Byzantine Vector Consensus in Complete Graphs

Nitin Vaidya

University of Illinois at Urbana-Champaign

Vijay Garg

University of Texas at Austin

Assumptions

- Complete graph of n processes
- f Byzantine faults
- Each process has d-dimensional vector input

d = 2

Inputs $\begin{pmatrix}0\\1\end{pmatrix}$ $\begin{pmatrix}0\\0\end{pmatrix}$ $\begin{pmatrix}1\\0\end{pmatrix}$ $\begin{pmatrix}1\\1\end{pmatrix}$

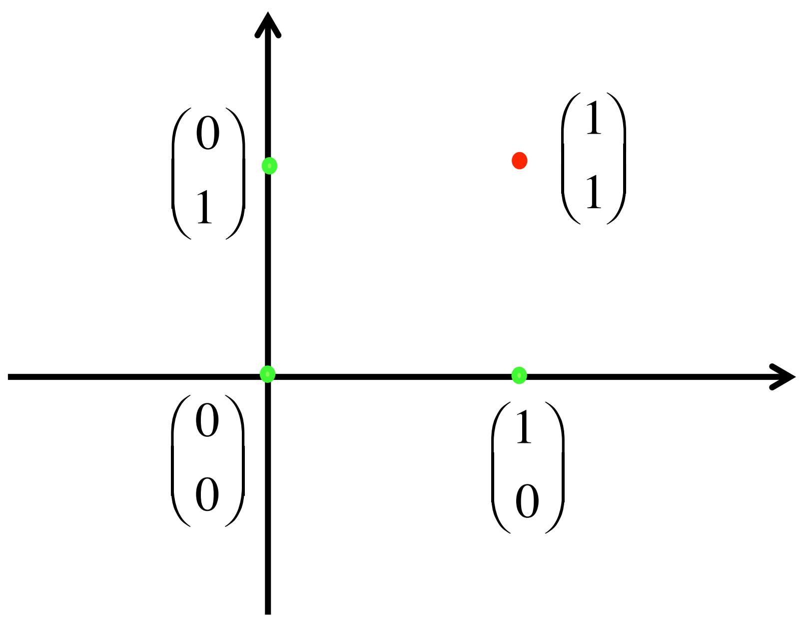

Exact Vector Consensus

- Agreement: Fault-free processes agree *exactly*
- Validity: Output vector in convex hull
  of inputs at fault-free processes
- Termination: In finite time

Inputs $\begin{pmatrix}0\\1\end{pmatrix}$ $\begin{pmatrix}0\\0\end{pmatrix}$ $\begin{pmatrix}1\\0\end{pmatrix}$ $\begin{pmatrix}1\\1\end{pmatrix}$

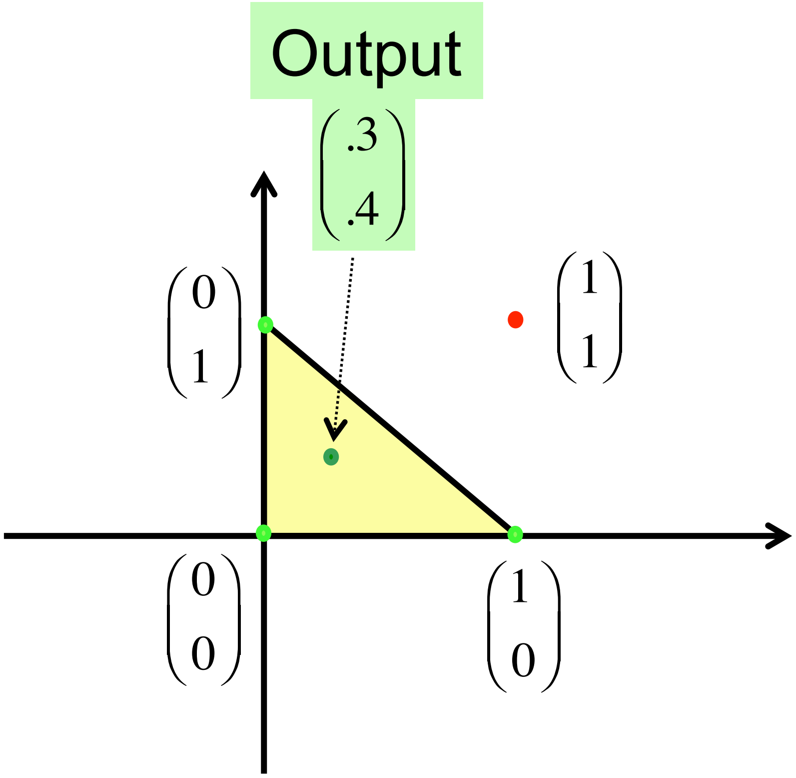

Approximate Vector Consensus

- ε-Agreement: output vector elements differ by ≤ ε
- Validity: Output vector in convex hull of inputs at fault-free processes
- Termination: In finite time

