

MULTITASKING CAPACITY: HARDNESS RESULTS AND IMPROVED CONSTRUCTIONS*

NOGA ALON[†], JONATHAN D. COHEN[‡], THOMAS L. GRIFFITHS[‡], PASIN MANURANGSI[§], DANIEL REICHMAN[¶], IGOR SHINKAR^{||}, TAL WAGNER[#], AND ALEXANDER YU^{††}

Abstract. We consider the problem of determining the maximal $\alpha \in (0, 1]$ such that every matching M of size k (or at most k) in a bipartite graph G contains an induced matching of size at least $\alpha|M|$. This measure was recently introduced in [N. Alon et al., *Adv. Neural Inf. Process. Syst.*, 2017, pp. 2097–2106] and is motivated by computational models in cognitive neuroscience as well as by modeling interference in radio and communication networks. We prove various hardness results for computing α either exactly or approximately. En route to our results, we also consider the maximum connected matching problem: determining the largest matching N in a graph G such that every two edges in N are connected by an edge. We prove a nearly optimal $n^{1-\varepsilon}$ hardness of approximation result (under randomized reductions) for connected matching in bipartite graphs (with both sides of cardinality n). Toward this end we define bipartite half-covers: a new combinatorial object that may be of independent interest. To our knowledge, the best previous hardness result for the maximum connected matching problem was that it is hard to approximate within some constant $\beta > 1$. Finally, we demonstrate the existence of bipartite graphs with n vertices on each side of average degree d , achieving $\alpha = 1/2 - \varepsilon$ for matchings of size sufficiently smaller than n/d . This nearly matches the trivial upper bound of $1/2$ on α which holds for any graph containing a path of length 3.

Key words. matching, connected matchings, hardness of approximation, wireless networks

AMS subject classification. 68R05

DOI. 10.1137/18M1224672

1. Introduction. A matching in an undirected graph G is a set of vertex disjoint edges. An induced matching M in G is a matching such that no two edges in M are connected by another edge in G . Matchings and induced matchings can be used to measure the capacity of a parallel network of processors. Here we study computational and combinatorial aspects of such a measure [2] arising from radio and wireless networks as well as computational neuroscience. Some of our findings build on a new hardness of approximation result concerning the problem of computing a maximal connected matching in a bipartite graph (see Definition 2.3 for a formal definition) which may be of independent interest.

*Received by the editors November 5, 2018; accepted for publication (in revised form) December 16, 2019; published electronically March 24, 2020. The U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. Copyright is owned by SIAM to the extent not limited by these rights.

<https://doi.org/10.1137/18M1224672>

Funding: The work of the first author was partially supported by NSF grant DMS-1855464, BSF grant 2018267, and the Simons Foundation.

[†]Tel Aviv University, Tel Aviv, Israel and Princeton University, Princeton, NJ 08544 (nogaa@tau.ac.il).

[‡]Princeton University, Princeton, NJ 08544 (jdc@princeton.edu, tomg@princeton.edu).

[§]Google Research. Work done while a Ph.D. student at the University of California, Berkeley, CA 94720 (pasin@berkeley.edu).

[¶]Worcester Polytechnic Institute, Worcester, MA 01609 (daniel.reichman@gmail.com).

^{||}Simon Fraser University, Burnaby, British Columbia, Canada V5A 1S6 (ishinkar@sfu.ca).

[#]MIT CSAIL, Cambridge, MA 02139 (talw@mit.edu).

^{††}Princeton University, Princeton, NJ 08544 (vengencemage@gmail.com).

When using matchings to study parallel and distributed systems, the object of study is oftentimes a set of units that transmit or receive information. For example, in the communication setting there is a bipartite network $G = (S, R, E)$ consisting of senders (S) and receivers (R).¹ Given a set of edges $E' = \{(s_i, r_i)\}_{1 \leq i \leq \ell} \subseteq E$ every sender $s_i \in S$ wishes to send a message to its neighbor $r_i \in R$. The assumption is that each sender s_i sends a message to a unique receiver, and in order for a receiver to successfully receive a message, he can have only a *single* incident edge carrying a message at a given time, as messages arriving on multiple incident edges create interference with each other. This is captured by a condition which we term the *matching condition*: A subset $E' \subset E$ can be used for concurrent interference-free communication if it forms a matching in G , i.e., no two edges in E' share a common vertex. However, in several communication settings, such as radio and wireless networks [4, 8, 1], a more constrained setting is considered: the senders cannot choose which edges to broadcast, but instead, if they choose to transmit, then they automatically broadcast on *all* their incident edges. This leads to the stronger *induced matching condition*: A subset $E' \subset E$ of edges can be used for concurrent interference-free communication if it forms an *induced* matching in G .

Similar interference assumptions, directed toward understanding multitasking constraints in neural systems, have been proposed in computational neuroscience [9, 11, 20, 19]. These works seek to understand the reason behind multitasking limitations: The limited ability of people to execute control-dependent processes concurrently, a central and ubiquitous finding in cognitive psychology [24]. Inspired by the parallel distributed processing framework [23], the main idea in these works is that such limitations arise from interference between computational units responsible for transmitting inputs to outputs and not, as commonly assumed, because of limited resources. [11, 20, 19] present a formal model to study multitasking where given a bipartite graph $G = (S, T, E)$ (task graph), every vertex $s \in S$ is associated with a set of inputs I_s , every vertex $t \in T$ is associated with a set of outputs O_t , and the edge (s, t) is associated with a function (“task”) $f_{s,t} : I_s \rightarrow O_t$. Every function $f_{s,t}$ is implemented by a neural network $N_{s,t}$. Given a set of ℓ edges $E' = (s_i, t_i)_{1 \leq i \leq \ell}$ the set of functions f_{s_i, t_i} can be performed concurrently (“multitasked”) if E' is an induced matching. The rationale for the matching condition for interference-free parallel processing is similar to the exclusive read exclusive write (EREW) assumption in parallel RAM: If the set of edges is not a matching, problems may occur as two different values may be stored simultaneously in the same output vertex in T . Alternatively, if two different tasks share the same input vertex in S , this may inhibit independent processing of these tasks as the input to both tasks has to be identical. The rationale for the induced matching assumption arises from the idea that if two tasks (s_1, t_1) and (s_2, t_2) (with $s_1 \neq s_2$ and $t_1 \neq t_2$) are performed, then if the tasks (s_1, t_2) or (s_2, t_1) exist they are performed automatically as well, interfering with computing f_{s_1, t_1} and f_{s_2, t_2} . We refer to [11, 20, 19, 2] for further study and justifications of this interference model. Finally we comment that this model assumes a selection mechanism which selects at a given moment which set of tasks (edges) are to be performed (see, for example, [20, 19]).

Based on these interference assumptions [11] suggested using the cardinality of a maximal induced matching in G to measure the parallel processing capacity of a task graph G . One potential issue with this definition is that a task graph can contain a “large” induced matching even though some subset of tasks allows for very poor

¹To simplify matters, we consider the synchronous setting where transmissions occur in discrete time slots.

multitasking. For example, consider a graph H consisting of two bipartite cliques $H_1 = K_{m,m}$ and $H_2 = K_{m,m}$, with H_1 connected to H_2 by an induced matching of cardinality $2m$. While H contains an induced matching of size $2m$, there are sets of m tasks (edges) out of which only a single edge can be executed without interference. Furthermore, by taking a perfect matching in H_1 one can find such a set of m “bad edges” that forms a matching in H .

In [2] a new measure has been proposed to capture how well task graphs allow for interference-free processing (Definition 1.1 below). The idea behind this measure is to consider a parameter $k \leq n$ and ask whether *every* matching M of size k (or of size at most k) contains a large *induced* matching $M' \subseteq M$. By considering every matching this measure is no longer agnostic to subgraphs that are “badly multitaskable,” such as bipartite cliques. Unless stated otherwise we will always assume that graphs are bipartite and that both sides of the bipartition have cardinality n .

DEFINITION 1.1. *Let $G = (A, B, E)$ be a bipartite graph, and let $k \in \mathbb{N}$ be a parameter. For $\alpha \in (0, 1]$ we say that G is a (k, α) -multitasker if for every matching M in G of size $|M| = k$, there exists an induced matching $M' \subseteq M$ such that*

$$|M'| \geq \alpha|M|.$$

Define $\alpha_k(G)$ to be the maximal α such that G is a (k, α) -multitasker if G contains a matching of size k , and define $\alpha_k(G) = 1$ if G does not contain a matching of size k . We call the parameter $\alpha_k(G) \in (0, 1]$ the multitasking capacity of G for matchings of size k .

Also, define $\alpha_{\leq k}(G) = \min_{1 \leq \ell \leq k} \alpha_\ell(G)$ and call it the multitasking capacity of G for matchings of size at most k .²

The parameters α_k and $\alpha_{\leq k}$ measure how resilient to interference G is. The larger these parameters are, the better G is considered as a multitasker. One motivation for this definition is the distinction between interference effects that result from a violation of the matching condition to those that result from a violation of the *induced* matching condition. That is, the above multitasking measure allows us to assess the fraction of tasks that can be performed concurrently conditioned on not violating the matching condition. We omit the dependence of α on G when it is clear from the context.

