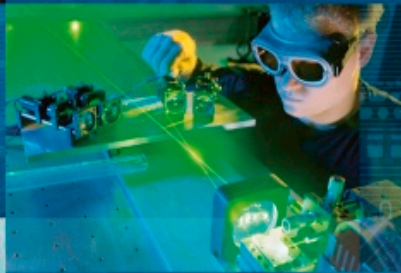
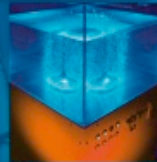


Which Graph Representation to Select for Static Graph-Algorithms on a CUDA-capable GPU

Thorsten Blaß and Michael Philippsen



Motivation (I)

Your task: write static graph algorithm on the GPU.

Preparation: check literature

- several data structures to represent graphs.
- all papers pick CRS data structure and focus on optimizing algorithmic aspects.

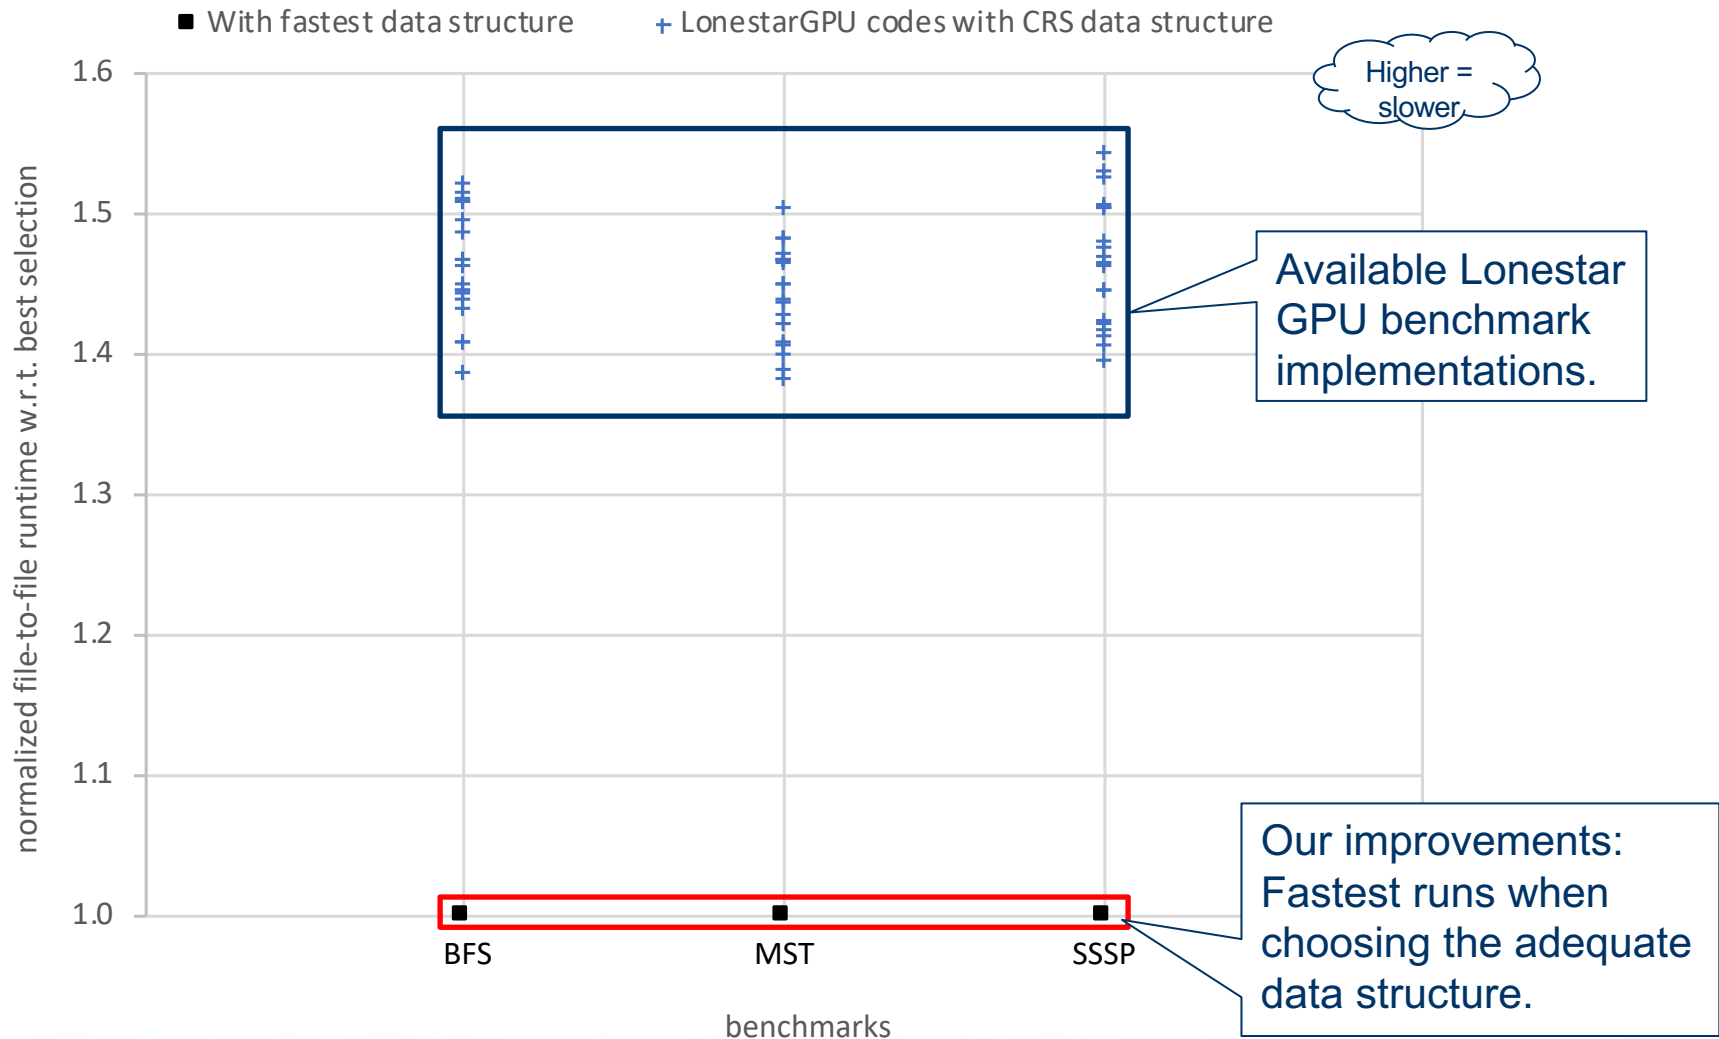
You wonder:

- Is data structure irrelevant for performance?
- Is it right to follow the crowd?

This paper: Graph representation matters!

Motivation (II)

- *Adequate* data structure speeds up graph algorithms.



Study Setup

- What is the *adequate* data structure?
- A total of 754,000 measurements:
 - 10 state-of-the-art static graph algorithms (do not modify the graph).
 - 19 input graphs.
 - 4 architecturally different Nvidia GPUs.
 - 10 graph data structures.
 - 3 widely used graph exchange formats.













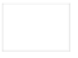





Study Setup – Benchmark Graph Algorithms

- 10 state-of-the-art implementations from recent research and benchmark suites with different characteristics.

	Thread per...	Granularity	Access Pattern	Termination	<i>weighted</i>	<i>directed</i>
APSP	edge	all	edge	worklist	✓	✓
BFS	node	subset	successors	recursion		
MIS	node	subset	random nodes	worklist		
			successors			
MST	node	all	successors	recursion	✓	
PR-n	node	all	predecessors	fixpoint		✓
PR-e	edge	all	edge	fixpoint		✓
SpMV	node	subset	successors	single pass	✓	
SSSP	node	subset	successors	worklist	✓	✓
WCC-n	node	subset	successors	fixpoint	✓	
WCC-e	edge	all	edge	fixpoint	✓	

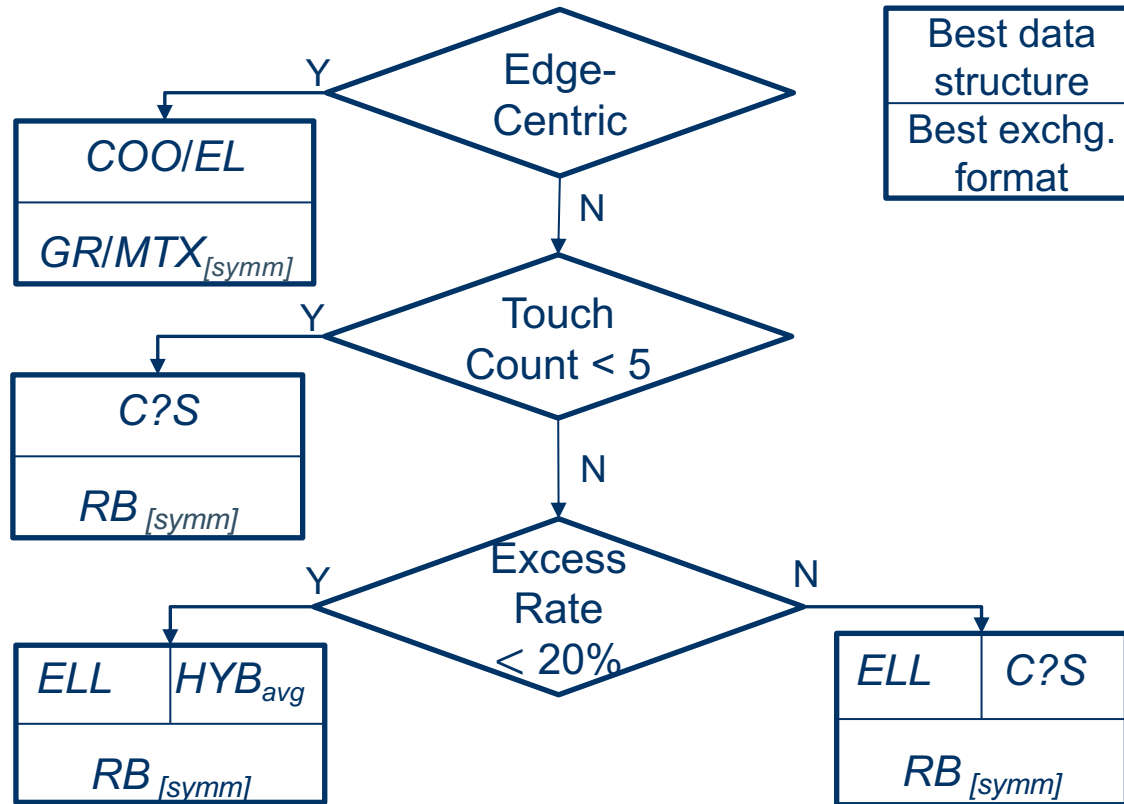
Study Setup – Input Graphs

- 19 input graphs with different characteristics.

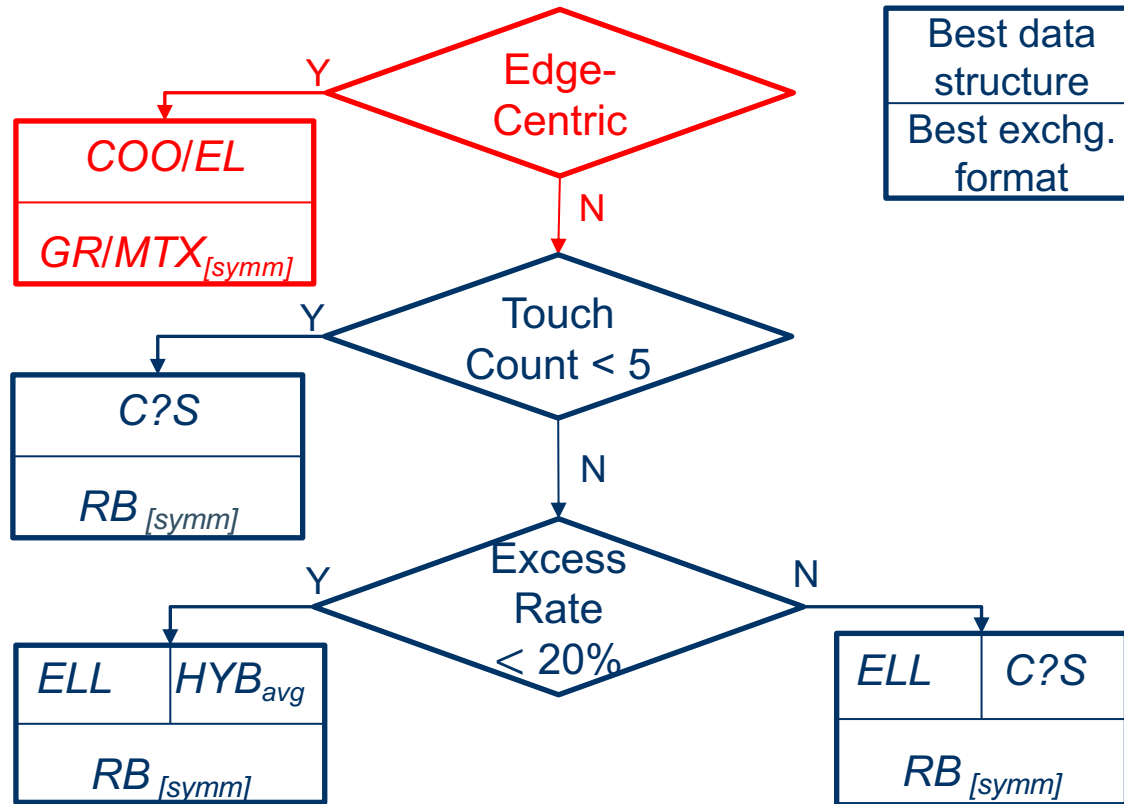
Graph (original name)	avail. Format	Nodes Edges	Degrees		Spyplot	Graph (original name)	avail. Format	Nodes Edges	Degrees		Spyplot
			avg. / max.	excess (%)					avg. / max.	excess (%)	
USA-road.d-USA	network <i>GR</i>	23.9M	2.46 / 9		it-2004	scale-free graph <i>MTX, RB</i>	41.2M	25.31 / 1,326k			
USA-road.d.-CAL		116.6M	1.79		indochina-2004		1.150.7M	6.99			
USA-road.d.-CRT		1.8M	2.46 / 8		kron_g500-logn21		7.4M	25.48 / 256.4k			
USA-road.d.-BAY		9.3M	1.83		amazon-2008		194.1M	7.16			
USA-road.d.-FLA		14.0M	2.48 / 9		twitter-retweet		2.1M	86.44 / 213.9k			
USA-road.d.-COL		68.5M	1.88		hollywood-2009		182.0M	8.06			
pwtk		mesh <i>MTX, RB</i>	0.2M	55.39 / 181			ljournal-2008	scale-free graph <i>MTX, RB</i>	0.7M	14.01 / 1k	
rgg-n-2-24-s0	5.9M		2.20	coAuthorsDBLP		5.1M	11.57				
msdoor	16.7M		15.97 / 40		soc-orkut	4.5M	11.87				
	265.1M		51.35		dblp-2010	1.1M	100.83 / 11.5k				
msdoor	0.4M		49.69 / 78		coAuthorsDBLP	113.8M	21.54				
	19.1M		56.03		soc-orkut	79.0M	22.91				
									0.3M	6.53 / 336	
							1.9M	26.98			
							3M	70.32 / 27.3k			
							106.3M	30.67			
							0.3M	4.95 / 238			
							1.6M	30.76			

