



Efficient Decentralized Federated Singular Vector Decomposition

Di Chai¹, Junxue Zhang¹, Liu Yang¹, Yilun Jin¹, Leye Wang²,
Kai Chen¹, and Qiang Yang¹

¹Hong Kong University of Science and Technology

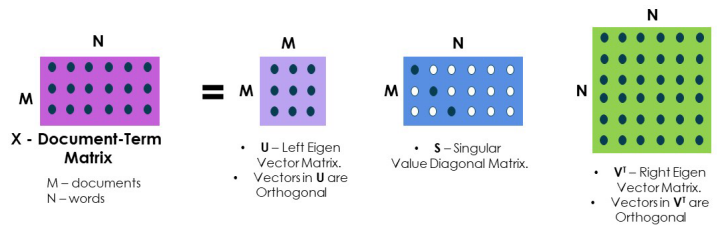
²Peking University



Outline

- **Introduction: Background and Motivation**
- **Excalibur's Matrix Protection**
- **Excalibur's Decentralized SVD Workflow**
- **Implementation and Evaluation**
- **Conclusion and Future Work**

Federated SVD is an Essential Primitive



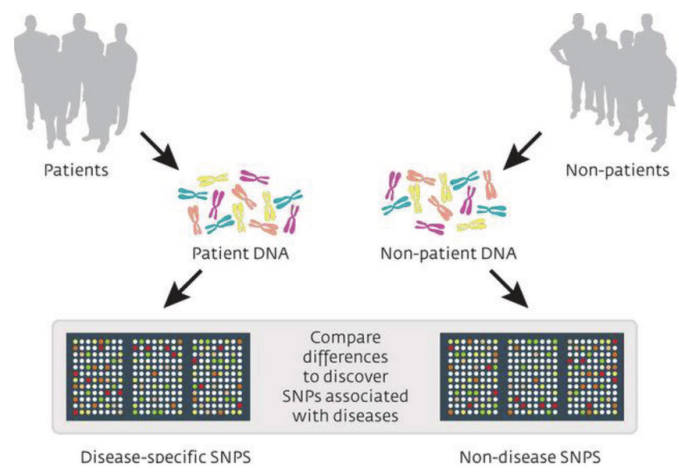
Latent semantic analysis / topic modeling

$$\min_w \|Xw - y\|_2^2 + \alpha \|w\|_2^2$$

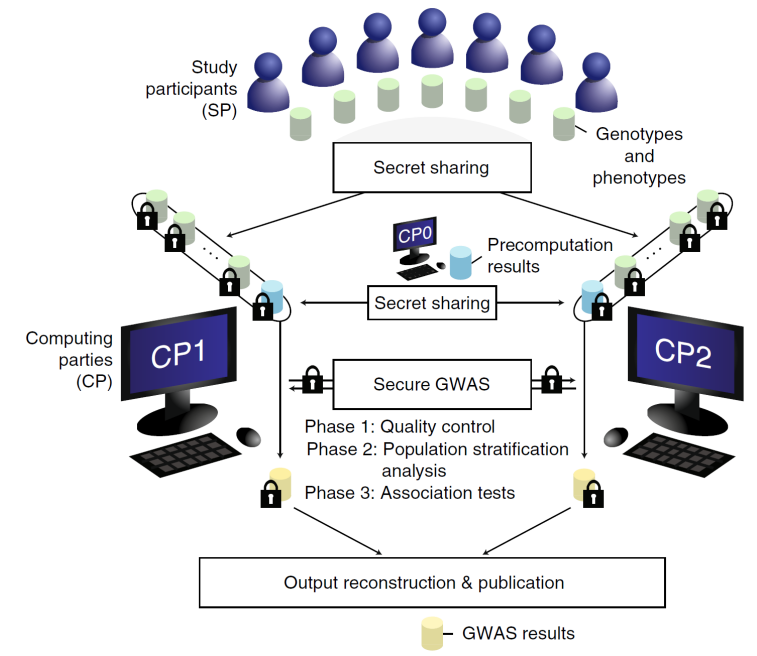
$$X = U\Sigma V^T$$

$$w = V(\Sigma^T \Sigma + \alpha I)^{-1} \Sigma^T U^T y$$

SVD works as the solver for LR.



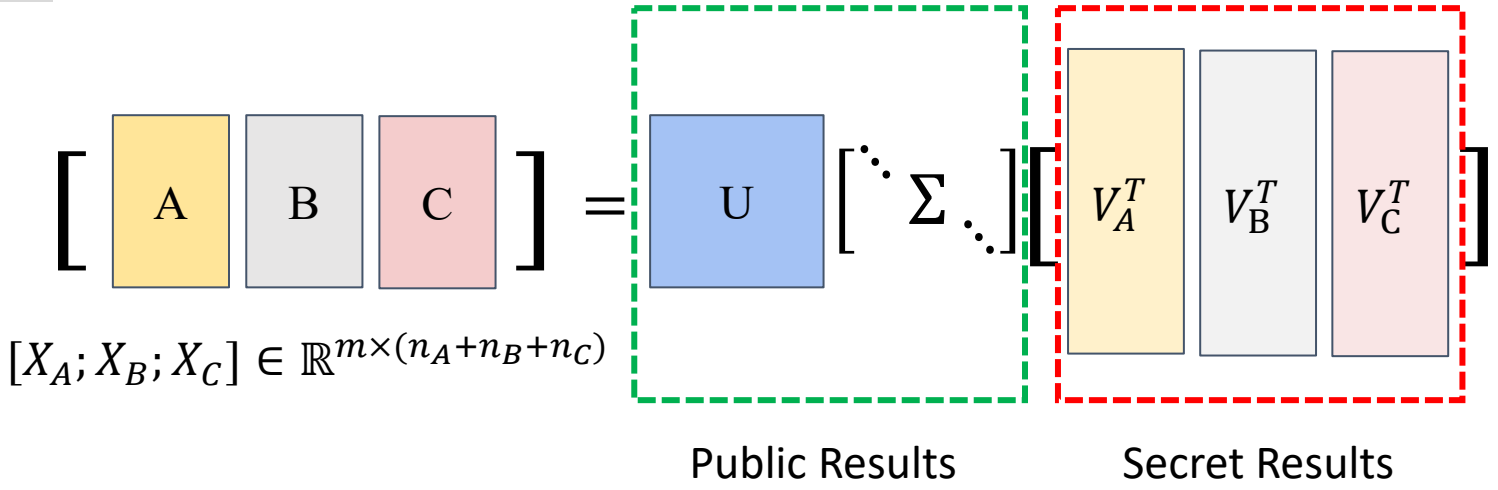
SVD-based Genome-Wide Association Studies (GWAS) require million-scale samples



