

Stochastic Local Search Algorithms for DNA Word Design

Dan C. Tulpan, Holger H. Hoos, and Anne E. Condon

Department of Computer Science
University of British Columbia
Vancouver, B.C., V6T 1Z4, Canada
{dctulpan,hoos,condon}@cs.ubc.ca
<http://www.cs.ubc.ca/labs/beta>

Abstract. We present results on the performance of a stochastic local search algorithm for the design of DNA codes, namely sets of equal-length words over the nucleotides alphabet $\{A, C, G, T\}$ that satisfy certain combinatorial constraints. Using empirical analysis of the algorithm, we gain insight on good design principles. We report several cases in which our algorithm finds word sets that match or exceed the best previously known constructions.¹

Keywords: DNA Word Design, Combinatorics, Stochastic Local Search, Coding Theory, Hamming Distance, GC Content

1 Introduction

The design of DNA code words, or sets of short DNA strands that satisfy combinatorial constraints, is motivated by the tasks of storing information in DNA strands used for computation or as molecular bar-codes in chemical libraries [3,4,8,20]. Good word design is important in order to minimise errors due to non-specific hybridisation between distinct words and their complements, to achieve a higher information density, and to obtain large sets of words for large-scale applications.

For the types of combinatorial constraints typically desired, there are no known efficient algorithms for design of DNA word sets. Techniques from coding theory have been applied to design of DNA word sets [4,10]. While valuable, this approach is hampered by the complexity of the combinatorial constraints on the word sets, which are often hard to reason about theoretically. For these reasons, heuristic approaches such as stochastic local search offer much promise in the design of word sets.

¹ This material is based upon work supported by the U.S. National Science Foundation under Grant No. 0130108, by the National Sciences and the Engineering Research Council of Canada.

Stochastic local search (SLS) algorithms strongly use randomised decisions while searching for solutions to a given problem. They play an increasingly important role for solving hard combinatorial problems from various domains of Artificial Intelligence and Operations Research, such as satisfiability, constraint satisfaction, planning, scheduling, and other application areas. Over the past few years there has been considerable success in developing stochastic local search algorithms as well as randomised systematic search methods for solving these problems, and to date, stochastic search algorithms are amongst the best known techniques for solving problems from many domains. Detailed empirical studies are crucial for the analysis and development of such high-performance stochastic search techniques.

Stochastic search methods have been used successfully for decades in the construction of good binary codes (see for example [7].) Typically, the focus of this work is in finding codes of size greater than the best previously known bound, and a detailed empirical analysis of the search algorithms is not presented. Stochastic search techniques have also been applied to the design of DNA word sets (see Section 3). However, some algorithmic details are not specified in these papers. In addition, while small sets of code words produced by the algorithms have been presented (and the papers make other contributions independent of the word design algorithms), little or no analysis of algorithm performance is provided. Consequently, it is not possible to extract general insights on the design of stochastic algorithms for code design or to compare their approaches in detail with other algorithms. Our goal is to understand what algorithmic principles are most effective in the application of stochastic local search methods to the design of DNA or RNA word sets (and more generally, codes over other alphabets, particularly the binary alphabet).

Towards this end, we describe a simple stochastic local search algorithm for design of DNA codes, and analyze its performance using an empirical methodology based on run-time distributions [15]. In this study, we have chosen to design word sets that satisfy one or more of the following constraints: Hamming distance (**HD**), GC content (**GC**), and reverse complement Hamming distance (**RC**). We define these constraints precisely in Section 2. Our reason for considering these constraints is that there are already some constructions for word sets satisfying these constraints, obtained using both theoretical and experimental methods, with which we can compare our results. (In follow-up work, we will apply our methods to other, more realistic constraints.)

Our algorithm, described in detail in Section 4, takes as input the desired word length and set size, along with a specification of which constraints the set should satisfy, and attempts to find a set that meets these requirements. The algorithm performs stochastic local search in a space of DNA word sets of fixed size that may violate the given constraints, using an underlying search strategy that is based on a combination of randomised iterative improvement and conflict-directed random walk. The basic algorithm is initialised with a randomly selected set of DNA words. Then, repeatedly a conflict, that is, a pair of words that violates a constraint, is selected and resolved by modifying one of the respective

words. The algorithm terminates if a set of DNA words that satisfies all given constraints is found, or if a specified number of iterations have been completed.

