Communication Cost of Distributed Computing

Abbas El Gamal

Stanford University

Brice Colloquium, 2009

Motivation

- Performance of distributed information processing systems often limited by communication

  Digital VLSI

  Multi-processors

  Data centers

  Peer-peer networks

  Sensor networks

  Networked mobile agents

- Purpose of communication is to make decision, compute function, coordinate action based on distributed data

How much communication is needed to perform such a task?

Wireless Video Camera Network

- Today's surveillance systems: Analog; costly; human operated
- Future systems: Digital; networked; self-configuring; automated detection, e.g., suspicious activity, localization, tracking, ...

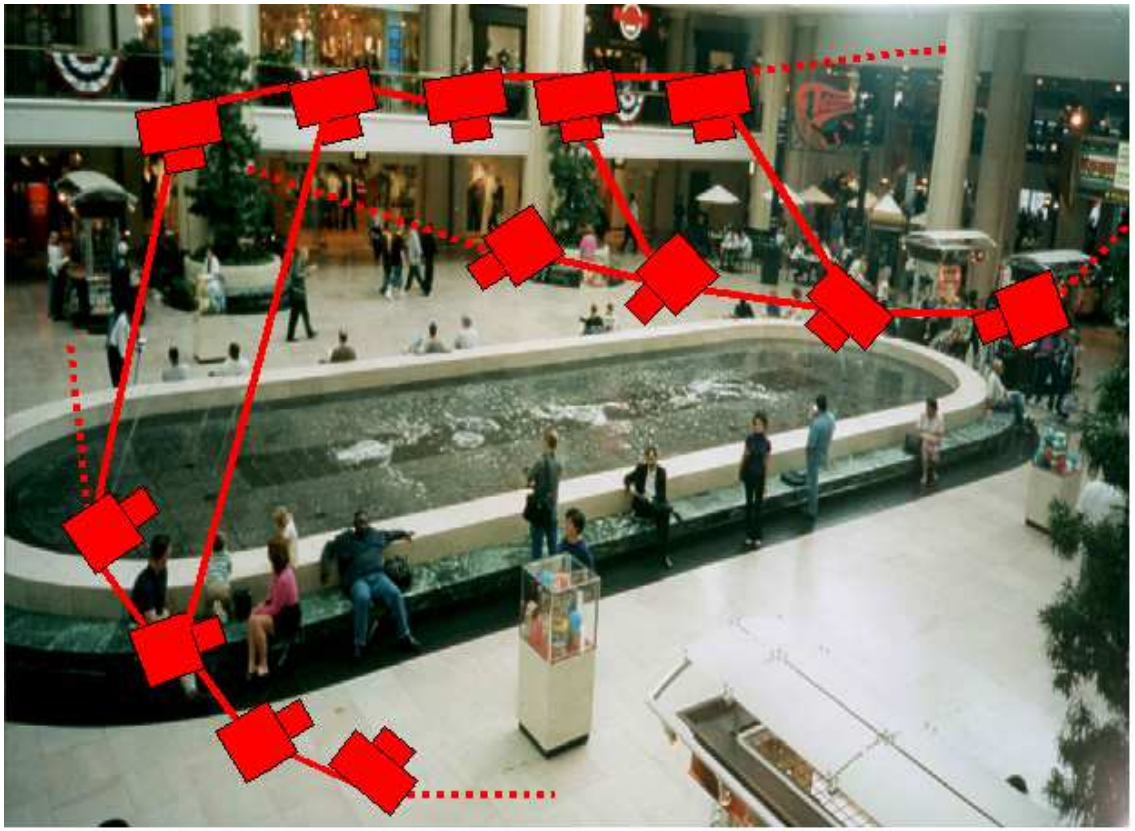

Sending all video data requires large communication bandwidth, energy

Distributed Computing Approach (YSEG 04, EEG 07)