- 4 architecturally different Nvidia GPUs
 - Titan XP, 12.2 GB, Pascal
 - Geforce GTX 980, 4.0 GB, Maxwell
 - Geforce GTX 680, 2.0 GB, Kepler
 - Geforce GTX 580, 1.0 GB, Fermi
- Only Titan XP can run all measurement configurations.
- All measurements scale w.r.t. nominal GPU performance.
→ Suffices to discuss the Titan XP measurements.

Decision Tree (file-to-file)

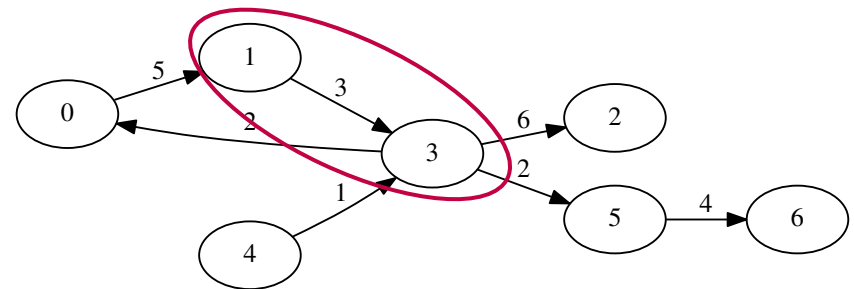


Decision Tree – First Layer (I)



Decision Tree – First Layer (II)

- First decision depends on whether the algorithm is edge-centric or node-centric.
 - Splits data structures into two groups.
- For edge-centric algorithms *COO* and *EL* perform best.
 - The other formats are 25% slower in our measurements.



- COO**ordinate format

- Stores source, target, and weight of an edge in different arrays.

$$\begin{array}{l} \text{src} = [0 \quad 1 \quad 3 \quad 3 \quad 3 \quad 4 \quad 5] \\ \text{trgt} = [1 \quad 3 \quad 0 \quad 2 \quad 5 \quad 3 \quad 6] \\ \text{weights} = [5 \quad 3 \quad 2 \quad 6 \quad 2 \quad 1 \quad 4] \end{array}$$

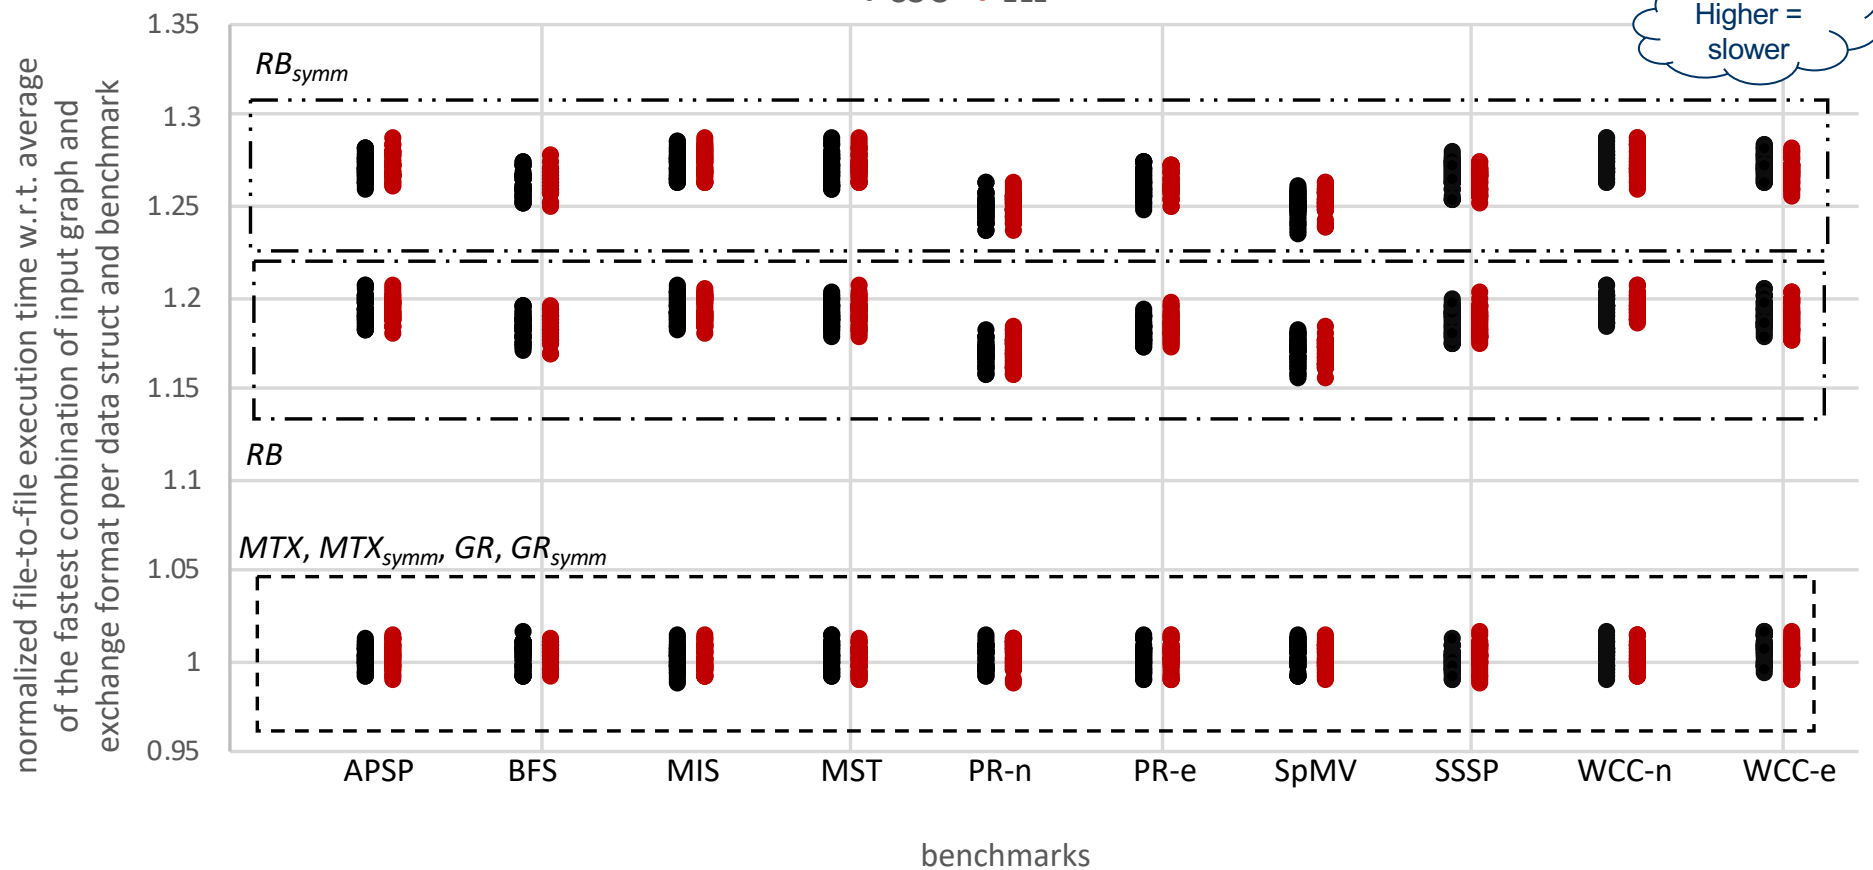
- Edge List** format

- Stores an edge as a struct in an array.

$$\text{edges} = \left[\left\{ \begin{array}{c} 0 \\ 1 \\ 5 \end{array} \right\}, \left\{ \begin{array}{c} 1 \\ 3 \\ 3 \end{array} \right\}, \left\{ \begin{array}{c} 3 \\ 0 \\ 2 \end{array} \right\}, \dots \right]$$

Decision Tree – First Layer (III)

- COO and *EL* show the same runtime behavior.
 - Either COO or *EL* can be picked for edge-centric algorithms.
 - Runtime can be minimized with MTX/GR exchange format.



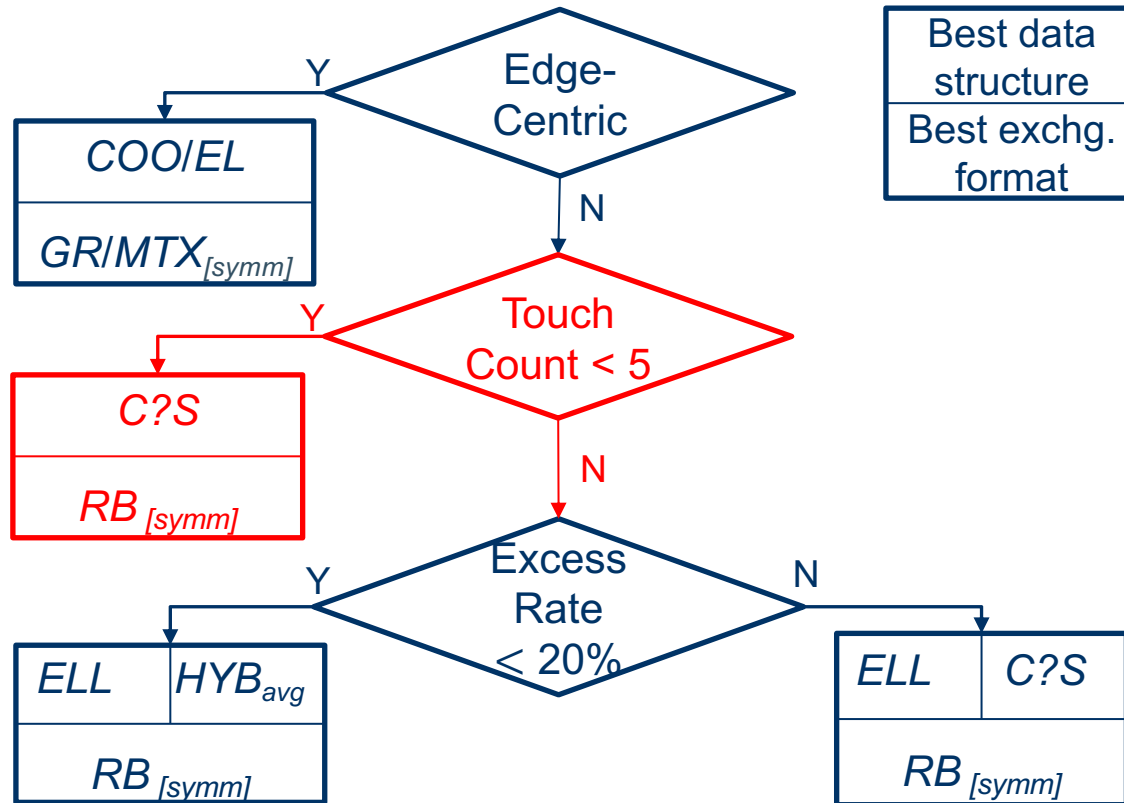
Decision Tree – First Layer (IV)

- The *MTX* and *GR* input file formats are dumps of the *COO* and *EL* formats.
- No costly conversion from file to representation in CPU memory.

	< <i>MTX header</i> >	< <i>GR header</i> >
Size of the array(s)	{ 7 7 7	<i>p sp</i> 7 7 7
	0 1 5	<i>a</i> 0 1 5
	1 3 3	<i>a</i> 1 3 3
	3 0 2	<i>a</i> 3 0 2
Encoded edges	{ 3 2 6	<i>a</i> 3 2 6
	3 5 2	<i>a</i> 3 5 2
	4 3 1	<i>a</i> 4 3 1
	5 6 4	<i>a</i> 5 6 4

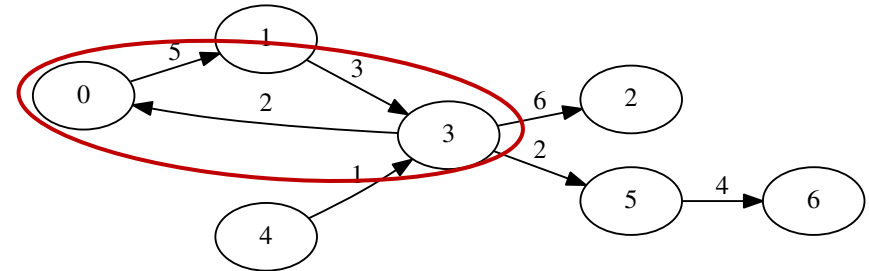
- The *RB* format is 20% slower in our measurements.

