BranchPipe: *Scalable Decision Trees for Stateful Processing at Line Rate*

Murayyiam Parvez, Annus Zulfiqar[1], Roman Beltiukov[2], Shir Landau Feibish[3], Walter Willinger[4], Arpit Gupta[2], Muhammad Shahbaz[1]

Purdue University, [1]University of Michigan, [2]UCSB, [3]Open University of Israel, [4]NIKSUN, Inc.

Abstract and Motivation

- Machine learning is increasingly being deployed in RMT switches to enable real-time, data-driven decision-making directly in the data plane
- We present **BrachPipe: a framework for scalable Decision Tree (DT) inference at line rate**, designed to support 5x more features than SOTA and scale effectively across millions of flows

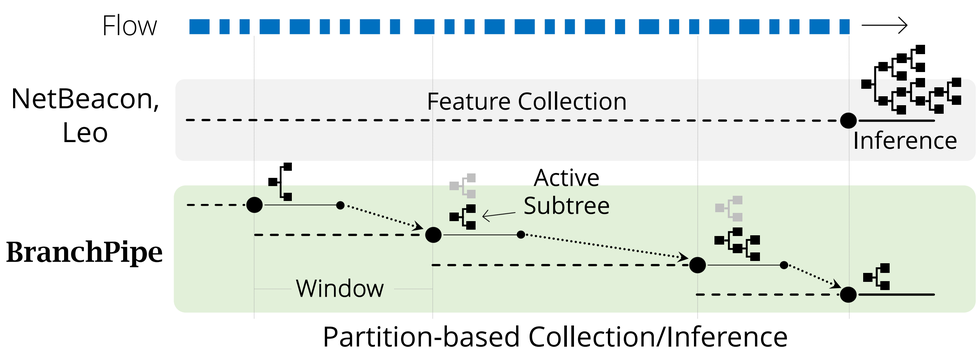

Figure 1: Comparison of in-network Decision Tree classification approaches

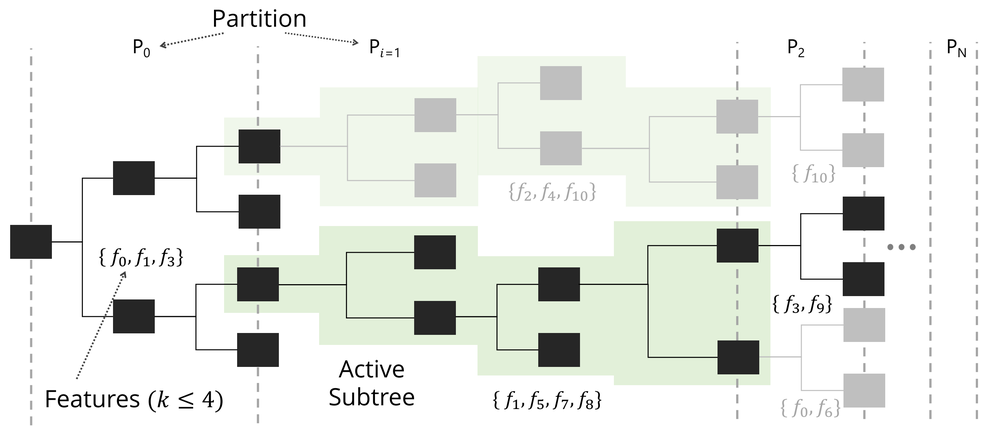

Figure 2: Partitions (P_i) comprising multiple subtrees, each with its own subset of top-k features

- By leveraging **recirculation**, programmable switches become a time-shared resource that can **dynamically store features** and perform **inference** only for the **active subtree** in each **partition**

