Detection of Network Flow Timestamp Reliability

Martin Žádník, Erik Šabík, Václav Bartoš

Intro

- Flow monitoring
  - Network visibility
  - Many applications

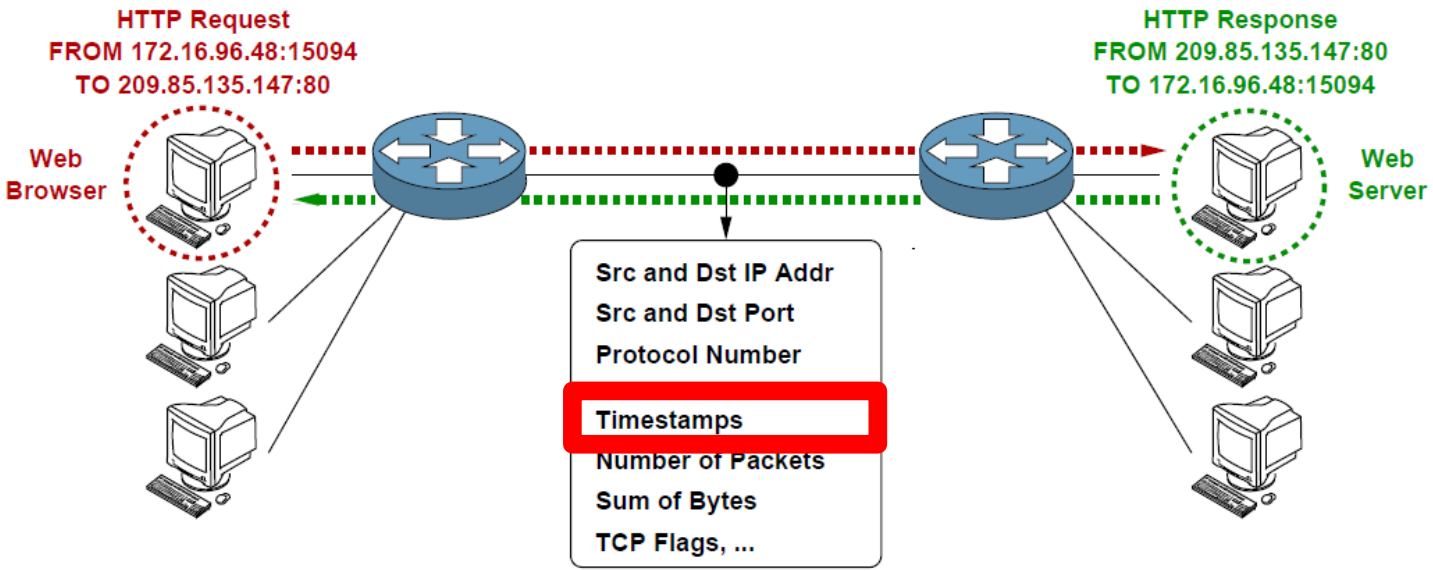

Flow start	uration	Proto	Src IP Addr:Port		Dst IP Addr:Port	Flags	Packets	Bytes
09:41:21.763	0.101	TCP	172.16.96.48:15094	->	209.85.135.147:80	.AP.SF	4	715
09:41:21.893	0.031	TCP	209.85.135.147:80	->	172.16.96.48:15094	.AP.SF	4	1594

Timestamp issues

- Flow measurement issues
  - Measurement artifacts in netflow data
  - Uncovering artifacts of flow measurement tools
  - One-way delay measurement based on flow data: Quantification and compensation of errors by exporter profiling

Timestamp issues

- Peeling away timing error in NetFlow data. Timestamp errors by design.