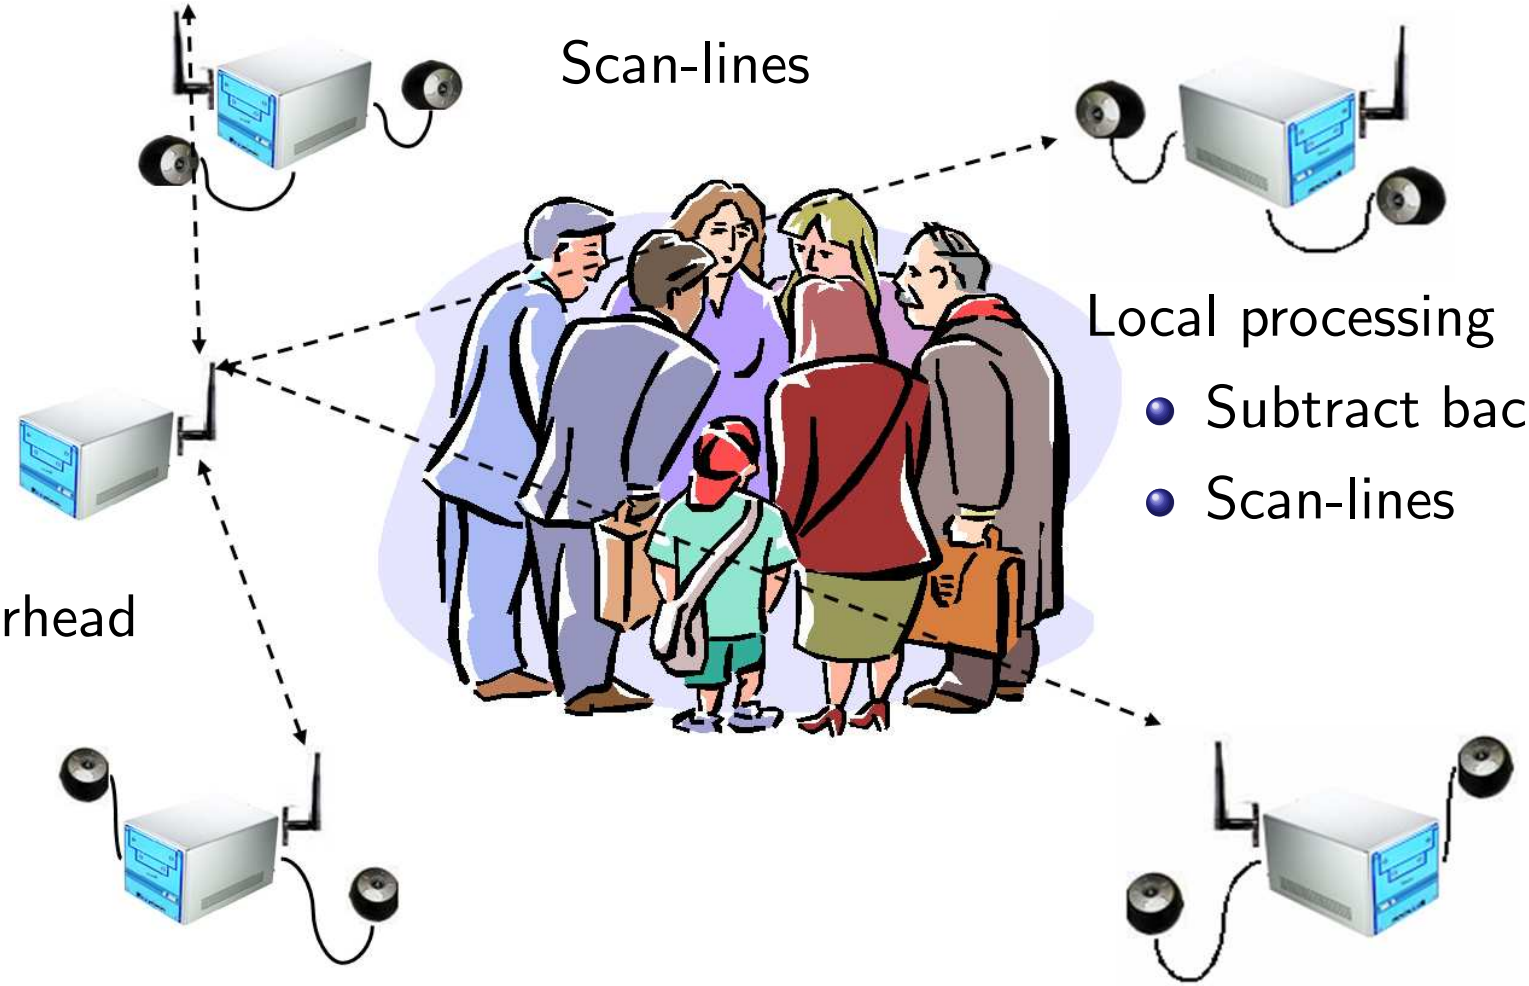

Local processing

- Subtract background
- Scan-lines

Experimental Setup

View of the setup

View from a camera

Top view of room

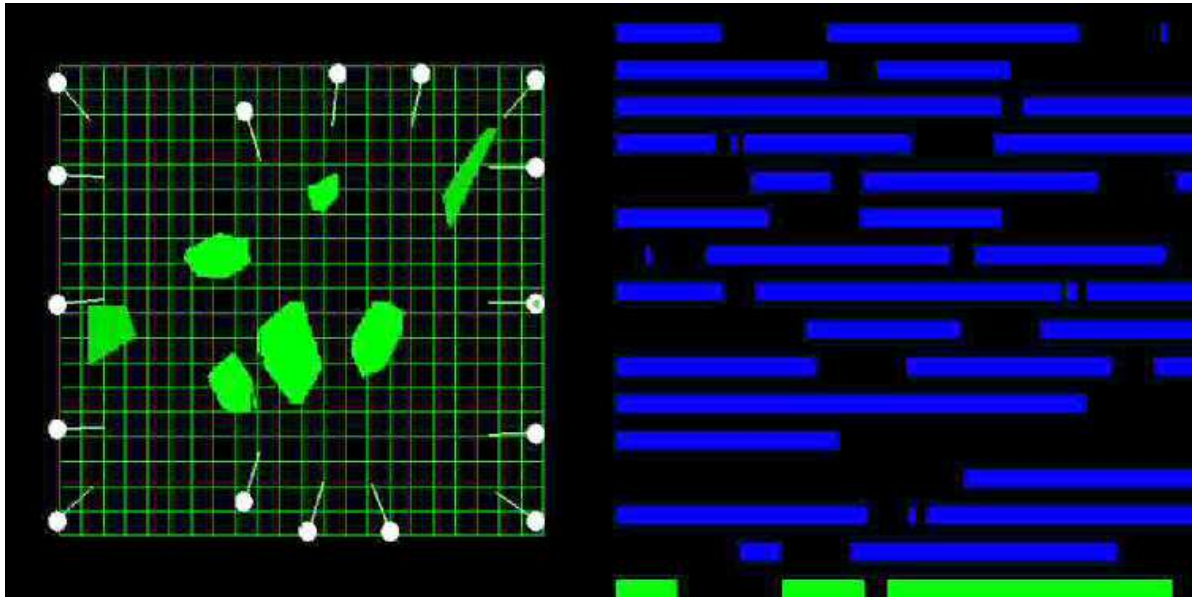

Scan-lines from 16 cameras

Problem Setup

- Network with m nodes. Node j has data x_j
- Node j wishes to *estimate* $g_j(x_1, x_2, \ldots, x_m)$, $j = 1, 2, \ldots, m$
- Nodes communicate and perform local computing

