Coupon collectors, q-binomial coefficients and the unsatisfiability threshold

Alexis C. Kaporis[1], Lefteris M. Kirousis[1], Yannis C. Stamatiou[1,2], Malvina Vamvakari[1,2], and Michele Zito[⋆3]

[1] University of Patras, Department of Computer Engineering and Informatics, Rio 265 00, Patras, Greece. e-mail: {kaporis,kirousis,stamatiu,mvamv}@ceid.upatras.gr

[2] Computer Technology Institute, Kolokotroni 3, 26321, Patras, Greece.

[3] Department of Computer Science, University of Liverpool, UK. e-mail: michele@csc.liv.ac.uk.

Abstract. The problem of determining the unsatisfiability threshold for random 3-SAT formulas consists in determining the clause to variable ratio that marks the (experimentally observed) abrupt change from almost surely satisfiable formulas to almost surely unsatisfiable. Up to now, there have been rigorously established increasingly better lower and upper bounds to the actual threshold value. An upper bound of 4.506 was announced by Dubois et al. in 1999 but, to the best of our knowledge, no complete proof has been made available from the authors yet. We consider the problem of bounding the threshold value from above using methods that, we believe, are of interest on their own right. More specifically, we explain how the method of *local maximum satisfying truth assignments* can be combined with results for coupon collector's probabilities in order to achieve an upper bound for the unsatisfiability threshold less than 4.571. Thus, we improve over the best, with an available complete proof, previous upper bound, which was 4.596. In order to obtain this value, we also establish a bound on the q-binomial coefficients (a generalization of the binomial coefficients) which may be of independent interest.

1 Introduction

Let ϕ be a random 3-SAT formula constructed by selecting uniformly and with replacement m clauses from the set of all possible clauses with three literals of three distinct variables. It has been observed experimentally that as the numbers n, m of variables and clauses respectively tend to infinity, while the ratio m/n tends to a constant r, the random formulas exhibit a *threshold* behaviour: if $r > 4.17$ (approximately) then almost all random formulas are unsatisfiable while the opposite is true if $r < 4.17$. The constant r is called the *density* of the formula. On the theoretical side, Friedgut [10] has proved that there exists a sequence γ_n such that for any $\epsilon > 0$, if $m/n \leq \gamma_n - \epsilon$ for sufficiently large n then the probability of a random formula being satisfiable approaches 0 while if $m/n \geq \gamma_n + \epsilon$ for sufficiently large n then this probability approaches 1 although it is not known if the sequence γ_n convergences to some constant value γ.

⋆ Research supported by EPSRC grant GR/L/77089.

Thus, finding the *exact* value of the threshold point or even proving that a threshold value exists is still a major problem in probability and complexity theory. Up to now, only upper and lower bounds have been rigorously established for the threshold value. The best lower bound has been recently proved by Achlioptas and Sorkin [1] and it is 3.26 while the currently best upper bound has been announced by Dubois et al. [7] and it is 4.506.

In this paper, we address the upper bound question for the unsatisfiability threshold from a new perspective that combines the idea of *local maximum* satisfying truth assignments proposed by Kirousis et al. [14], with the use of sharp estimates on some of the probabilities involved based on results about the so called coupon collector experiment (see for instance [17] and references thereafter). We obtain an upper bound of 4.571 thus improving over the best, with an available complete proof, previous upper bound (4.596 given in [12]). As a by-product of our proof, we also establish an upper bound to the q-binomial coefficients (a generalization of the binomial coefficients). Despite the extensive literature on q-binomial coefficients (see, e.g., [9, 11, 15]), no such bound was, to the best of our knowledge, known.

2 The method of local maxima

In this section, we will state briefly the methodology followed in [14] and obtain the starting upper bound on the probability that a random formula is satisfiable. Let $\mathcal{S}$ be the class of all truth assignments to n variables and $\mathcal{A}_n$ the (random) class of truth assignments that satisfy a random formula ϕ. For a given $A \in \mathcal{S}$, a *single flip sf* is the change in A of exactly one FALSE value to TRUE and by A^{sf} we denote the truth assignment that results from this change. We define as $\mathcal{A}_n^1 \subseteq \mathcal{A}_n$ the random class of truth assignments with the following two properties:

- $A \models \phi$,
- for every single flip *sf*, it holds $A^{sf} \not\models \phi$.

