Precise Data Center Traffic Engineering with Constrained Hardware Resources

Shawn Shuoshuo Chen[1]

Keqiang He[2], Rui Wang[3], Srinivasan Seshan[1], Peter Steenkiste[1]

[1] Carnegie Mellon University

[2] airbnb

[3] Google

What does a data center network look like?

What does a data center network look like?

Google data center in Douglas County, Georgia.

What does a data center network look like?

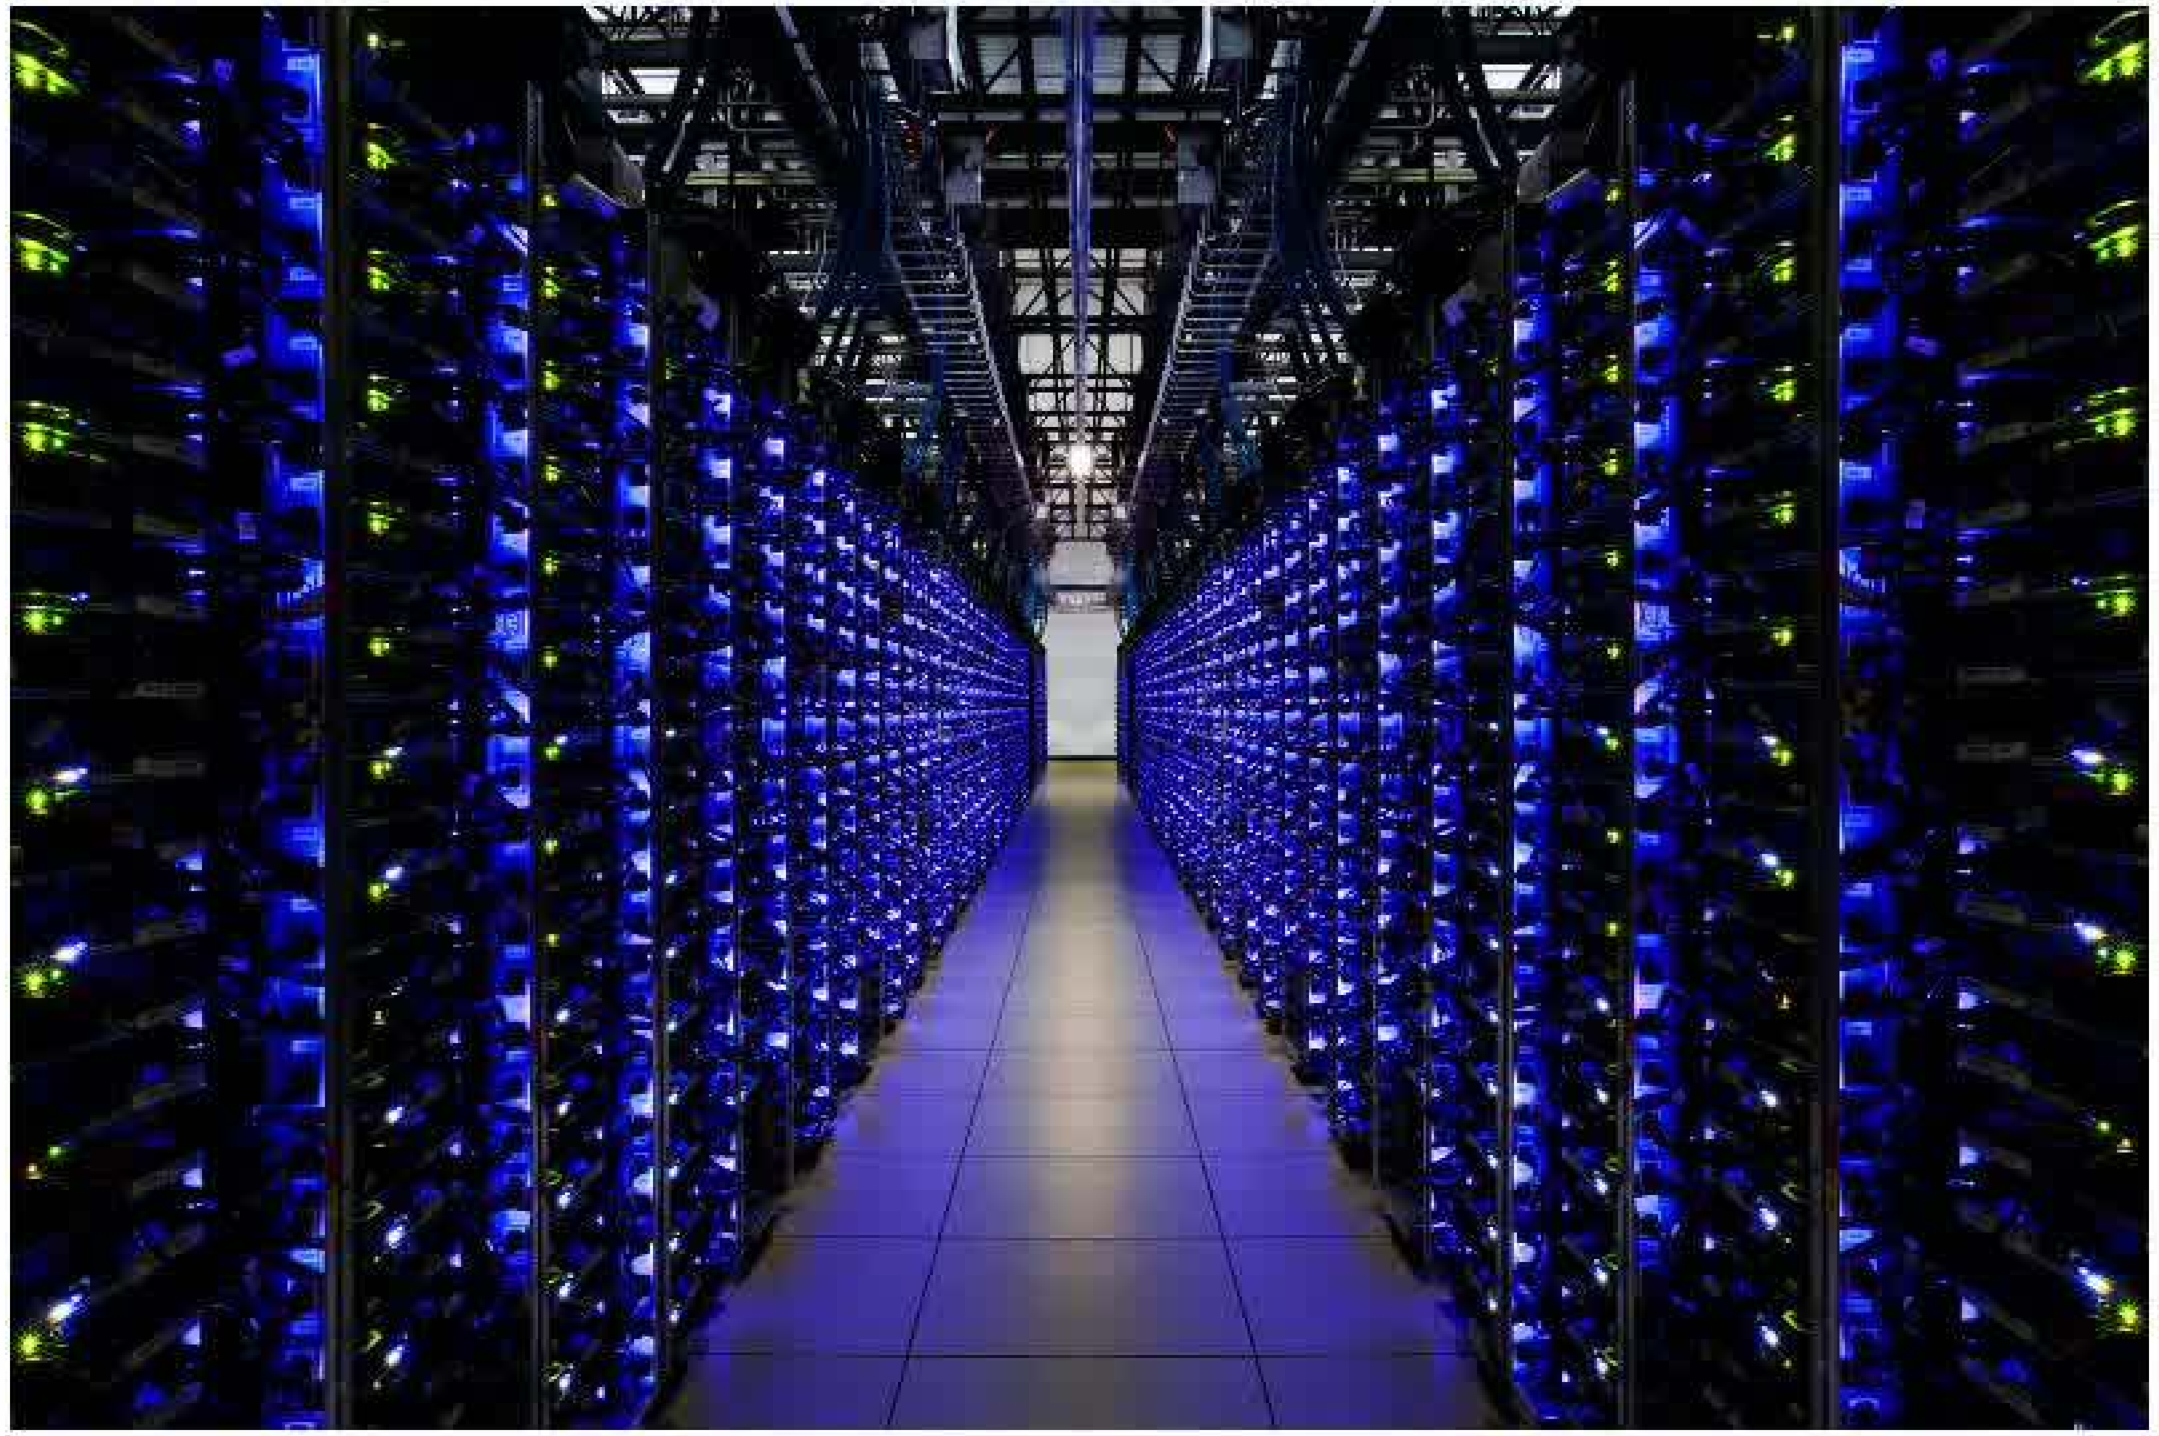

Racks

Google data center in Douglas County, Georgia.

What does a data center network look like?

Network wires

Racks

Google data center in Douglas County, Georgia.

What does a data center network look like?

Network wires

Racks

Google data center in Douglas County, Georgia.

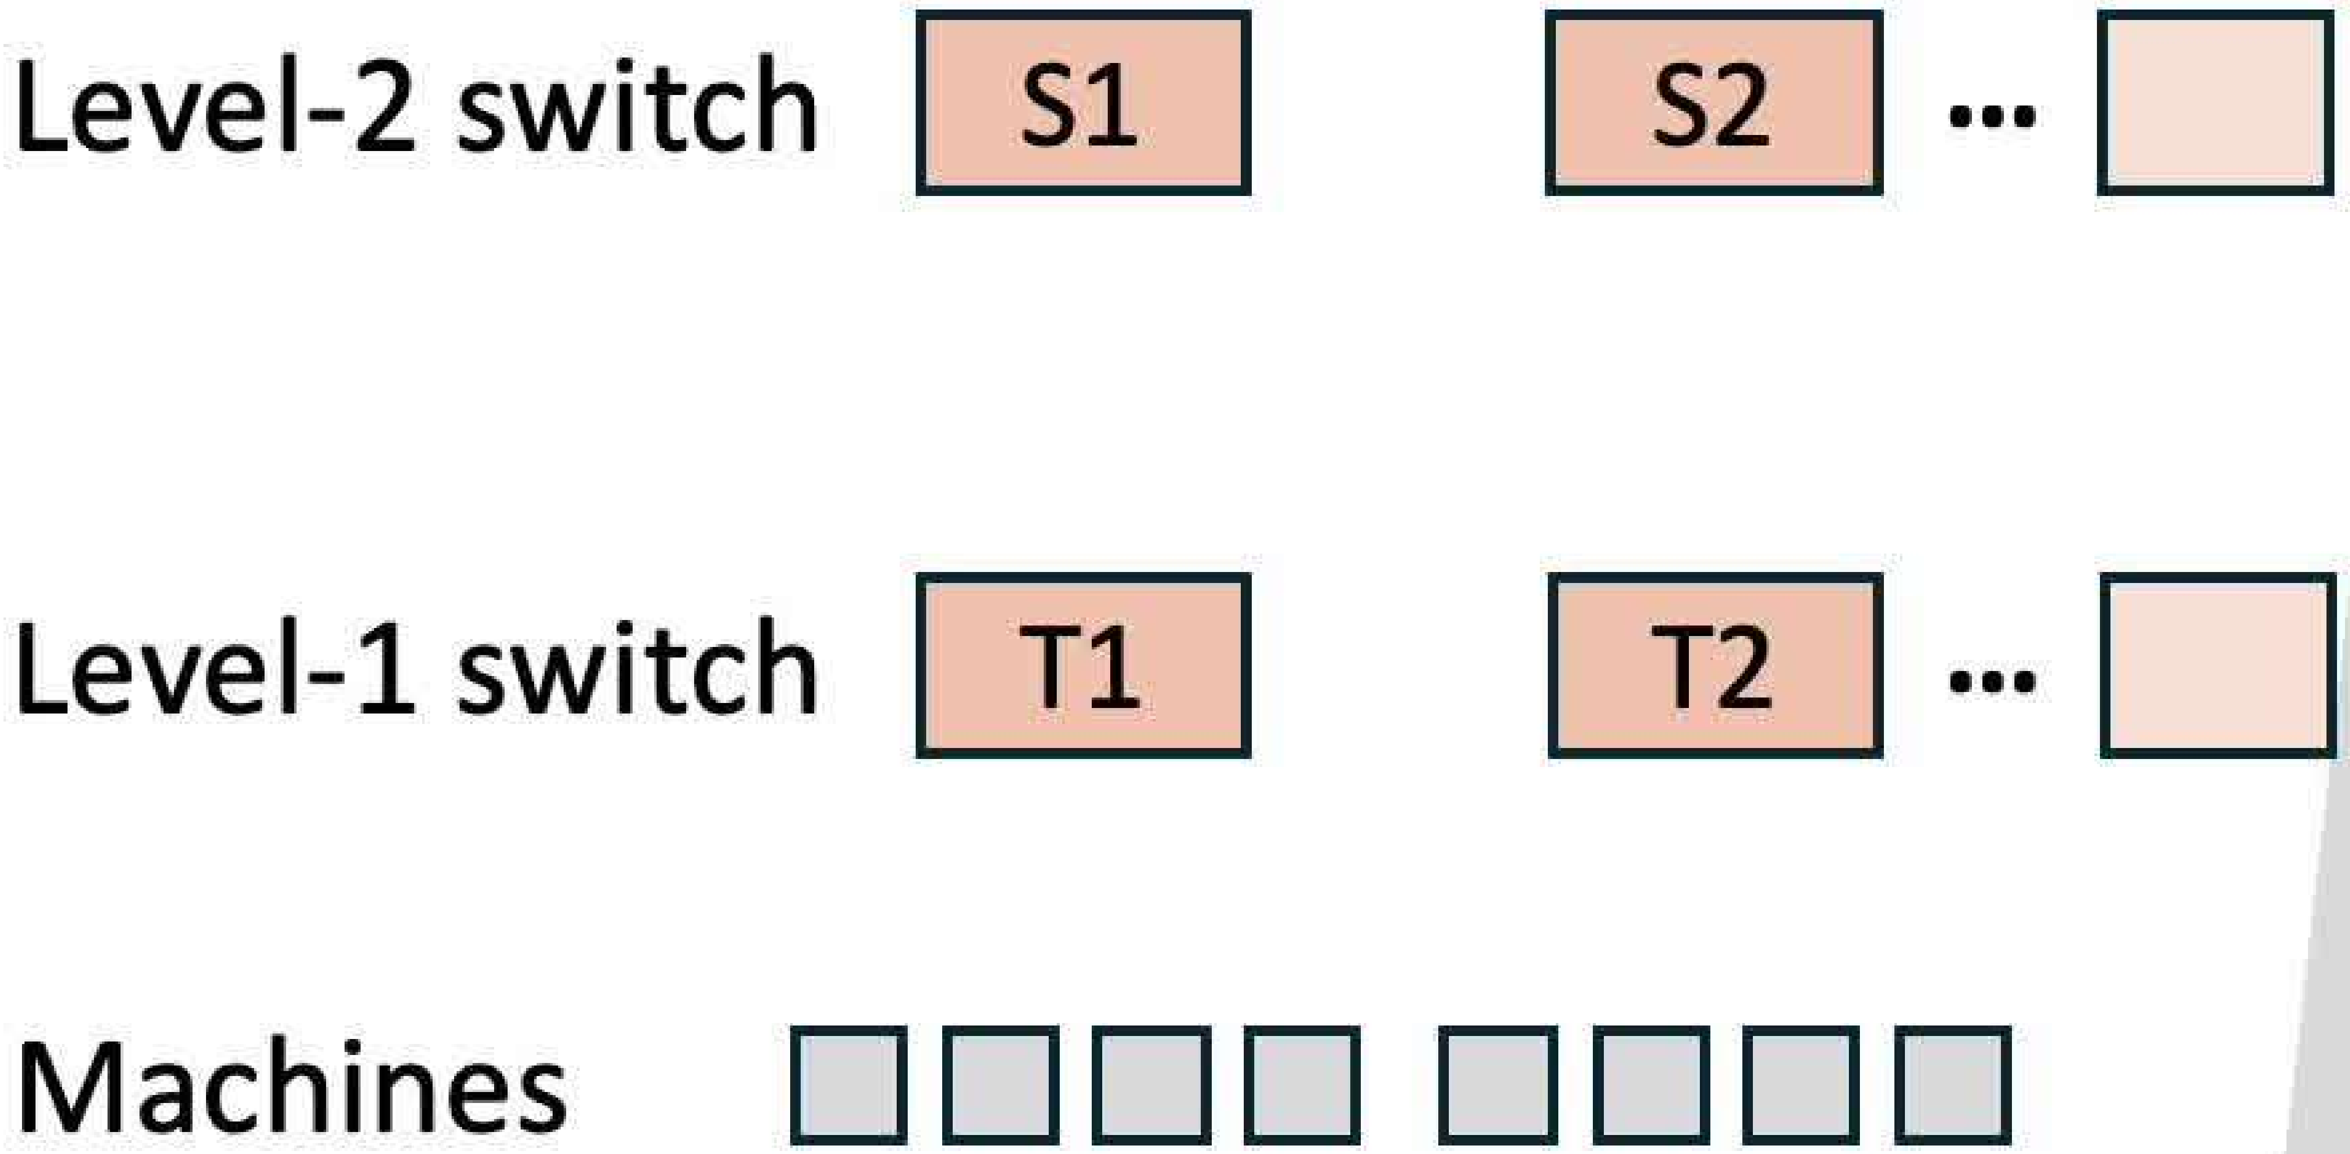

What does a data center network look like?

Google data center in Douglas County, Georgia.

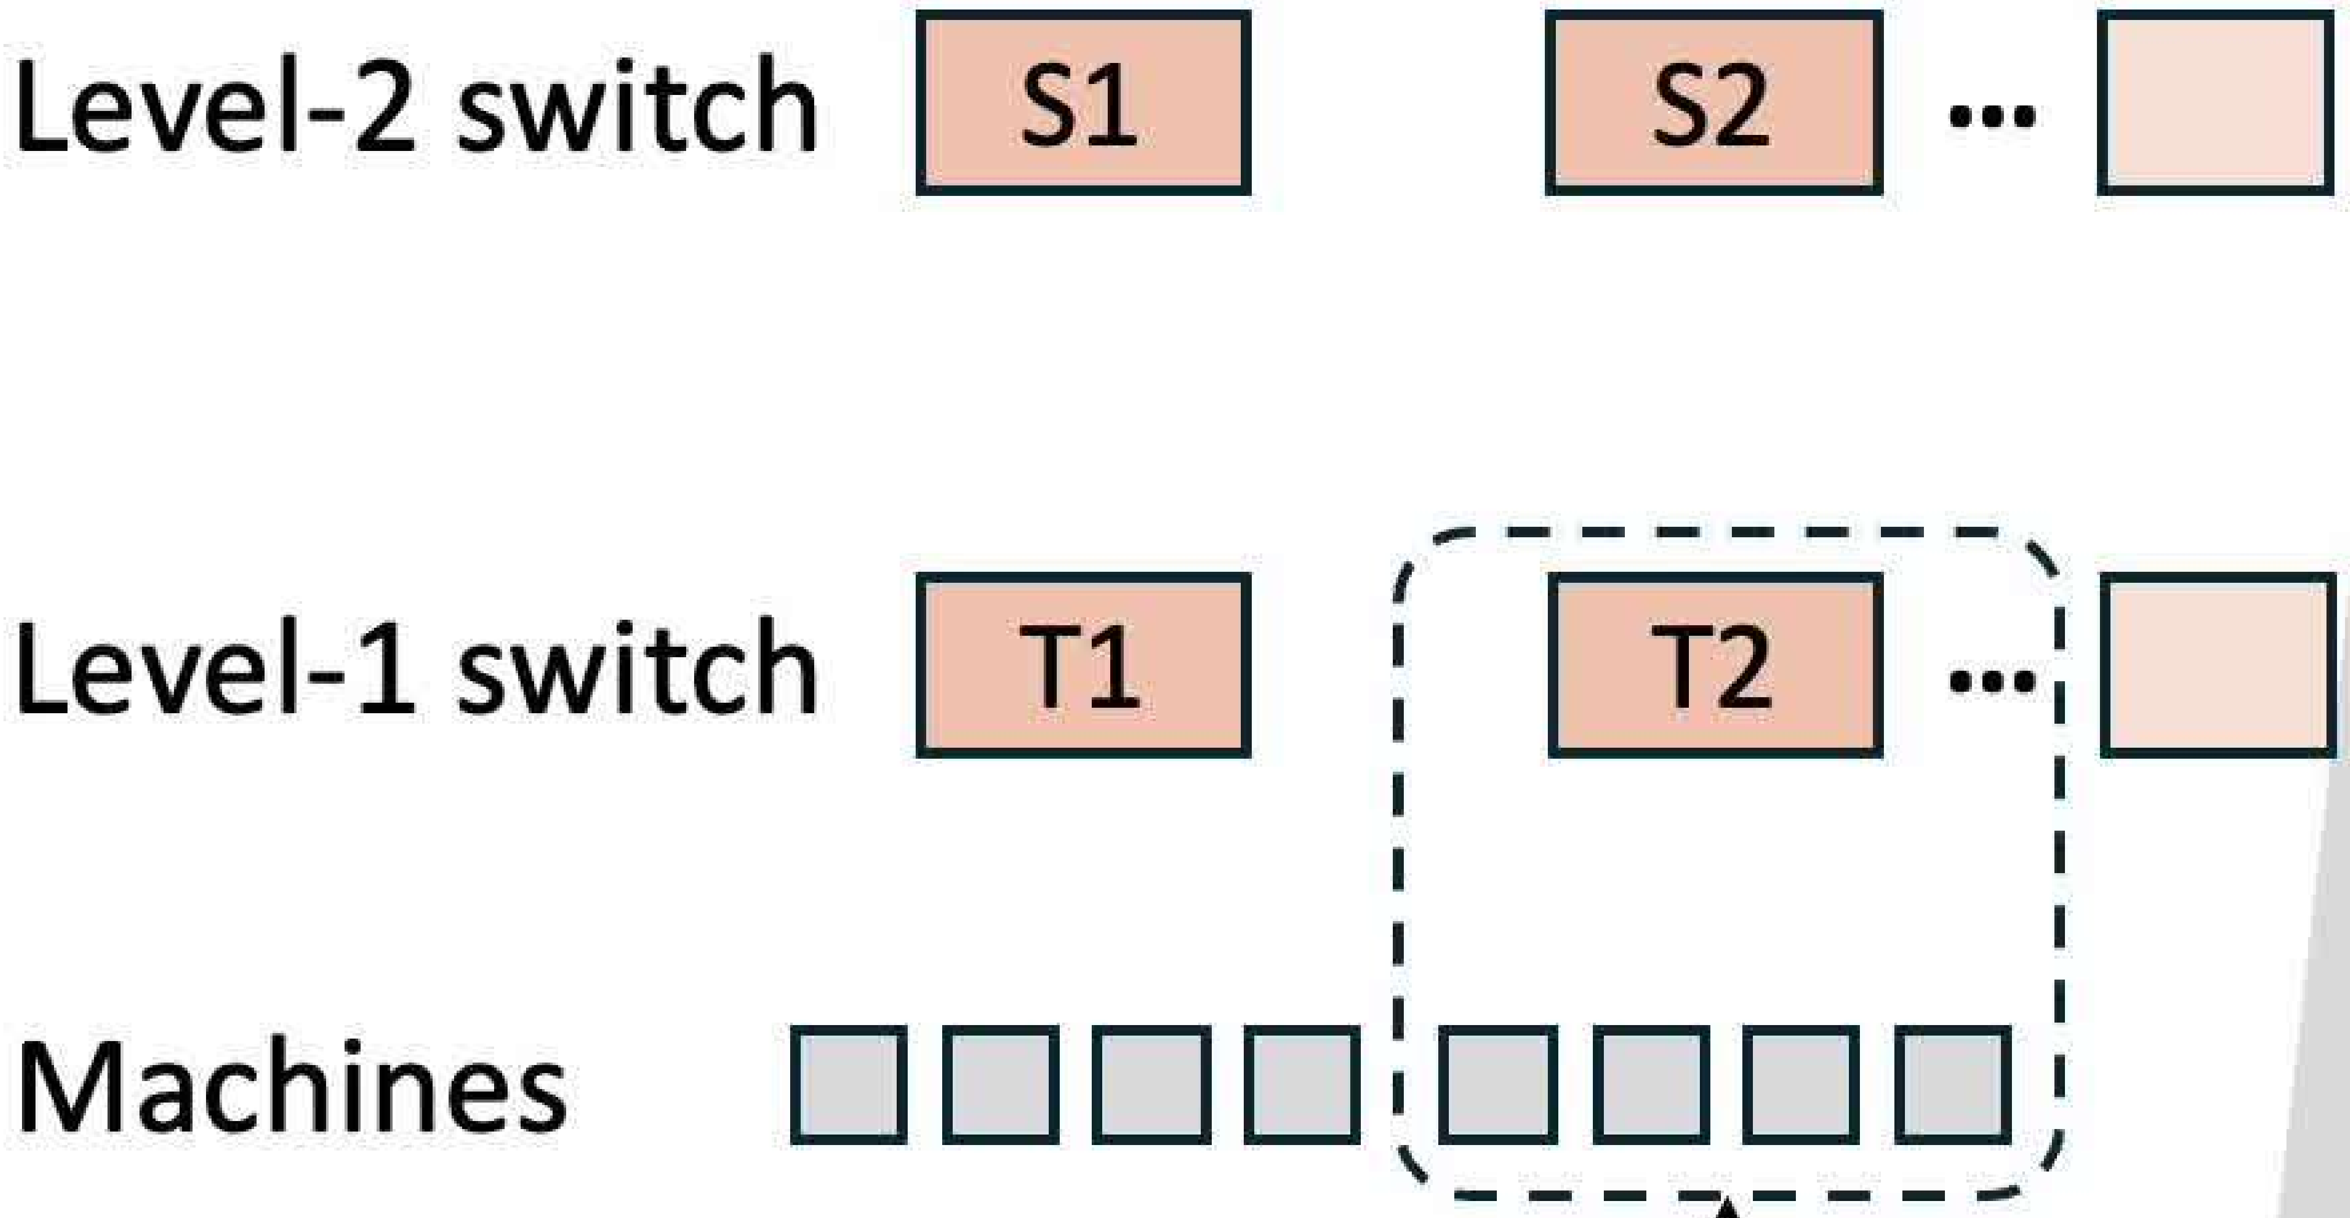

What does a data center network look like?

Network wires

Racks

Google data center in Douglas County, Georgia.

Level-2 switch

Level-1 switch

Machines

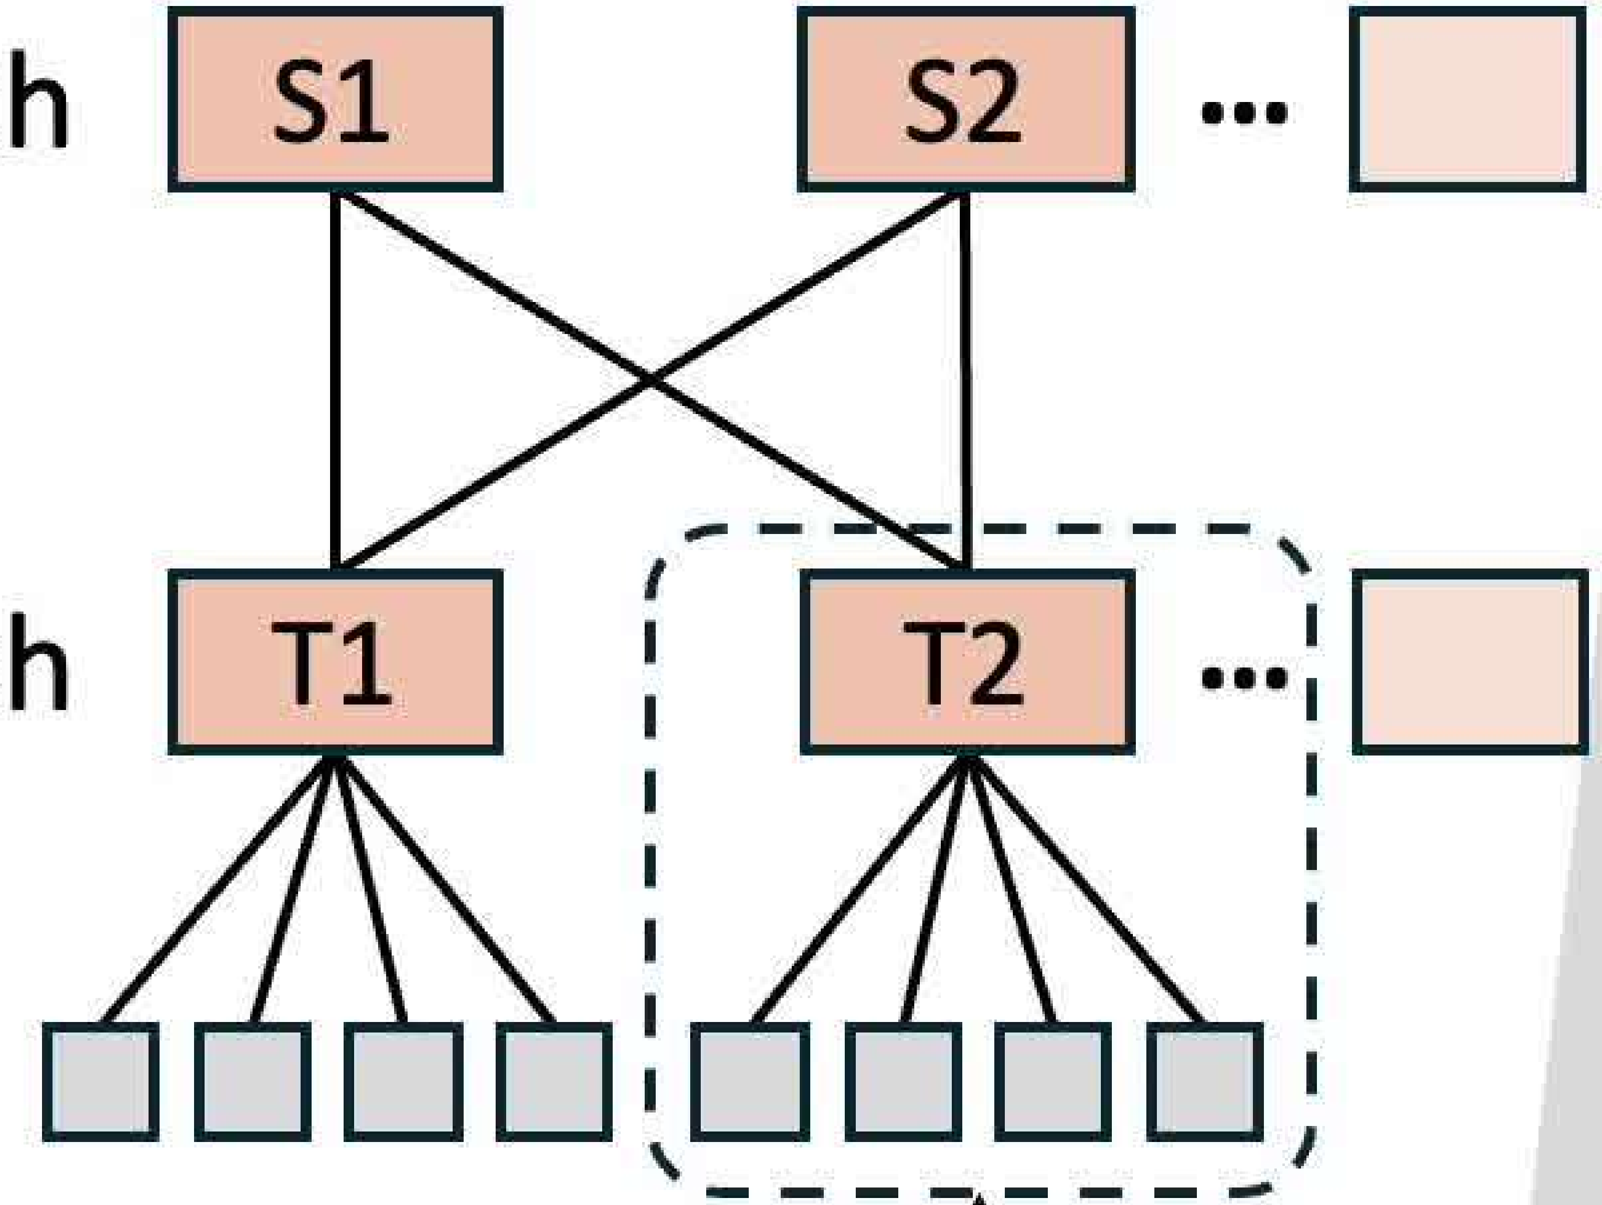

How does a traffic engineering (TE) system work?

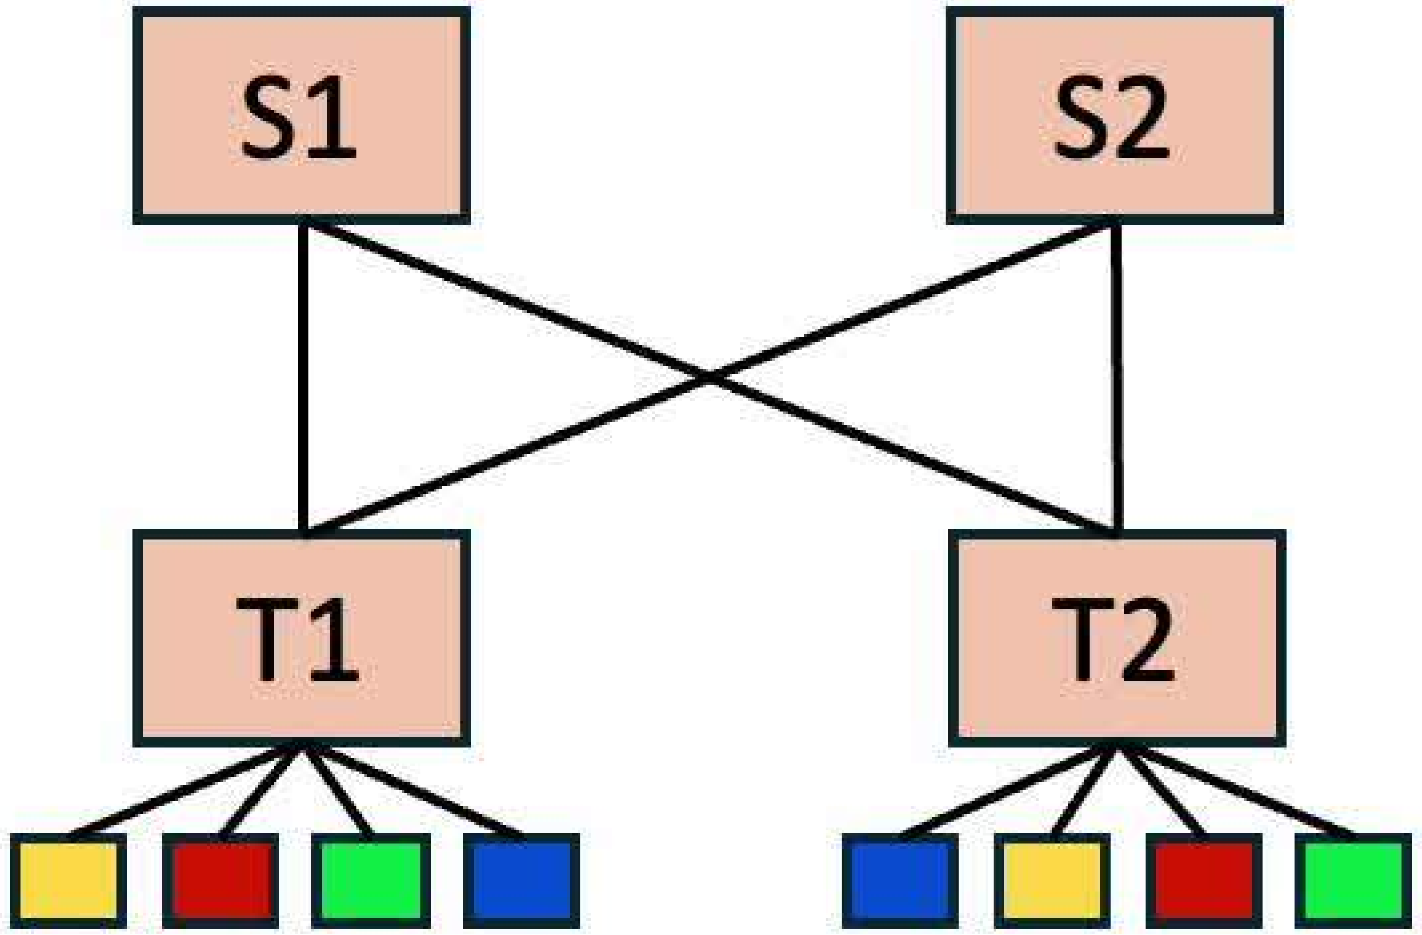

How does a traffic engineering (TE) system work?

