

1-1-2012

On the Hardness of Counting and Sampling Center Strings

Christina Boucher

University of California - San Diego

Mohamed Omar

Harvey Mudd College

Recommended Citation

Boucher, Christina and Omar, Mohamed, "On the Hardness of Counting and Sampling Center Strings" (2012). *All HMC Faculty Publications and Research*. Paper 805.

http://scholarship.claremont.edu/hmc_fac_pub/805

This Article - postprint is brought to you for free and open access by the HMC Faculty Scholarship at Scholarship @ Claremont. It has been accepted for inclusion in All HMC Faculty Publications and Research by an authorized administrator of Scholarship @ Claremont. For more information, please contact scholarship@cuc.claremont.edu.

On the Hardness of Counting and Sampling Center Strings

Christina Boucher, and Mohamed Omar

Abstract—Given a set S of n strings, each of length ℓ , and a non-negative value d , we define a *center string* as a string of length ℓ that has Hamming distance at most d from each string in S . The #CLOSEST STRING problem aims to determine the number of unique center strings for a given set of strings S and input parameters n , ℓ , and d . We show #CLOSEST STRING is impossible to solve exactly or even approximately in polynomial time, and that restricting #CLOSEST STRING so that any one of the parameters n , ℓ , or d is fixed leads to an FPRAS. We show equivalent results for the problem of efficiently sampling center strings uniformly at random.

Index Terms—biological sequence analysis, motif recognition, fully polynomial-time randomized approximation scheme (FPRAS), fully polynomial almost uniform sampler (FPAUS), computational complexity.



1 INTRODUCTION

Finding similar regions in multiple DNA, RNA, or protein sequences plays an important role in many applications, including universal PCR primer design [4], [16], [18], [26], genetic probe design [16], antisense drug design [16], [3], finding transcription factor binding sites in genomic data [27], determining an unbiased consensus of a protein family [1], and motif-recognition [16], [24], [25]. The CLOSEST STRING problem formalizes these tasks and can be defined as follows: given a set of n strings S of length ℓ over the alphabet Σ and parameter d , the aim is determine if there exists a string s that has Hamming distance at most d from each string in S . We refer to s as the *center string* and let $d(x, y)$ be the Hamming distance between strings x and y .

The CLOSEST STRING was first introduced and studied in the context bioinformatics by Lanctot *et al.* [16]. Frances and Litman *et al.* [11] showed the problem to be NP-complete, even in the special case when the alphabet is binary, implying there is unlikely to be a polynomial-time algorithm for this problem unless $P = NP$. Since its introduction, the investigation of efficient approximation algorithms and exact heuristics for the CLOSEST STRING problem has been thoroughly considered [9], [10], [12], [16], [17], [19], [20].

S is *pairwise bounded* if the Hamming distance for each pair strings in S is at most $2d$. The CLOSEST STRING problem reduces to separating pairwise bounded sets with a center string, and if so, finding

one such string, from those that do not. A set of strings S with at least one center string is a *motif set*; S is a *decoy set* if it is pairwise bounded but does not have a center string. We note that a center string for a given set S is not necessarily unique.

A related, uninvestigated problem is determining the computational difficulty in finding the number of center strings for a set of strings. In many biological applications, including the ones listed above, it is useful to identify the all possible center strings, rather than only determining whether one exists. Further, an important relationship between the number of unique center strings for a given set of strings S and the computational difficulty of solving the decision version of #CLOSEST STRING for an instance S has been shown. Specifically, empirical analysis demonstrates that for sufficiently large n , all motif sets are clustered together and are characterized as having one unique center string, which is a string of length ℓ containing the symbol that occurs most frequently at each position [2]. Imperative to the analytical explanation of this conjecture is the development of an algorithm to efficiently count the number of center strings for a given set of strings.

We give the formal description of this counting problem as follows:

#CLOSEST STRING

INSTANCE: Parameters n , ℓ and d and a set S of n , length ℓ strings from the alphabet Σ .

OUTPUT: The number of distinct strings taken from the alphabet Σ that have distance at most d from each string in S .

