Syotti: Scalable Bait Design for DNA Enrichment

Jarno Alanko^{*1,2}, Ilya Slizovskiy³, Daniel Lokshtanov⁴, Travis Gagie², Noelle Noyes³, and Christina Boucher^{†5}

¹Department of Computer Science, University of Helsinki

²Faculty of Computer Science, Dalhousie University

³Department of Veterinary Population Medicine, University of Minnesota

⁴Department of Computer Science at the University of California Santa Barbara

⁵Department of Computer and Information Science and Engineering, University of Florida

Abstract

Bait-enriched sequencing is a relatively new sequencing protocol that is becoming increasingly ubiquitous as it has been shown to successfully amplify regions of interest in metagenomic samples. In this method, a set of synthetic probes (“baits”) are designed, manufactured, and applied to fragmented metagenomic DNA. The probes bind to the fragmented DNA and any unbound DNA is rinsed away, leaving the bound fragments to be amplified for sequencing. This effectively enriches the DNA for which the probes were designed. Most recently, Metsky et al. (Nature Biotech 2019) demonstrated that bait-enrichment is capable of detecting a large number of human viral pathogens within metagenomic samples. In this work, we formalize the problem of designing baits by defining the Minimum Bait Cover problem, which aims to find the smallest possible set of bait sequences that cover every position of a set of reference sequences under an approximate matching model. We show that the problem is NP-hard, and that it remains NP-hard under very restrictive assumptions. This indicates that no polynomial-time exact algorithm exists for the problem, and that the problem is intractable even for small and deceptively simple inputs. In light of this, we design an efficient heuristic that takes advantage of succinct data structures. We refer to our method as SYOTTI. The running time of SYOTTI shows linear scaling in practice, running at least an order of magnitude faster than state-of-the-art methods, including the recent method of Metsky et al. At the same time, our method produces bait sets that are smaller than the ones produced by the competing methods, while also leaving fewer positions uncovered. Lastly, we show that SYOTTI requires only 25 minutes to design baits for a dataset comprised of 3 billion nucleotides from 1000 related bacterial substrains, whereas the method of Metsky et al. shows clearly super-linear running time and fails to process even a subset of 8% of the data in 24 hours. Our implementation is publicly available at <https://github.com/jnalanko/syotti>.

*alanko.jarno@gmail.com

†christinaboucher@ufl.edu

1 Introduction

Our understanding of microbial species has evolved at an impressive rate, but it is estimated that 99% of microorganisms cannot live outside their natural environments, and thus, cannot be cultured and sequenced [15]. This created a significant roadblock in studying such species. In the mid-2000s, metagenomic shotgun sequencing became widely available, which enabled the sequencing of all of the DNA within non-cultured samples—whether a swab of a person’s mouth or a soil sample. This enabled the large-scale study of microorganisms, and thus greatly expanded our knowledge in this field. However, many scientific questions focus on specific sequences such as antimicrobial resistance (AMR) genes [14], or human viral strains and substrains [4]. To address these questions, most of the sequence reads in a metagenomic dataset are irrelevant, and are typically eliminated from further consideration. Moreover, the percentage of such reads can be quite high; for instance in agricultural samples, 90 to 95% of reads are often eliminated because they are not of interest to the research question [14, 17, 9]. Furthermore, using shotgun metagenomics in these scenarios not only leads to undue monetary expense but also greatly lowers the sensitivity of detection. For example, Lee et al. [10] demonstrated that out of 31 viral strains identified, 11 were unidentified via traditional sequencing, and Noyes et al. [14] reported that 1,441 AMR genes within agricultural samples were undetected by standard metagenomic sequencing.

One way to address these issues is to move beyond sequencing the entire microbial population within a sample, and rather to select the sequences of interest for targeted sequencing. This can be accomplished through *targeted enrichment*, which uses *biotinylated cDNA bait molecules* to target specific regions from DNA libraries for sequencing. More specifically, baits (short synthetically created single-stranded cDNA molecules) bind to their DNA targets and are then captured within the sample using a magnet. Non-captured, unbound DNA fragments (i.e., non-targeted DNA) are then rinsed away. Thus, only bound—i.e., targeted—DNA is sequenced. Although non-targeted DNA is not completely eliminated from sequencing, it is greatly reduced. One of the first, and arguably the most critical steps within this process requires solving a computational problem: for a given set of target DNA sequences (e.g., set of genes or viral strains) and a specified bait length k , identify a set of baits such that there exists at least one bait in that set that binds to each position of every sequence in the database.

Initially one could expect that baits could be computationally designed by finding all unique k -length subsequences (k -mers) in the targeted DNA. This solution however is not feasible as the number of k -mers grows rapidly with increasing dataset size. Thus, two key challenges need to be addressed for effective bait design. First, the number of baits must be minimized, as the cost of targeted enrichment is proportional to the size of the bait set and there is an upper limit set by the bait manufacturer. For example, Agilent Technologies, Inc. sets a limit of approximately 220,000 baits; larger bait sets can be manufactured, but would need to be split into multiple sets, thus greatly increasing the cost and labor of purchasing and processing multiple bait sets per sample. Second, the baits do not bind exactly to a single DNA sequence identical to that bait, but will bind to any subsequence with some allowed number of mismatches between the bait and target DNA sequence (typically, over 70% of positions must match). For this reason, computationally designing effective bait sets is a challenging problem, which to the best of our knowledge, has not been formally defined or considered from an algorithmic perspective. Although there are a number of methods for designing baits—including MrBait [2], CATCH [13], and BaitFisher [12]—most methods are unable to scale or provide reasonable output to even moderate-sized sets of genomes.

We formalize this problem, which we call The Minimum Bait Cover Problem, and show it is NP-hard even for extremely restrictive versions of the problem. For example, the problem remains intractable even for an alphabet of size four (i.e., A, T, C, G), a single reference genome, a bait length that is logarithmic in the length of the reference genome, and Hamming distance of zero. In light of these hardness results, we provide an efficient heuristic based on the FM-index and demonstrate that it is capable of scaling to large sets of sequences. We refer to our method as SYOTTI* and compare it to all competing methods on three datasets: (1) MEGARes, which is a database of antimicrobial resistance (AMR) genes, (2) a set of genomes corresponding to substrains of *Salmonella enterica* subsp. *enterica*, *Escherichia coli*, *Enterococcus* spp. and *Campylobacter jejuni*, and (3) a set of viral genomes. We determined that SYOTTI is the only method capable of scaling to the second and third datasets. MrBait was able to process all datasets, but was over 20 times slower and produced over 40 times more baits than SYOTTI, making it unreasonable in practice. CATCH

*Finnish word for “Bait”, with the letter ö replaced with the letter o.

failed to process both the second and third dataset in 24 hours, and even failed to process prefixes of length 8% and 5% respectively in the same time limit. The curve of the running time of CATCH indicates that it would take at least 20 days to process the viral dataset, and probably much longer.