Decision Tree (file-to-file)



Decision Tree – Second Layer (I)

- Node-centric algorithms need a deeper analysis.



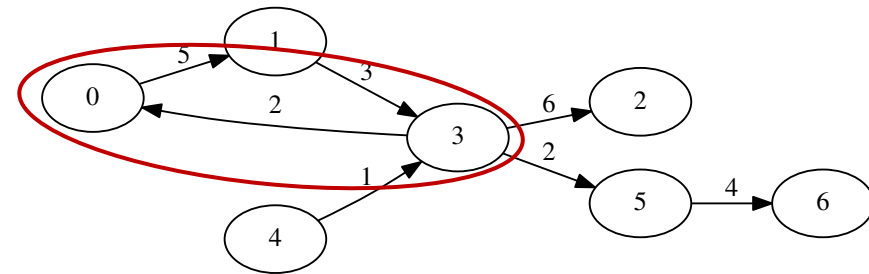
- Compressed Row Storage format**
 - Most memory efficient format.
 - Indirect addressing step necessary.

$$\begin{aligned}
 src_{ptr} &= [0 \quad 1 \quad 2 \quad \boxed{2} \quad 5 \quad 6 \quad 7 \quad 7] \\
 succ &= [1 \quad 3 \quad \boxed{0} \quad 2 \quad 5 \quad 3 \quad 6] \\
 weights &= [5 \quad 3 \quad \boxed{2} \quad 6 \quad 2 \quad 1 \quad 4]
 \end{aligned}$$

- Compressed Column Storage format**
 - Same as CRS but stores the predecessors of a node.

$$\begin{aligned}
 src_{ptr} &= [\boxed{0} \quad 1 \quad 2 \quad 3 \quad 5 \quad 5 \quad 7 \quad 8] \\
 pred &= [\boxed{3} \quad 0 \quad 3 \quad 1 \quad 4 \quad 3 \quad 5] \\
 weight &= [\boxed{2} \quad 5 \quad 6 \quad 3 \quad 1 \quad 2 \quad 4]
 \end{aligned}$$

Decision Tree – Second Layer (II)



■ ELL format

- Fixed row length reserved for successors.
- Well suited for vector processors.
- Memory efficiency depends on degree distribution of a graph.

<i>succ</i>	<i>weights</i>
$\begin{bmatrix} 1 & * & * \\ 3 & * & * \\ * & * & * \\ \mathbf{0} & 2 & 5 \\ 3 & * & * \\ 6 & * & * \\ * & * & * \end{bmatrix}$	$\begin{bmatrix} 5 & * & * \\ 3 & * & * \\ * & * & * \\ \mathbf{2} & 6 & 2 \\ 1 & * & * \\ 4 & * & * \\ * & * & * \end{bmatrix}$

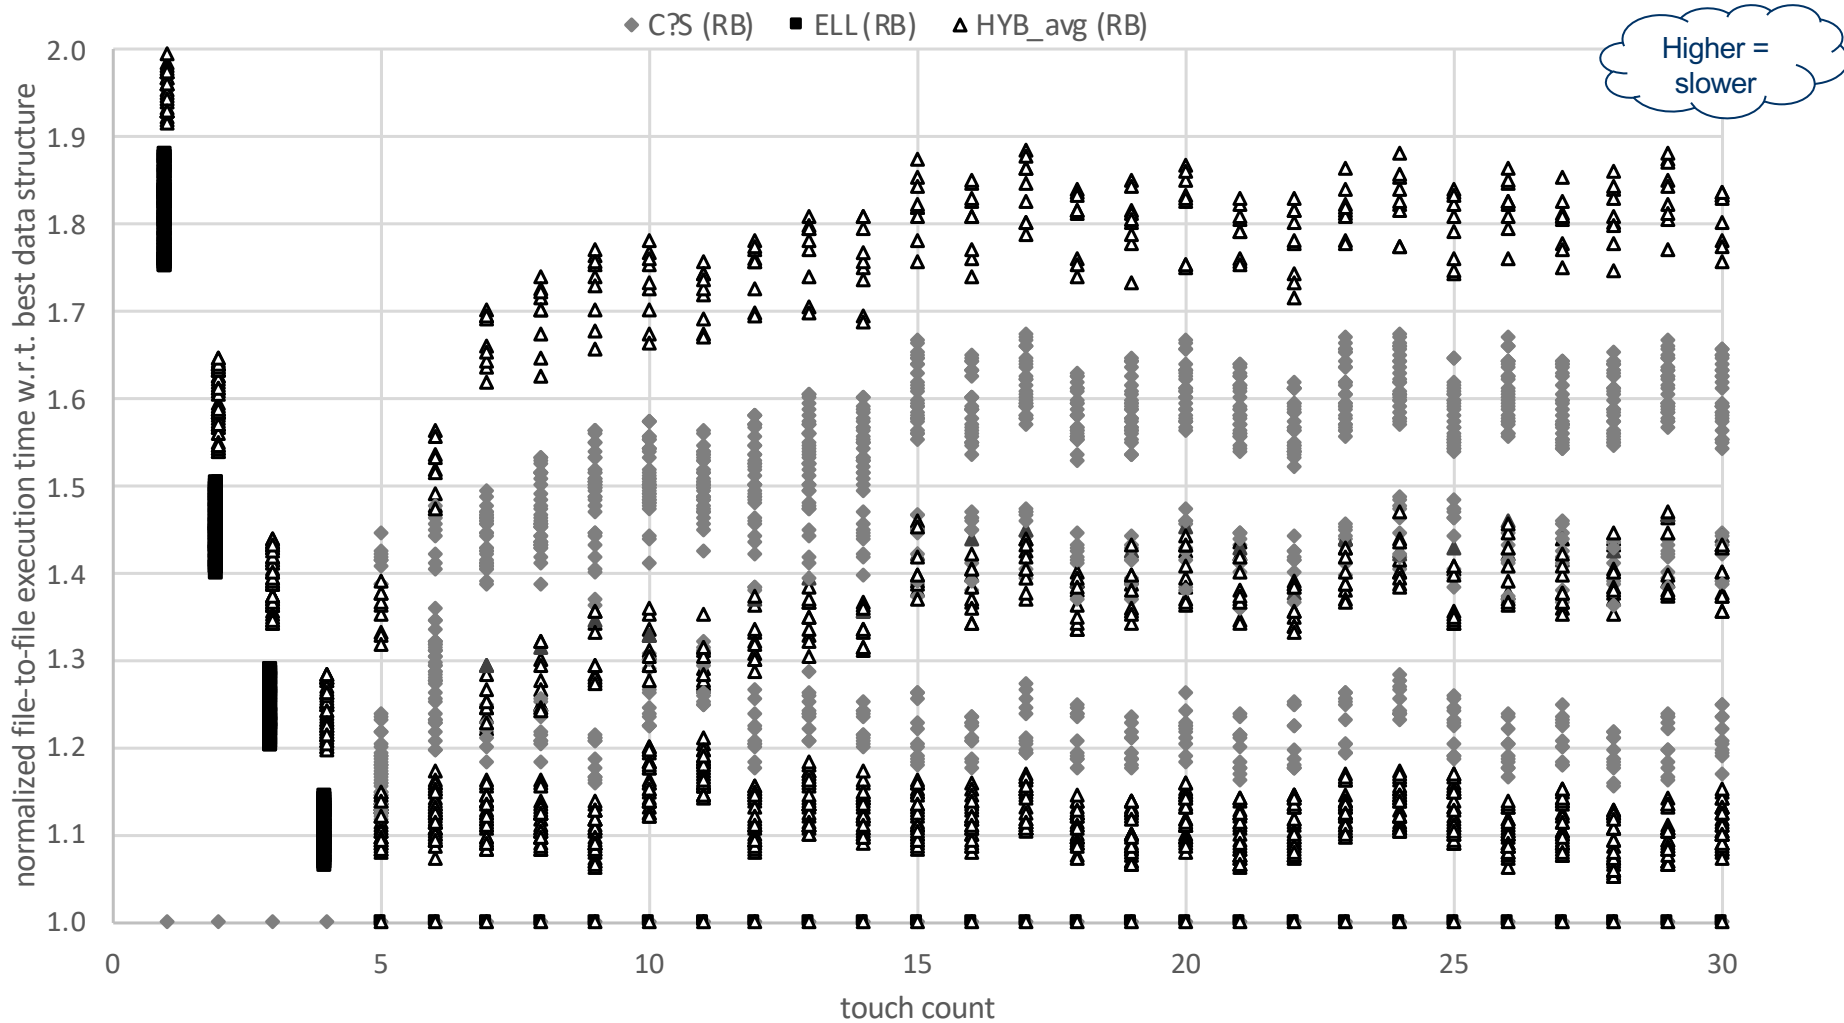
■ HYBRid format

- Better memory efficiency than *ELL*.
- Heuristics (avg, dstr) find smaller the row length.
- Additional successors are stored in *EL*.

<i>ell_v</i>	<i>ell_w</i>	<i>el</i>
$\begin{bmatrix} 1 \\ 3 \\ * \\ \mathbf{0} \\ 3 \\ 6 \\ * \end{bmatrix}$	$\begin{bmatrix} 5 \\ 3 \\ * \\ \mathbf{2} \\ 1 \\ 4 \\ * \end{bmatrix}$	$\left[\begin{matrix} \{3\} \\ \{2\} \\ \{6\} \end{matrix}, \begin{matrix} \{3\} \\ \{5\} \\ \{2\} \end{matrix} \right]$

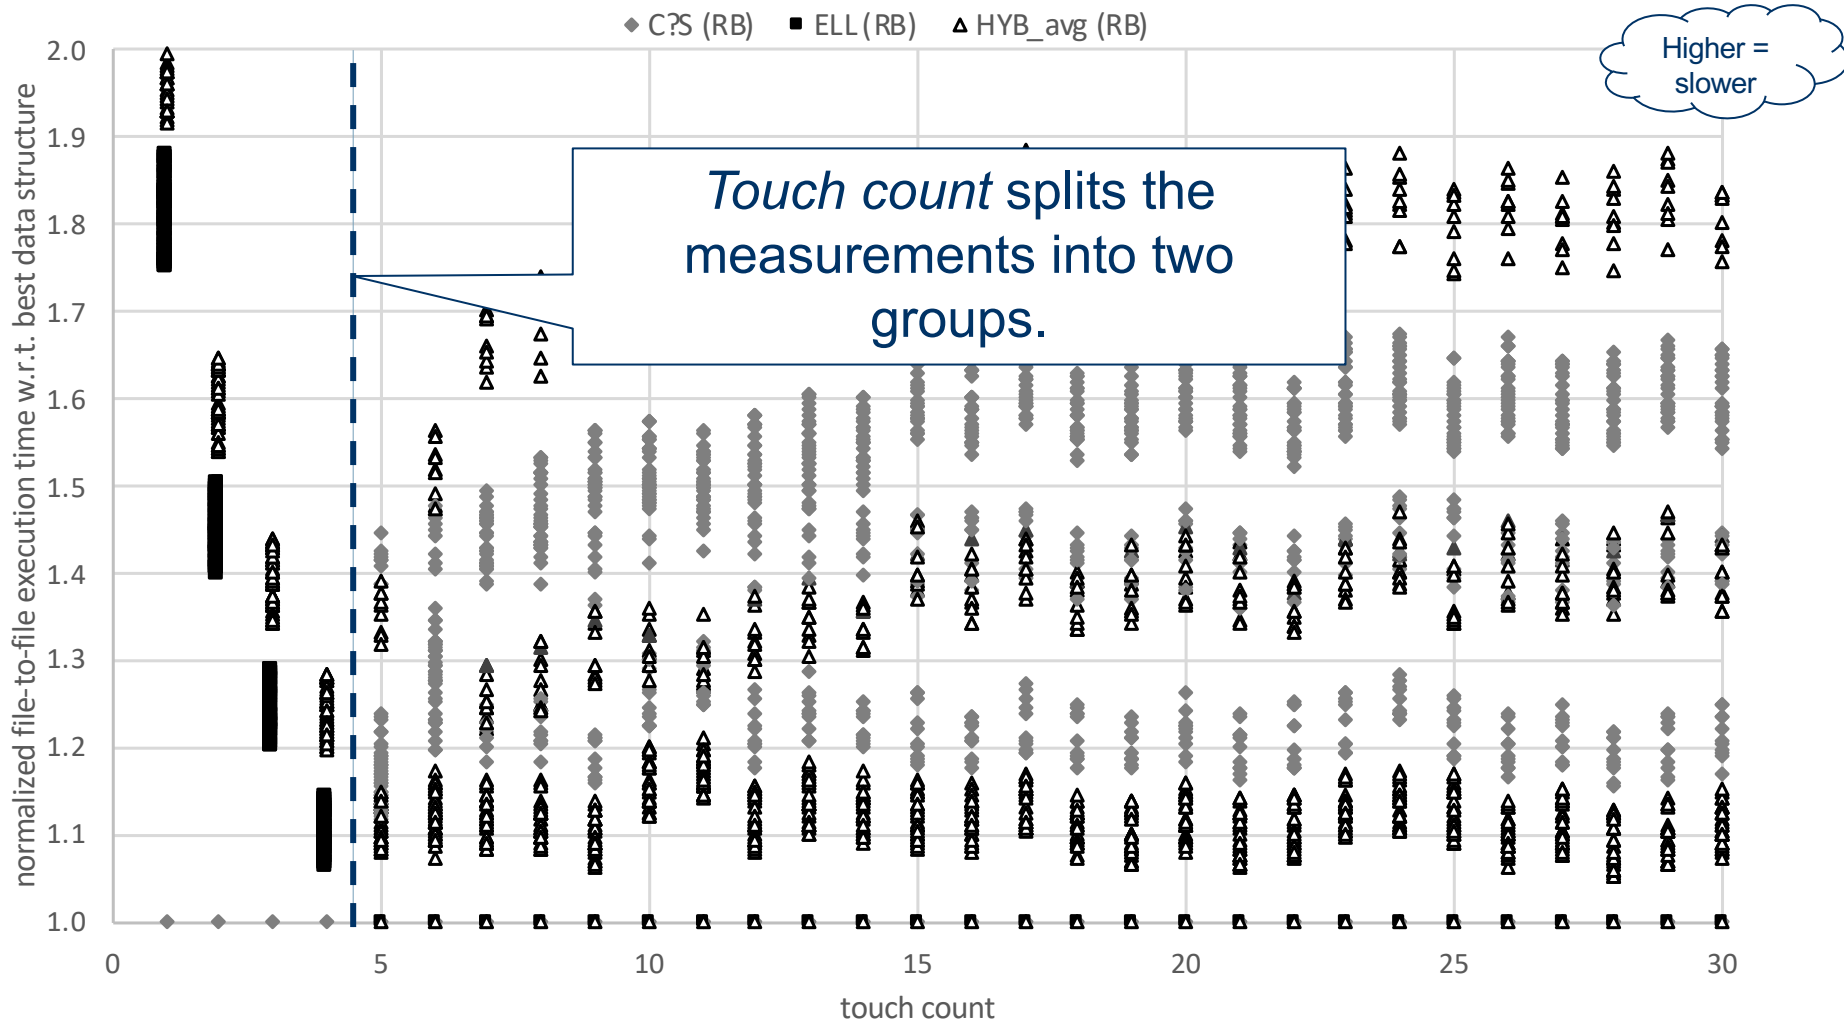
Decision Tree – Second Layer (III)

- Data structures cause costs for construction, shipping, ...