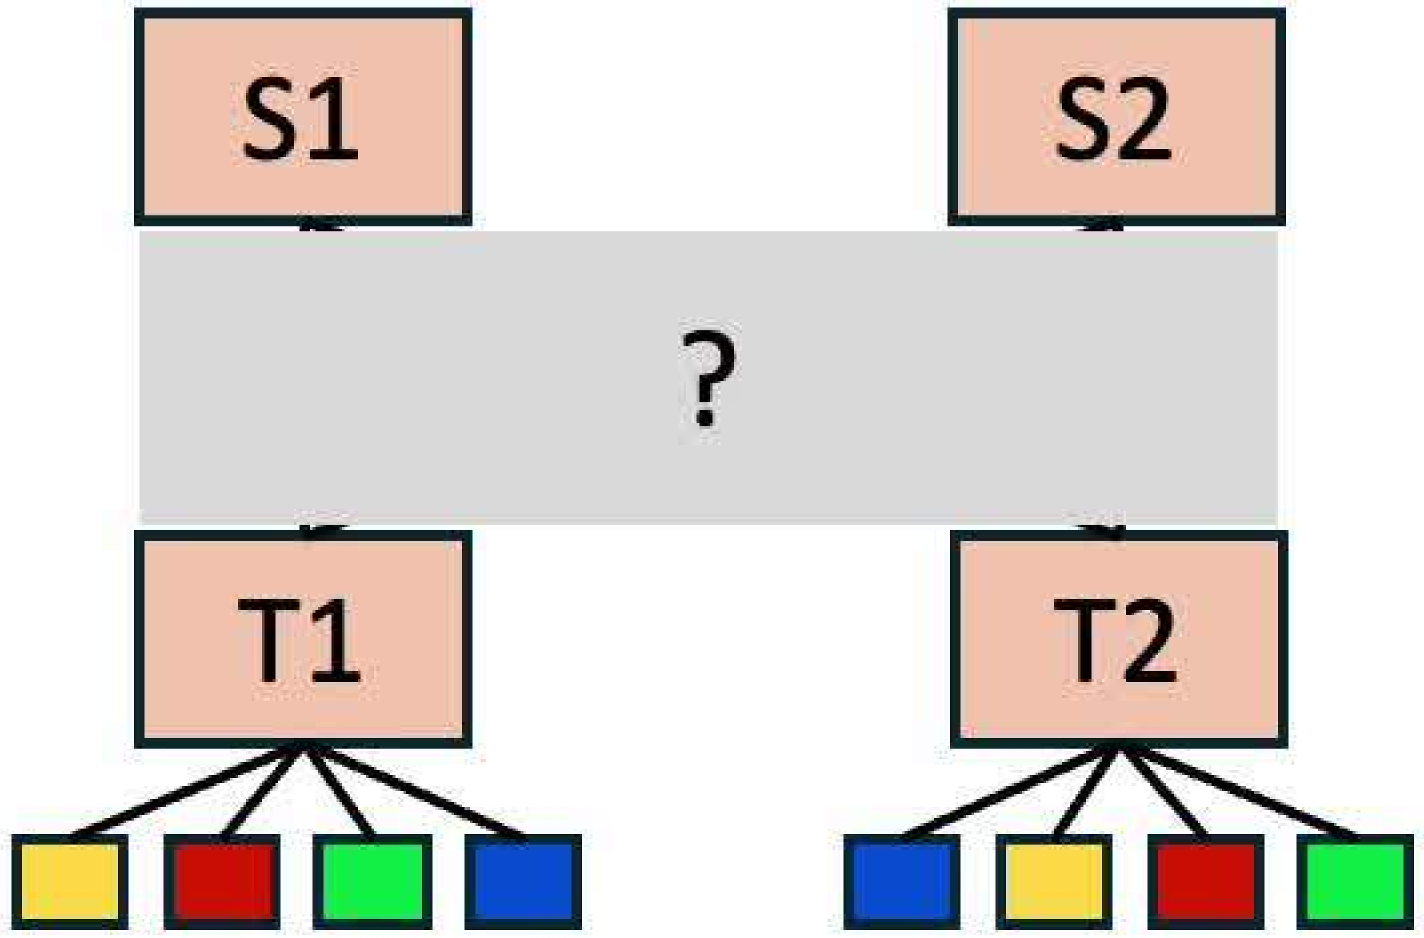

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

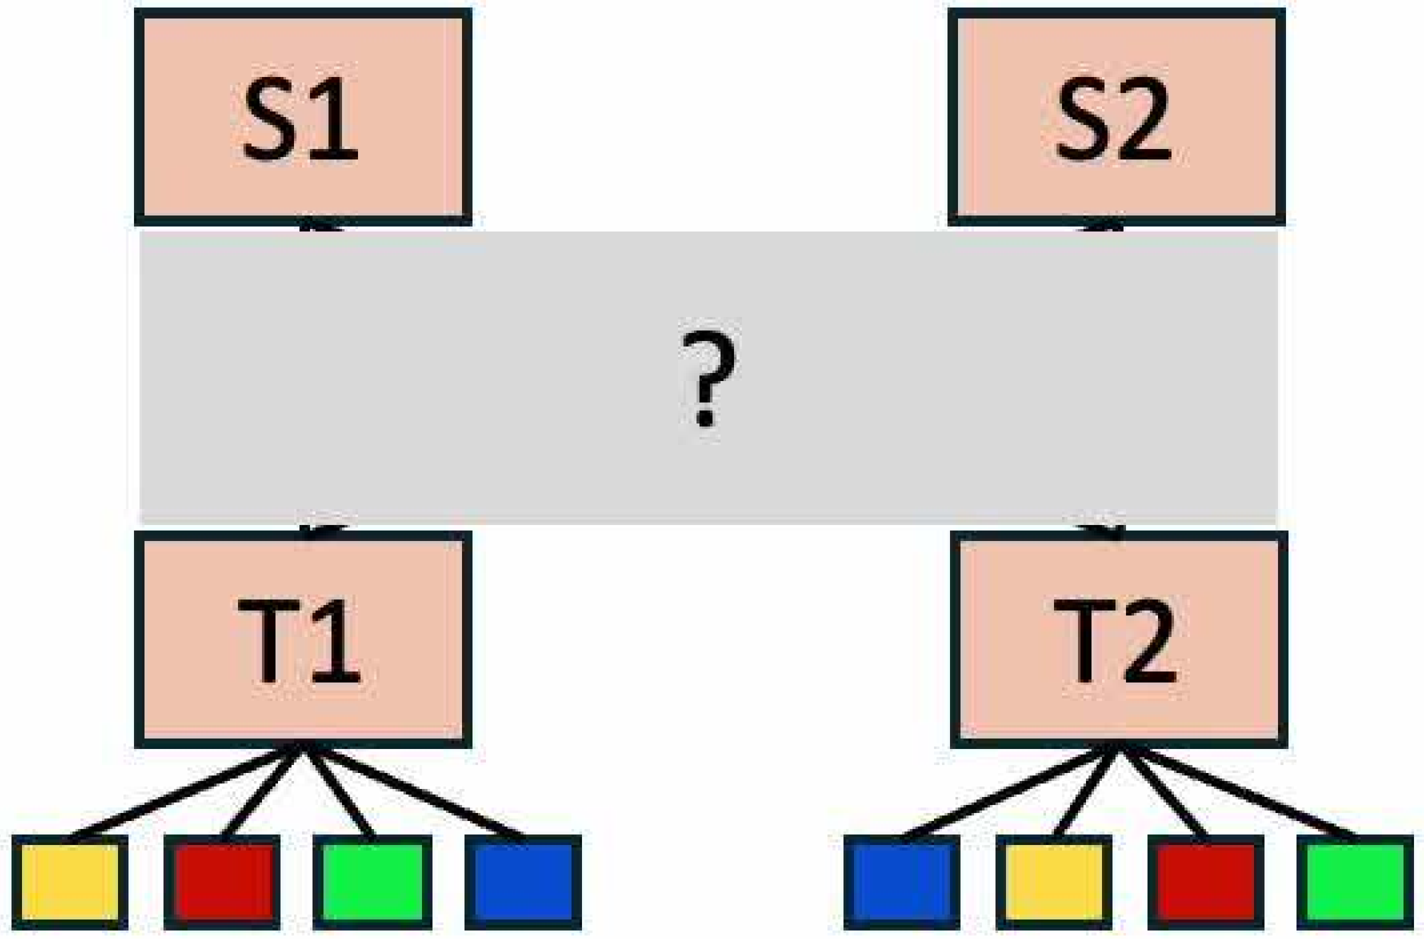

Carnegie Mellon University

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

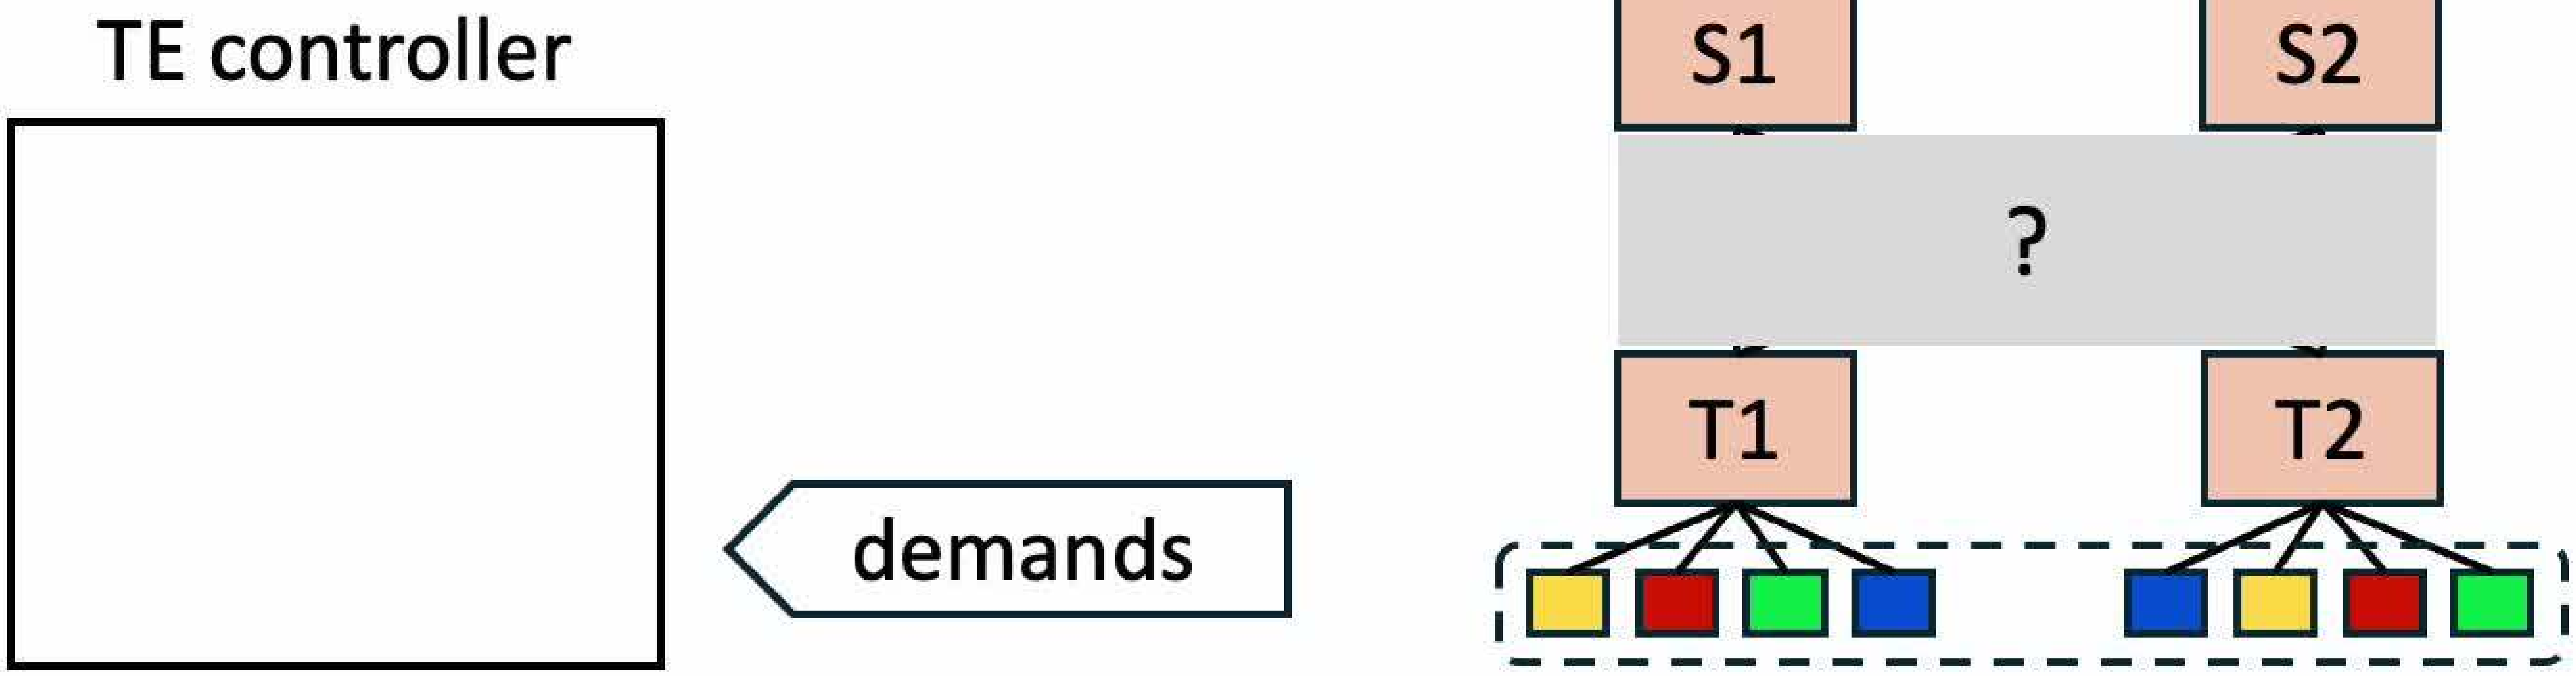

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

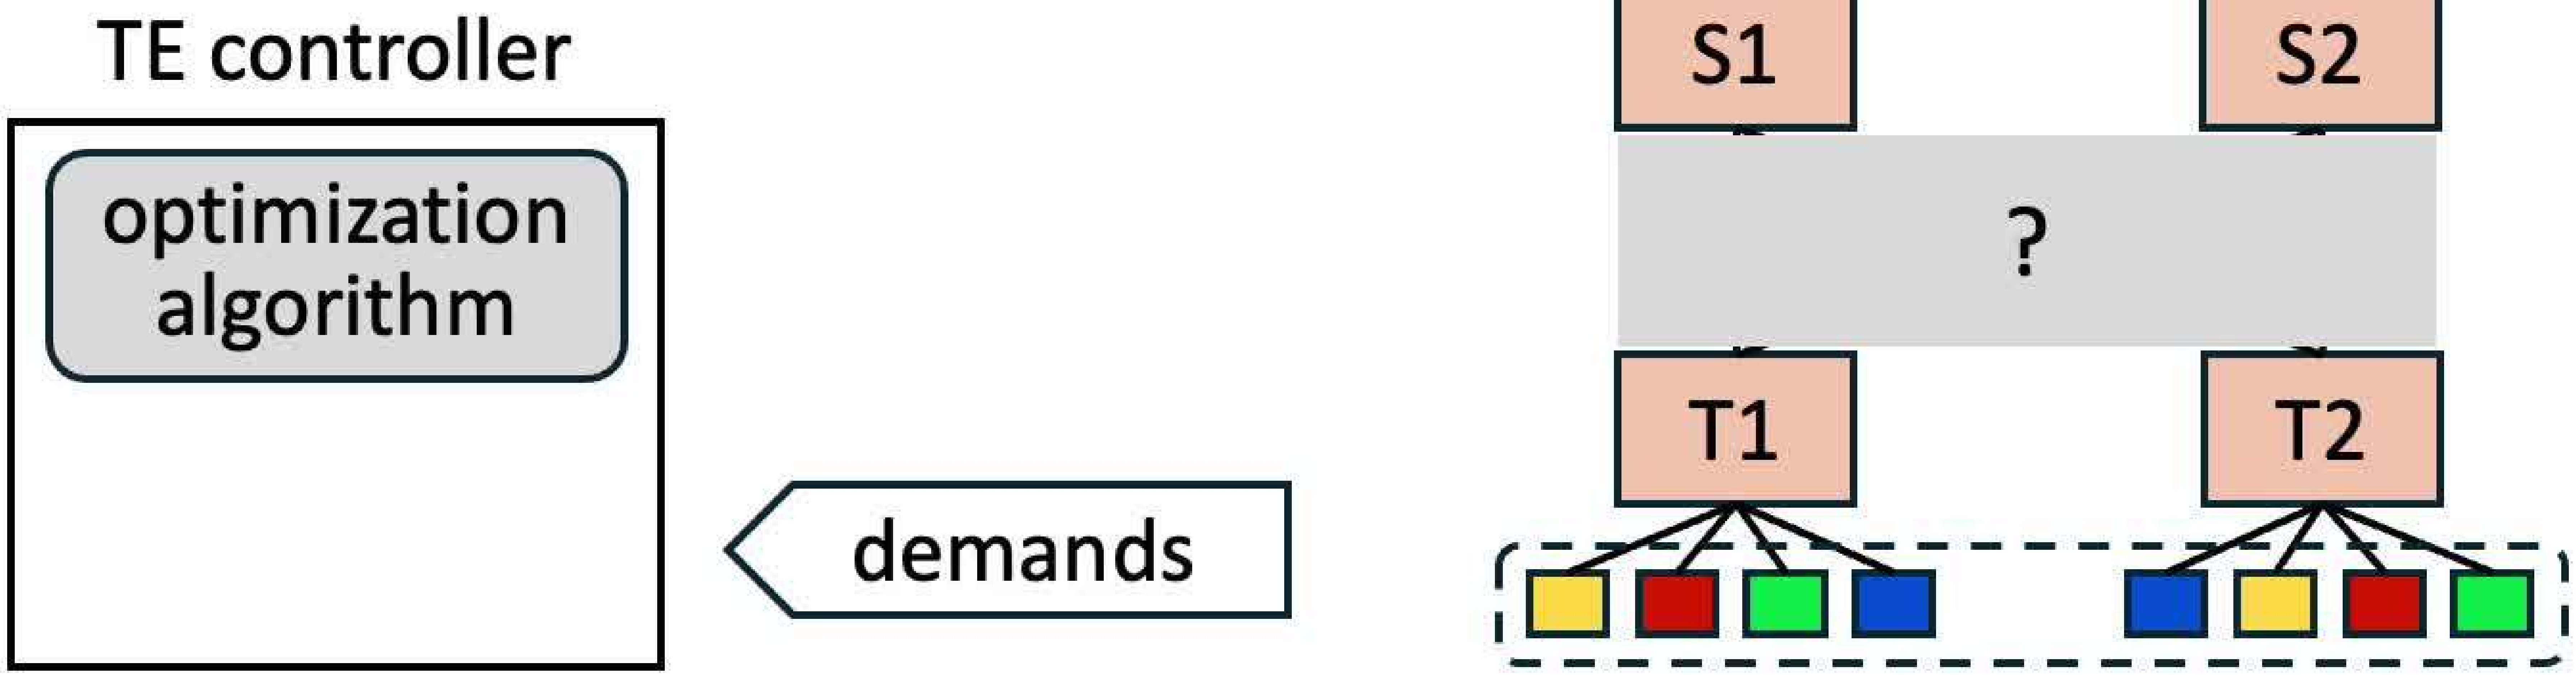

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

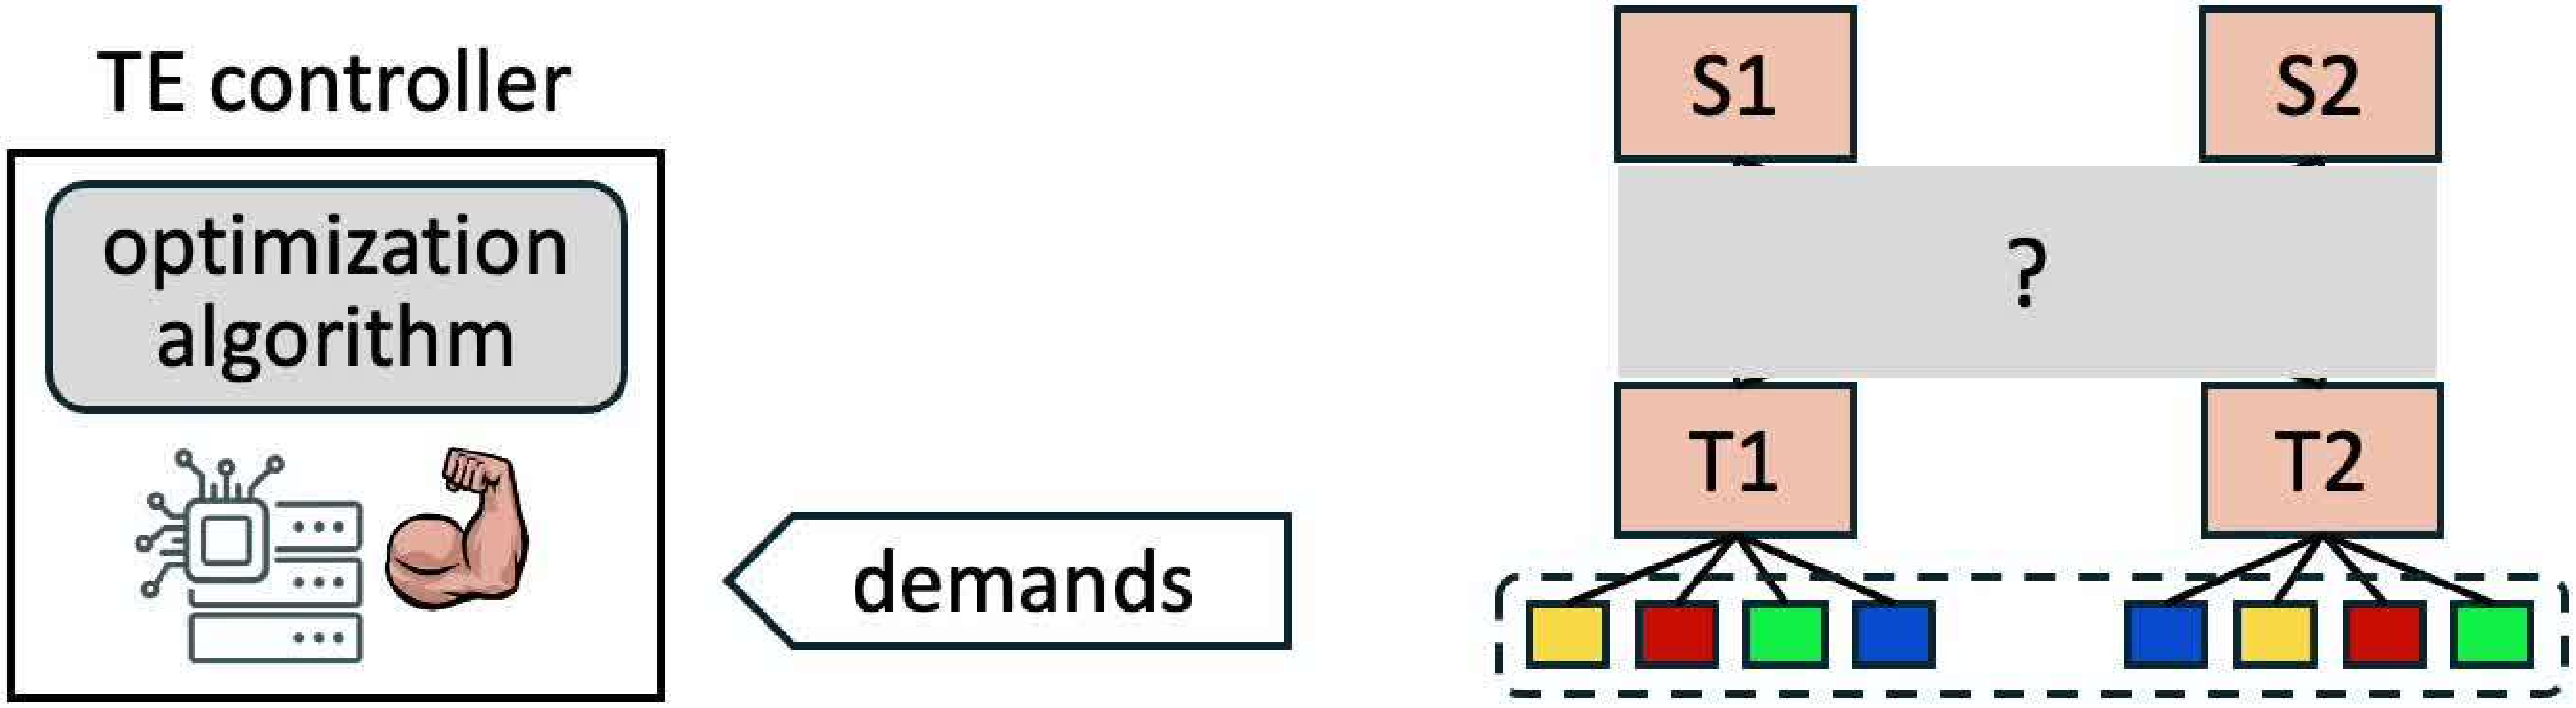

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

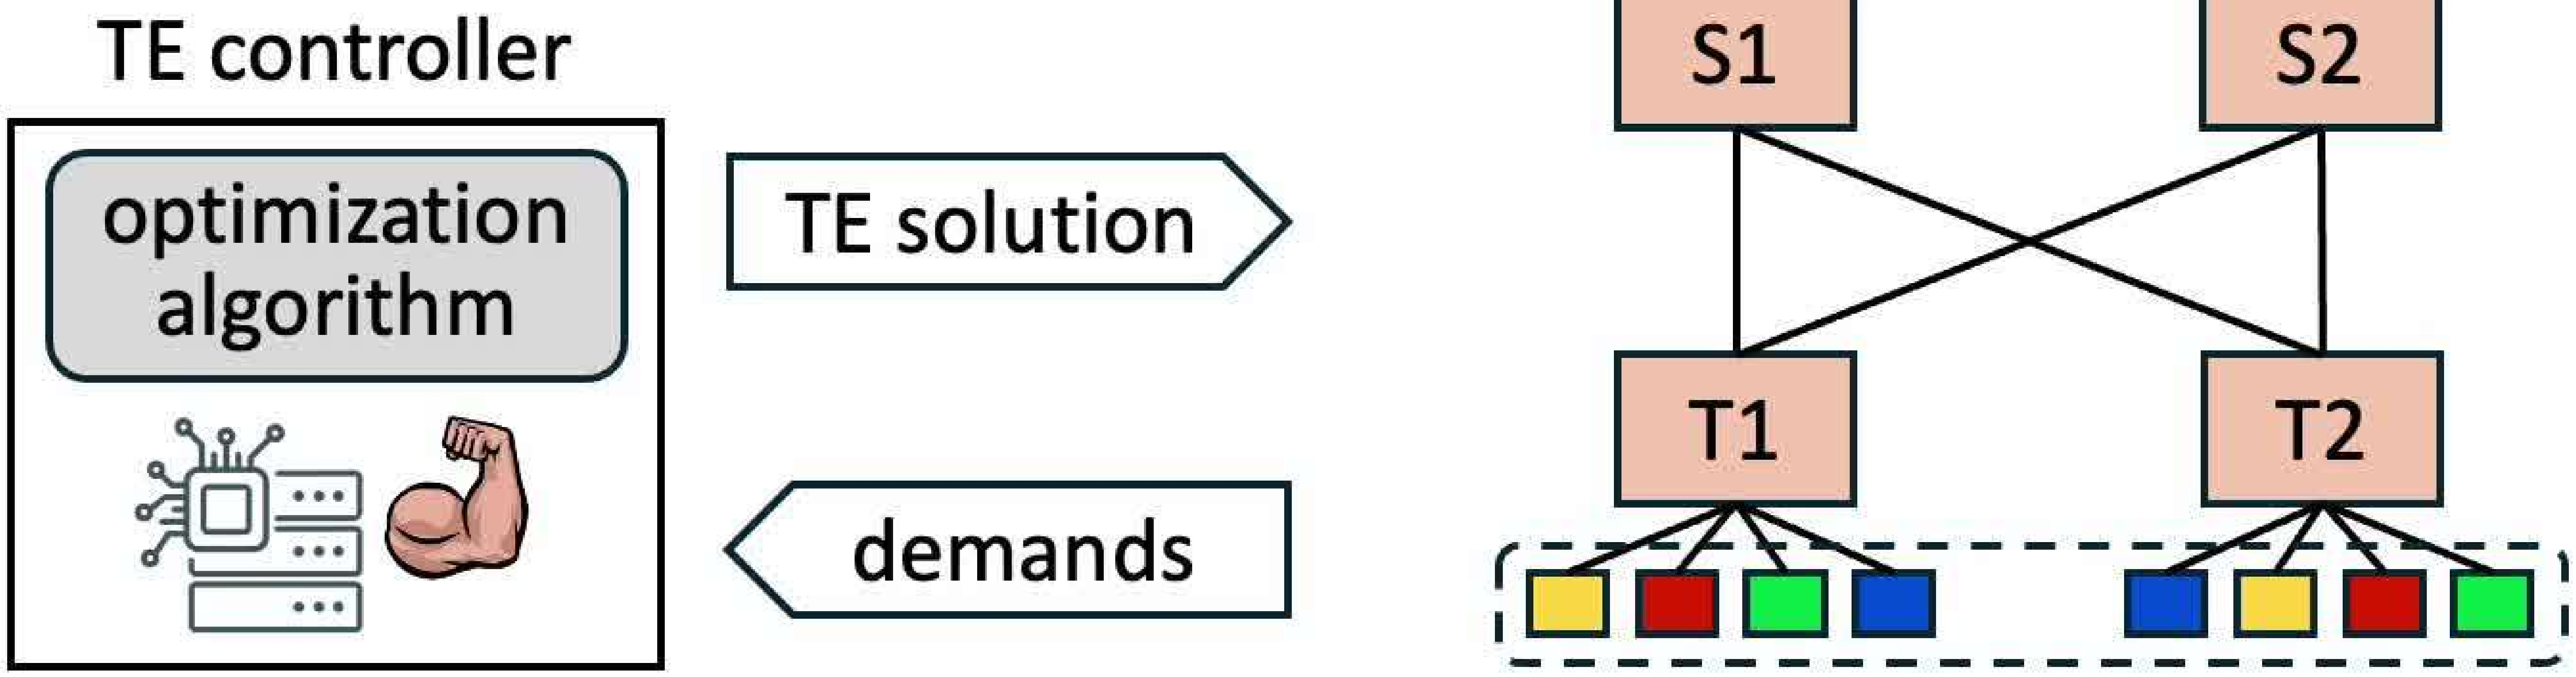

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation

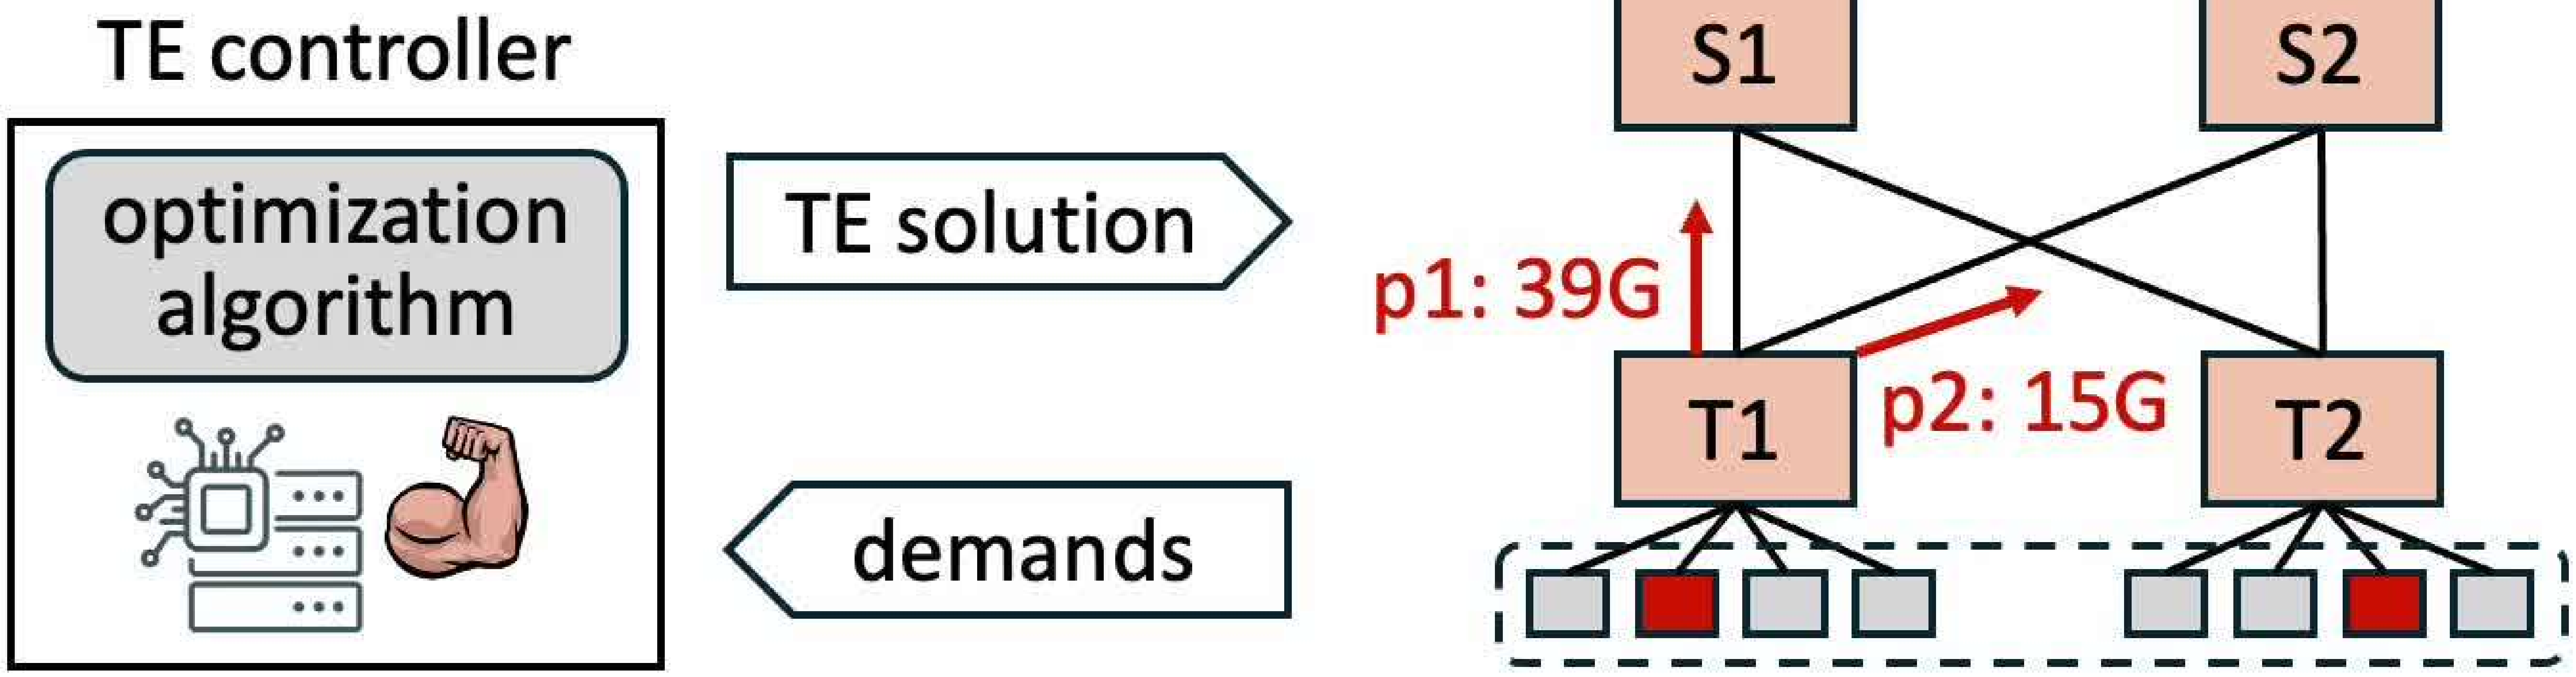

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation
- Repeatedly runs every other minute (e.g., Google's Orion TE controller)

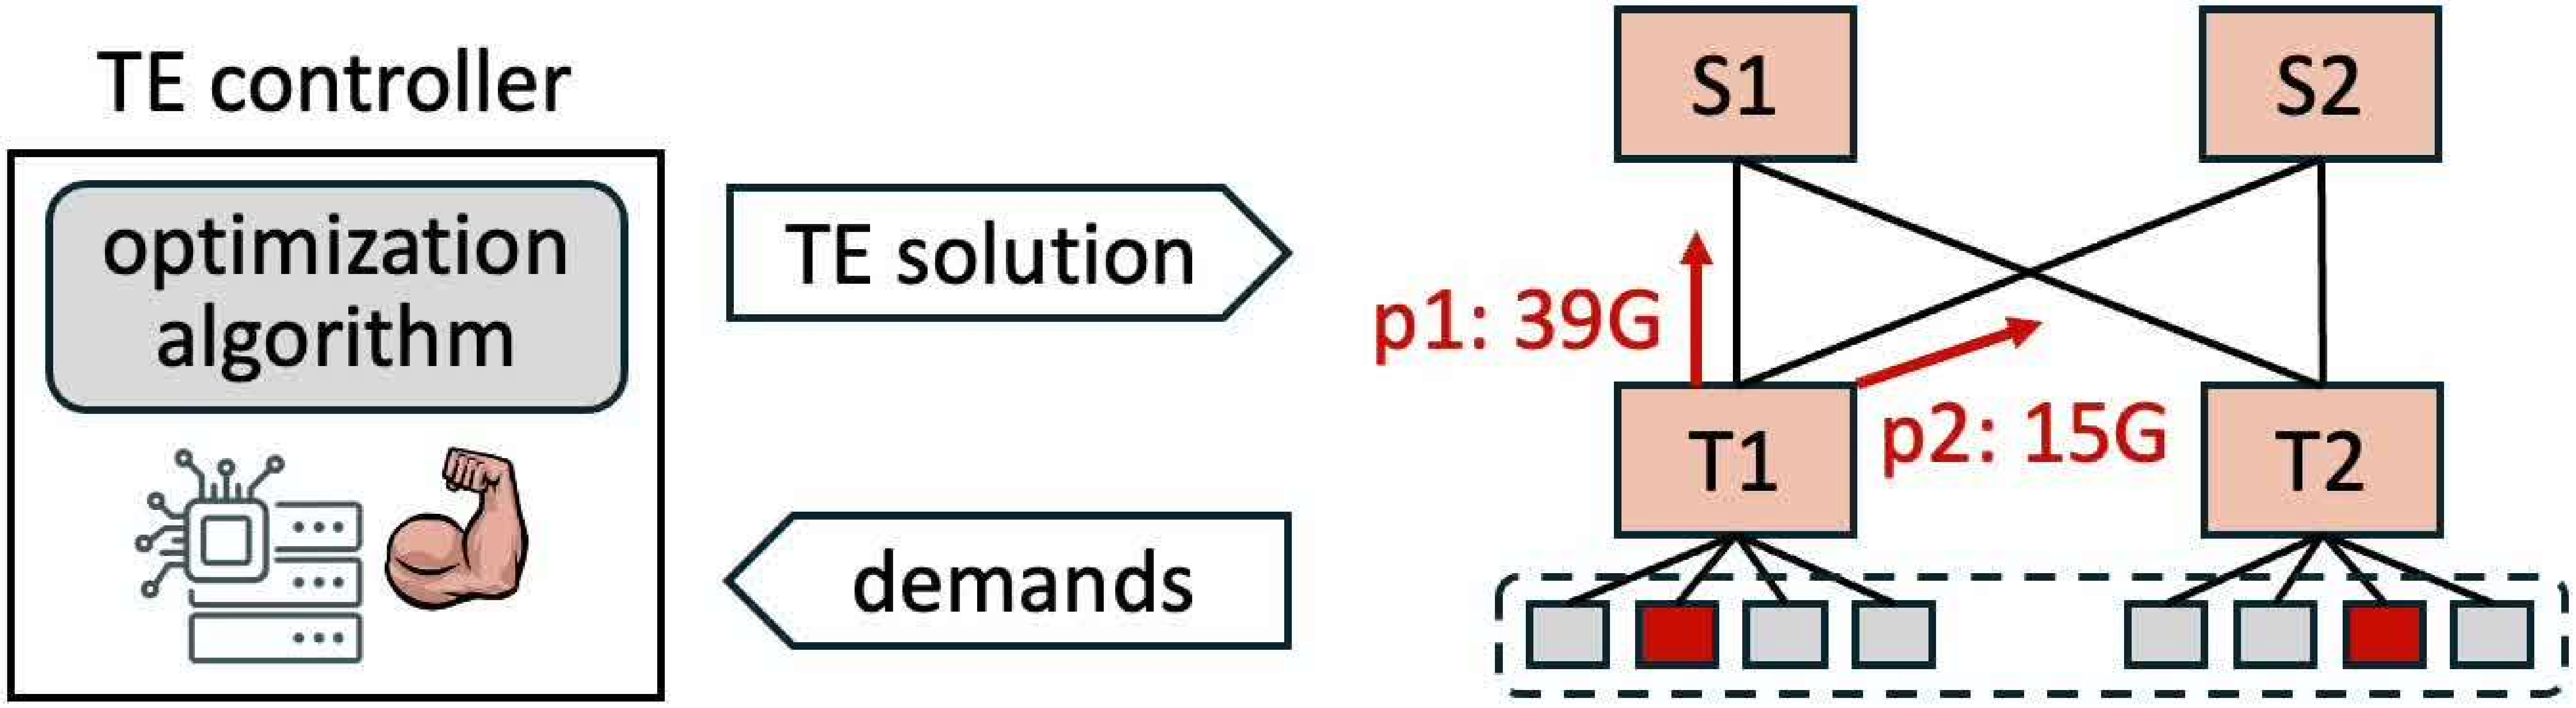

How does a traffic engineering (TE) system work?

- TE control loop: demand collection → solving → implementation
- Repeatedly runs every other minute (e.g., Google's Orion TE controller)

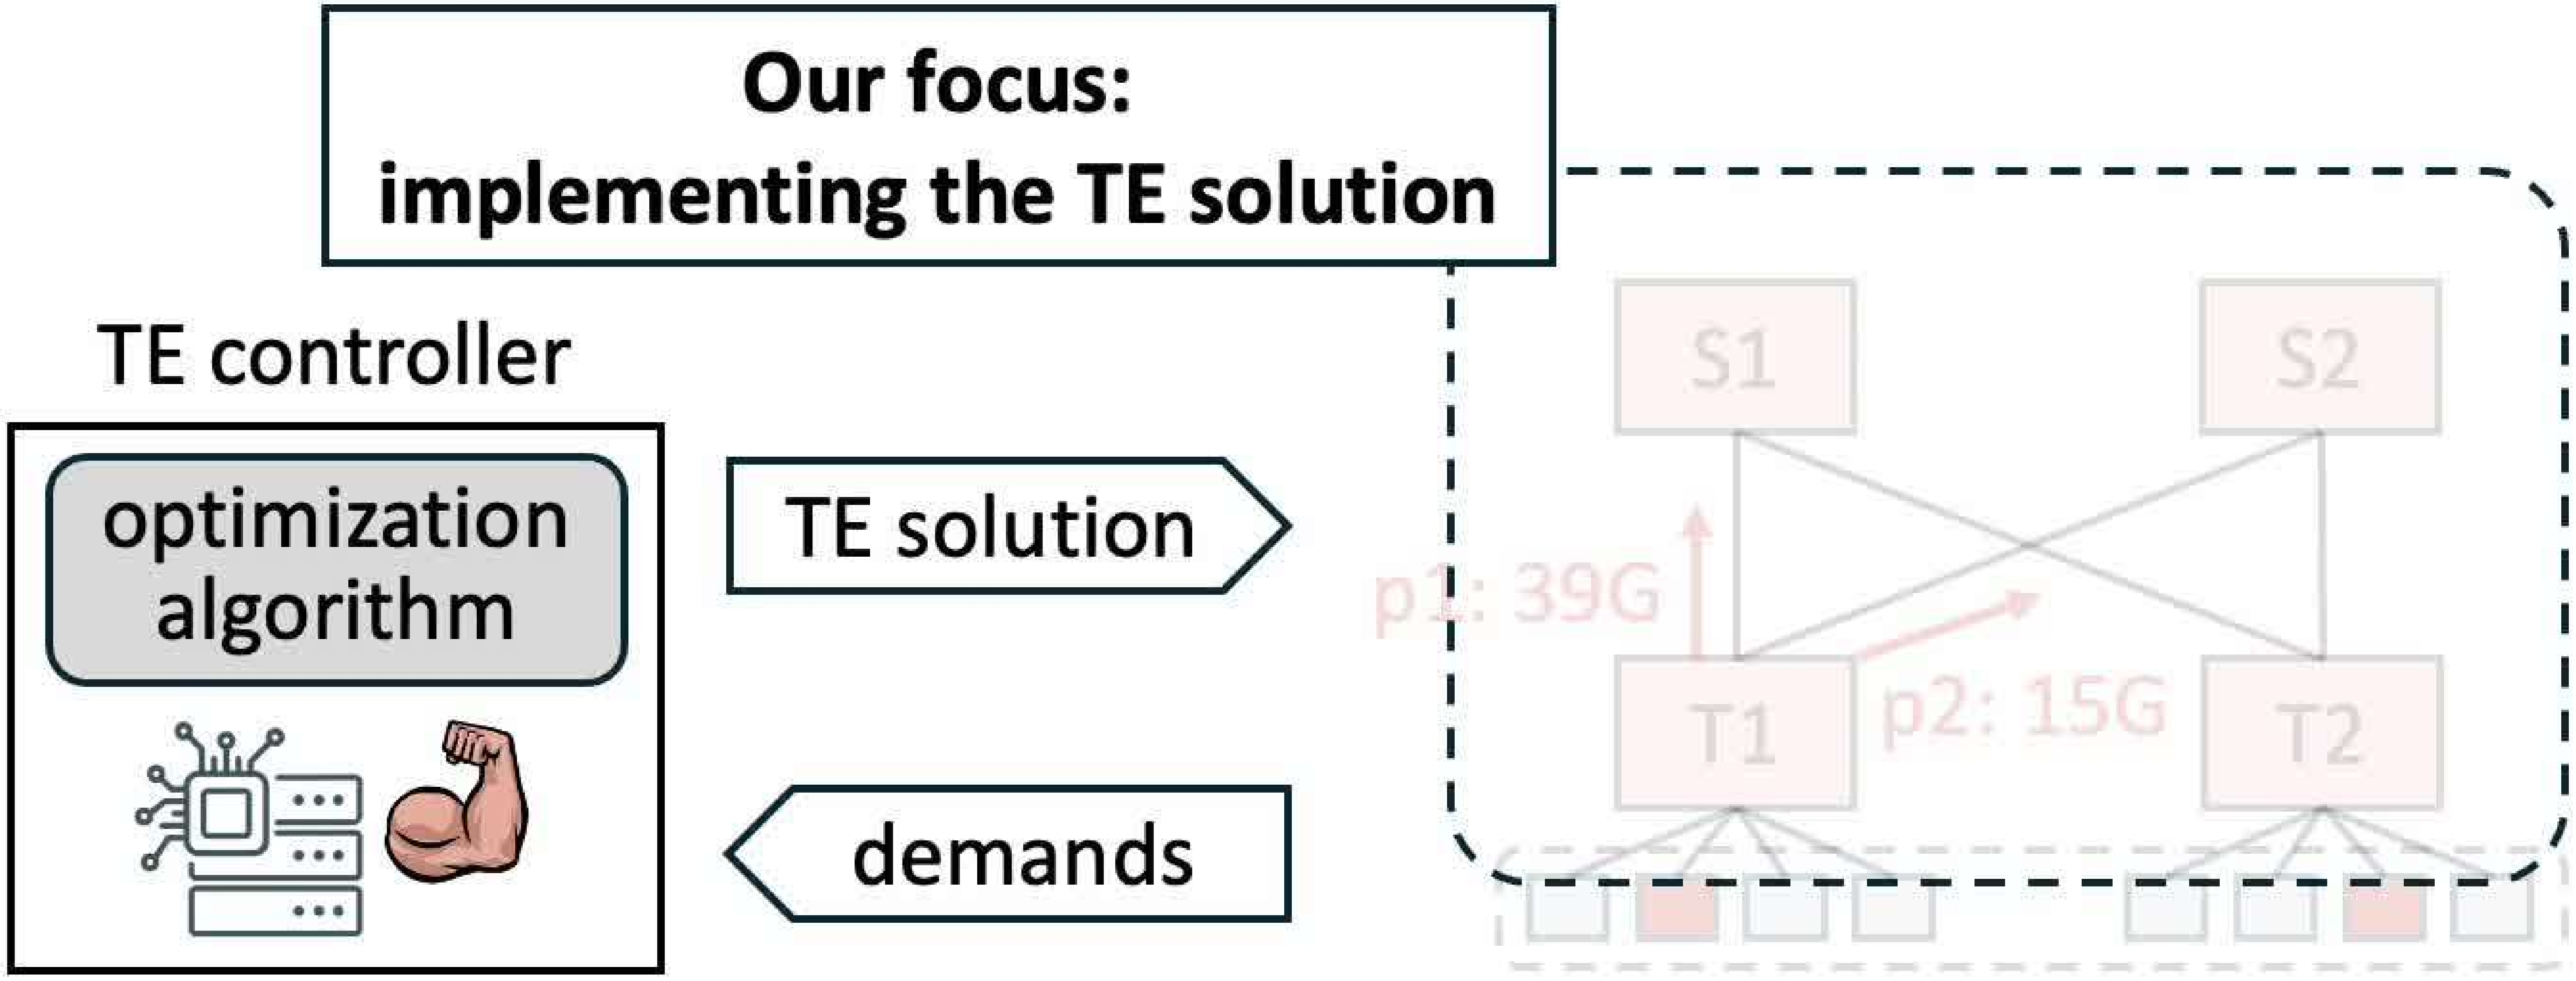

Precision loss is inherent in TE.

Switch T1 internal memory

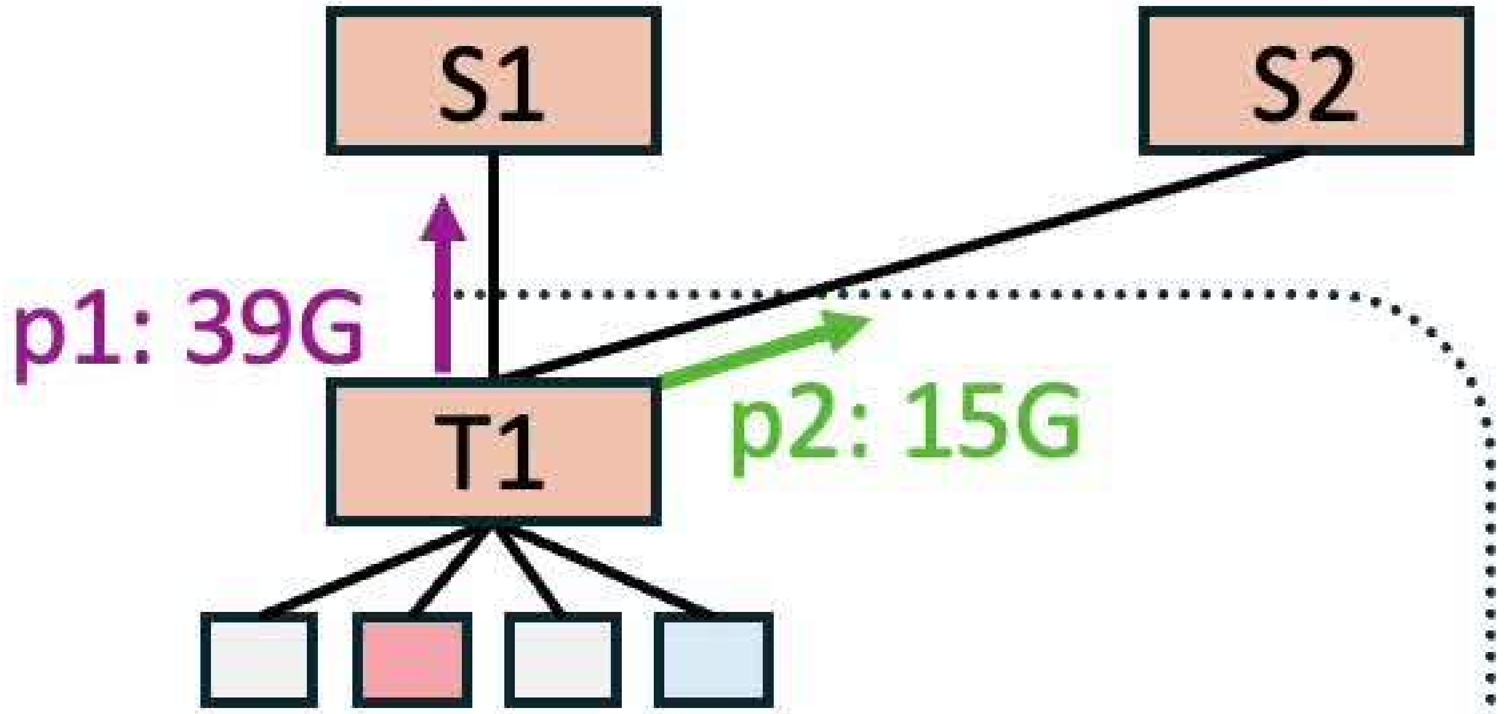

Precision loss is inherent in TE.

Switch T1 internal memory

Destination
<red dst IP>
<blue dst IP>
...

Group	
ID	Egress port
G1	p1
	p1
	...
	p1
	p2
	p2
	...
	p2
G2	...
...	

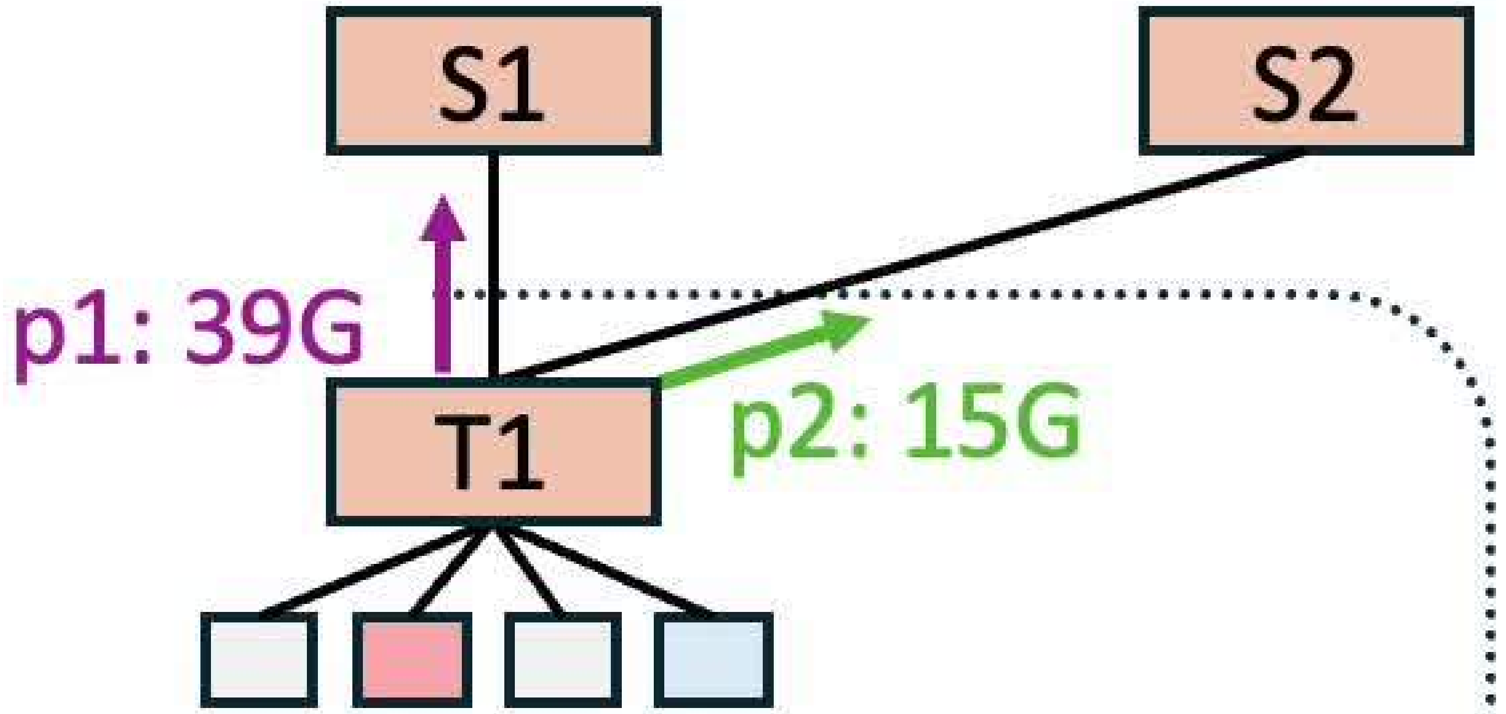

Precision loss is inherent in TE.

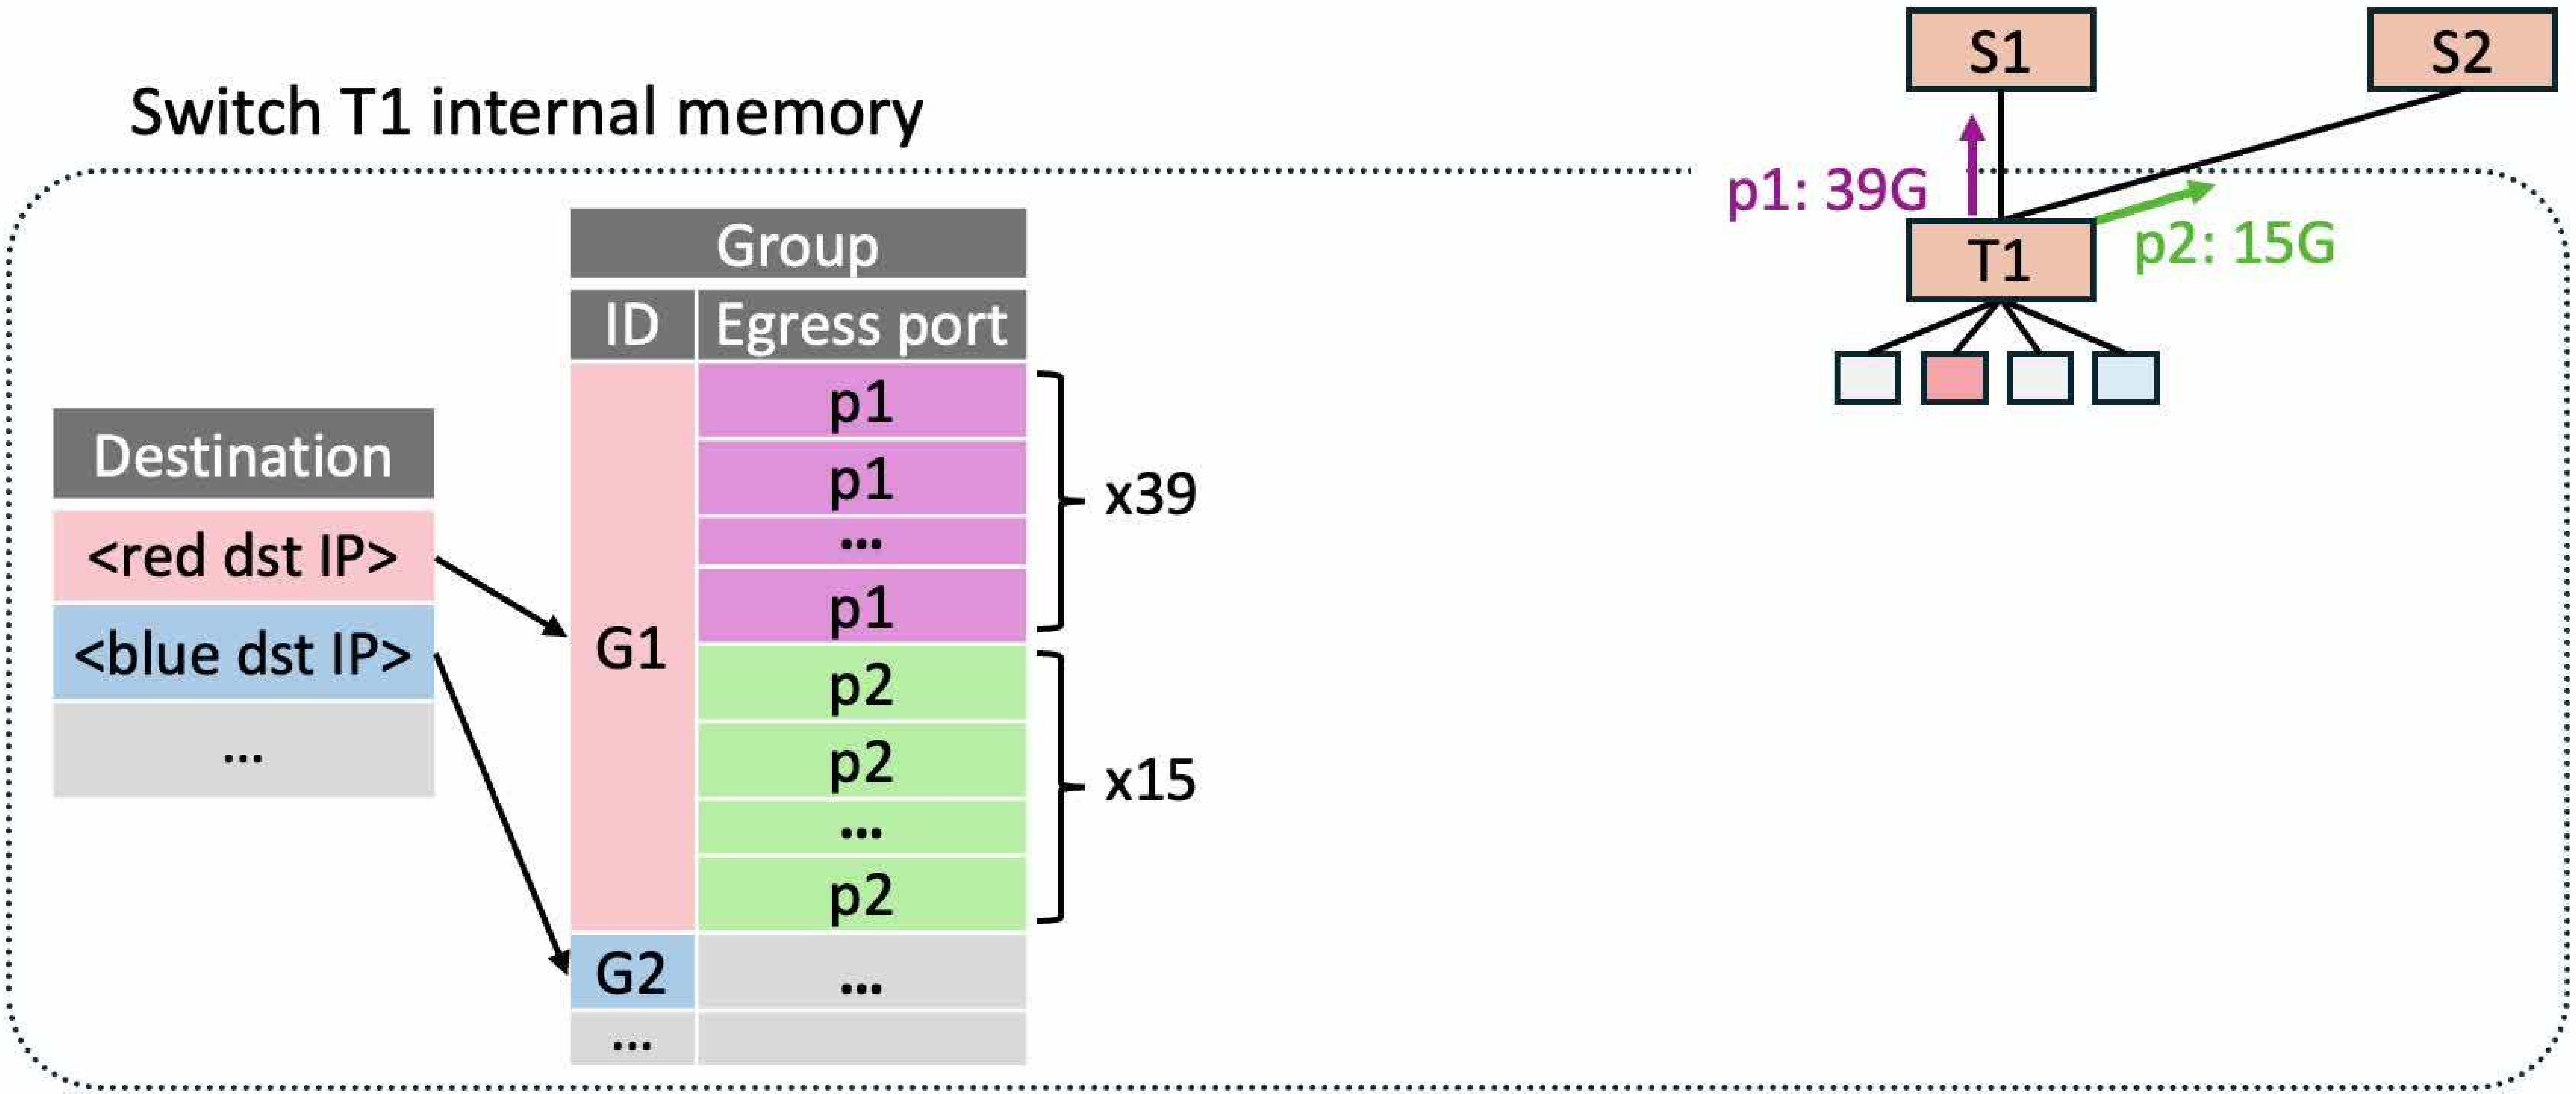

Precision loss is inherent in TE.

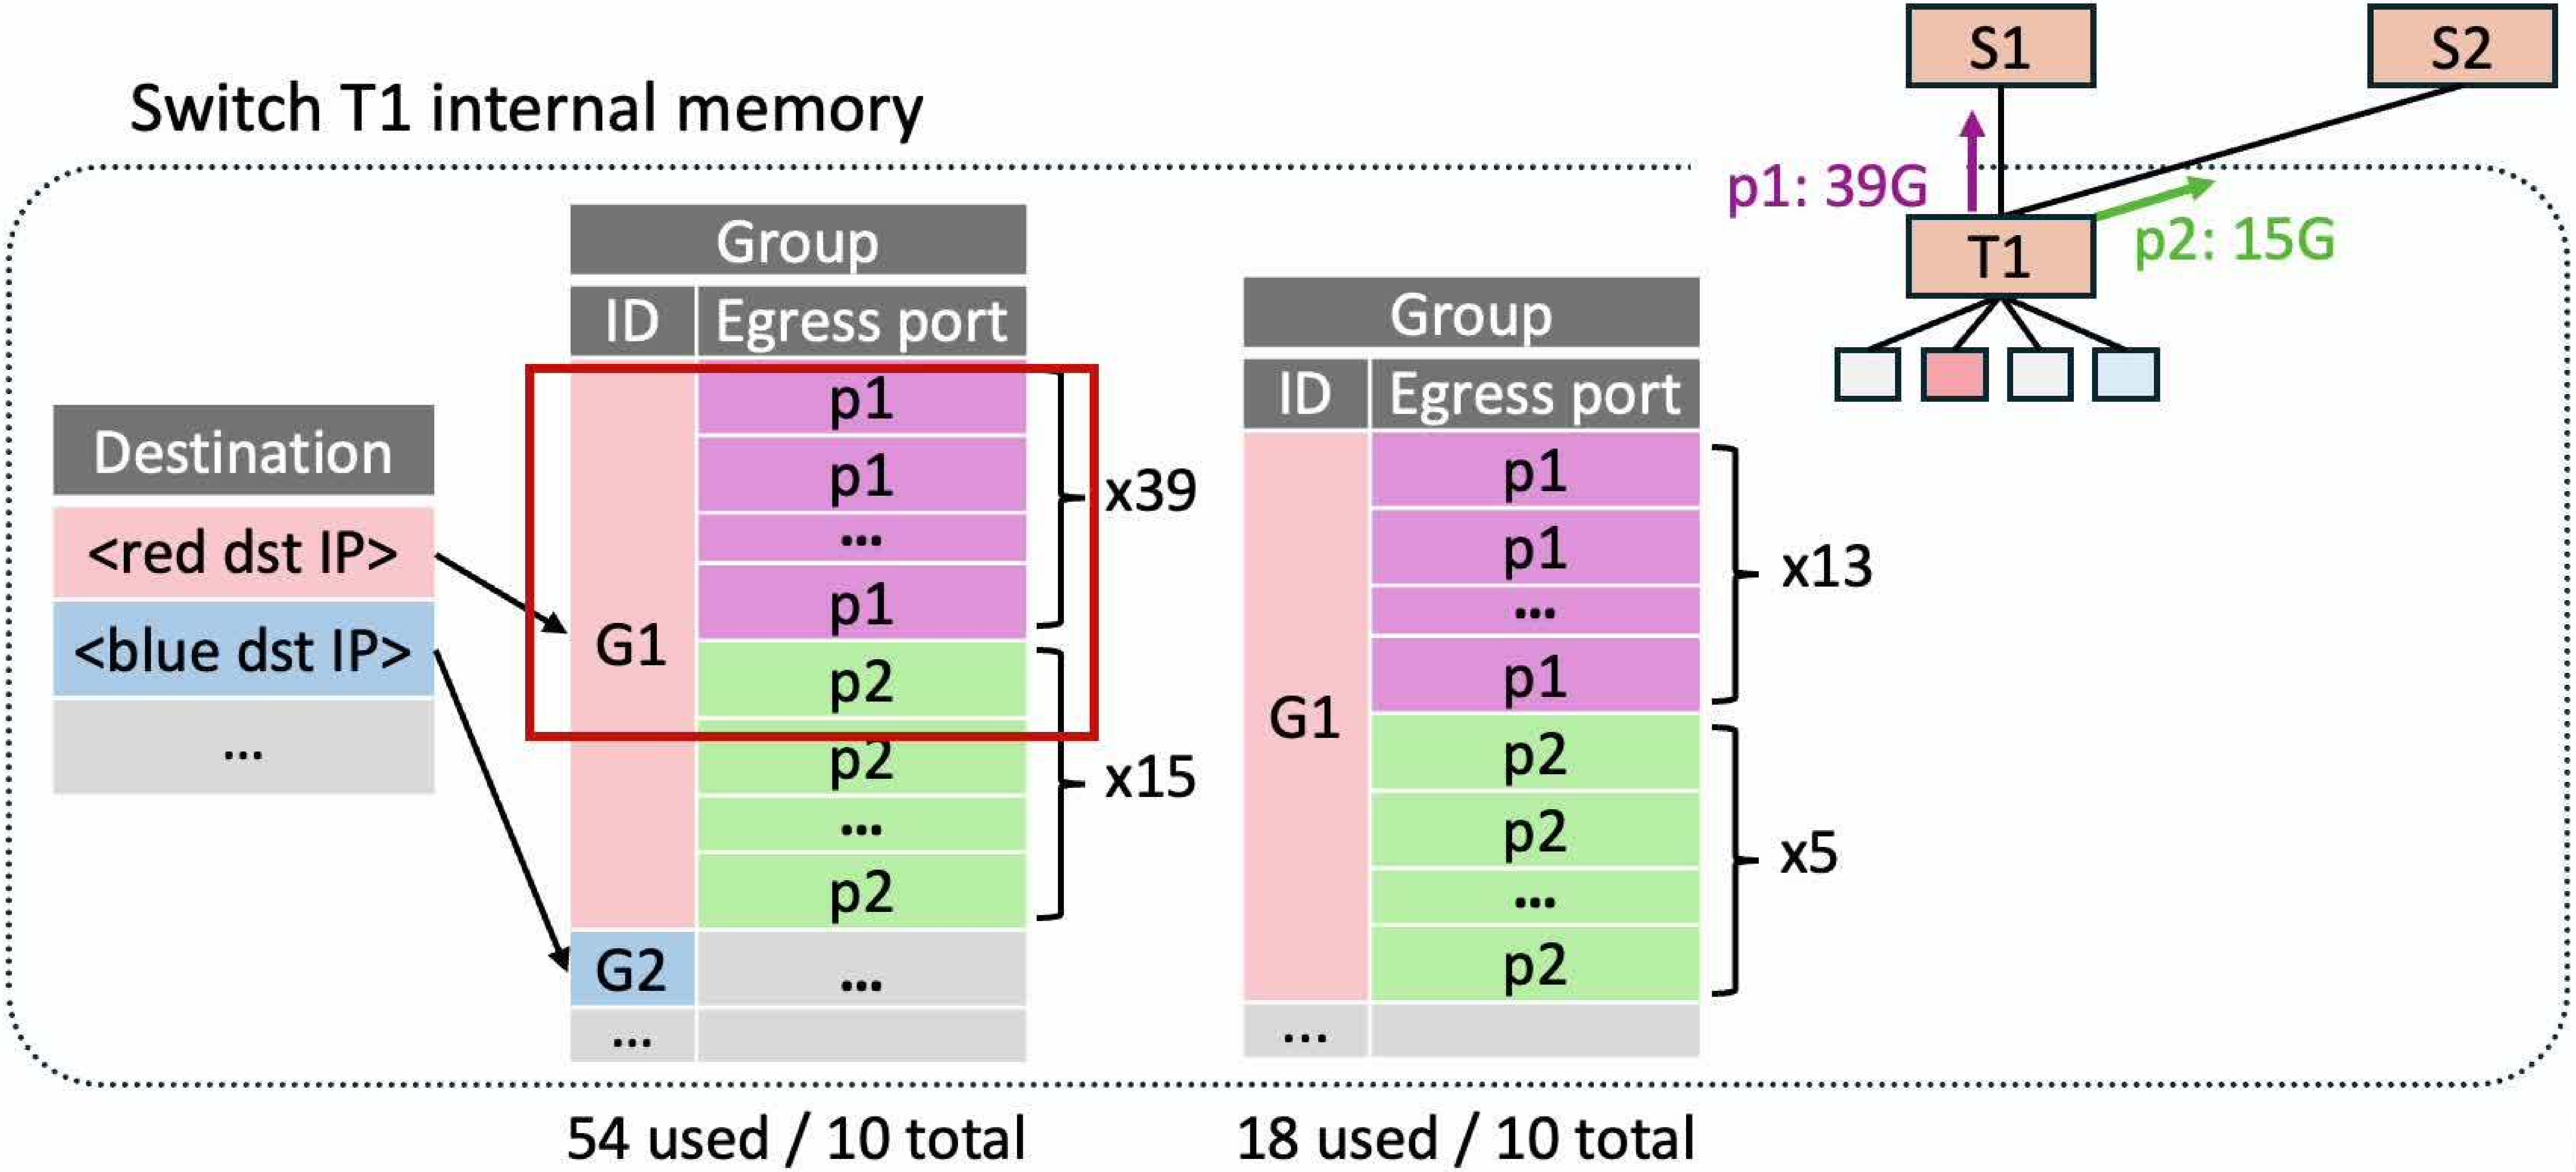

Precision loss is inherent in TE.

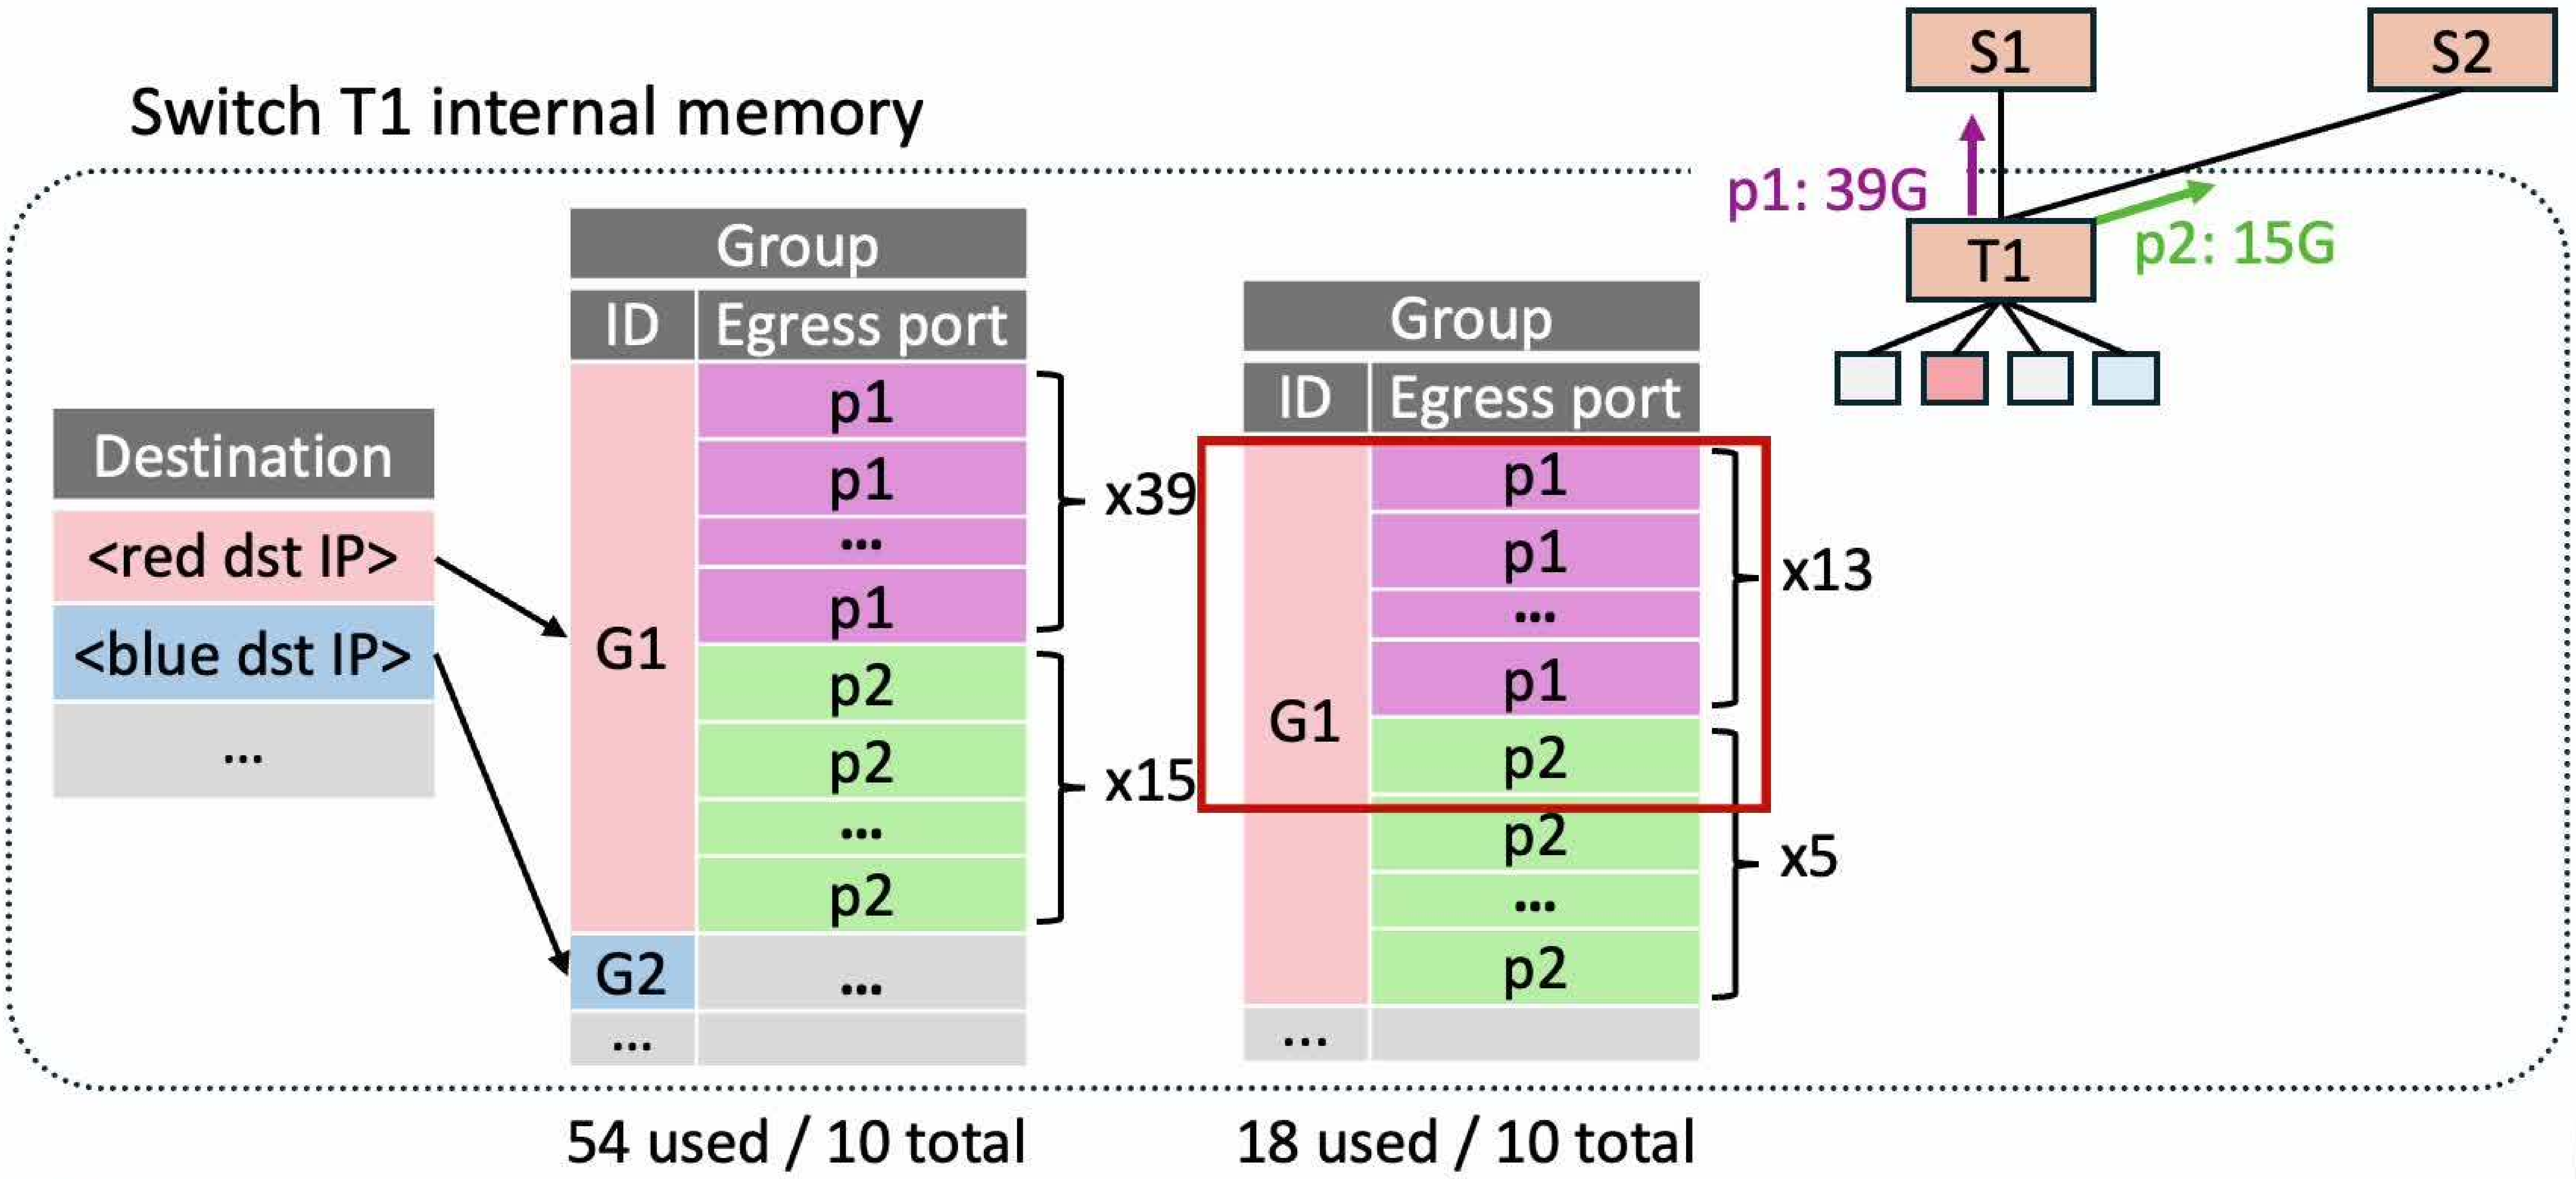

Precision loss is inherent in TE.

