

1 Congestion-Free Rerouting of Flows on DAGs*

2 **Saeed Akhoondian Amiri**

3 Max-Planck Institute of Informatics, Germany
4 samiri@mpi-inf.mpg.de

5 **Szymon Dudycz**

6 University of Wroclaw, Poland
7 szymon.dudycz@gmail.com

8 **Stefan Schmid**

9 University of Vienna, Austria
10 stefan_schmid@univie.ac.at

11 **Sebastian Wiederrecht**

12 TU Berlin, Germany
13 sebastian.wiederrecht@tu-berlin.de

14 — Abstract —

15 Changing a given configuration in a graph into another one is known as a reconfiguration problem.
16 Such problems have recently received much interest in the context of algorithmic graph theory.
17 We initiate the theoretical study of the following reconfiguration problem: How to reroute k
18 unsplittable flows of a certain demand in a capacitated network from their current paths to their
19 respective new paths, in a congestion-free manner? This problem finds immediate applications,
20 e.g., in traffic engineering in computer networks. We show that the problem is generally NP-hard
21 already for $k = 2$ flows, which motivates us to study rerouting on a most basic class of flow graphs,
22 namely DAGs. Interestingly, we find that for general k , deciding whether an unsplittable multi-
23 commodity flow rerouting schedule exists, is NP-hard even on DAGs. Our main contribution is
24 a polynomial-time (fixed parameter tractable) algorithm to solve the route update problem for a
25 bounded number of flows on DAGs. At the heart of our algorithm lies a novel decomposition of
26 the flow network that allows us to express and resolve reconfiguration dependencies among flows.

27 **2012 ACM Subject Classification** F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph
28 Theory

29 **Keywords and phrases** Unsplittable Flows, Reconfiguration, DAGs, FPT, NP-Hardness

30 **Digital Object Identifier** 10.4230/LIPIcs...118

31 **1** Introduction

32 Reconfiguration problems are combinatorial problems which ask for a transformation of one
33 configuration into another one, subject to some (reconfiguration) rules. Reconfiguration
34 problems are fundamental and have been studied in many contexts, including puzzles and
35 games (such as Rubik's cube) [24], satisfiability [15], independent sets [16], vertex coloring [9],
36 or matroid bases [17], to just name a few.

37 Reconfiguration problems also naturally arise in the context of networking applications
38 and routing. For example, a fundamental problem in computer networking regards the

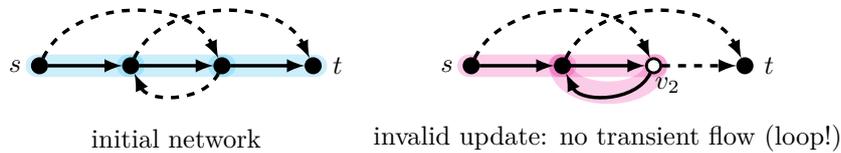
* The research of Saeed Amiri and Sebastian Wiederrecht was partly supported by the ERC consolidator grant DISTRUCT, agreement No 648527. Stefan Schmid was supported by the Danish VILLUM foundation project *ReNet*.



39 question of how to reroute traffic from the current path p_1 to a given new path p_2 , by
 40 changing the forwarding rules at routers (the *vertices*) one-by-one, while maintaining certain
 41 properties *during* the reconfiguration (e.g., short path lengths [7]). Route reconfigurations (or
 42 *updates*) are frequent in computer networks: paths are changed, e.g., to account for changes
 43 in the security policies, in response to new route advertisements, during maintenance (e.g.,
 44 replacing a router), to support the migration of virtual machines, etc. [13].

45 This paper initiates the study of a basic *multi-commodity flow rerouting problem*: how
 46 to reroute a set of *unsplittable flows* (with certain bandwidth demands) in a capacitated
 47 network, from their current paths to their respective new paths *in a congestion-free manner*.
 48 The problem finds immediate applications in traffic engineering [4], whose main objective
 49 is to avoid network congestion. Interestingly, while congestion-aware routing and traffic
 50 engineering problems have been studied intensively in the past [1, 10, 11, 12, 18, 19, 20, 22],
 51 surprisingly little is known today about the problem of how to reconfigure resp. *update* the
 52 routes of flows. Only recently, due to the advent of Software-Defined Networks (SDNs), the
 53 problem has received much attention in the networking community [3, 8, 14, 21].

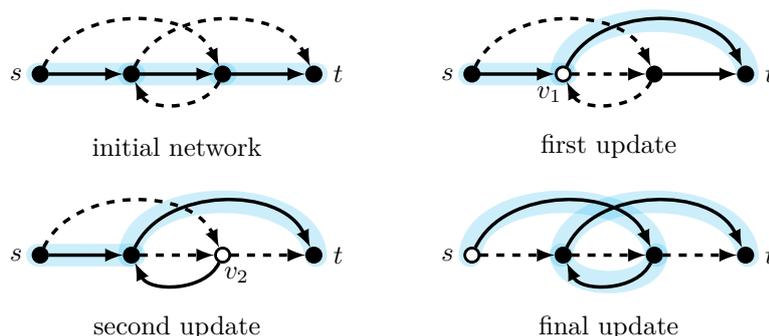
54 Figure 1 presents a simple example of the consistent rerouting problem considered in
 55 this paper, for just a *single* flow: the flow needs to be rerouted from the solid path to
 56 the dashed path, by changing the forwarding links at routers one-by-one. The example
 57 illustrates a problem that might arise from updating the vertices in an invalid order: if
 58 vertex v_2 is updated first, a forwarding loop is introduced: the transient flow from s to
 59 t becomes invalid. Thus, router updates need to be scheduled intelligently over time: A
 60 feasible sequence of updates for this example is given in Figure 2. Note that the example
 61 is kept simple intentionally: when moving from a single flow to multiple flows, additional
 62 challenges are introduced, as the flows may compete for bandwidth and hence interfere.



■ **Figure 1 Example:** We are given an initial network consisting of exactly one active flow F^o (solid edges) and the inactive edges (i.e., inactive forwarding rules) of the new flow F^u to which we want to reroute (dashed edges). Together we call the two flows an (update) pair $P = (F^o, F^u)$. Updating the outgoing edges of a vertex means activating all previously inactive outgoing edges of F^u , and deactivating all other edges of the old flow F^o . Initially, the blue flow is a valid (transient) (s, t) -flow. If the update of vertex v_2 takes effect first, an invalid (not transient) flow is introduced (in pink): traffic is forwarded in a loop, hence (temporarily) invalidating the path from s to t .

