Covering Tours and Cycle Covers with Turn Costs: Hardness and Approximation

Sándor P. Fekete and Dominik Krupke(✉)

Department of Computer Science, TU Braunschweig,
38106 Braunschweig, Germany
{s.fekete,d.krupke}@tu-bs.de

Abstract. We investigate a variety of geometric problems of finding tours and cycle covers with minimum turn cost, which have been studied in the past, with complexity and approximation results, and open problems dating back to work by Arkin et al. in 2001. Many new practical applications have spawned variants: For *full coverage*, every point has to be covered, for *subset coverage*, specific points have to be covered, and for *penalty coverage*, points may be left uncovered by incurring a penalty. We make a number of contributions. We first show that finding a minimum-turn (full) cycle cover is NP-hard even in 2-dimensional grid graphs, solving the long-standing open *Problem 53* in *The Open Problems Project* edited by Demaine, Mitchell and O'Rourke. We also prove NP-hardness of finding a *subset* cycle cover of minimum turn cost in *thin* grid graphs, for which Arkin et al. gave a polynomial-time algorithm for full coverage; this shows that their boundary techniques cannot be applied to compute exact solutions for subset and penalty variants.

On the positive side, we establish the first constant-factor approximation algorithms for all considered subset and penalty problem variants for very general classes of instances, making use of LP/IP techniques. For these problems with many possible edge directions (and thus, turn angles, such as in hexagonal grids or higher-dimensional variants), our approximation factors also improve the combinatorial ones of Arkin et al. Our approach can also be extended to other geometric variants, such as scenarios with obstacles and linear combinations of turn and distance costs.

1 Introduction

Finding roundtrips of minimum cost is one of the classic problems of theoretical computer science. In its most basic form, the objective of the *Traveling Salesman Problem (TSP)* is to minimize the total length of a single tour that covers all of a given set of locations. If the tour is not required to be connected, the result may be a *cycle cover*: a set of closed subtours that together cover the whole set.

A full version of this extended abstract can be found at [17].

P. Heggernes (Ed.): CIAC 2019, LNCS 11485, pp. 224–236, 2019.
https://doi.org/10.1007/978-3-030-17402-6_19

This distinction makes a tremendous difference for the computational complexity: while the TSP is NP-hard, computing a cycle cover of minimum total length can be achieved in polynomial time, based on matching techniques.

Evaluating the cost for a tour or a cycle cover by only considering its length may not always be the right measure. Figure 1 shows an example application, in which a drone has to sweep a given region to fight mosquitoes that may transmit dangerous diseases. As can be seen in the right-hand part of the figure, by far the dominant part of the overall travel cost occurs when the drone has to change its direction. (See our related video and abstract [10] for more details, and the resulting tour optimization.) There is an abundance of other related applied work, e.g., mowing lawns or moving huge wind turbines [8].

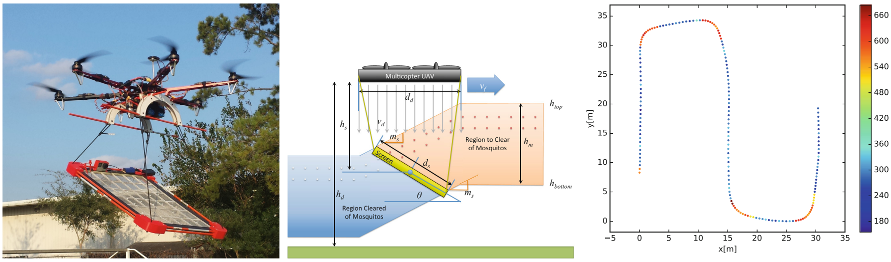

Fig. 1. (Left) A drone equipped with an electrical grid for killing mosquitoes. **(Middle)** Physical aspects of the flying drone. **(Right)** Making turns is expensive. See our related video [10].

For many purposes, two other variants are also practically important: for *subset coverage*, only a prespecified subset of locations needs to be visited, while for *penalty coverage*, locations may be skipped at the expense of an individual penalty. From the theoretical side, Arkin et al. [6] showed that finding minimum-turn tours in grid graphs is NP-hard, even if a minimum-turn cycle cover is given. The question whether a minimum-turn cycle cover can be computed in polynomial time (just like a minimum-length cycle cover) has been open for at least 17 years, dating back to the conference paper [5]; it has been listed for 15 years as *Problem 53* in *The Open Problems Project* edited by Demaine, Mitchell, and O'Rourke [15]. In Sect. 2 we resolve this problem by showing that computing a minimum-turn cycle cover in planar grid graphs is indeed NP-hard.

