

Probabilistic Encryption^a

- A deterministic cryptosystem can be broken if the plaintext has a distribution that favors the “easy” cases.
- The ability to forge signatures on even a vanishingly small fraction of strings of some length is a security weakness if those strings were the probable ones!
- A scheme may also “leak” *partial* information.
 - Parity of the plaintext, e.g.
- The first solution to the problems of skewed distribution and partial information was based on the QRA.

^aGoldwasser & Micali (1982). This paper “laid the framework for modern cryptography” (2013).

Shafi Goldwasser^a (1958–)



^aTuring Award (2013).

Silvio Micali^a (1954–)



^aTuring Award (2013).

Goldwasser and Micali



A Useful Lemma

Lemma 82 *Let $n = pq$ be a product of two distinct primes. Then a number $y \in Z_n^*$ is a quadratic residue modulo n if and only if $(y | p) = (y | q) = 1$.*

- The “only if” part:
 - Let x be a solution to $x^2 = y \pmod{pq}$.
 - Then $x^2 = y \pmod{p}$ and $x^2 = y \pmod{q}$ also hold.
 - Hence y is a quadratic modulo p and a quadratic residue modulo q .

The Proof (concluded)

- The “if” part:
 - Let $a_1^2 = y \pmod p$ and $a_2^2 = y \pmod q$.
 - Solve

$$x = a_1 \pmod p,$$

$$x = a_2 \pmod q,$$

for x with the Chinese remainder theorem (p. 486).

- As $x^2 = y \pmod p$, $x^2 = y \pmod q$, and $\gcd(p, q) = 1$, we must have $x^2 = y \pmod{pq}$.

The Jacobi Symbol and Quadratic Residuacity Test

- The Legendre symbol can be used as a test for quadratic residuacity by Lemma 69 (p. 554).
- Lemma 82 (p. 680) says this is *not* the case with the Jacobi symbol in general.
- Suppose $n = pq$ is a product of two distinct primes.
- A number $y \in Z_n^*$ with Jacobi symbol $(y | pq) = 1$ is a quadratic *nonresidue* modulo n when

$$(y | p) = (y | q) = -1,$$

because $(y | pq) = (y | p)(y | q)$.

The Setup

- Bob publishes $n = pq$, a product of two distinct primes, and a quadratic nonresidue y with Jacobi symbol 1.
- Bob keeps secret the factorization of n .
- Alice wants to send bit string $b_1b_2 \cdots b_k$ to Bob.
- Alice encrypts the bits by choosing a random quadratic residue modulo n if b_i is 1 and a random quadratic nonresidue (with Jacobi symbol 1) otherwise.
- So a sequence of residues and nonresidues are sent.
- Knowing the factorization of n , Bob can efficiently test quadratic residuacity and thus read the message.

The Protocol for Alice

- 1: **for** $i = 1, 2, \dots, k$ **do**
- 2: Pick $r \in Z_n^*$ randomly;
- 3: **if** $b_i = 1$ **then**
- 4: Send $r^2 \bmod n$; {Jacobi symbol is 1.}
- 5: **else**
- 6: Send $r^2 y \bmod n$; {Jacobi symbol is still 1.}
- 7: **end if**
- 8: **end for**

The Protocol for Bob

```
1: for  $i = 1, 2, \dots, k$  do  
2:   Receive  $r$ ;  
3:   if  $(r | p) = 1$  and  $(r | q) = 1$  then  
4:      $b_i := 1$ ;  
5:   else  
6:      $b_i := 0$ ;  
7:   end if  
8: end for
```

Semantic Security

- This encryption scheme is probabilistic.
- There are a large number of different encryptions of a given message.
- One is chosen at random by the sender to represent the message.
 - Encryption is a *one-to-many* mapping.
- This scheme is both polynomially secure and **semantically secure**.

What then do you call proof?
— Henry James (1843–1916),
The Wings of the Dove (1902)

Leibniz knew what a proof is.
Descartes did not.
— Ian Hacking (1973)

What Is a Proof?

- A proof convinces a party of a certain claim.
 - “ $x^n + y^n \neq z^n$ for all $x, y, z \in \mathbb{Z}^+$ and $n > 2$.”
 - “Graph G is Hamiltonian.”
 - “ $x^p = x \pmod p$ for prime p and $p \nmid x$.”
- In mathematics, a proof is a fixed sequence of theorems.
 - Think of it as a written examination.
- We will extend a proof to cover a proof *process* by which the validity of the assertion is established.
 - Recall a job interview or an oral examination.

