Precise Anomaly Detection in Network Flows

Sriharsha Gangam[*], Puneet Sharma[+], Sonia Fahmy[*]
Purdue University[*], HP Labs[+]
E-mail: sgangam@purdue.edu, puneet.sharma@hp.com, fahmy@purdue.edu

Problem

- Detect global heavy hitters or global iceberg flows
- Example applications: Network volume anomalies, DDoS attacks, port scanning attacks, SLA violations and worms
- Example query: Destination IPs that receive more than 1% of the total traffic

Naïve Solution

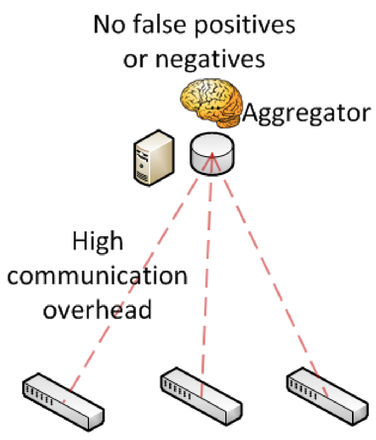

Sampling and Sketching Solutions

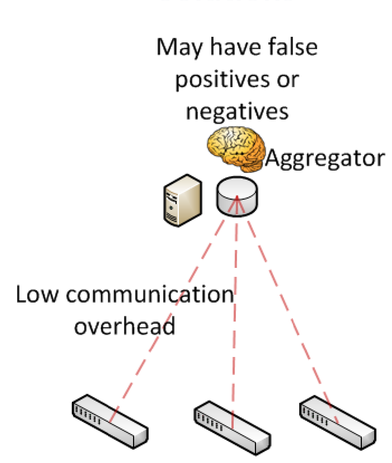

Lossy summary with on-demand recovery

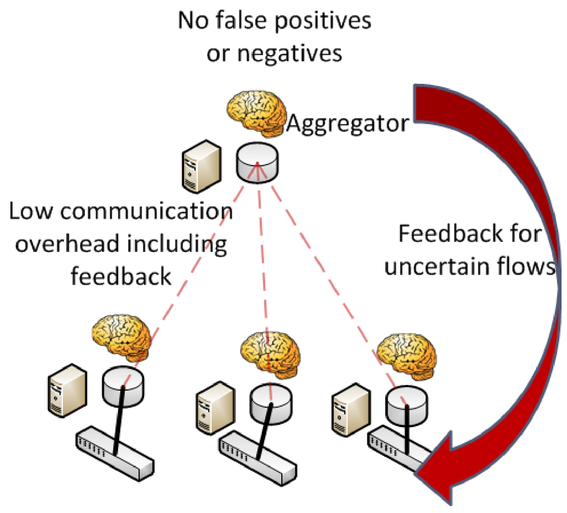

Challenges

- Large volumes of flow data for monitoring
- Sampling and sketching methods are lossy and have query agnostic parameters
- Tradeoff between accuracy and monitoring overhead

