THEORY OF COMPUTING, Volume 16 (12), 2020, pp. 1–18
www.theoryofcomputing.org

Optimality of Correlated Sampling Strategies

Mohammad Bavarian* Badih Ghazi† Elad Haramaty‡
Pritish Kamath§ Ronald L. Rivest¶ Madhu Sudan‖

Received January 15, 2018; Revised May 30, 2020; Published November 9, 2020

Abstract. In the *correlated sampling* problem, two players are given probability distributions P and Q, respectively, over the same finite set, with access to shared randomness. Without any communication, the two players are each required to output an element sampled according to their respective distributions, while trying to minimize the probability that their outputs disagree. A well known strategy due to Kleinberg–Tardos and Holenstein, with a close variant (for a similar problem) due to Broder, solves this task with disagreement probability at most $2\delta/(1+\delta)$, where δ is the total variation distance between P and Q. This strategy has been used in several different contexts, including sketching algorithms, approximation algorithms based on rounding linear programming relaxations, the study of parallel repetition and cryptography.

In this paper, we give a surprisingly simple proof that this strategy is essentially optimal. Specifically, for every $\delta \in (0,1)$, we show that any correlated sampling strategy incurs

*Work done while at MIT. Supported in part by NSF Award CCF-1420692.

†Work done while at MIT. Supported in part by NSF CCF-1420956, NSF CCF-1420692 and CCF-1217423.

‡Work done while at Harvard. Supported by NSF Award CCF-1565641.

§Work done while at MIT. Supported in part by NSF CCF-1420956 and NSF CCF-1420692.

¶This work was supported by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

‖This work was supported by NSF Awards CCF-1565641 and CCF-1715187, and a Simons Investigator Award.

ACM Classification: F.0, G.3

AMS Classification: 68Q99, 94A20, 68W15

Key words and phrases: distributions, sampling, correlated sampling, coupling, MinHash, communication complexity

 DOI: 10.4086/toc.2020.v016a012

a disagreement probability of essentially $2\delta/(1+\delta)$ on some inputs P and Q with total variation distance at most δ. This partially answers a recent question of Rivest.

Our proof is based on studying a new problem that we call *constrained agreement*. Here, the two players are given subsets $A \subseteq [n]$ and $B \subseteq [n]$, respectively, and their goal is to output an element $i \in A$ and $j \in B$, respectively, while minimizing the probability that $i \neq j$. We prove tight bounds for this question, which in turn imply tight bounds for correlated sampling. Though we settle basic questions about the two problems, our formulation leads to more fine-grained questions that remain open.

1 Introduction

In this paper, we study *correlated sampling*, a basic task, variants of which have been considered in the context of sketching algorithms [2], approximation algorithms based on rounding linear programming relaxations [7, 3], the study of parallel repetition [6, 12, 1] and cryptography [13].

The *correlated sampling problem* is defined as follows. Alice and Bob are given probability distributions P and Q, respectively, over a finite set Ω. Without any communication, using only shared randomness as the means to coordinate, Alice is required to output an element a distributed according to P and Bob is required to output an element b distributed according to Q. Their goal is to minimize the disagreement probability $\Pr[a \neq b]$, which we will compare with the total variation distance between P and Q, defined as

$$d_{\mathrm{TV}}(P,Q) \;:=\; \sup_{A \subseteq \Omega} P(A) - Q(A) \;=\; \frac{1}{2} \sum_{\omega \in \Omega} |P(\omega) - Q(\omega)|. \tag{1.1}$$

A *correlated sampling strategy* is formally defined below, where Δ_Ω denotes the set of all probability distributions over Ω and $(\mathcal{R}, \mathcal{F}, \mu)$ denotes the probability space corresponding to the randomness shared by Alice and Bob. $\mathcal{R}$ is the sample space, $\mathcal{F}$ is a σ-algebra over $\mathcal{R}$ and μ is a probability measure over $(\mathcal{R}, \mathcal{F})$. Even though Ω is finite, we allow $\mathcal{R}$ to be infinite. For simplicitly, we abuse notation and use $\mathcal{R}$ to denote both the sample space and the probability space.

Definition 1.1. A *correlated sampling strategy* for a finite set Ω with error $\varepsilon : [0,1] \to [0,1]$ is specified by a probability space $\mathcal{R}$ and a pair of functions[1] $f, g : \Delta_\Omega \times \mathcal{R} \to \Omega$, such that for all $P, Q \in \Delta_\Omega$ with $d_{\mathrm{TV}}(P,Q) \leq \delta$, the following hold.

$$\begin{array}{ll} \textbf{[Correctness]} & \{f(P,r)\}_{r\sim\mathcal{R}} = P \text{ and } \{g(Q,r)\}_{r\sim\mathcal{R}} = Q, \\ \textbf{[Error Guarantee]} & \Pr_{r\sim\mathcal{R}}[f(P,r) \neq g(Q,r)] \;\leq\; \varepsilon(\delta). \end{array}$$

In the above, $\{f(P,r)\}_{r\sim\mathcal{R}}$ denotes the push-forward measure, that is, the distribution of the random variable $f(P,r)$ over Ω, where $r \sim \mathcal{R}$ is the source of shared randomness. For simplicity, we will often not mention $\mathcal{R}$ explicitly when talking about correlated sampling strategies. While we defined correlated sampling strategies for finite sets only, it is possible to define it for infinite sets Ω; see Section 5 for a discussion. In this paper we consider finite sets Ω only, except where otherwise stated.

[1] We require both functions to be measurable in their second argument.

A correlated sampling strategy is notably different from the fundamental notion of a *coupling* (see, e. g., [14] for an introduction), where we require a *single* coupling function $h : \Delta_\Omega \times \Delta_\Omega \to \Delta_{\Omega\times\Omega}$ such that for any distributions P and Q it holds that the marginals of $h(P,Q)$ are P and Q respectively. In other words, a coupling function has the knowledge of both P and Q, whereas a correlated sampling strategy operates locally on the knowledge of P and on the knowledge of Q. It is well known that for any coupling function h, it holds that $\Pr_{(a,b)\sim h(P,Q)}[a \neq b] \geq d_{\mathrm{TV}}(P,Q)$ and that this bound is achievable. Observe that a correlated sampling strategy induces a coupling given as $\{(f(P,r), g(Q,r))\}_{r\sim\mathcal{R}}$. Thus, it follows that $\varepsilon(\delta) \geq \delta$. And yet a priori, it is unclear whether any non-trivial correlated sampling strategy can even exist, since the error ε is not allowed to depend on the size of Ω.

Somewhat surprisingly, there exists a simple strategy whose error can be bounded by roughly twice the total variation distance (and in particular does not degrade with the size of Ω). Variants of this strategy have been rediscovered multiple times in the literature, yielding the following theorem.

