



Reconfigurable routing in data center networks [☆]

David C. Kutner ^{ID}*, Iain A. Stewart

Department of Computer Science, Durham University, Upper Mountjoy Campus, Stockton Road, Durham DH1 3LE, UK

ARTICLE INFO

Keywords:

Algorithms
Complexity
Reconfigurable topologies
Optical circuit switches
Software-defined networking

ABSTRACT

A hybrid network is a static (electronic) network that is augmented with optical switches. The Reconfigurable Routing Problem (RRP) in hybrid networks is the problem of finding settings for the optical switches augmenting a static network so as to achieve optimal delivery of some given workload. The problem has previously been studied in various scenarios with both tractability and NP-hardness results obtained. However, the data center and interconnection networks to which the problem is most relevant are almost always such that the static network is highly structured (and often node-symmetric) whereas all previous results assume that the static network can be arbitrary (which makes existing computational hardness results less technologically relevant and also easier to obtain). In this paper, and for the first time, we prove various intractability results for RRP where the underlying static network is highly structured, for example consisting of a hypercube, and also extend some existing tractability results.

1. Introduction

The rapid growth of cloud computing applications has induced demand for new technologies to optimize the performance of data center networks dealing with ever-larger workloads. The data center topology design problem (that of finding efficient data center topologies) has been studied extensively and resulted in myriad designs (see, e.g., [7]). Advances in hardware, such as optical switches reconfigurable in milli- to micro-seconds, have enabled the development of reconfigurable topologies (see, e.g., [19]). These topologies can adjust in response to demand (*demand-aware* reconfigurable topologies) or vary configurations over time according to a fixed protocol (*demand-oblivious* reconfigurable topologies; see, e.g., [4]). So-called *hybrid* data center networks are a combination of a static topology consisting of, for example, electrical switches, and a demand-aware reconfigurable topology implemented, for example, with optical circuit switches or free space optics (see, e.g., [6,14,20,25]). An intuitive example of a simple reconfigurable topology is illustrated in Fig. 1.

The hybrid network paradigm combines the robustness guarantees of static networks with the ability of demand-aware reconfigurable networks to serve large workloads at very low cost. Consider, for example, the hybrid network shown in Fig. 2, and the configuration shown in Fig. 3. In the (unaugmented) static network, there are two possible paths along which a message from node b to node d may be routed: $b \rightarrow f \rightarrow h \rightarrow e \rightarrow d$ or $b \rightarrow f \rightarrow e \rightarrow d$. In the hybrid network as configured in Fig. 3, the path $b \rightarrow a \rightarrow c \rightarrow d$ (among others) is an option.¹

[☆] This article belongs to Section A: Algorithms, automata, complexity and games, Edited by Paul Spirakis.

* Corresponding author.

E-mail addresses: david.c.kutner@durham.ac.uk (D.C. Kutner), i.a.stewart@durham.ac.uk (I.A. Stewart).

¹ We denote by $u \rightarrow v$ the concatenation of a switch link from u to some switch, of the internal switch connection, and of a switch link to v from that switch.




Fig. 1. Basic model of an optical wireless data-center network, as described in [6,20,25]. Practical timescales for reconfiguration vary from milliseconds [20] to microseconds or nanoseconds [6,25].




Fig. 2. A hybrid network.




Fig. 3. An augmented network and its abstracted dynamic links.

Of particular interest to us is the question of how the reconfigurable (optical) portion of the network should be configured for some demand pattern, formalized by Foerster, Ghobadi and Schmid [12] as the RECONFIGURABLE ROUTING PROBLEM (RRP): in short, given a hybrid network (consisting of a static network and of some switches) and a workload, we wish to choose a *configuration* (setting of the switches) which results in an optimal delivery of the workload.

Crucially, existing hardness results are only valid when the static network is allowed to be arbitrary, which is almost never the case in practice where interconnection and data center network design is driven by symmetry, high connectivity, recursive decomposition, and so forth. For example: the popular switch-centric data center network Fat-Tree [3] is derived from a folded Clos network; the server-centric data center network DCell [17] is recursively-structured whereby at each level, a graph-theoretic matching of servers is imposed; and the server-centric data center network BCube [16] is recursively-structured with a construction based around a generalized hypercube. (It should be noted that there do exist examples of unstructured data center networks, such as Jellyfish [23] and Xpander [24] which utilize the theory of random graphs.) Many (but not all) NP-complete problems become tractable when the input is restricted to the graphs providing the communications fabric for data center networks and other interconnection networks. For example, Hamiltonian paths are often trivial to find in many interconnection networks; indeed, no finite connected vertex-transitive graph *without* a Hamiltonian path is known to exist (the Lovász Conjecture contends there is no such graph - see Section 4 of [22]). This motivates our investigation into how the complexity of RRP changes when we restrict to more structured and realistic networks. The question of the complexity of RRP for specific network topologies was specifically identified as an area for future work in [11].

In this paper, we establish for the first time hardness results for RRP that apply to various specific families of highly structured static networks such as, for example, the hypercubes. Our constructions are (perhaps not surprisingly) of a much more involved nature than has hitherto been the case.

2. Problem setting

The decision problem RECONFIGURABLE ROUTING PROBLEM considered in this paper is a proper restriction of that presented in prior work [11–13]. In this section, we provide technical detail to fully formalize our version of the problem, but also additionally provide sufficient framing to briefly review existing results and to identify the areas strengthened by our contribution.

We adopt the usual terminology of graph theory though we tend to use ‘nodes’ and ‘links’ when speaking about the components of reconfigurable networks and ‘nodes’ and ‘edges’ when dealing with (abstract) graphs. We denote the natural numbers by $\mathbb{N}$ (we include $0 \in \mathbb{N}$) and the non-negative rationals by $\mathbb{Q}_+$.

2.1. Hybrid networks, (re)configurations and (segregated) routing

A hybrid network $G(S)$ can be visualized as in Fig. 2, and consists of a static network G and some switches S augmenting it. A static network G can be abstracted as an undirected graph $G = (V, E)$ so that each static link $(u, v) \in E$ has some fixed weight $w \in \mathbb{Q}_+$ (reflecting a transmission cost) and is incident with internal ports of two distinct nodes of V . The number of internal ports of some node $v \in V$ is then exactly the degree of v in the abstracted graph G . We denote by S a set of switches augmenting the static network G with switch links joining switch ports of some switch to external ports of some of the nodes of V . Every switch link has weight 0 (we say more about switch link weights momentarily). Every switch $s \in S$ has at least two switch ports.

In general, the number of external ports of the nodes of a static network $G = (V, E)$ is variable, as is the number of switch ports of the switches of a hybrid network $G(S)$, and it may be the case that there is more than one switch link between a specific node and a specific switch. We assume that the switch links describe a bijection between the external ports and the switch ports; otherwise, there would be some unused ports, which we can safely ignore.

Given a hybrid network $G(S)$ and a switch $s \in S$, a switch matching N_s of s is a set of pairs of switch ports of s so that all switch ports involved are distinct. Each switch matching represents an internal setting of the switch and naturally yields a set of pairs of external ports of nodes where all such ports are distinct; we refer to a set of pairs of external ports obtained in this way as a node matching (note that this differs from the standard graph-theoretic notion of a matching). An illustration of a configured hybrid network is shown in Fig. 3: on the left side, switch matchings are represented as sets of arcs, and on the right side the corresponding node matching is shown as a set of dotted lines.

A configuration N is a set of switch matchings, one for each switch. A configuration straightforwardly encodes the corresponding node matchings. We say that (u, v) is a dynamic link in the configuration N (we sometimes write $(u, v) \in N$) if (u, v) appears in any node matching corresponding to N .

We allocate a fixed weight $\mu \in \mathbb{Q}_+$ to each internal port-to-port connection in a switch s . Although a dynamic link is an atomic entity, it can be visualized as consisting of a switch link followed by an internal port-to-port connection in s followed by another switch link. We denote by $G(N)$ the static network G augmented with the dynamic links (each of weight μ) resulting from the configuration N and we call $G(N)$ an augmented network. In the augmented network visualized in Fig. 3, for example: (a, b) is a dynamic link; (a, c) is a static link; and (e, h) is both a static link a dynamic link. Note that it is possible that an augmented network $G(N)$ is a multigraph.

The concepts defined above are driven by reconfigurable hardware technology such as optical switches, wireless (beamforming) and free-space optics, all of which establish port-to-port connections, i.e., switch matchings. The survey paper [14] provides some detail as regards the relationship between the emergent theoretical models and current opto-electronic technology.

2.2. Routing in hybrid networks

Consider again the example shown in Fig. 3. In the configuration shown, a message M from c to node e may be routed:

1. via static links only, along the path $\varphi_1 := c \rightarrow b \rightarrow f \rightarrow e$ with weight $3w$, or
2. via dynamic links only, along the path $\varphi_2 := c \rightarrow d \rightarrow h \rightarrow e$ with weight 3μ , or
3. via a combination of static and dynamic links, along the path $\varphi_3 := c \rightarrow d \rightarrow e$ with weight $\mu + w$.

Depending on the value of μ , any of the paths may minimize the cost to route M : if $\mu \geq 2w$ then φ_1 is optimal; if $\mu \leq \frac{w}{2}$ then φ_2 is optimal; and if $\mu \in [\frac{w}{2}, 2w]$ then φ_3 is an optimal. We may wish to bound the number of alternations allowed between optic and static links in any path a message takes; we capture this hardware requirement via a segregation parameter $\sigma \in \mathbb{N} \cup \{\infty\}$, as introduced in [13], that is the number of alternations between static and dynamic links. In the fully segregated case, $\sigma = 0$: messages may be routed either by static links only (as in φ_1) or by dynamic links only (as in φ_2). In the non-segregated case, $\sigma = \infty$ and there is no restriction on the number of alternations, so any path is admitted. Note φ_3 is admitted as a valid path to route M if and only if $\sigma \geq 1$. The dynamic link limit δ , like the segregation parameter σ , is a restriction on admissible flow-paths. Whereas σ describes the maximum number of alternations between static and dynamic links permitted, δ describes the maximum number of dynamic links any flow-path may use. In particular, when $\delta = 1$ every flow-path must contain at most one dynamic link.

Networks are expected to route many messages (of varying sizes) optimally at the same time. Given a hybrid network $G(S)$ we represent the set of all demands we must optimize for as a workload (matrix) D with entries $\{D[u, v] \in \mathbb{Q}_+ : u, v \in V\}$ providing the intended pairwise node-to-node workloads (each $D[u, u]$ is necessarily 0).

Given a configuration N and $u, v \in V$ for which $D[u, v] > 0$, we route the corresponding workload via a path in $G(N)$ from u to v in $G(N)$ so that this chosen flow-path $\varphi(u, v)$ has workload cost $D[u, v] \times wt_{G(N)}(\varphi(u, v))$, where the weight $wt_{G(N)}(\varphi(u, v))$ is the sum of the weights of the links of the flow-path $\varphi(u, v)$ (if $G(N)$ has both a static link (x, y) and a dynamic link (x, y) then we need to say which we are using in $\varphi(u, v)$). The total workload cost (of D under N) is defined as

$$\sum_{u, v \in V, D[u, v] > 0} D[u, v] \times wt_{G(N)}(\varphi(u, v)).$$

Our aim will be to find a configuration N in some hybrid network $G(S)$ and flow-paths in $G(N)$ for which the total workload cost of some workload matrix D is minimized. In an unrestricted scenario, we would choose any flow-path $\varphi(u, v)$ to be a flow-path of minimum weight from u to v in $G(N)$, the weight of which we denote by $wt_{G(N)}(u, v)$. When $\sigma \neq \infty$ we must also ensure the flow-path has at most σ alternations. We also have the analogous concepts $wt_G(\varphi(u, v))$ and $wt_G(u, v)$ where we work entirely in the static network G . Note that we often describe D by a weighted digraph, which we usually call D' , so that the node set is V and there is an edge (u, v) of weight $w > 0$ if, and only if, $D[u, v] = w$. We also refer to some $D[u, v] > 0$ as a demand (from u to v).

2.3. The reconfigurable routing problem

We are now in a position to introduce our protagonist:

RECONFIGURABLE ROUTING PROBLEM (σ) (RRP(σ))

Input: $(G, S, \mu, w, D, \kappa)$: D is a workload matrix for the hybrid network $G(S)$ with static (resp. dynamic) links all of weight w (resp. μ).

Question: Does $G(S)$ admit some configuration N such that the total workload cost of D under N (where the number of alternations for any path is bounded by σ) is at most κ ?

As previously alluded to, this setting is more expressive than we require for most of this paper, and more restrictive than the exact formalism considered in prior work [11–13]: in those works, w and μ are sometimes allowed to be functions of their endpoints rather than fixed constants. This provides much more expressivity; notably, their model loses no power when it is restricted to inputs where G is a complete graph and there is only one switch, since it is possible to simulate any other instance by assigning prohibitively large weights to any static edges and any pair of switch ports which should not be usable.

We now turn to the “realistic” networks we mentioned in our introduction. Henceforth unless otherwise specified, static link weights are all equal (and normalized to 1) and dynamic link weights are always some fixed constant $\mu \in \mathbb{Q}_+$. Also, there is a single switch and all nodes are connected to it with identical hardware. This is both practically relevant and intuitively realistic; see e.g. Fig. 1. Then the set of switches S of the hybrid network consists of just one switch, which is fully described by the number of switch links each node in the hybrid network has, which we call Δ_S . This is closely related to the maximum reconfigurable degree Δ_R from [13], which is an upper bound on the number of external ports per node. The resulting restriction of RRP can be formalized as follows:

Δ_S -SWITCHED RRP (σ)

Input: (G, μ, D, κ) : D is a workload matrix for the hybrid network $G(S)$ with static (resp. dynamic) links all have weight 1 (resp. μ) (where S consists of a single switch that every node in G is connected to exactly Δ_S times).

