


## 1 Decreasing Inter-ISP Traffic


- Inter-ISP traffic is a source of cost for ISPs
- Two directions to decrease inter-ISP traffic [1]:
  1. Locality-awareness based on
    - ISP-provided information (ALTO, P4P, “Oracle”)
    - Measurements or third-party provided information (ONO, BNS)
  2. P2P Caching



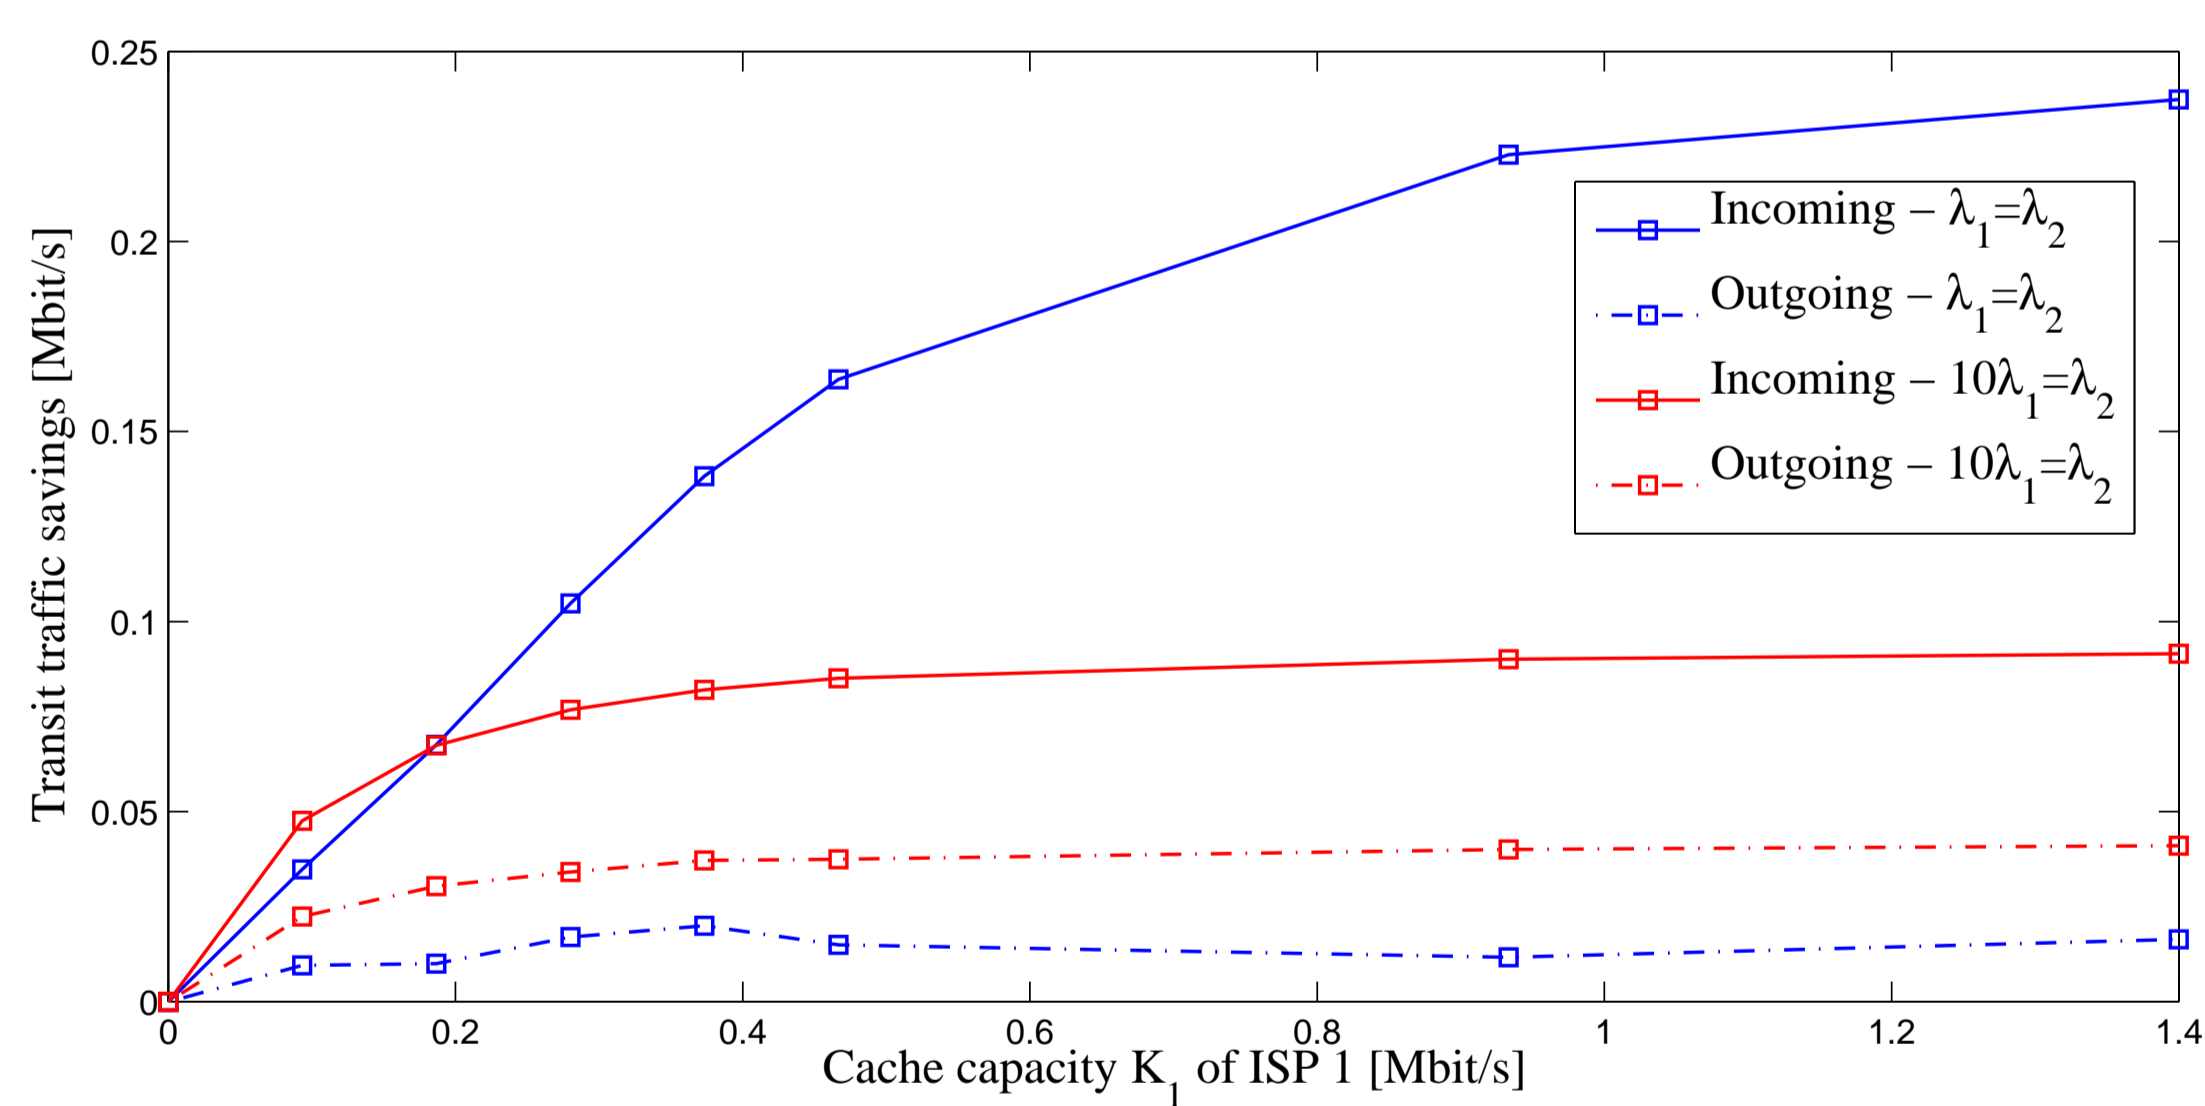
- Limiting factors of P2P caches
  1. Storage capacity  $\Rightarrow$  cache eviction
  2. Bandwidth
- Significant impact of P2P caching on both instantaneous transit traffic and system dynamic [2]
- Key question:
  - Should bandwidth be actively managed such as to minimize the inter-ISP traffic?