Theorem 1.2 ([2, 7, 6])**.** *For all $n \in \mathbb{Z}_{\geq 0}$, there exists a correlated sampling strategy for sets of size n, with error $\varepsilon : [0,1] \to [0,1]$, such that for all $\delta \in [0,1]$, it holds that*

$$\varepsilon(\delta) \;\leq\; \frac{2\cdot\delta}{1+\delta}. \tag{1.2}$$

Strictly speaking, Broder's paper [2] did not consider the general correlated sampling problem. Rather it introduced the *MinHash strategy*, which can be interpreted as a correlated sampling strategy for the special case where P and Q are *flat* distributions, i. e., they are uniform over some subsets of Ω. In particular, if $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$ are distributions that are uniform over sets $A, B \subseteq \Omega$, respectively, then the MinHash strategy gives an error probability of $1 - |A\cap B|/|A\cup B|$, also known as the *Jaccard distance* between A and B. In the special case when $|A| = |B|$, this is equivalent to the bound above.

The technique can be generalized to other (non-flat) distributions to get the bound in Theorem 1.2, thereby yielding a strategy due to Kleinberg–Tardos and Holenstein.[2] Several variants of this (sometimes referred to as "consistent sampling" protocols) have been used in applied work, including [10, 4, 9, 5].

Given Theorem 1.2, a natural and basic question is whether the bound on the error can be improved; the only lower bound we are aware of is the one that arises from coupling, namely $\varepsilon(\delta) \geq \delta$. This question was recently raised by Rivest [13] in the context of a new encryption scheme and was one of the motivations for this work. We give a surprisingly simple proof that the bound in Theorem 1.2 is essentially tight.

Theorem 1.3 (Main Result)**.** *For all $\delta, \gamma \in (0,1)$, for all sufficiently large n, any correlated sampling strategy for a set of size n with error $\varepsilon : [0,1] \to [0,1]$ satisfies*

$$\varepsilon(\delta) \;\geq\; \frac{2\cdot\delta}{1+\delta} - \gamma. \tag{1.3}$$

Organization of the paper. In Section 2, we prove Theorem 1.3. In Section 3, we consider the setting where Ω is of a fixed finite size, which was the question originally posed by Rivest [13]. In this regime, there turns out to be a surprising strategy that gets better error than Theorem 1.2 in a very special case.

[2] Strictly speaking, if P and Q are flat over subsets of different sizes, the above bound is weaker than that obtained from a direct application of the MinHash strategy. See Section 5 for a discussion.

However, it was conjectured in [13] that in fact a statement like Theorem 1.3 holds in every other case and we make progress on this conjecture by proving it in one such case. For completeness, in Section 4 we describe the correlated sampling strategies of Broder, Kleinberg–Tardos, and Holenstein, underlying Theorem 1.2. We conclude with some more observations and open questions in Section 5.

2 Lower bound on correlated sampling

In order to prove Theorem 1.3, we first introduce the *constrained agreement* problem, a relaxation of the correlated sampling problem. In this problem, Alice and Bob are given sets $A \subseteq \Omega$ and $B \subseteq \Omega$, respectively, where the pair (A,B) is sampled from some (known) distribution $\mathcal{D}$. Alice and Bob are required to output elements $a \in A$ and $b \in B$, respectively, such that the disagreement probability $\Pr_{(A,B)\sim D}[a \neq b]$ is minimized.

This can be viewed as a relaxation of the correlated sampling problem by first considering the case of *flat* distributions in Theorem 1.1 and relaxing the restrictions of $\{f(P,r)\}_{r\sim\mathcal{R}} = P$ and $\{g(Q,r)\}_{r\sim\mathcal{R}} = Q$ to only requiring that $f(P,r) \in \operatorname{supp}(P)$ and $g(Q,r) \in \operatorname{supp}(Q)$ for all $r \in \mathcal{R}$. This makes it a constraint satisfaction problem and we consider a distributional version of the same.

In the following definition, we use 2^Ω to denote the powerset of Ω.

Definition 2.1. A *constrained agreement strategy* for a finite set Ω and a probability distribution $\mathcal{D}$ over $2^\Omega \times 2^\Omega$ is specified by a pair of functions $f,g : 2^\Omega \to \Omega$ and has error $\operatorname{err}_{\mathcal{D}}(f,g)$ if the following hold.

$$\begin{array}{ll} \textbf{[Correctness]} & \forall A,B \subseteq \Omega \ : \ f(A) \in A \text{ and } g(B) \in B, \\ \textbf{[Error guarantee]} & \Pr_{(A,B)\sim\mathcal{D}}[f(A) \neq g(B)] \ =: \ \operatorname{err}_{\mathcal{D}}(f,g). \end{array}$$

Note that since the constrained agreement problem is defined with respect to a (known) probability distribution $\mathcal{D}$ on pairs of sets, we can require, without loss of generality, that the strategies (f,g) be deterministic (since any randomized strategy can be derandomized with no degradation in the error).

In order to prove Theorem 1.3, we characterize the optimal constrained agreement strategy in the special case when $\mathcal{D} = \mathcal{D}_p$ where every element $\omega \in \Omega$ is independently included in each of A and B with probability p.

Lemma 2.2. *For all $p \in [0,1]$, any constrained agreement strategy (f,g) for a finite set Ω and distribution $\mathcal{D} = \mathcal{D}_p$ over $2^\Omega \times 2^\Omega$, has error*

$$\operatorname{err}_{\mathcal{D}_p}(f,g) \geq \frac{2(1-p)}{2-p}.$$

Proof of Lemma 2.2. For ease of notation, let $\Omega = [n]$. Let (f,g) be a constrained agreement strategy. We will construct functions f^* and g^* such that

$$\operatorname{err}_{\mathcal{D}_p}(f,g) \ \geq \ \operatorname{err}_{\mathcal{D}_p}(f^*,g^*) \ \geq \ \frac{2(1-p)}{2-p}.$$

For every $i \in [n]$, let $\beta_i := \Pr_B[g(B) = i]$. Without loss of generality (by suitably permuting $[n]$), we can assume that $\beta_1 \geq \beta_2 \geq \cdots \geq \beta_n$. Since A and B are independently sampled in $\mathcal{D}_p$, it follows that when Bob's strategy is fixed to g, the strategy of Alice that results in the largest agreement probability is simply $f^*(A) := \operatorname{argmax}_{i\in A} \beta_i = \min\{i : i \in A\}$ for all $A \subseteq [n]$.

