Random 2-SAT Solution Components and a Fitness Landscape

Damien Pitman

Random 2-SAT Solution Components and a Fitness Landscape

Damien Pitman[†]

State University of New York, College at Cortland, USA

We describe a limiting distribution for the number of connected components in the subgraph of the discrete cube induced by the satisfying assignments to a random 2-SAT formula. We show that, for the probability range where formulas are likely to be satisfied, the random number of components converges weakly (in the number of variables) to a distribution determined by a Poisson random variable. The number of satisfying assignments or solutions is known to grow exponentially in the number of variables. Thus, our result implies that exponentially many solutions are organized into a stochastically bounded number of components. We also describe an application to biological evolution; in particular, to a type of fitness landscape where satisfying assignments represent viable genotypes and connectivity of genotypes is limited by single site mutations. The biological result is that, with probability approaching 1, each viable genotype is connected by single site mutations to an exponential number of other viable genotypes while the number of viable clusters is finite.

Keywords: 2-SAT, random graphs, components, clusters, fitness landscape, incompatibilitiies

1   Introduction

We are interested in the connective properties of a random subgraph of the discrete cube Q^n. This random subgraph Q_F is induced on those vertices that correspond to literal assignments that satisfy a 2-CNF formula $F = F_p$ that is randomly determined by letting each clause be included in F with probability p. Let N_n be the number of connected components in Q_F.

Theorem 1 *Let F, Q_F and N_n be randomly determined as above. If $p = c/2n$ and $c < 1$ is constant, then N_n converges weakly to $N = 2^\Psi$ where Ψ is a Poisson random variable with mean*

$$\lambda = -\frac{1}{2}\left(\ln\left(1 - c\right) + c\right).$$

In particular, the probability that there is a unique cluster converges to

$$e^{-\lambda} = \sqrt{(1 - c)e^c}.$$

[†]Email: `Damien.Pitman@Cortland.edu`

Our interest in the component structure is motived by the theory of *fitness landscapes*. The notion of fitness landscapes was introduced by a theoretical evolutionary biologist, Sewall Wright in [17] (see also [14, 11]). The study of fitness landscapes has proved extremely useful both in biology and well outside of it. In the standard interpretation, a fitness landscape is a relationship between a set of genes (or a set of quantitative characters) and a measure of fitness (e.g. viability, fertility, or mating success). In Wright's original formulation the set of genes (or quantitative characters) is the property of an individual. However, the notion of fitness landscapes can be generalized to the level of a mating pair, or even a population of individuals [11]. For further information one can look at the book *Fitness Landscapes and the Origin of Species* [11] or the paper [13], which is a sort of biological sister to our discussion here.

It turns out that k-SAT is also of interest to physicists working in statistical mechanics and it is somewhat surprising that the component structure for random 2-SAT has not been examined previously because the component structure for 3-SAT has been explored. *Replica symmetry breaking* is a method from statistical mechanics that is not strictly rigorous, but has been used to find two phase transitions in 3-SAT [5]. In one phase, all satisfying assignments form a single component in the n-cube. In the next, the number of components is exponential in n. In the final phase, there are no satisfying assignments. Since our result shows that the number of components is stochastically bounded for 2-SAT, it seems that the structures are quite different for 2-SAT and 3-SAT.

We organize the remainder as follows. In Section 2 we introduce notation and provide the necessary background and results in both 2-SAT and fitness landscapes. We also describe the connection between the two topics. Section 3 pertains to the number of clusters and contains the proof of our result.

2   Background

2.1   Fitness Landscapes

The particular model of fitness landscapes that we consider here begins with an idealized collection of all genotypes that are haploid and diallelic at n loci. A *haploid* organism is one for which there is only one copy of each chromosome. *Alleles* are the potential versions of a gene at a specified locus. In our model, for any specified n loci, there exists a list L of all the unordered pairs of strictly distinct alleles that are *incompatible*. Each genotype for these n loci is viable if none of its allele pairs is incompatible and inviable otherwise. Each list L is an instance of a fitness landscape on the given genotype space.

We are primarily interested in two questions relating to a given list L: how many viable genotypes are there and how well connected are they through single locus mutations? With regards to connectivity, we say that two genotypes are in the same cluster if and only if they are connected by a path of single locus mutations through viable genotypes. Thus, each cluster of viable genotypes corresponds to a collection of genotypes which can each evolve into any other genotype in the cluster without passing through an inviable genotype and without altering two or more alleles simultaneously. Since there are n loci and two alleles at each locus, our genotype space is in bijection with the set of binary strings of length n or the vertices of the discrete cube Q^n. In addition, since the only allowed mutations are single locus mutations (which are equivalent to bit flips) we can represent these mutations by edges in Q^n. The viable genotypes and mutations between them form an induced subgraph of Q^n and the clusters are the connected components.

2.2   2-SAT

We begin with n Boolean variables $\xi_1, \ldots, \xi_n$. For each $1 \le i \le n$, the two truth values of ξ_i are referred to as the *positive literal* and the *negative literal* of ξ_i. We will be using a non-standard notation to refer to these literals. We let 0 be the negative literal and 1 be the positive literal and then put the index in the subscript, i.e., for $1 \le i \le n$

$$0_i \text{ is the negative literal of } \xi_i$$
$$1_i \text{ is the positive literal of } \xi_i$$

The negation of any literal x is denoted $\overline{x}$. Thus, $\overline{0_i} = 1_i$ and $\overline{1_i} = 0_i$. Literals of distinct variables are said to be *strictly distinct*. An assignment of n strictly distinct literals to n Boolean variables is referred to as a truth assignment to the n variables.

A 2-clause $c = x \vee y$ is the disjunction of two literals x and y. Generally these literals need not be distinct, but here we assume x and y are strictly distinct. Thus, there are

$$2^2 \binom{n}{2} = 2n(n-1)$$

possible 2-clauses. Suppose we have some number $m \ge 0$ of 2-clauses $c_1, \ldots, c_m$. Then

$$F = c_1 \wedge \ldots \wedge c_m$$

is a 2-SAT formula. Again, we assume that the clauses are distinct. If there is some literal assignment to the n variables for which the formula F is TRUE, then it is said to be *satisfied* or SAT. Determining if a formula is satisfiable is known as the 2-SAT problem.

The 2-SAT problem is the simplest of the k-SAT problems. For any $k \ge 2$, a k-SAT formula is the conjunction of m disjunctions where the number of literals per disjunction is k. The k-SAT problem, like the 2-SAT problem, is to determine whether or not a given formula is satisfiable. Of course, one can consider other types of satisfiability problems as well. Satisfiability problems in general have been important in theoretical computer science since 1971, when Cook established that they are NP-complete [9]. The most studied type of satisfaction problem is k-SAT because most other models were proven to generate easy-to-satisy formulas. The 2-SAT problem, unlike k-SAT problems for $k > 2$, is not NP-complete. In fact, solutions can be found in linear time [3].

