

## 1 Title

2 Innocent until proven guilty: Privacy-preserving search over a central CODIS criminal database from the  
3 field

## 4 Authors

5 Jacob A. Blindenbach<sup>1,\*</sup>, Karthik A. Jagadeesh<sup>2,\*</sup>, Gill Bejerano<sup>3,4,5,6,7</sup>, David J. Wu<sup>1,7</sup>

## 6 Affiliations

7 \* These authors contributed equally to this work and will list themselves first on their CVs.

8  
9 <sup>1</sup> Department of Computer Science, University of Virginia, Charlottesville, VA

10 <sup>2</sup> Klarman Cell Observatory, Broad Institute of MIT and Harvard, Cambridge, MA

11 <sup>3</sup> Department of Computer Science, Stanford University, Stanford, CA

12 <sup>4</sup> Department of Developmental Biology, Stanford University, Stanford, CA

13 <sup>5</sup> Department of Pediatrics (Medical Genetics), Stanford University, Stanford, CA

14 <sup>6</sup> Department of Biomedical Data Science, Stanford University, Stanford, CA

15 <sup>7</sup> Corresponding authors: [bejerano@stanford.edu](mailto:bejerano@stanford.edu) (G.B) and [dwu4@virginia.edu](mailto:dwu4@virginia.edu) (D.J.W)

## 16 Abstract

17 The presumption of innocence (i.e., the principle that one is considered innocent until proven guilty) is a  
18 cornerstone of the criminal justice system in many countries, including the United States. DNA analysis is  
19 an important tool for criminal investigations<sup>1</sup>. In the U.S. alone, it has already aided in over half a million  
20 investigations using the Combined DNA Index System (CODIS) and associated DNA databases<sup>2</sup>. CODIS  
21 includes DNA profiles of crime scene forensic samples, convicted offenders, missing persons and more.  
22 The CODIS framework is currently used by over 50 other countries<sup>3</sup> including much of Europe, Canada,  
23 China and more. During investigations, DNA samples can be collected from multiple individuals who may  
24 have had access to, or were found near a crime scene, in the hope of finding a single criminal match<sup>4</sup>.  
25 Controversially, CODIS samples are sometimes retained from adults and juveniles despite *not* yielding  
26 any database match<sup>4-6</sup>. Here we introduce a cryptographic algorithm that finds any and all matches of a  
27 person's DNA profile against a CODIS database *without* revealing anything about the person's profile to  
28 the database provider. With our protocol, matches are immediately identified as before; however,  
29 individuals who do *not* match anything in the database retain their full privacy. Our novel algorithm runs  
30 in 40 seconds on a CODIS database of 1,000,000 entries, enabling its use to privately screen potentially-  
31 innocent suspects even in the field.

32

## 33 Introduction

34 DNA-based forensic analysis is a powerful tool used by law enforcement agencies around the world for  
35 solving crimes<sup>1-3</sup>. With today's technology, local police stations<sup>7</sup> and even agents in the field<sup>8</sup> can  
36 generate a suspect's DNA profile to search against central criminal DNA databases in an impressive 90  
37 minutes<sup>9,10</sup>. At the same time, the increased prominence of DNA-based forensics opens up new avenues  
38 for misuse including social control and racial profiling<sup>11,12</sup>.

39 In particular, the storage of DNA samples from potentially innocent individuals *permanently* links their  
40 genetic identities to criminal databases (without due process). Within the United States, these  
41 controversial practices have included collecting DNA from individuals who have been arrested but not  
42 convicted or even charged with a crime<sup>4</sup>, people who are not even arrested (so-called "stop-and-spit"  
43 and "swab-and-go" practices<sup>13</sup>), detained immigrants and asylum seekers<sup>6</sup>, and even children brought in  
44 for questioning<sup>5</sup>.

45 Here we develop a novel solution (Figure 1) whereby an agent in the field, using modest computational  
46 resources, could privately query a suspect's DNA profile against a large central database of DNA profiles.  
47 In less than 40 seconds, the agent learns whether the suspect's profile matches, based on CODIS rules,  
48 against any profile in the central database of 1,000,000 profiles, while learning nothing else about the  
49 contents of the central database. More importantly, the central database itself learns absolutely *nothing*  
50 about the DNA profile being searched. Any match discovered can be investigated further as before.  
51 However, should no match be made, the suspect's DNA profile, now exonerated, can be disposed of on  
52 the spot, with zero risk that the central database provider chose to retain it.

## 53 Results

### 54 CODIS system of DNA profiles and profile matches

55 The United States CODIS system originally established a set of 13 loci across the genome coinciding with  
56 short tandem repeats (STRs) as a method for comparing genetic data for identification purposes<sup>3</sup>. These  
57 were expanded to a set of 20 loci in 2017. Many countries have adopted a similar system with a  
58 combination of existing core STR loci and region-specific STR loci. For example, a European Union (EU)  
59 system of 16 loci, a UK system of 11 loci and a Chinese system of 20 loci have all been described<sup>14,15</sup> (see  
60 Online Methods and Supplementary Table 1). At each locus, individuals have 2 alleles, one inherited  
61 (with possible personal modification) from each parent. The allele at each locus is represented by a  
62 varying range and represents the number of repeats of a 2 to 6-character (base pair) generic sequence  
63 observed in the individual's genome. Based on frequency statistics collected by the FBI, the probability  
64 that two unrelated individuals share the same STR profile across all 13 core loci positions of the older US  
65 system is 1 in 575 trillion<sup>16</sup>. This probability further decreases with the expanded set of 20 loci  
66 introduced in 2017.

67 The CODIS system describes several ways to query a database of STR profiles. The standard and default  
68 method is a "high-stringency search with one mismatch," which requires that both of the alleles  
69 appearing in at least 19 out of the 20 STR loci between the query profile and the database profile match  
70 exactly<sup>17,18</sup>. Central CODIS databases holding thousands to millions of DNA profiles that can be used for  
71 such queries are maintained in the US at the national, state, and sometimes even municipal level<sup>2</sup> (as  
72 well as very similar databases and matching rules in dozens of other countries<sup>3</sup>).

### 73 [Encoding an STR profile for secure computation](#)

74 We encode an STR profile as a binary string. For each STR locus, we propose a public dictionary that  
75 maps every pair of STR alleles to a unique binary string. The number of bits used to encode each allele is  
76 determined based on the number of unique alleles expected at the locus<sup>14</sup>. We assume that all entries in  
77 the central database (unknown to the agent making the query) and the suspect's STR profile held by the  
78 agent (unknown to the central database) are encoded in this commonly-agreed upon manner.

### 79 [Secure protocol overview](#)

80 It suffices to construct a secure protocol for comparing a single entry in the central database with the  
81 suspect's profile. By running this protocol over and over against all entries in the central database, the  
82 agent will learn the indices of the complete set of matching records, and nothing else, while the central  
83 database will learn nothing (Figure 1C).

84 We start by fitting a computational model to the task, before later securing it: given two binary strings  
85 encoding two STR profiles as above, decide whether they correspond to a match according to the CODIS  
86 specification or not. Our key observation is that one can efficiently compute this using a compact  
87 deterministic finite automaton<sup>19</sup> (DFA). In a DFA, the computation begins at an initial state, and at each  
88 step of the computation, the DFA reads a bit of the input and advances the state. After reading all of the  
89 input bits, the DFA ends in either an "accepting" state or a "rejecting" state. For example, a DFA can be  
90 constructed to test for equality between two equal length bit-strings by defining two sets of states: a set  
91 of "matching" states and another for "mismatching" states. The computation begins in the "matching"  
92 set, and as each bit of the input is read, it is compared against the target bit. As long as the current state  
93 is in the matching set, if the two bits match, the program transitions to the next state in the matching  
94 set, and otherwise, it transitions to a state in the mismatching set. Once a single bit mismatches, one  
95 enters the mismatching set, from which all inputs lead only to the next state in the mismatching set. The  
96 input bit string is equal to the target string if the computation concludes in a state in the matching set,  
97 and is otherwise unequal (see Supplementary Figure 1). We use this DFA to check for matches at a single  
98 STR locus. A similar DFA can be used to test for equality with up to one mismatch (see Online Methods  
99 and Supplementary Figure 2). We use this one to compute a CODIS match of at least 19 of 20 STR loci.

100 In our setting, the central database owner constructs a DFA for each profile in its database. The input to  
101 the DFA is the agent-held suspect's STR profile. The DFA computation ends in an accepting state if the  
102 suspect's profile is a CODIS match to the central database profile; otherwise, the DFA computation ends  
103 in a rejecting state.

104 Our protocol now proceeds as follows. At the beginning, the agent knows the initial state of the DFA as  
105 well as the suspect's STR profile (hidden from the central database). The central database holds the DFA  
106 corresponding to a database entry (hidden from the agent). Using a cryptographic protocol called  
107 "oblivious transfer<sup>20,21</sup>," the agent and the central database now jointly perform the evaluation of the  
108 DFA on the agent's input STR profile.

109 Specifically, at each step of the DFA evaluation, the central database enumerates *all* possible states of  
110 the DFA that the agent might be in and all possible states the agent will end up in based on the next bit  
111 of the agent's input. These statements are of the form "if you are in state *X* and the next bit of your  
112 input is 0, then you will proceed to state *Y*." Using the oblivious transfer protocol, the agent can secretly  
113 choose to learn exactly one of these statements *without* revealing which one she chose to the database

114 server. In this case, the agent chooses the statement corresponding to her current state and the next bit  
115 of her input; this in turns reveals to the agent her next state in the DFA evaluation. The oblivious  
116 transfer protocol hides all of the other statements from the agent, so the agent cannot learn what  
117 would have happened had she been in a different state or had a different input.

