Many H-copies in graphs with a forbidden tree

Shoham Letzter*

Abstract

For graphs H and F, let $\mathrm{ex}(n, H, F)$ be the maximum possible number of copies of H in an F-free graph on n vertices. The study of this function, which generalises the well-studied Turán numbers of graphs, was initiated recently by Alon and Shikhelman. We show that if F is a tree then $\mathrm{ex}(n, H, F) = \Theta(n^r)$ for an (explicit) integer $r = r(H, F)$, thus answering one of their questions.

1 Introduction

Given graphs H and F with no isolated vertices and an integer n, let $\mathrm{ex}(n, H, F)$ be the maximum possible number of copies of H in an F-free graph on n vertices. This function was introduced recently by Alon and Shikhelman [1]. In the special case where $H = K_2$, this is the maximum possible number of edges in an F-free graph on n vertices, known as the *Turán number* of F, which is one of the main topics in extremal graph theory (see e.g. [21] for a survey).

A few instances of $\mathrm{ex}(n, H, F)$, with $H \neq K_2$, where studied prior to [1]. The first of these is due to Erdős [5] who determined $\mathrm{ex}(n, K_r, K_s)$ for all r and s (see also [2]).

A different example that has received considerable attention recently is $\mathrm{ex}(n, C_r, C_s)$ for various values of r and s. In 2008 Bollobás and Győri [3] showed that $\mathrm{ex}(n, K_3, C_5) = \Theta(n^{3/2})$, and their upper bound has been improved several times [1, 6]. Győri and Li [17] obtained upper and lower bounds on $\mathrm{ex}(n, K_3, C_{2k+1})$, that were subsequently improved by Füredi and Özkahaya [7] and by Alon and Shikhelman [1]. Moreover, the number $\mathrm{ex}(n, C_5, K_3)$ was calculated precisely [14, 18]. Very recently, Gishboliner and Shapira [13] determined $\mathrm{ex}(n, C_r, C_s)$, up to a constant factor, for all $r > 3$, and, additionally, they studied $\mathrm{ex}(n, K_3, C_s)$ for even r. Some additional more precise estimates for $\mathrm{ex}(n, C_r, C_s)$ are known (see [9, 15]).

There are, unsurprisingly, many more instances of studies of the function $\mathrm{ex}(n, H, F)$ or variations of it (e.g. when F is replaced by a family of graphs, or when the objects of interest are hypergraphs or posets rather than graphs); see, for example, [4, 8, 10, 11, 12, 19, 20, 22].

*ETH Institute for Theoretical Studies, ETH, 8092 Zurich; e-mail: shoham.letzter@eth-its.ethz.ch. Research supported by Dr. Max Rössler, the Walter Haefner Foundation and the ETH Zurich Foundation.

In this paper we shall be interested in the value of $\mathrm{ex}(n, H, T)$ when T is a tree. Alon and Shikhelman [1] showed that if H is also a tree then the following holds.

$$\mathrm{ex}(n, H, T) = \Theta(n^r) \text{ for some (explicit) integer } r = r(H, T). \tag{1}$$

See also [16] for the study of the special case where T and H are paths. Alon and Shikhelman asked if (1) still holds if only T is required to be a tree (and H is an arbitrary graph). Our main result answers this question affirmatively.

Theorem 1. *Let H be a graph and let T be a tree. Then there exists an integer $r = r(H, T)$ such that* $\mathrm{ex}(n, H, T) = \Theta(n^r)$.

We note that, as in Alon and Shikhelman's result for the case where H is also a tree, the integer $r = r(H, T)$ can be determined explicitly in terms of H and T; see Definition 3.

We present the proof in the Section 2, and conclude the paper in Section 3 with some closing remarks.

2 The proof

Our aim is to prove that $\mathrm{ex}(n, H, T) = \Theta(n^r)$ for a certain integer r. In order to describe this integer, we need the following two definitions.

Definition 2. Given a graph H, a subset $U \subseteq V(H)$ and an integer t, the (U, t)-*blow-up of* H is the graph obtained by taking t copies of H and identifying all the vertices that correspond to u, for each $u \in U$ (see Figure 1 for an example).