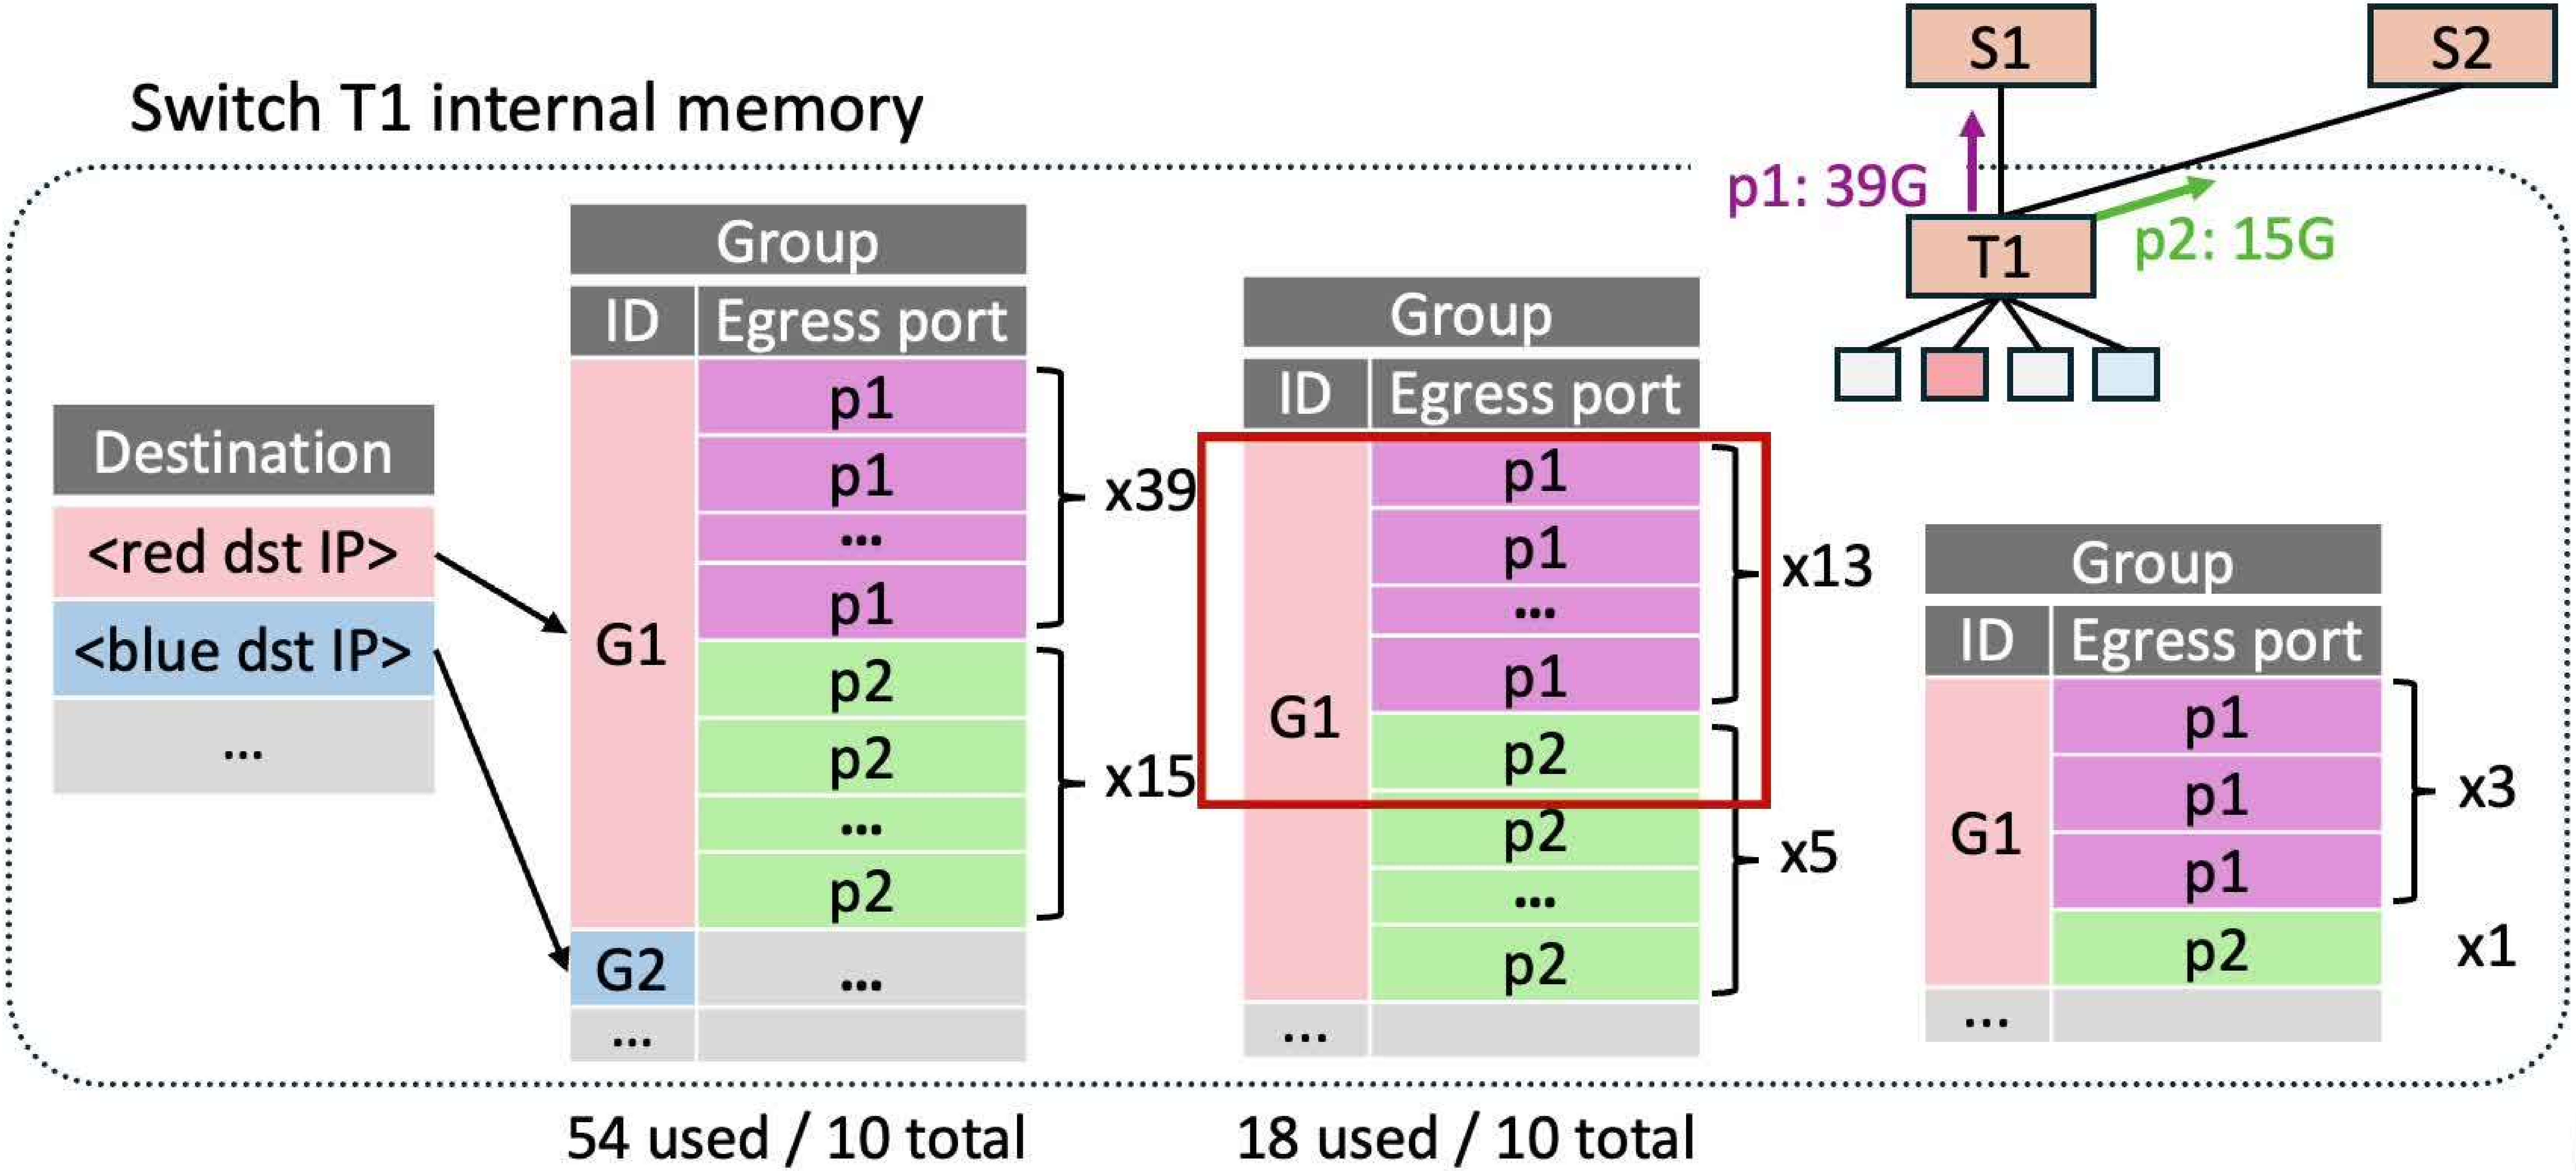

Precision loss is inherent in TE.

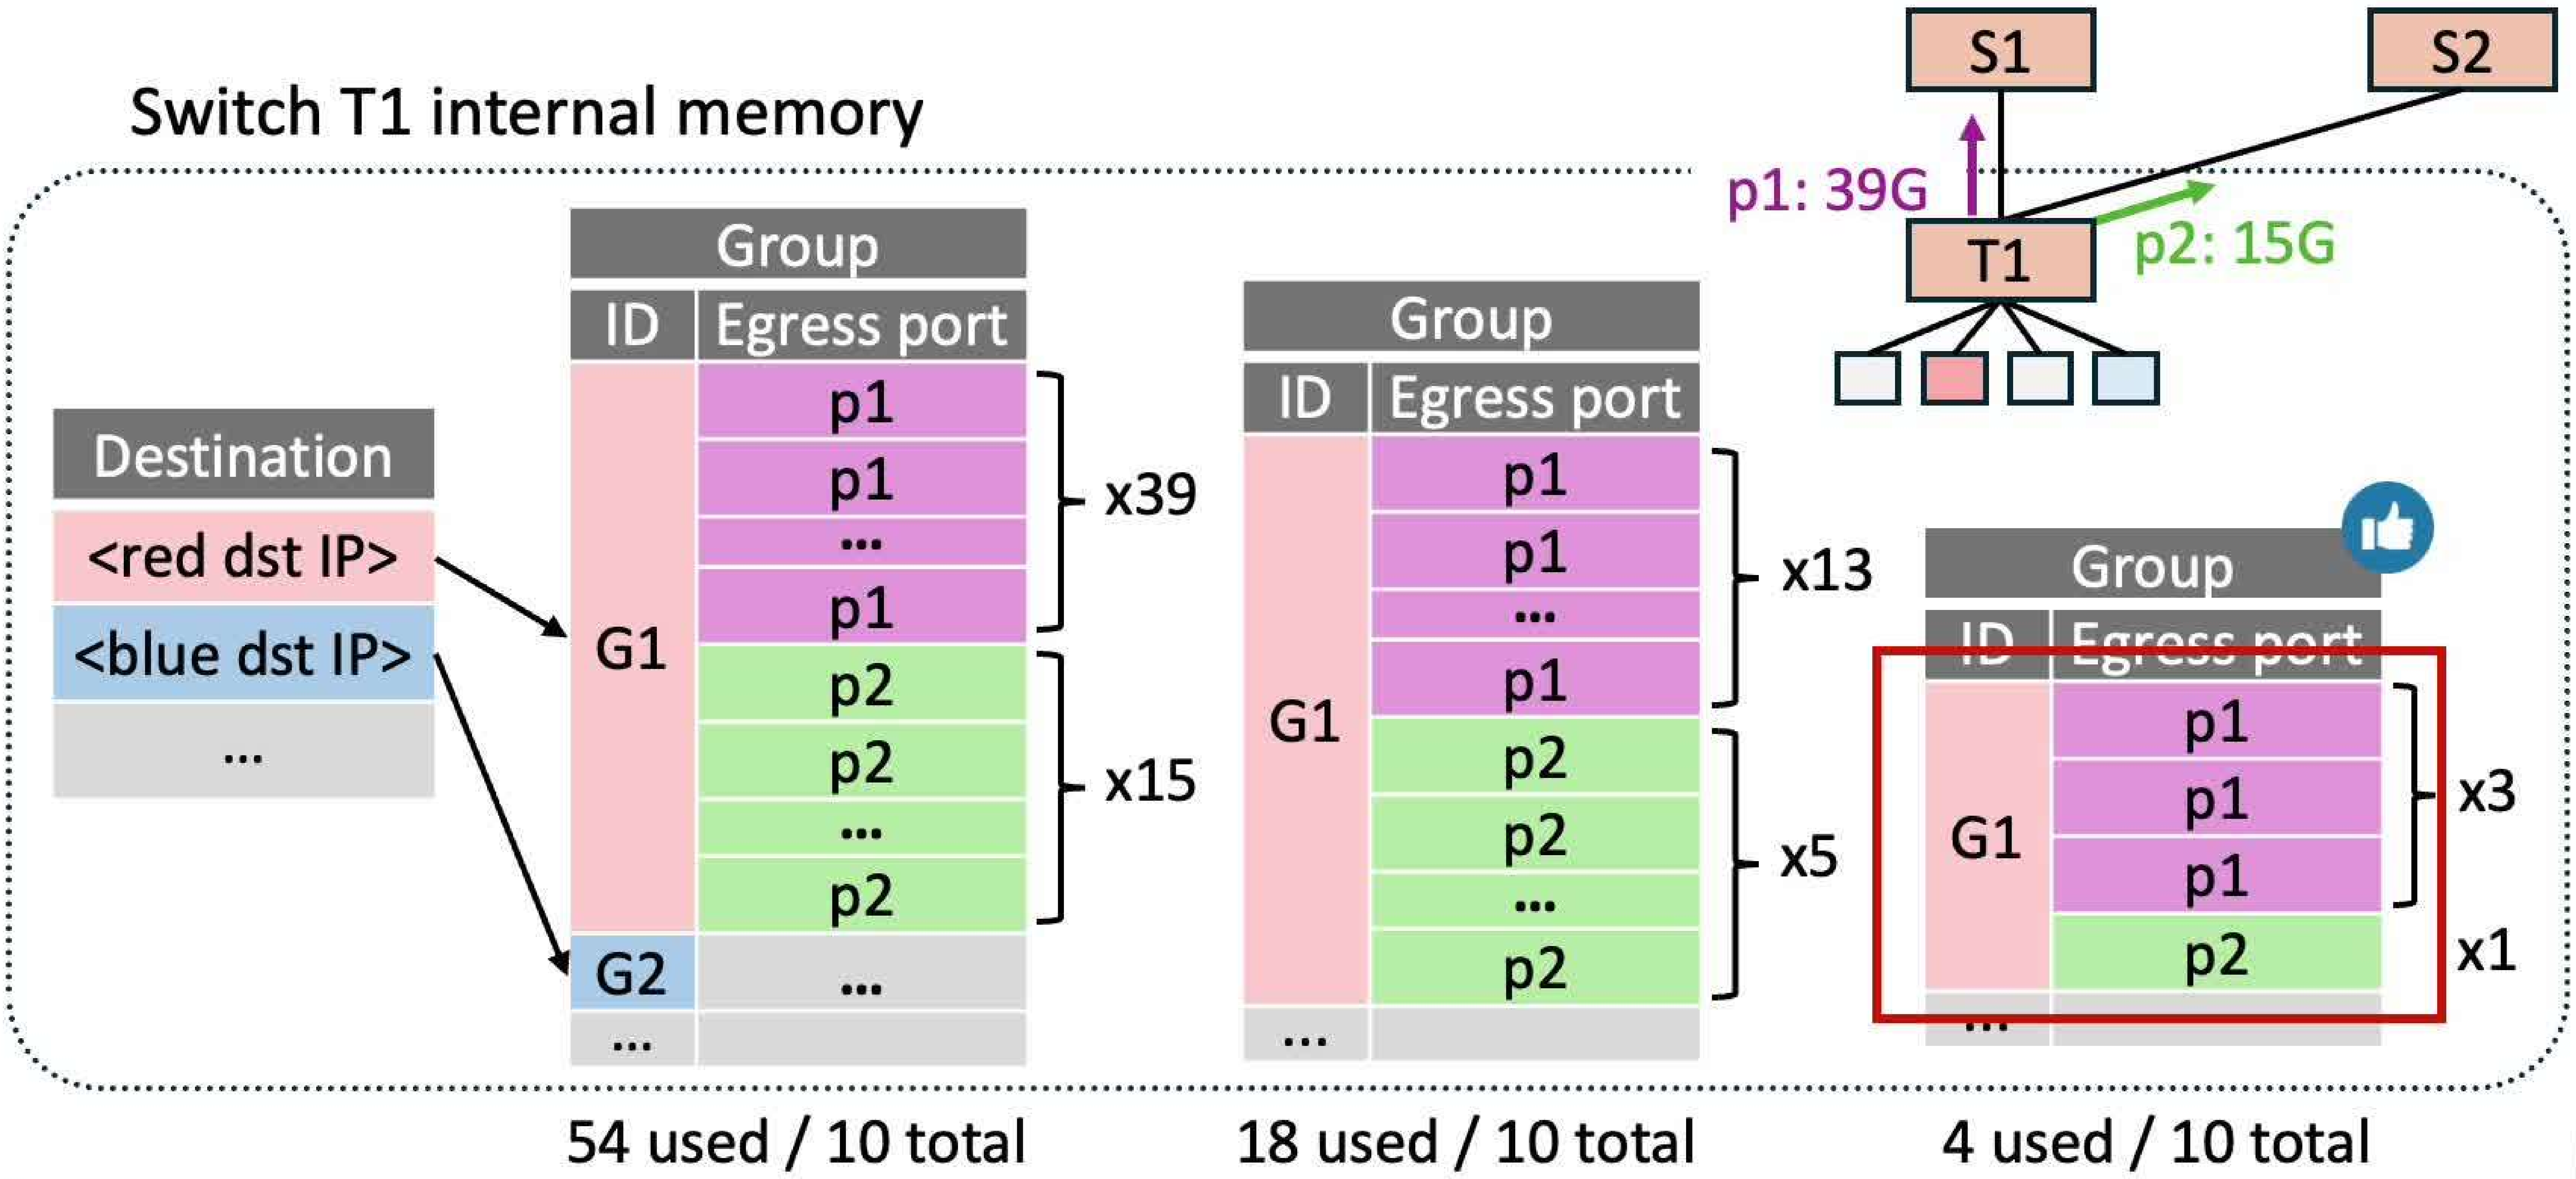

Carnegie Mellon University

Precision loss is inherent in TE.

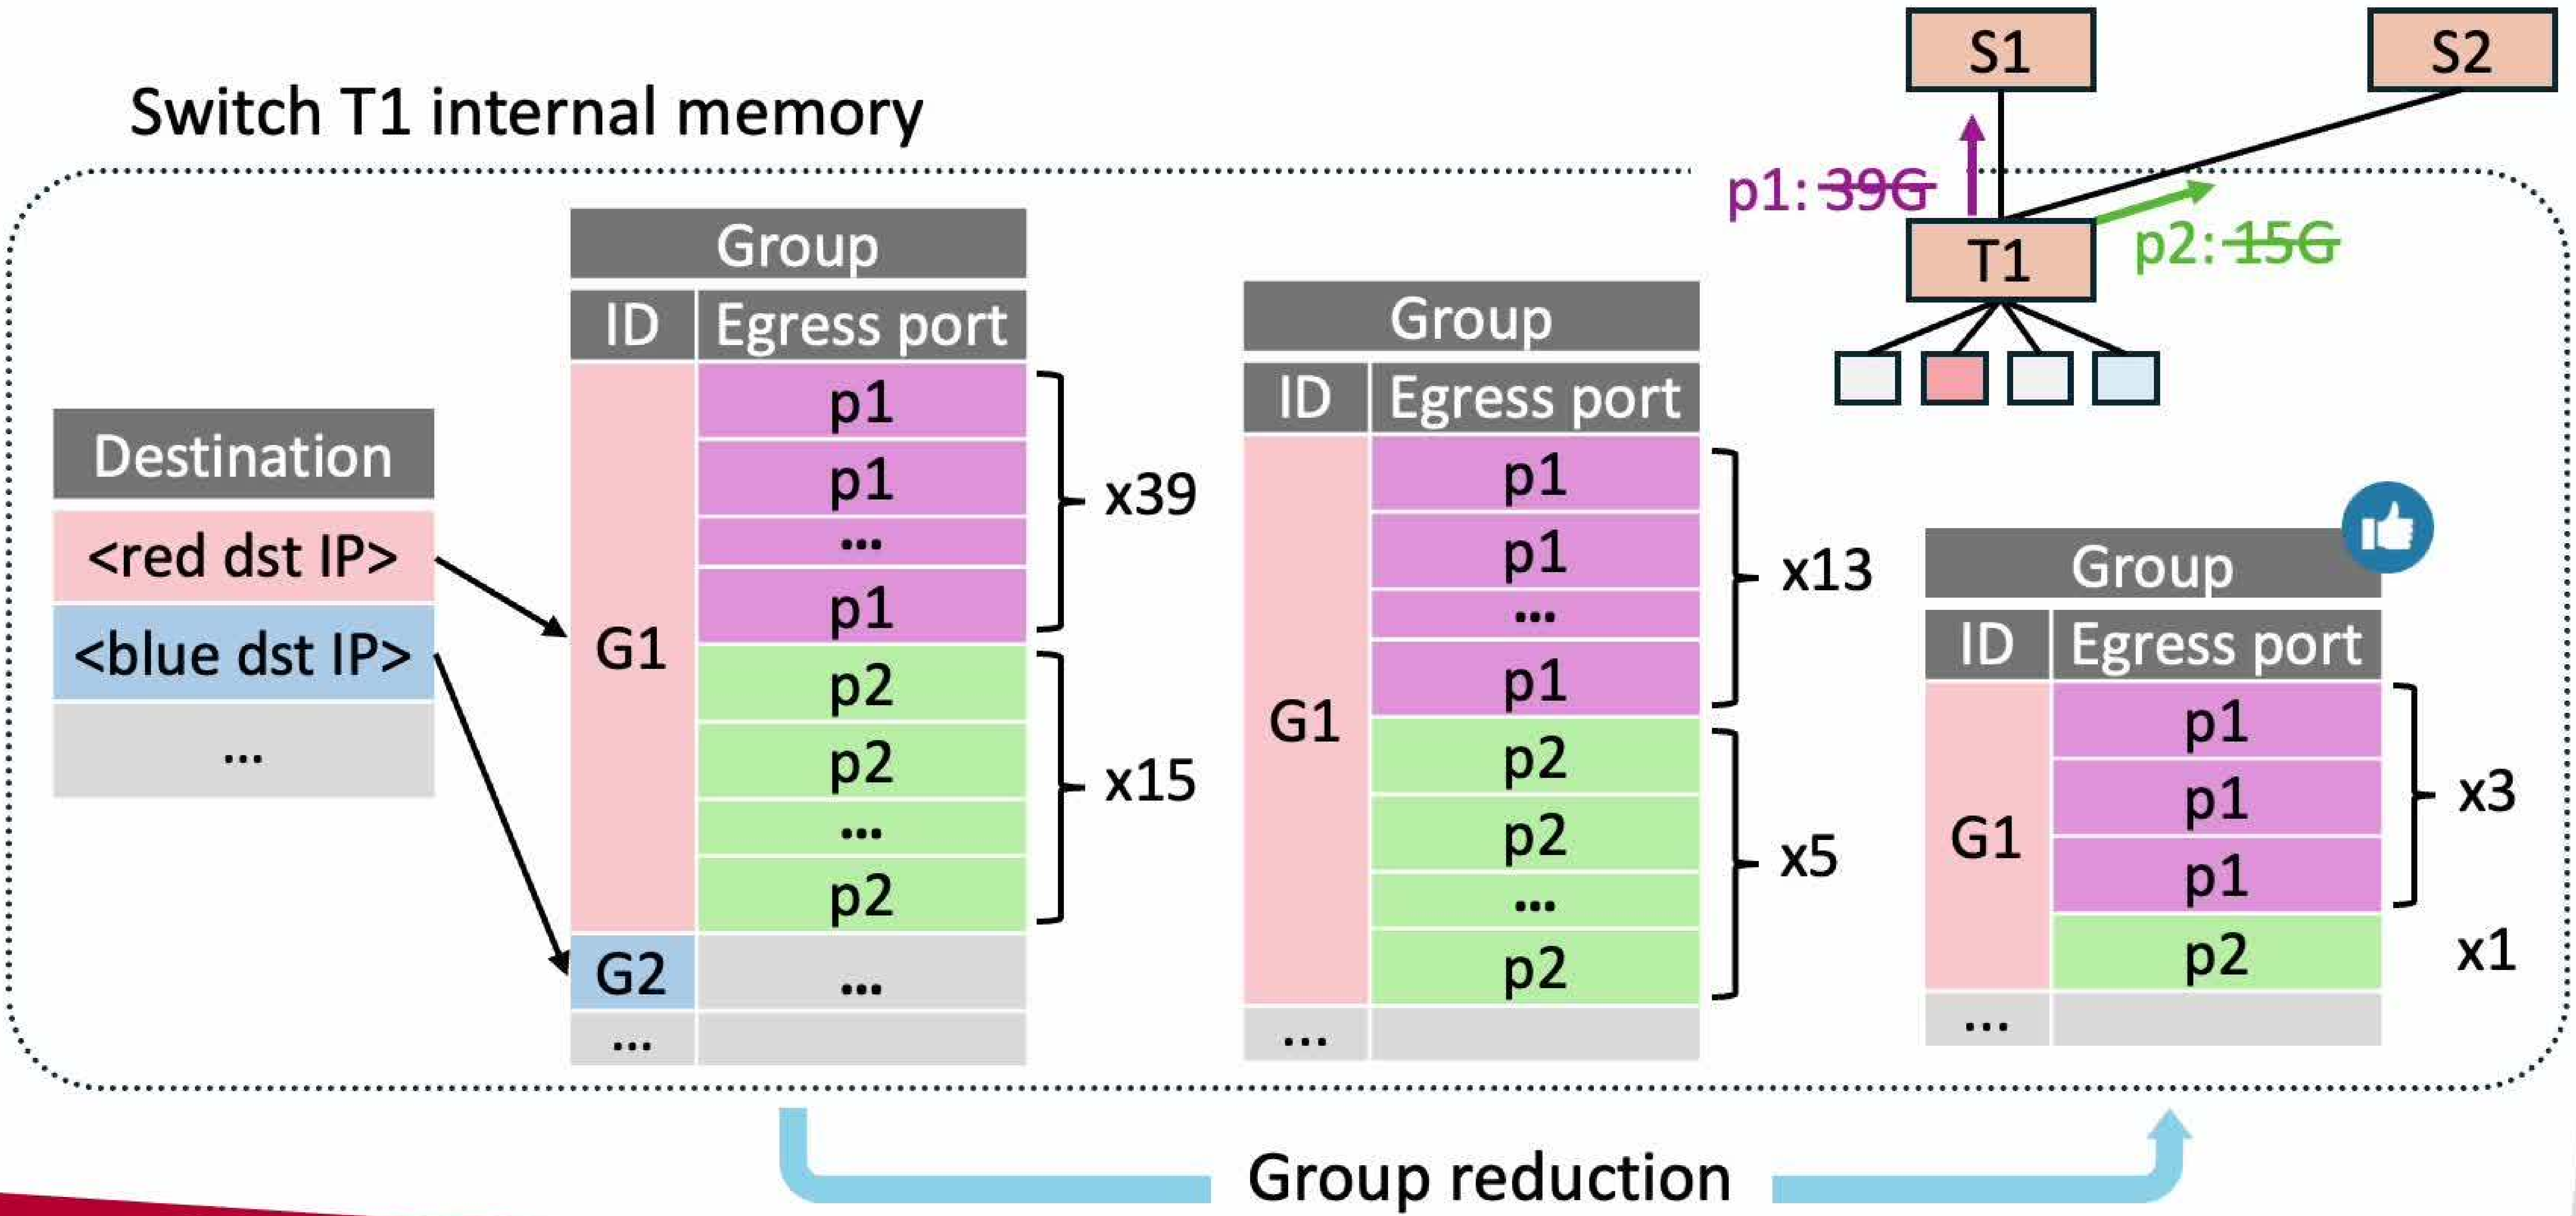

Precision loss is inherent in TE.

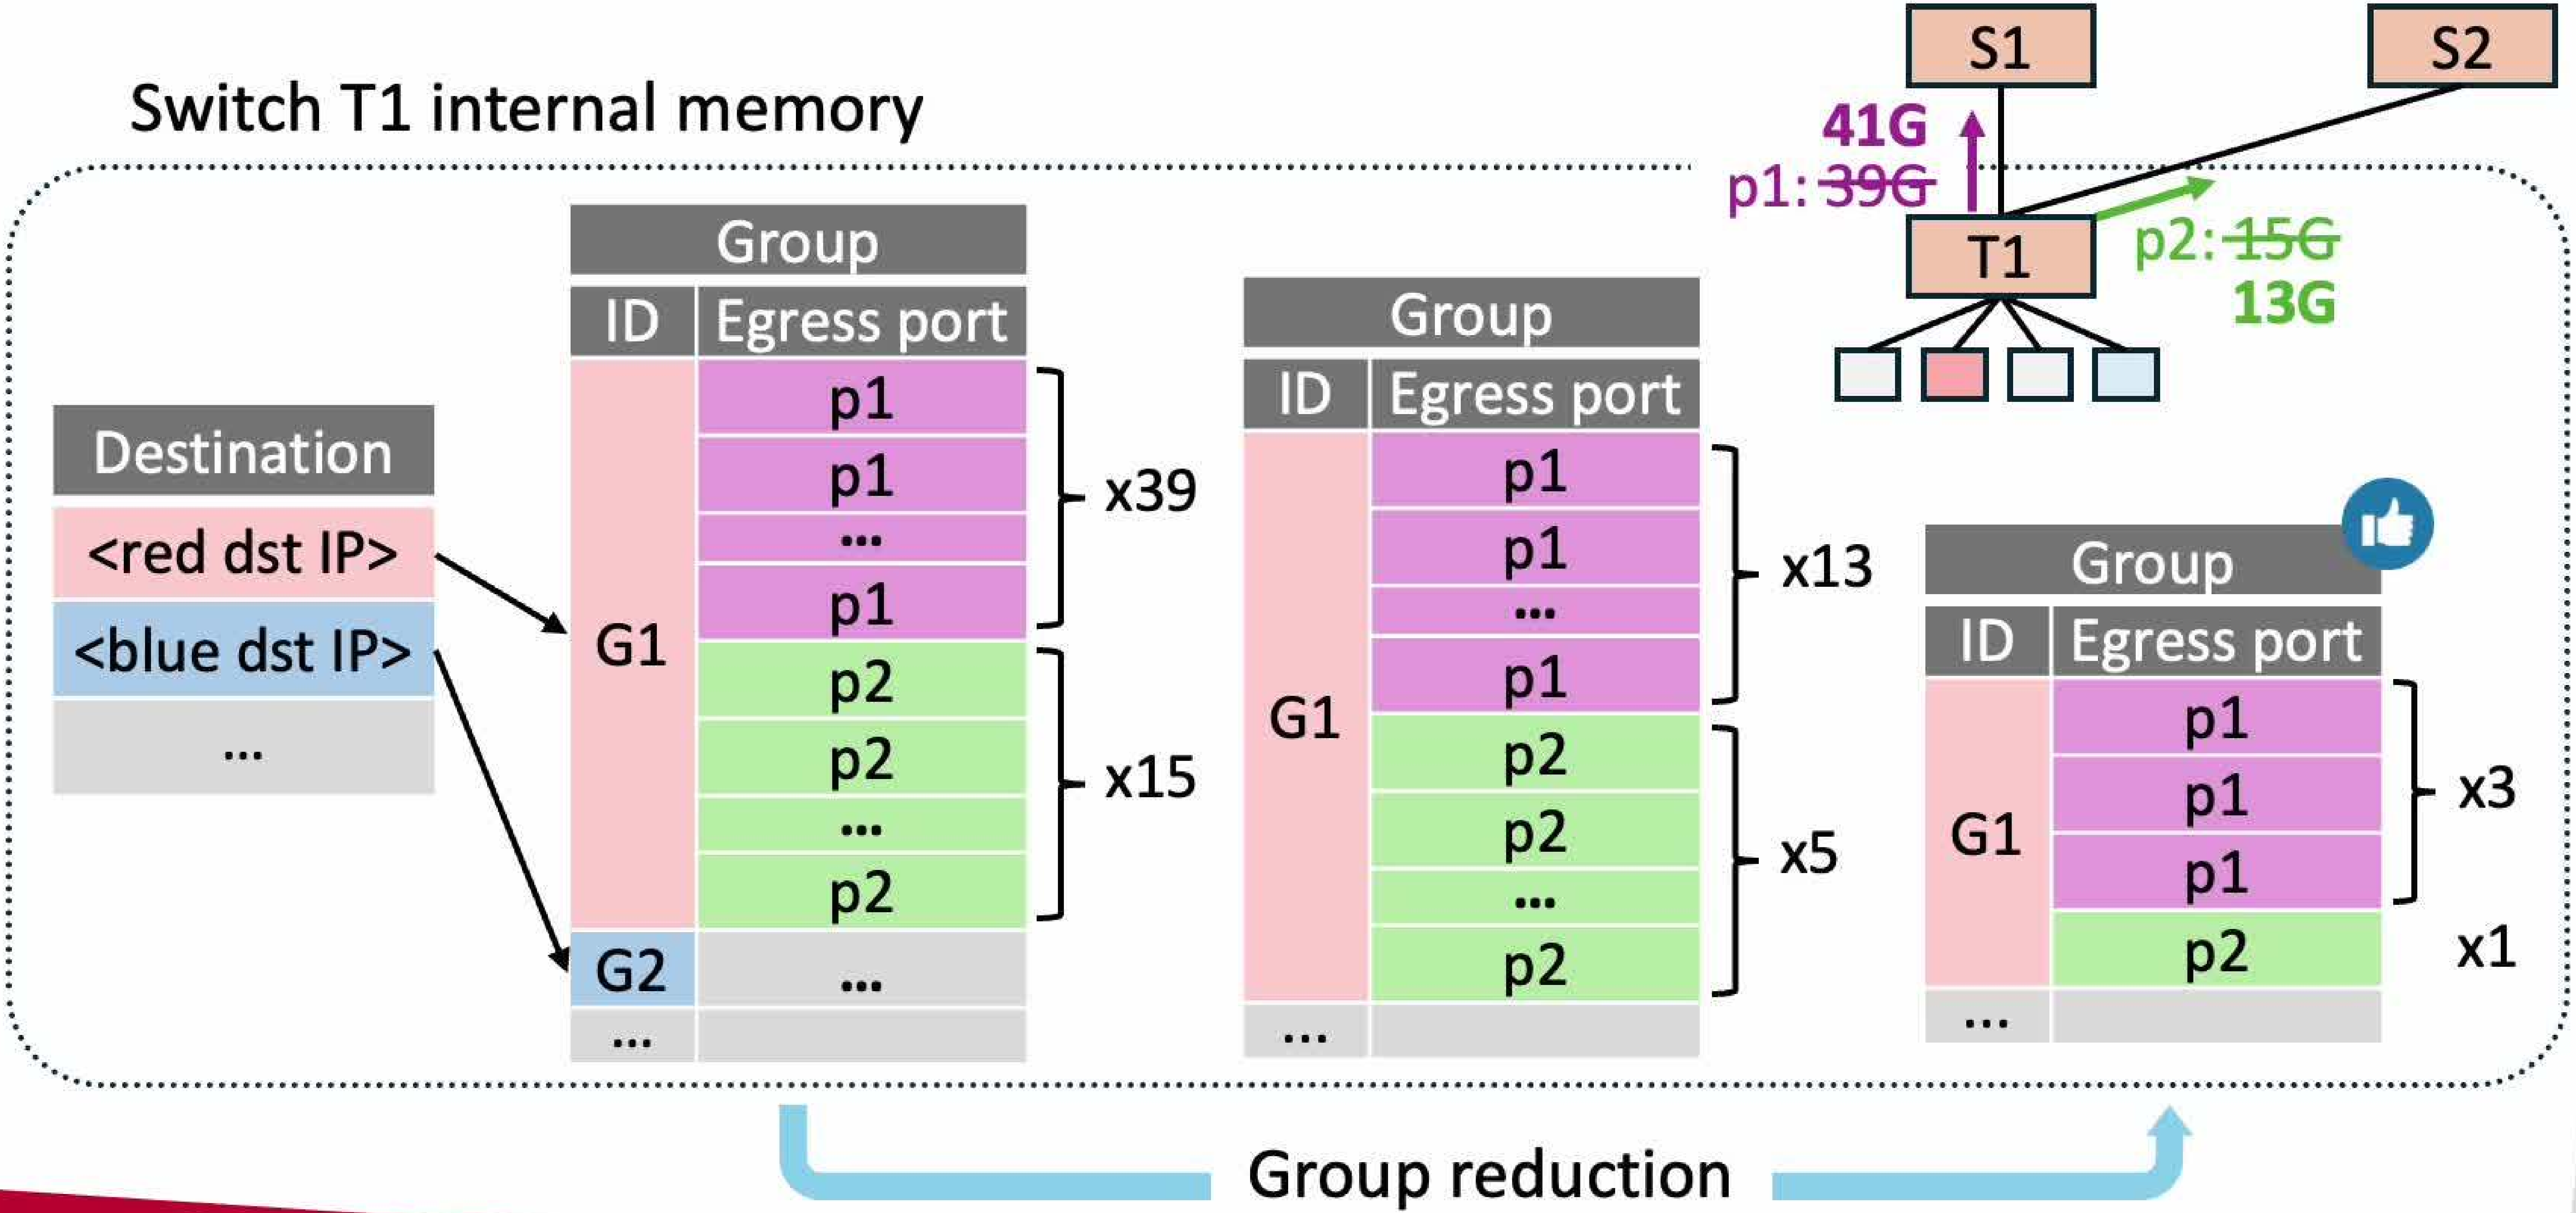

Precision loss is inherent in TE.

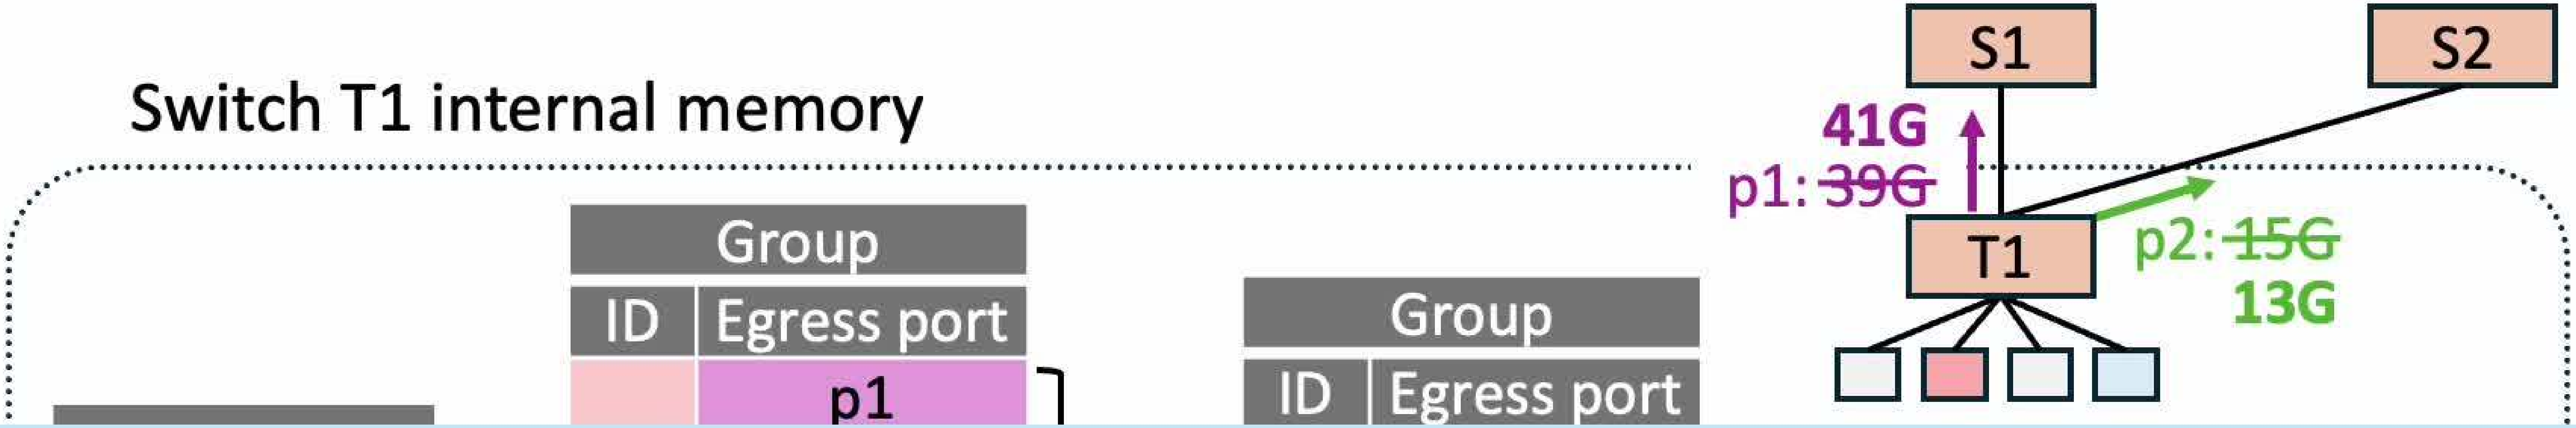

- **Precision loss** is the diff between TE solution and implementation.
- It's caused by limited hardware resources.

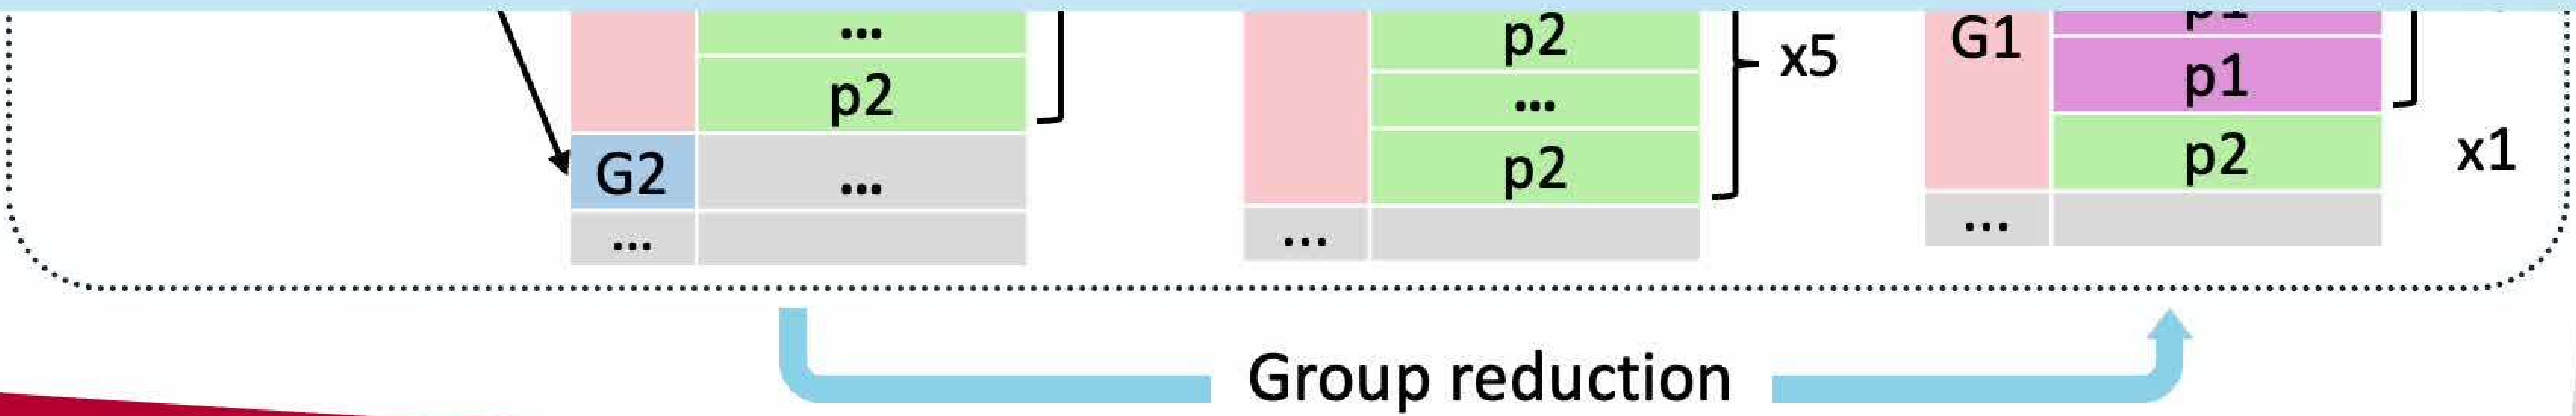

Precision loss accumulates in hierarchical networks.

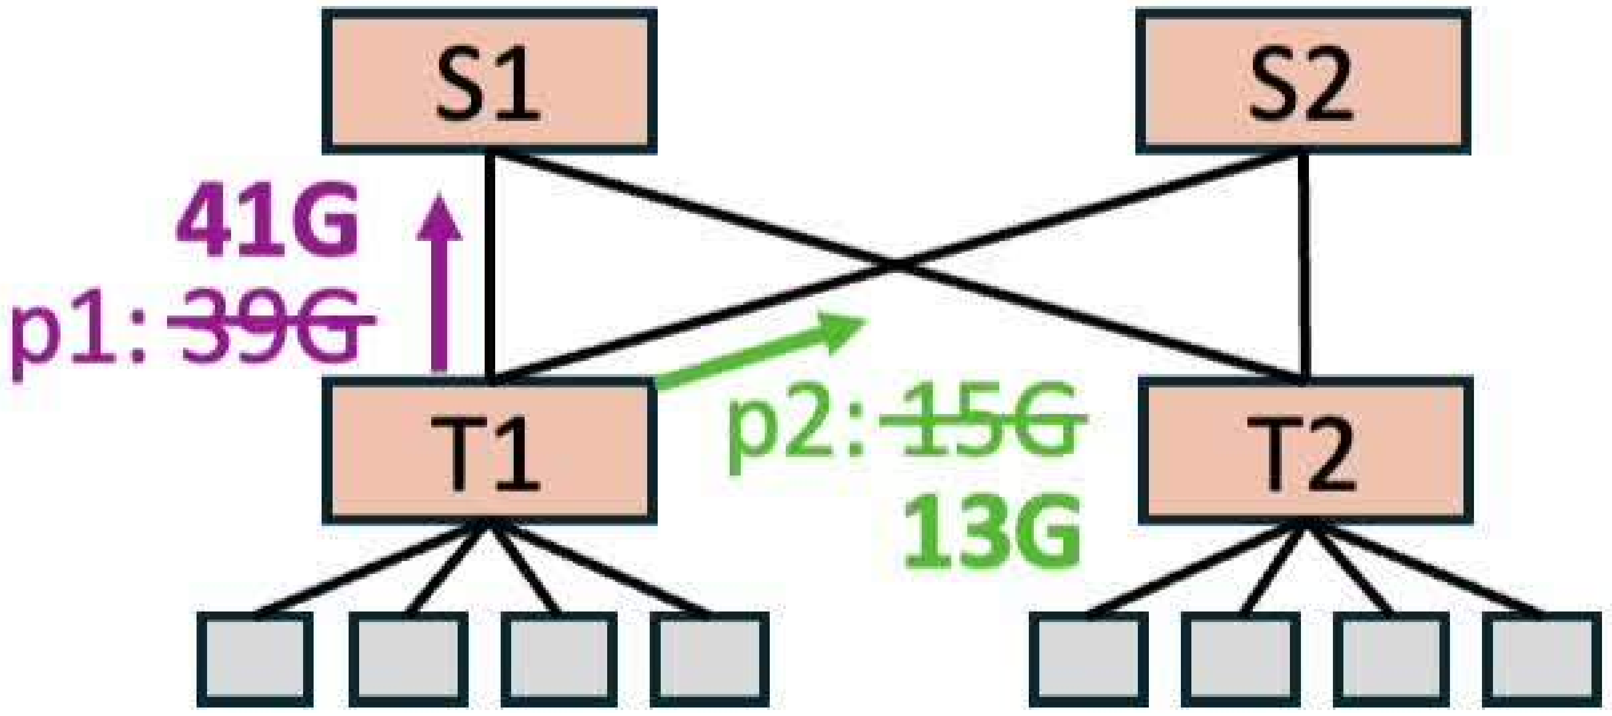

Carnegie Mellon University

Precision loss accumulates in hierarchical networks.

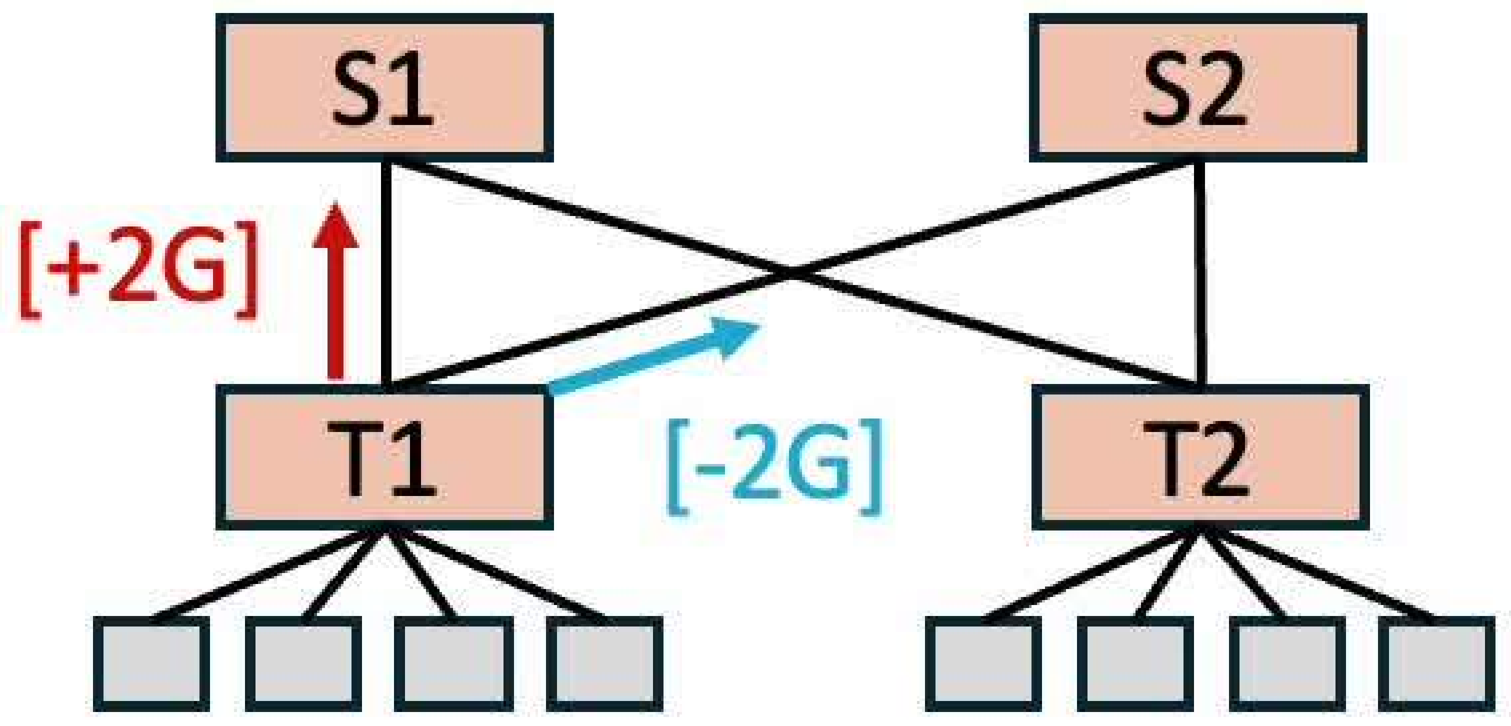

Precision loss accumulates in hierarchical networks.

