ICC: An Incentive-Compatible Inter-Cloud Communication Traffic Management Mechanism

Manos Dramitinos, George D. Stamoulis

Network Economic and Services Laboratory

Athens University of Economics and Business

http://nes.aueb.gr

CNSM – November 09, 2015

ATHENS UNIVERSITY
OF ECONOMICS
AND BUSINESS

Outline

- Motivation and problem definition
- Our approach: ICC
  - The main ideas, avoiding the technicalities as much as possible
  - Some interesting properties and tradeoffs
- Simulation model and results with real traces
- Material not covered in the paper
  - Implementation, results from network experiments, on-going work
- Conclusions

Internet Transit

- ISP sells access to *all destinations* in its routing table
  - Tier-1 providers: do not buy transit, maintain settlement-free peering **(P)** amongst them and sell global connectivity **(T)** to lower-tier ISPs

Source: Dr. Peering

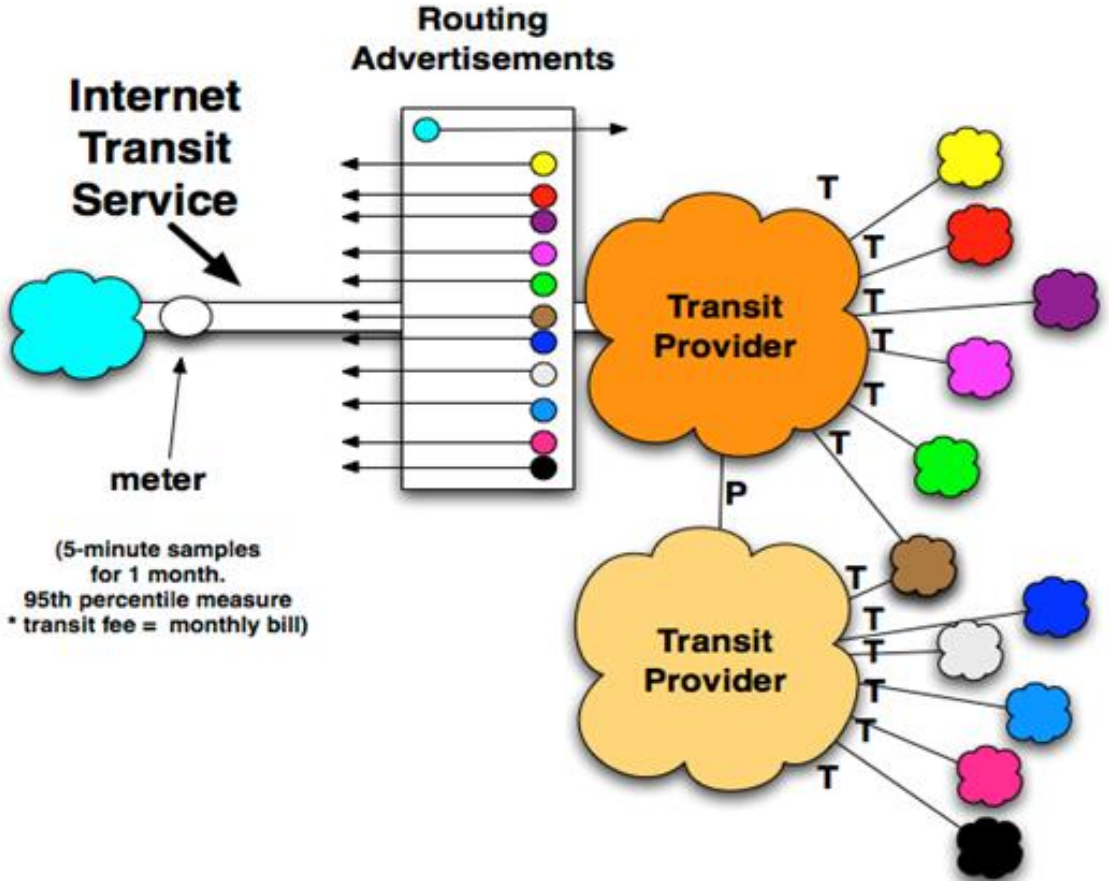

95th Percentile Explained

- Transit: metered service: 5-min samples, 95th percentile
- 95th percentile is typically significantly lower than average rate