## 2 The Impact of P2P Caching

- $\mathcal{I} = \{1, \dots, I\}$  set of ISPs,  $\mathcal{S} = \{1, \dots, S\}$  set of swarms
- $\lambda_{s,i}$  arrival rate of leecher of ISP  $i \in \mathcal{I}$  to swarm  $s \in \mathcal{S}$
- $\theta$  leechers' impatience,  $\gamma$  departing rate of seeders,  $\mu$  upload capacity of peers
- $Z_{i,s}(t) = (x_{i,s}(t), y_{i,s}(t))$  state of swarm  $s$  in ISP  $i$  as the number of leechers and seeders in ISP  $i$  in swarm  $s$  at time  $t$
- P2P cache:
  - $K_i < \infty$  upload capacity of cache in ISP  $i$ , it serves only peers in ISP  $i$



- $I_{i,s}(Z_s(t), k_{i,s}(t))$  incoming transit traffic rate in ISP  $i$
- The impact of caching on the transit traffic depends on the distribution of the peers among ISPs



## 3 Cache Capacity Allocation Policies

A cache capacity allocation *policy* of ISP  $i$  specifies  $k_i(t)$  from the set of feasible cache capacity allocations of ISP  $i$ :  $\mathcal{K}_i = \{k_i | \sum_{s \in \mathcal{S}} k_{i,s} \leq K_i\}$

- **No Policy:**
  - ISP  $i$  does not actively allocate its  $K_i$
- **Uniform Capacity Reservation:**
  - The same amount of cache capacity is reserved to every swarms:  $k_{i,s} = \frac{K_i}{|\mathcal{S}|}$
- **Greedy Traffic Minimization:**
  - “short term” approximation, does not consider the impact of  $k_i(t)$  on the evolution of  $Z(t)$

$$k_i(t) = \arg \min_{k_i \in \mathcal{K}_i} \sum_{s \in \mathcal{S}} I_{i,s}(Z_s(t_n), k_{i,s}) \text{ for } t_n < t < t_{n+1}$$

- every swarm with  $k_{i,s} \neq 0$  should provide equal marginal traffic saving at optimality
- **Priority Based Allocation Policies:**
  - Idea: priority to swarms based on the ratio  $r = \lambda_2/\lambda_1$
  - Approximated using the instantaneous ratio:  $\hat{r}_{i,s} = \frac{x_{i,s}(t)}{\sum_{j \neq i} z_{j,s}(t)}$
  - **Ratio Priority:** swarms with highest ratio have highest priority
  - **Inverse Ratio Priority:** swarms with lowest ratio have highest priority

## 4 Performance Evaluation - Simulations

- Incoming transit traffic estimations based on two assumptions [2]:
  1. Leechers with free download capacity compete with each other for the available upload rate
  2. The distribution of the sources of content downloaded in ISP  $i$  is proportional to the amount of upload rate exposed to leechers in ISP  $i$

