

Conditioning and covariance on caterpillars

Sarah R. Allen*

Ryan O'Donnell†

July 17, 2014

Abstract

Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be joint $\{\pm 1\}$ -valued random variables. It is known that conditioning on a random subset of $O(1/\epsilon^2)$ of them reduces their average pairwise covariance to below ϵ (in expectation). We conjecture that $O(1/\epsilon^2)$ can be improved to $O(1/\epsilon)$. The motivation for the problem and our conjectured improvement comes from the theory of global correlation rounding for convex relaxation hierarchies. We suggest attempting the conjecture in the case that $\mathbf{X}_1, \dots, \mathbf{X}_n$ are the leaves of an information flow tree. We prove the conjecture in the case that the information flow tree is a caterpillar graph (similar to a two-state hidden Markov model).

1 Introduction

Let $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ be a list of jointly distributed Boolean random variables taking values in $\{\pm 1\}$. We are interested in the quantity

$$\text{avg}_{\substack{\text{distinct pairs} \\ u, v \in [n]}} \{ |\mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v]| \} \in [0, 1].$$

For brevity we call this the *average covariance* of the random variables (absolute-value sign notwithstanding). It is a quantification of the extent to which the random variables are (pairwise) independent.

If the average covariance of $\mathbf{X}_1, \dots, \mathbf{X}_n$ is not small, then in some sense a “typical” \mathbf{X}_j contains a considerable amount of information about a sizeable fraction of the other \mathbf{X}_k 's. Then if we condition on \mathbf{X}_j , we might expect the variance of these other \mathbf{X}_k 's to decrease, thereby decreasing the overall average covariance. For $t \in \mathbb{Z}^+$, we introduce the following notation:

$$\text{avgCov}_t(\mathbf{X}) := \text{avg}_{\substack{J \subseteq [n] \\ |J|=t}} \text{avg}_{\substack{\text{distinct pairs} \\ u, v \in [n] \setminus J}} \left\{ \mathbf{E} \left[|\mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v]| \mid (\mathbf{X}_j)_{j \in J} \right] \right\}.$$

The intuitions just described lead to the idea that choosing large t should cause $\text{avgCov}_t(\mathbf{X})$ to become small. Indeed, the following has recently been proven [4, 6, 8]:

Theorem 1.1. *Let $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ be $\{\pm 1\}$ -valued random variables and let $0 < \epsilon \leq 1$. Then for some integer $0 \leq t \leq O(1/\epsilon^2)$ it holds that $\text{avgCov}_t(\mathbf{X}) \leq \epsilon$.*

We present the following conjecture, made jointly with Yuan Zhou.

Conjecture A. Theorem 1.1 holds with $O(1/\epsilon)$ in place of $O(1/\epsilon^2)$.

*Department of Computer Science, Carnegie Mellon University. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 0946825. srallen@cs.cmu.edu

†Department of Computer Science, Carnegie Mellon University. Supported by NSF grants CCF-0747250 and CCF-1116594. Part of this work performed at the Boğaziçi University Computer Engineering Department, supported by Marie Curie International Incoming Fellowship project number 626373. odonne11@cs.cmu.edu

Remark 1.2. In Theorem 1.1, by $t \leq O(1/\epsilon^2)$ we mean $t \leq C/\epsilon^2$ where C is a universal constant independent of $\mathbf{X}_1, \dots, \mathbf{X}_n$. (We also assume $n \geq C/\epsilon^2 + 2$.) However one *cannot* simply fix $t = \lceil C/\epsilon^2 \rceil$ independently of $\mathbf{X}_1, \dots, \mathbf{X}_n$; this would make Theorem 1.1 false (see Proposition 1.7). These comments apply equally to Conjecture A with $O(1/\epsilon)$ in place of $O(1/\epsilon^2)$.

Motivation for Theorem 1.1 and Conjecture A comes from the theory of rounding algorithms for convex relaxations of optimization problems; specifically, the “correlation rounding” technique for the Sherali–Adams and SOS hierarchies. In Section 1.2 we further discuss this motivation, as well as the importance of improving the bound $t \leq O(1/\epsilon^2)$ to $t \leq O(1/\epsilon)$.

We were led to make Conjecture A based on algorithmic optimism as well as being unable to find any counterexample refuting it. The following example (which we call the “homogeneous star”) is particularly instructive. Let $\mathbf{X}_0 \sim \{\pm 1\}$ be uniformly random and suppose $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ is a list of independent “ ρ -correlated” copies of \mathbf{X}_0 (where $\rho \in [0, 1]$). I.e., for each $j \in [n]$ we have $\mathbf{X}_j = \mathbf{X}_0 \mathbf{R}_j$, where $\mathbf{R}_1, \dots, \mathbf{R}_n$ are independent $\{\pm 1\}$ -valued random variables satisfying $\mathbf{E}[\mathbf{R}_j] = \rho$. By symmetry, all sets J in the definition of $\text{avgCov}_t(\mathbf{X})$ contribute equally to the average, so suppose we condition on $\mathbf{X}_1, \dots, \mathbf{X}_t$. It is not hard to check that the conditional average covariance of $\mathbf{X}_{t+1}, \dots, \mathbf{X}_n$ is then

$$\rho^2 \text{Var}[\mathbf{X}_0 \mid \mathbf{X}_1, \dots, \mathbf{X}_t].$$

If $\rho \leq \sqrt{\epsilon}$ then this quantity is automatically at most ϵ , even without conditioning. On the other hand, if $\rho \gg \sqrt{\epsilon}$ then we need to rely on the conditional variance above being small. It’s not difficult to show via a Hoeffding bound that this conditional variance is very small if (and only if) $t\rho^2 \gg 1$; i.e., $\rho \gg 1/\sqrt{t}$. Thus by taking t a little bigger than $1/\epsilon$, the case of $\rho \gg \sqrt{\epsilon}$ is handled as well. In other words, these rough calculations confirm (perhaps up to a log factor) that Conjecture A holds for the homogeneous star for every value of ρ . On the other hand, this example also implies that one cannot hope for an improved bound of $t < o(1/\epsilon)$ in Conjecture A.

1.1 Information flow trees

Being unable to prove Conjecture A, we turn to trying to prove it in a wide family of special cases. Specifically, we study the conjecture in the special case of *information flow trees* (which includes the homogeneous star example discussed above). Information flow trees have been studied in an extremely wide variety of contexts, under various names: in the theory of noisy communication and computation; in statistical physics (as the *Ising model* on trees); in biology (as *phylogenetic trees*); and in learning theory (as Markov networks/graphical models). See Evans et al. [5] for a number of results, and Mossel [7] for a survey.

Definition 1.3. An *information flow tree* $\mathcal{T} = (V, E, \rho)$ is an undirected tree graph (V, E) (with $|V| > 1$) together with a function $\rho : E \rightarrow [-1, 1]$ giving a *correlation* parameter for each edge. We think of \mathcal{T} as generating a collection of $\{\pm 1\}$ -valued random variables $(\mathbf{X}_v)_{v \in V}$, $(\mathbf{R}_e)_{e \in E}$ as follows: First, the random variables $\mathbf{R}_e \in \{\pm 1\}$ are chosen such that $\mathbf{E}[\mathbf{R}_e] = \rho(e)$, independently for all $e \in E$. Next, the random variables $(\mathbf{X}_v)_{v \in V}$ are collectively chosen so that $\mathbf{X}_u \mathbf{X}_v = \mathbf{R}_{(u,v)}$ holds for all $(u, v) \in E$, uniformly at random from the two possibilities.