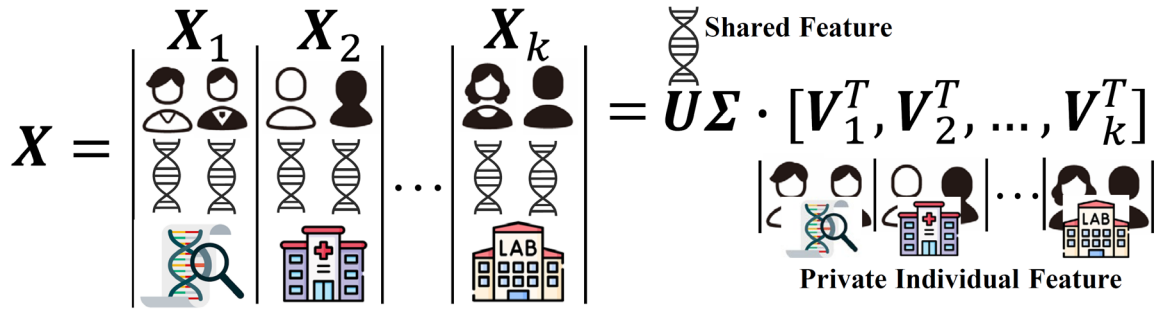
Real-world applications require combining different data sources!

Federated Singular Vector Decomposition (SVD) is an essential primitive to support many real-world **distributed application.**

Problem Definition of Federated SVD



Federated SVD factorizes matrix from multiple domains into public and secret parts.



An illustration of federated SVD in GWAS.

External Servers Downgrade the Privacy Protection

Most of the existing works rely on **external servers**.

Existing Studies	# of Servers	Job of Data Contributors	Job of Server(s)	Threats of Privacy Leakage
[7]	Two	$\mathbf{X}' = \mathbf{P}\mathbf{X}\mathbf{Q}$	<u>SVD</u> on \mathbf{X}'	Raw Data
[9]	Four	$\mathbf{X}\mathbf{X}^T, \mathbf{X}^T\mathbf{X}$	<u>SVD</u> on $\mathbf{X}\mathbf{X}^T, \mathbf{X}^T\mathbf{X}$	Raw Data
[10]	Three	Secret Sharing $\mathbf{X} = \mathbf{X}_a + \mathbf{X}_b$	<u>SVD</u> on $\mathbf{X}_a + \mathbf{X}_b$	Raw Data
[25]	One	Project \mathbf{X} $\mathbf{H} = \mathbf{X}\mathbf{G}$	<u>Orthogonal H</u> (<i>e.g.</i> , Gram-Schmidt)	Projections of Raw Data
[34]	One	Project $\mathbf{X}^T\mathbf{X}$ $\mathbf{Y} = \mathbf{X}^T\mathbf{X}\mathbf{Z}$	<u>Orthogonal Y</u> (<i>e.g.</i> , Gram-Schmidt)	Projections of Raw Data

The servers obtain excessive access to the private data and thus **significantly decreases the privacy protection**.

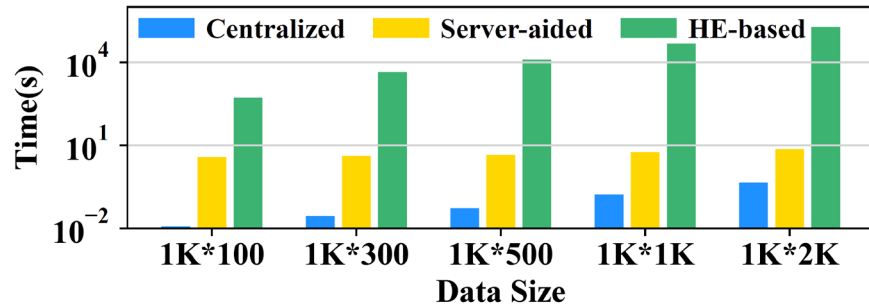
Intuitive ideas of enhancing privacy protection at the server side cannot work

- Pick a subset of the users as “servers”. The privacy issues remain in unselected users.
- Deploy TEE at the servers. The issue of distrust, particularly in the server-aided approach, poses a significant challenge.
- Leveraging HE at the servers. HE brings severe computational overhead (will discuss more).

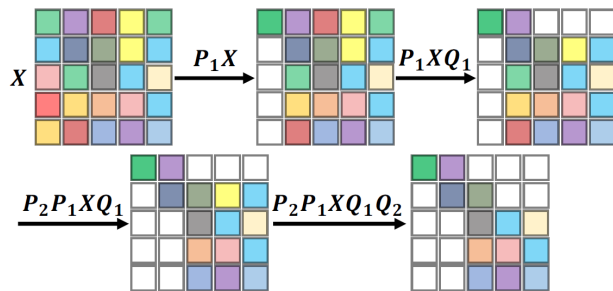
Efficient Decentralization is Challenging

Existing works have explored using **Homomorphic Encryption** to remove the servers but suffer from **significant efficiency issues**.

Computational Challenge



HE-based solution is 4~5 orders of magnitudes slower.



Sequential computation makes stacking more hardware less effective. Designing new hardware is also challenging.

Communication Challenge

Data Size	Communication Size	
	Before Encryption	After Encryption
$m = 1K$ $n = 50M$	372GB	11.6TB