Question: Does $G(S)$ admit some configuration N such that the total workload cost of D under N (where the number of alternations for any path is bounded by σ) is at most κ ?

3. Results

Table 1 shows a summary of hardness results from previous work as well as our three main intractability results. In general terms, we obtain NP-completeness for 2-SWITCHED RRP and 3-SWITCHED RRP on any fixed class of static networks of practical interest (defined more fully below) and for any value of σ . We then restrict our focus (and associated parameters) to the case where the static network is a hypercube when we establish the NP-completeness of 1-SWITCHED RRP($\sigma = 3$) in this setting; we conjecture that a similar construction can be used to establish hardness when $\sigma > 3$. We also, in Theorem 3, show that 1-SWITCHED RRP($\sigma = 0$) is solvable in polynomial time. The cases when $\sigma \in \{1, 2\}$ remain interesting open problems. Subsection 3.1 is devoted to the case of 1-SWITCHED RRP($\delta = 1$) (i.e. any flow-path must contain at most a single dynamic link) which entails at most 2 alternations, and we show this case is NP-complete for hypercubes, grids, and toroidal grids. The proof in that section is intended to facilitate the task of proving hardness for other interesting families of networks.

As is standard in NP-hardness proofs, we reduce from known NP-complete problems to instances of RRP; the challenge is that, due to the expansive scope of our theorems, we lose several “degrees of freedom” which are used for encoding hard instances in, e.g., [10–12]. Specifically, we may not make use of varying static or dynamic link weights to prohibit certain connections, nor encode any features of the input instance in the topology of the hybrid network $G(S)$. For example, in Lemma 1 [12], many small switches with

Table 1

Settings for some pre-existing hardness results for RRP. $|S|$ is the number of switches; Δ_R is the maximum number of external ports per node; σ is the segregation parameter and δ is the dynamic link limit; D is the workload matrix; n denotes the number of nodes in the instance.

Result	$ S $	Δ_R	σ/δ	D	link weights	notes
[11], Theorem 1	$\Theta(n)$ (or 1^a)	$\Theta(n)$	any $\sigma \geq 2$	sparse, all values 0 or 1	variable; $w \in [1, 100n^2]$ $\mu \in [1, 100n^2]$	Shown inapprox. within $\Omega(\log n)$
[12], Lemma 1	$\Theta(n)$	1			fixed; $w = \mu = 1$	All switches have 3 ports.
[12], Theorem 2	1					
[13], Theorems 4.1, 4.2		2	any $\sigma \geq 0$	dense, values in $\text{poly}(n)$	G is empty; there are no static links	
Theorem 1		3	sparse, values in $\text{poly}(n)$	fixed; $w = 1$ $\mu \in \Theta(\frac{1}{\text{poly}(n)})$	$G \in \mathcal{H}$, where $\mathcal{H}$ is any polynomial family of networks (incl. hypercubes, grids, cycles).	
Theorem 2	fixed; $w = 1$ $\mu \in \Theta(\frac{1}{\log(n)})$					
Theorem 4	1			$\sigma = 3$	fixed; $w = 1$ any $\mu \in (0, 1)$	G is a hypercube
Theorem 8			$\delta = 1$		G is a hypercube, grid, or toroidal grid	

^a By using variable μ with prohibitively large weights, it is possible to simulate many switches with just one.

two feasible configurations each are used to encode a truth assignment, and in Theorem 1 of [11] “bad” links are given weights of order $\Theta(n^2)$. Neither of these mechanisms can be leveraged to obtain hardness in our setting; in this sense, our hardness results are strictly stronger and also harder to obtain than those from [10–12]. We are constrained to choose a size for the network G , and then to encode the input instance in the demand matrix D .

Our first two results hold for a wide class of graph families, which may be of broader interest for the study of computational hardness in network problems. Rather than allowing arbitrary static networks in instances of RRP, we wish to force any such static network to come from a fixed family of networks where a *family of networks* $\mathcal{H}$ is an infinite sequence of networks $\{H_i : i \geq 0\}$ so that the size $|H_i|$ of any H_i is less than the size of H_{i+1} . However, we wish to control the sequence of network sizes. Consequently, we define a *polynomial family of networks* as being a family of networks $\mathcal{H} = \{H_i : i \geq 0\}$ where there exists a polynomial $p_{\mathcal{H}}(x)$ so that $|H_{i+1}| = p_{\mathcal{H}}(|H_i|)$, for each $i \geq 0$.² Note that given any $n \geq 0$, we can determine in time polynomial in n the smallest i such that $n \leq |H_i|$. As an example of a polynomial family of networks, consider the hypercubes; here, the polynomial $p_{\mathcal{H}}(x) = 2x$. Other examples include independent sets, complete graphs, cycles, complete binary trees and square grids, among many others. The sweeping generality of having a single construction which holds for any polynomial family $\mathcal{H}$ poses a challenge in our proofs of Theorems 1 and 2; we require that our constructed network $H(S)$ behaves identically when H is a connected (or even complete) graph and, at the opposite extreme, when H is disconnected (or even independent).

Theorem 1. *For any polynomial family of networks $\mathcal{H} = \{H_i : i \geq 0\}$, the problem 2-SWITCHED RRP restricted to instances (H, μ, D, κ) satisfying:*


- $H \in \mathcal{H}$ has size n
- the workload matrix D is sparse and all values in it are polynomial in n
- $\mu \in \Theta(\frac{1}{\text{poly}(n)})$ is fixed for all dynamic links

is NP-complete.

Proof. Note that 2-SWITCHED RRP as in the statement of the theorem is in NP as it can be straightforwardly reduced to an equivalent instance of the more general RRP, which is known to be in NP. We now build a polynomial-time reduction from the problem 3-MIN-BISECTION to our restricted version of RRP where the problem 3-MIN-BISECTION is defined as follows:

- instance of size n : a 3-regular graph $G = (V, E)$ on n nodes and an integer $k \leq n^2$

² Technically, we insist that there exists a polynomial Turing machine $\mathcal{M}$ which computes H_{i+1} on input H_i , for each $i \geq 0$, but this definition obfuscates the utility of this description.

Fig. 4. The graph D' .

- yes-instance: there exists a partition of V into two disjoint subsets A and B , each of size $\frac{n}{2}$, so that the set of edges incident with both a node in A and a node in B has size at most k ; that is, G has *bisection width* at most k .

Note that any 3-regular graph necessarily has an even number of nodes. The problem 3-MIN-BISECTION was proven to be NP-complete in [5] where it was also shown that approximating 3-MIN-BISECTION to within a constant factor approximation ratio entails the existence of a constant factor approximation algorithm for the more general and widely-studied problem MIN-BISECTION, which is defined as above but where G can be arbitrary and where A and B have sizes differing by at most one. Given an arbitrary instance $(G = (V, E), k)$ of size n of 3-MIN-BISECTION, we now build an instance (H, μ, D, κ) of 2-SWITCHED RRP. Moreover, we may assume that $k \leq \frac{n}{3} + 46$ as it was proven in [9] that every 3-regular graph has bisection width at most $\frac{n}{3} + 46$.

First, we define a weighted digraph $D' = (V', E')$ which will encode a description of our workload matrix D via: there is a directed edge (u, v) with weight w if, and only if, there is a node-to-node workload of w from u to v . Let $\bar{n}$ be the size of the network H_i where i is the smallest integer such that $n + 6n^2 + 2 \leq |H_i|$ and set $H = H_i$.

- The node set V' is taken as a disjoint copy of the node set V of G , which we also refer to as V , together with the set of nodes $V_c = \{x_i, y_i : -\frac{L}{2} \leq i \leq \frac{L}{2}\}$, where $L = 3n^2$ (recall, n is even), and another set of nodes U of size $\bar{n} - (n + 6n^2 + 2)$; so, $|V'| = \bar{n}$. We call every node of V_c a *chain-node*. For ease of presentation, we denote the chain-nodes $x_{\frac{L}{2}}$ and $x_{-\frac{L}{2}}$ by x^+ and x^- , respectively, and we define the chain-nodes y^+ and y^- analogously.
- The (directed) edge set E' consists of $E_\alpha \cup E_\beta \cup E_1$ where:
 - the set of *chain-edges* $E_\alpha = \{(x_i, x_{i+1}), (y_i, y_{i+1}) : 0 \leq i < \frac{L}{2}\} \cup \{(x_i, x_{i-1}), (y_i, y_{i-1}) : -\frac{L}{2} < i \leq 0\}$
 - the set of *star-edges* $E_\beta = \{(x_0, v) : v \in V\} \cup \{(y_0, v) : v \in V\}$
 - the set of *unit-edges* E_1 which is a copy of the edges E of G , but on our (copied) node set V and so that every edge is replaced by a directed edge of arbitrary orientation.

Note that the nodes of U are all isolated in D' and that $|V'| = \bar{n}$ (the nodes of U will play no role in the following construction). The workloads on the edges of E' are α , β or 1 depending upon whether the edge is a chain-edge from E_α , a star-edge from E_β or a unit-edge from E_1 , respectively, where we define $\alpha = 24n^6$ and $\beta = 6n^3$. If the directed edge (u, v) has weight α (resp. β , 1) in D' then we say that (u, v) is an α -demand (resp. β -demand, 1-demand). The digraph D' can be visualized as in Fig. 4. The grey rectangle denotes the nodes of V , the nodes of V_c appear along the top and the bottom and the dashed (resp. dotted, solid) directed edges depict the chain-edges (resp. star-edges, unit-edges). The nodes of U are omitted.

As stated earlier, our static network H is the network $H_i \in \mathcal{H}$ where $|H_i| = \bar{n}$. We refer to the node set of H as V' also and we refer to the subset of nodes within V' corresponding to V as V also. Since we are in the 2-switched setting, we have one switch s with $2|V'|$ ports so that every node of H is adjacent, via switch links, to exactly two ports of the switch. Hence, our switch set is $S = \{s\}$ and our hybrid network is $H(S)$. It is important to note that for any configuration N , any node of $H(N)$ can be adjacent to at most 2 other nodes via dynamic links (as $\Delta_S = 2$).

As can be seen, we have the graph $G = (V, E)$, the digraph $D' = (V', E')$ and the hybrid network $H(S)$ with node set V' . Although G , D' and $H(S)$ are disjoint in terms of node sets, we do not distinguish between, say, the node set V of G and the subset of nodes V of H . It should always be obvious as to which set we are referring to. We proceed similarly when we talk of specific nodes. As ever, we refer to ‘edges’ in graphs and ‘links’ in networks (but they are really one and the same).

We set the weight of any dynamic link as $\mu = \frac{1}{2L} = \frac{1}{6n^2}$ and the bound κ for the total workload cost as $\kappa = \kappa_\alpha + \kappa_\beta + \kappa_1$ where:

- $\kappa_\alpha = 24n^6$
- $\kappa_\beta = 3n^4 + \frac{n^3}{2} + n^2$
- $\kappa_1 = \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2}$.

The values of κ_α , κ_β and κ_1 have the following significance.

- Suppose that for every chain-edge (x_i, x_{i+1}) (resp. (x_i, x_{i-1}) , (y_i, y_{i+1}) , (y_i, y_{i-1})) of E_α , we force a dynamic link joining x_i and x_{i+1} (resp. x_i and x_{i-1} , y_i and y_{i+1} , y_i and y_{i-1}) in $H(N)$ and choose a corresponding flow-path serving this α -demand as consisting of this dynamic link (N is the resulting configuration from our chosen switch matching). The total workload cost of flow-paths serving α -demands is $2L\alpha\mu = 24n^6 = \kappa_\alpha$.
- Further, suppose that the dynamic links incident with nodes of V in $H(N)$ are chosen so that we have a path of dynamic links p_A from x^+ to either y^- or y^+ , involving the subset of nodes $A \subseteq V$, and a path of dynamic links p_B from x^- to y^+ or y^- , respectively, involving the subset of nodes $B \subseteq V$, so that both p_A and p_B have length $\frac{n}{2} + 1$. That is, we choose the dynamic links so that they form a cycle C (of length $n + 2L + 2$) in $H(N)$ covering exactly the nodes of V and V_c . Suppose that for any star-edge (x_0, v) (resp. (y_0, v)) of E_β , we choose the flow-path in $H(N)$ serving this star-edge as consisting entirely of dynamic links resulting from the shortest path in our cycle C from x_0 to v (resp. y_0 to v). The total workload cost of flow-paths corresponding to the star-edges is

$$4\mu\beta \sum_{i=1}^{\frac{n}{2}} \left(\frac{L}{2} + i \right) = 3n^4 + \frac{n^3}{2} + n^2 = \kappa_\beta.$$

- Further, suppose we choose the flow-path in $H(N)$ serving the 1-demand (u, v) (in E_1) to be a path of dynamic links within the cycle C of shortest length. If u and v both lie on p_A or both lie on p_B then the workload cost of this flow-path is at most $\mu(\frac{n}{2} - 1) = \frac{1}{6n}(\frac{1}{2} - \frac{1}{n})$, and if one of u and v lies on p_A with the other node lying on p_B then the workload cost of this flow-path is at most $\mu(\frac{n}{2} + L + 1) = \frac{1}{2} + \frac{1}{12n} + \frac{1}{6n^2}$. If the width of the bisection of G formed by A and B is at most k then the total workload cost of flow-paths corresponding to the unit-edges is at most

$$\left(\frac{3n}{2} - k \right) \mu \left(\frac{n}{2} - 1 \right) + k \mu \left(\frac{n}{2} + L + 1 \right) = \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2} = \kappa_1.$$

From above, we immediately obtain that if (G, k) is a yes-instance of 3-MIN-BISECTION then $(H(S), D, \mu, \kappa)$ is a yes-instance of RRP.