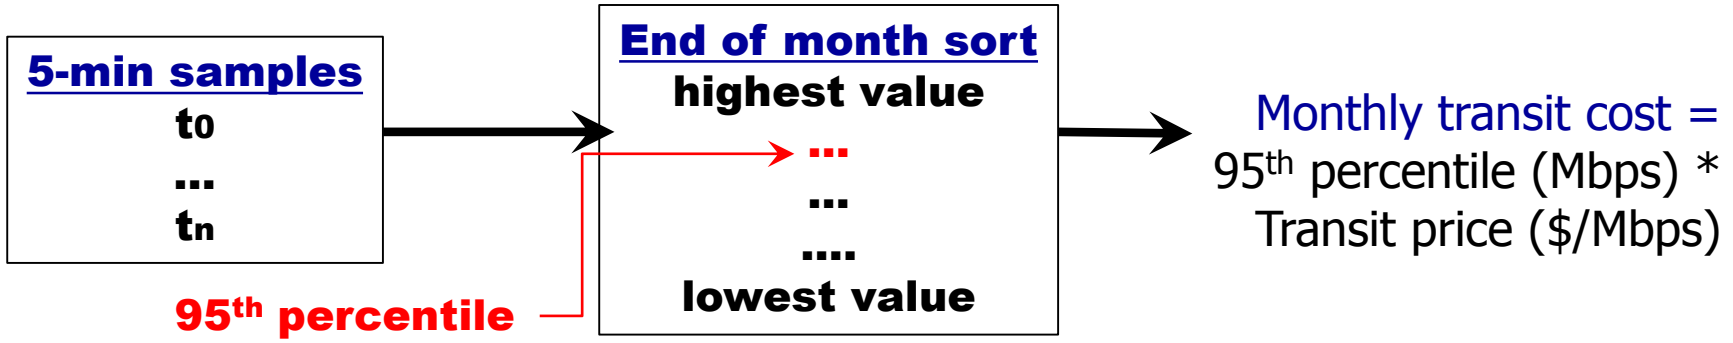

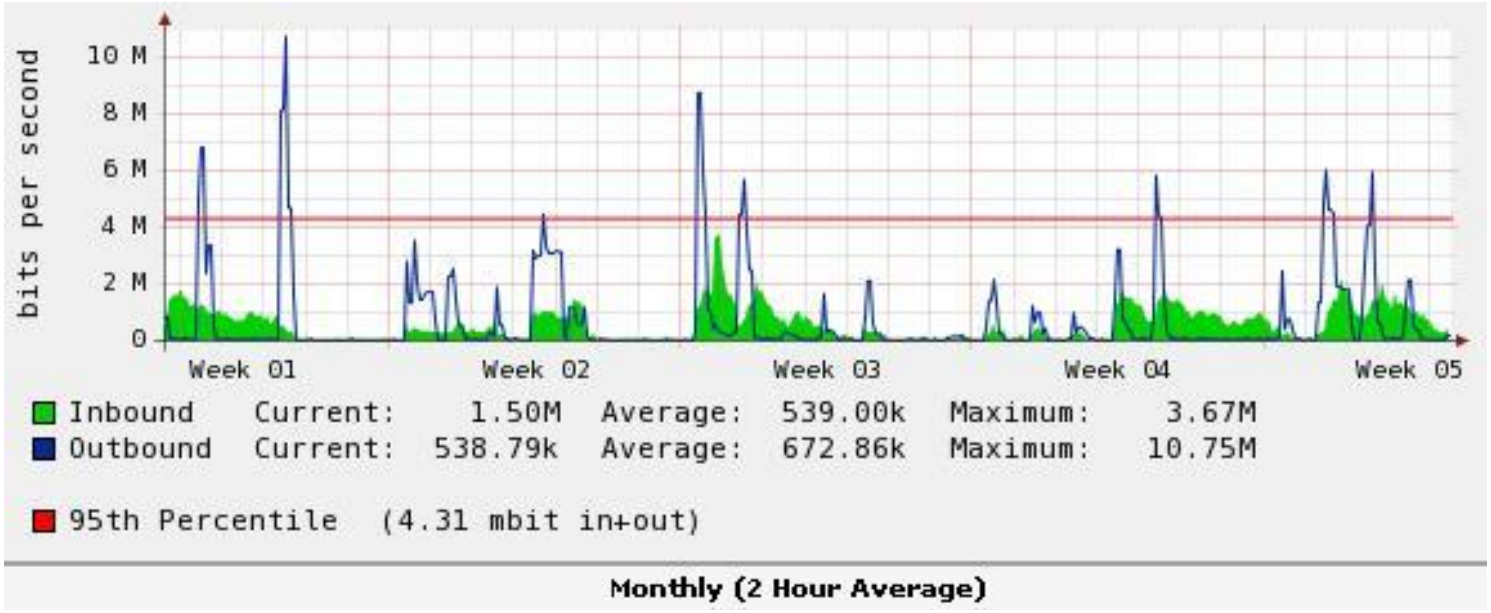

ICC Motivation

- Some traffic can be sent “later”/rate controlled without impact on end users (“Web back-office”)
  - Inter-cloud bulk data transfers (e.g. back-ups), CDN cache sync
  - Monetary compensation as incentive to business customer (Cloud, DC) to mark a portion of the traffic as “time-shiftable”
- ISP performs rate control over marked traffic only
  - … reducing traffic peaks and 95th percentile, i.e. transit cost
  - Win-win: Cost savings shared among ISP and his customers
  - Reasonable traffic management:
    - net-neutral, no throttling