In [2] several properties of $\alpha_{\leq n}(G)$ have been proven. For example, it was shown that $\alpha_{\leq n}(G) \leq \frac{9}{\sqrt{d}}$ for all d -regular graphs and that $\alpha_{\leq n}(G) \leq O((\frac{\log n}{d})^{1/3})$ for all graphs of average degree d . This upper bound supports a previous hypothesis [11] suggesting that there is an *inherent tradeoff* between density and multitasking capacity: for every task graph as the average degree diverges to infinity, the multitasking capacity inevitably decrease to 0.³ Observe that such a tradeoff does *not* hold when multitasking capacity is defined as the maximum cardinality of induced matching (normalized by the number of vertices of the graph) as a graph can have average degree $\Omega(n)$ and still contain an induced matching of size $\Omega(n)$. Finally, it was also shown in [2] how to construct graphs with desirable multitasking properties, namely, graphs for which $\alpha_{\leq k}(G) \geq \tau$ for $\tau = \Omega(1)$ provided that $k = O(n/d^{1+\tau})$, where d is the average degree of G .

The results in [2] leave several questions.

²Since we consider the minimum, the definition of α_k ensures that values of $r \leq k$ for which there is no matching of size r have no influence on $\alpha_{\leq k}(G)$.

³In the irregular case this tradeoff holds assuming the average degree satisfies $d \gg \log n$.

QUESTION 1.2. *Given a graph G and a parameter k , can we compute $\alpha_k(G)$ or $\alpha_{\leq k}(G)$ efficiently?*

Indeed, if we are to use $\alpha_k(G)$ or $\alpha_{\leq k}(G)$ to evaluate how prone to interference parallel architectures are, then a natural question is whether it is possible to compute or approximate these quantities in polynomial time. Hence to evaluate the usefulness of $\alpha_k(G)$ in graph-theoretic models of multitasking it is desirable to have efficient methods to compute $\alpha_k(G)$ exactly or approximately.

Another question is whether it is possible to construct multitaskers with near-optimal capacity. While [2] provides graphs with average degree d and $\alpha_{\leq k}(G) = \Omega(1)$ for $k \leq n/d^{O(1)}$, the best constant value of $\alpha_{\leq k}(G)$ they achieve is bounded away from the natural barrier $\alpha_{\leq k}(G) \leq 1/2$ (if a network contains a path of length 3, then trivially $\alpha_{\leq k}(G) \leq 1/2$ for all $k \geq 3$). We thus raise the following question.

QUESTION 1.3. *Is there an infinite family of graphs G_n of average degree d such that $\alpha_{\leq k}(G_n) \geq 1/2 - \varepsilon$ for arbitrarily small $\varepsilon > 0$ and $k \geq n/d^{f(\varepsilon)}$ for some function $f > 0$?*

Here we address these two questions. For Question 1.2 we show that under standard complexity theoretic assumptions $\alpha_k(G)$ and $\alpha_{\leq k}(G)$ cannot be computed efficiently, thus giving a negative answer to this question. Toward this end we give new hardness of approximation results for computing the size of a maximum connected matching (Definition 2.3) in bipartite graphs. For Question 1.3 we give a positive answer, by showing how to construct bipartite graphs G_n (each side in G_n has n vertices) of average degree d such that $\alpha_{\leq k}(G_n) \geq 1/2 - \varepsilon$ for arbitrarily small $\varepsilon > 0$ and $k \geq n/d^{f(\varepsilon)}$ for some function $f > 0$. Our proof is algorithmic in the sense that given any matching M in G_n with $|M| \leq k$ there is a simple polynomial time algorithm that recovers an induced matching $M' \subseteq M$, where $|M'| \geq (1/2 - \varepsilon)|M|$. We note that while for *regular* graphs it is proven in [2] that if $k \gg n/d$, then $\alpha_{\leq k} = o_d(1)$, we do not know if one can achieve $\alpha_{\leq k}$ bounded away from 0 for $k \gg n/d$ for irregular graphs and arbitrary⁴ d . Resolving this question is left for future work. Finally we show that $\alpha_k(G) = 1/2$ enforces k to be logarithmic in n (assuming the average degree of G is larger than 8): Namely, for every bipartite graph G with average degree $d > 8$ there is a matching M in G of size $O(\log n / \log d)$ that does not contain an induced matching of size $|M|/2$.

1.1. Our results. A useful notion in studying the computational hardness of computing the multitasking capacity is that of a *connected matching*, which is a matching in which every two edges are connected by a third edge. Connected matchings have been studied in several contexts, such as Hadwiger's conjecture [16, 22, 12]. Motivated by applications to other optimization problems [15], algorithms for finding connected matchings of maximum cardinality have been studied in special families of graphs such as chordal [6] and bipartite chordal graphs [15]⁵ and bipartite permutation graphs [13].

In section 3 we establish hardness of approximation for the size of the largest connected matching to within a factor of $n^{1-\varepsilon}$ assuming $\mathcal{NP} \neq \text{coRP}$. Previously, this problem was known to be \mathcal{NP} -hard to approximate within some constant factor [22] for general (nonbipartite) graphs. We also prove that deciding whether a bipartite graph $G = (A, B, E)$ with $|A| = |B| = n$ contains a connected matching of size n is \mathcal{NP} -hard.

⁴For $d = O(\log \log n)$ it is shown in [2] that there are graphs of average degree d with $\alpha_{\leq n}(G) \geq 1/3$.

⁵Observe that bipartite chordal graphs are not necessarily chordal. See [15] for details.

TABLE 1

Hardness of approximation results for computing the multitasking capacity. In each row, the stated variant of the multitasking capacity (either $\alpha_k(G)$ or $\alpha_{\leq k}(G)$) is hard to approximate under the stated assumption up to a multiplicative factor f , for the stated values of k and f .

Variant	Assumption	k	f (approximation factor)	Remarks
$\alpha_k(G)$	$\mathcal{P} \neq \mathcal{NP}$	n	$n^{1-\varepsilon}$ for any $\varepsilon > 0$	
$\alpha_k(G)$	$\mathcal{P} \neq \mathcal{NP}$	n	$O(d/\text{poly log } d)$	G has maximum degree d
$\alpha_{\leq k}(G)$	$\mathcal{NP} \neq \text{coRP}$	n	some constant	
$\alpha_{\leq k}(G)$	$\mathcal{NP} \neq \text{coRP}$	$n^{1-\varepsilon}$	arbitrarily large constant	
$\alpha_{\leq k}(G)$	ETH	$n^{1-1/\text{polyloglog}(n)}$	$n^{1/\text{polyloglog}(n)}$	

In section 4 we prove several hardness results for computing the multitasking capacity. To be more precise, we define the decision problem of computing the multitasking capacity as follows.

DEFINITION 1.4. Let MT be the problem of deciding whether for a given graph G , a positive integer $k \in \mathbb{N}$, and a rational number $\eta > 0$ it holds that $\alpha_k(G) \geq \eta$.

The problem MT belongs to the second level of the polynomial hierarchy, Π_2 , since the statement $\alpha_k(G) \geq \eta$ can be expressed as for all $M \exists M' P(G, k; M, M')$, where P is the predicate checking that M is a matching in G of size k , and M' is an induced matching (P is clearly computable in time $\text{poly}(|G|)$.) We note that it is not clear whether it belongs to \mathcal{NP} or to coNP , and in fact, we give evidence that MT belongs to neither of the classes. Specifically, we show that MT is both \mathcal{NP} -hard and coNP -hard; thus, if $\text{MT} \in \mathcal{NP} \cup \text{coNP}$, then the polynomial hierarchy collapses to the first level.

Furthermore, we show various hardness of approximation results for computing $\alpha_k(G)$ and $\alpha_{\leq k}(G)$. Most notably, we show under standard complexity theoretic assumptions that (1) $\alpha_n(G)$ is inapproximable to within $n^{1-\varepsilon}$ for any $\varepsilon > 0$, and, (2) $\alpha_{\leq k}(G)$ is inapproximable to within any constant for $k = n^{1-\varepsilon}$ for any $\varepsilon > 0$. Furthermore, under a stronger assumption, we improve the inapproximability ratio for $\alpha_{\leq k}(G)$ to $n^{1/\text{polyloglog}(n)}$ for $k = n^{1-1/\text{polyloglog}(n)}$. Our hardness results are summarized in Table 1.

In section 6, we prove the existence of multitaskers with near-optimal capacity. For integers d, n with $n \geq d$ and $\varepsilon \in (0, 1)$, we show how to construct multitasker graph G on $2n$ vertices with average degree d and $\alpha_{\leq k}(G) \geq 1/2 - \varepsilon$, where $k = \Omega(n/d^{1+O(1/\varepsilon)})$. In particular, for $d = n^{o(1)}$ this implies that ε can be taken to be $o(1)$, and thus $\alpha_{\leq k}(G)$ tends to its natural barrier $1/2$ as n grows.

1.2. Our techniques.

Hardness results. With respect to multitasking, connected matchings are the worst possible multitasking configuration for a matching of size k . In particular, it holds trivially that $\alpha_k(G) \geq 1/k$ and $\alpha_{\leq k}(G) \geq 1/k$, and the equality holds if and only if G contains a connected matching of size k . This fact, together with the extremal Ramsey-type bound on the size of independent sets, turns out to be instrumental in proving hardness results for computing the multitasking capacity.

