

Logarithmically larger deletion codes of all distances

Noga Alon ^{*} Gabriela Bourla [†] Ben Graham [‡] Xiaoyu He [§] Noah Kravitz [¶]

August 10, 2023

Abstract

The deletion distance between two binary words $u, v \in \{0, 1\}^n$ is the smallest k such that u and v share a common subsequence of length $n - k$. A set C of binary words of length n is called a k -deletion code if every pair of distinct words in C has deletion distance greater than k . In 1965, Levenshtein initiated the study of deletion codes by showing that, for $k \geq 1$ fixed and n going to infinity, a k -deletion code $C \subseteq \{0, 1\}^n$ of maximum size satisfies $\Omega_k(2^n/n^{2k}) \leq |C| \leq O_k(2^n/n^k)$. We make the first asymptotic improvement to these bounds by showing that there exist k -deletion codes with size at least $\Omega_k(2^n \log n/n^{2k})$. Our proof is inspired by Jiang and Vardy’s improvement to the classical Gilbert–Varshamov bounds. We also establish several related results on the number of longest common subsequences and shortest common supersequences of a pair of words with given length and deletion distance.

1 Introduction

The main goal of coding theory is to construct schemes for efficiently and faithfully communicating messages across a noisy channel. In this paper, we study a noise model proposed by Levenshtein [8] in which messages are finite binary words in $\{0, 1\}^n$ and the communication channel, a “deletion channel,” deletes a fixed number k of bits from the transmitted message; the locations of the deletions are unknown to the receiver. Deletion errors are a special case of “synchronization errors”, which are remarkably poorly understood compared to the better-studied noise models of bit flips and bit erasures.

Formally, for $n \geq k \geq 1$ a k -deletion code of length n is a collection $C \subseteq \{0, 1\}^n$ of binary words with the property that for any $y \in \{0, 1\}^{n-k}$, there is at most one $x \in C$ containing y as a subsequence. Equivalently, C is a k -deletion code if for all distinct $s, t \in C$, the longest common subsequence of s and t has length strictly smaller than $n - k$. We would like to determine the maximum size $D(n, k)$ of a k -deletion code of length n . In his seminal 1965 paper [8], Levenshtein established the upper and lower bounds

$$\Omega_k \left(\frac{2^n}{n^{2k}} \right) \leq D(n, k) \leq O_k \left(\frac{2^n}{n^k} \right). \quad (1)$$

For the case $k = 1$, Levenshtein used a construction of Varshamov and Tenengolts [11] to show that $D(n, 1) = \Theta(2^n/n)$, so the upper bound in (1) is asymptotically correct in this case. In contrast, despite a

^{*}Department of Mathematics, Princeton University, Princeton, NJ 08544, USA and Schools of Mathematical Sciences and Computer Science, Tel Aviv University, Tel Aviv, Israel. Email: nalon@math.princeton.edu. Research supported in part by NSF grant DMS-2154082

[†]Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: gboula@princeton.edu. Research supported by the math department’s undergraduate funding.

[‡]Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: bagraham@princeton.edu. Research supported by the math department’s undergraduate funding.

[§]Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: xiaoyuh@princeton.edu. Research supported by NSF Award DMS-2103154.

[¶]Department of Mathematics, Princeton University, Princeton, NJ 08544. Email: nkravitz@princeton.edu. Research supported by NSF GRFP Award DGE-2039656.

great deal of effort on many related questions in recent years, neither bound in (1) has been improved for any fixed $k \geq 2$. Since the breakthrough in [5], there has been some progress on constructing explicit 2-deletion codes that nearly match Levenshtein’s lower bound; see also [3, 4, 6, 10].

Our main result is a logarithmic improvement on the lower bound, which holds for all alphabet sizes. Write $D_\alpha(n, k)$ for the maximum size of a k -deletion code of length n over the fixed alphabet $[\alpha]$.¹

Theorem 1. *If $n \geq k \geq 2$ and $\alpha \geq 2$, then $D_\alpha(n, k) \geq \Omega_{\alpha, k}(\alpha^n \log n / n^{2k})$.*

Our proof is nonconstructive: We reduce the problem of finding large codes to the problem of finding a large independent set in the associated k -deletion graph $\Gamma_{n, k, \alpha}$. The graph $\Gamma_{n, k, \alpha}$ has vertex set $[\alpha]^n$, and two words are connected by an edge if they have a common subsequence of length at least $n - k$. We show that $\Gamma_{n, k, \alpha}$ is locally sparse, that is, contains few triangles. Theorem 1 then follows from standard lemmas about the independence number of locally sparse graphs. A similar application of local sparsity to coding theory appears in the work of Jiang and Vardy [7], who obtained the first asymptotic improvements on the Gilbert–Varshamov bounds.