Prover and Verifier

- There are two parties to a proof.
 - The **prover** (**Peggy**).
 - The **verifier** (**Victor**).
- Given an assertion, the prover's goal is to convince the verifier of its validity (**completeness**).
- The verifier's objective is to accept only correct assertions (**soundness**).
- The verifier usually has an easier job than the prover.
- The setup is very much like the Turing test.^a

^aTuring (1950).

Interactive Proof Systems


- An **interactive proof** for a language L is a sequence of questions and answers between the two parties.
- At the end of the interaction, the verifier decides whether the claim is true or false.
- The verifier must be a probabilistic polynomial-time algorithm.
- The prover runs an exponential-time algorithm.^a
 - If the prover is not more powerful than the verifier, no interaction is needed!

^aSee the problem to Note 12.3.7 on p. 296 and Proposition 19.1 on p. 475, both of the textbook, about alternative complexity assumptions without affecting the definition. Contributed by Mr. Young-San Lin (B97902055) and Mr. Chao-Fu Yang (B97902052) on December 18, 2012.

Interactive Proof Systems (concluded)

- The system decides L if the following two conditions hold for any common input x .
 - If $x \in L$, then the probability that x is accepted by the verifier is at least $1 - 2^{-|x|}$.
 - If $x \notin L$, then the probability that x is accepted by the verifier with *any* prover replacing the original prover is at most $2^{-|x|}$.
- Neither the number of rounds nor the lengths of the messages can be more than a polynomial of $|x|$.

An Interactive Proof



IP^a

- IP is the class of all languages decided by an interactive proof system.
- When $x \in L$, the completeness condition can be modified to require that the verifier accept with certainty without affecting IP .^b
- Similar things cannot be said of the soundness condition when $x \notin L$.
- Verifier's coin flips can be public.^c

^aGoldwasser, Micali, & Rackoff (1985).

^bGoldreich, Mansour, & Sipser (1987).

^cGoldwasser & Sipser (1989).

The Relations of IP with Other Classes

- $NP \subseteq IP$.
 - IP becomes NP when the verifier is deterministic and there is only one round of interaction.^a
- $BPP \subseteq IP$.
 - IP becomes BPP when the verifier ignores the prover's messages.
- $IP = PSPACE$.^b

^aRecall Proposition 41 on p. 331.

^bShamir (1990).

Graph Isomorphism

- $V_1 = V_2 = \{ 1, 2, \dots, n \}$.
- Graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ are **isomorphic** if there exists a permutation π on $\{ 1, 2, \dots, n \}$ so that $(u, v) \in E_1 \Leftrightarrow (\pi(u), \pi(v)) \in E_2$.
- The task is to answer if $G_1 \cong G_2$.
- No known polynomial-time algorithms.^a
- The problem is in NP (hence IP).
- It is not likely to be NP-complete.^b

^aThe recent bound of Babai (2015) is $2^{O(\log^c n)}$ for some constant c .

^bSchöning (1987).

GRAPH NONISOMORPHISM

- $V_1 = V_2 = \{ 1, 2, \dots, n \}$.
- Graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ are **nonisomorphic** if there exist *no* permutations π on $\{ 1, 2, \dots, n \}$ so that $(u, v) \in E_1 \Leftrightarrow (\pi(u), \pi(v)) \in E_2$.
- The task is to answer if $G_1 \not\cong G_2$.
- Again, no known polynomial-time algorithms.
 - It is in coNP, but how about NP or BPP?
 - It is not likely to be coNP-complete.^a
- Surprisingly, GRAPH NONISOMORPHISM \in IP.^b

^aSchöning (1987).

^bGoldreich, Micali, & Wigderson (1986).