Conversely, suppose that (H, μ, D, κ) is a yes-instance of RRP. Let N be a configuration (consisting of a switch matching of s) and let F be a collection of flow-paths that witness that the total workload cost for the workload matrix D is at most κ . Denote by N , also, the sub-network of $H(N)$ consisting solely of dynamic links. W.l.o.g. we may assume that every node of $H(N)$ is incident with exactly two dynamic links of N (by adding additional dynamic links that we do not actually use in any flow-path, if necessary).

Claim 1. *If $(u, v) \in E_\alpha$ (that is, (u, v) is a chain-edge in D') then (u, v) is a dynamic link in N ; so, the total workload cost of the flow-paths serving α -demands is exactly $\kappa_\alpha = 24n^6$.*

Proof. Suppose that $(x_i, x_{i+1}) \notin N$ (there is an analogous argument for each of (x_i, x_{i-1}) , (y_i, y_{i+1}) and (y_i, y_{i-1})). So, the flow-path of F serving the α -demand (x_i, x_{i+1}) consists of at least two links and has workload cost at least $2\mu\alpha$ (as 2μ is strictly less than 1 which is the weight of a static link). Hence, the total workload cost of flow-paths corresponding to the chain-edges is at least $\kappa_\alpha + \mu\alpha$. Denote the total workload cost of flow-paths serving the β -demands (resp. 1-demands) by $\bar{\kappa}_\beta$ (resp. $\bar{\kappa}_1$). So, $\kappa_\alpha + \mu\alpha + \bar{\kappa}_\beta + \bar{\kappa}_1 \leq \kappa$ with $\mu\alpha \leq \mu\alpha + \bar{\kappa}_\beta + \bar{\kappa}_1 \leq \kappa_\beta + \kappa_1$; that is, with $n^4 \leq \frac{n^3}{2} + n^2 + \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2}$. This yields a contradiction (so long as n is big enough) and the claim follows. $\square$

Claim 2. *The set of dynamic links N forms a cycle in $H(N)$ covering exactly the nodes of $V \cup V_c$ and on which every link from $\{(x_i, x_{i+1}), (y_i, y_{i+1}) : 0 \leq i < \frac{L}{2}\} \cup \{(x_i, x_{i-1}), (y_i, y_{i-1}) : -\frac{L}{2} < i \leq 0\}$ lies. Moreover, every flow-path of F serving a β -demand consists entirely of dynamic links.*

Proof. By Claim 1, if $(u, v) \in E_\alpha$ then $(u, v) \in N$ and the total workload cost of the flow-paths serving the α -demands is exactly $\kappa_\alpha = 24n^6$. Note that because every node of $H(N)$ is adjacent to at most 2 dynamic links, there are no other dynamic links incident with a node from $V_c \setminus \{x^+, x^-, y^+, y^-\}$.

Suppose that a flow-path of F serving a β -demand consists entirely of dynamic links. So, the workload cost of such a flow-path is at least $\mu\beta(\frac{L}{2} + 1) = \frac{3n^3}{2} + n$. Alternatively, if a flow-path of F serving a β -demand contains at least one static link then the workload cost of such a flow-path is at least $\beta = 6n^3$. Consequently, if at least one flow-path of F serving a β -demand contains a static link (of weight 1) then the total workload cost of flow-paths serving the β -demands is at least $(2n - 1)\mu\beta(\frac{L}{2} + 1) + \beta = (2n - 1)(\frac{3n^3}{2} + n) + 6n^3 =$

$3n^4 + \frac{9n^3}{2} + 2n^2 - n > 3n^4 + \frac{n^3}{2} + n^2 + \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2} = \kappa_\beta + \kappa_1 = \kappa - \kappa_\alpha$ which yields a contradiction. Thus, all flow-paths of F serving a β -demand consist entirely of dynamic links. In particular, N must be connected. As N is regular of degree 2 then N must, in fact, be a cycle covering the nodes of $V \cup V_c$. The claim follows. $\square$

By Claim 2, N consists of a path of dynamic links involving all links from $\{(x_i, x_{i+1}) : -\frac{L}{2} \leq i < \frac{L}{2}\}$ from x^- to x^+ concatenated with a path p_A of dynamic links from x^+ to either y^- or y^+ concatenated with a path of dynamic links involving all links from $\{(y_i, y_{i+1}) : -\frac{L}{2} \leq i < \frac{L}{2}\}$ from y^- or y^+ to y^+ or y^- concatenated with a path p_B of dynamic links from y^+ or y^- to x^- , respectively. W.l.o.g. we may assume that p_A runs from x^+ to y^+ and p_B from y^- to x^- .

We can now recalculate the total workload cost of the flow-paths serving the β -demands. Suppose that there are p nodes of V on p_A and so $n - p$ nodes of V on p_B . We may assume that no nodes of U lie on p_A or p_B : if any do, then removing all such nodes from those paths results in a strictly smaller total workload cost. The total workload cost of the flow-paths serving the β -demands is

$$\begin{aligned} & 2\mu\beta \sum_{i=1}^p \left(\frac{L}{2} + i\right) + 2\mu\beta \sum_{i=1}^{n-p} \left(\frac{L}{2} + i\right) \\ &= p\mu\beta L + (n-p)\mu\beta L + 2\mu\beta \sum_{i=1}^p i + 2\mu\beta \sum_{i=1}^{n-p} i \\ &= n\mu\beta L + (p(p+1) + (n-p)(n-p+1))\mu\beta \\ &= 3n^4 + n(2p^2 - 2np + n^2 + n) \\ &\geq 3n^4 + \frac{n^3}{2} + n^2 = \kappa_\beta \text{ (when } p \text{ takes the value } \frac{n}{2}) \end{aligned}$$

as the minimum value for $3n^4 + n(2p^2 - 2np + n^2 + n)$ is when $p = \frac{n}{2}$.

Suppose that the paths p_A and p_B have different lengths. From above, the total workload cost of the flow-paths serving the β -demands is at least $3n^4 + n(2(\frac{n}{2} + 1)^2 - 2n(\frac{n}{2} + 1) + n^2 + n) = 3n^4 + \frac{n^3}{2} + n^2 + 2n = \kappa_\beta + 2n$. Consequently, we must have that $2n \leq \kappa_1 = \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2}$. As $k \leq \frac{n}{3} + 46$ (from [9]), we have that $2n \leq \frac{n}{6} + 23 + \frac{1}{8} - \frac{1}{4n} + \frac{1}{9n} + \frac{46}{3n^2}$ which yields a contradiction. Hence, the paths p_A and p_B each have length $\frac{n}{2}$, with the total workload cost of the flow-paths corresponding to the unit-edges being at most κ_1 .

Let (u, v) be a unit-edge. From above: if u and v both lie on p_A or both lie on p_B and the flow-path serving the 1-demand (u, v) consists entirely of dynamic links then the workload cost of this flow-path lies between μ and $\mu(\frac{n}{2} - 1)$, i.e., $\frac{1}{6n^2}$ and $\frac{1}{6n}(\frac{1}{2} - \frac{1}{n})$; if u and v lie on p_A and p_B , respectively, or vice versa, and the flow-path serving the 1-demand (u, v) consists entirely of dynamic links then the workload cost of this flow-path lies between $\mu(L + 2)$ and $\mu(\frac{n}{2} + L + 1)$, i.e., $\frac{1}{2} + \frac{1}{3n^2}$ and $\frac{1}{2} + \frac{1}{12n} + \frac{1}{6n^2}$; and if the flow-path serving the 1-demand (u, v) contains a static link then the workload cost of this flow-path is at least 1. In particular, we may assume that any flow-path corresponding to a unit-edge consists entirely of dynamic links.

Let A (resp. B) be the nodes of V that appear on p_A (resp. p_B). So, (A, B) is a bisection of G . Suppose that this bisection has width $k + \epsilon$, for some $\epsilon \geq 1$. So there are $k + \epsilon$ unit-edges whose corresponding work-flows contribute collectively at least $(k + \epsilon)\mu(L + 2) \geq (k + 1)(\frac{1}{2} + \frac{1}{3n^2}) = \frac{k}{2} + \frac{1}{2} + \frac{1}{3n^2} + \frac{k}{3n^2} > \frac{k}{2} + \frac{1}{8} - \frac{1}{4n} + \frac{k}{3n^2} = \kappa_1$ to the total workload cost, which yields a contradiction. Consequently, G has bisection width at most k and we have that if (H, μ, D, κ) is a yes-instance of RRP then (G, k) is a yes-instance of 3-MIN-BISECTION. Our result follows as (H, μ, D, κ) can be constructed from (G, k) in time polynomial in n . $\square$

This result significantly strengthens Theorems 4.1 and 4.2 from [13]: there, $\text{RRP}(\Delta_R \geq 2, \sigma = 0)$ is shown to be NP-complete when the static network is an independent set, and the proof does not enable us to restrict the workload matrix D meaningfully. The main weakness of Theorem 1 is its reliance on μ being a polynomial factor smaller than any static link weight. This is actually related to the fact that a connected 2-regular network, as is $G(N)$ when G is an independent set and $\Delta_S = 2$, has diameter linear in the number of nodes n . A network of maximum degree 3, on the other hand, may have diameter logarithmic in n (e.g., a complete binary tree has this property) and we indeed show NP-completeness of $\text{RRP}(\Delta_S = 3)$ when $\mu = \Theta(\frac{1}{\log n})$.

Theorem 2. For any polynomial family of networks $\mathcal{H} = \{H_i : i \geq 0\}$, the problem 3-SWITCHED RRP restricted to instances (H, μ, D, κ) satisfying:

- $H \in \mathcal{H}$ has size n
- the workload matrix D is sparse and all values in it are polynomial in n
- $\mu \in \Theta(\frac{1}{\log(n)})$

is NP-complete.

Proof. As in the case of Theorem 1, membership of NP is straightforward. The problem RESTRICTED EXACT COVER BY 3-SETS (RXC3) is defined as follows:




Fig. 5. Attaching the gadgets.

- instance of size n : a finite set of *elements* $\mathcal{X} = \{x_i : 1 \leq i \leq 3n\}$, for some $n \geq 1$, and a collection C of 3-element subsets of $\mathcal{X}$, called *clauses*, where $C = \{c_j : 1 \leq j \leq 3n\}$ and where each $c_j = \{x_1^j, x_2^j, x_3^j\}$ so that every element of $\mathcal{X}$ appears in exactly three clauses of C
- yes-instance: C contains an *exact cover* C' for $\mathcal{X}$; that is, a subset $C' \subseteq C$ so that every element of $\mathcal{X}$ appears in exactly one clause of C' (hence, C' necessarily has size n).

The problem RXC3 is NP-complete as was proven in Theorem A.1 of [15]. Let $(\mathcal{X}, C)$ be an instance of RXC3 of size n . We will build an instance (H, μ, D, κ) of 3-SWITCHED RRP corresponding to $(\mathcal{X}, C)$.

First, we define a weighted graph D' so as to describe the workload D . Having built D' , we will orient every undirected edge so that: there is a directed edge (u, v) with weight w if, and only if, there is a node-to-node workload of w from u to v (we still refer to the resulting digraph as D'). Consequently, we will need one flow-path in our resulting hybrid network corresponding to each directed edge of $D' = (V', E')$. Our construction proceeds in stages.


- We begin with a complete binary tree T of depth d so that the number of leaves is 2^d with $2^{d-1} < n \leq 2^d$. If the leaves are $\{l_i : 1 \leq i \leq 2^d\}$ then we colour: every leaf l_i , where $n < i$, white; every leaf l_i , where $1 \leq i \leq n$, grey; and the root of T grey. Throughout, if we do not specify the colour of a node as grey or white then it is coloured black. Note that $d = \lceil \log_2(n) \rceil$.
- For every $x_i \in \mathcal{X}$, we create a unique node x_i and for every clause $c_j \in C$, we create a unique 2-path as follows: there are 3 nodes u_j, v_j and w_j together with the edges (u_j, v_j) and (v_j, w_j) and the node u_j is coloured grey.
- For each $1 \leq j \leq 3n$, if the clause $c_j = \{x_1^j, x_2^j, x_3^j\}$ then there are edges (x_1^j, u_j) , (x_2^j, w_j) and (x_3^j, w_j) .
- To every node coloured grey, we attach a unique gadget consisting of a 4-clique with one of its edges subdivided by a node, denoted z , so that there is an edge joining the gray node and the node z ; and to every node coloured white, we attach the same gadget except that we identify the white node and the node denoted z . These attachments can be visualized as in Fig. 5a and Fig. 5b where the white and gray nodes are as shown.
- There is an edge (r, x_i) , for all $1 \leq i \leq 3n$: call these the *root-edges*.

This completes the construction of the graph D' which can be visualized as in Fig. 6 where the root-edges are depicted as dashed. Note that the number of nodes in V' is $N = 6(2^d) + 28n + 4$ and the number of edges in E' is $9(2^d) + 43n + 6$, with $3n$ of these edges root-edges. Let D'' be D' with the root-edges removed. Note that every node of D'' has degree 3 except for the leaf nodes of $\{l_i : 1 \leq i \leq n\}$ and the nodes of $\{v_j : 1 \leq j \leq m\}$, all of which have degree 2 (we return to this comment presently). Let $E' = E'' \cup E^{root}$ where E^{root} is the set of root-edges (and so E'' is the edge set of D''). As regards the edge-weights, orient each of the edges of E' 'down' the graph D' as it is portrayed in Fig. 6 so as to obtain a digraph and: give each directed edge of E'' weight α , where $\alpha = 4n \log_2(n)$; and give each directed root-edge of the form (r, x_i) weight 1.

