Advanced Tools from Modern Cryptography

Lecture 2

First Tool: Secret-Sharing

Secret-Sharing

- Dealer encodes a message into n shares for n parties
  - Privileged subsets of parties should be able to reconstruct the secret

    Access Structure: Set of all privileged sets
  - View of an unprivileged subset should be independent of the secret
- Very useful
  - Direct applications (distributed storage of data or keys)
  - Important component in other cryptographic constructions
    - Secure multi-party computation
    - Attribute-Based Encryption
    - Leakage resilience ...

Threshold Secret-Sharing

- (n,t)-secret-sharing
  - Divide a message m into n shares $s_1,\ldots,s_n$, such that
    - any t shares are enough to reconstruct the secret
    - upto t-1 shares should have no information about the secret
- Recall last time: (2,2) secret-sharing

e.g., $(s_1,\ldots,s_{t-1})$ has the same distribution for every m in the message space

Threshold Secret-Sharing

Additive Secret-Sharing

- Construction: (n,n) secret-sharing
  - Message-space = share-space = G, a finite group
    - e.g. G = $\mathbb{Z}_2$ (group of bits, with xor as the group operation)
    - or, G = $\mathbb{Z}_2^d$ (group of d-bit strings)
    - or, G = $\mathbb{Z}_p$ (group of integers mod p)
  - Share(M):
    - Pick $s_1,\ldots,s_{n-1}$ uniformly at random from G
    - Let $s_n = -(s_1 + \ldots + s_{n-1}) + M$
  - Reconstruct($s_1,\ldots,s_n$): $M = s_1 + \ldots + s_n$
  - Claim: This is an (n,n) secret-sharing scheme [Why?]

PROOF

Additive Secret-Sharing: Proof

- Share(M):
  - Pick $s_1,\ldots,s_{n-1}$ uniformly at random from G
  - Let $s_n = M - (s_1 + \ldots + s_{n-1})$
- Reconstruct($s_1,\ldots,s_n$): $M = s_1 + \ldots + s_n$
- Claim: Upto n-1 shares give no information about M
- Proof: Let $T \subseteq \{1,\ldots,n\}$, $|T| = n-1$. We shall show that $\{ s_i \}_{i\in T}$ is distributed the same way (in fact, uniformly) irrespective of what M is.
  - For $T = \{1,\ldots,n-1\}$, true by construction. How about other T?
  - For concreteness consider $T = \{2,\ldots,n\}$. Fix any (n-1)-tuple of elements in G, $(g_1,\ldots,g_{n-1}) \in G^{n-1}$. To prove $\Pr[\ (s_2,\ldots,s_n)=(g_1,\ldots,g_{n-1})\ ]$ is same for all M.
  - Fix any M.
  - $(s_2,\ldots,s_n) = (g_1,\ldots,g_{n-1}) \Leftrightarrow (s_2,\ldots,s_{n-1}) = (g_1,\ldots,g_{n-2})$ and $s_1 = M-(g_1+\ldots+g_{n-1})$.
  - So $\Pr[\ (s_2,\ldots,s_n)=(g_1,\ldots,g_{n-1})\ ] = \Pr[\ (s_1,\ldots,s_{n-1})=(a,g_1,\ldots,g_{n-2})\ ]$, $a:=(M-(g_1+\ldots+g_{n-1})$
  - But $\Pr[(s_1,\ldots,s_{n-1})=(a,g_1,\ldots,g_{n-2})] = 1/|G|^{n-1}$, since $(s_1,\ldots,s_{n-1})$ uniform over G^{n-1}
  - Hence $\Pr[\ (s_2,\ldots,s_n)=(g_1,\ldots,g_{n-1})\ ] = 1/|G|^{n-1}$, irrespective of M. □

An Application

- Gives a "private summation" protocol (for commutative groups)

