A system capable of verifiably and privately screening global DNA synthesis

Carsten Baum[1,2], Jens Berlips[3], Walther Chen[3], Helena Cozzarini[3], Hongrui Cui[4], Ivan Damgård[1,*], Jiangbin Dong[5], Kevin M. Esvelt[3,6,*], Leonard Foner[3], Mingyu Gao[5,12], Dana Gretton[3,6], Martin Kysel[3], Juanru Li[4], Xiang Li[5], Omer Paneth[7], Ronald L. Rivest[7], Francesca Sage-Ling[3] , Adi Shamir[8], Yue Shen[10], Meicen Sun[11], Vinod Vaikuntanathan[7], Lynn Van Hauwe[3], Theia Vogel[3], Benjamin Weinstein-Raun[3], Yun Wang[10], Daniel Wichs[9], Stephen Wooster[3], Andrew C. Yao[3,5,12,*], Yu Yu[4,12], Haoling Zhang[10], and Kaiyi Zhang[4]

[1]Department of Computer Science, Aarhus University, Aarhus Centrum 8000, Denmark;
[2]DTU Compute, Technical University of Denmark, Kongens Lyngby 2800, Denmark;
[3]SecureDNA Foundation, Basel 4010, Switzerland;
[4]Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China;
[5]Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China;
[6]Media Lab, Massachusetts Institute of Technology, Cambridge, MA 02142, USA;
[7]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA;
[8]Department of Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel;
[9]Department of Computer Science, Northeastern University, Boston, MA 02115, USA;
[10]BGI Research, Changzhou 213299, China;
[11]Department of Political Science, Massachusetts Institute of Technology, Cambridge, MA 02142, USA;
[12]Shanghai Qi Zhi Institute, Shanghai 200232, China

The authors are listed in alphabetical order by surname.

* Corresponding authors. E-mails: ivan@cs.au.dk; andrewcyao@mail.tsinghua.edu.cn; esvelt@mit.edu

ABSTRACT

Printing custom DNA sequences is essential to scientific and biomedical research, but the technology can be used to manufacture plagues as well as cures. Just as ink printers recognize and reject attempts to counterfeit money, DNA synthesizers and assemblers should deny unauthorized requests to make viral DNA that could be misused. There are three complications. First, we don't need to quickly update printers to deal with newly discovered currencies, whereas we regularly learn of new potential pandemic viruses and other biological threats. Second, convincing counterfeit bills can’t be printed in small pieces and taped together, while preventing the distributed synthesis and subsequent re-assembly of controlled sequences will require tracking which DNA fragments have been ordered across all providers and benchtop devices while protecting legitimate customer privacy. Finally, counterfeiting can at worst undermine faith in currency, whereas unauthorized DNA synthesis could be used to deliberately cause pandemics. Here we describe SecureDNA, a free, privacy-preserving, and fully automated system capable of verifiably screening all DNA synthesis orders of 30+ nucleotides against an up-to-date database of controlled sequences, and its operational performance and specificity when applied to 67 million nucleotides of DNA synthesized by providers in the United States, Europe, and China.

Keywords: DNA synthesis screening, cryptography, privacy

INTRODUCTION

Custom DNA synthesis is foundational to biomedical research, underpinning everything from cancer immunotherapies to SARS-CoV-2 vaccines. However, the same technology can also be used to produce pathogens [1-3] (Fig. 1a). The first infectious virus to be assembled from synthetic DNA was generated in 2002 at a cost of more than \$10 per nucleotide [4]. Just over two decades later, the price has fallen more than a thousandfold, the number of individuals with the necessary skills has grown from dozens to many thousands [5], and a pandemic has directly and indirectly killed over 20 million people [6], demonstrating the scale of harm that could be inflicted by a virus generated from synthetic DNA.

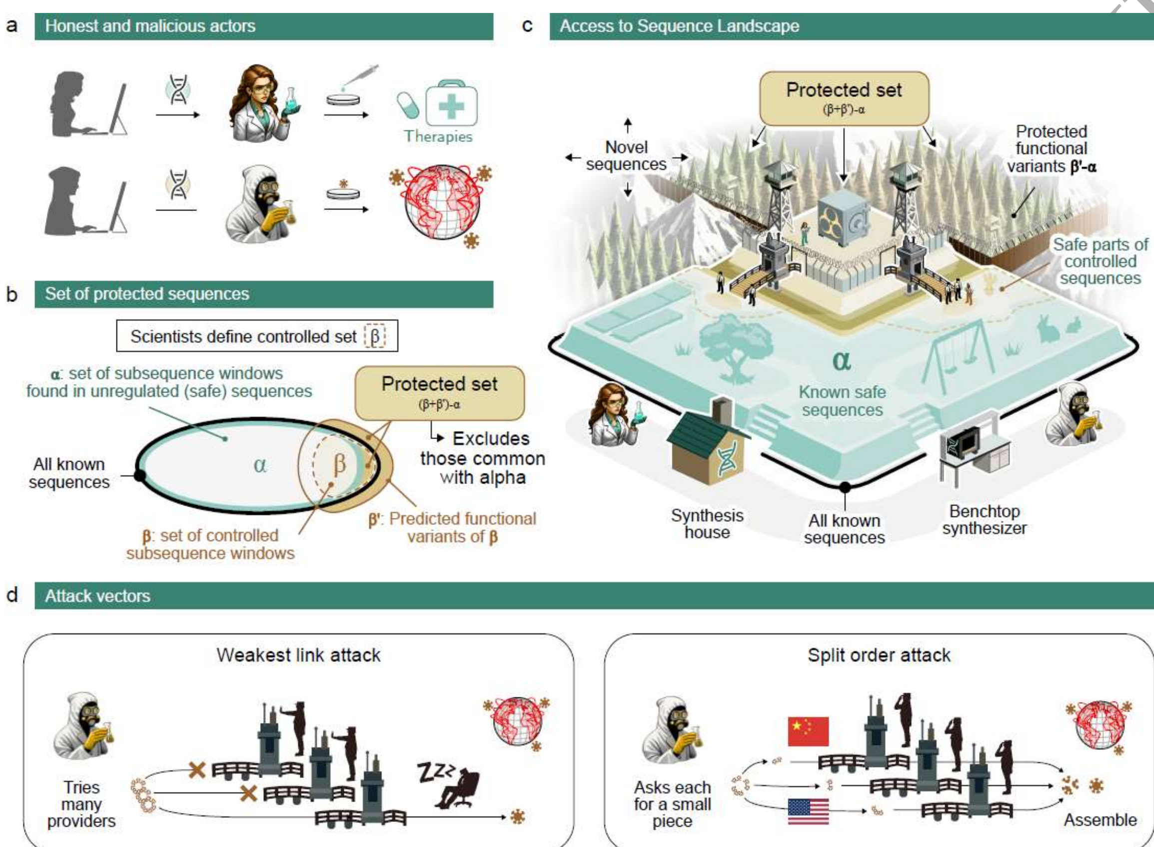

Figure 1. DNA synthesis screening ecosystem and security framework. (a) DNA synthesis enables therapeutic development (top), but also presents biosecurity risks (bottom). (b) DNA screening framed as a set-membership problem using configurable sets. α represents constant-length subsequences found in non-controlled (safe) sequences, including most organisms and common molecular biology tools. Scientists and regulators define the safe set α, a controlled set β, and compute functional variants β' using any suitable variant prediction method (e.g., MCMC sampling over predicted mutation tolerance scores; see [19]). The protected set (β+β')–α comprises subsequences from controlled sequences, or their variants, that do not appear in safe sequences. (c) Screening must cover both synthesis houses and benchtop synthesizers. Orders containing protected-set sequences require authorization. (d) A weakest link attack bypasses screening by finding a provider with inadequate security. A split-order attack evades screening by splitting controlled sequences into fragments that are either too short to screen or appear harmless in isolation. Both underscore the need for globally adopted screening of all assembly-viable sequence lengths.

