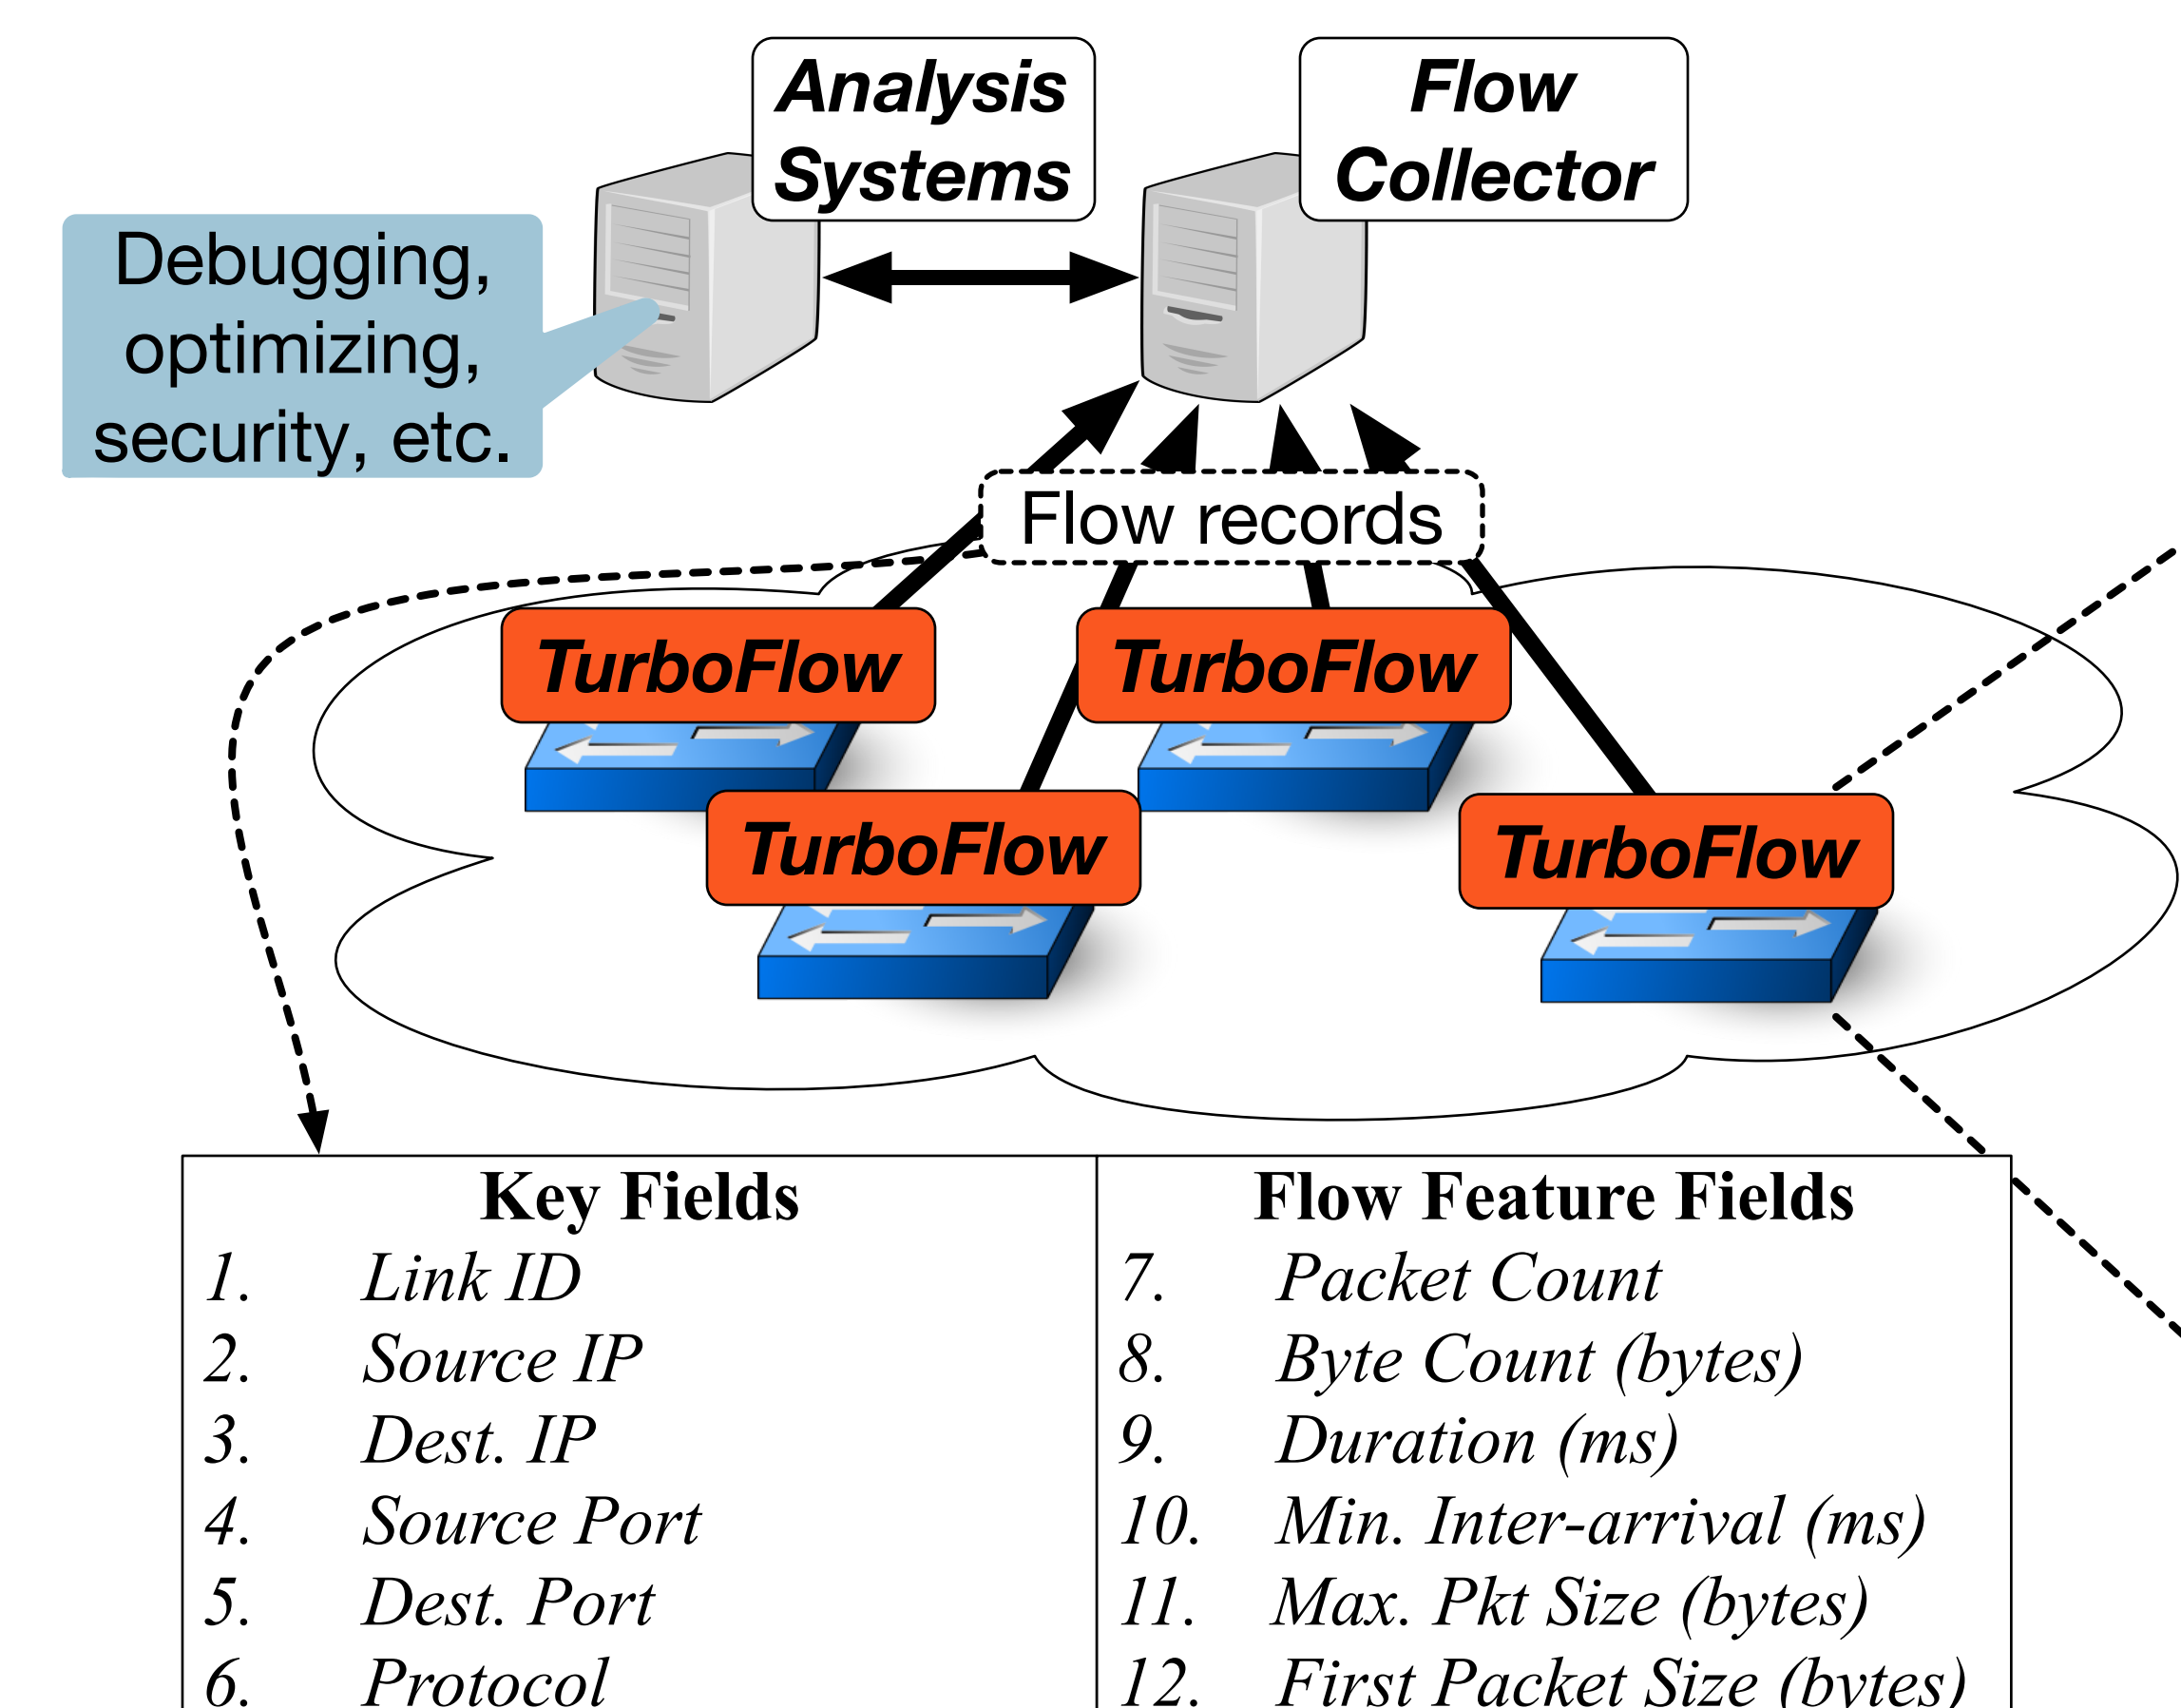


# TurboFlow: Accelerating Flow Record Generation with Commodity Switches

John Sonchack, Adam Aviv, Eric Keller, and Jonathan M. Smith

## Introduction

**TurboFlow** is a **scalable** flow