Construction of multitaskers. The starting point of our multitaskers with nearly optimal multitasking capacity is based on locally sparse graphs, similarly to [2]. They used the local sparsity with Turan’s lower bound on independent sets in graphs with a given average degree in order to establish the existence of sufficiently large independent sets (which translate to induced matchings). However, the use of Turan’s bound necessarily entails a constant loss, which makes the final multitasking capacity

bounded away from $1/2$. We circumvent this roadblock by also requiring that the graph has large girth and use this fact along with local sparsity in order to carefully construct for any matching M a matching $M' \subseteq M$ of size $(1/2 - \varepsilon)|M|$.

2. Preliminaries. All graphs considered in this work are undirected. A matching in a graph $G = (V, E)$ is a collection $M \subseteq E$ of vertex disjoint edges. We say that a vertex $v \in V$ is *covered* by M if it is one of the endpoints of an edge in M . We say that a matching M is *induced in G* if no two edges in M are connected by an edge in E , i.e., the vertices in M span only the edges in M and no other edges. Given a graph G and an edge $e = (u, v) \in E$, we define the *contraction* of e to be the operation that produced the graph $G \setminus e$, whose vertex set is $(V \cup v_e) \setminus \{u, v\}$, the vertex v_e is connected to all vertices in G neighboring u or v , and for all other vertices $x, y \in V \setminus \{u, v\}$, they form an edge in $G \setminus e$ if and only if they were connected in G . Contracting a set of edges, and in particular contracting a matching, means contracting the edges one by one in an arbitrary order.⁶ A connected graph G has radius r if r is the minimal number such that there exists a vertex v with every vertex in G of distance at most r from v .

Below we define two combinatorial optimization problems that we will relate to when proving hardness of approximation results for the parameters α_k and $\alpha_{\leq k}$.

DEFINITION 2.1. *Given an undirected graph G , an independent set in G is a set of vertices that spans no edges. The maximum independent set problem (MIS) is the problem of finding a maximum cardinality of an independent set in G .*

DEFINITION 2.2. *Given a graph $G = (V, E)$, we say that two disjoint subsets of the vertices $A, B \subseteq V$ form a bipartite clique (biclique) in G if $(a, b) \in E$ for all $a \in A$ and $b \in B$. We say that the biclique (A, B) is balanced if $|A| = |B|$. In the maximum balanced biclique problem we are given a bipartite graph G and a parameter k , and the goal is to decide whether G contains a balanced biclique with k vertices on each size.*

DEFINITION 2.3. *Given a graph G , a connected matching in G is a matching M such that every two edges in M are connected by an edge in G . We use $\nu_c(G)$ to denote the size of the maximum cardinality of a connected matching in G . In the connected matching problem, we are given graph G and parameter k and our goal is to determine whether $\nu_c(G) \geq k$.*

Given an optimization (minimization or maximization) problem Π over graphs, we denote by $OPT_{\Pi}(G) > 0$ the value of the optimal solution of Π for G . An algorithm A for a maximization (minimization) problem is said to achieve an approximation ratio $\rho > 1$ if for every input G the algorithm returns a solution $A(G)$ such that $OPT_{\Pi}(G) \geq A(G) \geq OPT_{\Pi}(G)/\rho$ (resp., $OPT_{\Pi}(G) \leq A(G) \leq \rho \cdot OPT_{\Pi}(G)$).

We assume familiarity with complexity classes such as \mathcal{NP} , $\text{co}\mathcal{NP}$, $\text{co}\mathcal{RP}$, Π_2 , and the polynomial-time hierarchy. Precise definitions of these terms are omitted and can be found, e.g., in [21].

3. Hardness results for maximum connected matchings. In this section, we prove hardness results for finding large connected matchings in graphs.

3.1. Hardness of approximating the size of a maximum connected matching. We start by showing an almost optimal hardness of approximation result for the connected matching problem.

⁶We remark that the graph obtained from contracting a set of edges, indeed, does not depend on the order.

THEOREM 3.1. *Given a bipartite graph G with n vertices on each side, it is \mathcal{NP} -hard to approximate $\nu_c(G)$ within a factor of $n^{1-\varepsilon}$ for any $\varepsilon > 0$ under a randomized polynomial time reduction.*

More precisely, given a bipartite graph G with n vertices on each side, it is \mathcal{NP} -hard to distinguish between the case where $\nu_c(G) \geq n^{1-\varepsilon}$ and the case where $\nu_c(G) \leq n^\varepsilon$ for any $\varepsilon > 0$.

A natural approach to prove hardness of approximation results for connected matching is to reduce the clique problem to it. Namely, given a graph $G = ([n], E_G)$ for which we wish to determine if G contains a k -clique, replace every vertex i by an edge $e_i = (u_i, v_i)$ and add two edges (u_i, v_j) and (u_j, v_i) for every edge (i, j) in G . Call the resulting graph after these transformations G' . While it is clear that a large clique in G translates to a large connected matching in G' , it is not clear that a large connected matching in G' implies a large clique in G . The difficulty is that a connected matching might contain “bad” edges of the form (u_i, v_j) , where $i \neq j$. An illustrative example is the case where $G = K_{n/2, n/2}$ is a biclique; in this case, the largest clique in G has size only 2 but the resulting graph G' contains a large connected matching of size as large as n .

To overcome this problem, we first observe that instead of adding both (u_i, v_j) and (u_j, v_i) to the graph G' for every edge (i, j) in G , it suffices to add only one of the two to retain a large connected matching in the YES case. Then, the insight is that, when we choose the edge to add independently at random for each (i, j) , we can control the number of bad edges in every connected matching in G' .

We formalize the described ideas below, starting with the main gadget of our reduction.

DEFINITION 3.2. *Fix $n \in \mathbb{N}$. A bipartite graph $HC_n = (A = \{u_1, \dots, u_n\}, B = \{v_1, \dots, v_n\}, E_H)$ is said to be a bipartite half-cover of K_n if (1) for every $\{i, j\} \subseteq [n]$, $(u_i, v_j) \in E_H$, or $(u_j, v_i) \in E_H$ and (2) for every $i \in [n]$, $(u_i, v_i) \notin E_H$.*

The reduction used in the proof of Theorem 3.1 uses the existence of such bipartite half-covers of K_n that do not contain a large connected matching. Such graphs can be easily constructed using a randomized algorithm as shown below.

CLAIM 3.3. *There is an $O(n)$ -time randomized algorithm that on input $n \in \mathbb{N}$ outputs a graph HC_n , which is a bipartite half-cover of K_n such that $\nu_c(HC_n) \leq O(\log n)$ with probability $1 - o(1)$.*

Proof. We construct HC_n by choosing for each $\{i, j\} \subseteq [n]$ to add to E_H either (u_i, v_j) or (u_j, v_i) independently with probability $1/2$. Clearly, HC_n is a bipartite half-cover of K_n . Below we show that $\nu_c(H) \leq O(\log n)$ with probability $1 - o(1)$. We prove this in two steps: first, we will prove the $O(\log n)$ upper bound on a special class of connected matching and, then, we will show that any connected matching contains a large (constant fraction) matching of this type.

Let $M \subseteq E_H$ be any matching in H . We say that the matching is *nonrepetitive* if, for each $i \in [n]$, at most one of u_i or v_i appears in M . We will now argue that with probability $1 - o(1)$, any connected nonrepetitive matching has size less than $D := 20 \log n$. To do so, consider any ordered tuple $(i_1, j_1, \dots, i_D, j_D)$, where $i_1, \dots, i_D, j_1, \dots, j_D$ are all distinct. The probability that $(u_{i_1}, v_{j_1}), \dots, (u_{i_D}, v_{j_D})$ is

a connected matching is at most

$$\begin{aligned} & \Pr[\forall 1 \leq k < \ell \leq D, (u_{i_k}, v_{j_\ell}) \in E_H \vee (u_{i_\ell}, v_{j_k}) \in E_H] \\ &= \prod_{1 \leq k < \ell \leq D} \Pr[(u_{i_k}, v_{j_\ell}) \in E_H \vee (u_{i_\ell}, v_{j_k}) \in E_H] \\ &= \prod_{1 \leq k < \ell \leq D} (3/4) = (3/4)^{D(D-1)/2}, \end{aligned}$$

where the first two equalities use the fact that $i_1, \dots, i_D, j_1, \dots, j_D$ are distinct, meaning that the events considered are all independent. Hence, by union bound over all such sequences, we can conclude that the probability that H contains a connected non-repetitive matching of size D is at most $n^{2D} \cdot (3/4)^{D(D-1)/2} = (n^2 \cdot (3/4)^{(D-1)/2})^D = o(1)$.

Finally, observe that any matching $M \subseteq E_H$ contains a nonrepetitive matching $M' \subseteq M$ of size at least $|M|/3$. Indeed, given a matching M we can construct M' iteratively by picking an arbitrary edge $e = (u_i, v_j) \in M$, remove e and all edges touching v_i or u_j from M , and add e to M' . We repeat this procedure until $M = \emptyset$. Since we add one edge to M' while removing at most three edges from M , we arrive at a nonrepetitive $M' \subseteq M$ of size at least $|M|/3$. As a result, the graph HC_n does not contain any connected matching of size at least $3D = O(\log n)$ with probability $1 - o(1)$. \square

Remark 1.