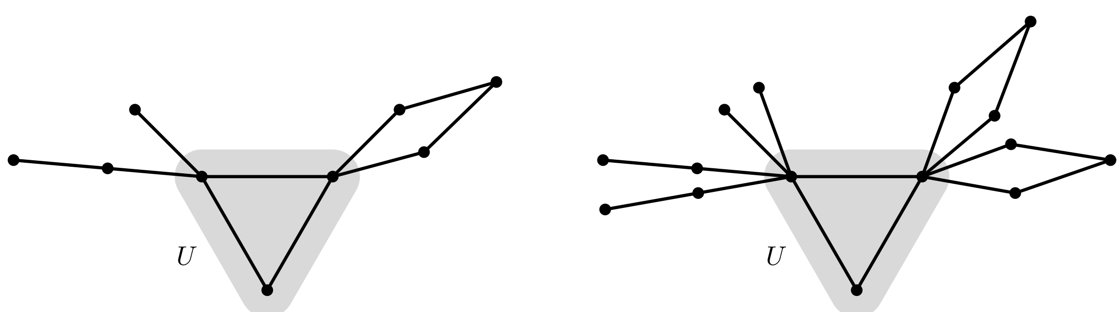

Figure 1: A graph H and a subset $U \subseteq V(H)$ and the $(U, 2)$-blow-up of H.

Definition 3. Given graphs H and T, let $r(H, T)$ be the maximum number of components in $H \setminus U$, over subsets $U \subseteq V(H)$ for which the $(U, |T|)$-blow-up of H is T-free.

In the following theorem we estimate $\mathrm{ex}(n, H, T)$, where T is a tree, in terms of the value $r(H, T)$. Note that Theorem 1 follows immediately.

Theorem 4. *Let H be a graph and let T be a tree. Then* $\mathrm{ex}(n, H, T) = \Theta(n^r)$, *where* $r = r(H, T)$.

The lower bound follows quite easily from the definition of $r(H, T)$, so the main work goes into proving the matching upper bound. In [1] Alon and Shikhelman proved the same statement under the additional assumption that H is a tree. In order to prove the upper bound, they showed that a graph G which is T-free and has at least $c \cdot n^r$ copies of H (for any integer r and a large constant c) contains a $(U, |T|)$-blow-up of H, for some $U \subseteq V(H)$ such that $H \setminus U$ has at least $r+1$ components. Since G is T-free, it follows that the $(U, |T|)$-blow-up is also T-free, which implies, by definition of $r(H, T)$, that G has fewer than $c \cdot n^{r(H,T)}$ copies of H, as required. Our ideas are somewhat similar, but we do not prove that G contains such a blow-up. Instead, we find a subgraph G' of G with many H-copies that behaves somewhat similarly to a $(U, |T|)$-blow-up of H, for some U for which the number of components of $H \setminus U$ is larger than r. We then show that if the blow-up contains a copy of T then so does G'. It again follows that the number of H-copies in G is smaller than $c \cdot n^{r(H,T)}$.

Proof of Theorem 4. Let $r = r(H, T)$, $h = |H|$, $t = |T|$ and $m = \mathrm{ex}(n, H, T)$. Our aim is to show that $m = \Theta(n^r)$.

We first show that $m = \Omega(n^r)$. Indeed, let $U \subseteq V(H)$ be such that $H \setminus U$ has r components and the (U, t)-blow-up of H is T-free. Let G be the $(U, n/h)$-blow-up of H. Note that G is T-free; indeed, otherwise, since any T-copy in G uses vertices from at most t copies of H, it would follow that the (U, t)-blow-up of H is not T-free. Additionally, the number of H-copies in G is at least $(n/h)^r$ since, for every component in $H \setminus U$, we can choose any of the n/h copies of it in G, and together with U this forms a copy of H.

The remainder of the proof will be devoted to proving the upper bound $m = O(n^r)$. Suppose to the contrary that $m \geq c \cdot n^r$, for a sufficiently large constant c. Let G be a T-free graph on n vertices with m copies of H.

Instead of studying G directly, we will consider a subgraph G' of G that has many H-copies and is somewhat similar to a (U, t)-blow-up of H for an appropriate U. We obtain the required subgraph in three steps.

First, we find an r-partite subgraph G_0 of G that has many H-copies. To achieve this goal, pick a label in $V(H)$ uniformly at random for each vertex in G. Denote by X the number of H-copies in G for which each vertex $u \in V(H)$ is mapped to a vertex in G that received the label u. It is easy to see that the $\mathbb{E}(X) = m/h^h$. It follows that there exists a partition $\{V_u\}_{u \in V(H)}$ of the vertices of G for which $X \geq m/h^h$. Fix such a partition, and denote by $\mathcal{H}_0$ the family of H-copies for which every $u \in V(H)$ is mapped to V_u (so $|\mathcal{H}_0| \geq m/h^h$). Let G_0 be the subgraph of G whose edge set is the collection of edges that appear in some H-copy in $\mathcal{H}_0$.

