Relativistic quantum cryptography

Jed Kaniewski

Centre for Quantum Technologies, NUS
QuTech, TU Delft

23 March 2015

Advances In Quantum Cryptography Workshop
Paris Centre for Quantum Computing

Outline

- **Two-party** cryptography

Outline

- **Two-party** cryptography
- Classical **non-communicating** (split) models

Outline

- **Two-party** cryptography
- Classical **non-communicating** (split) models
- Is **quantum** any useful?

Outline

- **Two-party** cryptography
- Classical **non-communicating** (split) models
- Is **quantum** any useful?
- Communication constraints from **relativity**

Outline

- Two-party cryptography
- Classical non-communicating (split) models
- Is quantum any useful?
- Communication constraints from relativity
- The simplest relativistic setup and three bit commitment protocols

Outline

- **Two-party** cryptography
- Classical **non-communicating** (split) models
- Is **quantum** any useful?
- Communication constraints from **relativity**
- The simplest relativistic setup and **three bit commitment protocols**
- Longer commitments? The trouble of **multiple** rounds...

Two-party cryptography

Two-party cryptography

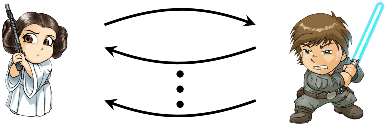

the protocol terminates
the outcome is correct

Two-party cryptography

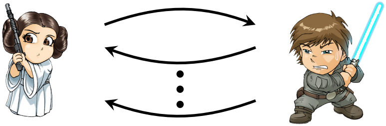

the protocol terminates
the outcome is correct

Two-party cryptography

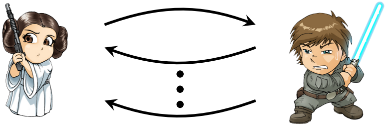

the protocol terminates
the outcome is correct

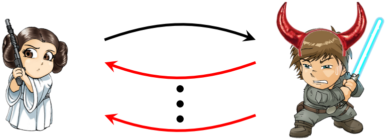

the **honest** party is **protected**

Two-party cryptography

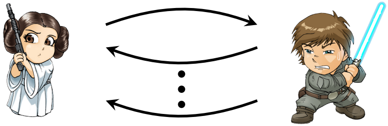

the protocol terminates
the outcome is correct

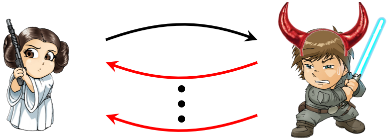

the **honest** party is **protected**

Examples: coin flipping, secure function evaluation, bit commitment

Bit commitment

Bit commitment

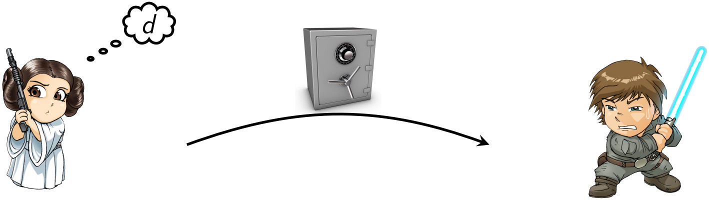

Commit phase

Bit commitment

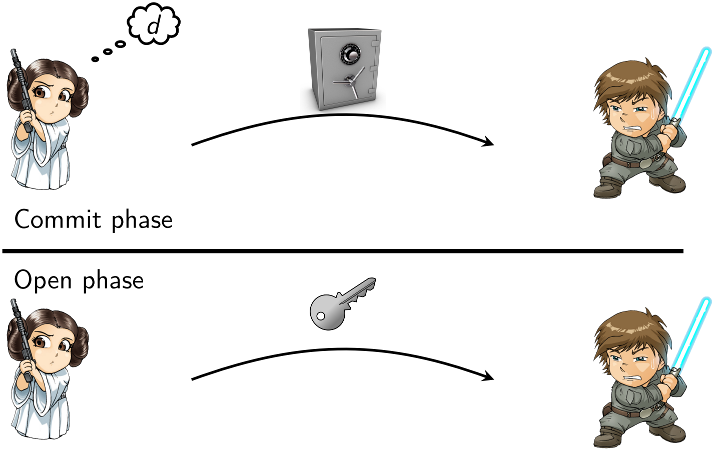

Bit commitment

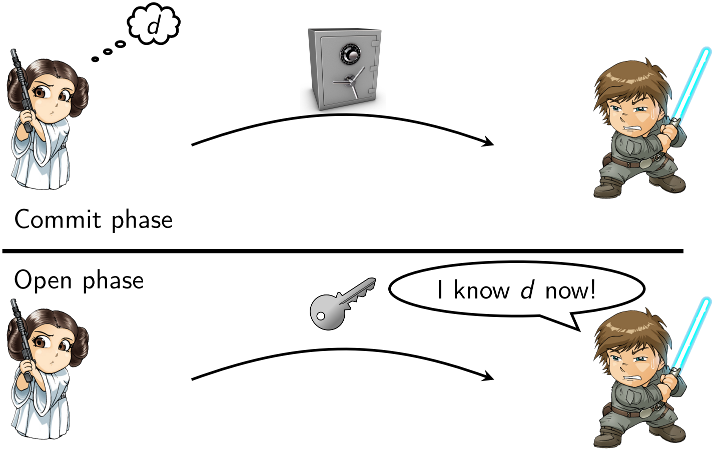

Bit commitment

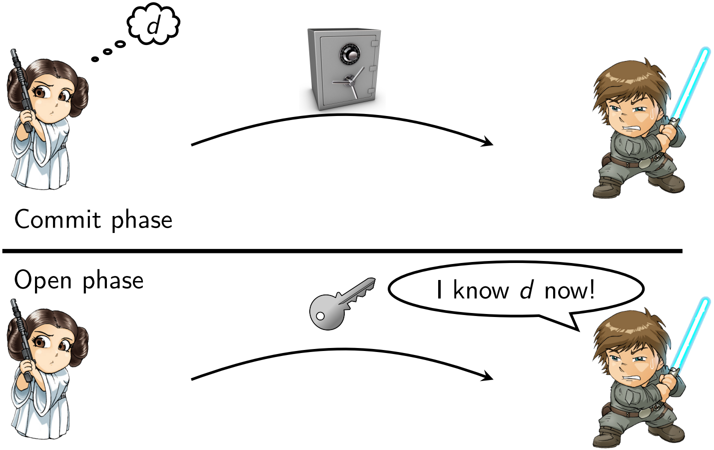

Correctness: both honest $\Longrightarrow$ Bob always accepts the commitment

Hiding: Alice honest $\Longrightarrow$ Bob does not know d before the open phase

Binding: Bob honest $\Longrightarrow$ $\exists$ only one value of d that Alice can unveil

The classical no-go

Theorem (Classical no-go)

*Any protocol that is **correct** and **hiding** allows Alice to **cheat perfectly**.*

Intuition: if at the end of the commit phase Bob is ignorant about d then for both values of d there must exist an opening strategy for Alice that will make him accept.

The classical no-go

Theorem (Classical no-go)

*Any protocol that is **correct** and **hiding** allows Alice to **cheat perfectly**.*

Intuition: if at the end of the commit phase Bob is ignorant about d then for both values of d there must exist an opening strategy for Alice that will make him accept.

How could **communication constraints** possibly help to avoid this?

Maybe cheating becomes difficult if it has to be coordinated between multiple non-communicating agents?

The classical no-go

Theorem (Classical no-go)

*Any protocol that is **correct** and **hiding** allows Alice to **cheat perfectly**.*

Intuition: if at the end of the commit phase Bob is ignorant about d then for both values of d there must exist an opening strategy for Alice that will make him accept.

How could **communication constraints** possibly help to avoid this?

Maybe cheating becomes difficult if it has to be coordinated between multiple **non-communicating** agents?

It only makes sense to split a party during their **"turn to cheat"**

Receiver split in the commit phase

Idea #1: maybe the combined information of Bob_1 and Bob_2 determines the commitment but ...?

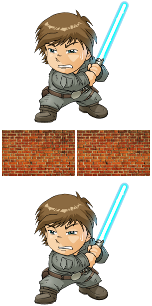

Receiver split in the commit phase

Idea #1: maybe the combined information of Bob_1 and Bob_2 determines the commitment but ...?

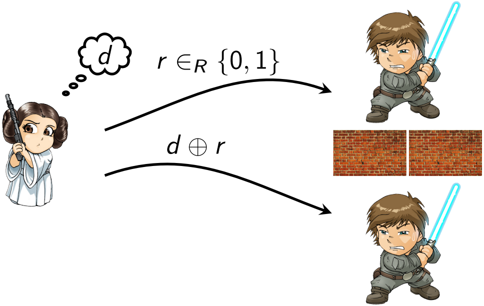

Receiver split in the commit phase

Idea #1: maybe the combined information of Bob_1 and Bob_2 determines the commitment but ...?

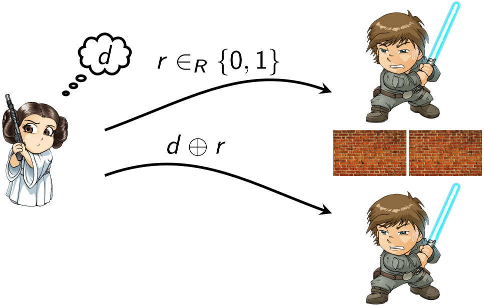

individual Bobs learn **nothing** about d as long as no communication is allowed

Receiver split in the commit phase

Idea #1: maybe the combined information of Bob_1 and Bob_2 determines the commitment but ...?