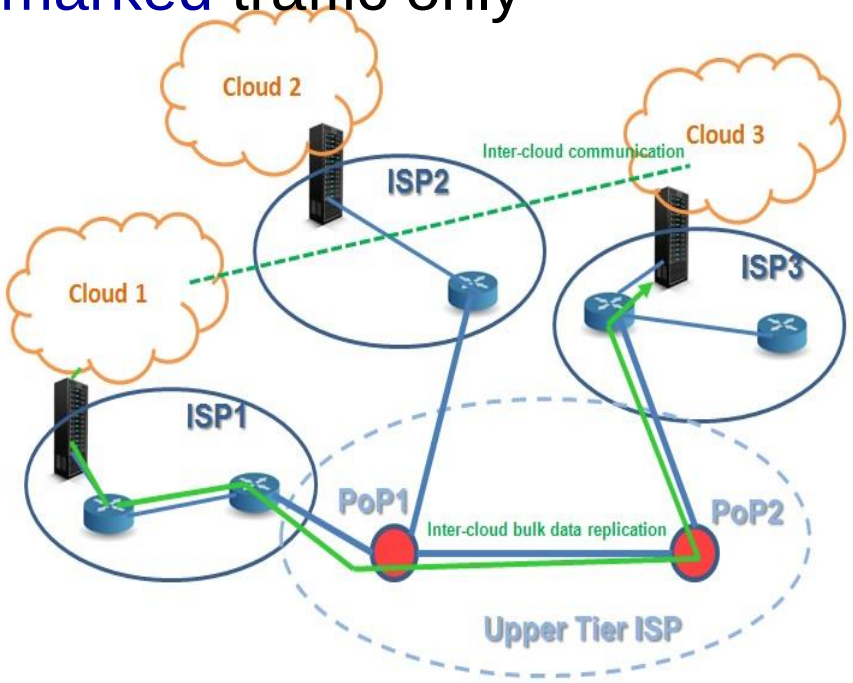

ICC: The Big Picture

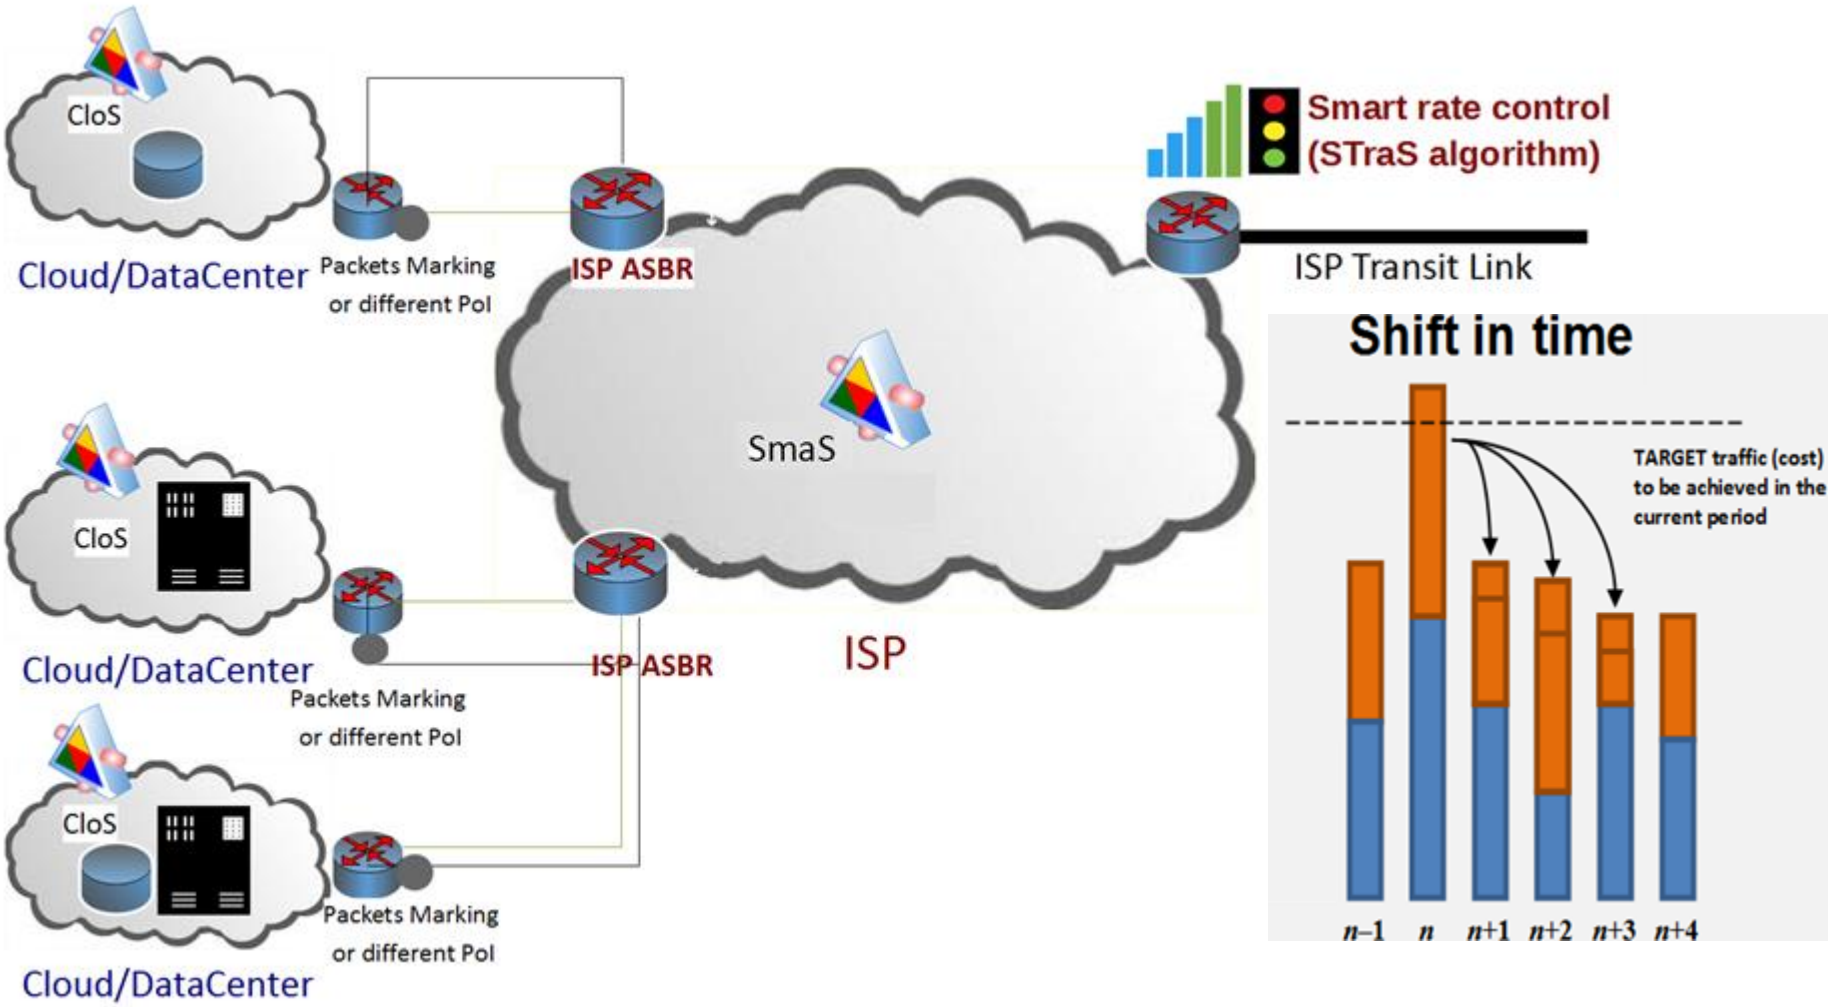

ICC: STraS Traffic Management

- ISP sets C_{target} for transit charge (target 95th percentile)
- Rationale: Water-filling so that C_{target} is not violated
  - Fine rate control of time-shiftable traffic over multiple epochs (*y=10*) within the 5-min interval
    - use *y* thresholds, one per epoch for the target rate
    - First *(y-1)* epochs: *E[rt]+time-shiftable <0.9*C_{target}*,
    - Last epoch: Look back, adjust rate so that C_{target} is not violated within the 5-min interval as a whole
    - Send with max rate at the last 5% 5-min intervals if possible, since this will not affect the transit cost (ICC_FA)
  - Alterative: Send with max rate at peak periods (ICC_FAP)
  - Evaluated using ISP traces under full information and using statistics-based traffic expectation (ICC_STATS)
  - Partly implemented over Juniper MX240 routers

ICC: Main Properties

- Granularity of control
- Simplicity, scalability
  - Traffic aggregates, no per-flow guarantees
- Implementability
- Optimization potential
  - Depends on the traffic mix
- Net neutrality
  - No uneven power distribution, no throttling
- Incentives
  - Pricing
- Built-in support for cloud layer optimizations
  - Optimal destination selection, not covered in this presentation

ICC Pricing and Incentives

- Main issues: Incentives, sustainability, simplicity
  - No unrealistic traffic monitoring and billing/accounting overheads
- We specify ISP to return a cut *p* of his savings:

$$discount_i = p \cdot ISPsavings \cdot \frac{vol_i}{\sum_{j \in I} vol_j}$$

  - Increasing in the volume of traffic to be managed by ICC
  - Simple, information required is readily available for the ISP
- Accurate computation of savings? Approximation with *cagr*
- What could *vol*$_i$ be?
  - Total? Volume at peak epochs? Volume actually shifted?
  - Last two definitions have adverse impact on incentives!
- A priori or ex post announcement? Use *cagr*

ICC Evaluation: Full Information (1/2)