The performance of this algorithm is primarily controlled by a so-called noise parameter that determines the probability of greedy vs. random conflict resolution. Interestingly, optimal settings of this parameter appear to be consistent for different problem instances (word set sizes) and constraints.

Our empirical results, reported in Section 5, show that the run-time distributions that characterise our algorithm's performance on hard word design problems often have a "fat" right tail. This indicates a stagnation of the underlying search process that severely compromises performance. As a first approach to overcome this stagnation effect, we extended our algorithm with a mechanism for diversifying the search by occasional random replacement of a small fraction of the current set of DNA words. Empirically, this modification eliminates the stagnation behavior and leads to substantial improvements in the performance of our algorithm.

We compared the sizes of the word sets obtainable by our algorithm with previously known word sets, starting with the already extensively studied case of word sets that satisfy the Hamming distance constraint only. Out of 42 tests, our algorithm was able to find a word set whose size matches the best known theoretical construction in 16 cases. We also did several comparisons with word sets that satisfy at least two of our three constraints, for which again previous results were available. Out of a total of 42 comparisons with previous results that satisfy Hamming distance and Reverse Complement constraints, we found word sets that improved on previous constructions in all but one case.

2 Problem Description

The DNA code design problem that we consider is: given a target k and word length n , find a set of k DNA words, each of length n , satisfying certain combinatorial constraints. A DNA word of length n is simply a string of length n over the alphabet $\{A, C, G, T\}$, and naturally corresponds to a DNA strand with the left end of the string corresponding to the 5' end of the DNA strand. The constraints that we consider are:

- **Hamming Distance Constraint (HD)**: for all pairs of distinct words w_1, w_2 in the set, $H(w_1, w_2) \geq d$. Here, $H(w_1, w_2)$ represents the Hamming distance between words w_1 and w_2 , namely the number of positions i at which the i th letter in w_1 differs from the i th letter in w_2 .
- **GC Content Constraint (GC)**: a fixed percentage of the nucleotides within each word is either G or C. Throughout, we assume that this percentage is 50%.
- **Reverse Complement Hamming Distance Constraint (RC)**: for all pairs of DNA words w_1 and w_2 in the set, where w_1 may equal w_2 , $H(w_1, wcc(w_2)) \geq d$. Here, $wcc(w)$ denotes the Watson-Crick complement of DNA word w , obtained by reversing w and then by replacing each A in w by T and vice versa, and replacing each C in w by G and vice versa.

We note that there are many alternative formulations of the problem of designing DNA words that better capture the chemical motivation of ensuring that the DNA words behave well in a laboratory experiment (see for example [11]). For example, most formulations insist that the Hamming distance between shifted copies of words is also above a threshold. Even better than designing for Hamming constraints is to directly take thermodynamic constraints into account, so that the designed words have uniformly high melting temperature and the melting temperature of mismatched pairs of words is relatively low [19].

However we believe that the simple constraints above are best suited to the goals of the work presented here, for the following reasons. First, in order to understand whether our algorithmic techniques can find large sets of words, it is very useful to be able to compare the set sizes output by our algorithm with previous constructions, and this is possible here. Second, there are cases where one does not need to take shifts into account in a word design, for example when spacers or word labels are used between words [10] or when relatively short words are concatenated to form word labels [4]. In these cases, designs according to the three constraints above are already directly useful. Third, it may be that in order to design sets of words that satisfy more realistic thermodynamic constraints, a good initial starting point would be a set designed using the simpler combinatorial constraints above (see for example [2]). Finally, it is our hope that once we have identified good stochastic search principles for DNA word design for the above constraints, these principles will also be effective in designing word sets that satisfy other constraints, and this is a research direction that we are actively pursuing.

