTENT RESOLUTION

In the content resolution process, a content consumption request issued by a CC is resolved by discovering the location of the requested content and is finally delivered to the actual content source to trigger the content transmission. A CC initiates the resolution process via a `Consume` message containing the ID of the desired content. The primary resolution procedure follows the same *provider route forwarding* rule in the publication process (i.e., the `Consume` message will be further forwarded to its provider(s) if the CRS cannot find the content entry in its local repository). If a tier-1 domain is not aware of the content location, the request is forwarded to all its neighboring tier-1 domains until the content consumption request is delivered to the identified content source. If the content is not found after the entire resolution process, an `Error` message is returned to the requesting CC indicating a resolution failure.

The *scoping functions* can also be applied in the resolution process, either embedded in the request from a CC or actively issued by a CRS for route optimization purposes during the content delivery phase. The function allows a CC to indicate preferred ISP network(s) as the source domain of the requested content. Specifically, a CC may use the `INCLUDE` option in `Consume` messages, which carry one or multiple IP prefixes to indicate from where he/she would like to receive the content.[1] Since a set of explicit IP prefixes for candidate content source is carried in the `Consume` message, the corresponding resolution process becomes straightforward: each intermediate CRS only needs to forward the request (splitting required in the presence of multiple non-adjacent IP prefixes) towards the targeted IP prefix(es) directly according to the underlying BGP routes. In case multiple interdomain routes are available towards a specific prefix, the most explicit one will be followed, as

In the content resolution process, a content consumption request issued by a CC is resolved by discovering the location of the requested content and is finally delivered to the actual content source to trigger the content transmission.

[1] It is not always required that CCs know the actual IP prefix of the domains they prefer but their local CRSs may be responsible for translating the region information (e.g., domain names) into IP prefixes through standard DNS services.