2.3   2-SAT as a Fitness Landscape

Let each Boolean variable ξ_i be associated with a genetic locus i and let the literals 0_i and 1_i represent the two alleles at that locus. Also, let us write that xy is an incompatibility if the alleles x and y are incompatible. Then, notice that if a formula F contains the clause $\overline{x} \vee \overline{y}$ then no satisfying assignment to n variables contains both x and y. Mapping each clause $\overline{x} \vee \overline{y}$ to the corresponding incompatibility xy leads to a bijection between the set of 2-SAT formulas and the set of lists of incompatibilities. It follows that if L is a list of incompatibilities and F is the 2-SAT formula with the corresponding clauses, then a genotype with no allele pairs on L is associated with a literal assignment that satisfies the formula F and vice versa. Consequently, we can let F refer to either the formula or the list of incompatibilities without distinction and we can think of the biological questions in terms of 2-SAT or the other way around.

2.4   Random 2-SAT and Results

Determining the probability that a randomly chosen formula F is satisfiable is a random 2-SAT problem. Throughout this paper, F is chosen by letting each possible clause be included independently with probability

$$p = \frac{c}{2n}.$$

where $c > 0$. In general, c may depend on n, but all results stated here hold only for constant c. The reason p is written as $c/2n$ is that there is a phase transition from satisfiable to unsatisfiable where the expected number of clauses is asymptotically equivalent to the number of literals $2n$.

Theorem 2 *Let $\epsilon > 0$ be constant and F be a formula chosen randomly as described above. Then the following statements hold with probability approaching 1 as n tends to infinity.*

- *If $c \geq 1 + \epsilon$ then F is not satisfiable.*

- *If $c \leq 1 - \epsilon$ then F is satisfiable.*

This is the fundamental result for random 2-SAT and was proved independently by three different groups around 1992 [8, 10, 12]. The model examined in these papers is actually slightly different, but is asymptotically equivalent to ours in that both exhibit the same phase transition at equivalent critical values. This equivalence and the best current results are described in [7], where $\epsilon > 0$ is allowed to vary with n and the size of the finite scaling window is described.

   We are also interested in the number of satisfying assignments in the SAT phase. The following result is proved in [13].

Theorem 3 *The number of satisfying assignments to a satisfiable 2-SAT formula is at least*

$$\exp((e^{-2c} + ce^{-4c})\log 2 - \delta)n)$$

where $c < 1$ and $\delta > 0$ are constants.

   And now for the context of random 2-SAT. Random satisfiability problems became popular as a way to study the hardness of "typical" formulas. A particularly interesting fact about random k-SAT problems is that all random k-SAT problems undergo a phase transition similar to the transition for 2-SAT. More on this can be found in [1], which is a nice paper on both the latest research on k-SAT and some general satisfiability context.

2.5   Digraphs

2.5.1   The Standard Digraph for 2-SAT

Every 2-SAT formula F can be associated to a directed graph D_F as defined by Aspvall et al [3]. The proofs of Theorem 2 and the remainder of our discussion here are based on this digraph.

- The vertex set of D_F is the set of literals $\{0_1, \ldots, 0_n, 1_1, \ldots, 1_n\}$.

- For each clause $x \vee y$ in F the directed edges $\overline{x} \to y$ and $\overline{y} \to x$ are included in D_F.

To clarify this definition, consider the clause $x \vee y$: a literal assignment including $\overline{x}$ must include y to be satisfying and an assignment with literal $\overline{y}$ must have x. Thus, the two edges in the definition correspond to the two implications naturally associated with the clause $x \vee y$: $x \Rightarrow \overline{y}$ and $y \Rightarrow \overline{x}$.

If there is a directed path from x to y in D_F it is denoted by $x \rightsquigarrow y$. Notice that this path is in D_F if and only if the path $\overline{y} \rightsquigarrow \overline{x}$ is also in D_F. Also, like the edges of D_F, these paths do correspond to logical implications corresponding to the satisfiability of F. The set of literals and edges that can be reached from x is referred to as the out-graph $D_F^+(x)$ and the literal set of D_F^+ is denoted

$$L^+(x) = L_F^+(x) = \{y : x \rightsquigarrow y\}.$$

These are the literals that must be present in any satisfying assignment with literal x. Of course, an assignment may have all the literals in $L^+(x)$ and yet still not be satisfying due to an unrelated clause. On the other hand, those literals that require x to be present in any satisfying assignment are denoted

$$L^-(x) = L_F^-(x) = \{y : y \rightsquigarrow x\}.$$

and the corresponding graph is called the in-graph of x and denoted $D^-(x)$.

We say that two vertices x and y in a digraph are *weakly connected* if there is an undirected path from x to y in the digraph, i.e., x and y are weakly connected in the digraph if they are connected in the undirected (multi)graph that is formed by replacing directed edges with undirected edges. The components of this undirected graph formed from D_F are called the *weak components* of D_F. Notice that x and y may be in the same weak component without having a directed path from x to y or from y to x.

Suppose that $y \in L_F^+(x) \cap L_F^-(x)$, so $x \rightsquigarrow y$ and $y \rightsquigarrow x$. In this case x and y are said to be *strongly connected*. Consider the set of literals

$$C_F(x) = \{y : x \rightsquigarrow y \rightsquigarrow x\}$$

that are strongly connected to x. The subgraph of D_F induced by these literals is a *strongly connected component*. We will refer to both the strongly connected component and its set of literals as the strong component of x. We follow the convention that $x \rightsquigarrow x$ for every vertex $x \in D_F$ so that every vertex is a member of a strong component and $C_F(x)$ is well defined for every literal and the set of all $2n$ literals can be partitioned into strong components of D_F. When a strong component consists of a single literal we will say that the strong component is *trivial*. Otherwise the strong component is considered to be *nontrivial*. The set of all strong components can be partially ordered. Let

$$C_F(x) \leq C_F(y) \quad \text{if} \quad x \rightsquigarrow y,$$

which also implies that $x' \rightsquigarrow y'$ for any $x' \in L_F^-(x) \supseteq C_F(x)$ and $y' \in L_F^+(y) \supseteq C_F(y)$.

For any set of literals M, we define

$$\overline{M} = \{\overline{y} : y \in M\}.$$

When $M \cap \overline{M} = \emptyset$, we say that M and $\overline{M}$ are *complementary*. Notice that if $y \in C_F(x)$, then $\overline{y} \in C_F(\overline{x})$ so $C_F(\overline{x}) = \overline{C_F(x)}$ for every literal x. We also see that

$$C_F(x) \cap \overline{C_F(x)} \neq \emptyset \iff C_F(x) = \overline{C_F(x)}.$$

If there is a literal x such that $C_F(x) = \overline{C_F(x)}$, then there are literals $y, z \notin \{x, \overline{x}\}$ (not necessarily distinct) such that

$$x \rightsquigarrow y \rightsquigarrow \overline{x} \rightsquigarrow z \rightsquigarrow x$$

and we say that x (and each $y \in C_F(x)$) is on a *contradictory cycle* in D_F. Notice that this cycle need not be simple (have the same number of literals as edges). Since the literal x depends upon its complement $\overline{x}$ to be satisfying and vice versa, neither literal can be assigned to the corresponding variable and there are no satisfying assignments. On the other hand, it can be shown that if there are no contradictory cycles, then there is at least one satisfying assignment (see [7] for example). This is perhaps the most fundamental result in the study of 2-SAT.

