Application-Driven Coding Theory Workshop, Simons Institute, Berkeley

Sparsity and Privacy in Distributed Matrix Multiplication

Rawad Bitar

Joint work with

Maximilian Egger, Marvin Xhemrishi and Antonia Wachter-Zeh

School of Information Computation and Technology
Technical University of Munich, Germany

March 8, 2024

Deutsche Forschungsgemeinschaft
German Research Foundation

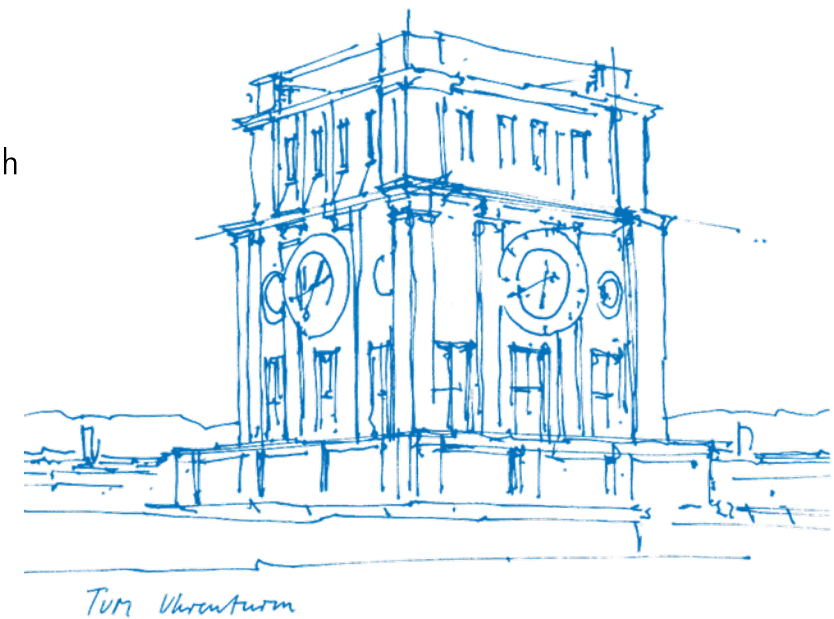

Tremendous Amount of Data Generated and Analyzed

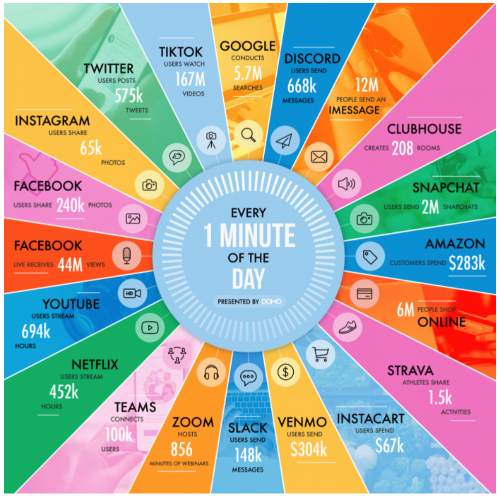

Figure: Data Created per Minute (2021) [1]

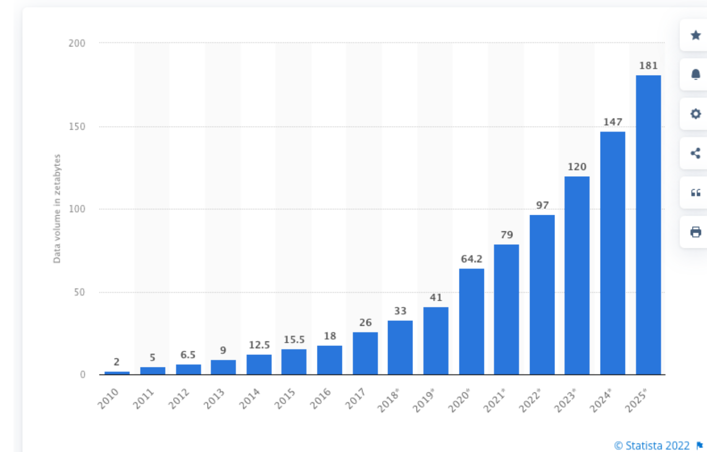

Figure: History of Worldwide Data (2021) [2]

Our main concerns:
Privacy and Efficiency in distributed learning

[1] https://dailyinfographic.com/how-much-data-is-generated-every-minute
[2] https://www.statista.com/statistics/871513/worldwide-data-created/

Distributed Learning Model

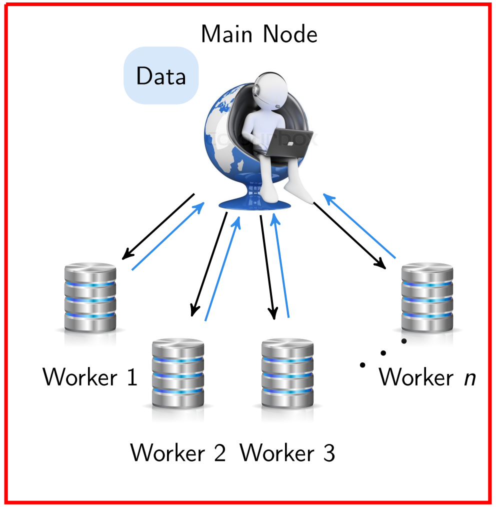

This talk

Main Node – Workers

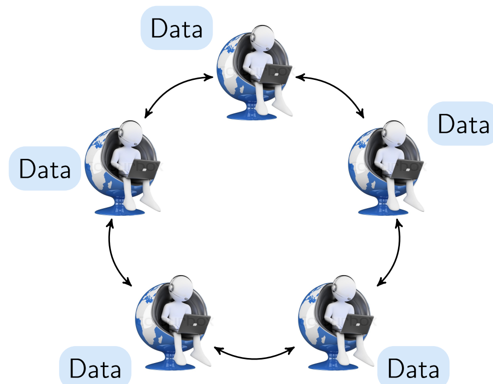

Decentralized Learning

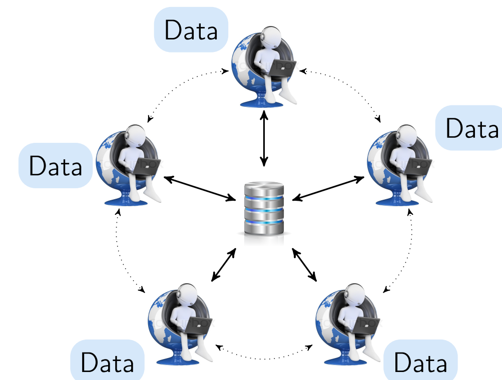

Federated Learning

Main Challenges in Distributed Learning

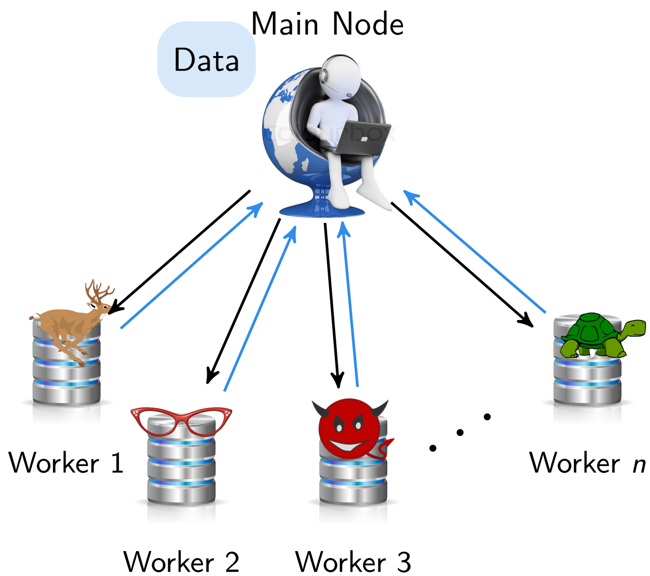

- **Stragglers:** *Slow or unresponsive* workers
- **Heterogeneity:** *Different time-varying* computing power of the workers
- **Privacy:** Workers *collude* to gain knowledge of main node's data
- **Security:** Workers are *malicious* and try to jam the computation
- **Efficiency:** Reduce *overall run-time* and *compute time* of the workers

