The Delicate Tradeoffs in BT-like Protocol Design: Performance VS Fairness

John C.S. Lui

Outline

- Background
- Mathematical Model
- Design Knob
- Simulation
- Conclusion

BitTorrent (BT) System

- A Peer-to-Peer (P2P) file distribution application, created by Bram Cohen.
- Designed to distribute large content (Linux distribution) without saturating servers and bandwidth resources.
- BitTorrent traffic accounts for ~35% of all traffic on the Internet today.
- Key idea of BT:
  - File is divided into small pieces
  - Choking algorithm to make peers cooperative

Characterizing Peers

- Peers in the system are *heterogeneous*
  - "*Resourceful peers*": peers with higher up/ down link bandwidth
  - "*Thin peers*": peers with lower up/down link bandwidth
- Peers in the system are *selfish*
  - Incentive Mechanism is necessary to prevent free-riding

Dilemma of Protocol Design

Dilemma of Protocol Design

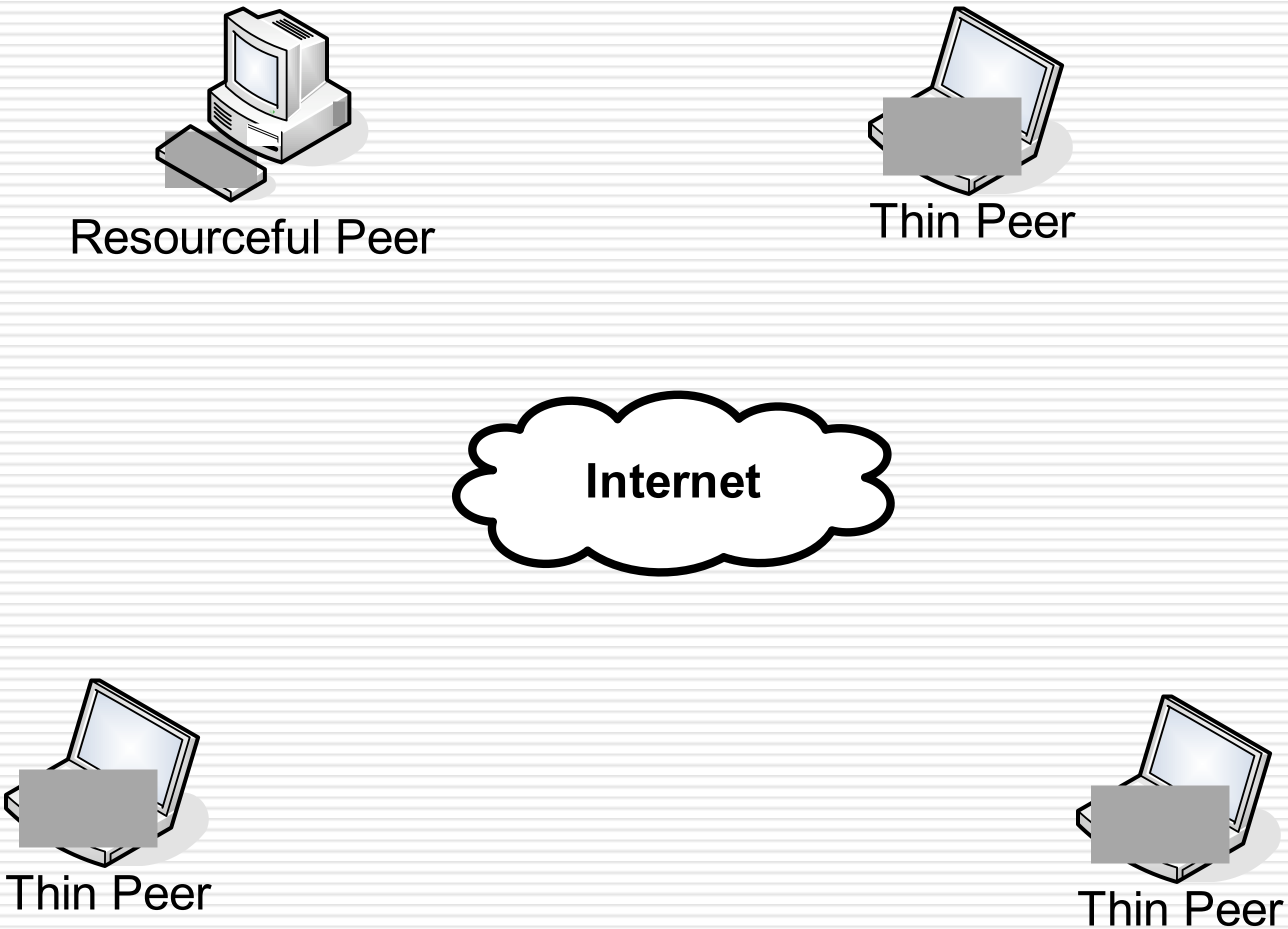

Dilemma of Protocol Design

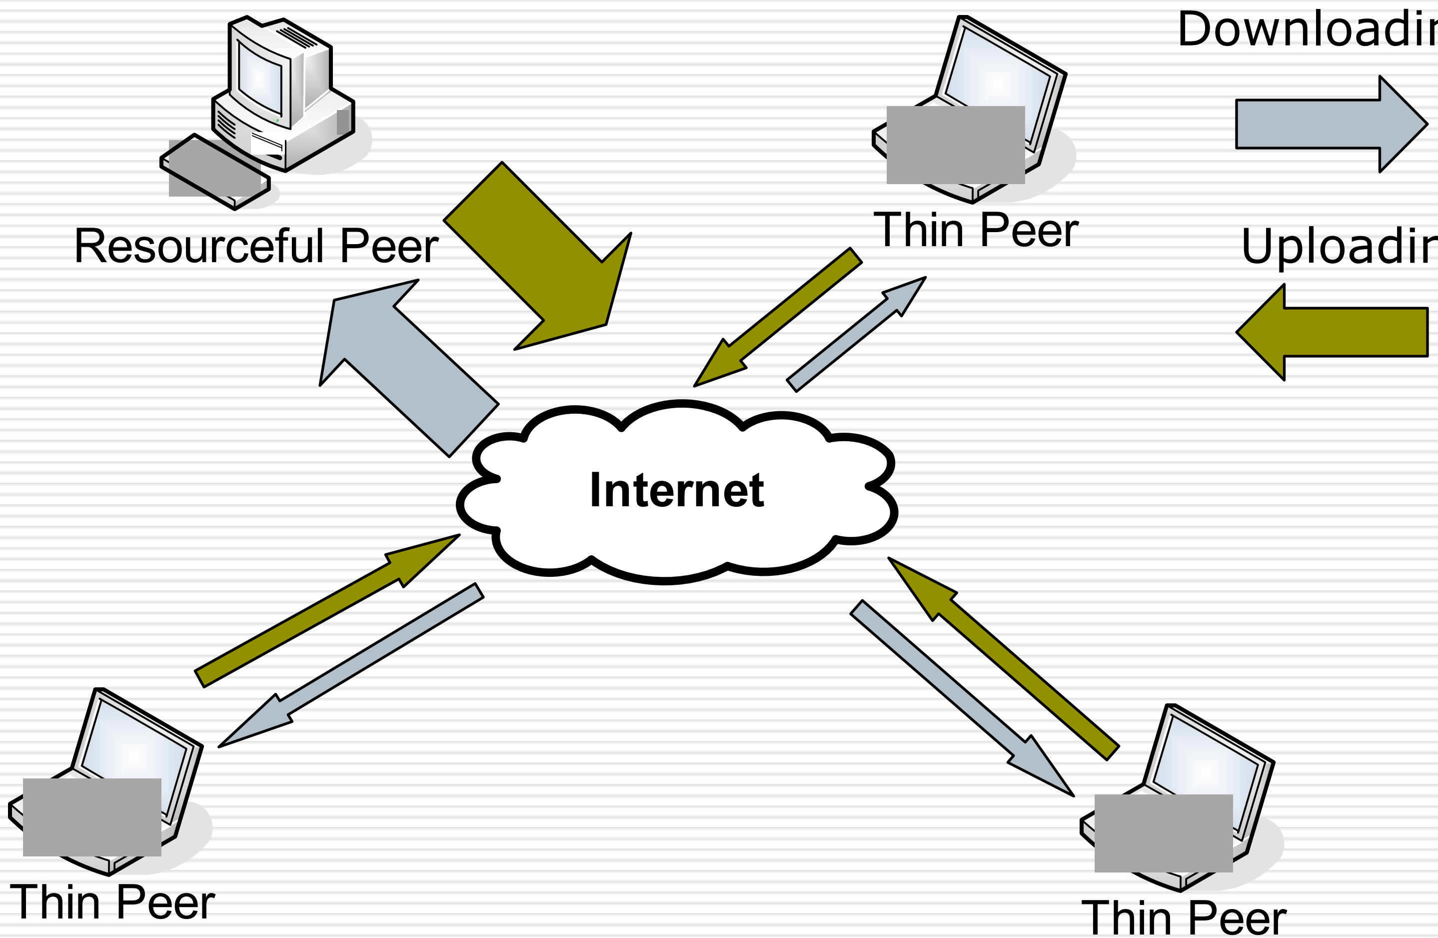

Dilemma of Protocol Design

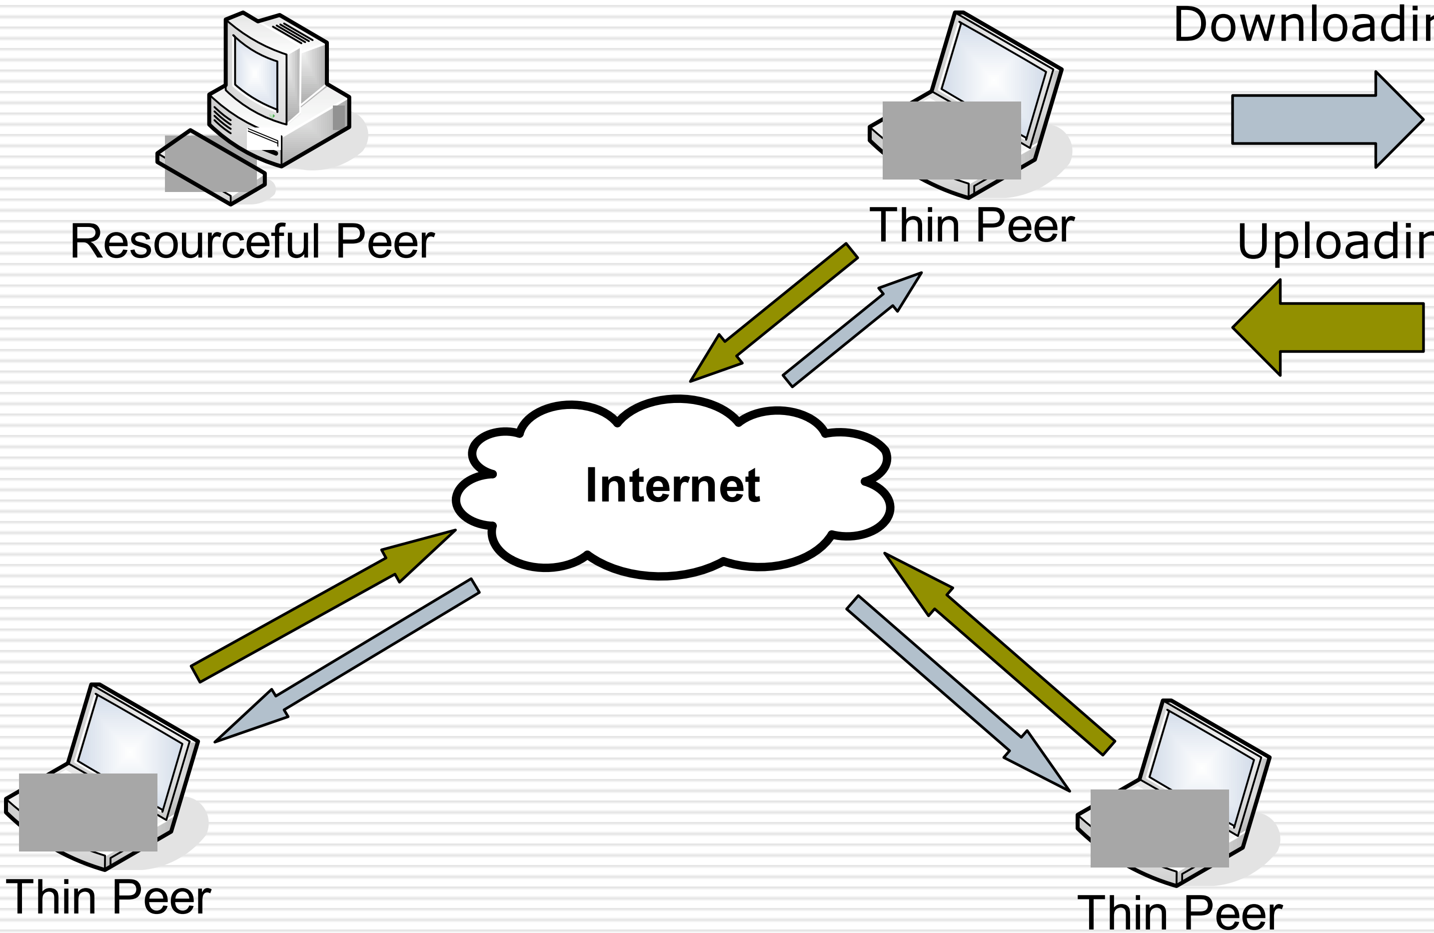

Dilemma of Protocol Design

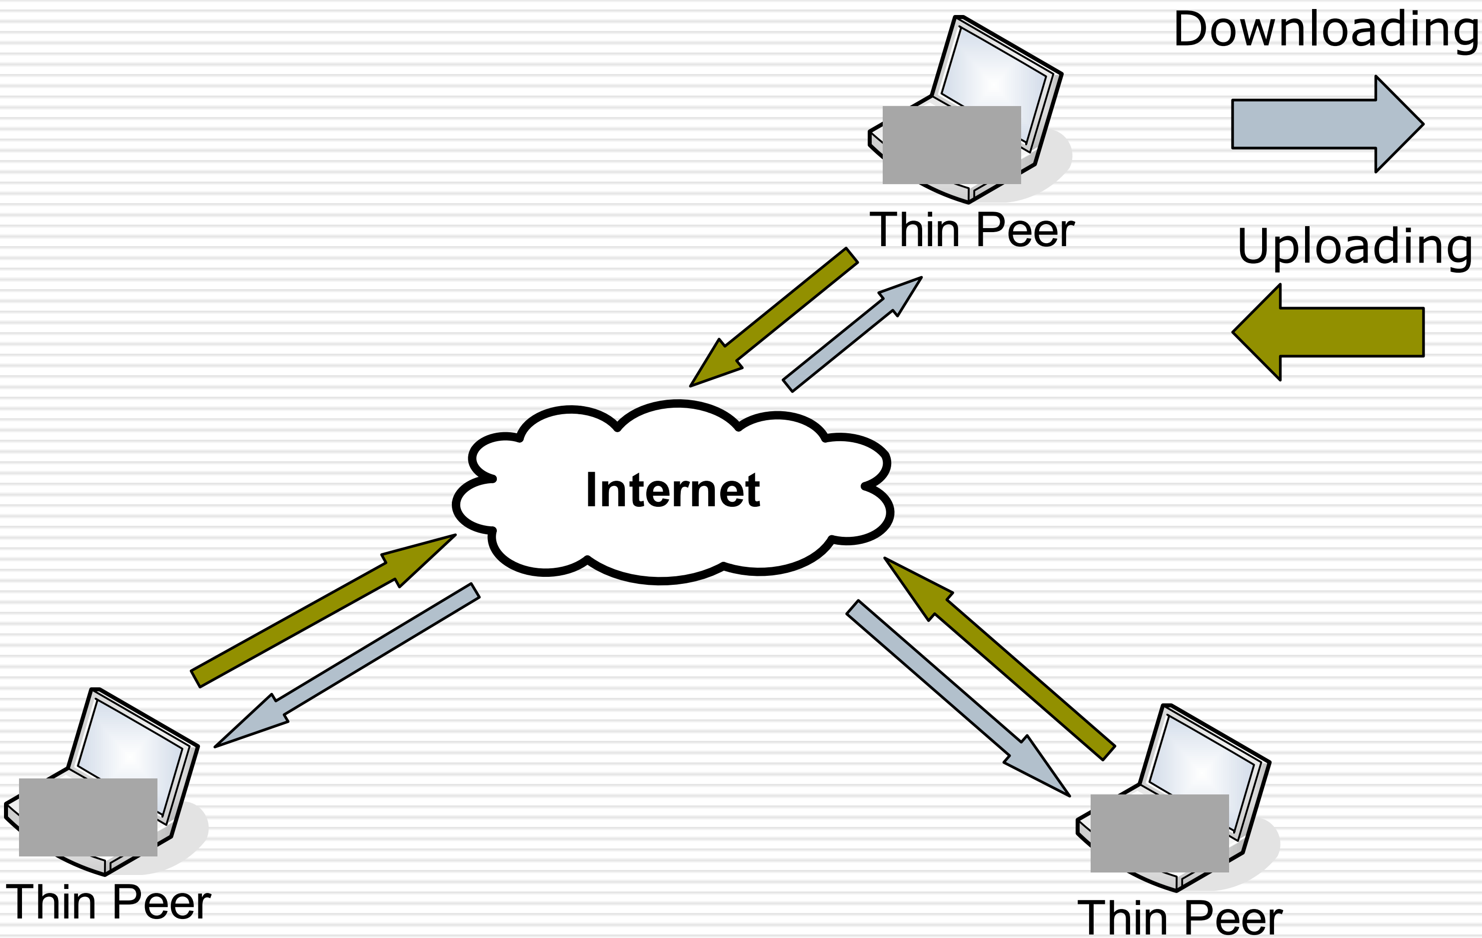

Dilemma of Protocol Design

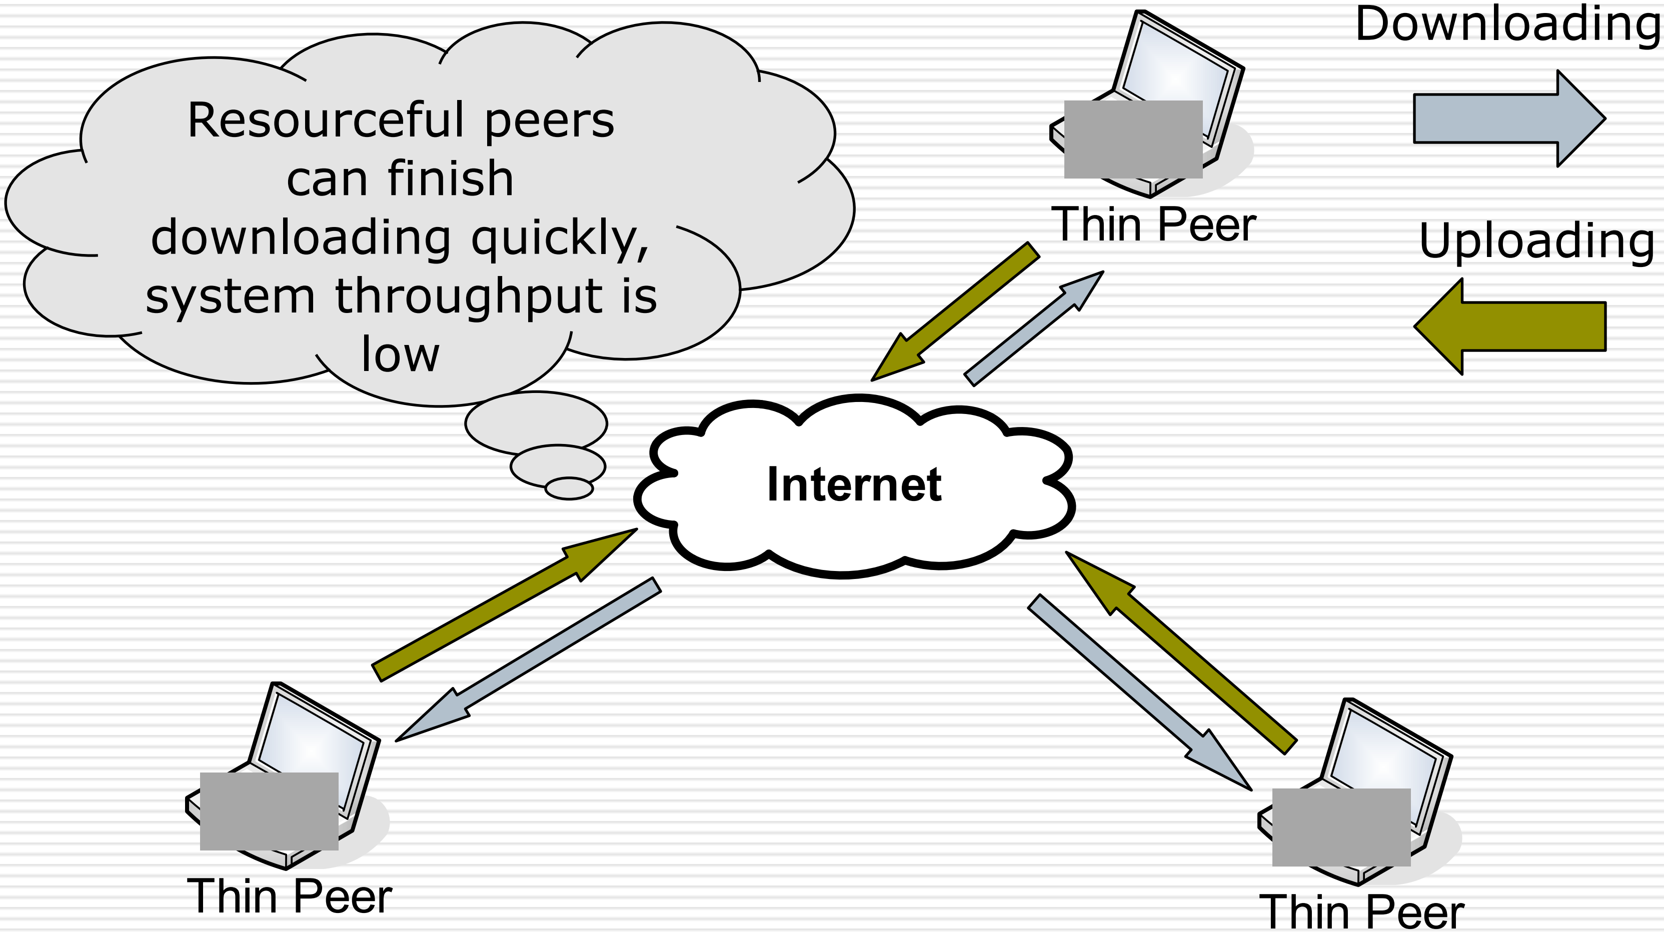

