

Self-Adjusting Grid Networks

Chen Avin^a, Ingo van Duijn^{b,*}, Maciej Pacut^c, Stefan Schmid^c

^a*School of Electrical and Computer Engineering, Ben Gurion University of the Negev*

^b*Department of Computer Science, Aalborg University*

^c*Faculty of Computer Science, University of Vienna*

Abstract

Emerging networked systems become increasingly flexible, reconfigurable, and “self-*”. This introduces an opportunity to adjust networked systems in a demand-aware manner, leveraging spatial and temporal locality in the workload for *online* optimizations. However, it also introduces a tradeoff: while more frequent adjustments can improve performance, they also entail higher reconfiguration costs. This paper studies self-adjusting *grid* networks in which frequently communicating nodes (e.g., virtual machines) are moved topologically closer in an online and demand-aware manner, striking a balance between the benefits and costs of reconfigurations. We show that the underlying algorithmic problem can be seen as a generalization of the classic dynamic list update problem known from self-adjusting data structures: in a network, requests can occur between *node pairs*. This pairwise optimization turns out to be significantly harder than the classical problem it generalizes. Our main result is a general $\Omega(\log n)$ lower bound even in a scenario where not only the self-adjusting network topology forms a grid but also the communication pattern (the demand graph); hence, in principle, in this scenario the demand can be served at constant cost, once it is learned. To demonstrate the challenge of adapting a network to pair-wise communication requests, we also discuss the 1-dimensional grid in more details and present an online algorithm that is $O(\log n)$ -competitive in this setting.

Keywords: Self-*, self-adjusting data structures, self-adjusting networks, competitive analysis, distributed algorithms, communication networks.

1. Introduction

Communication networks are becoming increasingly flexible, along three main dimensions: routing (enabler: software-defined networking), embedding (enabler: virtualization), and topology (enabler: reconfigurable optical technologies). In particular, the possibility to quickly reconfigure communication networks, e.g., by migrating (virtualized) communication endpoints (9) or

by reconfiguring the (optical) topology (16; 11), allows these networks to become *demand-aware*: i.e., to adapt to the traffic pattern they serve, in an online and “self-*” manner. In particular, in a *self-adjusting* network, frequently communicating node pairs can be moved *topologically closer*, saving communication costs (e.g., bandwidth, energy) and improving performance (e.g., latency, throughput).

However, today, we still do not have a good understanding yet of the algorithmic problems underlying self-

*Corresponding author: ingovanduijn@gmail.com

19 adjusting networks. The design of such algorithms faces
 20 several challenges. In particular, as the demand is often
 21 not known ahead of time, *online* algorithms are required
 22 to react to changes in the workload in a clever way;
 23 ideally, such online algorithms are “competitive” even
 24 when compared to an optimal offline algorithm which
 25 knows the demand ahead of time. Furthermore, online
 26 algorithms need to strike a balance between the benefits
 27 of adjustments (i.e., improved performance and/or re-
 28 duced costs) and their costs (i.e., frequent adjustments
 29 can temporarily harm consistency and/or performance,
 30 or come at energy costs).

31 The vision of self-adjusting networks is reminiscent of
 32 self-adjusting data structures such as *self-adjusting lists*
 33 and *splay trees*, which optimize themselves toward the
 34 workload. In particular, the *dynamic list update prob-*
 35 *lem*, introduced already in the 1980s by Sleator and
 36 Tarjan in their seminal work (22), asks for an online
 37 algorithm to reconfigure an unordered linked list data
 38 structure, such that a sequence of lookup requests is
 39 served optimally and at minimal reconfiguration costs
 40 (i.e., pointer rotations). It is well-known that a sim-
 41 ple *move-to-front* strategy, which immediately promotes
 42 each accessed element to the front of the list, is *dy-*
 43 *namically optimal*, that is, has a constant competitive
 44 ratio.

45 This paper initiates the study of a most basic self-
 46 adjusting *grid network*, which can be seen as a gen-
 47 eralization of the dynamic list update problem, along
 48 two dimensions: first, instead of considering a 1-
 49 dimensional list, we consider a d -dimensional grid;
 50 second and more importantly, while data structures
 51 serve requests originating from the front of the list (the
 52 “root”) to access data items, networks serve *commu-*

53 *nication* requests between *pairs of nodes*. The objec-
 54 tive in this *pairwise* generalization, is to move nodes
 55 (e.g., virtual machines) which currently communicate
 56 frequently, closer to each other, while accounting for
 57 reconfiguration costs.

58 1.1. Formal Model

59 We study the design of a self-adjusting network which
 60 optimizes itself toward the pairwise communication re-
 61 quests it serves. At the core of this model is a set V
 62 of communicating nodes and a static network N . One
 63 can think of the nodes in V as virtual machines, and
 64 the nodes of N as physical hosts; the virtual machines
 65 are hosted on the nodes of N (one virtual machine per
 66 host), and their communication is facilitated by the net-
 67 work links (edges) of N . An embedding (injection) of V
 68 into the nodes of N is called a *host configuration* and
 69 is denoted h ; the set of all configurations is denoted
 70 $C_{V \hookrightarrow N}$.

71 A sequence of communication requests $\sigma =$
 72 $(\sigma_0, \sigma_1, \dots, \sigma_m)$ is called *demand*, and can be
 73 served at a cost of the sum of its constituent requests.
 74 Specifically, a communication request is a pair of
 75 communicating nodes, and a configuration $h \in C_{V \hookrightarrow N}$
 76 can *serve* a communication request $(u, v) \in V \times V$ at
 77 cost $d_N(h(u), h(v))$, where d_N denotes the (hop) distance
 78 in N . See Figure 1 for an illustration of a network with
 79 communicating nodes configured. The cost of serving
 80 communication (4, 6) in the figure’s configuration is 2.
 81 Note that the demand (dashed lines) in the figure can
 82 actually be configured so that every communication
 83 request can be served at cost 1.

84 In our model, we are interested in minimizing the cost of
 85 the demand by allowing *host reconfigurations*. A recon-

86 figure is a transition from a configuration h to h' by a
 87 sequence of *host migrations*, where one migration com-
 88 prises two neighboring network nodes exchanging their
 89 embedded communicating nodes; the cost of a reconfig-
 90 uration is the minimal number of migrations necessary
 91 to implement it. In summary, this yields:

92 **Definition 1.** Given a finite demand $\sigma =$
 93 $(\sigma_0, \sigma_1, \dots, \sigma_m)$ and a sequence of configurations
 94 $h_0, h_1, \dots, h_m \in C_{V \rightarrow N}$. The cost of serving σ is the
 95 sum of serving each σ_i in h_i plus the reconfiguration
 96 cost between subsequent configurations h_i, h_{i+1} .

97 *The Pairwise Grid Update Problem.* In particular, we
 98 study the problem of designing a self-adjusting *grid*
 99 *network*: a network N whose topology forms a d -
 100 dimensional grid, or d -*grid*, for which we use the fol-
 101 lowing notation:

- 102 • A d -grid is a graph $N = (V_N, E)$ where $V_N \subset$
 103 $\mathbb{N} \times \dots \times \mathbb{N}$, with $E = \{((n_1, \dots, n_d), (m_1, \dots, m_d)) \mid$
 104 $\exists i \text{ s.t. } |n_i - m_i| = 1 \text{ and } \forall j \neq i, n_j = m_j\}$.
- 105 • The *length* of a grid is its largest dimension, i.e.,
 106 $\max_{(n_1, \dots, n_d) \in V} n_i$.
- 107 • A *subgrid* of a d -grid is any subset of the grid
 108 which forms a d' -grid, with $1 \leq d' \leq d$




Figure 1: A 2×3 grid network N (solid squares), with communicating nodes (dashed circles) configured on it.

109 **Definition 2 (PAIRWISE GRID UPDATE).** Let V , h , and σ

110 be as before, with N representing the d -dimensional grid
 111 (for some constant d), with all dimensions having length
 112 $n^{1/d}$.

113 The cost of serving a request $\sigma_i = (u, v) \in \sigma$ is given by
 114 $d_N(h(u), h(v)) = \|h(u), h(v)\|_1$, where $\|\cdot\|_1$ is the ℓ_1 norm
 115 (taxicab distance).

116 As a first step, we in this paper consider the PAIRWISE
 117 GRID UPDATE problem for the case where the demand
 118 has the same topology as the network: the communica-
 119 tion requests between nodes also form a grid topology.
 120 We will refer to this graph as the *demand graph* or, syn-
 121 nonymously, *request graph*. More formally, to denote the
 122 demand as a graph, we let $E_i = \{\sigma_1, \dots, \sigma_i\}$ be the first
 123 i requests of σ interpreted as a set of edges on V , so
 124 that the request graph $R(\sigma) = (V, E_m)$ models the entire
 125 demand σ .

