Persistent Data Sketching

Zhewei Wei[1], Ge Luo [2], Ke Yi[2], Xiaoyong Du[2], Ji-Rong Wen[1]
[1] *DEKE and School of Information, Renmin University of China*
[2] *Department of Computer Science and Engineering, The Hong Kong University of Science and Technology*
Contact: zhewei@ruc.edu.cn

Motivation and Problem Statement

Persistent Database:
- Microsoft Immortal DB, SNAP, Ganymed, Skippy and LIVE
- Allow queries on past version of the database
- Linear space: store all updates

(Ephemeral) Sketches:
- Count-Min sketch, AMS sketch
- Allow queries on current data
- Answer queries approximately with sub-linear space

Persistent Sketches:
- Allow queries on historical data < Historical query $(0, t]$ / Historical window query $(s, t]$
- Sub-linear space

An Illustrating Example

URL	actual count	estimation
/	1138896	1138970
/images/space.gif	1117634	1120050
/images/dot.gif	880322	880765
/images/hm_nbg.jpg	818126	818586
/images/home_intro.anim.gif	799697	800323

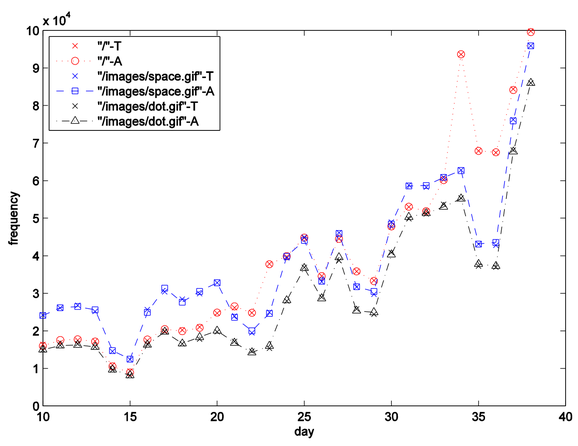

Tracking top-k items

- **1998 World Cup web site access log:**
(timestamp, IP address of the request, requested URL; size of the response, method).

- **Ephemeral sketch**
Track top-k URLs at the end of the stream

- **Persistent sketch**
Track how top-k frequencies change over time

Baseline Solution

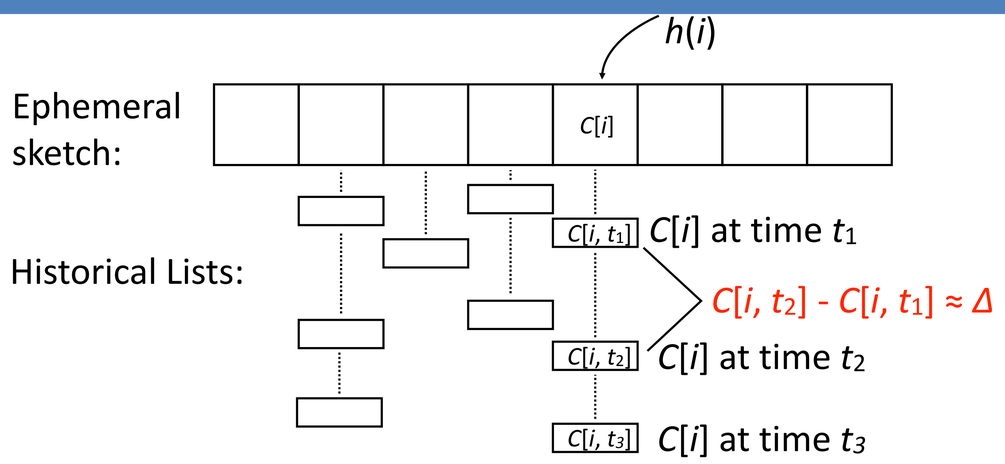

- Query for time t: find counter with timestamp closest to t
- Error: εn (ephemeral error) + Δ (persistent error)
- Space: proportional to $(1/\varepsilon + m/\Delta)$

PLA-based Count-Min Persistent Sketch

- Each counter is a discrete function according to timestamp
- (Count, time) points can be approximated with a line segment
- Approximate ratio: at most Δ

Piece-wise linear approximation (PLA):

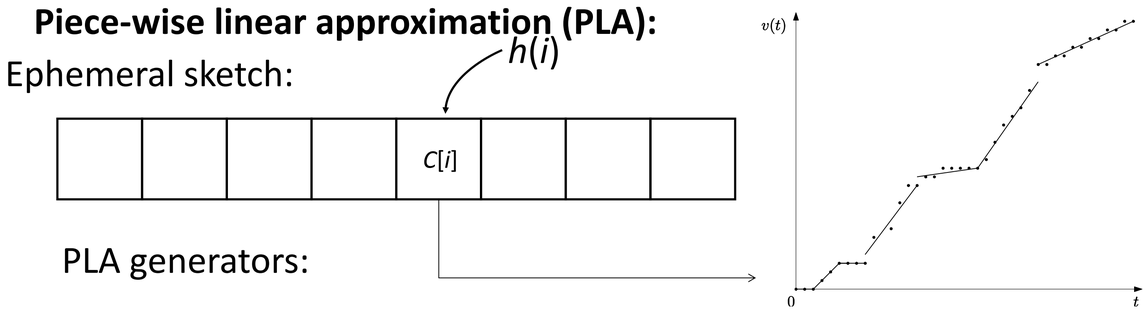

- Query for time t: return the value of the PLA function at t
- Error: εn (ephemeral error) + Δ (persistent error)
- Space: proportional to $(1/\varepsilon + m/\Delta^2)$ in random stream model
  - Random walk along the piece-wise linear function
  - Expected time to escape the linear function by Δ : Δ^2

