Information-Theoretic Security with Asymmetries

Tim Beyne[1] Yu Long Chen[12]

COSIC, KU Leuven

NIST

August 22, 2024

Hypothesis Testing

A Practical Example: COVID 19

H_0 : Patient should be placed in an ICU

vs

H_1 : Patient shouldn't be placed in an ICU

A Practical Example: COVID 19 (2)

- False positives (α): loss of patients
- False negatives (β): unnecessary expenses

Trade-off between α and β

A Practical Example: COVID 19 (2)

False positives (α): loss of patients

False negatives (β): unnecessary expenses

Trade-off between α and β

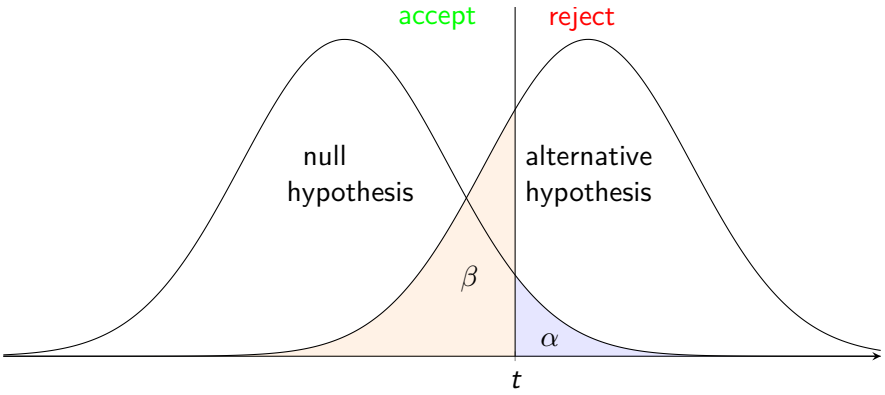

$1 - \beta =$ power of the test (function of α)

Another Example: CRYPTO Reviewing

H_0 : This paper should be accepted

vs

H_1 : This paper should be rejected

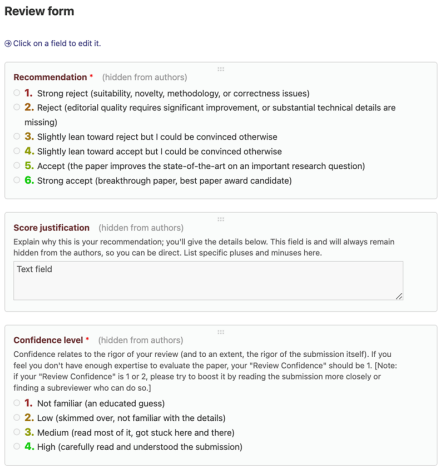

Review form

⊕ Click on a field to edit it.

Recommendation * (hidden from authors)

1. Strong reject (suitability, novelty, methodology, or correctness issues)
2. Reject (editorial quality requires significant improvement, or substantial technical details are missing)
3. Slightly lean toward reject but I could be convinced otherwise
4. Slightly lean toward accept but I could be convinced otherwise
5. Accept (the paper improves the state-of-the-art on an important research question)
6. Strong accept (breakthrough paper, best paper award candidate)

Score justification (hidden from authors)

Explain why this is your recommendation; you'll give the details below. This field is and will always remain hidden from the authors, so you can be direct. List specific pluses and minuses here.

Text field

Confidence level * (hidden from authors)

Confidence relates to the rigor of your review (and to an extent, the rigor of the submission itself). If you feel you don't have enough expertise to evaluate the paper, your "Review Confidence" should be 1. [Note: if your "Review Confidence" is 1 or 2, please try to boost it by reading the submission more closely or finding a subreviewer who can do so.]

1. Not familiar (an educated guess)
2. Low (skimmed over, not familiar with the details)
3. Medium (read most of it, got stuck here and there)
4. High (carefully read and understood the submission)