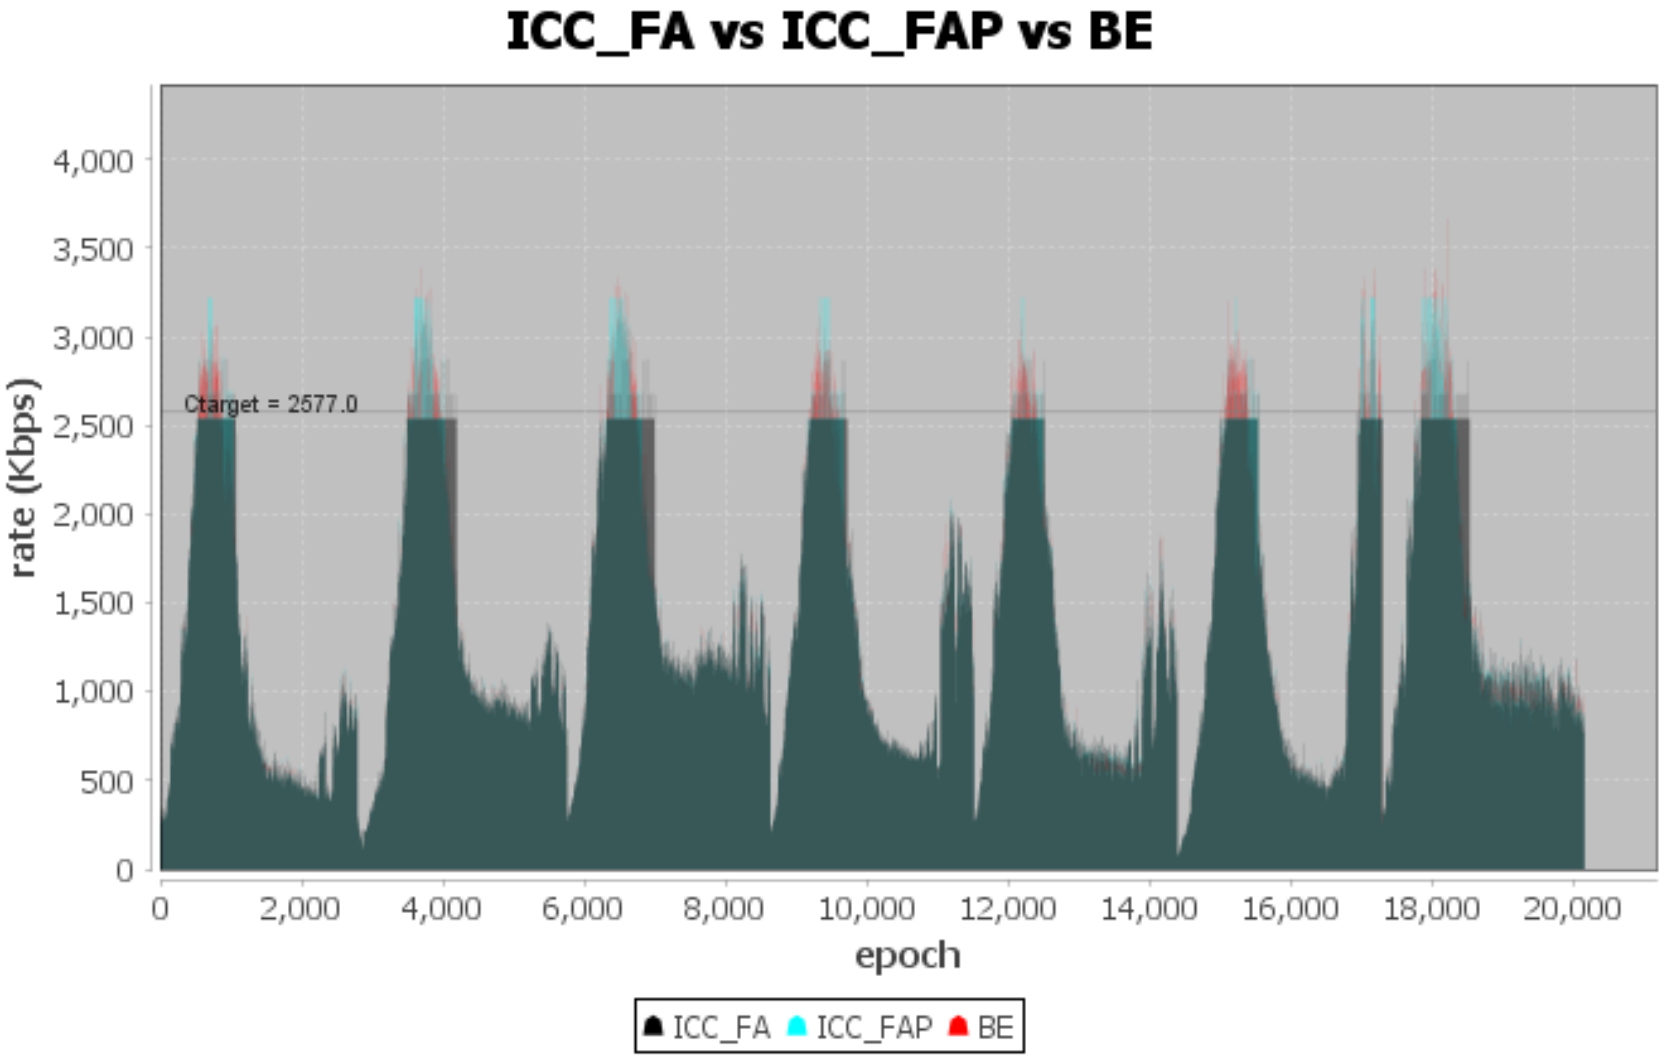

ICC Evaluation: Full Information (2/2)