Countless sampling and counting problems have been studied, including the sampling and counting versions of the following problems: matchings in a

- C. Boucher is with the Department of Computer Science and Engineering at the University of California, San Diego.
E-mail: cboucher@eng.ucsd.edu
- M. Omar is with the Department of Mathematics at University of California, Davis.
E-mail: momar@math.ucdavis.edu

graph [14], the graph-colouring [6], [13], [21], Hamiltonian path [7], independent set [8], and knapsack [5], [22].

This paper focuses on the computational difficulty of counting and sampling center strings exactly or approximately. To our knowledge this is the first consideration of this problem but is motivated by problems addressing the analysis and use of biological data. We show #CLOSEST STRING is #P-complete, implying it is in the complexity class of hard counting problems.

Given that this problem cannot be solved efficiently, we investigate if it can be reasonably approximated efficiently. Many #P-complete problems have a *fully-polynomial-time randomized approximation scheme (FPRAS)* which produces with high probability an approximation of arbitrarily small error in time that is polynomial with respect to both the size of the problem and accuracy. Jerrum *et al.* [15] showed that every #P-complete problem either has an FPRAS or is impossible to approximate. For sampling problems, the aim is to obtain a *fully polynomial almost uniform sampler (FPAUS)*, which outputs solutions that achieve an approximation to a given distribution of solutions. In absence of the existence of an FPRAS or FPAUS for a general counting or sampling problem, interest remains in showing an FPRAS or FPAUS exists for a restricted version of the problem (*i.e.*, when one of the problem parameters is fixed). We prove that although there does not exist an FPRAS for the general #CLOSEST STRING problem, the most natural restricted versions of #CLOSEST STRING lead to the existence of an FPRAS. Similarly, we show that although there does not exist an FPAUS for sampling center strings uniformly at random (*u.a.r.*), restricting interest to this sampling problem where one of the parameters is fixed leads to an FPAUS.

2 HARDNESS RESULTS

We show the #CLOSEST STRING problem is #P-complete, and that there does not exist an FPRAS for #CLOSEST STRING under reasonable complexity assumptions. A problem is #P-complete if and only if it is in #P and every problem in #P can be reduced to it by a polynomial-time counting reduction. To prove that #CLOSEST STRING is #P-hard, it is sufficient to show that #3-SAT, a #P-complete problem, can be reduced to #CLOSEST STRING. #3-SAT aims to determine for given a 3-CNF formula F , how many satisfying assignments exist for F . Let $X = \{x_1, \dots, x_n\}$ be a set of Boolean variables. A *literal* is either x_i or $\neg x_i$ for some i . We refer to a *3-clause* as a disjunction of three distinct literals, made of three different variables.

Proposition 1: #CLOSEST STRING is #P-complete.

Proof: We present the reduction of a single 3-clause and then extend it to a general 3-CNF formula. Given a 3-clause ω over the variables in X we define the string $s = s(1) \dots s(2n)00$ of length $2(n+1)$ by:

$$s(2i-1)s(2i) = \begin{cases} 00 & \text{if } \omega \text{ contains the literal } \neg x_i, \\ 11 & \text{if } \omega \text{ contains the literal } x_i, \\ 01 & \text{otherwise} \end{cases}$$

Note that s has exactly four positions i , $1 \leq i \leq n+1$ such that $s(2i-1) = s(2i)$, and the last two positions are equal to 00. Let $\phi: X \rightarrow \{0,1\}^{2(n+1)}$ be the one-to-one mapping from X onto the set of all binary strings of length $2(n+1)$ as defined above (*i.e.* the transformation from X to s given ω).

