Must the Communication Graph of MPC Protocols be an Expander?

Elette Boyle (IDC)
Ran Cohen (MIT & Northeastern)
Deepesh Data (UCLA)
Pavel Hubacek (Charles University)

Secure Multiparty Computation

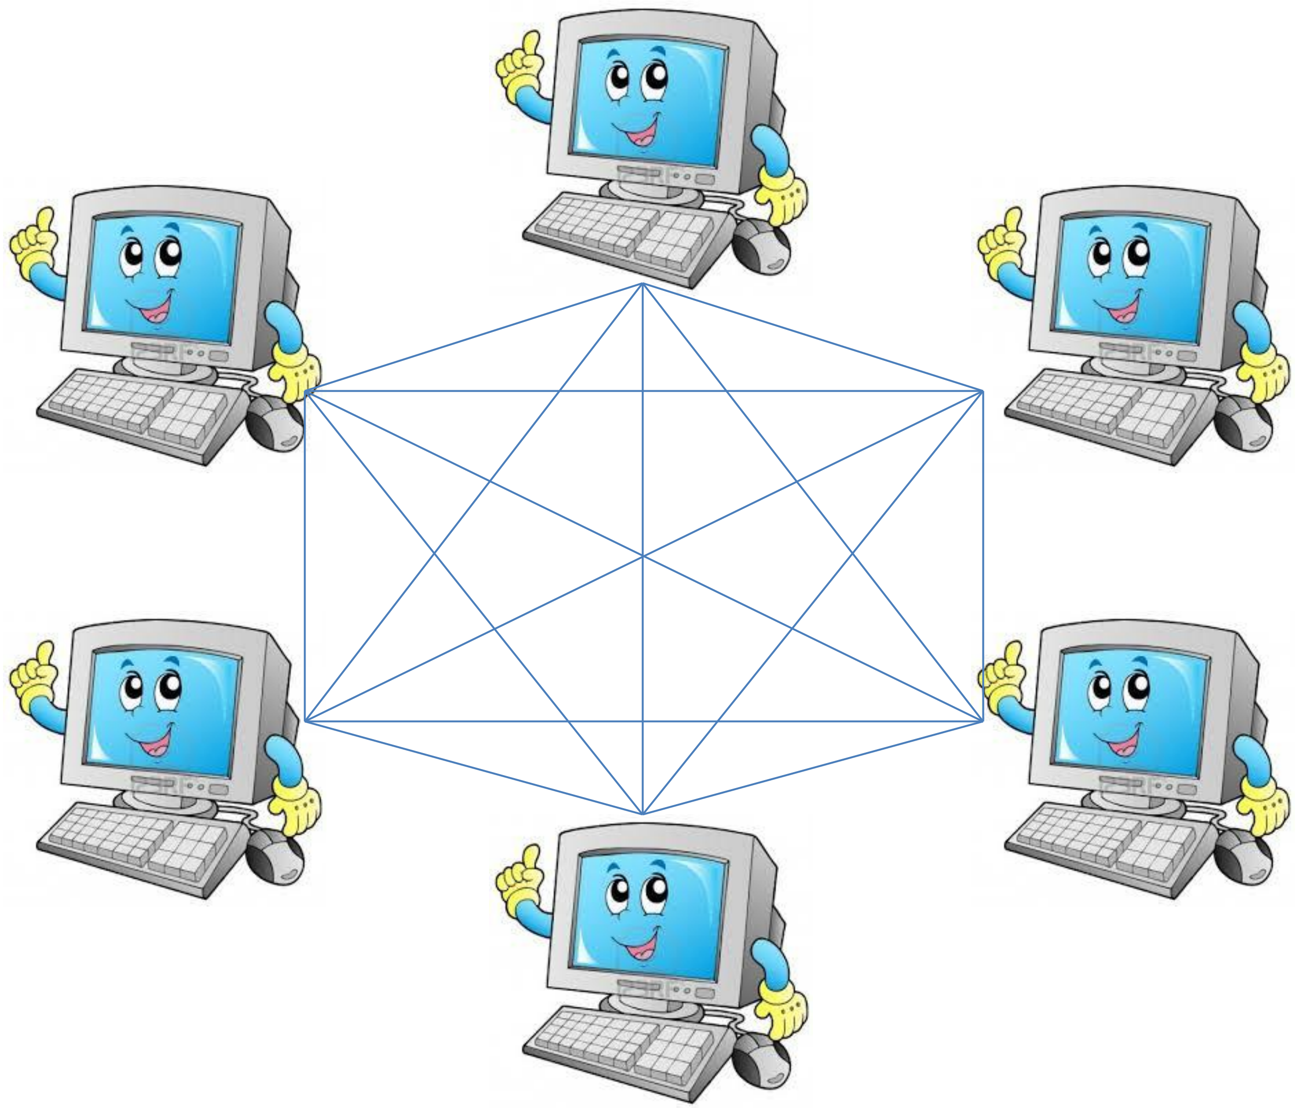

Classical Results

- Byzantine Agreement
  - [Pease, Shostak, Lamport'80]
  - [Lamport, Shostak, Pease'82]
  - [Dolev, Strong'83]
  - [Feldman, Micali'88]
  - [Garay, Moses'93]
  - …
- Secure Function Evaluation
  - [Yao'82/86]
  - [Goldreich, Micali, Wigderson'87]
  - [Ben-Or, Goldwasser, Wigderson'88]
  - [Chaum, Crepeau, Damgard'88]
  - [Rabin, Ben-Or'89]
  - …

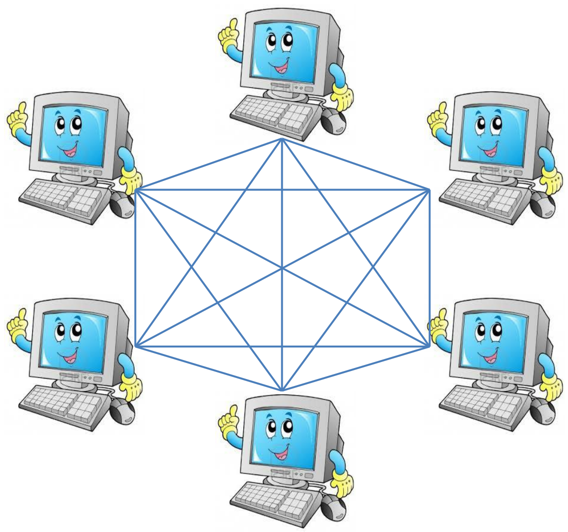

Everyone talks to everyone

Complete communication graph

Large-Scale MPC

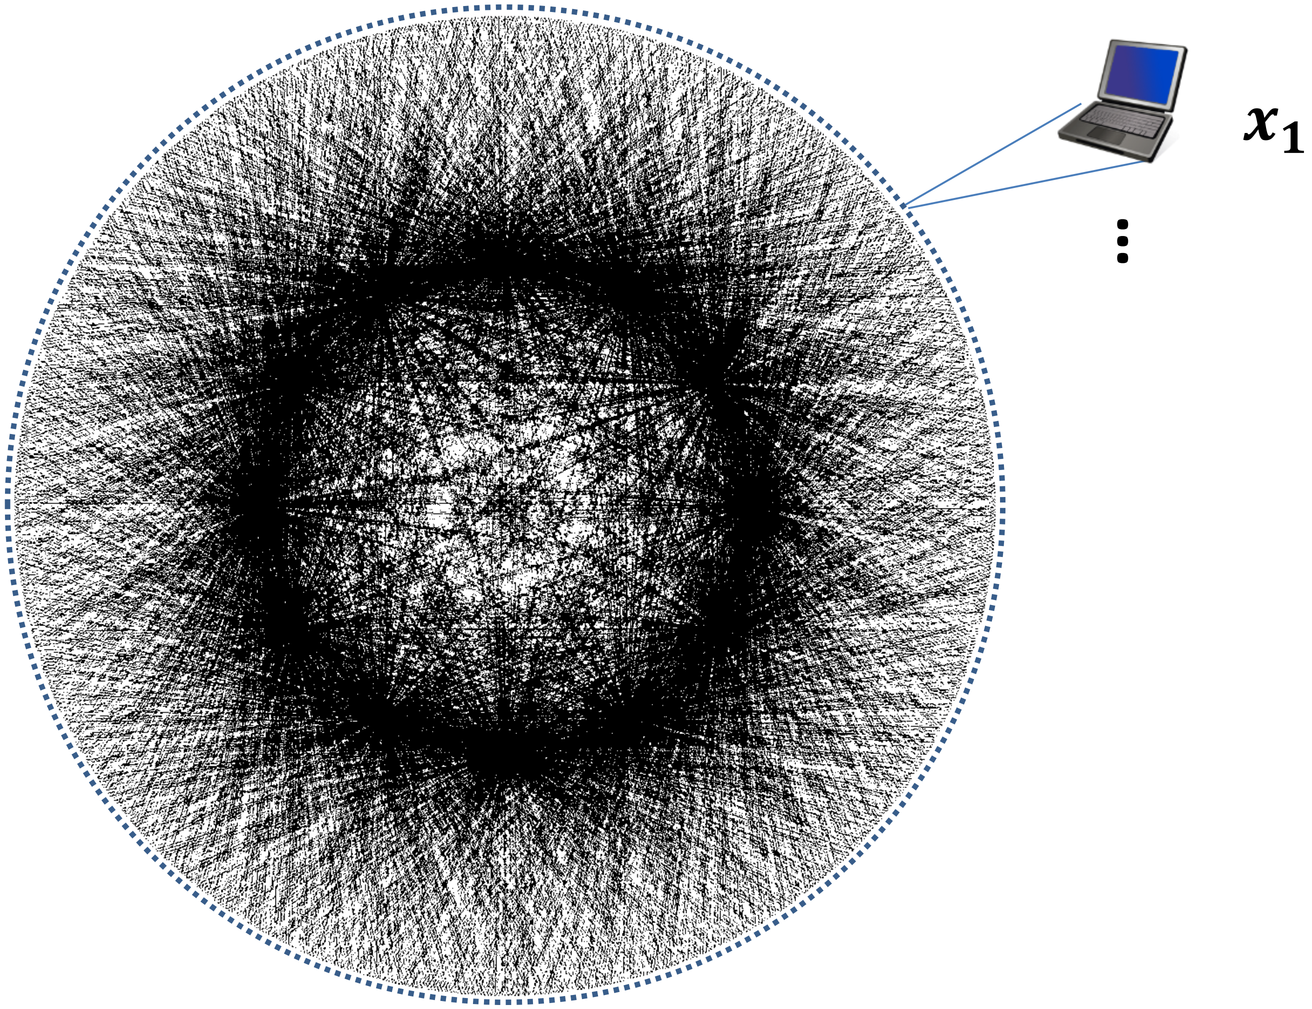

Can we use a ***sparse graph***?

Model #1: Fixed Partial Graph

The graph known ahead of time

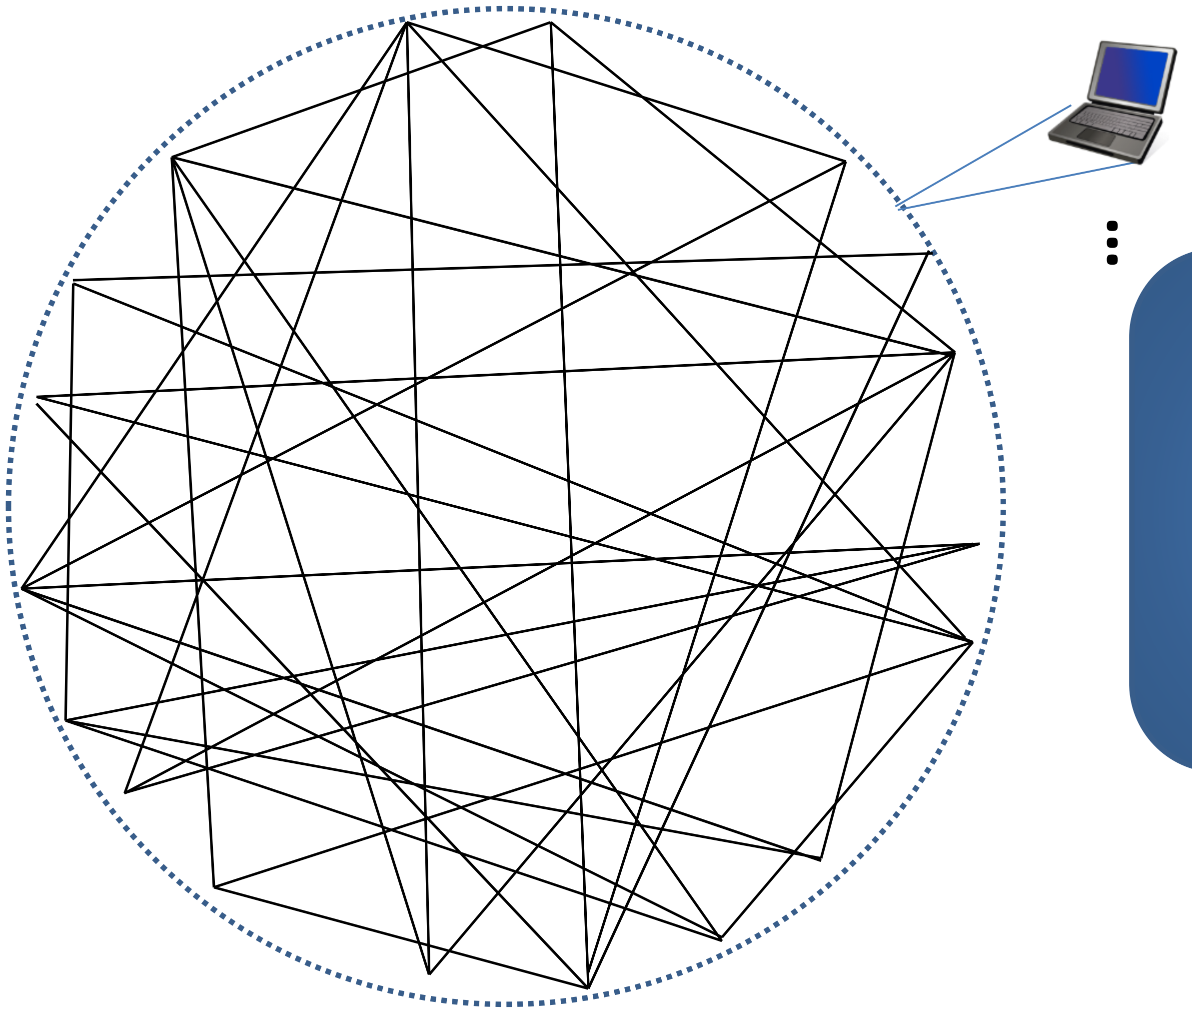

Corruptions based on the graph

Model #1: Fixed Partial Graph

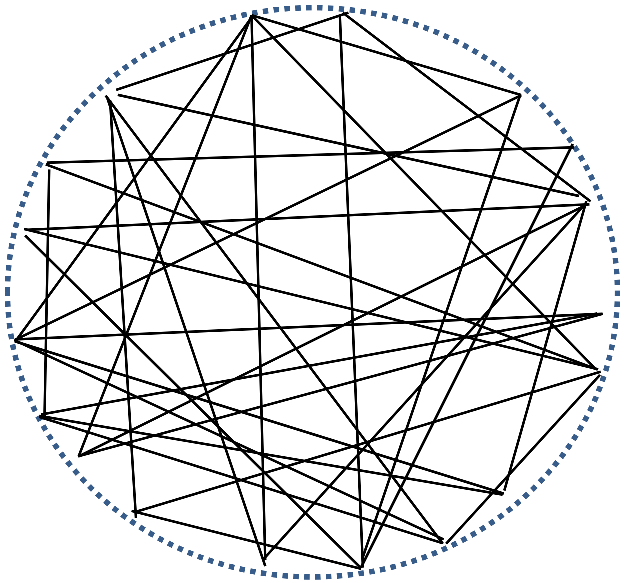

- **Lower bounds** for BA
  - Connectivity $t+1$ (without setup $2t+1$) [Dolev'82] [Fischer, Lynch, Merritt'85]
  - Comm. complexity $\Omega(n^2)$ [Dolev, Reischuk'82]
- **Weaker** correctness/privacy guarantees
  - Byzantine Agreement
    - [Dwork, Peleg, Pippenger, Upfal'86]
    - [Berman, Garay'90]
    - [Upfal'92]
  - Secure Function Evaluation
    - [Beimel'07] [Garay, Ostrovsky'08] [Halevi, Lindell, Pinkas'11]
    - [Chandran, Garay, Ostrovsky'12]
    - [Halevi, Ishai, Jain, Kushilevitz, Rabin'16]

