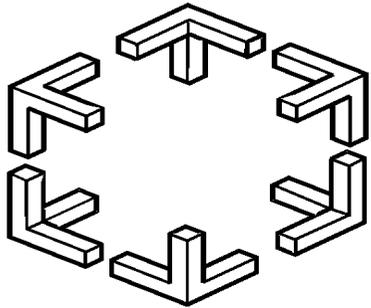


EuroSys 2020



Delegation Sketch: a Parallel Design with Support for Fast and Accurate Concurrent Operations

Charalampos Stylianopoulos, Ivan Walulya, Magnus Almgren,
Olaf Landsiedel and Marina Papatriantafilou

Chalmers University of Technology, Sweden



CHALMERS

Distributed Computing and Systems



Motivation

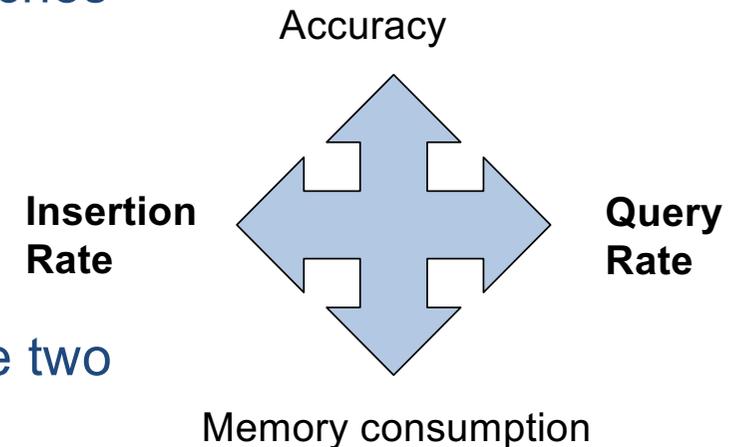
Example:

At any point in time, report how many packets from a given IP have been seen

- Exact answer \Rightarrow space proportional to number of unique IPs
- Approximate answer \Rightarrow constant space
 - Enough to summarize the input stream
 - **Sketches** are established tools for that

Challenges

- We focus on **how to parallelize sketches**
- Parallelization is necessary:
 - High-speed networks \Rightarrow Mops/sec on the sketch
 - Many-core platforms \Rightarrow underutilized in sketches
- 4-way tradeoff:
 - Applications require fast insertions **and** queries
 - E.g. intrusion detection, traffic scheduling
 - Most parallel approaches focus on one of the two (with the exception of recent work[1])



Our Work

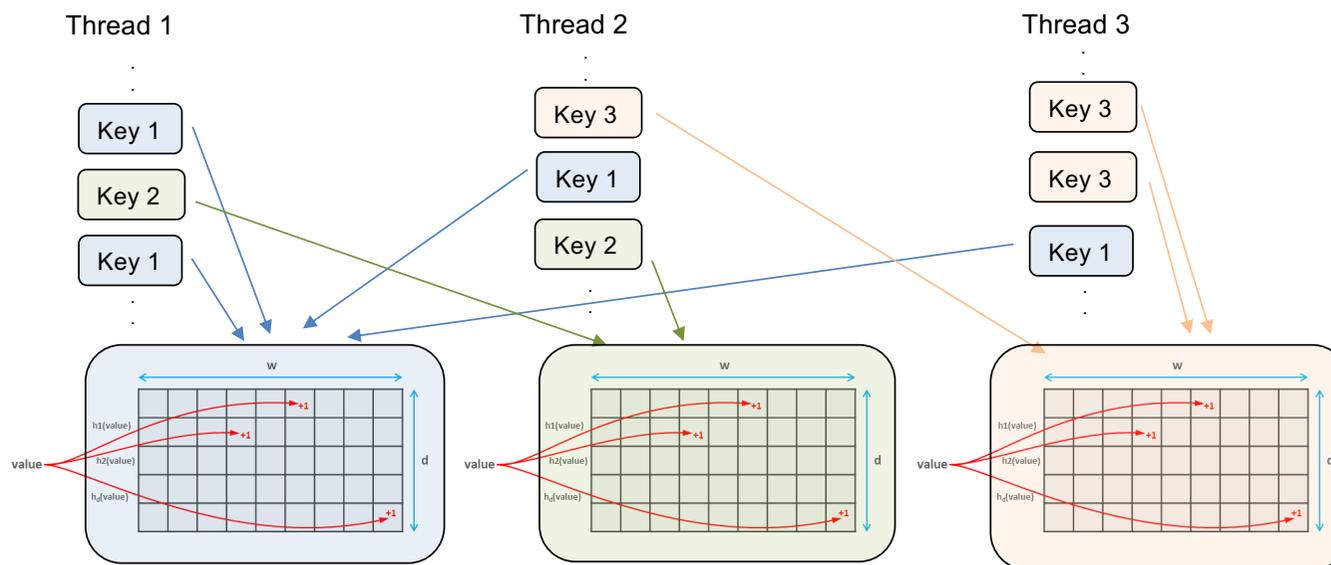
Delegation Sketch: A parallelization design for sketches