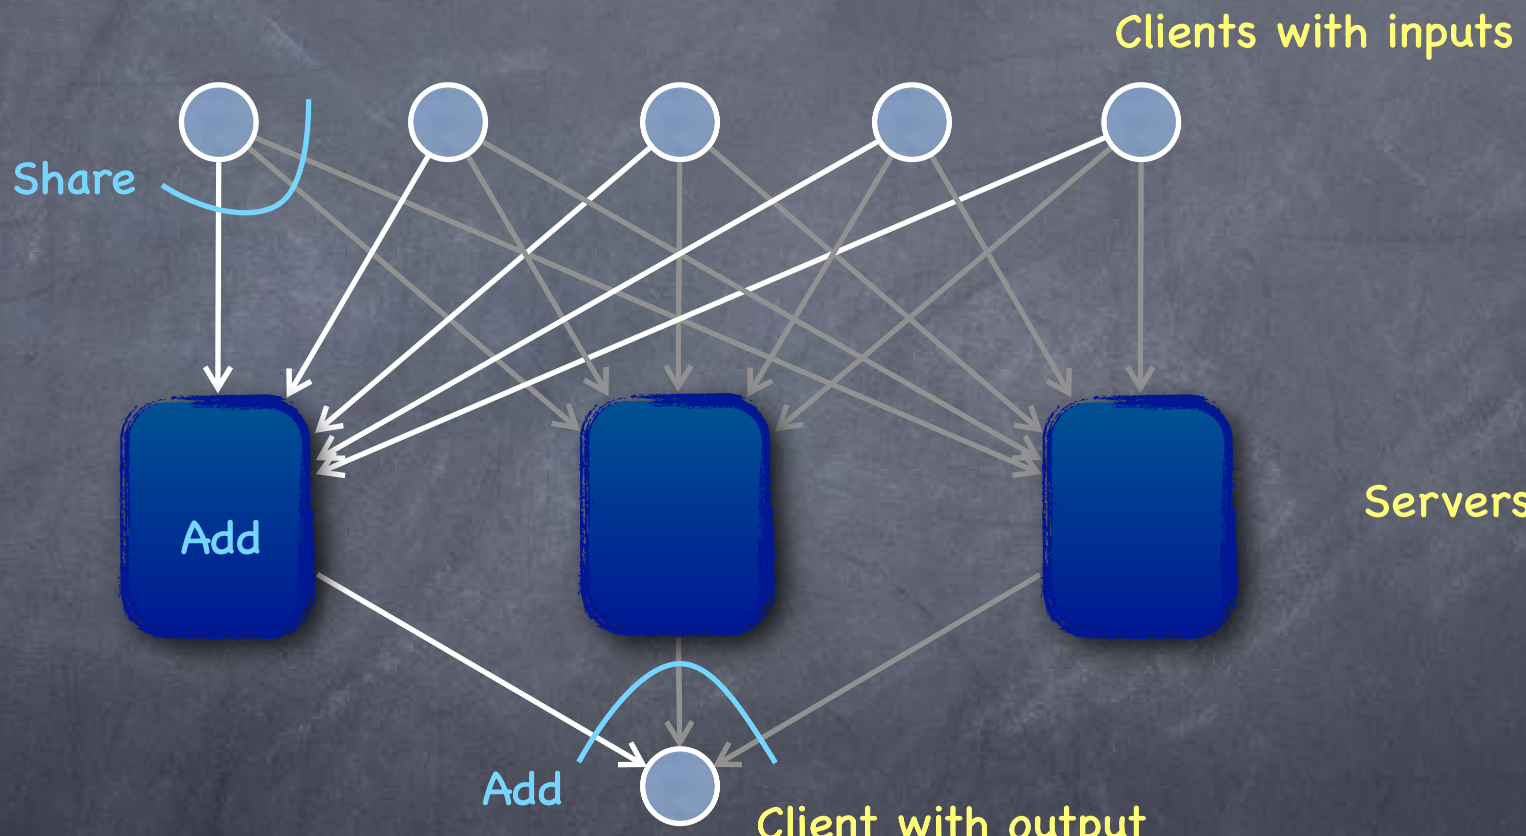

- "Secure against passive corruption" (i.e., no colluding set of servers/clients learn more than what they must) if at least one server stays out of the collusion