Indistinguishability

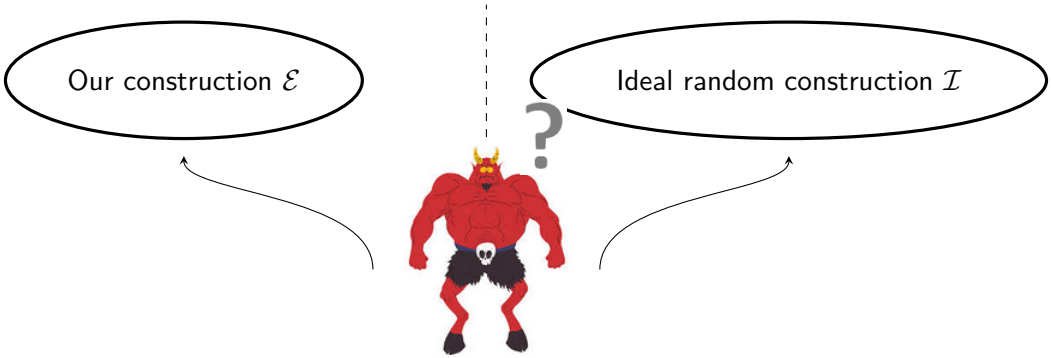

- Deterministic information-theoretical distinguisher
  = Hypothesis test between transcript distributions P and Q
- Traditionally we measure security by $\Delta(P;Q) = (1-\beta) - \alpha$

Asymmetrical Costs

- Cost of an attack is application-dependent: cost function $\mathscr{C}(\alpha, \beta)$
- Advantage corresponds to $\mathscr{C}(\alpha, \beta) = \alpha + \beta$
  $\rightarrow$ Minimized for $\alpha = \beta$

Asymmetrical Costs

- Cost of an attack is application-dependent: cost function $\mathscr{C}(\alpha, \beta)$
- Advantage corresponds to $\mathscr{C}(\alpha, \beta) = \alpha + \beta$
  $\rightarrow$ Minimized for $\alpha = \beta$

However ...

Asymmetrical Costs

- Cost of an attack is application-dependent: cost function $\mathscr{C}(\alpha, \beta)$
- Advantage corresponds to $\mathscr{C}(\alpha, \beta) = \alpha + \beta$
  $\rightarrow$ Minimized for $\alpha = \beta$

However ...

- Cost function is usually not symmetric in α and β
- Choice of null hypothesis matters
- Can we lower bound all possible cost functions?

- Solution: power bound

$$1 - \beta \leq f(\alpha)$$

Power bounds

Power bound

$$1 - \beta \leq f(\alpha)$$

Power bounds

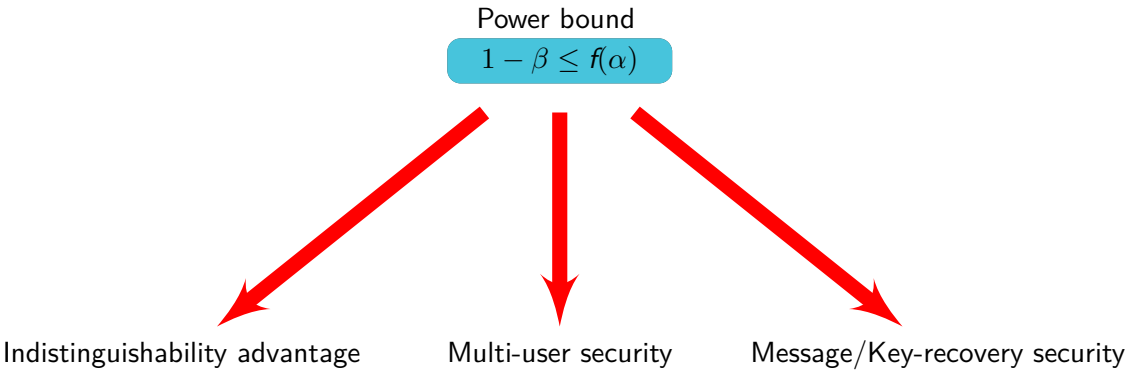

Only Modular Approach For Standard Model Multi-User Security

advantage bound
$1 - \beta - \alpha \leq \epsilon(q)$

?

mu advantage bound
$1 - \beta - \alpha \leq \epsilon(\sum_{i=1}^{u} q_i)$

power bound
$1 - \beta \leq \frac{\alpha}{1-\epsilon(q)}$

mu power bound
$1 - \beta \leq \max \prod_{i=1}^{u} \frac{\alpha}{1-\epsilon(q_i)}$