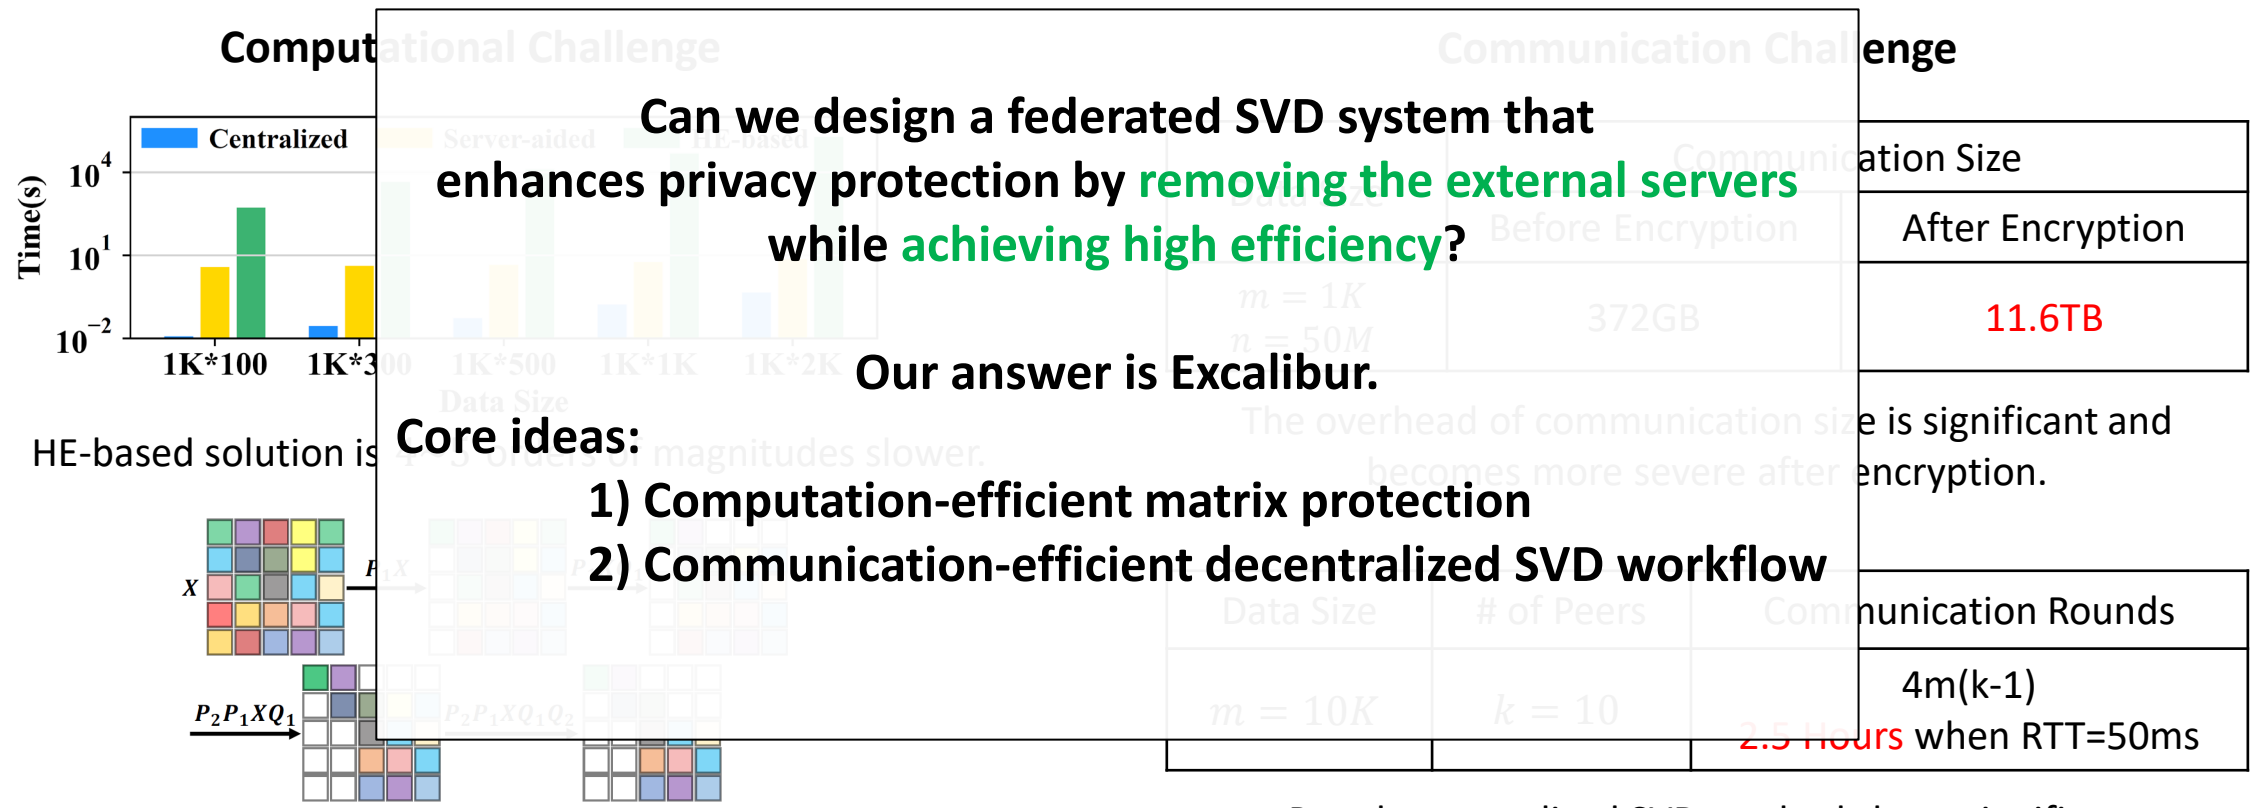
The overhead of communication size is significant and becomes more severe after encryption.

Data Size	# of Peers	Communication Rounds
$m = 10K$	$k = 10$	$4m(k-1)$ 2.5 Hours when RTT=50ms

Popular centralized SVD methods have significant overhead of communication rounds.

Efficient Decentralization is Challenging

Existing works have explored using **Homomorphic Encryption** to remove the servers but suffer from **significant efficiency issues**.



Sequential computation makes stacking more hardware less effective. Designing new hardware is also challenging.

Popular centralized SVD methods have significant overhead of communication rounds.

Outline

- Introduction
- **Excalibur's Matrix Protection (computation-efficient protection)**
 - **Threat Model and Security Goals**
 - **Multiplicative Matrix Sharing**
 - **Accelerating the Multiplicative Operations**
- Excalibur's Decentralized SVD Workflow (communication-efficient workflow)
- Implementation and Evaluation
- Conclusion and Future Work

Threat Model and Security Goals

Threat Model: We assume all peers are semi-honest.



- ✓ Strictly follow pre-defined protocol →
- ✗ But try to discover privacy during execution

Discussion

(1) What if they do not follow the protocol?

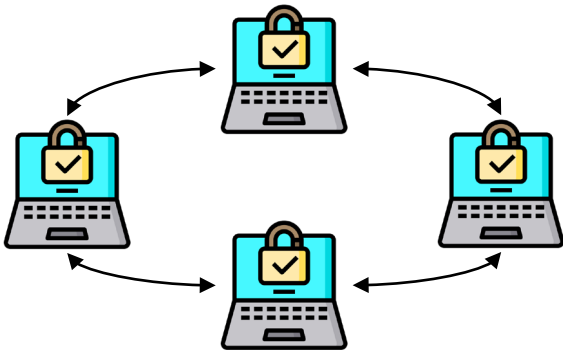
Check whether $UU^T = I, V_iV_i^T = I, X_i = U\Sigma V_i^T$

Security Definition: The system is secure if all intermediate results could be derived from final results.
(Having the same distribution in mathematical language)



(2) How to protect the final results?

Leveraging differential privacy.



