Dynamic Learning-based Link Restoration in Traffic Engineering with Archie

Wenlong Ding, Hong Xu

The Chinese University of Hong Kong, Hong Kong, China

IEEE INFOCOM

香港中文大學
The Chinese University of Hong Kong

Optical Restoration in Wide-Area Networks (WANs)

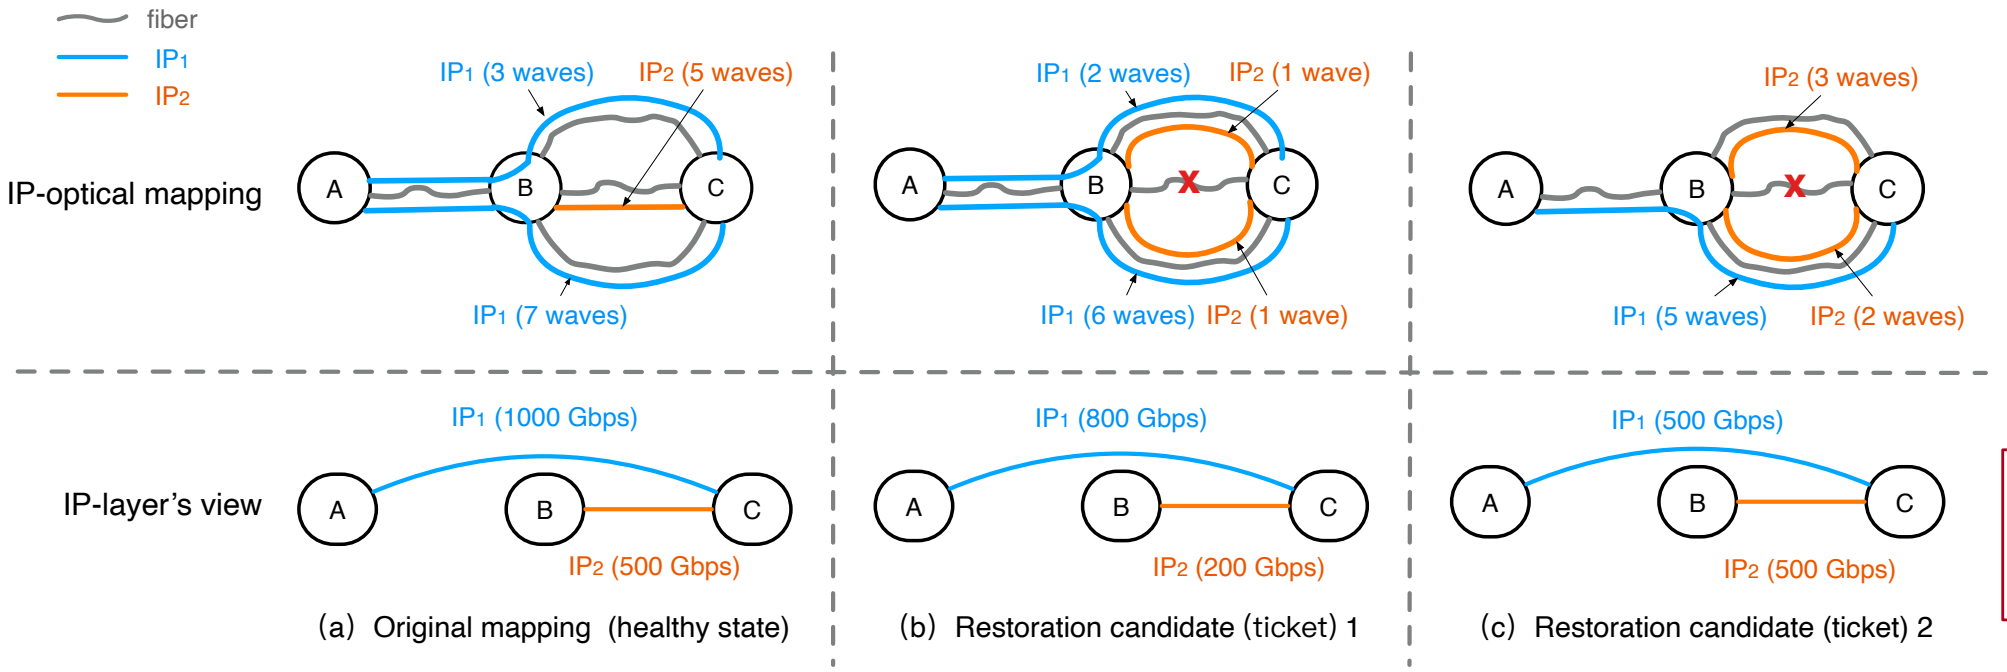

* 1 wave = 100 Gbps
* Each restoration candidate is recorded as a restoration ticket

- IP layer in WAN is constructed through IP-optical mapping
- When there is a fiber cut, there are many partial restoration candidates
- Traffic is routed on the IP layer in WAN

Best Restoration Ticket Selection Depends on Traffic Demand

- Which restoration candidate (ticket) leads to highest throughput?