Let $\bar{n}$ be the size of the network H_i where i is the smallest integer such that $N \leq |H_i|$ and we fix our static network as $H = H_i$. We name a subset of nodes of H as V' with the remaining nodes named U . $\Delta_S = 3$, so we have one switch s that has $3n$ switch ports and every node of H has 3 external ports; so, every node of H has exactly 3 switch links to s , and $S = \{s\}$. Finally, we define $\mu = \frac{1}{2 \lceil \log_2(n) \rceil}$ and $\kappa = \mu \alpha |E''| + 3\mu n (\lceil \log_2(n) \rceil + 3)$ so as to complete the construction of our instance (H, μ, D, κ) of 3-SWITCHED RRP.

Suppose that $(\mathcal{X}, C)$ is a yes-instance of RXC3. So, let $C' \subseteq C$ be an exact cover of $\mathcal{X}$. For every $(u, v) \in E''$, we choose (u, v) to be a dynamic link and ensure that every node of $\{v_j : 1 \leq j \leq 3n, c_j \in C'\}$ is joined via a dynamic link to a leaf node of $\{l_i : 1 \leq i \leq n\}$. If v_j is such that $c_j \in C'$ then denote the unique leaf node to which there is a dynamic link from v_j by $\bar{l}_j$. Denote the resulting switch matching of s as constituting the configuration N . Referring back to our comment on node degrees above, the dynamic links so added must form a matching in $H(N)$. For any workload of weight α in D , we choose the corresponding flow-path to consist of the dynamic link joining the two nodes in question. Consequently, for all of the workloads corresponding to directed edges of E'' , the total workload cost of the corresponding flow-paths is $\mu \alpha |E''|$. All of the remaining workloads involve the root and the nodes of $\{x_i : 1 \leq i \leq n\}$. Any one of these nodes x_i is such that the element x_i in a unique clause $c_j \in C'$. We choose the flow-path consisting entirely of dynamic links that routes from the root down the tree T to the leaf-node $\bar{l}_j$ and on to v_j and, through either u_j or w_j , on to x_i . This flow-path has (graph-theoretic) length $d + 3$ and consequently the total workload cost due to flow-paths corresponding to edges of E^{root} is exactly $3\mu n(d + 3) = 3\mu n(\lceil \log_2(n) \rceil + 3)$. Hence, the total workload cost due to all flow-paths is $\mu \alpha |E''| + 3\mu n(\lceil \log_2(n) \rceil + 3) = \kappa$ and (H, μ, D, κ) is a yes-instance of 3-SWITCHED RRP.

Conversely, suppose that (H, μ, D, κ) is a yes-instance of 3-SWITCHED RRP and that N is a configuration and F a set of flow-paths witnessing that the total workload cost is at most κ .

Fig. 6. The graph D' .

Claim 1. *It is necessarily the case that every directed edge $(u, v) \in E''$ is such that (u, v) is a dynamic link in $H(N)$.*

Proof. Suppose, for contradiction, that there is a directed edge $(u, v) \in E''$ so that (u, v) is not a dynamic link in N . By construction, the node-to-node workload of α from u to v contributes a cost of at least $2\mu\alpha$ to the total workload cost (note that the cost of a static link is 1 and $2\mu < 1$ when $n \geq 3$). So, the total workload cost from flow-paths corresponding to directed edges of E'' is at least $\mu\alpha(|E''| + 1)$. However, $\kappa = \mu\alpha|E''| + 3\mu n(\lceil \log_2(n) \rceil + 3)$ and so we must have that $\mu\alpha = 4\mu n \log_2(n) \leq 3\mu n(\lceil \log_2(n) \rceil + 3)$ which yields a contradiction when $n \geq 2^{10}$. The claim follows. $\square$

By Claim 1, the total workload cost of flow-paths corresponding to the directed edges of E'' is exactly $\mu\alpha|E''|$. Consequently, the total workload cost of flow-paths corresponding to the directed edges of E^{root} must be at most $3\mu n(\lceil \log_2(n) \rceil + 3)$.

Given the set of dynamic links as supplied by Claim 1, the only other possible dynamic links we might have involve the leaf nodes $\{l_i : 1 \leq i \leq n\}$, the nodes of $\{v_i : 1 \leq i \leq 3n\}$ and nodes of U . If we are looking for a path of dynamic links joining r and any node x_i then no matter how we might extend the set of dynamic links corresponding to the directed edges of E'' , we can never obtain a path of (graph-theoretic) length less than $d + 3$: a path of length d down the tree T from r to a leaf l followed by a path of length 3 of the form (l, v_j, u_j, x_i) or (l, v_j, w_j, x_i) .

We may assume U is not involved in any such path. Any path visiting some node $\bar{u} \in U$ can be shortened to a path of the former form by simply removing $\bar{u}$ from the path. This corresponds to a rewiring of the configuration N which decreases the total workload cost so, for example, the path $(r, \dots, l, \bar{u}, v_j, u_j, x_i)$ of length $n + 4$ can be shortened to $(r, \dots, l, v_j, u_j, x_i)$ of length $n + 3$ by removing the dynamic links $(l, \bar{u})$ and $(\bar{u}, v_j)$ from N and adding the dynamic link (l, v_j) to N . Clearly this does not increase the workload cost for any demand.

So, any flow-path corresponding to some directed edge $(r, x_i) \in E^{root}$ and consisting entirely of dynamic links has workload cost at least $\mu(d + 3) = \mu(\lceil \log_2(n) \rceil + 3)$. If a flow-path corresponding to some directed edge $(r, x_i) \in E^{root}$ contains a static link then the workload cost of this flow-path is at least 1 and $1 > \frac{1}{2} + \frac{3}{2\lceil \log_2(n) \rceil} = \mu(\lceil \log_2(n) \rceil + 3)$, when $n \geq 2^4$. As there are $3n$ flow-paths corresponding to directed edges of E^{root} , we must have that every flow-path corresponding to a directed edge of E^{root} consists entirely of dynamic links and has (graph-theoretic) length $\lceil \log_2(n) \rceil + 3$.

Consider the flow-path corresponding to the directed edge $(r, x_i) \in E^{root}$. Suppose that the dynamic link (l_k, v_j) appears on this flow-path (exactly one such dynamic link does). Note that there are at most n such dynamic links used in the flow-paths of F corresponding to the directed edges of E^{root} . Build the set of clauses $C' \subseteq C$ by including every such c_j in C' . Consequently, we have a subset $C' \subseteq C$ of at most n clauses. Moreover, as we can reach any x_i from r by a flow-path of $d + 3$ dynamic links, every $x_i \in \mathcal{X}$ appears in a clause of C' ; that is, $|C'| = n$, C' is an exact cover for $\mathcal{X}$ and $(\mathcal{X}, C)$ is a yes-instance of RXC3. As our instance (H, μ, D, κ) can be constructed from $(\mathcal{X}, C)$ in time polynomial in n , our result follows. $\square$

These results led us to consider the problem 1-SWITCHED RRP($\sigma = k$) for $k \geq 0$. By extending Lemma 1 due to [13], we show that this restriction leads to a tractable problem when either $\sigma = 0$ or the static network is a complete graph, in contrast with our NP-completeness results. That result makes use of the *dynamic link limit* δ - recall that when $\delta = 1$ every flow-path must contain at most one dynamic link. In combination with the constraint $\sigma = 0$, this entails that every flow path consists of either a single dynamic link (and no static links), or of one or more static links (and no dynamic links). The result below is proven via a maximum matching argument in [13].

Lemma 1 ([13, Theorem 3.1]). *The problem $RRP(\sigma = 0, \delta = 1)$ can be solved in polynomial-time when there is only one switch. Moreover, the problem remains solvable in polynomial-time even if we allow non-uniform weights for dynamic links.*

Theorem 3. *1-SWITCHED $RRP(\sigma = 0)$ is in P.*

Proof. Note that this proof holds even if we allow variable weights for static and optic links. Because each node v is connected to the switch exactly once, it is not possible that any vertex is incident to two optic links, and hence impossible for there to be any path consisting of two or more optic links. Consequently setting $\delta = 1$ introduces no new constraints; in this scenario a flow is permissible with $\sigma = 0$ if and only if it is permissible with $\sigma = 0$ and $\delta = 1$. Hence the input instance of 1-SWITCHED $RRP(\sigma = 0, \delta = 1)$ is a yes-instance if and only if the corresponding instance of $RRP(\sigma = 0, \delta = 1)$ is a yes-instance, and tractability follows from Lemma 1. $\square$

Corollary 1. *1-SWITCHED $RRP(\sigma = k)$ restricted to instances where the static network G is a complete graph is in P, for any $k \in \mathbb{N} \cup \{\infty\}$.*

Proof. If G is a complete graph with all edge weights equal (without loss take $wt(u, v) = 1 \forall (u, v) \in E$) then without loss under any configuration N , each demand $D[u, v]$ is routed via the flow-path $\phi(u, v)$ of minimum weight, which is either:

- a single static link from u to v with unit weight.
- a single dynamic link from u to v with weight μ .

It follows that setting $\sigma = 0$ introduces no new constraints, and then by Theorem 3 we have tractability. $\square$

Corollary 1 rules out the possibility that 1-SWITCHED RRP might be NP-complete for any polynomial graph family $\mathcal{H}$ (since such a claim would extend to the family of complete graphs) unless P equals NP. This leaves open the practically relevant case where $\Delta_S = 1$ and $\sigma > 0$ for *specific* topologies. We consequently consider the scenario where the static network is a hypercube and the segregation parameter $\sigma = 3$.

Theorem 4. *For any fixed $\mu \in (0, 1)$, the problem 1-SWITCHED $RRP(\sigma = 3)$ restricted to instances (H, μ, D, κ) satisfying:*

- $H \in \mathcal{Q}$, where $\mathcal{Q} := \{Q_d \mid d \in \mathbb{N}\}$ is the family of hypercubes
- the workload matrix D is sparse and all values in it are polynomial in n

is NP-complete.

Proof. Note that RRP as in the statement of the theorem is in NP on account of the numerical restrictions imposed. As in the proof of Theorem 2, we reduce from the problem RXC3. Let $(\mathcal{X}, C)$ be an instance of RXC3 of size n . W.l.o.g., we may assume that n is even. Let $m = \lceil \log_2(3n) \rceil$ (so, $3n \leq 2^m$). Our static network H is the hypercube Q_{8m} (so, $|Q_{8m}| = 2^{8m} = O(n^8)$) and there is exactly 1 switch link from every node to our switch s .

We begin by defining a weighted digraph $D' = (V, E')$ where V is the node set of our hypercube Q_{8m} . The digraph D' is to provide a description of our workload matrix. In order to define the directed edges of D' , we define some specific subsets of nodes of V . We explain the notation that we use as we proceed. By ‘distance’ we mean the distance (that is, shortest path length) in the hypercube Q_{8m} and by V_i (resp. $V_{\leq i}$) we mean the subset of nodes of V that are at distance exactly i (resp. at most i) from the root node $r = (1, 1, \dots, 1) = 1^{8m}$ (in general, we also denote nodes of V as bit-strings of length $8m$).

- Within the nodes of V_m , we define the subset $\tilde{P} = \{z\bar{z}1^{6m} : z \in \{0, 1\}^m\}$ (in general: yz denotes the concatenation of two bit-strings y and z , with z' denoting the concatenation of i copies of the bit-string z ; and $\bar{z}$ denotes the complement of the bit-string z). Note that $|\tilde{P}| = 2^m$.
- Within the nodes of V_{3m} , we define the subset $\tilde{C} = \{(z\bar{z})^3 1^{2m} : z \in \{0, 1\}^m\}$. Note that $|\tilde{C}| = 2^m$.
 - For any bit-string $z \in \{0, 1\}^{8m}$ and for any bit-string $y \in \{0, 1\}^i$ where $i < 8m$, we write $\oplus(z; y)$ to denote the bitwise exclusive OR of z and $0^{8m-i}y$ (so, $\oplus(z; y)$ differs from z on at most the last i bits). For any set U of bit-strings of length $8m$ and bit-string y of length $i < 8m$, we define $U^{\oplus y} = \{\oplus(z; y) : z \in U\}$.
 - For any $c \in \tilde{C}$, we define
 - * $c^{\oplus 001} = \oplus(c; 001)$
 - * $c^{\oplus 010} = \oplus(c; 010)$
 - * $c^{\oplus 100} = \oplus(c; 100)$
 and so we obtain sets of nodes $\tilde{C}^{\oplus 001}$, $\tilde{C}^{\oplus 010}$ and $\tilde{C}^{\oplus 100}$, each consisting of sets of neighbours of nodes in $\tilde{C}$ (note also that if $c_1, c_2 \in \tilde{C}$ are distinct then c_1 and c_2 are at distance at least 6 in Q_{8m} and any node of $\{c_1^{\oplus 001}, c_1^{\oplus 010}, c_1^{\oplus 100}\}$ and node of $\{c_2^{\oplus 001}, c_2^{\oplus 010}, c_2^{\oplus 100}\}$ are at distance at least 6 in Q_{8m}).
- Within the nodes of V_{4m} , we define the subset $\tilde{X} = \{(z\bar{z})^4 : z \in \{0, 1\}^m\}$. Note that $|\tilde{X}| = 2^m$.
 - For any $x \in \tilde{X}$, we define
 - * $x^{\oplus 01} = \oplus(x; 01)$

$$* \ x^{\oplus 10} = \oplus(x; 10)$$

and so we obtain sets of nodes $\tilde{X}^{\oplus 01}$ and $\tilde{X}^{\oplus 10}$, each consisting of sets of neighbours of nodes in $\tilde{X}$ (note that if $x_1, x_2 \in \tilde{X}$ are distinct then x_1 and x_2 are at distance at least 8 in Q_{8m} and any node of $\{x_1^{\oplus 01}, x_1^{\oplus 10}\}$ and node of $\{x_2^{\oplus 01}, x_2^{\oplus 10}\}$ are at distance at least 6 in Q_{8m}).

