iSwitch: Accelerating Distributed Reinforcement Learning with In-Switch Computing

Youjie Li Iou-Jen Liu Yifan Yuan
Deming Chen Alexander Schwing Jian Huang

University of Illinois at Urbana-Champaign
Electrical & Computer Engineering

AI Applications are Increasingly Operating in Dynamic Environments

Autonomous Driving

Robotics

Games

AI Applications are Increasingly Operating in Dynamic Environments

Autonomous Driving

Robotics

Games

Reinforcement Learning Empowers AI Applications to Take Real-Time Intelligent Actions

What is Reinforcement Learning?

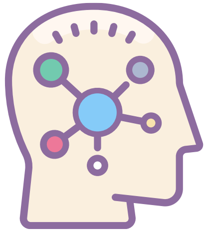

Agent

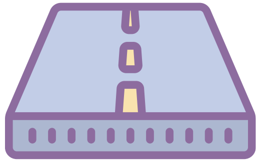

Environment

What is Reinforcement Learning?

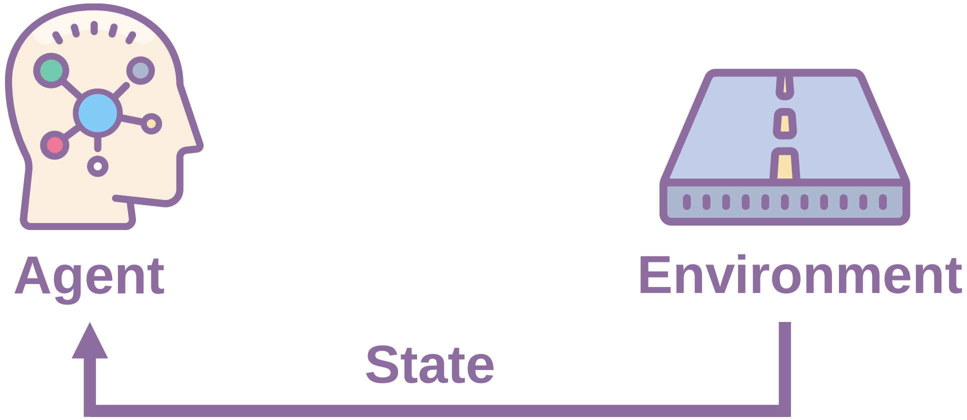

What is Reinforcement Learning?

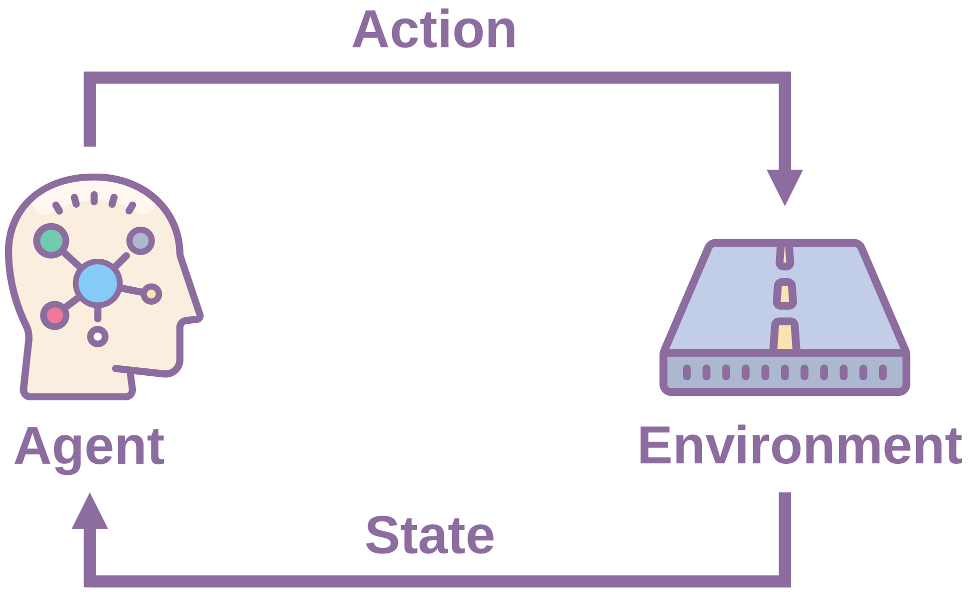

What is Reinforcement Learning?

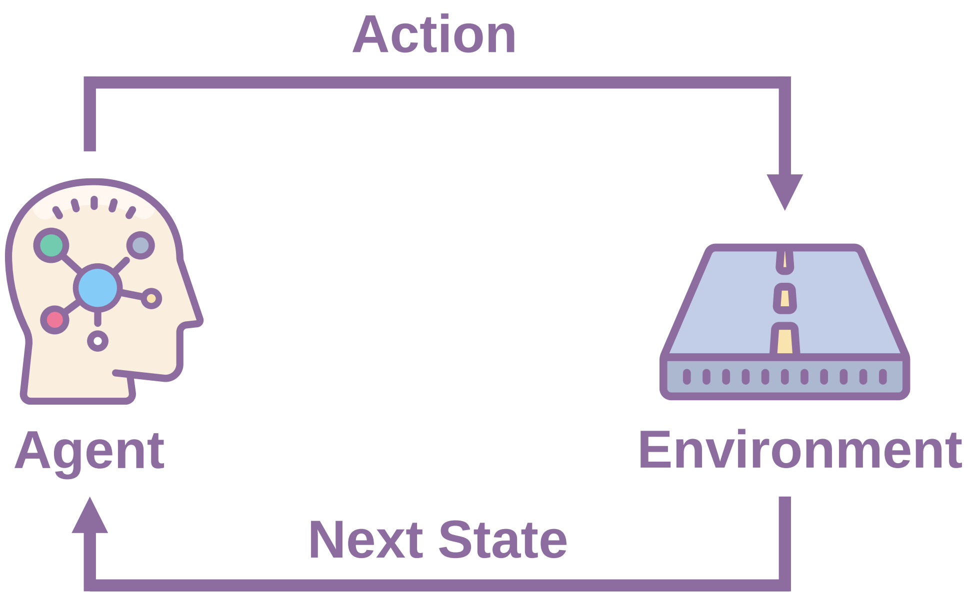

What is Reinforcement Learning?

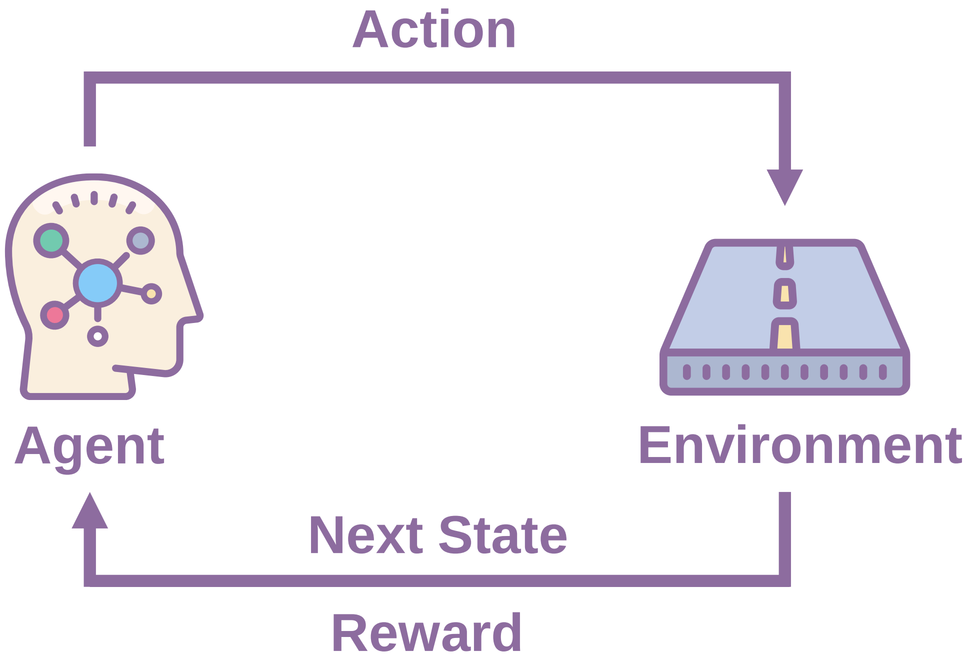

What is Reinforcement Learning?

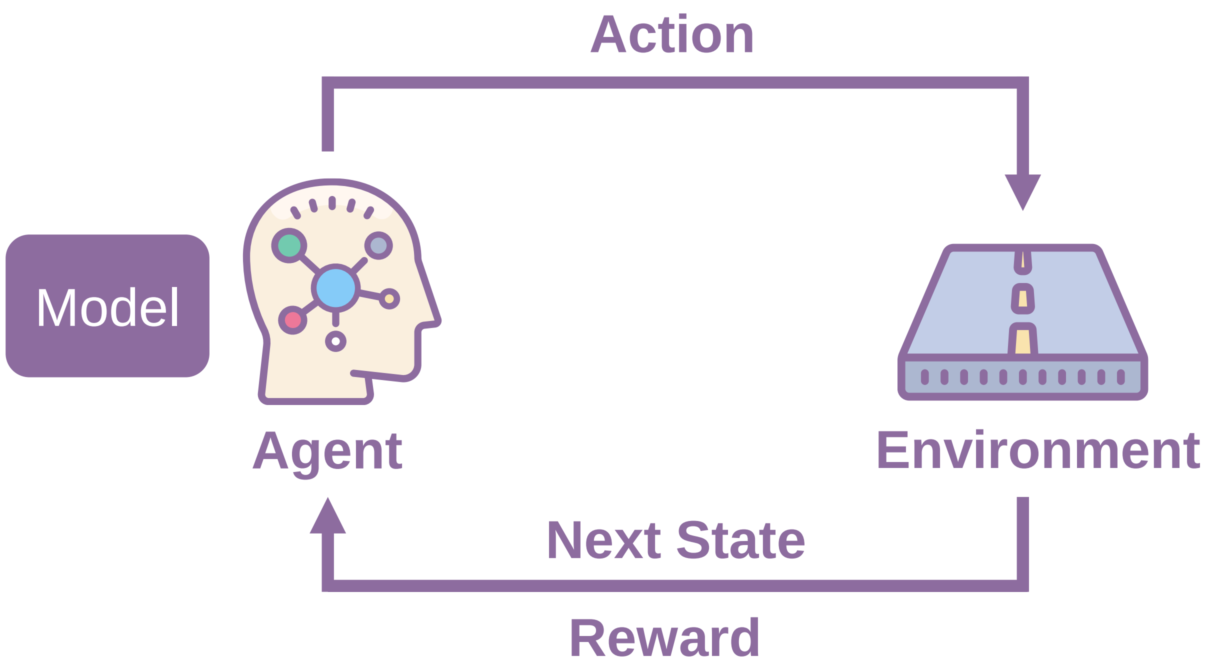

What is Reinforcement Learning?

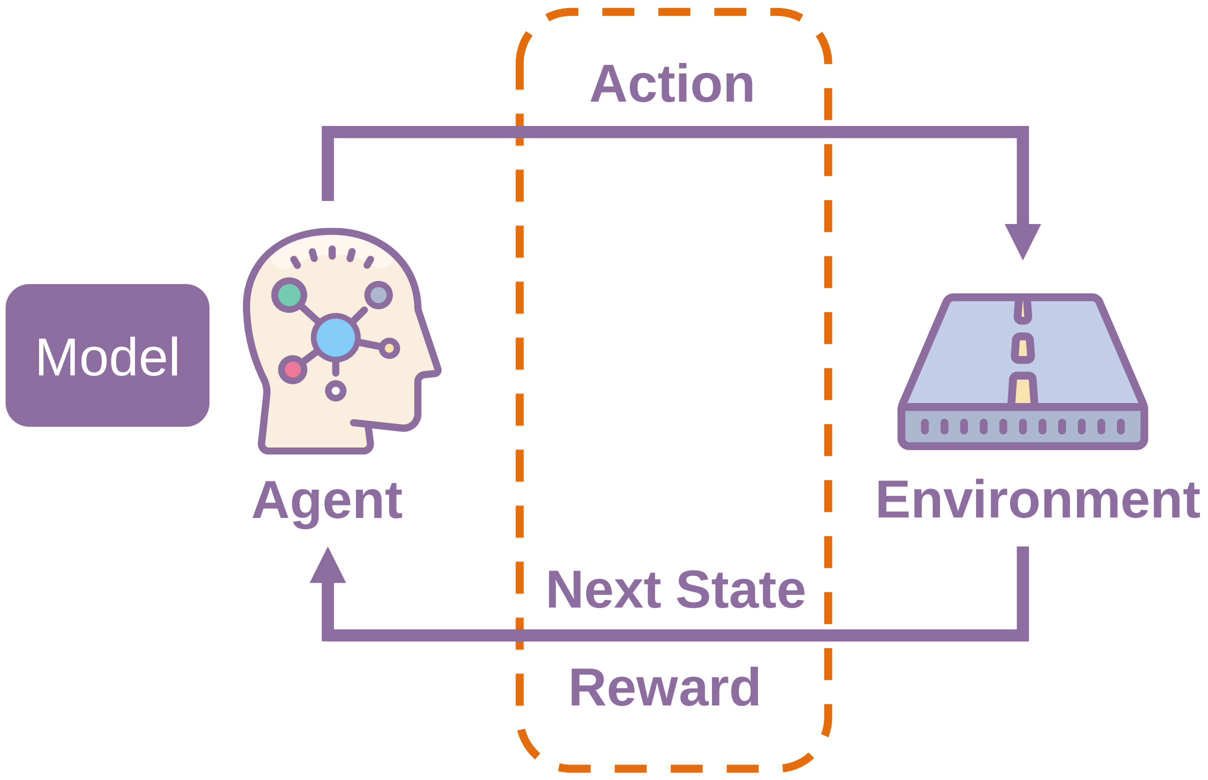

What is Reinforcement Learning?

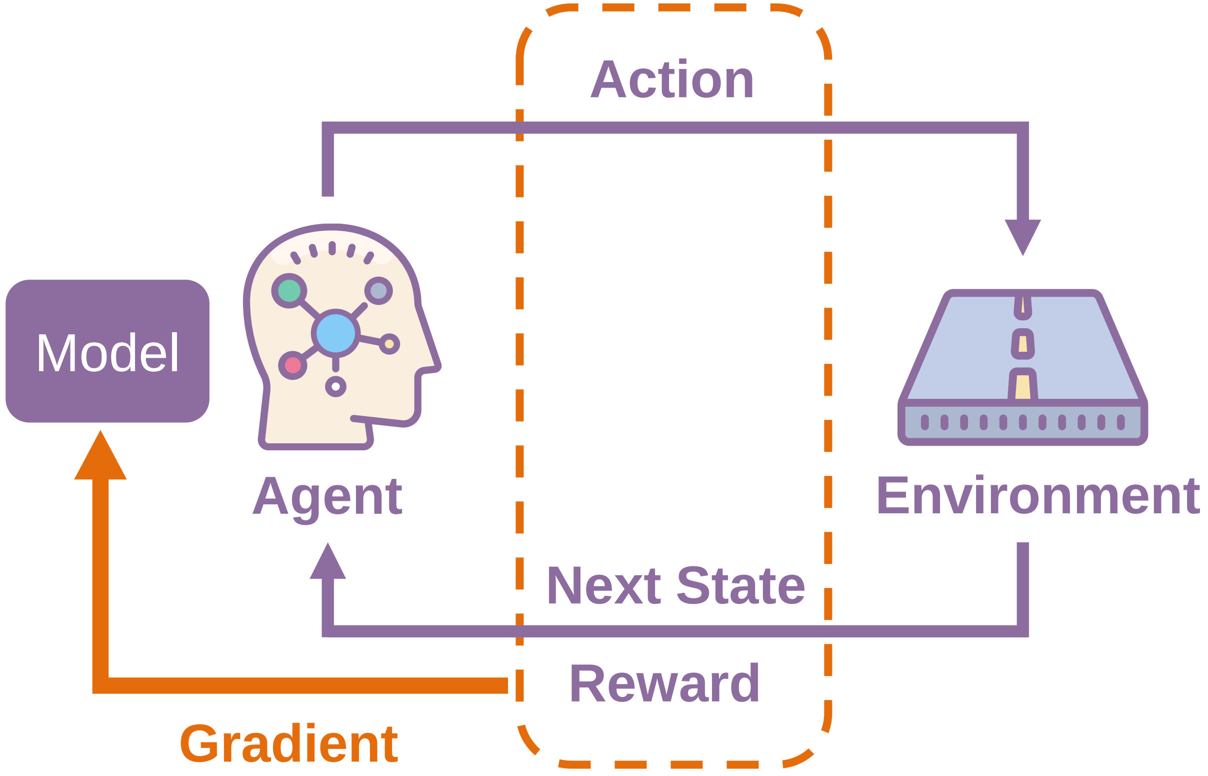

What is Reinforcement Learning?

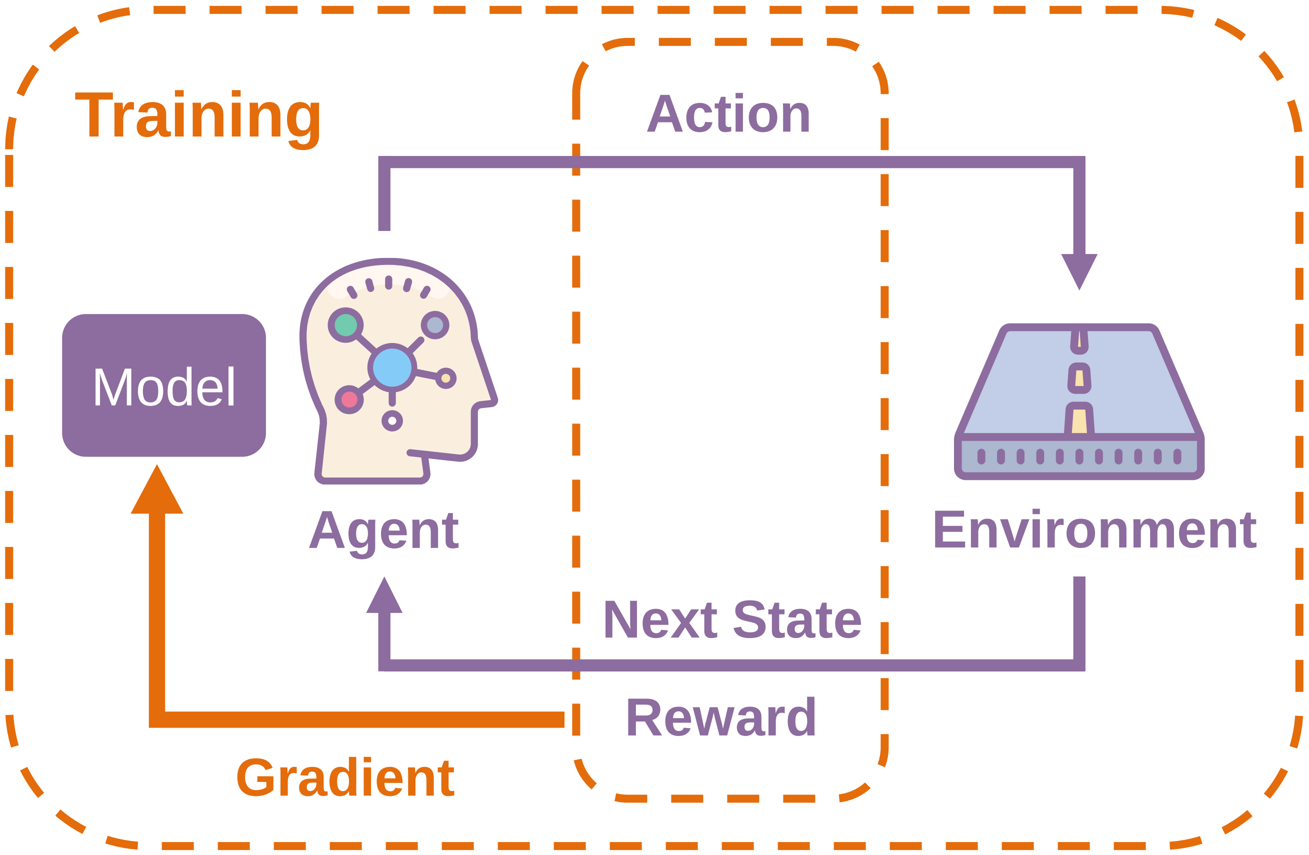

What is Reinforcement Learning?

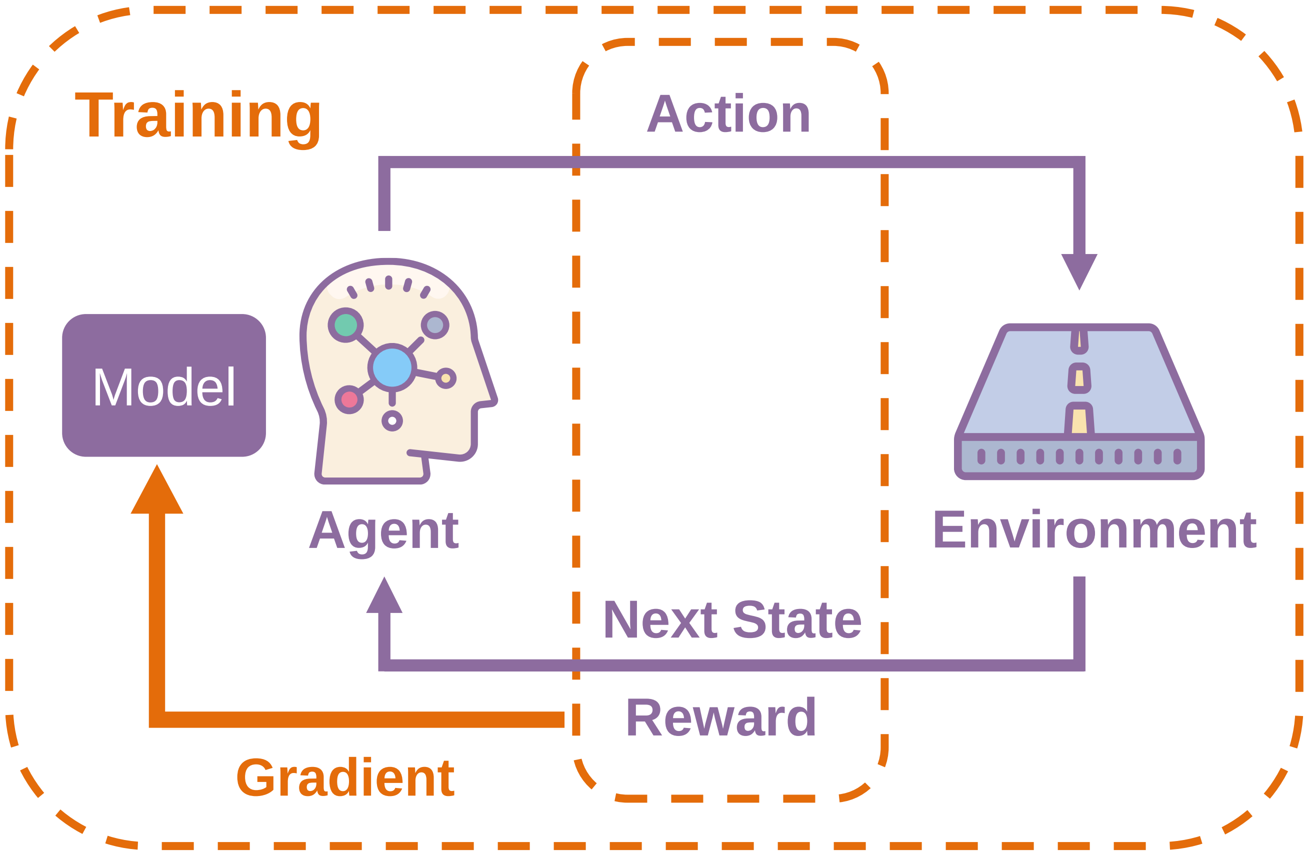

Train a Typical RL Agent on a Single GPU = **8** Days*

*Mnih, ICML'16

What is Reinforcement Learning?

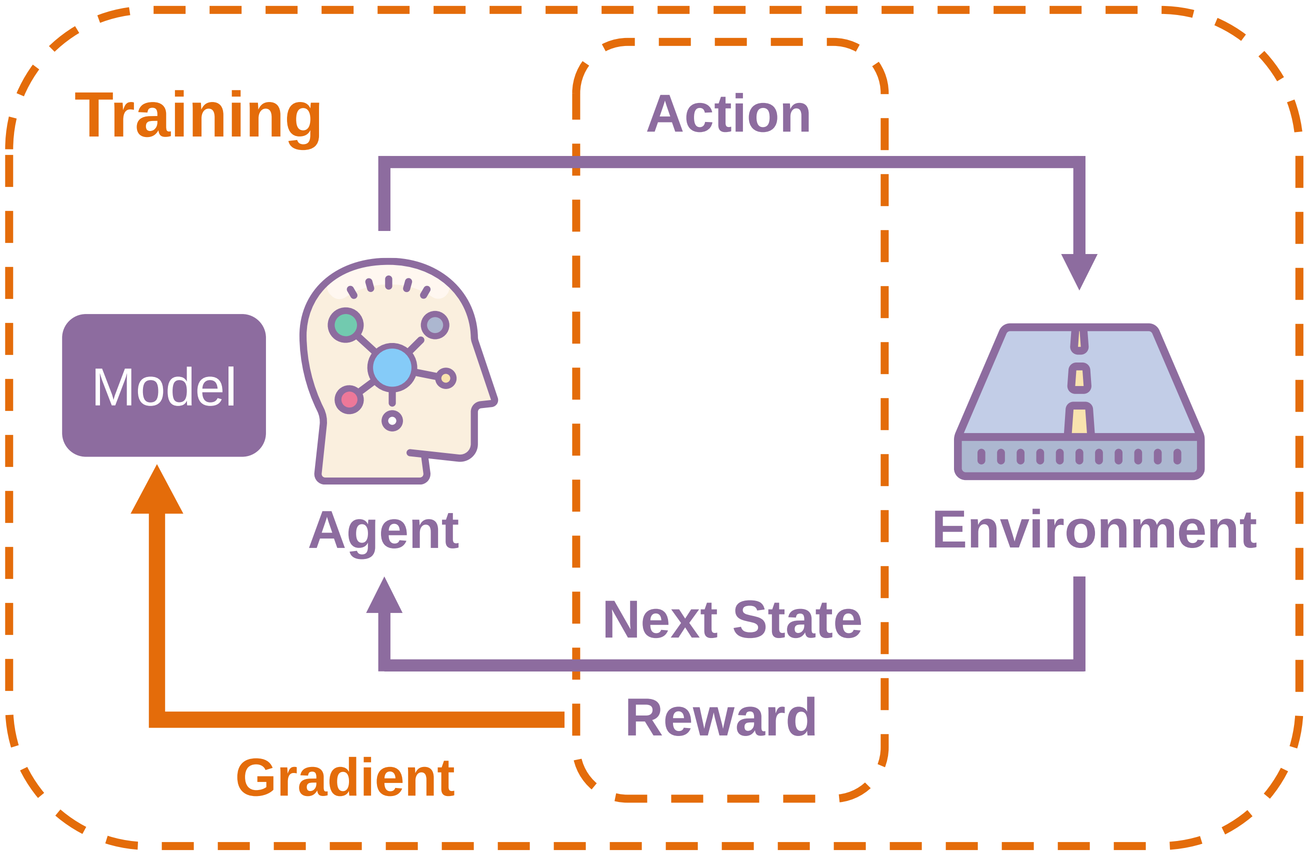

Train a Typical RL Agent on a Single GPU = **8** Days*

*Mnih, ICML'16

RL Requires Distributed Training for Improved Performance

Centralized Distributed RL Training: Parameter-Server Based

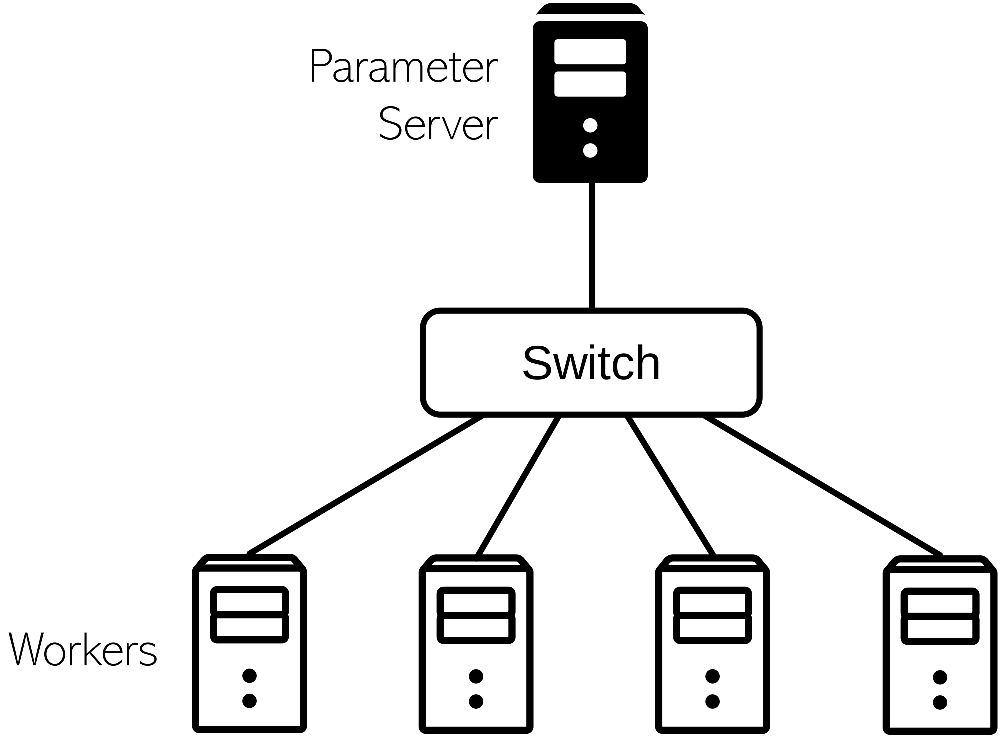

Centralized Distributed RL Training: Parameter-Server Based

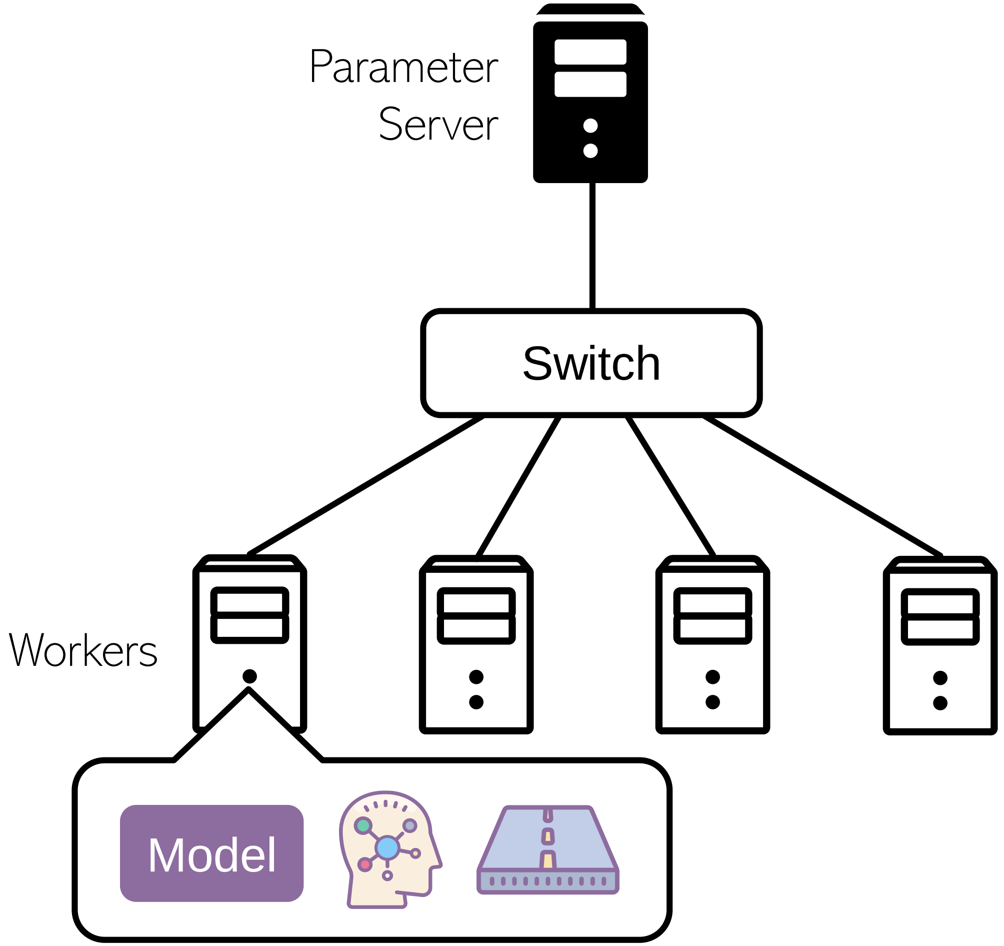

Centralized Distributed RL Training: Parameter-Server Based

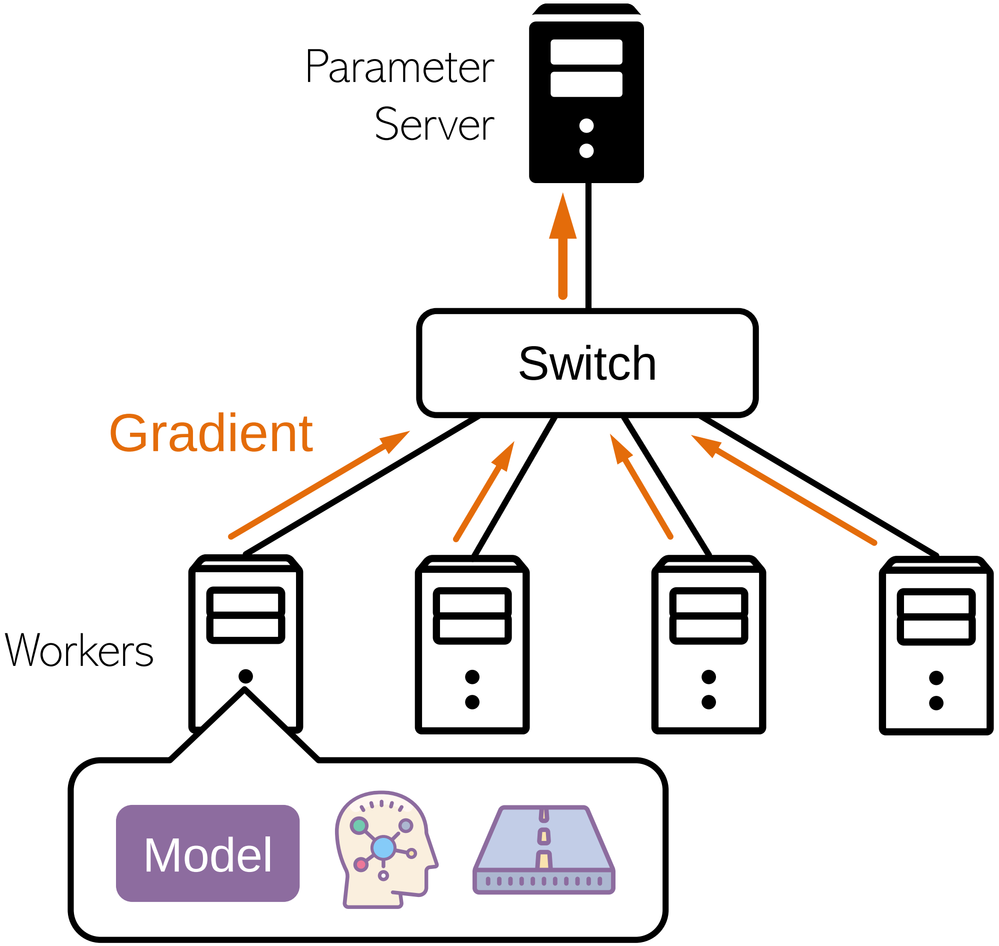

Centralized Distributed RL Training: Parameter-Server Based

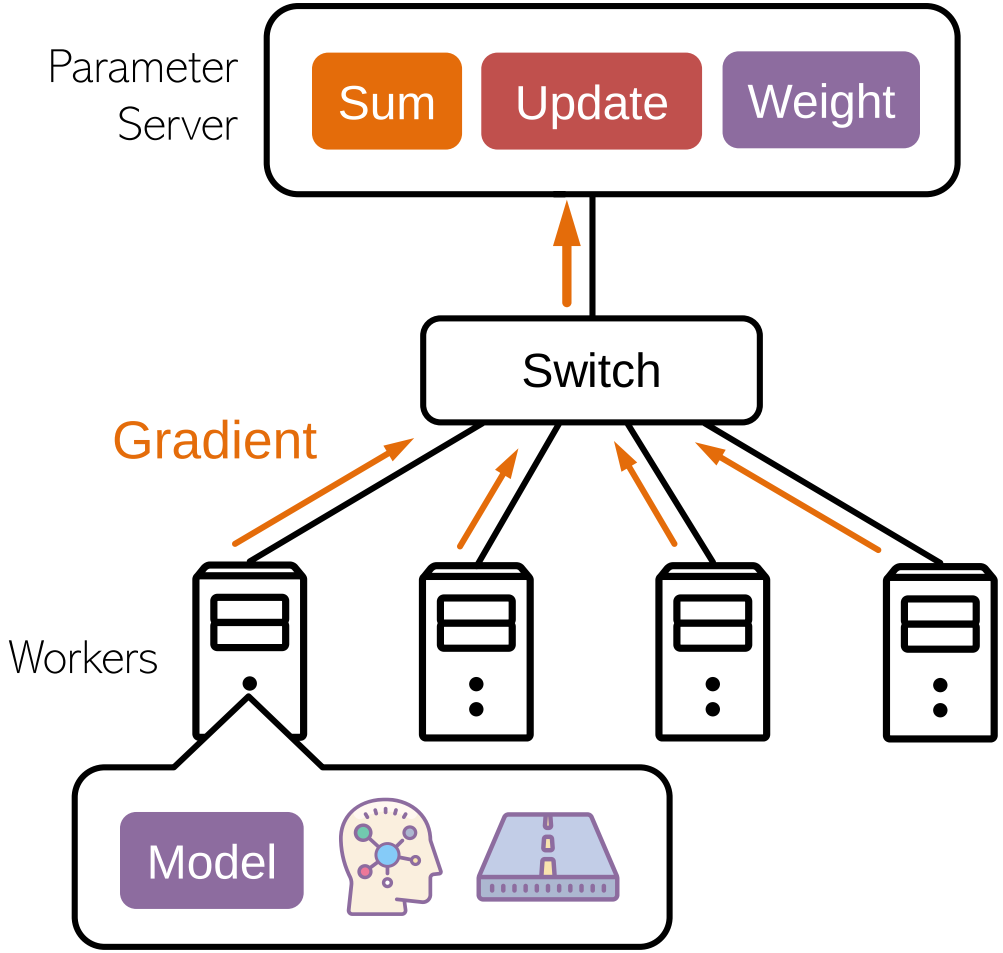

Centralized Distributed RL Training: Parameter-Server Based

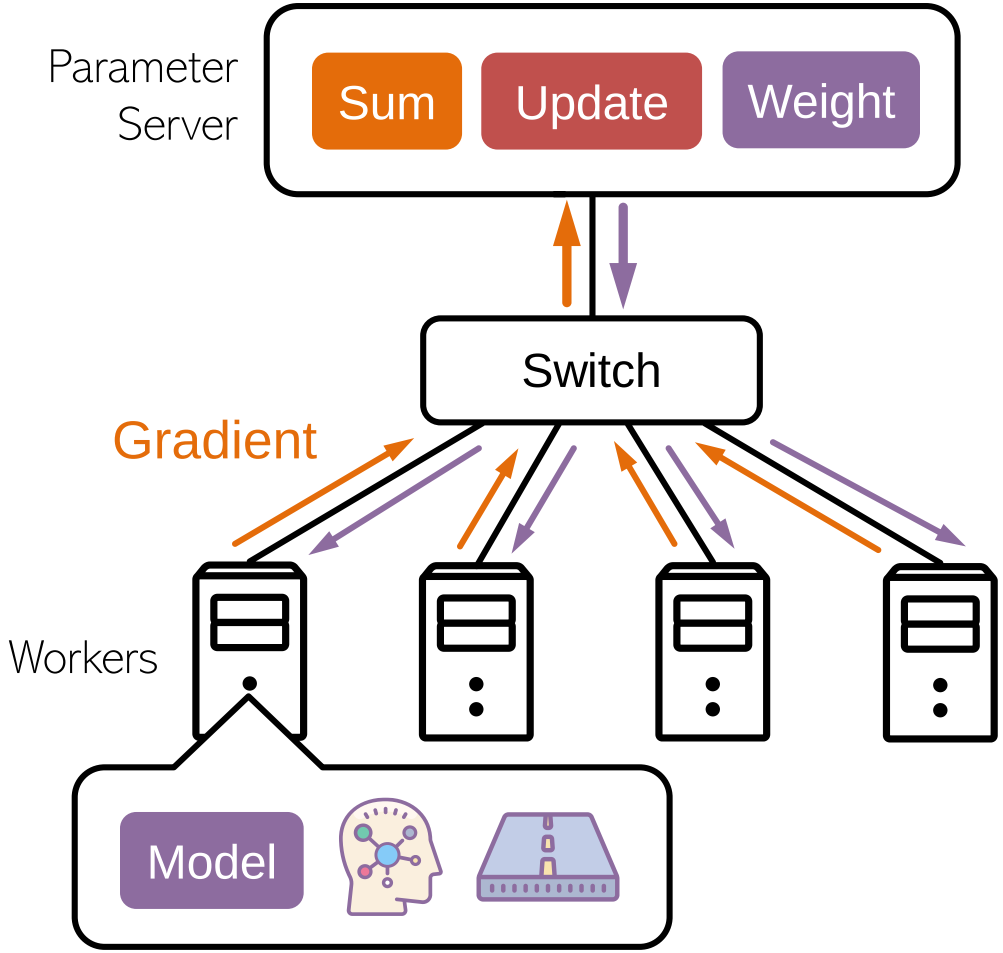

Centralized Distributed RL Training: Parameter-Server Based

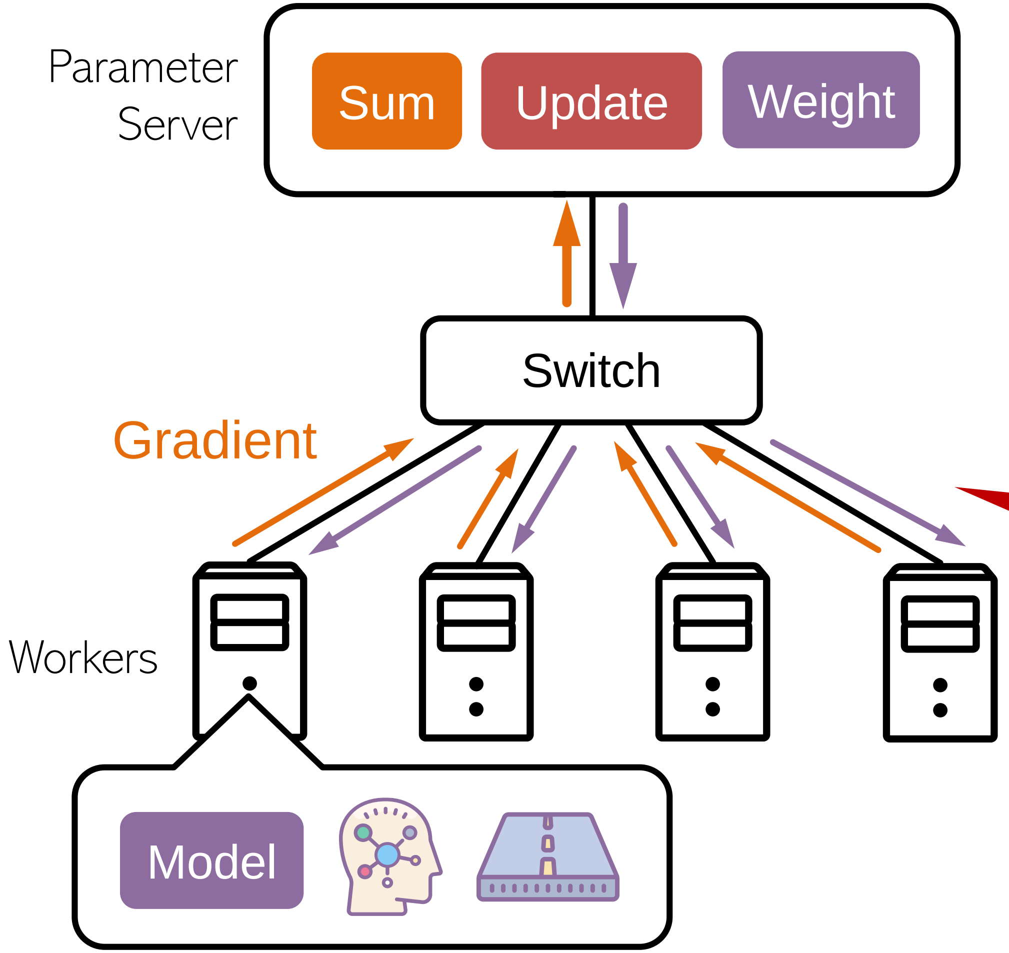

Multiple Network Hops

Centralized Distributed RL Training: Parameter-Server Based

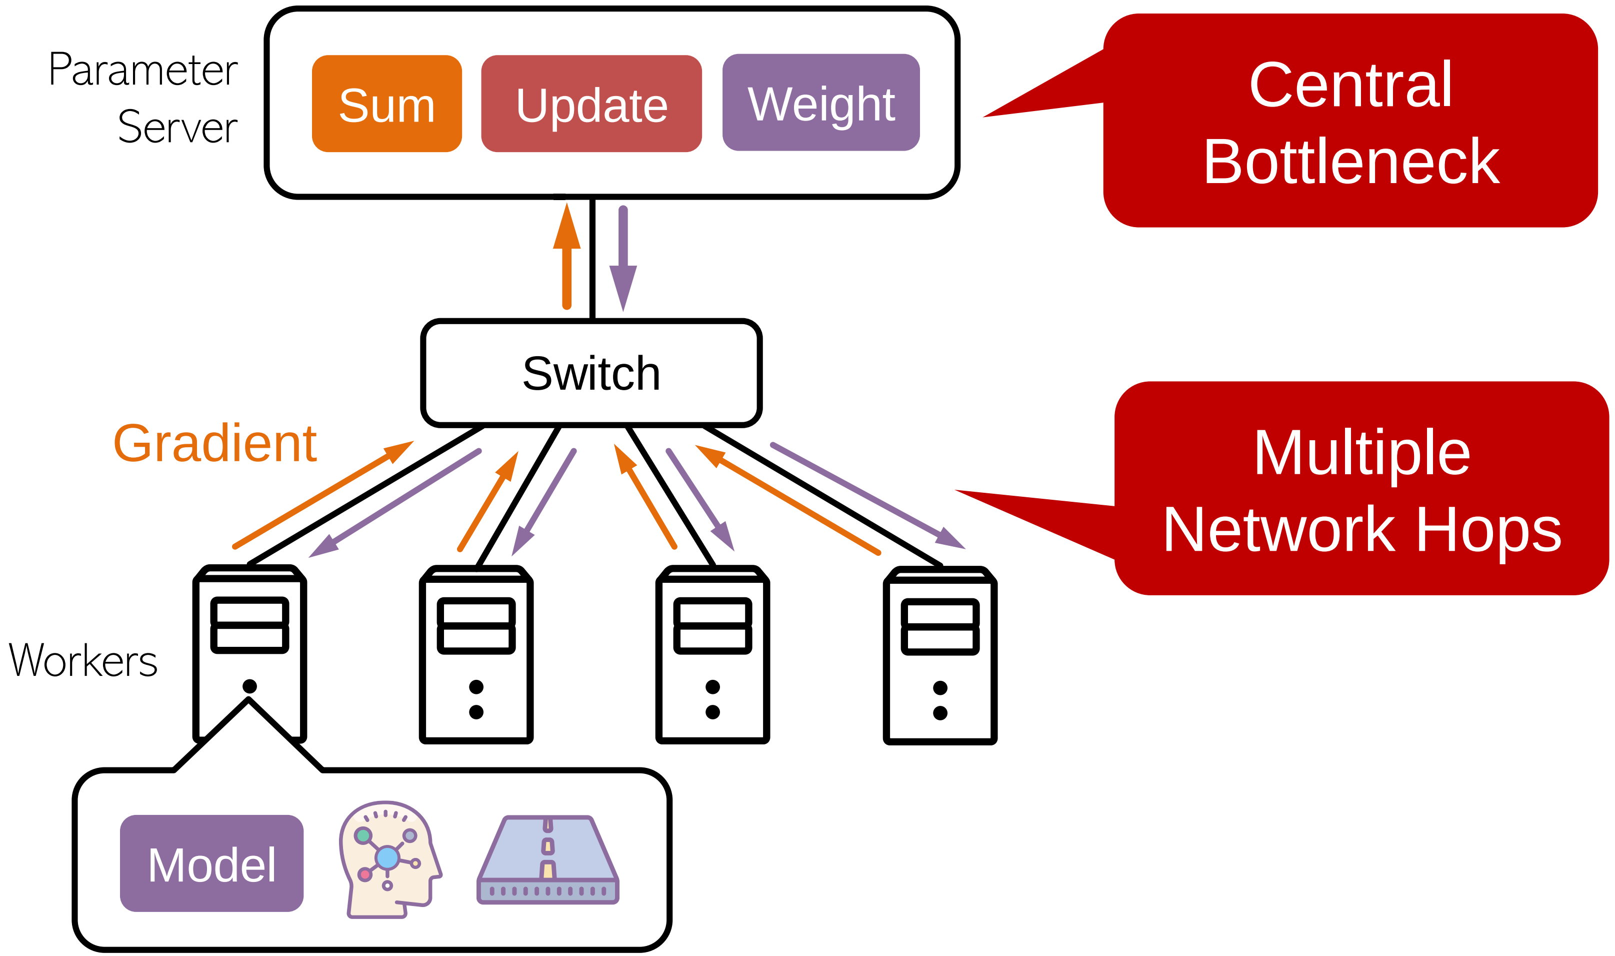

Decentralized Distributed RL Training: AllReduce Based

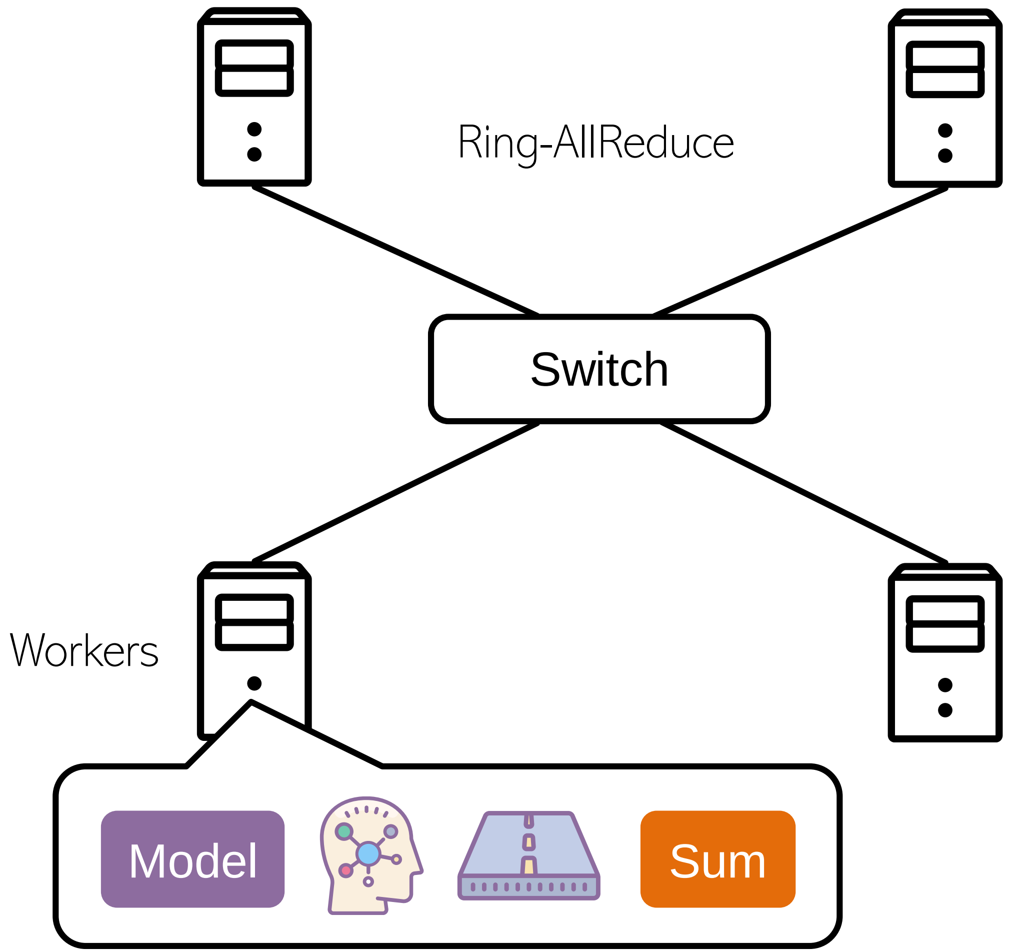

Decentralized Distributed RL Training: AllReduce Based

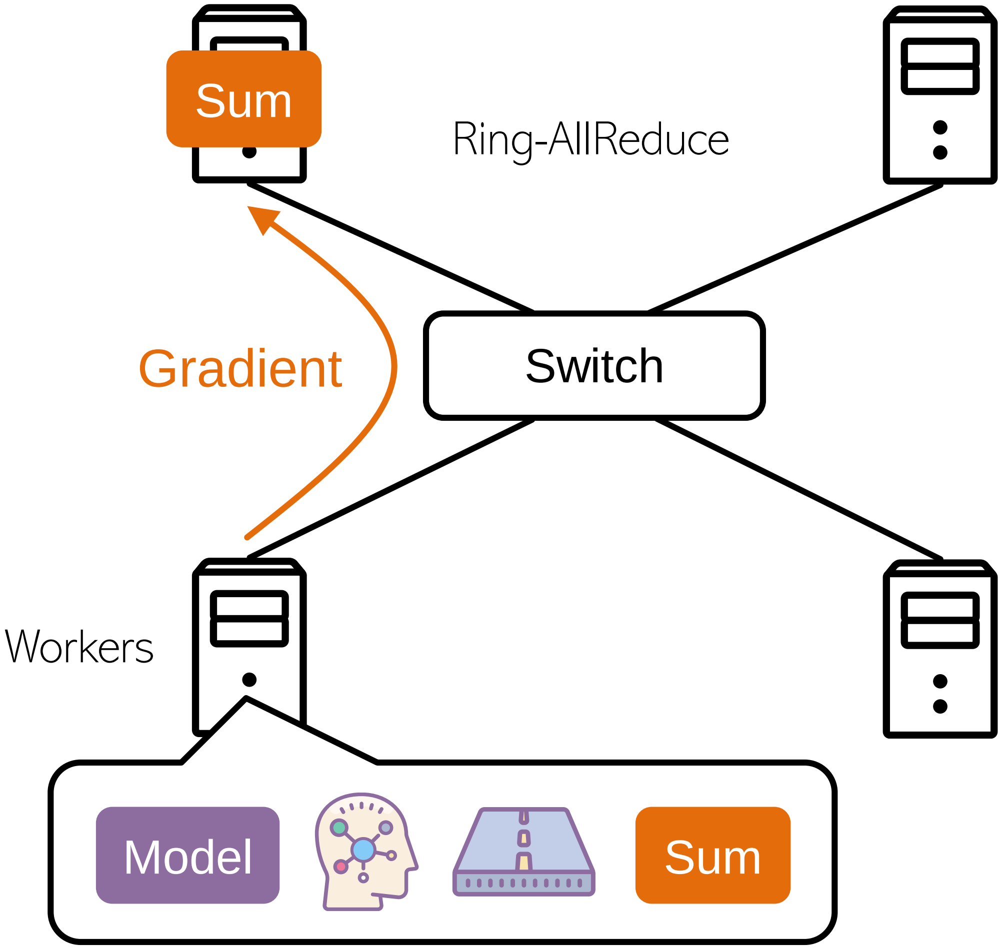

Decentralized Distributed RL Training: AllReduce Based

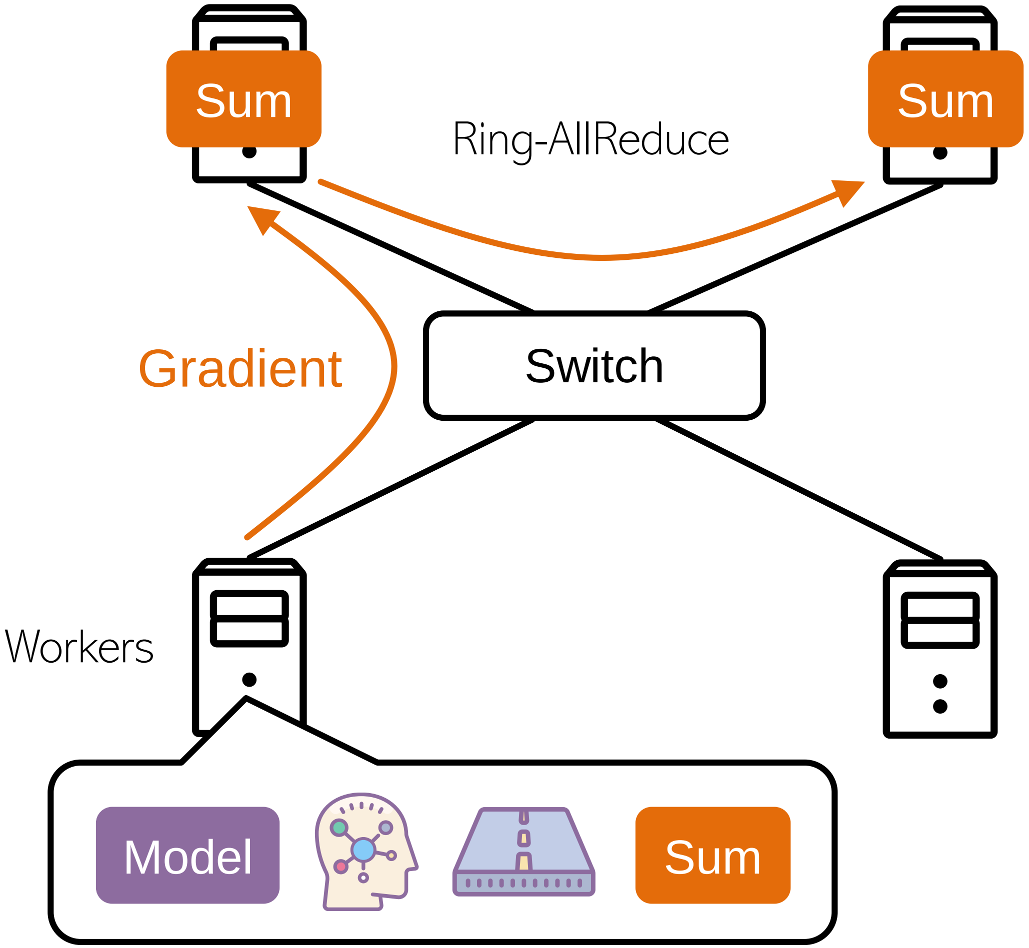

Decentralized Distributed RL Training: AllReduce Based

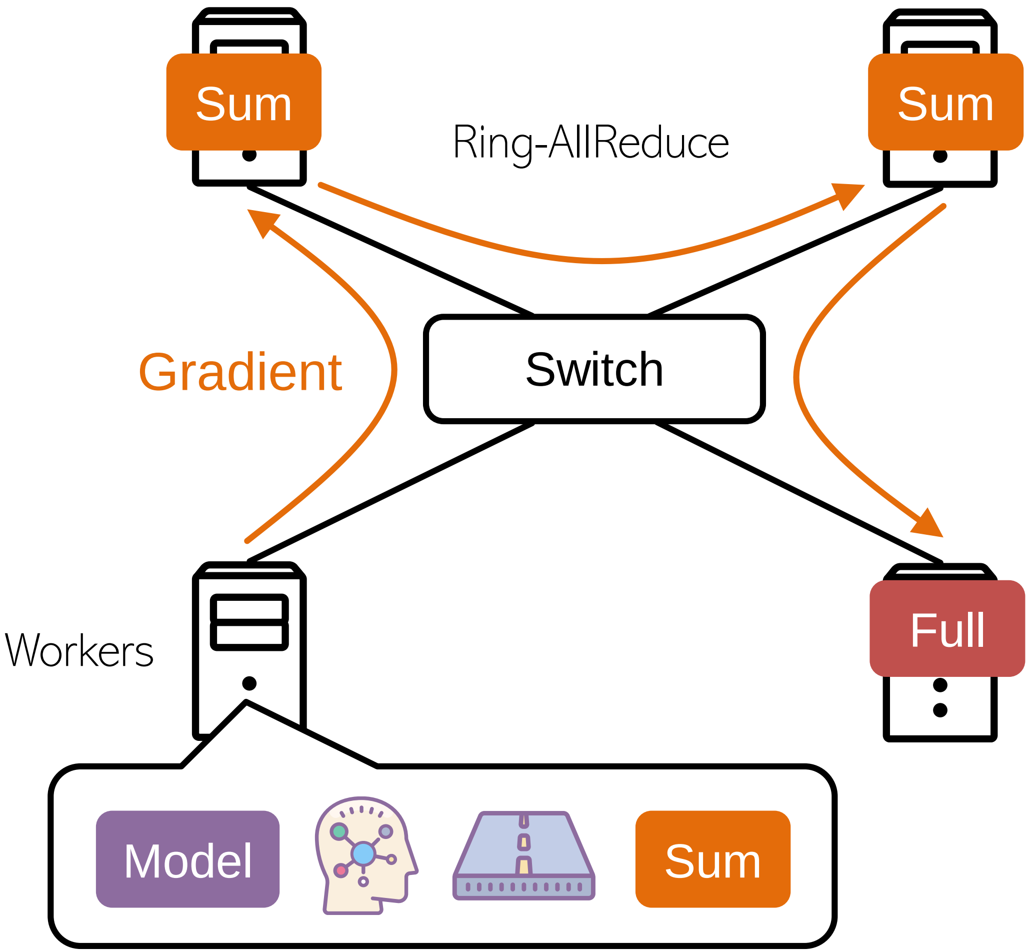

Decentralized Distributed RL Training: AllReduce Based

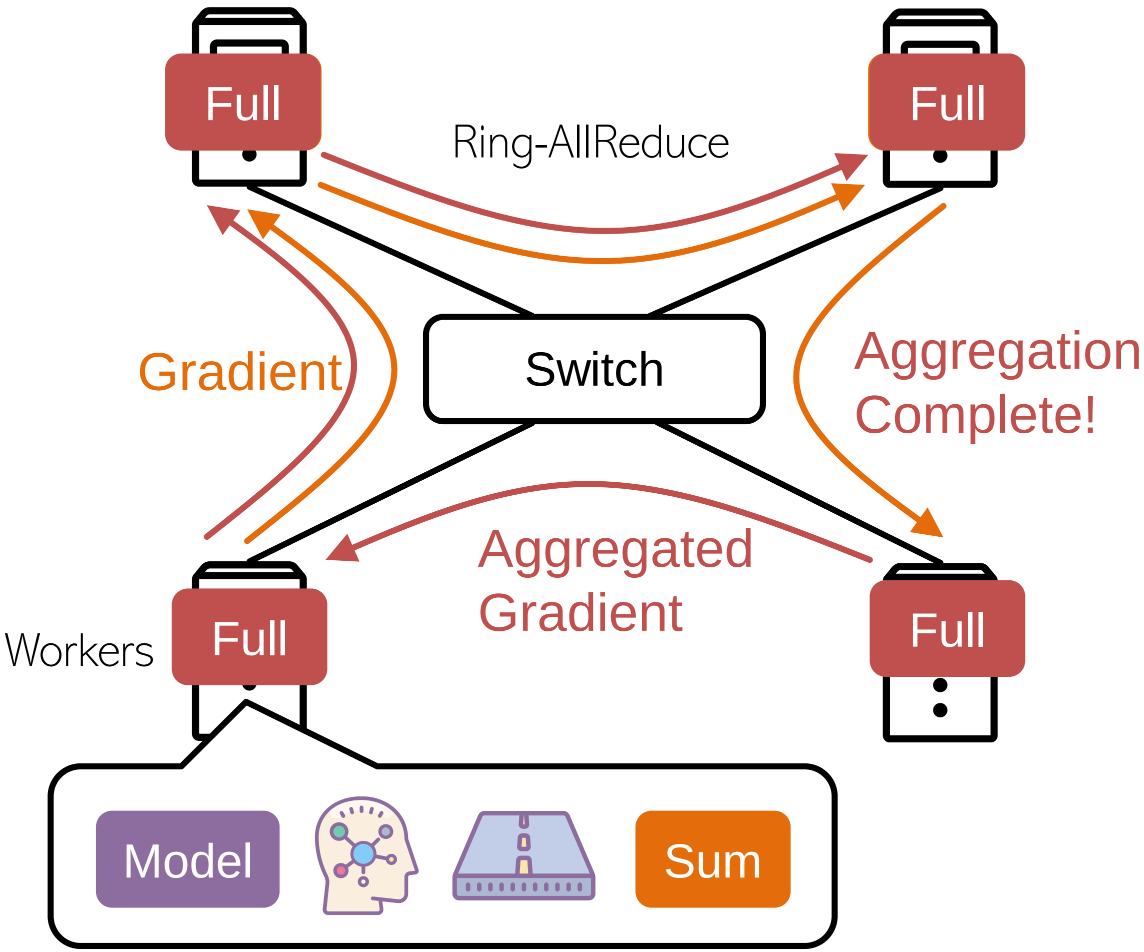

Decentralized Distributed RL Training: AllReduce Based

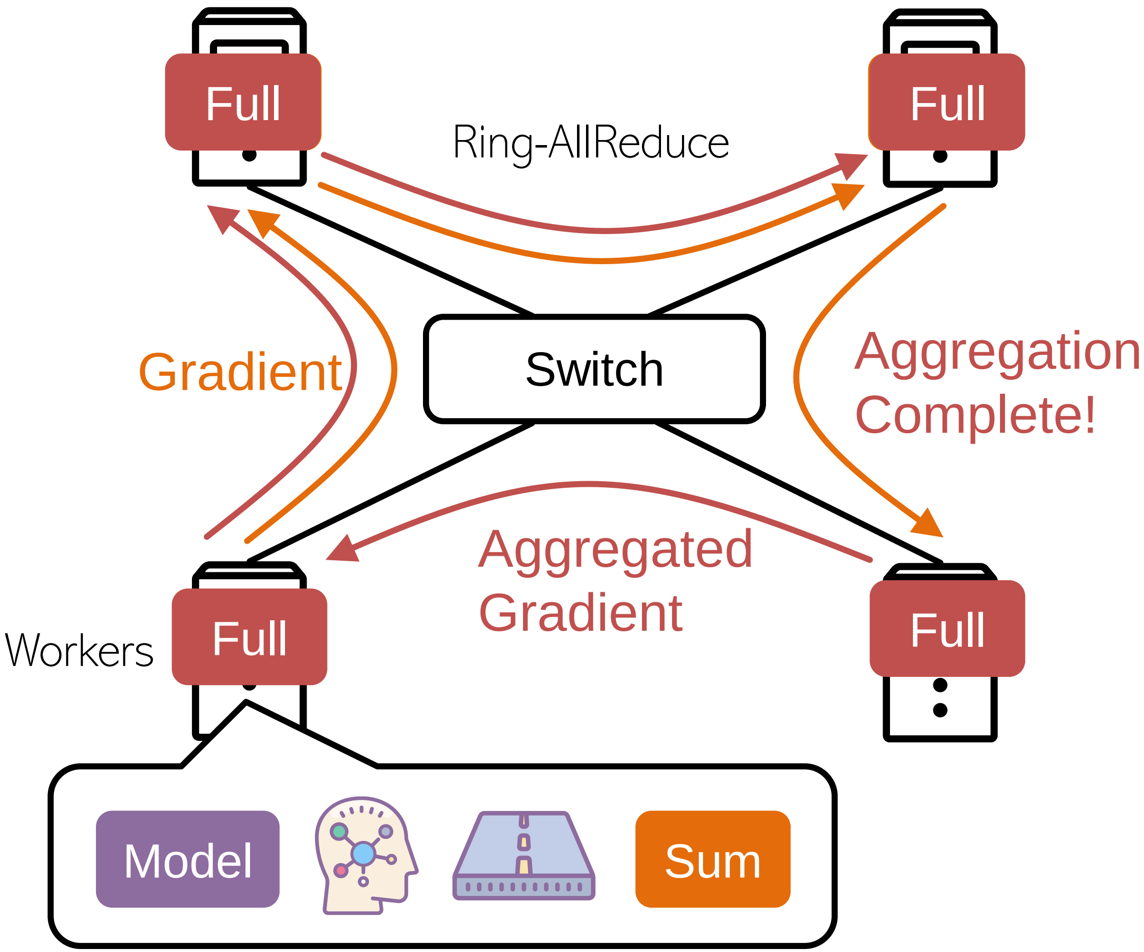

Excessive Network Hops

Network Communication is the Bottleneck in Distributed RL Training

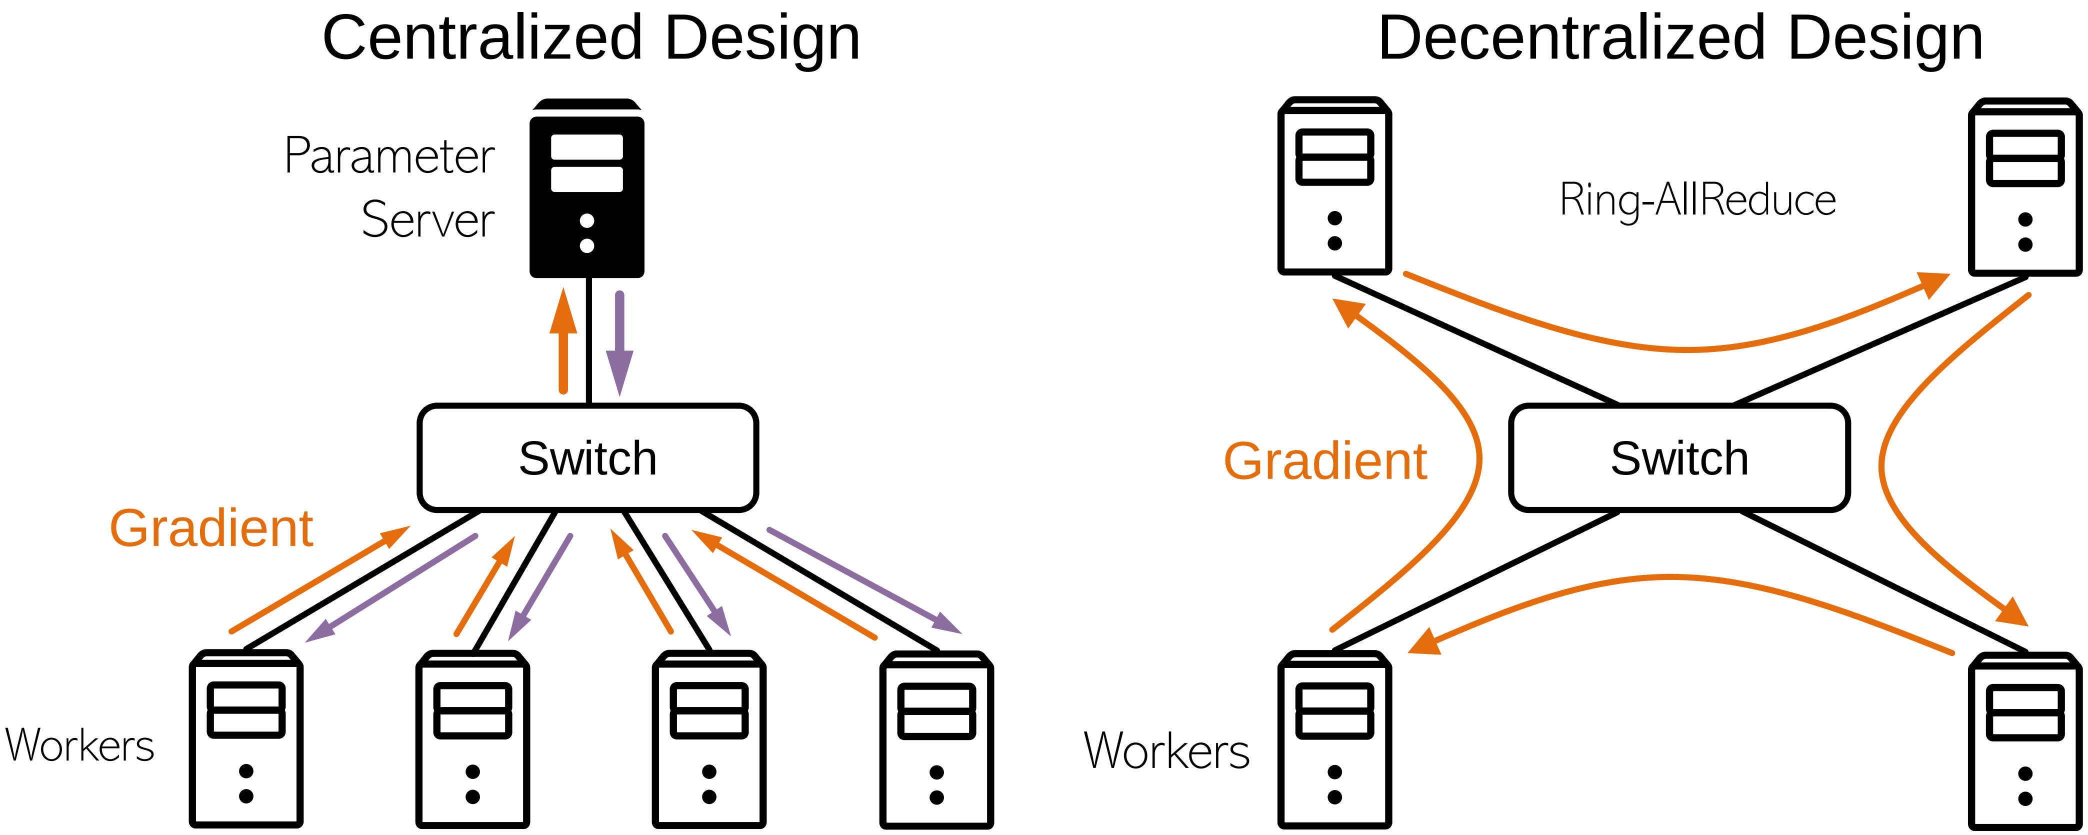

Network Communication is the Bottleneck in Distributed RL Training

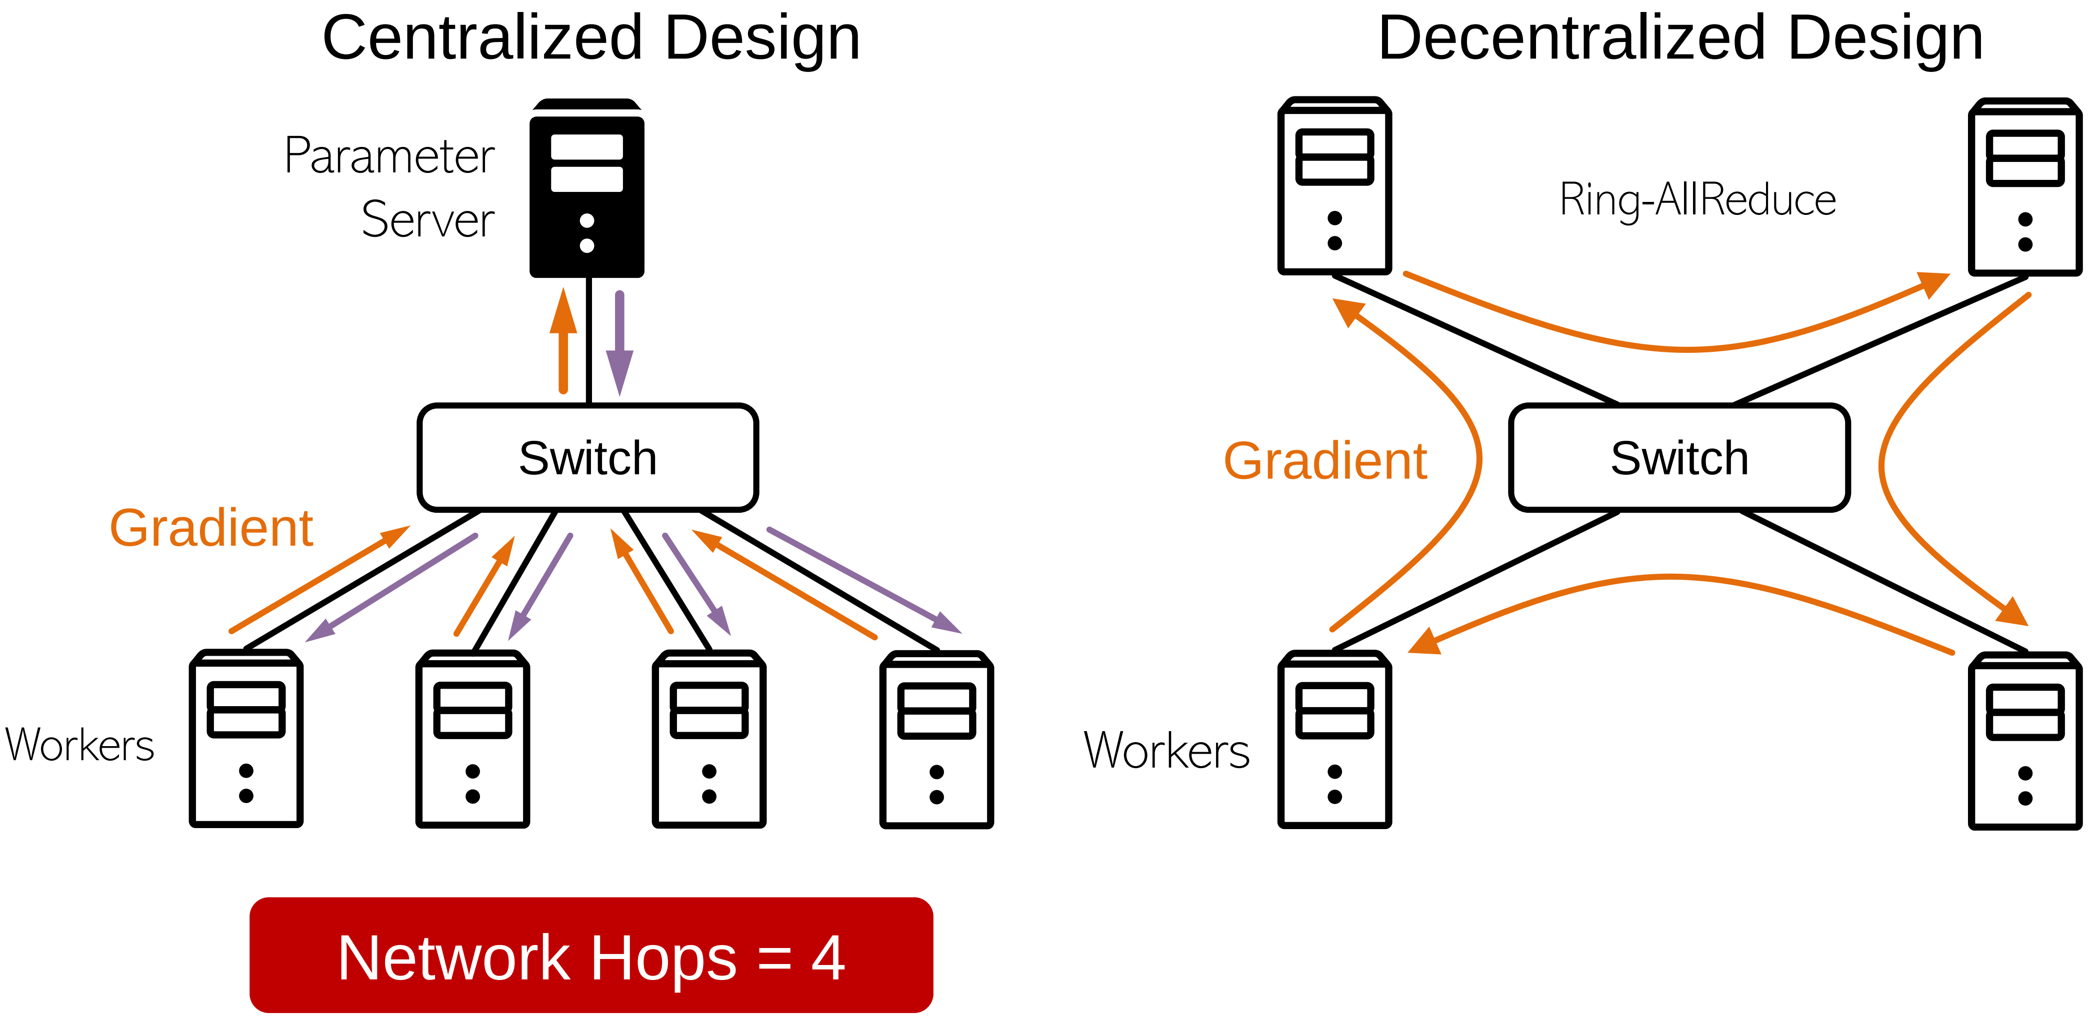

Network Communication is the Bottleneck in Distributed RL Training

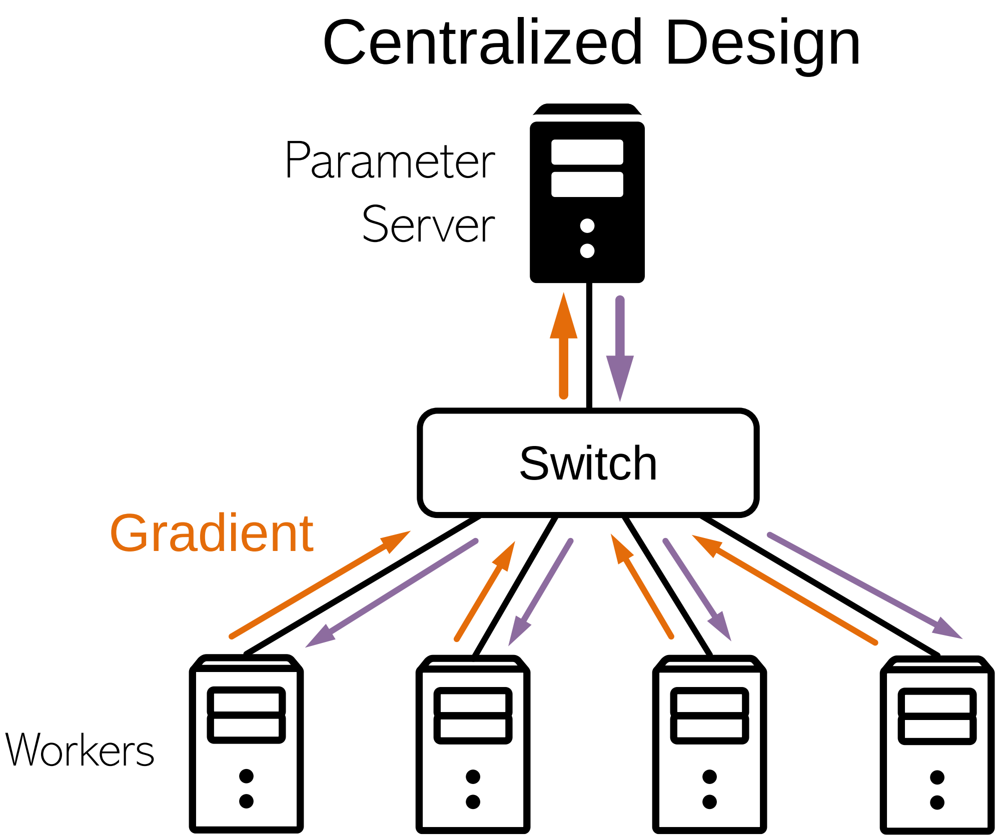

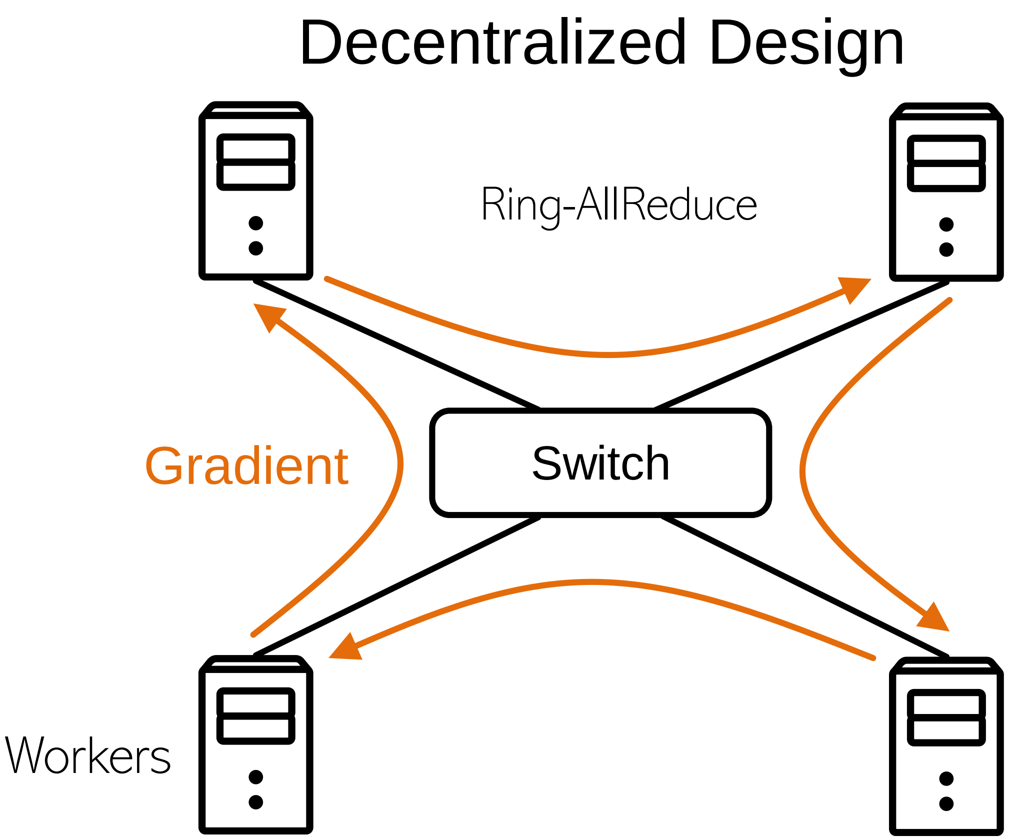

Network Hops = 4

Network Hops = 4N - 4

The Unique Characteristic of Distributed RL Training: Latency Critical

RL Benchmark	DQN-Atari	A2C-Atari	PPO-MuJoCo	DDPG-MuJoCo
Gradient Size	**6 MB**	**3 MB**	**40 KB**	**158 KB**
Training Iterations	**200 M**	**2 M**	**0.2 M**	**3 M**