Remark 1.4. An equivalent way to think of the (\mathbf{X}_v) random variables being generated is as follows: First, a vertex $r \in V$ is chosen to be the “root”. (This choice can be arbitrary, as it does not affect the final distribution.) Next, \mathbf{X}_r is chosen uniformly at random from $\{\pm 1\}$. Finally, the remaining random variables $(\mathbf{X}_v)_{v \neq r}$ are determined by “noisily propagating” \mathbf{X}_r ’s value along edges of the tree: if \mathbf{X}_u has been chosen, and $(u, v) \in E$, then \mathbf{X}_v is set to \mathbf{X}_u with probability $\frac{1}{2} + \frac{1}{2}\rho(u, v)$ and is set to $-\mathbf{X}_u$ otherwise. We add the remark that in the end, each \mathbf{X}_v is individually uniformly distributed on $\{\pm 1\}$.

Remark 1.5. When discussing information flow trees, we think of the vertex random variables \mathbf{X}_v as the main objects of interest, and the edge random variables \mathbf{R}_e merely as ancillary information used to construct the \mathbf{X}_v ’s. Furthermore, if $V = L \sqcup M$ is the partition of V into *leaf* vertices L and *internal* vertices M , we usually think of the *leaf random variables* $(\mathbf{X}_v)_{v \in L}$ as being “observable” and the internal random variables $(\mathbf{X}_v)_{v \in M}$ as being “hidden”.




Figure 1: An example of a caterpillar tree

In this paper we study the special case of Conjecture A in which $\mathbf{X}_1, \dots, \mathbf{X}_n$ are the leaf random variables of an information flow tree. Referring to Remark 1.2, in this case we conjecture it *is* possible to fix $t = \text{const}/\epsilon$ independently of $\mathbf{X}_1, \dots, \mathbf{X}_n$. Assuming we can fix t allows us to make a few more simplifications. Since

$$\text{avgCov}_t(\mathbf{X}) = \text{avg}_{\substack{U \subseteq [n] \\ |U|=t+2}} \left\{ \text{avgCov}_t \left((\mathbf{X}_k)_{k \in U} \right) \right\},$$

it follows that proving the conjecture in the $n = t + 2$ case suffices to prove it for general $n \geq t + 2$. And when $n = t + 2$, the experiment reduces to the following: we choose a random pair of leaves u and v , condition on *all* other leaf random variables \mathbf{X}_w , and then measure the (conditional) covariance of $\mathbf{X}_u, \mathbf{X}_v$. Thus we are led to the following conjecture (in which we write t instead of $t + 2$ for notational simplicity):

Conjecture B. Let \mathcal{T} be an information flow tree with leaf random variables $\mathbf{X}_1, \dots, \mathbf{X}_t$ (where $t \geq 2$). Then

$$\text{avg}_{\substack{\text{distinct pairs} \\ u, v \in [t]}} \mathbf{E} \left[\left| \text{Cov}[\mathbf{X}_u, \mathbf{X}_v] \right| \mid (\mathbf{X}_j)_{j \in [t] \setminus \{u, v\}} \right] \leq O(1/t). \quad (1)$$

We emphasize that Conjecture B implies Conjecture A in the case that $\mathbf{X}_1, \dots, \mathbf{X}_n$ are the leaves of an information flow tree, and is in fact slightly stronger in that the bound is $O(1/t)$ for all t , independently of $\mathbf{X}_1, \dots, \mathbf{X}_n$.

In Sections 2–5 we will give some results in the direction of proving Conjecture B; however, we are still unable to prove the conjecture. The main theorem that we *do* prove in this work is that Conjecture B holds for *caterpillars*.

Theorem C. Conjecture B holds when the underlying tree of \mathcal{T} is a caterpillar.

Here we are using the following standard graph-theoretic definition:

Definition 1.6. A *caterpillar graph* is a tree in which every vertex has distance at most 1 from a central *spine* (path). Equivalently, a caterpillar is a graph of pathwidth 1. An example of a caterpillar tree is depicted in Figure 1.

We remark that caterpillar graphs arise quite naturally in some of the contexts where information flow trees are studied; for example, in Hidden Markov Models, where the leaf random variables are observed and the spine random variables are hidden.

1.2 Motivation and previous work

Besides being a natural problem in information theory, Conjecture A is motivated by certain problems in the algorithmic theory of convex relaxation hierarchies. We give here a very high-level sketch of the connection, as developed in the following works: [1–4, 6, 8–10].

Consider a Boolean optimization problem such as Max-Cut on a graph $G = (V, E)$, where we write $V = [n]$; the task is to find a ± 1 assignment x_1, \dots, x_n to the vertices so as to minimize $\text{avg}_{(u,v) \in E} x_u x_v$. This is a non-convex (and NP-hard) optimization problem. A natural algorithmic approach is to relax it to an (efficiently-solvable) convex optimization problem and then argue that the relaxed solution can be “rounded”

to a genuine ± 1 assignment with approximately the same value. Two important families of such relaxations are the Sherali–Adams LP relaxation and the SOS (Lasserre–Parrilo) SDP relaxation. The families have a tunable “degree” parameter $t \in \mathbb{Z}^+$; as t increases, the convex relaxations become tighter and tighter but the running time for solving them increases like $n^{O(t)}$.

Roughly speaking, solving these relaxations yields an optimal solution to the original Max-Cut problem, except that instead of getting a ± 1 assignment x_1, \dots, x_n , one gets a collection of “fake degree- t ± 1 -valued random variables” $\mathbf{X}_1, \dots, \mathbf{X}_n$. In fact, these are not random variables at all; they are merely a list of numbers ρ_S for all $S \subseteq [n]$ with $|S| \leq t$. However, there is a promise that for each such S there exists a collection of true ± 1 -valued random variables $(\mathbf{Y}_v)_{v \in S}$ with $\mathbf{E}[\prod_{v \in S} \mathbf{Y}_v] = \rho_S$. Thus, being very imprecise, an algorithm can act as though it has true random variables $\mathbf{X}_1, \dots, \mathbf{X}_n$, as long as it only ever uses them in groups of at most t .

The objective function minimized by the convex relaxation is $\alpha := \text{avg}_{(u,v) \in E} \rho_{\{u,v\}}$. An algorithm would now like to take the fake random variables and produce a genuine ± 1 assignment x_1, \dots, x_n which has, say, $\text{avg}_{(u,v) \in E} x_u x_v \leq \alpha + \epsilon$. A simple idea for doing this is to draw x_j according to \mathbf{X}_j , independently for each $j \in [n]$. (This counts as using the fake random variables in groups of size 1 and is thus legal since $t \geq 1$.) However in doing this we will get $\mathbf{E}[x_u x_v] = \mathbf{E}[\mathbf{X}_u] \mathbf{E}[\mathbf{X}_v] = \rho_{\{u\}} \rho_{\{v\}}$, which need not bear any relationship to the quantities $\rho_{\{u,v\}}$ entering into the definition of α . What would be desirable is if we had $|\rho_{\{u,v\}} - \rho_{\{u\}} \rho_{\{v\}}| \leq \epsilon$ for all pairs (u, v) , or at least on average over all pairs. In other words, we wish for the “average covariance” (as defined at the beginning of Section 1) of the fake random variables $\mathbf{X}_1, \dots, \mathbf{X}_n$ to be smaller than some ϵ . Of course it need not be, but Conjecture A implies that it can be made so, provided we are allowed to condition on some $t \leq O(1/\epsilon)$ randomly chosen \mathbf{X}_j 's. In the end, using the Sherali–Adams or SOS relaxations with degree parameter t would allow us to do this in time $n^{O(1/\epsilon)}$.

Thus we see that the quantitative dependence in Conjecture A directly relates to the running time of algorithms based on “correlation rounding” of Sherali–Adams/SOS hierarchies. An example consequence of Conjecture A (see [10]) would be that the Sherali–Adams LP hierarchy provides an arbitrarily good multiplicative approximation to Max-Cut on n -vertex, ϵn^2 -edge graphs in time $n^{O(1/\epsilon)}$. This gives a very nice tradeoff between density and running time, one that works almost all the way down to the “sparse” regime (i.e., $O(n)$ edges). On the other hand, using the weaker Theorem 1.1, the running time becomes $n^{O(1/\epsilon^2)}$. This is only nontrivial when $\epsilon \gg n^{-1/2}$; i.e., for graphs with $\omega(n^{3/2})$ edges.