Model #2: Dynamic Partial Graph

Everyone *can* talk to everyone

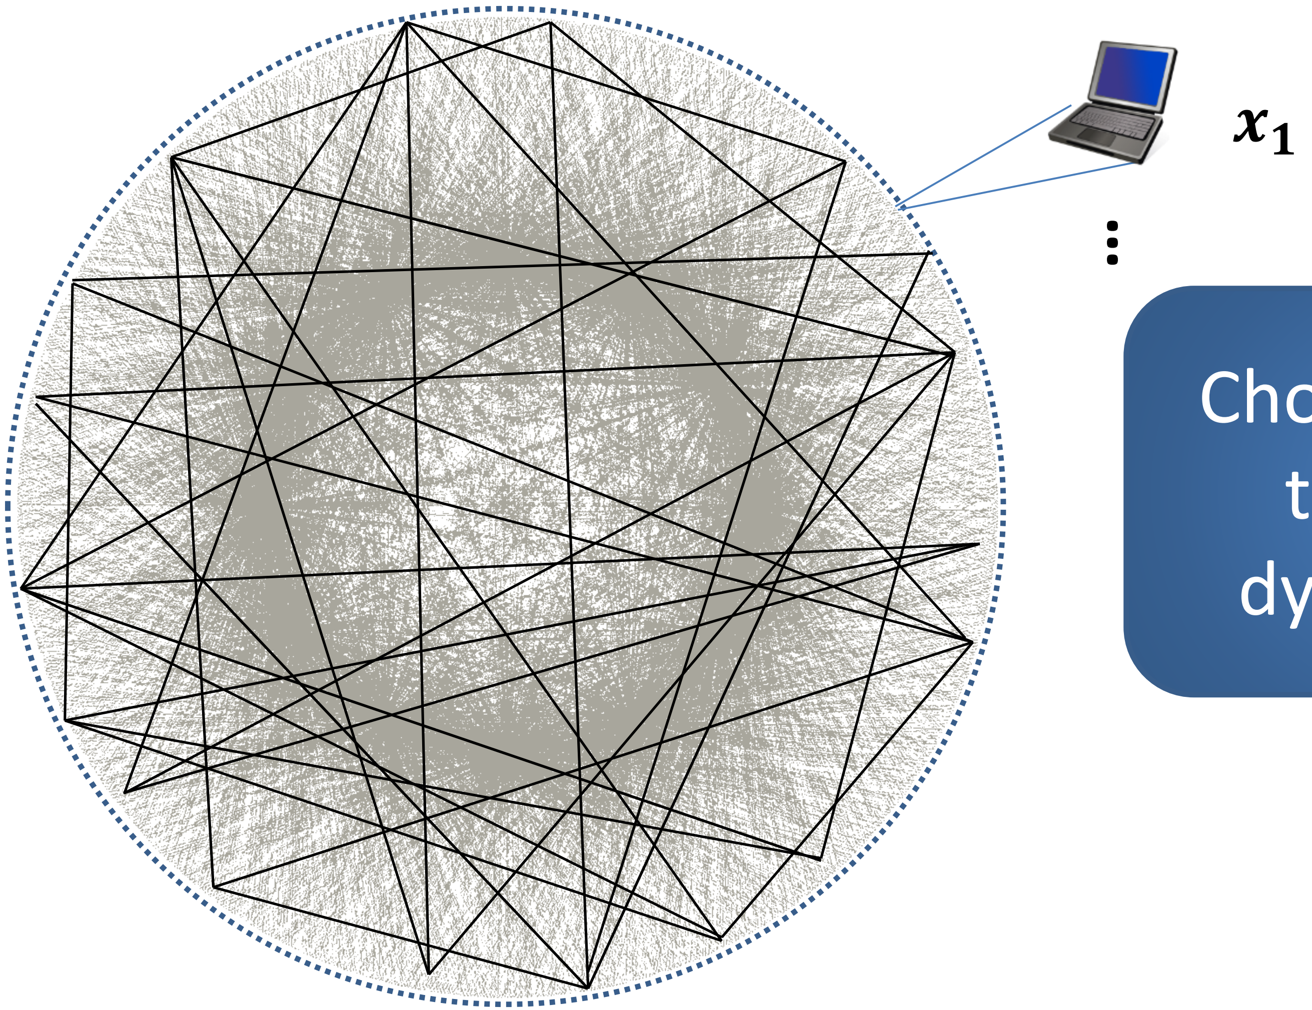

Choose whom to talk to dynamically

Model #2: Dynamic Partial Graph

- **Overcoming lower bounds** (BA)
  – Comm. complexity $\tilde{O}(n)$
- **Scalability** & low communication locality
  - ➢ Byzantine Agreement
    - [King, Saia, Sanwalani, Vee'06]
    - [Kapron, Kempe, King, Saia, Sanwalani'08]
    - [King, Saia'09] [King, Saia'10]
    - [Braud-Santoni, Guerraoui, Huc'13]
  - ➢ Secure Function Evaluation
    - [Dani, King, Movahedi, Saia'12]
    - [Boyle, Goldwasser, Tessaro'13]
    - [Chandran, Chongchitmate, Garay, Goldwasser, Ostrovsky, Zikas'15]
    - [Boyle, Chung, Pass'15]

max degree

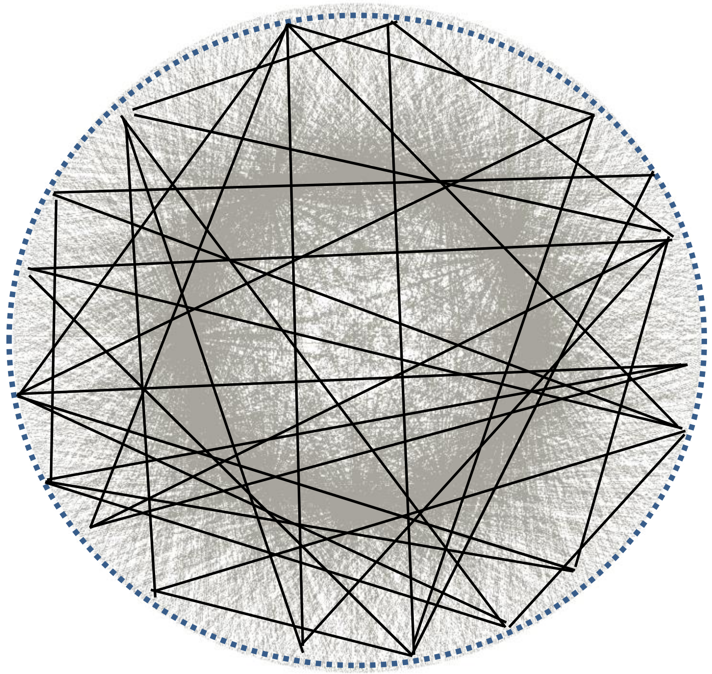

Partial Graph Models

Fixed Graph

- Strong lower bounds
  - $\Theta(n)$ connectivity
  - Comm. complexity $\Omega(n^2)$
- Well studied

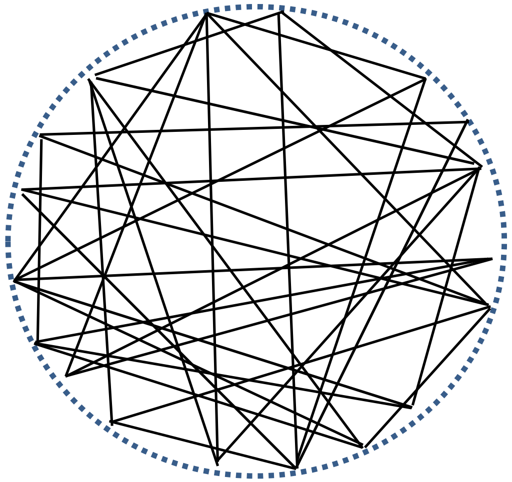

Dynamic Graph

- Overcoming lower bounds
  - Polylog locality
  - Comm. complexity $\tilde{O}(n)$
- **Less understood**

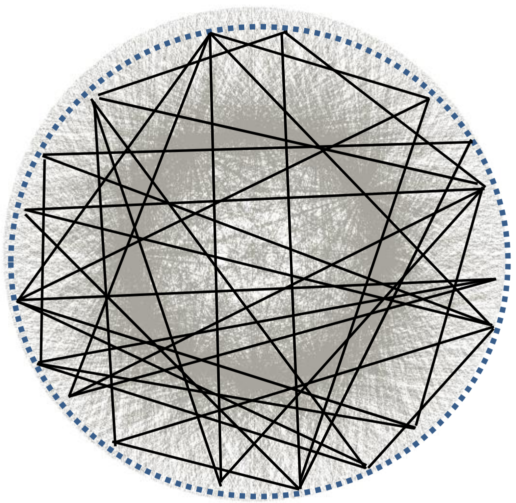

Main Question

What graph properties are
necessary
to support secure protocols?

Dynamic-Graph Model

Prior work:

This work:

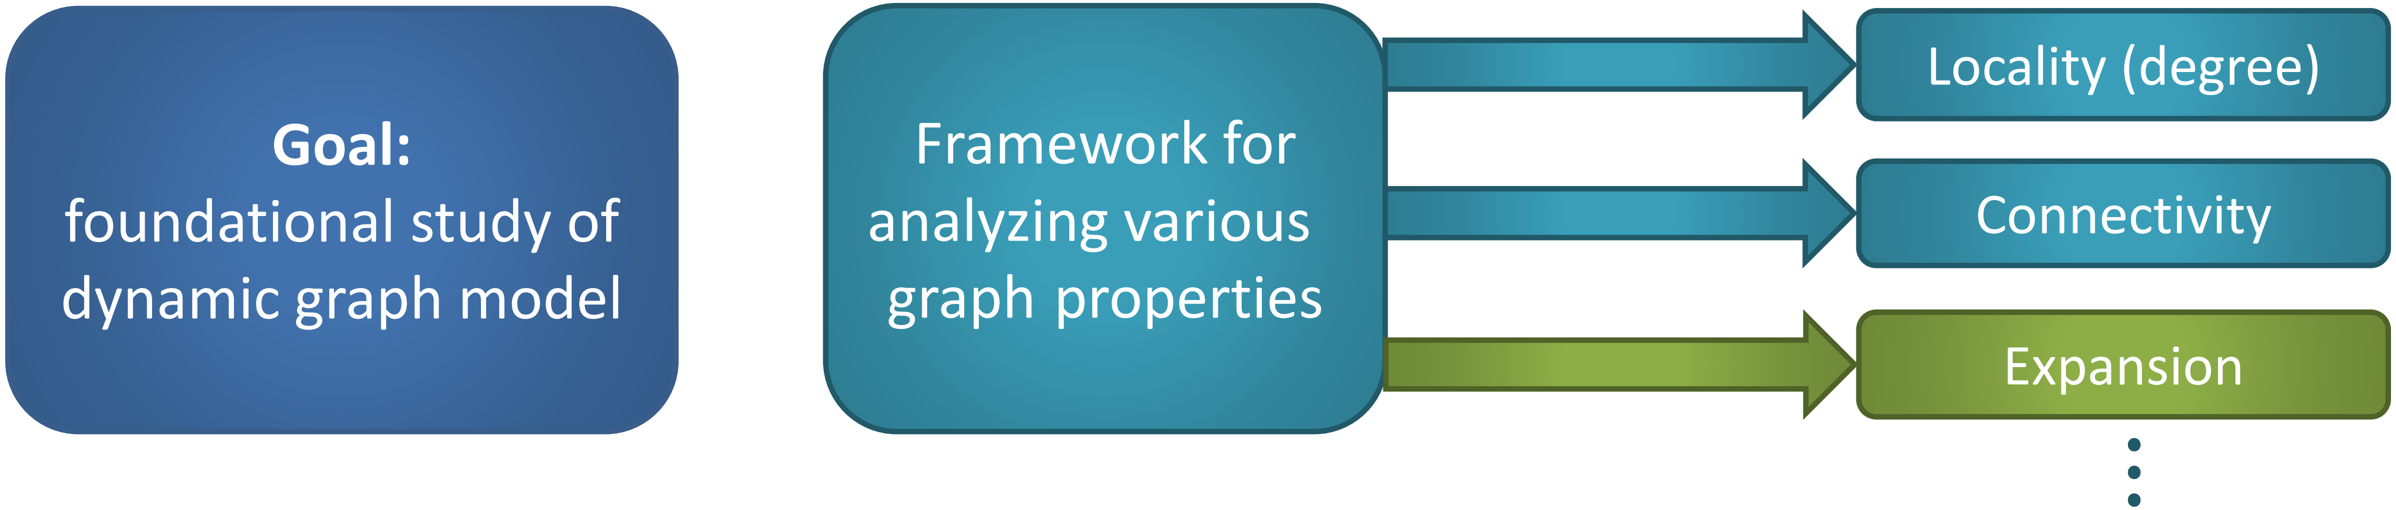

Expander Graph

"(Sparse) graph with strong connectivity properties"

All existing protocols induce expander graphs

- Classical protocols (complete graph)
- Protocols with low locality (dynamic partial graph)
  - E.g., every party randomly chooses its neighbors

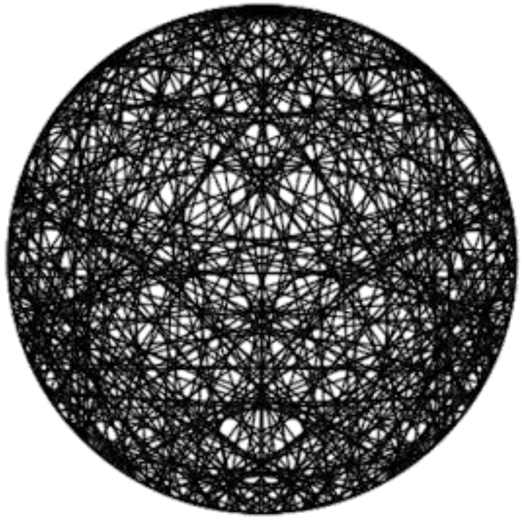

Expansion is natural (high connectivity, good mixing properties,...)

Expander Graph (2)

We focus on **edge expansion**

- Let $G = (V, E)$ be a graph of size $|V| = n$
- For every $S \subseteq V$ define $h(G, S) = \frac{|\mathrm{edges}(S, \bar{S})|}{|S|}$
- The **edge expansion ratio** of G is $h(G) = \min_{0<|S|\leq n/2} h(G, S)$
- $\{G_n\}_{n\in\mathbb{N}}$ is a **family of expander graphs** if $\exists \epsilon > 0$ s.t. $\forall n{:}\, h(G_n) \geq \epsilon$

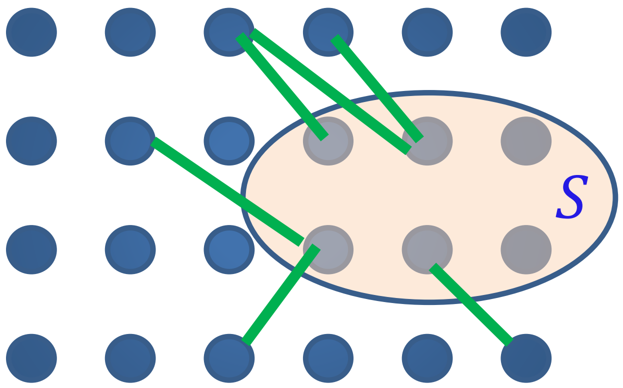

Example of Non-Expander Graph

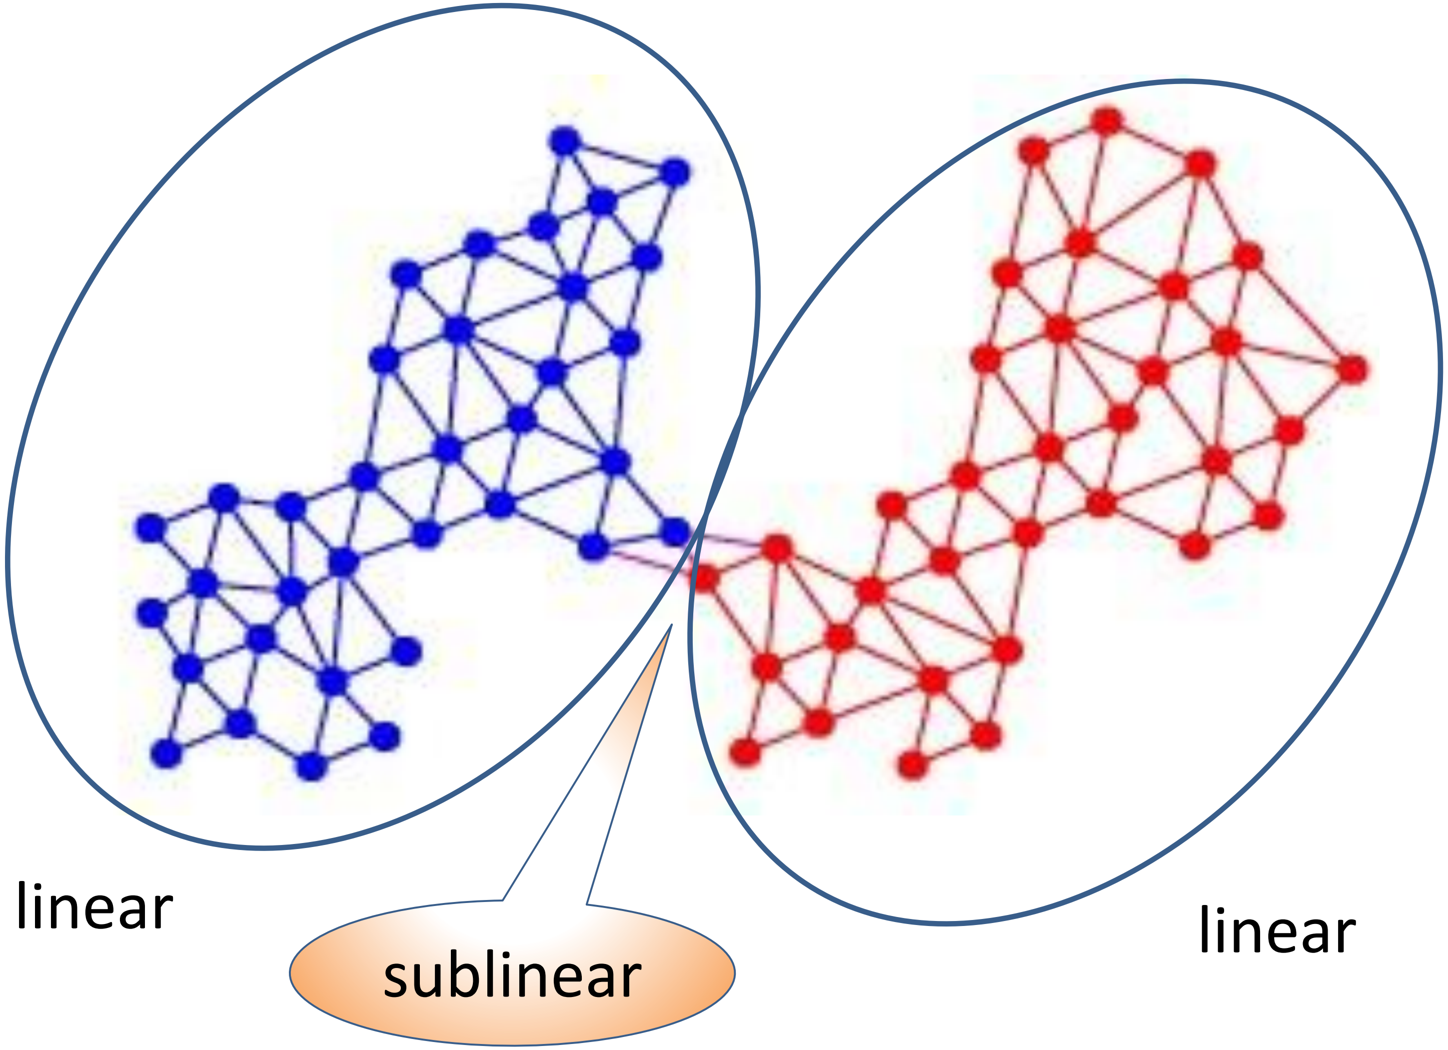

More Focused Question

Must the comm. graph of MPC protocols
(tolerating linear corruptions)
be an ***expander***?