Same definition to secure multi-party computation (SMC)

Multiplicative Matrix Sharing (MMS)

We protect the matrix with random non-singular matrices

$$X' = AXB = A[X_1, \dots, X_k] \begin{bmatrix} B_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & B_k \end{bmatrix}$$

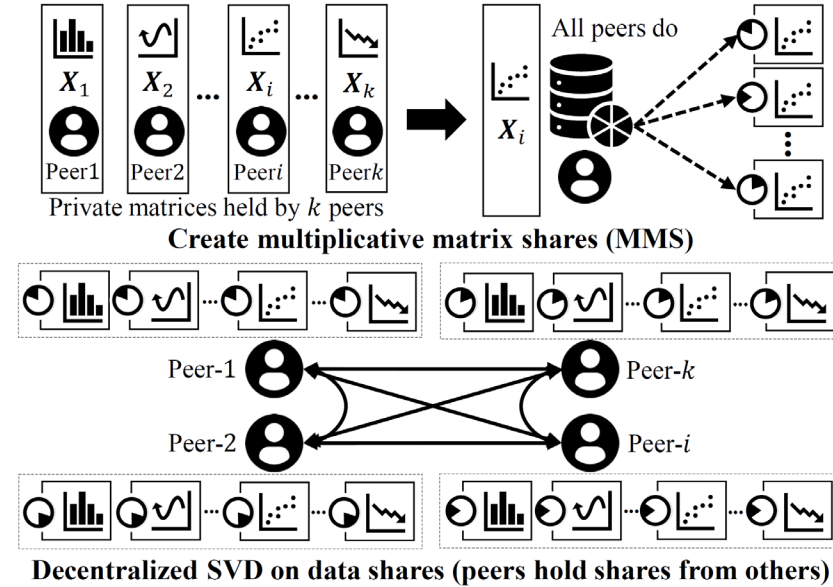
For peer- i

$$X'_i = \begin{bmatrix} A_1 \\ \vdots \\ A_k \end{bmatrix} X_i B_i = \begin{bmatrix} A_1 X_i B_i \\ \vdots \\ A_k X_i B_i \end{bmatrix} = \begin{bmatrix} S_i^1 \\ \vdots \\ S_i^k \end{bmatrix} \text{ One MMS}$$

Recovering X_i needs all the matrix shares

$$X_i = A^{-1} X'_i B_i^{-1} = [A_1^{-1}, \dots, A_k^{-1}] \begin{bmatrix} S_i^1 \\ \vdots \\ S_i^k \end{bmatrix} B_i^{-1} = \sum_{j=1}^k A_j^{-1} S_i^j B_i^{-1}$$

(A_j^{-1} is the columns of A^{-1})



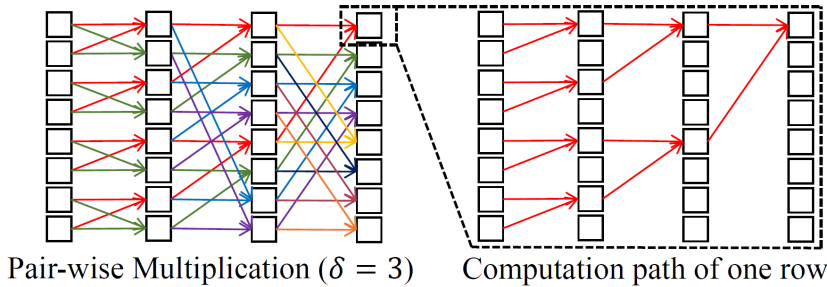
Theorem 1. Denote \mathbb{O}_n as the compact group of $n \times n$ orthogonal matrices under Haar measure, if we choose dense matrix $A \in \mathbb{O}_m$ and uniformly generate $B_i \in \mathbb{O}_{n_i}$, Excalibur produces federated SVD results with no accuracy loss and can satisfy the security defined in Definition 1 while the adversary can compromise up to $k - 1$ peers ($|C| = k - 1$).

Accelerating the Multiplicative Operations

Complexity of generating and applying random orthogonal matrices is $O(m^2n)$ or $O(n^2m)$.

How to efficiently support large-scale data?

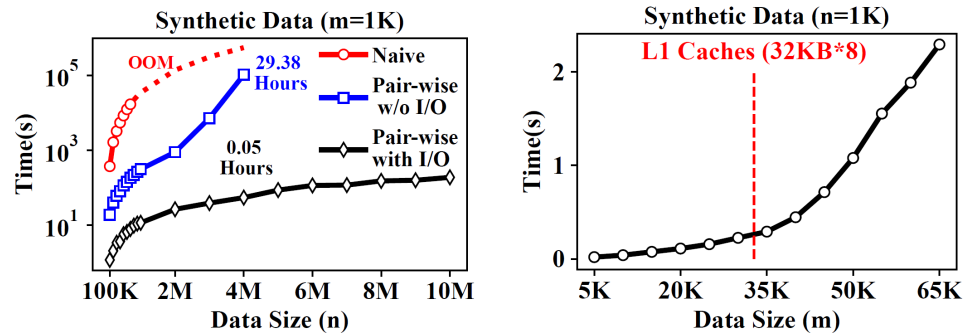
(1) Reduce algorithm complexity



$$O(m^2n_i) \rightarrow O(mn_i \log m)$$

Transfer A to a group of 2×2 rotations on random selected rows.

(2) Solve the I/O bottleneck



Process by columns instead of by rows.
Formulate each column into rectangular matrix if it exceed L1 cache.