126 Since the self-adjusting network and the demand have
 127 the same topology, in principle, there exists a configura-
 128 tion of V into the grid network such that any request can
 129 be served at cost 1. However, initially, the embedding
 130 may be arbitrarily far from optimal, and hence needs to
 131 be learned.

Online Competitive Ratio. Recall that the cost incurred
 by an algorithm A on σ is the sum of communication
 and reconfiguration costs. In the realm of online algo-
 rithms and competitive analysis, we compare an online
 algorithm ON to an offline algorithm OFF which has
 complete knowledge of σ ahead of time. We want to
 devise online algorithms ON which minimize the com-
 petitive ratio ρ :

$$\rho = \max_{\sigma} \frac{\text{cost}(ON(\sigma))}{\text{cost}(OFF(\sigma))}$$

1.2. Related Work

One important area of related work arises in the context of the dynamic list update problem. Since the groundbreaking work by Sleator and Tarjan on amortized analysis and self-adjusting data structures (22), researchers have explored many interesting variants of self-adjusting data structures, also using randomized algorithms (20), studying the power of lookaheads (1; 3), or devising offline algorithms (5; 19). The deterministic Move-To-Front (MTF) algorithm is known to optimally solve the standard formulation of the list update problem: it is constant competitive (22; 4). To the best of our knowledge, the exact competitive ratio in the randomized setting (against an oblivious adversary) is still an open problem (3; 23). The state-of-the-art randomized algorithm (3) makes an initial random choice between two known algorithms that have different worst-case request sequences, relying on the BIT (20) and TIMESTAMP (2) algorithms.

We also note that the self-adjusting network design problem for pairwise requests can be considered a special case of general online problems such as the online Metrical Task System (MTS) problems. However, given the exponential number of possible configurations, the competitive ratio of generic MTS algorithms will be high if applied to our more specific problem (at least according to the existing bounds). Furthermore, we note that in case the demand graph and the network have the same topology, the problem can be seen as a learning problem and is hence related to bandits theory (13); however, in our model, reconfigurations come at a cost.

There also exists work on self-adjusting networks whose (bounded-degree) topology can be adjusted,

e.g., (8; 10; 21; 18; 15). However, these approaches cannot be applied in our context, where only the mapping between the virtual machines and the physical servers can be changed, but not the network itself.

The paper closest to ours are by Avin et al. (6; 7) and by Olver et al. (17). In (6), the authors consider a problem related to self-adjusting grid networks, for a special type of input sequences which are chosen *i.i.d.* from a given pairwise distribution, namely a symmetric product distribution. For this model, the authors propose a local, distributed algorithm that is constant competitive for sufficiently long sequences; they also show that the problem is NP-hard. In (7), the authors present different swapping heuristics for migrating virtual machines on a hypercubic networks, aiming to aggressively collocate communicating nodes. Olver et al. (17) introduced the Itinerant List Update (ILU) problem: a relaxation of the classic dynamic list update problem in which the pointer no longer has to return to a home location after each request. The authors present an $\Omega(\log n)$ lower bound on the randomized competitive ratio and also describe a polynomial-time offline algorithm and prove that it achieves an approximation ratio of $O(\log^2 n)$. In contrast, we in our paper focus on online algorithms and demand graphs (where the edges are communication requests) which form a grid. In fact, we show that the lower bound $\Omega(\log n)$ even holds in this restrictive case, at least for deterministic algorithms. We also present an online algorithm which matches this bound in our model.

Bibliographic note. A preliminary version of this paper was presented at SSS 2019 (12). For this journal version we significantly revised the presentation, and also include additional technical results such as an NP-

200 hardness proof.

201 1.3. Contributions

202 This paper initiates the study of a class of basic self-
203 adjusting networks (d -grids), which optimize them-
204 selves toward the dynamically changing demand (re-
205 stricted to the same topology), while amortizing recon-
206 figuration cost. The underlying algorithmic problem is
207 natural and motivated by emerging reconfigurable com-
208 munication networks where virtual machines can be mi-
209 grated in a demand-aware manner. In the special case
210 of a 1-dimensional grid (a line), the problem can also
211 be seen as a pairwise generalization of the fundamental
212 dynamic list update problem. Our main result is a neg-
213 ative one: we show that unlike the classic dynamic list
214 update problem, which admits for constant-competitive
215 online algorithms, there is an $\Omega(\log n)$ lower bound on
216 the competitive ratio of any deterministic online algo-
217 rithm for the PAIRWISE GRID UPDATE problem, even if the
218 demand forms the same grid topology as the physical
219 network. We further show that the offline variant of the
220 problem is NP-complete. We also demonstrate the chal-
221 lenges of designing online algorithms by presenting an
222 online algorithm which is $O(\log n)$ -competitive for long
223 enough sequences on 1-grids (i.e., *lines*), where the de-
224 mand has the same topology.

225 1.4. Organization

226 The remainder of this paper is organized as follows. In
227 Section 2, we put the problem and its challenges into
228 perspective with respect to the list update problem. In
229 Section 3 we derive the lower bound for PAIRWISE GRID
230 UPDATE. Then, in Section 4 we present the upper bound
231 for PAIRWISE LIST UPDATE and the NP-completeness of its
232 offline variant. We conclude in Section 5.

233 2. From Classic to Pairwise List Update

234 To provide an intuition of the challenges involved in
235 designing online algorithms for PAIRWISE GRID UPDATE
236 problems and to put the problem into perspective, we
237 first revisit the classic list update problem and then dis-
238 cuss why similar techniques fail if applied to communi-
239 cating node *pairs*, i.e., scenarios where requests do not
240 come from fixed points in the network (e.g. the head of
241 a list network).

242 The (*dynamic*) *list update problem* (22) introduced by
243 Sleator and Tarjan over 30 years ago is one of the most
244 fundamental and oldest online problems: Given a set
245 of n elements stored in a linked list, how to update the
246 list over time such that it optimally serves a request se-
247 quence $\tau = (\tau_1, \tau_2, \dots)$ where for each i , $\tau_i \in V$ is an
248 arbitrary element stored in the list? The cost incurred by
249 an algorithm is the sum of the access costs (i.e. scanning
250 from the *front* of the list to the accessed element) and the
251 number of migrations (*swaps* in their terminology). As
252 accesses to the list elements start at the front of the list,
253 it makes sense to amortize high access costs by mov-
254 ing frequently accessed elements closer to the front
255 of the list. In fact, the well-known *Move-To-Front* (MTF)
256 algorithm even moves an accessed element to the front
257 *immediately*, and is known to be *constant competitive*:
258 its cost is at most a constant factor worse than that of
259 an optimal offline algorithm which knows the entire se-
260 quence τ ahead of time (22). Throughout the literature,
261 slightly different cost models have been used for the list
262 update problem. Generally, a *cursor* is located at the
263 head of the list at each request. Then, the algorithm can
264 perform two operations, each operation incurring unit
265 cost. i) *Move* the cursor to the left, or to the right, one
266 position; the element in the new position is referred to

267 as touched. ii) Swap the element at the cursor with the
 268 element one position to the left or right; the cursor also
 269 moves.

270 In the 1-dimensional pair-wise update problem, upon a
 271 request $\sigma_i = (s_i, t_i)$, the cursor is placed at s_i instead of
 272 the head of the list, and t_i needs to be looked up. To
 273 demonstrate the significance of this difference, we first
 274 present a paraphrased version of the proof by Tarjan and
 275 Sleator showing the dynamic optimality of MTF. After
 276 that, we showcase a simple access sequence differenti-
 277 ating the two problems.