Dilemma of Protocol Design

Dilemma of Protocol Design

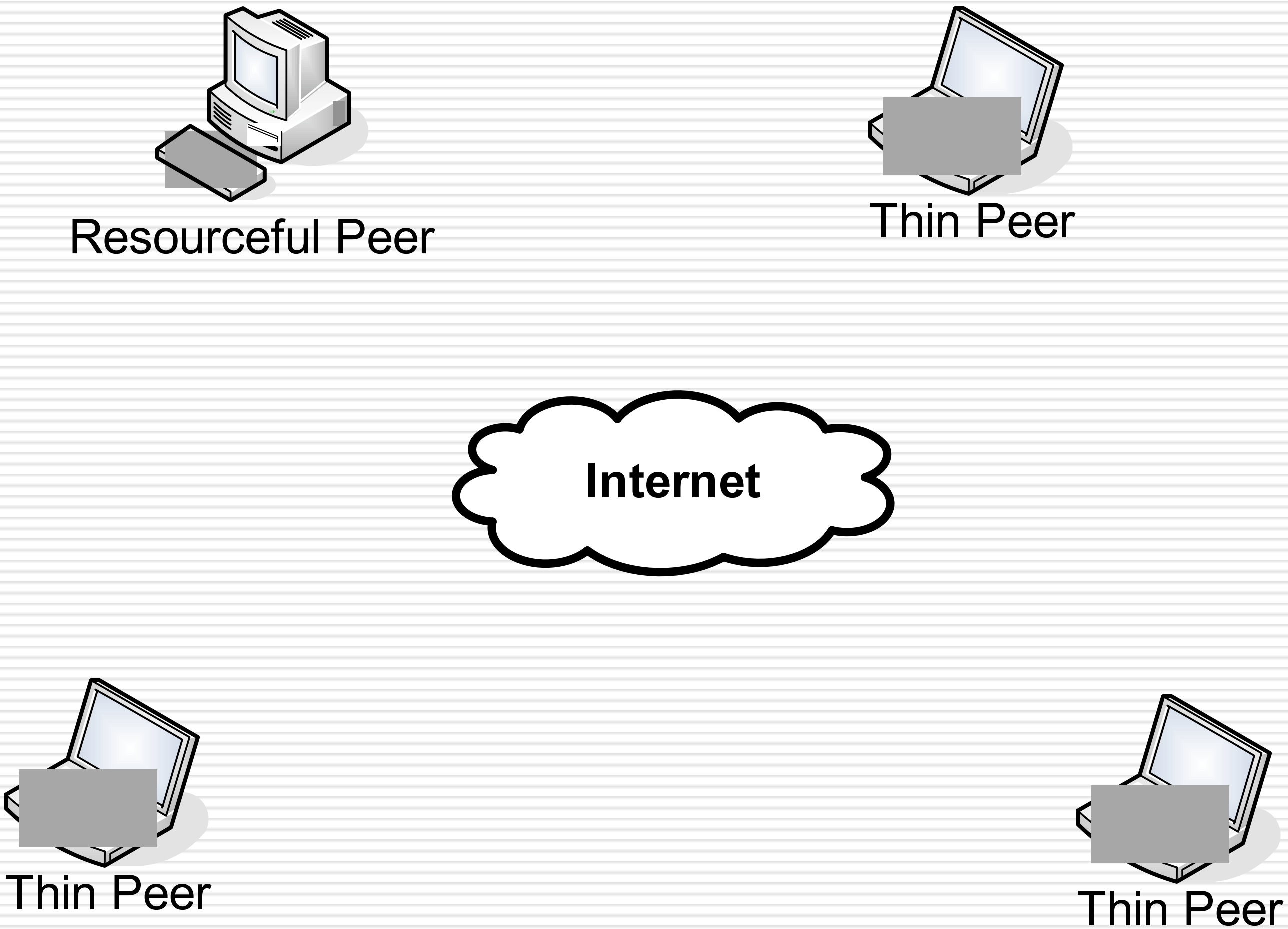

Dilemma of Protocol Design

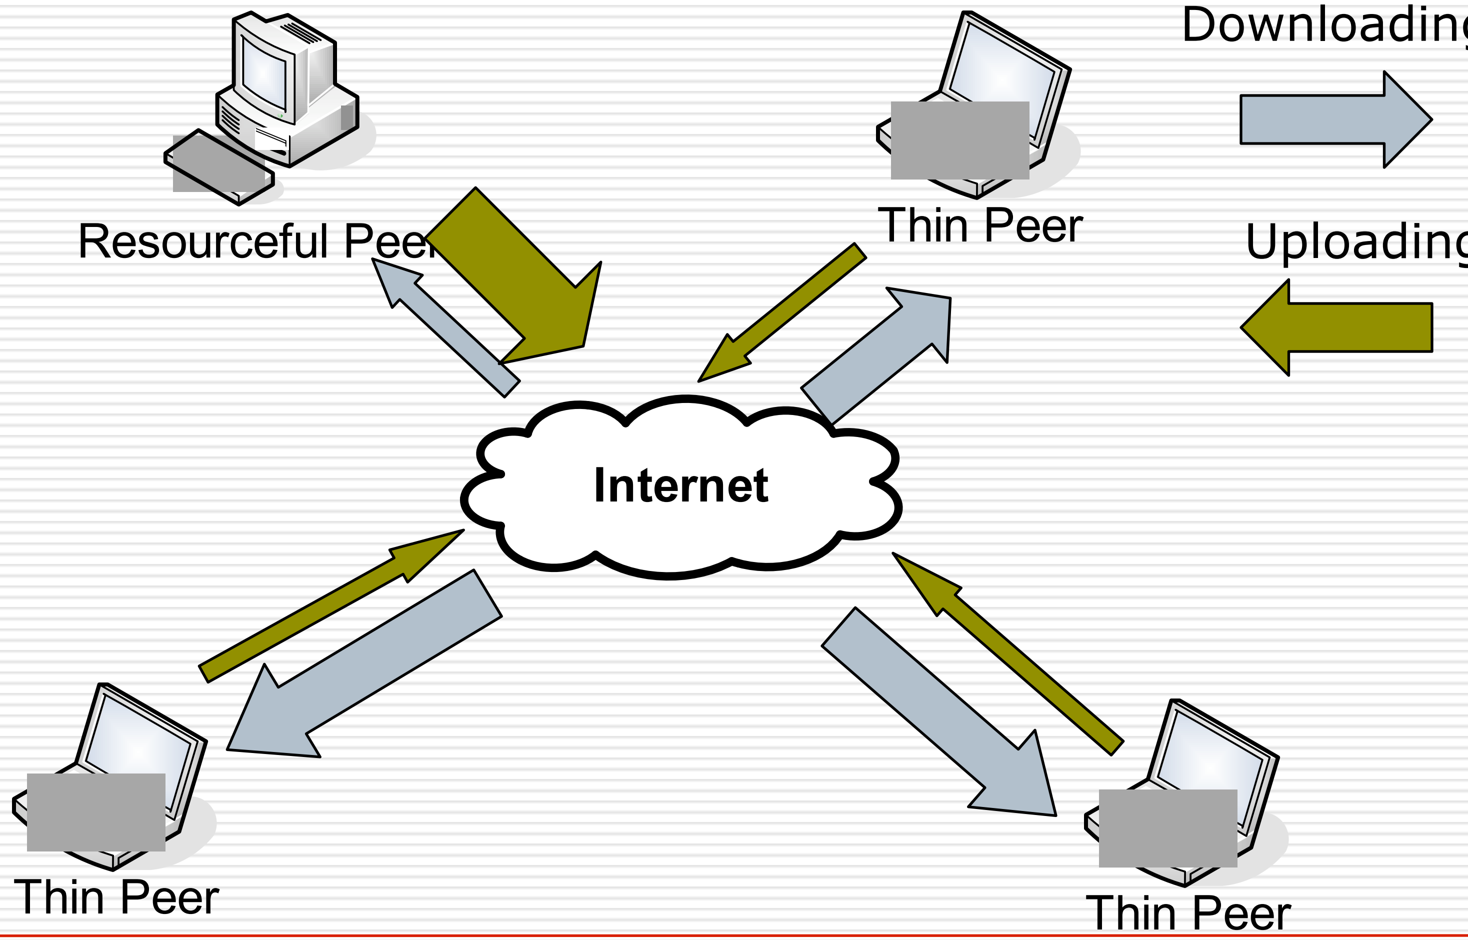

Dilemma of Protocol Design

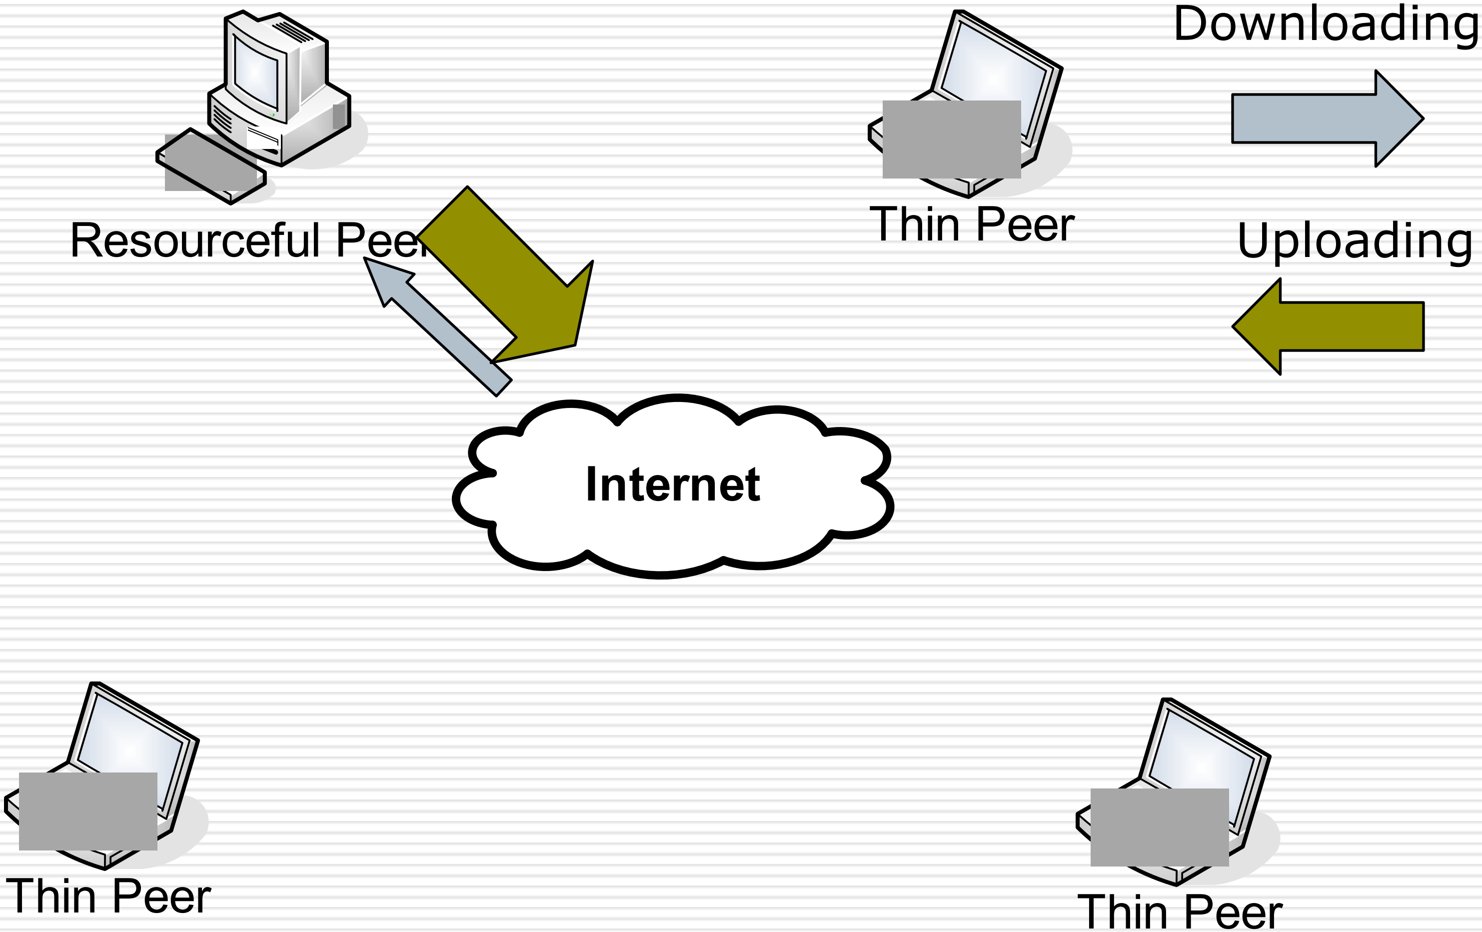

Dilemma of Protocol Design

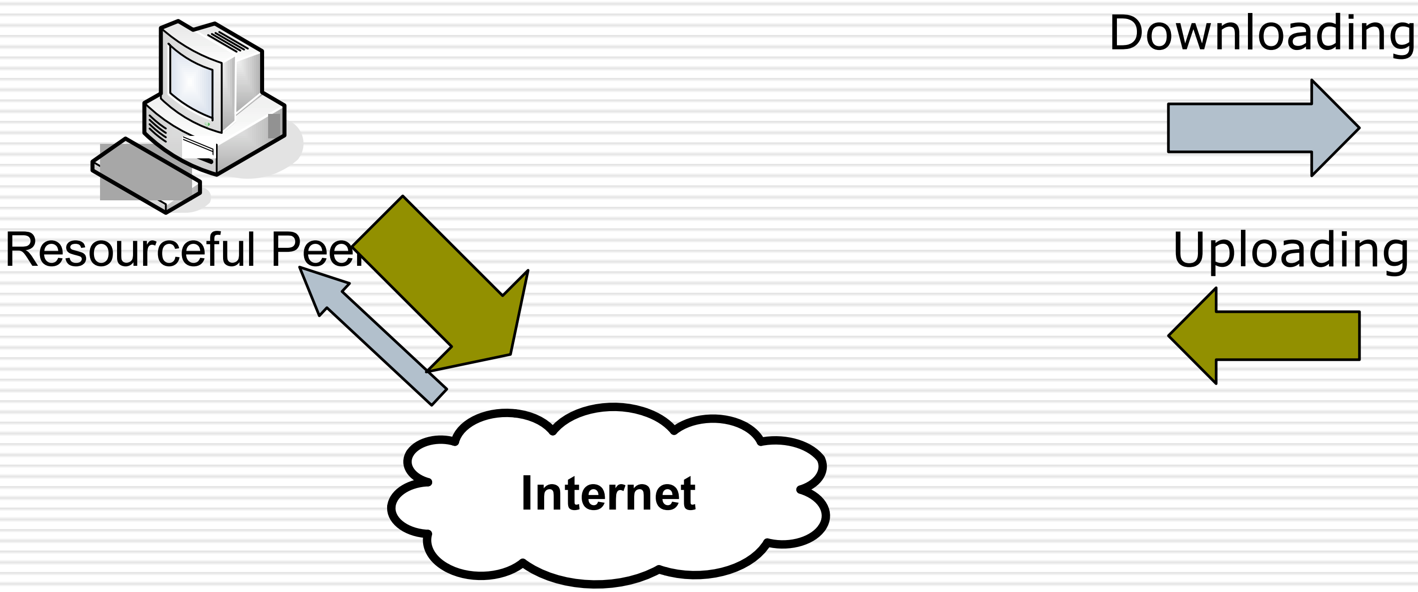

Dilemma of Protocol Design

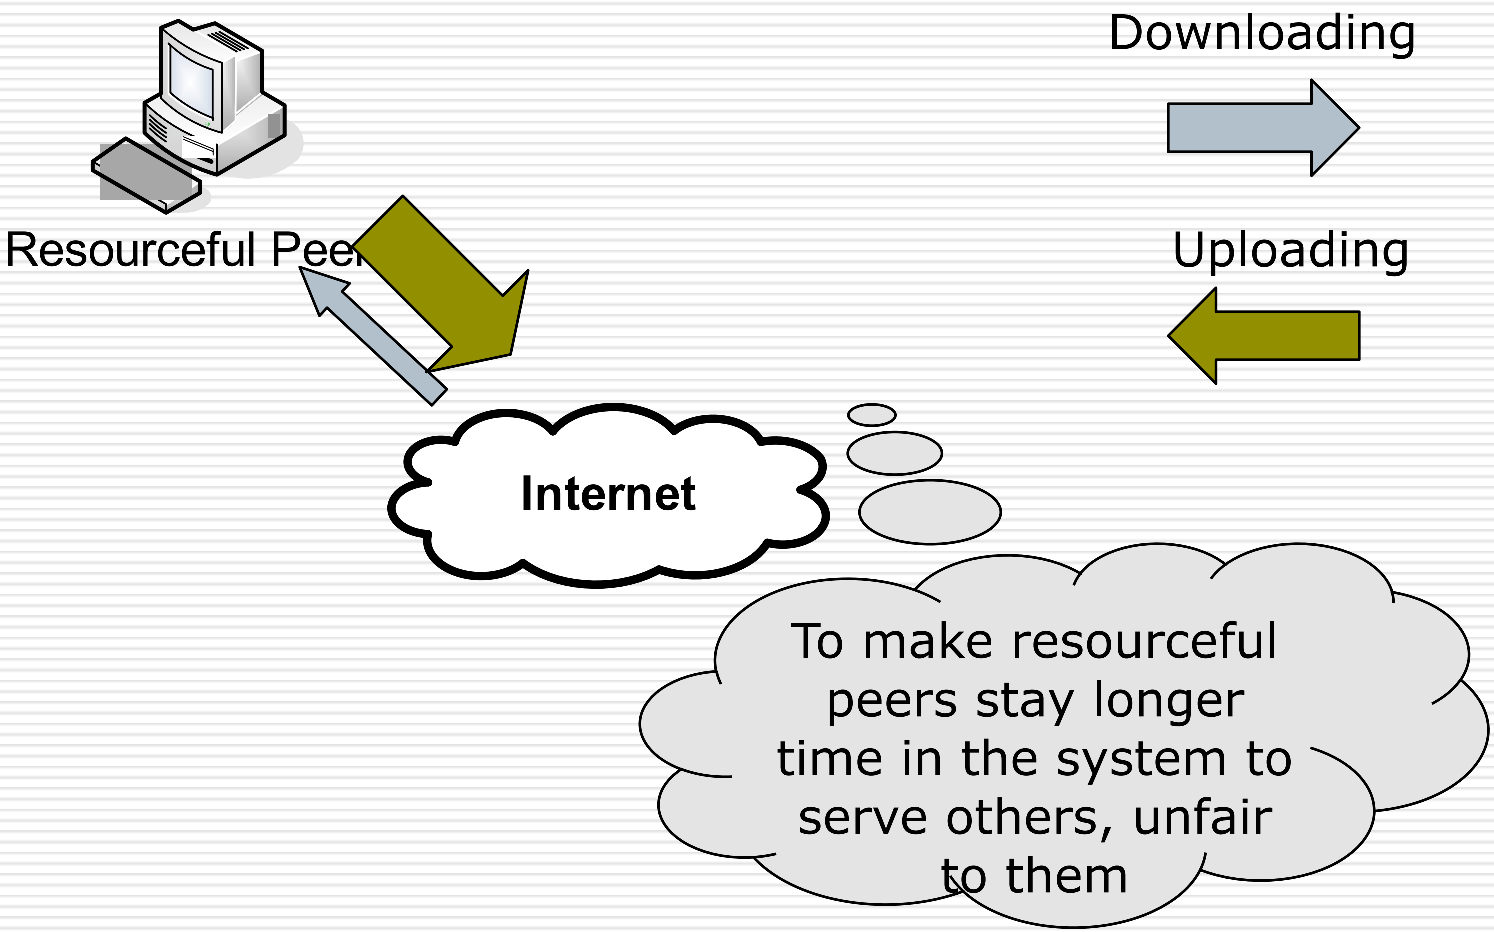

Design Object?

Design Object?

To encourage resourceful peers so that they can obtain higher download rate

Design Object?

To encourage resourceful peers so that they can obtain higher download rate

Versus

To make resourceful peers stay in system to improve the downloading of others

Mathematical Model

- N types of peers, for type-i:
  - Uploading capacity: U_i
  - Downloading capacity: D_i
  - Feasible uploading rate: $u_i \leq U_i$
  - Feasible downloading rate: $d_i \leq D_i$
  - Probability of a new peer to be type-i: p_i

Uplink Sharing[1]

- Limitation of system throughput is uploading
- Bottleneck is assumed not the network
- Lower bound to disseminate a file is studied in [1]
- Arrival and departure of peers are considered in our model

[1] J. Mundinger and et al, Analysis of Peer-to-Peer File Dissemination amongst users of different upload capacities.