It depends…

Main Results

Upper bound:

SFE protocols with **non-expander** graph (in **PKI model**):

- Static/adaptive corruptions
- Information-theoretic/computational security
- With/out polylog locality

Lower bound:

$\exists f$ s.t. every secure protocol for f induces an **expander**

- Adaptive corruptions, **CRS model**

Upper Bound:
Non-Expander Protocols

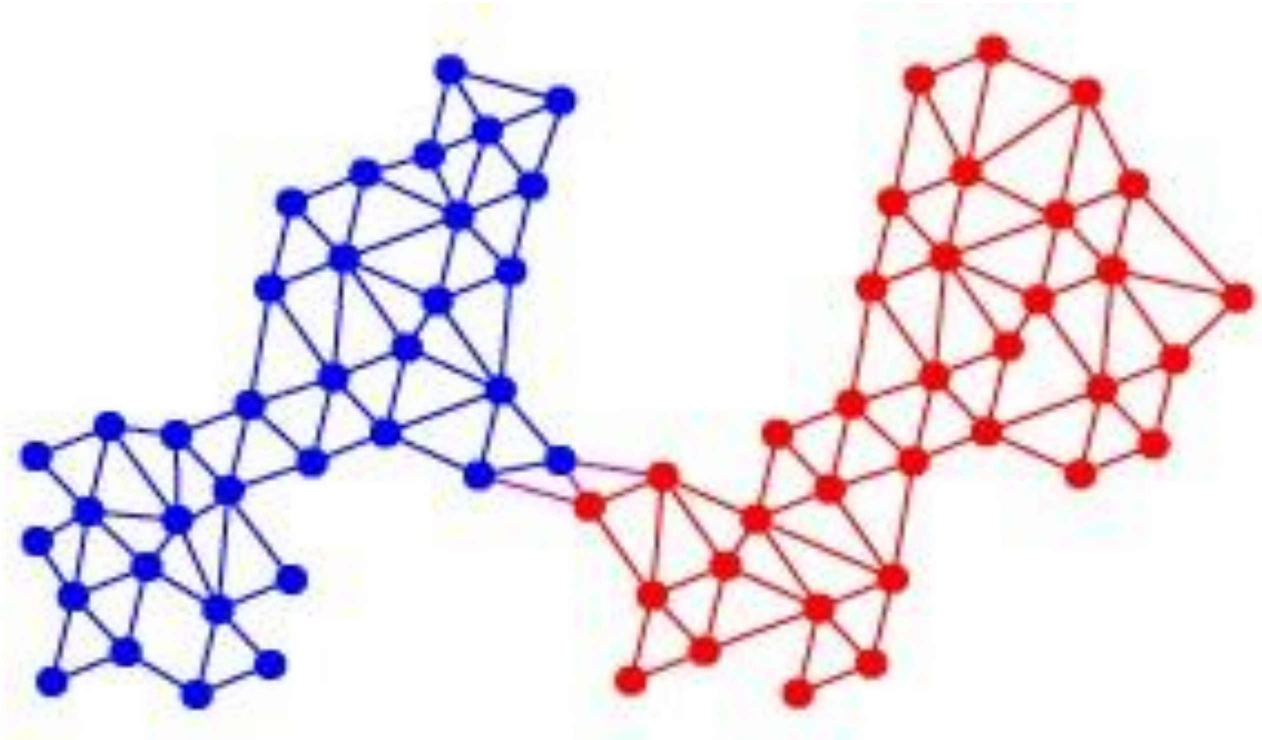

Theorem (Upper Bound)

Let f be n-party function and assume **digital signatures** exist

Then, $\exists$ protocol π in the **PKI model** such that

- π computes f tolerating $(1/4 - \epsilon)n$ static corruptions
- The communication graph of π is **not an expander**

Protocol Template

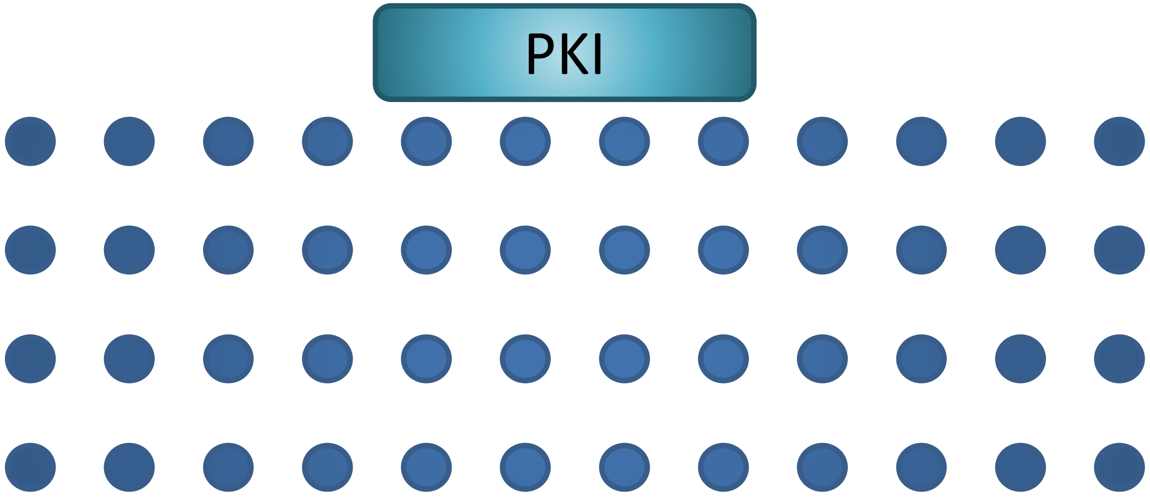

$\mathcal{P}_1 = \{P_1, \dots, P_{2m}\}$

Protocol Template

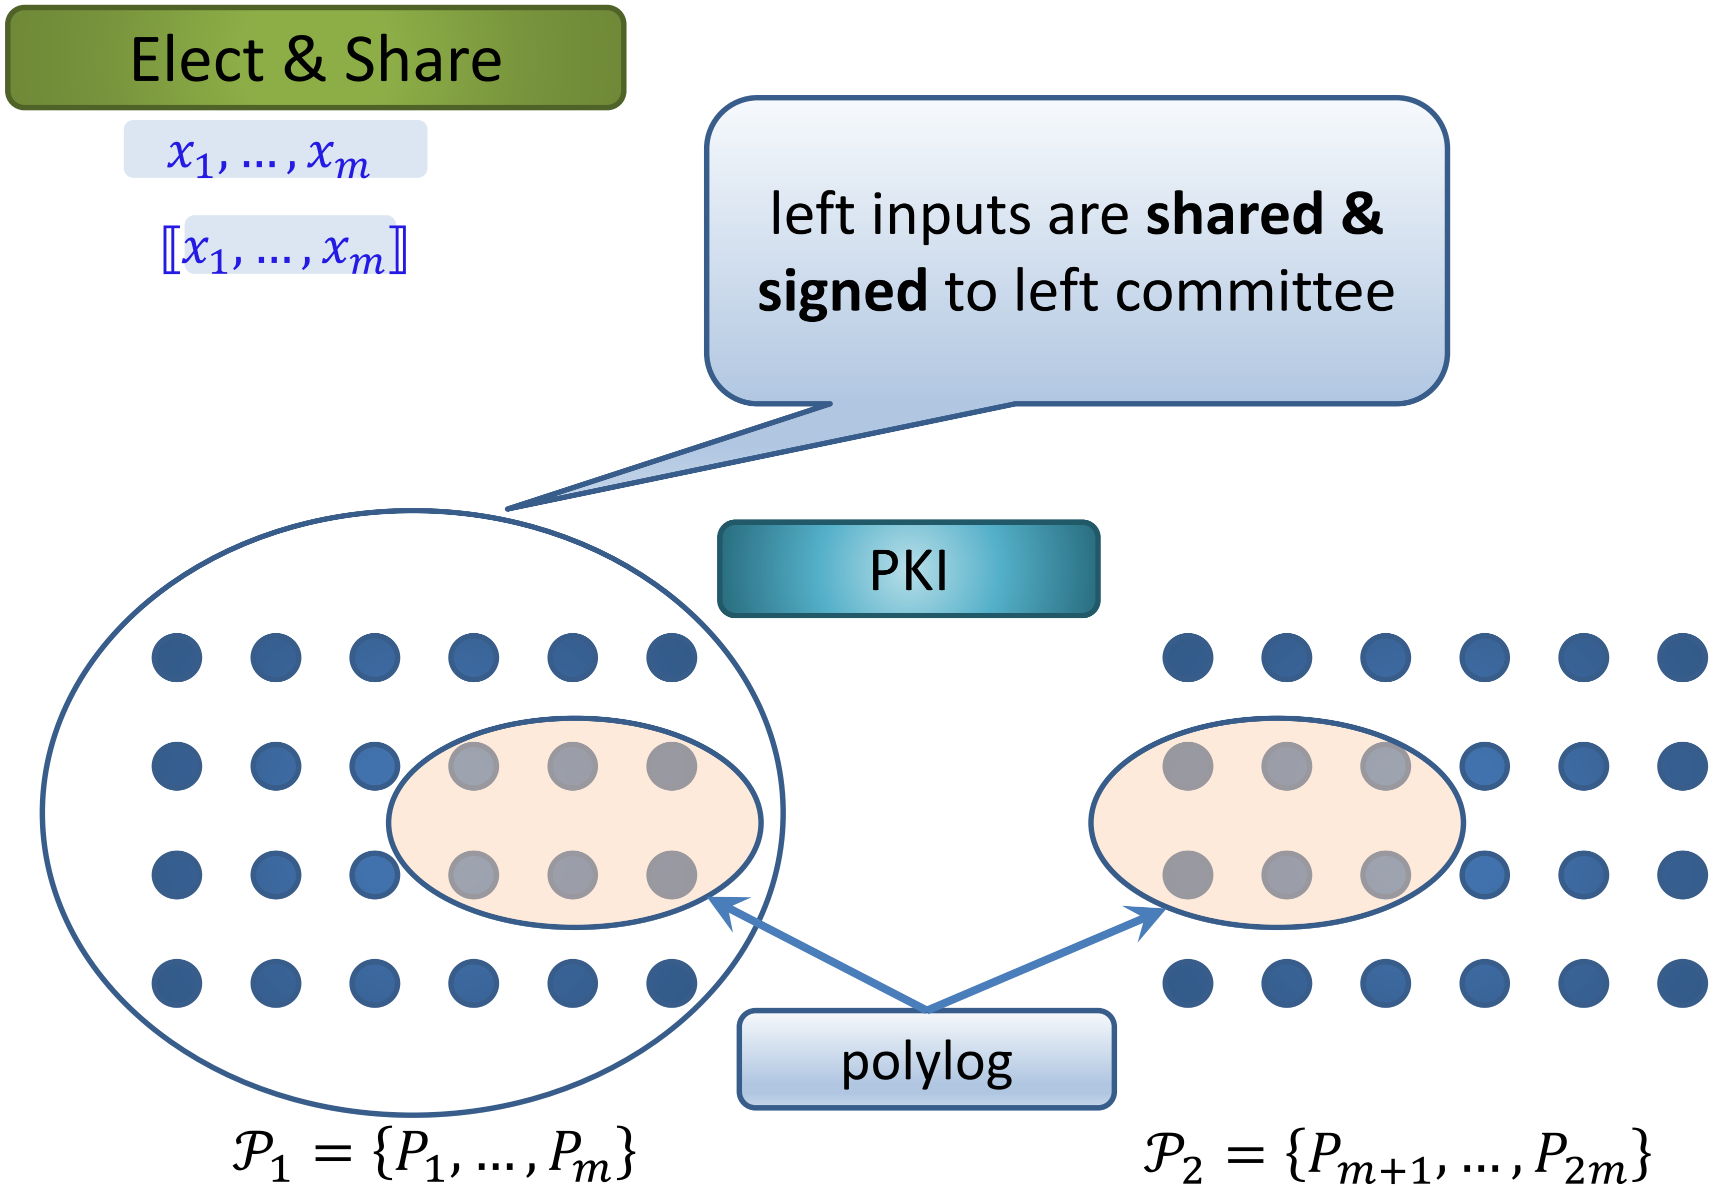

Protocol Template

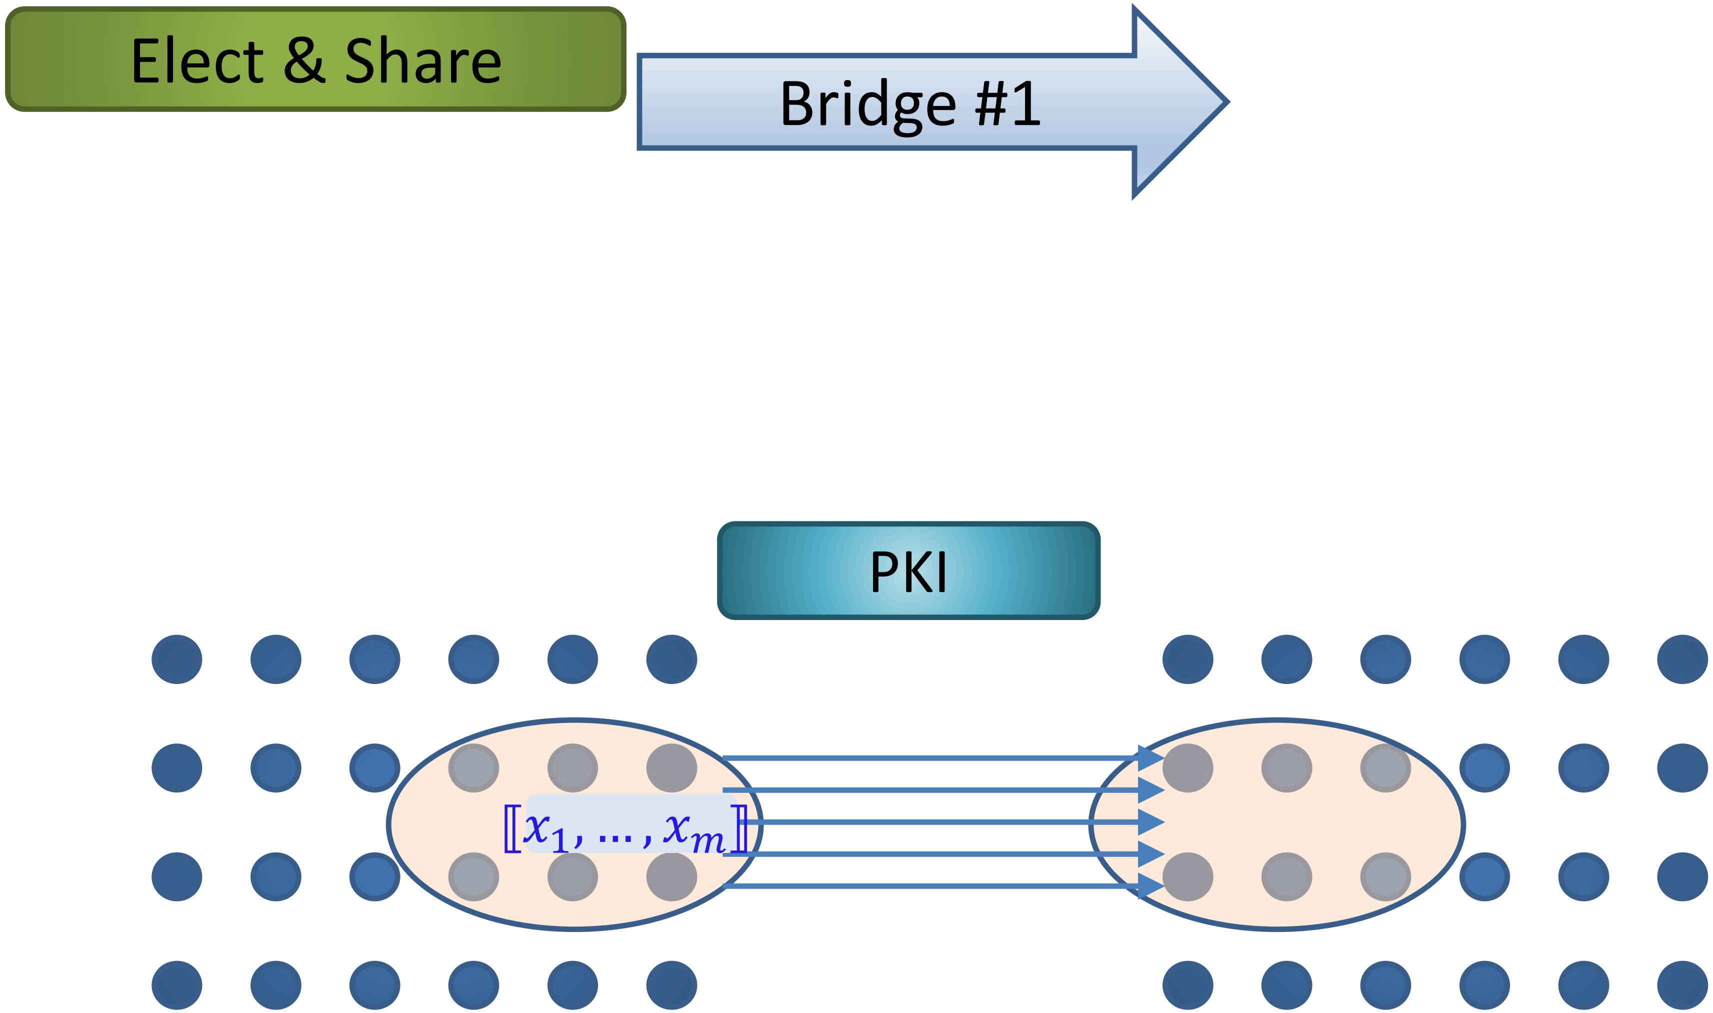

$\mathcal{P}_1 = \{P_1, \dots, P_m\}$ $\mathcal{P}_2 = \{P_{m+1}, \dots, P_{2m}\}$