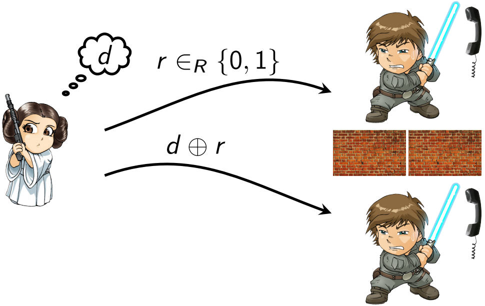

individual Bobs learn **nothing** about d as long as no communication is allowed

as soon as communication is allowed the commitment is unveiled

Receiver split in the commit phase

Idea #1: maybe the combined information of Bob_1 and Bob_2 determines the commitment but ...?

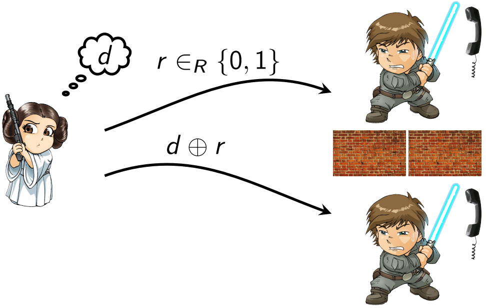

individual Bobs learn **nothing** about *d* as long as no communication is allowed

as soon as communication is allowed the commitment is unveiled

secret-sharing BC

Committer split in the open phase

Idea #2: maybe Alice_1 and Alice_2 find it difficult to coordinate the openings?

Committer split in the open phase

Idea #2: maybe Alice_1 and Alice_2 find it difficult to coordinate the openings?

Good start but...
...what is the **exact definition** of cheating in the split committer model?

Security for honest Bob as a game

1. Alice performs a **generic commit strategy**
2. Alice is **challenged** to open one of the bits with equal probabilities
3. Alice wins if Bob **accepts** the commitment

Security for honest Bob as a game

1. Alice performs a **generic commit strategy**
2. Alice is **challenged** to open one of the bits with equal probabilities
3. Alice wins if Bob **accepts** the commitment

Want: $\mathrm{p}_{\mathrm{win}} \leq \frac{1}{2} + \varepsilon$ for all strategies of dishonest Alice
Ideally, ε should be exponentially small in the number of bits exchanged

[Note that $2\,\mathrm{p}_{\mathrm{win}} = p_0 + p_1$ for $p_d =$ "probability that Alice successfully unveils d"
$\Longrightarrow$ equivalent to the usual requirement $p_0 + p_1 \leq 1 + 2\varepsilon$]

Security for honest Bob as a game

1. Alice performs a **generic commit strategy**
2. Alice is **challenged** to open one of the bits with equal probabilities
3. Alice wins if Bob **accepts** the commitment

Want: $\mathrm{p_{win}} \leq \frac{1}{2} + \varepsilon$ for all strategies of dishonest Alice
Ideally, ε should be exponentially small in the number of bits exchanged

[Note that $2\,\mathrm{p_{win}} = p_0 + p_1$ for $p_d =$ "probability that Alice successfully unveils d"
$\implies$ equivalent to the usual requirement $p_0 + p_1 \leq 1 + 2\varepsilon$]

Question: Who receives the challenge? **Both** Alices or just **one** of them?

Security for honest Bob as a game

1. Alice performs a **generic commit strategy**
2. Alice is **challenged** to open one of the bits with equal probabilities
3. Alice wins if Bob **accepts** the commitment

Want: $\mathrm{p_{win}} \leq \frac{1}{2} + \varepsilon$ for all strategies of dishonest Alice
Ideally, ε should be exponentially small in the number of bits exchanged

[Note that $2\,\mathrm{p_{win}} = p_0 + p_1$ for $p_d =$ "probability that Alice successfully unveils d"
$\implies$ equivalent to the usual requirement $p_0 + p_1 \leq 1 + 2\varepsilon$]

Question: Who receives the challenge? **Both** Alices or just **one** of them?

If **just one** (local command) then simple checking for consistency is sufficient.

If **both** (global command) then we need to try harder...

Committer split in the open phase

Strongly split committer (both commit and open phases):

$a \in_R \{0,1\}^n$

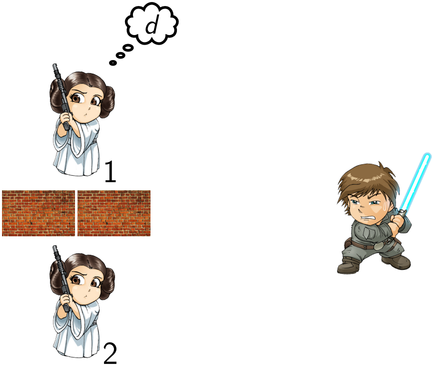

$b \in_R \{0,1\}^n$

Committer split in the open phase

Strongly split committer (both commit and open phases):

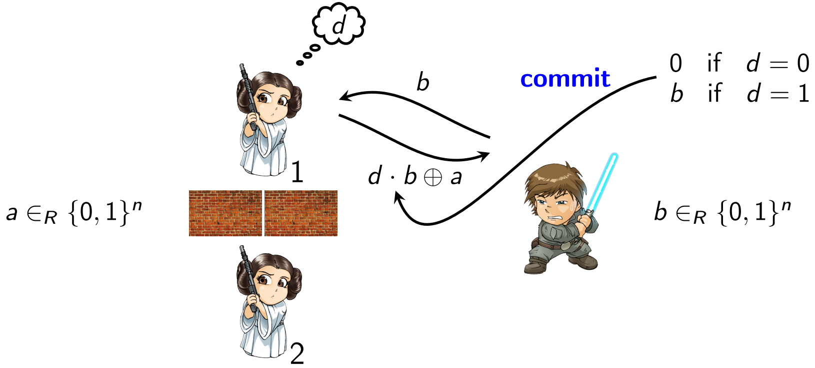

Bob learns nothing because the message is one-time-padded

Committer split in the open phase

Strongly split committer (both commit and open phases):

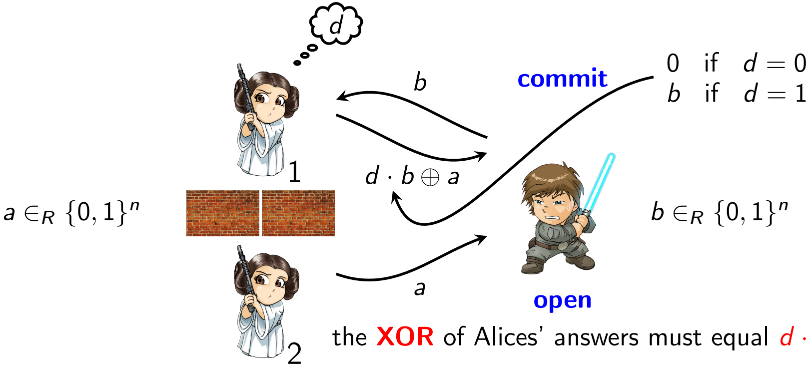

the **XOR** of Alices' answers must equal $d \cdot b$

Bob learns nothing because the message is one-time-padded

Committer split in the open phase

Strongly split committer (both commit and open phases):

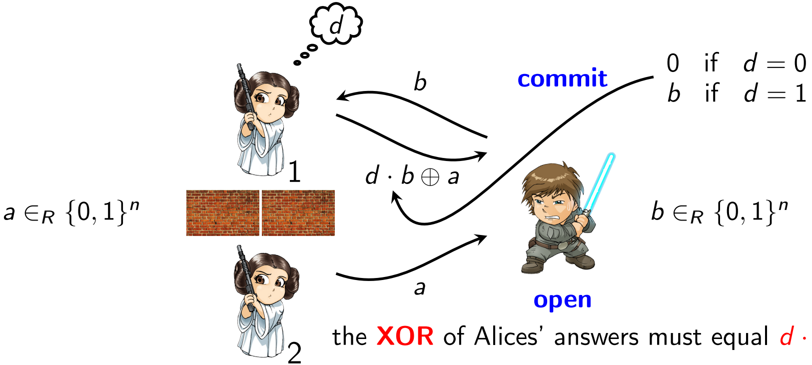

the **XOR** of Alices' answers must equal $d \cdot b$

Bob learns nothing because the message is one-time-padded

Intuition: Alice_2 cannot cheat because she does not know b

one-time pad BC (Ben-Or et al., Kent, Simard et al.)

Committer split in the open phase

Weakly split committer (only open phase):

Both Alices have **full information** about the commit phase and they can agree on a consistent cheating strategy; the **no-go still holds**.

In the classical case splitting at this stage does not make any difference because everything can be **copied...**

Going quantum?

In the **classical** world...

split model	BC possible?
split receiver	**yes** (secret-sharing BC)
weakly split committer	**no**
strongly split committer	**yes** (one-time pad BC)

Going quantum?

In the **classical** world...

split model	BC possible?
split receiver	**yes** (secret-sharing BC)
weakly split committer	**no**
strongly split committer	**yes** (one-time pad BC)

Does **quantum** make any difference?

Going quantum?

In the **classical** world...

split model	BC possible?
split receiver	yes (secret-sharing BC)
weakly split committer	**no**
strongly split committer	yes (one-time pad BC)

Does **quantum** make any difference?

Yes!

- **strongly** split committer: security proof for honest Bob against quantum adversaries for one-time pad BC necessary!
- **weakly** split committer: the no-go does not apply anymore!