Main Challenges in Distributed Learning

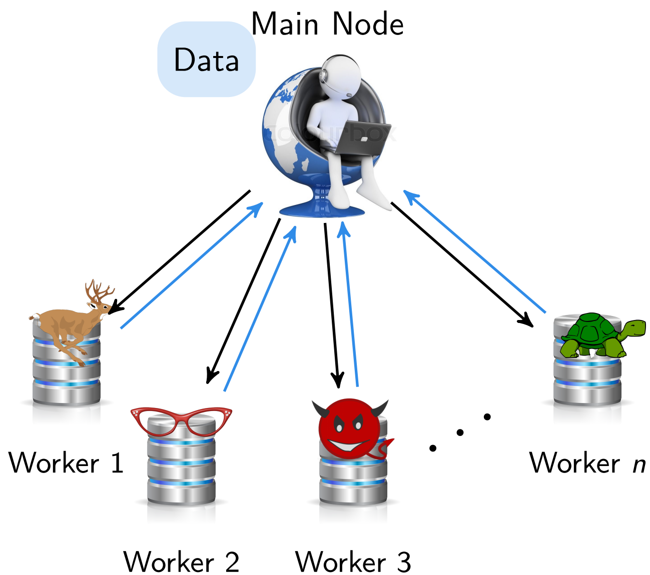

- **Stragglers:** *Slow or unresponsive* workers
- **Heterogeneity:** *Different time-varying* computing power of the workers
- **Privacy:** Workers *collude* to gain knowledge of main node's data
- **Security:** Workers are *malicious* and try to jam the computation
- **Efficiency:** Reduce *overall run-time* and *compute time* of the workers

In this talk

Efficiency (sparsity), privacy and stragglers.

System Model: Computation, Sparsity and Privacy

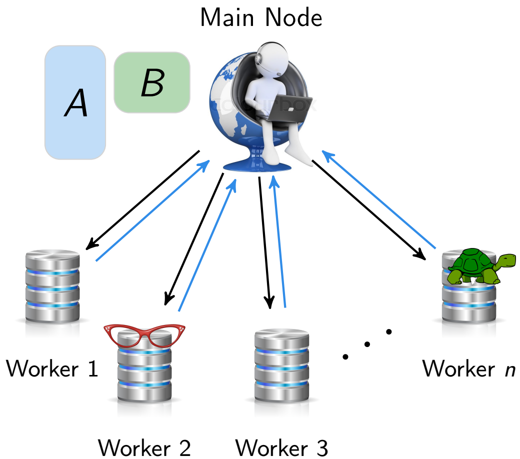

- **Data:** Sparse private matrices in $\mathbb{F}_q$

$$\Pr(A_{i,j} = a) = \begin{cases} s_A & \text{for } a = 0, \\ \frac{1 - s_A}{q - 1} & \text{otherwise} \end{cases}$$

- **Privacy:** IT privacy of A and B
- **No collusion:** Each worker eavesdrops alone
- **Stragglers:** Slow or unresponsive workers
- **Efficiency:** *sparsity* of matrices assigned to the workers

System Model: Computation, Sparsity and Privacy

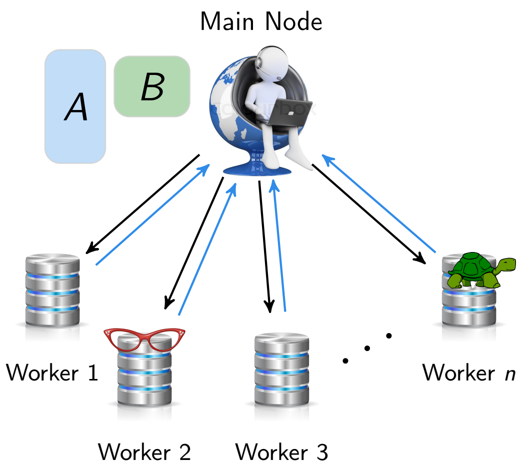

- **Data:** Sparse private matrices in $\mathbb{F}_q$

$$\Pr(A_{i,j} = a) = \begin{cases} s_\mathrm{A} & \text{for } a = 0, \\ \dfrac{1 - s_\mathrm{A}}{q-1} & \text{otherwise} \end{cases}$$

- **Privacy:** IT privacy of A and B
- **No collusion:** Each worker eavesdrops alone
- **Stragglers:** Slow or unresponsive workers
- **Efficiency:** *sparsity* of matrices assigned to the workers

Desired coding scheme

Encode A and B satisfying
- Privacy constraints
- Best sparsity in the codewords
- Straggler tolerance

Outline

Outline

Sparsity and Perfect IT Privacy

Trade-Off Between Sparsity and Privacy

- Sparse One-Time Pad
- Sparse Shamir Secret Sharing

Numerical Observations

Conclusion

Encoding and Privacy Measure

Information-Theoretic Privacy

Definition:

- Observation is statistically independent from the private data, i.e., $\mathrm{I}(\text{private data}; \text{observation}) = 0$

Assumptions:

\+ Adversary with unbounded computation power

\- Limited number of collusions

Encoding and Privacy Measure

Information-Theoretic Privacy

Definition:

- Observation is statistically independent from the private data, i.e., $\mathrm{I}(\text{private data}; \text{observation}) = 0$

Assumptions:

\+ Adversary with unbounded computation power

\- Limited number of collusions

Variations of Information-Theoretic privacy

- Perfect: $\mathrm{I}(\text{private data}; \text{observation}) = 0$ *Usual privacy measure*
- Strong: $\mathrm{I}(\text{private data}; \text{observation}) = \varepsilon \xrightarrow{\text{when the data is large}} 0$
- Weak: $\mathrm{I}(\text{private data}; \text{observation}) = \varepsilon > 0$

Encoding and Privacy Measure

Information-Theoretic Privacy

Definition:

- Observation is statistically independent from the private data, i.e., $\mathrm{I}(\text{private data}; \text{observation}) = 0$

Assumptions:

