

Genome analysis

Re-identification of individuals in genomic data-sharing beacons via allele inference

Nora von Thenen¹, Erman Ayday^{1,2,*} and A. Ercument Cicek ^{1,3,*}

¹Computer Engineering Department, Bilkent University, Ankara 06800, Turkey, ²Department of Electrical Engineering and Computer Science, Case Western Reserve University, Cleveland, OH 44106, USA and ³Computational Biology Department, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

*To whom correspondence should be addressed.

Associate Editor: John Hancock

Received on February 14, 2018; revised on May 31, 2018; editorial decision on July 12, 2018; accepted on July 18, 2018

Abstract

Motivation: Genomic data-sharing beacons aim to provide a secure, easy to implement and standardized interface for data-sharing by only allowing yes/no queries on the presence of specific alleles in the dataset. Previously deemed secure against re-identification attacks, beacons were shown to be vulnerable despite their stringent policy. Recent studies have demonstrated that it is possible to determine whether the victim is in the dataset, by repeatedly querying the beacon for his/her single-nucleotide polymorphisms (SNPs). Here, we propose a novel re-identification attack and show that the privacy risk is more serious than previously thought.

Results: Using the proposed attack, even if the victim systematically hides informative SNPs, it is possible to infer the alleles at positions of interest as well as the beacon query results with very high confidence. Our method is based on the fact that alleles at different loci are not necessarily independent. We use linkage disequilibrium and a high-order Markov chain-based algorithm for inference. We show that in a simulated beacon with 65 individuals from the European population, we can infer membership of individuals with 95% confidence with only 5 queries, even when SNPs with MAF <0.05 are hidden. We need less than 0.5% of the number of queries that existing works require, to determine beacon membership under the same conditions. We show that countermeasures such as hiding certain parts of the genome or setting a query budget for the user would fail to protect the privacy of the participants.

Availability and implementation: Software is available at http://ciceklab.cs.bilkent.edu.tr/beacon_attack.

Contact: erman@cs.bilkent.edu.tr or cicek@cs.bilkent.edu.tr

Supplementary information: [Supplementary data](#) are available at *Bioinformatics* online.

1 Introduction

Exciting times are on the horizon for the genomics field with the announcement of the precision medicine initiative (Collins and Varmus, 2015) which was followed by the \$55 million funding by NIH for the sequencing of a million individuals and AstraZeneca's project of sequencing two million individuals (Ledford, 2016). Even though such million-sized genomic datasets are invaluable resources for research, sharing the data is a big challenge due to re-identification risk. Several studies in the last decade have shown that

removal of personal identifiers from genomic data is not enough and that individuals can be re-identified using allele frequency information (Clayton, 2010; Homer *et al.*, 2008; Jacobs *et al.*, 2009; Sankararaman *et al.*, 2009; Visscher and Hill, 2009).

Genomic data-sharing beacons (referred to as beacons from now on) are the gateways that let users and data owners exchange information without—in theory—disclosing any personal information. A user who wants to apply for access to the dataset can learn whether individuals with specific alleles of interest are present in the beacon

through an online interface. More specifically, the user submits a query, asking whether a genome exists in the beacon with a certain nucleotide at a certain position, and the beacon answers ‘yes’ or ‘no’. Beacons are easy to set up systems that provide very restricted access to the stored data. The Beacon Project is an initiative by the Global Alliance for Genomics and Health (GA4GH) which creates policies to ensure standardized and secure sharing of genomic data.

Beacons make it difficult for adversaries to re-identify individuals due to several reasons. First, the data access policy is very restrictive and allows only presence/absence queries for nucleotides for specific positions in the genome. Given possibly large number of individuals in a beacon, a ‘yes’ answer can be due to several individuals and it cannot be tied to a specific person. Second, binary response scheme makes the system secure against attacks that make use of allele frequencies [e.g. (Wang *et al.*, 2009)]. However, it has been shown that such countermeasures are not sufficient to completely prevent the privacy threats raised by genomic data-sharing beacons.

In 2015, Shringarpure and Bustamante introduced a likelihood-ratio test (LRT) that predicts if an individual is in the beacon or not, by repeatedly querying the beacon for single-nucleotide polymorphisms (SNPs) of the victim (dubbed the SB attack) (Shringarpure and Bustamante, 2015). This attack is serious because inferring the membership of an individual in a beacon that is associated with a sensitive phenotype is equivalent to uncovering the sensitive phenotype about the victim. The SB attack does not use the allele frequencies and can compensate sequencing errors. They show that they could re-identify an individual in a beacon with 65 European individuals from the 1000 Genomes Project (Siva, 2008) with 250 queries (with 95% confidence). In their scheme, both the queries posed and the answers received from the beacon are assumed to be independent, therefore the hypothesis is tested based on a binomial test. Very recently, the work by Raisaro *et al.* showed that if the attacker has access to the minor allele frequencies (MAFs) of the population, she/he can sort the victim’s SNPs and query the SNPs starting from the one with the lowest MAF (dubbed the optimal attack) (Raisaro *et al.*, 2016). Unlike the SB attack, queries are not random in this case. As low MAF SNPs are more informative, Raisaro *et al.* show that fewer queries are needed to re-identify an individual. Furthermore, Raisaro *et al.* proposed countermeasures against re-identification attacks such as adding noise to the beacon results and assigning a budget to beacon members which limits the number of informative queries that can be asked on each member.

