Bounds on the Efficiency of "Black-Box" Commitment Schemes

Omer Horvitz* and Jonathan Katz**

Department of Computer Science, University of Maryland
College Park, MD 20742
{horvitz,jkatz}@cs.umd.edu

Abstract. Constructions of cryptographic primitives based on general assumptions (e.g., the existence of one-way functions) tend to be less efficient than constructions based on specific (e.g., number-theoretic) assumptions. This has prompted a recent line of research aimed at investigating the best possible efficiency of (black-box) constructions based on general assumptions. Here, we present bounds on the efficiency of statistically-binding commitment schemes constructed using black-box access to one-way permutations; our bounds are tight for the case of *perfectly*-binding schemes. We present the bounds in an extension of the Impagliazzo-Rudich model; that is, we show that any construction beating our bounds would imply the unconditional existence of a one-way function (from which a commitment scheme could be constructed "from scratch"). Our analysis is the first in the area to pertain directly to an information-theoretic component of the security notion.

1 Introduction

A central goal of modern cryptography has been to characterize the minimal assumptions needed to construct cryptographic tools and protocols. For example, we now know that one-way functions are sufficient for constructing pseudorandom generators (PRGs) [2, 18, 8, 9], universal one-way hash functions (UOWHFs) and digital signature schemes [14, 16], private-key encryption schemes [6] and commitment schemes [13]. In each of these cases, one-way functions are also known to be necessary [10, 16], making the characterization exact. While important from a theoretical point of view, the above constructions have had limited practical impact due to their inefficiency. In practice, more efficient constructions based on stronger assumptions tend to be used. Moreover, for all of the examples listed above, no constructions from general assumptions that improve on the efficiency of the original solutions are known.

The tension between general and efficient constructions has motivated a recent line of research, aimed at understanding the minimal assumptions needed for *efficient* constructions of various primitives. Put another way, we ask what is

* Supported by U.S. Army Research Office award DAAD19-01-1-0494.
** Supported by NSF CAREER award #0447075.

the best possible efficiency that can be achieved (for a particular primitive of interest) if we limit ourselves to constructions based on the weakest cryptographic assumptions. Initial work by Kim, Simon, and Tetali [12] demonstrated (non-tight) bounds on the efficiency of constructing universal one-way hash functions from one-way permutations. Extending their results, Gennaro, Gertner, Katz, and Trevisan [4] showed that known constructions of UOWHFs based on one-way permutations are in fact *optimal*; they also show efficiency bounds for the case of PRGs, private-key encryption schemes, and digital signatures based on one-way permutations, as well as for the case of public-key encryption schemes based on trapdoor permutations.

It is noted that in all the above bounds, analysis pertains directly to a *computational* aspect of the security notion at hand: for PRGs, it is the indistinguishability of the output from random; for UOWHFs – the computational hardness of finding a collision with a given element for a random member of the family; for encryption – the computational indistinguishability of encryptions of distinct messages; and for signatures – the computational hardness of generating a forgery. For commitment schemes, appealing to computational formulations of both the hiding and binding properties resists similar analysis. Taking either of the properties as *information-theoretic* introduces new challenges. Indeed, finding lower bounds on the efficiency of commitments from one-way permutations is left as an open problem in [4].

Our results. Here, we prove lower bounds on the efficiency of (black-box) constructions of statistically-binding, interactive commitment schemes based on one-way permutations. A commitment scheme for m-bit messages is an interactive protocol between a sender $\mathcal{S}$ and a receiver $\mathcal{R}$, where $\mathcal{S}$ takes as input a message $M \in \{0,1\}^m$. Call the view of $\mathcal{R}$ during an interaction with $\mathcal{S}$ the *commitment*. Informally, a statistically-binding commitment scheme satisfies two properties (defined more formally in Section 2.2): (1) *Hiding*: the distribution over commitments to M_0 is computationally indistinguishable from the distribution over commitments to M_1, for any M_0, M_1 and a possibly malicious receiver; and (2) (statistical) *Binding*: the probability (over coin tosses of the honest receiver $\mathcal{R}$) that there exist different M, M' and coins s, s' for $\mathcal{S}$ such that the corresponding commitments to M, M' are identical is at most ε_b. When $\varepsilon_b = 0$ we say the scheme is *perfectly binding*. A permutation $\pi : \{0,1\}^n \to \{0,1\}^n$ is *one-way with security* S if any circuit of size $\leq S$ inverts π on a fraction $\leq 1/S$ of its inputs.

We show that any black-box construction[1] of a statistically-binding commitment scheme based on one-way permutations must have the sender query its oracle $\Omega((m - \log(1 + 2^m \cdot \varepsilon_b))/\log S)$ times. For perfectly-binding schemes, the bound translates to $\Omega(m/\log S)$; for non perfectly-binding schemes, the bound becomes $\Omega(\log \varepsilon_b^{-1}/\log S)$ when $m > \log \varepsilon_b^{-1}$, and is essentially $\Omega(m/\log S)$ otherwise. We prove our bounds in the model of Gennaro, et al. [4] which, in turn, extends the model of Impagliazzo and Rudich [11, 15] (see [4] for a detailed com-

[1] Using the terminology of [17] (see also [4]), our bounds hold for *weak* black-box constructions and thus rule out stronger types of black-box constructions as well.

parison). In particular, we show that the existence of a construction beating our efficiency bound would imply the unconditional existence of a one-way function (and hence $P \neq NP$).

For the case of perfectly-binding schemes constructed from one-way permutations, our bound matches the efficiency achieved by the construction of Blum [1] as extended by Goldreich and Levin [8] (see also [5]). For statistically-binding schemes, our bound matches the efficiency achieved by the construction of Naor [13] for single-bit messages (setting $\varepsilon_b = 2^{-S}$). Closing the gap for the case of longer messages remains an interesting open problem.