Counting triangles is harder for the graph $\Gamma_{n, k, \alpha}$ than it is for the extremely symmetric setting studied by Jiang and Vardy, where the analogous graph is just a power of the Hamming cube. In contrast, $\Gamma_{n, k, \alpha}$ is not even regular. In order to overcome these difficulties, we restrict our attention to “pseudorandom” words in the graph that are related by “pseudorandom” sequences of insertion and deletion operations, for suitable notions of pseudorandomness.

In this work we also prove additional results about the number of longest common subsequences and shortest common supersequences of a pair of words, as a function of their lengths and deletion distance. These bounds, which were necessary in earlier versions of our proof of Theorem 1, are of independent interest and may be useful for future study of deletion codes and the structure of the graphs $\Gamma_{n, k, \alpha}$.

We denote the length of a word u by $|u|$. We say that the word w is a *subsequence* of the word u if w can be obtained from u by deleting some of the letters of u . If w is a subsequence of the words u and v , then we say that w is a *common subsequence* of u and v ; further, w is a *longest common subsequence* (or *LCS*) of u and v if it is a common subsequence of maximum length. We let $\text{LCS}(u, v)$ denote the length of an LCS of u and v . If u and v are words of the same length $|u| = |v| = n$, then we define the *deletion distance* between u and v to be $d(u, v) := n - \text{LCS}(u, v)$. One can define *shortest common supersequences* (or *SCS*’s) and the *insertion distance* analogously. It is well known that $\text{LCS}(u, v) + \text{SCS}(u, v) = |u| + |v|$ for all words u, v , so, in particular, deletion distance and insertion distance are identical; LCS’s and SCS’s are in this sense dual.

For words u and v , we define the *LCS multiplicity* $m_{\text{LCS}}(u, v)$ (respectively, *SCS multiplicity* $m_{\text{SCS}}(u, v)$) to be the number of distinct LCS’s (respectively, SCS’s) of u and v . The following simple inequality relating LCS and SCS multiplicity is probably known to experts, but we could not locate a reference in the literature.

Proposition 2. *For all words u, v , we have $m_{\text{LCS}}(u, v) \leq m_{\text{SCS}}(u, v)$.*

Our main result on LCS and SCS multiplicity is the following.

Theorem 3. *Let n, a, b be natural numbers with $n \geq a + b$. If u and v are words with lengths $n - a$ and $n - b$ (respectively) and $\text{SCS}(u, v) = n$ (equivalently, $\text{LCS}(u, v) = n - a - b$), then*

$$m_{\text{LCS}}(u, v) \leq m_{\text{SCS}}(u, v) \leq \binom{a + b}{a}.$$

The $a = b$ case can be phrased symmetrically as follows: If u, v are words of equal length with $d(u, v) = d$, then we have

$$m_{\text{LCS}}(u, v) \leq m_{\text{SCS}}(u, v) \leq \binom{2d}{d},$$

¹When $\alpha = 2$, we will sometimes instead work over the “usual” binary alphabet $\{0, 1\}$; this will not cause any confusion.

independent of the lengths of u, v . We also prove in the appendix that this theorem is tight in that for all choices of a and b and all sufficiently large n (in terms of a, b), there exists a pair of words u, v for which equality is attained in both inequalities.

The paper is organized as follows. We prove the main result Theorem 1 in Section 2; we prove Proposition 2 and Theorem 3 in Section 3; finally, we describe a family of pairs of words which attain equality in Theorem 3 in the appendix.

We use standard asymptotic notation, as follows. If $f(n), g(n) : \mathbb{N} \rightarrow \mathbb{R}$ are functions, then we write $f = O(g)$ to indicate that there is some constant $C > 0$ such that $|f(n)| \leq Cg(n)$ for all natural numbers n . If g is nonnegative, then we write $f = \Omega(g)$ to indicate that $g = O(f)$. We write $f = \Theta(g)$ if $f = O(g)$ and $g = O(f)$. Subscripts on O, Ω, Θ indicate that the implied constants C may depend on the subscripted parameters. All logarithms are base-2.

2 Proof of Theorem 1

In this section we prove Theorem 1 in two steps: We reduce the problem to counting triangles in the k -deletion graph $\Gamma_{n,k,\alpha}$, and then we approximate this triangle count. Observe that $D_\alpha(n, k)$ is by definition the independence number of $\Gamma_{n,k,\alpha}$. We need the following standard lemma of Ajtai, Komlós, and Szemerédi [1] on independence numbers of graphs with few triangles. See also Shearer [9] for a simpler argument and [2, pp. 336-337] for a very short proof. This line of work has led to several important developments in extremal graph theory and Ramsey theory.