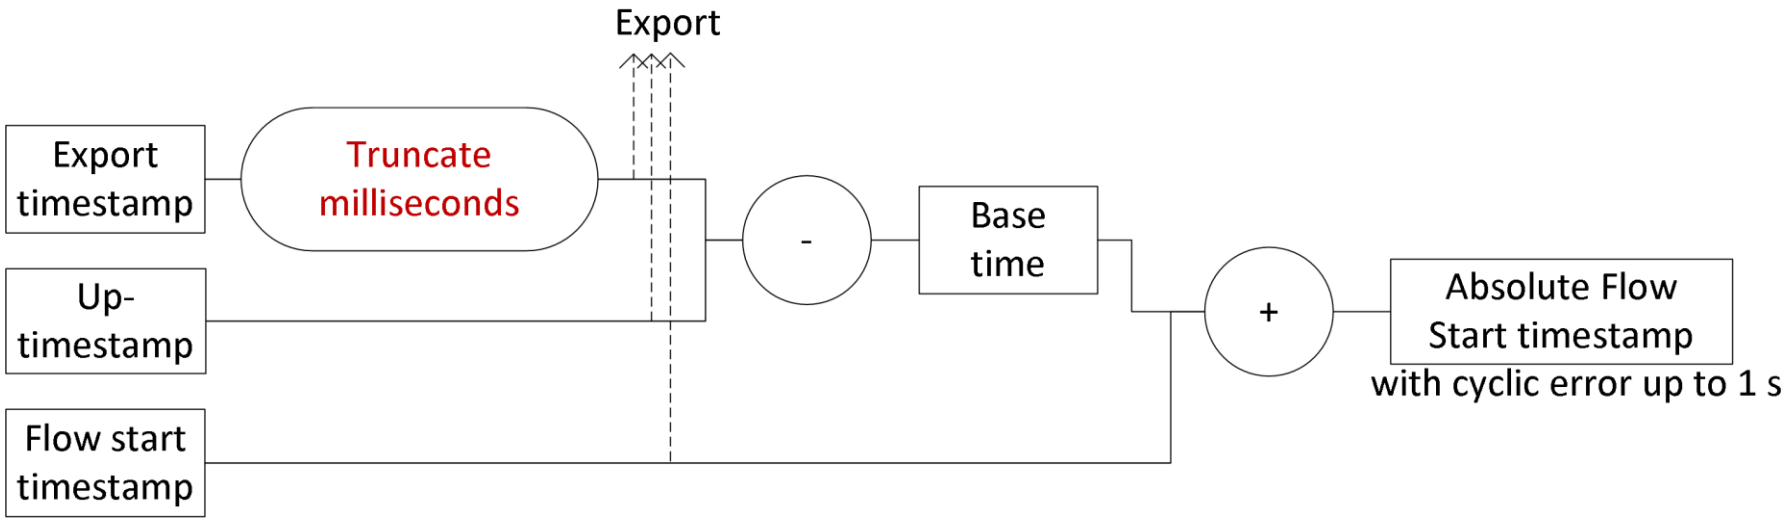

Timestamp issues

- Buffers
- Packet sampling
- Timestamp representation
- Deduplication

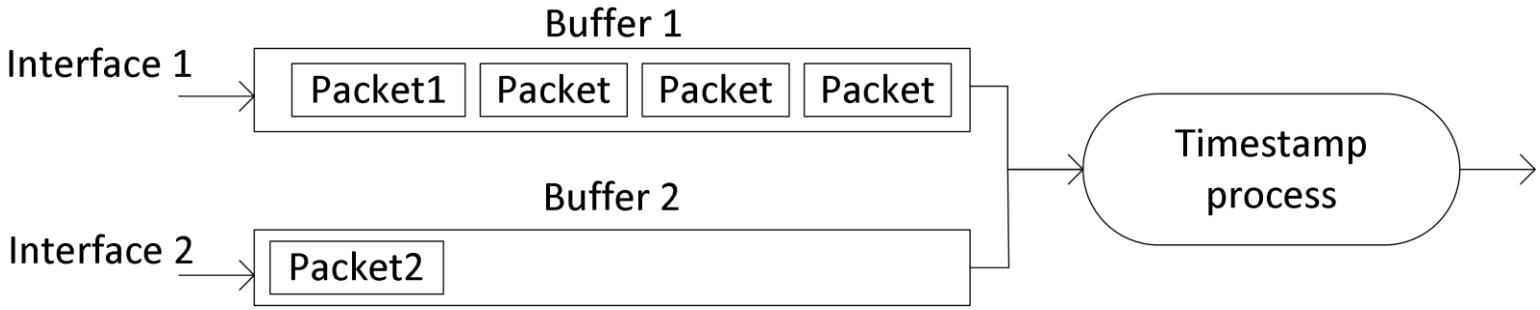

Detection

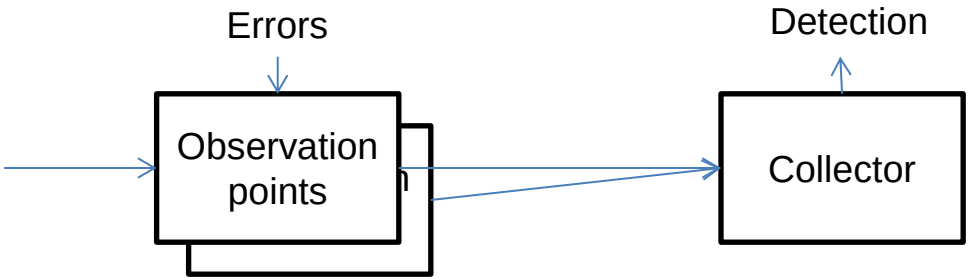

- Goal: Estimate the number of flow records with suspicious timestamps

- Input: flow records
- Output: percentage of reliable timestamps

Detection

- Mismatch between timestamps and port numbers

