Architecture for IPTv Distribution: Cooperative P2P and Multicast

Vijay Gopalakrishnan, TaeWon Cho, Rittwik Jana, Samrat Bhattacharjee, Divesh Srivastava, Yin Zhang* and K. K. Ramakrishnan

AT&T Labs Research, NJ

*University of Texas at Austin

May 2007

Video Distribution

- Video viewing is likely to be increasingly on-demand
  - Only news and sports events are likely to be delivered in real-time as “linear TV” over the long-term
- Video on-demand system in the service provider context currently organized as:
  - Set of media distribution servers in metropolitan cities that retain copies of on-demand content
  - User requests served on a unicast basis
    - Desirable characteristic: Provider can control perceived quality
    - But approach **does not scale**
- Multicast based approach for serving on-demand content highly desirable
  - But opportunity to exploit it while meeting user’s requirements for small start-up delay limited

Alternate Technologies on the Internet

- On-demand content is increasingly provided by mainstream content providers
  - ABC, NBC and other “mainstream” providers.
  - Depend on caching to overcome scaling problems while serving requests on a unicast basis
- Peer-to-peer technologies were being used by “hobbyists” until recently
  - Based on end-systems and Internet connectivity
  - Becoming more “mainstream” (e.g., BBC)
  - Likely to be increasingly used (e.g., ’08 Olympics)
- Peer-to-peer technologies have been predominantly download-and-play
  - BUT: increasingly being refined to provide streaming capability
  - But mechanisms do not exploit or need knowledge or control of network infrastructure
    - No SLAs can be made. Use end-system intelligence and buffering.
    - Start-up latency may still not be satisfactory for a “paying customer”

Serving the needs of “Entertainment”

- “Entertainment” content (Video and Audio): large amount of data, with real-time constraints
  - Place stringent demands on the network – capacity and latency – tight performance requirements in general.
- Peer-to-Peer technologies are attempting to serve this need across the entire spectrum – including “streaming” content to the user
  - But not tied to the infrastructure: can be quite inefficient in how they use the network.
- Users will migrate to “The single converged Network” when it can meet their needs
  - one which provides a satisfying experience.
- Entertainment is all about satisfying the viewer in such a way that they are absorbed in the story being told – not how it is delivered.
  - No significant Latency, Loss, Artifacts; Good sound, without clipping or interference.

Streaming with Peer-to-Peer Technologies

- "Traditional" P2P focused on "finding" the content and caching the content at peer end-points
  - Distribution of content across points in the network help to serve the needs of a distributed population
    - Driven by nodes that have indicated an interest in that piece of content
- Streaming content with P2P: solutions beginning to grapple with the traditional issues
  - Scheduling: which request to serve; when to make a request; prioritizing requests
  - Resource management: overcome limitations in bandwidth by having multiple peers serve "chunks" of content - stripe from peers
- Content Providers complement P2P technologies
  - Servers to complement peers serving up content to overcome capacity limitations
    - Unicast; multicast (different forms: cyclic; skyscraper etc.)

Our Approach for Video-on-Demand

- Unified approach to provide efficient support for VoD in a service provider environment using
  - Multicast
  - Caching
  - Peer-to-peer that is topology aware
- Good user experience
  - Fast start: Decouple user-perceived performance from popularity
  - Maintain quality – minimum (→ zero) user perceived interruptions while watching arbitrary length content
  - User experience should be limited only because of user client capability/storage
  - Make it easy for users to find the content of interest
- Service provider friendly
  - Scalable: Decouple performance from population of users
  - Efficient use of resources