The huge number of possible sets that must be explored in order to find a big set of words suggests the use of non-exhaustive, heuristic search algorithms for solving this type of problems.

3 Related Work

Deaton et al. [5,6] and Zhang and Shin [21] describe genetic algorithms for finding DNA codes that satisfy much stronger constraints than the HD and RC constraints, in which the number of mismatches must exceed a threshold even when the words are aligned with shift. However, they do not provide a detailed analysis of the performance of their algorithms. Hartemink et al. [12] used a computer algorithm for designing word sets that satisfy yet other constraints, in which a large pool (several billion) of strands were screened in order to determine whether they meet the constraints. Several other researchers have used computer algorithms to generate word sets (see for example [3]), but provide no details on the algorithms.

Some DNA word design programs are publicly available. The DNASequences-Generator program [9] designs DNA sequences that satisfy certain subword distance constraints and, in addition, have melting temperature or GC content within prescribed ranges. The program can generate DNA sequences *de novo*, or integrate partially specified words or existing words into the set. The PER-

MUTE program was used to design the sequences of Faulhammer et al. [8] for their RNA-based 10-variable computation.

4 A Stochastic Local Search Algorithm

Our basic stochastic local search algorithm performs a randomised iterative improvement search in the space of DNA word sets; this overall approach is similar to WalkSAT, one of the best known algorithms for the propositional satisfiability problem [18,13] and Casanova, one of the best-performing algorithms for the winner determination problem in combinatorial auctions [14]. All word sets constructed during the search process have exactly the target number of words. Additionally, when using the **GC** constraint, we ensure that all words in the given candidate set always have the prescribed GC content.

The other two combinatorial constraints considered in this paper, the **HD** and **RC** constraints, are binary predicates over DNA words that have to be satisfied between any pair of words from a given DNA set. The function our algorithm is trying to minimise is the number of pairs that violate the given constraint(s). Figure 1 gives a general outline of our basic algorithm. In the following, we will describe this algorithm in more detail.

The initial set of words is determined by a simple randomised process that generates any DNA word of length n with equal probability. If the GC content constraint is used, we check each word that is generated and only accept words with the specified GC content. This works well for 50% GC content and typical designs use a GC content close to 50%, but could be easily made more efficient for less balanced GC contents if needed.

Our implementation also provides the possibility to initialise the search with a given set of DNA words (such sets could be obtained using arbitrary other methods); if such a set has less than k words, it is expanded with randomly generated words such that a set of k words is always obtained. Note that the initial word set may contain multiple copies of the same word.

In each *step* of the search process (i.e. one execution of the inner for loop from Figure 1), first, a pair of words violating one of the Hamming distance constraints is selected uniformly at random. Then, for each of these words, all possible single-base modifications are considered. As an example of single-base modifications take the code word *ACTT* of length 4. A new code word *GCTT* can be obtained by replacing letter *A* from the first code word with letter *G*. When the GC content constraint is used, this is restricted to modifications that maintain the specified GC content. For a pair of words of length n without the GC content constraint, this yields $6n$ new words. (Some of these might be identical.)

With a fixed probability θ , one of these modifications is accepted uniformly at random, regardless of the number of constraint violations that will result from it. In the remaining cases, each modification is assigned a score, defined as the net decrease in the number of constraint violations caused by it, and a modification with maximal score is accepted. (If there are multiple such modifications, one of

```

procedure Stochastic Local Search for DNA Word Design
  input: Number of words ( $k$ ), word length ( $n$ ), set of constraints ( $C$ )
  output: Set  $S$  of  $m$  words that fully or partially satisfies  $C$ 
  for  $i := 1$  to  $\text{maxTries}$  do
     $S :=$  initial set of words
     $\hat{S} := S$ 
    for  $j := 1$  to  $\text{maxSteps}$  do
      if  $S$  satisfies all constraints then
        return  $S$ 
      end if
      Randomly select words  $w_1, w_2 \in S$  that violate one of the constraints
       $M_1 :=$  all words obtained from  $w_1$  by substituting one base
       $M_2 :=$  all words obtained from  $w_2$  by substituting one base
      with probability  $\theta$  do
        select word  $w'$  from  $M_1 \cup M_2$  at random
      otherwise
        select word  $w'$  from  $M_1 \cup M_2$  such that number of conflict violations
        is maximally decreased
      end with probability
      if  $w' \in M_1$  then
        replace  $w_1$  by  $w'$  in  $S$ 
      else
        replace  $w_2$  by  $w'$  in  $S$ 
      end if
      if  $S$  has no more constraint violations than  $\hat{S}$  then
         $\hat{S} := S;$ 
      end if
    end for
  end for
  return  $\hat{S}$ 
end Stochastic Local Search for DNA Word Design