A partial order can be defined on $\mathcal{S}$: a truth assignment A is smaller than a truth assignment A' iff there exists an i such that both A and A' assign the same value to all variables x_j, for all $j < i$ while A assigns FALSE to x_i and A' assigns TRUE to it. The random class $\mathcal{A}_n^1$ coincides with the set of satisfying truth assigments that are *local maxima* with respect to the partial order defined above among satisfying truth assignments that differ in one bit.

A more restricted random class of truth assignments results from $\mathcal{A}_n^1$ if we extend the scope of locality in obtaining a local maximum. A *double flip* is the change of exactly two variables x_i and x_j (with $i < j$) where x_i is changed from FALSE to TRUE and x_j from TRUE to FALSE. In analogy with single flips, by A^{df} we denote the truth assignment that results from A if we apply the double flip *df*. Let $\mathcal{A}_n^{2\sharp}$ be defined as the set truth of assignments A that have the following properties:

- $A \models \phi$,
- for all single flips *sf*, it holds $A^{sf} \not\models \phi$,
- for all double flips *df*, it holds $A^{df} \not\models \phi$.

Our starting point is the following inequality:

Lemma 1. [14]

$$\Pr[\phi \text{ is satisfiable}] \leq \mathrm{E}[|\mathcal{A}_n^{2\sharp}|] = \sum\nolimits_{A \in S} \Pr[\forall df\ A^{df} \not\models \phi, \forall sf\ A^{sf} \not\models \phi, A \models \phi]$$
$$= (7/8)^{rn} \sum\nolimits_{A \in \mathcal{S}} \Pr[\forall df\ A^{df} \not\models \phi, \forall sf\ A^{sf} \not\models \phi \mid A \models \phi]$$
$$= (7/8)^{rn} \sum\nolimits_{A \in \mathcal{S}} \Pr[\forall sf\ A^{sf} \not\models \phi \mid A \models \phi] \cdot \Pr[\forall df\ A^{df} \not\models \phi \mid A \in \mathcal{A}_n^1]. \quad (1)$$

In order to find an upper bound for the unsatisfiability threshold, it suffices to find the smallest possible value for r for which the right-hand side of (1) tends to 0. In the sections to follow, we will describe the sequence of steps that will lead us to the determination of an upper bound on the probabilities that appear in the third line of (1).

3 Coupon collectors and single flips

For notational convenience, we will consider a formula ϕ as a set of clauses. Thus, the expression $\phi \cap \mathcal{A}$, with $\mathcal{A}$ a set of clauses, has the meaning of set intersection with the additional requirement that a clause that appears in the intersection, appears as many times as it appears in ϕ.

Given a truth assignment A, and a variable x such that $A(x) = \text{FALSE}$, the set of *critical clauses* for x in A, $\mathcal{B}(A, x)$, is the set of clauses whose unique TRUE literal is $\neg x$. Note that $|\mathcal{B}(A, x)| = \binom{n-1}{2}$ and $\mathcal{B}(A, x) \cap \mathcal{B}(A, y) = \emptyset$ for $x \neq y$.

Assuming A sets k variables FALSE, the probability $\Pr[\forall sf\ A^{sf} \not\models \phi \mid A \models \phi]$ in (1) is the ratio between a function $N(n, m, k)$, counting the number of ways to build a formula with m clauses out of n variables containing at least one critical clause for each of the k *critical* variables, and the total number of ways to build a formula on m clauses out of n variables which is satisfied by A. Hence

$$\Pr[\forall sf\ A^{sf} \not\models \phi \mid A \models \phi] = \frac{N(n,m,k)}{\left[7\binom{n}{3}\right]^m}$$

If ϕ contains $l \in \{k, k+1, \ldots, m\}$ critical clauses, then

$$\Pr[\forall sf\ A^{sf} \not\models \phi \mid A \models \phi] = \sum\nolimits_{l=k}^{m} \frac{C(n,m,k,l)\ R(n,m,k,l)}{\left[7\binom{n}{3}\right]^m}$$

where $C(n, m, k, l)$ counts the number of ways of choosing l critical clauses so that at least one member of $\mathcal{B}(A, x)$ is chosen for each of the k critical variables and $R(n, m, k, l)$ counts the number of ways of filling up the remainder of ϕ with $m - l$ clauses that are true under A but not critical.

Lemma 2. *For any choice of the parameters* $R(n, m, k, l) = \left(7\binom{n}{3} - k\binom{n-1}{2}\right)^{m-l}$.