Let γ_r and γ_r^c be the strings which corresponds to the r -times concatenation of 10 and 01, respectively, *i.e.* $\gamma_r = 1010 \dots 10 \in \{0,1\}^{2r}$ and *i.e.* $\gamma_r^c = 0101 \dots 01 \in \{0,1\}^{2r}$. Further, let $\bar{S}_{r,i}$ be the set of four strings of length $2r$, which includes γ_r and γ_r^c with 01 and 10 replaced at positions $2i-1$ and $2i$. For example, $\bar{S}_{r,2}$ is equal to:

$$\{011001 \dots 01, 100110 \dots 10\}.$$

Lastly, we let $\bar{S}_r = \bar{S}_{r,1} \cup \bar{S}_{r,2} \cup \dots \cup \bar{S}_{r,r} \cup \gamma_r \cup \gamma_r^c$. Note that $|\bar{S}_r| = 2r+2$.

We claim that, for any $r > 2$, any string $s \in \{0,1\}^{2r}$ that has Hamming distance at most r from each string \bar{S}_r satisfies $s \in \{00, 11\}^r$. Assume to the contrary that there exists a string $s \in \{0,1\}^{2r}$ that has Hamming distance at most r from each string in \bar{S}_r , but there exists a position i , $1 \leq i \leq r$, such that $s(2i-1) \neq s(2i)$. Assume without loss of generality that $s(2i-1) = 0$ and $s(2i) = 1$. Since any position in s that matches γ_r mismatches γ_r^c , we have that these two positions in s mismatch γ_r and match γ_r^c . Now, consider the string x in $\bar{S}_{r,i}$ that is equal to γ_r^c , except for the positions $2i-1$ and $2i$. This string has Hamming distance $d+2$ from s , where d is the Hamming distance of γ_r^c from s . Note that all the mismatches counted in d come from positions other than $2i-1$ and $2i$. Since s can have Hamming distance at most r also from γ_r , it follows that $d > r-1$, otherwise, γ_r would have Hamming distance $2r-d > r$. We, therefore, get that the Hamming distance of x and s is $d+2 > r$, contradiction.

Given a 3-CNF formula $\Omega = \omega^1 \wedge \dots \wedge \omega^t$ over the variables x_1, \dots, x_n we consider the set S that includes the strings $\phi(\omega^1), \dots, \phi(\omega^t)$ and the strings in the set \bar{S}_{n+1} . We assume that $n \geq 2$ (otherwise, the problem can be solved in polynomial time). Hence,

$$S = \underbrace{\{\phi(\omega^1), \dots, \phi(\omega^t), \bar{S}_{n+1}\}}_{2(n+1)+2+t}.$$

We let $d_c = n+1$. We observe that this reduction can be performed in $O(n+t)$ time.

We now prove the following statements are identical:

- 1) There is an assignment to x_1, \dots, x_n that satisfies the clauses in Ω .

- 2) $\{S, d_c\}$ has a solution s which has Hamming distance at most $n + 1$ from the strings in S .

Let $A(x_1, \dots, x_n) \in \{0, 1\}^n$ be an assignment satisfying the clauses of Ω , and define the string $s_A = s_A(1) \dots s_A(2n)$ as follows:

$$s_A(2i-1)s_A(2i) = \begin{cases} 00 & \text{if } A(x_i) = 0, \\ 11 & \text{if } A(x_i) = 1, \end{cases}$$

Note that the string $s_A \cdot 00$ has Hamming distance at most $n + 1$ from each of the strings in S .

Next, let s be a string that has Hamming distance at most $n + 1$ from each of the strings in S . Note that since s has Hamming distance at most $n + 1$ from all the $2(n+1) + 2$ strings in \bar{S}_{n+1} we know that $s \in \{00, 11\}^r$ and $s(2n+1)s(2n+2) = 00$. Consider the assignment A_s to x_1, x_2, \dots, x_n corresponding to the string s , defined as follows: $A_s(x_i) = 0$ if $s(i)s(i+1) = 00$; otherwise, $A_s(x_i) = 1$. Observe that if s has Hamming distance at most $n + 1$ from a string in the set $\phi(\omega^j)$ then A_s satisfies clause ω^j . Hence, the number of satisfying formulae to Ω is parsimonious to the number of center strings to the set S . \square

Randomization can be quite powerful in achieving an approximation scheme for several #P-complete problems. Jerrum *et al.* show that the problem of counting the number of simple cycles in a directed graph is not approximable by proving that the existence of an almost uniform generator for this problem, implies the existence of a randomized polynomial time algorithm for determining the existence of a Hamiltonian cycle in a directed graph [15]. FPRAS is the subclass of #P counting problems whose answer, y , is approximable in the following sense: there exists a randomized algorithm that, with probability at least $1 - \delta$, approximates y to within an ϵ multiplicative factor in time polynomial in n (the input size), $1/\epsilon$, and $\log(1/\delta)$.