One-time pad BC – honest Bob

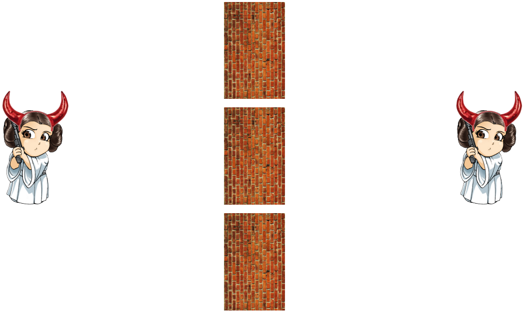

One-time pad BC – honest Bob

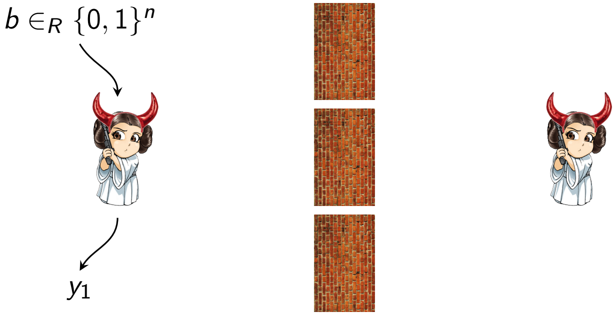

One-time pad BC – honest Bob

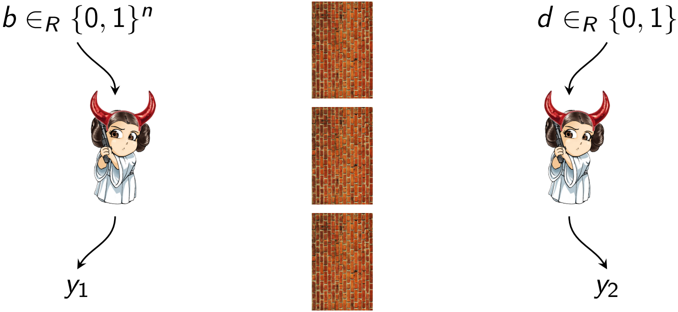

win **iff** $y_1 \oplus y_2 = d \cdot b$

One-time pad BC – honest Bob

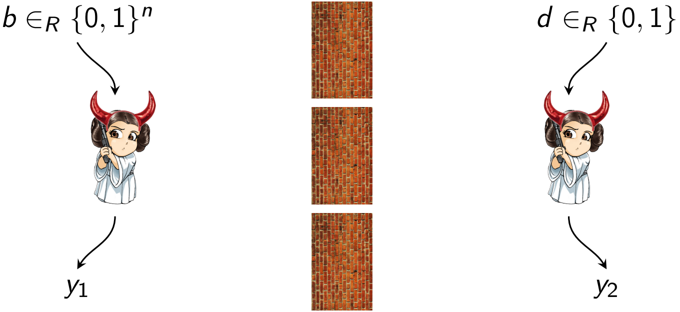

win **iff** $y_1 \oplus y_2 = d \cdot b$

Classically: $p_{win} = \frac{1}{2} + \frac{1}{2^n}$

Quantumly: $p_{win} \leq \frac{1}{2} + \frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{2^n}}$ [Sikora, Chailloux, Kerenidis'14]

One-time pad BC – honest Bob

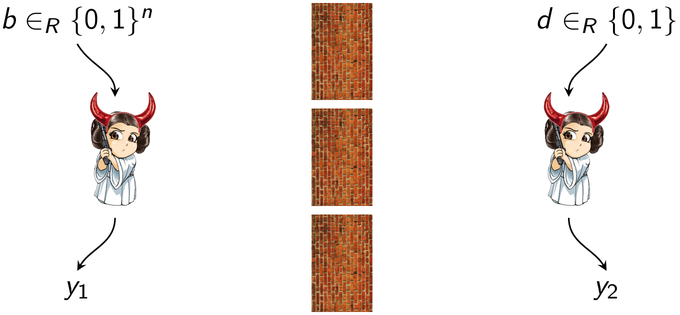

win **iff** $y_1 \oplus y_2 = d \cdot b$

Classically: $p_{win} = \frac{1}{2} + \frac{1}{2^n}$ (**tight**)

Quantumly: $p_{win} \leq \frac{1}{2} + \frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{2^n}}$ [Sikora, Chailloux, Kerenidis'14]

exponential decay conjectured to be (essentially) tight

One-time pad BC – honest Bob

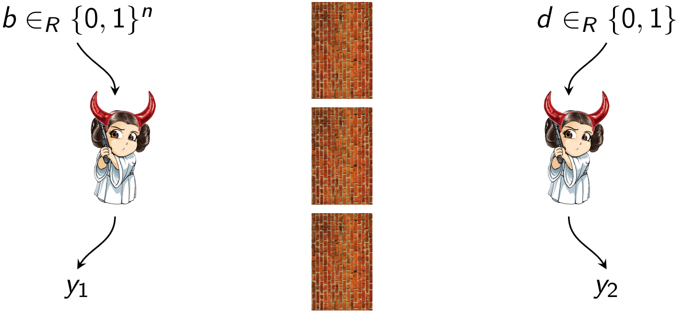

win **iff** $y_1 \oplus y_2 = d \cdot b$

Classically: $p_{win} = \frac{1}{2} + \frac{1}{2^n}$ (**tight**)

Quantumly: $p_{win} \leq \frac{1}{2} + \frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{2^n}}$ [Sikora, Chailloux, Kerenidis'14]

exponential decay conjectured to be (essentially) tight

quantum-classical gap

quantum adversary **strictly more** powerful

Weakly split committer with quantum

Weakly split committer with quantum

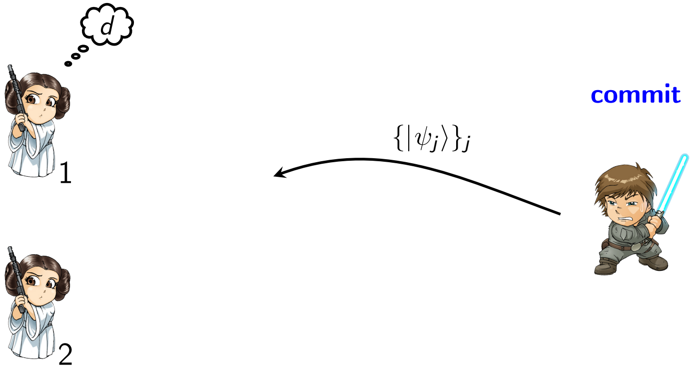

Weakly split committer with quantum

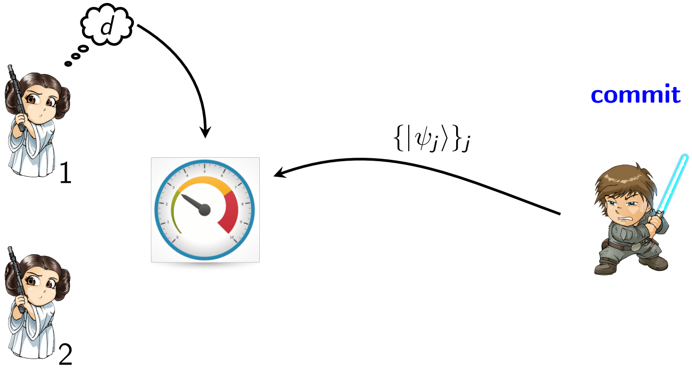

Weakly split committer with quantum

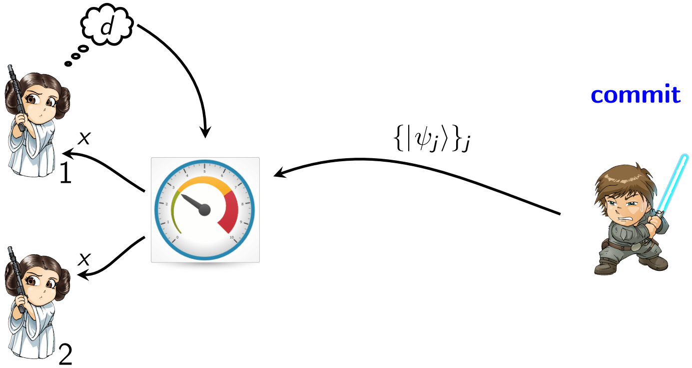

Weakly split committer with quantum

Weakly split committer with quantum

commit

Weakly split committer with quantum

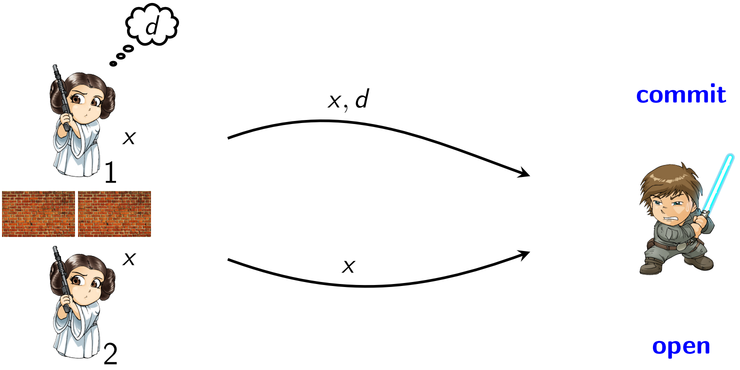

Weakly split committer with quantum

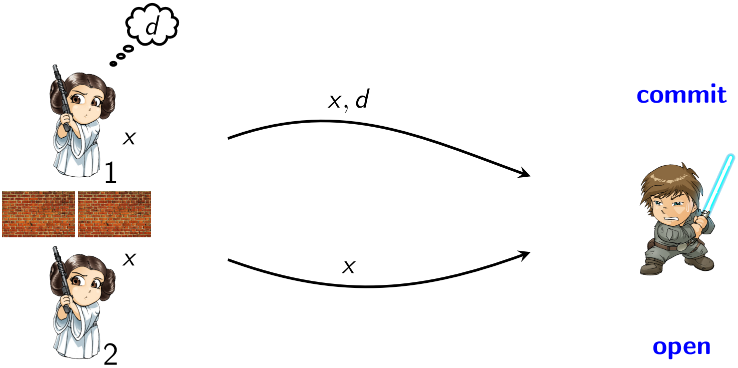

Alices cannot cheat because this would require both of them to know outcomes of **incompatible** measurements