63 **Contributions.** This paper initiates the algorithmic study of a fundamental unsplittable
 64 multicommodity flow rerouting problem. We present a rigorous formal model and show that
 65 the problem of rerouting flows in a congestion-free manner is NP-hard already for two flows
 66 on general graphs. This motivates us to focus on a most fundamental type of flow graphs,
 67 namely the DAG. The main results presented in this paper are the following:

- 68 1. Deciding whether a consistent network update schedule exists in general graphs is NP-hard,
 69 already for 2 flows.
- 70 2. For constant k , we present a linear-time (fixed parameter tractable) algorithm which
 71 finds a feasible update schedule on DAGs in time and space $2^{O(k \log k)} O(|G|)$, whenever
 72 such a consistent update schedule exists.
- 73 3. For general k , deciding whether a feasible schedule exists is NP-hard even on loop-free



■ **Figure 2** *Example*: We revisit the network of Figure 1 and reroute from F^o to F^u without interrupting the connection between s and t along a unique (transient) path (in blue). To avoid the problem seen in Figure 1, we first update the vertex v_1 in order to establish a shorter connection from s to t . Once this update has been performed, the update of v_2 can be performed without creating a loop. Finally, by updating s , we complete the rerouting.

74 networks (i.e., DAGs).

75 Against the backdrop that the problem of *routing* disjoint paths on DAGs is known to
 76 be $W[1]$ -hard [23] and computing routes *subject to congestion* even harder [1], our finding
 77 that the multicommodity flow *rerouting* problem is fixed parameter tractable on DAGs is
 78 intriguing.

79 **Technical Novelty.** Our algorithm is based on a novel decomposition of the flow graph
 80 into so-called *blocks*. This block decomposition allows us to express dependencies between
 81 flows, and we represent dependencies between blocks by a (directed) dependency graph D .
 82 The structure of D is sophisticated, hence to analyze it, we first construct a helper graph H .
 83 In our first main technical lemma, we show that if there is an independent set I in H , then
 84 the dependency graph that corresponds to the vertices of I is a DAG (Theorem 11). So we
 85 may concentrate on a subgraph of D with a simpler structure, which we use to prove the
 86 next main technical lemma: there is a congestion-free rerouting if and only if the maximum
 87 independent set in H is large enough (Lemma 14). We are left with the challenge that finding
 88 a maximum independent set is a hard problem, even in our very restricted graph classes. We
 89 hence carefully modify H to obtain a much simpler graph of *bounded pathwidth*, without
 90 losing any critical properties. Thanks to these lemmas, the proof of the main theorem will
 91 follow.

92 In addition to our algorithmic contributions, we present NP-hardness proofs. These
 93 hardness proofs are based on novel and non-trivial insights into the flow rerouting problem,
 94 which might be helpful for similar problems in the future.

95 2 Model and Definitions

96 The problem can be described in terms of edge capacitated directed graphs. In what follows,
 97 we will assume basic familiarity with directed graphs and we refer the reader to [5] for more
 98 background. We denote a directed edge e with head v and tail u by $e = (u, v)$. For an
 99 undirected edge e between vertices u, v , we write $e = \{u, v\}$; u, v are called endpoints of e .

100 A **flow network** is a directed capacitated graph $G = (V, E, s, t, c)$, where s is the *source*,
 101 t the *terminal*, V is the set of vertices with $s, t \in V$, $E \subseteq V \times V$ is a set of ordered pairs

102 known as edges, and $c: E \rightarrow \mathbb{N}$ a capacity function assigning a capacity $c(e)$ to every edge
103 $e \in E$.

104 Our problem, as described above is a multi-commodity flow problem and thus may have
105 *multiple* source-terminal pairs. To simplify the notation but without loss of generality, in
106 what follows, we define flow networks to have exactly one source and one terminal. In fact,
107 we can model any number of different sources and terminals by adding one super source
108 with edges of unlimited capacity to all original sources, and one super terminal with edges of
109 unlimited capacity leading there from all original terminals.

110 An (s, t) -flow F of capacity $d \in \mathbb{N}$ is a *directed path* from s to t in a flow network such
111 that $d \leq c(e)$ for all $e \in E(F)$. Given a family \mathcal{F} of (s, t) -flows F_1, \dots, F_k with demands
112 d_1, \dots, d_k respectively, we call \mathcal{F} a **valid flow set**, or simply **valid**, if $c(e) \geq \sum_{i: e \in E(F_i)} d_i$.

113 Recall that we consider the problem of how to reroute a current (old) flow to a new
114 (update) flow, and hence we will consider such flows in “update pairs”: An **update flow**
115 **pair** $P = (F^o, F^u)$ consists of two (s, t) -flows F^o , the *old flow*, and F^u , the *update flow*,
116 each of demand d . A graph $G = (V, E, \mathcal{P}, s, t, c)$, where (V, E, s, t, c) is a flow network,
117 and $\mathcal{P} = \{P_1, \dots, P_k\}$ with $P_i = (F_i^o, F_i^u)$, a family of update flow pairs of demand d_i ,
118 $V = \bigcup_{i \in [k]} V(F_i^o \cup F_i^u)$ and $E = \bigcup_{i \in [k]} E(F_i^o \cup F_i^u)$, is called **update flow network** if the
119 two families $\mathcal{P}^o = \{F_1^o, \dots, F_k^o\}$ and $\mathcal{P}^u = \{F_1^u, \dots, F_k^u\}$ are valid. For an illustration, recall
120 the initial network in Figure 2: The old flow is presented as the directed path made of solid
121 edges and the new one is represented by the dashed edges.

122 Given an update flow network $G = (V, E, \mathcal{P}, s, t, c)$, an **update** is a pair $\mu = (v, P) \in$
123 $V \times \mathcal{P}$. An update (v, P) with $P = (F^o, F^u)$ is *resolved* by deactivating all outgoing edges of
124 F^o incident to v and activating all of its outgoing edges of F^u . Note that at all times, there
125 is at most one outgoing and at most one incoming edge, for any flow at a given vertex. So
126 the deactivated edges of F^o can no longer be used by the flow pair P (but now the newly
127 activated edges of F^u can).