Lastly, we evaluated the coverage of the genomes with respect to the bait sets using our implementation of the evaluation method provided with CATCH [13]. On the smallest dataset (MEGARes), both SYOTTI and CATCH covered 100.0% of the nucleotide positions at least once, but MrBait covered only 96.4%. On the largest subset of bacteria genomes that all tools were able to run, SYOTTI produced 158 thousand baits, CATCH produced 241 thousand, and MrBait over 1 million; hence, the bait set of MrBait was too large to be deemed of any practical use. Lastly, we considered the full set of viral genomes and compared the baits of SYOTTI against the published and publicly available bait sets of size 250k, 350k and 700k published with CATCH. The coverage of the CATCH bait sets were 84.1%, 90.6% and 97.5% respectively. The coverage of SYOTTI was 99.5%, with the remaining 0.5% being due to unknown N-characters in the data.

2 The Minimum Bait Cover Problem

2.1 Preliminaries

We define a string S as a finite sequence of characters $S = S[1 \dots n]$ over an alphabet $\Sigma = \{c_1, \dots, c_\sigma\}$. We denote the length of a string S by $|S|$, and the empty string by ε . We denote by $S[1 \dots i]$ the i -th prefix of S , and by $S[i \dots n]$ the i -th suffix of S . Given two strings S and T , we say that S is lexicographically smaller than T if either S is a prefix of T or there exists an index $i \geq 1$ such that $S[1 \dots i] = T[1 \dots i]$ and $S[i+1]$ occurs before $T[i+1]$ in Σ . We denote this as $S \prec T$. Next, we define the *Hamming distance* between S and T (which have the same length) as the number of positions where S and T mismatch, namely $d(S, T) := |\{i : S[i] \neq T[i]\}|$. Lastly, we denote $S \circ T$ for the string concatenation of S and T . Given a set of strings, $\mathcal{S} = \{S_1, \dots, S_m\}$, we denote the total length $\sum_{S \in \mathcal{S}} |S|$ of all strings in $\mathcal{S}$ by $\|\mathcal{S}\|$.

For a string S and integer $\ell \leq |S|$ we denote by $\text{pre}(S, \ell)$ the prefix $S[1, \dots, \ell]$ of S of length ℓ . Similarly we denote by $\text{suf}(S, \ell)$ the suffix $S[\ell - \ell + 1, \dots, \ell]$ of S of length ℓ .

2.2 Hardness Results

We fix the length of the bait sequences to a constant L . We say a string X *covers* a position $i \leq |Y|$ in a string Y if there exists a j such that $i - |X| < j \leq i$ and $X = Y[j, j+1, \dots, j+|X|-1]$. More generally, X *distance θ -covers* (we will just say *θ -covers*) position i in Y if there exists a j such that $i - |X| < j \leq i$ and $d(X, Y[j, j+1, \dots, j+|X|-1]) \leq \theta$. Observe that X covers i if and only if X 0-covers i . A set $\mathcal{T}$ of strings θ -cover a string S if, for every position i in S there exists a string $T \in \mathcal{T}$ that θ -covers i . A set $\mathcal{T}$ of strings θ -cover a set $\mathcal{S}$ of strings if every string $S \in \mathcal{S}$ is θ -covered by $\mathcal{T}$. We are now ready to define the main problem considered in this paper, namely MINIMUM BAIT COVER.

MINIMUM BAIT COVER

Input: Here the input consists of two integer parameters $\theta \geq 0$ and $L > 0$ and a set $\mathcal{S}$ of n strings $\mathcal{S} = \{S_1, \dots, S_n\}$ over a finite alphabet Σ .

Question: What is the smallest possible set $\mathcal{T}$ of strings, each of length exactly L , such that $\mathcal{S}$ is θ -covered by $\mathcal{T}$.

In the closely related STRING COVER problem input is a single string S and a parameter $L > 0$. The task is to find a minimum size set $\mathcal{T}$ of strings, each of length exactly L , such that S is θ -covered by $\mathcal{T}$. Comparing the two problem definitions it is easy to see that STRING COVER is precisely equal to the special case of MINIMUM BAIT COVER when $n = 1$ and $\theta = 0$. Cole et al. [3] proved that STRING COVER is NP-hard even for $L = 2$. This immediately leads for the following hardness result for MINIMUM BAIT COVER.

Proposition 2.1. [3] (a) STRING COVER is NP-hard even for $L = 2$. (b) MINIMUM BAIT COVER is NP-hard, even for $n = 1$, $\theta = 0$ and $L = 2$.

Proposition 2.1 effectively rules out any hope of a polynomial time algorithm, or an efficient (Fixed Parameter Tractable or Slicewise Polynomial) parameterized algorithm with parameters n , θ , and L . At the

same time the instances of STRING COVER constructed in the reduction of Cole et al. [3] are strings over a large alphabet $|\Sigma|$. Since we are primarily interested in instances of MINIMUM BAIT COVER with $|\Sigma| = 4$, this might leave some hope that for small alphabets one can still obtain efficient parameterized algorithms. We now show that (almost) all of the hardness from Proposition 2.1 is retained even when $|\Sigma|$ is constant.

Theorem 2.2. (a) For every $k \geq 2$, STRING COVER is NP-hard with $|\Sigma| = k$ and $L = O(\log |S|)$. (b) For every $k \geq 2$, MINIMUM BAIT COVER is NP-hard, even for $|\Sigma| = k$, $\mathcal{S} = \{S\}$, $\theta = 0$ and $L = O(\log |S|)$.

Theorem 2.2 shows that unless $P = NP$ there can't exist an algorithm for MINIMUM BAIT COVER with running time $2^{O(L)} \|\mathcal{S}\|^{g(|\Sigma|, n, \theta)}$ for any function g of constant parameters $|\Sigma|, n$ and θ . This effectively rules out parameterized algorithms that exploit any of the most natural parameters one could expect to be small in relevant input instances. We remark that part (a) of Theorem 2.2 resolves in the affirmative a conjecture of Cole et al. [3] that STRING COVER remains NP-hard even with a constant size alphabet.

Proof of Theorem 2.2. We only prove (a), since (b) follows from (a) together with the fact that STRING COVER is the special case of MINIMUM BAIT COVER with $\theta = 0$ and $n = 1$. Fix an integer $k \geq 2$. We give a reduction from STRING COVER with $L = 2$ and unbounded alphabet size (i.e. $|\Sigma| \leq |S|$) to STRING COVER with alphabet $\Sigma_k = \{0, \dots, k-1\}$ of size exactly k .

We now describe the construction. Given as input an instance $(\Sigma, S, L = 2)$ of STRING COVER the reduction algorithm sets $\Sigma_k = \{0, \dots, k-1\}$, ℓ to be the smallest integer above $\lceil 100 \log_2(|\Sigma|) \rceil$ that is divisible by 3, and $L' = 2\ell$. The algorithm then computes a set of strings $\{S_a : a \in \Sigma\}$ by using Claim 2.3.