Contributions:

- **Concurrent Insertions and queries**, at high rates
- Maintains high accuracy and low memory consumption
- Scales better than state-of-the art on **hundreds of cores**

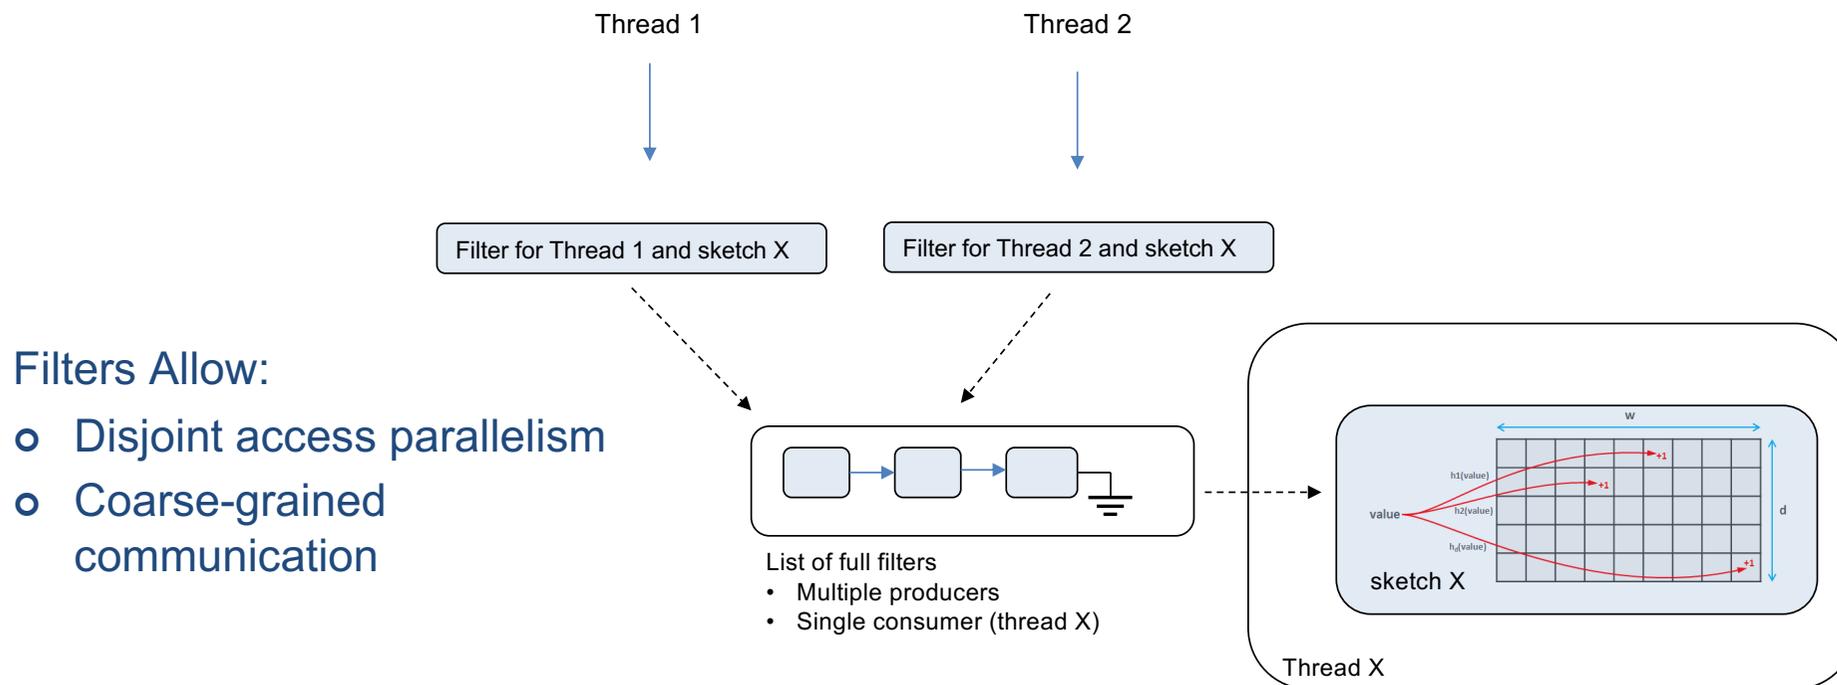
Delegation Sketch: Domain splitting

- Every key in the input domain is assigned an “owner” sketch and inserted there
- Queries are fast and accurate:
 - every key is in a specific sketch



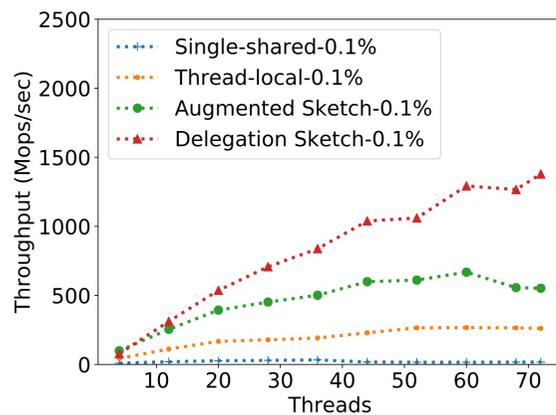
Delegation Sketch: Delegation and Combining

- Threads aggregate multiple keys into filters locally, without communication
- Filters are the units of synchronization
- Full filters are delegated to the “owner” sketch