+ Adversary with unbounded computation power
- Limited number of collusions

Encoding

- Draw random matrices R and S
- $\mathsf{A} \rightarrow f_{\mathsf{A}}(x) = \mathsf{A} + x\mathsf{R}$
- $\mathsf{B} \rightarrow g_{\mathsf{B}}(x) = \mathsf{B} + x\mathsf{S}$
- Assign $f_{\mathsf{A}}(\alpha_i)$ and $g_{\mathsf{B}}(\alpha_i)$ to worker i

Privacy guarantee

- Depends on how R and S are drawn

Variations of Information-Theoretic privacy

- Perfect: $\mathrm{I}(\text{private data}; \text{observation}) = 0$ *Usual privacy measure*
- Strong: $\mathrm{I}(\text{private data}; \text{observation}) = \varepsilon \xrightarrow{\text{when the data is large}} 0$
- Weak: $\mathrm{I}(\text{private data}; \text{observation}) = \varepsilon > 0$

Implication of Perfect Privacy

$$A \in \mathbb{F}_q^{s \times r} \longrightarrow \boxed{\begin{array}{l}\text{Secret Sharing} \\ (n, k, z) = (n, 2, 1)\end{array}}$$

Private matrix

Implication of Perfect Privacy

Randomness

$$\mathsf{R} \sim U(\mathbb{F}_q^{s\times r})$$

$$\mathsf{A} \in \mathbb{F}_q^{s\times r} \longrightarrow \boxed{\text{Secret Sharing } (n,k,z) = (n,2,1)}$$

Private matrix

Implication of Perfect Privacy

Randomness

$\mathrm{R} \sim U(\mathbb{F}_q^{s\times r})$

Private matrix

$\mathrm{A} \in \mathbb{F}_q^{s\times r}$

Secret Sharing

$(n, k, z) = (n, 2, 1)$

$\mathrm{A} + \mathrm{R} \sim U(\mathbb{F}_q^{s\times r})$

$\mathrm{A} + 2\mathrm{R} \sim U(\mathbb{F}_q^{s\times r})$

$\vdots$

$\mathrm{A} + n\mathrm{R} \sim U(\mathbb{F}_q^{s\times r})$

n Output shares

Implication of Perfect Privacy

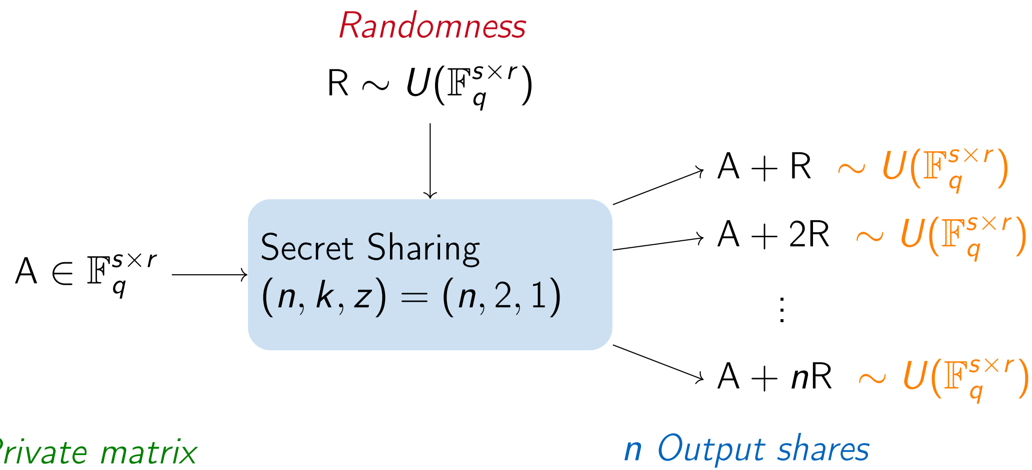

- In several applications, e.g., medical imaging, data is represented by sparse matrices (non-uniform)

Problem of perfect privacy

Output shares have uniform distribution $\Rightarrow$ Higher computation complexity.

Outline

Trading Off Sparsity vs. Privacy

- Insisting on *perfect* privacy does not allow sparsity

Lemma: fundamental tradeoff [BEWX24]

For $k = 2$ and $z = 1$, perfect privacy can be achieved if and only if the entries of R are i.i.d uniformly at random.

[1] [BEWX24] **R. Bitar**, M. Egger, A. Wachter-Zeh, and M. Xhemrishi, "Sparsity and privacy in secret sharing: A fundamental trade-off," *accepted in IEEE Transactions on Information Forensics and Security*, 2024

Trading Off Sparsity vs. Privacy

- Insisting on *perfect* privacy does not allow sparsity

Lemma: fundamental tradeoff [BEWX24]

For $k = 2$ and $z = 1$, perfect privacy can be achieved if and only if the entries of R are i.i.d uniformly at random.

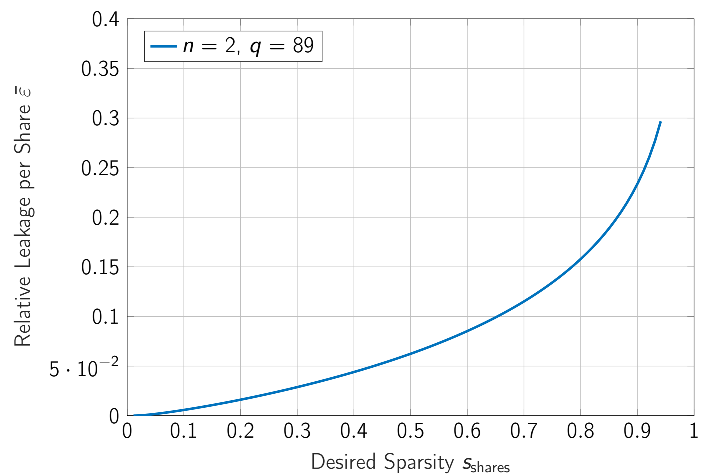

Figure: Relative leakage $\bar{\varepsilon} = \dfrac{\mathrm{I}(\mathrm{A} + x\mathrm{R}; \mathrm{A})}{\mathrm{H}(\mathrm{A})}$ as function of desired sparsity.

[1][BEWX24] **R. Bitar**, M. Egger, A. Wachter-Zeh, and M. Xhemrishi, "Sparsity and privacy in secret sharing: A fundamental trade-off," *accepted in IEEE Transactions on Information Forensics and Security*, 2024

Relax to Weak Privacy

Main Idea

Design R *dependently* on A, i.e., design a conditional PMF $\mathrm{P}_{\mathsf{R}|\mathsf{A}}(R_{ij} = r | A_{ij} = a)$.
$\Rightarrow$ This allows for sparsity, but leaks information about A.

Challenge

Given a desired sparsity of the shares, design R to get the smallest leakage.