While no credible and accessible pandemic viruses are publicly known, numerous well-intentioned research programs aim to identify viruses capable of causing new pandemics and share their genome sequences [7-11]; one maintains a public list of viruses ranked by threat level [12]. Future discoveries and advances in biological programming will identify other pandemic-class agents that can be generated from synthetic DNA. Because state actors can independently synthesize any DNA, the central security challenge is preventing proliferation.

The obvious response to this challenge—screening all synthetic DNA orders for a list of controlled sequences—was first recommended in 2006 [13]. Remarkably, members of the International Gene Synthesis Consortium (IGSC) voluntarily monitor an estimated 80% of global DNA synthesis, even though the relative cost of human-assisted screening is rising with volume [14,15].

Nevertheless, synthesis screening remains an unsolved global problem:

(1) 36 out of 38 providers shipped DNA fragments collectively sufficient to generate the 1918 pandemic influenza virus because they had no way to know that others were supplying the remaining fragments [16];
(2) Chatbots now suggest ways to evade screening, including ordering from non-IGSC providers [17];
(3) Oligonucleotides under 50 nucleotides can be assembled and used to generate pandemic threats [14];
(4) Scientific consensus on which sequences warrant control varies across regions and evolves over time;
(5) Most benchtop devices do not screen, on-device screening is insecure, and cloud-based screening [18] is not private.

The DNA industry is also at risk:

(1) Firms may be liable because they cannot verify that they performed best-in-practice screening;
(2) Rising demand may overwhelm expert screening with false alarms and costs may become prohibitive [14].

Solving these problems requires a free, centralized, and privacy-preserving screening system configurable for any defined set of controlled sequences and effective for synthesis houses and benchtops (Fig. 1b-c).

Current screening systems rely on fuzzy-match algorithms that compare orders to both controlled and unregulated sequences, attempting to determine which they more closely resemble. This approach faces fundamental limitations: it incorrectly flags more and more unregulated sequences as the sequence window length decreases, it requires human expert review to resolve these ambiguous matches and consequently can't be used by benchtop DNA synthesis devices, it can be evaded by introducing mutations that make controlled sequences appear more similar to unregulated genes, and it cannot reliably detect split orders distributed across providers in different nations (Fig. 1d).

Overcoming these limitations requires a new approach that can address two design challenges:

(1) *Bio-design*: translate the biological problem of sequence recognition into a computer science problem;
(2) *System-design*: implement an automated system capable of verifiably screening global DNA synthesis without delaying research or compromising customer or database privacy.

In a companion paper [19], we describe a candidate solution to the first challenge: searching for exact matches to short subsequences that are unique to controlled genes and functional variants. By identifying these diagnostic fragments in advance and verifying that they don't appear in unregulated sequences, it can achieve perfect specificity among known sequences (Fig. 1b-c). This precision eliminates the need for human expert review, while the inclusion of functional variants renders screening robust against evasion attempts.

Here we address the system design challenge (Fig. 2) using a novel application of cryptography to enable centralized and privacy-preserving detection. Our automated implementation, including a graphical user interface, screens at high speed and low cost while demonstrating very high specificity, and implements a certificate system permitting authorized laboratories to seamlessly access controlled DNA.

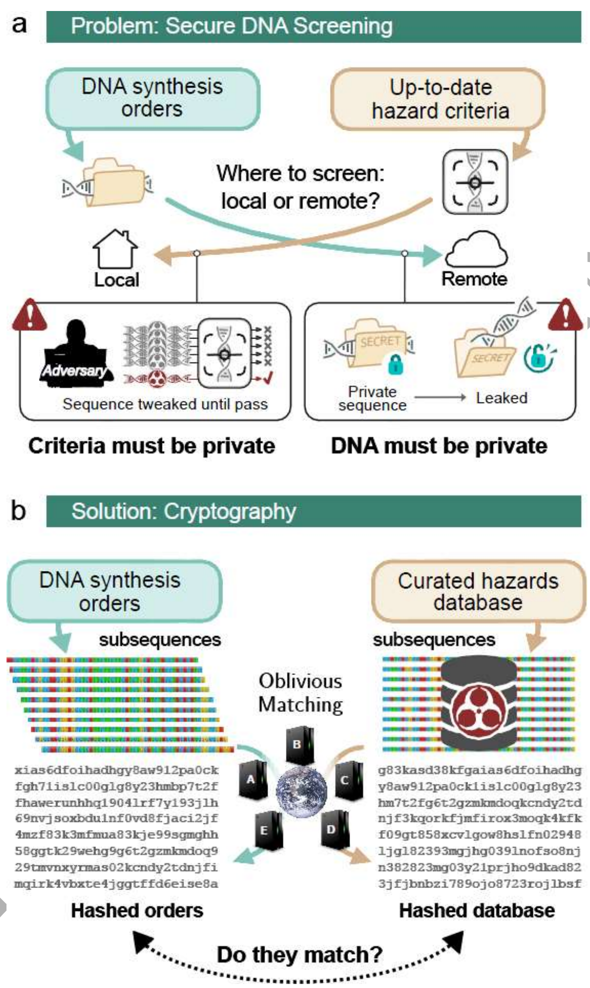

Figure 2. Securing DNA synthesis screening. (a) Secure and universal DNA synthesis screening requires a way to verifiably determine whether DNA synthesis orders correspond to controlled biological functions, including controlled sequences split across multiple providers and benchtop devices, without disclosing anything else about the orders. Disclosing a private customer order may compromise trade secrets, while leaking the criteria makes it possible to evade screening by introducing mutations that make a controlled sequence appear to be unregulated. These constraints cannot be satisfied simultaneously. (b) The SecureDNA system allows synthesizers and database contributors to *obliviously* perform one-way transformations of their subsequence windows, which can be directly compared. The database provides the synthesizer with a timestamped

verification that *n* windows sent by the synthesizer were screened against a particular database version and can detect split-order attacks if adopted by most providers.

RESULTS

Construction of a Controlled Sequence Database

Scientists and regulators designate a list of DNA and peptide sequences as controlled, and construct a database of such controlled sequences. These sequences are split into overlapping windows of constant length (e.g., the first 42 characters of the text representation of a pathogen genome, then those ranging from the 2nd to the 43rd, and so on); these are defined as the controlled set β. From some of these windows known or predicted to have significant biological functions, up to 10 million variants (changes to a few characters at a time) are generated that are predicted to be at least somewhat functional, comprising the variant set β′. These are compared against a much larger set of similar constant-length sequences known to appear in nature or research in non-controlled sequences, the safe set α. The database is populated with those windows and functional variants of controlled sequences that do not appear in unregulated sequences, $(\beta + \beta') - \alpha$, or, in set notation, $(\beta \cup \beta') \setminus \alpha$.