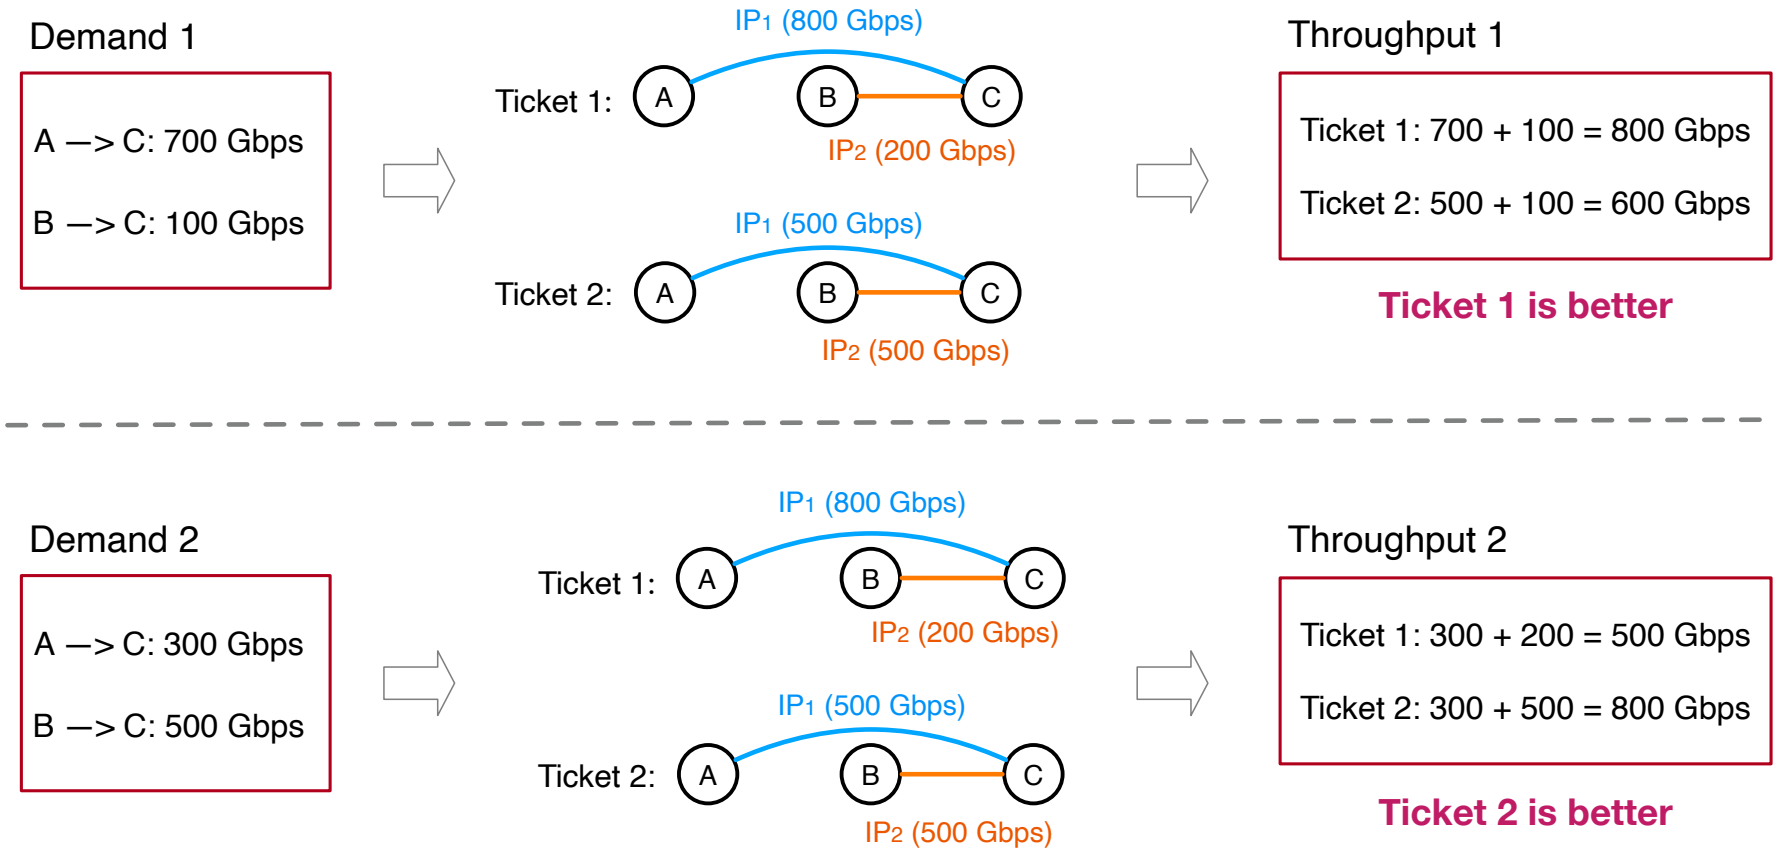

Existing Work: Candidate Ticket Set + Traffic Engineering (TE)

- Within a fixed IP-layer view, a flow (src-dst pair) can be routed among different candidate paths.
- The possible maximum throughput can be formulated as TE optimization formulation.

Determination Variable

W_f^k : traffic fraction of flow *f* on path *k*

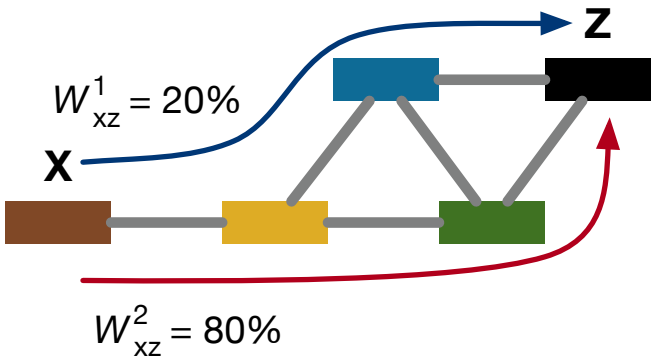

Demand for flow *f*

$$\max \sum_f \sum_k w_f^k D_f$$

Maximize throughput

$$\text{s.t.} \sum_k w_f^k \le 1, \forall f,$$

$$w_f^k \ge 0, \forall f, k,$$

Allocate non-negative traffic and total traffic assigned cannot exceed flow's demand

$$\sum_f \sum_k P_e^k w_f^k D_f \le C_e, \forall e.$$

Aggregated traffic on each Link does not exceed its capacity

Whether link *e* on path *k*

Capacity of link *e*

- Arrow (Sigcomm'21) select the best ticket by *formulating TE* on *different restoration tickets* and select the best with highest throughput using **instantaneous traffic** when fiber is cut.