(3) Local pre-processing for matrix B

$$X_i = R_i^T Q_i^T$$

$$[R_1^T, \dots, R_k^T] = U \Sigma [V_{R_1}^T, \dots, V_{R_k}^T]$$

$$O(n_i^2 m) \rightarrow O(m^2 n_i + m^3)$$

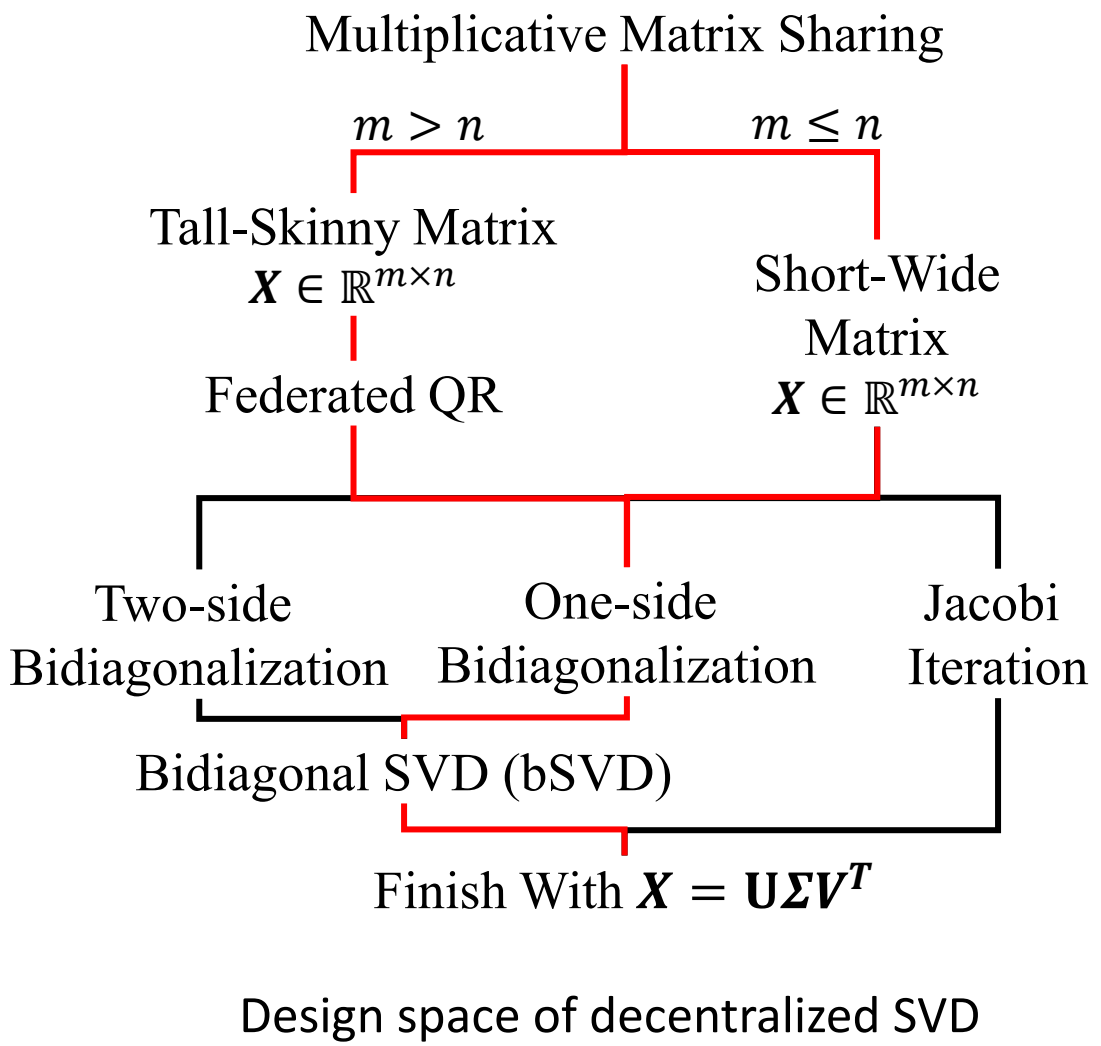
Local pre-process to reduce complexity when $m \ll n$.

With all the above optimizations, the MMS can efficiently support billion-scale data.

Outline

- Introduction
- Excalibur's Matrix Protection (computation-efficient protection)
- **Excalibur's Decentralized SVD Workflow (communication-efficient workflow)**
 - **Analyzing the Design Space**
 - **Overlapping the Pipelines**
- Implementation and Evaluation
- Conclusion and Future Work

Design Space of Decentralized SVD



Decentralized SVD Path	Short-Wide Matrix $X \in \mathbb{R}^{m \times n} \ m \leq n; \ k \text{ peers}$	
	CommAmount (Bandwidth)	CommTimes (Latency)
Jacobi iteration (1 iteration)	$3 \frac{k-1}{k} (m^2 - m)$ $\Rightarrow O(m^2)$	$m(m-1)(k-1)$ $\Rightarrow O(km^2)$
Two-side Bidiagonal + bSVD	$(k-1)(\frac{3}{2}m^2 + \frac{7}{2}m)$ $\Rightarrow O(km^2)$	$4m(k-1)$ $\Rightarrow O(km)$
One-side Bidiagonal + bSVD	$\frac{k-1}{k} (m^2 - m)$ $\Rightarrow O(m^2)$	$(6m-6)(k-1)$ $\Rightarrow O(km)$

- Jacobi iteration has significantly higher communication rounds.
- Two-side bidiagonalization is the popular method used in NumPy and LAPACK, but its communication size is large.
- The one-side bidiagonalization has the minimum communication complexity.

Overlapping the Pipelines

```

Algorithm 1: Excalibur's decentralized SVD workflow. (The three looped ring all-reduce are highlighted and we will reduce them to only one through overlapping the pipelines (§5.2).)

Input: Matrix  $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k]$  held by  $k$  peers, where  $\mathbf{X} \in \mathbb{R}^{m \times n}$ ,  $m \leq n$ ,  $\mathbf{X}_i \in \mathbb{R}^{m \times n_i}$ , and  $\sum_{i=1}^k n_i = n$ .
Output:  $\mathbf{U}, \Sigma, [\mathbf{V}_1^T, \mathbf{V}_2^T, \dots, \mathbf{V}_k^T]$  (i.e., SVD of  $\mathbf{X}$ )

1 Function DecSVD ( $\mathbf{X}$ ):
  // All peers run this function in parallel
  2  $\mathbf{U} \leftarrow I, c \leftarrow \text{MyPeerID}$   $\triangleright$  e.g.,  $c = 1$  for peer-1
  3 for  $i = 1, 2, \dots, m-2$  do
    4  $\mathbf{h} \leftarrow \text{RingAllReduce}(\mathbf{X}_{\mathbf{c}[i]} * \mathbf{X}_{\mathbf{c}[i+1]}^T)$ 
    // Apply reflector to  $\mathbf{X}$  and  $\mathbf{U}$ 
    5  $\mathbf{X}_{\mathbf{c}[i+1]} \leftarrow \text{house}(\mathbf{h}) \otimes \mathbf{X}_{\mathbf{c}[i+1]}$ 
    6  $\mathbf{U}[i+1] \leftarrow \text{house}(\mathbf{h}) \otimes \mathbf{U}[i+1]$ 
  7 end
  //  $\alpha$  contains the diagonal elements.
  //  $\beta$  contains the subdiagonal elements.
  8  $\alpha \leftarrow \{0\}^m, \beta \leftarrow \{0\}^{m-1}$ 
  9  $\alpha[1] \leftarrow \sqrt{\text{RingAllReduce}(\|\mathbf{X}_{\mathbf{c}[1]}\|_2^2)}$ 
  10  $\mathbf{V}_{\mathbf{c}[1]}^T[1] = \mathbf{X}_{\mathbf{c}[1]}/\alpha[1]$ 
  11 for  $i = 2, 3, \dots, m$  do
    12  $\beta[i-1] = \text{RingAllReduce}(\mathbf{X}_{\mathbf{c}[i]} * \mathbf{V}_{\mathbf{c}[i-1]})$ 
    13  $\mathbf{X}_{\mathbf{c}[i]} \leftarrow \mathbf{X}_{\mathbf{c}[i]} - \beta[i-1] * \mathbf{V}_{\mathbf{c}[i-1]}^T$ 
    14  $\alpha[i] \leftarrow \sqrt{\text{RingAllReduce}(\|\mathbf{X}_{\mathbf{c}[i]}\|_2^2)}$ 
    15  $\mathbf{V}_{\mathbf{c}[i]}^T[1] = \mathbf{X}_{\mathbf{c}[i]}/\alpha[i]$ 
  16 end
  17  $\mathbf{U}_{\mathbf{b}}, \Sigma, \mathbf{V}_{\mathbf{b}}^T \leftarrow \text{bSVD}(\alpha, \beta)$ 
  // Combine results together
  18  $\mathbf{U} \leftarrow \mathbf{U} * \mathbf{U}_{\mathbf{b}}$ 
  19  $\mathbf{V}_{\mathbf{c}}^T \leftarrow \mathbf{V}_{\mathbf{b}}^T * \mathbf{V}_{\mathbf{c}}^T$ 
  20 return  $\mathbf{U}, \Sigma, [\mathbf{V}_1^T, \mathbf{V}_2^T, \dots, \mathbf{V}_k^T]$ 
21 End Function