Next, since G_0 is T-free (as it is a subgraph of G), it is t-degenerate; fix an ordering $<$ of $V(G_0)$ such that every vertex u has at most t neighbours that appear after u in $<$. Each H-copy in $\mathcal{H}_0$ inherits an ordering of $V(H)$ from $<$. Denote by $<_H$ the most popular such ordering and let $\mathcal{H}_1$ be the subfamily of H-copies in $\mathcal{H}_0$ that received the ordering $<_H$ (so $|\mathcal{H}_1| \geq |\mathcal{H}_0|/h! \geq m/(h^h h!)$).

We now turn to the final step towards obtaining the required subgraph of G. Ideally, we would have liked to find a graph F, which is the union of $\Omega(m)$ distinct copies of H in $\mathcal{H}_1$, and satisfies the following: for every $uw \in E(H)$, either all vertices in V_u have small degree into V_w, or all vertices in V_u have much larger degree into V_w. Such a property would allow us to show that if a suitable (U,t)-blow-up of H contains a copy of T, then so does F. However, it is not clear if such a family of H-copies exists. Instead, we aim for a sequence of graphs $F_1 \supseteq \ldots \supseteq F_t$ (each of which is a union of a large collection of H-copies in $\mathcal{H}_1$) such that for every $uw \in E(H)$, either all vertices in V_u have small degree into V_w in the graph F_1, or all non-isolated vertices in F_i have much larger degree into V_w in the graph F_{i-1} for every $2 \le i \le t$. Such a sequence still allows us to find a copy of T in F_1, under the assumption that a certain (U,t)-blow-up of H contains a copy of T, using the fact that T is a tree. In order to find the required sequence of graphs, pick constants $t \ll c_0 \ll \ldots \ll c_{e(H)} \ll c$, and follow Procedure 1 below (see Figure 2 for an illustration of this procedure).

Procedure 1. Modifying $\mathcal{H}_1$

Set $\mathcal{H}_0^{(1)} = \mathcal{H}_1$.
Set E_0 to be the set of ordered pairs $\{uw : uw \in E(H), u >_H w\}$ (so $|E_0| = e(H)$).
Set $b = 0$ (b counts pairs (V_u, V_w) with bounded maximum degree in an appropriate graph).
Set $i = 1$ (i denotes the position in the sequence of t graphs we wish to generate).
while $b < e(H), i < t$ **do**
  For every $e = uw \in E_b$, let B_e be the set of vertices in V_u whose degree into V_w, with respect to $\mathcal{H}_b^{(i)}$, is at most c_b.
  if at least half the H-copies in $\mathcal{H}_b^{(i)}$ avoid $\bigcup_{e \in E_b} B_e$ **then**
    Set $\mathcal{H}_b^{(i+1)}$ to be the family of H-copies in $\mathcal{H}_b^{(i+1)}$ that avoid $\bigcup_{e \in E_b} B_e$.
    $i \leftarrow i + 1$.
  else
    Let $e \in E_b$ be such that at least $\frac{1}{2|E_b|}$ of the H-copies in $\mathcal{H}_b^{(i+1)}$ are incident with B_e.
    Set $\mathcal{H}_{b+1}^{(1)}$ to be the family of H-copies in $\mathcal{H}_b^{(i)}$ that are incident with B_e.
    Set $E_{b+1} = E_b \setminus \{e\}$.
    $b \leftarrow b + 1$, $i \leftarrow 1$.
  end if
end while

Note that the procedure ends either with $b = e(H)$ and $E_b = \emptyset$, or with $b \le e(H)$, $i = t$ and $|E_b| = e(H) - (b-1)$. Let $\bar{b}$ be the value of b at the end of the procedure. In the next claim we show that the latter case holds, i.e. $\bar{b} < e(H)$ (in other words, there is a pair (V_u, V_w) whose maximum degree in $\mathcal{H}_b^{(t)}$ is unbounded).

Claim 5. $\bar{b} < e(H)$.

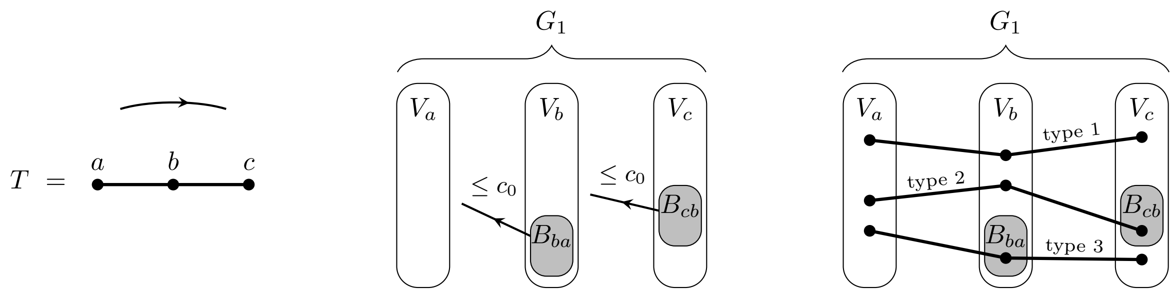

Figure 2: A simple example to illustrate Algorithm 1.