Proof. There are $7\binom{n}{3}$ clauses consistent with A. If A forces k variables to be critical there are k disjoint groups of $\binom{n-1}{2}$ critical clauses. □

Lemma 3. *For any choice of the parameters* $C(n,m,k,l) = \binom{m}{l}\left[k\binom{n-1}{2}\right]^l \mathrm{coupon}(l,k)$ *where* $\mathrm{coupon}(l,k)$ *is the probability that a coupon collector picks* k *distinct random coupons over* l *trials.*

Proof. Assume that there are k critical variables associated with a given assignment A. Moreover ϕ contains l critical clauses. There are $\binom{m}{l}$ ways of choosing l positions out of the m available. Also, there are $k\binom{n-1}{2}$ critical clauses. Therefore, if we do not distinguish among the non-critical clauses, there are $\binom{m}{l}\left[k\binom{n-1}{2}\right]^l$ ways of choosing a sequence of m clauses so that exactly l of them are critical. Since $C(n,m,k,l)$ counts the number of these which has at least one occurrence of a critical clause for each of the k critical variables, and since there are equal numbers of possible critical clauses for each variable, the ratio of these terms is the probability $\mathrm{coupon}(l,k)$. □

To be able to state the main result in this section we need to quote a result giving asymptotic approximations to the probabilities $\mathrm{coupon}(l,k)$.

Theorem 1. [4] *Let* $x = l/k$ *with* $l = \Theta(k)$. *For all* $x > 1$ *define* $g_1(x) =_{df} (\mathrm{e}^{r_0} - 1)\left(\frac{x}{\mathrm{e}r_0}\right)^x$ *where* r_0 *is the solution of* $\frac{r\mathrm{e}^r}{\mathrm{e}^r-1} = x$. *Also let* $g_1(1) = \mathrm{e}^{-1}$. *Then for all sufficiently large integer* k *and all* $x \geq 1$, $\mathrm{coupon}(l,k) \sim g_1(x)^k$.

The proof of the following theorem is entailed by the argument above, the use of the estimate given in Theorem 1 and Stirling's approximation to the various factorials involved. In the following result $F \asymp G$ denotes the fact that $\ln F \sim \ln G$. So for example $\binom{m}{l} \asymp \left[(\frac{rn}{l})^{\frac{l}{rn}}(\frac{rn}{rn-l})^{(1-\frac{l}{rn})}\right]^{rn}$.

Theorem 2. *The probability that a truth assignment* A *with* αn FALSE *values is a local maximum satisfies:*

$$\Pr[\forall sf\, A^{sf} \not\models \phi \mid A \models \phi] \asymp \sum_{l=\alpha n}^{rn} \left(\frac{3\alpha rn}{7l}\right)^l \left(\frac{(7-3\alpha)rn}{7(rn-l)}\right)^{rn-l} g_1\left(\frac{l}{\alpha n}\right)^{\alpha n}. \tag{2}$$

An important remark is that in the expression given in Theorem 2, two *polynomially large* factors have been omitted: one implicit in the relation "$\sim$" used in Theorem 1 and one related to the asymptotics of the binomial coefficients. However, for our goal of making a certain expression that contains (2) converge to 0, such factors are immaterial and what is required is an optimal estimate only for the exponential factors which is guaranteed by Theorem 1 and the asymptotics for the binomial coefficients given above.

4 Probability models for random formulas

A random 3-SAT formula ϕ with $m = rn$ clauses is most commonly formed according to one of the following probability models (Ω is the set of $8\binom{n}{3}$ possible 3-SAT clauses):

1. Select the m clauses of ϕ, drawing each clause uniformly and independently from Ω, with replacement (model G_{mm}).
2. Select the m clauses of ϕ, drawing each clause uniformly and independently from Ω, without replacement (model G_m).

3. With probability $p(n)$ each clause is chosen independendly of the others and with probability $p(n)$ for inclusion in ϕ (model G_p).

The probability that a random formula ϕ generated according to model G_m, G_{mm} or G_p belongs to a set Q defining some property, is denoted by $\Pr_m[\phi \in Q], \Pr_{mm}[\phi \in Q]$ and $\Pr_p[\phi \in Q]$ respectively. Notice that the probabilities in (1) are all in G_{mm} since the model we considered until now allows clause repetitions when forming a formula. We will now outline an argument showing that the second probability in the third line of (1) can be rewritten into the G_p model, in order to take advantage of the computation of this probability in G_p that has already been performed in [14].