So far we have $\mathrm{err}_{\mathcal{D}_p}(f,g) \geq \mathrm{err}_{\mathcal{D}_p}(f^*,g)$. We can repeat the same process again. For every $i \in [n]$, define $\alpha_i := \Pr_A[f^*(A) = i]$. Due to the specific choice of f^*, it holds that $\alpha_i = (1-p)^{i-1}p$ and hence $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_n$. Thus, when Alice's strategy is fixed to f^*, the strategy of Bob that results in the largest agreement probability is given by $g^*(B) = \operatorname{argmax}_{i\in B} \alpha_i = \min\{i : i \in B\}$ for all $B \subseteq [n]$. Hence, we get $\mathrm{err}_{\mathcal{D}_p}(f,g) \geq \mathrm{err}_{\mathcal{D}_p}(f^*,g) \geq \mathrm{err}_{\mathcal{D}_p}(f^*,g^*)$ where

$$\begin{aligned}
\mathrm{err}_{\mathcal{D}_p}(f^*,g^*) \;&:=\; 1 - \Pr_{(A,B)\sim\mathcal{D}_p}[f^*(A) = g^*(B)] \\
&=\; 1 - \sum_{i=1}^{n} \Pr_A[f^*(A) = i] \cdot \Pr_B[g^*(B) = i] \\
&=\; 1 - \sum_{i=1}^{n} (1-p)^{2\cdot(i-1)} \cdot p^2 \\
&\geq\; 1 - \frac{p}{2-p} \;=\; \frac{2(1-p)}{2-p}.
\end{aligned}$$

Thus, we conclude that

$$\mathrm{err}_{D_p}(f,g) \;\geq\; \frac{2(1-p)}{2-p}. \qquad \square$$

Before turning to the proof of Theorem 1.3, we note a couple of basic facts.

Fact 2.3. *For* flat *distributions $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$ with $A, B \subseteq \Omega$, it holds, that*

$$d_{\mathrm{TV}}(P,Q) \;=\; 1 - \frac{|A \cap B|}{\max\{|A|,|B|\}}.$$

The following concentration bound, due to Sergei Bernstein from the 1920s, is often referred to as Chernoff's or Hoeffding's bound. The Bernoulli random variable $X \sim \mathrm{Ber}(p)$ is a 0-1 random variable with $Pr[X = 1] = p$.

Fact 2.4 (See, e. g., Cor 4.6 in [11])**.** *For $X_1,\dots,X_n$ drawn i. i. d. from* $\mathrm{Ber}(p)$*, it holds for all $\tau > 0$, that*

$$\Pr\left[\left|\textstyle\sum_i X_i - pn\right| \geq \tau \cdot pn\right] \leq 2 \cdot \mathrm{e}^{-pn\tau^2/3}.$$

Proof of Theorem 1.3. Fix $\delta, \gamma \in (0,1)$. Assume, for the sake of contradiction, that for infinitely many values of n, there is a correlated sampling strategy (f^*, g^*) for a set of size n with error

$$\varepsilon(\delta) < \frac{2\cdot\delta}{1+\delta} - \gamma.$$

Let $\delta' \in (0,\delta)$ be such that

$$\frac{2\cdot\delta}{1+\delta} - \gamma < \frac{2\cdot\delta'}{1+\delta'} < \frac{2\cdot\delta}{1+\delta}. \tag{2.1}$$

Consider the distribution $\mathcal{D}_p$ over pairs (A,B) of subsets $A,B \subseteq [n]$ where each $i \in [n]$ is independently included in each of A and B with probability $p := 1-\delta'$. Thus, we have $\mathbb{E}[|A|] = \mathbb{E}[|B|] = p \cdot n$, and $\mathbb{E}[|A \cap B|] = p^2 \cdot n$. Moreover, by Theorem 2.4 with $\tau = n^{-0.2}$, we have that

$$\begin{aligned}
\Pr_A[||A| - pn| > p \cdot n^{0.8}] \;&\leq\; 2 \cdot \mathrm{e}^{-p \cdot n^{0.6}/3}, \\
\Pr_B[||B| - pn| > p \cdot n^{0.8}] \;&\leq\; 2 \cdot \mathrm{e}^{-p \cdot n^{0.6}/3}, \\
\Pr_{A,B}[||A \cap B| - p^2 n| > p^2 \cdot n^{0.8}] \;&\leq\; 2 \cdot \mathrm{e}^{-p^2 \cdot n^{0.6}/3}.
\end{aligned}$$

Hence, by the union bound and using $p^2 \leq p$, we get that with probability at least $1 - 6 \cdot \mathrm{e}^{-p^2 \cdot n^{0.6}/3}$, we have that $||A| - p \cdot n| \leq pn^{0.8}$, $||B| - p \cdot n| \leq pn^{0.8}$ and $||A \cap B| - p^2 \cdot n| \leq p^2 n^{0.8}$. Thus, for the distributions $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$, it holds with probability at least $1 - 6 \cdot \mathrm{e}^{-p^2 \cdot n^{0.6}/3}$ that

$$\begin{aligned}
d_{\mathrm{TV}}(P,Q) &= 1 - \frac{|A \cap B|}{\max\{|A|,|B|\}} \\
&\leq 1 - p + o_n(1) \\
&= \delta' + o_n(1) \\
&< \delta \quad \text{for sufficiently large } n.
\end{aligned}$$

The assumed strategy (f^*, g^*) is such that

$$\Pr_{r \sim \mathcal{R}}[f(P,r) \neq g(Q,r)] \leq \frac{2\delta}{(1+\delta)} - \gamma$$

when $d_{\mathrm{TV}}(P,Q) \leq \delta$ and at most 1 otherwise. In our random choice of the pair of distributions (P,Q), the probability of $d_{\mathrm{TV}}(P,Q) > \delta$ is at most $o_n(1)$. Thus,

$$\Pr_{(P,Q),\, r \sim \mathcal{R}}[f(P,r) \neq g(Q,r)] \leq \frac{2 \cdot \delta}{1+\delta} - \gamma + o_n(1)$$

when applied on the randomly sampled (P,Q). In particular, by averaging, there exists a deterministic constrained agreement strategy with no worse disagreement probability. That is,

$$\exists\, (f,g), \quad \mathrm{err}_{\mathcal{D}_p}(f,g) \;\leq\; \frac{2 \cdot \delta}{1+\delta} - \gamma + o_n(1). \tag{2.2}$$

But from Lemma 2.2 we have that,

$$\forall\, (f,g), \quad \mathrm{err}_{\mathcal{D}_p}(f,g) \;\geq\; \frac{2(1-p)}{2-p} = \frac{2 \cdot \delta'}{1+\delta'}. \tag{2.3}$$

Putting Equations (2.2) and (2.3) together contradicts Equation (2.1) for sufficiently large n. □

3 Correlated sampling over a fixed set of finite size

Theorem 1.3 establishes that the correlated sampling strategy underlying Theorem 1.2 is *nearly* optimal for Ω that is sufficiently large in size. However, it does not say that the strategy underlying Theorem 1.2 is exactly optimal for a fixed set of finite size. The quest for understanding optimality in this setting was motivated by a new encryption scheme proposed by Rivest [13]. But as we will see shortly, this quest is not entirely straightforward!

