Spring 2015

ECE544 Project 3: Information Centric Network – Content Routing Protocol Design

Profs. Raychaudhuri & Reininger

Background

This project will focus on the design and implementation of a simple "content routing" protocol where the network service involves delivery of named content files. In contrast to an IP network where packets are routed to a destination address (network port), a content network goes one step further by directly fetching a content file (as identified by a unique identifier) from the nearest location at which it is currently stored. Content may be stored in network-attached devices such as client PCs or servers as in the current Internet. The objective of the network design is to support a service by which a client device can query the network with a content identifier and this query is automatically routed to the nearest location of the content file, and the device storing the content then replies by sending back the content file to the requesting device.

Project Description

The goal of the project is to design and implement a routing protocol that supports efficient delivery of files distributed in the nodes of a network while satisfying the following requirements:

- Each content file will be specified by a unique content ID and can be generated at arbitrary times at any host. The network should support the announcement or discovery of new content files generated at a host such that other hosts can retrieve this file when they ask for by its content ID.

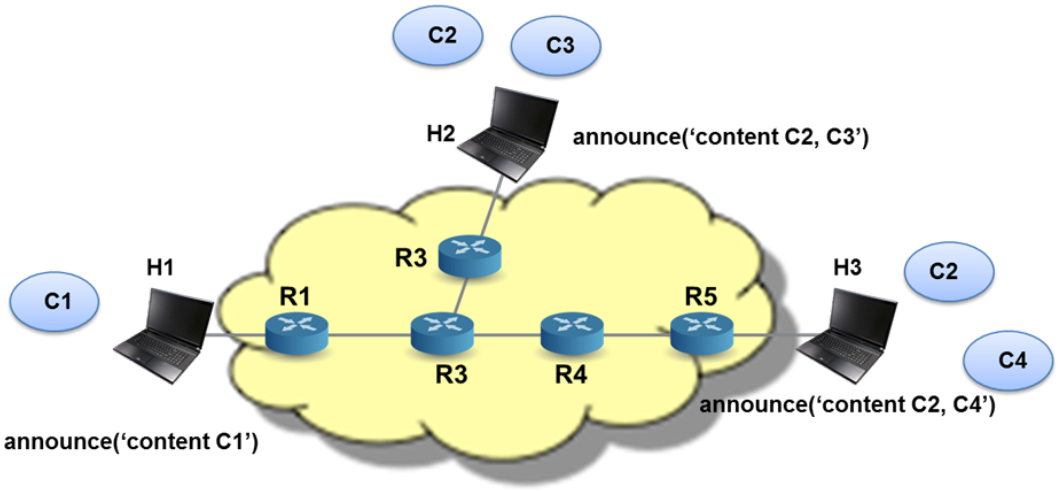

- The same content file can be present at multiple hosts. In such a case, the network must support the retrieval of content files from the best possible location (e.g. when host H1 requests content C2 it is retrieved from H2 because it is closer to H1).

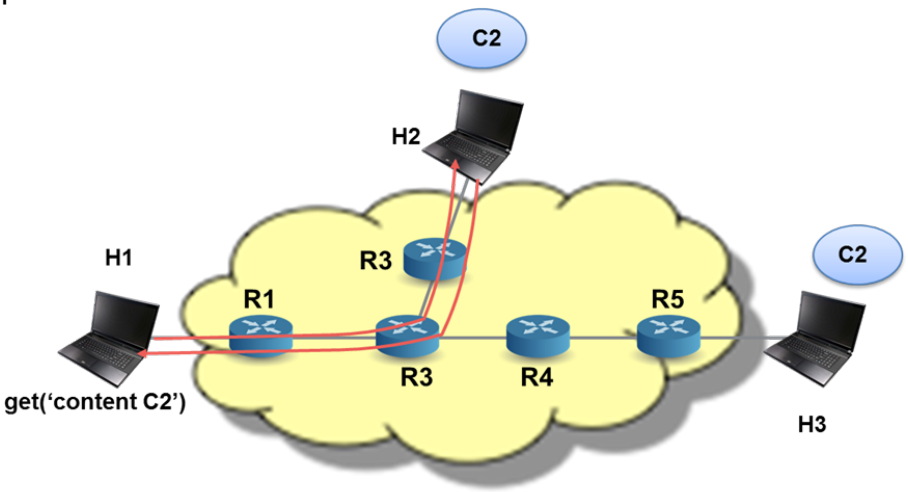

- The content files on a host may get deleted at arbitrary times. Similarly hosts can connect/disconnect from the network at arbitrary times. The network must support these dynamic changes in the distribution of contents (e.g. if content C2 is no longer available at H2 than the content will be retrieved from H3).

All the protocol aspects involved in the network will have to be covered such as: how content is advertised by host machines, how routing of queries is performed in order to reach content locations, how the reply containing a content file is routed back to the requesting source, what ARQ scheme is used for reliability either on links or on an end-to-end basis. Note that as shown in the figures, there are three types of entities to be considered in your design: content files identified by globally unique content IDs, host machines identified by port addresses to which they are connected ("locators") and routers which are connected by an arbitrary topology.

Requirements and Assumptions

The following set of requirements and assumptions have to be taken into consideration:

- The number of hosts connected to the network $N < 255$. Corresponding host address space needed is 0-255. A static ARP table (MAC ID to network address) may be assumed for simplicity.
- The number of content files is limited to $K < 255$ and each file has a preassigned name. Corresponding content ID space needed is 0-255. A static content name resolution table (content ID to content file name) can be assumed for simplicity.
- Content files are no bigger than 1466B units to be carried in a single packet.
- Location of contents might change over time (insertion and deletion of contents at different locations).
- Links connecting network elements have a probability of loss equal to *p.*

- The network service to be provided has the basic primitive get(content_ID), and should be realized as a packet datagram type design rather than a connection oriented design (i.e. the content-ID is explicitly provided in the packet datagram and used by the network to route the query).
- The actual protocol design approach is not specified and different teams may decide to use routing based on content identifiers alone or with both content identifiers and host addresses.
- Alternative designs will be compared in terms of control overhead costs and data path forwarding efficiency. Protocol complexity should also be considered when making design choices.
- No additional network elements such as directory servers or centralized gateways are to be used.

Project Guidelines

Each group is required to develop their own protocol specification proposal, and these will later be converged into a single specification using a process similar to an industry standards committee. Groups of 3 students will first present their proposals using a standard Powerpoint template. Groups with similar design characteristics will then join and present a unified proposal. Finally the whole class will have to reach an agreement on a common design (detailed in a specification document). The spec will be finalized through mailing list distribution and feedback with two student volunteers serving as committee chair and secretary respectively. The chair is responsible for conducting the meetings and the secretary is responsible for writing and distribution of the spec draft. The final draft will need to include the protocols spec and algorithms necessary to implement the specified network.

Finally, each initial group will have to submit a software implementation of the components belonging to the network.

Active participation in the committee meetings will be a factor in the final grade.