Fairness Metrics

- **"*share ratio*" of type-i peer:**

$$c_i = \frac{u_i}{d_i}$$

  - When share ratio = 1, type-i peer provides as much service as it receives
  - When share ratio = 0, free ridering
- **Fairness Index to measure share ratios of all peers:**

$$\mathcal{F} = \frac{(p_1 c_1 + \ldots + p_n c_n)^2}{p_1 c_1^2 + \ldots + p_n c_n^2} = \frac{1}{p_1 c_1^2 + \ldots + p_n c_n^2}$$

Performance Metric:

- In a P2P system, throughput is related to the peers staying in the system
- The service differentiation policy will affect the average downloading time.
- Average downloading time:

$$T = \frac{N_1 + \ldots + N_n}{\lambda} = \frac{p_1}{d_1} + \ldots + \frac{p_n}{d_n}.$$

To Achieve Optimal Average Downloading Time

- Solve the optimization problem:

$$\text{Min} \qquad T = \frac{p_1}{d_1} + \ldots + \frac{p_n}{d_n} \qquad \text{s.t.} \qquad p_1 \frac{U_1}{d_1} + \ldots + p_n \frac{U_n}{d_n} = 1,$$

$$0 \le d_i \le D_i, i = 1, \ldots, n$$

- The Solution:
  - Type-1:

$$d_1 = \frac{p_1 U_1}{1 - \sum_{i=2}^{n} p_i \frac{U_i}{D_i}},$$

Type-i:

$$d_i = D_i, \quad i = 2, \ldots, n.$$

- Insights:
  - First serve less resourceful peers as much as possible
  - Then serve most resourceful peers

To Achieve Optimal Fairness

- All peers have the same share ratio
- Rate assignment:

$$\text{Type-i peer:} \qquad d_i = u_i = U_i.$$

- Insights:
  - Every peer just gets as much as it contributes

To Achieve Max-min Fairness

- Rate assignment:

$$\text{Type } i \text{ peer:} \qquad d_i = d \equiv p_1 U_1 + \ldots + p_n U_n.$$

- Insights:
  - Every peer receives the same service

Three Rate Assignments

	Fairness Index	Av. Download Time
Optimal Performance	$\frac{1}{p_1(\frac{1-\sum_{i=2}^n p_i U_i/D_i}{p_1})^2 + \sum_{i=2}^n p_i(\frac{U_i}{D_i})^2}$.	$\frac{1}{U_1} + \sum_{i=2}^{n} \frac{p_i}{D_i} \frac{U_1 - U_i}{U_1}$
Max-min	$\frac{(\sum_{i=1}^n p_i U_i)^2}{\sum_{i=1}^n p_i U_i^2}$	$\frac{1}{p_1 U_1 + \ldots + p_n U_n}$.
Optimal Fairness	1	$\frac{p_1}{U_1} + \ldots + \frac{p_n}{U_n}$

Trade-off:

- In terms of average downloading time:

$$T_{opt} < T_{mm} < T_{fair}$$

- In terms of fairness

$$\mathcal{F}_{opt} < \mathcal{F}_{mm} < \mathcal{F}_{fair}$$

Numerical Illustration

- 2-Types of peers:

	Resourceful Peer	Thin Peer
U	4	2
D	5	6
p	0.4	0.6

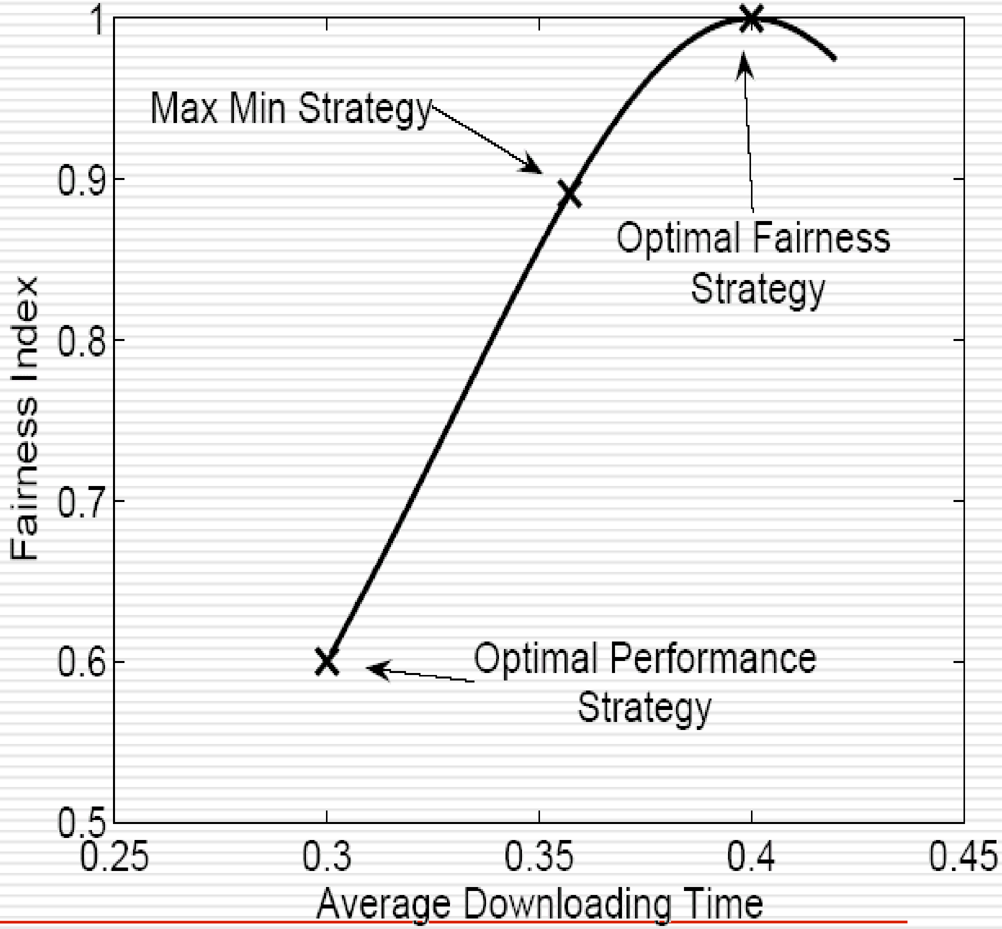

Implementation

- Feasible rate assignment can be realized by centralized algorithm
  - Require global knowledge
  - Require centralized scheduler
- Distributed algorithm?
  - Easy to implement
  - Easy to adjust fairness/performance

Two Uploading Strategies

- Selective uploading
  - Provide uploading service to the top n_s peers based on their downloading rates
  - Similar to ‘tit-for-tat’ used by BT
- Non-discriminative uploading
  - Randomly choose n_a peers to provide uploading
  - Similar to ‘optimistic-unchoking’ in BT

Selective Uploading

- Formulate the peer selection as a game.
- In Nash equilibrium, downloading rate of peer i:

$$d_i \approx u_i.$$

- the optimal fairness is achieved!

Non-discriminative Uploading

- Every peer randomly chooses n_a peers to serve
- The downloading rate of peer i:

$$d_i \approx \frac{\sum_{j \in N} u_j}{n} = \bar{u}$$

- Max-min fairness is achieved

Design Knob

- Use (n_s, n_a) as the design knob

$$d_i = \frac{n_s}{n_s + n_a} u_i + \frac{n_a}{n_s + n_a} \bar{u}.$$

- Official BT protocol:

$$n_s = 4, n_a = 1$$

- Official BT emphasis on fairness

Revisit Optimistic-unchoking

- Optimisitic-unchoking (OU) is more than the complement of ‘tit-for-tat’ to find potential connections
- OU is also an approach to improve the system performance