278 2.1. *Expositional Proof for Optimality of MTF*

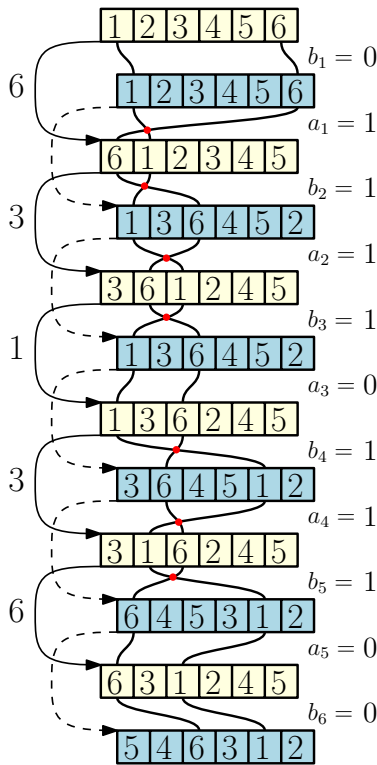


Figure 2: MTF (yellow) and A (blue) on $\tau = 6, 3, 1, 3, 6$

279 While the potential argument used to show dynamic opti-
 280 mality of the move-to-front strategy for the list access
 281 problem yields a very elegant and succinct proof (22),
 282 it lacks intuition which makes it difficult to generalize

283 the argument. The key idea in the potential argument
 284 is to compare the execution of MTF to the execution of
 285 an arbitrary algorithm A. The algorithm is fixed for the
 286 analysis, but any valid algorithm can be used, e.g. the
 287 optimal offline algorithm. The state (represented by a
 288 list) of MTF and A are juxtaposed at every access, com-
 289 paring how the order of elements in both lists differ.
 290 In fact, it is sufficient to only consider the relative or-
 291 der of two arbitrary but fixed elements, call them u and
 292 v . Consider the order of u and v in the state of A be-
 293 fore it performs the i th access. If this order is the same
 294 as in MTF before it performs the i th access, let $b_i = 0$
 295 and otherwise $b_i = 1$. Similarly, if their relative or-
 296 der is the same in MTF after its i th access, let $a_i = 0$
 297 and otherwise $a_i = 1$. This describes an inversion se-
 298 quence $b_1 a_1 b_2 a_2 \dots b_m a_m$. Figure 2 illustrates this for
 299 MTF and an arbitrarily chosen algorithm A on a sequence
 300 $\tau = 6, 3, 1, 3, 6$, with the inversions of 1 and 6 described
 301 by the sequence 01111011100.

302 Suppose that $\tau_i \in \{u, v\}$ and that MTF touches u and v
 303 while accessing τ_i . The proof by Tarjan and Sleator
 304 boils down to three observations.

305 **Observation 1.** MTF inverts u and v relative to A by
 306 accessing τ_i , i.e. $b_i \neq a_i$.

307 **Observation 2.** If $b_i = 0$, MTF and A agree on the order
 308 of u and v before τ_i . Since MTF touches both, A also
 309 touches both in order to access τ_i .

310 **Observation 3.** For $b_i = 1$, let $j < i$ be the largest index
 311 such that $b_j = 0$ or $a_j = 0$ (note that j exists because
 312 $b_1 = 0$). When $a_j = 0$, and thus $b_{j+1} = 1$, A inverts
 313 u and v and therefore must have touched both. When
 314 $b_j = 0$, and thus $a_j = 1$, MTF inverts u and v and one of
 315 them is τ_j . By Observation 2, if $b_j = 0$ and MTF touches

316 u and v to access τ_j , then A does as well.

317 The last observation is essentially the amortized argu-
 318 ment rephrased as a charging argument. We can now
 319 easily prove the dynamic optimality of MTF.

320 **Theorem 1** (Tarjan & Sleator). MTF is 4-competitive.

321 *Proof.* We prove that for all $\tau_i = v$ where MTF touches
 322 u , there is a move by A touching u . MTF first moves the
 323 cursor to τ_i , and then swaps τ_i to the front. Along the
 324 way it touches u twice, once with a move and once with
 325 a swap, incurring a cost of 2.

326 For $b_i = 0$ (resp. $b_i = 1$), we use Observation 2 (resp.
 327 3) to charge the cost to A touching u while accessing
 328 τ_i (resp. τ_j). By Observation 1, $b_i \neq a_i$, and thus for
 329 any $\tau_k \in \{u, v\}$ with $i < k$, the largest index $j' < k$ with
 330 $b_{j'} = 0$ or $a_{j'} = 0$ must be at least i , and therefore $j <$
 331 $i \leq j'$. This guarantees that MTF charges at most a cost
 332 of 4 to one move of A. Since all the cost incurred by MTF
 333 is charged to some move of A, the claim follows. \square

334 In the original work by Tarjan and Sleator, MTF is
 335 shown to be 2-competitive. This is because their cost
 336 model allows accessed elements to be moved to the
 337 front ‘for free’. If we allow this as well, the cursor
 338 touches u only once to access v , resulting in a factor
 339 2.

340 2.2. The Challenge of PAIRWISE LIST UPDATE

341 Generalizing dynamic list update to pairwise requests
 342 introduces a number of challenges already present in 1-
 343 dimensional grids, which render the problem more dif-
 344 ficult. First, the natural inversion argument no longer
 345 works: a reference point such as the front of the list is
 346 missing in the pairwise setting. This makes it harder to




Figure 3: A star graph used to construct a cyclic sequence of requests $\sigma_c = (c, v_1), (c, v_2), \dots, (c, v_{n-1}), (c, v_1), \dots$

347 relate algorithms to each other, and to perform a poten-
 348 tial function analysis of the competitive ratio. Second,
 349 for general request graphs $R(\sigma)$ (describing the demand
 350 pattern), an online algorithm needs to be able to essen-
 351 tially “recognize” certain patterns over time.

352 Regarding the latter, consider the set of nodes $V =$
 353 $\{v_1, \dots, v_n\}$ and let τ_c be a cyclic sequence: for all
 354 $\tau_i, \tau_{i+1} \in \tau_c$ with $\tau_i = v_j$ and $\tau_{i+1} = v_k$ it holds that
 355 $j+1 = k \pmod{n-1}$. From this we construct a similar se-
 356 quence σ_c for the pairwise problem, on the set of nodes
 357 $V \cup \{c\}$, with $\sigma_i = (c, \tau_i)$. This yields a star graph
 358 as denoted in Figure 3. An offline algorithm can effi-
 359 ciently serve the cyclic order by first embedding the ele-
 360 ments in the order v_1, \dots, v_k , and then moving the element
 361 c one position further after every request. If the cost of
 362 embedding the initial order is dominated by serving all
 363 requests, then the amortized cost is $O(1)$ per request (per
 364 cycle there are $n - 1$ moves of cost $O(1)$ and once c is
 365 moved a distance n). However, in the list update model,
 366 any sequence cycling through all elements is a worst-
 367 case sequence with $\Omega(n)$ per request. This demonstrates
 368 that a “dynamic cursor” can mean a factor n difference
 369 in cost. What the sequence σ_c also demonstrates, is that

370 aggregating elements around a highly communicative
 371 node is suboptimal; in the particular case of σ_c , it is
 372 this central node that needs to be moved.

373 Another pattern is a request sequence σ that forms a
 374 connected path in the request graph $R(\sigma)$ (describing
 375 the demand). When restricted to only these patterns, the
 376 1-dimensional pairwise problem corresponds to the *Itin-*
 377 *erant List Update Problem (ILU)* studied in (17). In this
 378 work it is shown that deriving non-trivial upper bounds
 379 on the competitive ratio already seems notoriously hard
 380 (even offline approximation factors are relatively high).
 381 Note that the star example can be expressed as a path,
 382 i.e. $\sigma'_c = (c, v_1), (v_1, c), (c, v_2), (v_2, c), (c, v_3), \dots$, demon-
 383 strating the significance of understanding simple request
 384 patterns for pairwise requests. This, among other, moti-
 385 vates us to specifically consider request graphs with a
 386 linear demand in this paper.

387 3. Lower Bound for PAIRWISE GRID UPDATE

388 This section derives a lower bound on the competitive
 389 ratio of any algorithm for PAIRWISE GRID UPDATE, where
 390 both the self-adjusting network as well as the demand
 391 feature a grid topology. The proof relies on a techni-
 392 cal statement which can be independently understood
 393 of its application in the lower bound. We therefore
 394 first present the technical part in Section 3.1, and then
 395 present the adversarial lower bound strategy in Sec-
 396 tion 3.2.

397 3.1. Technical Proof

398 **Theorem 2.** Let x_1, \dots, x_k and y_1, \dots, y_k be sequences
 399 of k nonnegative numbers, and let x (resp. y) denote
 400 $\sum_{i=1}^k x_i$ (resp. $\sum_{i=1}^k y_i$). Let the *weight* of an involu-

401 *tion*¹ over the indices $1, \dots, k$ be defined as $w(f) =$
 402 $\sum_{i=1}^k x_i y_{f(i)}$.

403 The average weight over all involutions is $\Omega(\frac{xy}{k})$.




Figure 4: A staircase of 8 points based on the sequences x_1, \dots, x_8 and y_1, \dots, y_8 . The values $w(2, 5)$, $w(8, 4)$, and $w(7, 7)$ are visualized as the area of rectangles highlighted in red, green, and blue respectively.

404 *Proof.* Let \mathcal{I}_k denote the set of all involutions on a set of
 405 k elements, where $|\mathcal{I}_k| = T(k)$ is given by the recurrence
 406 $T(k) = T(k-1) + (k-1)T(k-2)$ with $T(0) = T(1) = 1$.
 407 For every pair of distinct indices i, j , there are $T(k-2)$
 408 involutions $f \in \mathcal{I}_k$ such that $f(i) = j$ (namely for all
 409 involutions on the remaining $k-2$ indices). Similarly
 410 for every index i , there are $T(k-1)$ involutions such that
 411 $f(i) = i$. Thus, for every ordered pair of (not necessarily
 412 distinct) indices i, j there are *at least* $T(k-2)$ involutions
 413 with $f(i) = j$.