Theoretical Crypto-Design and Analysis

Suppose that a Customer places an order for sequence *s* with a Synthesizer, who wants to know whether it's safe to make, without disclosing *s*—which may be a trade secret—to any eavesdroppers (Fig. 3a). The Synthesizer can ask an up-to-date remote Database *D* of controlled DNA and peptide subsequences (Fig. 3b; Appendix A) to check whether any of the constant-length DNA and translated peptide subsequences found in *s* are present in *D*. The crypto-design challenge is to find a way for the Database to answer without 1) learning anything consequential about *s* or 2) revealing anything about *D* beyond conveying the yes/no answer, even if some of the parties among the Customer and Synthesizer are compromised.

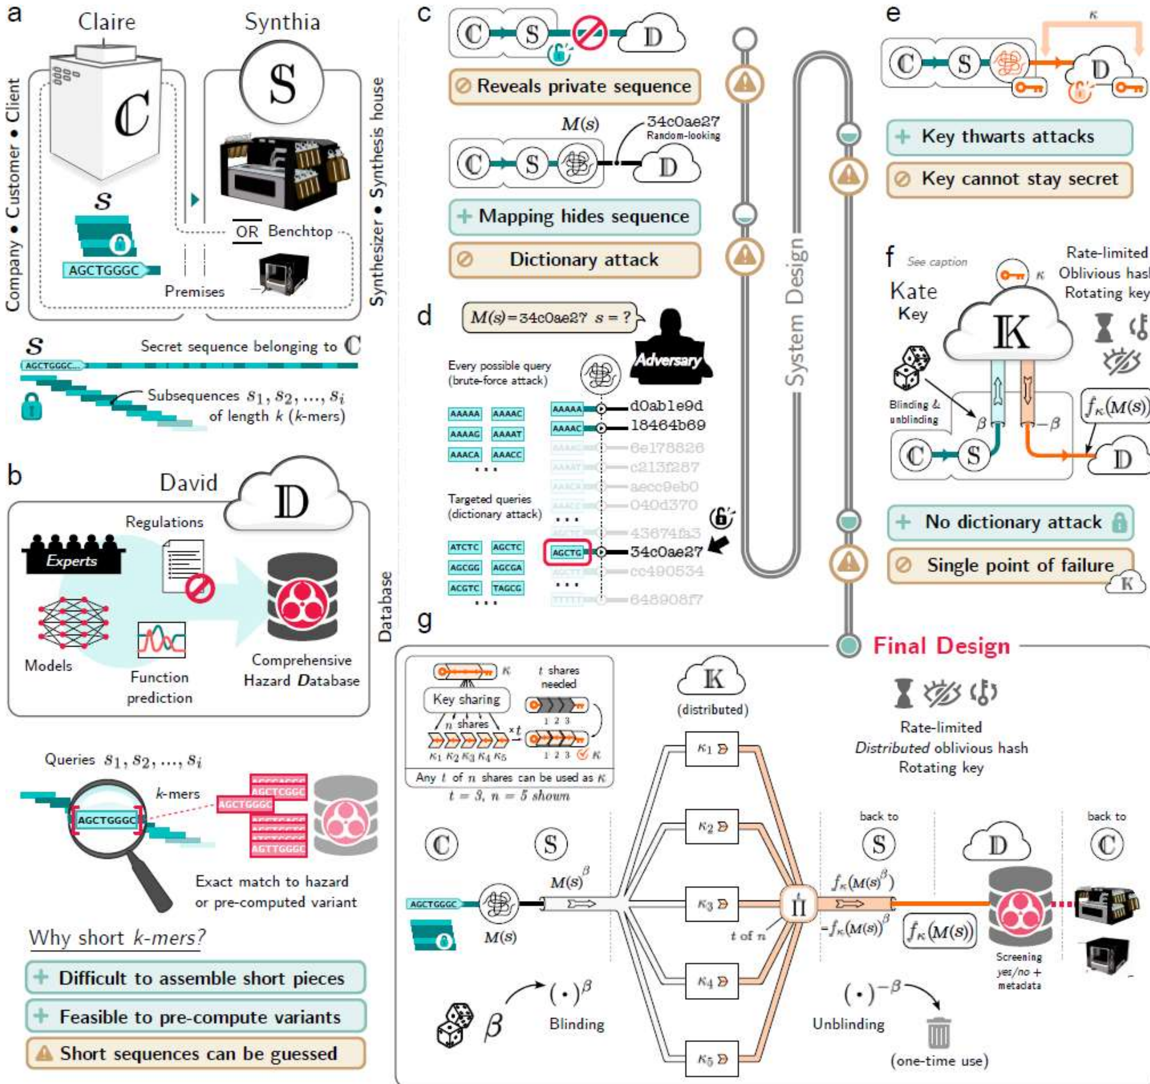

Figure 3. Cryptographic challenges and solutions for secure DNA synthesis screening. (a) Client Claire orders DNA from Synthesizer Synthia. A sequence s is split into all subsequences of length k (k-mers). (b) Database David contains k-mers from controlled sequences and predicted functional variants. Synthia sends k-mers to be screened, and David detects exact k-mer matches. Subsequence length k is small to prevent assembly of unscreened fragments. (c) Synthia cannot send s directly to David (as in current remote DNA screening systems) without violating Claire's privacy, but a cryptographic mapping $M(s)$ enables secure comparison. (d) Since inputs are known to be short, a standard hash is easily cracked using a brute-force or dictionary attack. (e) A keyed hash with a secret key κ can thwart dictionary attacks, but requires that κ be known to both Synthia and David, in which case Synthia could interrogate D and David could crack s. (f) Keyserver Kate helps Synthia compute keyed hash $f_\kappa(M(s))$ with oblivious cryptographic hashing. Synthia "blinds" the sequence s with a random value β before Kate applies κ, then unblinds it afterward, so no eavesdropper (including Kate) can learn s. Kate can limit the rate of evaluation of f_κ, thwarting dictionary attacks. However, Kate is a single point of failure: an adversary compromising Kate gains κ, and if Kate is offline, global synthesis stops. (g) In the final design, $f_\kappa(M(s))$ is distributed, making Kate's role robust. κ is split into n key shares among n separated servers. Evaluating $f_\kappa(M(s))$ requires a threshold number t of the n shares (n = 5, t = 3 shown). Up to (n – t) keyservers may be offline. κ is rotated, e.g., biweekly. Any adversary must compromise t servers simultaneously between key rotations to steal κ.

The simplest way for the Synthesizer and Database to privately determine whether any of their DNA or peptide subsequences are identical is to transform them using a one-way hash mapping M and compare the hash outputs $M(s)$ and $M(D)$ (Fig. 3c). But either of them could rapidly enumerate all possible short subsequences, letting them decrypt s

and D (Fig. 3d), even when using a shared key κ (Fig. 3e). To restore privacy, we introduce the Keyserver, who helps compute hashes using a frequently-changed key κ while rate-limiting requests to match plausible DNA synthesis speeds, which renders enumeration attacks infeasible (Fig. 3f). The Synthesizer and Keyserver can compute hash outputs *obliviously* using cryptographic techniques: the Keyserver learns nothing about the input sequences or resulting hashes, while the Synthesizer learns only the hash outputs.