In this paper, we introduce two new inference-based attacks that (i) carefully select the SNPs to be queried and predict query results of the beacon, and (ii) infer hidden or missing alleles of a victim’s genome. First, we show that if the queried locus is in linkage disequilibrium (LD) with others, it is enough to query for that particular allele, as the attacker can infer the answers of the other alleles with high confidence (Humbert *et al.*, 2013). We refer to this method as the query inference attack (QI-attack). Second, we introduce the genome inference attack (GI-attack) which recovers hidden parts of a victim’s genome by using a high-order Markov chain (Samani *et al.*, 2015).

We show that in a simulated beacon with 65 European individuals (CEU) from the HapMap Project (Gibbs *et al.*, 2003), our QI-attack requires 282 queries and our GI-attack requires only 5 queries on average to re-identify an individual, whereas the SB attack requires 19 525 queries and the optimal attack requires 415 queries, all at the 95% confidence level when the victim’s SNPs with MAFs < 0.03 are hidden. Therefore, the attacker models presented

here can efficiently work when certain regions in the genome of the victim are systematically hidden as a security countermeasure. The number of queries required by the SB and the optimal attacks substantially increase as more SNPs are concealed, while the GI-attack still requires only a few queries on average. Finally, we show that the QI-attack can still re-identify individuals despite the stringent query budget countermeasure proposed by Raisaro *et al.*, 2016 and the beacon censorship countermeasure proposed by Shringarpure and Bustamante.

We demonstrate that the beacons are more vulnerable than previously thought and that the proposed countermeasures in the literature still fail to protect the privacy of the individuals. The contributions of this paper can be summarized as follows:

- By inferring query results and alleles at certain positions, we show that it is possible to significantly decrease the number of required queries compared to other attacks in the literature: Shringarpure and Bustamante; Raisaro *et al.*
- We show that beacons are vulnerable even under a weaker adversary model, in which informative parts of a victim’s genome are concealed (such as all SNPs with an MAF less than a threshold).
- We discuss the feasibility and the effectiveness of the proposed countermeasures in the literature and show that using the presented attack models, the participants are still under risk.

The rest of the manuscript is organized as follows: we describe the methods in Section 2 and then present the results in Section 3. Section 4 discusses the results and the effectiveness of countermeasures proposed in the literature. Finally, we conclude in Section 5.

2 Materials and methods

In this section, we first describe attacker models in the literature [i.e. SB attack (Shringarpure and Bustamante, 2015) and the optimal attack (Raisaro *et al.*, 2016)] and then describe our proposed attacks. In our first proposed model, the attacker not only has access to MAFs of the victim’s population, but also can access or calculate the corresponding LD values from public resources (QI-attack). In the second model, the attacker has the same background knowledge as the QI-attack, and also has access to VCF files of people from the victim’s population from public sources (GI-attack). The four different attacker models [SB attack (Shringarpure and Bustamante, 2015), optimal attack (Raisaro *et al.*, 2016), QI-attack and GI-attack] are described in Figure 1. We consider two scenarios. Scenario 1 assumes the attacker has access to the full genome of the victim. In this case ‘full’ means that part of the DNA of the victim (e.g. a chromosome) is available in full and no locus is systematically hidden. Scenario 2 considers a more realistic and weaker attacker model. As publicly available genomic data is typically found partially, in this scenario, some SNPs are systematically hidden. That is, SNPs with $MAF < t$ are not available to the attacker.

2.1 Background: SB attack and optimal attack

Shringarpure and Bustamante proposed the SB attack, which queries a beacon for the victim’s heterozygous SNP positions. Queried SNPs are picked randomly and a LRT statistic is calculated. The null hypothesis (H_0) refers to the query genome not being in the beacon. Under the alternative hypothesis (H_1), the query genome is a member of the beacon. The attacker model is visualized in Figure 2a.









Adversary + Background	Scenario	Scenario 1 $t = 0$	Scenario 2 $t > 0$
 SB attack	Population of victim	AA + AT + AA	 + AT + AA
 Optimal attack	Population of victim, corresponding MAF	CT + GG + CC	 + GG + 
 QI - attack	Population of victim, corresponding MAF & LD	GT + CC	 + CC
 GI - attack	Population of victim, corresponding MAF, High-Order Correlation		CC + TA + TT

Fig. 1. Four attacker models: SB attack (Shringarpure and Bustamante, 2015), Optimal attack (Raisaro et al., 2016), QI-attack and GI-attack and their background knowledge for two scenarios are shown. In the first scenario $t = 0$ and in the second scenario $t > 0$, where t is the threshold up to which SNPs of the victim with an MAF $< t$ are hidden as a countermeasure. In Scenario 1, the attacker has access to the full genome of the victim (no hidden SNPs). In Scenario 2, SNPs with an MAF $< t$ are hidden and the attacker has partial access to the genome of the victim