In order to elaborate on this, it will be useful to formally define restricted versions of the correlated sampling strategy which are required to work only when the input pair (P,Q) is promised to lie in a given relation $\mathcal{G} \subseteq \Delta_\Omega \times \Delta_\Omega$.

Definition 3.1. For a finite set Ω and a relation $\mathcal{G} \subseteq \Delta_\Omega \times \Delta_\Omega$, a $\mathcal{G}$*-restricted correlated sampling strategy* with error ε is specified by a probability space $\mathcal{R}$, a pair of functions $f, g : \Delta_\Omega \times \mathcal{R} \to \Omega$ if the following hold for all pairs of distributions $(P,Q) \in \mathcal{G}$,

$$\begin{array}{ll} \textbf{[Correctness]} & \{f(P,r)\}_{r\sim\mathcal{R}} = P \text{ and } \{g(Q,r)\}_{r\sim\mathcal{R}} = Q, \\ \textbf{[Error Guarantee]} & \Pr_{r\sim\mathcal{R}}[f(P,r) \neq g(Q,r)] \ \leq\ \varepsilon. \end{array}$$

For example, letting $\mathcal{G}$ to be set of all pairs (P,Q) with $d_{\mathrm{TV}}(P,Q) \leq \delta$ essentially recovers the original setting of correlated sampling, for a fixed total variation distance bound between the input distributions. For the rest of this section, we will consider a special kind of $\mathcal{G}$-restriction corresponding to Alice and Bob having *flat distributions* over $\Omega = [n]$.

Definition 3.2. For all n, the relation $\mathcal{G}^n_{a,b,\ell} \subseteq \Delta_{[n]} \times \Delta_{[n]}$ is defined to consist of pairs (P,Q) of flat distributions corresponding to sets $A, B \subseteq [n]$ such that $P = \mathcal{U}(A)$, $Q = \mathcal{U}(B)$ and $|A| = a$, $|B| = b$, $|A \cap B| = \ell$. (For the relation to be non-empty, it is required that $\ell \leq \min\{a,b\}$ and $a+b-\ell \leq n$.)

Recall from Theorem 2.3, that for all $(P,Q) \in \mathcal{G}^n_{a,b,\ell}$ with $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$, is given by

$$d_{\mathrm{TV}}(P,Q) \ = \ 1 - \frac{|A \cap B|}{\max\{|A|,|B|\}} \ = \ 1 - \frac{\ell}{\max\{a,b\}}.$$

Moreover, the MinHash strategy applied on input pairs $(P,Q) \in \mathcal{G}^n_{a,b,\ell}$ has a disgreement probability

$$1 - \frac{|A \cap B|}{|A \cup B|} \ = \ 1 - \frac{\ell}{a+b-\ell}.$$

One might suspect that this is optimal for all values of n, a, b and ℓ. But rather surprisingly, in the very special case where $|A \cap B| = 1$ and $|A \cup B| = n$, Rivest [13] gave a strategy with smaller error probability than the above! While we don't know of any applications for this strategy itself, its purpose here is to illustrate that there can be strategies which do better than the MinHash strategy in some special cases.

Theorem 3.3 ([13])**.** *For all $a, b \in \mathbb{Z}_{\geq 1}$ there exists a $\mathcal{G}^{a+b-1}_{a,b,1}$-restricted correlated sampling strategy with error at most $1 - 1/\max\{a,b\}$.*

For completeness, we describe this strategy in Section 3.1. Note that for $(P,Q) \in \mathcal{G}^{a+b-1}_{a,b,1}$,

$$d_{\mathrm{TV}}(P,Q) \;=\; 1 - \frac{1}{\max\{a,b\}} \;<\; 1 - \frac{1}{a+b-1}.$$

This naturally leads to the question: *Is there a better correlated sampling strategy for larger intersection sizes?* In fact, the MinHash strategy was conjectured to be optimal in every other case (in particular, for all $\ell > 1$) by Rivest [13] as this is sufficient for proving the security of his proposed encryption scheme.

Conjecture 3.4 (Rivest)**.** *For every collection of positive integers $n \geq a, b \geq \ell \geq 2$ such that $n \geq a+b-\ell$, any $\mathcal{G}^n_{a,b,\ell}$-restricted correlated sampling strategy makes error at least $1-\ell/(a+b-\ell)$.*

As partial progress towards this conjecture, we prove that in the other extreme (as compared to Theorem 3.3), the above conjecture does hold. In particular, we show the following theorem in Section 3.2.

Theorem 3.5. *For all $a = b \geq 1$, $\ell = a-1$ and $n \geq a+1$, any $\mathcal{G}^n_{a,b,\ell}$-restricted correlated sampling strategy makes error at least $1-\ell/(a+b-\ell)$.*

3.1 Correlated sampling strategy of Rivest [13]

In order to prove Theorem 3.3, we recall Philip Hall's "Marriage Theorem."

Lemma 3.6 (P. Hall; see, e. g., [8])**.** *Fix a bipartite graph G on vertex sets L and R (with $|L| \leq |R|$). There exists a matching that entirely covers L if and only if for every subset $S \subseteq L$, we have that $|S| \leq |N_G(S)|$, where $N_G(S)$ denotes the set of neighbors in G of vertices in S.*

Proof of Theorem 3.3. First, let us consider the case where $a = b$. Let $\binom{[n]}{a}$ denote the set of all subsets $A \subseteq [n]$ with $|A| = a$. Consider the bipartite graph G on vertices $\binom{[n]}{a} \times \binom{[n]}{a}$, with an edge between vertices A and B if $|A \cap B| = 1$. It is easy to see that G is a-regular (since $n = 2a-1$). Iteratively using Theorem 3.6, we get that the edges of G can be written as a disjoint union of a matchings. Let us denote these as $M_1, M_2, \ldots, M_a$.

The $\mathcal{G}^{2a-1}_{a,a,1}$-restricted correlated sampling strategy of Alice and Bob is as follows: Use the shared randomness to sample a random index $r \in [a]$ and consider the matching M_r. If (A, B') is the edge present in M_r, then Alice outputs the unique element in $A \cap B'$. Similarly, if (A', B) is the edge present in M_r, then Bob outputs the unique element in $A' \cap B$. This strategy is summarized in Algorithm 1.

It is easy to see that both Alice's and Bob's outputs are uniformly distributed in A and B, respectively. Moreover, the probability that they output the same element, is exactly $1/a$, which is the probability of choosing the unique matching M_r which contains the edge (A,B) (i. e., enforcing $A' = A$ and $B' = B$).