We define the weighted directed edges E' of D' as follows. There are three weights: some $\beta > \max\{\frac{3n(m+1+2\mu)}{\mu}, \frac{3n(m+1+2\mu)}{1-\mu}\}$; some $\alpha > 15n\beta + 9n^2 + 9n$; and 1. The directed edges of E_β (resp. E_α , E_1) all have weight β (resp. α , 1) and $E' = E_\beta \cup E_\alpha \cup E_1$.

- The directed edges of E_β are derived using the structure of our instance $(\mathcal{X}, C)$ of RXC3.
 - For each $1 \leq j \leq 3n$, identify the clause $c_j \in C$ with the node $c_j = (\text{bin}_m(j-1)\text{bin}_m(j-1))^{31}2^m$ of the subset $\tilde{C}$ of V_{3m} , where in general $\text{bin}_i(j)$ is the binary representation of the natural number j as a bit-string of length i (so, in particular, $0 \leq j < 2^i$). Henceforth, we refer to these specific nodes of $\tilde{C}$ as the set of *clause nodes* $C = \{c_j : 1 \leq j \leq 3n\}$.
 - For each $1 \leq i \leq 3n$, identify the element x_i of $\mathcal{X}$ with the node $x_i = (\text{bin}_m(i-1)\text{bin}_m(i-1))^{41}$ of the subset $\tilde{X}$ of V_{4m} . Henceforth, we refer to these specific nodes of $\tilde{X}$ as the set of *element nodes* $X = \{x_i : 1 \leq i \leq 3n\}$.
- For any clause $c_j = \{x_1^j, x_2^j, x_3^j\}$ of C , there are directed edges $(c_j^{\oplus 001}, x_1^i)$, $(c_j^{\oplus 010}, x_2^i)$ and $(c_j^{\oplus 100}, x_3^i)$ in E_β and we refer to the nodes $c_j^{\oplus 001}$, $c_j^{\oplus 010}$ and $c_j^{\oplus 100}$ as the *associate clause nodes* of clause c_j or of clause node c_j . We define the *associate element nodes* of an element or an element node analogously. This constitutes all directed edges of E_β .
- We denote by $P = \{p_i : 1 \leq i \leq n\}$ a subset of n nodes of $\tilde{P}$ which we refer to as the set of *port nodes*. The directed edges of E_α are $E_\alpha^P \cup E_\alpha^W$ where:
 - $E_\alpha^P = \{(z, \bar{z}) : z \in V_{\leq m+5} \setminus P\}$
 - E_α^W is an arbitrary orientation of a specific perfect matching M (coming up) on the nodes of $W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})$ where W is defined as $V_{4m} \cup \bigcup_{i=1}^3 V_{4m-i} \cup \bigcup_{i=1}^3 V_{4m+i}$.
- Note that $|W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})|$ is even and so a perfect matching exists (recall, we assumed that n is even); however, as stated, we require that our perfect matching M has a specific property which we describe below.
- The set of directed edges E_1 is defined as $\{(r, x_i) : 1 \leq i \leq 3n\}$; that is, as $\{r\} \times X$.

Now for our perfect matching M .

Claim 1. *There is a perfect matching M of $W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})$ so that if (u, v) is in the matching then there is no static link joining u and v .*

Proof. Define a perfect matching M on W as follows. First, match every node in V_{4m} with the node v that differs from u in every bit; that is, $v = \bar{u}$.


Consider some node $z\bar{z}z\bar{z}z\bar{z}z$ of X . Its matched node is $\bar{z}z\bar{z}z\bar{z}z\bar{z}z$ which is either in X or outside X . Suppose it is the latter. Suppose also that some other node $w\bar{w}w\bar{w}w\bar{w}w$ is in X and its matched node $\bar{w}w\bar{w}w\bar{w}w\bar{w}w$ lies outside X . Note that the distance between the two matched nodes $z\bar{z}z\bar{z}z\bar{z}z$ and $w\bar{w}w\bar{w}w\bar{w}w$ is at least 8. We amend our matching M by removing all matched pairs of nodes of X and choose an arbitrary matching on the remaining matched nodes (note that there is an even number of such matched nodes). Any pair of nodes in M is such that there is no static link joining them and every node of $V_{4m} \setminus X$ is involved in M .

Extend M with a matching so that every node of $V_{4m-1} \setminus (X^{\oplus 01} \cup X^{\oplus 10})$ is matched with a node in $V_{4m+1} \setminus (X^{\oplus 01} \cup X^{\oplus 10})$. This is possible as $|V_{4m-1}| = |V_{4m+1}|$ and $|V_{4m-1} \cap (X^{\oplus 01} \cup X^{\oplus 10})| = |V_{4m+1} \cap (X^{\oplus 01} \cup X^{\oplus 10})|$. Further extend M with a matching so that: every node of V_{4m-2} (resp. V_{4m-3}) is matched with a node in V_{4m+2} (resp. V_{4m+3}). The claim follows. $\square$

The digraph D' can be visualized in Fig. 7. The rows are intended to illustrate nodes in different V_i 's, although the rows labelled $V_{3m-1/3m+1}$ and $V_{4m-1/4m+1}$ contain all nodes of $V_{3m-1} \cup V_{3m+1}$ and $V_{4m-1} \cup V_{4m+1}$, respectively (this is where the associate clause nodes and the associate element nodes lie, respectively). We have not shown all directed edges; just enough to give a flavour of the construction. Directed edges from E_β are shown as solid directed edges with the directed edges corresponding to the clause $c_1 = \{x_1, x_n, x_2\}$, for example, depicted (the grey nodes are the associate clause nodes in order $c_1^{\oplus 001}, c_1^{\oplus 010}, c_1^{\oplus 100}, c_2^{\oplus 001}, \dots$ from left to right). Directed edges from E_α are shown as dotted directed edges and the grey rectangle contains all nodes from W involved in directed edges of E_α . The white nodes within this rectangle are the element nodes and the associate element nodes (that is, the nodes of W that are not incident with directed edges from E_α). Finally, the dashed directed edges depict the directed edges from E_1 and the port nodes and clause nodes are also shown in white.

In order to complete our instance (H, μ, D, κ) , we define μ to be any fixed (rational) value in the interval $(0, 1)$ and $\kappa = \kappa_\beta + \kappa_\alpha + \kappa_1$ where:

- $\kappa_\beta = \frac{|E_\beta|}{3}(\mu\beta) + \frac{2|E_\beta|}{3}((\mu+1)\beta) = 3n\mu\beta + 6n(\mu+1)\beta$
- $\kappa_\alpha = |E_\alpha|\mu\alpha$
- $\kappa_1 = |E_1|(m+1+2\mu) = 3n(m+1+2\mu)$.

Fig. 7. The graph D' .

Suppose that there is an exact cover $C' = \{c_{j_i} : 1 \leq i \leq n\} \subseteq C$ of $\mathcal{X}$. Let the set of clause nodes $C' \subseteq C$ be $\{c_{j_i} : 1 \leq i \leq n\}$. We define a configuration N by choosing dynamic links as follows.

- For some arbitrary bijection f from P to C' , $\{(p, f(p)) : p \in P\}$ is a set of n dynamic links.
- For each clause $c_j = \{x_1^j, x_2^j, x_3^j\} \in C'$, there are dynamic links $(c_j^{\oplus 001}, x_1^j)$, $(c_j^{\oplus 010}, x_2^j)$ and $(c_j^{\oplus 100}, x_3^j)$; so, all nodes of X are incident with a dynamic link as are all associate clause nodes of $\{c_j^{\oplus 001}, c_j^{\oplus 010}, c_j^{\oplus 100} : c_j \in C'\}$. Note that these dynamic links result from the directed edges of E_β incident with the associate clause nodes of the clause nodes of C' .
- All directed edges of E_α result in dynamic links. In particular, there is no dynamic link from a node outside W to a node inside W except possibly incident with the element nodes and the associate element nodes; indeed, the element nodes are already incident with such 'external' dynamic links.
- Consider the node $x_1 \in X$. The element x_1 appears in two clauses of $C \setminus C'$, say c_{i_1} and c_{i_2} . Suppose, for example, that the element x_1 appears as the second element of clause c_{i_1} and as the first element of clause c_{i_2} . If so then include a dynamic link from each of $c_{i_1}^{\oplus 010}$ and $c_{i_2}^{\oplus 001}$ to $x_1^{\oplus 01}$ and $x_1^{\oplus 10}$. Alternatively, if, say, the element x_1 appears as the third element of clause c_{j_1} and as the third element of clause c_{j_2} then include a dynamic link from each of $c_{j_1}^{\oplus 100}$ and $c_{j_2}^{\oplus 100}$ to $x_1^{\oplus 01}$ and $x_1^{\oplus 10}$. Proceed similarly and analogously with all remaining element nodes of $X \setminus \{x_1\}$. On completion of this iterative process, there is a dynamic link from every associate clause node to a unique element node or associate element node; that is, these dynamic links depict a bijection from the associate clause nodes to the element nodes and the associate element nodes.

This constitutes the configuration N .

Consider some β -demand of our workload. For some node $c_j \in C'$, all such workloads originating at the nodes $c_j^{\oplus 001}$, $c_j^{\oplus 010}$ and $c_j^{\oplus 100}$ can be served via a flow-path consisting of a solitary dynamic link at workload cost $\mu\beta$. For some node $c_j \in C \setminus C'$, suppose that there is a demand $D[c_j^{\oplus 001}, x_i^{\oplus 01}] = \beta$. This demand exists because the element x_i is the first element of clause c_j . As the element x_i is the first element of clause c_j , there is a dynamic link from $c_j^{\oplus 001}$ to a neighbour of x_i in $\{x_i^{\oplus 01}, x_i^{\oplus 10}\}$, which w.l.o.g. we may

assume to be $x_i^{\oplus 01}$; hence, the demand can be served via the flow-path $c_j^{\oplus 001}, x_i^{\oplus 01}, x_i$ at workload cost $(\mu + 1)\beta$. Defining other flow-paths analogously means that we can find flow-paths corresponding to the directed edges of E_β the total workload cost of which is $3n\mu\beta + 6n(\mu + 1)\beta = \kappa_\beta$.

All α -demands can be served via a flow-path consisting of a solitary dynamic link at total workload cost $|E_\alpha|\mu\alpha = \kappa_\alpha$.

Consider some 1-demand $D[r, x_i] = 1$. Let the clause $c_j \in C'$ be such that x_i is an element of c_j ; hence, there is a dynamic link from some port node $p_{j'}$ of P to c_j and a dynamic link from some neighbour of c_j from $\{c_j^{\oplus 001}, c_j^{\oplus 010}, c_j^{\oplus 110}\}$ to x_i . Consequently, there is a static path from r to $p_{j'}$ of weight m , a dynamic link from $p_{j'}$ to the node c_j , a static link from c_j to the above neighbour in $\{c_j^{\oplus 001}, c_j^{\oplus 010}, c_j^{\oplus 110}\}$ and a dynamic link from this neighbour to x_i . Hence, the demand $D[r, x_i] = 1$ can be served via a flow-path of workload cost $m + 1 + 2\mu$ with all 1-demands served via flow-paths at a total workload cost $3n(m + 1 + 2\mu) = \kappa_1$. Consequently, our instance (H, μ, D, κ) is a yes-instance of 1-SWITCHED RRP($\sigma = 3$).

Conversely, suppose it is the case that (H, μ, D, κ) is a yes-instance of 1-SWITCHED RRP($\sigma = 3$) and that N is a configuration and F a set of flow-paths witnessing that the total workload cost is at most κ .

Claim 2. Every directed edge (u, v) of E_α is necessarily such that $(u, v) \in N$ and the total workload cost of flow-paths serving the α -demands is exactly κ_α .

Proof. Suppose that at least one of the α -demands is served via a flow-path at a workload cost of more than $\mu\alpha$; so, it must be at a workload cost of more than $(1 + \mu)\alpha$ as we cannot traverse a dynamic link followed immediately by another dynamic link (recall, $\Delta_S = 1$) and by Claim 1, if $D[u, v] = \alpha$ then there is no static link (u, v) . Hence, the total workload cost of flow-paths serving the α -demands is at least $(|E_\alpha| - 1)\mu\alpha + (1 + \mu)\alpha = \kappa_\alpha + \alpha$. We have that $\kappa_\beta + \kappa_1 = 3n\mu\beta + 6n(\mu + 1)\beta + 3n(m + 1 + 2\mu) < 3n\beta + 12n\beta + 3n(3n + 3) = 15n\beta + 9n^2 + 9n < \alpha$. Hence, the total workload cost of all flow-paths of F is strictly greater than κ which yields a contradiction and the claim follows. $\square$

Call the dynamic links corresponding to the directed edges of E_α the α -dynamic links.

Claim 3. Exactly $3n$ (resp. $6n$) β -demands are served by a flow-path of workload cost $\mu\beta$ (resp. $(1 + \mu)\beta$) exactly; that is, exactly $3n$ (resp. $6n$) flow-paths serving the β -demands consist of a dynamic link (resp. a dynamic link and a static link). Hence, the total workload cost of the flow-paths serving β -demands is exactly κ_β .

