# Shortest-Path Queries in Planar Graphs on GPU-Accelerated Architectures 

Guillaume Chapuis and Hristo Djidjev<br>Los Alamos National Laboratory, Los Alamos, NM 87545, USA<br>\{gchapuis,djidjev\}@lanl.gov


#### Abstract

We develop an efficient parallel algorithm for answering shortestpath queries in planar graphs and implement it on a multi-node CPU/GPU clusters. The algorithm uses a divide-and-conquer approach for decomposing the input graph into small and roughly equal subgraphs and constructs a distributed data structure containing shortest distances within each of those subgraphs and between their boundary vertices. For a planar graph with $n$ vertices, that data structure needs $O(n)$ storage per processor and allows queries to be answered in $O\left(n^{1 / 4}\right)$ time.


Keywords: shortest path problems, graph algorithms, distributed computing, GPU computing, graph partitioning

## 1 Introduction

Finding shortest paths (SPs) in graphs has applications in transportation, social network analysis, network routing, and robotics, among others. The problem asks for a path of shortest length between one or more pairs of vertices. There are many algorithm for solving SP problems sequentially. Dijkstra's algorithm [2] finds the distances between a source vertex $v$ and all other vertices of the graph in $O(m \log n)$ time, where $n$ and $m$ are the numbers of the vertices and edges of the graph, respectively. It can also be used to find efficiently the distance between a pair of vertices. This algorithm is nearly optimal (within a logarithmic factor), but has irregular structure, which makes it hard to implement efficiently in parallel. Floyd-Warshall's algorithm, on the other hand, finds the distances between all pairs of vertices of the graph in $O\left(n^{3}\right)$ time, which is efficient for dense ( $m=\Theta\left(n^{2}\right)$ ) graphs, has a regular structure good for parallel implementation, but is inefficient for sparse $(m=O(n))$ graphs such as planar graphs.

In this paper we are considering the query version of the problem. It asks to construct a data structure that will allow to answer any subsequent distance query fast. A distance query asks, given an arbitrary pair of vertices $v, w$, to compute $\operatorname{dist}(v, w)$. This problem has applications in web mapping services such as MapQuest and Google Maps. There is a tradeoff between the size of the data structure and the time for answering a query. For instance, Dijkstra's algorithm gives a trivial solution of the query version of the SP problem with (small) $O(n+m)$ space (for storing the input graph), but large $O(m \log n)$ query time
(for running Dijkstra's algorithm with a source the first query vertex). On the other end of the spectrum, Floyd-Warshall's algorithm can be used to construct a (large) $O\left(n^{2}\right)$ data structure (the distance matrix) allowing (short) $O(1)$ query time (retrieving the distance from the data base). However, for very large graphs, the $O\left(n^{2}\right)$ space requirement is impractical. We are interested in an algorithm that needs significantly less than than $O\left(n^{2}\right)$ space, but will answer queries faster than Disjkstra's algorithm. Our algorithm will use the structure of planar graphs for increased efficiency, as most road networks are planar or near-planar, and will also be highly parallelizable, making use of the features available in modern high-performance clusters and specialized processors such as the GPUs.

The query version for shortest path queries in planar graphs was proposed in [3] and after that different aspects of the problem were studied by multiple authors, e.g., 6189]. Here we present the first distributed implementation for solving the problem that is designed to make use of the potential for parallelism offered by GPUs. Our solution makes use of the fast parallel algorithm for computing shortest paths in planar graphs from [4], resulting in asymptotically faster and also shown to be efficient in practice.

## 2 Preliminaries

Given a graph $G$ with a weight $\mathrm{wt}(e)$ on each edge $e$, the length of a path $p$ is the sum of the weights of the edges of the path. The single-pair shortest path problem $(S P S P)$ is, given a pair $v, w$ of vertices of $G$, to find a path between $v$ and $w$, called shortest path (SP), with minimum length. The length of that path is called distance between $v$ and $w$ and is denoted as $\operatorname{dist}(v, w)$. For any subgraph $H$ of $G$, the distance between $v$ and $w$ in $H$ is denoted as $\operatorname{dist}_{H}(v, w)$. The single-source shortest path problem (SSSP) is to find SPs from a fixed vertex $v$ to all other vertices of $G$. Finally, the all-pairs shortest path problem (APSP) is to find SPs between all pairs of vertices. There are distance versions of SPSP, SSSP, and APSP, which are more commonly studied, where the objective is to compute the corresponding distances instead of SPs. Most distance algorithms allow the corresponding SPs to be retrieved in additional time proportional to the number of the edges of the path. In this paper, by SPSP, SSSP, and APSP we mean the distance versions of these problems.