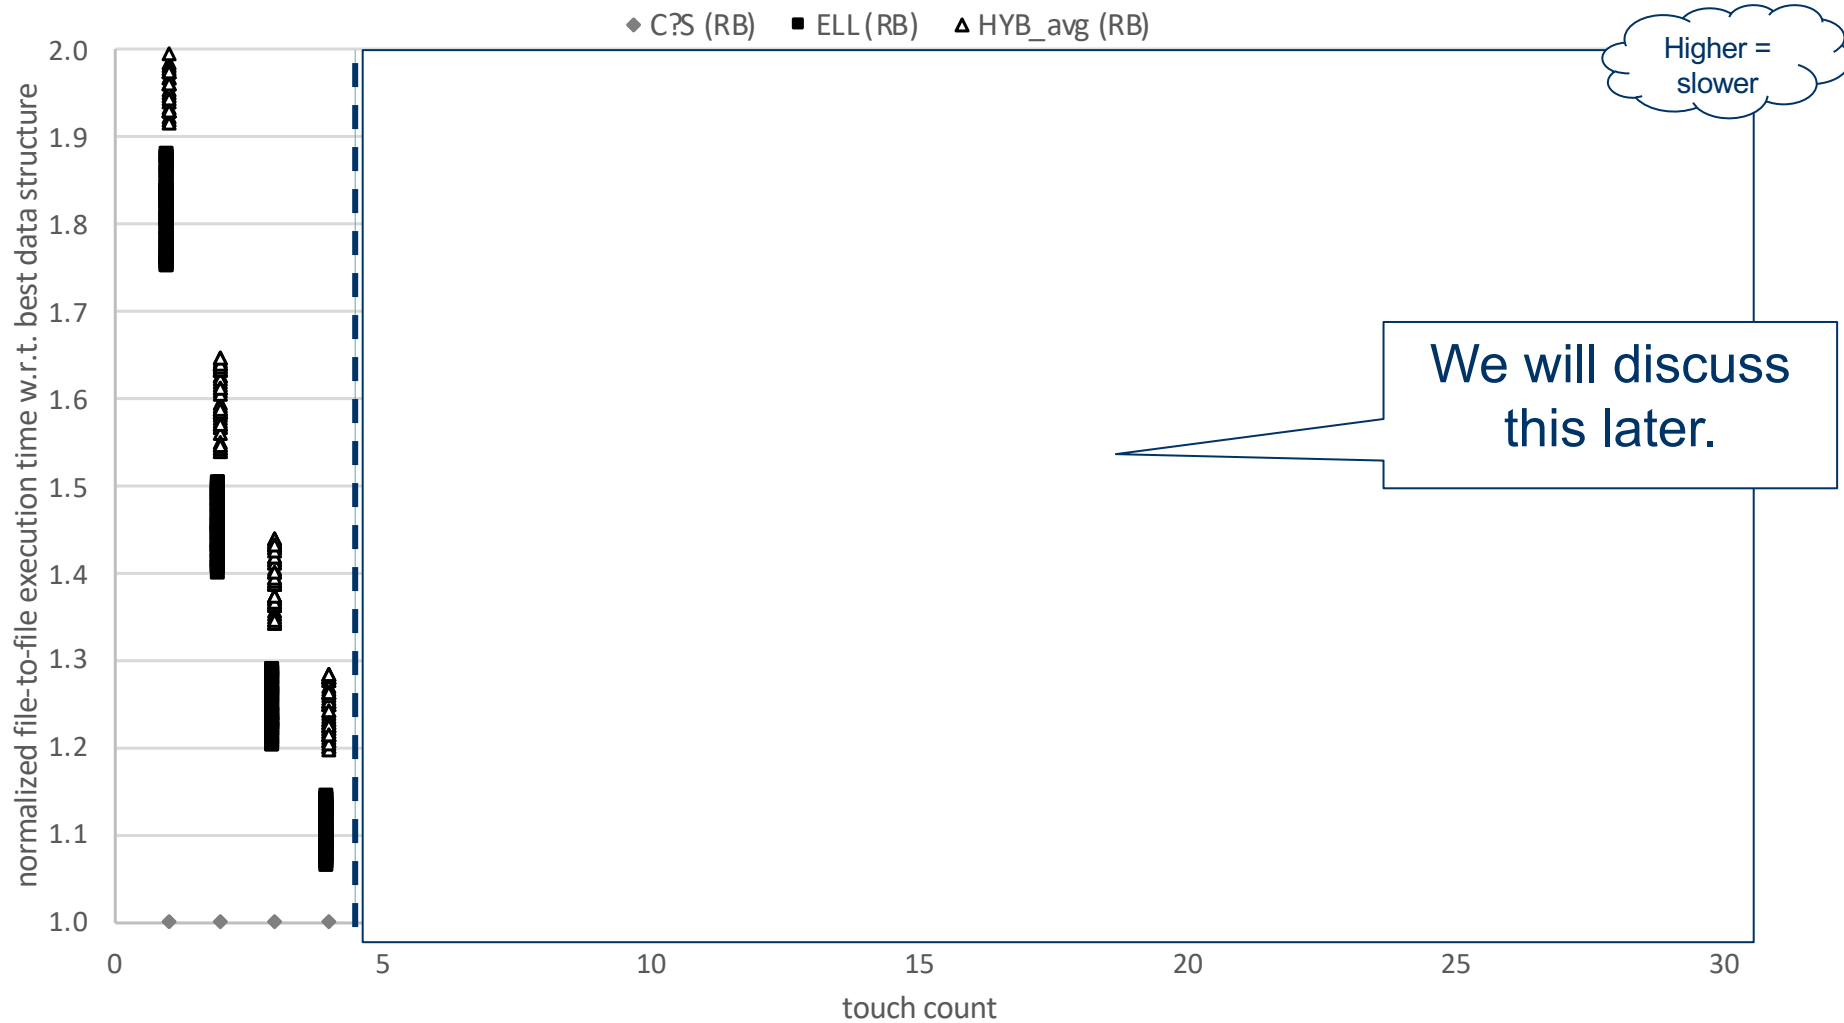
Decision Tree – Second Layer (III)

- Data structures cause costs for construction, shipping, ...



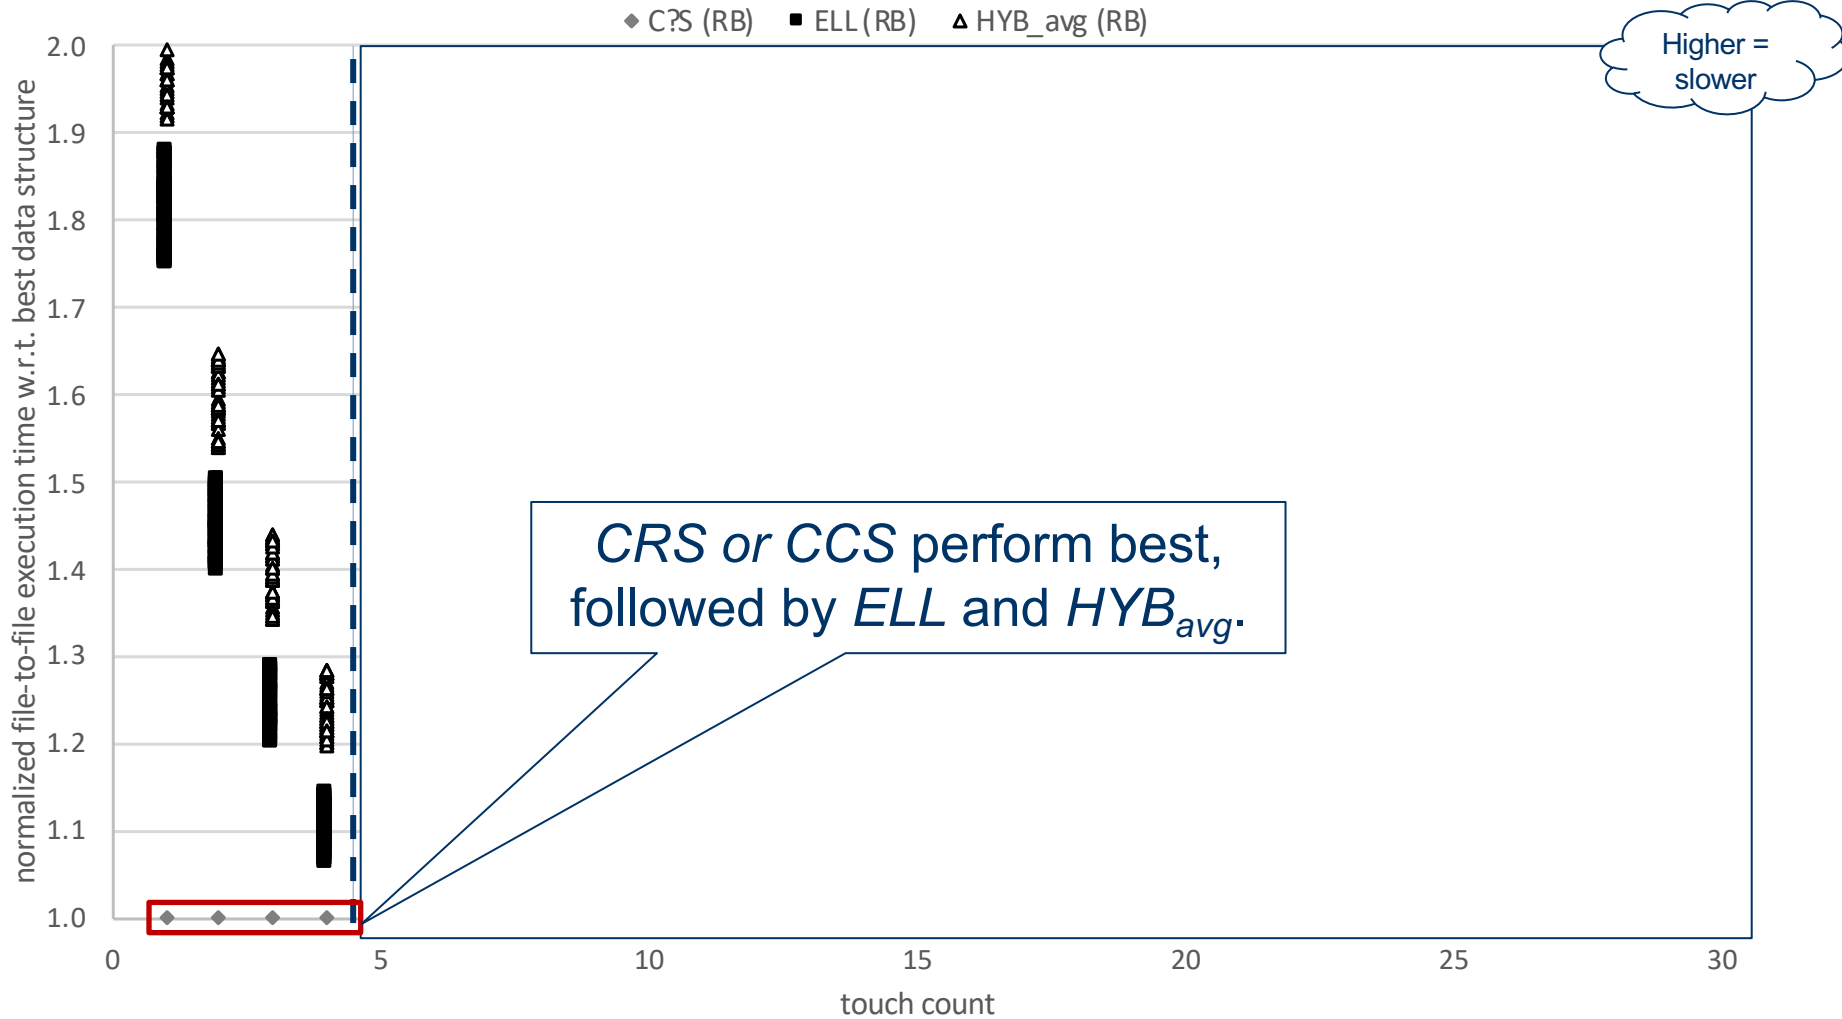
Decision Tree – Second Layer (III)

- Data structures cause costs for construction, shipping, ...