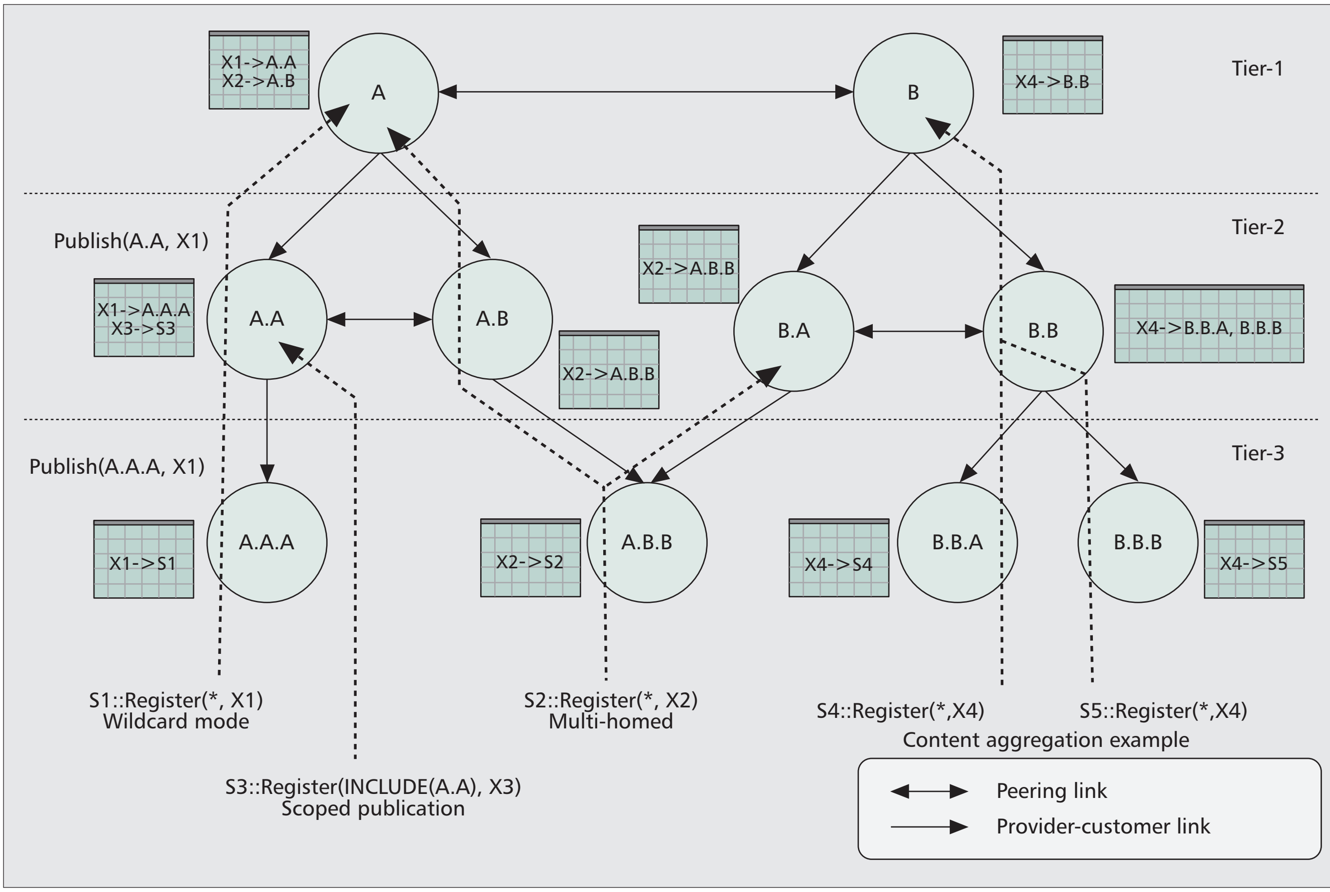

Figure 3. *Content publication process.*

is consistent with today's interdomain routing policy. In Fig. 4, CC **C1** issued a `Consume` message for content **X1** indicating its preference for content source in domain **A** or its customer domains. This `Consume` message is then explicitly forwarded towards **A** from **B** following the underlying BGP routing, but without splitting it to **C** despite that a copy of **X1** is also accessible from **C**'s customer domain **C.A**. This scoping-based content resolution path is illustrated with the solid line in the figure.

The *filtering* function in content resolution operations has complementary effect to *scoping*. Instead of specifying the preferred networks, the CC has the opportunity to indicate unwanted domains as possible sources of the desired content. The filtering function is enabled via the `EXCLUDE` option in `Consume` messages. It is important to note the fundamental difference in resolving content consumption requests with scoping and filtering functions. In contrast to the scoping scenario in which a content consumption request is explicitly routed towards the desired IP prefix(es) according to the BGP route, in the filtering case, each request is routed based on the business relationship between domains (similar to content publication operations). Consider again Fig. 4 with CC **C2** requesting **X1** with the exclusion of domain **C** and its customer domains. Since it is multihomed, the request is sent to both its providers **A.B** and **B.A** (see the dashed line in the figure). However, at the tier 1 level, domain **C** is excluded when resolving this request even though a copy of content **X1** can be found in the customer domain of **C**.

A wildcard in a content consumption request can be regarded as a special case whereby the CC does not have preferences on the geographical location of the content source. The wildcard-based resolution is illustrated in Fig. 4 via the request from consumer **C3** for content **X2** (dotted lines). We see that **B** splits the request to both **A** and **C** at the tier-1 level. Since only **A** has the record entry for **X2**, the request is resolved to **S2**.

Through these illustrations, we show that *bidirectional location independence* can be achieved in the sense that neither CCs nor CPs need to know a priori the explicit location of each other for content consumption. In particular, CCs may include *implicit* content scoping/filtering information when requesting content. The content resolution system then automatically identifies the server in the desired *area* that hosts the content. On the CP side, when content is published, scoping can be applied such that the content can only be accessed by consumers in the designated area in the Internet. We show in the following section that, thanks to the multicast-oriented content delivery mechanism, the content server is not aware of the explicit location of the active consumers of that content.