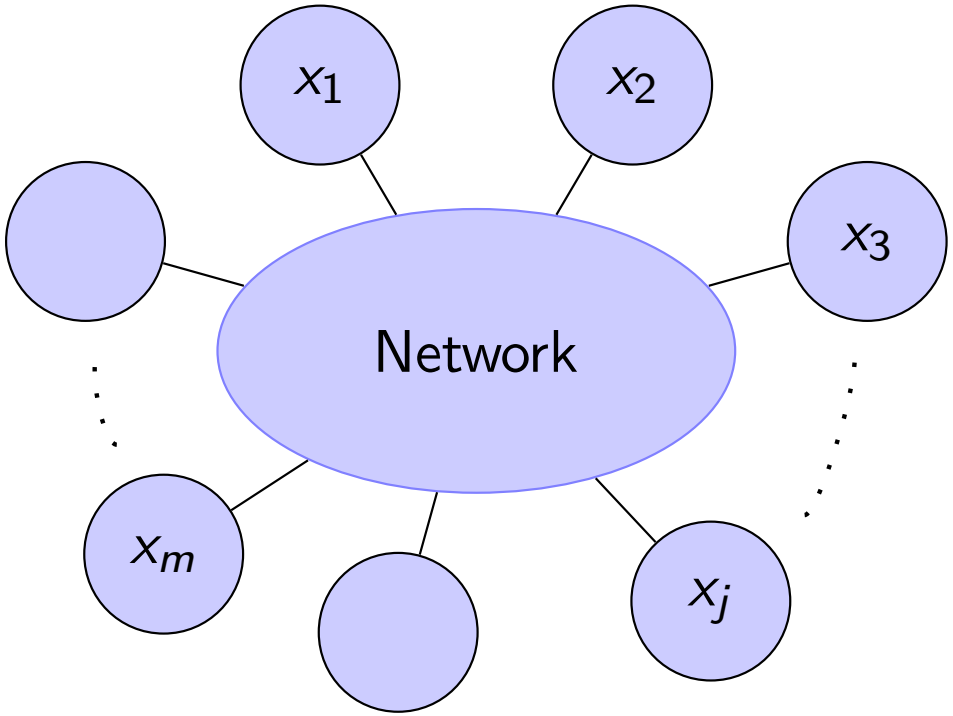

What is the minimum amount of communication needed?

Communication Cost of Distributed Computing

- Problem formulated and studied in several fields:
  - Computer science: Communication complexity; gossip
  - Information theory: Coding for computing; CEO problem
  - Control: Distributed consensus

- Formulations differ in:
  - Data model
  - Type of communication-local computing protocol
  - Estimation criterion
  - Metric for communication cost

Outline

Communication Complexity (Yao 79)

- Two nodes, two-way communication
- Node $j = 1, 2$ has binary k-vector x_j
- Both nodes wish to compute $g(x_1, x_2)$
- Use round robin communication and local computing protocol

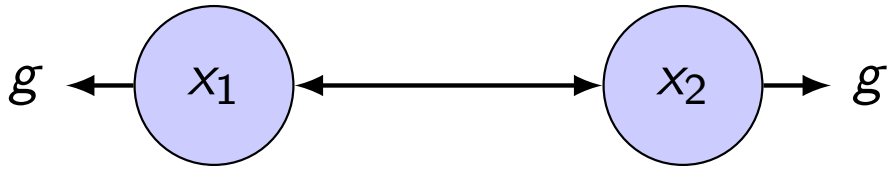

What is minimum number of bits of communication needed (communication complexity, $C(g)$) ?

Example (Equality Function)

$g(x_1, x_2) = 1$ if $x_1 = x_2$, and 0, otherwise

Upper bound: $C(g) \leq k + 1$ bits

For $k = 2$: $C(g) \leq 3$ bits

		x_2			
		00	01	10	11
	00	1	0	0	0
x_1	01	0	1	0	0
	10	0	0	1	0
	11	0	0	0	1

Lower bound: *Every* communication protocol partitions $\{0,1\}^k \times \{0,1\}^k$ into g-monochromatic rectangles, each having distinct *code*

If $r(g)$ is minimum number of such rectangles, then $C(g) \geq \log r(g)$

For $k = 2$, $r(g) = 8 \Rightarrow C(g) \geq 3$ bits

In general, $C(g) \geq k + 1 \Rightarrow C(g) = k + 1$ bits

Number-on-Forehead Game

- m players. Player j has binary k-vector x_j on her forehead
- Player j can see all numbers except her own, x_j
- Every player wants to know if $x_1 = x_2 = \ldots = x_m$ or not
- Players take turns in broadcasting to all other players
- Recall for $m = 2$, $C(g) = k + 1$
- What is $C(g)$ for $m \geq 3$?
- Consider the following protocol:
  - Player 1 broadcasts $A = 1$ if $x_2 = x_3 = \ldots = x_m$, $A = 0$, otherwise
  - If $A = 0$, then all players know that values are not equal
  - If $A = 1$ only player 1 does not know if all values are equal or not; Player 2 broadcasts $B = 1$ if $x_1 = x_3$, $B = 0$, otherwise

$C(g) = 2$ bits suffice independent of m and k!!

Overlap in information can significantly reduce communication

More on Communication Complexity

- Other definitions of communication complexity:
  - Average: Uniform distribution over node values
  - Randomized: Nodes use "coin flips" in communication
  - Non-deterministic: Non-zero probability of computing error
- Some references:

A. Yao, "Some complexity questions related to distributive computing," *STOC*, 1979

A. Orlitsky, A. El Gamal, "Average and randomized communication complexity," *IEEE Trans. on Info. Th.*, 1990

R. Gallager, "Finding parity in a simple broadcast network," *IEEE Trans. on Info. Th.*, 1988

E. Kushilevitz and N. Nisan, Communication Complexity, Cambridge University Press, 2006

Information Theory: Lossless Computing

- m node network. Node j observes source X_j
- X_j generates i.i.d. process $\{X_{j1}, X_{j2}, \ldots\}$, $X_{ji} \in \mathcal{X}$ a finite set
- Sources may be correlated: $(X_1, X_2, \ldots, X_m) \sim p(x_1, x_2, \ldots, x_m)$
- Node j wishes to finds an estimate $\hat{g}_{ji}$ of $g_j(X_{1i}, X_{2i}, \ldots, X_{mi})$ for each sample $i = 1, 2, \ldots$
- Block coding: Nodes communicate, perform local computing over blocks of n-samples

Lossless criterion

$$\hat{g}_j^n = g_j^n \text{ for all } j \in \{1, 2, \ldots, m\} \text{ with high probability } (whp)$$

- What is the minimum sum rate in bits/sample-tuple?
- Problem is in general open, even for 3-nodes

Lossless Compression

- Two nodes, one-way communication
- Sender node observes i.i.d. source $X \sim p(x)$
- Receiver node wishes to losslessly estimate $\hat{X}^n$ of X^n

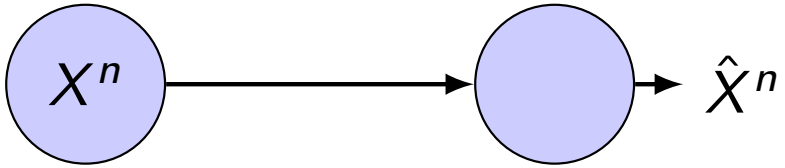

Theorem (Shannon 1948)

The minimum lossless compression rate is the entropy of X

$$H(X) := -\mathrm{E}\left(\log p(X)\right) \ \textit{bits/sample}$$

Example (Binary Source)