monitoring system that generates **high quality flow records**.



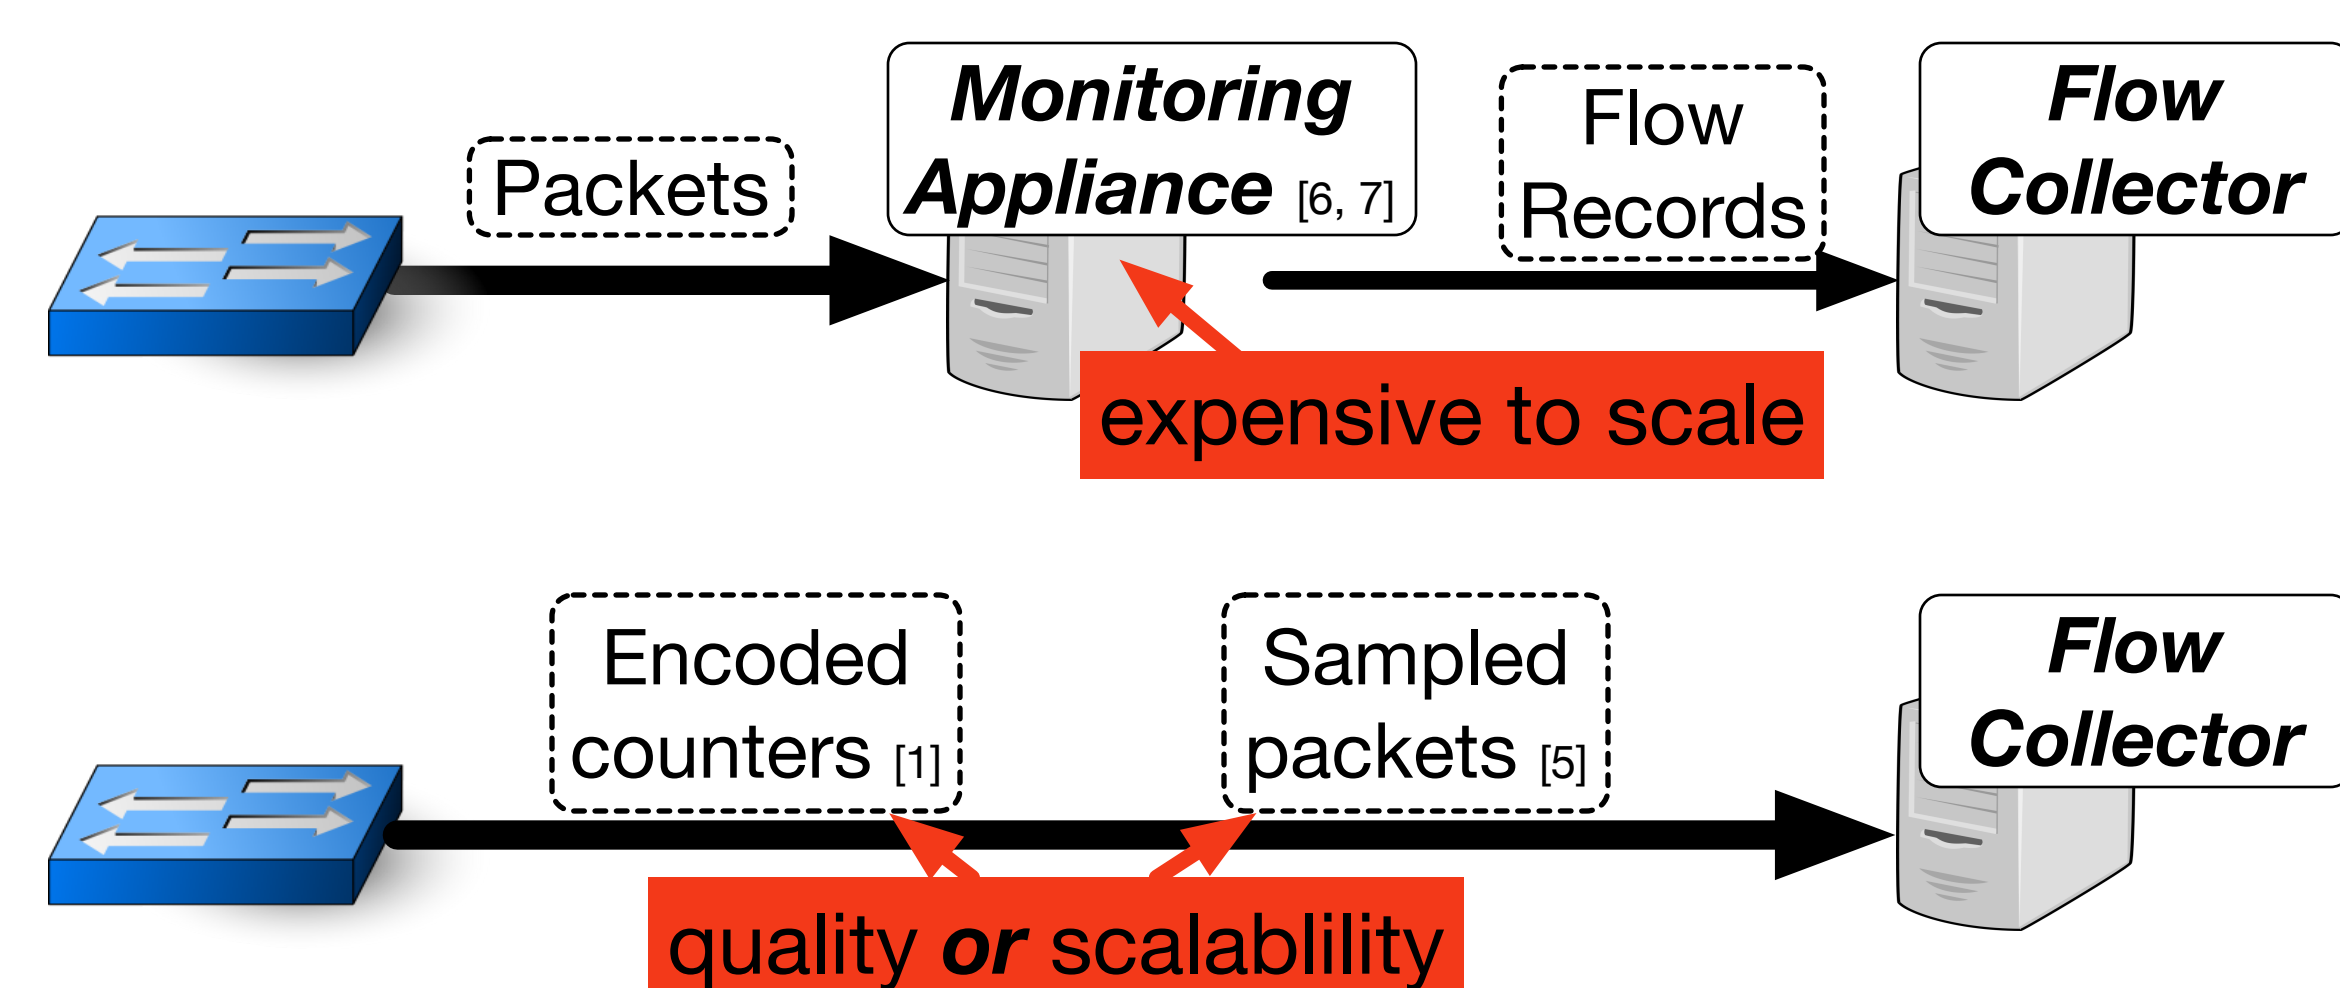
### Design Goals

**Scalability:** monitor high rate traffic from many vantage points.

**Record quality:** accurate flow records with many features.

## Background

Current systems do not meet both goals.



**Research Question:** Can we aggregate more effectively using next-generation switches with programmable forwarding engines [2,3,4]?

