[奨励講演]
インセンティブによるピア協力型 CDN の運用実現性向上

NS-OCS-PN研究会

2009/6/25

森†,上山†,原田††,長谷川†,川原†

†NTT サービスインテグレーション基盤研究所

†† NTT西日本研究開発センタ

Background

- The great popularity of large-scale video services on the Internet; e.g., YouTube, Youku
- Not only UGC (User-generated content) but professional content is attracting traffic; e.g, Hulu, BBC iPlayer
- Dominant of Internet traffic is now shifting to video from P2P file sharing.
  - Ratio of P2P traffic will decrease to 20% in 2013 (Cisco, May 2009)

→ Managing ultimately huge video traffic is an important and challenge task

Existing approach -- Peer-assisted CDN

- The idea:
  - Make use of resources of participating peers
    - 100,000 users x 1GB / 10 % of CPU → 10TB of storage capacity, 10,000 CPU power
  - Distribute the workload on several peers/locations → good scalability
- Effective to the large-scale video sharing services;
  - YouTube [imc07]
  - MSN Video [sigcomm06]
- Already deployed in the real world
  - Joost, BBC iPlayer, P2P-next

Drawback of Peer-assisted CDN

- Random peer selection
  - It does not consider underlying network topology
  - Traffic can be unnecessarily scattered
  - Increase cross-domain traffic, which in general requires cost to deliver

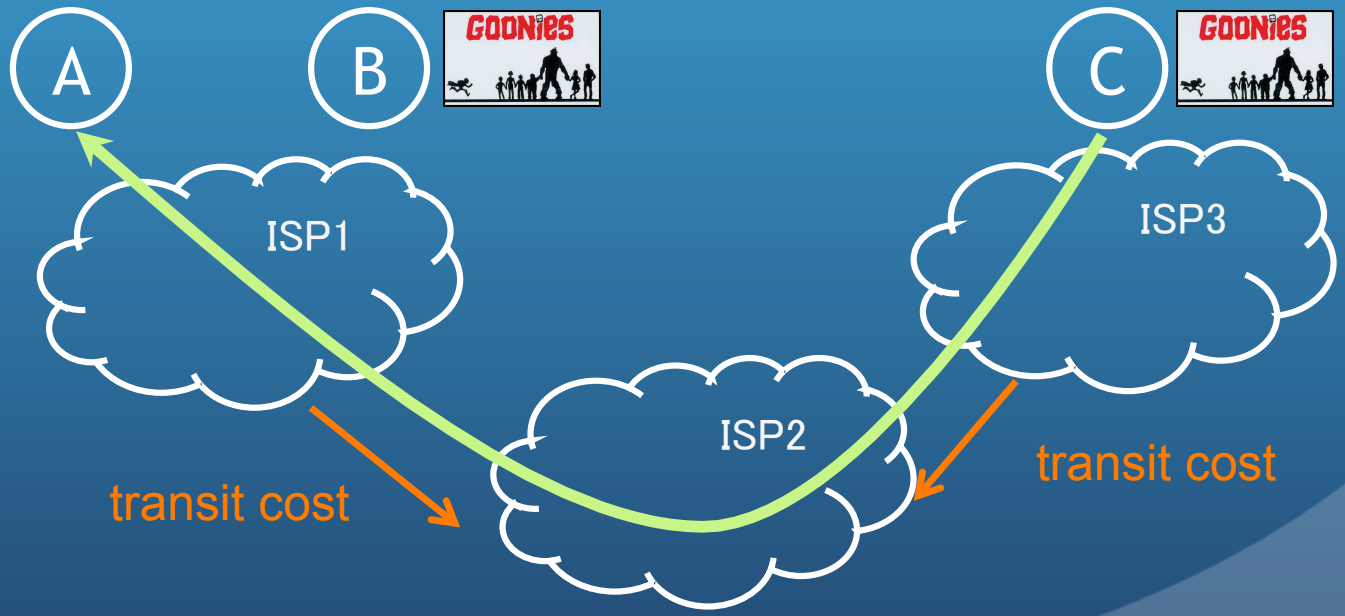

Existing approach
– Managed Peer-assisted CDN (MP-CDN)

- The idea:
  - Make use of "Oracle" to avoid inefficient peer selection
  - "optimize" traffic based on the knowledge collected by Oracle
- Extensively studied in the past year
  - P4P: SIGCOMM 2008
  - Taming the Torrent: SIGCOMM 2008
  - IETF ALTO WG