To screen for matches, the Database is given a table H comprising the oblivious keyed hashes of all elements in the plaintext database D, without learning its contents. Whenever the Synthesizer submits a hashed query, the Database can convey whether there are any matches to H without learning anything about the sequences apart from whether they match. This problem can be viewed as a special case of unbalanced private set intersection (PSI), or more precisely, a private membership testing (PMT) problem, where the querier holds a single sequence and the server holds a large set of sequences.

The Database provides the Synthesizer with a timestamped signature verifying that screening was performed against a specific version of H, allowing providers to demonstrate compliance with legal requirements and minimize liability. In addition, the Database can maintain privacy-preserving logs of which controlled subsequence hashes have been detected over time, which enables detection of split-order attacks when fragments of a controlled sequence are ordered across multiple providers. Overall, this architecture ensures that no single party needs to know the list of controlled sequences or which subsequence windows are defended with functional variants—a critical feature, as knowledge of defended windows would enable evasion. In the future, the system could allow nations to screen for emerging threats without publicly acknowledging them [19].

To summarize the baseline system:

(1) Scientists and Keyserver build and update Database H of obliviously hashed subsequences;
(2) Customer orders sequence s from Synthesizer;
(3) Synthesizer and Keyserver obliviously hash subsequences;
(4) Synthesizer sends hash outputs to Database;
(5) Database tells Synthesizer if any are found in H;
(6) If there are matches and Customer lacks authorization, Synthesizer refuses to make s.

Security Analysis of Baseline Screening

In the *semi-honest* model of cryptographic security, all parties follow the protocol, but may eavesdrop. Oblivious hashing ensures the Keyserver learns nothing about sequences or hashes, while the Synthesizer learns only whether matches exist—the minimum necessary for screening. If no controlled sequences are found, the Database learns only the total sequence length (which is biologically irrelevant) and possibly whether any subsequences are shared with other clients' queries. If s contains controlled sequences, the Database learns which subsequences match the controlled set from D, but cannot determine any other sequences in s or D.

In the malicious model, a party may deviate from the protocol to eavesdrop on or sabotage its execution. A fundamental premise, however, is that the Synthesizer seeks to avoid creating controlled sequences: no software-only mechanism can prevent a synthesizer that is not hardware-locked from generating controlled DNA. Commercial synthesis providers are therefore assumed to be incentivized to ensure that their internal infrastructure faithfully follows the screening workflow in order to comply with legal requirements and minimize liability. For benchtop synthesizers, stronger guarantees require hardware enforcement, such as ensuring that synthesis cannot begin until a valid screening result is returned and that the control path is protected using tamper-resistant mechanisms (e.g., techniques meeting NIST FIPS 140-2 Security Level 3 or 4 standards [20,21]). While a dishonest Synthesizer could disclose the query sequence s, rate-limiting prevents reconstruction of the Database D through repeated queries. A corrupt Database could disclose H and associated statistical correlations, but cannot interrogate s or recover D without assistance from the Keyserver (see Appendix B for a detailed analysis of information leakage). Conversely, a compromised Keyserver could use κ to hash all possible subsequences and thereby recover information about s and D.

Distributing the Key

The integrity of the Keyserver is so vital that we divide the key κ into n shares and distribute them across n Keyservers using Shamir secret sharing [22] to avoid a single point of failure (Fig. 3g). This architecture requires a threshold t Keyservers to jointly perform oblivious hashing: any subset of t Keyservers can jointly compute the keyed hash with Synthia, but any smaller coalition learns nothing about κ. The SecureDNA PMT protocol's main efficiency gain over traditional PSI comes from introducing this set of trusted third parties who pre-process parts of the computation offline. To prevent a malicious party from stealing κ by gradually corrupting a majority of the Keyservers, the system periodically updates key shares through a coordinated process with the same security guarantees as distributed oblivious hashing—no single party ever reconstructs κ. This achieves what cryptographers call *proactive security*: κ, and thereby all s and D, remain protected even if all Keyservers are eventually hacked as long as they are not compromised simultaneously. See Appendix C for the complete protocol.

System Design and Implementation

The database is constructed by partitioning controlled sequences into 30 or 42 nucleotide DNA windows and 20 amino acid peptide windows (β), which are tagged with geographical regions in which they are regulated. For a subset of windows, chosen pseudo-randomly, we stochastically include mutants and functional variants predicted using protein structure and sequence homology (β').

To ensure specificity among all known sequences, we check each candidate subsequence against a dataset of nucleotides and proteins derived from NCBI's nr/nt GenBank and GenPept databases and remove any that are also found in an unregulated sequence (i.e., remove any in α). This filtering process combined taxonomic analysis, keyword identification (e.g., "synthetic" or "recombinant"), and controlled sequence coverage metrics to improve database annotations. An analysis of controlled viral sequences revealed that nearly all contained thousands of k-mers ($k \geq 30$) unique to strains of the controlled virus, providing ample targets for effective screening.

The system can accommodate any arbitrarily defined set of controlled sequences, permitting customization by provider preference. By default, it detects controlled sequences defined in Australia Group, ITAR, or national regulations, viruses described as potential pandemic pathogens in the scientific literature, and viruses capable of human-to-human transmission or illness that could be misused. While the system returns an alert upon any unique match, it recommends denying synthesis for controlled viruses, potential pandemic viruses, genes encoding controlled toxins, and sequences that could confer controlled-pathogen capabilities to more accessible strains per the Australia Group's "endow or enhance pathogenicity" definition.

Split-Order Detection

Even modestly resourced actors can evade screening by splitting controlled sequences into multiple individually harmless fragments and ordering them from different providers, who have no way of knowing whether others are making the remaining pieces [16]. SecureDNA's primary defense against such split-order attacks is fragment-level screening at short sequence lengths, which substantially constrains assembly-viable evasion strategies. However, detecting accumulation of fragments associated with a particular controlled sequence across globally distributed providers remains important for identifying coordinated misuse attempts. Even if governments required logs of all controlled fragments ordered from their providers, nations would need to share data with one another to detect such attacks, which appears unlikely given diplomatic constraints. Because it is privacy-preserving, offered free of charge by a nonprofit in a neutral country, co-developed internationally, and can be used by centralized synthesis houses and benchtop synthesizers, SecureDNA could potentially be adopted widely enough to detect such attempts by monitoring which pieces have been ordered from different suppliers and notifying either suppliers or law enforcement of detected attacks, as appropriate. Providers can then strictly require evidence of authorization for any requests for the remaining pieces, including pausing any in-progress orders, thereby denying access to the adversaries attempting the split-order attack. To enable split-order detection, the system stores hashes of controlled subsequences in order and logs which ones have been detected, while learning nothing of which customer ordered which k-mers, only which provider they ordered it from. Statistical accumulation patterns over time can then be used to support escalation or trigger additional scrutiny, without requiring customer identity matching.