New Question: Ticket Selection for the Long Run

- Fiber cut always take long to repair (9 hours on average (Arrow Sigcomm'21))
- **Challenge 1: Traffic dynamics** in fiber repair time

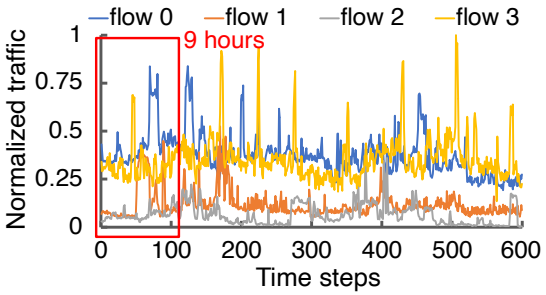

(a) Traffic on 4 selected flows

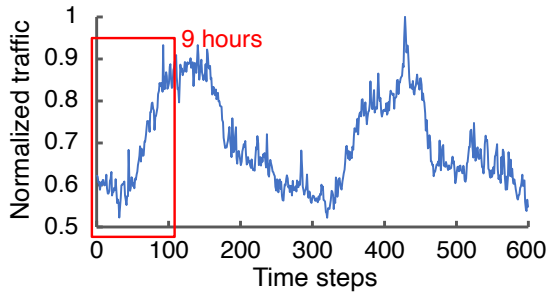

(b) Sum of the traffic on all flows

- **Challenge 2: Reconfiguration Overhead**
- Reconfiguring wavelengths takes non-negligible time
  - At least *O*(10) seconds with the latest hardware on a simple 4-node topology (Arrow Sigcomm'21)

Static Restoration: Calculate one ticket with traffic at fiber cut and use it for the whole repair time.

Constant Restoration: Change ticket every TE time step to cope with dynamic traffic.

Extreme ticket selection methods may lead to throughput loss

- **Balance traffic dynamic and reconfiguration overhead:** select restoration tickets every *T* time steps. (*T* is a short period (e.g. *T*=10) compared with whole repair time)

Strawman Solution for *T*-Step Ticket Selection

- Strawman solution: Demand Prediction

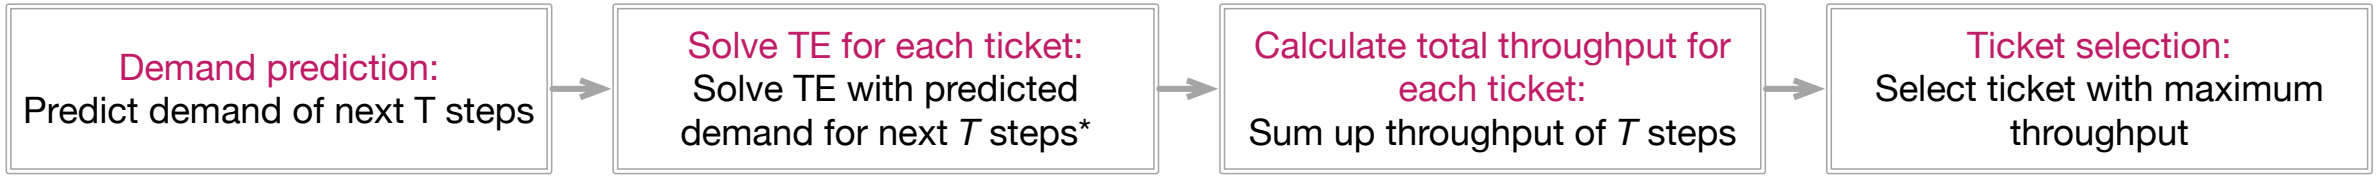

- Backwards of Demand Prediction:

1. Time consuming

- Requires to solve TE for *TZ* times (*Z* is the number of candidate tickets) to select one ticket.

2. Prediction inaccuracy

- Traffic prediction will be inaccurate for the long horizon even with state-of-the-art learning methods.

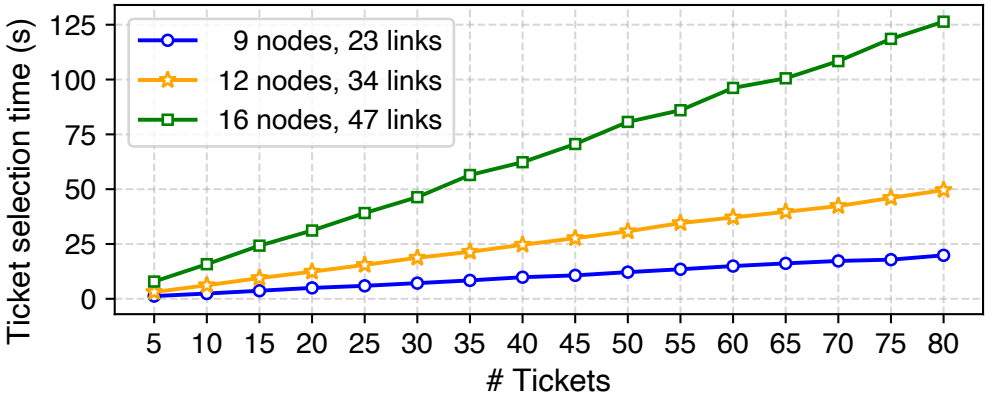

* Current TE always predict traffic 1-step ahead (Smore NSDI'18), we extend it to prediction of *T* steps.