The log-likelihood under the null hypothesis has been defined as

$$L_{H_0}(R) = \sum_{i=1}^n x_i \log(1 - D_N) + (1 - x_i) \log(D_N), \quad (1)$$

where R is the response set and D_N the probability that no individual in the beacon has the queried allele at that position. x_i is the answer of the beacon to the query at position i (1 for yes, 0 for no), and n is the total number of posed queries. Accordingly, the log-likelihood of the alternative hypothesis has been stated as

$$L_{H_1}(R) = \sum_{i=1}^n x_i \log(1 - \delta D_{N-1}) + (1 - x_i) \log(\delta D_{N-1}), \quad (2)$$

where D_{N-1} represents the probability of no individual except for the queried person having the queried SNP. δ represents a possible sequencing error. Finally, the LRT statistic is stated as follows:

$$\Lambda = nB + C \sum_{i=1}^n x_i, \quad (3)$$

where B and C are defined as $B = \log(D_N / \delta D_{N-1})$ and $C = \log(\delta D_{N-1} (1 - D_N / \delta D_{N-1}))$, respectively. The null hypothesis is rejected for any Λ that is less than a certain threshold.

The Optimal attack introduced by Raisaro et al. integrates publicly available MAF information into the attacker's background knowledge (Raisaro et al., 2016). In this attack, the victim's SNPs are sorted with respect to their MAFs. The beacon is queried starting from the first heterozygous SNP with the lowest MAF. The model of this attack is illustrated in Figure 2b. In this setting, the

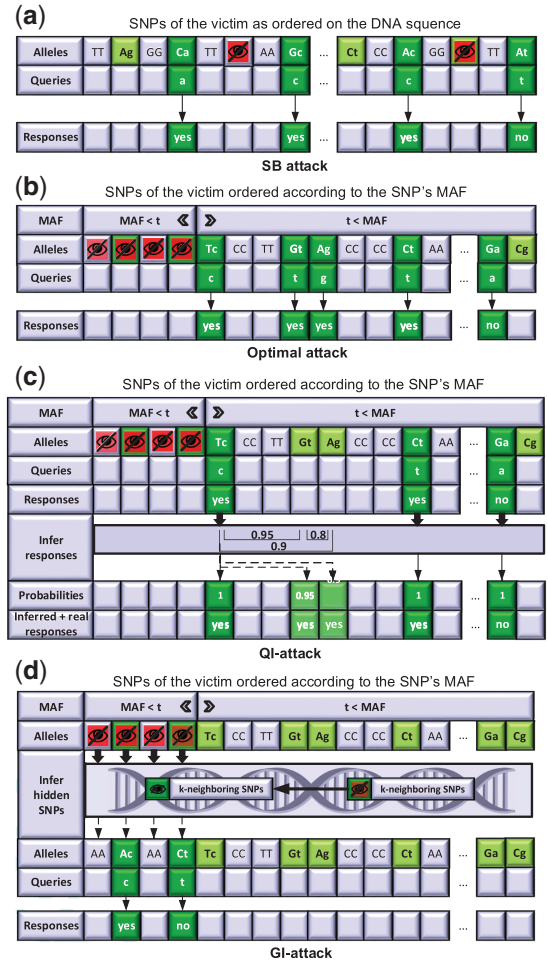


Fig. 2. System models of the four attacker models (a) SB attack (Shringarpure and Bustamante, 2015), (b) optimal attack (Raisaro et al., 2016), (c) QI-attack and (d) GI-attack. Upper-case letters represent the major allele at a SNP position and the lower-case letters the corresponding minor allele. The SB attack randomly selects the minor allele from heterozygous SNP positions of the victim and queries those. The Optimal attack first sorts the heterozygous SNPs regarding their MAFs and queries for the minor alleles starting with the lowest frequency. Depending on the threshold t , SNPs with MAF $< t$ are hidden and are not available to the attacker. The QI-attack extends the optimal attack by inferring beacon answers using LD correlations between SNP pairs. The GI-attack infers the hidden SNPs with MAFs $< t$, using a high-order Markov chain and queries the beacon for the minor alleles of those positions

computations of D_{N-1} and D_N depend on the queried position i and change at each query as shown as follows:

$$D_{N-1}^i = (1 - f_i)^{2N-2}, \quad (4)$$

$$D_N^i = (1 - f_i)^{2N}, \quad (5)$$

where f_i represents the MAF of the SNP at position i . Accordingly, Λ changes as follows:

$$\Lambda = \sum_{i=1}^n \log\left(\frac{D_N^i}{\delta D_{N-1}^i}\right) + \log\left(\frac{\delta D_{N-1}^i (1 - D_N^i)}{D_N^i (1 - \delta D_{N-1}^i)}\right) x_i. \quad (6)$$

2.2 Query inference attack

The QI-attack uses pairwise SNP correlations (LD) in order to infer the answers of unasked queries from previously answered queries.

