Demand-Aware Content Distribution

Srinivas Shakkottai
Texas A&M University

Hybrid content distribution

High level idea:

Use P2P dissemination to "assist" traditional client-server methods, e.g., content delivery network (CDN).

Key question:

How should the two methods be combined?

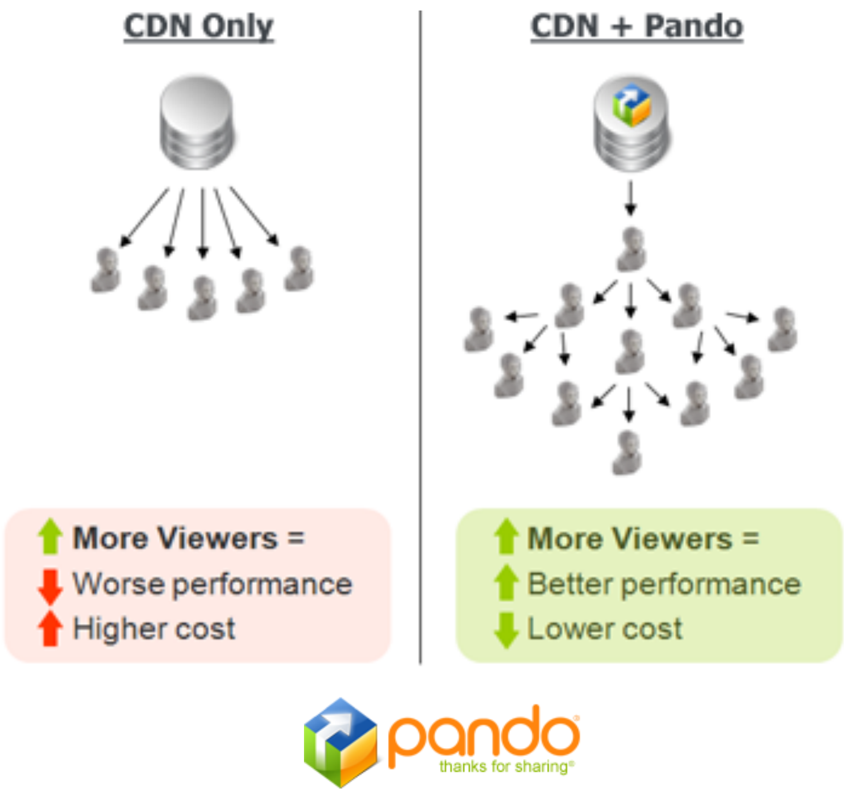

Outline

- Demand Evolution
- Service models: CDN, P2P, and hybrid
- Comparison
- File arrivals: heavy traffic and multiplexing
- Future work

Demand Evolution

Bass model (1969)

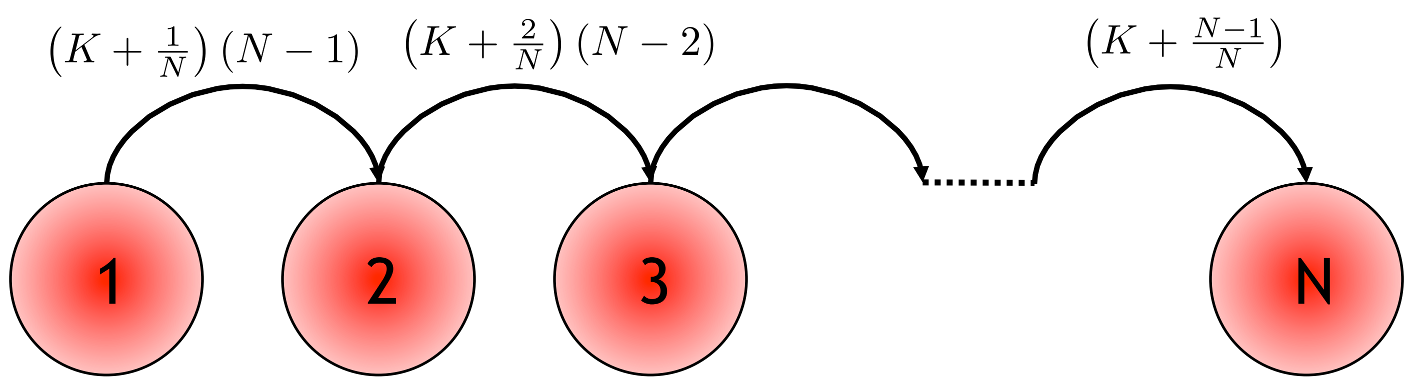

- Total user population of size N.
- Exponentially distributed transition rates.
- Effect of advertising captured by K.
- "Word-of-mouth" propagation of interest adds to transition rate.

