

Garbled Circuits


Andrew Yao

- Turing award winner
- PhD in physics and CS (second in 3 years)
- Taught at MIT, Stanford, Berkeley, Princeton
- Now runs "Yao Class" at Tsinghua
- Many influential early papers
 - ↳ Results on OWFs, PRGs, now standard
 - ↳ Too many other cool results to list

Today, we discuss one: "Yao's Millionaire's Problem"

↳ Introduced in talk in 1986

- Two millionaires sit on yacht in French Polynesia




- Intuitively, both learn $f(x, y)$, "but nothing else"
- Using this idea, can jointly compute $f(x, y)$ for any f s.t. nothing leaks except $f(x, y)$.
 - Compute RSA modulus s.t. neither knows factors
 - $x = DB$, $y =$ indices of rows $f(x, y) =$ sum of selected rows

Circuits

We need to represent f as a boolean circuit. (Wlog!)

→ AND, OR, NOT gates

→ Useful b/c evaluating a circuit is "data oblivious"



Yao's protocol does not hide f .

⇒ Wlog. let f be deterministic (when is this wlog?)

Yao's "Garbled Circuits"

Let us define what it means for MPC to be secure.

→ We are going to stick w/ "honest but curious" adv for now
↳ Later will discuss Full security.

Two parties $\{1, 2\}$ w/ inputs $\{x, y\}$ s.t. $|M| = |Y|$

Def'n Protocol π securely computes $f(x, y)$ is \exists ppt simulators S_1, S_2 s.t. $\forall x, y \in \{0, 1\}^n$

$$\{S_1(x, f(x, y))\} \stackrel{S(x, y)}{\approx} \{\text{view}_1^\pi(x, y)\} \stackrel{\text{output}(x, y)}{\approx} \{f(x, y)\}$$

(Same for party 2.)

Informally: Each party can simulate its view of the protocol execution given only

- 1) Its private input
 - 2) The value $f(x, y)$.
- } "What leaks"

$\Rightarrow \pi$ is secure if it leaks $f(x, y)$ and nothing else!

This is similar to ZK def'n's, except that we allow a little extra leakage.

Def'n π is correct if for all $x, y \in \{0, 1\}^n$


$$\{\text{output}^\pi(x, y)\} \stackrel{S}{\approx} \{f(x, y)\}.$$

[Note: Can also have players 1 and 2 receive distinct outputs that the other player doesn't learn, ... also not too hard.]

Garbled Circuits

High-level idea:

- * P_1 generates a "garbled" version of ckt $C(x,y)$
↳ hides all internal wire values
- * P_2 uses O.T. to fetch enough info to evaluate $C(x,y)$ without leaking y to P_1 ("Garbled input wires")
- * P_2 evaluates $C(x,y)$, sends result to P_1 .



First, let's see what P_1 does to garble C .

We need a semantically secure enc scheme s.t.

- 1) hard to generate "valid" cts w/o knowing key
- 2) given key, easy to check whether ct is valid.


Using PRF $F: \underbrace{\{0,1\}^n}_{\text{key}} \times \underbrace{\{0,1\}^n}_{\text{input}} \rightarrow \underbrace{\{0,1\}^{2n}}_{\text{output}}$

$$E(k, m) := \left\{ \begin{array}{l} r \parallel \text{PRF}(r, m) \\ (r, m \parallel 0^n \oplus F_k(r)) \end{array} \right.$$

$$D(k, ct = (r, c)) := \begin{cases} \text{if last } n \text{ bits of } F_k(r) \oplus c \neq 0: \text{ fail} \\ \text{else output first } n \text{ bits of } F_k(r) \oplus c \end{cases}$$

Garbled CKts

For each gate in the circuit we use 6 keys



We want the property that: given K_L^a, K_R^b , can compute only $K_{out}^{(a \text{ NAND } b)}$

a	b	$a \text{ NAND } b$
0	0	1
0	1	1
1	0	1
1	1	0

encrypt

$$C_{00} = E(K_L^0, E(K_R^0, K_{out}^1))$$

$$C_{01} = E(K_L^0, E(K_R^1, K_{out}^1))$$

$$C_{10} = E(K_L^1, E(K_R^0, K_{out}^1))$$

$$C_{11} = E(K_L^1, E(K_R^1, K_{out}^0))$$


shuffle

C_{10}
C_{11}
C_{00}
C_{01}

 } "Garbled gate"

Garbled Ckts

* To garble C , choose pair of keys for each wire in ckt ("0" key and "1" key)... some subtlety...



* Output gate has encryptions of 0/1 \rightarrow not keys in encrypted truth table

Claim Let $x \in \{0,1\}^n$, $y \in \{0,1\}^m$ be inputs
Claim Given a set of keys

$$(k_{1,1}^{x_1}, \dots, k_{1,n}^{x_n}), (k_{2,1}^{y_1}, \dots, k_{2,m}^{y_m})$$

corresponding to input wires of ckt, can compute $C(x,y)$.