2 Preliminaries

2.1 One-Way Functions and Permutations

Let A^f denote a circuit A with oracle access to the function f. We say that a function $f : \{0,1\}^n \rightarrow \{0,1\}^n$ is (S, ε)-*one-way* if for every circuit A of size $\leq S$ we have

$$\Pr_x[A^f(f(x)) \in f^{-1}(f(x))] \leq \varepsilon.$$

To reduce the number of parameters, we will call a function S-*hard* if it is $(S, 1/S)$-one way.

Let Π_t denote the set of all permutations over $\{0,1\}^t$. Then:

Theorem 1 ([4]). *For sufficiently large t, a random $\pi \in \Pi_t$ is $2^{t/5}$-hard with probability at least $1 - 2^{-2^{t/2}}$.*

Let $a||b$ denote the concatenation of strings a and b. For $t < n$, let $\Pi_{t,n}$ denote the subset of Π_n such that $\pi \in \Pi_{t,n}$ iff $\pi(a||b) = \hat{\pi}(a)||b$ for some $\hat{\pi} \in \Pi_t$. A corollary to Theorem 1 is that if $t = 5 \log S_p$, then for $n > t$, a randomly chosen $\pi \in \Pi_{t,n}$ is S_p-hard with high probability.

Corollary 1 ([4]). *For sufficiently large t and $n > t$, a random $\pi \in \Pi_{t,n}$ is $2^{t/5}$-hard with probability at least $1 - 2^{-2^{t/2}}$.*

2.2 Commitment Schemes

A *commitment scheme* for m-bit messages is defined by a pair of probabilistic, interactive algorithms $(\mathcal{S}, \mathcal{R})$. The inputs to $\mathcal{S}$, the *sender*, are a message $M \in \{0,1\}^m$ and randomness s, while the input to $\mathcal{R}$, the *receiver*, is randomness r. Let $\langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle$ denote the receiver's view of an interaction with the sender on the specified inputs; this view simply consists of the receiver's randomness and the messages it receives from the sender during the interaction (when the receiver makes queries to an oracle, the view additionally includes the answers it receives from the oracle). We call $\langle \mathcal{S}(M; s), \mathcal{R}(r) \rangle$ a *commitment*. A *decommitment* consists of a message and sender randomness. We say that a commitment $C = \langle \mathcal{S}(M;s), \mathcal{R}(r) \rangle$ can be *decommitted* to a message M' if there exists a string s' such that $\langle \mathcal{S}(M';s'), \mathcal{R}(r) \rangle = C$. For a message M, let

$$\langle \mathcal{S}(M), \mathcal{R} \rangle \stackrel{\text{def}}{=} \left\{ s, r \stackrel{R}{\leftarrow} \{0,1\}^* : \langle \mathcal{S}(M;s), \mathcal{R}(r) \rangle \right\}.$$

We say that two distributions $\mathcal{X}, \mathcal{Y}$ are (S, ε)-*indistinguishable*, and write $\mathcal{X} \stackrel{(S,\varepsilon)}{\approx} \mathcal{Y}$, if for every circuit A of size $\leq S$, we have

$$\left| \Pr_{x \in \mathcal{X}}[A(x) = 0] - \Pr_{x \in \mathcal{Y}}[A(x) = 0] \right| \leq \varepsilon.$$

Definition 1. *Let $(\mathcal{S}, \mathcal{R})$ be a commitment scheme for m-bit messages. We say that $(\mathcal{S}, \mathcal{R})$ is (S_h, ε_h)-*hiding *if for every circuit $\mathcal{R}^*$ of size $\leq S_h$, for all $M_0, M_1 \in \{0,1\}^m$, we have*

$$\langle \mathcal{S}(M_0), \mathcal{R}^* \rangle \stackrel{(S_h,\varepsilon_h)}{\approx} \langle \mathcal{S}(M_1), \mathcal{R}^* \rangle.$$

*We say that $(\mathcal{S}, \mathcal{R})$ is ε_b-*binding *if*

$$\Pr_r \left[\begin{array}{c} \exists M, M' \in \{0,1\}^m, s, s' \text{ s.t.} \\ \langle \mathcal{S}(M;s), \mathcal{R}(r) \rangle = \langle \mathcal{S}(M';s'), \mathcal{R}(r) \rangle \end{array} \right] \leq \varepsilon_b.$$

*We say that $(\mathcal{S}, \mathcal{R})$ is ε_b-*binding for an honest sender *if for all $M \in \{0,1\}^m$, we have*

$$\Pr_{s,r} \left[\begin{array}{c} \exists M' \neq M, s' \text{ s.t.} \\ \langle \mathcal{S}(M;s), \mathcal{R}(r) \rangle = \langle \mathcal{S}(M';s'), \mathcal{R}(r) \rangle \end{array} \right] \leq \varepsilon_b.$$

*If $\varepsilon_b = 0$ we say the scheme is *perfectly-binding*. Finally, we say that $(\mathcal{S}, \mathcal{R})$ is $(S_h, \varepsilon_h, \varepsilon_b)$-*secure *(resp., *secure for an honest sender*) if $(\mathcal{S}, \mathcal{R})$ is (S_h, ε_h)-hiding and ε_b-binding (resp., binding for an honest sender).*

*A *construction *of a commitment scheme for m-bit messages (based on one-way permutations) is a pair of oracle algorithms $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ such that, for all $\pi \in \Pi_n$, $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is a commitment scheme for m-bit messages. We say that $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ is $(S_p, S_h, \varepsilon_h, \varepsilon_b)$-*secure *(resp., *secure for an honest sender*) if for every $\pi \in \Pi_n$ that is S_p-hard, $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is (S_h, ε_h)-hiding and for every $\pi \in \Pi_n$ (even those not S_p-hard), $(\mathcal{S}^\pi, \mathcal{R}^\pi)$ is ε_b-binding (resp., binding for an honest sender).* $\diamondsuit$

We note that, in the terminology of [17], our definitions of security for commitment constructions fall in the category of *weak black-box* constructions. This is evident in that in the definition of the hiding property, the distinguishing algorithm is *not* given oracle access to π. As we prove lower bounds, our results apply to stronger notions of black-box constructions as well, and in particular to *semi-* and *fully black-box* constructions (see [17] for further details.)