118 At the very end of the protocol, the agent arrives at either an accepting state or a rejecting state, which  
119 indicates whether the STR profile she provided matched against the central database's profile or not. As  
120 described so far, the agent learns the full execution path in the DFA as well as whether her input profile  
121 matches the database profile or not. To ensure that the execution path does not leak additional  
122 information about the profiles on the database server, the central database additionally encrypts all of  
123 the intermediate steps of the computation (in a way that still enables the above evaluation procedure).  
124 Then, at each of the intermediate steps, the agent no longer knows where she is in the actual DFA  
125 execution. The central database only provides a single intelligible state: the very last state which reveals  
126 whether the two profiles match or not match.

127 By repeating the above protocol with each profile in the central database, while changing the encryption  
128 of intermediate states every run, the agent has learned just one thing: the database indices of any and  
129 all profile matches. She learns nothing else about any of the centrally-stored profiles. On the flip side,  
130 the central database learns nothing at all about the suspect's STR profile; at every step, they only  
131 provide an exhaustive list of "if you are here, and have this bit next, then go there." We have thus  
132 achieved the desired goal of Figure 1C. We encourage our readers to refer to the Online Methods for  
133 the full technical details and security analysis.

### 134 [Performance measurements](#)

135 With our protocol implementation, an agent in the field in Northern California can privately query a  
136 CODIS database containing 1 million STR records located 3,000 miles away in Northern Virginia in just 38  
137 seconds, using 180 MB of online communication. In our experiments, we represent each STR profile as a  
138 vector of 20 biallelic components (212 bits in total) based on the current US CODIS specification<sup>2</sup>. Our  
139 protocol additionally requires preparing 116 million oblivious transfer correlations<sup>20,21</sup> and 122 MB of  
140 client-side storage (see Online Methods and Supplementary Table 2). These can be generated in a  
141 separate preprocessing phase on commodity hardware in about a minute using existing state-of-the-art  
142 oblivious transfer extension protocols<sup>22</sup>. Since these correlations are independent of both the query and  
143 the database contents, they can be prepared concurrently with the 90 minutes needed for STR profile  
144 derivation<sup>9,10</sup>, and thus, contributes no extra latency.

145 We also measured the performance of our protocol on CODIS specifications reported for the UK<sup>14</sup>, the  
146 EU<sup>14,18</sup>, and China<sup>15</sup> (see Online Methods for specifications details). In all cases, the cost of the protocol is  
147 smaller or comparable to that of the US CODIS system (up to 40 seconds and under 200 MB of  
148 communication; see Table 1). To illustrate the scalability of our approach, we also measured the  
149 performance for a CODIS system with 40 loci and an encoding length of 492 bits, a setting based on a  
150 system previously tested by NIST<sup>23</sup>. Performing a CODIS search in this setting over a database of a million  
151 records completes in just 72 seconds and requires only 340 MB of communication. Thus, our protocol  
152 also scales favorably to future scenarios with an expanded set of STR loci.

153

154 We also tested the performance of our protocol on central databases of different sizes (Figure 2). Both  
155 the online bandwidth and execution time of our protocol scale linearly with the size of the STR profile (a  
156 function of both the number of loci and the number of bits needed to represent the allele values at each  
157 locus) and with the size of the central database. For example, our protocol can search over a 20-STR U.S.  
158 database with 10 million records in just under 6 minutes.

## 159 Discussion

160 DNA profiles are an extremely powerful tool in forensics and crime solving<sup>24</sup>. Our law enforcement  
161 agencies have a duty to both serve and protect their communities. How can they use DNA databases to  
162 find criminals while simultaneously protecting the privacy of innocent individuals who make up the  
163 majority of each society<sup>12</sup>? It is natural to think that searching a suspect's DNA profile against a master  
164 database either requires the database to see the suspect's profile, or for the agent holding the suspect  
165 profile to have a local copy of the master database. Both scenarios compromise privacy (Figure 1). Here  
166 we show a third way, whereby a field agent searches the remote master database without learning  
167 anything about it except any possible match. Moreover, the central database aids the search while  
168 learning *nothing* about the suspect's profile. Should the search come up empty, if the field agent  
169 disposes of the sample, they have both served their community and protected its privacy.

170 DNA profiling machines can now routinely produce a searchable profile in a matter of 90 minutes, and  
171 can even be carried to the field<sup>8</sup>. Integrating our privacy-preserving protocol with such a system adds  
172 minimal overhead and can easily fit into existing workflows. In fact, the computational requirements to  
173 run the protocol are so modest, it is likely they can be performed on a modern smartphone. The  
174 performance of our protocol also compares favorably against generic approaches for privacy-preserving  
175 computation. For instance, a direct implementation of a query protocol using Yao's garbled circuits<sup>25</sup>  
176 would require communicating, storing, and evaluating a circuit that is around 8 GB ( $\approx$  260 million AND  
177 gates) in size for a database with 1 million profiles (and increase to 80 GB of communication and storage  
178 for a 10 million entry database). In comparison, our protocol requires less than 125 MB of offline storage  
179 for the OT correlations (which can be generated in about a minute<sup>22</sup>) for a database of 1 million profiles.

180 Our implementation follows the CODIS guidelines for high-stringency matching (the default mode of  
181 searching)<sup>17,18</sup>. The deterministic finite automata at the root of our approach can be easily extended to  
182 also support moderate and low-stringency matches as well as partial match queries, with modest  
183 increases in computation time and communication. In fact, similar operations like paternity testing and  
184 ethnicity identification can also be formulated as a similar string-matching problem and implemented  
185 using a similar approach. In all cases, the correct answer is obtained while the input DNA profiles to the  
186 computation remain private.

187 Our protocol is relevant not only in the US, but also in any of dozens of countries that use a CODIS-like  
188 system<sup>3</sup>. It scales well with the size of the central database (Figure 2), on current hardware that will only  
189 get faster, and most importantly, it gives the agent in the field or local office, the ability to destroy an  
190 exonerated profile that has yielded no incriminating match. Whether they are instructed to do so or not  
191 is a civil rights matter that each country must resolve for itself<sup>11,12,26,27</sup>. The importance of our work is in  
192 showing that accurate practical implementations to enable these fundamental rights are already doable.

## 193 Online Methods

### 194 STR profiles and representation

195 The US National Institute of Standards and Technology (NIST) has enumerated all possible allele values  
196 for each STR core locus in the current US system<sup>14</sup>. While over 50 other countries reportedly use a local  
197 variant of the same CODIS system<sup>3</sup>, details of each country’s system, or even a list of these 50+ countries  
198 are not easily found. Possibly current versions of the EU and UK systems are described by NIST<sup>14</sup>, while a  
199 partial description of the Chinese system is provided in a paper co-authored by officials from the  
200 Chinese Ministry of Justice<sup>15</sup>. Similarly, the high-stringency matching rule with at most one mismatching  
201 locus is published for the US<sup>17</sup> and EU<sup>18</sup>. As the exact CODIS system parameters are immaterial for the  
202 essence of our proposed solution, when necessary, we infer the number of possible values at a locus  
203 based on the existing NIST standards, and also assume a default search configuration of high-stringency  
204 matching of all but one locus for all systems.

205 In 2017, the US 13 loci system was replaced by the 20 loci system<sup>3</sup>. It is unclear if or when this set will be  
206 expanded on. For purposes of illustrating the scalability of our solution we use a 40 loci system  
207 described in a NIST paper partly funded by the FBI<sup>23</sup>.

### 208 Security model

209 We work in a two-party setting where the server holds a profile  $\mathbf{v} = (v_1, \dots, v_n)$  and the client holds a  
210 query vector  $\mathbf{w} = (w_1, \dots, w_n)$ , where each of the components  $v_i, w_i \in \{0,1\}^\ell$  can be represented by  $\ell$ -  
211 bit strings. We define the “threshold matching” function  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  to be the following Boolean-valued  
212 function:

$$213 \quad \text{TM}_k(\mathbf{v}, \mathbf{w}) = \begin{cases} 1, & |\{v_i \neq w_i \mid 1 \leq i \leq n\}| \leq k \\ 0, & |\{v_i \neq w_i \mid 1 \leq i \leq n\}| > k \end{cases}$$

214 In words,  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  outputs 1 if the vectors  $\mathbf{v}$  and  $\mathbf{w}$  disagree on at most  $k$  components. In this work,  
215 we design a secure protocol such that at the end of the protocol, the client learns the value of  
216  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  and nothing more about the server’s profile  $\mathbf{v}$ , while the server learns nothing about the  
217 client’s query  $\mathbf{w}$ . Importantly, the client only learns whether the number of differing components  
218 between  $\mathbf{v}$  and  $\mathbf{w}$  is greater than  $k$  or not, but nothing about the exact number of differing components  
219 or the indices of the differing components.

220 We can naturally extend this notion to the setting where the server holds a database with many vectors  
221  $\{\mathbf{v}_1, \dots, \mathbf{v}_t\}$  and the client’s goal is to learn all of the indices  $i \in \{1, \dots, t\}$  where  $\text{TM}_k(\mathbf{v}_i, \mathbf{w}) = 1$ , but  
222 nothing more about the vectors  $\mathbf{v}_1, \dots, \mathbf{v}_t$ . Since a protocol for private evaluation of  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  for a  
223 single entry  $\mathbf{v}$  suffices for searching over a database of values (by repeating the protocol for each  
224 database entry  $\mathbf{v}_1, \dots, \mathbf{v}_t$ ), we focus on the single-instance setting for the remainder of this section.