A $k$-partition $\mathcal{P}$ of $G$ is a set $V_{1}, \ldots, V_{k}$ of subsets of $V(G)$, the set of the vertices of $G$, such that $V_{i} \cap V_{j}=\emptyset$ if $i \neq j$ and $\bigcup_{i=1}^{k} V_{i}=V(G)$. We call the subgraphs of $G$ induced by $V_{i}$ components of $\mathcal{P}$. The boundary of the partition consists consists of all vertices of $G$ that have at least one neighbor in a different component. We denote by $B G(G)$ or simply by $B G$ the subgraph of $G$ induced by the boundary vertices. For any $C \in \mathcal{P}$, we denote by $B(C)$ the set of all boundary vertices that are from $C$. For any planar graph of $n$ vertices and bounded $(O(1)$ as a function of $n$ ) vertex degree one can find in $O(n)$ time a $k$-partition $\mathcal{P}$ with $|B(C)|=O(\sqrt{n / k})$ for each component $C \in \mathcal{P}$.

## 3 Algorithm overview and analysis

Our algorithm works in two modes: preprocessing mode, during which a data structure is computed that allows efficient SP queries, and the query mode that uses that data structure to compute the distance between a query pair of vertices. We assume that the input is a planar graph $G$ of $n$ vertices and bounded vertex degree and the cluster has $p$ nodes.

### 3.1 Preprocessing mode

The preprocessing algorithm (Algorithm 1) has three phases. During the first phase (line 1), the graph is partitioned and each component is assigned to a distinct cluster node. During the second phase (lines 2-5), the APSP problem is solved for each component $C$ independently and in parallel and the computed distance matrix $\operatorname{APSP}(C)$ is stored at the same node. Finally, in the third phase (lines 6-10), the boundary graph $B G$ is constructed and the APSP is solved for $B G$. That computation is done distributedly such that the distances from vertex $v \in B G$ to all other vertices of $B G$ are computed at the node containing $v$, by using Dijkstra's algorithm [2]. The computed distance matrix is stored at the node that has done the computations. Hence, at the end of the algorithm, the node $N(C)$ contains two matrices: one containing the SP distances in $C$ and the other containing all SP distances in $B G$ with source a vertex in $B G \cap C$.

One can think of $B G$ as a compressed version of $G$ where the non-boundary vertices are removed, but are implicitly represented in $B G$ by the information encoded in its edge weights. Note however that the distances $\operatorname{APSP}(C)$ (and the corresponding edge weights of $B G$ ) are not distances in $G$; the reason is that a shortest path between two vertices $v$ and $w$ from $C$ might pass through vertices not in $C$. Hence the following fact is non-trivial.

Lemma 1. [4] For any two vertices $v, w \in B G$ the distance between $v$ and $w$ in $B G$ is equal to the distance between $v$ and $w$ in $G$.

We will next estimate the time and space (memory) required to run the algorithm. As $G$ is planar and of bounded vertex degree (as a function of $n$ ), it can be divided in $O(n)$ time into $k$ parts so that each part has no more than $(n / k)$ vertices and $O(\sqrt{n / k})$ boundary vertices [5]. We will estimate the requirements of each phase. Since the maximum amount of coarse-grained parallelism of Algorithm 1 is $\min \{p, k\}$, we assume without loss of generalization that $p \leq k$.

Phase 1 requires $O(n)$ running time and $O(n)$ space [5].
The complexity of Phase 2 is dominated by the time for computing distances in line 3. We assume that we are using the algorithm from 4] that can be implemented efficiently on a GPU-accelerated architecture and has complexity $O\left(N^{9 / 4}\right)$. Then Phase 2 requires $O\left((k / p)(n / k)^{9 / 4}\right)=O\left(n^{9 / 4} /\left(p k^{5 / 4}\right)\right)$ time and $k O((n / k) \sqrt{n / k})=O\left(n^{3 / 2} / k^{1 / 2}\right)$ total space. The space per processor is $k O\left((n / k)^{2}\right)=O\left(n^{2} / k\right)$.

