Content-oriented Networking (CON)

Ted "Taekyoung" Kwon
SNU

Mar-8-2011
SNU

outline

- IP networking vs. Content networking
- CANA (Content-aware Network Architecture)
- SCAN (Scalable Content Routing for Content-Aware Networking)

Why Internet has problems?

- Original motivation for Internet is to share computing resources
  - Remote login, file transfer
- Hence the Internet communication model is static host-to-host conversations
- It is ossifying
  - Mobility, security, accountability,,…

Why content-oriented networking (CON)?

- Internet traffic is already content-oriented
  - CDN, Edge Caching, multimedia, P2P, web, …
- Users/applications care “what to receive”
  - They don’t care “from where”
- Storage cost is getting cheaper sharply
  - Compared to networking cost
- Other advantages of CON

IP networking

- Lookup-by-name
  - DNS: Indirection from name to locator
    - DNS is extendible and highly available
    - Distributed design, caching
  - Host/link availability concern
- Delivery inefficiency concern
- Locators can be aggregated
  - Network prefix
    - Currently 350k+
  - Routing scalability is better than CON

Content-oriented networking (CON)

- Route-by-name
  - No indirection, better availability
  - Content name (or ID) is a routing entry
  - Huge scalability concern
- Global-scale and systematic CON may not be feasible
  - NDN, TRIAD
    - Some aggregation by using URL-like names
  - DONA, PSIRP
    - Flat names for persistency
- Better delivery performance
  - Exploiting multiple sources, multiple paths/interfaces
  - Potential opportunities for data explosion

outline

- IP networking vs. Content networking
- CANA (Content-aware networking architecture)
- SCAN (Scalable Content Routing for Content-Aware Networking)

What is an IP address?

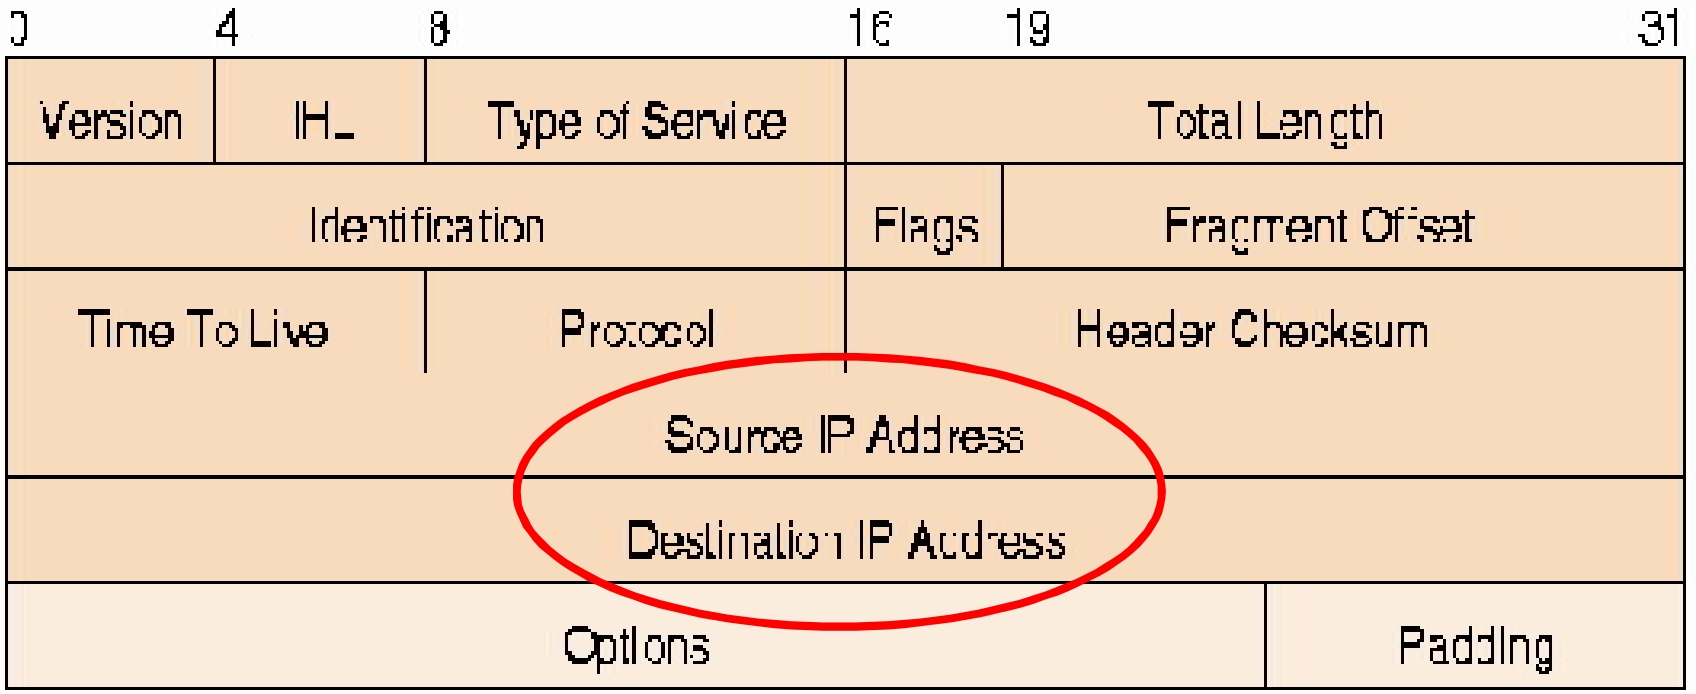

IP address

- An IP address originally indicates the endpoint
  - End-to-end principle
  - Serves as both locator and identifier

- Current role of IP address
  - Not endpoint
    - NAT, tunneling, overlay,…
  - Not identifier
    - Mobility, multi-homing,…

Then what should be an IP address?

- Just locator
  - Not identifier
- Locater of next transit point
  - NAT, tunneling,…
  - Some agents
    - E.g. mobility agent in mobile IP solutions
- Transit-by-transit
    - Not end-to-end

Wait, where is the endpoint identifier?