225 Throughout this work, we assume that the client and the server are semi-honest or “honest-but-  
226 curious;” namely, both the client and the server will follow the protocol as described, but may try to  
227 infer additional information about the other party’s private inputs based on the messages they receive  
228 in the protocol. We note that our protocol actually ensures privacy of the client’s query even against a  
229 *malicious* database server that may arbitrarily deviate from the protocol execution, provided that the  
230 underlying “oblivious transfer” protocol we use (see below) is secure against a malicious sender.

## 231 Oblivious transfer and the OT-hybrid model

232 An oblivious transfer (OT) protocol<sup>20,21,28</sup> is a two-party protocol between a sender and a receiver. In a 1-  
233 out-of- $k$  OT protocol on  $t$ -bit messages, the sender holds  $k$  messages  $m_1, \dots, m_k \in \{0,1\}^t$  while the  
234 receiver holds an index  $i \in \{1, \dots, k\}$ . At the end of the protocol, the receiver learns the  $i^{\text{th}}$  message  $m_i$   
235 and nothing else about any of the remaining messages, while the sender learns nothing. To facilitate the  
236 analysis of our protocol, we will work in the “OT-hybrid” model where we assume that the parties have  
237 access to a *trusted party* that implements the above 1-out-of- $k$  OT functionality<sup>20</sup>. We can then replace  
238 the trusted party with a cryptographic implementation of a 1-out-of- $k$  OT protocol<sup>28</sup>. If our protocol  
239 provides semi-honest security in the OT-hybrid model, and we instantiate the OT protocol with a  
240 cryptographic protocol that is secure against semi-honest adversaries, then the overall protocol is also  
241 secure against semi-honest adversaries (*without* relying on any trusted party)<sup>29</sup>.

## 242 Oblivious transfer correlations

243 We can significantly reduce the *online* cost of oblivious transfer by first precomputing *input-independent*  
244 “oblivious transfer correlations” in an offline (or preprocessing phase). Because the correlations are  
245 input-independent, they can be precomputed without knowledge of the client’s query or the server’s  
246 database. These OT correlations can be generated efficiently using a technique called OT extension in a  
247 separate input-independent preprocessing step<sup>30</sup> (this can even be done with low communication using  
248 a recent approach called silent OT extension<sup>31</sup>). Alternatively, they can be generated ahead of time by a  
249 trusted dealer or a secure hardware platform (observe that in both of these settings, the party  
250 generating the correlations does *not* need to know anything about the query or the database entries).

251 Very briefly, a 1-out-of- $k$  OT correlation for  $t$ -bit messages consists of the following: (1) a tuple of  $k$   
252 random values  $r_1, \dots, r_k \in \{0,1\}^t$  for the server; and (2) a random index  $\beta \in \{1, \dots, k\}$  together with the  
253 value  $r_\beta$  for the receiver. Once the client and server have this OT correlation, it is straightforward to  
254 implement a 2-message 1-out-of- $k$  OT (on an arbitrary collection of sender messages  $m_1, \dots, m_k \in$   
255  $\{0,1\}^t$  and receiver index  $i \in \{1, \dots, k\}$ ). We recall the construction below:

- 256 • **Receiver message:** The receiver sends the index  $j = i + \beta \pmod k$  to the sender. Observe that  
257 since  $\beta$  is uniformly random over  $\{1, \dots, k\}$  and unknown to the sender, this message perfectly  
258 hides the receiver’s index  $i$ . For notational convenience, we will consider the output of  
259 arithmetic modulo  $k$  to be a value between 1 and  $k$  (as opposed to 0 and  $k - 1$ ).
- 260 • **Sender response:** On input an index  $j \in \{1, \dots, k\}$  from the receiver, the sender computes the  
261 blinded message  $c_i \leftarrow m_i \oplus r_{j-i \pmod k}$ , where  $\oplus$  denotes the bitwise exclusive-or operator  
262 (i.e., bitwise xor). The sender sends the blinded messages  $c_1, \dots, c_k$  to the receiver. Since  
263  $r_1, \dots, r_k$  are uniform over  $\{0,1\}^t$  and the receiver knows exactly one of these values (i.e.,  $r_\beta$ ),  
264  $k - 1$  out of the  $k$  values are perfectly hidden from the receiver.
- 265 • **Receiver reconstruction:** The receiver computes its message as  $c_i \oplus r_\beta$ .

266 Correctness of the protocol follows from the following simple relation:

$$267 \quad c_i \oplus r_\beta = m_i \oplus r_{j-i \pmod k} \oplus r_\beta = m_i \oplus r_{i+\beta-i \pmod k} \oplus r_\beta = m_i \oplus r_\beta \oplus r_\beta = m_i.$$

268 To summarize, a 1-out-of- $k$  OT correlation on  $t$ -bit values yields a 1-out-of- $k$  OT on  $t$ -bit messages in  
269 two rounds of interaction where the total communication is  $\lceil \log k \rceil + kt$  bits. The only necessary  
270 computation is integer arithmetic and bitwise operations, so this is a very lightweight protocol.

271 **Representing threshold matching as a deterministic finite automaton**

272 Given two vectors  $\mathbf{v} = (v_1, \dots, v_n)$  and  $\mathbf{w} = (w_1, \dots, w_n)$  where  $v_i, w_i \in \{0,1\}^\ell$ , our objective is to  
 273 securely compute the threshold matching function  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  defined above. In our setting, we assume  
 274 the database server holds  $\mathbf{v}$  while the client holds  $\mathbf{w}$ . The starting point of our design is to express the  
 275 computation of  $\text{TM}_k$  as a composition of two deterministic finite automaton (DFAs):

- 276 • The first DFA checks equality of  $\ell$ -bit strings. Namely, for a string  $v \in \{0,1\}^\ell$ , the machine  
 277 computes the function  $g_v(w)$  that outputs 1 if  $v = w$  and 0 otherwise.
- 278 • The second DFA computes a threshold function. Namely, for a target sequence of bits  
 279  $a_1, \dots, a_n \in \{0,1\}$  and a threshold  $k$ , the machine computes the function  $h_{(a_1, \dots, a_n), k}(b_1, \dots, b_n)$   
 280 that outputs 1 if  $a_i \neq b_i$  on at most  $k$  indices  $1 \leq i \leq n$  and 0 otherwise.

281 By definition, the threshold matching function  $\text{TM}_k$  can now be expressed as

$$282 \quad \text{TM}_k(\mathbf{v}, \mathbf{w}) = h_{(1,1, \dots, 1), k}(g_{v_1}(w_1), \dots, g_{v_n}(w_n)).$$

283 We note that while we can construct a *single* DFA that combines both functionalities, decomposing the  
 284 computation into two separate steps enables a more efficient privacy-preserving protocol.

285 We now show how to construct simple DFAs for computing the functions  $g_v$  and  $h_{(a_1, \dots, a_n), k}$ . For both  
 286 functions, we design a “layered” DFA, which can be more efficiently computed in a privacy-preserving  
 287 manner. First, a DFA consist of a tuple  $M = (Q, \Sigma, \delta, q_0, S)$ , where  $Q$  denote the set of states,  $\Sigma$  is the  
 288 alphabet,  $\delta: Q \times \Sigma \rightarrow Q$  is the state-transition function,  $q_0 \in Q$  is the start state, and  $S \subseteq Q$  is the set of  
 289 accepting states. On input  $x = x_1x_2 \dots x_n \in \Sigma^n$ , the output  $M(x)$  is 1 if  $\delta(q_{n-1}, x_n) \in S$  where  $q_i =$   
 290  $\delta(q_{i-1}, x_i)$  for all  $1 \leq i < n$ , and 0 otherwise. Finally, we say that  $M = (Q, \Sigma, \delta, q_0, S)$  is a layered DFA if  
 291 the following properties hold:

- 292 • The set of states  $Q$  can be partitioned into  $\cup_{i=0}^n Q_i$  where  $Q_0, \dots, Q_n$  are pairwise disjoint,  $Q_0 =$   
 293  $\{q_0\}$ , and  $S \subseteq Q_n$ .
- 294 • The state-transition function  $\delta$  can be decomposed into a collection of functions  $(\delta_1, \dots, \delta_n)$   
 295 where  $\delta_i: Q_{i-1} \times \Sigma \rightarrow Q_i$  and  $\delta(q, \sigma) = \delta_i(q, \sigma)$  for all  $1 \leq i \leq n, q \in Q_{i-1}, \sigma \in \Sigma$ .

296 In words, a layered DFA is one whose states can be partitioned into a collection of  $n + 1$  pairwise  
 297 disjoint sets (i.e., “layers”)  $Q_0, \dots, Q_n$ . On input  $x \in \Sigma^n$ , the state of the DFA after reading the first  $i$  bits  
 298 of  $x$  is contained in layer  $i$  (i.e., in the set  $Q_i$ ). We now describe how to represent  $g_v$  and  $h_{(a_1, \dots, a_n), k}$  as  
 299 layered DFAs:

- 300 • Let  $g_v(w)$  denote the function that outputs 1 if  $v = w$  and 0 otherwise, where  $v, w \in \{0,1\}^\ell$ . It  
 301 is easy to construct a layered DFA for the equality function. The DFA consists of two branches,  
 302 each with  $\ell$  states: an “accept” branch that corresponds to a matching input and a “reject”  
 303 branch that corresponds to a non-matching input. Evaluation begins on the accept branch and  
 304 successively compares the bits of  $w$  to the bits of  $v$  (encoded in the DFA transitions). If a  
 305 mismatch is encountered, the DFA transitions to the reject branch. The output is 1 if the final  
 306 state is on the accept branch and 0 otherwise. We illustrate this in Figure 1. Thus, for  $v \in \{0,1\}^\ell$ ,  
 307 the function  $g_v$  can be computed by a layered DFA with  $\ell + 1$  layers, each containing up to 2  
 308 states. Note that the vector  $v$  is entirely encoded in the transition function  $\delta$  of the DFA. Our



309 protocol for privacy-preserving layered DFA evaluation assumes that the topology of  $M$  is public,  
310 but will hide the transition function  $\delta$  (and thus, the value of  $v$ ) from the client.

