Two-Round Zero Knowledge and Proof Auditors

IPAM Workshop

January 10, 2002

Cynthia Dwork
Microsoft Research – SVC

Larry Stockmeyer
IBM Almaden

Vocabulary: NP, Witness

$L \in$ NP if there is a nondeterministic polynomial time machine M that accepts exactly the strings $x \in L$.

"Languages with short proofs of membership"

3-colorable graphs, Hamiltonian graphs, satisfiable Boolean formulas . . .

P $\neq$ NP $\Rightarrow$ sometimes:
checking proof easier than finding it

Given M accepting L, a witness w for $x \in L$ is an accepting computation path of M on x.

Vocabulary:
Interactive Proofs [B, GMR]

Interactive Proof System (IP) for L: $P(x)$ and ppt $V(x)$ interact. V accepts if $x \in L$. If $x \notin L$, then a cheating P^* can convince V to (erroneously) accept with only negligible prob.

IP = PSPACE [S]

Proofs vs. Arguments: Power of Prover

Proofs: arbitrary

Arguments: probabilistic poly time

Our work: Proofs for $L \in$ NP; good prover only needs ppt + a witness.

Vocabulary:
Zero Knowledge [GMR]

V^* learns only that $x \in L$, but not "why."
Formalized by a ppt simulator that, without a witness, constructs simulated conversations that can'be be distinguished from real ones in pptime.

Intuition: "witness protection"

$$\text{real}(x, w) \approx_{\text{poly}} \text{simulated}(x)$$

$\Rightarrow$ real conversations protect w well

IP = ZK [IY]

Two-Round Zero Knowledge $\forall L \in$ NP
Prover is *resource-bounded* during protocol

Time

Implicit in argument systems

Explicit in DNS, DN

Size/Advice

Implicit *e.g.* in FS – no huge sorted table of trapdoor commitment schemes and associated trapdoors

State of the Art: Rounds

Feige and Shamir: 4 rounds suffice.

Dwork and Naor: Using timing and moderately hard functions, 3 rounds suffice.

One round (one single $P \longrightarrow V$ message) is insufficient in standard and time-bounded prover models (cf. NIZKs)

Vocabulary: Witness-Indistinguishability, Zaps

Witness Indistinguishability (WI): $\forall$ ptime V^*, $\forall x \in L$, $\forall w_1, w_2 \in w(x)$, $\forall$ aux. inputs z to V^*, $(P, V^*)(x, w_1, z) \approx_{\text{poly}} (P, V^*)(x, w_2, z)$.

Zap: 2-round public-coin WI IP in which the first round $(V \longrightarrow P : \rho)$ is a string of (supposedly) random bits (which can be fixed once and for all).

Flawed Prototype: $L \in NP$

"Icky-Poly" f requires time $q(|x|)$ on x.

Committing PPKC Generator GK
Common: 1^k, t of size k^{c_1}.

$V \longrightarrow P$: $x \in_R \{0,1\}^k$, ρ

$P \longrightarrow V$:

$E \in_R \mathsf{GK}(1^k)$, $x^* \in_R \{0,1\}^k$, $\beta \in_R E(x^*)$;
Send: E, zap("$t \in L \vee \beta \in E(f(x))$"), β

Timing Constraints: Response time $<< q(k)$.

V accepts iff P's reply is timely and the zap is acceptable.

Intuition for Zero-Knowledge

$V \longrightarrow P$: $x \in_R \{0,1\}^k$, ρ

$P \longrightarrow V$:

$E \in_R \mathsf{GK}(1^k)$, $x^* \in_R \{0,1\}^k$, $\beta \in_R E(x^*)$;

Send: E, zap("$t \in L \vee \beta \in E(f(x))$"), β

Zap("$t \in L \vee \beta \in E(f(x))$"), where $\beta \in E(x^*)$ and witness is for $t \in L$, looks like:

Zap("$t \in L \vee \beta \in E(f(x))$"), where $\beta \in E(f(x))$ and witness is for $t \in L$, looks like:

Zap("$t \in L \vee \beta \in E(f(x))$"), where $\beta \in E(f(x))$ and witness is for $\beta \in E(f(x))$.

Intuition(?!) for Soundness:

f requires time $q(|x|)$ on x.

$V \longrightarrow P$: $x \in_R \{0,1\}^k$, ρ

$P \longrightarrow V$:

$E \in_R \mathsf{GK}(1^k)$, $x^* \in_R \{0,1\}^k$, $\beta \in_R E(x^*)$;

Send: E, zap("$t \in L \vee \beta \in E(f(x))$"), β

Timing Constraints: Response time $<< q(k)$.