Protocol Template

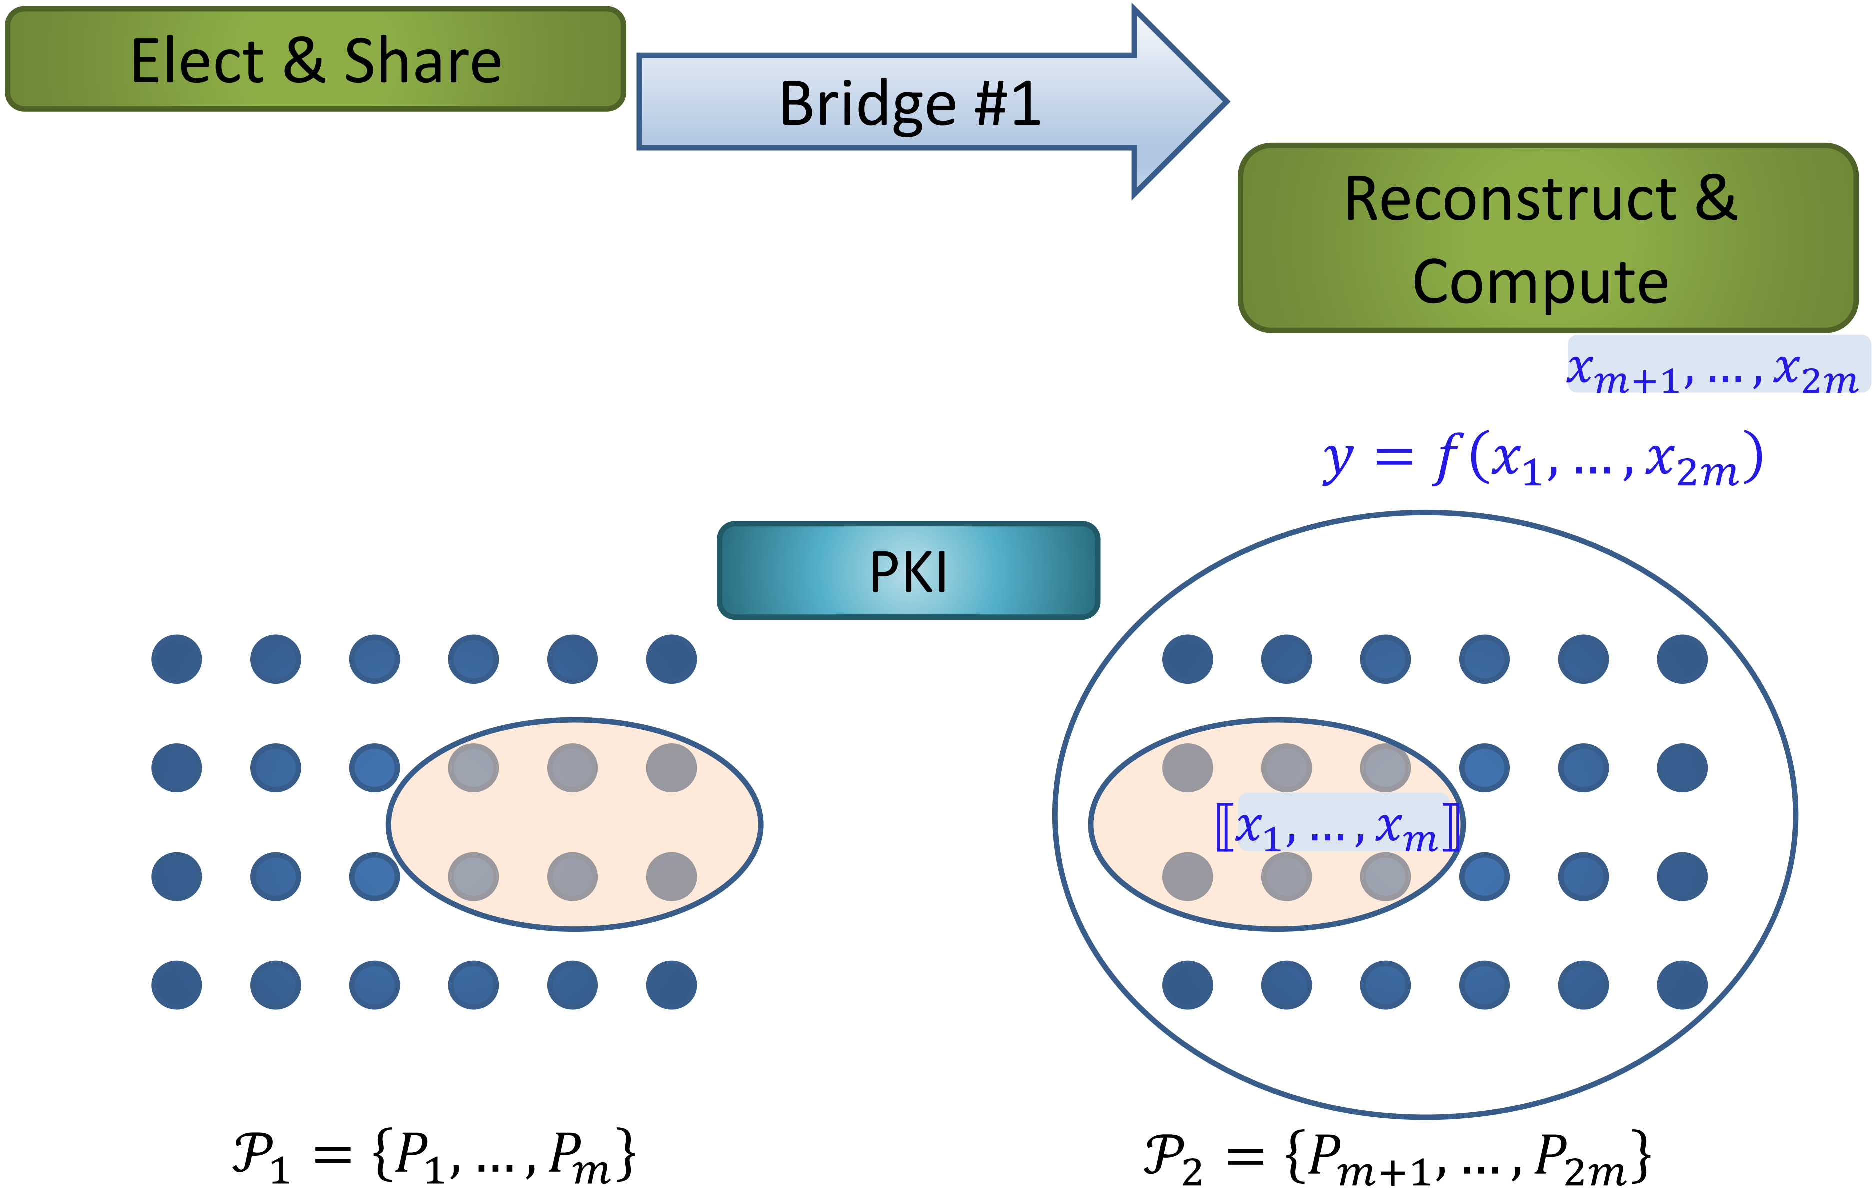

Protocol Template

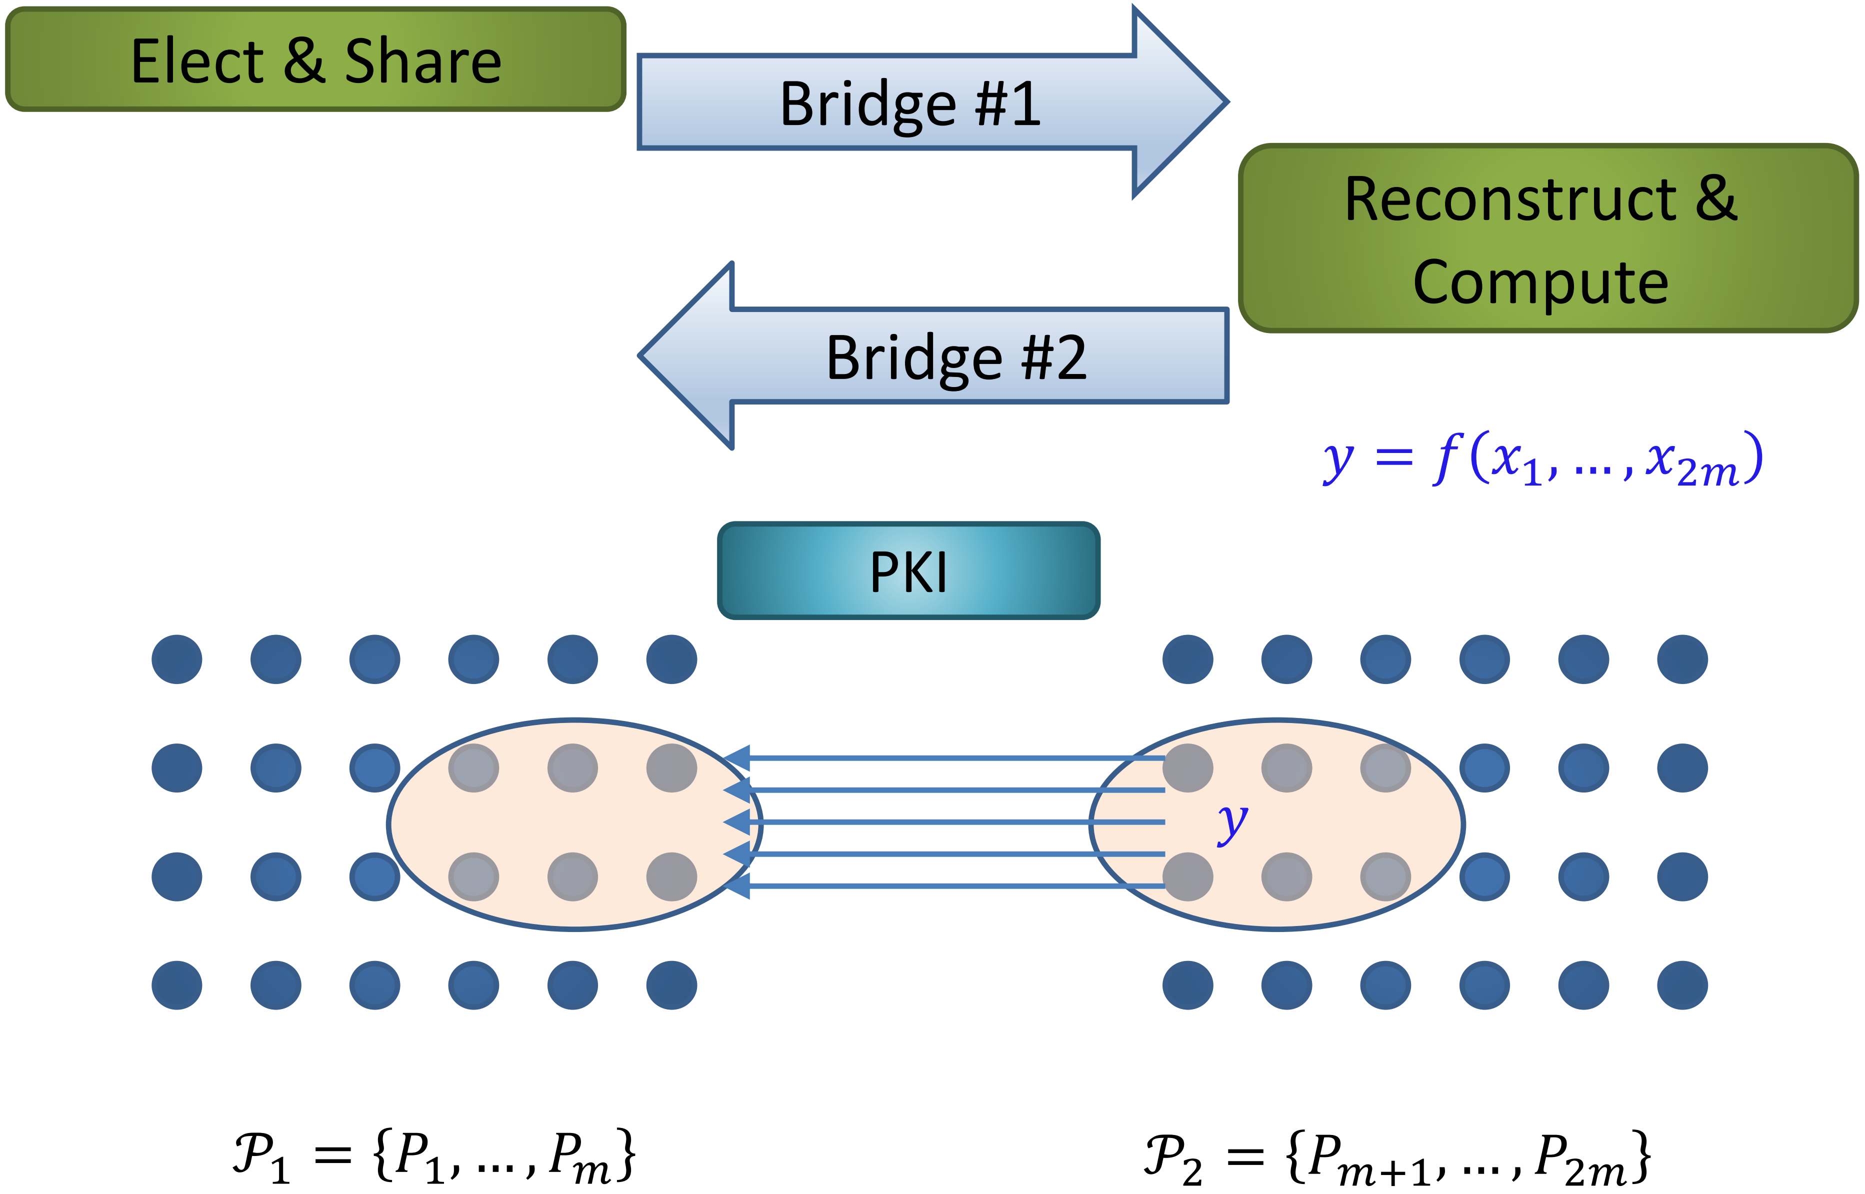

Protocol Template

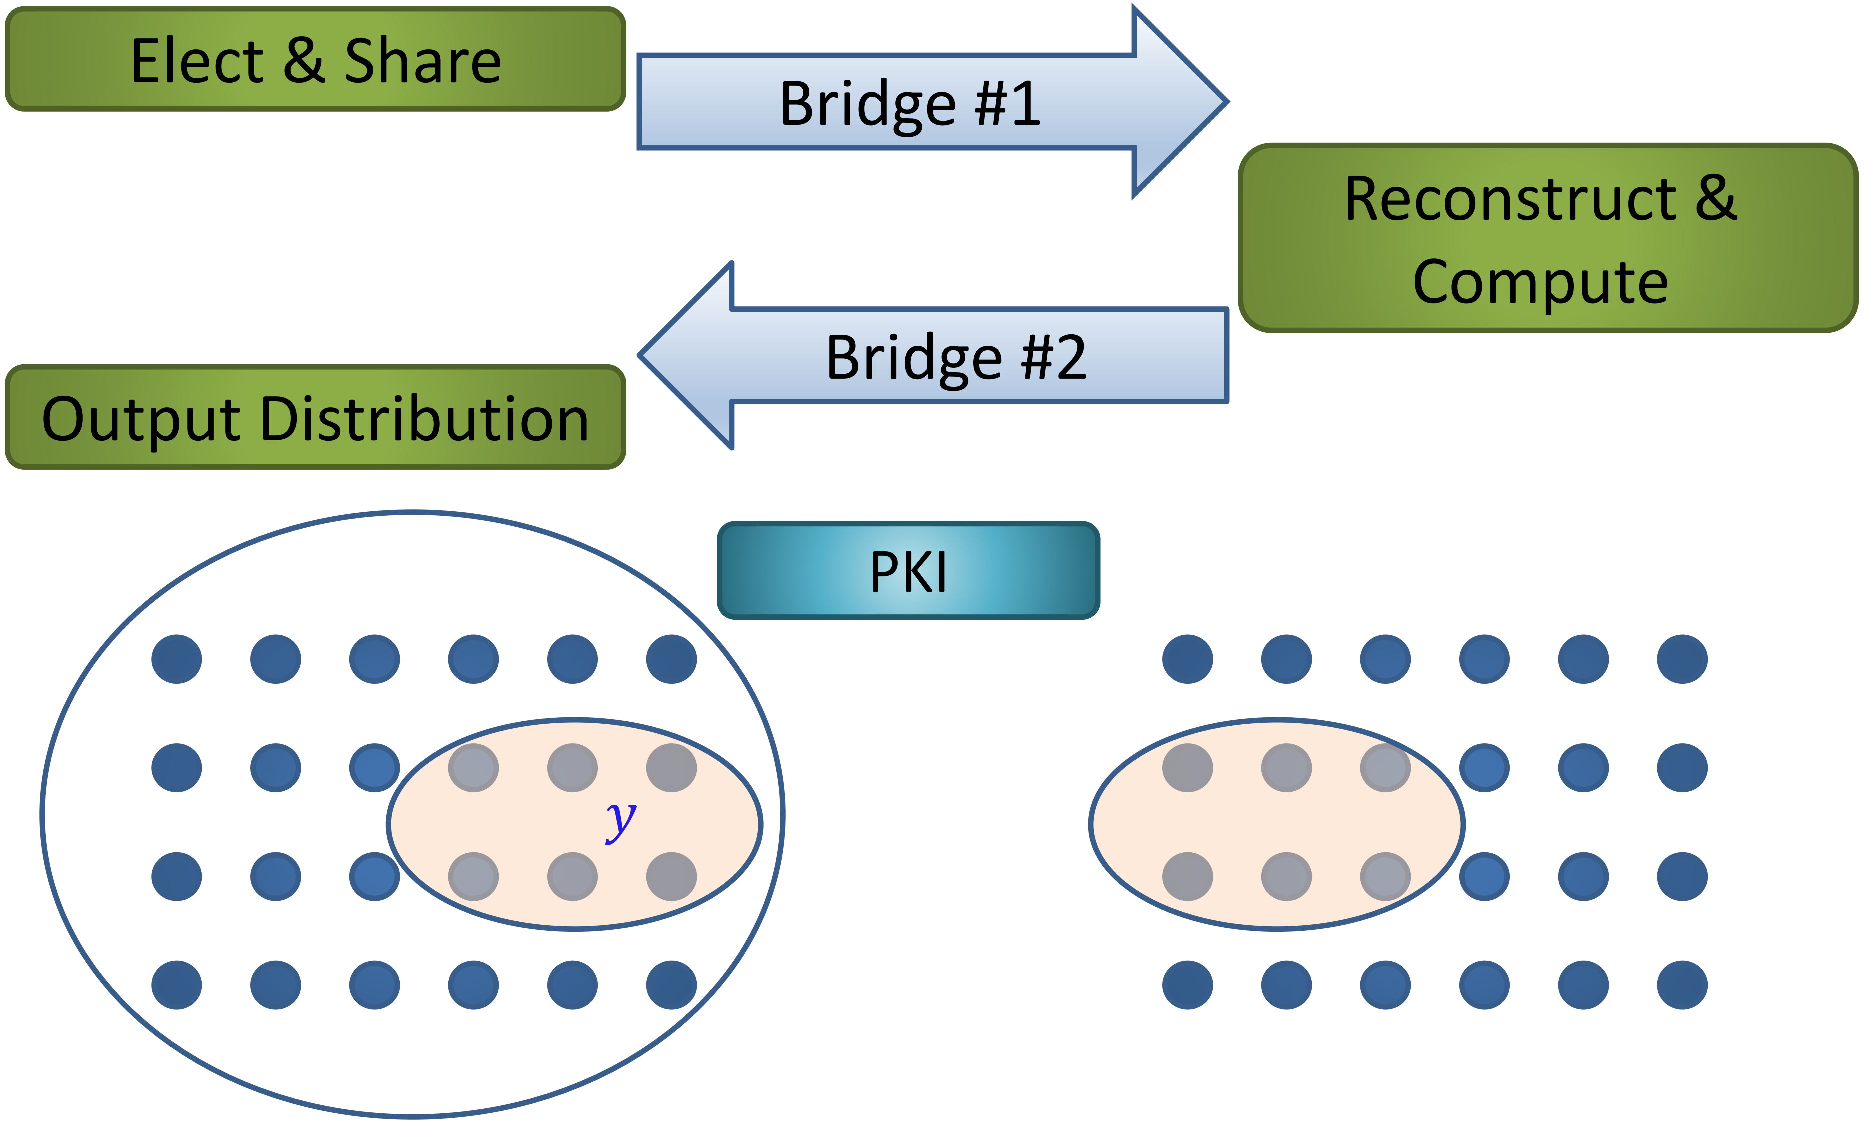

$\mathcal{P}_1 = \{P_1, \dots, P_m\}$ $\mathcal{P}_2 = \{P_{m+1}, \dots, P_{2m}\}$