We end this section by commenting on the Raghavendra–Tan proof [8] of Theorem 1.1. They study the analog $\text{avgInfo}_t(\mathbf{X})$ of $\text{avgCov}_t(\mathbf{X})$, in which $|\mathbf{Cov}(\mathbf{X}_u, \mathbf{X}_v)|$ is replaced by the *mutual information*, $I(\mathbf{X}_u; \mathbf{X}_v) \geq 0$. They deduce very simply from the definitions that for any $0 < T < n - 1$,

$$\sum_{t=0}^{T-1} \text{avgInfo}_t(\mathbf{X}) \leq 1.$$

This means that there exists a $t < T$ such that $\text{avgInfo}_t(\mathbf{X}) \leq 1/T$. The basic relationship $|\mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v]| \leq \sqrt{2} \sqrt{I(\mathbf{X}_u; \mathbf{X}_v)}$ lets them complete the proof Theorem 1.1 with a bound of $t \leq 2/\epsilon^2$. Thus we see that proving Conjecture A requires surmounting a familiar difficulty: the quadratic relationship between L_1 -distance and KL-distance.

Finally, while it's tempting to think that $\text{avgCov}_t(\mathbf{X})$ and $\text{avgInfo}_t(\mathbf{X})$ should be decreasing functions of t (thereby allowing us to fix t independently of $\mathbf{X}_1, \dots, \mathbf{X}_n$ in Theorem 1.1 and Conjecture A), this is not the case.

Proposition 1.7. *For any fixed integer $T \in \mathbb{Z}^+$, there exist random variables $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$, $n = T + 2$, such that $\text{avgCov}_t(\mathbf{X}) = 0$ for $t < T$ but $\text{avgCov}_T(\mathbf{X}) = 1$ (and similarly for avgInfo_t).*

Proof. We simply define $\mathbf{X}_1, \dots, \mathbf{X}_{T+2}$ to be uniformly random conditioned on $\mathbf{X}_1 \mathbf{X}_2 \cdots \mathbf{X}_{T+2} = 1$. Then consider any $J \subset [T+2]$ and any outcome of $(\mathbf{X}_j)_{j \in J}$. If $|J| < T$ then the remaining \mathbf{X}_k 's are (conditionally) pairwise independent. On the other hand, if $|J| = T$ then the remaining pair $(\mathbf{X}_u, \mathbf{X}_v)$ is either uniform on $\{(+1, +1), (-1, -1)\}$ or uniform on $\{(+1, -1), (-1, +1)\}$; in either case, the (conditional) covariance is 1. \square

1.3 Organization of this paper

In Section 2, we describe some basic transformations on information flow trees that preserve the joint distribution on the leaf random variables. These allow us to make certain convenient assumptions about the structure of our information flow trees in subsequent sections. Section 3 contains an explicit formula for the covariance of two leaves in an information flow tree conditioned on some outcome of the other leaves. In Section 4 we demonstrate that this expression is nondecreasing as a function of edge correlations along the spine. This essentially lets us reduce a caterpillar tree to an inhomogeneous star; we analyze the latter in Section 5. Finally, the proof of Theorem C is given in Section 6.

2 Information flow tree equivalences

Given an information flow tree, there are several ways it can be modified so that the joint distribution of its leaf random variables does not change. Since Conjecture B and Theorem C are only concerned with the leaf random variables, and not the “internal” random variables, we are free to make such modifications. We use the following definition:

Definition 2.1. Let \mathcal{T} and \mathcal{T}' be information flow trees, generating random variables $(\mathbf{X}_v)_{v \in V}$ and $(\mathbf{X}'_{v'})_{v' \in V'}$. Further, assume that V and V' have the same set of leaves, L (though \mathcal{T} and \mathcal{T}' may otherwise have different tree topologies and correlation functions). We say \mathcal{T} and \mathcal{T}' are *equivalent* if $(\mathbf{X}_\ell)_{\ell \in L}$ and $(\mathbf{X}'_\ell)_{\ell \in L}$ have the same joint distribution.

In this section we describe some transformations on general information flow trees that put them into simpler, equivalent forms.

The first two transformations allow us to assume without loss of generality that $\rho(e) \geq 0$ for all edges e , except possibly for edges incident to leaves. (In fact, allowing correlations in $[-1, 0)$ is not really an essential aspect of our model; the reader will not lose much by simply assuming $\rho \geq 0$ always.)

Lemma 2.2. *Let $\mathcal{T} = (V, E, \rho)$ be an information flow tree and let $w \in V$ be an internal vertex. Let $\mathcal{T}' = (V, E, \rho')$ be the information flow tree that is the same as \mathcal{T} except with $\rho'(e) = -\rho(e)$ for all edges e incident on w . Then \mathcal{T} and \mathcal{T}' are equivalent.*

Proof. Let $(\mathbf{X}_v)_{v \in V}$ and $(\mathbf{R}_e)_{e \in E}$ be the random variables generated by \mathcal{T} . Define:

$$\mathbf{R}'_e = \begin{cases} -\mathbf{R}_e & \text{if } e \text{ is incident to } w, \\ \mathbf{R}_e & \text{otherwise;} \end{cases} \quad \mathbf{X}'_v = \begin{cases} -\mathbf{X}_v & \text{if } v = w, \\ \mathbf{X}_v & \text{otherwise.} \end{cases}$$

It's easy to see that $(\mathbf{R}'_e)_{e \in E}$ has the correct joint distribution for \mathcal{T}' . It's then easy to see that $(\mathbf{X}'_v)_{v \in V}$ has the correct joint distribution for \mathcal{T}' . Since w is not a leaf, we have $\mathbf{X}_\ell = \mathbf{X}'_\ell$ for all leaves ℓ . Thus \mathcal{T} and \mathcal{T}' are equivalent. □

Lemma 2.3. *For every information flow tree $\mathcal{T} = (V, E, \rho)$, there is an equivalent one $\mathcal{T}' = (V, E, \rho')$ in which $\rho'(e) \geq 0$ for all “internal” edges e ; i.e., for edges e not touching a leaf.*

Proof. Given \mathcal{T} , choose a root vertex $r \in V$ arbitrarily. The idea is that in a top-down fashion starting from r , we “fix” all negative internal edges. Specifically, we apply the following procedure:

for $j = 1, 2, 3, \dots$,
 for each non-leaf vertex w at distance j from r ,
 if the parent edge e of w has $\rho(e) < 0$,
 apply the transformation from Lemma 2.3 to w .

It is easy to see that this procedure terminates with an equivalent information flow tree \mathcal{T}' in which all internal edges have a nonnegative correlation value. □

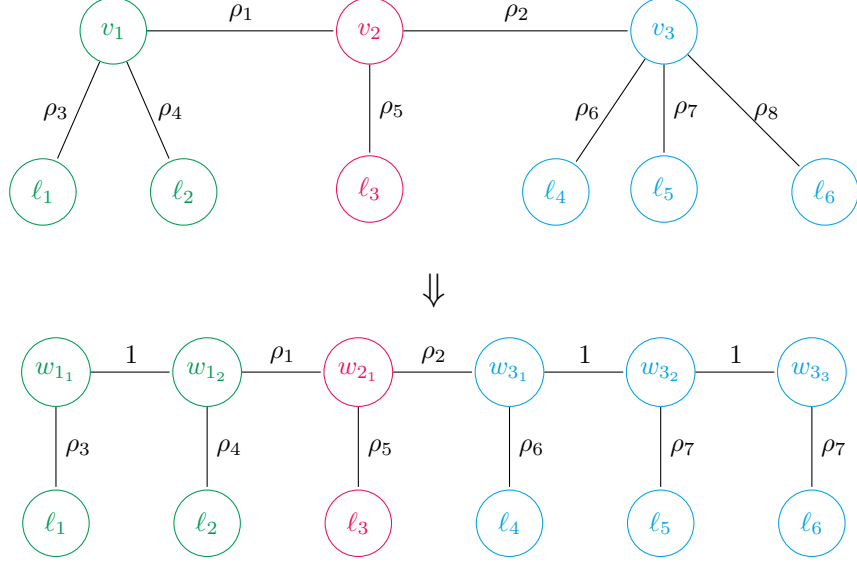


Figure 2: Applying the transformation of Lemma 2.7 to v_1 and v_3

Correctness of the next two transformations follows from the fact that if $\mathbf{R}_1, \mathbf{R}_2$ are independent $\{\pm 1\}$ -valued random variables with expectations ρ_1, ρ_2 , then $\mathbf{R}_1 \mathbf{R}_2$ is a $\{\pm 1\}$ -valued random variable with expectation $\rho_1 \rho_2$:

Lemma 2.4. *Suppose $\mathcal{T} = (V, E, \rho)$ is an information flow tree, $v \in V$ has degree 2, and (e_1, e_2) is the length-two path of edges through v . Modify \mathcal{T} by deleting v and replacing (e_1, e_2) with a single edge e satisfying $\rho(e) = \rho(e_1)\rho(e_2)$. Then the resulting information flow tree is equivalent to the original \mathcal{T} .*

Lemma 2.5. *Let $\mathcal{T} = (V, E, \rho)$ be an information flow tree and let $e \in E$. Modify \mathcal{T} by splitting e into a length-two path (e_1, e_2) with $\rho(e_1)\rho(e_2) = \rho(e)$. Then the resulting information flow tree is equivalent to the original \mathcal{T} .*

The next two transformations are similar and use the fact that along a path in which all correlations are 1, the vertex random variables are always equal.

Lemma 2.6. *Let $\mathcal{T} = (V, E, \rho)$ be an information flow tree and let $P = (V', E')$ be a connected subgraph of (V, E) in which $\rho(e) = 1$ for all $e \in E'$. (For us, P will typically be a path.) Assume V' does not contain leaves of V . Then the information flow tree gotten from \mathcal{T} by contracting V' into a single vertex is equivalent to \mathcal{T} .*

Lemma 2.7. *Let $\mathcal{T} = (V, E, \rho)$ be an information flow tree and let $v \in V$. For any $m \in \mathbb{Z}^+$, suppose we delete v and replace it with a path (w_1, \dots, w_m) whose edges are assigned correlation 1 by ρ . For each edge $e = (u, v)$ formerly attached to v , we replace it with $e = (u, w_i)$ for an arbitrarily chosen $i \in [m]$. Then the resulting information flow tree is equivalent to the original \mathcal{T} . An example of such a transformation is depicted in Figure 2.*

The next transformation allows us to convert to trees of maximum degree 3. Indeed, we can say slightly more.

Lemma 2.8. *For each information flow tree \mathcal{T} , there is an equivalent one \mathcal{T}' in which the underlying graph has maximum degree 3. Indeed, we can take \mathcal{T}' to be a rooted binary tree in which each internal node has exactly 2 children.*

Proof. Suppose v is a vertex in \mathcal{T} of degree $d > 3$. We apply Lemma 2.7 to v ; by taking $m = d$ when doing so, we have room to attach each of v 's neighbors to a different w_i . As a result, each w_i will have degree at most 3. Repeating this for each vertex of degree exceeding 3, we get an equivalent tree \mathcal{T}' of degree at most 3. By applying Lemma 2.5 to an arbitrary edge e , we can root the tree at newly created vertex. Finally, any vertices of degree 2 (other than the root) that remain can be eliminated using Lemma 2.4. \square

Finally, we will use the last lemma to simplify general caterpillar trees.

Definition 2.9. We say a caterpillar is *simple* if the spine has at least two vertices, and each vertex of the spine is attached to exactly one leaf.

Lemma 2.10. *For each information flow caterpillar \mathcal{T} , there is an equivalent information flow simple caterpillar \mathcal{T}' .*

Proof. We apply Lemma 2.8 to \mathcal{T} , taking care with one step: When replacing a high-degree spine vertex v , we insert the new path (w_1, \dots, w_m) as part of the spine, with v 's spine neighbor(s) attached appropriately at w_1 or w_m . Finally, the resulting binary tree will not quite be a simple caterpillar because the spine node furthest from the root will have two leaf children. To fix this we can simply take either of these two leaf edges and split it using Lemma 2.5, creating one more spine vertex. \square

Remark 2.11. It is also possible to convert any information flow tree \mathcal{T} into an “essentially” equivalent one $\mathcal{T}' = (V', E', \rho')$ which is *homogeneous* — meaning ρ' is a constant — and which has maximum degree 3. Since we won't use this, we merely sketch the conversion. Given \mathcal{T} , we can assume it has maximum degree 3 using the first part of Lemma 2.8. Next, fix $\rho' = -(1 - \delta)$ for some very small $\delta > 0$. Now for each edge $e \in E$, replace it with a path of length $k \in \mathbb{Z}^+$ so that $(\rho')^k$ is as close as possible to $\rho(e)$. Since we can make δ arbitrarily small, we can get all approximations $(\rho')^k \approx \rho(e)$ simultaneously as close as desired, yielding an “essentially” equivalent tree \mathcal{T}' . Note that we can't quite ensure all vertices of \mathcal{T}' have exactly two children: merging the degree-2 vertices of \mathcal{T}' would spoil its homogeneity property.

3 A formula for conditional covariance on general trees

Suppose $\mathbf{X}_{v_1}, \mathbf{X}_{v_m}$ are two vertex random variables in an information flow tree. Prior to any conditioning, it's easy to see that their covariance is equal to the product of $\rho(e)$ along the edges e joining v_1 and v_m . In this section we generalize this to a formula for their *expected* covariance when conditioned on any event that is comprised of several conditionally independent events.

Theorem 3.1. *Let $\mathcal{T} = (V, E, \rho)$ be a information flow tree, with associated vertex random variables $(\mathbf{X}_v)_{v \in V}$. Fix any path $P = (v_1, \dots, v_m)$ of vertices in \mathcal{T} , where $m \in \mathbb{Z}^+$. We think of P as partitioning \mathcal{T} into a sequence of subtrees \mathcal{T}_i , with \mathcal{T}_i rooted at v_i . For notational simplicity we write*

$$\mathbf{X}_i = \mathbf{X}_{v_i}, \quad \bar{\mathbf{X}} = (\mathbf{X}_1, \dots, \mathbf{X}_m), \quad \rho_i = \rho(v_i, v_{i+1}).$$

Let L_i be any event depending only on the random variable outcomes in \mathcal{T}_i . Write $L = L_1 \wedge L_2 \wedge \dots \wedge L_m$, and assume L has nonzero probability. Then

$$\text{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L] = \frac{\prod_{i=1}^{m-1} \rho_i \prod_{i=1}^m \Pr[L_i \mid \mathbf{X}_i = +1] \Pr[L_i \mid \mathbf{X}_i = -1]}{\Pr[L]^2} \quad (2)$$

Proof. Introducing the notation

$$\lambda_i^\pm = \Pr[L_i \mid \mathbf{X}_i = \pm 1], \quad (3)$$

we want to show that

$$\Pr[L]^2 \text{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L] = \prod_{i=1}^{m-1} \rho_i \prod_{i=1}^m \lambda_i^+ \lambda_i^-. \quad (4)$$