- How about using some other identifier?
- General identifier requirements
  - Unique
  - Routable/locatable
  - Persistent
    - Location-independent
- We choose uniqueness and routability

How about Content identifier (CID)?

- CID will fill the fading role of the IP address
  - Host-independent endpoint identifier
- Globally routable and Unique
- Domain name (or public IP address) + port number (or its hint)
  - Static content, e.g. http://www.nytimes.com/logo.jpg
  - Dynamic content, e.g. 20.30.40.50:4000

Content-aware Network Architecture (CANA)

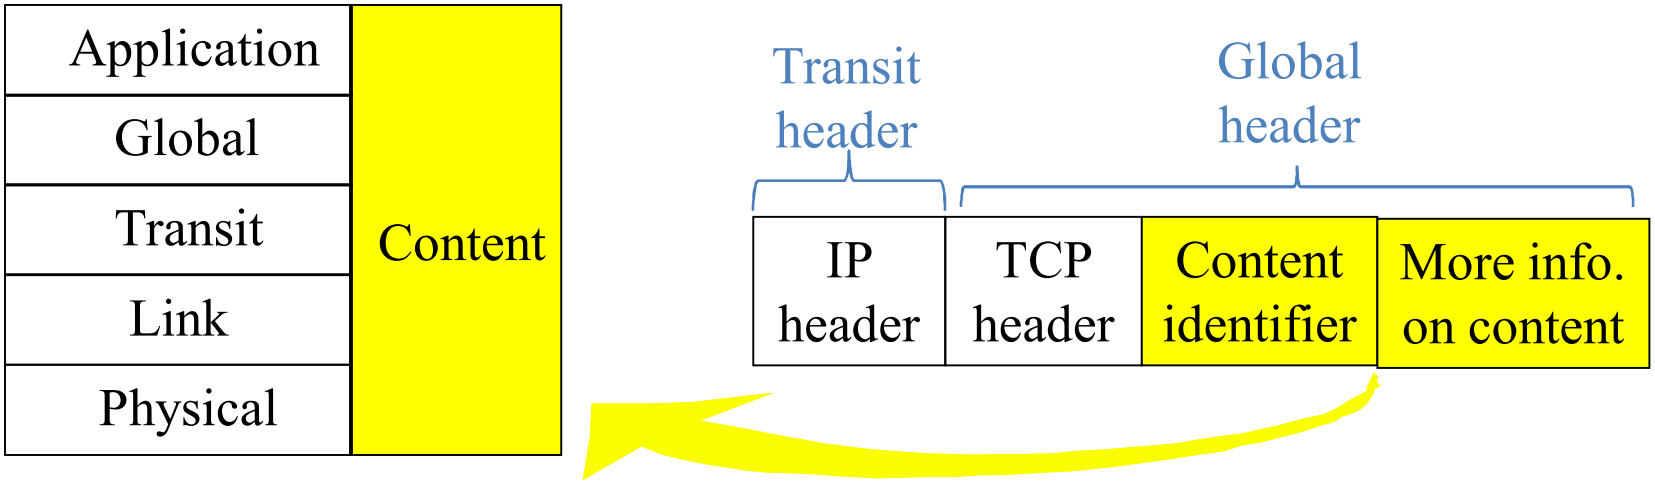

- Network layer is renamed as transit layer
- Transport layer is extended to global layer
  - CID is added
- CID is locatable and unique
- Additional content info (e.g. bit rate, chunk index) helps other layers
  - Deep packet inspection is assumed for other layers

CANA: Host side

- New model for IP subnet:
  - solicitor vs. agent
    - An access router becomes an agent
    - Solicitor and its agent communicate in a content-oriented fashion
  - An agent contacts DNS
    - Solicitors cannot
  - solicitor cannot contact server directly

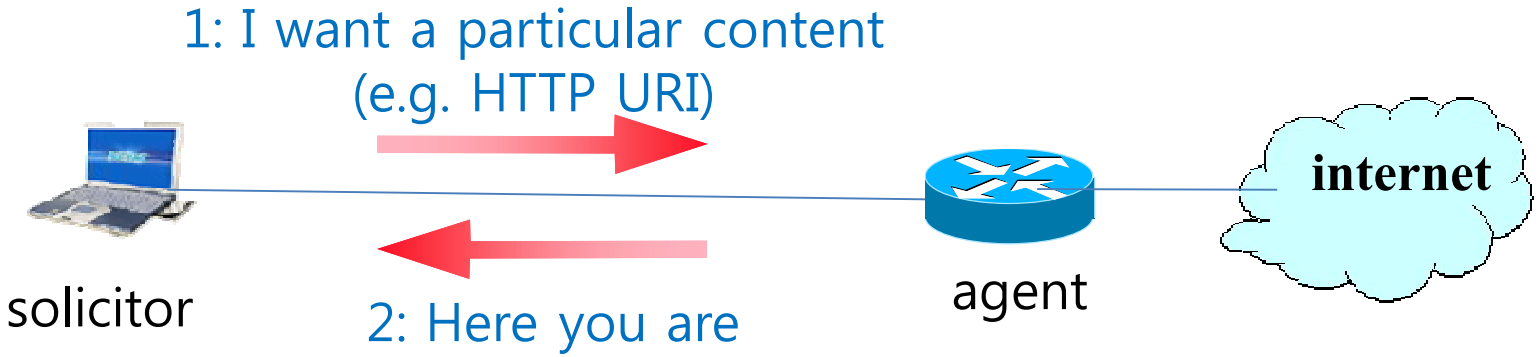

Other aspects of agents

- Flash crowd can be dealt with by caching content at agents
- NAT does not matter
- First line of defense
  - Supervise users by looking at content requests
  - Better accountability

CANA: Publisher side

- Registers its hostname with the DNS
  - Agent's IP address
- Publisher and agent will communicate in a content-oriented fashion