Protocol Template

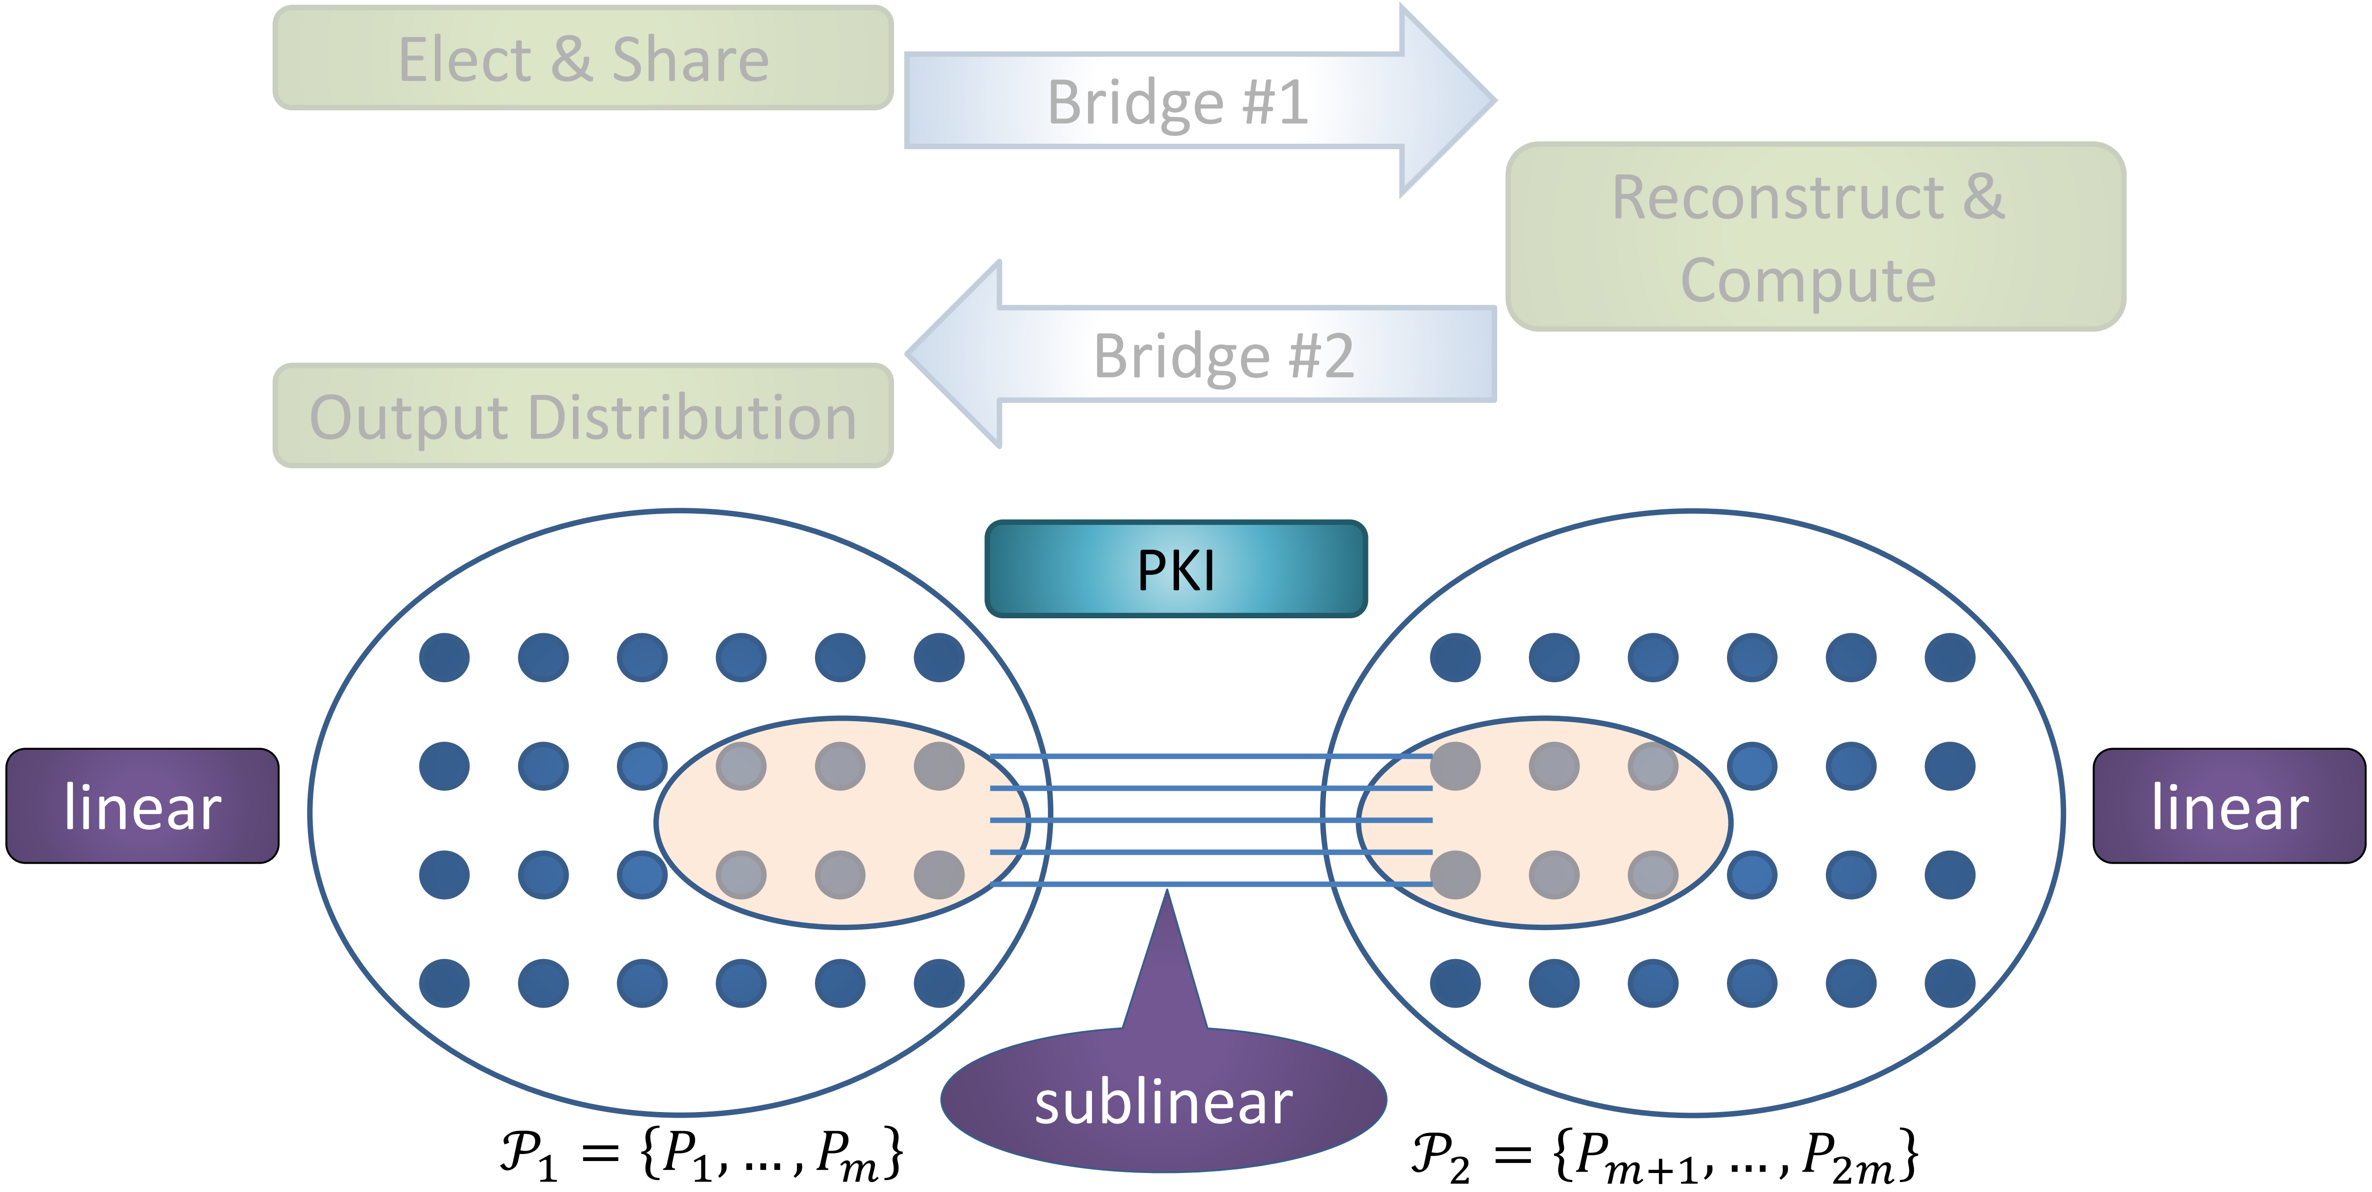

Corollaries (Static Corruptions)

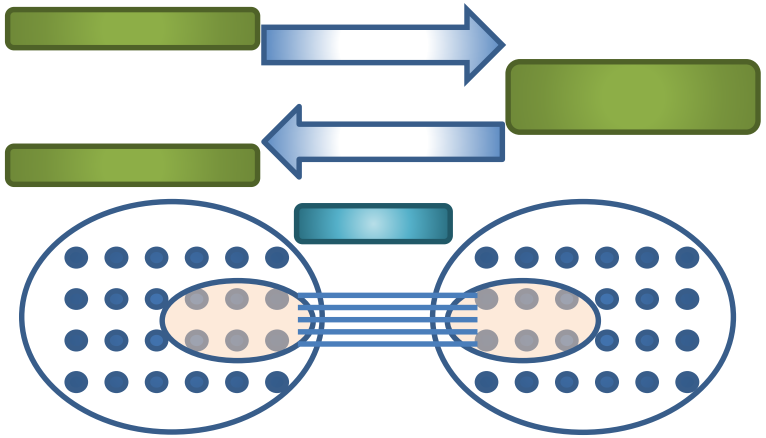

- Computational (PKI model)
  - $t = (1/4 - \epsilon)n$, assuming OWF using [Beaver, Micali, Rogaway'90]
  - $t = (1/6 - \epsilon)n$, with **polylog locality**, assuming OWF using [Boyle, Goldwasser, Tessaro'13]
  - $t = (1/4 - \epsilon)n$, with **polylog locality**, stronger assumptions using [Chandran, Chongchitmate, Garay, Goldwasser, Ostrovsky, Zikas'15]
- Information-theoretic (PKI for IT signatures)
  - $t = (1/4 - \epsilon)n$, using [Rabin, Ben-Or'89]
  - $t = (1/12 - \epsilon)n$, with **polylog locality**, **[This work]**

Adaptive Corruptions

Can the protocol template support *adaptive* corruptions?

- **Problem:** $\mathcal{A}$ sees messages between committees
- **Solution:** use **hidden channels** [CCGGOZ'15]
  $\mathcal{A}$ is unaware of messages between honest parties

Inherent for **this template** and **low-locality** protocols

- **Problem:** committees are known - can be fully corrupted
- **Solution:** hide the committees
  Every member only learns **one** corresponding partner

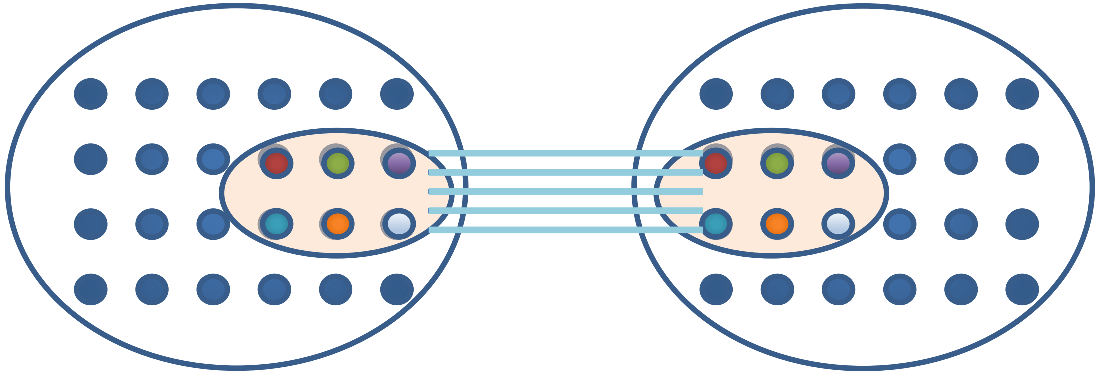

Corollaries (Adaptive Corruptions)

- Computational (PKI model)
  - $t = (1/8 - \epsilon)n$, assuming OWF, using [Damgard, Ishai'05]
  - $t = (1/8 - \epsilon)n$, with **polylog locality**, stronger assumptions, using [Chandran, Chongchitmate, Garay, Goldwasser, Ostrovsky, Zikas'15]
- Information-theoretic (using IT signatures)
  - $t = (1/8 - \epsilon)n$, using [Cramer, Damgard, Dziembowski, Hirt, Rabin'99]

Lower Bound:
Protocols that must be Expanders

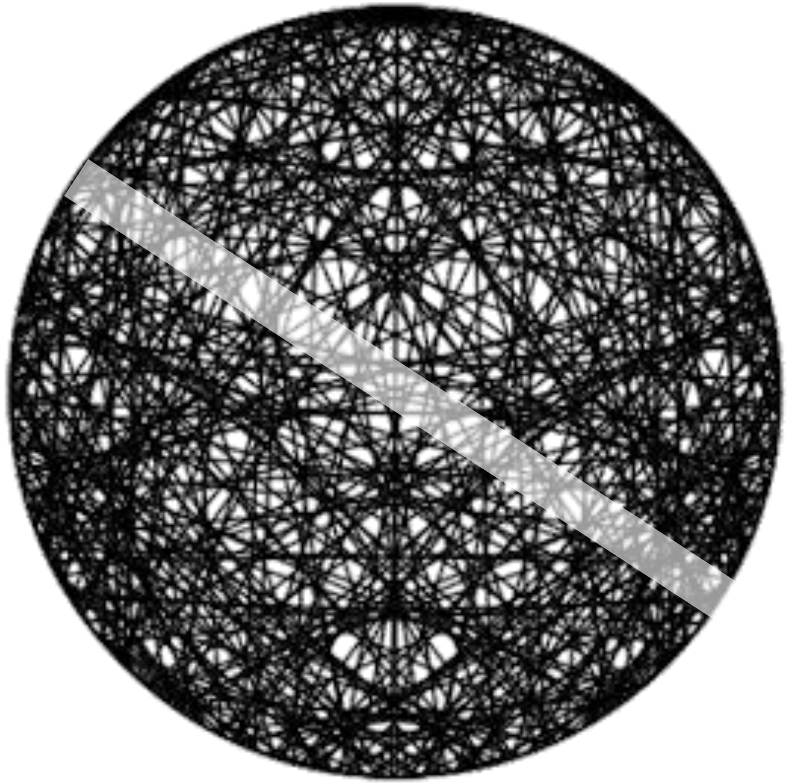

Lower Bound

The setting:

- Adaptive adversary
- Common Reference String (CRS)
- Private (visible) channels

Parallel broadcast (aka interactive consistency [PSL'80]):

- Every party broadcasts $x_i \in \{0,1\}^n$
- Common output is $(y_1, \dots, y_n)$, if P_i is honest $y_i = x_i$

Theorem (Lower Bound)

Let π be **parallel broadcast** protocol tolerating PPT adversary adaptively corrupting $\beta \cdot n$ parties (for any constant $\beta > 0$)

Then, there are **no sublinear cuts** in the communication graph of π

In particular, π is an **expander**

Lower Bound – isn't it trivial?

- **Idea**: linear corruptions, sublinear cut – corrupt the "bridge"
- **Problem 1:** the location is unknown ahead of time
- **Problem 2:** maybe one side is fully corrupt
  Need to separate two honest parties

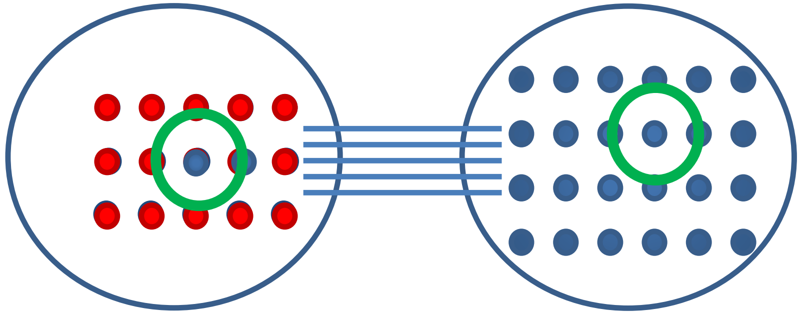

- **Idea:** wait until the location of the cut is known
- **Problem:** this is too late – information already crossed over

Our approach:
Gradually learn the location of cut while blocking information flow

Proof Idea (Very High Level)

- Can focus on $\beta < 1/3$ [PSL'80]
- Execute π over **random** inputs
- Assume there exists $\alpha(n)$-cut (sublinear)

Phase 1: Isolate a random party until its degree is n/c (c is const depends on β)

Phase 2: Block all messages between every U_i and U_j

After Phase 1:

- With noticeable probability all nodes have degree n/c
- Can efficiently find $(\alpha(n), n/c)$-partition of the graph

Partition $\{U_1, \dots, U_c\}$ of nodes

- $|U_i| \geq n/c$
- $|\text{edges}(U_i, U_j)| \leq \alpha(n)$
- "basis" for $\alpha(n)$-cuts