BC by transmitting measurement outcomes (Kent)

The simplest relativistic set-up

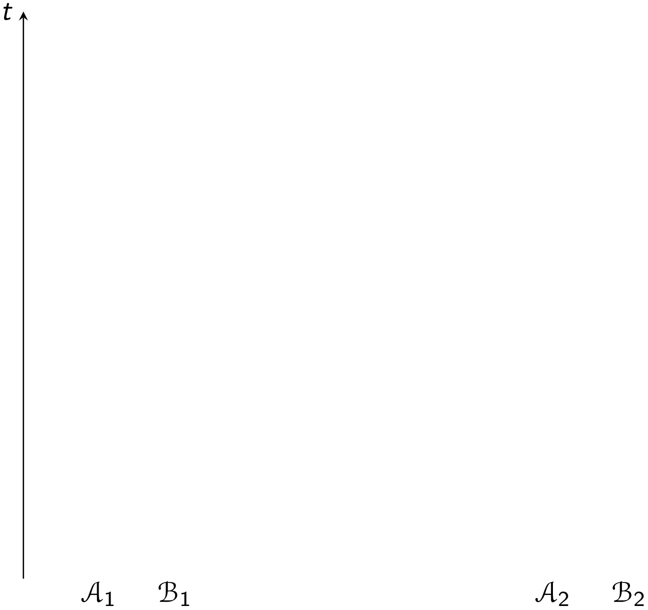

The simplest relativistic set-up

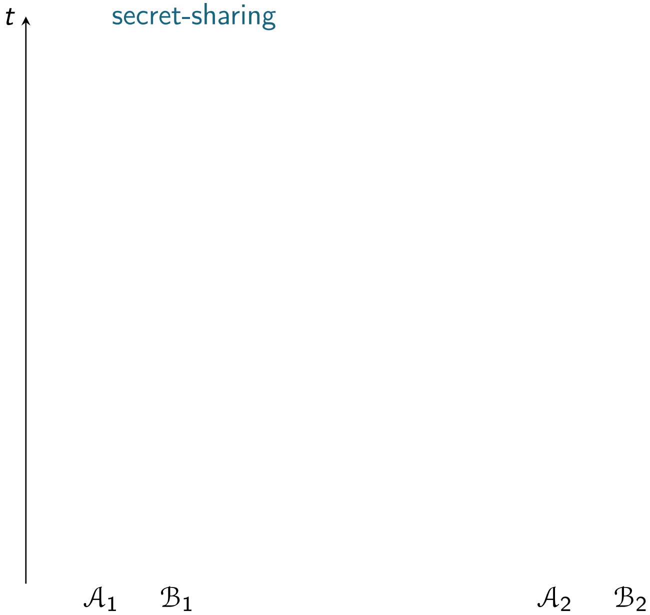

The simplest relativistic set-up

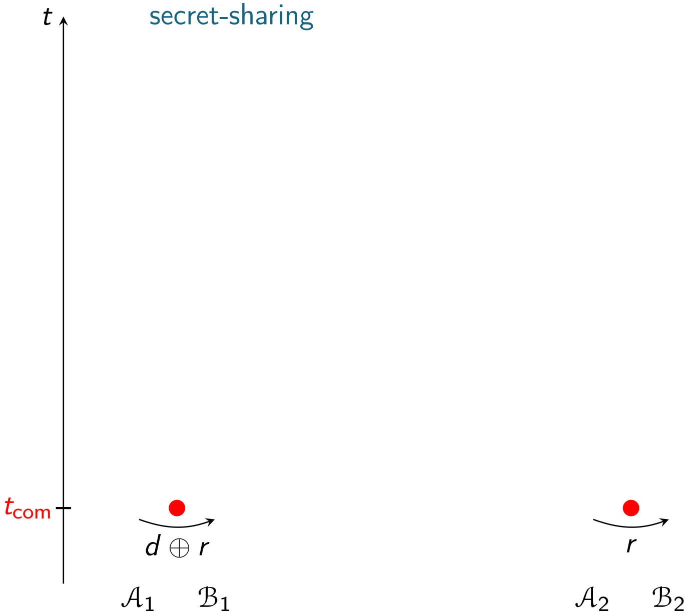

The simplest relativistic set-up

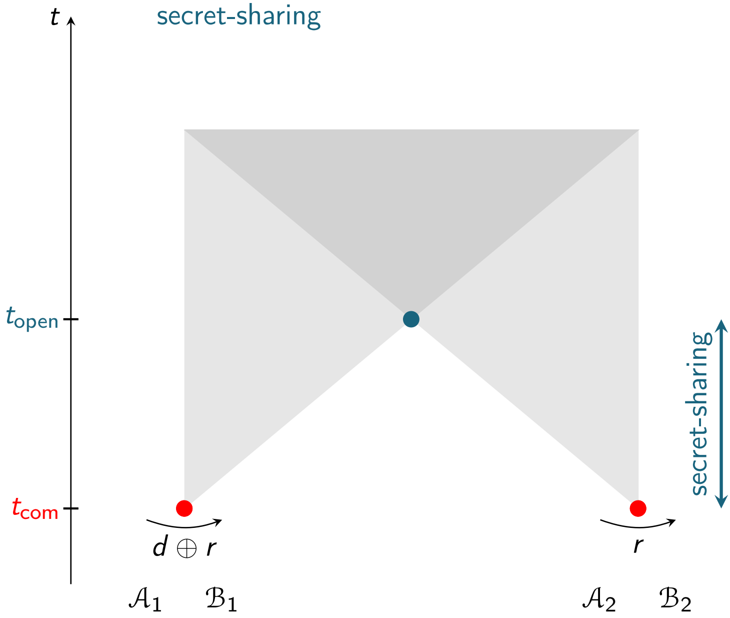

The simplest relativistic set-up

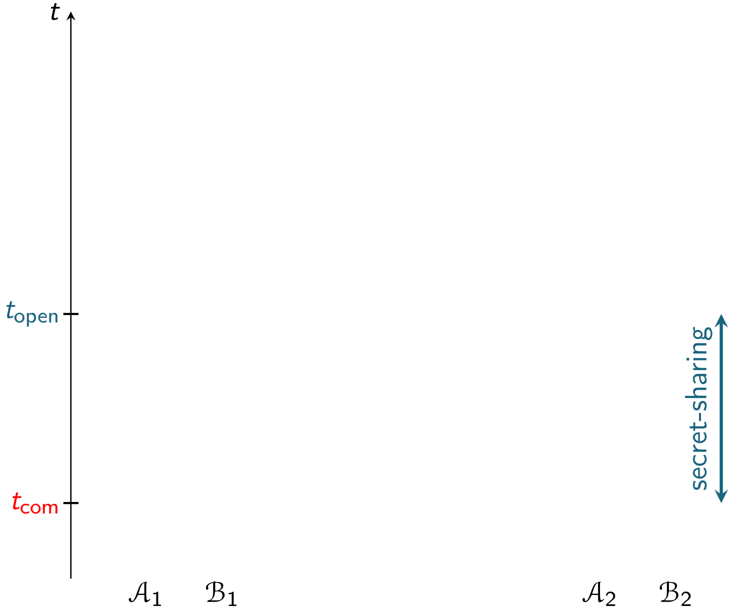

The simplest relativistic set-up

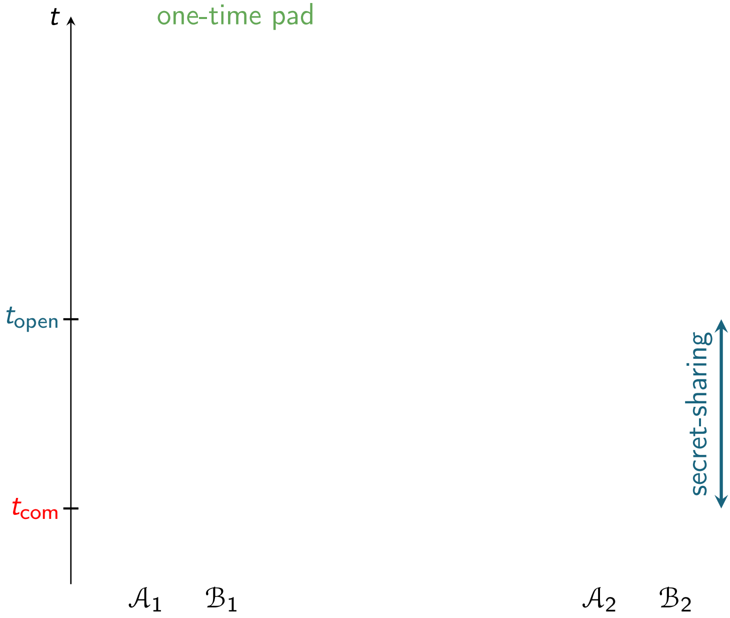

The simplest relativistic set-up

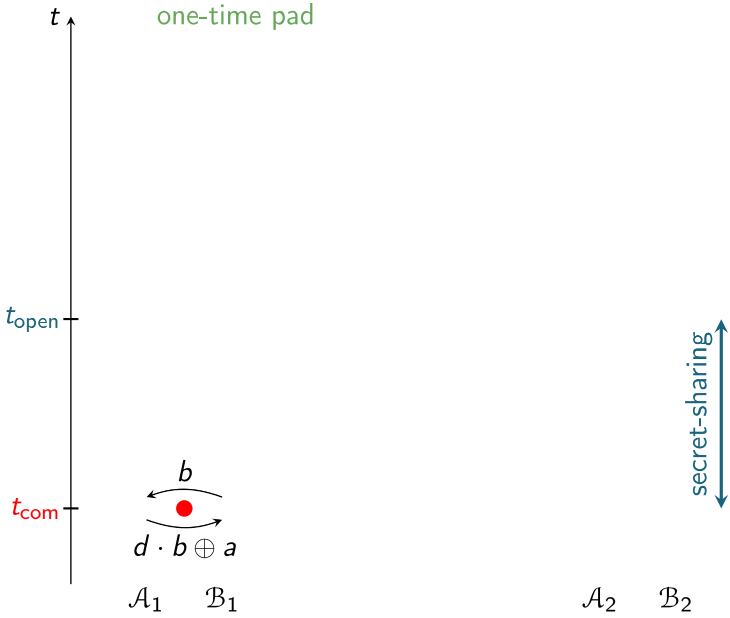

The simplest relativistic set-up

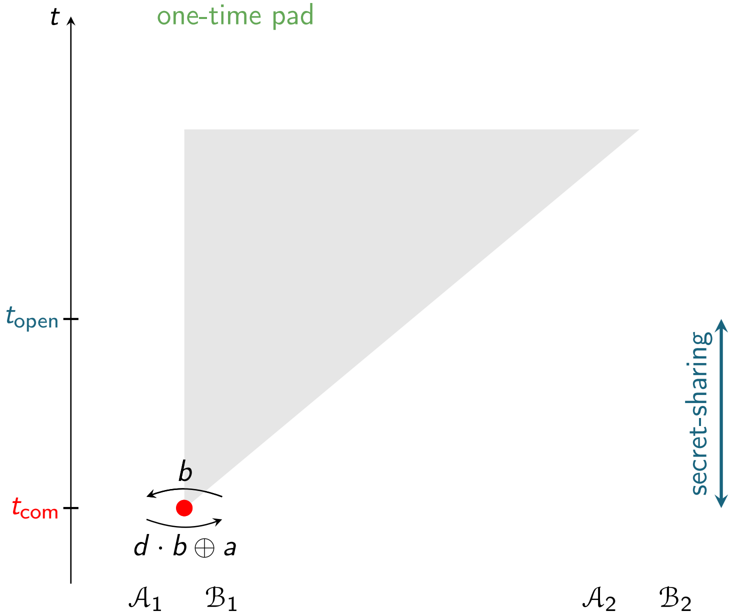

The simplest relativistic set-up

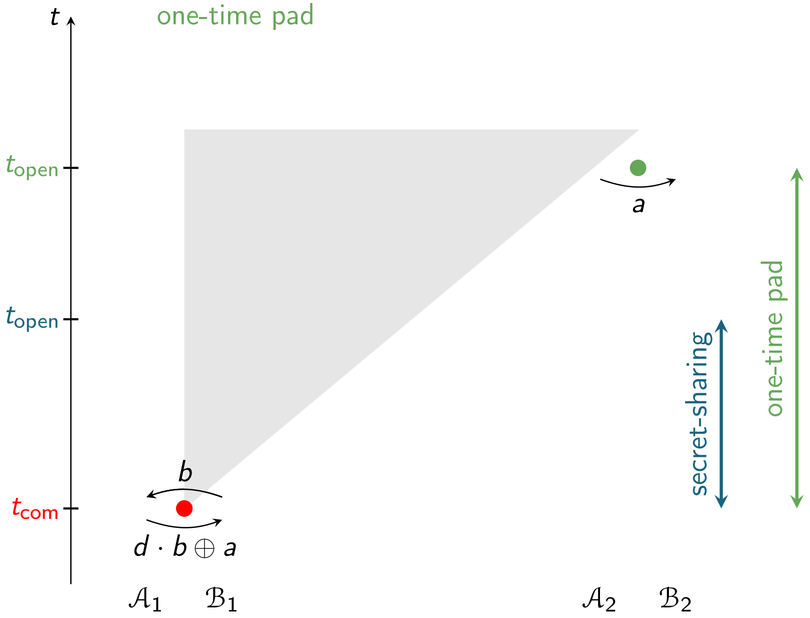

The simplest relativistic set-up