Theorem 4 *A formula F is satisfiable if and only if there is no contradictory cycle in D_F.*

2.5.2   The Digraph of the Partial Ordering of Strong Components O_F

Since a formula F is satisfiable exactly when there is no literal x such that $C_F(x) = \overline{C_F(x)}$, we see that having a satisfiable formula is equivalent to having every strong component of the digraph being a member of a complementary pair $\{C_F(x), \overline{C_F(x)}\}$. It follows that, given any satisfiable formula F and any strong component C of D_F, a necessary (but not sufficient) condition for an assignment to be satisfying is that it must have either have all the literals of C or else all of the literals of $\overline{C}$. This means that, while there are $2^{|C|}$ potentially satisfying (partial) assignments to the variables of C, we know that there are actually only two possible combinations of the literals on these variables that might correspond to satisfying assignments. To indicate the assignment of all literals in C (or $\overline{C}$) we will simply state that we have assigned C (or $\overline{C}$). Now suppose that there are s complementary strong component pairs determined by a satisfiable formula F, and $\{C_1, \overline{C_1}\} \ldots \{C_s, \overline{C_s}\}$ is a listing of these. Then every satisfying assignment must satisfy

$$(C_1 \vee \overline{C_1}) \wedge \ldots \wedge (C_s \vee \overline{C_s}). \tag{1}$$

It follows that there are no more than 2^s assignments that satisfy F. In this manner, it makes sense to think of complementary strong components as a tool for reducing the dimension of the solution space.

Continuing in this vein, we find it helpful to consider the directed graph corresponding to the partial order on the strong components of D_F, which we will denote by O_F. The vertex set of O_F is the set of strong components of D_F and there are edges $X \to Y$ in O_F exactly when X and Y are strong components of D_F such that there exist literals x in X and y in Y such that $x \to y$ is an edge of D_F. Thus the edges of O_F correspond to implications that are represented in D_F. Moreover, the implications corresponding to O_F in conjunction with the statement (1) form a statement that is logically equivalent to F. In much of Section 3 we will effectively ignore the internal structure of the strong components that determines (1) and focus our attention on the ordering of the strong components represented in O_F.

We will relax our language a bit and write that a collection of strong components $\mathcal{C} = \{C_1, \ldots, C_i\}$ is a partial assignment to indicate that all the literals in $\cup_i C_i$ have been assigned. Similarly, we will use $L^+(\mathcal{C})$ to refer to $L^+(\cup_i C_i)$, and so on. Given two partial assignments $\mathcal{C}$ and $\mathcal{D}$, we notice that $\mathcal{C} \cup \mathcal{D}$ is a partial assignment without contradiction if and only if $L^+(\mathcal{C}) \cup L^+(\mathcal{D})$ is strictly distinct. A related result is

Lemma 5 *If $\mathcal{X}$ is a collection of strong components such that $L^+(\mathcal{X}) = \mathcal{X}$ and $\mathcal{Y}$ is any collection of strong components that is strictly distinct from $\mathcal{X}$, then $L^+(\mathcal{Y}) \cap \overline{\mathcal{X}} = \emptyset$. Thus, under the given premises, $\mathcal{X} \cup \mathcal{Y}$ can be assigned without contradiction if and only if $L^+(\mathcal{Y})$ is strictly distinct.*

Proof: If this were not the case, then there are strong components $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ such that $Y \leq \overline{X}$, which is equivalent to $X \leq \overline{Y}$. Since $\mathcal{X} = L^+(\mathcal{X})$ this means that $\overline{Y} \in \mathcal{X}$; which means that $\mathcal{X}$ and $\mathcal{Y}$ were not strictly distinct. □ This result allows us to make the assignment $\mathcal{X}$

for any collection of strong components that satisfies $\mathcal{X} = L^+(\mathcal{X})$ and then consider partial assignments to the remaining variables without consideration of the variables of $\mathcal{X}$.

Clearly, we are interested in the structure of the partial order O_F. A *chain* is a totally ordered subset of a partially ordered set, which means that any two elements in the chain are comparable in order. An *antichain* is a subset $\mathcal{A}$ of a partially ordered set such that any two elements in $\mathcal{A}$ are incomparable. A *maximal* (anti)chain is an (anti)chain that is not a proper subset of any other (anti)chain. A *maximum* (anti)chain is an (anti)chain that has cardinality at least as large as every other (anti)chain. The *height* of a partial order is the cardinality of a maximum chain. The *width* of a partial order is the cardinality of a maximum antichain. We will use the term *complementary antichains* to refer to a pair $\{\mathcal{A}, \overline{\mathcal{A}}\}$ where $\mathcal{A}$ is a strictly distinct collection of strong components where each element in $\mathcal{A}$ is incomparable to every other element in $\mathcal{A} \cup \overline{\mathcal{A}}$. We will refer to the cardinality of a maximum complementary antichain as the *complementary width* of O_F.

Lemma 6 *Suppose F is satisfiable. If $\{\mathcal{A}, \overline{\mathcal{A}}\}$ is a pair of complementary antichains in O_F where $\mathcal{A} = \{C_1, \ldots, C_l\}$, and $\mathcal{X} = \{X_1, \ldots, X_l\}$ is any partial assignment where $X_i \in \{C_i, \overline{C_i}\}$ for $1 \leq i \leq l$, then there exists a satisfying assignment containing $\mathcal{X}$.*

Proof: Notice that $L^+(\mathcal{X})$ is strictly distinct. Indeed, if not, then there is a component Z such that there exist X_i and X_j from $\mathcal{X}$ such that $X_i \leq Z$ and $X_j \leq \overline{Z}$. And so we have $\overline{Z} \leq \overline{X_i}$, which gives us $X_j \leq \overline{X_i}$. Since this contradicts the complementary nature of $\mathcal{A}$, the outgraph is strictly distinct and we can select all the literals of $L^+(\mathcal{X})$ for a partial assignment $\mathcal{A}$ that is without contradictions. Notice that $\mathcal{A} = L^+(\mathcal{A}) = L^+(\mathcal{X})$, so $\overline{\mathcal{A}} = L^-(\overline{\mathcal{A}})$. Thus, once we have selected $\mathcal{A}$, we can not choose a strong component Y from $L^-(\overline{\mathcal{A}})$ (because we have already chosen the complements). This means that the choice between any strong component pair $\{Y, \overline{Y}\}$ beyond $\mathcal{A}$ can be made independent of $\mathcal{A}$. Thus, we can simply invoke the satisfiability of F to be sure that there exists some satisfying assignment containing $\mathcal{A}$ which of course contains $\mathcal{X}$. □