```

Fig. 1. Outline of a general stochastic local search procedure for DNA word design.

them is chosen uniformly at random.) Note that using this scheme, in each step of the algorithm, exactly one base in one word is modified.

The parameter θ , also called the noise parameter, controls the greediness of the search process; for high values of θ , constraint violations are not resolved efficiently, while for low values of θ , the search has more difficulties to escape from local optima of the underlying search space.

Throughout the run of the algorithm, the best candidate solution encountered so far, i.e., the DNA word set with the fewest constraint violations, is memorised. Note that even if the algorithm terminates without finding a valid set of size k , a valid subset can always be obtained by iteratively selecting pairs of words that violate a Hamming distance constraint and removing one of the two words

involved in that conflict from the set. Hence, a word set of size k with t constraint violations can always be reduced to a valid set of at least size $k - t$.

Hamming distances between words and/or their reverse complements are not recomputed in each iteration of the algorithm; instead, these are computed once after generating the initial set, and updated after each search step. This can be done very efficiently, since any modification of a single word can only affect the Hamming distances between this word and the $k - 1$ remaining words in the set.

The outer loop of our algorithm can perform multiple independent *runs* of the underlying search process. In conjunction with randomised search initialisation, multiple independent runs can yield better performance; essentially, this is the case if there is a risk for the search process to stagnate, i.e., to get stuck in a region of the underlying search space that it is very unlikely to escape from. In this context, the parameter *maxSteps*, that controls the number of steps after which the search is aborted and possibly restarted from a new initial word set, can have an important impact on the performance of the algorithm; this will become apparent in our experimental results presented in Section 5.

The simple SLS algorithm presented above can easily be enhanced with an additional diversification mechanism that is based on occasional random replacements of small subsets of code words. In this extended algorithm, after performing a search step, i.e., after each iteration of the inner loop of Figure 1, we check whether more than *nsteps* steps have been performed since the last improvement over the best candidate solution found so far. Whenever this is the case, we take it as an indication for search stagnation and perform a diversification step by replacing a fixed fraction *fr* of the code words with words that are newly generated at random (this is done exactly as during search initialisation). After this we continue as in the simple SLS algorithms; we generally make sure that between any two diversification steps at least *nsteps* simple search steps are performed. The values *nsteps* and *fr* are additional parameters of the extended algorithm.

5 Results

To evaluate the performance of our SLS algorithm (and its variants), we performed two types of computational experiments. Detailed analyses of the run-time distributions of our algorithm on individual problem instances were used to study the behavior of the algorithm and the impact of parameter settings. For these empirical analyses, the methodology of [15] for measuring and analyzing run-time distributions (RTDs) of Las Vegas algorithms was used. Run-time was measured in terms of search steps, and absolute CPU time per search step was measured to obtain a cost model of these search steps. The other type of experiment used the optimised parameter settings obtained from the detailed analyses for obtaining DNA word sets of maximal size for various word lengths and combinatorial constraints.

Introducing noise in our algorithmic approach, i.e. using probabilistic moves when taking decisions, provides robustness to our algorithm and allows it to escape from local minima encountered during search. Thorough experimentation




Fig. 2. RTDs for SLS algorithm. Left side: HD and GC constraints (different set sizes $k \in \{100, 120, 140\}$); right side: HD, GC, RC constraints (SLS with and without random replacement).