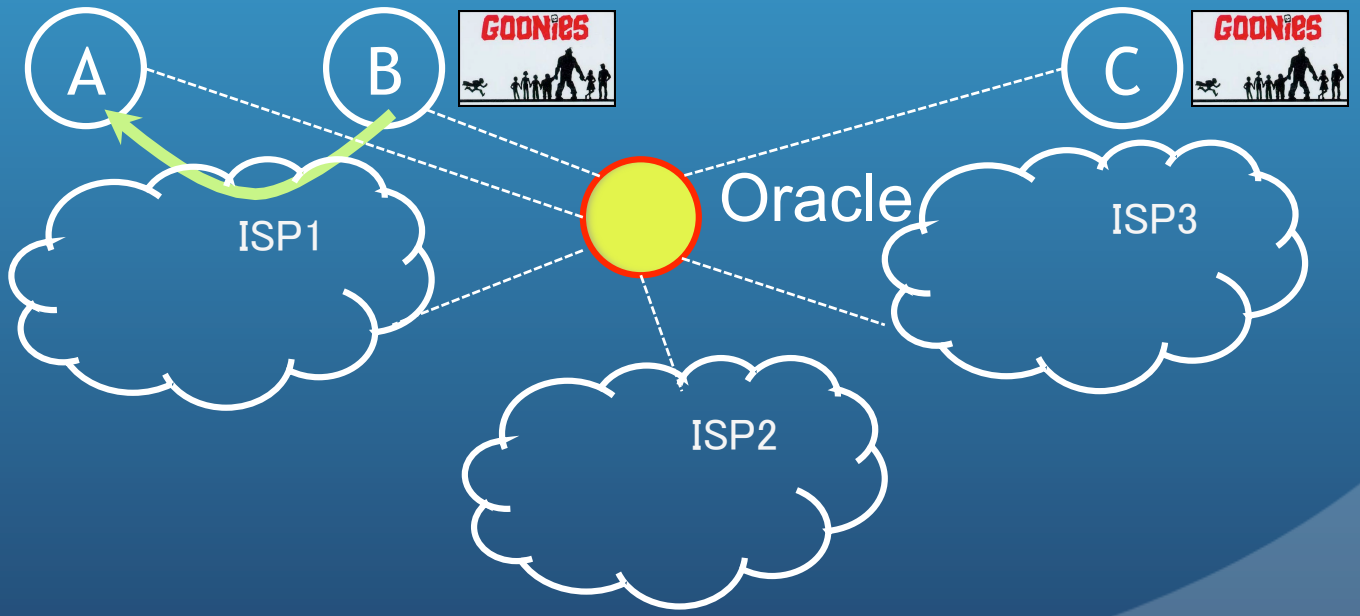

Open issue of MP-CDN

- MP-CDN works gracefully in theory or in a controlled environment
- There have been *no* general studies that address how peers can be incentivized in MP-CDN
- Question:
  - What is the motivation for peer nodes to participate in the system and contribute their resources?
- Some apps have such function embedded implicitly, e.g., BBC iPlayer, however, will end-users be willing to accept the implicit exploiting situation in general?

Our solution: a new business model

- ISP manages Oracle (PM server) and provides users with explicit incentive if they are cooperative to the system
  - Incentive can be virtual currency or some “points” that can be used in the system
  - Incentive can be fixed charge or calculated charge
- ISP provides CDN platform and Content provider and end-users use it
- Principle: End-users would prefer candy (incentive) rather than whip (bandwidth cap)

Key Idea

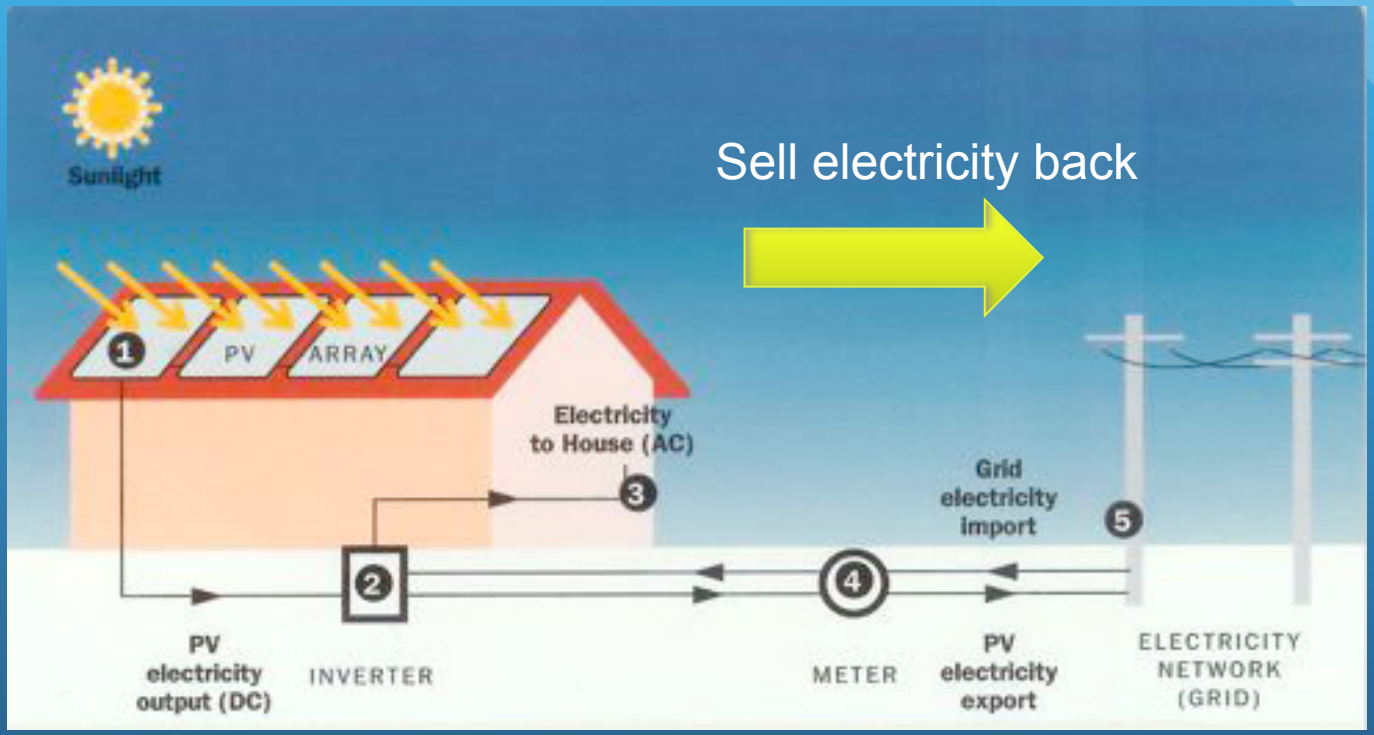
Sunlight
Sell electricity back
1
PV ARRAY
Electricity to House (AC)
3
Grid electricity import
5
2
4
PV electricity output (DC)
INVERTER
METER
PV electricity export
ELECTRICITY NETWORK (GRID)