ε = 0.04

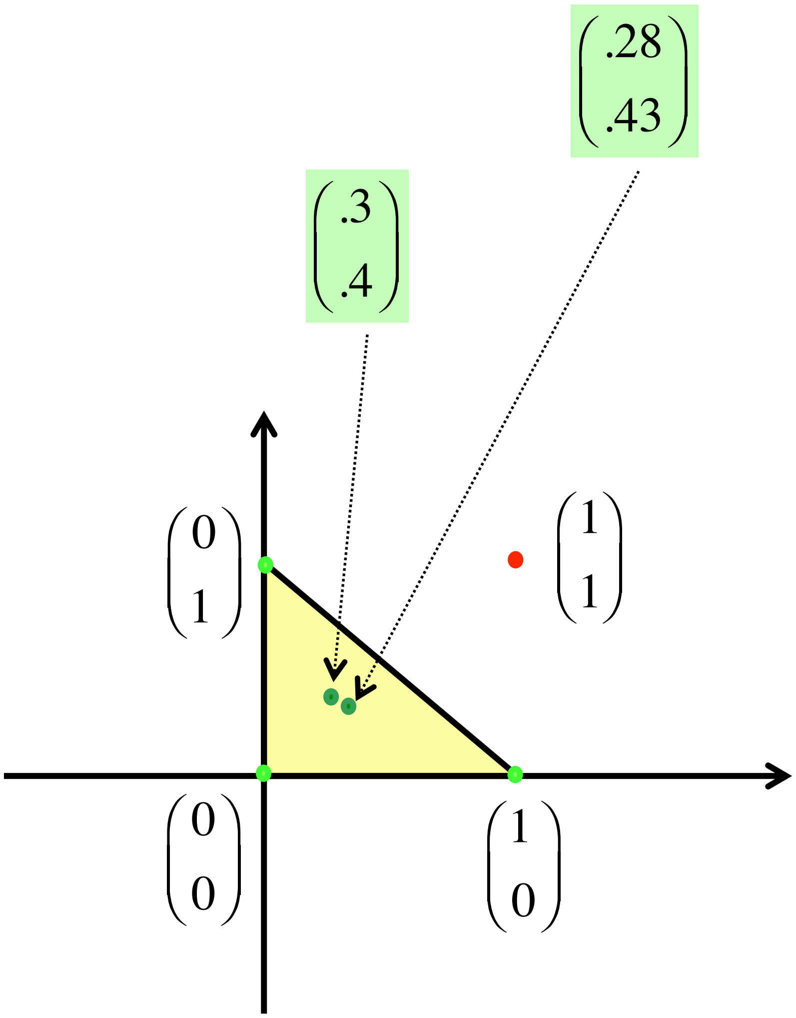

$\begin{pmatrix} .28 \\ .43 \end{pmatrix}$
$\begin{pmatrix} .3 \\ .4 \end{pmatrix}$
$\begin{pmatrix} 0 \\ 1 \end{pmatrix}$
$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$
$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$
$\begin{pmatrix} 1 \\ 0 \end{pmatrix}$

Traditional Consensus Problem

- Special case of vector consensus : $d = 1$

- Necessary & sufficient condition for complete graphs:

$$n \geq 3f + 1$$

in synchronous [Lamport,Shostak,Pease]
& asynchronous systems [Abraham,Amit,Dolev]

Results

Necessary and Sufficient Conditions (Complete Graphs)

- Exact consensus in synchronous systems

$$n \geq \max(3, d+1)\, f + 1$$

- Approximate consensus in asynchronous systems

$$n \geq (d+2)\, f + 1$$

STOC 2013

Similar results for asynchronous systems

Hammurabi Mendes & Maurice Herlihy

Talk Outline

	Necessity	Sufficiency
Synchronous	max(3,d+1) f +1	max(3,d+1) f +1
Asynchronous	(d+2) f +1	(d+2) f +1