In this model, the attacker uses the LD value of a SNP pair to calculate the correlation of two minor alleles at the corresponding loci. The correlation is equal to the probability of the two minor alleles occurring together. Let p_2 be the MAF of SNP A (with minor allele a) and q_2 be the MAF of SNP B (with minor allele b). Assuming A and B are in LD, the probability of two major or two minor alleles in these loci occurring together increases. This can be calculated as $Pr(ab) = p_2q_2 + D$, where D represents the strength of the correlation of the two SNPs (see [Supplementary Material, Part A](#) for details). On this basis, the attacker constructs a SNP network that uses weighted, directed edges between SNPs in high LD (see [Supplementary Fig. SA.1](#)). The weight corresponds to the probability of two minor alleles occurring together. [Figure 2c](#) illustrates this model. First, the attacker selects the SNPs to be queried. This step is identical to the optimal attack and leads to a set of candidate SNPs S to be queried, starting from the lowest MAF SNP_i . Second, if any non-queried SNP_j in S is a neighbor of SNP_i in the SNP network, the attacker infers the result of the query and does not pose a query for SNP_j . In the following, we present the null and the alternative hypotheses in this model which also integrates the inference error.

$$L_{H_0}(R) = \sum_{i=1}^n \left(x_i \log(1 - D_N^i) + (1 - x_i) \log(D_N^i) \right) + \sum_{j=1}^m \gamma x_i \log(1 - D_N^j) + \gamma(1 - x_i) \log(D_N^j) \quad (7)$$

$$L_{H_1}(R) = \sum_{i=1}^n \left(x_i \log(1 - \delta D_{N-1}^i) + (1 - x_i) \log(\delta D_{N-1}^i) \right) + \sum_{j=1}^m \gamma x_i \log(1 - \delta D_{N-1}^j) + \gamma(1 - x_i) \log(\delta D_{N-1}^j) \quad (8)$$

where n is the number of posed queries, m is the number of neighbors that can be inferred for each posed query x_i , and γ corresponds to the confidence of the inferred answer, obtained from the SNP network. Λ is then determined as

$$\Lambda = \sum_{i=1}^n \left(\log\left(\frac{D_N^i}{\delta D_{N-1}^i}\right) + \log\left(\frac{\delta D_{N-1}^i(1 - D_N^i)}{D_N^i(1 - \delta D_{N-1}^i)}\right) x_i \right) + \sum_{j=1}^m \log\left(\frac{D_N^j}{\delta D_{N-1}^j}\right) + \log\left(\frac{\delta D_{N-1}^j(1 - D_N^j)}{D_N^j(1 - \delta D_{N-1}^j)}\right) \gamma x_i. \quad (9)$$

By not querying the beacon for answers that can be inferred with high confidence, this model requires less number of queries compared to the optimal attack, while achieving the same response set. For more detail, see [Supplementary Material, Part B](#).

2.3 Genome inference attack

Individuals may publicly share their genomes by taking necessary precautions, such as hiding their sensitive SNP positions with MAFs $< t$ (Scenario 2 in [Fig. 1](#)). The GI-attack performs allele inference to recover hidden SNP positions and infers alleles at the victim's hidden loci. Note that Scenario 1 ([Fig. 1](#)) is not applicable to the GI-attack, since in that scenario, the attacker can access SNPs with low MAFs. The attacker uses a high-order Markov chain to model SNP correlations as described by Samani *et al.*

The model of this attack is illustrated in [Figure 2d](#). Depending on the threshold t , the attacker infers SNP positions with $MAF < t$ that are not available in the victim's VCF file. Based on the victim's genome sequence, the attacker calculates the likelihood of the victim

having a heterozygous position at the chosen SNP position i as follows

$$P_k(SNP_i) = P(SNP_i | SNP_{i-1}, SNP_{i-2}, \dots, SNP_{i-k}), \quad (10)$$

where k is the order of the Markov chain. In order to use a high-order Markov chain to infer hidden SNPs, genome sequences from public sources such as the 1000 Genomes project or HapMap can be used to train the model. Such publicly available genome datasets are typically available with the population information about its anonymized participants. In such a case, we use a dataset that is consistent with the victim's population to build our high-order model. If the population information is not available in a dataset, it can be extracted by using ancestry inference techniques. Accordingly, Samani *et al.* define the k th-order model as

$$P_k(SNP_i) = \begin{cases} 0 & \text{if } F(SNP_{i-k,i-1}) = 0 \\ \frac{F(SNP_{i-k,i})}{F(SNP_{i-k,i-1})} & \text{if } F(SNP_{i-k,i-1}) > 0, \end{cases} \quad (11)$$

