Approximating the Unsatisfiability Threshold of Random Formulas (Extended abstract)

Lefteris M. Kirousis[1] Evangelos Kranakis[2] Danny Krizanc[2]

[1] University of Patras, Department of Computer Engineering and Informatics, Rio, 265 00 Patras, Greece
[2] Carleton University, School of Computer Science, Ottawa, ON K1S 5B6, Canada.

Abstract. Let ϕ be a random Boolean formula that is an instance of 3-SAT. We consider the problem of computing the least real number κ such that if the ratio of the number of clauses over the number of variables of ϕ strictly exceeds κ, then ϕ is almost certainly unsatisfiable. By a well known and more or less straightforward argument, it can be shown that $\kappa \leq 5.191$. This upper bound was improved by Kamath, Motwani, Palem, and Spirakis to 4.758, by first providing new improved bounds for the occupancy problem. There is strong experimental evidence that the value of κ is around 4.2. In this work, we define, in terms of the random formula ϕ, a decreasing sequence of random variables such that if the expected value of any one of them converges to zero, then ϕ is almost certainly unsatisfiable. By letting the expected value of the first term of the sequence converge to zero, we obtain, by simple and elementary computations, an upper bound for κ equal to 4.667. From the expected value of the second term of the sequence, we get the value 4.598. In general, by letting the expected value of further terms of this sequence converge to zero, one can, if the calculations are performed, obtain even better approximations to κ. This technique generalizes in a straightforward manner to k-SAT, for $k > 3$.

1 Introduction

Let ϕ be a random 3-SAT formula on n Boolean variables $x_1, \ldots, x_n$. Let m be the number of clauses of ϕ. The clauses-to-variables ratio of ϕ is defined to be the number m/n. We denote this ratio by r. The problem we consider in this paper is to compute the least real number κ such that if r strictly exceeds κ, then the probability of ϕ being satisfiable converges to 0 as n approaches infinity. We say in this case that ϕ is asymptotically almost certainly unsatisfiable. Experimental evidence suggests that the value of κ is around 4.2. Moreover, experiments suggest that if r is strictly smaller than κ, then ϕ is asymptotically almost certainly satisfiable. Thus, experimentally, κ is not only the lower bound for unsatisfiability, but it is a threshold value where, "suddenly", probabilistically certain unsatisfiability yields to probabilistically certain satisfiability (for a review of the experimental results see [11]).

In the literature for this problem, the most common model for random 3-SAT formulas is the following: from the space of clauses with *exactly three* literals of three *distinct* variables from $x_1, \ldots, x_n$, uniformly, independently, and with replacement select m clauses that form the set of conjuncts of ϕ (thus a clause may be selected more than once). We adopt this model in this paper, however, the results can be generalized to any of the usual models for random formulas. The total number N of all possible clauses is $8\binom{n}{3}$, and given a truth assignment A, the probability that a random clause is satisfied by A is $7/8$. Also, given three distinct variables x_i, x_j, x_k, there is a unique clause on the variables x_i, x_j, x_k which is *not* satisfied by A. There are $\binom{n}{3}$ such clauses, and they constitute exactly the set of clauses not satisfied by A.

A proposition stating that if r exceeds a certain constant, then ϕ is asymptotically almost certainly unsatisfiable has as immediate corollary that this constant is an upper bound for κ. We use this observation in our technique to improve the upper bound for κ.

A well known "first moment" argument shows that

$$\kappa \leq \log_{8/7} 2 = 5.191.$$

To prove it, observe that the expected value of the number of truth assignments that satisfy ϕ is $2^n (7/8)^{rn}$, then let this expected value converge to zero and use Markov's inequality (this argument is expanded below). According to Chvátal and Reed [4], this observation is due to Franco and Paull [6], Simon et al. [14], Chvátal and Szemerédi [5], and possibly others.

Let $\mathcal{A}_n$ be the set of all truth assignments on the n variables $x_1, \ldots x_n$, and let $\mathcal{S}_n$ be the set of truth assignments that satisfy the random formula ϕ. The cardinality $|\mathcal{S}_n|$ is thus a random variable. Also, for an instantiation ϕ of the random formula, let $|\mathcal{S}_n(\phi)|$ denote the number of truth assignments that satisfy ϕ. (A word of caution: in order to avoid overloading the notation, we use the same symbol ϕ to denote the random formula and an instantiation of it.) We give below a rough outline of the simplest case of our technique.