T is a path on three vertices, with vertex order $a <_T b <_T c$; the vertices of G_1 are partitioned into sets V_a, V_b, V_c and we are interested in T-copies where x is mapped to V_x for $x \in \{a, b, c\}$.

By definition of G_1, vertices in V_a have degree at most t into V_b, and vertices in V_b have degree at most t into V_c. The set B_{ba} consists of vertices in V_b with small degree (at most c_0) into V_a; B_{cb} is defined similarly.

In the first iteration of the procedure (when $b = 0$), we distinguish three types of T-copies: copies that avoid $B_{cb} \cup B_{ba}$ (type 1); copies that are incident with B_{cb} (type 2); and copies that are incident with B_{ba} (type 3). We keep T-copies of one of the types, depending on which one is most common. We either repeat this step (if we chose to keep the type 1 vertices) or we proceed to the next iteration of the procedure (with $b = 1$).

Proof. Let $\mathcal{F} := \mathcal{H}^{(1)}_{\overline{b}}$, and let F be the corresponding graph. Note that, as c is large,

$$|\mathcal{F}| \geq \left(\frac{1}{2e(H)}\right)^{t \cdot e(H)} |\mathcal{H}_1| \geq \left(\frac{1}{2e(H)}\right)^{t \cdot e(H)} \frac{1}{h^h h!} \cdot m > \frac{1}{\sqrt{c}} \cdot m.$$

Suppose that $\overline{b} = e(H)$. Then, for every $uw \in E(H)$, every vertex in V_u sends at most $c_{\overline{b}}$ edges into V_w (with respect to F).

Let a be the number of connected components in H. Note that the $(\emptyset, t)$-blow-up of H is T-free (it is a disjoint union of copies of H, and we may assume that H is T-free, as otherwise $m = 0$ and we are done immediately), and has a components. Thus, by Definition 3, we have $a \leq r$.

In order to upper-bound the number of H-copies in $\mathcal{F}$, let U be a set of vertices in H that contains exactly one vertex from each component. Trivially, there are at most n^a ways to map each vertex $u \in U$ to a vertex in V_u. Fix such a mapping. Let w be a vertex in H with a neighbour $u \in U$, and suppose that u is mapped to $x \in V(F)$. Since w is mapped to one of the neighbours in V_w of x, there are at most $c_{\overline{b}}$ vertices that w can be mapped to. Similarly, if w is in distance d from a vertex $u \in U$, there are at most $(c_{\overline{b}})^d$ vertices that w can be mapped to. By choice of U, every vertex in H is in distance at most h from some vertex in U, hence there are at most $(c_{\overline{b}})^{h^2}$ ways to complete the embedding of U to an H-copy in $\mathcal{F}$. In total, we find that $|\mathcal{F}| \leq (c_{\overline{b}})^{h^2} \cdot n^a < \sqrt{c} \cdot n^a$.

Putting the two bounds on $|\mathcal{F}|$ together, we have $m < c \cdot n^a \leq c \cdot n^r$, a contradiction to the assumption on m. It follows that $\bar{b} < e(H)$, as desired. $\square$

From now on, we may assume that $\bar{b} < e(H)$, which means that $\mathcal{H}_{\bar{b}}^{(i)}$ has been defined for every $i \in [t]$. Write $\mathcal{F}_i = \mathcal{H}_{\bar{b}}^{(i)}$, and denote by F_i the graph formed by taking the union of all H-copies in $\mathcal{F}_i$. Let D be the directed graph on vertex set $V(H)$ with edges $\{uw, wu : uw \in E(H)\}$ (so each edge in H is replaced by two directed edges, one in each direction). We 2-colour the edges of D: colour the edges in $E_{\bar{b}}$ red and colour the remaining edges blue. (Note that if uw is red then wu is blue.) Denote the graph of blue edges by D_B and the graph of red edges by D_R. By definition of $\mathcal{F}_i$ using Algorithm 1, one can check that

(a) $G \supseteq G_1 \supseteq F_1 \supseteq \ldots \supseteq F_t$.

(b) If $uw \in D_B$, all vertices in V_u have degree at most $c_{\bar{b}-1}$ into V_w in F_1.

(c) If $uw \in D_R$, all non-isolated vertices in V_u with respect to F_i have degree at least $c_{\bar{b}}$ into V_u in, for every $2 \leq i \leq t$.