311 • Let  $h_{(a_1, \dots, a_n), k}(b_1, \dots, b_n)$  be the threshold function that outputs 1 if there are at most  $k$  indices  
312  $1 \leq i \leq n$  where  $a_i \neq b_i$  and 0 otherwise. It is also straightforward to construct a layered DFA  
313 for  $h_{(a_1, \dots, a_n), k}$ . Namely, we consider a DFA with  $k + 2$  branches. We label each branch with an  
314 integer  $j = 0, 1, \dots, k, k + 1$ , corresponding to the number of mismatches encountered thus far.  
315 Evaluation begins on branch 0 and whenever the DFA reads in an input bit  $b_i \neq a_i$ , then the  
316 state transitions from branch  $j$  to branch  $j + 1$ ; otherwise, evaluation continues on branch  $j$ . All  
317 state transitions on the final branch  $j = k + 1$  (corresponding to an input that differs on more  
318 than  $k$  indices), remain on the branch irrespective of the input bit. We illustrate this in  
319 Supplementary Figure 2 (for the case where  $k = 1$ ). Thus, the function  $h_{(a_1, \dots, a_n), k}$  can be  
320 implemented by a layered DFA with  $n + 1$  layers, each containing up to  $k + 2$  states. Note that  
321 the sequence of bits  $(a_1, \dots, a_n)$  is entirely encoded in the transition function  $\delta$  of the DFA. The  
322 topology of the DFA is a function of the dimension  $n$  and the threshold  $k$ , both of which will be  
323 assumed to be publicly known in our protocol; our final protocol will require that the target bits  
324  $(a_1, \dots, a_n)$  be hidden, which holds as long as the protocol for layered DFA evaluation hides  $\delta$ .

325 It is easy to see that the two layered DFAs described above can be combined into a single layered  
326 DFA for computing  $\text{TM}_k$ . However, the number of layers in the resulting DFA will be the *product* of  
327 the number of layers in the two underlying DFAs. When developing our cryptographic protocol for  
328 privacy-preserving layered DFA evaluation, both the round complexity and the communication  
329 complexity scales with the number of layers. Thus, privately-evaluating two smaller DFAs yields a  
330 significantly more efficient protocol compared with evaluating a single larger layered DFA.

### 331 Privacy-preserving evaluation of a layered DFA

332 We now describe how to use oblivious transfer to construct a privacy-preserving protocol for computing  
333  $\text{TM}_k$ . We leverage our implementation of the threshold matching function as a *layered* DFA to enable a  
334 more efficient protocol. While there already exist protocols for private evaluation of general (i.e., not  
335 necessarily layered) DFAs<sup>32,33</sup>, in this work, we show that the layered structure of our DFAs are  
336 amenable to a simpler and direct OT-based evaluation procedure (without needing to additionally rely  
337 on heavier cryptographic primitives such as homomorphic encryption).

338 Let  $M = (Q, \Sigma, \delta, q_0, S)$  be a layered DFA: namely,  $Q = \bigcup_{i=0}^n Q_i$  and  $\delta = (\delta_1, \dots, \delta_n)$ . We will label the  
339 states  $Q_i$  in layer  $i$  by indices  $1, 2, \dots, |Q_i|$ . In our model, the server holds the layered DFA  $M$  while the  
340 client holds the input  $x \in \Sigma^n$ . At the end of the protocol, the server should learn nothing while the client  
341 should learn the value  $M(x)$ . In addition, we assume that the *topology* of the DFA  $M$  is publicly-known  
342 (i.e., the client knows the number of layers in  $M$  as well as the number of states in each layer of  $M$ ).  
343 What is hidden is the transition function  $\delta$ . This assumption is true for the setting we consider in this  
344 work: here, the CODIS algorithm itself is public (and this determines the topology of the DFA), but the  
345 records within the server's database determine the exact transition function, which is precisely what our  
346 protocol hides.

347 Our protocol relies on the following simple observation: the transition function  $\delta_i$  can be described by a  
348 truth table of size  $|Q_i| \cdot |\Sigma|$ . This yields the following general approach for evaluating  $M$  privately:

349 • The client initializes her state to  $s_0 \leftarrow q_0 \in Q_0$ .

350 • Let  $x \in \Sigma^n$  be the client's input. We maintain the following invariant. At the beginning of round  
 351  $i = 1, \dots, n - 1$ , the client knows the state  $s_{i-1} \in Q_{i-1}$  of the DFA after reading the prefix  
 352  $x_1 \dots x_{i-1}$ . The client wants to learn  $s_i := \delta_i(r_{i-1}, x_i)$ . To do so, the client and server run a 1-out-  
 353 of- $(|Q_i| |\Sigma|)$  OT where the receiver's input is  $(s_{i-1}, x_i)$  and the server associates the message  
 354  $y_i = \delta_i(q, \sigma)$  with index  $(q, \sigma) \in Q_{i-1} \times \Sigma$  is. At the end of this protocol, the client learns the  
 355 state  $s_i = \delta_i(s_{i-1}, x_i) \in Q_i$  of the DFA after reading  $x_1 \dots x_i$ .  
 356 • After  $n - 1$  rounds, the client learns the state  $s_{n-1} \in Q_{n-1}$ . For the final round, the client and  
 357 server run a 1-out-of- $(|Q_{n-1}| |\Sigma|)$  OT protocol where the client's input is  $(s_{n-1}, x_n)$  and the  
 358 server associates the value 1 with the index  $(q, \sigma) \in Q_{n-1} \times \Sigma$  if  $\delta_n(q, \sigma) \in S$  (i.e.,  $\delta_{n-1}(q, \sigma)$  is  
 359 an accepting state) and value 0 with the remaining indices  $(q, \sigma)$  where  $\delta_n(q, \sigma) \notin S$ .  
 360 Correctness of the above protocol follows by correctness of the underlying OT protocol. In particular, OT  
 361 is used to iteratively evaluate the layered transitioned functions  $\delta_1, \dots, \delta_n$ . Moreover, privacy for the  
 362 client's input is ensured by security of the OT protocol since the server's view in the entire protocol  
 363 execution only consists of its view in the OT queries (which hide the client's input).

364 However, this protocol does *not* provide privacy for the server. Namely, the client learns the sequence  
 365 of states corresponding in the DFA evaluation on her input, which could reveal information about the  
 366 state-transition functions. Consider for instance a layered DFA with  $k$  branches where the node in  
 367 branch  $j$  of layer  $i$  is always labeled with the index  $j$ . Then, the sequence of states the client obtains by  
 368 executing the above protocol completely reveals which branch of the computation her input takes at  
 369 each step in the DFA evaluation. This in turn leaks information about the transition function  $\delta_i$  (e.g., the  
 370 client learns whether  $\delta_i(s_{i-1}, \sigma)$  outputs a state in the *same* branch or in a *different* branch). In the case  
 371 of our threshold matching DFA, this in turn reveals to the client partial matches between her query and  
 372 the server's database entry.

373 The above attack shows that additional randomization is *necessary* to ensure security for the server. The  
 374 fix is simple: the server simply *blinds* the index of each state in each round of the protocol. The full  
 375 protocol is described below:

- 376 1. The client initializes her current state to  $s_0 \leftarrow 1$ . The server initializes  $\alpha_0 \leftarrow 0$ .
- 377 2. On round  $i = 1, \dots, n - 1$  of the protocol, the client makes an OT query on the index  $(s_{i-1}, x_i)$ .  
 378 On round  $i$  of the protocol, the server samples a random blinding factor  $\alpha_i \leftarrow \{1, \dots, |Q_i|\}$ . It  
 379 prepares a table of size  $|Q_i| \cdot |\Sigma|$  where the message  $y_i$  associated with index  $(q, \sigma) \in Q_{i-1} \times \Sigma$   
 380 is  $y_i = \delta_i(q - \alpha_{i-1} \bmod |Q_{i-1}|, \sigma) + \alpha_i \bmod |Q_i|$ . The server uses this collection of  $|Q_i| \cdot |\Sigma|$  as  
 381 its set of messages in the OT protocol. Let  $s_i \in \{1, \dots, |Q_i|\}$  be the client's output in the OT  
 382 protocol.
- 383 3. On the final round of the protocol, the client and server run a 1-out-of- $(|Q_{n-1}| |\Sigma|)$  protocol  
 384 where the client's input is  $(s_{n-1}, x_n)$  and the server associates the message 1 with index  
 385  $(q, \sigma) \in Q_{n-1} \times \Sigma$  if  $\delta_n(q - \alpha_{n-1} \bmod |Q_{n-1}|, \sigma) \in S$  and 0 otherwise.