By definition, the expected value of the number of satisfying truth assignments of a random formula, i.e., $\mathbf{E}[|\mathcal{S}_n|]$, satisfies the following relation

$$\mathbf{E}[|\mathcal{S}_n|] = \sum_{\phi} \left(\mathbf{Pr}[\phi] \cdot |\mathcal{S}_n(\phi)| \right). \tag{1}$$

On the other hand, the probability of a random formula being satisfiable is given by the equation:

$$\mathbf{Pr}[\text{the random formula is satisfiable}] = \sum_{\phi} \left(\mathbf{Pr}[\phi] \cdot I_\phi \right), \tag{2}$$

where

$$I_\phi = \begin{cases} 1 \text{ if } \phi \text{ is satisfiable,} \\ 0 \text{ otherwise.} \end{cases} \tag{3}$$

From equations (1) and (2) the following Markov's inequality follows immediately:

$$\mathbf{Pr}[\text{the random formula is satisfiable}] \leq \mathbf{E}[|\mathcal{S}_n|]. \tag{4}$$

It is easy to find a condition on κ under which $\mathbf{E}[|\mathcal{S}_n|]$ converges to zero. Such a condition, by Markov's inequality (4), implies that ϕ is asymptotically almost certainly unsatisfiable (this elementary technique is known as the "first moment method"). However, as in the right-hand side of equation (1) we may have small probabilities multiplied with large cardinalities, such a condition may be unnecessarily strong for guaranteeing only that ϕ is almost certainly unsatisfiable. In this work, instead of considering the random class $\mathcal{S}_n$ that may have a large cardinality for certain instantiations of the random formula with small probability, we consider a subset of it obtained by taking truth assignments that satisfy a local maximality condition. Thus, the condition obtained by letting the expected value of this new class converge to zero is weakened, and consequently, the upper bound for κ is lowered.

As we show in the next section, the bound for κ obtained by this sharpened first moment technique is equal to 4.667. This improves the previous best bound due to Kamath, Motwani, Palem, and Spirakis [9] of 4.758, which was obtained by non-elementary means. Moreover our method is not computational, i.e. it does not use any mechanical computations that do not have provable accuracy and correctness (the fact that in our method we use a computer program to find a solution of an equation with *one* unknown does not render our proof computational, because the algorithms that find solutions to such equations have provable accuracy). The bound that Kamath et al. [9] attain with a non-computational proof is equal to 4.87.

In Section 3 we show how to further improve the bound to 4.598 by defining an even smaller subset of $\mathcal{S}_n$. This is achieved by increasing the range of locality when selecting the local maxima that represent $\mathcal{S}_n$. Actually, we define a decreasing sequence of subsets of $\mathcal{S}_n$ by selecting from $\mathcal{S}_n$ truth assignments that satisfy a condition of local maximality with increasing range of locality. From this sequence, if we perform the calculations, we can obtain a sequence of improving approximations to κ. In the last section, we discuss the case of letting this range of locality become unboundedly large.

Moreover, our bounds can be possibly improved even further if one uses not the Markov type inequality mentioned above, but an analog of the "harmonic mean formula" given by Aldous [3], and then apply the technique that is used in Kamath et al. [9]. This is discussed in the last section. Finally, our method readily generalizes to k-SAT, for $k > 3$.

2 Single Flips

Recall, $\mathcal{A}_n$ is the class of all truth assignments, and $\mathcal{S}_n$ is the random class of truth assignments that satisfy a random formula ϕ. We now define a class even smaller than $\mathcal{S}_n$.

Definition 1. For a random formula ϕ, $\mathcal{S}_n^\sharp$ is defined to be the random class of truth assignments A such that (i) $A \models \phi$, and (ii) any assignment obtained from A by changing exactly one FALSE value of A to TRUE does not satisfy ϕ.

Notice that the truth assignment with all its values equal to TRUE vacuously satisfies condition (ii) of the previous definition. Consider the lexicographic ordering among truth assignments, where, as usual, the value FALSE is considered smaller than TRUE and the values of variables with higher index are of lower priority in establishing the way two assignments compare. It is not hard to see that $\mathcal{S}_n^\sharp$ is the set of elements of $\mathcal{S}_n$ that are local maxima in the lexicographic ordering of assignments, where the neighborhood of determination of local maximality is the set of assignments that differ from A in at most one position.