128 For any set of updates $U \subset V \times \mathcal{P}$ and any flow pair $P = (F^o, F^u) \in \mathcal{P}$, $G(P, U)$ is the
129 update flow network consisting exactly of the vertices $V(F^o) \cup V(F^u)$ and the edges of P
130 that are active after resolving all updates in U .

131 As an illustration, after the second update in Figure 2, one of the original solid edges is
132 still not deactivated. However, already two of the new edges have become solid (i.e., active).
133 So in the picture of the second update, the set $U = \{(v_1, P), (v_2, P)\}$ has been resolved.

134 We are now able to determine, for a given set of updates, which edges we can and
135 which edges we cannot use for our routing. In the end, we want to describe a process of
136 reconfiguration steps, starting from the *initial state*, in which no update has been resolved,
137 and finishing in a state where the only active edges are exactly those of the new flows, of
138 every update flow pair.

139 The flow pair P is called **transient** for some set of updates $U \subseteq V \times \mathcal{P}$, if $G(P, U)$
140 contains a unique valid (s, t) -flow $T_{P,U}$. If there is a family $\mathcal{P} = \{P_1, \dots, P_k\}$ of update flow
141 pairs with demands d_1, \dots, d_k respectively, we call \mathcal{P} a **transient family** for a set of updates
142 $U \subseteq V \times \mathcal{P}$, if and only if every $P \in \mathcal{P}$ is transient for U . The family of transient flows after
143 all updates in U are resolved is denoted by $\mathcal{T}_{\mathcal{P},U} = \{T_{P_1,U}, \dots, T_{P_k,U}\}$.

144 We again refer to Figure 2. In each of the different states, the transient flow is depicted
145 as the light blue line connecting s to t and covering only solid (i.e., active) edges.

146 An **update sequence** $(\sigma_i)_{i \in [|V \times \mathcal{P}|]}$ is an ordering of $V \times \mathcal{P}$. We denote the set of updates
147 that is resolved after step i by $U_i = \bigcup_{j=1}^i \sigma_j$, for all $i \in [|V \times \mathcal{P}|]$.

148 **► Definition 1 (Consistency Rule).** Let σ be an update sequence. We require that for any
149 $i \in [|V \times \mathcal{P}|]$, there is a family of transient flow pairs $\mathcal{T}_{\mathcal{P},U_i}$.

150 To ease the notation, we will denote an update sequence $(\sigma)_{i \in [|V \times \mathcal{P}|]}$ simply by σ and
 151 for any update (u, P) we write $\sigma(u, P)$ for the the position i of (u, P) within σ . An update
 152 sequence is **valid**, if every set $U_i, i \in [|V \times \mathcal{P}|]$, obeys the consistency rule.

153 We note that this consistency rule models and consolidates the fundamental properties
 154 usually studied in the literature, such as congestion-freedom [8] and loop-freedom [21].

155 ► **Definition 2** (*k*-NETWORK FLOW UPDATE PROBLEM). Given an update flow network G
 156 with k update flow pairs, is there a feasible update sequence σ ?

157 3 On Hardness of 2-Flow Update in General Graphs

158 It is easy to see that for an update flow network with a single flow pair, feasibility is always
 159 guaranteed. However, it turns out that for two flows, the problem becomes hard in general.

160 ► **Theorem 3.** *Deciding whether a feasible network update schedule exists is NP-hard already*
 161 *for $k = 2$ flows.*

162 The proof, briefly sketched in the following, is by reduction from 3-SAT. Let C be any
 163 3-SAT formula with n variables and m clauses. Denote the variables by X_1, \dots, X_n and the
 164 clauses by C_1, \dots, C_m . The resulting update flow network is denoted by $G(C)$. Assume that
 165 the variables are ordered by their indices, and their appearance in each clause respects this
 166 order.

167 We create 2 update flow pairs, a blue one $B = (B^o, B^u)$ and a red one $R = (R^o, R^u)$,
 168 both of demand 1. The pair B contains gadgets corresponding to the variables. The order in
 169 which the edges of each of those gadgets are updated will correspond to assigning a value to
 170 the variable. The pair R on the other hand contains gadgets representing the clauses: they
 171 have edges that are “blocked” by the variable edges of B . Therefore, we will need to update
 172 B to enable the updates of R .

173 4 Rerouting Flows in DAGs

174 We now consider the flow rerouting problem when the underlying flow graph is acyclic. In
 175 particular, we identify an important substructure arising for flow-pairs in acyclic graphs,
 176 which we call *blocks*. These blocks will play a major role in both the hardness proof and the
 177 algorithm presented in this section.

178 Let $G = (V, E, \mathcal{P}, s, t, c)$ be an acyclic update flow network, i.e., we assume that the
 179 graph (V, E) is a DAG. Let \prec be a topological order on the vertices $V = \{v_1, \dots, v_n\}$.
 180 Let $P_i = (F_i^o, F_i^u)$ be an update flow pair of demand d and let $v_1^i, \dots, v_{\ell_i}^i$ be the induced
 181 topological order on the vertices of F_i^o ; analogously, let $u_1^i, \dots, u_{\ell_i}^i$ be the order on F_i^u .
 182 Furthermore, let $V(F_i^o) \cap V(F_i^u) = \{z_1^i, \dots, z_{k_i}^i\}$ be ordered by \prec as well.

183 The subgraph of $F_i^o \cup F_i^u$ induced by the set $\{v \in V(F_i^o \cup F_i^u) \mid z_j^i \prec v \prec z_{j+1}^i\}, j \in$
 184 $[k_i - 1]$, is called the j th *block* of the update flow pair F_i , or simply the j th *i*-*block*. We will
 185 denote this block by b_j^i .

186 For a block b , we define $\mathcal{S}(b)$ to be the *start of the block*, i.e., the smallest vertex w.r.t. \prec ;
 187 similarly, $\mathcal{E}(b)$ is the *end of the block*: the largest vertex w.r.t. \prec .

188 Let $G = (V, E, \mathcal{P}, s, t, c)$ be an update flow network with $\mathcal{P} = \{P_1, \dots, P_k\}$ and let \mathcal{B} be
 189 the set of its blocks. We define a binary relation $<$ between two blocks as follows. For two
 190 blocks $b_1, b_2 \in \mathcal{B}$, where b_1 is an i -block and b_2 a j -block, $i, j \in [k]$, we say $b_1 < b_2$ (b_1 is
 191 *smaller than* b_2) if one of the following holds.

