

Milling a Graph with Turn Costs: A Parameterized Complexity Perspective*

Mike Fellows^{1,**}, Panos Giannopoulos^{2,***}, Christian Knauer³,
Christophe Paul^{4,†}, Frances Rosamond^{1,‡}, Sue Whitesides^{5,§}, and Nathan Yu⁶

¹ PCRU, Office of DVC(Research), University of Newcastle, Australia

Michael.Fellows@newcastle.edu.au, mathgypsie@yahoo.com

² Institut für Informatik, Freie Universität Berlin, Berlin, Germany

panos@mi.fu-berlin.de

³ Institut für Informatik, Universität Bayreuth, Bayreuth, Germany

christian.knauer@uni-bayreuth.de

⁴ NRS - LIRMM, Montpellier, France

Christophe.Paul@lirmm.fr

⁵ Department of Computer Science, University of Victoria, Canada

sue@uvic.ca

⁶ nuo.nathan.yu@gmail.com

Abstract. The DISCRETE MILLING problem is a natural and quite general graph-theoretic model for geometric milling problems: Given a graph, one asks for a walk that covers all its vertices with a minimum number of *turns*, as specified in the graph model by a 0/1 turncost function f_x at each vertex x giving, for each ordered pair of edges (e, f) incident at x , the *turn cost* at x of a walk that enters the vertex on edge e and departs on edge f . We describe an initial study of the parameterized complexity of the problem.

1 Introduction

We study the parameterized complexity of the following problem:

DISCRETE MILLING

Instance: A simple graph $G = (V, E)$ and for each vertex x , a *turncost function* f_x indicating whether a *turn* is required, with $f_x : E(x) \times E(x) \rightarrow \{0, 1\}$, where $E(x)$ is the set of edges incident on x .

* Research initiated at the 6th McGill - INRIA Barbados Workshop on Computational Geometry in Computer Graphics, 2007.

** Research supported by the Australian Research Council through the ARC Centre of Excellence in Bioinformatics and Discovery Project DP0773331.

*** Research supported by the German Science Foundation (DFG) under grant Kn 591/3-1.

† Research conducted while the author was on sabbatical at McGill University, School of Computer Science, Canada. Supported by the project ANR-06-BLAN-0148.

‡ Research supported by the Australian Research Council.

§ Research supported by NSERC and FQRNT.

Question: Is there a walk making at most k turns that visits every vertex of G ? The GRID MILLING problem restricts the input to *grid graphs*: rectilinearly plane-embedded graphs that are subgraphs of the integral grid, with the natural turncost function.

Related Work. DISCRETE MILLING was introduced by Arkin *et al.* [1] as a graph model for studying geometric milling problems with turn costs and other constraints. Such problems are common in manufacturing applications such as numerically controlled machining and automatic tool path generation; see Held [6] for a survey. In DISCRETE MILLING, a solution path must visit a set of vertices that are connected by edges representing the different directions (“channels”) that the “cutter” can take. Arkin *et al.* studied a restricted version of the problem where incident edges to a vertex x are *paired* in the cost function f_x in the sense that for each incident edge e there is at most one incident edge f such that $f_x(e, f) = 0$, and symmetric: if $f_x(e, f) = 0$ then $f_x(f, e) = 0$. Here, we consider also a more general version that allows an arbitrary 0/1 turncost function at each vertex.

Arkin *et al.* [1] showed that DISCRETE MILLING is NP-hard (even for grid graphs) and described a constant-factor approximation algorithm for minimizing the number of turns in a solution walk. They also described a PTAS for the case where the cost is a linear combination of the length of the walk and the number of turns. No PTAS is known for the case of turn costs only.

Results. We start by showing that GRID MILLING is fixed-parameter tractable when parameterized by the numbers of turns. For this, two approaches are presented: one that is based on monadic second-order logic of graphs of bounded treewidth, and another, more practical one, based on dynamic programming on branch decompositions. Generalizing the former approach, we give an FPT result for DISCRETE MILLING, parameterized by (k, t, d) , where k is the number of turns, t is the tree-width of the input graph G , and d is the maximum degree of G . We then explore whether this positive result can be further strengthened. However: DISCRETE MILLING, even in its restricted version, is $W[1]$ -hard when parameterized by (k, p) , where k is the number of turns and p is the path-width of G (and therefore also when parameterized by (k, t)). Our negative result provides one of the few problems known to be $W[1]$ -hard when parameterized by pathwidth.

Definitions and Preliminaries. We will assume that the basic ideas of parameterized complexity theory and bounded tree-width algorithmics up through the basic form of Courcelle’s Theorem and monadic second-order logic (MSO) are known to the reader. For background on these topics, see [3,5,7]. Details of routine deployments of MSO in the proofs of our theorems (that can be laborious in full formality) are relegated to the full version of the paper due to space limitations.