```

The three **all-reduce communications** in the for loops **become the bottleneck**.

$$\alpha_i^2 = \sum_{j=1}^k \|\mathbf{X}_j^i\|_2^2 = \sum_{j=1}^k \|\mathbf{X}_j^i - \beta_{i-1} (\mathbf{V}_j^{i-1})^T\|_2^2$$

Overlapping the communications is challenging due to data dependency.

$$\alpha_i = \sqrt{\theta_1^i - 2(\theta_2^i)^2 + (\theta_2^i)^2 \theta_3^i}, \quad \beta_{i-1} = \theta_2^i$$

We find the underlying shared components of these communications.

3	for $i = 1, 2, \dots, m-2$ do
10	$\mathbf{v}_{\mathbf{c}[1]} = \mathbf{X}_{\mathbf{c}[1]}/\alpha[1]$
11	for $i = 2, 3, \dots, m$ do

Merge for loops via pipeline parallel.

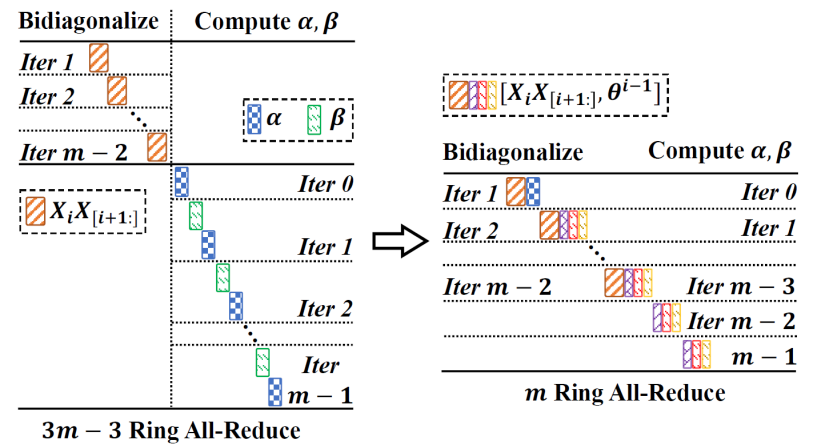


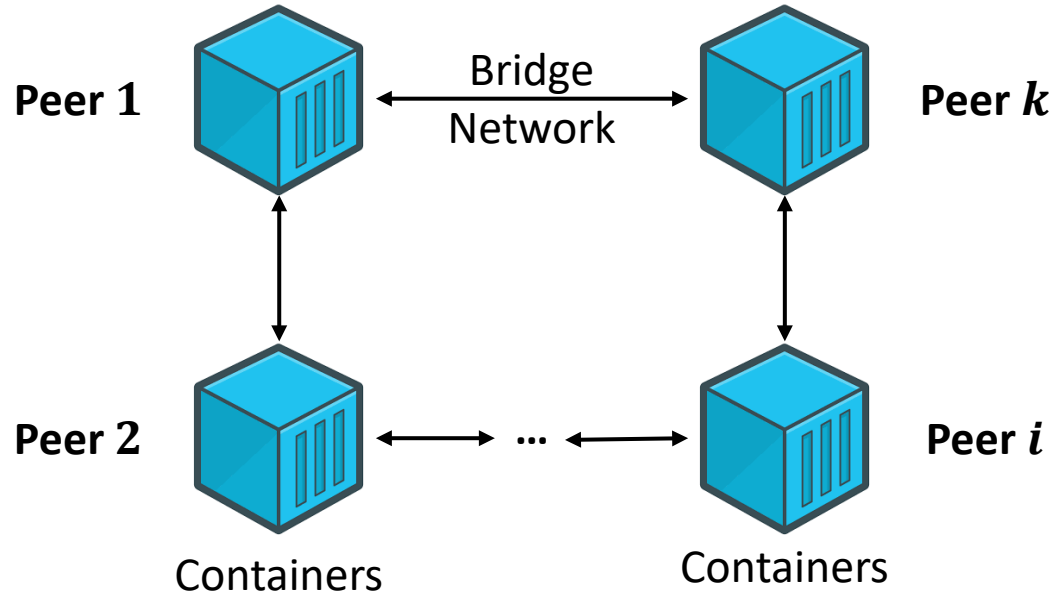
Figure 7: Overlapping the pipeline reduces the number of ring all-reduce communication from $3m - 3$ to m , i.e., approximately reduced by 66%.

Reducing 66% communication cost.

Outline

- Introduction
- Excalibur's Matrix Protection
- Excalibur's Decentralized SVD Workflow
- **Implementation and Evaluation**
- Conclusion and Future Work

Implementation and Testbed



Each container is assigned with 4 Cores and 64GB RAM
 Default network setting: 1Gbps bandwidth and 50ms RTT

We implement a fully functional prototype using C/C++. The system runs in double precision, i.e., 64 bits.

- We use BLAS and LAPACKE from Intel MKL as the major library.
- For operations not included in existing library, we implement from scratch and use OMP and AVX2 for parallelism.
- To support large-scale data that cannot fit into memory, we create memory-mapped files and offload the data to NVMe SSD.

Evaluation

Datasets

We have used four datasets in the evaluation: MNIST, Wine, ML100K, and synthetic data.

Baselines

- 1) FedSVD: state-of-the-art (SOTA) server-aided federated SVD.
- 2) SF-PCA: SOTA multi-key HE-based solution.
- 3) FATE and SecureML: widely used federated linear regression (LR) systems.

Tasks

SVD task and its three applications:

Principal components analysis (PCA), latent semantic analysis (LSA), and LR.

