Practical Large-scale Distributed Key Generation

John Canny, Stephen Sorkin

{jfc,ssorkin}@cs.berkeley.edu

Computer Science Division

UC Berkeley

Motivation

- Homomorphism-based computation cheap and useful.
- Key generation the limiting factor.
- Broadcast not feasible.
- Result: Trade slight chance of failure for performance.

Outline

- Motivation
- Security model
- Threshold Cryptography Refresher
- Polynomial-based Key Generation
- Matrix-based Key Generation
- Sparse matrix-based Key Generation

Security Model

- Reliable point-to-point links.
- No broadcast channel (implement with Byzantine Agreement).
- Static adversary.
- Common stream of randomness.

Discrete Log Threshold Cryptography

- n players: $P_1, \ldots, P_n$.
- Each player has a share of the private key x.
- Any $t+1$ able to sign or decrypt.
- Public keys g, g^x known to everyone.

Key Generation with a Dealer

- Dealer chooses degree t polynomial.
  (any $t+1$ evaluation points allows for interpolation)
- Distribute $f(i)$ to P_i.
- Define the secret to be $x = f(0)$.
- With $t+1$ players we know $f(1), f(2), \ldots, f(t+1)$.
- Interpolate to find $f(0)$.
- Compromising dealer will reveal key!

Requirements for DKG

Fewer than t adversarially controlled players.

Correctness:

(C1) All subsets of $t+1$ shares provided by honest players define the same unique secret key x.

(C2) All honest parties have the same value of the public key $y = g^x \mod p$, where x is the unique secret guaranteed by (C1).

(C3) x is uniformly distributed in Z_q (and hence y is uniformly distributed in the subgroup generated by g).

Secrecy:

(S1) The adversary can learn no information about x except for what is implied by the value $y = g^x \mod p$.

Polynomial DKG

Pedersen91, GJKR99

- Each player, P_i, picks random degree t polynomial, f_i.
- P_i commits to the coefficients of f_i.
- P_i shares $f_i(j)$ with P_j.
- Define global secret poly $f(\cdot) \triangleq \sum_{\{i | i \text{ is valid}\}} f_i(\cdot)$.
- The secret key is $f(0)$.
- Any $t+1$ players can find f_i and hence f.

Efficiency of Polynomial DKG

For $t = n/2$, player P_i must:

- send $O(n)$ point-to-point messages.
- broadcast $O(n)$ commitments.
- receive $O(n^2)$ messages.
- check validity of n shares.

Polynomials as Matrices

The Vandermonde matrix makes polynomial evaluation the same as matrix multiplication.

$$V = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & & n \\ \vdots & & \ddots & \vdots \\ 1 & 2^{m-1} & \cdots & n^{m-1} \end{bmatrix}$$

Premultiplying by row vector of coefficients yields row vector of evaluations.

Intuition

- What if P_i broadcasts to a much smaller group?
- Call this group Q_i, the checking group for P_i.
- For $n^{-\kappa}$ chance of failure, need $|Q_i| = \Omega(\kappa \log n)$.
- Only $n^{-\kappa+1}$ chance of failure for all groups.
- What about P_i's secret?
- Shared with only $\Theta(\log n)$ players.
- If more than $\Theta(\log n)$ degrees of freedom, recovery impossible.

Requirements for DKG

With high probability, for threshold γ:

(C1) All shares provided by honest players define the same unique secret key x, or no key at all.

(C2) All honest parties have the same value of the public key $y = g^x \mod p$, where x is the unique secret guaranteed by (C1).

(C3) x is uniformly distributed in Z_q (and hence y is uniformly distributed in the subgroup generated by g).

(C4) Almost all subsets of $(\gamma + \epsilon)n$ players can recover the key.

(S1) An adversary who corrupts fewer than $(\gamma - \epsilon)n$ players can learn no information about x except for what is implied by the value $y = g^x \mod p$.

Sparse Secret, Sparse Matrix

Each secret has k consecutive non-zero elements.

