

Speeding up compact planar graphs by using shallower trees[†]

Alexander Iribarra-Cortés^{1,2} **José Fuentes-Sepúlveda**¹
Diego Seco^{1,2} **Roberto Asín**¹

¹Department of Computer Science, Universidad de Concepción, Chile

²Millennium Institute for Foundational Research on Data, Chile

March 24, 2022

[†]This work was funded by ANID Millennium Science Initiative Program - Code ICN17_002 (1st and 3rd authors), PFCHA/Doctorado Nacional/2021-21211768 (1st author), PAI grant 77190038 (2nd author), and CYTED grant 519RT0579 (3rd author).

- 1 Introduction
- 2 Related work
 - Compact data structures
 - Succinct ordinal trees
 - Planar embeddings
- 3 Our contribution
 - Effect of the trees' topology
 - Computing shallower spanning trees
 - Alternative data structures
- 4 Experimental analysis
- 5 Conclusions and future work

- 1 Introduction
- 2 Related work
 - Compact data structures
 - Succinct ordinal trees
 - Planar embeddings
- 3 Our contribution
 - Effect of the trees' topology
 - Computing shallower spanning trees
 - Alternative data structures
- 4 Experimental analysis
- 5 Conclusions and future work

Introduction

Problem: Compact representation of planar graphs.

- A planar graph can be embedded into the plane with no edges crossing each other.
- A planar graph may have more than one planar embedding.
- Two planar embeddings of the same graph are differentiated by the orientation or order of their edges.






Figure: Two planar embeddings of the same planar graph.

Problem

- In some applications a specific embedding may be of special interest.
- A planar embedding is associated with a primal graph (**vertices**) and a dual graph (**faces**).
- It is wanted to be able to support queries of the primal graph and dual graph.



1 Introduction

2 Related work

- Compact data structures
- Succinct ordinal trees
- Planar embeddings

3 Our contribution

- Effect of the trees' topology
- Computing shallower spanning trees
- Alternative data structures

4 Experimental analysis

5 Conclusions and future work

Sequence of bits B supporting the following operations:

- $access(B, i)$: Return the value of the i -th bit.
- $rank_v(B, i)$: Return the amount of bits with value v until position i .
- $select_v(B, i)$: Return the position of the i -th occurrence of v .


$B = \underline{0110} \ 1101 \ 1010 \ 1100$

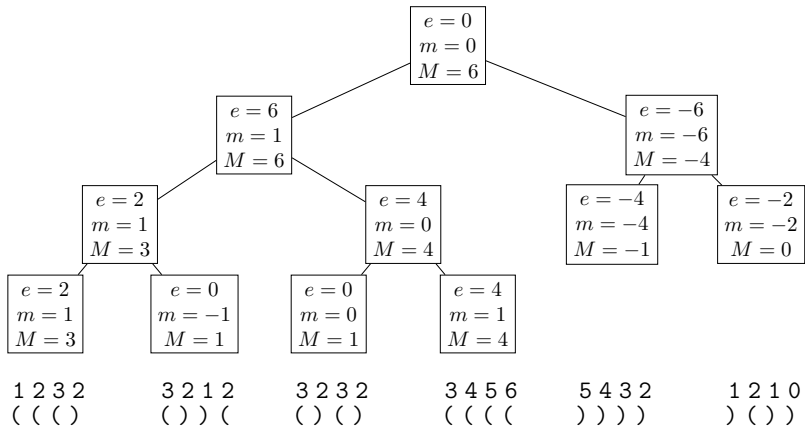
$$rank_1(B, 4) = 3$$

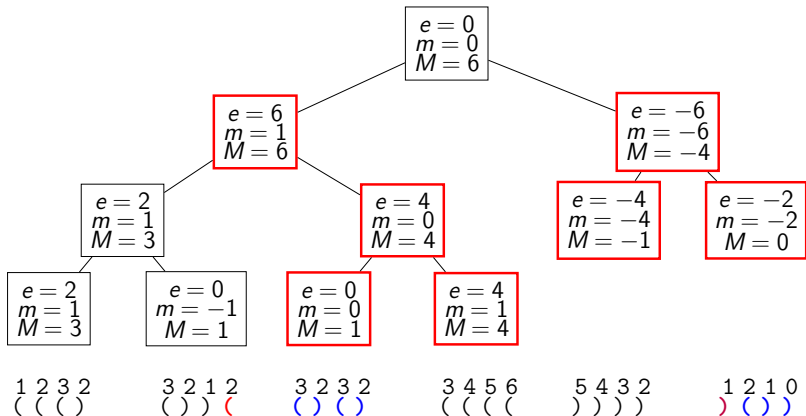
$$select_0(B, 3) = 6$$

Fully Functional (Navarro and Sadakane, 2014)

- Supports multiple operations (e.g. *match*, *enclose*, *height*, *etc*).
- Uses the auxiliary data structure *Range min-Max Tree*. It divides the sequence into blocks.
- Leaves store information about the blocks, while internal nodes store aggregated information about their children.