We now prove:

Lemma 2. *The following Markov type inequality holds for* $\mathcal{S}_n^\sharp$*:*

$$\mathbf{Pr}[\text{the random formula is satisfiable}] \leq \mathbf{E}[|\mathcal{S}_n^\sharp|]. \tag{5}$$

PROOF From the previous definition we easily infer that if an instantiation ϕ of the random formula is satisfiable, then $\mathcal{S}_n^\sharp(\phi) \neq \emptyset$. (Recall that $\mathcal{S}_n^\sharp(\phi)$ is the instantiation of the random class $\mathcal{S}_n^\sharp$ at the instantiation ϕ.) We also have that

$$\mathbf{Pr}[\text{the random formula is satisfiable}] = \sum_{\phi} \left(\mathbf{Pr}[\phi] \cdot I_\phi\right),$$

where

$$I_\phi = \begin{cases} 1 \text{ if } \phi \text{ is satisfiable,} \\ 0 \text{ otherwise.} \end{cases} \tag{6}$$

On the other hand,

$$\mathbf{E}[|\mathcal{S}_n^\sharp|] = \sum_{\phi} \left(\mathbf{Pr}[\phi] \cdot |\mathcal{S}_n^\sharp(\phi)|\right).$$

The lemma now immediately follows from the above. □

We also have the following:

Lemma 3. *The expected value of the random variable* $|\mathcal{S}_n^\sharp|$ *is given by the formula*

$$\mathbf{E}[|\mathcal{S}_n^\sharp|] = (7/8)^{rn} \sum_{A \in \mathcal{A}_n} \mathbf{Pr}[A \in \mathcal{S}_n^\sharp \mid A \in \mathcal{S}_n]. \tag{7}$$

PROOF First observe that the random variable $|\mathcal{S}_n^\sharp|$ is the sum of indicator variables and then condition on $A \models \phi$ (recall, r is the number of clauses-to-number-of-variables ratio of ϕ, so $m = nr$). □

We call a change of *exactly one* FALSE value of a truth assignment A to TRUE a *single flip*. The number of possible single flips, which is of course equal to the number of FALSE values of A, is denoted by $sf(A)$. The assignment obtained by applying a single flip sf on A is denoted by A^{sf}.

We now prove that

Theorem 4. *The expected value* $\mathbf{E}[|\mathcal{S}_n^\sharp|]$ *is at most* $(7/8)^{rn}(2 - e^{-3r/7} + o(1))^n$. *It follows that the unique positive solution of the equation*

$$(7/8)^r(2 - e^{-3r/7}) = 1,$$

is an upper bound for κ *(this solution is less than 4.667).*

PROOF Fix a single flip sf_0 on A and assume that $A \models \phi$. Observe that the assumption that $A \models \phi$ excludes $\binom{n}{3}$ clauses from the conjuncts of ϕ, i.e., there remain $7\binom{n}{3}$ clauses to choose the conjuncts of ϕ from. Consider now the clauses that are not satisfied by A^{sf_0} and contain the flipped variable. There are $\binom{n-1}{2}$ of them. Under the assumption that $A \models \phi$, in order to have that $A^{sf_0} \not\models \phi$, it is necessary and sufficient that at least one of these $\binom{n-1}{2}$ clauses be a conjunct of ϕ. Therefore, for each of the m clause selections for ϕ, the probability of being one that guarantees that $A^{sf_0} \not\models \phi$ is $\binom{n-1}{2}/7\binom{n}{3} = 3/(7n)$. Therefore, the probability that $A^{sf_0} \not\models \phi$ (given that $A \models \phi$) is equal to $1-(1-3/(7n))^m$. Now, there are $sf(A)$ possible flips for A. The events that ϕ is not satisfied by the assignment A^{sf} for *each* single flip sf (under the assumption that $A \models \phi$) refer to disjoint sets of clauses. Therefore, it can be easily seen [2] that the dependencies among them are such that:

$$\mathbf{Pr}[A \in \mathcal{S}_n^\sharp \mid A \models \phi] \leq \left(1 - \left(1 - \frac{3}{7n}\right)^m\right)^{sf(A)} = \left(1 - e^{-3r/7} + o(1)\right)^{sf(A)}. \quad (8)$$

Remember, $sf(A)$ is equal to the number of FALSE values of A. Therefore, by equation (7) and by Newton's binomial formula, $\mathbf{E}[|\mathcal{S}_n^\sharp|]$ is bounded above by $(7/8)^{rn}(2 - (1 - 3/(7n))^{rn})^n$, which proves the first statement of the theorem.