For Phase 3, the number of the vertices of $B G$ is $k O(\sqrt{n / k})=O(\sqrt{n k})$ and the number of the edges is $k O\left((\sqrt{n / k})^{2}\right)=O(n)$. One execution of line 8 (for one
component $C)$ takes $(k / p)|B G \cap C||E(B G)| \log (|B G|)=(k / p) O(\sqrt{n / k}) O(n \log n)$ time and $O(n)$ space. The space needed for one iteration of Step 9 is $\mid B G \cap$ $C\left||B G|=O(\sqrt{n / k} \sqrt{n k})=O(n)\right.$. Hence Phase 3 requires $O\left((k / p) n^{3 / 2} / k^{1 / 2} \log n\right)$ $=O\left(n^{3 / 2} k^{1 / 2} \log n / p\right)$ time and $O(n k / p)$ space per processor.

Summing up the requirements for Phases 1,2 , and 3, we get $O\left(n^{9 / 4} /\left(p k^{5 / 4}\right)+\right.$ $\left.n^{3 / 2} k^{1 / 2} \log n / p\right)$ ) time and $O\left(n+n^{2} /(p k)+n k / p\right)$ space per processor needed for Algorithm 1. Assuming space is more important in this case than time (since nodes have limited memory), we find that $k=n^{1 / 2}$ minimizes the function $n^{2} / k+n k$. Hence we have the following result.

Lemma 2. With $k=\left\lceil n^{1 / 2}\right\rceil$ and $p \leq k$, Algorithm 1 runs in $O\left(n^{7 / 4} \log n / p\right)$ time and uses $O\left(n^{3 / 2} / p\right)$ space per processor. With $p=k$, the time and space are $O\left(n^{5 / 4}\right)$ and $O(n)$, respectively.

The time bound of Lemma 2 is conservative as it doesn't take into account our use of fine-grain parallelism due to multi-threading, e.g., by the GPUs.

```
Algorithm 1 Preprocessing algorithm
Input: A planar graph \(G\)
Output: A data structure for efficient shortest path queries in \(G\)
                            /* Partitioning */
    1: Construct a \(k\)-partition \(\mathcal{P}\) of \(G\) and assign each component \(C\) to a distinct node
    \(N(C)\)
        /* Solve the APSP problem for each component */
    for all components \(C \in \mathcal{P}\) do in parallel
        Solve APSP for \(C\) and save the distances in a table \(\operatorname{APSP}(C)\)
        For each pair of boundary vertices \(v, w \in C\) define edge \((v, w)\), if not already in
            \(G\), and assign a weight \(\operatorname{wt}(v, w)=\operatorname{dist}_{C}(v, w)\)
        end for
            /* Solve the APSP problem for the boundary graph */
    Define a boundary graph \(B G\) with vertices all boundary vertices of \(G\) and edges as
    defined in the previous step and store it at each node
    for all components \(C \in \mathcal{P}\) do in parallel
        Solve SSSP in \(B G\) for each vertex of \(C \cap B G\)
        Store the distances from all vertices of \(C \cap B G\) to all vertices of \(B G\) in a
            table \(\operatorname{APSP}_{B G}(C)\)
        end for
```