Design of BranchPipe

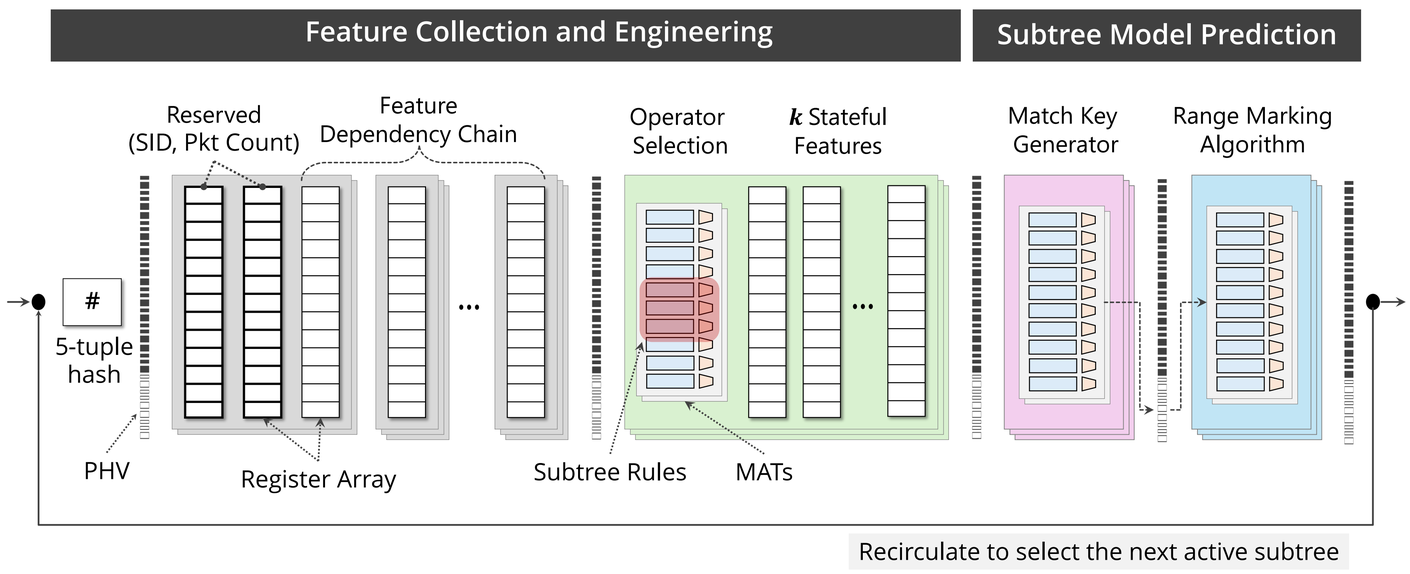

Figure 3: A high-level view of BranchPipe's inference architecture

BranchPipe framework first uses the **subtree ID** as match key to **collect** the top-k **features** for each subtree, **dynamically updating** them at runtime as the active subtree changes. It uses the **subtree inference** to select the **next subtree** until a label is found, using **recirculation** to reset state and activate the next subtree in the switch.

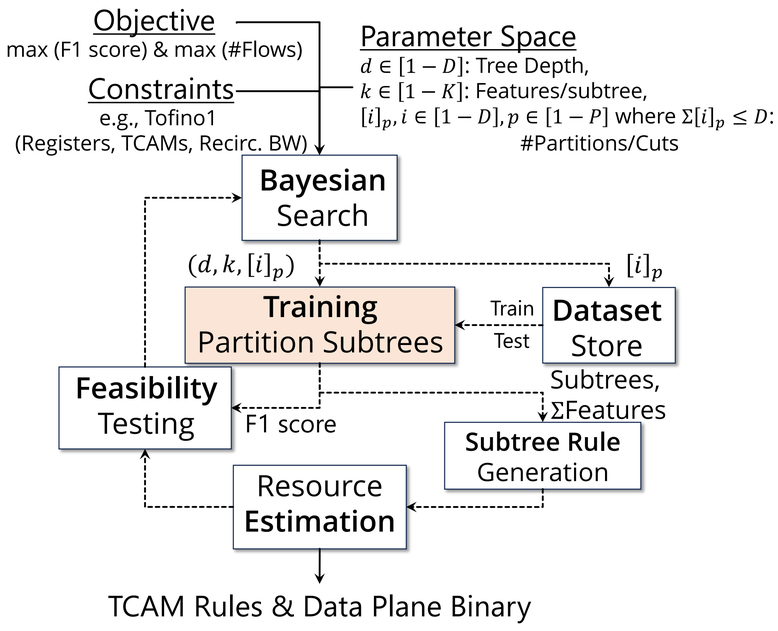

Figure 4: BranchPipe's model design search workflow

BranchPipe uses **Bayesian Optimization** to efficiently train DTs, tuning **feature count, tree depth,** and **partitioning strategy** to find Pareto-optimal models balancing accuracy and flow count.