This raises the need for approximation algorithms. In Sect. 3, we present a technique based on Integer Programming (IP) formulations and their Linear Programming (LP) relaxations. Based on polyhedral results and combinatorial modifications, we prove constant approximation for all problem variants.

1.1 Related Work

Milling with Turn Costs. Arkin et al. [5,6] introduce the problem of milling (i.e., "carving out") with turn costs. They show hardness of finding an optimal tour, even in *thin* 2-dimensional grid graphs (which do not contain an induced 2×2 subgraph) with a given optimal cycle cover. They give a 2.5-approximation algorithm for obtaining a cycle cover, resulting in a 3.75-approximation algorithm for tours. The complexity of finding an optimal cycle cover in a 2-dimensional grid graph was established as *Problem 53* in *The Open Problems Project* [15].

Maurer [23] proves that a cycle *partition* with a minimum number of turns in grid graphs can be computed in polynomial time and performs practical experiments for optimal cycle covers. De Assis and de Souza [14] computed a provably optimal solution for an instance with 76 vertices. For the abstract version on graphs (in which "turns" correspond to weighted changes between edges), Fellows et al. [20] show that the problem is fixed-parameter tractable by the number of turns, tree-width, and maximum degree. Benbernou [11] considered milling with turn costs on the surface of polyhedrons in the 3-dimensional grid. She gives a corresponding 8/3-approximation algorithm for tours.

Note that the theoretical work presented in this paper has significant practical implications. As described in our forthcoming conference paper [18], the IP/LP-characterization presented in Sect. 3 can be modified and combined with additional algorithm engineering techniques to allow solving instances with more than 1000 pixels to provable optimality (thereby expanding the range of de Assis and de Souza [14] by a factor of 15), and computing solutions for instances with up to 300,000 pixels within a few percentage points (thereby showing that the practical performance of our approximation techniques is dramatically better than the established worst-case bounds).

For mowing problems, i.e., covering a given area with a moving object that may leave the region, Stein and Wagner [25] give a 2-approximation algorithm on the number of turns for the case of orthogonal movement. If only the traveled distance is considered, Arkin et al. [7] provide approximation algorithms for milling and mowing.

Angle and Curvature-Constrained Tours and Paths. If the instances are in the $\mathbb{R}^2$ plane and only the turning angles are measured, the problem is called the *Angular Metric Traveling Salesman Problem*. Aggarwal et al. [3] prove hardness and provide an $O(\log n)$ approximation algorithm for cycle covers and tours that works even for distance costs and higher dimensions. As shown by Aichholzer et al. [4], this problem seems to be very hard to solve optimally with integer programming. Fekete and Woeginger [19] consider the problem of connecting a point set with a tour for which the angles between the two successive edges are constrained. Finding a curvature-constrained shortest *path* with obstacles has been shown to be NP-hard by Lazard et al. [22]. Without obstacles, the problem is known as the *Dubins path* [16] that can be computed efficiently. With complexity depending on the types of obstacles, Boissonnat and Lazard [12], Agarwal et al. [1], and Agarwal and Wang [2] provide polynomial-time algorithms when possible or $1 + \epsilon$ approximation algorithms otherwise. Takei et al. [26] consider the solution of the problem from a practical perspective.

Related Combinatorial Problems. Goemans and Williamson [21] provide an approximation technique for constrained forest problems and similar problems that deal with penalties. In particular, they provide a 2-approximation algorithm for *Prize-Collecting Steiner Trees* in general symmetric graphs and the *Penalty Traveling Salesman Problem* in graphs that satisfy the triangle inequality. An introduction into approximation algorithms for prize-collecting/penalty problems, k-MST/TSP, and minimum latency problems is given by Ausiello et al. [9].

1.2 Preliminaries

The angular metric traveling salesman problem resp. cycle cover problem ask for a cycle resp. set of cycles such that a given set P of n points in $\mathbb{R}^d$ is covered and the sum of turn angles in minimized. A cycle is a closed chain of segments and covers the points of the segments' joints. A cycle has to cover at least two points. The turn angle of a joint is the angle difference to 180°. In the presence of polygonal obstacles, cycles are not allowed to cross them. We consider three coverage variants: Full, subset, and penalty. In full coverage, every point has to be covered. In subset coverage, only points in a subset $S \subseteq P$ have to be covered (which is only interesting for grid graphs). In penalty coverage, no point has to be covered but every uncovered point $p \in P$ induces a penalty $c(p) \in \mathbb{Q}_0^+$ on the objective value. Optionally, the objective function can be a linear combination of distance and turn costs.

