

1 Fine-Grained Complexity of k -OPT in 2 Bounded-Degree Graphs for Solving TSP

3 Édouard Bonnet

4 ENS Lyon, LIP, Lyon, France

5 edouard.bonnet@ens-lyon.fr

6 Yoichi Iwata

7 National Institute of Informatics, Tokyo, Japan

8 yiwata@nii.ac.jp

9 Bart M.P. Jansen

10 Eindhoven University of Technology, Eindhoven, The Netherlands

11 b.m.p.jansen@tue.nl

12 Łukasz Kowalik 

13 Institute of Informatics, University of Warsaw, Poland

14 kowalik@mimuw.edu.pl

15 Abstract

16 The TRAVELING SALESMAN PROBLEM asks to find a minimum-weight Hamiltonian cycle in an
17 edge-weighted complete graph. Local search is a widely-employed strategy for finding good solutions
18 to TSP. A popular neighborhood operator for local search is k -opt, which turns a Hamiltonian
19 cycle \mathcal{C} into a new Hamiltonian cycle \mathcal{C}' by replacing k edges. We analyze the problem of determining
20 whether the weight of a given cycle can be decreased by a k -opt move. Earlier work has shown
21 that (i) assuming the Exponential Time Hypothesis, there is no algorithm that can detect whether
22 or not a given Hamiltonian cycle \mathcal{C} in an n -vertex input can be improved by a k -opt move in
23 time $f(k)n^{o(k/\log k)}$ for any function f , while (ii) it is possible to improve on the brute-force running
24 time of $\mathcal{O}(n^k)$ and save linear factors in the exponent. Modern TSP heuristics are very successful
25 at identifying the *most promising* edges to be used in k -opt moves, and experiments show that
26 very good global solutions can already be reached using only the top- $\mathcal{O}(1)$ most promising edges
27 incident to each vertex. This leads to the following question: can improving k -opt moves be found
28 efficiently in graphs of bounded degree? We answer this question in various regimes, presenting new
29 algorithms and conditional lower bounds. We show that the aforementioned ETH lower bound also
30 holds for graphs of maximum degree three, but that in bounded-degree graphs the best improving
31 k -move can be found in time $\mathcal{O}(n^{(23/135+\epsilon_k)k})$, where $\lim_{k \rightarrow \infty} \epsilon_k = 0$. This improves upon the
32 best-known bounds for general graphs. Due to its practical importance, we devote special attention
33 to the range of k in which improving k -moves in bounded-degree graphs can be found in *quasi-linear*
34 time. For $k \leq 7$, we give quasi-linear time algorithms for general weights. For $k = 8$ we obtain a
35 quasi-linear time algorithm when the weights are bounded by $\mathcal{O}(\text{polylog } n)$. On the other hand, based
36 on established fine-grained complexity hypotheses about the impossibility of detecting a triangle in
37 edge-linear time, we prove that the $k = 9$ case does not admit quasi-linear time algorithms. Hence
38 we fully characterize the values of k for which quasi-linear time algorithms exist for polylogarithmic
39 weights on bounded-degree graphs.

40 **2012 ACM Subject Classification** Theory of computation \rightarrow Graph algorithms analysis; Theory of
41 computation \rightarrow Parameterized complexity and exact algorithms

42 **Keywords and phrases** traveling salesman problem, k -OPT, bounded degree

43 **Digital Object Identifier** [10.4230/LIPIcs.ESA.2019.21](https://doi.org/10.4230/LIPIcs.ESA.2019.21)

44 **Funding** Yoichi Iwata: Supported by JSPS KAKENHI Grant Number JP17K12643.

45 Bart M.P. Jansen: Supported by ERC Starting Grant 803421 ReduceSearch.

46 Łukasz Kowalik: Supported by ERC Starting Grant TOTAL (grant agreement No 677651).



© Édouard Bonnet, Yoichi Iwata, Bart M.P. Jansen, Łukasz Kowalik;
licensed under Creative Commons License CC-BY

27th Annual European Symposium on Algorithms (ESA 2019).

Editors: Michael A. Bender, Ola Svensson, and Grzegorz Herman; Article No. 21; pp. 21:1–21:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

47 **Acknowledgements** This research was initiated at the Shonan Meeting *Parameterized Graph Al-*
 48 *gorithms & Data Reduction: Theory Meets Practice* held during March 4-8, 2019 in Shonan Village
 49 Center, Japan. Yoichi Iwata thanks Kaggle Traveling Santa 2018 Competition for motivating him
 50 to study practical TSP heuristics. He also thanks Shogo Murai for valuable discussion about the
 51 possibility of faster k -OPT algorithms.

52 **1 Introduction**

53 **1.1 Motivation**

54 The TRAVELING SALESMAN PROBLEM (TSP) hardly needs an introduction; it is one of the
 55 most important problems in combinatorial optimization, which asks to find a Hamiltonian
 56 cycle of minimum weight in an edge-weighted complete graph. Local search is widely used
 57 in practical TSP solvers [10, 11]. The most commonly used neighborhood is a k -move (or
 58 k -opt move). A k -move on a Hamiltonian cycle \mathcal{C} is a pair (E^-, E^+) of edge sets such that
 59 $E^- \subseteq E(\mathcal{C})$, $|E^-| = |E^+| = k$ and $(\mathcal{C} \setminus E^-) \cup E^+$ is also a Hamiltonian cycle. Marx [13]
 60 showed that finding an improving k -move (i.e., a k -move that results in a lighter Hamiltonian
 61 cycle) is W[1]-hard parameterized by k , and this result was refined by Guo et al. [6] to
 62 obtain an $f(k)n^{\Omega(k/\log k)}$ lower bound under the Exponential Time Hypothesis (ETH). For
 63 small values of k , the current fastest running time is $\mathcal{O}(n^k)$ for $k = 2, 3$ (by exhaustive
 64 search), $\mathcal{O}(n^3)$ for $k = 4$ [4], and $\mathcal{O}(n^{3.4})$ for $k = 5$ [3]. Moreover, de Berg et al. [4] and
 65 Cygan et al. [3] showed that improving the running time to $\mathcal{O}(n^{3-\epsilon})$ for $k = 3$ or $k = 4$
 66 implies a breakthrough result of $\mathcal{O}(n^{3-\delta})$ -time algorithm for ALL-PAIRS SHORTEST PATHS.

67 From the hardness shown by the theoretical studies, it seems that local search can be
 68 applied only to small graphs. Nevertheless, state-of-the-art local search TSP solvers can deal
 69 with large graphs with tens of thousands of vertices. This is mainly due to the following two
 70 heuristics.

- 71 1. They sparsify the input graph by picking the top- d important incident edges for each
 72 vertex according to an appropriate importance measure. For example, Lin-Kernighan [12]
 73 picks the top-5 nearest neighbors, and its extension LKH [8] picks the top-5 α -nearest
 74 neighbors, where the α -distance of an edge is the increase of the Held-Karp lower
 75 bound [7] by including the edge. The empirical evaluation by Helsgaun [8] showed that
 76 the sparsification by the α -nearest neighbors can preserve almost optimal solutions.
- 77 2. They mainly focus on *sequential* k -moves. In general, $E^- \cup E^+$ is a set of edge-disjoint
 78 closed walks, each of which alternately uses edges in E^- and E^+ . If it consists of a single
 79 closed walk, the move is called sequential. Graphs of maximum degree d with n vertices
 80 have at most $n(2(d-2))^{k-1}$ sequential k -moves (n choices for the starting point, 2 choices
 81 for the next edge in E^- , and at most $d-2$ choices for the next edge in E^+), which
 82 is linear in n when considering d and k as constants. On the other hand, linear-time
 83 computation of non-sequential k -moves appears non-trivial. Lin-Kernighan does not
 84 search for non-sequential moves at all, and after it finds a local optimum, it applies special
 85 non-sequential 4-moves called *double bridges* to get out of the local optimum. LKH-2 [9]
 86 improves Lin-Kernighan by heuristically searching for non-sequential moves during the
 87 local search.