386  
 387 **Correctness.** Let  $z_0 = 1, z_1, \dots, z_n$  be the sequence of states in the evaluation of  $M(x)$ . Namely,  $z_i =$   
 388  $\delta_i(z_{i-1}, x_i)$  for each  $1 \leq i \leq n$ . First, we show that for each  $0 \leq i \leq n$ ,  $s_i - \alpha_i = z_i \bmod |Q_i|$ . We  
 389 proceed inductively. The claim trivially holds for  $i = 0$ . By correctness of the OT protocol,

$$390 \quad s_{i+1} = \delta_{i+1}(s_i - \alpha_i \bmod |Q_i|, x_{i+1}) + \alpha_{i+1} \bmod |Q_{i+1}| = \delta_{i+1}(z_i \bmod |Q_i|, x_{i+1}) + \alpha_{i+1} \bmod |Q_{i+1}|.$$

391 Thus,

$$392 \quad s_{i+1} - \alpha_{i+1} = \delta_{i+1}(z_i, x_{i+1}) \bmod |Q_{i+1}| = z_{i+1} \bmod |Q_{i+1}|,$$

393 and the claim holds. In the final OT, the output is 1 if

$$394 \quad \delta_n(s_{n-1} - \alpha_{n-1} \bmod |Q_{n-1}|, \sigma) = \delta_n(z_{n-1}, x_n) = z_n \in S,$$

395 and 0 otherwise. Similarly, the output of  $M(x)$  is 1 if and only if  $z_n \in S$ , and correctness follows.

396 **Security.** We show that in the OT-hybrid model, the above protocol is secure in the presence of semi-  
 397 honest adversaries. Recall that in the OT-hybrid model, the client and server interact with an “ideal” OT  
 398 functionality (i.e., the client sends an index  $i \in \{1, \dots, k\}$  to the OT functionality, the server sends a  
 399 collection of messages  $m_1, \dots, m_k$ , and the OT functionality replies to the client with  $m_i$ ). If we then  
 400 instantiate the OT with any concrete protocol that is secure against semi-honest adversaries, we obtain  
 401 a secure protocol for evaluating layered DFAs. We consider client and server security separately:

- 402 • **Client Security (Query Privacy).** In the OT-hybrid model, the server only provides inputs to the  
 403 ideal OT functionality. It does not receive any message from either the client or the ideal OT  
 404 functionality. Thus, the view of the server is *independent* of the client’s query so security for the  
 405 client against a semi-honest server is immediate.
- 406 • **Server Security (Function Privacy).** In the OT-hybrid model, the client receives a sequence of  
 407 values  $s_1, \dots, s_n$  from the ideal OT functionality. From the above correctness analysis, we have  
 408 that  $s_n = M(x) \in \{0,1\}$ . By construction (and correctness of the OT protocol), each of the  
 409 intermediate indices  $s_i$  for  $1 \leq i < n$  satisfies  $s_i = \delta_i(s_{i-1}, x_i) + \alpha_i \bmod |Q_i|$ . Here, the server  
 410 samples each  $\alpha_i$  uniformly from the set  $\{1, \dots, |Q_i|\}$ , and in particular, independently of  $\delta_i$ . This  
 411 means that the value of each  $s_i$  is uniform and independent over  $\{1, \dots, |Q_i|\}$ . (Formally, we can  
 412 construct a simulator that on input  $M(x)$  outputs random indices  $s_i \leftarrow \{1, \dots, |Q_i|\}$  for  $1 \leq i < n$   
 413 and  $s_n = M(x)$ . The output of this simulator is identically distributed as the receiver’s view in  
 414 the protocol. Thus, the protocol provides perfect security in the OT-hybrid model.)

415 **Efficiency.** For a layered DFA with  $n + 1$  layers, the protocol as described requires  $n$  rounds of  
 416 communication, where on round  $1 \leq i \leq n$ , the client and server perform a single 1-out-of- $(|Q_{i-1}||\Sigma|)$   
 417 OT on  $\lceil \log |Q_i| \rceil$ -bit messages. Using precomputed OT correlation, the total communication in bits is then

$$418 \quad \sum_{i=1}^n \lceil \log(|Q_{i-1}||\Sigma|) \rceil + \sum_{i=1}^n |Q_{i-1}||\Sigma| \lceil \log |Q_i| \rceil.$$

419

## 420 [Secure evaluation of the threshold matching function](#)

421 We now describe how to use our layered DFAs evaluation protocol to obtain a secure protocol for the  
 422 threshold matching function  $\text{TM}_k(\mathbf{v}, \mathbf{w})$ . As discussed above, we can write  $\text{TM}_k(\mathbf{v}, \mathbf{w}) =$

423  $h_{(1,1,\dots,1),k}(g_{v_1}(w_1), \dots, g_{v_n}(w_n))$  and also represent the functions  $g_{v_i}$  and  $h_{(1,1,\dots,1),k}$  as layered DFAs.

424 We can leverage our protocol for secure layered DFA evaluation to evaluate  $g_{v_i}$  and  $h_{(1,1,\dots,1),k}$ .

425 However, directly applying our privacy-preserving protocol for layered DFAs does not yield a secure  
 426 protocol for evaluating  $\text{TM}_k$  since such a protocol would leak the intermediate values

427  $g_{v_1}(w_1), \dots, g_{v_n}(w_n)$  to the client. This means the client would learn the *exact* set of components that

428 match between its vector  $\mathbf{w}$  and the server’s vector  $\mathbf{v}$  (irrespective of whether the threshold is satisfied

429 or not). We address this by *blinding* the output of the equality-checking circuits.

430 As currently defined  $g_{v_i}(w_i)$  outputs 1 if  $v_i = w_i$  and 0 otherwise. We can blind the equality bit by  
 431 having the server sample a uniform bit  $a_i \leftarrow \{0,1\}$  and define  $g'_{v_i,a_i}(w_i) = a_i$  if  $v_i = w_i$  and  $1 - a_i$   
 432 otherwise. We make a few simple observations:

- 433 • If we have a layered DFA for computing  $g_{v_i}$ , then the same DFA (from Figure 1) can be used to  
 434 compute  $g'_{v_i,a_i}$  by swapping the accept and reject states in the final layer of the DFA.
- 435 • Let  $\beta_i \leftarrow g'_{v_i,a_i}(w_i)$ . By construction, this means that  $\beta_i \oplus a_i \oplus 1 = g_{v_i}(w_i)$ . If the client and  
 436 server apply the privacy-preserving protocol for evaluating a layered DFA to  $g'_{v_i,a_i}$ , then at the  
 437 end of the protocol, the client learns  $\beta_i$  (and nothing more) while the server learns nothing. In  
 438 this case,  $\beta_i$  is uniformly random (it is perfectly hidden by  $a_i$ ), and thus, the client does not learn  
 439 anything about the value of  $g_{v_i,a_i}(w_i)$ . (More precisely, we say that the client and the server  
 440 have a “secret sharing” of the negated equality bit  $g_{v_i,a_i}(w_i) \oplus 1$ .)
- 441 • The client and server use the privacy-preserving protocol for computing layered DFAs to  
 442 compute the quantity

$$443 \quad h_{(a_1, \dots, a_n), k} \left( g'_{v_1, a_1}(w_1), \dots, g'_{v_n, a_n}(w_n) \right) = h_{(1, 1, \dots, 1), k} \left( g_{v_1}(w_1), \dots, g_{v_n}(w_n) \right).$$

444 The overall protocol is now given as follows:

- 445 1. The server begins by sampling  $a_1, \dots, a_n \leftarrow \{0,1\}$ . For each  $i = 1, \dots, n$ , the client and server  
 446 execute the protocol for layered DFA evaluation for the function  $g'_{v_i, a_i}$ . On the  $i^{\text{th}}$  iteration of  
 447 the protocol, the client provides  $w_i$  as its input. Let  $b_1, \dots, b_n \in \{0,1\}$  be the client’s output in  
 448 the  $n$  protocol executions.
- 449 2. The client and server execute the protocol for layered DFA evaluation for the function  
 450  $h_{(a_1, \dots, a_n), k}$  where the client uses  $b_1, \dots, b_n$  as its input. The output  $z$  is the result of the threshold  
 451 matching function  $\text{TM}_k(\mathbf{v}, \mathbf{w})$ .

452 **Correctness.** Correctness of the above protocol follows via correctness of the underlying layered DFA  
 453 evaluation protocol. Take any two inputs  $\mathbf{v}$  and  $\mathbf{w}$ . Then, we have that  $b_i = g'_{v_i, a_i}(w_i)$  for all  $1 \leq i \leq n$ .  
 454 The client’s output at the end of the computation is

$$455 \quad h_{(a_1, \dots, a_n), k}(b_1, \dots, b_n) = h_{(a_1, \dots, a_n), k} \left( g'_{v_1, a_1}(w_1), \dots, g'_{v_n, a_n}(w_n) \right) = h_{(1, 1, \dots, 1), k}(g_{v_1}(w_1), \dots, g_{v_n}(w_n)),$$

456 which is precisely the value of  $\text{TM}_k(\mathbf{v}, \mathbf{w})$ .