In the following, we introduce the *discretized angular metric*, by considering for every point $p \in P$ a set of ω possible orientations (and thus, 2ω possible directions) for a trajectory through p. We model this by considering for each $p \in P$ a set O_p of ω infinitely short segments, which we call *atomic strips*; a point is covered if one of its segments is part of the cycle, see Fig. 2. The corresponding selection of atomic strips is called *Atomic Strip Cover*, i.e., a selection of one $o \in O_p$ for every $p \in P$.

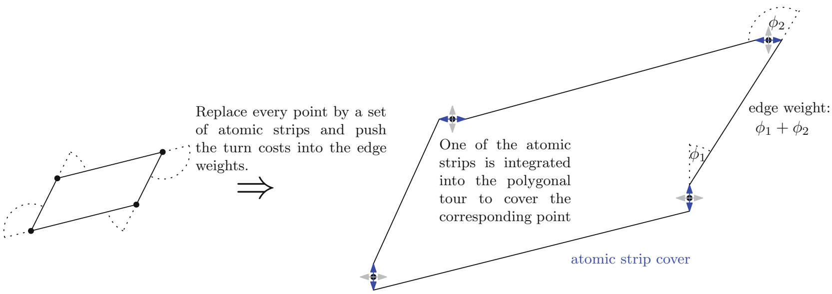

Fig. 2. Transforming an angular metric TSP instance and solution to an instance based on atomic strips, which can be considered infinitely small segments.

The atomic strips induce a weighted graph $G_O(V_O, E_O)$ with the endpoints of the atomic strips as vertices and the connections between the endpoints as edges. The weight of an edge in G_O equals the connection costs, in particular the turn costs on the two endpoints. Thus, the cycle cover problem turns into finding an Atomic Strip Cover with the minimum-weight perfect matching on its induced subgraph. As the cost of connections in it depends on *two* edges in the original graph, we call this generalized problem (in which the edge weights do not have to be induced by geometry) the *semi-quadratic cycle cover problem*.

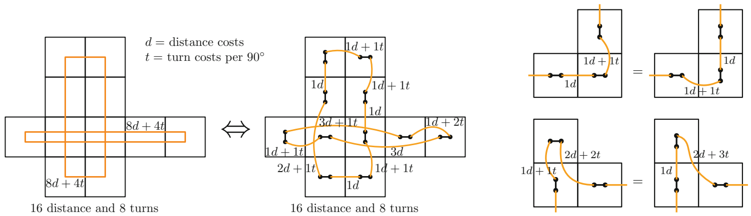

Fig. 3. **(Left)** From an optimal cycle cover (dotted) we can extract an Atomic Strip Cover (thick black), such that the matching (orange) induces an optimal solution. **(Right)** For turns it does not matter if we choose the horizontal or vertical atomic strip. (Color figure online)

It is important to note that the weights do not satisfy the triangle inequality; however, a direct connection is not more expensive than a connection that includes another atomic strip, giving rise to the following *pseudo-triangle inequalities*.

$$\forall v_1, v_2 \in V_O, w_1w_2 \in O_p, p \in P: \begin{array}{l} \text{cost}(v_1v_2) \leq \text{cost}(v_1w_1) + \text{cost}(w_2v_2) \\ \text{cost}(v_1v_2) \leq \text{cost}(v_1w_2) + \text{cost}(w_1v_2) \end{array} \quad (1)$$

Our model allows the original objective function to be a linear combination of turn and distance costs, as it does not influence Eq. (1). Instances with polygonal obstacles for 2-dimensional geometric instances are also possible (however, for 3D, the corresponding edge weights can no longer be computed efficiently). A notable special case are *grid graphs* that arise as vertex-induced subgraphs of the infinite integer orthogonal grid. In this case, a point can only be covered straight, by a simple 90° turn, or by a 180° u-turn. We show grid graphs as polyominoes in which vertices are shown as *pixels*. We also speak of the number of *simple turns* (u-turns counting as two) instead of turn angles. More general grid graphs can be based on other grids, such as 3-dimensional integral or hexagonal grids.