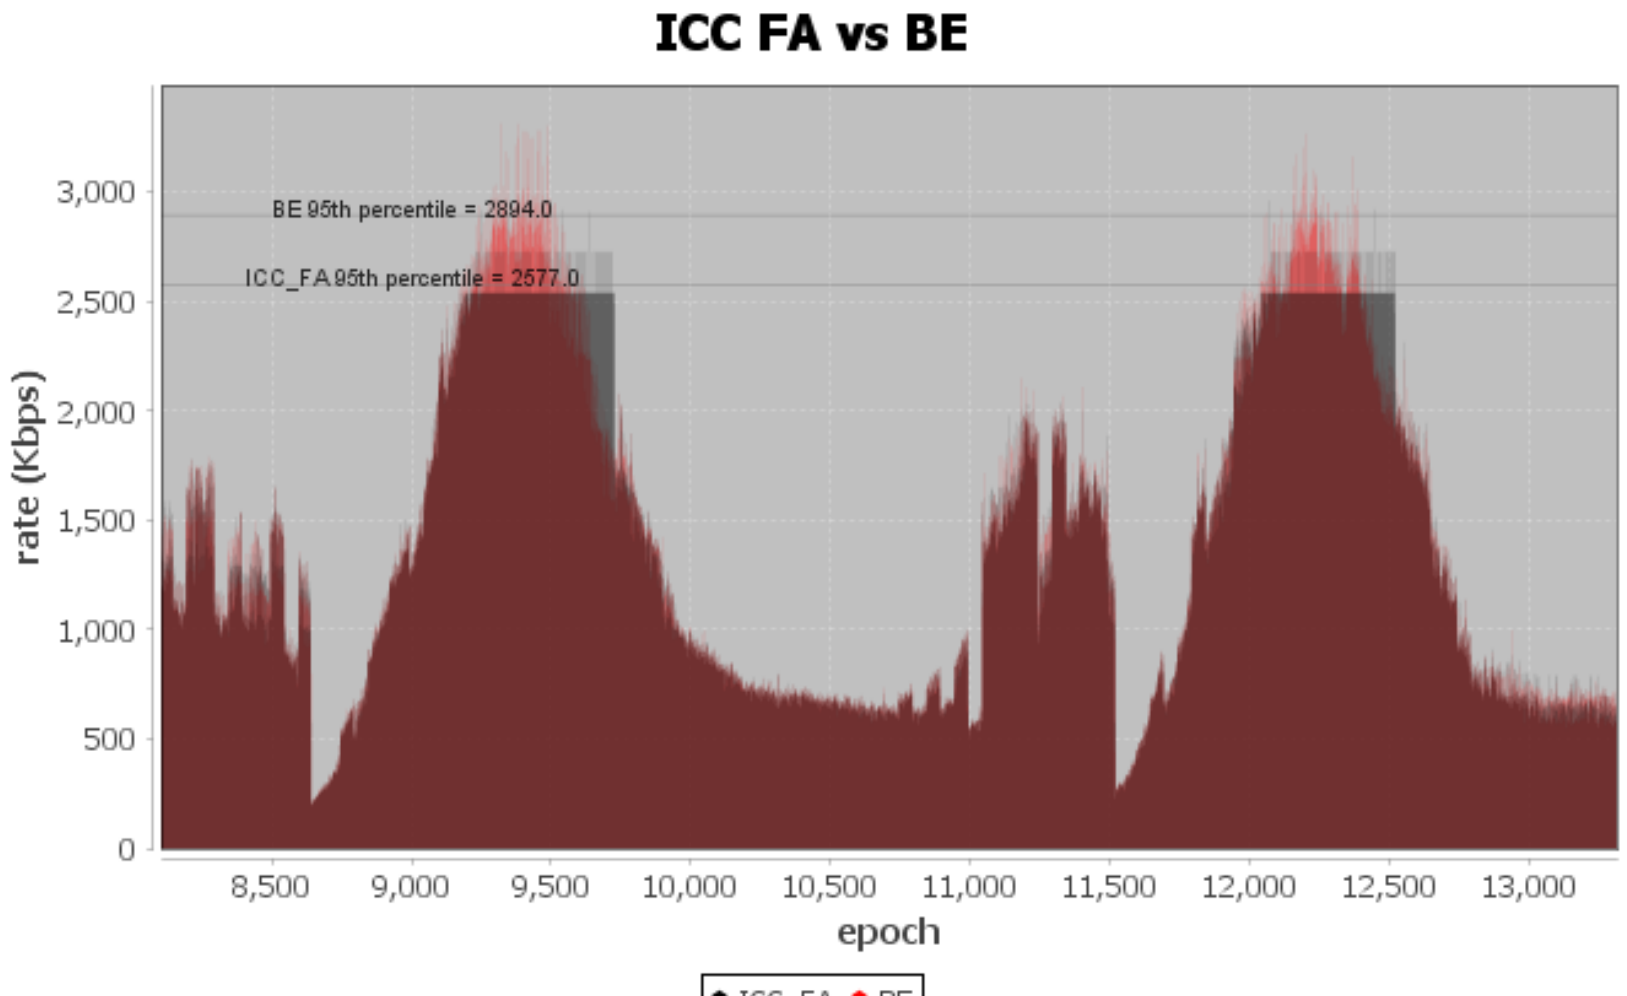

ICC Evaluation: ICC_STATS (1/3)

- First day of week used as training sample
- All rate control thresholds set to 1 (error-prone)
  - ICC always improves 95th percentile compared to Best Effort – no ICC

Attained 95^{th} Percentile ($epoch\_tholds[1, \ldots, y] = 1$)					
BE - no ICC	ICC_FA	ICC_FAP	ICC_STATS	C_{target}	ICC_STATS deviation(%)
4674870	3285000	3285000	3607089	3285000	9.80
4674870	3559000	3559000	3847349	3559000	8.10
4674870	3833000	3833000	4077187	3833000	6.37
4674870	4107000	4107000	4260687	4107000	3.74
4674870	4381000	4381000	4417069	4381000	0.82
4674870	4654000	4654000	4653527	4654000	-0.01
					Average: 4.80%

ICC Evaluation: ICC_STATS (2/3)

- All thresholds set to same value

TABLE II. ICC_STATS SENSITIVITY ANALYSIS

Attained 95^{th} Percentile			
$epoch\_tholds[] = 0.85$			
BE - no ICC	C_{target}	ICC_STATS	ICC_STATS deviation(%)
4674870	3285000	3268679	-0.50
4674870	3559000	3437469	-3.41
4674870	3833000	3627181	-5.36
4674870	4107000	3834504	-6.63
4674870	4381000	4036428	-7.86
4674870	4654000	4212442	-9.48
			Average: -5.54%
$epoch\_tholds[] = 0.90$			
BE - no ICC	C_{target}	ICC_STATS	ICC_STATS deviation(%)
4674870	3285000	3381226	2.92
4674870	3559000	3582225	0.65
4674870	3833000	3797732	-0.92
4674870	4107000	4011633	-2.32
4674870	4381000	4201133	-4.11
4674870	4654000	4356671	-6.39
			Average: -1.69%

ICC Evaluation: ICC_STATS (3/3)

$epoch\_tholds[] = 0.91$

BE - no ICC	C_{target}	ICC_STATS	ICC_STATS deviation(%)
4674870	3285000	3410791	3.82
4674870	3559000	3609757	1.42
4674870	3833000	3831876	-0.03
4674870	4107000	4047543	-1.45
4674870	4381000	4221327	-3.64
4674870	4654000	4384063	-5.80
			Average: -0.94%