457 **Security.** Security of the protocol also follows by security of the underlying protocol for evaluating  
 458 layered DFAs. As before, we consider client and server security separately:

- 459 • **Client Security (Query Privacy).** The threshold matching protocol consists of  $n + 1$  invocations  
 460 of the protocol for layered DFA evaluation. Security of our protocol for layered DFA evaluation  
 461 ensures that the server does not learn anything about the client’s input in any of these protocol  
 462 executions, and security follows. (More formally, security of the layered DFA evaluation protocol  
 463 implies that the server’s view can be simulated *without* knowledge of the client’s input or  
 464 output, and correspondingly, the server’s view in the threshold matching protocol can be  
 465 simulated by simply running the simulator for the underlying layered DFA evaluation protocol).
- 466 • **Server Security (Database Privacy).** Security for the server follows similarly from security of the  
 467 underlying layered DFA evaluation protocol. In the above protocol, the client learns the values

468 of  $b_1, \dots, b_n$  and  $z$ . From the correctness analysis, we have that  $z = \text{TM}_k(\mathbf{v}, \mathbf{w})$ . Next, by  
 469 security of the layered DFA evaluation protocol (as applied to the computation of  $h_{(a_1, \dots, a_n), k}$ ),  
 470 the client learns no additional information other than the value  $z$ . In particular, the protocol  
 471 does not leak any information about the values of  $a_1, \dots, a_n$  to the client. By correctness of the  
 472 layered DFA evaluation protocol, each bit  $b_i$  the client obtains satisfies  $b_i = g_{v_i}(w_i) \oplus a_i \oplus 1$ .  
 473 Security of the layered DFA evaluation protocol (as applied to the computation of  $g'_{v_i, a_i}$ ) implies  
 474 that the client learns nothing more about  $v_i, a_i$  other than the value  $b_i = g_{v_i}(w_i) \oplus a_i \oplus 1$ . In  
 475 particular, this means that the client does not learn anything about the value of  $a_i$  from the  
 476 layered DFA protocol executions. Since the server samples  $a_i \leftarrow \{0, 1\}$  to be a uniformly random  
 477 bit, the distribution of  $b_i$  is also uniform (even conditioned on the client's view of the protocol  
 478 execution, which is independent of  $a_i$ ). As such, the client's view in the protocol execution can  
 479 be described by a sequence of uniform random bits  $b_1, \dots, b_n$  together with the output  $z =$   
 480  $\text{TM}_k(\mathbf{v}, \mathbf{w})$ . Thus, the client does not learn anything more about  $\mathbf{v}$  other than the value  
 481  $\text{TM}_k(\mathbf{v}, \mathbf{w})$ . (More formally, we can simulate the client's view of the protocol based on the value  
 482 of  $z = \text{TM}_k(\mathbf{v}, \mathbf{w})$  and then sampling the bits  $b_1, \dots, b_n \leftarrow \{0, 1\}$ ).

483 **Efficiency.** We now analyze the cost of privately computing  $\text{TM}_k(\mathbf{v}, \mathbf{w})$  for two  $n$ -dimensional  
 484 vectors  $\mathbf{v}, \mathbf{w}$  with  $n$ -bit components where each component  $v_i, w_i \in \{0, 1\}^\ell$  is an  $\ell$ -bit string.

- 485 • **Equality checking.** Computing the (blinded) equality-check functions  $g'_{v_i, a_i}(w_i)$  requires  
 486 evaluating a DFA with  $\ell + 1$  layers, where each layer contains 2 nodes. This requires  $\ell$   
 487 rounds of communication and  $6\ell$  bits of communication.
- 488 • **Thresholding.** Computing the threshold function  $h_{(a_1, \dots, a_n), k}$  requires evaluating a DFA with  
 489  $n + 1$  layers, where each layer contains  $k + 2$  nodes. This requires  $n$  rounds of  
 490 communication and  $n(1 + (2k + 5)\lceil \log(k + 2) \rceil)$  bits of communication.

491 Finally, we note that the equality checks can be conducted in *parallel*. Thus, the round complexity of the  
 492 complete protocol is  $n + \ell$  and the total communication (in bits) is

$$493 \quad n((2k + 5)\lceil \log(k + 2) \rceil + 6\ell + 1) = O(nk \log k + n\ell).$$

494 We can achieve a tradeoff between communication complexity and round complexity by having each  
 495 DFA read multiple bits of the input at a time. This is equivalent to considering a DFA with a larger  
 496 alphabet (i.e.,  $\Sigma = \{0, 1\}^m$  if the DFA is reading  $m$  bits of the input for each state transition). Reading  
 497 multiple bits of the input per state transition decreases the number of layers in the DFA (thus decreasing  
 498 the round complexity of the private layered DFA evaluation protocol), but increases the size of the  
 499 transition table between each layer (thus increasing the communication complexity of the protocol). In  
 500 particular, if we consider DFAs that read  $m$ -bits of input per transition, then the overall round  
 501 complexity of the protocol becomes  $\lceil \frac{\ell}{m} \rceil + \lceil \frac{n}{m} \rceil$  while the communication complexity (in bits) is

$$502 \quad \lceil \frac{\ell}{m} \rceil n(m + 1 + 2^{m+1}) + \lceil \frac{n}{m} \rceil (m + (2^m(k + 2) + 1)\lceil \log(k + 2) \rceil).$$

503 When  $\ell$  and  $m$  are even, then setting  $m = 2$  (i.e., reading 2 bits at a time) yields a protocol with smaller  
 504 round complexity compared to the  $m = 1$  setting (by roughly a factor of 2) *and* slightly smaller  
 505 communication complexity (compared to both the case where  $m = 1$  and all  $m > 2$ ). This is the setting  
 506 we use in our implementation.

## 507 [Implementation and evaluation](#)

508 We implemented the protocol in C++. For our benchmarks, we conducted experiments using two  
509 Amazon EC2 instances (M4.2xlarge). Each instance has an 8-core 2.4 GHz Intel Xeon E5-2676 v3  
510 (Haswell) processor and 32 GB of memory. The client instance is located on the West Coast (Northern  
511 California region) while the server is located on the East Coast (Northern Virginia region) to simulate a  
512 wide-area network (WAN). The network latency between the two instances is roughly 60ms and the  
513 bandwidth is roughly 25 MB/s. We use a single-threaded execution environment for all of our  
514 experiments.

## 515 [Code availability](#)

516 All code will be made publicly available on Github for non-commercial use upon publication.

## 517 [Author Contributions](#)

518 JAB, KAJ, GB, and DJW designed the study, analyzed results, and wrote the manuscript. JAB wrote  
519 software for the analysis with input from KAJ, GB, and DJW.

## 520 [Competing Interests](#)

521 The authors declare no competing interests.

## 522 [Acknowledgements](#)

523 We thank Benton Case and Dan Boneh for helpful discussions in an early phase of this project and Aviv  
524 Regev for support (K.A.J.). This work was also supported by the Joint University Microelectronics  
525 Program (JUMP) Undergraduate Research Initiative (J.A.B), the Stanford A.I. Lab (G.B), NSF CNS-1917414  
526 (D.J.W) and a University of Virginia SEAS Research Innovation Award (D.J.W).

527

## 528 References

- 529 1. Arnaud, C. Thirty years of DNA forensics: How DNA has revolutionized criminal  
530 investigations. *Chemical and Engineering News* vol. 95 16–20 (2017).
- 531 2. CODIS - NDIS Statistics. *Federal Bureau of Investigation*  
532 <https://www.fbi.gov/services/laboratory/biometric-analysis/codis/ndis-statistics>.
- 533 3. Combined DNA Index System (CODIS). *Federal Bureau of Investigation*  
534 <https://www.fbi.gov/services/laboratory/biometric-analysis/codis>.
- 535 4. Ransom, J. & Southall, A. ‘Race-Biased Dragnet’: DNA From 360 Black Men Was  
536 Collected to Solve Vetrano Murder, Defense Lawyers Say. *The New York Times* (2019).
- 537 5. Ransom, J. & Southall, A. N.Y.P.D. Detectives Gave a Boy, 12, a Soda. He Landed in a  
538 DNA Database. *The New York Times* (2019).
- 539 6. Dickerson, C. U.S. Government Plans to Collect DNA From Detained Immigrants. *The New*  
540 *York Times* (2019).
- 541 7. Murphy, H. Coming Soon to a Police Station Near You: The DNA ‘Magic Box’. *The New*  
542 *York Times* (2019).
- 543 8. Crowley, M. How Commandos Could Quickly Confirm They Got Their Target. *The New*  
544 *York Times* (2019).
- 545 9. What is Rapid DNA? *ANDE Rapid DNA* <https://www.ande.com/what-is-rapid-dna/>.
- 546 10. RapidHIT ID System for Human Identification - US.  
547 [https://www.thermofisher.com/us/en/home/industrial/forensics/human-](https://www.thermofisher.com/us/en/home/industrial/forensics/human-identification/forensic-dna-analysis/dna-analysis/rapidhit-id-system-human-identification.html)  
548 [identification/forensic-dna-analysis/dna-analysis/rapidhit-id-system-human-](https://www.thermofisher.com/us/en/home/industrial/forensics/human-identification/forensic-dna-analysis/dna-analysis/rapidhit-id-system-human-identification.html)  
549 [identification.html](https://www.thermofisher.com/us/en/home/industrial/forensics/human-identification/forensic-dna-analysis/dna-analysis/rapidhit-id-system-human-identification.html).
- 550 11. Wee, S.-L. China Is Collecting DNA From Tens of Millions of Men and Boys, Using U.S.  
551 Equipment. *The New York Times* (2020).
- 552 12. Joly, Y. *et al.* Establishing the International Genetic Discrimination Observatory. *Nat. Genet.*  
553 **52**, 466–468 (2020).
- 554 13. NYPD’s ‘Knock-and-Spit’ DNA Database Makes You a Permanent Suspect. *Newsweek*  
555 [https://www.newsweek.com/police-dna-database-nypd-swab-testing-collection-new-york-](https://www.newsweek.com/police-dna-database-nypd-swab-testing-collection-new-york-1326722)  
556 [1326722](https://www.newsweek.com/police-dna-database-nypd-swab-testing-collection-new-york-1326722) (2019).
- 557 14. NIST. Core STR Loci Used in Human Identity Testing.  
558 <https://strbase.nist.gov/coreSTRs.htm>.
- 559 15. Wang, Z. *et al.* Developmental Validation of the Huaxia Platinum System and application in  
560 3 main ethnic groups of China. *Sci. Rep.* **6**, 31075 (2016).
- 561 16. Norrgard, K. Forensics, DNA Fingerprinting, and CODIS. *Nat. Educ.* **1**, 35 (2008).