Corollary 7 *Suppose F is satisfiable. If k is the complementary width of O_F, then there are at least 2^k assignments that satisfy F.*

Proof: Let $\mathcal{A}$ be a maximum complementary antichain. We observe that there are 2^k distinct partial assignments $\mathcal{X}$ as described in Lemma 6 that correspond to $\mathcal{A}$. Thus, Lemma 6 ensures the existence of 2^k distinct assignments that satisfy F. □

3   The Number of Components

The purpose of our discussion here is ultimately to study the number of connected components of Q_F. Besides our innate mathematical interest in structure, we are also interested in connected components of Q_F for biological reasons. Recall that in our biological model, literals correspond to alleles, so complementary strong components represent alternate strategies for survival. If we anthropomorphize evolution,

then we can think of the evolution of an organism as a sequence of choices between these alternate strategies. We imagine the set of satisfying assignments as the collection of all possible viable strategies for survival. We are interested in the topology of Q_F as opposed to a more dense graph because allowing for edges only between assignments that differ on a single variable corresponds with allowing evolutionary pathways through single allele mutations only. If one considers our primary result, Theorem 1 in light of Theorem 3, we see that exponentially many solutions are organized into a stochastically bounded number of components. In the biological model, this means that each viable genotype is connected by single site mutations to an exponential number of other viable genotypes; whereas the number of these clusters, which are not connected to each other is very likely to be small.

We begin our study in Section 3.1 by considering the structure of Q_F for a fixed formula F. We will see that this structure is somewhat difficult to describe when a formula F gives rise to a complicated partial order O_F. In Section 3.2 we consider the random picture and prove Theorem 1. The proof of this result is based on the fact that in a random setting where satisfiability is asymptotically almost sure, the partial order O_F is asymptotically almost surely an antichain.

3.1   The Deterministic Setting

Consider two satisfying assignments $\mathcal{U}$ and $\mathcal{V}$ in Q_F that differ by only one strong component pair $\{C, \overline{C}\}$. We take the perspective that moving from $\mathcal{U}$ to $\mathcal{V}$ is achieved by changing the literals of C in $\mathcal{U}$ to the literals of its complement $\overline{C}$ in $\mathcal{V}$. We will use the notation $\mathcal{V} = \mathcal{U} \circledast C$ to denote this evolution. In other words, $\mathcal{U} \circledast C = (\mathcal{U} - C) \cup \overline{C}$. In general, two satisfying assignments $\mathcal{U}$ and $\mathcal{V}$ may differ by several strong component pairs. We will write $\mathcal{X} = \mathcal{U} - \mathcal{V}$ to denote the collection of strong components in $\mathcal{U}$ whose complements are in $\mathcal{V}$. We will also write $\mathcal{V} = \mathcal{U} \circledast \mathcal{X}$ to indicate that $\mathcal{V}$ is the assignment that results from starting with $\mathcal{U}$ and changing $\mathcal{X}$ to $\overline{\mathcal{X}}$.

Lemma 8 *If $\mathcal{U}$ and $\mathcal{V}$ are any two F satisfying assignments, then there exist satisfying assignments $\mathcal{U}_0, \mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_l$ such that $\mathcal{U}_0 = \mathcal{U}$, $\mathcal{U}_l = \mathcal{V}$ and $\mathcal{U}_{i+1} = \mathcal{U}_i \circledast C_i$ for $0 \leq i < l$ where $C_1, \ldots, C_l$ is some ordering of the strong components of $\mathcal{U} - \mathcal{V}$.*

Proof: We begin by making the following definitions:

$$\mathcal{C} = \mathcal{U} - \mathcal{V}$$
$$P = \text{ the graph of the partial order of } \mathcal{C} \cup \overline{\mathcal{C}}$$
$$\mathcal{B} = \{C \in \mathcal{C} \mid L_P^-(C) = \emptyset\}.$$

Notice that $\mathcal{B}$ is defined so that for any $\mathcal{X} \subseteq \mathcal{B}$, we have $L_P^+(\overline{\mathcal{X}}) = \overline{\mathcal{X}}$. Thus, if $\mathcal{Y} = \mathcal{U} - \mathcal{X}$, then Lemma 5 applies and $\mathcal{U} \circledast \mathcal{X}$ is satisfying for any $\mathcal{X} \subseteq \mathcal{B}$. Thus, given any ordering of the components in $\mathcal{C}$, we can generate a sequence of satisfying assignments from $\mathcal{U}$ to $\mathcal{U} \circledast \mathcal{B}$ following this order. To be concrete, let $C_1, C_2, \ldots, C_l$ be the components in $\mathcal{B}$ in the order induced by the original order of the literals, i.e., C_1 has the least literal among all literals in $\mathcal{B}$ and C_2 has the least literal among all literals in $\mathcal{B} - C_1$ and so on. Then, each of $\mathcal{U}_0, \mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_m$ is a satisfying assignment, where $\mathcal{U}_0 = \mathcal{U}$, $\mathcal{U}_m = \mathcal{U} \circledast \mathcal{B}$ and $\mathcal{U}_{i+1} = \mathcal{U}_i \circledast C_i$ for $0 \leq i < m$.

We can complete our proof via induction on the depth of P. Our base case is that P has depth one. In this case, $\mathcal{B} = \mathcal{C}$ and $\mathcal{U}_m = \mathcal{V}$, so we are done. If the depth of P is k, then the depth of the partial order

on $\mathcal{C} - \mathcal{B}$ is $k-1$, so we can apply the inductive hypothesis to $\mathcal{U}_m$ and $\mathcal{V}$ and concatenate the satisfying assignments found there to the satisfying assignments $\mathcal{U}_0, \mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_m$ and we have our proof. $\qquad \square$

Corollary 9 *Suppose F is a satisfiable 2-formula and that $\mathcal{U}$ and $\mathcal{V}$ are any two distinct satisfying assignments in Q_F. Then $\mathcal{U}$ and $\mathcal{V}$ are connected in Q_F if and only if $\mathcal{U} - \mathcal{V}$ consists entirely of trivial strong components.*

Proof: We use the notation from the proof of Lemma 8. In the case that $\mathcal{U} - \mathcal{V}$ consists entirely of trivial strong components, the sequence $\mathcal{U}_0, \mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_m$ is a path in Q_n. Since Q_F is induced from Q_n, it follows that this path connects $\mathcal{U}$ to $\mathcal{V}$ in Q_F.