We conducted simulation experiments based on a real domain-level Internet sub-topology rooted at a tier-1 ISP network (AS7018). This four-tier network topology is extracted (with

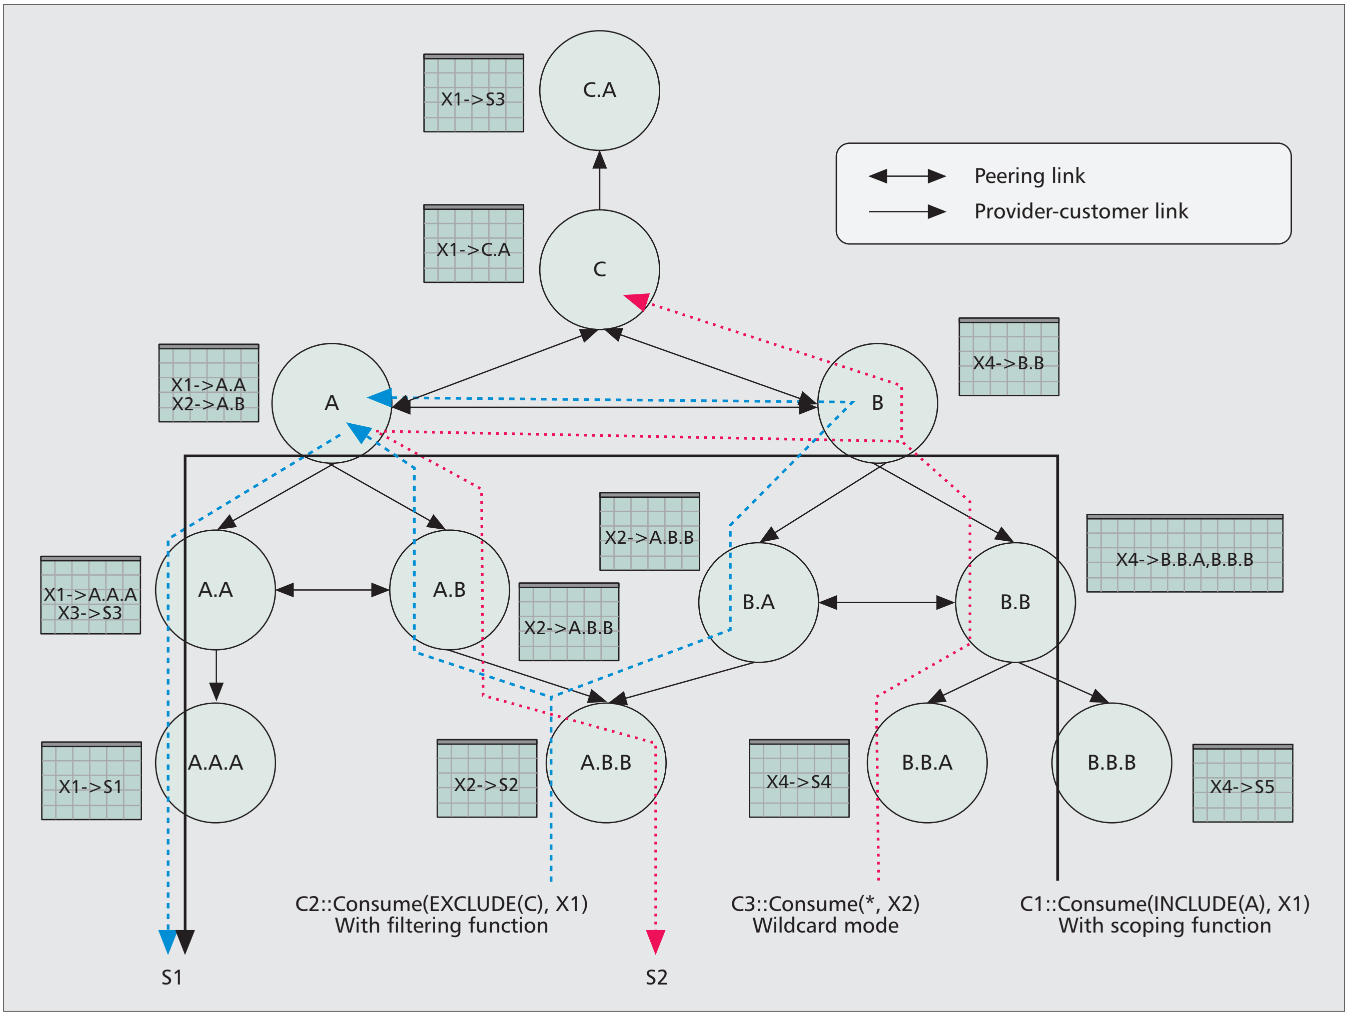

Figure 4. *Content resolution in scoping, filtering, and wildcard modes.*

aggregation) from the CAIDA dataset [9], with explicit business relationships between neighboring nodes. Content sources and consumers were randomly distributed in the domains of this topology. According to our results, the average length of the content resolution paths between individual CCs and resolved content sources is 4.4 domain-level hops (i.e., the content is on average 4.4 autonomous systems [ASs] away according to the resolution paths). This is a very good result and also consistent with the general observation that Internet interdomain sessions are of similar length based on Border Gateway Protocol (BGP) routing and the power-law interdomain Internet topology.