88 This state of affairs raises the following questions: what is the complexity of finding
 89 improving k -moves in bounded-degree graphs? How does the complexity scale with k , and
 90 can it be done efficiently for small values of k ? Since improving *sequential moves* can be
 91 found in linear time for fixed k and d , to answer these questions we have to investigate
 92 non-sequential k -moves in bounded-degree graphs.

1.2 Our contributions

We classify the complexity of finding improving k -moves in bounded-degree graphs in various regimes. We present improved algorithms that exploit the degree restrictions using the structure of k -moves, treewidth bounds, color-coding, and suitable data structures. We also give new lower bounds based on the Exponential Time Hypothesis (ETH) and hypotheses from fine-grained complexity concerning the complexity of detecting triangles. To state our results in more detail, we first introduce the two problem variants we consider; a weak variant to which our lower bounds already apply, and a harder variant which can be solved by our algorithms.

k -OPT DETECTION

Parameter: k .

Input: An undirected graph G , a weight function $w: E(G) \rightarrow \mathbb{Z}$, an integer k , and a Hamiltonian cycle $\mathcal{C} \subseteq E(G)$.

Question: Can \mathcal{C} be changed into a Hamiltonian cycle of strictly smaller weight by a k -move?

The related optimization problem k -OPT OPTIMIZATION is to compute, given a Hamiltonian cycle in the graph, a k -move that gives the largest cost improvement, or report that no improving k -move exists. With this terminology, we describe our results.

We show that k -OPT DETECTION is unlikely to be fixed-parameter tractable on bounded-degree graphs: we give a new constant-degree lower-bound construction to show that there is no function f for which k -OPT DETECTION on subcubic graphs with weights $\{1, 2\}$ can be solved in time $f(k) \cdot n^{o(k/\log k)}$, unless ETH fails. Hence the running time lower bound for general graphs by Guo et al. [6] continues to hold in this very restricted setting. While the degree restriction does not make the problem fixed-parameter tractable, it is possible to obtain faster algorithms. By adapting the approach of Cygan et al. [3], exploiting the fact that the number of sequential moves is linear in n in bounded-degree graphs, and proving a new upper bound on the pathwidth of an k -edge even graph, we show that k -OPT OPTIMIZATION in n -vertex graphs of maximum degree $\mathcal{O}(1)$ can be solved in time $\mathcal{O}(n^{(23/135+\epsilon_k)k}) = \mathcal{O}(n^{(0.1704+\epsilon_k)k})$, where $\lim_{k \rightarrow \infty} \epsilon_k = 0$. This improves on the behavior for general graphs, where the current-best running time [3] is $\mathcal{O}(n^{(1/4+\epsilon_k)k})$.

Since quasi-linear running times are most useful for dealing with large inputs, we perform a fine-grained analysis of the range of k for which improving k -moves can be found in time $\mathcal{O}(n \text{ polylog } n)$ on n -vertex graphs. Observe that in the bounded-degree setting, the number of edges m is $\mathcal{O}(n)$. We prove lower bounds using the hypothesis that detecting a triangle in an unweighted graph cannot be done in nearly-linear time in the number of edges m , which was formulated in several ways by Abboud and Vassilevska Williams [1, Conjectures 2–3]. By an efficient reduction from TRIANGLE DETECTION, we show that an algorithm with running time $\mathcal{O}(n \text{ polylog } n)$ for 9-OPT DETECTION in subcubic graphs with weights $\{1, 2\}$ implies that a triangle in an m -edge graph can be found in time $\mathcal{O}(m \text{ polylog } m)$, contradicting popular conjectures. We complement these lower bounds by quasi-linear algorithms for all $k \leq 8$ to obtain a complete dichotomy for the case of integer weights bounded by $\mathcal{O}(\text{polylog } n)$. When the weights are not bounded, we obtain quasi-linear time algorithms for all $k \leq 7$, leaving open only the case $k = 8$.

1.3 Organization

Preliminaries are presented in Section 2. In Section 3 we give faster XP algorithms for varying k . By refining these ideas, we give quasi-linear-time algorithms for $k \leq 8$ in Section 4. Section 5 gives the reduction from TRIANGLE DETECTION to establish a superlinear lower

135 bound on subcubic graphs for $k = 9$. In Section 6 we describe the lower bound for varying k .

136 **2 Preliminaries**


137 Given a graph G edge-weighted by $w: E(G) \rightarrow \mathbb{Z}$, and a subset $F \subseteq E(G)$ of its edges,
 138 $w(F) := \sum_{e \in F} w(e)$. A k -move on a Hamiltonian cycle \mathcal{C} is pair (E^-, E^+) of edge sets such
 139 that $|E^-| = |E^+| = k$ and $(\mathcal{C} \setminus E^-) \cup E^+$ is also a Hamiltonian cycle. A k -move is called
 140 *improving* if $w((\mathcal{C} \setminus E^-) \cup E^+) < w(\mathcal{C})$, or equivalently and more simply $w(E^+) < w(E^-)$.
 141 A necessary condition for a pair (E^-, E^+) to be a k -move is that the multiset of endpoints
 142 of E^- is equal to the multiset of endpoints of E^+ . An exchange (E^-, E^+) that satisfies this
 143 condition is called a k -swap. We say that a k -swap *results* in the graph $(\mathcal{C} \setminus E^-) \cup E^+$. Note
 144 that a k -swap always results in a spanning disjoint union of cycles. A k -swap resulting in a
 145 graph with a single connected component is therefore a k -move. An *infeasible* k -swap is a
 146 k -swap which is not a k -move.

147 We say that a k -swap (E^-, E^+) *induces* the graph $E^- \cup E^+$. As a slight abuse of
 148 notation, a k -swap will sometimes directly refer to this graph. A k -swap (E^-, E^+) such
 149 that all edges $E^- \cup E^+$ are visited by a single closed walk alternating between E^- and E^+
 150 is called *sequential*. In particular, in a simple graph, every 2-swap is sequential. One can
 151 notice that an infeasible (sequential) 2-swap results in a disjoint union of exactly two cycles.
 152 A k -move can always be decomposed into sequential k_i -swaps (with $\sum k_i = k$) but some
 153 k -moves cannot be decomposed into sequential k_i -moves. The quantity $w(E^-) - w(E^+)$ is
 154 called the *gain* of the swap (E^-, E^+) . We distinguish *neutral* swaps, with gain 0, *improving*
 155 swaps, with strictly positive gain, and *worsening* swaps, with strictly negative gain.

156 For an integer n , we denote $[n] = \{1, \dots, n\}$. A k -embedding (or shortly: *embedding*) is
 157 an increasing function $f: [k] \rightarrow [n]$. A *connection k -pattern* (or shortly: *connection pattern*)
 158 is a perfect matching in the complete graph on the vertex set $[2k]$. A pair (f, M) where f is
 159 a k -embedding and M is a connection k -pattern, is an alternative description of a k -swap.
 160 Indeed, let e_1, \dots, e_n be subsequent edges of \mathcal{C} . Then, $E^- = \{e_{f(i)}: i \in [k]\}$. Vertices of the
 161 connection pattern correspond to endpoints of E^- , i.e., vertices $2i - 1, 2i \in [2k]$ correspond
 162 to the left and right (in the clockwise order) endpoint of $e_{f(i)}$, respectively. Thus, edges of
 163 the connection pattern correspond to a set E^+ of $|M|$ edges in G . We say that a k -swap
 164 (E^-, E^+) *fits into* M if there is an embedding f such that (f, M) describes (E^-, E^+) . Note
 165 that every pair of an embedding and a connection pattern (f, M) describes exactly one
 166 swap (E^-, E^+) . Conversely, for a swap (E^-, E^+) the corresponding embedding f is also
 167 unique (and determined by E^-). However, in case E^- contains incident edges, the swap fits
 168 into more than one matching M (see Fig. 1). See [3] for a more formal description of the
 169 equivalence.