- 192 i $\mathcal{S}(b_1) \prec \mathcal{S}(b_2)$,
 193 ii if $\mathcal{S}(b_1) = \mathcal{S}(b_2)$ then $b_1 < b_2$, if $\mathcal{E}(b_1) \prec \mathcal{E}(b_2)$,
 194 iii if $\mathcal{S}(b_1) = \mathcal{S}(b_2)$ and $\mathcal{E}(b_1) = \mathcal{E}(b_2)$ then $b_1 < b_2$, if $i < j$.

195 Let b be an i -block and P_i the corresponding update flow pair. For a feasible update sequence
 196 σ , we will denote the round $\sigma(\mathcal{S}(b), P_i)$ by $\sigma(b)$. We say that an i -block b is *updated*, if all
 197 edges in $b \cap F_i^u$ are active and all edges in $b \cap F_i^o \setminus F_i^u$ are inactive. We will make use of a
 198 basic, but important observation on the structure of blocks and how they can be updated.
 199 This structure is the key to our flow reconfiguration algorithm (presented below), as it allows
 200 us to consider the update of blocks as a whole, rather than vertex-by-vertex.

201 **► Lemma 4.** *Let b be a block of the flow pair $P = (F^u, F^o)$. Then in a feasible update*
 202 *sequence σ , all vertices (resp. their outgoing edges belonging to P) in $F^u \cap b - \mathcal{S}(b)$ are*
 203 *updated strictly before $\mathcal{S}(b)$. Moreover, all vertices in $b - F^u$ are updated strictly after $\mathcal{S}(b)$*
 204 *is updated.*

205 **► Lemma 5.** *Let G be an update flow network and σ a valid update sequence for G . Then*
 206 *there exists a feasible update sequence σ' which updates every block in consecutive rounds.*

207 Recall that G is acyclic and every flow pair in G forms a single block. Let σ be a feasible
 208 update sequence of G . We suppose in σ , every block is updated in consecutive rounds
 209 (Lemma 5). For a single flow F , we write $\sigma(F)$ for the round where the last edge of F was
 210 updated.

211 4.1 Updating k -Flows in DAGs is NP-complete

212 We first show that if k is part of the input, the congestion-free flow reconfiguration problem
 213 is even hard on the DAG. Hence the algorithm presented in the following is essentially tight.
 214 To prove the theorem, we use a polynomial time reduction from the 3-SAT problem.

215 **► Theorem 6.** *Finding a feasible update sequence for k -flows is NP-complete, even if the*
 216 *update graph G is acyclic.*

217 4.2 Linear Time Algorithm for Constant Number of Flows on DAGs

218 By Theorem 6 we cannot hope to find a polynomial time algorithm that finds a feasible
 219 update sequence. However, if the problem is parameterized by the number k of flows, a
 220 rerouting sequence can be computed in FPT-linear time if the update graph is acyclic. In
 221 this subsection we describe an algorithm to solve the network update problem on DAGs in
 222 time $2^{O(k \log k)} O(|G|)$, for arbitrary k . In the remainder of this section, we assume that every
 223 block has at least 3 vertices (otherwise, postponing such block updates will not affect the
 224 solution).

225 We say a block b_1 *touches* a block b_2 (denoted by $b_1 \succ b_2$) if there is a vertex $v \in b_1$ such
 226 that $\mathcal{S}(b_2) \prec v \prec \mathcal{E}(b_2)$, or there is a vertex $u \in b_2$ such that $\mathcal{S}(b_1) \prec u \prec \mathcal{E}(b_1)$. If b_1
 227 does not touch b_2 , we write $b_1 \not\succeq b_2$. Clearly, the relation is symmetric, i.e., if $b_1 \succ b_2$ then
 228 $b_2 \succ b_1$.

229 For some intuition, consider a drawing of G which orders vertices w.r.t. \prec in a line.
 230 Project every edge on that line as well. Then two blocks touch each other if they have a
 231 common segment on that projection.

232 **Proof Sketch:** Before delving into details, we provide the main ideas behind our algorithm.
 233 We can think about the update problem on DAGs as follows. Our goal is to compute a
 234 feasible update order for the (out-)edges of the graph. There are at most k flows to be

235 updated for each edge, resulting in $k!$ possible orders and hence a brute force complexity of
 236 $O(k!^{|G|})$ for the entire problem. We can reduce this complexity by considering blocks instead
 237 of edges.

238 The update of a given i -block b_i might depend on the update of a j -block sharing at least
 239 one edge of b_i . These dependencies can be represented as a directed graph. If this graph
 240 does not have any directed cycles, it is rather easy to find a feasible update sequence, by
 241 iteratively updating sink vertices.

242 There are several issues here: First of all these dependencies are not straight-forward
 243 to define. As we will see later, they may lead to representation graphs of exponential size.
 244 In order to control the size we might have to relax our definition of dependency, but this
 245 might lead to a not necessarily acyclic graph which will then need further refinement. This
 246 refinement is realized by finding a suitable subgraph, which alone is a hard problem in general.
 247 To overcome the above problems, we proceed as follows.

248 Let $\text{TouchSeq}(b)$ contain all feasible update sequences for the blocks that touch b : still a
 249 (too) large number, but let us consider them for now. For two distinct blocks b, b' , we say
 250 that two sequences $s \in \text{TouchSeq}(b), s' \in \text{TouchSeq}(b')$ are *consistent*, if the order of any
 251 common pair of blocks is the same in both s, s' . If for some block b , $\text{TouchSeq}(b) = \emptyset$, there
 252 is no feasible update sequence for G : b cannot be updated.

253 We now consider a graph H whose vertices correspond to elements of $\text{TouchSeq}(b)$, for
 254 all $b \in \mathcal{B}$. Connect all pairs of vertices originating from the same $\text{TouchSeq}(b)$. Connect
 255 all pairs of vertices if they correspond to inconsistent elements of different $\text{TouchSeq}(b)$. If
 256 (and only if) we find an independent set of size $|\mathcal{B}|$ in the resulting graph, the update orders
 257 corresponding to those vertices are mutually consistent: we can update the entire network
 258 according to those orders. In other words, the update problem can be reduced to finding an
 259 independent set in the graph H .