The strategy in the general case of $a \neq b$ is obtained by a simple reduction to the case above. Suppose w.l.o.g. that $a > b$. Alice and Bob get sets $A \subseteq [n]$ and $B \subseteq [n]$ such that $|A| = a$, $|B| = b$ and $|A \cap B| = 1$ and $A \cup B = [n]$. We extend the universe by adding $(a-b)$ dummy elements to get a universe of size $(2a-1)$ (note, $n = a+b-1$). Moreover, whenever Bob gets set B, he extends it to B' by adding all the dummy elements to B and thus $|B'| = a$ while having $|A \cap B'| = 1$ and $|A \cup B| = 2a-1$. Now, Alice and Bob can use the $\mathcal{G}^{2a-1}_{a,a,1}$-restricted correlated sampling strategy from above on the input pair (A, B'). This achieves an error of $1 - 1/a = 1 - 1/\max\{a,b\}$. However, Bob's output is uniformly distributed over

Algorithm 1: Rivest's strategy [13]

Alice's input: $A \subseteq [n]$

Bob's input: $B \subseteq [n]$

$\mathcal{G}$-restriction: $|A| = |B| = a$, $|A \cap B| = 1$ and $A \cup B = [n]$, i. e., $n = a + b - 1$

Preprocessing: Let G be the bipartite graph on vertices $\binom{[n]}{a} \times \binom{[n]}{a}$, with an edge between vertices A and B if $|A \cap B| = 1$. Decompose the edges of G into a disjoint matchings $M_1, \ldots, M_a$.

Shared randomness: Index $r \sim \mathcal{U}([a])$

Strategy:

- Let (A, B') and (A', B) be edges present in M_r.
- $f(A, r) :=$ unique element in $A \cap B'$.
- $g(B, r) :=$ unique element in $A' \cap B$.

B' and not B. To fix this, Bob can simply output a uniformly random element of B whenever the above strategy requires him to return an element of $B' \setminus B$. It is easy to see that this doesn't change the error probability. □

3.2 Proof of Theorem 3.5

Proof of Theorem 3.5. Let $A, B \subseteq [n]$ be such that $a = |A| = |B| = |A \cap B| + 1$ and let $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$. For simplicity, we can assume without loss of generality that $A \cup B = [n]$. Thus, $n = a + 1$ and $\ell = a - 1$. Assume for the sake of contradiction that there is a $\mathcal{G}^{a+1}_{a,a,a-1}$-correlated sampling strategy with disagreement probability $< 1 - \ell/(2a - \ell) = 2/n$. Let $\mathcal{D}$ be the uniform distribution over pairs (A, B) of subsets of $[n]$ satisfying $A \cup B = [n]$ and $|A| = |B| = |A \cap B| + 1$. Note that $\mathcal{D}$ is not a product distribution over (A, B), unlike in Lemma 2.2, which is what makes it challenging to analyze. By an averaging argument, there is a deterministic strategy pair (f, g) such that,

$$\Pr_{(A,B)\sim\mathcal{D}}[f(A) \neq g(B)] < \frac{2}{n}. \tag{3.1}$$

Let

$$i \ := \ \operatorname*{argmax}_{\ell \in [n]} \left| \left\{ A \in \tbinom{[n]}{n-1} : f(A) = \ell \right\} \right| \tag{3.2}$$

be the element that is most frequently output by Alice's strategy f, and denote its number of occurences by

$$k \ := \ \left| \left\{ A \in \binom{[n]}{n-1} : f(A) = i \right\} \right|. \tag{3.3}$$

We consider three different cases depending on the value of k:

(i) If $k \le n-3$, then consider a subset $B \subseteq [n]$ with $|B| = n-1$. For any value of $f(B) \in B$, the conditional error probability $\Pr[f(A) \ne g(B) \,|\, B]$ is at least $2/(n-1)$. Averaging over all such B, we get a contradiction to Equation (3.1).

(ii) If $k = n-2$, let $A_1 \ne A_2$ be the two subsets of $[n]$ with $|A_1| = |A_2| = n-1$ such that $f(A_1) \ne i$ and $f(A_2) \ne i$. For all $B \subseteq [n]$ with $|B| = n-1$ such that $B \ne A_1$ and $B \ne A_2$, the conditional error probability $\Pr[f(A) \ne g(B) \,|\, B]$ is at least $2/(n-1)$. Note that there are $n-2$ such sets B, and that either A_1 or A_2 is the set $[n] \setminus \{i\}$. If $B = [n] \setminus \{i\}$, then the conditional disagreement probability $\Pr[f(A) \ne g(B) \,|\, B]$ is at least $(n-2)/(n-1)$. Averaging over all B, we get that

$$\Pr_{(A,B)\sim\mathcal{D}}[f(A) \ne g(B)] \ge \left(\frac{2}{n-1}\right) \cdot \left(\frac{n-2}{n}\right) + \left(\frac{n-2}{n-1}\right) \cdot \left(\frac{1}{n}\right) \ge \frac{2}{n},$$

where the last inequality holds for all $n \ge 2$. This contradicts Equation (3.1).

(iii) If $k = n-1$, then the only subset A_1 of $[n]$ with $|A_1| = n-1$ and such that $f(A_1) \ne i$ is $A_1 = [n] \setminus \{i\}$. For all $B \ne A_1$, the conditional error probability $\Pr[f(A) \ne g(B) \,|\, B]$ is at least $1/(n-1)$. On the other hand, if $B = A_1$, then the conditional error probability is equal to 1. Averaging over all B, we get that

$$\Pr_{(A,B)\sim\mathcal{D}}[f(A) \ne g(B)] \ \ge \ \left(\frac{1}{n-1}\right) \cdot \left(\frac{n-1}{n}\right) + 1 \cdot \left(\frac{1}{n}\right) \ = \ \frac{2}{n},$$

which contradicts Equation (3.1). □

Remark. In [7], the correlated sampling strategy is used to give a randomized rounding procedure for a linear program. The factor 2 loss in the correlated sampling strategy translates into an integrality gap of at most 2. In fact, they also prove that the integrality gap is roughly tight. As pointed out by an anonymous reviewer, their proof essentially establishes Theorem 3.5 under the assumption that $f = g$.

4 Correlated sampling strategies of [2, 7, 6]

For sake of completeness, we describe the correlated sampling strategies of Broder and of Kleinberg–Tardos and Holenstein, thereby proving Theorem 1.2.