170 The notion of a connection pattern can be extended to represent k' -swaps, for $k' < k$, as
 171 follows. Note that a matching N in the complete graph on the vertex set $[2k]$ corresponds to
 172 an $|N|$ -swap if and only if there is a set $\iota(N) \subseteq [k]$ such that $V(N) = \{2i - 1, 2i: i \in \iota(N)\}$.
 173 For a set $X \subseteq [k]$, by $M[X]$ we denote the swap N such that $\iota(N) = X$. We say that a
 174 connection pattern M *decomposes* into swaps N_1, \dots, N_t when $M = \uplus_{i=1}^t N_i$ and each N_i is
 175 a connection pattern of a swap. The notion of fitting extends to k' -swaps in the natural way.

176 Consider a connection pattern N of a swap, for $V(N) \subseteq [2k]$. We call N *sequential* if
 177 $N \cup \{\{2i - 1, 2i\}: i \in \iota(N)\}$ forms a simple cycle. In particular, every connection pattern
 178 can be decomposed into sequential connection patterns of (possibly shorter) swaps. The
 179 correspondence between sequential swaps and sequential connection pattern is somewhat
 180 delicate, so let us explain it in detail.



■ **Figure 1** A sequential swap (left) which fits two connection patterns (center, right). The pattern in the center is not sequential, while the pattern on the right is sequential. On the left the solid red edges are in E^- , the dashed green edges are in E^+ , and the thin black edges are the remaining edges of the Hamiltonian cycle \mathcal{C} . In the central and right pictures, the dashed green edges form some connection patterns.

181 Let N be a sequential connection pattern, $V(N) \subseteq [2k]$. Recall that for every embedding
 182 f there is exactly one $|N|$ -swap (E^-, E^+) that fits into N . Clearly, this swap is sequential,
 183 since every edge in $\{\{2i - 1, 2i\} : i \in \iota(N)\}$ corresponds to an edge of E^- and every edge in
 184 N corresponds to an edge in E^+ . Thus the resulting set of edges $E^- \cup E^+$ forms a single
 185 closed walk. In particular, if the image of f contains two neighboring indices $i, i + 1 \in [n]$, the
 186 closed walk is not a simple cycle.

187 Conversely, it is possible that a sequential swap fits into a connection pattern which is
 188 not sequential, see Fig. 1 for an example. However, every sequential ℓ -swap (E^-, E^+) fits at
 189 least one sequential connection pattern. This sequential connection pattern is determined by
 190 the closed walk which certifies the sequentiality of the swap. Indeed, let $E^- = \{e_{i_1}, \dots, e_{i_\ell}\}$,
 191 where i_1, \dots, i_ℓ is an increasing sequence. Let $v_0, \dots, v_{2\ell-1}$ be the closed walk alternating
 192 between E^- and E^+ , in particular assume that $E^- = \{v_i v_{i+1} : i \text{ is even}\}$. Consider any
 193 $i = 0, \dots, \ell - 1$ and the corresponding edge $e_{i_j} = v_{2i} v_{2i+1}$ in E^- , for some $j \in [\ell]$. If v_{2i}
 194 is the left endpoint of e_{i_j} , we put $w_{2i} = 2j - 1$ and $w_{2i+1} = 2j$, otherwise $w_{2i} = 2j$ and
 195 $w_{2i+1} = 2j - 1$. Then $w_0, \dots, w_{2\ell-1}$ is a simple cycle and $N = \{w_i w_{i+1} : i \text{ is odd}\}$ is a
 196 sequential connection pattern. By construction, (E^-, E^+) fits N , as required. Keeping in
 197 mind the nuances in the notions of sequential swaps and corresponding sequential connection
 198 patterns, for simplicity, we will often just say ‘a sequential swap M ’ for a matching M , instead
 199 of the more formal ‘a sequential connection pattern M of a swap’.

200 Fix a connection pattern M and let $f: S \rightarrow [n]$ be a partial embedding, for some $S \subseteq [k]$.
 201 For every $j \in S$, let v_{2j-1} and v_{2j} be the left and right endpoint of $e_{f(j)}$, respectively. We
 202 define

$$203 \quad E_f^- = \{e_{f(i)} \mid i \in S\},$$

$$204 \quad E_f^+ = \{\{v_{i'}, v_{j'}\} \mid i, j \in S, i' \in \{2i - 1, 2i\}, j' \in \{2j - 1, 2j\}, \{i', j'\} \in M\}.$$

206 Then, $\text{gain}_M(f) = w(E_f^-) - w(E_f^+)$.

207 3 Fast XP algorithms

208 For every fixed integers k and d , the number of sequential k -swaps in a graph of maximum
 209 degree d is $\mathcal{O}(n)$, and we can enumerate all of them in the same running time. Therefore,
 210 we can find the best improving k -move that can be decomposed into at most c sequential

211 k -swaps in $\mathcal{O}(n^c)$ time. Because c is at most $\lfloor \frac{k}{2} \rfloor$, we obtain an $\mathcal{O}(n^{\lfloor \frac{k}{2} \rfloor})$ -time algorithm for
 212 k -OPT OPTIMIZATION. In what follows, we will improve this naive algorithm. Below we
 213 present a relatively simple algorithm which exploits the range tree data structure [15] and
 214 achieves running time roughly the same as the more sophisticated algorithm of Cygan et
 215 al. [3] for general graphs.

216 ► **Theorem 1.** *For every fixed integers k , c , and d , there is an $\mathcal{O}(n^{\lceil \frac{c}{2} \rceil})$ polylog n -time
 217 algorithm to compute the best improving k -move that can be decomposed into c sequential
 218 swaps in graphs of maximum degree d .*

219 **Proof.** When $c = 1$, we can use the naive algorithm. Suppose $c \geq 2$ and let $h := \lceil \frac{c}{2} \rceil$.

220 For each possible connection pattern M consisting of c sequential swaps, we find the
 221 best embedding as follows. Let $M = \bigcup_{i=1}^c N_i$, where each N_i corresponds to a sequential
 222 swap. We split M into two parts $M_L = \bigcup_{i=1}^h N_i$ and $M_R = \bigcup_{i=h+1}^c N_i$ and we define
 223 $L = \bigcup_{i=1}^h \iota(N_i)$ and $R = \bigcup_{i=h+1}^c \iota(N_i)$. Note that $L \uplus R = [k]$. Let $f_L: L \rightarrow [n]$ and
 224 $f_R: R \rightarrow [n]$ be embeddings of L and R , respectively. The union of these two embeddings
 225 results in an embedding of $[k]$ if and only if the following conditions hold.

- 226 ■ For each $i \in [k-1]$ with $i \in L$ and $i+1 \in R$, $f_L(i) < f_R(i+1)$ holds.
- 227 ■ For each $i \in [k-1]$ with $i \in R$ and $i+1 \in L$, $f_R(i) < f_L(i+1)$ holds.