Since P^* can't compute $f(x)$ in time $< q(|x|)$, it can't find $E, \beta \in E(f(x))$. Thus, if the zap is accepted, the first clause, "$t \in L$," must be true.

Difficulties

f requires time $q(|x|)$ on x.

$V \longrightarrow P : x \in_R \{0,1\}^k, \rho$

$P \longrightarrow V :$

$E \in_R \mathsf{GK}(1^k)$, $x^* \in_R \{0,1\}^k$, $\beta \in_R E(x^*)$;

Send: E, $\mathsf{zap}(\text{“}t \in L \vee \beta \in E(f(x))\text{”})$, β

Timing Constraints: Response time $<< q(k)$.

• Current techniques: zap goes through Cook-Levin theorem $\Rightarrow |\mathsf{zap}| > q(|x|)$. For now, assume magically short zap.

• Don't know how to prove intuition. Maybe finding E, $\beta \in E(f(x))$ easy? Long standing open problem in cryptography!

Key Insight

f requires time $q(|x|)$ on x.

$P \longrightarrow V$:

$E \in_R \mathsf{GK}(1^k)$, $x^* \in_R \{0,1\}^k$, $\beta \in_R E(x^*)$;

Send: E, zap("$t \in L \vee \beta \in E(f(x))$"), β

Suppose perfect cheating: $\exists t \notin L \, \forall x$ P^* quickly produces E, β, and acceptable zap.

Then $\forall x$ there is a short proof of $y = f(x)$:

computation of P^*

decryption key D for E

computation $D(\beta)$

(miraculously short) zap

(don't need V's computation: it will accept)

The Point

If statements "$y = f(x)$" have no short proofs, that is, checking is no easier than computing (and assuming miraculously short zaps), then perfect cheating is impossible!

Remark: Since issue is checking rather than computing, P^* can even be nondeterministic.

Proof Auditor formalizes "proof".

Proof Auditor for f (Approximate Definition)

$\mathsf{Aud}(x, y, z) \in \{0, 1\}$

$(\exists z\ \mathsf{Aud}(\mathsf{x}, \mathsf{y}, \mathsf{z}) = 1) \Leftrightarrow (\mathsf{y} = \mathsf{f}(\mathsf{x}))$

Recall Difficulties

• Current techniques: zap goes through Cook's theorem $\Rightarrow |\text{zap}| > q(|x|)$.

• Current techniques: don't know how to prove intuition. Maybe finding E, $\beta \in E(f(x))$ easy? Long standing open problem in cryptography!

Proof auditor addresses unproven intuition.

Length of zap (with clause $\beta \in E(f(x))$) addressed through a mixture of prover preprocessing and limited malleability of encryptions.

Strategy for Resource-Bounded Provers

Our protocol involves a particular function f.

Tie existence of a cheating prover using specified resources and succeeding with a specified probability to existence of a proof auditor with corresponding resource and correctness bounds.

Assume (for time) or prove (for advice) that no auditor for f exists with these bounds.

Executive Summary of Results

Definition of proof auditor.

Assume: $\exists f$ with no fast auditor AND completely malleable sem. sec. cryptosystem AND standard crypto
Conclude: $NP \subseteq TZKIP(2)$

Assume: $\exists f \in \mathcal{LIN}$ with no fast auditor AND standard crypto (including QRA)
Conclude: $NP \subseteq TZKIP(2)$

Fact: $\exists f \in \mathcal{LIN}$ with no *small* auditor.
Assume: standard crypto (including QRA)
Conclude: $NP \subseteq SZKIP(2)$