CONTENT DELIVERY

In CURLING, content delivery paths are enforced in a receiver-driven multicast manner that needs state maintenance based on content IDs. As described, content consumption requests are resolved through a sequence of CRSs according to either the business relationships between ISPs (in *wildcard* and filtering modes) or the BGP reachability information on the scoped source prefix (in scoping mode). In both cases, once a CRS has forwarded the content consumption request to its next hop counterpart, it needs to configure the local CaRs that will be involved in the delivery of content back from the potential server. Specifically, once a CRS receives a content consumption request from its counterpart in the previous hop domain and forwards it towards the next hop CRS, it needs to correspondingly install the content ID state at the local egress and ingress border CaRs connecting to the two neighboring domains.[2] The determination of ingress/egress CaRs for each content consumption request is purely based on the BGP reachability information across networks. Within each domain, the communication between the non-physically connected ingress and egress CaRs can be achieved either by establishing intradomain tunnels that traverse non-content-aware core IP routers, or natively through the content-centric network routing protocols [1]. Therefore, the actual domain-level content delivery path is effectively the reverse path followed by the delivery of the original content consumption request. It is worth mentioning that CRSs do not directly constitute the content delivery paths, in which case the configuration interaction between the CRS and local ingress/egress CaRs is necessary.

Let us take Fig. 5 for illustration. We assume that currently CC **C1** (attached to domain 2.1/16) is consuming live streaming content **X** from server **S** (attached to domain 1.2.1/24). The content

[2] *In case of a failed content resolution, content states temporally maintained at CaRs can be either timed out or explicitly torn down by the local CRS.*

According to our results, the content traffic traversing higher tier 1 and 2 ISPs can be reduced by 8.7 percent through peering route switching, and in particular by a substantial 28.1 percent in tier-1 ISPs. This is beneficial given that less traffic traverses tier-1 domains through relatively long paths.