For a graph G , let $\text{tw}(G)$ be its treewidth. We assume that all graphs G are simple (no loops or multiple edges). A *walk* $W = [x_0, \dots, x_l]$ on a graph $G = (V, E)$ is a sequence of vertices such that every pair x_i, x_{i+1} of consecutive

vertices of the sequence are adjacent (we use $x_i x_{i+1}$ to refer to the edge between them). The *turn cost* of a walk W is defined as

$$\text{tc}(W) = \sum_{i=1}^{l-1} f_{x_i}(x_{i-1}x_i, x_i x_{i+1}).$$

A walk that visits every vertex of a graph is termed a *covering walk*. Note that in DISCRETE MILLING a solution covering walk may visit a vertex *many times*.

2 Grid Milling is Fixed-Parameter Tractable

We prove here that GRID MILLING is FPT for parameter k , the number of turns. We first argue that instances with large tree-width are no-instances.

Lemma 1. *Let $G = (V, E)$ be a connected grid graph with $\text{tw}(G) > 6k - 5$. Then G does not contain a $(k - 2)$ -turn covering walk.*

Proof. We show that G contains k vertices that have pairwise different x - and y -coordinates. Then, any covering walk needs to take at least one turn between any two such vertices, and thus it needs at least $k - 1$ turns in total.

Since G is planar and $\text{tw}(G) > 6k - 5$, by the Excluded Grid Theorem for planar graphs (c.f. [5]), it has a $(k \times k)$ -grid H as a minor. H contains $k/2$ vertex-disjoint consecutively nested cycles. Since taking minors can destroy or merge cycles but not create completely new ones, in the “pre-images” (under the operation taking minor) of these cycles there must be $k/2$ vertex-disjoint subgraphs in G , each containing a cycle. Thus, G contains a set \mathcal{C} of $k/2$ nested vertex-disjoint cycles. Consider a straight line L of unit slope that intersects the innermost cycle of \mathcal{C} at two vertices (grid points). L must also intersect every other cycle at at least two vertices. This produces a set of at least k vertices in G with the claimed property.

2.1 FPT via MSO

We give an MSO-based approach first, which also serves as a good starting point for our result on the general DISCRETE MILLING problem.

We associate to the grid graph G an annotated graph $\mathcal{M}(G)$: we simply regard the horizontal edges as being of one type, and the vertical edges as being of a second type. Equivalently, we can think of G as presented to us with a partition of the edge set: $G = (V, E_h, E_v)$. The idea is to show that the property of being a yes-instance of the problem can be expressed as an MSO property of $\mathcal{M}(G)$.

Intuitively, G has a k -covering walk if and only if there exist a start vertex v_0 , turn vertices v_1, \dots, v_k , an end vertex v_{k+1} , and sets of vertices S_0, \dots, S_k , such that:

- (i) the graph induced by S_i , $i = 0, \dots, k$, is a monochromatic path, i.e. a path whose edges are all either in E_h or in E_v ,

- (ii) the path induced by S_i starts at v_i and ends at v_{i+1} , and
- (iii) $V = \cup S_i$, i.e. all vertices of G are covered.

This can be straightforwardly formalized in MSO.

Lemma 2. *Let $G = (V, E_h, E_v)$ be a grid graph. The property of having a k -covering walk on G is expressible in MSO.*

Easily, $\mathcal{M}(G)$ has treewidth bounded as a function of $\text{tw}(G)$, and from Lemmata 1, 2 and Courcelle’s Theorem we get:

Theorem 1. *GRID MILLING is FPT with respect to k (number of turns).*

2.2 Dynamic Programming on a Sphere Cut Decomposition

A *branch decomposition* (T, μ) of a graph G is an unrooted ternary tree T together with a bijection μ between the leaves of T and the edge set of G . For an edge e of T , let T_1, T_2 be the two subtrees obtained by removing e , and let G_1, G_2 be the subgraphs of G induced by the edges of the leaves of T_1 and T_2 respectively. The *middle set* of e is defined as $\text{mid}(e) = V(G_1) \cap V(G_2)$. The *width* of a branch decomposition is the maximum size of a middle set over all edges. The *branchwidth* $\text{bw}(G)$ of G is the minimum width over all possible branch decompositions.

A *noose* of a plane graph G is a simple closed curve on the plane that intersects G only at vertices and every face at most once. The noose separates the plane into two regions, which have the noose as a common boundary, and G into two subgraphs, each lying inside one of the regions; the subgraphs meet only at vertices on the noose. A *sphere cut (sc-) decomposition* (T, μ) is a branch decomposition with the property that for every edge e of T there is a noose of G such that its two corresponding subgraphs are the subgraphs G_1, G_2 associated with e . Note that the noose intersects G in $\text{mid}(e)$. An sc-decomposition of a plane graph G (with no degree-1 vertices) of width $\text{bw}(G)$ can be constructed in $O(n^3)$ time [2].