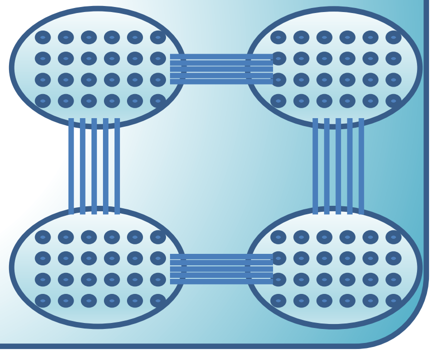

Phase 1

- Choose a random party P_{i^*}
- Block all outgoing messages
- **Important**: all parties must be unaware of the attack
  - Simulate P_{i^*} on random input to all other (**red execution**)
  - Simulate honest execution towards P_{i^*} (**blue execution**)

parties might change behavior
start talking faster to/from P_{i^*}

cannot work
with PKI

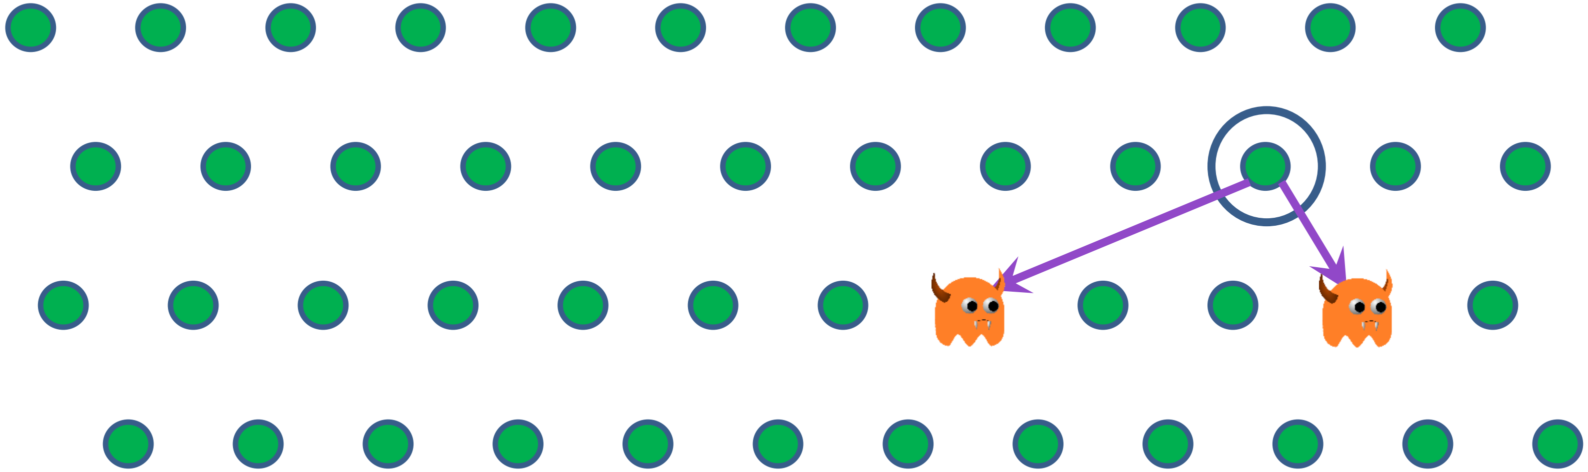

Phase 1

- Choose a random party P_{i^*}
- Block all outgoing messages
- **Important**: all parties must be unaware of the attack
  - Simulate P_{i^*} on random input to all other (**red execution**)
  - Simulate honest execution towards P_{i^*} (**blue execution**)

parties might change behavior start talking faster to/from P_{i^*}

cannot work with PKI

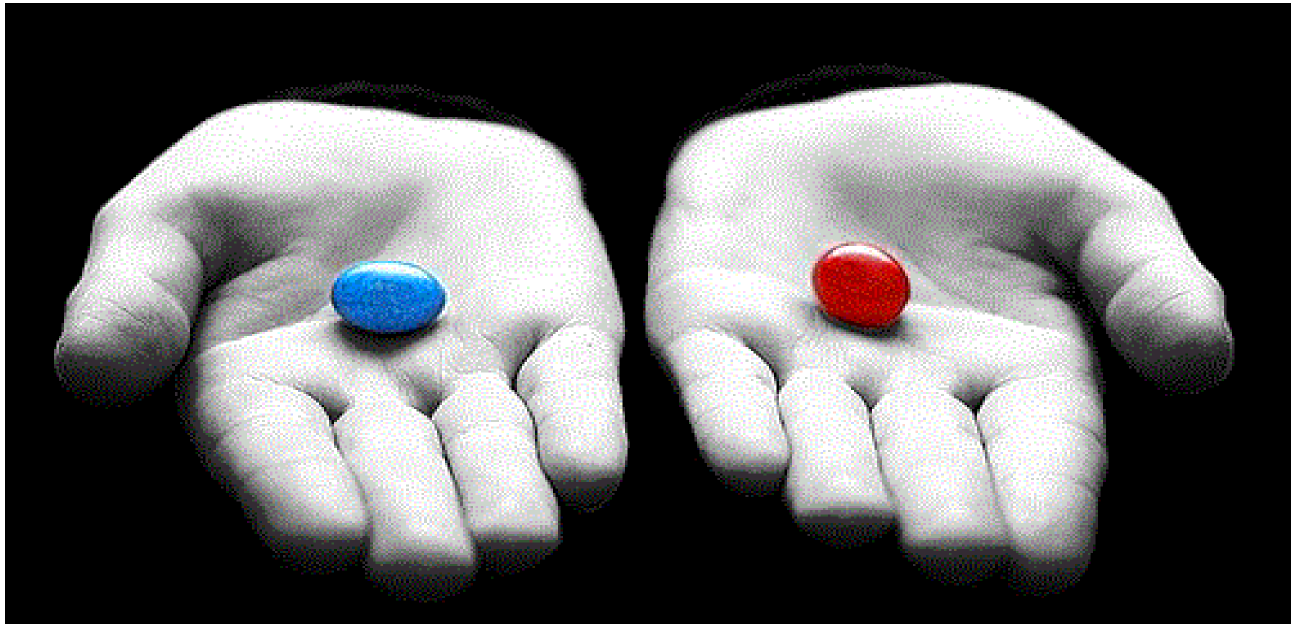

Blue Execution

Goal: make P_{i^*} think he runs in an honest (virtual) execution

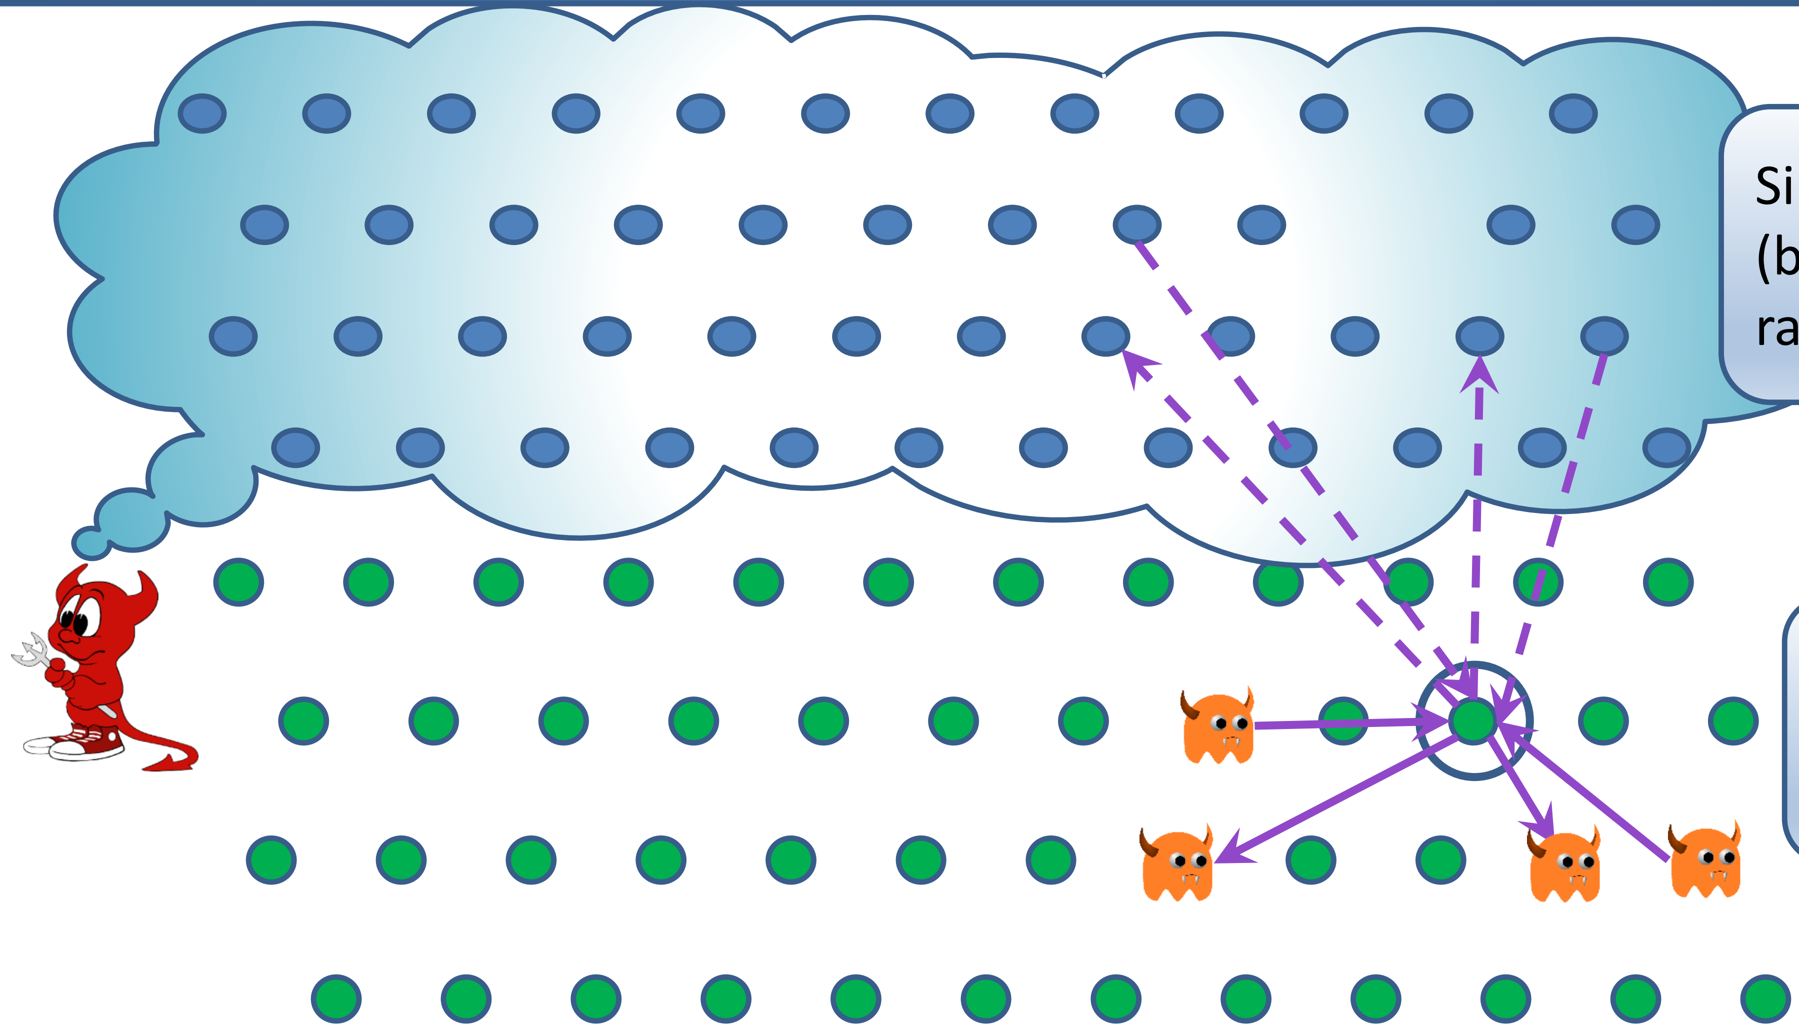

Simulate all parties (but P_{i^*}) on random $\tilde{x}_i$ and $\tilde{r}_i$

P_{i^*} runs on (real) input x_{i^*} and r_{i^*}

Red Execution

Goal: trick other honest parties to think there is no attack

All $\{P_i\}_{i \neq i^*}$ run on (real) x_i and r_i

Simulate $\tilde{P}_{i^*}$ on random $\tilde{x}_{i^*}$ and $\tilde{r}_{i^*}$

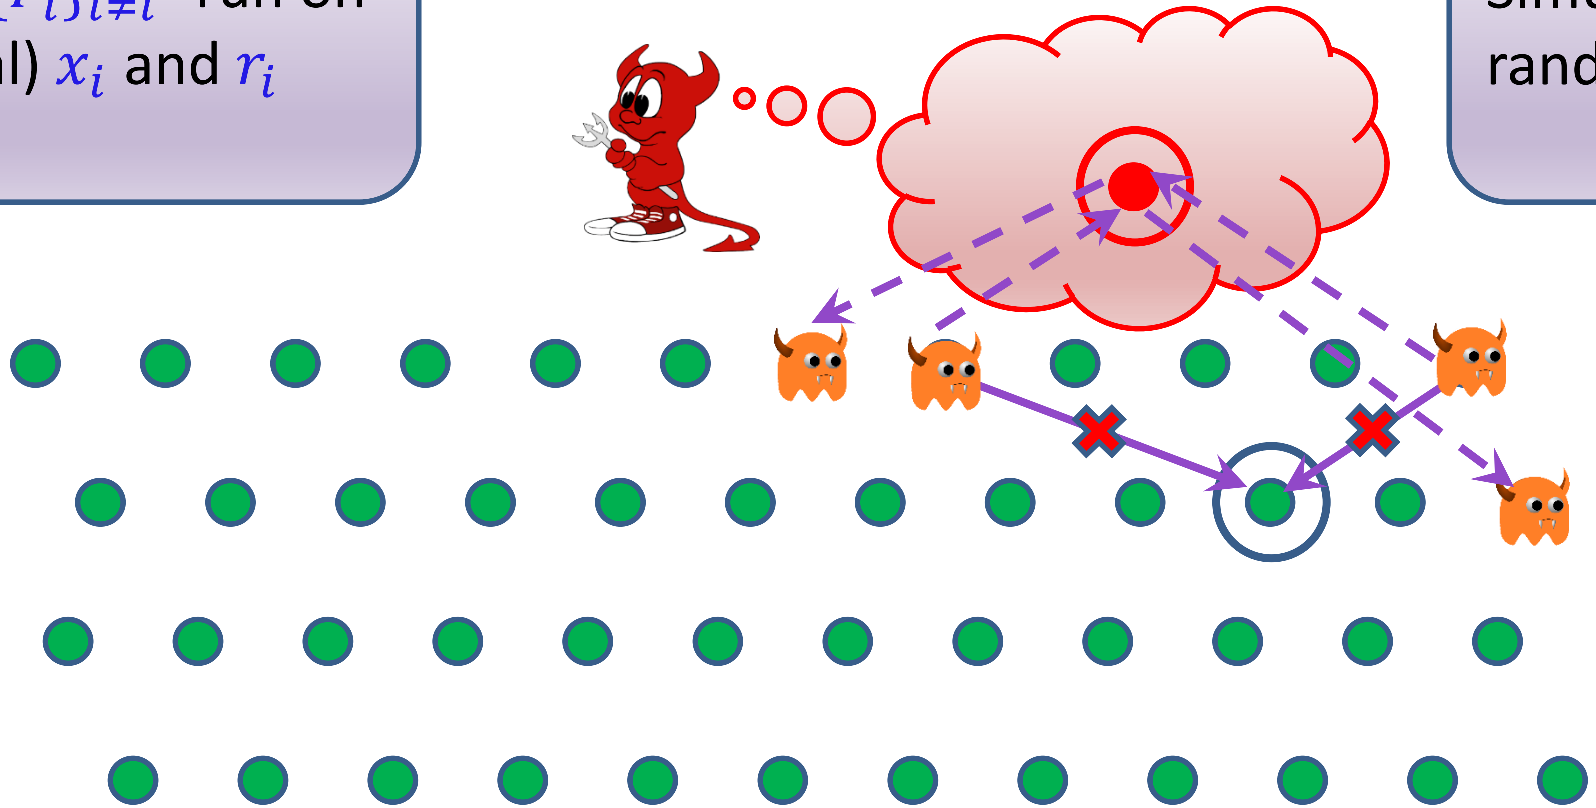

Phase 1 - Summary

- Both **red** and **blue** executions are distributed as **independent honest executions over random inputs**
- Continue until P_{i^*} has $\beta n/4$ neighbors in **both** executions
  - wp $1/n^2$ party P_{i^*} is **last** to have degree $\beta n/4$ in both
  - $\Rightarrow$ All parties have degree $\geq n/c$ where c depends on β

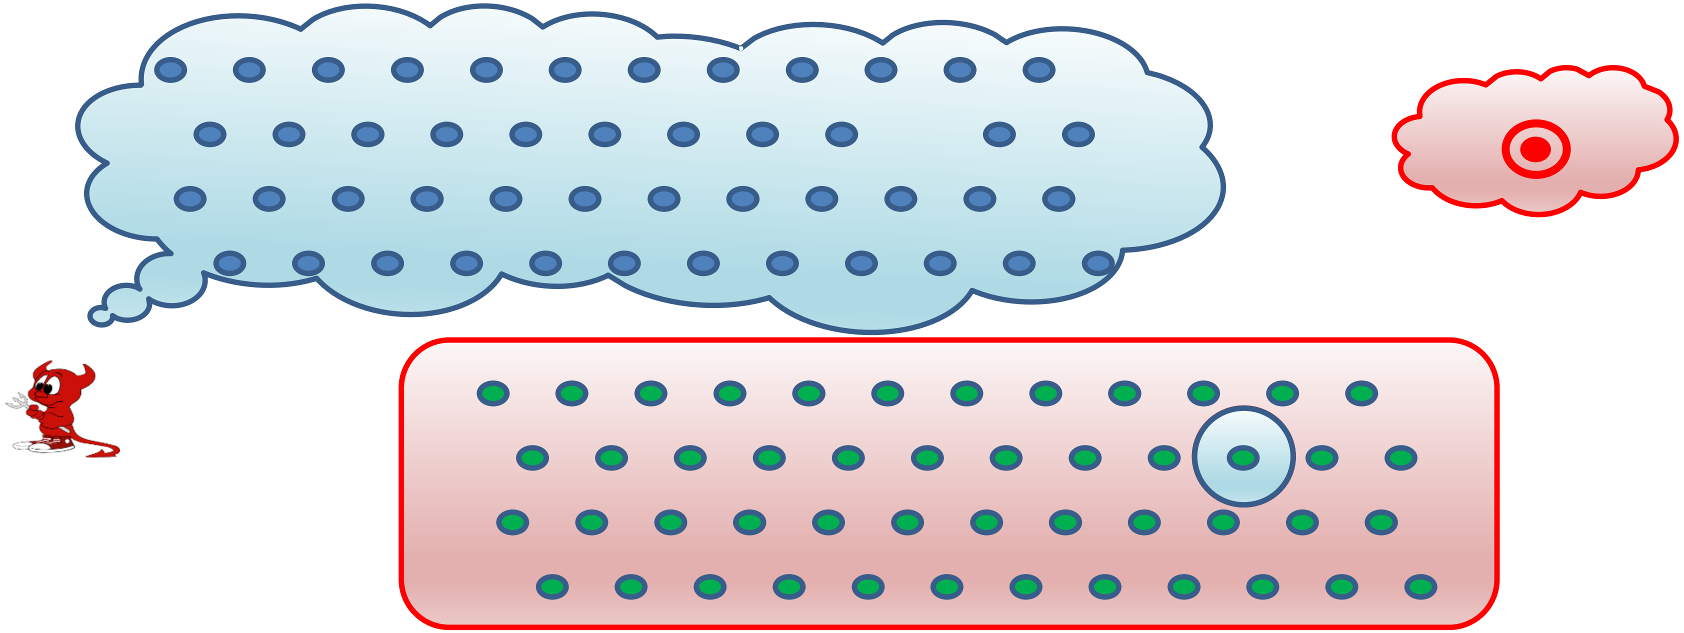

Graph-Theoretic Pause

Theorem (Linear degree ⇒ constant number of sublinear cuts):