(not reported here due to restricted space) shows that the setting of the noise parameter, θ , has a substantial effect on the performance of our algorithm: the time required for solving a given problem instance can easily vary more than an order of magnitude depending on the noise parameter setting used. Somewhat surprisingly, empirically optimal settings of θ are consistently close to 0.2 for different problem instances and sizes; consequently, this value was used for all experimental results reported here.

The absolute CPU time for each search step was measured on a PC with 2 1GHz Pentium III CPUs, 512 MB cache and 1GB RAM running Red Hat Linux 7.1 (kernel 2.4.9). The obtained values range between .0006 CPU seconds for **HD** and **GC** constraints and .0079 CPU seconds for **HD** and **RC** constraints, both using set size $k = 100$, word length $n = 8$, and Hamming distance $d = 4$.

5.1 RTD Analysis

To study and characterise the behavior of our algorithm, we measured RTDs from 1000 runs (i.e. $maxTries = 1000$) of the algorithm applied to individual problem instances, using extremely high settings of the cutoff parameter, $maxSteps$, to ensure that a solution was found in each run without using random restarts. For each run, we recorded the number of search steps (i.e. executions of the inner for loop of Figure 1) required for finding a solution. From this data, the RTD gives the probability of success as a function of the number of search steps performed.

As can be seen in Figure 2 (left side), for given constraints, n , and d , the time required for obtaining a word set of size k with a fixed probability p increases with k for all p . Furthermore, for high p , this increase is much more dramatic than for low p , as can be seen in the “fat” right tails of the RTDs. This indicates that our simple SLS algorithm, even when using optimised noise settings, can suffer from stagnation behavior that compromises its performance for long run-times. The easiest way to overcome this stagnation effect is to restart the algorithm after a fixed number $maxSteps$ of search steps. Optimal values of $maxSteps$ can be easily

determined from the empirical RTD data (see [15]). When solving a problem instance for the first time, however, this method is not applicable (unless the RTD can be estimated *a priori*), and the actual performance depends crucially on the value used for *maxSteps*.

It is therefore desirable to find different mechanisms for ameliorating or eliminating the observed search stagnation. The random replacement variant of our simple SLS algorithm described at the end of the previous section was designed to achieve this goal, and in many cases substantially reduces the stagnation effect (Figure 2 shows a typical example), leading to a significant improvement in the robustness and performance of our algorithm. However, as can be shown by comparing the empirical RTDs with appropriately fitted exponential distributions (see Figure 2: right side), even this modification does not completely eliminate the stagnation behavior.

5.2 Quality of Word Sets Obtained by Our Algorithm

Using the improvements and parameter settings obtained from the detailed analysis described above, we used our SLS algorithm to compute word sets for a large number of DNA code design problems. The corresponding results can be summarised as follows (detailed results are tabulated in the Appendix):

HD constraint only. There is a significant body of work on the construction of word sets over an alphabet of size 4 that satisfy the Hamming distance constraint only, with the best bounds summarised by Bogdanova et al. [1]. We used our algorithm to design word sets in which the word lengths ranged from 4 to 12 with the minimum Hamming distance between pairs of words ranging from 4 to 10. Out of 42 tests, our algorithm was able to find a word set whose size matches the best known theoretical construction in 16 cases.

HD+RC constraints. In this case, we considered word lengths ranging from 4 to 12, and the minimum Hamming and reverse complement Hamming distances between pairs of words ranging from 4 to 10. Out of 42 tests, our algorithm was able to find a word set whose size matched or exceeded the theoretical results of Marathe et al. [17] in 41 cases.

HD+GC constraints. We obtained word sets satisfying the Hamming distance and the 50% GC content constraints for words length ranging from 4 up to 20, and the minimum Hamming distance between pair of words ranging from 2 to 10. Compared to the results reported by Corn et al. [16], we obtained bigger sets in almost all cases except two particular combinations, namely word length 8 and Hamming distance 4 and word length 12 and Hamming distance 6.