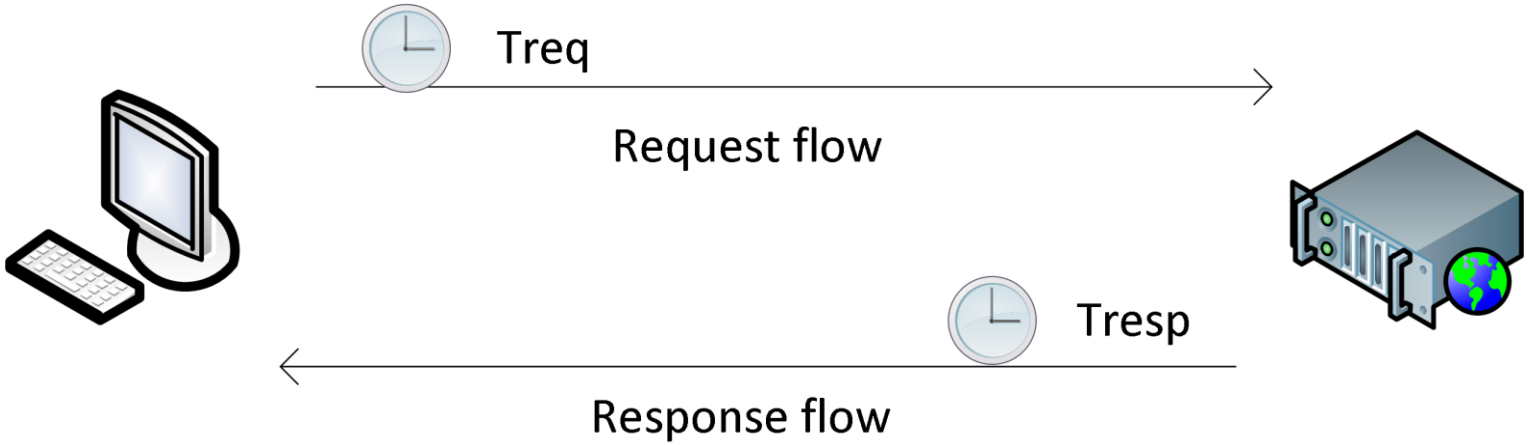

- *T*req < *T*resp when ReqDstPort is well known

Detection

- Only subset of flows may be used for timestamp evaluation – *T*
  - Request and response flows
  - TCP and SYN flags set in both directions
  - Timestamps are not equal

Detection

- Only subset of flows may be used for port heuristic – *P*
  - TCP or UDP
  - Port < 1024