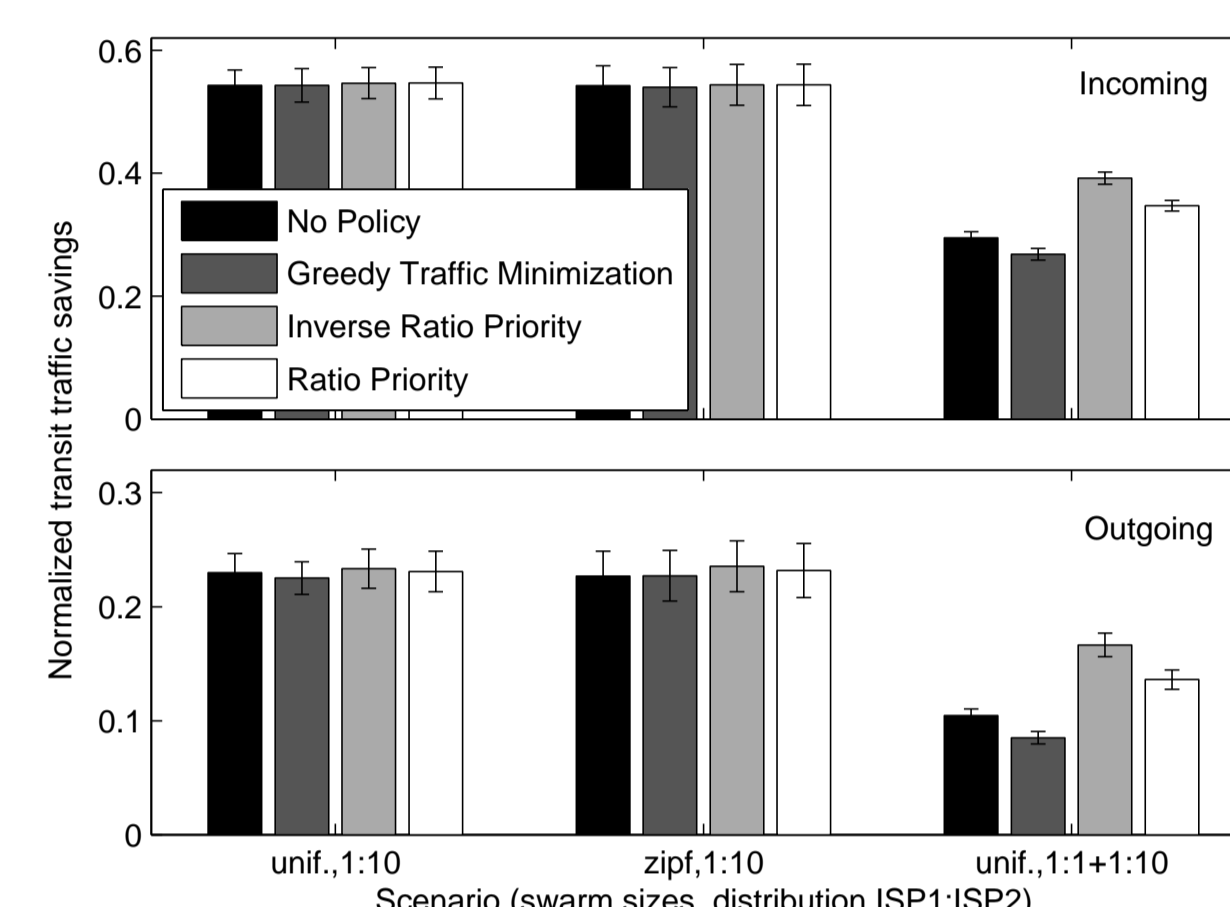
$$I_i(z_s, k_i) = D_i^r \left( \frac{\sum_{j \neq i} u_j^P}{u_i^{PL} + \sum_{j \neq i} u_j^P} \right)$$

$$D_i^r = f(k_i, u_i^{PL}, (u^P)_i) \text{ total receiving rate of leechers in ISP } i$$

- $u_i^P = \mu(\eta x_i + y_i)$  publicly available upload rate in ISP  $i$  available to leechers in every ISP
- $u_i^{PL} = \max[0, \mu(\eta(x_i - 1) + y_i)]$  publicly available upload rate in ISP  $i$  available to a local leecher