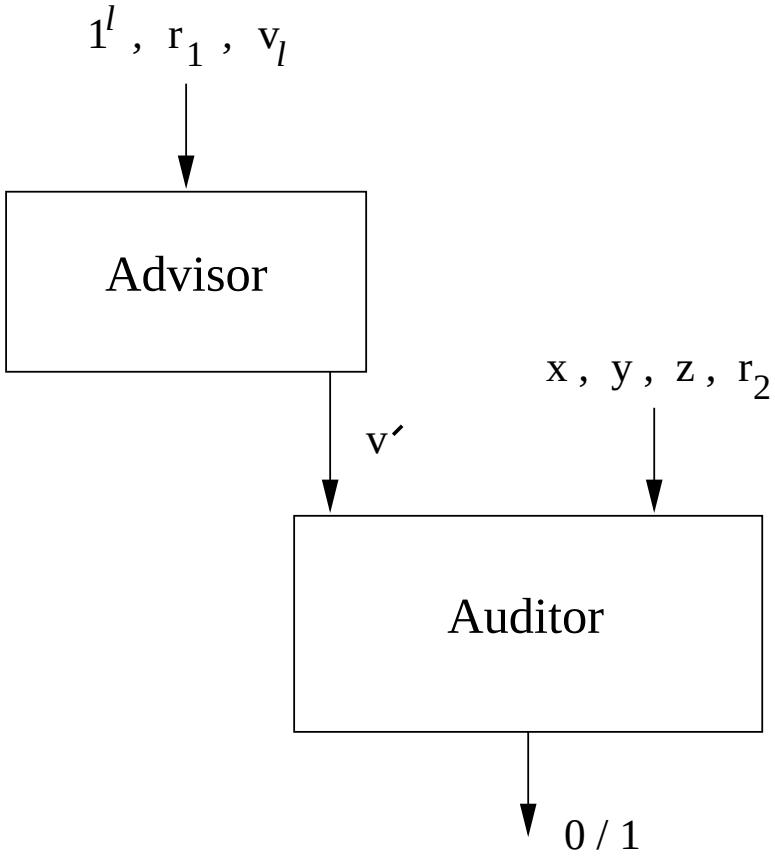

Auditor $\mathcal{A}$: Components $adv(\mathcal{A})$ and $aud(\mathcal{A})$

i.o. Proof Auditors for f

- Sequence of *advice strings* v_ℓ
- Deterministic *advisor* $adv(\mathcal{A})$

  $adv(\mathcal{A})(1^\ell, r_1, v_\ell) = v'$
- Deterministic *auditor* $aud(\mathcal{A})$

  $aud(\mathcal{A})(x, y, z, r_2, v') \in \{0, 1\}$

  z "proof" that $y = f(x)$, $\ell = |x|$
- Correctness bound $\chi(\ell)$

i.o. ℓ: $\Pr_{x,r_1,r_2}[E] \geq \chi(\ell)$; $E =$

$\forall y\big((\exists z : aud(\mathcal{A})(x, y, z, r_2, v') = 1) \Leftrightarrow y = f(x)\big)$

where $v' = adv(\mathcal{A})(1^\ell, r_1, v_\ell)$

"infinitely often" (over ℓ)

"input-output" (checking "$y = f(x)$")

The Goldwasser-Micali PPKC

$Q_N(z) = 1$ iff z is a quadratic residue mod N.

Public Key: N where $N = pq$ is a product of two certified primes congruent to 3 mod 4. Note: $Q_N(-1) = 0$.

Private Key: p, q and their certificates.

Encryption: Bit-by-bit:

Encryptions of 0: $r^2 \bmod N$, $r \in_R \mathbf{Z}_N^*$

Encryptions of 1: $-r^2 \bmod N$, $r \in_R \mathbf{Z}_N^*$.

Decryption of Ciphertext c:

$1 - Q_N(c) = 1 - Q_p(c) \cdot Q_q(c)$.

XOR-malleability: given two 1-bit ciphertexts α and β, $\alpha\beta \bmod N$ is the encryption of the exclusive-or of the decryptions of α and of β.

Protocol Preliminaries

GK(1^k): GM w/ certified $p, q \equiv 3 \bmod 4$

$L_K = \{\text{valid encryption keys}\} \in NP$

Easy to check E/D pairs

$\mathcal{LIN}$ = linear fns f, $\forall \ell\, f_\ell : \text{GF}(2)^\ell \to \text{GF}(2)^\ell$.

Fix $f \in \mathcal{LIN}$, $\oplus$-circuit size $s(\ell) = \ell^2$.

Let g be a function such that $\forall E \in L_K$:

$\alpha \in E(x) \Rightarrow g(E, \alpha) \in E(f(x))$

g computable in time $O(s(\ell) \cdot \text{poly}(k))$

Since GM encryption is "$\oplus$-malleable," define $g(E, \alpha)$, where $\alpha \in E(x)$, $|x| = \ell$, by "pushing" α through a fixed $\oplus$-circuit that computes f_ℓ.

Simplified Protocol (not sound)

Common: 1^k, $t \in L \cap \{0,1\}^{k^{c_1}}$(?), $\ell = k^d$.