Claim 2.3. There exists an algorithm that given as input a set Σ , runs in time polynomial in $|\Sigma|$, and outputs a set $\{S_a : a \in \Sigma\}$ of strings over the alphabet Σ_k , all of length $L' = 2\ell$, that satisfy the following properties:

1. For distinct characters $a, b \in \Sigma$ we have $\text{pre}(S_a, \ell/3) \neq \text{pre}(S_b, \ell/3)$,
2. for distinct characters $a, b \in \Sigma$ we have $\text{suf}(S_a, \ell/3) \neq \text{suf}(S_b, \ell/3)$,
3. for every pair of characters $a, b \in \Sigma$ (including the case when $a = b$) and every $r \geq \ell/3$ we have that $\text{pre}(S_a, r) \neq \text{suf}(S_b, r)$.

Proof. Let us first observe that the total number of strings in $(\Sigma_k)^\ell$ is $\Theta(2^{\log_2 k 100 \log_2 |\Sigma|}) = \Theta(|\Sigma|^{100 \log_2 k})$ and therefore upper bounded by a polynomial in $|\Sigma|$. The algorithm initializes a list of strings $\mathcal{L}$ to contain all strings in $(\Sigma_k)^\ell$. The algorithm then removes from $\mathcal{L}$ all strings T such that there exists some $r \geq \ell/3$ such that $\text{pre}(T, r) \neq \text{suf}(T, r)$. For every $r \geq \ell/3$, the number of strings $T \in (\Sigma_k)^\ell$ such that the prefix of T of length r is equal to the suffix of T of length r is at most $k^{\frac{2\ell}{3}}$ (because fixing the $\frac{2\ell}{3}$ first characters uniquely determines the remaining $\frac{\ell}{3}$ characters). Thus, the total number of strings removed from the list in this initial cleaning step is at most $\ell \cdot k^{\frac{2\ell}{3}}$.

The algorithm then iterates over the characters $a \in \Sigma$ one by one. When considering the character a the algorithm picks an arbitrary string still on the list $\mathcal{L}$, calls it S_a , and removes it from the list $\mathcal{L}$. It then goes over all strings on the list $\mathcal{L}$ and removes all strings T such that $\text{pre}(T, \ell/3) = \text{pre}(S_a, \ell/3)$, or $\text{suf}(T, \ell/3) = \text{suf}(S_a, \ell/3)$, or there exists an $r \geq \ell/3$ such that $\text{pre}(T, r) = \text{suf}(S_a, r)$ or $\text{suf}(T, r) = \text{pre}(S_a, r)$.

There are at most $k^{\frac{2\ell}{3}}$ strings T such that $\text{pre}(T, \ell/3) = \text{pre}(S_a, \ell/3)$ (since $\frac{\ell}{3}$ of the characters of S are uniquely determined by $\text{pre}(S_a, \ell/3)$). Similarly, there are at most $k^{\frac{2\ell}{3}}$ strings T such that $\text{suf}(T, \ell/3) = \text{suf}(S_a, \ell/3)$, and at most $\ell \cdot k^{\frac{2\ell}{3}}$ strings T such that there exists an $r \geq \ell/3$ such that $\text{pre}(T, r) = \text{suf}(S_a, r)$ or $\text{suf}(T, r) = \text{pre}(S_a, r)$. Therefore, in each iteration (i.e. after selecting one string S_a) the algorithm removes at most $4\ell \cdot k^{\frac{2\ell}{3}}$ strings from $\mathcal{L}$.

There are $|\Sigma| \leq 2^{\frac{\ell}{100}}$ iterations of the algorithm in total. Thus, in each iteration of the algorithm there are at least $k^\ell - 4|\Sigma|\ell k^{\frac{2\ell}{3}} \geq \frac{k^\ell}{2}$ strings from $\mathcal{L}$ to choose the next string S_a from. Thus, the algorithm successfully selects a string $S_a \in \mathcal{L}$ for every character $a \in \Sigma$.

Finally we argue that the set of strings $\{S_a : a \in \Sigma\}$ satisfy properties 1, 2 and 3. Suppose for distinct characters a, b we have that $\text{pre}(S_a, \ell/3) = \text{pre}(S_b, \ell/3)$ or $\text{suf}(S_a, \ell/3) = \text{suf}(S_b, \ell/3)$, or there exists an $r \geq \ell/3$ such that $\text{pre}(S_a, r) = \text{suf}(S_b, r)$ or $\text{suf}(S_a, r) = \text{pre}(S_b, r)$. Without loss of generality S_a was selected before S_b , and then S_b would have been removed from $\mathcal{L}$ in the cleaning step immediately after S_a was

selected. This contradicts that S_b was selected from the list $\mathcal{L}$. Further, suppose that there exists an $r \geq \ell/3$ such that $\text{pre}(S_a, r) = \text{suf}(S_a, r)$ for some string S_a . But then S_a would have been removed in the initial cleaning step, again contradicting that S_a was selected from $\mathcal{L}$. This completes the proof of the claim. $\square$

The reduction algorithm now constructs the string S' from S by replacing every character $a \in \Sigma$ with the corresponding string S_a . The reduction algorithm outputs the instance (Σ', S', L') . This concludes the construction. We claim that for every integer t , there exists a set $\mathcal{T}$ of t strings of length $L = 2$ that cover S if and only if there exists a set $\mathcal{T}'$ of t strings of length L' that cover S' .

We begin by proving the forward direction of the above claim. Let $\mathcal{T}$ be a set of t strings of length L that cover S . We define $\mathcal{T}' = \{S_a S_b : ab \in \mathcal{T}\}$. Clearly $|\mathcal{T}'| = t$. To see that $\mathcal{T}'$ covers S' consider an arbitrary position p' in the string S' and let $p = \lceil \frac{p'}{\ell} \rceil$. Since $\mathcal{T}$ covers p there is a string $ab \in \mathcal{T}$ that covers p in S . But then $S_a S_b \in \mathcal{T}'$ covers all positions $\{\ell \cdot p - \ell + 1, \ell \cdot p - \ell + 2, \ell \cdot p\}$ in S' . However $p' \in \{\ell \cdot p - \ell + 1, \ell \cdot p - \ell + 2, \ell \cdot p\}$ and therefore $\mathcal{T}'$ covers p' . Since p' was an arbitrarily chosen position we conclude that $\mathcal{T}'$ covers S' .

We now prove the reverse direction: if there exists a set $\mathcal{T}'$ of t strings of length L' that cover S' then there exists a set $\mathcal{T}$ of t strings of length $L = 2$ that cover S . Without loss of generality, every string $T \in \mathcal{T}'$ is a substring of S' of length 2ℓ (otherwise $\mathcal{T}' \setminus \{T\}$ also covers S'). Thus $T = \text{suf}(S_a, i) \circ S_b \circ \text{pre}(S_c, \ell - i)$ for some $i \in \{1, \dots, \ell\}$ and $a, b, c \in \Sigma$. Note that if $i = \ell$ then $T = S_a S_b$ for $a, b \in \Sigma$. Next we prove a claim about strings on this form.