Decision Tree – Second Layer (III)

- Data structures cause costs for construction, shipping, ...



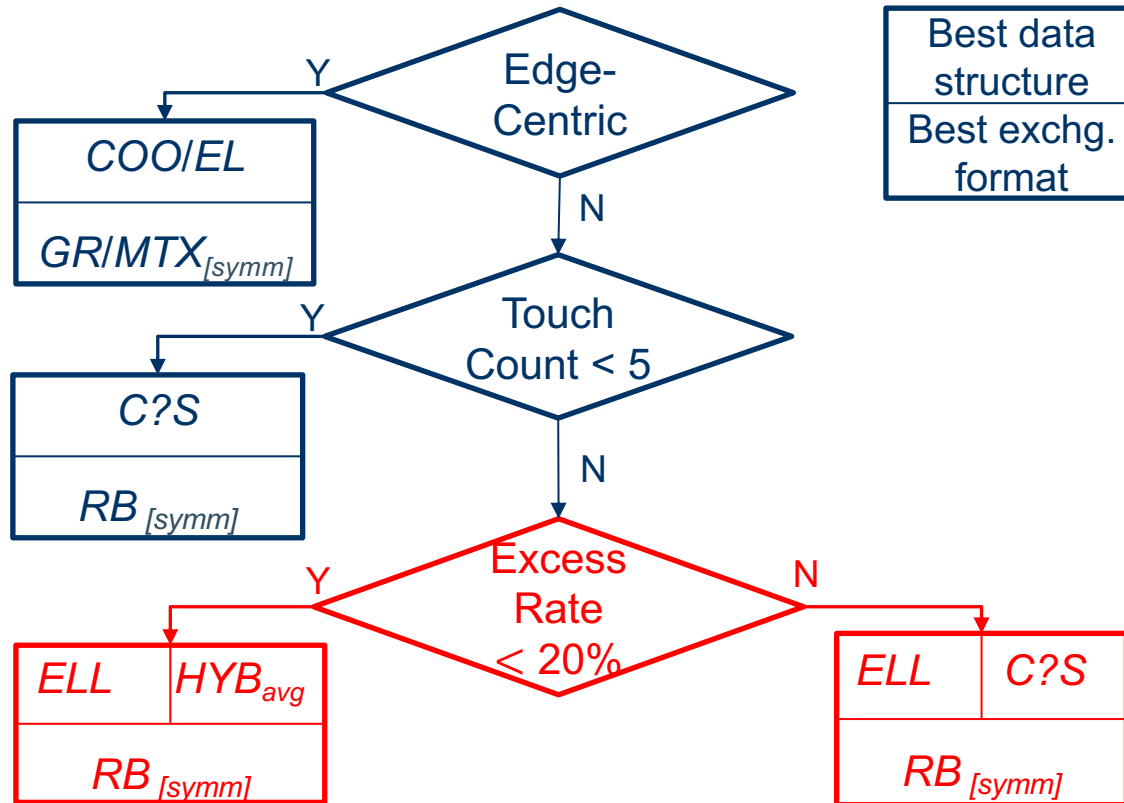
Decision Tree – Second Layer (IV)

- The Rutherford Boeing (RB) format is a dump of the *CRS* format.

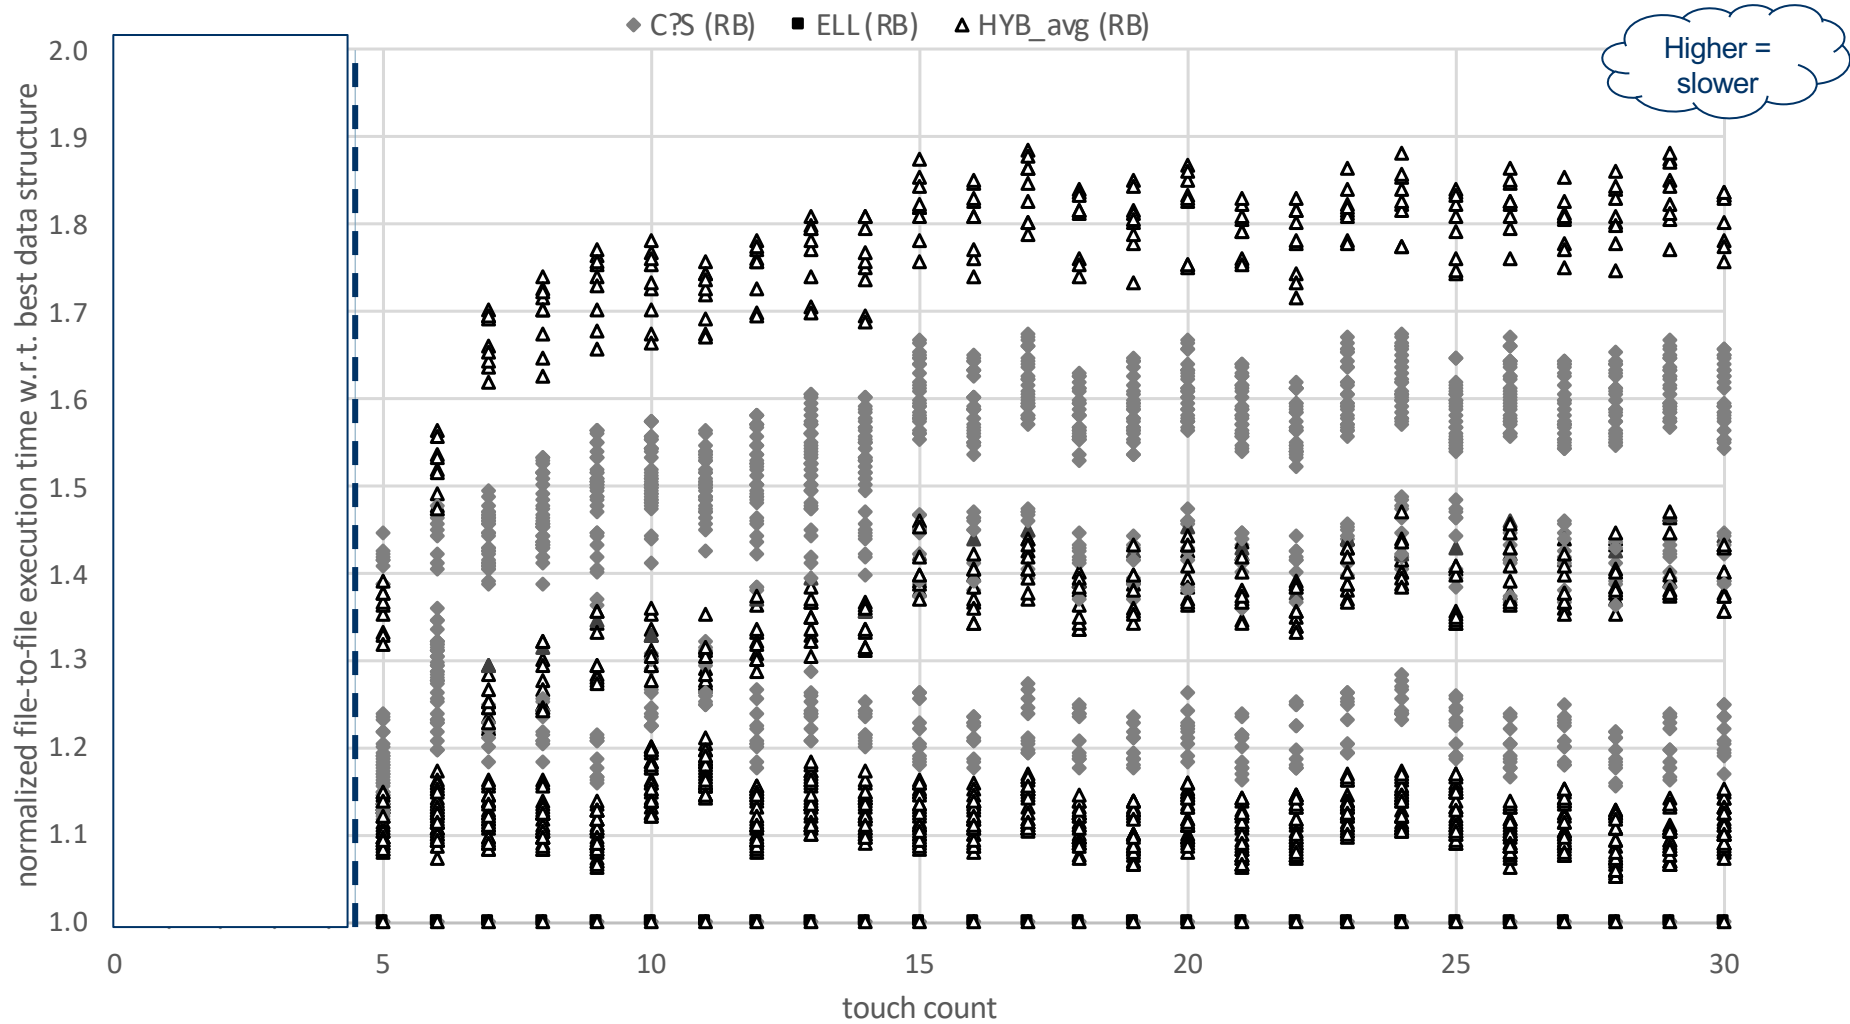
8	2	2	2	}	<i>Information about the structure of the encoded graph and the file.</i>
<i>iga</i>	7	7	7		
(11 4)	(11 3)			}	<i>src_{ptr}</i>
0	1	2	2		
5	6	7		}	<i>succ/pred</i>
1	3	0	2		
5	3	6		}	<i>weights</i>
5	3	2	6		
2	1	4			

- Other file formats are 40% slower in our measurements.

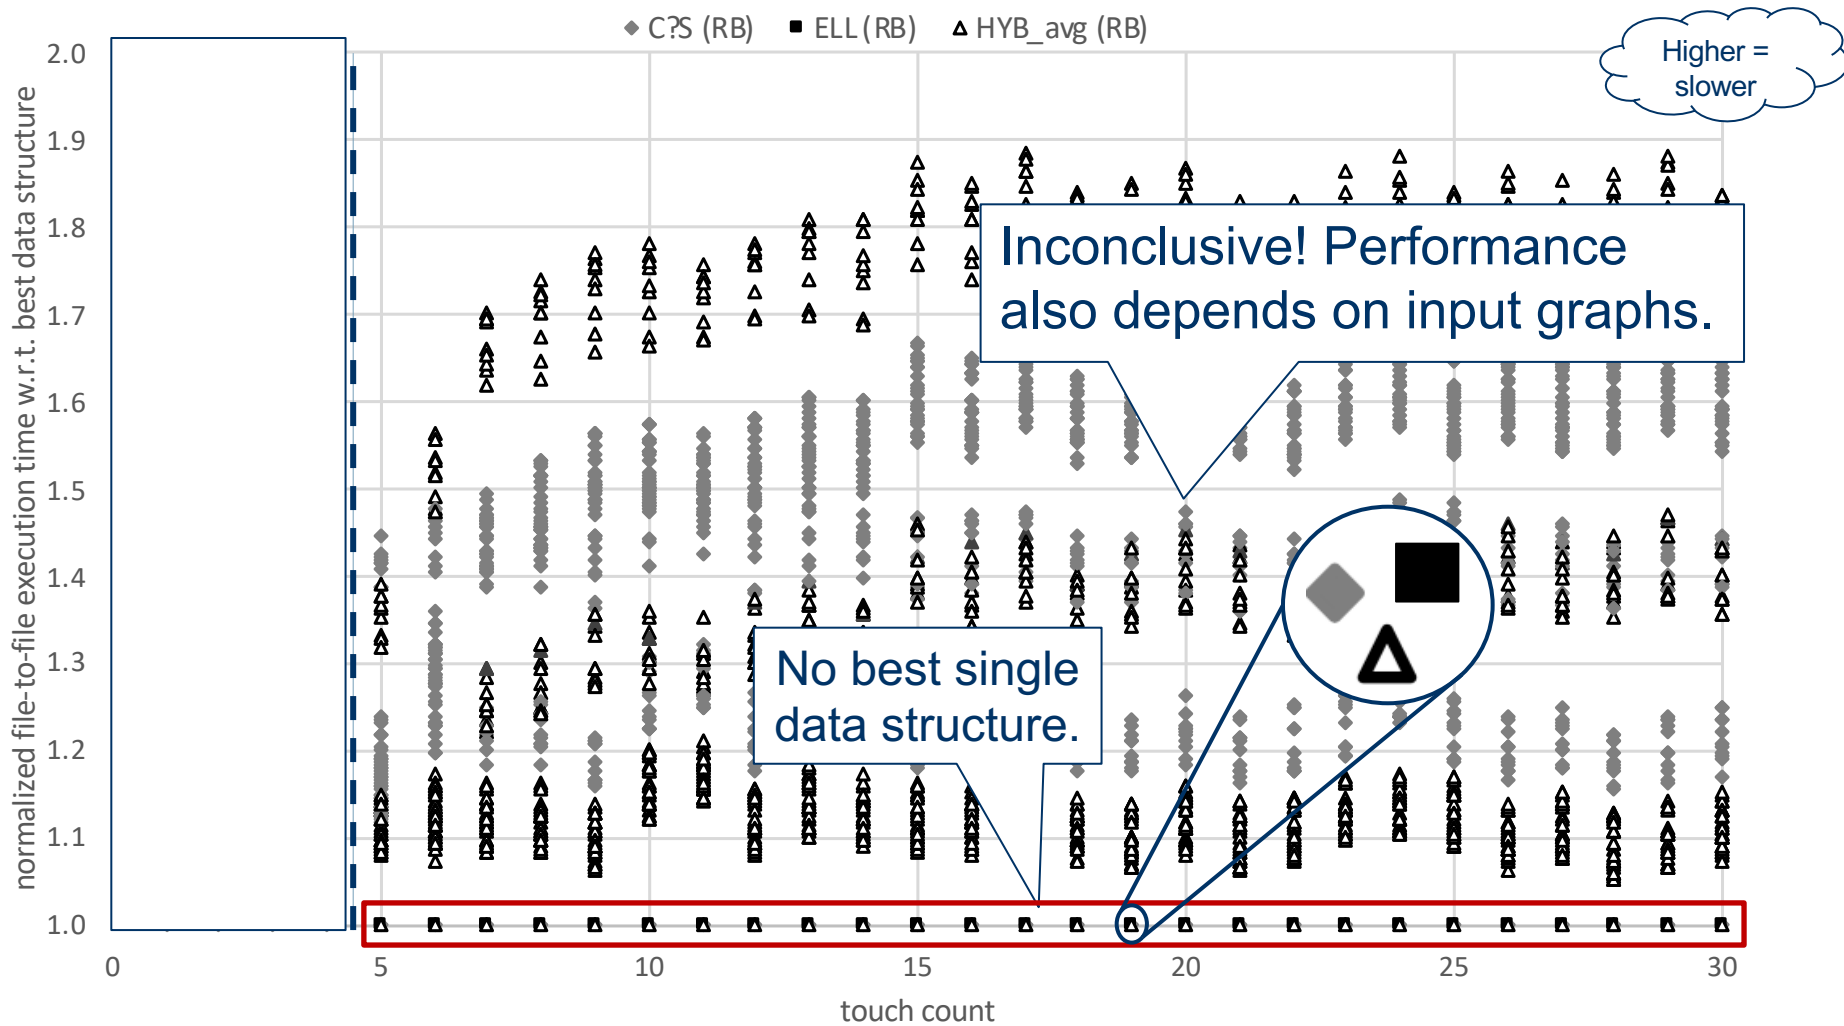
Decision Tree (file-to-file)