$P(1^k)$ Precomputation ($O(\ell k)$ bits, ind. of t):

$E \in_R \mathsf{GK}(1^k)$; $x^* \in_R \{0,1\}^\ell$;

$\alpha \in_R E(x^*)$, $\beta := g(E, \alpha)$

$V \longrightarrow P$: $x \in_R \{0,1\}^\ell$ and ρ

$P \longrightarrow V$: E, α, β, zap$(t \in L \vee \alpha \in E(x))$

Accept iff zap accepted and $\beta = g(E, \alpha)$.

If V accepts, then $t \in L \vee \beta \in E(f(x))$.

The miraculously short zap!

For timing-based version: (1) zap performed as ℓ zaps (one/bit of α) and (2) timeliness test: $\ell k^c << \ell^2$

Full Protocol: Advice-Bounded Prover

Common: 1^k, $t \in L \cap \{0,1\}^k(?)$, $\ell = k^d$.

P Precomputation ($O(\ell k)$ bits):

$E_1, E_2 \in_R \mathsf{GK}(1^k)$; $x^* \in_R \{0,1\}^\ell$;

$\alpha_i \in_R E_i(x^*)$, for $i = 1,2$;

$\beta_i := g(E_i, \alpha_i)$, for $i = 1,2$

$V \longrightarrow P$: $x \in_R \{0,1\}^\ell$ and $\rho = \rho_1 \cdot \rho_2$

$P \longrightarrow V$: $E_1, E_2, \alpha_1, \alpha_2, \beta_1, \beta_2$

zap$(E_1 \in L_K \vee E_2 \in L_K)$ using ρ_1

zap$(t \in L \vee \alpha_1 \in E_1(x))$ using ρ_2

zap$(t \in L \vee \alpha_2 \in E_2(x))$ using ρ_2

Accept iff all zaps accepted, and $\beta_i = g(E_i, \alpha_i)$ for $i = 1,2$.

Cheating Provers and Auditors

Cheating P^* causing V to accept ∞ $t_\ell \notin L$ with prob. $> \chi(\ell)$ yields i.o. proof auditor for f.

$$\mathit{pre}(P^*) \leftrightarrow \mathit{adv}(\mathcal{A}),\ \mathit{pro}(P^*) \leftrightarrow \mathit{aud}(\mathcal{A})$$

Advice, time, and correctness bounds of the auditor are related to those of P^*.

Main Theorem

length of "theorem" + "witness"	$O(k^{c_1})$
bit decryption time	$O(k^{c_2})$
verification of $(D, w(E, D))$	$O(k^{c_3})$
length of $(D, w(E, D))$	$O(k^{c_4})$

$\exists P^*$ and ∞K : $\forall k \in K\, \exists t_k, w_k$ with $t_k \notin L$ s.t. $\Pr[(P^*, V)\langle t_k, w_k\rangle = 1] \geq \varepsilon(k)$ $\Rightarrow$ $\exists$ i.o. proof auditor $\mathcal{A}$ for f, where $k = \lceil \ell^{1/d} \rceil$, with bounds:

$$\begin{aligned}
\text{advisor advice} &= O(A_1(k) + k^{c_1}) \\
\text{advisor time} &= O(T_1(k) + k^{c_1}) \\
\text{auditor advice} &= O(A_2(k) + k^{c_1}) \\
\text{auditor time} &= O(a\ell k^c + \ell k^{c_2} + k^{c_3}) \\
\text{proof length} &= O(k^{c_4}) \\
\text{correctness} &= \varepsilon(k) - 2 \cdot 4^{-k}.
\end{aligned}$$

Advice-Bounded Provers

$\mathcal{LIN}$ = linear fns f
$\forall \ell \, \forall f \in \mathcal{LIN}_\ell \quad f : \mathrm{GF}(2)^\ell \to \mathrm{GF}(2)^\ell$
each $f \in \mathcal{LIN}_\ell$ represented by $\ell \times \ell$ bit matrix

Theorem (roughly):
If $\delta, \varepsilon > 0$ are sufficiently small constants, then most $(1 - 2^{-\ell^2(1-\delta-3\varepsilon-6\ell^{-1})}$ fraction) functions in $\mathcal{LIN}_\ell$ have no proof auditor with auditor advice bound $\delta\ell^2$ and correctness bound $2^{-\varepsilon\ell}$.

- Advice ℓ^2: perfect auditor for any $f \in \mathcal{LIN}_\ell$
- Correctness bound $2^{-\ell}$ is trivial:
  check $x = y = 0^\ell$

Technical Issues

For $\delta > 0$ sufficiently small, most functions in $\mathcal{LIN}_\ell$ have no proof auditor with auditor advice bound $\delta\ell^2$ and correctness bound $\chi = 2^{-\delta\ell+1}$.