Lemma 4 ([1], Lemma 5). *For any $\varepsilon > 0$ and any graph G on $N \geq 1$ vertices with average degree d containing $T < Nd^{2-\varepsilon}$ triangles, we have $\alpha(G) \geq \Omega_\varepsilon((N/d) \log d)$.*

It follows easily that we can replace average degree d by maximum degree Δ in the lemma above, and this is the form we will use. The graph $\Gamma_{n,k,\alpha}$ has $N = [\alpha]^n$ vertices and maximum degree $\Delta = O_{\alpha,k}(n^{2k})$, since from any given vertex $u \in [\alpha]^n$, a neighbor v can be obtained by choosing k letters of u to delete in at most $\binom{n}{k}$ ways and then k letters to insert in at most $\binom{n}{k}\alpha^k$ ways. Thus, if we want to use Lemma 4 to prove that $D_\alpha(n, k) = \Omega_{\alpha,k}(\alpha^n \log n/n^{2k})$, it suffices to show that the number of triangles in $\Gamma_{n,k,\alpha}$ is $O_{\alpha,k}(\alpha^n n^{4k-\varepsilon})$ for some $\varepsilon > 0$.

We will actually prove the sharper bound that $\Gamma_{n,k,\alpha}$ has $O_{\alpha,k}(\alpha^n n^{3k}(\log n)^k)$ triangles. This estimate is tight up to the logarithmic factor. It will be convenient to focus our attention on “pseudorandom” words, as follows. If $u \in [\alpha]^n$ is a word of length n and $S \subseteq [n]$ is a subset, let u_S denote the subword of u indexed by S . If $I = [x, y]$ is an interval, then we call u_I a *subinterval* of u . For $1 \leq \lambda \leq n$, we say that $u \in [\alpha]^n$ is λ -*nonrepeating* if $u_I \neq u_J$ for all pairs of distinct intervals $I, J \subseteq [n]$ of length λ ; u is λ -*repeating* otherwise. By the first-moment method, if $\lambda > (2 + \varepsilon) \log n$ for some $\varepsilon > 0$, then almost all words of length n are λ -nonrepeating (see the proof of Lemma 7 for the formal proof of this fact).

Next we introduce notation for a sequence of insertion and deletion operations. Let $u \in [\alpha]^n$, $t \in \{\text{del}, \text{ins}_1, \text{ins}_2, \dots, \text{ins}_\alpha\}$ and $i \in [0, n]$, where i is not allowed to be 0 if $t = \text{del}$ (since the 0-th letter of u , which does not exist, cannot be deleted). We write $f_{i,t}(u)$ for the word obtained from u by deleting u_i , if $t = \text{del}$, inserting an x after u_i , if $t = \text{ins}_x$. Here, “inserting after u_0 ” means inserting before u_1 .

Definition 5. Fix nonnegative integers n and ℓ . Let $I = (i_\ell, i_{\ell-1}, \dots, i_1)$ be a nonincreasing sequence of nonnegative integers $n \geq i_\ell \geq i_{\ell-1} \geq \dots \geq i_1 \geq 0$, and let $T = (t_\ell, \dots, t_1) \in \{\text{del}, \text{ins}_1, \text{ins}_2, \dots, \text{ins}_\alpha\}^\ell$ be a sequence of insertion/deletion types. We further require that if $t_j = \text{del}$ then $i_j \neq 0$ (the 0-th letter of u cannot be deleted) and $i_{j-1} < i_j$ (we do not operate on an already-deleted letter). We then call the pair (I, T) a *sequence of ℓ insertions and deletions*, and we write

$$f_{I,T}(u) := (f_{i_1,t_1} \circ f_{i_2,t_2} \circ \dots \circ f_{i_\ell,t_\ell})(u)$$

for the composition of the operations f_{i_ℓ,t_ℓ} through f_{i_1,t_1} applied to a word $u \in [\alpha]^n$.

Whenever one obtains a word v from u by inserting and deleting letters, one can reorder these operations to find a sequence (I, T) of insertions and deletions such that $v = f_{I, T}(u)$. Note that, because the elements of I are nonincreasing, an earlier operation cannot shift the location of a later operation. In particular, i_j is not only the position in $(f_{i_{j+1}, t_{j+1}} \circ \dots \circ f_{i_\ell, t_\ell})(u)$ at which the operation f_{i_j, t_j} is applied, but also the original position in u at which the operation occurs. This lets us refer unambiguously to the “position” i_j in u of each operation f_{i_j, t_j} .

We say that an element i of a set $I \subseteq [0, n]$ is λ -isolated if $\lambda < i < n - \lambda$ and no other element $j \in I$ satisfies $|j - i| \leq 2\lambda$. We are now ready to prove our key lemma.