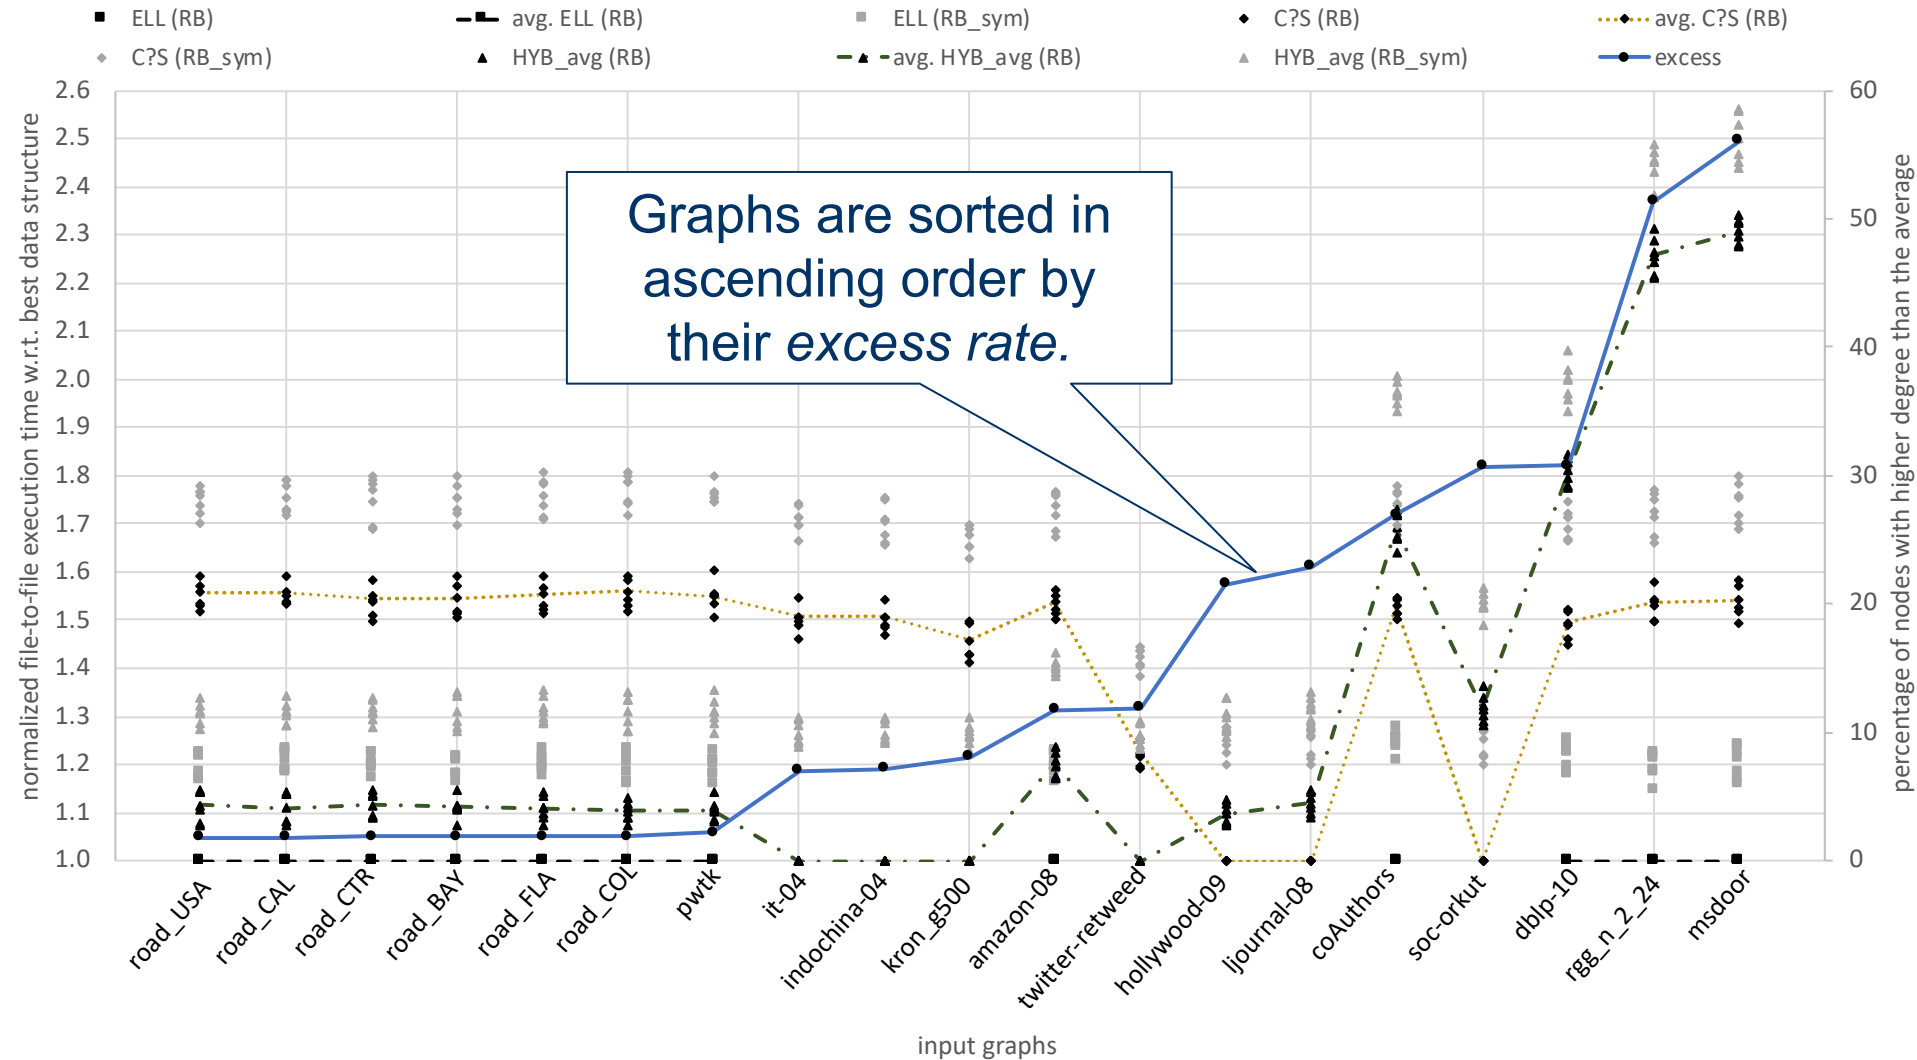
Decision Tree – Third Layer (I)



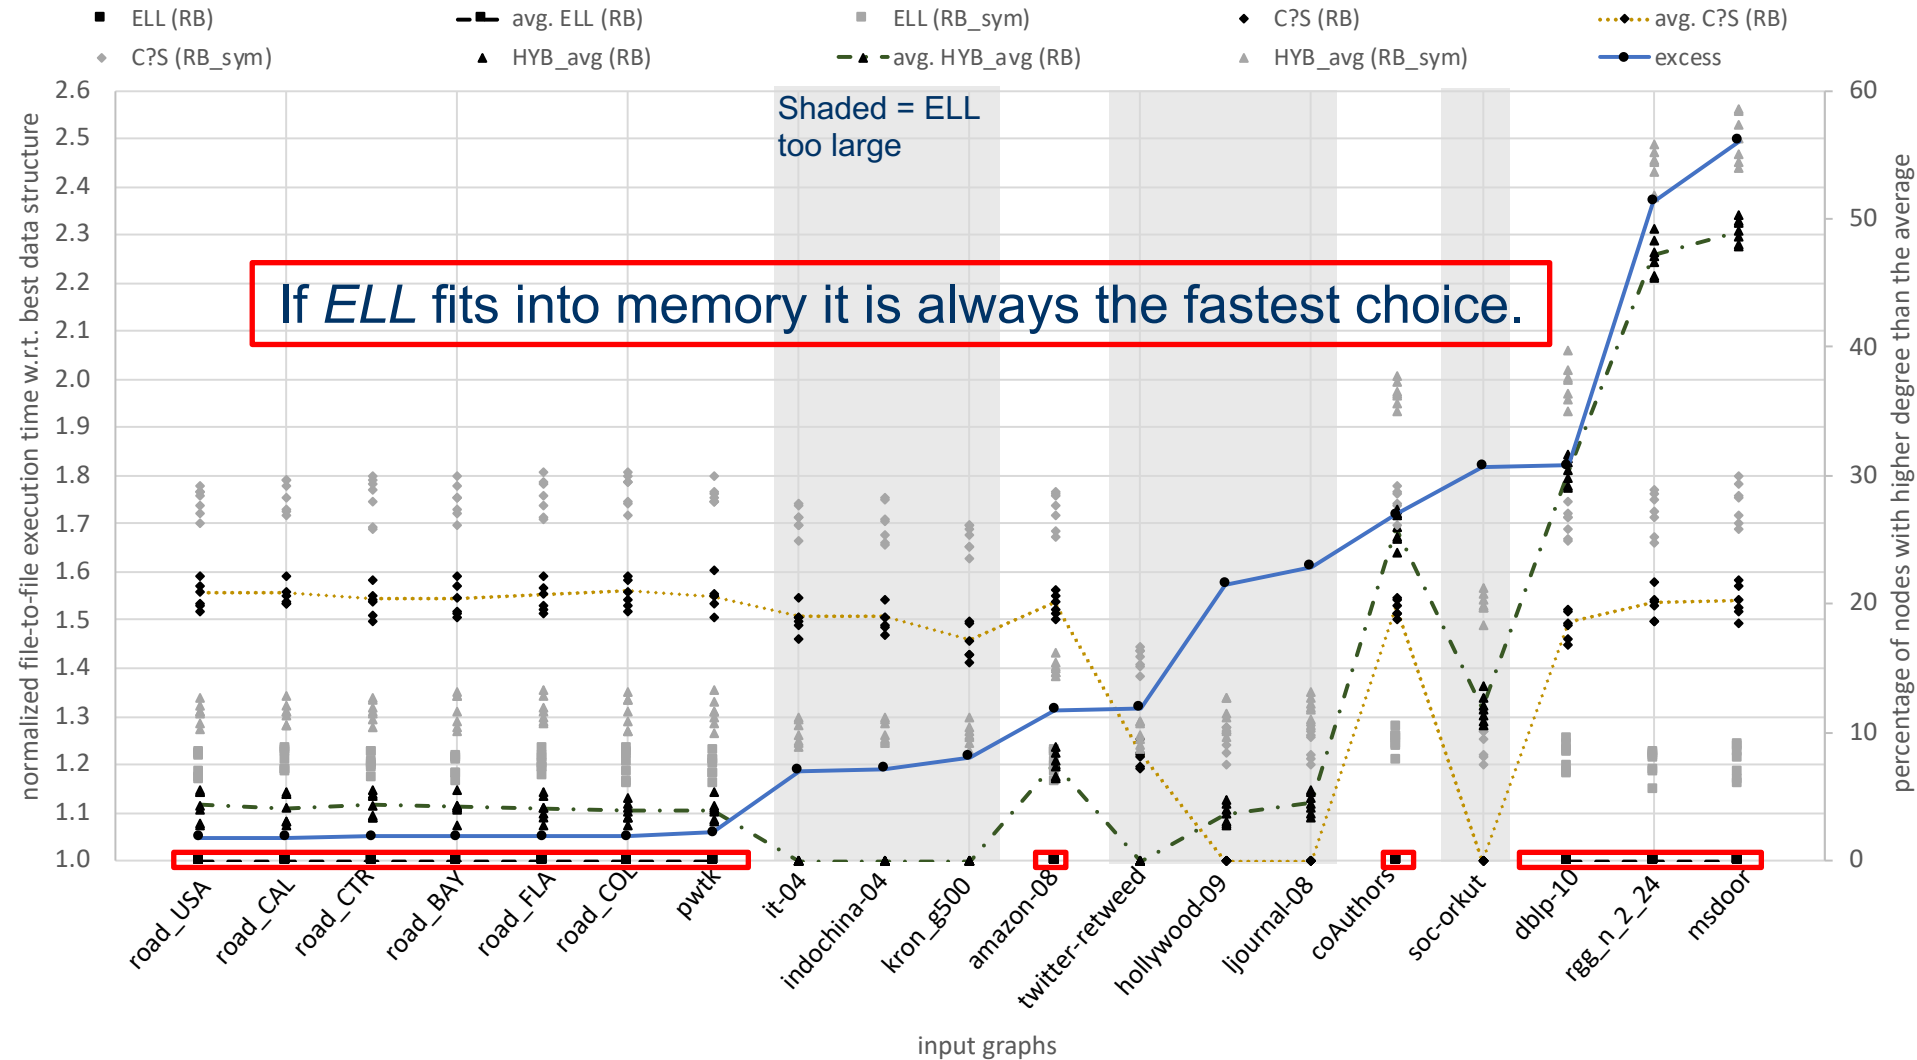
Decision Tree – Third Layer (I)