* Assume that publication is already done

Intra-domain

- Publisher's agent will be contacted by the solicitor's agent
- Publisher's agent will receive the content from the publisher
  - Will relay the content to the host via the host's agent
- Agents can cache contents

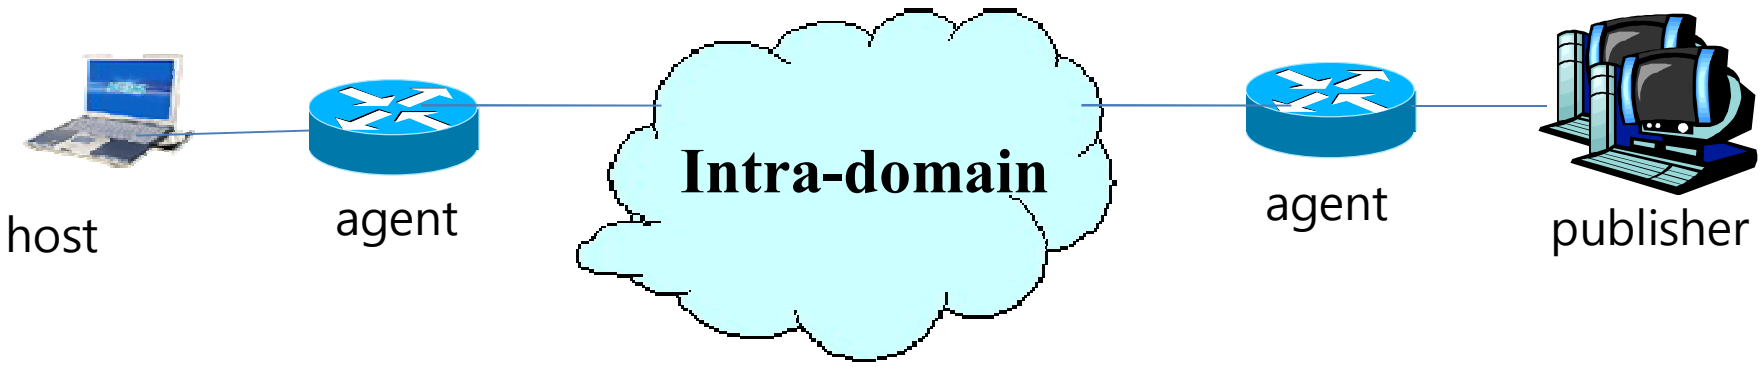

Inter-domain (Next stage)

- Gateway A requests the content to gateway B
- Gateway B will get the content from agent of publisher
  - Then relay the content to gateway A
- Gateway A will relay the content to the agent of the host
- Gateways can cache contents

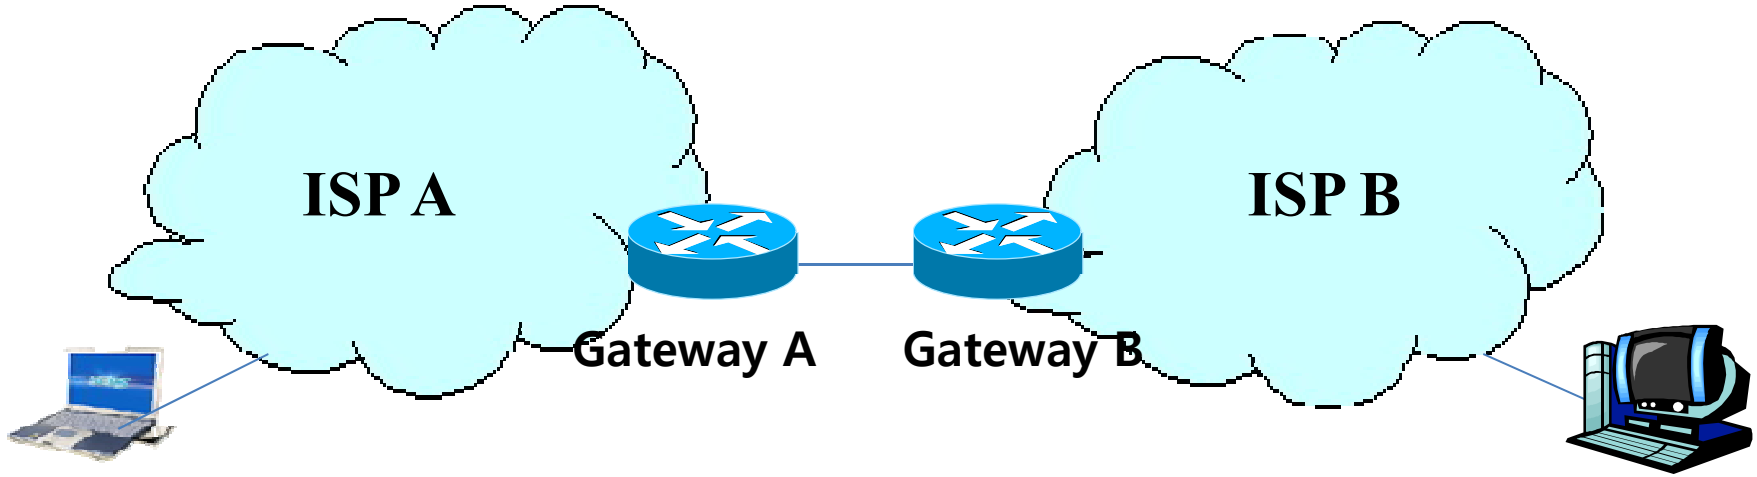

Content-aware routers (CARs)

- Legacy routers look at IP address in transit header
- CARs look at CID in global header as well
- CARs can participate in content relaying
  - CARs can cache contents