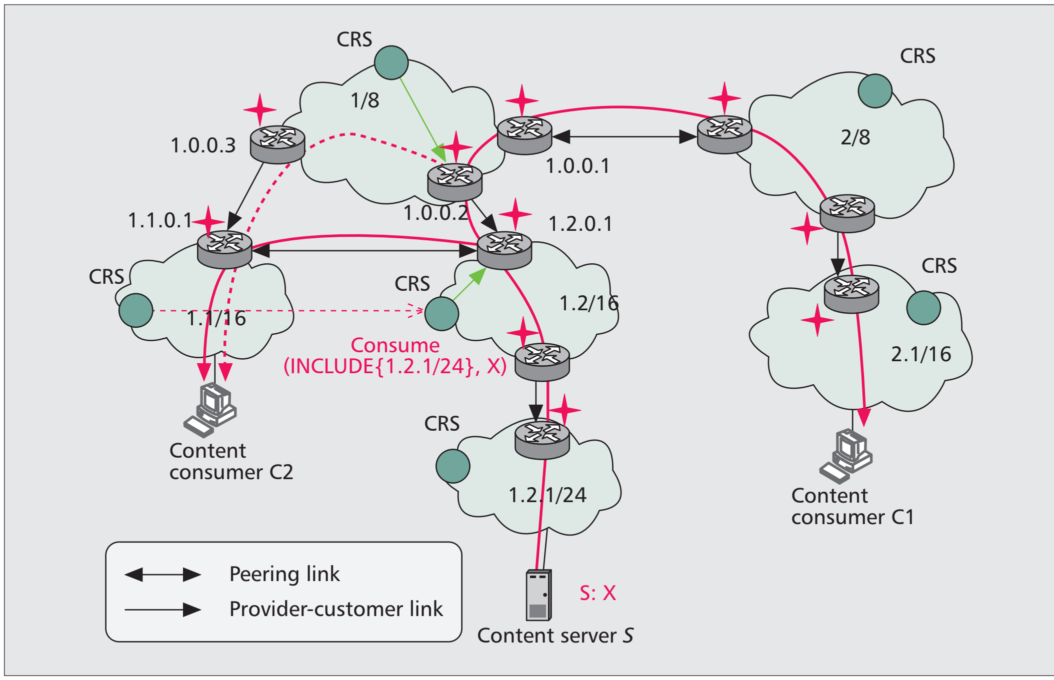

Figure 5. *Multicast-based content delivery process.*

delivery path traverses a sequence of intermediate domains, and each of the corresponding ingress/egress CaRs is associated with a star that indicates the content state maintained for content delivery. As mentioned, these states are configured by the local CRSs during the content resolution phase. Now CC **C2** (attached to domain 1.1/16) issues a consumption request for the same content. Upon receiving the request, the local CRS forwards it to its provider counterpart in domain 1/8, as it is unaware of the content source location. Since the CRS in 1/8 knows that content flow for **X** is being injected into the local network via the originally configured ingress CaR 1.0.0.2, it then updates its outgoing next-hop CaR list by adding a new egress 1.0.0.3 leading towards CC **C2**. Thus, a new branch is established from CaR 1.0.0.2 which is responsible for delivering the content back to the new consumer **C2** (the dashed line), but without any further content resolution process.

The proposed content delivery operation is also supported by a routing optimization technique for path switching from *provider* routes to *peering* routes. In the figure, once the CRS in domain 1.1/16 noticed that the content flow with source address belonging to prefix 1.2.1/24 has been injected into the local domain via ingress CaR 1.1.0.1 via the provider route, and it also knows from the local BGP routing information that there exists a peering route towards the content source, it then issues a new *scoping-based* content consumption request: `Consume(INCLUDE{1.2.1/24},X)` to the CRS in domain 1.2/16 in the peering route towards the source. Upon receiving the request, the CRS in 1.2/16 updates the local CaR 1.2.0.1 by adding a new outgoing next-hop CaR 1.1.0.1. As a result, a new branch via the peering route is established towards **C2**. Once the ingress CaR 1.1.0.1 has received the content via the interface connecting to 1.2.0.1, it prunes the old branch via the provider route (the dashed line). This content delivery path optimization effectively reduces content traffic within top-tier ISP networks and also possibly reduces the content delivery cost for customer domains. Of course, this operation is not necessary if a CRS is allowed to send content consumption requests to its peering counterparts (in addition to the provider direction) during the resolution phase. However, such an option will incur unnecessarily higher communication overhead in disseminating content consumption requests, especially when the peering route does not lead to any source that holds the requested content.

We are also interested in the actual benefit from such inter-domain routing optimization techniques for cost-efficient content delivery across the Internet, especially from the view point of tier-1 ISPs that constitute the Internet core. We used the same domain-level topology as previously described for evaluating the corresponding performance. According to our results, the content traffic (in terms of the number of media sessions) traversing higher tier 1 and 2 ISPs can be reduced by 8.7 percent through peering route switching, and in particular by a substantial 28.1 percent in tier-1 ISPs. This is beneficial given that less traffic traverses tier-1 domains through relatively long paths.

Discussion on Scalability

The domain-level hop-by-hop content resolution strategy presented follows a similar style to that proposed in [3]. However, through the new scoping and filtering functions, our architecture provides the necessary flexibility for both CPs and CCs to publish/request content at/from their desired area(s). The scalability of the system, thus, is dependent on the amount and popularity of content in each CRS, with the most *vulnerable* CRSs being those that maintain the highest number of popular content entries. This is in contrast with intuition that the most strained CRSs will be the tier-1 ones, since content publications and requests may often not reach the tier-1 level based on our approach. Again, we take BBC iPlayer as an example where both the content publication and consumption requests are restricted to IP prefixes only from the United Kingdom. In addition to that, local domain policies may also override the default publication route (see **S2** in Fig. 3).