228 We can efficiently compute a pair of embeddings satisfying these conditions using an ortho-
 229 gonal range maximum data structure as follows. Let $\{l_1, \dots, l_p\} = \{i: l_i \in L \text{ and } l_i + 1 \in R\}$
 230 and let $\{r_1, \dots, r_q\} = \{i: r_i - 1 \in R \text{ and } r_i \in L\}$. We first enumerate all the $|L|$ -swaps that
 231 fit into M_L and all the $|R|$ -swaps that fit into M_R , in $\mathcal{O}(n^h)$ time. For each such $|L|$ -swap
 232 (f_L, M_L) , we create a $(p+q)$ -dimensional point $(f_L(l_1), \dots, f_L(l_p), f_L(r_1), \dots, f_L(r_q))$ with
 233 a priority gain $_{M_L}(f_L)$, and we collect these points into a data structure. It stores $\mathcal{O}(n^h)$
 234 points. For each $|R|$ -swap (f_R, M_R) , we query for the embedding f_L of maximum priority
 235 satisfying $f_L(l_i) < f_R(l_i + 1)$ for every $i \in [p]$ and $f_R(r_i - 1) < f_L(r_i)$ for every $i \in [q]$, and
 236 we answer the pair maximizing the total gain, i.e., the sum $\text{gain}_{M_L}(f_L) + \text{gain}_{M_R}(f_R)$. Using
 237 the range tree data structure [15], each query takes $\mathcal{O}(\log^{p+q} n^h) = \mathcal{O}(\text{polylog } n)$ time, so
 238 the total running time is $\mathcal{O}(n^h \text{ polylog } n)$. ◀

239 Since $c \leq \lfloor \frac{k}{2} \rfloor$ we get the following corollary.

240 ► **Corollary 2.** *For all fixed integers k and d , k -OPT OPTIMIZATION in graphs of maximum
 241 degree d can be solved in time $\mathcal{O}(n^{\lceil \frac{k-1}{4} \rceil})$ polylog n .*

242 Let us take another look at the proof of Theorem 1. Recall that for merging embeddings
 243 f_L and f_R , we were interested only in values $f_L(i)$ for $i \in L$ such that $i+1 \in R$ or $i-1 \in R$.
 244 The embeddings of the remaining elements of L were forgotten at that stage, but we knew
 245 that it is possible to embed them and we stored the gain of embedding them. This suggests
 246 the following, different approach. We decompose the connection pattern into sequential
 247 swaps and we scan the swaps in a carefully chosen order. Assume we scanned t swaps already
 248 and there are $c-t$ swaps ahead. Assume that only $p \ll t$ of the t ‘boundary’ swaps interact
 249 with the remaining $c-t$ swaps, where two swaps N_1 and N_2 interact when there is $i \in \iota(N_1)$
 250 such that $i-1 \in \iota(N_2)$ or $i+1 \in \iota(N_2)$. Then it suffices to compute, for every embedding
 251 f_L of the p swaps, the gain of the best (i.e., giving the highest gain) embedding g_L of the t
 252 swaps, such that f_L matches g_L on the boundary swaps. This amounts to $\mathcal{O}(n^p)$ values to
 253 compute, since each sequential swap can be embedded in $\mathcal{O}(n)$ ways, if k and the maximum
 254 degree are $\mathcal{O}(1)$. The idea is to (1) compute these values fast (in time linear in their number)
 255 using analogous values computed for the prefix of $t-1$ swaps, (2) find an order of swaps

so that p is always small, namely $p \leq (23/135 + \epsilon_k)k$. The readers familiar with the notion of pathwidth recognize that p here is just the pathwidth of the graph obtained from the path $1, 2, \dots, k$ by identifying vertices in the set $\iota(N)$ for every sequential swap N in M , and that (2) is just dynamic programming over path decomposition. The resulting algorithm is summarized in Theorem 3, and due to space limits, its formal proof is skipped here and will be included in the full version.

► **Theorem 3.** *For all fixed integers k and d , k -OPT OPTIMIZATION in graphs of maximum degree d can be solved in time $\mathcal{O}(n^{(23/135 + \epsilon_k)k}) = \mathcal{O}(n^{(0.1704 + \epsilon_k)k})$, where $\lim_{k \rightarrow \infty} \epsilon_k = 0$.*

4 Fast algorithms for small k

Note that the algorithm for k -OPT OPTIMIZATION from Corollary 2 is quasi-linear for $k \leq 5$. In this section we extend the quasi-linear-time solvability to $k \leq 7$ for k -OPT DETECTION. Under an additional assumption of bounded weights, we are able to reach quasi-linear time for $k = 8$ as well, but the details of this part are deferred to the full version because of space constraints. To be precise, in the $k = 7$ case we prove the following stronger statement than just finding an arbitrary improving k -move.

► **Theorem 4.** *For $k \leq 7$, there is a quasi-linear-time algorithm to compute the best improving k -move in bounded-degree graphs under the assumption that there are no improving k' -moves for $k' < k$.*

We say that a connection pattern M of k -swaps is *reducible* if it can be decomposed into two moves. Note that if M is improving, then at least one of the two moves is improving, contradicting the assumption of Theorem 4.

▷ **Observation 5.** If there are no improving k' -moves for $k' < k$, then no improving k -swap fits into a reducible connection pattern.

Before we formulate our algorithm, we need two lemmas. We can prove these lemmas by case analysis, and because of the space constraints, their proofs are skipped here and will be included in the full version. Let $M[X]$ and $M[Y]$ be two swaps in a connection pattern M , for some disjoint $X, Y \subseteq [k]$. *Interaction* between $M[X]$ and $M[Y]$ is any $i \in [k-1]$ such that $i \in X$ and $i+1 \in Y$ or $i \in Y$ and $i+1 \in X$.

► **Lemma 6.** *For any $k \geq 6$, there is no feasible and irreducible connection k -pattern that contains two 2-swaps that interact at least twice.*

Let M be a connection pattern, i.e., a perfect matching on vertices $[2k]$. We say that M' is obtained from M by swapping i and $i+1$, for $i \in [k]$, when M' is obtained from M by swapping the mates of $2i-1$ and $2i+1$ and swapping the mates of $2i$ and $2i+2$.

► **Lemma 7.** *Let M be a feasible irreducible connection k -pattern. Assume that M decomposes into three sequential swaps $M[X]$, $M[Y]$, and $M[Z]$, such that $|X| = |Y| = 2$. If there is exactly one index $i \in [k-1]$ with $i \in X$ and $i+1 \in Y$ or $i \in Y$ and $i+1 \in X$, the connection pattern M' obtained from M by swapping i and $i+1$ is either feasible or reducible.*

Now we are ready to describe the algorithm from Theorem 4 (see also Pseudocode 1). For each feasible and irreducible connection k -pattern M , we compute the best embedding as follows. If M consists of at most two sequential swaps, we can use the algorithm in Theorem 1. Otherwise, M consists of three sequential swaps $M[X]$, $M[Y]$, $M[Z]$ such that

Pseudocode 1 Quasi-linear-time algorithm for $k \leq 7$

```

1: for each feasible irreducible connection  $k$ -pattern  $M$  do
2:   if  $M$  consists of at most two sequential swaps then
3:     Apply the algorithm in Theorem 1.
4:   else
5:     Let  $M = M[X] \uplus M[Y] \uplus M[Z]$  where  $|X| = |Y| = 2$  and  $|Z| = k - 4$ .
6:     if there are no interactions between  $X$  and  $Y$  then
7:       for each embedding  $f_Z$  for  $Z$  do
8:         Independently compute the best embeddings  $f_X$  for  $X$  and  $f_Y$  for  $Y$ .
9:       else
10:        Relax the constraint  $f_X(i) < f_Y(i+1)$  to  $f_X(i) \neq f_Y(i+1)$ .
11:        for each embedding  $f_Z$  for  $Z$  do
12:          Compute the best pair  $(f_X, f_Y)$  satisfying the relaxed constraints.