The results of Jerrum *et al.* [15] imply that every #P-complete problem exhibits an FPRAS or is not approximable. Given a #P-complete problem an important question to answer is if there exists a FPRAS for the problem – since the existence implies the approximability of the problem. Unfortunately, the existence of a FPRAS for #CLOSEST STRING is unlikely.

Observation 1: There is no FPRAS for #CLOSEST STRING, unless NP = RP.

Consider a problem π and I be an instance of the problem π , and let $\#(I)$ denote the number of solutions for I . an FPRAS can be used to distinguish between the case where $\#(I)$ is equal to zero and when $\#(I)$ is greater than zero; hence, providing a randomized polynomial time algorithm for the decision version of the problem π . Therefore, π must be contained in the class BPP. Since it is unlikely that BPP equal to NP, all NP-complete problems are believed not to contain an FPRAS [23, page 309]. Since CLOSEST

STRING problem is NP-complete [11], there exists no FPRAS for #CLOSEST STRING, unless BPP = NP.

The notions of counting and sampling are closely related. Jerrum *et al.* established the equivalence between the existence of an FPRAS and an FPAUS; namely for self-reducible problems there exists an FPRAS if and only if there exists an FPAUS [15]. It follows that we have the following negative result concerning the approximability of sampling center strings

Observation 2: There is no FPAUS for sampling center string uniformly at random, unless NP = RP.

3 COUNTING AND SAMPLING WITH FIXED PARAMETERS

Observation 2 is evidence that determining the number of center strings is computationally difficult to approximate and therefore, to achieve progress on the existence of an FPRAS we consider the problem where one of the parameters is fixed. We first determine if the decision problem corresponding to the restricted version of #CLOSEST STRING can be solved in polynomial-time, since otherwise we could show there does not exist an FPRAS by using similar argument to that for Observation 2. In order for the existence of an FPRAS to be possible for some restricted version of the #CLOSEST STRING problem, the corresponding decision problem has to be *fixed parameter tractable (FPT)*, meaning there exists an algorithm that is exponential only in the size of a fixed parameter while polynomial in the remaining, unfixed parameters.

CLOSEST STRING is trivially FPT when the parameter ℓ is fixed since the enumeration algorithm that tries all possible length ℓ strings is polynomial-time when ℓ is fixed. Gramm *et al.* prove CLOSEST STRING is FPT when n or d is fixed [12]. These results imply that restricting #CLOSEST STRING such that at least one of ℓ , d or n is fixed leads to a problem that may have an FPRAS.

When ℓ is fixed the enumeration algorithm that attempts all $|\Sigma|^\ell$ strings is an FPRAS. The $O(d(|\Sigma|^\ell \epsilon)^d)$ algorithm that determines which strings from the set of strings that have distance at most d from s_1 proves the existence of an FPRAS when d is fixed. These enumeration algorithms prove the existence of a FPAUS for the problem of sampling center strings u.a.r. when ℓ or d is fixed. Next we show there exists both an FPRAS and an FPAUS for the respective problems of counting and sampling center strings when the number of strings is fixed. In fact, we show a stronger result: that there exists an exact polynomial-time algorithm for counting and sampling center strings when n is fixed.

Proposition 2: When the number of strings is fixed and Σ is the binary alphabet there exists a polynomial-time algorithm for #CLOSEST STRING and for sampling center strings u.a.r.