Indeed, (a) and (c) follow easily from the definition of the procedure. To see (b), if uw is blue, then either $u <_H w$ which implies that vertices in V_u have at most t edges into V_w in G_1, or the edge uw was originally in E_0, but was removed at some point before the final iteration of the procedure, which implies that every vertex in V_u sends at most c_b edges into V_w in F_b, for some $b < \bar{b}$.

We shall use the following properties of F_i and $\mathcal{F}_i$.

Claim 6. *The following two properties hold for $2 \leq i \leq t$.*

(i) *every non-isolated vertex in F_i is contained in an H-copy in $\mathcal{F}_{i-1}$,*

(ii) *let uw be a red edge in D and let $S = \bigcup_{v:\text{ there is a blue path from } v \text{ to } w} V_v$. Then for every non-isolated vertex $x \in V_u$ there is a collection of t copies of H in $\mathcal{F}_{i-1}$ that contain x and whose intersections with S are pairwise vertex-disjoint.*

Proof. The first property follows immediately from the definition of F_i as the union of H-copies in $\mathcal{F}_i$: if a vertex is non-isolated in F_i it is also non-isolated in F_{i-1}, and thus it must be contained in some H-copy in $\mathcal{F}_{i-1}$.

Now let us see why the second property holds. Note that the directed edge uw is in $E_{\bar{b}}$ as uw is a red edge in D. Thus, by definition of $\mathcal{F}_i$, any non-isolated vertex $x \in V_u$ sends at least $c_{\bar{b}}$ edges into V_w in the graph F_{i-1}. This means that there is a collection of at least $c_{\bar{b}}$ copies of H in $\mathcal{F}_{i-1}$ that contain x, each of which uses a different edge from x to V_u; denote this family of H-copies by $\mathcal{F}$. We claim that every H-copy in $\mathcal{F}$ intersects at most $h \cdot (c_{\bar{b}-1})^h$ other H-copies in $\mathcal{F}$ in S. Indeed, there are at most h ways to choose an intersection point; suppose that the intersection is in $y \in V_v \subseteq S$. By choice of S, there is a path $(v_0 = v, v_1, \ldots, v_k = w)$ from v to w in D_B. This means that the degree into $V_{v_{j+1}}$ (with respect to F_{i-1}) of any vertex in V_{v_j} is at most $c_{\overline{b}-1}$. Thus, there are at most $(c_{\overline{b}-1})^k \leq (c_{\overline{b}-1})^h$ vertices in V_w that can be in the same H-copy in $\mathcal{F}$ as y. Since each H-copy in $\mathcal{F}$ uses a different vertex of V_w, it follows that at most $(c_{\overline{b}-1})^h$ copies of H in $\mathcal{F}$ contain y, and in total there are at most $h(c_{\overline{b}-1})^h$ copies of H in $\mathcal{F}$ that intersect any single H-copy in $\mathcal{F}$. Since the total number of H-copies in $\mathcal{F}$ is $c_{\overline{b}} \geq t \cdot (h \cdot (c_{\overline{b}-1})^h + 1)$, there is a collection of t copies of H in $\mathcal{F}$ whose intersections with S are pairwise disjoint, as required. $\square$

We now wish to find a particular subset $U \subseteq V(H)$ such that the (U, t)-blow-up of H behaves similarly to the sequence of graphs $F_1, \ldots, F_t$. The set U will be defined in terms of a certain set $A \subseteq V(H)$, which we define now. Let $\mathcal{P}$ be a partition of $V(H)$ into strongly connected components according to D_B. Pick a set $A \subseteq V(H)$ that satisfies the following properties.

(a) every vertex in D_B is reachable from A, i.e. for every $u \in D_B$ there is a blue path from A to u,

(b) $|A|$ is minimal among sets that satisfy (a),

(c) among sets that satisfy (a) and (b), A maximises

$$\sum_{u \in A} (\#\text{ vertices reachable from } u). \tag{2}$$

Let W be the set of vertices in $V(H)$ that are in the same part of $\mathcal{P}$ as one of the vertices in A, and let $U = V(H) \setminus W$. In the following two claims we list some useful properties of A, U and W.

Claim 7. *The following properties hold.*

(i) *A contains at most one vertex from each part of $\mathcal{P}$,*

(ii) *there are no edges of D between distinct parts of $\mathcal{P}$ that are contained in W,*

(iii) *there are no blue edges from U to W.*

Proof. Property (i) clearly holds because of the minimality of $|A|$ and the fact that for every part $X \in \mathcal{P}$, the set of vertices reachable from X is the same as the set of vertices reachable from any individual vertex $x \in X$.