- 562 17. Frequently Asked Questions on CODIS and NDIS. *Federal Bureau of Investigation*  
563 <https://www.fbi.gov/services/laboratory/biometric-analysis/codis/codis-and-ndis-fact-sheet>.
- 564 18. ENFSI DNA Working Group. Matching rules. in *DNA DATABASE MANAGEMENT*  
565 *REVIEW AND RECOMMENDATIONS* 22–25 (ENFSI, 2017).
- 566 19. Hopcroft, J. E., Motwani, R. & Ullman, J. D. *Introduction to Automata Theory, Languages,*  
567 *and Computation*. (Pearson, 2006).
- 568 20. Kilian, J. Founding Cryptography on Oblivious Transfer. in *STOC* 20–31 (1988).
- 569 21. Rabin, M. O. How To Exchange Secrets with Oblivious Transfer. *IACR Cryptol. EPrint*  
570 *Arch.* **2005**, 187 (2005).
- 571 22. Kolesnikov, V., Kumaresan, R., Rosulek, M. & Trieu, N. Efficient Batched Oblivious PRF  
572 with Applications to Private Set Intersection. in *ACM CCS* 818–829 (2016).
- 573 23. O’Connor, K. L., Butts, E., Hill, C. R., Butler, J. & Vallone, P. Evaluating the effect of  
574 additional forensic loci on likelihood ratio values for complex kinship analysis. in  
575 *Proceedings of the 21st International Symposium on Human Identification* 10–14 (2010).
- 576 24. Doleac, J. L. The Effects of DNA Databases on Crime. *Am. Econ. J. Appl. Econ.* **9**, 165–201  
577 (2017).
- 578 25. Yao, A. C.-C. Protocols for Secure Computations. in *FOCS* 160–164 (1982).
- 579 26. Hazel, J. W., Clayton, E. W., Malin, B. A. & Slobogin, C. Is it time for a universal genetic  
580 forensic database? *Science* **362**, 898–900 (2018).
- 581 27. Joly, Y., Marrocco, G. & Dupras, C. Risks of compulsory genetic databases. *Science* **363**,  
582 938–940 (2019).
- 583 28. Naor, M. & Pinkas, B. Oblivious Transfer and Polynomial Evaluation. in *STOC* 245–254  
584 (1999).
- 585 29. Canetti, R. Security and Composition of Multiparty Cryptographic Protocols. *J Cryptol.* **13**,  
586 143–202 (2000).
- 587 30. Ishai, Y., Kilian, J., Nissim, K. & Petrank, E. Extending Oblivious Transfers Efficiently. in  
588 *CRYPTO* 145–161 (2003).
- 589 31. Boyle, E. *et al.* Efficient Two-Round OT Extension and Silent Non-Interactive Secure  
590 Computation. in *ACM CCS* 291–308 (2019).
- 591 32. Troncoso-Pastoriza, J. R., Katzenbeisser, S. & Celik, M. U. Privacy preserving error resilient  
592 dna searching through oblivious automata. in *ACM CCS* 519–528 (2007).
- 593 33. Sasakawa, H. *et al.* Oblivious Evaluation of Non-deterministic Finite Automata with  
594 Application to Privacy-Preserving Virus Genome Detection. in *Workshop on Privacy in the*  
595 *Electronic Society (WPES)* 21–30 (2014).

596



597 Tables and Figures

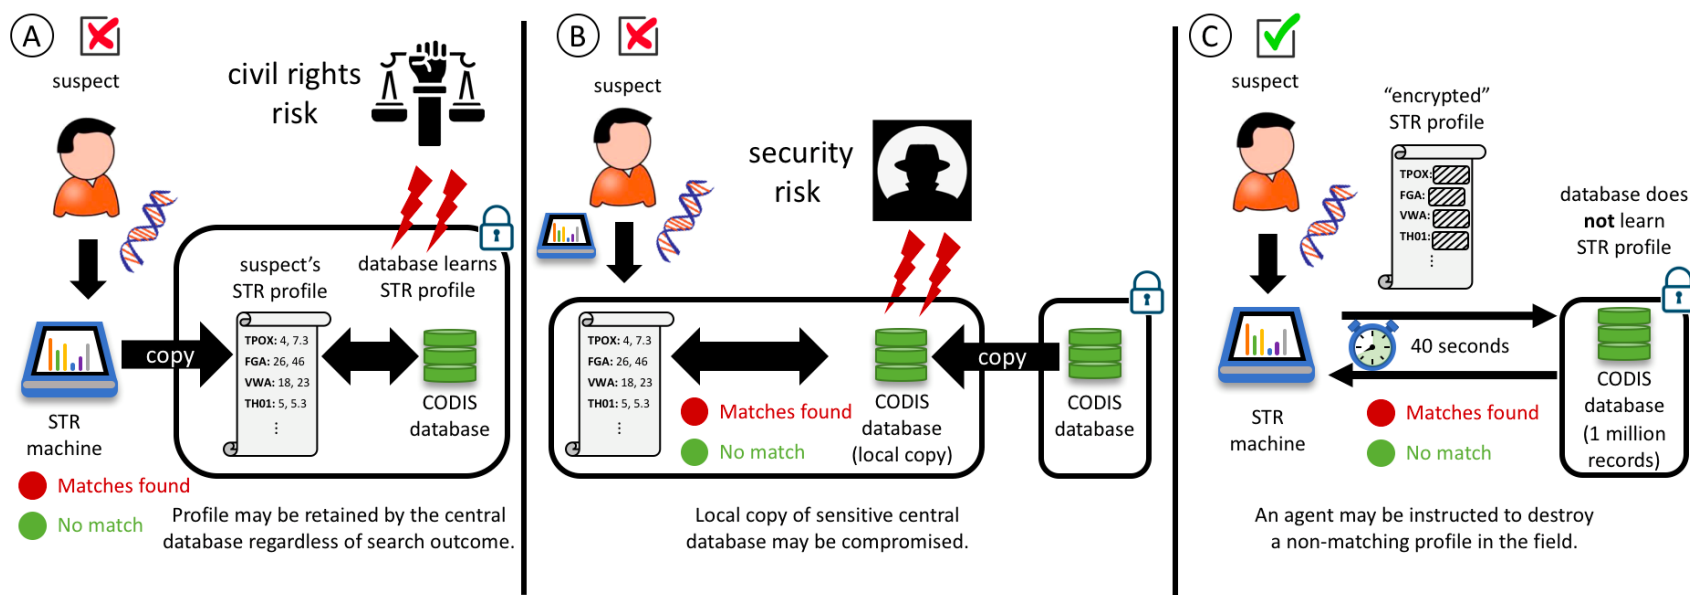
598 Table 1

<b>CODIS System</b>	<b>End-to-End Time (seconds)</b>	<b>Bandwidth (MB)</b>
<b>US System, current (20 STRs)</b>	38.4	172.4
<b>US System, pre-2017 (13 STRs)</b>	27.7	114.9
<b>UK-like System (11 STRs)</b>	23.4	97.3
<b>EU-like System (16 STRs)</b>	34.1	154.5
<b>Chinese-like System (20 STRs)</b>	40.0	189.0