Business incentives also present a natural load distribution mechanism for our system. We foresee ISPs charging higher publication tariffs for popular content published at higher tier domains (with tier-1 domains being the most expensive) which can be potentially accessed by a higher number of consumers. This mechanism forms a business tussle from the CPs' point of view when provision of wider access is coupled with higher monetary cost. Instead, a CP may strategically replicate content to multiple lower-tier regional ISPs (by applying scoping functions there) in which they believe their content will be locally popular.

Finally, our system allows aggregation in two ways. First, as illustrated in Fig. 3 for **S4** and **S5**, the record for the same content can be merged during the publication process among CRSs. Second, a block of sequential content IDs should be allocated to *interrelated* content so that they can be published in one single process. This rule exploits the fact that a specific CP usually offers content with some relationship with each other (e.g., all episodes of a television series). This allows for coarser granularity in the publication process whereby the CP can send only one `Publish` message to publish all the related content. The local CRS still assigns a unique content ID for each content, but the IDs are sequentially connected. The onwards publication process will only involve the entire block of IDs rather than the individual content records, especially towards high-tier ISPs.

Conclusions

In this article, we present CURLING, a new content-based Internet architecture that supports content publication, resolution and delivery. Content providers can cost-efficiently publish content based on its expected popularity in different regions by scoping its publication while content consumers can express their location preferences by scoping/filtering their content consumption requests. The processes are devised so that both sides are oblivious of their counterpart's location, resulting in a bi-directional location independence paradigm, but without sacrificing content providers' and consumers' location preferences. The proposed route optimization mechanism enhances the efficiency of content delivery by using content states established during the resolution process and initiating content delivery path switching; it mimics as such interdomain multicast delivery, which has seen very slow deployment until now.

Acknowledgments

This work was undertaken under the Information Society Technologies (IST) COMET project, which is partially funded by the Commission of the European Union. We would also like to thank our project partners who have implicitly contributed to the ideas presented here.

References

[1] V. Jacobson *et al.*, "Networking Named Content," *ACM CoNEXT '09*, 2009, pp. 1–12.

[2] P. Jokela *et al.*, "LIPSIN: Line Speed Publish/Subscribe Inter-networking," *Proc. ACM SIGCOMM '09*, Barcelona, Spain, Aug. 2009.

[3] T. Koponen *et al.*, "A Data-Oriented (and Beyond) Network Architecture," *Proc. ACM SIGCOMM '07*, Kyoto, Japan, Aug. 2007.

[4] P. Francis and R. Gummadi, "IPNL: A NAT-Extended Internet Architecture," *Proc. ACM SIGCOMM '01*, San Diego, CA, Aug. 2001, pp. 69–80.

[5] I. Stoica *et al.*, "Internet Indirection Infrastructure," *Proc. ACM SIGCOMM '02*, Pittsburgh, PA, Aug. 2002, pp. 73–86.

[6] D. Clark *et al.*, "FARA: Reorganizing the Addressing Architecture," *Proc. ACM SIGCOMM, FDNA Wksp.*, Aug. 2003.

[7] M. Caesar *et al.*, "Routing on Flat Labels," *Proc. ACM SIGCOMM '06*, Pisa, Italy, Sept. 2006, pp. 363–74.

[8] P. Mockapetris, "Domain Names — Concepts and Facilities," IETF RFC 1034, Nov. 1987.

[9] CAIDA dataset; http://www.caida.org/research/topology/#Datasets.

Biographies

Wei Koong Chai (wchai@ee.ucl.ac.uk) is a research fellow in the Department of Electronic and Electrical Engineering, University College London (UCL), United Kingdom. He received a B.Eng. degree in electrical engineering from Universiti Teknologi Malaysia in 2000 and an M.Sc. (Distinction) in communications, networks, and software, and a Ph.D. in electronic engineering, both from the University of Surrey, United Kingdom, in 2002 and 2008, respectively. His research interests include content-centric networks, QoS and service differentiation, resource optimization, and traffic engineering.

Ning Wang (n.wang@surrey.ac.uk) is a lecturer at the University of Surrey. He received his B.Eng. (Honors) degree from Changchun University of Science and Technology, P.R. China, in 1996, his M.Eng. degree from Nanyang University, Singapore, in 2000, and his Ph.D. degree from the University of Surrey in 2004. His research interests include Internet content delivery techniques, end-to-end QoS provisioning, and traffic engineering.