For (ii), suppose that there is an edge uw in D with u and w belonging to distinct parts of $\mathcal{P}$ that are contained in W; without loss of generality uw is blue. If we remove from A the vertex from the same part of $\mathcal{P}$ as w, we obtain a smaller set that still satisfies (a) above, a contradiction to the minimality of A.

Now suppose that Property (iii) does not holds, i.e. there is a blue edge uw with $u \in U$ and $w \in W$. Let A' be the set obtained from A by removing the vertex w' that is in the same part of $\mathcal{P}$ as w and adding u. Note that every vertex that is reachable from A is also reachable from A'. Moreover, every vertex that is reachable from w' is also reachable from u, but u is not reachable from w', because otherwise u and w' would have been in the same strongly connected component, and hence in the same part of $\mathcal{P}$. It follows that

$$\sum_{u\in A'}(\#\text{ vertices reachable from }u) > \sum_{u\in A}(\#\text{ vertices reachable from }u),$$

a contradiction to the maximality property of A. $\square$

Claim 8. $|A| > r$.

Proof. Suppose that $|A| \le r$. As in the proof of Claim 5, there are at most $n^{|A|}$ ways to embed A in $V(F_1)$ in such a way that every $a \in A$ is sent to V_a. Fix such an embedding, and let $u \in V(H)$. Because there is a blue path from A to u (by (a) in the definition of A), there are at most $(c_{\bar{b}-1})^h$ vertices that u could be mapped to which may form an H-copy in $\mathcal{F}_1$ together with the vertices that A is mapped to. Thus, in total there are at most $(c_{\bar{b}-1})^{h^2} \cdot n^r$ copies of H in $\mathcal{F}_1$. As in the proof of Claim 5, this implies that there are fewer than $c \cdot n^r$ copies of H in G, a contradiction. $\square$

Let Γ be the (U,t)-blow-up of H (see Definition 2 and Figure 1). Denote its vertices by $U \cup \left(\bigcup_{i\in[t]} W_i\right)$, where the W_i's are copies of the set W (so $\Gamma[U \cup W_i]$ induced a copy of H for every $i \in [t]$). For every vertex x in Γ, denote by $\phi(x)$ the vertex in H that it corresponds to. By Claim 7 (i) and (ii), $H \setminus U$ consists of $|A| > r$ components. Because $r = r(H,T)$ (see Definition 3), Γ contains a copy of T.

Consider a specific embedding of T in Γ. Let $\{X_1,\ldots,X_k\}$ be a partition of $V(T)$, such that for every $i \in [k]$ the subgraph $T[X_i]$ is a maximal non-empty subtree of T that is contained either in W_j, for some j, or in U. We assume, for convenience, that the ordering is such that there is an edge between X_i and $X_1 \cup \ldots \cup X_{i-1}$ for every $i \in [k]$; in fact, there would be exactly one such edge as T is a tree. By choice of the X_i's and by definition of Γ, this edge must be an edge between some set W_j and U.

Our final aim is to show that G contains a copy of T, a contradiction to the assumptions on G. We reach the required contradiction by proving the following claim.

Claim 9. *For every $i \in [k]$ there is a copy of $T[X_1 \cup \ldots \cup X_i]$ in $F_{t-(i-1)}$ such that x is mapped to $V_{\phi(x)}$ for every $x \in X_1 \cup \ldots \cup X_i$.*

Proof. We prove the statement by induction on i. For $i = 1$, the statement can easily be seen to hold, by picking any H-copy in $\mathcal{F}_t$, and mapping each vertex of X_1 to the corresponding vertex in the copy of H.

Now suppose that the statement holds for i; let $f_i : X_1 \cup \ldots \cup X_i \to V(F_{t-(i-1)})$ be the corresponding mapping of the vertices. Now, there are two possibilities to consider: $X_{i+1} \subseteq U$ or $X_{i+1} \subseteq W_j$ for some j.

Let us consider the first possibility. Let uw be the edge between $X_1 \cup \ldots \cup X_i$ and X_{i+1}, where $u \in U$ and $w \in W_j$ for some j (so $u \in X_{i+1}$ and $w \in X_1 \cup \ldots \cup X_i$). We may assume that $f_i(w)$ is non-isolated in $F_{t-(i-1)}$. Indeed, if $|X_1 \cup \ldots \cup X_i| \geq 2$, this is clear since $T[X_1 \cup \ldots \cup X_i]$ spans a tree. Otherwise, we must have that $i = 1$ and $|X_1| = 1$, but then we can choose $f_1(w)$ to be a non-isolated vertex in V_w with respect to F_t. As $f_i(w)$ is non-isolated, by Claim 6 (and the fact that $i \leq k \leq t$) there is an H-copy in $\mathcal{F}_{t-i}$ that contains $f_i(w)$; denote the corresponding embedding by $g : V(H) \to V(F_{t-i})$. We define $f_{i+1} : X_1 \cup \ldots \cup X_{i+1} \to V(F_{t-i})$ simply by

$$f_{i+1}(x) = \begin{cases} f_i(x) & x \in X_1 \cup \ldots \cup X_i \\ g(x) & x \in X_{i+1}. \end{cases}$$