CANA operations: Content Request Message

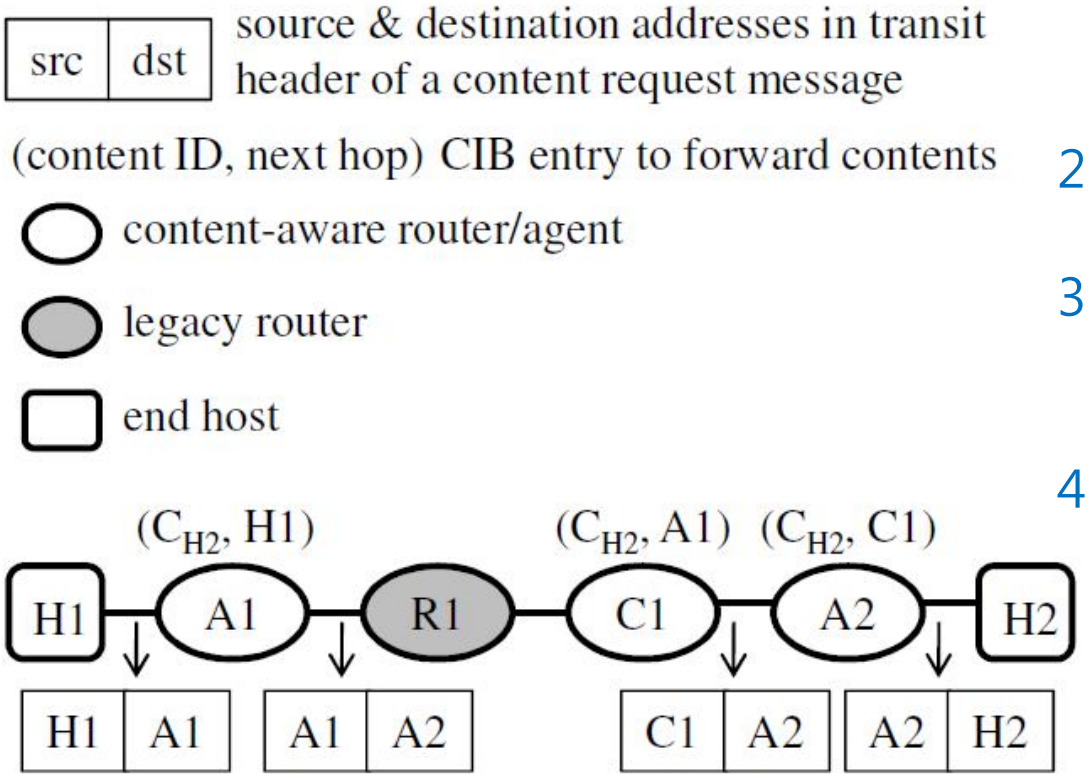

1. H1 sends a content request message to A1, its src:dst is H1:A1 (the content C_{H2} belongs to H2)
2. A1 makes a CIB entry (C_{H2}, H1) to deliver content data
3. The content request message now has A1:A2 as src:dst IP addresses
4. C1 makes a CIB entry (C_{H2}, A1) to deliver content data

.....

- As content request message traverses, a content info base (CIB) entry is set up backwards to relay content data

outline

- IP networking vs. Content networking
- CANA (Content-aware Network Architecture)
- SCAN (Scalable Content Routing for Content-Aware Networking)

IP vs. Content networking

- Inefficiency in TCP/IP networking
  - Cannot know closer copies of the content
  - Don't do parallel transmissions
- Content networking
  - Scalability, reachability issues
- Our Solution: A hybrid approach
  - IP routing: default routing for reachability
  - Content routing: opportunistic routing for efficiency (closer & multiple copies)

SCAN Overview

- Content routers (C-routers) do IP and content routing
  - Each content has a content identifier (CID)
  - Local content table (LCT): cached content files
  - Content routing table (CRT): CIDs of the content files in neighbor C-routers
  - SCAN propagates the information of the contents by Bloom Filter (BF) to mitigate the routing scalability issue

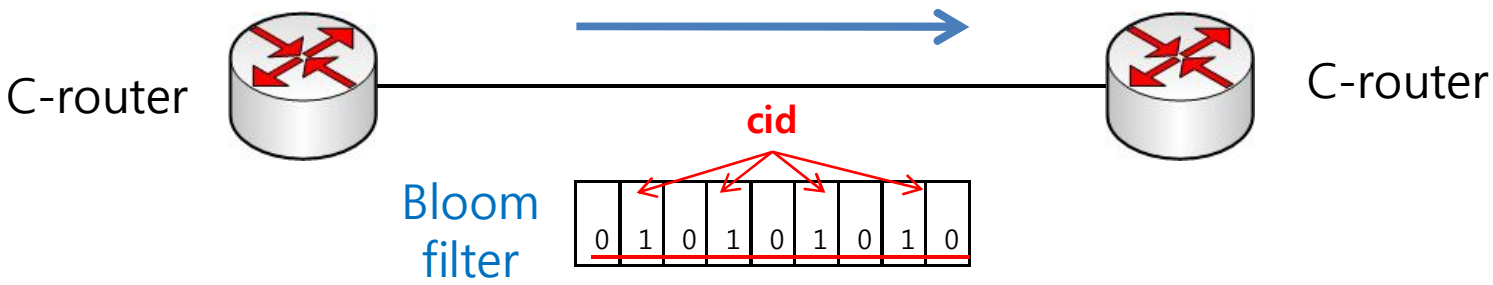

Bloom Filter (BF) issues

- More bits of the BF may be set to 1
  - As the number of content files increases
- C-router will decay the bits of a BF probabilistically before exchanging content routing table (CRT) info
  - E.g. if decaying prob. is 0.5, around the half of the bits 1 will be randomly set to 0

SCAN Operations (1/2)