Let G be a grid graph (with no degree-1 vertices) and (T, μ) a sphere cut decomposition of G with minimum width. Since $\text{tw}(G) = \Theta(\text{bw}(G))$, from Lemma 1 we can assume that $\text{bw}(G) = O(k)$. We root T as explained in [2]. For a node v of T , other than the root, let O_v be the noose corresponding to the edge between v and its parent. Let G_v be the subgraph of G that is associated with this edge and induced by the leaves of the subtree rooted at v , and let Δ_v be the region where G_v lies into. A covering walk of G induces a sequence of paths in Δ_v . Their union covers all vertices (of G_v) in the interior of Δ_v but not necessarily all its boundary vertices, i.e., the vertices on O_v . Note that any vertex or edge might be used more than once.

For a node v of T we define a table S_v of subproblems as follows: Let $C_v \subseteq O_v \cap V(G_v)$ be a set of boundary vertices and $D_v \subseteq E(G_v)$ be the set of edges with at least one endpoint in C_v . Also, let $Q_v = (e_1, e'_1, \dots, e_l, e'_l)$ be a sequence of edges with $e_i \in D_v$ for $i = 1, \dots, l$ and some $l \in \mathcal{N}^*$. We want to compute $S_v(C_v, Q_v)$, which is a set $\mathcal{P} = \{P_1, \dots, P_{|Q_v|}\}$ of paths satisfying the following:

- (i) The first (last) edge of P_i is e_i (e'_i),
- (ii) P_i is contained in G_v ,
- (iii) the union of the paths in \mathcal{P} covers all vertices in C_v along with all the interior vertices of G_v , and
- (iv) the total number of turns of all paths in \mathcal{P} is minimum.

We compute $S_v(C_v, Q_v)$ for all possible combinations of C_v and Q_v .

The dynamic program proceeds from the leaves to the root of T . Let x, y be the children of v . Observe that G_v is the union of the two subgraphs G_x, G_y drawn in the regions Δ_x, Δ_y (corresponding to the nooses O_x and O_y). Thus, a path in $S_v(C_v, Q_v)$ consists of a sequence of path segments in Δ_x and Δ_y . We compute $S_v(C_v, Q_v)$ by enumerating all possible combinations of $S_x(C_x, Q_x)$ and $S_y(C_y, Q_y)$ that form paths in Q_v . Note that when joining two paths an increase of the total turn cost by one might be needed.

Since $\text{bw}(G) = O(k)$, there are $O(k)$ vertices on every noose and $2^{O(k)}$ possible subsets C_v . Since the degree of every vertex of G is at most 4, we have that $|D_v| = O(k)$. The crucial observation for bounding the size of Q_v is the following: any covering walk that uses an edge more than $k + 1$ times makes at least $k + 1$ turns. Thus, any edge of D_v cannot appear in Q_v more than $k + 1$ times, and so $|Q_v| = O(k^2)$. There are $O(2^k \cdot k^{k^2}) = O(2^{O(k^2 \log k)})$ possible sequences Q_v , thus each S_v has $O(2^k \cdot 2^{O(k^2 \log k)}) = O(2^{O(k^2 \log k)})$ entries, and since T has $O(n)$ nodes, the total time of the algorithm is $O(2^{O(k^2 \log k)})n$.

Theorem 2. GRID MILLING can be solved in $O(2^{O(k^2 \log k)}n + n^3)$ time.

3 Extending Tractability

What makes the GRID MILLING problem FPT? A few properties of grid graphs might lead us to tractable generalizations: (i) yes-instances must have bounded treewidth, (ii) vertices in grid graphs have bounded degree, and (iii) the turn-cost function is pairing and symmetric.

We are naturally led to three questions, by relaxing these conditions:

- What is the complexity of DISCRETE MILLING parameterized by (k, t, d) , where k is the number of turns, t is a treewidth bound, and d is a bounded on maximum degree?
- What is the complexity of DISCRETE MILLING parameterized by (k, t) ?
- What is the complexity of DISCRETE MILLING parameterized by (k, d) ?

In the remainder of this paper, we answer the first two. The third question remains open.

Theorem 3. DISCRETE MILLING is FPT for parameter (k, t, d) , where k is the number of turns, t the tree-width of the graph G and d is the maximum degree of G .