Lemma 6. *Let $n, k, \lambda \geq 1$, and let $u, v \in [\alpha]^n$ be λ -nonrepeating words such that $v = f_{I, T}(u)$ for some sequence of operations (I, T) . If the number of λ -isolated elements of I is at least $2k + 1$, then $d(u, v) > k$.*

Proof. We may pick $2k + 1$ of the λ -isolated terms of I and call them $j_{2k+1} > j_{2k} > \dots > j_1$; let the corresponding terms of T be t_{2k+1}, \dots, t_1 . Note that since the operations of (I, T) are applied in decreasing order of index, these operations f_{j_s, t_s} are applied in decreasing order of s as well. For each $\lambda < j < n - \lambda$, write $L(j) := [j - \lambda, j - 1]$ and $R(j) := [j + 1, j + \lambda]$ for the length- λ intervals in $[n]$ immediately to the left and right of j . The definition of λ -isolation implies that the $4k + 2$ intervals

$$L(j_1), R(j_1), L(j_2), R(j_2), \dots, L(j_{2k+1}), R(j_{2k+1})$$

are pairwise disjoint intervals of length λ . Moreover, no insertion or deletion operations occur in any of the corresponding $4k + 2$ subintervals of u (since the λ -isolation assumption ensures that the j_i 's are far apart). Since u and v are λ -nonrepeating, each of these $4k + 2$ words appears exactly once as a subinterval of u and once as a subinterval of v .

The key observation is that when we apply f_{j_s, t_s} , we either insert or delete a single letter between $u_{L(j_s)}$ and $u_{R(j_s)}$. Inside u , these two subintervals appear with exactly one letter between them, and no other insertion or deletion operations happen nearby. Thus, $u_{L(j_s)}$ and $u_{R(j_s)}$ appear in v , and the number of letters in v between the unique appearances of $u_{L(j_s)}$ and $u_{R(j_s)}$ is either 0 (if a letter was deleted by f_{j_s, t_s}) or 2 (if a letter was inserted by f_{j_s, t_s}).

Assume for the sake of contradiction that $d(u, v) \leq k$, and let (I', T') be a sequence of at most k insertions and at most k deletions such that $v = f_{I', T'}(u)$. Since $|I'| \leq 2k$, there exists some $1 \leq s \leq 2k + 1$ for which I' is disjoint from the entire length- $(2\lambda + 1)$ subinterval $[j_s - \lambda, j_s + \lambda]$. It follows that $u_{[j_s - \lambda, j_s + \lambda]}$ appears unaltered as a subinterval of v , and in particular the unique copy of $u_{L(j_s)}$ in v and the unique copy of $u_{R(j_s)}$ in v have exactly one letter between them. This contradicts the key observation in the previous paragraph, so we conclude that $d(u, v) > k$, as desired. \square

The next lemma lets us upper-bound the number of triangles in $\Gamma_{n, k, \alpha}$ and, more generally, the number of triples $(u, v, w) \in ([\alpha]^n)^3$ with prescribed values of $d(u, v), d(v, w), d(w, u)$.

Lemma 7. *Let $n \geq a \geq b \geq c \geq 1$. The number of triples $(u, v, w) \in ([\alpha]^n)^3$ with $d(u, v) \leq a$, $d(v, w) \leq b$, and $d(w, u) \leq c$ is $O_{a, \alpha}(\alpha^n n^{a+b+c} (\log n)^{b+c-a})$.*

Proof. Say that a triple $(u, v, w) \in ([\alpha]^n)^3$ is *good* if $d(u, v) \leq a$, $d(v, w) \leq b$, and $d(w, u) \leq c$. Note that $d(u, v) \leq d(v, w) + d(w, u) \leq b + c$ by the Triangle Inequality, so all good triples (u, v, w) satisfy $d(u, v) \leq b + c$, and we may restrict our attention to the regime $a \leq b + c$.

Let $\lambda = 10a \log n$, and observe that the probability of a uniformly random $u \in [\alpha]^n$ being λ -repeating is at most $\binom{n}{2} n^{-10a} \leq n^{-8a}$. Thus, the total number of such exceptional words is at most $\alpha^n n^{-8a}$. For each $u \in [\alpha]^n$, there are at most $O_a(n^{2a})$ words v at distance at most a and at most $O_a(n^{2c})$ words w at distance at most c , so there are at most $O_a(n^{2a+2c}) \leq O_a(n^{4a})$ good triples (u, v, w) for each choice of fixed u (and likewise for each fixed choice of v or w). We find that the total number of good triples containing a λ -repeating word is $\alpha^n n^{-8a} \cdot O_a(n^{4a}) = o(\alpha^n)$, which is negligible. It remains to bound the number of good triples consisting of λ -nonrepeating words.

It suffices to prove that every λ -nonrepeating u lies in at most $O_a(n^{a+b+c}(\log n)^{a+b-c})$ good triples (u, v, w) with v, w both λ -nonrepeating. Note that a good triple (u, v, w) is uniquely determined by the data of u and sequences $(I, T), (I', T')$ of insertion or deletion operations for which $w = f_{I, T}(u)$ and $v = f_{I', T'}(w)$. Since $d(u, w) \leq c$ and $d(w, v) \leq b$, we may choose (I, T) to have length at most $2c$ and (I', T') to have length at most $2b$. Since $f_{I', T'}(f_{I, T}(u)) = v$, we can “combine” the insertions and deletions of (I, T) and (I', T') to obtain a sequence (I'', T'') of insertions and deletions of length $|I''| = |I| + |I'| \leq 2b + 2c$ such that $f_{I'', T''}(u) = v$. Furthermore, there are only $O_a(1)$ choices of (I, T) and (I', T') that produce each such sequence (I'', T'') . Thus, for a given u , the number of good triples (u, v, w) is at most $O_a(1)$ times the number of ways to pick a sequence (I'', T'') of at most $2b + 2c$ total insertions and deletions such that $v = f_{I'', T''}(u)$ is λ -nonrepeating and $d(u, v) \leq a$.