1. We remark that a deterministic polynomial time construction of such graphs would imply that the hardness result in Theorem 3.1 holds under a deterministic reduction (as opposed to the randomized reduction, currently stated).
2. We comment that there is a connection between Ramsey graphs and half-cover of K_n with small $\nu_c(HC_n)$. Specifically, if we can deterministically construct half-cover for K_n with $\nu_c(HC_n) \leq f(n)$, then we can deterministically construct n -vertex $(f(n) + 1)$ -Ramsey graphs. This is because we can think of half-cover HC_n as a bichromatic K_n , where (i, j) for $i < j$ is colored red if $(u_i, v_j) \in E_H$ and it is colored blue otherwise (i.e., $(u_j, v_i) \in E_H$). It is easy to check that any monochromatic clique of size r in K_n implies a connected matching of size $r - 1$ in HC_n . While there are explicit constructions of Ramsey graphs, it is unclear (to us) how to construct such half-cover from these constructions.
3. Using a different approach we can show that it is \mathcal{NP} -hard to compute $\nu_c(G)$ under a *deterministic* reduction. See Appendix A for details.

3.1.1. Proof of Theorem 3.1. With the gadget from Claim 3.3 we are ready to prove Theorem 3.1. This is done in the following claim.

CLAIM 3.4. *Let $G = (V_G = [n], E_G)$ be an n -vertex graph, and let $H = (A = \{u_1, \dots, u_n\}, B = \{v_1, \dots, v_n\}, E_H)$ be a balanced bipartite graph. Let $G \boxplus H = (A, B, E_{G \boxplus H})$ be the balanced bipartite graph with n vertices on each side, where (1) for every $\{i, j\} \subseteq [n]$, $(u_i, v_j) \in E_{G \boxplus H}$ if and only if $(u_i, v_j) \in E_H$ and $(i, j) \in E_G$, and (2) for every $i \in [n]$, $(u_i, v_i) \in E_{G \boxplus H}$.*

Then, for any such G we have $\nu_c(G \boxplus H) \leq \omega(G) + 3\nu_c(H)$, where $\omega(G)$ denotes the clique number of G . Furthermore, if H is a bipartite half-cover of K_n , then $\omega(G) \leq \nu_c(G \boxplus H)$.

Claim 3.4 immediately implies Theorem 3.1. Indeed, by [14, 27] given an n -vertex graph G it is NP-hard to decide between the case where $\omega(G) \geq n^{1-\varepsilon/2}$ and the case

where $\omega(G) \leq n^{\epsilon/2}$. Therefore, we can define a randomized reduction that given an n -vertex graph G constructs (with high probability) HC_n , the bipartite half-cover of K_n , with $\nu_c(HC_n) \leq O(\log n)$, and outputs $G \boxplus H$, which can be clearly constructed in time that is linear in the size of G . In the YES case, if $\omega(G) \geq n^{1-\epsilon/2}$, then by the “furthermore” part of Claim 3.4 we have $\nu_c(G \boxplus HC_n) \geq \omega(G) \geq n^{1-\epsilon/2}$, and in the NO case, if $\omega(G) \leq n^{\epsilon/2}$, then by Claim 3.4 we have $\nu_c(G \boxplus HC_n) \leq \omega(G) + \nu_c(HC_n) \leq n^{\epsilon/2} + O(\log n)$. This completes the proof of Theorem 3.1.

We now turn to the proof of Claim 3.4.

Proof of Claim 3.4. First, we will show that $\nu_c(G \boxplus H) \leq \omega(G) + 3\nu_c(H)$. Let $M \subseteq E_{G \boxplus H}$ be any connected matching in $G \boxplus H$. We partition M into two disjoint sets M_{\parallel} and M_{\times} , where $M_{\parallel} = M \cap \{(u_i, v_i) \mid i \in [n]\}$ and $M_{\times} = M \setminus M_{\parallel}$. We will show that $|M_{\parallel}| \leq \omega(G)$ and $|M_{\times}| \leq 3\nu_c(H)$.

To show that $|M_{\parallel}| \leq \omega(G)$, suppose that $M_{\parallel} = \{(u_{i_1}, v_{i_1}), \dots, (u_{i_t}, v_{i_t})\}$. By the definition if (u_i, v_i) is connected to $(u_{i'}, v_{i'})$ in $G \boxplus H$, then $(i, i') \in E_G$. Therefore, $\{i_1, \dots, i_t\}$ induces a clique in G and $\omega(G) \geq t = |M_{\parallel}|$ follows.

Next, we show that $|M_{\times}| \leq 3\nu_c(H)$. Let us first define nonrepetitive matching in the same way as in the proof of Claim 3.3. Using the same argument as in that proof, we can conclude that M_{\times} contains a nonrepetitive connected matching $M'_{\times} \subseteq M_{\times}$ of size at least $|M_{\times}|/3$. We claim that M'_{\times} is also a connected matching in H . Indeed, since every edge in M'_{\times} belongs to E_H , the nonrepetitiveness implies that any pair of edges in M'_{\times} is connected by an edge that also belongs to E_H . As a result, we can conclude that $|M_{\times}| \leq 3|M'_{\times}| \leq 3\nu_c(H)$.

Combining the above two bounds yields $\nu_c(G \boxplus H) \leq \omega(G) + 3\nu_c(H)$ as desired.

Finally, assume that H is a bipartite half-cover of K_n . For any clique $C \subseteq V_G$ in G , it is not hard to see that the matching $M_C = \{(u_i, v_i) : i \in C\}$ is a connected matching in $G \boxplus H$. Indeed, for each distinct $i, j \in C$ we have either $(u_i, v_j) \in E_H$ or $(u_j, v_i) \in E_H$ (from definition of bipartite half-cover of K_n), and hence either (u_i, v_j) or (u_j, v_i) belongs to $E_{G \boxplus H}$. Therefore, $\nu_c(G \boxplus H) \geq \omega(G)$, which completes our proof. \square

3.2. Hardness of finding a connected perfect matching. In this section we show that given a bipartite graph G with n vertices on each side, it is \mathcal{NP} -hard to find a connected matching of size n .

THEOREM 3.5. *Given a bipartite graph $G = (A, B, E)$ with $|A| = |B| = n$ it is \mathcal{NP} -hard to determine whether $\nu_C(G) = n$.*

Proof. By Theorem 3.1 given a graph $G = (A, B, E_G)$ with N vertices of each side it is \mathcal{NP} -hard to decide whether G contains a connected matching of size $k = N^{1-\epsilon}$. Consider the reduction that given a graph $G = (A, B, E_G)$ outputs $H = (A \cup A', B \cup B', E_H)$ as follows. The sets A' and B' are two disjoint sets that are also disjoint from A, B with $|A'| = |B'| = N - k$. The set of edges E_H is defined as $E_H = E_G \cup \{(i, j) : i \in A', j \in B \cup B'\} \cup \{(i, j) : i \in A \cup A', j \in B'\}$. That is, the graph H contains the graph G as the induced graph on the vertices $A \cup B$, and in addition, every vertex in A' is connected to all vertices in $B \cup B'$, and every vertex in B' is connected to all vertices in $A \cup A'$.

The graph H is a balanced bipartite graph with $n = 2N - k$ vertices on each side. We claim that $\nu_C(G) = k$ if and only if $\nu_C(H) = n$.

In one direction, suppose that G has a connected matching $M_G = \{e_1, \dots, e_k\}$ of size k . We construct a matching M' of size $2N - k$ as follows. For each vertex $v \in A \cup B$ not covered by M_G , we pick a distinct element $w_v \in A' \cup B'$ that is a neighbor of v .

Define a matching in H to be $M' = M \cup N$, where $N = \{(v, w_v) : v \in V(G) \setminus V(M_G)\}$. By the construction of H , each edge in N is connected to every other edge in M' using an edge between A' and B' . Every pair of edges in M_G are connected since M_G is a connected matching in G . Thus, M' is a connected matching of size n in H .

Conversely, suppose H has a connected matching M_H of size n . Then, there must be a submatching $M \subseteq M_H$ of size $|M| = k$ such that no edge in M contains a vertex in $A' \cup B'$. Thus, M is a matching in G , and since M_H is a connected matching so is M . It follows that G has a connected matching of size k , as required. \square

4. Hardness results for computing $\alpha_k(G)$. In this section we study the computational complexity both of the decision problem MT as well as the problem of computing $\alpha_k(G)$ exactly or approximately. We first show an almost optimal inapproximability result for $\alpha_n(G)$, which is stated and proved below.

THEOREM 4.1. *For any $\varepsilon > 0$, given a bipartite graph G with n vertices in each part, it is \mathcal{NP} -hard to approximate $\alpha_n(G)$ within a factor $n^{1-\varepsilon}$.*

Furthermore, given a bipartite graph G with n vertices in each part, where the degree of each vertex is at most d it is \mathcal{NP} -hard to approximate $\alpha_n(G)$ within a factor $O(\frac{d}{\log^4(d)})$ and it is \mathcal{UG} -hard to approximate $\alpha_n(G)$ within a factor $O(\frac{d}{\log^2(d)})$.