$epoch\_tholds[] = 0.92$

BE - no ICC	C_{target}	ICC_STATS	ICC_STATS deviation(%)
4674870	3285000	3434814	4.56
4674870	3559000	3639854	2.27
4674870	3833000	3866020	0.86
4674870	4107000	4080928	-0.63
4674870	4381000	4253913	-2.90
4674870	4654000	4404175	-5.37
			Average: -0.20%

$epoch\_tholds[] = 0.93$

BE - no ICC	C_{target}	ICC_STATS	ICC_STATS deviation(%)
4674870	3285000	3454156	5.15
4674870	3559000	3671885	3.176
4674870	3833000	3900517	1.76
4674870	4107000	4107957	0.02
4674870	4381000	4282814	-2.24
4674870	4654000	4421640	-4.99

Implementation

- Partly over Juniper MX240 routers
- Implementation by Rafal Stankiewicz and Zbigniew Dulinski (AGH)
- Using router hierarchical policer;
- Traffic generated by Spirent
- SmartenIT Deliverable 2.5

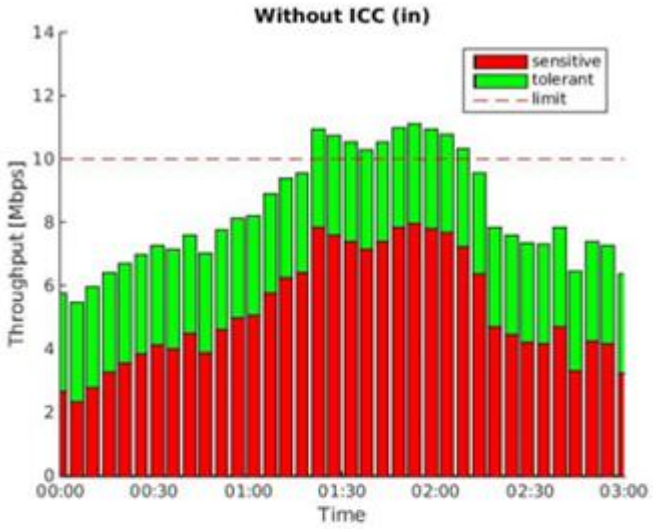

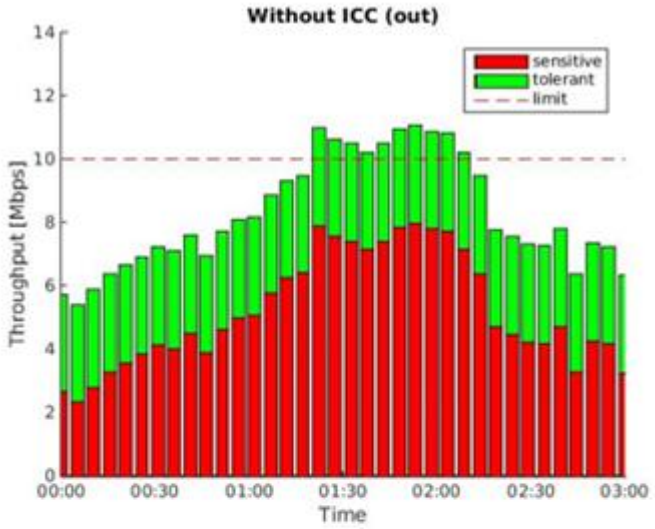

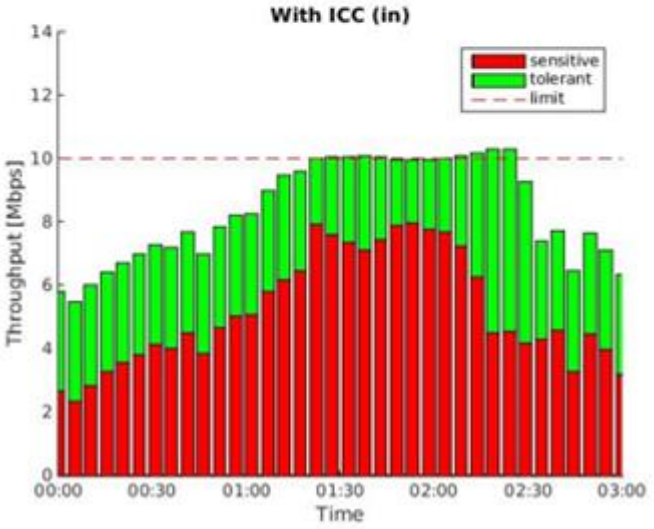

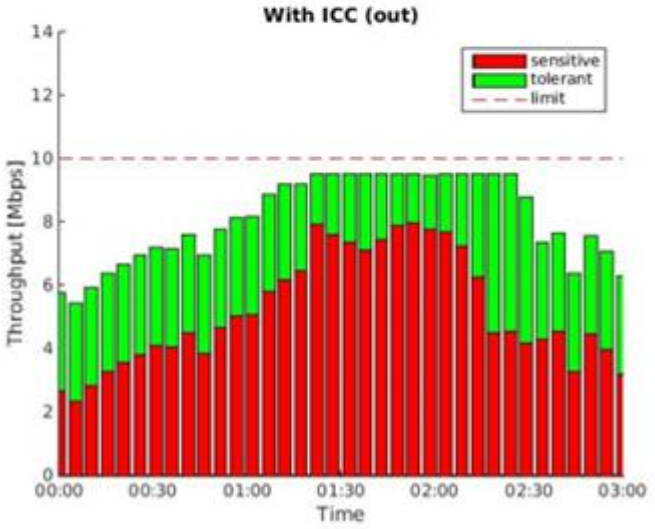

On-going work

- Combine ICC with DTM, AUEB-AGH collaboration