$X \sim \mathrm{Bern}(p)$: $H(p) := -p \log p - (1-p)\log(1-p) \in [0,1]$

Random binning proof (Cover 1975)

Ket facts:

- If x^n is generated i.i.d. $\sim p(x)$, then by law of large numbers, its empirical distribution is close to $p(x)$, i.e. it is typical, *whp*
- There are $\doteq 2^{nH(X)}$ typical sequences

Proof outline:

- Randomly assign every sequence $x^n \in \mathcal{X}^n$ to one of 2^{nR} *bins*
  Node assignments revealed to both nodes

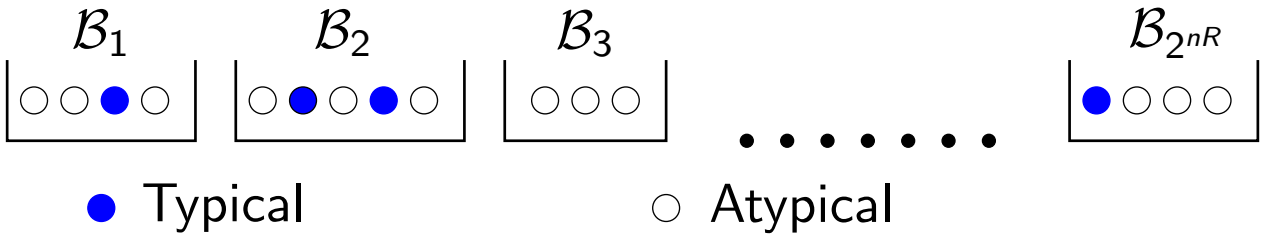

- Sender: Upon observing x^n, send bin index (nR bits)
- Receiver: If x^n is *unique* typical sequence in bin, set $\hat{x}^n = x^n$; otherwise declare an error
- Probability of correct decoding:
  Since x^n is typical *whp*, bin has ≥ 1 typical sequences *whp*
  If $R > H(X)$, bin has 1 typical x^n *whp* $\Rightarrow \hat{X}^n = X^n$ *whp*

Lossless Compression with Side Information

- Two nodes, one way communication
- Sender node has i.i.d. source X, receiver node has i.i.d. source Y
- $(X, Y) \sim p(x, y)$
- Receiver node wishes to estimate X^n losslessly

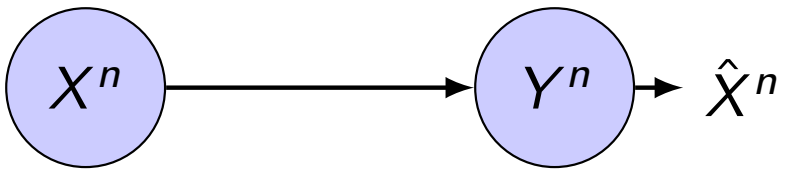

- $H(X)$ sufficient; Can we do better?
- If sender also knows Y^n, then minimum rate is conditional entropy

$$H(X|Y) := -\mathrm{E}_{X,Y}\left(\log p(X|Y)\right) \leq H(X)$$

Theorem (Slepian-Wolf 1973)

Minimum rate with side information at receiver only is $H(X|Y)$

- Proof uses random binning

Example (Doubly Symmetric Binary Sources)

$Y \sim \mathrm{Bern}(1/2)$ and $Z \sim \mathrm{Bern}(p)$ be independent, $X = Y \oplus Z$

Receiver has Y, wishes to losslessly estimate X

Minimum rate: $H(X|Y) = H(Y \oplus Z|Y) = H(Z) = H(p)$

For example, if $Z \sim \mathrm{Bern}(0.01)$, $H(p) \approx 0.08$
Versus $H(X) = H(1/2) = 1$ if correlation is ignored

Remarks

Shannon theorem holds for error-free compression (e.g., Lempel-Ziv)

Slepian-Wolf does not hold in general for error-free compression

Lossless Computing with Side Information

- Receiver node wishes to compute $g(X, Y)$
- Let R_g^* be minimum rate

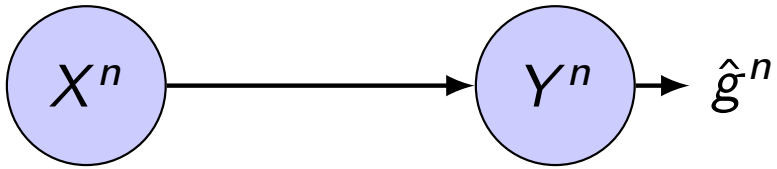

Upper and lower Bounds

$$R_g^* \leq H(X|Y)$$

$$R_g^* \geq H(g(X, Y)|Y)$$

- Bounds sometimes coincide

Example (Mod-2 Sum)

$Y \sim \mathrm{Bern}(1/2)$, $Z \sim \mathrm{Bern}(p)$ independent, $X = Y \oplus Z$
Function: $g(X, Y) = X \oplus Y = Z$
Lower bound: $H(X \oplus Y | Y) = H(Z|Y) = H(p)$
Upper bound: $H(X|Y) = H(p)$
$\Rightarrow R_g^* = H(p)$

- Bounds do not coincide in general

Example

$X = (V_1, V_2, \ldots, V_{10})$, V_j i.i.d. $\mathrm{Bern}(1/2)$, $Y \sim \mathrm{Unif}\{1, 2, \ldots, 10\}$
Function: $g(X, Y) = V_Y$
Lower bound: $H(V_Y|Y) = 1$ bit
Upper bound: $H(X|Y) = 10$ bits
Can show that $R_g^* = 10$ bits

Theorem (Orlitsky, Roche 2001)

$$R_g^* = H_G(X|Y)$$

H_G is conditional entropy of characteristic graph of X, Y, and g

Generalizations:

- Two way, two rounds:

  A. Orlisky, J. Roche, "Coding for computing," *IEEE Trans. on Info. Th.*, 2001

- Infinite number of rounds:

  N. Ma, P. Ishwar, "Two-terminal distributed source coding with alternating messages for function computation," *ISIT*, 2008

Lossless Computing: $m = 3$

- 2-sender, 1-receiver network, one-way communication
- Receiver wishes to estimate a function $g(X, Y)$ losslessly
- What is the minimum *sum rate* in bits/sample-pair?