A 2-Round Algorithm

- 1: Victor selects a random $i \in \{1, 2\}$;
- 2: Victor selects a random permutation π on $\{1, 2, \dots, n\}$;
- 3: Victor applies π on graph G_i to obtain graph H ;
- 4: Victor sends (G_1, H) to Peggy;
- 5: **if** $G_1 \cong H$ **then**
- 6: Peggy sends $j = 1$ to Victor;
- 7: **else**
- 8: Peggy sends $j = 2$ to Victor;
- 9: **end if**
- 10: **if** $j = i$ **then**
- 11: Victor accepts; $\{G_1 \not\cong G_2.\}$
- 12: **else**
- 13: Victor rejects; $\{G_1 \cong G_2.\}$
- 14: **end if**

Analysis

- Victor runs in probabilistic polynomial time.
- Suppose $G_1 \not\cong G_2$.
 - Peggy is able to tell which G_i is isomorphic to H , so $j = i$.
 - So Victor always accepts.
- Suppose $G_1 \cong G_2$.
 - No matter which i is picked by Victor, Peggy or any prover sees 2 *identical* copies.
 - Peggy or any prover with exponential power has only probability one half of guessing i correctly.
 - So Victor erroneously accepts with probability $1/2$.
- Repeat the algorithm to obtain the desired probabilities.

Knowledge in Proofs

- Suppose I know a satisfying assignment to a satisfiable boolean expression.
- I can convince Alice of this by giving her the assignment.
- But then I give her more knowledge than is necessary.
 - Alice can claim that she found the assignment!
 - Login authentication faces essentially the same issue.
 - See
www.wired.com/wired/archive/1.05/atm_pr.html
for a famous ATM fraud in the U.S.

Knowledge in Proofs (concluded)

- Suppose I always give Alice random bits.
- Alice extracts no knowledge from me by any measure, but I prove nothing.
- Question 1: Can we design a protocol to convince Alice (the knowledge) of a secret without revealing anything extra?
- Question 2: How to define this idea rigorously?

Zero Knowledge Proofs^a

An interactive proof protocol (P, V) for language L has the **perfect zero-knowledge** property if:

- For every verifier V' , there is an algorithm M with expected polynomial running time.
- M on any input $x \in L$ generates the same probability distribution as the one that can be observed on the communication channel of (P, V') on input x .

^aGoldwasser, Micali, & Rackoff (1985).

Comments

- Zero knowledge is a property of the prover.
 - It is the robustness of the prover against attempts of the verifier to extract knowledge via interaction.
 - The verifier may deviate arbitrarily (but in polynomial time) from the predetermined program.
 - A verifier cannot use the transcript of the interaction to convince a third-party of the validity of the claim.
 - The proof is hence not transferable.

Comments (continued)

- Whatever a verifier can “learn” from the specified prover P via the communication channel could as well be computed from the verifier alone.
- The verifier does not learn anything except “ $x \in L$.”
- Zero-knowledge proofs yield no knowledge in the sense that they can be constructed by the verifier who believes the statement, and yet these proofs do convince him.

Comments (continued)

- The “paradox” is resolved by noting that it is *not* the transcript of the conversation that convinces the verifier.
- But the fact that this conversation was held “on line.”
- *Computational* zero-knowledge proofs are based on complexity assumptions.
 - M only needs to generate a distribution that is computationally indistinguishable from the verifier’s view of the interaction.

Comments (concluded)

- If one-way functions exist, then zero-knowledge proofs exist for every problem in NP.^a
- If one-way functions exist, then zero-knowledge proofs exist for every problem in PSPACE.^b
- The verifier can be restricted to the honest one (i.e., it follows the protocol).^c
- The coins can be public.^d
- The digital money Zcash (2016) is based on zero-knowledge proofs.

^aGoldreich, Micali, & Wigderson (1986).

^bOstrovsky & Wigderson (1993).

^cVadhan (2006).

^dVadhan (2006).

Quadratic Residuacity

- Let n be a product of two distinct primes.
- Assume extracting the square root of a quadratic residue modulo n is hard without knowing the factors.
- We next present a zero-knowledge proof for the input

$$x \in Z_n^*$$

being a quadratic residue.