Evaluation and Results

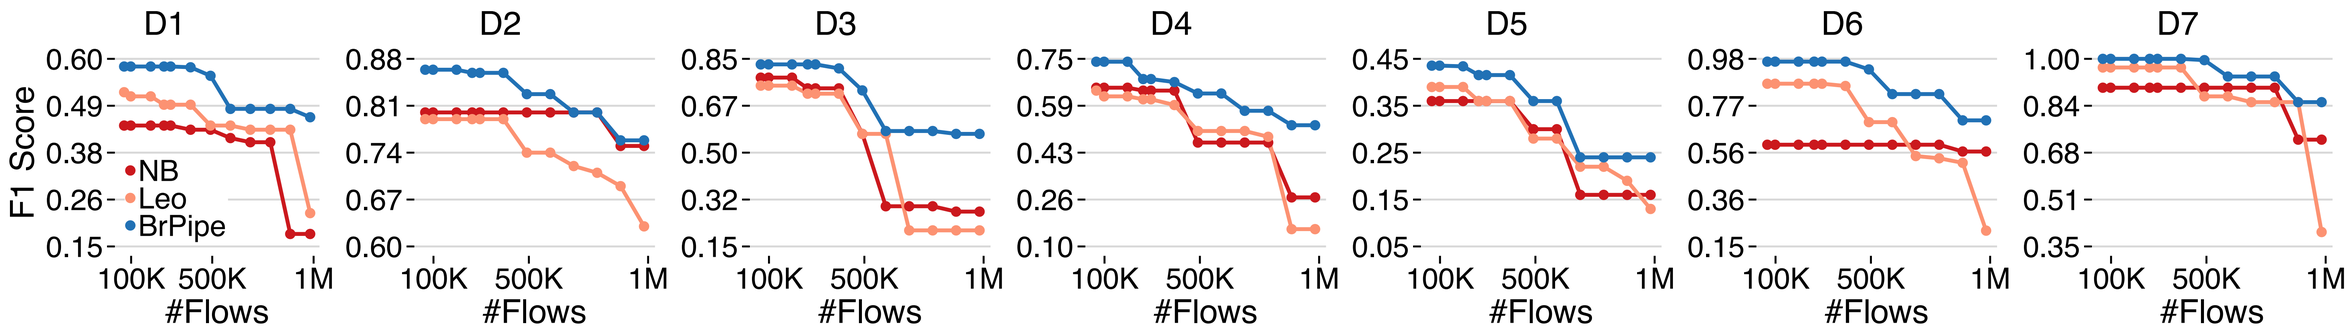

(a) Pareto frontier of BranchPipe vs. baselines, indicating the best F1 score for a given number of flows in the data plane

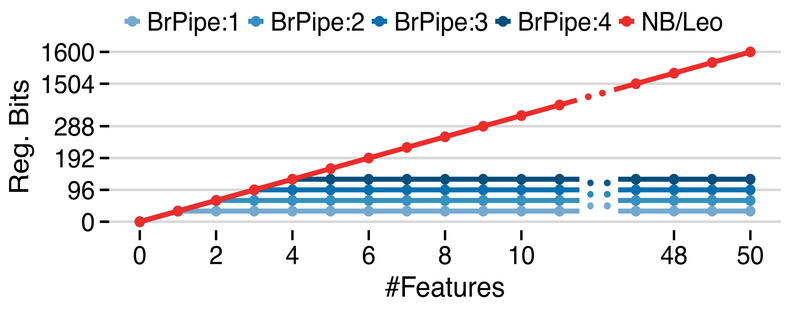

(b) Register size (bits) vs number of features supported by each model. BranchPipe:*k* is a partitioned tree with k features per subtree.

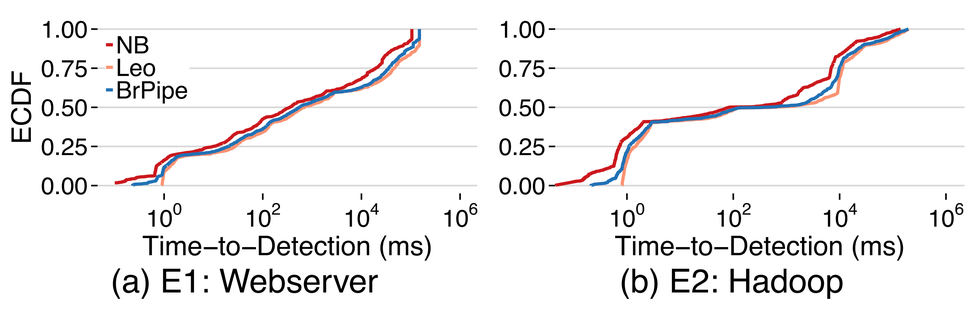

(c) Flow-level time-to-detection (TTD) of D3 for two datacenter environments (E1-E2)

Dataset	Recirculation Bandwidth (Mbps)	
	E1: Webserver	E2: Hadoop
D1	2.93 ± 2.44	5.99 ± 3.51
D2	6.01 ± 4.01	12.32 ± 5.76
D3	3.58 ± 3.21	7.33 ± 4.62

(d) Maximum recirculation bandwidth (Mbps) when processing the datasets (D1–3) for two datacenter environments (E1–2)

Contact us: parvezm@purdue.edu