HD+RC+GC constraints. We obtained word sets for word lengths ranging from 4 to 12, and the minimum Hamming and reverse complement Hamming distances between pairs of words ranging in length from 4 to 10. There was no body of previous results that we could compare with except in one case, namely for words of length 8 with 50% GC content that satisfy the Hamming and reverse-complement constraints with at least 4 mismatches between pairs of words. For this case, Frutos et al. [10] constructed a set of 108 words of length 8. We were not able to obtain any 108 DNA word sets satisfying these constraints, using

our algorithm when initialised with a random initial set. The biggest set found had 92 code words. However, when we initialised our algorithm with the set of 108 words obtained by Frutos et al., along with one additional randomly chosen word, we were able to obtain a valid set of size 109. Continuing this incremental improvement procedure, we were able to construct sets containing 112 words in less than one day of CPU time. Interestingly, this set has the same general “template-map” structure as the set of Frutos et al.

6 Conclusions

We have presented a new stochastic local search algorithm for DNA word design, along with empirical results that characterise its performance and indicate its ability to find high-quality sets of DNA words satisfying various combinations of combinatorial constraints.

In future work, we will examine ways to further improve our algorithm. One possibility is to consider a more diverse neighbourhood instead of the rather small neighbourhood that we now explore (which is based on single base modifications). We conjecture that, particularly when the GC constraint is imposed, the current small neighbourhood structure limits the effectiveness of the algorithm. Another possibility is to consider more complex stochastic local search strategies, which we expect to achieve improved performance that will likely lead to larger word sets.

In another direction of future work, we will test our methods on different word design constraints, including those based on thermodynamic models. While modified or additional constraints can be quite easily accommodated by relatively minor modifications of our current SLS algorithm and its variants, different algorithmic strategies might be more efficient for solving these word design problems.

Finally, it would be interesting to see if better theoretical design principles can be extracted from the word sets that we have now obtained empirically. In the construction of classical codes, theory and experiment are closely linked, and we expect that the same can be true for the construction of DNA codes.

References

1. Galina T. Bogdanova, Andries E. Brouwer, Stoian N. Kapralov & Patric R. J. Ostergard, *Error-Correcting Codes over an Alphabet of Four Elements*, Andries Brouwer (aeb@cwi.nl)
2. A. Ben-Dor, R. Karp, B. Schwikowski, and Z. Yakhini, “Universal DNA tag systems: a combinatorial design scheme,” Proc. RECOMB 2000, ACM, pages 65-75.
3. R.S. Braich, C. Johnson, P.W.K. Rothemund, D. Hwang, N. Chelyapov, and L.M. Adleman, “Solution of a satisfiability problem on a gel-based DNA computer,” Preliminary Proc. Sixth International Meeting on DNA Based Computers, Leiden, The Netherlands, June, 2000.
4. S. Brenner and R. A. Lerner, “Encoded combinatorial chemistry,” Proc. Natl. Acad. Sci. USA, Vol 89, pages 5381-5383, June 1992.