It also follows that $\mathbf{E}[|\mathcal{S}_n^\sharp]$ converges to zero for values of r that strictly exceed the unique positive solution of the equation $(7/8)^r(2-e^{-3r/7}) = 1$. By Lemma 2, this solution is an upper bound for κ. As it can be seen by any program that computes roots of equations with accuracy of at least four decimal digits (we used Maple [13]), this solution is less than 4.667. □

The generalization of the previous result to the case of k-SAT, for an arbitrary $k \geq 3$ is immediate:

Theorem 5. *For the case of k-SAT* ($k \geq 3$)*, the expected value* $\mathbf{E}[|\mathcal{S}_n^\sharp]$ *is at most* $((2^k - 1)/2^k)^{rn}(2 - e^{-kr/(2^k-1)} + o(1))^n$. *It follows that the unique positive solution of the equation*

$$\left(\frac{2^k - 1}{2^k}\right)^r (2 - e^{-kr/(2^k-1)}) = 1,$$

is an upper bound for κ *(as defined for k-SAT).*

3 The General Method and Double Flips

In this section, we generalize the previous method to an arbitrary range of locality when selecting the subset of $\mathcal{S}_n$. We start with a definition:

Definition 6. Given a random formula ϕ and a nonnegative integer l, $\mathcal{A}_n^l\,(l \leq n)$ is defined to be the random class of truth assignments A such that (i) $A \models \phi$, and (ii) any assignment that differs from A in at most l variables and is lexicographically strictly larger than A does not satisfy ϕ.

Observe that $\mathcal{S}_n$ of the previous section, i.e., the class of truth assignments satisfying the random formula is, in the notation of the previous definition, equal to $\mathcal{A}_n^0$, whereas $\mathcal{S}_n^\sharp$ is equal to $\mathcal{A}_n^1$. In general, $\mathcal{A}_n^l$ is the subset of $\mathcal{S}_n$ that consists of the lexicographic local maxima of it where the neighborhood of locality for an assignment A is the set of assignments that differ from A in at most l places. Moreover, obviously, $\mathcal{A}_n^l$ is a sequence of classes which is decreasing relative to l (with respect to set inclusion).

Now, exactly as in Lemma 2, it can be proved that:

Lemma 7. *The following Markov type inequalities hold for the classes $\mathcal{A}_n^l$:*

$$\mathbf{Pr}[\phi \text{ is satisfiable}] = \mathbf{E}[|\mathcal{A}_n^n|] \leq \mathbf{E}[|\mathcal{A}_n^{n-1}|] \leq \cdots \leq \mathbf{E}[|\mathcal{A}_n^1|] \leq \mathbf{E}[|\mathcal{A}_n^0|]. \tag{9}$$

It follows from the above that for a fixed l, by letting $\lim_n \mathbf{E}[|\mathcal{A}_n^l|] = 0$, we obtain upper bounds for κ which decrease as l increases. In other words, if r_l denotes the infimum of the values of r that make the expression $\mathbf{E}[|\mathcal{A}_n^l|]$ converge to zero (as $n \to \infty$), then r_l is an upper bound for κ, and the larger l is, the better the bound. We concetrate below on the case $l = 2$.

A change of exactly two values of a truth assignment A that gives a truth assignment which is lexicographically strictly larger than A must be of one of the following kinds: (1) a change of the value FALSE of a variable to TRUE and a change of the value TRUE of a higher indexed variable to FALSE, or (2) a change of two variables both of value FALSE to TRUE. From these two possible kinds of changes, we consider only the first, since the calculations become easier, while the final result remains the same. We call such changes *double flips.* Define A^{df} and $df(A)$ in a way analogous to the single flip case (notice that if A is considered as a sequence of the Boolean values 0 and 1, then $df(A)$ is equal to the number of order inversions as we move along A from high-indexed variables to low-indexed ones, i.e. from right to left). Let $\mathcal{A}_n^{2\sharp}$ be the set of assignments A such that $A \models \phi$ and for all single flips sf, $A^{sf} \not\models \phi$ and for all double flips df, $A^{df} \not\models \phi$. It can be easily seen that $\mathcal{A}_n^2$ is a subset of $\mathcal{A}_n^{2\sharp}$ (in general a proper one, because in the definition of $\mathcal{A}_n^{2\sharp}$ we did not take into account the changes of kind (2)). Therefore a value of r that makes the expected value $\mathbf{E}[|\mathcal{A}_n^{2\sharp}|]$ converge to zero is, by Lemma 7, an upper bound for κ. Actually, it can be proved that both $\mathbf{E}[|\mathcal{A}_n^{2\sharp}|]$ and $\mathbf{E}[|\mathcal{A}_n^2|]$ converge to zero for the same values of r, but we will not use this fact, so we omit its proof.