Fluid model

- Total user population of size N (infinitely divisible)
- $I(0)$: initial number of interested users
- Effect of advertising captured by K.
- Interested users select other users at random

$$\frac{dI(t)}{dt} = \left(K + \frac{I(t)}{N}\right)(N - I(t))$$

Advertising

Fraction of interested users

Random user is not interested

Single file demand model

- This demand model is a version of the *Bass model* with only word of mouth propagation.
- Solution:

$$I(t) = \frac{N I(0) e^t}{N - I(0) + I(0) e^t}$$

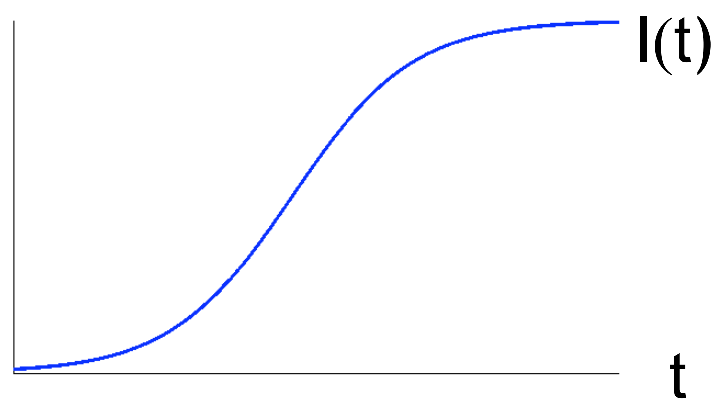

Propagation in Power Law Graphs

- Thresholds for virus spread on networks, Draief et al.
- The Effect of Network Topology on the Spread of Epidemics, Ganesh et al.

- Interested users never leave, so demand is not modulated by supply.

Data from CoralCDN

- CoralCDN is a distributed network running on PlanetLab.
- Duplicates popular files, http://www.cnn.com.nyud.net
- Data on multiple popular video files on the Asian Tsunami courtesy M.Freedman.