Our Design: Archie, An End-to-End Learning-Based Method

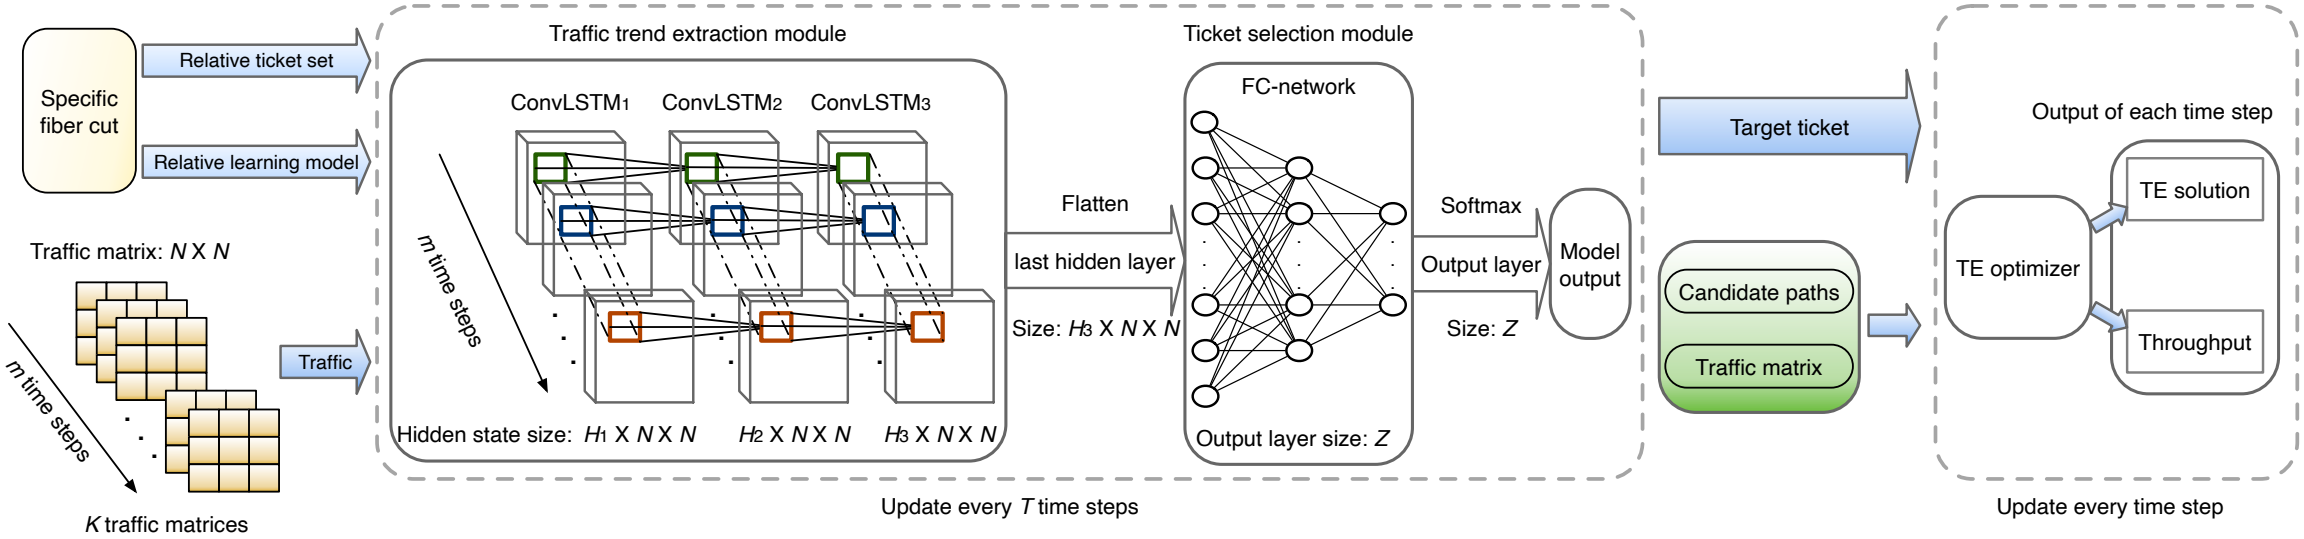

(a) *T*-step restoration ticket selector (ConvLSTM based classifier)

(b) Normal TE system

- Models are *trained* ***offline*** and only conduct *inference* ***online***.
- **Label generation** of *supervised learning*: Using the actual traffic demand of next *T* steps.

Intuitive advantages of Archie:

1. **Faster decision time**: one pass of the model inference instead of solving TE optimizations *TZ* times.
2. **Robust to inherent prediction inaccuracy**: end-to-end learning copes with traffic uncertainties as long as the selected ticket is correct, instead of requiring accurate demand prediction for all *T* steps.
   - We will show insights why Archie outperforms Demand Prediction with analysis evaluation afterwards.

Evaluation Setups

- Evaluated topologies and traffic traces

Topology	# Nodes	# Fibers	# IP links	# Traffic matrices (Train + Test)
Arpanet	9	10	23	2880+120
Abilene	12	15	34	
Airtel	16	26	47	
GRnet	37	42	101	

- Evaluation setups
  - *Ticket setup:* Z=30 by default for each fiber cut scenario
  - *TE setup:* K=4 shortest path for each flow, one time step is 5 minute
  - *Fiber cut scenario:* Random one fiber cut

Performance of Archie: Throughput Loss

- Settings: *T* is fixed to be 10, reconfiguration overhead is not considered currently.