Ioannis Psaras (i.psaras@ee.ucl.ac.uk) is a research fellow in the Department of Electronic and Electrical Engineering, UCL. He received a diploma in electrical and computer engineering, and a Ph.D. degree from Democritus University of Thrace, Greece, in 2004 and 2008, respectively. He won the Ericsson Award of Excellence in Telecommunications in 2004. He has worked at DoCoMo Eurolabs and Ericsson Eurolab. His research interests include congestion control, delay-tolerant networks, and user- and content-centric networks.

George Pavlou (g.pavlou@ee.ucl.ac.uk) holds the Chair of Communication Networks at the Department of Electronic and Electrical Engineering, UCL. Over the last 20 years he has undertaken and directed research in networking, network management, and service engineering, having extensively published in those areas. He has contributed to ISO, ITU-T, IETF, and TMF standardization activities, and has been instrumental in a number of key European and U.K. projects that produced significant results. He is currently the technical manager of COMET.

Chaojiong Wang (c.wang@surrey.ac.uk) holds a B.Sc. in computing and information technology from the University of Surrey, and an M.Sc. in computer science from Oxford University. He is a Ph.D. student in electrical engineering at the Centre for Communication Systems Research (CCSR), University of Surrey.

The proposed route optimization mechanism enhances the efficiency of content delivery by using content states established during the resolution process and initiating content delivery path switching; it mimics as such inter-domain multicast delivery which that has seen very slow deployment until now.

GERARDO GARCIA DE BLAS (ggdb@tid.es) is a research engineer in the IP Network Technologies group in Telefónica I+D. He holds a Master's in telecommunications engineering from the Technical University of Madrid, Spain. Since 2002 he has worked in Telefónica I+D, with his main focus on network evolution, network planning, and traffic analysis. He is currently coordinating the architecture definition in the EU FP7 project COMET.

FRANCISCO JAVIER RAMON-SALGUERO (fjrs@tid.es) received his Master's in telecommunications engineering from Málaga University, Spain, in 2000, and his Master's in economics at UNED, Spain, in 2006. With Telefónica I+D since 2000, currently he heads the IP Network Technologies group, coordinating research on long-term evolution of Internet technologies. Since 2010 he has led the FP7-COMET consortium, focused on providing to the future Internet a unified approach to content location, access, and delivery with the appropriate network resources.

LEI LIANG (l.liang@surrey.ac.uk) is a research fellow at CCSR at the University of Surrey, from which he received his B.E. degree in 1998, M.S. degree in 2001, and Ph.D. degree in 2005. With strong knowledge in multiparty communications, network performance measurement and analysis, IP networking QoS, IP over satellite networks, IP multicast over satellite network, and IP network security, he has been heavily involved in 14 EU and U.K. projects since 2001.

SPIROS SPIROU (spis@intracom.com) is a senior engineer at Intracom Telecom, Greece, working on IPTV, content-aware networking, and network management. Previously, he was a research associate on data acquisition and grid computing at NCSR Demokritos, Greece, and CERN, Switzerland. He holds a B.Sc. (1997) in computer engineering, a postgraduate diploma (1998) in mathematics, and an M.Sc. (2000) in neural networks. He chaired the European Future Internet Socio-Economics Task Force.

ANDRZEJ BEBEN (abeben@tele.pw.edu.pl) received M.Sc. and Ph.D. degrees in telecommunications from Warsaw University of Technology (WUT), Poland, in 1998 and 2001, respectively. Since 2001 he has been an assistant professor with the Institute of Telecommunications at WUT, where he is a member of the Telecommunication Network Technologies research group. His research areas include IP networks (fixed and wireless), content-aware networks, traffic engineering, simulation techniques, measurement methods, and testbeds.

ELEFTHERIA HADJIOANNOU (eleftheriah@prime-tel.com) received a Dipl.-Ing. degree from the Electrical Engineering and Computer Science Department, Aristotle University of Thessaloniki, and an M.B.A. from the University of Macedonia, Greece. Previously, she was a research associate at the Information System Laboratory (ISlab), University of Macedonia, where she was involved in a number of projects relevant to eGovernment and eParticipation. Currently, she is a member of the R&D Department at Primetel PLC, an alternative telecommunication provider in Cyprus.