Consider again an arbitrary but fixed truth assignment A. As all probabilities which will undergo a change in the probabilistic model are conditional on $A \models \phi$, in the considerations below we assume that the universe of all clauses is restricted to those that are satisfied by A, and consequently that $p(n) = \frac{rn}{7\binom{n}{3}} \sim \frac{6r}{7n^2}$.

First let "NoRep" be the event that ϕ has no two clauses identical and let $\overline{\mathrm{NoRep}}$ its complement. Then, because the order of the number of all possible clauses is $\Theta(n^3)$ and the order of the number of the clauses contained in ϕ is $\Theta(n)$, $\lim_{n\to\infty} \Pr_{mm}[\overline{\mathrm{NoRep}}] = 0$.

Now let Q_1 and Q_2 be two arbitrary events such that the following two conditions, which we call *regularity* conditions hold:

– For some $\epsilon > 0$ and for all n, $\ln\left(\Pr_{mm}[Q_2|Q_1]\right) < -\epsilon$, i.e. $\Pr_{mm}[Q_2|Q_1]$ is bounded away from 1.
– $\lim_{n\to\infty} \Pr_{mm}[\overline{\mathrm{NoRep}}|Q_1, Q_2] = \lim_{n\to\infty} \Pr_{mm}[\overline{\mathrm{NoRep}}|Q_1] = 0$.

Notice that the events we consider in this paper have probabilities (conditional or not) that are exponentially small, so the first of the two regularity conditions is satisfied. Also, the second regularity condition is true when Q_1 and Q_2 are the events $A \in \mathcal{A}_n^1$ and $\forall df\ A^{df} \not\models \phi$, respectively. Indeed both these events and their conjunction are negatively correlated with $\overline{\mathrm{NoRep}}$, so $\Pr_{mm}[\overline{\mathrm{NoRep}}|Q_1] \le \Pr_{mm}[\overline{\mathrm{NoRep}}] \to 0$ and similarly for $\Pr_{mm}[\overline{\mathrm{NoRep}}|Q_1, Q_2]$. To prove the negative correlation claim for, say, Q_1 and $\overline{\mathrm{NoRep}}$, observe that the correlation claim is equivalent to $\Pr_{mm}[Q_1|\mathrm{NoRep}] \ge \Pr_{mm}[Q_1]$, which in turn is equivalent to $\Pr_m[Q_1] \ge \Pr_{mm}[Q_1]$. This last inequality is intuitively obvious under the assumption that $A \models \phi$, because the probability to get at least a critical clause for each critical variable of the satisfying truth assignment A increases when the clauses of the formula are assumed to be different. For a formal proof of this for general increasing and reducible properties (like Q_1 and Q_2), we refer to [13]. Therefore, the second regularity condition is also true for the probabilities we will consider below.

Under the above regularity conditions, we have that:

$$\Pr_m[Q_2|Q_1] \asymp \Pr_{mm}[Q_2|Q_1]. \tag{3}$$

Indeed,

$$\Pr_m[Q_2\,|Q_1] = \Pr_{mm}[Q_2\,|Q_1, \mathrm{NoRep}] = \frac{\Pr_{mm}[Q_2\,|Q_1] - \Pr_{mm}[Q_2 \wedge \overline{\mathrm{NoRep}}\,|Q_1]}{1 - \Pr_{mm}[\overline{\mathrm{NoRep}}\,|Q_1]}$$
$$= \Pr_{mm}[Q_2\,|Q_1]\frac{1 - \Pr_{mm}[\overline{\mathrm{NoRep}}\,|Q_2, Q_1]}{1 - \Pr_{mm}[\overline{\mathrm{NoRep}}\,|Q_1]}.$$

Now first taking logarithms, then dividing both sides with $\ln(\Pr_{mm}[Q_2|Q_1])$ and finally letting $n \to \infty$, we get the required inequality (the regularity conditions are needed in the computation of the limits).

On the other hand, it follows easily from Theorem II.2 (iii) in [3] that:

$$\Pr_m[Q_2 \,|Q_1] \leq 3m^{1/2}\Pr_p[Q_2 \,|Q_1]. \tag{4}$$

The inequality above has been proved by Bollobas for an arbitrary unconditional event. In general, it might not be true for conditional events. However, if both Q_2 and Q_1 are monotone increasing (i.e. $\Pr_{m_2}[Q_i] \geq \Pr_{m_1}[Q_i]$ for any $m_2 \geq m_1$, for both $i = 1, 2$), then it still holds true. For an informal explanation why this is indeed so, first observe that the conditional Q_1, being monotone increasing, "forces more clauses" into the formula in the variable-length G_p model. Since Q_2 is also monotone increasing, this has as a consequence that the conditional probability of Q_2 in G_p deviates even further to the right from the corresponding conditional probability in the fixed-length model G_m. A formal proof of this, based in the four-functions theorem of Ahlswede and Daykin [8] (see [2] for a nice presentation) will be given in the full paper. Finally, notice that the events Q_2 and Q_1, for which we apply this inequality below are trivially monotone increasing.