Video Distribution: Environment

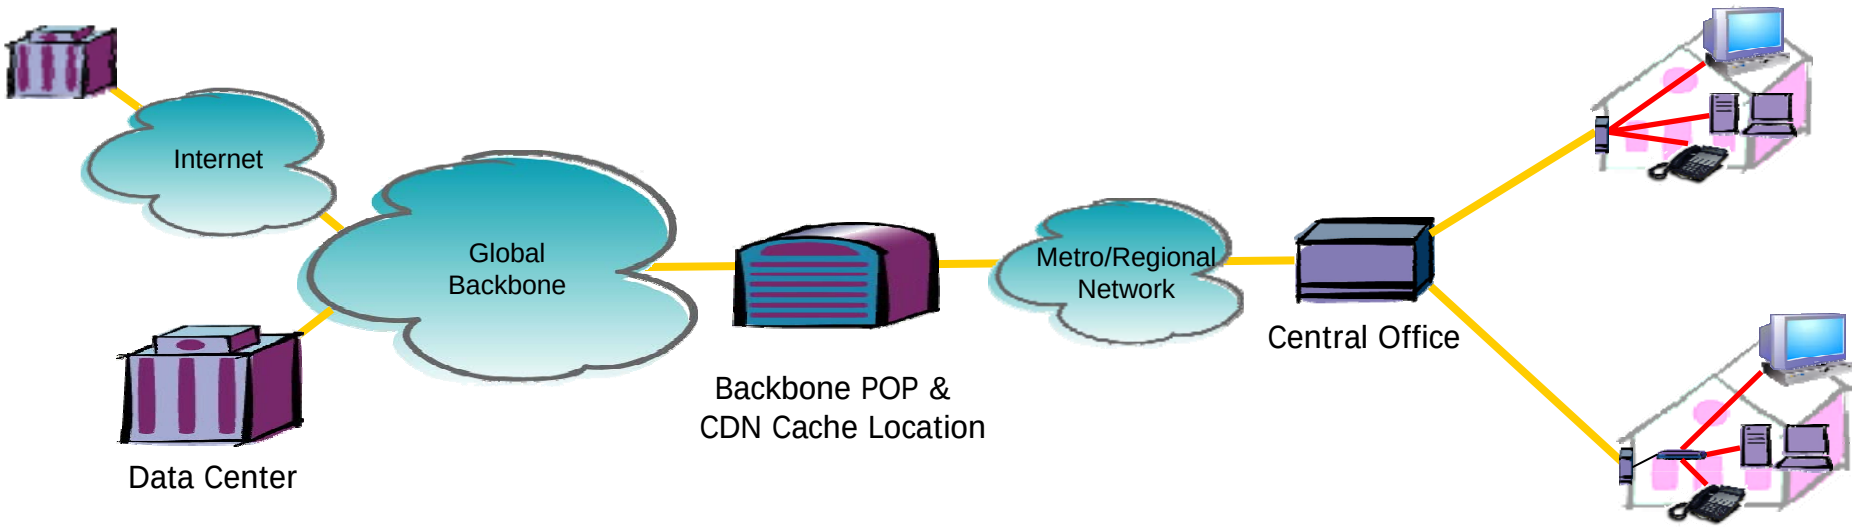

- With traditional P2P or unicast, traffic traverses significant portion of backbone
  - Network optimized solutions are desirable
    - Multicast for live content is very desirable
    - CDN caches for popular non-live content
  - We are investigating peer-assisted near-VoD that uses multicast and caches that understand and exploit the topology

Goals of our Architecture

- Use Network Resources Efficiently
  - Use multicast wherever possible
  - Reduce server load
  - Use peer capability to store and serve popular content whenever possible
  - Leverage storage and intelligence in the clients when possible; server and network wherever necessary; caching
  - Exploit popularity of content to achieve efficiency and optimize user experience
- User experience should be limited only because of user client capability/storage
  - Isolate bandwidth and server capacity limitations from the clients as much as possible
- Early stages of design and implementation of a prototype to help us understand the issues

Long-term: Need to Handle Meta-Data

- IP as the medium for Video Distribution gives us the opportunity to enrich the viewing experience for the consumer
- Thousands of content providers for video and multimedia streams worldwide
  - Topics may be of local, regional, national and international interest
- Desire: serve diverse needs of communities
  - Distribute programming of interest to different ethnic groups; programming that may be generated worldwide
- Enable integration of video with other media
  - E.g., integrate a browser to provide related information for a viewer of a current program
  - Enable launching of video related to text consumer is reading on browser

Searching for Programming

- Search for linear TV by
  - Interactive channel guide that we typically see with video distributed on other media
  - Name of a TV show
  - Name of a person involved in the TV show
- Search for a video-on-demand
  - Name of a movie
  - Name of a person (director; actor)
- Desire:
  - Ability to search in flexible ways, using fine-grained specification of interest
  - Ability to subscribe to programming of interest, especially as more and more content moves to being on-demand