Broder's Min Hash Strategy. Consider the case of *flat distributions*, where the distributions P and Q are promised to be of the following form: there exist $A, B \subseteq [n]$ such that $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$. In this case, it is easy to show that the strategy given in Algorithm 2 achieves an error probability of $1 - |A \cap B|/|A \cup B|$. Since π is a random permutation, $f(A, \pi)$ is uniformly distributed over A and $g(B, \pi)$ is uniformly distributed over B. Let i_0 be the smallest index such that $\pi(i_0) \in A \cup B$. The probability that $\pi(i_0) \in A \cap B$ is exactly $|A \cap B|/|A \cup B|$, and this happens precisely when $f(A, \pi) = g(B, \pi)$. Hence, we get the claimed error probability.

The correlated sampling strategy of [7, 6] follows a similar approach.

Algorithm 2: MinHash strategy [2]
Alice's input: $A \subseteq [n]$ **Bob's input:** $B \subseteq [n]$ **Shared randomness:** a random permutation $\pi : [n] \to [n]$ **Strategy:** • $f(A,\pi) = \pi(i_A)$, where i_A is the smallest index such that $\pi(i_A) \in A$. • $g(B,\pi) = \pi(i_B)$, where i_B is the smallest index such that $\pi(i_B) \in B$.

Proof of Theorem 1.2. Given a finite set Ω and probability distributions P and Q over Ω, define

$$A := \{(\omega, p) \in \Omega \times [0,1] : p < P(\omega)\} \quad \text{and} \quad B := \{(\omega, q) \in \Omega \times [0,1] : q < Q(\omega)\}.$$

Also for all $\omega \in \Omega$, define $A_\omega := A \cap (\{\omega\} \times [0,1])$ and $B_\omega := B \cap (\{\omega\} \times [0,1])$.

The strategy of [7, 6] can be intuitively understood as follows. Alice and Bob use the MinHash strategy on inputs A and B over $\Omega \times [0,1]$, to obtain elements (ω_A, p_A) and (ω_B, p_B), respectively, and simply output ω_A and ω_B, respectively. However, this by itself is not well defined since $\Omega \times [0,1]$ is not a finite set. Nevertheless, the MinHash strategy can be modified to instead have a (countably) infinite sequence of points sampled i. i. d. from the uniform distribution over $\Omega \times [0,1]$, instead of a permutation π. This strategy is summarized in Algorithm 3.

Let μ be the uniform distribution over $\Omega \times [0,1]$. Observe that $\mu(A) = \mu(B) = 1/|\Omega|$ and for all $\omega \in \Omega$, we have $\mu(A_\omega) = P(\omega)/|\Omega|$ and $\mu(B_\omega) = Q(\omega)/|\Omega|$. Similar to the analysis of the MinHash strategy, for Alice's chosen index i_A, we have (ω_{i_A}, r_{i_A}) is uniform over A. Thus, $\Pr[f(P,\pi) = \omega]$ is precisely $\mu(A_\omega)/\mu(A) = P(\omega)$. Thus, $f(P,\pi)$ is distributed according to P and similarly, $g(B,\pi)$ is distributed according to Q. Finally, $\Pr[f(P,\pi) = g(Q,\pi)] \geq \Pr[i_A = i_B]$. To bound this probability, note that $\mu(A \cap B) = (1-\delta)/|\Omega|$ and $\mu(A \cup B) = (1+\delta)/|\Omega|$.

$$\Pr[f(P,\pi) = g(Q,\pi)] \;\geq\; \Pr[i_A = i_B] \;=\; \frac{\mu(A \cap B)}{\mu(A \cup B)} \;=\; \frac{1-\delta}{1+\delta} \;=\; 1 - \frac{2\delta}{1+\delta}.$$

We can ignore the possibility that no index i_A exists satisfying $(\omega_{i_A}, r_{i_A}) \in A$ (similarly for B) since this happens with probability 0. □

5 Discussion and open questions

An immediate open question is to resolve Conjecture 3.4. We reflect on some further open questions that are raised by the results discussed in this paper.

5.1 Case of *negatively correlated* sets

In the context of Conjecture 3.4, even in the setting where the set sizes are allowed to vary slightly, our knowledge is somewhat incomplete. Lemma 2.2 shows optimality of the MinHash strategy when

Algorithm 3: The strategy of Kleinberg–Tardos and Holenstein [7, 6]

Alice's input: $P \in \Delta_\Omega$; let $A := \{(\omega, p) \in \Omega \times [0,1] : p < P(\omega)\}$

Bob's input: $Q \in \Delta_\Omega$; let $B := \{(\omega, q) \in \Omega \times [0,1] : q < Q(\omega)\}$

Shared randomness: An infinite sequence $\pi = ((\omega_1, r_1), (\omega_2, r_2), \ldots)$ where each (ω_i, r_i) is i. i. d. sampled uniformly from $\Omega \times [0,1]$.

Strategy:

- $f(P, \pi) := \omega_{i_A}$, where i_A is the smallest index such that $(\omega_{i_A}, r_{i_A}) \in A$
- $g(Q, \pi) := \omega_{i_B}$, where i_B is the smallest index such that $(\omega_{i_B}, r_{i_B}) \in B$

$(A,B) \sim \mathcal{D}_p$. In this case, A and B are independent and each of them is p-biased, so $|A| \approx p \cdot n$, $|B| \approx p \cdot n$ and $|A \cap B| \approx p^2 \cdot n$. A simple reduction to Lemma 2.2 also implies the optimality of the MinHash strategy in the case where A and B are *positively correlated*. Specifically for parameters $\alpha > p$, consider the following distribution $\mathcal{D}_{p,\alpha}$ on pairs (A,B) of subsets of $[n]$, where we first sample $S \subseteq [n]$ by independently including each element of $[n]$ with probability p/α, and then independently including every $i \in S$ in each of A and B with probability α. In this case, $|A| \approx p \cdot n$, $|B| \approx p \cdot n$ and $|A \cap B| \approx \alpha p \cdot n > p^2 \cdot n$. Even if we reveal S to both Alice and Bob, Lemma 2.2 implies a lower bound of $2\delta/(1+\delta)$ on the error probability (for large enough n). It is unclear if the optimality holds even in the case where A and B are *negatively correlated*, i. e., when $|A| \approx p \cdot n$, $|B| \approx p \cdot n$ and $|A \cap B| \ll p^2 \cdot n$.

5.2 Fine-grained understanding of $\mathcal{G}$-restricted correlated sampling

As alluded to in the Introduction, in the setting where P and Q are flat distributions on subsets of Ω of different sizes, there is a strategy with lower error than provided in Theorem 1.2. In particular, for $P = \mathcal{U}(A)$ and $Q = \mathcal{U}(B)$ where $|A| \neq |B|$, the MinHash strategy gives an error probability of

$$1 - \frac{|A \cap B|}{|A \cup B|} \tag{5.1}$$

(which is the Jaccard distance between A and B). However, naïvely using the strategy of Kleinberg–Tardos and Holenstein would give an error probability of

$$1 - \frac{|A \cap B|}{|A \cup B| + \big||A| - |B|\big|}, \tag{5.2}$$

which is higher than the Jaccard distance when $|A| \neq |B|$. This implies that the strategy of Kleinberg–Tardos and Holenstein is not always optimal. Thus, it will be interesting to identify the right measure that captures the minimum error of a general $\mathcal{G}$-restricted correlated sampling strategy.