The Unique Characteristic of Distributed RL Training: Latency Critical

RL Benchmark	DQN-Atari	A2C-Atari	PPO-MuJoCo	DDPG-MuJoCo
Gradient Size	**6 MB**	**3 MB**	**40 KB**	**158 KB**
Training Iterations	**200 M**	**2 M**	**0.2 M**	**3 M**

DNN Benchmark	AlexNet-ImageNet	ResNet50-ImageNet	VGG16-ImageNet	MLP-MNIST
Gradient Size	**250 MB**	**100 MB**	**525 MB**	**4 MB**
Training Iterations	**320 K**	**600 K**	**370 K**	**10 K**

The Unique Characteristic of Distributed RL Training: Latency Critical

RL Benchmark	DQN-Atari	A2C-Atari	PPO-MuJoCo	DDPG-MuJoCo
Gradient Size	**6 MB**	**3 MB**	**40 KB**	**158 KB**
Training Iterations	**200 M**	**2 M**	**0.2 M**	**3 M**

DNN Benchmark	AlexNet-ImageNet	ResNet50-ImageNet	VGG16-ImageNet	MLP-MNIST
Gradient Size	**250 MB**	**100 MB**	**525 MB**	**4 MB**
Training Iterations	**320 K**	**600 K**	**370 K**	**10 K**

88x Smaller Gradient Size

158x More Iterations

The Unique Characteristic of Distributed RL Training: Latency Critical

RL Benchmark	DQN-Atari	A2C-Atari	PPO-MuJoCo	DDPG-MuJoCo
Gradient Size	**6 MB**	**3 MB**	**40 KB**	**158 KB**
Training Iterations	**200 M**	**2 M**	**0.2 M**	**3 M**

DNN Benchmark	AlexNet-ImageNet	ResNet50-ImageNet	VGG16-ImageNet	MLP-MNIST
Gradient Size	**250 MB**	**100 MB**	**525 MB**	**4 MB**
Training Iterations	**320 K**	**600 K**	**370 K**	**10 K**

88x Smaller Gradient Size

158x More Iterations

Distributed RL Training is Latency Critical

Quantifying the Network Overhead in Distributed RL Training

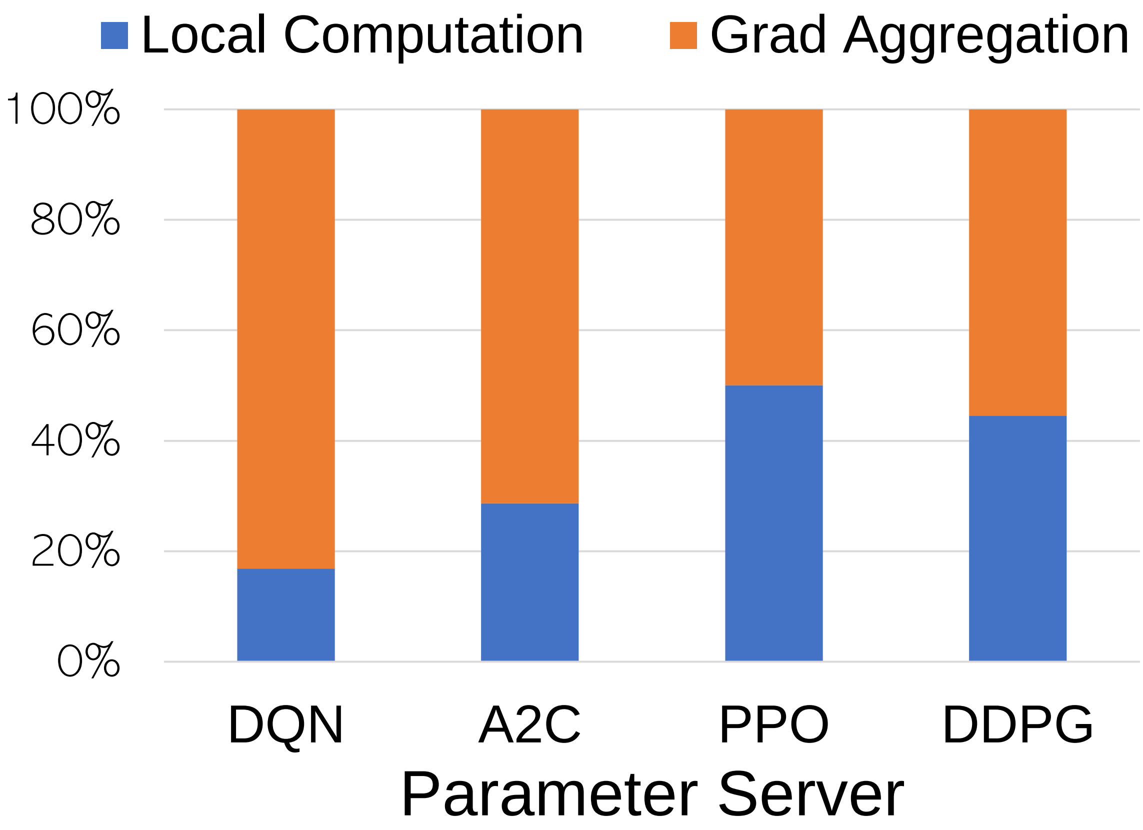

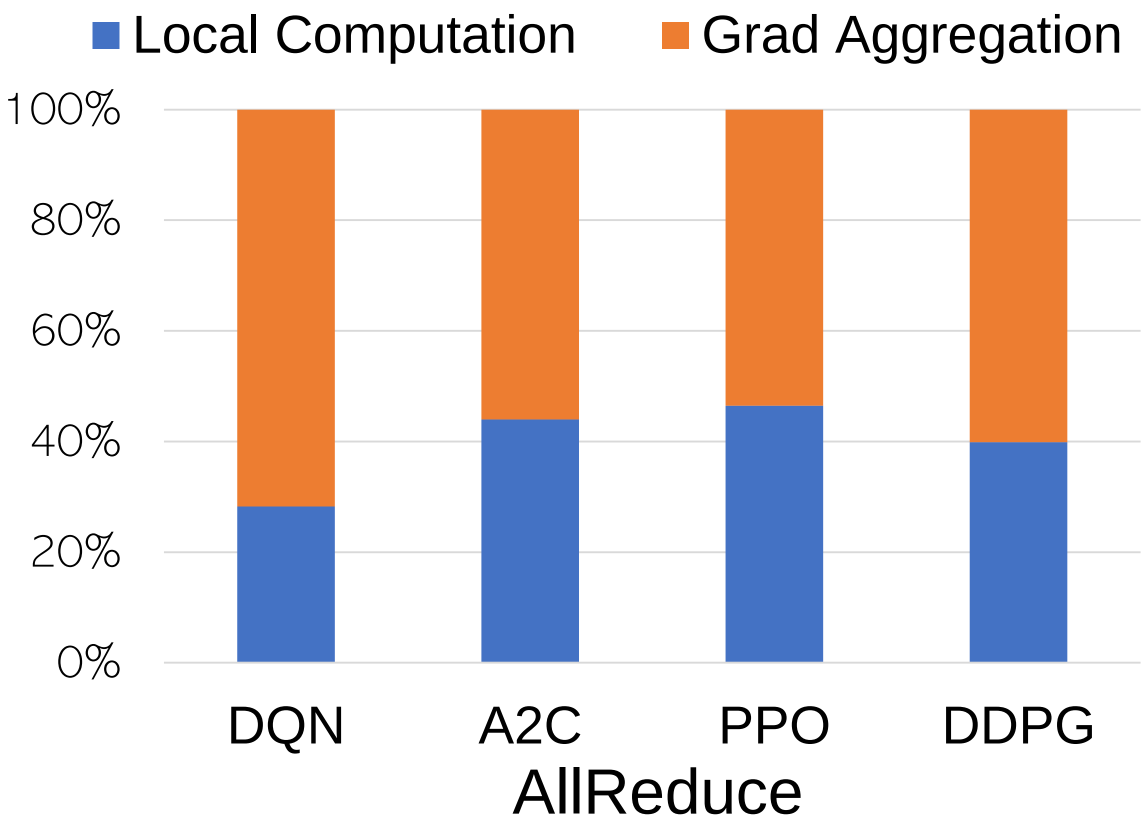

Quantifying the Network Overhead in Distributed RL Training

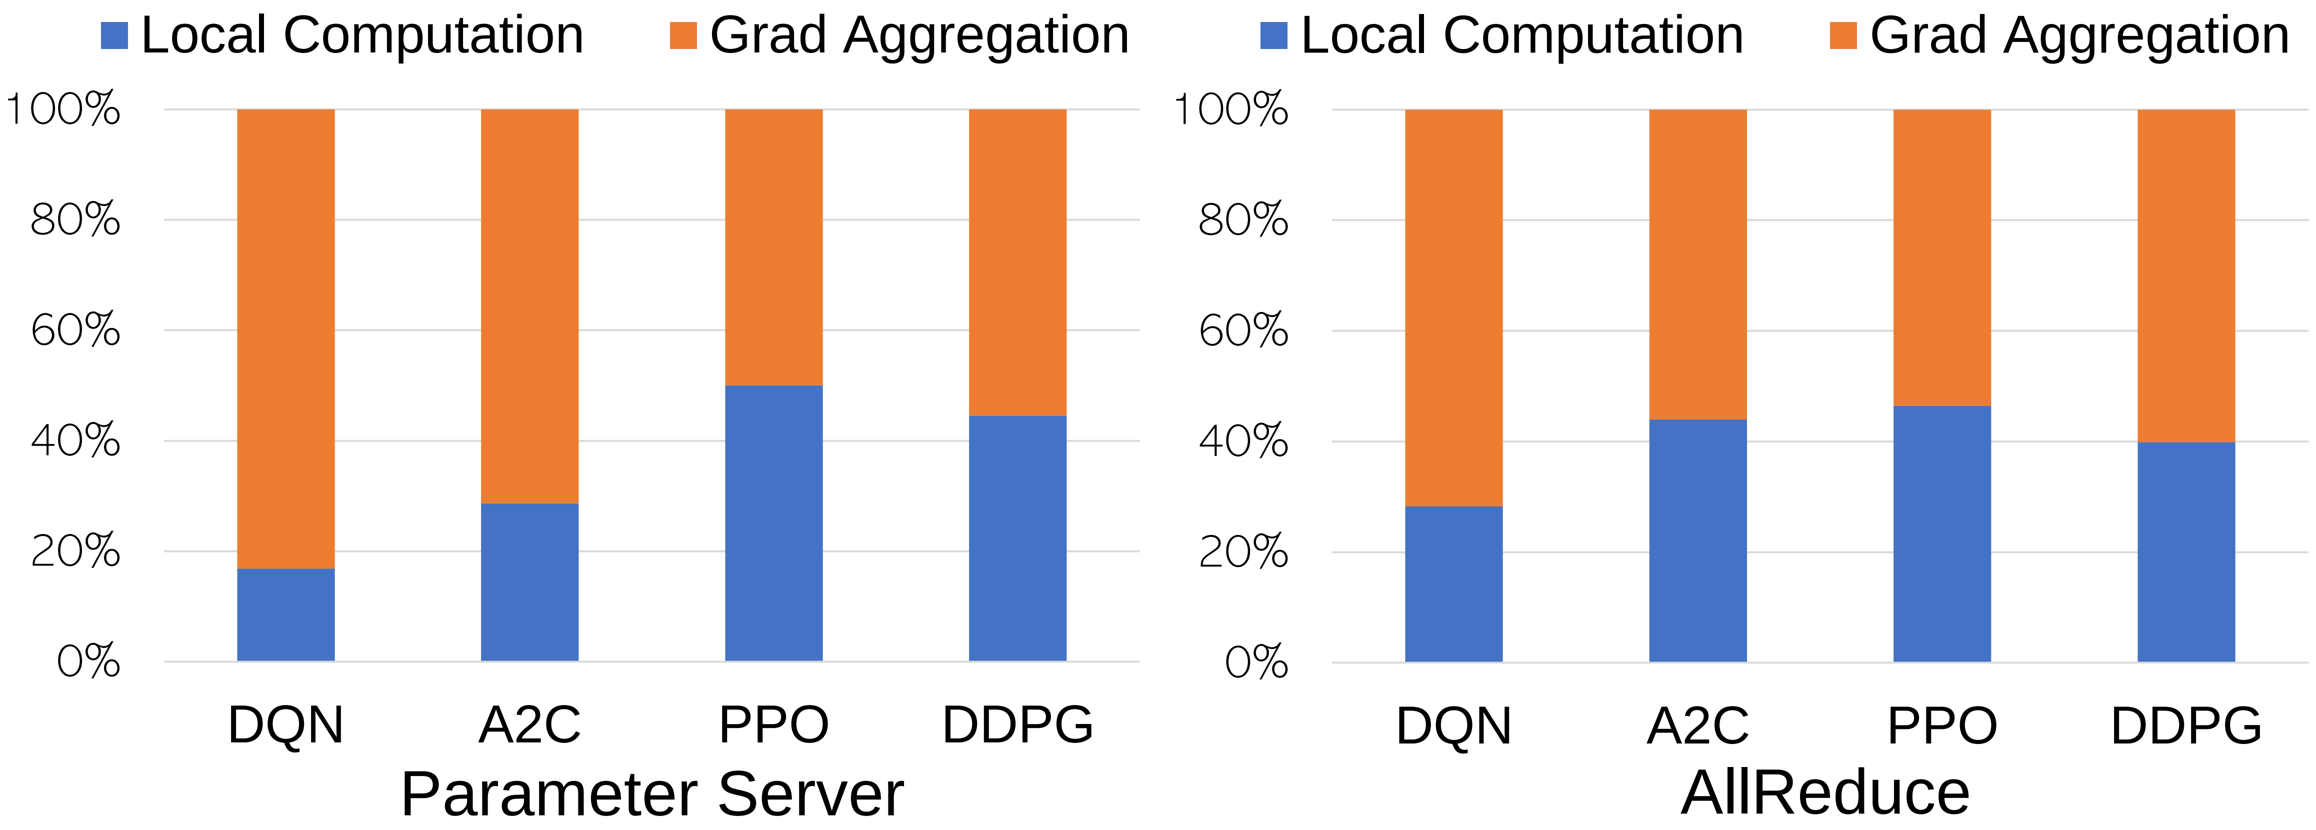

Gradient Aggregation over the Network Dominates the Training Time **(50~83%)**

Quantifying the Network Overhead in Distributed RL Training

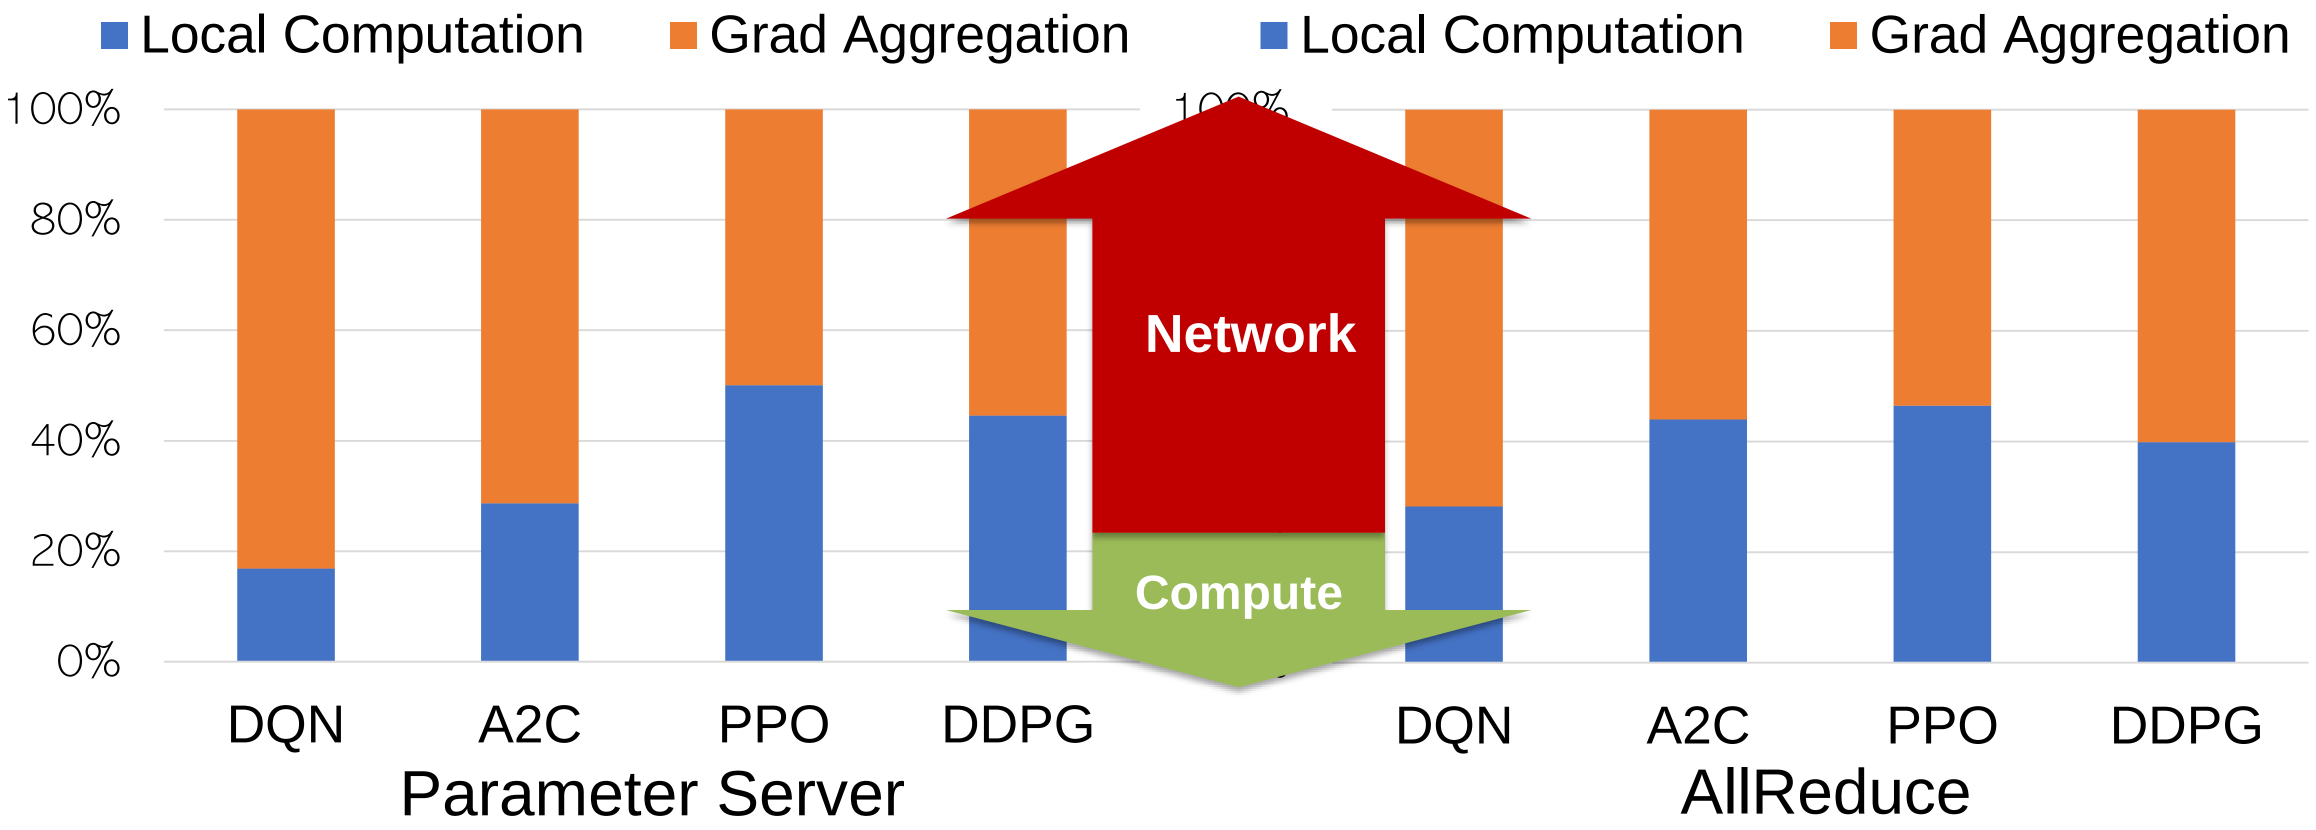

Gradient Aggregation over the Network Dominates the Training Time **(50~83%)**

In-Switch Acceleration: A New Distributed Computing Paradigm

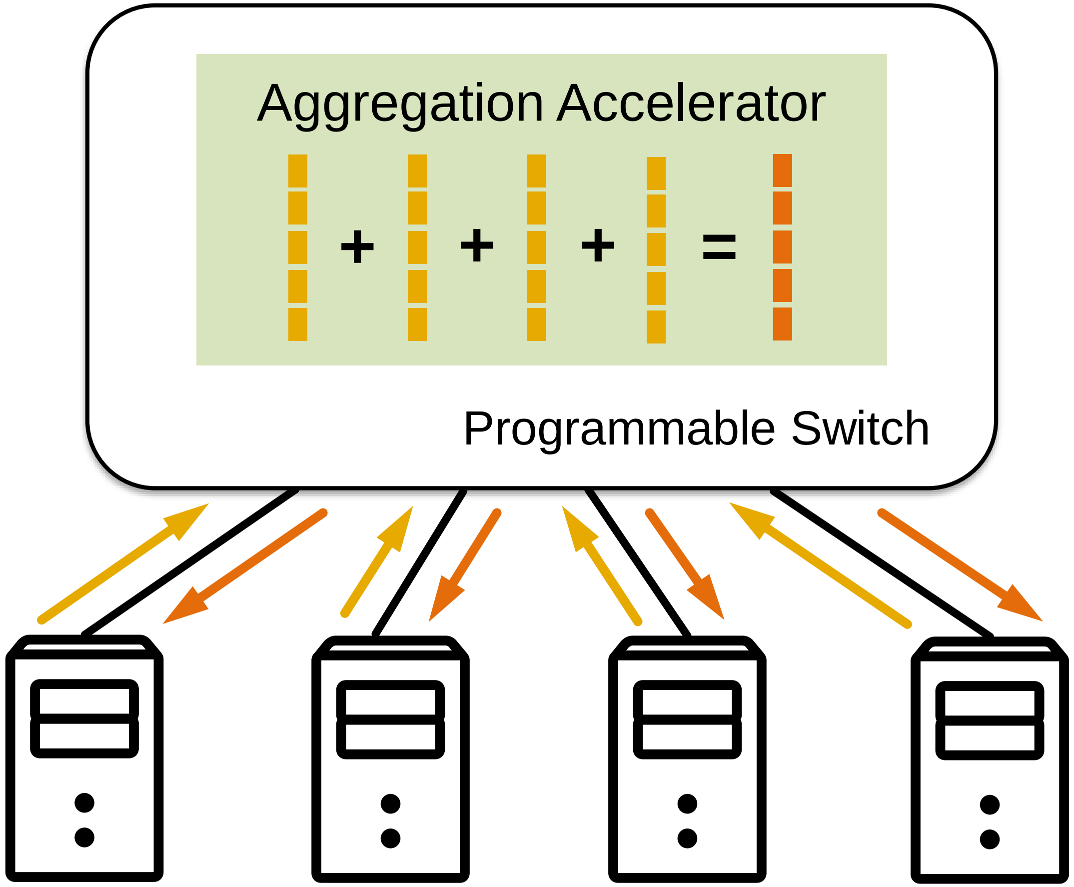

In-Switch Acceleration: A New Distributed Computing Paradigm

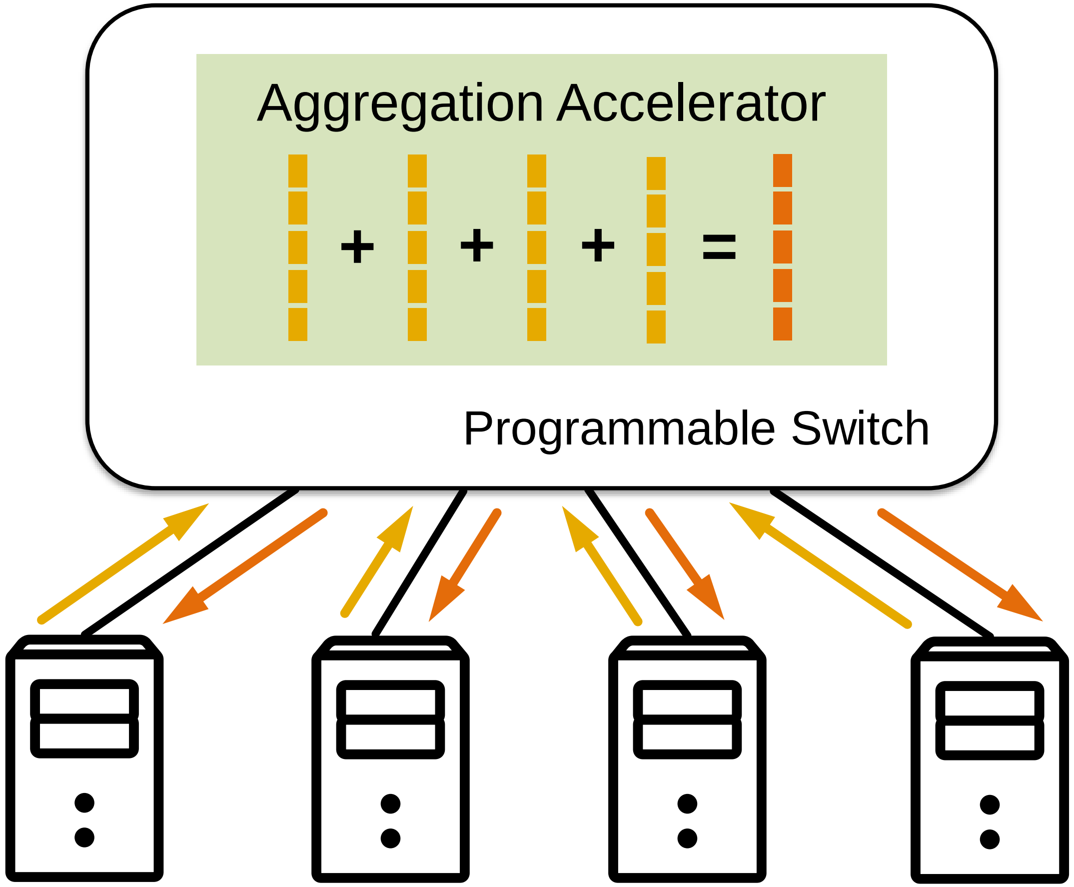

Performance

Reduce End-to-End Network Latency

In-Switch Acceleration: A New Distributed Computing Paradigm

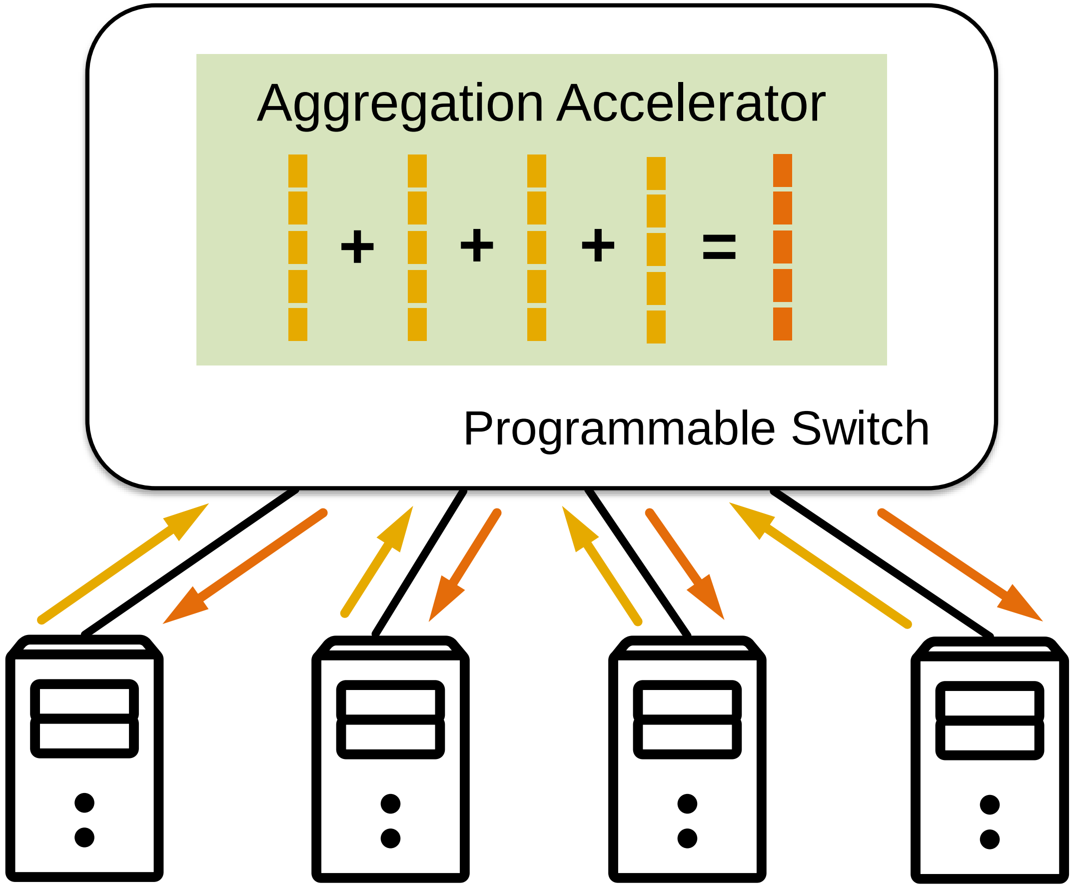

Performance — Reduce End-to-End Network Latency

Programmability — Hardware-Algorithm Co-Design

In-Switch Acceleration: A New Distributed Computing Paradigm

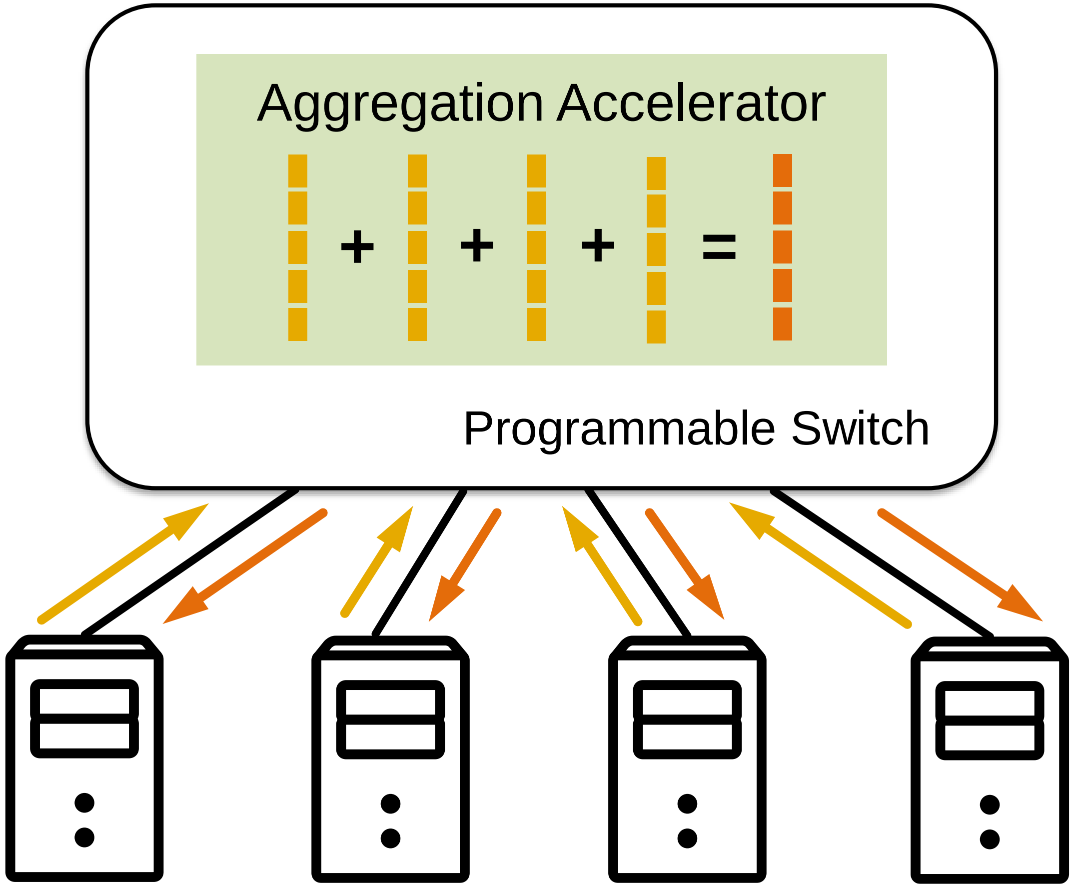

Performance	Reduce End-to-End Network Latency
Programmability	Hardware-Algorithm Co-Design
Scalability	Scale Training at Rack Scale