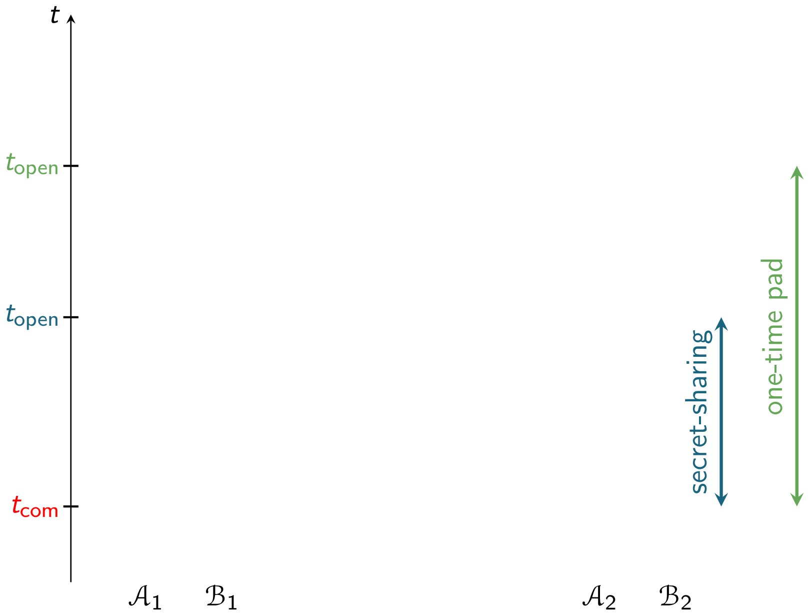

The simplest relativistic set-up

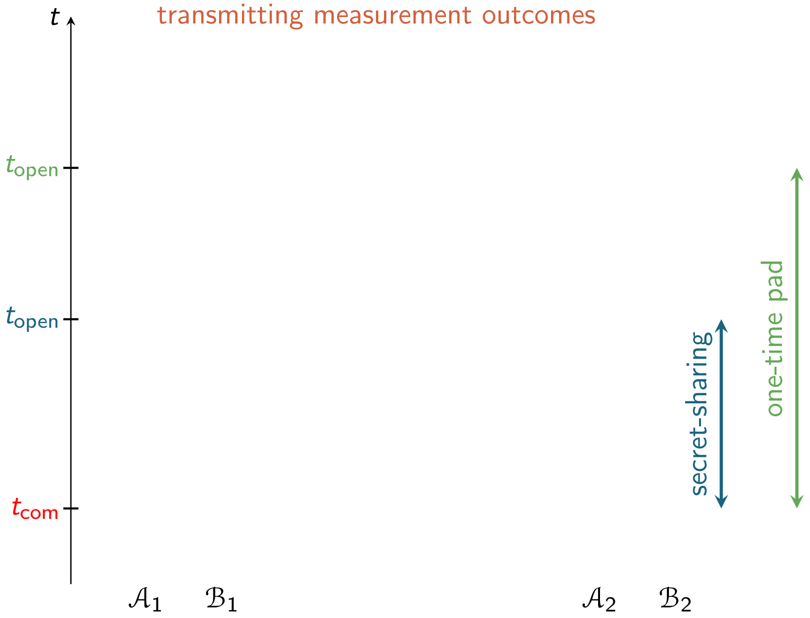

The simplest relativistic set-up

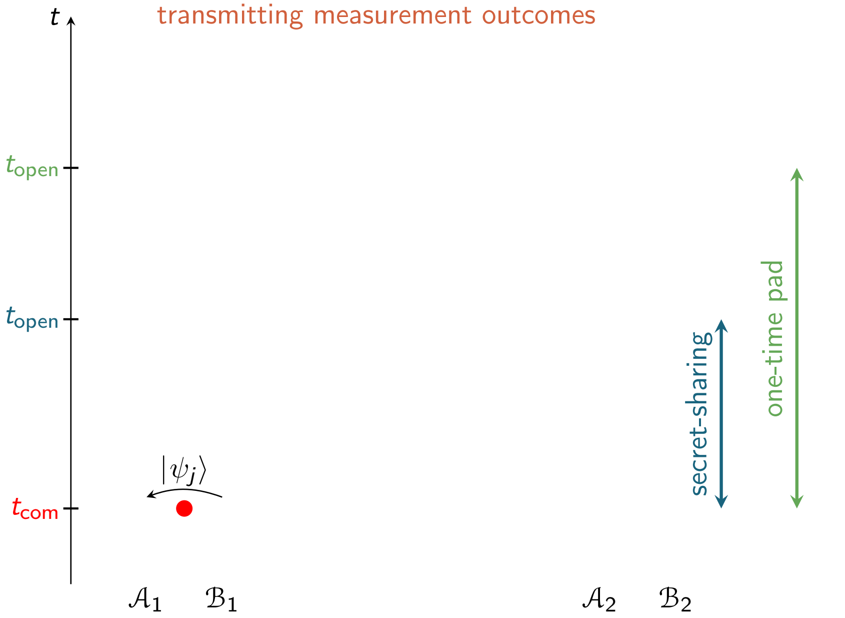

The simplest relativistic set-up

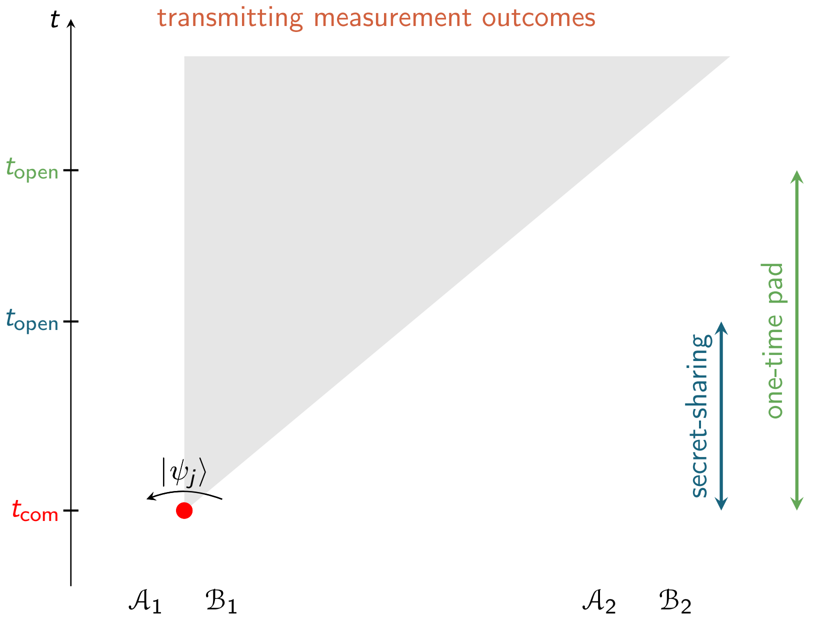

The simplest relativistic set-up

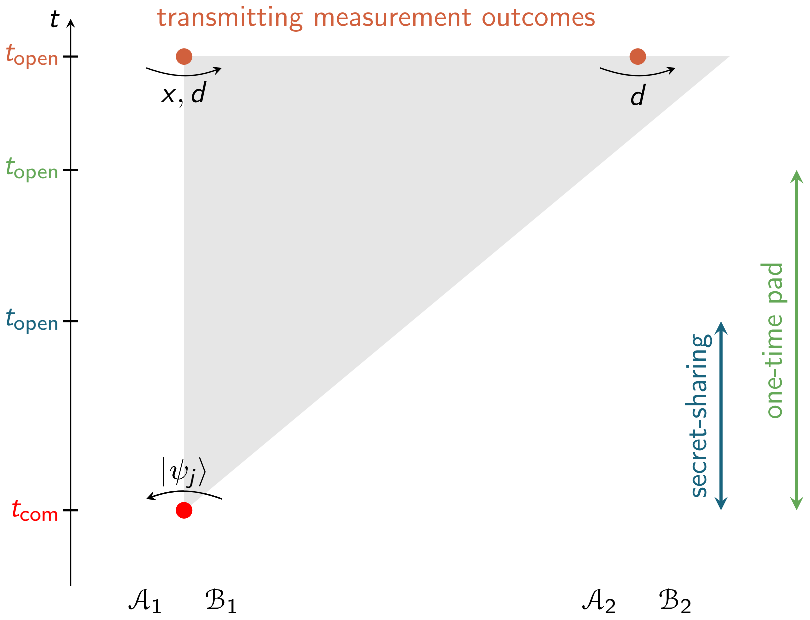

The simplest relativistic set-up

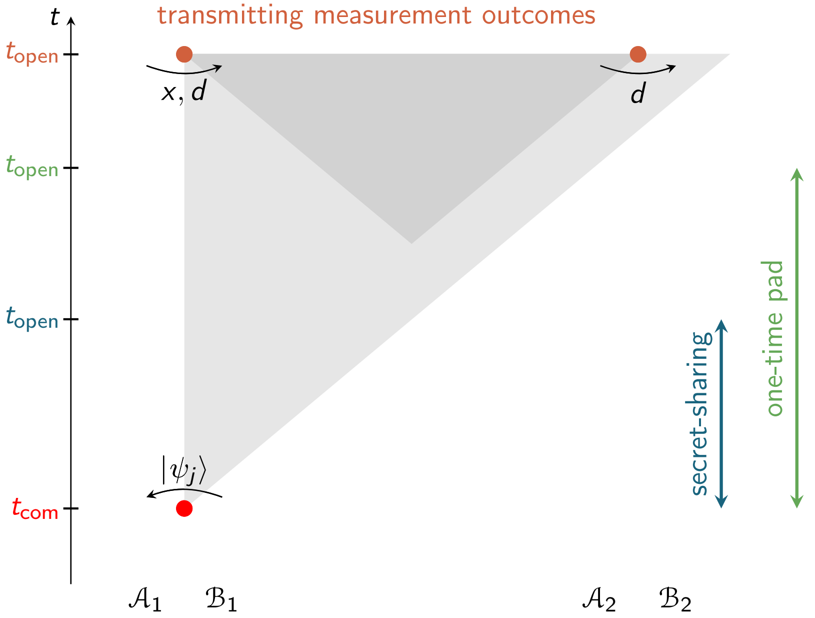

The simplest relativistic set-up

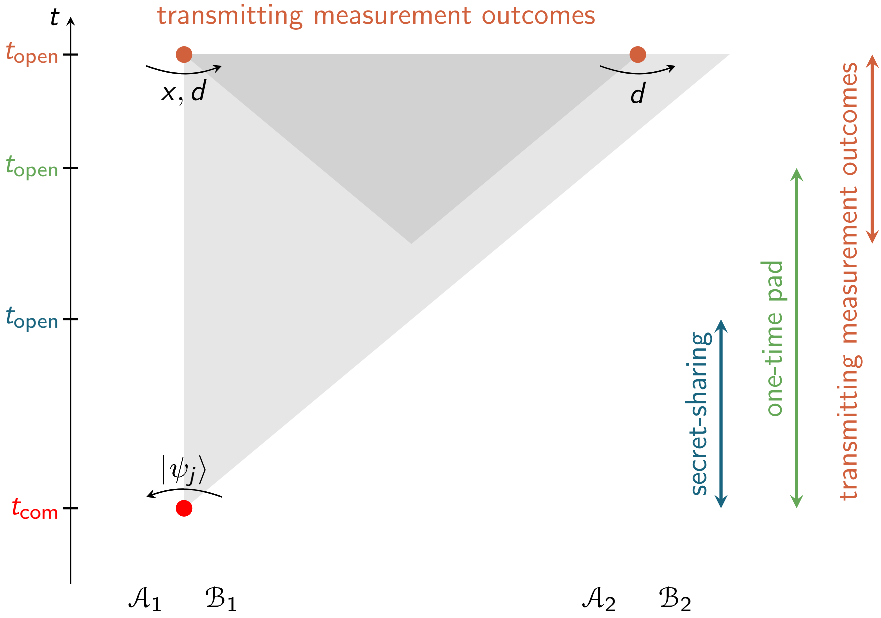

Timed commitments

- **Secret-sharing BC** essentially "opens itself".
- **One-time pad BC** must be opened before a certain time, after that it expires without revealing any information.
- **BC by transmitting measurement outcomes** can be opened any time but the commitment is only valid for a fixed period before the opening.