Zero-Knowledge Proof of Quadratic Residuacity

- 1: **for** $m = 1, 2, \dots, \log_2 n$ **do**
- 2: Peggy chooses a random $v \in Z_n^*$ and sends $y = v^2 \bmod n$ to Victor;
- 3: Victor chooses a random bit i and sends it to Peggy;
- 4: Peggy sends $z = u^i v \bmod n$, where u is a square root of x ; $\{u^2 \equiv x \bmod n.\}$
- 5: Victor checks if $z^2 \equiv x^i y \bmod n$;
- 6: **end for**
- 7: Victor accepts x if Line 5 is confirmed every time;

A Useful Corollary of Lemma 82 (p. 680)

Corollary 83 *Let $n = pq$ be a product of two distinct primes. (1) If x and y are both quadratic residues modulo n , then $xy \in Z_n^*$ is a quadratic residue modulo n . (2) If x is a quadratic residue modulo n and y is a quadratic nonresidue modulo n , then $xy \in Z_n^*$ is a quadratic nonresidue modulo n .*

- Suppose x and y are both quadratic residues modulo n .
- Let $x \equiv a^2 \pmod{n}$ and $y \equiv b^2 \pmod{n}$.
- Now xy is a quadratic residue as $xy \equiv (ab)^2 \pmod{n}$.

The Proof (concluded)

- Suppose x is a quadratic residue modulo n and y is a quadratic nonresidue modulo n .
- By Lemma 82 (p. 680), $(x | p) = (x | q) = 1$ but, say, $(y | p) = -1$.
- Now xy is a quadratic nonresidue as $(xy | p) = -1$, again by Lemma 82 (p. 680).

Analysis

- Suppose x is a quadratic residue.
 - Then x 's square root u can be computed by Peggy.
 - Peggy can answer all challenges.
 - Now,

$$z^2 \equiv (u^i)^2 v^2 \equiv (u^2)^i v^2 \equiv x^i y \pmod{n}.$$

- So Victor will accept x .

Analysis (continued)

- Suppose x is a quadratic nonresidue.
 - Corollary 83 (p. 708) says if a is a quadratic residue, then xa is a quadratic nonresidue.
 - As y is a quadratic residue, $x^i y$ can be a quadratic residue (see Line 5) only when $i = 0$.
 - Peggy can answer only one of the two possible challenges, when $i = 0$.^a
 - So Peggy will be caught in any given round with probability one half.

^aLine 5 ($z^2 \equiv x^i y \pmod{n}$) cannot equate a quadratic residue z^2 with a quadratic nonresidue $x^i y$ when $i = 1$.

Analysis (continued)

- How about the claim of zero knowledge?
- The transcript between Peggy and Victor when x is a quadratic residue can be generated *without* Peggy!
- Here is how.
- Suppose x is a quadratic residue.^a
- In each round of interaction with Peggy, the transcript is a triplet (y, i, z) .
- We present an efficient Bob that generates (y, i, z) with the same probability *without* accessing Peggy's power.

^aThere is no zero-knowledge requirement when $x \notin L$.

Analysis (concluded)

- 1: Bob chooses a random $z \in Z_n^*$;
- 2: Bob chooses a random bit i ;
- 3: Bob calculates $y = z^2 x^{-i} \bmod n$;^a
- 4: Bob writes (y, i, z) into the transcript;

^aRecall Line 5 on p. 707: Victor checks if $z^2 \equiv x^i y \bmod n$.

Comments

- Assume x is a quadratic residue.
- For (y, i, z) , y is a random quadratic residue, i is a random bit, and z is a random number.
- Bob cheats because (y, i, z) is *not* generated in the same order as in the original transcript.
 - Bob picks Peggy's answer z first.
 - Bob then picks Victor's challenge i .
 - Bob finally patches the transcript.

Comments (concluded)

- So it is not the transcript that convinces Victor, but that *conversation with Peggy is held “on line.”*
- The same holds even if the transcript was generated by a cheating Victor’s interaction with (honest) Peggy.
- But we skip the details.^a

^aOr apply Vadhan (2006).

Zero-Knowledge Proof of 3 Colorability^a

- 1: **for** $i = 1, 2, \dots, |E|^2$ **do**
- 2: Peggy chooses a random permutation π of the 3-coloring ϕ ;
- 3: Peggy samples encryption schemes randomly, commits^b them, and sends $\pi(\phi(1)), \pi(\phi(2)), \dots, \pi(\phi(|V|))$ *encrypted* to Victor;
- 4: Victor chooses at random an edge $e \in E$ and sends it to Peggy for the coloring of the endpoints of e ;
- 5: **if** $e = (u, v) \in E$ **then**
- 6: Peggy reveals the colors $\pi(\phi(u))$ and $\pi(\phi(v))$ and “proves” that they correspond to their encryptions;
- 7: **else**
- 8: Peggy stops;
- 9: **end if**

^aGoldreich, Micali, & Wigderson (1986).