5. R. Deaton, R. C. Murphy, M. Garzon, D. R. Franceschetti, and S. E. Stevens, Jr., "Good encodings for DNA-based solutions to combinatorial problems," Proc. DNA Based Computers II, DIMACS Workshop June 10-12, 1996, L. F. Landweber and E. B. Baum, Editors, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 44, 1999, pages 247-258.
6. R. Deaton, M. Garzon, R. C. Murphy, J. A. Rose, D. R. Franceschetti, and S. E. Stevens, Jr., "Genetic search of reliable encodings for DNA-based computation," Koza, John R., Goldberg, David E., Fogel, David B., and Riolo, Rick L. (editors), Proceedings of the First Annual Conference on Genetic Programming 1996.
7. A. A. El Gamal, L. A. Hemachandra, I. Shperling, and V. K. Wei, "Using simulated annealing to design good codes," IEEE Transactions on Information Theory, Vol. IT-33, No. 1, January 1987.
8. Faulhammer, D., Cukras, A. R., Lipton, R.J., and L. F. Landweber, "Molecular computation: RNA solutions to chess problems," Proc. Natl. Acad. Sci. USA, 97: 1385-1389.
9. U. Feldkamp, W. Banzhaf, H. Rauhe, "A DNA sequence compiler," Poster presented at the 6th International Meeting on DNA Based Computers, Leiden, June, 2000. See also <http://ls11-www.cs.uni-dortmund.de/molcomp/Publications/publications.html> (visited November 11, 2000).
10. A. G. Frutos, Q. Liu, A. J. Thiel, A. M. W. Sanner, A. E. Condon, L. M. Smith, and R. M. Corn, "Demonstration of a word design strategy for DNA computing on surfaces," Nucleic Acids Research, Vol. 25, No. 23, December 1997, pages 4748-4757.
11. M. Garzon, R. J. Deaton, J. A. Rose, and D. R. Franceschetti, "Soft molecular computing," Preliminary Proc. Fifth International Meeting on DNA Based Computers, June 14-15, MIT, 1999, pages 89-98.
12. A.J. Hartemink, D.K. Gifford, and J. Khodor, "Automated constraint-based nucleotide sequence selection for DNA computation," 4th Annual DIMACS Workshop on DNA-Based Computers, Philadelphia, Pennsylvania, June 1998
13. Holger H. Hoos, *Stochastic Local Search - Methods, Models, Applications*, infix-Verlag, Sankt Augustin, Germany, ISBN 3-89601-215-0, 1999.
14. H.H. Hoos and C. Boutilier, "Solving combinatorial auctions using stochastic local search," In *Proceedings of the Seventeenth National Conference on Artificial Intelligence*, pages 22-29, Austin, TX, 2000.
15. H.H. Hoos and T. Stützle, "Evaluating Las Vegas Algorithms — Pitfalls and Remedies," In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, 1998, pages 238-245.
16. M. Li, H-J. Lee, A. E. Condon, and R. M. Corn, "DNA Word Design Strategy for Creating Sets of Non-interacting Oligonucleotides for DNA Microarrays," *Langmuir*, 18, 2002, pages 805-812.
17. A. Marathe, A. Condon, and R. Corn, "On combinatorial DNA word design," *J. Computational Biology*, 8:3, 2001, 201-220.
18. D. McAllester, H. Kautz, and B. Selman, "Evidence for invariants in local search," *Proc. AAAI-97*, 321-326, Providence, RI, 1997.
19. J. A. Rose, R. Deaton, D. R. Franceschetti, M. Garzon, and S. E. Stevens, Jr., "A statistical mechanical treatment of error in the annealing biostep of DNA computation," Special program in DNA and Molecular Computing at the Genetic and Evolutionary Computation Conference (GECCO-99), Orlando, FL., July 13-17, 1999, Morgan Kaufmann.

20. H. Yoshida and A. Suyama, "Solution to 3-SAT by breadth first search," Proc. 5th Intl. Meeting on DNA Based Computers, M.I.T. 1999, pages 9-20.
21. B-T. Zhang and S-Y. Shin, "Molecular algorithms for efficient and reliable DNA computing," Proc. 3rd Annual Genetic Programming Conference, Edited by J. R. Koza, K. Deb, M. Dorigo, D.B. Fogel, M. Garzon, H. Iba, and R. L. Riolo, Morgan Kaufmann, 1998, pages 735-742.

Appendix

In this appendix, we report the sizes of word sets obtained by the SLS algorithm presented in this paper. This data illustrates in detail the performance of our algorithm and will be useful as a baseline for future work. Most of the word sets are publicly available at <http://www.cs.ubc.ca/~dctulpan/papers/dna8/tables>.

The data on our word sets is grouped in four tables showing the results for different constraint combinations as follows: **HD** (Table 1), **HD+GC** (Table 2), **HD+RC** (Table 3) and **HD+RC+GC** constraints (Table 4). In these tables, entries in **bold face** match or exceed the best previously known lower bounds. All other values are smaller than known lower bounds for corresponding (n, d) values. Dashes ("—") indicate cases of limited or no interest for DNA word design, for which we did not run our algorithm.