Proof. We describe how an instance of the DISCRETE MILLING problem, consisting of G and the turncost functions, can be represented by an annotated

digraph $\mathcal{M}(G)$, that allows us to use MSO logic to express a property that corresponds to the question that the DISCRETE MILLING problem asks. The proof therefore consists of three parts: (1) a description of $\mathcal{M}(G)$, (2) necessary and sufficient criteria regarding $\mathcal{M}(G)$, for the instance of DISCRETE MILLING to be a yes-instance, and (3) the expression of these criteria in MSO logic.

Let $G = (V, E)$ be the graph of the DISCRETE MILLING instance. The vertex set of the digraph $\mathcal{M}(G)$ is $\mathcal{V} = \mathcal{V}_1 \cup \mathcal{V}_2$ where

$$\mathcal{V}_1 = \{l[v] : v \in V\} \quad \text{and} \quad \mathcal{V}_2 = \{t[e] : e \in E\} \cup \{t'[e] : e \in E\}$$

Intuitively (see Figure 1), we “keep a copy” of the vertex set V of G , mnemonically “ $l[v]$ ” for v , as a vertex location we might be during a solution walk in G . Each edge e of G is replaced by two vertices $t[e]$ and $t'[e]$ that represent a “state” in a solution walk: traversing e in one direction or the other. In order to distinguish the directions, consider that the vertex set V of G is linearly ordered. Let $e = uv \in E$ with $u < v$ in the ordering. Our convention will be that $t[e]$ represents a traversal of e from u to v , and that $t'[e]$ represents a traversal of e in the direction from v to u . Thus each edge e of G is represented by two vertices in $\mathcal{M}(G)$.

In describing arcs of the digraph model $\mathcal{M}(G)$ we will use the notation $x \cdot y$ to denote an arc from x to y . The arc set of the digraph $\mathcal{M}(G)$ is

$$\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{A}_3 \cup \mathcal{A}_4 \cup \mathcal{A}_5$$

where

$$\begin{aligned} \mathcal{A}_1 &= \{t[e] \cdot l[v] : e = \{u, v\} \in E \text{ with } u < v\} \\ \mathcal{A}_2 &= \{l[u] \cdot t[e] : e = \{u, v\} \in E \text{ with } u < v\} \\ \mathcal{A}_3 &= \{t'[e] \cdot l[u] : e = \{u, v\} \in E \text{ with } u < v\} \\ \mathcal{A}_4 &= \{l[v] \cdot t'[e] : e = \{u, v\} \in E \text{ with } u < v\} \end{aligned}$$

Let \mathcal{A}' denote the union of these four sets of arcs. Intuitively, the arcs of \mathcal{A}' just “attach” the vertices of the digraph that represent edges in G to the vertices of the digraph that represent the endpoints of the edge, so that the orientations of the arcs are compatible with the interpretation of a vertex of \mathcal{V}_2 as representing, say, a traversal of the edge uv in the direction from u to v ; the vertex therefore has an arc *to it* from $l[u]$ and an arc *from it* to $l[v]$. An inspection of Figure 1 will help to clarify.

The arc set \mathcal{A}_5 is more complicated to write down formally. Its mission is to record the possibilities for cost-free passages through vertices of a solution walk in G . Suppose a is an arc in \mathcal{A}' . Then a is *to* or *from* either a $t[e]$ vertex, for some e , or a $t'[e]$ vertex for some e . Let $\epsilon(a)$ be defined to be this edge e of G . This is well-defined. We can then define

$$\begin{aligned} \mathcal{A}_5 &= \{x \cdot y : x, y \in \mathcal{V}_2, \exists z = l[v] \in \mathcal{V}_1 \text{ and } \exists a, b \in \mathcal{A}' \\ &\quad \text{with } a = x \cdot z \text{ and } b = z \cdot y \text{ and } f_v(\epsilon(a), \epsilon(b)) = 0\} \end{aligned}$$

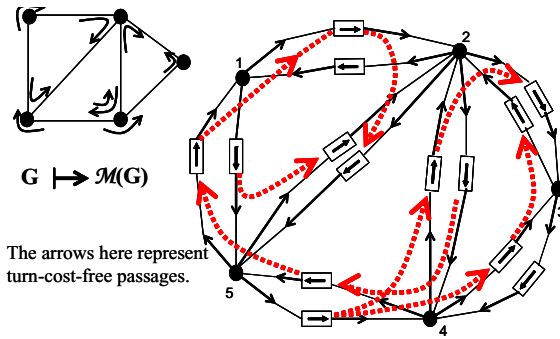


Fig. 1. The arrows drawn near the G vertices represent the turncost functions, indicating the zero-cost possibilities. These become arcs in the digraph $\mathcal{M}(G)$.