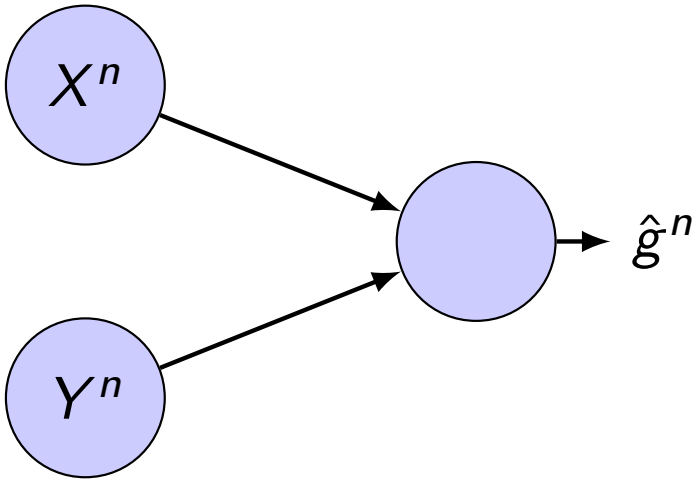

- Problem is in general open

Distributed Lossless Compression

- Let $g(X, Y) = (X, Y)$

Theorem (Slepian-Wolf)

The minimum achievable sum rate is the joint entropy

$$\begin{aligned} H(X, Y) &= H(X) + H(Y|X) \\ &= H(Y) + H(X|Y) \end{aligned}$$

- Same as if X, Y were jointly encoded!
- Achievability uses random binning
- Theorem can be generalized to m nodes

Distributed Mod-2 Sum Computing

- Let $Y \sim \mathrm{Bern}(1/2)$, $Z \sim \mathrm{Bern}(p)$ independent, $Y = X \oplus Z$
- Receiver wishes to compute $g(X, Y) = X \oplus Y = Z$
- Slepian-Wolf gives sum rate of $H(X, Y) = 1 + H(p)$
- Can we do better?

Theorem (Körner, Marton 1979)

Minimum sum rate is: $2H(p)$

- Minimum sum rate can be $<< 1 + H(p)$
- Achievability uses random linear coding:
  - Randomly generate $k \times n$ binary matrix A
  - Senders: Send binary k-vectors AX^n, AY^n
  - Receiver: Receiver node computes $AX^n \oplus AY^n = AZ^n$
  - If $k > nH(p)$, AZ^n *uniquely* determines Z^n *whp*

Information Theory: Lossy Computing

- Node j observes i.i.d. source X_j
- Every node wishes to estimate $g(X_{1i}, X_{2i}, \ldots, X_{mi})$, $i = 1, 2, \ldots$
- Nodes communicate and perform local computing
- As in lossless case, we use block codes

MSE distortion criterion

$$\frac{1}{mn} \sum_{j=1}^{m} \sum_{i=1}^{n} \mathrm{E}\left[(\hat{g}_{ij} - g_i)^2\right] \leq D$$

- What is the minimum sum rate $R^*(D)$ bits/sample-tuple?
- Problem is in general open (even for 2-nodes)

Lossy Compression

- Two nodes, one-way communication
- Sender node observes i.i.d. Gaussian source $X \sim \mathrm{N}(0, P)$
- Receiver node wishes to estimate $\hat{X}$ to prescribed MSE distortion D

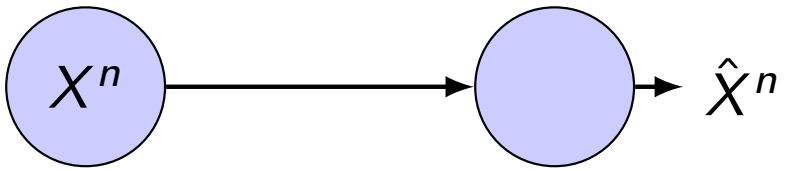

- What is the minimum required rate $R(D)$?

Theorem (Shannon 1949)

Let $d := D/P$ be normalized distortion. The rate distortion function is

$$R(D) = \begin{cases} \frac{1}{2}\log\left(\frac{1}{d}\right) & d \leq 1 \\ 0 & d > 1 \end{cases}$$

Lossy Compression with Side Information

- Two nodes, one-way communication
- (X, Y) Gaussian; 0 mean; average power P; correlation coeff. ρ
- Receiver wishes to estimate X to MSE distortion D

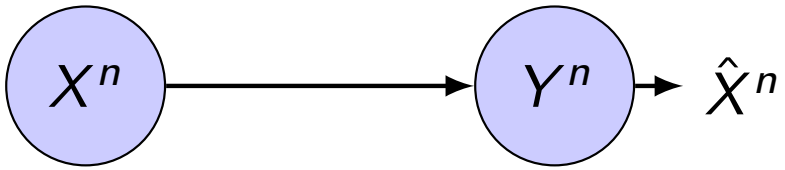

Theorem (Wyner-Ziv 1976)

$$R(D) = \begin{cases} \frac{1}{2}\log\left(\frac{(1-\rho^2)}{d}\right) & d < (1-\rho^2) \\ 0 & \textit{otherwise} \end{cases}$$

$d = D/P$

Lossy Averaging with Side Information

Receiver node now wishes to estimate $(X+Y)/2$ to MSE distortion D

Proposition

$$R(D) = \begin{cases} \frac{1}{2}\log\left(\frac{(1-\rho^2)}{d}\right) & d < (1-\rho^2) \\ 0 & \text{otherwise} \end{cases}$$

$d = 4D/P$

- Same rate as sending $X/2$ with MSE distortion D

For two-way; both nodes wish to estimate average with MSE D

Proposition (Su, El Gamal 2009)

$$R(D) = \begin{cases} \log\left(\frac{(1-\rho^2)}{d}\right) & d < (1-\rho^2) \\ 0 & \text{otherwise} \end{cases}$$

- Same as two *independent* one-way rounds

Distributed Lossy Compression

- (X, Y) jointly Gaussian
- Receiver wishes to estimate X, Y each with MSE distortion D
- What is the sum rate distortion function $R(D)$?