Proof. The proof is by a reduction from the MIS problem. Given an n vertex graph $H = (U_H, E_H)$ instance of the MIS we construct a bipartite graph G as follows. Denote the vertices of H by $U_H = \{u_1, u_2, \dots, u_n\}$. Then the vertices of the bipartite graph $G = (V_G = A \cup B, E_G)$ are defined by $A = \{v_i : i \in [n]\}$ and $B = \{v'_i : i \in [n]\}$, and the edges of G are $E_G = \{(v_i, v'_i) : i \in [n]\} \cup \{(v_i, v'_j) : i < j \wedge (u_i, u_j) \in E_H\}$. Note that the only perfect matching in G , i.e., a matching of size n , is the matching $N = \{(v_i, v'_i) : i \in [n]\}$. Indeed, suppose there exists another matching M with $|M| = n$. Then M has at least one edge of the form $e = (v_i, v'_j)$ with $i < j$ and suppose that e is such that i is minimal (where the minimum is taken with respect to all edges not in N). If any edge in M covers v'_i , then it cannot belong to N as M is a matching. By the definition of E_G there cannot be an edge in M that covers v'_i by the minimality of i . As all vertices of H must be matched in order for $|M| = n$, we get a contradiction showing that N is indeed the unique matching of size n .

We claim that H contains an independent set of size at least α if and only if $\alpha_n(G) \geq \frac{\alpha}{n}$. Indeed, a set $I \subseteq V_H$ is an independent set in H if and only if $M' = \{(v_i, v'_i) : i \in I\}$ is an induced matching contained in M . Hence if H contains an independent set of size α , then M contains an induced matching of size α . Conversely, if M contains an induced matching of size α , then H has an independent set of size α . It is well known that for any $\delta < 1/2$ it is \mathcal{NP} -hard to distinguish between n -vertex graphs that contain an independent set of size at least $n^{1-\delta}$ (YES case) and graphs that do not contain an independent set of size at least n^δ (NO case) [14, 27]. By the reduction described above it is \mathcal{NP} -hard to distinguish between a bipartite graph G' with sides of cardinality n satisfying $\alpha_n(G') \geq n^{1-\delta}/n = n^{-\delta}$ to a graph G'' satisfying $\alpha_n(G'') \leq n^\delta/n = n^{\delta-1}$ as this would enable distinguishing between the YES and NO cases described above. The result now follows by taking δ to equal $\varepsilon/2$.

The result for graphs of maximum degree d follows by noting that if the maximal degree of H is at most d , then the maximal degree of G is upper bounded by $d + 1$. Therefore, since it is \mathcal{NP} -hard to approximate MIS in graphs of maximum degree d within a factor of $O(\frac{d}{\log^4(d)})$ [7] and \mathcal{UG} -hard to approximate MIS in graphs of maximum degree d within a factor of $O(\frac{d}{\log^2(d)})$ [3], the analogous hardness computing α_n also follows. \square

We remark that by adding isolated vertices to the graph, the above hardness result also implies hardness of approximating $\alpha_k(G)$ to within factor of $k^{1-\varepsilon}$ for every $\varepsilon > 0$ and every $k \geq n^\delta$ for any constant $\delta \in (0, 1)$.

Recall the decision problem MT from Definition 1.4. As mentioned in the introduction, MT clearly belongs to the class Π_2 . We show the following.

THEOREM 4.2. *The decision problem MT is \mathcal{NP} -hard and $\text{co}\mathcal{NP}$ -hard.*

Proof of Theorem 4.2. By Theorem 4.1 it follows that there is a reduction from any problem in \mathcal{NP} that produces a graph G and a parameter $k = n$ such that in the YES case $\alpha_k(G) \geq 1/n^\varepsilon$, and in the NO case $\alpha_k(G) \leq 1/n^{1-\varepsilon}$. In particular, this implies that MT is \mathcal{NP} -hard.

In order to prove that MT is $\text{co}\mathcal{NP}$ -hard we use Theorem 3.5. Indeed, observe that $\alpha_n(G) \leq 1/n$ if and only if G contains a connected matching of size n , and hence there is a reduction from any problem in \mathcal{NP} that produces a graph G and $k = n$ such that in the YES case $\alpha_k(G) \leq 1/k$, and in the NO case $\alpha_k(G) \geq 2/k$. This completes the proof of Theorem 4.2. \square

Using Theorem 3.5, we demonstrate that it is unlikely that MT belongs to $\mathcal{NP} \cup \text{co}\mathcal{NP}$.

COROLLARY 4.3. *If the decision problem MT belongs to $\mathcal{NP} \cup \text{co}\mathcal{NP}$, then the polynomial-time hierarchy collapses to the first level.*

Indeed, this follows from the fact that if $\mathcal{NP} \subseteq \text{co}\mathcal{NP}$, then $\mathcal{NP} = \text{co}\mathcal{NP}$ (see, e.g., [21, Proposition 10.2]), and hence the polynomial hierarchy collapses to the first level.

We end this section with several remarks.

1. Note that the proof of Theorem 4.1 shows that the problem of computing $\alpha_n(G)$ is \mathcal{NP} -hard on graphs with n vertices on each side even if G contains a unique perfect matching.
2. Note also that the hardness result in Theorem 4.1 for bounded degree graphs cannot hold for d regular graphs (as opposed to graphs with degree at most d) unless $P = \mathcal{NP}$. This is because in [2] it is shown that $\alpha_n(G) \leq O(1/\sqrt{d})$ for every d -regular graph G . In particular, this implies that it is easy to approximate $\alpha_n(G)$ within a factor of $O(\sqrt{d})$ for d -regular graphs.

5. Hardness results for computing $\alpha_{\leq k}(G)$. Here we prove that it is hard to calculate the parameter $\alpha_{\leq k}(G)$.

5.1. Hardness results for computing $\alpha_{\leq n}(G)$. We first consider the $k = n$ case.

THEOREM 5.1. *Given a bipartite graph $G = (A, B, E)$ with $|A| = |B| = n$, it is \mathcal{NP} -hard to compute $\alpha_{\leq n}(H)$.*

Proof. It is immediate that $\alpha_{\leq n}(H) \geq 1/n$ and that equality holds if and only if H contains a connected matching of size n . The theorem follows from Theorem 3.5. \square

We proceed and consider approximating $\alpha_{\leq n}(G)$.

THEOREM 5.2. *Unless $\mathcal{NP} = \text{co}\mathcal{RP}$, there is no polynomial algorithm for approximating $\alpha_{\leq n}(H)$ within some constant factor.*

Proof. We first use the fact that it is \mathcal{NP} -hard to distinguish between n -vertex graphs with cliques of size $b \cdot n$ to graphs with no clique of size $a \cdot n$, where a, b are some constants satisfying $1/2 < a < b < 1$. Indeed it is well known that there are

$a, b \in (0, 1)$ such that it is \mathcal{NP} -hard to distinguish between n -vertex graphs with cliques of size $b \cdot n$ and graphs with no clique of size $a \cdot n$ (e.g., [14]). The fact now follows by taking a graph G of n vertices, adding to it a clique of size n , and connecting all vertices in this clique to all vertices of G .

Given a graph G apply the reduction in Claim 3.4 (with H being the random graph described in Claim 3.3) and call the resulting graph G' . If there is a clique G of size $b \cdot n$, then clearly $\alpha_{\leq n}(G') \leq \frac{b}{n}$. Suppose there is no clique of size $a \cdot n$ in G . Then by Claim 3.3, with high probability there is no connected matching in G' of size greater than $(a + \delta) \cdot n$, where $\delta > 0$ can be taken to be arbitrarily small. It follows that for $c > a + \delta$, every connected matching in G contains a induced matching of size at least 2. Therefore, for $(a + \delta) < c < 1$ we have that conditioned on the existence of a matching of size k , $\alpha_k(G') = \frac{2}{cn} > \frac{1}{(a+\delta) \cdot n}$. Indeed, $\frac{2}{c} > \frac{1}{a+\delta}$ as $a + \delta > 1/2$. As for $k < (a + \delta)n$ it clearly holds that $\alpha_k(G') > \frac{1}{(a+\delta)n}$, and we have that in this case $\alpha_{\leq n}(G') = \frac{a+\delta}{n}$. This implies that approximating $\alpha_{\leq n}(H)$ within a ratio smaller than $\frac{b}{a+\delta}$ in polynomial time would allow one to determine whether G contains a clique of size $b \cdot n$ or no clique of size $a \cdot n$. Taking δ such that $\frac{b}{a+\delta} > 1$ concludes the proof. \square

5.2. Hardness results for computing $\alpha_{\leq k}(G)$ for $k < n$. We now turn to the problem of proving hardness of approximation results for $\alpha_{\leq k}(G)$ for $k < n$; for certain values of k , we show that $\alpha_{\leq k}(G)$ is \mathcal{NP} -hard to approximate to within any constant factor under randomized reduction. One approach to prove this is to use the reduction in Theorem 4.1. However, this approach does not seem to work, as it allows one to consider also matchings that contain “diagonal edges” of the form (u_i, v'_j) and it is not clear how to apply the analysis in Theorem 4.1 to such matchings. Instead, we build upon the hardness of the connected matching problem given in Theorem 3.1. We claim that the reduction in Theorem 3.1 shows that it is hard to approximate $\alpha_{\leq k}(G)$ for $k = n^{1-\varepsilon}$. Note that in the YES case, if $\nu_c(G) = k \geq n^{1-\varepsilon}$, then $\alpha_{\leq k}(G) = 1/k$. The NO case is a bit subtle, and it is, a priori, not clear why $\nu_c(G) \leq n^\varepsilon$ implies that any matching of size at most k contains a large induced matching. We resolve this problem using the following Ramsey-theoretic fact (see, e.g., [5, 10]).