Range min Tree (Grossi and Ottaviano, 2015)


- Variation of the *Range min-Max Tree* in which the maximum is not stored.
- It uses other precomputed tables to speed up the search inside a block.
- Faster than the *Range min-Max Tree* in practice.

Compact planar embeddings

Year, Author(s)	Memory (bits)	Supports operations
1984, Turán	$4m$	No
1989, Jacobson	$O(m)$	Yes
1995, Keeler and Westbrook	$3.58m + O(1)$	No
2001, Munro and Raman	$2m + 8n + o(n)$	Yes
2010, Blelloch and Farzan	$3.58m + o(m)$	Yes
2020, Ferres <i>et al.</i>	$4m + o(m)$	Yes

Ferres *et al.* representation.

- Based on Turán's representation.
- Built with a spanning tree T of the primal graph.




Turán's representation

A counterclockwise DFS traversal over G is performed:

- When an edge $e \in T$ is visited, the symbol '(' is written the first time, and ')' the second time. This path is followed.
- When an edge $e \notin T$ is visited, the symbol '[' is written the first time, and ']' the second time.

4 bits per edge.




([((]) ([[)) (])]

Ferres *et al.* representation

The sequence is decomposed in three parts, exploiting the fact that $G - T$ is homologous to a spanning tree over the dual T' .

- A : A bitvector encoding the interleaving between T and T' .
- B : The tree T encoded as a balanced parentheses sequence.
- B^* : The tree T' encoded as a balanced parentheses sequence.
- The spanning tree T used in the representation is computed with a DFS traversal.



([((]) ([[])) (])]


A 1 0 1 1 0 1 1 0 0 1 1 1 1 0 1 0

B ((() ())) ()

B* [] [[]]

Spanning trees T and T'

- The spanning tree T is mostly used for queries on the primal graph.
- The spanning tree T' is mostly used for queries on the dual graph.
- Changing the topology of one of the trees will change the topology of the other tree.



	([(()	([[))	())]		
A	1	0	1	1	0	1	1	0	0	1	1	1	1	0	1	0
B	((()	()))	())	())))
B*	[)))))	[[)))]]]		

1 Introduction

2 Related work

- Compact data structures
- Succinct ordinal trees
- Planar embeddings

3 Our contribution

- Effect of the trees' topology
- Computing shallower spanning trees
- Alternative data structures

4 Experimental analysis

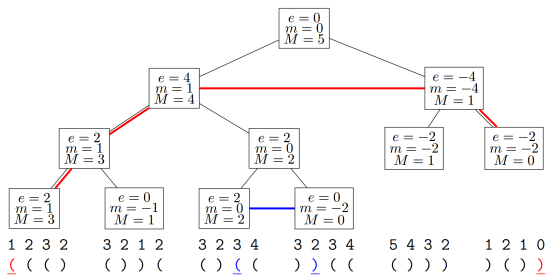
5 Conclusions and future work

Effect of the trees' topology

- Time complexity of *matching* query is $O(\log \frac{b-a}{B})$ when the matching parentheses are in positions a and b (Navarro, 2016, Section 7.2).

Effect of the trees' topology

- Time complexity of *matching* query is $O(\log \frac{b-a}{B})$ when the matching parentheses are in positions a and b (Navarro, 2016, Section 7.2).
- Therefore, to speed up the operations we must reduce the distances.



Computing shallower spanning trees

We evaluated the following methods to compute the spanning trees

Computing shallower spanning trees

We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.

Computing shallower spanning trees

We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.
- $\text{BFS}_{\text{primal}}$: BFS traversal over the vertices of the primal graph.

We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.
- $\text{BFS}_{\text{primal}}$: BFS traversal over the vertices of the primal graph.
- BFS_{dual} : BFS traversal over the faces of the dual graph.

We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.
- $\text{BFS}_{\text{primal}}$: BFS traversal over the vertices of the primal graph.
- BFS_{dual} : BFS traversal over the faces of the dual graph.
- DFS_x : DFS traversal over the vertices of the primal graph, limiting the height to x times the height obtained with the method $\text{BFS}_{\text{primal}}$.


We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.
- $\text{BFS}_{\text{primal}}$: BFS traversal over the vertices of the primal graph.
- BFS_{dual} : BFS traversal over the faces of the dual graph.
- DFS_x : DFS traversal over the vertices of the primal graph, limiting the height to x times the height obtained with the method $\text{BFS}_{\text{primal}}$.
- **Alt**: Alternate computation of BFS traversal over the primal graph and the dual graph.