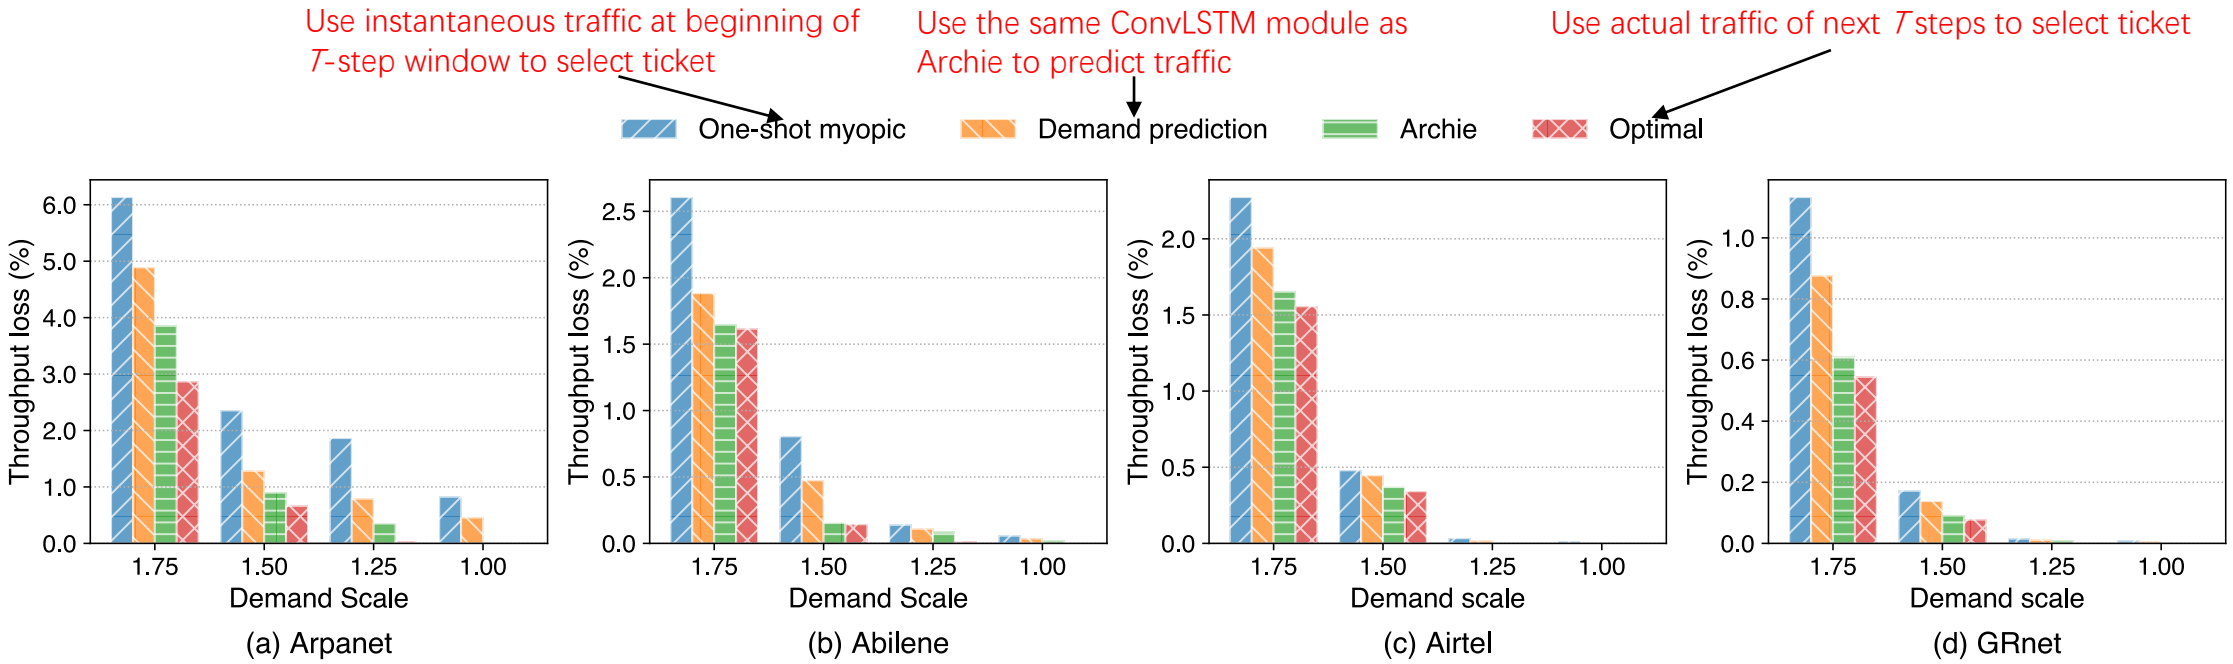

- Near-to-optimal performance: additional throughput loss only range from 0% to 0.98% at most.
- Reduce 48.5% and 27.1% throughput loss compared to One-shot myopic and Demand prediction, respectively.
- Performance are robust to all demand scales and topology scales.

Performance of Archie: Ticket Selection Time

Topology	T	Ticket selection time (seconds)		
		One-shot myopic	Demand prediction	Archie
Arpanet	5	0.729	3.644	0.01875
	10	0.721	7.102	0.01900
	15	0.732	10.982	0.01894
Abilene	5	1.845	9.221	0.02544
	10	1.852	18.001	0.02501
	15	1.847	27.705	0.02539
Airtel	5	4.746	23.712	0.03011
	10	4.745	48.003	0.02996
	15	4.751	71.265	0.03008
GRnet	5	51.974	259.182	0.04355
	10	52.123	522.313	0.04441
	15	51.988	779.320	0.04483

- Neat-to-zero selection time: within 45ms in all our test cases.
- Speedups are 362x and 3598x compared to One-shot myopic and Demand prediction on average.
- Good scalability: Growth speed of decision time in Archie is slow when the topology scales.

Performance of Archie: Candidate Ticket Number

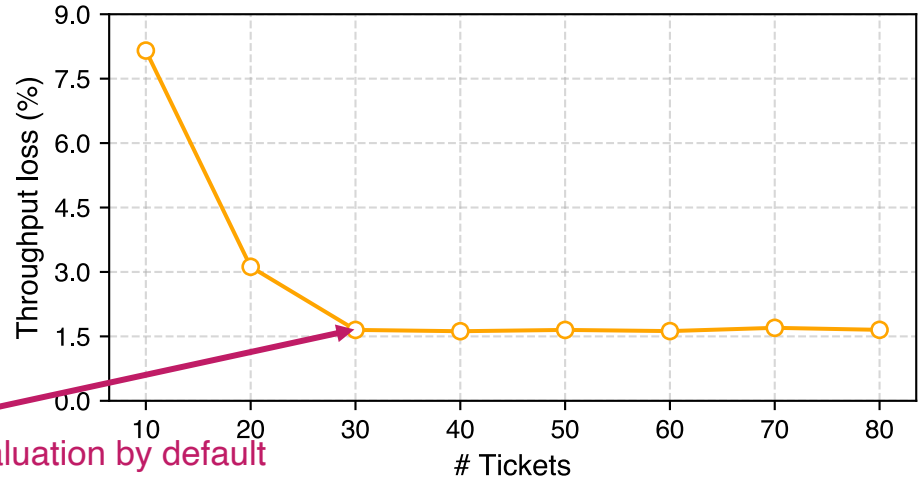

Elect 30 candidate tickets in evaluation by default

- In Abilene topology

- When tickets are not enough (≤ 30), ticket set may not cover a good ticket, adding tickets may improve performance.
- When exceeding 30 tickets, performance stop improving obviously. More tickets increase burden of model preparation.