where $F(SNP_{i,j})$ is the frequency of occurrence of the sequence that contains SNP_i to SNP_j . The SNPs are ordered according to their physical position on the genome. The model works by comparing the SNPs in $SNP_{i,j}$ which are prior to SNP_i on the genome sequence to the same SNP positions in the training dataset. If the training set contains other genomes with the same SNP sequence and these sequences are followed by a heterozygous position, we can calculate the probability of SNP_i being heterozygous for our victim. As an example, the victim's 4th-order SNP sequence is [AA, AT, CC, TT]. We would now like to determine whether the following SNP_i , that is hidden in the VCF file at hand, is likely to be a heterozygous position. Therefore, we identify other genomes in the training dataset with the same sequence and compute the frequency of this sequence being followed by a heterozygous position. That is, [AA, AT, CC, TT] \rightarrow [AG]. As a result, we can determine the probability of the four SNPs being followed by a heterozygous position, which we can use to query the beacon.

If the calculated likelihood of the victim having a heterozygous position is high enough (in this case equal to 1), the attacker queries the beacon for the inferred SNP position, starting from the SNP with the lowest MAF.

3 Results

To evaluate our attacks, we tested our methods on (i) a simulated beacon and compared our results with the SB attack ([Shringarpure and Bustamante, 2015](#)) and the optimal attack ([Raisaro *et al.*, 2016](#)) (Section 3.1), and (ii) the beacons of the beacon-network (<http://www.beacon-network.org>) operated by GA4GH Beacon-Network and compared our results with the optimal attack ([Raisaro *et al.*, 2016](#)) (Section 3.2).

3.1 Re-identification on a simulated beacon

In this section, we evaluated the performance of the four attacks on a simulated beacon with 65 people from the CEU population of the HapMap dataset. While testing for the alternative hypothesis, we used 20 randomly-picked people from the beacon. For the null hypothesis, we used 40 additional people from the same population of the HapMap project. The CEU population is the population of choice because previous works [SB attack ([Shringarpure and Bustamante, 2015](#)) and optimal attack ([Raisaro *et al.*, 2016](#))] have also been evaluated on this population. The LD scores, allele frequencies and genotype data were also obtained from the CEU

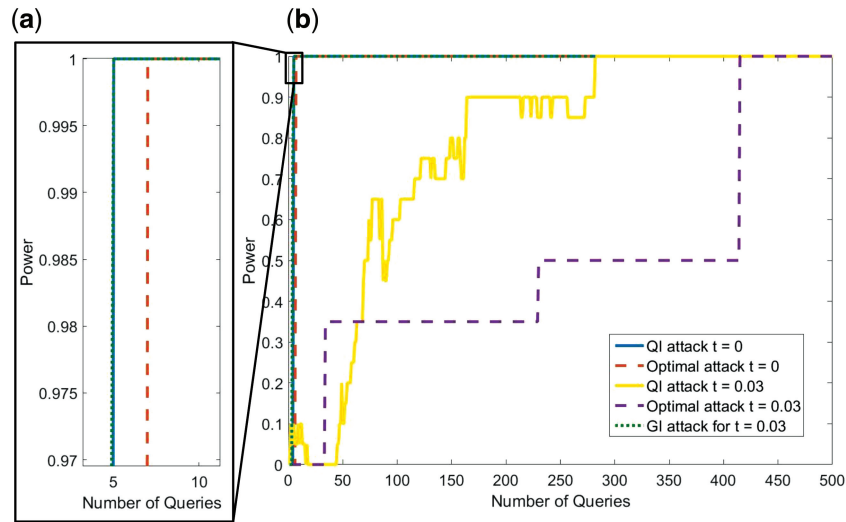


Fig. 3. (a) Close-up of the power curves, where number of queries < 10 . (b) Power curves of the optimal attack (Raisaro et al., 2016), the QI-attack, and the GI-attack for different thresholds of t on a beacon with 65 members constructed with individuals from the CEU dataset of the HapMap project. t indicates the threshold up to which SNPs with an $\text{MAF} < t$ are hidden as a countermeasure

dataset of the HapMap project (Gibbs et al., 2003). For the GI-attack, we used a 4th-order Markov chain (see Supplementary Material, Part C for details of selecting the order).

We show the power curves for the optimal, the QI-attack and the GI-attack each at 5% false positive rate in Figure 3 and the number of queries needed to receive the first negative response in Table 1. We empirically build the null hypothesis. That is, we determine the distribution of Λ under the null hypothesis using the 40 people who are not in the beacon. When Λ is less than a threshold, the null hypothesis is rejected. Similar to Raisaro et al., we reject the null hypothesis when $\Lambda < t_x$. We find the threshold t_x from the null hypothesis with $\alpha = 0.05$ (corresponding to 5% false positive rate). The power $1 - \beta$ is then the proportion of the individuals in the control set having a Λ value, where $\Lambda < t_x$. See Supplementary Material, Part D for more information on the power calculation.

We observed that the SB attack requires the highest number of queries (1400–56 800). The QI-attack requires 30% less number of queries on average compared to the optimal attack. The GI-attack requires only five queries for all tested thresholds of t .