Only Modular Approach For Standard Model Multi-User Security

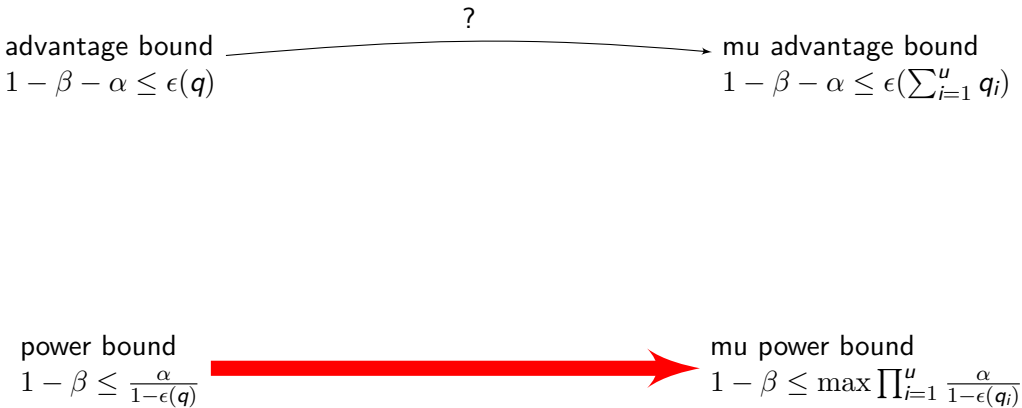

Only Modular Approach For Standard Model Multi-User Security

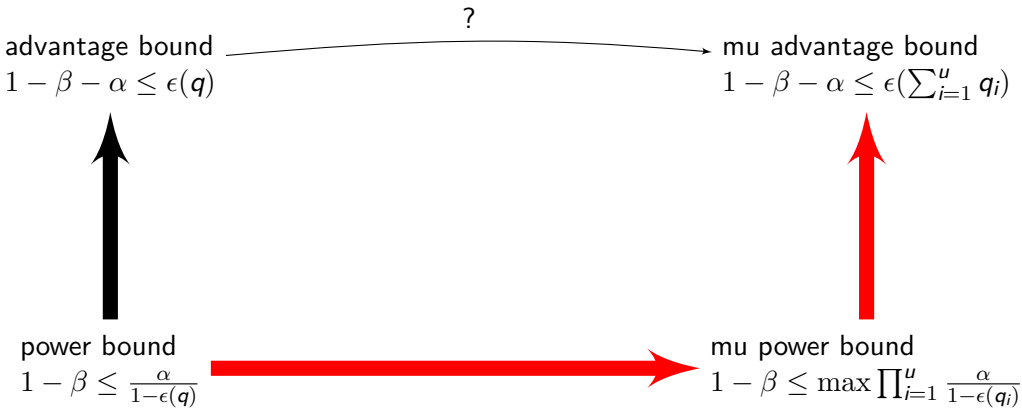

What's the Catch?

- Is it more difficult to prove power bounds?
- In many cases it is not; we show exising proofs can be adapted:
  - Hybrid arguments carry over
  - H-coefficient method carries over
- Examples in our paper:
  - PRP-PRF switching lemma
  - Even-Mansour (EM)
  - Sum-of-Permutations (SoP)

Application: PRP-PRF Switching Lemma

- Difference between the worlds: collision event
- Replace the PRF transcript distribution P by $P_{\mathcal{T}\setminus B}$ with $B \subseteq \mathcal{T}$

$$1 - \beta \leq \frac{\alpha}{1 - \varepsilon}$$

- $B =$ collision event of PRF with $P(B) \leq \varepsilon = \frac{q(q-1)}{2N}$

Application: PRP-PRF Switching Lemma (2)

- Power bound (null hypothesis = outputs from PRF):

$$1 - \beta \leq \frac{\alpha}{1 - \frac{q(q-1)}{2N}}$$

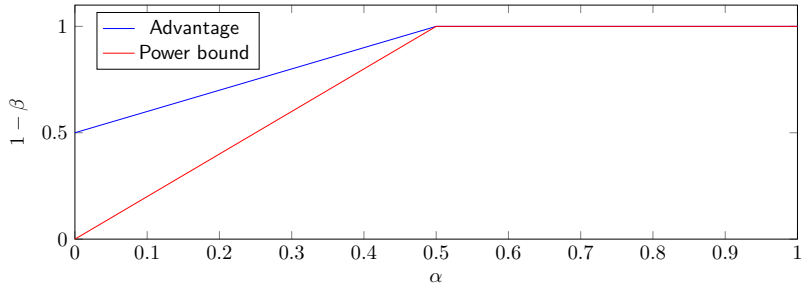

- Swapping of null and alternative hypothesis (null hypothesis = outputs from PRP)

$$1 - \beta \leq \frac{q(q-1)}{2N} + \left(1 - \frac{q(q-1)}{2N}\right) \alpha$$

Observation of a collision $\rightarrow$ null hypothesis can be rejected with certainty