We evaluated the following methods to compute the spanning trees

- $\text{DFS}_{\text{primal}}$: DFS traversal over the vertices of the primal graph.
- $\text{BFS}_{\text{primal}}$: BFS traversal over the vertices of the primal graph.
- BFS_{dual} : BFS traversal over the faces of the dual graph.
- DFS_x : DFS traversal over the vertices of the primal graph, limiting the height to x times the height obtained with the method $\text{BFS}_{\text{primal}}$.
- Alt: Alternate computation of BFS traversal over the primal graph and the dual graph.
- DegHeur: Heuristical traversal using a degree ordered heap.

Computing shallower spanning trees



Distances in B



Distances in B^*

Figure: Distance histograms using different methods on the dataset *tiger_map*.

Average distances

Dataset	DFS _{primal}		BFS _{primal}		BFS _{dual}		Alt		DFS ₂		DegHeur	
	<i>B</i>	<i>B</i> *	<i>B</i>	<i>B</i> *	<i>B</i>	<i>B</i> *	<i>B</i>	<i>B</i> *	<i>B</i>	<i>B</i> *	<i>B</i>	<i>B</i> *
PE1M	263.96	345.84	0.57	5.65	3.11	1.30	0.57	5.65	2.34	6.42	4.78	12.30
PE5M	1261.52	1612.13	1.17	12.07	6.03	2.68	1.17	12.07	5.12	13.22	7.39	33.49
PE10M	2547.27	3265.26	1.86	20.55	9.69	4.29	1.86	20.55	7.01	16.41	11.88	47.00
PE25M	6440.49	8246.12	2.87	32.22	15.16	6.66	2.87	32.22	11.23	25.67	19.39	102.99
tiger_map	1882.50	2482.99	0.27	26.06	5.80	1.96	0.27	26.06	2.03	28.15	2.07	44.82

Table: Average parentheses distance for B and B^* . Values are multiplied by 10^{-3} .

Using alternative data structures

- We tested using newer CDS for balanced parentheses sequences, specifically the range min tree (Grossi and Ottaviano, 2015).
- We used both the implementation available in the SUCCINCT library and a variation of the range min-max tree of the SDSL library which doesn't store the maximums*.

*This is not a full implementation of the range min tree

- 1 Introduction
- 2 Related work
 - Compact data structures
 - Succinct ordinal trees
 - Planar embeddings
- 3 Our contribution
 - Effect of the trees' topology
 - Computing shallower spanning trees
 - Alternative data structures
- 4 Experimental analysis
- 5 Conclusions and future work

- We modified Ferres *et al.* implementation to plug in the approaches already described.
- The base implementation uses the C++ SDSL library (Gog et al., 2014).
- We also tested using the range min tree implementation in the SUCCINCT library.

- The code was compiled using GCC 4.8.4 with the optimization flag `-O3`.
- We measured time using the `clock` function.
- The experiments were carried out on a machine with Intel Core i7-3820 processor clocked at 3.60 GHz. This processor has per-core L1 and L2 caches of 32KB and 256KB respectively, and a shared cache of 10MB. The machine runs Linux 3.13.0-86-generic and has 32GB of DDR3 RAM.

We used both synthetic and real datasets with different numbers of nodes, which are available at

<http://www.inf.udec.cl/~jfuentess/datasets/graphs.php>

Dataset	Vertices	Edges
PE1M	1.000.000	2.999.978
PE5M	5.000.000	14.999.983
PE10M	10.000.000	29.999.979
PE25M	25.000.000	74.999.979
tiger_map	19.785.187	43.903.023

Datasets used in the experimental analysis



Implementation	Description
succinct	Uses only CDS from succinct
SDSL	Uses only CDS from SDSL
minTree	Uses the <i>range min tree</i> from succinct, and SDSL on the remaining CDS
v5	Variation that uses <i>rank_support_v5</i> from SDSL for <i>rank</i> support
SDSL _{min}	Modification of the <i>range min max tree</i> from SDSL which doesn't store the maximums

- We evaluated operations *degree*, *listing*, *face* and *dfs*.
- All the operations were performed on the primal graph and dual graph.
- We averaged 15 repetitions over all vertices/edges except for operation *dfs*, in which we averaged over 5 starting vertices.
- We measured using three starting edges for the construction, and we report the median of the averages.