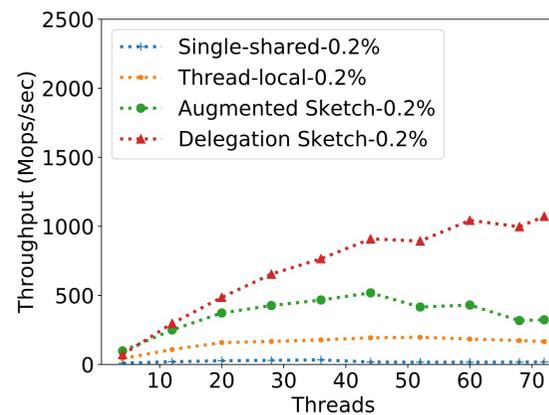


Evaluation Results

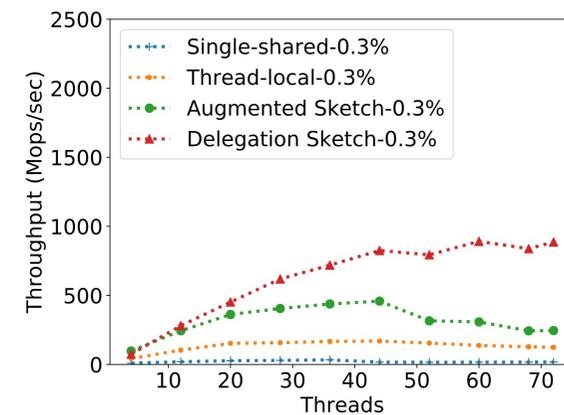
Parameters: # threads & query rate



0.1% queries



0.2% queries



0.3% queries

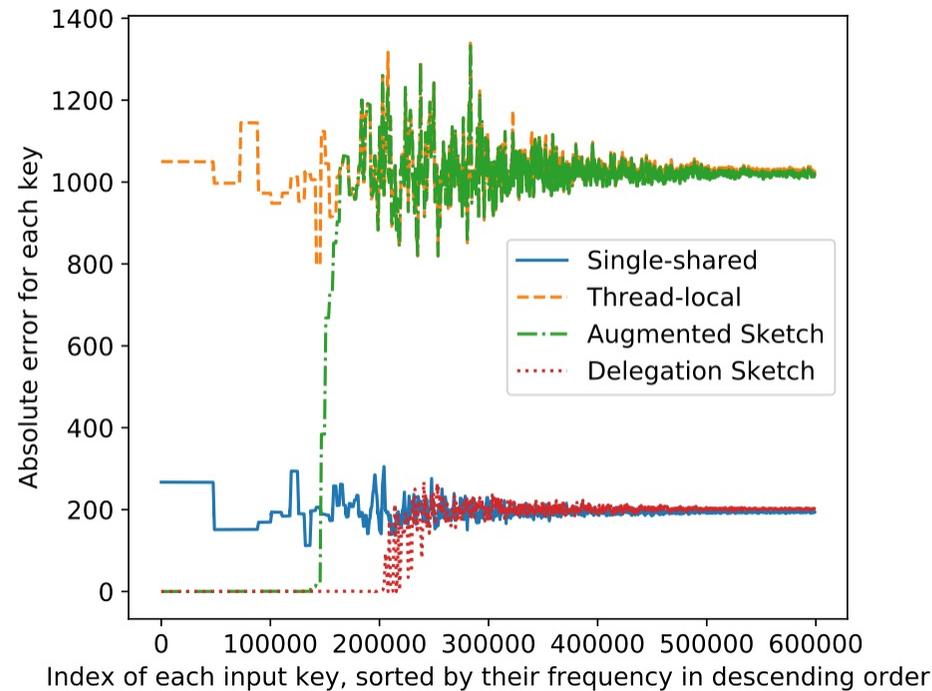
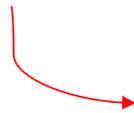
Delegation Sketch:

- **2-4X higher relative speedup** at higher query rates
- **Better scaling** at higher query rates

Evaluation Results

Accuracy