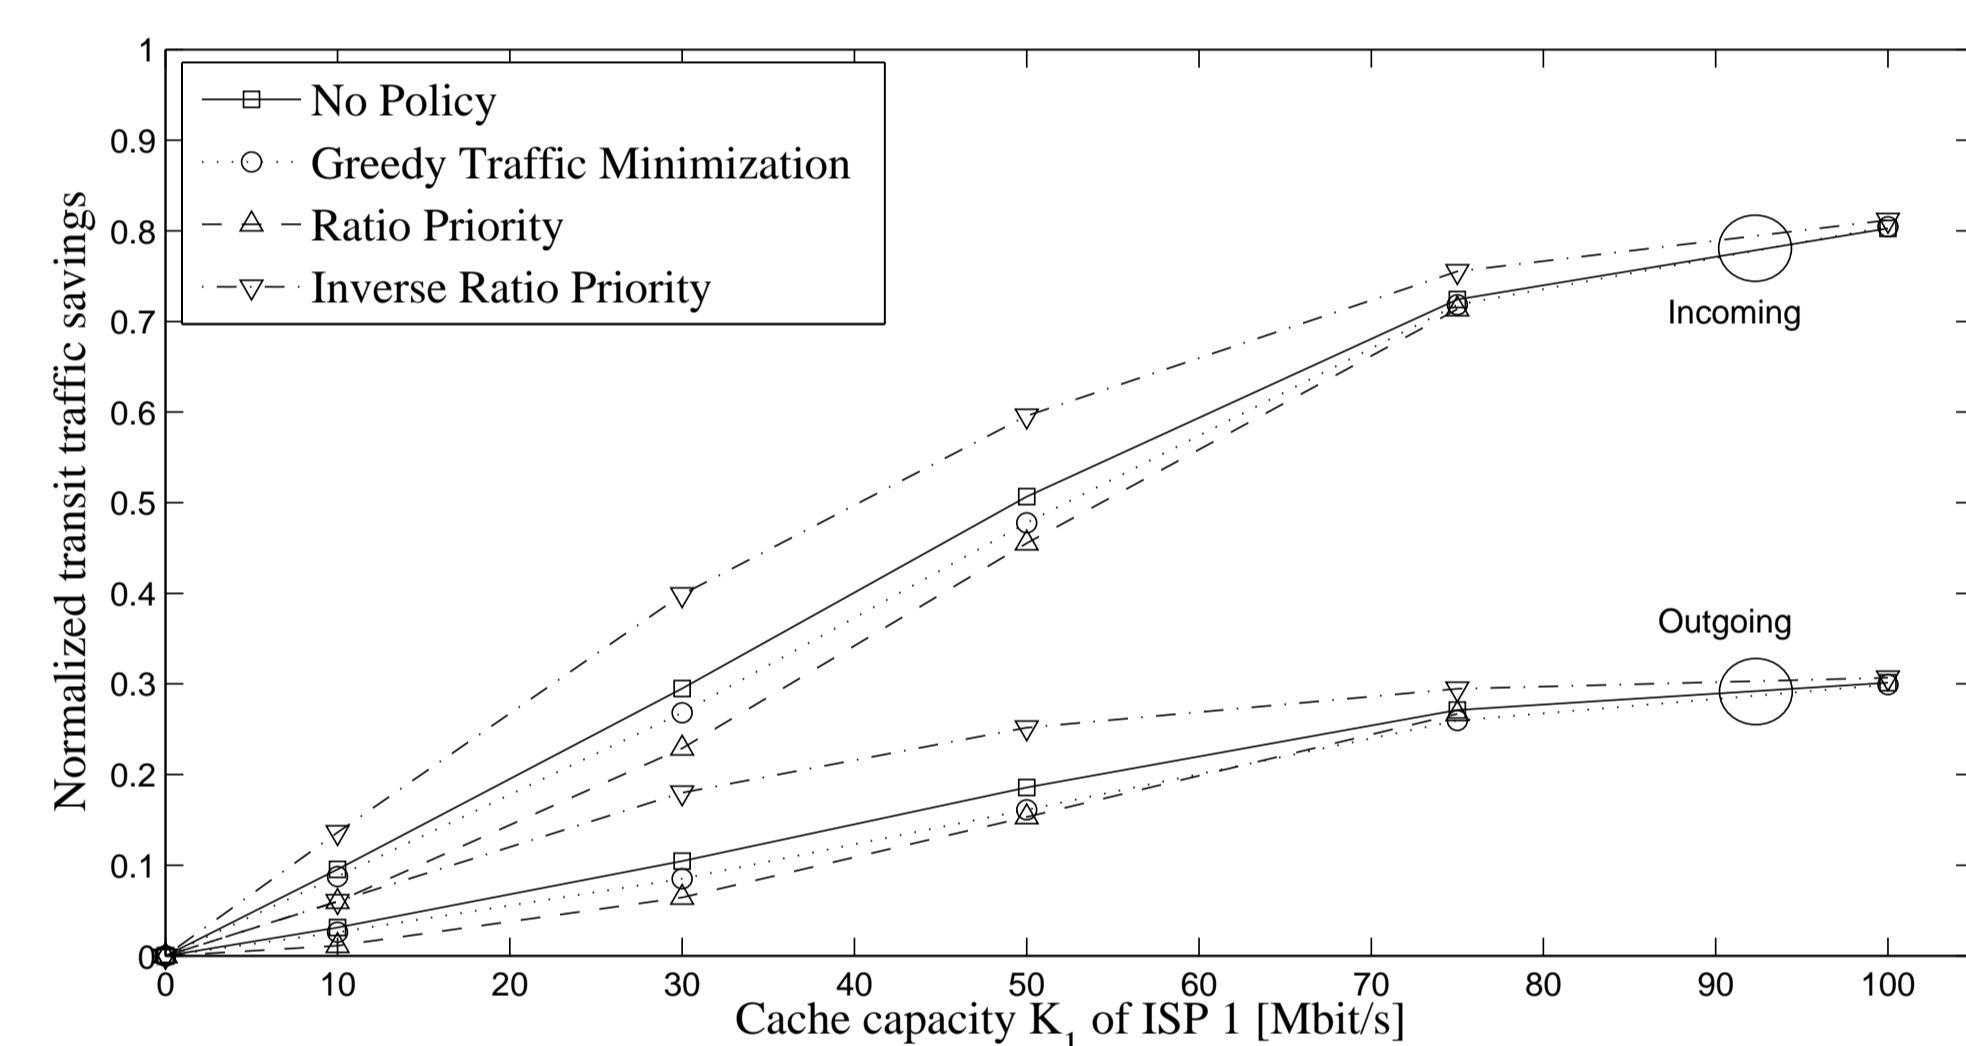
- Simulations using ProtoPeer [3]