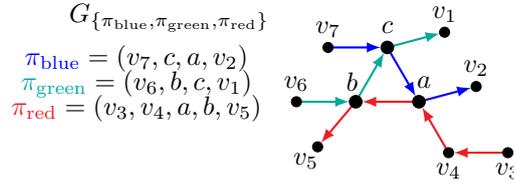
260 However, there are two main issues with this approach. First, H can be very large. A
 261 single $\text{TouchSeq}(b)$ can have exponentially many elements. Accordingly, we observe that we
 262 can assume a slightly different perspective on our problem: we linearize the lists $\text{TouchSeq}(b)$
 263 and define them sequentially, bounding their size by a function of k (the number of flows).
 264 The second issue is that finding a maximum independent set in H is hard. The problem
 265 is equivalent to finding a clique in the complement of H , a $|\mathcal{B}|$ -partite graph where every
 266 partition has bounded cardinality. We can prove that for an n -partite graph where every
 267 partition has bounded cardinality, finding an n -clique is NP-complete. So, in order to solve
 268 the problem, we either should reduce the number of partitions in H (but we cannot) or
 269 modify H to some other graph, further reducing the complexity of the problem. We do the
 270 latter by trimming H and removing some extra edges, turning the graph into a very simple
 271 one: a graph of *bounded path width*. Then, by standard dynamic programming, we find the
 272 independent set of size $|\mathcal{B}|$ in the trimmed version of H : this independent set matches the
 273 independent set I of size $|\mathcal{B}|$ in H (if it exists). At the end, reconstructing a correct update
 274 order sequence from I needs some effort. As we have reduced the size of $\text{TouchSeq}(b)$ and
 275 while not all possible update orders of all blocks occur, we show that they suffice to cover all
 276 possible feasible solutions. We provide a way to construct a valid update order accordingly.

277 With these intuitions in mind, we now present a rigorous analysis. Let $\pi_{S_1} = (a_1, \dots, a_{\ell_1})$
 278 and $\pi_{S_2} = (a'_1, \dots, a'_{\ell_2})$ be permutations of sets S_1 and S_2 . We define the *core* of π_{S_1} and
 279 π_{S_2} as $\text{core}(\pi_{S_1}, \pi_{S_2}) := S_1 \cap S_2$. We say that two permutations π_1 and π_2 are *consistent*,
 280 $\pi_1 \approx \pi_2$, if there is a permutation π of symbols of $\text{core}(\pi_1, \pi_2)$ such that π is a subsequence
 281 of both π_1 and π_2 .

282 The **dependency graph** is a labelled graph defined recursively as follows. The depend-
 283 ency graph of a single permutation $\pi = (a_1, \dots, a_\ell)$, denoted by G_π , is a directed path
 284 v_1, \dots, v_ℓ , and the label of the vertex $v_i \in V(G_\pi)$ is the element a with $\pi(a) = i$. We denote
 285 by $\text{Labels}(G_\pi)$ the set of all labels of G_π .

286 Let G_Π be a dependency graph of the set of permutations Π and $G_{\Pi'}$ the dependency
 287 graph of the set Π' . Then, their union (by identifying the same vertices) forms the dependency
 288 graph $G_{\Pi \cup \Pi'}$ of the set $\Pi \cup \Pi'$. Note that such a dependency graph is not necessarily acyclic
 289 (see Figure 3).

290 We call a permutation π of blocks of a subset $\mathcal{B}' \subseteq \mathcal{B}$ *congestion free*, if the following
 291 holds: it is possible to update the blocks in π in the graph $G_{\mathcal{B}}$ (the graph on the union of
 292 blocks in \mathcal{B}), in order of their appearance in π , without violating any edge capacities in $G_{\mathcal{B}}$.
 293 Note that we do not respect all conditions of our *Consistency Rule* (Definition 1) here.



■ **Figure 3** *Example:* The dependency graph of three pairwise consistent permutations π_{blue} , π_{green} and π_{red} . Each pair of those permutation has exactly one vertex in common and with this the cycle (a, b, c) is created. With such cycles being possible, a dependency graph does not necessarily contain sink vertices. To get rid of them, we certainly need some more refinements.

294 In the approach we are taking, one of the main advantages we have is the nice properties of
 295 blocks when it comes to updating. The following algorithm formalizes the procedure already
 296 described in Lemma 5. The correctness follows directly from said lemma. Let $P = (F^o, F^u)$
 297 be a given flow pair.

298

299 **Algorithm 1. Update a Free Block b**

- 300 1. Resolve (v, P) for all $v \in F^u \cap b - \mathcal{S}(b)$.
- 301 2. Resolve $(\mathcal{S}(b), P)$.
- 302 3. Resolve (v, P) for all $v \in (b - F^u)$.
- 303 4. For any edge in $E(b \cap F^u)$ check whether d_{F^u} together with the other loads currently
 304 active on e exceed $c(e)$. If so output: *Fail*.

305 ► **Lemma 7.** *Let π be a permutation of the set $\mathcal{B}_1 \subseteq \mathcal{B}$. Whether π is congestion free can be
 306 determined in time $O(k \cdot |G|)$.*

307 The smaller relation defines a total order on all blocks in G . Let $\mathcal{B} = \{b_1, \dots, b_{|\mathcal{B}|}\}$ and
 308 suppose the order is $b_1 < \dots < b_{|\mathcal{B}|}$.

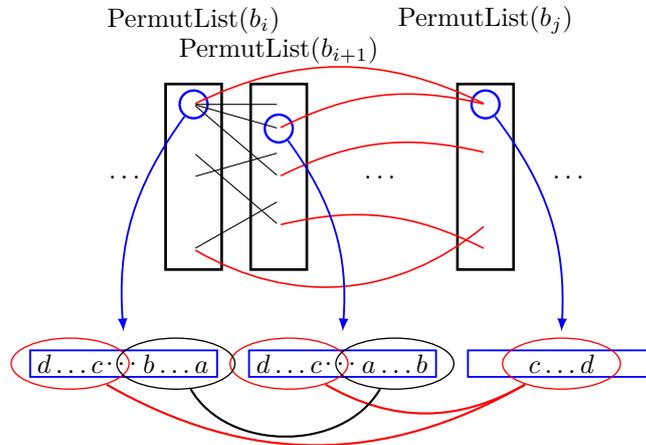
309 We define an auxiliary graph H which will help us find a suitable dependency graph for
 310 our network. We first provide some high-level definitions relevant to the construction of
 311 the graph H only. Exact definitions will follow in the construction of H , and will be used
 312 throughout the rest of this section.