By Lemma 6, the assumption $d(u, v) \leq a$ implies that at most $2a$ of the elements of I'' are λ -isolated. We claim that the total number of ways to pick such an I'' is at most $O_a(n^{a+b+c}(\log n)^{b+c-a})$. Indeed, we can define an equivalence relation \sim on the elements of I'' by setting $i \sim j$ if $|i - j| \leq 2\lambda$ and then taking the transitive closure. Let Q denote the number of equivalence classes. There are at most $2a$ equivalence classes of size 1 coming from λ -isolated elements, and $y \leq 2$ equivalence classes of size 1 coming from elements $i \in I''$ satisfying $i \leq \lambda$ or $i \geq n - \lambda$, which we call *boundary* equivalence classes. Hence, altogether $Q \leq 2a + y + (2b + 2c - 2a - y)/2 = a + b + c + y/2$. There are at most n^{Q-y} ways to choose the minimal elements of the non-boundary equivalence classes, λ^y ways to choose the minimal elements of the boundary equivalence classes, and then $(2\lambda)^{2b+2c-Q}$ ways to choose the remaining elements of I'' . The quantity $n^{Q-y}(2\lambda)^{y+2b+2c-Q}$ is at most $n^{a+b+c}(2\lambda)^{b+c-a}$, and multiplying by $a + b + c = O_a(1)$ (for the possible values of Q and y) establishes the claim. Finally, there are at most $(\alpha + 1)^{2b+2c} = O_{a,\alpha}(1)$ ways to pick T'' , and this completes the proof. \square

The proof of Theorem 1 is now immediate.

Proof of Theorem 1. By Lemma 7 with $a = b = c = k$, the number T of triangles in $\Gamma_{n,k,\alpha}$ satisfies $T = O_{\alpha,k}(\alpha^n n^{3k}(\log n)^k)$. Applying Lemma 4 with $N = \alpha^n$ and $\Delta = O_{\alpha,k}(n^{2k})$, we find that

$$D(n, k) = \Omega_{\alpha,k}(\alpha^n \log n / n^{2k}),$$

as desired. \square

3 LCS and SCS Multiplicity

In this section, we prove Proposition 2 and Theorem 3. Before proving Proposition 2, which says that the LCS multiplicity is always smaller than or equal to the SCS multiplicity, we set up one piece of notation. Suppose u is a word of length n which contains the word w of length ℓ as a subsequence. Then there is at least one subset $S \subseteq [n]$ of size ℓ such that $u_S = w$, and there may be several such subsets. We define $\text{left}(u, w)$ to be the smallest of these subsets according to the lexicographic ordering; that is, we choose S to have the smallest possible smallest element, and we break ties by looking at the second-smallest element, and so on. We can think of $\text{left}(u, w)$ as describing the position of the “left-most” copy of w in u .

Proof of Proposition 2. Let $|u| = m$, $|v| = n$, and $\text{LCS}(u, v) = \ell$, and note that $\text{SCS}(u, v) = m + n - \ell$. We define an injective map φ from the set of LCS's of u, v to the set of SCS's of u, v , as follows. Fix an LCS w of u, v . We now construct an SCS y of u, v one letter at a time.

To illustrate the idea, consider $u = 1011$ and $v = 0101$. There are two choices of an LCS for u and v , namely, $w = 101$ and $w = 011$. For the first choice $w = 101$, we mark its left-most copy in $u = \underline{1}011$ and $v = 0\underline{1}01$, and “glue” u and v together along these copies to obtain $\varphi(w) = 0\underline{1}0\underline{1}1$. For the second choice $w = 011$, we mark its left-most copy in $u = 1\underline{0}11$ and $v = 0\underline{1}0\underline{1}$ and glue to obtain $\varphi(w) = 1\underline{0}1\underline{0}1$.

Here is the formal description of the algorithm. To begin, initialize two indices $i = j = 1$ to track our current indices in u and v , respectively. For each $1 \leq k \leq m + n - \ell$, define y_k according to the following algorithm:

- (i) If $i \notin \text{left}(u, w) \cup \{m + 1\}$, then let $y_k = u_i$ and increment i .
- (ii) If $i \in \text{left}(u, w) \cup \{m + 1\}$ and $j \notin \text{left}(v, w) \cup \{n + 1\}$, then let $y_k = v_j$ and increment j .
- (iii) If $i \in \text{left}(u, w)$ and $j \in \text{left}(v, w)$, then let $y_k = u_i$ (which is also equal to v_j), and increment both i and j . (An easy induction shows that the r -th time this third possibility occurs, we have $u_i = v_j = w_r$.)