FACT 5.3. *Let G be an n -vertex graph not containing a clique of size $k + 1$ and suppose $k \geq 2 \log n$. Then G contains an independent set of size at least $s = \log n / \log(k / \log n)$.*

Coupled with Theorem 4.1 we prove the following result.

THEOREM 5.4. *For any constants $\varepsilon \in (0, 1/2)$ and $\rho > 1$, it is \mathcal{NP} -hard (under randomized reduction) to approximate $\alpha_{\leq k}(G)$ within a factor of ρ on bipartite graphs with n vertices on each side for $k = n^{1-\varepsilon}$.*

Proof. By Theorem 3.1 given a bipartite graph G it is \mathcal{NP} -hard to distinguish between the case where $\nu_c(G) \geq n^{1-\varepsilon}$ and the case where $\nu_c(G) \leq n^\delta$ for $\delta = 1/(2\rho)$.

For the YES case if $\nu_c(G) \geq n^{1-\varepsilon}$, then clearly $\alpha_{\leq k}(G) = 1/k$ for $k = n^{1-\varepsilon}$.

In the NO case suppose that $\nu_c(G) \leq n^\delta$, and consider an arbitrary matching M of size s with $s \leq k$. If $s < 2\delta k$, then clearly M contains an induced matching of size at least $s/(2\delta k)$. Otherwise, contract all edges in M . Denote by $H(M)$ the subgraph induced by the s contracted vertices. Observe that a subset of vertices in $H(M)$ forms a clique if and only if their corresponding edges in G form a connected matching. Otherwise, by the assumption that $\nu_c(G) \leq n^\delta$ we get that $H(M)$ contains no clique of size n^δ . Hence, by Fact 5.3 we conclude that $H(M)$ contains an independent set of size at least $\frac{\log s}{\log(n^\delta / \log s)} \geq \frac{1}{2\delta}$ (assuming n is sufficiently large).

Therefore, given a bipartite graph G with n vertices on each side, and $k = n^{1-\epsilon}$ it is \mathcal{NP} -hard to distinguish between the YES case of $\alpha_{\leq k}(G) \leq 1/k$ and the NO case of $\alpha_{\leq k}(G) \geq 1/(2\delta k) = \rho/k$. This concludes the proof. \square

We can achieve stronger hardness results under stronger assumptions than \mathcal{NP} -hardness. Recall that the exponential time hypothesis (ETH) postulates that no algorithm of running time $2^{o(n)}$ can decide whether an n -variable conjunctive normal form formula has a satisfying assignment. Assuming ETH we have the following hardness result.

THEOREM 5.5. *Assuming ETH there exists a k such that given $H = (A, B, E)$ with $|A| = |B| = n$ there is no polynomial time algorithm that approximates $\alpha_{\leq k}(H)$ within a factor of $n^{(1/\log \log n)^c}$, where $c > 0$ is a universal constant independent of n .*

We will rely on the following simple lower bound on independent sets in graphs of average degree d_{avg} due to Turan.

LEMMA 5.6. *Every n -vertex graph with average degree d_{avg} contains an independent set of size at least $\frac{n}{d_{avg}+1}$.*

Proof of Theorem 5.5. It is known [18] that assuming ETH for $k = n^{1-1/\text{polyloglog}(n)}$ there is no polynomial algorithm that distinguishes between the case where H contains a bipartite clique with t vertices on each side (YES case) from the case where every subgraph contained in H with $k' \leq k$ vertices satisfies $|E(H)| \leq \binom{k'}{2}/n^{(1/\log \log n)^c}$ (NO case). In the first case $\alpha_{\leq k}(H) = 1/k$. In the second case, given a matching M with $|M| = k$ and $k' \leq k$ we claim that M contains an induced matching of size $\Omega(\max((k'n^{-(1/\log \log n)^c}, 1))$. The claim is trivially true if $k' \leq n^{(1/\log \log n)^c}$, and hence assume $k' > n^{(1/\log \log n)^c}$. Let $H(M)$ be the graph induced on M and let $H'(M)$ be the graph obtained after all edges in M are contracted. Clearly the average degree of $H'(M)$ is $O(k'n^{-(1/\log \log n)^c})$ (see [2, Lemma 2.1]), and hence by Lemma 5.6 it contains an independent set I' of size $\Omega(n^{(1/\log \log n)^c})$. It is easily verified that this independent set corresponds to an induced matching contained in M whose size is $\Omega(n^{(1/\log \log n)^c})$. Therefore every matching of size at most $k' \leq k$ contains an induced matching of size $\Omega(\lceil k'n^{-(1/\log \log n)^c} \rceil)$, which implies that $\alpha_{\leq k}(H) = \Omega(n^{(1/\log \log n)^c}/k)$. It follows that if we could approximate $\alpha_{\leq k}(H)$ within a factor better than $\Omega(n^{(1/\log \log n)^c})$ in polynomial time, then we could distinguish between the YES and NO cases described above. This concludes the proof. \square

6. Improved construction of multitaskers. In this section we prove the following theorem.

THEOREM 6.1. *Let $d \leq n$ be positive integers such that n is sufficiently large, and let $\epsilon \in (0, 1)$ be such that $\epsilon \geq \frac{20 \log d}{\log n}$. Then, there is a bipartite graph G with n vertices on each side and average degree at least $d/2$, such that $\alpha_{\leq k}(G) \geq 1/2 - \epsilon$ for $k = (\frac{1}{101e^\epsilon})^{4/\epsilon} \cdot \frac{n}{d^{1+8/\epsilon}} = \frac{n}{d^{1+O(1/\epsilon)}}$.*

For the proof of Theorem 6.1 we need the following lemma. We remark that a similar result also appears in [2] (proof of Theorem 4.14 in the arXiv version).

LEMMA 6.2. *Let $G = (A, B, E)$ be a balanced bipartite graph, and let g be the girth of G . Let $t \in \mathbb{N}$ be such that for every subset of vertices $T \subseteq A \cup B$ satisfying $|T \cap A| = |T \cap B| \leq s \leq t$ it holds that $|E(T)| \leq (2 + \beta/g)s$ edges for some $\beta > 0$. Then $\alpha_{\leq t}(G) \geq \frac{1}{2} - \frac{1+\beta}{g}$.*

Proof. Let $G = (A, B, E)$ with $|A| = |B| = n$ that satisfies the assumptions in the lemma and let M be a matching in G of size $s \leq t$. We show that M contains an induced matching M' of size at least $(\frac{1}{2} - \frac{1+\beta}{g})|M|$.

Let F be the graph whose vertices correspond to the s edges of M , and two vertices in F are connected if the corresponding edges are connected by an edge in G . We show below that F contains an independent set on nearly half of its vertices. By the assumptions of the claim, the girth of F is at least $g/2$, and any set of s of its vertices spans at most $(1 + \beta/g)s$ edges. Construct an independent set in F as follows. As long as F contains a vertex of degree at most 1 add it to the independent set and omit it and its unique neighbor from F . Suppose that this process stops with h vertices. This implies that the independent set so far has at least $(s - h)/2$ vertices. If $h = 0$, we are done, as the independent set has at least $s/2$ vertices. Otherwise, in the induced subgraph of F on the remaining h vertices the minimum degree is at least 2 and the average degree is at most $2 + 2\beta/g$. Hence it contains at most $2\beta h/g$ vertices of degree at least 3. Omit these vertices. The remaining graph is a union of paths and cycles, which may contain odd cycles, but all cycles in it are of length at least $g/2$. Therefore this part contains an independent set of size at least $\frac{1}{2}(1 - 2\beta/g) \cdot (1 - 2/g)h$, which together with the $(s - h)/2$ vertices obtained in the initial process results in an independent set of size at least

$$\begin{aligned} \frac{s-h}{2} + \frac{1}{2}(1-2\beta/g) \cdot (1-2/g)h &> \frac{s-h}{2} + \frac{1}{2}(1-2\beta/g-2/g)h \\ &> \frac{s}{2} - \frac{1+\beta}{g}h \geq \left(\frac{1}{2} - \frac{1+\beta}{g}\right)s, \end{aligned}$$

as required. \square

We can now prove Theorem 6.1.

Proof. We start with a random bipartite graph G' with n vertices on each side, in which each edge is included independently with probability $p = d/n$. The following two claims prove the properties required in order to apply Lemma 6.2.