Recall that if (\mathbf{Y}, \mathbf{Z}) is a pair of random variables, $\mathbf{Cov}[\mathbf{Y}, \mathbf{Z}] = \frac{1}{2} \mathbf{E}[(\mathbf{Y} - \mathbf{Y}')(\mathbf{Z} - \mathbf{Z}')]$, where $(\mathbf{Y}', \mathbf{Z}')$ denotes an independent copy of (\mathbf{Y}, \mathbf{Z}) . Substituting this into (4), the identity we want to prove becomes

$$\begin{aligned} \prod_{i=1}^{m-1} \rho_i \prod_{i=1}^m \lambda_i^+ \lambda_i^- &= \Pr[L]^2 \cdot \frac{1}{2} \mathbf{E}[(\mathbf{X}_1 - \mathbf{X}'_1)(\mathbf{X}_m - \mathbf{X}'_m) \mid L] \\ &= \Pr[L]^2 \cdot \sum_{x, x' \in \{\pm 1\}^m} \Pr[\bar{\mathbf{X}} = x \mid L] \Pr[\bar{\mathbf{X}}' = x' \mid L] \cdot \frac{1}{2}(x_1 - x'_1)(x_m - x'_m) \\ &= \sum_{x, x' \in \{\pm 1\}^m} \Pr[\bar{\mathbf{X}} = x, L] \Pr[\bar{\mathbf{X}}' = x', L] \cdot \frac{1}{2}(x_1 - x'_1)(x_m - x'_m). \end{aligned} \quad (5)$$

We will prove (5) by induction on m . The base case, $m = 1$, is

$$\lambda_1^+ \lambda_1^- = \sum_{x, x' \in \{\pm 1\}} \Pr[\mathbf{X}_1 = x, L_1] \Pr[\mathbf{X}'_1 = x', L_1] \cdot \frac{1}{2}(x - x')^2. \quad (6)$$

To verify this, note that when $x = x'$ the summand in (6) is zero and when $x \neq x'$ the summand in (6) is $2 \Pr[\mathbf{X}_1 = +1, L_1] \Pr[\mathbf{X}_1 = -1, L_1]$. Thus the whole sum in (6) is indeed

$$\begin{aligned} &4 \Pr[\mathbf{X}_1 = +1, L_1] \Pr[\mathbf{X}_1 = -1, L_1] \\ &= 4(\Pr[L_1 \mid \mathbf{X}_1 = +1] \Pr[\mathbf{X}_1 = +1]) (\Pr[L_1 \mid \mathbf{X}_1 = -1] \Pr[\mathbf{X}_1 = -1]) = 4\lambda_1^+ \cdot \frac{1}{2} \cdot \lambda_1^- \cdot \frac{1}{2} = \lambda_1^+ \lambda_1^-. \end{aligned}$$

We now assume (5) holds for a given $m \in \mathbb{Z}^+$ and prove it for $m + 1$. Thus we need to show

$$\begin{aligned} \prod_{i=1}^m \rho_i \prod_{i=1}^{m+1} \lambda_i^+ \lambda_i^- &= \sum_{\substack{x \in \{\pm 1\}^m \\ x' \in \{\pm 1\}^m}} \sum_{\substack{x_{m+1} \in \{\pm 1\} \\ x'_{m+1} \in \{\pm 1\}}} \Pr[\bar{\mathbf{X}} = x, \mathbf{X}_{m+1} = x_{m+1}, L, L_{m+1}] \cdot \\ &\quad \Pr[\bar{\mathbf{X}}' = x', \mathbf{X}'_{m+1} = x'_{m+1}, L, L_{m+1}] \cdot \frac{1}{2}(x_1 - x'_1)(x_{m+1} - x'_{m+1}), \end{aligned} \quad (7)$$

Because of the information flow tree structure we have

$$\begin{aligned} \Pr[\bar{\mathbf{X}} = x, \mathbf{X}_{m+1} = x_{m+1}, L, L_{m+1}] &= \Pr[\bar{\mathbf{X}} = x, L] \Pr[\mathbf{X}_{m+1} = x_{m+1}, L_{m+1} \mid \mathbf{X}_m = x_m] \\ &= \Pr[\bar{\mathbf{X}} = x, L] \cdot \left(\frac{1}{2} + \frac{1}{2}\rho_m x_m x_{m+1}\right) \cdot \lambda_{m+1}^{x_{m+1}}, \end{aligned}$$

and similarly for x', x'_{m+1} . Thus the right-hand side of (7) is

$$\begin{aligned} &\sum_{\substack{x \in \{\pm 1\}^m \\ x' \in \{\pm 1\}^m}} \left(\Pr[\bar{\mathbf{X}} = x, L] \Pr[\bar{\mathbf{X}}' = x', L] \cdot \frac{1}{2}(x_1 - x'_1) \right. \\ &\quad \cdot \sum_{\substack{x_{m+1} \in \{\pm 1\} \\ x'_{m+1} \in \{\pm 1\}}} \left. \left(\frac{1}{2} + \frac{1}{2}\rho_m x_m x_{m+1}\right) \cdot \lambda_{m+1}^{x_{m+1}} \cdot \left(\frac{1}{2} + \frac{1}{2}\rho_m x'_m x'_{m+1}\right) \cdot \lambda_{m+1}^{x'_{m+1}} \cdot (x_{m+1} - x'_{m+1}) \right). \end{aligned} \quad (8)$$

Regarding the inner sum here (i.e., the second line in (8)), there is no contribution when $x_{m+1} = x'_{m+1}$; by a little algebra, the contribution from the two $x_{m+1} \neq x'_{m+1}$ summands is

$$\frac{1}{2} \lambda_{m+1}^+ \lambda_{m+1}^- \left((1 + \rho_m x_m)(1 - \rho_m x'_{m+1}) - (1 - \rho_m x_m)(1 + \rho_m x'_{m+1}) \right) = \lambda_{m+1}^+ \lambda_{m+1}^- \rho_m (x_m - x'_{m+1}).$$

Thus (8) (equivalently, the right-hand side of (7)) is

$$\rho_{m+1} \lambda_{m+1}^+ \lambda_{m+1}^- \sum_{x, x' \in \{\pm 1\}^m} \Pr[\bar{\mathbf{X}} = x, L] \Pr[\bar{\mathbf{X}}' = x', L] \cdot \frac{1}{2}(x_1 - x'_1)(x_m - x'_m) = \prod_{i=1}^m \rho_i \prod_{i=1}^{m+1} \lambda_i^+ \lambda_i^-,$$

by the induction hypothesis (5). □

We present an equivalent way to write (2) in the following corollary.

Corollary 3.2. *Suppose we are in the setting of Theorem 3.1. Then for any $x \in \{\pm 1\}^m$ and its bitwise negation $-x$,*

$$\mathbf{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L] = \prod_{i=1}^{m-1} \rho_i \cdot \frac{\Pr[\bar{\mathbf{X}} = x \mid L]}{\Pr[\bar{\mathbf{X}} = x]} \cdot \frac{\Pr[\bar{\mathbf{X}} = -x \mid L]}{\Pr[\bar{\mathbf{X}} = -x]}.$$

Proof. Beginning with (2), we have

$$\begin{aligned} \mathbf{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L] &= \prod_{i=1}^{m-1} \rho_i \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \mathbf{X}_i = +1]}{\Pr[L]} \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \mathbf{X}_i = -1]}{\Pr[L]} \\ &= \prod_{i=1}^{m-1} \rho_i \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \mathbf{X}_i = x_i]}{\Pr[L]} \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \mathbf{X}_i = -x_i]}{\Pr[L]}. \end{aligned}$$

By virtue of the information flow tree structure we have $\Pr[L_i \mid \bar{\mathbf{X}} = x] = \Pr[L_i \mid \mathbf{X}_i = x_i]$. Thus the above equals

$$\prod_{i=1}^{m-1} \rho_i \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \bar{\mathbf{X}} = x]}{\Pr[L]} \cdot \frac{\prod_{i=1}^m \Pr[L_i \mid \bar{\mathbf{X}} = -x]}{\Pr[L]} = \prod_{i=1}^{m-1} \rho_i \cdot \frac{\Pr[L \mid \mathbf{X} = x]}{\Pr[L]} \cdot \frac{\Pr[L \mid \mathbf{X} = -x]}{\Pr[L]}.$$

The proof is now completed by applying Bayes' theorem. □

4 A monotonicity property

Though the formula in Theorem 3.1 is quite precise, we will only use it in a rather “soft” way, to show a certain monotonicity property. Suppose we are in the setting of Theorem 3.1 and that L denotes a certain outcome for all the leaf random variables in the tree (these are the events of main interest for us). Assume for simplicity that the “path correlations” $\rho_1, \dots, \rho_{m-1}$ are all nonnegative. Then the formula from Theorem 3.1 implies that $\mathbf{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L]$ is also nonnegative, a fact that does not seem obvious a priori. We’ll in fact be interested in the *expected value* of this conditional covariance, over all the outcomes L .