For convenience we define a *staircase* of points $p_j =$
 $(\sum_{i=1}^j x_i, \sum_{i=1}^j y_i)$. Observe that we can subdivide the

¹A function f such that $f(f(x)) = x$ for all x . One can think of it as a matching that allows self-pairings.

rectangle defined by p_k and the origin into k^2 axis-aligned rectangles, so that the area of every such rectangle corresponds to the weight of one ordered pair of indices (see Figure 4). Since every ordered pair of indices appears in at least $T(k-2)$ involutions, their weight (and thus the corresponding rectangle), contributes at least $T(k-2)$ times in the sum of weights over all involutions. This means that the area xy of the complete rectangle contributes $T(k-2)$ times to that sum:

$$\begin{aligned} \frac{\sum_{f \in \mathcal{I}_k} w(f)}{|\mathcal{I}_k|} &\geq \frac{T(k-2) \cdot \sum_{i=1}^k \sum_{j=1}^k w(i, j)}{T(k)} \\ &= \frac{T(k-2)}{T(k)} \cdot xy \end{aligned}$$

To lower bound $\frac{T(k-2)}{T(k)}$, we first define $R(n) = \frac{T(n)}{T(n-1)}$ and observe that this definition is equivalent to the one in Lemma 1:

$$\begin{aligned} R(n) &= \frac{T(n)}{T(n-1)} \\ &= \frac{T(n-1) + (n-1)T(n-2)}{T(n-1)} \\ &= 1 + (n-1) \frac{T(n-2)}{T(n-1)} \\ &= 1 + \frac{n-1}{R(n-1)} \end{aligned}$$

Since $\frac{T(n)}{T(n-2)} = R(n)R(n-1)$, we can use Lemma 1 to lower bound $\frac{T(k-2)}{T(k)}$ by:

$$\begin{aligned} \frac{T(k-2)}{T(k)} &= \frac{1}{R(k)R(k-2)} \\ &\geq \frac{1}{(1 + \sqrt{k+1})(1 + \sqrt{k-1})} = \Theta\left(\frac{1}{k}\right) \end{aligned}$$

thus yielding an average weight of $\Theta(\frac{xy}{k})$ over all involutions. \square

Lemma 1. Let $R(n) = 1 + \frac{n-1}{R(n-1)}$ with $R(1) = 1$; for all $n \geq 1$ it holds that:

$$\sqrt{n} \stackrel{(i)}{\leq} R(n) \stackrel{(ii)}{<} 1 + \sqrt{n+1}$$

Proof. The proof is by induction on n , with base case $\sqrt{1} \leq R(1) < 1 + \sqrt{1+1}$. From $R(n) < 1 + \sqrt{n+1}$ we conclude:

$$\begin{aligned} \sqrt{n+1} &= 1 + \frac{n}{1 + \sqrt{n+1}} \\ &\stackrel{(ii)}{<} 1 + \frac{(n+1)-1}{R(n)} = R(n+1) \end{aligned}$$

And from $\sqrt{n} \leq R(n)$ we conclude:

$$\begin{aligned} R(n+1) &= 1 + \frac{(n+1)-1}{R(n)} \\ &\stackrel{(i)}{\leq} 1 + \frac{n}{\sqrt{n}} = 1 + \sqrt{n} \\ &< 1 + \sqrt{(n+1)+1} \end{aligned}$$

\square

3.2. Adversarial Lower Bound

Theorem 3. The competitive ratio $\rho = \max_{\sigma} \frac{\text{cost}(ON(\sigma))}{\text{cost}(OFF(\sigma))}$ for PAIRWISE GRID UPDATE, with $|\sigma| = \Omega(n^{1+\frac{1}{d}})$, is at least $\Omega(\log n)$. This bound holds for arbitrarily long sequences, but if $|\sigma| = \Theta(n^{1+\frac{1}{d}})$, it even holds if the request graph is a d -grid.

To prove this, we consider an arbitrary online algorithm ON for PAIRWISE GRID UPDATE. The main idea is to have an adaptive online adversary construct a sequence σ_{ON} that depends on the algorithm ON . The adversary constructs σ_{ON} so that the resulting request graph $R(\sigma_{ON})$ is a d -grid. Because an offline algorithm knows $R(\sigma_{ON})$ in advance, it can immediately configure it at cost $O(n^{1+\frac{1}{d}})$ and serve all requests at optimal cost of 1; since $|\sigma| = \Omega(n^{1+\frac{1}{d}})$, the configuration cost of $O(n^{1+\frac{1}{d}})$ is negligible. We show that the online algorithm is forced to essentially reconfigure its layout $\log n$ times, resulting in the desired ratio. To facilitate our analysis, we use a notion of the *distortion* of an embedding reminiscent of the one used in the Minimum Linear Arrangement (MLA) (14) problem.

Definition 3. Given a request graph (V, E) with $E \subseteq V \times V$, let

$$E^+ = \{(u, v) \mid d_G(u, v) < \infty\}$$

denote the transitive closure of E .

For $h \in C_{V \rightarrow N}$, let $d_h(E)$ denote the *distortion* of E , which is defined as:


$$d_h(E) = \sum_{(u,v) \in E^+} d_N(h(u), h(v))$$

438 To build σ_{ON} , the adversary gradually commits to the
 439 edges of $R(\sigma_{ON})$. Having already requested $\sigma_1, \dots, \sigma_i$,
 440 then depending on the distortion the adversary:


441 **Option 1:** requests $\sigma_{i+1} = \arg \max_{(u,v) \in E_i} d_N(h(u), h(v))$
 442 (largest distorted previous request)

443 **Option 2:** reveals a new batch of edges $M \subset V \times V$

444 From these two options, the adversary's strategy be-
 445 comes clear; Option 1 forces the highest possible cost
 446 to ON based on E_i and h , and Option 2 introduces new
 447 communication edges to force an increase in distortion.
 448 What is left to show is how the value of $d_h(E_i)$ comes
 449 into play, and which edges the adversary reveals in Op-
 450 tion 2. The adversary reveals at most dn edges (since the
 451 final request graph is a d -dimensional grid), and the size
 452 of every subsequent batch of edges is essentially halved,
 453 resulting in $\Theta(\log n)$ batches. After each batch, for ON
 454 to remain optimal it must permute its layout at cost
 455 $\Omega(n^{1+\frac{1}{d}})$, totaling a cost of $\Omega(n^{1+\frac{1}{d}} \log n)$ for all batches
 456 combined. To ensure that $R(\sigma_{ON})$ is a subgrid, the par-
 457 tial request graph E_i (i.e., the set of revealed edges) al-
 458 ways comprises a set of disjoint subgrids. Therefore,
 459 the adversary only reveals sets of edges that concate-
 460 nate two subgrids in E_i along their length. Initially E_i is
 461 empty and the corresponding subgrids are all singleton
 462 sets of $u \in V$.



(a) Two embeddings h and h' of a set of edges. Even though both embeddings embed only a single edge suboptimally, the distortion of $d_h(E)$ is bigger than $d_{h'}(E)$ because more paths in E_i cross e_1 than e_2 .



(b) A visualization of $d_h(E_i \cup M)$: the list graph N , E_i (solid) and M (dashed) are sets of edges, configured on N by h (dotted). The sum of length of the configured edges $h((E_i \cup M)^+)$ is the distortion $d_h(E_i \cup M)$.

Figure 5: Illustrations of distortion on a 1-grid (list)

463 To help decide which edges to reveal, we use the dis-
 464 tortion to associate a cost to batches of edges that the
 465 adversary can commit to. Let $M \subseteq V \times V \setminus E_i$ be any
 466 set of edges such that the graph $(V, E_i \cup M)$ comprises
 467 a set of disjoint subgrids. For a configuration h of ON ,
 468 the set M induces a distortion of $d_h(E_i \cup M)$, as shown
 469 in Figure 5b. We show that for any embedding that ON
 470 chooses, the adversary can find a set M so that the dis-
 471 tortion is large.

472 **Lemma 2.** Let N be a d -grid graph, and $E \subseteq V \times V$ a set
 473 of edges so that the graph $G = (V, E)$ induces k disjoint
 474 subgrids of the same shape. For every $h \in C_{V \rightarrow N}$, there
 475 exists a set $M \subseteq V \times V$ such that $d_h(E \cup M) = \Omega(\frac{n^{2+\frac{1}{d}}}{k})$
 476 and $(V, E \cup M)$ contains a set of at least $k/2$ disjoint

477 subgrids of the same shape.