Compared to the monotonically increasing behavior of the power curves for the optimal attack, the power curve for the QI-attack shows a zig-zag behavior. This is because t_x is recalculated at each posed query and the Λ values change based on the number of inferred queries.

The threshold t of hidden SNPs significantly affects the performance of the attacks. As t increases, more common SNPs are available to the attacker which means that the likelihood of another individual in the beacon having the same allele increases. When the beacon was queried for each of the 40 people who are not in the beacon, the SB attack was not able to receive a ‘no’ response with 100 000 queries, (i) for four people when SNPs with an $\text{MAF} < 0.04$ were hidden and (ii) for 12 people when SNPs with an $\text{MAF} < 0.05$ were hidden. Therefore, it was not possible to correctly determine beacon membership for all test individuals to reach 100% power for larger t values. Compared to the GI-attack, the optimal and the QI-attack required a significantly higher amount of queries to determine beacon membership and reach 100% power. The GI-attack successfully determined the correct status for all 40 individuals despite the high threshold of t with only a few queries.

Table 1. Average number of queries needed to receive the first negative response for the SB attack (Shringarpure and Bustamante, 2015), the optimal attack (Raisaro et al., 2016), the QI-attack and the GI-attack for different thresholds of t on a beacon with 65 members constructed with 40 case individuals from the CEU dataset of the HapMap project

t	# of queries			
	SB attack	Optimal attack	QI-attack	GI-attack
0	1 418	3	3	NA
0.03	19 525	270	160	2
0.05	56 759	1495	1031	2

Note: t indicates the threshold up to which SNPs with an $\text{MAF} < t$ are hidden. As the GI-attack concentrates on inferring hidden parts of the genome, we do not consider $t = 0$ (nothing is hidden) for the GI-attack.

3.2 Re-identification on existing beacons

We tested our methods on the beacons of the beacon-network. We selected an individual from the Personal Genomes Project (PGP) (Person’s id: PGP180/hu2D53F2) (Church, 2005) as the victim. To determine if this person is a member of the beacons, we applied the SB attack as ground truth as detailed in Supplementary Material, Part E. For the QI-attack, we used the same SNP network as for the simulated beacon in Section 3.1 (based on the CEU population of HapMap). The Markov chain of the GI-attack was trained on the CEU population of the HapMap (Gibbs et al., 2003) dataset. We again used a 4th-order Markov chain.

The beacons can return an empty response, that is, the beacon has no information at that position, a ‘no’-response, and a ‘yes’-response. We consider two cases for the evaluation of the query results. In the first case, an empty answer is treated as a ‘no’ (results shown in Table 2), in the second case an empty answer is not treated as a ‘no’, as it is also possible that the beacon has a different copy of the victim’s genome (results shown in Supplementary Table SF.1 in Supplementary Material, Part F). As the results are similar, we concentrate on the first case.

Unlike all other beacons, the 1000 Genome Project beacon required fewer number of queries for re-identification as t is increased. Note that the victim’s SNPs are sorted based on the CEU population’s allele frequencies. Thus, SNPs that we query are not necessarily the

Table 2. Number of queries required to receive a ‘no’ within 1000 queries to existing beacons using an individual from PGP (Church, 2005) when $t = \{0, 0.03, 0.05\}$ for the optimal attack (Raisaro et al., 2016), the QI-attack and the GI-attack

Beacon name	$t =$	Optimal attack 0 0.03 0.05	QI-attack 0 0.03 0.05	GI-attack 0.03 0.05
Known VARIants		— — —	— — —	— —
Broad institute		2 2 2	2 2 2	1 1
1000 Genomes project		4 3 2	4 3 2	3 3
Cafe CardioKit		— — —	— — —	— —
Wellcome trust		—	—	—
Sanger institute		1 1 1	1 1 1	1 1
NCBI		— — —	— — —	— —
ICGC		1 — —	1 — —	1 1
AMPLab		20 45 73	20 40 73	39 39
1000 Genomes				
Project phase 3		20 130 250	20 116 250	48 48

Note: Here, empty answers (i.e. the beacon has no information about the queried locus in the underlying dataset and returns neither a ‘no’ nor a ‘yes’) are not considered as a ‘no’ response. ‘—’ means in no ‘no’ was found in 1000 queries.

rarest in the queried beacon, which can explain this behavior. Furthermore, the SNP network used is also based on the CEU population and therefore, does not include all SNPs of the victim’s genome.

The GI-attack performed as expected, that is constant over the two tested thresholds of t and outperformed the optimal attack (Raisaro et al., 2016) as well as the QI-attack for $t > 0$. For the 1000 Genomes Beacon the GI-attack required the same amount of queries as the other attacks, as the number of queries needed is already very low.

In summary, for six of the nine tested beacons, we were able to determine that the victim is not a member of the beacons. For the Known VARIants (Kaviar), the Cafe CardioKit, and the NCBI, it was not possible within 1000 queries (Fig. 5). Overall, we observed that the experiments on real beacon support our findings in Section 3.1. That is, the optimal and the QI-attack need more queries as t increases, the GI-attack is stable over all thresholds, and the QI-attack requires less queries than the optimal attack.