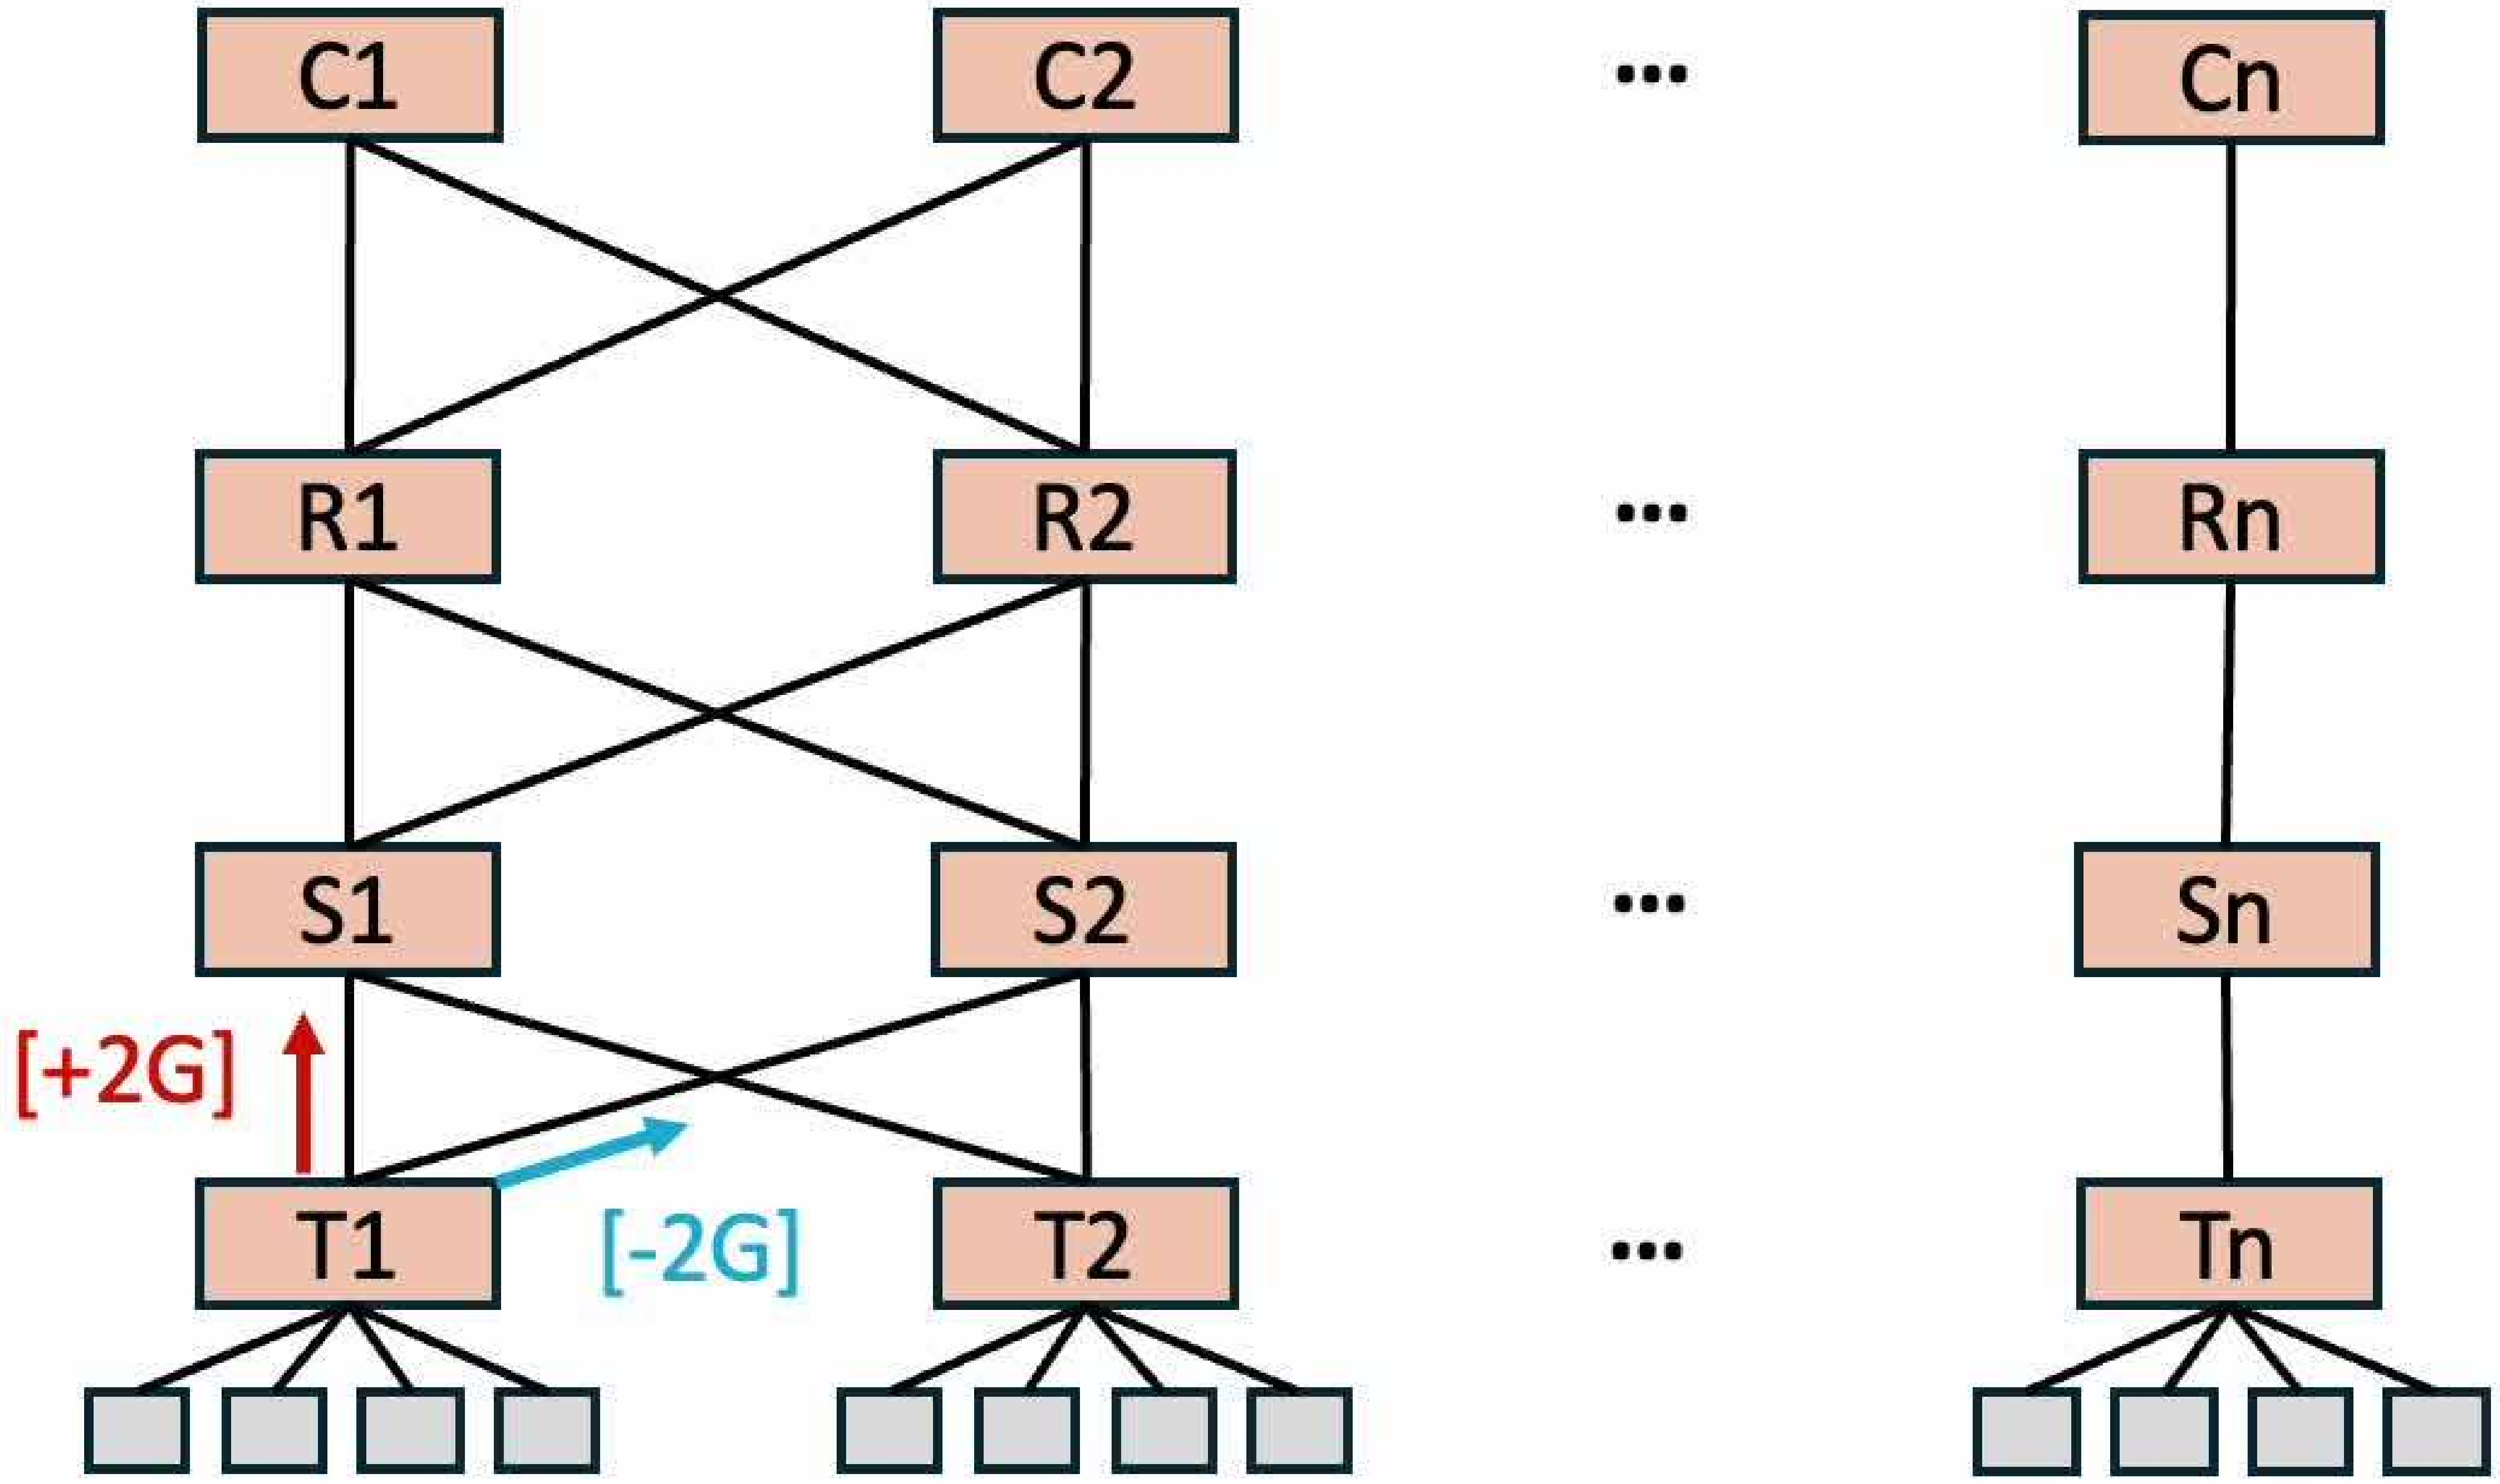

Precision loss accumulates in hierarchical networks.

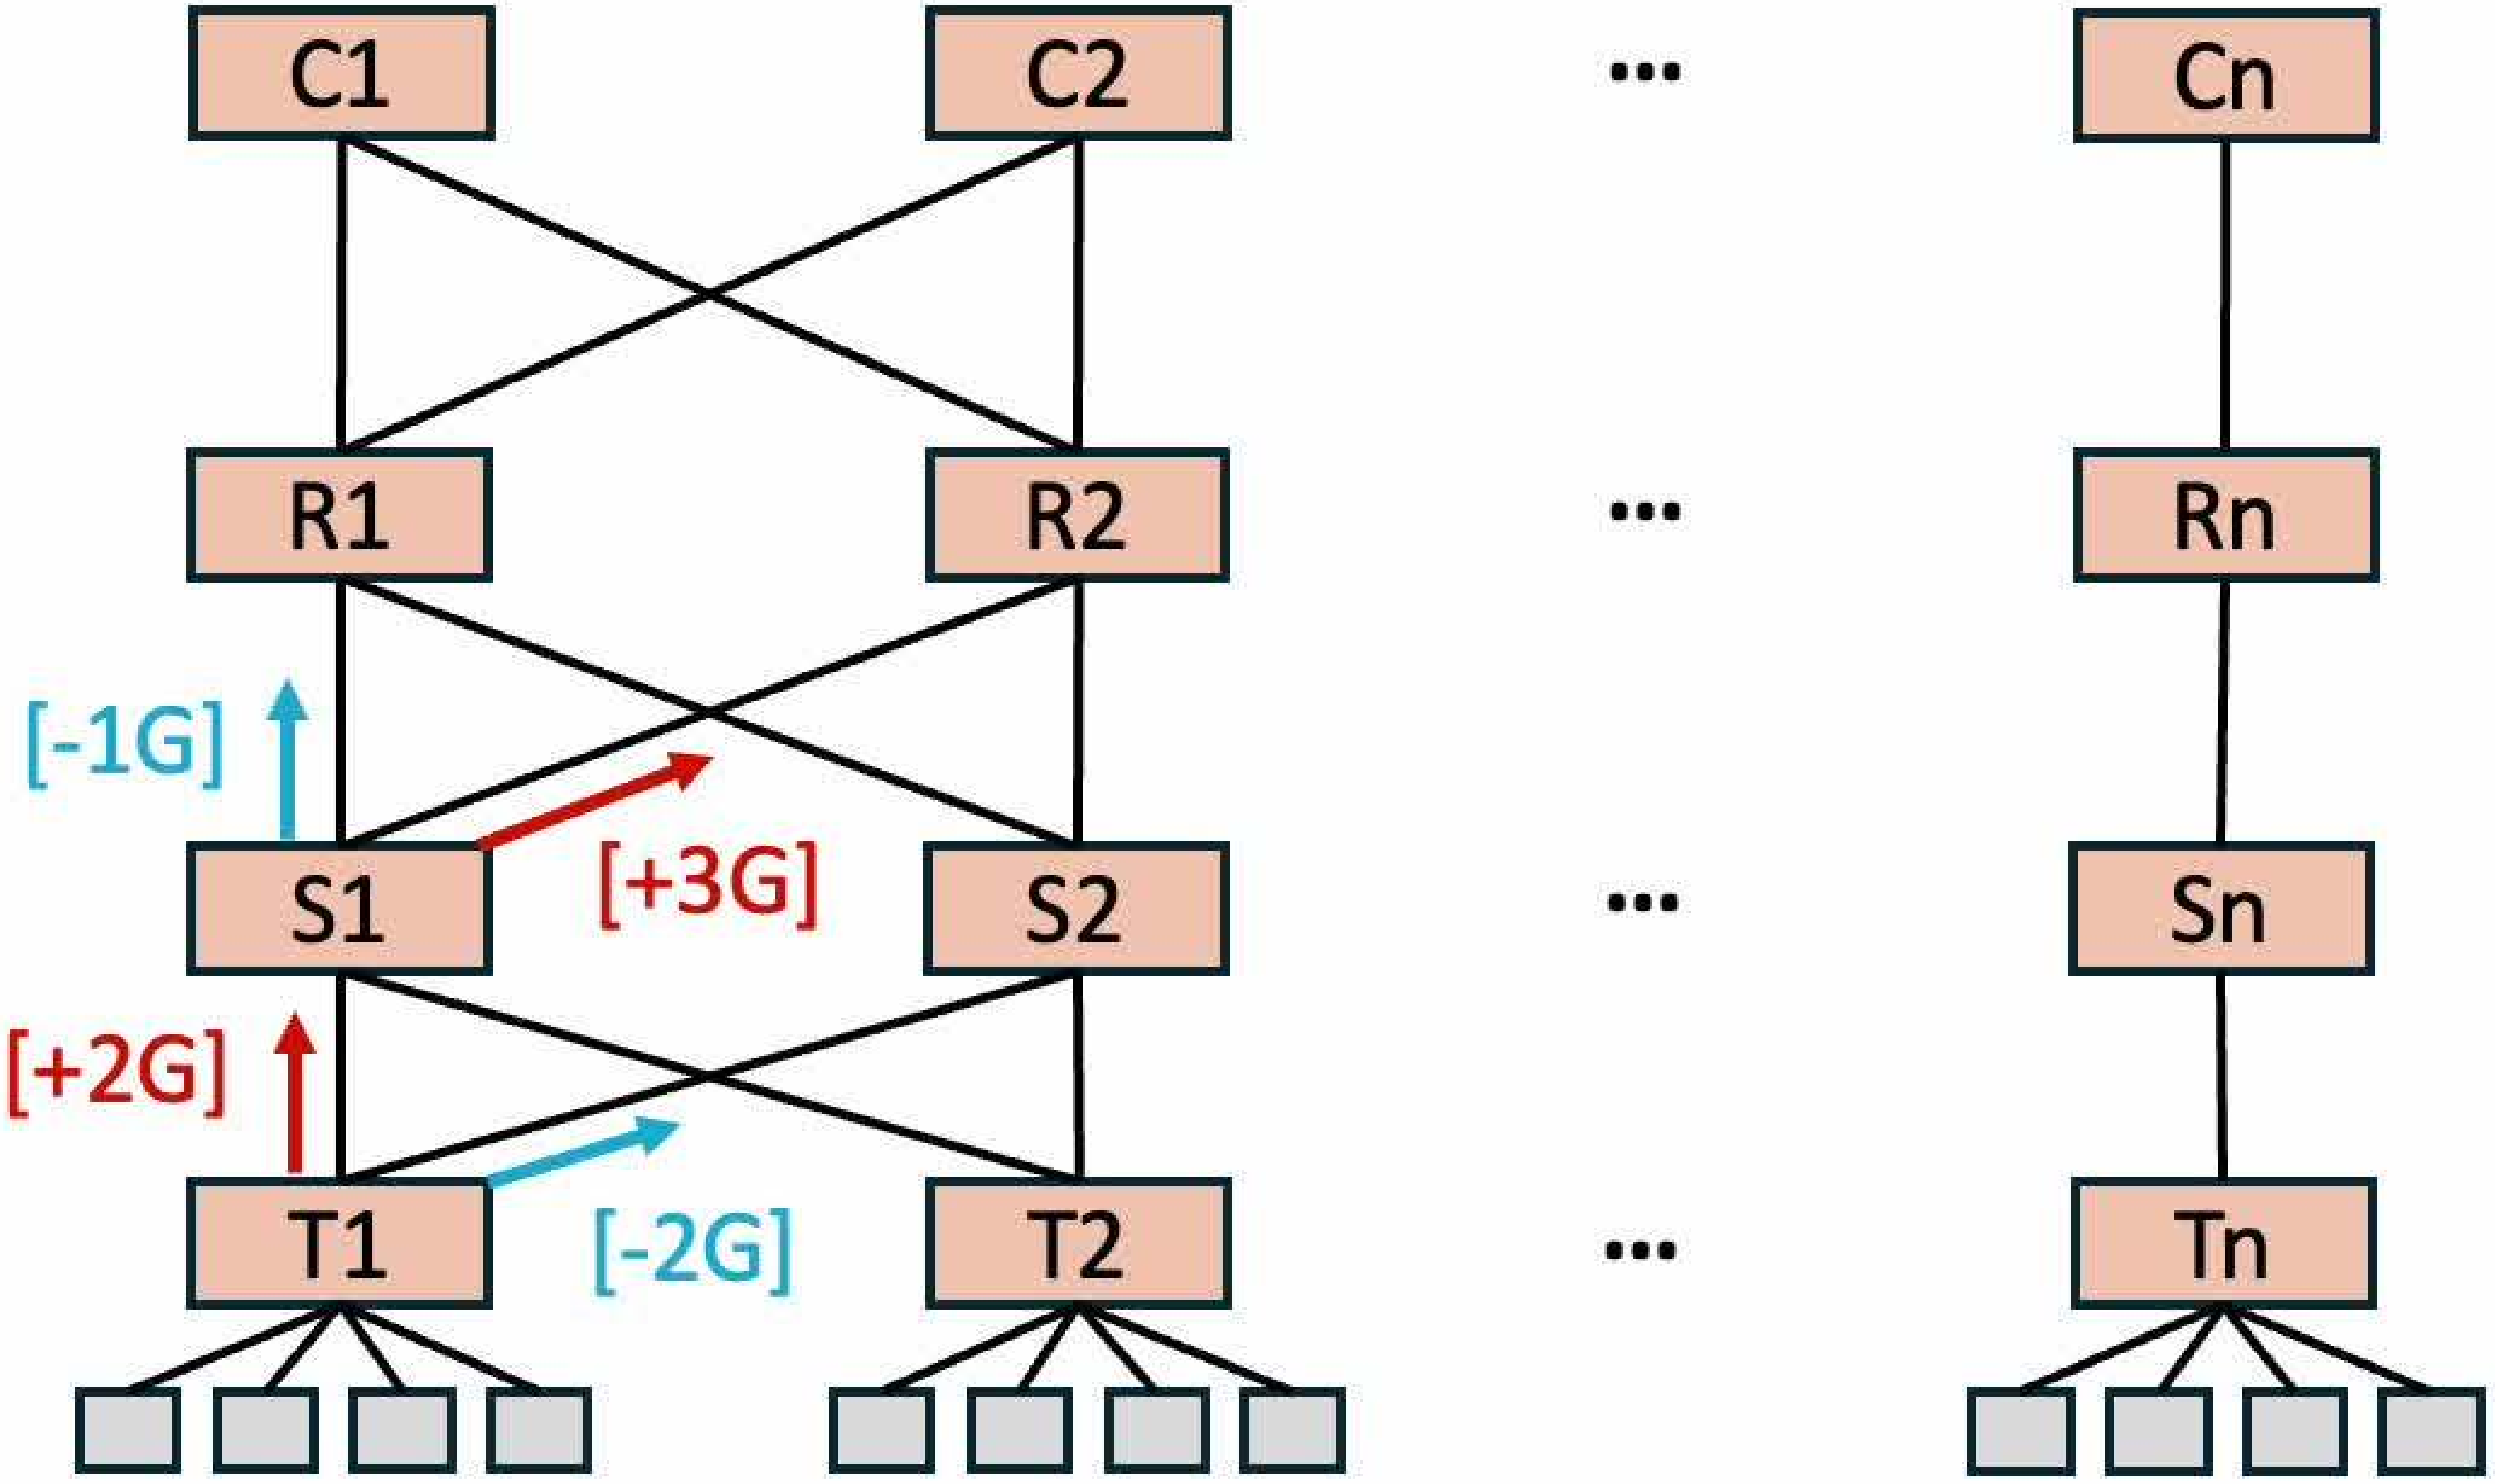

Carnegie Mellon University

Precision loss accumulates in hierarchical networks.

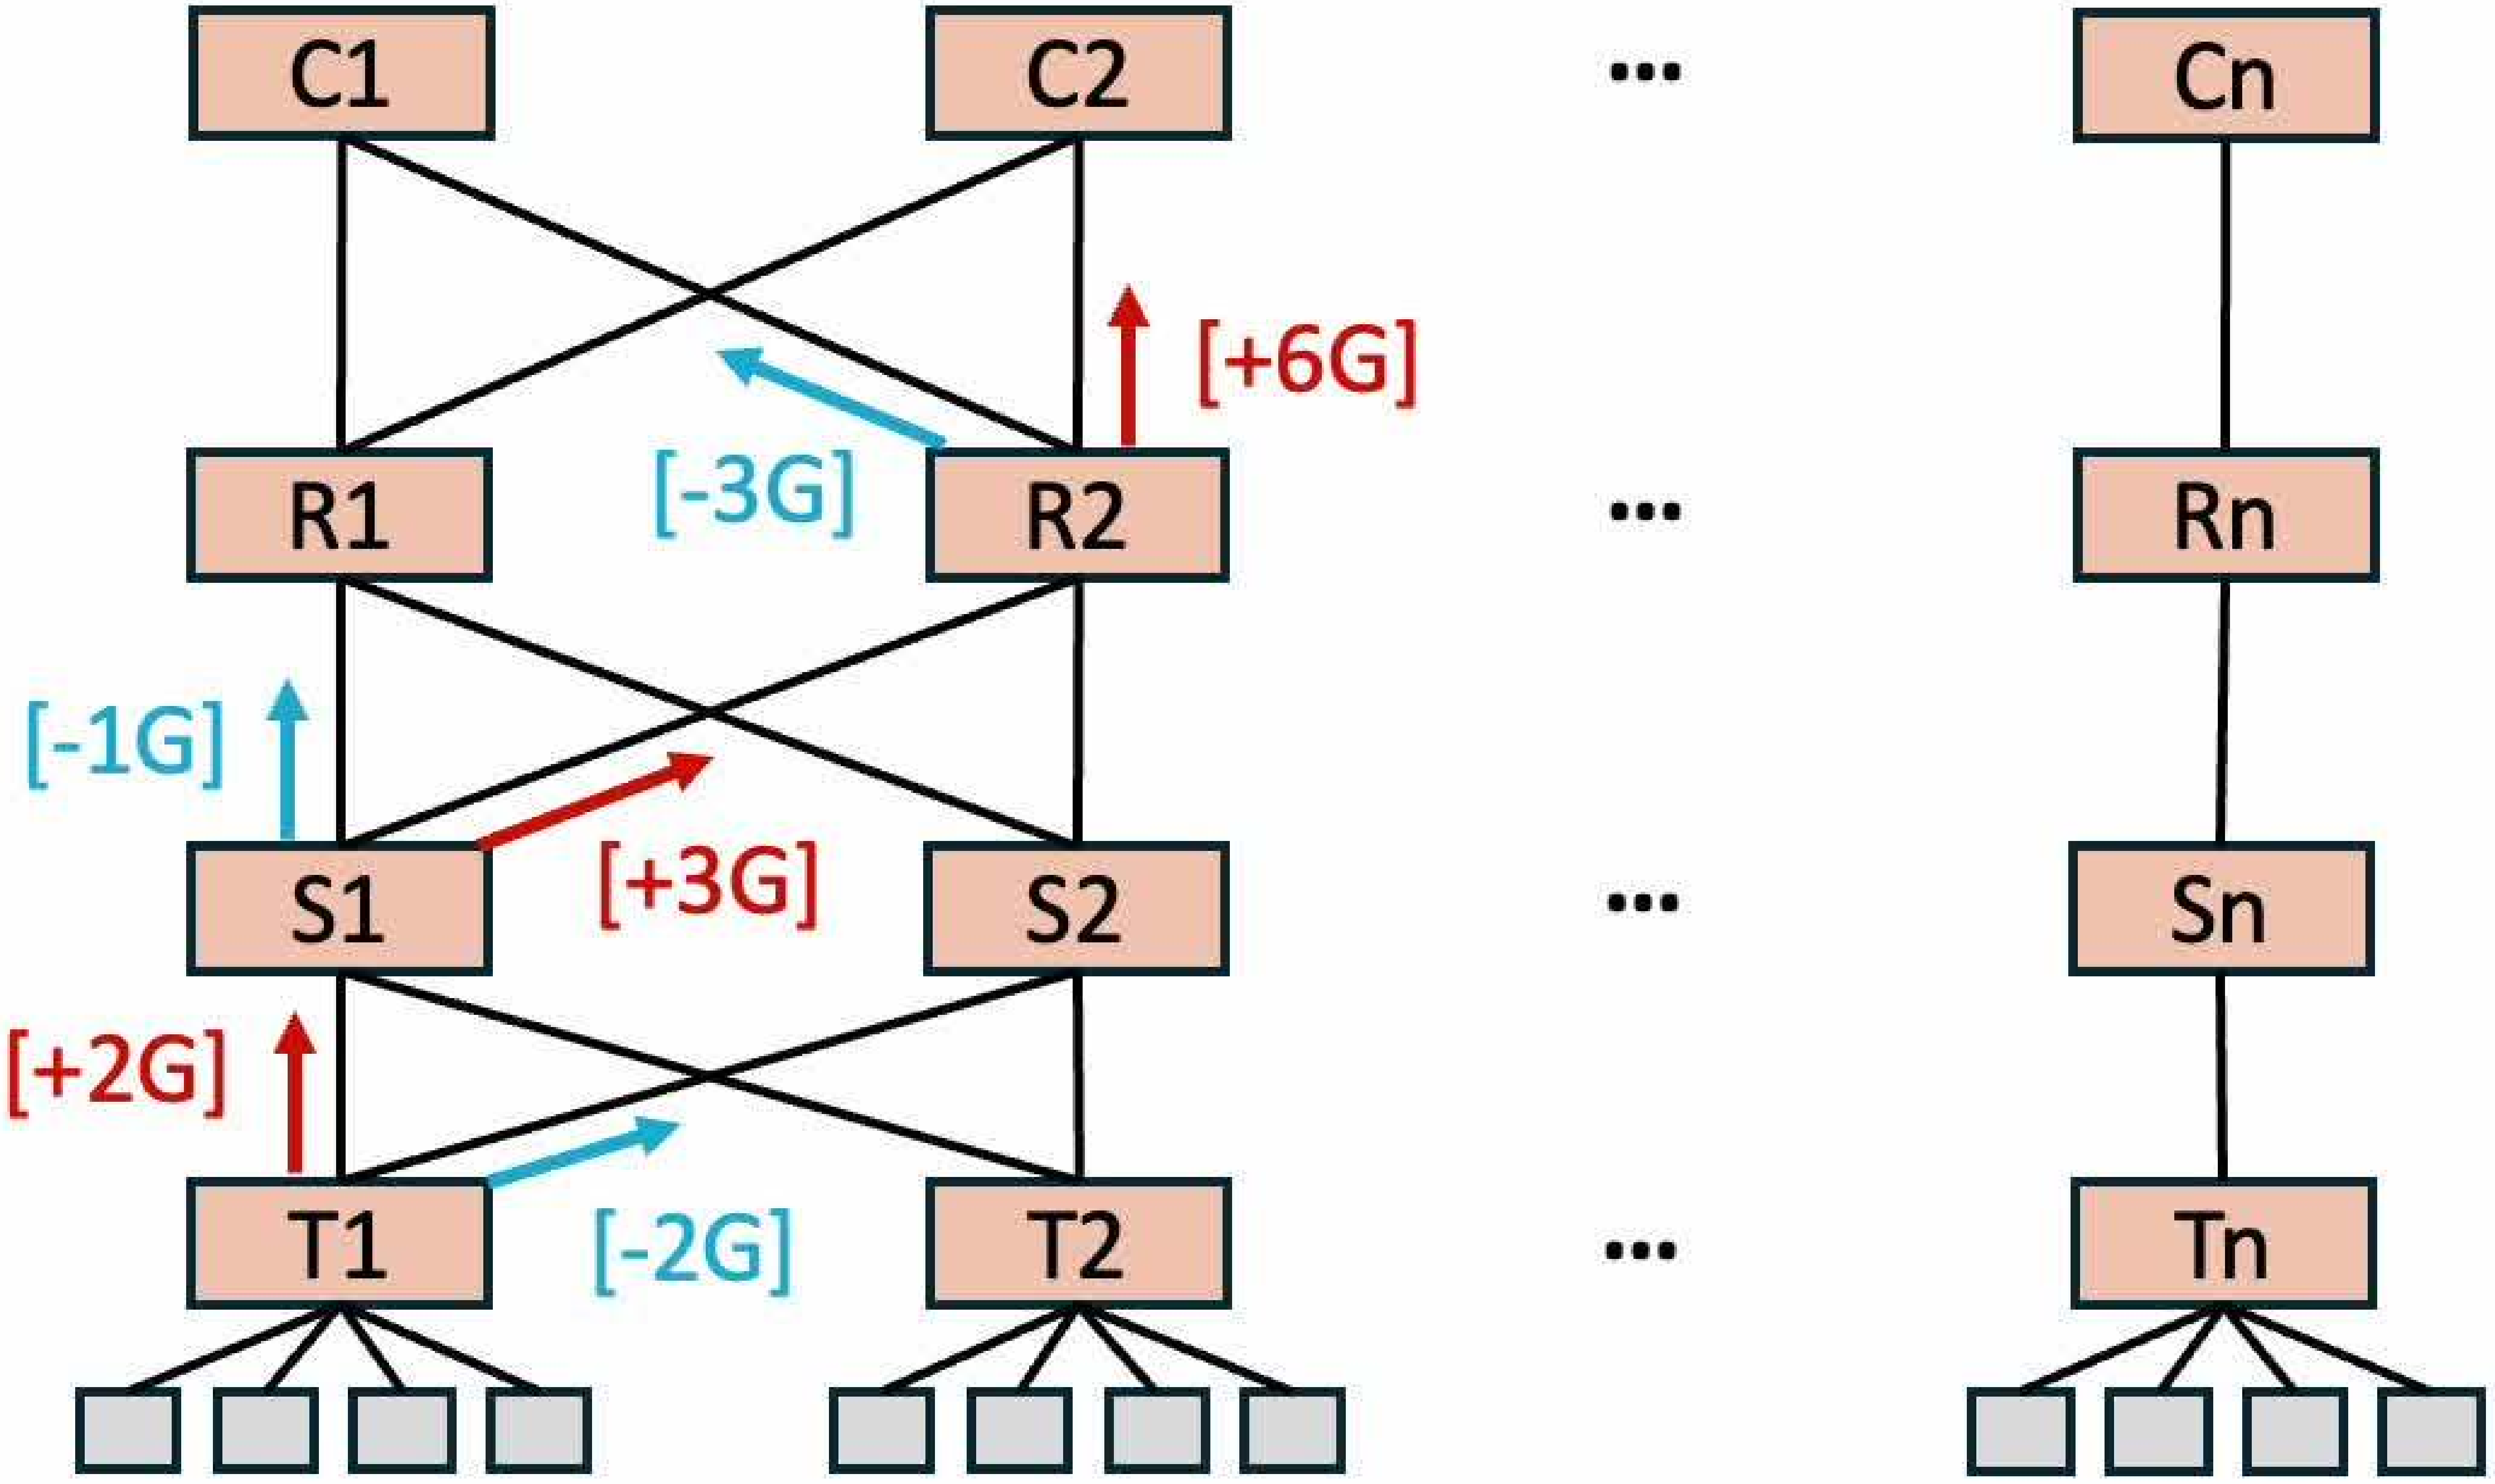

Carnegie Mellon University

Real world precision loss can be severe.

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

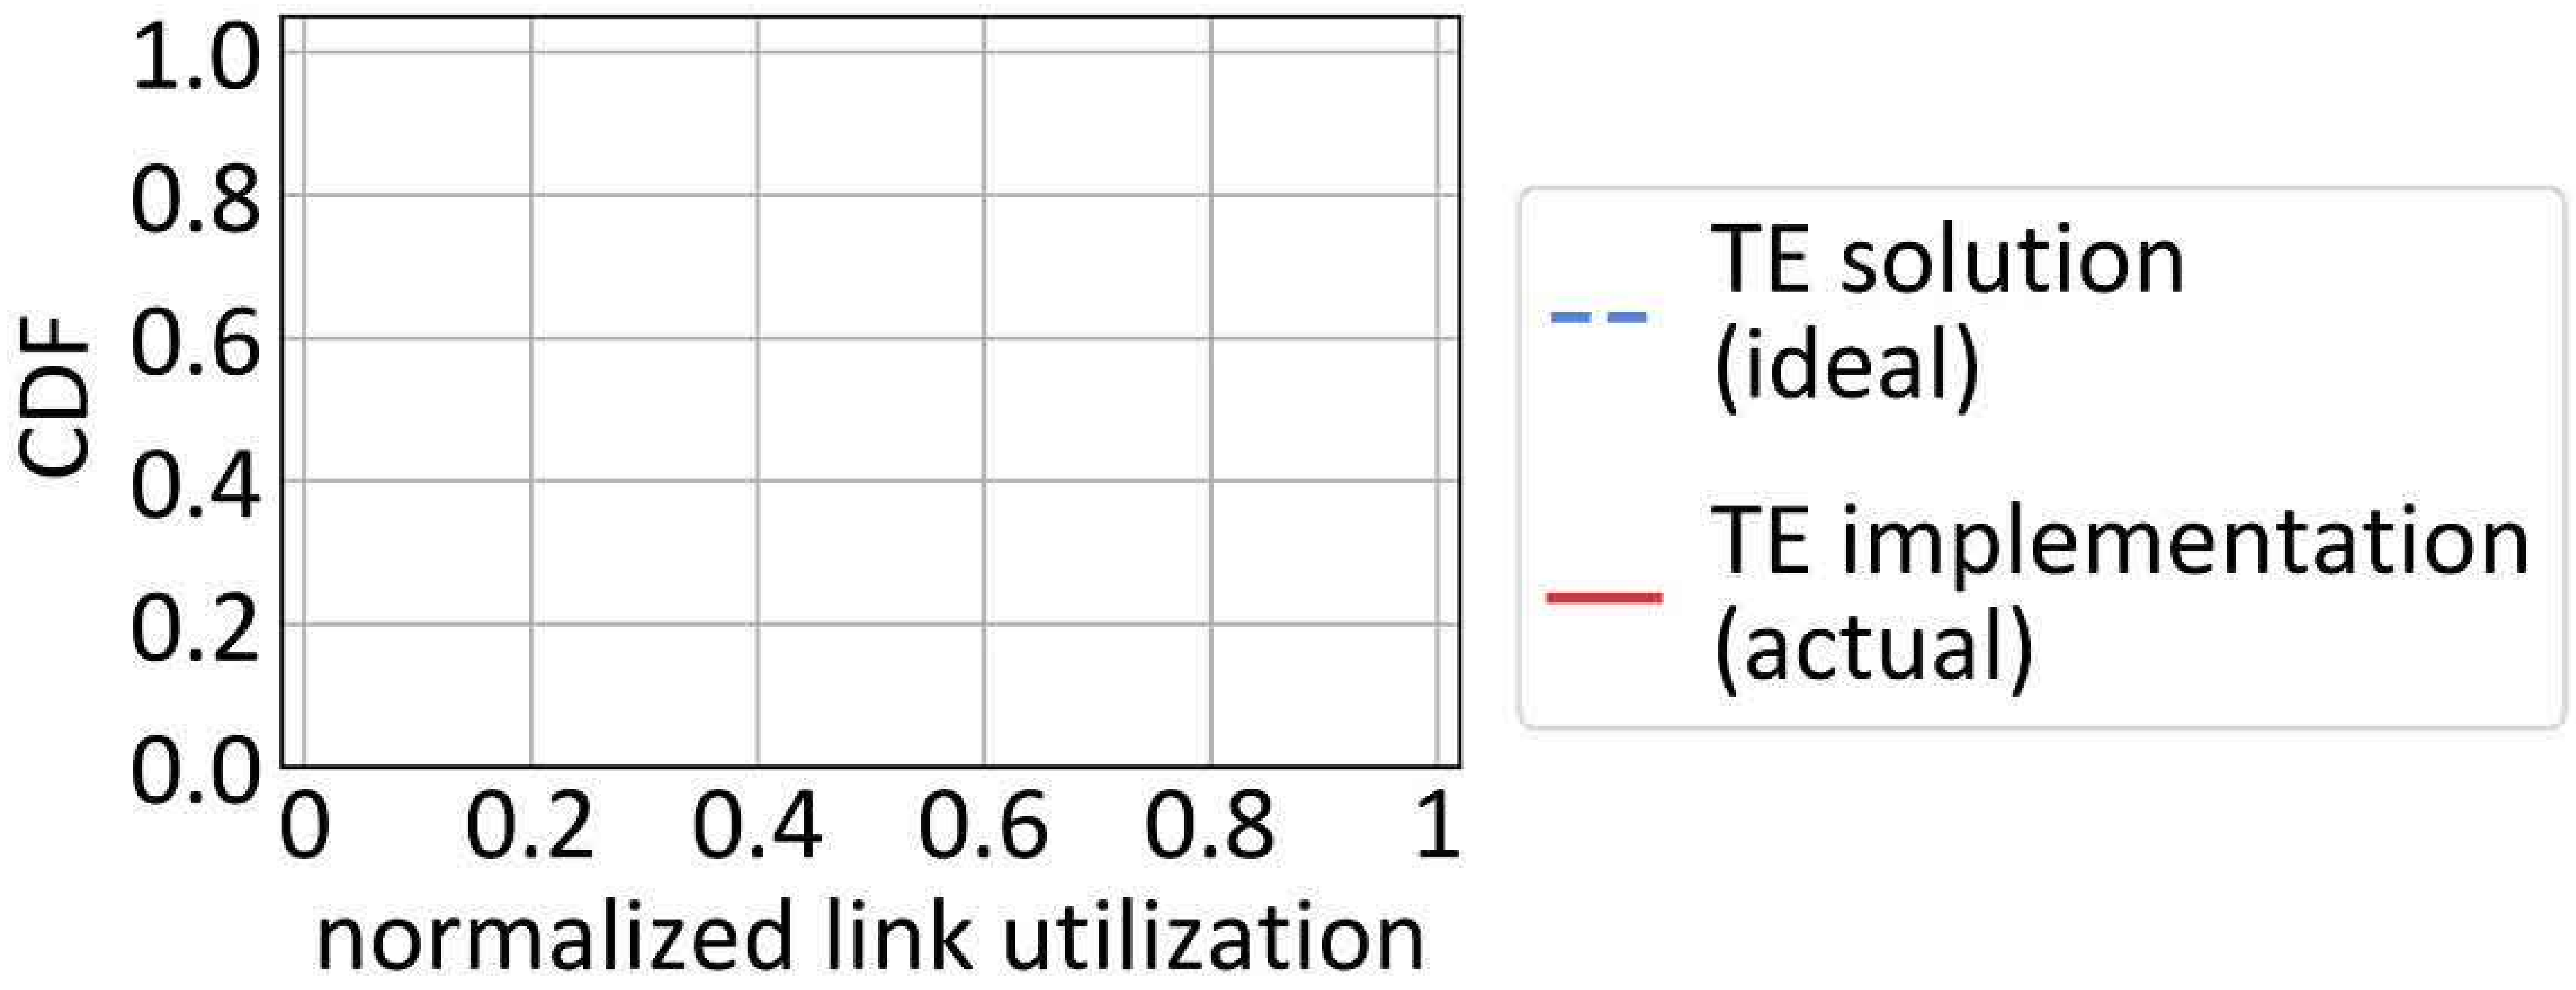

Carnegie Mellon University

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

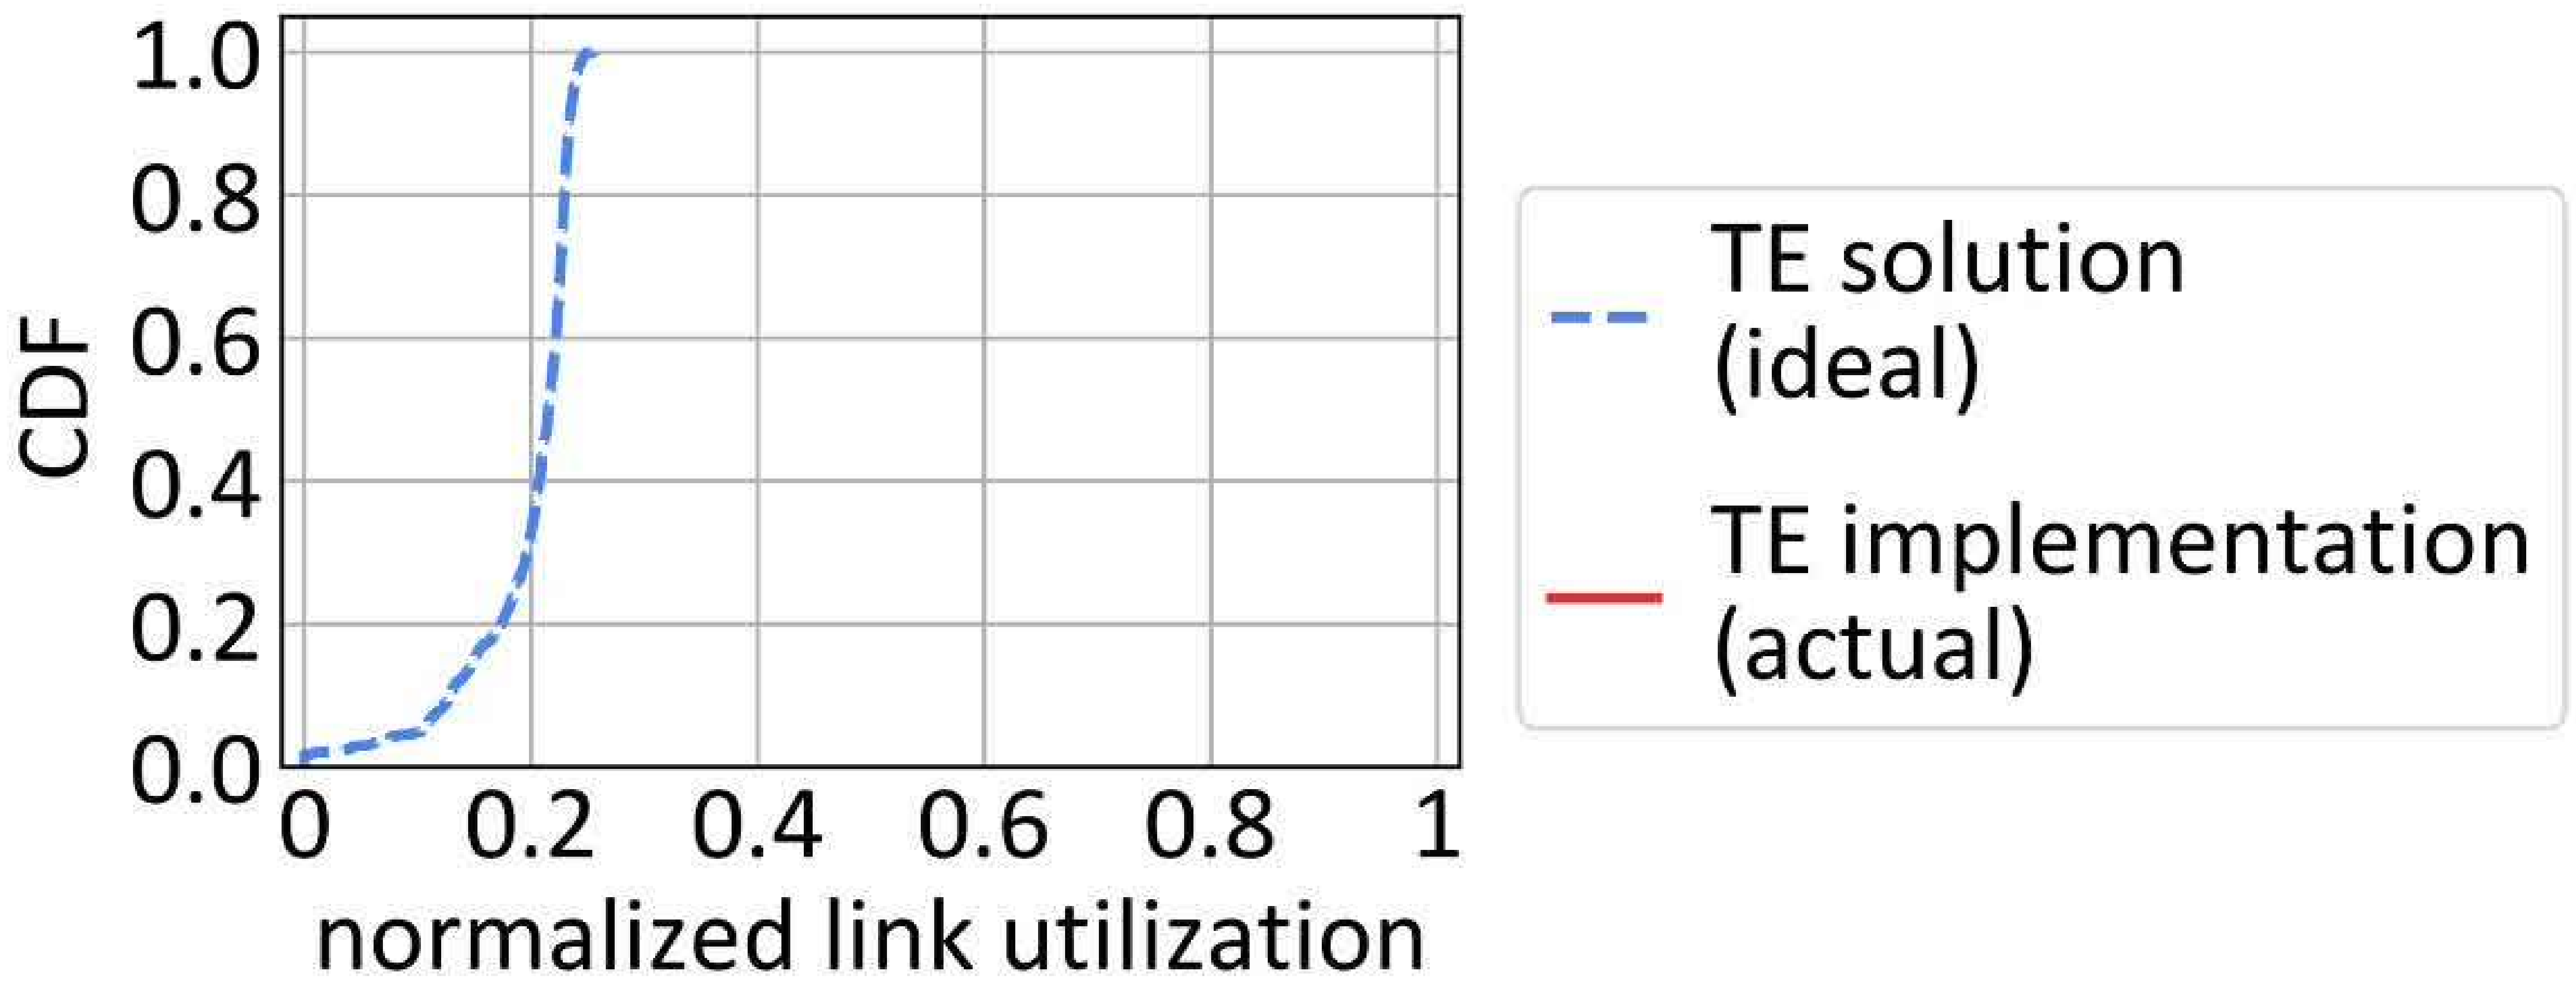

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

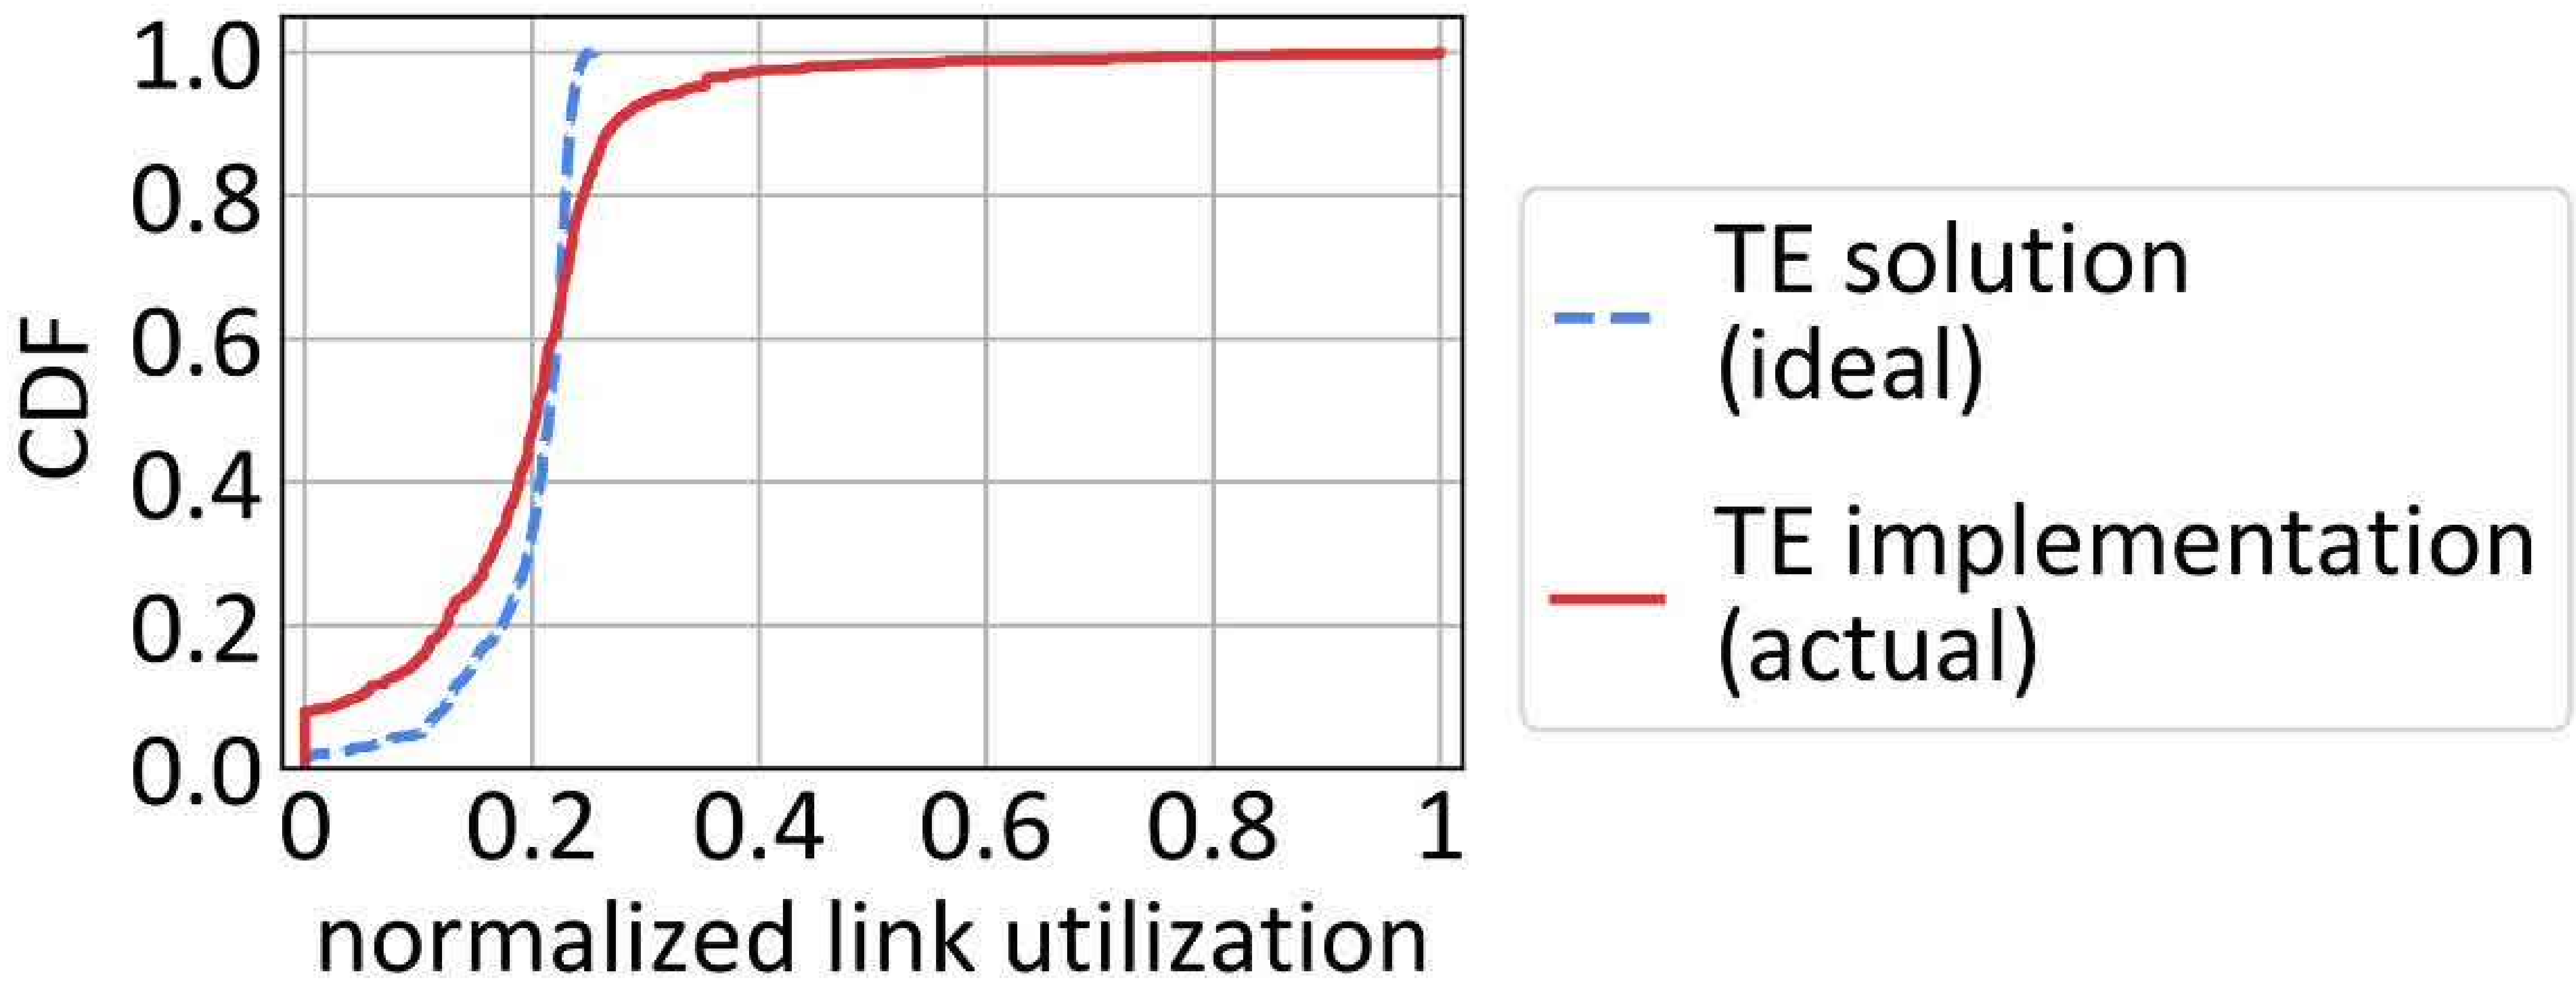

Carnegie Mellon University

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

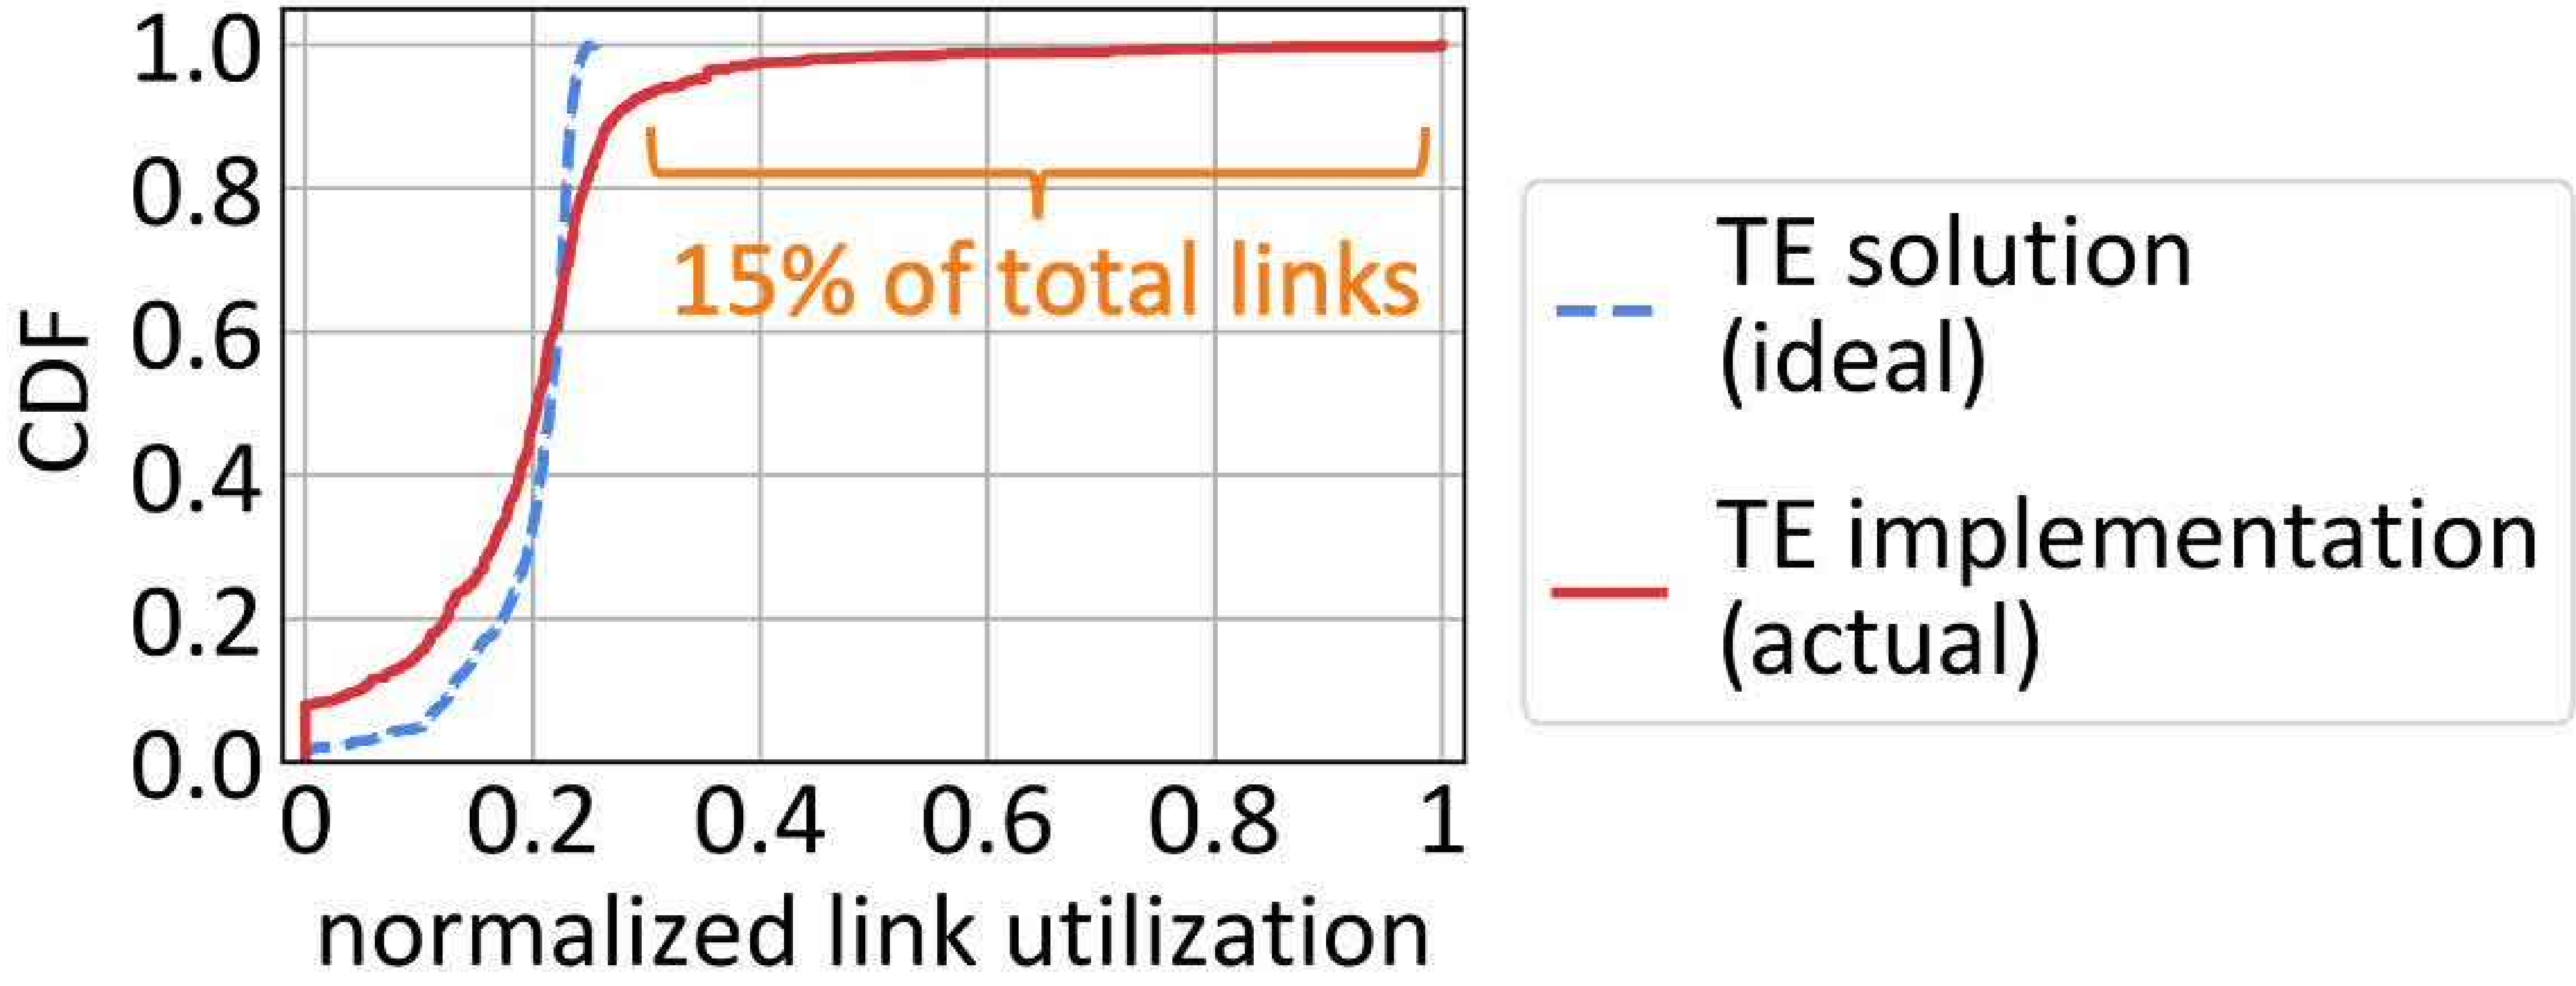

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.