In order to show that f_{i+1} is an embedding with the required properties, we need to show that it has the following three properties: it maps edges in $T[X_1 \cup \ldots \cup X_{i+1}]$ to edges in F_{t-i}; $f_{i+1}(x) \in V_{\phi(x)}$ for every $x \in X_1 \cup \ldots \cup X_{i+1}$; and f_{i+1} is injective.

We first show that f_{i+1} preserves edges. This follows because f_i and g preserve edges (this holds for g by definition, and holds for f_i because it sends edges of $T[X_1 \cup \ldots \cup X_i]$ to edges of $F_{t-(i-1)}$ which is a subgraph of F_{t-i}) so edges inside $X_1 \cup \ldots \cup X_i$ and inside X_{i+1} are mapped to edges in F_{t-i}, and moreover by choice of g the only edge between these two sets is mapped to an edge of F_{t-i}.

Next, we note that for every $x \in X_1 \cup \ldots \cup X_{i+1}$, we have $f_{i+1}(x) \in V_{\phi(x)}$. This is because this holds for both f_i (by assumption) and g (as g corresponds to an H-copy in $\mathcal{F}_{t-i}$).

Finally, we show that f_{i+1} is injective. As both f_i and g are injective, it suffices to show that $g(x) \neq f_i(y)$ for every $x \in X_{i+1}$ and $y \in X_1 \cup \ldots \cup X_i$. This holds because $\phi(x) \neq \phi(y)$ (since x is in U, it is the only vertex in $X_1 \cup \ldots \cup X_{i+1}$ with $\phi(x) = x$) and because x and y are mapped to $V_{\phi(x)}$ and $V_{\phi(y)}$, respectively, and these two sets are disjoint.

Now we consider the second possibility, namely that $X_{i+1} \subseteq W_j$ for some j. Let uw be the edge between $X_1 \cup \ldots \cup X_i$ and X_{i+1}, where $u \in U$ and $w \in W_j$ (so $w \in X_{i+1}$). By Claim 7 (iii), the edge uw is red. Hence, by Claim 6, there is a collection of t copies of H in $\mathcal{F}_{t-i}$ that contain $f_i(u)$ and whose intersections with $S = \bigcup_{v\text{: there is a blue path from } v \text{ to } w} V_v$ are pairwise vertex-disjoint. As $|X_1 \cup \ldots \cup X_i| < t$, it follows that there is an H-copy in $\mathcal{F}_{t-i}$ that contains $f_i(w)$ and whose intersection with S is disjoint of $f_i(X_1 \cup \ldots \cup X_i)$; denote the corresponding embedding of H by $g : V(H) \to V(F_{t-i})$. As before, define $f_{i+1} : X_1 \cup \ldots \cup X_{i+1} \to V(F_{t-i})$ by

$$f_{i+1}(x) = \begin{cases} f_i(x) & x \in X_1 \cup \ldots \cup X_i \\ g(x) & x \in X_{i+1}. \end{cases}$$

As before, f_{i+1} maps edges of $T[X_1 \cup \ldots X_{i+1}]$ to edges of F_{t-i}, and it sends every $x \in X_1 \cup \ldots \cup X_{i+1}$ to $V_{\phi(x)}$. Moreover, by choice of g and since $g(X_{i+1}) \subseteq S$, we find that $g(X_{i+1})$ and $f_i(X_1 \cup \ldots \cup X_i)$ are disjoint. Since f_i and g are both injective, it follows that f_{i+1} is injective. This completes the proof of the induction step, and thus of the claim. $\square$

By taking $i = k$ in the previous claim, we find that $F_{t-(k-1)}$ contains a copy of T. But $F_{t-(k-1)} \subseteq F_1 \subseteq G$ (note that $k \leq t$), so G has a copy of T, a contradiction. It follows that the number of H-copies in G is at most $c \cdot n^{r(H,T)}$, as required. $\square$

3 Conclusion

In this paper we determined, up to a constant factor, the function $\mathrm{ex}(n, H, T)$ for any tree T. We note that the assumption that T is a tree was crucial in our proof to work, but it was used only in the proof of Claim 9 (where we made use of the fact that there is exactly one edge between $X_1 \cup \ldots \cup X_i$ and X_{i+1}).