The goal of this section is to show that this expected covariance can only increase if one of the “path correlations” ρ_i is increased. Though this fact seems “intuitive”, it’s also not a priori obvious; we’ve only been able to prove it with the aid of Theorem 3.1. Note that this monotonicity property is not immediately obvious from the formula in Theorem 3.1, since the expressions $\Pr[\bar{\mathbf{X}} = (\pm 1, \dots, \pm 1) \mid L]$ have an implicit dependence on the ρ_i ’s.

Theorem 4.1. *In the setting of Theorem 3.1, assume that $\rho_1, \dots, \rho_{m-1} \in [0, 1]$. Let $\bar{\mathbf{Y}} = (\mathbf{Y}_1, \dots, \mathbf{Y}_\ell)$ be the leaf random variables of \mathcal{T} . Then*

$$\mathbf{E} \left[\left[\mathbf{Cov}[\mathbf{X}_1, \mathbf{X}_m] \mid \mathbf{Y} \right] \right] \tag{9}$$

is a nondecreasing function of each ρ_i .

Proof. Let $y \in \{\pm 1\}^\ell$ be any potential outcome for the leaf random variables, and let L_y denote the event that $\bar{\mathbf{Y}} = y$. Then it’s easy to see that L_y has the factorizable form described in Theorem 3.1. (A slight annoyance is that it’s possible to have $\Pr[L_y] = 0$. However this can only happen if some pair $\mathbf{Y}_i, \mathbf{Y}_j$ is fully correlated; i.e., $\mathbf{Cov}[\mathbf{Y}_i, \mathbf{Y}_j] = \pm 1$. In this case, letting \mathbf{X}_i denote one of the ancestors of $\mathbf{Y}_i, \mathbf{Y}_j$ on path P , we have that \mathbf{X}_i is fully determined by *every* possible outcome y . In turn, this means \mathbf{X}_1 and \mathbf{X}_m are *independent* conditioned on every possible outcome y ; i.e., the random variable $\mathbf{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid \mathbf{Y}]$ is

identically 0. Then (9) is trivially a nondecreasing function of the ρ_i 's. Thus we may henceforth assume that no pair $\mathbf{Y}_i, \mathbf{Y}_j$ is fully correlated and hence that $\Pr[L_y] \neq 0$ for all $y \in \{\pm 1\}^\ell$, no matter what the ρ_i 's are.)

Rewriting identity (4), Theorem 3.1 equivalently states that for any $y \in \{\pm 1\}^\ell$,

$$\Pr[L_y] \cdot \left| \text{Cov}[\mathbf{X}_1, \mathbf{X}_m \mid L_y] \right| = \frac{\prod_{i=1}^{m-1} \rho_i \prod_{i=1}^m \lambda_i^+ \lambda_i^-}{\Pr[L_y]}. \quad (10)$$

(We are able to insert the absolute-value sign on the left because, as noted earlier, the right-hand side is evidently nonnegative.) By definition, our quantity of interest (9) is the sum of (10) over all $y \in \{\pm 1\}^\ell$. We'll in fact show that for *every* $y \in \{\pm 1\}^\ell$ and every $j \in [m-1]$, the quantity (10) is a nondecreasing function of ρ_j .

In the numerator of (10) we have that $\prod_{i \neq j} \rho_i$ is a nonnegative constant independent of ρ_j . The same is true of $\prod_{i=1}^m \lambda_i^+ \lambda_i^-$: by the definition (3), each $\lambda_i^{\pm 1}$ represents a probability that depends on y but not on any of the ρ_i 's. Thus it remains to show that

$$\frac{\rho_j}{\Pr[L_y]} = \frac{\rho_j}{\Pr[\bar{\mathbf{Y}} = y]} \quad (11)$$

is a nondecreasing function of ρ_j . Note that $\Pr[\bar{\mathbf{Y}} = y]$ implicitly depends on all of the ρ_i 's; in fact, it's a *linear* function of each of them. To see this, note that $\rho_j = \rho(v_j, v_{j+1})$ enters into the generation of \mathcal{T} 's random variables only through the edge random variable $\mathbf{R}_{v_j, v_{j+1}}$; thus we can write

$$\Pr[\bar{\mathbf{Y}} = y] = \left(\frac{1}{2} + \frac{1}{2}\rho_j\right) \Pr[\bar{\mathbf{Y}} = y \mid \mathbf{R}_{v_j, v_{j+1}} = +1] + \left(\frac{1}{2} - \frac{1}{2}\rho_j\right) \Pr[\bar{\mathbf{Y}} = y \mid \mathbf{R}_{v_j, v_{j+1}} = -1],$$

where the two conditional probabilities on the right do not depend on ρ_j . Thus we can express (11) as

$$\frac{\rho_j}{\Pr[\bar{\mathbf{Y}} = y]} = \frac{\rho_j}{b + c\rho_j} \quad (12)$$

for some numbers b, c not depending on ρ_j . Now a function of this form, $\frac{\rho_j}{b + c\rho_j}$, is nondecreasing if and only if $b \geq 0$; i.e., if and only if the denominator in (12) is nonnegative for $\rho_j = 0$. But indeed this quantity is nonnegative, being a probability. \square

We end this section by observing that although we have shown that (9) is an increasing function of the ‘‘path correlations’’ ρ_i , we actually expect it to be a *decreasing* function of $|\rho(e)|$ for all edges e *not* on the path between \mathbf{X}_{v_1} and \mathbf{x}_{v_m} . The intuition is that increasing one such $|\rho(e)|$ gives more information about its ancestor random variable \mathbf{X}_{v_i} on the path P . In turn, this should decrease the expected covariance between \mathbf{X}_{v_1} and \mathbf{X}_{v_m} . As an example, if v_i had just a single edge (v_i, ℓ) hanging off it, and $\rho(v_i, \ell)$ were increased to 1, then observing the leaf random variable \mathbf{X}_ℓ would determine \mathbf{X}_{v_i} completely. Thus \mathbf{X}_{v_1} and \mathbf{X}_{v_m} would become independent (covariance-0) conditioned on any observed outcome for \mathbf{X}_ℓ .

5 The inhomogeneous star

To motivate the result in this section, let's recall Conjecture B. Suppose we are given any information flow tree \mathcal{T} and we would like to upper-bound the expected covariance of some *particular* pair of leaves $\mathbf{Y}_u, \mathbf{Y}_v$. As we'll see, it's easy to reduce this to analyzing the expected covariance of the leaves' parents, call them $\mathbf{X}_{v_1}, \mathbf{X}_{v_m}$. Next, our monotonicity result Theorem 4.1 implies that this expected covariance can only increase if all edge-correlations along the path between $\mathbf{X}_{v_1}, \mathbf{X}_{v_m}$ were increased to 1. In this case, by Lemma 2.6 we could equivalently think of the entire path as being contracted into one internal random variable \mathbf{X}_0 .

Suppose now that the original tree was a caterpillar—in fact, by Lemma 2.10 we can assume it was a simple caterpillar. After contracting the path, the collection \mathcal{L} of leaves that were originally ‘‘between’’ \mathbf{Y}_u and \mathbf{Y}_v now hang directly off of \mathbf{X}_0 . The two parts of the caterpillar ‘‘to the outside’’ of \mathbf{X}_{v_1} and \mathbf{X}_{v_m} also hang off of \mathbf{X}_0 as gangly caterpillar-subtrees — but we plan on ignoring them. We only intend to analyze

the “inhomogeneous star” formed by \mathbf{X}_0 and \mathcal{L} . The hope will be that if there is “squared correlation” along the edges to \mathcal{L} , then conditioning on them will typically leave \mathbf{X}_0 with very small variance; equivalently, \mathbf{X}_{v_1} and \mathbf{X}_{v_m} will have very small covariance.