Claim 2.4. *If $\text{suf}(S_a, i) \circ S_b \circ \text{pre}(S_c, \ell - i) = \text{suf}(S_{a'}, i') S_{b'} \text{pre}(S_{c'}, \ell - i')$ then $i = i'$ and $b = b'$. Furthermore, if $i \geq \ell/2$ then $a = a'$, otherwise $c = c'$.*

Proof. Suppose that $\text{suf}(S_a, i) \circ S_b \circ \text{pre}(S_c, \ell - i) = \text{suf}(S_{a'}, i') S_{b'} \text{pre}(S_{c'}, \ell - i')$, and assume for contradiction that $i \neq i'$. Without loss of generality we have $i < i'$.

We have that: $\text{suf}(S_{a'}, i' - i) = \text{pre}(S_b, i' - i)$, $\text{suf}(S_b, \ell - i' + i) = \text{pre}(S_{b'}, \ell - i' + i)$, and $\text{suf}(S_{b'}, i' - i) = \text{pre}(S_c, i - i)$. Furthermore we have that $|\text{suf}(S_{a'}, i' - i)| + |\text{suf}(S_b, \ell - i' + i)| + |\text{suf}(S_{b'}, i' - i)| = \ell + i' - i \geq \ell$. Thus, at least one of $\text{suf}(S_{a'}, i' - i)$, $\text{suf}(S_b, \ell - i' + i)$, or $\text{suf}(S_{b'}, i' - i)$ has length at least $\ell/3$. However, either one of $\text{suf}(S_{a'}, i' - i)$, $\text{suf}(S_b, \ell - i' + i)$, or $\text{suf}(S_{b'}, i' - i)$ having length at least $\ell/3$ contradicts Property 3 of the set $\{S_a : a \in \Sigma\}$ of strings. We conclude that $i = i'$. But then $\text{suf}(S_a, i) = \text{suf}(S_{a'}, i)$, $S_b = S_{b'}$ and $\text{pre}(S_c, \ell - i) = \text{pre}(S_{c'}, \ell - i)$. Property 1 implies that $b = b'$. If $i \geq \ell/2$ then Property 2 implies that $a = a'$. If $i < \ell/2$ then $\ell - i \geq \ell/3$ and therefore Property 1 implies that $c = c'$. $\square$

For every string $T \in \mathcal{T}'$ we have that $T = \text{suf}(S_a, i) \circ S_b \circ \text{pre}(S_c, \ell - i)$ for some $i \in \{1, \dots, \ell\}$ and $a, b, c \in \Sigma$. By Claim 2.4 i and b are uniquely determined by T . If $i \geq n/2$ then by Claim 2.4 a is also uniquely determined by T . In this case we add ab to the set $\mathcal{T}$. If $i < \ell/2$ then by Claim 2.4 c is also uniquely determined by T . In that case we add bc to the set $\mathcal{T}$. In either case we add precisely one string of length 2 to $\mathcal{T}$ for each $T \in \mathcal{T}'$. Thus $|\mathcal{T}| \leq |\mathcal{T}'|$ and it remains to prove that $\mathcal{T}$ covers S .

Consider an arbitrary position p of S and let $abcde = S[p - 2, p - 1, p, p + 1, p + 2]$. Let $p' = p\ell - \lfloor \frac{\ell}{2} \rfloor$, and let $T \in \mathcal{T}'$ be such that $T = S[x, x + 1, \dots, x + L' - 1]$ where $p' \in \{x, x + 1, \dots, x + L' - 1\}$. We have the following cases:

1. $T = \text{suf}(S_a, i) \circ S_b \circ \text{pre}(S_c, \ell - i)$, where $\ell - i \geq \ell/2$, or
2. $T = \text{suf}(S_b, i) \circ S_c \circ \text{pre}(S_d, \ell - i)$, where $i \in 1, \dots, \ell$, or
3. $T = \text{suf}(S_c, i) \circ S_d \circ \text{pre}(S_e, \ell - i)$, where $i \geq \ell/2$.

In the first case $bc \in \mathcal{T}$ and $S[p - 1, p] = bc$, in the second case $bc \in T$ and $S[p - 1, p] = bc$ or $cd \in T$ and $S[p, p + 1] = cd$, while in the third case $cd \in \mathcal{T}$ and $S[p, p + 1] = cd$. In either case the position p is covered by $\mathcal{T}$, and so S is covered by $\mathcal{T}$.

We remark that in the argument above, if $p \in \{1, 2, |S| - 1, |S|\}$ then some of a, b, d, e are not properly defined. However this only restricts which of the cases 1, 2 and 3 may apply, one of them must still apply (for characters from a, b, c, d, e that are well defined). This concludes the proof. $\square$

Algorithm 1 The SYOTTI algorithm.

Input: Sequences $S_1, \dots, S_n$, bait length L , Hamming distance θ , seed length k

Output: Set of bait sequences covering $S_1, \dots, S_n$

```

1: Construct the FM-index and the generalized suffix array  $GSA$  of  $S_1 \dots, S_n$  ▷ Preprocessing
2: Baits  $\leftarrow \emptyset$ 
3: for  $i = 1..n$  do
4:   for  $j = 1..|S_i|$  do
5:     if  $S_i[j]$  is not marked as covered then
6:        $B \leftarrow S_i[j..j + L - 1]$ 
7:       Baits  $\leftarrow$  Baits  $\cup \{B\}$ 
8:       MarkMatches( $B$ )
9:       MarkMatches(ReverseComplement( $B$ ))
10: return Baits.
11:
12: function MARKMATCHES( $B$ ):
13:   for  $i = 1..|B| - k + 1$  do ▷ For all seed starting positions
14:      $x \leftarrow B[i..i + k - 1]$ 
15:      $[\ell, r] \leftarrow$  FMindex.Search( $x$ ) ▷ GSA-interval of the seed  $x$ 
16:     for  $j = \ell..r$  do
17:        $t, p \leftarrow GSA[j]$  ▷ Locate the seed occurrence  $x = S_t[p..p + k - 1]$ 
18:        $B' \leftarrow S_t[p - i + 1..p - i + L]$  ▷ Extend from the occurrence
19:       if  $d(B', B) \leq \theta$  then ▷ Hamming distance at most  $\theta$ .
20:         Mark  $S_t[p - i + 1..p - i + L]$  as covered
```

2.3 The Syotti Algorithm

Our algorithm requires an index that can search for substrings in the input sequences $S_1, \dots, S_n$ and report the locations of the exact matches. For this purpose, we use the combination of an FM-index [6] and a generalized suffix array GSA (the suffix array of a set of strings [18]). Given a substring x , the FM-index is able to compute an interval $[\ell, r]$ such that the entries $GSA[\ell], \dots, GSA[r]$ give all starting positions of x in $S_1, \dots, S_n$. That is, the entries of GSA in the interval are pairs (t, p) such that $S_t[p..p + |x| - 1] = x$. We refer the interested reader to Mäkinen et al. [11] for a detailed exposition of the techniques involved.