Challenges of In-Switch Acceleration

No Impact on Regular Switch Functions

Challenges of In-Switch Acceleration

No Impact on Regular Switch Functions

Limited On-Chip Resources

Challenges of In-Switch Acceleration

No Impact on Regular Switch Functions

Limited On-Chip Resources

Scale with More Switches and Nodes

Basics of Programmable Switch

Control Plane

Data Plane

Basics of Programmable Switch

Control Plane

Data Plane

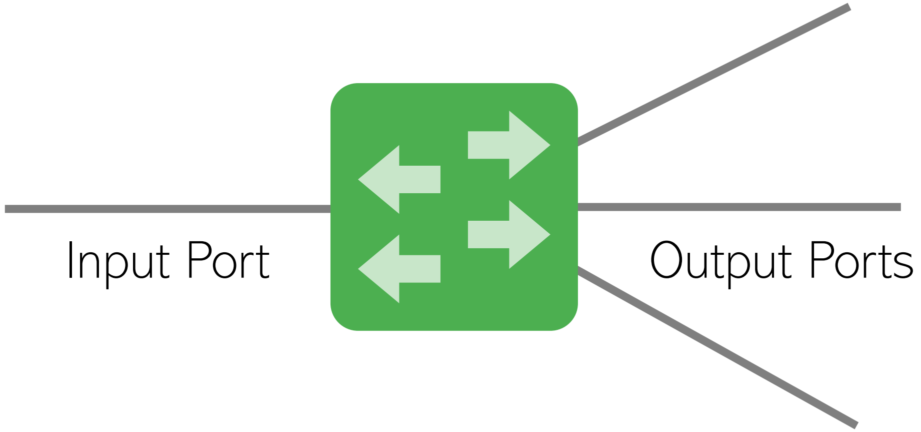

Basics of Programmable Switch

Control Plane

Data Plane

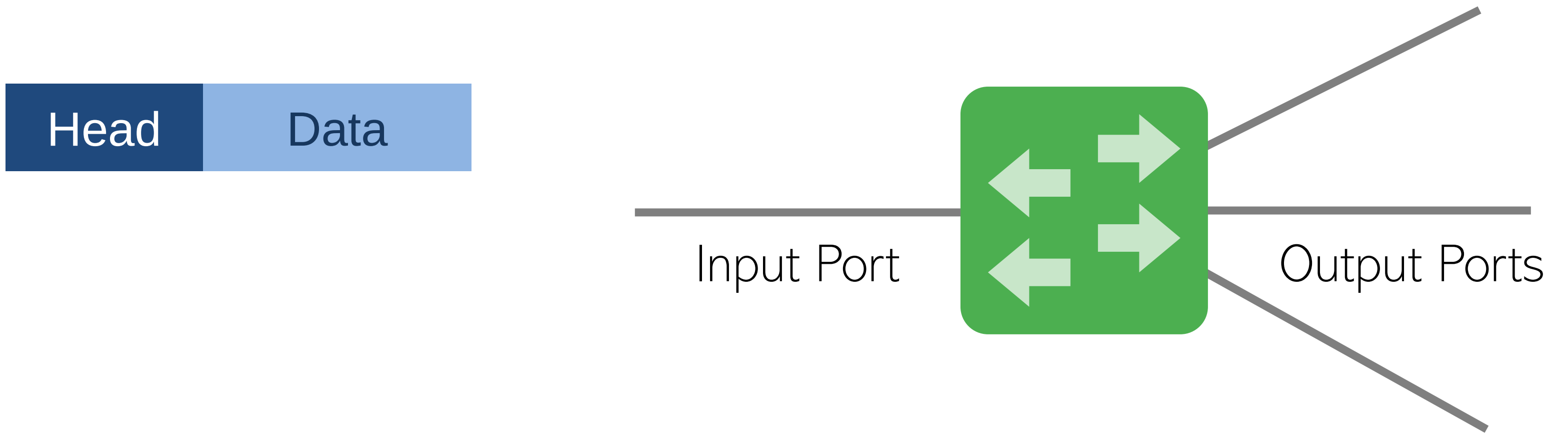

Basics of Programmable Switch

Control Plane

Data Plane

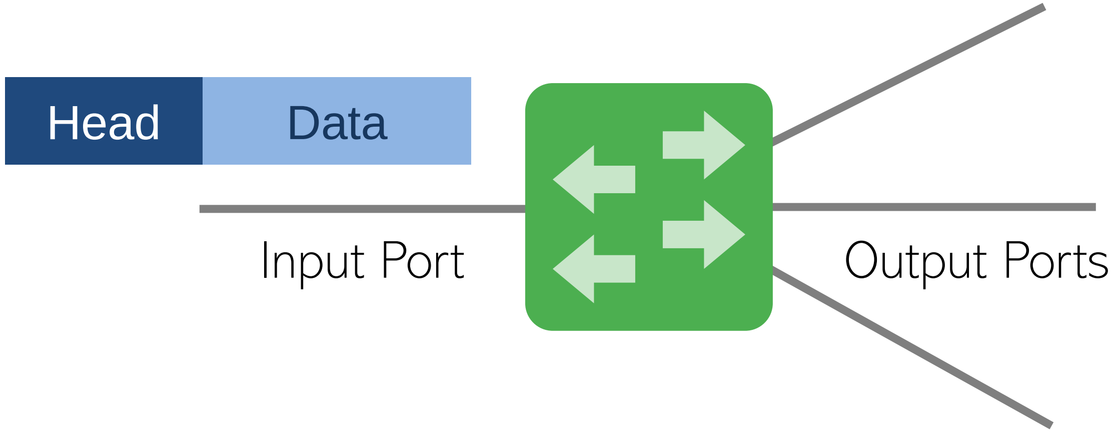

Basics of Programmable Switch

Control Plane

Data Plane

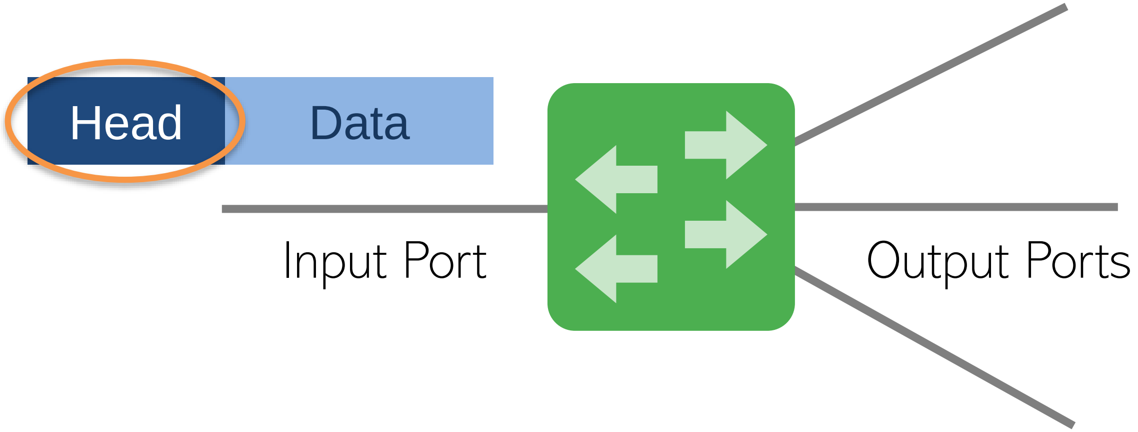

Basics of Programmable Switch

Control Plane

Data Plane

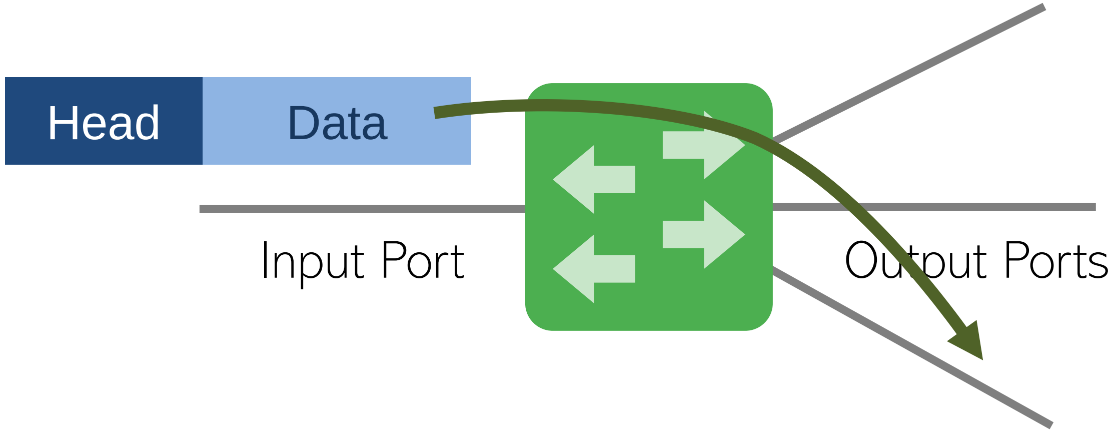

Basics of Programmable Switch

Control Plane

Data Plane

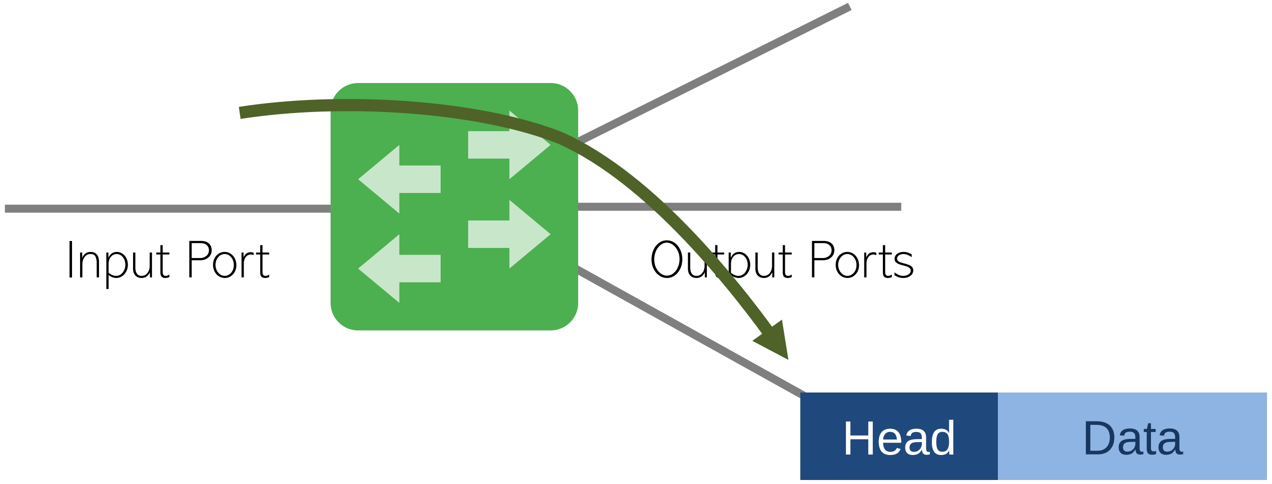

Basics of Programmable Switch

Control Plane

Data Plane

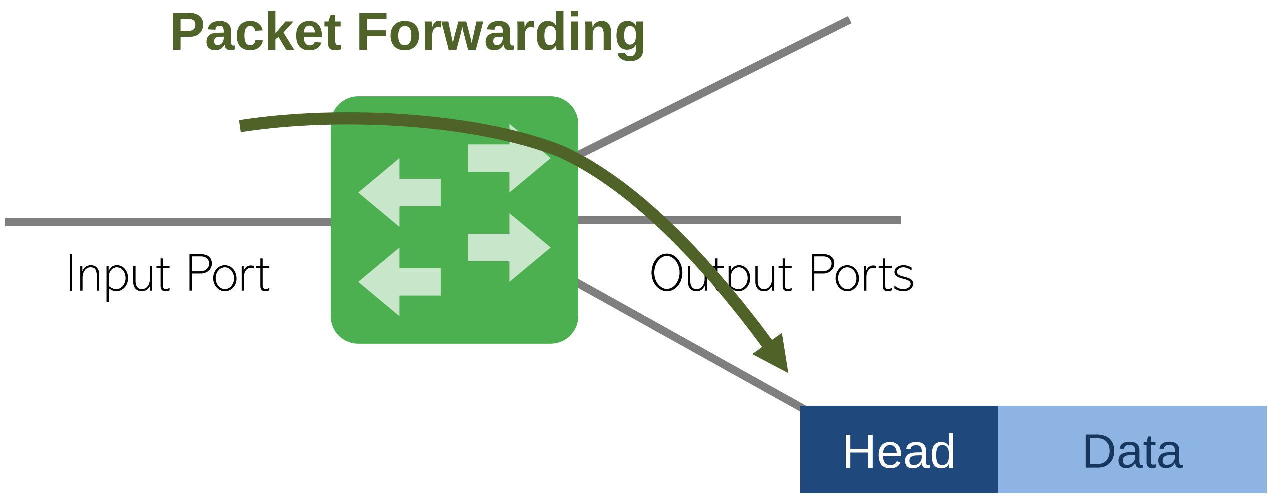

Basics of Programmable Switch

Control Plane

Forwarding Control

Data Plane

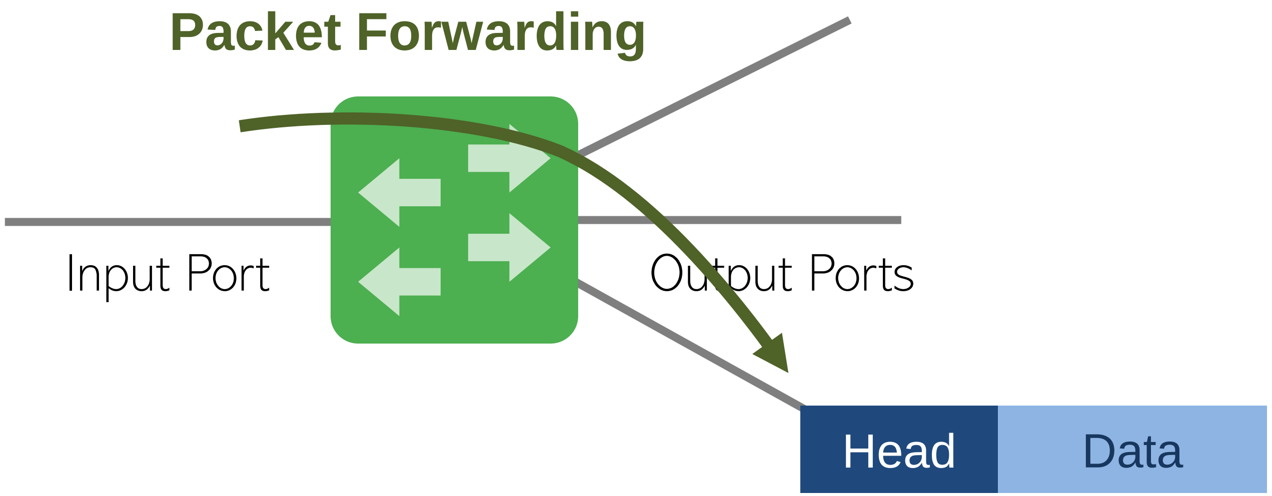

Basics of Programmable Switch

Control Plane

Forwarding Control

System Configuration

Data Plane

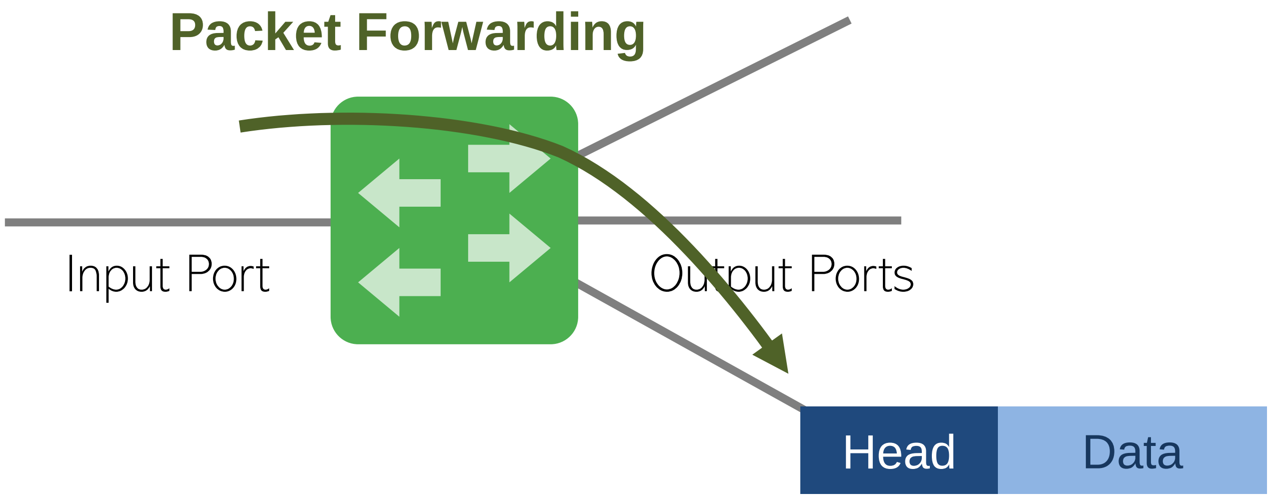

Basics of Programmable Switch

Control Plane

Forwarding Control

System Configuration

…

Data Plane

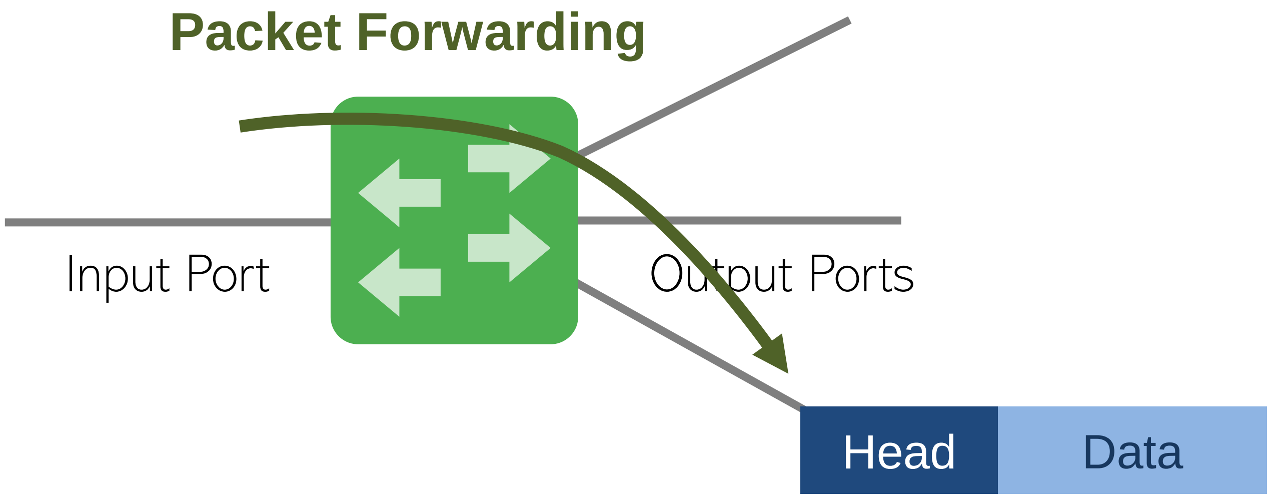

Basics of Programmable Switch

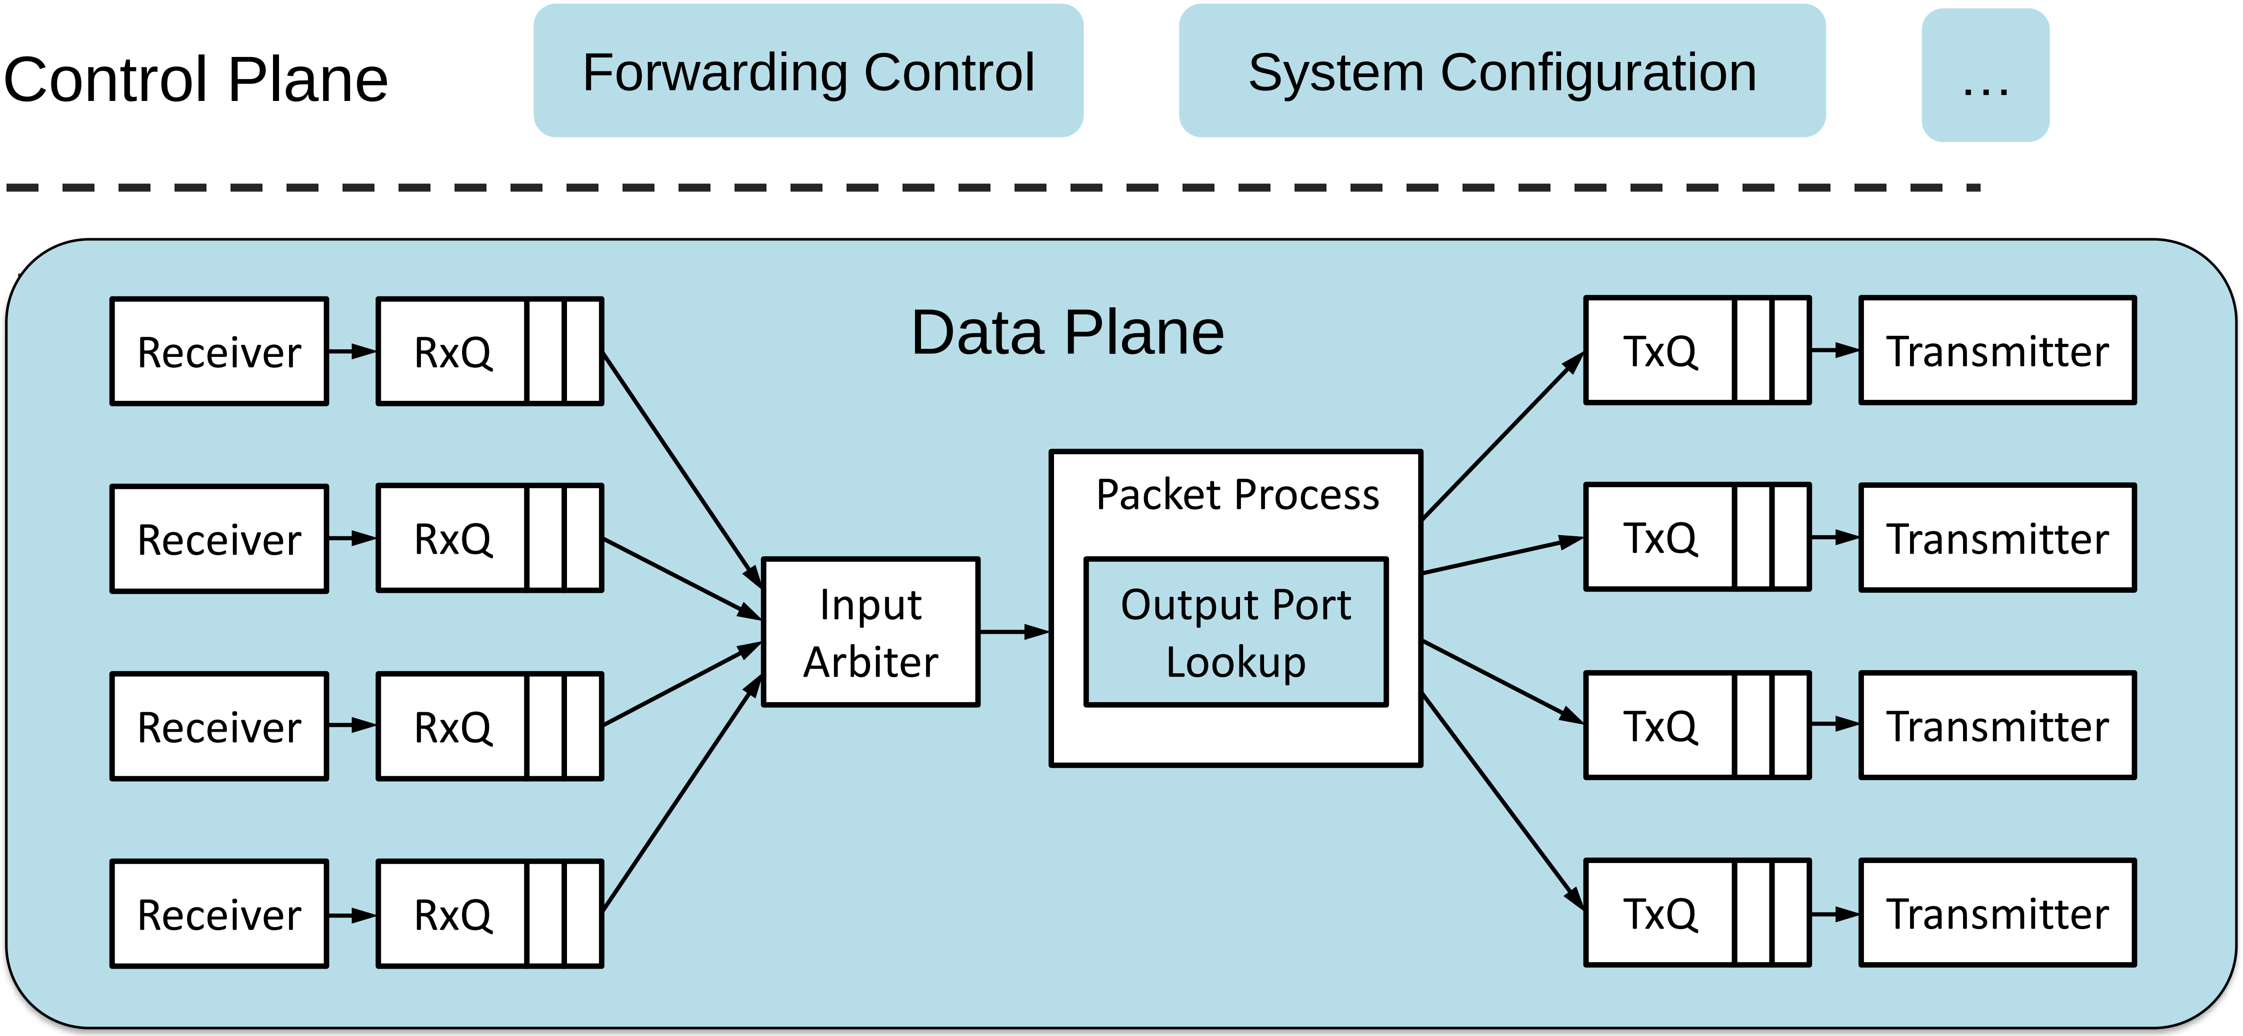

Basics of Programmable Switch

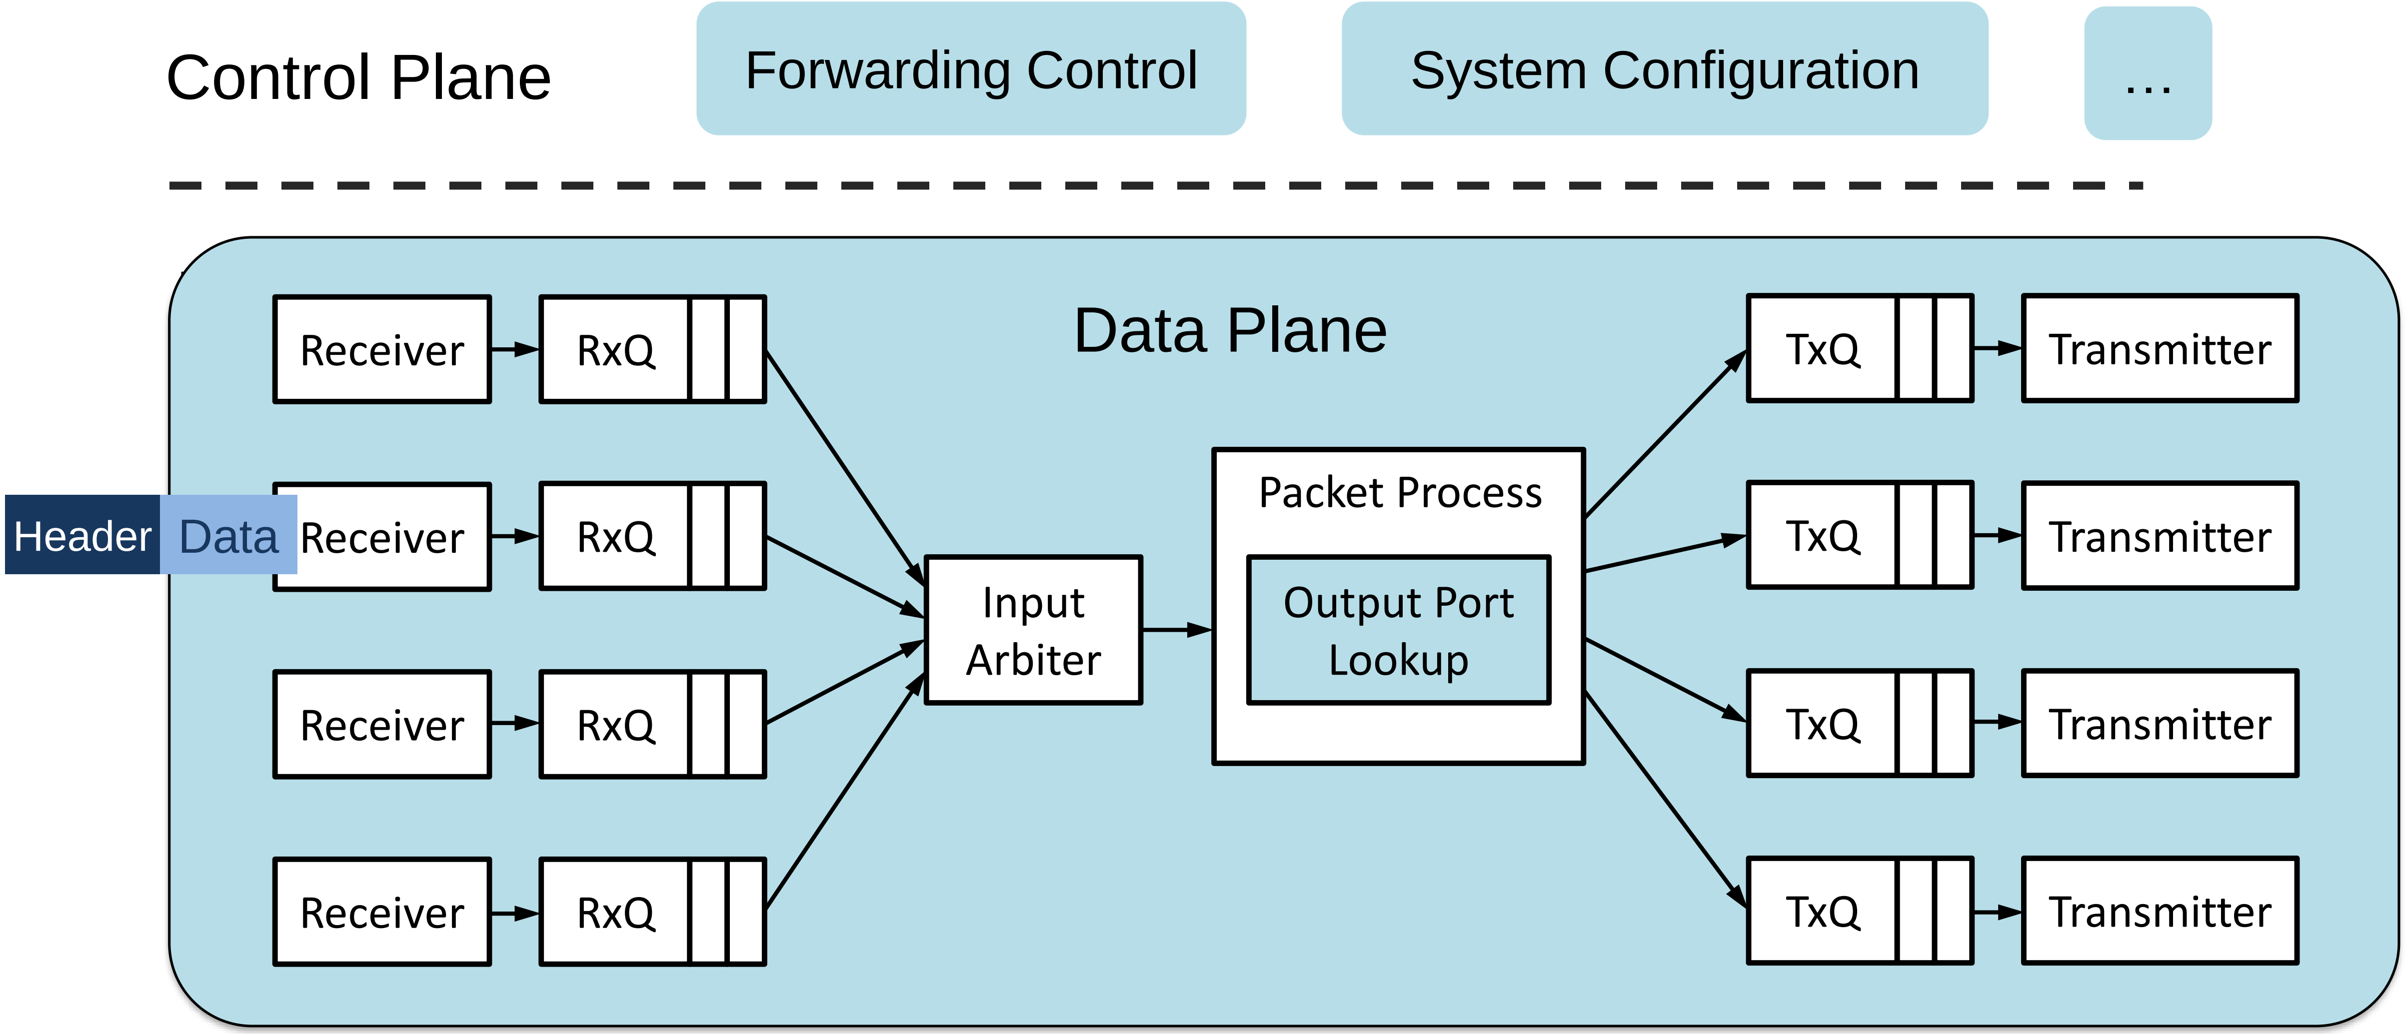

Basics of Programmable Switch

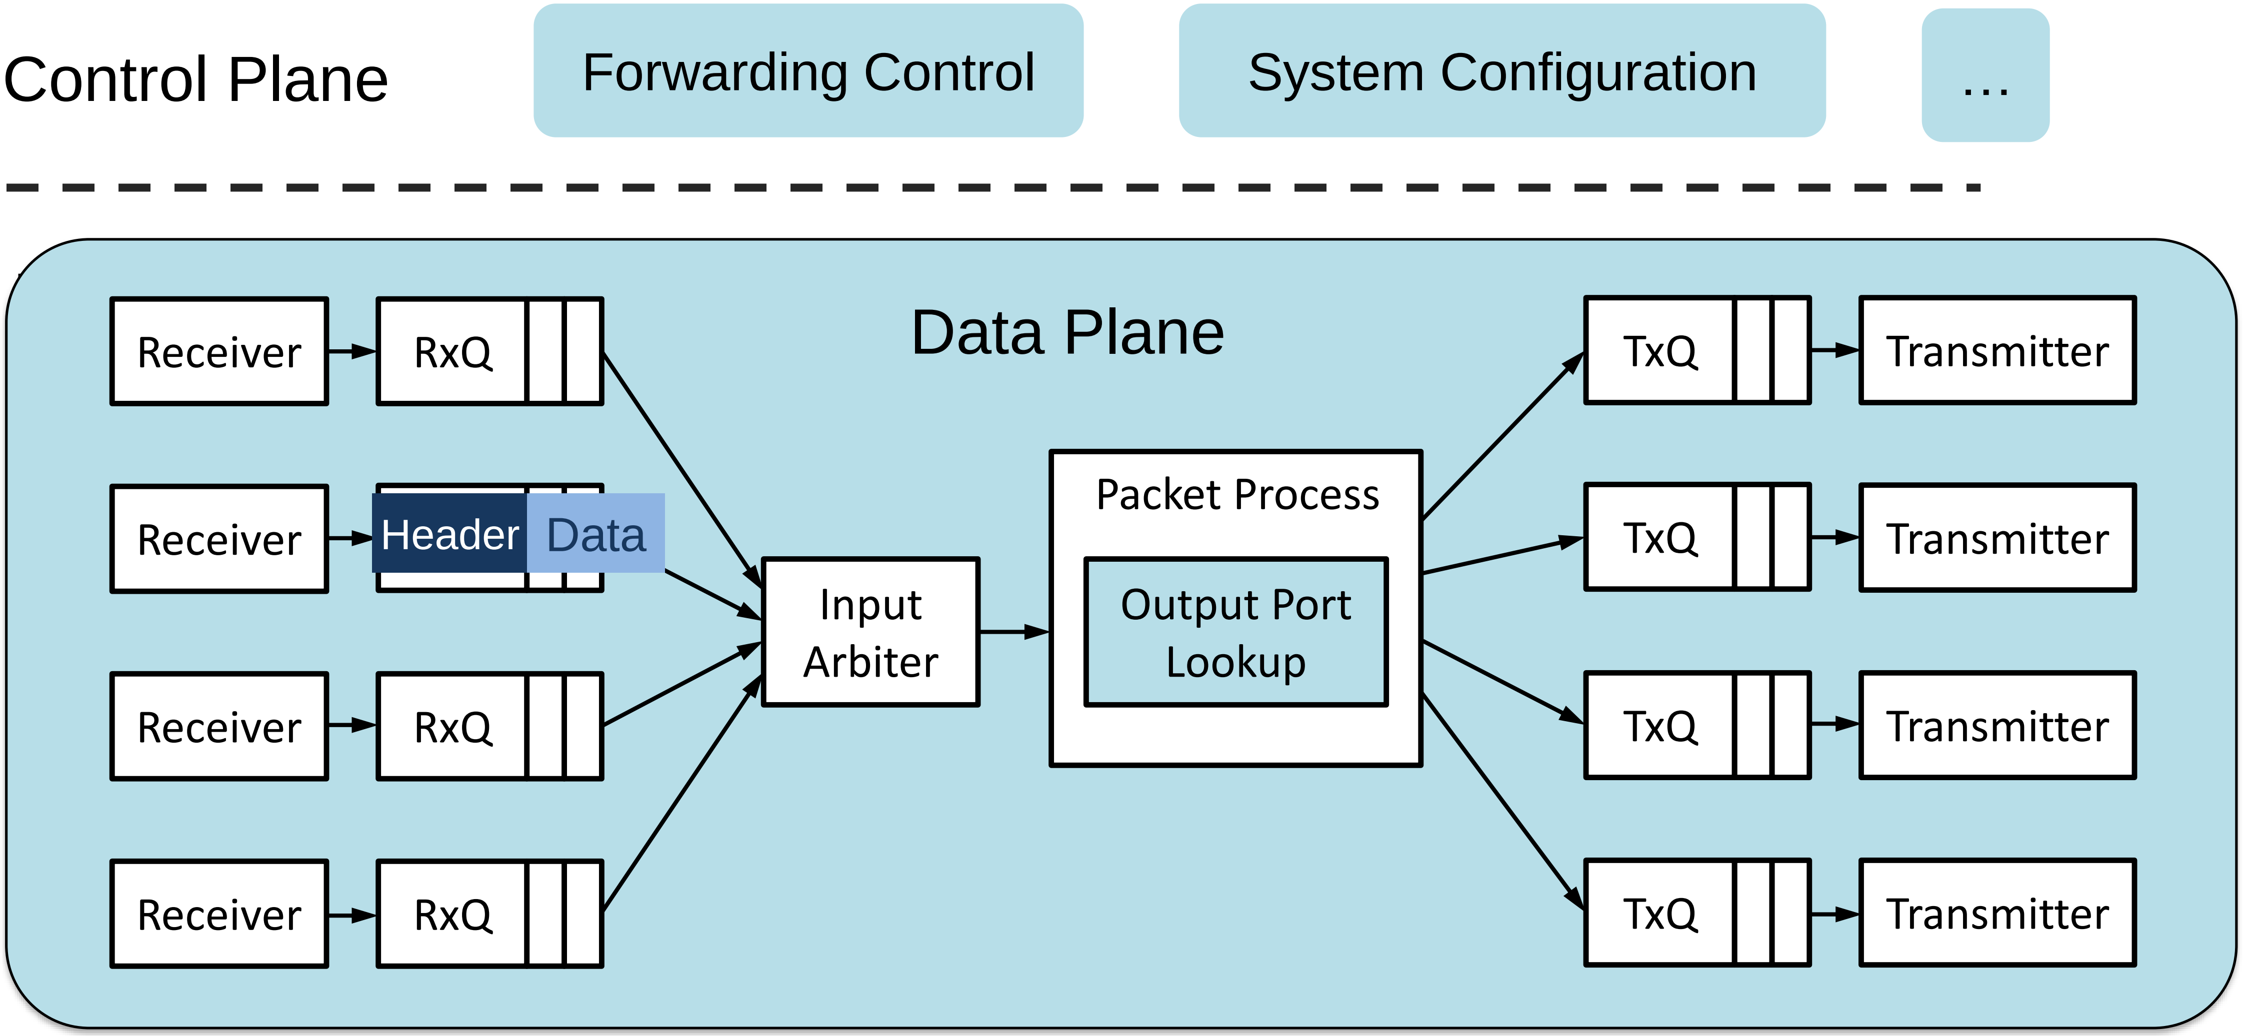

Basics of Programmable Switch

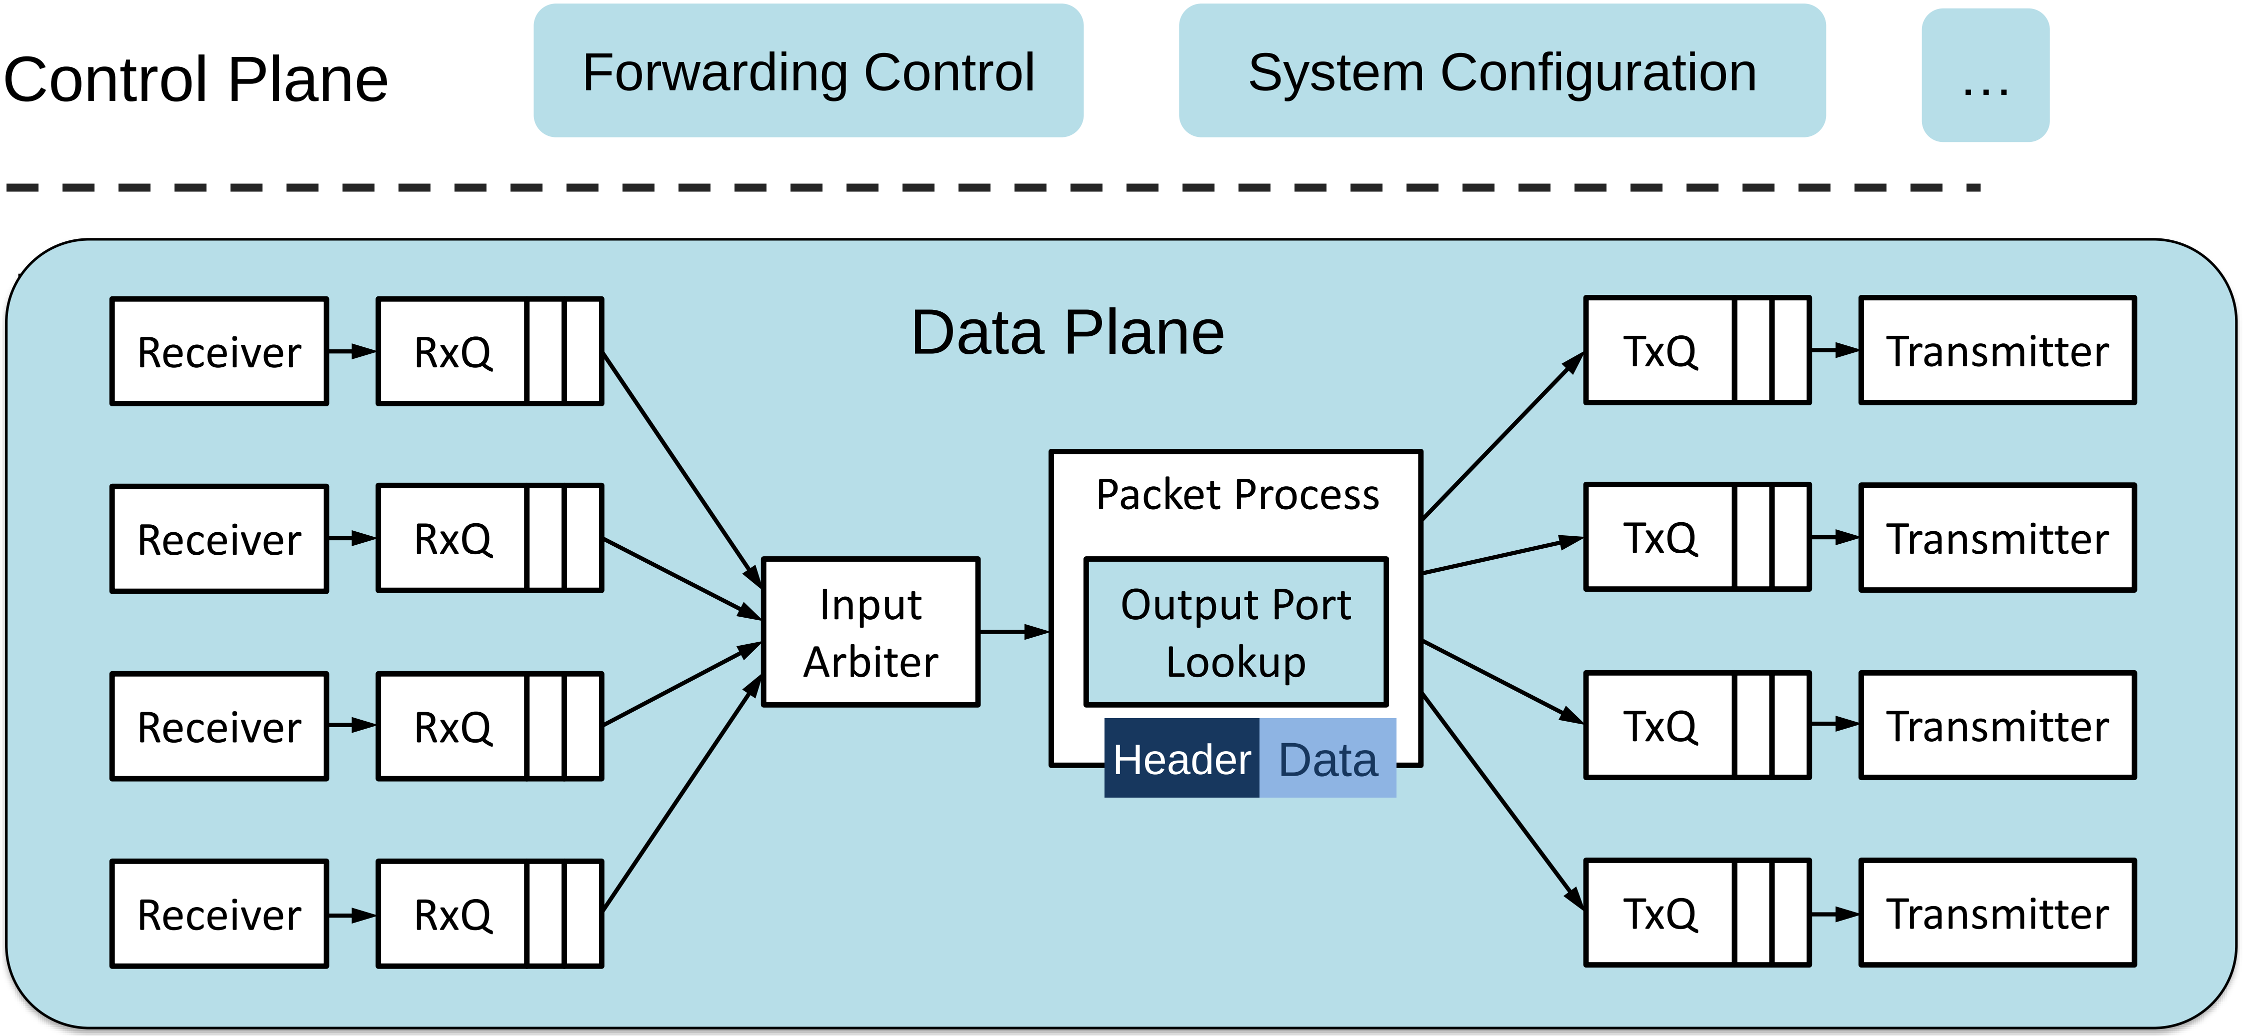

Basics of Programmable Switch

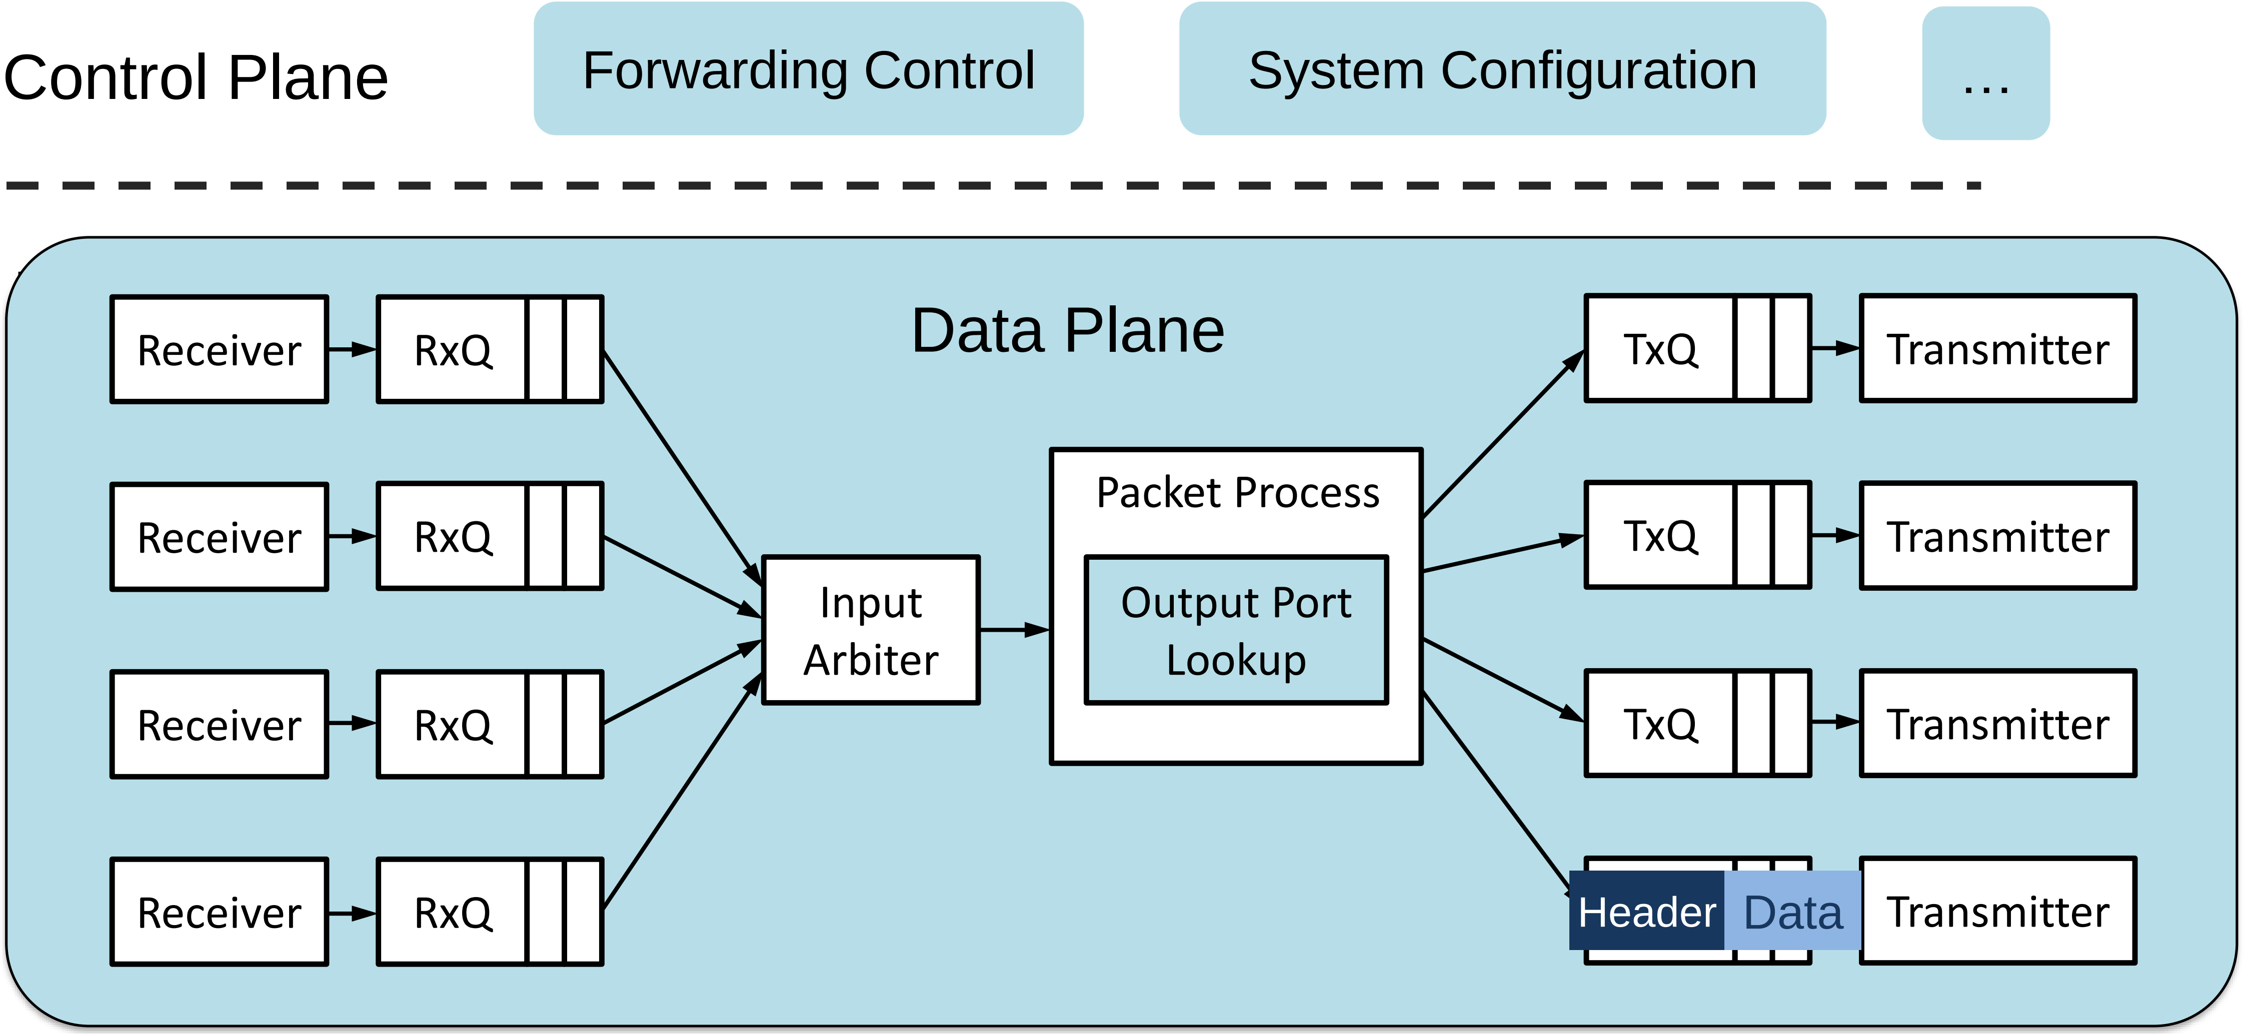

Basics of Programmable Switch

Control Plane

Forwarding Control

System Configuration

…

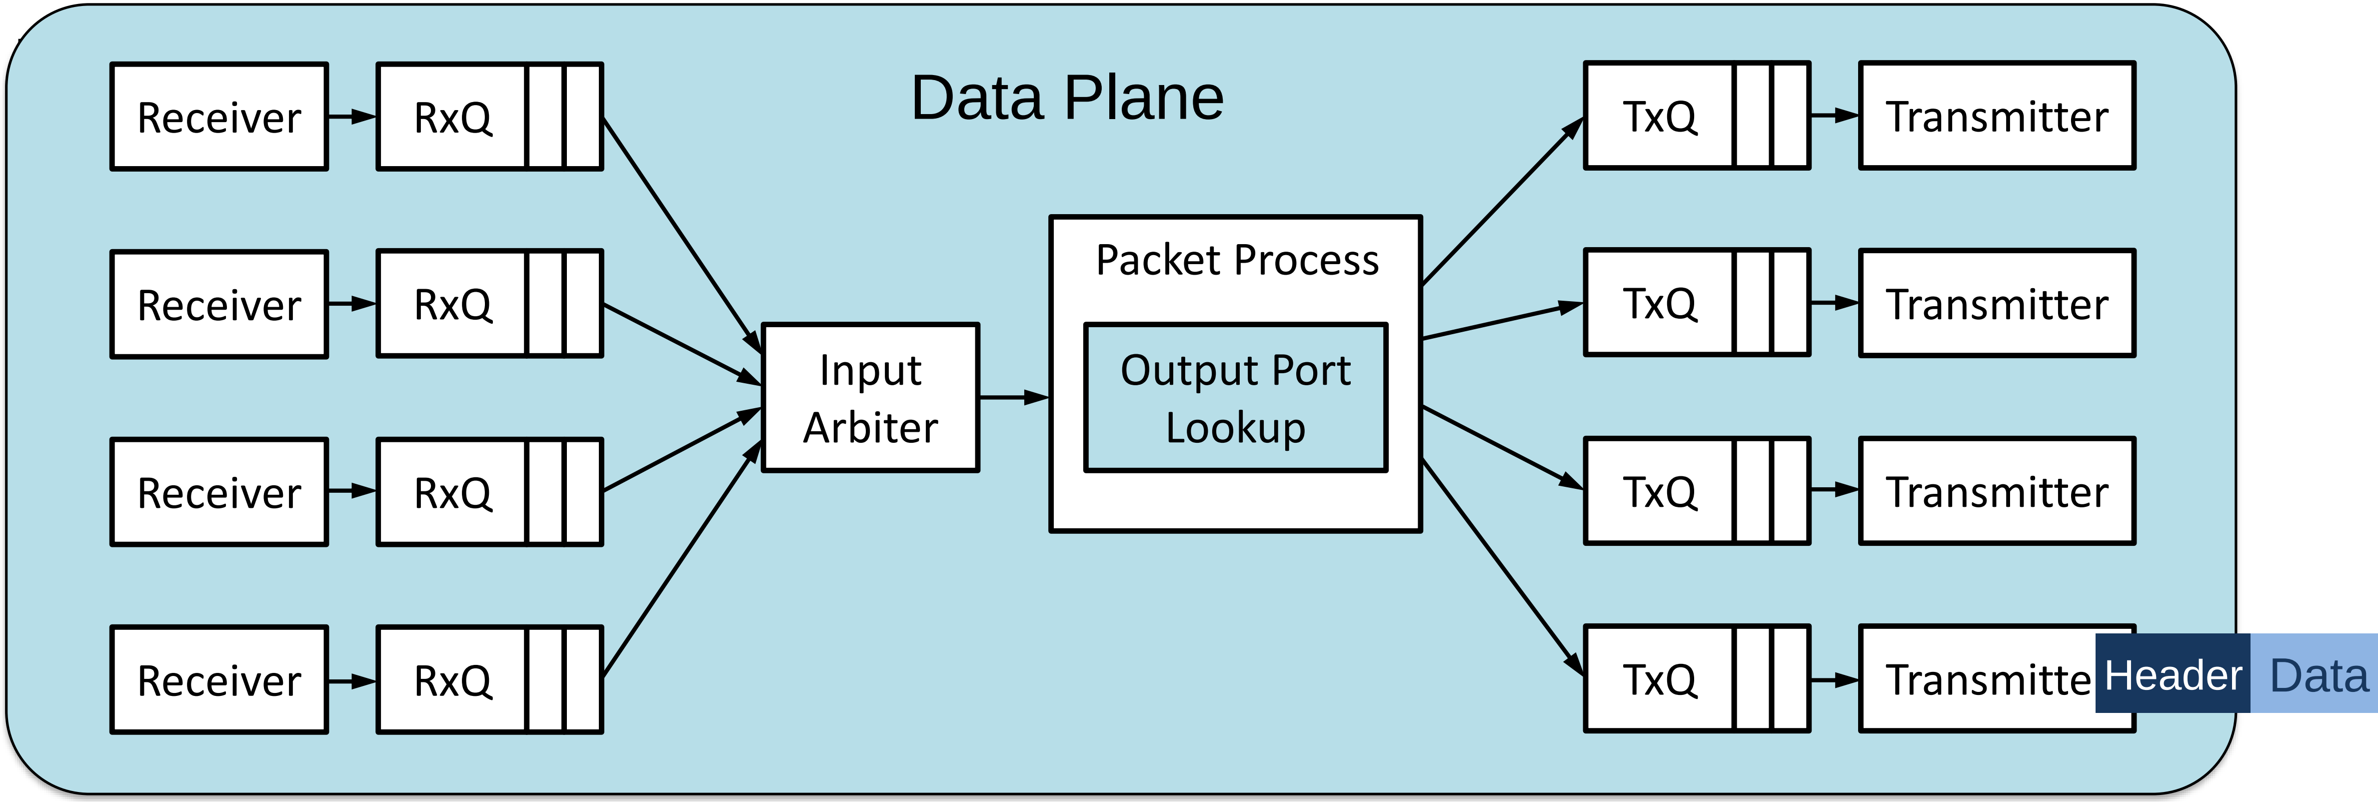

Integrating Aggregation Accelerator into the Programmable Switch

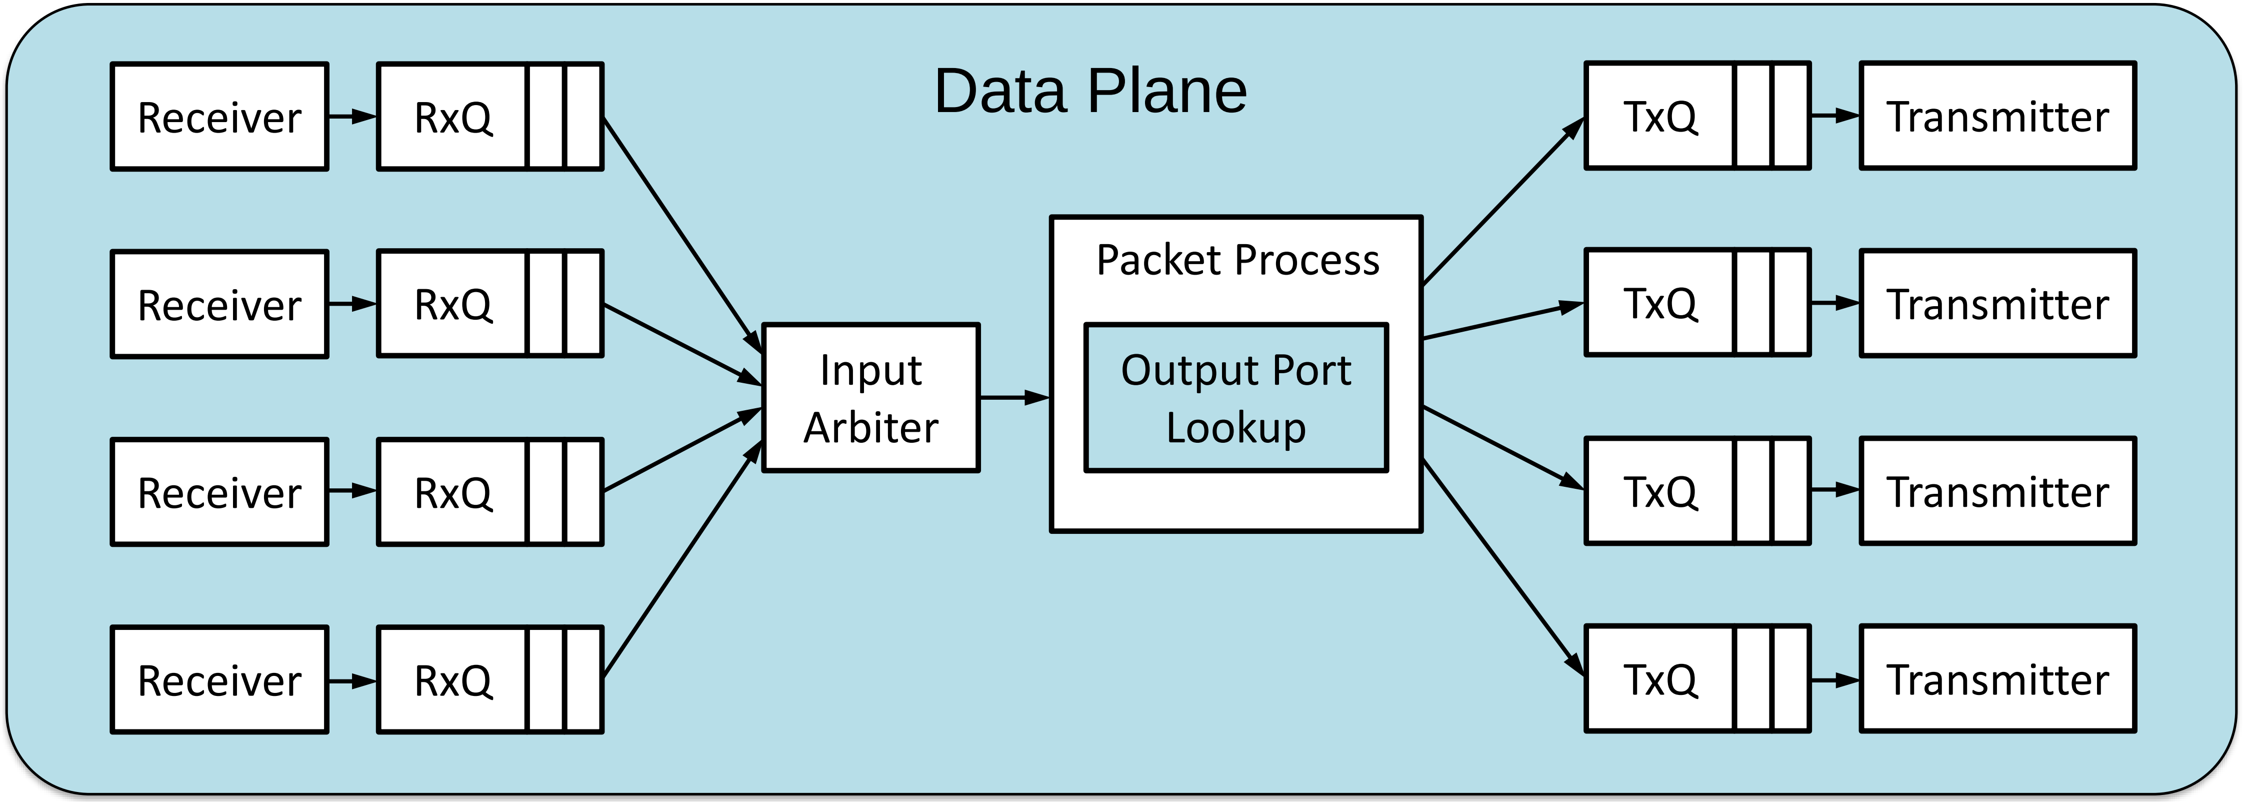

Integrating Aggregation Accelerator into the Programmable Switch