Estimation

- Timestamp reliability estimate *e* is correlation of *T* and *P* in overlap

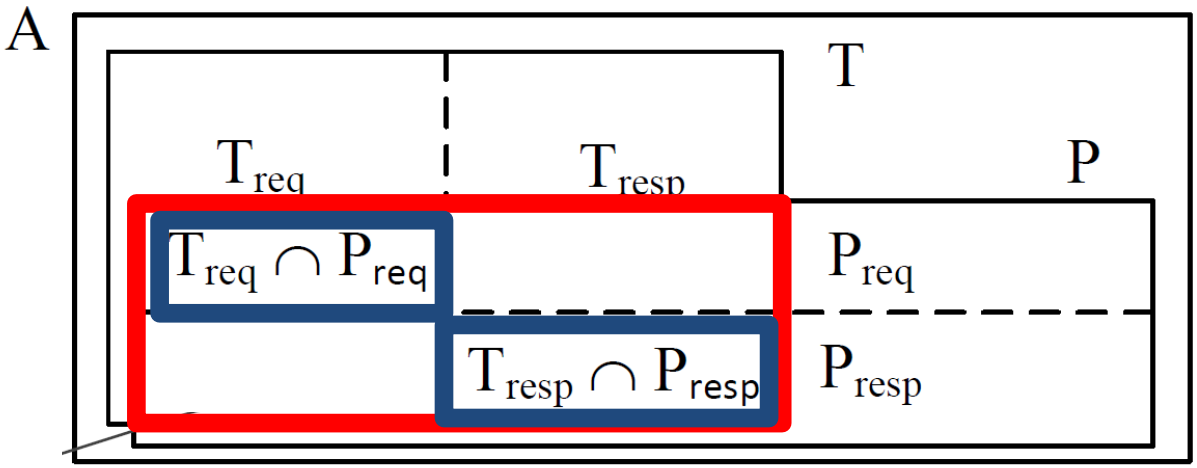

$$e = \frac{|T_{req} \cap P_{req}| + |T_{resp} \cap P_{resp}|}{|T \cap P| + |T_{equal}|}$$

Penalty

Estimation

- If *e* is 100% then both heuristics are inline and timestamps of other flows (such as UDP) are deemed correct
- If *e* is close to 50% then timestamps are deemed not reliable
- If *e* drops to 0% then negative correlation

- Corner condition the overlap should contain at least 5% of all flows

Data traces

- 3 data traces

	Flows [mil.]	Packets [mil.]	Bytes [bil.]	$\lvert T_{equal} \rvert$	e
data set A - Aconet	266	10379	8887	0.02%	100%
data set B - VUT	190	7595	6668	0.04%	54%
data set C - Mawi	8	58	27	0.7%	89%