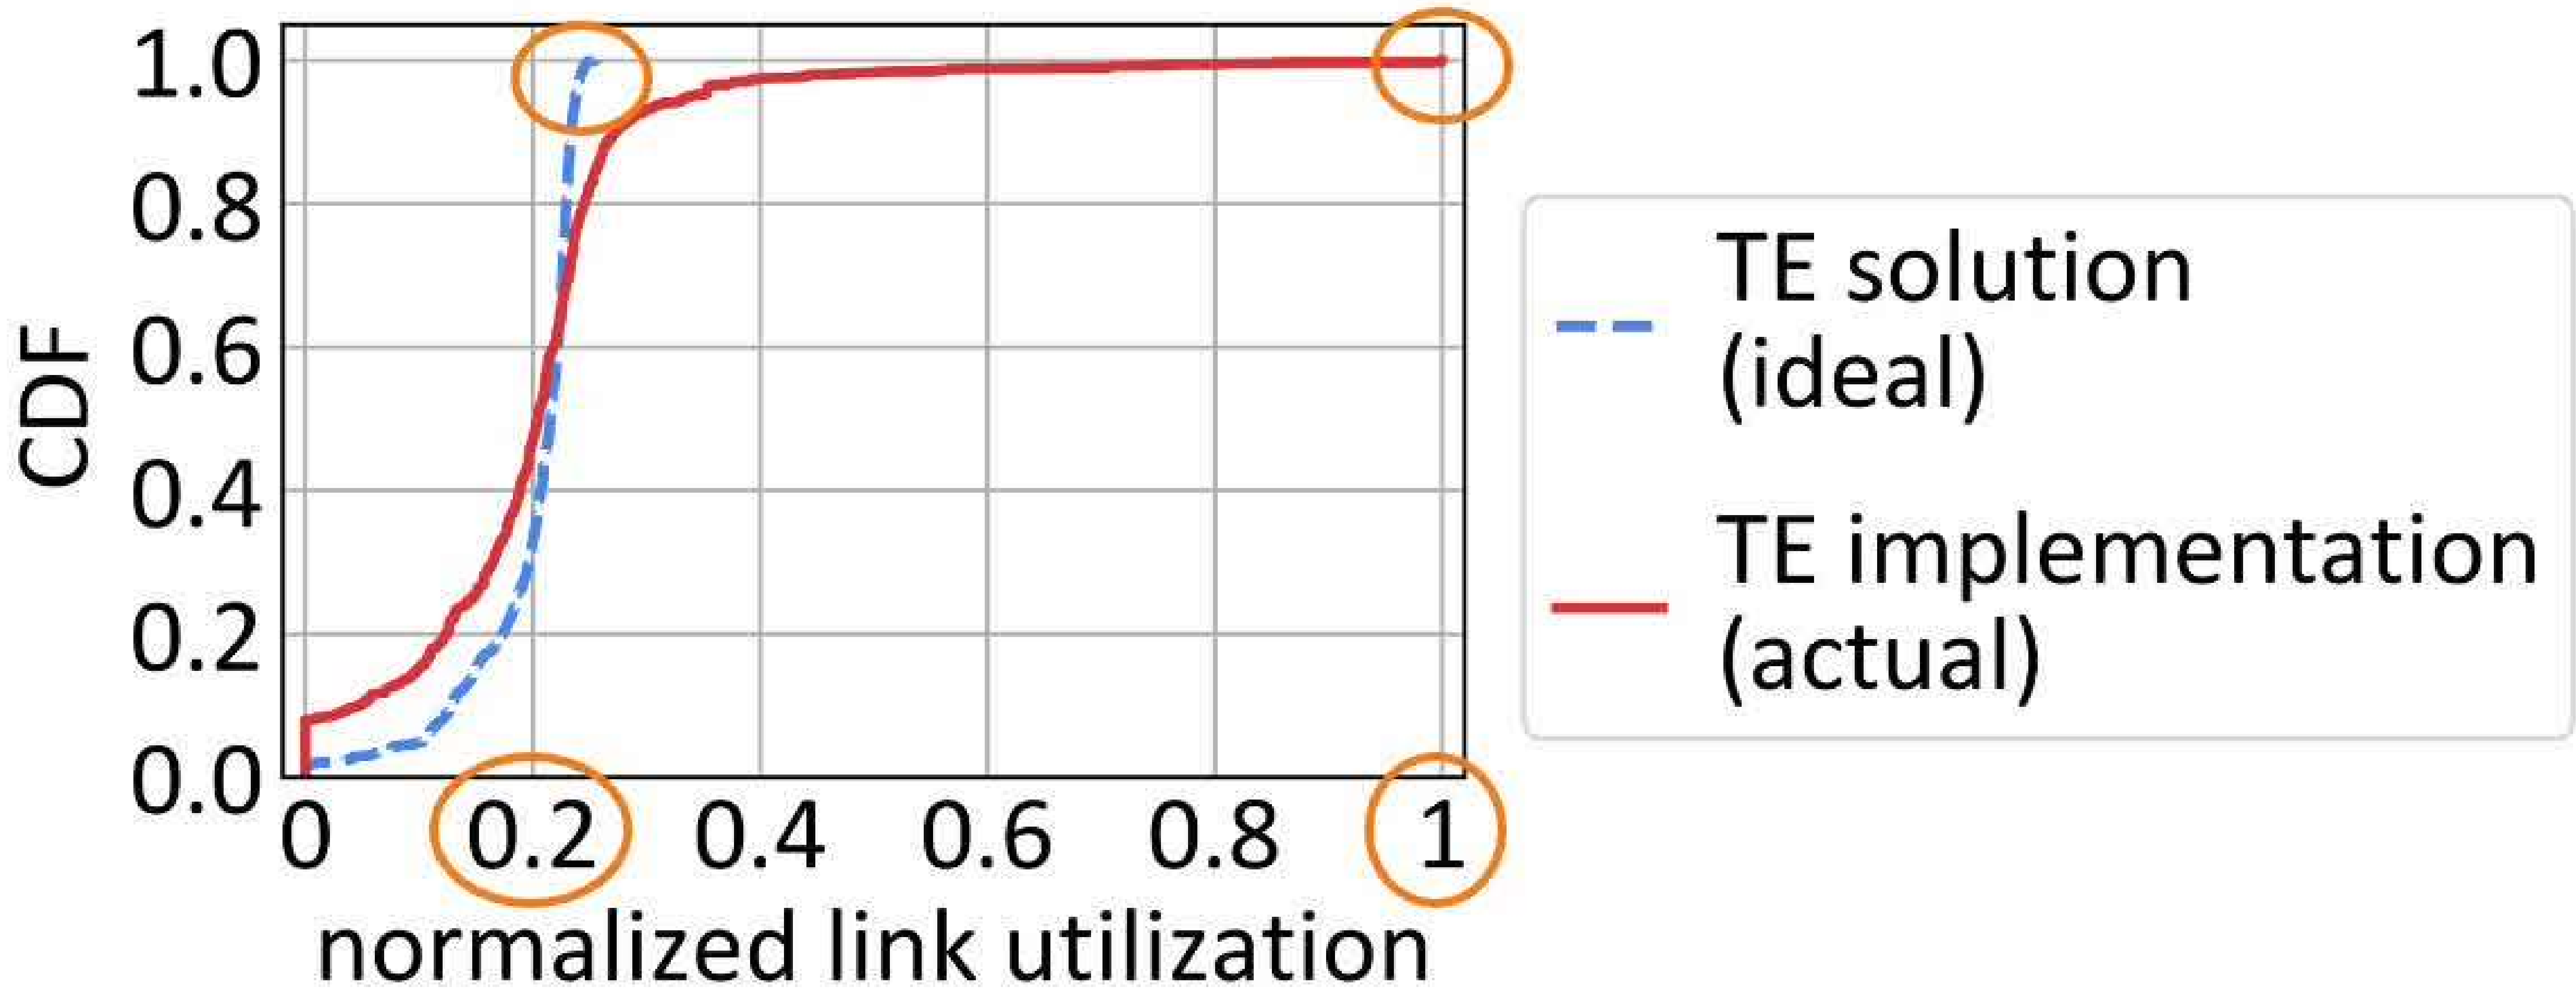

Carnegie Mellon University

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.
- Actual link utilization up to 5x higher than ideal.

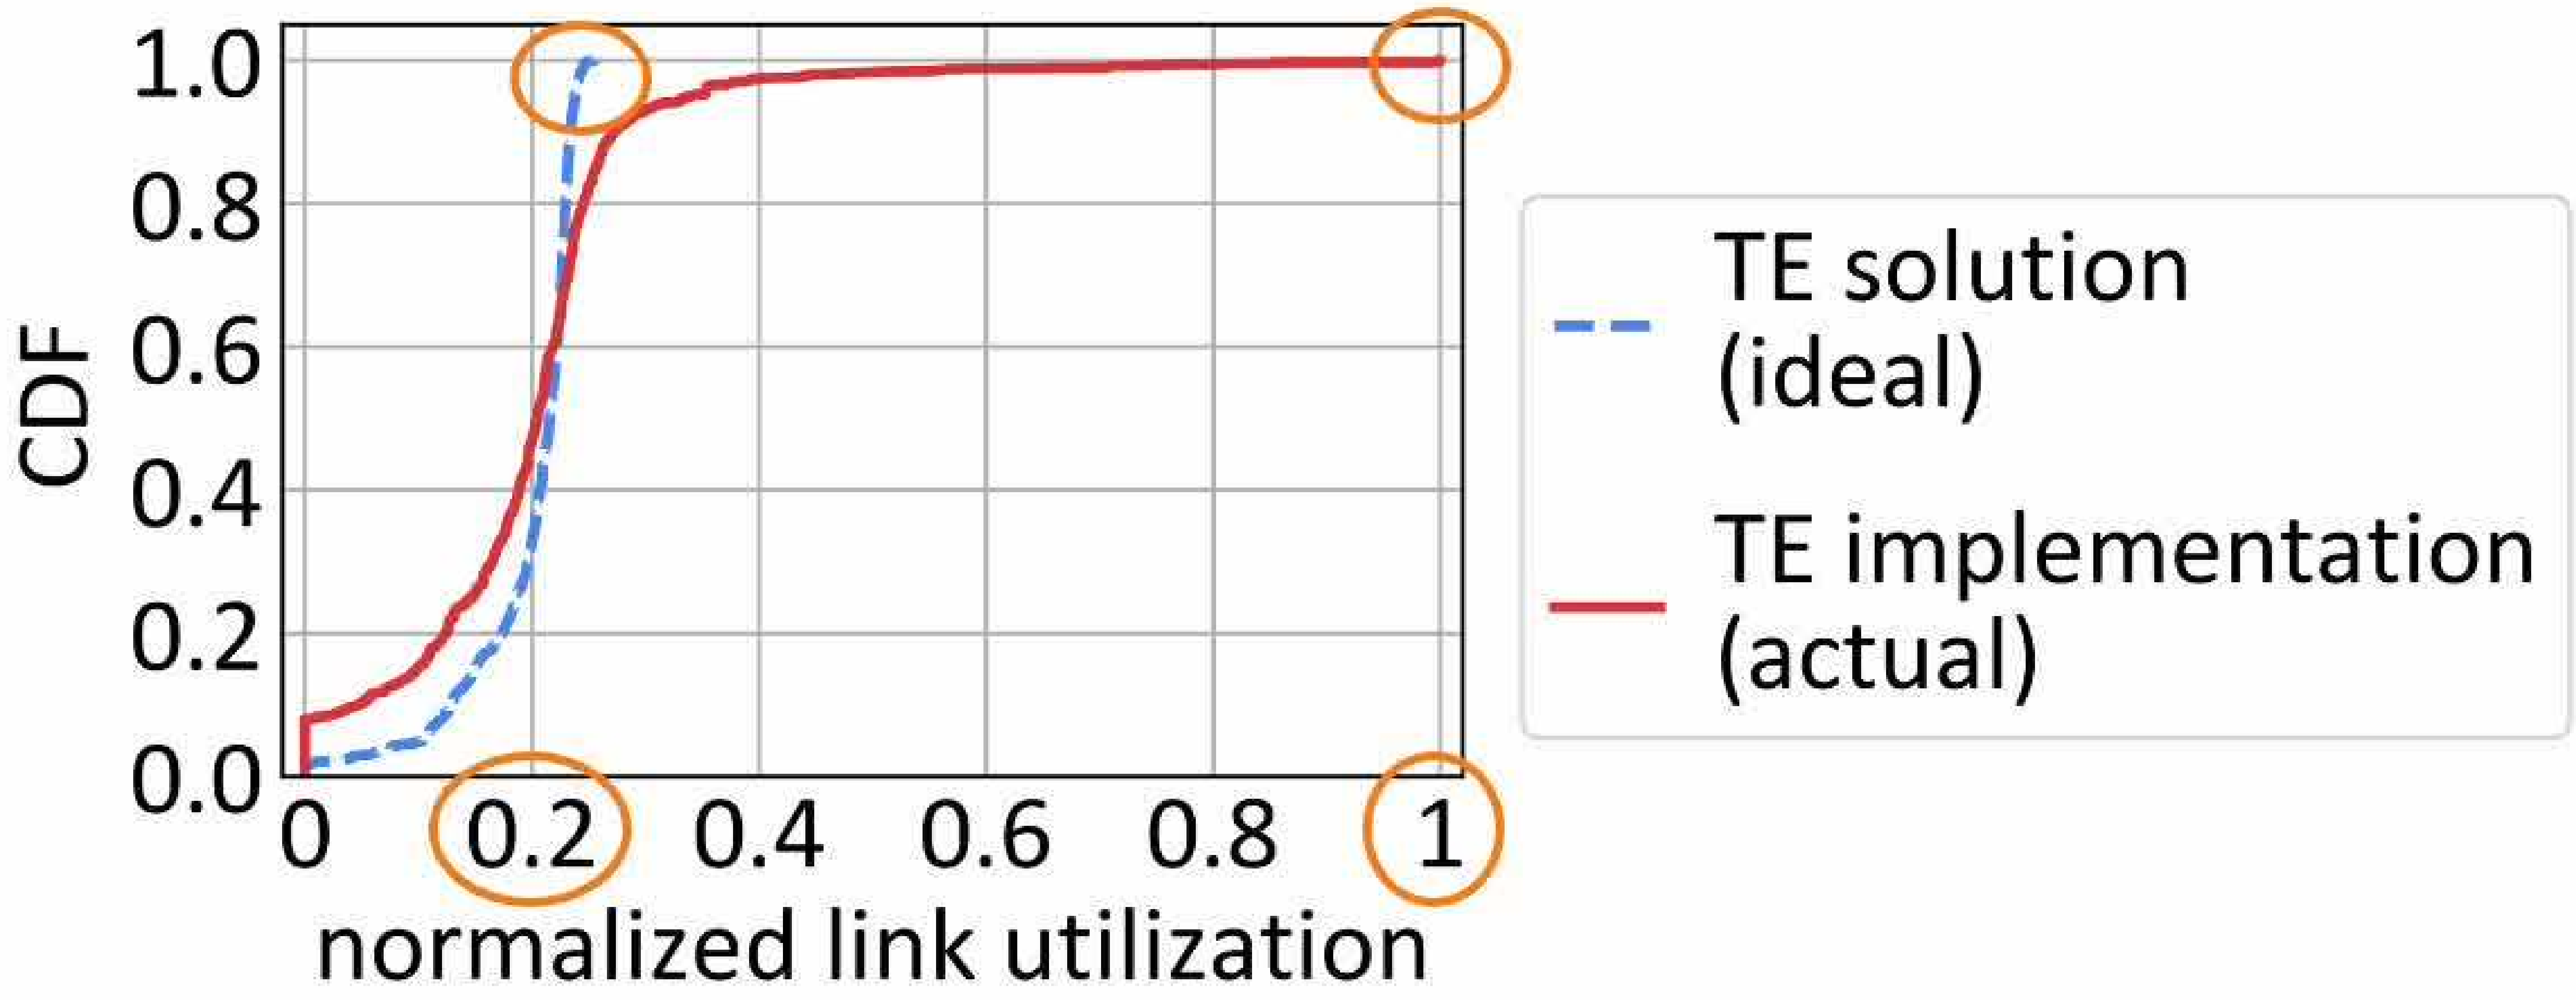

Carnegie Mellon University

Real world precision loss can be severe.

- Flow completion time @p99 up by 40% if link utilization higher by 10%.
  - Above numbers change with different baseline link utilizations.
  - Congested links cause packet loss and retransmission, delay transfer.
- Lower link utilization hugely benefits flow completion time.
- Actual link utilization up to 5x higher than ideal.

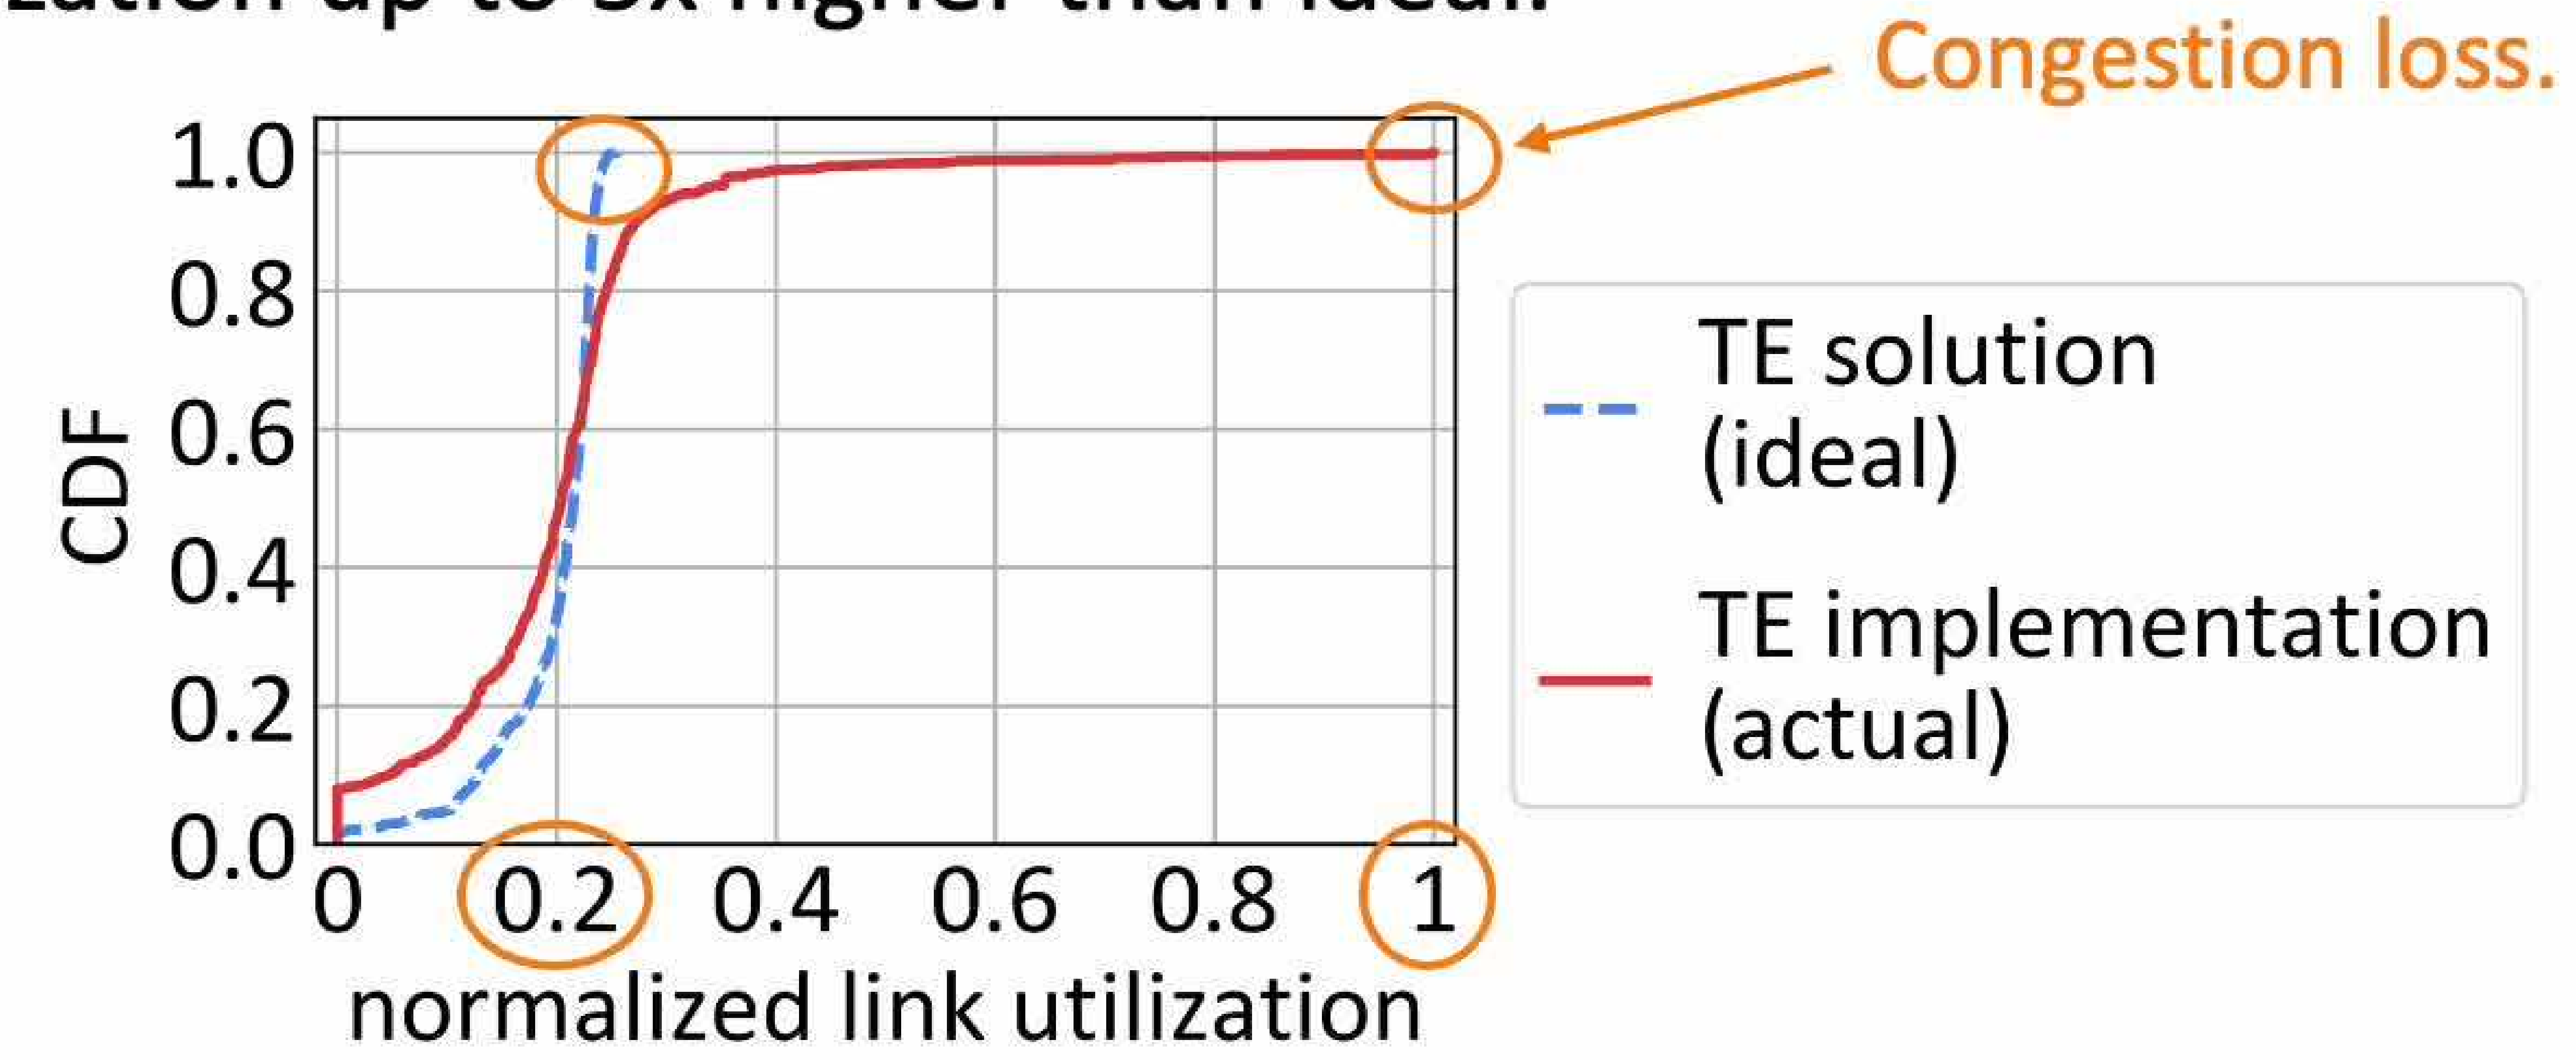

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).

Group	
ID	Egress port
G1	p1
	p2
	…
	p64
G2	p1
	p2
…	
GN	p1
	…
	p1
	p2

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **<u>scale</u>**
  - # groups scales with network size.

Group	
ID	Egress port
G1	p1
	p2
	…
	p64
G2	p1
	p2
…	
GN	p1
	…
	p1
	p2

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**

Group	
ID	Egress port
G1	p1
	p2
	...
	p64
G2	p1
	p2
...	
GN	p1
	...
	p1
	p2

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**

Group	
ID	Egress port
G1	p1
	p2
	...
	p64
G2	p1
	p2
...	
GN	p1
	...
	p1
	p2

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**
  - Skewed weights are hard to reduce.

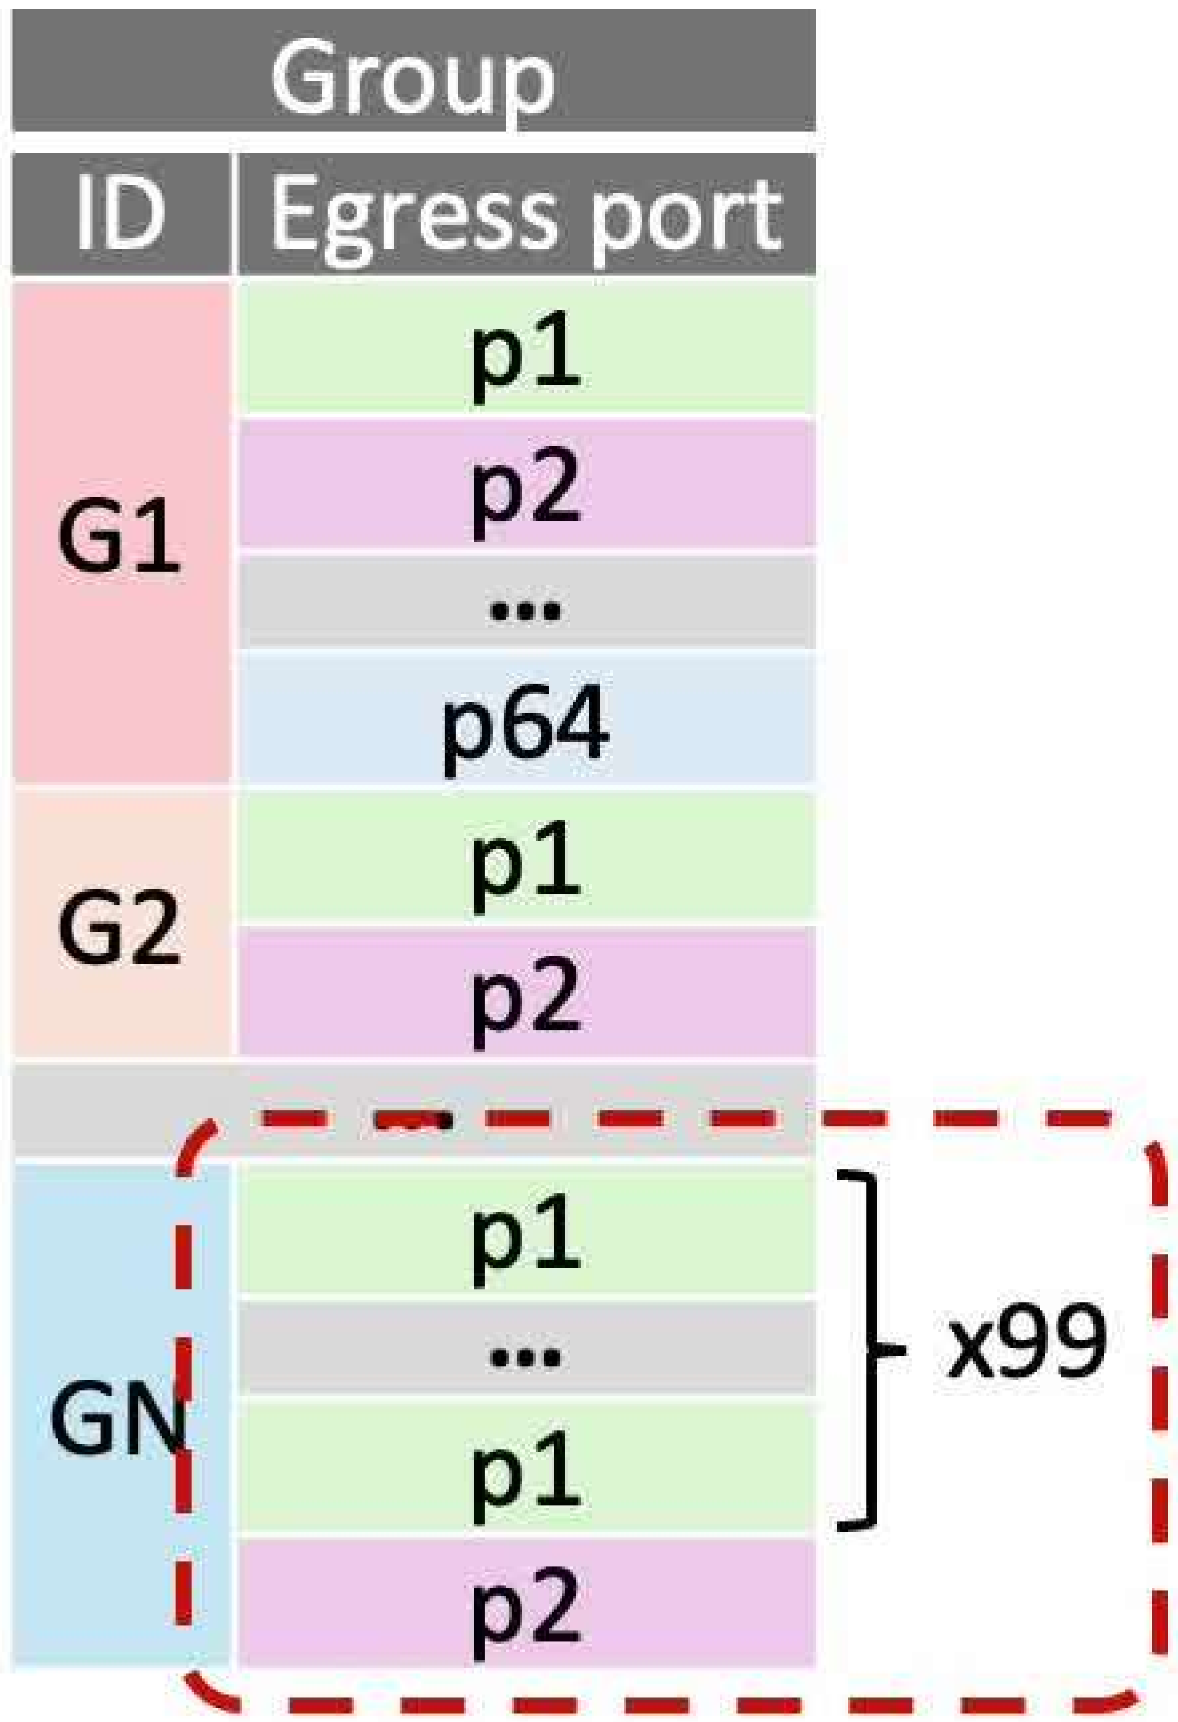

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**
  - Skewed weights are hard to reduce.
- Cause 4: **heterogeneity**

Group	
ID	Egress port
G1	p1
	p2
	...
	p64
G2	p1
	p2
...	
GN	p1
	...
	p1
	p2

Carnegie Mellon University

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**
  - Skewed weights are hard to reduce.
- Cause 4: **heterogeneity**
  - Old gen. switches have small memory space.

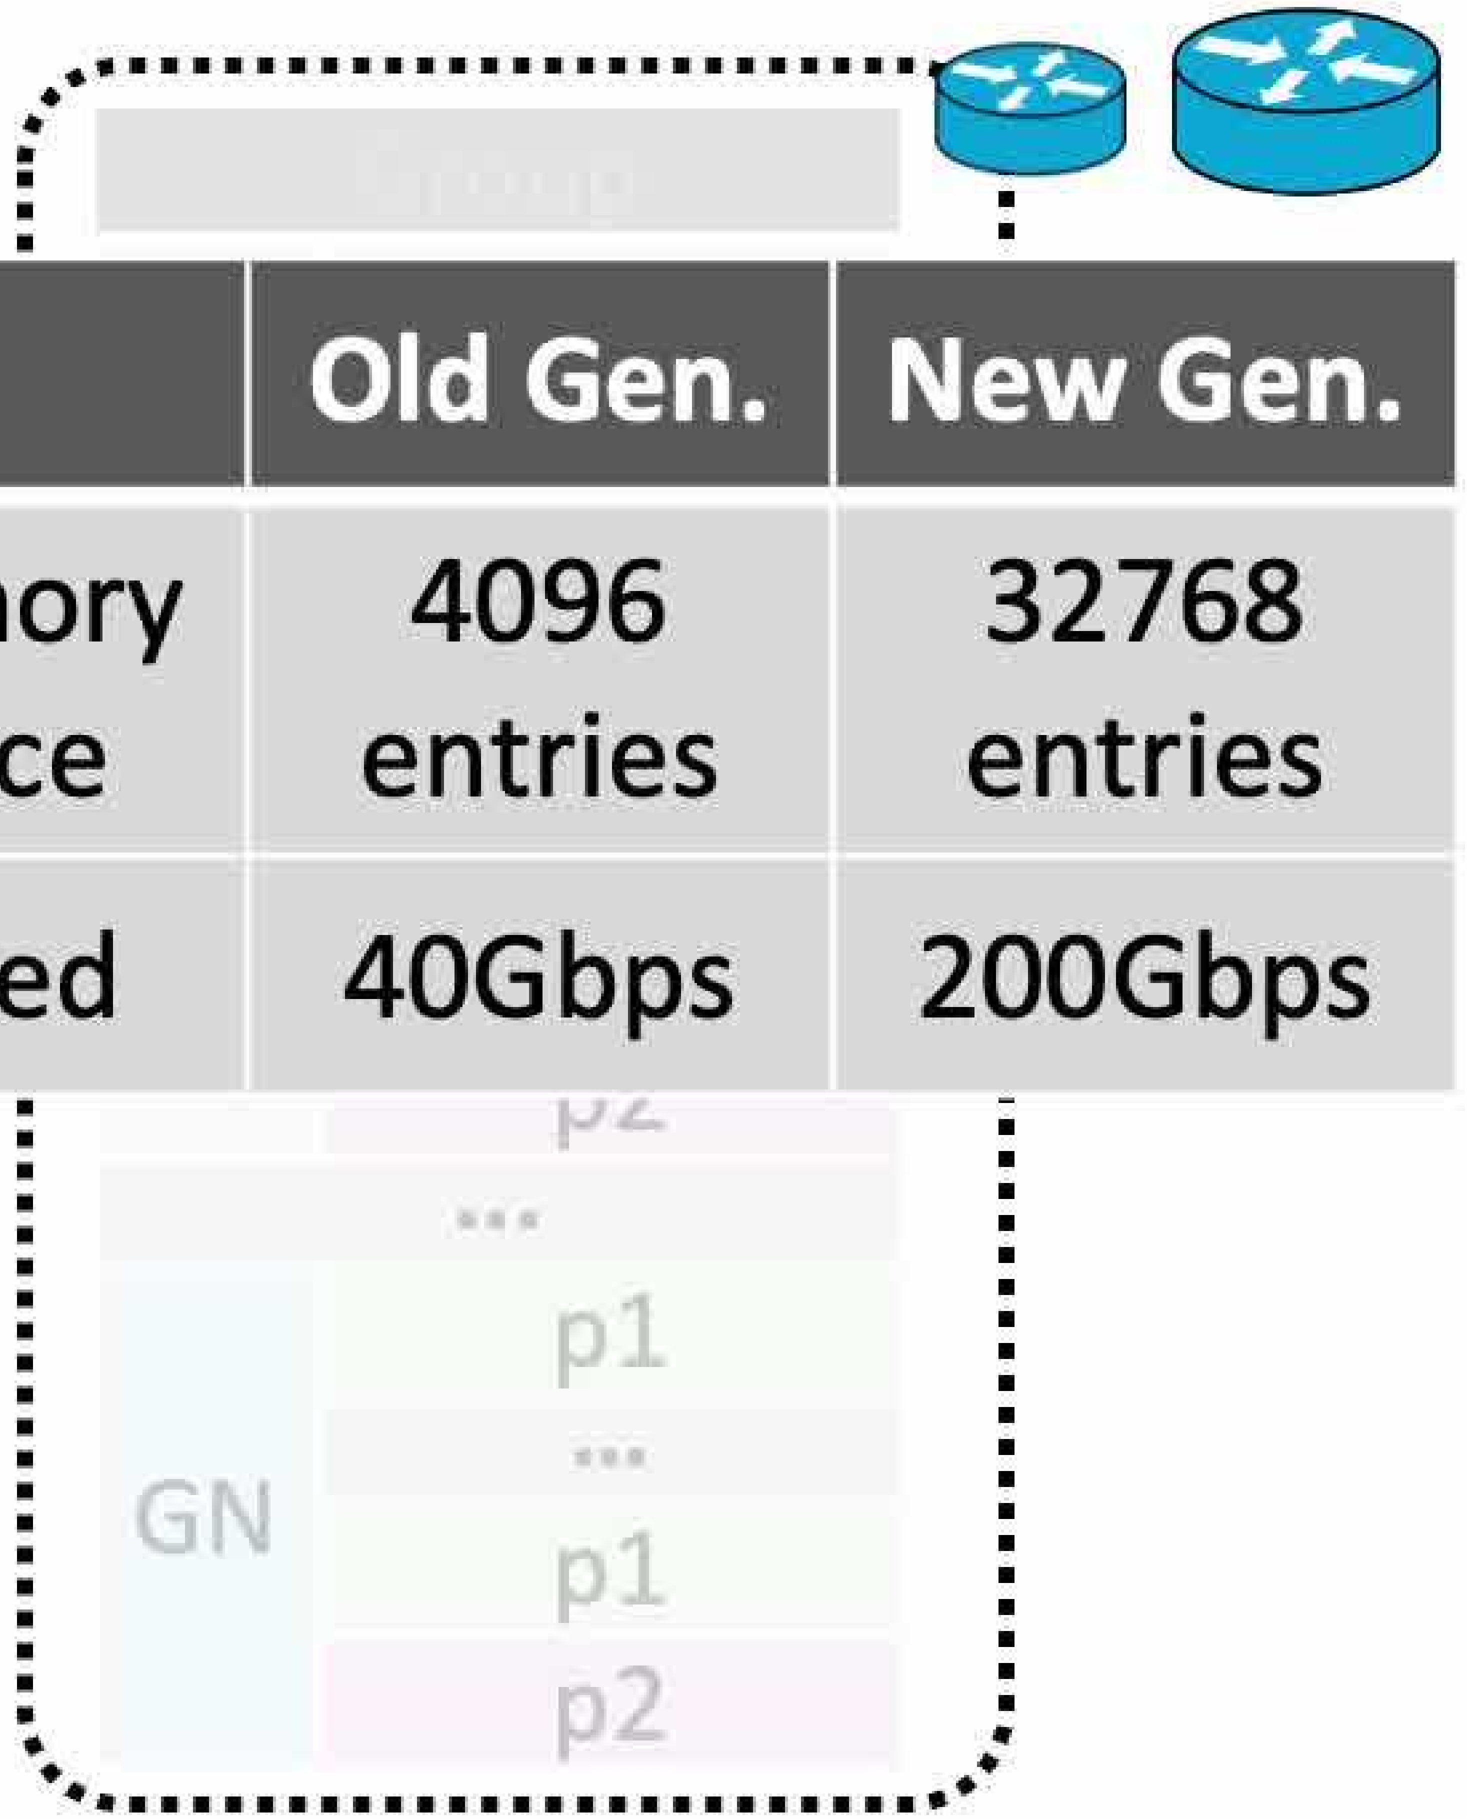

	Old Gen.	New Gen.
Memory space	4096 entries	32768 entries
Speed	40Gbps	200Gbps

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**
  - Skewed weights are hard to reduce.
- Cause 4: **heterogeneity**
  - Old gen. switches have small memory space.
  - Mixed speed leads to larger weights.
- Cause 5: **cascading effect**

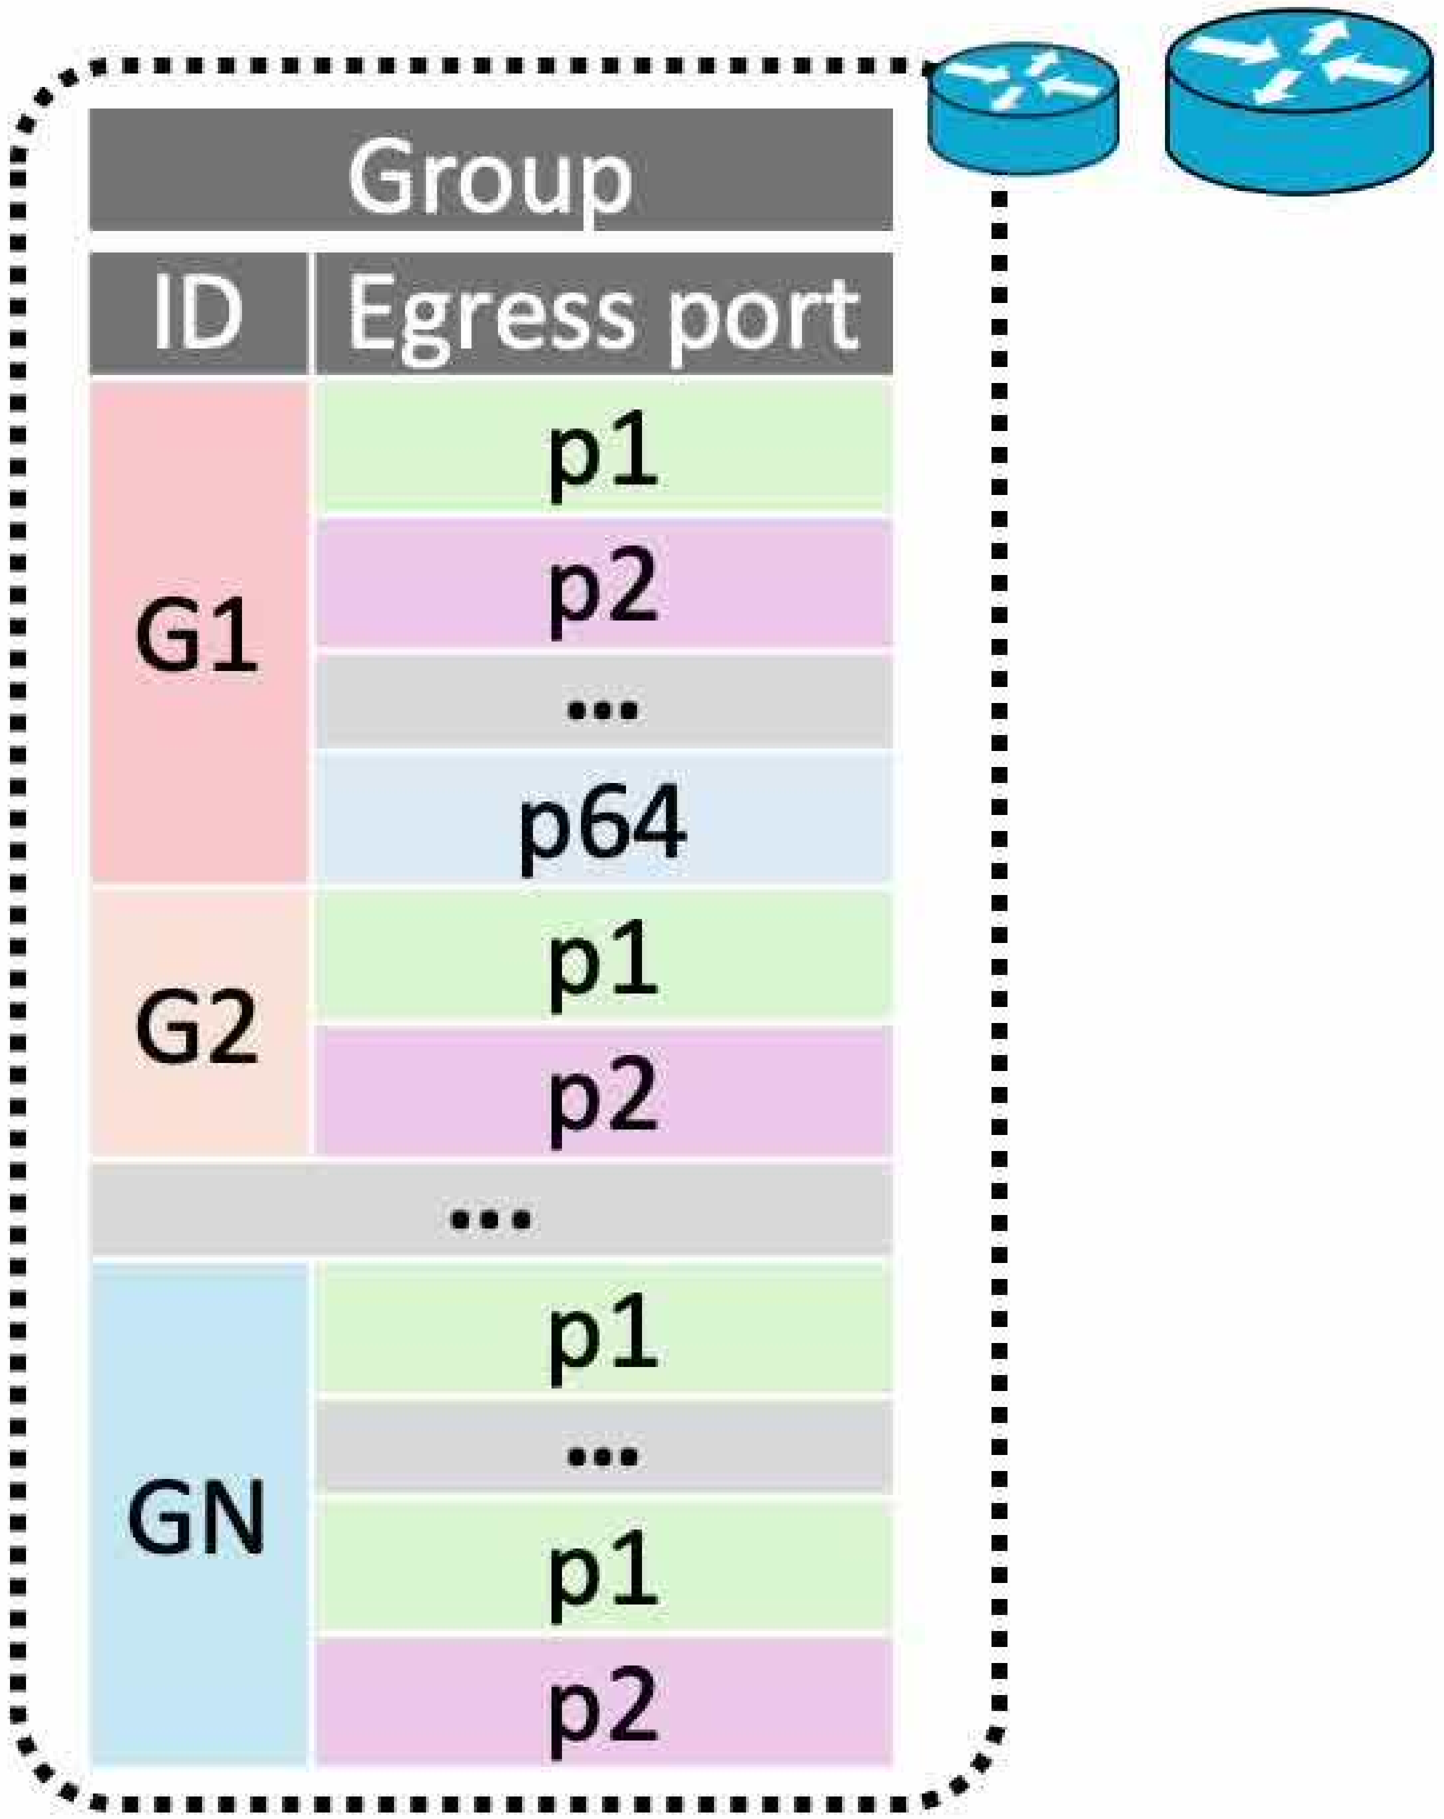

5 root causes that exacerbate precision loss.

- Group space usage depends on (# groups, # ports/group, port weights).
- Cause 1: **scale**
  - # groups scales with network size.
- Cause 2: **path diversity**
  - TE uses many ports/group.
- Cause 3: **skewed weights**
  - Skewed weights are hard to reduce.
- Cause 4: **heterogeneity**
  - Old gen. switches have small memory space.
  - Mixed speed leads to larger weights.
- Cause 5: **cascading effect**
  - Precision loss multiplies in multi-tier networks.

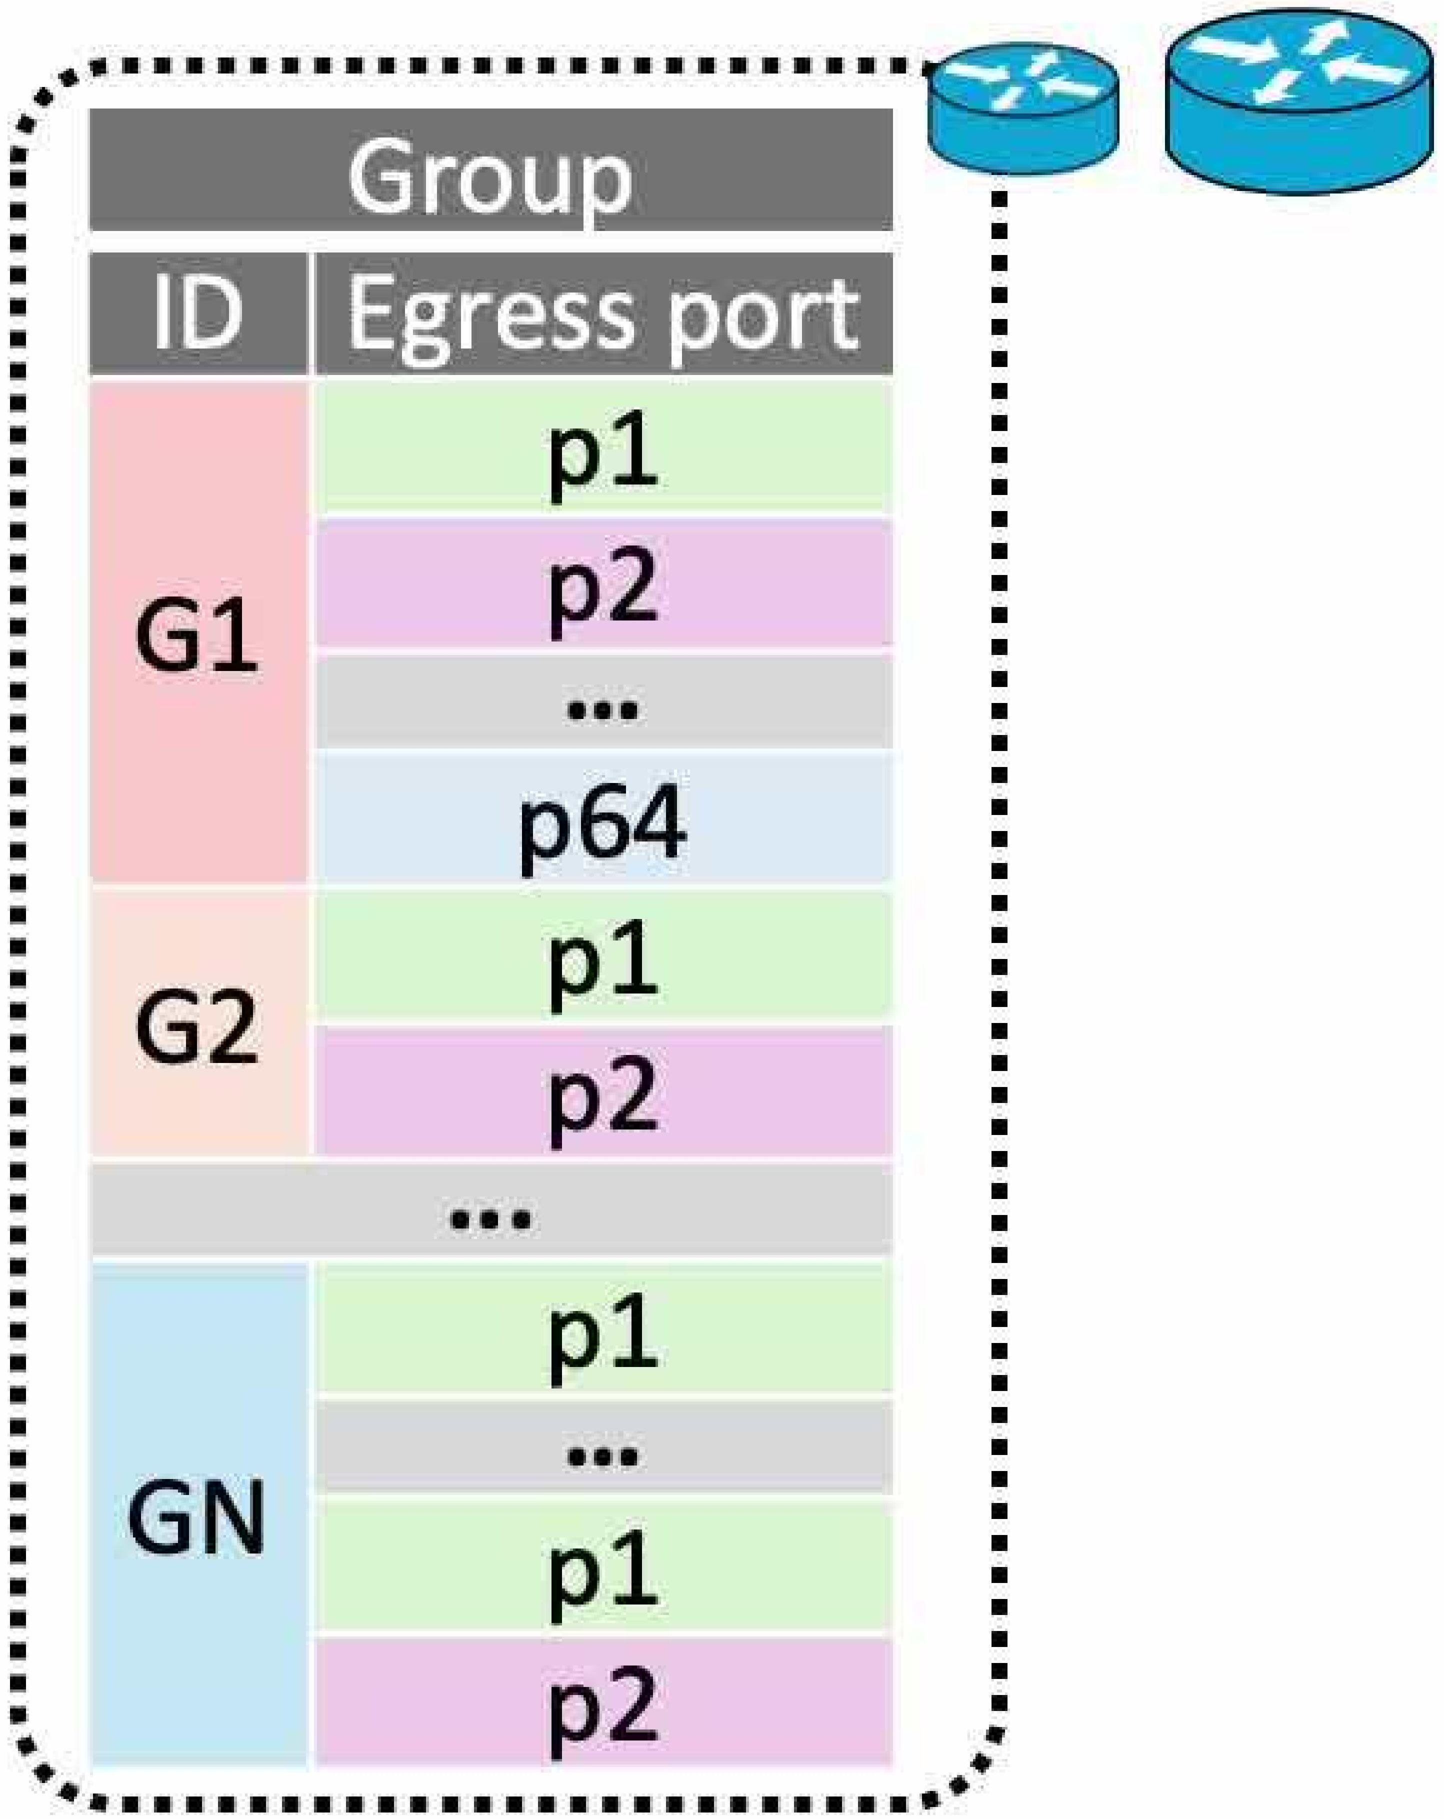

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities

Groups can become identical post reduction.