It would, of course, be interesting to sharpen our result by determining $\mathrm{ex}(n, H, T)$ completely, or at least asymptotically. While this may be hopeless in general, in some special cases this task may not be out of reach. For example, Alon and Shikhelman [1] consider the special case where $H = K_h$ for some $h < t$ and $t = |T|$. They ask if the n-vertex graph, which is the union of $\lfloor n/t \rfloor$ disjoint cliques of size t, and perhaps one smaller clique on the remainder maximises the number of copies of K_h among all T-free graphs on n vertices. This question generalises a question of Gan, Loh and Sudakov [8], who considered the case where T is a star on t vertices. In other words, they were interested in maximising the number of cliques of size h among n-vertex graphs with maximum degree smaller than t. They proved that the aforementioned construction of disjoint cliques is the unique extremal example when $n \leq 2t$, thus proving a conjecture of Engbers and Galvin [4]. The question whether this construction is best for larger values of n remains open.

For other questions regarding the value of $\mathrm{ex}(n, H, F)$, where F need not be a tree, see [1].

Acknowledgements

I would like to thank the anonymous referee for their useful suggestions.

References

[1] N. Alon and C. Shikhelman, *Many T copies in H-free graphs*, J. Combin. Theory Ser. B **121** (2016), 146–172.

[2] B. Bollobás, *On complete subgraphs of different orders*, Math. Proc. Cambridge Philos. Soc. **79** (1976), 19–24.

[3] B. Bollobás and E. Győri, *Pentagons vs. triangles*, Discrete Math. **308** (2008), 4332–4336.

[4] J. Engbers and D. Galvin, *Counting independent sets of a fixed size in graphs with a given minimum degree*, J. Graph Theory **76** (2014), 149–168.

[5] P. Erdős, *On the number of complete subgraphs contained in certain graphs*, Magyar Tud. Akad. Mat. Kutató Int. Közl. **7** (1962), 459–464.

[6] B. Ergemlidze, E. Györi, A. Methuku, and N. Salia, *A note on the maximum number of triangles in a C_5-free graph*, Electron. Notes Discrete Math. **61** (2017), 395–398.

[7] Z. Füredi and L. Özkahaya, *On the maximum number of triangles in special families of graphs and hypergraphs*, in preparation.

[8] W. Gan, P. Loh, and B. Sudakov, *Maximizing the number of independent sets of a fixed size*, Combin Probab Comput **24** (2015), 521–527.

[9] D. Gerbner, E. Győri, A. Methuku, and M. Vizer, *Generalized Turán problems for even cycles*, Acta Math. Univ. Comenianae **87** (2019), 723–728.

[10] D. Gerbner, B. Keszegh, and B. Patkos, *Generalized forbidden subposet problems*, preprint, arXiv:1701.05030.

[11] D. Gerbner, A. Methuku, and M. Vizer, *Generalized Turán problems for disjoint copies of graphs*, Discr. Math. (2019).

[12] D. Gerbner and C. Palmer, *Counting copies of a fixed subgraph in F-free graphs*, European J. Combin. **82** (2019), 103001.

[13] L. Gishboliner and A. Shapira, *A Generalized Turán Problem and its Applications*, Int. Math. Res. Notices, to appear.

[14] A. Grzesik, *On the maximum number of five-cycles in a triangle-free graph*, J. Combin. Theory Ser. B **102** (2012), 1061–1066.

[15] A. Grzesik and B. Kielak, *On the maximum number of odd cycles in graphs without smaller odd cycles*, Acta Math. Univ. Comen. **88** (2019), 755–758.

[16] E. Győri, N. Salia, C. Tompkins, and O. Zamora, *The maximum number of P_ℓ copies in P_k-free graphs*, preprint, arXiv:1803.03240.

[17] E. Győri and H. Li, *The maximum number of triangles in C_{2k+1}-free graphs*, Combin. Probab. Comput. **21** (2012), 187–191.

[18] H. Hatami, J. Hladký, D. Král', S. Norine, and A. Razborov, *On the number of pentagons in triangle-free graphs*, J. Combin. Theory Ser. A **120** (2013), 722–732.

[19] R. Luo, *The maximum number of cliques in graphs without long cycles*, J. Combin. Theory Ser. B **128** (2018), 219–226.

[20] J. Ma and Y. Qiu, *Some sharp results on the generalized Turán numbers*, European J. Combin. **84** (2020), 103026.

[21] M. Simonovits, *Paul Erdős' influence on extremal graph theory*, The mathematics of Paul Erdős, II, Algorithms Combin., vol. 14, Springer, Berlin, 1997, pp. 148–192.

[22] J. Wang, *The shifting method and generalized Turán number of matchings*, preprint, arXiv:1812.01832.