Now, (3) (4) and (1) imply that $\Pr_{mm}[\phi$ is satisfiable$]$ is at most (ignoring polynomial factors):

$$\left(\frac{7}{8}\right)^{rn} \sum_{A \in \mathcal{S}} \Pr_{mm}[\forall sf\, A^{sf} \not\models \phi \mid A \models \phi] \cdot \Pr_p[\forall df\, A^{df} \not\models \phi \mid A \in \mathcal{A}_n^1]. \tag{5}$$

We are now in a position to use the probability calculations performed earlier in this paper using the coupon collector analogy for the method of local maxima with the probability calculations in [14] in order to derive a value for r that makes the right-hand side of (5) converge to 0. It is perhaps interesting to note that $\Pr_p[A \models \phi] \not\asymp \Pr_{mm}[A \models \phi]$ (the former is larger). Therefore, as both $\Pr_p[A \models \phi]$ and $\Pr_{mm}[A \models \phi]$ are easily computable, it was advantageous to retain in the first factor of the right hand side of (5) the value of $\Pr[A \models \phi](= (7/8)^{rn})$ computed in the model G_{mm} rather than replace it with its value in G_p ($\Pr_p[A \models \phi] \asymp (e^{-(1/8)rn})$. Also, $\Pr_p[\forall sf\, A^{sf} \not\models \phi \mid A \models \phi] \not\asymp \Pr_{mm}[\forall sf\, A^{sf} \not\models \phi \mid A \models \phi]$. The coupon collector analogy helped us exploiting the advantage of the model G_{mm} for the probability of single flips. We were unable to do the same for the computation of $\Pr_{mm}[\forall df\, A^{df} \not\models \phi \mid A \in \mathcal{A}_n^1]$, so we resorted to the "easier", but worse, model G_p (in contrast, by Equation (3), the models G_{mm} and G_m are not asymptotically distinguishable, ignoring polynomial factors).

5 Calculations

Let *sf*(A) denote the number of FALSE values assigned by a truth assignment A. We define the following functions of r:

$$\begin{aligned} u &= e^{-r/7} \\ z &= -\frac{6u^6 \ln(1/u)}{1-u^3} + \frac{18u^9 \ln^2(1/u)}{(1-u^3)^2}\phi_2\left(\frac{6u^6 \ln(1/u)}{1-u^3}\right) \end{aligned} \tag{6}$$

$$X(sf(A)) = \Pr_{mm}[\forall sf\, A^{sf} \not\models \phi \mid A \models \phi] \quad (7)$$

$$Y = 1 + z\frac{1}{n} + o\left(\frac{1}{n}\right), \quad \text{(notice that } z < 0\text{)}, \quad (8)$$

where $\phi_2(t)$ is the smallest root of $\phi_2(t) = e^{t\phi_2(t)}$ and can also be expressed by means of the Lambert $\mathcal{W}$ function [6]. Although ϕ_2 is defined only in the interval $[0, e^{-1}]$, we have verified that the argument to ϕ_2 in the definition of z in (6) lies in this interval for all $r \geq 0$. In [14], it was proved that

$$\Pr_p[\forall df\, A^{df} \not\models \phi \mid A \in \mathcal{A}_n^1] \leq Y^{df(A)}.$$

Therefore, using Equations (7) and (8), expression (5) may be written as follows:

$$\left(\frac{7}{8}\right)^{rn} \sum\nolimits_{A \in \mathcal{A}_n} X(sf(A)) Y^{df(A)}. \quad (9)$$