478 To prove this lemma, we use Theorem 2.

Lemma 2. Let $L_1, \dots, L_k \subseteq E$ be the uniform subgrids in G . For all pairs (i, j) , let (L_i, L_j) denote any set of edges so that $L_i \cup L_j \cup (L_i, L_j) = L_i \oplus L_j$ is a concatenation of L_i and L_j along their length, which is well defined since the subgrids have the same shape (if $L_i = L_j$ then $L_i \oplus L_j = \emptyset$). For any involution f on the subgrids we have:

$$2d_h(E \cup \{(L_i, L_{f(i)}) \mid i \neq f(i)\}) \geq \sum_{i=1}^k d_h(L_i \oplus L_{f(i)}). \quad (1)$$

The factor 2 is necessary because for i such that $i \neq f(i)$, the term $d_h(L_i \oplus L_{f(i)})$ appears twice in the sum. Now partition N into three subgrid slabs: a bottom slab $X = \{(n_1, \dots, n_d) \in N \mid n_1 \leq \lceil n/3 \rceil\}$, a top slab $Y = \{(n_1, \dots, n_d) \in N \mid \lfloor 2n/3 \rfloor \leq n_1\}$, and the centre slab $C = N \setminus (X \cup Y)$. Let $h_X(L_i)$ (resp. $h_Y(L_i)$) denote the number of elements of L_i that h maps onto X (resp. Y). Every two vertices u, v so that $h(u) \in X$ and $h(v) \in Y$ are by construction at least $\Theta(n^{\frac{1}{d}})$ apart in N , and therefore we can lower bound $d_h(L_i \oplus L_j)$ by:

$$d_h(L_i \oplus L_j) \geq \Theta(n^{\frac{1}{d}}) \cdot h_X(L_i)h_Y(L_j) \quad (2)$$

For an involution f drawn uniformly at random, Theorem 2 gives us a bound on the expected value of the following:

$$\mathbf{E} \left(\sum_{i=1}^k h_X(L_i)h_Y(L_{f(i)}) \right) = \Omega \left(\frac{\lceil n/3 \rceil^2}{k} \right) \quad (3)$$

Therefore, there exists an involution f for which we

have:

$$\begin{aligned} 2d_h(E \cup \bigcup_i (L_i, L_{f(i)})) &\stackrel{(1)}{\geq} \sum_{i=1}^k d_h(L_i \oplus L_{f(i)}) \\ &\stackrel{(2)}{\geq} \Theta(n^{\frac{1}{d}}) \cdot \sum_{i=1}^k h_X(L_i)h_Y(L_{f(i)}) \\ &\stackrel{(3)}{\geq} \Theta(n^{\frac{1}{d}}) \cdot \Omega(n^2/k) = \Omega \left(\frac{n^{2+\frac{1}{d}}}{k} \right) \end{aligned}$$

479 Since this holds for any choice of (L_i, L_j) , we can pick
480 them so that $(V, E \cup \bigcup_i (L_i, L_{f(i)}) \mid i \neq f(i))$ contains at
481 least $k/2$ disjoint subgrids with the same shape.

482

□

483 This lemma (and the proof) partially reveals how the ad-
484 versary commits to a new batch of edges in Option 2; it
485 can essentially pick a random matching between sub-
486 grids and concatenate them so they form bigger sub-
487 grids. For a clean scheme of building up subgrids, we
488 assume w.l.o.g. that $n = 2^{dp}$. This means that the size
489 $n^{\frac{1}{d}}$ of the grid is 2^p , meaning that the adversary can
490 perfectly build up hypercube-shaped subgrids of size
491 $2^1, 2^2, \dots, 2^p$. However, a hypercube comprises 2^d half-
492 sized hypercubes, whereas Lemma 2 matches *pairs* of
493 subgrids. Since a hypercube needs to be cut in d di-
494 mensions to yield said half-sized hypercubes, we can
495 conversely also build up the hypercubes of size 2^{i+1} by
496 d rounds of concatenating subgrids composed of hyper-
497 cubes of size 2^i .

498 Next we show the precondition for the adversary to opt
499 for Option 1, including a lower bound on the corre-
500 sponding cost imposed on ON .

Lemma 3. Let N be a d -grid, $h \in C_{V \rightarrow N}$ a config-
501 uration, and $E \subseteq V \times V$ a set of edges so that the
502 graph $G = (V, E)$ has n/ℓ disjoint subgrids of size ℓ .
503

504 If $d_h(E) = \Omega(\ell n^2)$, then there exists an edge $(u, v) \in E$
 505 such that $d_h(u, v) = \Omega(n/\ell)$.

506 *Proof.* There are $n/\ell \cdot \binom{\ell}{2} = O(\ell n)$ distinct connected
 507 pairs of vertices in G , meaning that the average distort-
 508 tion of the shortest path between each pair is $\frac{\Omega(\ell n^2)}{O(\ell n)} =$
 509 $\Omega(n)$. The highest distortion is at least the average, and
 510 every path in G has length at most ℓ . On this path, there
 511 must exist an edge with distortion $\Omega(n/\ell)$, since if all
 512 edges have a distortion of $o(n/\ell)$, the total would be
 513 $o(n)$. \square

514 Combined, Lemma 2 and Lemma 3 imply that the ad-
 515 versary can either request an edge at cost $\Omega(n/\ell)$, or in-
 516 crease the distortion to $\Omega(\ell n^2)$ by revealing a new batch
 517 of edges. The final ingredient is a lower bound on how
 518 much cost the adversary can impose on ON in between
 519 these batches.

520 **Lemma 4.** Let N be a d -grid, and $E \subset V \times V$ a set
 521 of communication edges forming disjoint subgrids. If
 522 $h, h' \in C_{V \leftrightarrow N}$ are two embeddings that differ only in the
 523 order of two adjacent elements u and v , then $d_h(E) \leq$
 524 $d_{h'}(E) + 2\ell$, where ℓ is the size of the largest sublist in E .

525 *Proof.* Consider all shortest paths in E that end in u . At
 526 most ℓ paths ending in u (or v) are reduced by 1, and
 527 therefore $d_h(E) - d_{h'}(E) \leq 2\ell$. \square

528 Combining the previous lemmata, we can prove the
 529 main technical result.

530 **Lemma 5.** For every online algorithm A , there
 531 is a sequence σ_{ON} of length $\Theta(n^{1+\varepsilon})$ such that
 532 $\text{cost}(ON(\sigma_{ON})) = \Omega(\varepsilon n^{1+\frac{1}{d}} \log n)$, for $0 < \varepsilon \leq 1$. Fur-
 533 thermore, the resulting request graph $R(\sigma_{ON})$ is a sub-
 534 grid.

535 *Proof.* W.l.o.g. assume that $n = 2^{dp}$ for some integer p
 536 so that the hypercube concatenation scheme described
 537 earlier can be employed.

538 Consider the situation right after a batch of edges is
 539 revealed, where all subgrids have cardinality ℓ . By
 540 Lemma 2 this implies that the distortion is $\Omega(\ell n^2)$. Let
 541 $\sigma = \sigma_i, \sigma_{i+1}, \dots, \sigma_{i+\ell n}$ be the requests obtained by re-
 542 peatedly requesting the edge in E_i with largest distort-
 543 tion. There are two situations:

- Throughout serving σ , the distortion is always at
 544 least $\Omega(\ell n^2)$. Then by Lemma 3 each σ_j , $i \leq j \leq$
 545 $i + \ell n$ incurred a cost of $\Omega(n/\ell)$, at total cost $\Omega(n^2)$.
- By serving σ , ON halves the distortion, thus reduc-
 546 ing it by at least $\Omega(\ell n^2)$. Then, since by Lemma 4
 547 every swap reduces the distortion by at most 2ℓ ,
 548 ON must have used at least $\Omega(n^2)$ swaps.

549 This argument holds for each batch of edges revealed.
 550 The adversary stops when the subgrids have cardinality
 $2^{\varepsilon \log n}$, yielding a sequence σ_{ON} with

$$|\sigma_{ON}| = \sum_{\ell \in \{2^0, \dots, 2^{\varepsilon \log n}\}} \ell n = \Theta(n^{1+\varepsilon})$$

551 and $\text{cost}(\sigma_{ON}) = \Omega(\varepsilon n^{1+\varepsilon} \log n)$. By Lemma 3, the ad-
 552 versary only requests edges that are introduced using the
 553 matching from Lemma 2. Any set of edges introduced
 554 by the latter Lemma concatenates two already existing
 555 subgrids, hence $R(\sigma_{ON})$ is a subgrid. \square

556 To wrap up the proof for Theorem 3, we conclude by
 557 showing that for any online algorithm ON , the sequence
 558 σ_{ON} can be solved in $\Theta(n^{1+\frac{1}{d}})$ by an optimal offline al-
 559 gorithm.