5.3 Correlated sampling for infinite spaces

While this paper dealt with correlated sampling for finite sets Ω, it might also be interesting to study it for infinite sets. This needs to be defined carefully in a measure theoretic sense, which could be done as follows. Consider a measure space $(\Omega, \mathcal{F}, \mu)$, where Ω is the sample space, $\mathcal{F}$ is a σ-algebra over Ω and μ is a finite measure on $(\Omega, \mathcal{F})$. Let $\Delta_{(\Omega,\mathcal{F},\mu)}$ be the set of all probability measures over $(\Omega, \mathcal{F})$ that are absolutely continuous with respect to μ. The respective inputs of Alice and Bob are probability measures P and Q in $\Delta_{(\Omega,\mathcal{F},\mu)}$. A correlated sampling strategy for $(\Omega, \mathcal{F}, \mu)$ is given by a pair of functions $f, g : \Delta_{(\Omega,\mathcal{F},\mu)} \times \mathcal{R} \to \Omega$, where f and g are required to be measurable in their second argument $r \in \mathcal{R}$. In order to define the error guarantee in terms of $\Pr_{r\sim\mathcal{R}}[f(P,r) \neq g(Q,r)]$, however, we require that the event $\{(\omega, \omega) : \omega \in \Omega\}$ be measurable in $(\Omega \times \Omega, \mathcal{F} \otimes \mathcal{F})$. This is true, for example, when Ω is a separable metric space equipped with the standard Borel algebra (see Chapters 3, 4 in [14]). We will assume this to be the case in the discussion henceforth and it might be useful to keep in mind a concrete example such as $\Omega = [0, 1]$, equipped with the Lebesgue measure.

To the best of our knowledge, it remains open whether there exists a correlated sampling strategy for general measure spaces $(\Omega, \mathcal{F}, \mu)$ with any *non-trivial* error bound, that is, to get $\varepsilon(\delta) < 1$ for all $\delta < 1$. This is in sharp contrast to coupling, where any two probability measures P and Q with $d_{\mathrm{TV}}(P,Q) = \delta$ over $(\Omega, \mathcal{F})$ can be coupled with a disagreement probability of at most δ.

Suppose the inputs P and Q are promised to be such that the corresponding Radon–Nikodym derivatives (a. k. a. *densities*) $dP/d\mu$ and $dQ/d\mu$ are bounded everywhere by a known constant c. Then it is possible to generalize the strategy of Kleinberg–Tardos and Holenstein (Algorithm 3) and get the same error guarantee as in Theorem 1.2; this can be done by using μ instead of the uniform measure on Ω and replacing $[0, 1]$ by $[0, c]$.

However, the problem gets challenging if there is no promised upper bound on the Radon–Nikodym derivatives. One explanation for why this challenge is not faced in obtaining a coupling is because knowing both P and Q, we can always take $\mu' = (P + Q)/2$ as a measure with respect to which both P and Q are absolutely continuous and more strongly, the Radon–Nikodym derivatives $dP/d\mu'$ and $dQ/d\mu'$ are never greater than 2. On the other hand, for correlated sampling, the players do not have access to such a common μ'.

It might also be interesting to study a generalized notion of error in correlated sampling strategies where we wish to minimize $\mathbb{E}_{r\sim\mathcal{R}}[d(f(P,r), g(Q,r))]$ for some metric $d : \Omega \times \Omega \to \mathbb{R}_{\geq 0}$ over Ω. The error guarantee studied in this paper corresponds to the discrete metric $d(x,y) = \mathbb{1}\{x \neq y\}$. For $\Omega \subseteq \mathbb{R}$, such as $\Omega = [0, 1]$, we might alternatively want to consider $d(x,y) = |x - y|$. Since a correlated sampling strategy induces a coupling, this notion of error can never be lower than the Wasserstein distance $W_1(P,Q)$ (also known as Earth-Mover distance) between the distributions P and Q. To the best of our knowledge, it remains open in this setting, whether correlated sampling strategies can get an error that is never larger than some function of $W_1(P,Q)$.

Acknowledgements

We thank anonymous ToC reviewers for their feedback that has significantly helped improve the presentation of this paper. We are also grateful to Laci Babai for detailed comments that helped reorganize the paper and brought in more clarity to the discussion about correlated sampling for infinite spaces.

References

[1] BOAZ BARAK, MORITZ HARDT, ISHAY HAVIV, ANUP RAO, ODED REGEV, AND DAVID STEURER: Rounding parallel repetitions of unique games. In *Proc. 49th FOCS*, pp. 374–383. IEEE Comp. Soc., 2008. [doi:10.1109/FOCS.2008.55] 2

[2] ANDREI Z. BRODER: On the resemblance and containment of documents. In *Proc. Compression and Complexity of Sequences (SEQUENCES'97)*, pp. 21–29. IEEE Comp. Soc., 1997. [doi:10.1109/SEQUEN.1997.666900] 2, 3, 10, 11

[3] MOSES S. CHARIKAR: Similarity estimation techniques from rounding algorithms. In *Proc. 34th STOC*, pp. 380–388. ACM Press, 2002. [doi:10.1145/509907.509965] 2

[4] SREENIVAS GOLLAPUDI AND RINA PANIGRAHY: A dictionary for approximate string search and longest prefix search. In *15th Internat. Conf. on Information and Knowledge Managment (CIKM'06)*, pp. 768–775. ACM Press, 2006. [doi:10.1145/1183614.1183723] 3

[5] BERNHARD HAEUPLER, MARK MANASSE, AND KUNAL TALWAR: Consistent weighted sampling made fast, small, and easy, 2014. [arXiv:1410.4266] 3

[6] THOMAS HOLENSTEIN: Parallel repetition: simplifications and the no-signaling case. *Theory of Computing*, 5(8):141–172, 2009. Preliminary version in STOC'07. [doi:10.4086/toc.2009.v005a008] 2, 3, 10, 11, 12

[7] JON KLEINBERG AND ÉVA TARDOS: Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. *J. ACM*, 49(5):616–639, 2002. Preliminary version in FOCS'99. [doi:10.1145/585265.585268] 2, 3, 10, 11, 12

[8] JACOBUS HENDRICUS VAN LINT AND RICHARD MICHAEL WILSON: *A Course in Combinatorics*. Cambridge Univ. Press, 2001. [doi:10.1017/CBO9780511987045] 8