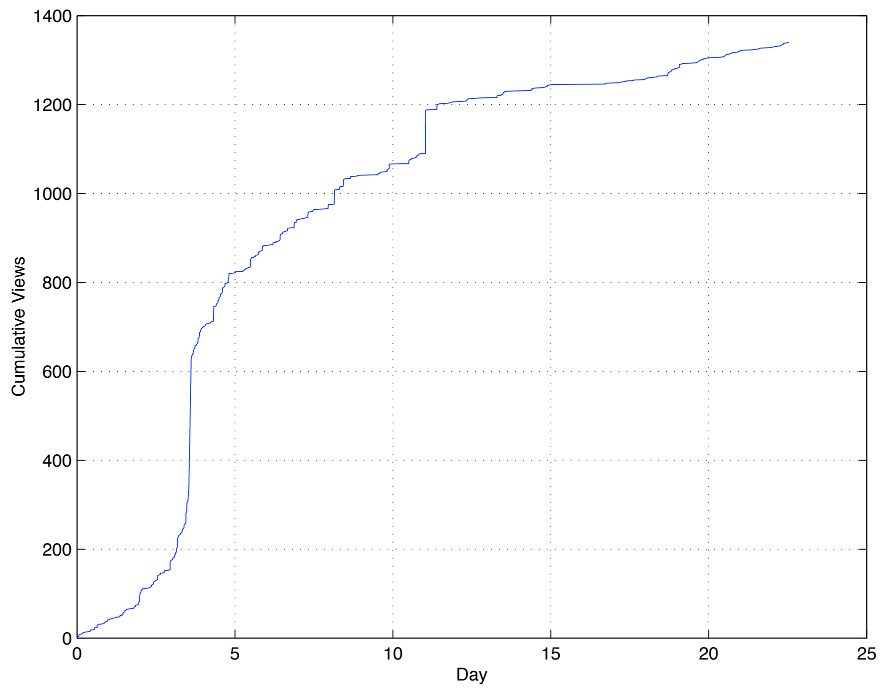

Supplying Demand

Service models

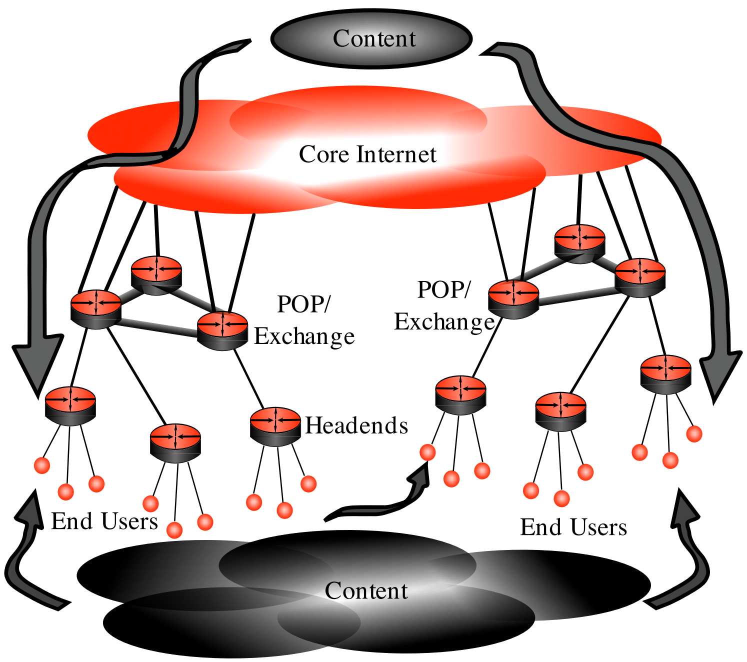

- *CDN*: Use a bank of servers
- *P2P*: Use peer-to-peer dissemination
- *Hybrid*: Use both

Which has the best delay performance as N scales?

- $P(t)$ denotes *cumulative service up to* t.
- Work conserving service assumed:

→ total delay = *area between* $I(t)$ *and* $P(t)$.