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \\ \dfrac{1 - p_1}{q - 1}, & r \neq 0, \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \\ p_3, & r = -a \\ \dfrac{1 - p_2 - p_3}{q - 2}, & r \notin \{0, -a\}. \end{cases}$$

[1] [XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(Sparisty of R)} \\ \dfrac{1-p_1}{q-1}, & r \neq 0, \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \\ p_3, & r = -a \\ \dfrac{1-p_2-p_3}{q-2}, & r \notin \{0, -a\}. \end{cases}$$

[1] [XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(Sparisty of } \mathrm{R}\textit{)} \\ \frac{1-p_1}{q-1}, & r \neq 0, \quad \textit{(iid non-zero values in } \mathrm{R}\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \\ p_3, & r = -a \\ \frac{1-p_2-p_3}{q-2}, & r \notin \{0, -a\}. \end{cases}$$

[1][XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(Sparisty of } \mathrm{R}\textit{)} \\ \dfrac{1-p_1}{q-1}, & r \neq 0, \quad \textit{(iid non-zero values in } \mathrm{R}\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \quad \textit{(keeping non-zero values in } \mathrm{A} + \mathrm{R}\textit{)} \\ p_3, & r = -a \\ \dfrac{1-p_2-p_3}{q-2}, & r \notin \{0, -a\}. \end{cases}$$

[1][XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(Sparisty of R)} \\ \dfrac{1 - p_1}{q - 1}, & r \neq 0, \quad \textit{(iid non-zero values in R)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \quad \textit{(keeping non-zero values in A + R)} \\ p_3, & r = -a \quad \textit{(Sparsity in A + R)} \\ \dfrac{1 - p_2 - p_3}{q - 2}, & r \notin \{0, -a\}. \end{cases}$$

[1][XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(Sparisty of } \mathrm{R}\textit{)} \\ \dfrac{1-p_1}{q-1}, & r \neq 0, \quad \textit{(iid non-zero values in } \mathrm{R}\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_2, & r = 0 \quad \textit{(keeping non-zero values in } \mathrm{A}+\mathrm{R}\textit{)} \\ p_3, & r = -a \quad \textit{(Sparsity in } \mathrm{A}+\mathrm{R}\textit{)} \\ \dfrac{1-p_2-p_3}{q-2}, & r \notin \{0, -a\}. \quad \textit{(iid non-zero values in } \mathrm{A}+\mathrm{R}\textit{)} \end{cases}$$

[1][XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Sparse One-time Pad

Constuction: Sparse One-time Pad [XEB21]

Use the shares as R and A + R. Design R as follows:

$$\Pr\{R_{ij}=r|A_{ij}=0\}=\begin{cases} p_1, & r=0 \quad \textit{(Sparisty of } \mathrm{R}\textit{)} \\ \dfrac{1-p_1}{q-1}, & r\neq 0, \quad \textit{(iid non-zero values in } \mathrm{R}\textit{)} \end{cases}$$

$$\Pr\{R_{ij}=r|A_{ij}=a\}=\begin{cases} p_2, & r=0 \quad \textit{(keeping non-zero values in } \mathrm{A}+\mathrm{R}\textit{)} \\ p_3, & r=-a \quad \textit{(Sparsity in } \mathrm{A}+\mathrm{R}\textit{)} \\ \dfrac{1-p_2-p_3}{q-2}, & r\notin\{0,-a\}. \quad \textit{(iid non-zero values in } \mathrm{A}+\mathrm{R}\textit{)} \end{cases}$$

Proposition: Sparsity as function of the PMF

$$s_{\mathrm{R}}=p_1 s+p_2(1-s),$$
$$s_{\mathrm{A+R}}=p_1 s+p_3(1-s).$$

[1][XEB21] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Minimizing the Leakage

Minimizing Entry-Wise Leakage

Let $\mathcal{P}$ be the set of all q^2 values of $\mathrm{P}_{\mathsf{R}|\mathsf{A}}$, then the optimal leakage is

$$\begin{aligned} \mathrm{L}_{\mathsf{opt}} &= \min_{\mathcal{P}} \mathrm{I}_q\left(\mathsf{R};\mathsf{A}\right) + \mathrm{I}_q\left(\mathsf{A}+\mathsf{R};\mathsf{A}\right) \\ &= \min_{\mathcal{P}} \mathsf{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{R}}\right) + \mathsf{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{A}+\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{A}+\mathsf{R}}\right), \end{aligned}$$

and is subject to valid PMF and desired sparsities.

Minimizing the Leakage

Minimizing Entry-Wise Leakage

Let $\mathcal{P}$ be the set of all q^2 values of $\mathrm{P}_{\mathsf{R}|\mathsf{A}}$, then the optimal leakage is

$$\begin{aligned}\mathrm{L}_{\mathsf{opt}} &= \min_{\mathcal{P}} \mathrm{I}_q\left(\mathsf{R};\mathsf{A}\right) + \mathrm{I}_q\left(\mathsf{A}+\mathsf{R};\mathsf{A}\right)\\ &= \min_{\mathcal{P}} \mathrm{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{R}}\right) + \mathrm{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{A}+\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{A}+\mathsf{R}}\right),\end{aligned}$$

and is subject to valid PMF and desired sparsities.

- **Constrained Convex Optimization**
  - $\rightarrow$ For desired s_{R} and $s_{\mathsf{A}+\mathsf{R}}$, we solve convex optimization $\min_{\mathcal{P}} \mathrm{L}(p_1, p_2, p_3)$ analytically.
  - $\rightarrow$ Solution is given by root finding of degree three polynomial.
  - $\rightarrow$ For small q, numerical results are the same as optimizing over q^2 values of $\mathrm{P}_{\mathsf{R}|\mathsf{A}}$.

Minimizing the Leakage

Minimizing Entry-Wise Leakage

Let $\mathcal{P}$ be the set of all q^2 values of $\mathrm{P}_{\mathsf{R}|\mathsf{A}}$, then the optimal leakage is

$$\begin{aligned}\mathrm{L}_{\text{opt}} &= \min_{\mathcal{P}} \mathrm{I}_q\left(\mathsf{R};\mathsf{A}\right) + \mathrm{I}_q\left(\mathsf{A}+\mathsf{R};\mathsf{A}\right)\\ &= \min_{\mathcal{P}} \mathsf{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{R}}\right) + \mathsf{D}_{\mathsf{KL}}\left(\mathrm{P}_{\mathsf{A},\mathsf{A}+\mathsf{R}} \| \mathrm{P}_{\mathsf{A}}\mathrm{P}_{\mathsf{A}+\mathsf{R}}\right),\end{aligned}$$

and is subject to valid PMF and desired sparsities.

- **Constrained Convex Optimization**
  - $\rightarrow$ For desired s_{R} and $s_{\mathsf{A}+\mathsf{R}}$, we solve convex optimization $\min_{\mathcal{P}} \mathrm{L}(p_1, p_2, p_3)$ analytically.
  - $\rightarrow$ Solution is given by root finding of degree three polynomial.
  - $\rightarrow$ For small q, numerical results are the same as optimizing over q^2 values of $\mathrm{P}_{\mathsf{R}|\mathsf{A}}$.

$\Rightarrow$ Results in optimal privacy guarantees, i.e., minimal leakage.

Setting of Partly-Trusted/Untrusted Workers

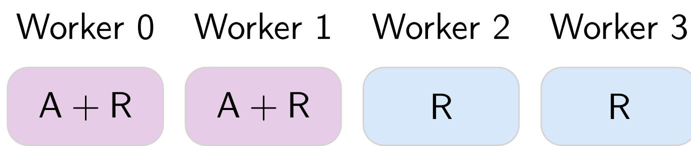

Figure: Two non-communicating clusters. One completely untrusted, one partially trusted.