- Content Routing

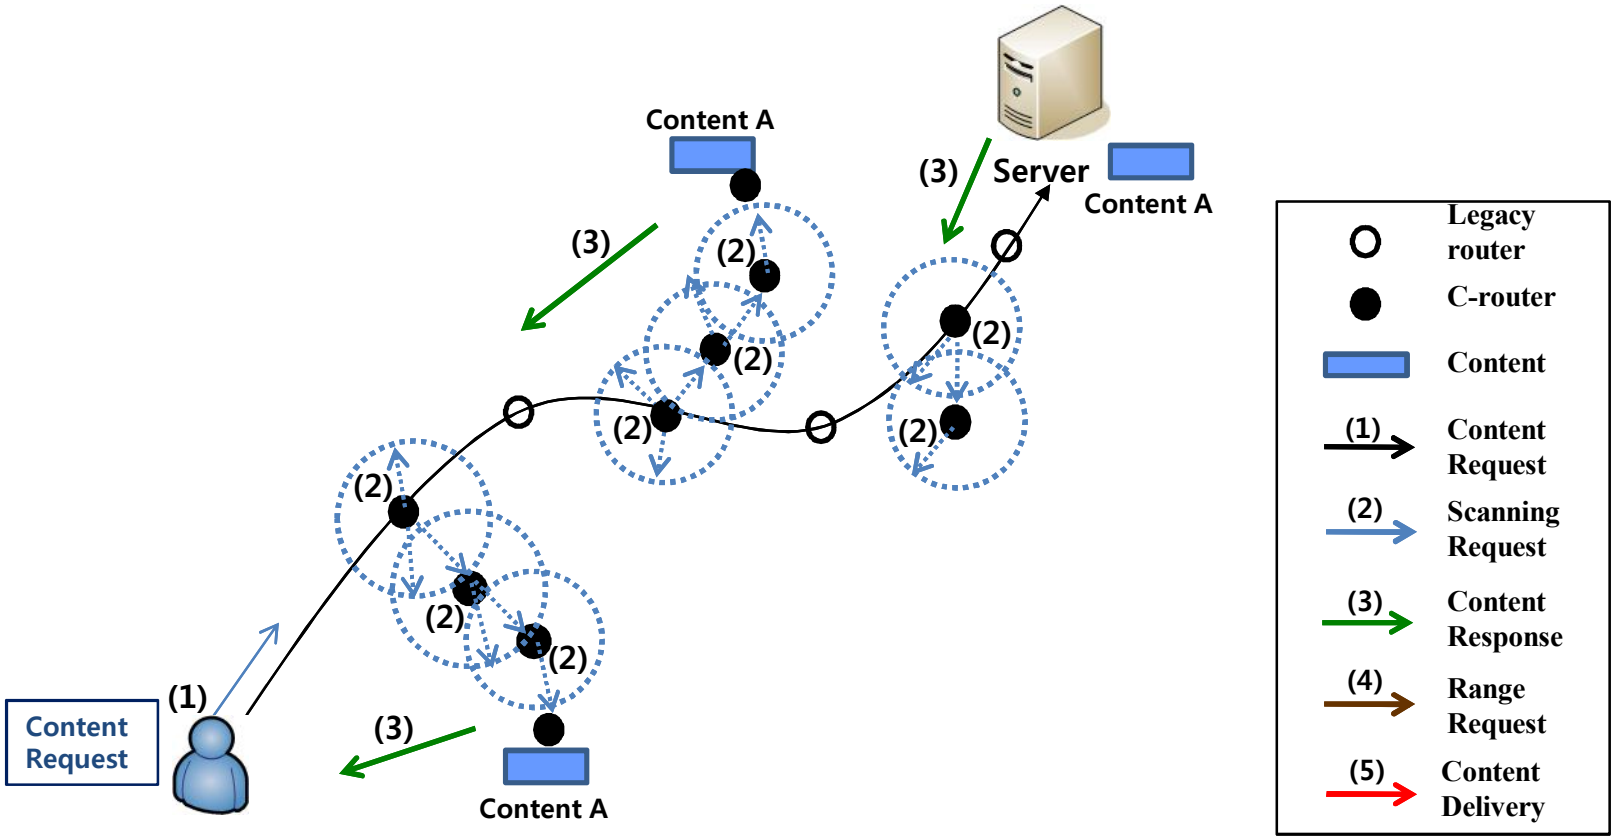

SCAN Operations (2/2)

- Content Delivery (an example)