599

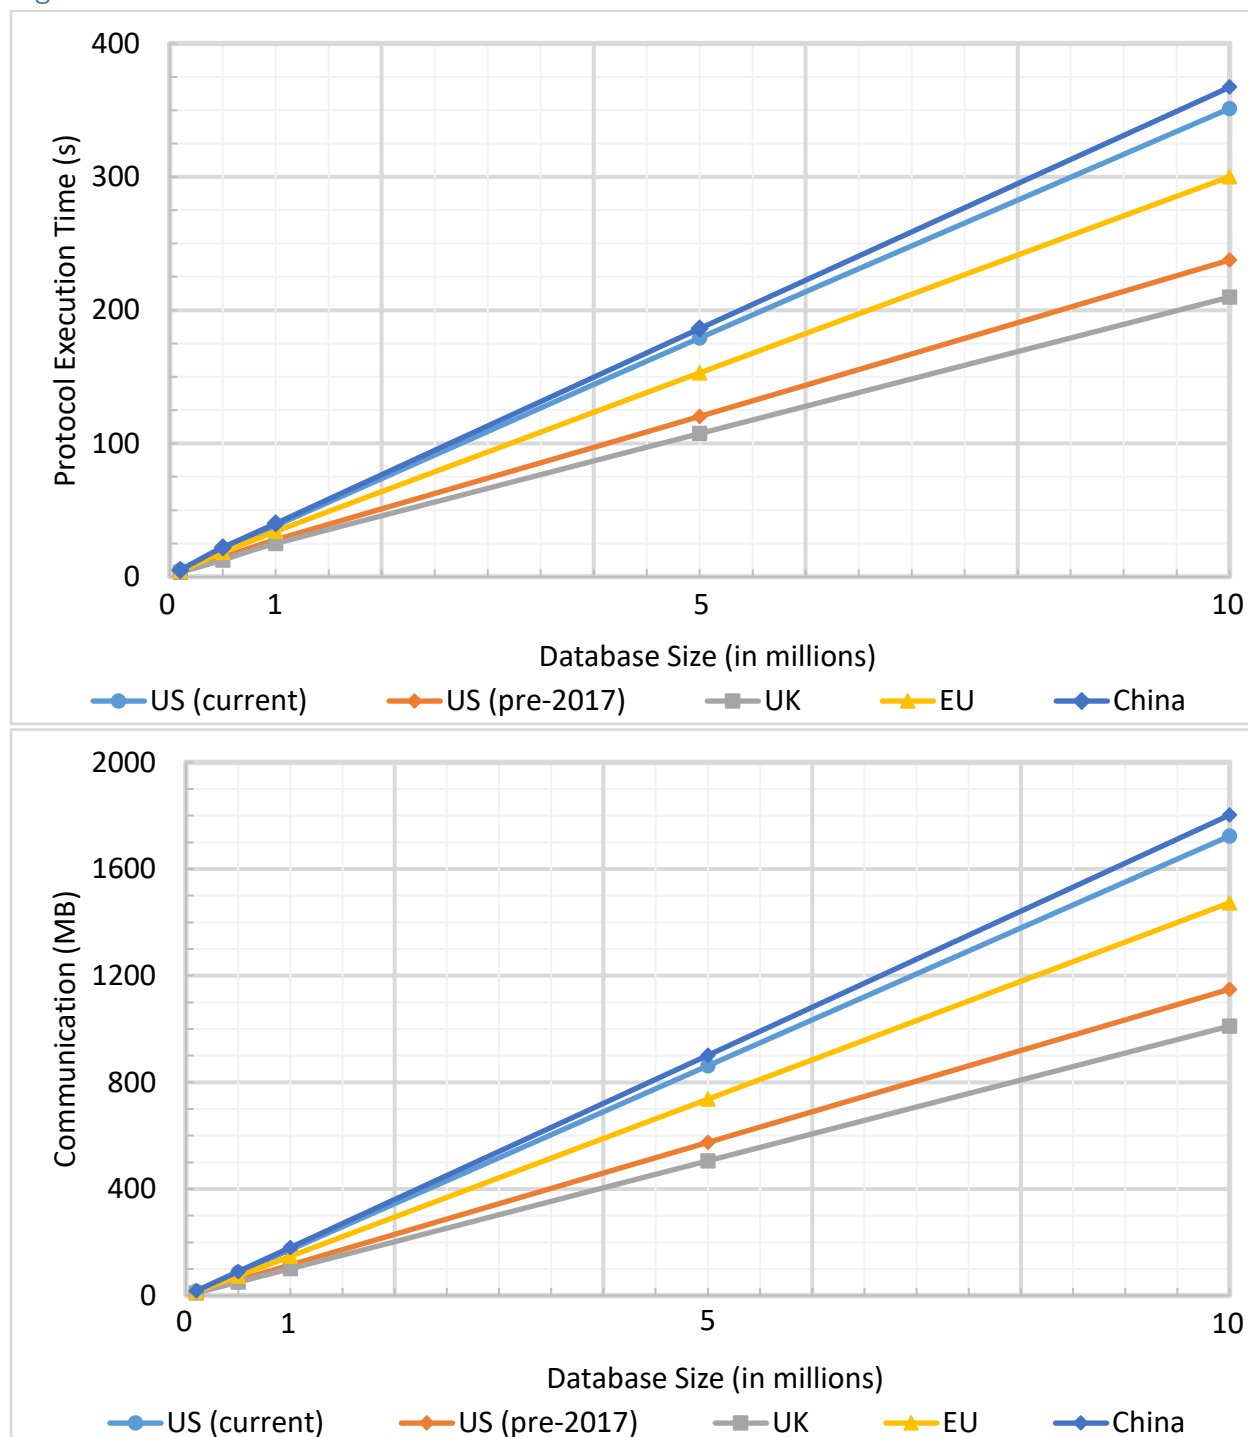
600 **Table 1. Runtime and network communication needed to privately compare a suspect CODIS profile in**  
601 **the field to a central office database far away.** End-to-end protocol execution time and communication  
602 required to privately query a central database of 1,000,000 entries for CODIS systems deployed at  
603 different countries (see Online Methods). Here, the client and server are Amazon EC2 instances, with  
604 the client located on the West Coast of the U.S. and the server located on the East Coast of the U.S.

Figure 1



**Figure 1. A novel privacy-preserving CODIS DNA profile matching protocol.** Rapid STR profiling technologies enables genetic testing and matching from the field. They provide a valuable tool for crime solving but raise significant civil rights concerns regarding data retention and racial profiling. (A) Today, STR profiles collected from potentially innocent individuals are sent to a central CODIS database to check for matches. Here, the central database learns the full query profile and has the option to retain it, irrespective of whether the search yields any match or not. (B) To try and provide anonymity to exonerated (unmatched) profiles, one may load a private copy of the central database onto every field device. This way, an unmatched suspect profile may be destroyed in the field to retain suspect privacy. However, this approach risks exposure of all or parts of the sensitive central database to malicious parties who get ahold of a field device. (C) Our privacy-preserving search protocol enables a new approach where agents can still query a central CODIS database as in (A), but in a way that completely hides the query from the central database. The agent still learns the outcome of the query as before. However, an innocent profile, not matching anything in the database may safely be destroyed in the field. The central database can no longer store it, as it has learned nothing about it through the query.

Figure 2



**Figure 2. Single query search performance against an entire CODIS database as a function of database size.** The number of STR loci and number of precision bits assumed for each system are described in Supplementary Table 1.

## Supplementary Tables and Figures

### Supplementary Table 1

<b>CODIS System</b>	<b># of Loci</b>	<b># of Bits</b>	<b>Bits per Locus Breakdown</b>
<b>US System, current</b>	20	212	10 bits (17 loci); 14 bits (3 loci)
<b>US System, pre-2017</b>	13	142	10 bits (10 loci); 14 bits (3 loci)
<b>UK-like System</b>	11	126	10 bits (7 loci); 14 bits (4 loci)
<b>EU-like System</b>	16	178	10 bits (12 loci); 14 bits (3 loci); 16 bits (1 locus)
<b>Chinese-like System</b>	20	224	10 bits (14 loci); 14 bits (6 loci)
<b>NIST 40 Loci System</b>	40	492	10 bits (17 loci); 14 bits (23 loci)


**Supplementary Table 1. Approximate CODIS system specifications for select countries.** Number of STR loci and number of bits used to encode alleles at each locus for a handful of the 50+ countries using the CODIS system (See Online Methods).

## Supplementary Table 2

CODIS System	Precomputation Size (MB)		# of OT Correlations (in millions)
	Client	Server	
US System, current (20 STRs)	122	923	116
US System, pre-2017 (13 STRs)	87	616	78
UK-like System (11 STRs)	71	574	73
EU-like System (16 STRs)	104	763	96
Chinese-like System (20 STRs)	122	969	122
NIST 40 Loci System (40 STRs)	260	2105	266


**Supplementary Table 2. Offline precomputation cost for each setting.** Number of oblivious transfer (OT) correlations, and the memory footprint of the OT correlations for the client and server needed to implement a single CODIS search query against a database with 1,000,000 records. Using state-of-the-art OT extension protocols<sup>22</sup>, it is possible to setup  $2^{24} > 16$  million OT correlations in 7.5 seconds over a wide-area network. For system specifications, see Online Methods.

### Supplementary Figure 1



**Supplementary Figure 1. Layered DFA for equality test.** This DFA computes the equality-check function  $g_v(\mathbf{w})$  that outputs 1 if  $\mathbf{v} = \mathbf{w}$  and 0 otherwise. In particular, for a vector  $\mathbf{v} = (v_1, \dots, v_n) \in \{0,1\}^n$ , this DFA only accepts the input  $\mathbf{w} = (w_1, \dots, w_n) \in \{0,1\}^n$  where  $v_i = w_i$  for all  $1 \leq i \leq n$ . We use this DFA to decide whether there is a match at a single STR locus. If we denote the single start state as “layer 0”, the two states one can arrive at from layer 0 after reading the first bit as “layer 1”, etc. we see that this DFA has  $n + 1$  layers, such that after reading  $i$  bits, it can only be in one of the two states in layer  $i$ .

## Supplementary Figure 2



**Supplementary Figure 2. Layered DFA for thresholding.** This DFA computes the threshold function  $h_{(a_1, \dots, a_n), k}$  for the  $k = 1$  case. Namely,  $h_{(a_1, \dots, a_n), k}(b_1, \dots, b_n)$  outputs 1 if  $a_i = b_i$  for all but at most  $k$  indices  $1 \leq i \leq n$ . In other words, for any sequence of bits  $(a_1, \dots, a_n) \in \{0, 1\}^n$ , this DFA accepts if the input  $b_1, \dots, b_n$  satisfies  $b_i = a_i$  for all but at most one index  $i$ . For instance, in this work, we use this DFA to decide whether a DNA profile matches against a database record on at least 19 out of 20 loci (i.e., the setting where  $k = 1$  and  $n = 20$ ) as well as the other configurations. Here, the  $i^{\text{th}}$  input bit  $b_i \in \{0, 1\}$  is the (blinded) equality bit denoting whether there is a match in the  $i^{\text{th}}$  STR locus (between the agent's query and the central database's record). In our protocol, this (blinded) equality bit is computed using the equality-test DFA from Supplementary Figure 1. The bits  $a_1, \dots, a_n$  in the function description  $h_{(a_1, \dots, a_n), k}$  are the blinding values chosen by the server. Recall that the blinding is introduced to hide from the client all information on whether there was a match at STR locus  $i$  between the database server's profile and the client's query. The client only learns whether her query matches the record or not, and nothing more. Much like Supplementary Figure 1, this DFA has  $n + 1$  layers, such that after reading  $i$  bits, the computation can only be in one of the (at most) 3 states of layer  $i$ .