Proof of Theorem 3. Let ON be any online algorithm
 solving PAIRWISE GRID UPDATE. Apply Lemma 5 with

$\varepsilon = 1/d$, yielding $\text{cost}(ON(\sigma_{ON})) = \Omega(n^{1+\frac{1}{d}} \log n)$. Since σ_{ON} is a d -grid, an offline algorithm can embed this graph at (worst case optimal) cost $\Theta(n^{1+\frac{1}{d}})$, and serve every request at optimal cost $O(1)$. This yields $\text{cost}(OFF(\sigma_{ON})) = \Theta(n^{1+\frac{1}{d}})$, and thus

$$\rho = \frac{\text{cost}(ON(\sigma))}{\text{cost}(OFF(\sigma))} = \Omega(\log n)$$

560 In order to make this bound hold for arbitrary long se-
 561 quences, we slightly modify the adversary. After every
 562 $\Theta(n^{1+\frac{1}{d}})$ requests it serves, it can reconfigure to a new
 563 grid at cost $O(n^{1+\frac{1}{d}})$, and repeat the argument to force
 564 cost of $\Omega(n^{1+\frac{1}{d}} \log n)$ to ON for the subsequent $\Theta(n^{1+\frac{1}{d}})$
 565 requests. \square

566 4. Online Algorithms and Complexity

567 While our main contribution is the derived lower bound,
 568 we in this section also initiate the study of upper
 569 bounds, shedding light on the underlying algorithmic
 570 challenges. We present a $O(\log n)$ -competitive online
 571 algorithm **GREAD** for requests issued according to a 1-
 572 dimensional grid (i.e., a linear order), and discuss its
 573 distributed implementation. We also prove that the of-
 574 fline variant of **PAIRWISE LIST UPDATE**, the 1-dimensional
 575 grid update problem, is already NP-complete.

576 4.1. An Upper Bound

This section presents a $O(\log n)$ -competitive online algorithm for **PAIRWISE LIST UPDATE**. Our main technical lemma shows that the total cost spent on learning the optimal embedding never exceeds $O(n^2 \log n)$. We propose a simple greedy algorithm that identifies a *locally* 589 *optimal* embedding, and always moves towards this em- 590 bedding. Observe that a set of k sublists can be embed- 591 ded perfectly on a line graph in at most $2^k k!$ ways (they 592

are permuted in some order, and every list has at most two orientations). Given a configuration $h \in C_{V \rightarrow N}$ of the lists, we define the locally optimal embedding to be an optimal embedding one that takes the fewest number of reconfigurations to reach, starting at h . Formally, if $\text{opt}(E)$ is the set of optimal embeddings of a set edges, then the h -optimal embedding of E is

$$h[E] = \arg \min_{h' \in \text{opt}(E)} \sum_{v \in V} |h(v) - h'(v)|$$

With such a configuration we associate the cost:

$$\Phi_h[E] = \sum_{v \in V} |h(v) - h[E](v)|$$

577 Let **GREAD** be the algorithm (it **GRE**edily **AD**joins sub-
 578 lists), that upon seeing a new edge σ_i , *immediately*
 579 moves to the embedding $h[E_i \cup \{\sigma_{i+1}\}]$.

580 For each E_i , let $\mathcal{V}(E_i)$ be the connected components
 581 of (V, E_i) , so that $\mathcal{V}_\sigma = \cup_{1 \leq i \leq m} \mathcal{V}(E_i)$ is the set of all
 582 sublists induced by σ . This naturally defines a binary
 583 tree $T_\sigma = (\mathcal{V}_\sigma, E_\sigma)$: for every first occurrence σ_i of
 584 $(u, w) \in E_m$ connecting two sublists U, W in $R(E_i)$, there
 585 are two corresponding edges $(U, U \cup W)$ and $(W, U \cup W)$
 586 in E_σ . For every $\sigma_i \in E_m$, **GREAD** incurs some cost
 587 for reconfiguring, and the following lemma bounds this
 588 cost.

Lemma 6. Let h be an optimal embedding of E_i , and let σ_{i+1} be an edge connecting two sublists U and W of E_i . It holds that

$$\Phi_h[E_i \cup \{\sigma_{i+1}\}] \leq n \cdot \min(|U|, |W|)$$

Proof. Since E_i is optimally embedded by h , we simply need to move the smaller of U and W into its correct location so that $E_i \cup \{\sigma_{i+1}\}$ is optimally embedded. This requires every element in the smaller list to be

593 moved at most n locations, therefore $\Phi_h[E_i \cup \{\sigma_{i+1}\}] \leq$
 594 $n \min(|U|, |W|)$. \square

For a node $U \in \mathcal{V}_\sigma$, let $\text{left}(U)$ and $\text{right}(U)$ denote
 U 's left and right child respectively. Further, let $w(U)$
 denote the number of nodes in the subtree rooted at U .
 Observe that for any binary tree with nodes N , it holds
 that

$$\sum_{v \in N} \min(w(\text{left}(v)), w(\text{right}(v))) \leq |N| \log |N|$$

Theorem 4. For any σ , with $|\sigma| = m$, such that $|E_m| = k$
 and $R(\sigma)$ is a list,

$$\text{cost}(\text{GREAD}(\sigma)) = O(m + nk \log k)$$

Proof. Let h_i denote the configuration after request σ_1 ,
 and let h_0 denote the trivial optimal initial embedding.
 Then the total cost of GREAD is the sum of reconfiguring
 after every σ_i plus accessing every request at cost 1:

$$\begin{aligned} \text{cost}(\text{GREAD}(\sigma)) - m &= \sum_{i=0}^m \Phi_{h_i}[E_i \cup \{\sigma_{i+1}\}] \\ &\leq \sum_{U \in \mathcal{V}_\sigma} n \min(w(\text{left}(U)), w(\text{right}(U))) \\ &\leq nk \log k \end{aligned}$$

595 \square

596 To give an example of GREAD, we can consider the se-
 597 quence derived for our lower bound in Section 3. When
 598 applied to this sequence, after each batch of revealed
 599 edges, GREAD spends $O(n^2)$ reconfiguring to perfectly
 600 embed the current demand, yielding a total $O(n^2 \log n)$.
 601 This also means that for this sequence, GREAD is $\log n$
 602 competitive.

4.2. Distributed Implementation of GREAD

604 To execute GREAD in a distributed environment, we have
 605 to address several aspects: i) how to perform (local)
 606 routing, ii) how to ensure nodes know to which (tem-
 607 porary) sublist they belong, together with the size of the
 608 sublist, and iii) how to perform the reconfiguration and
 609 merging of two sublists. We address these issues one by
 610 one.

Routing: The basic problem with routing is that the
 611 source nodes do not know the location of the destina-
 612 tion, since initially there is no *sense of direction*. To
 613 overcome this problem each source initiates an expo-
 614 nential search on *both* sides of the line network when
 615 it first needs to communicate with a destination. This
 616 will guarantee that the cost of the *first* route request will
 617 be $O(i)$ for a destination that is i hops away on the line
 618 network. Note that this is proportional to the cost of any
 619 algorithm. According to GREAD the cost of all future re-
 620 quests will be 1 since after the first communication re-
 621 quest the source and destination are reconfigured to be
 622 neighbors.
 623

Sublist: During the execution each node maintains the
 624 following information: A bit that indicates if it is at the
 625 *end* of a sublist (a node is at the end of a sublist if it has
 626 less than two neighbors from that sublist). If it is at the
 627 end of the list then the node maintains the size of the list
 628 (up to $\log n$ bits).
 629

Reconfiguration: Basically GREAD merges two sublists
 630 by swapping the shorter list toward the longer sublist.
 631 Note that this happens only on the *first* routing request
 632 from a source to destination. This can be done in a dis-
 633 tributed manner in the following way. On the first rout-
 634 ing request, the source (which must be an *end* node)
 635

636 attaches the size of its sublist to the message. The des- 670
 637 tination (which also must be an *end* node), upon receiv- 671
 638 ing the request, answers to the source with the size of 672
 639 its own sublist (initially set to one). It is then clear to 673
 640 both the source and destination which sublist needs to 674
 641 move toward which sublist and what will be the size 675
 642 of the merged sublist. Then, both source and destina- 676
 643 tion send messages within their sublist informing the
 644 other *ends* of the sublist of the size of the merged list.
 645 Now, w.l.o.g., assume the destination needs to move to-
 646 ward the source. The destination then starts performing
 647 swaps (with its neighbor that is not on its current list)
 648 toward the source. This process ends when both the
 649 destination is a neighbor of the source and the source
 650 is a neighbor of its previous neighbor on its list. Before
 651 starting the swaps the destination informs its neighbor
 652 (which in turn informs its neighbor and so on) to *follow*
 653 *up* after it with similar swaps. It can be observed that af-
 654 ter this process the two list will be merged into a larger
 655 list and both *ends* will know the sizes of the new sub-
 656 list. The cost of the reconfiguration is $O(n \min(|U|, |W|))$
 657 where U and W are the two sublists involved in the
 658 merging.