For the converse, suppose that there is a nontrivial strong component $C \in \mathcal{U} - \mathcal{V}$. Consider the subcubes of Q^n formed by fixing the literals corresponding to C and $\overline{C}$ respectively, but varying over both literals on all other variables. Since the pair is nontrivial, any assignment in one subcube differs from any assignment in the other subcube for at least two variables. Moreover, there are no satisfying assignments outside these subcubes because every satisfying assignment has to have all the literals of C or else all the literals of $\overline{C}$. Since $\mathcal{U}$ and $\mathcal{V}$ are in different subcubes they are not connected in Q_F. $\qquad \square$

It is important to realize that Lemma 8 and Corollary 9 apply only under the assumption that $\mathcal{U}$ and $\mathcal{V}$ are both satisfying assignments. It is certainly possible that $\mathcal{U}$ is a satisfying assignment, while the assignment $\mathcal{V}$ is unsatisfying even though $\mathcal{U}$ and $\mathcal{V}$ differ by only one strong component pair. A few example formulas and the corresponding graphs are provided in Figure 1. In the figure, we refer to incompatibilities as opposed to disjunctions, so we remind the reader that an incompatibility xy is equivalent to the disjunction $\overline{x} \vee \overline{y}$. Incompatibilities are useful when studying the structure of Q_F because we simply remove the subcubes of Q^n that are specified by the list of incompatibilities.

From Corollary 9 we see that to study the component structure of Q_F one must understand the partial order on the nontrivial strong components. In particular, if there are no nontrivial strong components in D_F on which two assignments differ, then the components are not in distinct components of Q_F. From Corollary 7, we see that we will have satisfying assignments with every combination of strong components from a maximum complementary antichain and its complement. But this tells us nothing about the component structure of Q_F. If we want to count the number of components in Q_F, then we need to identify satisfying assignments in distinct components of Q_F. To accomplish this, we will look at complementary antichains of nontrivial strong components. We will use the term *nontrivial complementary width* to refer to the greatest cardinality among the complementary antichains consisting entirely of nontrivial strong components of D_F. We will use the term *nontrivial height* to refer to the greatest number of nontrivial strong components within a single chain of D_F. The term *nontrivial complementary height* will mean the greatest number of nontrivial strong components in a maximal chain that is complementary. Notice that we first make a list of all maximal chains and then consider only those that are complementary to find one with a maximum number of nontrivial strong components.

Theorem 10 *Suppose that F is a satisfiable formula with nontrivial width j. Also, let k be the number of complementary pairs of nontrivial strong components determined by F and let N be the number of connected components in Q_F. Then $2^j \leq N \leq 2^k$. Furthermore*

1. *If the nontrivial height of O_F is one, then $N = 2^k$.*

2. *If the nontrivial height of O_F is greater than one, then $2^j \leq N < 2^k$.*

Fig. 1: A sequence of digraphs and the corresponding subgraphs of Q^n for $n = 4$. Dimensions of Q^n are ordered left/right, down/up, front/back, in/out. The subcubes of nonsatisfying assignments are shaded. For $F_1 = \{0_1 1_2, 0_2 1_3\}$, D_{F_1} has no cycles and Q_{F_1} is connected. A complementary pair $\{\{0_1, 0_2, 0_3\}, \{1_1, 1_2, 1_3\}\}$ that is not order related arises in $F_2 = F_1 \cup 1_1 0_2$ and Q_{F_2} is disconnected. In $F_3 = F_2 \cup 0_1 0_2$, the complementary pair from F_2 becomes order related and we see that one of the two components from Q_{F_2} is lost in Q_{F_3}.

3. *If there are complementary weak components containing chains of nontrivial height greater than one, then $2^j < N < 2^k$.*

Proof: The weak upper bound follows directly from Corollary 9. The weak lower bound follows directly from Corollaries 7 and 9. Result 1 follows from the fact that $j = k$ if there are no order relations between nontrivial strong components. On the other hand, if there are nontrivial strong components C_1 and C_2 such that $C_1 < C_2$, then $j < k$. For the remainder of the proof, we suppose that this is the case.

Since $C_1 < C_2$ there is no satisfying assignment with C_1 and $\overline{C_2}$, so $N < 2^k$. Thus, the strict upper bound in 2 and 3 holds. Notice that this allows for $C_2 = \overline{C_1}$; in which case there is no satisfying assignment with C_1. All that remains is to show that the strict lower bound holds in the case that C_1 and C_2 are strictly distinct and C_1 and C_2 are not related to $\overline{C_1}$ or $\overline{C_2}$.

We consider Lemma 5 and use the notation there for $\mathcal{X}$. We notice that $L^+(C_1)$ is strictly distinct, so if we set $\mathcal{X} = L^+(C_1)$, then we see that there will be a satisfying assignment with C_1 and C_2. Similarly, if we set $\mathcal{X} = L^+(C_2) \cup L^+(\overline{C_1})$ we find another satisfying assignment. The same can also be said for $\mathcal{X} = L^+(\overline{C_2})$. Thus, we have satisfying assignments with each of $\overline{C_1 C_2}$, $\overline{C_1} C_2$, and $C_1 C_2$, but not with $C_1 \overline{C_2}$. We clearly have that these three possible combinations fall strictly between two and four. We also notice that any satisfying assignment must have one of these three combinations. It follows that $2^j < N < 2^k$. □

3.2   The Random Setting

In this subsection, we prove the primary result of this paper, Theorem 1. Throughout, F and consequently D_F, are randomly determined as described in Section 1. When a strong component is comprised of exactly $i \geq 2$ edges and i literals, we will refer to the strong component as a *simple cycle of order i*, or more succinctly, as an *i-cycle*. The proof of Theorem 1 follows easily from Theorem 10 and the following lemmas.

Lemma 11 *With probability $1 - O(n^{-1})$, every nontrivial strong components is a simple cycle.*

Lemma 12 *With probability $1 - O(n^{-1})$, the set of all nontrivial strong components is an antichain.*

Lemma 13 *The distribution on complementary simple cycle pairs converges to a Poisson distribution with parameter $\lambda = -(\ln(1 - c) + c)/2$.*

Proof: (of Theorem 1 under the assumption of Lemmas 11, 12 and 13.) If Lemmas 11 and 12 are true, then it follows from Theorem 10 that we are a.a.s. that the number of components in Q_F is 2^{X_n} where X_n is the random number of complementary simple cycle pairs. If Lemma 13 is true then X_n converges weakly to a Poisson random variable Ψ with parameter $\lambda = -(\ln(1 - c) + c)/2$. Together, these imply that the number of connected components in Q_F converges weakly to 2^Ψ, and that the probability that there is a unique cluster converges to $e^{-\lambda} = \sqrt{(1 - c)e^c}$. □

All that remains is to show that Lemmas 11, 12 and 13 are true. Before beginning these proofs, we will attempt to provide some context for the results in this section. From a combinatorial perspective (or even simply from a descriptive perspective), Theorem 10 is quite unsatisfactory. We would like to know what happens when the nontrivial height of P is not one. A cursory examination of various partial orders on nontrivial strong components of D_F leads to the conclusion that categorizing all the possible scenarios

would be quite involved and is certainly beyond our scope here. The good news is that Lemmas 11 and 12 show that from a probabilistic perspective, none of these complexities are important because, as the number of variables becomes large, we are asymptotically almost sure that the nontrivial height will be one.