Let $G = ([n], E)$ with linear degree n/c and let $\alpha(n) \in o(n)$

1. There is an $(\alpha(n), n/c)$-partition $\Gamma = \{U_1, \dots, U_m\}$ of the nodes s.t.
   - $m \leq c$
   - $|U_i| \geq n/c$
   - $\left|\mathrm{edges}(U_i, U_j)\right| \leq \alpha(n)$
   - Γ is a "basis" for $\alpha(n)$-cuts
2. The number of $\alpha(n)$-cuts is constant ($\leq 2^{c-1}$)
3. Γ can be found in polynomial time

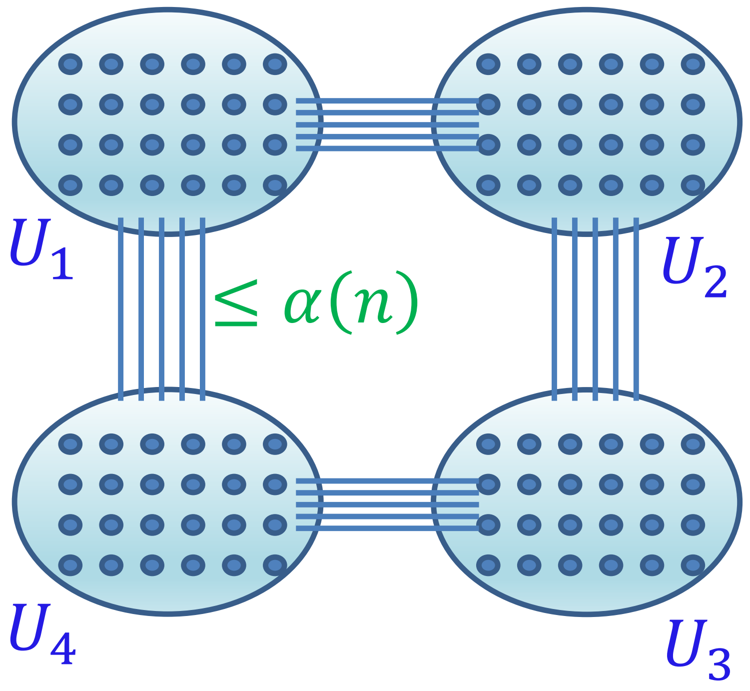

Back to the Attack - Phase 2

- With prob. $1/n^2$ every party has linear degree n/c
- Find $(\alpha(n), n/c)$-partition $\Gamma = \{U_1, \dots, U_m\}$
- Block messages between every U_i and U_j
  - Stop blocking if $|\mathrm{edges}(U_i, U_j)| \geq \alpha(n)$
  - Never corrupt P_{i^*}

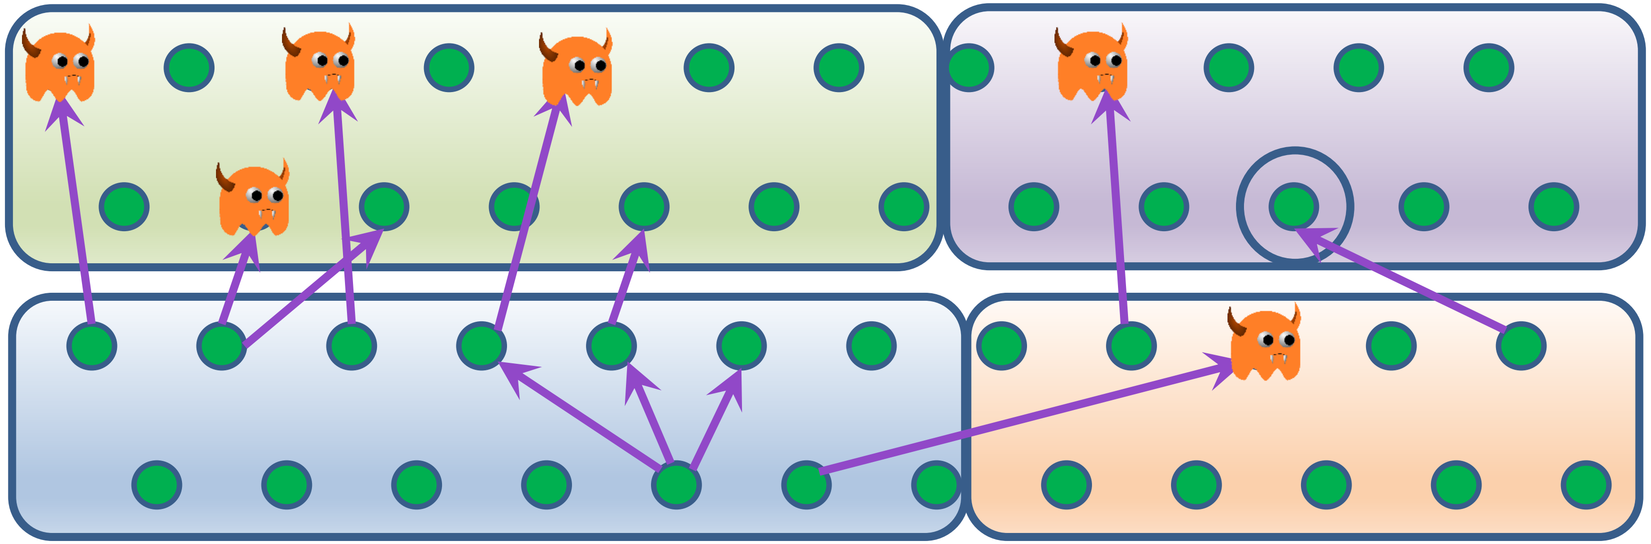

Where do we stand

- P_{i^*} is honest $\Rightarrow$ by **correctness** all honest parties output $y_{i^*} = x_{i^*}$
- By assumption $\exists$ an $\alpha(n)$-cut at the end
- Phase 1: messages across the cut **independent** of x_{i^*}
- Phase 2: **no messages** across the cut

Phase 2: $o(n)$ corruptions
Phase 1: $o(n)$ **blue** corruptions in $\bar{S}$
Phase 1: **linear red** corruptions in $\bar{S}$

Does this imply that some honest parties output $\boldsymbol{y_{i^*} \neq x_{i^*}}$?

Problem 1: maybe the **entire side** of the cut is corrupt?

Problem 2: maybe **information** is flowing by other means?