Scaling Issues with Large Numbers of Content Providers offering content

- Information Scale: Producers and Consumers face challenges
  - Large number of producers (publishers; data sources)
  - Even larger number of consumers (subscribers, users querying/looking for content)
  - Large number of information producers makes it difficult for a consumer to know where to find relevant information
  - Significant challenge: “whom to ask” and “whom to tell”
- XML becoming ubiquitous format for information exchange
  - With XML: easier to find information of interest and extract data
  - Keyword queries; Structured queries
- Annotating videos: e.g., with XML tags (MPEG)
  - Enable more elaborate searches
    - Fine grained specification of desired content
  - Enable combinations of “publish-subscribe” and “search-view” of content

XTreeNet: Meta-data and media-data

- Meta-data describes the media-file
  - Generated from closed caption, speech recognition, DVD subtitles
  - Publisher can be a media source (NBC) or second-hand producer (Miracle)
- Network connects clients to publishers using CD
  - Content Descriptors (CDs) act like “indexes” in a distributed data base environment
  - CDs decouple producers from the consumers
  - CD can be keyword (“Britney Spears”) or XML schema path (“/title/nightly news”)
  - Multicast meta-data over multiple core based trees
    - Different cores for different CDs – to reduce traffic concentration

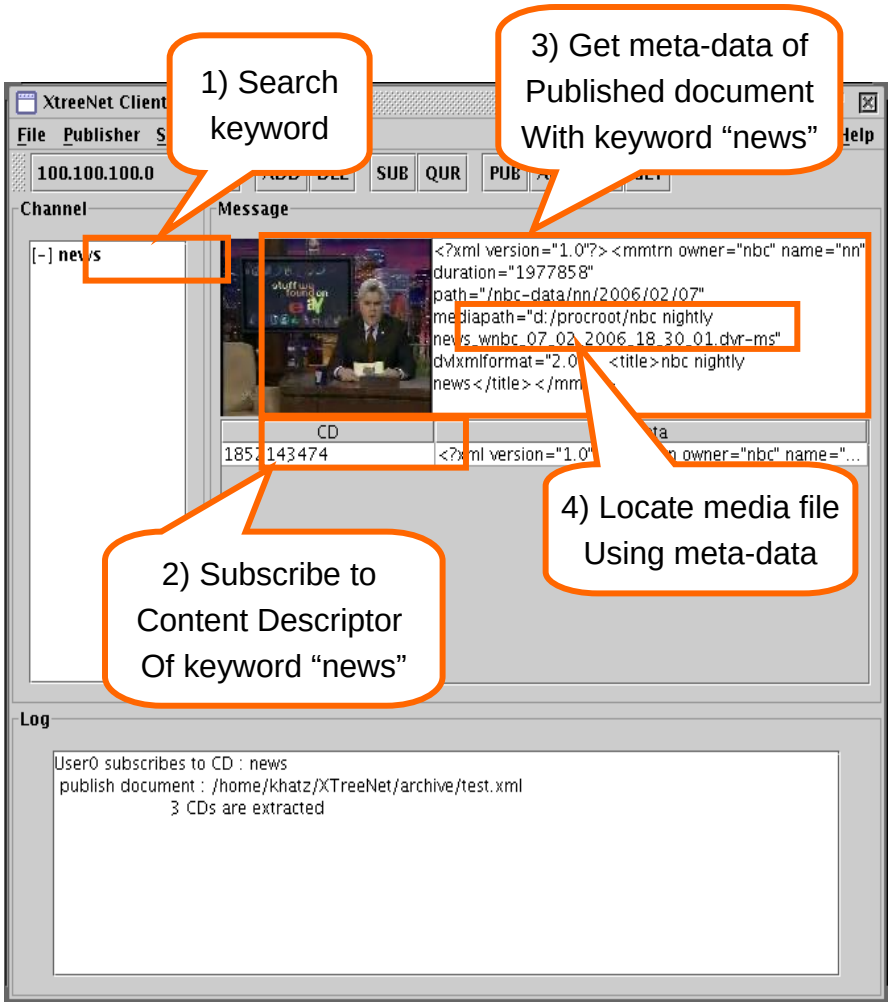

XTreeNet Client