659 4.3. NP-Completeness of the Offline Variant

660 We present a reduction from the classic MINIMUM LINEAR 694
 661 ARRANGEMENT (MLA) problem to the offline variant of 695
 662 PAIRWISE LIST UPDATE. Throughout this section, we refer
 663 to the offline PAIRWISE LIST UPDATE as DLU. The MLA 696
 664 problem can be seen as a static variant of the DLU prob- 697
 665 lem, where we choose a static network configuration at 698
 666 the beginning for free, and we serve all requests without 699
 667 further reconfiguration. 700

668 In (17), the authors considered the ITINERANT LIST UP-
 669 DATE problem, which is closely related to DLU. They 702

suggested that MLA is equivalent to ITINERANT LIST UP-
 DATE, where many identical copies of the MLA graph
 arrive. In this section, we highlight the technical chal-
 lenge for showing the equivalency to MLA (that is valid
 for both DLU and ITINERANT LIST UPDATE). Then, we in-
 troduce necessary modifications to this idea to show the
 NP-completeness of DLU.

677 The idea of repeating multiple MLA graphs faces the
 678 following technical challenge. If we are not diligent
 679 with the request order, the cost of DLU may decrease
 680 in comparison to MLA by interleaving requests with
 681 reconfigurations. To see that, consider the example of
 682 transforming of the MLA instance in form of a star with
 683 a central vertex c and 5 other vertices to the instance of
 684 DLU. Two optimal MLA solutions of cost 9 exist: plac-
 685 ing c at any of two graph centers (vertices at positions 3
 686 and 4). Assume c is placed at position 3, and in the
 687 chosen edge order, the requests to the nodes at 1 and 2
 688 precede the requests to the nodes at 5 and 6. Then,
 689 this ordering enables the optimal DLU solution to ob-
 690 tain a lower cost than the optimal MLA by performing
 691 a single reconfiguration. Moving c to position 4 after
 692 serving requests to nodes at 1 and 2 decreases the dis-
 693 tance to both nodes at 5 and 6 at the cost of a single
 694 reconfiguration. The optimal DLU cost is then 8, which
 695 is smaller than the optimal MLA cost 9.

696 We show a reduction from MLA to DLU that mitigates
 697 the problem of edge ordering entirely. To this end, we
 698 map any vertex of MLA to a group of DLU nodes that
 699 stay adjacent due to a large number of inter-group re-
 700 quests. We map the MLA edges to requests between the
 701 central nodes of the corresponding groups.

702 We repeat the MLA request pattern several times, called

703 rounds. We show that there exists a low-cost round
 704 where groups stay in a fixed order — and we recon-
 705 struct the solution to MLA from it. To determine the
 706 sufficient number of rounds, we introduce the following
 707 helper lemma.

708 **Lemma 7.** Fix any integers A and D . Consider a se-
 709 quence of non-negative integers with the average upper-
 710 bounded by A . We mark at most D elements of the
 711 sequence as *defective*. If an element is defective, it
 712 equals 0. Then the following holds: if length of the
 713 sequence is at least $D(A + 1) + 1$, then a non-defective
 714 element no larger than A exists.

715 *Proof.* Assume the opposite, i.e., that the sequence of
 716 length n , where $n \geq D \cdot (A + 1) + 1$ consists only of
 717 defective elements and elements larger than A . Then,
 718 the sum of the sequence is at least $(n - D) \cdot (A + 1)$ for
 719 $D \leq (n - 1)/(A + 1)$. This sum is at least $A \cdot n + 1$, thus the
 720 average is strictly larger than A , a contradiction. \square

721 **Theorem 5.** Offline DLU is NP-complete.

722 *Proof.* We reduce MLA to (offline) DLU. Given any in-
 723 teger k and any instance I_{MLA} of MLA, we construct an
 724 instance I_{DLU} of DLU. We show that there exists a so-
 725 lution to I_{MLA} with cost at most k if and only if there
 726 exists a solution to I_{DLU} with cost at most Thr , where
 727 Thr is bounded by a polynomial of k and $|I_{MLA}|$. First,
 728 we describe the construction, and then we prove (\Rightarrow) ,
 729 followed by (\Leftarrow) .

730 In the following, we describe the construction of I_{DLU} .
 731 Fix an integer k and an instance $I_{MLA} = \langle V, E \rangle$. We
 732 denote $n = |V|$ and $m = |E|$. The construction consists
 733 of the following parts.

Groups of vertices. We order V arbitrarily, and for each
 734 $v \in V$ we produce a group of adjacent g nodes, where g
 735 is the smallest odd number that exceeds k^3 . We denote
 736 the sequence of nodes of the group by $G(v)$, and we dis-
 737 tinguish its $(g/2 + 1)$ -th (central) node by $c(v)$. In total
 738 we construct $n \cdot g$ nodes.

Rounds of requests. We produce the requests in $R :=$
 739 $\lfloor (n^2 \cdot (k + 1) + 1)/(1 - 1/k) \rfloor$ rounds, each consisting
 740 of identical requests. We elaborate on the choice for R
 741 later in this proof. In each round, for each edge of MLA
 742 we produce a single inter-group request that we follow
 743 with a sequence of intra-group requests.

Inter-group requests. In each round, we order MLA
 744 edges arbitrarily, and for each edge $\{u, v\} \in E$ we pro-
 745 duce a request $\langle c(u), c(v) \rangle$. After each such request, we
 746 issue a sequence of intra-group requests.

Intra-group requests. After each inter-group request
 747 in each round, we produce $S := Rk(g - 1) + n^2g^2 + 1$
 748 requests between each pair of adjacent vertices
 749 $\langle G(v)_{(i)}, G(v)_{(i+1)} \rangle$ in each group. These requests enforce
 750 the preservation of the structure of each group: the adja-
 751 cency of consecutive nodes and the position of the cen-
 752 tral node. (The value S is equal to $Thr_I + Thr_E + 1$ that
 753 we introduce in a moment.)

The feasibility threshold. Recall that we perform
 754 a reduction between two decision problems. The cost
 755 threshold Thr for a solution to DLU consists of three
 756 parts: $Thr := Thr_I + Thr_E + Thr_S$.

- 757 1. The initial reconfiguration budget is $Thr_I := n^2g^2$,
 758 and it allows for a full reconfiguration of all nodes.
- 759 2. The budget for serving inter-group requests is
 760 $Thr_E := Rk(g - 1)$, and it equals to k times the dis-
 761 tance between the centers of the groups, multiplied

767 by the number of rounds R .
768 3. The budget for intra-group requests is $Thr_S :=$
769 $RSn(g - 1)$, and it allows for S requests between
770 adjacent nodes of each group at distance 1, multi-
771 plied by the number of rounds R .

772 First, we show the implication (\Rightarrow). We take any solu-
773 tion S_{MLA} to MLA with cost at most k , and construct the
774 DLU instance as follows. Before we serve any request,
775 we reconfigure groups of nodes according to the vertex
776 arrangement in S_{MLA} , and we serve all requests in this
777 configuration. The reconfiguration cost is bounded by
778 Thr_I , as we perform one reconfiguration of all nodes.
779 The structure of groups is preserved (consecutive nodes
780 of each group are adjacent), hence inter-group requests
781 cost exactly Thr_S . If MLA incurs the cost c_e for an
782 edge e , then for it we incur the cost $c_e \cdot (g - 1)$ in each of
783 R rounds for the corresponding request. As the cost of
784 MLA edges sums to k , the cost of inter-group requests
785 sum to Thr_E . In total, the constructed solution has cost
786 Thr , and the claim holds.

787 Next, we show the implication (\Leftarrow). Fix a solution to I
788 of cost at most Thr . We say that a configuration is
789 *well-aligned* if each pair of consecutive nodes of each
790 group is adjacent. We say that a round is *defective*, if
791 a group exchange happens during this round. We say
792 that a round is *excessive*, if the cost of inter-cluster re-
793 quests exceeds $k(g - 1)$ during this round.

794 In the following, we express the necessary conditions
795 for the existence of non-defective and non-excessive
796 round. Then, we construct a solution to MLA from the
797 group order in this round, and show that its cost is at
798 most k .