Idea: By induction on gate index in ckt:

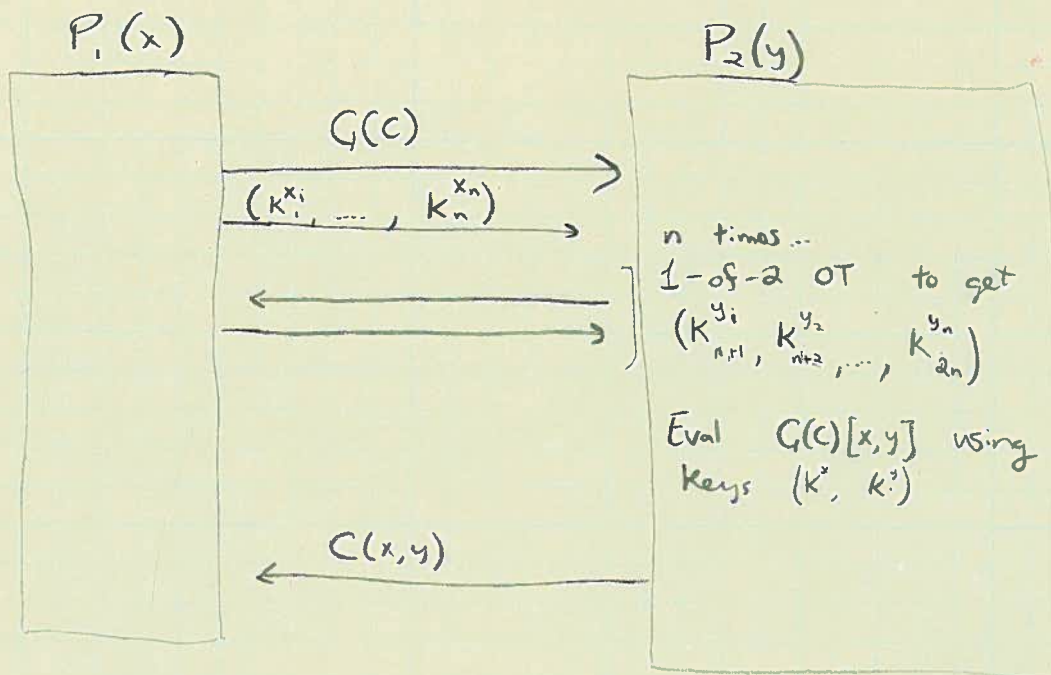
Base case: Given key to output gate, can recover output value.

Induction: Given keys (k_L^x, k_R^y) to gate i , can decrypt ct in truth table corresponding to gate $i-1$.

Now have keys for all subsequent gates
 \Rightarrow Can compute $C(x)$ by induction hypothesis.

Efficient

Garbled Cts for MPC



Correctness: If P_1, P_2 honest, both parties learn $C(x, y)$.

Security Pf Idea:

P_1 is dishonest: need Sim s.t. $\text{Sim}(x, f(x, y)) \stackrel{c}{\approx} \{\text{view}_{P_1}^*(x, y)\}$

- Simulate OTs using OT simulator
- Can output $f(x, y)$ as given to Sim

P_2 is dishonest:

- Can generate a bogus GC that looks like the real one
- Bogus GC outputs $f(x, y)$ on all inputs
- ↳ Generate bogus GC, then use hybrid arg to show its comp indist to real
- Simulation of OT step comes from OT protocol.

Malicious Security

Q: Why are we not done?

A: Many problems. * E.g. P_1 sends garbling of ckt $\hat{C} \neq C$. Output of \hat{C} might leak all of P_2 's secret input!

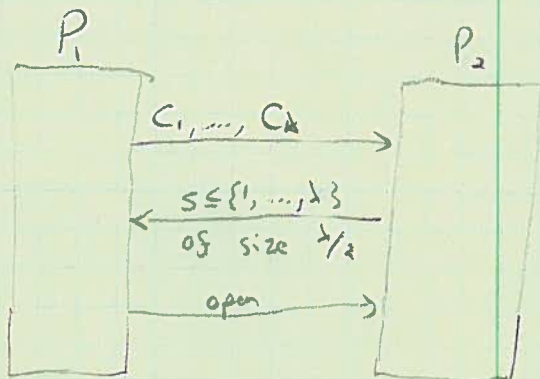
* OT protocol might not be fully secure either!

Proving security against fully malicious advs is very subtle.... "Cryptographers rarely sleep well"

We show one piece of a full maliciously secure GC.

"Cut-and-Choose"

- Very simple (but surprising!) technique
- P_1 creates λ GCs
- P_2 asks P_1 to "open" $\lambda/2$ of them
- P_1 sends all labels for subset S .
- P_1 checks all opened ckts are good
- P_1 is convinced that a majority of the remaining ckts are good.



Soundness

- If cheating, P_1 wins, $\geq \frac{\lambda}{4}$ ckts. must be bad.
- Then in step 2, P_2 opens a bad ckt w.p.

$$\begin{aligned} \Pr[\text{all opened good}] &= \Pr[C_1 \text{ good} \cap C_2 \text{ good} \cap \dots \cap C_\lambda \text{ good}] \\ &= \Pr[C_1 \text{ good}]^{\lambda/2} \\ &= \left(\frac{3}{4}\right)^{\lambda/2} \leftarrow \text{negligible in } \lambda \end{aligned}$$

\Rightarrow Cut-and-choose is insufficient... need to worry about labels too!

Optimizations

Many tricks to reduce size of GCs.

Examples: "Free XOR" [KS08]

↳ No need to store garbled XOR gates

"Random Permutation instead of DDF [Bellare et al 2013]

↳ Faster (?) on AES hw

"Point and Permute" [Malkhi et al 2004]

↳ add extra info to labels to avoid 4 decryption attempts per gate

⋮

Many other nice tricks

⇒ JustGarble reports ~ 25 cycles/gate to evaluate
 ~ 250 to garble
64 bytes/gate ckt size

→ Practical for smallish ckts in HBC setting
↳ getting closer to practical in malicious setting too
(see David's paper)