Benefit of *T*-Step Dynamic Restoration

- Settings: Reconfiguration time is set to be 1/30 steps (10s).

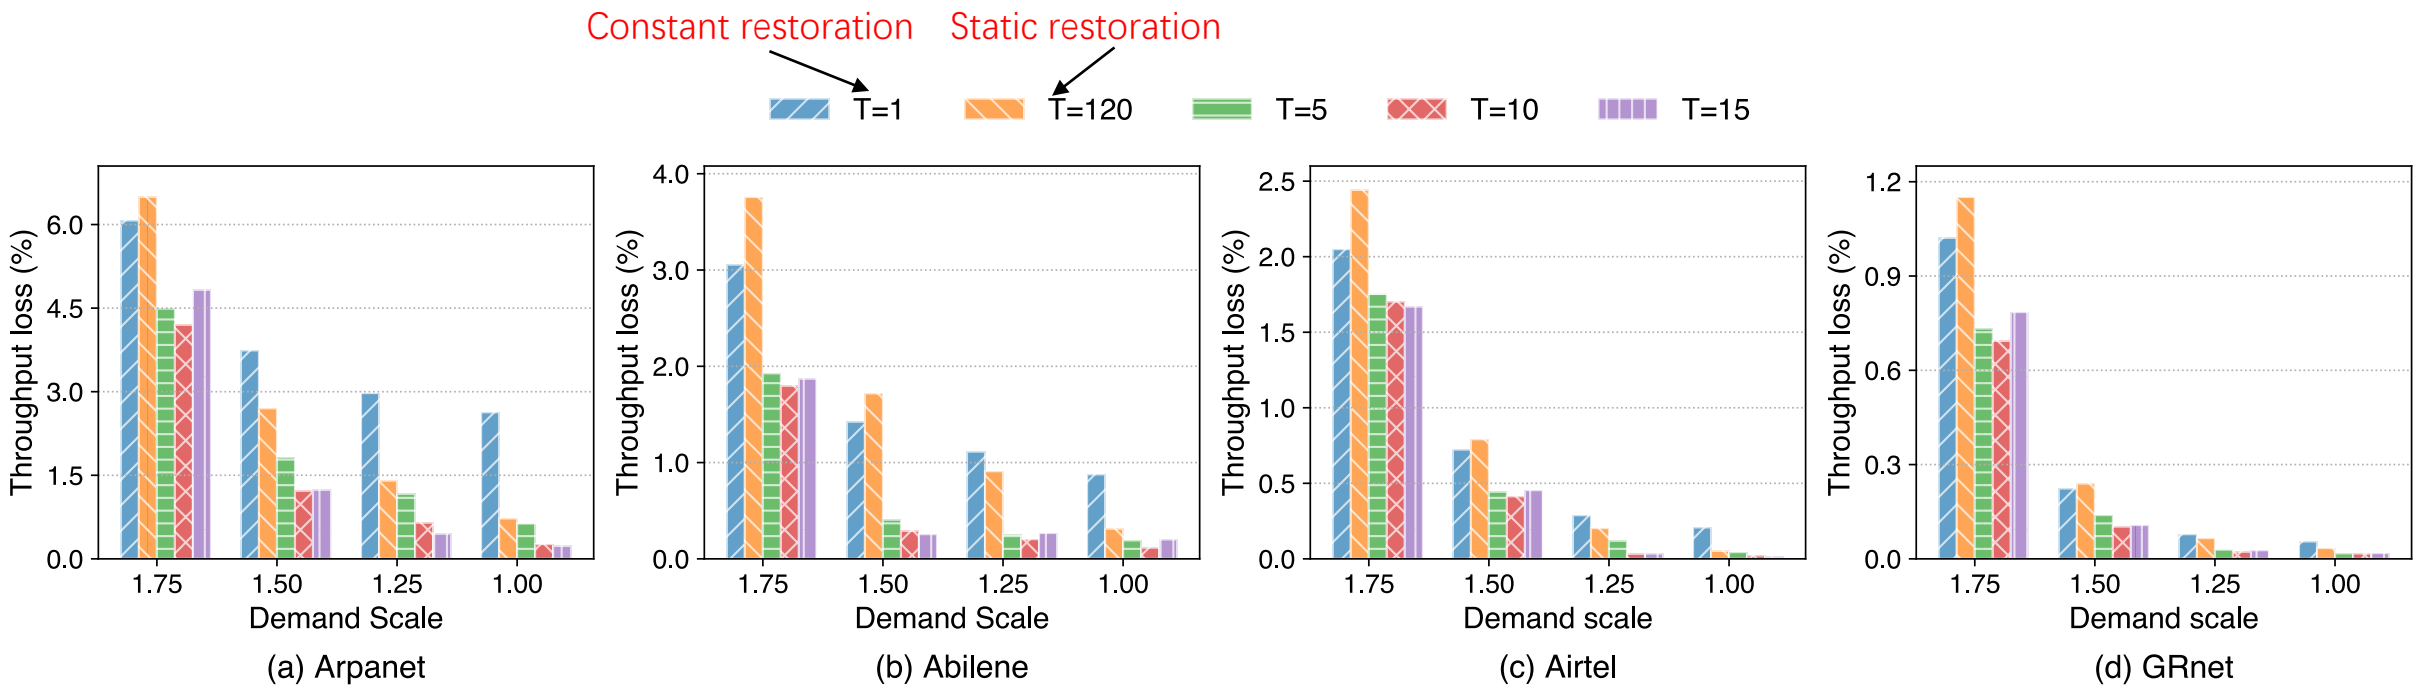

- Archie under best *T* reduce throughput loss by 64.7% and 59.6% compared to constant and static restoration.
- Archie with there moderate settings of *T* (5, 10, 15) has small performance difference less than 0.2%.

Offline Model Preparation Time of Archie

- Settings:
  - $T = 10$, $Z = 30$, concurrent threads $P = 192$ for generating $O(1000)$ labels. Needs $O\left(\frac{1000TZ}{P}\right)$ times TE solving.
  - One GPU for model training for all one fiber cut scenarios.