^bContributed by Mr. Ren-Shuo Liu (D98922016) on December 22, 2009.

```
10:  if the “proof” provided in Line 6 is not valid then
11:    Victor rejects and stops;
12:  end if
13:  if  $\pi(\phi(u)) = \pi(\phi(v))$  or  $\pi(\phi(u)), \pi(\phi(v)) \notin \{1, 2, 3\}$  then
14:    Victor rejects and stops;
15:  end if
16: end for
17: Victor accepts;
```

Analysis

- If the graph is 3-colorable and both Peggy and Victor follow the protocol, then Victor always accepts.
- Suppose the graph is not 3-colorable and Victor follows the protocol.
- Let e be an edge that is *not* colored legally.
- Victor will pick it with probability $1/m$ per round, where $m = |E|$.
- Then however Peggy plays, Victor will reject with probability at least $1/m$ per round.

Analysis (concluded)

- So Victor will accept with probability at most

$$(1 - m^{-1})^{m^2} \leq e^{-m}.$$

- Thus the protocol is a valid IP protocol.
- This protocol yields no knowledge to Victor as all he gets is a bunch of random pairs.
- The proof that the protocol is zero-knowledge to *any* verifier is intricate.^a

^aBut no longer necessary because of Vadhan (2006).

Comments

- Each $\pi(\phi(i))$ is encrypted by a different cryptosystem in Line 3.^a
 - Otherwise, the coloring will be revealed in Line 6.
- Each edge e must be picked randomly.^b
 - Otherwise, Peggy will know Victor's game plan and plot accordingly.

^aContributed by Ms. Yui-Huei Chang (R96922060) on May 22, 2008

^bContributed by Mr. Chang-Rong Hung (R96922028) on May 22, 2008

Approximability

All science is dominated by
the idea of approximation.
— Bertrand Russell (1872–1970)

Just because the problem is NP-complete
does not mean that
you should not try to solve it.
— Stephen Cook (2002)

Tackling Intractable Problems

- Many important problems are NP-complete or worse.
- **Heuristics** have been developed to attack them.
- They are **approximation algorithms**.
- How good are the approximations?
 - We are looking for theoretically *guaranteed* bounds, not “empirical” bounds.
- Are there NP problems that cannot be approximated *well* (assuming $NP \neq P$)?
- Are there NP problems that cannot be approximated *at all* (assuming $NP \neq P$)?

Some Definitions

- Given an **optimization problem**, each problem instance x has a set of **feasible solutions** $F(x)$.
- Each feasible solution $s \in F(x)$ has a cost $c(s) \in \mathbb{Z}^+$.
 - Here, cost refers to the quality of the feasible solution, not the time required to obtain it.
 - It is our **objective function**: total distance, number of satisfied clauses, cut size, etc.

Some Definitions (concluded)

- The **optimum cost** is

$$\text{OPT}(x) = \min_{s \in F(x)} c(s)$$

for a minimization problem.

- It is

$$\text{OPT}(x) = \max_{s \in F(x)} c(s)$$

for a maximization problem.

Approximation Algorithms

- Let (polynomial-time) algorithm M on x returns a feasible solution.
- M is an ϵ -**approximation algorithm**, where $\epsilon \geq 0$, if for all x ,

$$\frac{|c(M(x)) - \text{OPT}(x)|}{\max(\text{OPT}(x), c(M(x)))} \leq \epsilon.$$

- For a minimization problem,

$$\frac{c(M(x)) - \min_{s \in F(x)} c(s)}{c(M(x))} \leq \epsilon.$$

- For a maximization problem,

$$\frac{\max_{s \in F(x)} c(s) - c(M(x))}{\max_{s \in F(x)} c(s)} \leq \epsilon. \quad (17)$$

Lower and Upper Bounds

- For a minimization problem,

$$\min_{s \in F(x)} c(s) \leq c(M(x)) \leq \frac{\min_{s \in F(x)} c(s)}{1 - \epsilon}.$$

- For a maximization problem,

$$(1 - \epsilon) \times \max_{s \in F(x)} c(s) \leq c(M(x)) \leq \max_{s \in F(x)} c(s). \quad (18)$$

Lower and Upper Bounds (concluded)