4 Discussion

Recent works by Shringarpure and Bustamante and Raisaro et al. have shown, that beacon servers fail at protecting their members’ privacy. As beacons are often associated with a certain phenotype, the membership identification of an individual could leak sensitive information. They proposed countermeasures such as (i) user budget, (ii) adding noise and (iii) increasing beacon size to improve the security level of existing beacons.

In this work, we have shown that beacon membership can be detected with even a lower number of queries and with high confidence, despite strict countermeasures. Overcoming the proposed countermeasures is possible by including publicly available information such as MAF, LD and VCF files [from e.g. HapMap (Gibbs et al., 2003) or 1000 Genomes Project (Siva, 2008)] into the attacker model. Previous works in the field of genomics and privacy have shown that it is possible to increase the success rate of genomic re-identification attacks by including LD information into the attacker model. Namely, Wang et al. showed that individuals can be re-identified by using (i) publicly available SNP-to-disease correlation information, and (ii) SNPs in LD. Humbert et al. showed how LD can be used to build a framework to reconstruct the genomes of people using the genome of a family member.

The success of our QI-attack depends significantly on the structure of the underlying SNP network. The larger and denser the network becomes, the more query responses can be inferred. Additionally, the strength of the SNP correlations is an important factor. In this work, we included SNP pairs that are in strong LD (i.e. $r^2 > 0.7$) in our SNP network to limit inference error.

The GI-attack shows that even if genomes do not contain any SNPs with low MAFs, individuals’ privacy is not ensured, as it is possible to infer these loci using information from publicly available datasets [e.g. HapMap (Gibbs et al., 2003) or 1000 Genomes Project (Siva, 2008)]. Additionally, the GI-attack still performs as good even when the attacker trains the high-order Markov chain on a different population than the victim’s.

Our experiments on a simulated beacon (Section 3.1) and existing beacons (Section 3.2) show that as the threshold up to which SNPs of the victim with an $MAF < t$ are hidden (t) increases, our attacks require fewer queries than existing attacks [SB attack (Shringarpure and Bustamante, 2015) and optimal attack (Raisaro et al., 2016)]. Table 2 shows that for the existing beacons the number of queries needed increases as t increases and that the margins are even larger compared to the simulated beacon (Table 1).

Several countermeasures against re-identification attacks have been proposed in the literature. Shringarpure and Bustamante discusses the following: (i) increasing the beacon size, (ii) sharing only small genomic regions, (iii) using single population beacons, (iv) not publishing the metadata of a beacon and (v) adding control samples to the beacon dataset (Shringarpure and Bustamante, 2015). Lately, Al Aziz et al. (2017) proposed two algorithms which are based on randomizing the response set of the beacons with the goal of protecting beacon members’ privacy while maintaining the efficacy of the beacon servers.

Raisaro et al. have analyzed the behavior of the beacon when applying three different countermeasures. First, they propose the beacon should only respond ‘yes’ for an allele if multiple samples have it. The second countermeasure adds noise to the responses. However, this countermeasure significantly reduces the utility of the dataset. Instead, the beacon could return an empty answer. Second, they discuss assigning a query budget per sample. That is, every member of the beacon is assigned with a certain budget that is reduced if a query to the beacon matches the sample. As an example, if a user queries the beacon for allele A in position 1000 of chromosome 21, then the budget of every member with an allele A in that position is reduced. The amount of the budget reduction is determined based on the risk of the query, where the lower the allele frequency of the queried allele is, the higher the risk becomes. The budget is calculated as $b_i = \log(p)$, where Raisaro et al. use $P = 0.05$. The risk then is calculated as $r_i = -\log(1 - D_N^i)$. If the budget of a beacon member is depleted, the beacon stops including the member into the beacon responses. We argue that adding noise to beacon answers makes the system useless due to the significant decrease in utility and should not be applied. We show that an attacker using the QI-attack can overcome this countermeasure. For instance, in our simulated beacon as described in Section 3.1, an attacker using the optimal attack needs seven queries to re-identify the victim [individual ‘NA12272’ of the HapMap project (Gibbs et al., 2003)], when no SNPs are hidden. However, the beacon would start giving false responses after six queries as the budget would be depleted, which means the attack would fail. By using the QI-attack, an attacker would only need five queries. Therefore, a query budget that is merely based on the SNPs’ MAFs and that does not consider SNP correlations would fail to protect an individual’s privacy. An attacker using the QI-attack would not exhaust the budget but still be able to determine the victim’s beacon membership.




Fig. 4. Power curves for the QI-attack for varying beacon sizes ($t=0$). All beacons contain only CEU individuals and only chromosome 4 is used for inference

Using the QI-attack, we tested how the size and the diversity beacon affect the privacy breach. First, we repeated our power analysis on the CEU population, while varying the size of the beacon as 45, 65, 85 and 105. Figure 4 shows that increasing beacon size also increases the number of queries needed to achieve 100% power (5% FDR).