Linear Secret-Sharing

- Another look at additive secret-sharing

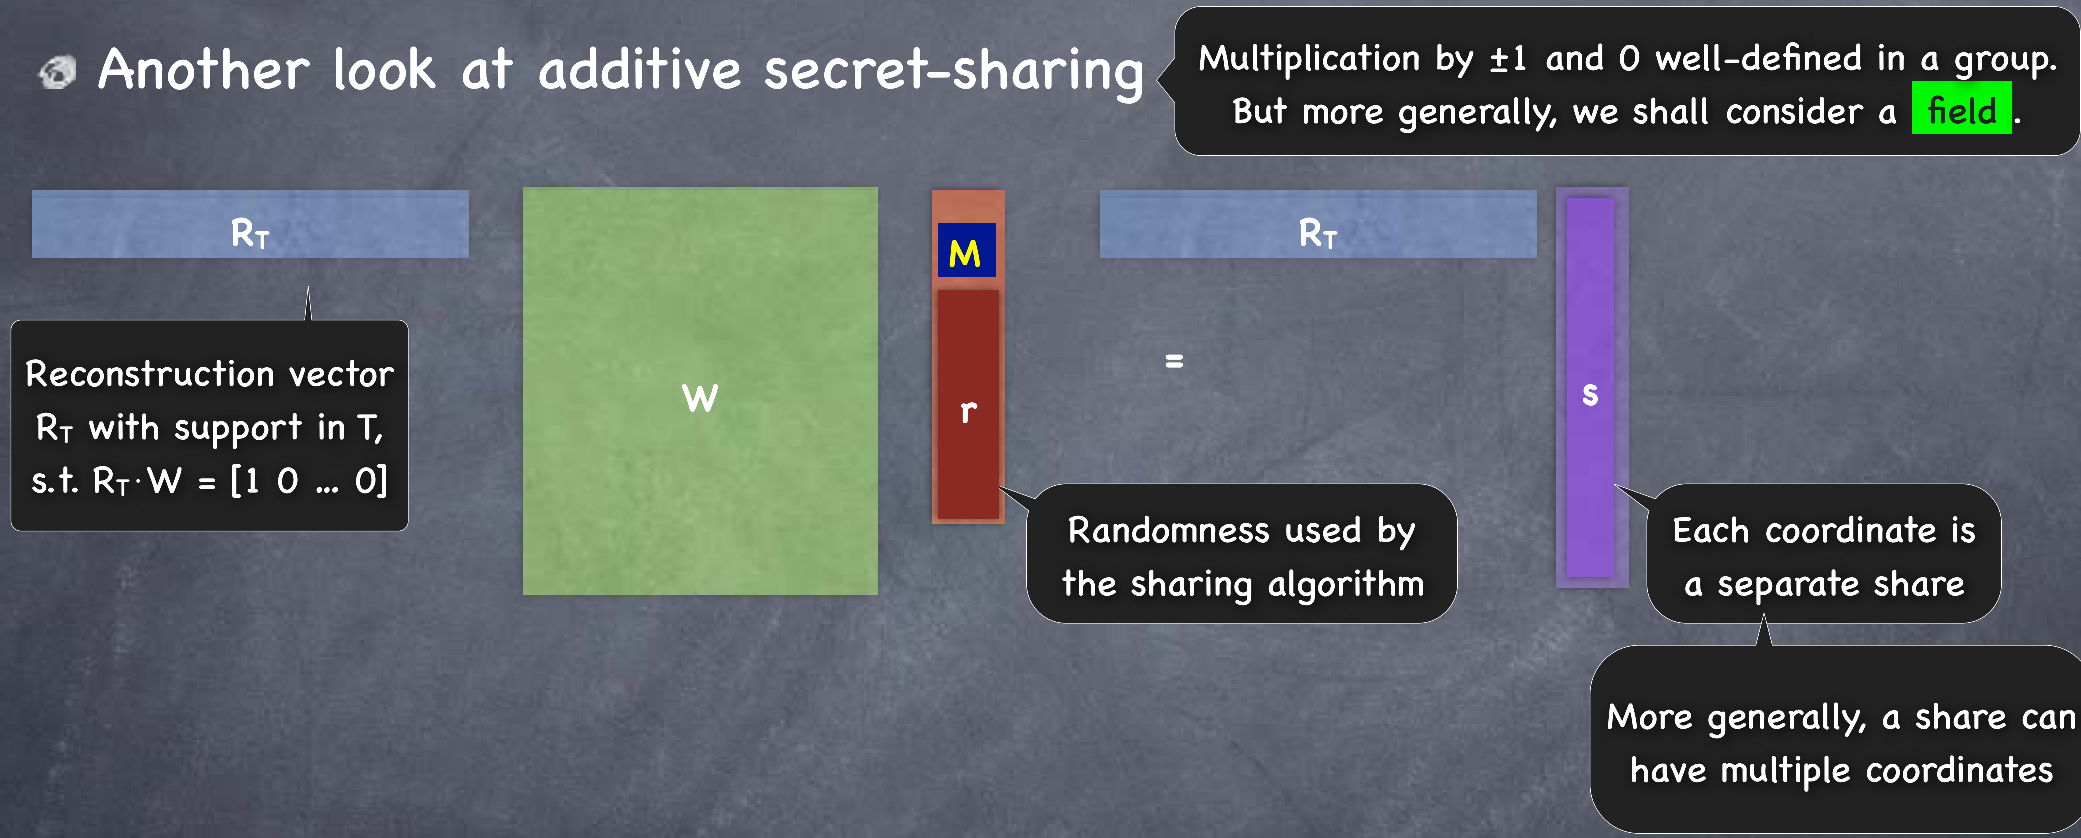

- Linear Secret-Sharing over a field: message and shares are field elements
- Reconstruction by a set $T \subseteq [n]$: solve the message from given shares
  - i.e., solve $W_T \begin{bmatrix} M \\ r \end{bmatrix} = s_T$ for M

Security of Linear Secret-Sharing

- Claim: Every such linear scheme is a secure secret-sharing scheme for some access structure
- Suppose $T \subseteq [n]$ s.t. M not uniquely reconstructible from $\underline{s}_T$
  - i.e., solution space (of $\underline{z}$) for $W_T \cdot \underline{\mathbf{z}} = \underline{s}_T$ contains at least two points with distinct values α and β for M
  - Then, $\forall \gamma \in F$, the solution space has a point with $M=\gamma$ (e.g., linear combination of the above points with factors $(\gamma-\beta)/(\alpha-\beta)$ and $(\alpha-\gamma)/(\alpha-\beta)$)
  - Therefore, for any $\gamma \in F$, can add equation $M=\gamma$ and get a solution space of dimension k equal to the nullity of the system
    - i.e., with $M=\gamma$, exactly $|F|^k$ choices of randomness $\underline{r}$ that give $\underline{s}_T$
  - i.e., for all $\underline{s}_T$ and γ, $\Pr[\text{view}=\underline{s}_T \mid M=\gamma] = |F|^k/|F|^{t-1}$

Threshold Secret-Sharing

- Construction: (n,2) secret-sharing
- Message-space = share-space = F, a finite field (e.g. integers mod prime)
  - Share(M): pick random r. Let $s_i = r \cdot a_i + M$ (for $i=1,\ldots,n < |F|$)
  - Reconstruct(s_i, s_j): $r = (s_i-s_j)/(a_i-a_j)$; $M = s_i - r \cdot a_i$
  - Each s_i by itself is uniformly distributed, irrespective of M [Why?]
  - "Geometric" interpretation
    - Sharing picks a random "line" $y = f(x)$, such that $f(0)=M$. Shares $s_i = f(a_i)$.
    - s_i is independent of M: exactly one line passing through (a_i,s_i) and $(0,M')$ for any secret M'
    - But can reconstruct the line from two points!

a_i are n distinct, non-zero field elements

Since a_i^{-1} exists, exactly one solution for $r \cdot a_i+M=d$, for every value of d