First, we construct the aforementioned data structures from the input sequences. Then, we initialize the set of bait sequences to an empty set, and initialize a bit vector for each sequence, of the same length as that sequence. The bit vector signifies which positions are covered by the bait sequences. Next, we perform a linear scan of each sequence, checking the bit vector at each position. When we come across a position that is not yet marked as covered, we add the substring of length L starting from that position into the set of bait sequences and update the bit vectors for all positions covered by the new bait. More specifically, we search for positions in all the sequences to which the new bait has Hamming distance less than or equal to θ . This search is done with a seed-and-extend heuristic based on the exact matching index as follows. For each seed k -mer x in the new bait sequence B , we find all occurrences of x in our input sequences. For each occurrence of x we consider the L -length sequence B' that contains x in the same position as in B . If the Hamming distance of B and B' is at most θ , we update the bit vector; otherwise, we move onto the next occurrence of x . After processing all seeds of B , we repeat the process for the reverse complement of B .

We give the pseudocode in Algorithm 1. There are some implementation details omitted from the pseudocode to avoid obfuscating the main idea of the algorithm. The string indices can go past the ends of the strings, but these corner cases are handled easily. That is, if a bait would run past the end of the string, we shift it back inside the string, and we ignore candidate matches from seeds that run past either end of the string. Also, if implemented directly as in the pseudocode, the seed-and-extend algorithm often ends up comparing the bait against the same position from multiple seeds in the same bait. To avoid this, we store the starting positions of the candidate matches in a set data structure that keeps only the distinct starting positions. We then run the Hamming distance computation on each distinct candidate match.

3 Experiments

We implemented SYOTTI in the C++ programming language. We use the Divsufsort and SDSL libraries [8] for the construction of the FM-index and the generalized suffix array required by the SYOTTI algorithm. The implementation of SYOTTI is publicly available at <https://github.com/jnalanko/syotti>.

The main competitor for SYOTTI is the method CATCH by Metsky et al. [13]. CATCH reduces the bait design problem to an instance of the classic NP-hard set cover problem, and finds an approximate solution by greedily picking baits in decreasing order of profitability. Two other potential competitors are MrBait [2] and BaitFisher [12]. MrBait employs a bait tiling strategy followed by optional post-processing filtering, and BaitFisher works by clustering sequences using a multiple alignment and tiling the consensus sequences of the clusters with baits.

We evaluate our method and the competing methods on three datasets: (1) **MEGARES**, which is a database containing 7,868 AMR genes of total length 8,106,325 bp, publicly available at the MEGARes website: <https://megares.meglab.org/download>. (2) **VIRAL**, which is a set of 422,568 viral genomes from CATCH [13], of total length 1,257,789,768 bp, publicly available at: <https://github.com/broadinstitute/catch>. (3) **BACTERIA**, which is a custom (see Appendix A) set of 1,000 bacterial genomes representing four foodborne pathogens containing 98,984 sequences of total length 3,040,260,476 bp.

As advised by the bait manufacturer Agilent, we set the bait length to 120 and tolerated 40 mismatches. The seed length in SYOTTI was set to 20. We ran CATCH, MrBait, and SYOTTI on these datasets on a server with two Intel Xeon Gold 5122 CPUs at 3.60GHz for total of 8 physical cores and up to 16 threads with hyperthreading, 503GB of available memory and a 3TB NVMe SSD for storage. The time is measured in wall-clock seconds. The experiments do not include BaitFisher because it requires a multiple sequence alignment of all the input sequences as input, but our sequences do not all align well to each other. Feeding multiple alignments computed with Clustal Omega [19] to BaitFisher produced extremely poor baits sets that had more nucleotides than the original sequences.

CATCH was run with the command line options `-pl 120 -m 40 -l 120 -ps 60`. We also experimented with the options `--cluster-and-design-separately 0.15` and `--cluster-from-fragments 10000`, but did not see significant improvement on the results or the run time. MrBait was run with the options `-A 120 -s tile=0`. MrBait also offers post-processing filtering using all-pairs global alignment from the external tool VSEARCH [16]. This was feasible only for MEGARES because the time complexity of all-pairs alignment is quadratic. SYOTTI was run with options `-L 120 -d 40 -r -t 16 -g 20`.

We restricted all methods to using 24 hours of time, and the available 503 GB of memory, and 3 TB of disk space. None of the methods exceeded the memory and disk limits before being terminated by the 24 hour time limit. We analyzed the coverage of the produced bait sets by mapping the baits to the reference sequences using the seed-and-extend method with seed length 20, tolerating up to 40 mismatches. This is the same method of analysis as the one used in the coverage analyser of CATCH. We did not use the CATCH analyser because it does not scale to our larger datasets. To verify the comparability of the two approaches, we ran them both on small inputs and found negligible difference.

3.1 Results on Antimicrobial Resistance Genes

We ran all methods on exponentially increasing subsets of the MEGARes database that consists of 7,868 AMR genes. All methods were capable of being run on the complete dataset. The memory, running time, and number of baits is given in Table 1. The scaling of the time usage is plotted in Figure 1.

Even though the worst case running time of SYOTTI is at least quadratic in the length of the input, the running time behaved approximately linearly. MrBait also behaved approximately linearly, but up to 40 times slower than SYOTTI. On the other hand, CATCH showed an upward-bending growth curve, which we believe is due to having to update the profitability scores of candidate baits after every iteration, resulting in a quadratic time complexity in practice as well as in the worst case.

The size of bait sets of CATCH and SYOTTI had negligible difference, with the bait set of SYOTTI being consistently slightly smaller. MrBait produced 2-3 times larger bait sets than the other two. The memory usage of SYOTTI was the smallest of the three, however, the memory usage for all methods was less than 2 GB. The slope of the memory of MrBait was smaller than SYOTTI, which indicates that the memory of SYOTTI would grow past MrBait if the dataset size was still increased. We believe that MrBait's good

	SYOTTI			CATCH			MrBait		
Input length (bp)	Time (mm:ss)	Memory (MB)	Baits (#)	Time (mm:ss)	Memory (MB)	Baits (#)	Time (mm:ss)	Memory (MB)	Baits (#)
125271	00:00	6	820	00:01	78	824	00:02	72	987
254294	00:00	8	1372	00:04	104	1392	00:04	73	1981
508459	00:00	11	2604	00:09	158	2635	00:09	76	3983
1032802	00:00	20	4633	00:28	260	4742	00:19	80	8093
2090517	00:00	35	7901	01:30	423	8121	00:37	88	16374
4187569	00:01	67	13099	05:37	735	13489	01:12	101	32764
8106325	00:03	125	20976	19:13	1250	21771	02:19	128	63428

Table 1: Results on the dataset MEGARES (without VSEARCH filtering). The seconds are rounded down. Rows where all tools ran in less than one second have been removed. The full table is in Appendix B.

Figure 1: Time scaling on all three datasets.

memory scaling is due to its use of a disk-based SQLite database.