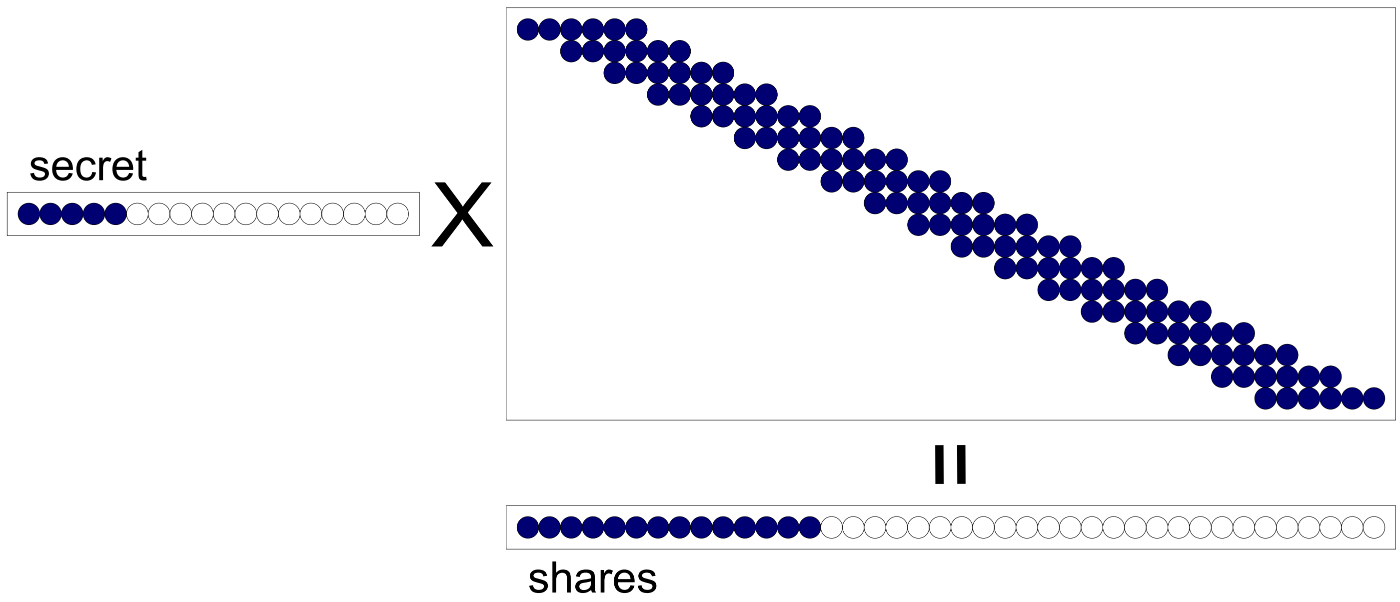

- Premultiply by secret, get vector of $\sim 2k$ non-zero shares.
- Sum of secrets is the global secret.
- Sum of shares are shares of global secret.

Missing Shares

If only a subset of the shares can be used:

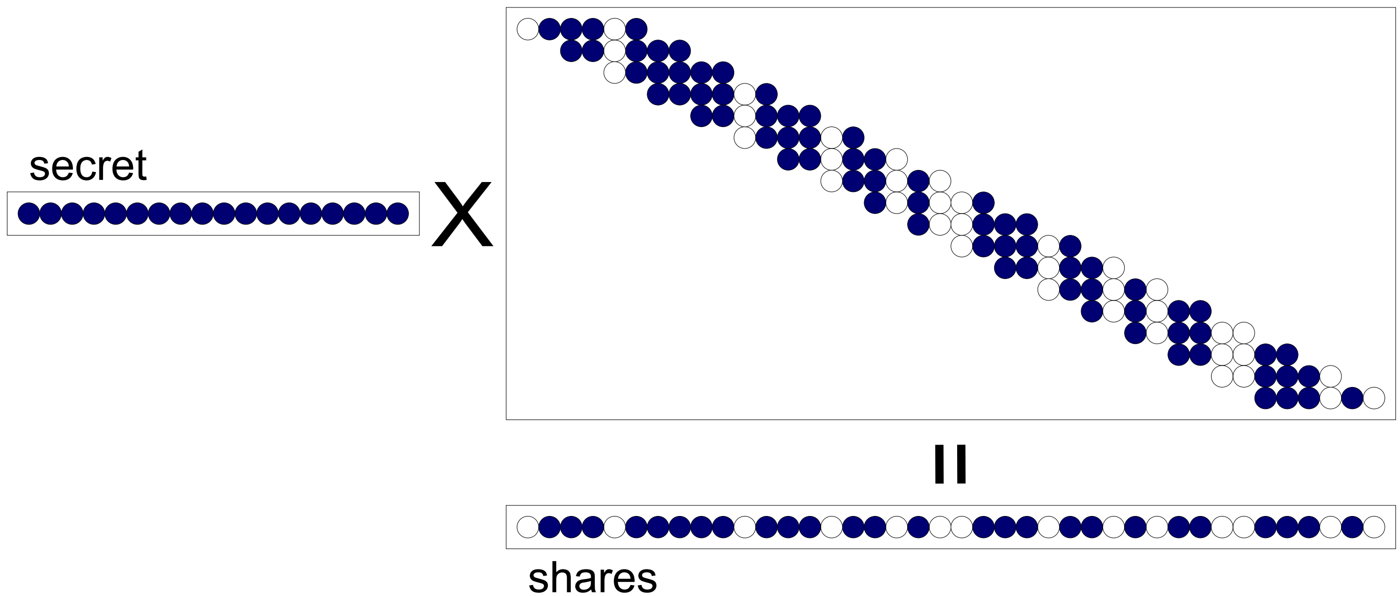

- Secret must still satisfy this smaller set of linear constraints.
- Are there enough to find the secret?

Recovery

When is recovery possible?

- Each column of the evaluation matrix represents one player's share.
- The sum of all players' secrets can be recovered if the submatrix has full rank.

Proof sketch

- Construct non-singular matrix incrementally as columns added.
- Failure if no more non-zero elements in a given row.
- We have ℓ chances to get a non-zero element.
- $\frac{1}{2} + \epsilon$ chance of getting any given column.
- Process identical to a reflecting random walk.

In practice (Band width)

For $\epsilon = 1/10$:

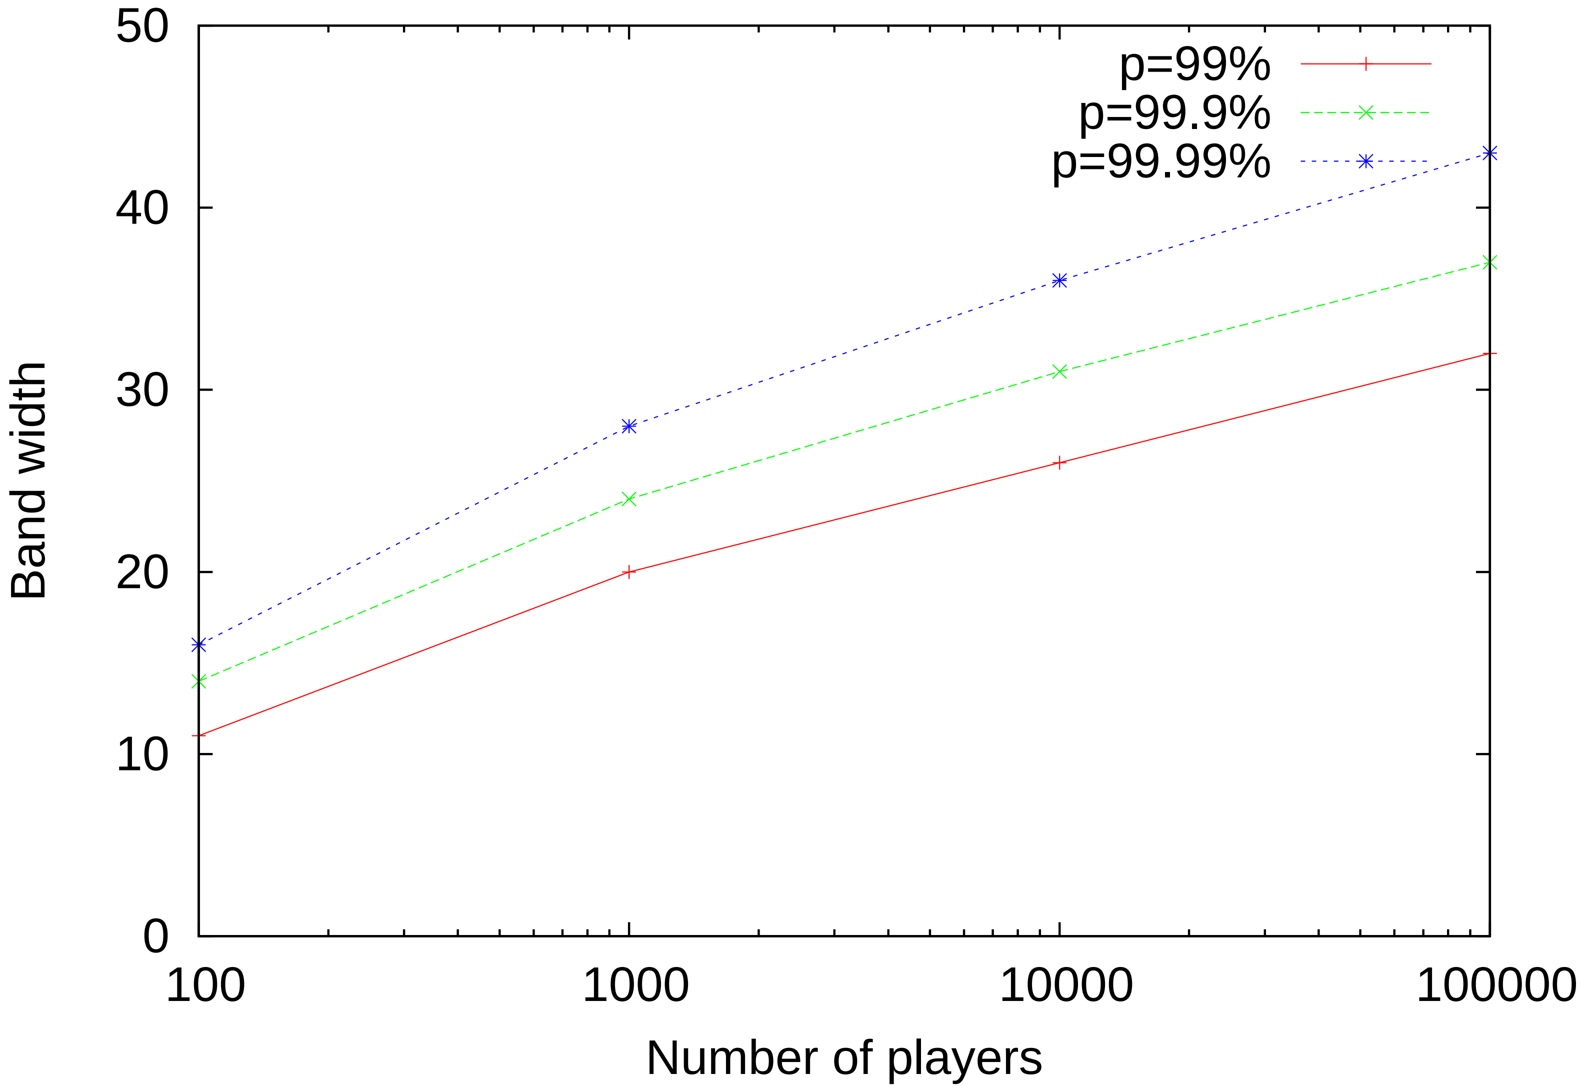

In practice (Group size)

For $\epsilon = 1/10$:

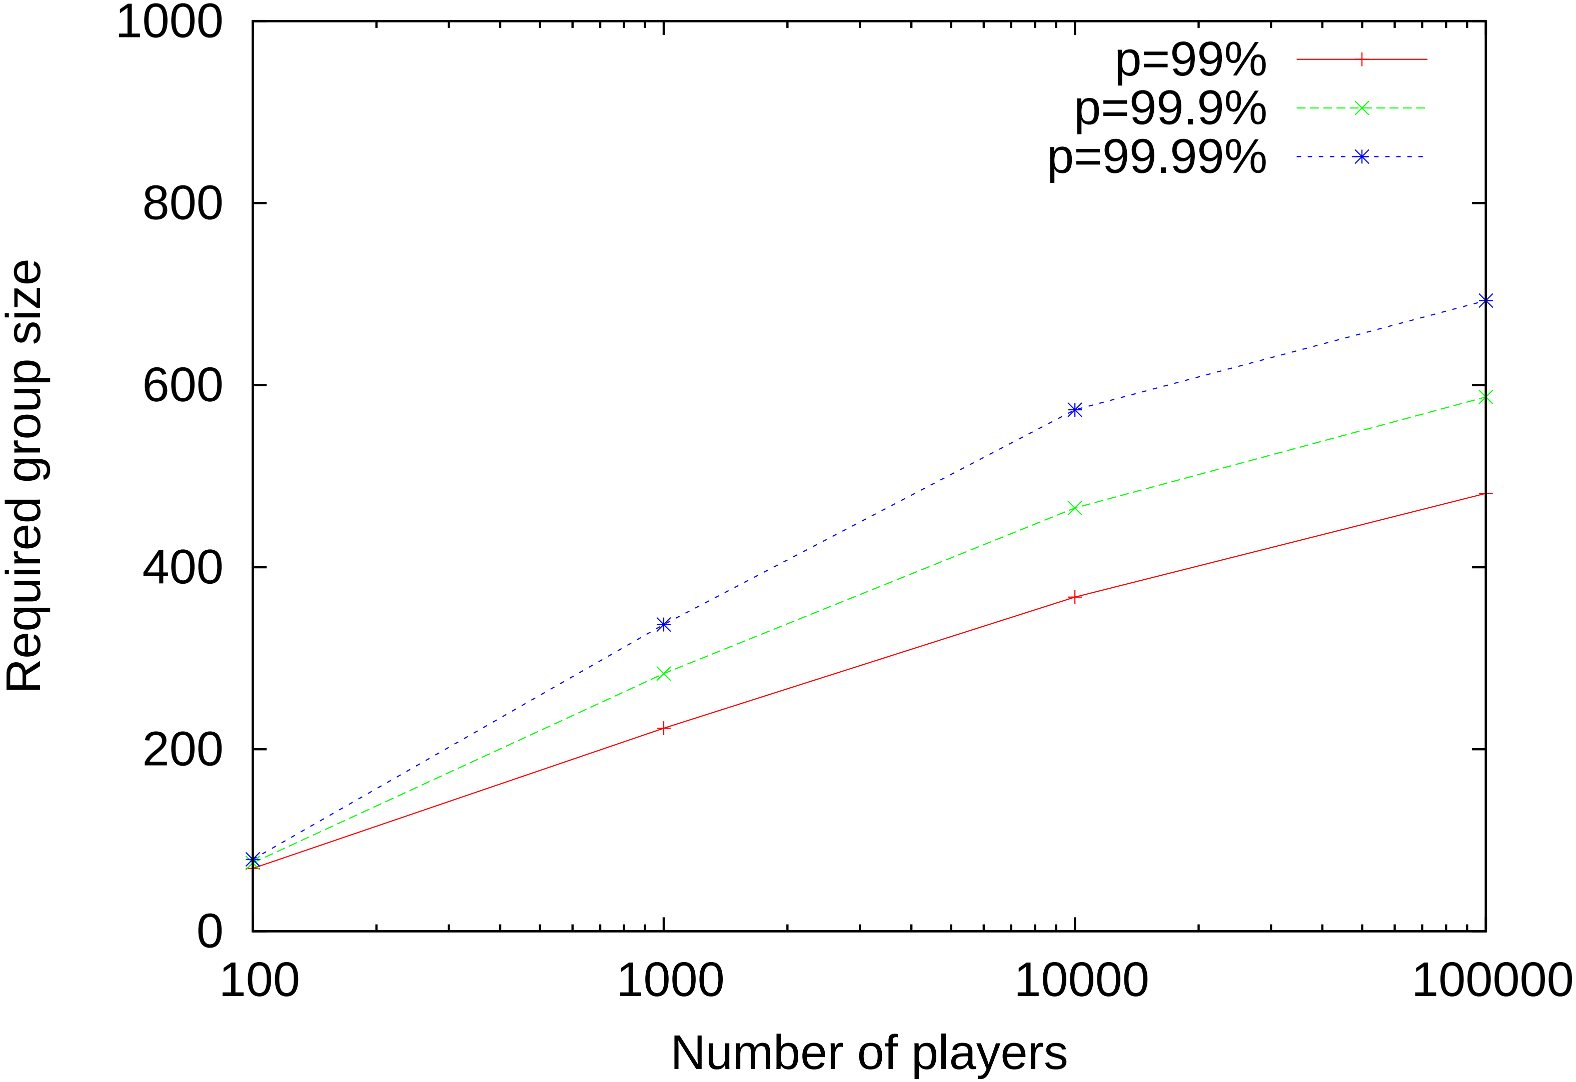

What do we get?

- Broadcast to only $\Theta(\log n)$.
- Checking only $\Theta(\log n)$ other players.
- Slight chance of failure.
- Not as sharp threshold.