### 3.2 Query mode

The query algorithm (Algorithm 2) is based on the fact that if $C_{1} \neq C_{2}$, then any path between $v_{1}$ and $v_{2}$ should cross both $B\left(C_{1}\right)$ and $B\left(C_{2}\right)$. Let $\pi$ be a shortest path between $v_{1}$ and $v_{2}$. Then $\pi$ can be divided into three parts: from $v_{1}$ to a vertex $b_{1}$ from $B\left(C_{1}\right)$, from $b_{1}$ to a vertex $b_{2}$ on $p$ from $B\left(C_{2}\right)$, and from $b_{2}$ to $v_{2}$. Vertices $b_{1}$ and $b_{2}$ minimizing the length of $p$ are found as follows: in
the loop on lines 2-7, for each $b_{2}$ an optimal $b_{1}$ and $\operatorname{dist}\left(v_{1}, b_{2}\right)$ are found; in lines $10-12$ an optimal $b_{2}$ is found.

```
Algorithm 2 Query algorithm
Input: Vertices \(v_{1}, v_{2}\) of \(G\), a \(k\)-partition \(\mathcal{P}\) of \(G\), tables \(\operatorname{APSP}(C)\) and \(\operatorname{APSP}_{B G}(C)\)
    for all \(C \in \mathcal{P}\)
Output: \(\operatorname{dist}\left(v_{1}, v_{2}\right)\)
    Determine components \(C_{1}\) and \(C_{2}\) such that \(v_{1} \in C_{1}, v_{2} \in C_{2}\)
    for all vertices \(b_{2} \in B\left(C_{2}\right)\) do in parallel
                    /* Compute \(\operatorname{dist}\left(v_{1}, b_{2}\right)\) */
        \(\operatorname{dist}\left(v_{1}, b_{2}\right)=\infty\)
        for all vertices \(b_{1} \in B\left(C_{1}\right)\) do
            \(\operatorname{dist}\left(v_{1}, b_{2}\right)=\min \left\{\operatorname{dist}\left(v_{1}, b_{2}\right), \operatorname{dist}_{C_{1}}\left(v_{1}, b_{1}\right)+\operatorname{dist}_{B G}\left(b_{1}, b_{2}\right)\right\}\)
        end for
    end for
    If \(N\left(C_{1}\right) \neq N\left(C_{2}\right)\) then transfer the column of \(\mathrm{SP}\left(C_{2}\right)\) corresponding to \(v_{2}\) from
    \(N\left(C_{2}\right)\) to \(N\left(C_{1}\right)\).
                    /* Now we can compute \(\operatorname{dist}\left(v_{1}, v_{2}\right) * /\)
    \(\operatorname{dist}\left(v_{1}, v_{2}\right)=\infty\)
    for all vertices \(b_{2} \in B\left(C_{2}\right)\) do
        \(\operatorname{dist}\left(v_{1}, v_{2}\right)=\min \left\{\operatorname{dist}\left(v_{1}, v_{2}\right), \operatorname{dist}\left(v_{1}, b_{2}\right)+\operatorname{dist}_{C_{2}}\left(b_{2}, v_{2}\right)\right\}\)
    end for
    If \(C_{1}=C_{2}\) then \(\operatorname{dist}\left(v_{1}, v_{2}\right)=\min \left\{\operatorname{dist}\left(v_{1}, v_{2}\right), \operatorname{dist}_{C_{1}}\left(v_{1}, v_{2}\right)\right\}\), where the distance
    dist \(_{C_{1}}\left(v_{1}, v_{2}\right)\) is taken from \(\operatorname{APSP}\left(C_{1}\right)\).
```

Lemma 3. Algorithm 2 correctly computes $\operatorname{dist}\left(v_{1}, v_{2}\right)$ and its running time is $O\left(n^{1 / 4}\right)$ with $k=\left\lceil n^{1 / 2}\right\rceil$ and $p \geq\left\lceil n^{1 / 4}\right\rceil$.

Proof. Let $\pi$ be a shortest path between $v_{1}$ and $v_{2}$, let $C_{1} \neq C_{2}$, and let $b_{1}$ be the first vertex along $\pi$ that is on $B\left(C_{1}\right)$ and $\pi_{1}$ be the subpath of $\pi$ from $v_{1}$ to $b_{1}$, let $\pi_{2}$ be the last vertex along $\pi$ that is on $B\left(C_{2}\right)$ and $\pi_{2}$ be the subpath of $\pi$ from $b_{1}$ to $b_{2}$, and let $\pi_{3}$ be the subpath of $\pi$ from $b_{2}$ to $v_{2}$. Then $\pi_{1}$ is entirely in $C_{1}$ and hence $\operatorname{dist}_{C_{1}}\left(v_{1}, b_{1}\right)=\operatorname{dist}_{G}\left(v_{1}, b_{1}\right)$ (note, however, that the distances in $\operatorname{APSP}\left(C_{1}\right)$ from $v_{1}$ to other vertices from $B\left(C_{1}\right)$ may not be correct). Similarly, $\operatorname{dist}_{C_{2}}\left(b_{2}, v_{2}\right)=\operatorname{dist}_{G}\left(b_{2}, v_{2}\right)$. Finally, $\operatorname{dist}_{B G}\left(b_{1}, b_{2}\right)=\operatorname{dist}_{G}\left(b_{1}, b_{2}\right)$ by Lemma 1 . Hence lines 5 and 11 use correct values for computing the distances between $v_{1}$ and $b_{2}$ and between $b_{2}$ and $v_{2}$.