Lastly, we evaluated the coverage of the baits. We note that the most optimal solution has a coverage value of 1, indicating sufficient likelihood of bait binding, but without redundancy (i.e., inefficiency) in the bait set. This is important for two reasons: first, multiple baits for the same position could create bait-bait interference during the binding process; and second, each bait is relatively costly, with a fixed upper limit on the total number of baits that can be manufactured. Coverage was analyzed by mapping the baits to the reference sequences with a seed-and-extend approach with a seed length of 20. The coverage of a position is defined as the number of mapped baits spanning that position with at most 40 mismatches. Figure 2 shows the mean coverage in each sequence of the database. Both SYOTTI and CATCH covered 100.0% of the nucleotide positions at least once, but MrBait covered only 96.4%. The bait set designed by MrBait was extremely inefficient, with many positions being covered by hundreds of baits (Figure 2). The coverage efficiencies for CATCH and SYOTTI were comparable, with a slight advantage for SYOTTI (Figure 2). Filtering the MrBait baits with VSEARCH and adding the command line parameters `-f pw=0.9,0.9` to MrBait, resulted in 22,230 baits with coverage 95.9%, but the run time increased drastically from 2min 19s to 4h 22min 36s.

3.2 Results on Bacterial Strains and Substrains

The aim of our second experiment was to study the scalability of the methods on pangenomes of clinically relevant bacterial species. Lacking a standard reference database of this type, we built our own, incorporating available sequences from foodborne pathogens *Salmonella enterica*, *Campylobacter jejuni*, *Escherichia coli*, *Enterococcus faecalis*. The purpose of the dataset is to cover as much of the known sequence diversity of these species as possible. The data was carefully selected and filtered to be suitable for the downstream application of foodborne pathogen detection. The full details of the process of compiling the dataset are

Figure 2: Average sequence coverage on the MEGARES dataset with all three methods (left) and with just CATCH and SYOTTI (right). The sequences are in the order of the MEGARES database. The coverage plot with all tools after filtering MrBait baits with VSEARCH is in Appendix B.

	SYOTTI			CATCH			MrBait		
Input length	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits
15920441	00:00:08	237	72035	00:29:32	4096	90182	00:04:48	195	132408
30786116	00:00:14	454	96595	01:00:27	7575	134112	00:10:00	316	256041
62502135	00:00:27	917	123541	02:26:56	14234	181825	00:18:44	572	519838
125063199	00:00:52	1831	157818	08:43:14	24269	240747	00:38:48	1076	1040174
254576853	00:01:45	3723	188813	>24h	NA	NA	01:20:03	2071	2117436
505422833	00:03:31	7387	223931	NA	NA	NA	02:39:03	4056	4203752
1003934029	00:07:11	14673	267890	NA	NA	NA	05:18:28	7980	8349888
2018459352	00:15:50	29496	324797	NA	NA	NA	10:30:37	15697	16788084
3040260476	00:24:52	44425	366761	NA	NA	NA	16:13:43	23615	25286576

Table 2: Running time (Time), peak memory usage (Memory), and number of baits (baits) for increasingly larger subsets of sequences from the BACTERIA dataset. The times are in the format hh:mm:ss. “NA” signifies that the dataset surpassed 24 hours of running time. Rows where all tools took less than 30 minutes have been removed to save space. The full data table is in Appendix B

available in Appendix A. We note that this database and bait set from SYOTTI will be used in follow-up studies that aim to amplify these bacteria genomes in samples taken from food production facilities, in order to test for foodborne pathogens.

To measure the scalability of the methods, we ran them on subsets of the data containing 1,2,4,8,16 ... 65536 sequences, as well as the full dataset consisting of all 98,984 sequences. The results are shown in Table 2. SYOTTI and MrBait were able to process all inputs, but CATCH hit the 24 hour time limit on inputs larger than 250 million base pairs.

All three tools produced bait sets of similar size for the smaller prefixes. But as the length of the prefix increased, the bait sets produced by SYOTTI became significantly smaller than those produced by the other two tools. On the largest prefix that all tools were able to run, SYOTTI produced 158 thousand baits, CATCH 241 thousand, and MrBait 1040 thousand. This showcases the ability of SYOTTI to deal with multiple bacterial pangenomes at once. On the full dataset, SYOTTI produced 367 thousand baits, whereas MrBait produced 25 million. Given the 220,000 limit provided by Agilent, two bait kits could be manufactured to include all our bait sequences, but over 113 bait kits would have to be manufactured for MrBait (which would cost more than 1 million USD for a single sample), making this bait set highly impractical.

The running times and memory usage of the tools showed similar trends to the first experiment on AMR

	SYOTTI			CATCH			MrBait		
Input length	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits
26238	00:00:00	5	148	00:00:00	55	164	00:00:01	71	218
9672491	00:00:03	145	1103	01:19:09	2878	1065	00:02:43	130	67054
66970374	00:00:37	1011	1640	> 24h	NA	NA	00:18:59	577	522683
91129015	00:00:55	1370	6016	NA	NA	NA	00:25:36	705	680298
93144457	00:00:56	1400	6959	NA	NA	NA	00:28:41	734	696193
103290818	00:01:00	1549	11737	NA	NA	NA	00:32:01	799	773385
155140833	00:01:28	2321	28543	NA	NA	NA	00:45:40	1154	1160119
230198416	00:04:04	3424	60539	NA	NA	NA	01:09:07	1549	1557350
564924375	00:07:45	8350	103401	NA	NA	NA	02:45:21	3975	4139176
1040580227	00:12:39	15391	174742	NA	NA	NA	05:14:59	7164	7757360
1257789768	00:16:41	18613	226751	NA	NA	NA	06:24:44	8519	9225038

Table 3: Running time (Time), peak memory usage (Memory), and number of baits (baits) for increasingly larger subsets of sequences from the VIRAL dataset. The times are in the format hh:mm:ss. “NA” signifies that the dataset surpassed 24 hours of running time.

Figure 3: Coverage of the 684k baits generated by SYOTTI for the VIRAL dataset, versus and the published bait sets of sizes 250k, 350k and 700k from CATCH. The dashed line shows coverage 1.0.

sequences. SYOTTI is consistently the fastest by an order of magnitude. MrBait was eventually the most efficient in terms of peak memory, thanks to its SQLite database on disk, but it pays a large price for the disk-based approach in the form of a significantly slower running time.

We observe that very good trade-offs between coverage and number of baits are available by halting the greedy algorithm of SYOTTI before it reaches 100% coverage. Figure 4 shows the coverage as a function of the number of baits selected. For example, taking just the first 200k baits out of 367k already results in 98.8% coverage, whereas taking a *random* subset of 200k baits of the full set resulted in only 62.7% coverage. If long gaps in coverage are undesirable, we can add baits manually to patch the long gaps. For example, if we halt the greedy algorithm at 200k baits and then switch to filling the gaps such that the maximum gap length is 50 bp, we obtain a set of 275k baits with 99.6% coverage and with no gap longer than 50. Figure 4 shows the coverage plot for the 200k bait set after gap filling.