313 Recall that \mathcal{B} is the set of all blocks in G . We define another set of blocks \mathcal{B}' and
 314 initialize it as \mathcal{B} ; the construction of H is iterative, and in each iteration, we eliminate a
 315 block from \mathcal{B}' . At the end of the construction of H , \mathcal{B}' is empty. For every block $b \in \mathcal{B}'$,

we also define the set $\text{TouchingBlocks}(b)$ of blocks which touch the block b , note that this set is dynamically defined: it depends on \mathcal{B}' . Another set which is defined for every block b is the set $\text{PermutList}(b)$; this set actually corresponds to a set of vertices, each of which corresponds to a valid congestion free permutation of blocks in $\text{TouchingBlocks}(b)$. Clearly if $\text{TouchingBlocks}(b)$ does not contain any congestion-free permutation, then $\text{PermutList}(b)$ is an empty set. As we already mentioned, every vertex $v \in \text{PermutList}(b)$ comes with a **label** which corresponds to some congestion-free permutation of elements of $\text{TouchingBlocks}(b)$. We denote that permutation by $\text{Label}(v)$.

Construction of H : We recursively construct a labelled graph H from the blocks of G as follows.

- i Set $H := \emptyset$, $\mathcal{B}' := \mathcal{B}$, $\text{PermutList} := \emptyset$.
- ii For $i := 1, \dots, |\mathcal{B}|$ do
 - 1 Let $b := b_{|\mathcal{B}|-i+1}$.
 - 2 Let $\text{TouchingBlocks}(b) := \{b'_1, \dots, b'_t\}$ be the set of blocks in \mathcal{B}' touched by b .
 - 3 Let $\pi := \{\pi_1, \dots, \pi_t\}$ be the set of congestion free permutations of $\text{TouchingBlocks}(b)$.
 - 4 Set $\text{PermutList}(b) := \emptyset$.
 - 5 For $i \in [t]$ create a vertex v_{π_i} with $\text{Label}(v_{\pi_i}) = \pi_i$ and set $\text{PermutList}(b) := \text{PermutList}(b) \cup v_{\pi_i}$.
 - 6 Set $H := H \cup \text{PermutList}(b)$.
 - 7 Add edges between all pairs of vertices in $H[\text{PermutList}(b)]$.
 - 8 Add an edge between every pair of vertices $v \in H[\text{PermutList}(b)]$ and $u \in V(H) - \text{PermutList}(b)$ if the labels of v and u are inconsistent.
 - 9 Set $\mathcal{B}' := \mathcal{B}' - b$.



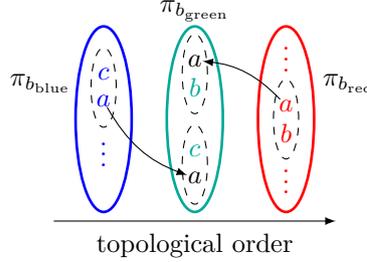
■ **Figure 4** *Example:* The graph H consists of vertex sets $\text{PermutList}(b_i)$, $i \in [|\mathcal{B}|]$, where each such partition contains all congestion free sequences of the at most k iteratively chosen touching blocks. In the whole graph, we then create edges between the vertices of two such partitions if and only if the corresponding sequences are inconsistent with each other, as seen in the three highlighted sequences. Later we will distinguish between such edges connecting vertices of neighbouring partitions (w.r.t. the topological order of their corresponding blocks), $\text{PermutList}(b_i)$ and $\text{PermutList}(b_{i+1})$, and partitions that are further away, $\text{PermutList}(b_i)$ and $\text{PermutList}(b_j)$. Edges of the latter type, depicted as red in the figure, are called long edges and will be deleted in the trimming process of H .

We have the following lemmas based on our construction.

118:10 Congestion-Free Rerouting of Flows on DAGs

341 ► **Lemma 8.** For Item (ii2) of the construction of H , $t \leq k$ holds.

342 ► **Lemma 9 (Touching Lemma).** Let $b_{j_1}, b_{j_2}, b_{j_3}$ be three blocks (w.r.t. $<$) where $j_1 < j_2 < j_3$.
 343 Let b_z be another block such that $z \notin \{j_1, j_2, j_3\}$. If in the process of constructing H , b_z is in
 344 the touch list of both b_{j_1} and b_{j_3} , then it is also in the touch list of b_{j_2} .



► **Figure 5 Example:** Select one of the permutations of length at most k from every $\text{PermutList}(b)$. These permutations obey the Touching Lemma. Taking the three permutations from the example in Figure 3, we can see that the Touching Lemma forces a to be in the green permutation as well. Assuming consistency, this would mean a to come *before* b and *after* c . Hence $a <_{\pi_{\text{green}}} b$ and $b <_{\pi_{\text{green}}} a$, a contradiction. So if our permutations are derived from H and are consistent, we will show that cycles cannot occur in their dependency graph.

345 For an illustration of the property described in the Touching Lemma, see Figure 5: it
 346 refers to the dependency graph of Figure 3. This example also points out the problem with
 347 directed cycles in the dependency graph and the property of the Touching Lemma, which is
 348 crucial for Observation 10 and Theorem 11.

349 We prove a series of lemmas in regard to the dependency graph of elements of H , to
 350 establish the base of the inductive proof for Lemma 13.

351 ► **Observation 10.** Let π be a permutation of a set S . Then the dependency graph G_π does
 352 not contain a cycle.

353 ► **Lemma 11.** Let π_1, π_2 be permutations of sets S_1, S_2 such that π_1, π_2 are consistent. Then
 354 the dependency graph $G_{\pi_1 \cup \pi_2}$ is acyclic.

355 In the next lemma, we need a closure of the dependency graph of permutations which we
 356 define as follows.

357 ► **Definition 12 (Permutation Graph Closure).** The *permutation graph closure*, or simply
 358 *closure*, of a permutation π is the graph G_π^+ obtained from taking the transitive closure of
 359 G_π , i.e. its vertices and labels are the same as G_π and there is an edge (u, v) in G_π^+ if there
 360 is a path starting at u and ending at v in G_π . Similarly the *permutation graph closure* of a
 361 set of permutations $\Pi = \{\pi_1, \dots, \pi_n\}$ is the graph obtained by taking the union of $G_{\pi_i}^+$'s (for
 362 $i \in [n]$) by identifying vertices of the same label.