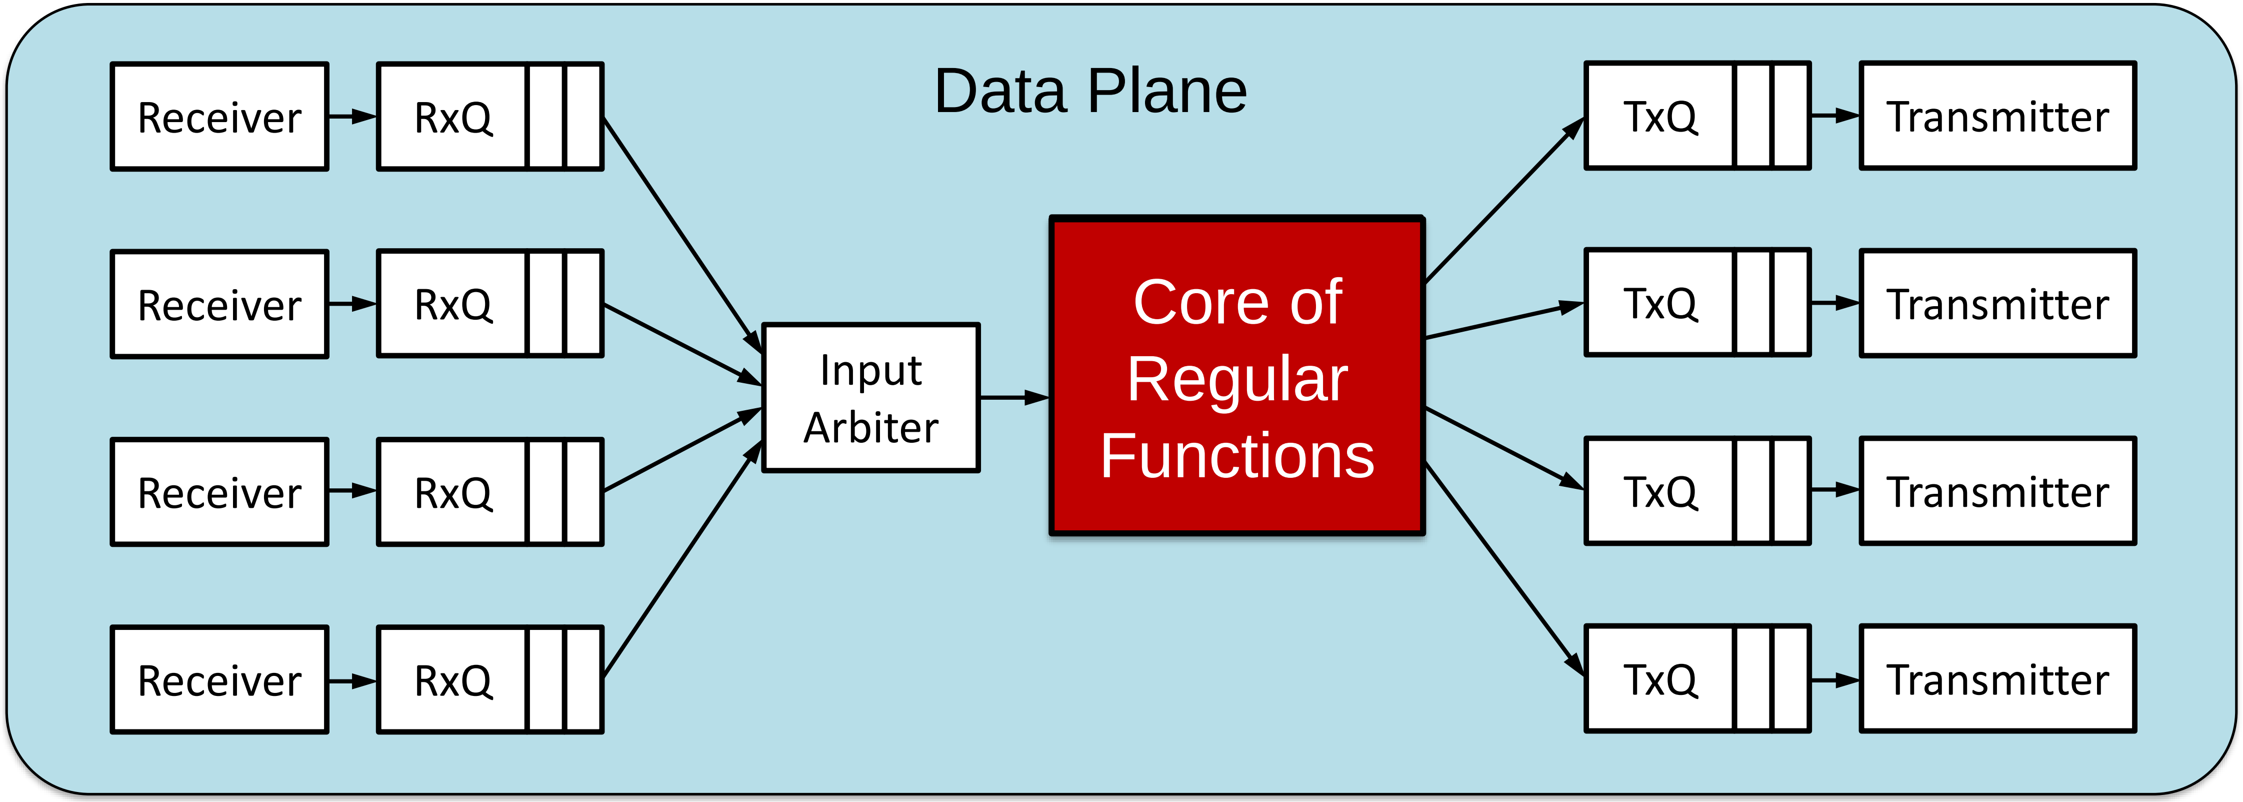

Integrating Aggregation Accelerator into the Programmable Switch

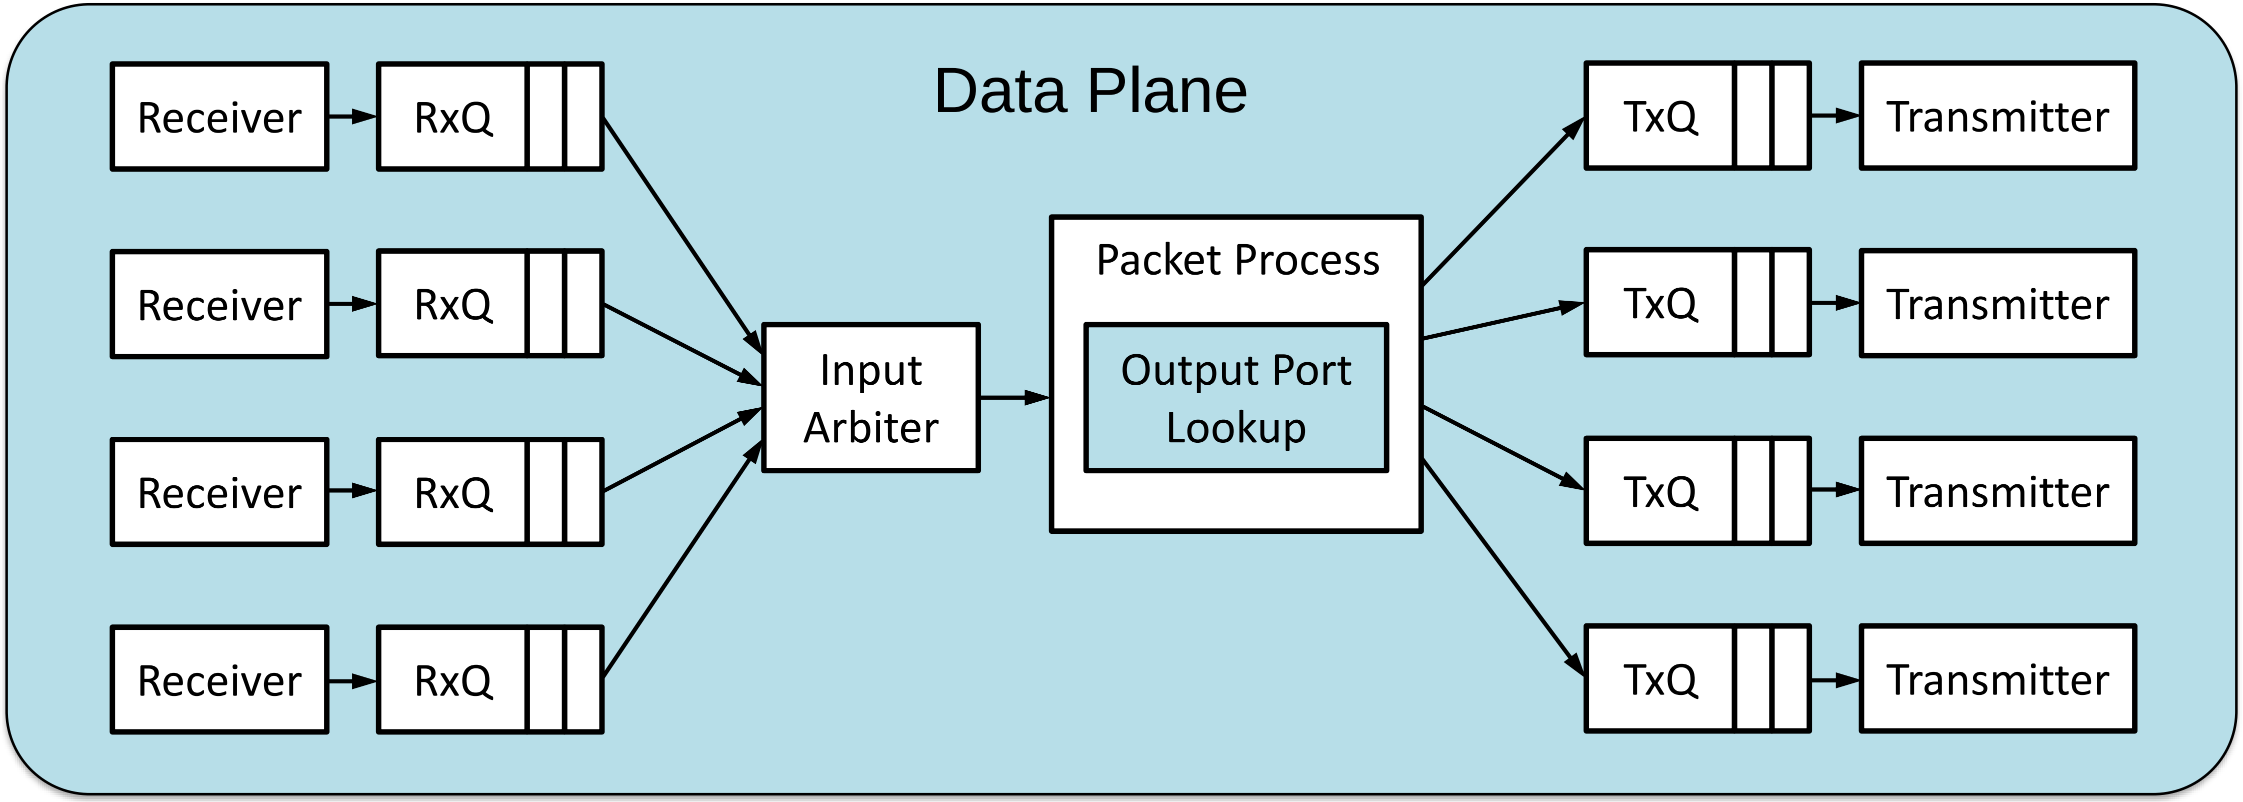

Integrating Aggregation Accelerator into the Programmable Switch

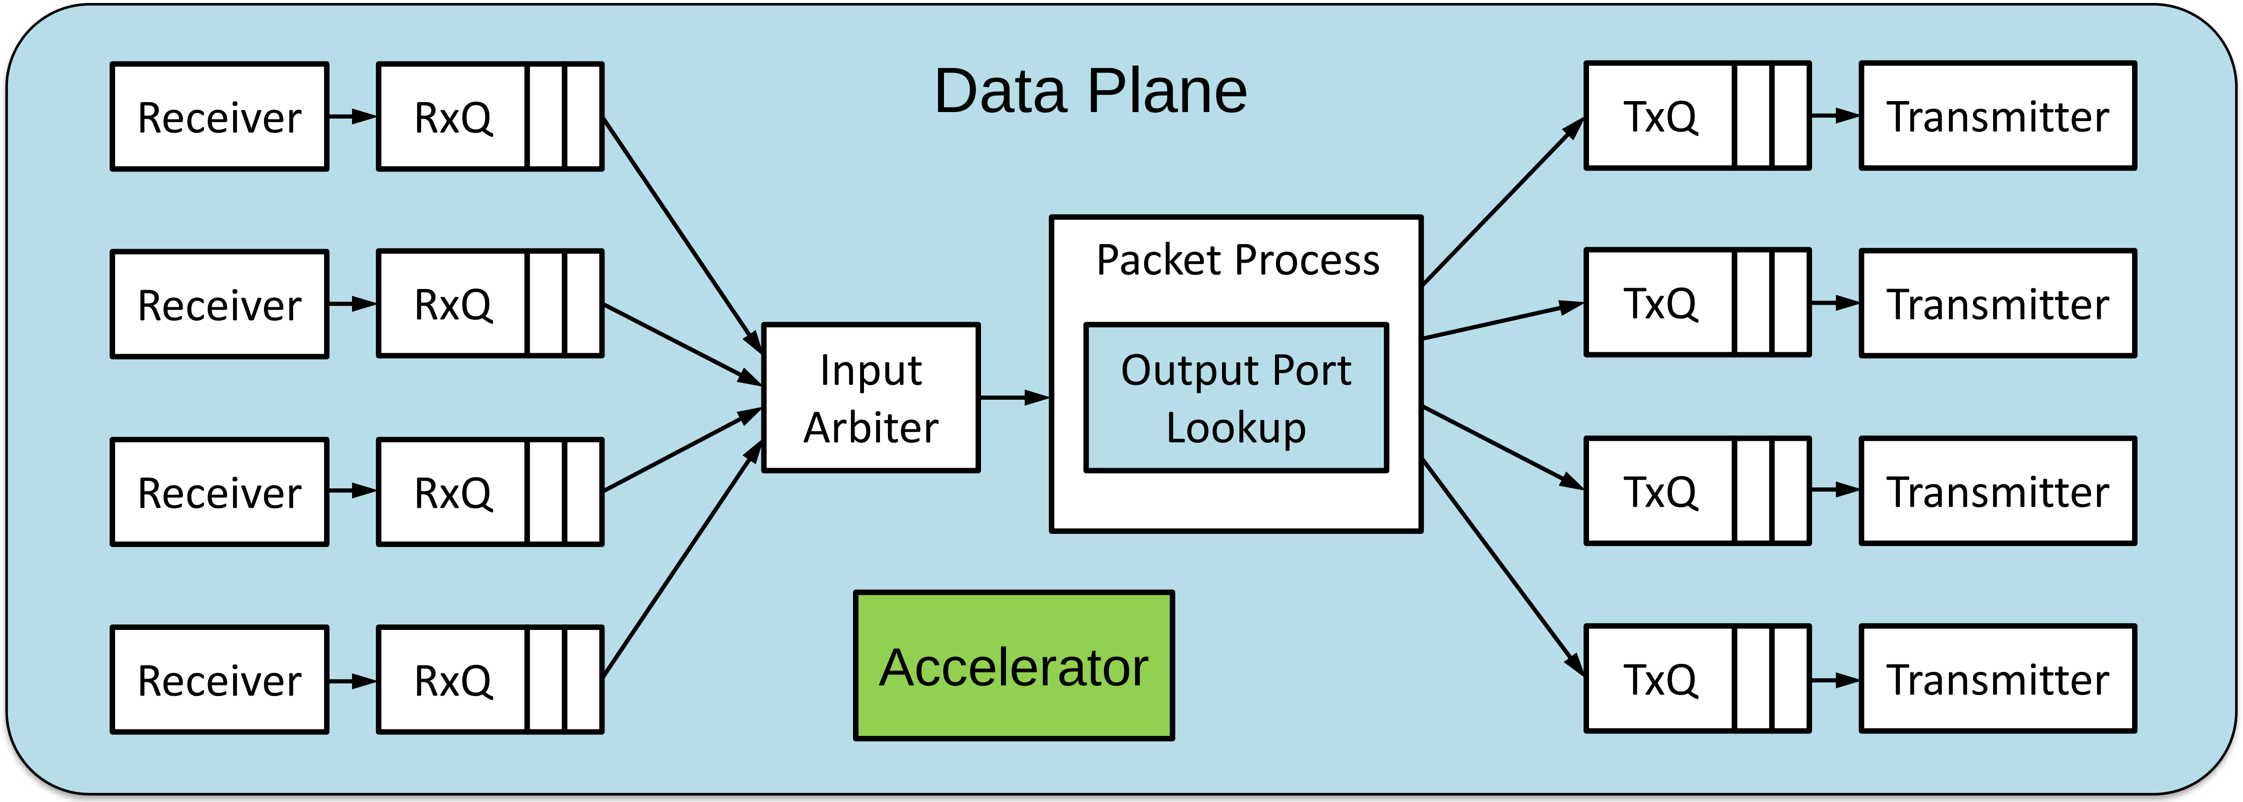

Integrating Aggregation Accelerator into the Programmable Switch

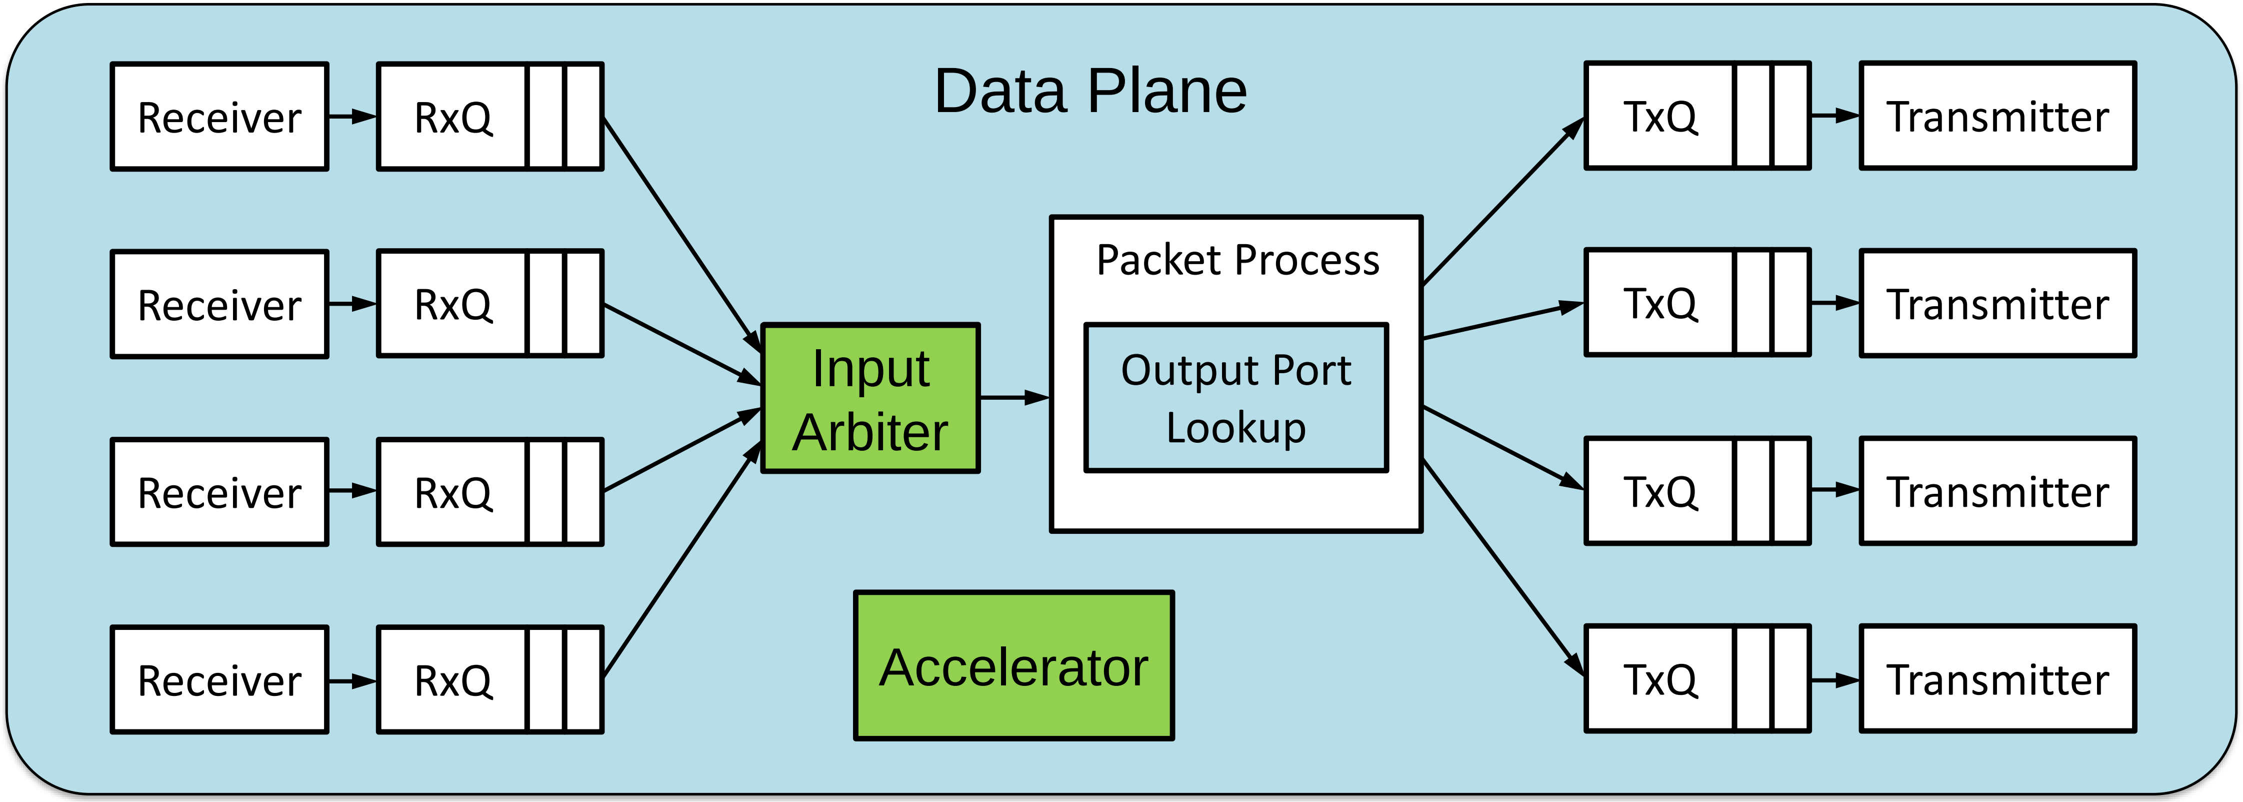

Integrating Aggregation Accelerator into the Programmable Switch

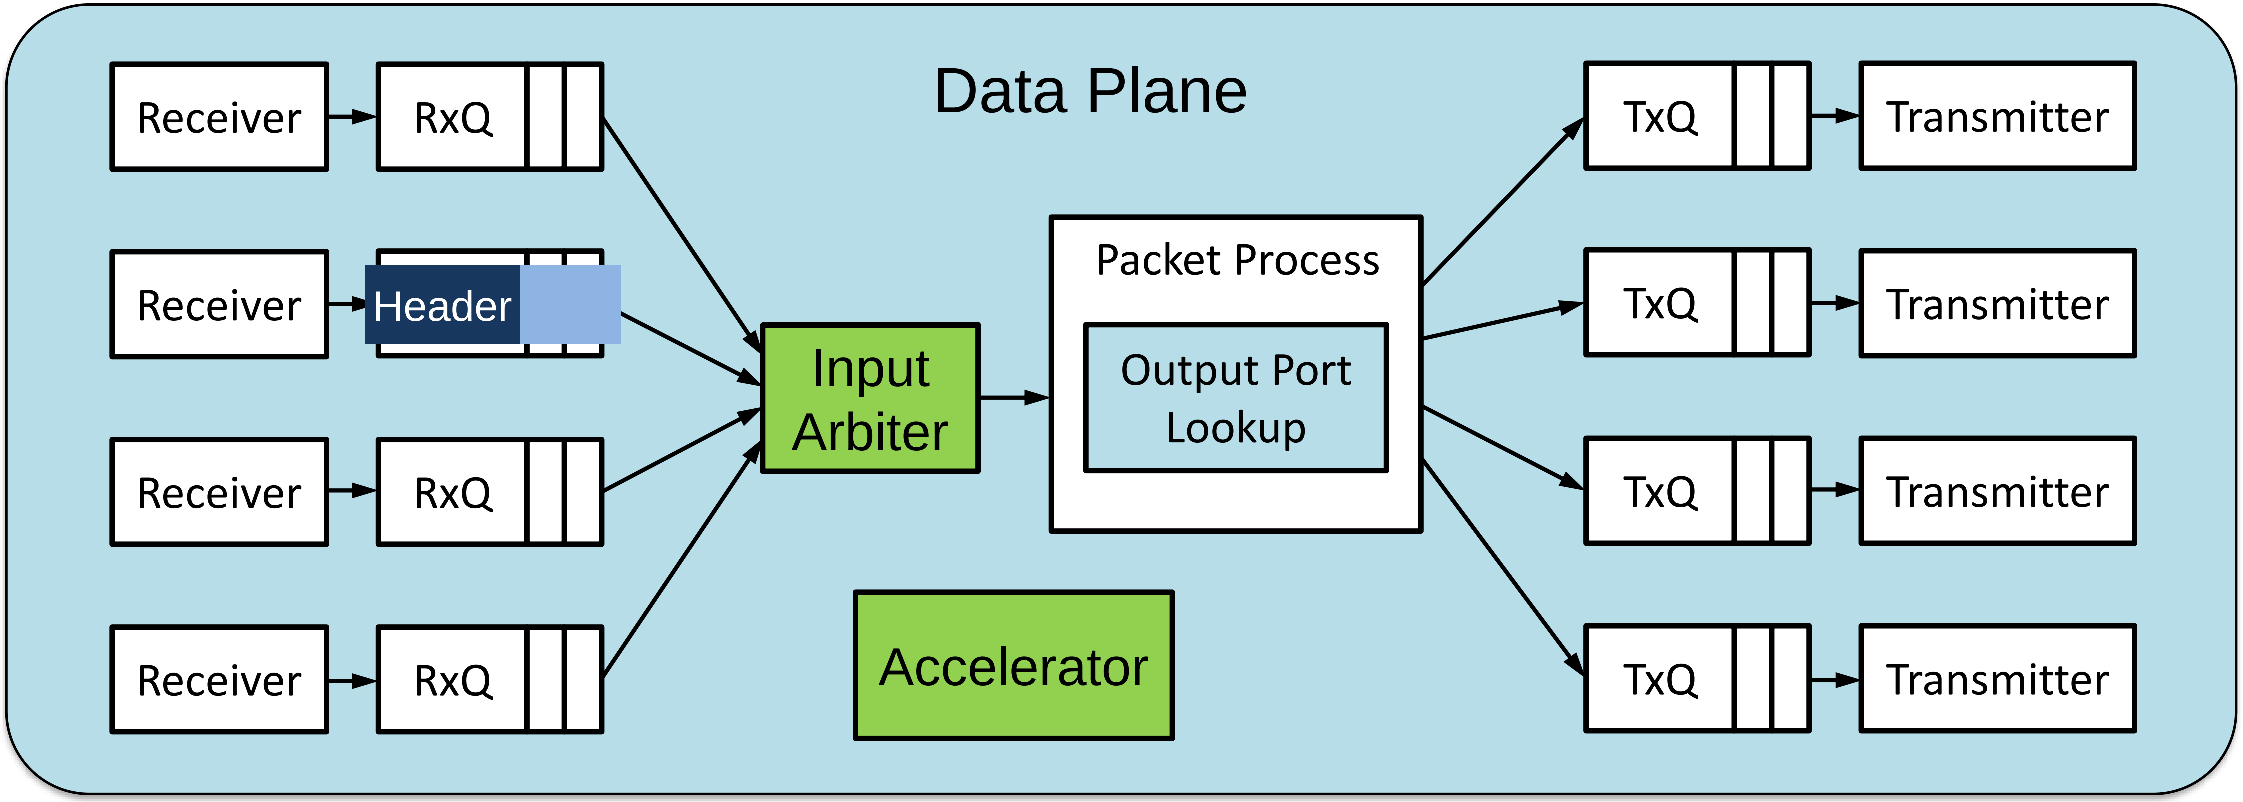

Integrating Aggregation Accelerator into the Programmable Switch

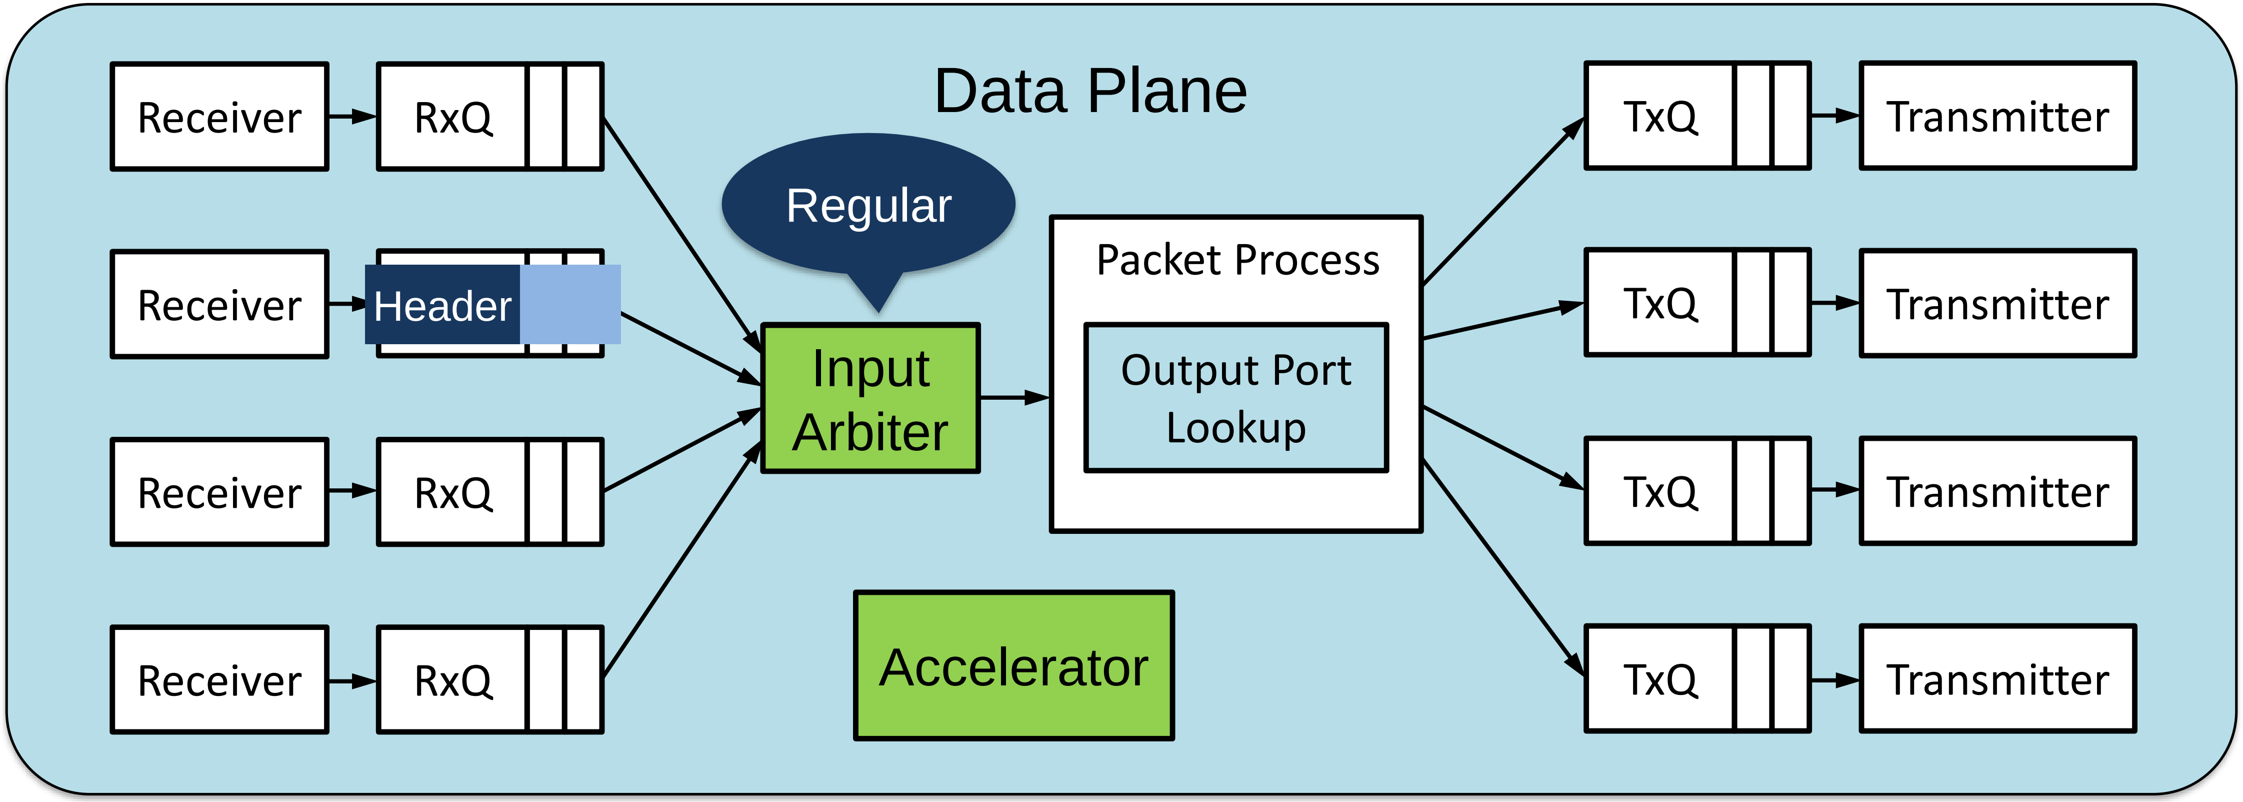

Integrating Aggregation Accelerator into the Programmable Switch

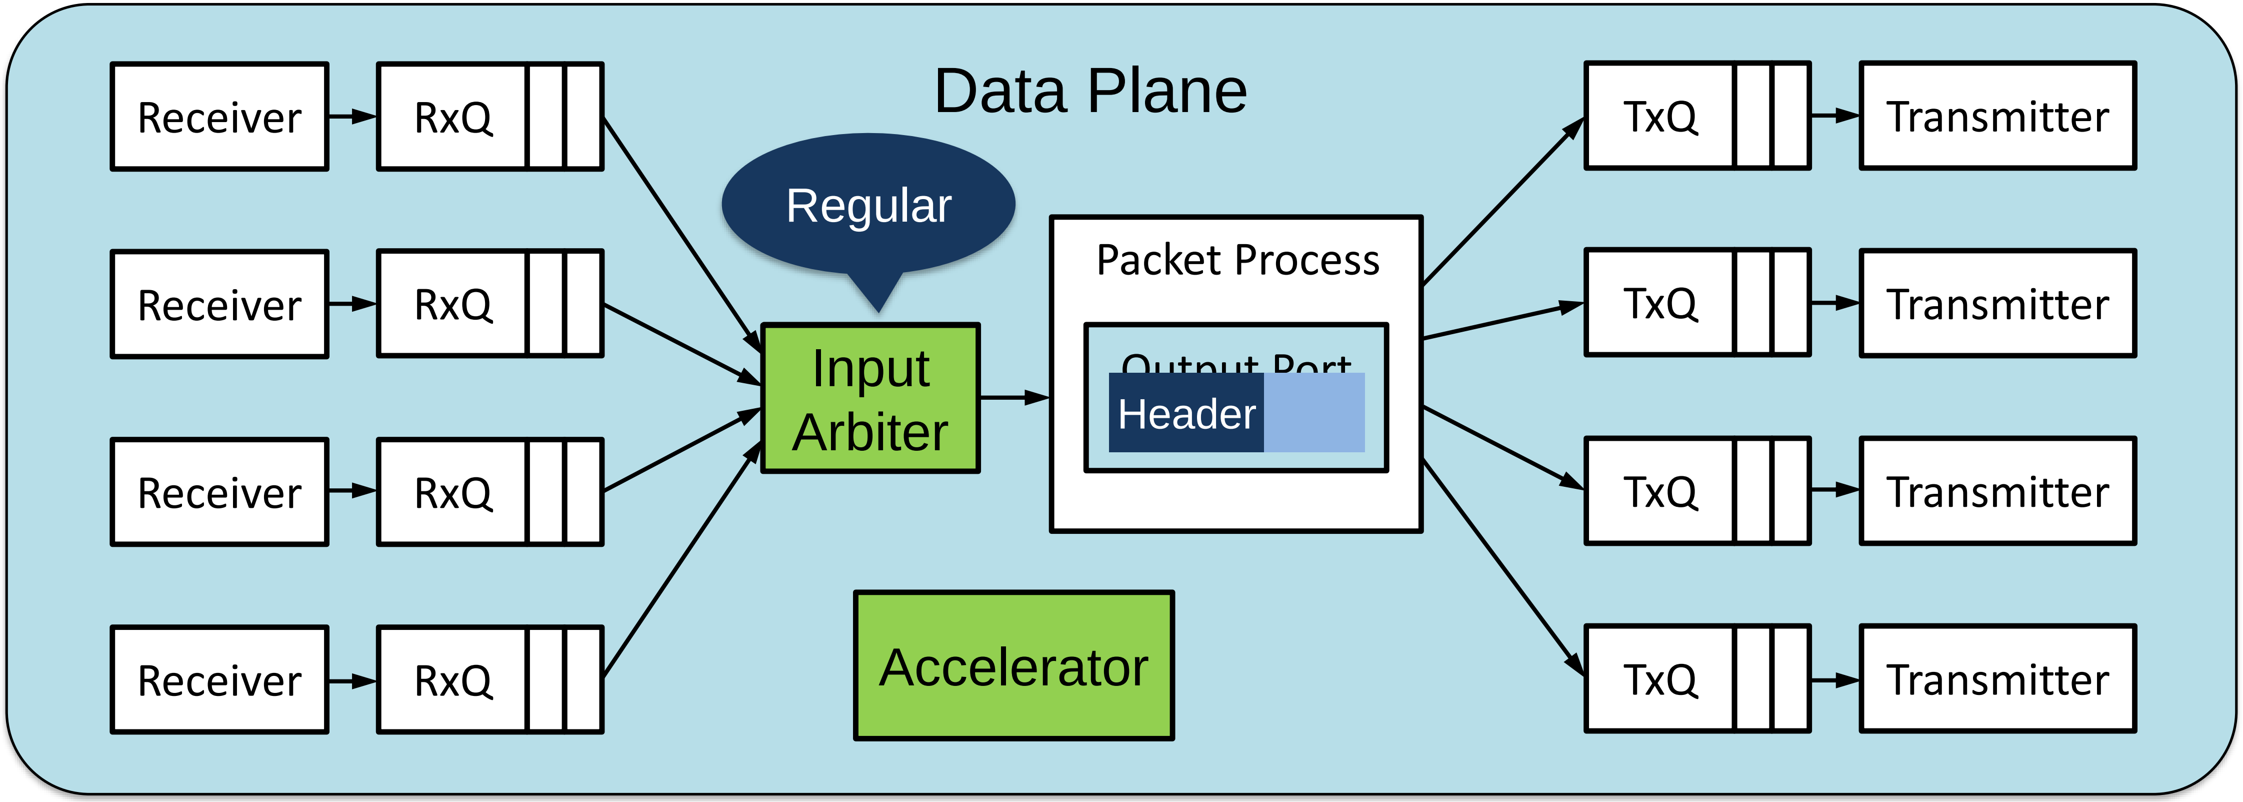

Integrating Aggregation Accelerator into the Programmable Switch

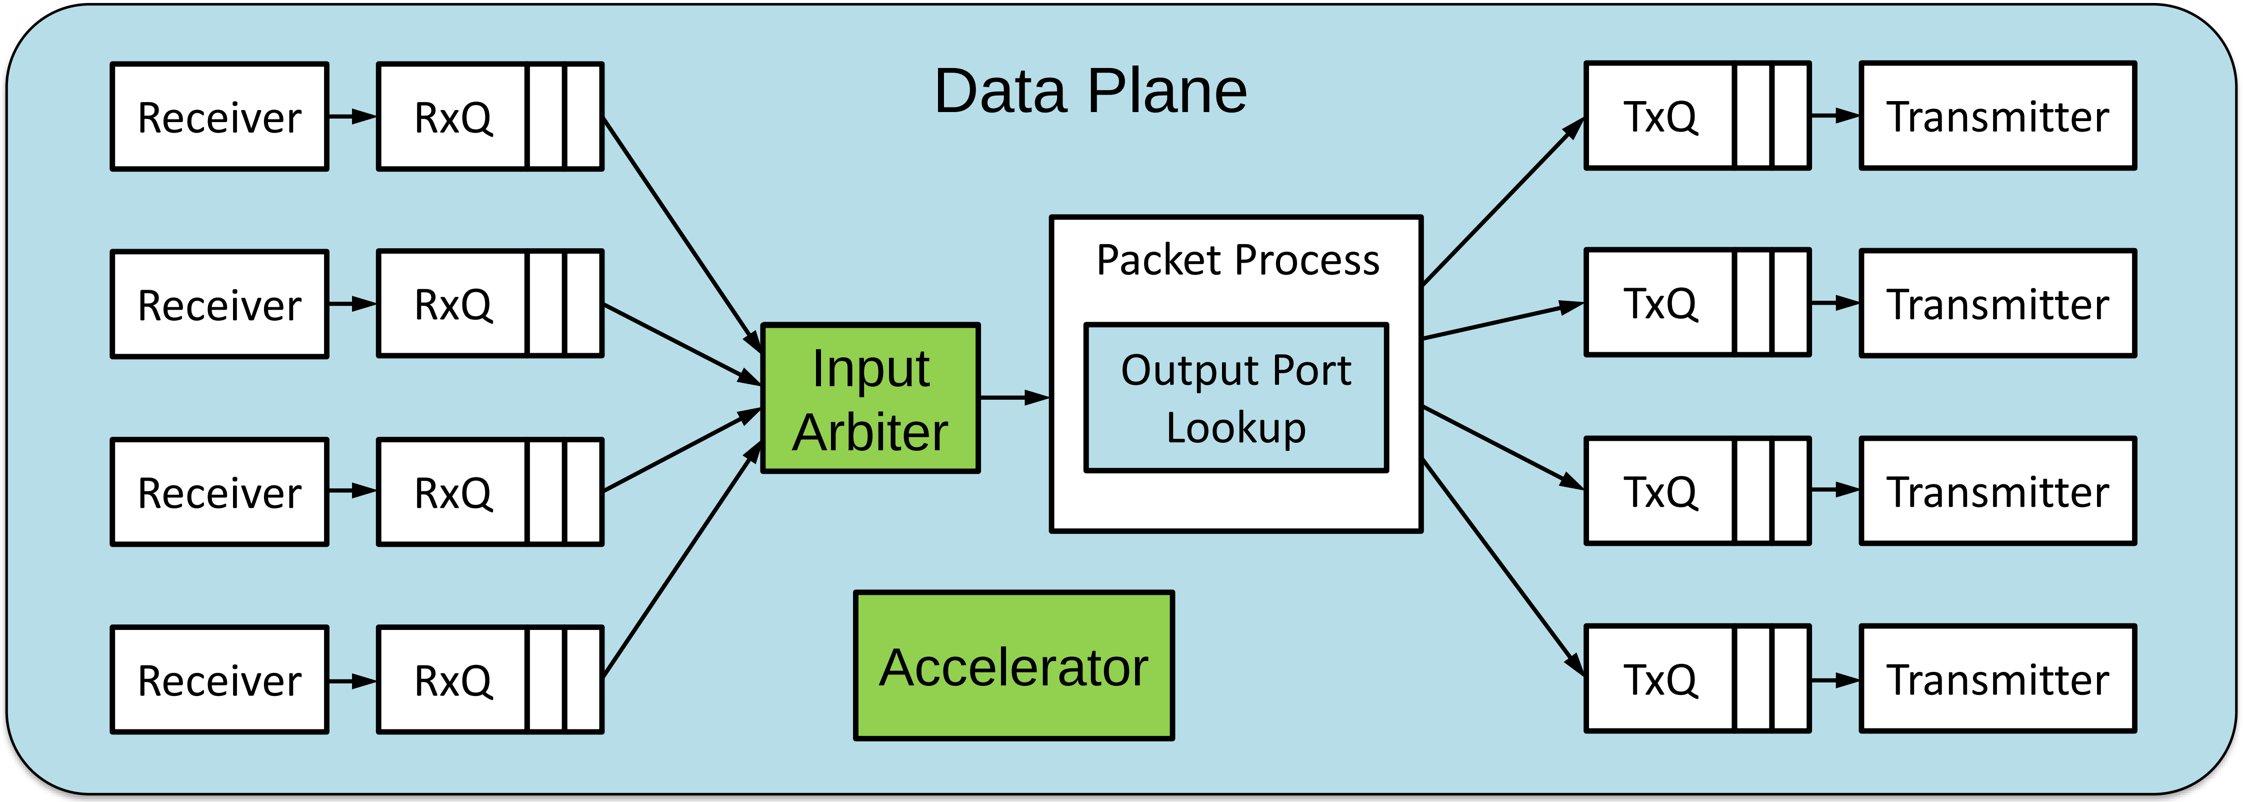

Integrating Aggregation Accelerator into the Programmable Switch

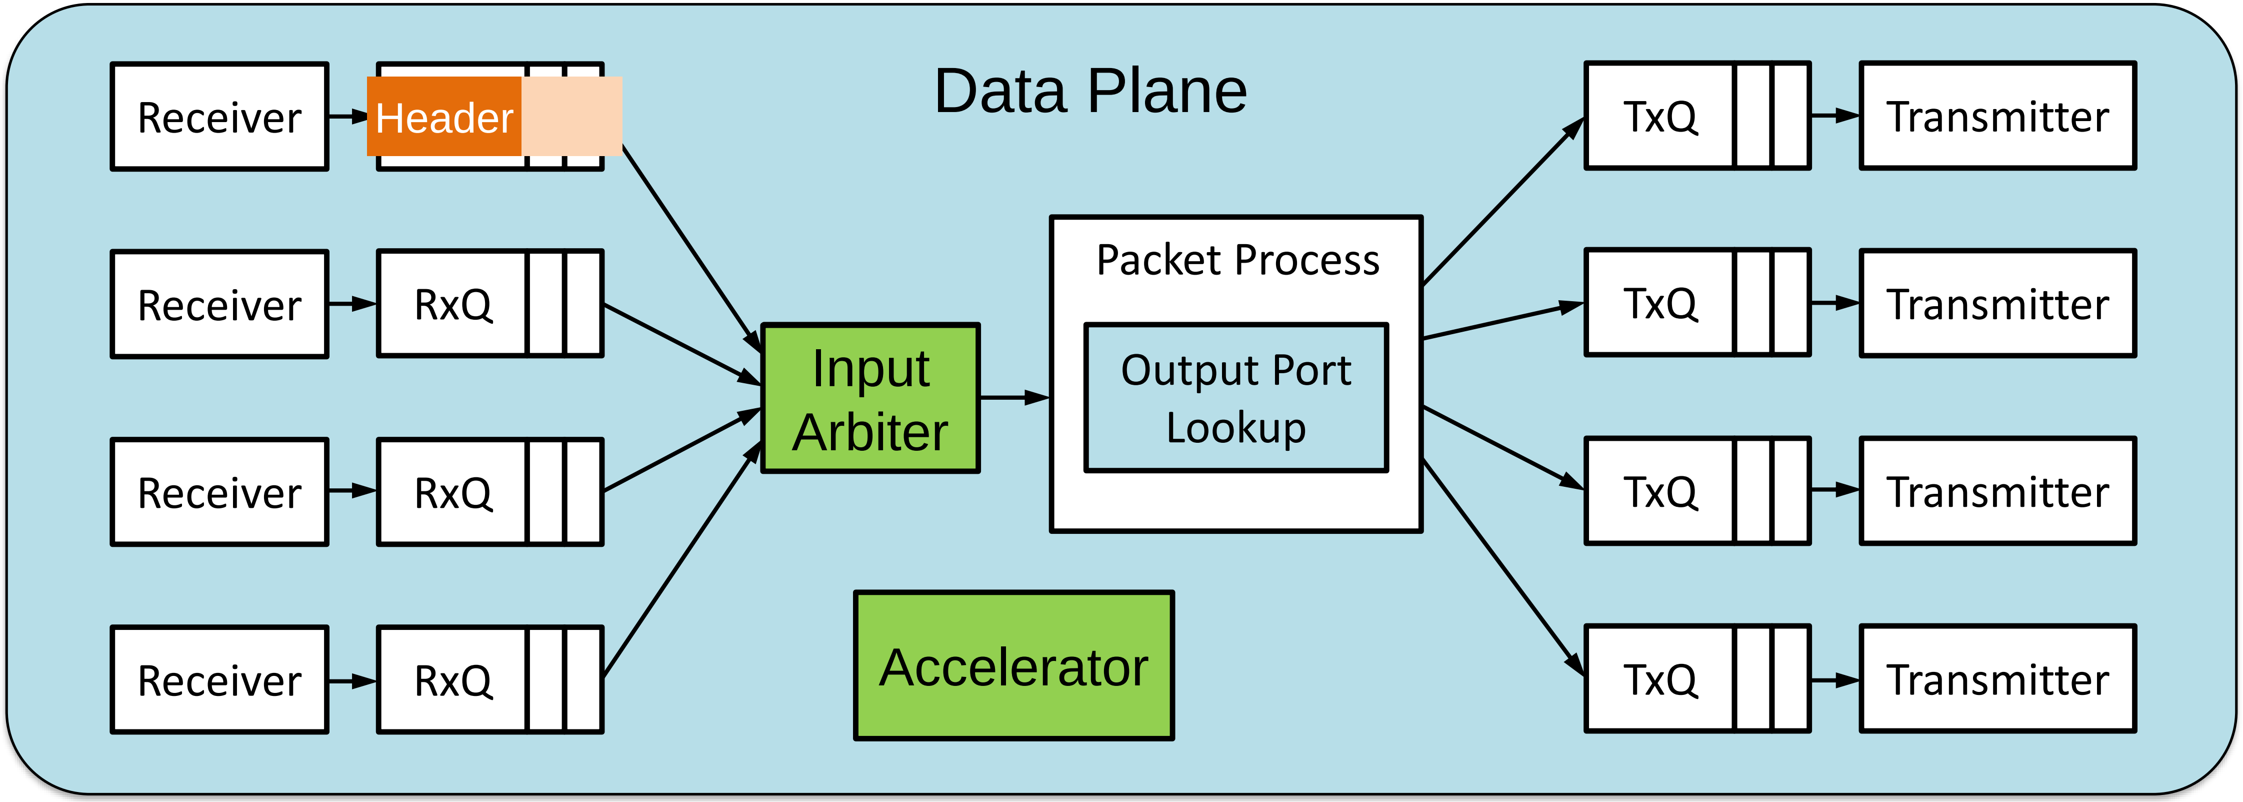

Integrating Aggregation Accelerator into the Programmable Switch

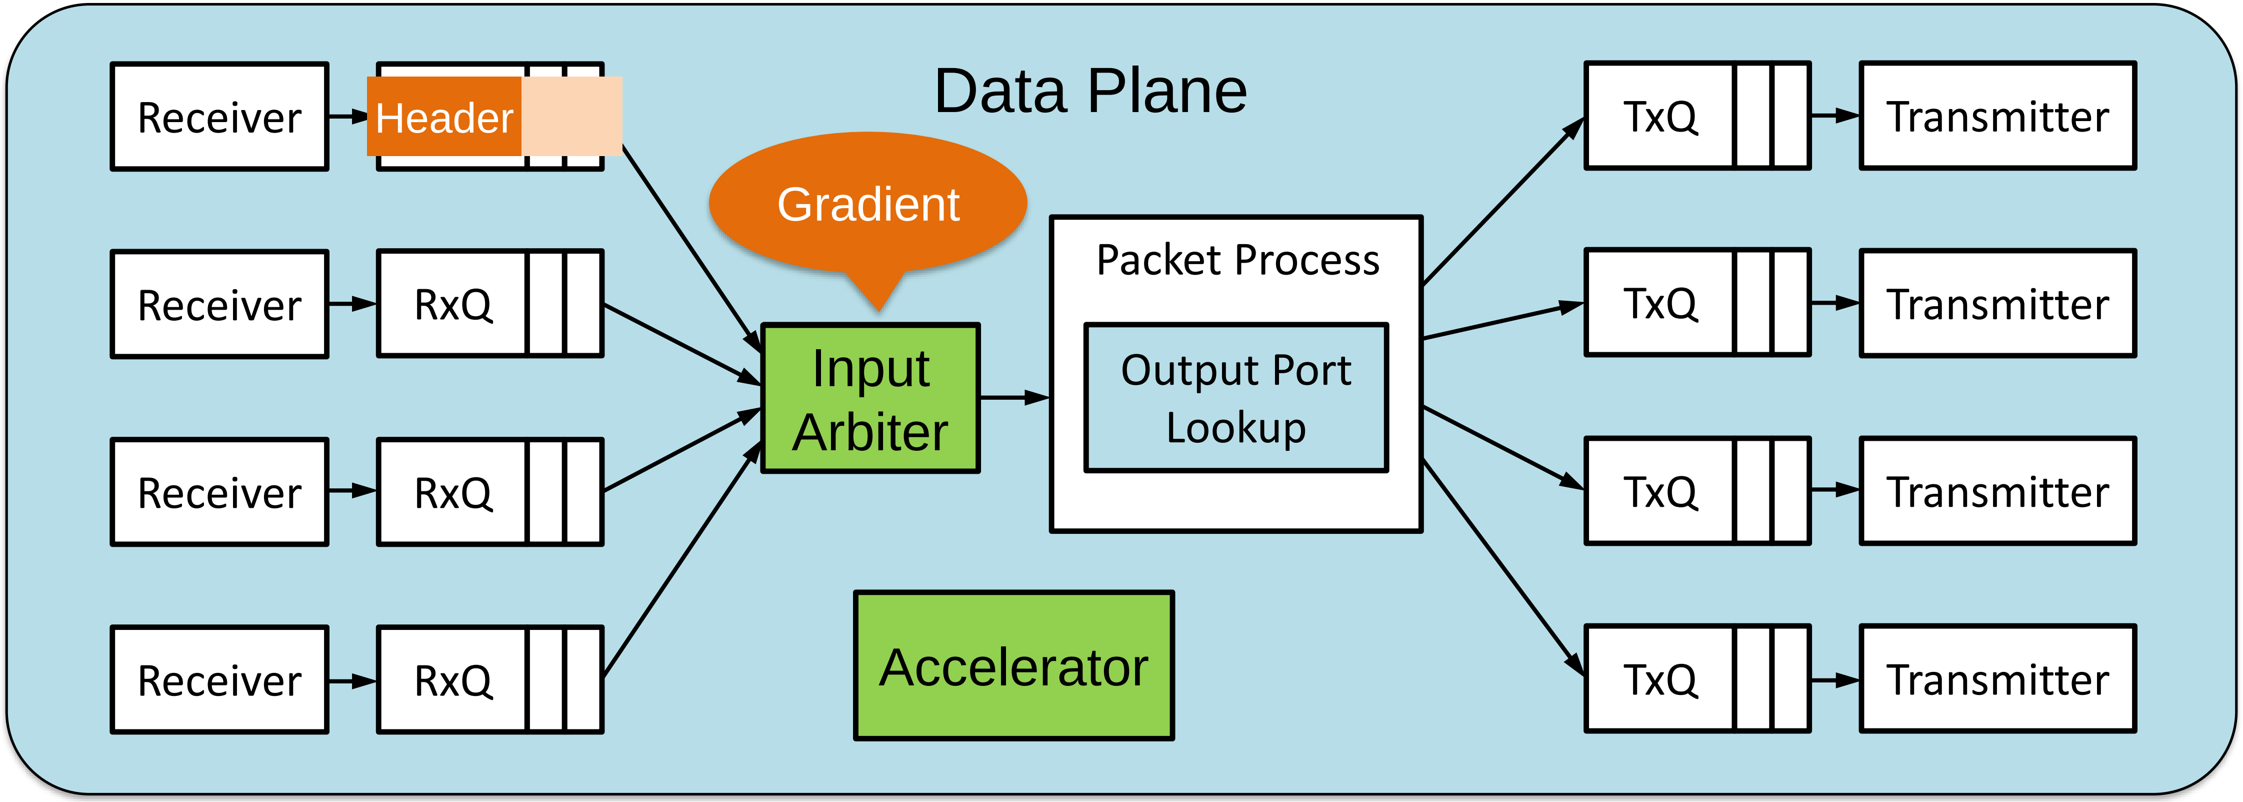

Integrating Aggregation Accelerator into the Programmable Switch

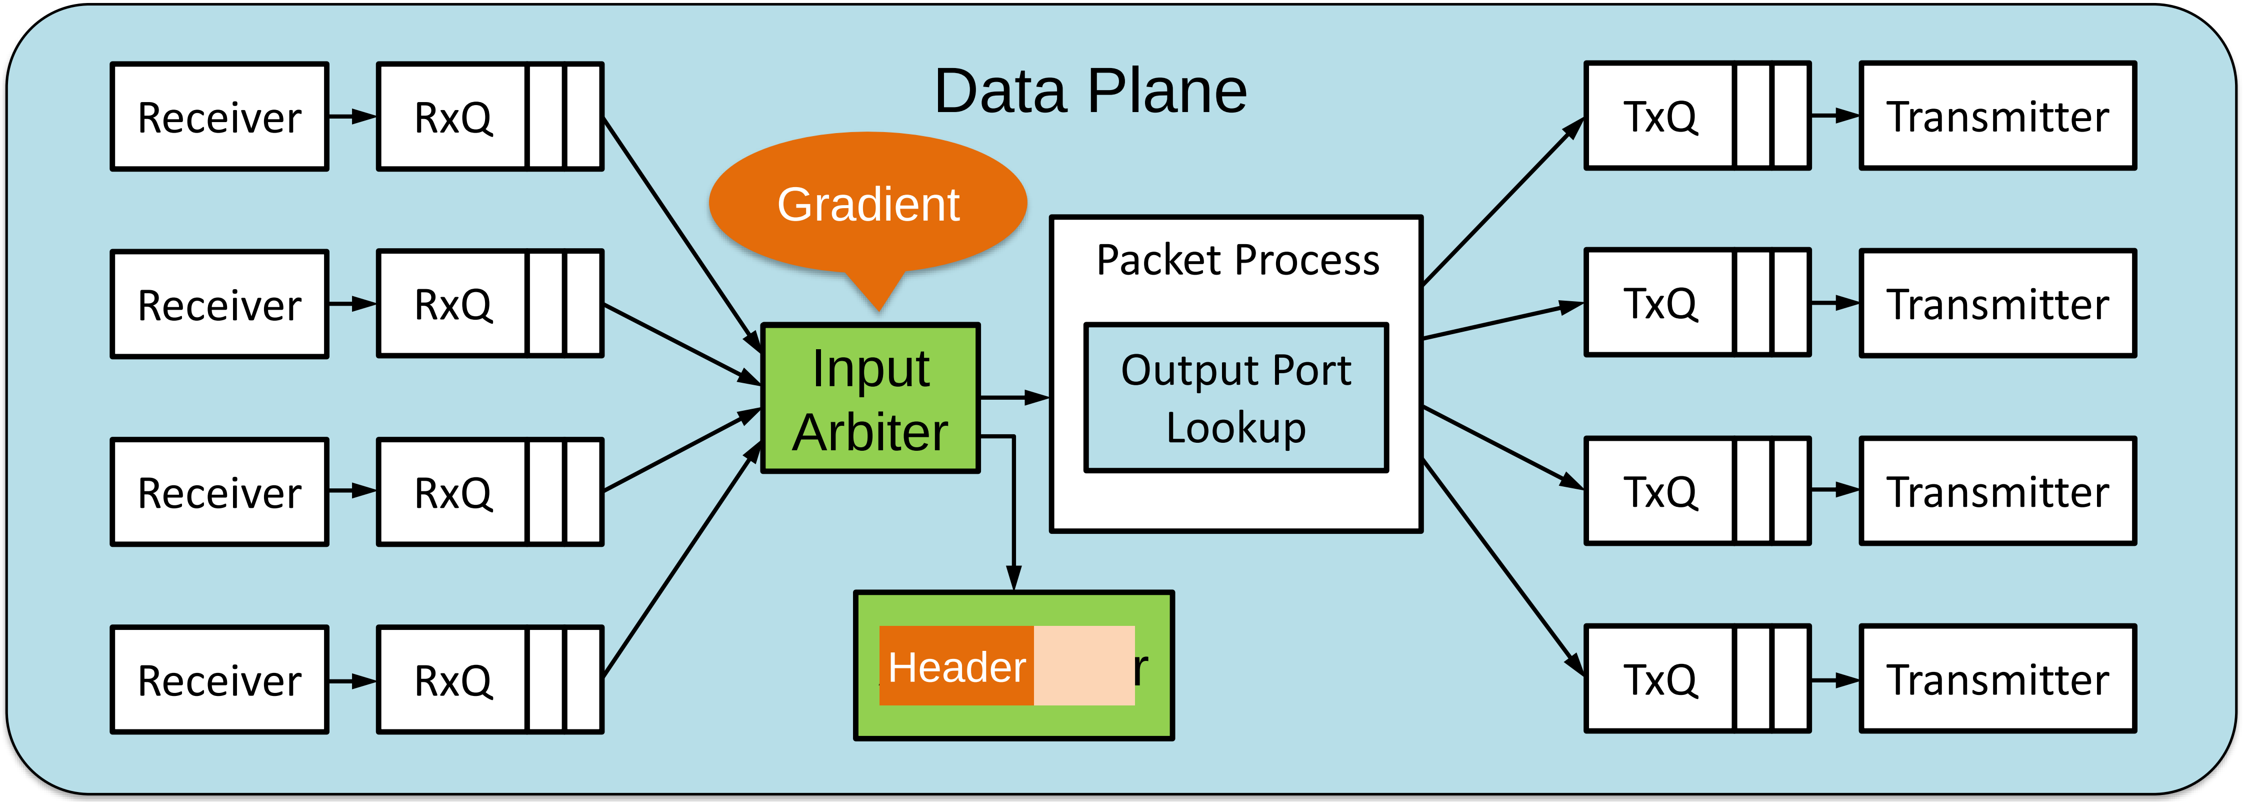

Integrating Aggregation Accelerator into the Programmable Switch

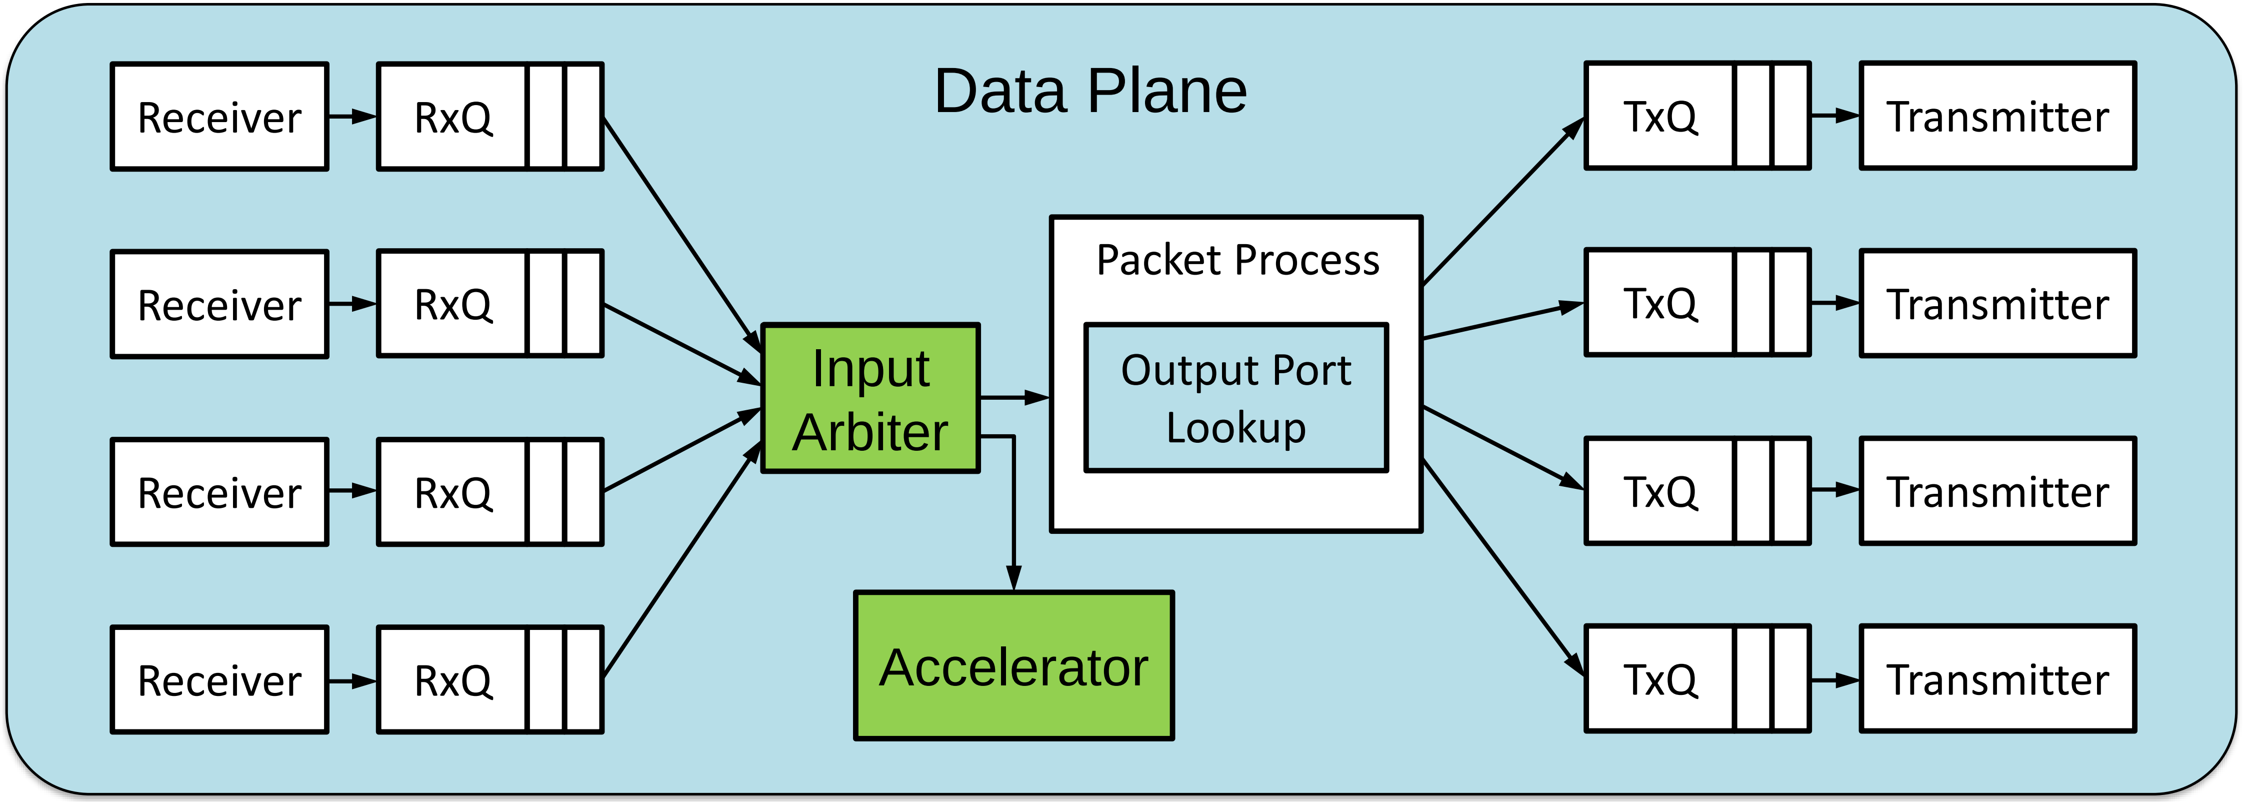

Integrating Aggregation Accelerator into the Programmable Switch

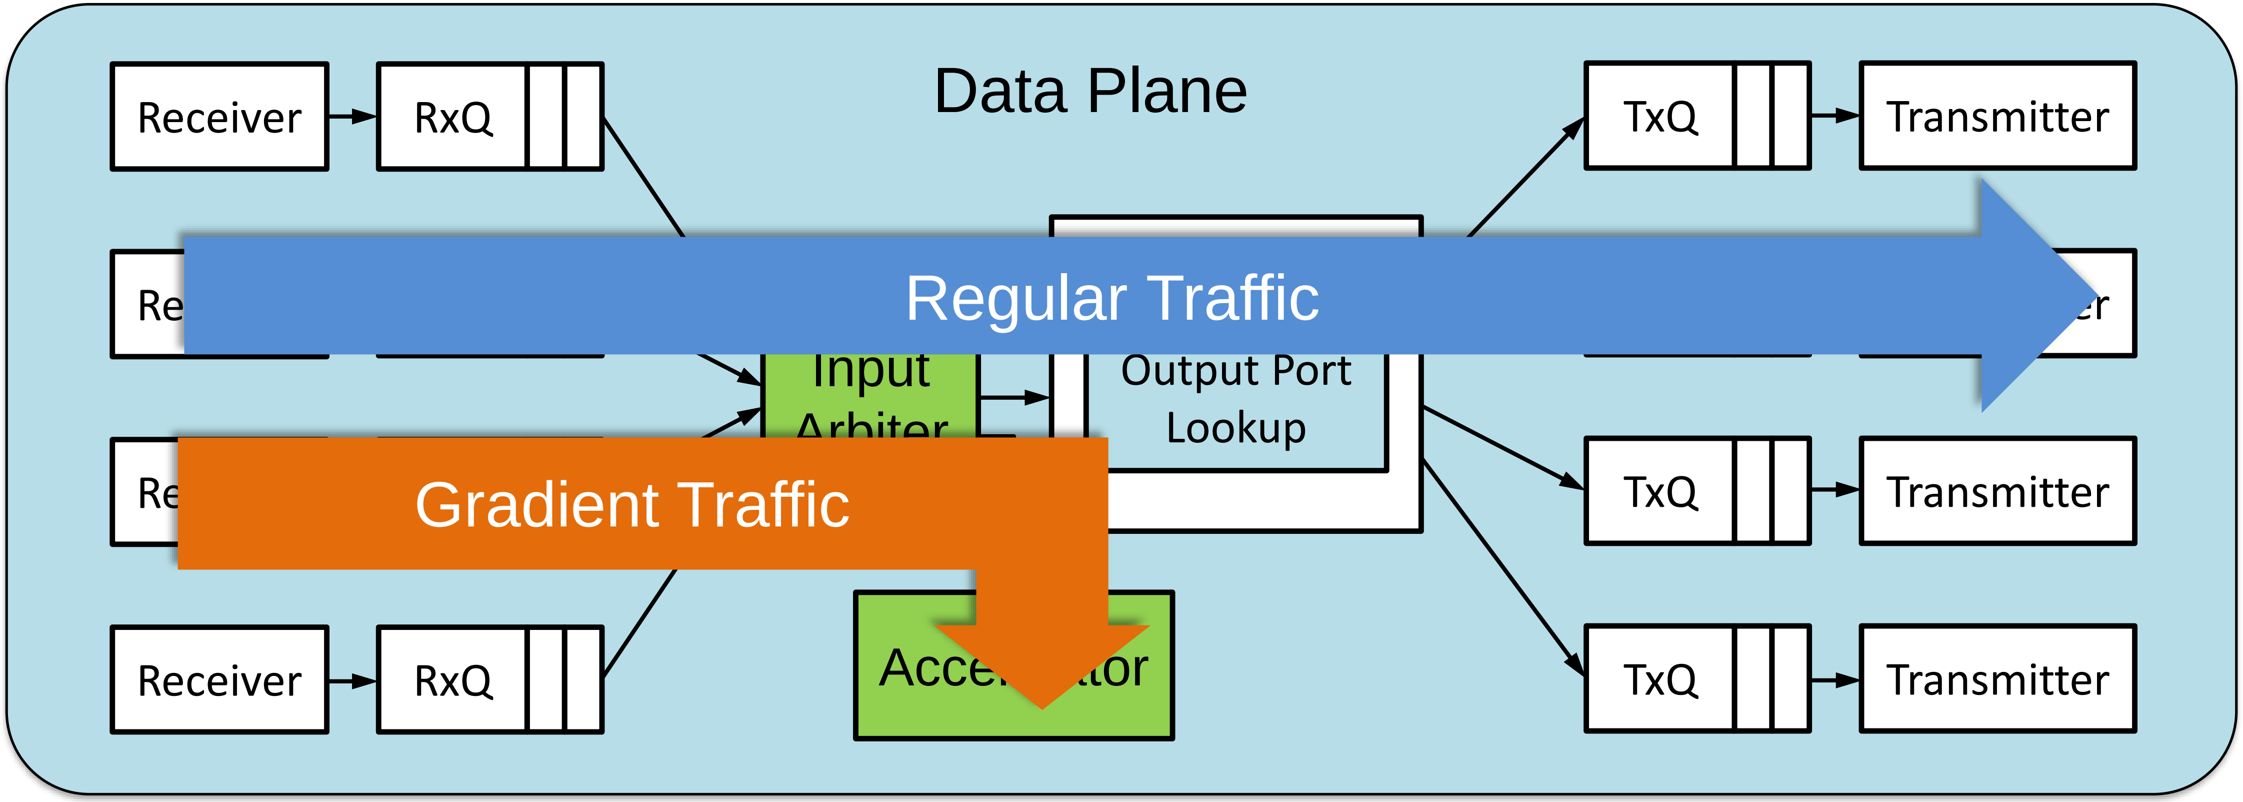

Integrating Aggregation Accelerator into the Programmable Switch

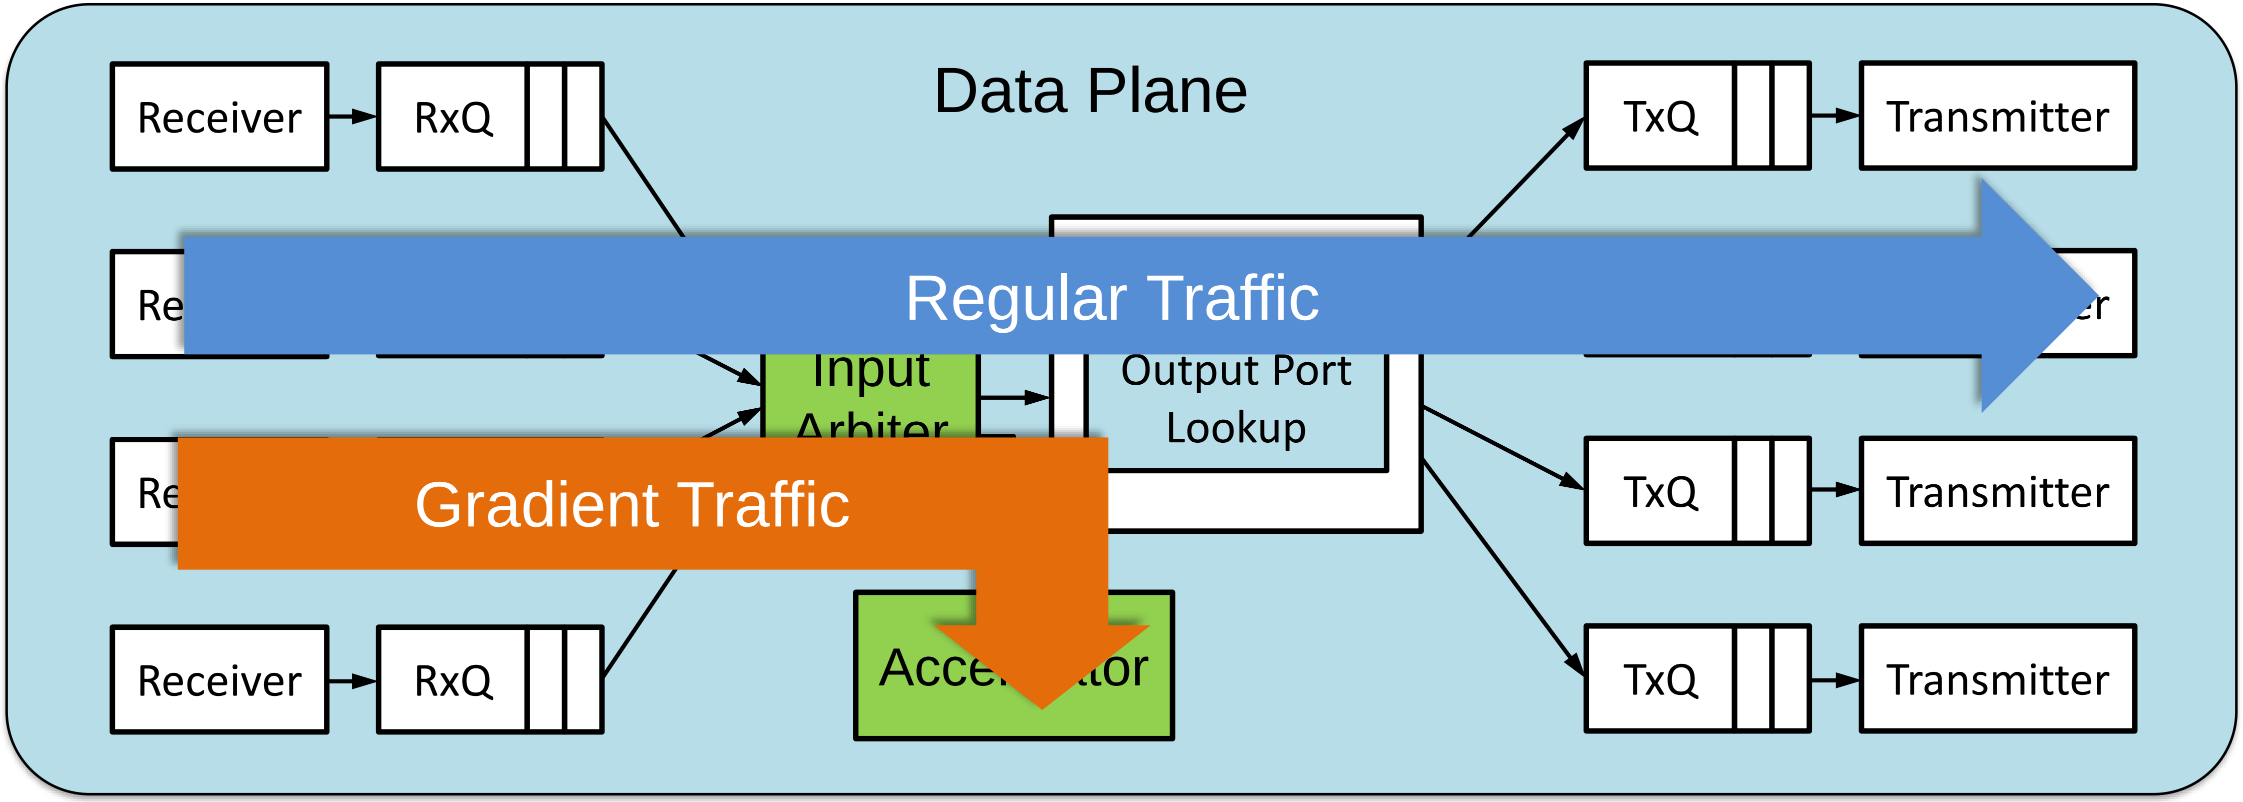

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

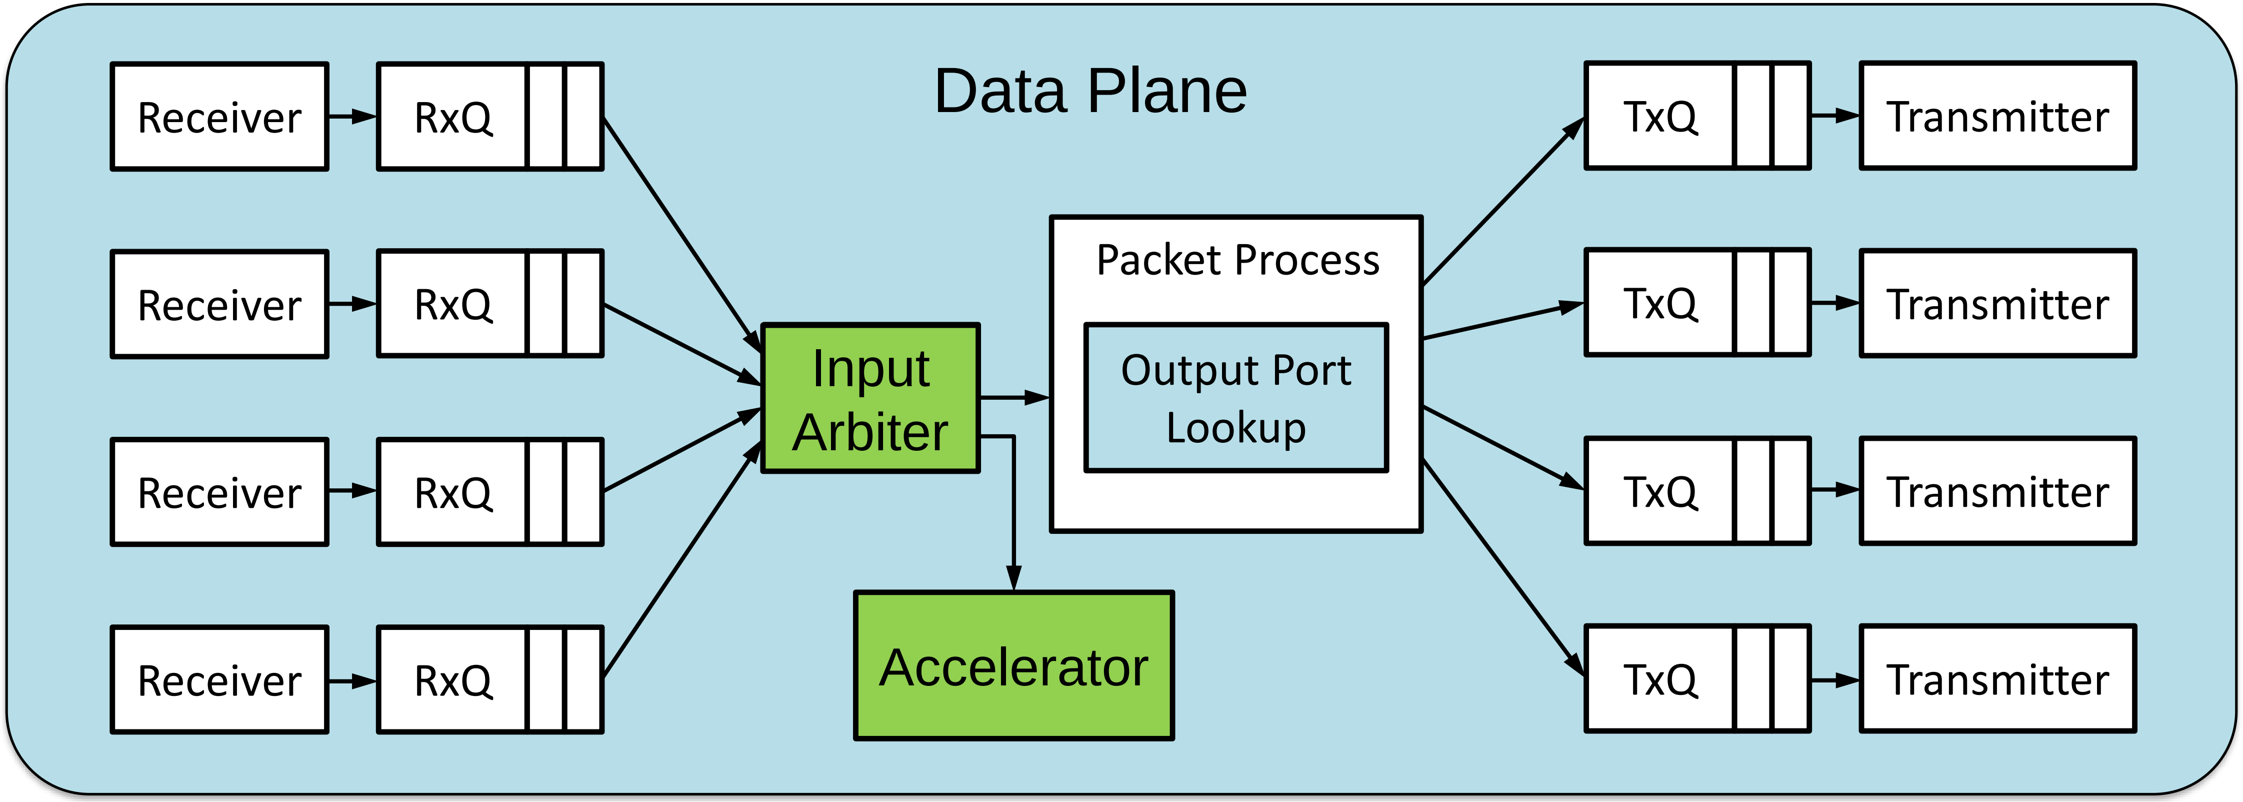

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

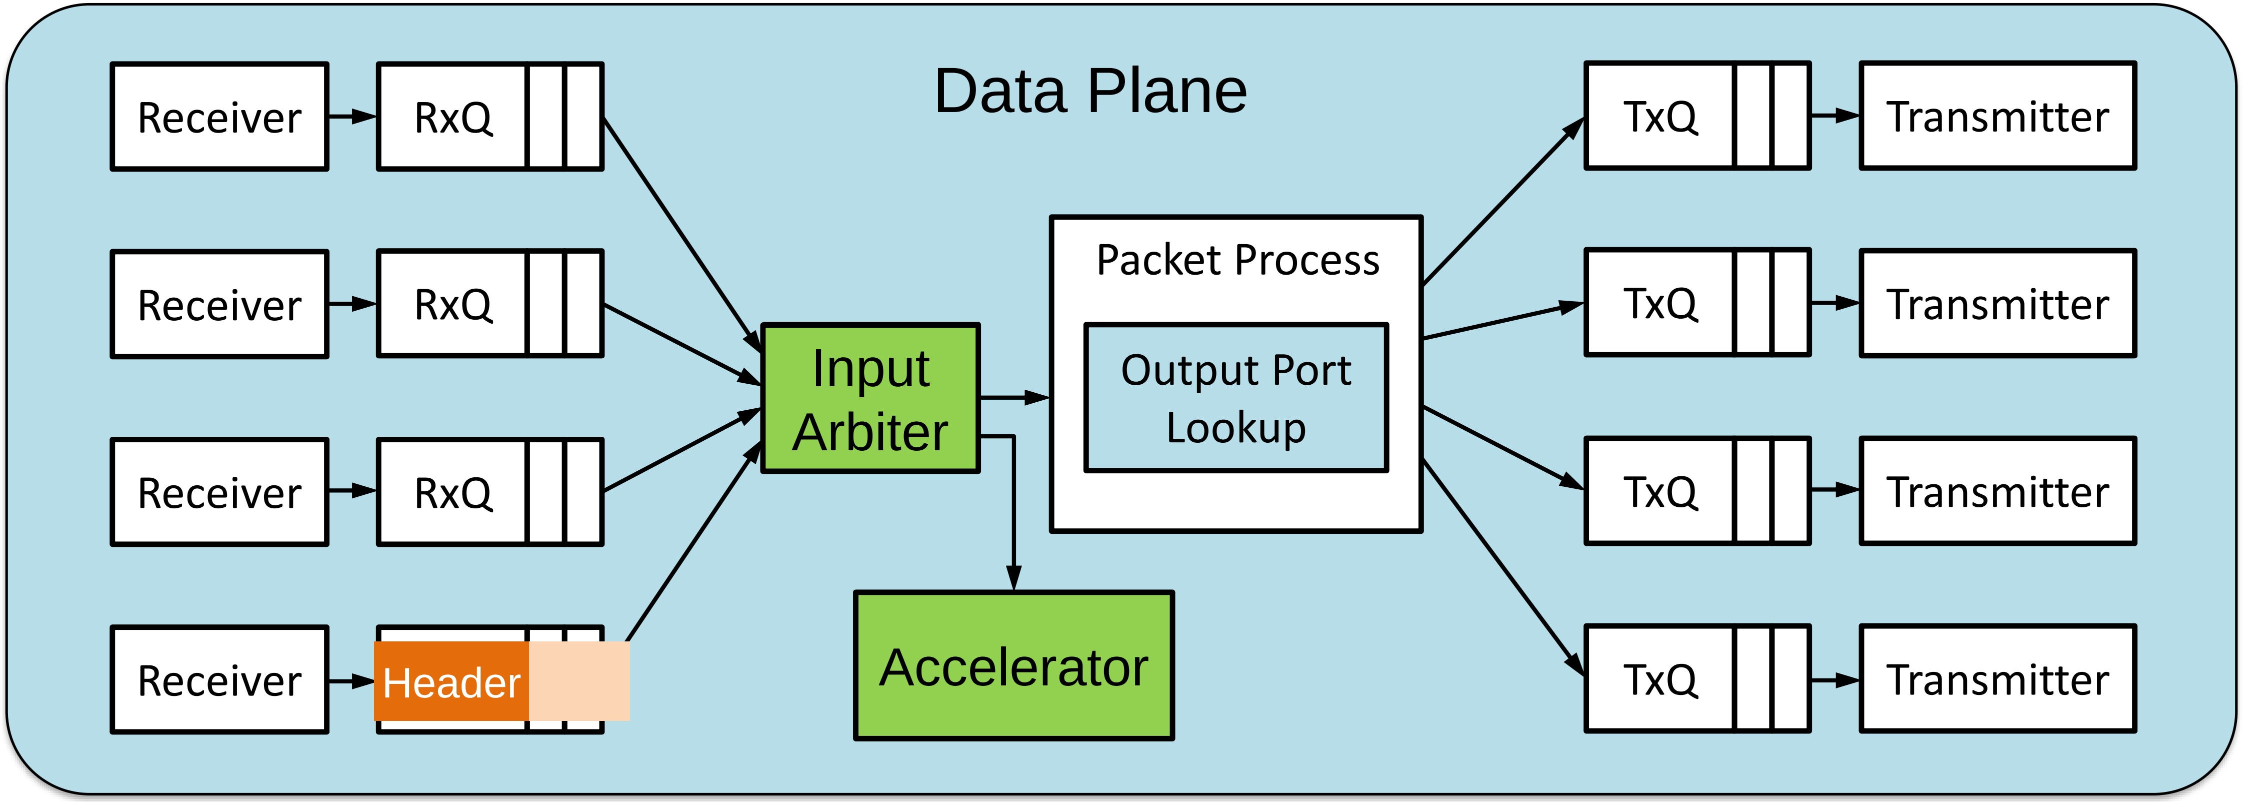

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

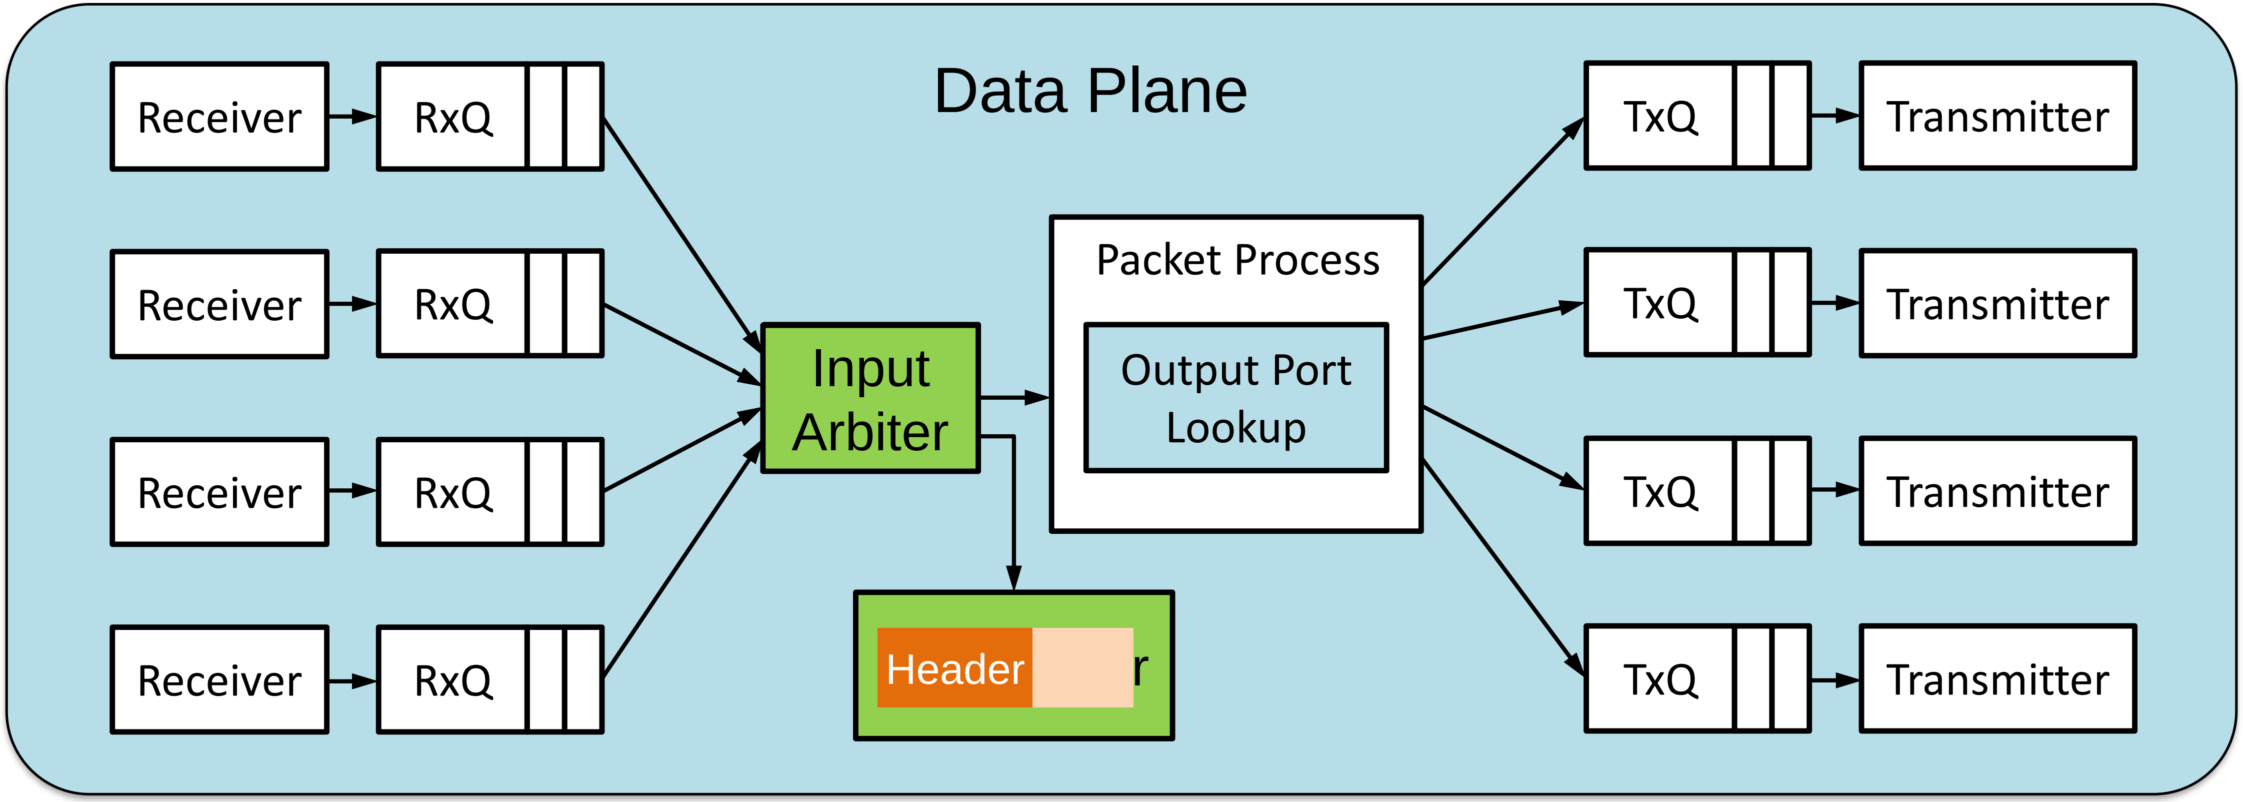

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

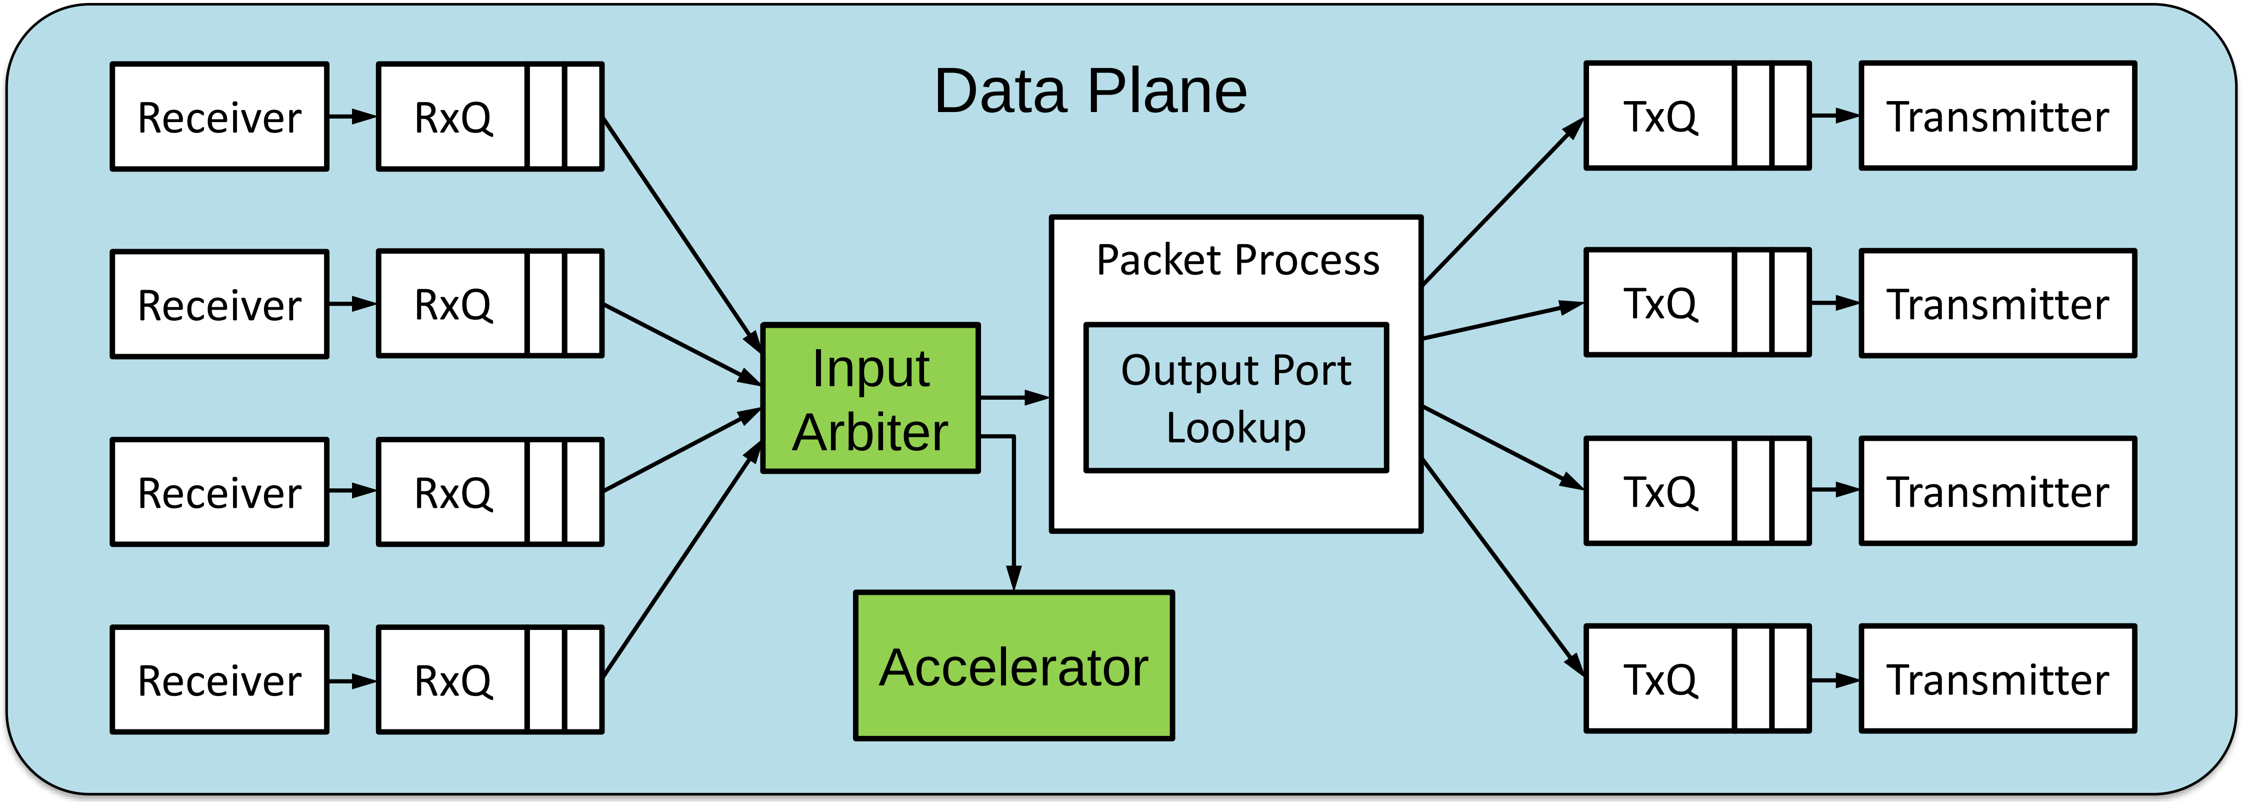

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

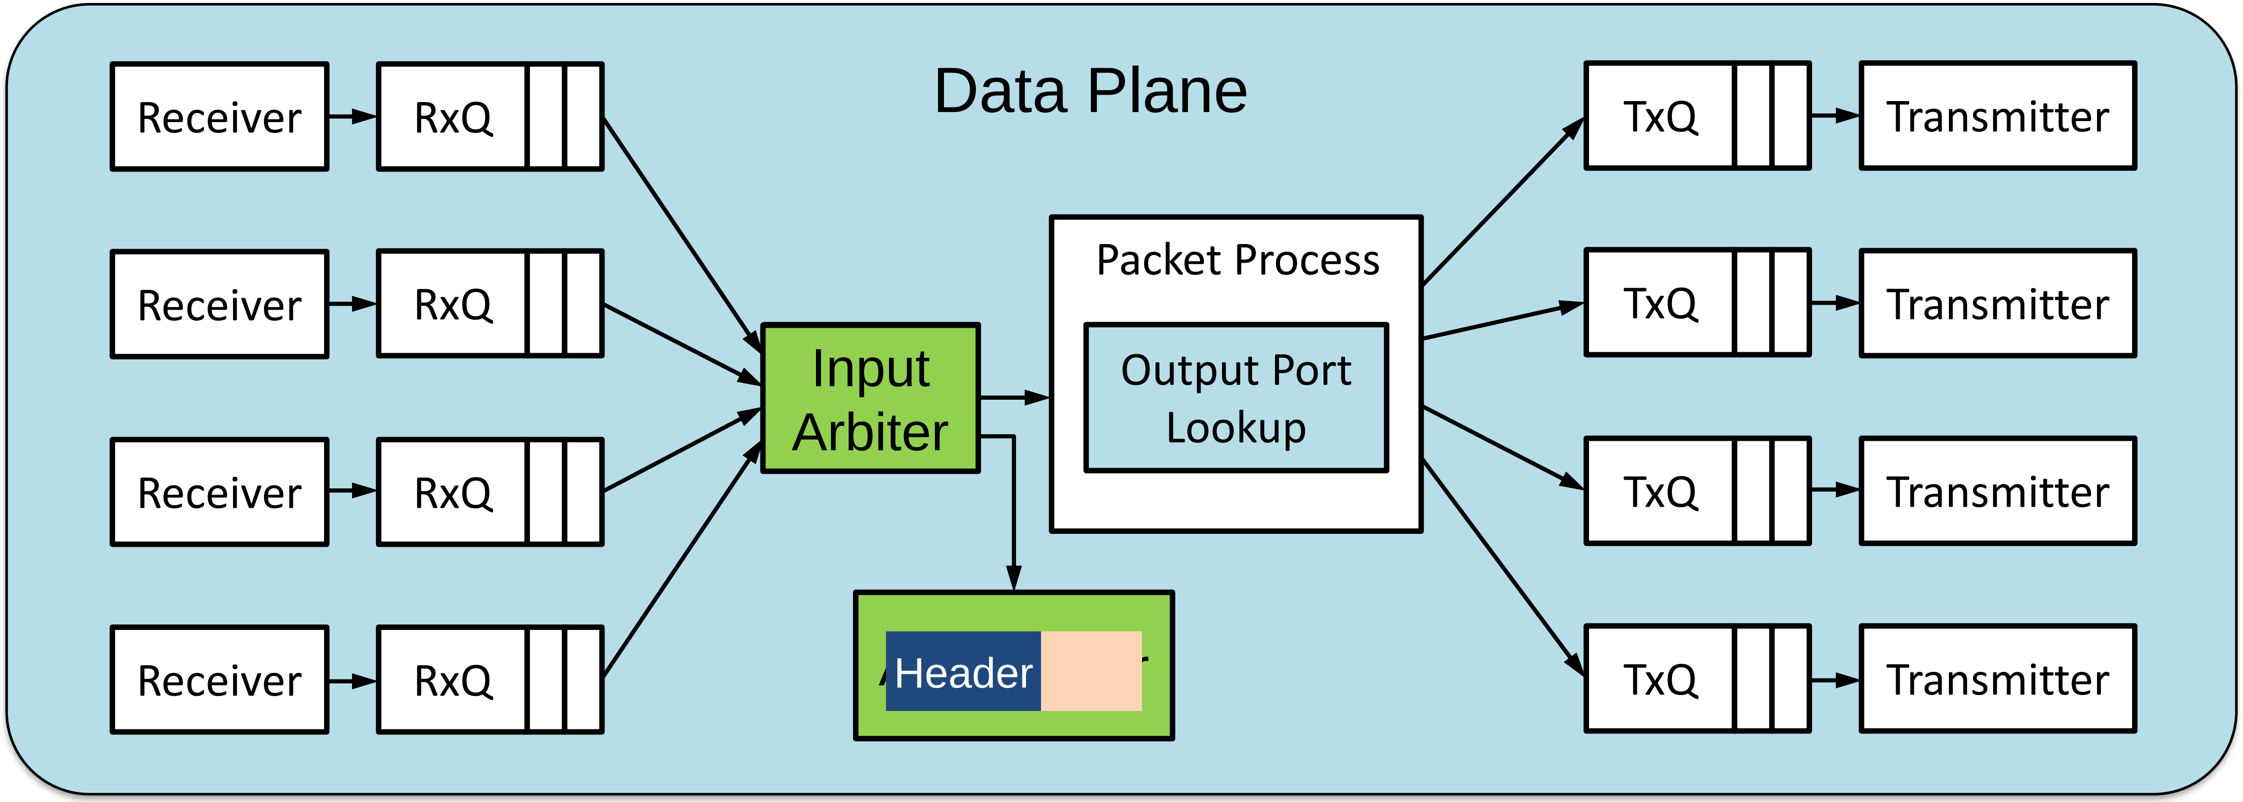

Hardware Acceleration Isolated From Regular Switch Function

Integrating Aggregation Accelerator into the Programmable Switch

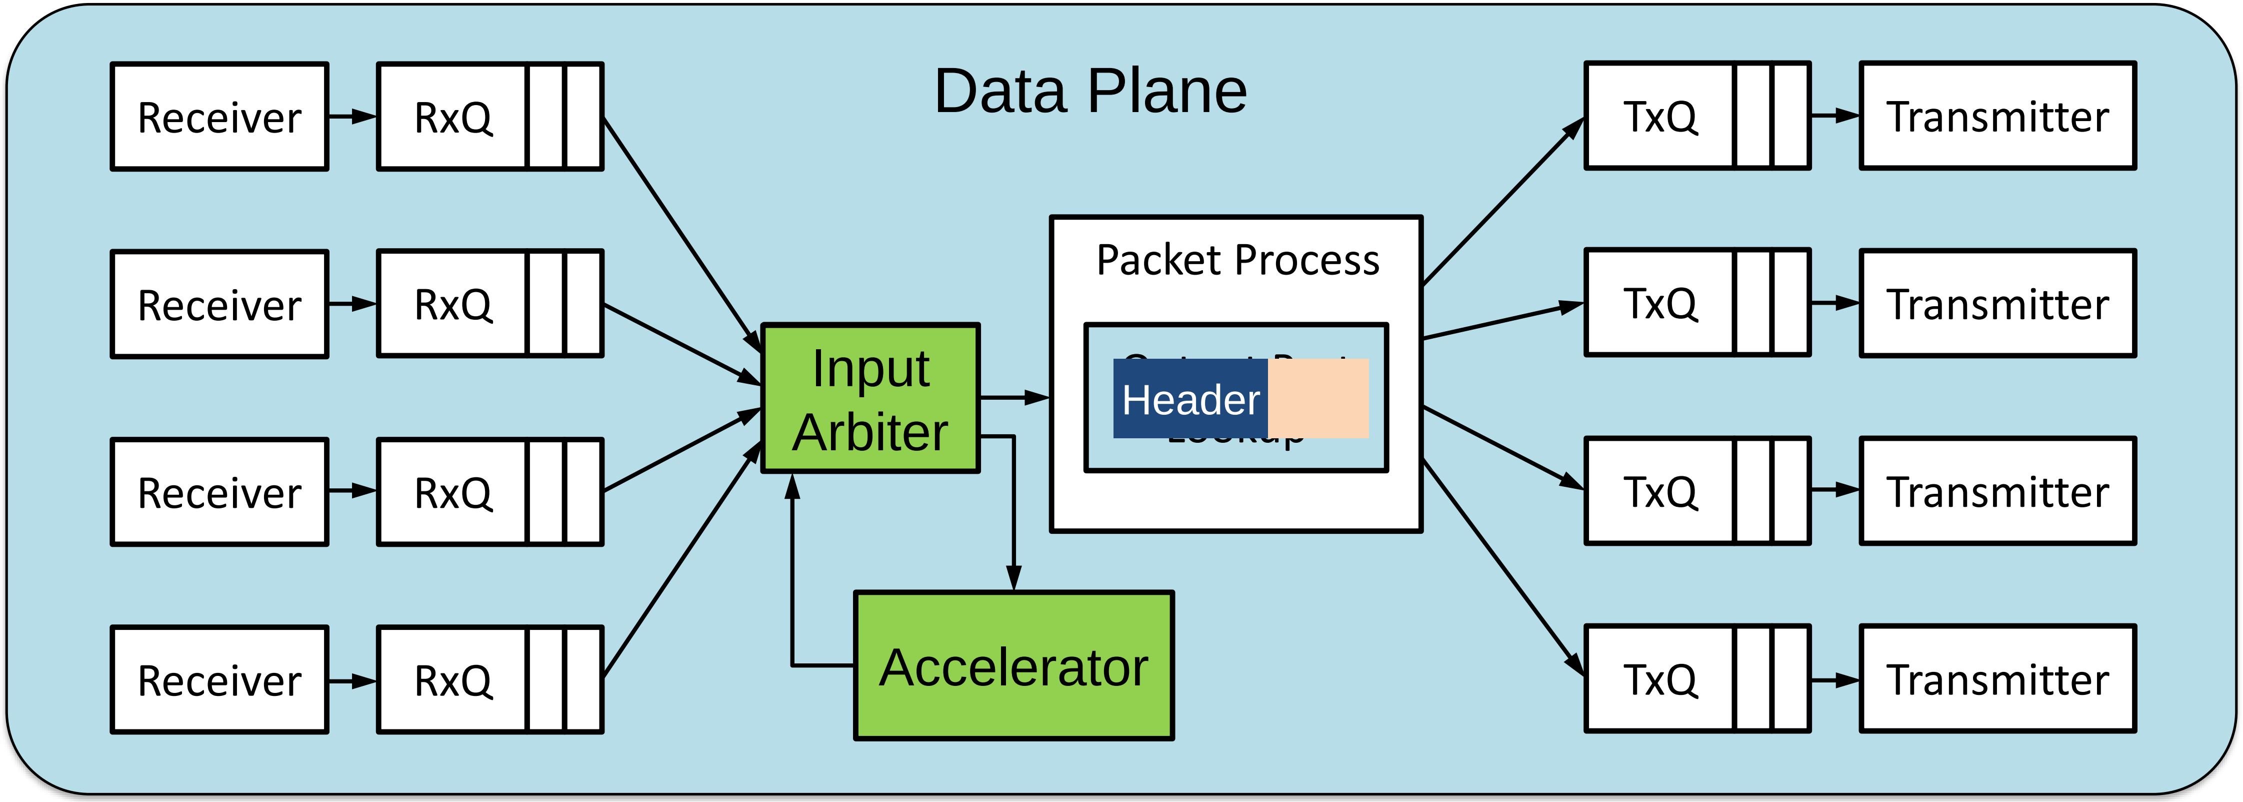

Hardware Acceleration Isolated From Regular Switch Function