3.3 Results on Viral Strains and Substrains

For our final experiment, we downloaded the set of human viral pangenomes published in the Github repository of CATCH in 2018. We studied scaling of the tools by shuffling the list of files and ran the tools on prefixes containing 1, 2, 4, 8, 16, ..., 512 files, as well as the whole dataset of 608 files.

SYOTTI was again the fastest tool by an order of magnitude, producing 227k baits on the full dataset.

Figure 4: Left: Fraction of nucleotides covered by SYOTTI as the algorithm progresses on the full **BACTERIA** dataset. Right: average coverage of the reference sequences in **BACTERIA** after taking the first 200k baits produced by SYOTTI and filling gaps to maximum length 50. The dashed line shows coverage 1.0.

MrBait again showed slow but approximately linear scaling, but the bait set was 40 times larger than the one produced by SYOTTI. Surprisingly, CATCH itself was unable to process prefixes longer than 1% of the full dataset within the time limit of 24 hours. However, the Github page of CATCH contains pre-designed bait sets with 250k, 350k and 700k baits of length 75 for the dataset. The sets were designed by the authors of CATCH by running the tool separately for each species, and pooling the baits together with different parameters, optimizing the combined bait sets in the process.

To show CATCH in the best possible light, we compared the baits of SYOTTI against the published and optimized bait sets of CATCH. We used SYOTTI to generate baits of length 75, matching the bait length of the published data sets. The pooling method used with CATCH uses different values for the mismatch tolerance for different species to optimize the coverage. As SYOTTI does not support a varying mismatch tolerance, the tolerance of SYOTTI was set to 5, which is the maximum value in the parameters selected in Supplementary Table 1 of the manuscript of CATCH [13]. With these settings, SYOTTI generated 684k baits.

We analyzed the coverage of the bait sets allowing up to 5 mismatches. Coverage of the CATCH bait sets of size 250k, 350k and 700k were 84.1%, 90.6% and 97.5% respectively. The coverage of SYOTTI was 99.5%, with the remaining 0.5% being due to unknown N-characters in the data. The coverage plots of the bait set with 684k baits versus the CATCH bait sets are in Figure 3. Compared to the 700k baits of CATCH, SYOTTI had mostly lower coverage than CATCH (i.e., higher efficiency), except for one notable stretch of the input which finished with 60-fold coverage, where CATCH was able to keep the coverage to only 10-fold. This stretch represents the 59,686 rotavirus A strains in the database. Taking the first 250k baits generated by SYOTTI results in a coverage of 96.5%.

4 Conclusion

In this paper, we provide a formulation of designing baits and demonstrate that the problem is NP-hard even for deceptively simple instances, and provide a heuristic that works efficiently in practice. While both SYOTTI and CATCH use a kind of greedy heuristic, the heuristic of SYOTTI is designed to be much more efficiently implementable. We conclude by suggesting some areas that warrant future work. First, a more sophisticated binding model than the Hamming distance could be plugged into the heuristic at Line 19 of Algorithm 1. Second, while our results rule out any reasonable fixed parameter tractable algorithms, there is potential that the problem admits an approximation algorithm. However, we conjecture that the problem remains hard even for a constant approximation. Third, other artifacts of the laboratory process are left for consideration, including considering the GC-content of the baits; designing baits in a manner that avoids contaminants; and designing baits to avoid bait-to-bait binding. Lastly, computationally designing baits in a manner that models and controls off-target binding is worth consideration.

References

- [1] Food safety and inspection service (FSIS). Serotypes profile of Salmonella isolates from meat and poultry products January 1998 through December 2014. United States Department of Agriculture. 2015. https://www.fsis.usda.gov/sites/default/files/media_file/2020-10/Salmonella-Serotype-Annual-2014.pdf. Retrieved on 28. November 2021.
- [2] T.K. Chafin, M.R. Douglas, and M.E. Douglas. MrBait: universal identification and design of targeted-enrichment capture probes. *Bioinformatics*, 34(24):4293–4296, 2018.
- [3] R. Cole, C.S. Iliopoulos, M. Mohamed, W.F. Smyth, and L. Yang. The Complexity of the Minimum k-Cover Problem. *J. Autom. Lang. Comb.*, 10(5/6):641–653, 2005.
- [4] X. Deng et al. Metagenomic sequencing with spiked primer enrichment for viral diagnostics and genomic surveillance. *Nature Microbiol.*, 5:443–454, 2020.
- [5] B. Devleeschauwer, S. M. Pires, I. Young, A. Gill, and S. E. Majowicz. Associating sporadic, foodborne illness caused by shiga toxin-producing escherichia coli with specific foods: a systematic review and meta-analysis of case-control studies. *Epidemiol. Infect.*, 147:e235, 2019.
- [6] P. Ferragina and G. Manzini. Indexing compressed text. *J. of the ACM*, 52:552–581, 2005.
- [7] National Center for Biotechnology Information. NCBI pathogen detection project [Internet]. *Bethesda, MD. National Library of Medicine (US)*, 2016.
- [8] S. Gog, T. Beller, A. Moffat, and M. Petri. From theory to practice: Plug and play with succinct data structures. In *Proc. SEA*, pages 326–337, 2014.
- [9] A.K. Guitor, A.R. Raphenya, J. Klunk, M. Kuch, B. Alcock, M. G. Surette, A. G. McArthur, H. N. Poinar, and G. D. Wright. Capturing the resistome: a targeted capture method to reveal antibiotic resistance determinants in metagenomes. *Antimicrob. Agents Chemother.*, 64(1), 2019.
- [10] J.S. Lee, R. S. Mackie, T. Harrison, B. Shariat, T. Kind, T. Kehl, M. Löchelt, C. Boucher, and S. Van-deWoude. Targeted enrichment for pathogen detection and characterization in three felid species. *J. Clin. Microbiol.*, 55(6):1658–1670, 2017.
- [11] V. Mäkinen, D. Belazzougui, F. Cunial, and A. I. Tomescu. *Genome-Scale Algorithm Design: Biological Sequence Analysis in the Era of High-Throughput Sequencing*. Cambridge University Press, 2015.
- [12] C. Mayer et al. BaitFisher: a software package for multispecies target DNA enrichment probe design. *Mol. Biol. and Evol.*, 33(7):1875–1886, 2016.
- [13] H.C. Metsky et al. Capturing sequence diversity in metagenomes with comprehensive and scalable probe design. *Nature Biotechnol.*, 37(2):160–168, 2019.
- [14] N.R. Noyes, M.D. Weinroth, J.K. Parker, C.A. Dean, S.E. Lakin, R.A. Raymond, P. Rovira, E. Doster, Z. Abdo, J.M. Martin, K.L. Jones, J. Ruiz, C.A. Boucher, K.E. Belk, and P.S. Morley. Enrichment allows identification of diverse, rare elements in metagenomic resistome-virulome sequencing. *Microbiome*, 5:142, 2017.
- [15] Z. Palkova. Multicellular microorganisms: laboratory versus nature. *EMBO Rep.*, 5(5):470–476, 2004.
- [16] T. Rognes, T. Flouri, B. Nichols, C. Quince, and F. Mahé. VSEARCH: a versatile open source tool for metagenomics. *PeerJ*, 4:e2584, 2016.
- [17] S. Rubiola, F. Chiesa, A. Dalmaso, P. Di Ciccio, and T. Civera. Detection of antimicrobial resistance genes in the milk production environment: Impact of host dna and sequencing depth. *Front. Microbiol.*, 11:1983, 2020.