This should come as no surprise to someone well versed in random graph theory because of the similarity between D_F and a standard directed random graph on $2n$ vertices where each edge appears with probability p. One can also relate D_F to a standard undirected random graph $G_{p'}$ where $p' = 2p - p^2$ as was done in [7]. It is well known [6, 4, 15] that for $p = O(n^{-1})$, the number of simple cycle pairs of any fixed length in directed or undirected random graphs converges weakly to a Poisson distribution. What we must show is that the total number of simple cycle pairs of all lengths converges weakly to the Poisson distribution with parameter $\lambda = -(\ln(1 - c) + c)/2$ when $p = c/2n$ and $c < 1$. While this result doesn't seem to be published for the 2-SAT digraph, it has been shown [16, 2] that in the standard random graph, G_{p*} where $p* = c/n$, the total number of simple cycles converges weakly to a Poisson distribution with parameter $-\frac{1}{2}(\ln(1-c)+c+c^2/2)$. The proof we provide for D_F closely resembles the proof in Section 4.3 of [2].

We now establish our notation and make some preliminary observations. Let $\Gamma_{i,n}$ be the set of (unordered) complementary i-cycle pairs on the $2n$ literals and define $\Gamma_n = \bigcup_{i=2} \Gamma_{i,n}$. Then, for each $\alpha \in \Gamma_n$ let

$$I_\alpha = \begin{cases} 1 & \text{if } \alpha \in D_F \\ 0 & \text{else} \end{cases} \qquad \text{and} \qquad p_\alpha = E(I_\alpha).$$

We remind that reader that for any literals $y_1, y_2, \ldots, y_i$ we can make the following statement about cycles in D_F:

$$y_1 \to y_2 \to \ldots \to y_i \to y_1 \iff \overline{y}_1 \leftarrow \overline{y}_2 \leftarrow \ldots \leftarrow \overline{y}_i \leftarrow \overline{y}_1.$$

Since every $\alpha \in \Gamma_{i,n}$ is a complementary pair of directed cycles on strictly distinct literals, $p_\alpha = p^i$ for all $\alpha \in \Gamma_{i,n}$. We next define

$$X_{i,n} = \sum_{\alpha \in \Gamma_{i,n}} I_\alpha \qquad \text{and} \qquad \mu_{i,n} = E(X_{i,n}).$$

Again, by the complementary nature of D_F, we get

$$\mu_{i,n} = \sum_{\alpha \in \Gamma_{i,n}} p_\alpha = \sum_{\alpha \in \Gamma_{i,n}} p^i = \frac{(n)_i}{i} 2^{i-1} p^i = \frac{(n)_i}{n^i} \frac{c^i}{2i}.$$

Notice that for any fixed i, $(n)_i/n^i = (1 - 1/n)(1 - 2/n) \ldots (1 - (i-1)/n)$. Thus, $\mu_{i,n}$ monotonically increases with n and converges to $c^i/2i$. Next, we define

$$X_n = \sum_{i=2}^{n} X_{i,n} \qquad \text{and} \qquad \lambda_n = E(X_n).$$

Notice that X_n is the random variable that counts the total number of simple cycle pairs in D_F. Thus, λ_n is the expected total number of simple cycle pairs in D_F. If $c < 1$, then

$$\lambda_n = \sum_{i=2}^{n} \mu_{i,n} < \sum_{i=2}^{\infty} \frac{c^i}{2i} = -\frac{1}{2}(\ln(1 - c) + c).$$

This immediately implies that if $c < 1$ and $\omega = \omega(n) \to \infty$, then

$$P\left(\sum_{i=\omega}^{n} X_{i,n} > 0\right) \leq E\left(\sum_{i=\omega}^{n} X_{i,n}\right) \xrightarrow{n} 0. \tag{2}$$

Thus, the probability that the length of the longest cycle will diverge with n converges to 0. Nevertheless, our proof technique will require that we control cycles of unbounded length. We discuss the reason for this after completing this proof. For now, we define for any $2 \leq m \leq n$

$$T_{m,n} = \sum_{i=2}^{m} X_{i,n} \qquad \text{and} \qquad \lambda_{m,n} = E(T_{m,n}).$$

We also introduce the notation $\Psi_{m,n}$ for the Poisson random variable with mean $\lambda_{m,n}$. Now that we have established some notation we proceed with the proofs.

Proof: (of Lemma 11.) We observe that any strong component of order $i \geq 3$ that is not a simple cycle contains at least a simple cycle α and a path β with first and last literals in α. Suppose $\alpha \in \Gamma_{i,n}$ and β is a path with $j > 0$ edges and both endpoints in α. Then the expected number of such ordered pairs (α, β) in D_F is bounded above by

$$\binom{n}{i} 2^{i-1}(i-1)! \binom{n-i}{j-1} 2^{j-1}(j-1)! i^2 p^{i+j} < i(2n)^{i+j-1} p^{i+j} = \frac{ic^{i+j}}{2n}.$$

Since $c < 1$, the sum over $1 \leq j \leq n-i$ and then over $2 \leq i \leq n$ is $O(n^{-1})$. We have actually counted at least two ordered pairs for each of the potential nonsimple components because exactly one of the paths of α with endpoints in common with β will form a cycle with β, which will be counted as some cycle α' with the other path as β'. Thus, for any D_F, the number of nonsimple strong components is not more than one half of the number of ordered pairs considered here. Thus, the expected number of ordered pairs we have estimated is at least twice the expected number of nonsimple strong components, which is greater than the probability that there exists a nonsimple strong component. $\qquad\square$

Proof: (of Lemma 12.) Notice that if there are strong components A and C such that $A \leq C$, then there are directed cycles $\alpha \subseteq A$ and $\gamma \subseteq C$ such that α and γ are connected by a directed path β. Similar to the proof of Lemma 12, it is sufficient to show that the expected number of such ordered triples (α, β, γ) is $O(n^{-1})$. To identify (α, β, γ) we select α and mark a literal of α that is common to β; and then select the remaining literals of β in order, the last of which is the first literal of γ; and then select the remaining literals of γ in order. Assuming that α is an i-cycle, β is a path with $j > 0$ edges, and γ is a k-cycle, we have must choose $i+j+k-1$ literals in order. Thus, there are no more than $(2n)_{i+j+k-1} < (2n)^{i+j+k-1}$ such triples. Notice that there are necessarily $i + j + k$ edges in the subgraph corresponding to (α, β, γ). Thus, $c^{i+j+k-1}p$ is an upper bound for the expected number of such configurations. Since $c < 1$ we can sum over i, j, and k to get the desired result. $\qquad\square$

Proof: (of Lemma 13) We will prove the weak convergence of X_n to Ψ by demonstrating that for any subset A of the positive integers we can make $|P(X_n \in A) - P(\Psi \in A)|$ arbitrarily small by taking n to

be large enough. An upper bound for $|P(X_n \in A) - P(\Psi \in A)|$ is provided by

$$|P(X_n \in A) - P(T_{m,n} \in A)| \tag{3}$$

$$+ |P(T_{m,n} \in A) - P(\Psi_{m,n} \in A)| \tag{4}$$

$$+ |P(\Psi_{m,n} \in A) - P(\Psi \in A)|. \tag{5}$$

Thus, to prove the weak convergence of X_n to Ψ it is enough to show that we can choose n and m (which may depend on n) to make (3), (4) and (5) arbitrarily small.