Error for each key
(lower is better)



Delegation sketch:

- As accurate as the most accurate baseline (single-shared)

All keys in the stream

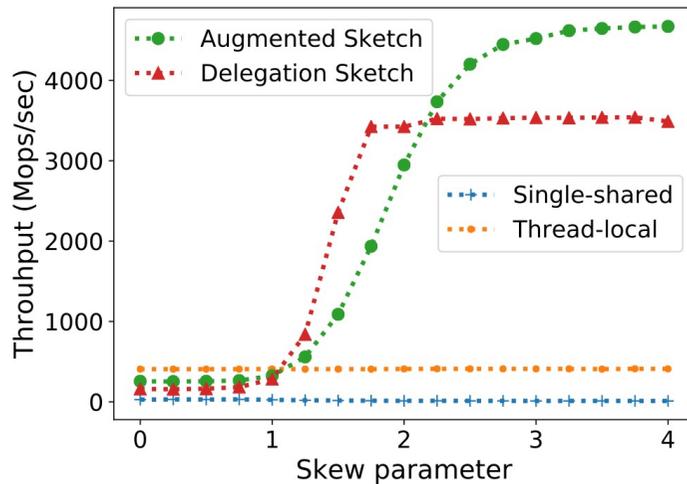
Delegation Sketch

- A parallelization scheme for high-rate traffic summarization
- Supports both **insertions and queries**
- Up to 2-4X higher throughput than next fastest baseline
- As accurate as the most accurate baseline