Evaluation

- Swap timestamps of flows

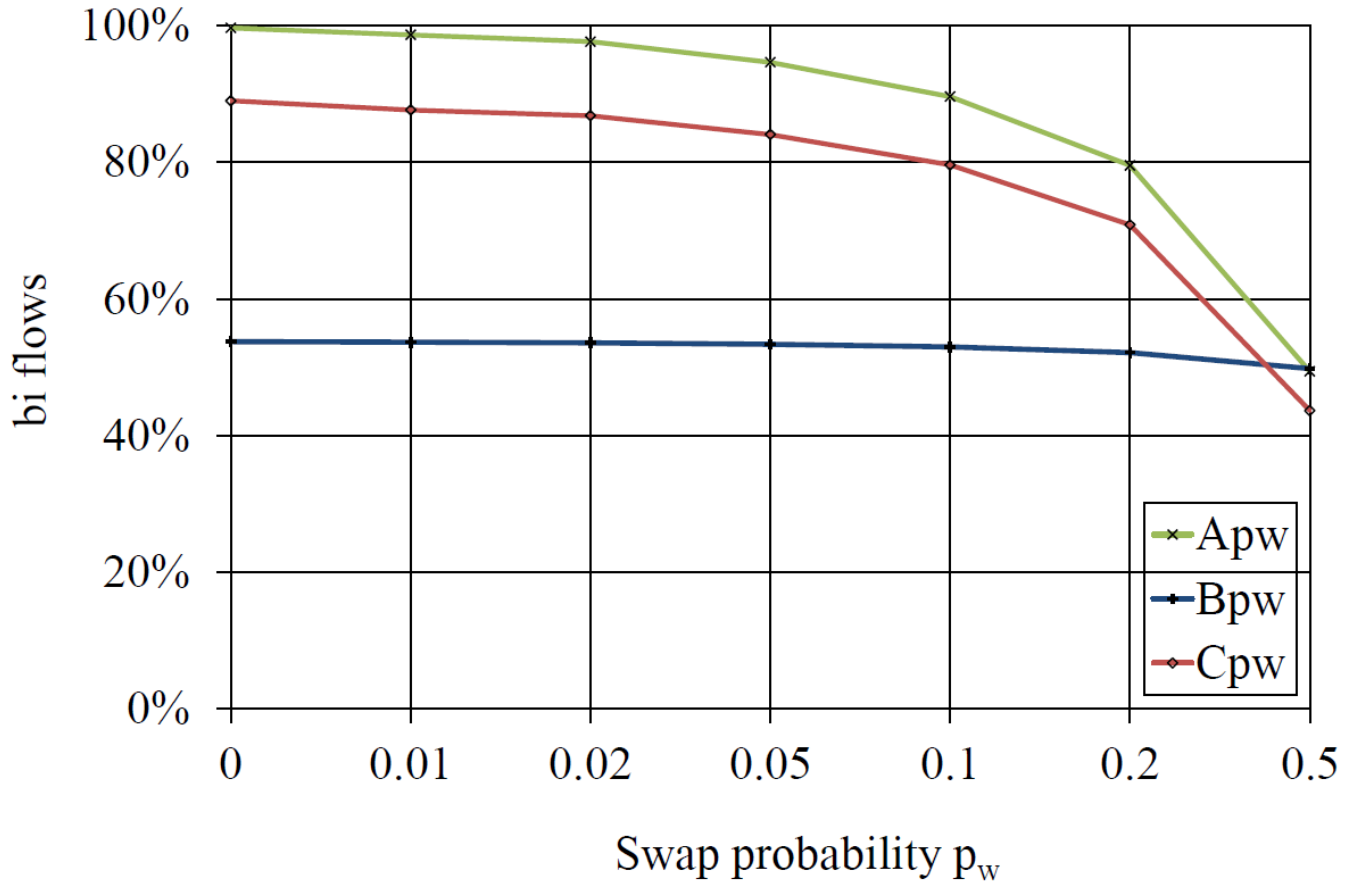

Evaluation

- Equalize timestamps that are closer than

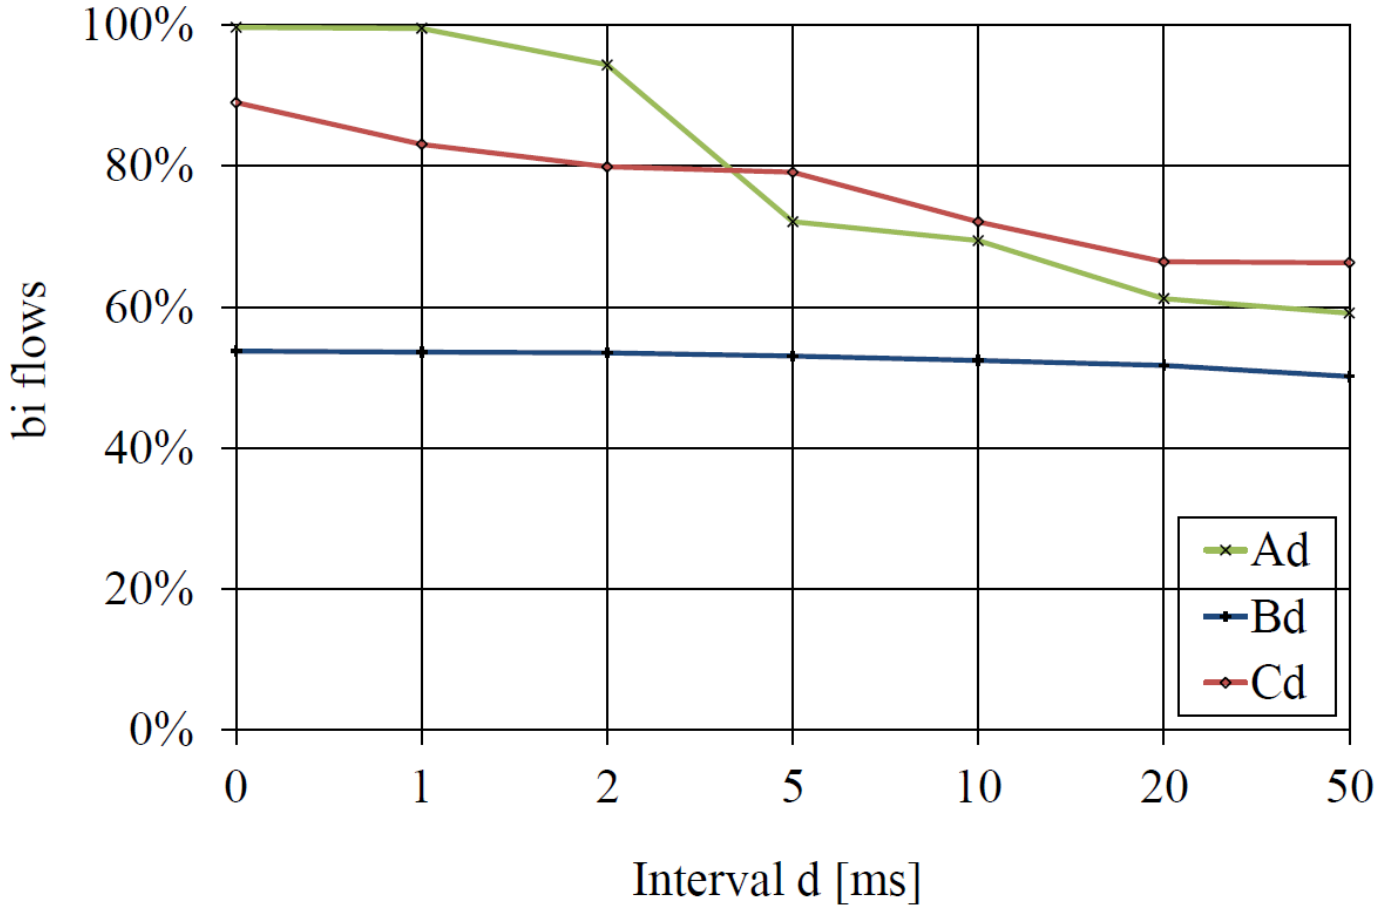