The number of times the algorithm falls into cases (i), (ii), (iii) above are (respectively) $m - \ell$, $n - \ell$, and ℓ , so the algorithm terminates at exactly $i = m + 1$, $j = n + 1$, with $m + n - \ell$ well-defined letters $y_1, \dots, y_{m+n-\ell}$. Define $y := y_1 y_2 \cdots y_{m+n-\ell}$. We have $|y| = m + n - \ell$ and y contains u, v as subsequences, so y is in fact an SCS of u, v . Finally, let $\varphi(w) = y$.

It remains to show that φ is injective, i.e., that w can be recovered from $\varphi(w)$. Note that item (iii) occurs if and only if $u_i = v_j$, so, working from $k = 1$ to $k = m + n - \ell$, we can determine the set K of indices k 's for which item (iii) occurs. By the parenthetical remark in item (iii), we get $w = \varphi(w)_K$, as needed. \square

As promised, we now prove Theorem 3 on the sharp upper bound for SCS multiplicity (and by extension LCS multiplicity). In fact, we establish a more general upper bound. If u and v are words, then we can order all of the common supersequences of u and v by inclusion and study the minimal common supersequences under this partial ordering. Note that the SCS's of u, v are always minimal common supersequences of u, v . The converse, however, is not always true: For instance, if $u = 1000$ and $v = 0001$, then the unique SCS of u, v is 10001, and the common supersequence 0001000 does not contain any proper subsequence containing u and v .

Lemma 8. *Let n, a, b be natural numbers with $n \geq a + b$. If u and v are words with length $n - a$ and $n - b$ (respectively) and $\text{SCS}(u, v) = n$, then the number of minimal common supersequences of u and v is at most $\binom{a+b}{a}$.*

Proof. We proceed by induction on $a + b$. The base case $a = b = 0$ is trivial. We now perform the induction step. If u, v have a common prefix, then every minimal supersequence of u, v must also share this prefix. By removing any common prefix of u, v , we may assume that u, v have different first letters. The key observation is that every minimal common supersequence of u, v is of the form

$$u_1 x \quad \text{or} \quad v_1 y,$$

where x is a minimal common supersequence of $u_{[2, n-a]}$ and v and y is a minimal common supersequence of u and $v_{[2, n-b]}$. The result now follows from Pascal's Identity for binomial coefficients, namely, $\binom{j}{i} + \binom{j}{i+1} = \binom{j+1}{i+1}$ for natural numbers $0 \leq i \leq j - 1$. \square

As mentioned above, Theorem 3 follows immediately from the observation that every SCS is a minimal common supersequence. We now show that the methods in Section 2 can be used to prove that most words $u, v \in [\alpha]^n$ at a given distance have a unique SCS and LCS.

Proposition 9. *If $n \geq k \geq 1$, then the number of pairs $u, v \in [\alpha]^n$ with $d(u, v) = k$ and $m_{\text{SCS}}(u, v) > 1$ is $O_{\alpha, k}(\alpha^n n^{2k-1} \log n)$.*

Proof. If $u, v \in [\alpha]^n$ and $d(u, v) = k$, then there exists a sequence of $2k$ operations (I, T) for which $v = f_{I, T}(u)$. There are at most $O_{\alpha, k}(n^{2k})$ choices of (I, T) of length $2k$. We say that (u, v) is *exceptional* if $m_{\text{SCS}}(u, v) > 1$.

Let $\lambda = 3 \log n$. The probability of a uniformly random $u \in [\alpha]^n$ being λ -repeating is at most $\binom{n}{\lambda} 2^{-\lambda} \leq n^{-1}$. The number of exceptional pairs (u, v) for which either u or v is λ -repeating is thus at most $O_{\alpha, k}(\alpha^n n^{2k-1})$, so it remains to count exceptional pairs where both u and v are λ -nonrepeating.

Suppose now that I is λ -separated. For each $\lambda < j < n - \lambda$, write $I(j) := [j - \lambda, j + \lambda]$, $L(j) := [j - \lambda, j - 1]$, and $R(j) := [j + 1, j + \lambda]$. Since I is λ -separated, we can partition $u = z_0 u_{I(j_1)} z_1 u_{I(j_2)} \cdots z_{2k-1} u_{I(j_{2k})} z_{2k}$ into subintervals (where the z_s 's may be empty). Furthermore, since the operations in (I, T) operate only within the $u_{I(j_s)}$'s, v can also be partitioned into $v = z_0 v_{I(j_1)} z_1 v_{I(j_2)} \cdots z_{2k-1} v_{I(j_{2k})} z_{2k}$ where the intermediate subintervals z_s are the same as in u . We now claim that the only SCS of u and v is the word $w = z_0 w_1 z_1 \cdots w_{2k} z_{2k}$ where w_s is the unique SCS of $u_{I(j_s)}$ and $v_{I(j_s)}$ (which differ by one letter).