If $C_{1}=C_{2}$ (line 13), then a shortest path between $v_{1}$ and $v_{2}$ may or may not leave $C_{1}$. In the first case lines 1-12 compute correctly $\operatorname{dist}\left(v_{1}, v_{2}\right)$, in the second case $\operatorname{APSP}\left(C_{1}\right)$ contains the correct distance.

The loop on lines 5-10 takes time $\left|B\left(C_{1}\right)\right|\left|B\left(C_{2}\right)\right| / p=O(\sqrt{n / k} \sqrt{n / k} / p)=$ $O(n /(p k))$, for $p \leq \min \left\{k,(n / k)^{1 / 2}\right\}$. If $k=n^{1 / 2}$ and $p=n^{1 / 4}$ (the maximum value for which the formula applies), that time becomes $O\left(n^{1 / 4}\right)$. The loop in lines 10-12 takes time $O\left((n / k)^{1 / 2}\right)=O\left(n^{1 / 4}\right)$ for $k=n^{1 / 2}$.

Note that using the methodology of [3], a more complex implementation of Algorithm 2 can reduce the query time to logarithmic. Note also that computation in lines 2-7 can be overlapped with transferring of data in line 8 thereby saving time (upto a factor of two).

## 4 Implementation details

In this section, we describe how the preprocessing and query modes are implemented on a hybrid CPU-GPU cluster. We use a distance matrix to represented both the input graph $G$ and the output. Such a 2 -dimensional matrix contains in cell $(i, j)$ the value of the distance from vertex $i$ to vertex $j$. Initially, cell $(i, j)$ contains $\mathrm{wt}(i, j)$ if an edge $(i, j)$ is present in $G$, or infinity otherwise. These values are updated as the algorithm progresses. At the end of the algorithm, cell $(i, j)$ contains $\operatorname{dist}(i, j)$.

In phase 1 of the preprocessing mode, we construct a $k$-partition of $G$ using the METIS library [7]. Based on that partition, we reorder the vertices of $G$ so that vertices from the same component have consecutive indices and boundary vertices of each components have the lowest indices - see Figure 1 .


Fig. 1: Distance matrix after reordering of the vertices. Vertices from the same component are stored contiguously starting with boundary vertices. Red submatrices are also part of the boundary distance matrix. Grey submatrices do not generate any computations in preprocessing mode.


Fig. 2: The distances required to compute $\operatorname{dist}(v, w)$, shown in green, are scattered in three submatrices: two diagonal ones, for component $I$ and for component $J$, and a non-diagonal submatrix $(I, J)$.

In phase 2, we compute the shortest distances within each of the components. For $k$ components, this phase gives a total $k$ independent tasks that can be executed in parallel. Computations at this phase are already balanced across
nodes as components contain roughly the same number of vertices and the APSP algorithm from [4] ensures the same $O\left(N^{9 / 4}\right)$ complexity with respect to the number of nodes.

Finally, phase 3 consists in computing the shortest distances within the boundary graph using Dijkstra's algorithm. Computations at this phase may be imbalanced between nodes for two reasons. First, the number of boundary vertices in two components may differ and, second, the complexity of Dijkstra's algorithm does not solely depend on the number of vertices in the graph, but also on the number of edges, which may vary even more than the number of vertices between two components' boundary graphs.

In the query mode, we are interested in finding $\operatorname{dist}(v, w)$, where $v$ and $w$ are from components $I$ and $J$, respectively. The required values for that computation are scattered in three submatrices, as illustarted in Figure 2 For such a query, assuming $k=p$, node $i$, holding the required values from diagonal submatrix $I$ and non-diagonal submatrix $(I, J)$, will be in charge of the computations. Required values from diagonal submatrix $J$ are held by node $j$ and need to be transfered to node $i$.

## 5 Experimental evaluation

In this section we describe experiments designed to test our algorithm and its implementation. Specifically, we are going to test the strong scaling properties by running our code on a fixed graph size and a varying number $p$ of cluster nodes and number $k$ of components. All computations are run on a 300 node cluster. Each cluster node is comprised of $2 \times$ Eight-Core Intel Xeon model E5-2670 @ 2.6 GHz and two GPGPU Nvidia Tesla M2090 cards connected to PCIe-2.0 x16 slots. In order to make full use of the available GPUs, each node is assigned at least two graph components so that the two associated diagonal submatrices can be computed simultaneously on the two GPUs.