Sampling-based AMS Persistent Sketch

- Estimating (self) join size: $\Sigma_{i\in[n]}$ ($C[i]$ + error of Δ)^2
- Bias will amplify error significantly

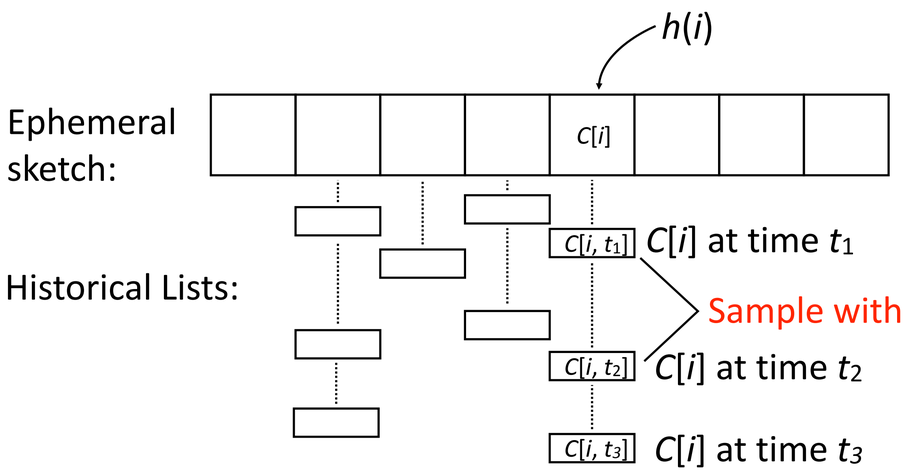

- Given a query time t
  - Set $C[i]$ to be $C[i, t_k] + \Delta - 1$ if $C[i, t_k]$ is the last sampled counts precedes t, and 0 if there is no sample preceding t.
  - Unbiased estimator for $C[i]$ at time t
- Error: εF_2 (ephemeral error) + $(\Delta/\varepsilon)^2$ (persistent error)
- Space: proportional to $(1/\varepsilon + m/\Delta)$

Experimental Results

- 7,000,000 requests from the 1998 World Cup web site access log
- Built sketches on two attributes

Attributes: Requested URL / IP address of the request

Point query

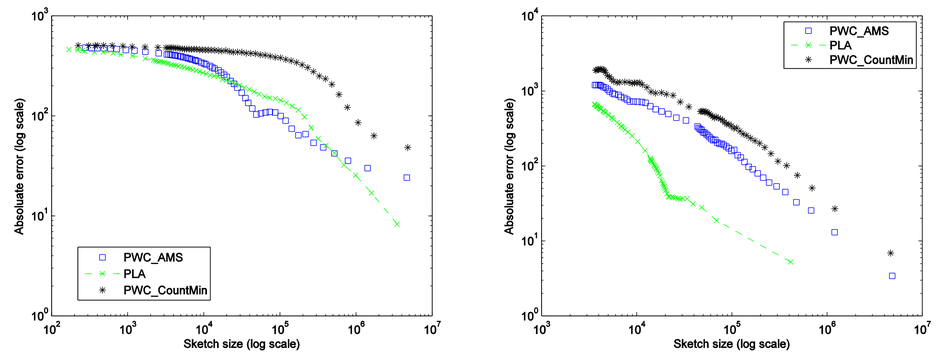

Heavy hitters query

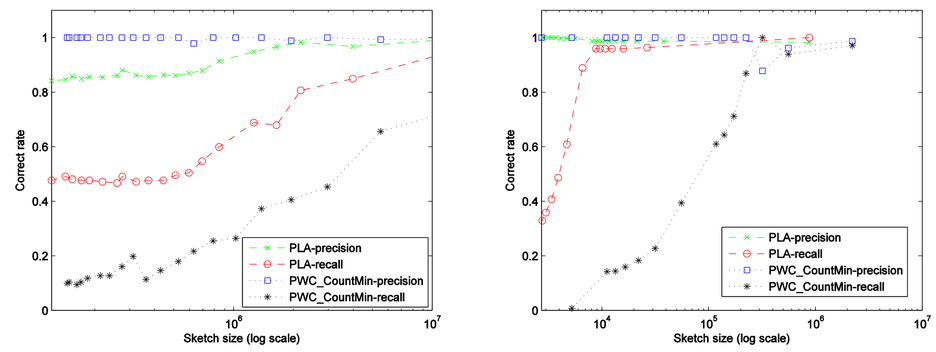

Self join size query