Now in analogy to Lemma 3 we have

Lemma 8.

$$\mathbf{E}[|\mathcal{A}_n^{2\sharp}|] = (7/8)^{rn} \sum_{A \in \mathcal{A}_n} \mathbf{Pr}[A \in \mathcal{A}_n^1 \mid A \models \phi] \cdot \mathbf{Pr}[A \in \mathcal{A}_n^{2\sharp} \mid A \in \mathcal{A}_n^1]. \quad (10)$$

Therefore, by the remarks in the beginning of the current section, an upper bound for κ can be found by computing a value (the smaller the better) for r for which the right-hand side of the equality above converges to zero. We will do this in two steps. First we will compute an upper bound for the second factor in the terms of the sum in the equality above (the first factor has been computed in the previous section); then we will find an upper bound for $\mathbf{E}[|\mathcal{A}_n^{2\sharp}|]$ which will be a closed expression of r and n. Letting this closed expression converge to zero with n, we will get an equation in terms of r that gives the required bound for κ.

To compute an upper bound for the second factor of the sum, we will make use of the Janson's inequality [8], which gives an estimate for the probability of the intersection of dependent events. We give the details in the first subsection of the present section. The computations that will then give a closed expression that is an upper bound for $\mathbf{E}[|\mathcal{A}_n^{2\sharp}|]$ are carried out in the second subsection.

3.1 Probability Calculations

In this subsection, we compute the probability $\mathbf{Pr}[A \in \mathcal{A}_n^{2\sharp} \mid A \in \mathcal{A}_n^1]$ (this expression appears in the right-hand side of equation (10)). We condition, for the rest of the section, on $A \models \phi$. It is also convenient to introduce the following notation to be used in the sequel: for a variable x_i, x_i^A is the literal x_i if the value of x_i in A is TRUE, and it is the literal $\neg x_i$, otherwise.

First, fix a double flip df_0. Then we have:

Lemma 9. *The following holds:*

$$\mathbf{Pr}[A^{df_0} \not\models \phi \mid A \in \mathcal{A}_n^1] = 1 - \frac{6e^{-6r/7} r}{7(1 - e^{-3r/7})} \frac{1}{n} + o\left(\frac{1}{n}\right). \quad (11)$$

PROOF Assume without loss of generality that df_0 changes the values of x_1 and x_2 and that these values are originally FALSE and TRUE, respectively. Also let sf_0 be the *unique* single flip that changes a value which is also changed by df_0. In this case, sf_0 is the flip that changes the value of x_1 from FALSE to TRUE.

Notice that because all single flips that are distinct from sf_0 change values which are not changed by df_0, the dependencies are such that:

$$\mathbf{Pr}[A^{df_0} \not\models \phi \mid A \in \mathcal{A}_n^1] \leq \mathbf{Pr}[A^{df_0} \not\models \phi \mid A^{sf_0} \not\models \phi].$$

Actually, it can be proved that asymptotically with n, the previous inequality is an equality. To compute the "negated" probability in the right-hand side of the above inequality, we proceed as follows:

In the previous section (proof of Theorem 4), we proved that $\mathbf{Pr}[A^{sf_0} \not\models \phi] = 1 - (1 - 3/(7n))^m$. We now first compute the "positive" (with respect to A^{df_0}) probability:

$$\mathbf{Pr}[A^{df_0} \models \phi \wedge A^{sf_0} \not\models \phi].$$

Observe that in order to have that $A^{df_0} \models \phi$, any clause that contains *at least one* of the literals $\neg x_1, x_2$ and its remaining literals belong to $\neg x_i^A$, $i > 2$, *must not be* among the conjuncts of ϕ. The number of these clauses is equal to $2\binom{n-2}{2} + n - 2 = (n-2)^2$. However the additional requirement that $A^{sf_0} \not\models \phi$, in conjunction with the requirement that $A^{df_0} \models \phi$, makes necessary that at least one clause that contains *both* $\neg x_1, \neg x_2$ and one of $\neg x_i^A$, $i > 2$, *is* among the conjuncts of ϕ (the number of such clauses is $n-2$). The probability for these events to occur simultaneously is equal to

$$\left(1 - \frac{(n-2)^2}{7\binom{n}{3}}\right)^m \cdot \left(1 - \left(1 - \frac{n-2}{7\binom{n}{3} - (n-2)^2}\right)^m\right).$$

This last expression gives the probability $\mathbf{Pr}[A^{df_0} \models \phi \wedge A^{sf_0} \not\models \phi]$.