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities

Groups can become identical post reduction.

Different groups contribute to the overall precision loss differently.

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities	Our heuristics
Groups can become identical post reduction.	**Group Sharing**
Different groups contribute to the overall precision loss differently.	
Not all ports in a group need to be preserved.	

Carnegie Mellon University

3 new heuristics for group reduction algorithms

Group Sharing: de-duplicate & reuse identical groups.

Opportunities	Our heuristics
Groups can become identical post reduction.	**Group Sharing**
Different groups contribute to the overall precision loss differently.	
Not all ports in a group need to be preserved.	

3 new heuristics for group reduction algorithms

Table Carving: allocate space to each group proportional to its traffic volume.

Opportunities	**Our heuristics**
Groups can become identical post reduction.	**Group Sharing**
Different groups contribute to the overall precision loss differently.	**Table Carving**
Not all ports in a group need to be preserved.	

3 new heuristics for group reduction algorithms

Group Pruning: prune select ports from a group to enable size reduction.

Opportunities	**Our heuristics**
Groups can become identical post reduction.	**Group Sharing**
Different groups contribute to the overall precision loss differently.	**Table Carving**
Not all ports in a group need to be preserved.	**Group Pruning**

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities	**Our heuristics**	Root causes
Groups can become identical post reduction.	**Group Sharing**	scale heterogeneity
Different groups contribute to the overall precision loss differently.	Table Carving	
Not all ports in a group need to be preserved.	Group Pruning	

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities	**Our heuristics**	Root causes
Groups can become identical post reduction.	Group Sharing	scale
		heterogeneity
Different groups contribute to the overall precision loss differently.	**Table Carving**	path diversity
Not all ports in a group need to be preserved.	**Group Pruning**	skewed weights

Carnegie Mellon University

3 new heuristics for group reduction algorithms

- Recap: need to optimize both per-group size & total group size per switch.
- Opportunities lie in where the current approach (TableFitting [EuroSys'14]) falls short.

Opportunities	Our heuristics	Root causes
Groups can become identical post reduction.	**Group Sharing**	scale
Different groups contribute to the overall precision loss differently.	**Table Carving**	heterogeneity
Not all ports in a group need to be preserved.	**Group Pruning**	path diversity
		skewed weights
		cascading effect

Carnegie Mellon University

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)

Carnegie Mellon University

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search

Carnegie Mellon University

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

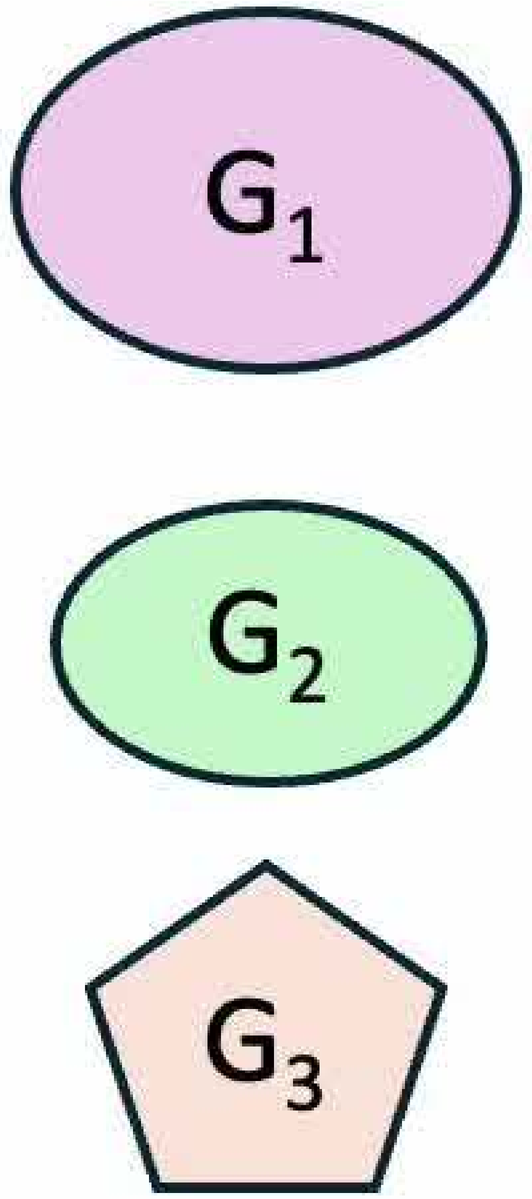

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

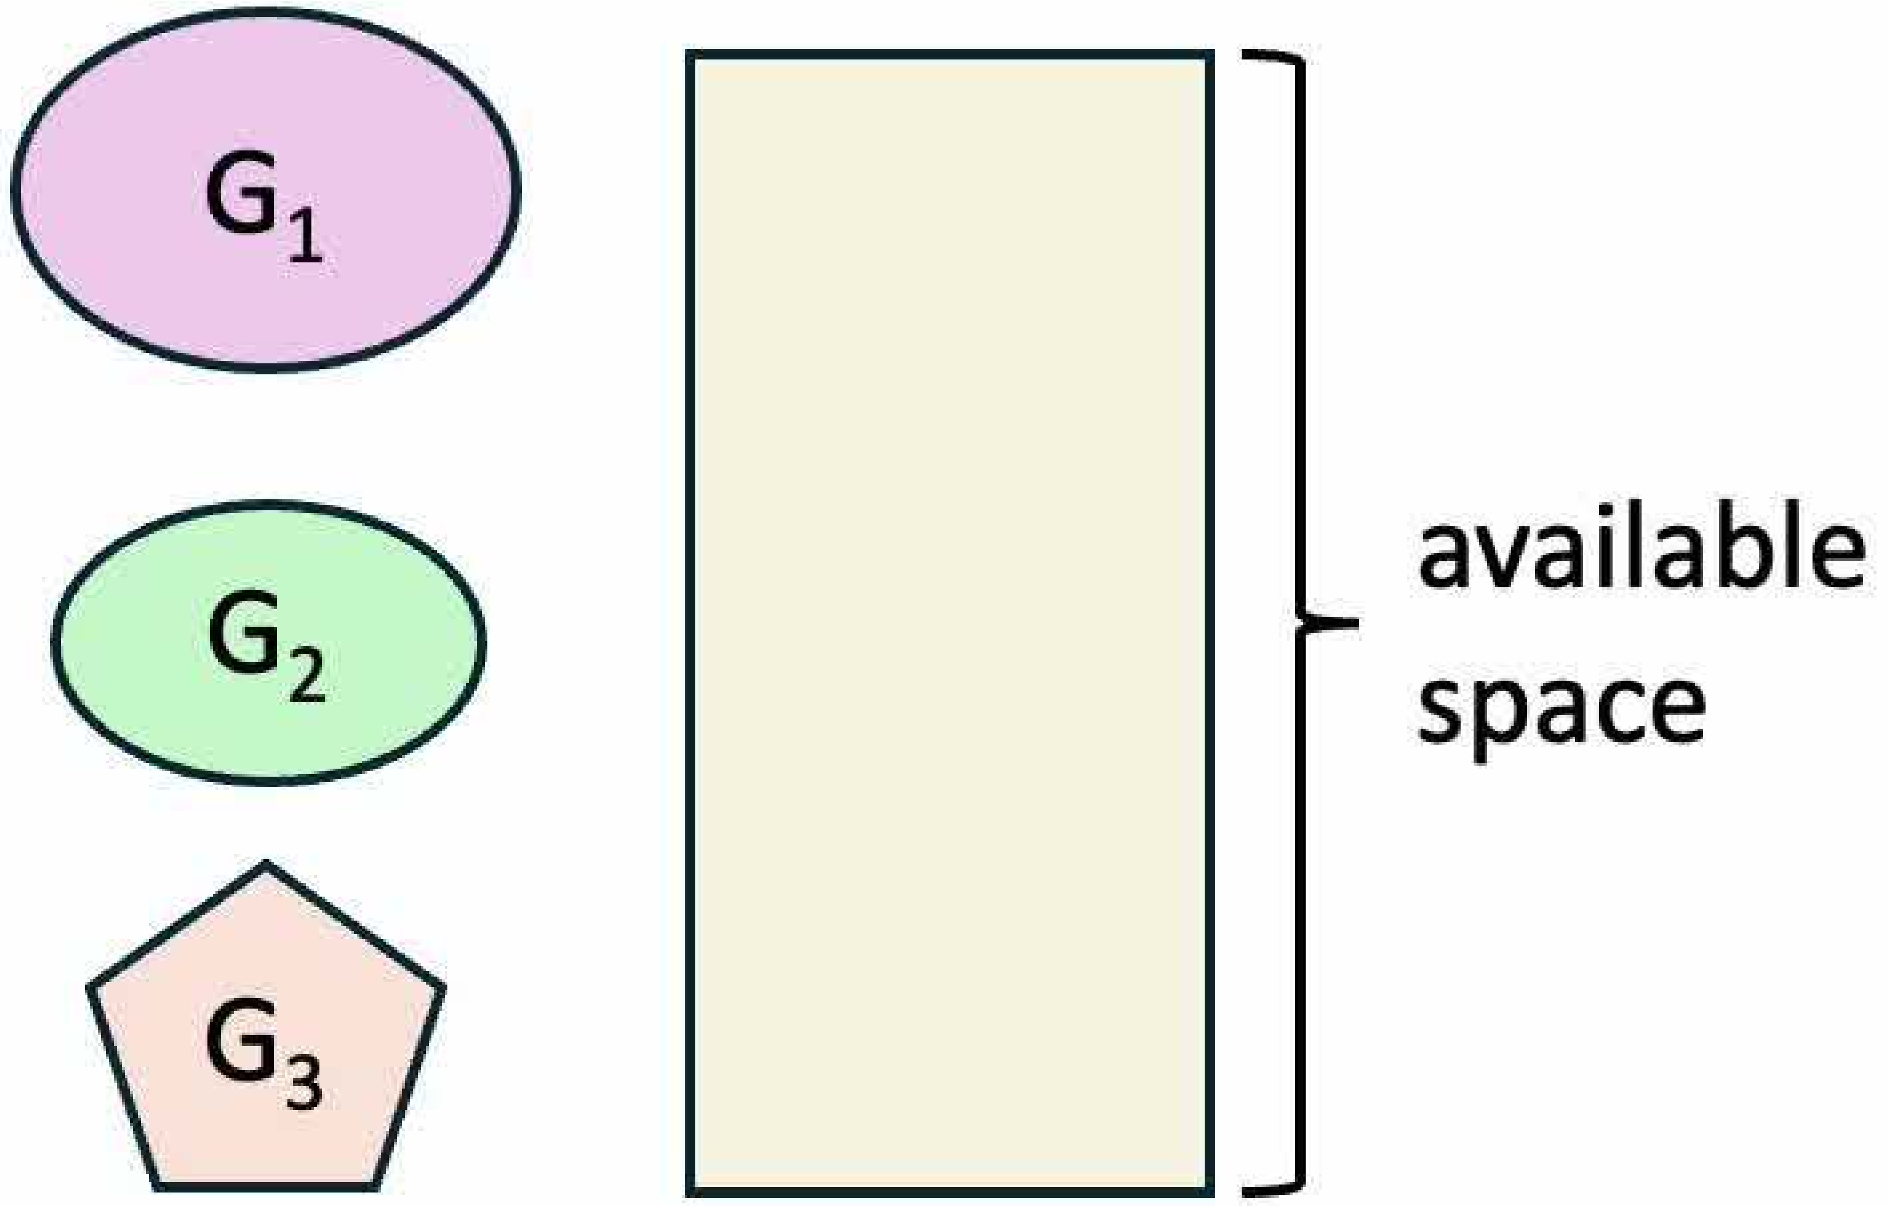

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

Step 1: Table Carving

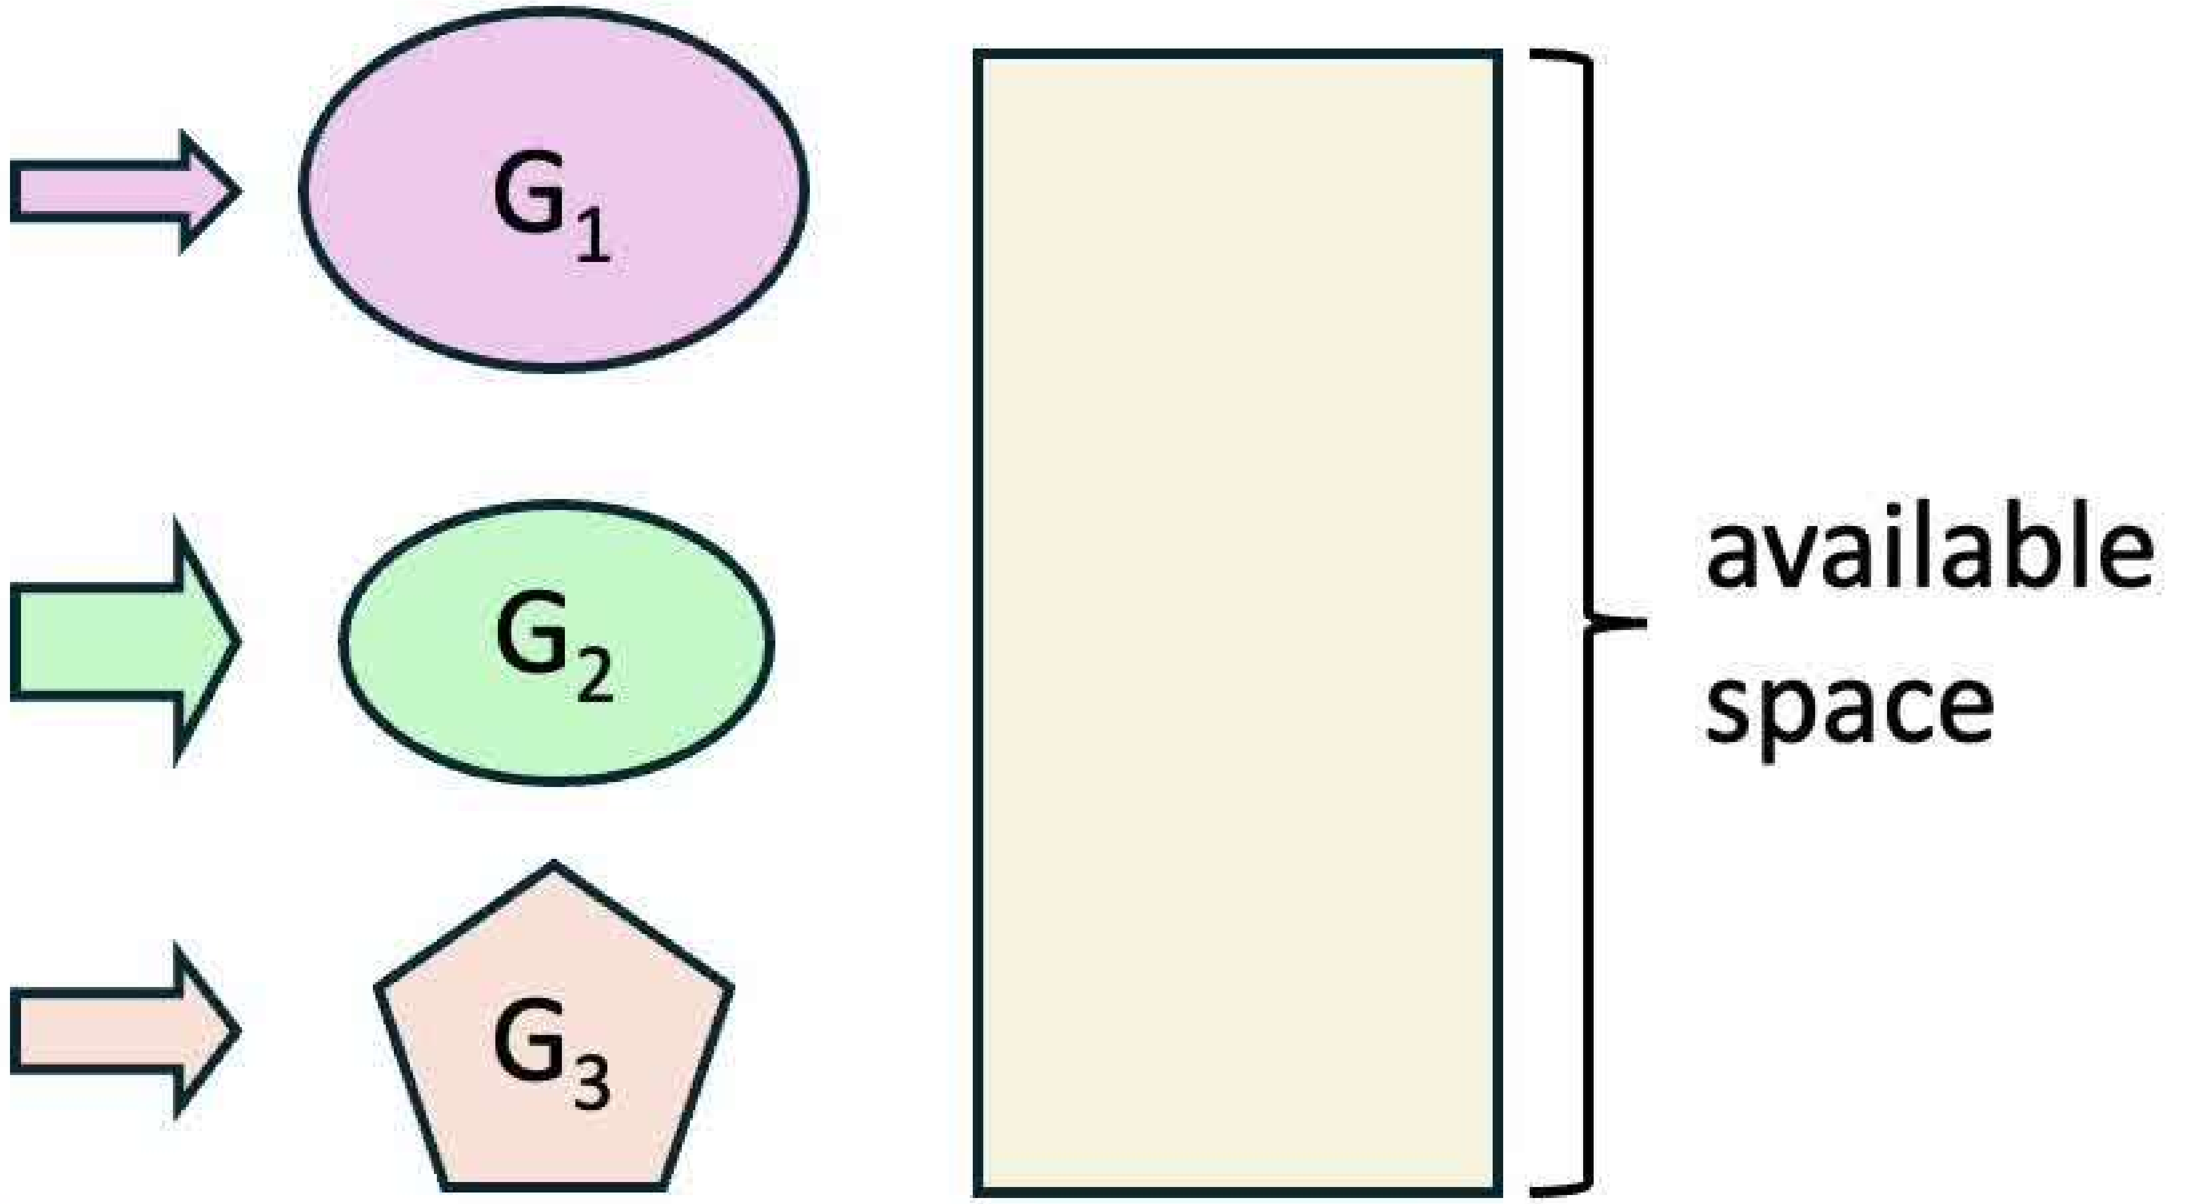

Carnegie Mellon University

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

Step 1: Table Carving

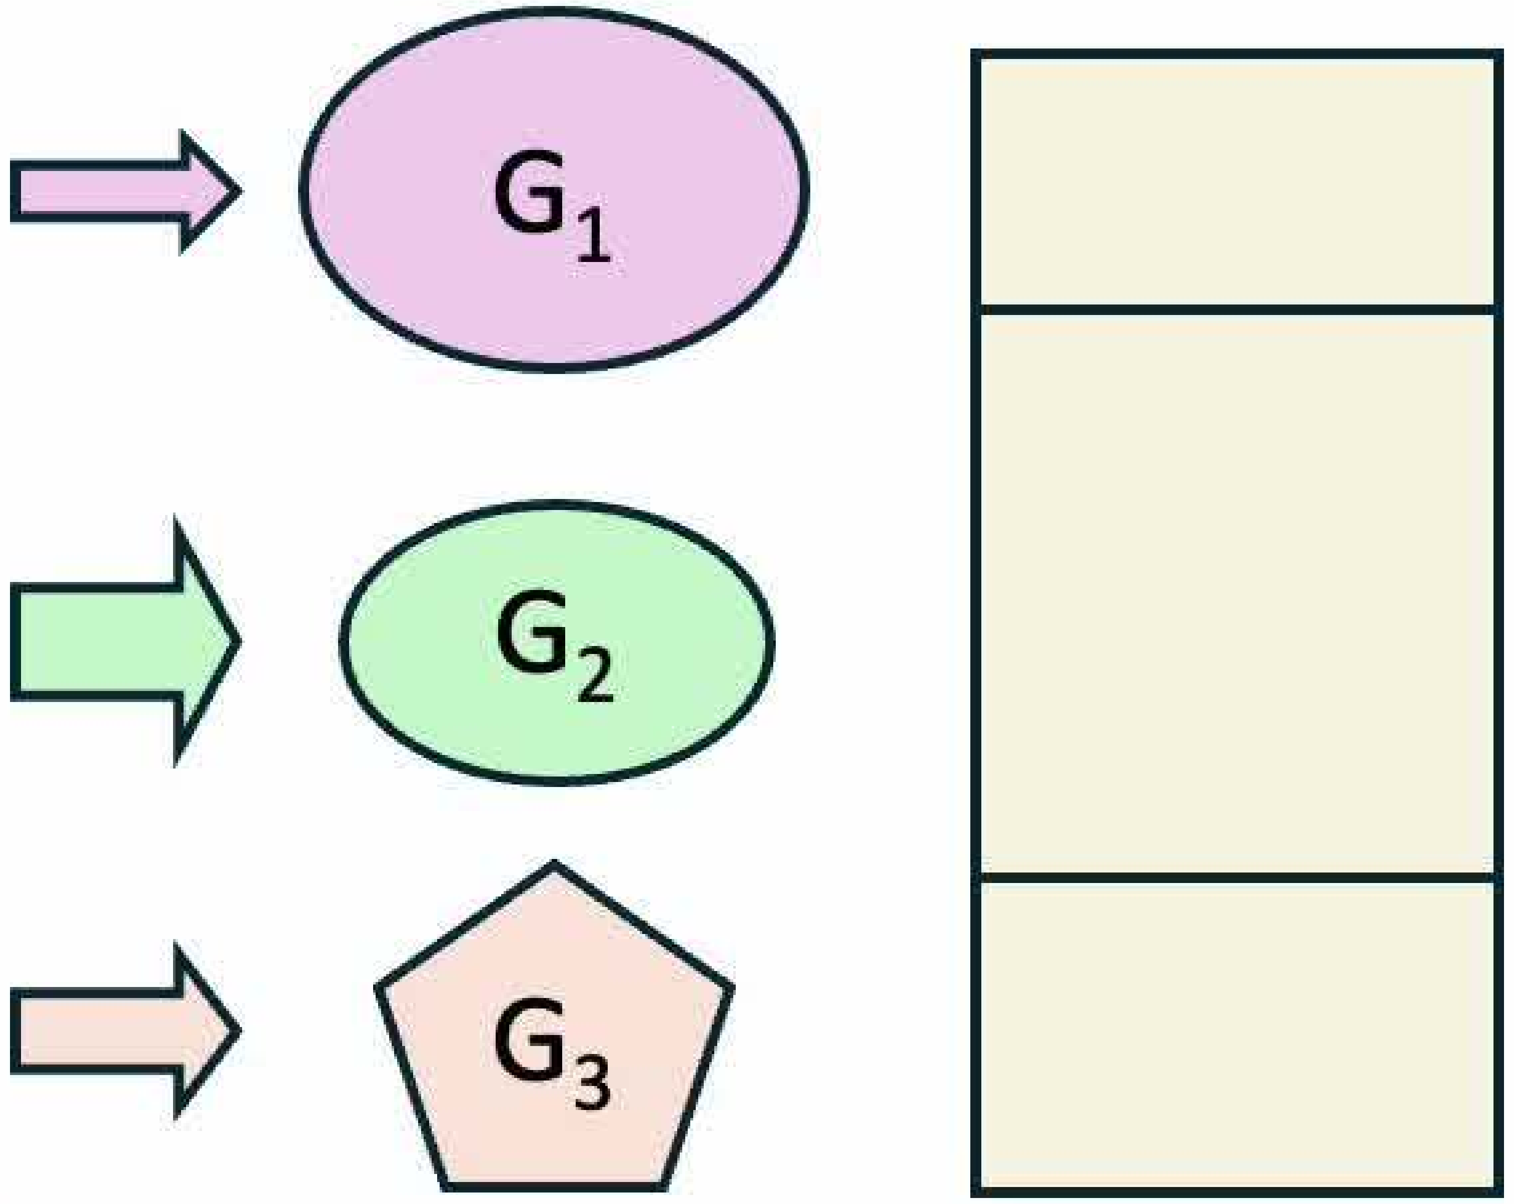

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

Step 2: parallel single-group reduction
(MIP or greedy)
invokes Group Pruning

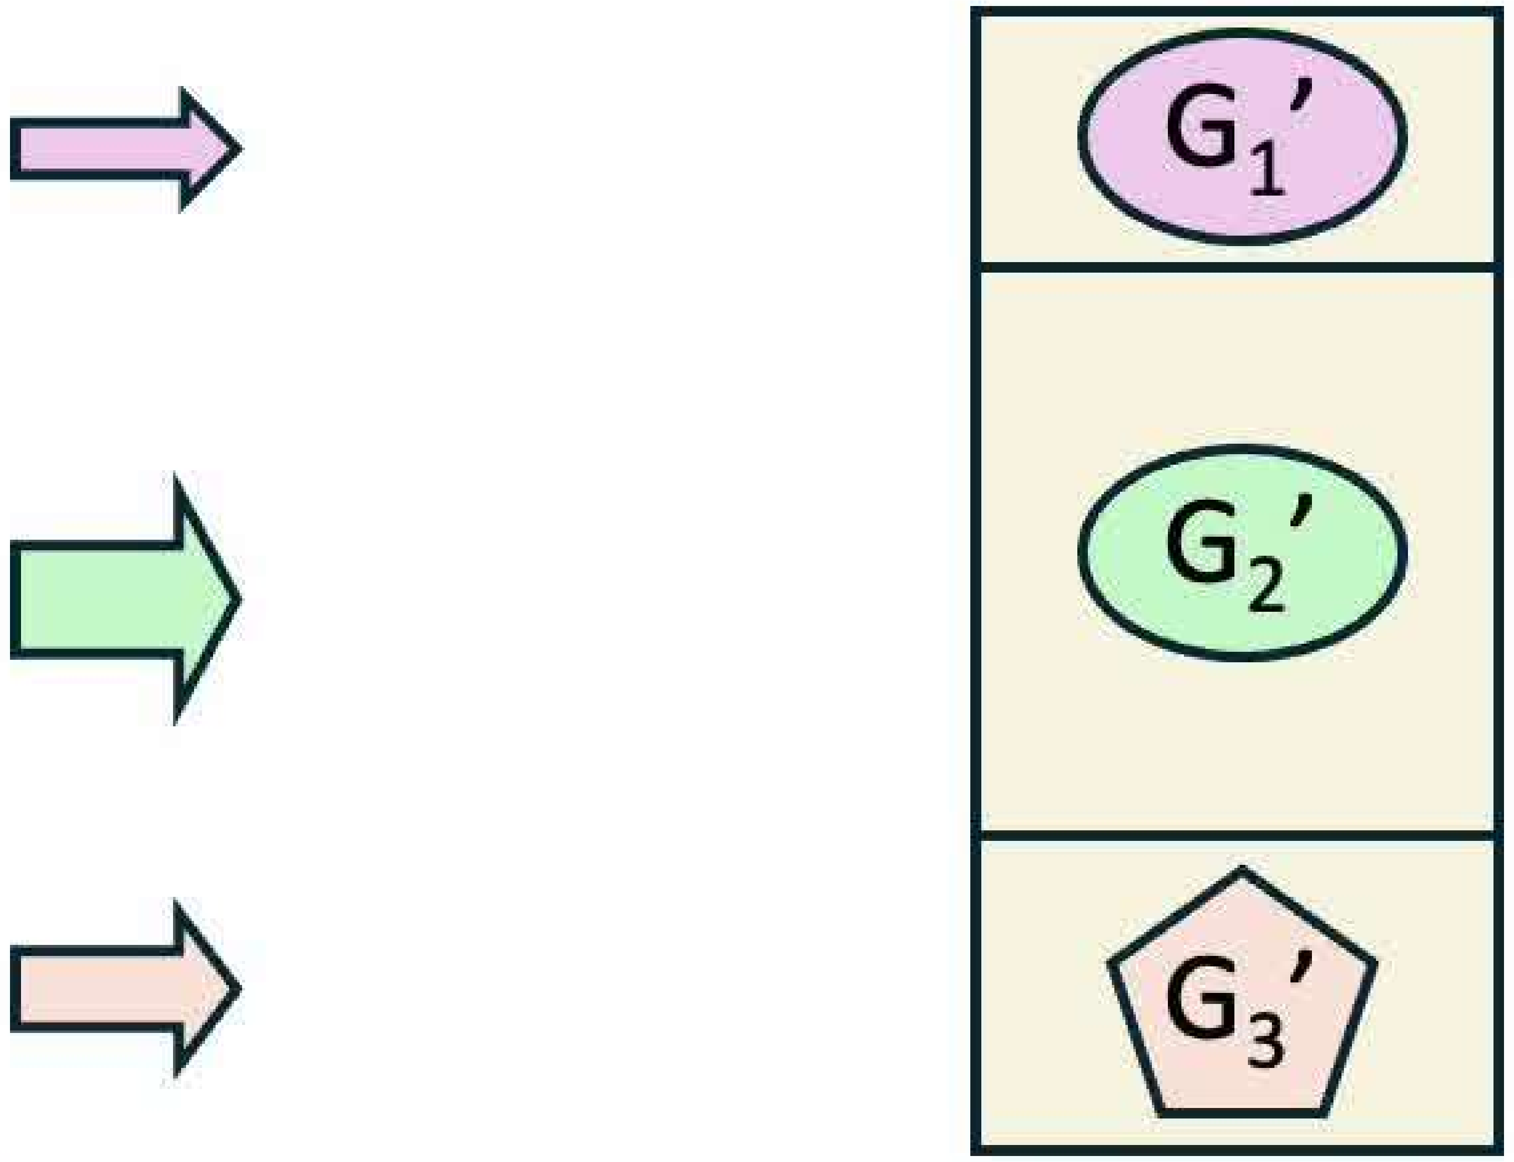

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

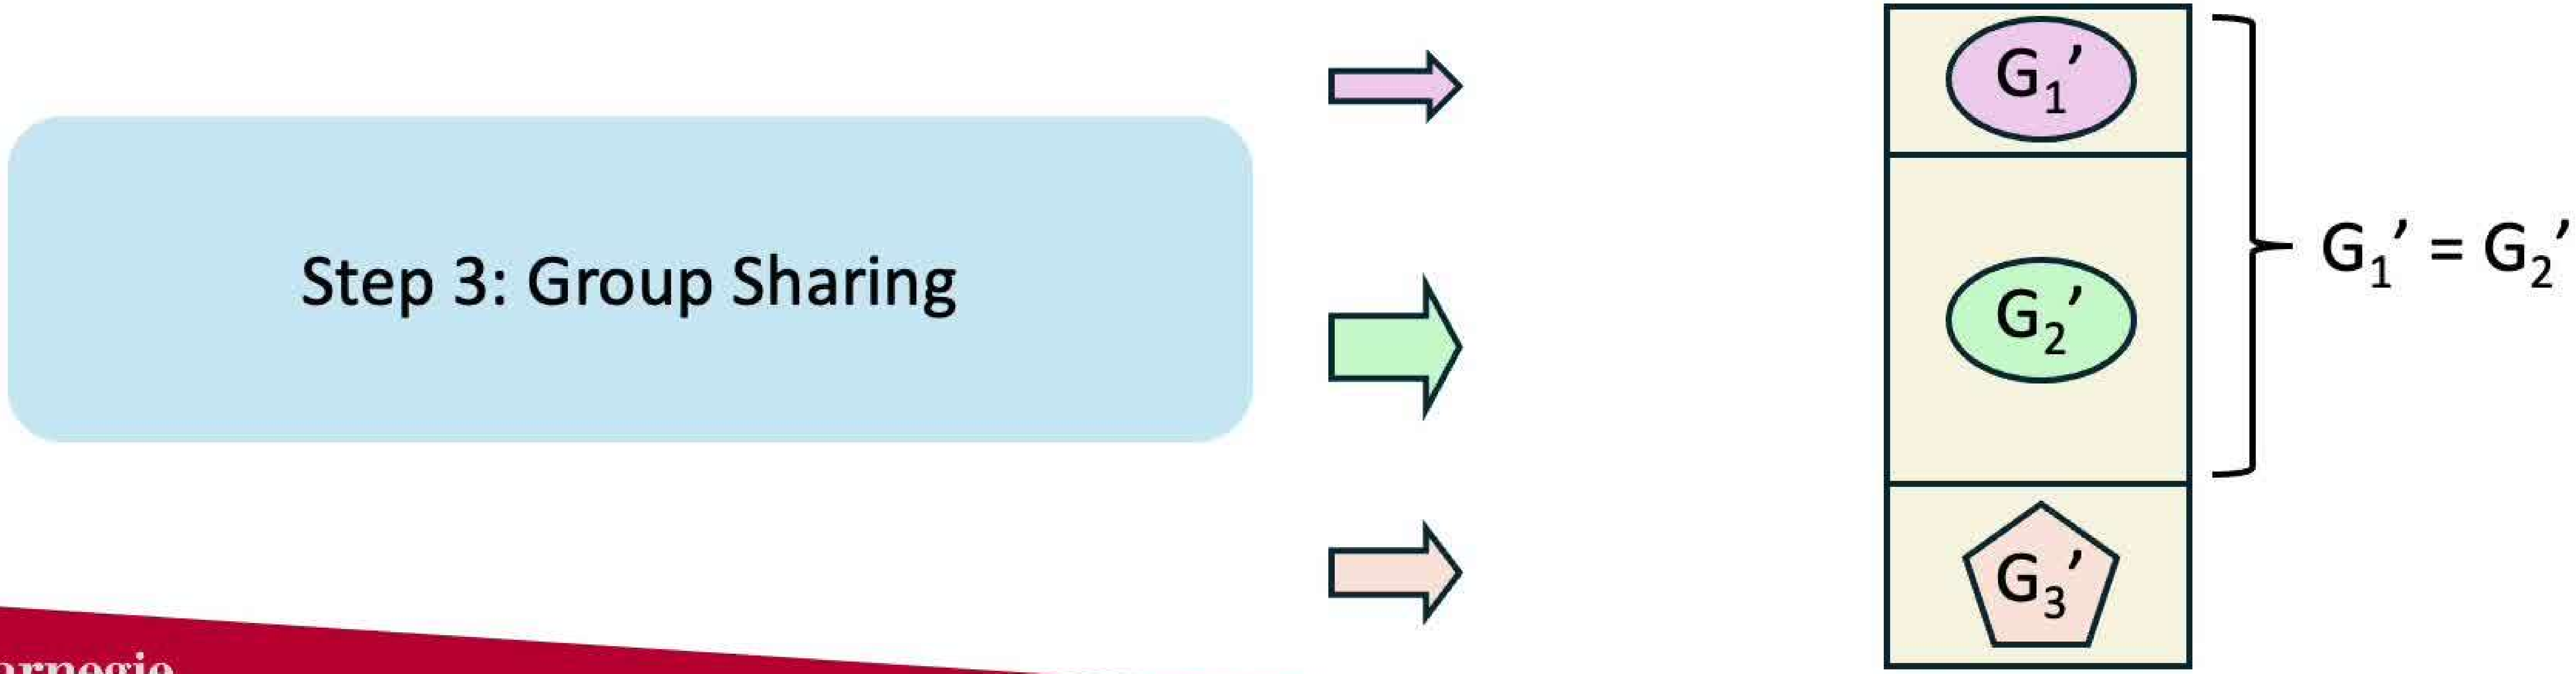

Carnegie Mellon University

New group reduction algorithms: DMIR & IGR

	Direct Mixed-Integer Reduction (**DMIR**)	Iterative Greedy Reduction (**IGR**)
Core algorithm	mixed-integer programming	greedy search
Optimality	optimal	less optimal
Execution speed	slow	fast

Step 3: Group Sharing

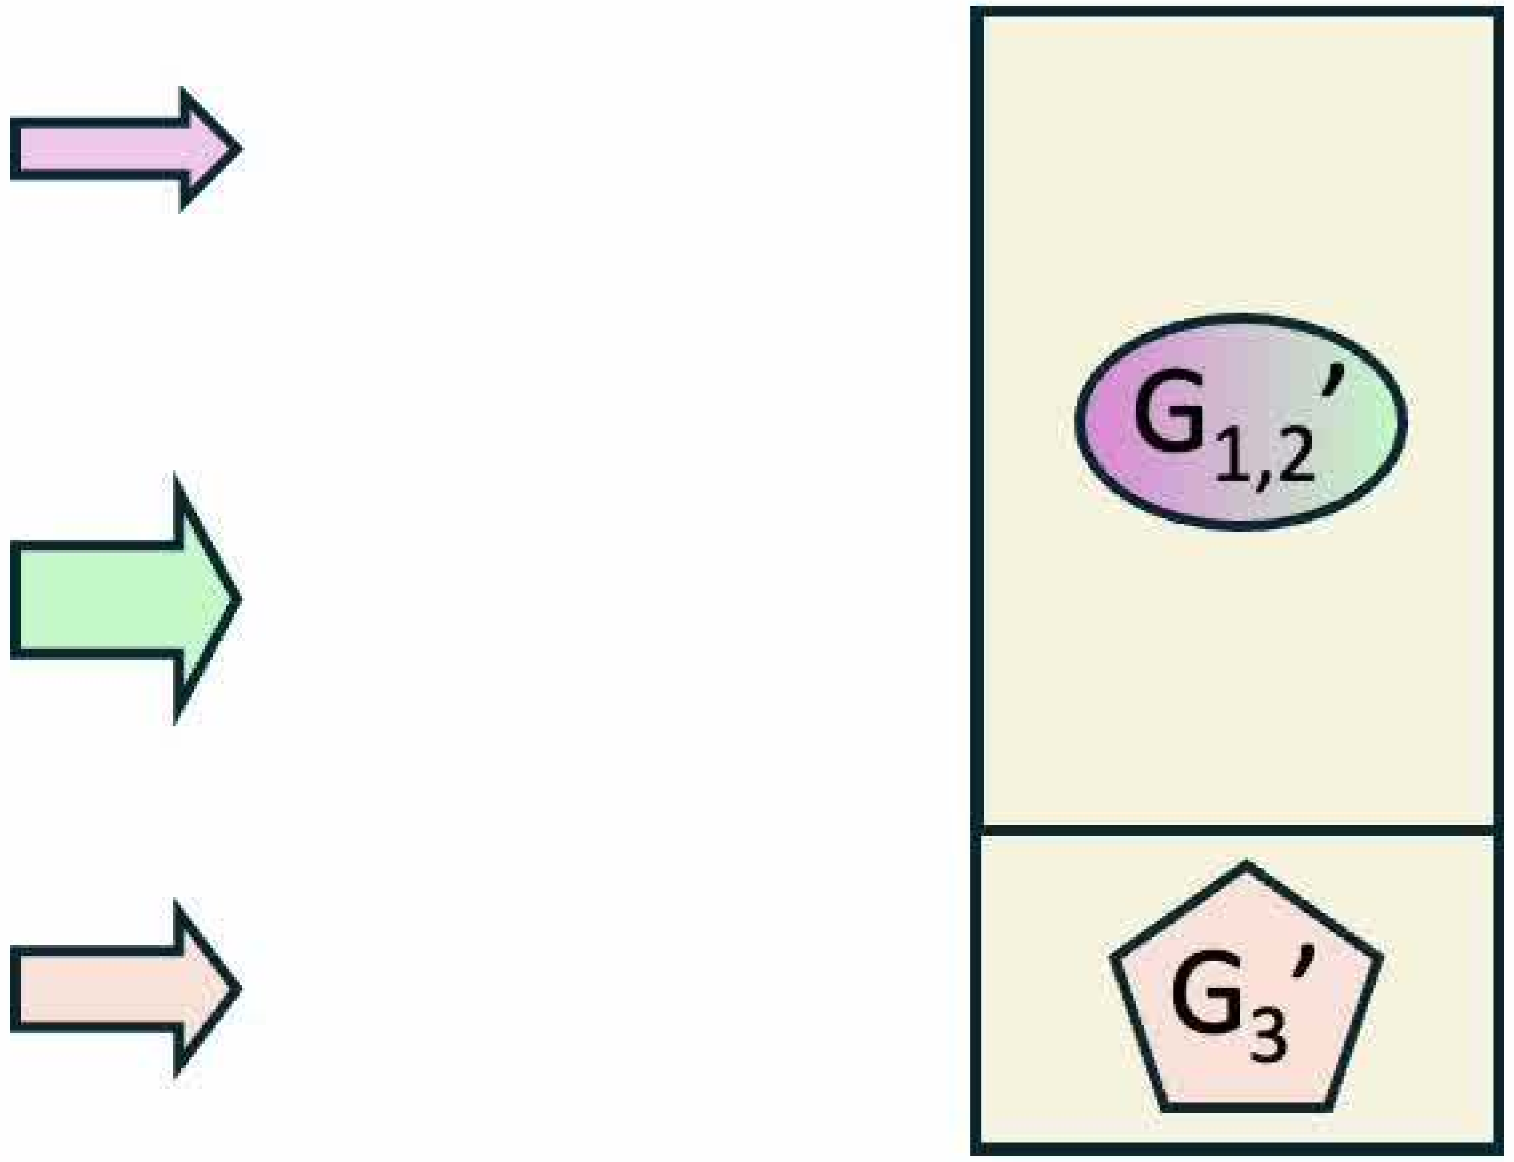

DMIR & IGR are more precise than current work.

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization

- Application-level metric: flow completion time (FCT)

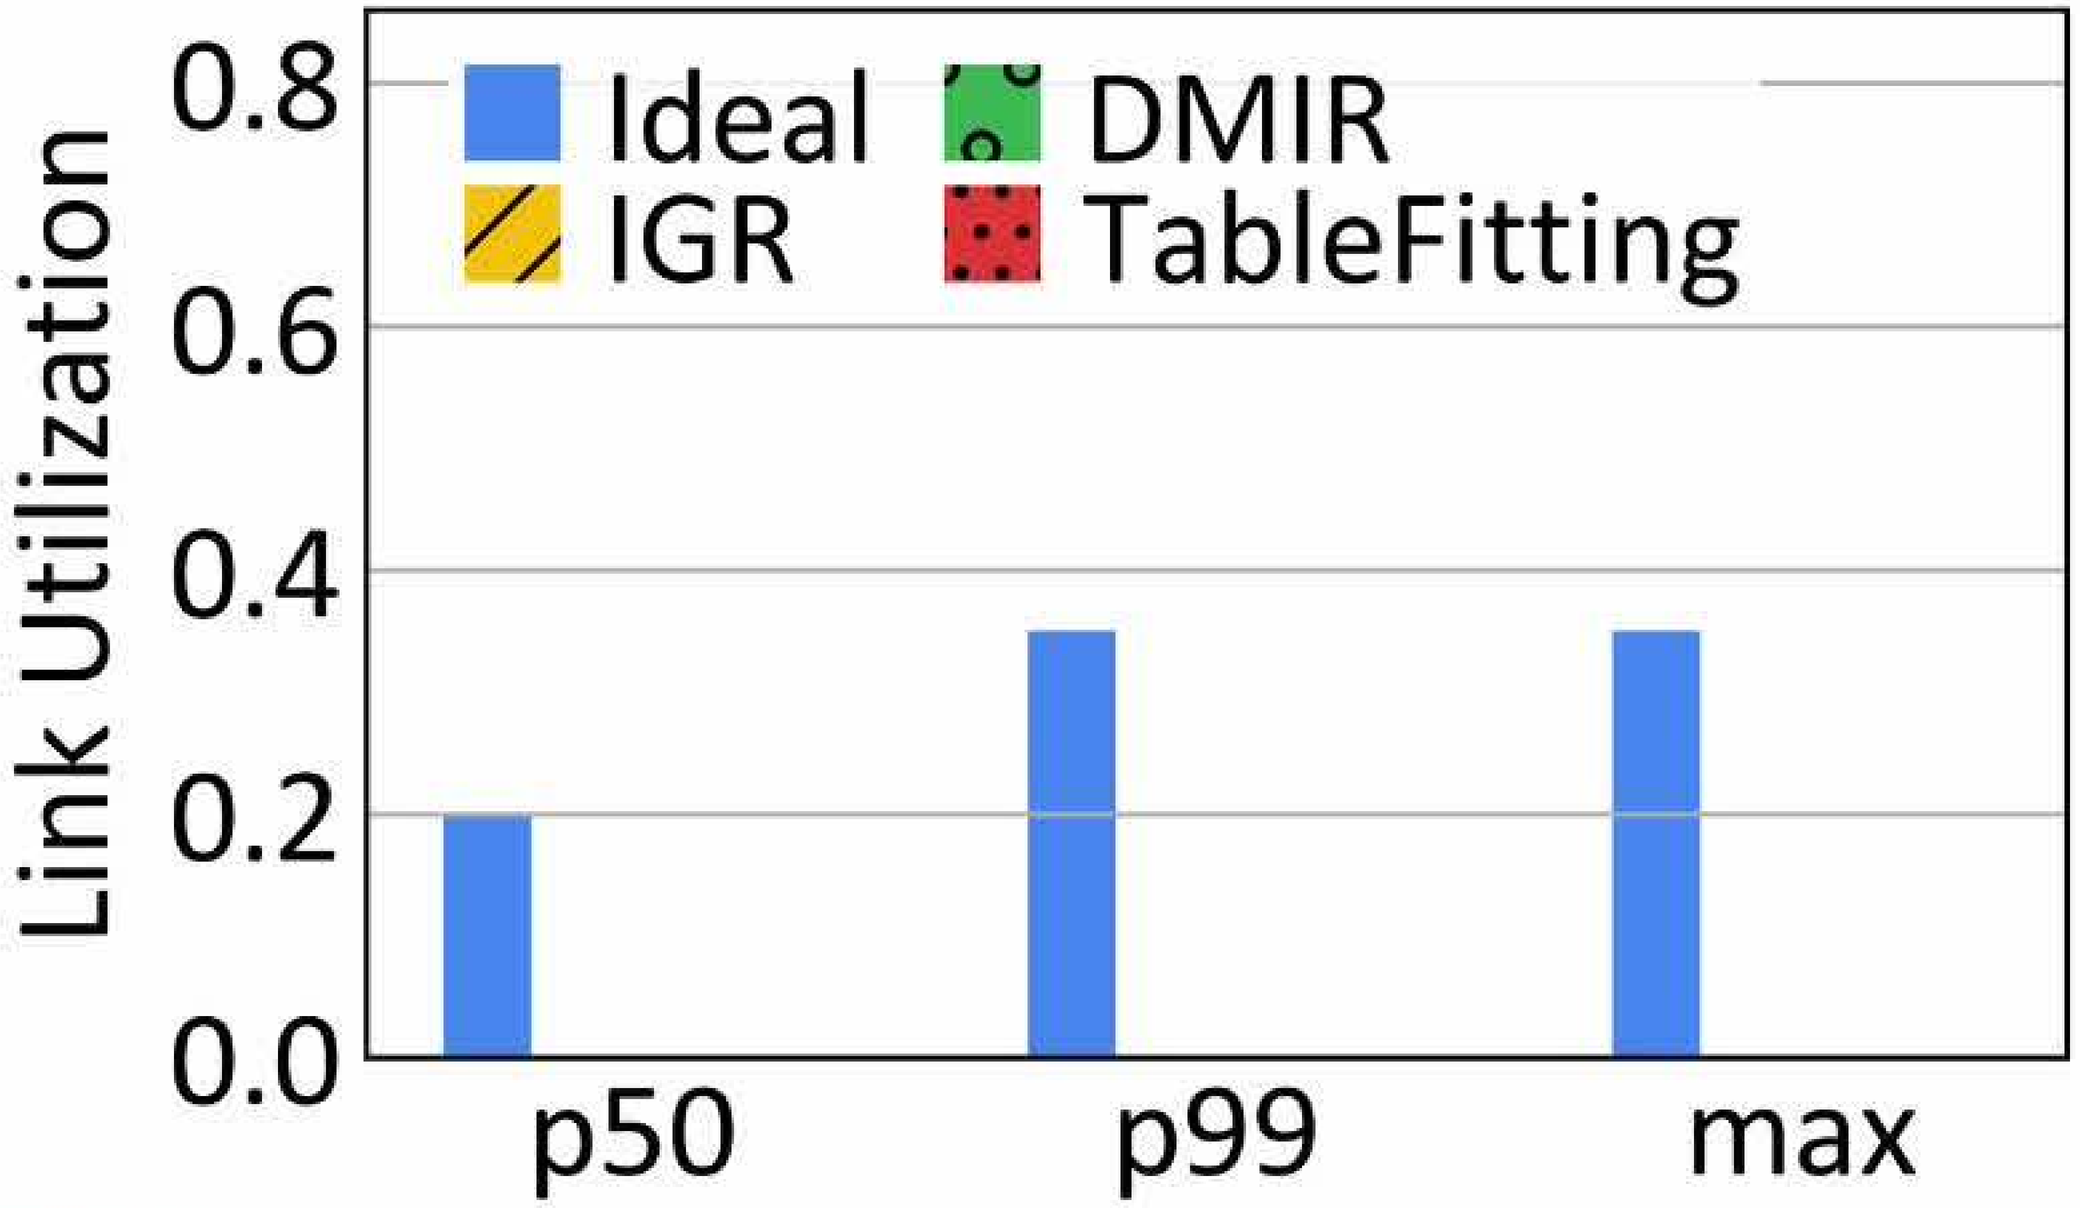

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
- Application-level metric: flow completion time (FCT)

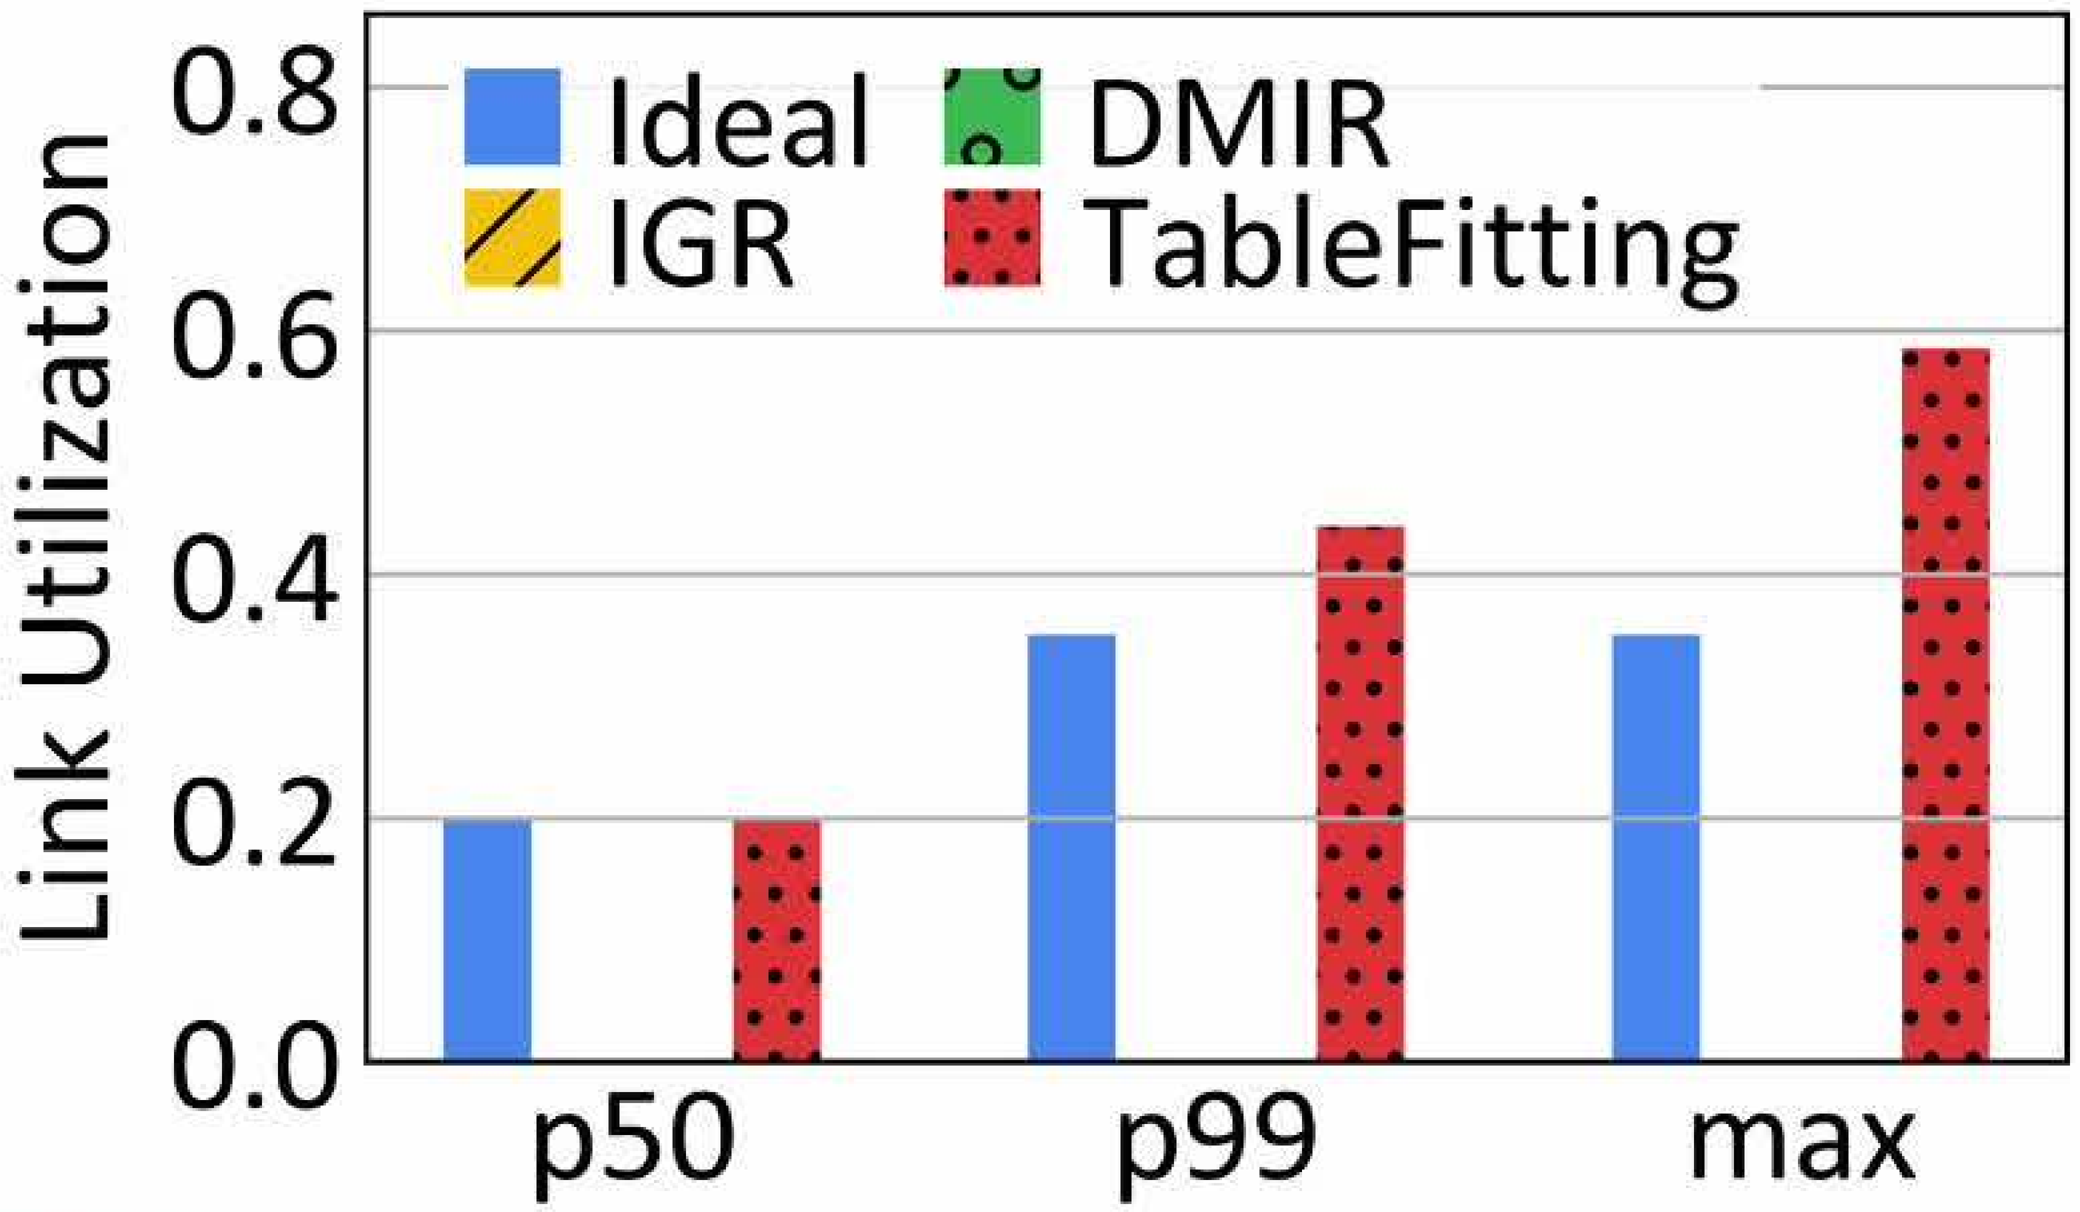

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization

- Application-level metric: flow completion time (FCT)