Business model

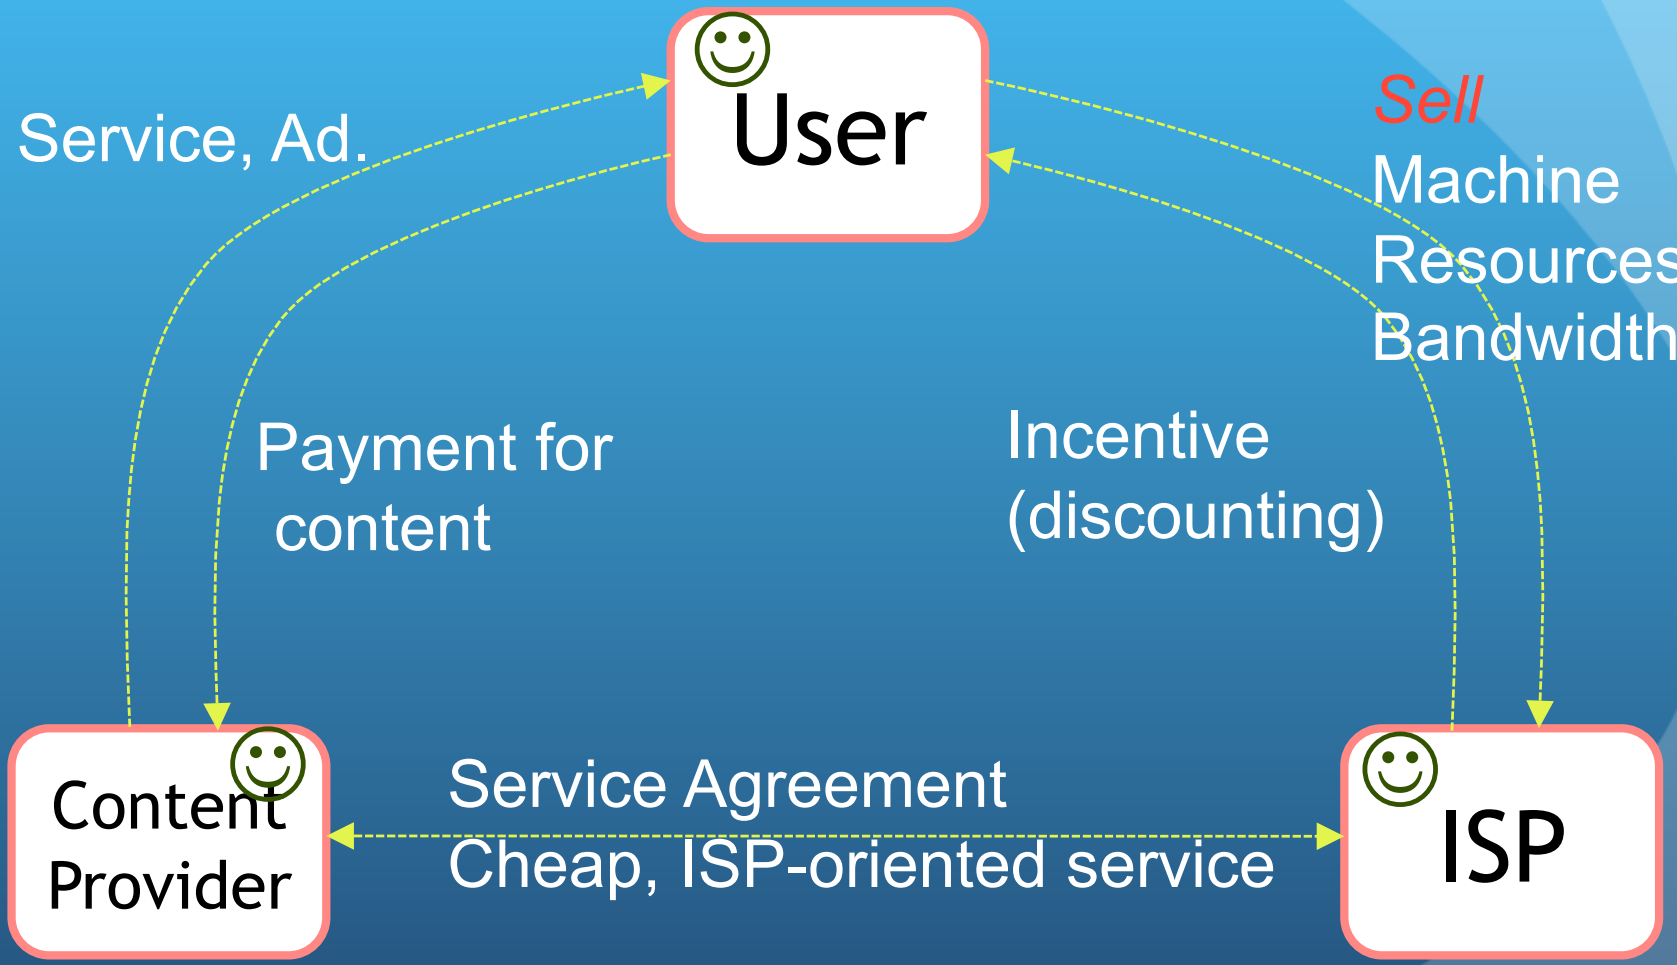

Can establish “Win-win-win” situation

Model of MP-CDN in an ISP

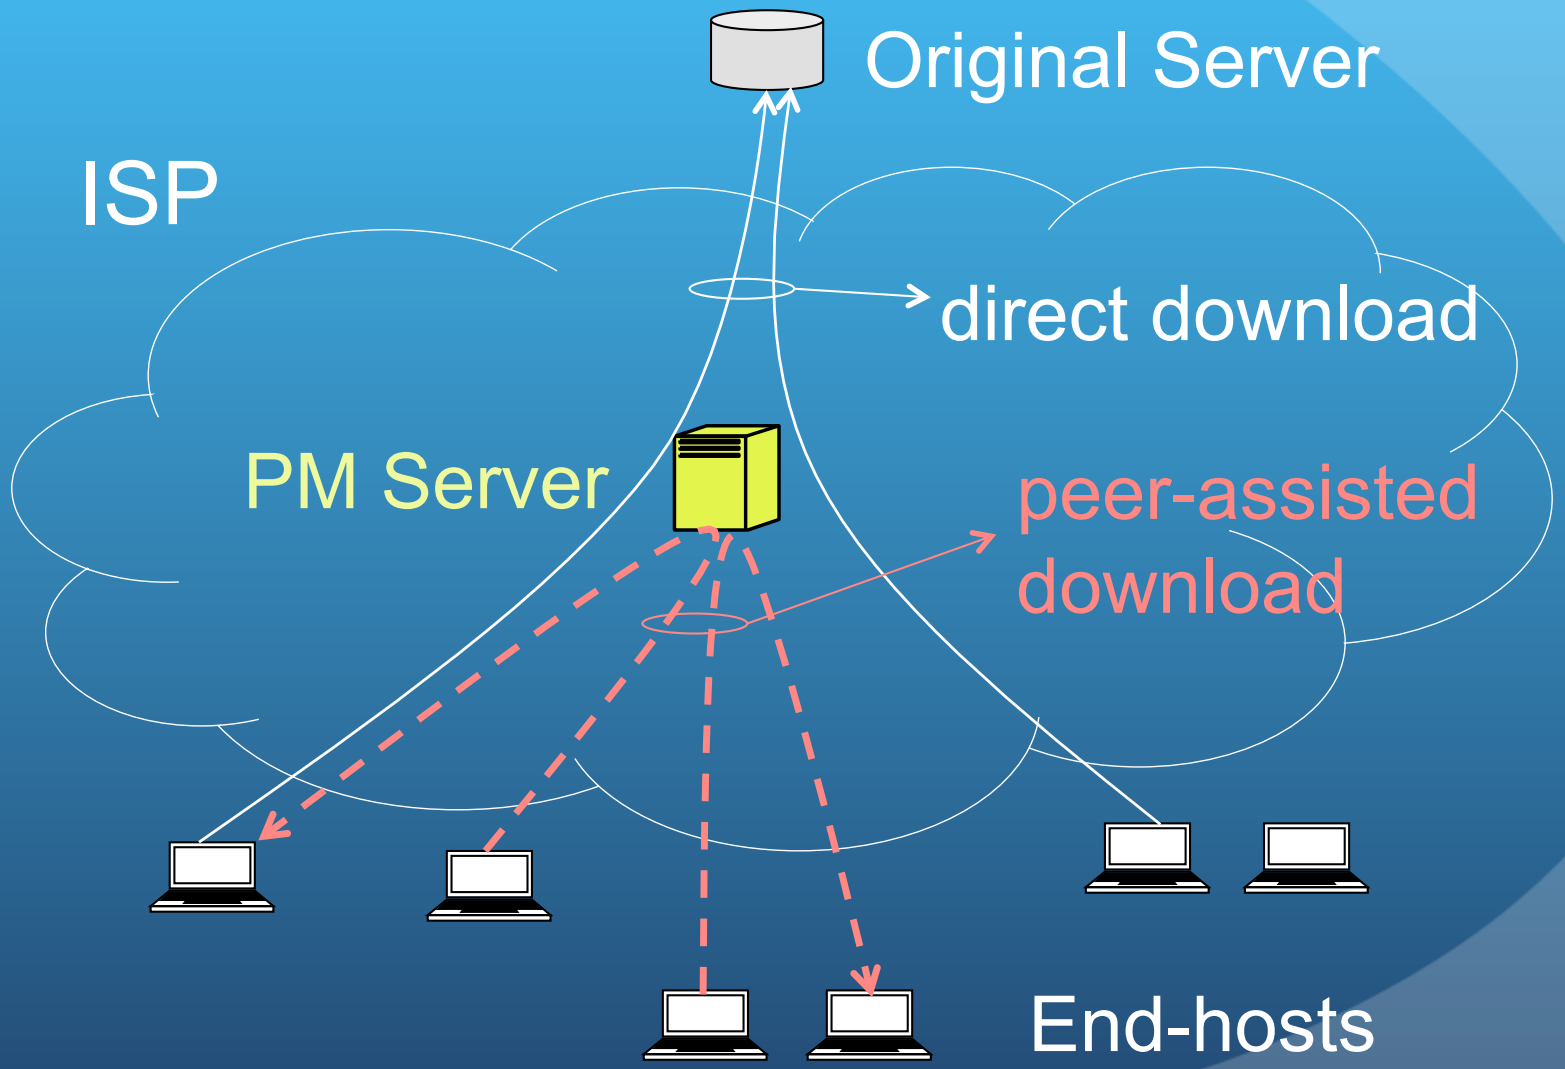

Roles of PM (peer mgmt) server

- Acts as “virtual cache server”
  - Keeps the list of peer nodes and their files
  - Storage space is given by peer nodes
- Acts as “Oracle”
  - Select peers according to the underlay network information
- Provides “AAA” functions
  - Accounting, authentication, and authorization

Analysis of the model

- Show the intrinsic trade-off between cache performance and cost for incentive
  - increase incentive
    → increase # of participating users
    → improve the cache performance → save the traffic cost
    BUT…
    → increase the cost for incentive as well
- Study how external factors such as #of users, #of files, storage capacity of each user, affect the cache performance.
- Goal: To obtain design implications

Simulation setup

- Nodes are identical
- Content requests arrive with the Poisson process
- Nodes keep content files with LFU cache algorithm
- The nodes and content files are fixed (no churn)
- There are no resource constraints on bandwidth and CPU of nodes

File access pattern

- Stretched Exponential Distribution (Discrete Weibull)
- Realistic model of modern web workload

$$p_i = y_i / \sum_i y_i$$

$$y_i = (-a \log(i) + b)^{1/c}$$

$$a = x_0^c, \text{ and } b = y_1^c$$

Acceptance of Incentive

- Logit model
- Given indentive of x, a node becomes cooperative with the probability:

$$p(x) = \frac{1}{1 + e^{-(\beta_1 + \beta_2 x)}}$$

Simulation Setup cont'

- 10 independent experiments for each parameter setting
- N: # of end-hosts
- m: # of content files
- S: cache capacity of each node (# of files)
- Simulation time T = 10000
  - Corresponds to a month in real time

Asymptotical property of system behavior (N=100, m=1000, S=5)