Performance Evaluation 1: Nash Equilibrium

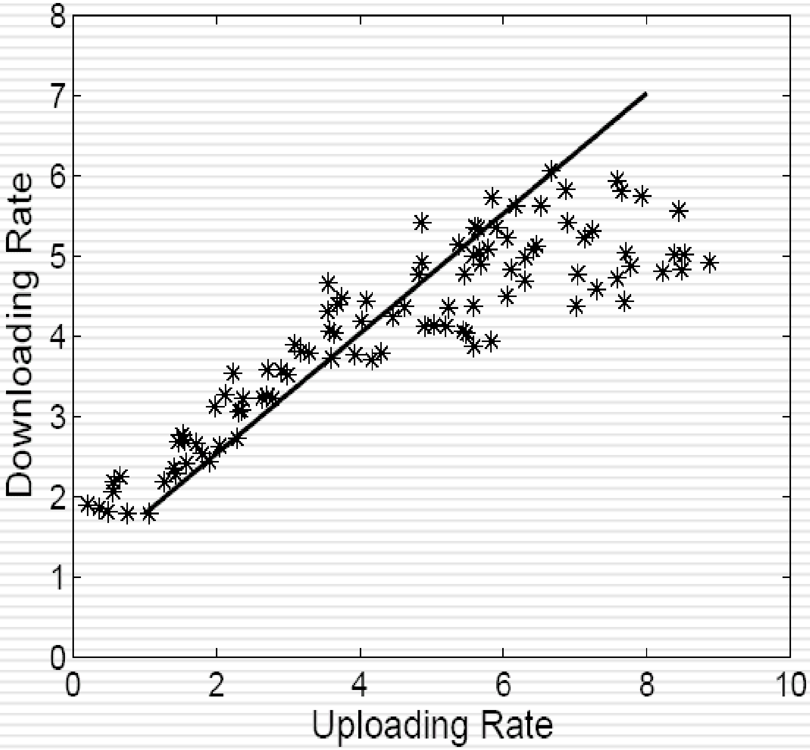

(a) Simulation

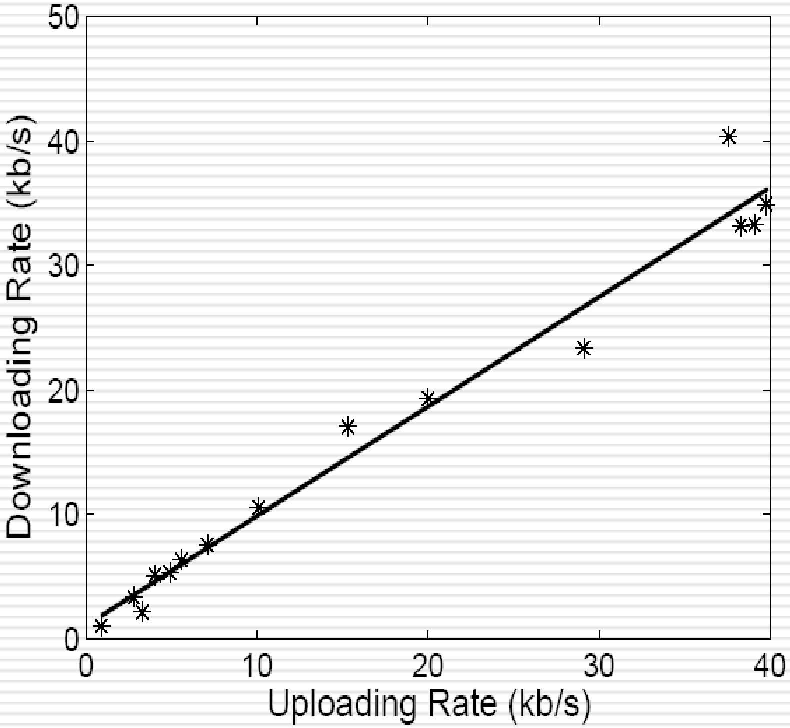

(b) Real BT's Measurements

Performance Evaluation: Design Knob

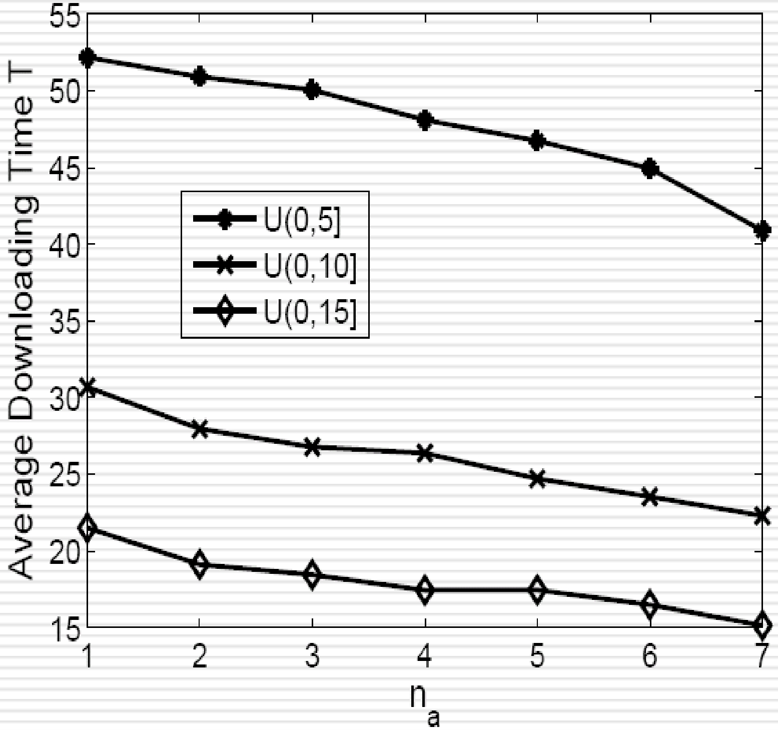

(a) n_a vs T

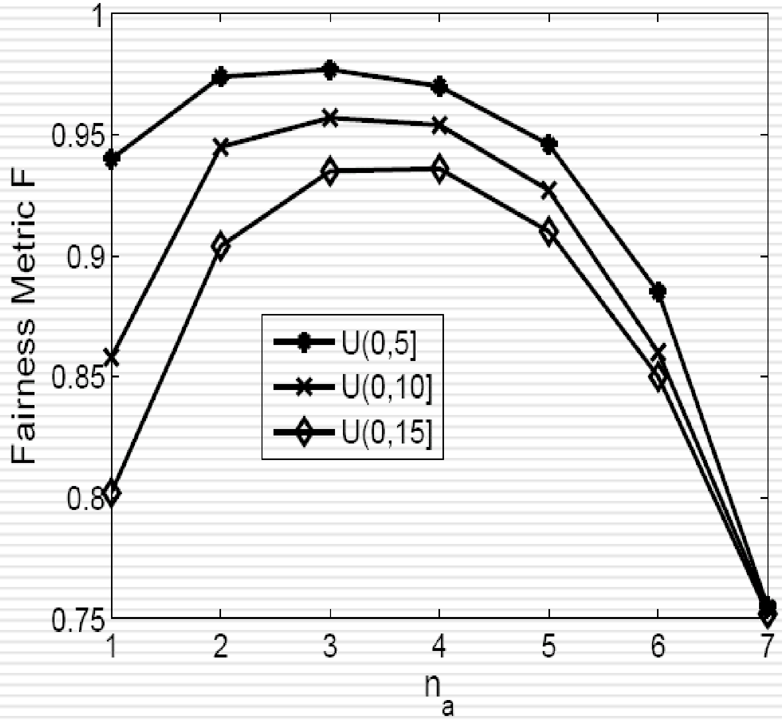

(b) n_a vs $\mathcal{F}$

Conclusion

Conclusion

- The trade-off between performance and fairness for a BT-like file sharing protocol

Conclusion

- The trade-off between performance and fairness for a BT-like file sharing protocol
- BT protocol is only one particular point in the whole design space

Conclusion

- The trade-off between performance and fairness for a BT-like file sharing protocol
- BT protocol is only one particular point in the whole design space
- Deeper understanding of “tit-for-tat” and “optimistic-unchoking” used by BT

Conclusion

- The trade-off between performance and fairness for a BT-like file sharing protocol
- BT protocol is only one particular point in the whole design space
- Deeper understanding of "tit-for-tat" and "optimistic-unchoking" used by BT
- Design knob to adjust performance and fairness of the system

The End…

- Q & A