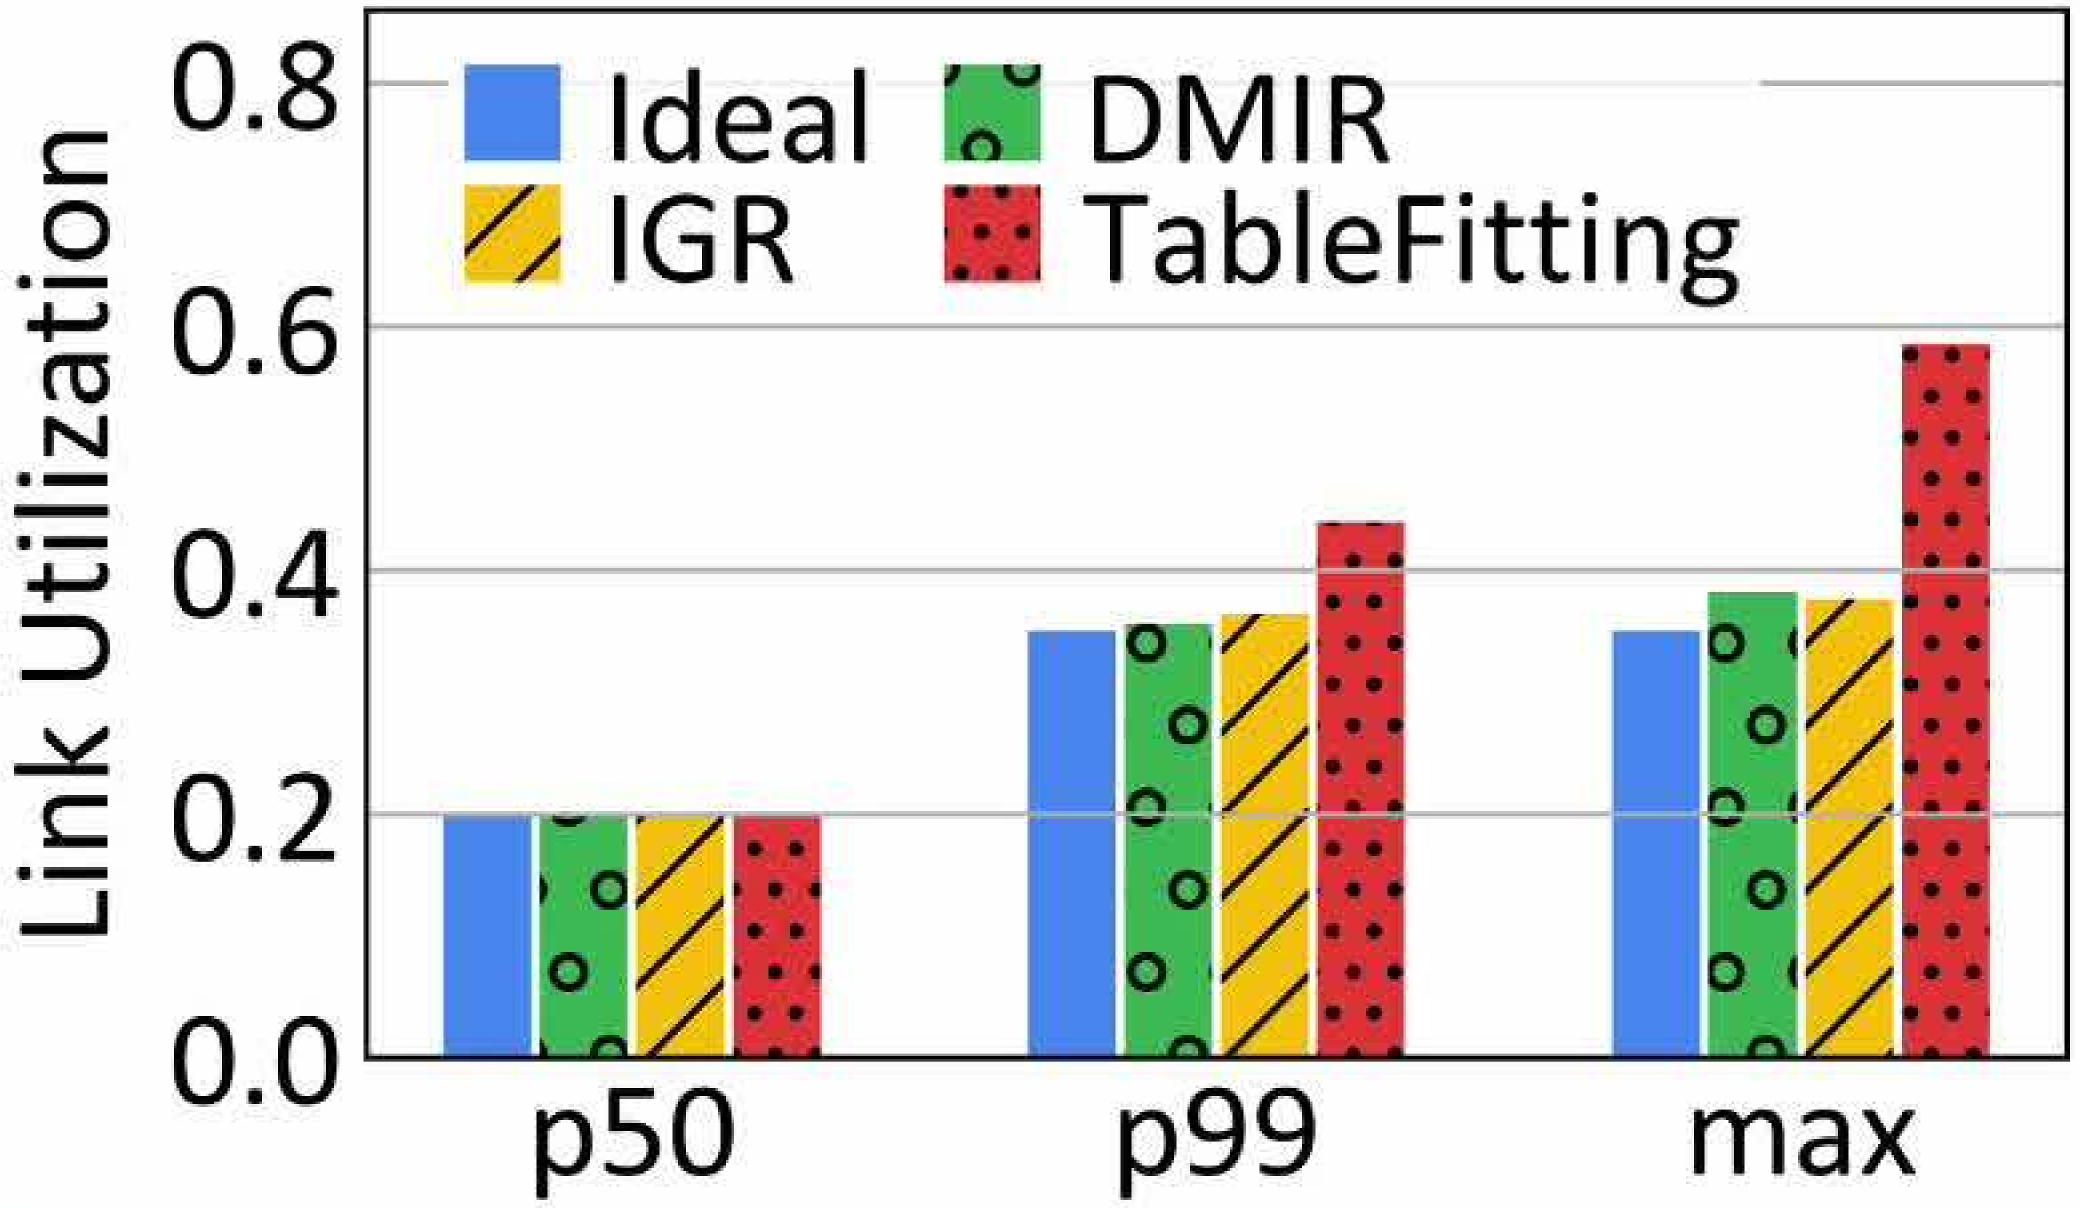

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)

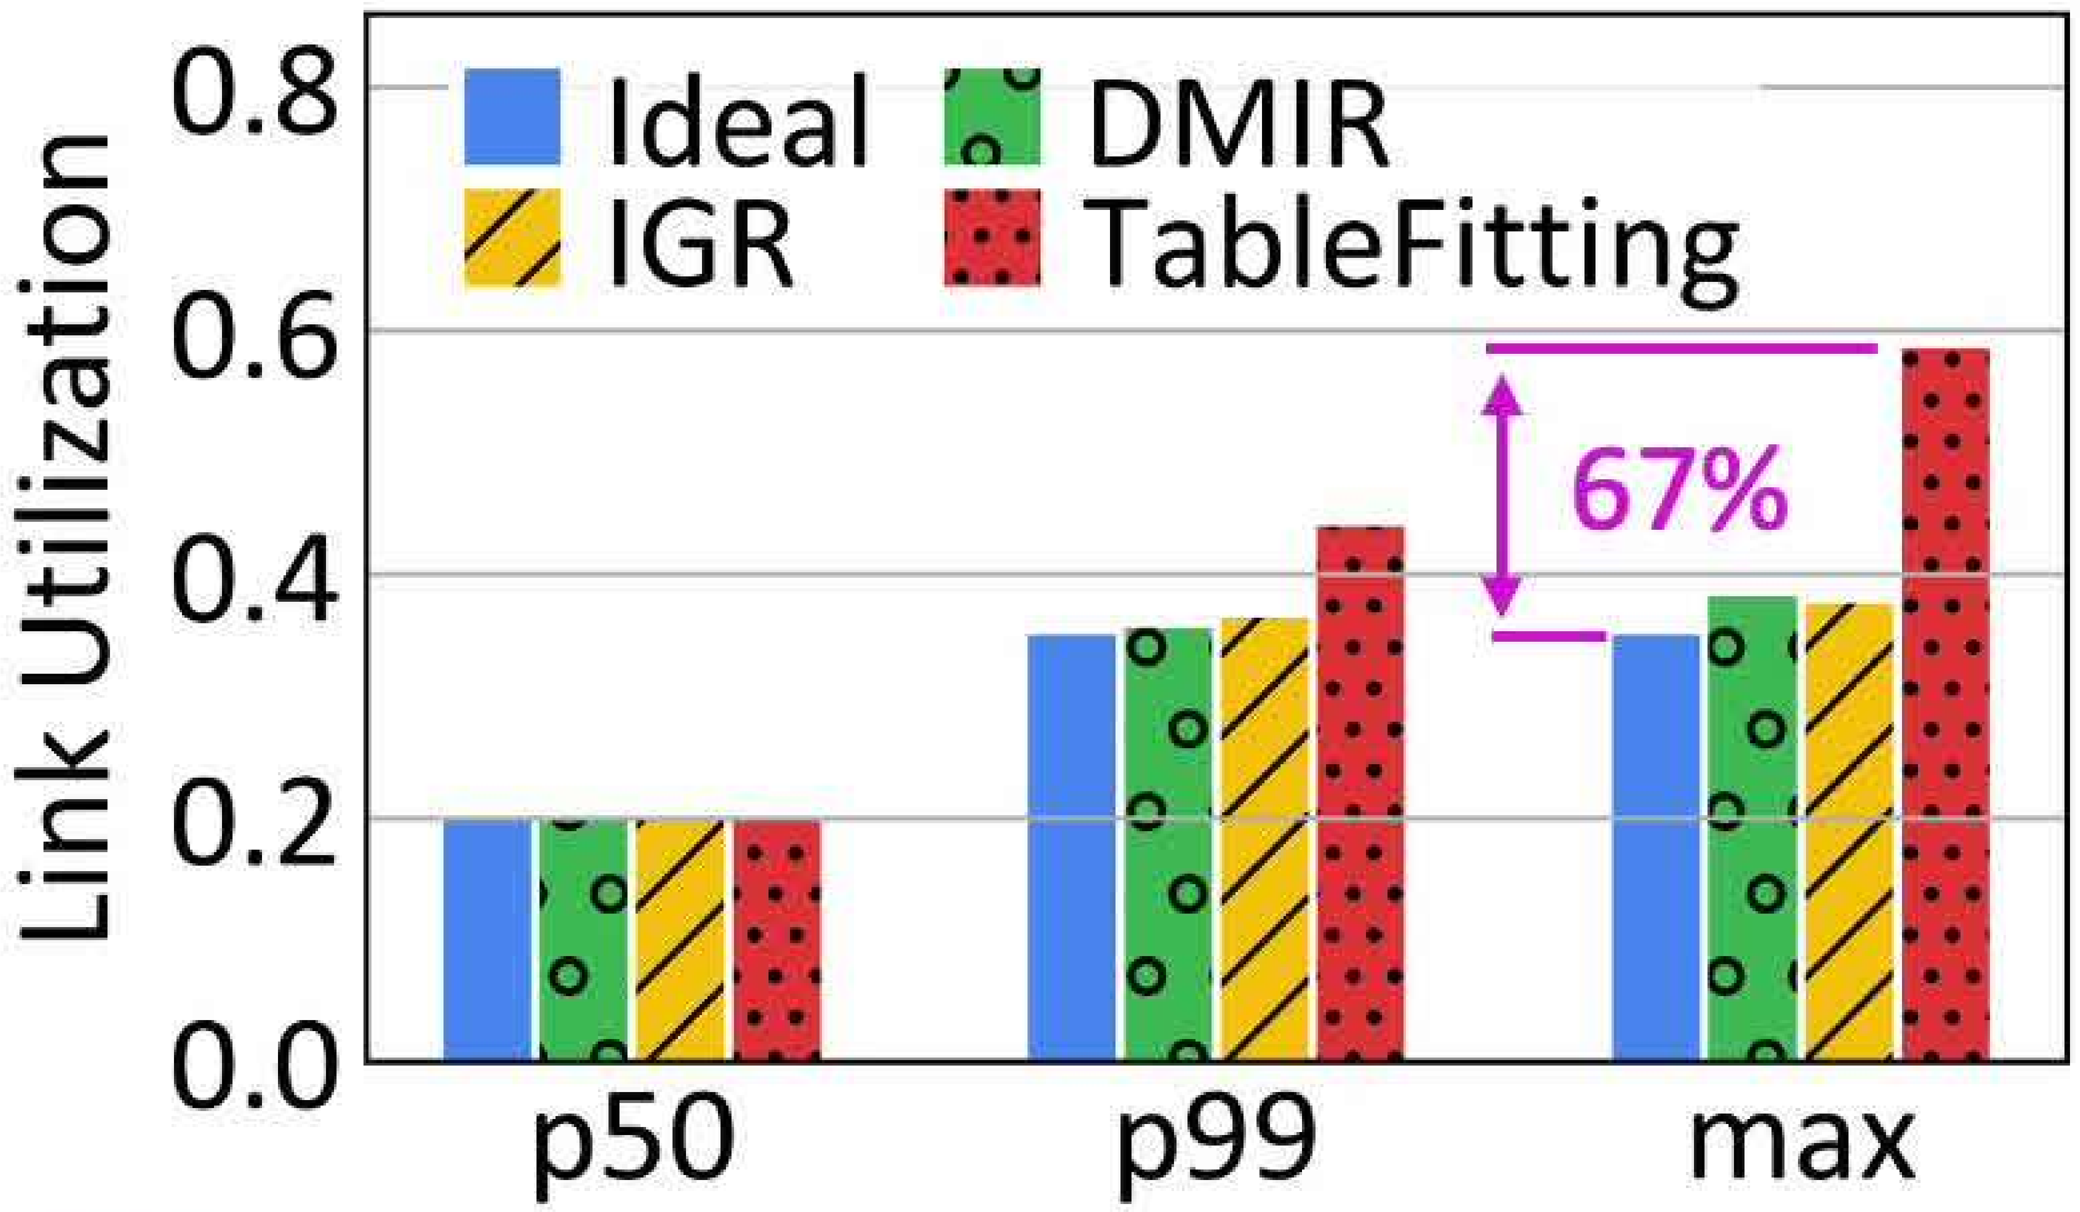

Carnegie Mellon University

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)

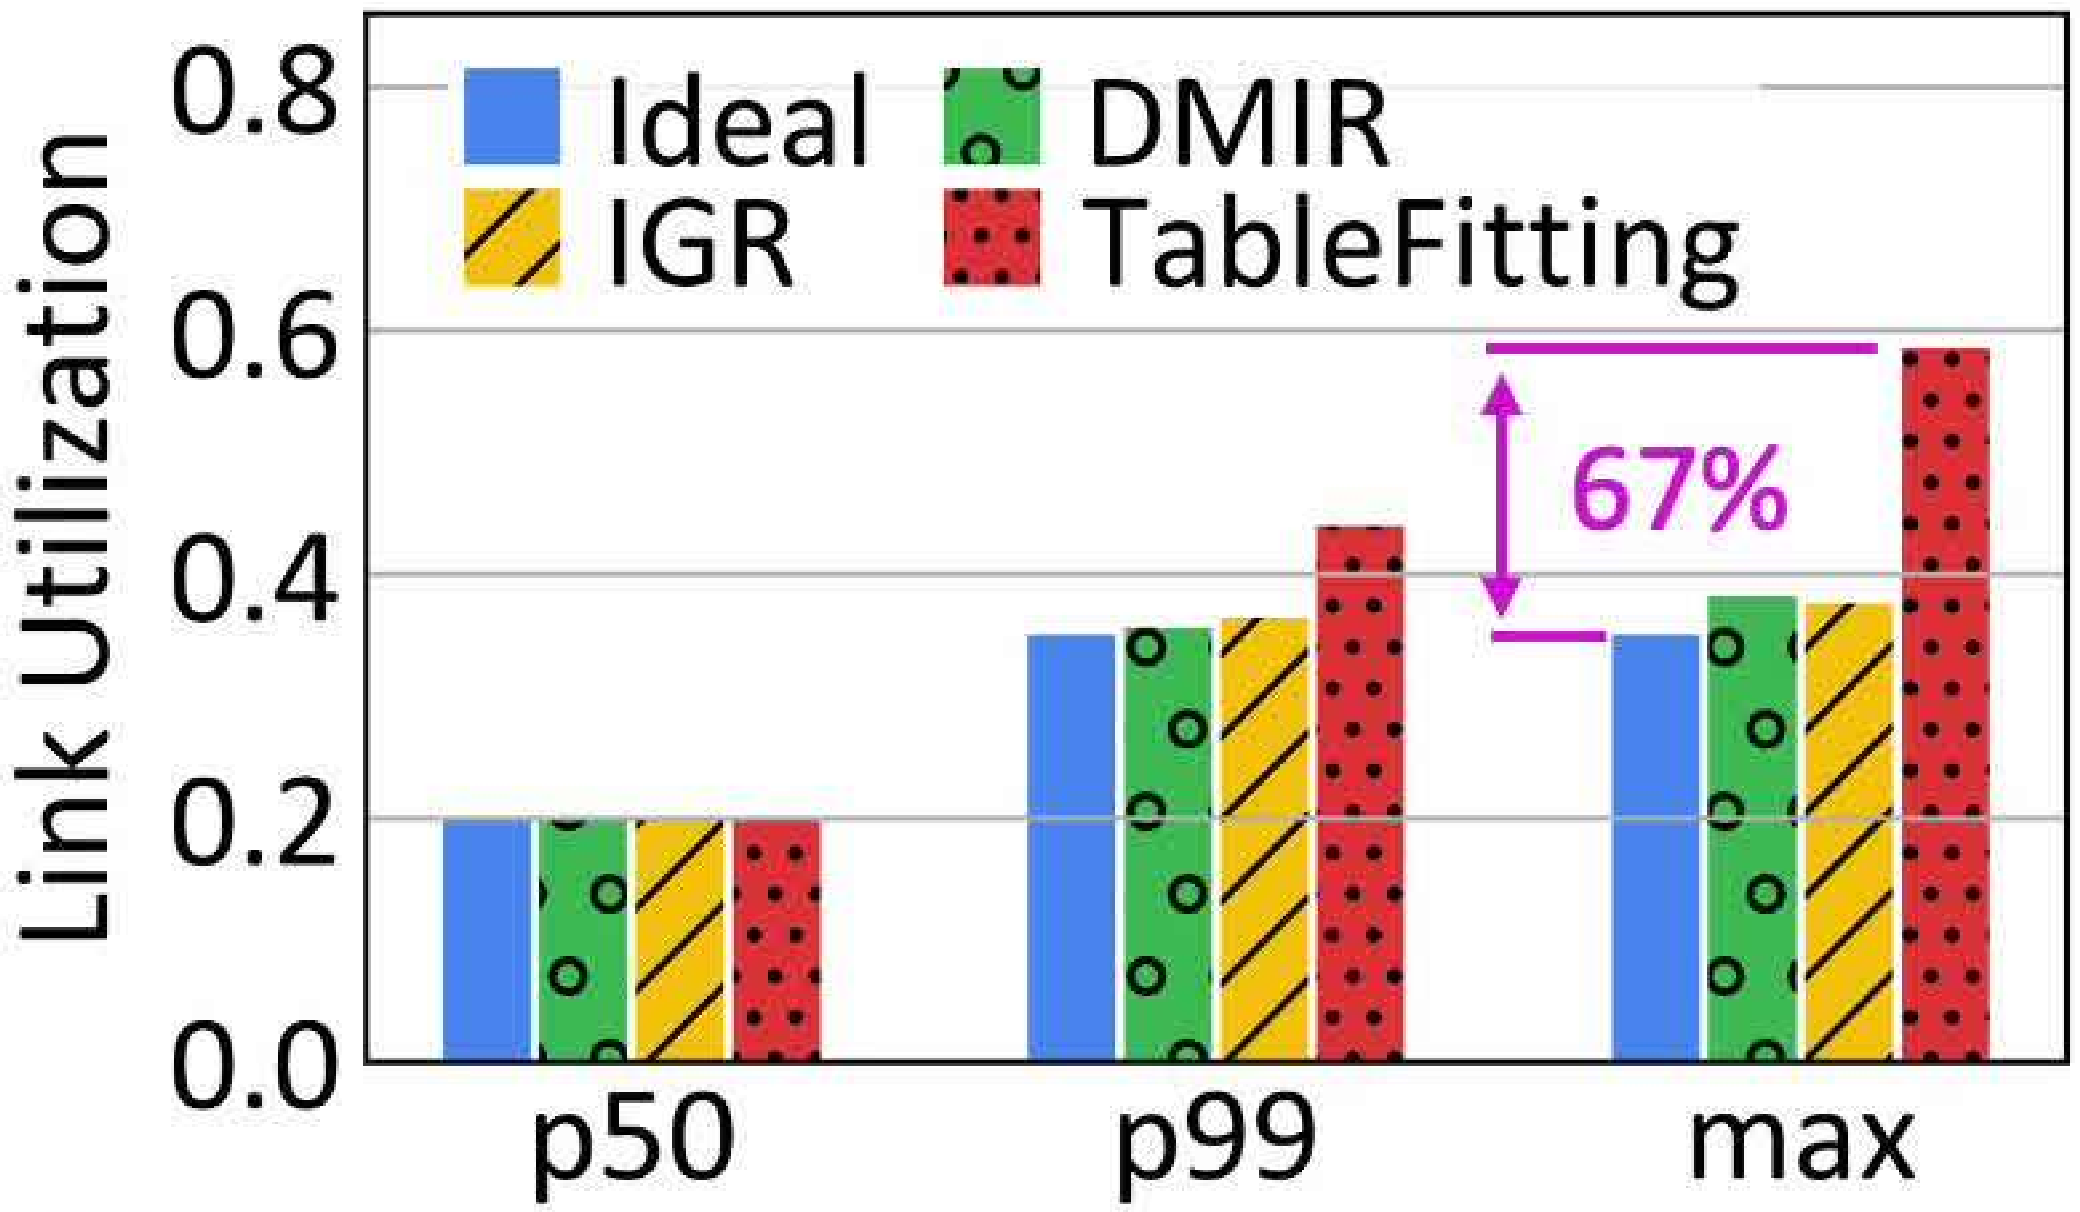

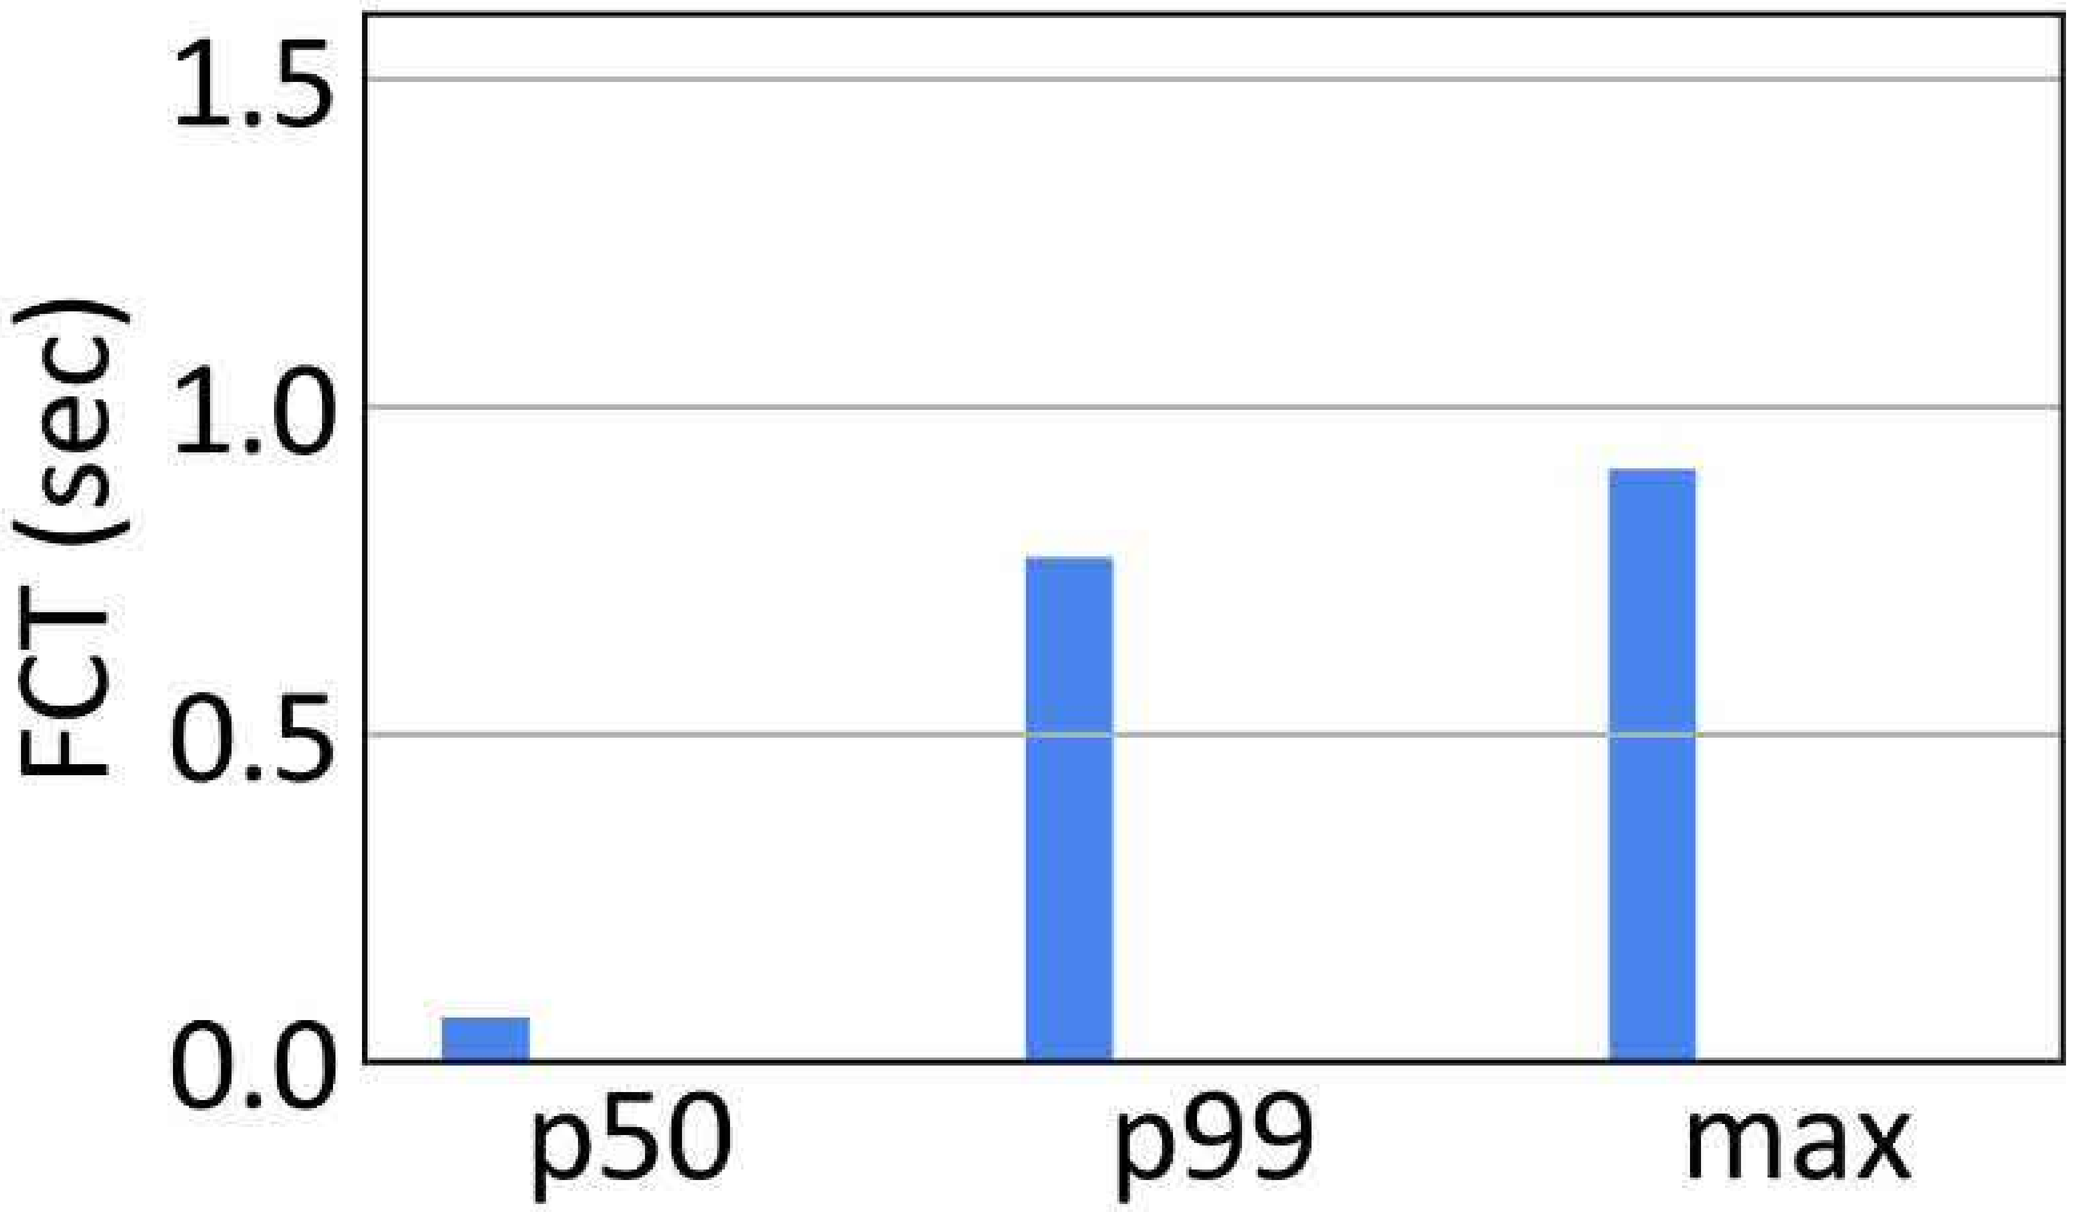

Carnegie Mellon University

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)

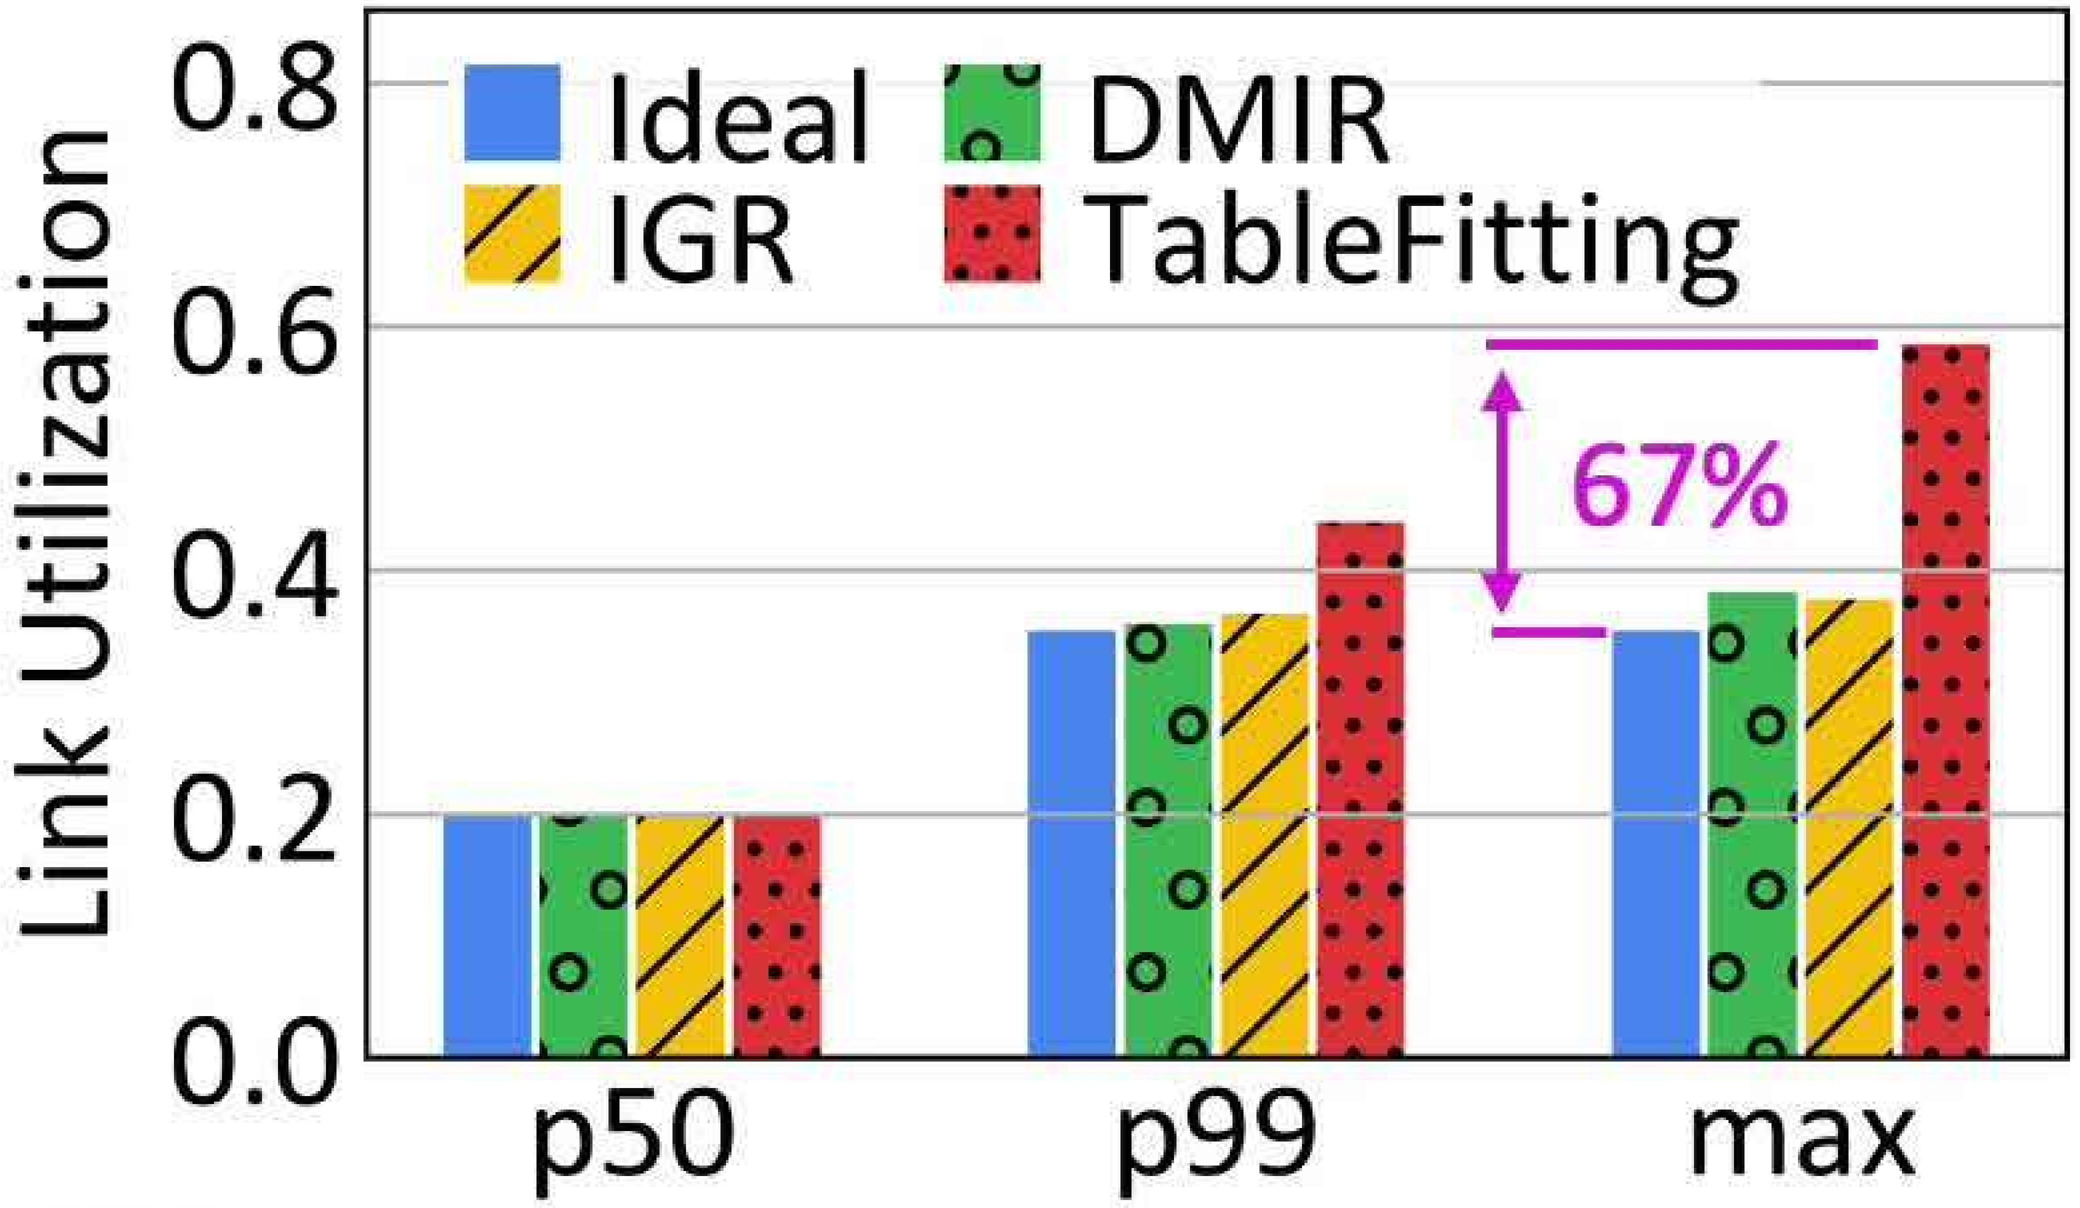

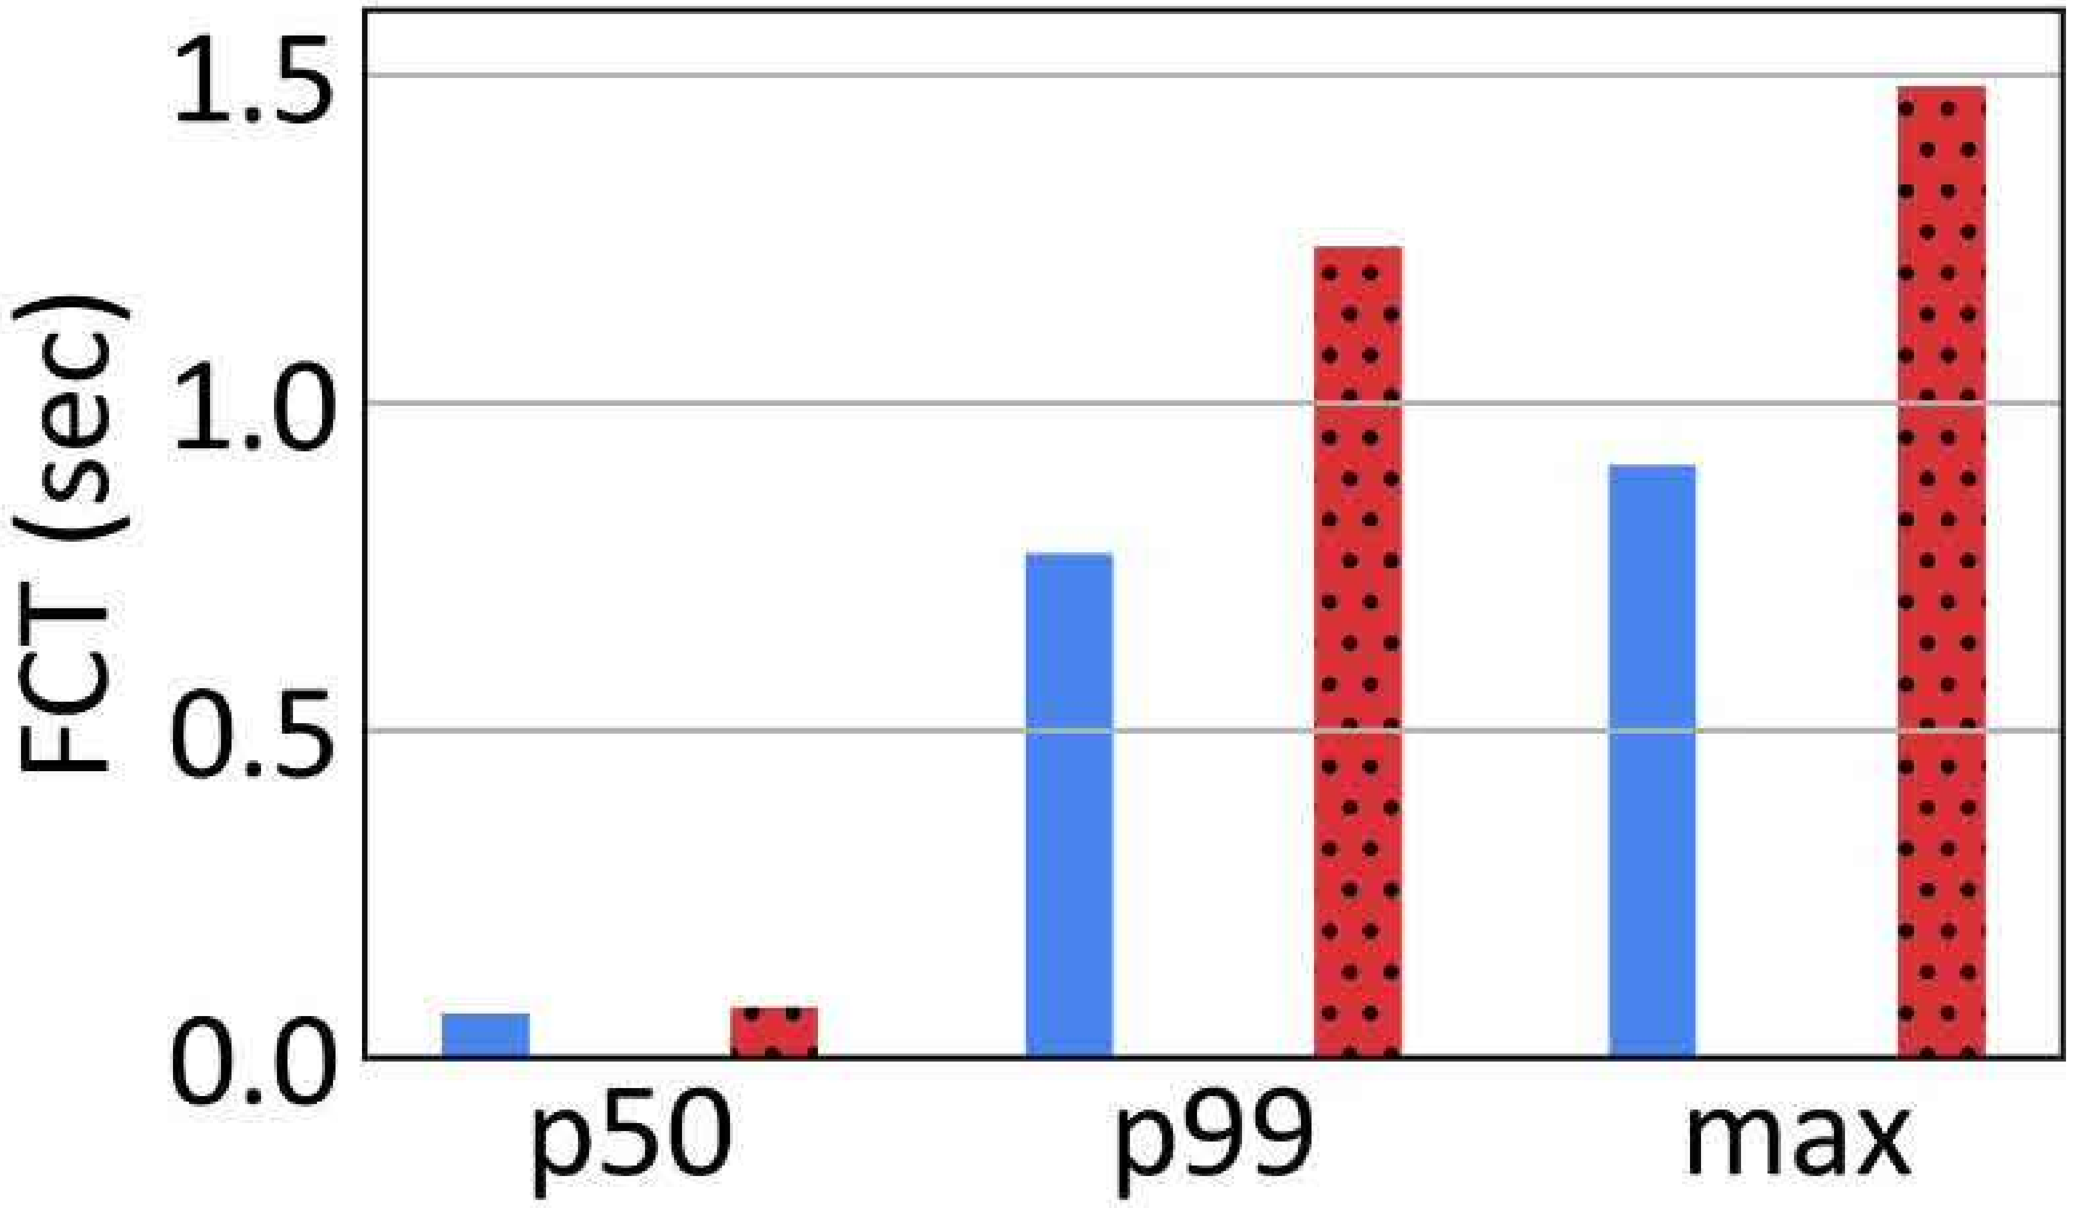

Carnegie Mellon University

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)

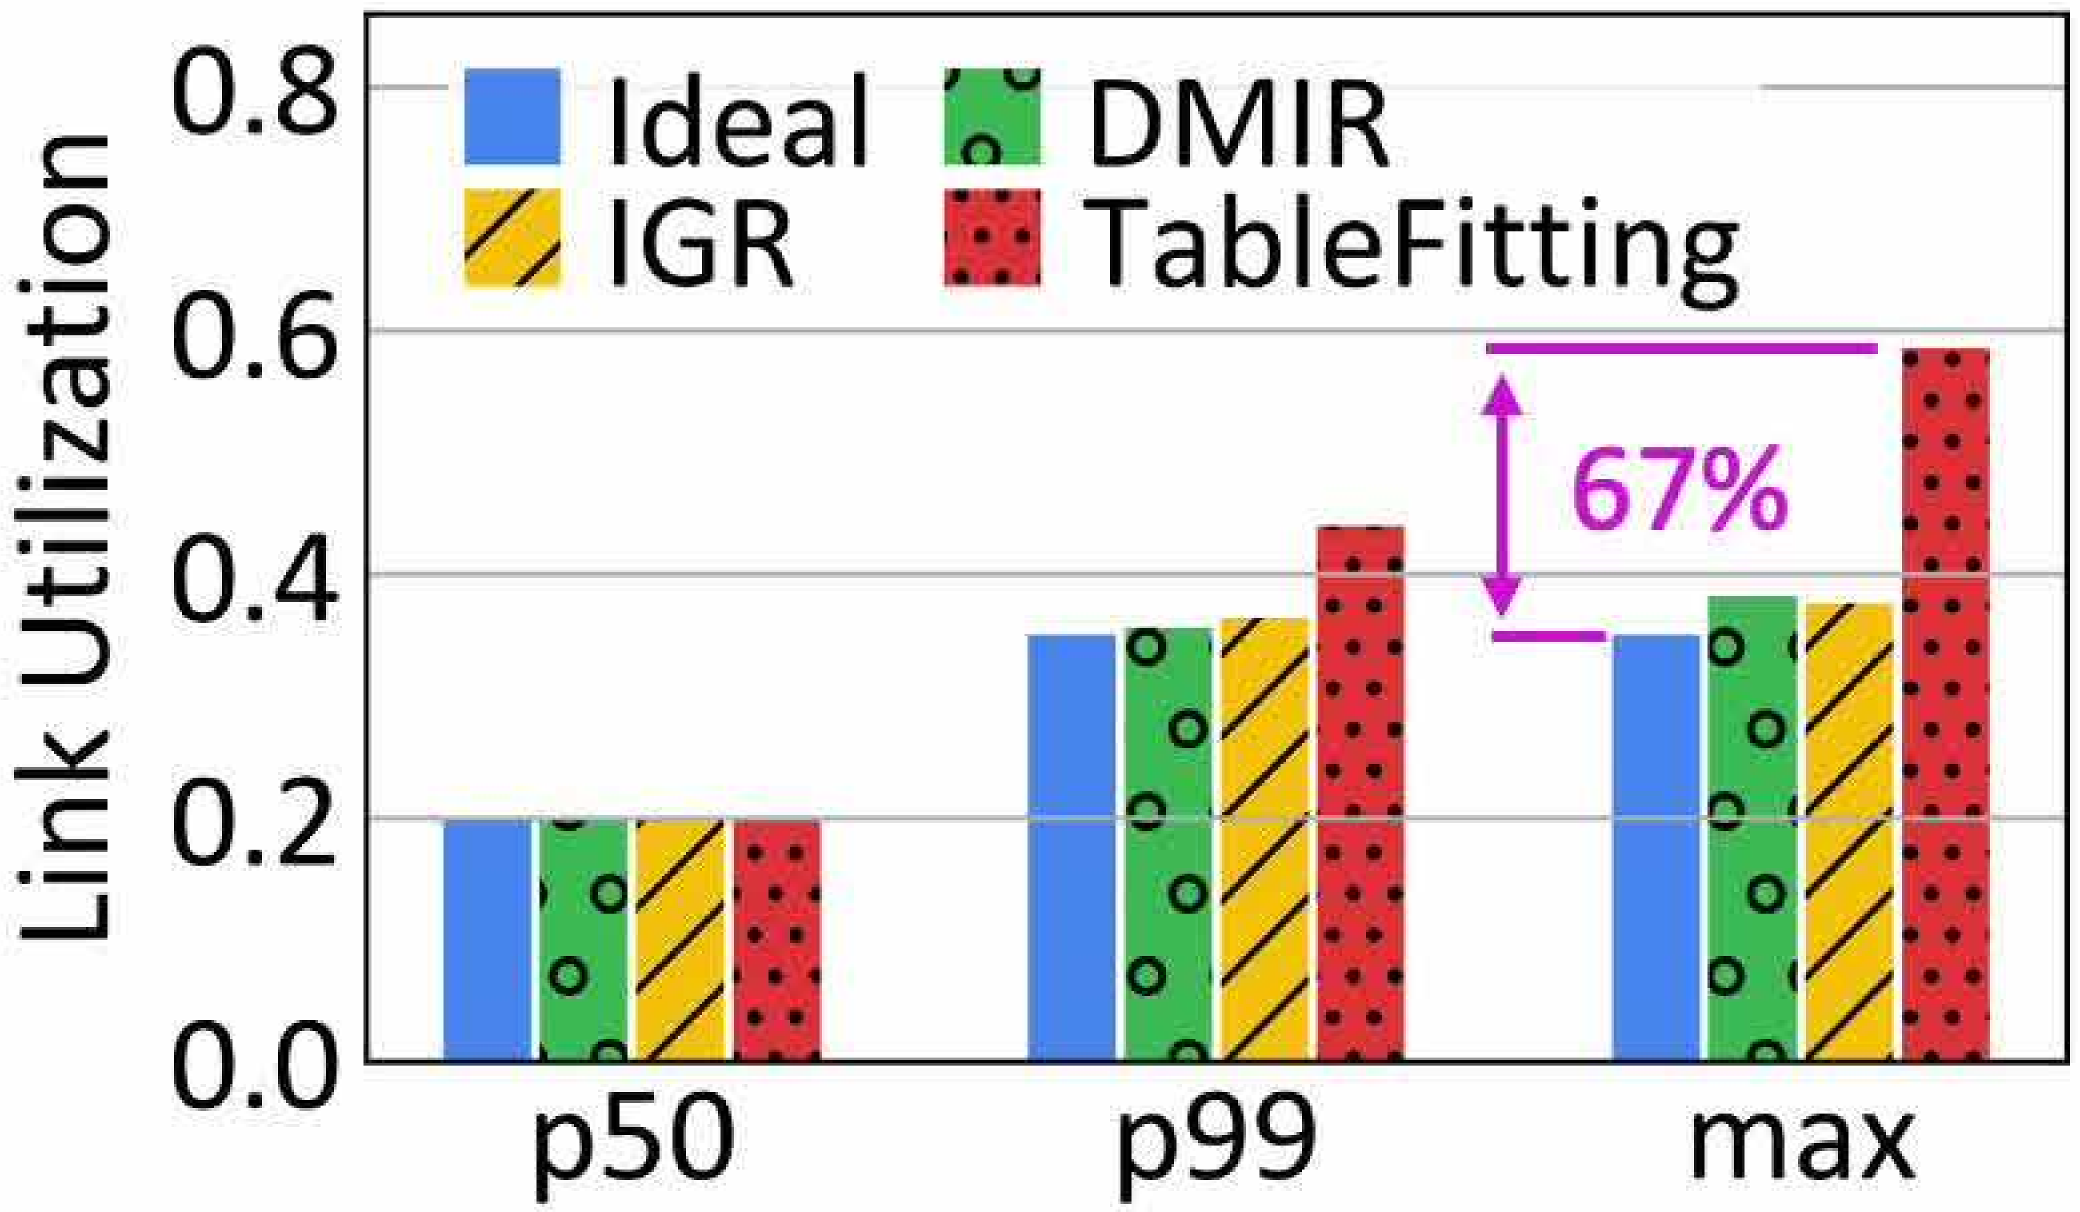

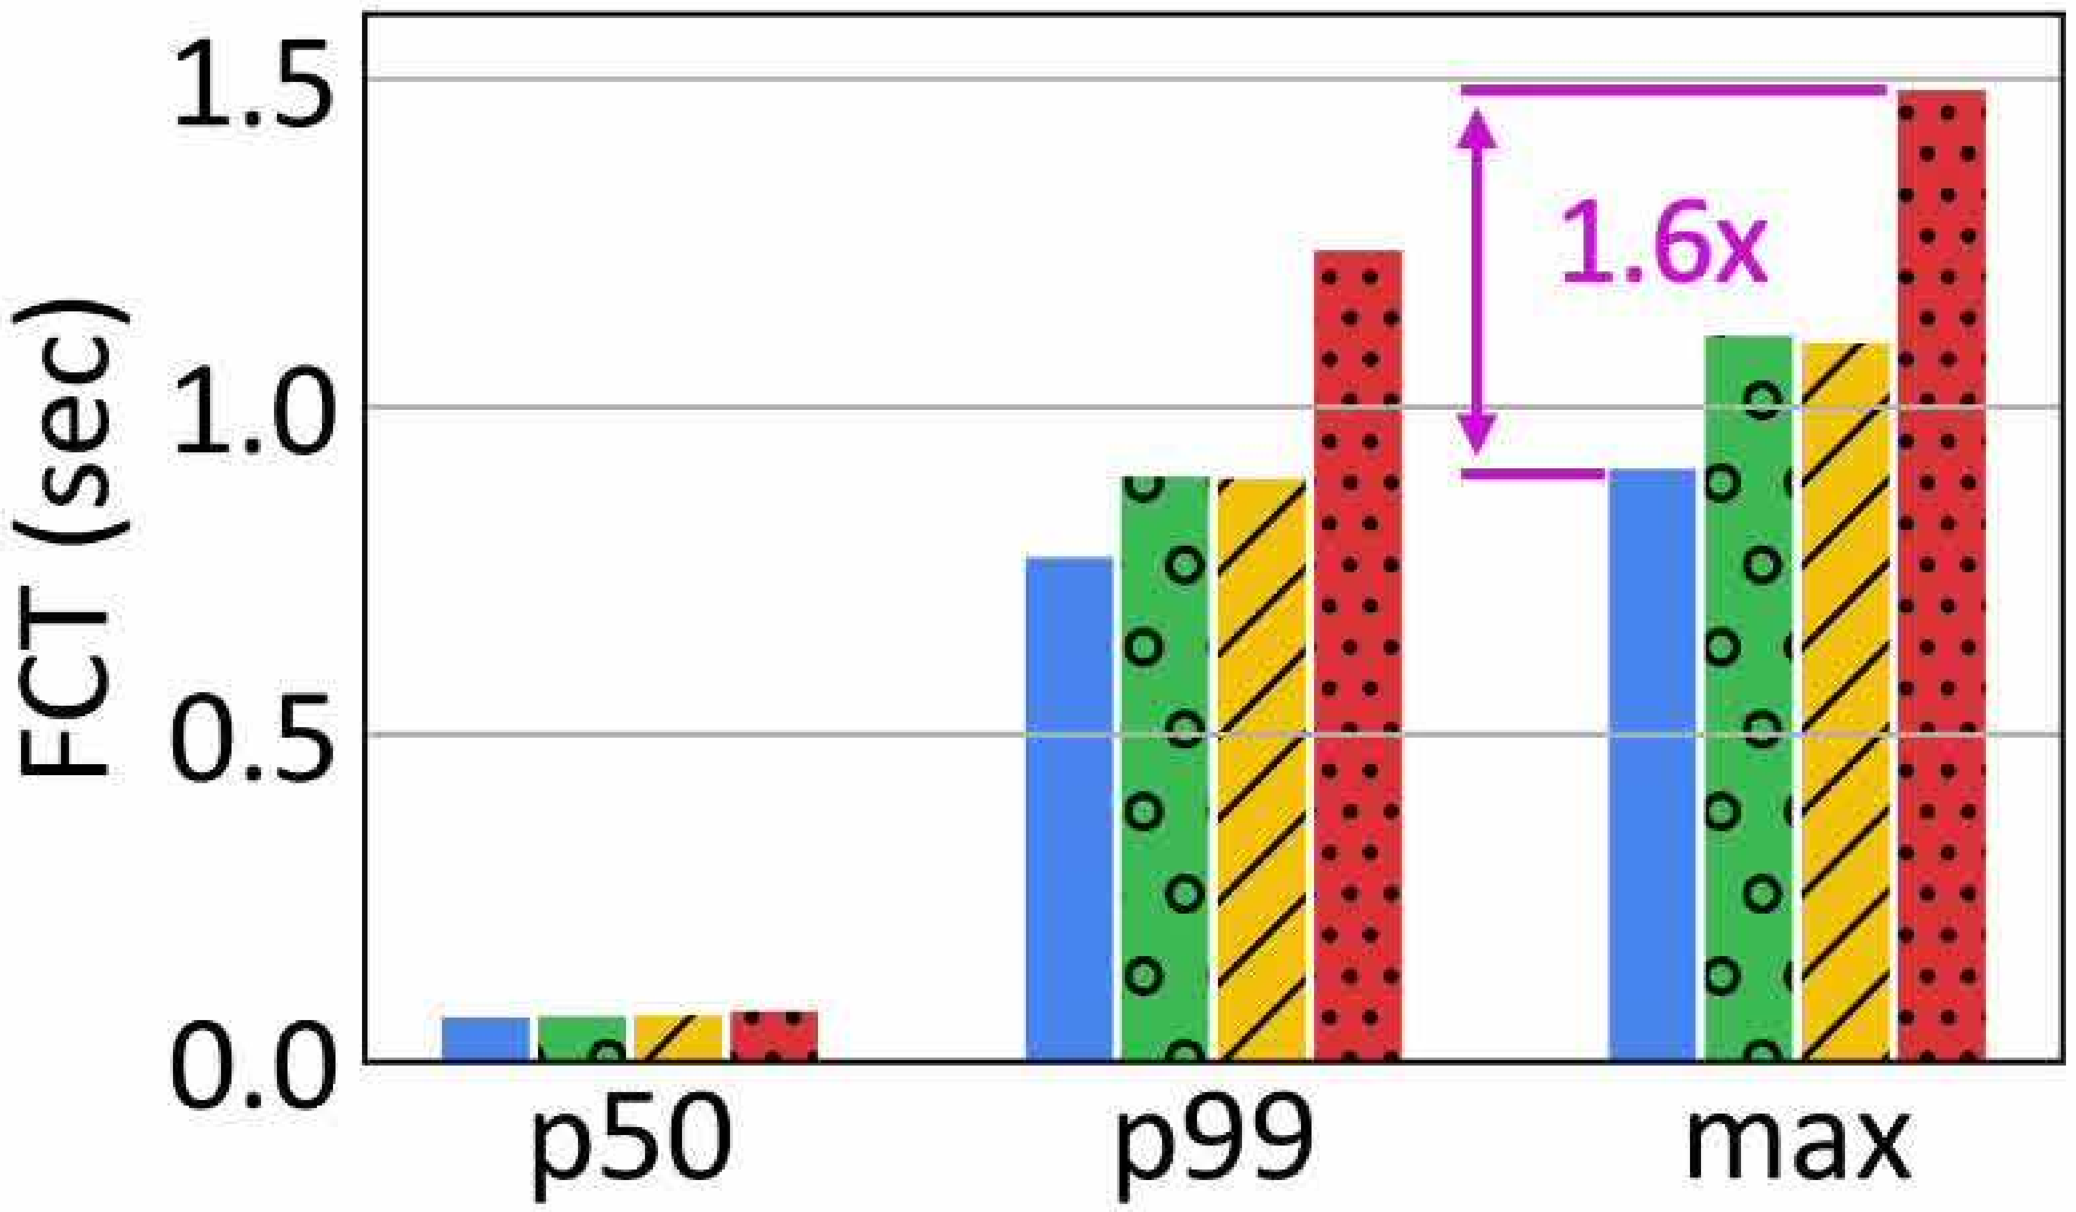

Carnegie Mellon University

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)
  - [worst case] DMIR & IGR **1.2x longer** vs. TableFitting **1.6x longer**