Service model I: C-D

Installed server capacity: C users per unit time

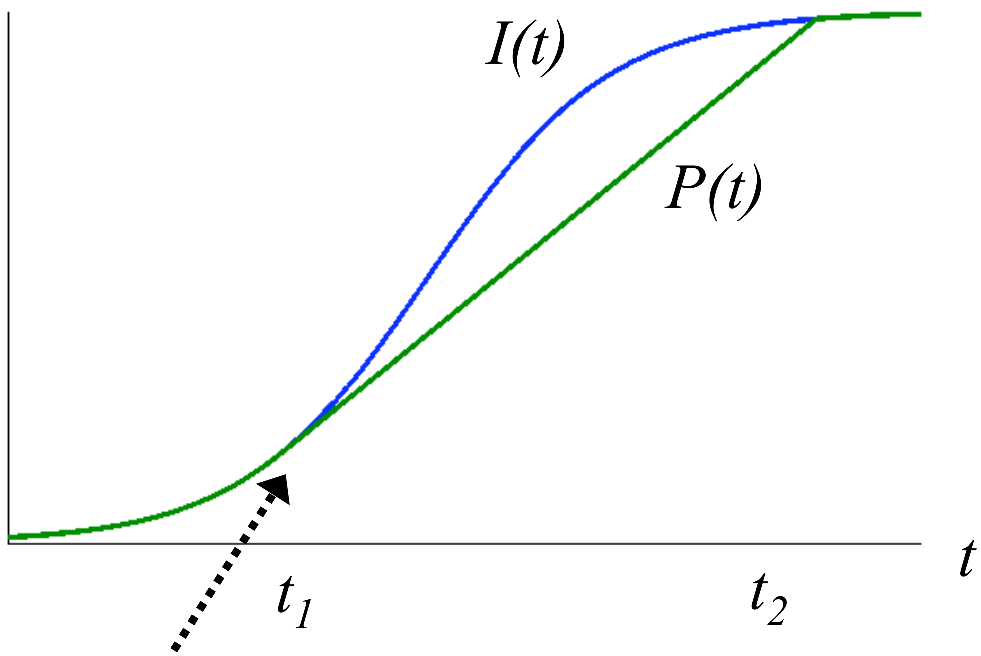

Service follows interest as long as $dI/dt < C$*, i.e., until* t_1 *…*

Service model I: CDN

Installed server capacity: C users per unit time

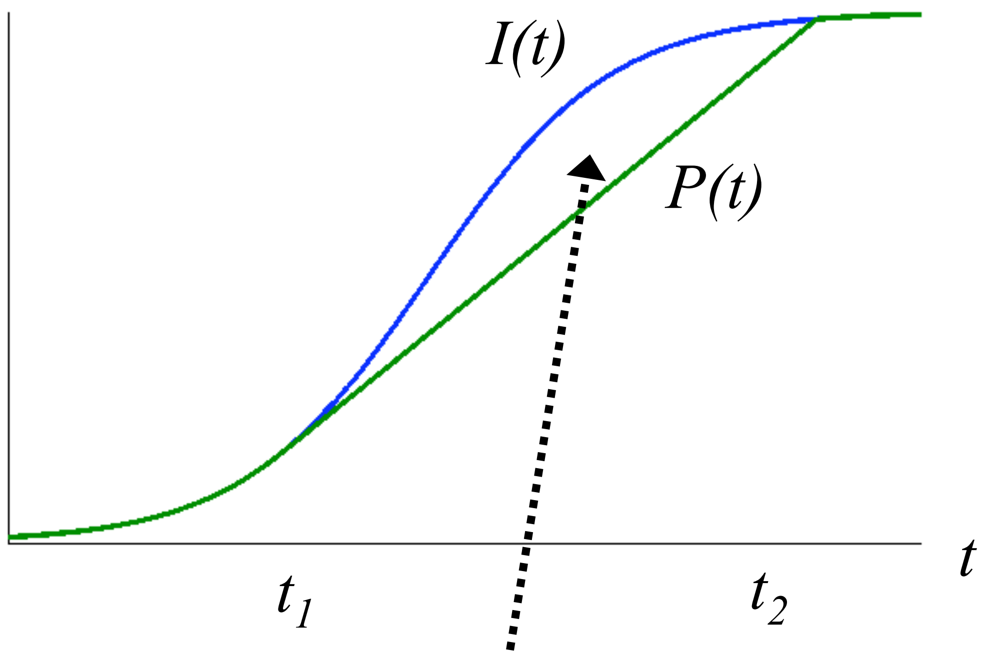

… after which interested users have to wait (until t_2).

Service model I: CDN

Proposition:

$P(t) = I(t)$ for $t \in [\,0,\, t_1]$, and $t \in [t_2, \infty)$, and
$P(t) < I(t)$ for $t \in (t_1, t_2)$, where

$t_1 \in \Theta(\ln(C/I(0)))$; $I(t_1) \in \Theta(C)$, and
$t_2 \in \Theta(N/C)$; $I(t_2) \in \Theta(N)$

Further, the area between $I(t)$ and $P(t)$ scales as $\Theta(N^2/C)$.

Service model II: P2P

- Model motivated by Bass diffusion
- Assume that “efficiency of sharing” given by parameter ν

$$\frac{dP(t)}{dt} = \nu(I(t) - P(t))\frac{P(t)}{N}$$

- Random peer selection
- Can be solved explicitly

Service model II: P2P

Comparison of interest and service curves:

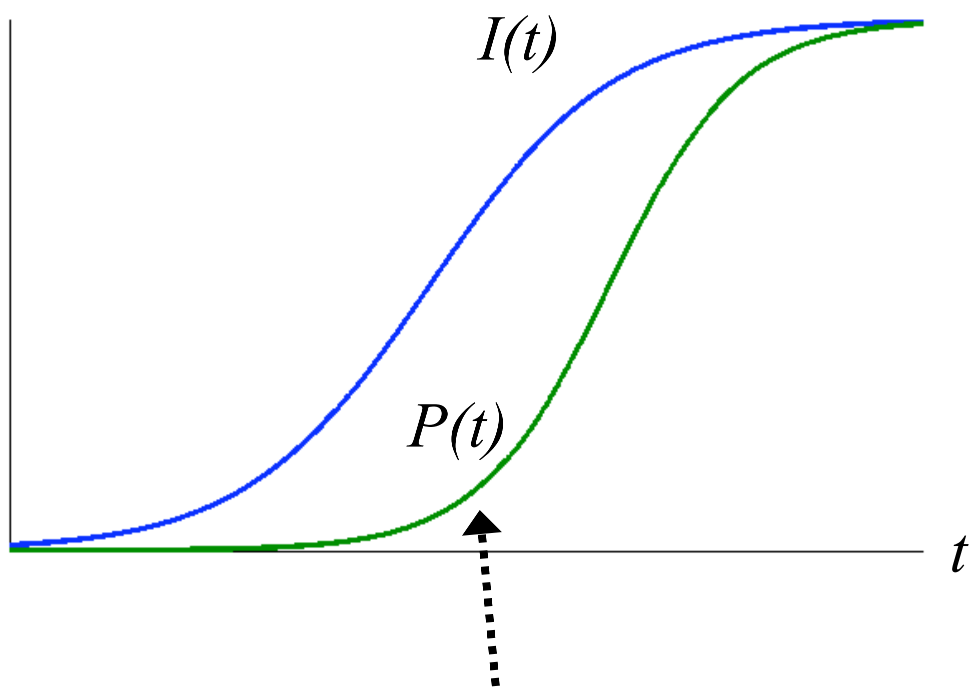

At time $t = \ln N$, $I(t) \sim N$ while $P(t) \sim 0$ ...