799 We claim that the solution reaches a well-aligned con-
800 figuration at least once during each sequence of inter-
801 group requests. To see that, note that the minimum cost
802 of serving inter-group requests throughout all rounds is
803 Thr_S , and this is achievable by serving all requests over
804 one hop. If the solution would not reach a well-aligned
805 configuration, a pair of nodes requested S times is not
806 adjacent, and this additionally incurs the cost at least
807 $S = Thr_I + Thr_E + 1$, and the solution exceeds the cost
808 Thr already.

809 This allows to reason about the cost of serving any intra-
810 cluster request. Consider a round where no groups ex-
811 change positions. For this fixed order of the groups, by
812 $p(u)$ we denote the index of the group u in this order.
813 Then, we claim that each request $\langle c(u), c(v) \rangle$ costs at
814 least $|p(u) - p(v)|(g - 1)$. If the cost would be lower, then
815 the distance between $c(u)$ and $c(v)$ must be decreased by
816 migrations. However, this means these must move back
817 to guarantee the well-aligned position for the upcoming
818 inter-group requests.

819 Some rounds may be defective, i.e., the solution may
820 move to a configuration with a different group order dur-
821 ing that round. This requires moving an entire group,
822 and costs at least g^2 . The defective rounds may be al-
823 lowed by exchanges paid from budgets Thr_I and Thr_E .
824 Note that the budget Thr_S is tight: the cost of inter-
825 group requests is at least Thr_S .

826 We claim that at most $n^2 + R/k^2$ rounds may be defective
827 (note that this is dependant on the number of rounds). At
828 most n^2 of these can be paid from Thr_I . This can happen
829 if the solution does not reconfigure groups at the begin-
830 ning, but defers some of the exchanges to later rounds.
831 Additionally, some exchanges may be paid from Thr_E :
832 for each k^2 rounds, the budget Thr_E allows for an addi-

833 tional exchange.

834 Now, we lower-bound the number of rounds R suffi-
835 cient to guarantee the existence of non-defective, non-
836 excessive round. Note that on average, the cost of intra-
837 group requests in each round is bounded by $k(g - 1)$.
838 We claim no cost for intra-group requests in defective
839 rounds. Note that in a non-defective round, each intra-
840 group request cost is divisible by $g - 1$. Hence, to sim-
841 plify calculations, we divide the cost of each round and
842 the average cost by $g - 1$.

843 To justify the choice of R , we use Lemma 4.3. We ap-
844 ply it to the sequence of (scaled) costs of rounds, and
845 conclude that it is sufficient that the length of the round
846 sequence (R) satisfies the inequality $R \geq (n^2 + R/k^2) \cdot$
847 $(k + 1) + 1$. We note that our choice of R satisfies this
848 criterion. This inequality also justifies the choice of
849 $g = \Omega(k^3)$. Crucially, the rate of growth of the budget
850 Thr_E must allow for an extra reconfiguration no more
851 frequently than every k^2 rounds. Otherwise, no positive
852 R would satisfy the inequality.

853 The solution to MLA reconstructed from the node order
854 in a non-defective, non-excessive round has the cost at
855 most k . The reconstruction is correct, as the round is
856 non-defective, and has cost at most k as the round is
857 non-excessive.

858 The reduction is polynomial: the number of nodes and
859 requests and the value Thr are polynomials of m, n
860 and k . Finding the round without reconfigurations is
861 linear in terms of the number of requests. \square

862 5. Conclusion

863 We presented a lower bound of $\Omega(\log n)$ on the online
864 competitiveness of PAIRWISE GRID UPDATE, even under

865 very restricted and seemingly simple communication
866 patterns. We further initiated the discussion of online
867 algorithms for such restricted scenarios and presented
868 a deterministic $\Theta(\log n)$ -competitive algorithm. We be-
869 lieve that our work opens several interesting directions
870 for future research. In particular, it would be interest-
871 ing to shed light on the competitive ratio achievable in
872 more general network topologies and request patterns.
873 It would further be interesting to study randomized al-
874 gorithms and to generalize our cost model, for example
875 also accounting for congestion aspects.

876 **Acknowledgments.** Research supported by the Euro-
877 pean Research Council (ERC), consolidator grant Ad-
878 justNet (grant agreement No. 864228).

879 References

- 880 [1] S. Albers. A competitive analysis of the list update problem
881 with lookahead. *Theoretical Computer Science*, 197(1-2):95–
882 109, 1998.
- 883 [2] S. Albers. Improved randomized on-line algorithms for the list
884 update problem. *SIAM Journal on Computing*, 27(3):682–693,
885 1998.
- 886 [3] S. Albers, B. Von Stengel, and R. Werchner. A combined bit and
887 timestamp algorithm for the list update problem. *Information*
888 *Processing Letters*, 56(3):135–139, 1995.
- 889 [4] S. Albers and J. Westbrook. Self-organizing data structures. In
890 *Online algorithms*, pages 13–51. Springer, 1998.
- 891 [5] C. Ambühl. Offline list update is np-hard. In *European Sym-
892 posium on Algorithms*, pages 42–51. Springer, 2000.
- 893 [6] C. Avin, M. Borokhovich, B. Haeupler, and Z. Lotker. Self-
894 adjusting grid networks to minimize expected path length. In
895 *International Colloquium on Structural Information and Com-
896 munication Complexity*, pages 36–54. Springer, 2013.
- 897 [7] C. Avin, O. Dunay, and S. Schmid. Strategies for traffic-aware
898 vm migration. In *Proc. 6th IEEE/ACM International Conference*
899 *on Utility and Cloud Computing (UCC)*, December 2013.
- 900 [8] C. Avin, A. Hercules, A. Loukas, and S. Schmid. Towards
901 communication-aware robust topologies. *ArXiv Technical Re-
902 port*, 2017.

- 903 [9] C. Avin, A. Loukas, M. Pacut, and S. Schmid. Online balanced
904 repartitioning. In *International Symposium on Distributed Com-*
905 *puting*, pages 243–256. Springer, 2016.
- 906 [10] C. Avin, K. Mondal, and S. Schmid. Demand-aware network
907 design with minimal congestion and route lengths. In *Proc. IEE*
908 *INFOCOM*, 2019.
- 909 [11] C. Avin and S. Schmid. Toward demand-aware networking: A
910 theory for self-adjusting networks. In *ACM SIGCOMM Com-*
911 *puter Communication Review (CCR)*, 2018.
- 912 [12] C. Avin, I. van Duijn, and S. Schmid. Self-adjusting linear net-
913 works. In *Proc. 21st International Symposium on Stabilization,*
914 *Safety, and Security of Distributed Systems (SSS)*, 2019.
- 915 [13] S. Bubeck, N. Cesa-Bianchi, et al. Regret analysis of stochastic
916 and nonstochastic multi-armed bandit problems. *Foundations*
917 *and Trends® in Machine Learning*, 5(1):1–122, 2012.
- 918 [14] J. Díaz, J. Petit, and M. Serna. A survey of graph layout prob-
919 lems. *ACM Computing Surveys (CSUR)*, 34(3):313–356, 2002.
- 920 [15] S. Huq and S. Ghosh. Locally self-adjusting skip graphs. In
921 *Proc. IEEE 37th International Conference on Distributed Com-*
922 *puting Systems (ICDCS)*, pages 805–815, 2017.
- 923 [16] M. Ghobadi et al. Projector: Agile reconfigurable data center
924 interconnect. In *Proc. ACM SIGCOMM*, pages 216–229, 2016.
- 925 [17] N. Olver, K. Pruhs, K. Schewior, R. Sitters, and L. Stougie. The
926 itinerant list update problem. In *13th Workshop on Models and*
927 *Algorithms for Planning and Scheduling Problems*, page 163,
928 2017.
- 929 [18] B. Peres, O. Souza, O. Goussevskaia, S. Schmid, and C. Avin.
930 Distributed self-adjusting tree networks. In *Proc. IEE INFO-*
931 *COM*, 2019.
- 932 [19] N. Reingold and J. Westbrook. Off-line algorithms for the list
933 update problem. *Information Processing Letters*, 60(2):75–80,
934 1996.
- 935 [20] N. Reingold, J. Westbrook, and D. D. Sleator. Randomized com-
936 petitive algorithms for the list update problem. *Algorithmica*,
937 11(1):15–32, 1994.
- 938 [21] S. Schmid, C. Avin, C. Scheideler, M. Borokhovich, B. Hae-
939 upler, and Z. Lotker. Splaynet: Towards locally self-adjusting
940 networks. *IEEE/ACM Transactions on Networking (ToN)*, 2016.
- 941 [22] D. D. Sleator and R. E. Tarjan. Amortized efficiency of list up-
942 date and paging rules. *Communications of the ACM*, 28(2):202–
943 208, 1985.
- 944 [23] B. Teia. A lower bound for randomized list update algorithms.
945 *Information Processing Letters*, 47(1):5–9, 1993.