We first consider (5). From (2), we see that when m diverges with n, $\lambda_{m,n}$ will become arbitrarily close to λ. Thus, we have the weak convergence of $\Psi_{m,n}$ to Ψ and (5) must vanish for any m that diverges with n.

We now focus on (3). Notice that (3) is bounded above by $P(X_n \neq T_{m,n})$, which is in turn bounded above by the expected number of events that distinguish X_n from $T_{m,n}$. Notice that by (2) the probability that there are any i-cycles for $i > m$ vanishes as m diverges. Notice also that by Lemma 11 we have for any m, that the probability that there is a strong component that is not a simple cycle is $O(n^{-1})$. Lastly, by Lemma 12 the probability that there are two nontrivial strong components that are ordered is $O(n^{-1})$ for any m. Thus, the probability that $X_n \neq T_{m,n}$ vanishes as n and m diverge.

The supremum over all A of (4) is half the total variation between the law of $T_{m,n}$ and $\Psi_{m,n}$. Thus, the bound for (4) and the proof of Lemma 13 is complete provided that the following lemma is true.

Lemma 14 *There exists an $m = m(n)$ that is unbounded such that the total variation between the law of $T_{m,n}$ and Poisson random variable with mean $\lambda_{m,n}$ is $O(n^{-1})$.*

Proof: There is a family of methods called the Stein-Chen or Chen-Stein method for bounding the total variation between two distributions. A thorough treatment of these methods can be found in [4]. Of particular interest here is Corollary 2.C.5, on pg. 26. of [4], which we restate using our current notation and the following definition. Let $\Gamma_n^m = \bigcup_{i=2}^{m} \Gamma_{i,n}$.

Lemma 15 *Suppose that for each $\alpha \in \Gamma_n^m$, Γ_n^m can be partitioned into Γ_α^0, Γ_α^i such that I_α and I_β are independent for all $\beta \in \Gamma_\alpha^i$. Then the total variation between the laws of $T_{m,n}$ and $\Psi_{m,n}$ is not more than*

$$\min\{1, \lambda_{m,n}^{-1}\} \left(\sum_{\alpha \in \Gamma_n^m} p_\alpha^2 + \sum_{\alpha \in \Gamma_n^m, \beta \in \Gamma_\alpha^0} (p_\alpha p_\beta + E(I_\alpha I_\beta)) \right).$$

For any α, there certainly is a set Γ_α^0 of all complementary cycle pairs that are distinct from α and yet not independent of α. Moreover, since $\lambda_{m,n}$ is bounded below, our proof will be complete when we can find an unbounded $m = m(n)$ for which the sums are both $O(n^{-1})$.

Consider the first sum

$$\sum_{\alpha \in \Gamma_n^m} p_\alpha^2 = \sum_{i=2}^{m} \binom{n}{i} 2^{i-1} (i-1)! p^{2i} < \sum_{i=2}^{\infty} \frac{1}{n^i}.$$

Since this sum is $O(n^{-2})$ for any $m \leq n$, we can address the second sum. If I_α and I_β are not independent, then they share some edge and the event that $I_\alpha = 1$ increases the probability that $I_\beta = 1$, so $E(I_\alpha I_\beta) > p_\alpha p_\beta$ and we need only show convergence of

$$\sum_{\alpha \in \Gamma_n^m, \beta \in \Gamma_\alpha^0} E(I_\alpha I_\beta). \tag{6}$$

Let $\alpha \in \Gamma_{i,n}$ and let β be such that the edges of β that are shared with those of α are found on j disjoint paths in each member of the pair and such that there are k literals on each member of β which are not on the paths common to both α and β. Then, the probability that any such fixed pair (α, β) is present in D_F is p^{i+j+k}. Notice that there are j more edges than literals in such an (α, β) configuration. This is what causes the expectation in question to converge to 0. Now for the details.

We will count the number of such pairs by first fixing α, and then examining the number of possible shared paths. Notice that by choosing $2j$ literals in one member of α we fix the endpoints for one of two possible collections of j shared paths. Thus, the number of ways to choose j paths from an i-cycle is $2\binom{i}{2j}$ where $1 \leq j \leq \lfloor i/2 \rfloor$. This fixes j pairs of complementary paths from the members of α that will be shared with β, but it does not determine which of the complementary paths will be chosen for a given member of β. There are 2^j binary sequences that are in one-to-one correspondence with the possible assignments of j paths to the first member of β. Thus there are $2^{j+1}\binom{i}{2j}$ ways to select the j paths from α that will be shared with one of the members of β.

Next we consider the k additional literals in each member of β. There are no more than $n - 2j$ variables that do not correspond to the shared paths. Moreover, β is a complementary pair of simple cycles, which must have strictly distinct literals in each member of the pair. Thus, there are at most $\binom{n-2j}{k}2^k$ possible ways to choose the remaining literals for one of the members of β.

Together, the literals and paths make for $j + k$ objects to be ordered on a cycle, and there are $(j+k-1)!$ ways to do this. Notice that we did not make use of the restriction that two paths are not allowed to be consecutive in this ordering if there is only one edge between them on α. Thus we are overestimating once again. The choices above, along with the $(n)_i 2^{i-1}/i$ choices for α provide an upper bound for the total number of ordered pairs (α, β). There is just one last bit of counting to mention. We have double counted everything. That is, when we started selecting the k additional literals for one of the members of β, it matters which of the members these k literals go with; but we could get the same structure by selecting the complementary k literals and the complementary member of α. Thus, all together we have $i + j + k - 1$ total factors of 2, which we will group together with the powers of p.

The computations on the following page are mostly self explanatory, but we mention that to obtain (12) and (13) we begin with the Robbins-Feller Stirling approximation and then simplify a bit to get

$$j! < \sqrt{2\pi j}\left(\frac{j}{e}\right)^j e^{1/(12j)} < 4\sqrt{j}\left(\frac{j}{e}\right)^j.$$