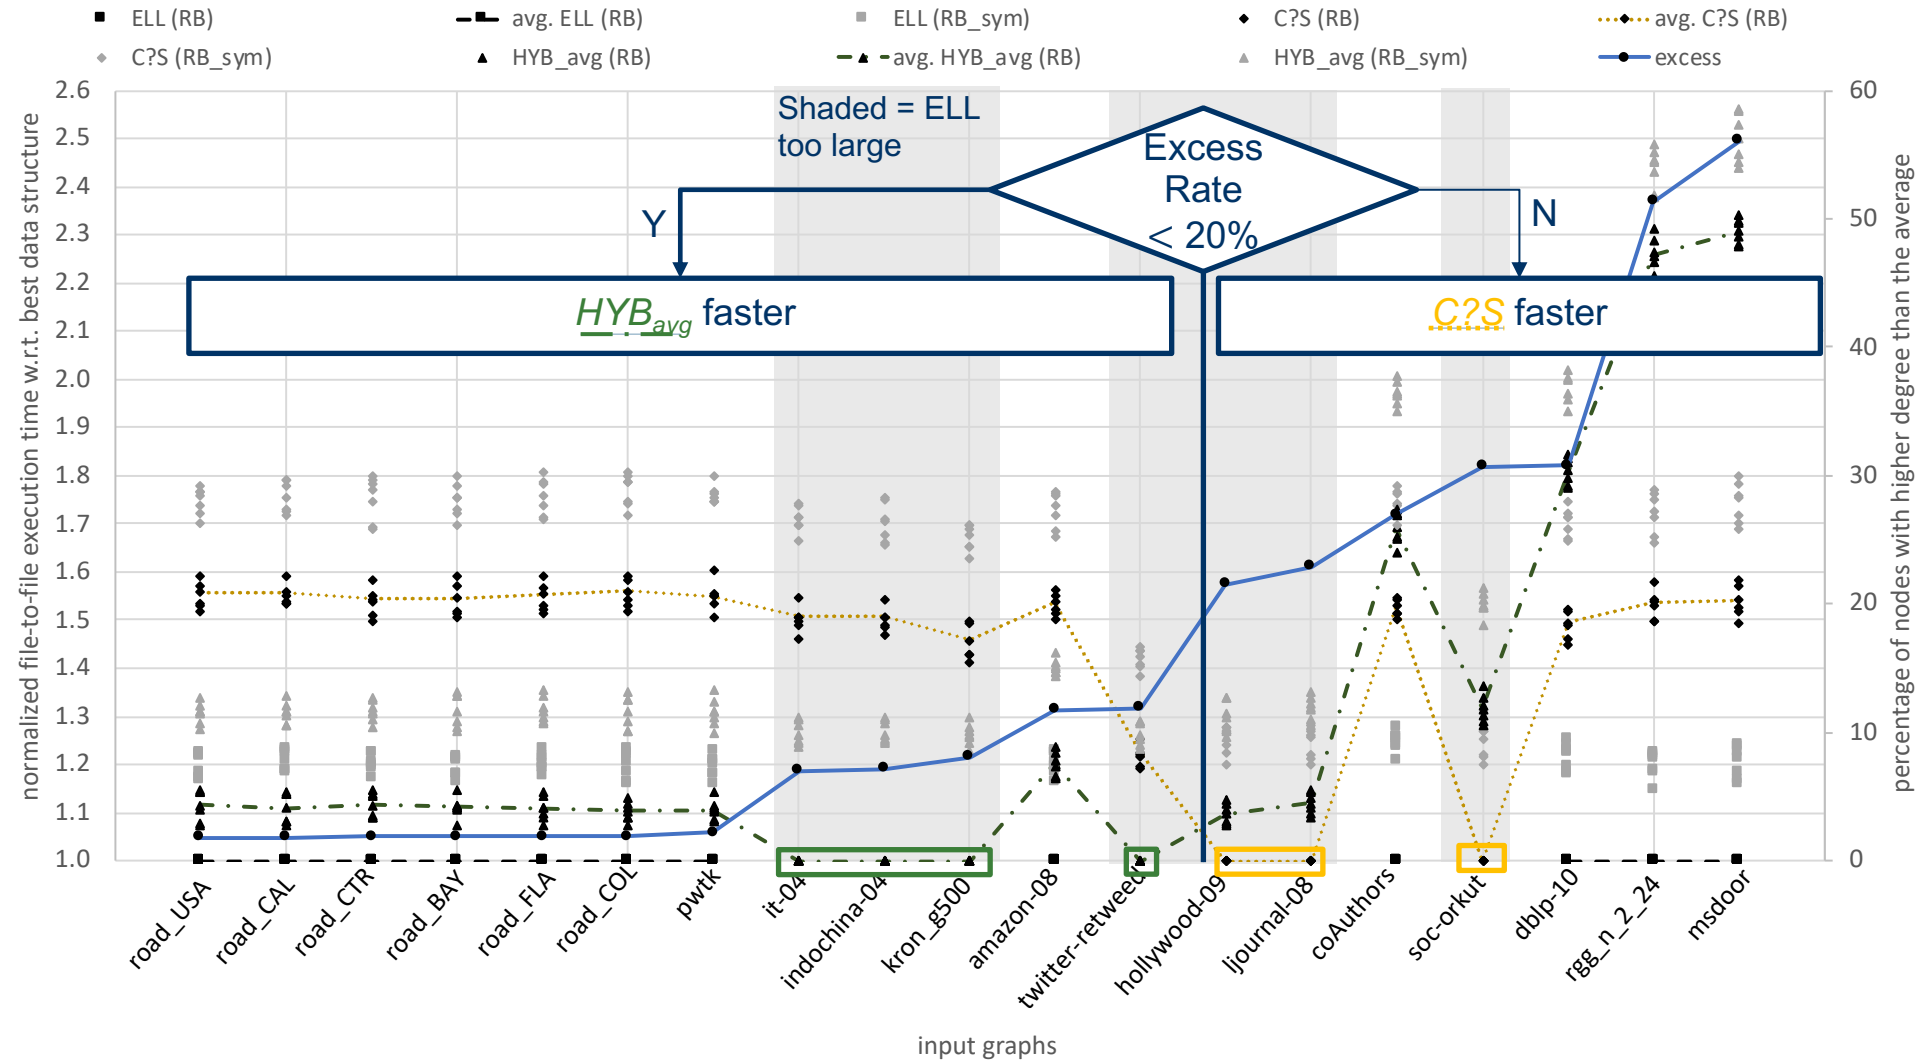
Decision Tree – Third Layer (II)



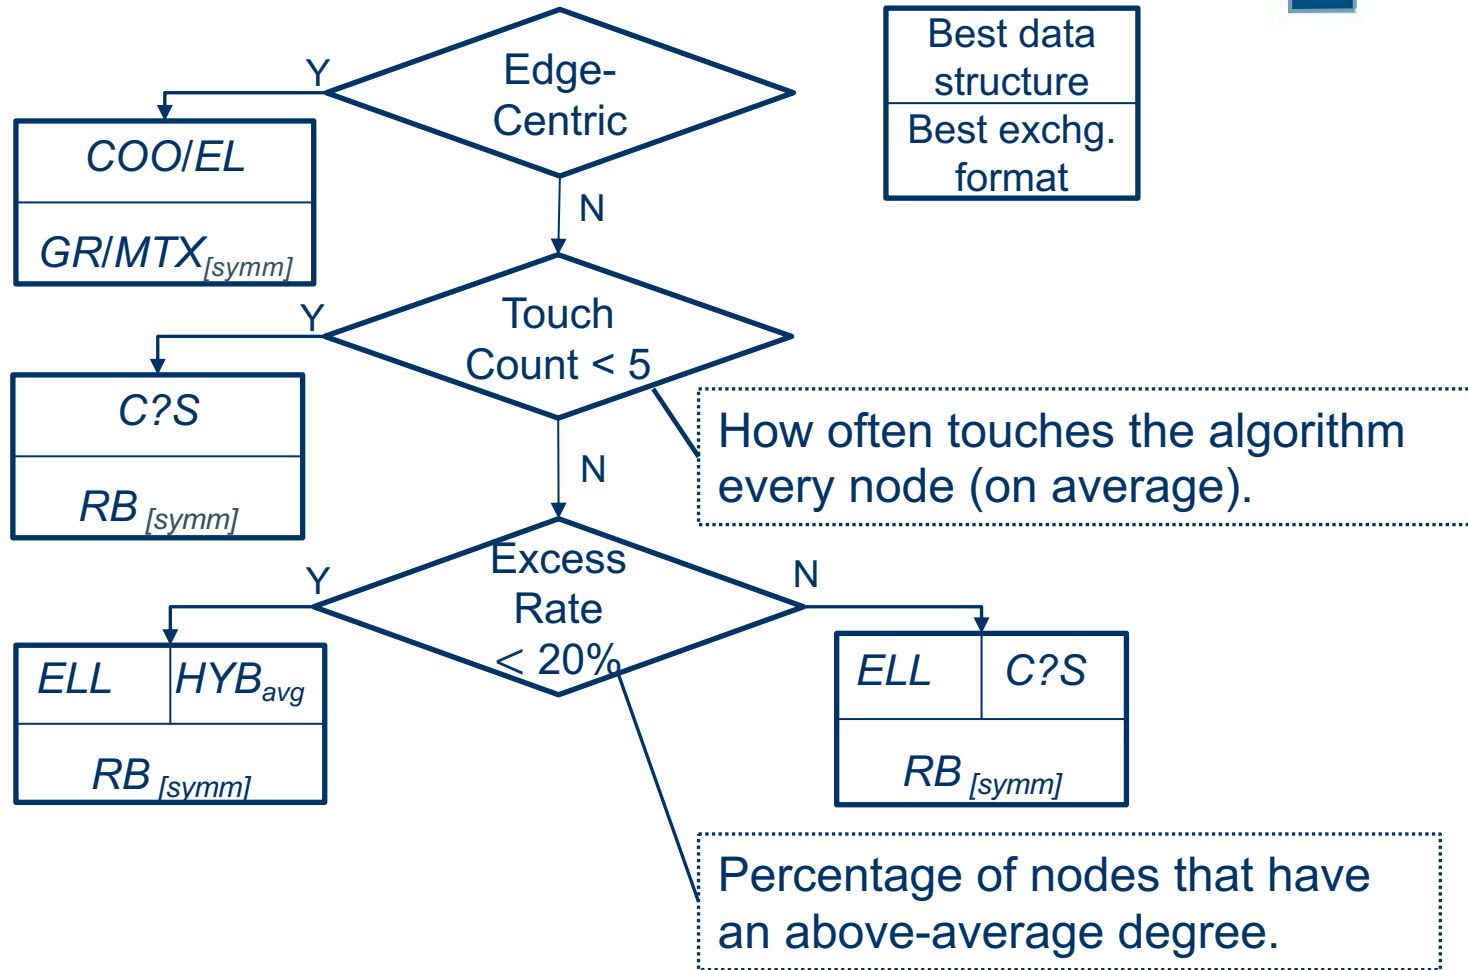
Decision Tree – Third Layer (II)



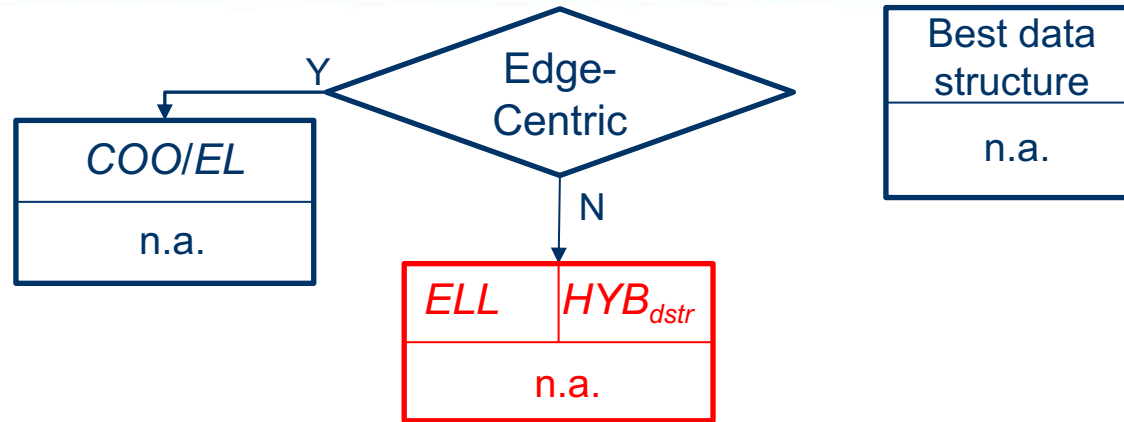
Decision Tree – Third Layer (II)