Minimum turn cycle covers in grid graphs can be modeled as a semi-quadratic cycle cover problem with $\omega = 2$ and edge weights satisfying Eq. (1). One of the atomic strips represents being in a horizontal orientation (with an east and a west heading vertex) and the other being in a vertical orientation (with a north and a south heading vertex). The cost of an edge is as follows; see Fig. 3: Every vertex is connected to a position and a direction. The cost is the cheapest transition from the position and direction of the first vertex to the position and opposite heading of the second vertex (this is symmetric and can be computed efficiently). We can easily transform a cycle cover in a grid graph into one based on atomic strips and vice versa; see Fig. 3 (left). For each pixel we choose one of its transitions. If it is straight, we select the equally oriented strip; otherwise it does not matter, see Fig. 3 (right). With more atomic strips we can also model more general grid graphs such as hexagonal or 3-dimensional grid graphs with three atomic strips.

1.3 Our Contribution

We provide the following results.

- We resolve *Problem 53* in *The Open Problems Project* [15] by proving that finding a cycle cover of minimum turn cost is NP-hard, even in the restricted case of grid graphs. We also prove that finding a subset cycle cover of minimum turn cost is NP-hard, even in the restricted case of *thin* grid graphs, in which no induced 2×2 subgraph exists. This differs from the case of full coverage in thin grid graphs, which is known to be polynomially solvable [6].
- We provide a general IP/LP-based technique for obtaining $2 * \omega$ approximations for the semi-quadratic (penalty) cycle cover problem if Eq. (1) is satisfied, where ω is the maximum number of atomic strips per vertex.
- We show how to connect the cycle covers to minimum turn tours to obtain a 6 approximation for full coverage in regular grid graphs, 4ω approximations for full tours in general grid graphs, $4\omega + 2$ approximations for (subset) tours, and $4\omega + 4$ for penalty tours.

To the best of our knowledge, this is the first approximation algorithm for the subset and penalty variant with turn costs. For general grid graphs our techniques yields better guarantees than than the techniques of Arkin et al. who give a factor of $6 * \omega$ for cycle covers and $6 * \omega + 2$ for tours. In practice, our approach also yields better solutions for regular grid graphs, see [18].

2 Complexity

Problem 53 in *The Open Problems Project* asks for the complexity of finding a minimum-turn (full) cycle cover in a 2-dimensional grid graph. This is by no means obvious: large parts of a solution can usually be deduced by local information and matching techniques. In fact, it was shown by Arkin et al. [5,6] that the full coverage variant in *thin* grid graphs (which do not contain a 2×2 square, so every pixel is a boundary pixel) is solvable in polynomial time. In this section, we prove that finding a *full* cycle cover in 2-dimensional grid graphs with minimum turn cost is NP-hard, resolving *Problem 53*. We also show that *subset* coverage is NP-hard even for *thin* grid graphs, so the boundary techniques by Arkin et al. [5,6] do not provide a polynomial-time algorithm.

Theorem 1. *It is NP-hard to find a cycle cover with a minimum number of* 90° *turns (*180° *turns counting as two) in a grid graph.*

The proof is based on a reduction from *One-in-three 3SAT* (1-in-3SAT), which was shown to be NP-hard by Schaefer [24]: for a Boolean formula in conjunctive normal form with only three literals per clause, decide whether there is a truth assignment that makes exactly one literal per clause `true` (and exactly two literals `false`). For example, $(x_1 \vee x_2 \vee x_3) \wedge (\overline{x_1} \vee \overline{x_2} \vee \overline{x_3})$ is not (1-in-3) satisfiable, whereas $(x_1 \vee x_2 \vee x_3) \wedge (\overline{x_1} \vee \overline{x_2} \vee \overline{x_4})$ is satisfiable.

See full version [17] for details, Fig. 5 for representing the one-clause formula $x_1+x_2+x_3 = 1$ with its three possible 1-in-3 solutions, and Fig. 4 for the instance $x_1+x_2+x_3 = 1 \wedge \overline{x_1}+\overline{x_2}+\overline{x_4} = 1 \wedge \overline{x_1}+x_2+\overline{x_3} = 1$. For every variable we have a gadget consisting of a gray gadget and a zig-zagging, high-cost path of blue pixels. A cheap solution traverses a blue path once and connect the ends through the remaining construction of gray and red pixels. Such *variable cycles* (highlighted in red) must either go through the upper () or lower () lane of the variable gadget; the former corresponds to a `true`, the later to a `false` assignment of the corresponding variable. A *clause gadget* modifies a lane of all three involved variable gadgets. This involves the gray pixels that are covered by the green cycles; we can show that they do not interfere with the cycles for covering the blue and red pixels, and cannot be modified to cover them. Thus, we only have to cover red and blue pixels, but can pass over gray pixels, too.