2.3 Pairwise-Independent Hashing

Let H be a family of functions mapping m-bit strings to m'-bit strings. We assume that the following can be done in time polynomial in m: (1) selecting a function $h \in H$ uniformly at random; (2) given $h \in H$ and $x \in \{0,1\}^m$, evaluating $h(x)$; and (3) given h^*, deciding whether $h^* \in H$ or not. We say that H is a *pairwise-independent hash family* (following [3]) if for any distinct $x_1, x_2 \in \{0,1\}^m$ and any $y_1, y_2 \in \{0,1\}^{m'}$ we have:

$$\Pr_{h \in H}[h(x_1) = y_1 \wedge h(x_2) = y_2] = 2^{-2m'}.$$

3 The Lower Bound

Let $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ be an $(S_p, S_h, \varepsilon_h, \varepsilon_b)$-secure construction of a commitment scheme for m-bit messages (based on one-way permutations). We prove that unless $\mathcal{S}$ queries its oracle at least $\Omega((m - \log(1 + 2^m \cdot \varepsilon_b))/\log S_p)$ times, there exists (constructively) a commitment scheme $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ secure for an honest sender which *does not* require access to any oracle (i.e., the scheme is secure unconditionally). Strengthening a result of Impagliazzo and Luby [10] (cf. Lemma 1), this implies the unconditional existence of a one-way function.

We describe here the intuition behind our proof, assuming for ease of exposition that $(\mathcal{S}, \mathcal{R})$ is a non-interactive, perfectly-binding scheme. As in [4], our starting point is that a random $\pi \in \Pi_{t,n}$ (for $t = \Theta(\log S_p)$) is S_p-hard with all but negligible probability (cf. Corollary 1). So, consider a scheme in which the sender $\mathcal{S}'$ simulates a random $\pi \in \Pi_{t,n}$ for $\mathcal{S}$, using independent random coins[2] y. Note that decommitment here consists of both s and y.

It is quite straightforward to show that the resulting scheme still satisfies hiding. Binding, however, is another matter. Note first that we need only show binding for an *honest* sender. That is, we may assume that s and y have been chosen honestly (i.e., uniformly at random) by $\mathcal{S}'$. Letting P denote the pairs of t-bit prefixes of oracle queries made by $\mathcal{S}'$ during the computation and t-bit prefixes of corresponding answers, the question then is whether there exist s', y' (with associated query/answer prefixes P') that form a decommitment of C to a different message. We first observe, informally, that having $P' = P$ will not help: this is because the original scheme is binding for any choice of π. In fact, this further implies that for any choice of P' there exists at most one message to which C can be decommitted. Said another way, this means that C can be decommitted to at most

$$2^{2t|P|} = 2^{2tq}$$

different messages (where q is the number of oracle queries made by the original sender). Although this clearly violates binding, it does limit the space of possible

[2] This can easily be done by selecting random t-bit answer-prefixes to new t-bit query-prefixes, as needed.

messages *somewhat* as long as $2^{2tq} < m$. We now show how to "bootstrap" this to achieve "full" binding with noticeable probability.

The idea is to modify the scheme as follows: the sender $\bar{\mathcal{S}}$ now chooses a function h uniformly at random from a pairwise-independent hash family mapping m-bit strings to m-bit strings. On message M, $\bar{\mathcal{S}}$ now runs $\mathcal{S}'$ *twice* (simulating a random π as before), first on input M and then on input $h(M)$. The commitment also includes h. Hiding for the derived scheme follows easily, as before. As for binding, denote the two sub-commitments of $\bar{\mathcal{S}}$ by C_1 and C_2. We have said already that for each of these commitments there are sets S_1, S_2, each of size at most $2^{2t|P|} < 2^m$, such that C_1 might potentially be decommitted to any message from S_1 and similarly for C_2, S_2. But now the commitment $C_1\|C_2\|h$ can only be decommitted as a different message if there exists an $M' \in S_1$ (with $M' \neq M$) for which $h(M') \in S_2$. The crux of our proof is to show that, with constant probability over choice of h (and for certain values of $|P|$), such an M' does not exist. This immediately implies that binding holds with constant probability.

3.1 A Technical Lemma

We begin by showing that the existence of a commitment scheme secure for honest senders implies the existence of a one-way function. Although the result can be derived from [10], we give a simple and more direct proof here.

Lemma 1. *Let ($\mathcal{S}$,$\mathcal{R}$) be a commitment scheme for m-bit messages which is $(S_h, \varepsilon_h, \varepsilon_b)$-secure for an honest sender. Let $S_{\mathcal{S}}, S_{\mathcal{R}}$ be the sizes of the circuits computing $\mathcal{S}, \mathcal{R}$, respectively. Then for any $\ell \in \mathbb{N}$, there exists an $(S_h - \ell(S_{\mathcal{S}} + S_{\mathcal{R}}), \ell\varepsilon_h + 2\varepsilon_b^\ell)$-one-way function.*

Proof. Via a standard hybrid argument, running ℓ independent interactions between $\mathcal{S}$ and $\mathcal{R}$ on the same, fixed message yields an $(S_h, \ell\varepsilon_h, \varepsilon_b^\ell)$-secure commitment scheme for m-bit messages. Let $(\mathcal{S}_\ell, \mathcal{R}_\ell)$ denote this scheme.