Setting of Partly-Trusted/Untrusted Workers

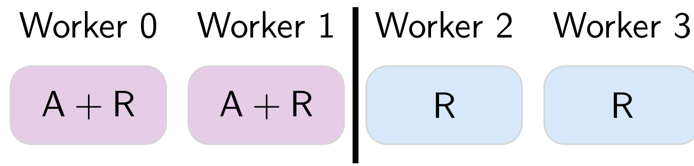

Figure: Two non-communicating clusters. One completely untrusted, one partially trusted.

Setting of Partly-Trusted/Untrusted Workers

Worker 0	Worker 1	Worker 2	Worker 3
A + R	A + R	R	R

Figure: Two non-communicating clusters. One completely untrusted, one partially trusted.

- Choose $p_1 = p_2 = p_3 = p$ such that $\mathrm{I}_q\left(\mathsf{A}+\mathsf{R};\mathsf{A}\right) = 0$
- Sparsity of the shares become

$$s_{\mathsf{R}} = p\frac{(sq-1)}{q-1} + \frac{(1-s)}{q-1}, \quad \text{and} \quad s_{\mathsf{A}+\mathsf{R}} = p\,.$$

- Choose p to satisfy the desired sparsity constraint

Sparse $(n, 2, 1)$ Secret Sharing

Constuction: Sparse Secret Sharing [EXWB24]

Use the encoding polynomial $f_{\mathrm{A}}(x) = \mathrm{A} + x\mathrm{R}$. Choose n distinct non-zero symbols $\alpha_1, \ldots, \alpha_n$ from $\mathbb{F}_q$. Share i is the evaluation $f(\alpha_i)$. Design the entries of R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \\ \dfrac{1-p_1}{q-1}, & r \neq 0, \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_s, & r \in \{-\dfrac{a}{\alpha_i}\}_{i \in [n]} \\ \dfrac{1-p_s}{q-1}, & r \notin \{-\dfrac{a}{\alpha_i}\}_{i \in [n]}. \end{cases}$$

[1][EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparse $(n, 2, 1)$ Secret Sharing

Constuction: Sparse Secret Sharing [EXWB24]

Use the encoding polynomial $f_{\mathrm{A}}(x) = \mathrm{A} + x\mathrm{R}$. Choose n distinct non-zero symbols $\alpha_1, \ldots, \alpha_n$ from $\mathbb{F}_q$. Share i is the evaluation $f(\alpha_i)$. Design the entries of R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(sparsity of } R\textit{)} \\ \frac{1 - p_1}{q - 1}, & r \neq 0, \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_s, & r \in \{-\frac{a}{\alpha_i}\}_{i \in [n]} \\ \frac{1 - p_s}{q - 1}, & r \notin \{-\frac{a}{\alpha_i}\}_{i \in [n]}. \end{cases}$$

[1][EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparse $(n, 2, 1)$ Secret Sharing

Constuction: Sparse Secret Sharing [EXWB24]

Use the encoding polynomial $f_{\mathrm{A}}(x) = \mathrm{A} + x\mathrm{R}$. Choose n distinct non-zero symbols $\alpha_1, \ldots, \alpha_n$ from $\mathbb{F}_q$. Share i is the evaluation $f(\alpha_i)$. Design the entries of R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(sparsity of } R\textit{)} \\ \frac{1 - \mathrm{p}_1}{\mathrm{q} - 1}, & r \neq 0, \quad \textit{(iid non-zero values in } R\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_s, & r \in \{-\frac{a}{\alpha_i}\}_{i \in [n]} \\ \frac{1 - p_s}{q - 1}, & r \notin \{-\frac{a}{\alpha_i}\}_{i \in [n]}. \end{cases}$$

[1][EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparse $(n, 2, 1)$ Secret Sharing

Constuction: Sparse Secret Sharing [EXWB24]

Use the encoding polynomial $f_{\mathrm{A}}(x) = \mathrm{A} + x\mathrm{R}$. Choose n distinct non-zero symbols $\alpha_1, \ldots, \alpha_n$ from $\mathbb{F}_q$. Share i is the evaluation $f(\alpha_i)$. Design the entries of R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(sparsity of } R\textit{)} \\ \dfrac{1 - p_1}{q - 1}, & r \neq 0, \quad \textit{(iid non-zero values in } R\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} \mathrm{p_s}, & r \in \{-\dfrac{a}{\alpha_i}\}_{i \in [n]} \quad \textit{(zero in } A + \alpha_i R\textit{)} \\ \dfrac{1 - p_s}{q - 1}, & r \notin \{-\dfrac{a}{\alpha_i}\}_{i \in [n]}. \end{cases}$$

[1][EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparse $(n, 2, 1)$ Secret Sharing

Constuction: Sparse Secret Sharing [EXWB24]

Use the encoding polynomial $f_{\mathrm{A}}(x) = \mathrm{A} + x\mathrm{R}$. Choose n distinct non-zero symbols $\alpha_1, \ldots, \alpha_n$ from $\mathbb{F}_q$. Share i is the evaluation $f(\alpha_i)$. Design the entries of R as follows:

$$\Pr\{R_{ij} = r | A_{ij} = 0\} = \begin{cases} p_1, & r = 0 \quad \textit{(sparsity of } R\textit{)} \\ \dfrac{1 - p_1}{q - 1}, & r \neq 0, \quad \textit{(iid non-zero values in } R\textit{)} \end{cases}$$

$$\Pr\{R_{ij} = r | A_{ij} = a\} = \begin{cases} p_s, & r \in \{-\frac{a}{\alpha_i}\}_{i \in [n]} \quad \textit{(zero in } A + \alpha_i R\textit{)} \\ \dfrac{1 - \mathrm{p_s}}{\mathrm{q} - 1}, & r \notin \{-\frac{a}{\alpha_i}\}_{i \in [n]}. \quad \textit{(iid non-zeros in } A + \alpha_i R\textit{)} \end{cases}$$

[1] [EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparsity of our Sparse Secret Sharing

Lemma: Sparsity of the shares [EXWB24]

Given a matrix A with sparsity s_{A}, the sparsity s_{share} of the shares is expressed as

$$s_{\mathrm{share}} = p_1 s_{\mathrm{A}} + p_s(1 - s_{\mathrm{A}}).$$

[1] [EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Sparsity of our Sparse Secret Sharing

Lemma: Sparsity of the shares [EXWB24]

Given a matrix A with sparsity s_{A}, the sparsity s_{share} of the shares is expressed as

$$s_{\text{share}} = p_1 s_{\mathrm{A}} + p_s(1 - s_{\mathrm{A}}).$$

✓ Sparsity *increases* with p_1 and p_s, e.g., $p_1 = 1, p_s = 1$ *maximum* sparsity

× So does the *information leakage* $\mathrm{I}(\mathrm{A} + x\mathrm{R}; \mathrm{A})$, e.g., $p_1 = 1, p_s = 1 \Rightarrow$ R is a multiple of $-\mathrm{A}$

[1] [EXWB24] M. Egger, M. Xhemrishi, A. Wachter-Zeh, and R. **Bitar**, "Sparse and private distributed matrix multiplication with straggler tolerance," in *IEEE International Symposium on Information Theory (ISIT)*, 2024