To this end, we must connect the ends of the blue paths; as it turns out, the formula is satisfiable if and only if we can perform this connection in a manner that also covers one corresponding red pixel with at most two extra turns.

For subset cover we can also show hardness for *thin* grid graphs. Arkin et al. [5,6] exploits the structure of these graphs to compute an optimal minimum-turn cycle cover in polynomial time. If we only have to cover a subset of the vertices, the problem becomes NP-hard again. The proof is inspired by the construction of Aggarwal et al. [3] for the angular-metric cycle cover problem and significantly simpler than the one for full coverage. See full version [17] for proof details.

Theorem 2. *The minimum-turn subset cycle cover problem is NP-hard, even in thin grid graphs.*

3 Approximation Algorithms

3.1 Cycle Cover

Now we describe a 2ω-approximation algorithm for the semi-quadratic (penalty) cycle cover problem with ω atomic strips per point if the edge weights satisfy Eq. (1). We focus on the full coverage version, as the penalty variant can be modeled in full coverage (with the same ω and while still satisfying Eq. (1)), by adding for every point $p \in P$ two further points that have a zero cost cycle only including themselves and a cycle that also includes p with the cost of the penalty.

Our approximation algorithm proceeds as follows. We first determine an atomic strip cover via linear programming. Computing an optimal atomic strip

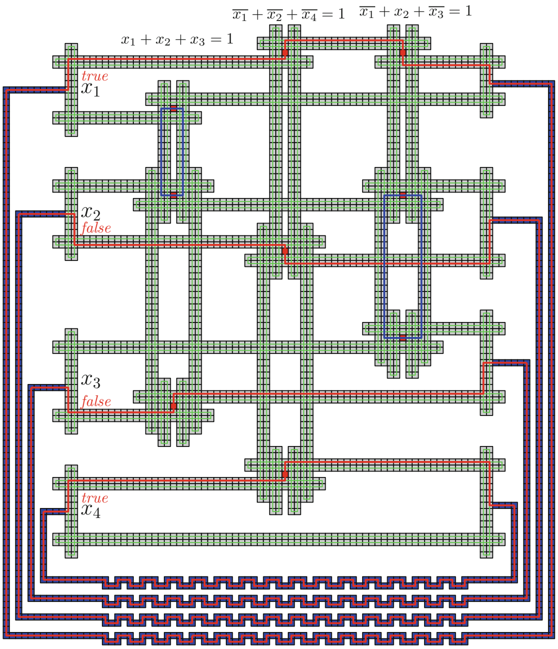

Fig. 4. Representing the *1-in-3SAT*-formula $x_1 + x_2 + x_3 = 1 \wedge \overline{x_1} + \overline{x_2} + \overline{x_4} = 1 \wedge \overline{x_1} + x_2 + \overline{x_3} = 1$. (Color figure online)

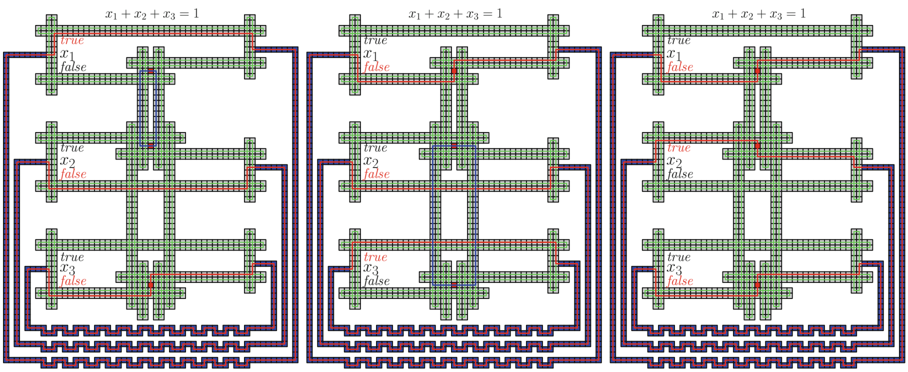

Fig. 5. Construction for the one-clause formula $x_1 + x_2 + x_3 = 1$ and three possible solutions. Every variable has a cycle traversing a zig-zagging path of blue pixels. A variable is `true` if its cycle uses the upper path ( ) through green/red pixels, `false` if it takes the lower path ( ). For covering the red pixels, we may use two additional turns. This results in three classes of optimal cycle covers, shown above. If we use the blue 4-turn cycle to cover the upper two red pixels, we are forced to cover the lower red pixel by the x_3 variable cycle, setting x_3 to `false`. The variable cycles of x_1 and x_2 take the cheapest paths, setting them to `true` or `false`, respectively. The alternative to a blue cycle is to cover all three red pixel by the variable cycles, as in the right solution. (Color figure online)