### Reconfigurable ASICs



>1 Tbps

### Network Processors



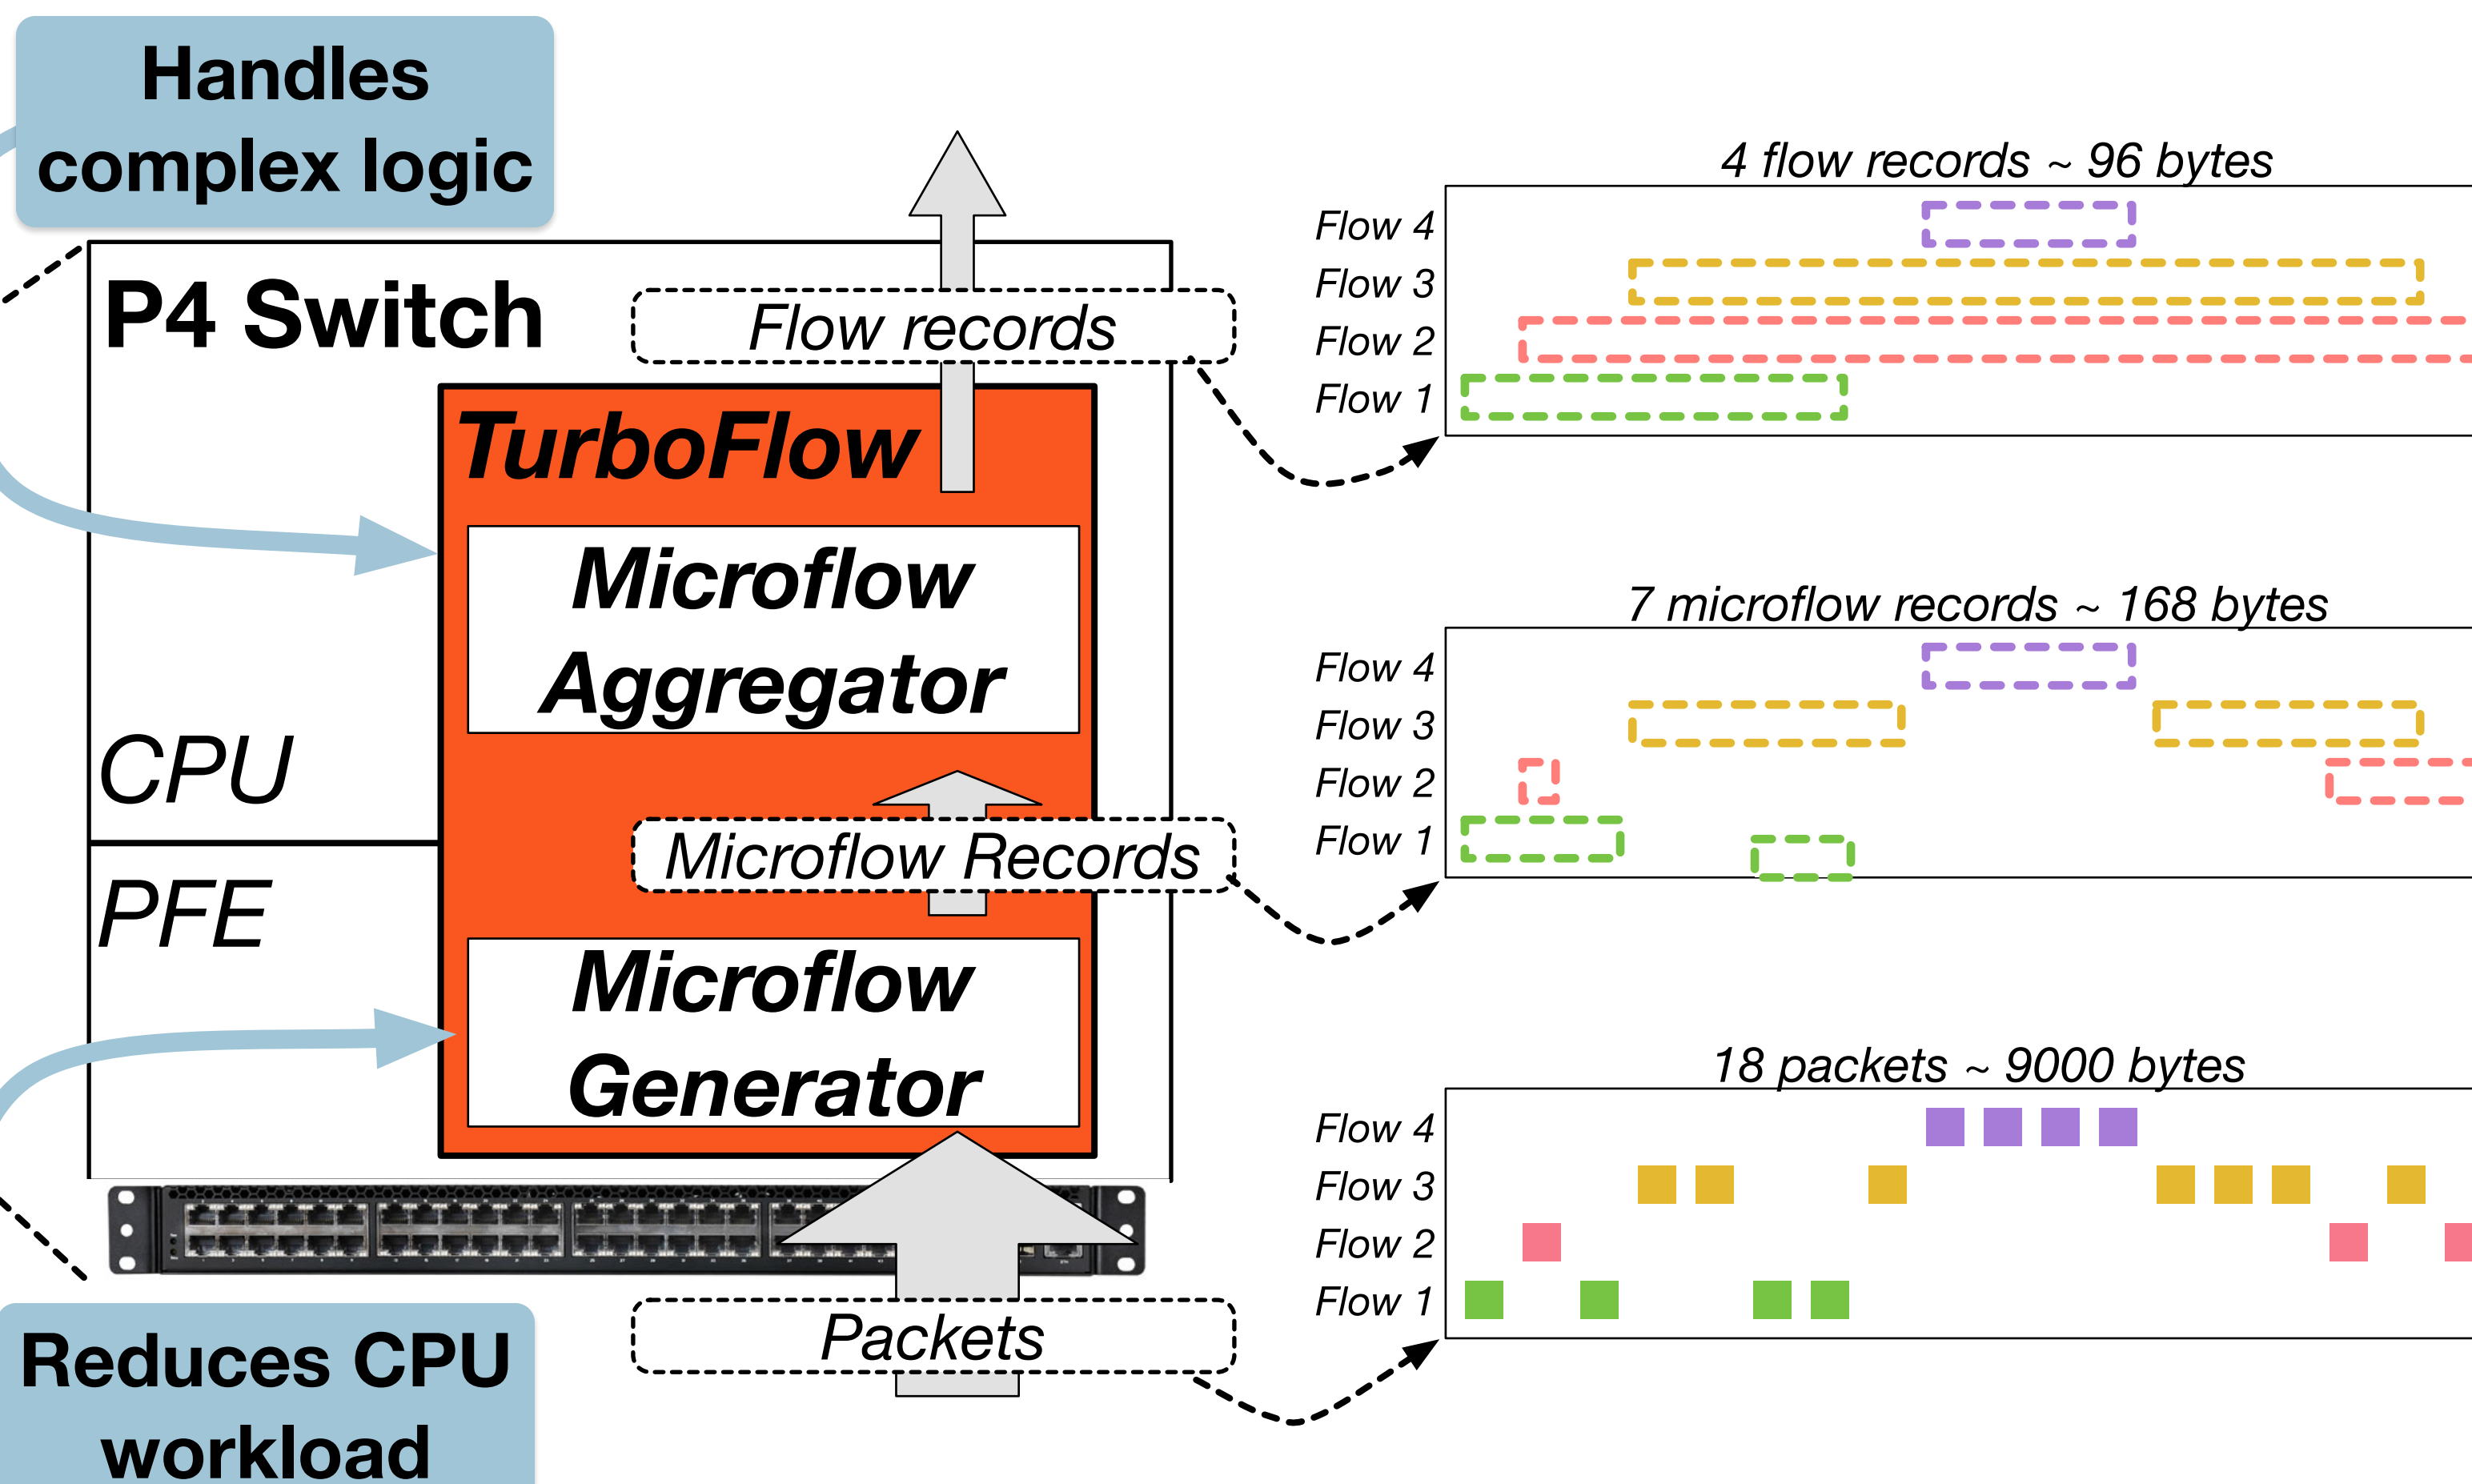
40 Gbps - 1 Tbps

**Design challenge:** PFEs have complex architectures that *restrict stateful functions*.

- Few memory updates per packet
- Parallel memory banks
- Parallel cores
- Limited code space