Theorem: Shares with minimum leakage

Given a desired sparsity s_{shares}, the leakage $\mathrm{I}(\mathrm{A} + x\mathrm{R}; \mathrm{A})$ of the n shares is *minimized* by setting $p_s = p_s^\star$ as the real root of the polynomial $\sum_{j=0}^{n+1} b_j p_s^j$ in p_s that satisfies $0 \leq p_s(1 - s_{\mathrm{A}}) \leq \min\{s_{\text{shares}}, \frac{1}{n}\}$, for $s_1 \triangleq s_{\text{share}}/(1-s)$, $s_2 \triangleq (s_{\mathrm{A}} - s_{\text{shares}})/(1 - s_{\mathrm{A}})$ and $c \triangleq (q-1)/(q-n)^n$ and

$$\begin{aligned}
b_{n+1} &= -1 - c(-n)^n \\
b_n &= c\big(s_1(-n)^n - n(-n)^{n-1}\big) - s_2 \\
b_i &= c\left(s_1\binom{n}{i}(-n)^i - \binom{n}{i-1}(-n)^{i-1}\right), \forall i \in [n-1] \\
b_0 &= cs_1.
\end{aligned}$$

Then, p_1 is computed as

$$p_1^\star = \frac{s_{\text{shares}} - p_s^\star(1 - s_{\mathrm{A}})}{s_{\mathrm{A}}}.$$

Proof Idea

To prove that the values $p_s^\star$ and $p_1^\star$ minimize the leakage, we do the following

- **Assume** sparsity is given and is *same* for all shares
- Prove that the leakage is a convex function of the conditional PMF $\mathrm{P}_{\mathsf{R}|\mathsf{A}}(R_{ij} = r | A_{ij} = a)$
- Find the leakage as function of p_s and p_1 for our construction
- Solve the non-linear convex optimization problem using Lagrange multipliers

Reducing the Computation Load

Group 0	Group 1	Group 2	Group 3
$f_0(x), g_0(x)$	$f_1(x), g_1(x)$	$f_2(x), g_2(x)$	$f_3(x), g_3(x)$

- Divide the matrices A and B into m smaller chunks such that $\frac{n}{m} = \sigma + 3$
- Compute and assign evaluations of $f_i(x)$ and $g_i(x)$ to workers of group i, each encoding a chunk of A and B

[1][DFHJCG19] S. Dutta, M. Fahim, F. Haddadpour, H. Jeong, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 278–301, 2019

Reducing the Computation Load

Group 0	Group 1	Group 2	Group 3
$f_0(x), g_0(x)$	$f_1(x), g_1(x)$	$f_2(x), g_2(x)$	$f_3(x), g_3(x)$

- Divide the matrices A and B into m smaller chunks such that $\frac{n}{m} = \sigma + 3$
- Compute and assign evaluations of $f_i(x)$ and $g_i(x)$ to workers of group i, each encoding a chunk of A and B

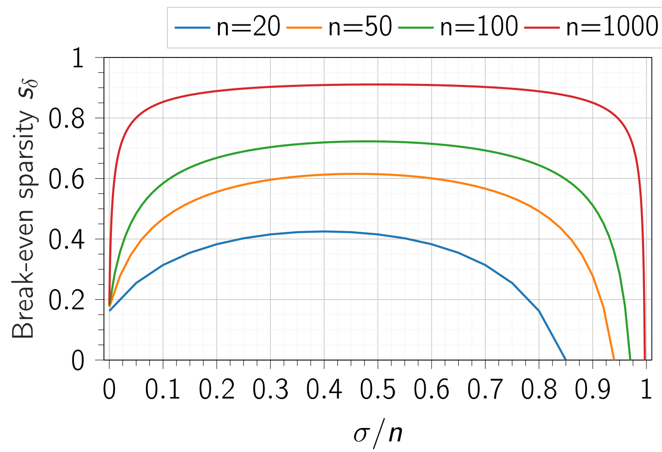

Figure: Sparsity values above which our scheme is beneficial over [DFHJCG19] polynomial codes.

[1][DFHJCG19] S. Dutta, M. Fahim, F. Haddadpour, H. Jeong, V. Cadambe, and P. Grover, "On the optimal recovery threshold of coded matrix multiplication," *IEEE Transactions on Information Theory*, vol. 66, no. 1, pp. 278–301, 2019

Outline

Leakage vs Scheme Parameters

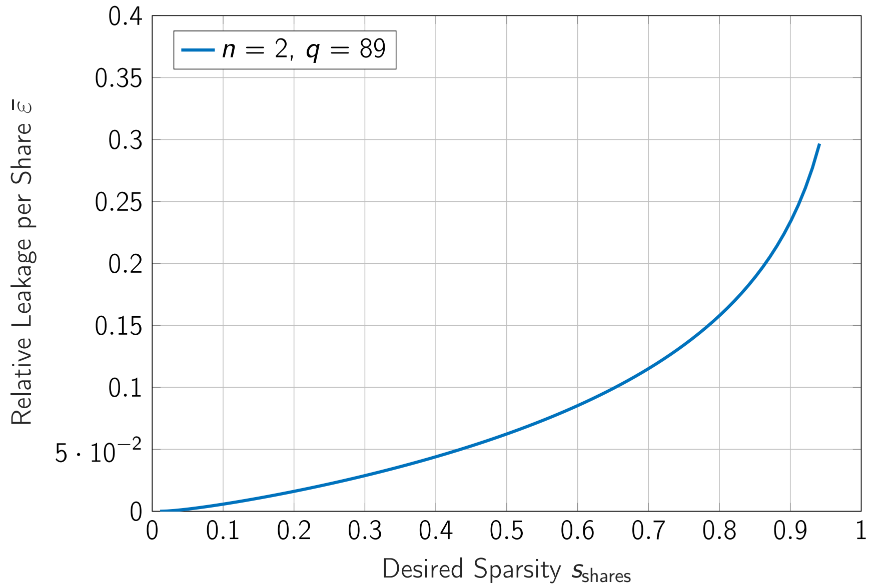

Figure: Relative leakage $\bar{\varepsilon} = \dfrac{\mathrm{I}(\mathsf{A} + x\mathsf{R}; \mathsf{A})}{\mathrm{H}(\mathsf{A})}$ as function of desired sparsity, number of shares n and field size q.

Leakage vs Scheme Parameters

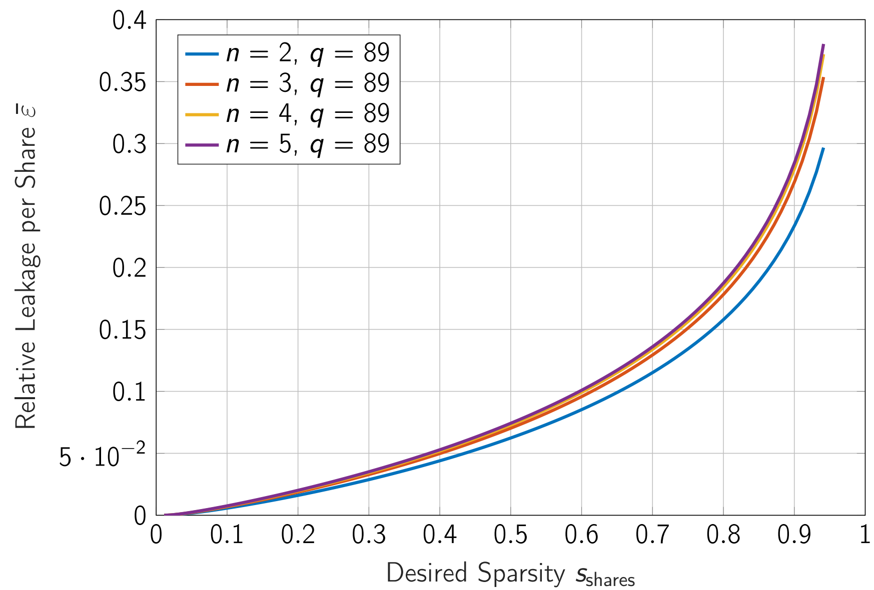

Figure: Relative leakage $\bar{\varepsilon} = \dfrac{\mathrm{I}(\mathsf{A} + x\mathsf{R}; \mathsf{A})}{\mathrm{H}(\mathsf{A})}$ as function of desired sparsity, number of shares n and field size q.