Let $f(M, s, r) \stackrel{\text{def}}{=} \langle \mathcal{S}_\ell(M; s), \mathcal{R}_\ell(r) \rangle$. We claim that f is $(S_h - \ell(S_{\mathcal{S}} + S_{\mathcal{R}}), \ell\varepsilon_h + 2\varepsilon_b^\ell)$-one-way. Assume the contrary. Then there exists a circuit B of size at most $S_h - \ell(S_{\mathcal{S}} + S_{\mathcal{R}})$ such that

$$\mathsf{Succ}_{B,f}^{\mathsf{owf}} \stackrel{\text{def}}{=} \Pr_{M,s,r}[B(f(M,s,r)) \in f^{-1}(f(M,s,r))] > \ell\varepsilon_h + 2\varepsilon_b^\ell.$$

We use B to construct a circuit A that violates the hiding property of $(\mathcal{S}_\ell, \mathcal{R}_\ell)$. On input (M_0, M_1, C), A computes $(M', s', r') \leftarrow B(C)$, and checks whether $f(M', s', r') \stackrel{?}{=} C$ and whether $M' \stackrel{?}{=} M_0$. If both hold, A outputs 0; otherwise, it outputs 1. Note that $|A| = |B| + \ell(S_{\mathcal{S}} + S_{\mathcal{R}}) \leq S_h$.

Let $\mathsf{Bad} \stackrel{\text{def}}{=} \{(M, s, r) | \exists M' \neq M, s' : \langle \mathcal{S}_\ell(M; s), \mathcal{R}_\ell(r) \rangle = \langle \mathcal{S}_\ell(M'; s'), \mathcal{R}_\ell(r) \rangle\}$. In what follows, note that if $(M', s', r') \in f^{-1}(f(M, s, r))$ then $r' = r$, as r is included in the receiver's view. We have:

$$
\begin{aligned}
&\Pr_{\substack{M_0,M_1\\ C\in\langle\mathcal{S}_\ell(M_0),\mathcal{R}_\ell\rangle}}[A(M_0,M_1,C)=0] \\
&\quad = \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_0,s,r)):\\ (M',s',r')\in f^{-1}(f(M_0,s,r))\bigwedge M'=M_0\end{matrix}\right] \\
&\quad \geq \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_0,s,r)):\\ (M',s',r')\in f^{-1}(f(M_0,s,r))\bigwedge(M_0,s,r)\notin\mathsf{Bad}\end{matrix}\right] \\
&\quad = \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_0,s,r)):\\ (M',s',r')\in f^{-1}(f(M_0,s,r))\end{matrix}\right] \\
&\qquad - \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_0,s,r)):\\ (M',s',r')\in f^{-1}(f(M_0,s,r))\bigwedge(M_0,s,r)\in\mathsf{Bad}\end{matrix}\right] \\
&\quad \geq \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_0,s,r)):\\ (M',s',r')\in f^{-1}(f(M_0,s,r))\end{matrix}\right] - \Pr_{\substack{M_0,M_1\\ s,r}}[(M_0,s,r)\in\mathsf{Bad}] \\
&\quad \geq \mathsf{Succ}^{\mathsf{owf}}_{B,f}-\varepsilon_b^\ell \quad > \quad \ell\varepsilon_h+\varepsilon_b^\ell.
\end{aligned}
$$

Furthermore, we have:

$$
\begin{aligned}
&\Pr_{\substack{M_0,M_1\\ C\in\langle\mathcal{S}_\ell(M_1),\mathcal{R}_\ell\rangle}}[A(M_0,M_1,C)=0] \\
&\quad = \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_1,s,r)):\\ (M',s',r')\in f^{-1}(f(M_1,s,r))\bigwedge M'=M_0\end{matrix}\right] \\
&\quad \leq \Pr_{\substack{M_0,M_1\\ s,r}}\left[\begin{matrix}(M',s',r')\leftarrow B(f(M_1,s,r)):\\ (M',s',r')\in f^{-1}(f(M_1,s,r))\bigwedge(M_1,s,r)\in\mathsf{Bad}\end{matrix}\right] \\
&\quad \leq \Pr_{\substack{M_0,M_1\\ s,r}}[(M_1,s,r)\in\mathsf{Bad}]\leq\varepsilon_b^\ell.
\end{aligned}
$$

Putting everything together, we have:

$$
\left|\Pr_{\substack{M_0,M_1\\ C\in\langle\mathcal{S}_\ell(M_0),\mathcal{R}_\ell\rangle}}[A(M_0,M_1,C)=0]-\Pr_{\substack{M_0,M_1\\ C\in\langle\mathcal{S}_\ell(M_1),\mathcal{R}_\ell\rangle}}[A(M_0,M_1,C)=0]\right|>\ell\varepsilon_h.
$$

But this implies that there exist two messages M_0, M_1 for which A can distinguish $\langle\mathcal{S}_\ell(M_0),\mathcal{R}_\ell\rangle$ from $\langle\mathcal{S}_\ell(M_1),\mathcal{R}_\ell\rangle$ with probability $>\ell\varepsilon_h$, contradicting the hiding property of $(\mathcal{S}_\ell,\mathcal{R}_\ell)$. ■

3.2 Main Result

Theorem 2. *Let $(\mathcal{S}^{(\cdot)},\mathcal{R}^{(\cdot)})$ be an $(S_p,S_h,\varepsilon_h,\varepsilon_b)$-secure construction of a commitment scheme for m-bit messages that expects an oracle $\pi\in\Pi_n$. Let $t=$*

$5 \log S_p$. If $\mathcal{S}$ makes $q_{\mathcal{S}} \leq (m - 2 - \log(1 + 2^{m+1} \cdot \varepsilon_b))/4t$ queries to its oracle, and $\varepsilon_h \leq 1/8 - 2^{-S_p}$, then there exists a commitment scheme for m-bit messages which is $(S_h, 1/4, 1/4)$-secure for an honest sender (without access to any oracle).