CLAIM 6.3. *Let g be an even integer such that $2/\varepsilon \leq g \leq 4/\varepsilon$. Then, with probability $1 - \frac{2}{n^{0.3}} \geq 0.99$ the number of cycles of length at most g is upper bounded by \sqrt{n} .*

Proof. The expected number of cycles of length up to g is upper bounded by

$$\sum_{s=2}^{g/2} \binom{n}{s}^2 (s!)^2 p^{2s} \leq \sum_{s=2}^{g/2} (np)^{2s} \leq \sum_{s=2}^{2/\varepsilon} d^{2s} \leq 2d^{4/\varepsilon}.$$

In particular, for $\varepsilon \geq \frac{20 \log d}{\log n}$ the expected number of cycles of length up to g is at most $2d^{4/\varepsilon} \leq 2n^{1/5}$. The claim follows by Markov's inequality. \square

CLAIM 6.4. *With probability 0.99, every subgraph of G' with at most $(\frac{1}{101e^5})^{4/\varepsilon} \cdot n/d^{1+8/\varepsilon}$ vertices on each side has average degree at most $(2 + \varepsilon/4)$.*

Proof. Let s be an integer satisfying $1 \leq s \leq (\frac{1}{101e^5})^{4/\varepsilon} \cdot n/d^{1+8/\varepsilon}$. By the union bound over all subsets of G' with s vertices on each side, the probability that G' contains a balanced subgraph with s vertices on each side and average degree at least

$(2 + \varepsilon/4)$ is

$$\begin{aligned} \binom{n}{s}^2 \binom{s^2}{(2 + \varepsilon/4)s} p^{(2+\varepsilon/4)s} &\leq \left(\frac{n\varepsilon}{s}\right)^{2s} \cdot (s\varepsilon)^{(2+\varepsilon/4)s} \cdot \left(\frac{d}{n}\right)^{(2+\varepsilon/4)s} \\ &\leq \left(\frac{e^5 d^{2+\varepsilon/4} s^{\varepsilon/4}}{n^{\varepsilon/4}}\right)^s \leq \left(\frac{1}{101}\right)^s. \end{aligned}$$

By taking the union bound over all values of s we get that the probability that G' contains a dense induced subgraph is at most $\sum_{s=1}^{\infty} \left(\frac{1}{101}\right)^s = 0.01$, as required. \square

By Chernoff bound with probability 0.99, G' contains at least $0.9dn$ edges. Therefore, with probability 0.97 the latter event occurs together with the events in the two foregoing lemmas.

Let $g \in [\frac{2}{\varepsilon}, \frac{4}{\varepsilon}]$ be an even integer, as in Claim 6.3. We remove an edge from each cycle of length at most g , thus removing at most \sqrt{n} edges, so that the average degree remains at least $d/2$. The resulting graph G satisfies the conditions of Lemma 6.2 with $g \in [\frac{2}{\varepsilon}, \frac{4}{\varepsilon}]$ and $t = (\frac{1}{101e^5})^{4/\varepsilon} \cdot n/d^{1+8/\varepsilon}$, and hence $\alpha_t(G) \geq 1/2 - 2/g \leq 1/2 - \varepsilon$, as required. This concludes the proof of Theorem 6.1. \square

Remark 2. We note that if we consider $\alpha_{\leq n}(G)$ instead of $\alpha_{\leq n/d^{1+O(1/\varepsilon)}}(G)$, then for the construction in the proof of Theorem 6.1 it holds that $\alpha_{\leq n}(G) = O(\frac{\ln d}{d} + O(1/\sqrt{n}))$ with high probability. Indeed, it can be shown that prior to deletions G' has a matching of size $\Omega(n)$ and no induced matching of size larger than $O(\frac{\ln d}{d}n)$ with high probability. Therefore, since removing \sqrt{n} edges can increase the size of any induced matching by at most \sqrt{n} , we get that the entire construction satisfies $\alpha_{\leq n}(G) = O(\frac{\ln d}{d} + 1/\sqrt{n})$.

6.1. Is $\alpha_k(G) = 1/2$ attainable? The foregoing positive result obtains $\alpha_{\leq k}(G) = 1/2 - \varepsilon$ for k as large as $\Theta(n/d^{1+O(1/\varepsilon)})$, approaching the natural barrier $1/2$. One may wonder whether $1/2$ can be attained exactly, and for which values of k . We now show the following limitation.

THEOREM 6.5. *There is an absolute positive constant d_0 such that for $n \geq d \geq d_0$ and $k \geq \log n / \log d + O(1)$, every graph G with n vertices on each side and average degree d has $\alpha_{\leq k}(G)$ strictly smaller than $1/2$. This is tight up to the leading constant 1 as there is a graph G' with n vertices on each side and average degree d so that for $k' = 0.4999 \log n / \log d$, $\alpha_{\leq k'}(G) = 1/2$.*

Note that by Theorem 6.1, for a fixed constant d there are graphs G with n vertices in each side and average degree d for which $\alpha_{\leq k}(G)$ is arbitrarily close to $1/2$ for $k = \Omega(n)$. The theorem above shows, however, that even for a logarithmic k , exactly $1/2$ is not attainable.

For the proof we need two results. The first is the following theorem of Verstraëte.

THEOREM 6.6 (Verstraëte [26]). *Let $r \geq 2$ be a natural number and let G be a bipartite graph of average degree at least $4r$ and girth at least g . Then there exist cycles of $(g/2 - 1)r$ consecutive even lengths in G . Moreover, the shortest of these cycles has length at most twice the radius of G .*

The second is a special case of a result of Kostochka and Pyber [17]; see [26] for its simple proof.

LEMMA 6.7. *Let G be a graph on n vertices with at least $bn^{1+1/t}$ edges, where $b \geq 1$. Then G contains a subgraph of average degree at least b and radius at most t .*

We proceed with the proof of Theorem 6.5.

Proof. The tightness, proved in [2], is simply the known existence of a d -regular bipartite graph G' with n vertices on each side and girth at least $\log n / \log d$. Indeed, any set vertices which are the endpoints of at most k' edges spans a forest in G' and any matching M in a forest contains an induced matching of size at least $|M|/2$.

To prove the main part of the theorem observe that one obstacle to obtaining $\alpha_{\leq k}(G) = 1/2$ is the existence of a short cycle of length 2 modulo 4. Indeed, consider a cycle of length $\ell = 2k$, where k is an odd integer. It is straightforward to check that picking every other edge of the cycle yields a matching M of size k , in which the largest induced matching contained in M has size $(k-1)/2 = (\frac{1}{2} - \frac{1}{2k})|M|$. Hence, a graph G containing such a cycle has $\alpha_{\leq k}(G)$ strictly less than $1/2$. It thus suffices to show that every bipartite graph with n vertices on each side and average degree $d > 8$ contains such a cycle for $k = (1 + o(1)) \log n / \log d + 1$.

Let G be such a graph. Let t satisfy $8n^{1+1/t} = dn$. Note that such t exists as long as $d > 8$. By solving for t we get $t = \frac{\log n}{\log d - 3} = (1 + o(1)) \log n / \log d$. By Lemma 6.7 with $b = 8$, G contains a subgraph with average degree at least 8 and radius at most t . By Theorem 6.6 with $r = 2$ and $g = 4$, this subgraph contains cycles of two consecutive even lengths, both of length at most $2t + 2 = 2(1 + o(1)) \log n / \log d + 2$. One of two consecutive even lengths must equal 2 modulo 4. Thus, G contains a cycle of length $\ell \leq 2(1 + o(1)) \log n / \log d + 2$ with $\ell \equiv 2 \pmod{4}$. As explained above, this implies that for $k = \ell/2$, $\alpha_{\leq k}(G) < 1/2$, completing the proof. \square

7. Conclusion and future directions. We have studied the computational complexity of computing $\alpha_k(G)$, a parameter that arises in radio networks and models of multitasking in cognitive neuroscience. We find it noteworthy that two independent lines of research lead to similar models of interference hinting that ideas from radio networks can be useful for theoretical neuroscience and vice versa. Furthermore, our study reveals that algorithmic as well as combinatorial questions (such as the existence of graphs with certain combinatorial properties) are relevant to connectionist models of cognition. We hope that future work will reveal more connections between such models, theoretical computer science, and combinatorics.

While we have shown that computing $\alpha_k(G)$ is intractable, our results do not rule out the existence of an efficient constant factor approximation algorithm for $\alpha_{\leq n}(G)$, which could potentially be used in computer simulations and in analyzing behavioral and neuroscientific data. Whether such an algorithm exists is an interesting direction for future study.

Our multitasking model (notably the induced matching condition) assumes that performing several tasks in parallel can have detrimental effects, when the different tasks interfere destructively with one another. This can arise if they draw on a shared set of representations that must be put to competing uses at the same time. For example, interference effects such as those in [11] arise when incongruent stimuli must be processed, which demands different, competing responses. However, it is important to recognize that performing tasks in parallel does not necessarily lead to interference and can even have a positive effect. This is the case if the tasks share representations that favor the same or compatible responses (that is, they benefit by constructive interference). Such mutual positive interactions between interactive parallel processes have been demonstrated in [25].

We conclude with several specific questions arising from this work.

- We believe that for d -regular graphs the upper bound $\alpha_{\leq n}(G) \leq 9/\sqrt{d}$ is not tight. It is an open problem whether for all d -regular graphs it holds that

$\alpha_{\leq n}(G) \leq o(1/\sqrt{d})$. In particular, it is consistent with our knowledge that $\alpha_{\leq n}(G) = O(\frac{\log d}{d})$ holds for all d -regular graphs.

- It would be interesting to see if the $n^{1-\epsilon}$ hardness of approximation result for the size of a largest connected matching can be obtained assuming $\mathcal{P} \neq \mathcal{NP}$ (that is, under a deterministic reduction). In particular, it would be interesting to find efficient and deterministic constructions of bipartite half-covers with maximal connected matching upper bounded by $n^{o(1)}$.
- For d -regular graphs it is proven in [2] that for $k \gg n/d$ it holds that $\alpha_{\leq k}(G) = o(1)$. It is an open question whether the same holds for irregular graphs (with average degree $d \gg \log \log n$).
- Finally, in many situations we are interested in multitasking a “small” number of tasks. This raises the question of computing (or approximating) α_k in the setting of fixed-parameter algorithms. That is, given an n -vertex graph G , can we compute α_k in time $f(k) \cdot \text{poly}(n)$, where $\text{poly}(n)$ is independent of k and $f(k)$ is some function of k independent of n .