Proof. All β -demands are from a unique associate clause node to an element node. As $\Delta_S = 1$, at most $3n$ dynamic links can be incident with an element node. So, at most $3n$ β -demands are served by a flow-path consisting of a solitary dynamic link. If a β -demand is served by a flow-path that does not consist of a solitary dynamic link then the flow-path has workload cost at least $(1 + \mu)\beta$.

Suppose there are less than $3n$ β -demands that are served by a flow-path consisting of a solitary dynamic link. So, the total workload cost of flow-paths serving β -demands is at least $(3n - 1)\mu\beta + (6n + 1)(1 + \mu)\beta = 9n\mu\beta + 6n\beta + \beta$. By Claim 2, $9n\mu\beta + 6n\beta + \beta \leq \kappa_\beta + \kappa_1 = 3n\mu\beta + 6n(\mu + 1)\beta + 3n(m + 1 + 2\mu)$; that is, $\beta \leq 3n(m + 1 + 2\mu)$ which yields a contradiction. So, there are exactly $3n$ β -demands served by a flow-path consisting of a solitary dynamic link.

Suppose that there is a β -demand served by a flow-path of workload cost neither $\mu\beta$ nor $(1 + \mu)\beta$. Such a flow-path has workload cost at least $\gamma\beta$ where $\gamma = \min\{2, 1 + 2\mu\}$. As before, the total workload cost of flow-paths serving β -demands is at least $3n\mu\beta + (6n - 1)(1 + \mu)\beta + \gamma\beta = 9n\mu\beta + 6n\beta + (\gamma - 1 - \mu)\beta$. By Claim 2, $9n\mu\beta + 6n\beta + (\gamma - 1 - \mu)\beta \leq \kappa_\beta + \kappa_1 = 3n\mu\beta + 6n(\mu + 1)\beta + 3n(m + 1 + 2\mu)$; that is, $(\gamma - 1 - \mu)\beta \leq 3n(m + 1 + 2\mu)$. If $\gamma = 2$ then $\beta \leq \frac{3n(m+1+2\mu)}{1-\mu}$; and if $\gamma = 1 + 2\mu$ then $\beta \leq \frac{3n(m+1+2\mu)}{\mu}$. Whichever is the case, we obtain a contradiction. Hence, there are exactly $6n$ β -demands served by a flow-path of workload cost $(1 + \mu)\beta$. Each of these flow-paths consists of a dynamic link and a static link and the claim follows. $\square$

By Claim 3, each element node x is incident via a dynamic link to a unique associate clause node of some clause c so that the element $x \in \mathcal{X}$ is in the clause $c \in C$. Consider some β -demand served by a flow-path of weight $1 + \mu$; that is, a flow-path f consisting of a static link and a dynamic link. By Claim 2, every node of $W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})$ is incident with a dynamic link incident with some other node of $W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})$. So, no node of $W \setminus (X \cup X^{\oplus 01} \cup X^{\oplus 10})$ can appear on f . Also, no associate clause node is adjacent via a static link to any other associate clause node. Hence, f must be of the form c', x', x_i , where c' is an associate clause node, of some clause c , that is adjacent via a dynamic link to the associate element node x' which is adjacent via a static link to the element node x and where the element $x \in \mathcal{X}$ is in the clause $c \in C$. In particular:

- for every $x_i \in \mathcal{X}$, if $x_i \in c_1^i, c_2^i, c_3^i$, with $c_1^i, c_2^i, c_3^i \in C$, then there are 3 dynamic links incident with a unique node of $\{x_i, x_i^{\oplus 01}, x_i^{\oplus 10}\}$ and incident with exactly one node of each of the sets of associate clause nodes of c_1^i, c_2^i and c_3^i
- if there is a dynamic link from an element node x or one of its associate element nodes to an associate clauses node of some clause c then the element $x \in \mathcal{X}$ is in the clause $c \in C$.

Call these dynamic links the β -dynamic links. Also, by Claim 3, the total workload cost of the flow-paths serving the β -demands is κ_β . Our aim now is to show that there are n clauses with the property that every associate clause node of any of these clauses is incident with a β -dynamic link incident with an element node. Doing so would result in our instance of RXC3 being a yes-instance.

Consider some flow-path f that services some 1-demand. Our preliminary aim is to show that every such flow-path f has weight at least $m + 1 + 2\mu$.

Suppose that f involves 2 dynamic links. Hence, the structure of f is s^*ds^+d or ds^+ds^* , where s (resp. d) denotes a static (resp. dynamic) link and $*$ (resp. $+$) denotes at least 0 (resp. at least 1) occurrence (in a regular-language style); recall that $\sigma = 3$.

- If f has structure s^*ds^+d then the second dynamic link must be of the form (c', x) , where c' is an associate clause node and x is an element node. If the first dynamic link is not an α - or β -dynamic link then the initial prefix of static links has weight at least m . Hence, the total weight of f is at least $m + 1 + 2\mu$. Alternatively, suppose that the first dynamic link is an α - or β -dynamic link.
 - If it is an α -dynamic link and is of the form $(z, \bar{z})$, for some $z \in V_{\leq m} \setminus P$, then the total weight of f is at least $i + \mu + 5m - i - 1 + \mu$, for some $0 \leq i \leq m$; that is, at least $5m - 1 + 2\mu > m + 1 + 2\mu$.
 - If it is an α -dynamic link and is of the form (u, v) , for some $u, w \in W$, then the total weight of f is at least $4m - 1 + \mu + 1 + \mu = 4m + 2\mu > m + 1 + 2\mu$.
 - If it is a β -dynamic link then the total weight of f is at least $3m - 1 + \mu + 1 + \mu = 3m + 2\mu > m + 1 + 2\mu$.
- If f has structure ds^+ds^* then the first dynamic link must be $(r, \bar{r})$.
 - Suppose that the first sub-path of static links has weight less than or equal to m with the second dynamic link being an α -dynamic link of the form $(\bar{z}, z)$. Then the total weight of f is at least $\mu + i + \mu + 4m - i$, for some $0 \leq i < m$; that is, $4m + 2\mu > m + 1 + 2\mu$.
 - Suppose that the first sub-path of static links has weight exactly m with the second dynamic link not an α - or β -dynamic link. As all nodes of W are already incident with a dynamic link, we must have that the weight of f is at least $\mu + m + \mu + 4 = m + 4 + 2\mu > m + 1 + 2\mu$.
 - Suppose that the first sub-path of static links has weight greater than m . The weight of f must be at least $\mu + m + 1 + \mu + 1 = m + 2 + 2\mu > m + 1 + 2\mu$.

Suppose that f involves 1 dynamic link; so, the structure of f is s^*ds^* . If the dynamic link is an α - or β -dynamic link then no matter which dynamic link this is, the weight of f is greater than $3m - 1 + \mu > m + 1 + 2\mu$. If the dynamic link is not an α - or β -dynamic link then the weight of f is at least $m + \mu + 4 > m + 1 + 2\mu$, as this dynamic link cannot be incident with any node of W . Finally, if f does not involve a dynamic link then the weight of f is at least $4m > m + 1 + 2\mu$. Hence, every flow-path serving some 1-demand has weight at least $m + 1 + 2\mu$ and when it has weight exactly $m + 1 + 2\mu$, it takes the form of a static path from r to some port node of P , augmented with a dynamic link to some neighbour of an associate clause node, augmented with a static link to the associate clause node, and augmented with a dynamic link from the associate clause node to some element node. As the total workload cost of all flow-paths serving 1-demands is at most $\kappa_1 = 3n(m + 1 + 2\mu)$, every such flow-path has weight exactly $m + 1 + 2\mu$ and is therefore of the form just described.

There are $3n$ 1-demands yet only n dynamic links from a port node to some neighbour of an associate clause node. Hence, these n neighbours of associate clause nodes must lie on $3n$ flow-paths. Consequently, these neighbours of associate clause nodes must actually be clause nodes. Let these clause nodes be $C' = \{c_{j_1}, c_{j_2}, \dots, c_{j_n}\}$. Thus, we have that there is a dynamic link from every associate clause node of each c_{j_i} to some element node and we have a subset of clauses $C' = \{c_{j_1}, c_{j_2}, \dots, c_{j_n}\}$ so that every element of $\mathcal{X}$ lies in exactly one clause of C' ; that is, $(\mathcal{X}, C)$ is a yes-instance of RXC3. The result follows as our instance (H, μ, D, κ) can clearly be constructed from $(\mathcal{X}, C)$ is time polynomial in n . $\square$

We emphasize the relevance of the hypercube as a prototypical model of interconnection networks (see, e.g., [16]) and the fact that we obtain hardness here for any choice of fixed dynamic link weight μ between 0 and 1.

Taken together our results comprehensively establish the computational hardness of RRP in practically relevant settings. In particular, we establish that the problem remains intractable in several cases where the demand matrix is sparse, the hybrid network is highly structured (in fact node-symmetric) and the weights of links depend only on their medium.

3.1. The case of $\delta = 1$

This section is devoted to the restriction of RRP where the dynamic link limit δ is set to 1 - where any flow-path must use no more than a single dynamic link. We shall require substantially different techniques from those we have used up until now so to prove our main result in the remainder of the paper, which entails that the problem 1-SWITCHED RRP($\sigma = 1$) is NP-complete for various graph classes including hypercubes, grids and toroidal grids. The intention is that this section provides a convenient template for proving hardness of the problem for other practically interesting classes beyond the ones explicitly considered here.

3.1.1. Additional definitions

We begin with some basic definitions from graph theory. Let $G = (V, E)$ be a simple undirected graph. We define the *open neighbourhood* of a vertex v to be $N(v) := \{u : (u, v) \in E\}$, and its *closed neighbourhood* $N[v] := N(v) \cup \{v\}$. Likewise for any set of vertices S , we define $N[S] := \cup_{v \in S} N[v]$ and $N(S) := N[V] \setminus S$. A *dominating set* in G is a set of vertices S such that each vertex in G is either in S or adjacent to a vertex in S . The *domination number* of G , denoted $\gamma(G)$, is the least number of vertices in any dominating set of G . DOMINATING SET is the decision problem asking, for input G and k , whether $\gamma(G) \leq k$.

Where $S \subseteq V$ is a set of vertices in the graph, we denote $G[S]$ the subgraph of G induced by S . That is, $G[S]$ has S as its set of vertices and as edges exactly those edges of G with both endpoints incident to a vertex in S . Also, $\text{dist}_G(u, v)$ is the number of edges on the shortest path between u and v in G .

Definition 1 (*Ball, Sphere*). Where G is a graph, we denote $B_G(v, r)$ (resp. $S_G(v, r)$) the *ball* (resp. *sphere*) with center v and radius r in G . Formally these are sets of vertices defined as:

$$B_G(v, r) := \{u : \text{dist}_G(u, v) \leq r\}$$

$$S_G(v, r) := \{u : \text{dist}_G(u, v) = r\}$$

We shall also make use of the following (likely non-standard) definitions.

Definition 2 (*Restriction of decision problems to a graph class*). Let Π be a decision problem which takes one or many inputs, exactly one of which is a simple undirected graph. We denote by $\Pi(\mathcal{G})$ the restriction of Π to the graph class $\mathcal{G}$. That is, $\Pi(\mathcal{G})$ has as yes-instances (resp. no-instances) exactly those yes-instances (resp. no-instances) of Π where the graph portion of the input belongs to $\mathcal{G}$.

We introduce a new graph problem, which is subtly different from the classic DOMINATING SET, and which has the property that it may remain hard even when the graph portion of the input is highly structured. This subtlety will become important presently - note that the property is precisely that which we wish to study for 1-SWITCHED RRP($\delta = 1$).

PARTIAL DOMINATION

Input: Simple undirected graph $G = (V, E)$; set $T \subseteq V$; integer k .

Question: Is $(G[T], k)$ a yes-instance of DOMINATING SET? Equivalently, is there some set $X \subseteq T$ with $|X| \leq k$ such that $T \subseteq N_G[X]$?

3.1.2. Hardness of PARTIAL DOMINATION on (toroidal) grids and hypercubes

We now show that PARTIAL DOMINATION is NP-complete for several graph classes of interest to us, namely grids and hypercubes. This computational hardness (together with another graph class property we shall come to later) provides the foundation for our proof of Theorem 8. We note that DOMINATING SET is trivially tractable for grids, though the same cannot be said of hypercubes; even $\gamma(Q_{10})$ is unknown [21]. This subsection leverages several results from the literature, which more or less straightforwardly yield the desired results.

Theorem 5 ([8] Theorem 5.1). *DOMINATING SET(Induced subgraphs of grids) is NP-complete.*

Note that the corollary below relies on some subtle properties of the proof applied in [8] (namely, that the construction provided explicitly describes an embedding into some grid). It is in general NP-hard, given some graph, to determine whether it is an induced subgraph of a grid (and also to produce its vertices' coordinates in a grid - see [18] and references therein).

Corollary 2. *PARTIAL DOMINATION(Grids) is NP-complete.*

We now turn to the other graph class of interest to us - the hypercubes Q . Although it is straightforward to show that every grid is the induced subgraph of some hypercube, we require for technical reasons that, more strongly, the hypercube in question is at most polynomially larger than the contained grid. Fortunately, we are able to rely here on the extensive literature surrounding the so-called *snake-in-the-box* problem, which consists in finding large induced cycles in hypercubes. The following result belongs to that body of work.

Theorem 6 (Abbott and Katchalski [1,2]). *For any $d \geq 2$, the hypercube Q_d contains an induced path on $\frac{77}{256}2^d$ vertices. Given d , the coordinates of such a path may be produced in time polynomial in 2^d .*

In their work, Abbott and Katchalski state that there is an induced cycle on **strictly more than** $\frac{77}{256}2^d$ vertices, which entails an induced path on exactly $\frac{77}{256}2^d$ vertices. The authors do not discuss the runtime of their construction in their work. However, it is clear from their proof that their description of the induced cycle can be realized as a recursive algorithm with a runtime as stated in the theorem. Their result has the following consequence, which shall be useful to us towards proving the computational hardness of PARTIAL DOMINATION(Q).