Results (tiger_map)

Method	Queries on the primal				Queries on the dual			Space
	dfs	degree	neighbors	face	dfs	degree	neighbors	
BFS _{primal} minTree	13.00	0.195	0.540	<u>0.608</u>	14.09	0.244	0.610	30.83
BFS _{primal} SDSL	16.59	0.142	0.621	0.718	18.89	0.197	0.724	30.25
BFS _{primal} SDSL _{min}	16.10	<u>0.137</u>	0.594	0.688	18.05	0.191	0.688	29.78
DFS _{primal} minTree	14.44	0.206	0.617	0.687	15.79	0.255	0.689	30.83
DFS _{primal} SDSL	27.43	0.205	1.177	1.309	29.67	0.268	1.265	30.25
DFS _{primal} SDSL _{min}	24.29	0.188	1.013	1.133	26.39	0.247	1.102	29.78
DFS _{primal} SDSL _{min+v5}	26.11	0.203	1.064	1.198	28.39	0.263	1.158	<u>27.72</u>
BFS _{dual} minTree	13.46	0.199	0.561	0.620	<u>13.98</u>	0.241	<u>0.598</u>	30.83
BFS _{dual} SDSL	18.18	0.156	0.677	0.767	17.89	0.179	0.664	30.25
BFS _{dual} SDSL _{min}	17.41	0.149	0.638	0.728	17.26	<u>0.176</u>	0.639	29.78
DegHeur minTree	<u>12.77</u>	0.197	<u>0.538</u>	<u>0.608</u>	14.06	0.241	0.601	30.83
DegHeur SDSL	16.59	0.139	0.640	0.763	18.90	0.194	0.729	30.25
DegHeur SDSL _{min}	15.93	<u>0.137</u>	0.612	0.719	18.05	0.187	0.689	29.78


Trade-off (tiger_map)



Trade-off (tiger_map)



Trade-off *dfs* primal



Trade-off *neighbors* dual

- 1 Introduction
- 2 Related work
 - Compact data structures
 - Succinct ordinal trees
 - Planar embeddings
- 3 Our contribution
 - Effect of the trees' topology
 - Computing shallower spanning trees
 - Alternative data structures
- 4 Experimental analysis
- 5 Conclusions and future work

Conclusions

- The topology of the spanning trees used in compact representations of planar graphs may have an important impact in the performance of the CDS in practice.

Conclusions

- The topology of the spanning trees used in compact representations of planar graphs may have an important impact in the performance of the CDS in practice.
- Using shallower trees reduces the average distance between the pair of parentheses encoding a node in a BPS. Therefore, CDS such as the *range min-max tree* can perform faster.

Conclusions

- The topology of the spanning trees used in compact representations of planar graphs may have an important impact in the performance of the CDS in practice.
- Using shallower trees reduces the average distance between the pair of parentheses encoding a node in a BPS. Therefore, CDS such as the *range min-max tree* can perform faster.
- Using new implementations of CDS for BPS may allow further improvements.
- Overall, our new methods can save up to 8% of space while queries may be twice as fast.

Conclusions

- The topology of the spanning trees used in compact representations of planar graphs may have an important impact in the performance of the CDS in practice.
- Using shallower trees reduces the average distance between the pair of parentheses encoding a node in a BPS. Therefore, CDS such as the *range min-max tree* can perform faster.
- Using new implementations of CDS for BPS may allow further improvements.
- Overall, our new methods can save up to 8% of space while queries may be twice as fast.
- In general, queries on the primal are faster when we use the degree heuristic or a BFS traversal on the primal graph.

Conclusions

- The topology of the spanning trees used in compact representations of planar graphs may have an important impact in the performance of the CDS in practice.
- Using shallower trees reduces the average distance between the pair of parentheses encoding a node in a BPS. Therefore, CDS such as the *range min-max tree* can perform faster.
- Using new implementations of CDS for BPS may allow further improvements.
- Overall, our new methods can save up to 8% of space while queries may be twice as fast.
- In general, queries on the primal are faster when we use the degree heuristic or a BFS traversal on the primal graph.
- Similarly, queries on the dual are faster when the BFS is performed on the dual graph.

- To find a method which gives a balanced point for queries on the primal and dual graph.

- To find a method which gives a balanced point for queries on the primal and dual graph.
- To identify new parameters that affect the performance of the data structure.

- To find a method which gives a balanced point for queries on the primal and dual graph.
- To identify new parameters that affect the performance of the data structure.
- To study the effect of the topology on balanced parentheses representations that are not based on trees.

Speeding up compact planar graphs by using shallower trees[†]

Alexander Iribarra-Cortés^{1,2} José Fuentes-Sepúlveda¹
Diego Seco^{1,2} Roberto Asín¹

¹Department of Computer Science, Universidad de Concepción, Chile

²Millennium Institute for Foundational Research on Data, Chile

March 24, 2022

[†]This work was funded by ANID Millennium Science Initiative Program - Code ICN17_002 (1st and 3rd authors), PFCHA/Doctorado Nacional/2021-21211768 (1st author), PAI grant 77190038 (2nd author), and CYTED grant 519RT0579 (3rd author).