Furthermore the following equality can be derived [14] by induction on n:

$$\sum\nolimits_{A \in \mathcal{A}_n} X(sf(A)) Y^{df(A)} = \sum_{k=0}^{n} \binom{n}{k}_Y X(k), \quad (10)$$

where $\binom{n}{k}_q$ denotes the *q-binomial* or *Gaussian* coefficients (see [11]). From expression (9), the definition of $X(sf(A))$ in (7), Equality (10) and Theorem 2, we deduce that $\Pr_{mm}[\phi$ is satisfiable$]$ is at most (omitting a polynomial factor)

$$\left(\frac{7}{8}\right)^{rn} \sum_{k=0}^{n} \sum_{l=k}^{rn} \binom{n}{k}_Y E\left(\frac{k}{n}, \frac{l}{rn}, r, n\right), \quad (11)$$

where

$$E(\alpha, \beta, r, n) = \left[\left(\frac{3\alpha}{7\beta}\right)^{\beta r} \left(\frac{7-3\alpha}{7(1-\beta)}\right)^{r - r\beta} g_1\left(\frac{r\beta}{\alpha}\right)^{\alpha}\right]^n$$

We will now consider an arbitrary term of the double sum that appears in (11) and examine for which values of r it converges to 0. If we find a condition on r that forces all such terms to converge to 0, then the whole sum will converge to 0 since it contains polynomially many terms, all of which vanish exponentially fast. This technique avoids the problem of finding a closed-form upper bound for the sum itself. However, in order to handle an arbitrary term, we need an upper bound for the q-binomial coefficients. To establish such a bound one will need the following standard result:

Lemma 4. [18] *Let $f(z) = \sum_{i=0}^{\infty} f_i z^i$ be the generating function for the sequence f_i, $i \geq 0$. Then if $f(z)$ is analytic in $|z| < R$ and if $f_i \geq 0$ for all $i \geq 0$, then for any t, $0 < t < R$, and any $n \geq 0$, it holds that $f_n \leq t^{-n} f(t)$.*

Using this lemma, we can prove the following:

Theorem 3. *Let $\binom{n}{\alpha n}_q$ denote the q-binomial coefficients for α real in $(0, 1)$ and αn an integer. Then the following inequality holds:*

$$\binom{n}{\alpha n}_q \leq 2q^{-\binom{n}{2}} x_0^{-n} e^{\frac{1}{\ln q}[\mathrm{dilog}(1+x_0) - \mathrm{dilog}(1+x_0 q^{n-1})]} \quad (12)$$

where $x_0 = \frac{1-q^{\alpha n}}{q^{\alpha n} - q^{n-1}}$ and $\mathrm{dilog}(x) = \int_1^x \frac{\ln t}{1-t} dt$.

Proof. For the ordinary generating function of $q^{\binom{i}{2}}\binom{n}{i}_q$ it holds [5, p.118]

$$\sum_{i=0}^{n} q^{\binom{i}{2}}\binom{n}{i}_q x^n = \prod_{i=1}^{n}(1+xq^{i-1}) = e^{\sum_{i=1}^{n}\ln(1+xq^{i-1})}$$
$$= (1+x)\cdot e^{\sum_{i=2}^{n}\ln(1+xq^{i-1})}.$$

Since $\ln(1+xq^{i-1})$ is decreasing in i,

$$\sum_{i=0}^{n} q^{\binom{i}{2}}\binom{n}{i}_q x^i \leq (1+x)\cdot e^{\int_1^n \ln(1+xq^{i-1})di} = (1+x)\cdot e^{\frac{1}{\ln q}[\mathrm{dilog}(1+x)-\mathrm{dilog}(1+xq^{n-1})]}.$$

Applying Lemma 4 and using the fact that $x \leq 1$, we obtain the inequality

$$q^{\binom{i}{2}}\binom{n}{i}_q \leq x^{-i}(1+x)\cdot e^{\frac{1}{\ln q}[\mathrm{dilog}(1+x)-\mathrm{dilog}(1+xq^{n-1})]}. \tag{13}$$

The above inequality holds for any value of $x \in (0,1)$. Therefore, we may optimize it by choosing the value $x_0 = \frac{1-q^i}{q^i-q^{n-1}}$ that minimizes the expression on the right-hand side of (13). From this we obtain

$$\binom{n}{i}_q \leq 2q^{-\binom{i}{2}}x_0^{-i}e^{\frac{1}{\ln q}[\mathrm{dilog}(1+x_0)-\mathrm{dilog}(1+x_0q^{n-1})]}. \tag{14}$$

which gives the required inequality by setting $i = \alpha n$. □

Setting $q = Y = 1 + z/n$ in (12) and using the approximation $\ln(1+z/n) \sim z/n$, as $n \to \infty$, the following can be derived:

$$\binom{n}{\alpha n}_q \leq 2\left[\left(\frac{1}{x_0}\right)^{\alpha}\cdot e^{-\frac{\alpha^2 z}{2}+\frac{1}{z}[\mathrm{dilog}(1+x_0)-\mathrm{dilog}(1+x_0e^z)]}\right]^n, \tag{15}$$

where $x_0 = \frac{1-e^{\alpha z}}{e^{\alpha z}-e^z}$, which is expedient in the proof of the following:

Theorem 4. *An arbitrary term of the double sum in (11) is asymptotically (ignoring polynomial multiplicative factors) bounded from above by:*

$$T_r(\alpha,\beta)^n = \left[\left(\frac{3\alpha}{7\beta}\right)^{\beta r}\left(\frac{7-3\alpha}{7(1-\beta)}\right)^{r-r\beta} g_1\left(\frac{r\beta}{\alpha}\right)^{\alpha}\frac{e^{-\frac{\alpha^2 z}{2}+\frac{1}{z}[\mathrm{dilog}(1+x_0)-\mathrm{dilog}(1+x_0e^z)]}}{x_0{}^{\alpha}}\right]^n$$

where, $\alpha = \frac{k}{n}$, $\beta = \frac{l}{rn}$, $x_0 = \frac{1-e^{\alpha z}}{e^{\alpha z}-e^z}$, z *as given in (6).*

An immediate consequence of this result is that the smallest value of r for which $T_r(\alpha,\beta)$ is smaller than 1 for all $\alpha \in (0,1)$ and $\beta \in (\alpha/r, 1)$ is an upper bound for the unsatisfi ability threshold.

We fi nally claim that for any value of r, the expression $\ln T_r(\alpha,\beta)$ is a convex function of α,β over the domain $\mathcal{D} = \{\alpha,\beta \in [0,1]$ and $\frac{\beta r}{\alpha} \geq 1\}$. Therefore we will compute its *unique* maximum value for $r = 4.571$ and $(\alpha,\beta) \in \mathcal{D}$. Due to the complexity of the expression for $\ln T_r(\alpha,\beta)$, we maximized it numerically using a Maple [16] implementation of *Downhill Simplex*. This implementation is based on the method and the code described in [19] and it is freely distributed by F.J. Wright in his Web page [21]. Using the plots of $T_r(\alpha, \beta)$ we obtained with Maple, we chose as a starting set of values for the downhill simplex algorithm the values $(\alpha, \beta) = (0.42, 0.21)$ and we set the accuracy and the scale parameters equal to 10^{-50}. In addition, we set the *Digits* parameter of Maple (accuracy of floating point numbers) equal to 100. We ran the algorithm and it returned as the maximum value of $\ln T_r(\alpha, \beta)$ the value -0.0000884. Additionally, we computed all the partial derivatives of $\ln T_r(\alpha, \beta)$ at the point where downhill simplex claims that it has located the maximum and they were found to be numerically equal to 0. Therefore, this provides additional support that at this point the function attains its maximum. As a final check, we generated 30000 random points close to the point at which downhill simplex finds the maximum of $\ln T_r(\alpha, \beta)$ and we confirmed that the value of $\ln T_r(\alpha, \beta)$ is not above the value returned by the method. All these considerations show that -0.0000884 is a global maximum of $\ln T_r(\alpha, \beta)$, which establishes the value $r = 4.571$ as an upper bound to the unsatisfiability threshold.

6 Discussion

We derived an upper bound for the unsatisfiability threshold that improves over all previously proved upper bounds, except the one announced in [7]. We looked at the problem from a new perspective, by combining the method of local maximum satisfying truth assignments proposed in [14] with the sharp estimates on some of the probabilities involved based on the coupon collector experiment. In addition, we gave a relationship between two conditional probability spaces for generating random formulas that allowed us to use probability calculations performed in the easier to handle probability model according to which each of the clause is selected independently of the others with some fixed probability to appear in the formula. As a final ingredient, we proved a tight upper bound for the q-binomial coefficients that may be of interest in its own right. Our approach showed that the unsatisfiability threshold is less than 4.571. This bound improves over the best previous upper bound with a complete proof (4.596, see [12]). Dubois et al. in [7] have announced the value 4.506 but to the best of our knowledge, no complete proof is available yet from the authors. Nevertheless, we believe that our approach contains elements of a separate value and interest that might be useful in another context or in other applications: exact (in the exponential order) computation of the first probability in the last line of (1) using the coupon collector problem, relationships between conditional probability models and an upper bound for the q-binomial coefficients. And even though we used a numerical method for maximizing our function in order to show that for $r = 4.571$ it is strictly below 1 for every legal value of its parameters, our proof that the function is convex and the observation that its derivatives are bounded, renders our proof essentially rigorous since the Downhill Simplex method is certain to find a *global* maximum within *guaranteed* accuracy. We also believe that our approach of using the occupancy problem for the accurate computation of the first probability in (1) can be extented in order to give a sharp estimate also for the second probability in the last line of (1). To this end, we are currently working on extending the coupon collector approach or some similar scheme to model *double flips*. If this is accomplished, then it is conceivable that an analogue of Theorem 2 can be proved that will enable further improvements on the value obtained in this paper. There is also the question of combining this approach with the idea of "typical formulas" proposed in [7], thus obtaining still better bounds, under 4.5.