Incoming and outgoing transit traffic savings for the scenarios described in the table:



Scenario	Identical swarms (s)	$\frac{\lambda_2}{\lambda_1}$	$\frac{\lambda_{2,s}}{\lambda_{1,s}}$
unif., 1:10	1,...,12	1/12	10
zipf, 1:10		$\propto \frac{1}{s}$	10
unif., 1:1+1:10	1,...,10	1/12	10
	11,12	1/12	1

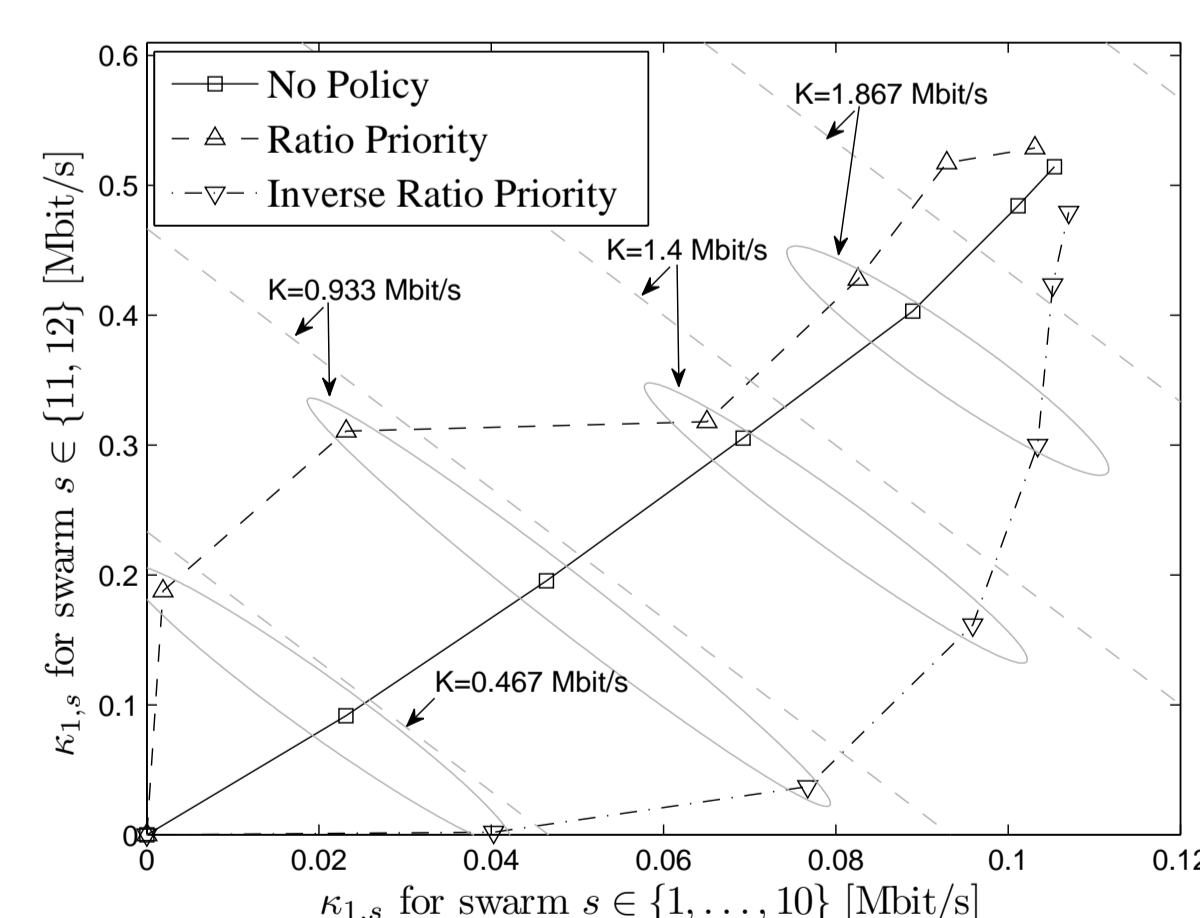
Incoming and outgoing transit traffic saving for the *unif., 1:1+1:10* scenario:



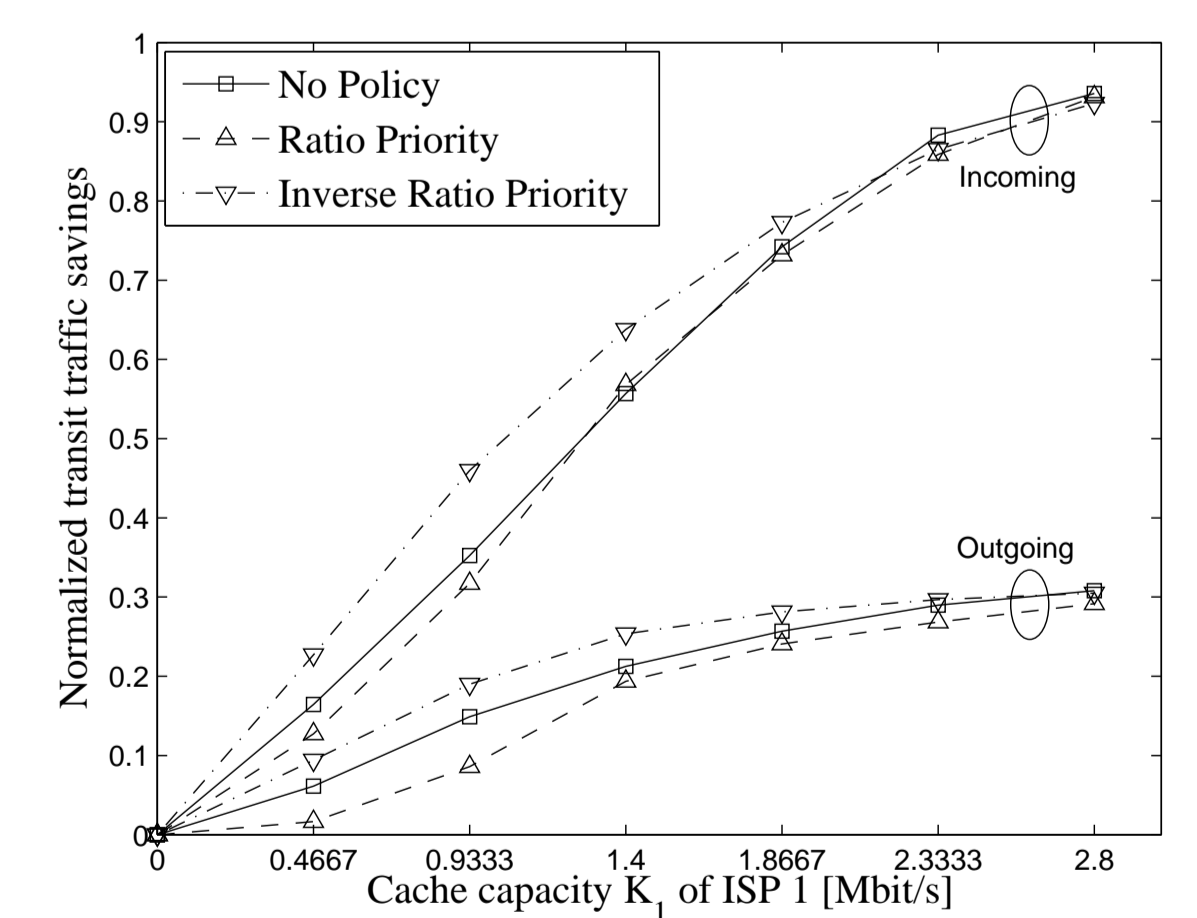
## 5 Experimental Validation

- Experiments on ca. 500 Planet-lab nodes using Bit-Torrent:
  - File-size 3.5 MB
  - Peers' upload capacity 23 kbit/s, download capacity 373 kbit/s
  - *unif., 1:1+1:10* scenario
  - Cache: 12 peers running on the same Linux machine
  - Cache capacity allocation policies implemented using Linux Traffic Control (tc)

Experimental results for the *unif., 1:1+1:10* scenario:



Cache capacity allocations of ISP 1



Normalized transit traffic saving of ISP  $i$

## References

- [1] G. Dán, T. Hossfeld, S. Oeschner, P. Cholda, R. Stankiewicz, I. Papafili, G. D. Stamoulis, “Interaction Patterns between P2P Content Distribution Systems and ISPs”, in *IEEE Communications Magazine*, vol. 49, num. 5, May 2011, pp. 222-230.
- [2] F. Lehrieder, G. Dán, T. Hossfeld, S. Oeschner, V. Singeorzan, “The Impact of Caching on BitTorrent-like Peer-to-peer Systems”, in *10th IEEE International Conference on Peer-to-Peer Computing 2010 - IEEE P2P 2010*, Delft, Netherlands, 2010, pp. 6978.
- [3] W. Galuba, K. Aberer, Z. Despotovic, W. Kellerer “ProtoPeer: A P2P Toolkit Bridging the Gap between Simulation and Live Deployment”, in *Proceedings of International Conference on Simulation Tools and Techniques*, Mar. 2009.