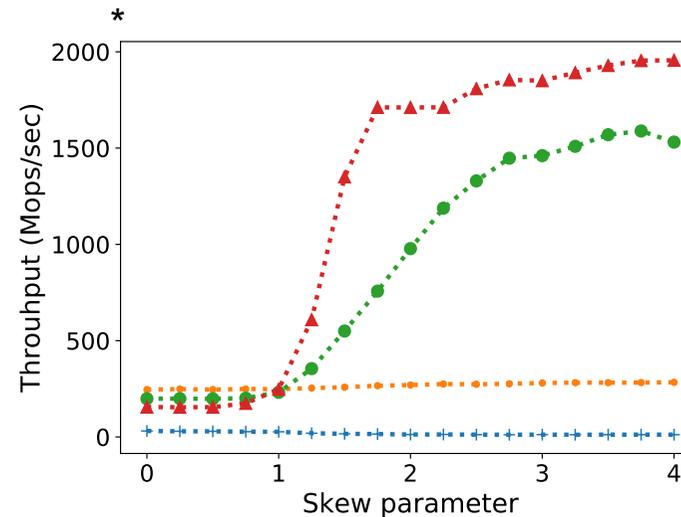
Backup slides

Evaluation Results

Parameter: input skew



0.0% queries



0.1% queries

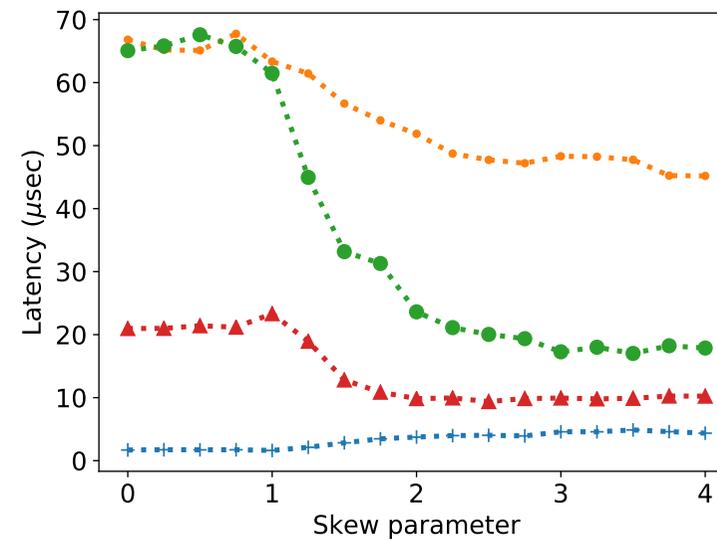
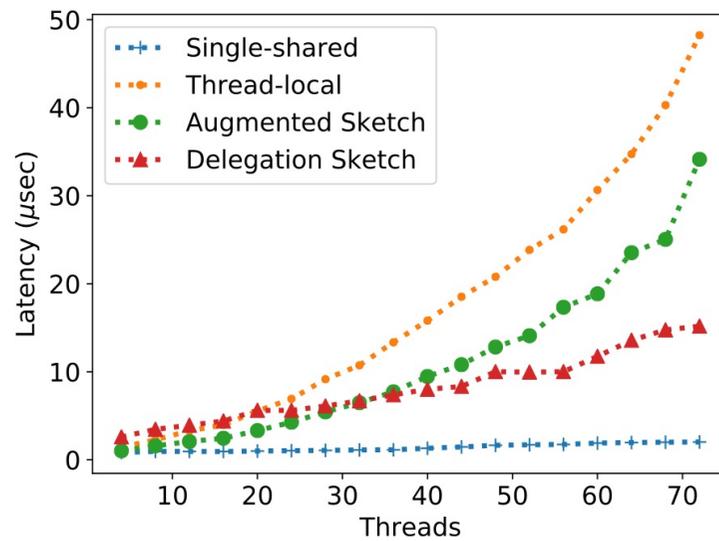
*Note: different y-axis range

Delegation Sketch:

- o **High throughput** at medium-high skew, due to filters

Evaluation Results

Query Latency



Accuracy

Accuracy bounds for Delegation Sketch:

$$f(i) \leq \hat{f}(i) \leq f(i) + \epsilon N' \text{ with probability } 1 - \delta$$

$f(i)$: the true frequency of key i

$\hat{f}(i)$: the reported estimate for key i

$\epsilon = e/w$ (w = number of buckets)

N' = number of keys that hash to the same sketch

$\delta = e^{-d}$, d = number of rows

Accuracy