## TurboFlow

**TurboFlow** generates full, high quality flow records *entirely on the switch* by tightly coupling a switch's programmable forwarding engine (PFE) and CPU.



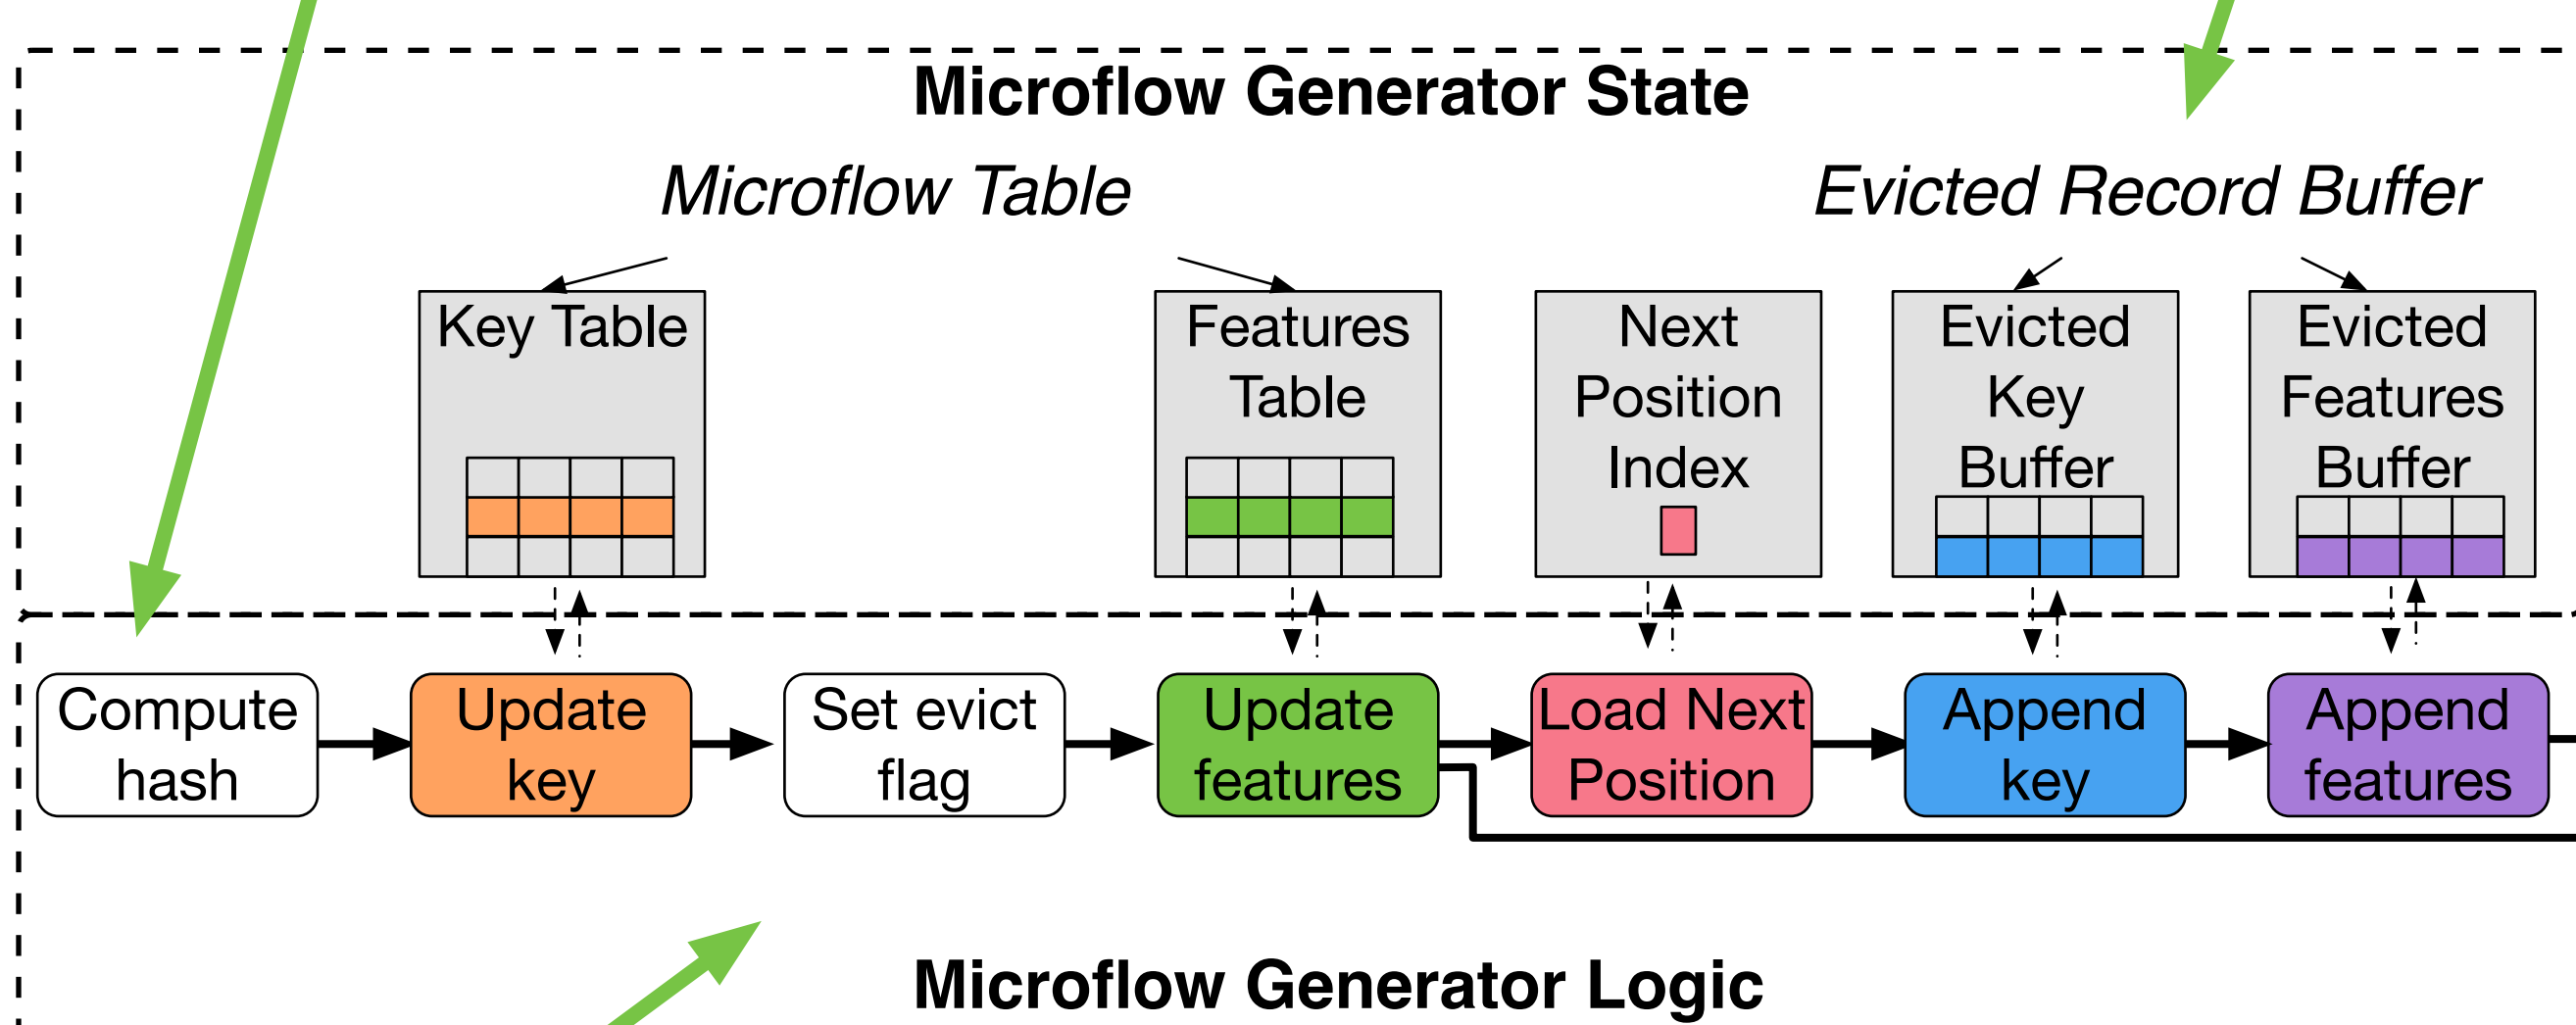
The microflow aggregator tracks flows with a standard hash table.