- Leakage increases with n
- Leakage decreases with q
- Leakage increase with n is less emphasized for large q

Leakage vs Scheme Parameters

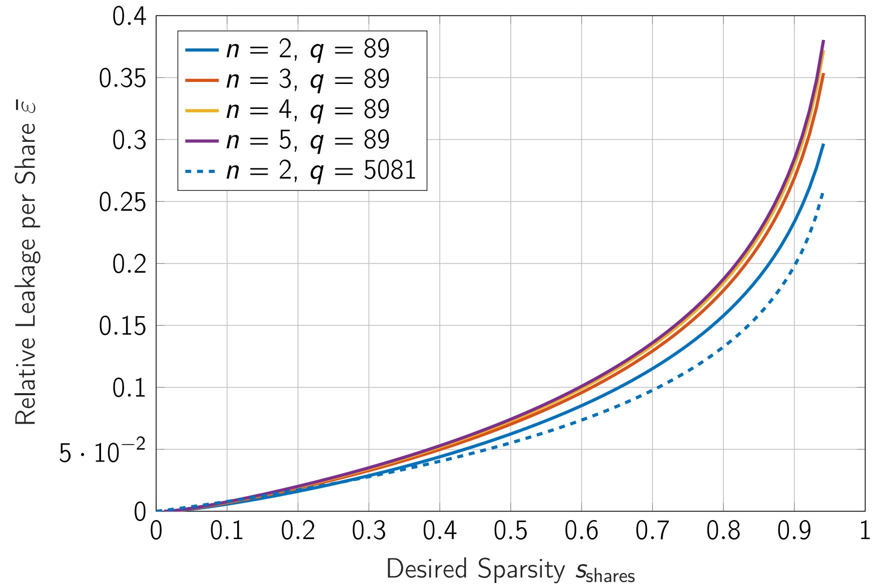

Figure: Relative leakage $\bar{\varepsilon} = \dfrac{\mathrm{I}(A + xR; A)}{\mathrm{H}(A)}$ as function of desired sparsity, number of shares n and field size q.

- Leakage increases with n
- Leakage decreases with q
- Leakage increase with n is less emphasized for large q

Leakage vs Scheme Parameters

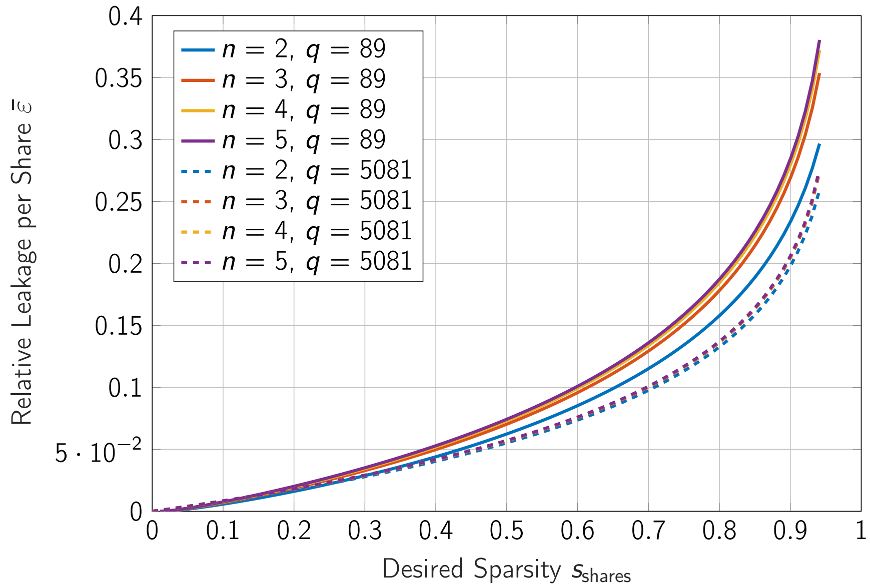

Figure: Relative leakage $\bar{\varepsilon} = \dfrac{\mathrm{I}(\mathsf{A} + x\mathsf{R}; \mathsf{A})}{\mathrm{H}(\mathsf{A})}$ as function of desired sparsity, number of shares n and field size q.

- Leakage increases with n
- Leakage decreases with q
- Leakage increase with n is less emphasized for large q

Why Same Sparsity for all Shares?

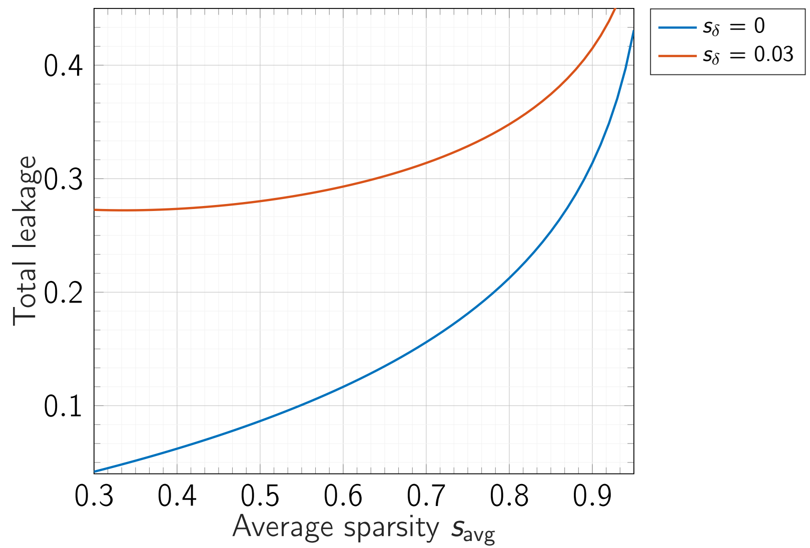

Figure: Optimal element-wise total leakage over different s_{avg} with varying s_δ for $q = 256$ and $s = 0.95$.

Lemma: Optimal sparsity for two shares [XEB22]

Sparse secret sharing with shares R and A + R give the minimal total leakage when $s_\delta \triangleq s_{\text{A+R}} - s_{\text{R}} = 0$.