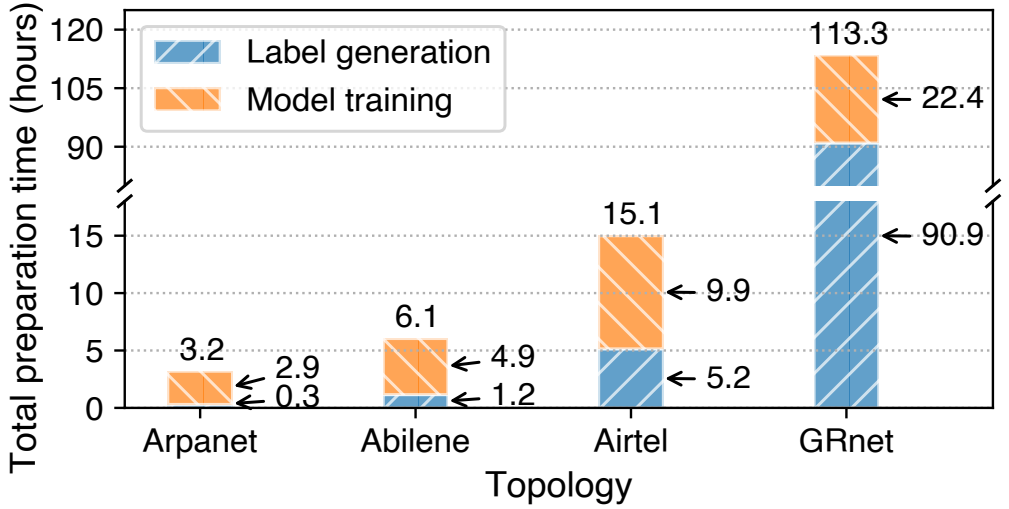

- Offline time investment is acceptable for real-world use:
  - Common medium topology (Abilene) can finish preparation within 6 hours.
  - Large topology (GRnet) can finish within 5 days.
- Time can be further reduced by leveraging more GPUs and concurrent threads to *enhance parallelism*.

Insight from Archie: Spatial Feature Analysis

Spatial Feature: Does Archie pay more attention to specific flows?

- **Method: Occlusion analysis** (Occludes evaluated traffic input part and evaluate the performance gap)
  - Clear demand to 0 for occluded flows.

- Evaluate traffic epochs whose ground truth is the following tickets.
- Occlude all flows traversing some links in the topology (IP1, IP2, IP3), respectively.

Topology	Information item	$ticket_1$			$ticket_2$		
		IP_1	IP_2	IP_3	IP_1	IP_2	IP_3
G_1	# Wavelengths	4	1	1	3	2	1
G_2	# Wavelengths	8	2	2	2	8	2

- Ticket Selection: occluded past traffic
- Throughput Calculation: original future traffic

- Performance Result of Occlusion

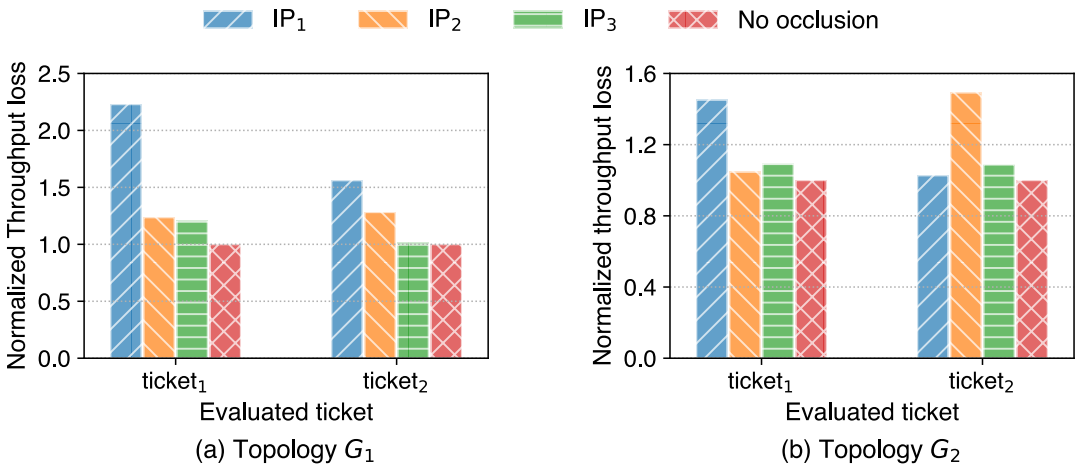

- For ticket 1 in G1, Traffic flows traversing IP1 are the more important than others (in future traffic).
- Archie assign more restored wavelengths to the corresponding links of the flows.

- Archie can identify parts of critical future flows, thus assign more restored wavelengths to corresponding links.

Insight from Archie: Spatial Feature Analysis

Spatial Feature: Does Archie's ability to identify critical flows makes it superior to Demand Prediction?

- Introduce the same flow importance in Archie to Demand Prediction.
  - Modify the MSE loss when training:

$$J = \frac{1}{T \times F} \sum_{t}^{T} \sum_{f}^{F} \alpha_f (y_{ft} - h_{ft})^2$$

  - Traditional MSE: α_f=1 for all flow *f*. (Unweighted)

- Weight assignment example

Topology	Information item	ticket$_1$			ticket$_2$		
		IP$_1$	IP$_2$	IP$_3$	IP$_1$	IP$_2$	IP$_3$
G_1	# Wavelengths	4	1	1	3	2	1
	Weights$_1$ (α_f)	2.0	1.0	1.0	2.0	1.5	1.0
	Weights$_2$ (α_f)	3.0	1.0	1.0	3.0	2.0	1.0
G_2	# Wavelengths	8	2	2	2	8	2
	Weights$_1$ (α_f)	2.0	1.0	1.0	1.0	2.0	1.0
	Weights$_2$ (α_f)	3.0	1.0	1.0	1.0	3.0	1.0