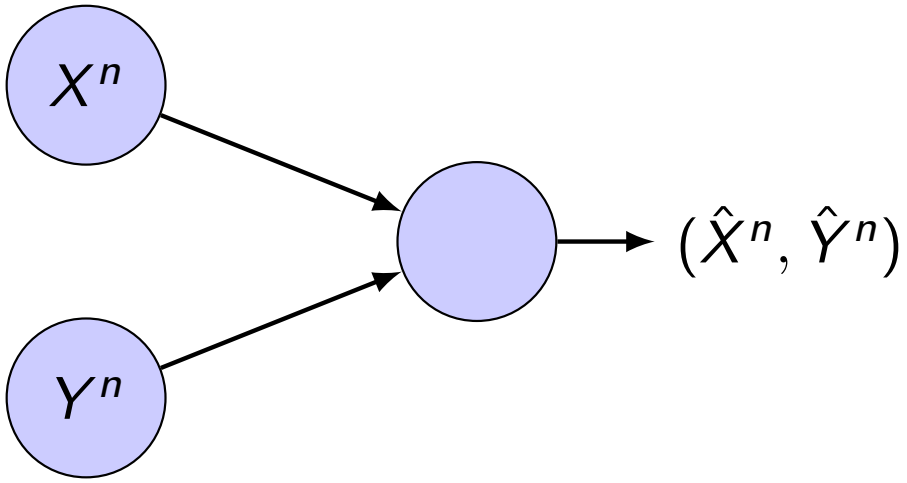

- Problem recently solved (Oohama 97; Wagner,Tavildar,Viswanath 08)
- Optimal sum rate known also for $g = \alpha X + (1-\alpha)Y$, $\alpha \in [0,1]$ (CEO, μ-sum (Oohama 05; Prabhakaran, Tse, Ramchandran 04))

Distributed Consensus: Motivation

Distributed coordination, information aggregation in distributed systems

- Flocking and schooling behaviors in nature
- Spread of rumors, epidemics
- Coordination of mobile autonomous vehicles
- Load balancing in parallel computers
- Finding file sizes, free memory size in peer-to-peer networks
- Information aggregation in sensor networks

In many of these scenarios, communication is asynchronous, subject to link and node failures and topology changes

Distributed Averaging

- m-node network. Node j observes real-valued scalar x_j
- Every node wishes to estimate the average $(1/m)\sum_{j=1}^{m} x_j$
- Communication performed in *rounds*, for example,
  - 2 nodes selected in each round
  - They exchange information and perform local computing

- Distributed Protocol: Communication and local computing do *not* depend on node identity
  - Synchronous (deterministic): Node-pair selection predetermined
  - Asynchronous (Gossip): Node-pairs selected at random

How many rounds are needed to estimate the average (time-to-consenus)?

Recent Related Work

- D. Kemp, A. Dobra, J. Gehrke, "Gossip-based computation of aggregate information," *FOCS*, 2003
- L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *CDC,* 2003
- S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, "Gossip algorithms: design, analysis and applications," *INFOCOM,* 2005 gossip, convergence time as function of graph connectivity

Example: Centralized Protocol

- Node 1 acts as cluster-head
- Rounds 1–4: Cluster-head receives values from other nodes
- The cluster head computes the average
- Rounds 5–8: Cluster-head sends average to other nodes
- The number of rounds for m-node network: $2m - 2$
- Computation is perfect

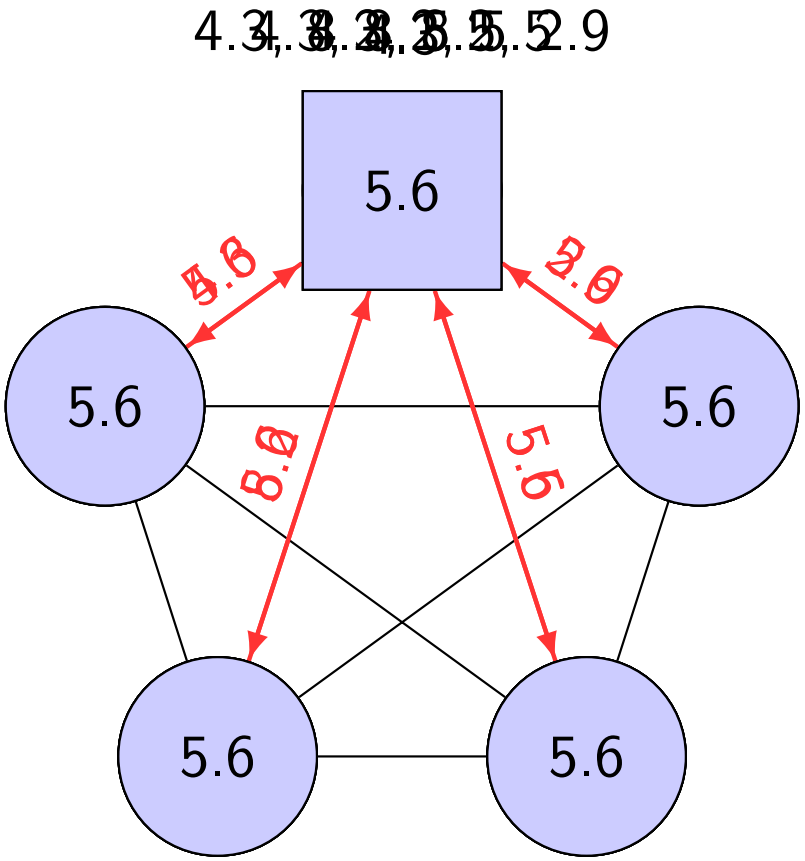

Example: Gossip Protocol

- At each round, a node-pair (j, k) is randomly selected
- Nodes exchange values and average them
- Computation is not perfect: MSE= 0.00875
- In general, *cannot* achieve perfect computation in finite number of rounds
- Estimates converge to true average