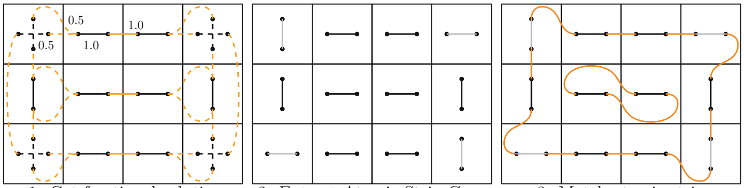

Fig. 6. Example of the approximation algorithm for a simple full cycle cover instance in a grid graph. First the fractional solution of the integer program (2)–(4) is computed. Strips and edges with value 0 are omitted, while dashed ones have value 0.5. Then the dominant (i.e., highest valued) atomic strips of this solution are selected. Finally, a minimum weight perfect matching on the ends of the atomic strips is computed. (Recall that atomic strips only have an but no length, so the curves in the corner indicate simple 90° turns.)

cover is NP-hard; we can show that choosing the *dominant* strips for each pixel in the fractional solution, i.e., those with the highest value, suffices to obtain provable good solutions. As a next step, we connect the atomic strips to a cycle cover, using a minimum-weight perfect matching. See Fig. 6 for an illustration.

We now describe the integer program whose linear programming relaxation is solved to select the dominant atomic strips. It searches for an optimal atomic strip cover that yields a perfect matching of minimum weight. To satisfy Eq. (1), transitive edges (connections implied by multiple explicitly given edges) may need to be added, especially loop-edges (which are not used in the final solution). The IP does not explicitly enforce cycles to contain at least two points: all small cycles consist only of transitive edges that implicitly contain at least one further atomic strip/point. For the usage of a matching edge $e = vw \in E_O$, we use the Boolean variable $x_e = x_{vw}$. For the usage of an atomic strip $o = vw \in O_p, p \in P$, we use the Boolean variable $y_o = y_{vw}$.

$$\min \quad \sum_{e \in E_O} \text{cost}(e) x_e \tag{2}$$

$$\text{s.t.} \quad \sum_{vw \in O_p} y_{vw} = 1 \qquad p \in P \tag{3}$$

$$2x_{vv} + \sum_{\substack{e \in E_O(v)\\ e \neq vv}} x_e = 2x_{ww} + \sum_{\substack{e \in E_O(w)\\ e \neq ww}} x_e = y_{vw} \quad p \in P, vw \in O_p \tag{4}$$

We minimize the cost of the used edges, with Eq. (3) forcing the selection of one atomic strip per pixel (atomic strip cover) and Eq. (4) ensuring that exactly the vertices (endpoints) of the selected atomic strips are matched, with loop edges counting double due to their two ends.

Theorem 3. *Assuming edge weights that satisfy Eq. (1), there is a 2ω-approximation for semi-quadratic (penalty) cycle cover.*

Proof. Consider the described fractional atomic strip cover and matching of the integer program, which is a lower bound on the optimal cycle cover. We now show that we can transform this solution to a matching of the dominant strips with at most 2ω times the value. First we modify the solution such that exactly the dominant strips are used. In the current solution, the dominant strips are already used with at least $\frac{1}{\omega}$, so multiplying the solution by ω ensures a full usage of them. Now we can remove all superfluous strip usages by replacing two fractional matching edges that go through such a strip by a directly connecting matching edge without increasing the cost. This can create loop matching edges (assume these to have the same cost as the two edges they replace); these can easily be removed later. After this, we are left with a matching polytope that is half-integral (based on the same proof as for Theorem 6.13 in the book of Cook et al. [13]). Thus, we can assume our matching to be half-integral and double it to obtain an integral solution with double usages of strips. These double usages can be removed the same way as before while remaining integral. Whole redundant cycles may be removed on this way. We are now left with a feasible matching of the dominant strips that has at most 2ω times the cost of the original fractional solution, giving us the desired upper bound. More details on this proof can be found in the full version [17].