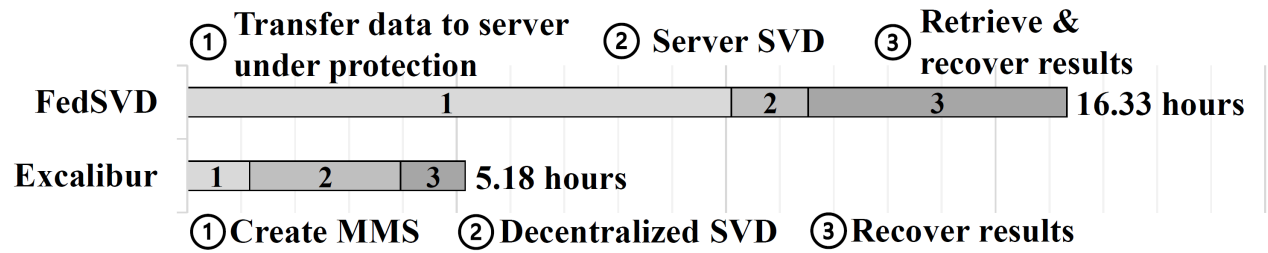
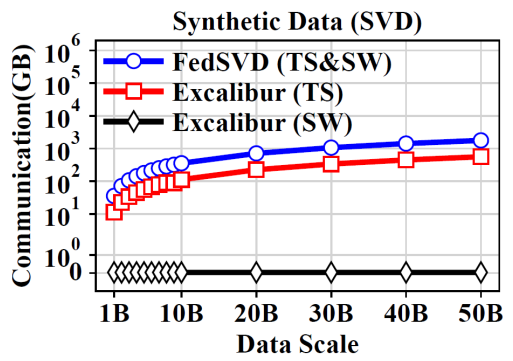
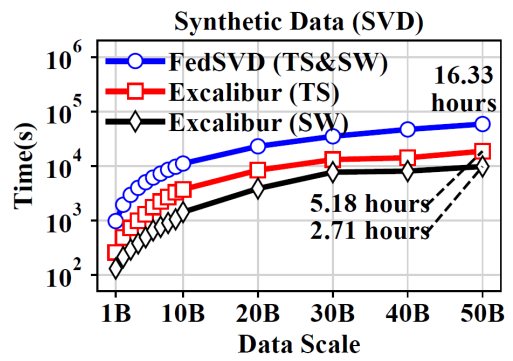
Accuracy Evaluation

Method	Wine	MNIST	ML100K	Synthetic
FedSVD	1.44×10^{-13}	2.66×10^{-10}	2.56×10^{-12}	5.67×10^{-15}
Excalibur	3.56×10^{-14}	2.15×10^{-13}	3.76×10^{-15}	2.96×10^{-17}

Table 3: Reconstruction error of SVD on three real-world datasets and synthetic data.

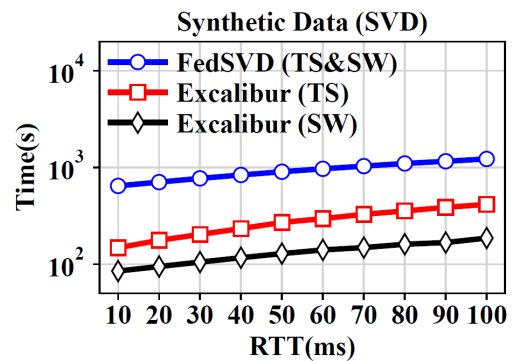
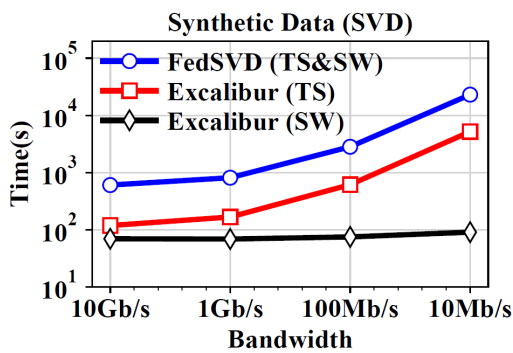
Evaluation

Efficiency on SVD Task



(a) End-to-end time comparison on SVD task under 1~50 billion data.

(b) Communication cost on SVD task under 1~50 billion data.



(c) Impact of network bandwidth on SVD efficiency using 1 billion data.

(d) Impact of network latency on SVD efficiency using 1 billion data.

Compared to the SOTA server-aided system, Excalibur not only removes the external servers but also achieves better efficiency.

- Excalibur is $3.1 \times \sim 6.0 \times$ faster than FedSVD.
- Excalibur reduces more than 68.4% amount of communication.

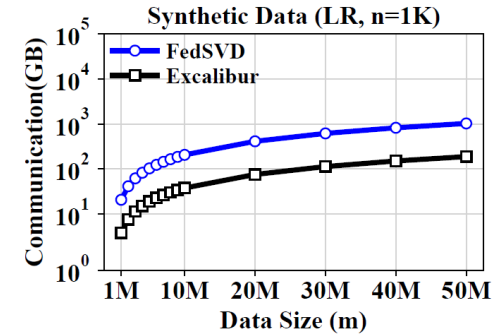
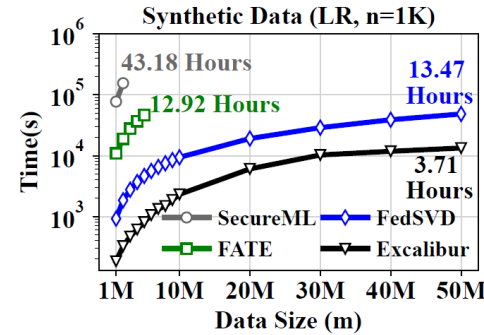
Evaluation

Efficiency on SVD Applications

Solution	Data	CPU Cores	Time	Single Core Throughput
SF-PCA[17]	45.6M (760×60K)	72	2.22 Hours	0.28 M/H
Excalibur	10000M (1K×10M)	24	0.063 Hours	6595.55 M/H

Comparing Excalibur with SF-PCA on PCA application, while computing the top-5 principal components (bandwidth=1Gb/s, RTT=20ms, six peers).

- Compared to the SOTA HE-based system that attempted to remove the servers, Excalibur is far more efficient and has $> 23000 \times$ larger throughput.
- Comparing to two widely used federated LR systems: FATE and SecureML, Excalibur is 100x and 1000x faster, respectively.

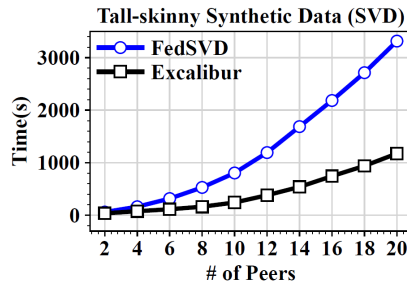


(a) Efficiency of Excalibur, FedSVD, FATE, and SecureML on LR. (b) Communication size of Excalibur and FedSVD on LR.