We regard $\mathcal{M}(G)$ as an annotated digraph, in the sense that there are two kinds of vertices, those of \sqsubseteq_1 and those of \sqsubseteq_2 , and two kinds of arcs, those of \mathcal{A}' and those of \mathcal{A}_5 .

The rest of the proof will show that that the question of whether G admits a covering walk making at most k turns is represented by a property of the annotated digraph $\mathcal{M}(G)$ that can be expressed in MSO logic. However, before proceeding to that, it is important to verify that if the treewidth of G is bounded by t , then the treewidth of $\mathcal{M}(G)$ is bounded by a function of the parameter. This depends crucially on the fact that the maximum degree of G is part of our compound parameterization.

Suppose $\mathcal{T}(G)$ is a tree-decomposition of G of width at most t . We can describe a bounded width tree-decomposition \mathcal{T}' of $\mathcal{M}(G)$ as follows. Without confusion, henceforth in this argument consider $\mathcal{M}(G)$ as an undirected graph by forgetting all arc orientations. Use the same bag-indexing tree for \mathcal{T}' as for $\mathcal{T}(G)$. Suppose $B \subseteq V$ is a bag of $\mathcal{T}(G)$. Replace B with the union of the closed neighborhoods of the vertices of \mathcal{V}_1 corresponding to the vertices of B , in $\mathcal{M}(G)$. It is easy to check that all the axioms for a tree-decomposition hold, and that the treewidth of $\mathcal{M}(G)$ is therefore bounded by $2dt$.

In a digraph $D = (V, A)$, by a *purposeful set of arcs* (S, s, t) we refer to a set of arcs $S \subseteq A$, together with two distinguished vertices $s, t \in V$. We say that a purposeful set of arcs (S, s, t) is *walkable* if there is a directed walk W in D from s to t such that the set of arcs traversed by W (possibly repeatedly) is S .

Now consider how the information about G and its turncost functions is represented in $\mathcal{M}(G)$. A k -turn covering walk W in G that starts at a vertex s and ends at a vertex t is described by the information:

- (1) a sequence of $k + 2$ vertices: $s = x_0, x_1, \dots, x_{k+1} = t$, and
- (2) a sequence of $k+1$ subwalks W_0, \dots, W_k where for $i = 0, \dots, k$, W_i is a turncost-free walk from x_i to x_{i+1} , that has the property that every vertex of G is visited on at least one of the subwalks.

Let $\mathcal{D}(G)$ be the subdigraph of $\mathcal{M}(G)$ induced by the vertices of \mathcal{V}_2 . A turncost-free walk in G corresponds to a directed walk in $\mathcal{D}(G)$, and vice versa, by the definition of \mathcal{A}_5 .

Claim 1. G admits a k -turn covering walk if and only if:

- (1) there are $k + 2$ vertices x_0, \dots, x_{k+1} of \mathcal{V}_1 in $\mathcal{M}(G)$, and
- (2) $k + 1$ purposeful sets of arcs (S_i, s_i, t_i) in $\mathcal{D}(G)$, $0 \leq i \leq k$, such that
 - (i) (S_i, s_i, t_i) is walkable in $\mathcal{D}(G)$ for $i = 0, \dots, k$,
 - (ii) there is an arc in \mathcal{A}' from t_i to $x_{i+1} \in \mathcal{V}_1 = V$, and from x_i to s_i , for $i = 0, \dots, k$,
 - (iii) for every vertex $x \in \mathcal{V}_1 = V$, there is some index j , $0 \leq j \leq k$, and an arc $a = u \cdot v \in S_j$, such that there is an arc in \mathcal{A}' in either direction, between x and u or v .

In one direction, the claim is easy: given a k -turn covering walk W in G , it is naturally factored into $k + 1$ turncost-free subwalks W_i , and each traversal of an edge of G in a subwalk W_i corresponds in $\mathcal{M}(G)$ to a visit to a vertex of \mathcal{V}_2 , thus the sequence of edge transversals of W_i in G corresponds 1:1 with a sequence (y_0, \dots, y_{m+1}) of vertices of \mathcal{V}_2 in $\mathcal{M}(G)$. Because W_i is turncost-free, by the definition of \mathcal{A}_5 , there is an arc in $\mathcal{D}(G)$ from y_i to y_{i+1} for $i = 0, \dots, m$. We take the set of arcs to be S_i , $s_i = y_0$ and $t_i = y_{m+1}$, giving us (1) and (2) in a well-defined manner. It is straightforward to check that the conditions hold. For example, the assumption that W is a covering walk in G yields the last condition.