Decision Tree (file-to-file)

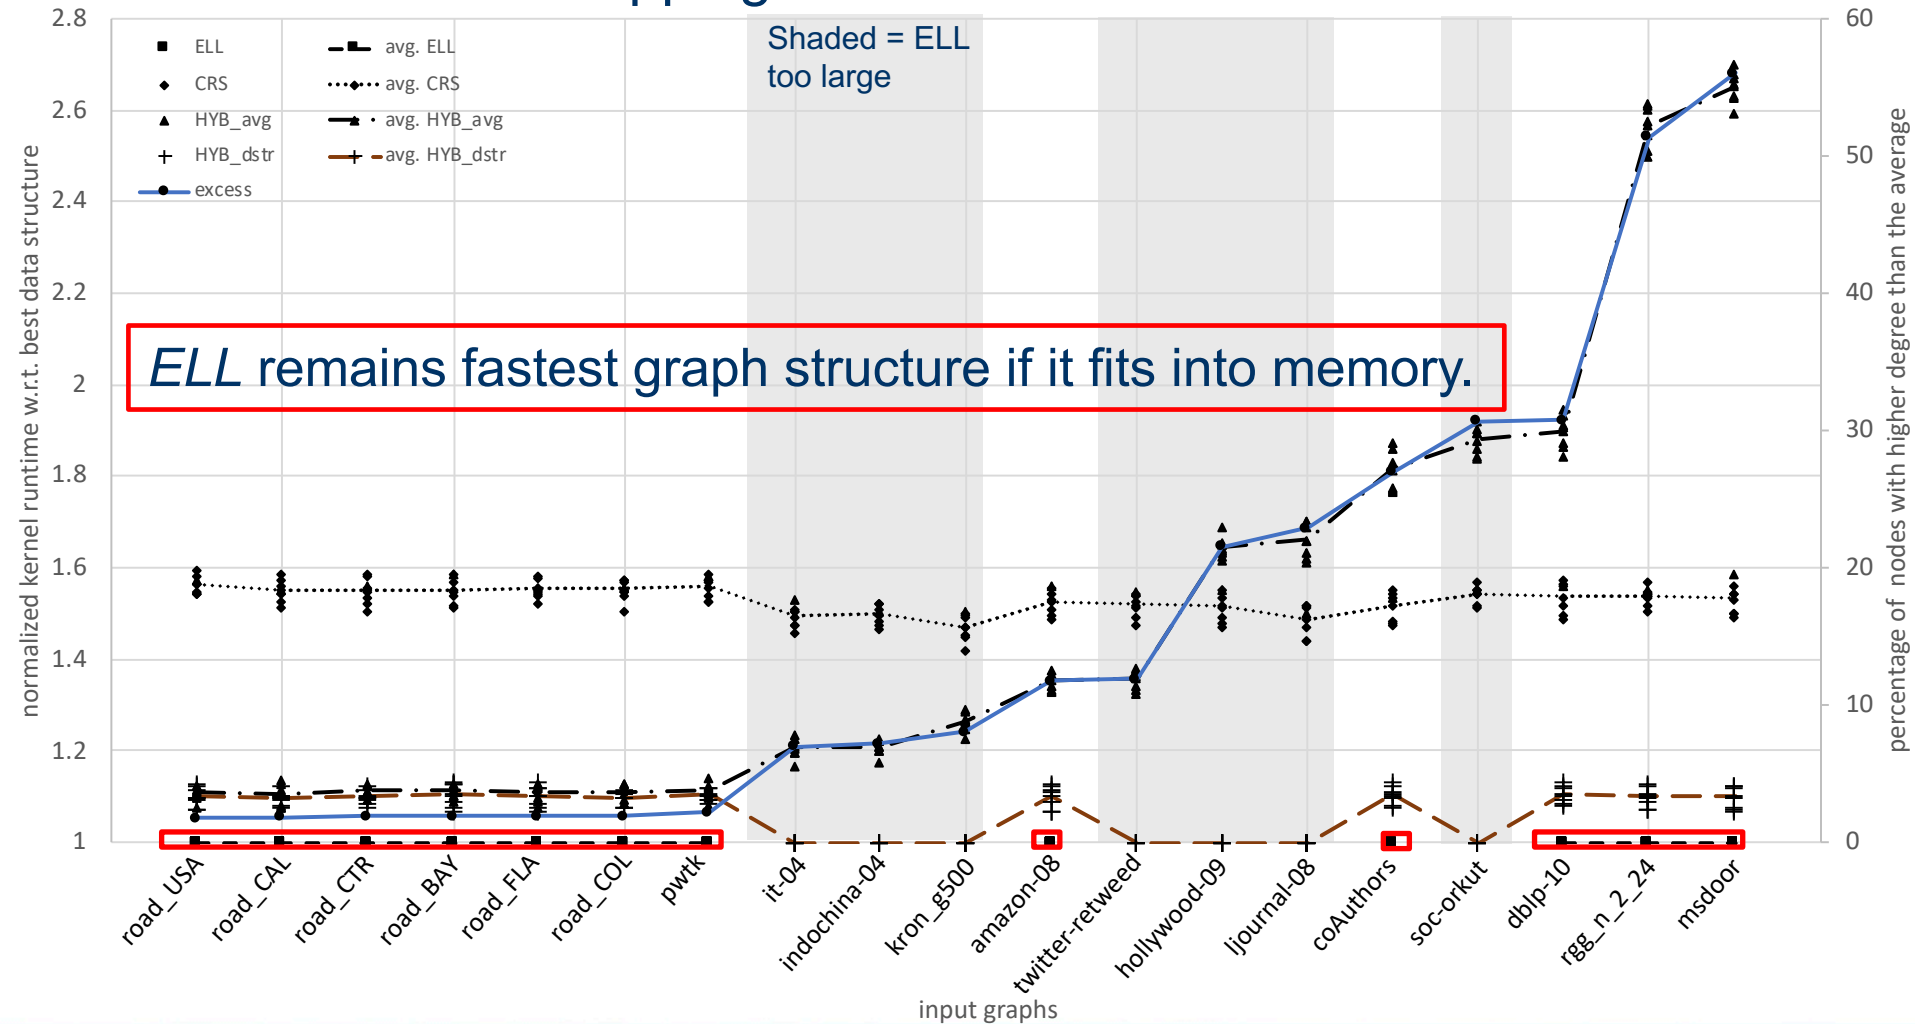


Decision Tree – Kernel only



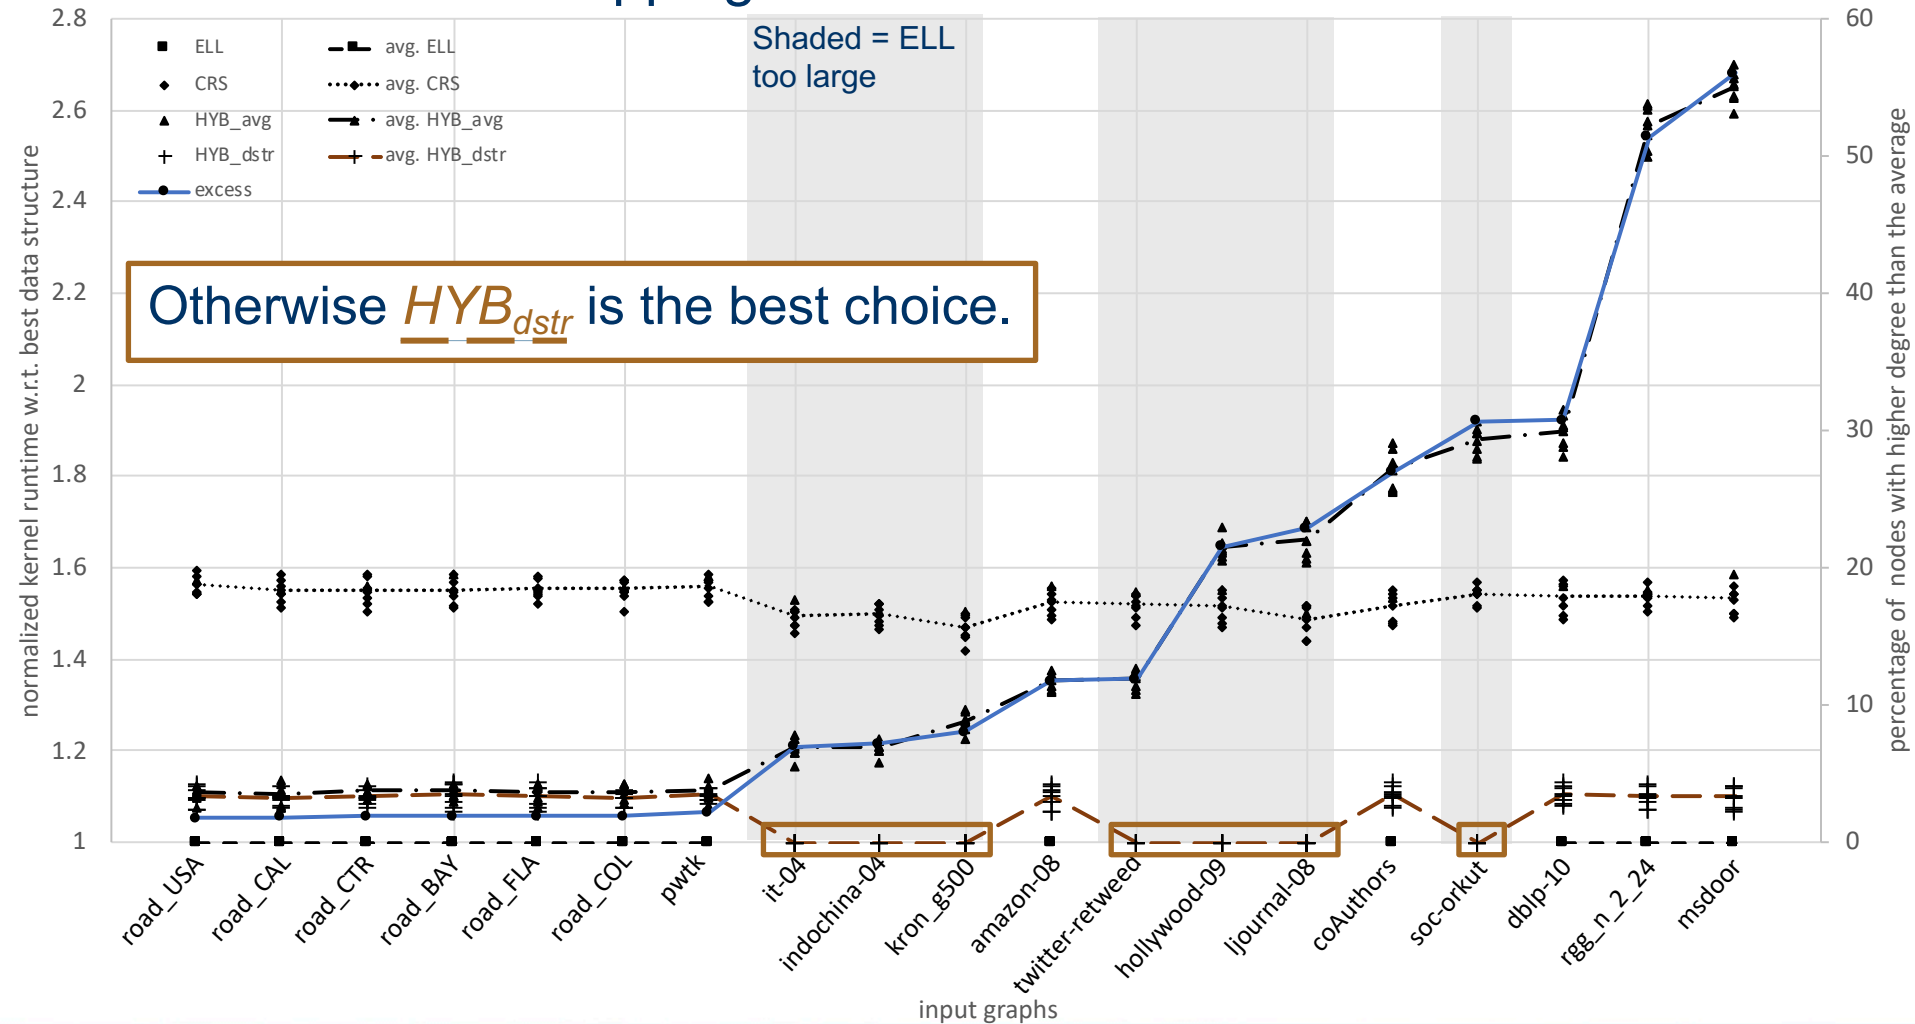
Kernel-only results

- Without I/O, data structure construction, and CPU \leftrightarrow GPU shipping.



Kernel-only results

- Without I/O, data structure construction, and CPU \leftrightarrow GPU shipping.



Conclusion

- There is no single best graph data structure for GPUs.
- Three decision layers:
 - Edge-centric vs. node-centric
 - Touch count threshold around 5
 - Excess rate threshold around 20%
- Our Decision Tree can help developers to pick an adequate data structure for their use case.
- Choosing the adequate data structure can speed up graph algorithms by up to 45%.

Conclusion

- There is no single best graph data structure for GPUs.
- Three decision layers:
 - Edge-centric vs. node-centric
 - Touch count threshold around 5
 - Excess rate threshold around 20%
- Our Decision Tree can help developers to pick an adequate data structure for their use case.
- Choosing the adequate data structure can speed up graph algorithms by up to 45%.

Thank you for your
attention!
Any questions?