System Operation

Each synthesizer or provider running the open-source SecureDNA *synthclient* program converts FASTA files from sequence orders into all eligible subsequences and translations (assuming the canonical genetic code; alternative genetic codes do not enable access to transmissible pandemic agents). Screening proceeds in both forward and reverse-compliment directions across the query sequence [19] and, for benchtop synthesizers, includes base permutations to prevent evasion through reagent bottle swapping (Fig. S1). This step negligibly reduces specificity and actually improves performance, as all 48 possible hash combinations (including all possible base permutations in both directions) are consolidated. After local processing, *synthclient* contacts the keyservers to perform oblivious hashing and sends the results to the database for comparison. Screening completes faster than most web pages load (Fig. 4a), providing an immediate 'accepted', 'alert', or 'denied' decision along with a signed, timestamped verification that screening was performed using a specific database version. The system also generates a visual mapping of any matches to public controlled sequences and notes relevant export control restrictions (Fig. 4b).

Deployment and Performance

The production system operates with five keyservers distributed across multiple geographic zones to maximize fault tolerance, requiring any three to process requests. Under normal load, the system completes key rotation every 1-2 weeks, with the capacity to recompute the entire database H in three days if needed. The primary computational bottleneck remains elliptic curve group operations, requiring 10-1000 microseconds per window depending on architecture.

Despite modest computational resource requirements (~$5-8K/year, Methods), the system demonstrates robust performance. *Synthclient* processes 1100 nucleotides per second per CPU thread, or about 2200 nt/second per hyperthreaded core, on provider hardware. Keyservers achieve 4400 nt/second on 4-thread CPUs. The database, using cloud SSD storage, handles 41,500 nt/second per disk, and over 74,000 nt/second in a physical machine with NVMe storage, with constant-time lookup regardless of database size up to 15x the size of the current database. Most providers need only a single $2,000 desktop computer or can screen for pennies per day using on-demand cloud servers, while benchtop synthesizers can operate with a $50 Raspberry Pi 4B (Table S1). For customers with very large oligo-manufacturing needs, we have demonstrated 40 million nucleotides in under 10 minutes at a customer cost of 61 US cents for cloud-based clients. The current hardware configuration can screen more than 4 trillion nucleotides annually—sufficient for projected global synthesis volume through 2029. As SecureDNA expands, we plan to increase security and redundancy by adding servers (Appendix D). Operating fees estimated at $15-30K/yr (depending on degree of redundancy) for servers housed in a colocation facility constitute the only recurring usage cost borne by SecureDNA (Appendix D), and are sufficient to support an order of magnitude more screening than the estimated size of the current market for both genes and oligos.

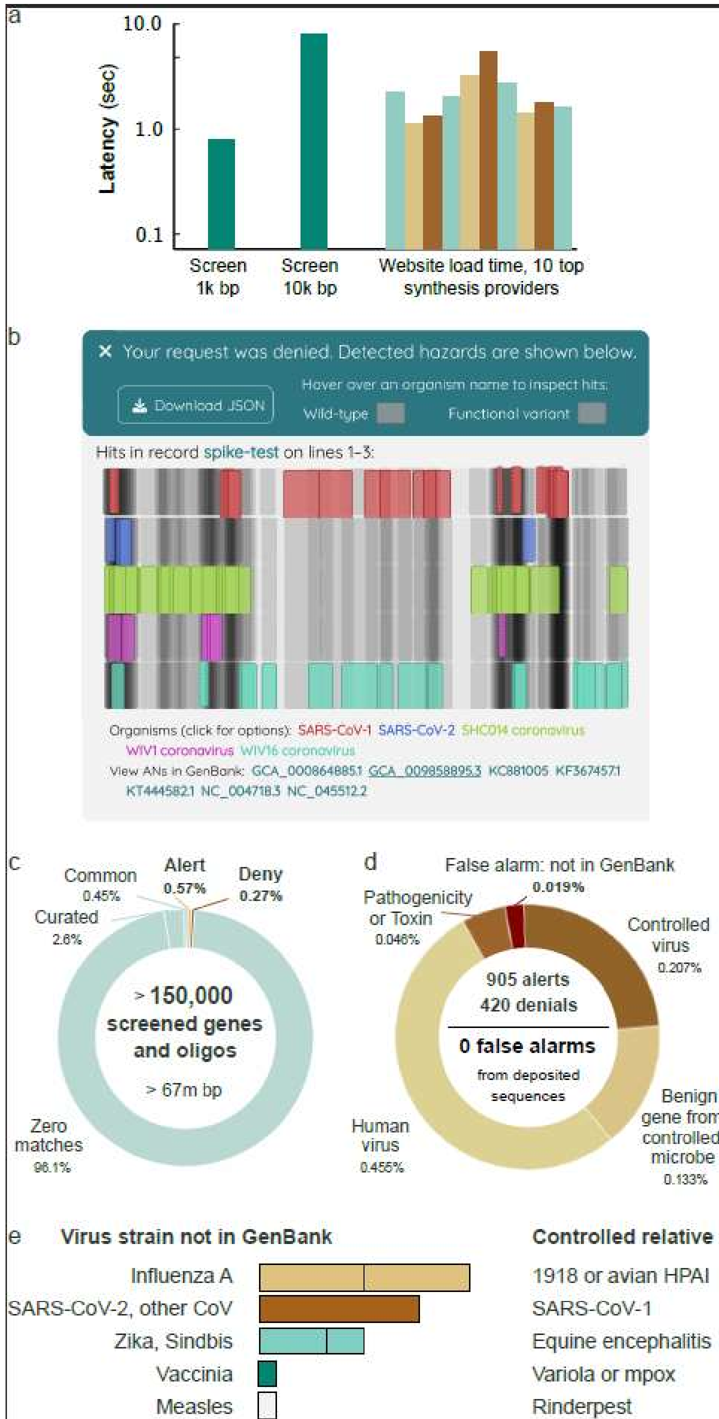

Figure 4. SecureDNA latency, output and specificity on real-world synthetic genes. (a) A comparison of the time required to screen 1,000- and 10,000-bp orders to the typical provider website response times. SecureDNA does not meaningfully delay order placement and can return results to customers in real-time. (b) Output showing matches to wild-type and variant windows from controlled pathogens with hover-over interface for inspection, including GenBank references and flagged sequences from human viruses and benign genes from regulated bacteria. The depicted chimera was denied because it includes *k*-mers unique to SARS-CoV-1. (c) Analyzing >67 million nucleotides of synthesized genes from multiple providers with SecureDNA flagged 0.57% of genes (yellow) and denied 0.27% (brown). Fragments present in both controlled (β or β') and safe (α) sequences are Curated, while sequences unique to β but often used in biotechnology (e.g., P2A peptides) are Common. Both represent matches but not alerts. (d) Iterative manual BLAST revealed zero false alarms from known sequences: all 30 false alarms were from strains of related viruses not present in GenBank. (e) Probable identities of the 30 false alarms (1 in 5,000 orders) flagged due to sequence windows not present in GenBank.

Quantifying Specificity using Real-World Synthesis Data

To measure specificity in a real-world context, we screened anonymized orders for over 150,000 genes and oligonucleotides synthesized by providers in the United States, Europe, and China, which together comprise over 61 million nucleotides generated for biological research and 6 million for DNA storage.

Over 99% of sequences passed screening, and an overwhelming majority (96%) featured zero matches. Among sequences with one or more matches (4%), two-thirds (2.6%) were found in unregulated sequences and were not flagged (Fig. 4c). One-third of those remaining (0.45%) were flagged as "Common" because they matched one of the many frequently used sequences in biotechnology that originated in a regulated organism or a human-infecting virus. Examples include the porcine teschovirus 2A peptide and the cytomegalovirus enhancer/promoter.