In the relativistic scenario nothing can be **permanently** secure...
It is not clear how powerful these primitives are...

Can we increase the commitment time by requiring multiple rounds of communication?

Relativistic scenario – a closer look

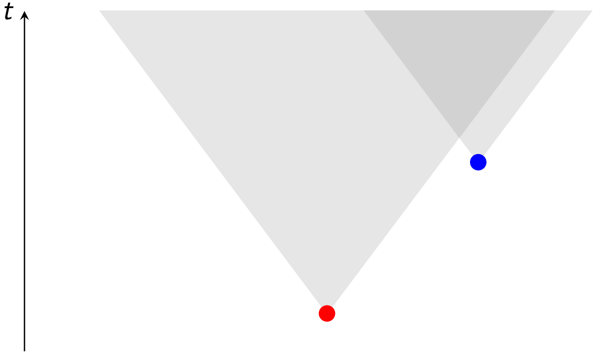

Relativistic scenario – a closer look

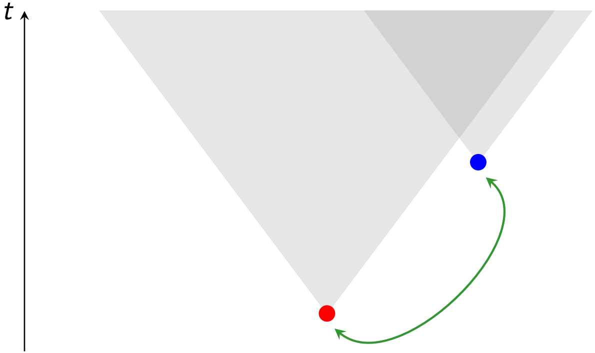

space-like separated $\Longrightarrow$ no communication

Relativistic scenario – a closer look

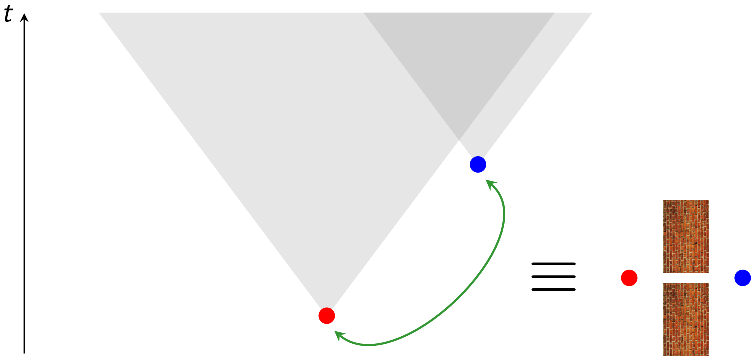

space-like separated $\Longrightarrow$ no communication

Relativistic scenario – a closer look

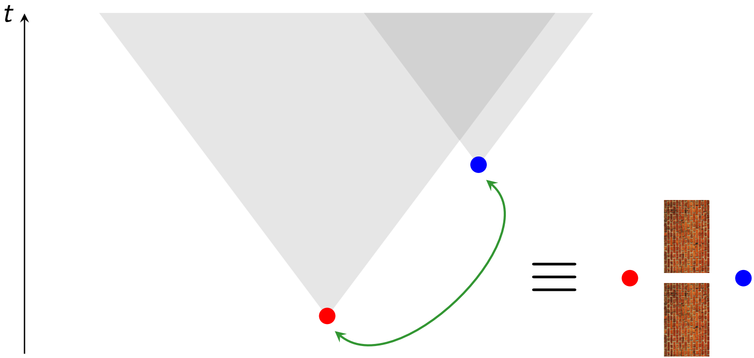

space-like separated $\implies$ no communication

For two rounds (classical or quantum)