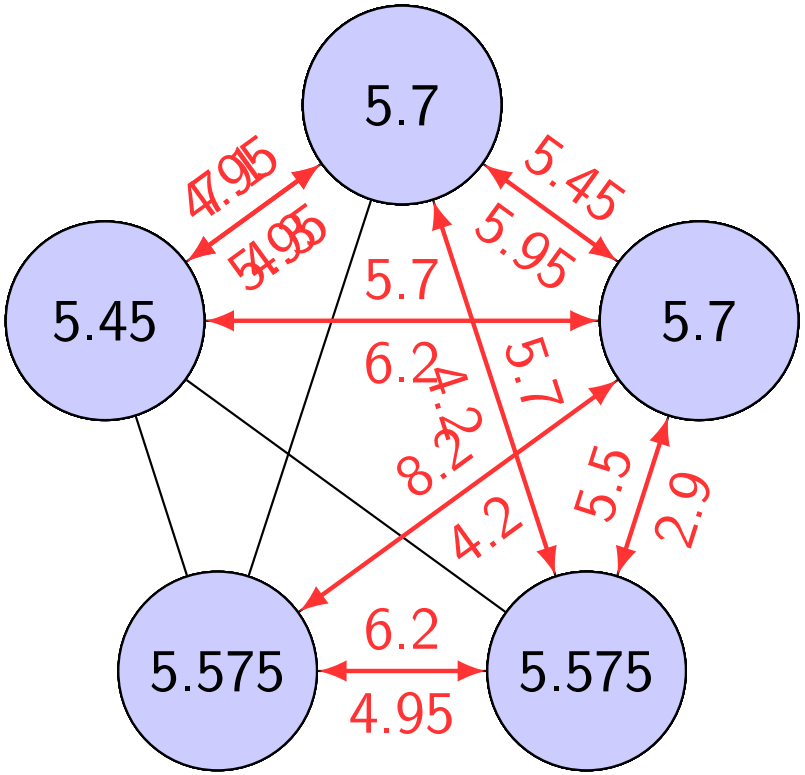

Issues

- Communicating/ computing with infinite precision not realistic
- Number of rounds not good measure of communication cost

- Quantized consensus:

  L. Xiao, S. Boyd, S.J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Distrib. Comput.*, 2007

  M. Yildiz, A. Scaglione, "Coding with side information for rate constrained consensus," *IEEE Trans. on Sig. Proc.*, 2008

  O. Ayaso, D. Shah, M. Dahleh, "Information theoretic bounds on distributed computatiton," preprint, 2008

- We recently studied lossy computing formulation:

  H. Su, A. El Gamal,"Distributed lossy averaging," *ISIT*, 2009

Distributed Lossy Averaging (Su, El Gamal 2009)

- Network with m nodes
- Node j observes Gaussian i.i.d. source $X_j \sim \mathrm{N}(0,1)$
- Assume sources are independent
- Each node wishes to estimate the average $g^n := (1/m)\sum_{j=1}^{m} X_j^n$
- Rounds of two-way, node-pair communication; block codes used

Network rate distortion function

$R^*(D)$ is minimum sum rate such that

$$\frac{1}{mn}\sum_{j=1}^{m}\sum_{i=1}^{n} \mathrm{E}\left[(\hat{g}_{ij} - g_i)^2\right] \leq D$$

- $R^*(D)$ known only for $m = 2$

Cutset Lower Bound on $R^*(D)$

Super-node
$P = \frac{m-1}{m^2}$

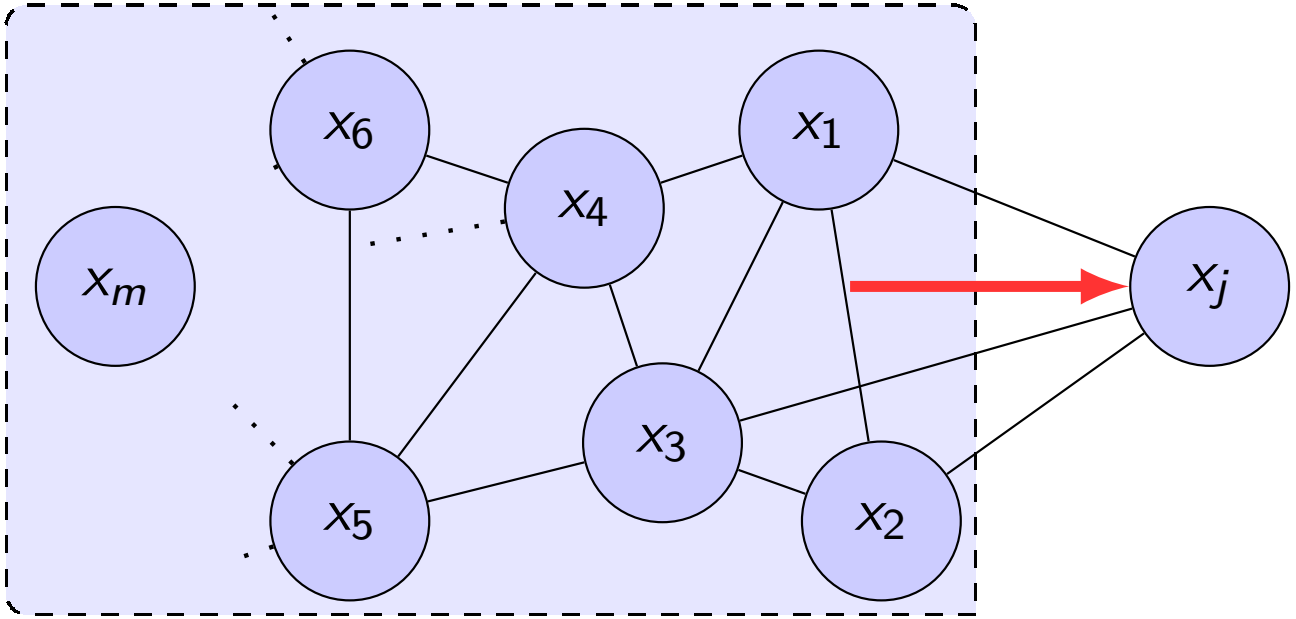

Theorem (Cutset Lower Bound (Su, El Gamal 20090))

$$R^*(D) \geq \frac{m}{2} \log \left(\frac{m-1}{m^2 D} \right) \quad \textit{for } D < (m-1)/m^2$$

- Bound is tight for $m = 2$
- Can achieve within factor of 2 using centralized protocol for large m

Gossip Protocol

- Two nodes are selected at random in each round
- Use weighted-average-and-compress scheme
- Before communication, node j sets $\hat{g}_j^n(0) = X_j^n$
- At round $t+1$: Suppose nodes j, k selected:
  - Nodes exchange lossy estimates of their states $\hat{g}_j^n(t)$, $\hat{g}_k^n(t)$ at normalized distortion d
  - Nodes update their estimates:

$$\hat{g}_j^n(t+1) = \frac{1}{2}\hat{g}_j^n(t) + \frac{1}{2(1-d)}\hat{\hat{g}}_k^n(t),$$
$$\hat{g}_k^n(t+1) = \frac{1}{2}\hat{g}_k^n(t) + \frac{1}{2(1-d)}\hat{\hat{g}}_j^n(t)$$