Service model II: P2P

Comparison of interest and service curves:

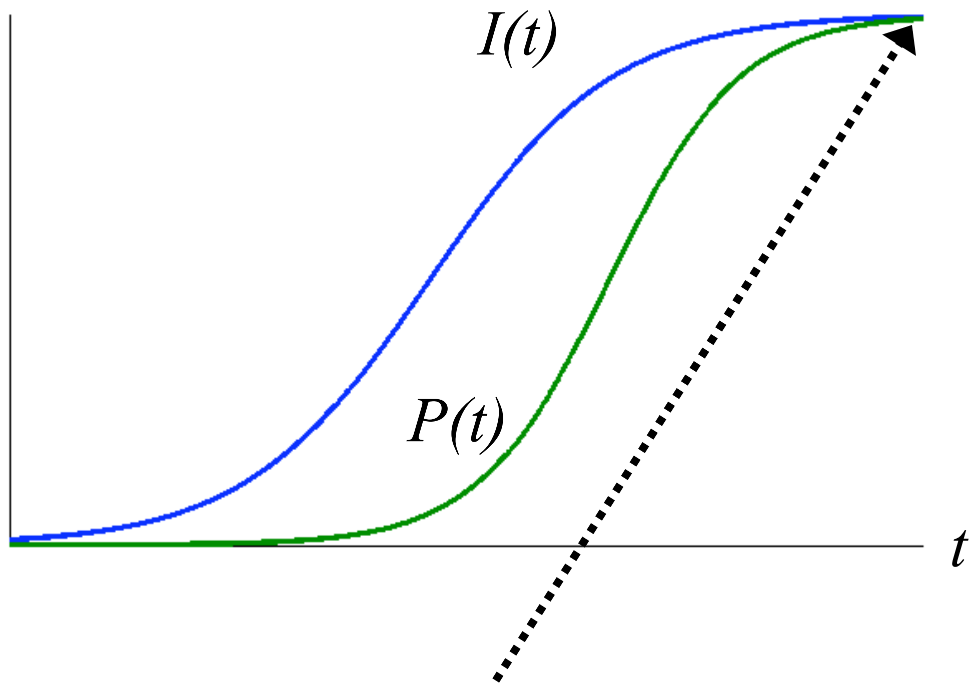

… but by time $t = 2 \ln N$, $I(t) \sim N$ and $P(t) \sim N$.

Service model II: P2P

Proposition:

$P(t) \approx I(t)$ for $t \geq 2 \ln N$.

Further, the area between the interest and service curves scales as $\Theta(N \ln(N/P(0)))$.

Service model III: Hybrid

- CDN does well *until interest overloads servers*
- P2P does well *once installed user base is large*

- Consider a hybrid scheme where:
  - CDN used until $t_1 = \Theta(\ln (C / I(0)))$
  - P2P used thereafter

Service model III: Hybrid

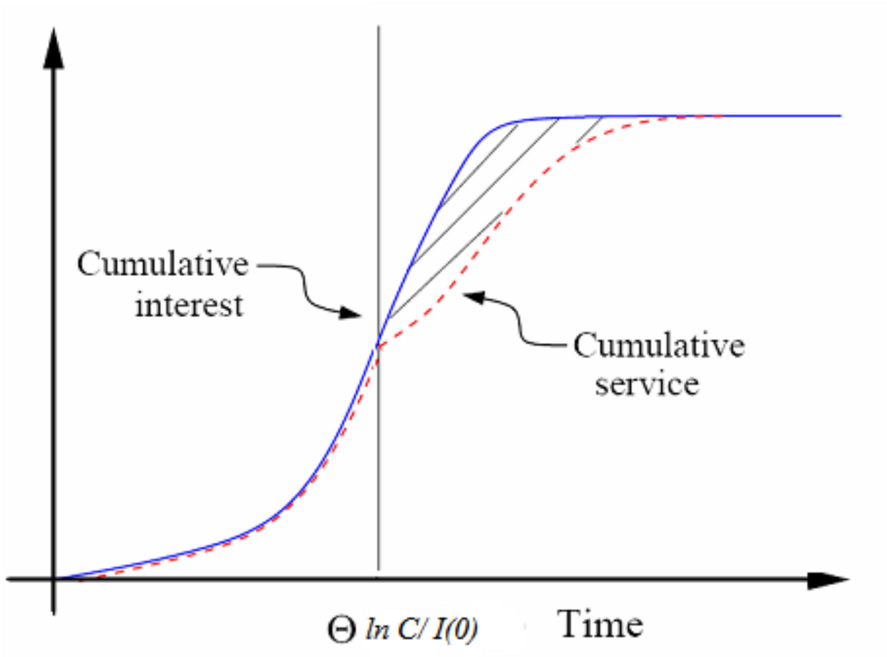

Proposition:

For the hybrid scheme, the area between the interest and service curves scales as $O(\ N \ln(N/C)\ )$

if $C = o(N)$.

Comparison

Per user delay is:

$\Theta(N/C)$ for CDN;

$\Theta(\ \ln(N/P(0))\ )$ for P2P;

$O(\ \ln(N/C)\ )$ for hybrid.

Choice of dissemination method will depend on cost structure of capacity. We now develop an example to study this.

Example

Per user delay using $C = N / \ln N$:

$\Theta(\ln N)$ for CDN;

$\Theta(\ln N)$ for P2P;

$O(\ln \ln N)$ for hybrid.

Capacity gain of $\ln N$ or equivalently, delay gain of the same order.

C-D versus P2P

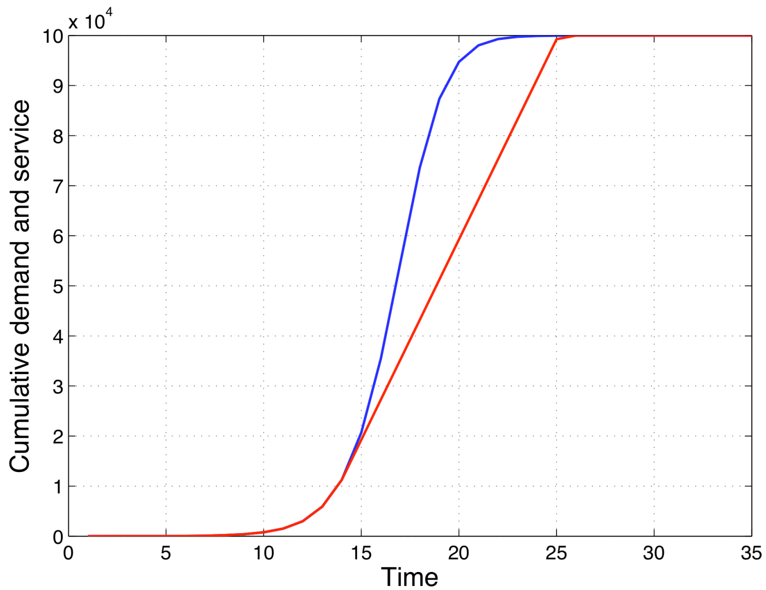

Centralized Distribution

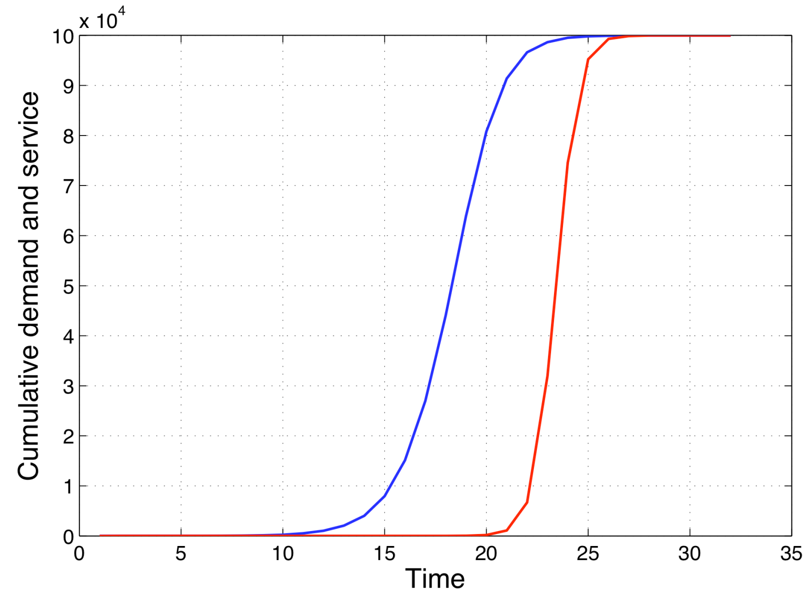