SecureDNA identified genes with at least one nucleotide or peptide match to a controlled organism or human virus in 0.85% of sequences (Fig. 4c). Each was investigated by running BLAST on the entire gene sequence, then on any subsequences that were not clearly identified in the first attempt, and manually labeling (Fig. 4d). Approximately half (0.455%) were from unregulated viruses capable of infecting humans, for which several providers have requested alerts; others (0.133%) matched benign genetic sequences from regulated microbes that could not confer pathogenicity upon more easily obtained related or avirulent strains. With respect to the controlled sequences included in the evaluation corpus (e.g., U.S. Select Agents and Australia Group pathogens), we observed zero false negatives.

Absent an exemption token (see below), we would have recommended denying just 0.272% of all ordered genes: 0.046% encoded a toxin or conferred pathogenicity upon a more easily obtained avirulent microbe, 0.207% were unique to a controlled or potential pandemic virus, and 0.019% (30 gene fragments) were from strains of unregulated human-infecting viruses not deposited in Genbank (Fig. 4d), all related to controlled viruses.

Crucially, we observed zero false alarms from known sequences in GenBank: the system perfectly distinguished between all known sequences. The 30 false alarms from sequences absent from GenBank were all from strains of related human-infecting viruses or deliberately created mutants that had never been deposited. We estimate that at current DNA synthesis rates, this false alarm rate would correspond to just ~4 wrongly denied sequences globally each day. By comparison, we anticipate approximately 50 legitimate authorization requests daily for controlled sequences, meaning more than 90% of all authorization requests would be for actual controlled sequences rather than false alarms. Nevertheless, it is imperative for authorized laboratories that work with these pathogens—or with controlled agents—to gain access to the DNA that they need to advance human knowledge and develop therapies, without delay.

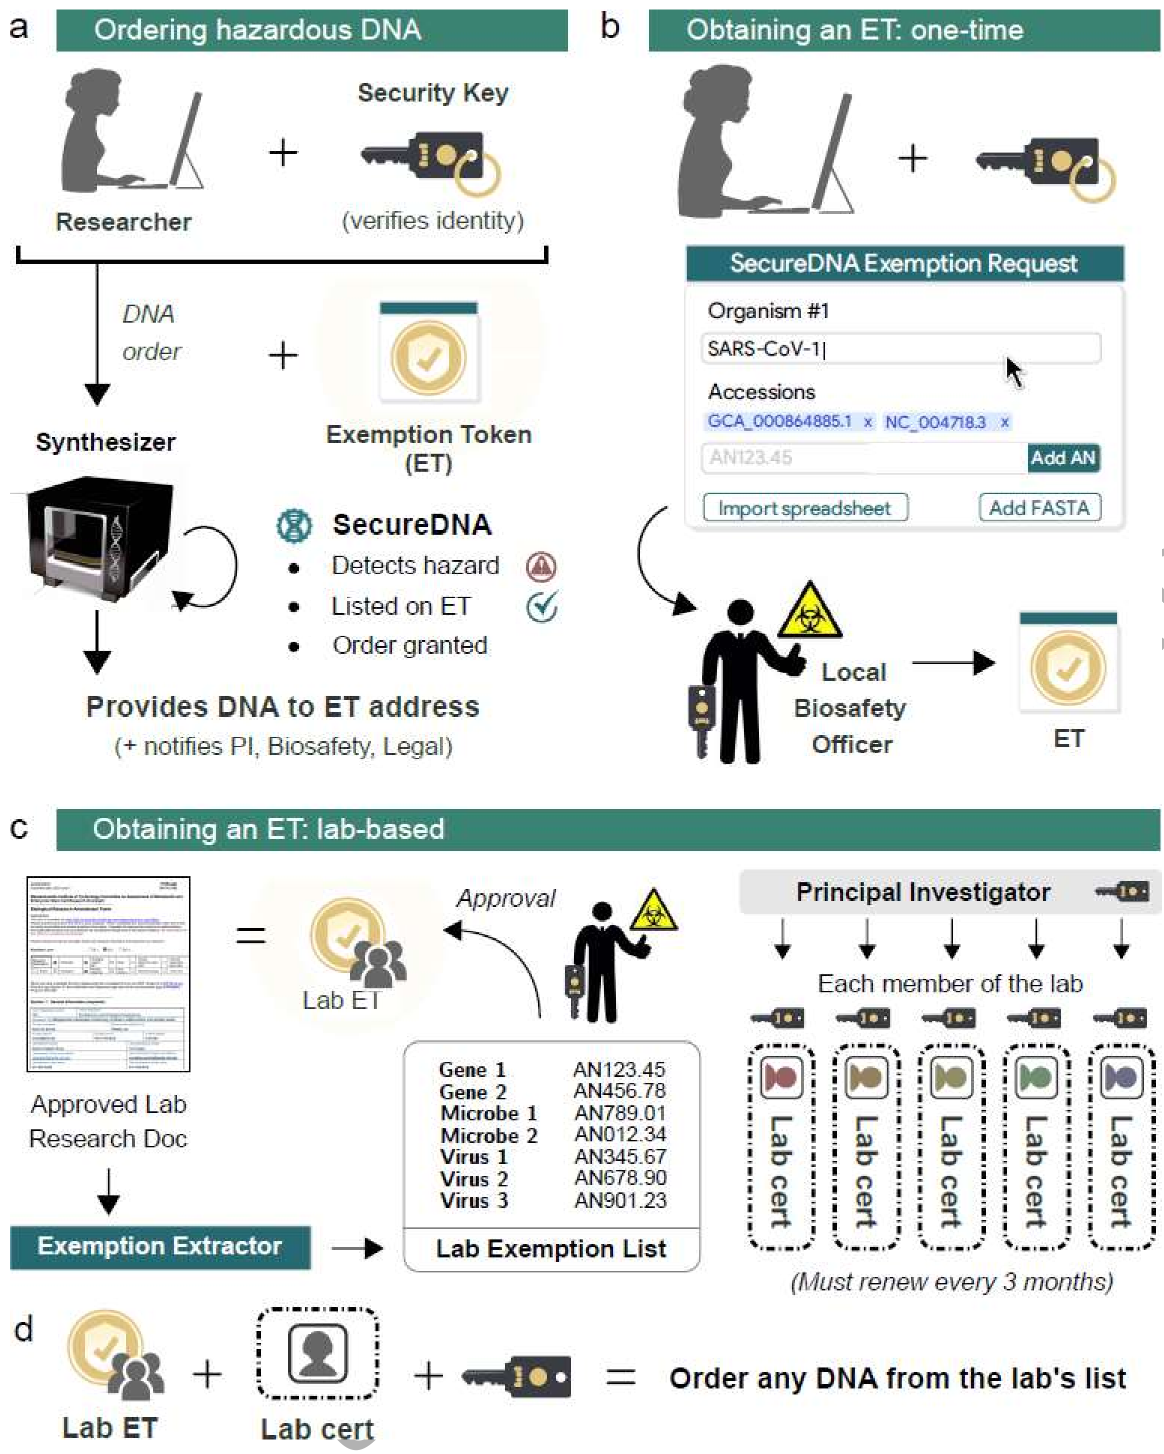

Figure 5. SecureDNA exemption lists for automated customer screening. (a) Researchers with a hardware security key can submit an *exemption token* (ET) with their order to bypass denials for any controlled sequence listed on the ET. (b) To obtain a one-time ET, researchers can specify a controlled agent from a drop-down table, enter a GenBank accession number, or import a DNA sequence, then send it to their local biosafety officer to be approved and signed using the institution's certificate. (c) Labs can obtain an ET that will work for all members by using an extraction tool on their research registration document to obtain a list of genes, microbes, and viruses listed in the document, which can then be approved by the biosafety officer. The lab principal investigator can grant "lab certificates" to each member. (d) A lab ET, lab cert, and matching security key can obtain DNA from any gene or organism the lab is approved to work with, including not-yet-deposited strains and predicted functional variants.