Along with each maximal word set size found for our algorithm, we also report the corresponding median number of steps spent by our algorithm in order to find that set (in square brackets, [...]), typically specified in thousands of search steps). Each table entry is based on between 5 and 20 runs of our algorithm; the precise number of runs was varied with problem size and solution quality; median run times are taken over successful runs only.

$n \setminus d$	4	5	6	7	8	9	10
4	4 [.01k]	-	-	-	-	-	-
5	16 [1.5k]	4 [.02k]	-	-	-	-	-
6	64 [80k]	9 [15k]	4 [.02k]	-	-	-	-
7	78 [70k]	23 [75k]	8 [.5k]	4 [.04k]	-	-	-
8	219 [65k]	55 [40k]	19 [50k]	5 [.03k]	4 [.09k]	-	-
9	616 [285k]	138 [65k]	41 [50k]	15 [50k]	5 [.1k]	4 [.12k]	-
10	1783 [471k]	358 [312k]	95 [130k]	32 [75k]	12 [5k]	5 [.5k]	4 [.32k]
11	5598 [523k]	967 [592k]	227 [165k + .1k]	70 [36k]	27 [545k]	10 [15k]	4 [.05k]
12	>10000 [24k]	2689 [263k]	578 [300k]	156 [735k]	53 [390k]	23 [532k + .5k]	9 [.5k + 5k]

Table 1. Empirical bounds on **HD** quaternary codes.

<i>4-2</i>	<i>5-2</i>	<i>6-3</i>	<i>7-3</i>	<i>8-4</i>	<i>9-4</i>	<i>10-5</i>	<i>11-5</i>	<i>12-6</i>
48 [.3k]	142 [93k]	85 [860k]	230 [650k]	209 []	400 [72k]	256 [250k]	620 [48k]	410 []
<i>13-6</i>	<i>14-7</i>	<i>15-7</i>	<i>16-8</i>	<i>17-8</i>	<i>18-9</i>	<i>19-9</i>	<i>20-10</i>	
990 [36k]	500 [5k]	1571 [1000k]	966 [242k]	2355 [745k]	1200 [15k]	3451 [396k]	2193 [596k]	-

Table 2. Empirical bounds on **(HD,GC)** quaternary codes for different $n-d$ combinations.

$n \setminus d$	4	5	6	7	8	9	10
4	2 [.02k]	-	-	-	-	-	-
5	4 [.02k]	2 [.005k]	-	-	-	-	-
6	28 [28k]	4 [.1k]	2 [.02k]	-	-	-	-
7	40 [254k]	11 [3k]	2 [.005k]	2 [.05k]	-	-	-
8	112 [850k]	27 [7k]	10 [31k]	2 [.01k]	2 [.03k]	-	-
9	314 [677k]	72 [142k]	20 [37k]	8 [16k]	2 [.02k]	2 [.04k]	-
10	938 [702k]	180 [386k]	49 [287k]	16 [495k]	8 [11.01k]	2 [.01k]	2 [.02k]
11	2750 [117k]	488 [257k]	114 [145k]	35 [6k]	12 [1k]	5 [2k]	2 [.05k]
12	>8000 [72k]	1340 [400k]	290 [327k]	79 [236k]	27 [712.5k]	11 [828k]	4 [.2k]

Table 3. Empirical bounds on **(HD,RC)** quaternary codes.

$n \setminus d$	4	5	6	7	8	9	10
4	2 [.001k]	-	-	-	-	-	-
6	11 [3k]	2 [.003k]	2 [.006k]	-	-	-	-
8	92* [4000k]	19 [1.2k]	7 [10k]	2 [.01k]	2 [.02k]	-	-
10	640 [406k]	127 [170.5k]	37 [38k]	11 [134k]	5 [1.3k]	2 [.02]	1 [.005k]
12	5685 [455k]	933 [531k]	210 [121.5k]	59 [77k]	21 [217k]	9 [341.5k]	3 [1.5k]

Table 4. Empirical bounds on **(HD,RC,GC)** quaternary codes. For the *d entry we found a better bound, namely 112 code words as described in Section 5.2.