Proof: The goal is to give an integer linear programming (ILP) formulation such that the number of variables depends only on the value of n . Let $S = \{s_1, \dots, s_n\}$ denote a set of n binary strings, each of length ℓ , and denote each string s_j as $s_j(1) \cdots s_j(\ell)$. Let \mathcal{C}_S denote the set of center strings for the set S . Given a set of n strings of length ℓ , we can think of these strings as a $n \times \ell$ matrix. We refer to the columns of an instance of #CLOSEST STRING as the columns of a matrix. There are 2^n possible number of unique columns. Using the column types, we show how #CLOSEST STRING restricted to the binary alphabet can be formulated as an ILP with 2^n variables.

Let $\vec{b} = [b_1, \dots, b_n]^T$ correspond to one particular column type, and let $\mathcal{P}(\vec{b}) = \{i \mid (s_i(1), s_i(2), \dots, s_i(n)) = \vec{b}\}$ (i.e. $\mathcal{P}(\vec{b})$ is the set of positions in S which are equal to \vec{b}). Therefore, $|\mathcal{P}(\vec{b})|$ is equal to the number of positions in S that are equal to \vec{b} .

For a string $u = u(1), \dots, u(\ell)$, let $\rho(\vec{b})$ be the number of positions that are equal to \vec{b} where u is equal to 0 (i.e. $j \in \mathcal{P}(\vec{b})$ and $u(j) = 0$). Hence, $|\mathcal{P}(\vec{b})| - \rho(\vec{b})$ is the number of positions that are equal to \vec{b}

Therefore, u has distance at most d from s_i if and only if

$$\sum_{\vec{b}} b_i \rho(\vec{b}) + (1 - b_i)(|\mathcal{P}(\vec{b})| - \rho(\vec{b})) \leq d.$$

Thus, \mathcal{C}_S is nonempty if and only if there is a feasible integer solution to

$$\sum_{\vec{b}} (2b_i - 1)\rho(\vec{b}) + (1 - b_i)|\mathcal{P}(\vec{b})| \leq d \quad (1 \leq i \leq n) \quad (1)$$

$$0 \leq \rho(\vec{b}) \leq |\mathcal{P}(\vec{b})| \quad (2)$$

for the variables $\rho(\vec{b})$.

Assuming there is a solution, we know the number of strings u corresponding to each such solution. It is exactly

$$\prod_{\vec{b}} \binom{|\mathcal{P}(\vec{b})|}{\rho(\vec{b})} \quad (3)$$

where the product is over all feasible values of the variables $\rho(\vec{b})$.

To sample the strings in \mathcal{C}_S

(i) generate random values of each of the numbers $\rho(\vec{b})$, for each \vec{b} , with the correct probabilities — proportional to the formula in (3)

(ii) sample exactly uniformly from the vectors corresponding to the given $\rho(\vec{b})$ by choosing subsets of given size uniformly at random.

This immediately shows that for fixed n there is a polynomial-time algorithm for perfect sampling, since one can run through all possible values of the set of variables $\rho(\vec{b})$ (there are at most ℓ values for each of these, hence at most ℓ^{2^n} values in total) and

for each one, compute the value of (3). Then choose between one of these polynomially many values with the required probability, and then perform step (ii). \square

A slight modification of the proof for the previous proposition leads to the following stronger result that eliminates the requirement that the alphabet is binary. See the Appendix for the details of the proof.

Proposition 3: When the number of strings is fixed there exists a polynomial-time algorithm for #CLOSEST STRING and for sampling center strings u.a.r.

4 CONCLUSION

Counting and sampling from a specific distribution is a well-studied area in discrete mathematics and theoretical computer science that has been useful for the study of combinatorial problems. The problem of counting and sampling center strings has a natural application to several bioinformatic problems, including motif-recognition. We prove #CLOSEST STRING is #P-complete and does not exist a FPRAS, the problem of sampling center strings u.a.r. does not have a FPAUS, and any natural restriction of these counting and sampling problems yields an FPRAS and FPAUS, respectively. This work suggests some open areas of study, including developing a more efficient sampling and counting algorithms, investigating the existence of a rapidly mixing chain for more restricted sampling problems, or proving hardness results that show the non-existence of a rapidly mixing chain when a single parameter is fixed.