For the strong-scaling experiment, the graph size is fixed to 256 k vertices. Preprocessing and queries are run with increasing numbers of nodes ranging from 4 to 64 . Each node handles 2 components (one per available GPU); therefore the number of components $k$ ranges from 8 to 128 .

Figure 3 shows the run times for the preprocessing mode. For low numbers of nodes and thus low values of $k$, preprocessing time is dominated by step 2 the computation of the shortest distances within each component - since lower $k$ values means larger components. For higher numbers of nodes and thus higher values of $k$, preprocessing time becomes dominated by step 3 - the computation of the boundary graph - as more components mean higher numbers of incident edges and thus larger boundary graphs. Note that while the figure seems to show supralinear speedup, that is not the case (and similarly for the memory usage). The reason is that, with increasing the number of processors $p$, the number $k$ of parts is increased too (as it is tied to $p$ in this implementation) and hence the complexity of the algorithm is also reduced.


Fig. 3: Preprocessing run times for a fixed graph size of 256 k vertices and increasing number of nodes.

Fig. 4: Peak memories and run times for 10k queries for a fixed graph size of 256 k vertices and increasing number of parts/processors.

Figure 4 shows the query times and peak memory usage per node. The run times are given for 10,000 queries from random sources to random targets. Note that in the query mode only fine-grain (node-level) parallelism is used, while multiple nodes are still needed for distributed storage and, optionally, to handle multiple queries in parallel (not implemented in the current version). For the memory usage, the optimal value for $k$, theoretically expected to be $\sqrt{n}-$ or 512 for this instance - is not reached in this experiment since $k$ only goes up to 128 . We can however see that peak memory usage per node is still dropping with increasing values of $k$ up to 128 . The query times in the figure vary from about 2 milliseconds per query for $k=8$ to 0.25 milliseconds for $k=128$. Compared with the Boost library implementation of Dijkstra's algorithm, our implementation answers queries on the largest instances about 1000 times faster.

## 6 Conclusion

We developed and implemented a distributed algorithm for shortest path queries in planar graphs with good scalability. It allows answering SP queries in $O\left(n^{1 / 4}\right)$ time by using $O(\sqrt{n})$ processors with $O(n)$ space per processor and $O\left(n^{5 / 4}\right)$ preprocessing time. Our implementation on 300 node CPU-GPU cluster has preprocessing time of less than 10 seconds using 32 or more nodes and 0.025 milliseconds per query using two nodes. Interesting tasks for future research is implementing a version allowing parallel queries and reducing the query time of the implementation to $O(\log n)$ by using properties of graph planarity.

## References

1. Chen, D.Z., Xu, J.: Shortest path queries in planar graphs. pp. 469-478. STOC '00, ACM, New York, NY, USA (2000)
2. Dijkstra, E.W.: A note on two problems in connexion with graphs. Numerische mathematik 1(1), 269-271 (1959)
3. Djidjev, H.: Efficient algorithms for shortest path queries in planar digraphs. In: Graph-Theoretic Concepts in Computer Science, Lecture Notes in Computer Science, vol. 1197, pp. 151-165 (1997)
4. Djidjev, H., Thulasidasan, S., Chapuis, G., Andonov, R., Lavenier, D.: Efficient multi-GPU computation of all-pairs shortest paths. In: IPDPS'2014. pp. 360-369 (2014)
5. Frederickson, G.N.: Planar graph decomposition and all pairs shortest paths. J. ACM 38(1), 162-204 (1991)
6. Hutchinson, D., Maheshwari, A., Zeh, N.: An external memory data structure for shortest path queries (extended abstract). pp. 51-60. COCOON'99, Springer-Verlag, Berlin, Heidelberg (1999)
7. Karypis, G., Kumar, V.: Multilevel k-way partitioning scheme for irregular graphs. Journal of Parallel and Distributed computing 48(1), 96-129 (1998)
8. Kowalik, L.: Shortest path queries in planar graphs in constant time. In: STOC03. pp. 143-148. ACM (2003)
9. Mozes, S., Sommer, C.: Exact distance oracles for planar graphs. pp. 209-222. SODA '12, SIAM (2012)