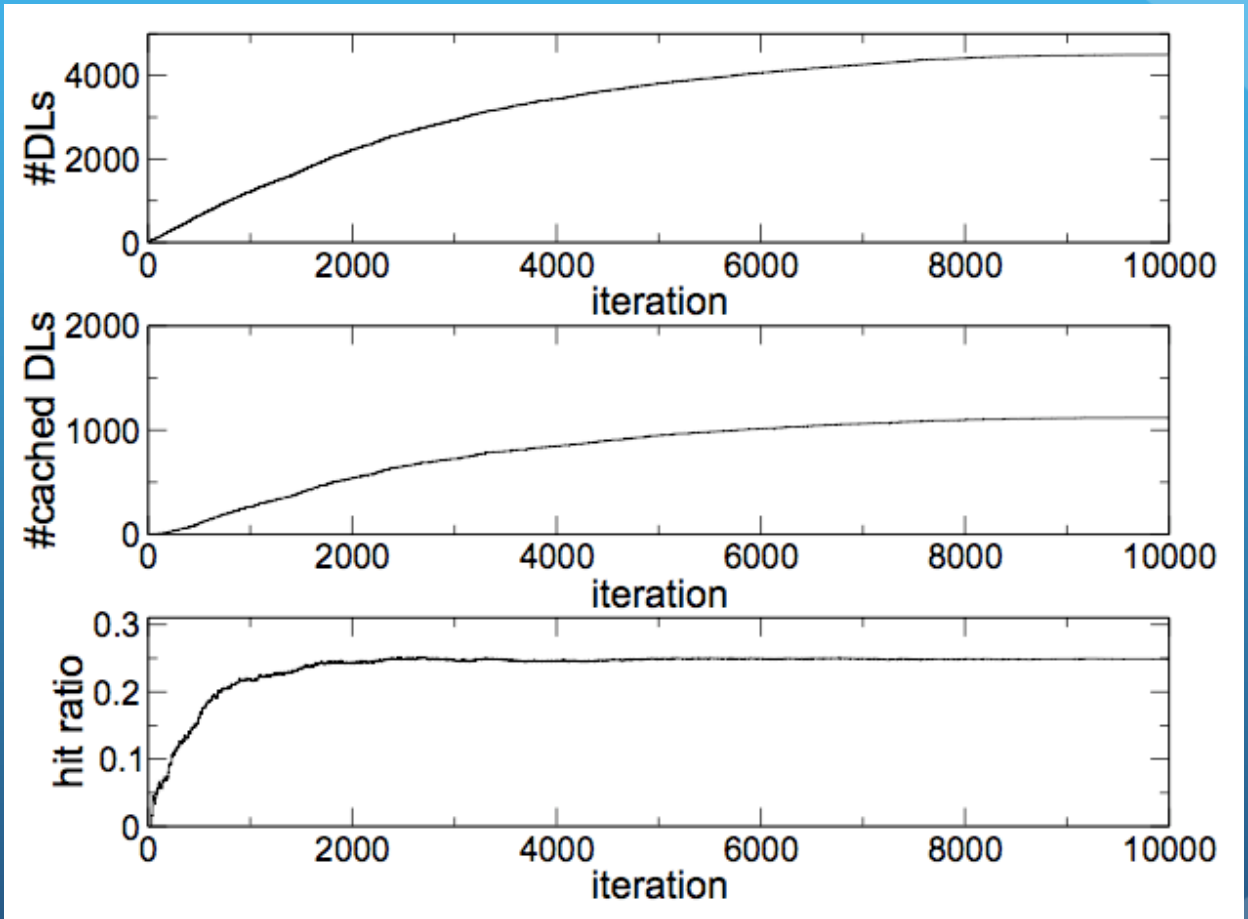

Role of incentive in the system

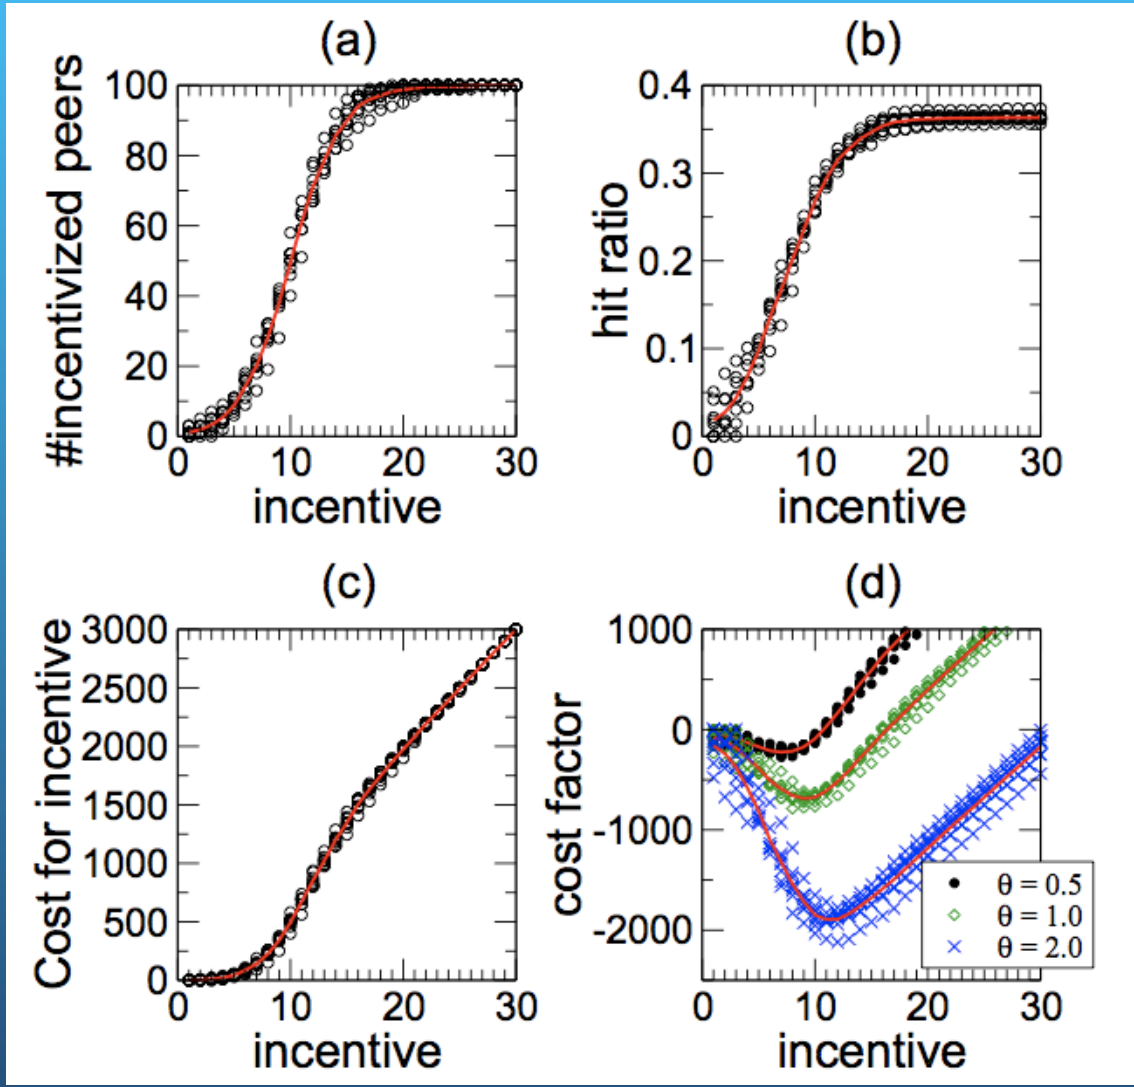

Cost factor
= Cost for incentive
- θ Cost of traffic

The effect of # of nodes (N)