Developing Light-Weight Accelerator for Aggregation

In-Switch
Accelerator

Developing Light-Weight Accelerator for Aggregation

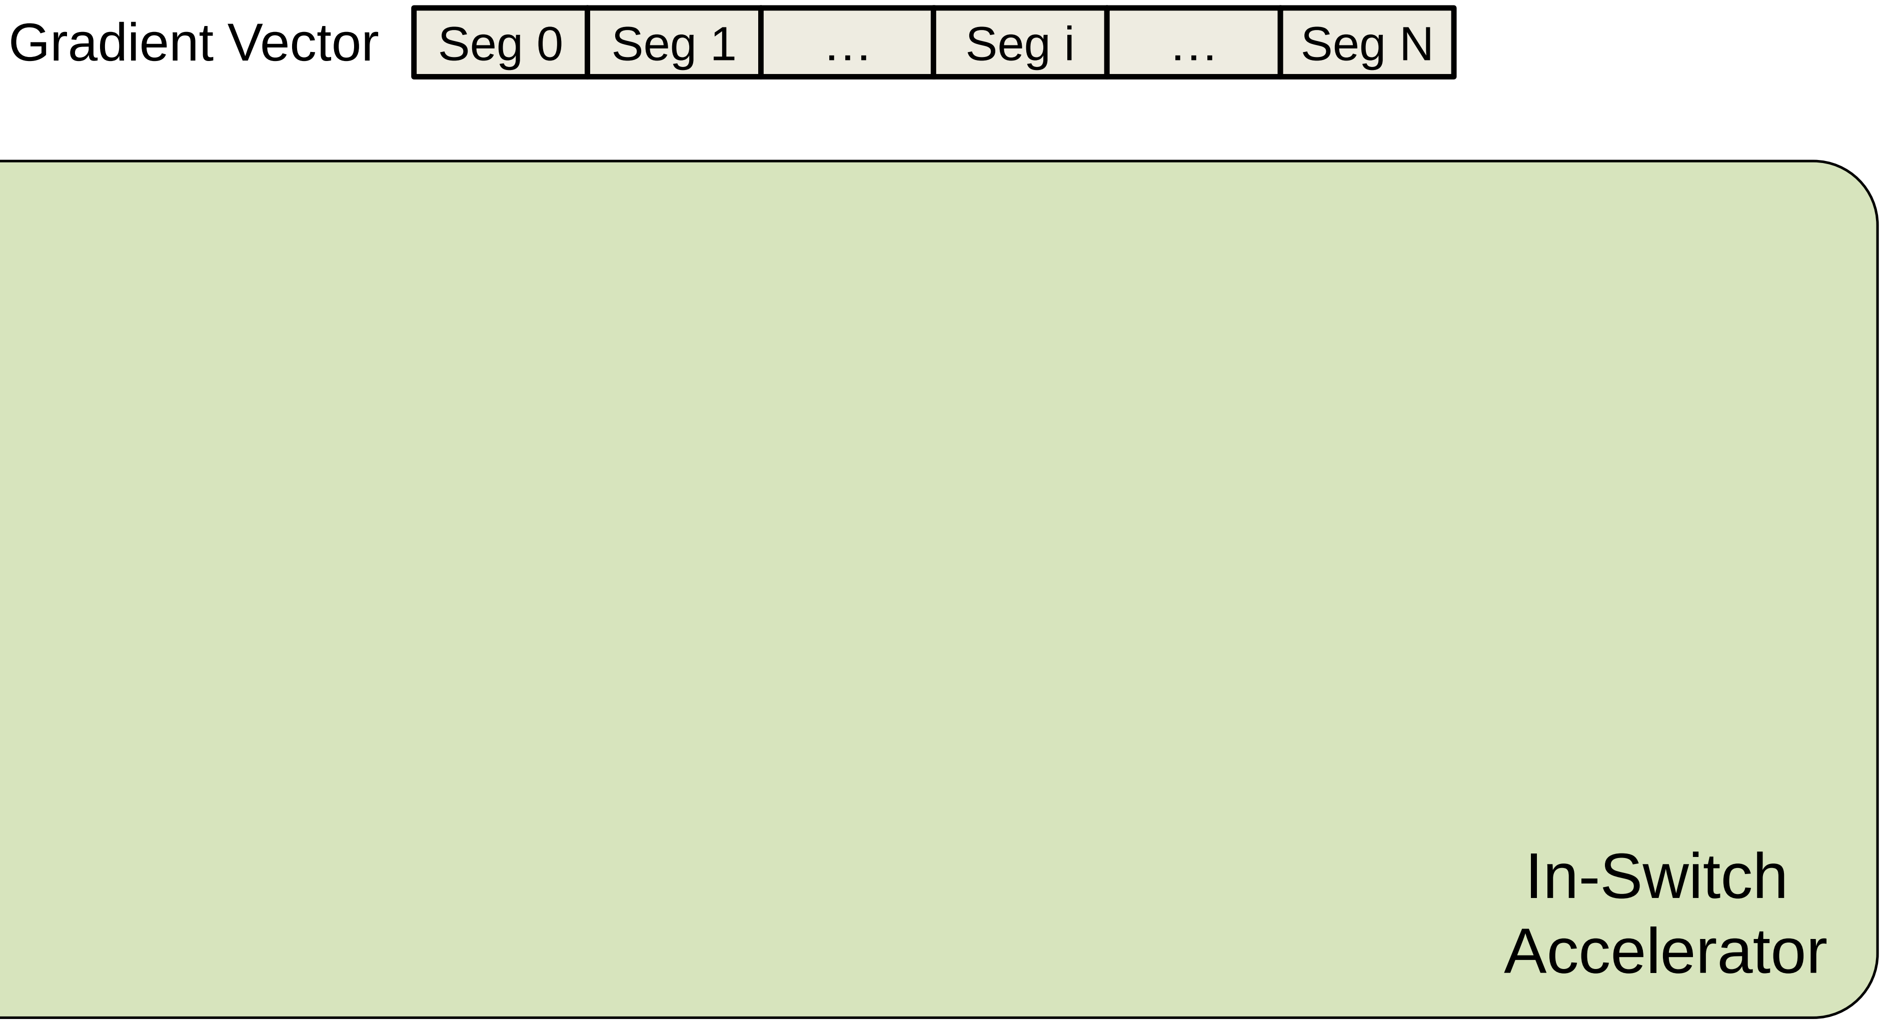

Developing Light-Weight Accelerator for Aggregation

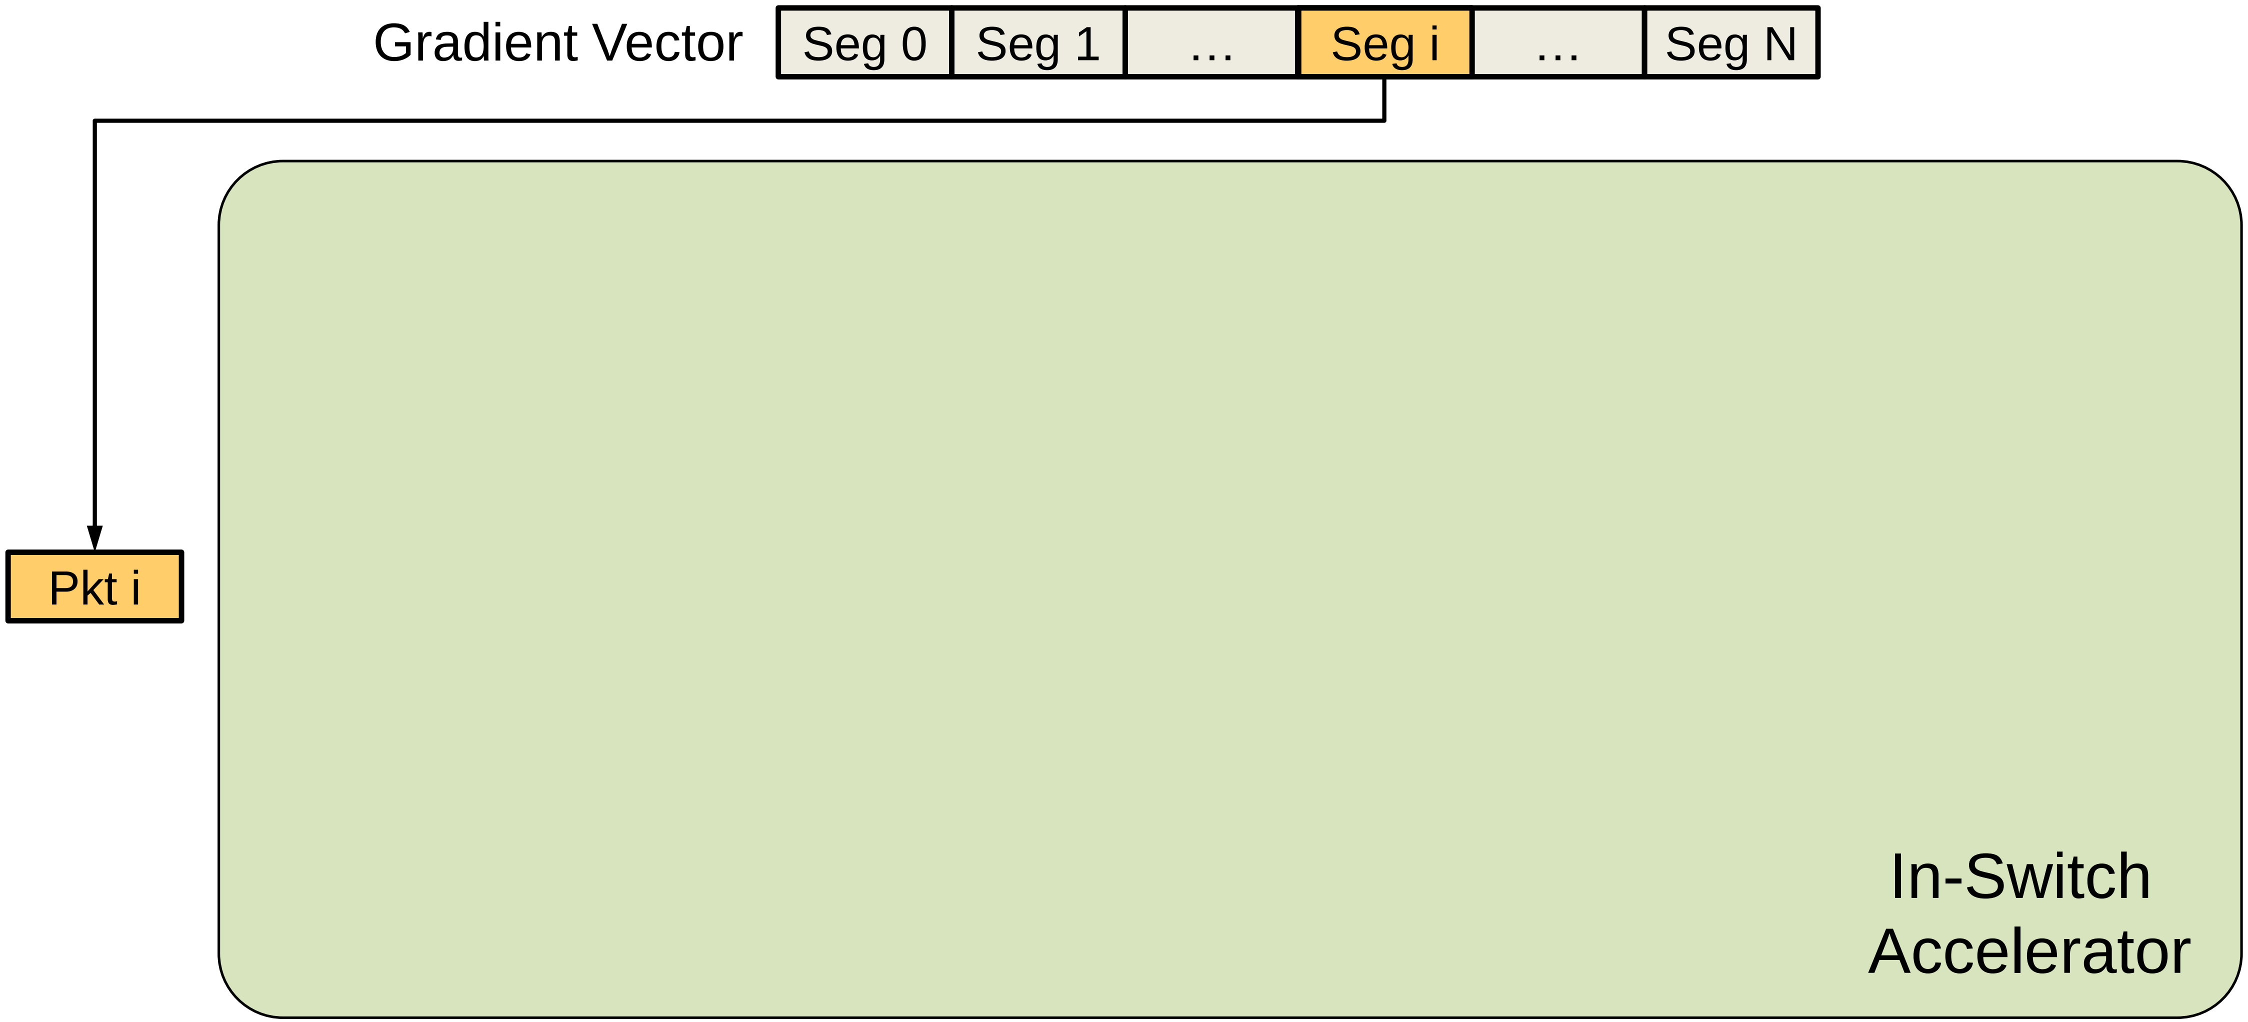

Developing Light-Weight Accelerator for Aggregation

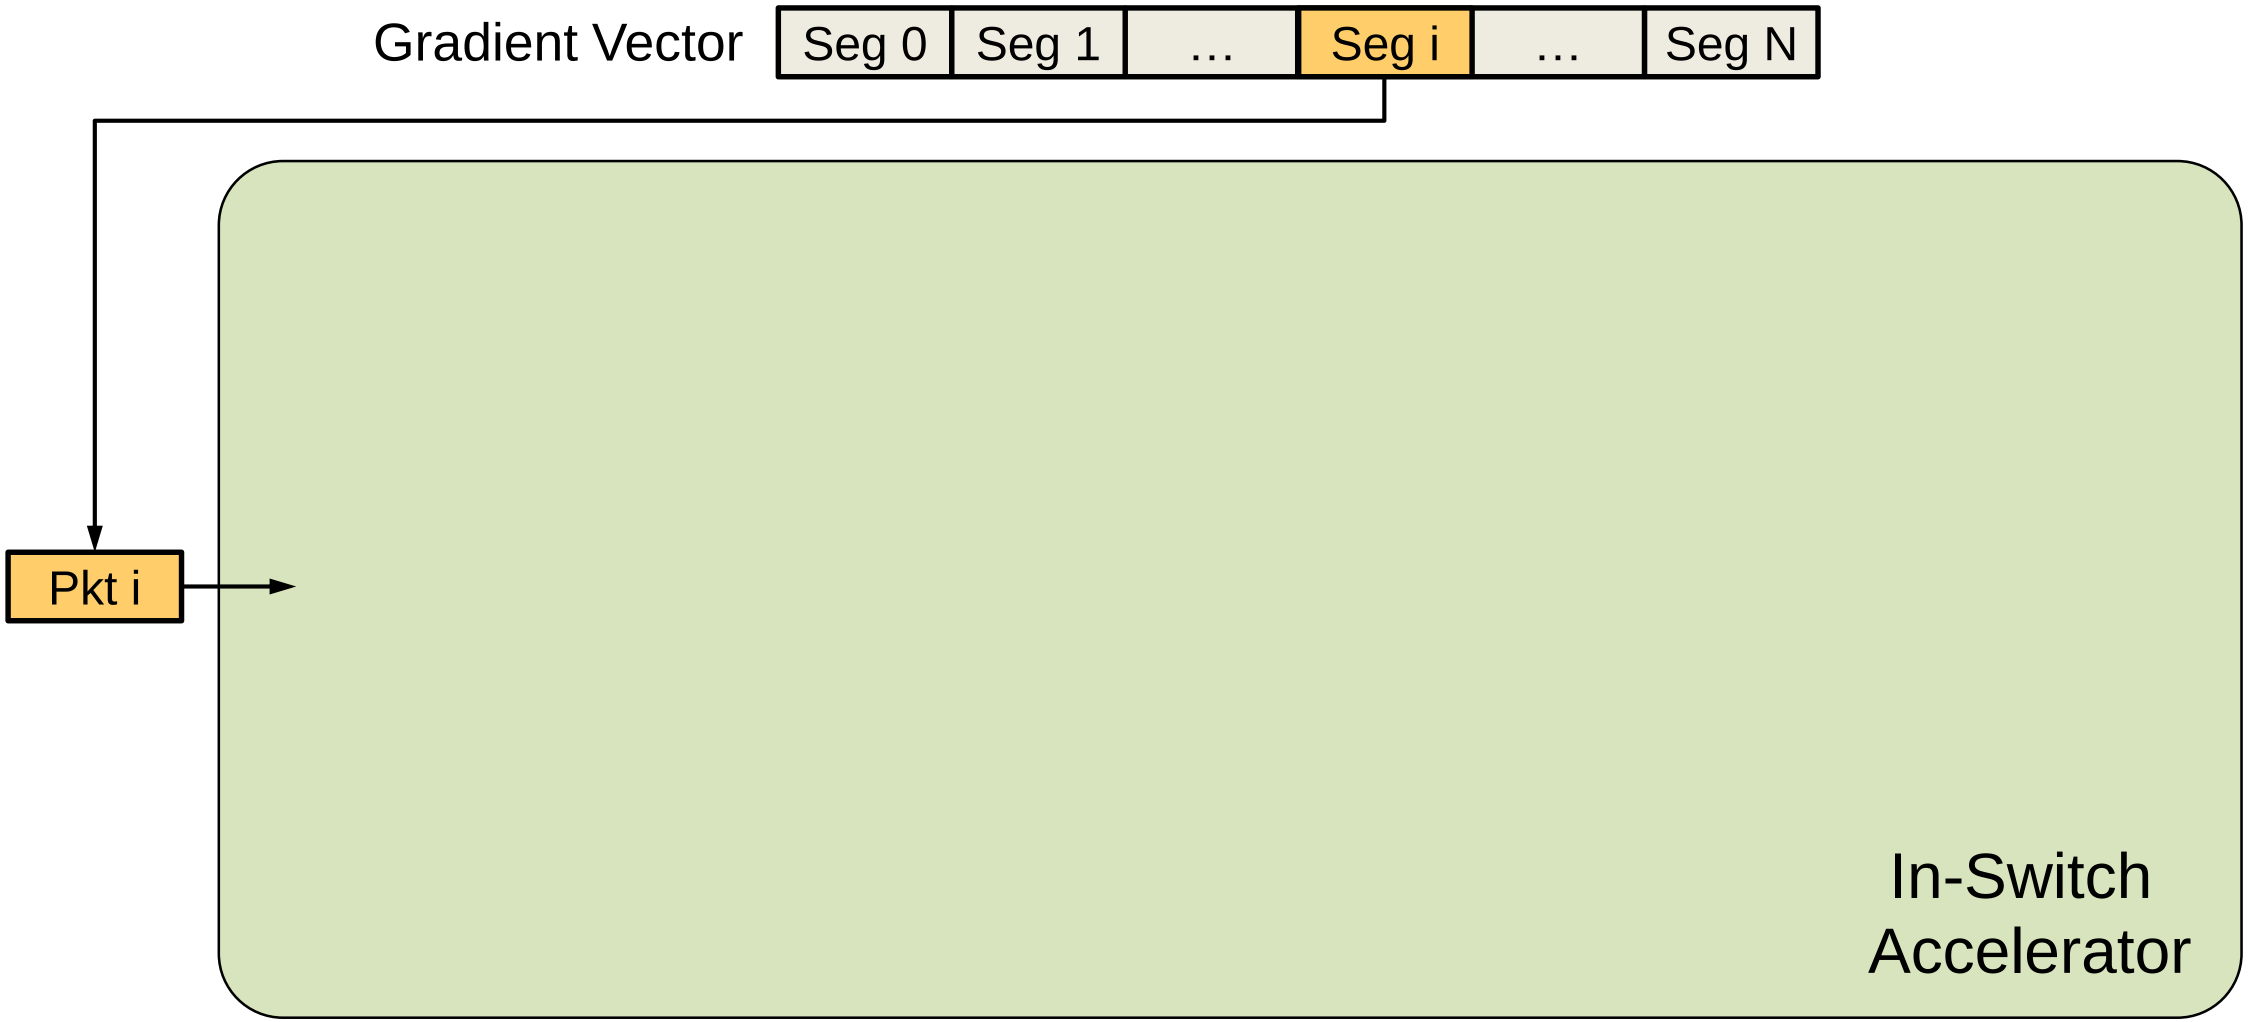

Developing Light-Weight Accelerator for Aggregation

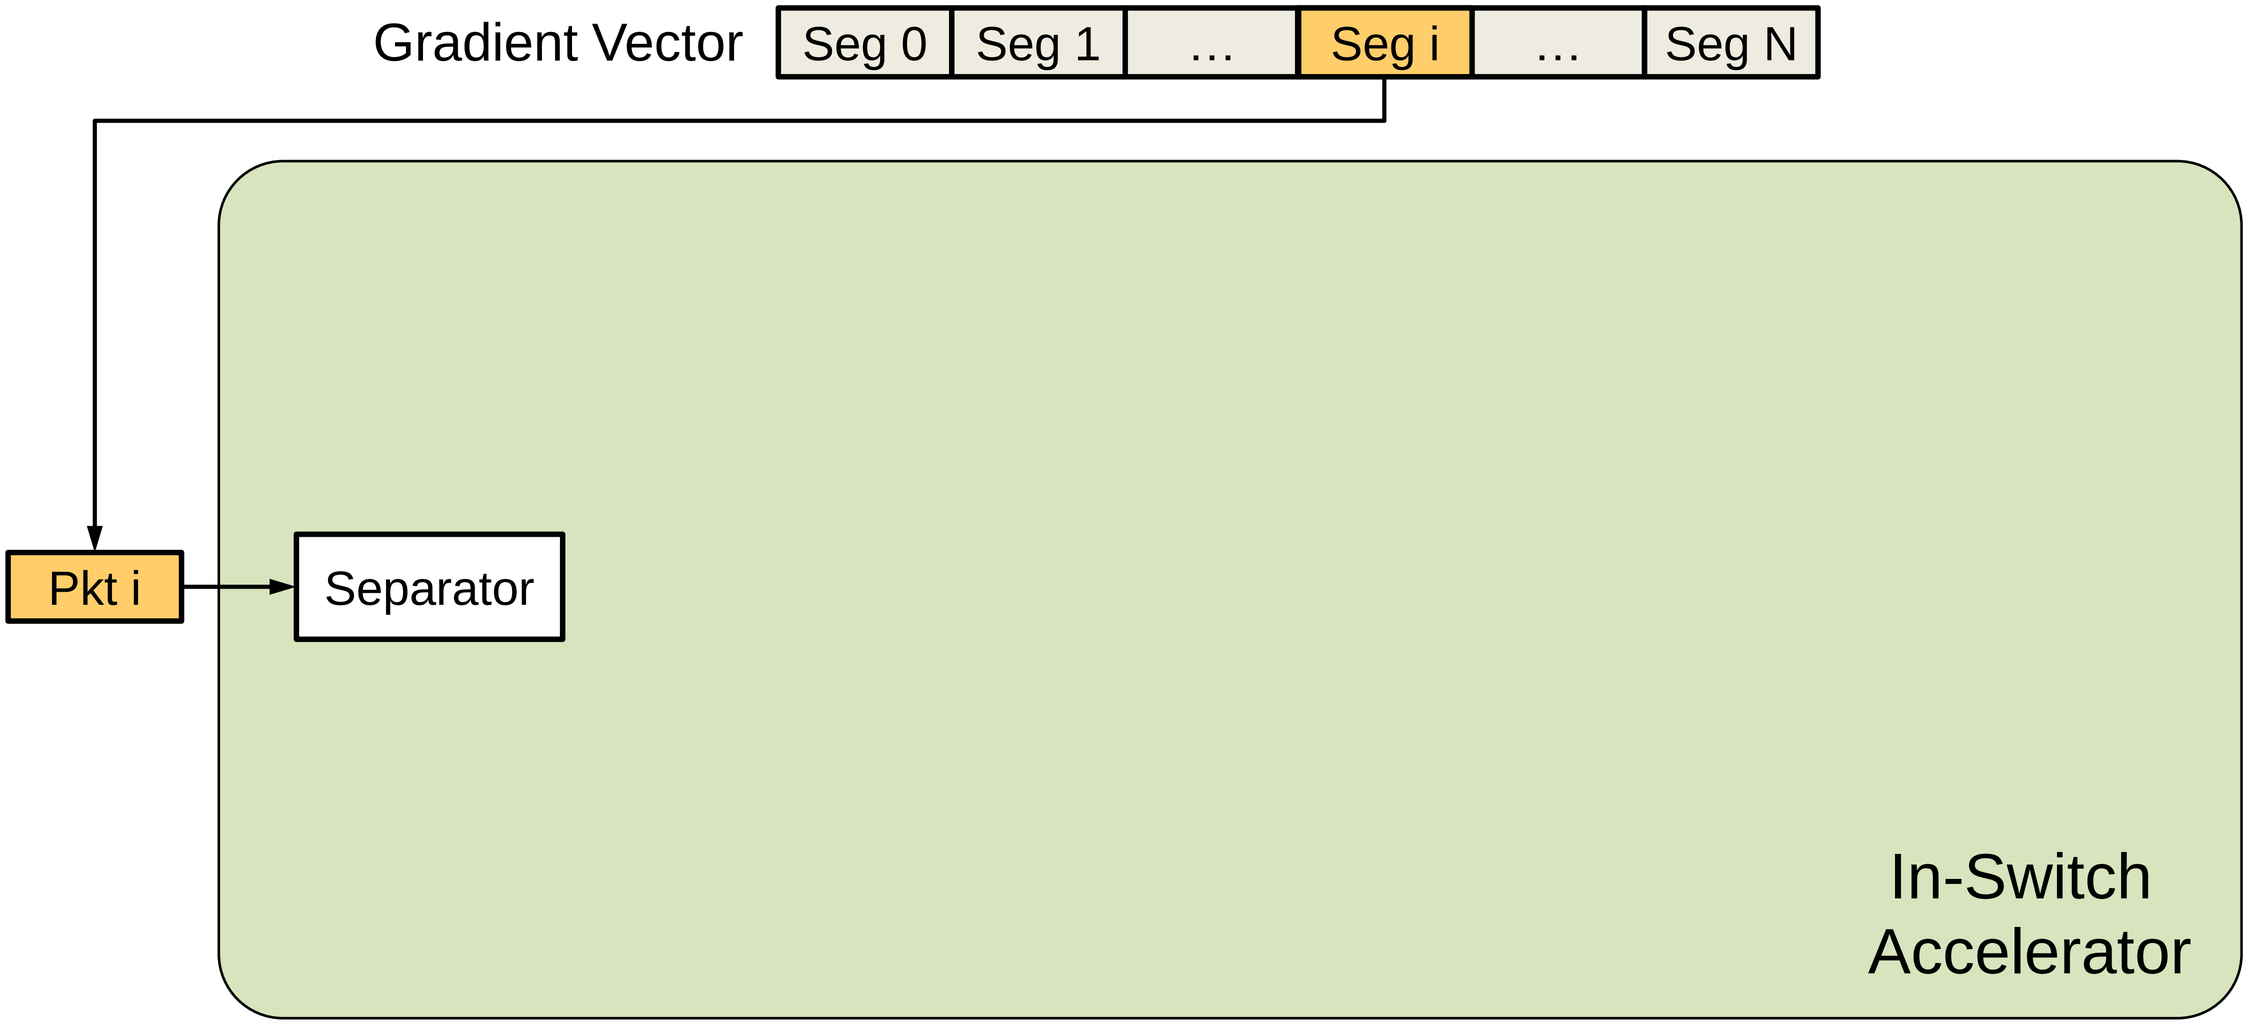

Developing Light-Weight Accelerator for Aggregation

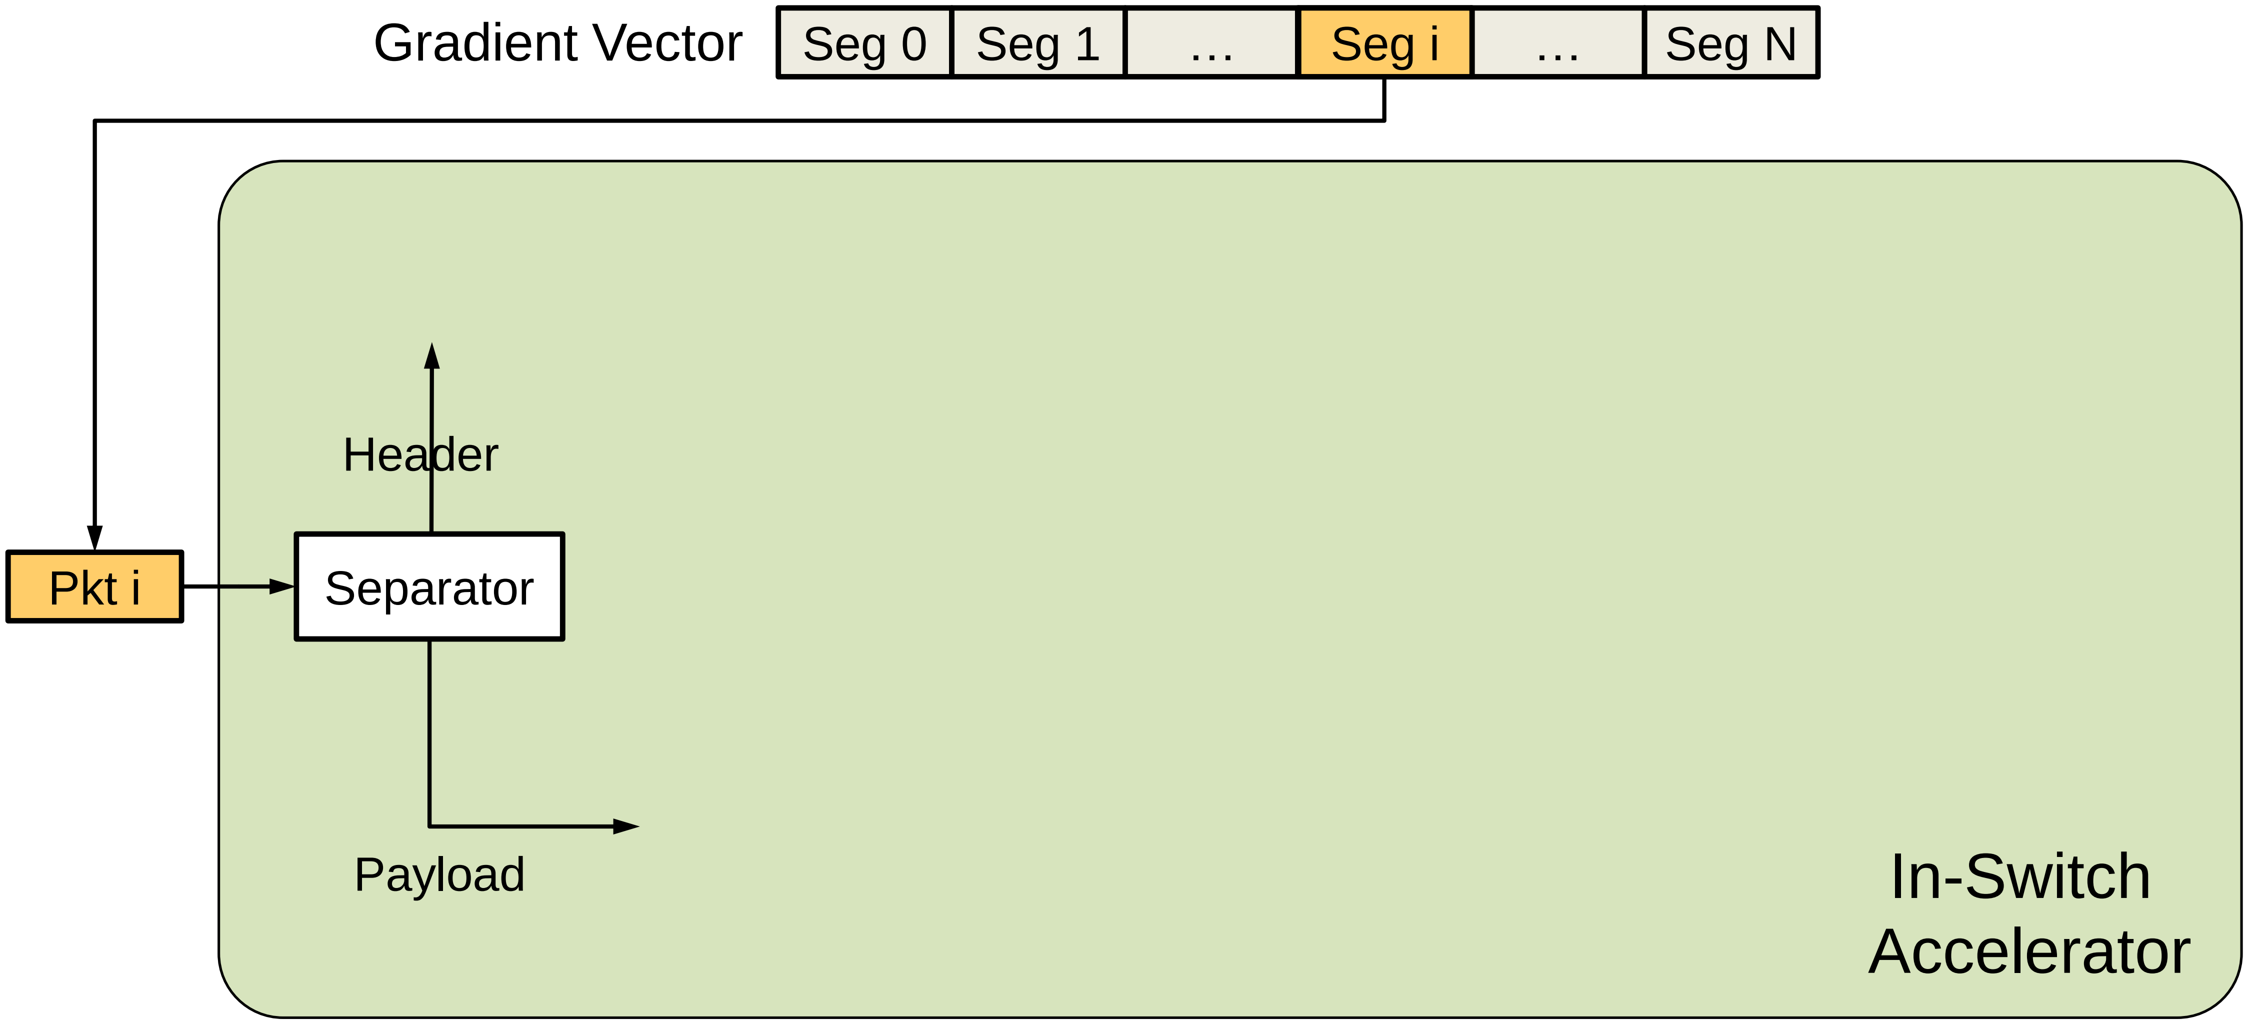

ECE ILLINOIS

Developing Light-Weight Accelerator for Aggregation

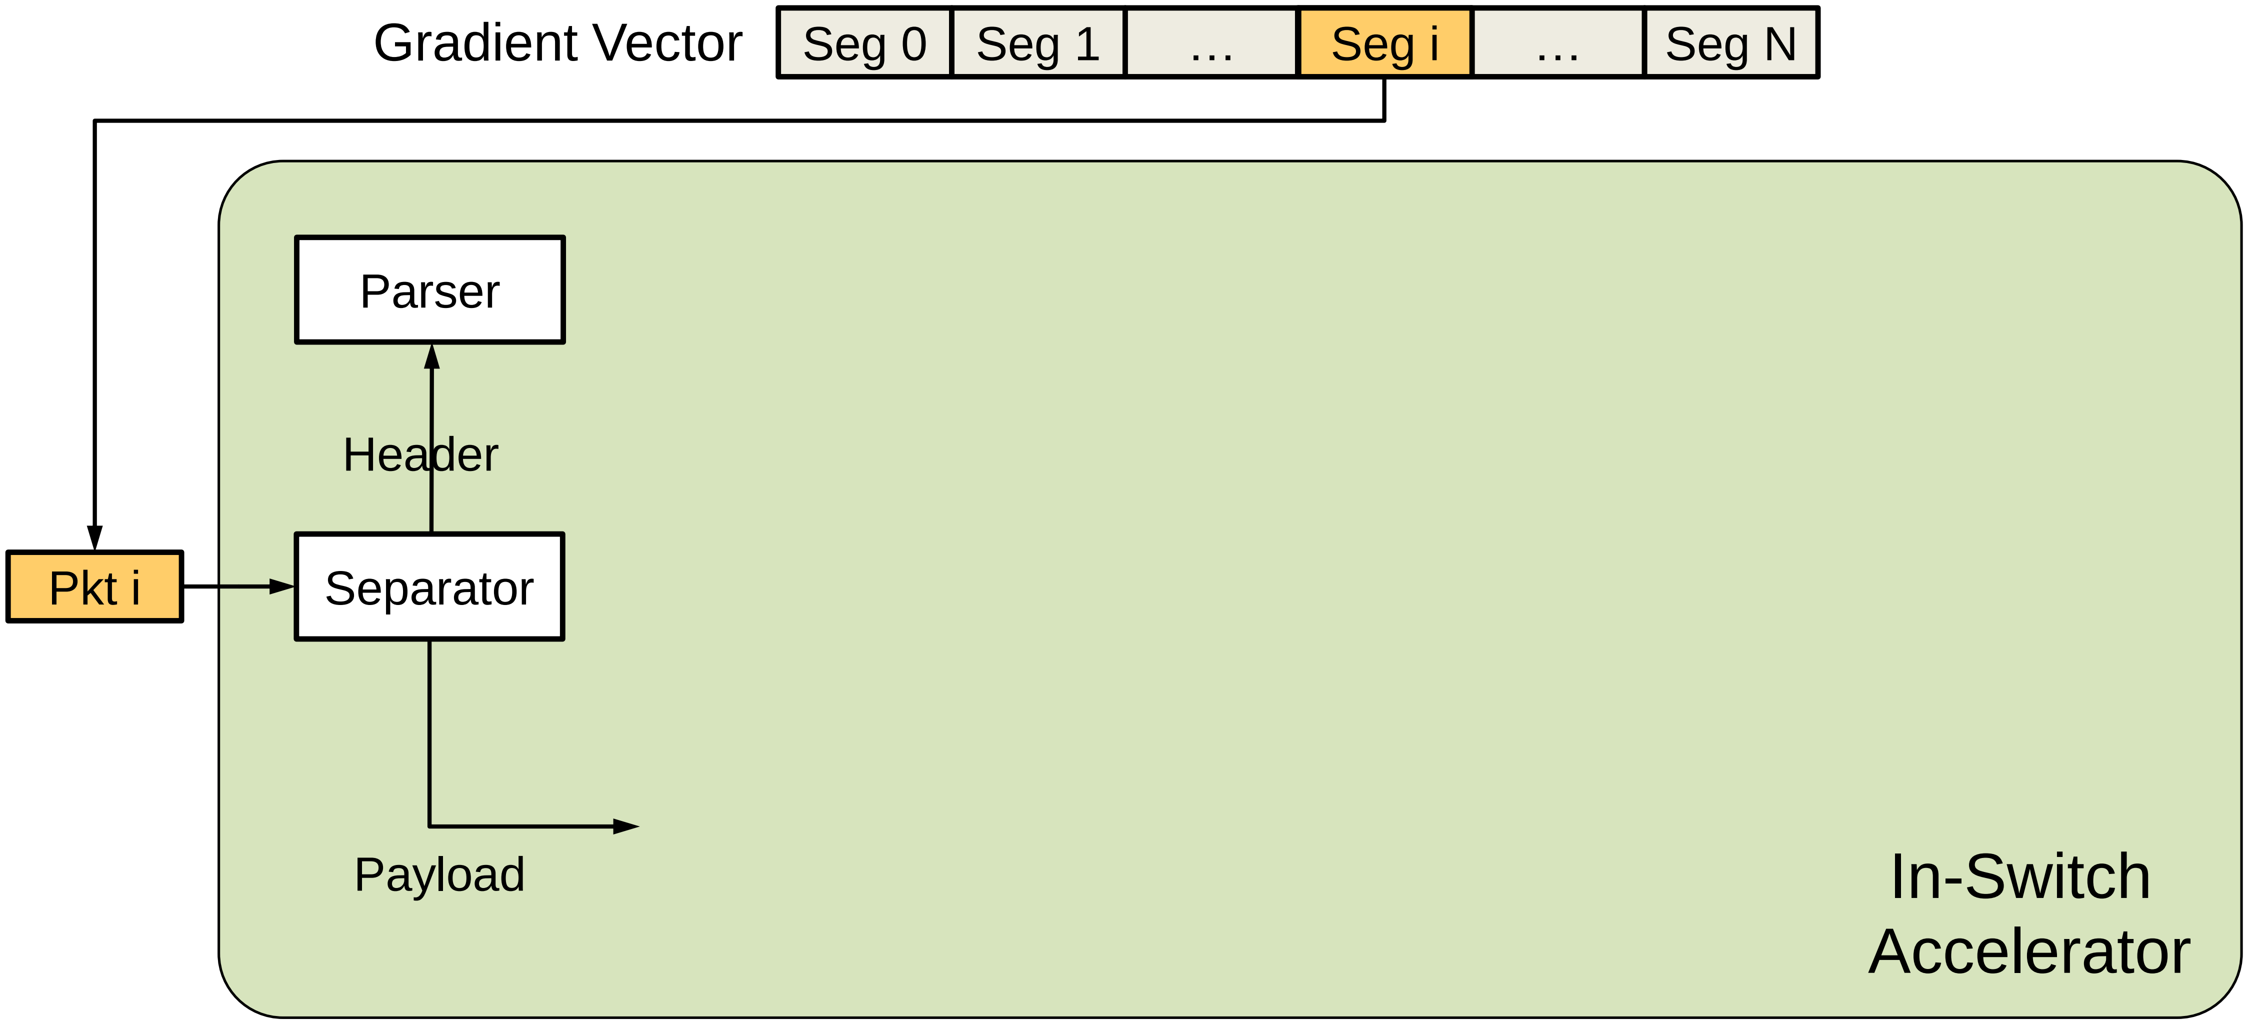

Developing Light-Weight Accelerator for Aggregation

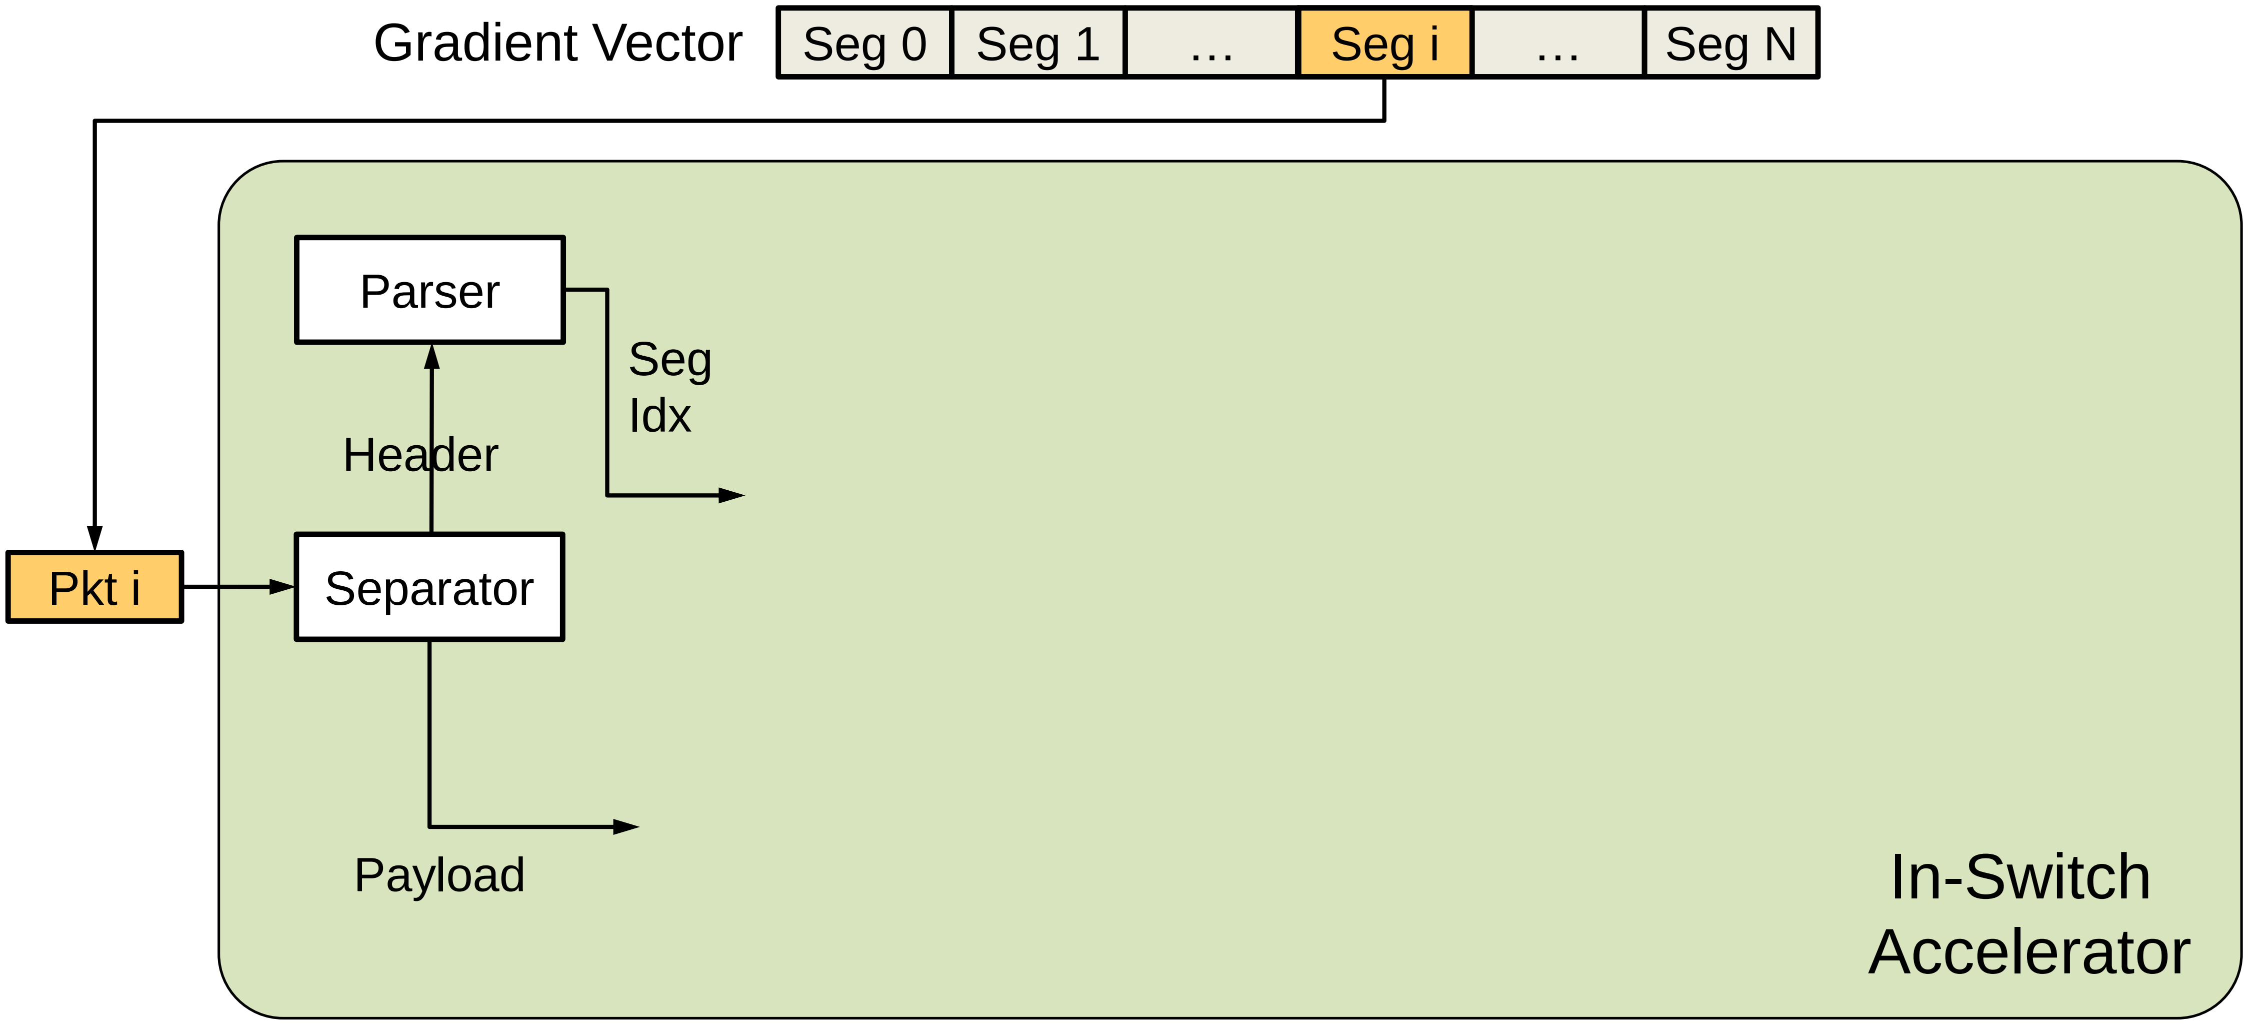

Developing Light-Weight Accelerator for Aggregation

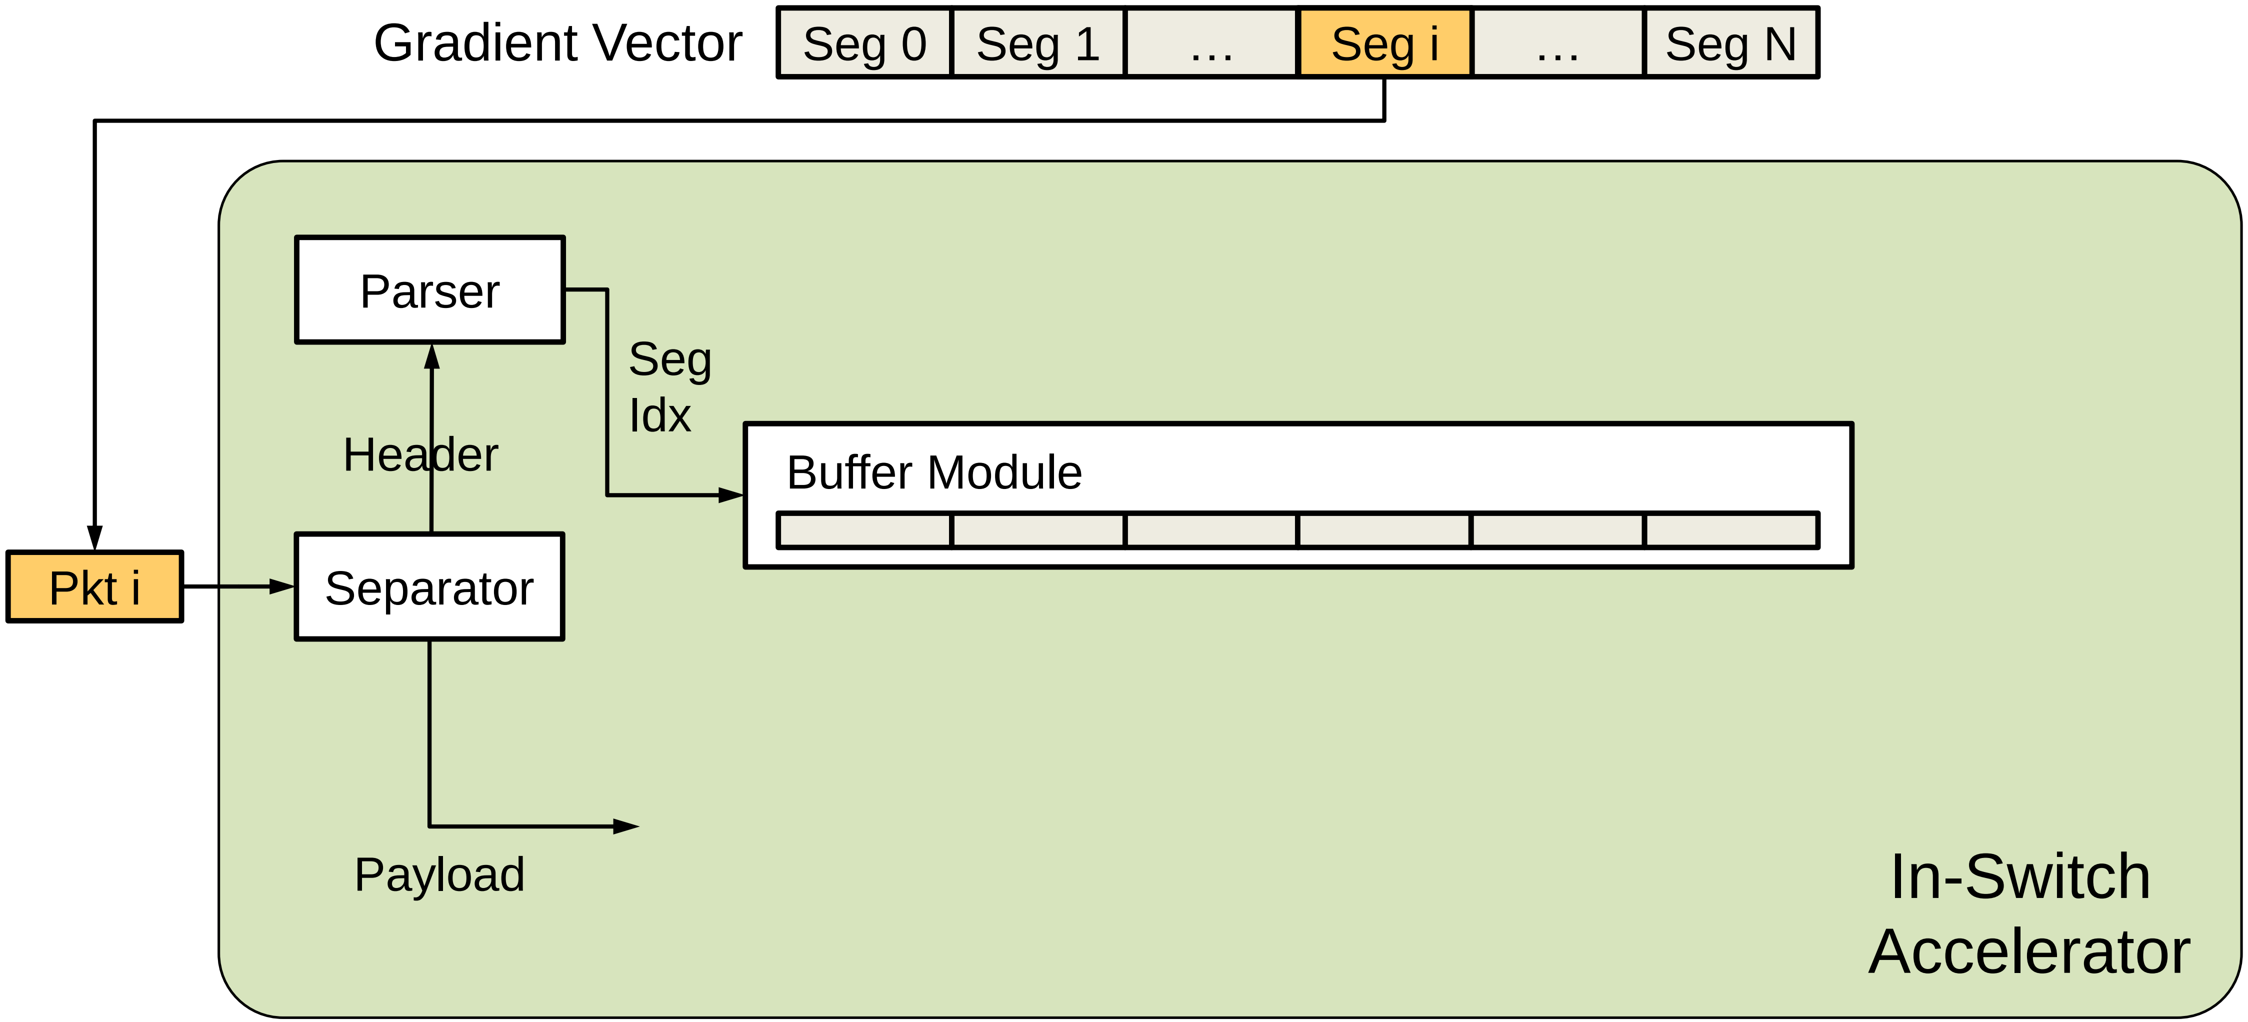

Developing Light-Weight Accelerator for Aggregation

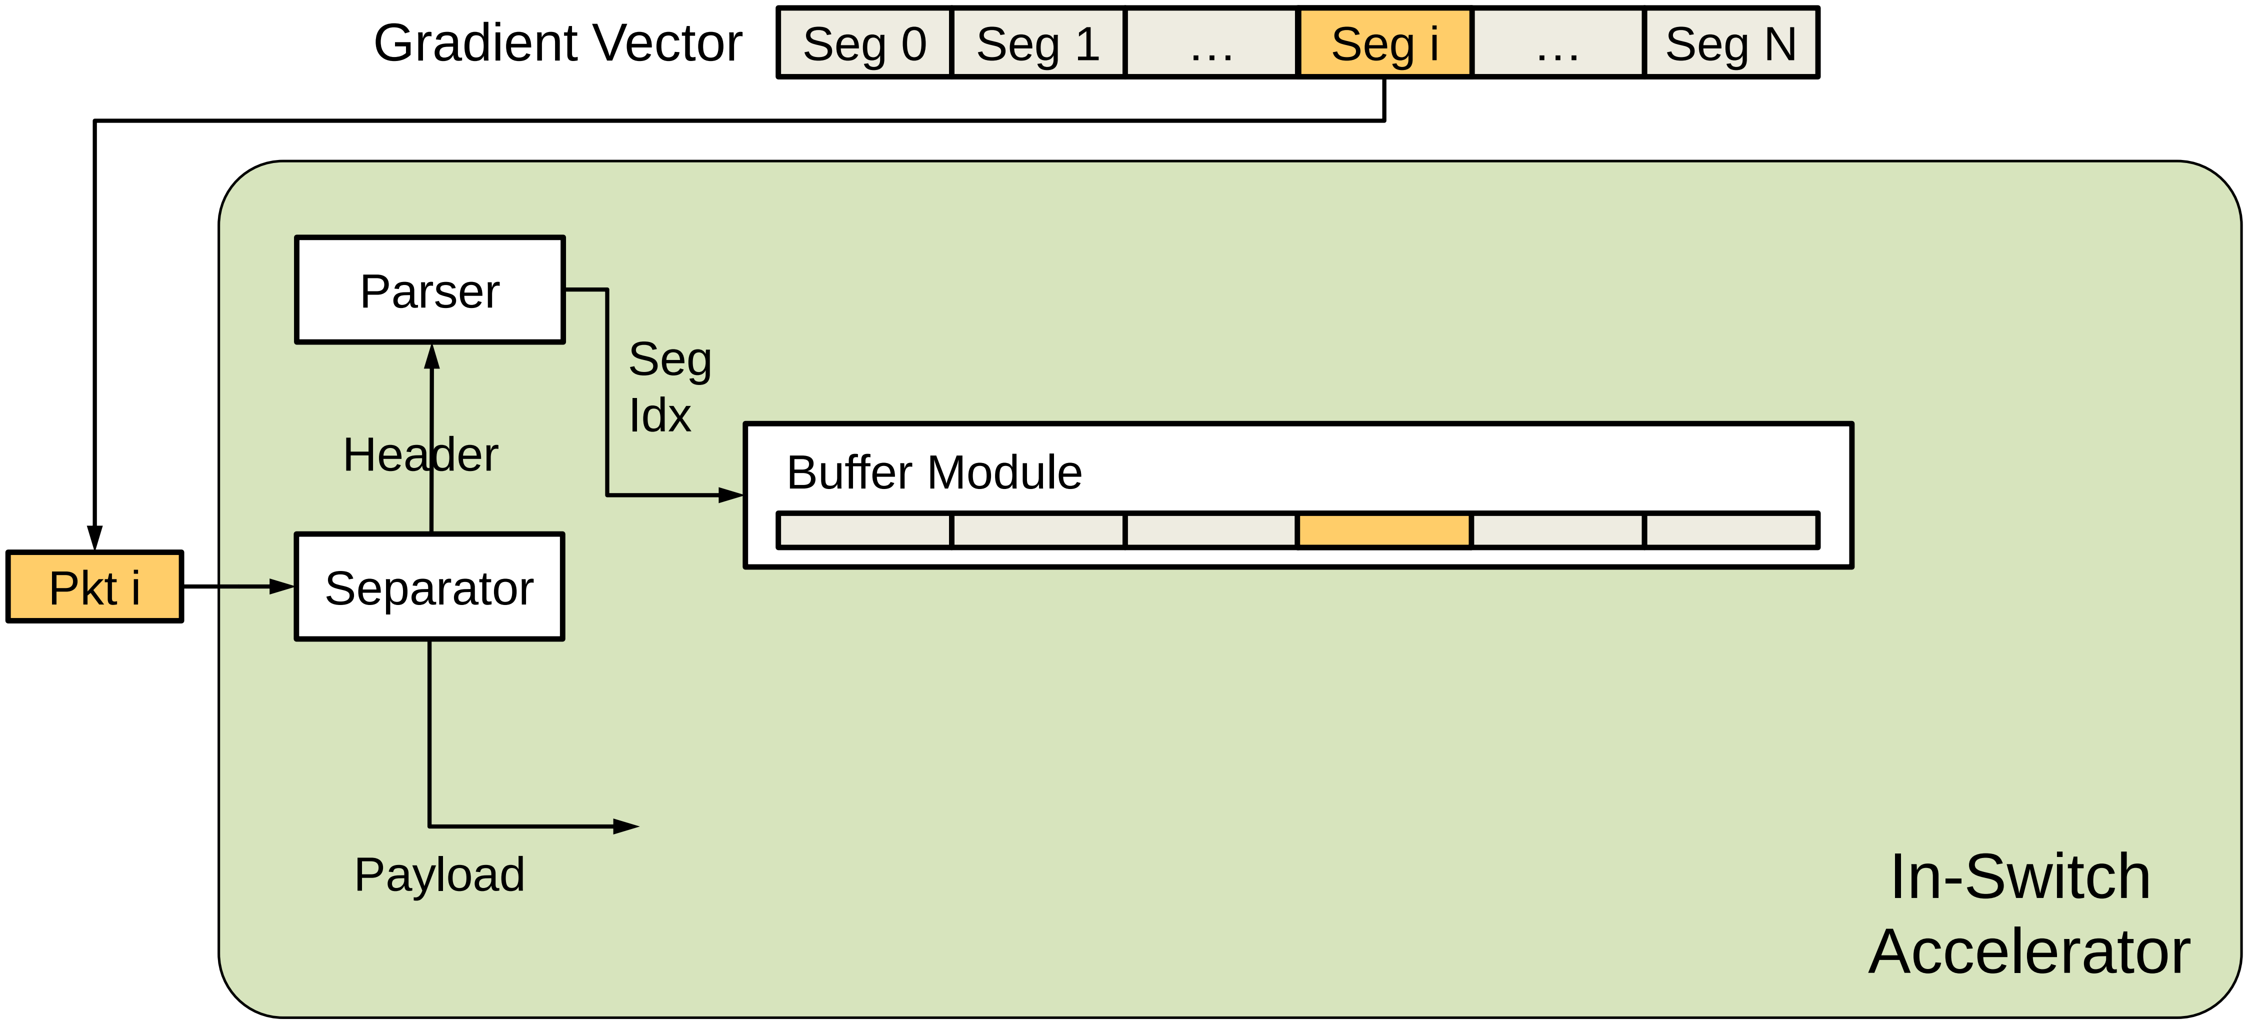

Developing Light-Weight Accelerator for Aggregation

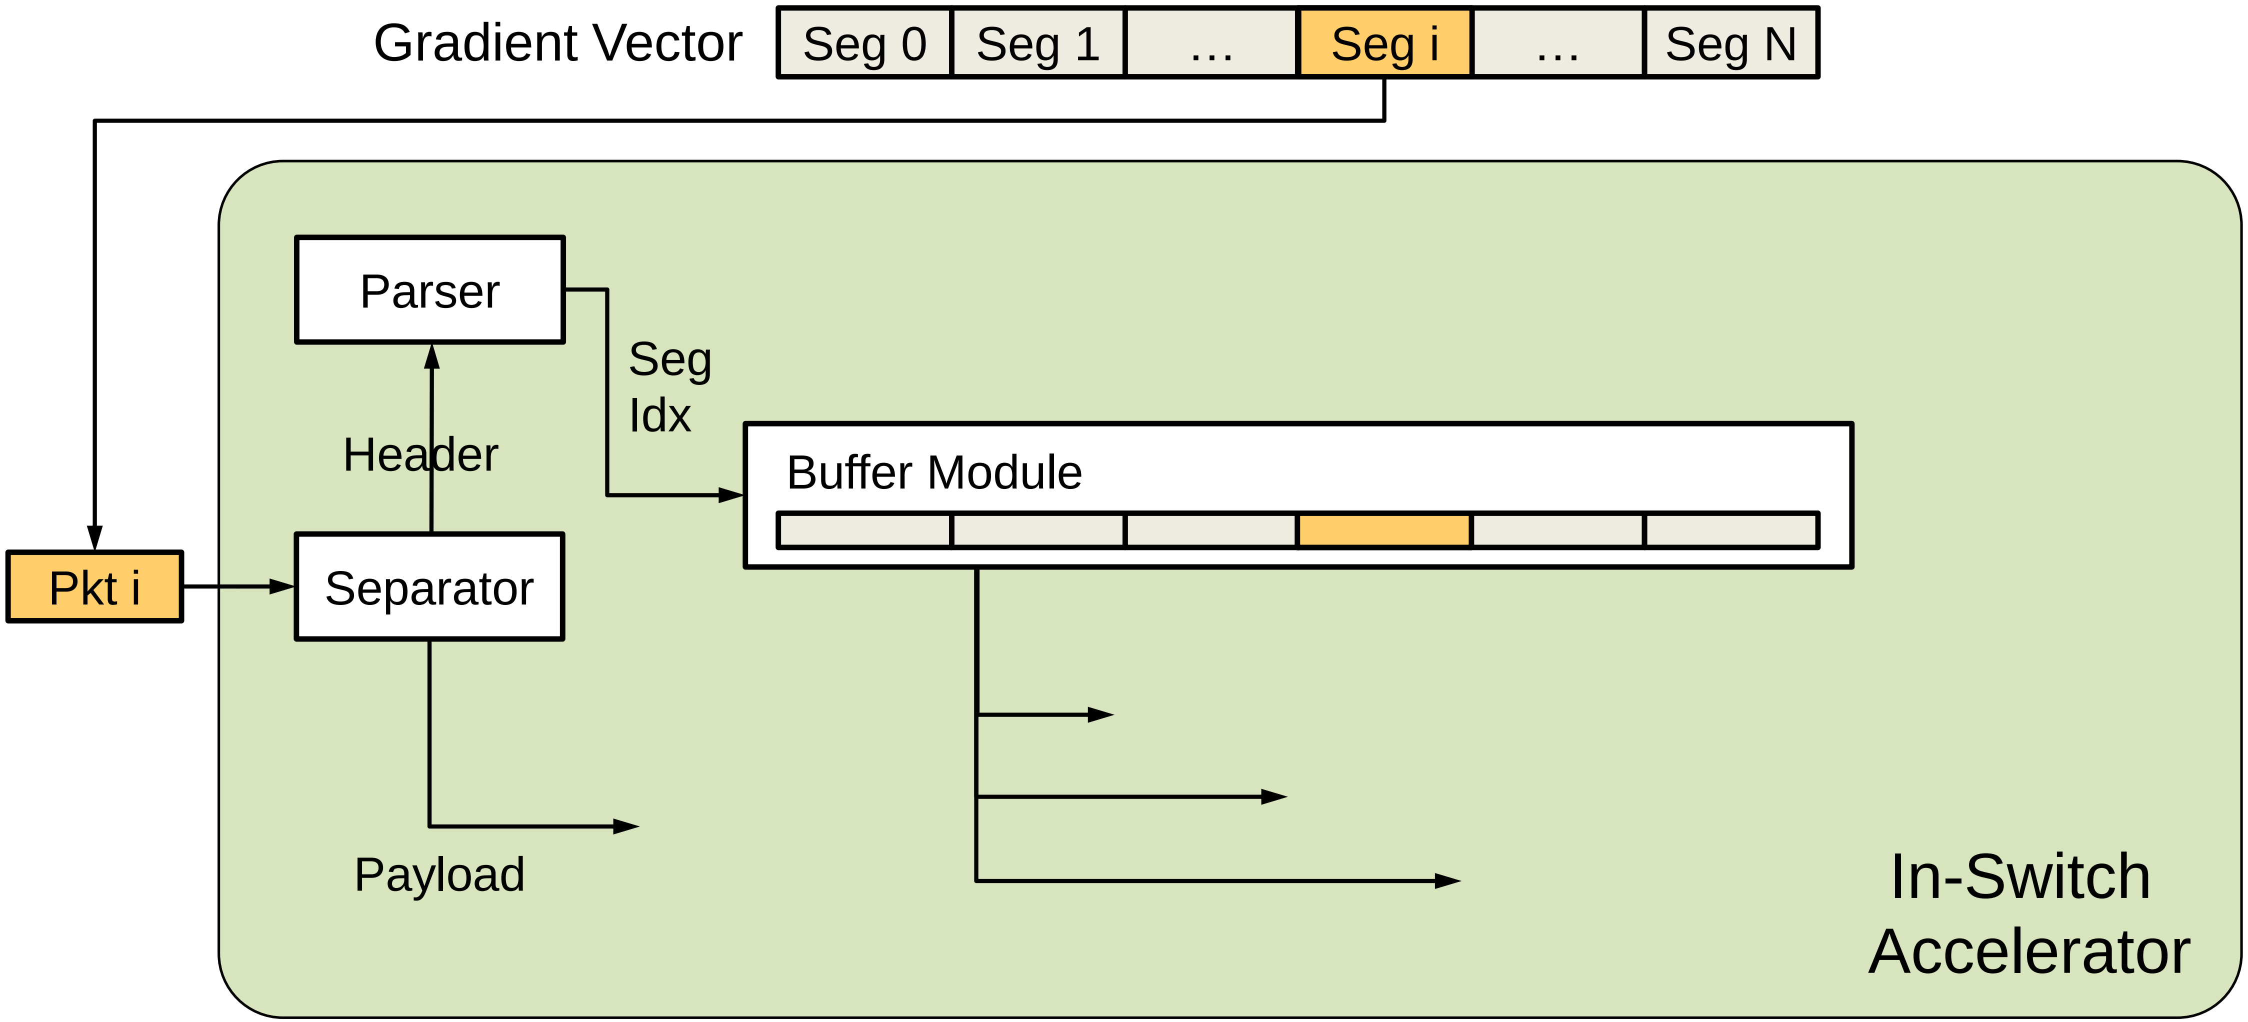

Developing Light-Weight Accelerator for Aggregation

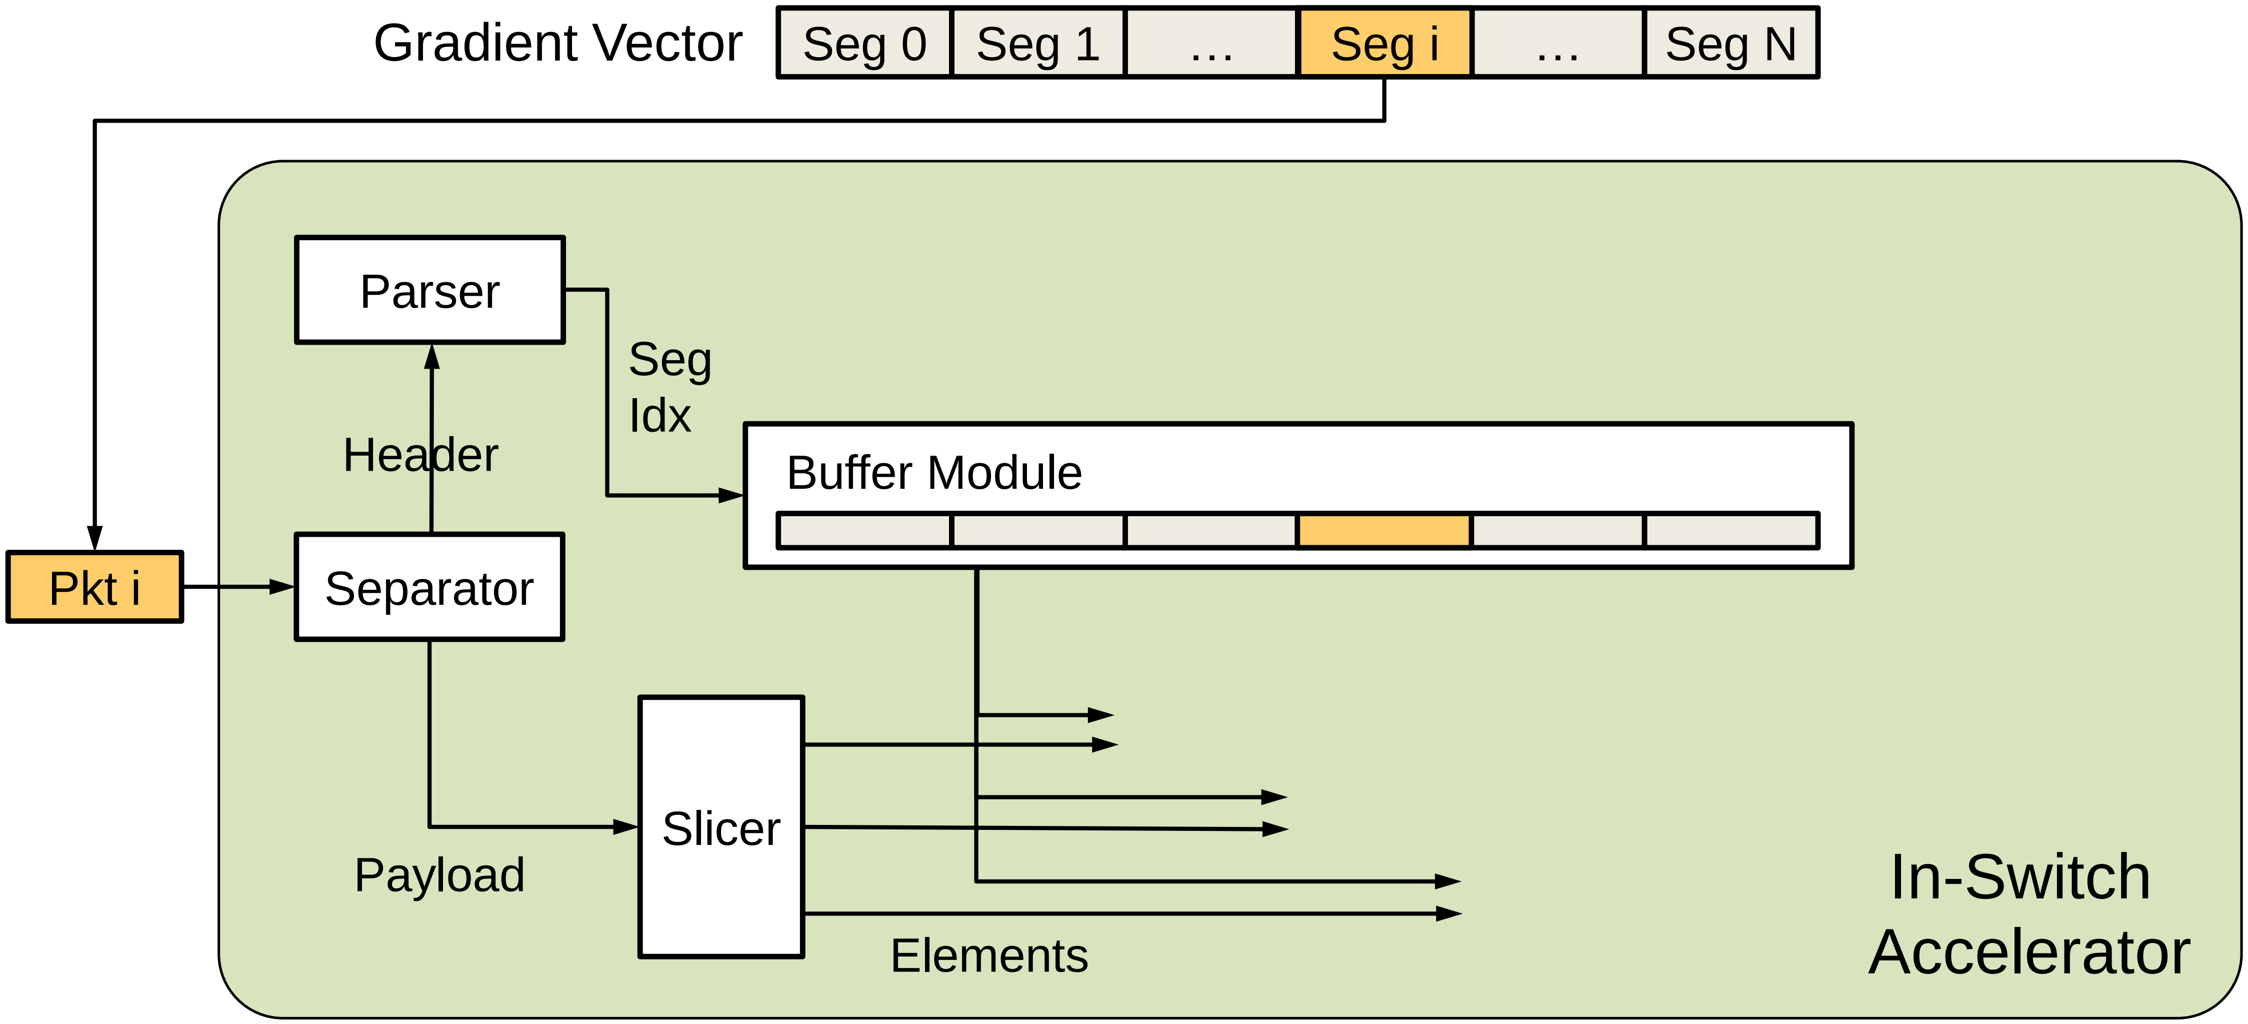

Developing Light-Weight Accelerator for Aggregation

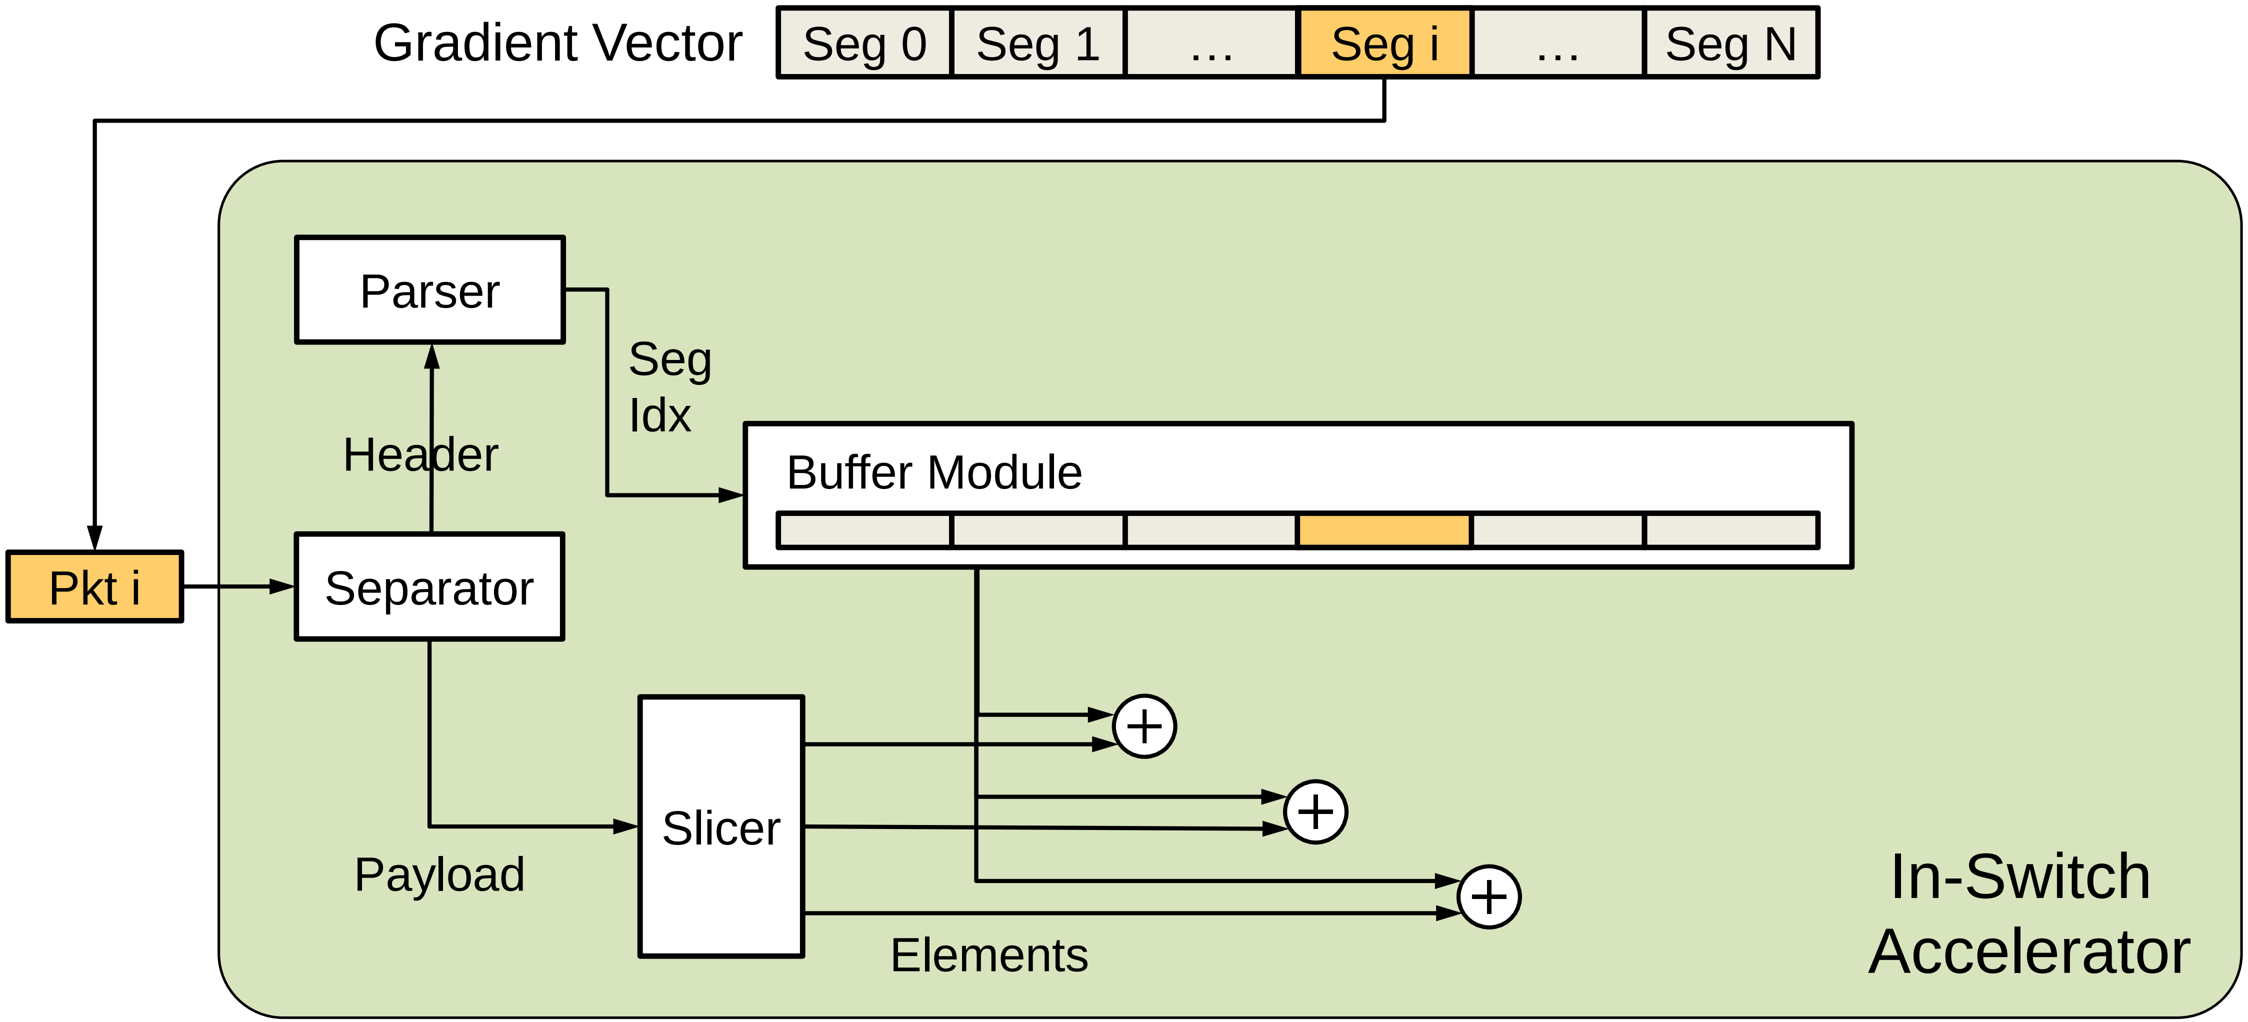

Developing Light-Weight Accelerator for Aggregation

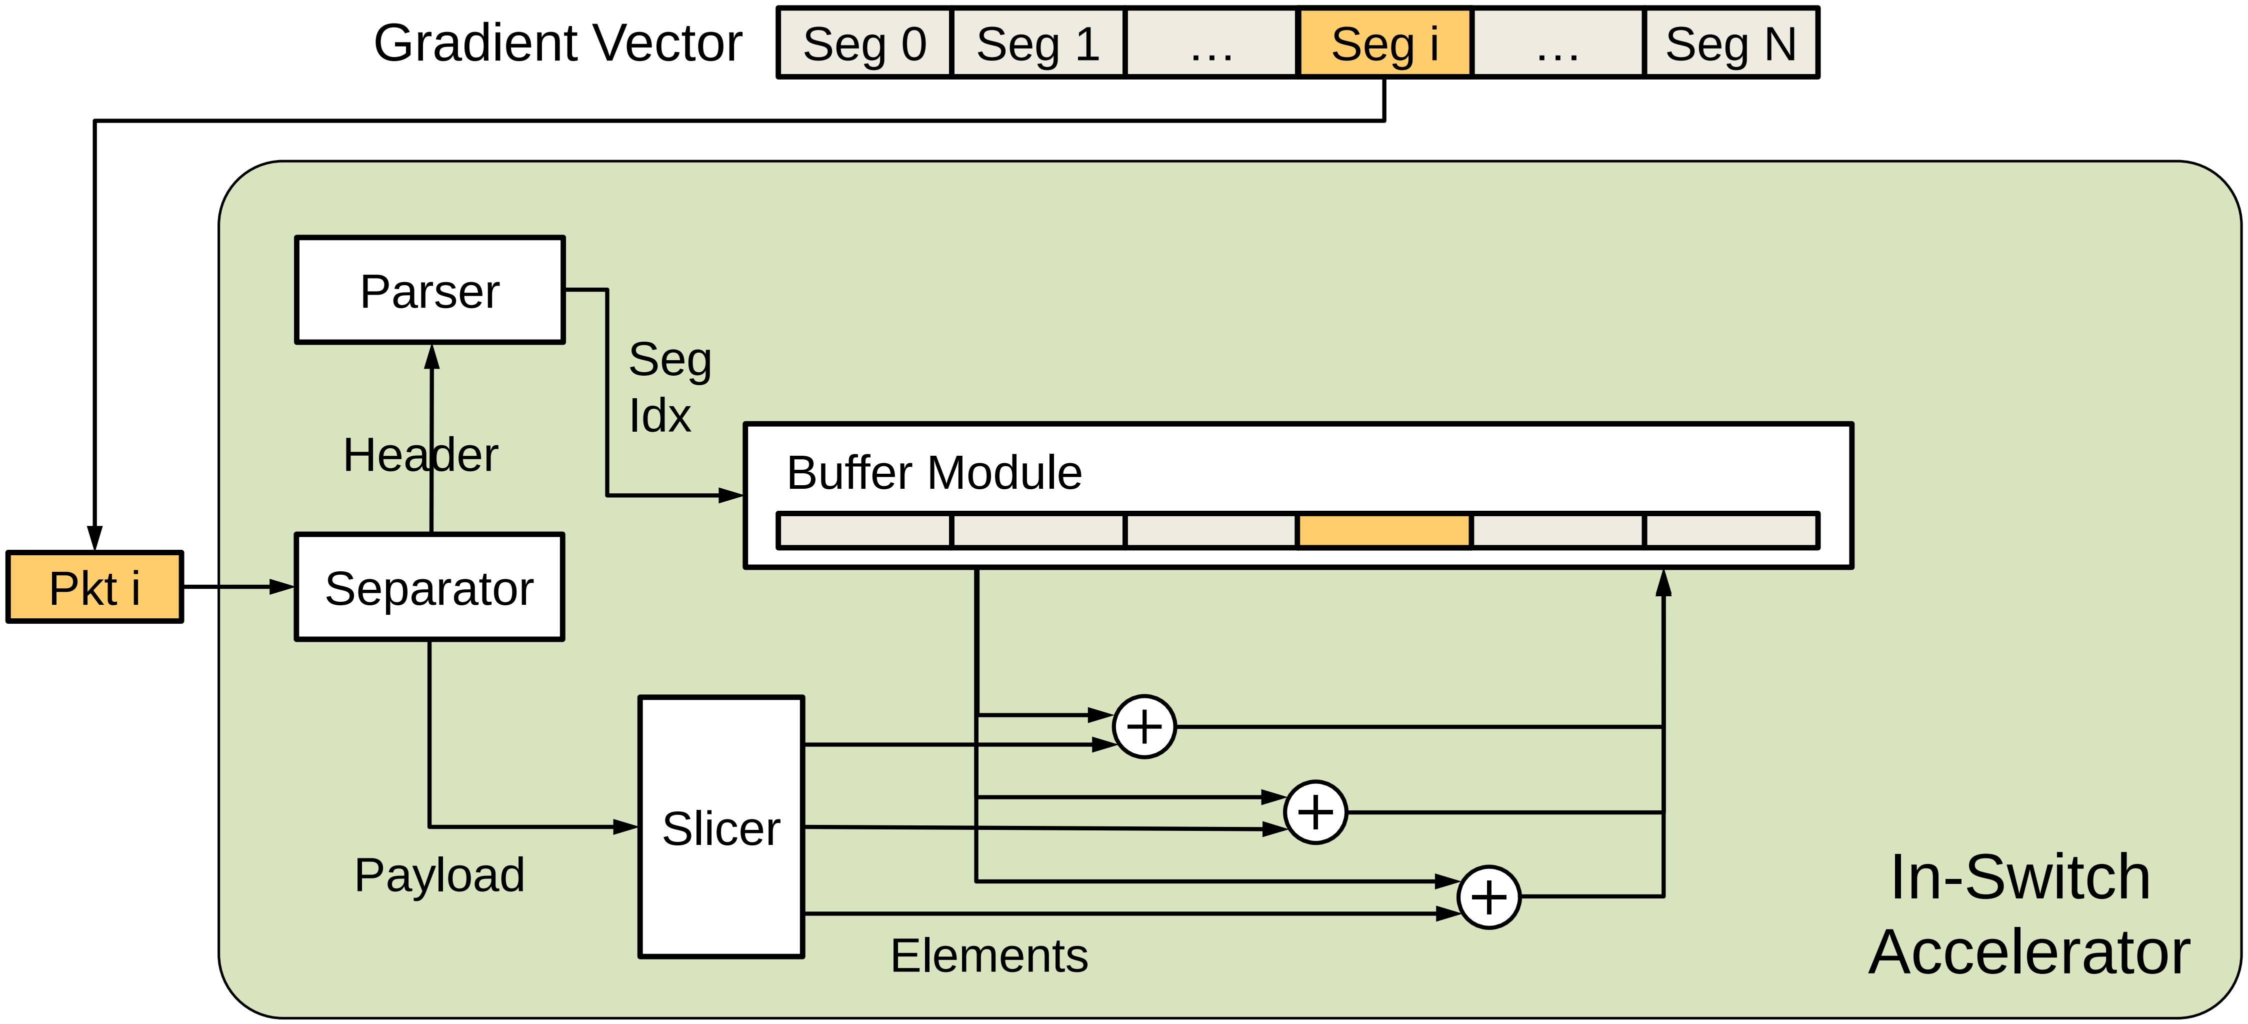

Developing Light-Weight Accelerator for Aggregation

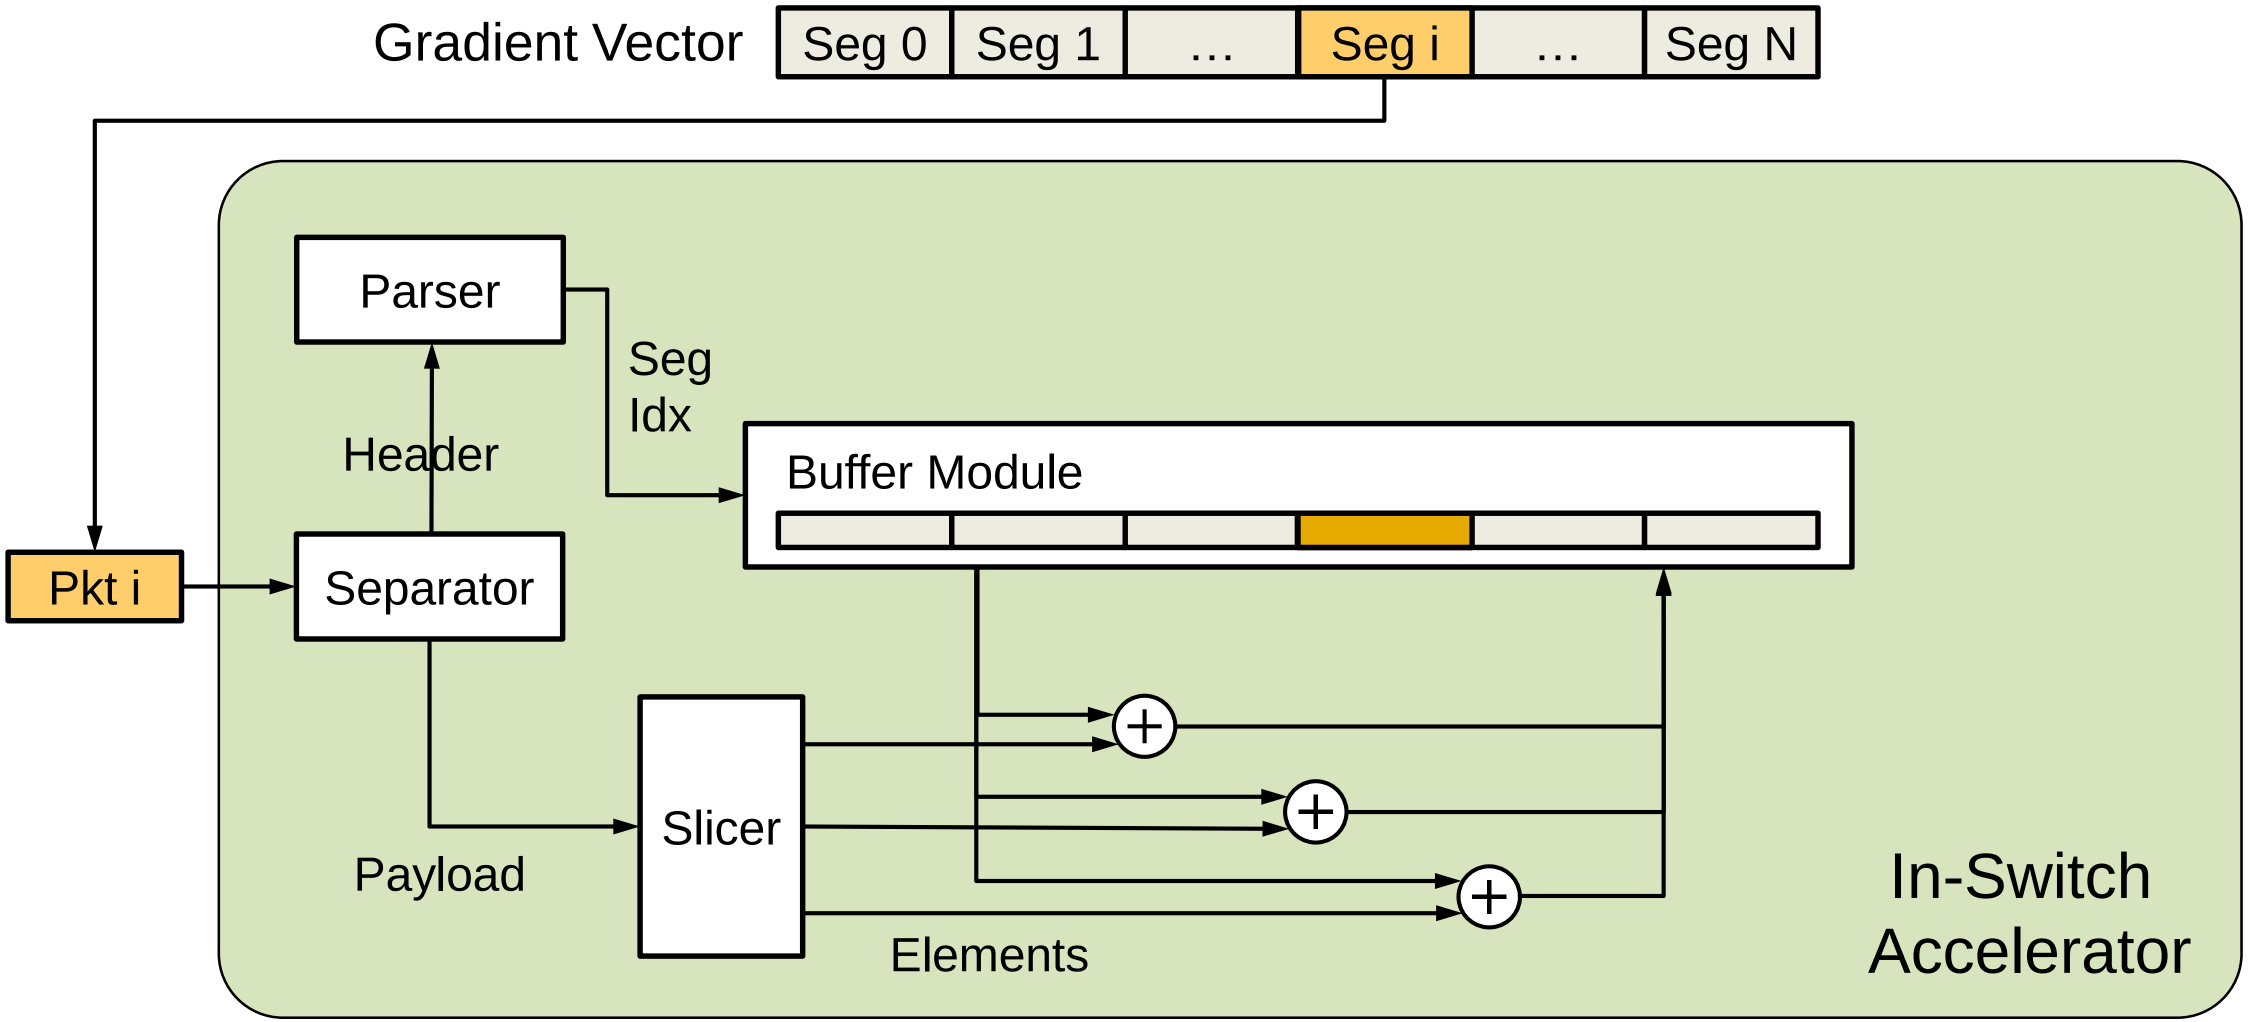

Developing Light-Weight Accelerator for Aggregation

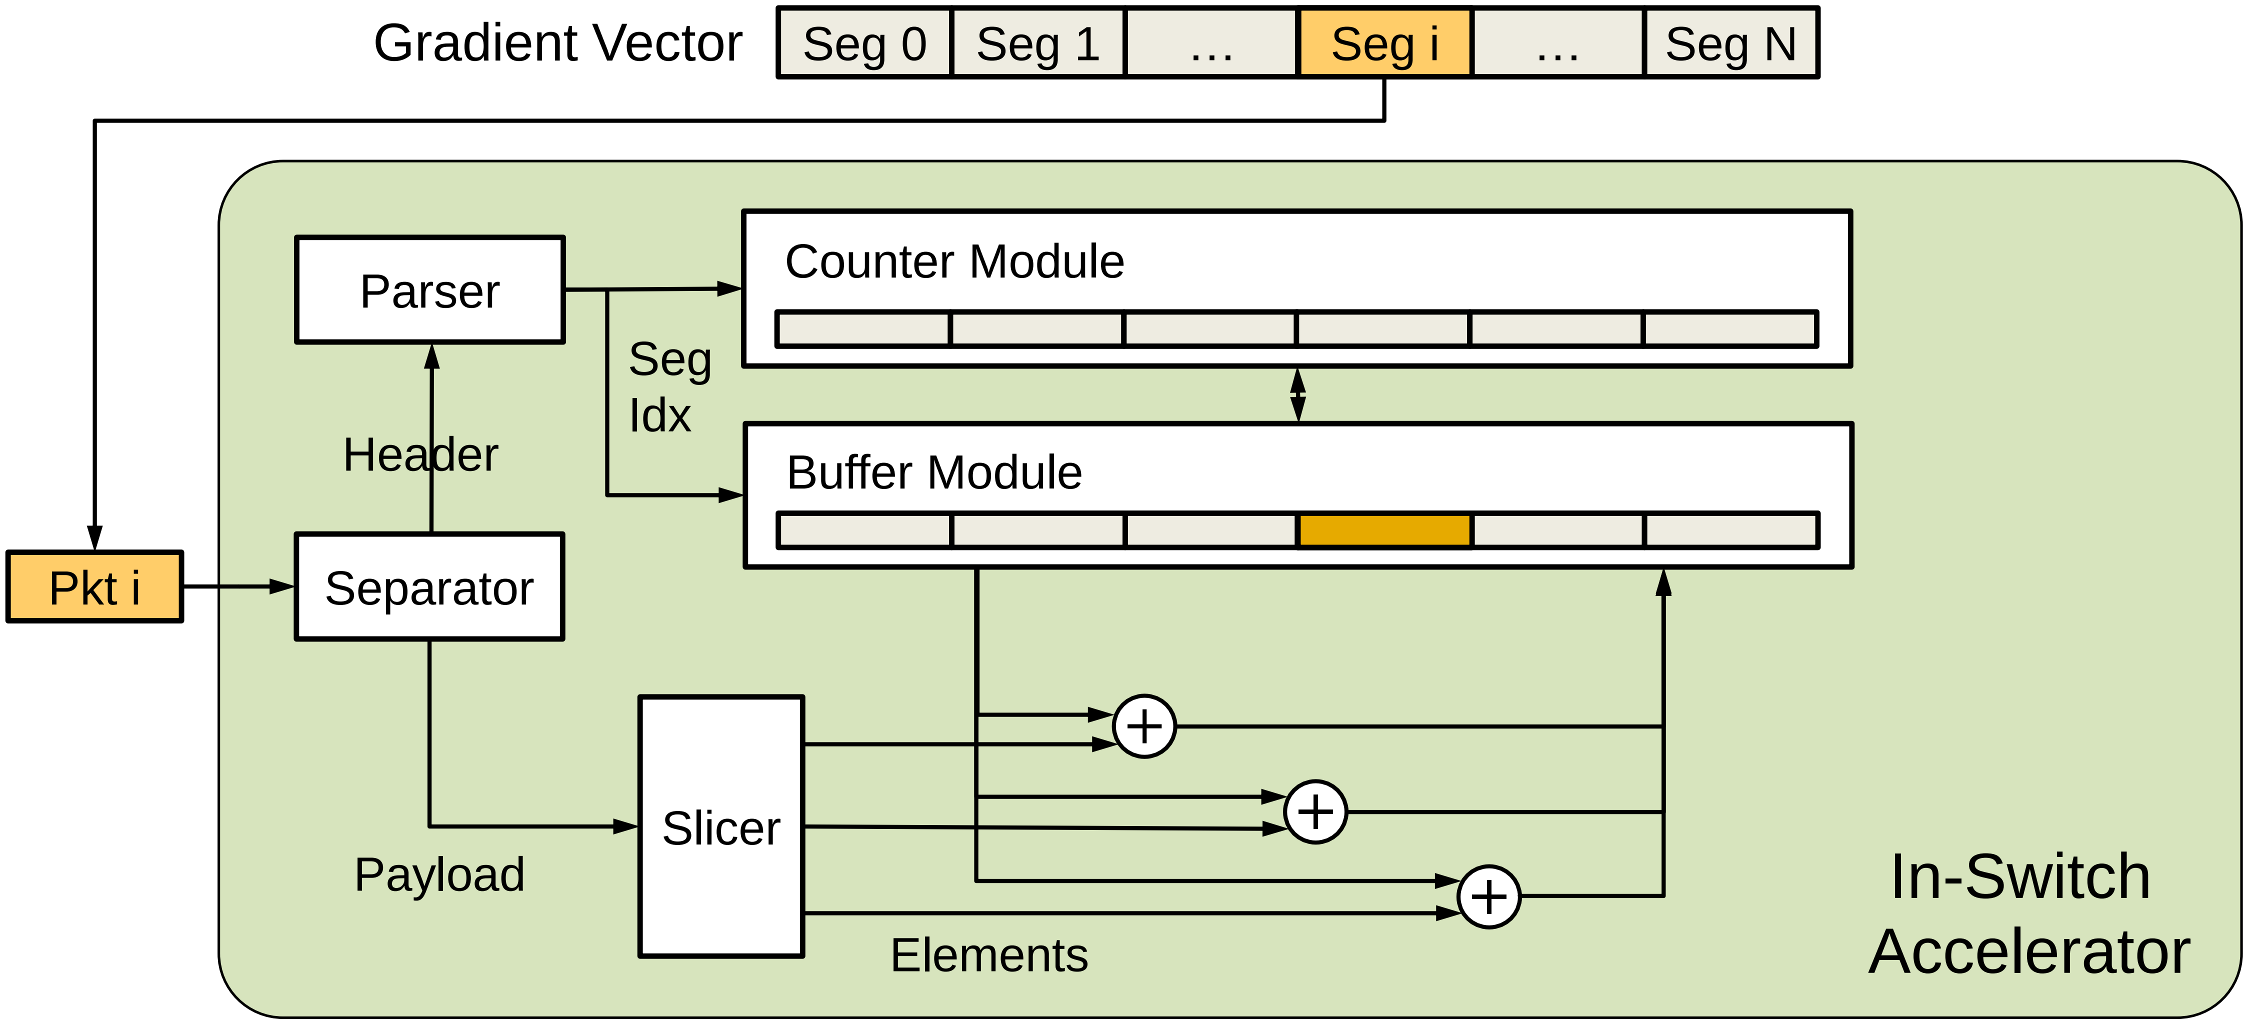

Developing Light-Weight Accelerator for Aggregation

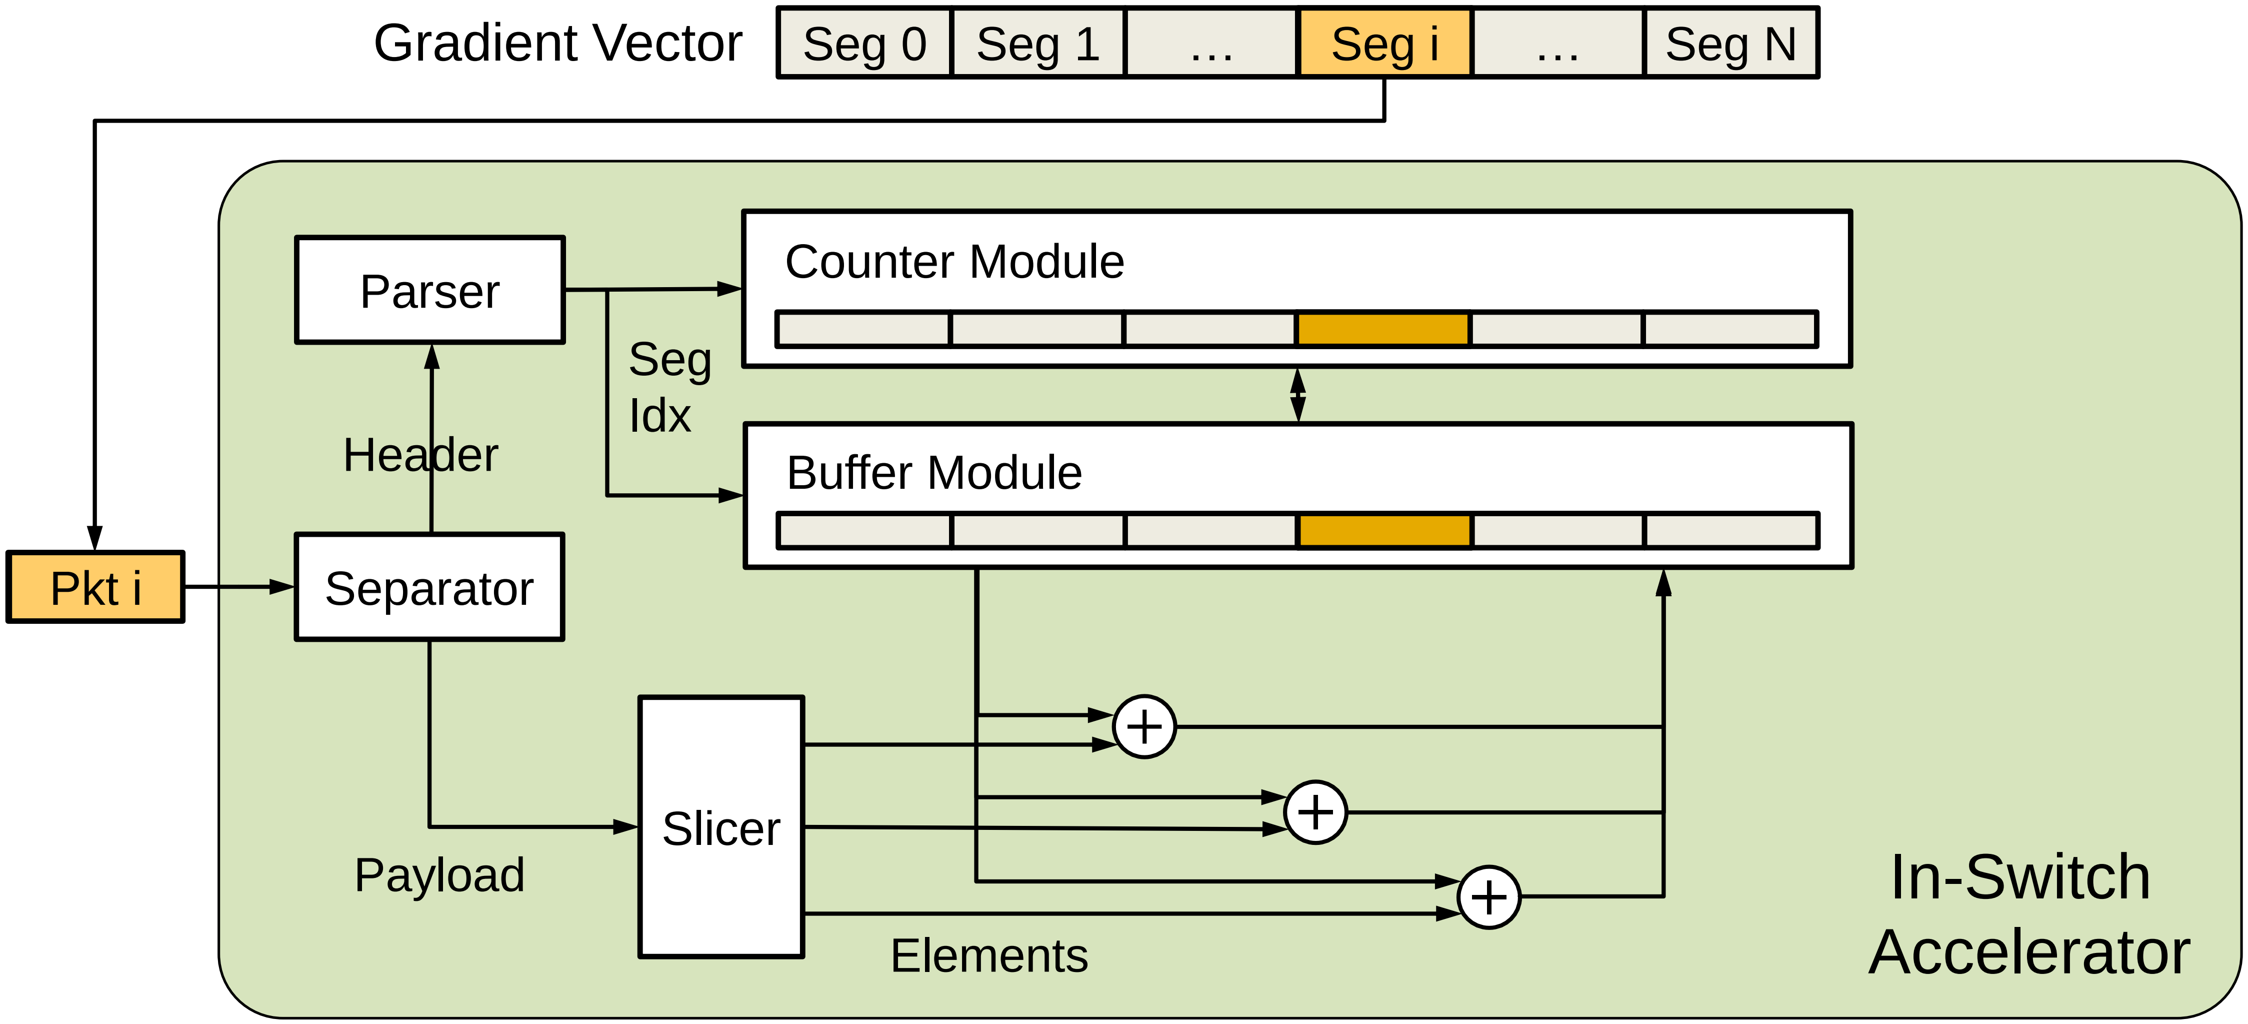

Developing Light-Weight Accelerator for Aggregation

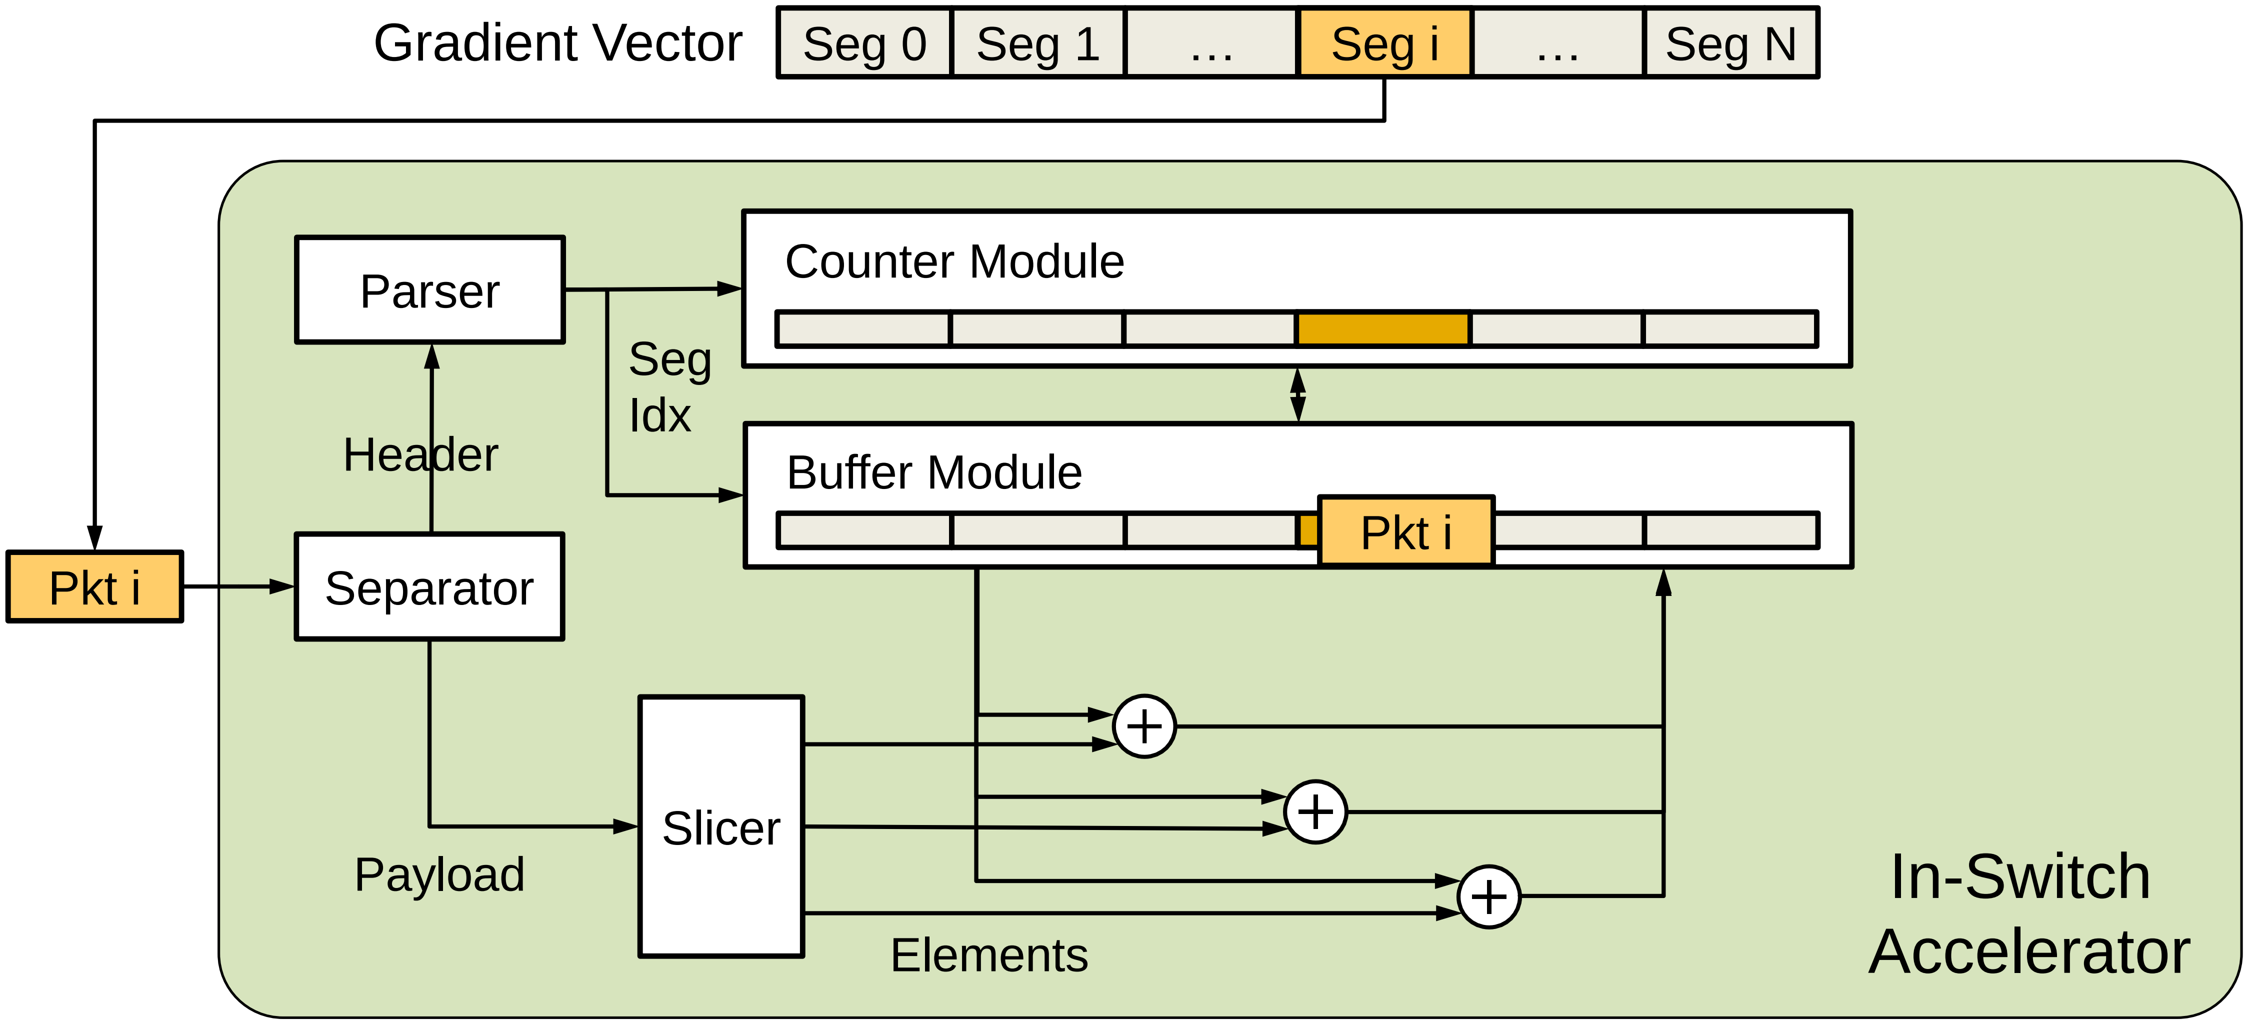

Developing Light-Weight Accelerator for Aggregation

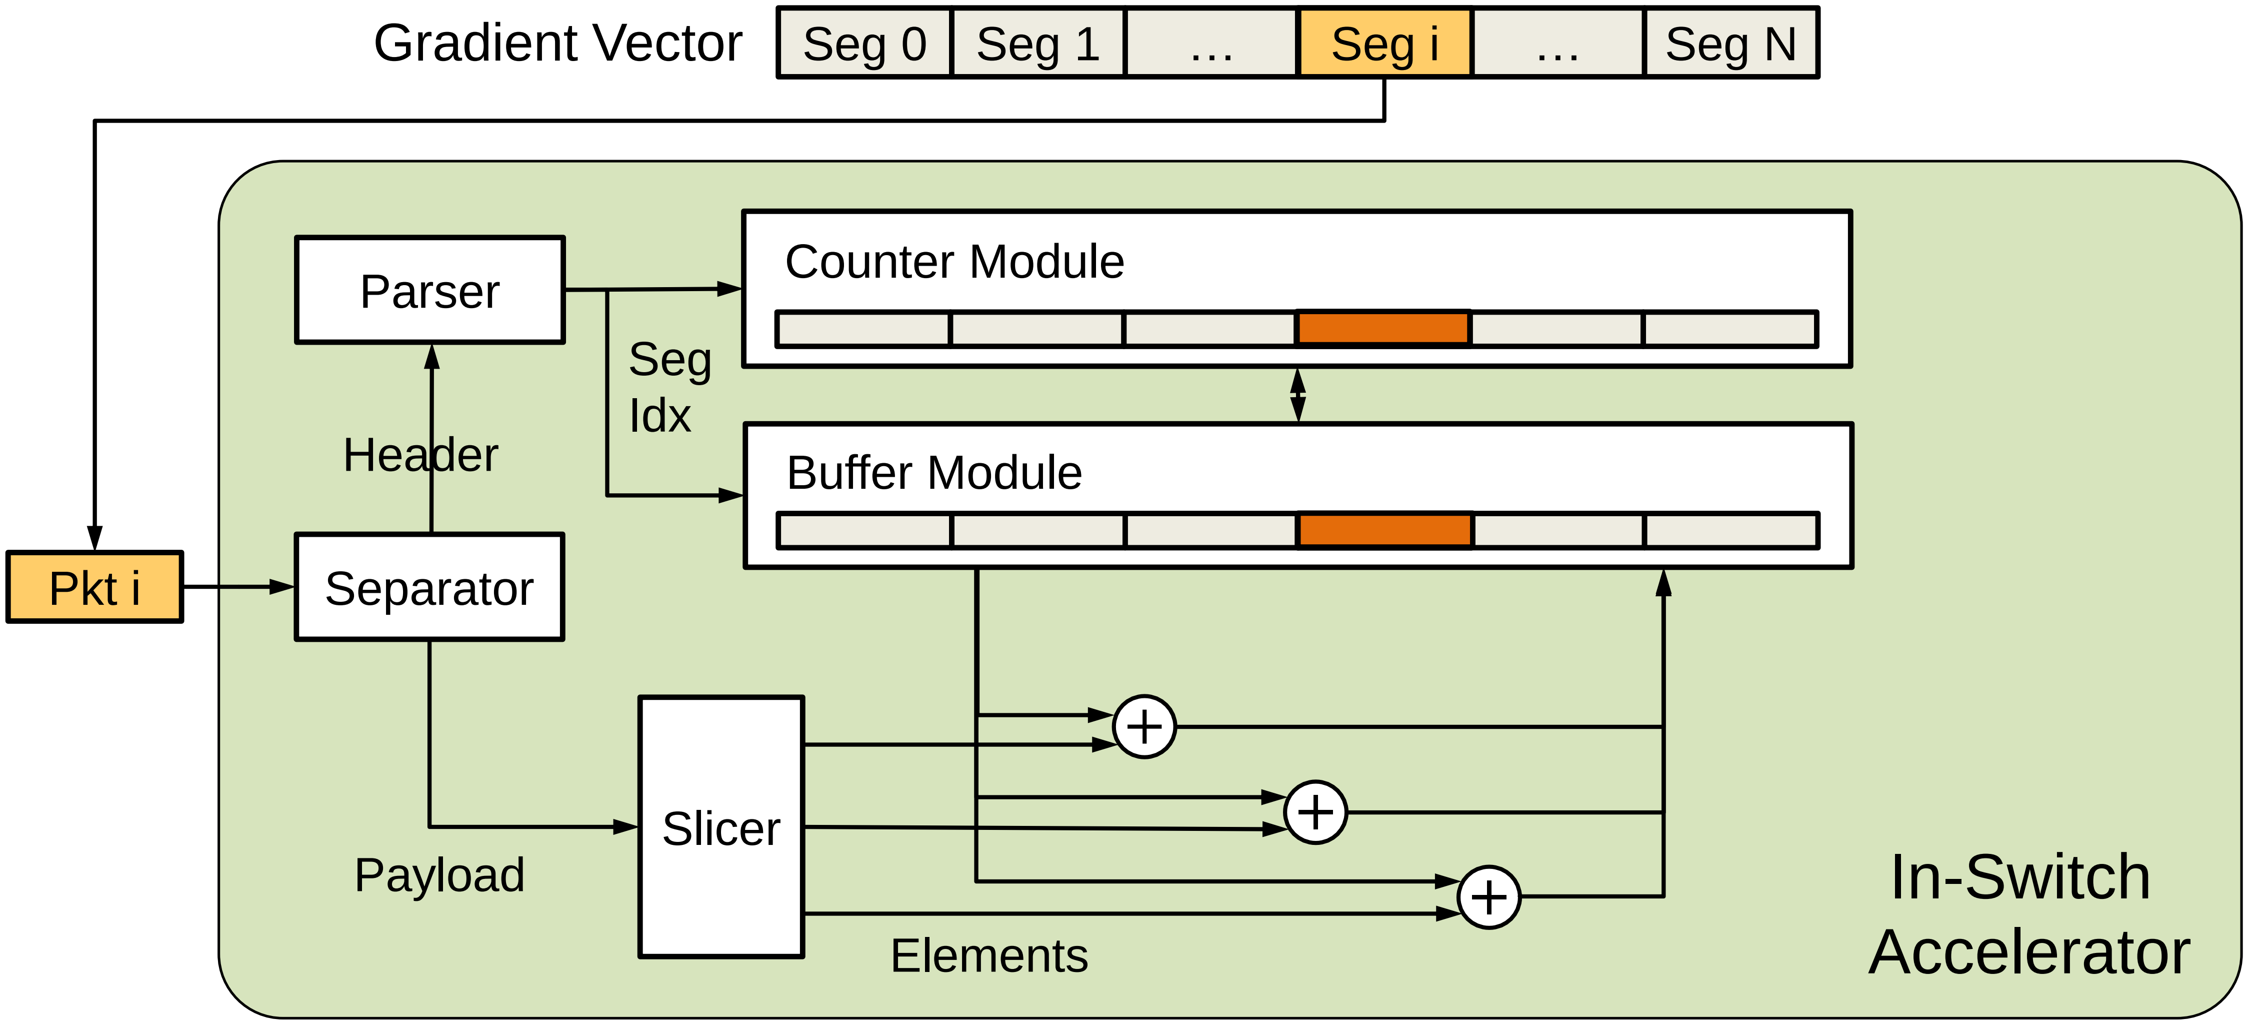

Developing Light-Weight Accelerator for Aggregation

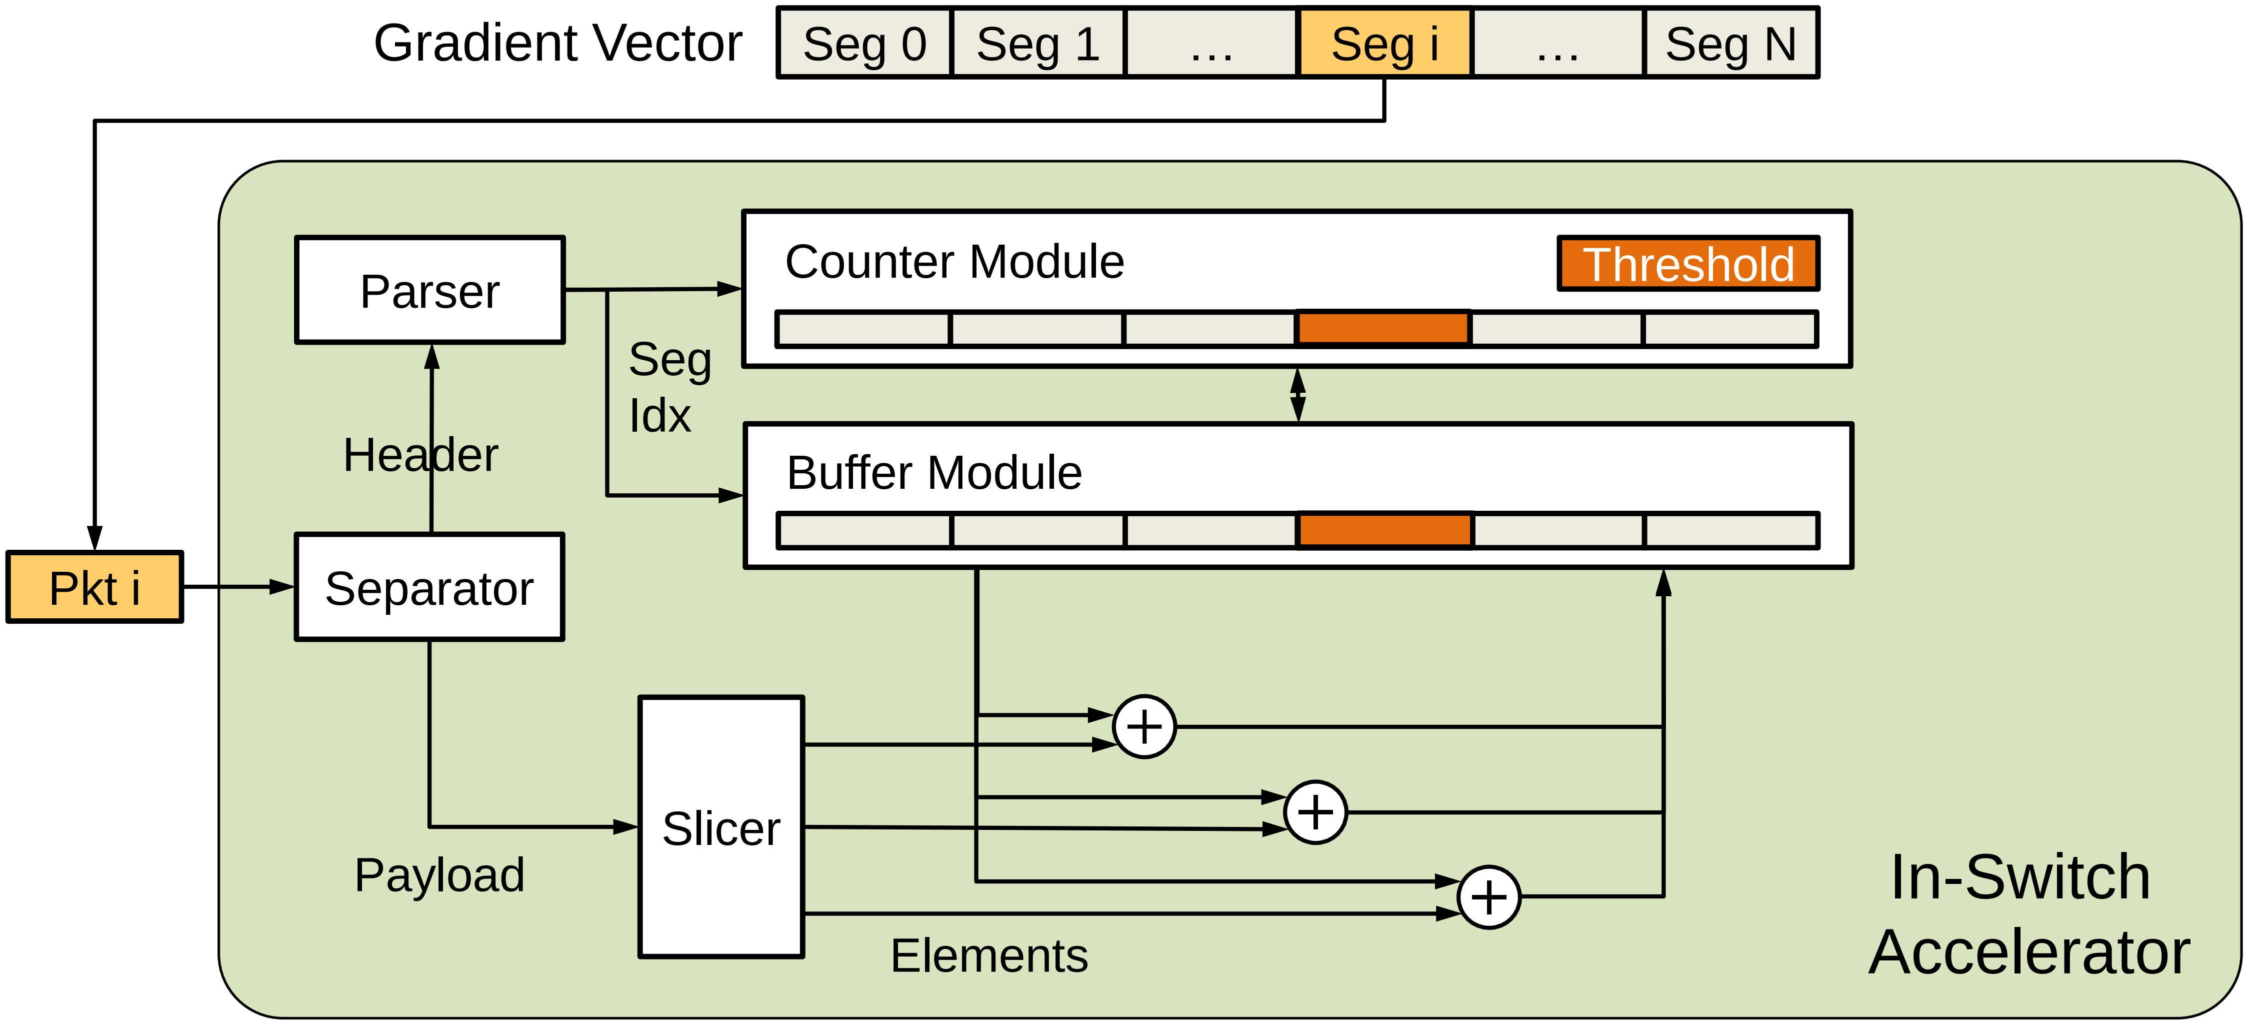

Developing Light-Weight Accelerator for Aggregation

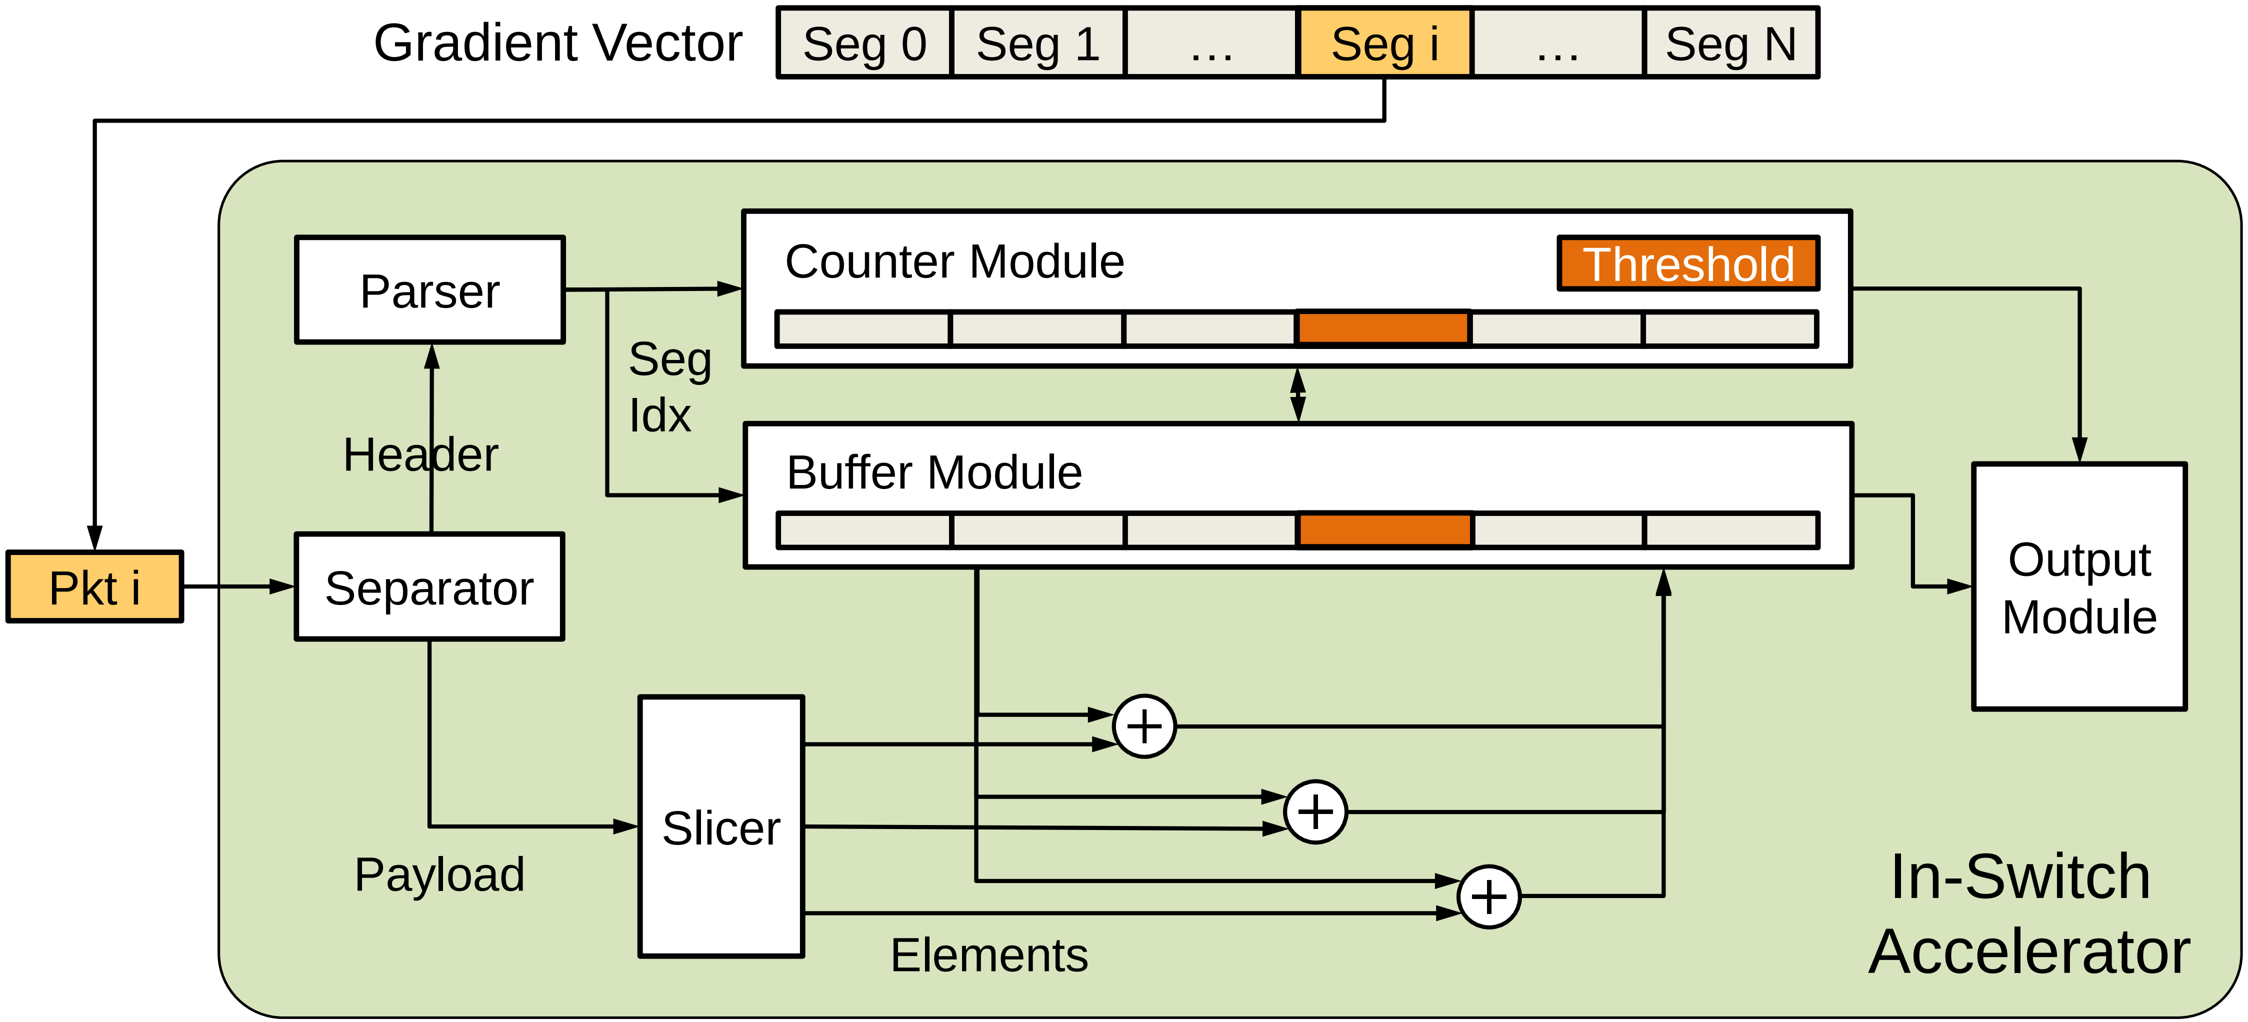

Developing Light-Weight Accelerator for Aggregation

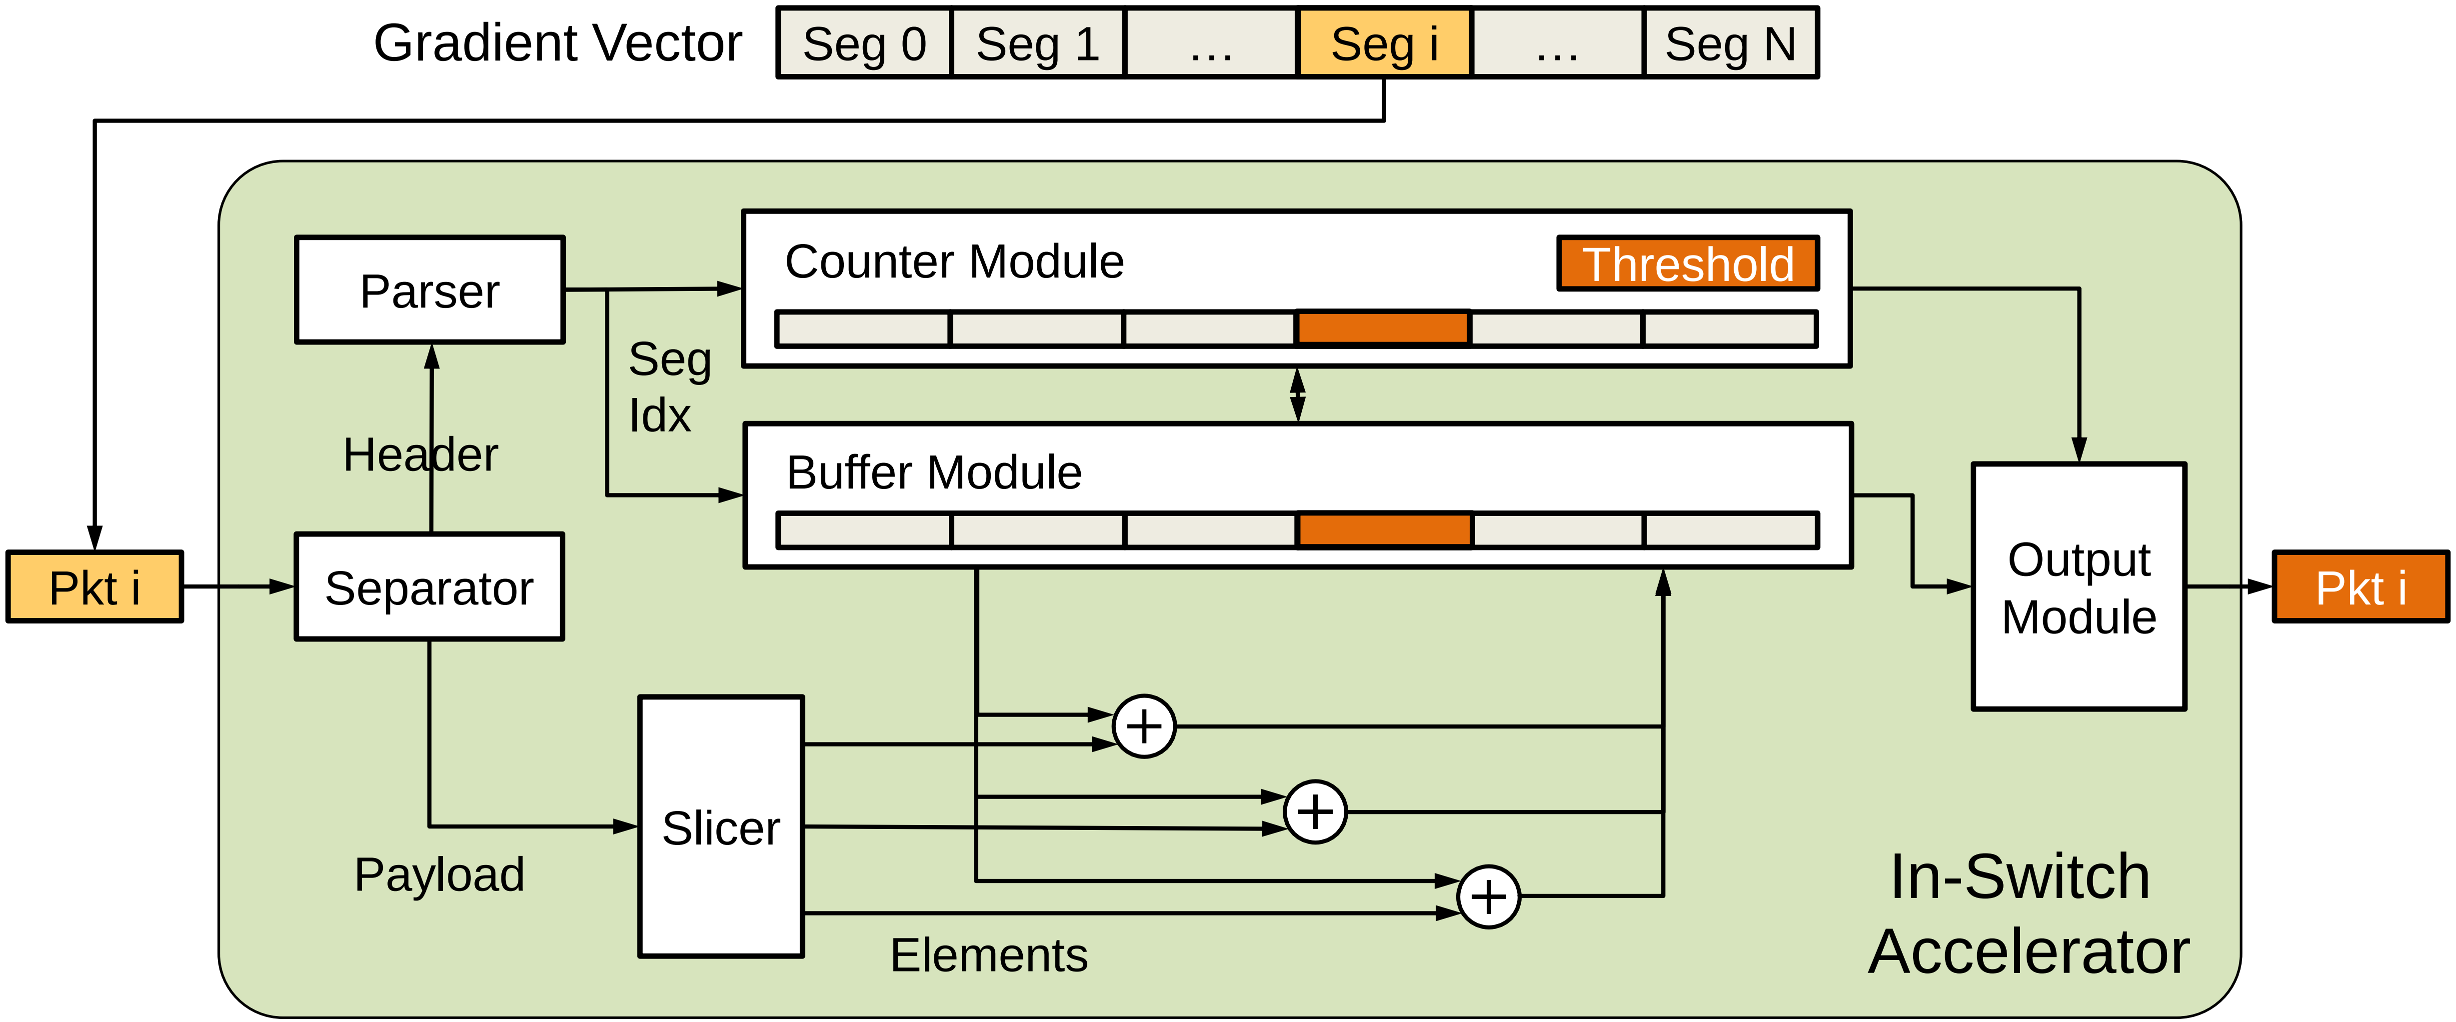

ECE ILLINOIS

Developing Light-Weight Accelerator for Aggregation

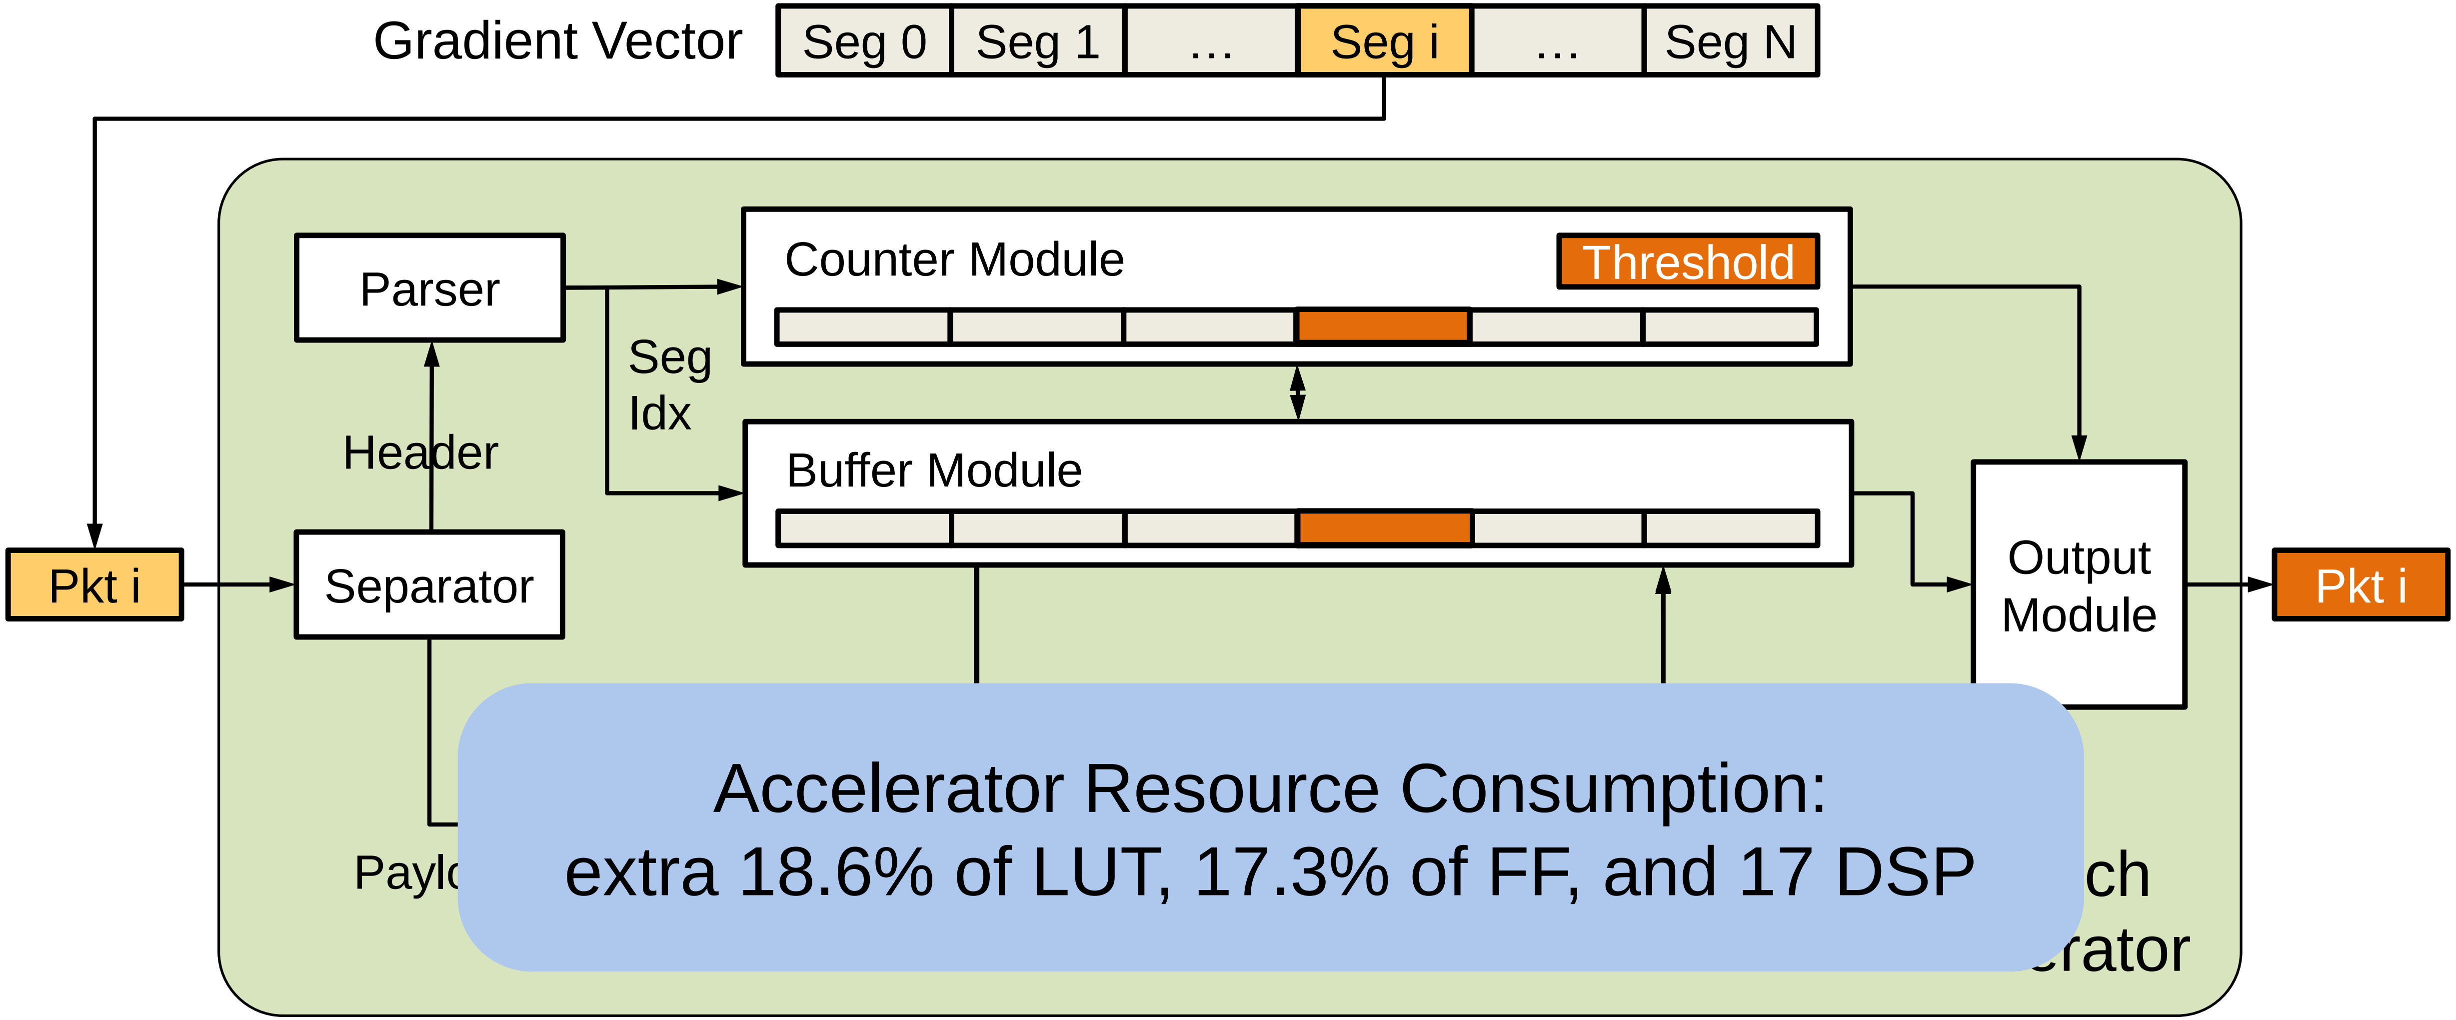

