Flowtree: Enabling Distributed Flow Summarization at Scale

Said Jawad Saidi, MPI-Informatics
Damien Foucard, TU Berlin
Georgios Smaragdakis, TU Berlin
Anja Feldmann, MPI-Informatics / UDS

Motivation

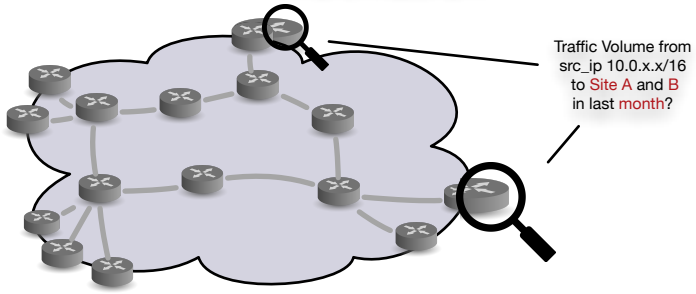

Network Management requires Network Flow Monitoring

Popular tools: NetFlow, sFlow, IPFIX

Challenges

Scalability

1. Increasing number of devices
2. Increasing storage and transfer requirements (over time and across sites)
3. Flow capture transfer is restricted: e.g. by regulation

Query Processing

1. Near real-time
2. Distributed nature of queries
3. Support of a query language
4. Interactive

Generalized Flows using Wildcards

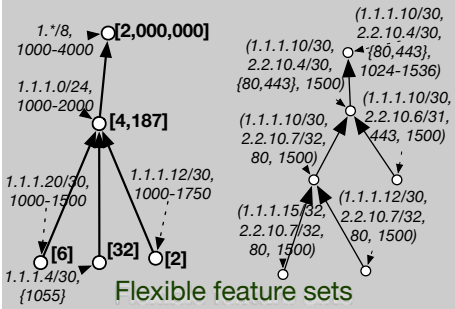

Flexible feature sets

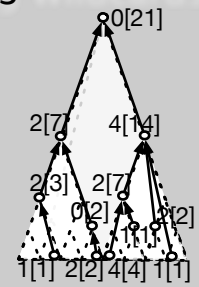

One level sums the level below

1. Captures most essential features of flows
2. Supports arbitrary feature sets hierarchies

Flowtree: Self-adjusting Data Structure

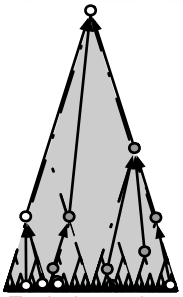

Early insertion of ancestors

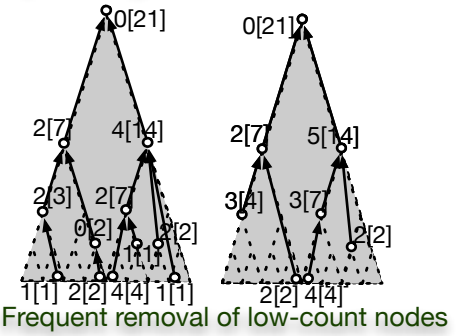

Frequent removal of low-count nodes

Beyond heavy hitters: summaries at different aggregation levels

Flowtree Operators: Merge, Diff, Compress

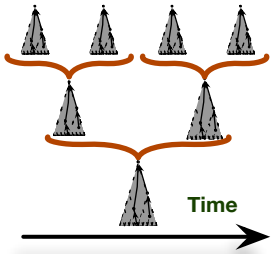

over time

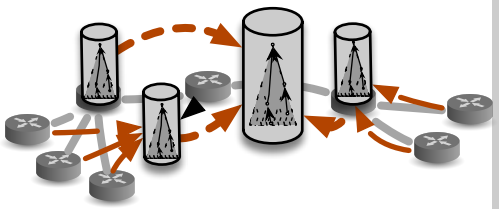

across sites

Adaptability: when volume is too high, reduce resolution

Building Near Real-time Query Processing

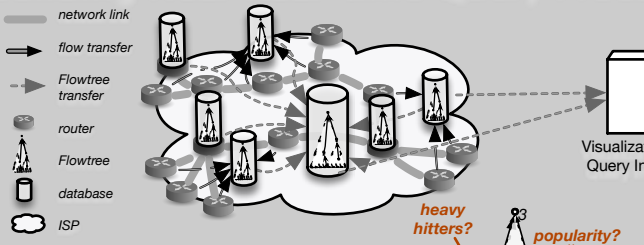

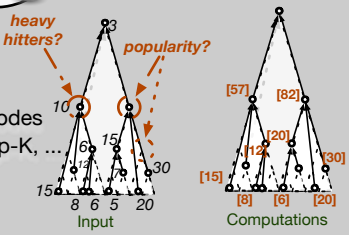

1. Approximations for prefixes that are not nodes
2. Multiple dimensions: Over time, space, Top-K, ...
3. SQL-like interface

Evaluation

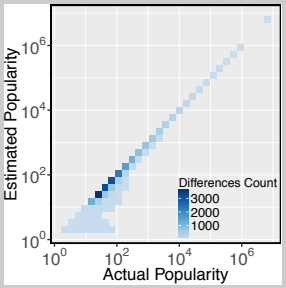

Accurate estimations even for less popular nodes

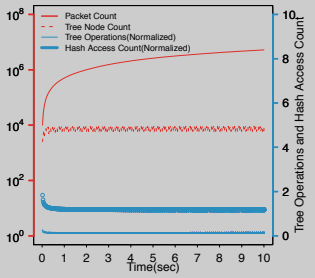

1. Fast convergence
2. Low space usage

Flowtree is accurate, fast and lightweight (up to 95% space saving)

Summary and Outlook

Summary

- Light and self-adjusting data structure for flow monitoring
- Flexible: Various feature sets, broad range of queries
- Efficient: High accuracy, Quick answers, low memory footprint, Fast convergence

Outlook

- Ongoing deployment of a full-fledged system at an IXP and an ISP

Related Work

[1] Cormode et al. "Finding hierarchical heavy hitters in streaming data." TKDD 2008
[2] Basat et al. "Constant time updates in hierarchical heavy hitters." SIGCOMM 2017
[3] Tilmans et al. "Stroboscope: Declarative Network Monitoring on a Budget." NSDI 2018
[4] Yuan et al. "Quantitative Network Monitoring with NetQRE." SIGCOMM 2017
[5] Li et al. "FlowRadar: A Better NetFlow for Data Centers." NSDI 2016