Applying Lemma 1, this implies the existence of a one-way function (without access to any oracle).

Proof. We construct non-interactive commitment scheme $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ for m-bit messages. The construction makes use of a procedure $\mathcal{SIM}$ that simulates a random permutation in $\Pi_{t,n}$, as follows [4]: $\mathcal{SIM}$ takes as input a list L. On a query $x \| x'$, where $|x| = t$, $\mathcal{SIM}$ checks whether there exists a y such that $(x, y) \in L$. If so, $\mathcal{SIM}$ returns $y \| x'$. Otherwise, it picks uniformly at random a $y \in \{0,1\}^t \setminus \{\hat{y} \mid \exists \hat{x} : (\hat{x}, \hat{y}) \in L\}$, adds (x, y) to L, and returns $y \| x'$. As usual, we let $\mathcal{SIM}(L; y)$ denote an execution of $\mathcal{SIM}$ on input L and randomness y. We let $\mathcal{SIM}_y$ denote $\mathcal{SIM}(\emptyset; y)$.

Let H be a pairwise-independent family of functions from m-bit strings to m-bit strings. Define $\bar{\mathcal{S}}$ as follows. On input a message $M \in \{0,1\}^m$, $\bar{\mathcal{S}}$ chooses uniformly at random $h \in H$ and values $s_1, r_1, y_1, s_2, r_2, y_2$. It then computes $C_1 = \langle \mathcal{S}^{\mathcal{SIM}_{y_1}}(M; s_1), \mathcal{R}^{\mathcal{SIM}_{y_1}}(r_1) \rangle$ and $C_2 = \langle \mathcal{S}^{\mathcal{SIM}_{y_2}}(h(M); s_2), \mathcal{R}^{\mathcal{SIM}_{y_2}}(r_2) \rangle$. The resulting commitment is $C_1 \| C_2 \| h$.[3] Decommitment consists of all the random coins used by $\bar{\mathcal{S}}$. We claim that $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is $(S_h, 1/4, 1/4)$-secure for an honest sender. This follows from the following two lemmata.

Lemma 2. *$(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is $(S_h, 1/4)$-hiding.*

Proof (of lemma). The hiding property of $(\mathcal{S}^{(\cdot)}, \mathcal{R}^{(\cdot)})$ guarantees that for any $\pi \in \Pi_n$ that is S_p-hard, for any circuit B of size $\leq S_h$, and for any $M_0, M_1 \in \{0,1\}^m$, we have

$$\left| \Pr_{C \in \langle \mathcal{S}^{\pi}(M_0), \mathcal{R}^{\pi} \rangle}[B(C) = 0] - \Pr_{C \in \langle \mathcal{S}^{\pi}(M_1), \mathcal{R}^{\pi} \rangle}[B(C) = 0] \right| \leq \varepsilon_h.$$

A straightforward hybrid argument shows that for any $\pi_1, \pi_2 \in \Pi_n$ that are S_p-hard, for any circuit B of size $\leq S_h$, and for any $M_0, M_1 \in \{0,1\}^m$, we have

$$\left| \Pr_{\substack{h \in H \\ C_1 \in \langle \mathcal{S}^{\pi_1}(M_0), \mathcal{R}^{\pi_1} \rangle \\ C_2 \in \langle \mathcal{S}^{\pi_2}(h(M_0)), \mathcal{R}^{\pi_2} \rangle}}[B(C_1 \| C_2 \| h) = 0] - \Pr_{\substack{h \in H \\ C_1 \in \langle \mathcal{S}^{\pi_1}(M_1), \mathcal{R}^{\pi_1} \rangle \\ C_2 \in \langle \mathcal{S}^{\pi_2}(h(M_1)), \mathcal{R}^{\pi_2} \rangle}}[B(C_1 \| C_2 \| h) = 0] \right| \leq 2\varepsilon_h.$$

Corollary 1 shows that a random $\pi \in \Pi_{t,n}$ is S_p-hard except with probability at most $2^{-S_p^{5/2}} \leq 2^{-S_p}$. Using the union bound and a simple averaging argument,

[3] The permutations simulated by $\mathcal{SIM}$ in the computations of C_1, C_2 may be different. The theorem can be strengthened (improving the bounds on ε_h) by having $\mathcal{SIM}$ provide a consistent simulation for both computations. We forgo this for simplicity.

we get that for any circuit B of size $\leq S_h$ and for any $M_0, M_1 \in \{0,1\}^m$,

$$\left| \Pr_{\substack{\pi_1,\pi_2 \in \Pi_{t,n} \\ h \in H \\ C_1 \in \langle \mathcal{S}^{\pi_1}(M_0), \mathcal{R}^{\pi_1} \rangle \\ C_2 \in \langle \mathcal{S}^{\pi_2}(h(M_0)), \mathcal{R}^{\pi_2} \rangle}} [B(C_1 \| C_2 \| h) = 0] - \Pr_{\substack{\pi_1,\pi_2 \in \Pi_{t,n} \\ h \in H \\ C_1 \in \langle \mathcal{S}^{\pi_1}(M_1), \mathcal{R}^{\pi_1} \rangle \\ C_2 \in \langle \mathcal{S}^{\pi_2}(h(M_1)), \mathcal{R}^{\pi_2} \rangle}} [B(C_1 \| C_2 \| h) = 0] \right| \leq 2\varepsilon_h + 2^{1-S_p}.$$