Accelerator Resource Consumption:
extra 18.6% of LUT, 17.3% of FF, and 17 DSP

Aggregating Gradient at Packet-Level for Improved Parallelism

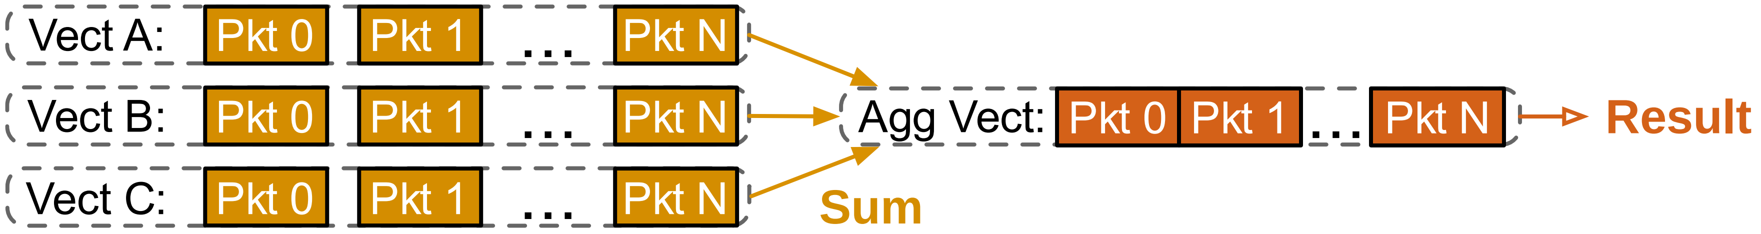

Conventional Vector-Level Aggregation

Aggregating Gradient at Packet-Level for Improved Parallelism

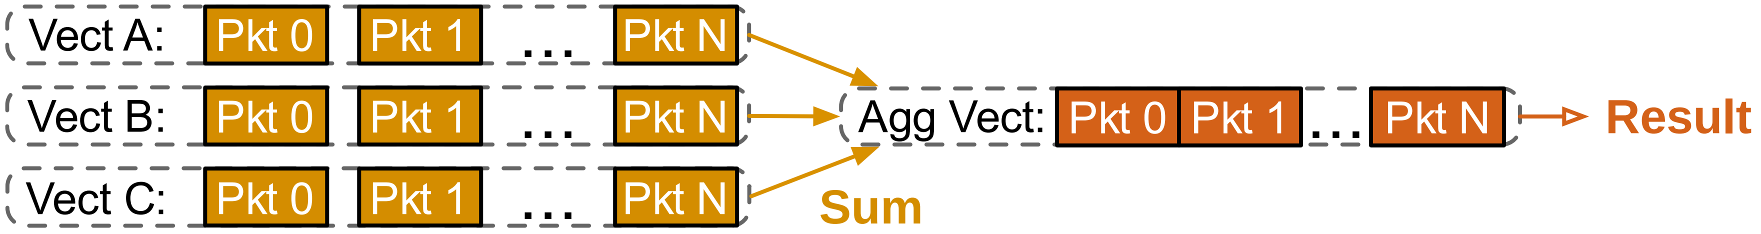

Conventional Vector-Level Aggregation

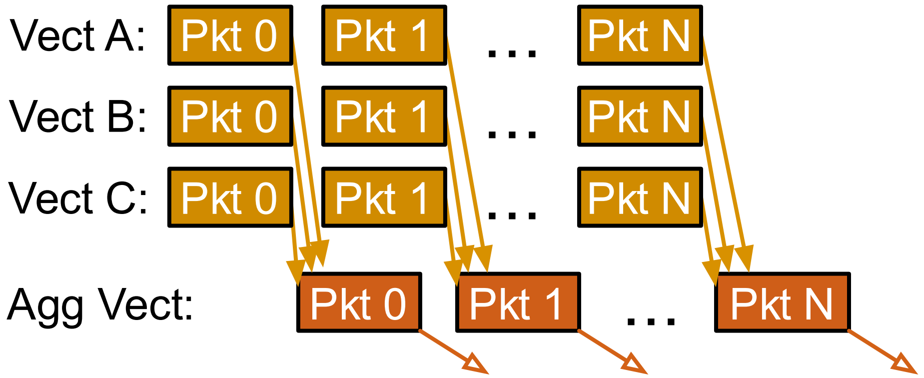

Packet-Level Aggregation in Our iSwitch

Aggregating Gradient at Packet-Level for Improved Parallelism

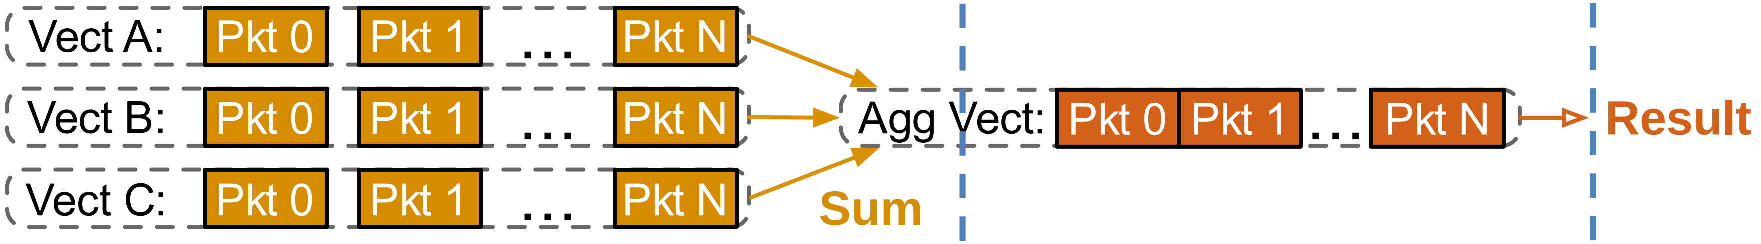

Conventional Vector-Level Aggregation

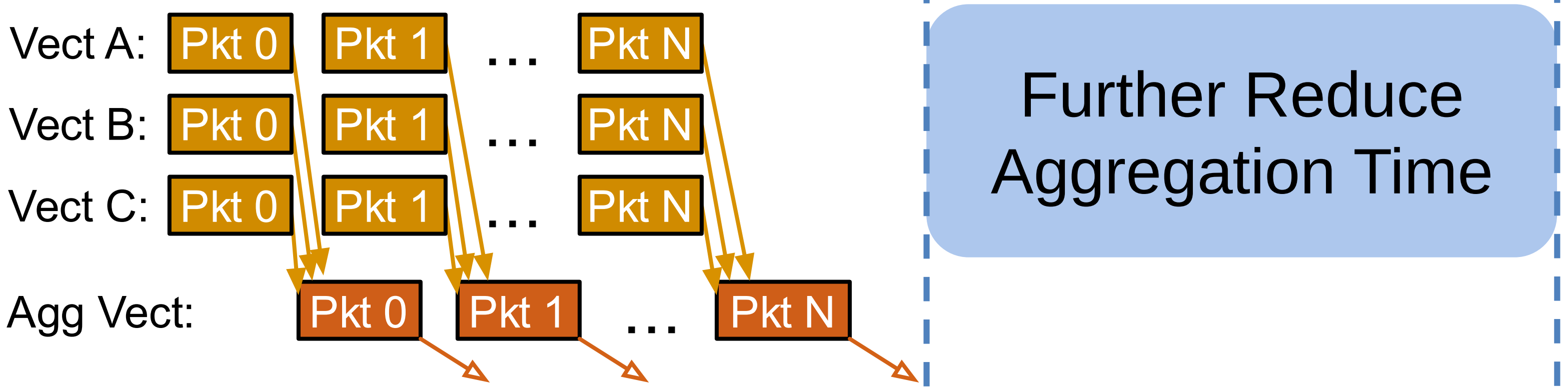

Packet-Level Aggregation in Our iSwitch

Extending Network Protocol for In-Switch Computing

Regular Packet:

ETH	IP	UDP	Application Data

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

ETH	IP	UDP	Application Data

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

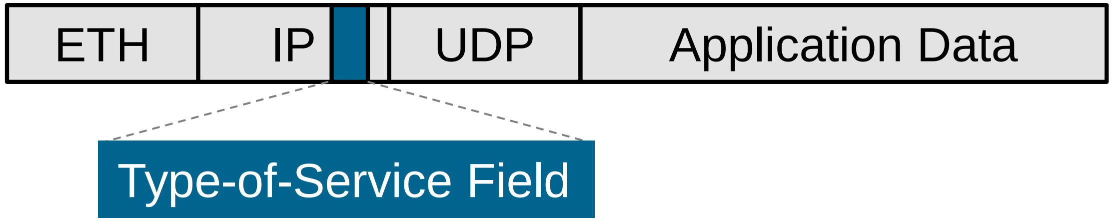

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

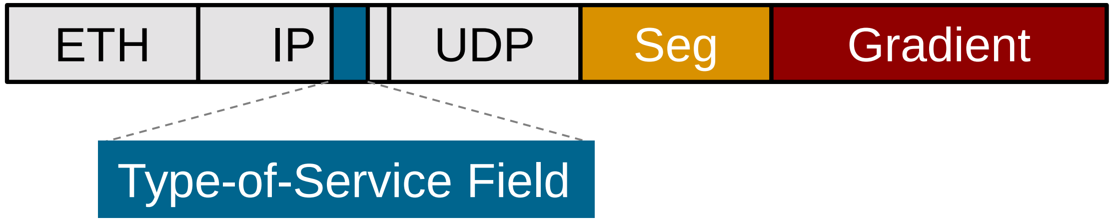

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

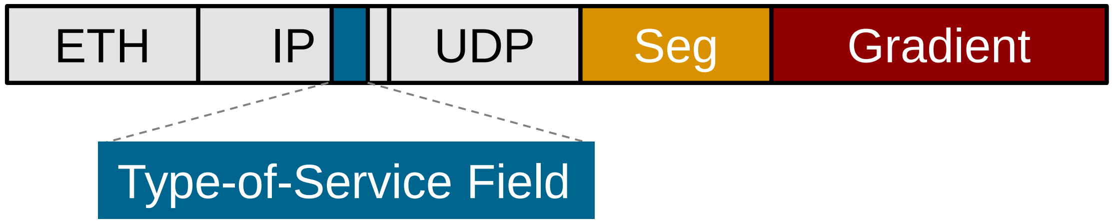

Control Packet of iSwitch:

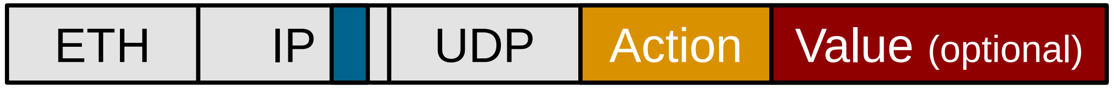

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

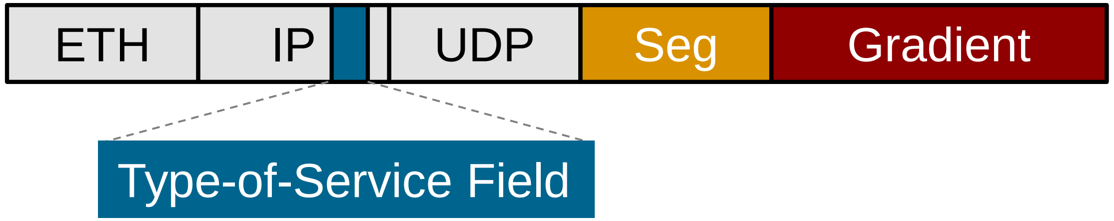

Control Packet of iSwitch:

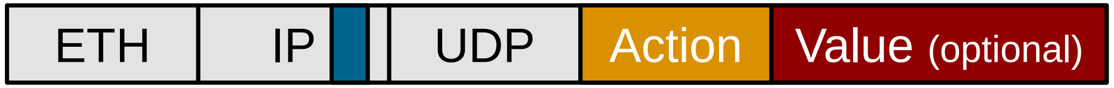

Action	Description
Join	Join the training job
Leave	Leave the training job
Reset	Clear the accelerator on the switch
SetH	Set aggregation threshold H on switch
FBcast	Force broadcast a segment on switch
Help	Request a lost data packet for a worker
Ack	Confirm the success of some actions

Extending Network Protocol for In-Switch Computing

Data Packet of iSwitch:

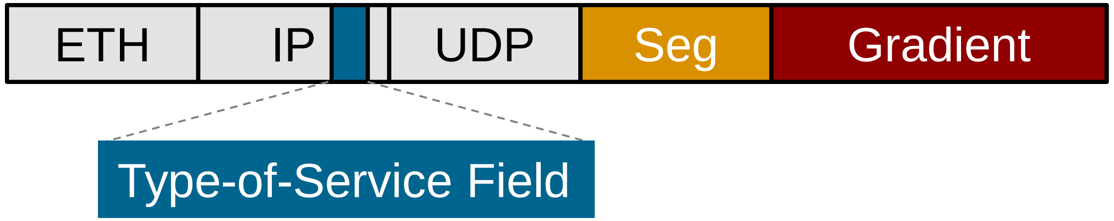

Control Packet of iSwitch:

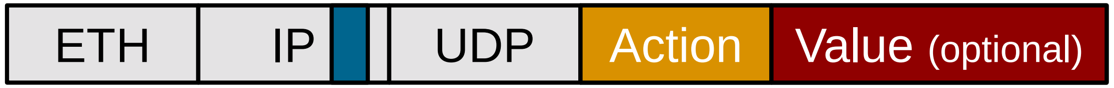

Action	Description
Join	Join the training job
Leave	Leave the training job
Reset	Clear the accelerator on the switch
SetH	Set aggregation threshold H on switch
FBcast	Force broadcast a segment on switch
Help	Request a lost data packet for a worker
Ack	Confirm the success of some actions