- Performance Result

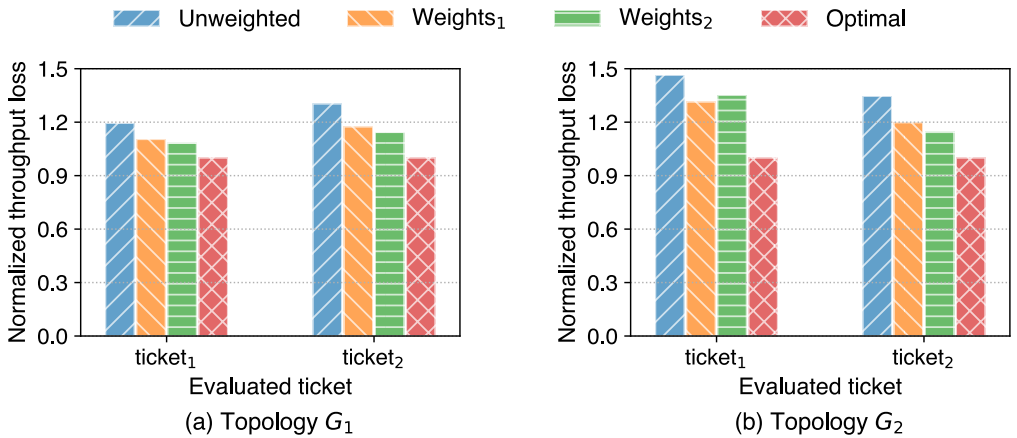

- Archie's ability to identify critical future flows is one of the reasons why it outperforms Demand prediction.
  - Performance: Optimal > Weight 1, Weight 2 > Unweighted

Insight from Archie: Temporal Feature Analysis

Temporal Feature: Does Archie identify to any special traffic patterns?

- Method: LSTMVis (check important time interval)

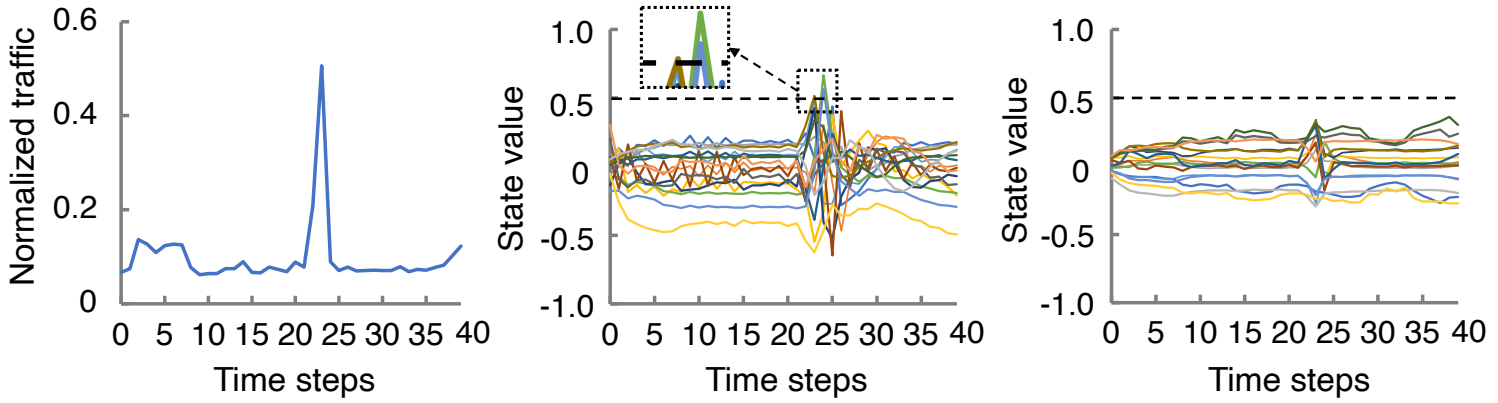

(a) Traffic demand (b) States of Archie (c) States of Demand Prediction

Interval from 20 to 25 has a traffic spike

Archie pays attention to this traffic spike

Demand Prediction does not

- Input past traffic length: 40
- Hidden states dimension of ConvLSTM: 16

- Archie can identify traffic spikes!

Insight from Archie: Spatial Feature Analysis

Temporal Feature: Does this traffic spike feature help ticket selection?

- Evaluate the performance without spike identification feature for Archie
  - Use ConvLSTM module parameters in Demand Prediction for Archie and freeze these parameters.
  - Retrain FC network in Archie to obtain a new model, it is called ***Hybrid***.

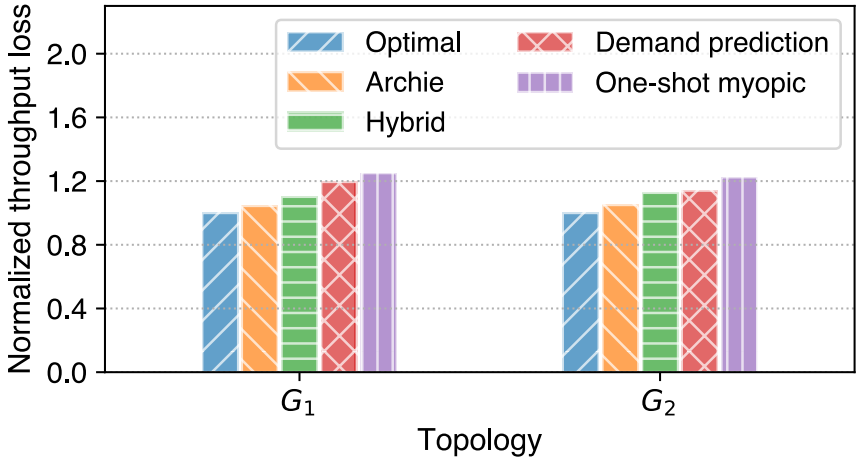

- Traffic spike feature in Archie does help ticket selection
  - **Archie outperforms Hybrid:** Identify traffic spikes
  - **Hybrid outperforms Demand Prediction:** Archie has other advantages for ticket selection (Identify vital flows).

Thank you!

Q&A