```


297 $X \uplus Y \uplus Z = [k]$, $|X| = |Y| = 2$ and $|Z| = k - 4$. For each embedding $f_X : X \rightarrow [n]$ of
298 $X = \{i, j\}$ we create a 2-dimensional point $(f_X(i), f_X(j))$ with priority $\text{gain}_X(f_X)$ and we
299 put all the points in a range tree data structure D_X [15]. We build an analogous data
300 structure for Y . Next, for each embedding f_Z for Z , we compute the best pair of embeddings
301 (f_X, f_Y) for X and Y as follows.

302 If there are no interactions between X and Y , we can find the best pair in $\mathcal{O}(\text{polylog } n)$
303 time by independently picking the best embeddings for X and Y by querying the range trees
304 D_X and D_Y . Indeed, first note that there is no index $i \in [k-1]$ such that $X = \{i, i+1\}$ because
305 in such a case, both the 2-swap and the remaining $(k-2)$ -swap have to be feasible (similarly
306 for Y). Since there are no interactions between X and Y , we must have $i-1 \in Z \cup \{0\}$ and
307 $i+1 \in Z \cup \{k+1\}$ for every $i \in X \cup Y$. To find the best embedding f_X of $X = \{i, j\}$, we query
308 D_X with the constraints $f_Z(i-1) < f_X(i) < f_Z(i+1)$ and $f_Z(j-1) < f_X(j) < f_Z(j+1)$,
309 where we define $f_Z(0) := 0$ and $f_Z(k+1) := n+1$. We proceed analogously for Y .

310 Finally, assume there are interactions between X and Y , so from Lemma 6, there is exactly
311 one interaction. W.l.o.g. $i \in X$ and $i+1 \in Y$. Note that $i-1 \in Z \cup \{0\}$ and $i+2 \in Z \cup \{k+1\}$.
312 We first relax the constraint $f_Z(i-1) < f_X(i) < f_Y(i+1) < f_Z(i+2)$, where we define
313 $f_Z(0) := 0$ and $f_Z(k+1) := n+1$, to three constraints $f_Z(i-1) < f_X(i) < f_Z(i+2)$,
314 $f_Z(i-1) < f_Y(i+1) < f_Z(i+2)$, and $f_X(i) \neq f_Y(i+1)$. We then drop the disturbing
315 inequality constraint $f_X(i) \neq f_Y(i+1)$ by color-coding¹. We color each vertex in $[n]$ in red
316 or blue, and we independently pick the best embedding for X (resp. Y) that uses only red
317 (resp. blue) vertices. By using a family of perfect hash functions [5], we can construct a set
318 of $\mathcal{O}(\log^2 n)$ colorings such that, for every pair of embeddings f_X and f_Y , there is at least
319 one coloring that colors all the vertices in f_X red and all the vertices in f_Y blue.

320 We now obtain the best pair of embeddings (f_X, f_Y) satisfying the relaxed constraints.
321 If the obtained k -swap is not improving, we immediately know that there are no improving
322 k -moves that fit into M . If it is improving and satisfies the original constraint, we are
323 done. Finally, if it is improving but does not satisfy the original constraint, it fits into the
324 connection pattern M' that is obtained from M by swapping i and $i+1$. By Lemma 7, M'

¹ Instead of color-coding, we can adapt the range tree to support orthogonal range maximum queries with an additional constraint of the form $x \neq i$ by keeping one additional point in each node. With this approach, we can avoid the additional $\log^2 n$ factor. Because this paper does not focus on optimizing the polylog n factor, we do not touch on the details.



■ **Figure 2** An instance of TRIANGLE DETECTION

325 is either feasible or reducible. Because no improving k -swaps fit into reducible connection
 326 patterns, M' has to be feasible. We therefore obtain a k -move that is as good as the best
 327 k -move that fits into M . This completes the proof of Theorem 4.

328 We finally consider the case of $k = 8$. Note that, because Lemma 6 and 7 do not
 329 assume $k \leq 7$, the above algorithm can also compute the best improving k -move that can be
 330 decomposed into three sequential swaps of size $(2, 2, k - 4)$ for any fixed k under the same
 331 assumption. Moreover, any connection patterns of 8-moves consisting of four 2-swaps are
 332 reducible because it always induces a pair of two 2-swaps that interact at least twice. The
 333 remaining case for $k = 8$ is only when the 8-move can be decomposed into three sequential
 334 swaps of size $(2, 3, 3)$. In order to tackle this case, we exploit the bounded-weight assumption
 335 as follows. For each connection pattern $M = M[X] \uplus M[Y] \uplus M[Z]$ with $|X| = 2$ and
 336 $|Y| = |Z| = 3$, and for each embedding f_Z for Z , we want to compute the best pair of
 337 embeddings f_X for X and f_Y for Y . When all the weights are integers from $[W]$, the gain
 338 of $(f_X, M[X])$ is an integer from $[-2W, 2W]$, and the gain $(f_Y, M[Y])$ is an integer from
 339 $[-3W, 3W]$. We therefore have only $\mathcal{O}(W^2)$ pairs of gains. By guessing the pair of gains, the
 340 query of finding the *best* pair can be reduced to the query of finding an *arbitrary* pair, and
 341 the latter query can be efficiently answered by adapting the range tree. This leads to the
 342 following algorithm, whose detailed description is skipped here and will be included in the
 343 full version.

344 ► **Theorem 8.** *When all the weights are integers from $[W]$, there is an $\mathcal{O}(W^2 n \text{ polylog } n)$ -*
 345 *time algorithm to compute the best improving 8-move under the assumption that there are no*
 346 *improving k' -moves for $k' < 8$.*

347 5 Lower bound for $k = 9$

348 The starting point for our reduction is the following problem (see Fig. 2 for an exemplary
 349 instance).

350 TRIANGLE DETECTION **Parameter:** $m := |E(H)|$.
Input: An undirected graph H whose vertex set $V(H)$ is partitioned into $A \cup B \cup C$.
Question: Is there a triple $(a, b, c) \in A \times B \times C$ such that $\{ab, ac, bc\} \subseteq E(H)$?

351 We assume without loss of generality that A , B , and C are three independent sets, so that
 352 finding such a triple is equivalent to finding a triangle in the graph H . By simple reductions
 353 that incur only a constant blow-up in the number of vertices and edges, this problem is
 354 equivalent to determining whether a graph has a triangle or not.

355 ▷ **Assumption 1** (Triangle hypothesis [1]). There is a fixed $\delta > 0$ such that, in the Word RAM
 356 model with words of $\mathcal{O}(\log n)$ bits, any algorithm requires $m^{1+\delta-o(1)}$ time in expectation to
 357 detect whether an m -edge graph contains a triangle.

358 It should be noted that one can solve TRIANGLE DETECTION in time $\mathcal{O}(n^\omega)$ where n is
 359 the number of vertices and $\omega \leq 2.373$ is the best-known exponent for matrix multiplication.
 360 Alon et al. [2] found an elegant win-win argument to solve TRIANGLE DETECTION in time
 361 $\mathcal{O}(m^{\frac{2\omega}{\omega+1}})$: the 3-vertex paths in which the middle vertex has degree less than $m^{\frac{\omega-1}{\omega+1}}$ can be
 362 listed in time $\mathcal{O}(m \cdot m^{\frac{\omega-1}{\omega+1}}) = \mathcal{O}(m^{\frac{2\omega}{\omega+1}})$, and for each, one can check if they form a triangle,
 363 whereas the number of vertices of degree greater than $m^{\frac{\omega-1}{\omega+1}}$ is at most $m^{\frac{2}{\omega+1}}$, so one can
 364 detect a triangle in time $\mathcal{O}(m^{\frac{2\omega}{\omega+1}})$ in the subgraph that they induce. After more than two
 365 decades, this is still the best worst-case running time (when $n^\omega = \Omega(m^{\frac{2\omega}{\omega+1}})$). This suggests
 366 that the triangle hypothesis is likely to hold. Moreover, if one thinks that the above scheme
 367 yields the best possible running time and that ω will eventually reach 2, then exponent $4/3$
 368 could be the *right answer* for TRIANGLE DETECTION parameterized by the number of edges.
 369 The following is implied by [1, Conjecture 2] (since $\omega \geq 2$), in the regime $m = \Theta(n^{3/2})$ (so
 370 that $\mathcal{O}(n^2)$ and $\mathcal{O}(m^{4/3})$ coincide).