Conversely, suppose we have (1) and (2) in $\mathcal{M}(G)$. By the second condition, each S_i is walkable. By the definition of \mathcal{A}_5 , a directed walk for S_i in $\mathcal{D}(G)$ corresponds to a turncost-free walk W_i in G . The third condition insures that the subwalks W_i in G can be sequenced into a k -turn walk W , where the turns occur at the vertices x_i by the first condition. W is covering in G by the fourth condition, yielding Claim 1.

Claim 2. Consider a digraph $D = (V, A)$ equipped with distinguished vertices s and t (allowing $s = t$). The property: “there exists a directed walk from s to t that traverses (allowing repetition) every arc in A ” (that is, (A, s, t) is walkable) is expressible in MSO logic.

We first argue that (A, s, t) is walkable if and only if there is a directed path P in D from s to t , such that every arc $a \in A$ either is an arc of P , or belongs to a strongly connected subdigraph D_a that includes a vertex of P . We then argue (in Appendix A) that this property is expressible in MSO logic in a straightforward manner.

Given such a directed path $P = (s = x_0, \dots, x_m = t)$ in D , we can describe a walk W that traverses every arc of A as follows. By the *arcs of P* we refer to the set of arcs

$$A[P] = \{x_0x_1, x_1x_2, \dots, x_{m-1}x_m\}$$

The walk has m phases, one for each vertex x_i of the path P . Partition the arcs of $A - A[P]$ into m classes A_0, \dots, A_m where for $i = 1, \dots, m$ every arc $a = uv \in A_i$

belongs to a strongly connected subdigraph D_a that includes the vertex x_i . Such a partition exists, by the supposed property of P . There is a directed path in D_a from x_i to u , and from v to x_i , by the strong connectivity of D_a , and so there is a directed cycle in D_a that includes both a and x_i . Include this cycle in W , starting from x_i and returning to x_i , for each arc $a \in A_i$. Increment i , take the arc from x_i to x_{i+1} and repeat this for $i = 0, \dots, m$.

Now suppose that there is a directed walk W in D from s to t that traverses every arc in A . If there is a vertex v that is visited more than once, then we can find a shorter walk W' that, considered as a sequence of arc transversals, is a subsequence of the sequence of arc transversals of W . Therefore, by downward induction, there is a directed path P from s to t , with no repeated internal vertex visits, that considered as a sequence of arc transversals, is a subsequence of the sequence of arc transversals of W . But then, every arc a traversed in the walk W (that is, every arc $a \in A$), that is not an arc of P , must belong to a subwalk W' of W that begins and ends at a vertex of P . The vertices visited by W' therefore induce a strongly connected subdigraph containing a vertex of P .

The second part of the proof of Claim 2 is to argue the property we have identified is expressible in MSO logic. The first subtask is to describe an MSO predicate that expresses that there is a directed path P in D from s to t , quantified on the sets of vertices and arcs that form the path:

$$dipath(s, t) = \exists U(\subseteq V)\exists B(\subseteq A) : \dots,$$

where the remainder of the predicate expresses that in the subdigraph $D' = (U, B)$:

- s has outdegree 1 and indegree 0
- t has indegree 1 and outdegree 0
- every vertex of U not s or t has indegree 1 and outdegree 1
- for every partition of U into U_1 and U_2 such that $s \in U_1$ and $t \in U_2$, there is a vertex $u \in U_1$ and a vertex $v \in U_2$ with an arc in B from u to v .

Being able to express that there is a directed path from s to t leads easily to an MSO predicate for strong connectivity of a subdigraph described by a set of vertices and a set of arcs. An MSO predicate for *walkability* of a set of arcs A relative to s and t is easily (but somewhat tediously) constructed on the basis of the structural characterization of Claim 2, using the predicates for the existence of an s - t path, and for strongly connected subdigraphs. An MSO formula to complete the proof of Theorem 3 is then trivial to construct by writing out Claim 1 in the formalism.

4 Discrete Milling is Hard for Bounded Pathwidth

In this section, we see that the maximum degree restriction implicit in the parameterization for our positive result is key to tractability for this problem. In the restricted version of the DISCRETE MILLING problem the turncost functions are pairing and symmetric. This is a significant assumption, but the outcome is still

negative, and the following result very much strengthens, in the parameterized setting, the NP-completeness result of Arkin *et al.* [1].

Theorem 4. DISCRETE MILLING (*with pairing and symmetric turn cost functions*) $W[1]$ -hard, with respect to (k, p) , where k is the number of turns and p is a bound on pathwidth.

Proof. The fpt-reduction is from MULTICOLOR CLIQUE, using an edge representation strategy, such as described, for example, in [4,8].