ACKNOWLEDGMENT

The authors would like to thank Nick Wormald for his discussions and insights concerning the results presented in this paper. We gratefully acknowledge research support of the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- [1] A. Ben-Dor, G. Lancia, J. Perone, and R. Ravi. *Banishing bias from consensus strings*. In proceedings of 8th CPM, pages 247–261, 1997.
- [2] C. Boucher and D.G. Brown. *Detecting motifs in a large data set: applying probabilistic insights to motif finding*. In proceedings of 1st BCoB, pages 139–150, 2008.
- [3] X. Deng, G. Li, Z. Li, B. Ma, and L. Wang. *Genetic design of drugs without side-effects*. SIAM Journal on Computing, 32(4):1073–1090, 2003.
- [4] J. Dopazo, A. Rodríguez, J.C. Sáiz, and F. Sobrino. *Design of primers for PCR amplification of highly variable genomes*. Computer Applications in the Biosciences, 9:123–125, 1993.
- [5] M. Dyer. *Approximate counting by dynamic programming*. In proceedings of 35th STOC, pages 693–699, 2003.
- [6] M. Dyer and A. Frieze. *Randomly colouring graphs with lower bounds on girth and maximum degree*. In Proceedings of 42nd FOCS, pages 579–587, 2001.
- [7] M. Dyer, A. Frieze, and M. Jerrum. *Approximately counting hamilton paths and cycles in dense graphs*. SIAM Journal on Computing, 27(5):1262–1272, 1998.
- [8] M. Dyer, A. Frieze, and M. Jerrum. *On counting independent sets in sparse graphs*. SIAM Journal on Computing, 31(5):1527–1541, 2002.

- [9] M.R. Fellows, J. Gramm, and R. Niedermeier. *On the parameterized intractability of closest substring and related problems*. In proceedings of 19th STACS, pages 262–273, 2002.
- [10] M.R. Fellows, J. Gramm, and R. Niedermeier. *On the parameterized intractability of motif search problems*. *Combinatorica*, 26:141–167, 2006.
- [11] M. Frances and A. Litman. *On covering problems of codes*. *Theoretical Computer Science*, 30(2):113–119, 1997.
- [12] J. Gramm, R. Niedermeier, and P. Rossmanith. *Fixed-parameter algorithms for closest string and related problems*. *Algorithmica*, 37(1):25–42, 2003.
- [13] T.P. Hayes and E. Vigoda. *A non-markovian coupling for randomly sampling colorings*. In proceedings of 44th FOCS, pages 618–627, 2003.
- [14] M.R. Jerrum and A. Sinclair. *Approximating the permanent*. *SIAM Journal on Computing*, 18(6):1149–1178, 1989.
- [15] M.R. Jerrum, L.G. Valiant, and V. Vazirani. *Random generation of combinatorial structures from a uniform distribution*. *Theoretical Computer Science*, 43, 1986.
- [16] J.K. Lanctot, M. Li, B. Ma, S. Wang, and L. Zhang. *Distinguishing string selection problems*. *Information and Computation*, pages 41–55, 2003.
- [17] M. Li, B. Ma, and L. Wang. *Finding similar regions in many strings*. *Journal of Computer and System Sciences*, 65(1):73–96, 2002.
- [18] K. Lucas, M. Busch, S. Össinger, and J.A. Thompson. *An improved microcomputer program for finding gene- and gene family-specific oligonucleotides suitable as primers for polymerase chain reactions or as probes*. *Computer Applications in the Biosciences*, 7:525–529, 1991.
- [19] B. Ma. *A polynomial time approximation scheme for the closest substring problem*. In proceedings of 11th CPM, pages 99–107, 2000.
- [20] B. Ma and X. Sun. *More efficient algorithms for closest string and substring problems*. In proceedings of 12th RECOMB, pages 396–409, 2008.
- [21] M. Molloy. *The glauher dynamics on colorings of a graph with high girth and maximum degree*. In proceedings of 36th STOC, pages 91–98, 2002.
- [22] B. Morris and A. Sinclair. *Random walks on truncated cubes and sampling 0-1 knapsack solutions*. In proceedings of 40th FOCS, pages 230–240, 1999.
- [23] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [24] G. Pavesi, G. Mauri, and G. Pesole. *An algorithm for finding signals of unknown length in DNA sequences*. *Bioinformatics*, 17:S207–S214, 2001.
- [25] P. Pevzner and S. Sze. *Combinatorial approaches to finding subtle signals in DNA strings*. In proceedings of 8th ISMB, pages 269–278, 2000.
- [26] V. Proutski and E.C. Holme. *Primer master: A new program for the design and analysis of PCR primers*. *Computer Applications in the Biosciences*, 12:253–255, 1996.
- [27] M. Tompa, N. Li, T. L. Bailey, G. M. Church, B. De Moor, E. Eskin, A. V. Favorov, M. C. Frith, Y. Fu, W. J. Kent, et al. *Assessing computational tools for the discovery of transcription factor binding sites*. *Nature Biotechnology*, 23(1):137–144, 2005.