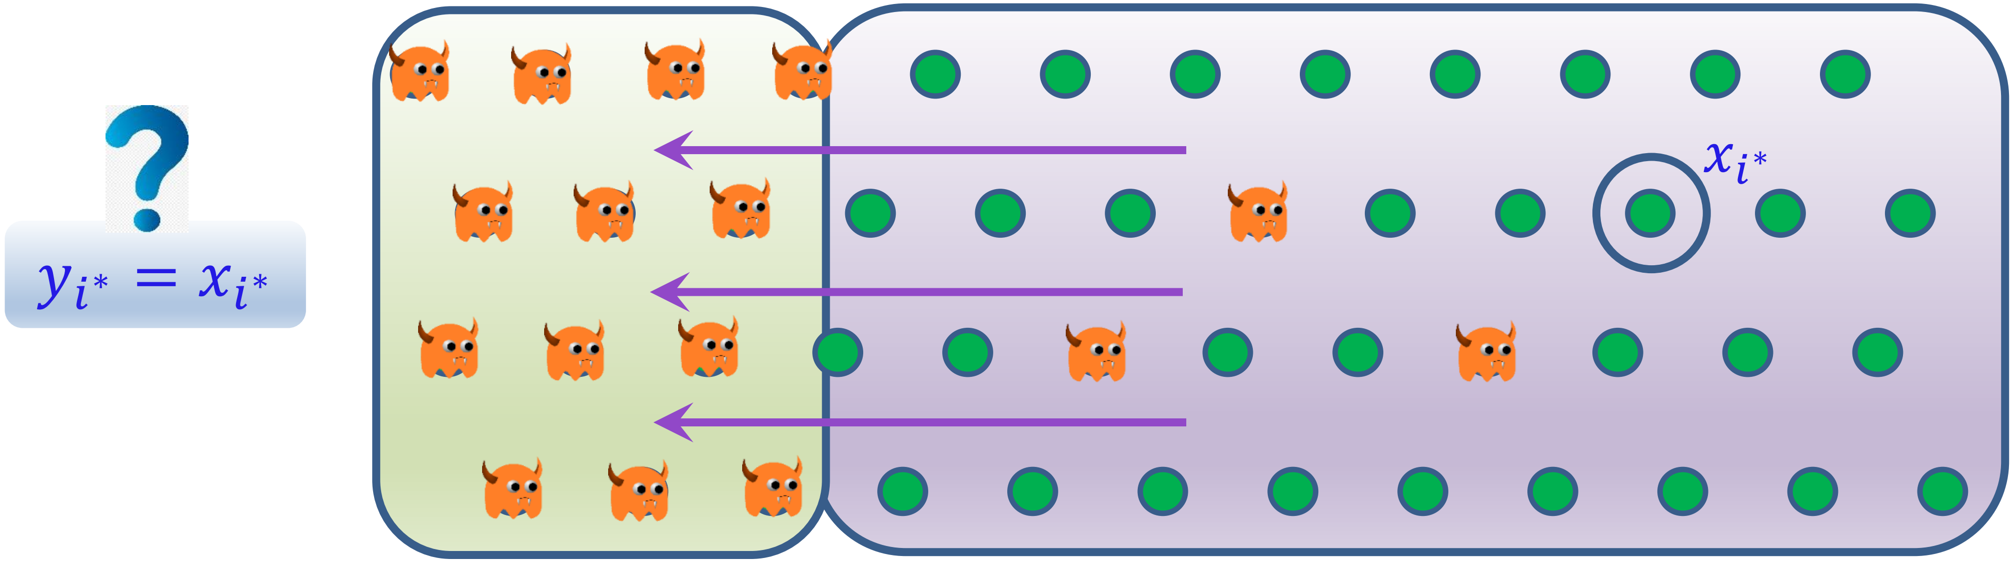

Problem 1: Guaranteeing Honest Party Across the Cut

We **DO NOT** guarantee honest party across the cut

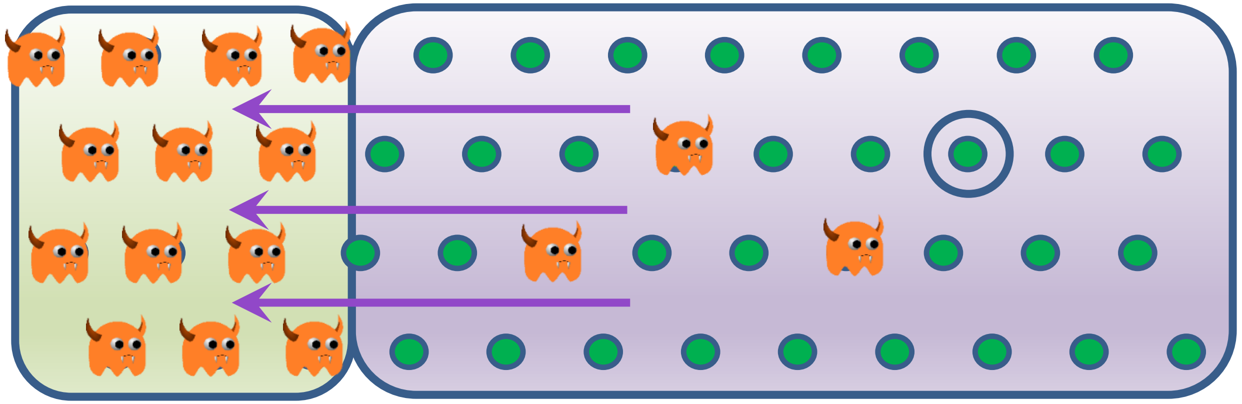

Instead, define dual adversary $\mathcal{A}_{\mathrm{dual}}$

- Only P_{i^*} is corrupt in Phase 1
- Emulate its behavior as if being attacked

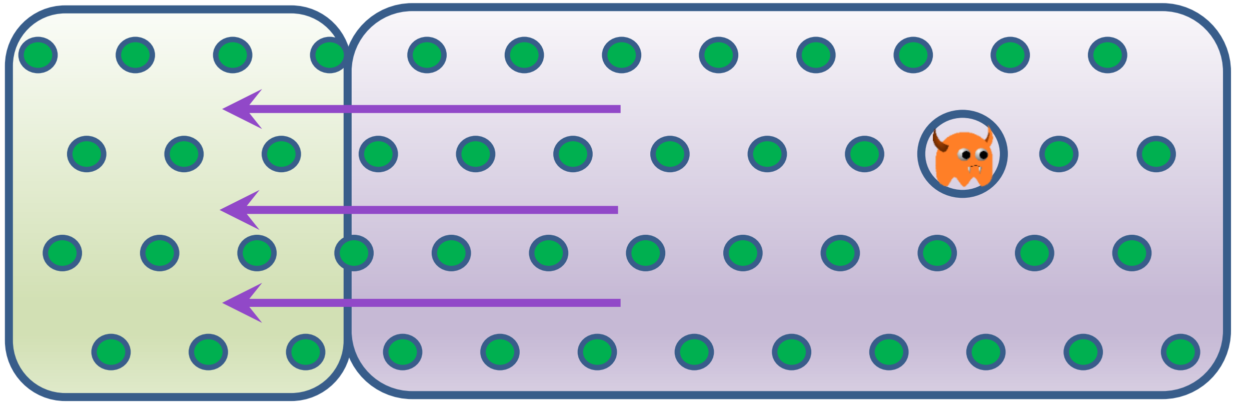

Dual Adversary

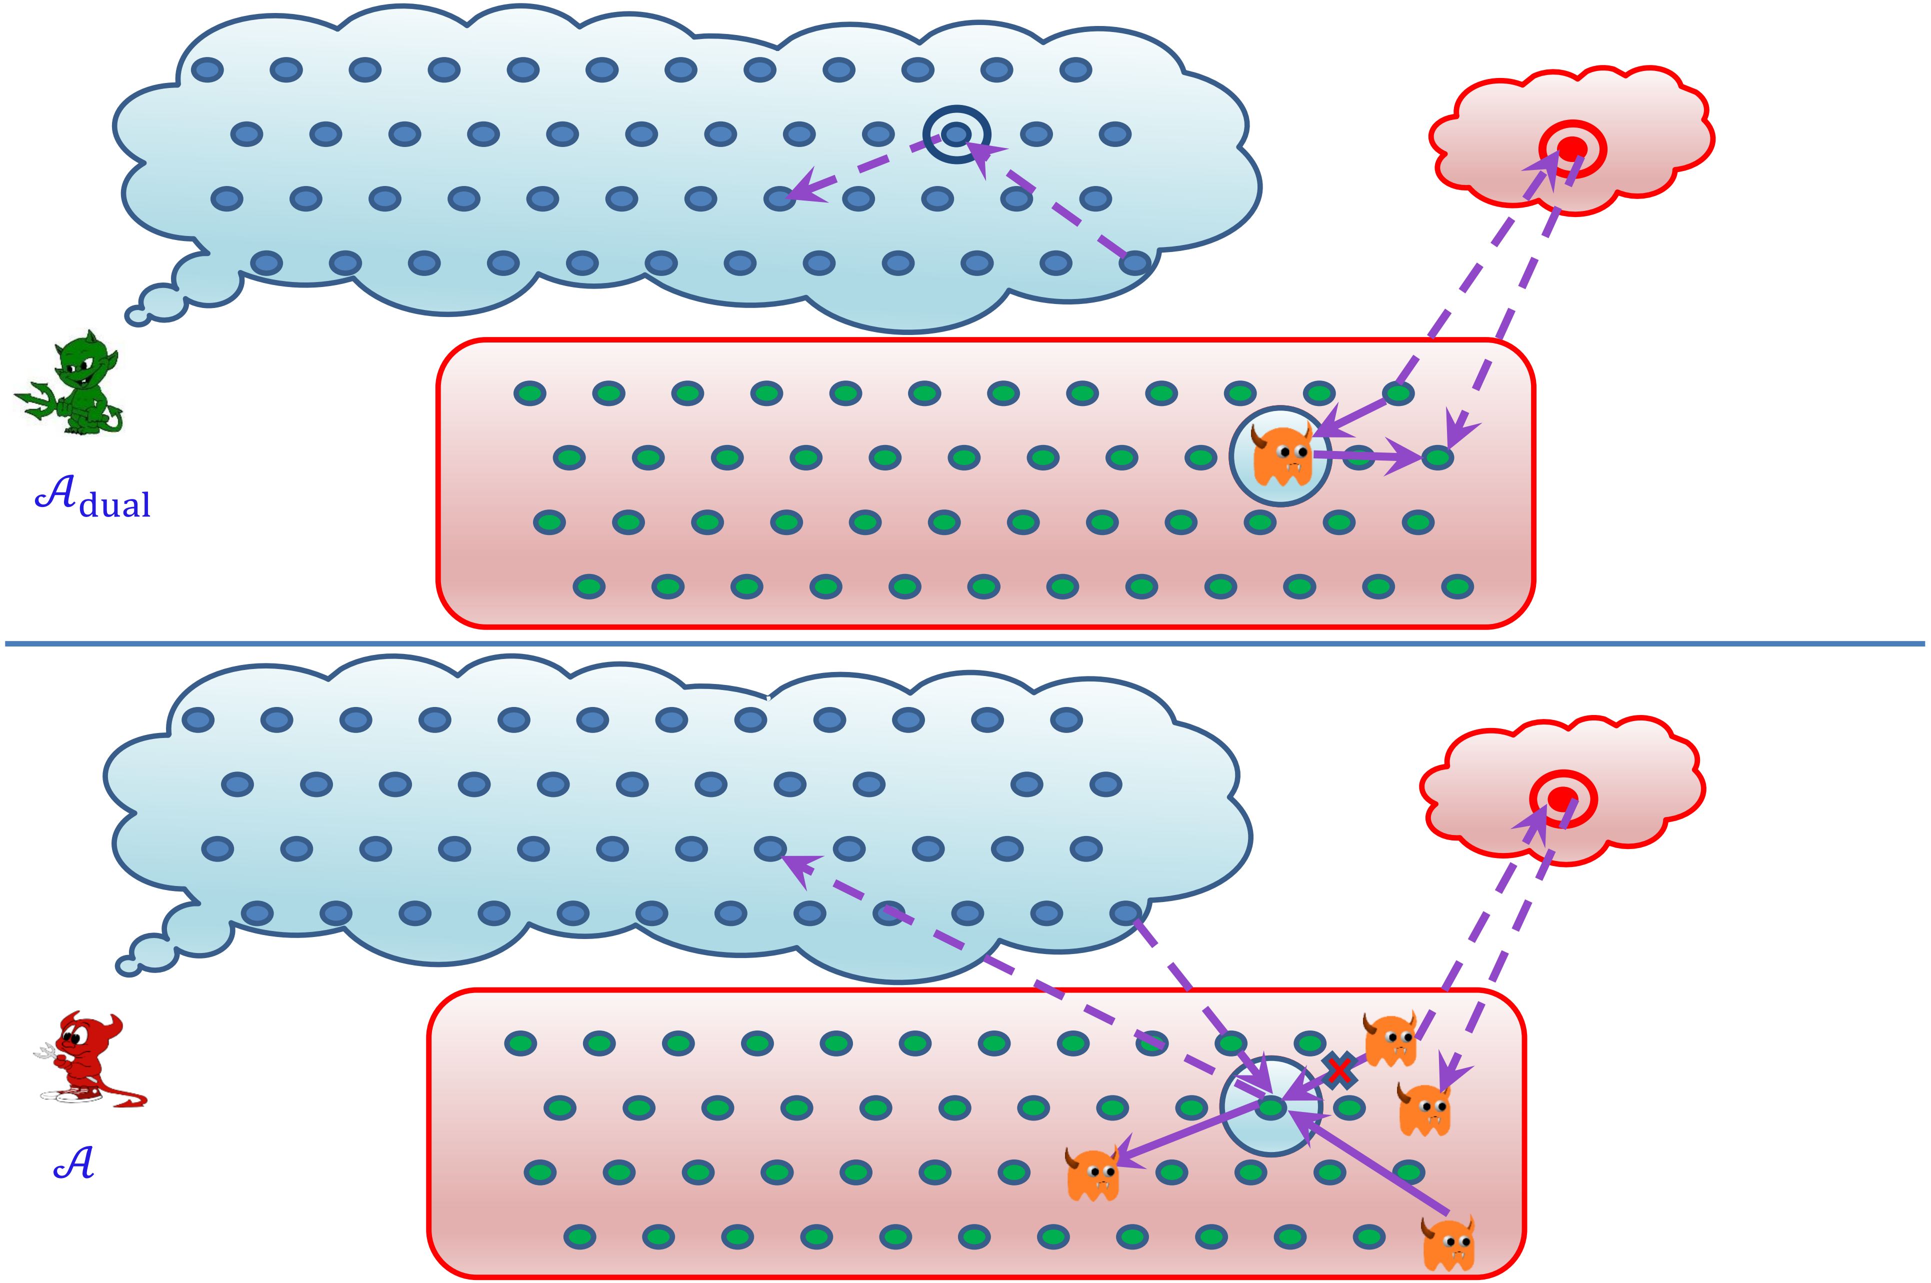

Guaranteeing Honest Party Across the Cut

1) With $\mathcal{A}$, party P_{i^*} is honest $\Rightarrow$ the common output is $y_{i^*} = x_{i^*}$
2) Some honest parties have same view under attacks of $\mathcal{A}$ and $\mathcal{A}_{\text{dual}}$ $\Rightarrow$ such parties output $y_{i^*} = x_{i^*}$ also with $\mathcal{A}_{\text{dual}}$
3) By correctness **all** honest parties output the same y_{i^*} with $\mathcal{A}_{\text{dual}}$
4) With $\mathcal{A}_{\text{dual}}$ there exists honest party across the cut

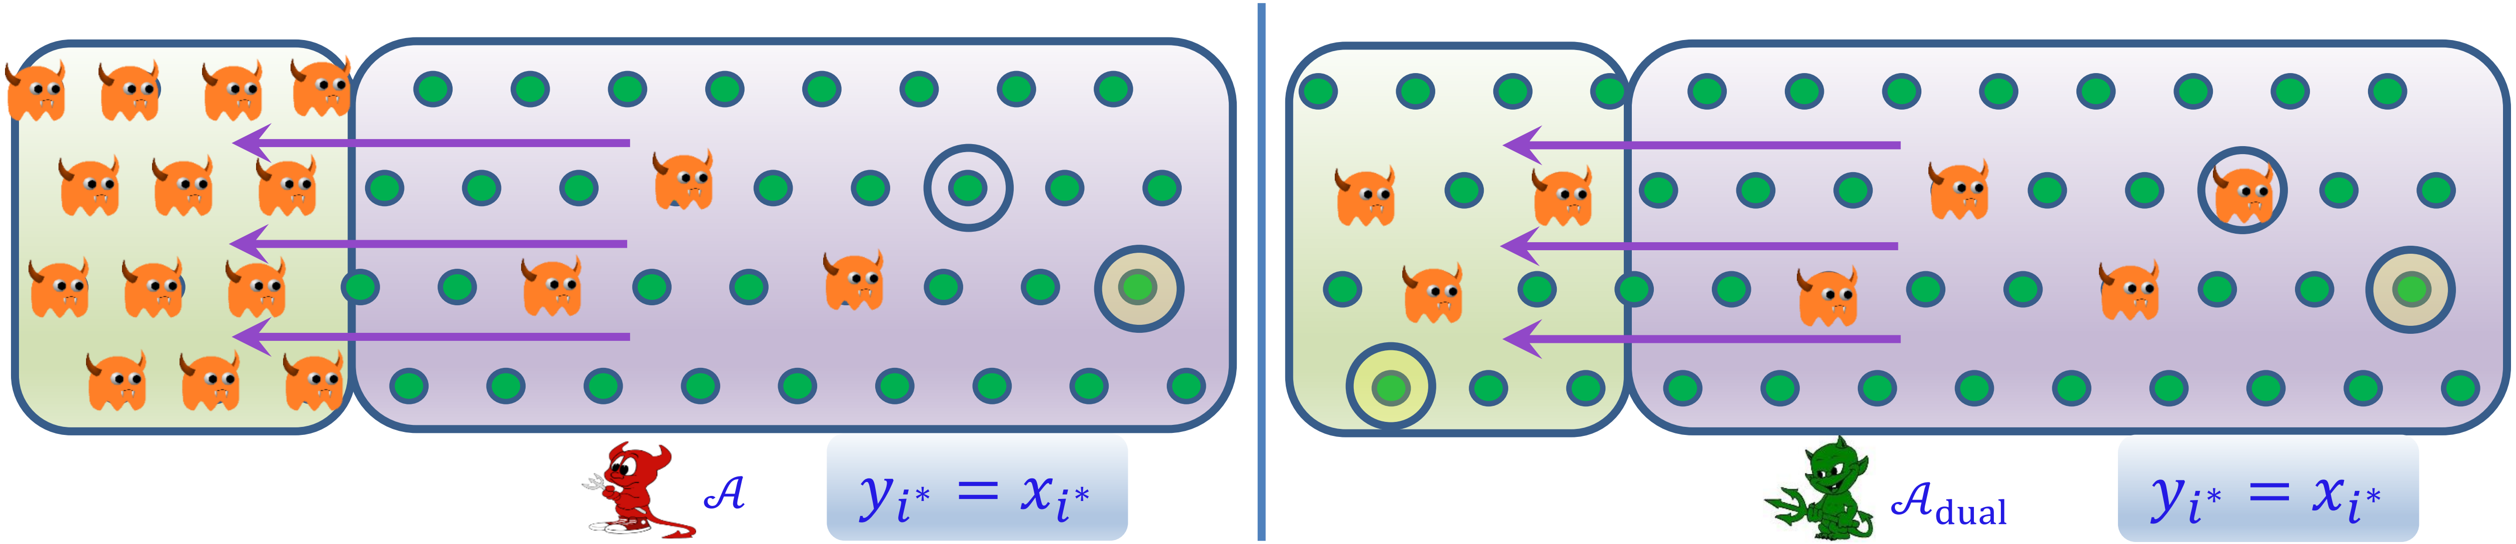

Guaranteeing Honest Party Across the Cut

1) With $\mathcal{A}$, party P_{i^*} is honest $\Rightarrow$ the common output is $y_{i^*} = x_{i^*}$
2) Some honest parties have same view under attacks of $\mathcal{A}$ and $\mathcal{A}_{\text{dual}}$ $\Rightarrow$ such parties output $y_{i^*} = x_{i^*}$ also with $\mathcal{A}_{\text{dual}}$
3) By correctness **all** honest parties output the same y_{i^*} with $\mathcal{A}_{\text{dual}}$
4) With $\mathcal{A}_{\text{dual}}$ there exists honest party across the cut

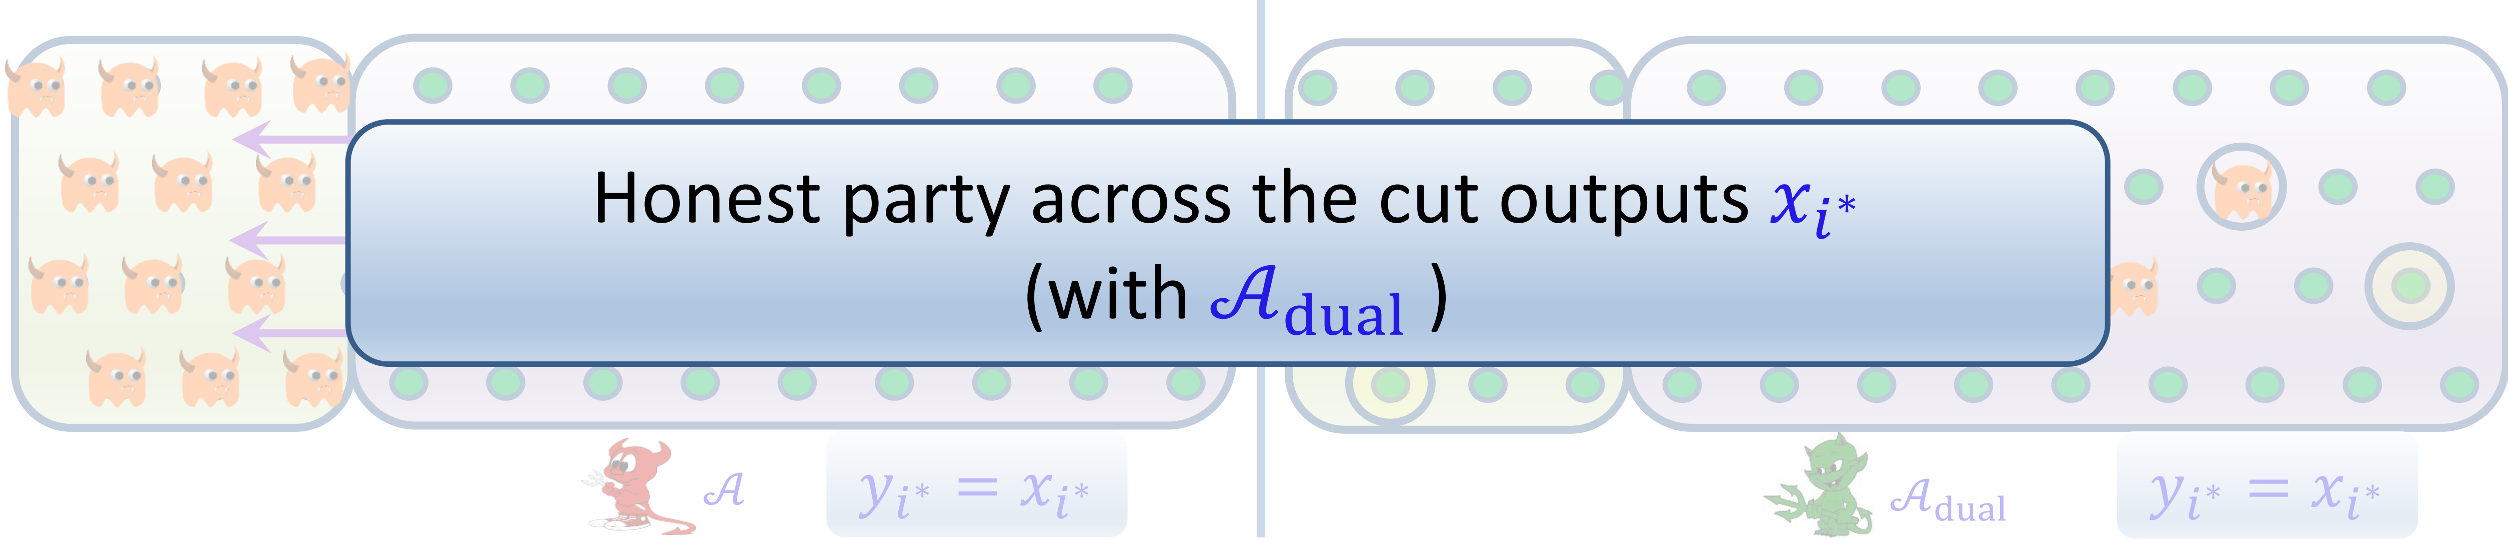

Honest party across the cut outputs x_{i^*}
(with $\mathcal{A}_{\text{dual}}$)

Problem 2: Bounding Information on x_{i^*}

1) The input x_{i^*} is a random n-bit string
2) Let $(S, \bar{S})$ be the $\alpha(n)$-cut at the end of the protocol
3) End of Phase 1: $\text{view}_{\text{Honest}}(\bar{S})$ is function of **red** execution (ind. of x_{i^*})
4) End of Phase 2: **only** new info is identity of cut $(S, \bar{S})$ (all else is simulatable)
5) Graph-theoretic Thm: $\exists$ at most 2^{c-1} possible cuts (c bits of info)
6) $\Rightarrow H(x_{i^*} | \text{view}_{\text{Honest}}(\bar{S})) \geq n - c$

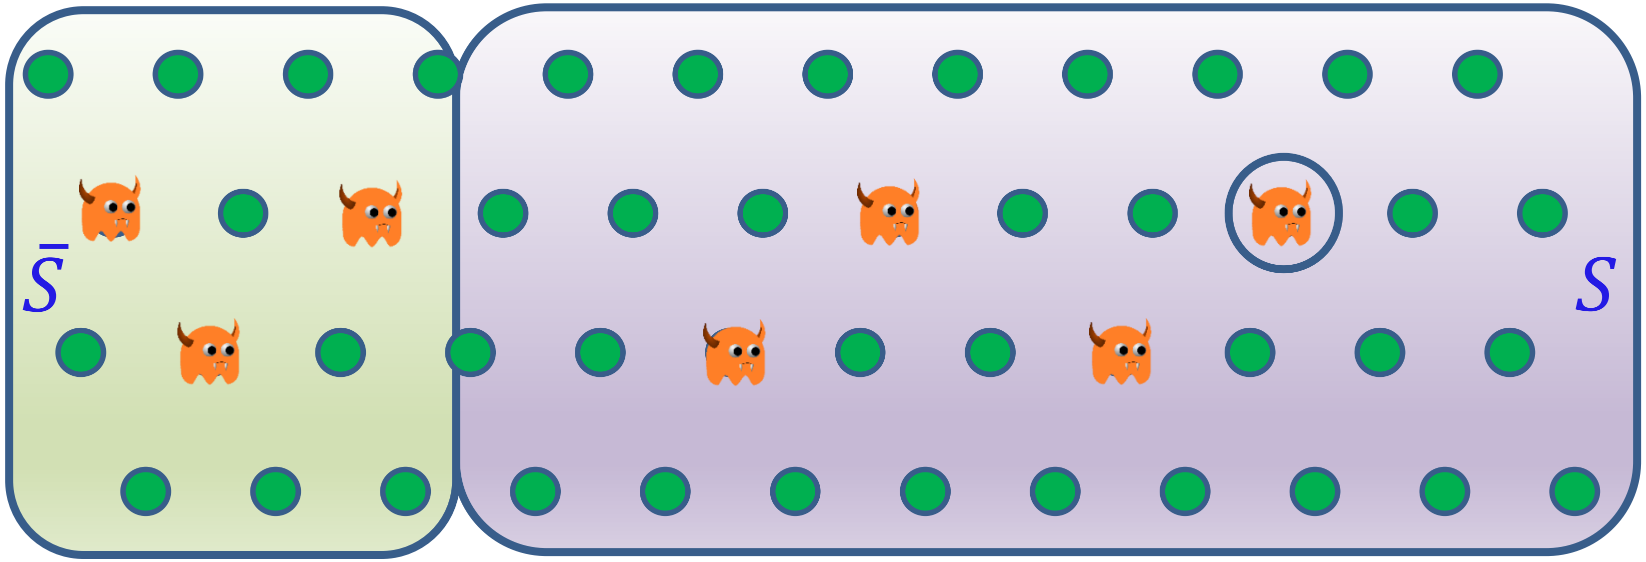

Problem 2: Bounding Information on x_{i^*}

1) The input x_{i^*} is a random n-bit string
2) Let $(S, \bar{S})$ be the $\alpha(n)$-cut at the end of the protocol
3) End of Phase 1: $\mathrm{view}_{\mathrm{Honest}}(\bar{S})$ is function of **red** execution (ind. of x_{i^*})
4) End of Phase 2: **only** new info is identity of cut $(S, \bar{S})$ (all else is simulatable)
5) Graph-theoretic Thm: $\exists$ at most 2^{c-1} possible cuts (c bits of info)
6) $\Rightarrow H(x_{i^*} | \mathrm{view}_{\mathrm{Honest}}(\bar{S})) \geq n - c$

Honest party across the cut outputs x_{i^*}
(with $\mathcal{A}_{\mathrm{dual}}$)

Contradiction

Recap of the Attack

The adv $\mathcal{A}$

Isolate random P_{i^*}

Identify the partition

Block crossing messages

P_{i^*} is honest → $\Pr[y_{i^*} = x_{i^*}] \geq 1 - negl$

no honest party across the cut

Info is conveyed across the cut

The adv $\mathcal{A}_{\mathbf{dual}}$

Corrupt random P_{i^*} & emulate phase 1

Identify the partition

Block crossing messages

$\exists$ party w/ same view → $\Pr[y_{i^*} = x_{i^*}] \geq 1/poly$

$\exists$ honest party across the cut → $\Pr[y_{i^*} = x_{i^*}] = negl$

x_{i^*} has high entropy → $\Pr[y_{i^*} = x_{i^*}] = negl$

Summary

Initiate a **foundational study** of dynamic graph model

Upper bound:

SFE protocols with **non-expander** graph (in **PKI model**):

- Static/adaptive corruptions
- Information-theoretic/computational security
- With/out polylog locality

Lower bound:

$\exists f$ s.t. every secure protocol for f induces an **expander**

- Adaptive corruptions, **CRS model**

Open Questions

- Fill the **gap** between **upper** & **lower** bounds
  - Adaptive corruptions

• Trusted setup (PKI) • Hidden channels	• No setup • Private (visible) channels

- What **other graph properties** are necessary for MPC?
- New connection between **graph theory** and **MPC**
  - Necessity of expansion $\Rightarrow$ new comm. complexity lower bounds?

Thank You