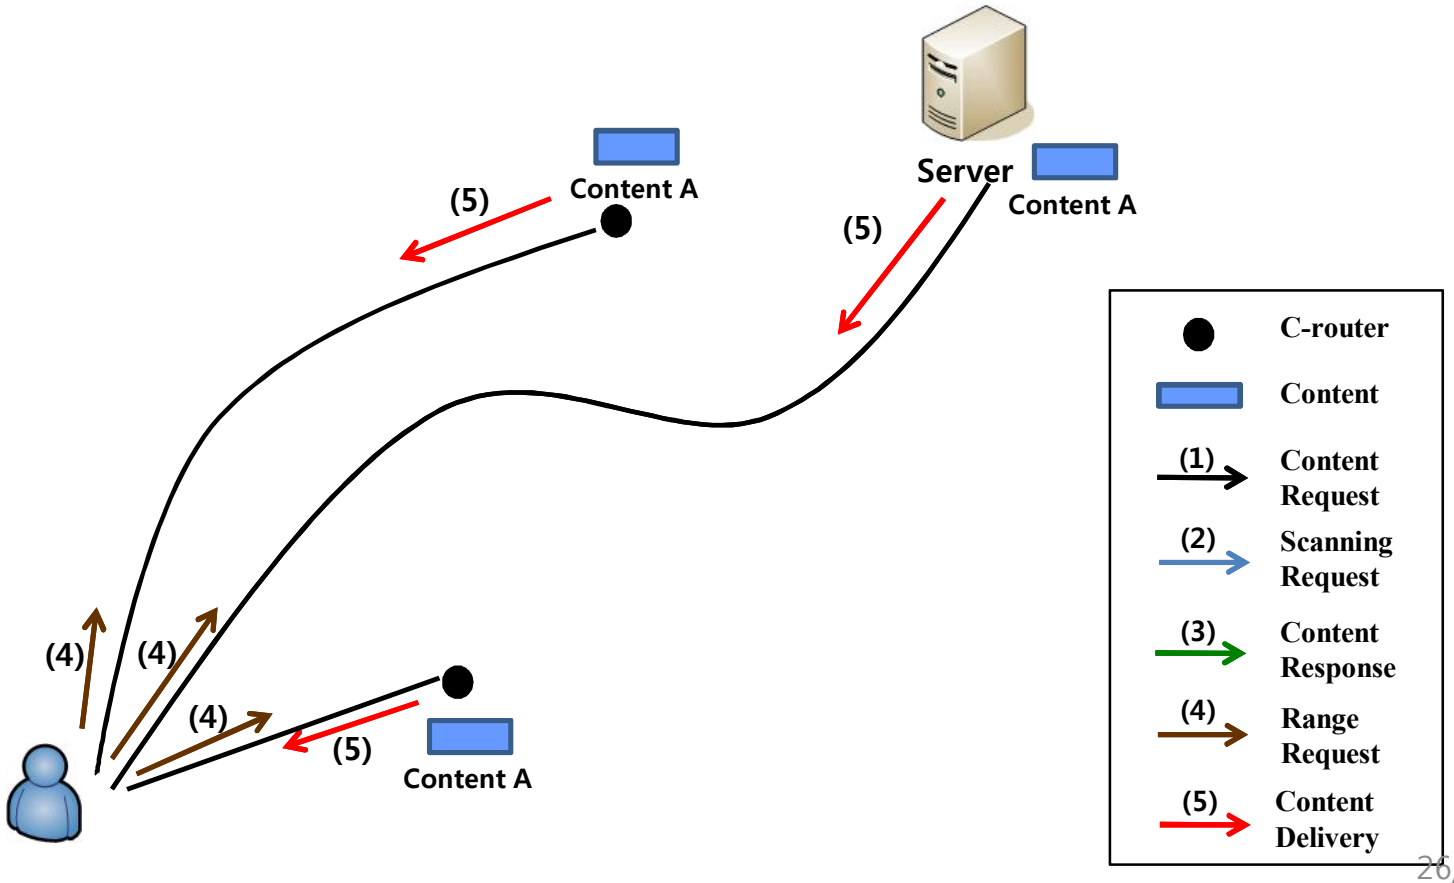

Forwarding Decision

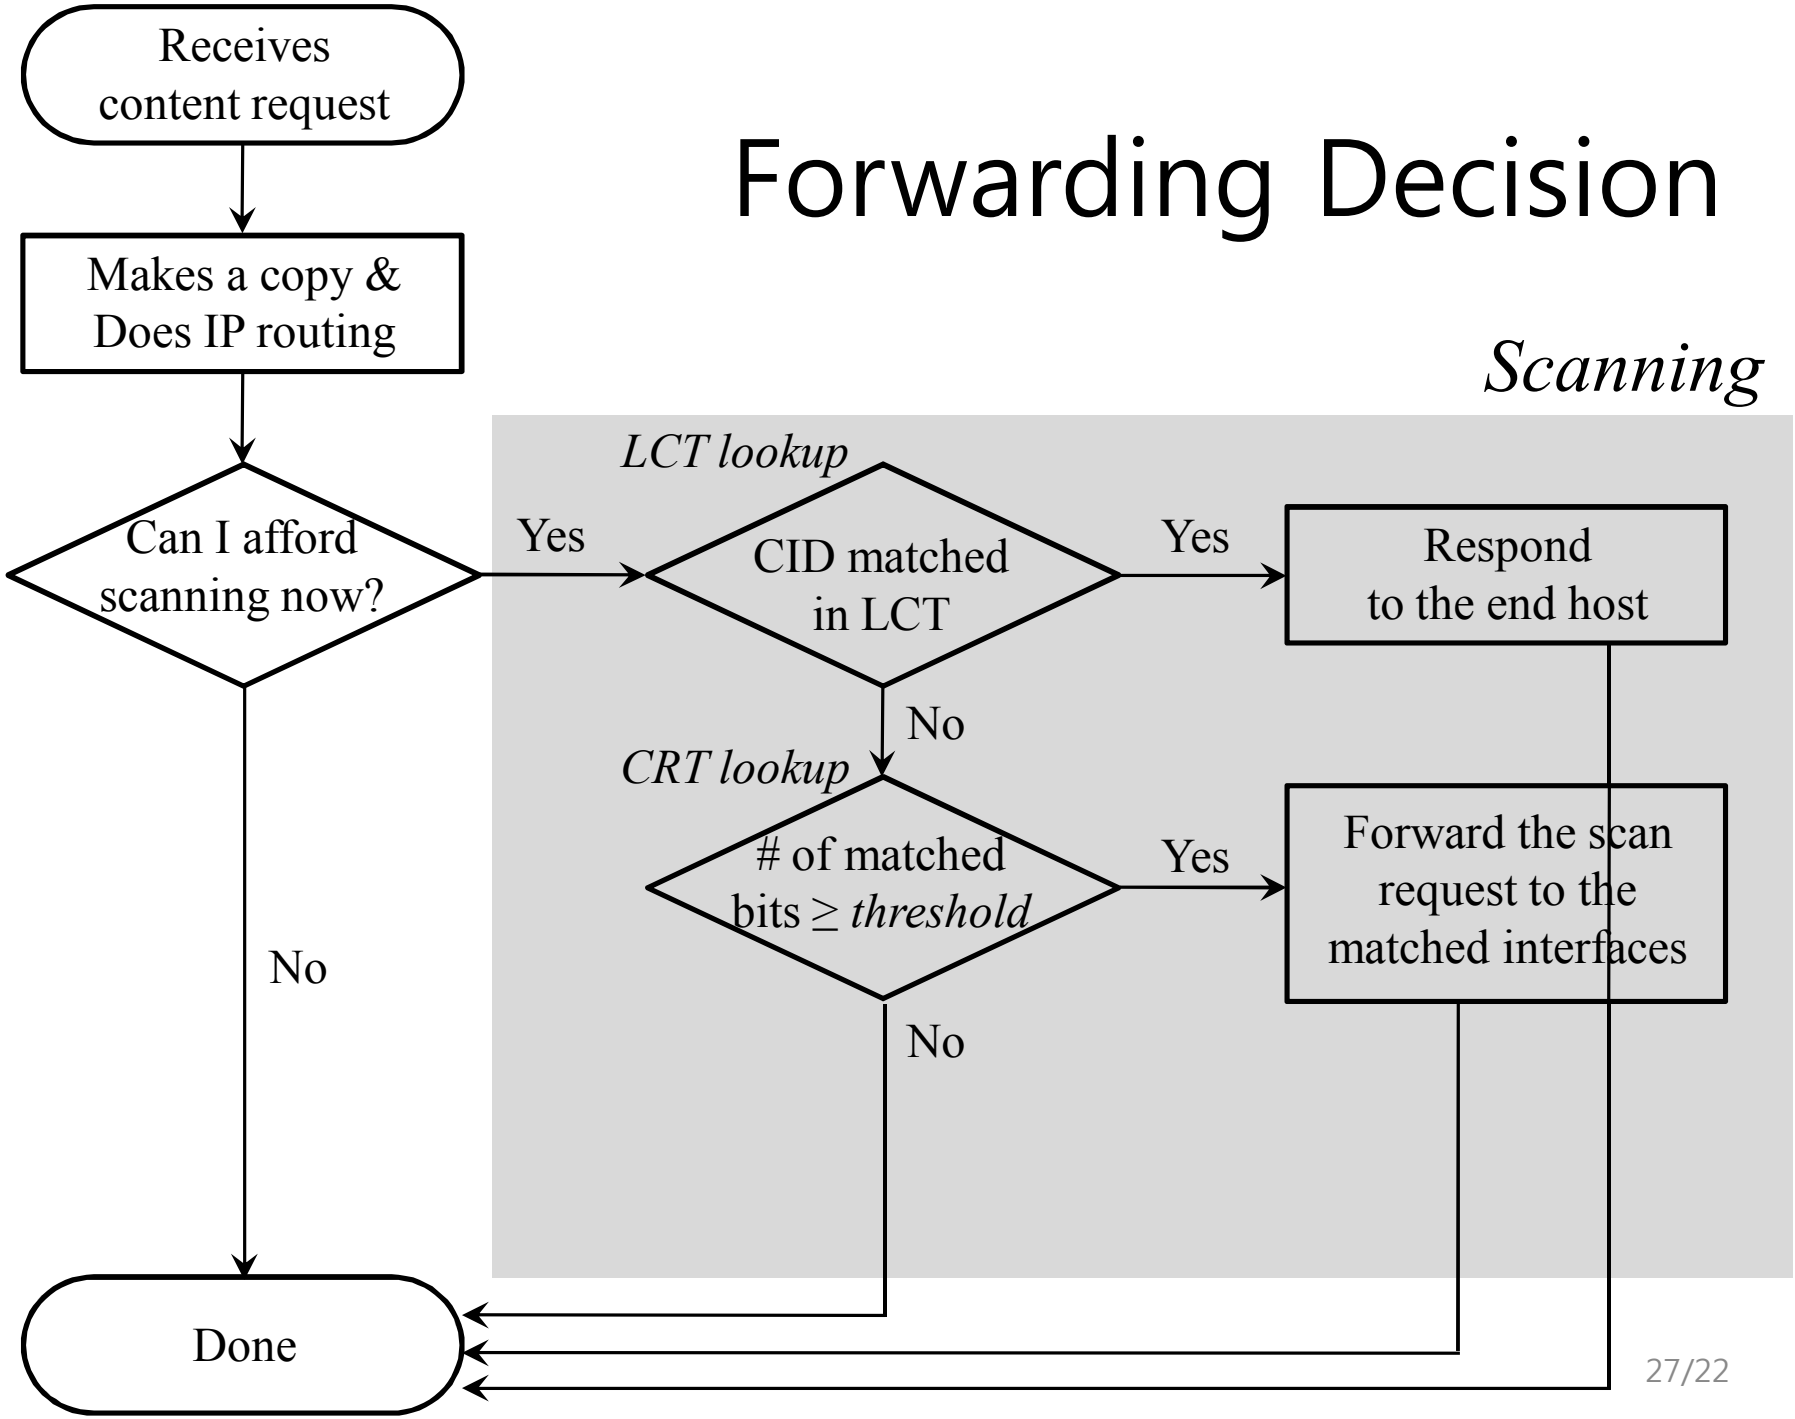

Simulation Setup

- GT-ITM: 1 transit and 5 stub domains
- 1*5+5*20 C-routers and 1,000 end hosts
- Total 20,000 content files
  - 10,000 different content files
  - top 10% have multiple copies: avg. 10 copies
- Content file size: 1GB
  - Each C-router has 100 files

- SCAN vs. IP routing, IP with caching, SCAN w/o BF (BF size C-info), SCAN-full

Simulation Results (1/2)

	SCAN	SCAN-FULL	SCAN w/o BF	IP w/ Caching	IP Routing
Average Number of Hops	3.2	3.2	4.5	4.7	6.8