From the above, it follows that

$$\mathbf{Pr}[A^{df_0} \not\models \phi \mid A^{sf_0} \not\models \phi] = 1 - \frac{6e^{-6r/7} r}{7(1 - e^{-3r/7})} \frac{1}{n} + o\left(\frac{1}{n}\right).$$

That concludes the proof. □

Unfortunately, we cannot just multiply the probabilities in the previous lemma to compute $\mathbf{Pr}[A \in \mathcal{A}_n^{2\sharp} \mid A \in \mathcal{A}_n^1]$, because these probabilities are not independent. This is so because two double flips may have variables in common. Fortunately, we can apply Janson's inequality [8] that gives an estimate for the probability of the intersection of dependent events. For a detailed presentation of this theorem we refer to the 2nd edition of Spencer's book [15]. In our case we will apply a variation tailored to our needs. Below we follow as closely as possible the notation of [15].

Let I be the class of all double flips (i.e. $I = df(A)$), and let $\mathcal{J}_i, i \in I$ be the finite family of the "positive" events $A^{df} \models \phi$, $df \in I$, conditioned on $A \in \mathcal{A}_n^1$. Let $i \sim j$ denote that i and j are double flips that they have a common variable to be flipped and that $i \neq j$. Let $\Delta = \sum_{i \sim j} \mathbf{Pr}[\mathcal{J}_i \wedge \mathcal{J}_j]$, and finally let $\epsilon \geq \mathbf{Pr}[\mathcal{J}_i], \forall i \in I$. Then it follows that:

$$\mathbf{Pr}[\wedge_{i \in I}(\neg \mathcal{J}_i)] \leq \left(\prod_{i \in I} \mathbf{Pr}[\neg \mathcal{J}_i]\right) \cdot e^{\Delta/[2(1-\epsilon)]}. \qquad \text{(Janson's inequality)}$$

To prove the above inequality, it suffices to follow the proof of the first form of Janson's inequality in [15] using, at the very beginning of that proof, the inequality:

$$\mathbf{Pr}[\neg \mathcal{J}_i \mid \wedge_{l=1}^{i-1} \neg(\mathcal{J}_l)] \leq \mathbf{Pr}[\neg \mathcal{J}_i \mid \wedge_{l=1}^{d} \neg(\mathcal{J}_l)], \tag{12}$$

where, under a renumbering, the flips $1, \ldots, d$ ($d \leq i$) are the ones that share with flip i a common variable to be flipped, while the flips $d+1, \ldots, i-1$ are the ones that do not. As it can be easily seen, the last inequality holds for the particular J_is we are considering. Using it in the very beginning of the proof in [15], makes unnecessary the use of the correlation condition assumed there.

We now conclude, making use of the above variant of Janson's inequality, that:

$$\mathbf{Pr}[A \in \mathcal{A}_n^{2\sharp} \mid A \in \mathcal{A}_n^1] \le \left(\mathbf{Pr}[A^{df_0} \not\models \phi \mid A \in \mathcal{A}_n^1]\right)^{df(A)} \cdot e^{\Delta/[2(1-\epsilon)]}, \qquad (13)$$

where Δ and ϵ are defined above. By equation (11), $\epsilon = o(1)$, so we ignore ϵ. The computation of Δ is a bit tedious. In the following lemmata, we give the results of the various steps in this computation, hoping that the interested (and patient) reader can carry them out by herself. The method to be used is very similar to that of the proof Lemma 9. In order to save a little more on notation, we set

$$u = e^{-r/7}.$$

Lemma 10. *Let df_0 and df_1 be two double flips that share the variable that they change from* FALSE *to* TRUE. *Then*

$$\mathbf{Pr}[A^{df_0} \models \phi, A^{df_1} \models \phi \mid A \in \mathcal{A}_n^1] = \frac{6u^9 \ln(1/u)}{1-u^3}\frac{1}{n^2} + o\left(\frac{1}{n^2}\right). \qquad (14)$$

Lemma 11. *Let df_0 and df_1 be two double flips that share the variable that they change from* TRUE *to* FALSE. *Then*

$$\mathbf{Pr}[A^{df_0} \models \phi, A^{df_1} \models \phi \mid A \in \mathcal{A}_n^1] = \frac{36u^9 \ln^2(1/u)}{(1-u^3)^2}\frac{1}{n^2} + o\left(\frac{1}{n^2}\right). \qquad (15)$$

Now observe that the number of pairs of flips described in Lemma 10 is at most $df(A)\cdot(n - sf(A))$, while the number of pairs described in Lemma 11 is at most $df(A)\cdot sf(A)$. Also, it is easy to see that the probability in Lemma 10 is smaller than the probability in Lemma 11. Therefore, we obtain the estimate:

$$\Delta \le df(A)\cdot\left(\frac{36u^9 \ln^2(1/u)}{(1-u^3)^2}\frac{1}{n} + o\left(\frac{1}{n}\right)\right).$$