P2P Distribution

Hybrid Scheme

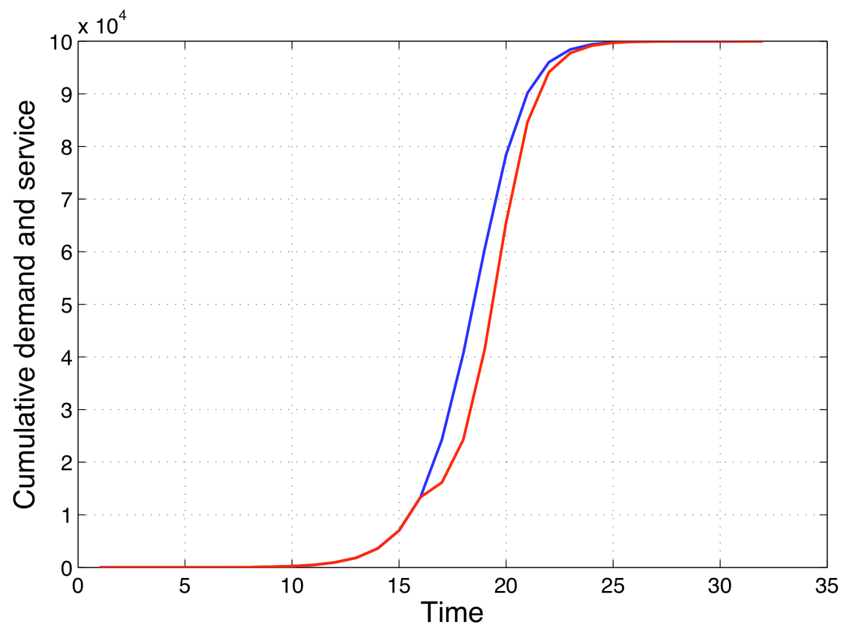

- Combines initial centralized distribution with later use of P2P.
- Central server is used only to "boost".
- Early estimate of total population allows us to determine "switching point" to guarantee an average delay.

Simultaneous use of C-D and P2P

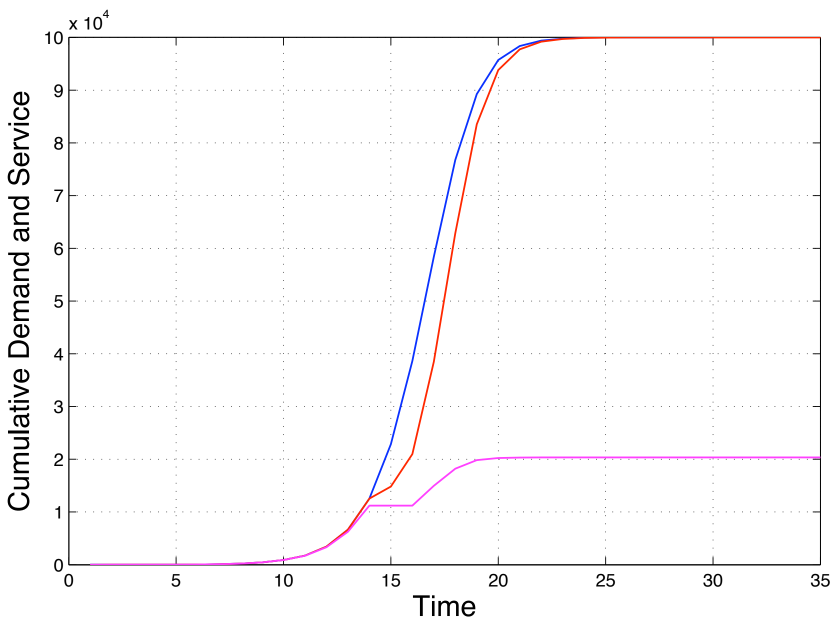

- Why have a distinct threshold?
- Use both C-D ad P2P initially → P2P has no effect.
- Use C-D to “boost” if required in the latter phase → C-D has no effect.

Dynamic File Arrivals

Data from CoralCDN

- CDN has to handle multiple files.
- Load binned using per minute binning.
- Traffic is bursty.