Automating Customer Screening and Permissions

Members of the International Gene Synthesis Consortium spend considerable time and effort screening their customers for legitimacy. In some cases, providers contact the customer's local biosafety authority to ask whether the customer is authorized to work with a particular controlled organism.

An ideal screening system would automatically grant researchers access to DNA corresponding to pathogens that their biosafety authority has already approved them to work with. SecureDNA allows researchers to submit an "exemption token" (ET) with their order to obtain any permitted controlled DNA (Fig. 5a). A one-time token contains a public key identification (PKI) certificate instructing the database to ignore matches corresponding to exempted controlled sequences, but only for requests originating from the user. This allows researchers to swiftly obtain DNA corresponding to any controlled gene, organism, or set of raw sequences by requesting an exemption from their biosafety officer, who approves it using their own security key and a certificate issued by a higher biosafety authority, such as one at the national level (Fig. 5b, Fig. S2).

To facilitate seamless ordering, we wrote software to help biosafety authorities convert approved laboratory research registration documents into exemption tokens for permitted genes and organisms, then issue the lab an ET (Fig. 5c). The principal investigator can separately grant certificates to lab members, enabling them to automatically bypass SecureDNA denials of sequences on the lab exemption list as long as they are shipped to the lab's address (Fig. 5d). In all cases, a physical hardware authentication key ($20-$30) or time-based one-time password (TOTP) from an authentication app is required to validate the user's identity, and placing an order for an exempted controlled sequence notifies the principal investigator, the biosafety officer, and the organization's legal contact.

For example, suppose that a laboratory wants to develop a broad-spectrum influenza vaccine and needs access to sequences from 1918 influenza. Using the exemption request tool, they can select "1918 influenza," then send the request to their biosafety authority (Fig. S3). Once approved, every lab member can use the ET and their personal certificate proving lab affiliation to ship any DNA sequences corresponding to that strain to the lab's registered address, or synthesize them on a benchtop device. Researchers can also request exemptions covering lists of sequences, such as oligonucleotide libraries for deep mutational scanning [23] and directed evolution [24]. Importantly, the database tracks which pieces of controlled sequences have been synthesized by any provider, with or without exemptions, in order to detect split-order attacks.

As a final precaution, whenever an order for a public controlled sequence is automatically approved using an exemption list, the SecureDNA database notifies the lab's principal investigator, biosafety authority, and institutional legal department for record-keeping purposes. Because biosafety vetting and approval are available as a commercial service [25], researchers anywhere in the world can access the benefits of seamless and secure ordering via automated customer screening.

DISCUSSION

Progress in the life sciences is demonstrably vulnerable to public mistrust [26]. A pandemic deliberately caused by a malicious actor would almost certainly trigger a backlash and draconian research restrictions. Safeguarding the promise of biotechnology requires a way to ensure controlled DNA is only shipped to legitimate laboratories [27,28], but the complexity of the problem and the expense of regulation have deterred international action.

The advent of free, automated, benchtop-compatible, and privacy-preserving DNA synthesis screening (Table S2) may allow nations currently hesitant to place their own companies at a competitive disadvantage to begin regulating the sector and clarify liability in the event of misuse. Indeed, signatories of the Biological Weapons Convention may be obligated to require free screening under Article IV, which states that parties must "take any necessary measures to prohibit and prevent the development, production, stockpiling, acquisition, or retention of the agents, toxins, weapons, equipment and means of delivery… within the territory of such State, under its jurisdiction or under its control anywhere."

While DNA synthesis screening may prevent widespread access to credible pandemic agents for many years, advances in *de novo* protein design [29-31] will gradually undermine its effectiveness. Design models can already generate functional equivalents of binding proteins; while this presumably includes toxins, they are far less relevant to international security than pandemic agents. However, *de novo* biodesign tools may eventually be capable of generating allosteric or catalytic proteins sufficient to enhance natural pathogens or even produce novel pandemic agents. Requiring that requested novel biomolecules be demonstrably consistent with a declared research intent could offer limited protection, but such intent-based checks are easily circumvented by benign explanations and do not scale to complex or exploratory research. Moreover, function and folding prediction tools that rely on a complete and intact sequence will not be able to detect such threats if the resulting synthetic genes are ordered in pieces, among other evasive strategies (Fig. S4). Finally, denying all unknown DNA as a blanket strategy would preclude routine use of optimized linkers, spacers, randomized regions, and other legitimate novel sequences, effectively banning a wide range of legitimate research.

Controlling access to protein and genome design tools via APIs and logging controlled designs for screening purposes, which could be done for nearly all users even if source code was shared among all registered developers, may help prevent redesigned versions of controlled agents from evading detection. To further mitigate security risks, leaders in the protein design community have called for the retention of cryptographic records of DNA synthesis orders [32], which could deter malicious actors by reliably identifying the source of the harmful DNA after the fact. Systems running SecureDNA could readily store order hash records, with old keys divided among trusted authorities after rotations and released only following a catastrophic event, to enable forensic screening of past synthesis activity while maintaining day-to-day privacy [32].

Collectively, our results suggest that SecureDNA can provide free, private, reliable, and verifiably up-to-date nucleic acid synthesis screening at a scale sufficient to meet global demand, with negligible false alarms and a largely automated authorization process. If supported by favorable regulatory and liability policies, hardware integration into next-generation synthesizers, and subsidized trade-in programs, near-universal screening could dramatically reduce unauthorized access to pandemic-class agents without delaying research. By preventing the conversion of controlled blueprints into dangerous pathogens, we can safeguard biotechnology and the world from the threat of deliberate pandemics.

CODE AND DATA AVAILABILITY

The open-source SecureDNA software is available at https://github.com/SecureDNA/. The privacy of customer data from historical DNA synthesis orders is protected by legal agreements signed by the relevant providers. We are not aware of any open data sources, and encourage interested parties to reach out to providers directly to discuss the possibility of signing non-disclosure agreements.

SUPPLEMENTARY DATA

Supplementary data are available at NSR online.

ACKNOWLEDGEMENTS

We are deeply grateful to several anonymous industry partners, and to B. Mueller for coordination.

FUNDING

Financial support was obtained from the Open Philanthropy Project (to MIT, Aarhus University, and SecureBio), an anonymous philanthropist from the mainland of China (to Tsinghua University), the Aphorism Foundation (to MIT and SecureBio), and Effective Giving (to MIT and SecureBio). The funders had no role in study design, data collection, data analysis, data interpretation, or writing of the report. To preserve international neutrality, no government funds were used to support this project.