- [18] Fei Shi. Suffix arrays for multiple strings: A method for on-line multiple string searches. In Joxan Jaffar and Roland H. C. Yap, editors, *Concurrency and Parallelism, Programming, Networking, and Security: Second Asian Computing Science Conference, ASIAN '96, Singapore, December 2-5, 1996, Proceedings*, volume 1179 of *Lecture Notes in Computer Science*, pages 11–22. Springer, 1996.
- [19] F. Sievers and D.G. Higgins. Clustal omega. *Curr. Protoc. Bioinformatics*, 48(1):3–13, 2014.
- [20] R. Zaheer et al. Surveillance of Enterococcus spp. reveals distinct species and antimicrobial resistance diversity across a One-Health continuum. *Sci. Rep.*, 10(1):1–16, 2020.

A Details on how the Bacterial Dataset was Constructed

The pathogen set used was selected to represent whole genomes of public-health relevant foodborne isolates identified in food, animals, and human patients (clinical) using the NCBI Pathogen Detection tool [7], a repository developed for the curation of pathogenic bacterial genomic sequences. Salmonellae genomes were chosen to represent the serovars of greatest public health concern [1], including Enteritidis, Typhimurium, Infantis, Kentucky, Montevideo, and Dublin, representing an initial pool of 76,638 genomes and 3,486 SNP clusters. Shigatoxigenic Escherichia coli (STEC) isolates were chosen to include the O157:H7 serovars as well as the non-O157 STEC group (O26, O45, O103, O111, O121, and O145) including H-untypable, H-pending, H-undetermined, and non-motile types [5]. The resulting initial pool of STEC isolates represented 4,964 genomes and 1,164 SNP clusters. No specific subtype or serovar was chosen for Campylobacter jejuni (initial pool representing 60,227 genomes and 4,591 SNP clusters) or the Enterococci [20] (initial pool of 24,221 genomes and 1,774 SNP clusters of faecium and faecalis).

All isolates were then further filtered according to location (known U.S. and U.S. territories only), source (either from a clinical sample or from a domesticated animal or fowl collected pre- or post-slaughter, pre- or post-processing, and / or pre- or immediately post-fabrication stage in the food production system), collection date (2018–2020), sequence identity (i.e., availability of isolate-level WGS accession), and availability of SNP cluster assignment. Isolates were excluded if they originated from environmental samples (e.g., soil, water, sewage), ready-to-eat (RTE) products, or if any metadata variables were designated as ‘unknown’, ‘pending’, or left blank.

Stratified random sampling was conducted to select 1,000 total isolate genomes representative of the relative proportions of genera and serotypes in the final filtered pool of isolates (Salmonella enterica ssp. enterica: 322; STEC: 60; Campylobacter jejuni: 580; and Enterococci: 42). Corresponding WGS assemblies for all 1,000 isolates containing either a GCA or GCF prefix were retrieved from NCBI using the EDirect tool, and integrity of transferred data was ensured by evaluating file md5 checksum values.

B Full Experimental Data Tables and Figures

All coverage plots, all time/mem plots. Full tables in case we delete some rows corresponding to the very smallest data prefixes.

	SYOTTI			CATCH			MrBait		
Input length	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits
498	00:00	5	5	00:00	49	5	00:01	71	4
1320	00:00	5	12	00:00	49	12	00:00	71	10
2937	00:00	5	26	00:00	50	26	00:00	71	22
6672	00:00	5	59	00:00	50	59	00:00	71	52
14244	00:00	5	115	00:00	52	118	00:00	71	111
31758	00:00	5	249	00:00	56	253	00:01	71	249
60670	00:00	5	434	00:00	63	437	00:01	71	477
125271	00:00	6	820	00:01	78	824	00:02	72	987
254294	00:00	8	1372	00:04	104	1392	00:04	73	1981
508459	00:00	11	2604	00:09	158	2635	00:09	76	3983
1032802	00:00	20	4633	00:28	260	4742	00:19	80	8093
2090517	00:00	35	7901	01:30	423	8121	00:37	88	16374
4187569	00:01	67	13099	05:37	735	13489	01:12	101	32764
8106325	00:03	125	20976	19:13	1250	21771	02:19	128	63428

Table 4: Running time, memory usage and number of baits for increasing larger subsets of the MEGARES dataset. Times are in the format mm:ss.

Figure 5: Memory scaling on all three datasets.

Figure 6: Bait set size scaling on all three datasets.

	SYOTTI			CATCH			MrBait		
Input length	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits	Time	Memory (MB)	Baits
16794	00:00:00	5	140	00:00:00	52	140	00:00:01	71	139
27058	00:00:00	5	226	00:00:00	55	226	00:00:00	71	224
164612	00:00:00	7	1373	00:00:02	95	1373	00:00:03	72	1369
242069	00:00:00	7	2017	00:00:03	112	2017	00:00:05	73	2012
344651	00:00:00	9	2875	00:00:05	146	2875	00:00:06	74	2862
991610	00:00:00	19	8189	00:00:29	312	8264	00:00:19	80	8244
1572067	00:00:00	27	12821	00:01:01	485	12975	00:00:30	84	13067
3470929	00:00:02	55	25351	00:03:47	1022	27378	00:01:07	98	28857
8038466	00:00:04	122	48658	00:13:27	2158	56069	00:02:35	133	66861
15920441	00:00:08	237	72035	00:29:32	4096	90182	00:04:48	195	132408
30786116	00:00:14	454	96595	01:00:27	7575	134112	00:10:00	316	256041
62502135	00:00:27	917	123541	02:26:56	14234	181825	00:18:44	572	519838
125063199	00:00:52	1831	157818	08:43:14	24269	240747	00:38:48	1076	1040174
254576853	00:01:45	3723	188813	>24h	NA	NA	01:20:03	2071	2117436
505422833	00:03:31	7387	223931	NA	NA	NA	02:39:03	4056	4203752
1003934029	00:07:11	14673	267890	NA	NA	NA	05:18:28	7980	8349888
2018459352	00:15:50	29496	324797	NA	NA	NA	10:30:37	15697	16788084
3040260476	00:24:52	44425	366761	NA	NA	NA	16:13:43	23615	25286576

Table 5: Running time, memory usage and number of baits for increasing larger subsets of the BACTERIA dataset. The times are in the format hh:mm:ss.

Figure 7: MEGARES coverage with VSEARCH filtering for MrBait.