Since $\mathcal{SIM}$ perfectly simulates a random $\pi \in \Pi_{t,n}$, we have

$$\left| \Pr_{\substack{y_1,y_2 \\ h \in H \\ C_1 \in \langle \mathcal{S}^{\mathcal{SIM}_{y_1}}(M_0), \mathcal{R}^{\mathcal{SIM}_{y_1}} \rangle \\ C_2 \in \langle \mathcal{S}^{\mathcal{SIM}_{y_2}}(h(M_0)), \mathcal{R}^{\mathcal{SIM}_{y_2}} \rangle}} [B(C_1 \| C_2 \| h) = 0] - \Pr_{\substack{y_1,y_2 \\ h \in H \\ C_1 \in \langle \mathcal{S}^{\mathcal{SIM}_{y_1}}(M_1), \mathcal{R}^{\mathcal{SIM}_{y_1}} \rangle \\ C_2 \in \langle \mathcal{S}^{\mathcal{SIM}_{y_2}}(h(M_1)), \mathcal{R}^{\mathcal{SIM}_{y_2}} \rangle}} [B(C_1 \| C_2 \| h) = 0] \right| \leq 2\varepsilon_h + 2^{1-S_p}.$$

But that precisely means that

$$\left| \Pr_{C \in \langle \bar{\mathcal{S}}(M_0), \mathcal{R}^* \rangle} [B(C) = 0] - \Pr_{C \in \langle \bar{\mathcal{S}}(M_1), \mathcal{R}^* \rangle} [B(C) = 0] \right| \leq 2\varepsilon_h + 2^{1-S_p} \leq 1/4$$

for any $\mathcal{R}^*$ and any B of size $\leq S_h$, where the last inequality uses the assumption that $\varepsilon_h \leq 1/8 - 2^{-S_p}$. The hiding property therefore holds as claimed. $\square$

Lemma 3. *$(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is 1/4-binding for an honest sender.*

Proof (of lemma). (See also Fig. 1.) Fix an arbitrary $M \in \{0,1\}^m$. Define

$$C(M, s, r, y) \stackrel{\text{def}}{=} \langle \mathcal{S}^{\mathcal{SIM}_y}(M; s), \mathcal{R}^{\mathcal{SIM}_y}(r) \rangle.$$

Since $(\bar{\mathcal{S}}, \bar{\mathcal{R}})$ is non-interactive, we are interested in the following probability:

$$\mathsf{NoBind} \stackrel{\text{def}}{=} \Pr_{\substack{h \in H \\ s_1, r_1, y_1 \\ s_2, r_2, y_2}} \begin{bmatrix} \exists M' \neq M, s_1', y_1', s_2', y_2' \text{ s.t.} \\ C(M', s_1', r_1, y_1') = C(M, s_1, r_1, y_1), \\ C(h(M'), s_2', r_2, y_2') = C(h(M), s_2, r_2, y_2) \end{bmatrix}$$

(here, we use the fact that r_1, r_2 and h are explicit in the receiver's view).

Let Perm_t^q denote the set of injective functions from t-bit strings to t-bit strings over domains of size q. Let $\mathsf{queries}_{\mathcal{S}}(M, s, r, y) \in \mathsf{Perm}_t^{q_{\mathcal{S}}}$ denote the t-bit prefixes of the queries/answers made by $\mathcal{S}$ to its oracle during the execution of $C(M, s, r, y)$ (i.e., an oracle query $a \| a'$ by $\mathcal{S}$ is answered with $b \| a'$ in $C(M, s, r, y)$ iff $(a, b) \in \mathsf{queries}_{\mathcal{S}}(M, s, r, y)$). Letting $q_{\mathcal{R}}$ denote the number of queries $\mathcal{R}$ makes to its oracle, we define $\mathsf{queries}_{\mathcal{R}}(M, s, r, y) \in \mathsf{Perm}_t^{q_{\mathcal{R}}}$ similarly.

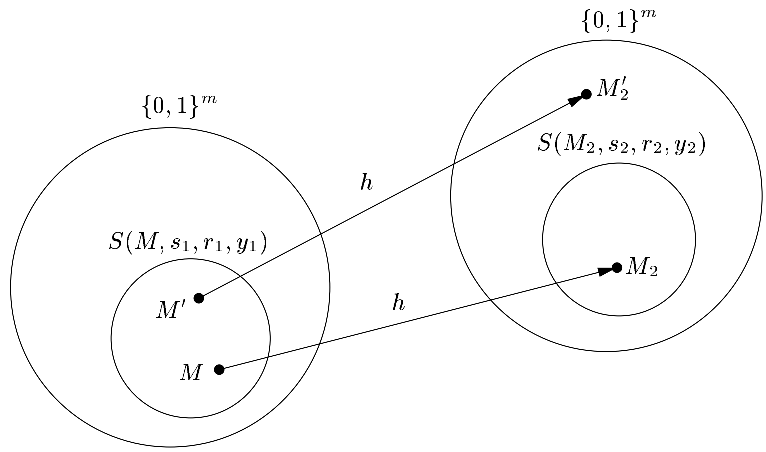

Fig. 1. $\bar{\mathcal{S}}$ uses the given scheme to generate commitments C_1 to M and C_2 to $M_2 = h(M)$ (simulating random permutations for $(\mathcal{S}, \mathcal{R})$, as needed). C_1 (resp., C_2) may be decommitted to any of the messages in the set $S(M, s_1, r_1, y_1)$ (resp., $S(h(M), s_2, r_2, y_2)$). But for a small number of sender oracle-queries, the sizes of these sets are small, and the probability that a random h maps a message from one to the other can be shown to be less than a constant.