3.2 Tours

A given cycle cover approximation can be turned into a tour approximation at the expense of an additional constant factor. Because every cycle involves at least two points and a full rotation, we can use classic tree techniques known for TSP variants to connect the cycles and charge the necessary turns to the involved cycles. We sketch the basic ideas; see full version [17] for details.

Theorem 4. *Assuming validity of Eq. 1 we can establish the following approximation factors for tours.*

(i) Full tours in regular grid graphs: 6*-approximation.*
(ii) Full tours in generalized grid graphs: 4ω*-approximation.*
(iii) Subset tours in (generalized) grid graphs: $(4\omega + 2)$*-approximation.*
(iv) Geometric full tours: $(4\omega + 2)$*-approximation.*
(v) Penalty tours (in grid graphs and geometric): $(4\omega + 4)$*-approximation.*

These results also hold for objective functions that are linear combinations of length and turn costs.

Proof. It is crucial that (1) a cycle always has a turn cost of at least 360°, (2) two intersecting cycles can be merged with a cost of at most 360°, and (3) two cycles intersecting on a 180° turn can be merged without additional cost.

(i) For full tours in grid graphs, greedily connecting cycles provides a tour with at most 1.5 times the turn cost of the cycle cover, while a local optimization can be exploited to limit the length to 4 times the optimum, as shown by Arkin et al. [5].

(ii) In a cycle cover for (generalized) grid graphs, there are always at least two cycles with a distance of one, while every cycle has a length of at least 2; otherwise the cycle cover is already a tour. This allows iteratively merging cycles at cost at most as much as a cheapest cycle; the total number of merges is less than the number of cycles.

(iii) and (iv) For subset coverage in grid graphs or full coverage in the geometric case, we need to compute the cheapest paths between any two cycles, ignoring the orientations at the ends. First connect all intersecting cycles, charging the cost on the vanishing cycles. The minimum spanning tree on these edges is a lower bound on the cost of the tour. Doubling the MST connects all cycles with the cost of twice the MST, the cost of the cycle cover, and the turn costs at the end of the MST edges, which can be charged to the cycles.

(v) Penalty tours can be approximated in a similar manner. Instead of an MST, we use a Price-Collecting Steiner Tree, which is a lower bound on an optimal penalty tour. We use a 2-approximation for the PCST [21], as it is NP-hard. We achieve a cost of twice the 2-approximation of the PCST, the cost of the penalty cycle cover, and the cost of its cycles again for charging the connection costs. The penalties of the points not in the cycle cover are already paid by the penalty cycle cover.

□

4 Conclusions

We have presented a number of theoretical results on finding optimal tours and cycle covers with turn costs. In addition to resolving the long-standing open problem of complexity, we provided a generic framework to solve geometric (penalty) cycle cover and tours problems with turn costs.

As described in [10], the underlying problem is also of practical relevance. As it turns out, our approach does not only yield polynomial-time approximation algorithms; enhanced by an array of algorithm engineering techniques, they can be employed for actually computing optimal and near-optimal solutions for instances of considerable size in grid graphs. Further details on these algorithm engineering aspects will be provided in our forthcoming paper [18].

References

1. Agarwal, P.K., Biedl, T.C., Lazard, S., Robbins, S., Suri, S., Whitesides, S.: Curvature-constrained shortest paths in a convex polygon. SIAM J. Comp. **31**(6), 1814–1851 (2002)
2. Agarwal, P.K., Wang, H.: Approximation algorithms for curvature-constrained shortest paths. SIAM J. Comp. **30**(6), 1739–1772 (2000)
3. Aggarwal, A., Coppersmith, D., Khanna, S., Motwani, R., Schieber, B.: The angular-metric traveling salesman problem. SIAM J. Comp. **29**(3), 697–711 (1999)