Evaluation

- Impact of sampling

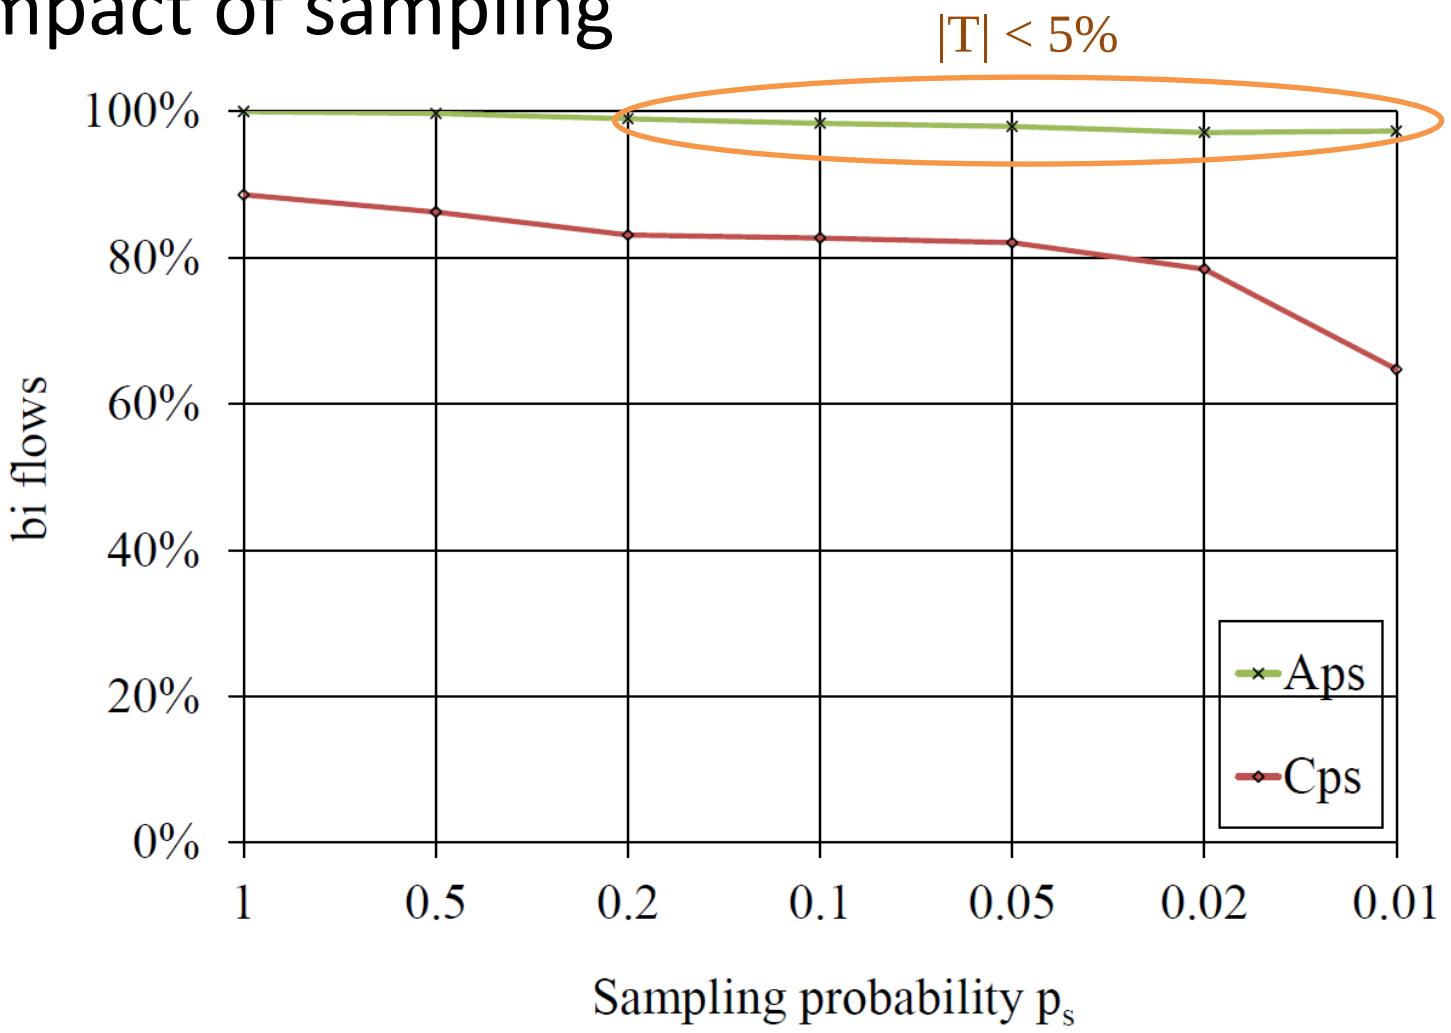

Summary

- Estimation *e* reflects reliability of timestamps well
- Let's utilize *e* for driving bi flow orientation

Biflow algorithm

- Utilize timestamps whenever
  - timestamps can be utilized
  - timestamps are reliable