The following lemma concerning the inhomogeneous star uses well-known ideas, but as we do not have a reference for the exact statement, we give a proof.

Lemma 5.1. *Let \mathcal{T} be an information flow tree comprising a “star center” vertex with random variable \mathbf{X}_0 , as well as m leaf vertices with random variables denoted $\mathbf{Y}_1, \dots, \mathbf{Y}_m$. We allow \mathcal{T} to contain additional vertices not mentioned. Write $\rho_i \in [-1, 1]$ for the correlation between \mathbf{X}_0 and \mathbf{Y}_i , and write $\alpha = \sum_{i=1}^m \rho_i^2$. Then*

$$\mathbf{E} [\mathbf{Var}[\mathbf{X}_0 \mid \bar{\mathbf{Y}}] \leq 4 \exp(-\alpha/2),$$

where $\bar{\mathbf{Y}}$ denotes $(\mathbf{Y}_1, \dots, \mathbf{Y}_m)$.

Proof. Let us define the random variable

$$\mathbf{S} = \mathbf{S}(\bar{\mathbf{Y}}) = \text{sgn}(\rho_1 \mathbf{Y}_1 + \dots + \rho_m \mathbf{Y}_m).$$

(Take $\text{sgn}(0) = +1$ for definiteness.) For each $y \in \{\pm 1\}^m$ let’s write

$$p(y) = \Pr[\mathbf{X}_0 \neq \mathbf{S} \mid \bar{\mathbf{Y}} = y].$$

Once we condition on $\bar{\mathbf{Y}} = y$, the random variable \mathbf{S} becomes some fixed sign $s \in \{\pm 1\}$, and the random variable \mathbf{X}_0 takes on some conditioned distribution, call it \mathbf{Z} . Now since \mathbf{Z} is a $\{\pm 1\}$ -valued random variable we have

$$\mathbf{Var}[\mathbf{Z}] = 4 \Pr[\mathbf{Z} = +1] \Pr[\mathbf{Z} = -1] \leq 4 \Pr[\mathbf{Z} \neq s],$$

no matter what s is. Thus

$$\mathbf{Var}[\mathbf{X}_0 \mid \bar{\mathbf{Y}} = y] \leq 4p(y),$$

and so

$$\mathbf{E} [\mathbf{Var}[\mathbf{X}_0 \mid \bar{\mathbf{Y}}] \leq 4 \mathbf{E}[p(\bar{\mathbf{Y}})] = 4 \Pr[\mathbf{X}_0 \neq \mathbf{S}].$$

It thus remains to show that

$$\exp(-\alpha/2) \geq \Pr[\mathbf{X}_0 \neq \mathbf{S}] \geq \Pr[\mathbf{X}_0(\rho_1 \mathbf{Y}_1 + \dots + \rho_m \mathbf{Y}_m) \leq 0].$$

The two cases $\mathbf{X}_0 = \pm 1$ are symmetric, so we may assume $\mathbf{X}_0 = +1$. Then $\mathbf{Y}_1, \dots, \mathbf{Y}_m$ are independent $\{\pm 1\}$ -valued random variables with $\mathbf{E}[\mathbf{Y}_i] = \rho_i$, and we wish to show that

$$\Pr[\rho_1 \mathbf{Y}_1 + \dots + \rho_m \mathbf{Y}_m \leq 0] \leq \exp(-\alpha/2).$$

This follows immediately from Hoeffding’s inequality, applied to the random variables $\rho_i \mathbf{Y}_i \in [-\rho_i, \rho_i]$. \square

6 Proof of Theorem C

In this section we prove Theorem C. Let $\mathcal{T} = (V, E, \rho)$ be an information flow caterpillar tree with $t \geq 2$ leaves. By Lemma 2.10 we may assume \mathcal{T} is a simple caterpillar. By Lemma 2.3 we may assume that ρ has a nonnegative value on all spine edges of \mathcal{T} . We write $\mathbf{X}_1, \dots, \mathbf{X}_t$ for the vertex random variables along \mathcal{T} ’s spine and $\mathbf{Y}_1, \dots, \mathbf{Y}_t$ for the leaf random variables, with e_i denoting the edge between \mathbf{X}_i and \mathbf{Y}_i . We write $\mathbf{R}_i = \mathbf{X}_i \mathbf{Y}_i$ for the edge random variable that \mathcal{T} associates to e_i , and we write $\rho_i = \rho(e_i) = \mathbf{E}[\mathbf{R}_i]$. See Figure 3 for a depiction of this.

Recall that we wish to show

$$\text{avg}_{\substack{\text{distinct pairs} \\ u, v \in [t]}} \mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{Y}_u, \mathbf{Y}_v] \right| \mid (\mathbf{Y}_j)_{j \in [t] \setminus \{u, v\}} \right] \leq O(1/t). \quad (13)$$




Figure 3: A Simple Information Flow Caterpillar Tree

Let us suppose for some time that the pair $u, v \in [t]$ is fixed. For brevity we'll write $\mathcal{Y} = (\mathbf{Y}_j)_{j \in [t] \setminus \{u, v\}}$ for the leaf random variables other than $\mathbf{Y}_u, \mathbf{Y}_v$. Then

$$\begin{aligned} \mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{Y}_u, \mathbf{Y}_v] \right| \mid \mathcal{Y} \right] &= \mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{X}_u \mathbf{R}_u, \mathbf{X}_v \mathbf{R}_v] \right| \mid \mathcal{Y} \right] = \mathbf{E} \left[\left| \rho_u \rho_v \mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v] \right| \mid \mathcal{Y} \right] \\ &= |\rho_u| |\rho_v| \mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v] \right| \mid \mathcal{Y} \right] \end{aligned} \quad (14)$$

where the second equality uses that $\mathbf{R}_u, \mathbf{R}_v$ are independent of $(\mathbf{X}_u, \mathbf{X}_v, \mathcal{Y})$.

Given u, v , let \mathcal{T}' denote \mathcal{T} with edges e_u, e_v deleted. We may apply our monotonicity result Theorem 4.1 to \mathcal{T}' , with P being the spine path between \mathbf{X}_u and \mathbf{X}_v . (Note that we earlier arranged for all spine edges to have nonnegative correlation, as required for Theorem 4.1.) We conclude that *if* the edge correlations along P were raised to 1, this could only increase the quantity $\mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v] \right| \mid \mathcal{Y} \right]$ appearing in (14). We could further upper-bound this quantity as follows: Write \mathcal{T}_{uv} for the modification of \mathcal{T}' in which P is contracted to a single vertex with random variable called \mathbf{X}_0 (as in Lemma 2.6). Then by applying the inhomogeneous star result, Lemma 5.1 to \mathcal{T}_{uv} , we would get

$$\mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{X}_u, \mathbf{X}_v] \right| \mid \mathcal{Y} \right] = \mathbf{E} \left[\mathbf{Var}[\mathbf{X}_0] \mid \mathcal{Y} \right] \leq 4 \exp(-\alpha(u, v)/2),$$

where

$$\alpha(u, v) := \sum \{ \rho_i^2 : i \text{ is between } u \text{ and } v \}.$$

Putting these observations together, we conclude that for a fixed pair $u, v \in [t]$,

$$\mathbf{E} \left[\left| \mathbf{Cov}[\mathbf{Y}_u, \mathbf{Y}_v] \right| \mid (\mathbf{Y}_j)_{j \in [t] \setminus \{u, v\}} \right] \leq |\rho_u| |\rho_v| \cdot 4 \exp(-\alpha(u, v)/2).$$

Thus to complete the proof of (13) we need to show

$$(*) := \text{avg}_{\substack{\text{distinct pairs} \\ u, v \in [t]}} \{ |\rho_u| |\rho_v| \cdot \exp(-\alpha(u, v)/2) \} \leq O(1/t). \quad (15)$$

This is now simply a combinatorial problem concerning the list of numbers ρ_1, \dots, ρ_t .