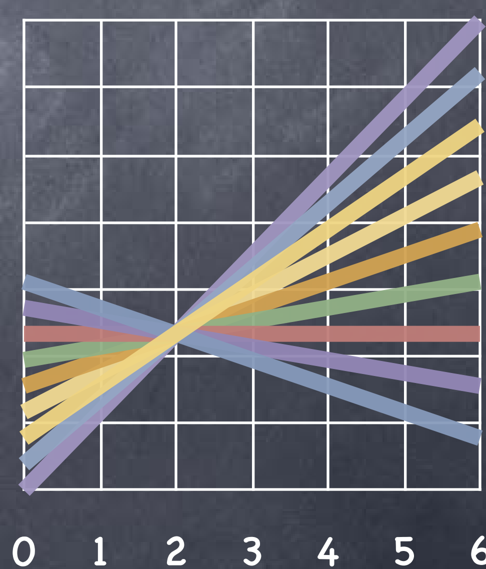

Threshold Secret-Sharing

Shamir Secret-Sharing

- (n,t) secret-sharing in a (large enough) field F
- Generalizing the geometric/algebraic view: instead of lines, use polynomials
  - Share(m): Pick a random degree t-1 polynomial f(X), such that f(0)=M. Shares are $s_i = f(a_i)$.
    - Random polynomial with f(0)=M: $c_0 + c_1X + c_2X^2 + \ldots + c_{t-1}X^{t-1}$ by picking c_0=M and $c_1,\ldots,c_{t-1}$ at random.
  - Reconstruct($s_1,\ldots,s_t$): Lagrange interpolation to find M=c_0
    - Given t points can reconstruct the polynomial. Given < t points, for any M', there are polynomials which pass through (0,M')
- Secrecy: Shamir's scheme is linear!

Linearity of Shamir Secret-Sharing

- Shamir's scheme is a linear secret-sharing scheme

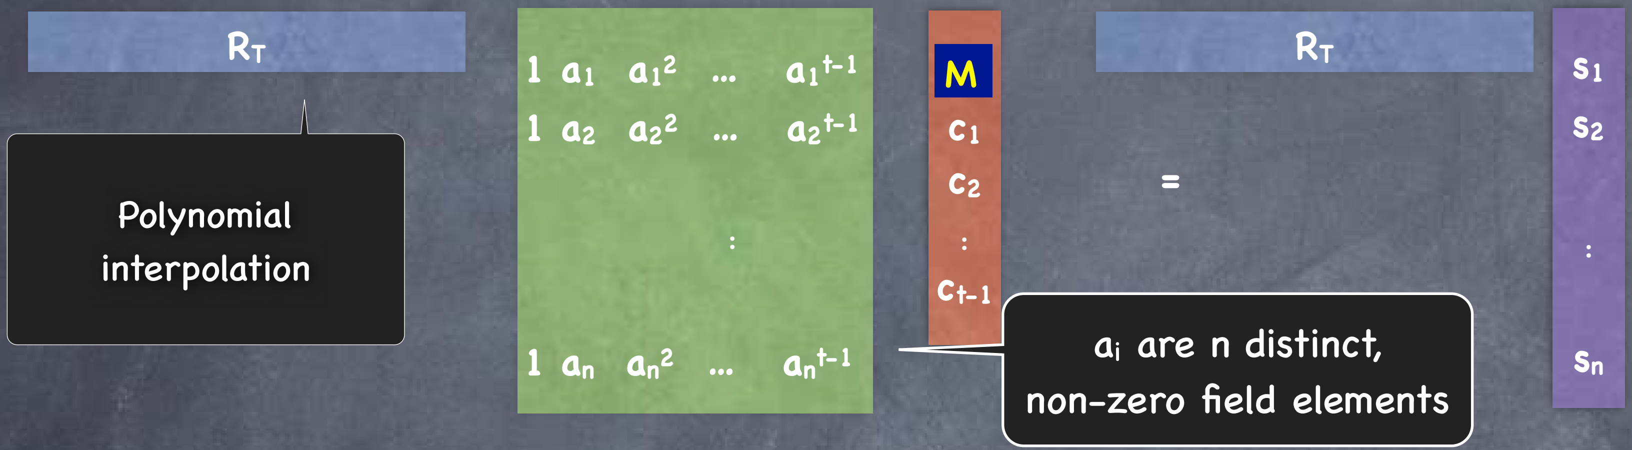

- Which sets $T \subseteq [n]$ can reconstruct? i.e., T s.t. W_T spans [1 0 ... 0]?
  - W_T spans [1 0 ... 0] iff $|T| \geq t$
    - For $|T|=t$, W_T is a Vandermonde matrix, and is a basis for $\mathbb{F}^t$
    - For $|T| < t$, can add a row [1 0 ... 0] and (optionally) more rows of the form $[1\ a\ a^2 \ldots a^t]$ to get a Vandermonde matrix. So [1 0 ... 0] is independent of the rows of W_T
- Secrecy: guaranteed for any linear secret-sharing scheme