With $\delta\ell^2$ bits can represent only a vanishing fraction of functions, as $|\mathcal{LIN}_\ell| = 2^{\ell^2}$. But:

- The auditor needn't be correct for *all* x; sufficient to be correct on a set X_f containing only $2^{-\varepsilon\ell}$ fraction of x's.
- Even on $x \in X_f$, auditor needn't be correct with probability 1.
- Possibly different X_f for each f.

Some Ideas from the Proof

$|\mathrm{GF}(2)^\ell| = 2^\ell$.

Fix an advice string v'. Parameters for the discussion:

$$\chi = 2^{-\delta\ell+1} \qquad |v'| = \delta\ell^2$$

$$c = \ell - \lceil \delta\ell \rceil \qquad p = 2^{c-\ell}$$

So $2p = 2^{-\lceil \delta\ell \rceil+1} \leq 2^{-\delta\ell+1} = \chi$ and we can forget about one of the parameters (χ).

Auditor advice v' *handles* f if
$\Pr_{x,r}[\forall y((\exists z \mathcal{U}(x,y,z,r,v') = 1) \Leftrightarrow y = f(x))] \geq 2p$

If v' handles f then $\exists$ p-fraction $X_f \subset \mathrm{GF}(2)^\ell$ on which $\mathcal{U}$ does well (over choice of r): $\forall x \in X_f$
$\Pr_r[\forall y((\exists z \mathcal{U}(x,y,z,r,v') = 1) \Leftrightarrow y = f(x))] \geq p$

y is possible for x:
$\exists f$ handled by v' s.t. $x \in X_f$ and $y = f(x)$
Few possible y's means v' handles few f's.

Claim: $\forall x \exists \leq p^{-1}$ possible values y for x.

Suppose $f_1(x) \neq f_2(x)$. If

$$\forall y((\exists z \mathcal{U}(x, y, z, r, v') = 1) \Rightarrow y = f_1(x))$$

then

$$\forall y((\exists z \mathcal{U}(x, y, z, r, v') = 1) \not\Rightarrow y = f_2(x))$$

Finish by pigeonhole argument over r's, using (1) if y is possible for x then for some f handled by v', $x \in X_f$ and $y = f(x)$ and (2) $\forall x \in X_f$:

$\Pr_r[\forall y((\exists z \mathcal{U}(x, y, z, r, v') = 1) \Leftrightarrow y = f(x))] \geq p$.

Each X_f contains a linearly independent set S of at least $c \sim \log |X_f|$ values.

By linearity, f completely determined by its value on ℓ elements. If all f's handled by v' had the same X_f, then there would be only $\ell - c$ values with which to play. Thus, the number of functions handled by v' would be bounded by:

$$(p^{-1})^c (2^\ell)^{\ell - c}$$

- p^{-1}: upper bound on possible y's $\forall x \in S$
- c: lower bound on $|S|$
- 2^ℓ: # choices for unconstrained values
- $\ell - c$: upper bound on number of linearly independent values to play with (i.e., not in S)

$$\mathcal{I} = \{I \subset \mathrm{GF}(2)^\ell \mid \mathcal{I} \text{ is l.i. and } |\mathcal{I}| = c\}$$

$\exists \mathcal{J} \subset \mathcal{I}$ such that $|\mathcal{J}| \leq 2^{\ell c - (c-1)^2 + \ell}$ and for every f handled by v': $\exists I \in \mathcal{J}\, \forall x \in I\ x \in X_f$.

That is, there is a small collection of c-sized sets of linearly independent values such that, if v' handles f, then there is an I in the collection such that $\mathcal{U}$ does well on all $x \in I$ for f with advice v'.

Something Surprising
(Joint with Li Zhang)

Common Input $N = pq$; ζ, where $|\zeta| = |N|^2$

Prover Input $\phi(N)$

Prover Precomputation: $z = \zeta \bmod \phi(N)$

$P \longrightarrow V$:

Choose: $x \in_R \mathbf{Z}_N^*$.

Send: $x, x^z, x^{z^2}, x^{z^3}, \ldots$ (all mod N)

V checks that each element is obtained by raising its predecessor to ζ, modulo N.

Hope (only!): without $\phi(N)$ there is no short ($O(|N|)$) advice string a ptime cheating prover can find enabling it to carry out the protocol.

Open Questions

One-Round Something?

Physical model for bounds on *pro*(P^*) size?

Lower bounds on auditor time?

Total malleability?

Length-preserving functions computable in Θ(icky-poly) time, with substantially more $\oplus$ operations than ANDs/ORs?