We solve the problem as follows. First, we'd like to switch u and v to being drawn *without* replacement. Note that

$$(*) = \mathbf{E}_{\substack{\mathbf{u}, \mathbf{v} \sim [t] \\ \text{uniformly, independently}}} \left[\left\{ \begin{array}{ll} |\rho_{\mathbf{u}}| |\rho_{\mathbf{v}}| \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2) & \text{if } \mathbf{u} \neq \mathbf{v}, \\ (*) & \text{if } \mathbf{u} = \mathbf{v}. \end{array} \right\} \right]$$

Since $|\rho_{\mathbf{u}}| |\rho_{\mathbf{v}}| \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2) \in [0, 1]$ always, and since $\mathbf{Pr}[\mathbf{u} = \mathbf{v}] = 1/t$, the above differs from

$$\mathbf{E}_{\substack{\mathbf{u}, \mathbf{v} \sim [t] \\ \text{uniformly, independently}}} \left[|\rho_{\mathbf{u}}| |\rho_{\mathbf{v}}| \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2) \right] \quad (16)$$




Figure 4: A small example illustrating the indices of $U_k(\mathbf{v})$. The label on edge e is ρ_e^2 .

by at most $2/t$. Thus to show (15), it suffices to upper-bound (16) by $O(1/t)$. To do this, we first apply Cauchy–Schwarz, obtaining

$$\begin{aligned} \mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} [|\rho_{\mathbf{u}}| |\rho_{\mathbf{v}}| \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2)] &\leq \sqrt{\mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} [\rho_{\mathbf{u}}^2 \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2)]} \sqrt{\mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} [\rho_{\mathbf{v}}^2 \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2)]} \\ &= \mathbf{E}_{\mathbf{u} \sim [t]} [\rho_{\mathbf{u}}^2 \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2)], \end{aligned} \quad (17)$$

the last equality because \mathbf{u} and \mathbf{v} are symmetrically distributed. Let’s introduce the following events:

$$A_0 = “\alpha(\mathbf{u}, \mathbf{v}) \in [0, 1]”, \quad A_k = “\alpha(\mathbf{u}, \mathbf{v}) \in [2^{k-1}, 2^k]”, \quad k \in \mathbb{Z}^+.$$

Using the fact that $\sum_{k \geq 0} \mathbf{1}_{A_k} \equiv 1$, we have that (17) equals

$$\mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} \left[\rho_{\mathbf{u}}^2 \cdot \sum_{k \geq 0} \mathbf{1}_{A_k} \cdot \exp(-\alpha(\mathbf{u}, \mathbf{v})/2) \right] \leq \mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} \left[\rho_{\mathbf{u}}^2 \cdot \sum_{k \geq 0} \mathbf{1}_{A_k} \cdot e^{1/4} \exp(-2^{k-2}) \right].$$

Here we essentially lower-bounded $\alpha(\mathbf{u}, \mathbf{v})/2$ by 2^{k-2} on the event that A_k occurs — except, that is not quite correct when $k = 0$; this why we included the factor $e^{1/4}$, to cover the $k = 0$ case. Thus it remains to show

$$\sum_{k \geq 0} \exp(-2^{k-2}) \cdot \mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} [\rho_{\mathbf{u}}^2 \cdot \mathbf{1}_{A_k}] \leq O(1/t). \quad (18)$$

To show this, let’s consider a fixed integer $k \geq 0$ and imagine that in the expectation, $\mathbf{v} \sim [t]$ is chosen first. Once \mathbf{v} is chosen, we define interval $U_k^-(\mathbf{v}) \subseteq [t]$ to be the set of all possible $\mathbf{u} < \mathbf{v}$ such that the event A_k occurs. We define $U_k^+(\mathbf{v})$ similarly, but for $\mathbf{u} > \mathbf{v}$. Figure 4 shows a small example. Denote the union of $U_k^-(\mathbf{v})$ and $U_k^+(\mathbf{v})$ by $U_k(\mathbf{v})$.

Furthermore, we must have

$$\sum_{u \in U_k(\mathbf{v})} \rho_u^2 = \sum_{u \in U_k^+(\mathbf{v})} \rho_u^2 + \sum_{u \in U_k^-(\mathbf{v})} \rho_u^2 \leq 2^k + 2^k = 2^{k+1} \quad \forall k \geq 0.$$

It follows that we have the upper bound

$$\mathbf{E}_{\mathbf{u}, \mathbf{v} \sim [t]} [\rho_{\mathbf{u}}^2 \cdot \mathbf{1}_{A_k}] = \mathbf{E}_{\mathbf{v} \sim [t]} \left[\sum_{u \in U_k(\mathbf{v})} \Pr[\mathbf{u} = u] \rho_u^2 \right] = (1/t) \mathbf{E}_{\mathbf{v} \sim [t]} \left[\sum_{u \in U_k(\mathbf{v})} \rho_u^2 \right] \leq (1/t) \mathbf{E}_{\mathbf{v} \sim [t]} [2^{k+1}] = 2^{k+1}/t.$$

Substituting this into (18), it remains to observe that indeed

$$\sum_{k \geq 0} \exp(-2^{k-2}) \cdot 2^{k+1} \leq O(1).$$

The proof of Theorem C is complete.

7 Conclusions

Lacking any directions for proving the main Conjecture A, we believe that Conjecture B (the case of general information flow trees) is a good place to start. Having proved Theorem C (the case of caterpillars), a natural next case to consider is an information flow tree with the property that each leaf is at distance at most *two* from a central spine. By the transformations in Section 2, it suffices to consider the case that each spine node has a single edge hanging off it, which in turn has an inhomogeneous star hanging off it. Perhaps some of the “reconstruction” results from [5] in terms of effective electrical resistance could be of use here. Another interesting special case of Conjecture B that one could try to resolve is that of a complete binary tree in which all edge correlations have the same value ρ . (This is the most heavily-studied information flow tree.) We believe that this case satisfies Conjecture B by a wide margin for all ρ , even with a sub-inverse-polynomial bound in place of $O(1/t)$. Perhaps the formula in our Theorem 3.1 could help prove this.

Acknowledgments

We thank Yuan Zhou for several early discussions on the topic of this paper.

References

- [1] Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better balance by being biased: A 0.8776-approximation for max bisection. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 277–294. SIAM, 2013.
- [2] Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing*, pages 307–326. ACM, 2012.
- [3] Boaz Barak, Jonathan A. Kelner, and David Steurer. Rounding sum-of-squares relaxations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 31–40. ACM, 2014.
- [4] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 472–481. IEEE, 2011.
- [5] William Evans, Claire Kenyon, Yuval Peres, and Leonard J. Schulman. Broadcasting on trees and the Ising model. *The Annals of Applied Probability*, 10(2):410–433, 2000.
- [6] Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 482–491. IEEE, 2011.
- [7] Elchanan Mossel. Survey: Information flow on trees. In *Graphs, morphisms and statistical physics*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, pages 155–170. AMS, Providence, RI, 2004.
- [8] Prasad Raghavendra and Ning Tan. Approximating CSPs with global cardinality constraints using SDP hierarchies. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 373–387. SIAM, 2012.
- [9] Thomas Rothvoß. The Lasserre hierarchy in approximation algorithms. *Lecture Notes for the MAPSP 2013 Tutorial*, 2013.
- [10] Yuichi Yoshida and Yuan Zhou. Approximation schemes via Sherali-Adams hierarchy for dense constraint satisfaction problems and assignment problems. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 423–438. ACM, 2014.