Hybrid Arguments

- Advantage bound:

$$\Delta(P;Q) \leq \Delta(P;X) + \Delta(X;Q)$$

- Power bound:
  - Every distinguisher from P to X with α_1 and β_1 satisfies $1 - \beta_1 \leq f(\alpha_1)$
  - Every distinguisher from X to Q with α_2 and β_2 satisfies $1 - \beta_2 \leq g(\alpha_2)$
  - All distinguishers from P to Q

$$1 - \beta \leq g(f(\alpha))$$

Application to CTR Mode

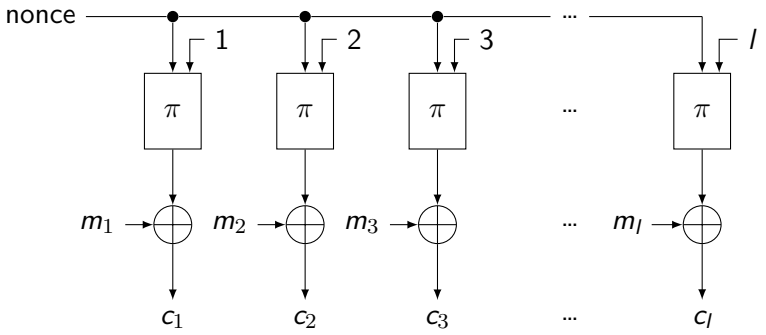

$$1 - \beta \leq g(f(\alpha)) = \frac{\alpha}{1 - \frac{\sigma(\sigma-1)}{2N}}$$

Application to CTR Mode

$$1 - \beta \leq g(f(\alpha)) = \frac{\alpha}{1 - \frac{\sigma(\sigma-1)}{2N}}$$

- Security in statistical distance
- Multi-user security bound
- Security bounds in other models e.g. message recovery

Application to CTR Mode Message Recovery

- An unknown message can take 2^b possible values
- Message-recovery security: adversary must output a list of $\alpha 2^b$ candidate values for unknown message
- Probability that distinguisher incorrectly guess the message: $\leq \alpha 2^b / 2^b = \alpha$
- Probability that message-recovery attack succeeds:

$$1 - \beta \leq \frac{\alpha}{1 - \frac{\sigma(\sigma-1)}{2N}}$$

Application to CTR Mode Message Recovery (2)

- The adversary has to uniquely recover b bits $\rightarrow \alpha = 2^{-b}$ $\qquad \varepsilon = \frac{\sigma(\sigma-1)}{2N}$
- Advantage bound: probability of recovering the message is
$$2^{-b} + \varepsilon$$
- Power bound: probability of recovering the message is
$$2^{-b}\big(1 + \varepsilon + \mathcal{O}(\varepsilon^2)\big)$$

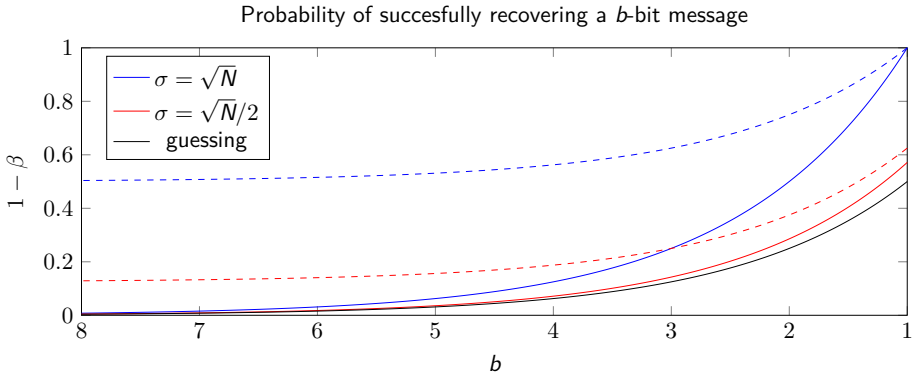

Conclusion

New results

- Power bounds instead of advantage bounds
- Hybrid arguments and H-coefficient method in power model + application to PRP-PRF, EM, and SoP
- Methods to convert single-user power bounds into multi-user power bounds

Future research

- Adapting other security notions to power bound model
- Adapting other proof techniques to power bound
- Prove power bound of constructions outside of symmetric-key cryptography

Conclusion

New results

- Power bounds instead of advantage bounds
- Hybrid arguments and H-coefficient method in power model + application to PRP-PRF, EM, and SoP
- Methods to convert single-user power bounds into multi-user power bounds

Future research

- Adapting other security notions to power bound model
- Adapting other proof techniques to power bound
- Prove power bound of constructions outside of symmetric-key cryptography

Thank you for your attention!