- ϵ ranges between 0 (best) and 1 (worst).
- For minimization problems, an ϵ -approximation algorithm returns solutions within

$$\left[\text{OPT}, \frac{\text{OPT}}{1 - \epsilon} \right].$$

- For maximization problems, an ϵ -approximation algorithm returns solutions within

$$[(1 - \epsilon) \times \text{OPT}, \text{OPT}].$$

Approximation Thresholds

- For each NP-complete optimization problem, we shall be interested in determining the *smallest* ϵ for which there is a polynomial-time ϵ -approximation algorithm.
- But sometimes ϵ has no minimum value.
- The **approximation threshold** is the greatest lower bound of all $\epsilon \geq 0$ such that there is a polynomial-time ϵ -approximation algorithm.
- By a standard theorem in real analysis, such a threshold exists.^a

^aBauldry (2009).

Approximation Thresholds (concluded)

- The approximation threshold of an optimization problem is anywhere between 0 (approximation to any desired degree) and 1 (no approximation is possible).
- If $P = NP$, then all optimization problems *in NP* have an approximation threshold of 0.
- So assume $P \neq NP$ for the rest of the discussion.

Approximation Ratio

- ϵ -approximation algorithms can also be measured via the **approximation ratio**:^a

$$\frac{c(M(x))}{\text{OPT}(x)}.$$

- For a minimization problem, the approximation ratio is

$$1 \leq \frac{c(M(x))}{\min_{s \in F(x)} c(s)} \leq \frac{1}{1 - \epsilon}. \quad (19)$$

- For a maximization problem, the approximation ratio is

$$1 - \epsilon \leq \frac{c(M(x))}{\max_{s \in F(x)} c(s)} \leq 1. \quad (20)$$

^aWilliamson & Shmoys (2011).

Approximation Ratio (concluded)

- Suppose there is an approximation algorithm that achieves an approximation ratio of θ .
 - For a minimization problem, it implies a $(1 - \theta^{-1})$ -approximation algorithm by Eq. (19).
 - For a maximization problem, it implies a $(1 - \theta)$ -approximation algorithm by Eq. (20).

NODE COVER

- NODE COVER seeks the smallest $C \subseteq V$ in graph $G = (V, E)$ such that for each edge in E , at least one of its endpoints is in C .
- A heuristic to obtain a good node cover is to iteratively move a node with the *highest degree* to the cover.
- This turns out to produce an approximation ratio of^a

$$\frac{c(M(x))}{\text{OPT}(x)} = \Theta(\log n).$$

- So it is not an ϵ -approximation algorithm for any constant $\epsilon < 1$ (see p. 733).

^aChvátal (1979).

A 0.5-Approximation Algorithm^a

- 1: $C := \emptyset$;
- 2: **while** $E \neq \emptyset$ **do**
- 3: Delete an arbitrary edge $[u, v]$ from E ;
- 4: Add u and v to C ; {Add 2 nodes to C each time.}
- 5: Delete edges incident with u or v from E ;
- 6: **end while**
- 7: **return** C ;

^aGavril (1974).


Analysis

- It is easy to see that C is a node cover.
- C contains $|C|/2$ edges.^a
- No two edges of C share a node.^b
- *Any* node cover C' must contain at least one node from *each* of the edges of C .
 - If there is an edge in C both of whose ends are outside C' , then C' will not be a cover.

^aThe edges deleted in Line 3.

^bIn fact, C as a set of edges is a *maximal* matching.

Analysis (continued)



Analysis (concluded)


- This means that $\text{OPT}(G) \geq |C|/2$.
- The approximation ratio is hence

$$\frac{|C|}{\text{OPT}(G)} \leq 2.$$

- So we have a 0.5-approximation algorithm.^a
- And the approximation threshold is therefore ≤ 0.5 .

^aRecall p. 733.

The 0.5 Bound Is Tight for the Algorithm^a



^aContributed by Mr. Jenq-Chung Li (R92922087) on December 20, 2003. **König's theorem** says the size of a maximum matching equals that of a minimum node cover in a bipartite graph.

Remarks

- The approximation threshold is at least^a

$$1 - \left(10\sqrt{5} - 21\right)^{-1} \approx 0.2651.$$

- The approximation threshold is 0.5 if one assumes the **unique games conjecture** (UGC).^b
- This ratio 0.5 is also the lower bound for any “greedy” algorithms.^c

^aDinur & Safra (2002).

^bKhot & Regev (2008).

^cDavis & Impagliazzo (2004).