Relativistic $\equiv$ Two isolated provers

$\implies$ well-studied and understood

Relativistic scenario – a closer look

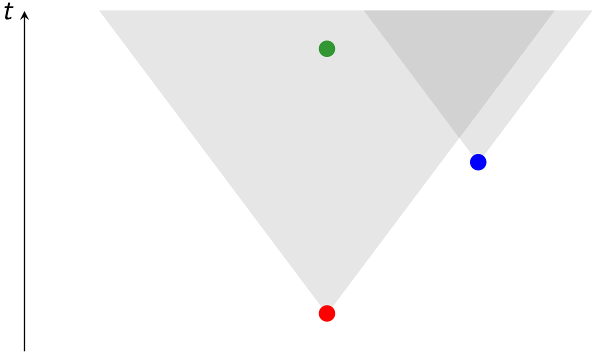

Relativistic scenario – a closer look

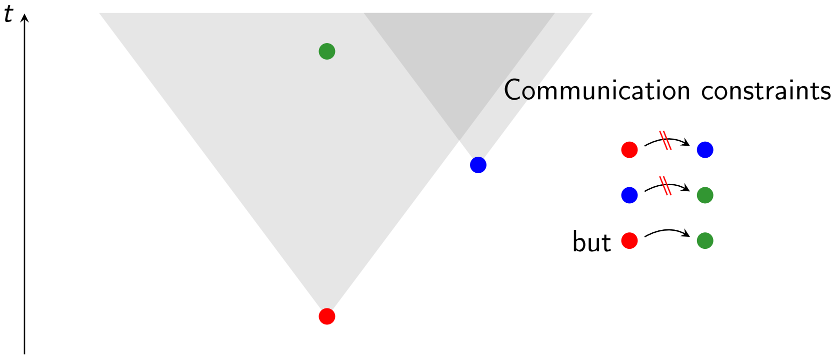

Relativistic scenario – a closer look

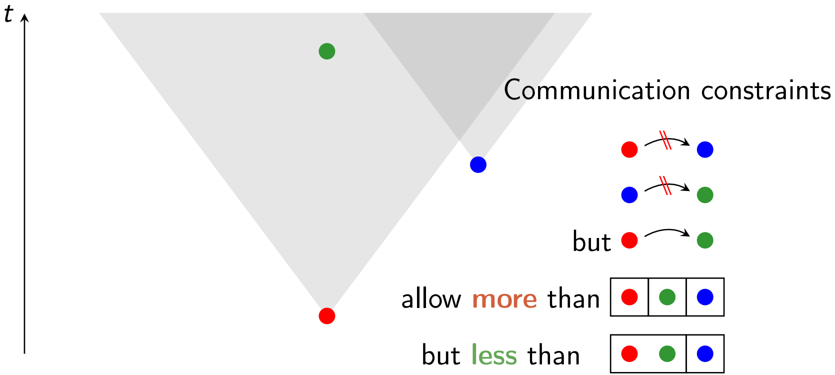

Relativistic scenario – a closer look

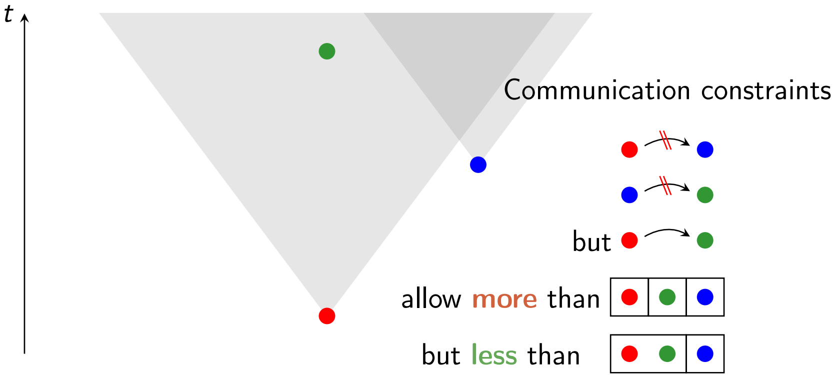

No **simple** description in terms
of non-communication models...
Security analysis likely to be **hard...**

A new multi-round protocol [Lunghi et al.]

$a_k, b_k \in_R \{0,1\}^n$
consecutive rounds must
be **space-like** separated

A new multi-round protocol [Lunghi et al.]

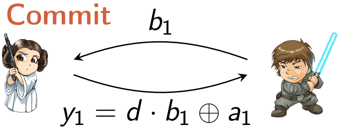

$a_k, b_k \in_R \{0,1\}^n$
consecutive rounds must be **space-like** separated

A new multi-round protocol [Lunghi et al.]

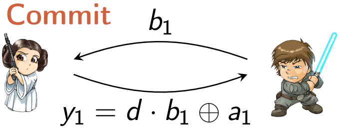

$a_k, b_k \in_R \{0,1\}^n$

consecutive rounds must be space-like separated

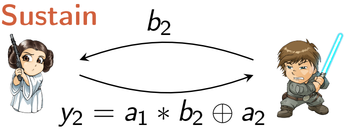

A new multi-round protocol [Lunghi et al.]

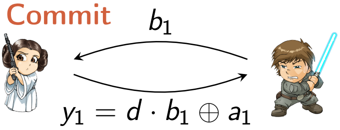

$a_k, b_k \in_R \{0,1\}^n$
consecutive rounds must be **space-like** separated

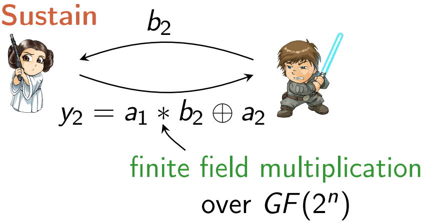

A new multi-round protocol [Lunghi et al.]

Commit

b_1

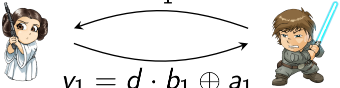

$y_1 = d \cdot b_1 \oplus a_1$

⋮

b_m

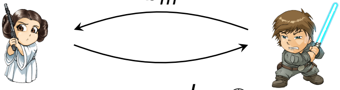

$y_m = a_{m-1} * b_m \oplus a_m$

$a_k, b_k \in_R \{0,1\}^n$

consecutive rounds must be **space-like** separated

Sustain

b_2

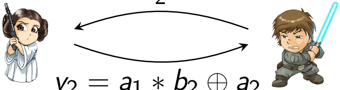

$y_2 = a_1 * b_2 \oplus a_2$

finite field multiplication over $GF(2^n)$

⋮

A new multi-round protocol [Lunghi et al.]

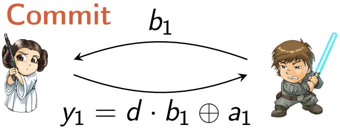

⋮

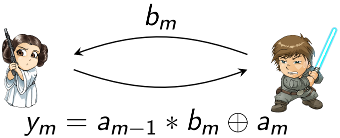

$a_k, b_k \in_R \{0,1\}^n$
consecutive rounds must be **space-like** separated

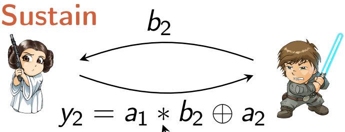

finite field multiplication over $GF(2^n)$

⋮

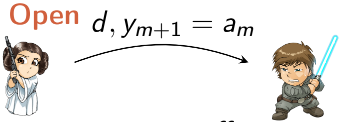

accept iff
$V(d, b_1, y_1, \ldots, b_m, y_m, y_{m+1}) = 1$

acceptance predicate

A new multi-round protocol [Lunghi et al.]

$a_k, b_k \in_R \{0,1\}^n$

consecutive rounds must be **space-like** separated

Commit b_1

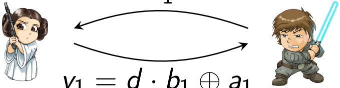

$y_1 = d \cdot b_1 \oplus a_1$

Sustain b_2

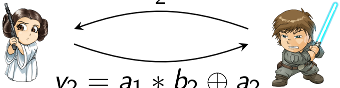

$y_2 = a_1 * b_2 \oplus a_2$

finite field multiplication over $GF(2^n)$

⋮

b_m

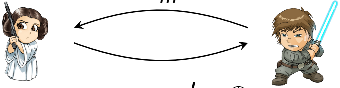

$y_m = a_{m-1} * b_m \oplus a_m$

Security for **honest Alice** guaranteed by the XOR

⋮

Open $d, y_{m+1} = a_m$

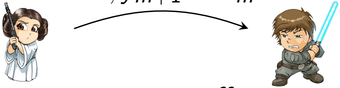

accept iff

$V(d, b_1, y_1, \ldots, b_m, y_m, y_{m+1}) = 1$

acceptance predicate

A new multi-round protocol [Lunghi et al.]