Then we use $j \leq i/2 \leq m$ to get

$$\frac{j!}{(2j+2)(2j+1)}\left(\frac{c}{n(1-c)}\right)^j < \frac{4\sqrt{j}}{(2j+2)(2j+1)}\left(\sqrt{\frac{cm}{2e(1-c)n}}\right)^{2j} < \left(\sqrt{\frac{cm}{2e(1-c)n}}\right)^{2j}.$$

$$\sum_{\alpha \in \Gamma_n, \beta \in \Gamma_\alpha} E(I_\alpha I_\beta) < \sum_{i=2}^{m} \sum_{j=1}^{\lfloor i/2 \rfloor} \sum_{k=0}^{m} \frac{(n)_i}{2i} \binom{i}{2j} \binom{n-2j}{k} (k+j-1)!(2p)^{i+j+k} \tag{7}$$

$$< \sum_{i=2}^{m} \frac{c^i}{2i} \sum_{j=1}^{\lfloor i/2 \rfloor} \binom{i}{2j} \frac{c^j}{n^j} (j-1)! \sum_{k=0}^{\infty} \binom{k+j-1}{k} c^k \tag{8}$$

$$= \sum_{i=2}^{m} \frac{c^i}{2i} \sum_{j=1}^{\lfloor i/2 \rfloor} \binom{i}{2j} \frac{c^j}{n^j} (j-1)! \left(\frac{1}{1-c} \right)^j \tag{9}$$

$$= \frac{c}{2n(1-c)} \sum_{i=2}^{m} \frac{c^i}{i} \sum_{j=1}^{\lfloor i/2 \rfloor} \frac{i(i-1)}{2j(2j-1)} \binom{i-2}{2(j-1)} (j-1)! \left(\frac{c}{n(1-c)} \right)^{j-1} \tag{10}$$

$$= \frac{c}{2n(1-c)} \sum_{i=2}^{m} (i-1)c^i \sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i-2}{2j} \frac{j!}{(2j+2)(2j+1)} \left(\frac{c}{n(1-c)} \right)^j \tag{11}$$

$$< \frac{c}{2n(1-c)} \sum_{i=2}^{m} (i-1)c^i \sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i-2}{2j} \left(\frac{j}{e} \right)^j \left(\frac{c}{n(1-c)} \right)^j \tag{12}$$

$$< \frac{c}{2n(1-c)} \sum_{i=2}^{m} (i-1)c^i \sum_{j=0}^{\lfloor i/2 \rfloor - 1} \binom{i-2}{2j} \left(\sqrt{\frac{cm}{2e(1-c)n}} \right)^{2j} \tag{13}$$

$$< \frac{c}{2n(1-c)} \sum_{i=2}^{m} (i-1)c^i \sum_{j=0}^{i-2} \binom{i-2}{j} \left(\sqrt{\frac{cm}{2e(1-c)n}} \right)^j \tag{14}$$

$$= \frac{1}{n} \cdot \frac{c}{2(1-c)} \sum_{i=2}^{m} (i-1)c^i \left(1 + \sqrt{\frac{cm}{2e(1-c)n}} \right)^{i-2} \tag{15}$$

If we look at expression (15) we see that given any positive constant $c < 1$ we can choose $m = m(n)$ so that for large enough n, we have m/n sufficiently small to satifty

$$c \left(1 + \sqrt{\frac{cm}{2e(1-c)n}} \right) < 1.$$

Once this inequality is satisfied, the sum over i in the last line above will be bounded above by a constant, and it follows that the expectation for this $m = m(n)$ will be $O(n^{-1})$, which proves Lemma 14 .   □

As was previously stated, Lemma 13 follows from Lemma 14 and this is all that remained to prove Theorem 1.   □

Corollary 16 *If* $0 < c < c_{max}$, *where* $c_{max} \approx 0.6375$ *is the solution to*

$$c \left(1 + \sqrt{\frac{c}{2e(1-c)}} \right) = 1,$$

then the rate of convergence of X_n to $\Psi_{n,n}$ is $O(n^{-1})$.

Proof: If $c < c_{\max}$, then the proof of Lemma 13 applies with $m = n$ and we have a rate of convergence of $O(n^{-1})$ for X_n to $\Psi_{n,n}$. $\qquad\square$

4 Discussion

The proof provided here of Lemma 13 relies on the use of truncated random variables and also fails to provide a bound on the distance between X_n and $\Psi_{n,n}$ for a significant range of c values. If the proof of Lemma 14 could be improved to provide some rate of convergence $f(n)$ that holds for $m = n$ and all $c < 1$, then there would be no need for truncated random variables and the overall rate of convergence would be no slower than $f(n)$. However, it may be difficult to find such a proof because there is reason to believe that the Chen-Stein method can not be applied to this problem. In [2], they provide numerical evidence (pg. 413) that $\sum_{\alpha \in \Gamma_n, \beta \in \Gamma_\alpha} E(I_\alpha I_\beta)$ diverges with $m = n$ for sufficiently large values of c. If this is true, then for these values of c, the second moment of X_n diverges with n and the Chen-Stein method will not be able to provide a rate of convergence for X_n to $\Psi_{n,n}$.

Acknowledgements

References

[1] Dimitris Achlioptas and Yuval Peres. The threshold for random k-SAT is $2^k \log 2 - O(k)$. *J. Amer. Math. Soc.*, 17(4):947–973 (electronic), 2004.

[2] R. Arratia, L. Goldstein, and L. Gordon. Poisson approximation and the chen-stein method. *Statistical Science*, 5(4):403–424, 1990.

[3] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified Boolean formulas. *Inform. Process. Lett.*, 8(3):121–123, 1979.

[4] A. D. Barbour, L. Holst, and S. Janson. *Poisson Approximation*. Oxford University Press, 1992.

[5] G. Biroli, R. Monasson, and M. Weigt. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J.*, B 14:551–568, 2000.

[6] B. Bollobás. *Random graphs*, volume 73. Cambridge Univ Pr, 2001.

[7] Béla Bollobás, Christian Borgs, Jennifer T. Chayes, Jeong Han Kim, and David B. Wilson. The scaling window of the 2-SAT transition. *Random Structures Algorithms*, 18(3):201–256, 2001.

[8] V. Chvátal and B. Reed. Mick gets some (the odds are on his side). *Proc. 33rd Symp Foundations of Computer Science*, pages 620–627, 1992.

 [9] Stephen A. Cook. The complexity of theorem-proving procedures. In *STOC '71: Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158, New York, NY, USA, 1971. ACM.

[10] W. Fernandez de la Vega. On random 2-SAT. unpublished manuscript, 1992.

[11] S. Gavrilets. *Fitness Landscapes and the Origin of Species*. Monographs in Population Biology. Princeton University Press, 2004.

[12] Andreas Goerdt. A threshold for unsatisfiability. In *Mathematical foundations of computer science 1992 (Prague, 1992)*, volume 629 of *Lecture Notes in Comput. Sci.*, pages 264–274. Springer, Berlin, 1992.

[13] Janko Gravner, Damien Pitman, and Sergey Gavrilets. Percolation on fitness landscapes: Effects of correlation, phenotype, and incompatibilities. *Journal of Theoretical Biology*, 248(4):627 – 645, 2007.

[14] S. A. Kauffman. *The origins of order*. Oxford University Press, Oxford, 1993.

[15] I. Palasti. On the threshold distribution function of cycles in a directed random graph. *Studia Scientiarum Mathematicarum Hungarica*, 6:67–73, 1971.

[16] L. Takács. On the limit distribution of the number of cycles in a random graph. *Journal of Applied Probability*, pages 359–376, 1988.

[17] S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones, editor, *Proceedings of the Sixth International Congress on Genetics*, volume 1, pages 356–366, Austin, Texas, 1932.