The $4k$ intervals

$$L(j_1), R(j_1), L(j_2), R(j_2), \dots, L(j_{2k}), R(j_{2k})$$

are disjoint intervals of length λ . By the definitions of u , v , I , and T , the subintervals $u_{L(j_s)}$ and $u_{R(j_s)}$ each appear in u exactly once, and for each s , $u_{L(j_s)}$ and $u_{R(j_s)}$ are separated by one letter u_{j_s} . In v , they also appear exactly once each, and for each s , $u_{L(j_s)}$ and $u_{R(j_s)}$ are separated by zero or two letters in v , depending on whether f_{j_s, t_s} performs an insertion or a deletion.

Let w be a shortest common supersequence of u and v , so the length of w is $|u| + k$. We can form w from u by a sequence of k insertions (I'_w, T'_w) , and v from w by a sequence of k deletions (I''_w, T''_w) . In particular, the combined operation (I_w, T_w) for which $f_{I_w, T_w} = f_{I''_w, T''_w} \circ f_{I'_w, T'_w}$ is a sequence of $2k$ operations for which $v = f_{I_w, T_w}(u)$. Using the observations in the previous paragraph, we can “read off” from the distance between the copies of $u_{L(j_s)}$ and $u_{R(j_s)}$ inside v that I_w must have exactly one element in each $I(j_s)$. Unwinding the definition of (I_w, T_w) , this means w must be of the form $w = z_0 w_1 z_1 \cdots w_{2k} z_{2k}$ where w_s is the unique SCS of $u_{I(j_s)}$ and $v_{I(j_s)}$, as desired.

This proves the claim and shows that $m_{\text{SCS}}(u, v) = 1$ if u and v are λ -nonrepeating and I is λ -separated. The number of choices of I not λ -separated is at most $O_{\alpha, k}(n^{2k-1} \lambda) = O_{\alpha, k}(n^{2k-1} \log n)$, so the total number of exceptional pairs is at most $O_{\alpha, k}(\alpha^n n^{2k-1} \log n)$, as desired. \square

Acknowledgments. We are grateful to Venkatesan Guruswami for helpful conversations, and to the anonymous referees for comments that improved the presentation of this paper.

References

- [1] Ajtai, M., Komlós, J. and Szemerédi, E. (1980). A note on Ramsey numbers, *J. Combinatorial Theory, Ser. A* **29**, 354–360.
- [2] Alon, N. and Spencer, J. H. (2016). **The Probabilistic Method**, Fourth edition, Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ.
- [3] Brakensiek, J., Guruswami, V., and Zbarsky, S. (2016). Efficient low-redundancy codes for correcting multiple deletions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 1884–1892.
- [4] Bukh, B., Guruswami, V. and Håstad, J. (2017). An improved bound on the fraction of correctable deletions. *IEEE Trans. Inform. Theory* **63**, 93–103.
- [5] Guruswami, V. and Håstad, J. (2021). Explicit two-deletion codes with redundancy matching the existential bound. *IEEE Trans. Inform. Theory* **67**, 6384–6394.
- [6] Guruswami, V., He, X., and Li, R. (2021). The zero-rate threshold for adversarial bit-deletions is less than $1/2$. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 727–738.
- [7] Jiang, T. and Vardy, A. (2004). Asymptotic improvement of the Gilbert-Varshamov bound on the size of binary codes. *IEEE Trans. Inform. Theory* **50**, 1655–1664.

- [8] Levenshtein, V. (1965). I. Binary codes capable of correcting deletions, insertions, and reversals. Dokl. Akad. Nauk SSSR **163** 845–848; translated as Soviet Physics Dokl. **10**, 707–710. (In Russian.)
- [9] Shearer, J. B. (1983). A note on the independence number of triangle-free graphs, *Discrete Math.* **46**, 83–87.
- [10] Sima, J. and Bruck, J. (2019). Optimal k -deletion correcting codes. In *2019 IEEE International Symposium on Information Theory (ISIT)*, 847–851.
- [11] Varshamov, R. R., and Tenengolts, G. M. (1965). A code which corrects single asymmetric errors, *Automatika i Telemekhanika* **161**, 288–292. (In Russian.)

Appendix

In this appendix, we construct a family of pairs of words achieving equality in Theorem 3. It is easy to find pairs of words achieving equality for the SCS bound. For instance, we can take $u = 0^a$, $v = 1^b$; then $\text{SCS}(u, v) = a + b$ and $m_{\text{SCS}}(u, v) = \binom{a+b}{a}$ since the SCS's of u, v are precisely the words containing a 0's and b 1's. To find longer words achieving equality with the same values of a, b , simply append a fixed word w (for instance, $w = 0^c$) to the right of both u, v .

It seems that there is no similarly simple example achieving equality for the LCS bound, and our construction requires a delicate induction. If u is a (nonempty) word of length n and m is a natural number, then we define $u^{(m)}$ to be the prefix of length m of the infinite word $uuu\dots$. For instance, $(01)^{\langle 7 \rangle} = 0101010$ and $(0110)^{\langle 3 \rangle} = 011$. Our extremal example is as follows.