References

1. D. Achlioptas, G. B. Sorkin, "Optimal Myopic Algorithms for Random 3-SAT" in: *Proc. 41st Annual Symposium on Foundations of Computer Science (FOCS)*, 2000. pp 590–600.
2. N. Alon, J. H. Spencer, P. Erdős, *The Probabilistic Method*, John Wiley and Sons, 1992.
3. B. Bollobás, *Random Graphs*, Academic Press, London, 1985.
4. V. Chvátal, "Almost all graphs with $1.44n$ edges are 3-colourable," *Random Structures and Algorithms* 2, pp 11–28, 1991.
5. L. Comtet, *Advanced Combinatorics; The Art of Finite and Infinite Expansions*. D. Reidel, 1974.
6. R.M. Corless, G.H. Gonnet, D.E.G. Hare, D.J. Jeffrey, D.E. Knuth, "On the Lambert W function," manuscript, Computer Science Department, University of Waterloo.
7. O. Dubois, Y. Boufkhad, J. Mandler, "Typical random 3-SAT formulae and the satisfiability threshold," in: *Proc. 11th Symposium on Discrete Algorithms (SODA)*, pp 126–127, 2000.
8. C. M. Fortuin, R. W. Kasteleyn, J. Ginibre, Correlation inequalities on some partially ordered sets, *Commun. Math. Physics* 22, pp. 89–103, 1971.
9. N.J. Fine, *Basic Hypergeometric Series and Applications*, Mathematical Surveys and Monographs, Number 27, 1980.
10. E. Friedgut, appendix by J. Bourgain, "Sharp thresholds of graph properties, and the k-sat problem," *J. Amer. Math. Soc.* 12, pp 1017–1054, 1999.
11. G. Gasper, M. Rahman, *Basic Hypergeometric Series*, Encyclopedia of Mathematics and its Applications, vol 35, Cambridge University Press, Cambridge, 1990.
12. S. Janson, Y.C. Stamatiou, M. Vamvakari, "Bounding the unsatisfiability threshold of random 3-SAT," *Random Structures and Algorithms*, 17, pp 108–116, 2000.
13. L.M. Kirousis, Y.C. Stamatiou, *An inequality for reducible, increasing properties of randomly generated words*, Tech. Rep. TR-96.10.34, C.T.I., University of Patras, Greece, 1996.
14. L.M. Kirousis, E. Kranakis, D. Krizanc, Y.C. Stamatiou, "Approximating the unsatisfiability threshold of random formulas," *Random Structures and Algorithms* 12, pp 253–269, 1998.
15. D.E. Knuth, *Fundamental Algorithms*, The Art of Computer Programming, 2nd ed., 1973.
16. M.B. Monagan, K.O. Geddes, K.M. Heal, G. Labahn, S.M. Vorkoetter, *Programming Guide*, Maple V Release 5, Springer–Verlag, 1998.
17. R. Motwani, P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
18. A.M. Odlyzko, "Asymptotic Enumeration Methods," in: R.L. Graham, M. Grötschel, and L. Lovász, eds. *Handbook of Combinatorics*, Chapter 22, 1063–1229, Elsevier, 1995.
19. W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed., Cambridge University Press, 1993.
20. A.E. Taylor, W.R Mann, *Advanced Calculus*, 3rd ed., John Wiley & Sons, 1983.
21. F.J. Wright: `http://centaur.maths.qmw.ac.uk/Computer_Algebra/`. A Maple (V-5) implementation of *Downhill Simplex* based on code given in [19].
22. M. Zito, Randomised Techniques in Combinatorial Algorithmics, PhD Thesis, Department of Computer Science, University of Warwick, November 1999.