More General Access Structures

- Idea: For arbitrary monotonic access structure $\mathcal{A}$, there is a "basis" $\mathcal{B}$ of minimal sets in $\mathcal{A}$. For each S in $\mathcal{B}$ generate an (|S|,|S|) sharing, and distribute them to the members of S.
  - Works, but very "inefficient"
    - How big is $\mathcal{B}$? (Say when $\mathcal{A}$ is a threshold access structure)

      $|\mathcal{B}| = \binom{n}{t}$
    - Total share complexity = $\Sigma_{S \in \mathcal{B}}$ |S| field elements. (Compare with Shamir's scheme: n field elements in all.)

      $t \cdot \binom{n}{t}$
  - More efficient schemes known for large classes of access structures

More General Access Structures

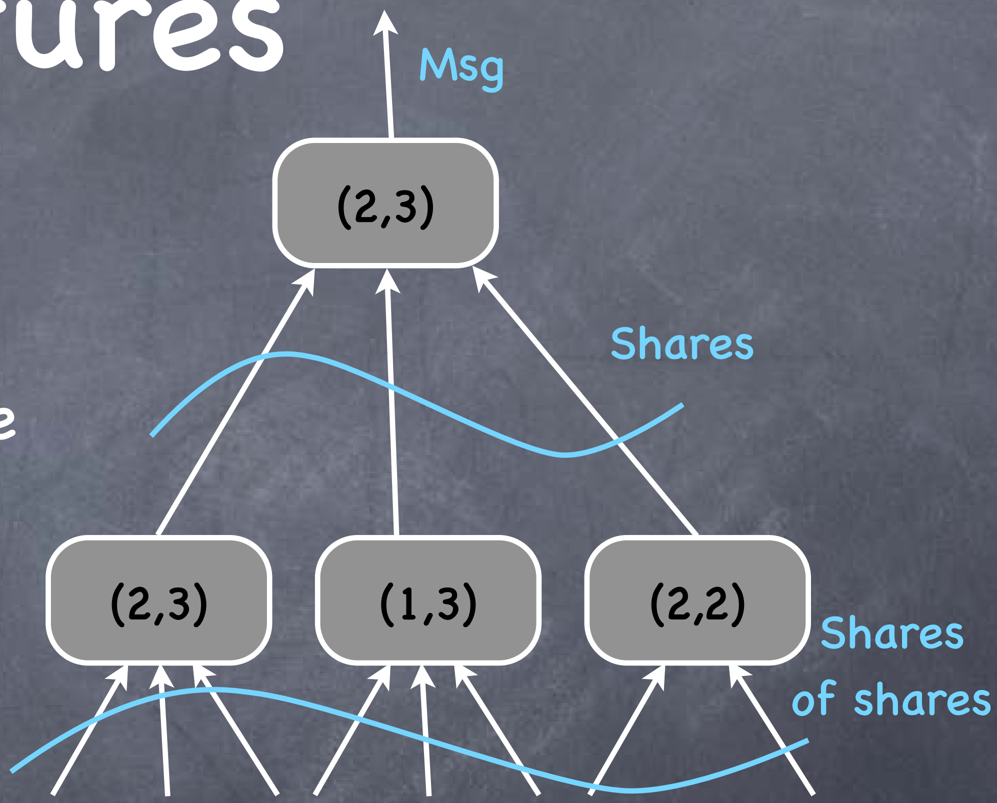

- A simple generalization of threshold access structures
  - A threshold tree to specify the access structure
  - Can realize by recursively threshold secret-sharing the shares
- Note: linear secret-sharing
- Fact: Access structures that admit linear secret-sharing are those which can be specified using "monotone span programs"

Today

- Secret-sharing schemes
  - (n,t) Threshold secret-sharing
    - Additive sharing for (n,n)
    - Shamir secret-sharing for all (n,t)
      - Optimal (ideal) when message-space is a field with more than n elements
  - Linear secret-sharing