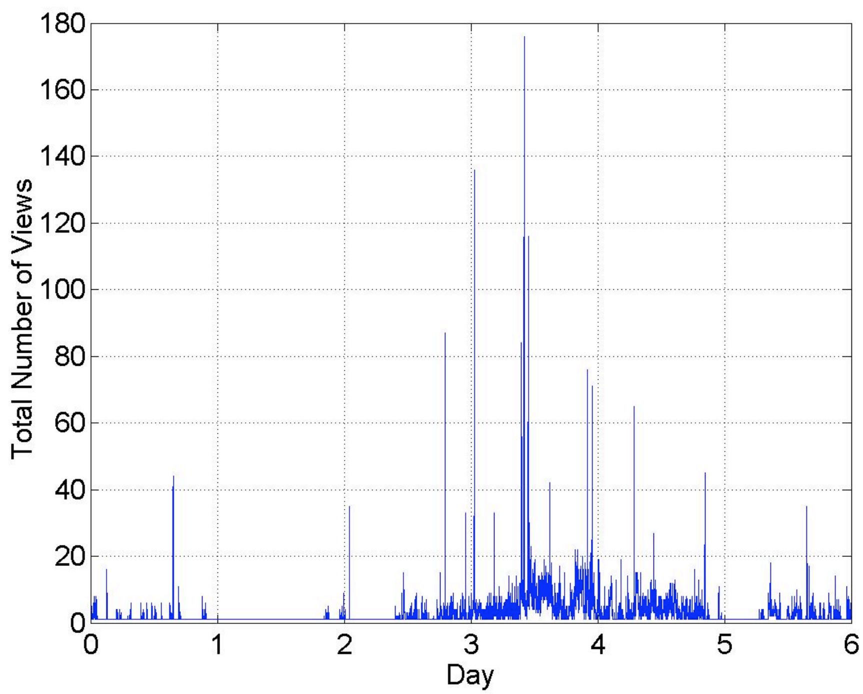

File Arrivals

Suppose now that a content distributor uses a CDN to simultaneously handle *dynamic file arrivals*.

Consider a *flow level fluid limit* where

λ = arrival rate of files per unit time.

N = Number of potentially interested users in each file.

What is the minimum capacity required in order to give an average per user delay guarantee d?

Multiple files: Hybrid Approach

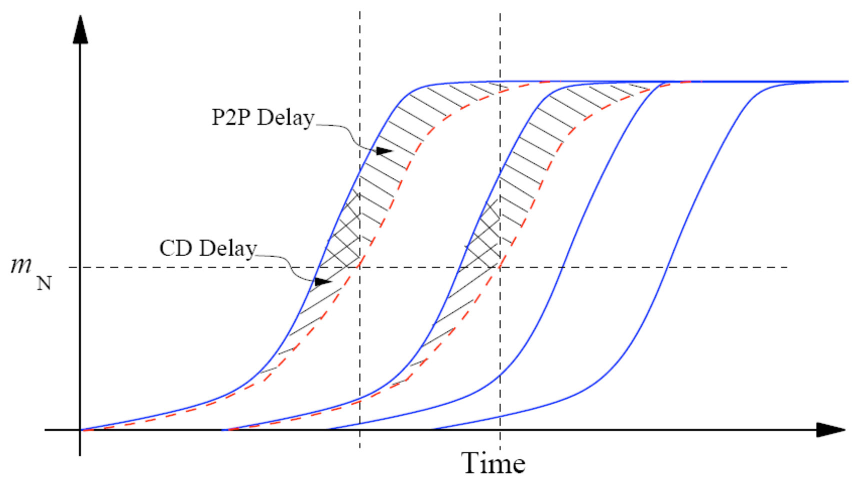

- The available capacity is multiplexed among different files.
- Say we serve m_N users for each file using centralized distribution.
- Minimum required capacity is $C_N = \lambda\, m_N$.

Multiple files: Hybrid Approach

Proposition: (heavy traffic or not?)

$$C_N = \begin{cases} \lambda N e^{-d} & if \ \lim_{N\to\infty} \sqrt{N} d e^{-d} = \infty \\ \bar{C}_N & else \end{cases}$$

Use a diffusion approximation of an M/D/1 process. Example: If $d = \ln \ln N$, then the heavy traffic regime applies.

In case of small desired delay, the P2P phase delay dominates, and "ideal" multiplexing of available capacity may be achieved.

Conclusions and ongoing work

- *Key insight:*
  It is possible to quantify the benefit of CDN-assisted P2P dissemination for large system scalings.

- *Ongoing work:*
  Incentivise users to stay.
  Handling varied topology effects.
  Use the QoS expressions as input to algorithm design.

Long Links and Incentives

- Each ISP has an incentive to keep traffic within its infrastructure.
- Exist P2P algorithms that reveal only a subset of content instances to peers.
- Need to create long-links to other ISPs on a need basis.
- In other words, the navigability of the network needs to change based on demand.


Thank you!