[9] MARK MANASSE, FRANK MCSHERRY, AND KUNAL TALWAR: Consistent weighted sampling. Technical Report MSR-TR 2010-73, 2010. 3

[10] UDI MANBER: Finding similar files in a large file system. In *USENIX Winter Tech. Conf. (WTEC'94)*, pp. 1–10. USENIX Assoc., 1994. Link at ACM DL. Available as U. Arizona CS TR93-33. 3

[11] MICHAEL MITZENMACHER AND ELI UPFAL: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge Univ. Press, 2005. 5

[12] ANUP RAO: Parallel repetition in projection games and a concentration bound. *SIAM J. Comput.*, 40(6):1871–1891, 2011. Preliminary version in STOC'08. [doi:10.1137/080734042] 2

[13] RONALD L. RIVEST: Symmetric Encryption via Keyrings and ECC, 2016. Slide presentation, available on author's home page. 2, 3, 4, 7, 8, 9

[14] HERMANN THORISSON: *Coupling, Stationarity, and Regeneration*. Springer, 2000. 3, 13

AUTHORS

Mohammad Bavarian [about the author]
Software Engineer
Rubrik Inc.
Palo Alto, CA, USA
mobavarian@gmail.com
https://bavarian.dev/

Badih Ghazi [about the author]
Research Scientist
Google Research
Mountain View, CA, USA
badihghazi@gmail.com
https://sites.google.com/view/badihghazi/home

Elad Haramaty [about the author]
Research Scientist
Amazon Inc.
Tel Aviv, Israel
seladh@gmail.com

Pritish Kamath [about the author]
Postdoctoral Scholar
Toyota Technological Institute at Chicago
Chicago, IL, USA
pritish@alum.mit.edu
https://pritishkamath.github.io/

Ronald L. Rivest [about the author]
Institute Professor
Massachusetts Institute of Technology
Cambridge, MA, USA
rivest@mit.edu
https://people.csail.mit.edu/rivest/

Madhu Sudan [about the author]
Gordon McKay Professor of Computer Science
Harvard John A. Paulson School of Engineering and Applied Sciences
Cambridge, MA, USA
madhu@cs.harvard.edu
http://madhu.seas.harvard.edu/

ABOUT THE AUTHORS

MOHAMMAD BAVARIAN obtained his Ph. D. under Prof. Madhu Sudan from MIT in 2017 working on quantum computing and complexity theory. He now works as a software developer at Rubrik Inc. and spends his free time thinking about startups, business, and engineering. While not busy with above, he also dabbles in programming competitions and Machine Learning. He enjoys the beautiful palm trees and sunshine of California, but still sometimes misses the city life of the East Coast.

BADIH GHAZI is a Research Scientist in the Algorithms team at Google. His current research interests include algorithmic aspects of machine learning, differential privacy, error-correcting codes and communication under uncertainty. He completed his Ph. D. in 2018 at the Electrical Engineering and Computer Science (EECS) department at MIT where he was advised by Professors Madhu Sudan and Ronitt Rubinfeld. Previously, he got his M.S. in EECS also from MIT, and his B.Eng. in Computer and Communications Engineering from the American University of Beirut. During his graduate studies, he was awarded an IBM Ph. D. Fellowship and an MIT Irwin and Joan Jacobs Presidential Fellowship.

ELAD HARAMATY is a Research Scientist at Amazon Israel. Previously, he held postdoctoral researcher positions at Harvard University (hosted by Madhu Sudan) and at Northeastern University (hosted by Emanuele Viola). He earned his Ph. D. from the Technion Israel Institute of Technology under the guidance of Amir Shpilka. His research interests lie in a broad range of areas of theoretical computer science, especially in algebraic complexity. In his work he has studied mostly the structure, testability and applications of polynomials and algebraic codes. He nowadays mostly works on various aspects of Search and Machine Learning.

PRITISH KAMATH is a postdoctoral scholar at the Toyota Technological Institute at Chicago. He completed his Ph. D. at MIT, co-advised by Professors Madhu Sudan and Ronitt Rubinfeld. Prior to that, he completed a B. Tech. in Computer Science at IIT Bombay in 2012 after which he was a Research Fellow at Microsoft Research India until 2013, where he was mentored by Neeraj Kayal. He has broad interests in complexity theory and has worked in areas touching upon communication complexity, information theory, Boolean and algebraic circuit complexity and proof complexity and most recently is also interested in foundational aspects of machine learning. He likes to juggle multiple things in life; sometimes on a bicycle.

RONALD L. RIVEST is an Institute Professor in the Department of Electrical Engineering and Computer Science at MIT. He is a member of MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL), a member of the lab's Theory of Computation Group and is a leader of its Cryptography and Information Security Group.

Ron has current research interests in cryptography, computer and network security, voting systems, and algorithms. In the past he has also worked extensively in the area of machine learning.

Ron is a coauthor (with Thomas Cormen, Charles Leiserson, and Clifford Stein) of the well-known text Introduction to Algorithms, published by MIT Press. Over 500,000 copies of this text have been sold. It has been translated into 12 languages.

Ron is an inventor of the RSA public-key cryptosystem. He has extensive experience in cryptographic design and cryptanalysis, and has published numerous papers in these areas. He is a founder of RSA Data Security. (RSA was bought by Security Dynamics; the combined company was renamed to RSA Security, and later purchased by EMC), and is also a co-founder of Verisign and of Peppercoin.

Ron is a member of the CalTech/MIT Voting Technology Project. He served 2004–2009 on the Technical Guidelines Development Committee (TGDC), advisory to the Election Assistance Commission, developing recommendations for voting system certification standards; he was chair of the TGDC's Computer Security and Transparency Subcommittee. He also serves on the Board of the Verified Voting Foundation. He is a member of a Scantegrity team developing and testing voting systems that are verifiable "end-to-end." He has worked extensively on statistical post-election tabulation audits, of both the "risk-limiting audit" and "Bayesian" flavors.

Ron is a member of the Center for Science of Information.

MADHU SUDAN is a Gordon McKay Professor in the John A. Paulson School of Engineering and Applied Sciences at Harvard University, where he has been since 2015. Madhu got his Bachelors degree (apparently in Technology) from IIT Delhi in 1987 where he was introduced to theoretical computer science by Sachin Maheshwari. He got his Ph. D. from U. C. Berkeley in 1992 under the supervision of Umesh Vazirani. Madhu subsequently spent time at, and was even paid by, IBM, MIT and Microsoft.

While Madhu's research explores communication and computational complexity, he prefers simplicity, and is especially proud of his exposition (with Peter Gemmell) on the Berlekemp-Welch decoding algorithm and his exposition (with David Xiang) of the analysis of the Lempel-Ziv compression algorithm.