363 In the above definition, note that if Π is a set of permutations, then $G_\Pi \subseteq G_\Pi^+$. The
 364 following lemma generalizes Theorem 11 and Observation 10 and uses them as the base of
 365 its inductive proof.

366 ► **Lemma 13.** Let $I = \{v_{\pi_1}, \dots, v_{\pi_\ell}\}$ be an independent set in H . Then the dependency
 367 graph G_Π , for $\Pi = \{\pi_1, \dots, \pi_\ell\}$, is acyclic.

368 **Proof.** Instead of working on G_Π , we can work on its closure G_Π^+ as defined above. First
 369 we observe that every edge in G_Π also appears in G_Π^+ , so if there is a cycle in G_Π , the same
 370 cycle exists in G_Π^+ .

371 We prove that there is no cycle in G_Π^+ . By Theorem 11 and Observation 10 there is no
 372 cycle of length at most 2 in G_Π^+ ; otherwise there is a cycle in G_Π which consumes at most
 373 two consistent permutations.

374 For the sake of contradiction, suppose G_Π^+ has a cycle and let $C = (a_1, \dots, a_n) \subseteq G_\Pi^+$ be
 375 a shortest cycle in G_Π^+ . By Theorem 11 and Observation 10 we know that $n \geq 3$.

376 In the following, because we work on a cycle C , whenever we write any index i we consider
 377 it w.r.t. its cyclic order on C , in fact $i \bmod |C| + 1$. So for example, $i = 0$ and $i = n$ are
 378 identified as the same indices; similarly for $i = n + 1, i = 1$, etc.

379 Recall the construction of the dependency graph where every vertex $v \in C$ corresponds
 380 to some block b_v . In the remainder of this proof we do not distinguish between the vertex v
 381 and the block b_v .

382 Let π_v be the label of a given vertex $v \in I$. For each edge $e = (a_i, a_{i+1}) \in C$, there is a
 383 permutation π_{v_i} such that (a_i, a_{i+1}) is a subsequence of π_{v_i} and additionally the vertex v_i is in
 384 the set I . So there is a block b^i such that π_{v_i} is a permutation of the set $\text{TouchingBlocks}(b^i)$.

385 The edge $e = (a_i, a_{i+1})$ is said to *represent* b^i , and we call it the representative of π_{v_i} .
 386 For each i we fix one block b^i which is represented by the edge (a_i, a_{i+1}) (note that one edge
 387 can represent many blocks, but here we fix one of them). We define the set of those blocks
 388 as $B^I = \{b^1, \dots, b^\ell\}$ and state the following claim.

389 *Claim 1.* For every two distinct vertices $a_i, a_j \in C$, either there is no block $b \in B^I$ such that
 390 $a_i, a_j \in \text{TouchingBlocks}(b)$ or if $a_i, a_j \in \text{TouchingBlocks}(b)$ then (a_i, a_j) or (a_j, a_i) is an
 391 edge in C . Additionally $|B^I| = |C|$.

392 By the above claim we have $\ell = n$. W.l.o.g. suppose $b^1 < b^2 < \dots < b^n$. There is an $i \in [n]$
 393 such that (a_{i-1}, a_i) represents b^1 , we fix this i .

394 *Claim 2.* If (a_{i-1}, a_i) represents b^1 then (a_{i-2}, a_{i-1}) represents b^2 .

395 Similarly we can prove the endpoints of the edges, that have a_i as their head, are in b^2 .

396 *Claim 3.* If (a_{i-1}, a_i) represents b^1 then (a_i, a_{i+1}) represents b^2 .

397 By Claims 2 and 3 we have that both (a_{i-2}, a_{i-1}) and (a_i, a_{i+1}) represent b^2 hence by
 398 Claim 1 they are the same edge. Thus there is a cycle on the vertices a_{i-1}, a_i in G_Π^+ and this
 399 gives a cycle in G_Π on at most 2 consistent permutations which is a contradiction according
 400 to Theorem 11. ◀

401 The following lemma is the key to establish a link between independent sets in H and
 402 feasible update sequences of the corresponding update flow network G .

403 ▶ **Lemma 14.** *There is a feasible sequence of updates for an update network G on k flow pairs,
 404 if and only if there is an independent set of size $|\mathcal{B}|$ in H . Additionally if the independent
 405 set $I \subseteq V(H)$ of size $|\mathcal{B}|$ together with its vertex labels are given, then there is an algorithm
 406 which can compute a feasible sequence of updates for G in $O(k \cdot |G|)$.*

407 With Lemma 14, the update problem boils down to finding an independent set of size $|\mathcal{B}|$
 408 in H .

409 Finding an independent set of size $|\mathcal{B}|$ in H is a hard problem already on very restricted
 410 class families. Hence, we trim H to avoid the above problem. We will use the special
 411 properties of the touching relation of blocks. We say that an edge $e \in E(H)$ is *long*, if one
 412 end of e is in $\text{PermutList}(b_i)$, and the other in $\text{PermutList}(b_j)$ where $j > i + 1$. The *length*
 413 of e is $j - i$. Delete all long edges from H to obtain the graph R_H . We prove the following
 414 lemmas.

415 ► **Lemma 15.** *There is an algorithm which computes R_H in time $O((k \cdot k!)^2 |G|)$.*

416 ► **Lemma 16.** *H has an independent set I of size $|\mathcal{B}|$ if, and only if, I is also an independent
 417 set of size $|\mathcal{B}|$ in R_H .*

418 R_H is a much simpler graph compared to H , which helps us find a large independent set
 419 of size $|\mathcal{B}|$ (if exists). We have the following lemma.

420 ► **Lemma 17.** *There is an algorithm that finds an independent set I of size exactly $|\mathcal{B}|$ in R_H
 421 if such an independent set exists; otherwise it outputs that there is no such an independent
 422 set. The running time of this algorithm is $O(|R_H|)$.*

423 Our main theorem is now a corollary of the previous lemmas and algorithms.

424 ► **Theorem 18.** *There is a linear time FPT algorithm for the network update problem on
 425 an acyclic update flow network G with k flows (the parameter), which finds a feasible update
 426 sequence, if it exists; otherwise it outputs that there is no feasible solution for the given
 427 instance. The algorithm runs in time $O(2^{O(k \log k)} |G|)$.*

428 5 Conclusion

429 This paper initiated the study of a natural and fundamental reconfiguration problem: the
 430 congestion-free rerouting of unsplittable flows. Interestingly, we find that while *computing*
 431 disjoint paths on DAGs is $W[1]$ -hard [23] and finding routes under congestion as well [1],
 432 *reconfiguring* multicommodity flows is fixed parameter tractable on DAGs. However, we also
 433 show that the problem is NP-hard for an arbitrary number of flows.