The microflow generator tracks flows with a specialized hash table that evicts microflows to the CPU when collisions occur.

The microflow generator maps efficiently to PFE architectures.

State access is efficient to synchronize using hash value

State partitions across parallel memory banks

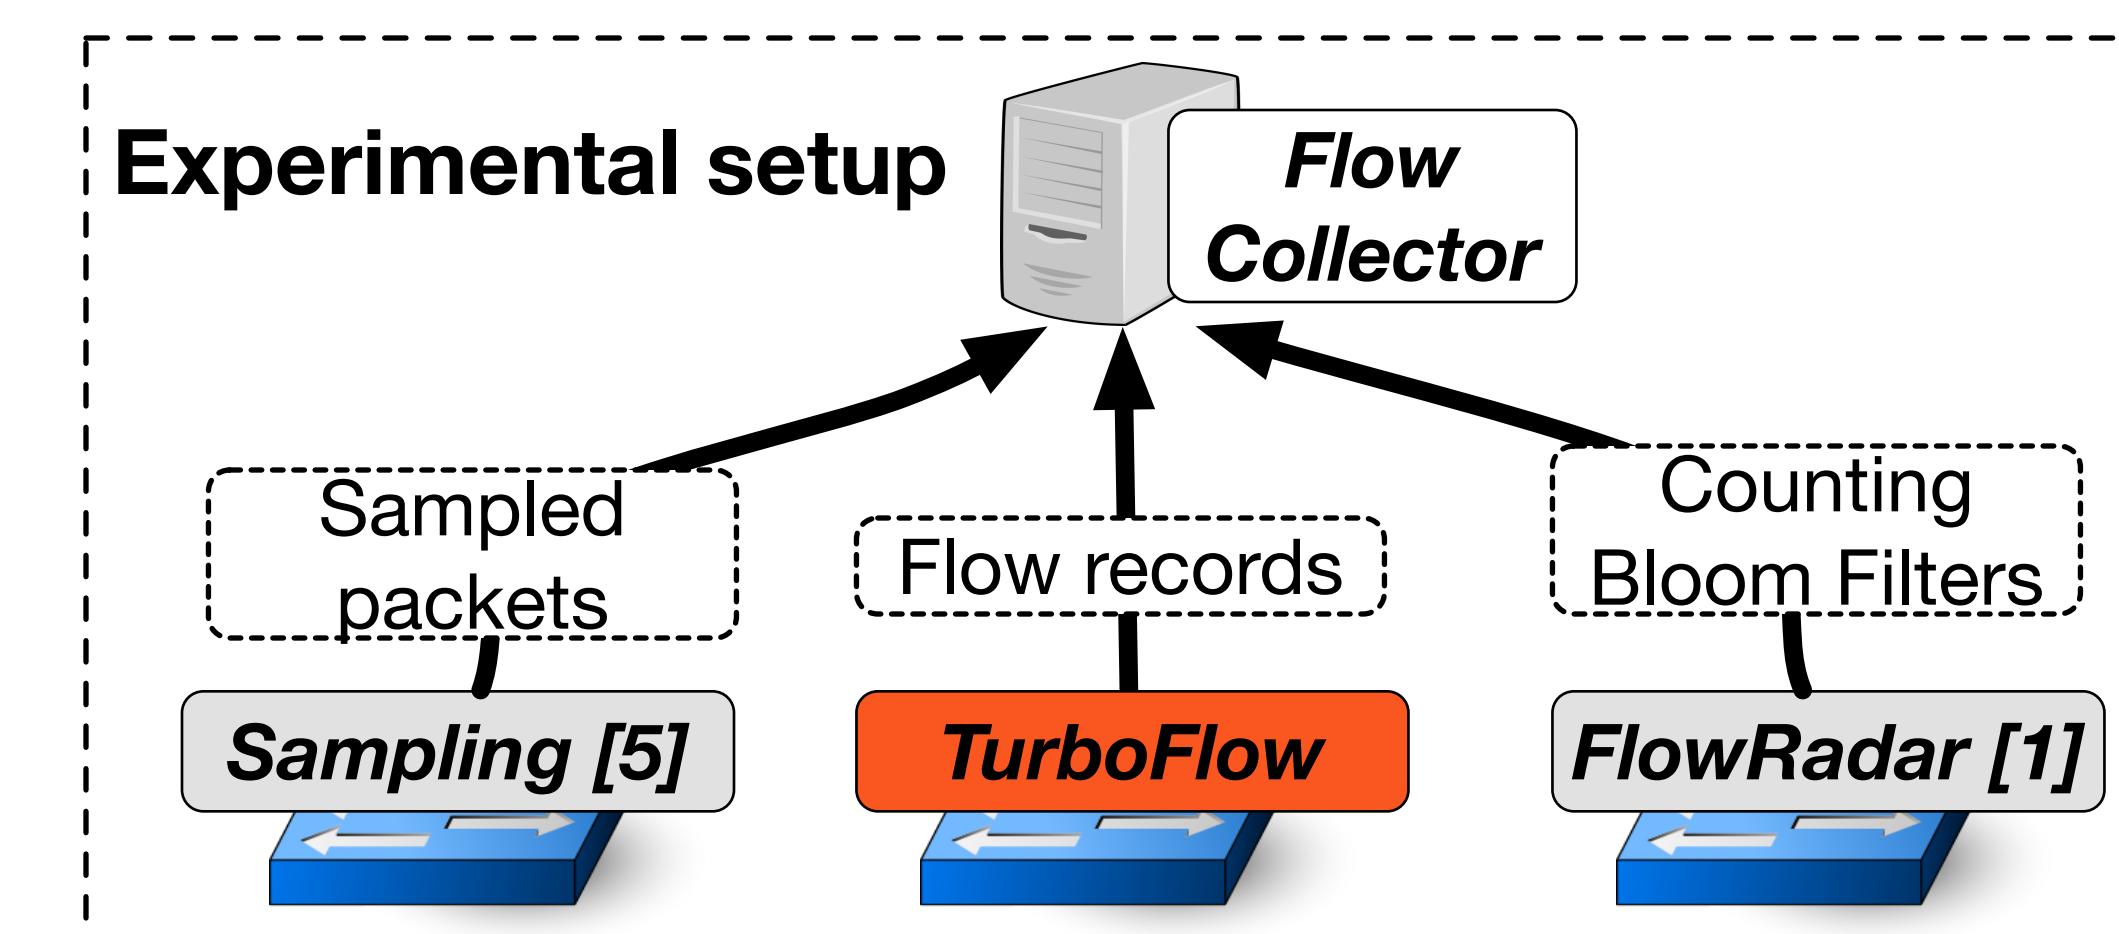


Compiles to short PFE function

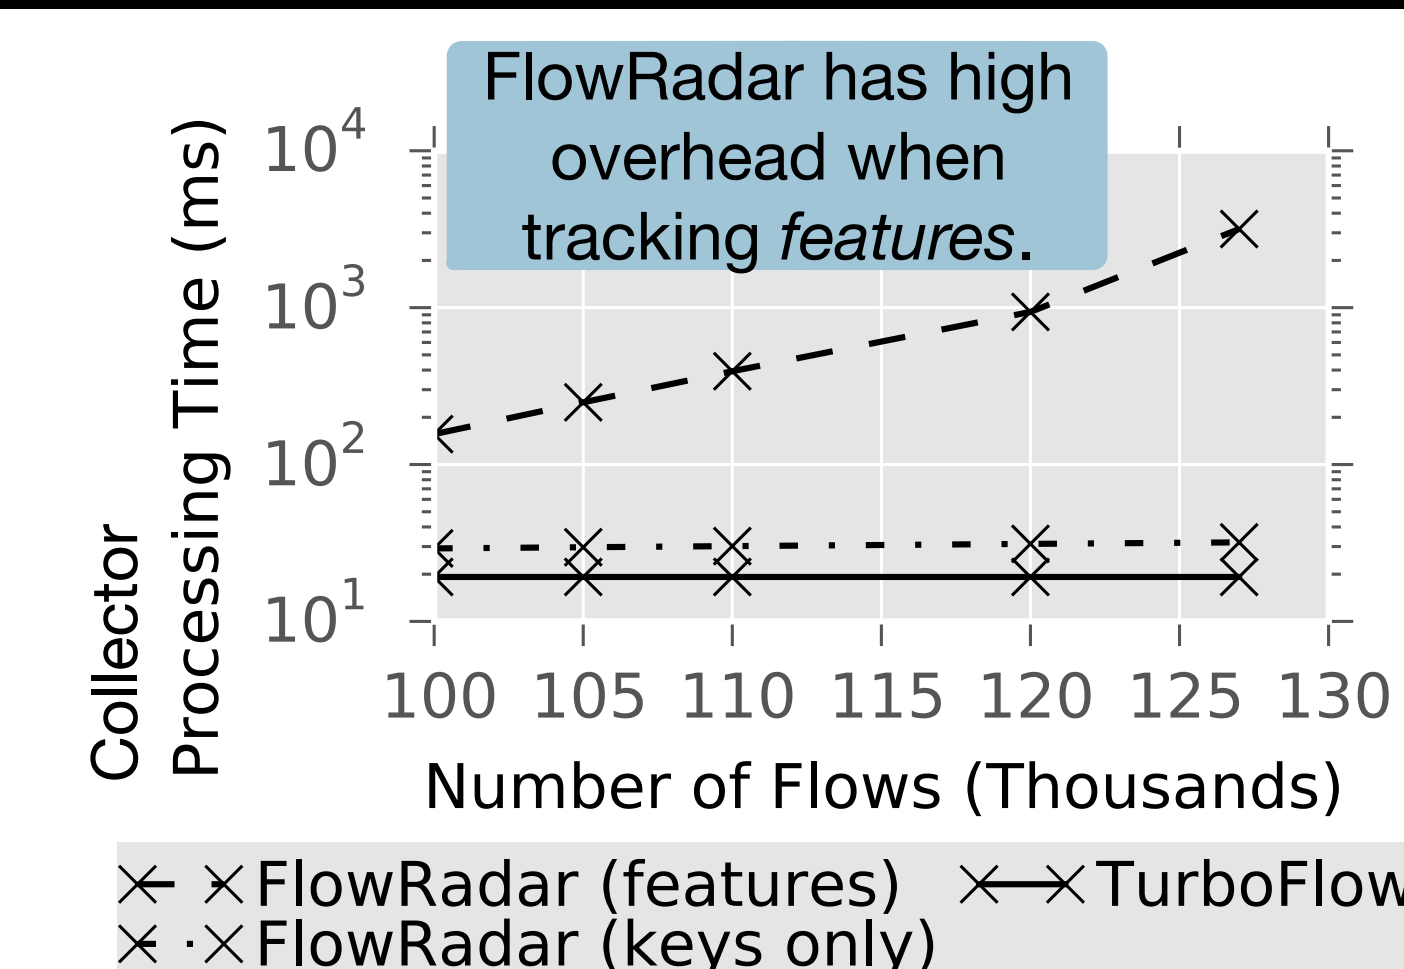
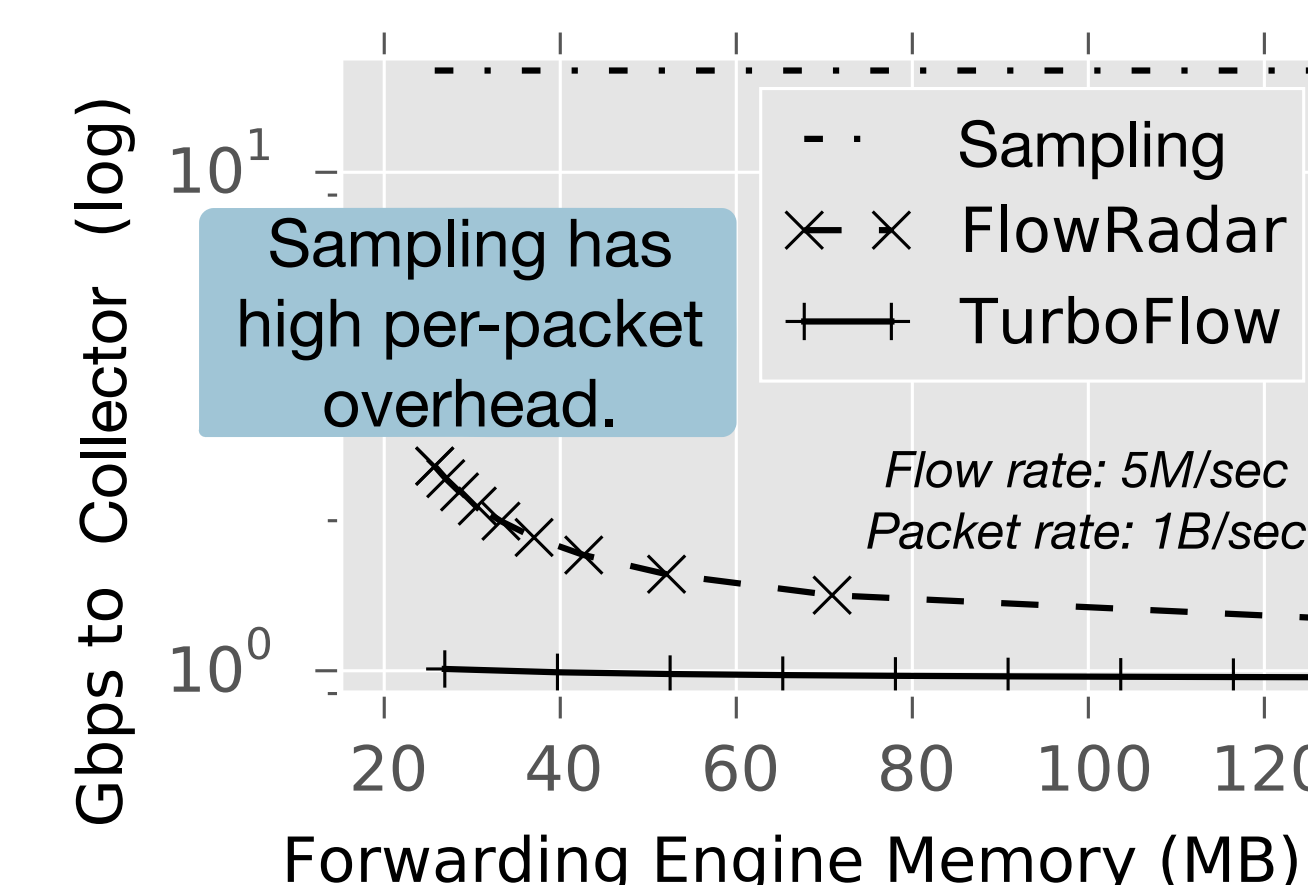
Small, constant number of memory operations