371 \triangleright **Assumption 2.** In the Word RAM model with words of $\mathcal{O}(\log n)$ bits, any algorithm
 372 requires $m^{4/3-o(1)}$ time in expectation to detect whether an m -edge $\Theta(m^{2/3})$ -node graph
 373 contains a triangle.

374 We show that SUBCUBIC 9-OPT DETECTION parameterized by the number of vertices
 375 is as hard as TRIANGLE DETECTION parameterized by the number of edges, by providing
 376 a linear-time reduction from the latter to the former. In light of Theorem 4, this implies
 377 that BOUNDED-DEGREE 8-OPT DETECTION is the only remaining open case where a quasi-
 378 linear algorithm is not known but also not ruled out by a standard fine-grained complexity
 379 assumption.

380 \blacktriangleright **Lemma 9.** *There is an $\mathcal{O}(m)$ -time reduction from TRIANGLE DETECTION on m -edge
 381 graphs to SUBCUBIC 9-OPT DETECTION on $\mathcal{O}(m)$ -vertex undirected graphs with edge weights
 382 in $\{1, 2\}$.*

383 **Proof.** From a tripartitioned instance of TRIANGLE DETECTION $H = (A \cup B \cup C, E(H))$
 384 with m edges, we build a subcubic graph G with $\Theta(m)$ vertices, an edge-weight function
 385 $w : E(G) \rightarrow \{1, 2\}$, and a Hamiltonian cycle \mathcal{C} . From \mathcal{C} , there is a swap of up to 9 edges (i.e.,
 386 up to 9 deletions and the same number of additions) which results in a lighter Hamiltonian
 387 cycle if and only if H has a triangle.

388 Overall construction of G .

389 We will build G by adding chords to the cycle \mathcal{C} . Henceforth, a *chord* is an edge of G
 390 which is not in \mathcal{C} . It is helpful to think \mathcal{C} as a (subdivided) triangle whose three sides
 391 correspond to A , B , and C , which we call the A -side (left), B -side (right), and C -side
 392 (bottom), respectively. We will only name the edges of G (and not the vertices), since the
 393 problem is more efficiently described in terms of edges. We will define some sequential
 394 3-swaps (we recall that a sequential i -swap is a closed walk of length $2i$ alternating edges
 395 of $E(\mathcal{C})$ and edges of $E(G) \setminus E(\mathcal{C})$). Eventually, all the edges that are not in a described
 396 sequential 3-swap are incident to a vertex of degree 2, making them undeletable. (One can
 397 also enforce that by subdividing every irrelevant edge once.)

398 The improving 9-move, should there be a triangle abc in H , will consist of a sequence of
 399 three 3-swaps. More precisely, it consists of one improving 3-swap, which splits \mathcal{C} into three
 400 cycles respectively containing:

401 **(1)** a part of the vertex gadget of some $a \in A$,

402 (2) the part of the B -side below the vertex gadget of b , as well as the C -side, and
 403 (3) the part of the B -side above the vertex gadget of some $b \in N_H(a) \cap B$.
 404 This decreases the total weight by 1. Then a neutral 3-swap reconnects (1) and (2) together,
 405 but also detaches (4) a part of the vertex gadget of some $c \in N_H(a) \cap C$. Finally a neutral
 406 3-swap glues (3), (1)+(2), and (4) together, provided $bc \in E(H)$. This results in a new
 407 Hamiltonian cycle of length $w(\mathcal{C}) - 1$.

408 There will be relatively few edges of weight 2. To simplify the presentation, every edge is
 409 of weight 1, unless specified otherwise. Let \vec{H} be the directed graph obtained from H by
 410 orienting its edges from A to B , from B to C , and from C to A . Note that finding a directed
 411 triangle in \vec{H} is equivalent to finding a triangle in H .

412 Vertex scopes, extended scopes, and nested chords.

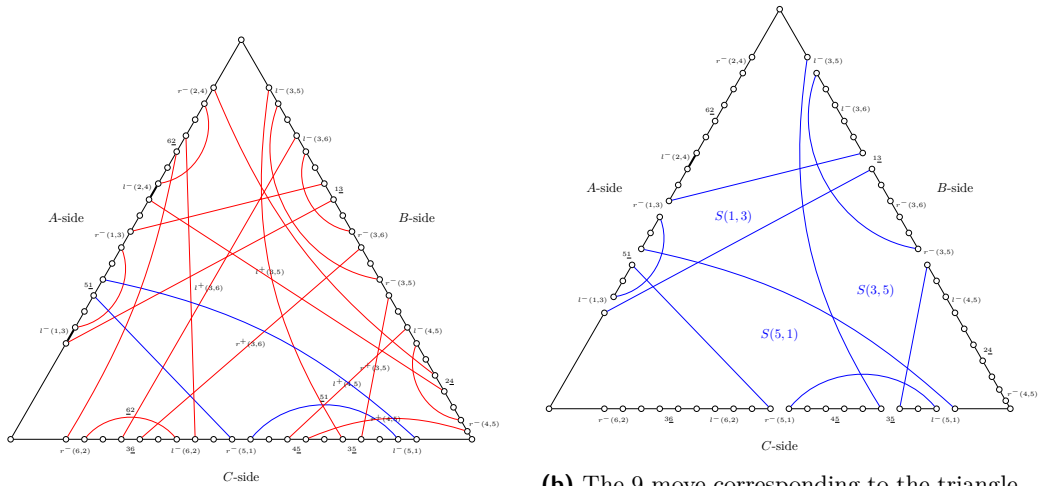
413 For $(X, Y) \in \{(A, B), (B, C), (C, A)\}$, we set $Z := \{A, B, C\} \setminus \{X, Y\}$ and we do the
 414 following as a preparatory step to encode the arcs of \vec{H} . Each vertex $v \in X$ is given
 415 a (pairwise vertex-disjoint) subpath I_v of \mathcal{C} , called the *extended scope* of v , with $|I_v| :=$
 416 $6(|N_H(v) \cap Y|) + 3(|N_H(v) \cap Z|) - 1$ vertices. We think of I_v as being displayed from left to
 417 right with the leftmost vertex of index 1, and the rightmost one of index $|I_v|$. The extended
 418 scopes of the vertices of A , B , and C occupy respectively the A -side, B -side, and C -side. In
 419 what follows, it will be more convenient to have a *circular* notion of *left* and *right*. Starting
 420 from the bottom corner of the A -side, and going clockwise to the top corner of the A -side,
 421 then down to the bottom corner of the B -side, the relative left and right within the A -side
 422 and the B -side coincide with the usual notion as displayed in Figure 3a. But then closing
 423 the loop from the right corner of the C -side to its left corner, left and right are switched: the
 424 closer to the bottom corner of A (resp. B), the more “right” (resp. “left”).