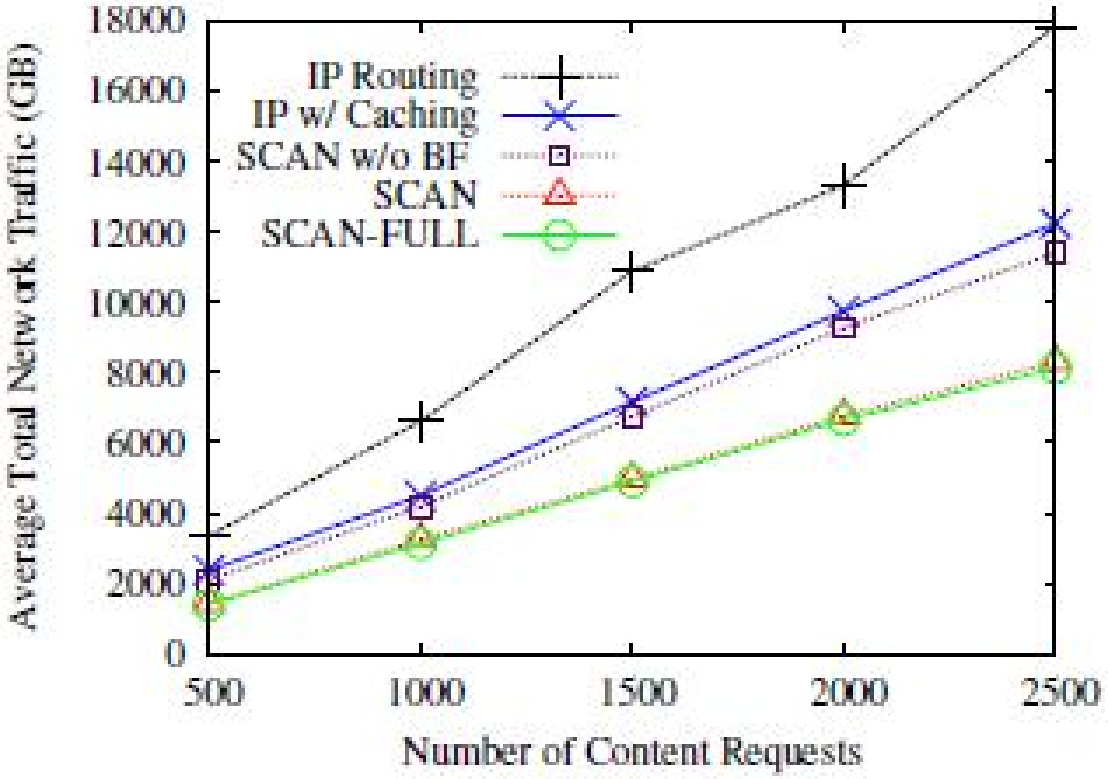

(a) Network traffic reduction

Simulation Results (2/2)

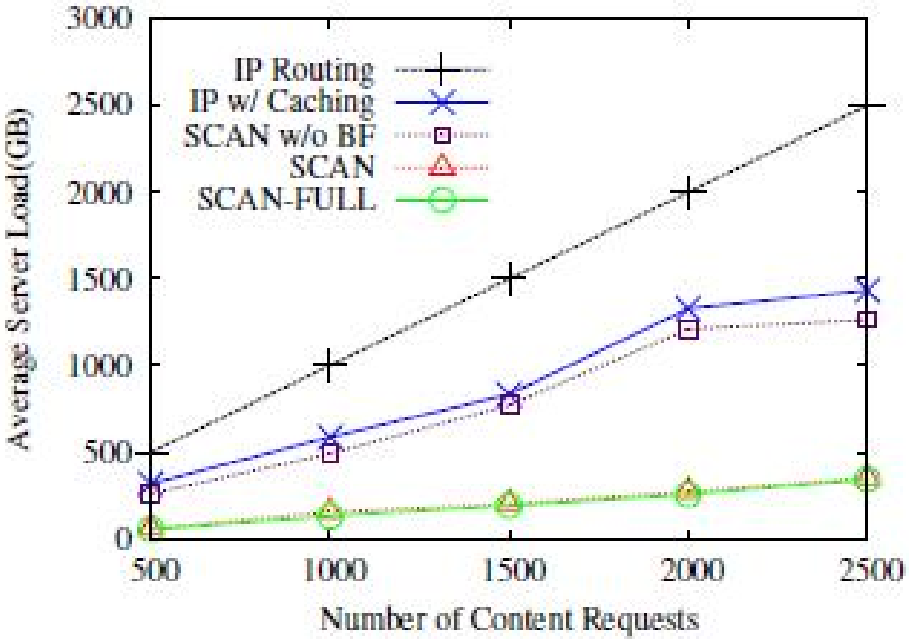

(b) Original server load reduction

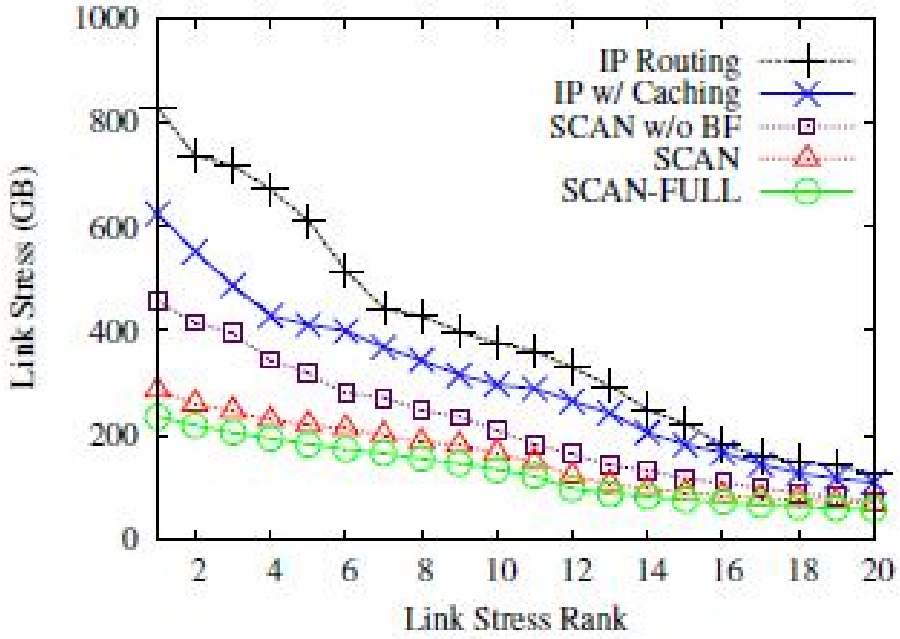

(c) Load balancing among links

Discussions

tkkwon@snu.ac.kr