From this, by inequality (13), it follows that:

$$\mathbf{Pr}[A \in \mathcal{A}_n^{2\sharp} \mid A \in \mathcal{A}_n^1] \le \left(1 - \frac{6u^6 \ln(1/u)}{1-u^3}\frac{1}{n} + \frac{18u^9 \ln^2(1/u)}{(1-u^3)^2}\frac{1}{n} + o\left(\frac{1}{n}\right)\right)^{df(A)}. \qquad (16)$$

It is easy to see that the expression at the base of the right-hand side of the above inequality is at most 1, for $u \in (0,1)$. Now, by equations (8), (10), and (16), we get that:

$$\mathbf{E}[|\mathcal{A}_n^{2\sharp}|] \le (7/8)^{rn} \sum_A X^{sf(A)} Y^{df(A)}, \qquad (17)$$

where

$$X = 1 - u^3 + o(1) \qquad (18)$$

and

$$Y = 1 - \frac{6u^6 \ln(1/u)}{1-u^3}\left(1 - \frac{3u^3 \ln(1/u)}{1-u^3}\right)\frac{1}{n} + o\left(\frac{1}{n}\right). \qquad (19)$$

In the next subsection, we give an estimate for the sum in inequality (17).

3.2 Estimates

Lemma 12. *If* $0 \le X^2 \le Y \le 1$, *then*

$$\sum_A X^{sf(A)} Y^{df(A)} \le \prod_{i=0}^{n-1} (1 + XY^{i/2}). \tag{20}$$

Notice that in our case the condition $X^2 \le Y$ holds, as by equations (18) and (19), we have that $Y = 1 + o(1)$ and $X = 1 - u^3 + o(1)$. The easiest way to prove the inequality in the lemma is first to show that

$$\sum_A X^{sf(A)} Y^{df(A)} = \sum_{k=0}^{n} \binom{n}{k}_Y X^k,$$

where

$$\binom{n}{k}_q = \frac{(1-q^n)(1-q^{n-1})\cdots(1-q^{n-k+1})}{(1-q^k)(1-q^{k-1})\cdots(1-q^1)}$$

are the so called q-nomial or Gauss coefficients (see Knuth's book [12], page 64), and then proceed inductively on n. Complete information on such techniques can be found in a book on basic or q-hypergeometric series by Gasper and Rahman [7]. A direct proof is also possible, but it is rather involved. We do not give the details, as they do not offer anything new to our problem (for a proof see [10]).

Now, recall that:

$$u = e^{-r/7}, \tag{21}$$

$$X = 1 - u^3 + o(1), \tag{22}$$

and

$$Y = 1 - \frac{6u^6 \ln(1/u)}{1-u^3} \left(1 - \frac{3u^3 \ln(1/u)}{1-u^3}\right) \frac{1}{n} + o\left(\frac{1}{n}\right). \tag{23}$$

Set also $Z = n \ln Y$ and observe that from equation (23) it follows that:

$$Z = -\frac{6u^6 \ln(1/u)}{1-u^3} \left(1 - \frac{3u^3 \ln(1/u)}{1-u^3}\right) + o(1). \tag{24}$$

Our estimate for $\mathbf{E}[|\mathcal{A}_n^{2\sharp}|]$ will be given in terms of the dilogarithm function (see the book by Abramowitz and Stegun [1]) which is defined as:

$$\mathrm{dilog}(x) = -\int_1^x \frac{\ln(t)}{t-1} dt.$$

Finally, let $df\_eq(r)$ be the expression that we get if we substitute in

$$\ln(7/8) r(Z/2) + \mathrm{dilog}(1+X) - \mathrm{dilog}(1+Xe^{Z/2})$$

the values of X and Z without their asymptotic terms and then set $u = e^{-r/7}$ (it will shortly become clear why we introduce the above expression of X, Z and r).

We now state the concluding result:

Theorem 13. *If $df\_eq(r) < 0$, then $\lim_n \mathbf{E}[|\mathcal{A}_n^{2\sharp}|] = 0$, and therefore*

$$\lim_n \mathbf{Pr}[\phi \textit{ is satisfiable}] = 0.$$

It follows that $\kappa < 4.598$.