iSwitch extension will NOT affect regular network functions

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

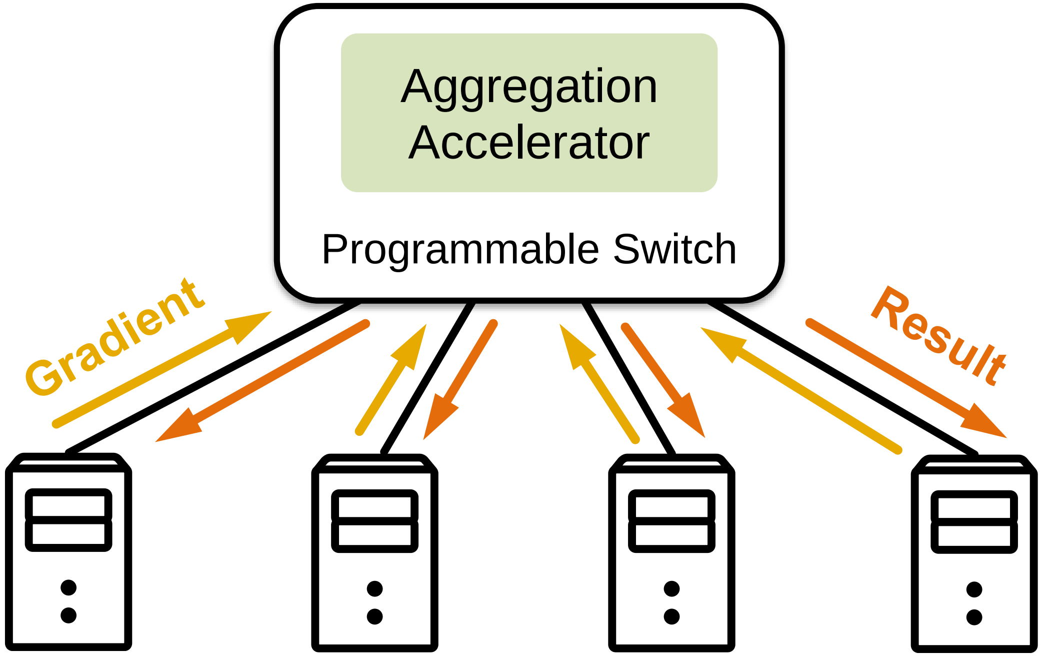

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

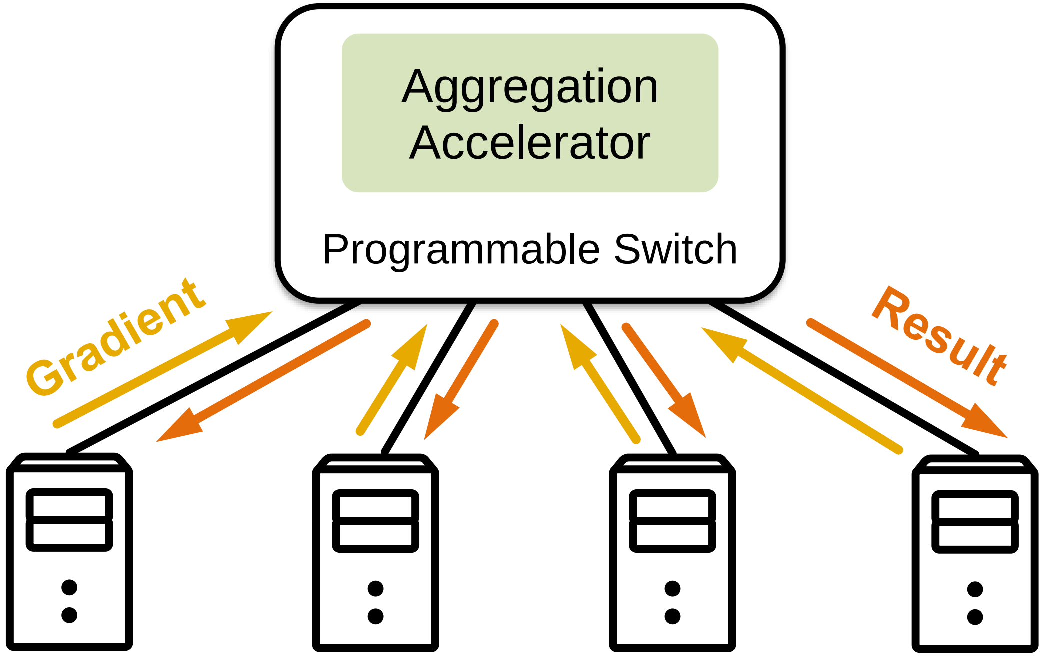

In-Switch Acceleration Directly Applies

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

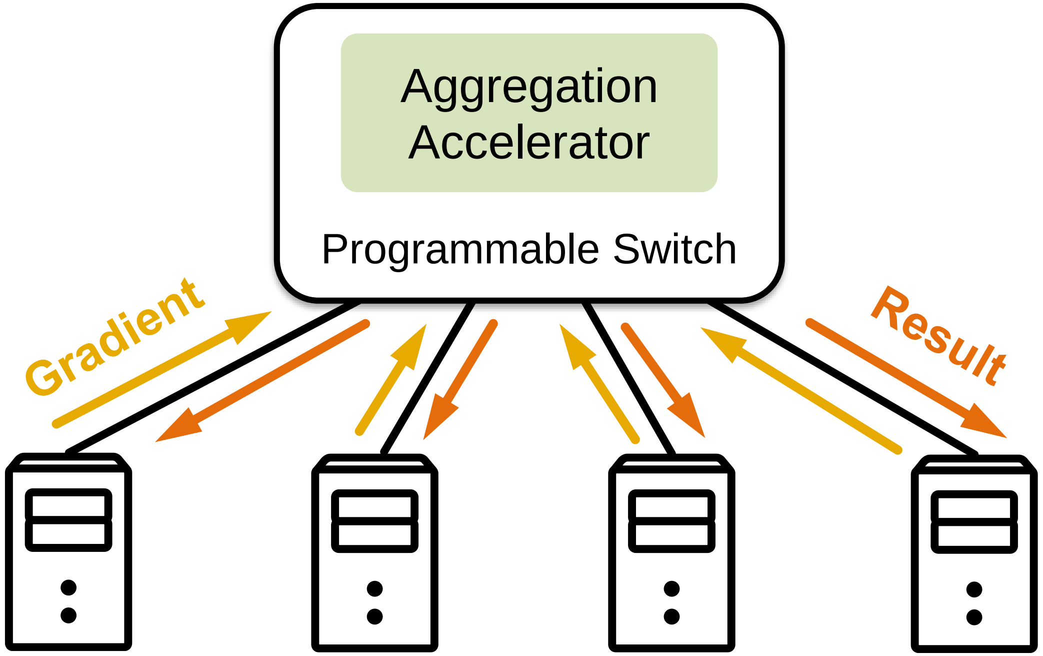

In-Switch Acceleration Directly Applies

Asynchronous Distributed Training

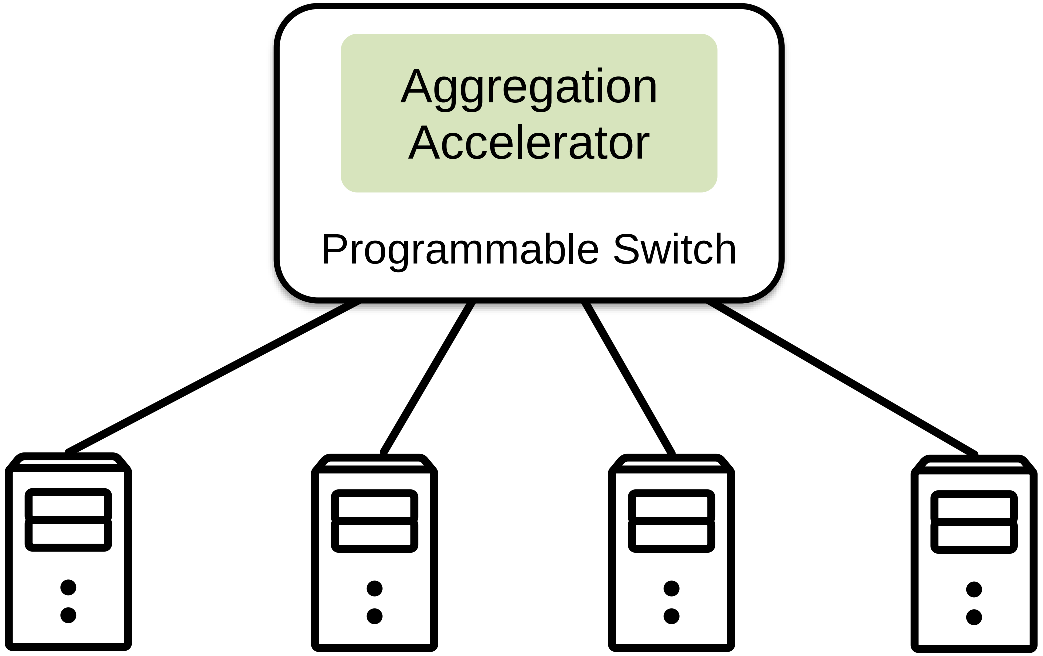

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

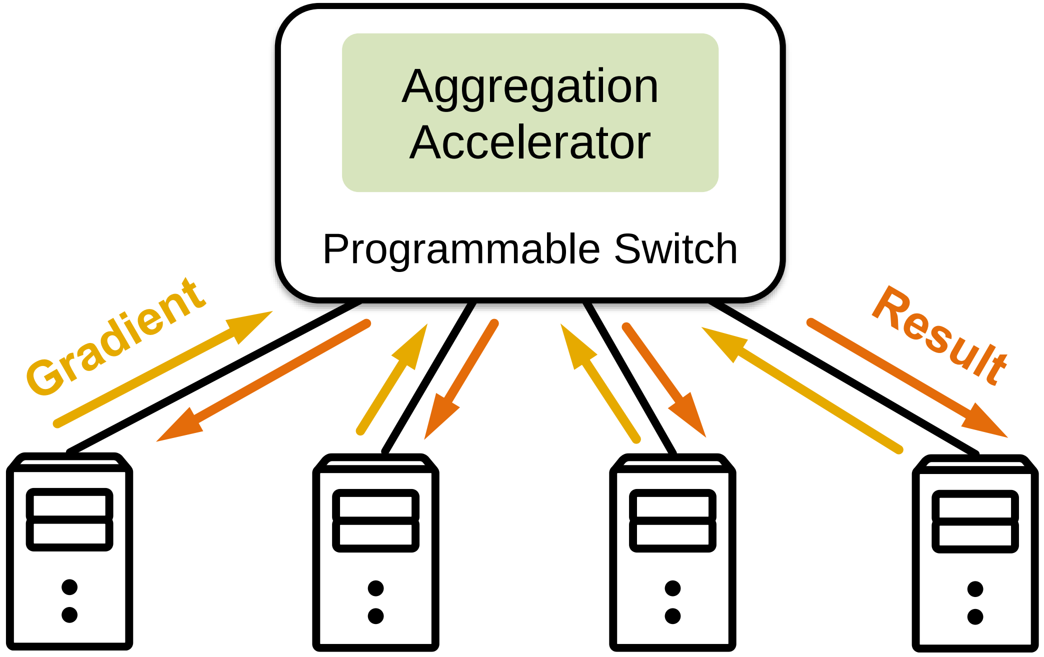

In-Switch Acceleration Directly Applies

Asynchronous Distributed Training

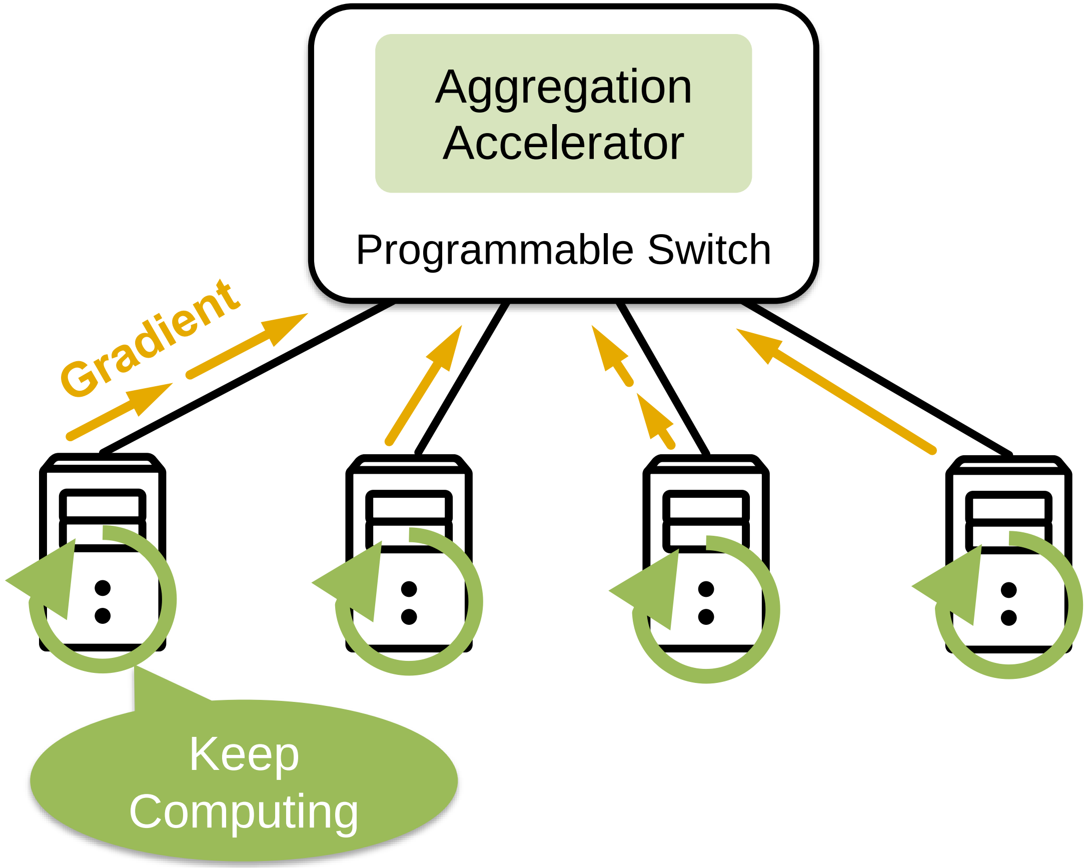

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

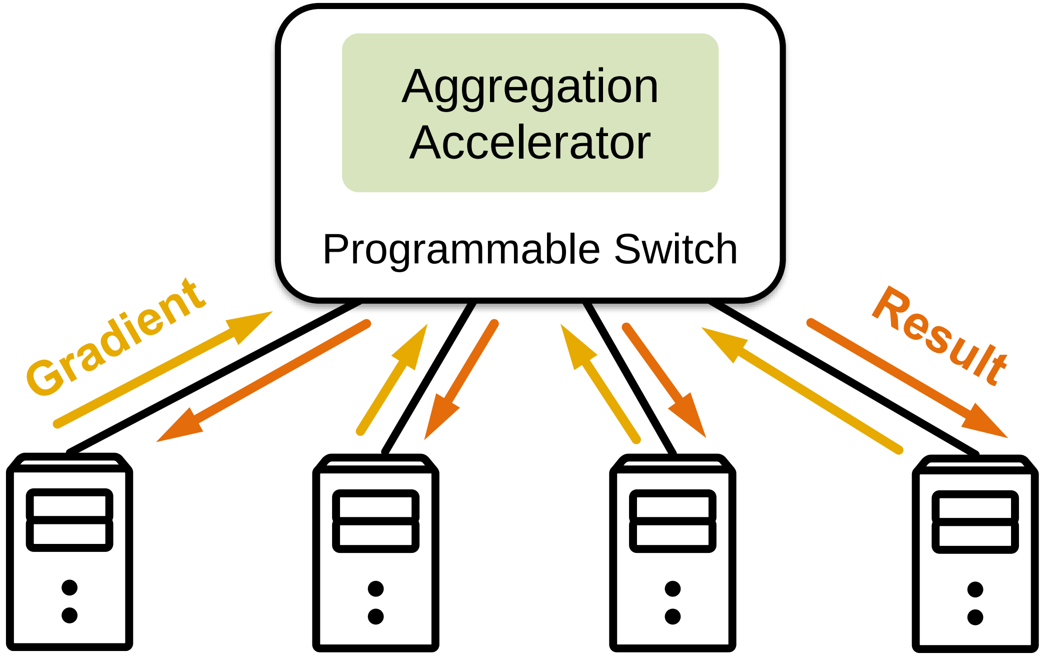

In-Switch Acceleration Directly Applies

Asynchronous Distributed Training

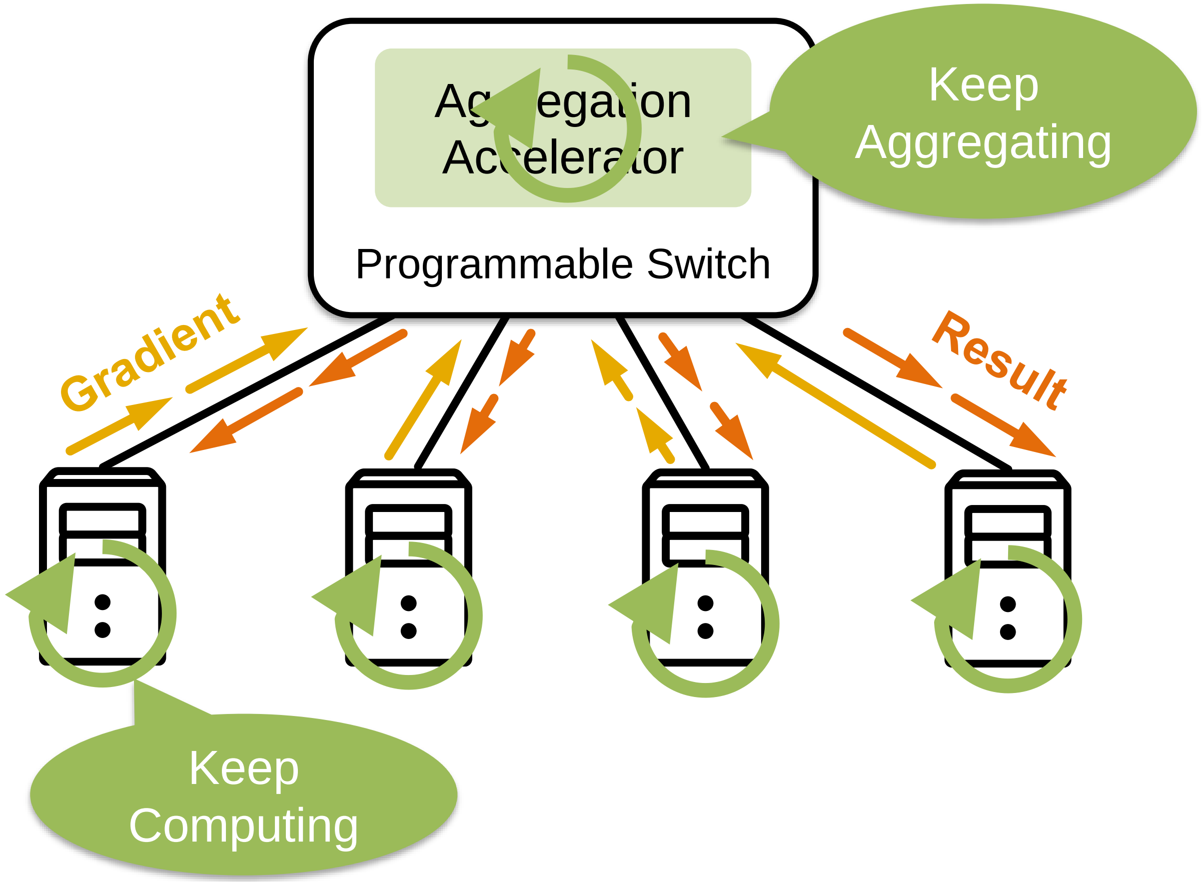

Supporting Different (Sync & Async) Training Execution Modes

Synchronous Distributed Training

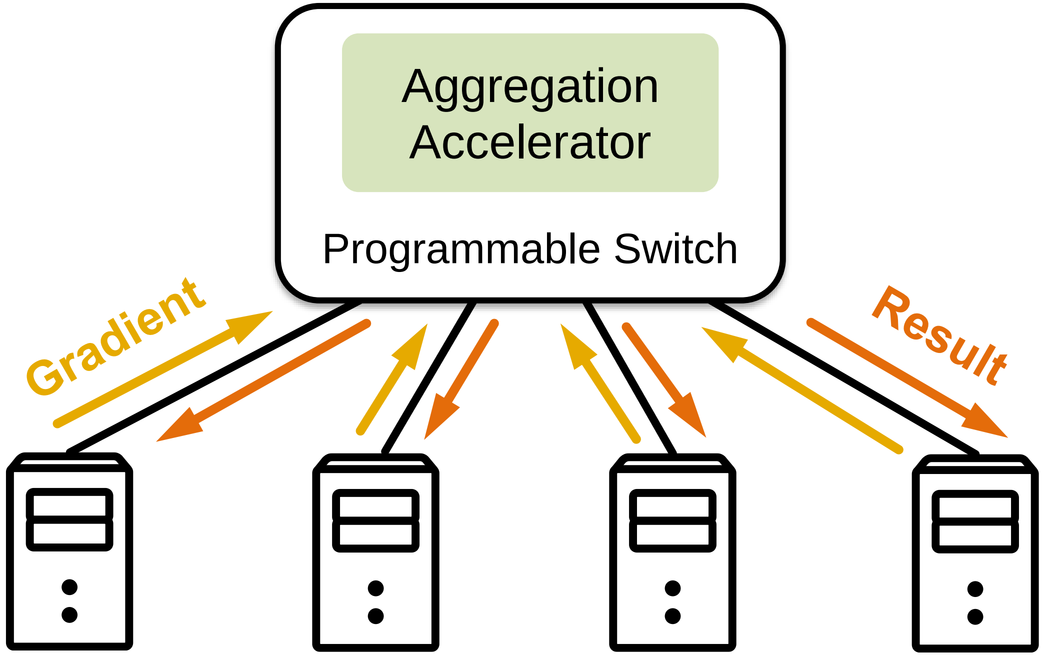

Asynchronous Distributed Training

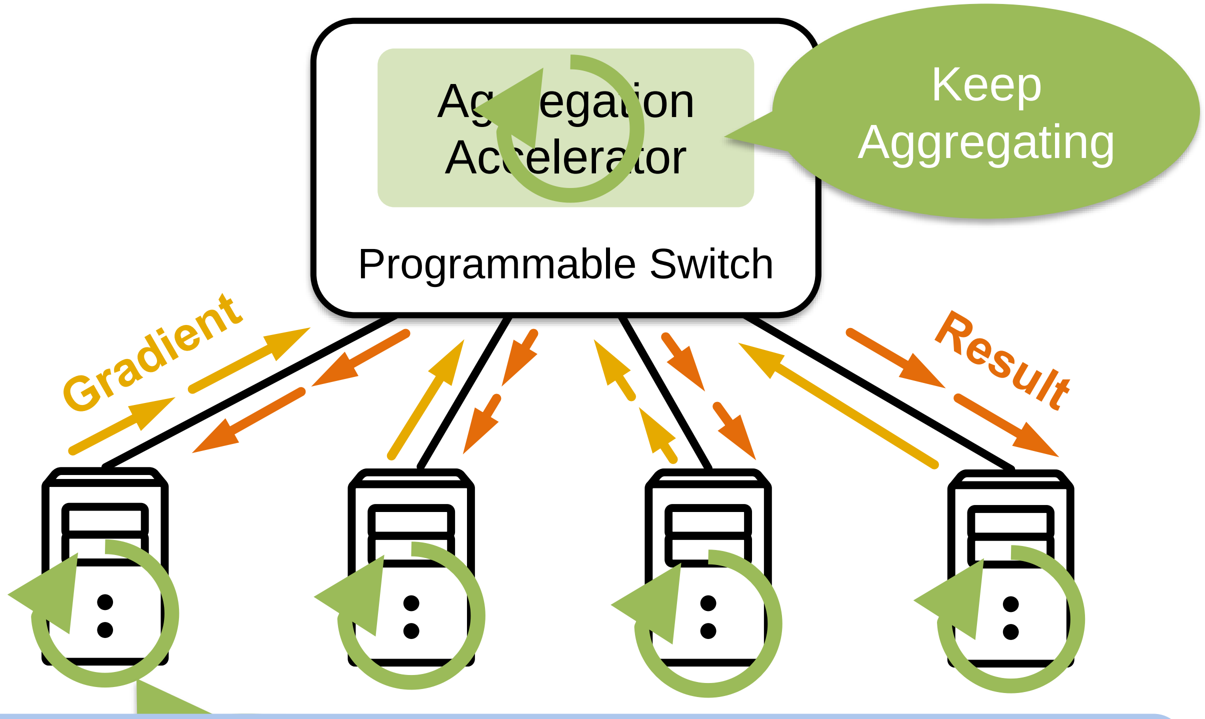

In-Switch Acceleration Directly Applies

HW/Algo Co-Design For Improved Parallelism

Scaling In-Switch Computing in Rack-Scale Data Centers

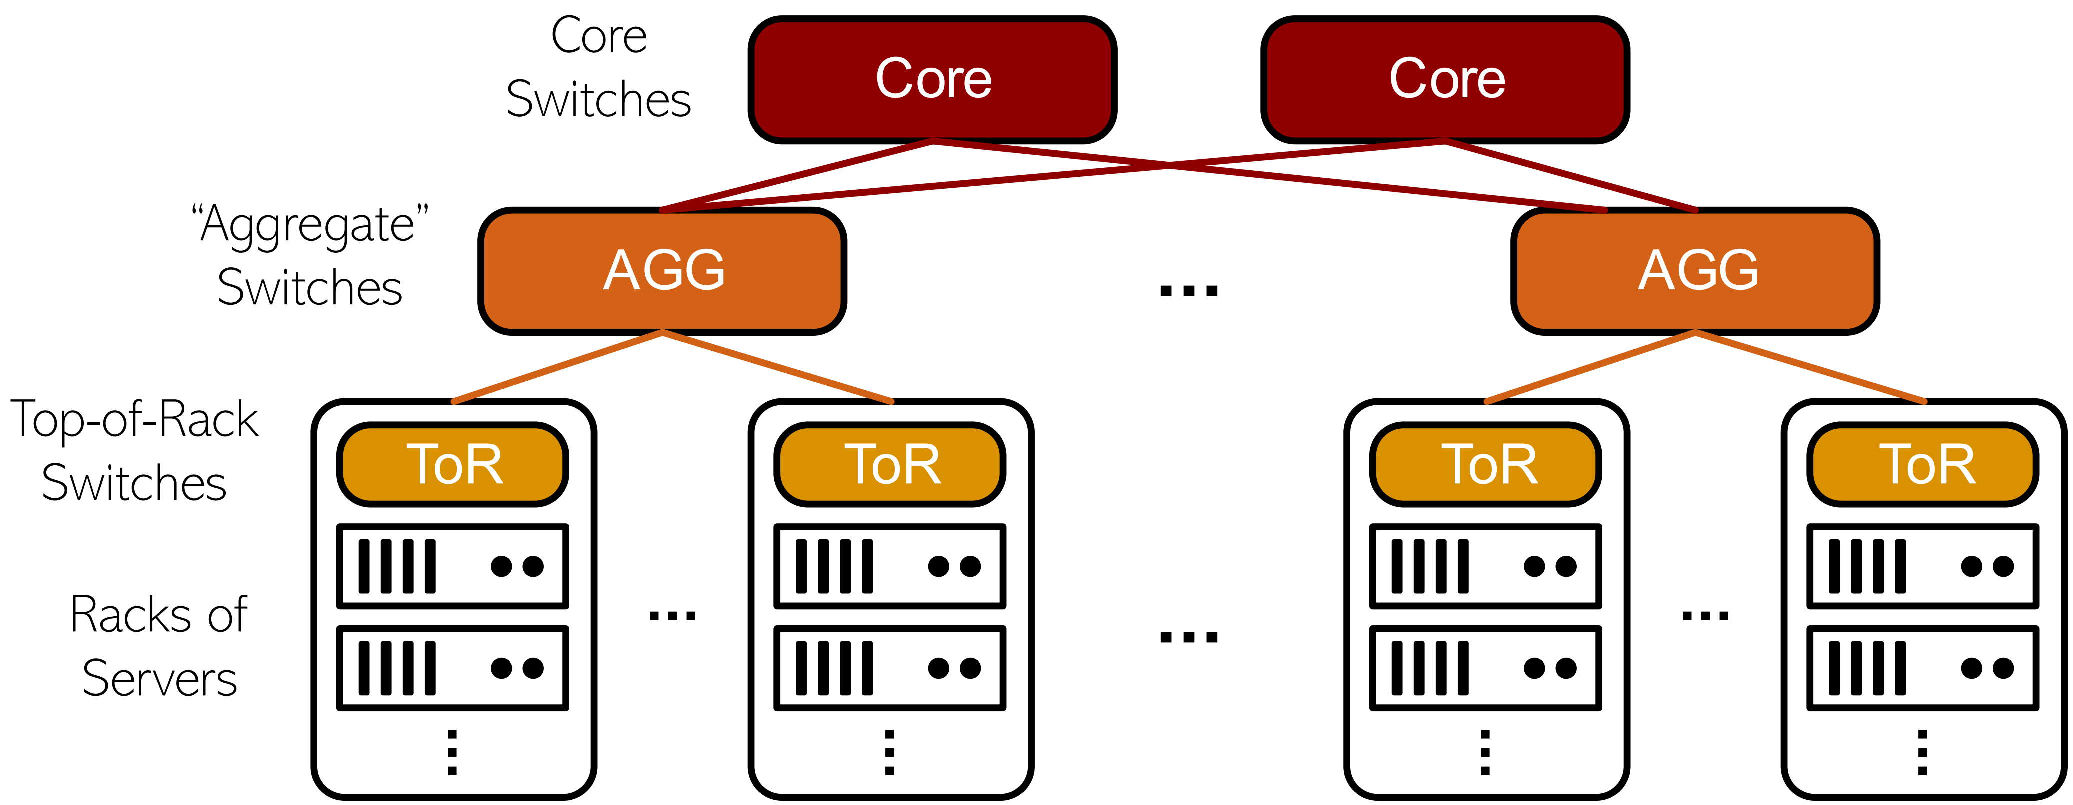

The Typical Network Architecture at Data Center

Scaling In-Switch Computing in Rack-Scale Data Centers

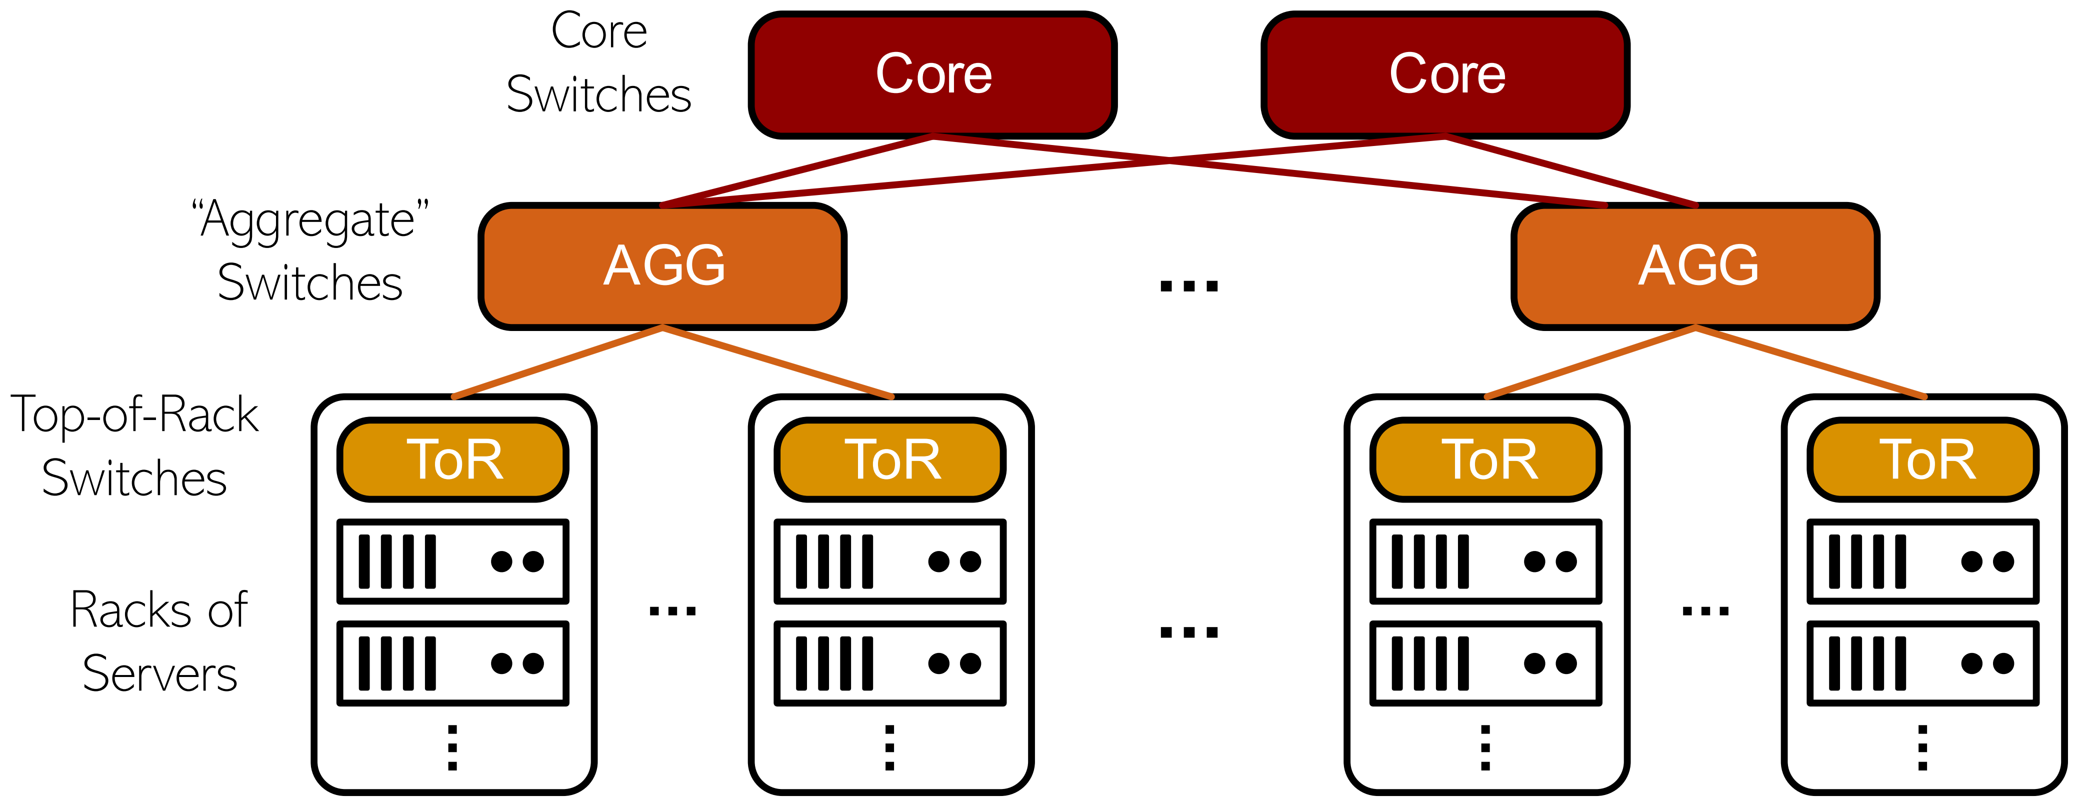

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

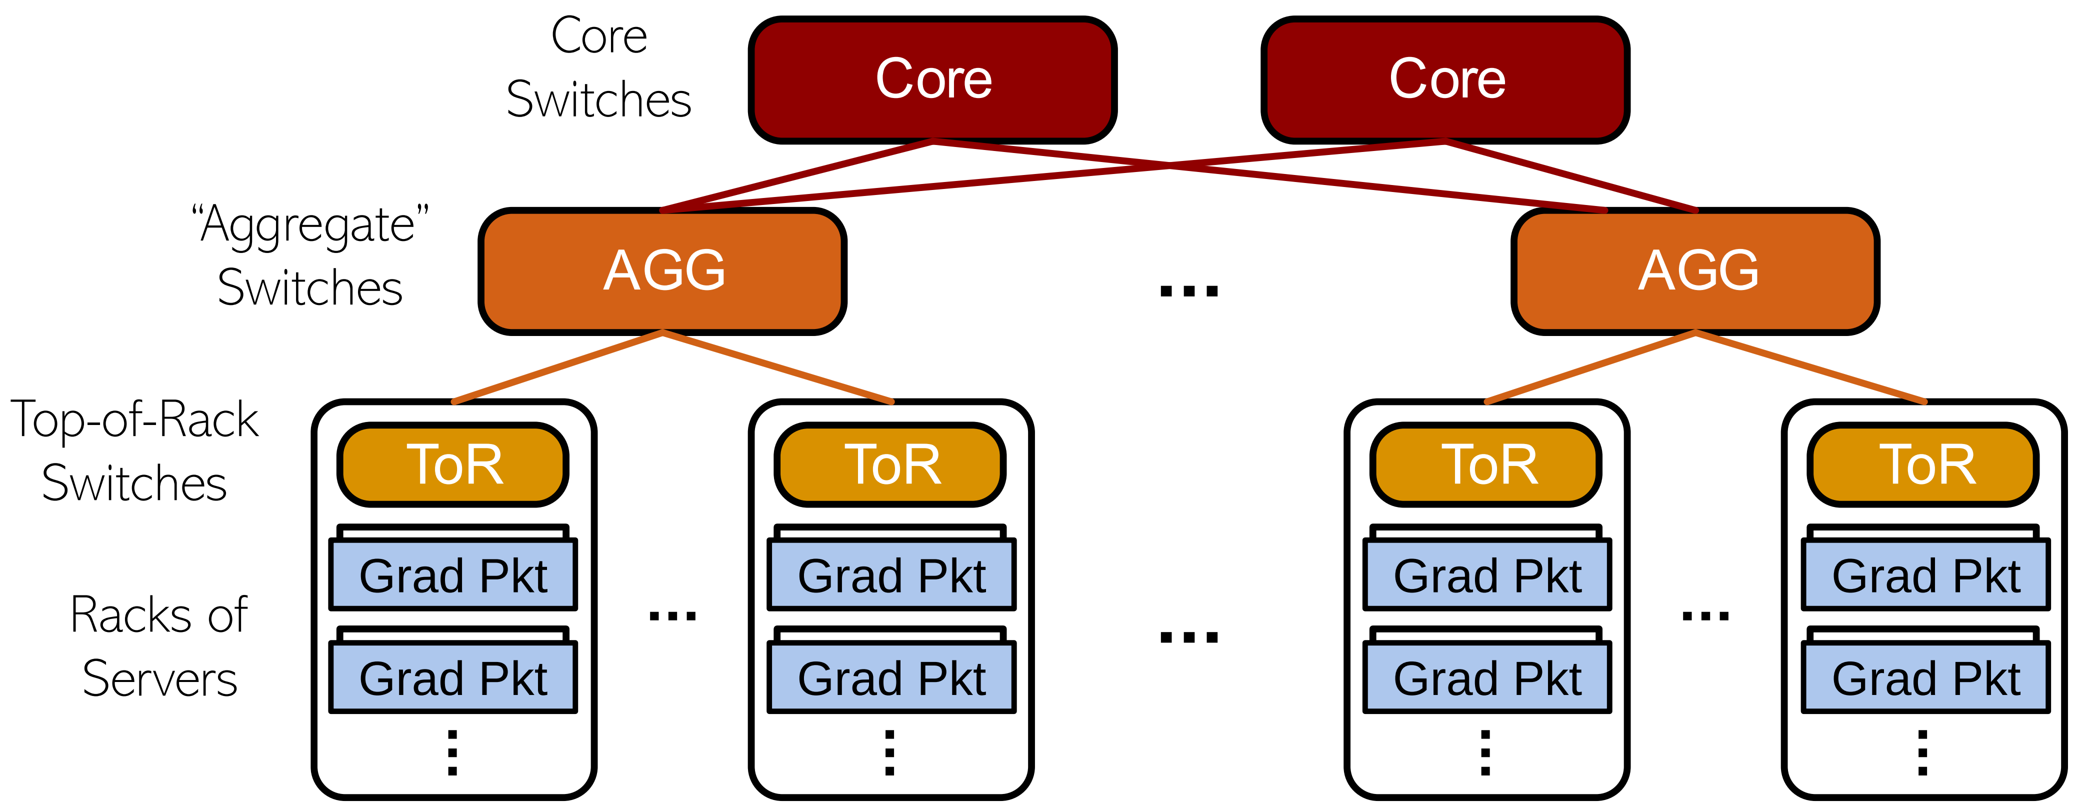

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

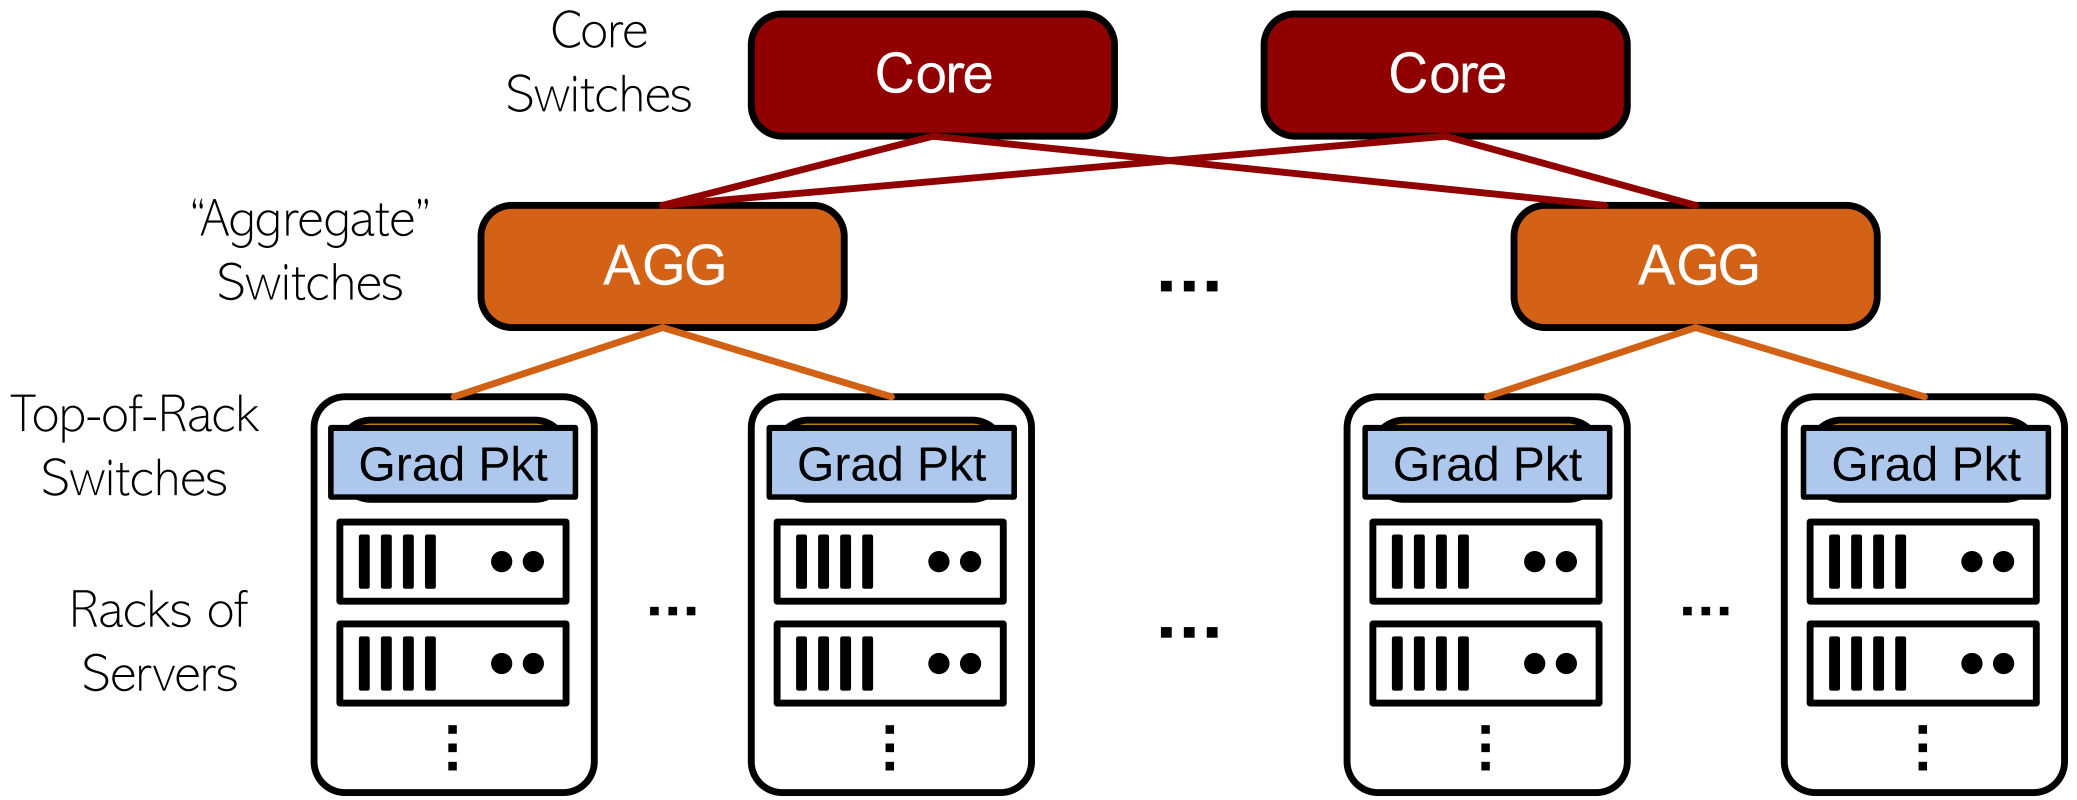

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

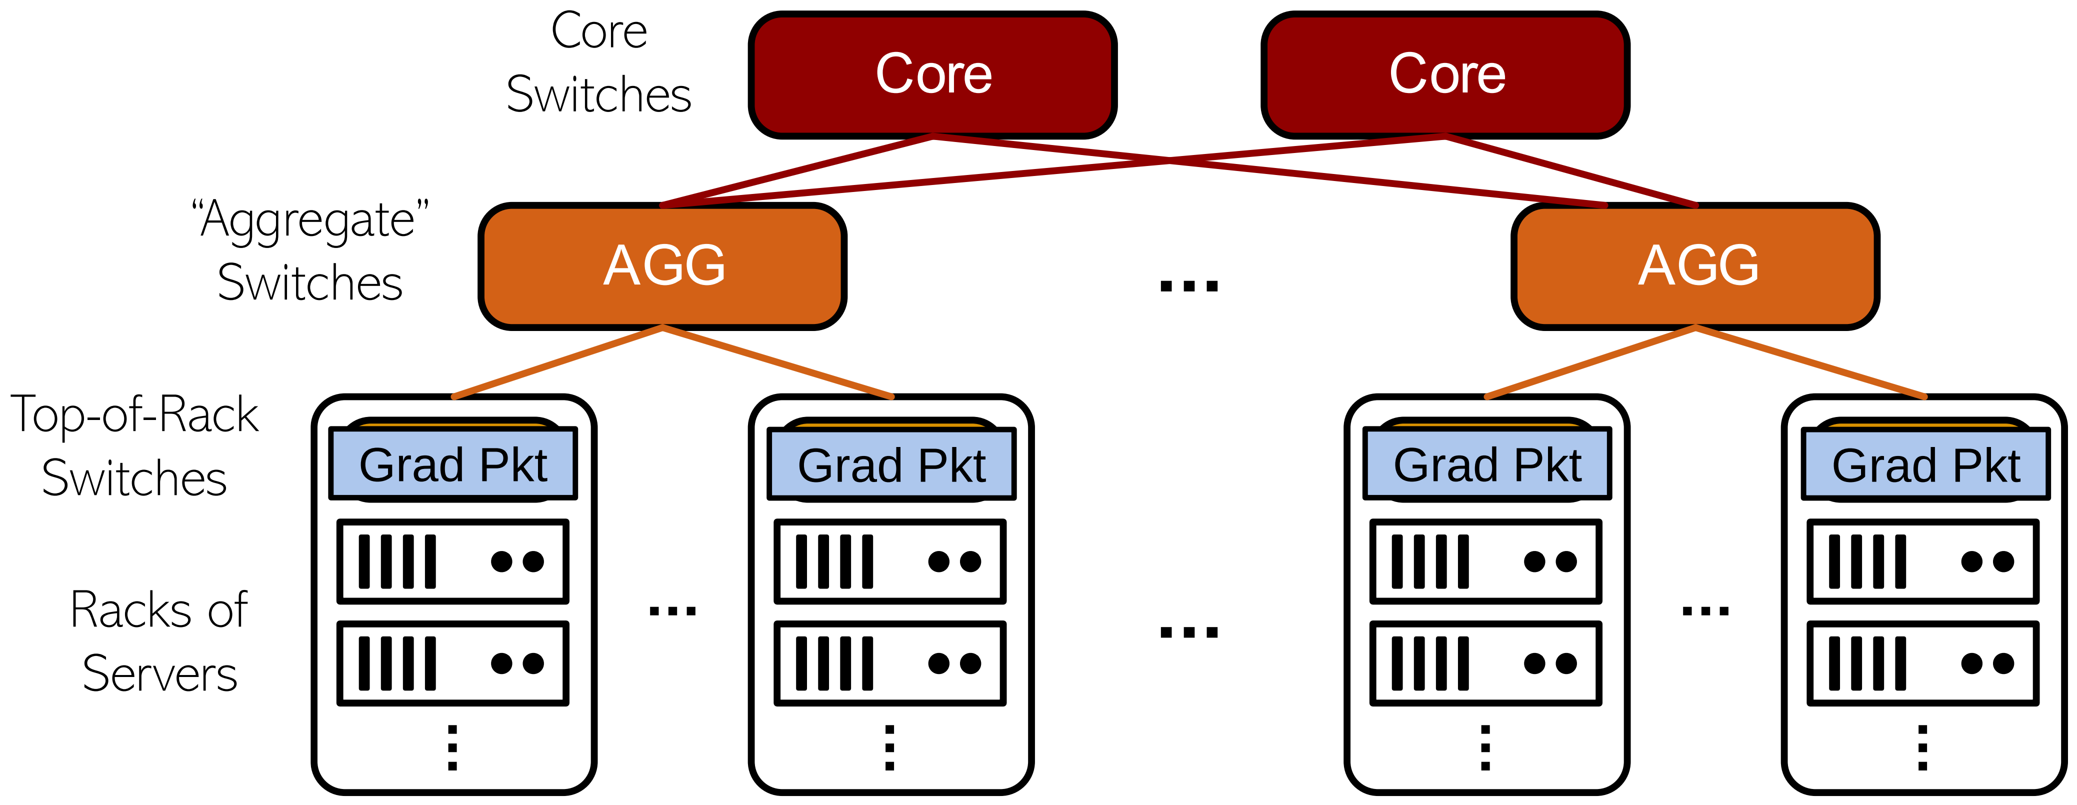

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

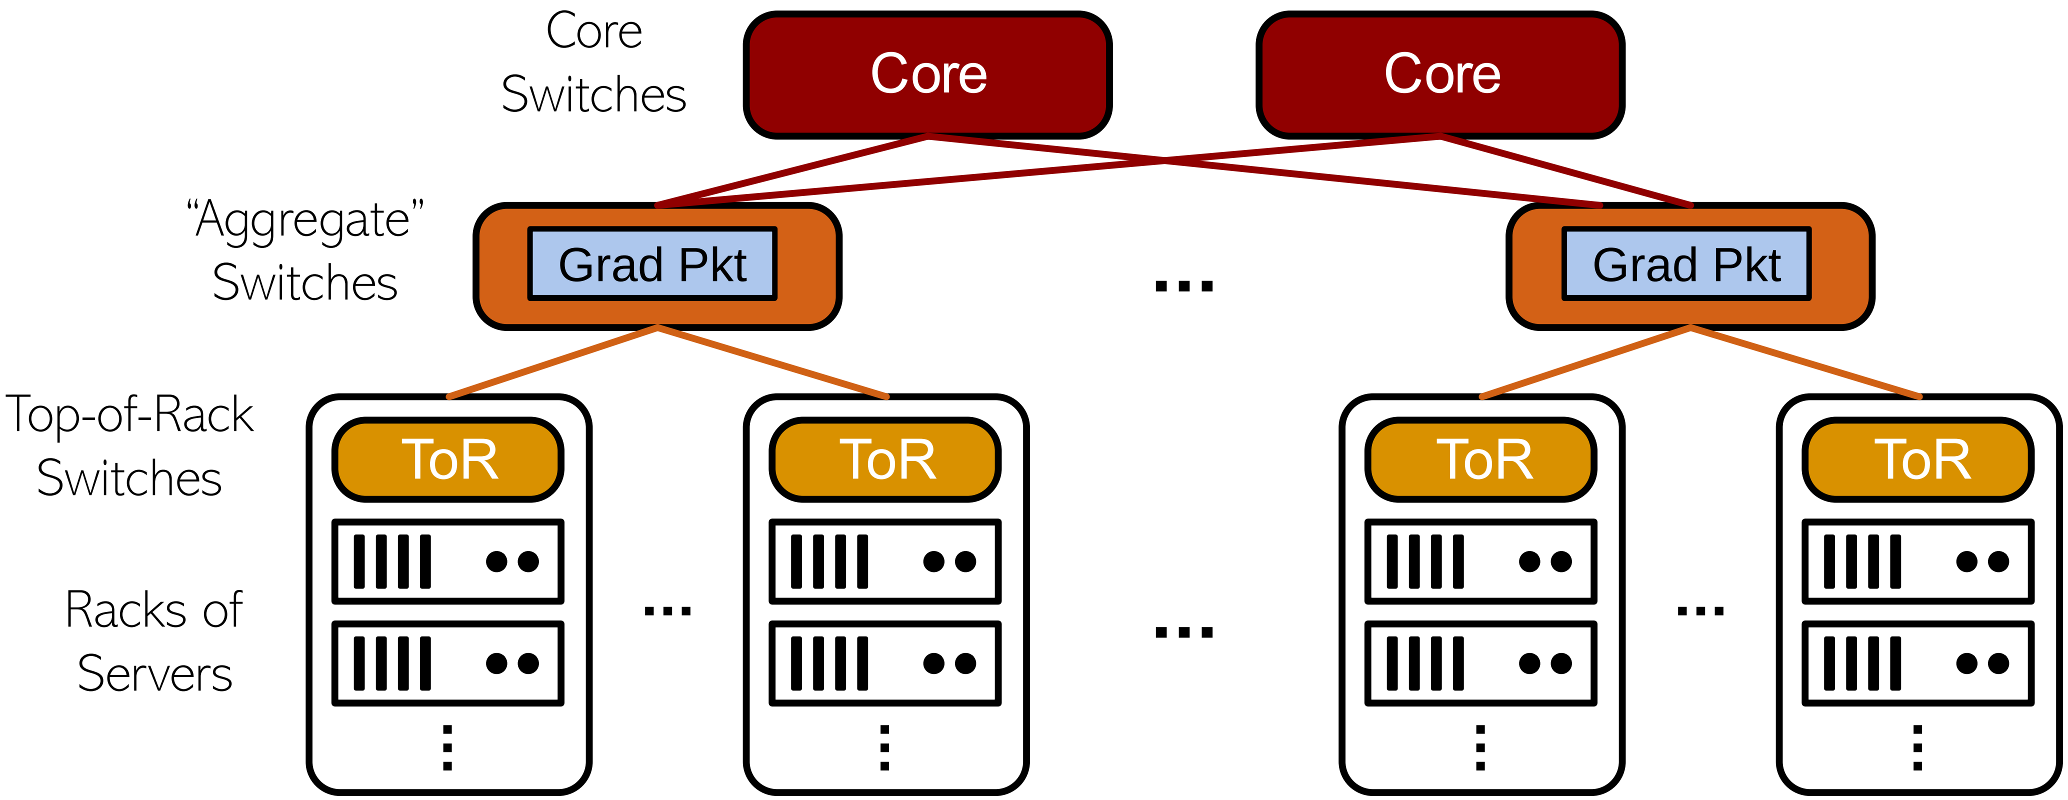

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

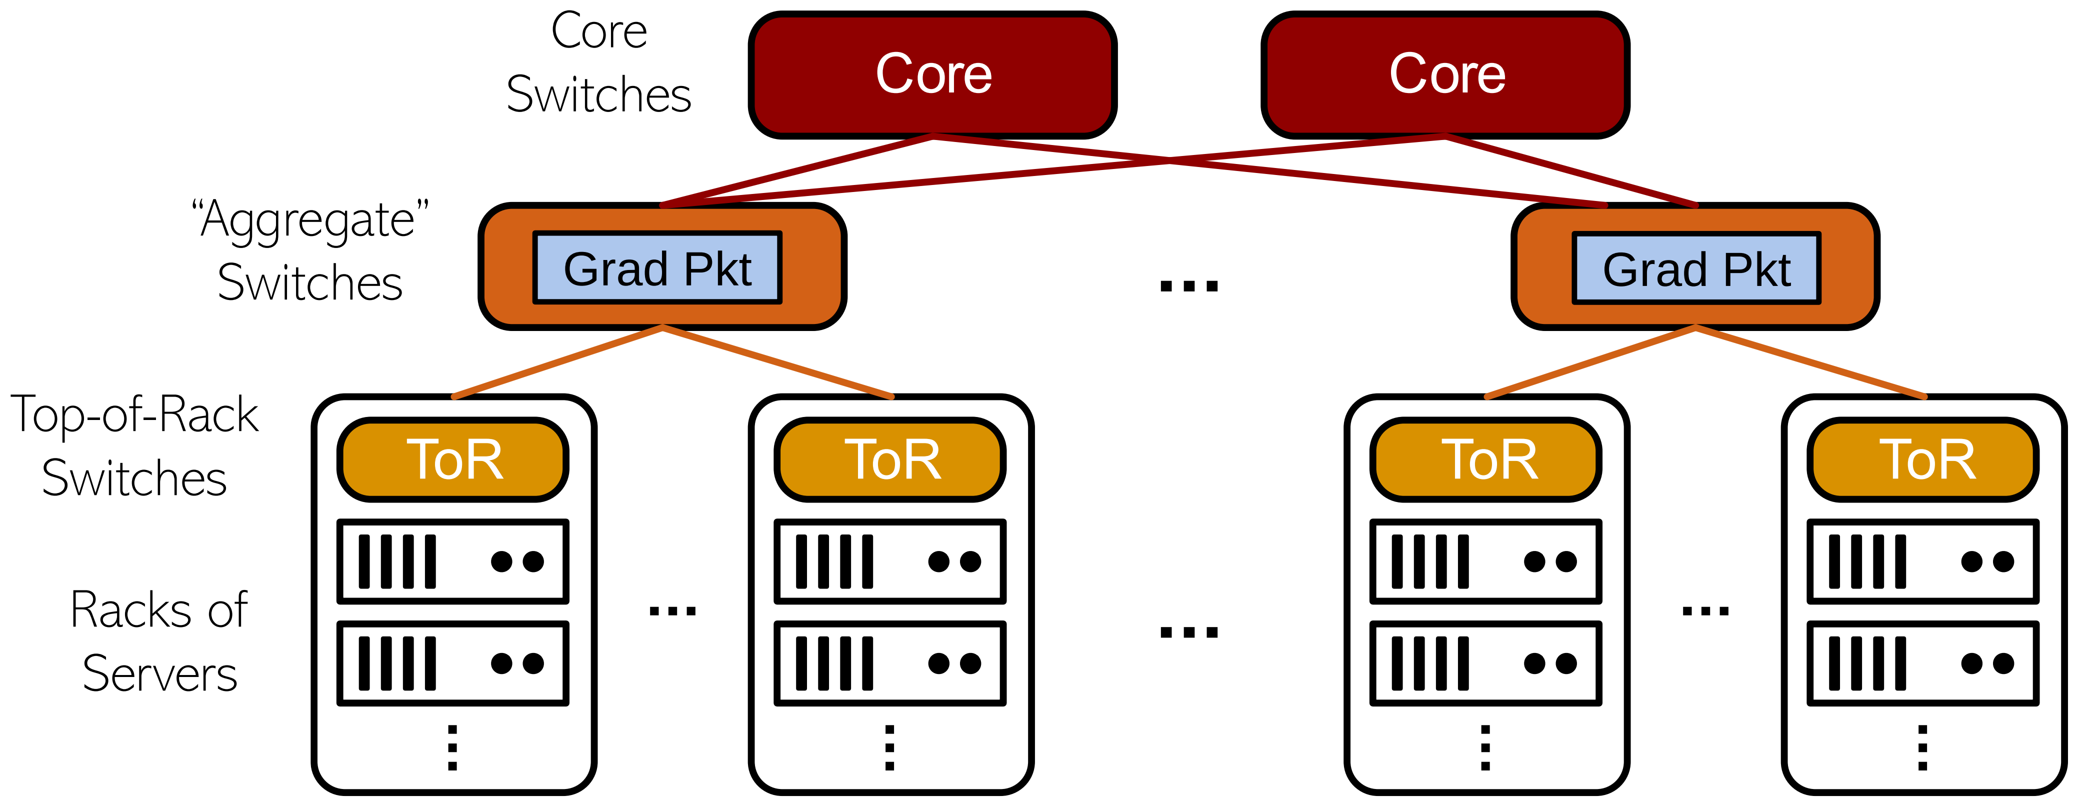

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

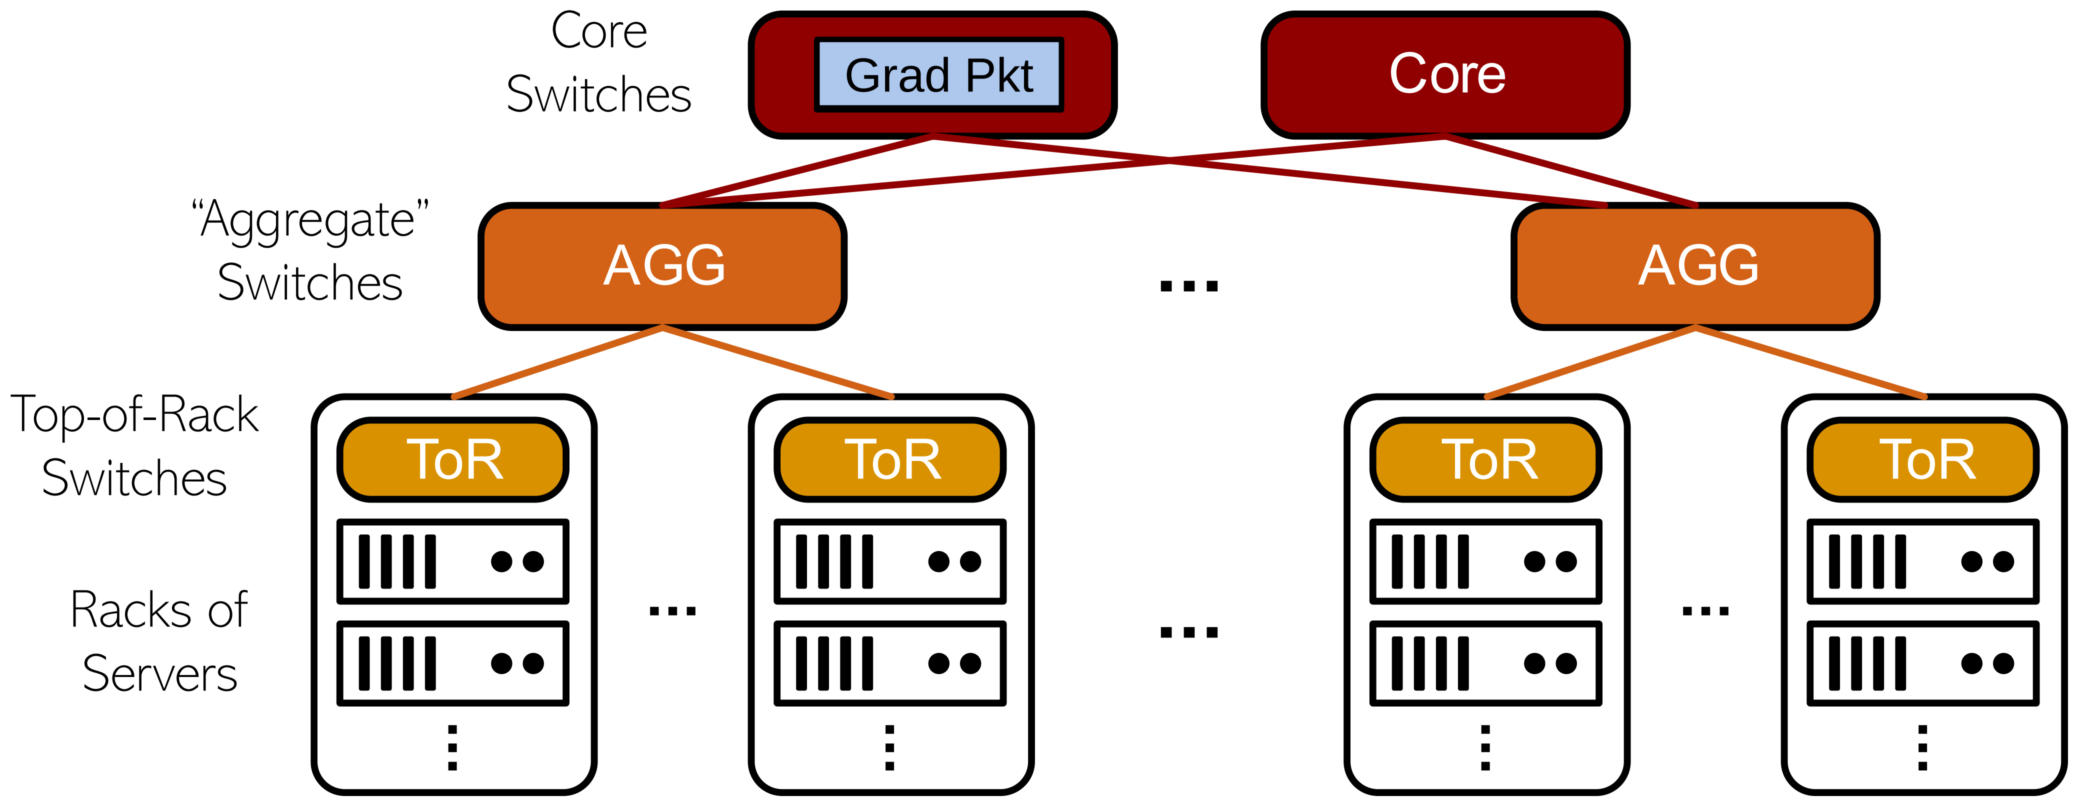

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

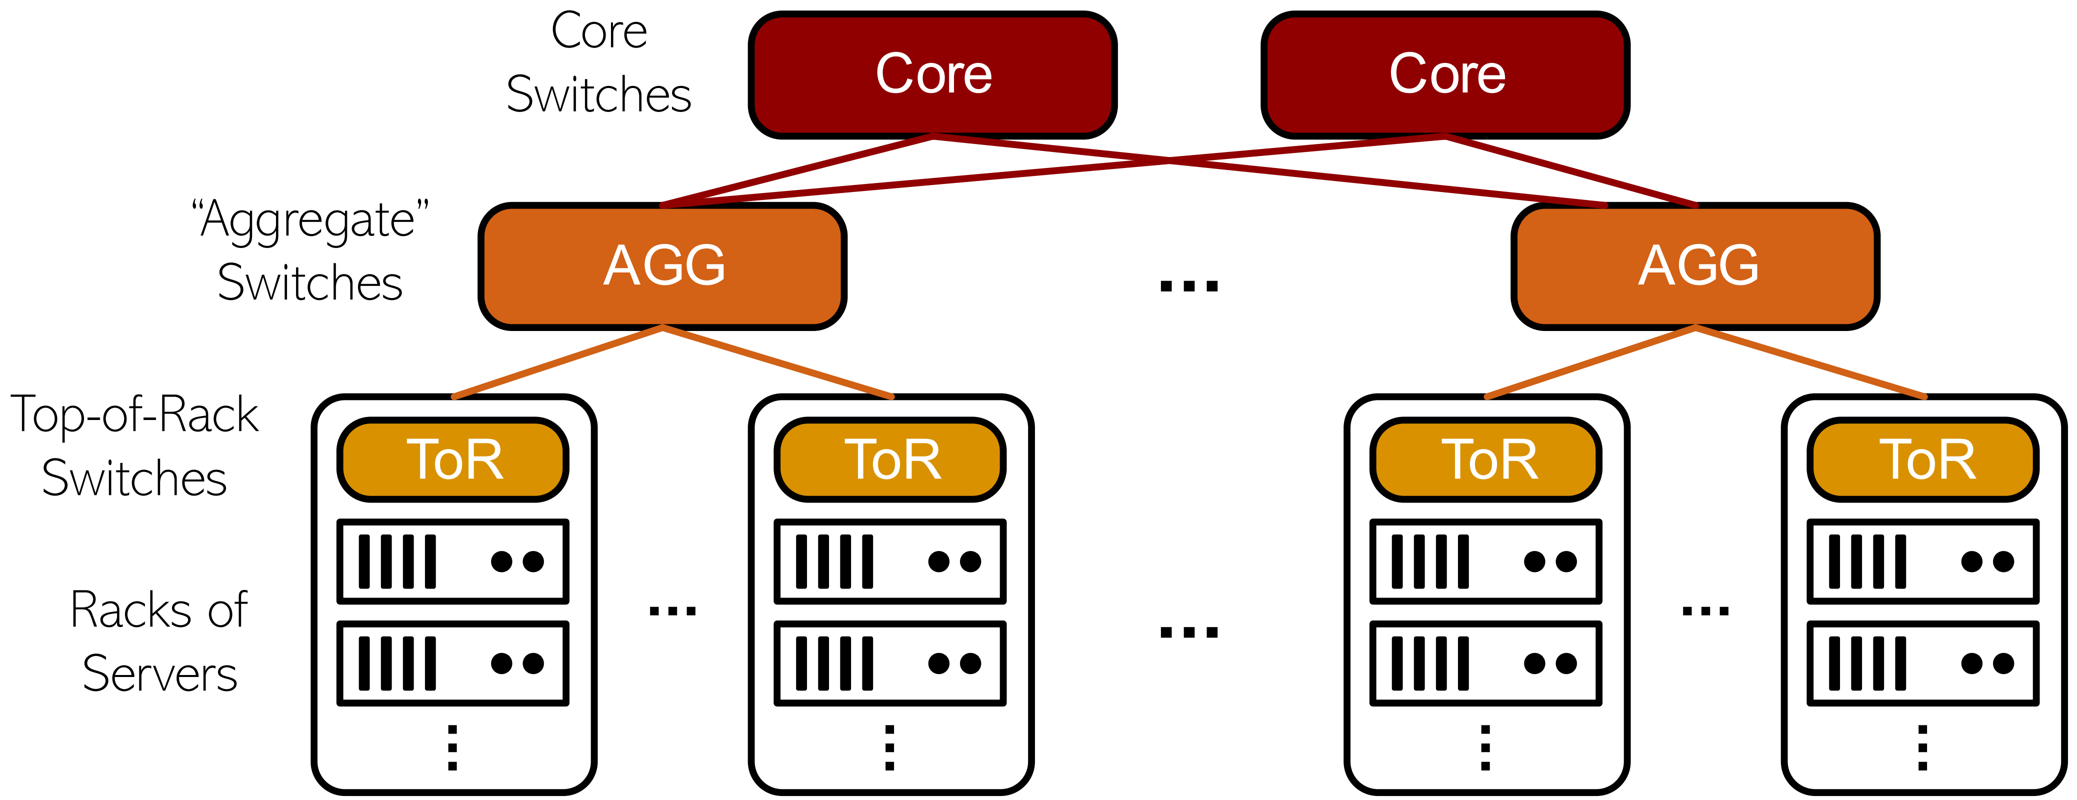

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