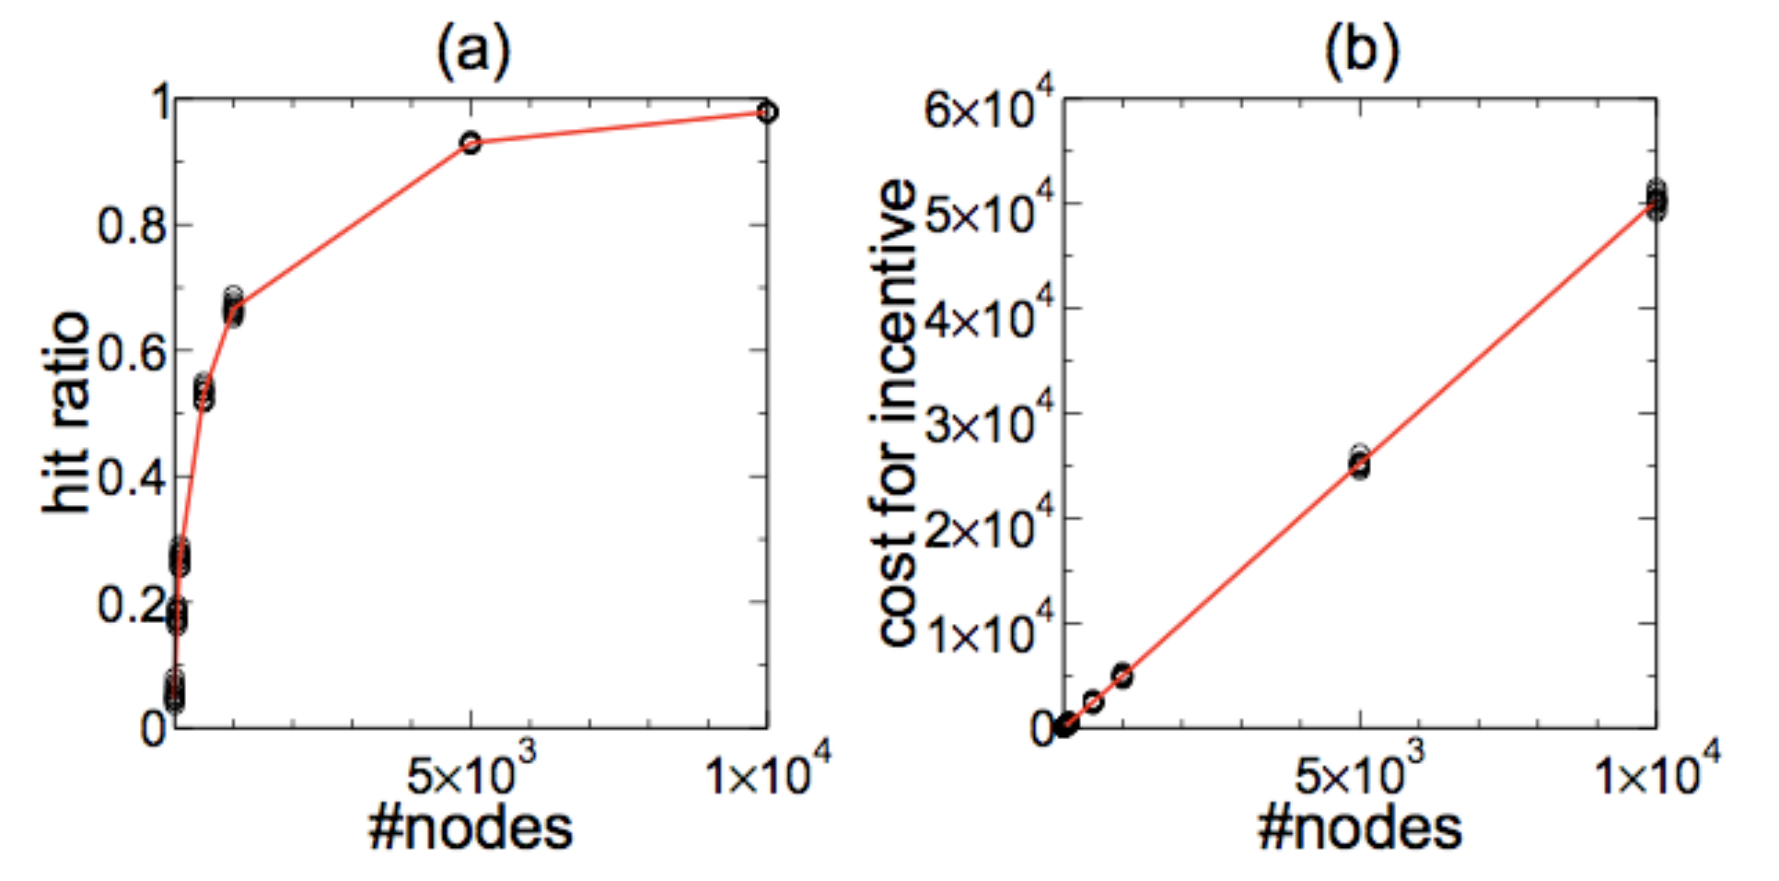

Incentive x = 10

Other external factors

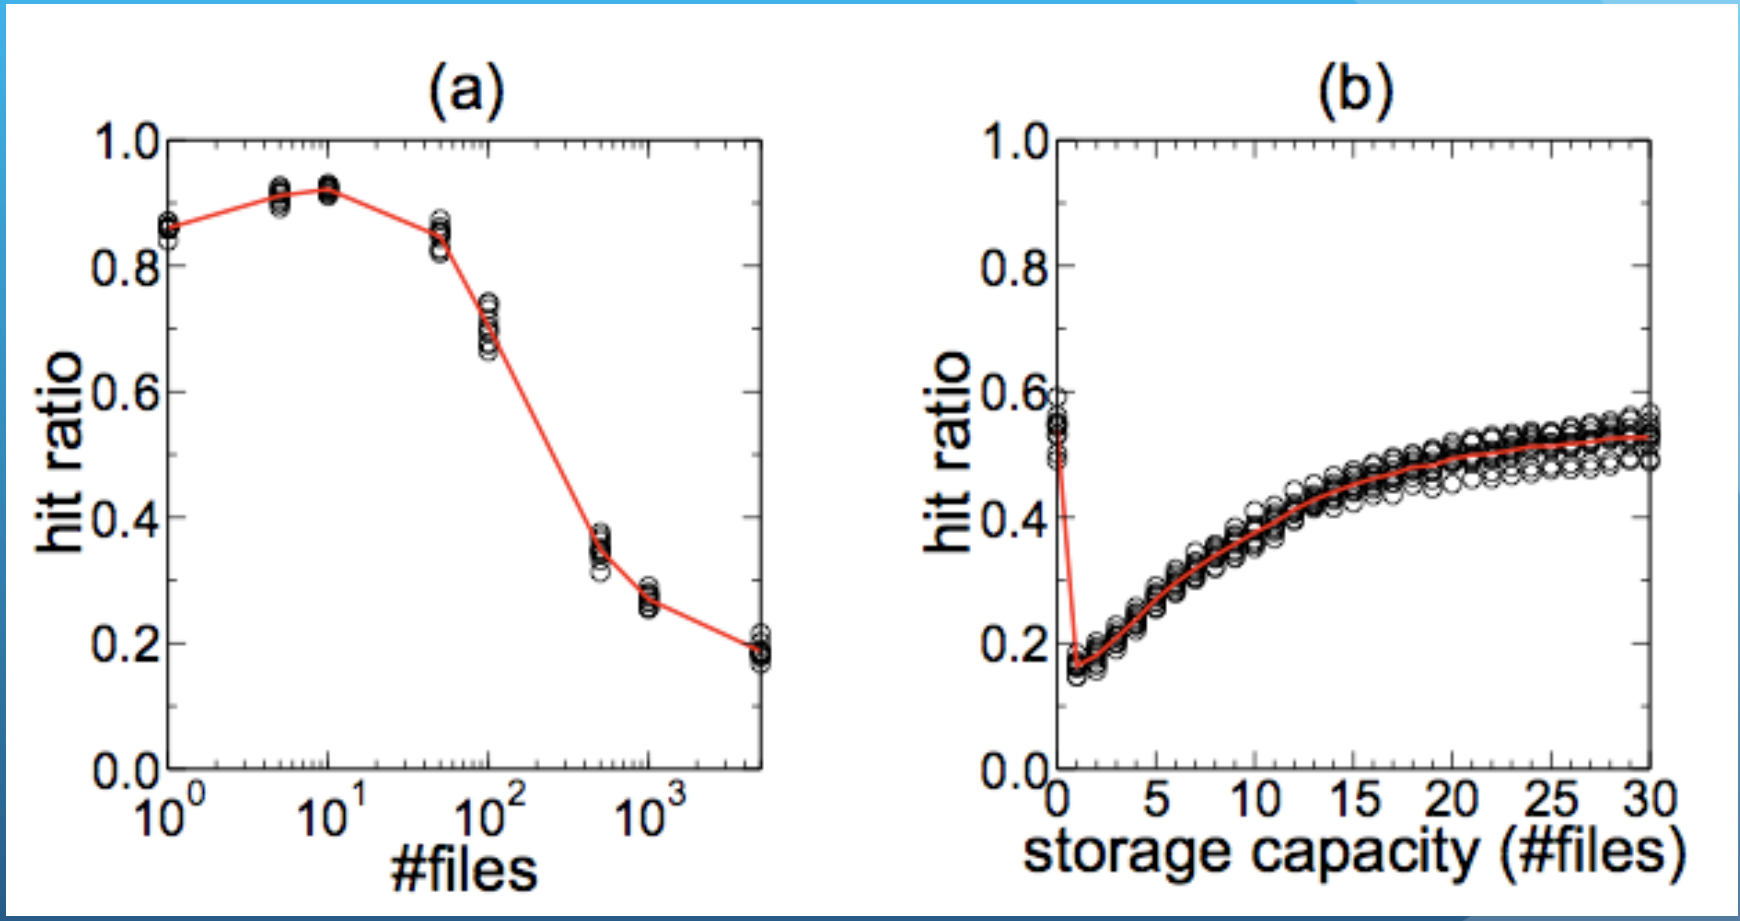

Design implications

- There exists *optimal* amount of incentive (should be designed carefully)
- It is better to keep # of nodes in a P2P NW small
- It is better to keep # of distinct content files small
- User storage capacity can be fairly small

Other issues:

- Scalability
  - Avoid making PM server be a single point of failure
- Underlying network structure
  - Upload bandwidth bottleneck (CATV)
- Privacy
  - Introduce some randomness in the peer selection
- More efficient content delivery
  - Introduce the pipelining model like BitTorrent, i.e., files are chopped into pieces and transferred simultaneously

Conclusion and future work

- A new business model – selling bandwidth back to ISP
  - Solve the incentive problem
- Design implications through the simulation analysis
- Discussed the potential issues that are essential to make the MP-CDN deployable
- Studying more realistic model, e.g., heterogeneous setting is for our future work