## Evaluation

TurboFlow provides excellent scalability and high quality records.



### Flow Collector Scalability



Full flow records from the switch → minimal flow collector overhead, high record quality

### Switch Overhead

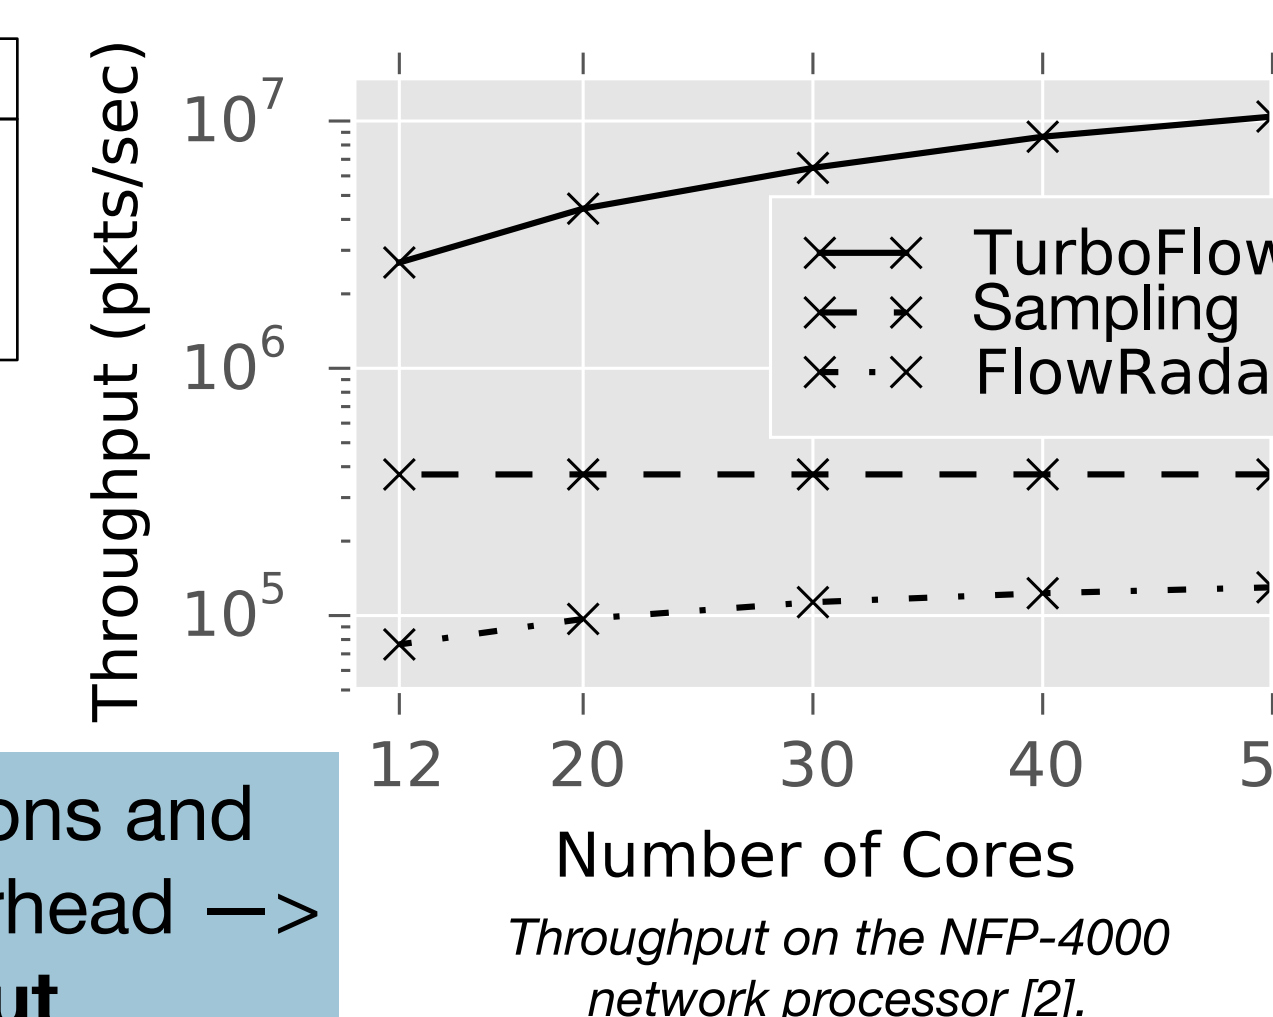
#### Switch PFEs

Metric	TurboFlow	FlowRadar	Sampling
# Memory Operations	13	35	2
% RMT [3] Stages	43.7%	does not fit	3.3%
% Banzai [4] Units	4.3%	11.6%	0.3%

Pipeline usage in programmable ASIC models, compiled with Domino [4].

Fewer memory operations → less pipeline utilization

Fewer memory operations and less synchronization overhead → higher throughput




#### Switch CPUs

Switch CPU cores	2 cores	3 cores	4 cores
TurboFlow	30.71 M	46.53 M	53.12 M
FlowRadar	7.47 M	7.47 M	7.47 M
Sampling	1.62 M	1.62 M	1.62 M

Sustained flow rate using 24MB of PFE memory.

Greater PFE aggregation → higher flow rates and more efficient use of switch CPU



### Citations

- [1] Y. Li, R. Miao, C. Kim, and M. Yu. Flowradar: a better netflow for data centers. NSDI, 2016.
- [2] Netronome. Agilio Cx Intelligent Server Adapters <https://www.netronome.com/products/agilio-cx/>.
- [3] P. Bosshart et al. Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN. SIGCOMM CCR, 2013.
- [4] Sivaraman, Anirudh, et al. Packet transactions: High-level programming for line-rate switches. SIGCOMM, 2016.
- [5] Panchen, Sonia, et al. InMon corporation's sFlow: A method for monitoring traffic in switched and routed networks. 2001.
- [6] Cisco. Netflow Generation Appliance 3340. [http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/netflow-generation-3000-series-appliances/data\\_sheet\\_c78-720958.html](http://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/netflow-generation-3000-series-appliances/data_sheet_c78-720958.html)
- [7] Argus. Argus — Auditing Network Activity. <https://www.qosient.com/argus/>