To see the effect of diversity of the beacon on the privacy breach, we created new simulated beacons of different populations. That is, we first selected 65 individuals from the Mexican (MEX) population and 65 individuals from the Yoruba Nigerian (YRI) population. Then, we added these separately on top of the simulated CEU beacon, for which the results were reported in Figure 3, to obtain (CEU+MEX), (CEU+YRI) and (CEU+MEX+YRI) beacons. Figure 5 shows that adding YRI population into the CEU beacon reduces the power of the attack, while adding MEX population does not affect the number of required queries to reach to 100% power (FDR = 5%). Comparing (CEU+MEX) and (CEU+MEX+YRI) beacons shows that the number of required queries is eight times more when three populations are mixed (40). Comparing (CEU+YRI) and (CEU+MEX+YRI) beacons shows that the number of required queries is slightly less for the three-way mixture, which indicates that YRI population contains different variants than MEX and CEU.

Among the countermeasures mentioned above, increasing the size and diversity of the beacon are shown to be effective in increasing privacy while fully preserving the utility. However, despite the increase in the number of required queries, the attacks are still applicable. Budget countermeasure can be effective, but again, we show that attack models proposed here can get around the budget. Also, the utility decreases significantly when many individuals in the beacon are removed due to budget depletion. One possible countermeasure could be assigning budgets to users rather than beacon participants. This requires having users sign up for the beacon with institutional accounts and agree to the terms. This would let the data owner monitor and restrict user activity without removing people from beacons' answers, and hence without decreasing utility.

5 Conclusion

Throughout the course of this work, we showed that data-sharing beacons are sensitive to re-identification attacks. Additionally, we showed that countermeasures that do not consider the MAFs and correlations of SNPs fail to protect the beacon members' privacy. Furthermore, even if individuals apply countermeasures before releasing their genome, such as systematically hiding SNPs with low MAFs, their privacy still could be at stake. Therefore, new countermeasures are needed to ensure privacy of individuals.




Fig. 5. Power curves of the QI-attack for (CEU+MEX), (CEU+YRI) and (CEU+MEX+YRI) beacons ($t=0$). Each population has 65 individuals in each beacon, so the beacons contain 130, 130 and 195 individuals, respectively. Only chromosome 4 is used for the experiment

Funding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 707135.

Conflict of Interest: none declared.

References

- Al Aziz,M.M. *et al.* (2017) Aftermath of bustamante attack on genomic beacon service. *BMC Med. Genomics*, **10**, 43.
- Church,G.M. (2005) The personal genome project. *Mol. Syst. Biol.*, **1**, E1.
- Clayton,D. (2010) On inferring presence of an individual in a mixture: a Bayesian approach. *Biostatistics*, **11**, 661–673.
- Collins,F.S. and Varmus,H. (2015) A new initiative on precision medicine. *New Engl. J. Med.*, **372**, 793–795.
- Gibbs,R.A. *et al.* (2003) The international hapmap project. *Nature*, **426**, 789–796.
- Homer,N. *et al.* (2008) Resolving individuals contributing trace amounts of dna to highly complex mixtures using high-density snp genotyping microarrays. *PLoS Genet.*, **4**, e1000167.
- Humbert,M. *et al.* (2013) Addressing the concerns of the lacks family: quantification of kin genomic privacy. In: *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, pp. 1141–1152. ACM, New York, NY, USA.
- Jacobs,K.B. *et al.* (2009) A new statistic and its power to infer membership in a genome-wide association study using genotype frequencies. *Nature Genet.*, **41**, 1253–1257.
- Ledford,H. (2016) Astrazeneca launches project to sequence 2 million genomes. *Nature*, **532**, 427.
- Raisaro,J.L. *et al.* (2016) Addressing beacon re-identification attacks: quantification and mitigation of privacy risks. *J. Am. Med. Inform. Assoc.*, **1**, 1–1.
- Samani,S.S. *et al.* (2015) Quantifying genomic privacy via inference attack with high-order SNV correlations. In: *2015 IEEE Security and Privacy Workshops (SPW)*, pp. 32–40. IEEE.
- Sankararaman,S. *et al.* (2009) Genomic privacy and limits of individual detection in a pool. *Nature Genet.*, **41**, 965–967.
- Shringarpure,S.S. and Bustamante,C.D. (2015) Privacy risks from genomic data-sharing beacons. *Am. J. Hum. Genet.*, **97**, 631–646.
- Siva,N. (2008) 1000 genomes project. *Nature Biotechnol.*, **26**, 256.
- Visscher,P.M. and Hill,W.G. (2009) The limits of individual identification from sample allele frequencies: theory and statistical analysis. *PLoS Genet.*, **5**, e1000628.
- Wang,R. *et al.* (2009) Learning your identity and disease from research papers: information leaks in genome wide association study. In: *Proceedings of the 16th ACM Conference on Computer and Communications Security*, pp. 534–544. ACM.