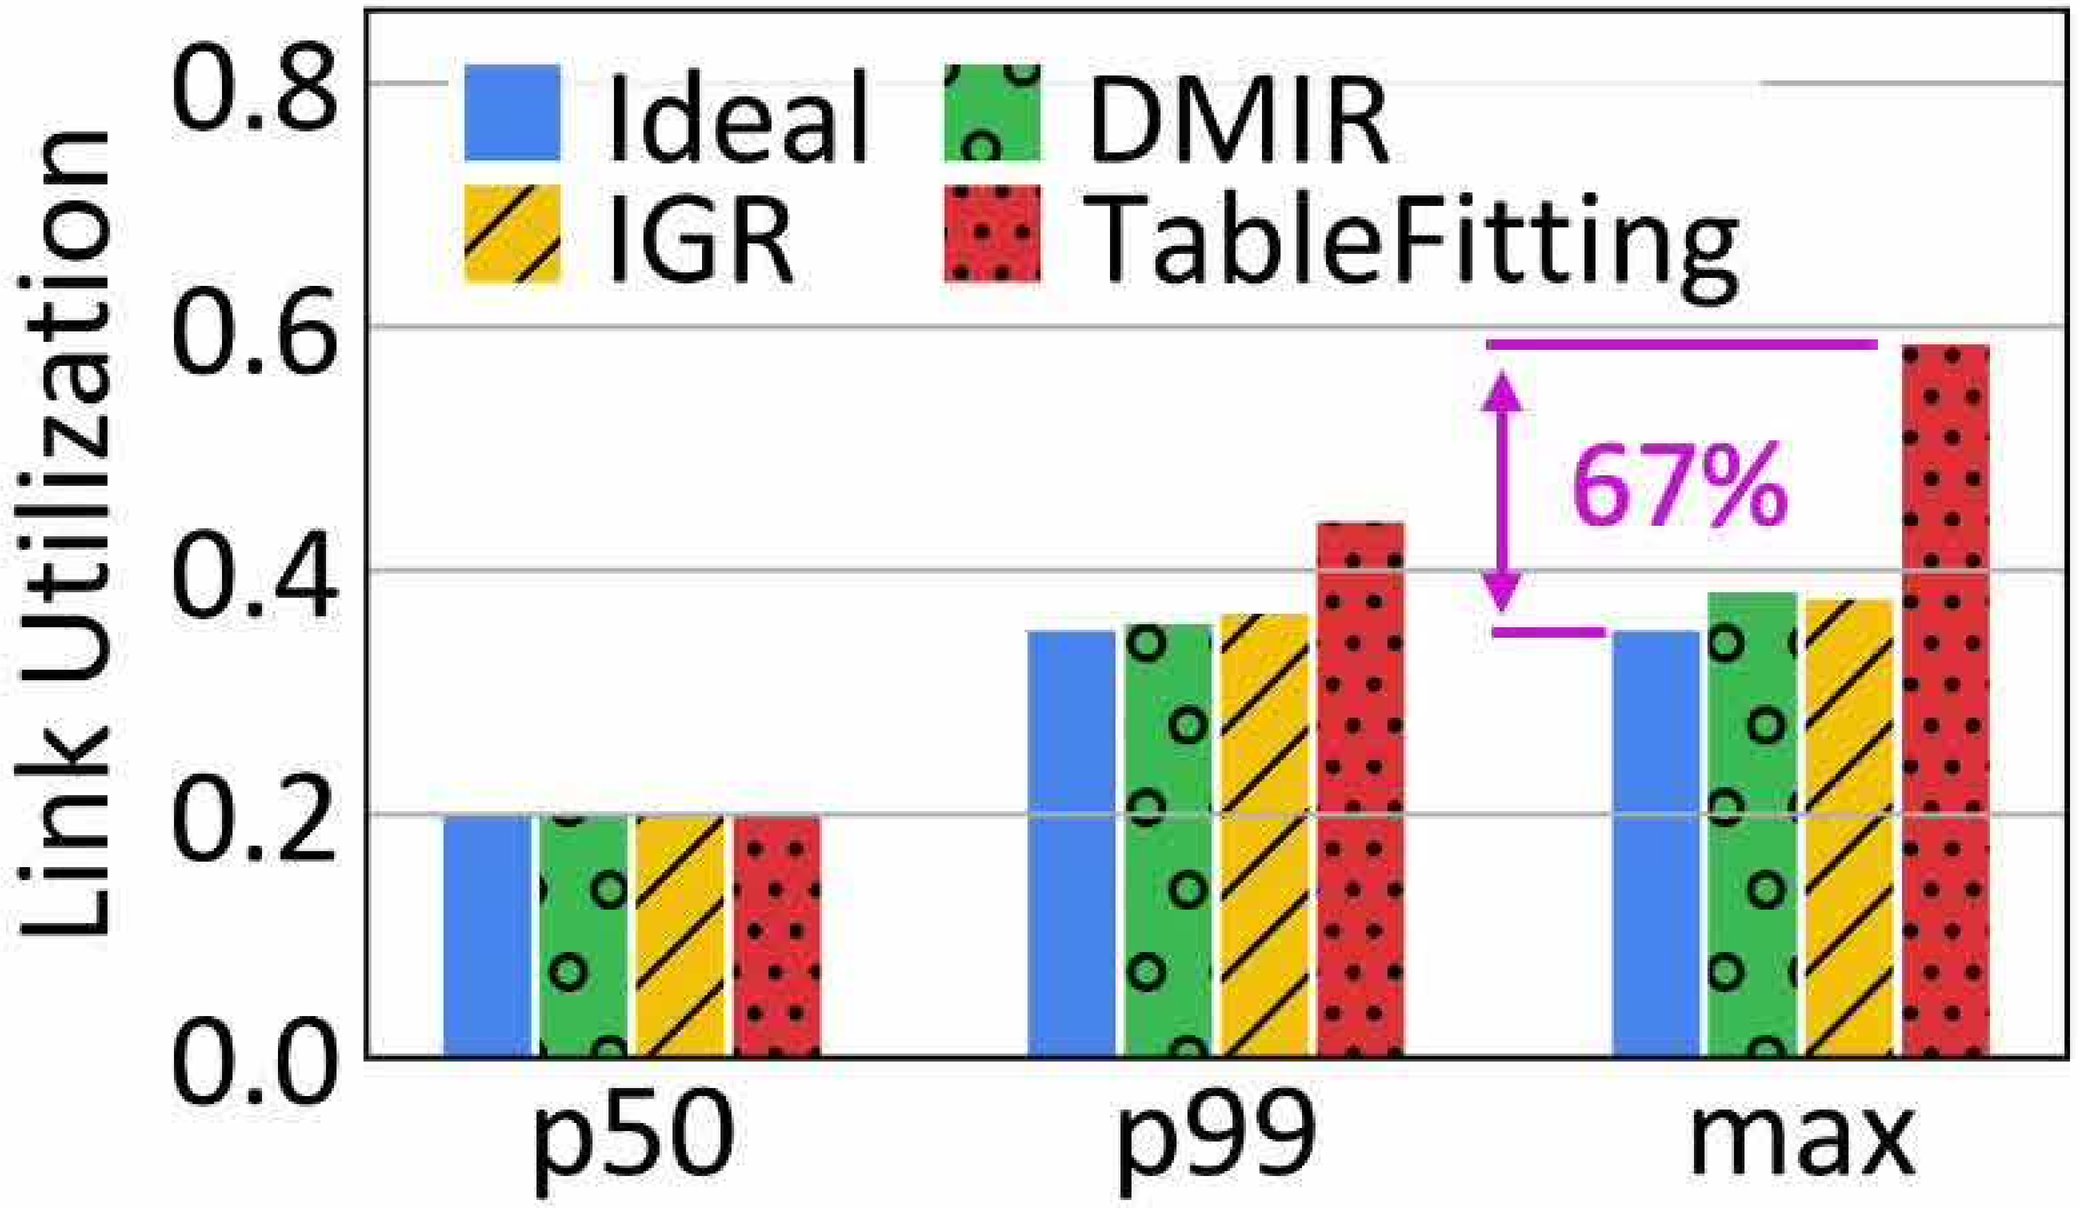

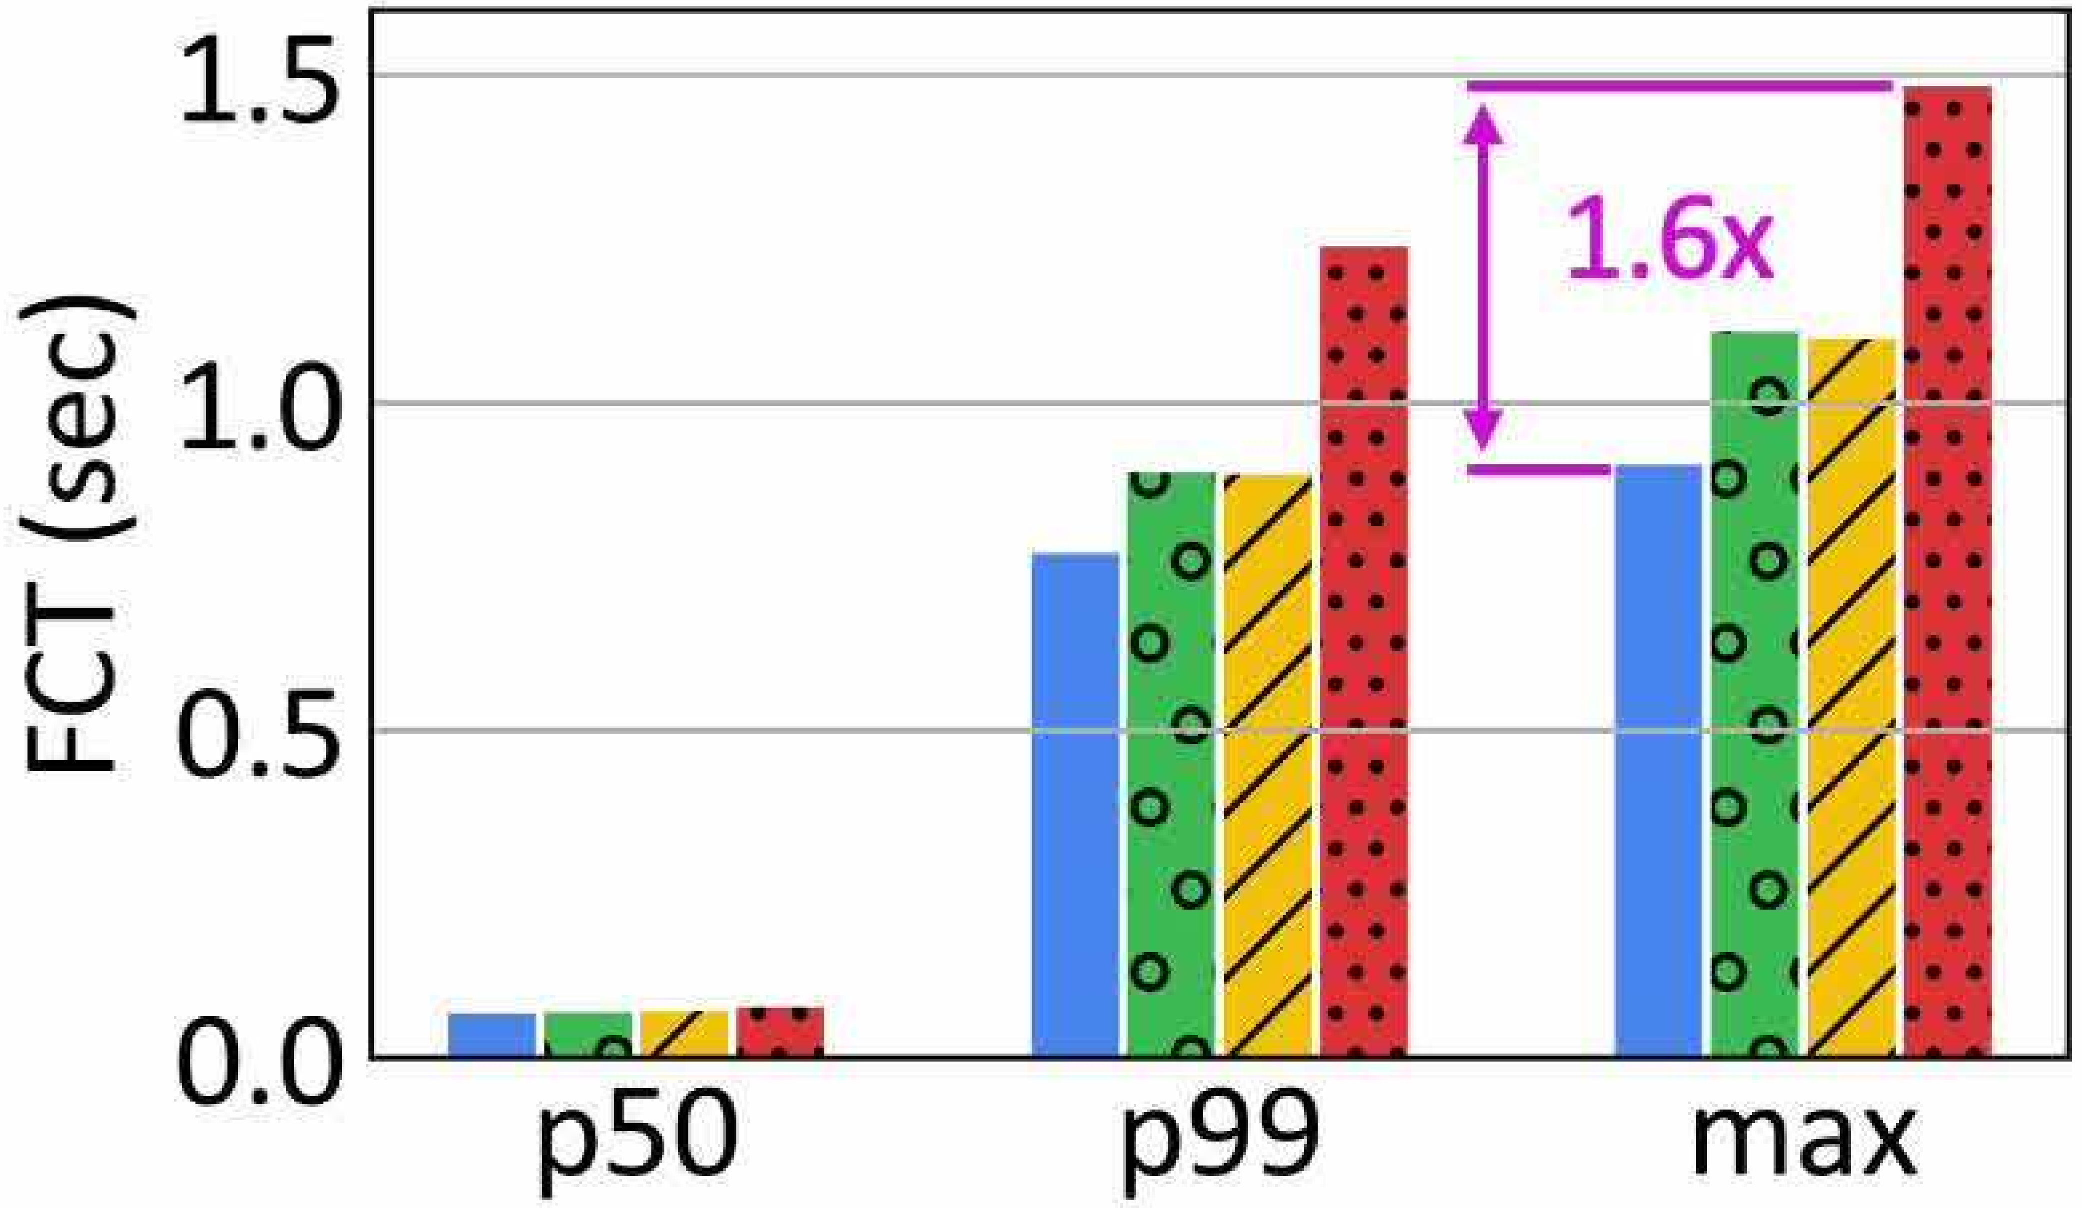

Carnegie Mellon University

DMIR & IGR are more precise than current work.

- Network-level metric: link utilization
  - [worst case] DMIR & IGR **7% error** vs. TableFitting **67% error**
- Application-level metric: flow completion time (FCT)
  - [worst case] DMIR & IGR **1.2x longer** vs. TableFitting **1.6x longer**

DMIR outperforms IGR in certain challenging scenarios, see our paper for details.

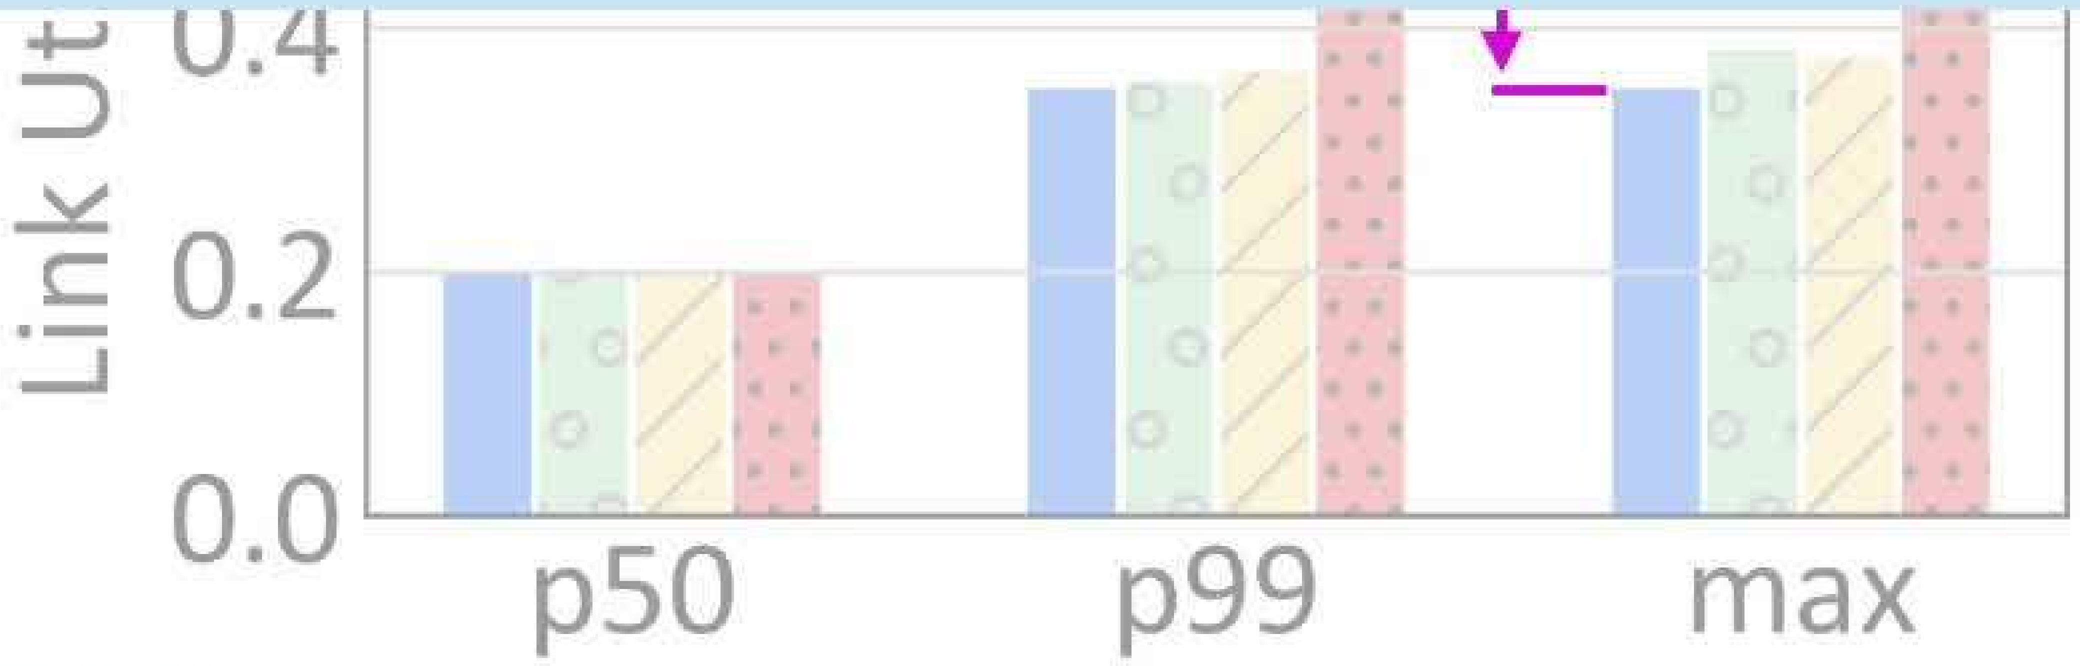

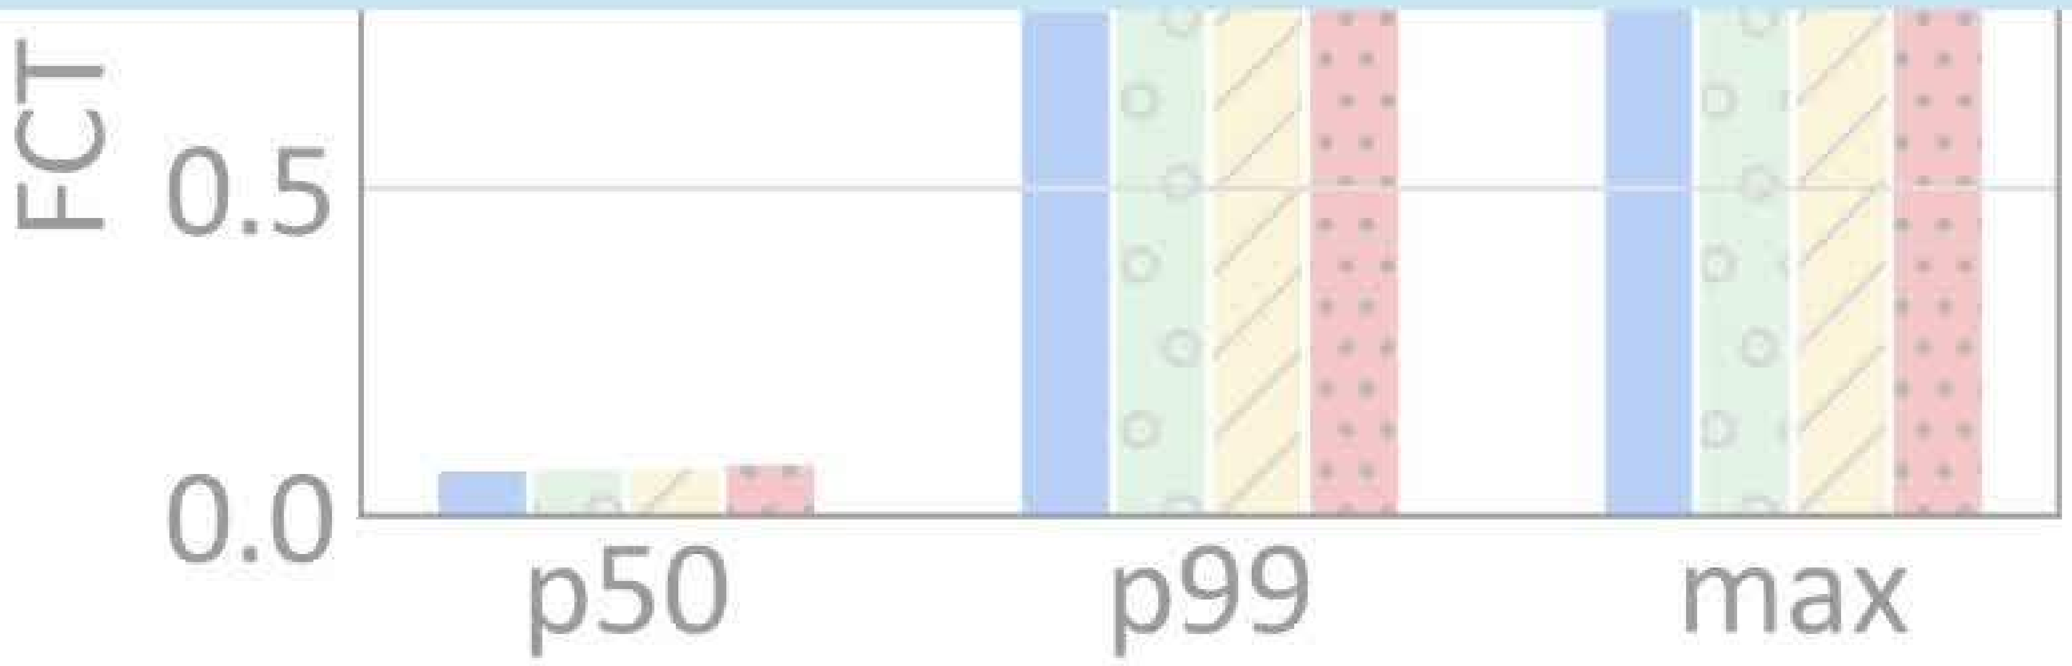

IGR runs faster than current work & DMIR.

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch

Carnegie Mellon University

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch

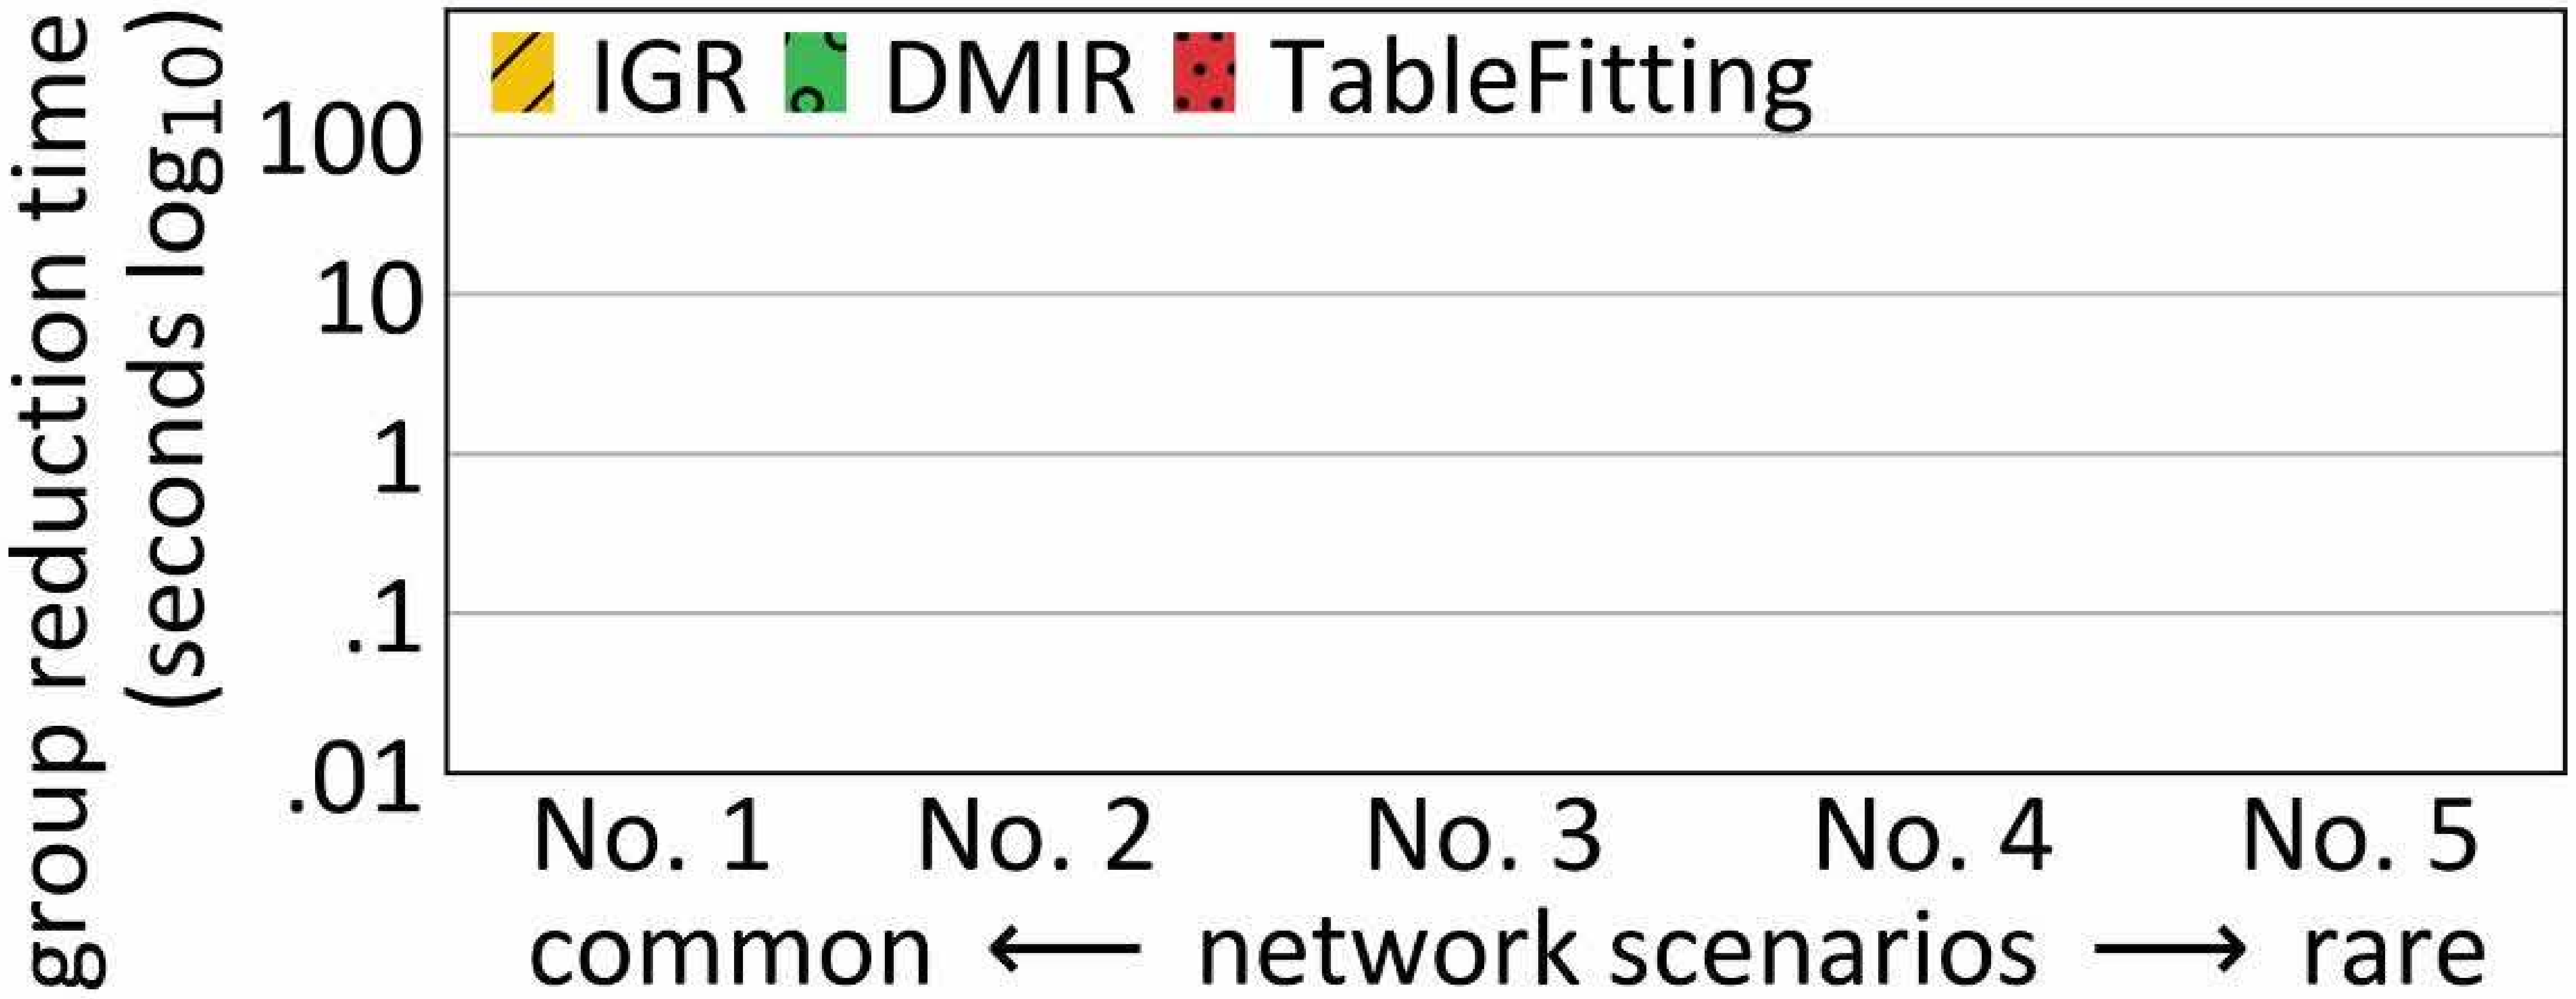

Carnegie Mellon University

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch

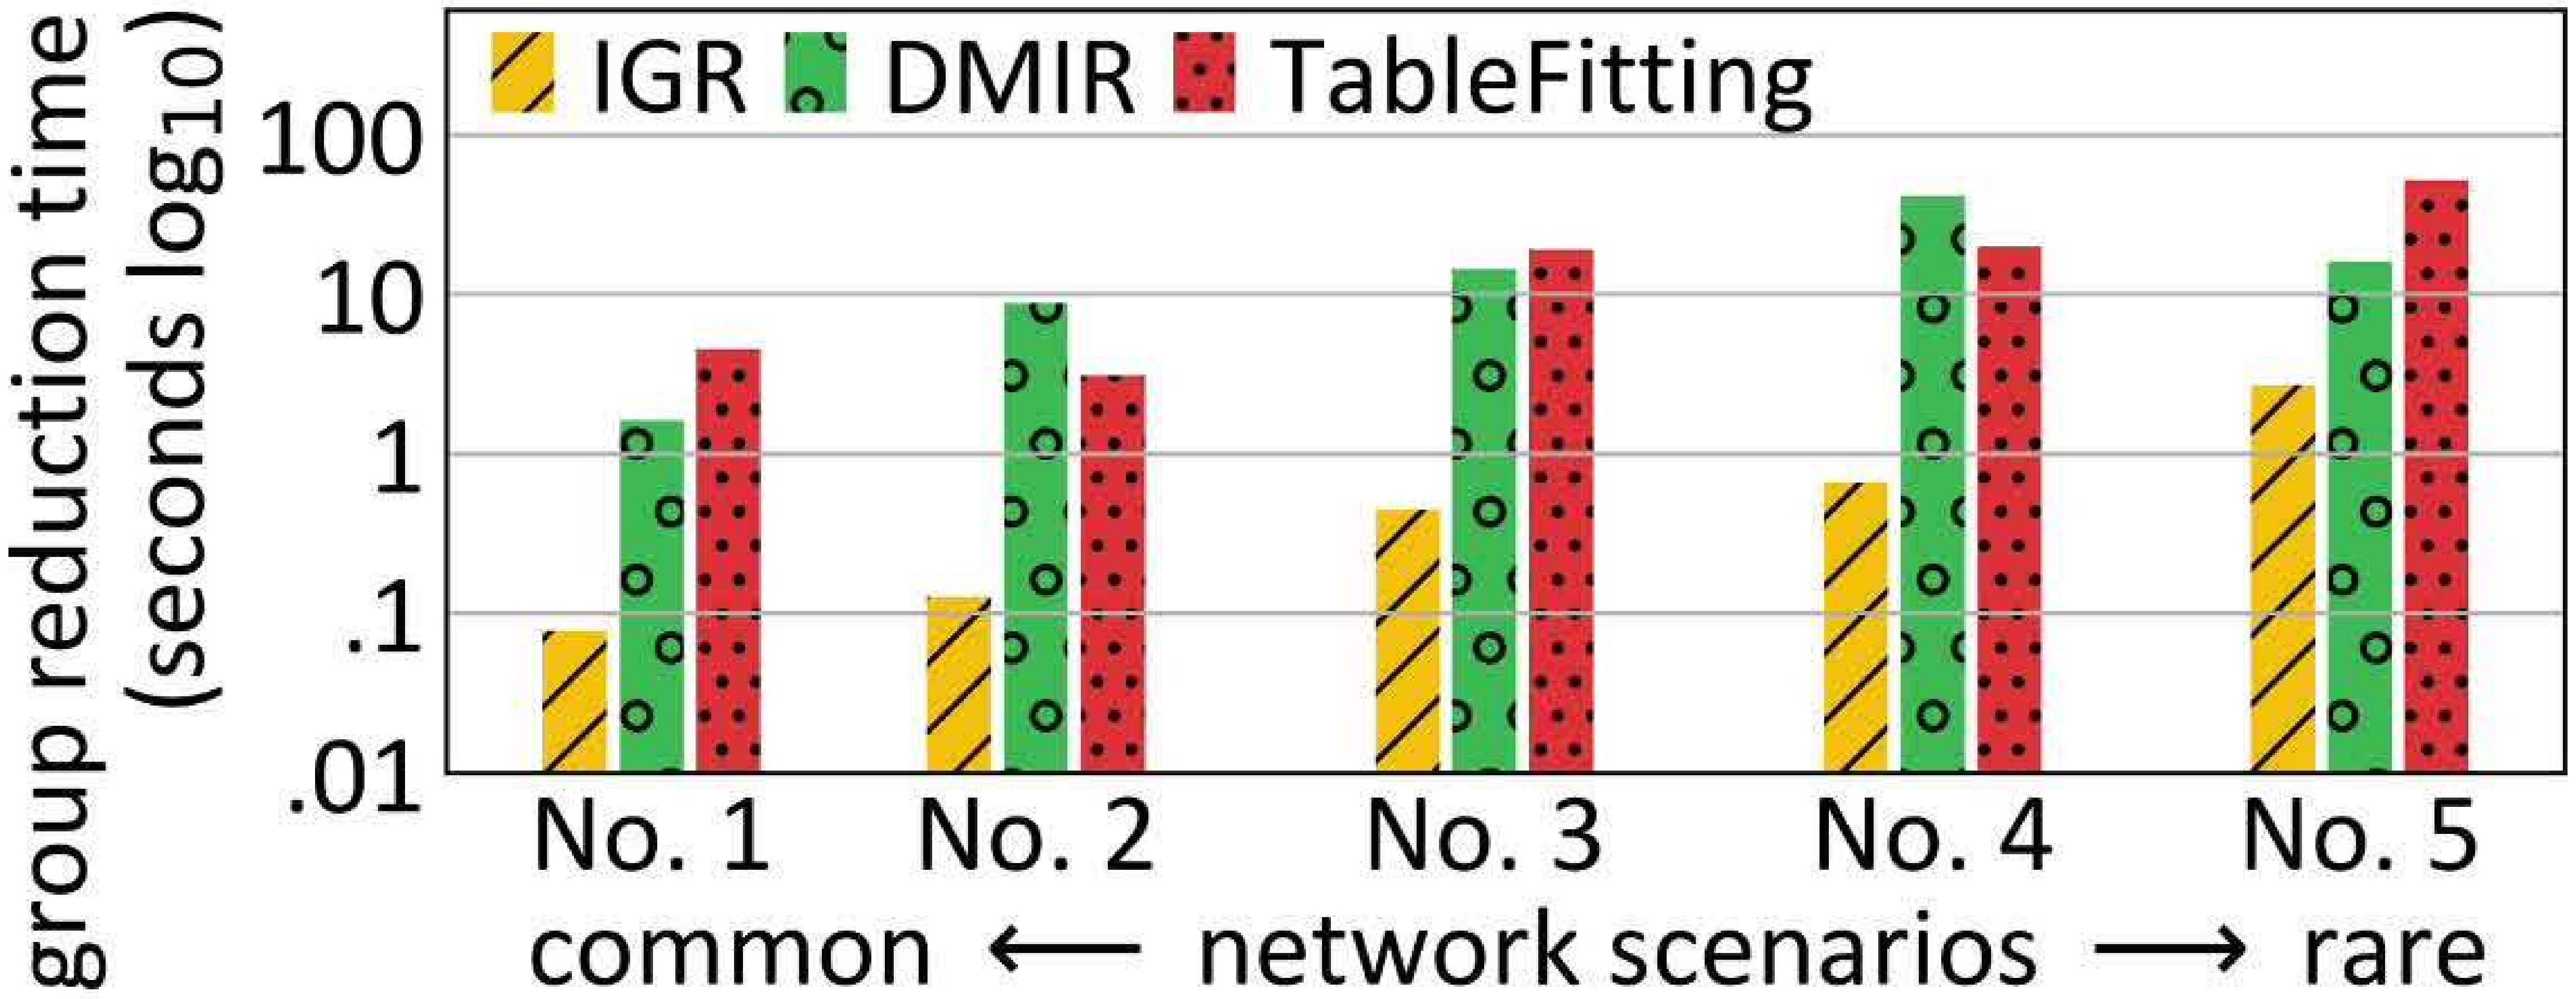

Carnegie Mellon University

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch
- IGR outperforms TableFitting by 1-2 orders of magnitude.

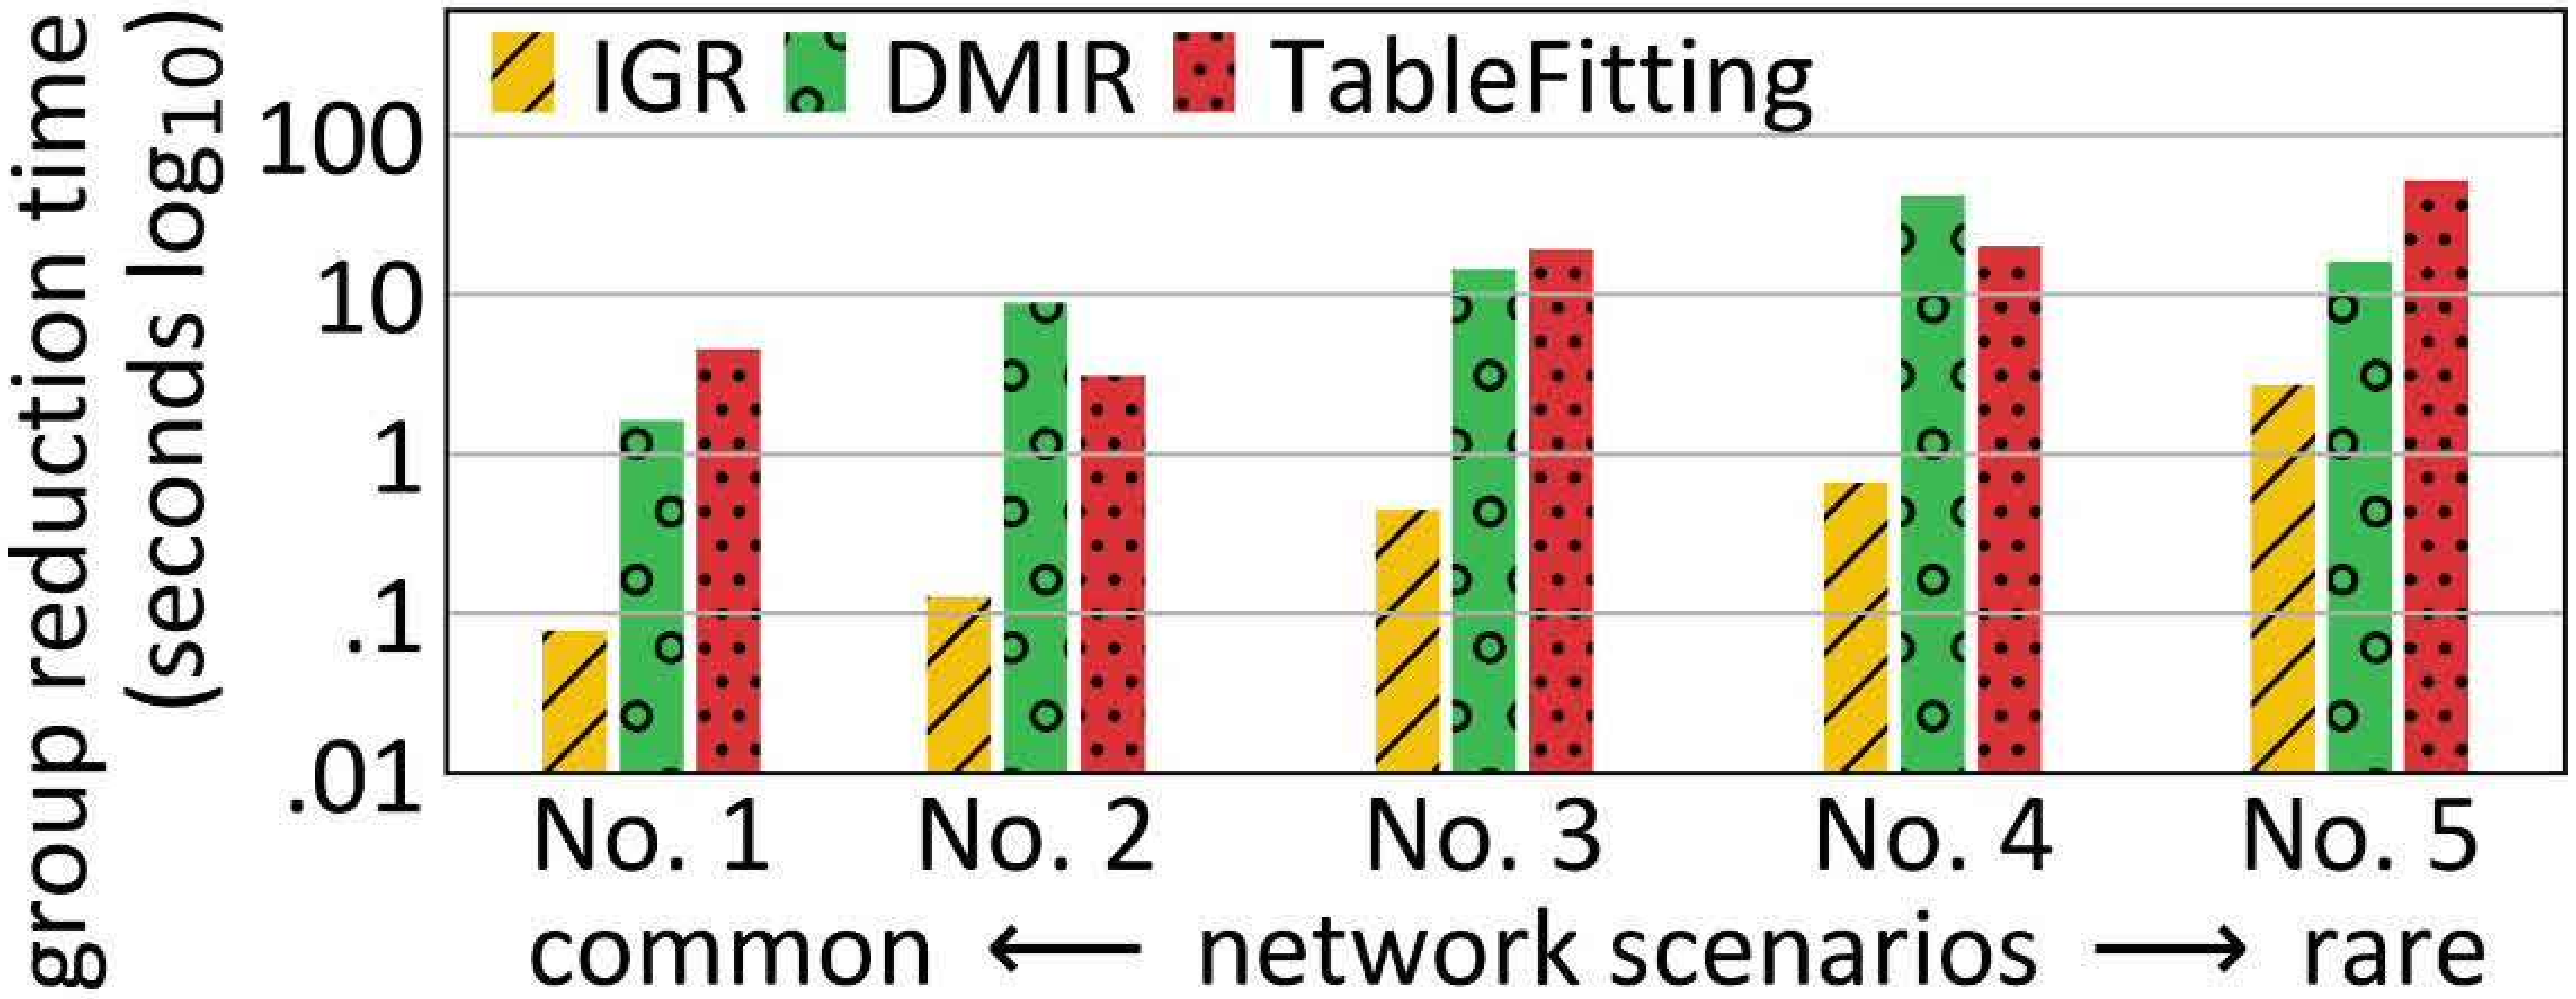

Carnegie Mellon University

IGR runs faster than current work & DMIR.

- Metric: average time to reduce a batch of groups on a switch
- IGR outperforms TableFitting by 1-2 orders of magnitude.
- DMIR is on par with TableFitting.

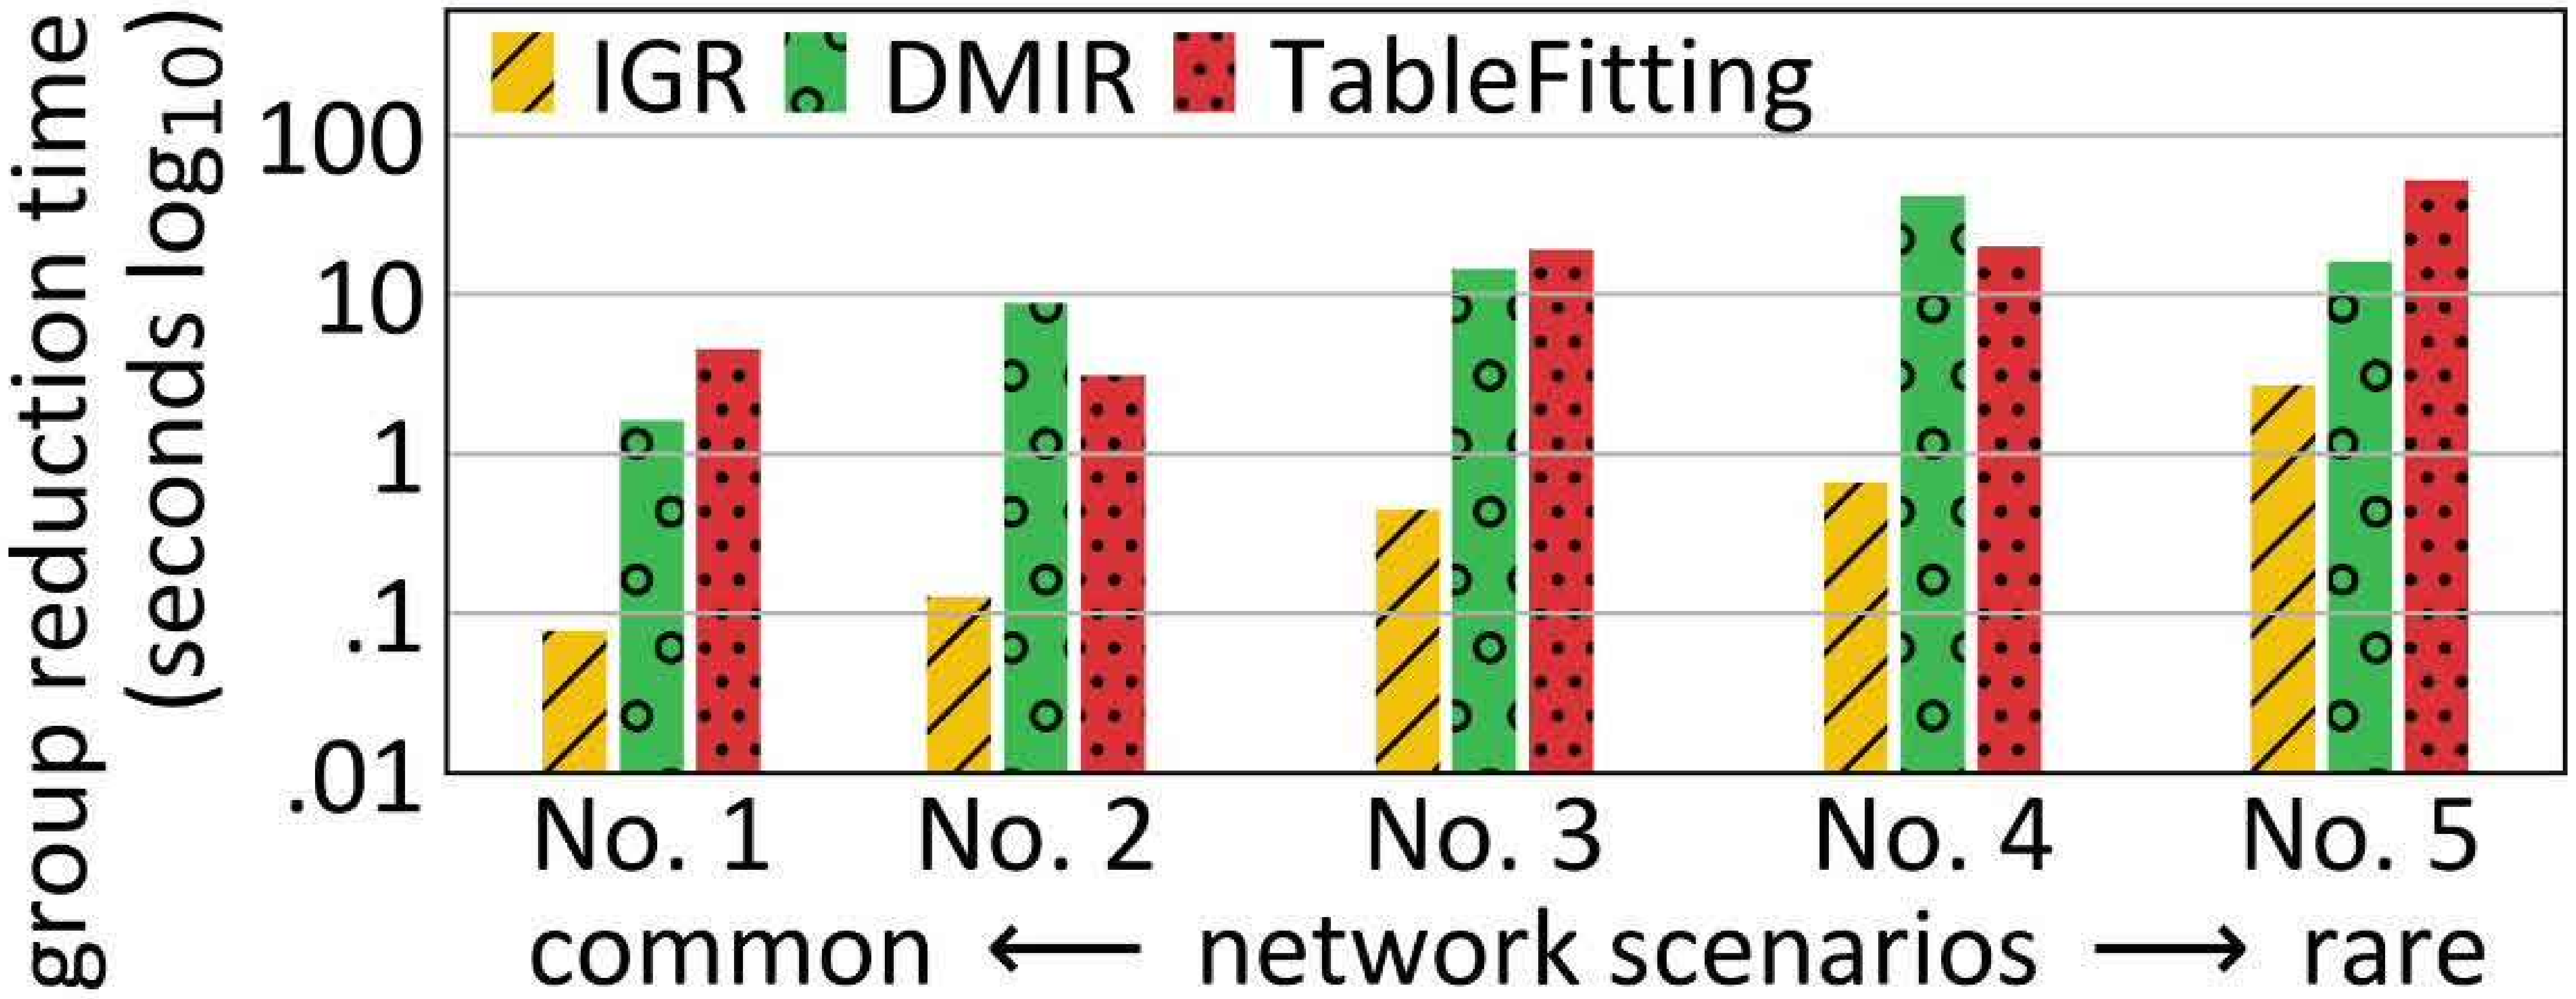

Summary

- Precision loss is inherent in TE with limited hardware resources.
  - It leads to load imbalance & traffic loss.
- We design 2 group reduction algorithms that when compared to the current approach
  - **reduce precision loss by 10x**.
  - **reduce FCT by 26%**.
  - **run up to 10x faster**.
- Use IGR for responsiveness, DMIR for challenging scenarios.

Check out this project online!

Contact us for questions:
shuoshuc@cs.cmu.edu

Carnegie Mellon University
airbnb
Google