Suppose $G = (V, E)$ has V partitioned into color classes $C_i, i = 1, \dots, r$. The MULTICOLOR CLIQUE problem asks whether G contains a r -clique consisting of one vertex from each color class C_i . We assume that each color class of G has size n [4]. The color-class partition of V induces a partition of E into $\binom{r}{2}$ classes $E_{\{i,j\}}$, for $1 \leq i \neq j \leq r$:

$$E_{\{i,j\}} = \{e \in E : \exists u \in C_i \text{ and } \exists v \in C_j \text{ with } e \text{ incident on } u \text{ and } v \}.$$

We can also assume that all these edge-partition classes $E_{\{i,j\}}$ have the same size m . We index the vertices and edges of G as follows:

$$C_i = \{v(i, q) : 1 \leq q \leq n\} \quad \text{for } i = 1, \dots, r$$

$$E_{\{i,j\}} = \{e(\{i, j\}, l) : 1 \leq l \leq m\} \quad \text{for } 1 \leq i \neq j \leq r.$$

To refer to the incidence structure of G , we define functions $\pi_{\{i,j\}}^i(l)$ and $\pi_{\{i,j\}}^j(l)$ as follows:

$$\pi_{\{i,j\}}^i(l) = q : \text{ the edge } e(\{i, j\}, l) \text{ is incident on } v(i, q)$$

$$\pi_{\{i,j\}}^j(l) = q \text{ the edge } e(\{i, j\}, l) \text{ is incident on } v(j, q),$$

so the edge $e(\{i, j\}, l)$ is incident to $v(i, \pi_{\{i,j\}}^i(l))$ and $v(j, \pi_{\{i,j\}}^j(l))$.

We describe the construction of a graph G' , together with the sets S_v of turn-free pairs of edges for the vertices v of G' . We first describe the vertices of G' , and then specify a set of paths on these vertices. The edge set of the multi-graph G' is the (abstract) disjoint union of the sets of edges of these abstractly-defined paths, and it is understood that each path is turn-free, so that (for the most part), the sets S_v of turn-free pairs of v -incident edges for the vertices v of G' are implicit in these *generating paths* of G' .

The vertex set V' for G' is the union of the sets $V_0 \cup V_1 \cup V_2 \cup V_3 \cup V_4$,

$$V_0 = \{\sigma, \tau\}$$

$$V_1 = \{t[i, j] : 1 \leq i \neq j \leq r\}$$

$$V_2 = \{s[i, j] : 1 \leq i \neq j \leq r\}$$

$$V_3 = \{c[i, j, u] : 1 \leq i \neq j \leq r, 1 \leq u \leq n\}$$

$$V_4 = \{p[i, j, l] : 1 \leq i \neq j \leq r, 1 \leq l \leq m\}.$$

Thus $|V_1| = |V_2| = 2\binom{r}{2}$, $|V_3| = 2n\binom{r}{2}$ and $|V_4| = 2m\binom{r}{2}$.

The edge set of G' is (implicitly) described by a generating set of paths \mathcal{P} (two paths for each edge of G), together with a few more edges:

$$\mathcal{P} = \{P[i, j, e(\{i, j\}, l)] : 1 \leq i \neq j \leq r, 1 \leq l \leq m, \},$$

where the path $P[i, j, e(\{i, j\}, l)]$ (1) starts at the vertex $p[i, j, l]$; (2) next visits $s[i, j]$; (3) then visits the vertices $c[i, j, u]$, except for $u = \pi_{\{i, j\}}^i(l)$ (the *exceptional vertex* of this block), in *consecutive order*, meaning that the vertices are visited by increasing index u , modified by skipping the exceptional vertex; (4) then visits the vertex $c[i, j^*, \pi_{\{i, j\}}^i(l)]$, where j^* is defined to be $j + 1$, unless $j + 1 = i$, when $j^* = j + 2$, or $j = r$ and $i \neq 1$, when $j^* = 1$, or $j = r$ and $i = 1$, when $j^* = 2$; and then (5) ends at the vertex $t[i, j]$.

Intuitively, there are two paths in \mathcal{P} corresponding to each edge of G . If we fix i and consider that there are $r - 1$ blocks of vertices (each block consisting of n vertices, corresponding to the vertices of C_i), then what a path $P[i, j, e(\{i, j\}, l)]$ (corresponding to the l^{th} edge of $E_{\{i, j\}}$) does is “hit” every vertex of its “own” $\{i, j\}^{\text{th}}$ block, except the vertex $c[i, j^*, \pi_{\{i, j\}}^i(l)]$ of the block corresponding to the vertex of C_i to which the indexing edge of G is incident, and in the “next block” in a circular ordering of the $r - 1$ blocks established by the definition of j^* , does the complementary thing: in this “next block” it hits *only* the vertex corresponding to the vertex of C_i to which the indexing edge is incident in G , and then ends at $t[i, j]$.