Commit

b_1

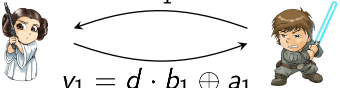

$y_1 = d \cdot b_1 \oplus a_1$

$\vdots$

b_m

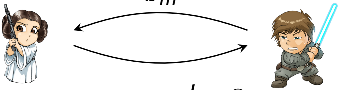

$y_m = a_{m-1} * b_m \oplus a_m$

Security for **honest Alice** guaranted by the XOR

Security for **honest Bob** more complicated...

$a_k, b_k \in_R \{0, 1\}^n$

consecutive rounds must be **space-like** separated

Sustain

b_2

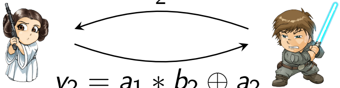

$y_2 = a_1 * b_2 \oplus a_2$

finite field multiplication over $GF(2^n)$

$\vdots$

Open $d, y_{m+1} = a_m$

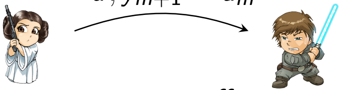

accept iff

$V(d, b_1, y_1, \ldots, b_m, y_m, y_{m+1}) = 1$

acceptance predicate

A new multi-round protocol – honest Bob

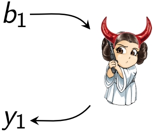

A new multi-round protocol – honest Bob

A new multi-round protocol – honest Bob

A new multi-round protocol – honest Bob

A new multi-round protocol – honest Bob

Non-trivial causal constraints make the analysis very hard...
Classically: shared randomness doesn't help; deterministic strategies "flatten" the causal structure to give a multi-prover model

A new multi-round protocol – honest Bob

A new multi-round protocol – honest Bob

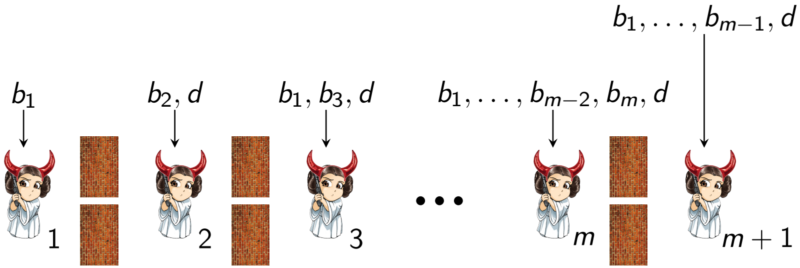

A new multi-round protocol – honest Bob

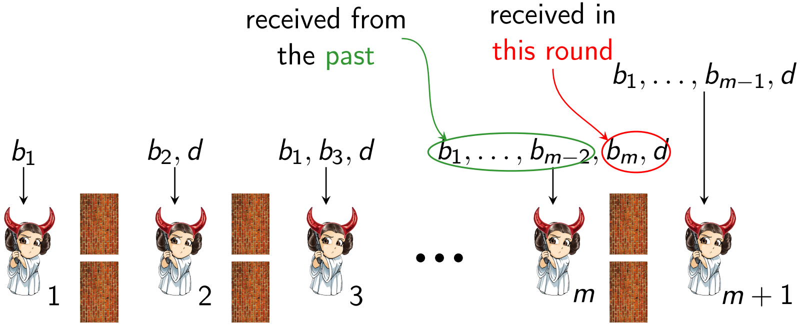

A new multi-round protocol – honest Bob

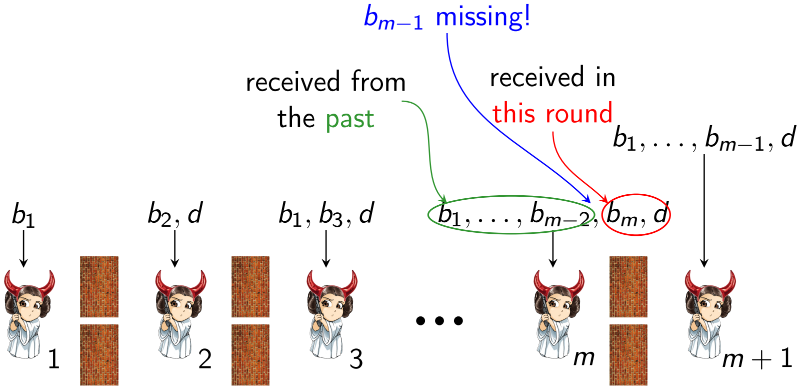

A new multi-round protocol – honest Bob

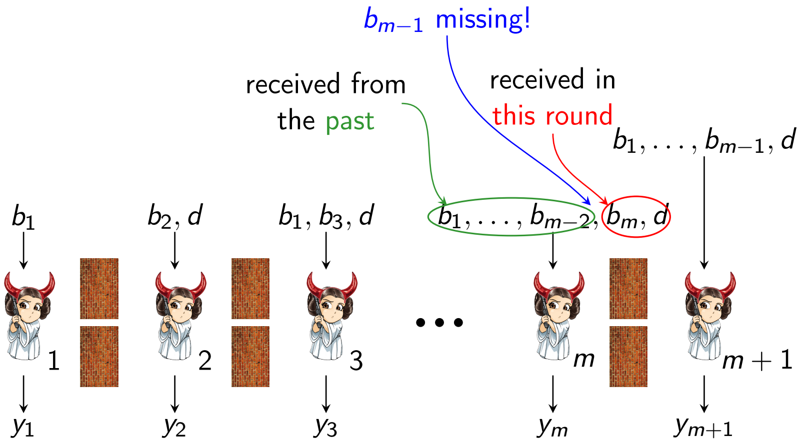

check whether $V(d, b_1, y_1, \ldots, b_m, y_m, y_{m+1}) = 1$

A new multi-round protocol – honest Bob

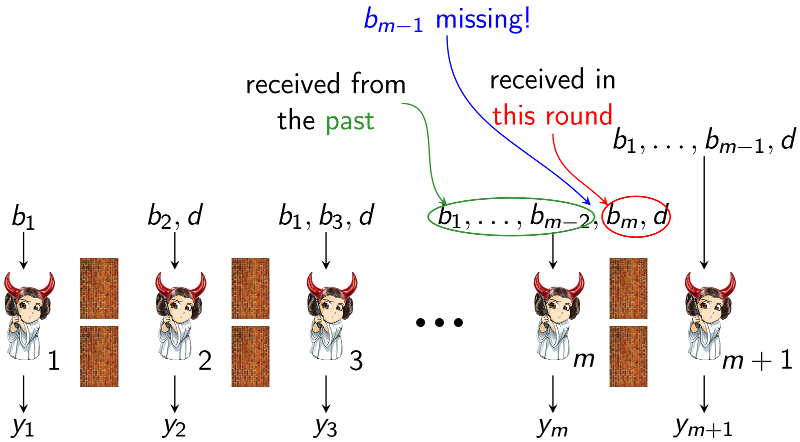

check whether $V(d, b_1, y_1, \ldots, b_m, y_m, y_{m+1}) = 1$

this reduction is **exact** – same optimal winning probability

A new multi-round protocol – honest Bob

Conclusions:

- End up with a **complicated** game of $m+1$ **non-communicating** players; exact cheating probability is hard to calculate.
- Can be relaxed to a very simple-looking problem of computing a certain function in the **"Number on the Forehead"** model. For $m=2$ it is exactly the finite-field generalisation of CHSH.
- Equivalent to counting the **number of zeroes** of a certain family of **multivariate polynomials** over finite field $GF(2^n)$.

A new multi-round protocol – honest Bob

Final result: Security for honest Bob with $\varepsilon \approx 2^{-n/2^m}$.

- Security **deteriorates drastically** as m increases.

A new multi-round protocol – honest Bob

Final result: Security for honest Bob with $\varepsilon \approx 2^{-n/2^m}$.

- Security **deteriorates drastically** as m increases.
- In **principle**, an arbitrary long commitment is possible (at the price of very large n).

A new multi-round protocol – honest Bob

Final result: Security for honest Bob with $\varepsilon \approx 2^{-n/2^m}$.

- Security **deteriorates drastically** as m increases.
- In **principle**, an arbitrary long commitment is possible (at the price of very large n).
- In **practice**, technology puts a limit on n so the commitment time is limited.

A new multi-round protocol – honest Bob

Final result: Security for honest Bob with $\varepsilon \approx 2^{-n/2^m}$.

- Security **deteriorates drastically** as m increases.
- In **principle**, an arbitrary long commitment is possible (at the price of very large n).
- In **practice**, technology puts a limit on n so the commitment time is limited.
- Looks very similar to **communication complexity lower bounds** for this model: $\Omega(\frac{n}{2^m})$.

Thanks for you attention!

Finite-field, multiprover generalisation of CHSH

$\mathbb{F}_q$ – finite field of size q, X_1, X_2 drawn uniformly at random. What are the best local functions that simulate the $X_1 X_2$ (can we argue that this is the "hardest" function to simulate?), i.e. we are trying to maximise

$$\Pr[X_1 X_2 = f_1(X_1) + f_2(X_2)].$$

Trivial strategy gives $\frac{1}{q}$, some probabilistic arguments might give $\frac{\log q}{q}$ but by connecting it to some algebraic geometry problem one can show that there exists strategy that achieves $\Omega(q^{-2/3})$ (see Bavarian and Shor). Unfortunately, no explicit strategies are known.
This is exactly what we get for $m = 2$, for more we are trying to satisfy

$$\prod_{k=1}^{m} X_k = \sum_{k=1}^{m} f_k(X_{[m]\setminus\{k\}}),$$

which is the number on the forehead model.