425 Each vertex $v \in X$ has $|N_H(v) \cap Y|$ nested chords spaced out every three vertices. More
 426 precisely, the second vertex of I_v is adjacent to the penultimate, the fifth to the one of index
 427 $|I_v| - 4$, the eighth to the one of index $|I_v| - 7$, and so on, until $|N_H(v) \cap Y|$ chords are
 428 drawn. Each of these chords is associated to an edge $vy \in E(\{v\}, Y)$, and is denoted by \underline{vy} .
 429 A vertex just to the right of the left endpoint, or just to the left of the right endpoint, of
 430 such a chord will remain of degree 2 in G . This is the case of the vertices of index 3, 6, ...
 431 and $|I_v| - 2, |I_v| - 5, \dots$ in I_v . We call $l^-(v, y)$ (resp. $r^-(v, y)$) the edge of I_v incident to
 432 both the left endpoint of \underline{vy} and the vertex just to its left (resp. right endpoint of \underline{vy} and
 433 the vertex just to its right). Both endpoints of $l^-(v, y)$ and of $r^-(v, y)$ will eventually have
 434 degree 3 in G .

435 The chord linking the most distant vertices in I_v is called the *outermost* chord, while
 436 the one linking the closest pair is called the *innermost* chord. We also say that a chord e is
 437 *wider* than a chord e' if e links a farther pair on I_v than e' does. The central path $J_v \subset I_v$
 438 on $|I_v| - (6|N_H(v) \cap Y| - 4) = 3(|N_H(v) \cap Z| + 1)$ vertices, surrounded by the innermost
 439 chord, is called the *scope* of v . We map in one-to-one correspondence the edges of $E(\{v\}, Z)$
 440 to every three edges of J_v starting from the third edge (that is, the third, sixth, and so on).
 441 Note that we have the exact space to do so, since $|J_v| = 3(|N_H(v) \cap Z| + 1)$. We denote by
 442 \underline{zv} the edge in J_v corresponding to the edge $vz \in E(\{v\}, Z)$.

443 Encoding the arcs of \vec{H} .

444 The last step to encode the arcs of \vec{H} , or equivalently the edges of H , is the following.
 445 Keeping the notations of the previous paragraphs, for every edge $xy \in E(X, Y)$, we add two
 446 chords (of weight 1): one chord $l^+(x, y)$ between the left endpoint of $l^-(x, y)$ and the right



(a) The construction for the instance of Figure 2. Edges of \mathcal{C} are in black, chords are in red, bold edges are the ones with weight 2. The three chords in blue are the edges to add to perform the neutral 3-swap $S(5, 1)$ of $S(C, A)$.

(b) The 9-move corresponding to the triangle 135 results in a Hamiltonian cycle using one less edge of weight 2. Note that after the swaps $S(1, 3)$ and $S(5, 1)$ are performed, the only 3-swap that can reconnect the three cycles into one, is $S(3, 5)$, implying the existence of the edge 35, and thereby of the triangle 135.

■ Figure 3 Illustration of the reduction (left) and of a potential solution (right).

447 endpoint of xy and one chord $r^+(x, y)$ between the right endpoint of $r^-(x, y)$ and the left
 448 endpoint of xy . We finish the construction of G (and \mathcal{C}) by subdividing each edge between
 449 consecutive extended scopes once, to make the resulting edges undeletable. The edges $l^-(a, b)$
 450 for $(a, b) \in A \times B$ get weight 2, while all the other edges of $E(G)$ get weight 1. This finishes
 451 the construction of (G, w, \mathcal{C}) . See Figure 3a for an illustration.

452 **Improving and neutral 3-swaps.**

453 For each $(x, y) \in E(\vec{H})$, we denote by $S(x, y)$ the 3-swap $(\{xy, l^-(x, y), r^-(x, y)\}, \{xy, l^+(x, y),$
 454 $r^+(x, y)\})$. For $(X, Y) \in \{(A, B), (B, C), (C, A)\}$, we define the set of 3-swaps $S(X, Y) :=$
 455 $\bigcup_{xy \in E(X, Y)} S(x, y)$, and $\mathcal{S} := S(A, B) \cup S(B, C) \cup S(C, A)$.

456 Note that all the 3-swaps of $S(A, B)$ are improving. They gain 1 since $l^-(a, b)$ has weight
 457 2 for any $(a, b) \in A \times B$. On the other hand, all the 3-swaps of $S(B, C)$ and $S(C, A)$ are
 458 neutral. The edges added in swaps of \mathcal{S} partition the chords of G , and the open neighborhood
 459 of the six vertices involved in every swap are six vertices of degree 2 in G . Therefore, all the
 460 possible swaps are in the set \mathcal{S} , they are on vertex-disjoint sets of vertices, and any move is
 461 a sequence of 3-swaps of \mathcal{S} .

462 The vertices of \mathcal{C} are incident to at most one chord. Hence the graph G is subcubic. It
 463 has $\sum_{v \in V(H)} 1 + |I_v| \leq 9|E(H)| + |V(H)| = \Theta(m)$ vertices and (G, w, \mathcal{C}) takes $\Theta(m)$ -time
 464 to build. To summarize, we defined a linear reduction from TRIANGLE DETECTION with
 465 parameter m to SUBCUBIC 9-OPT DETECTION with parameter n . So a quasi-linear algorithm
 466 for the latter would yield an unlikely quasi-linear algorithm for the former. We now check
 467 that the reduction is correct.

468 **A triangle in H implies an improving 9-move for (G, w, \mathcal{C}) .**

469 Let abc be a triangle in H . In particular, all three swaps $S(a, b)$, $S(b, c)$, and $S(c, a)$ exist.
 470 Performing these three 3-swaps results in a spanning union of (vertex-disjoint) cycles, whose
 471 total weight is $w(\mathcal{C}) - 1$. Indeed $S(a, b)$ is swap of weight -1 , while $S(b, c)$, and $S(c, a)$ are
 472 both neutral.

473 We thus only need to show that the three swaps result in a connected graph (hence,
 474 Hamiltonian cycle of lighter weight). By performing the 3-swap $S(a, b)$, we create three
 475 components: (1) one on a vertex set $K_{a,b}$ such that $J_a \subseteq K_{a,b} \subseteq I_a$, (2) one containing
 476 the scopes of vertices of the B -side to the right (lower part) of the scope of b , and (3) one
 477 containing the scopes of vertices of the B -side to the left (upper part) of the scope of b . Then
 478 the swap $S(c, a)$ glues (1) and (2) together, but also disconnects (4) a cycle on a vertex set
 479 $K_{c,a}$ such that $J_c \subseteq K_{c,a} \subseteq I_c$. At this point, there are three cycles: (3), (1)+(2), and (4).
 480 It turns out that the 3-swap $S(b, c)$ deletes exactly one edge in each of these three cycles:
 481 $b\bar{c}$ in (4), $l^-(b, c)$ in (3), and $r^-(b, c)$ in (1)+(2). Therefore, $S(b, c)$ reconnects these three
 482 components into one Hamiltonian cycle.

483 **An improving k -move for (G, w, \mathcal{C}) with $k \leq 9$ implies a triangle in H .**

484 We assume that there is an improving k -move $\mathcal{M} = (E^-, E^+)$ for (G, w, \mathcal{C}) with $k \leq 9$.
 485 Being improving, the k -move has to contain at least one improving 3-swap of $S(A, B)$.
 486 Let $S(a, b)$ be a 3-swap of $S(A, B)$ in \mathcal{M} such that for every other (improving) 3-swap
 487 $S(a, b')$ in \mathcal{M} , the chord $\underline{ab'}$ is wider than \underline{ab} . Since $S(a, b)$ exists, it holds in particular that
 488 $ab \in E(H)$. Performing $S(a, b)$ results in the union of three cycles: (1) on a vertex set $K_{a,b}$
 489 with $J_a \subseteq K_{a,b} \subseteq I_a$, and cycles (2) and (3) as described in the previous paragraph.