Appendix A. \mathcal{NP} -hardness of computing the maximum connected matching of a graph. In this section given a bipartite graph it is NP-hard to compute $\nu_C(G)$ exactly under a deterministic polynomial time reduction. This is as opposed to the randomized reduction given in Theorem 3.1. We remark that [22] proved this result for the nonbipartite case. Our proof is an adaptation of their proof for the bipartite case.

THEOREM A.1. *It is \mathcal{NP} -hard to determine given a bipartite graph $G = (A, B)$ and a parameter k whether G contains a connected matching of size k .*

Proof. We reduce the biclique to the problem of determining if $\nu_C(G) = k$. Recall that a biclique $G' = (C', D')$ in a bipartite graph G is a subgraph G' of G such that every vertex in C' is connected to every vertex in D' . A biclique (C', D') is balanced if $|C'| = |D'|$ the biclique problem is the following: given a bipartite graph $G = (A, B)$ (we assume that $|A| = |B|$) and integer k , is there a biclique (A', B') with $A' \subseteq A, B' \subseteq B$, and $|A'| = |B'| = k$? This problem is well known to be \mathcal{NP} -complete.

Given a bipartite graph $G = (A, B)$ with $|A| = |B| = n$, form a new graph H as follows. Initialize $H_1 = (A_1, B_1)$ to equal G and we call this the copy of G inside H_1 . Then add a new set A' of n vertices such that (A_1, A') forms a biclique, and add a new set B' of n vertices such that (B_1, B') forms a biclique. Initialize another graph $H_2 = (A_2, B_2)$ to be a biclique with $|A_2| = |B_2| = n$ (where A_2, B_2 are disjoint from $A_1 \cup B_1 \cup A' \cup B'$). Add an edge between every vertex of $(A_1 \cup B')$ and every vertex of B_2 , and add an edge between every vertex of $(B_1 \cup A')$ and every vertex of A_2 . The resulting (bipartite) graph is $H = (A_1 \cup B' \cup A_2, B_1 \cup A' \cup B_2)$.

Consider a connected matching M in H . Let $M_A \subseteq M$ be the set of all edges in M contained in the biclique (A_1, A') and let $M_B \subseteq M$ be the set all edges in M contained in the biclique (B_1, B') , and let $M_r = M - (M_A \cup M_B)$. Then $|M| = |M_A| + |M_B| + |M_r|$. Let $X_A \subseteq A_1$ denote the set of vertices in A_1 being an endpoint of an edge in M_A , and let X_B be analogously defined with respect to B_1 and M_B . Since M is a connected matching, (X_A, X_B) is a biclique. We also have $|M_r| \leq 2n - \max\{|X_A|, |X_B|\}$, which implies $|M| \leq 2n + \min\{|X_A|, |X_B|\}$, where we have used $|X_A| = |M_A|, |X_B| = |M_B|$. Thus, if G has a connected matching of size $2n + k$, then $\min\{|X_A|, |X_B|\} \geq k$, which means that there is a biclique of size k .

Conversely, if G contains a biclique (R, S) of size k , we can easily form a connected matching M in H of size $2n + k$. To construct M , we take k edges M_A in (A, A') with $X_A = R$, we take k edges M_B in (B, B') with $X_B = S$, we take $n - k$ edges matching the $n - k$ vertices of $A_1 - X_A$ with $n - k$ vertices $B'_2 \subseteq B_2$, we take $n - k$ edges matching the $n - k$ vertices of $B_1 - X_B$ with $n - k$ vertices $A'_2 \subseteq A_2$, and we take k edges matching $A_2 - A'_2$ with $B_2 - B'_2$.

Thus, G contains a biclique of size k if and only if H contains a connected matching of size $2n + k$. This completes the proof. \square

REFERENCES

- [1] N. ALON, A. MOITRA, AND B. SUDAKOV, *Nearly complete graphs decomposable into large induced matchings and their applications*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing, ACM, 2012, pp. 1079–1090.
- [2] N. ALON, D. REICHMAN, I. SHINKAR, T. WAGNER, S. MUSSLICK, J. D. COHEN, T. GRIFFITHS, B. DEY, AND K. ÖZCIMDER, *A graph-theoretic approach to multitasking*, Adv. Neural Inf. Process. Syst., 2017, pp. 2100–2109.
- [3] P. AUSTRIN, S. KHOT, AND M. SAFRA, *Inapproximability of vertex cover and independent set in bounded degree graphs*, in Proceedings of the 24th Annual IEEE Conference on Computational Complexity, 2009, IEEE, 2009, pp. 74–80.
- [4] Y. BIRK, N. LINIAL, AND R. MESHULAM, *On the uniform-traffic capacity of single-hop interconnections employing shared directional multichannels*, IEEE Trans. Inform. Theory, 39 (1993), pp. 186–191.
- [5] R. BOPPANA AND M. M. HALLDÓRSSON, *Approximating maximum independent sets by excluding subgraphs*, BIT, 32 (1992), pp. 180–196.
- [6] K. CAMERON, *Connected matchings*, in Combinatorial Optimization—Eureka, You Shrink!, Springer, New York, 2003, pp. 34–38.
- [7] S. O. CHAN, *Approximation resistance from pairwise-independent subgroups*, J. ACM, 63 (2016), 27.
- [8] I. CHLAMTAC AND S. KUTTEN, *On broadcasting in radio networks—Problem analysis and protocol design*, IEEE Trans. Commun., 33 (1985), pp. 1240–1246.
- [9] J. D. COHEN, K. DUNBAR, AND J. L. MCCLELLAND, *On the control of automatic processes: A parallel distributed processing account of the stroop effect*, Psychol. Rev., 97 (1990), pp. 332–361.
- [10] P. ERDOS AND G. SZEKERES, *A combinatorial problem in geometry*, Compos. Math., 2 (1935), pp. 463–470.
- [11] S. F. FENG, M. SCHWEMMER, S. J. GERSHMAN, AND J. D. COHEN, *Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors*, Cogn. Behav. Neurosci., 14 (2014), pp. 129–146.
- [12] Z. FÜREDI, A. GYÁRFÁS, AND G. SIMONYI, *Connected matchings and Hadwiger’s conjecture*, Combin. Probab. Comput., 14 (2005), pp. 435–438.
- [13] P. A. GOLOVACH, P. HEGGERNES, P. VAN’T HOF, AND C. PAUL, *Hadwiger number of graphs with small chordality*, in Proceedings of the International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, New York, 2014, pp. 201–213.
- [14] J. HÅSTAD, *Some optimal inapproximability results*, J. ACM, 48 (2001), pp. 798–859.
- [15] A. S. JOBSON, A. E. KÉZDY, AND S. C. WHITE, *Connected matchings in chordal bipartite graphs*, Discrete Optim., 14 (2014), pp. 34–45.
- [16] K.-I. KAWARABAYASHI, M. D. PLUMMER, AND B. TOFT, *Improvements of the theorem of Duchet and Meyniel on Hadwiger’s conjecture*, J. Combin. Theory Ser. B, 95 (2005), pp. 152–167.
- [17] A. KOSTOCHKA AND L. PYBER, *Small topological complete subgraphs of “dense” graphs*, Combinatorica, 8 (1988), pp. 83–86.
- [18] P. MANURANGSI, *Almost-polynomial ratio ETH-hardness of approximating densest k -subgraph*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2017, pp. 954–961.
- [19] S. MUSSLICK, B. DEY, K. ÖZCIMDER, M. PATWARY, T. L. WILLKE, AND J. D. COHEN, *Controlled vs. automatic processing: A graph-theoretic approach to the analysis of serial vs. parallel processing in neural network architectures*, in Proceedings of the 38th Annual Meeting of the Cognitive Science Society (CogSci), 2016, pp. 1547–1552.

- [20] S. MUSSLICK, B. DEY, K. OZCIMDER, M. PATWARY, T. L. WILLKE, AND J. D. COHEN, *Parallel processing versus efficiency of representation in neural networks*, in Proceedings of the 16th Neural Computation and Psychology Workshop, 2016.
- [21] C. H. PAPANITRIOU, *Computational Complexity*, John Wiley and Sons, New York, 2003.
- [22] M. D. PLUMMER, M. STIEBITZ, AND B. TOFT, *On a special case of Hadwiger's conjecture*, *Discuss. Math. Graph Theory*, 23 (2003), pp. 333–363.
- [23] D. E. RUMELHART, J. L. MCCLELLAND, AND P. R. GROUP, *Parallel Distributed Processing*, Vol. 1, MIT Press, Cambridge, MA, 1987.
- [24] R. M. SHIFFRIN AND W. SCHNEIDER, *Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory*, *Psychol. Rev.*, 84 (1977), pp. 127–190.
- [25] J. T. TOWNSEND AND M. J. WENGER, *A theory of interactive parallel processing: New capacity measures and predictions for a response time inequality series*, *Psychol. Rev.*, 111 (2004), pp. 1003–1035.
- [26] J. VERSTRÆTE, *On arithmetic progressions of cycle lengths in graphs*, *Combin. Probab. Comput.*, 9 (2000), pp. 369–373.
- [27] D. ZUCKERMAN, *Linear degree extractors and the inapproximability of max clique and chromatic number*, in Proceedings of the 38th Annual ACM Symposium on Theory of Computing, ACM, 2006, pp. 681–690.