Say $C(M, s, r, y)$ is *good for* $P \in \mathsf{Perm}_t^{qs}$ if there do not exist distinct values M', M'' along with s', s'', y', y'' such that

- $C(M', s', r, y') = C(M'', s'', r, y'') = C(M, s, r, y)$; and
- $\mathsf{queries}_{\mathcal{S}}(M', s', r, y') = \mathsf{queries}_{\mathcal{S}}(M'', s'', r, y'') = P$.

Say $C(M, s, r, y)$ is *good* if it is good for all $P \in \mathsf{Perm}_t^{qs}$. Now, if $C(M, s, r, y)$ is not good for some P, and $M', M'', s', s'', y', y''$ are the witnesses, then it is also the case that

$$\mathsf{queries}_{\mathcal{R}}(M', s', r, y') = \mathsf{queries}_{\mathcal{R}}(M'', s'', r, y'') = \mathsf{queries}_{\mathcal{R}}(M, s, r, y),$$

as queries/answers of the receiver to its oracle are explicit in the receiver's view. Furthermore, $P \cup \mathsf{queries}_{\mathcal{R}}(M, s, r, y)$ may be extended to a permutation in $\Pi_{t,n}$. Therefore, for any M, s, y and P, we have

$$\begin{aligned}
&\Pr_r[C(M, s, r, y) \text{ is not good for } P] \\
&\leq \Pr_r \begin{bmatrix} \exists \pi \in \Pi_{t,n} \text{ extending P} \cup \mathsf{queries}_{\mathcal{R}}(M, s, r, y) \\ \exists M', M'', s', s'' \text{ s.t.} \\ \langle \mathcal{S}^\pi(M'; s'), \mathcal{R}^\pi(r) \rangle = \langle \mathcal{S}^\pi(M''; s''), \mathcal{R}^\pi(r) \rangle \end{bmatrix} \leq \varepsilon_b,
\end{aligned}$$

following the binding property of $(\mathcal{S}, \mathcal{R})$. Applying the union bound over all $\binom{2^t}{qs} \cdot \prod_{i=0}^{qs-1}(2^t - i) \leq 2^{2tqs}$ permutations in Perm_t^{qs}, we obtain, for any M, s, y:

$$\Pr_r[C(M, s, r, y) \text{ is not good}] \leq 2^{2tqs} \varepsilon_b.$$

Assume $C \stackrel{\text{def}}{=} C(M, s, r, y)$ is *good*. Then there are at most 2^{2tq_S} possible messages to which C may be decommitted. This is because for any $P \in \mathsf{Perm}_t^{q_S}$, if $C(M', s', r, y') = C(M'', s'', r, y'') = C$ and $\mathsf{queries}_{\mathcal{S}}(M', s', r, y') = \mathsf{queries}_{\mathcal{S}}(M'', s'', r, y'') = P$, then it must be the case that $M' = M''$. Let $S(M, s, r, y)$ denote this set.

Viewed this way, we have:

$$\mathsf{NoBind} \leq \Pr_{\substack{h\in H\\ s_1,r_1,y_1\\ s_2,r_2,y_2}}\left[\begin{array}{c|c} \exists M' \neq M \text{ s.t.} & C(M, s_1, r_1, y_1), \\ M' \in S(M, s_1, r_1, y_1) & C(h(M), s_2, r_2, y_2) \\ h(M') \in S(h(M), s_2, r_2, y_2) & \text{are good} \end{array}\right] + 2^{2tq_S+1} \cdot \varepsilon_b,$$

where the right term represents the probability that either of $C(M, s_1, r_1, y_1)$ or $C(h(M), s_2, r_2, y_2)$ is not good. Continuing with the left term, we have:

$$\begin{aligned}
&\mathsf{LeftTerm} \\
&= \Pr_{\substack{h\in H\\ s_1,r_1,y_1\\ s_2,r_2,y_2}}\left[\begin{array}{c|c} \exists M' \in S(M, s_1, r_1, y_1) \setminus M & C(M, s_1, r_1, y_1), \\ \exists M_2' \in S(h(M), s_2, r_2, y_2) \text{ s.t.} & C(h(M), s_2, r_2, y_2) \\ h(M') = M_2' & \text{are good} \end{array}\right] \\
&= \sum_{M_2} \Pr_{\substack{h\in H\\ s_1,r_1,y_1\\ s_2,r_2,y_2}}\left[\begin{array}{c|c} \exists M' \in S(M, s_1, r_1, y_1) \setminus M & C(M, s_1, r_1, y_1), \\ \exists M_2' \in S(M_2, s_2, r_2, y_2) \text{ s.t.} & C(M_2, s_2, r_2, y_2) \\ h(M') = M_2', h(M) = M_2 & \text{are good} \end{array}\right] \\
&\leq \sum_{\substack{M_2\\ \hat{s}_1,\hat{r}_1,\hat{y}_1 \text{ s.t.}\\ C(M,\hat{s}_1,\hat{r}_1,\hat{y}_1) \text{ good},\\ \hat{s}_2,\hat{r}_2,\hat{y}_2 \text{ s.t.}\\ C(M_2,\hat{s}_2,\hat{r}_2,\hat{y}_2) \text{ good},\\ M'\in S(M,\hat{s}_1,\hat{r}_1,\hat{y}_1)\\ M_2'\in S(M_2,\hat{s}_2,\hat{r}_2,\hat{y}_2)}} \Pr_{\substack{h\in H\\ s_1,r_1,y_1\\ s_2,r_2,y_2}}\left[\begin{array}{c|c} s_1 = \hat{s}_1, r_1 = \hat{r}_1, y_1 = \hat{y}_1 & C(M, s_1, r_1, y_1), \\ s_2 = \hat{s}_2, r_2 = \hat{r}_2, y_2 = \hat{y}_2 & C(M_2, s_2, r_2, y_2) \\ h(M') = M_2', h(M) = M_2 & \text{are good} \end{array}\right] \\
&= 2^{-2m} \cdot \sum_{\substack{M_2\\ \hat{s}_1,\hat{r}_1,\hat{y}_1 \text{ s.t.}\\ C(M,\hat{s}_1,\hat{r}_1,\hat{y}_1) \text{ good},\\ \hat{s}_2,\hat{r}_2,\hat{y}_2 \text{ s.t.}\\ C(M_2,\hat{s}_2,\hat{r}_2,\hat{y}_2) \text{ good},\\ M'\in S(M,\hat{s}_1,\hat{r}_1,\hat{y}_1)\\ M_2'\in S(M_2,\hat{s}_2,\hat{r}_2,\hat{y}_2)}} \Pr_{\substack{s_1,r_1,y_1\\ s_2,r_2,y_2}}\left[\begin{array}{c|c} s_1 = \hat{s}_1, r_1 = \hat{r}_1, y_1 = \hat{y}_1 & C(M, s_1, r_1, y_1), \\ s_2 = \hat{s}_2, r_2 = \hat{r}_2, y_2 = \hat{y}_2 & C(M_2, s_2, r_2, y_2) \\ & \text{are good} \end{array}\right],
\end{aligned}$$

using the pairwise independence of H. Since $|S(M, s, r, y)| \leq 2^{2tq_S}$, we obtain

$$\mathsf{LeftTerm} \leq 2^{-2m} \cdot 2^m \cdot 2^{4tq_S} \leq 2^{4tq_S - m}.$$

Putting it together, we have

$$\mathsf{NoBind} \leq 2^{4tq_S - m} + 2^{2tq_S+1} \cdot \varepsilon_b \leq 2^{4tq_S} \cdot (2^{-m} + 2\varepsilon_b) \leq 1/4,$$

where the last inequality is due to the assumption that $q_{\mathcal{S}} \leq (m - 2 - \log(1 + 2^{m+1} \cdot \varepsilon_b))/4t$. □

This completes the proof of the theorem. ■

Acknowledgments. We thank Virgil Gligor and Chiu-Yuen Koo.

References

1. M. Blum. Coin Flipping by Phone. In *24th IEEE Computer Conference (CompCon)*, pp. 133–137, 1983. (See also *SIGACT News*, vol. 15(1), 1983.)
2. M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. In *SIAM J. Computing*, vol. 13(4), pp. 850–864, 1984.
3. J. Carter and M. Wegman. Universal Classes of Hash Functions. In *Journal of Computer and System Sciences*, vol. 18, pp. 143–154, 1979.
4. R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan. Bounds on the Efficiency of Generic Cryptographic Constructions. *SIAM J. Computing*, to appear.
5. O. Goldreich. Foundations of Cryptography, vol. 1: Basic Tools. Cambridge University Press, 2001.
6. O. Goldreich, S. Goldwasser, and S. Micali. On the Cryptographic Applications of Random Functions. In *Advances in Cryptology — CRYPTO 84*, LNCS vol. 263, Springer-Verlag, pp. 276–288, 1985.
7. O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. In *J. ACM*, vol. 33(4), pp. 792–807, 1986.
8. O. Goldreich and L. Levin. Hard-Core Predicates for any One-Way Function. In *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 25–32, 1989.
9. J. Håstad, R. Impagliazzo, L. Levin, and M. Luby. A Pseudorandom Generator From any One-Way Function. In *SIAM J. Computing*, vol. 28(4), pp. 1364–1396, 1999.
10. R. Impagliazzo and S. Luby. One-Way Functions are Essential for Complexity-Based Cryptography. In *30th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 230–235, 1989.
11. R. Impagliazzo and S. Rudich. Limits on the Provable Consequences of One-Way Permutations. In *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 44–61,1989.
12. J.H. Kim, D.R. Simon, and P. Tetali. Limits on the Efficiency of One-Way Permutation-Based Hash Functions. In *40th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 535–542, 1999.
13. M. Naor. Bit Commitment Using Pseudorandomness. In *J. Cryptology*, vol. 4(2), pp. 151–158, 1991.
14. M. Naor and M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. In *21st ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 33–43, 1989.
15. S. Rudich. Limits on the Provable Consequences of One-Way Functions. Ph.D. thesis, University of California at Berkeley, 1988.
16. J. Rompel. One-Way Functions are Necessary and Sufficient for Secure Signatures. In *22nd ACM Symposium on Theory of Computing (STOC)*, ACM, pp. 387–394, 1990.
17. O. Reingold, L. Trevisan, and S. Vadhan. Notions of Reducibility Between Cryptographic Primitives. In *1st Theory of Cryptography Conference*, LNCS vol. 2951, Springer-Verlag, pp. 1–20, 2004.
18. A.C.-C. Yao. Theory and Application of Trapdoor Functions. In *23rd IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 80–91, 1982.
19. A.C.-C. Yao. How to Generate and Exchange secrets. In *27th IEEE Symposium on Foundations of Computer Science (FOCS)*, IEEE, pp. 162–167, 1986.