AUTHOR CONTRIBUTIONS

A.C.Y., D.G., K.M.E., L.F., and R.R. conceived the study; Y.Y., C.B., and M.G. wrote the cryptographic analysis with advice from A.C.Y., I.D., R.R., D.W., V.V., and A.S.; D.G. and K.M.E. created the figures, A.C.Y., M.G., Y.Y., C.B., I.D., K.L, J.L., X.L., Y.W., and H.Z. contributed to system design; software systems architecture & security was led by L.F.; the software implementation was overseen by L.F. and J.B. and performed by L.F., J.B., T.V., M.K., W.C., F.S-L., L.V.H., S.W., and B.W-R.; database curation was led by T.V. and L.F. with input from K.M.E. and D.G.; performance assessments were performed by L.F. and D.G.; specificity analyses were conducted by D.G., L.F., T.V., and K.M.E.; and the certificate design and implementation were led by L.F. and F.S-L. All authors contributed to writing the paper.

Conflict of interest statement. K.M.E. and D.G. are authors of PCT/US2021/014814 filed by the Massachusetts Institute of Technology. All authors share an interest in preventing future pandemics.

REFERENCES

1. Neumann G, Ozawa M, Kawaoka Y. Reverse genetics of influenza viruses. *Methods Mol Biol* 2012; **865**: 193–206.

2. Maroun J, Muñoz-Alía M, Ammayappan A *et al*. Designing and building oncolytic viruses. *Future Virol* 2017; **12**: 193–213.

3. Xie X, Lokugamage KG, Zhang X *et al*. Engineering SARS-CoV-2 using a reverse genetic system. *Nat Protoc* 2021; **16**: 1761–84.

4. Cello J, Paul AV, Wimmer E. Chemical synthesis of poliovirus cDNA: generation of infectious virus in the absence of natural template. *Science* 2002; **297**: 1016–8.

5. Esvelt KM. Delay, detect, defend: Preparing for a future in which thousands can release new pandemics. *Geneva Papers* 2022.

6. Mathieu E, Ritchie H, Rodés-Guirao L, *et al*. COVID-19 Pandemic. *Our World in Data* https://ourworldindata.org/coronavirus (20 January 2026, date last accessed).

7. Carroll D, Daszak P, Wolfe ND *et al*. The Global Virome Project. *Science* 2018; **359**: 872–4.

8. Warren CJ, Yu S, Peters DK *et al*. Primate hemorrhagic fever-causing arteriviruses are poised for spillover to humans. *Cell* 2022; **185**: 3980–91.

9. Sun H, Li H, Tong Q *et al*. Airborne transmission of human-isolated avian H3N8 influenza virus between ferrets. *Cell* 2023; **186**: 4074–84.

10. Hou YJ, Chiba S, Leist SR *et al*. Host range, transmissibility and antigenicity of a pangolin coronavirus. *Nat Microbiol* 2023; **8**: 1820–33.

11. Thadani NN, Gurev S, Notin P *et al*. Learning from prepandemic data to forecast viral escape. *Nature* 2023; **622**: 818–25.

12. Grange ZL, Goldstein T, Johnson CK *et al*. Ranking the risk of animal-to-human spillover for newly discovered viruses. *Proc Natl Acad Sci U S A* 2021; **118**: e2002324118.

13. Bügl H, Danner JP, Molinari RJ *et al*. DNA synthesis and biological security. *Nat Biotechnol* 2007; **25**: 627–9.

14. Diggans J, Leproust E. Next Steps for Access to Safe, Secure DNA Synthesis. *Front Bioeng Biotechnol* 2019; **7**: 86.

15. Beal J, Clore A, Manthey J. Studying pathogens degrades BLAST-based pathogen identification. *Sci Rep* 2023; **13**: 5390.

16. Edison R, Toner S, Esvelt KM. Evaluating the adequacy of DNA synthesis screening. https://drive.google.com/file/d/1hNUnU8i2yubt5uesmmV17aTJXhYYDgTY/view?usp=drive_link (20 January

2026, date last accessed).

17. Soice EH, Rocha R, Cordova K *et al*. Can large language models democratize access to dual-use biotechnology? [preprint] arXiv: 2306.03809.

18. Nouri A, Chyba CF. DNA synthesis security. *Methods Mol Biol* 2012; **852**: 285–96.

19. Gretton D, Wang B, Edison R *et al*. Exact-match search with functional variant prediction enables automated DNA screening [preprint]. bioRxiv: 2024.03.20.585782.

20. National Institute of Standards and Technology. NIST.FIPS.140-2 Security Requirements for Cryptographic Modules. https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-2.pdf (20 January 2026, date last accessed).

21. National Institute of Standards and Technology. FIPS1402IG-Implementation Guidance for FIPS 140-2 and the Cryptographic Module Validation Program. https://csrc.nist.gov/csrc/media/projects/cryptographic-module-validation-program/documents/fips140-2/fips1402ig.pdf (20 January 2026, date last accessed).

22. Shamir A. How to share a secret. *Commun ACM* 1979; **22**: 612–3.

23. Fowler DM, Fields S. Deep mutational scanning: a new style of protein science. *Nat Methods* 2014; **11**: 801–7.

24. Packer MS and Liu DR. Methods for the directed evolution of proteins. *Nat Rev Genet* 2015; **16**: 379–94.

25. Biosafety Buyer's Guide - Biosafety Consultants. https://biosafetybuyersguide.org/consultants.html (20 January 2026, date last accessed).

26. Lewis T. The Quest to Overcome Gene Therapy's Failures. *Nature* 2021; doi: 10.1038/d41586-021-02734-w.

27. Carter S and DiEuliis D. Mapping the Synthetic Biology Industry: Implications for Biosecurity. *Health Secur* 2019; **17**: 403–6.

28. DiEuliis D. Revisiting the Digital Biosecurity Landscape. In: Greenbaum, D. (eds.). *Cyberbiosecurity*. Springer, Cham, 2023, 71-8.

29. Watson JL, Juergens D, Bennett NR *et al*. *De novo* design of protein structure and function with RFdiffusion. *Nature* 2023; **620**: 1089–100.

30. Alamdari S, Thakkar N, Berg R *et al*. Protein generation with evolutionary diffusion: sequence is all you need [preprint]. bioRxiv: 2023.09.11.556673.

31. Chen B, Cheng X, Li P *et al*. xTrimoPGLM: Unified 100B-Scale Pre-trained Transformer for Deciphering the Language of Protein [preprint]. arXiv: 2401.06199.

32. Baker D and Church G. Protein design meets biosecurity. *Science* 2024; **383**: 349.

33. Office of the National Cyber Director. Back to the Building Blocks: A Path Toward Secure and Measurable Software (The White House, 2024). https://www.whitehouse.gov/wp-content/uploads/2024/02/Final-ONCD-Technical-Report.pdf (20 January 2026, date last accessed).

34. Titus AJ, Flower A, Hagerty P *et al*. SIG-DB: Leveraging homomorphic encryption to securely interrogate privately held genomic databases. *PLoS Comput Biol* 2018; **14**: e1006454.

35. The SecureDNA cryptography team. Hiding dangerous DNA in plain sight. *submitted*.

36. Baum C, Berlips J, Chen W *et al*. Efficient Maliciously Secure Oblivious Exponentiations [preprint]. Cryptology ePrint Archive: 2024/1613.