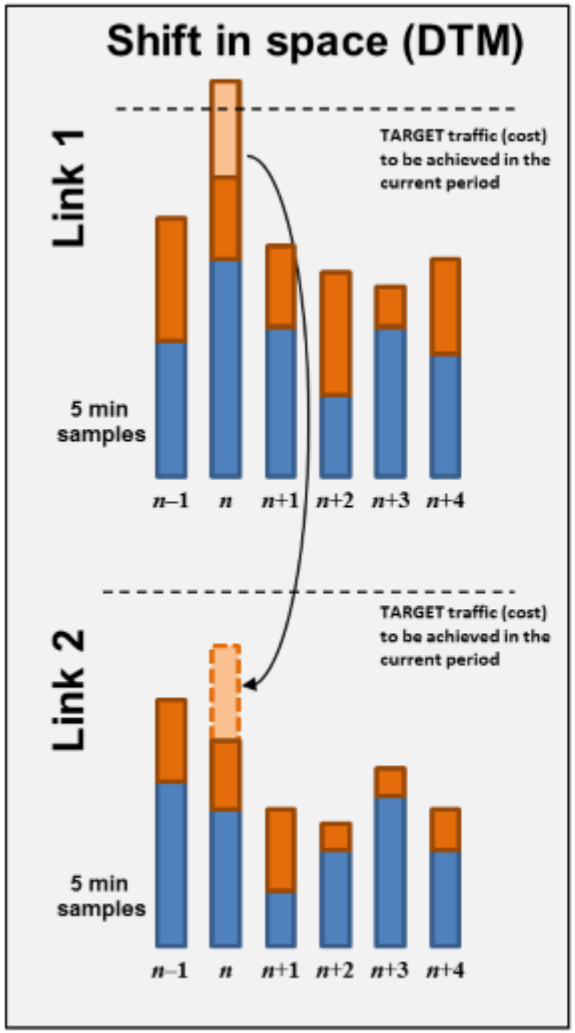

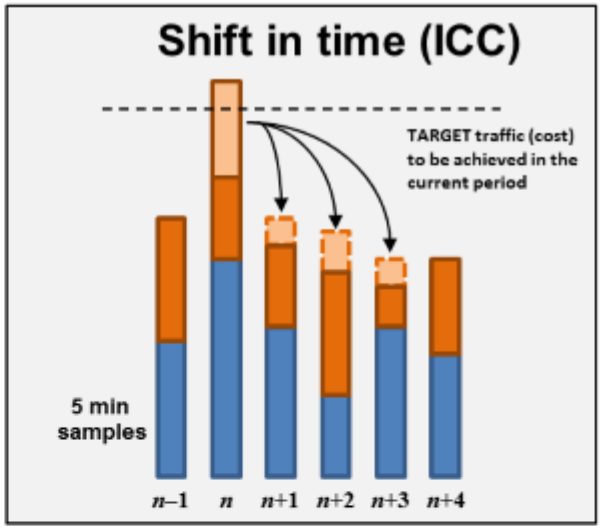

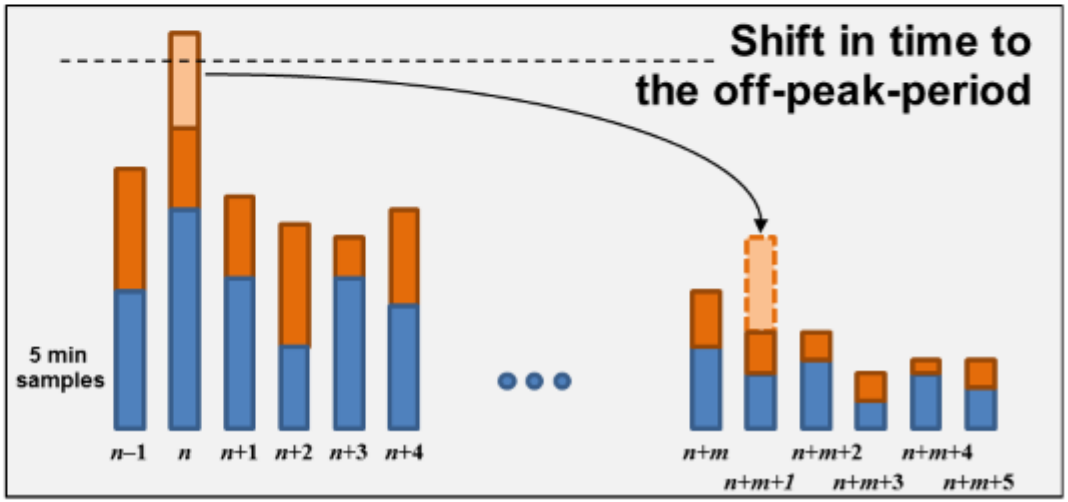

Conclusions

- STraS mechanism operates on top of Best Eff. Internet
  - Requires 1) marking/different PoI and 2) buffering or rate control
  - Operates in smaller time scales than most other mechanisms
  - Provides built-in support for Cloud/DC/Cloud Federations
  - Follows incentive compatibility and design-for-tussle principles
- Operates on the ISP transit link, incentive-compatible
  - Predictable low extra delay for “delay-tolerant” traffic
- Simulation model and initial results with real traces:
  - ISP can reduce transit charge and balance the network load under good knowledge of traffic and right selection of *Ctarget*
  - Business customers also benefit
  - Indirect benefit for real-time flows (not captured or measured)
- Potential to combine it with other proposed mechanisms

Thank you for your attention!

More information:
http://nes.aueb.gr