Corollary 3. *For any $d \geq 2$, the hypercube Q_{2d} contains an induced grid on $(\frac{77}{256}2^d)^2$ vertices. Given d , the coordinates of such a grid may be produced in time polynomial in 2^d .*

Proof. Given d , produce coordinates of a path $P = \{p_1, p_2, \dots, p_\ell\}$ of length $\ell = \frac{77}{256}2^d$ in Q_d by applying Theorem 6. Note that each (node) p_i is a bitstring of length d exactly. Then the set $\{p_i p_j : i, j \in [\ell]\}$ is an induced grid of size $(\frac{77}{256}2^d)^2$ in Q_{2d} , and the result follows. $\square$

It remains to combine the above results.

Theorem 7. *PARTIAL DOMINATION(Q) is NP-complete.*

Proof. Let (G, T, k) be an instance of PARTIAL DOMINATION(Grids). We assume without loss of generality that G is an $n \times n$ grid (if necessary, by extending G in some dimension and retaining the same set T). We shall denote each vertex in G by $v_{i,j}$ with $i, j \in [n]$. Let $d = \lceil \log_2(\frac{256n}{77}) \rceil$. Applying Corollary 3, we may efficiently produce a mapping from $V(G)$ (vertices of the $n \times n$ grid) to $V(Q_{2d})$. Denote this mapping f . Let $G' = Q_{2d}$, and $T' = \{f(t) : t \in T\}$. Then $G'[T'] = G[T]$ (so $\gamma(G'[T']) = \gamma(G[T])$ also) and (G', T', k) is an instance of PARTIAL DOMINATION(Q). The construction of G' and T' is feasible in polynomial time, and the result follows. $\square$

A reader interested in showing that some other graph class $\mathcal{G}$ (e.g. subcubic graphs) is hard for 1-SWITCHED RRP may substitute the above for a proof that PARTIAL DOMINATION($\mathcal{G}$) is NP-complete. As briefly alluded to earlier, this is one of two properties we shall require of a graph class in the proof of Theorem 8; we turn to the second property presently.

3.1.3. Lunar graph classes

We have chosen to adopt a celestial metaphor to aid intuition, since in defining the class we make use of spheres, balls, large distances, and vast differences in size. The reader may find Figs. 8 and 9 helpful in illustrating the definition.

Definition 3 (*Moon, Planet, Sun*). We say a graph G is the *moon* in some graph H if:

- There is some set $M \subset V(H)$ with $H[M] = G$. We also call M the *moon* in H .
- There is some $o \in V(H)$ and integer r such that the set $\tilde{P} := S(o, r)$ is a sphere of cardinality at least $|M|$ and $P := B(o, r)$ is the ball with the same center and radius. We call P the *planet* and $\tilde{P}$ the *planet surface* (or, briefly, *surface*) in H .
- $N[M] \cap N[P] = \emptyset$ (the planet and the moon are far apart).
- $|\tilde{P}| \geq |M|$ (the surface is bigger than the moon).
- There is some set $S \subset V(H)$ (the *sun* in H) such that $|S| \geq |N[P]| + |N[M]|$ (the sun is bigger than the moon and planet together) and $N[S]$ is disjoint from $N[P]$ (the sun and planet are far apart). Moreover any node in S is at distance at least $2r$ from any node in M (the sun and moon are very far apart).

We continue with an astronomical theme for the naming of the graph class of interest itself:

Definition 4 (*lunar graph class*). We say a graph class $\mathcal{L}$ is *lunar* in a graph class $\mathcal{H}$ if, for each $G \in \mathcal{L}$, there exists some $H \in \mathcal{H}$ such that G is the moon in H . We further require that such a graph H (and vertex sets $M, \tilde{P}, S$ within it) can be constructed in polynomial time from G (which entails that H is at most polynomially larger than G). We say a graph class $\mathcal{L}$ is *lunar in itself* (or simply *lunar*) if the above holds for $\mathcal{H} = \mathcal{L}$.

For this definition to be useful, we still need to show that it holds for those graph classes which we are interested. The previously hypothesized reader interested in proving NP-completeness of 1-SWITCHED RRP ($\delta = 1$) restricted to, e.g., subcubic graphs, may find the following a useful blueprint to prove that subcubic graphs are lunar.

Lemma 2. *The class of grid graphs is lunar in itself.*

Proof. The reader may find the illustrative example in Fig. 8 helpful. Let G be some $n \times m$ grid (w.l.o.g. $n \geq m$, so $|G|$ is polynomial in n).

We shall use the fact that a sphere of radius $r \geq 1$ in a grid (which does not spill over the grid's boundary) has size $4r$.

Let $r = \lceil \frac{nm}{4} \rceil$. Let $x = 2(n + 2r + 1)$ and let $y = 2(m + 2r)$. We choose H to be the $x \times y$ grid. Then we identify:

- The moon: $M = V(G)$ (vertices of the $n \times m$ grid are also vertices of $x \times y$ grid).
- The planet: $o = (n + r, m + r)$, $\tilde{P} = S(o, r)$, $P = B(o, r)$.
- The sun: $S = \{(i, j) : i \geq n + 2r + 3\}$.

It is easy to verify that the neighbourhoods of the moon, planet and sun are disjoint; that the sun and moon are at distance at least $2r$; that the size of the planet exceeds that of the moon; and that the size of the sun exceeds that of the moon and planet's neighbourhoods combined. Since our construction may be carried out in polynomial time, the result follows. $\square$




Fig. 8. Illustration of our construction where $n = 5, m = 3, r = 4$. The 28×11 grid contains the 5×3 grid as a moon. Marked are: the moon M (square vertices); $o = (n + r, m + r) = (7, 9)$, together with the planet $P = B(o, r)$ (o and disk and circle vertices) and its surface $\tilde{P} = S(o, r)$ (circle vertices); the sun $S = \{(i, j) : i \leq 8\}$ (cross vertices).

The $x \times y$ grid is an induced subgraph of the $2x \times 2y$ toroidal grid, with the notable property that every shortest path in the former remains a shortest path in the latter. Applying this fact, the proof above can straightforwardly be adapted to obtain the following corollary:

Corollary 4. *The class of grid graphs is lunar in the class of toroidal grid graphs.*

Note that the class of toroidal grid graphs is *not* lunar in itself. We now turn to the protagonist of this paper, the class of hypercubes.

Lemma 3. *The class of hypercubes Q is lunar in itself.*

Proof. Note that we reuse some of the notation from our proof of Theorem 4. Let $G \in Q$ be some hypercube of dimension d , i.e. $G = Q_d$ for some d . We define:

- Let $\ell = 6d + 9$ and $H = Q_\ell$.
- Let $M := \{0^{5d+9}x : x \in Q_d\}$.
- Let $o := 0^{4d+6}1^{2d+3}$ and $r = d$. Recall we denote the planet $P = B(o, r)$ and its surface $\tilde{P} = S(o, r)$.
- Let $S := \{\bar{x} : x \in N[P] \cup N[M]\}$. Intuitively, the sun nodes are the reflection of the neighbourhoods of planet nodes and moon nodes.

The illustration in Fig. 9 will be helpful in verifying the following:

- $H[M] = G$.
- The sphere $\tilde{P} = S(o, r)$ has cardinality at least $|Q_d|$ (easy to see by considering $\{o \oplus 0^{4d+9}\bar{x}x : x \in Q_d\}$, which clearly contains only vertices at distance exactly d from o and so is a subset of $\tilde{P}$).
- $N[M]$ and $N[P]$ are disjoint.
- $N[M]$ and $N[P]$ contain only vertices with at most $3d + 4$ ones.
- $|S| \geq |N[P]| + |N[M]|$ (by applying both two points above).
- $N[S]$ is disjoint from $N[M]$ and $N[P]$.
- Moreover any node in S is at distance at least $2r$ from any node in M .

Note that Q_ℓ has size $2^{6d+9} = 2^9 \cdot (2^d)^6$ which is polynomial in the size of Q_d . The result follows. $\square$

3.1.4. The main result

We are now able to state and prove the main result of this subsection.

Theorem 8. *For any fixed $\mu \in (0, 1)$, and any graph classes $\mathcal{L}$ and $\mathcal{G}$ with $\mathcal{L}$ lunar in $\mathcal{G}$, the problem $\text{PARTIAL DOMINATION}(\mathcal{L})$ is polynomially reducible to the problem $1\text{-SWITCHED RRP}(\delta = 1)$ restricted to instances (H, μ, D, κ) satisfying:*

- $H \in \mathcal{G}$, and
- the workload matrix D is sparse and all values in it are polynomial in $|H|$.




Fig. 9. Illustration of our proof hypercubes are lunar, showing the weights (number of ones in any vertex label) for different sets of vertices. The midsection of the hypercube (separating majority-1 vertices from majority-0 vertices) is shown as a dashed line.




Fig. 10. Illustration of the demands described in our reduction, for the instance $(G, T, 4)$ of PARTIAL DOMINATION(Grid graphs). Nodes of T (and T' in H) are shown as black squares; other nodes of G (and M in H) are shown as boxes. Other nodes are shown as in Fig. 8 earlier. Moonlight demands are shown as solid blue curves. The sunlit nodes $\tilde{S}$ are those in orange shaded regions (note exactly 4 nodes of $\tilde{P}$ are not sunlit - these are the moonlit nodes $\tilde{M}$). Sunlit demands are drawn as dashed orange arcs (for clarity, only a few are shown). (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

Proof. We are given an instance (G, T, k) of PARTIAL DOMINATION($\mathcal{L}$), with $G \in \mathcal{L}$, a set of vertices $T \subseteq V(G)$, and integer k . We shall produce an instance (H, μ, D, κ) of 1-SWITCHED RRP($\delta = 1$). μ is some prescribed value between 0 and 1, as in the theorem statement.

The graph H , along with vertex sets $M, \tilde{P}, P, S$, the vertex o , and the integer r , are obtained by applying the definition of a lunar class. We denote by T' the set of vertices in H to which T is mapped, so that $H[M][T] = H[T'] = G[T]$.

In order to describe our demands D , we identify some useful sets of vertices in H :

- *Target vertices* are exactly the set $T' = \{t'_1, t'_2, \dots, t'_{|T|}\}$.
- The set of moonlit nodes $\tilde{M} = \{\tilde{m}_1, \tilde{m}_2, \dots, \tilde{m}_k\}$ is some arbitrary subset of $\tilde{P}$ with cardinality k exactly.
- The set of sunlit nodes $\tilde{S} := N[M] \cup N[P] \setminus (T \cup \tilde{M})$.

The intention is that our construction will ensure that, in any configuration N of interest to us, the moonlit nodes $\tilde{M}$ (resp. sunlit nodes $\tilde{S}$) will be connected by a dynamic link to moon nodes M (resp. sun nodes S).

Consider a new graph: the complete bipartite graph with $\tilde{S} \cup S$ as vertices and $\tilde{S} \times S$ as edges. Let E_S be an arbitrary maximum matching in this bipartite graph. Note that $|S| \geq |\tilde{S}|$ by construction, and so $|E_S| = |\tilde{S}|$ exactly and each sunlit vertex is incident to exactly one edge in E_S . Note that by construction $E_S \cap E(H) = \emptyset$, because none of the sunlit nodes are adjacent to any sun node (by applying a property of lunar graphs). That is, if (u, v) is an edge in E_S then u and v are not adjacent in H .

Let $\alpha = (r + \mu + 1)|T|$. Our demand matrix D is fully described by the following:

- $D[u, v] = \alpha$ for each $(u, v) \in E_S$ (sunlight demands),
- $D[o, t] = 1$ for each $t \in T$ (moonlight demands), and
- all other entries of D are zero.

It remains for us to define κ . As before, we define this as $\kappa_\alpha + \kappa_1$, with $\kappa_\alpha = \mu\alpha|E_S|$ and $\kappa_1 = |T|(r + \mu) + |T| - k$. We note that $\alpha > \kappa_1$. This completes our construction of the instance (H, μ, D, κ) .




Fig. 11. Detail of an optimal configuration N for the instance shown in 10. Only moonlight demands (solid blue arcs) and dynamic links serving them (dashed red lines) are shown. Note that all moonlight demands are served at cost exactly κ_1 .

Claim 4. If (G, T, k) is a yes-instance of PARTIAL DOMINATION then (H, μ, D, κ) is a yes-instance of 1-SWITCHED RRP ($\delta = 1$).

Proof. An illustration of an optimal configuration is shown in Fig. 11. Let X be a dominating set of $G[T]$ of cardinality k . Denote also by $X' = \{x'_1, x'_2, \dots, x'_k\}$ the corresponding set of vertices in H .

Let $N := E_S \cup \{(x'_i, \check{m}_i) : 1 \leq i \leq k\}$. Clearly, under N each sunlight demand is served at cost exactly $\mu\alpha$ and so all sunlight demands cumulatively are served at cost $|E_S|\mu\alpha = \kappa_\alpha$.

Further, the k demands from each node x'_i to o are each served at cost exactly $\mu + r$ (by the path $x'_i \rightarrow \check{m}_i \rightsquigarrow_r o$), and the moonlight demand for each of the $|T| - k$ other nodes $t \in T$ is served at cost exactly $1 + \mu + r$ (by the path $t \rightsquigarrow_1 x'_i \rightarrow \check{m}_i \rightsquigarrow_r o$). The claim follows. $\square$

Conversely, suppose that (H, μ, D, κ) is a yes-instance of 1-SWITCHED RRP ($\delta = 1$) and that N is a configuration and F a set of flow-paths witnessing that the total workload cost is at most κ .