At this stage of the construction, the edges of G' are partitioned into (turn-free) paths that run between vertices of V_1 and vertices of V_4 , where the latter have degree 1 (so far) and the vertices of V_1 have degree m (so far). We complete the construction of G' by adding a few more edges, specifying a few more turn-free pairs as we do so.

(A) Add edges between the pairs of vertices $p[i, j, l]$ and $p[j, i, l]$ for all $1 \leq i \neq j \leq r$ and $1 \leq l \leq m$. After these edges are added, we have reached a stage where all vertices in V_4 have degree 2 (and they will have degree 2 in G'). For each vertex of V_4 we make the pair of incident edges a turn-free pair.

Note that for any instance of the DISCRETE MILLING problem, the edge set is naturally and uniquely partitioned into maximal turn-free paths. At this stage of the construction, these paths all run between $t[i, j]$ and $t[j, i]$ for $1 \leq i < j \leq r$.

(B) Add some edges between the vertices of $V_0 \cup V_1$. Let \leq_{lex} denote the lexicographic order on the set (of pairs of indices) $\mathcal{I} = \{[i, j] : 1 \leq i < j \leq r\}$. Let $[i, j]^*$ denote the immediate successor of $[i, j]$ in the ordering of \mathcal{I} by \leq_{lex} . For $[i, j] \in \mathcal{I}$, let $\text{rev}[i, j] = [j, i]$. We add the edges (using the notation $u \cdot v$ for the creation of an edge between u and v):

- $t[\text{rev}[i, j]] \cdot t[[i, j]^*]$ for $1 \leq i < j \leq r$ and $[i, j] \neq [r, r - 1]$,
- $\sigma \cdot t[1, 2]$, and
- $t[r, r - 1] \cdot \tau$.

We do not specify any further turn-free pairs of vertex co-incident edges beyond (A) or implicit by being internal to the generating paths \mathcal{P} of G' . That completes the description of G' .

To complete the proof, we need to show that: (1) the graph G' will admit a k -turn covering walk, where $k = 2\binom{r}{2}$, if and only if G has a multicolor r -clique; and (2) G' has path-width at most $6\binom{r}{2} + 4$. For reasons of space, the arguments will appear in the full version of the paper.

5 Open Problems

We have studied the parameterized complexity of (several versions of) the discrete milling problem with turn costs and gave an initial classification with respect to several parameterizations. We believe that there is good motivation to study “highly structured” graph problems, that is problems involving a graph together with “other information”, since they are often able to engage applications better than simple graph problems. Our FPT results are impractical, but can they be improved? Our dynamic programming approach for GRID MILLING is a first step. In particular, it would be interesting to know if DISCRETE MILLING parameterized by (k, t, d) admits a polynomial kernel [7]. Our negative result provides one of the very few natural examples of a parameterized graph problem that is $W[1]$ -hard, parameterized by pathwidth. Another notable open question is whether DISCRETE MILLING parameterized by (k, d) , is FPT or $W[1]$ -hard. Our suspicion is that it is $W[1]$ -hard, but will require an even more elaborate reduction than the hardness result described here.

References

1. Arkin, E., Bender, M., Demaine, E., Fekete, S., Mitchell, J., Sethia, S.: Optimal covering tours with turn costs. *SIAM J. Computing* 35(3), 531–566 (2005)
2. Dorn, F., Penninkx, E., Bodlaender, H., Fomin, F.: Efficient exact algorithms on planar graphs: exploiting sphere cut branch decompositions. In: Brodal, G.S., Leonardi, S. (eds.) *ESA 2005*. LNCS, vol. 3669, pp. 95–106. Springer, Heidelberg (2005)
3. Downey, R.G., Fellows, M.R.: *Parameterized Complexity*. Springer, Heidelberg (1999)
4. Fellows, M., Fomin, F., Lokshtanov, D., Rosamond, F., Saurabh, S., Szeider, S., Thomassen, C.: On the complexity of some colorful problems parameterized by treewidth. In: Dress, A.W.M., Xu, Y., Zhu, B. (eds.) *COCOA*. LNCS, vol. 4616, pp. 366–377. Springer, Heidelberg (2007)
5. Flum, J., Grohe, M.: *Parameterized Complexity Theory*. Springer, Heidelberg (2006)
6. Held, M.: *On the Computational Geometry of Pocket Machining*. LNCS, vol. 500. Springer, Heidelberg (1991)
7. Niedermeier, R.: *Invitation to Fixed Parameter Algorithms*, vol. 31. Oxford University Press, Oxford (2006)
8. Szeider, S.: Not so easy problems for tree decomposable graphs. In: *International Conference on Discrete Mathematics (ICDM)*, pp. 161–171 (2008) (invited talk)