Proposition 10. *For every $c \geq 1$, the words $u = (10)^{\langle 4c-2 \rangle}$, $v = (0110)^{\langle 4c-2 \rangle}$ satisfy $d(u, v) = c$ and $m_{\text{LCS}}(u, v) = \binom{2c}{c}$.*

Let us explain why Proposition 10 provides equality cases for Theorem 3 for all choices of a, b . Suppose a_0, b_0 are given, and consider the words u, v produced by the $c = a_0 + b_0$ case of Proposition 10. Since equality in Theorem 3 is achieved for $(a, b) = (c, c)$ by u, v , we see that equality is also achieved for all of the other pairs of words considered in the inductive argument of Theorem 3 (which can easily be run directly with LCS's rather than passing through SCS's). For instance, the words $u_{[2, n-a]}, v$ are an equality case of Theorem 3 for $(a, b) = (c, c-1)$, and the words $u, v_{[2, n-b]}$ are an equality case of Theorem 3 for $(a, b) = (c-1, c)$. Continuing in this manner, we eventually reach an equality case for $(a, b) = (a_0, b_0)$, as needed.

To prove Proposition 10, we recursively compute $\text{LCS}(u, v)$ and $m_{\text{LCS}}(u, v)$ for all words $u = (10)^{\langle a \rangle}$, $v = (0110)^{\langle b \rangle}$ with b even. We introduce the notation $\ell(a, b) := \text{LCS}((10)^{\langle a \rangle}, (0110)^{\langle b \rangle})$ and $m(a, b) := m_{\text{LCS}}((10)^{\langle a \rangle}, (0110)^{\langle b \rangle})$. We begin by computing $\ell(a, b)$.

Lemma 11. *For $a, b \geq 0$ with b even, we have*

$$\ell(a, b) = \begin{cases} a & \text{if } a \leq b/2 \\ \frac{b}{2} + \lfloor \frac{2a-b}{4} \rfloor & \text{if } \frac{b}{2} < a \leq \frac{3b}{2} \\ b & \text{if } a > 3b/2. \end{cases}$$

Proof. The pairs (a, b) with $a \leq 2$ or $b = 0$ can be checked by hand, so we restrict our attention to $a \geq 3$ and $b \geq 2$. Note that every LCS of u, v is, according to its first letter, of the form

$$1x \quad \text{or} \quad 01y,$$

where x is an LCS of $(01)^{\langle a-1 \rangle}, (1001)^{\langle b-2 \rangle}$ and y is an LCS of $(01)^{\langle a-3 \rangle}, (1001)^{\langle b-2 \rangle}$. Exchanging the roles of 0 and 1, we find that

$$\text{LCS}((01)^{\langle a-1 \rangle}, (1001)^{\langle b-2 \rangle}) = \text{LCS}((10)^{\langle a-1 \rangle}, (0110)^{\langle b-2 \rangle}) = \ell(a-1, b-2),$$

and likewise $\text{LCS}((01)^{\langle a-3 \rangle}, (1001)^{\langle b-2 \rangle}) = \ell(a-3, b-2)$. It follows that

$$\ell(a, b) = \max\{1 + \ell(a-1, b-2), 2 + \ell(a-3, b-2)\},$$

and it is not difficult to check that the function defined in the lemma statement is the unique function satisfying this recurrence and the same initial conditions. \square

It remains to compute $m(a, b)$.

Lemma 12. *For $a, b \geq 0$ with b even, we have*

$$m(a, b) = \begin{cases} \binom{b/2}{(2a-b)/4} & \text{if } 2a \equiv b \pmod{4} \\ \binom{b/2+1}{(2a-b+2)/4} & \text{if } 2a \equiv b+2 \pmod{4}. \end{cases}$$

Proof. As in Lemma 11, we deal separately with the small cases where $a = 0$ or $b \leq 2$. Otherwise, following the same case distinction as in Lemma 11, we find that

$$m(a, b) = m(a-1, b-2) \cdot \mathbb{1}_{1+\ell(a-1, b-2) \geq 2+\ell(a-3, b-2)} + m(a-3, b-2) \cdot \mathbb{1}_{1+\ell(a-1, b-2) \leq 2+\ell(a-3, b-2)},$$

where $\mathbb{1}$ is the 0-1 indicator function of its argument. Using the exact values of ℓ from Lemma 11, we can rewrite this equation as

$$m(a, b) = \begin{cases} m(a-1, b-2) & \text{if } a \leq b/2 \\ m(a-1, b-2) + m(a-3, b-2) & \text{if } b/2 < a < \frac{3b}{2} \\ m(a-3, b-2) & \text{if } a \geq \frac{3b}{2}, \end{cases}$$

and the lemma follows from induction and Pascal's Identity. \square

Taking $a = b = 4c - 2$ in the previous two lemmas gives Proposition 10.