4. Aichholzer, O., Fischer, A., Fischer, F., Meier, J.F., Pferschy, U., Pilz, A., Staněk, R.: Minimization and maximization versions of the quadratic travelling salesman problem. Optimization **66**(4), 521–546 (2017)
5. Arkin, E.M., Bender, M.A., Demaine, E.D., Fekete, S.P., Mitchell, J.S.B., Sethia, S.: Optimal covering tours with turn costs. In: Proceedings of 12th ACM-SIAM Symposium Discrete Algorithms (SODA), pp. 138–147 (2001)
6. Arkin, E.M., Bender, M.A., Demaine, E.D., Fekete, S.P., Mitchell, J.S.B., Sethia, S.: Optimal covering tours with turn costs. SIAM J. Comp. **35**(3), 531–566 (2005)
7. Arkin, E.M., Fekete, S.P., Mitchell, J.S.B.: Approximation algorithms for lawn mowing and milling. Comp. Geom. **17**(1–2), 25–50 (2000)
8. Astroza, S., Patil, P.N., Smith, K.I., Bhat, C.R.: Transportation planning to accommodate needs of wind energy projects. In: Transportation Research Board-Annual Meeting (2017). Article 17–05309
9. Ausiello, G., Bonifaci, V., Leonardi, S., Marchetti-Spaccamela, A.: Prize-collecting traveling salesman and related problems. In: Gonzalez, T.F. (ed.) Handbook of Approximation Algorithms and Metaheuristics. Chapman and Hall/CRC, Boca Raton (2007)
10. Becker, A.T., Debboun, M., Fekete, S.P., Krupke, D., Nguyen, A.: Zapping Zika with a mosquito-managing drone: Computing optimal flight patterns with minimum turn cost. In: Proceedings of 33rd Symposium on Computational Geometry (SoCG), pp. 62:1–62:5 (2017). Video at https://www.youtube.com/watch?v=SFyOMDgdNao
11. Benbernou, N.M.: Geometric algorithms for reconfigurable structures. Ph.D. thesis, Massachusetts Institute of Technology (2011)
12. Boissonnat, J., Lazard, S.: A polynomial-time algorithm for computing a shortest path of bounded curvature amidst moderate obstacles. In: Proceedings of 12th Symposium on Computational Geometry (SoCG), pp. 242–251 (1996)
13. Cook, W., Cunningham, W., Pulleyblank, W., Schrijver, A.: Combinatorial Optimization. Wiley-Interscience, Hoboken (1997)
14. de Assis, I.R., de Souza, C.C.: Experimental evaluation of algorithms for the orthogonal milling problem with turn costs. In: Pardalos, P.M., Rebennack, S. (eds.) SEA 2011. LNCS, vol. 6630, pp. 304–314. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20662-7_26
15. Demaine, E.D., Mitchell, J. S. B., O'Rourke, J.: The Open Problems Project. http://cs.smith.edu/~orourke/TOPP/
16. Dubins, L.E.: On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. Am. J. Math. **79**(3), 497–516 (1957)
17. Fekete, S.P., Krupke, D.: Covering tours and cycle covers with turn costs: hardness and approximation. arXiv (2018). http://arxiv.org/abs/1808.04417
18. Fekete, S.P., Krupke, D.: Practical methods for computing large covering tours and cycle covers with turn cost. In: Proceedings of 21st SIAM Workshop on Algorithm Engineering and Experiments (ALENEX) (2019)
19. Fekete, S.P., Woeginger, G.J.: Angle-restricted tours in the plane. Comp. Geom. **8**, 195–218 (1997)
20. Fellows, M., Giannopoulos, P., Knauer, C., Paul, C., Rosamond, F.A., Whitesides, S., Yu, N.: Milling a graph with turn costs: a parameterized complexity perspective. In: Thilikos, D.M. (ed.) WG 2010. LNCS, vol. 6410, pp. 123–134. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16926-7_13
21. Goemans, M.X., Williamson, D.P.: A general approximation technique for constrained forest problems. SIAM J. Comp. **24**(2), 296–317 (1995)

22. Lazard, S., Reif, J., Wang, H.: The complexity of the two dimensional curvature-constrained shortest-path problem. In: Proceedings of 3rd Workshop on the Algorithmic Foundations of Robotics (WAFR), pp. 49–57 (1998)
23. Maurer, O.: Winkelminimierung bei Überdeckungsproblemen in Graphen. Diplomarbeit, Technische Universität Berlin (2009)
24. Schaefer, T.J.: The complexity of satisfiability problems. In: Proceedings of 10th ACM symposium on Theory of computing (STOC), pp. 216–226 (1978)
25. Stein, C., Wagner, D.P.: Approximation algorithms for the minimum bends traveling salesman problem. In: Aardal, K., Gerards, B. (eds.) IPCO 2001. LNCS, vol. 2081, pp. 406–421. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45535-3_32
26. Takei, R., Tsai, R., Shen, H., Landa, Y.: A practical path-planning algorithm for a simple car: a Hamilton-Jacobi approach. In: Proceedings of 29th American Control Conference (ACC), pp. 6175–6180 (2010)