PROOF From inequalities (17) and (20), we conclude that in order to have

$$\lim_{n\to\infty} \mathbf{E}[|\mathcal{A}_n^{2\sharp}|] = 0,$$

it is sufficient to show that the expression

$$(7/8)^{rn}\left(\prod_{i=0}^{n-1}(1+XY^{i/2})\right)$$

converges to zero. Raising this last expression to the power $1/n$, then taking the logarithm, and finally making the standard approximation of a sum by an integral (for the case of a decreasing function), we conclude that a sufficient condition for $\lim_n \mathbf{E}[|\mathcal{A}_n^{2\sharp}|] = 0$, is that:

$$r\ln(7/8) + \lim_n\left((1/n)\int_{-1}^{(n-1)/2} \ln(1+XY^{\tau/2})d\tau\right) = 0.$$

However,

$$\int \ln(1+XY^{\tau/2})d\tau = -\frac{\mathrm{dilog}(1+XY^{\tau/2})}{\ln(Y^{1/2})}.$$

The first assertion of the theorem now follows by elementary calculations taking into account that $Y^{n/2} = e^{Z/2}$ and that $Y = 1 + o(1)$. The second assertion follows by Lemma 7. The estimate for κ is obtained by computing the unique positive solution of the equation $df\_eq(r) = 0$. We obtained the value 4.598 by using Maple [13]. □

Acknowledgments

Research performed while the first author was visiting the School of Computer Science of Carleton University and was partially supported by NSERC (Natural Sciences and Engineering Research Council of Canada), and by a grant from the University of Patrasfor sabbatical leaves. The second and third author were supported in part by grants from NSERC (Natural Sciences and Engineering Research Council of Canada).

We thank Dimitris Achlioptas and Mike Molloy for reading an earlier draft of this paper and noticing an error it contained. The discussions that followed were very helpful to us. We also thank Paul Spirakis. His remarks were crucial in improving our notation. We are grateful to Mizan Rahman for the very informative discussions concerning the sum $\sum_A X^{sf(A)}Y^{df(A)}$ and its relation to basic hypergeometric series.

References

1. M. Abramowitz and I. E. Stegun (eds.), *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables*, 10th printing, U.S. Department of Commerce, National Bureau of Standards, Washington, 1972.
2. D. Achlioptas, L. M. Kirousis, E. Kranakis, D. Krizanc, M. S.O. Molloy, A Correlation Inequality and Its Application to a Word Problem, Technical Report 96-11, School of Computer Science, Carleton University.
3. D.J. Aldous, "The harmonic mean formula for probabilities of unions: applications to sparse random graphs," *Discrete Mathematics* 76, pp 167–176, 1989.
4. V. Chvátal and B. Reed, "Mick gets some (the odds are on his side)," *Proc. 33rd IEEE Symposium on Foundations of Computer Science*, pp 620–627, 1992.
5. V. Chvátal and E. Szemerédi, "Many hard examples for resolution," *Journal of the Association for Computing Machinery* 35, pp 759–768, 1988.
6. J. Franco and M. Paull, "Probabilistic analysis of the Davis-Putnam procedure for solving the satisfiability problem," *Discrete Applied Mathematics* 5, pp 77–87, 1983.
7. G. Gasper and M. Rahman, *Basic Hypergeometric Series*, Encyclopedia of Mathematics and its Applications, vol 35, Cambridge University Press, Cambridge, 1990.
8. S. Janson, "Poisson approximation for large deviations," *Random Structures and Algorithms* 1, pp 221–230, 1990.
9. A. Kamath, R. Motwani, K. Palem, and P. Spirakis, "Tail bounds for occupancy and the satisfiability threshold conjecture," *Random Structures and Algorithms* 7, pp 59–80, 1995. Also in: *Proc. 35th FOCS*, IEEE, pp 592–603, 1994
10. L. Kirousis, E. Kranakis, and D. Krizanc, *An Upper Bound for a Basic Hypergeometric Series*, Carleton University, School of Computer Science, Technical Report TR-96-07, 1996.
11. S. Kirkpatrick and B. Selman, "Critical behavior in the satisfiability of random Boolean expressions," *Science* 264, pp 1297–1301, 1994.
12. D. Knuth, *Fundamental Algorithms*, The Art of Computer Programming, vol. 1, 2nd edition, Addison-Wesley, Reading, Massachusetts, 1973.
13. D. Redfern, *The Maple Handbook: Maple V Release 3*, Springer-Verlag, New York, 1994.
14. J.-C. Simon, J. Carlier, O. Dubois, and O. Moulines, "Étude statistique de l' existence de solutions de problèmes SAT, application aux systèmes-experts," *C.R. Acad. Sci. Paris. Sér. I Math.* 302, pp 283–286, 1986.
15. J. H. Spencer, *Ten Lectures on the Probabilistic Method*, 2nd edition, SIAM, Philadelphia, 1994