- Final estimate for node j is $\hat{g}_j^n(T)$
- Rate at each round is $2 \times (1/2)\log(1/d)$
- Update equations reduce to gossip when $d = 0$, rate $\to \infty$

Gossip Protocol

- Let $\mathrm{E}(R(D))$ be expected sum rate over node-pair selection
- $R^*(D) \leq \mathrm{E}(R(D))$
- Can show that $\mathrm{E}(R(D))$ larger than cutset bound by roughly a factor of $\log m$
- Cutset bound achievable within factor to 2 using centralized protocol

Price of gossiping is roughly $\log m$

- Protocol does not exploit build-up in correlation to reduce rate

Effect of Using Correlation ($m = 50$)

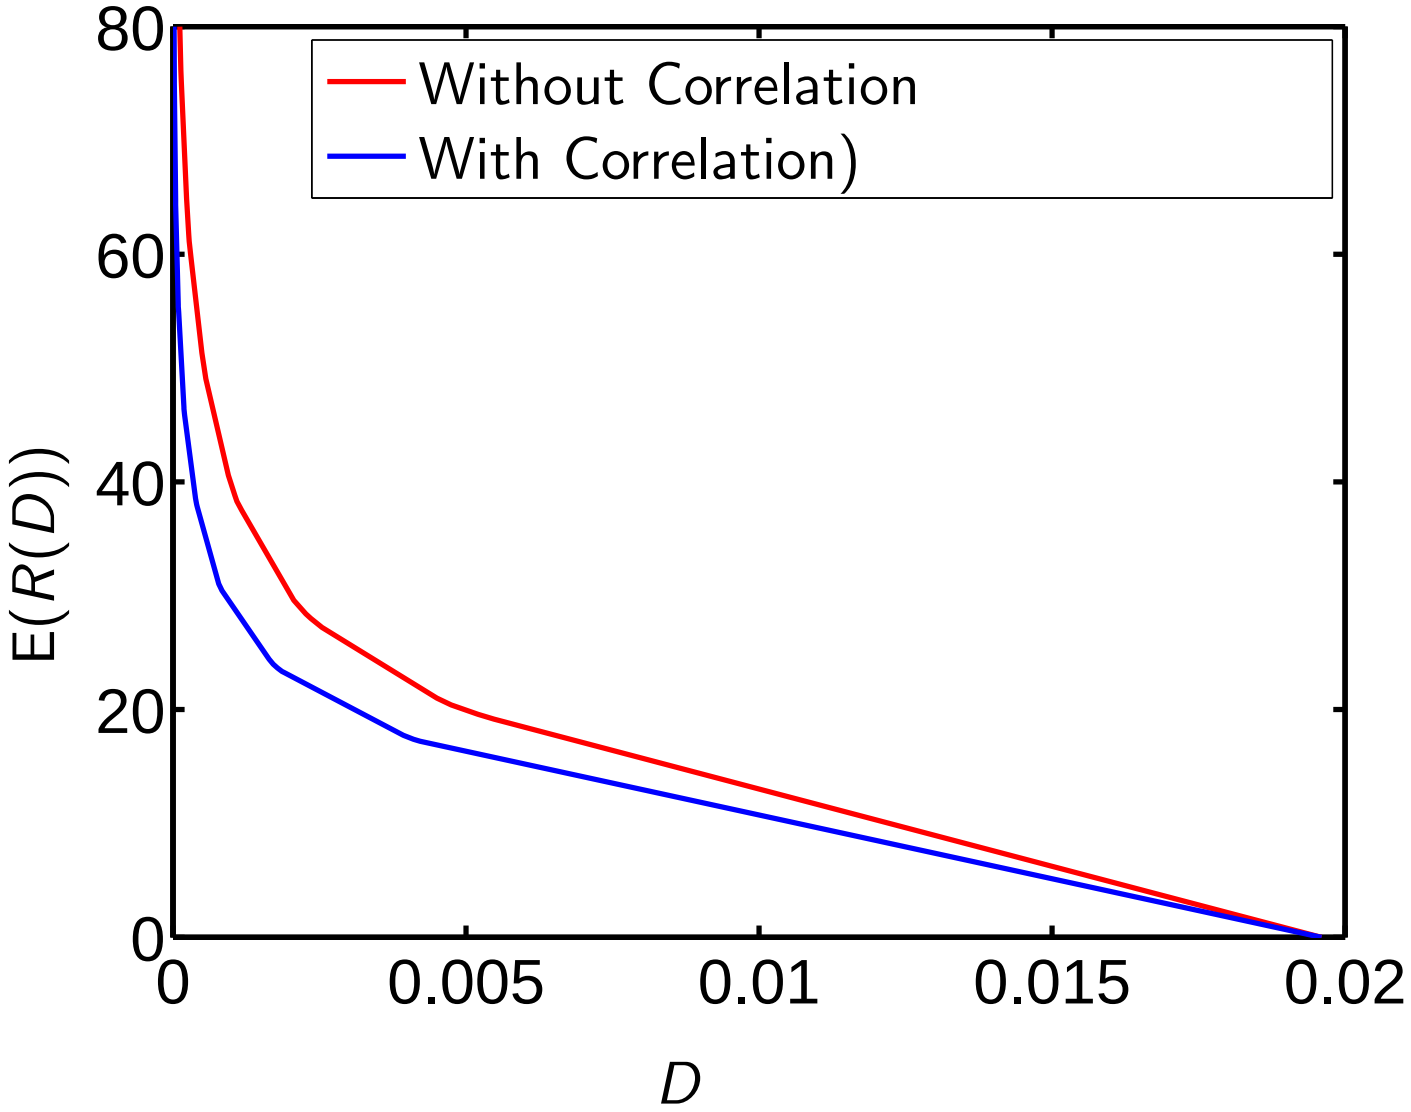

Summary

- Communication cost of distributed computing studied in several fields

	Source Model	Coding/ Computing	Estimation Criterion	Communication Cost
Comm. Comp.	discrete, Scalar	scalar	zero error	bits
Info. theory	random	block	lossless Lossy	bits/sample-tuple
Dist. Consensus	continuous, scalar	scalar	MSE	rounds

- Many cool results, but much remains to be done

Thank You