434 In future work, it will be interesting to chart a more comprehensive landscape of the
 435 computational complexity for the network update problem. In particular, it would be
 436 interesting to know whether the complexity can be reduced further, e.g., to $2^{O(k)} O(|G|)$.
 437 More generally, it will be interesting to study other flow graph families, especially more
 438 sparse graphs or graphs of bounded DAG width [2, 6]. Finally, besides feasibility, it remains
 439 to study algorithms to efficiently compute *short* schedules.

440 **Acknowledgements.** We would like to thank Stephan Kreutzer, Arne Ludwig and Roman
 441 Rabinovich for discussions on this problem.

442 — References —

- 443 1 Saeed Akhoondian Amiri, Stephan Kreutzer, Dániel Marx, and Roman Rabinovich. Rout-
 444 ing with congestion in acyclic digraphs. In *41st International Symposium on Mathematical
 445 Foundations of Computer Science, MFCS*, pages 7:1–7:11, 2016.
- 446 2 Saeed Akhoondian Amiri, Stephan Kreutzer, and Roman Rabinovich. Dag-width is pspace-
 447 complete. *Theor. Comput. Sci.*, 655:78–89, 2016.
- 448 3 Saeed Akhoondian Amiri, Arne Ludwig, Jan Marcinkowski, and Stefan Schmid. Transi-
 449 ently consistent sdn updates: Being greedy is hard. In *23rd International Colloquium on
 450 Structural Information and Communication Complexity, SIROCCO*, 2016.

- 451 4 D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow. Rsvp-te: Extensions
452 to rsvp for lsp tunnels. In *RFC 3209*, 2001.
- 453 5 Jørgen Bang-Jensen and Gregory Gutin. *Digraphs - theory, algorithms and applications*.
454 Springer, 2002.
- 455 6 Dietmar Berwanger, Anuj Dawar, Paul Hunter, Stephan Kreutzer, and Jan Obdržálek. The
456 dag-width of directed graphs. *J. Comb. Theory, Ser. B*, 102(4):900–923, 2012.
- 457 7 Paul Bonsma. The complexity of rerouting shortest paths. *Theoretical computer science*,
458 510:1–12, 2013.
- 459 8 Sebastian Brandt, Klaus-Tycho Förster, and Roger Wattenhofer. On Consistent Migration
460 of Flows in SDNs. In *Proc. 36th IEEE International Conference on Computer Communica-*
461 *tions (INFOCOM)*, 2016.
- 462 9 Luis Cereceda, Jan Van Den Heuvel, and Matthew Johnson. Finding paths between 3-
463 colorings. *Journal of graph theory*, 67(1):69–82, 2011.
- 464 10 Chandra Chekuri, Alina Ene, and Marcin Pilipczuk. Constant congestion routing of sym-
465 metric demands in planar directed graphs. In *43rd International Colloquium on Automata,*
466 *Languages, and Programming, ICALP*, 2016.
- 467 11 Chandra Chekuri, Sreeram Kannan, Adnan Raja, and Pramod Viswanath. Multicommod-
468 ity flows and cuts in polymatroidal networks. *SIAM J. Comput.*, 44(4):912–943, 2015.
- 469 12 Shimon Even, Alon Itai, and Adi Shamir. On the complexity of timetable and multicom-
470 modity flow problems. *SIAM J. Comput.*, 5(4):691–703, 1976.
- 471 13 Klaus-Tycho Foerster, Stefan Schmid, and Stefano Vissicchio. Survey of consistent network
472 updates. In *ArXiv Technical Report*, 2016.
- 473 14 Klaus-Tycho Förster, Ratul Mahajan, and Roger Wattenhofer. Consistent Updates in
474 Software Defined Networks: On Dependencies, Loop Freedom, and Blackholes. In *Proc.*
475 *15th IFIP Networking*, 2016.
- 476 15 Parikshit Gopalan, Phokion G Kolaitis, Elitza Maneva, and Christos H Papadimitriou. The
477 connectivity of boolean satisfiability: computational and structural dichotomies. *SIAM*
478 *Journal on Computing*, 38(6):2330–2355, 2009.
- 479 16 Robert A Hearn and Erik D Demaine. Pspace-completeness of sliding-block puzzles and
480 other problems through the nondeterministic constraint logic model of computation. *The-*
481 *oretical Computer Science*, 343(1-2):72–96, 2005.
- 482 17 Takehiro Ito, Erik Demaine, Nicholas Harvey, Christos Papadimitriou, Martha Sideri, Ry-
483 uhei Uehara, and Yushi Uno. On the complexity of reconfiguration problems. *Algorithms*
484 *and Computation*, pages 28–39, 2008.
- 485 18 Ken-ichi Kawarabayashi, Yusuke Kobayashi, and Stephan Kreutzer. An excluded half-
486 integral grid theorem for digraphs and the directed disjoint paths problem. In *Proc. Sym-*
487 *posium on Theory of Computing (STOC)*, pages 70–78, 2014.
- 488 19 Jon M. Kleinberg. Decision algorithms for unsplittable flow and the half-disjoint paths
489 problem. In *Proc. 30th Annual ACM Symposium on Theory of Computing (STOC)*, pages
490 530–539, 1998.
- 491 20 Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems
492 and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- 493 21 Arne Ludwig, Jan Marcinkowski, and Stefan Schmid. Scheduling loop-free network updates:
494 It’s good to relax! In *Proc. ACM PODC*, 2015.
- 495 22 Martin Skutella. Approximating the single source unsplittable min-cost flow problem. In
496 *Proc. IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 136–145, 2000.
- 497 23 Aleksandrs Slivkins. Parameterized tractability of edge-disjoint paths on directed acyclic
498 graphs. *SIAM Journal on Discrete Mathematics*, 24(1):146–157, 2010.
- 499 24 Jan van den Heuvel. The complexity of change. *Surveys in combinatorics*, 409(2013):127–
500 160, 2013.