[1][XEB22] M. Xhemrishi, M. Egger, and **R. Bitar**, "Efficient private storage of sparse machine learning data," in *IEEE Information Theory Workshop (ITW), Invited paper*, 2022

Matrices with Correlated Entries

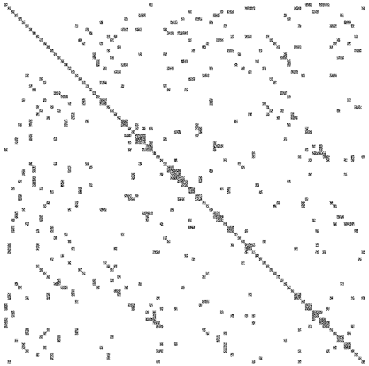

(a) Matrix A with $s \approx 0.94$

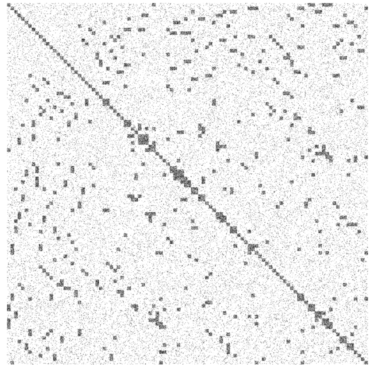

(b) Share $f(\alpha_i)$ of A with sparsity $s_{\text{share}} \approx 0.85$

Figure: A depiction of the impact of correlated entries on the privacy guarantee.

- Naively encoding matrices with correlated entries using our sparse secret sharing may leak more information than desired

Matrices with Correlated Entries

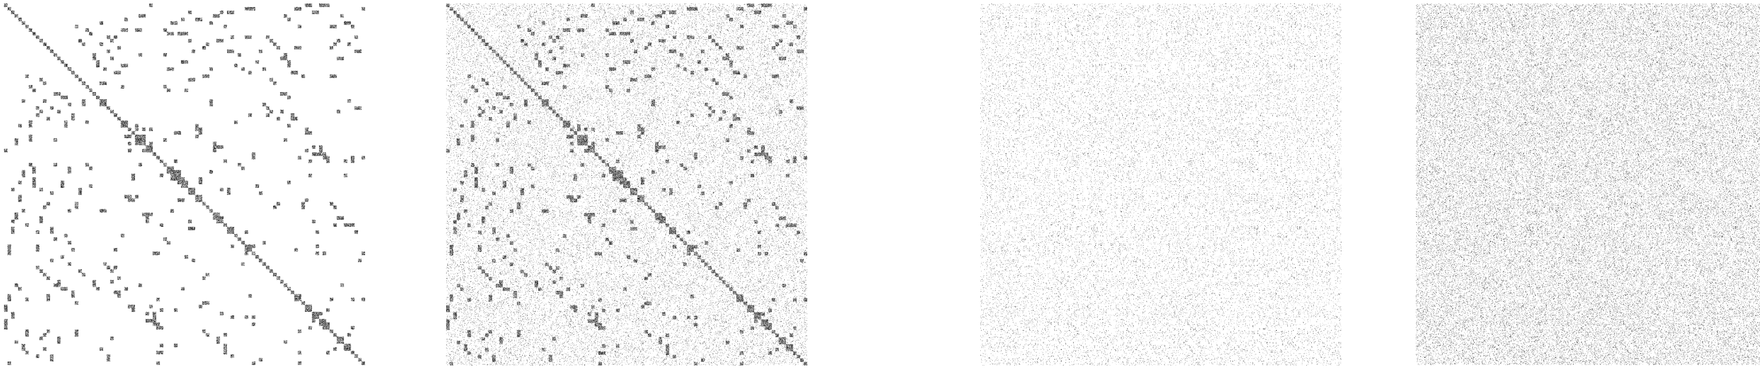

(a) Matrix A with $s \approx 0.94$

(b) Share $f(\alpha_i)$ of A with sparsity $s_{\text{share}} \approx 0.85$

(c) Matrix A′ after permutation

(d) Share $f(\alpha_i)$ of A′ with sparsity $s_{\text{shares}} \approx 0.85$

Figure: A depiction of the impact of correlated entries on the privacy guarantee.

- Naively encoding matrices with correlated entries using our sparse secret sharing may leak more information than desired
- Our approach is to randomly permute the entries

Outline

Summary and Future Directions

Summary

- Private and sparse matrix-matrix multiplication with no collusions
- Fundamental trade-off between sparsity and privacy
- Optimal solution under i.i.d entries of A for multiple shares with same sparsity
- Privacy improves with q and small n
- Extra care is needed for matrices with correlated entries

Summary and Future Directions

Summary

- Private and sparse matrix-matrix multiplication with no collusions
- Fundamental trade-off between sparsity and privacy
- Optimal solution under i.i.d entries of A for multiple shares with same sparsity
- Privacy improves with q and small n
- Extra care is needed for matrices with correlated entries

Future Directions

- Improve the rate of sparsity-preserving secret sharing schemes, i.e., $k > 2$
- Sparse secret sharing with collusions, i.e., $z > 1$
- Beyond i.i.d entries of the matrices

Munich Workshop on Coding and Cryptography

- Focus on post-quantum cryptography and privacy-preserving machine learning.
- Dates: April 8 – 10, 2024.
- Takes place after the Munich Workshop on Shannon Coding Techniques.

ISIT Satellite Workshop on DNA-based Data Storage

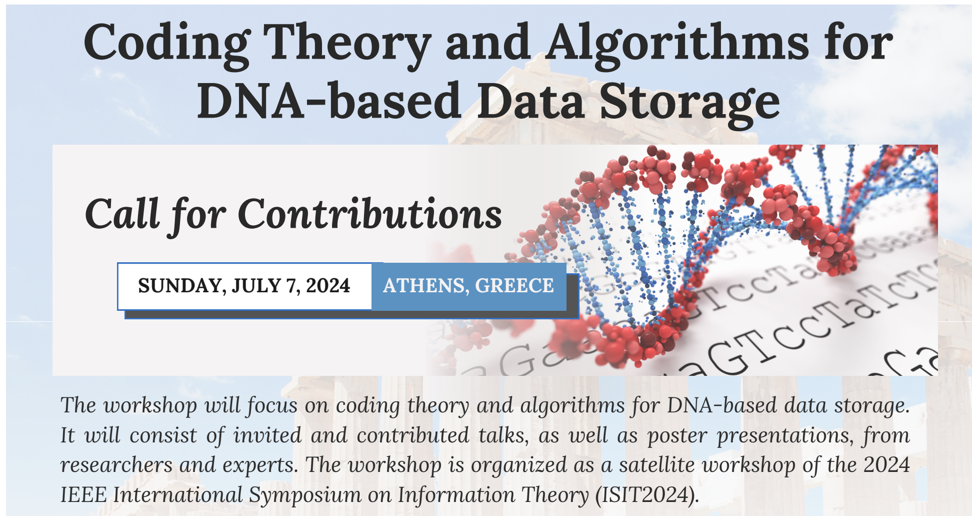

- Jointly organized with Dave Landsman from the DNA Data Storage Alliance.
- Contribution deadline: April 15, 2024.
- Designed to foster collaboration.

Thank you for your attention!

Questions?

Figure: https://arxiv.org/abs/2306.15134

Further Questions?

Rawad Bitar
Technical University of Munich
rawad.bitar@tum.de