Claim 5. Each sunlit node has a line of sight (optic link) to a sun node. Formally, every edge (u, v) of E_S is necessarily such that $(u, v) \in N$ and the total workload cost of flow-paths serving the α -demands is exactly κ_α .

Proof. Suppose that at least one of the α -demands is served via a flow-path at a workload cost of more than $\mu\alpha$; so, it must be at a workload cost of more than $(1 + \mu)\alpha$ as we cannot traverse a dynamic link followed immediately by another dynamic link (recall, $\Delta_S = 1$) and by our construction, if $D[u, v] = \alpha$ then there is no static link (u, v) . Hence, the total workload cost of flow-paths serving the α -demands is at least $(|E_\alpha| - 1)\mu\alpha + (1 + \mu)\alpha = \kappa_\alpha + \alpha$. We have that $\kappa_1 < (r + \mu + 1)|T| = \alpha$. Hence, the total workload cost of all flow-paths of F is strictly greater than κ which yields a contradiction and the claim follows. $\square$

Claim 6. Each moonlight demand is served at cost at least $\mu + r$. Furthermore, at most k moonlight demands are served at cost $\mu + r$ exactly.

Proof. Consider some moonlight demand; necessarily it has form $D[t, o] = 1$ for some $t \in T$.

We first show each moonlight demand is served at cost at least $\mu + r$. Suppose for contradiction that $D[t, o]$ is served at cost strictly less than $\mu + r$. Since t and o are at distance at least $r + 3$, the flow-path from t to o must make use of some dynamic link $u \rightarrow v$ at cost μ . The static portion of the path therefore must have cost at most $r - 1$, which entails that $v \in B(o, r - 1) \subsetneq \check{S}$ (that is, v is a sunlit node). Applying Claim 5 v is connected by dynamic link to some sun node, so $u \in S$ necessarily. Then the path from t to u has length at least $2r$, yielding the desired contradiction.

Now observe that $D[t, o]$ can be served at cost exactly $\mu + r$ if and only if the dynamic link $t \rightarrow \check{m}$ exists, for some $\check{m} \in \check{M}$. Since $|\check{M}| = k$ exactly by construction, we obtain that this may be the case for at most k nodes in T . The claim follows. $\square$

Since $\kappa_1 = (r + \mu)(k) + (r + \mu + 1)(|T| - k)$, we immediately obtain that exactly k moonlight demands are served at cost $r + \mu$ and all remaining moonlight demands are served at cost $r + \mu + 1$ exactly.

Claim 7. The set $X = \{u : (u, v) \in N \text{ and } v \in \check{M} \text{ is a moonlit node}\}$ is a dominating set of $H[M]$.

Proof. Suppose for contradiction that there is some vertex $t \in T'$ which is neither in X nor adjacent to any vertex in X . We show that the cost of serving the demand $D[y, o] = 1$ is strictly greater than $r + \mu + 2$. Denoting arbitrary nodes $m' \in N[M]$, $p' \in N[P]$, $\tilde{p} \in \tilde{P}$, this flow is routed either:

- Via static links only, along a path of length at least $r + 3 > r + \mu + 1$: $y \rightsquigarrow_{\geq 1} m' \rightsquigarrow_{\geq 1} p' \rightsquigarrow_1 \tilde{p} \rightsquigarrow_r o$, or
- Via static links and one dynamic link $u \rightarrow v$ with:
 - v in $\tilde{P}$, at cost at least $r + \mu + 2 > r + \mu + 1$: $y \rightsquigarrow_{\geq 2} u \rightarrow v \rightsquigarrow_r o$
 - u and v both outside $N[P]$, at cost at least $r + \mu + 2 > r + \mu + 1$: $y \rightsquigarrow_{\geq 0} u \rightarrow v \rightsquigarrow_{\geq 1} p' \rightsquigarrow_1 \tilde{p} \rightsquigarrow_r o$
 - either u or v in $N[P] \setminus \tilde{P}$ (and the other in S), at cost at least $2r + \mu$ (recall all vertices in S are distance at least $2r$ from any vertex in M).

This contradicts our earlier claim (that all moonlight demands are served at cost at most $r + \mu + 1$) and the result follows. $\square$

We note that the construction described takes polynomial time, and the main result follows. $\square$

Applying our earlier results on lunar graph classes (Lemmas 3 and 2, Corollary 4) together with Theorem 7 we obtain that RRP remains hard even when the number of dynamic links admitted on any path is limited to 1.

Corollary 5. *1-SWITCHED RRP($\delta = 1$)(G) is NP-complete if G is: the class of hypercubes; the class of grid graphs; or the class of toroidal grid graphs.*

We emphasize again the value of hypercubes as a prototypical model of interconnection networks, and additionally note that grids and toroidal grids exhibit additional properties which may have been expected to yield a tractable setting, such as planarity and bounded degree. The restriction to $\delta = 1$ in our setting also implies a restriction to $\sigma = 2$ (as 3 alternations along a path would entail a minimum 2 dynamic links along the same) but the converse does not hold; we expect that the case of 1-SWITCHED RRP($\sigma = 2, \delta = 2$) can be shown to be intractable through similar proof techniques, but leave this for future work.

4. Discussion and future work

Taken together, our results comprehensively establish the computational hardness of RRP in practically relevant settings. We establish that the problem remains intractable in several cases where the demand matrix is sparse, the hybrid network is highly structured (in fact node-symmetric) and the weights of links depend only on their medium. Furthermore, in all of our hardness results, the instrument used to “express” NP-completeness is the demand matrix D . In the real world, the computational workload for the network is generally expected to vary significantly with time, unlike the network’s hardware, which (in addition to its structural properties already discussed) does not rapidly change. Our results are in this sense closely relevant to the hardness of the real world reconfigurable routing problem.

We take this opportunity to identify some specific questions we have left open, as well as several more general avenues for future work in this area. First, it would be interesting to study the restriction of the problem to cases where Δ_S is greater than 1 and μ is a fixed constant. Results in this setting would “bridge the gap” between Theorems 1 and 2, and Theorems 4 and 8. Analogously, there is a gap for 1-SWITCHED RRP on hypercubes between $\sigma = 0$ (which is solvable in polynomial time) and $\sigma = 3$ (which is an intractable case). The complexity of the problem with $\sigma = 1$ and $\sigma = 2$ remains open for hypercubes (note that results for arbitrary networks do exist when $\sigma = 2$, as shown in Table 1). Our Theorem 8 in some sense sits between these two open cases, since the restriction to $\sigma = 1$ entails a restriction to $\delta = 1$, which itself entails a restriction to $\sigma = 2$.

Secondly, the present work considers only exact computation. In [11] the authors establish inapproximability within $\Omega(\log n)$ for RRP in a more permissive setting (making use of variable link weights). However, the empty solution (there are no dynamic links and all demands are routed through the static network only) is a $\frac{\log n}{\mu}$ -approximation for Δ_S -SWITCHED RRP on hypercubes. (This follows straightforwardly from hypercubes having logarithmic diameter.) It would be interesting to see what (in)approximability results can be derived in our model with fixed link weights, with and without restrictions to realistic topologies.

Lastly, parameterized algorithms may provide more fine-grained insights into the computational complexity of reconfigurable routing. Our Theorems 1 and 2 establish that structural parameters of the static network, such as treewidth, are insufficient to yield fixed-parameter tractable (fpt) algorithms (unless $P=NP$). However, it would be interesting to see whether it is possible to obtain an fpt algorithm by additionally parameterizing by the sum of the demand matrix D ; some structural parameters for the digraph representation of the demands, D' ; the dynamic link weight μ ; or a combination of these.

CRediT authorship contribution statement

David C. Kutner: Conceptualization, Investigation, Writing – original draft, Writing – review & editing. **Iain A. Stewart:** Conceptualization, Investigation, Supervision, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] Harvey L. Abbott, M. Katchalski, On the construction of snake in the box codes, *Util. Math.* 40 (1991) 97–116.
- [2] Harvey L. Abbott, Meir Katchalski, On the snake in the box problem, *J. Comb. Theory, Ser. B* 45 (1) (1988) 13–24.
- [3] Mohammad Al-Fares, Alexander Loukissas, Amin Vahdat, A scalable, commodity data center network architecture, *Comput. Commun. Rev.* 38 (2008) 63–74.
- [4] Chan Avin, Stefan Schmid, Toward demand-aware networking: a theory for self-adjusting networks, *Comput. Commun. Rev.* 48 (2019) 31–40.
- [5] Piotr Berman, Marek Karpinski, Approximation hardness of bounded degree MIN-CSP and MIN-BISECTION, in: *Proc. of 29th Int. Colloq. on Automata, Languages and Programming (ICALP)*, 2002, pp. 623–632.
- [6] Charidimos Chaintoutis, Behnam Shariati, Adonis Bogris, Paul V. Dijk, Chris G.H. Roeloffzen, Jerome Bourderionnet, Ioannis Tomkos, Dimitris Syvridis, Free space intra-datacenter interconnects based on 2d optical beam steering enabled by photonic integrated circuits, *Photonics* 5 (3) (2018).
- [7] Tao Chen, Xiaofeng Gao, Chen Guihai, The features, hardware, and architectures of data center networks: a survey, *J. Parallel Distrib. Comput.* 96 (2016) 45–74.
- [8] Brent N. Clark, Charles J. Colbourn, David S. Johnson, Unit disk graphs, *Discrete Math.* 86 (1–3) (1990) 165–177.
- [9] L.H. Clark, R.C. Entringer, The bisection width of cubic graphs, *Bull. Aust. Math. Soc.* (1988) 389–396.
- [10] Thomas Fenz, Klaus-Tycho Foerster, Stefan Schmid, Anaïs Villedieu, Efficient non-segregated routing for reconfigurable demand-aware networks, in: *Proc. of IFIP Networking Conf.*, IEEE Press, 2019, pp. 1–9.
- [11] Thomas Fenz, Klaus-Tycho Foerster, Stefan Schmid, Anaïs Villedieu, Efficient non-segregated routing for reconfigurable demand-aware networks, *Comput. Commun.* 164 (2020) 138–147.
- [12] Klaus-Tycho Foerster, Manya Ghobadi, Stefan Schmid, Characterizing the algorithmic complexity of reconfigurable data center architectures, in: *Proc. of Symp. on Architectures for Networking and Communications Systems (ANCS)*, ACM Press, 2018, pp. 89–96.
- [13] Klaus-Tycho Foerster, Maciej Pacut, Stefan Schmid, On the complexity of non-segregated routing in reconfigurable data center architectures, *Comput. Commun. Rev.* 49 (2019) 2–81.
- [14] Klaus-Tycho Foerster, Stefan Schmid, Survey of reconfigurable data center networks: enablers, algorithms, complexity, *ACM SIGACT News* 50 (2019) 62–79.
- [15] Teofilo F. Gonzalez, Clustering to minimize the maximum intercluster distance, *Theor. Comput. Sci.* (1985) 293–306.
- [16] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, Songwu Lu, BCube: a high performance, server-centric network architecture for modular data centers, *Comput. Commun. Rev.* 39 (2009) 63–74.
- [17] Chuanxiong Guo, Haitao Wu, Kun Tan, Lei Shi, Yongguang Zhang, Songwu Lu, DCell: a scalable and fault-tolerant network structure for data centers, in: *Proc. of ACM SIGCOMM Conf. on Data Communication*, 2008, pp. 75–86.
- [18] Siddharth Gupta, Guy Sa’ar, Meirav Zehavi, Grid recognition: classical and parameterized computational perspectives, *J. Comput. Syst. Sci.* 136 (2023) 17–62.
- [19] Matthew Nance Hall, Kalsu-Tycho Foerster, Stefan Schmid, Ramakrishnan Durairajan, A survey of reconfigurable optical networks, *Opt. Switching Netw.* 41 (2021) 100621.
- [20] Navid Hamedazimi, Zafar Qazi, Himanshu Gupta, Vyas Sekar, Samir R. Das, Jon P. Longtin, Himanshu Shah, Ashish Tanwer, Firefly: a reconfigurable wireless data center fabric using free-space optics, in: *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM ’14, Association for Computing Machinery, New York, NY, USA, 2014, pp. 319–330.
- [21] OEIS Foundation Inc., Entry A000983 in the on-line encyclopedia of integer sequences, Published electronically at <https://oeis.org/A000983>, 2024.
- [22] Igor Pak, Radoš Radoičić, Hamiltonian paths in Cayley graphs, *Discrete Math.* 309 (17) (2009) 5501–5508, Generalisations of de Bruijn Cycles and Gray Codes/Graph Asymmetries/Hamiltonicity Problem for Vertex-Transitive (Cayley) Graphs.
- [23] Ankit Singla, Chi-Yao Hong, Lucian Popa, P. Brighten Godfrey, Jellyfish: networking data centers randomly, in: *Proc. of 9th USENIX Conf. on Networked Systems Design and Implementation*, 2012, pp. 225–238.
- [24] Asaf Valadarsky, Gal Shahaf, Michael Dinitz, Michael Schapira, Xpander: towards optimal-performance datacenters, in: *Proc. of 12th Int. Conf. on Emerging Networking Experiments and Technologies*, 2016, pp. 205–219.
- [25] Shaojuan Zhang, Xuwei Xue, Eduward Tangdionga, Nicola Calabretta, Low-latency optical wireless data-center networks using nanoseconds semiconductor-based wavelength selectors and arrayed waveguide grating router, *Photonics* 9 (3) (2022).