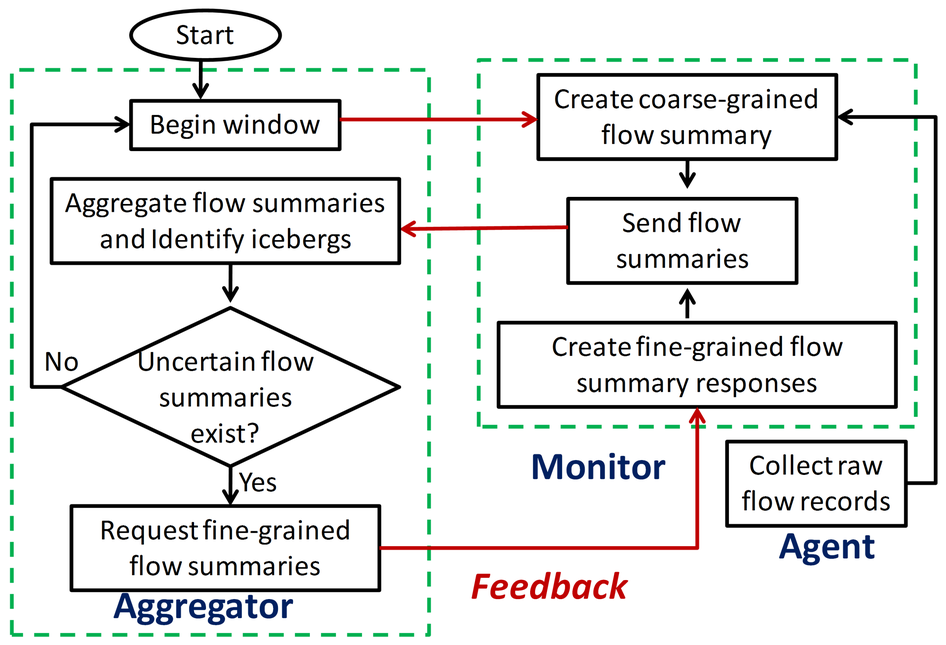

Approach

- Use of in-network processing units (blades)
- Adaptive aggregation: Send lossy summaries with on-demand recovery
- Query and traffic aware network monitoring

Comparison of Monitoring Paradigms

	NetFlow	MIND	Pegasus	OpenFlow	ProgME
Programmable/Feedback based monitoring			✓	✓	✓
Support Offline Queries	✓	✓	✓		
On-demand query processing	✓		✓		
Deployable in commodity hardware	✓	✓	✓	✓	
Low overhead/High accuracy			✓	✓	✓

Evaluation and Future Work

- Precisely detect all global icebergs (zero false positives or negatives)
- Low bandwidth overhead for Abilene Netflow traces, enterprise network sFlow traces, and PlanetLab's traffic traces
- Extension to TM and flow trajectory estimation

Abilene Trace

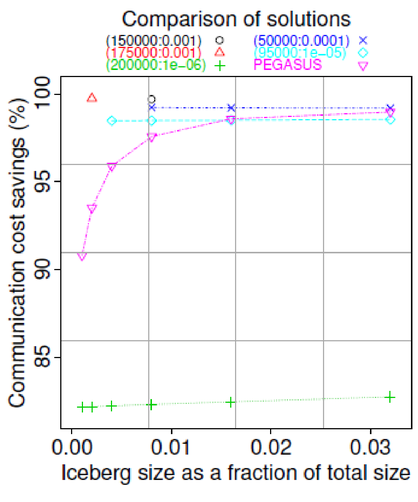

Iceberg size as a fraction of total size	Sampling parameter	Sketch parameter
0.001	200000	1e-06
0.002	175000	0.001
0.004	95000	1e-05
0.008	150000	0.001
0.016	50000	0.0001
0.032	200000	0.0001

sFlow Trace

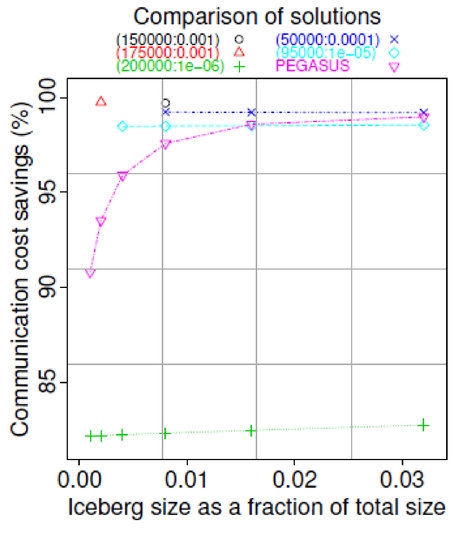

Iceberg size as a fraction of total size	Sampling parameter	Sketch parameter
0.001	50000	0.0001
0.002	175000	0.0001
0.004	190000	0.001
0.008	200000	0.0001
0.016	200000	0.001
0.032	200000	0.001