APPENDIX

Proposition 3 When the number of strings is fixed there exists a polynomial-time algorithm for #CLOSEST STRING and for sampling center strings u.a.r.

Proof: The goal is to give an ILP formulation such that the number of variables depends only on the value of n . Let $S = \{s_1, \dots, s_n\}$ denote a set of n strings from the alphabet Σ , each of length ℓ , and denote each string s_j as $s_j(1) \cdots s_j(\ell)$. Let \mathcal{C}_S denote the set of center strings for the set S . Let α be a letter in Σ . The n th Bell number is the number of partitions of a set of size n . Without loss of generality, we assume that the first string is equal to α^ℓ , since any set of

strings can be trivially converted to an equivalent set where this is true. Using the same terminology defined in the proof of Proposition 2, there exists at most $B_n \leq n!$ unique column types, where B_n is n th Bell number.

Let $\vec{b} = [b_1, \dots, b_n]^T$ correspond to one particular column type, and let $\mathcal{P}(\vec{b}) = \{i \mid (s_i(1), s_i(2), \dots, s_i(n)) = \vec{b}\}$. Let $|\mathcal{P}(\vec{b})|$ be equal to the number of positions in S that are equal to \vec{b} . For a string $u = u(1), \dots, u(\ell)$, let $\rho(\vec{b}, \nu)$ be the number of positions that are equal to \vec{b} where u is equal to ν and $\nu \in \Sigma$.

Hence, \mathcal{C}_S is nonempty if and only if there is a feasible integer solution to

$$\sum_{\vec{b}} \sum_{\nu \in (\Sigma - \nu(\vec{b}, i))} \rho(\vec{b}, \nu) \leq d \quad (1 \leq i \leq n) \quad (4)$$

$$0 \leq \rho(\vec{b}, \nu) \leq \mathcal{P}(\vec{b}) \quad (5)$$

for the variables $\rho(\vec{b}, \nu)$, where $\nu(\vec{b}, i)$ is symbol of string i at column \vec{b} . Sampling and counting the solutions to this ILP can be done equivalently to sampling and counting the solutions to the ILP given in the proof of Proposition 2. \square

Christina Boucher Christina Boucher received her PhD degree from the University of Waterloo in 2010. She started working in the field of bioinformatics as a research assistant at the Montreal General Hospital where she helped analyze HIV-1 protease inhibitor drugs. She is now a NSERC postdoctoral fellow at the University of California, San Diego in Pavel Pevzner's lab working on DNA assembly and mass spectrometry.

Mohamed Omar Mohamed Omar is completing his PhD degree from the University of California, Davis, partially supported by an NSERC postgraduate fellowship. He will be a Harry Bateman Instructor (postdoctoral fellow) in the Department of Mathematics at the California Institute of Technology beginning September, 2011, working on algebraic graph theory, and algebraic geometry applied to combinatorics and optimization.