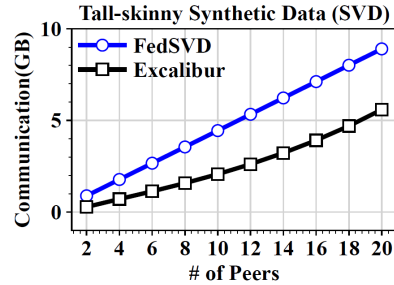
Comparing Excalibur with FedSVD, FATE, and SecureML on LR application.

Evaluation

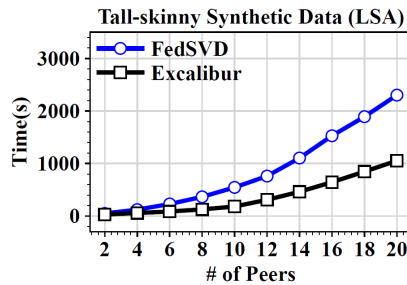
Scalability on SVD and LSA



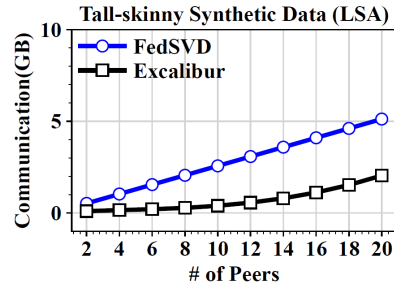
(a) End-to-end time consumption when increasing the # of peers.



(b) Comm cost (each peer) when increasing the # of peers.



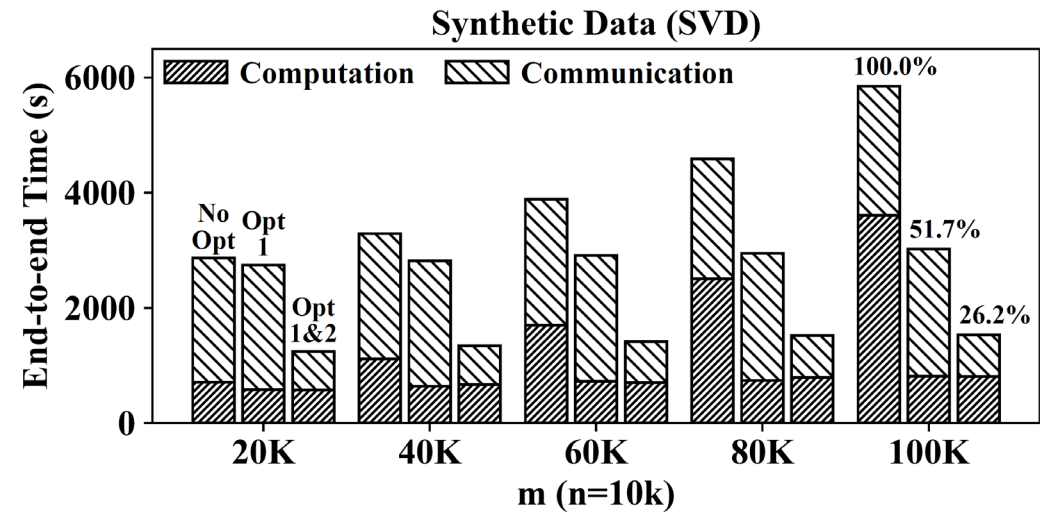
(c) End-to-end time consumption when increasing the # of peers.



(d) Comm cost (each peer) when increasing the # of peers.

We evaluate Excalibur's scalability when increasing the number of peers, assuming all peers hold the same amount of data and test the efficiency when more peers join the federation.

Effectiveness of the Optimizations

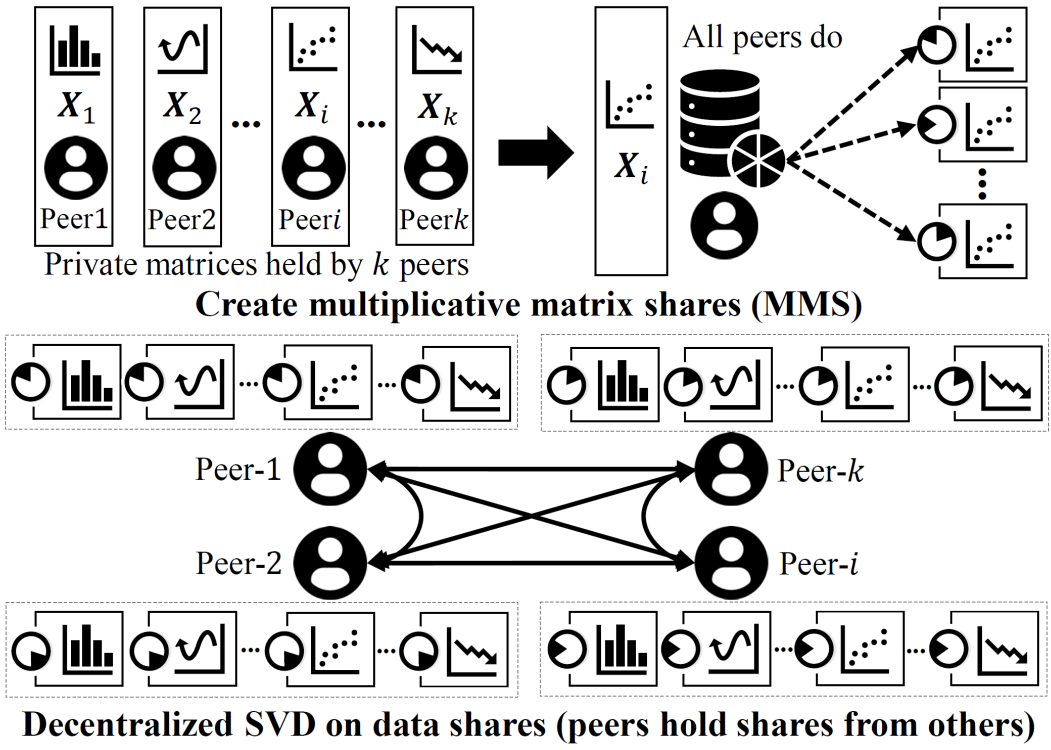


Measuring the effectiveness of system optimizations in Excalibur, while NoOpt means no optimization, Opt1 is optimizing the multiplicative operations in MMS, and Opt2 is overlapping pipelines to reduce communication rounds.

Outline

- Introduction
- Excalibur's Matrix Protection
- Excalibur's Decentralized SVD Workflow
- Implementation and Evaluation
- **Conclusion and Future Work**

Conclusion and Future Work



Conclusion

In this paper, we propose Excalibur, an efficient decentralized federated SVD that **not only eliminates the privacy concerns** caused by external servers **but also can efficiently decompose large-scale matrices.**

Future Work

- How to update the results when more peers are joining in the computation?
- How to utilize the matrix protection in this paper in other scenarios? For example, the secure model inference.



Thanks



Artifact available at: <https://github.com/Di-Chai/Excalibur>