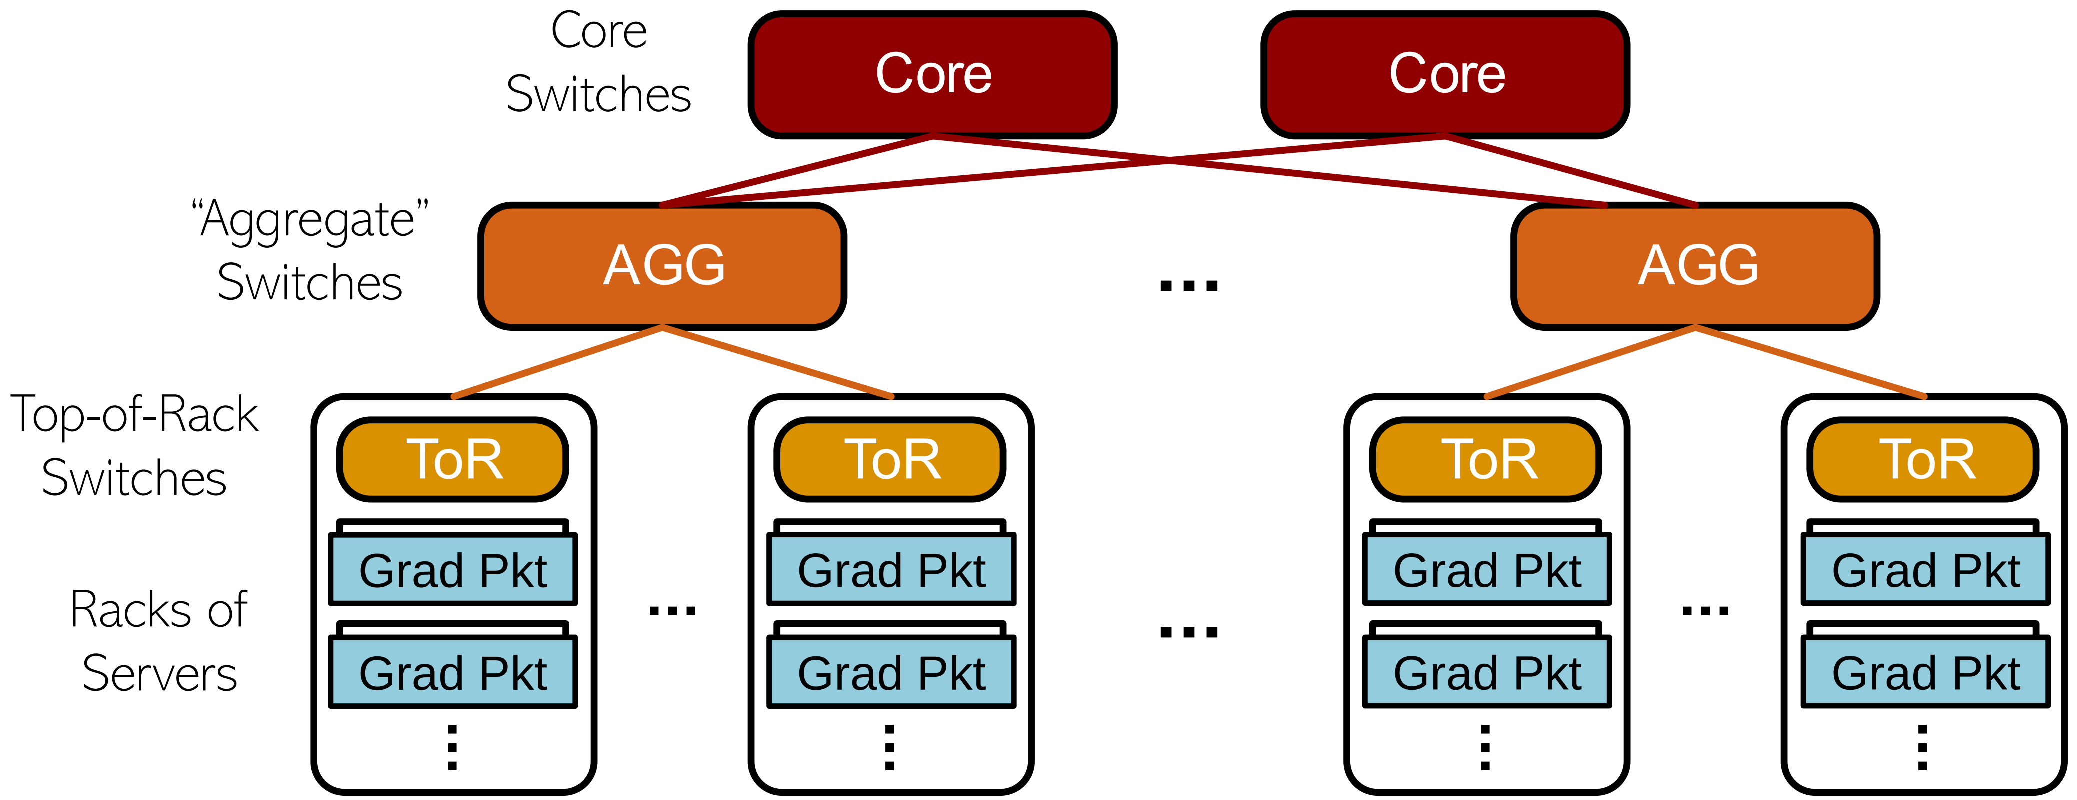

The Hierarchical Aggregation of iSwitch

Scaling In-Switch Computing in Rack-Scale Data Centers

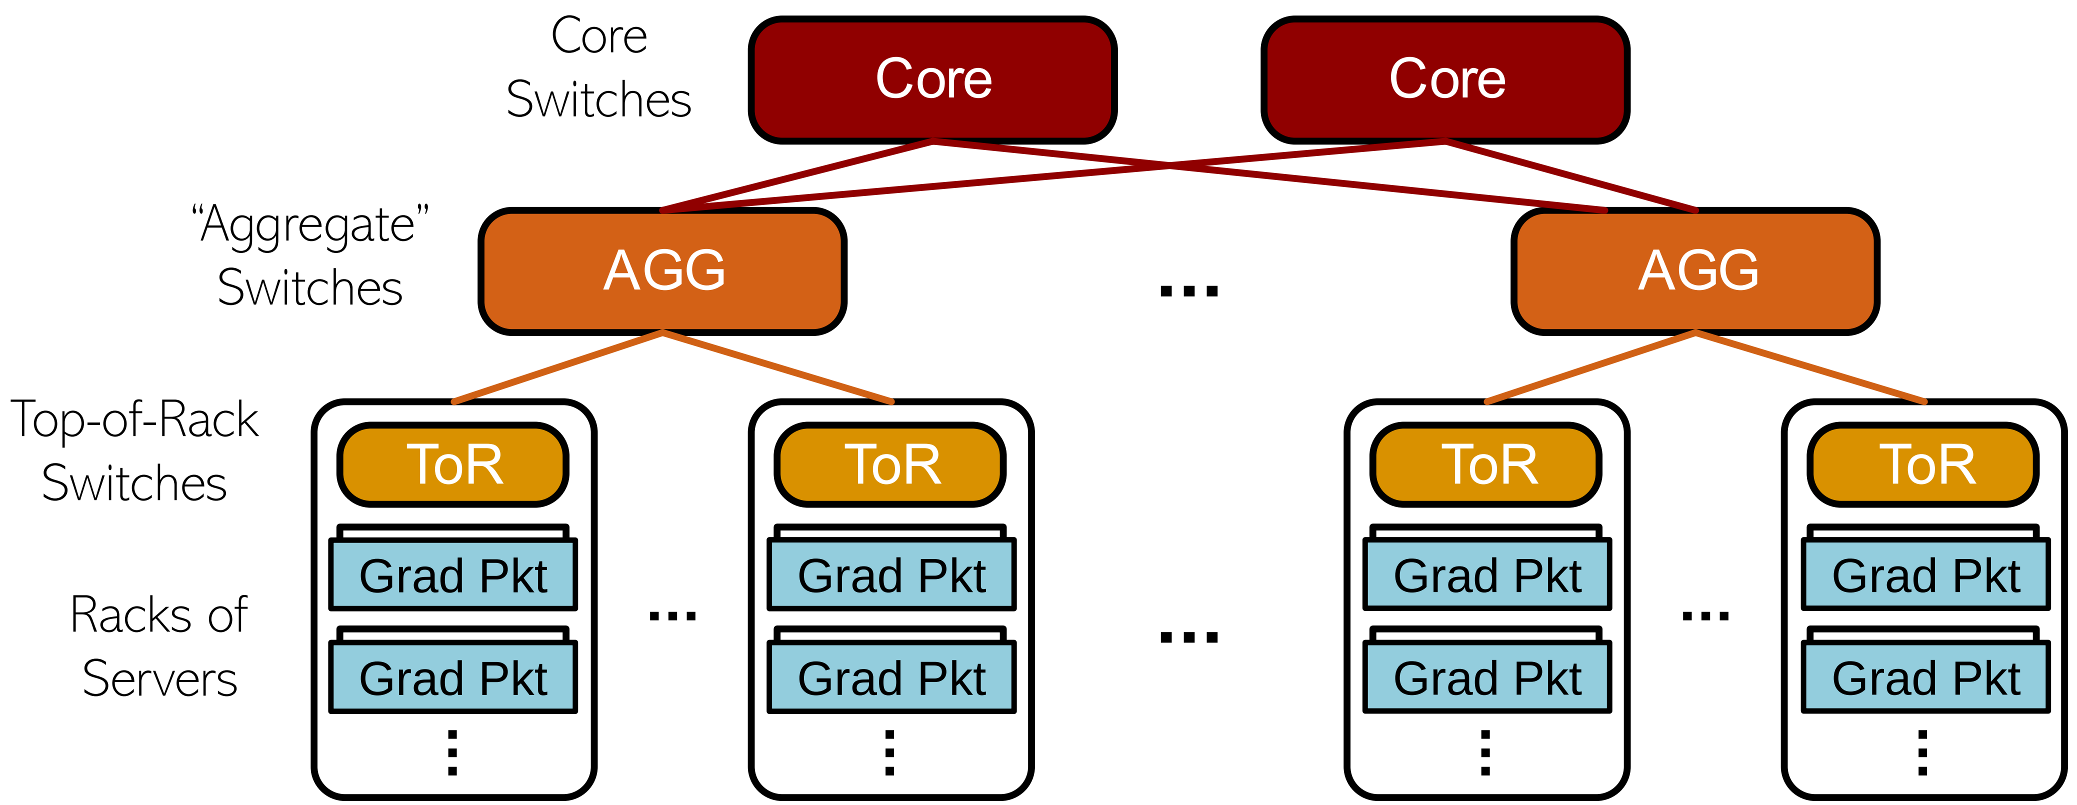

No Additional Cost or Topology Change for Scaling In-Switch Computing

In-Switch Computing Implementation

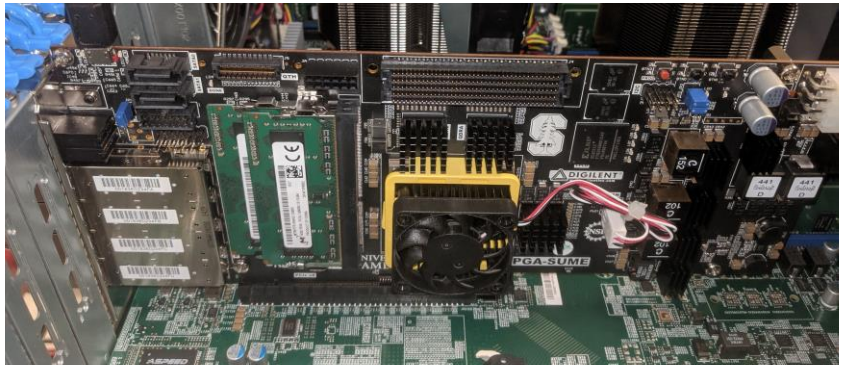

NetFPGA-SUME Board

GPU Cluster

RL Training Benchmarks

DQN

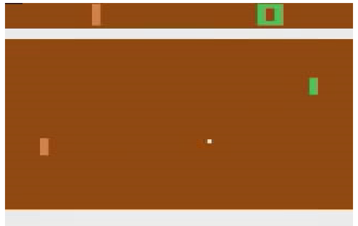

A2C

PPO

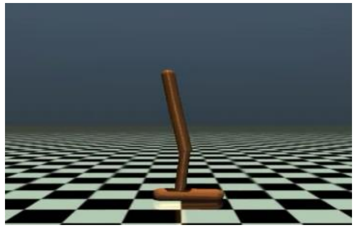

DDPG

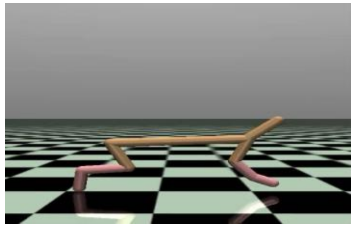

Reducing the End-to-End Training Time with iSwitch

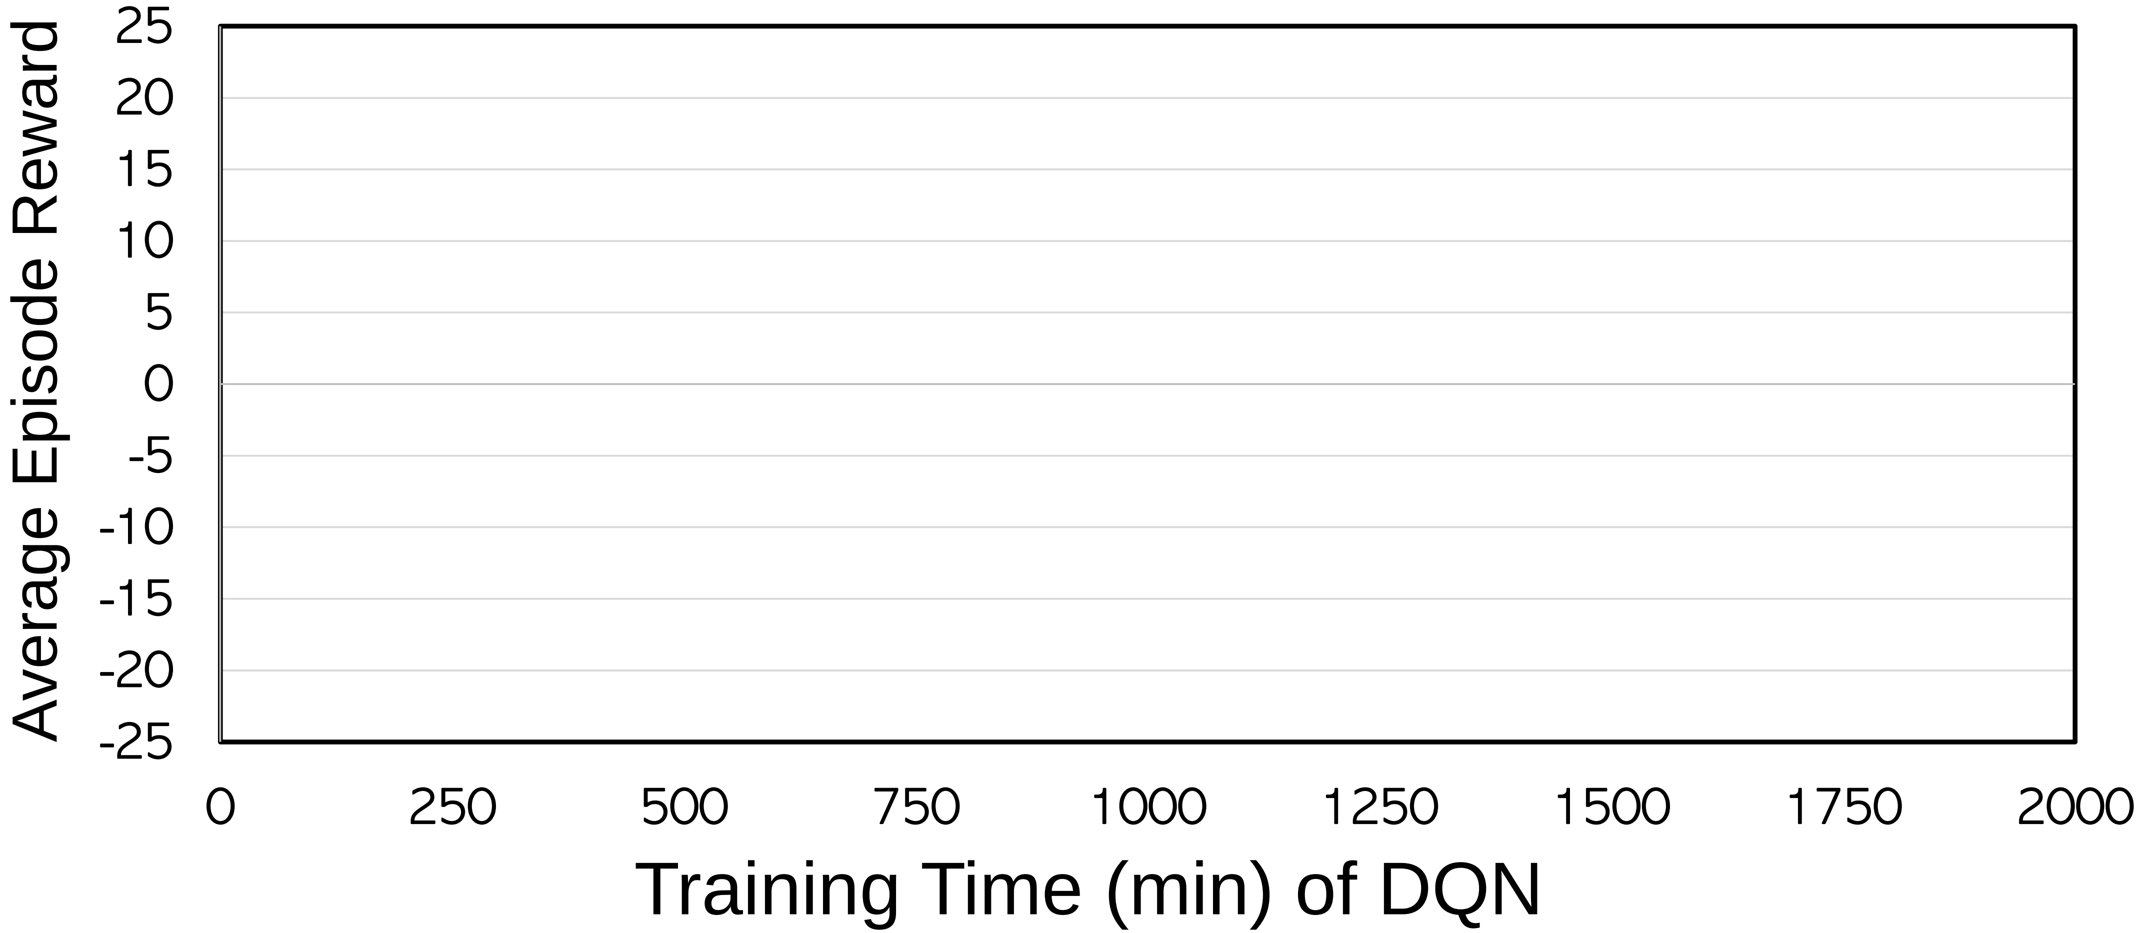

Reducing the End-to-End Training Time with iSwitch

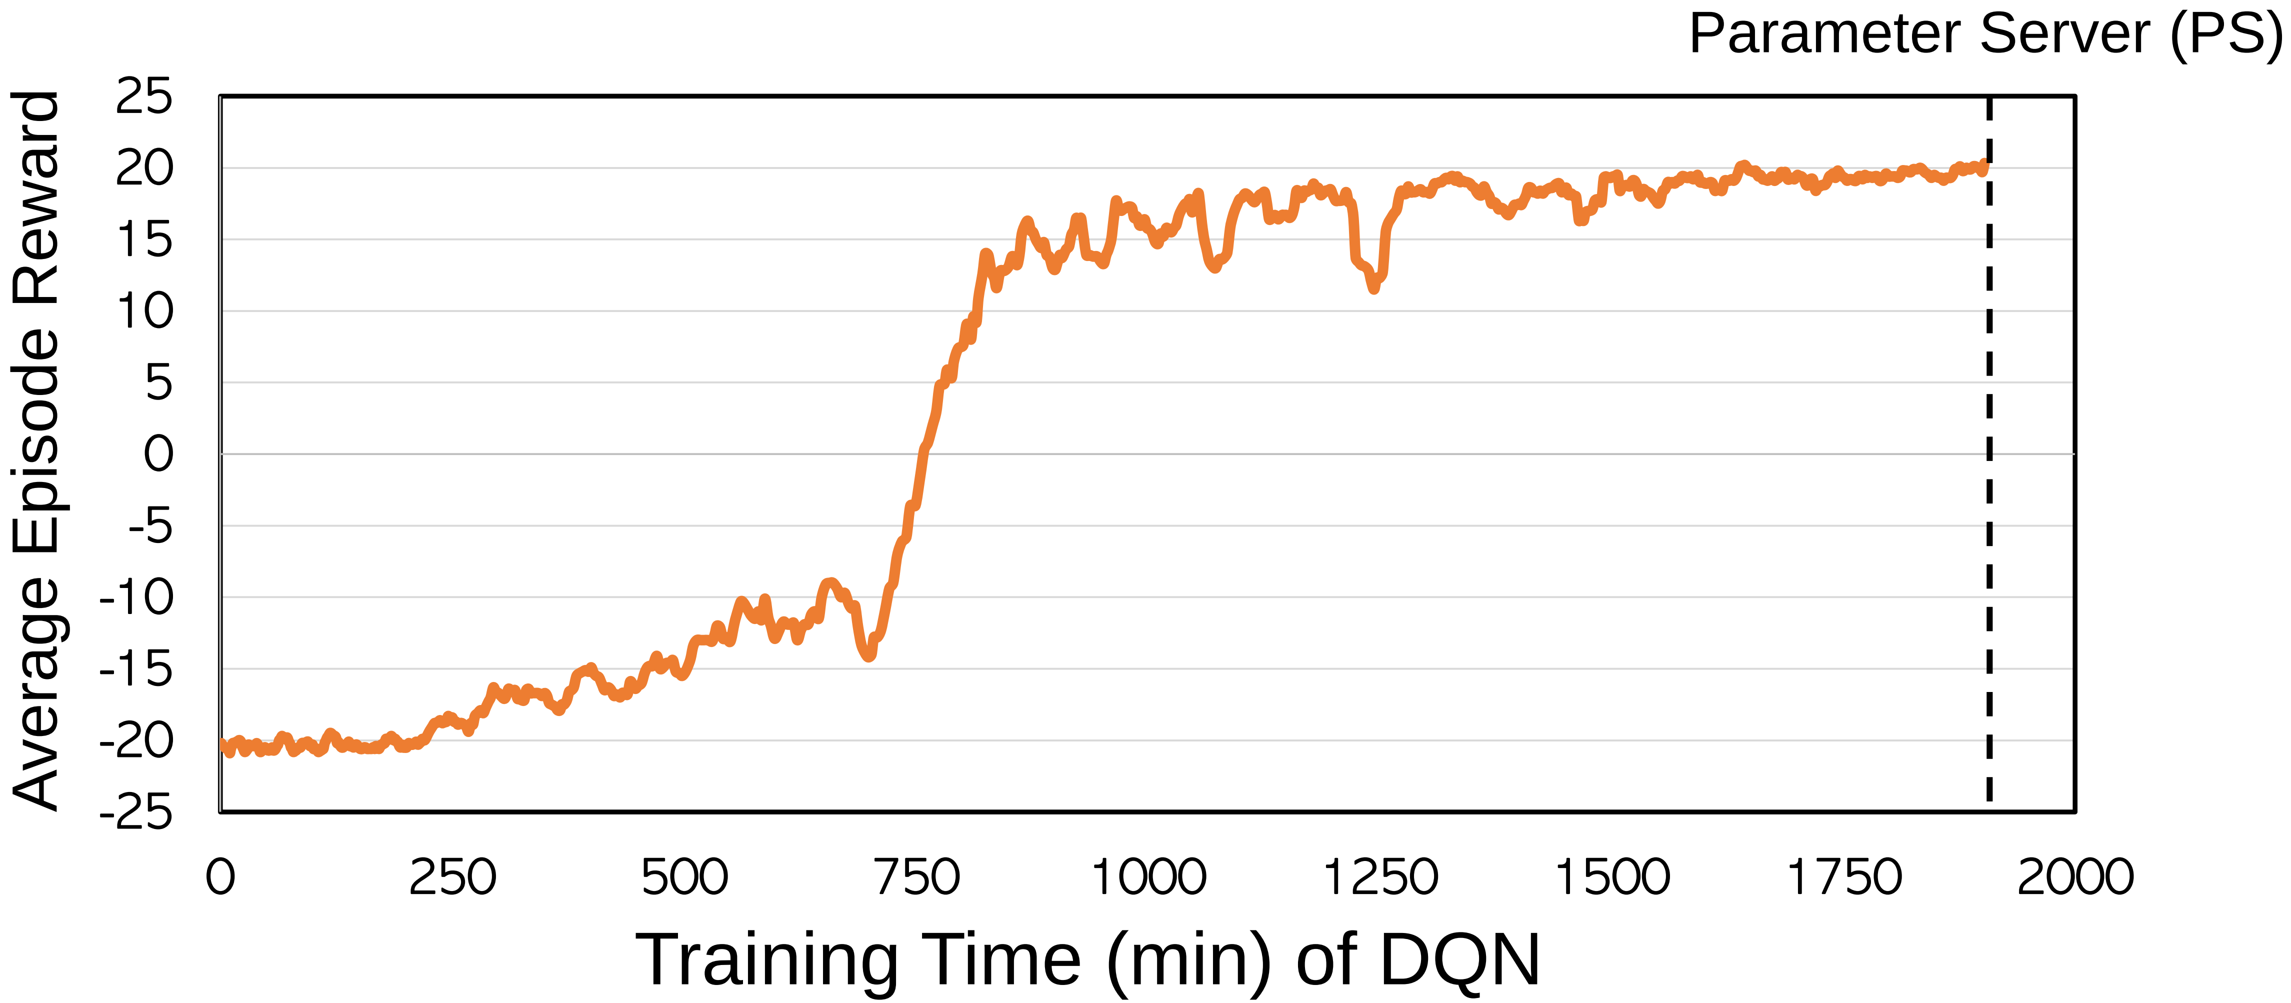

Reducing the End-to-End Training Time with iSwitch

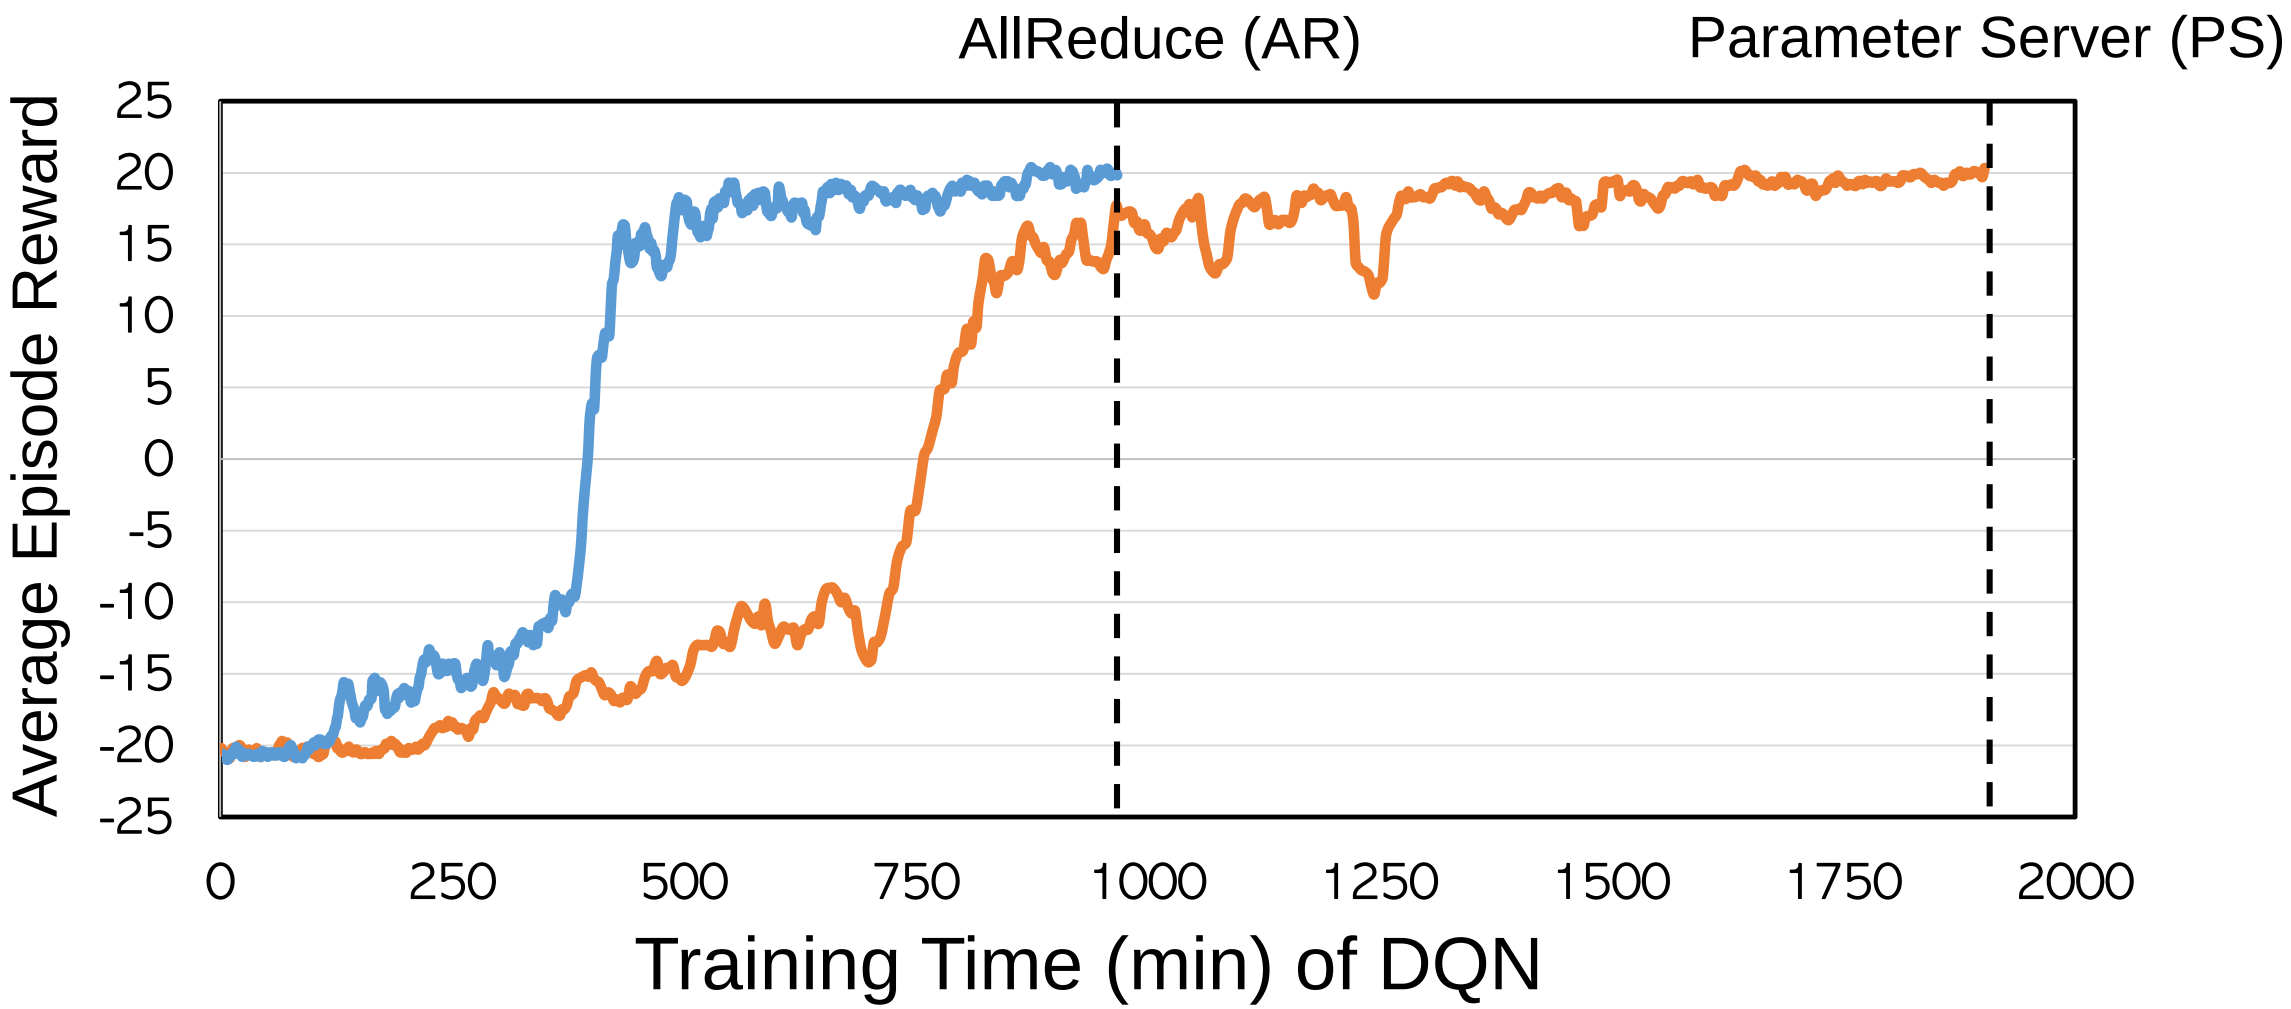

Reducing the End-to-End Training Time with iSwitch

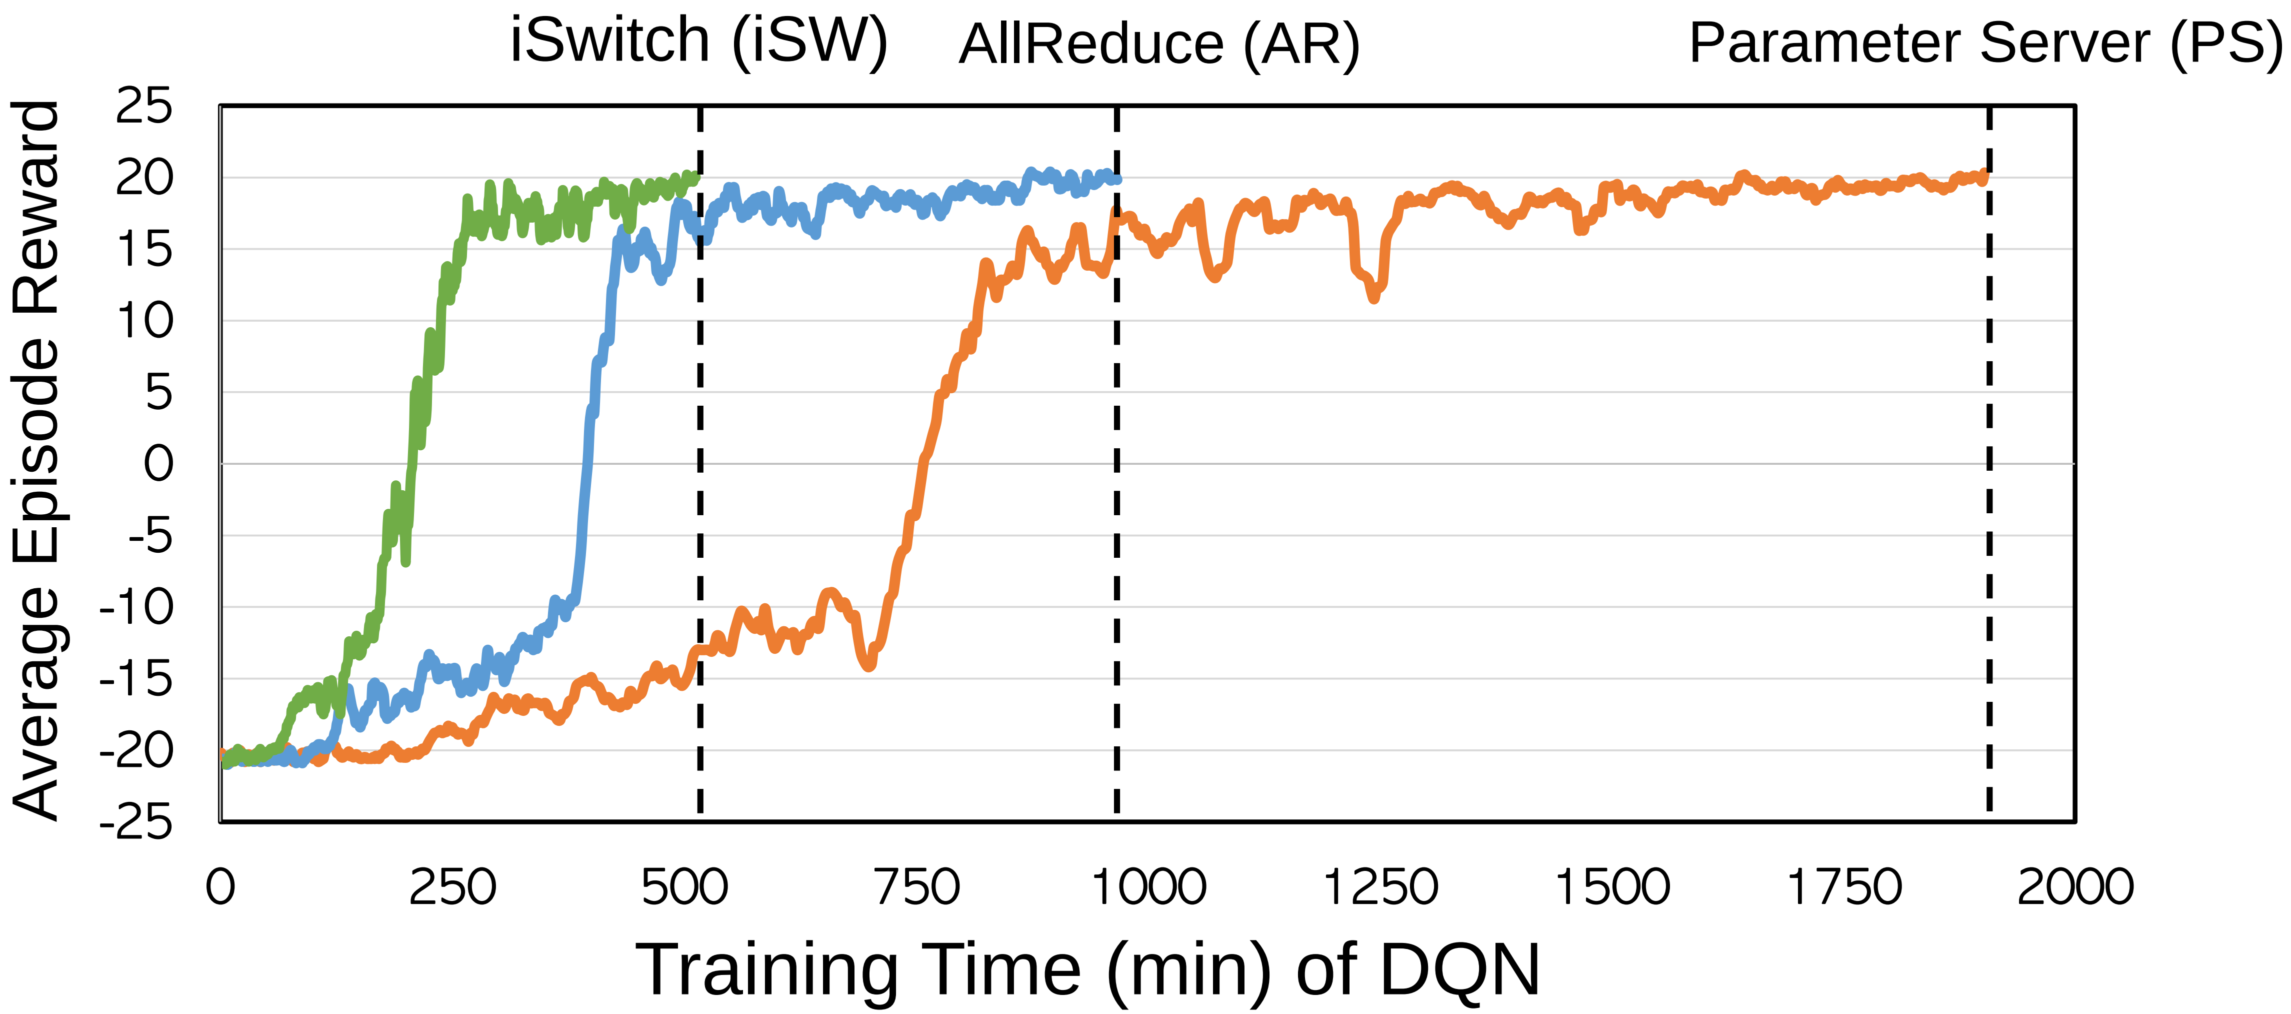

Reducing the End-to-End Training Time with iSwitch

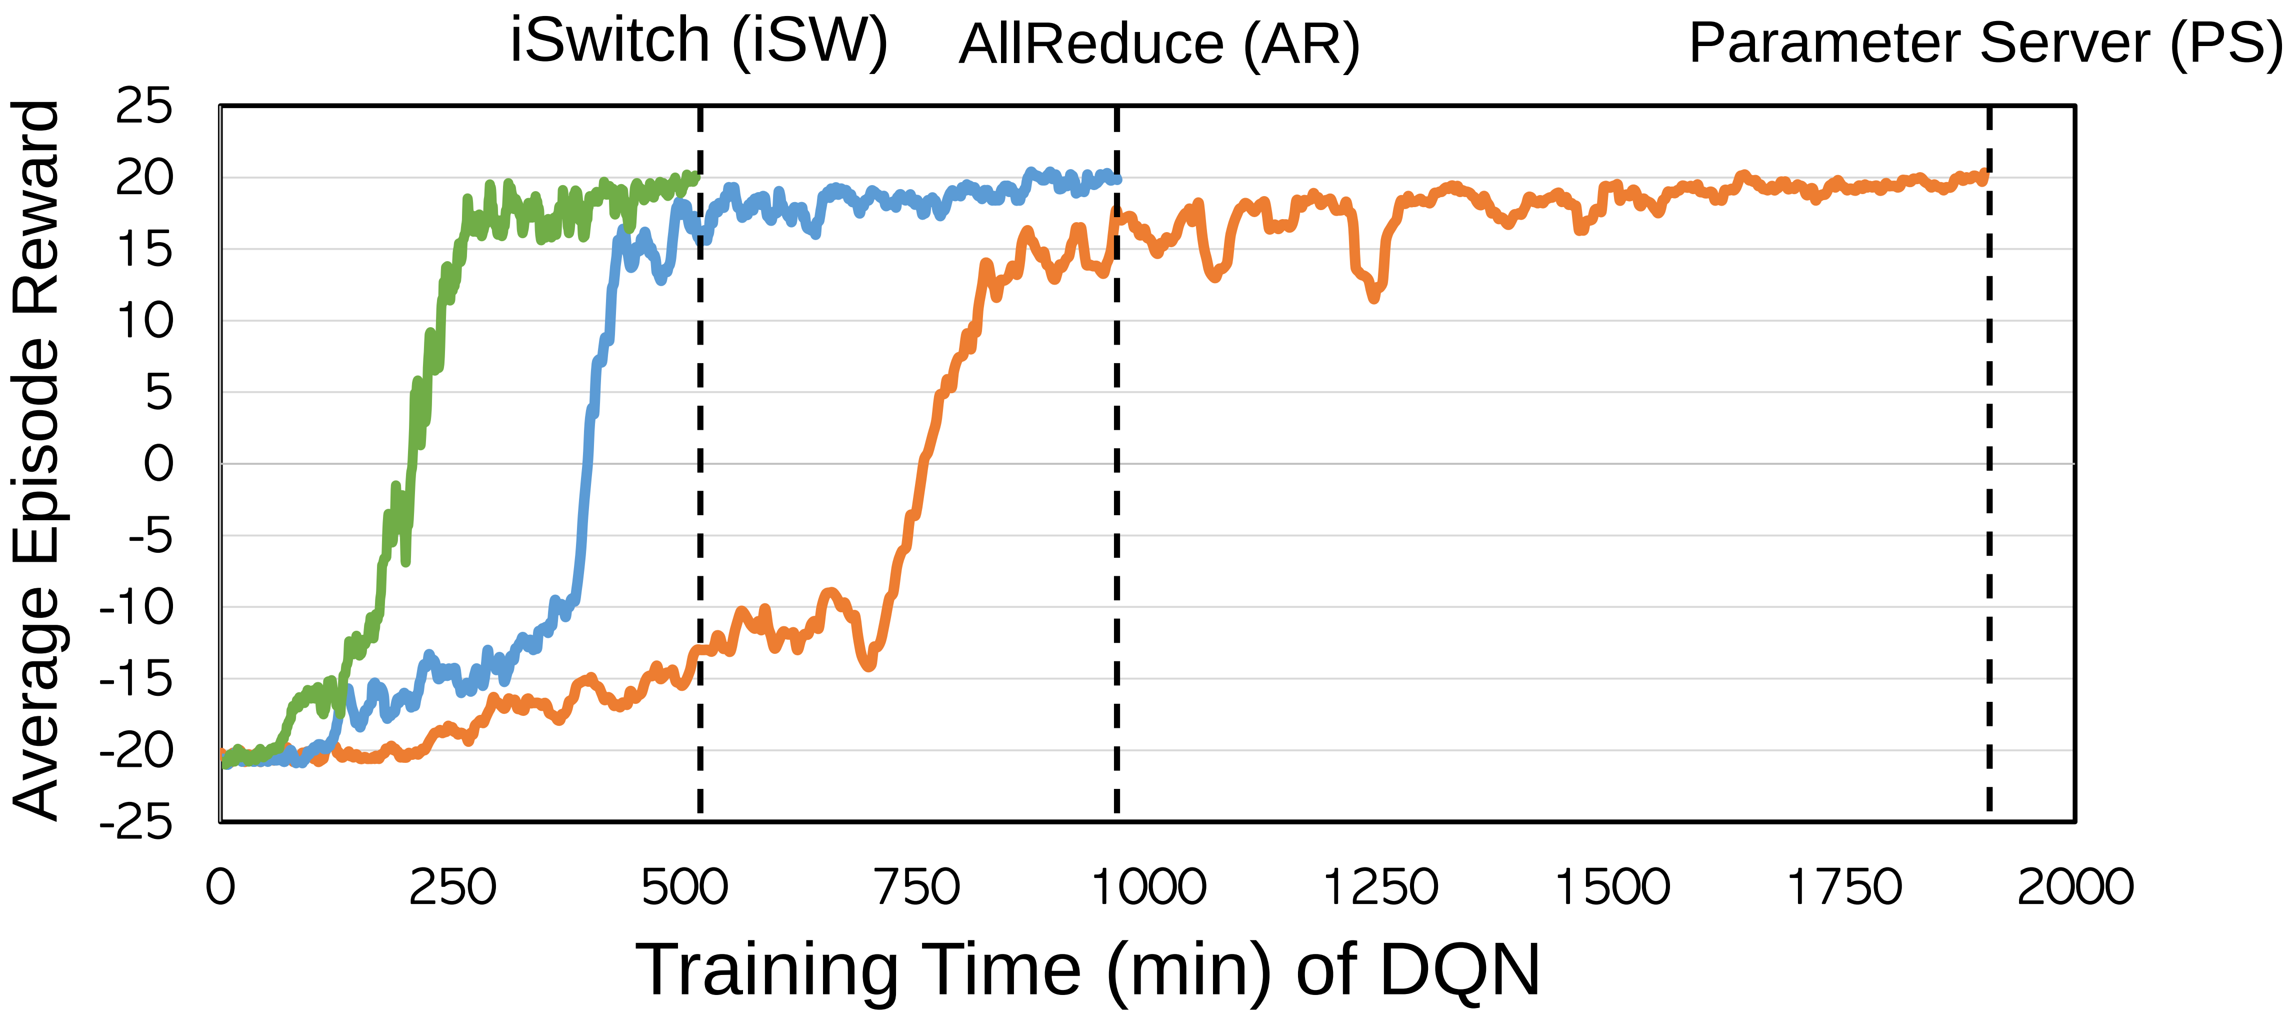

Reducing the End-to-End Training Time with iSwitch

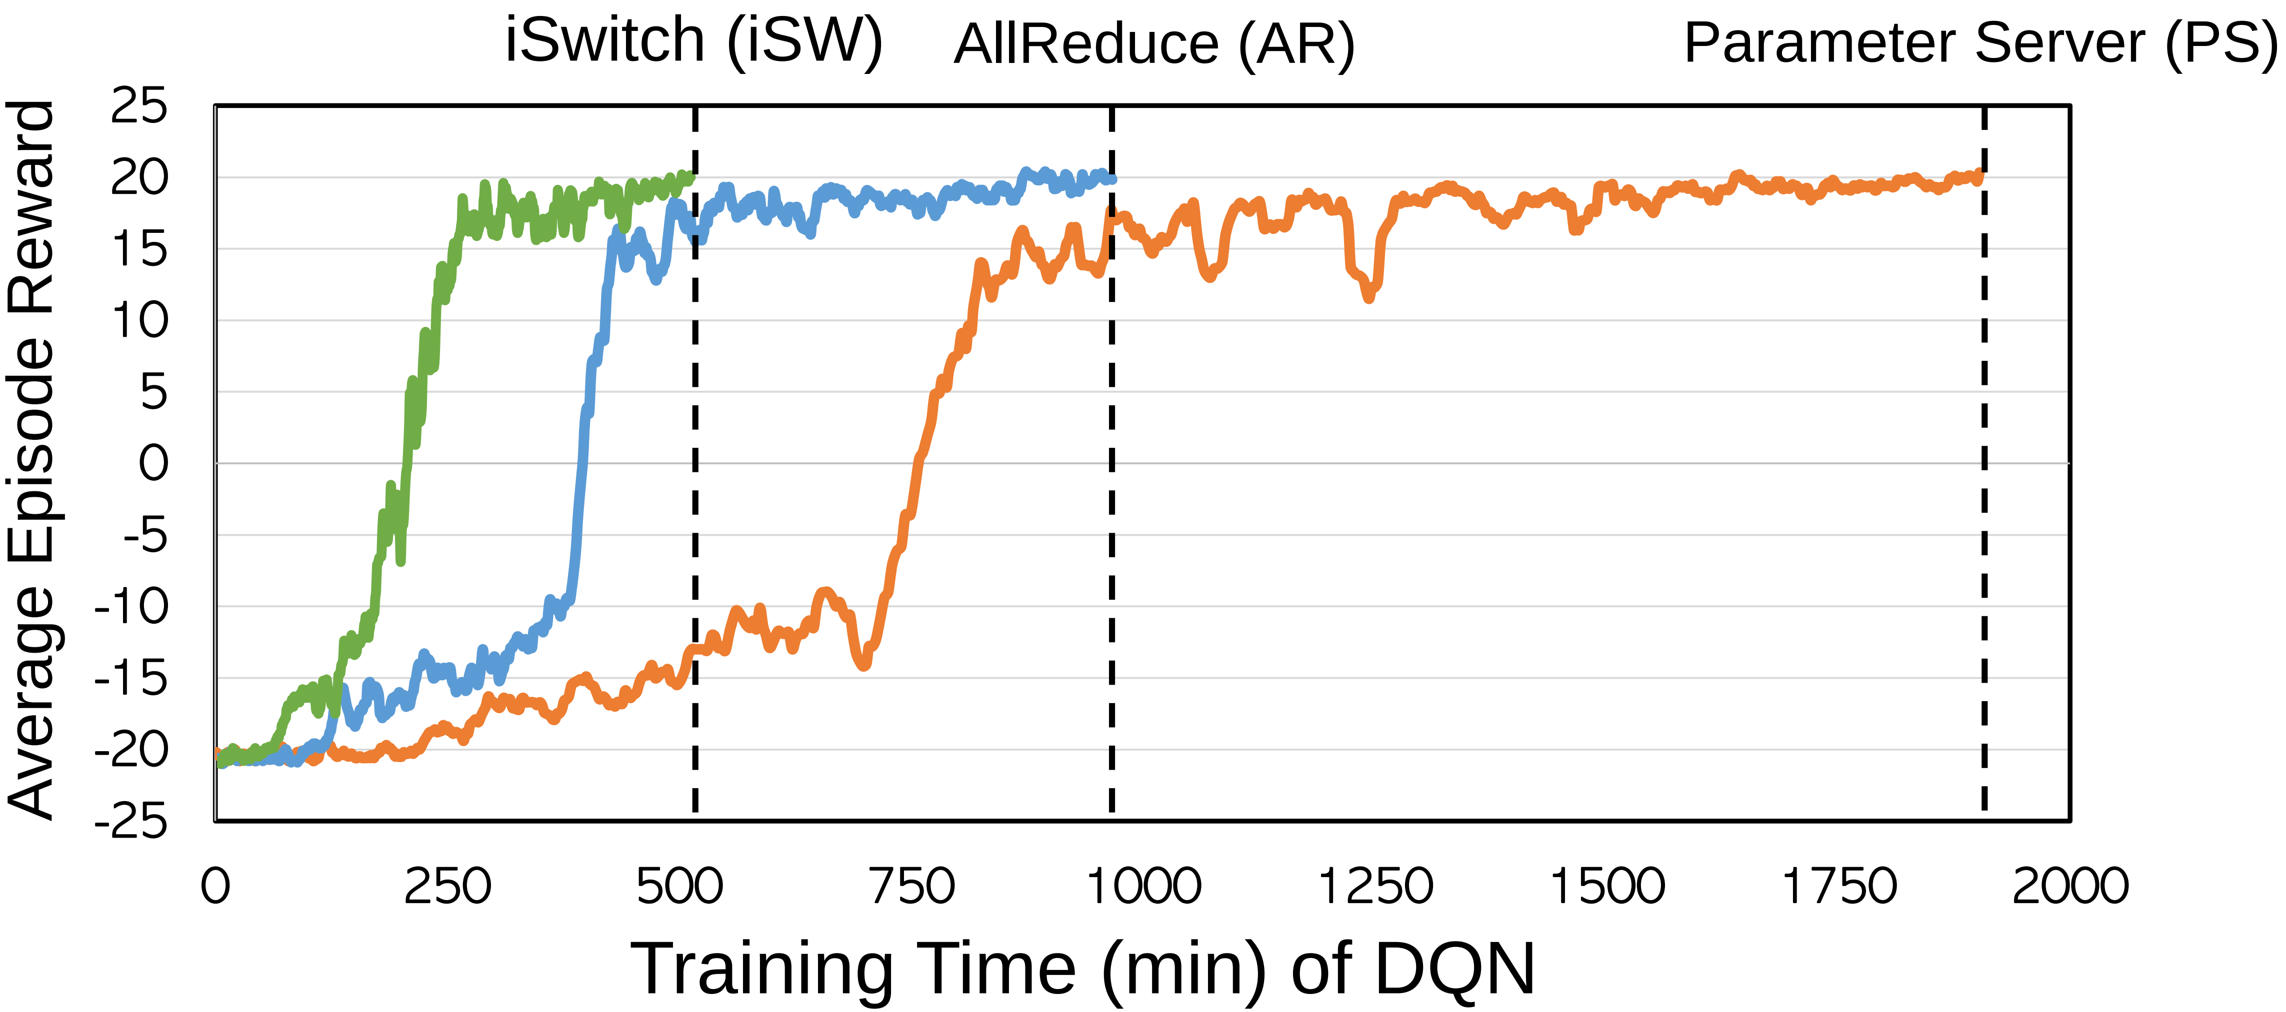

Reducing the End-to-End Training Time with iSwitch

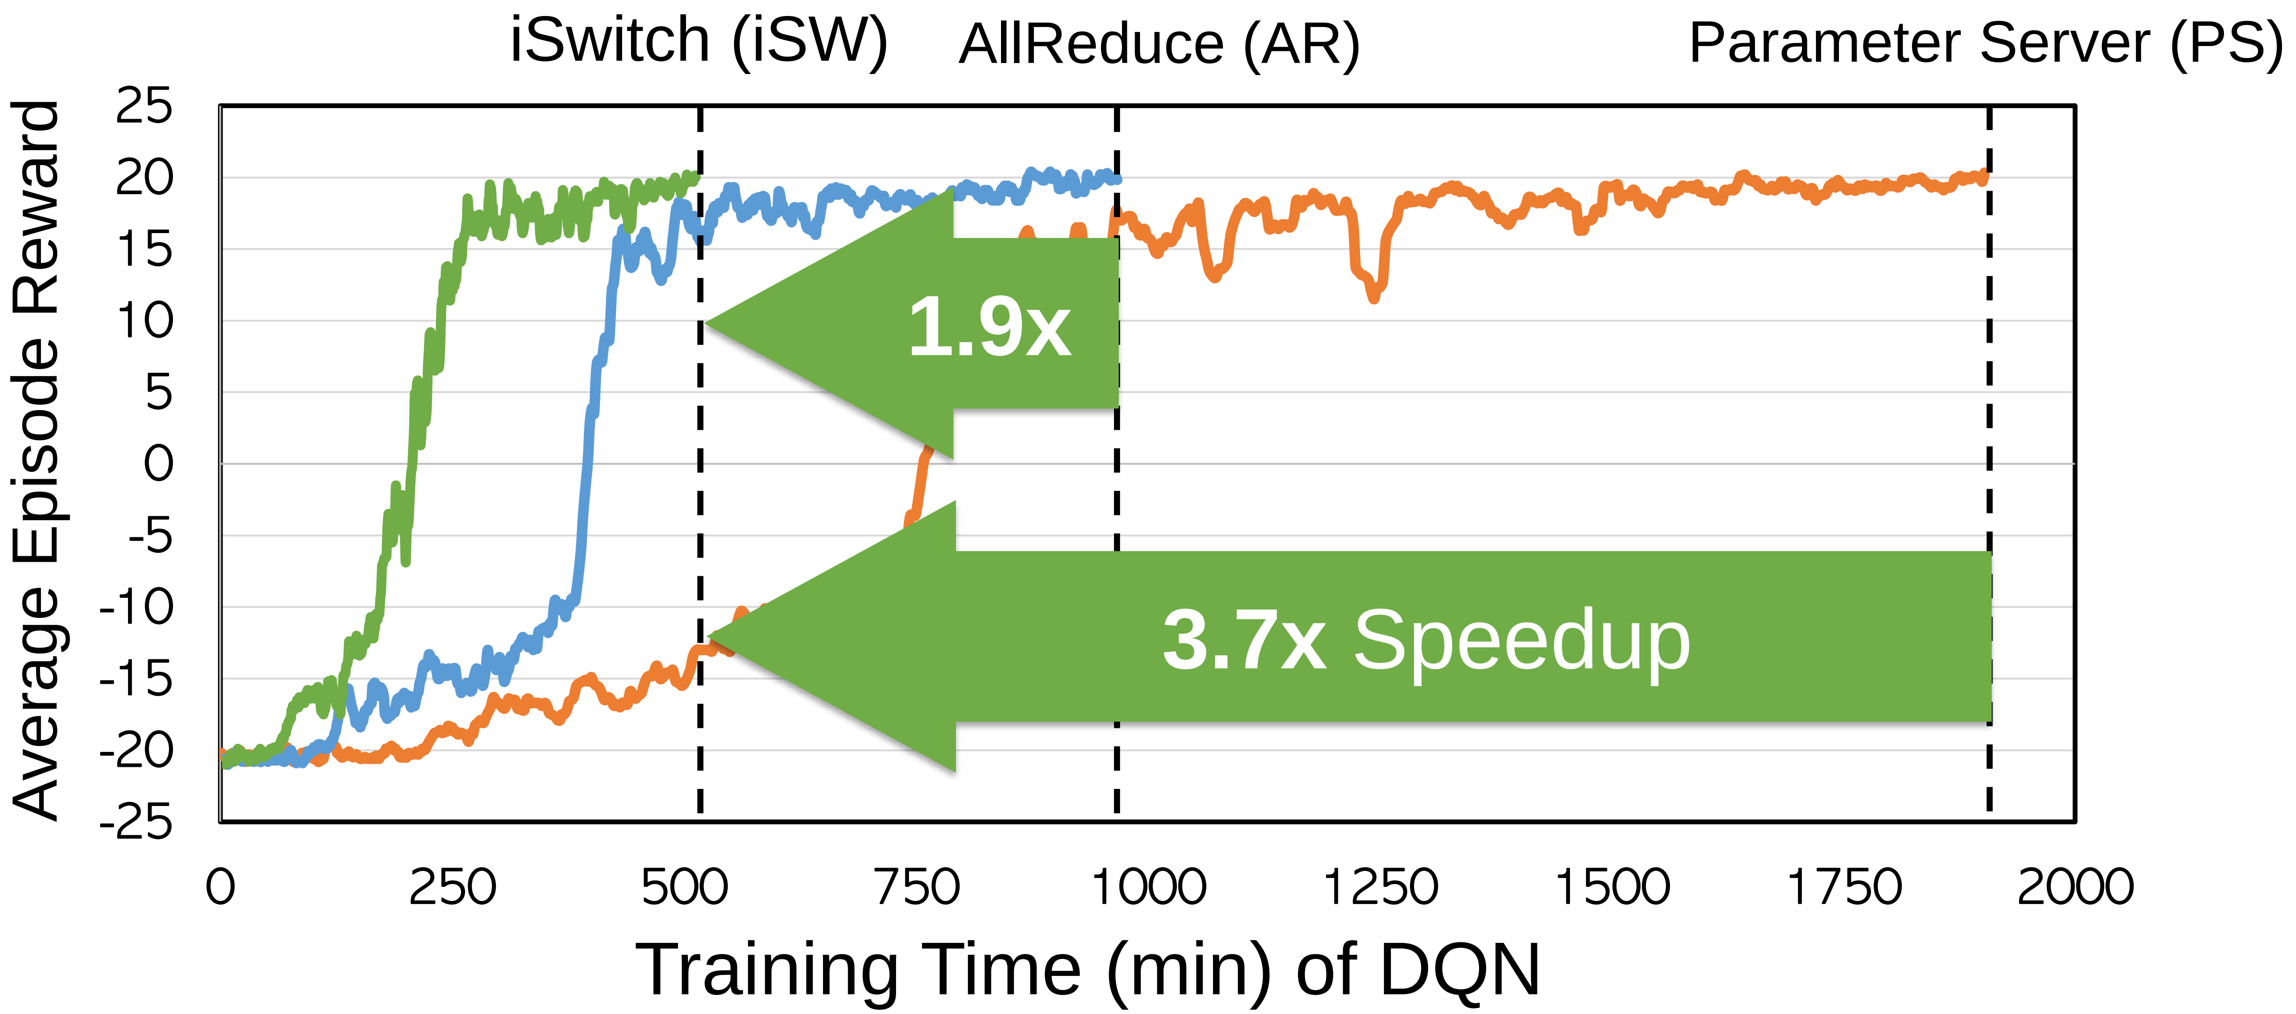

Performance Breakdown for Each Training Iteration

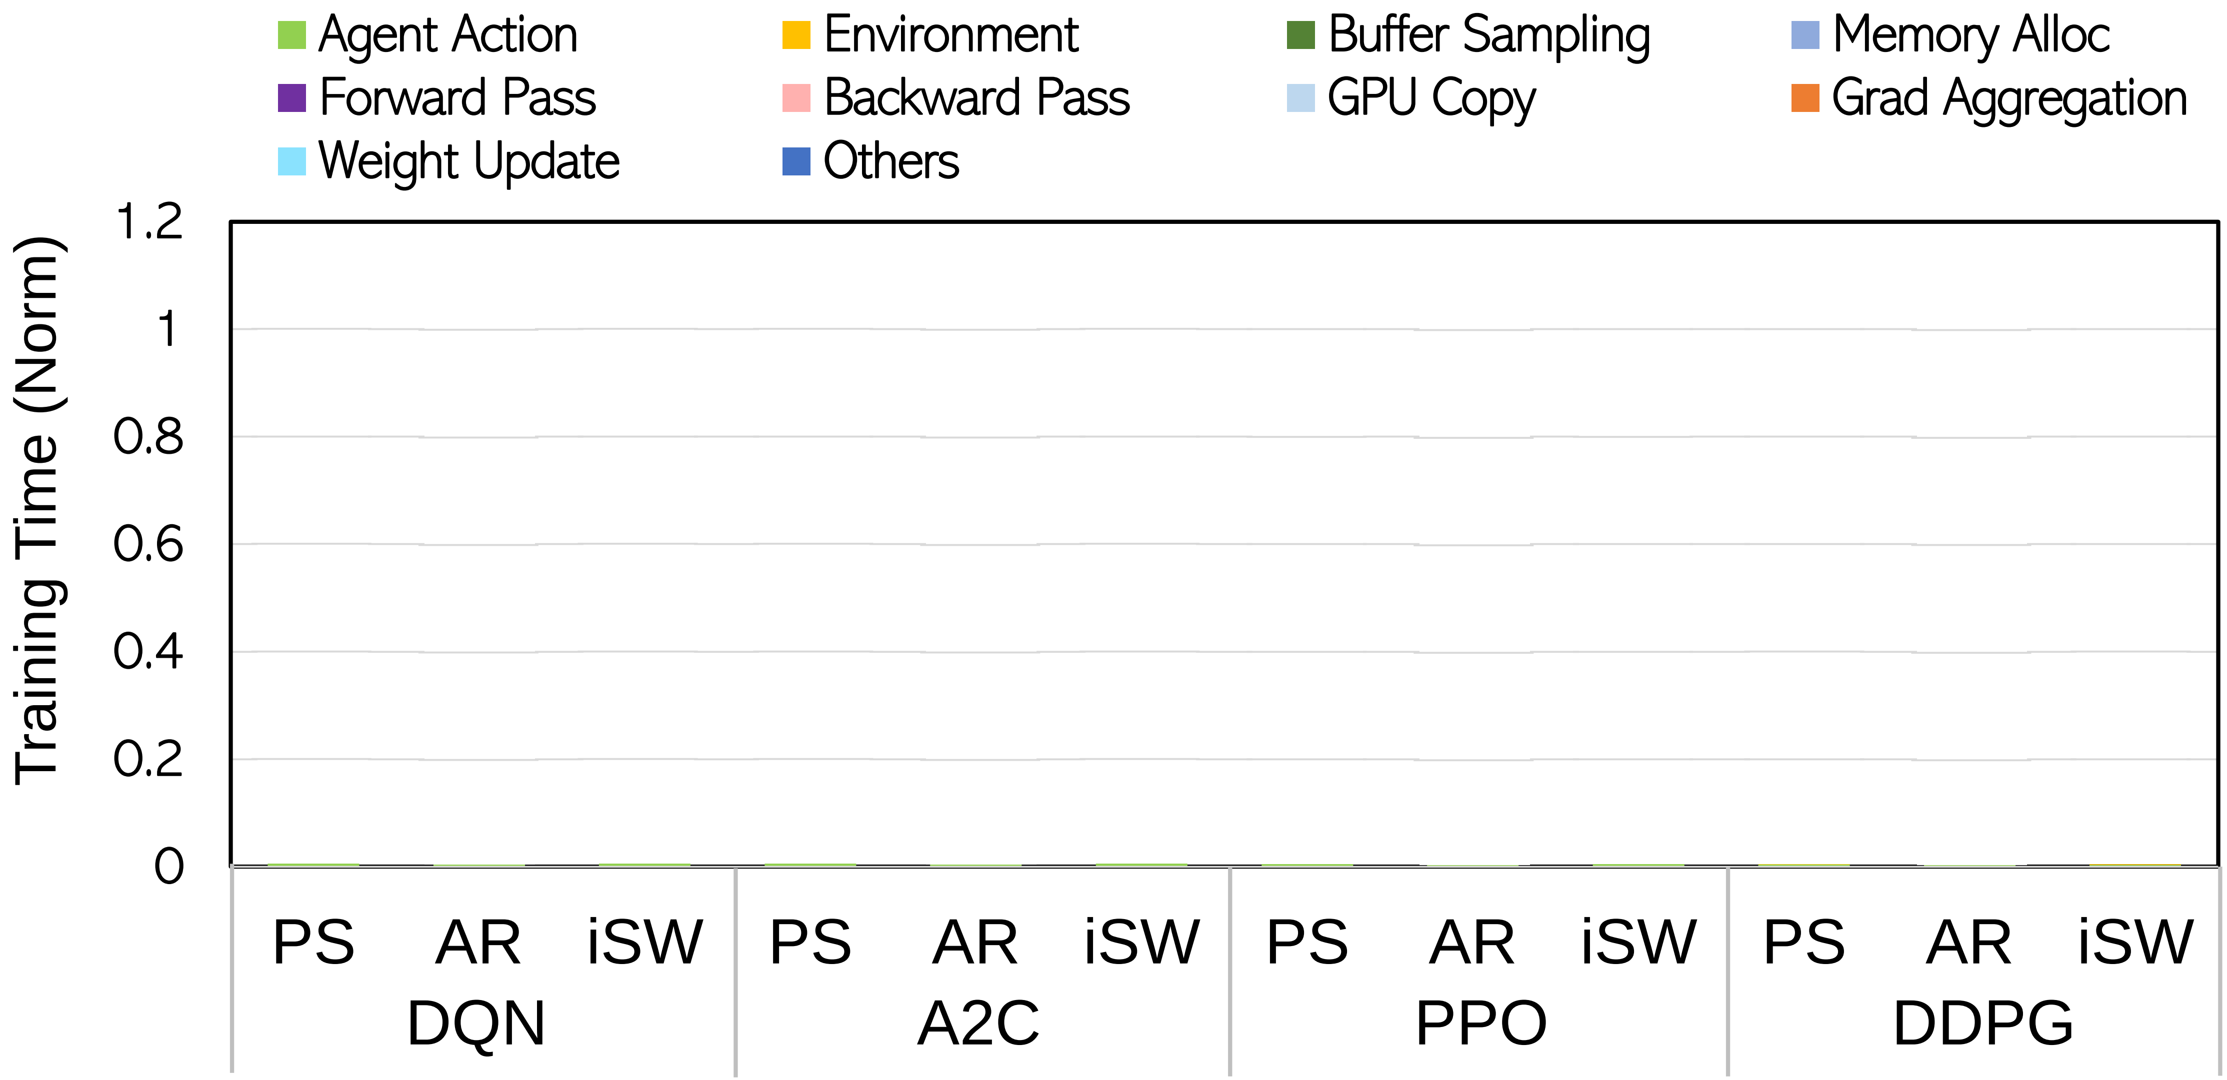

Performance Breakdown for Each Training Iteration

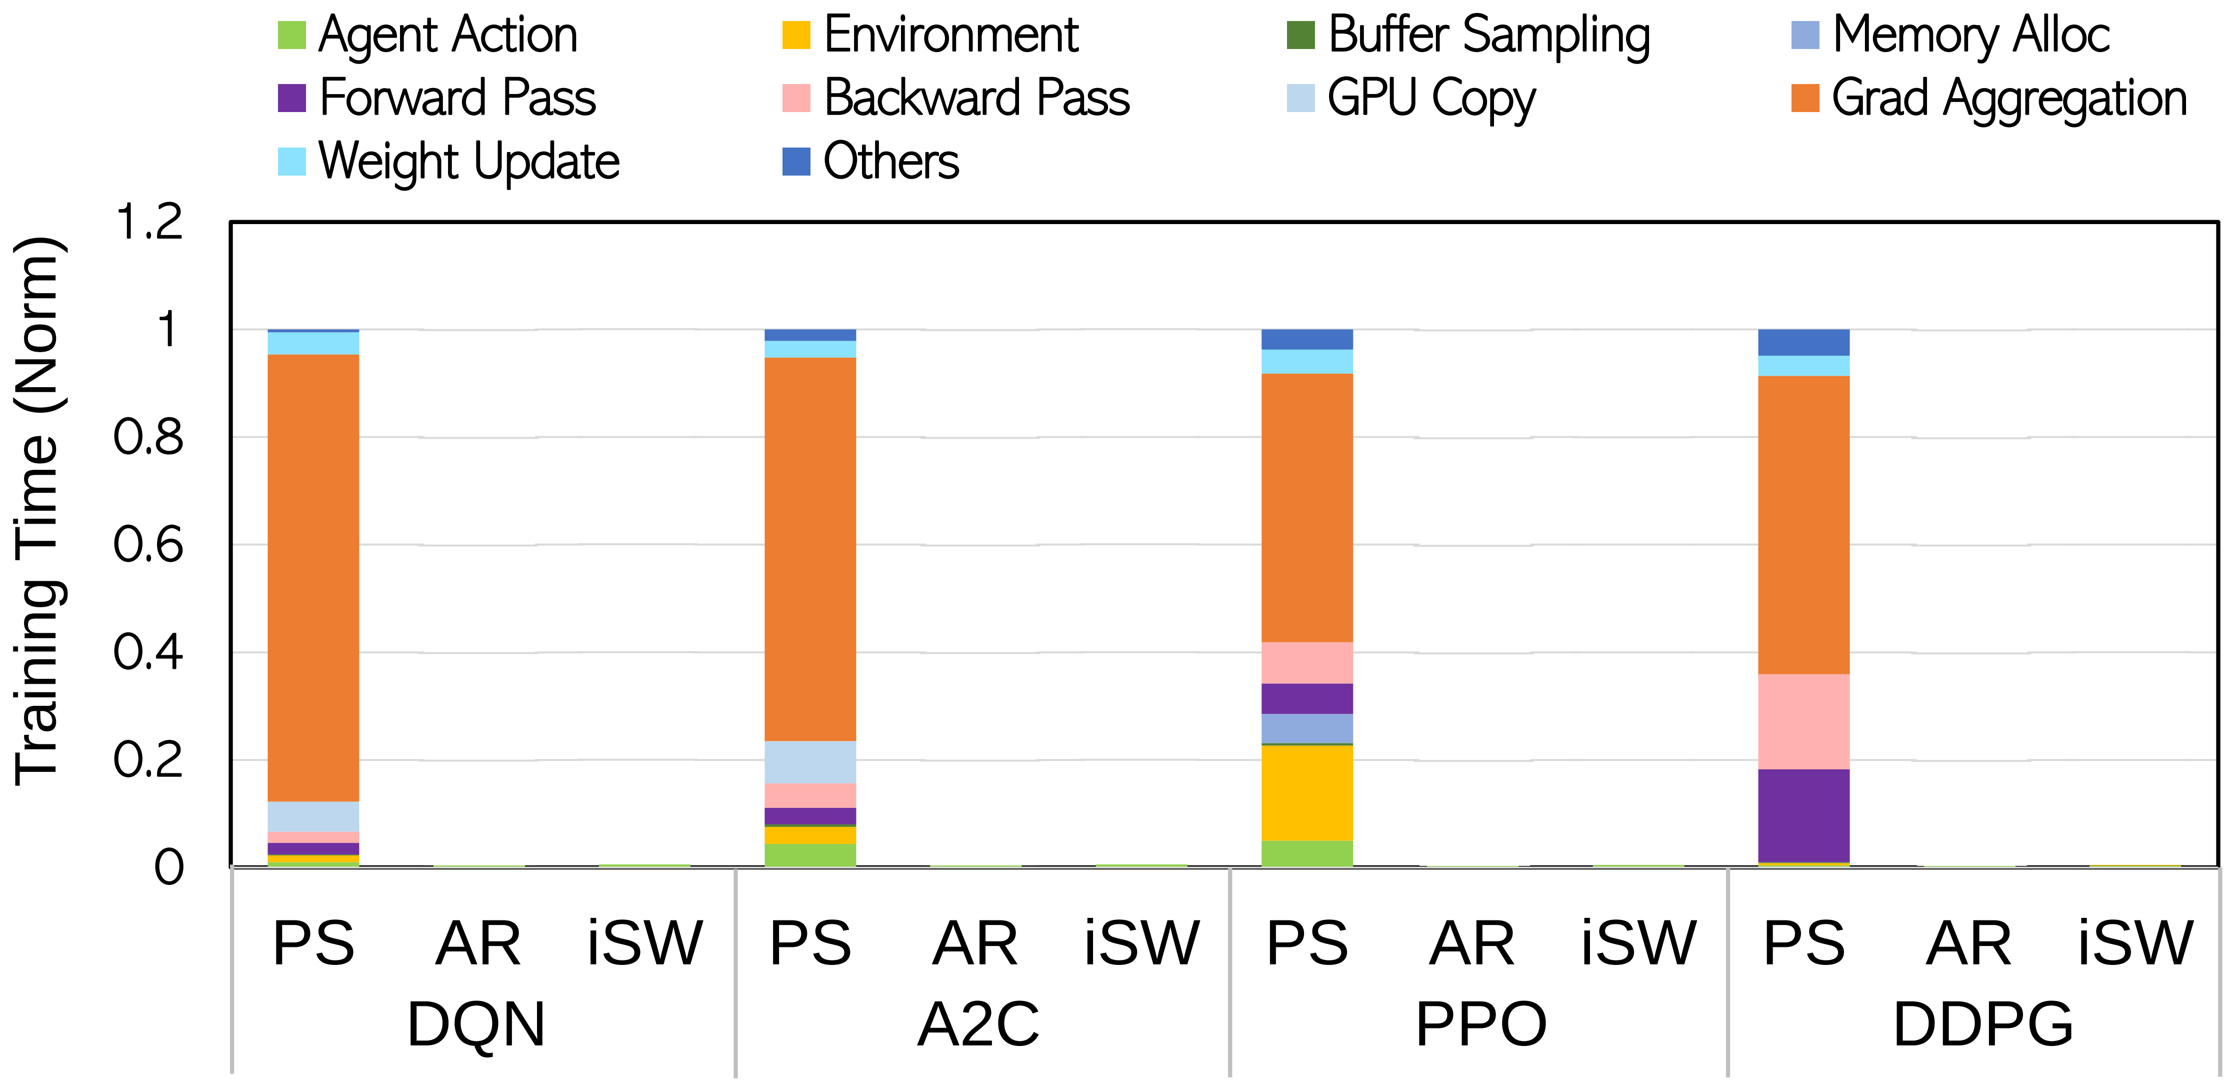

Performance Breakdown for Each Training Iteration

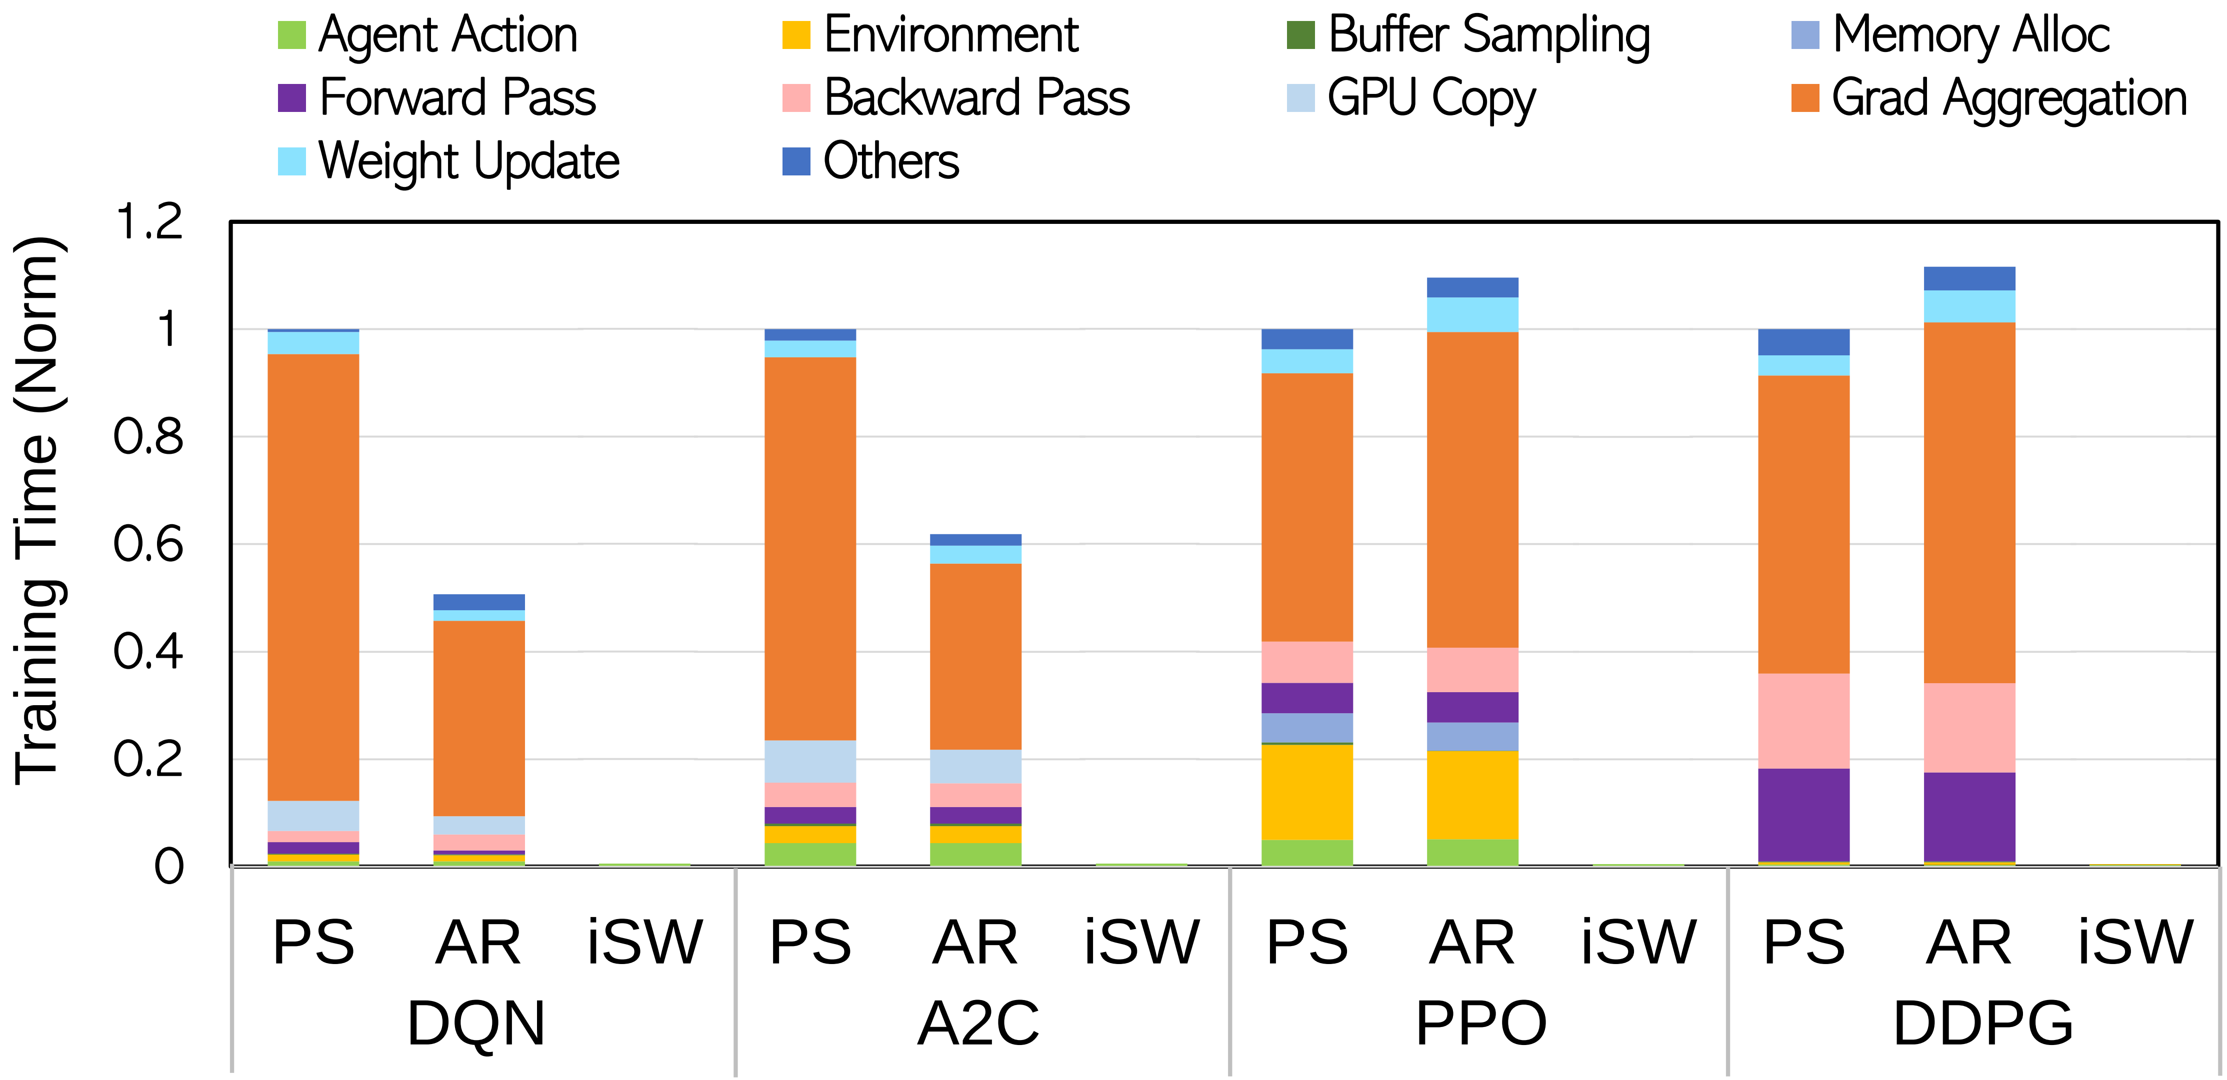

Performance Breakdown for Each Training Iteration

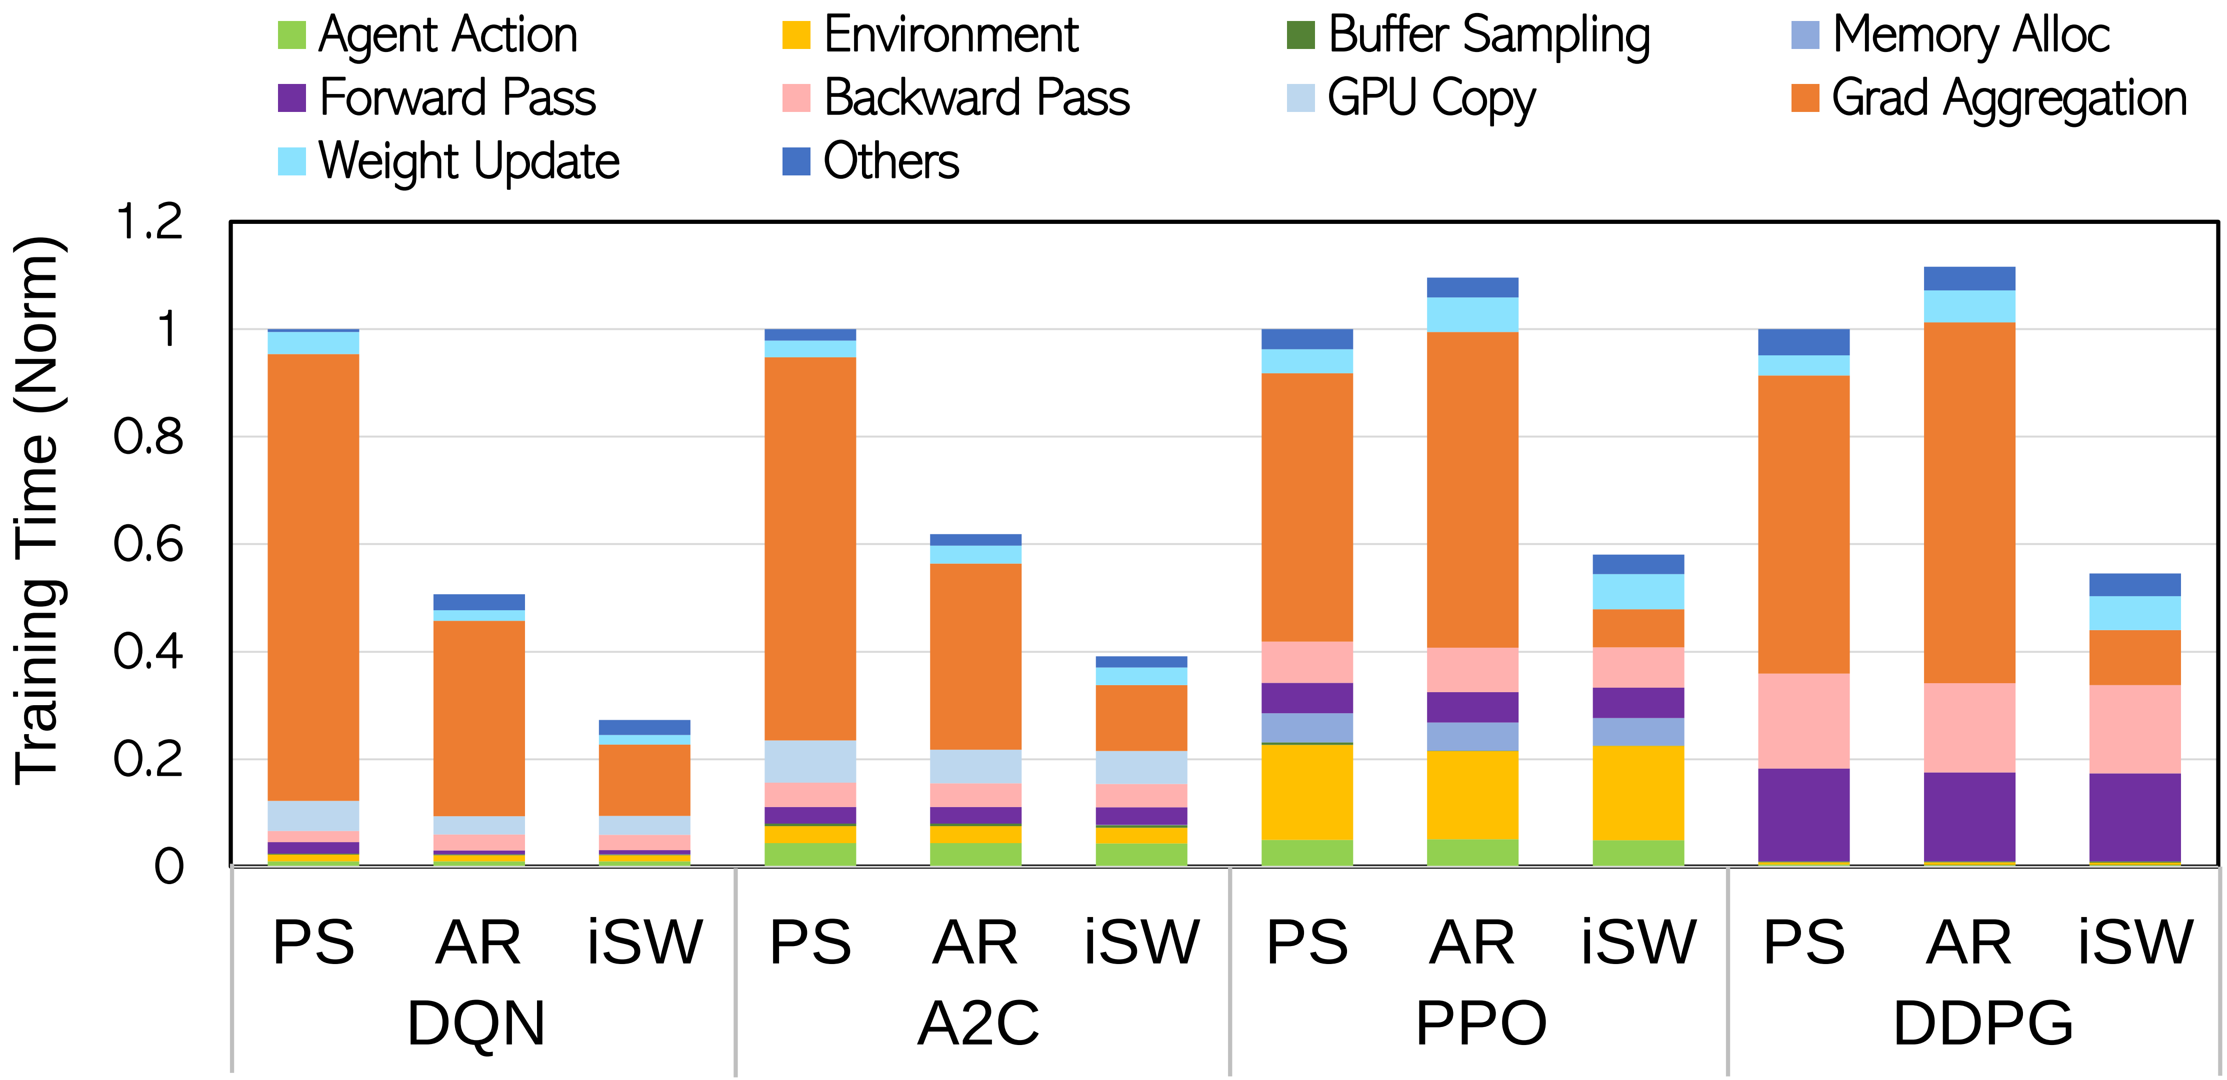

Improved Training Scalability with In-Switch Computing

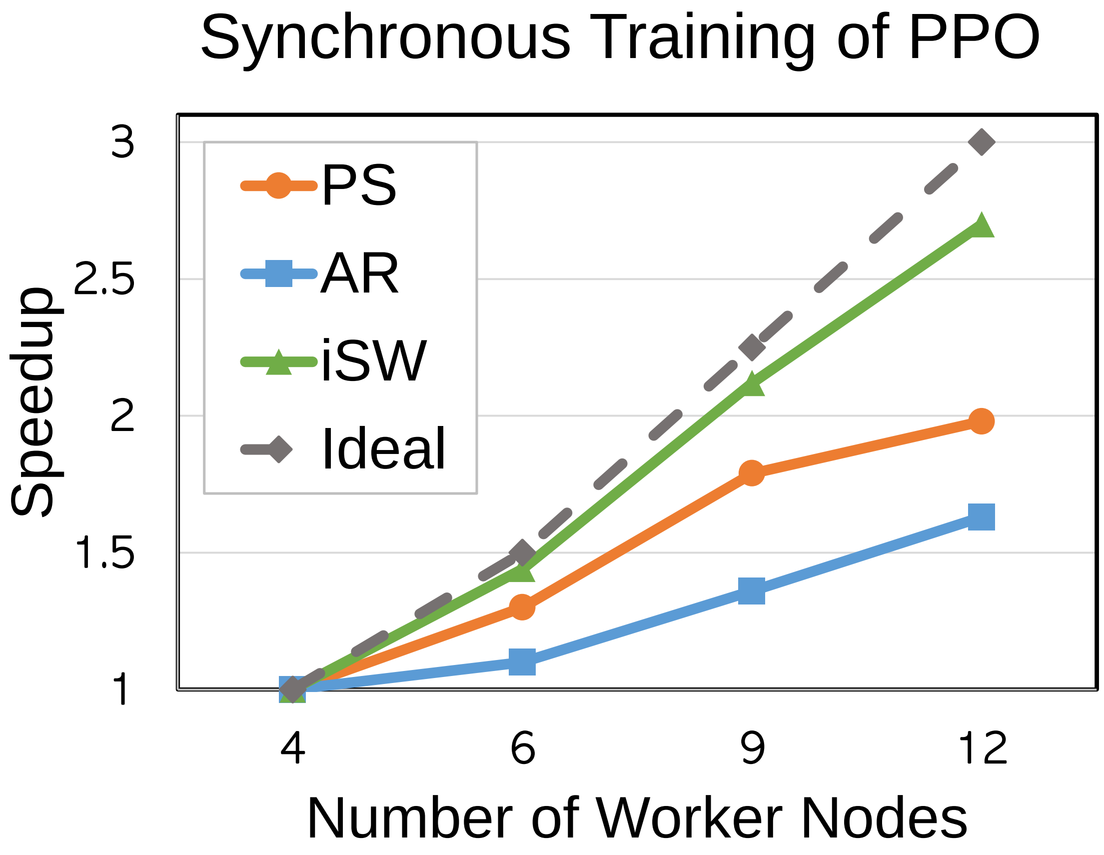

Improved Training Scalability with In-Switch Computing

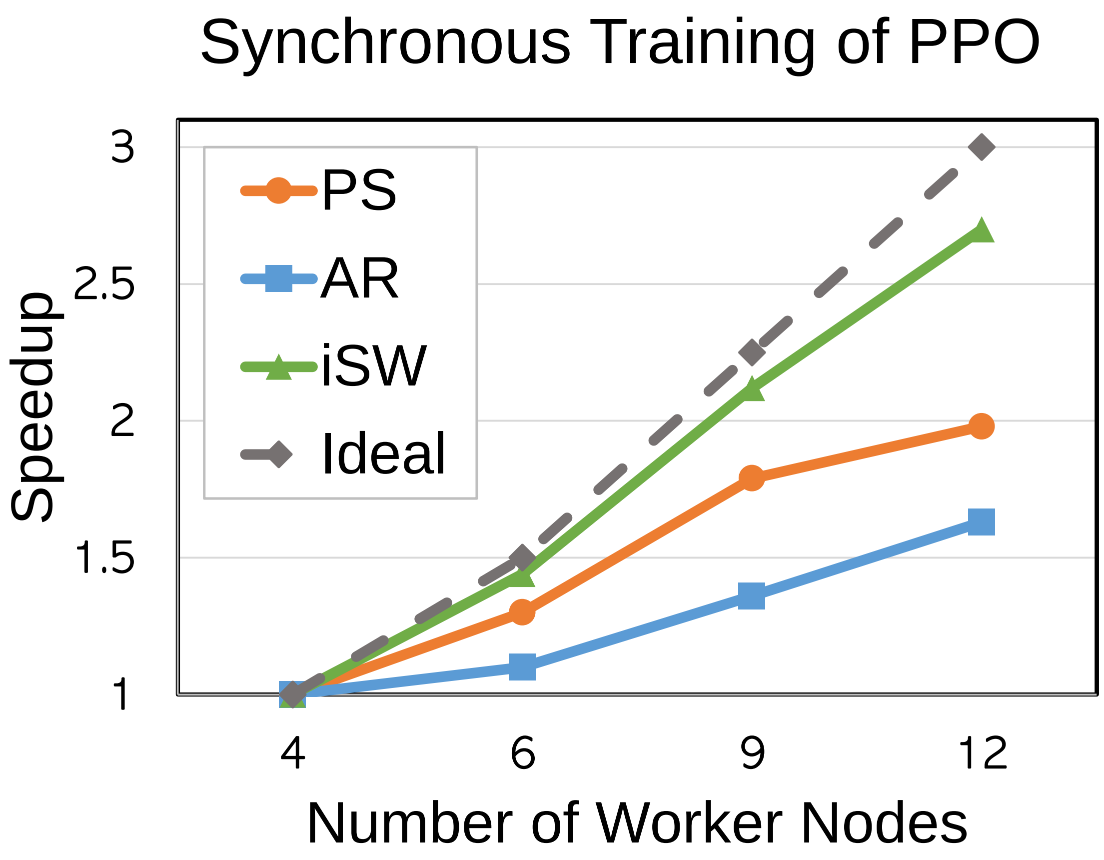

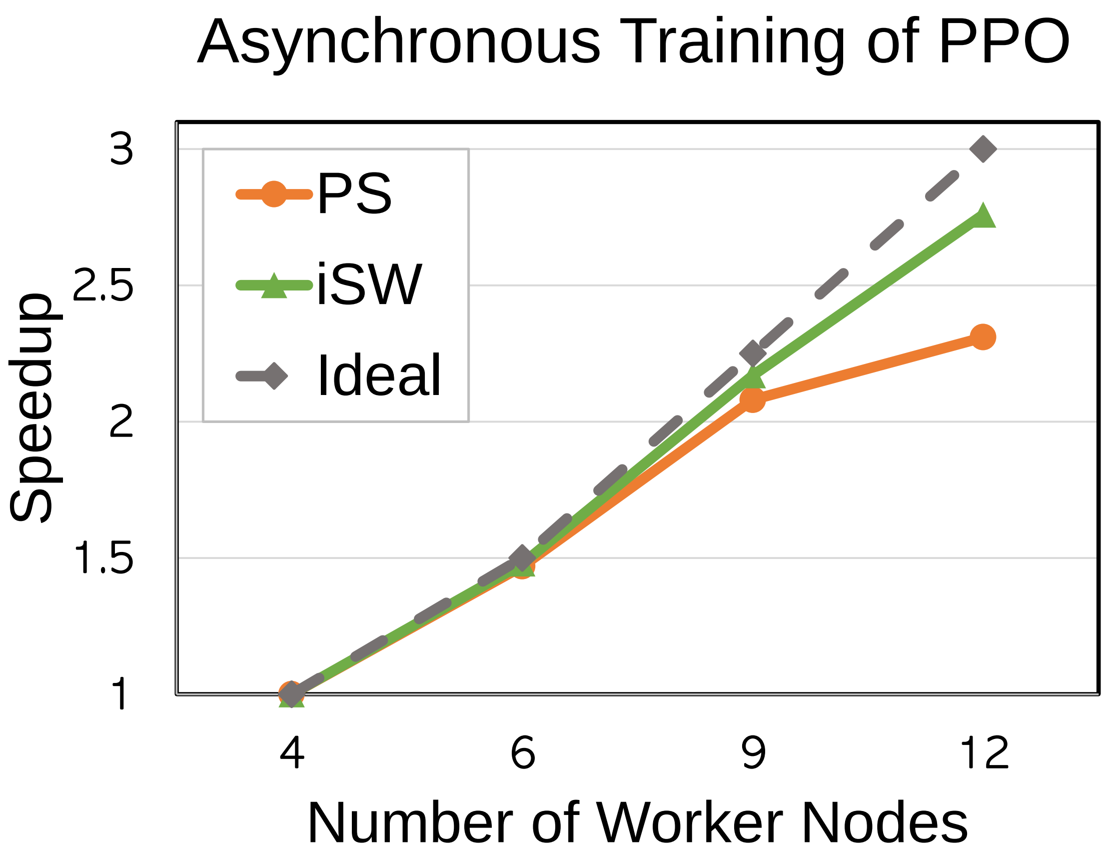

Improved Training Scalability with In-Switch Computing

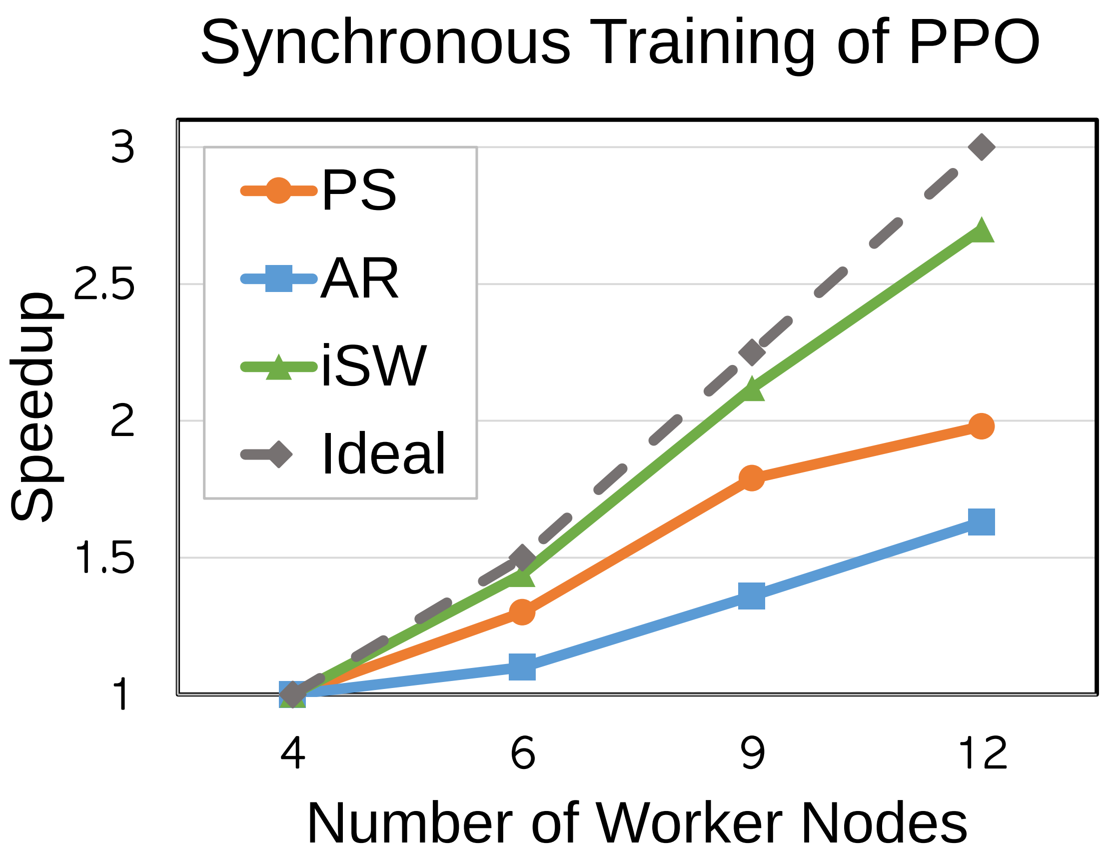

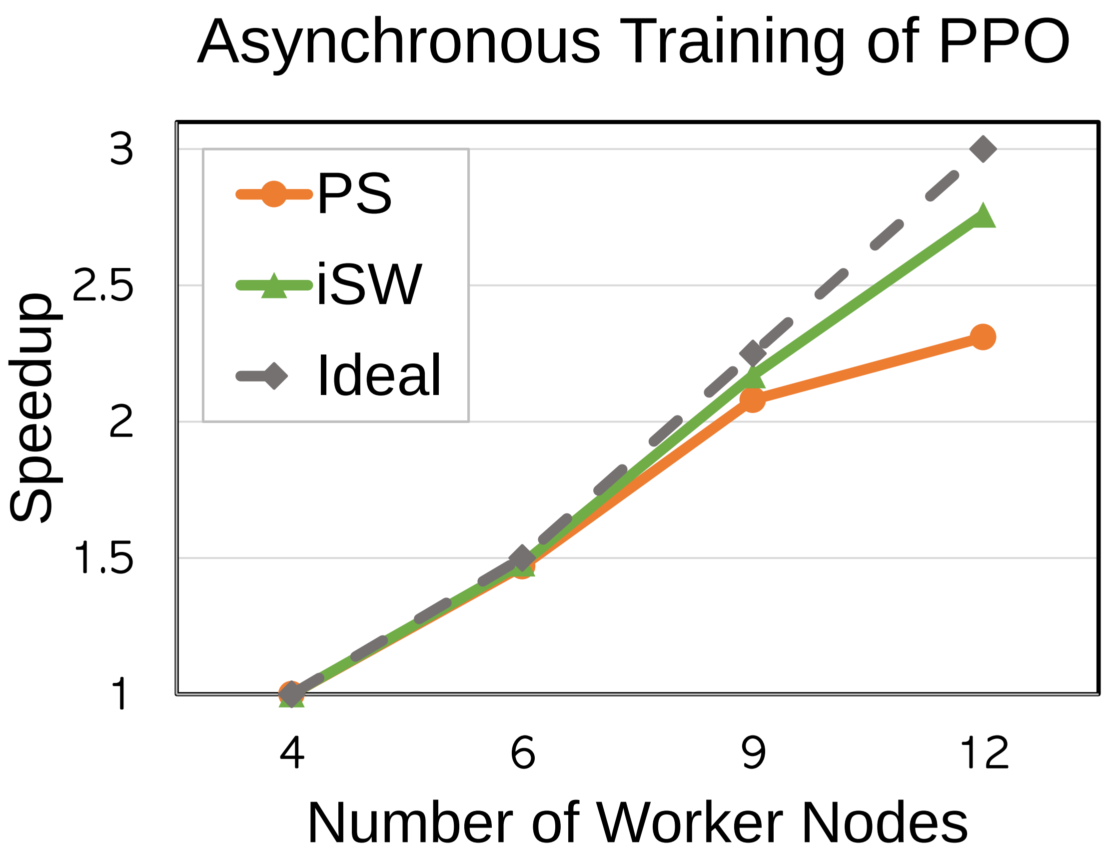

Close-to Linear Speedup for Both Training Modes

In-Switch Computing Summary

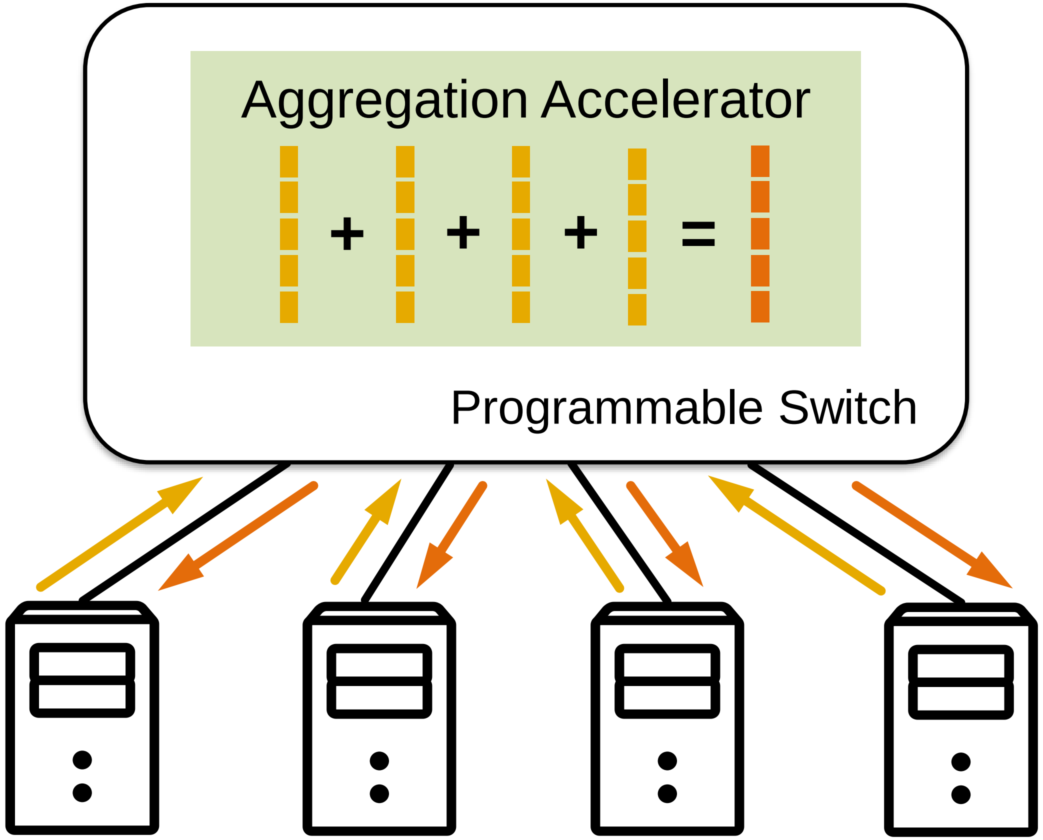

3.7x Speedup for Both Sync/Async Training

Scales at **Rack-Scale** Clusters

Thanks!

Youjie Li
li238@Illinois.edu

Iou-Jen Liu   Yifan Yuan

Deming Chen   Alexander Schwing   Jian Huang