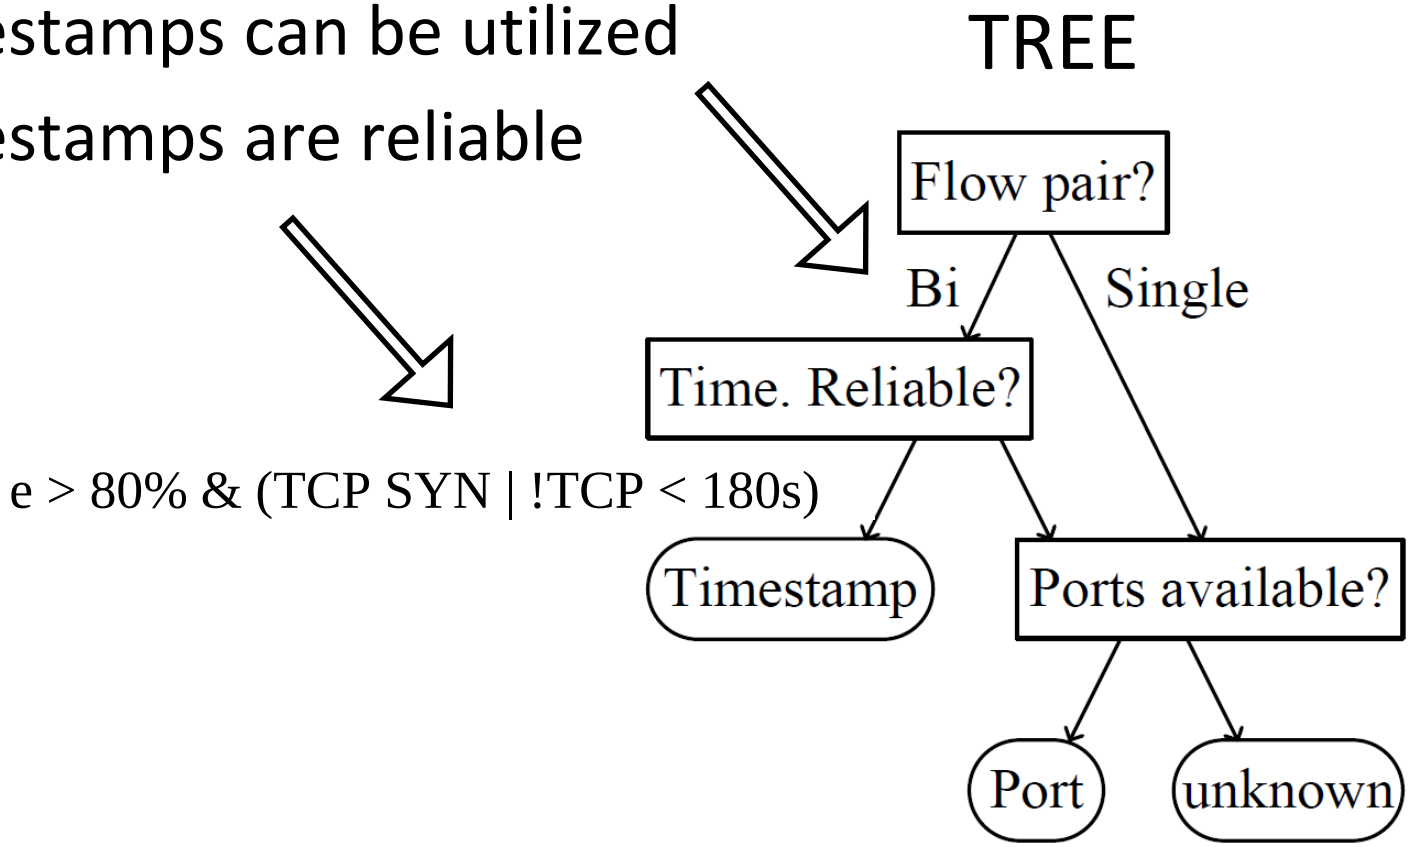

Results

Flow type	Classified by	Flows [mil]	PORT	TREE
Single flow	port	134	60%	60%
	unknown		40%	40%
Bi. flow	port	132	88%	39%
	timestamp		0%	57%
	unknown		12%	4%
Bi. flow	errors		8%	0%

Results

- Swap timestamps, observe orientation

		A_{p_w}					
	A	0.5	0.2	0.1	0.05	0.02	0.01
port	39%[t]	88%[p]	88%[p]	39%[t]	39%[t]	39%[t]	39%[t]
timestamp	57%[t]	0%[p]	0%[p]	57%[t]	57%[t]	57%[t]	57%[t]
unknown	4%[t]	12%[p]	12%[p]	4%[t]	4%[t]	4%[t]	4%[t]
errors	0%[t]	8%[p] 29%[t]	8%[p] 12%[t]	6%[t]	3%[t]	1%[t]	1%[t]

Conclusion

- For each flow exporter decide if timestamps are reliable
- Self-adapt timestamp utilization based on reliability
- Future work
  - Write NfSen patch