490 By the choice of b , the only remaining swaps of \mathcal{M} touching $K_{a,b}$ are in $S(C, A)$. So
 491 \mathcal{M} has to contain a neutral 3-swap $S(c, a)$ for some $c \in C$. This implies that $ac \in E(H)$.
 492 Performing this swap results in three cycles: (3), (1)+(2), and (4), as described above. To
 493 reconnect all three components into one Hamiltonian cycle, the 3-swap has to delete exactly
 494 one edge in (3), (1)+(2), and (4). The only 3-swap that does so is $S(b, c)$. This finally implies
 495 that $bc \in E(H)$. Thus abc is a triangle in H . ◀

496 We obtain the following theorem as a direct consequence of the previous lemma.

497 ▶ **Theorem 10.** SUBCUBIC 9-OPT DETECTION *requires time:*

- 498 (1) $n^{1+\delta-o(1)}$ for a fixed $\delta > 0$, under the triangle hypothesis, and
 499 (2) $n^{4/3-o(1)}$, under the strong triangle hypothesis,
 500 in expectation, even in undirected graphs with edge weights in $\{1, 2\}$.

501 If we use general integral weights and not just $\{1, 2\}$, we can show a stronger lower
 502 bound, by reducing from NEGATIVE EDGE-WEIGHTED TRIANGLE. Again, we can assume
 503 that the instance is partitioned into three sets A, B, C , and we look for a triangle abc
 504 such that $w'(ab) + w'(bc) + w'(ac) < 0$, where w' gives an integral weight to each edge.
 505 A truly subcubic (in the number of vertices) algorithm for this problem would imply one
 506 for ALL-PAIRS SHORTEST PATHS, which would be considered a major breakthrough. The
 507 assumption that such an algorithm is not possible is called the APSP hypothesis.

508 We only change the above construction in the weight of the edges $l^-(x, y)$. Now each edge
 509 $l^-(x, y)$ gets weight $-w'(xy)$. From a NEGATIVE EDGE-WEIGHTED TRIANGLE-instance with
 510 n vertices, we obtain an equivalent instance of SUBCUBIC 9-OPT DETECTION with $\mathcal{O}(n^2)$
 511 vertices, in time $\mathcal{O}(n^2)$. So we derive the following.

512 ► **Theorem 11.** SUBCUBIC 9-OPT DETECTION requires time $n^{3/2-o(1)}$, under the APSP
 513 hypothesis.

514 6 Lower bound for varying k

515 In this section we describe the main ideas behind the lower bound for k -OPT DETECTION
 516 in subcubic graphs for varying k . The details are deferred to the full version due to space
 517 restrictions. The overall approach is similar to the lower bound of Guo et al. [6], in that
 518 we give a linear-parameter reduction from the k -PARTITIONED SUBGRAPH ISOMORPHISM
 519 problem parameterized by the number of edges k . Marx [14] proved that, assuming the
 520 Exponential Time Hypothesis, the problem cannot be solved in time $f(k) \cdot n^{o(k/\log k)}$ for any
 521 function f .

522 The instance created in the reduction of Guo et al. [6] may contain vertices of arbitrarily
 523 large degrees. To obtain such a reduction to k -OPT DETECTION in subcubic graphs, an
 524 essential ingredient is a *choice gadget* with terminal pairs $(x_0, y_0), \dots, (x_\ell, y_\ell)$ which enforces
 525 that sufficiently cheap Hamiltonian cycles that enter at x_i , must leave via the corresponding y_i .
 526 The gadget can be implemented by suitable weight settings and vertices of degree at most three.
 527 This gadget allows us to enforce synchronization properties, which enforce that an improved
 528 Hamiltonian cycle first selects which vertices to use in the image of the subgraph isomorphism,
 529 and then selects incident edges for each selected vertex. By carefully coordinating the gadgets,
 530 this allows us to implement the hardness proof by an edge selector strategy. It leads to a
 531 proof of the following theorem.

532 ► **Theorem 12.** *There is no function f for which k -OPT DETECTION on n -vertex graphs of*
 533 *maximum degree 3 with edge weights in $\{1, 2\}$ can be solved in time $f(k) \cdot n^{o(k/\log k)}$, unless*
 534 *ETH fails.*

535 We remark that the lower bound also holds for *permissive* local search algorithms which
 536 output an improved Hamiltonian cycle of arbitrarily large Hamming distance to the starting
 537 cycle \mathcal{C} , if a cheaper cycle exists in the k -OPT neighborhood of \mathcal{C} .

538 — References —

- 539 1 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower
 540 bounds for dynamic problems. In *Proc. 55th FOCS*, pages 434–443. IEEE Computer Society,
 541 2014. doi:10.1109/FOCS.2014.53.
- 542 2 Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles.
 543 *Algorithmica*, 17(3):209–223, 1997. doi:10.1007/BF02523189.
- 544 3 Marek Cygan, Lukasz Kowalik, and Arkadiusz Socala. Improving TSP tours using dynamic
 545 programming over tree decompositions. In *Proc. 25th ESA*, volume 87 of *LIPICs*, pages
 546 30:1–30:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.
 547 ESA.2017.30.
- 548 4 Mark de Berg, Kevin Buchin, Bart M. P. Jansen, and Gerhard J. Woeginger. Fine-grained
 549 complexity analysis of two classic TSP variants. In *ICALP*, volume 55 of *LIPICs*, pages
 550 5:1–5:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- 551 5 Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with $O(1)$
 552 worst case access time. *J. ACM*, 31(3):538–544, 1984. doi:10.1145/828.1884.
- 553 6 Jiong Guo, Sepp Hartung, Rolf Niedermeier, and Ondrej Suchý. The parameterized complexity
 554 of local search for TSP, more refined. *Algorithmica*, 67(1):89–110, 2013. doi:10.1007/
 555 s00453-012-9685-8.

- 556 7 Michael Held and Richard M. Karp. The traveling-salesman problem and minimum spanning
557 trees: Part II. *Math. Program.*, 1(1):6–25, 1971.
- 558 8 Keld Helsgaun. An effective implementation of the Lin-Kernighan traveling salesman heur-
559 istic. *European Journal of Operational Research*, 126(1):106 – 130, 2000. doi:10.1016/
560 S0377-2217(99)00284-2.
- 561 9 Keld Helsgaun. General k -opt submoves for the Lin-Kernighan TSP heuristic. *Math. Program.*
562 *Comput.*, 1(2-3):119–163, 2009.
- 563 10 D. S. Johnson and L. A. McGeoch. Experimental analysis of heuristics for the STSP. In
564 G. Gutin and A. Punnen, editors, *The Traveling Salesman Problem and its Variations*, pages
565 369–443. Kluwer Academic Publishers, Dordrecht, 2002.
- 566 11 D.S. Johnson and L.A McGeoch. The traveling salesman problem: A case study in local
567 optimization. In E. Aarts and J.K. Lenstra, editors, *Local search in combinatorial optimization*,
568 pages 215–310. Wiley, Chichester, 1997.
- 569 12 S. Lin and Brian W. Kernighan. An effective heuristic algorithm for the traveling-salesman
570 problem. *Operations Research*, 21(2):498–516, 1973. doi:10.1287/opre.21.2.498.
- 571 13 Dániel Marx. Searching the k -change neighborhood for TSP is $W[1]$ -hard. *Oper. Res. Lett.*,
572 36(1):31–36, 2008. doi:10.1016/j.orl.2007.02.008.
- 573 14 Dániel Marx. Can you beat treewidth? *Theory of Computing*, 6(1):85–112, 2010. doi:
574 10.4086/toc.2010.v006a005.
- 575 15 Franco P. Preparata and Michael Ian Shamos. *Computational Geometry - An Introduction*.
576 Texts and Monographs in Computer Science. Springer, 1985.