Synchronous Systems: $n \geq \max(3, d+1) f + 1$ necessary

- $n \geq 3f + 1$ necessary due to Lamport, Shostak, Pease

Synchronous Systems:
$n \geq \max(3, d+1) f + 1$ necessary

- $n \geq 3f + 1$ necessary due to Lamport, Shostak, Pease

- Proof of $n \geq (d+1) f + 1$ by contradiction …

suppose that

$f = 1$

$n \leq (d+1)$

$n \leq d+1 = 3$ when $d = 2$

- Three fault-free processes, with inputs shown below

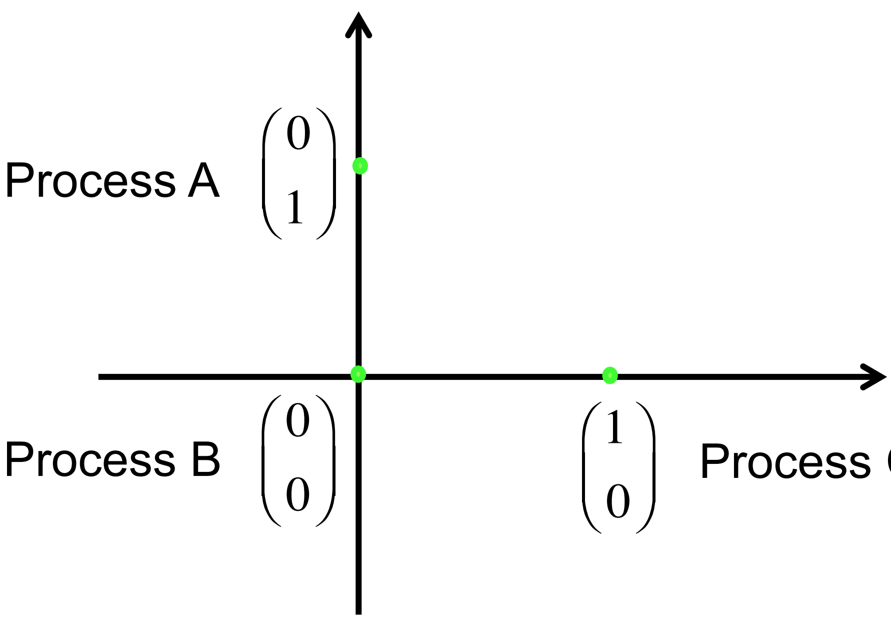

Process A's Viewpoint

- If B faulty : output on green segment (for validity)

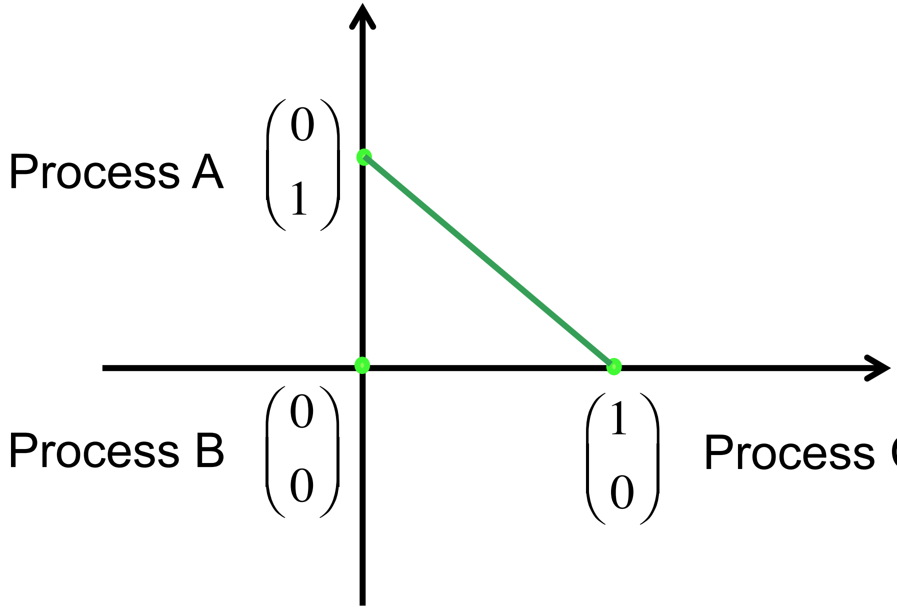

Process A's Viewpoint

- If B faulty : output on green segment (for validity)
- If C faulty : output on red segment

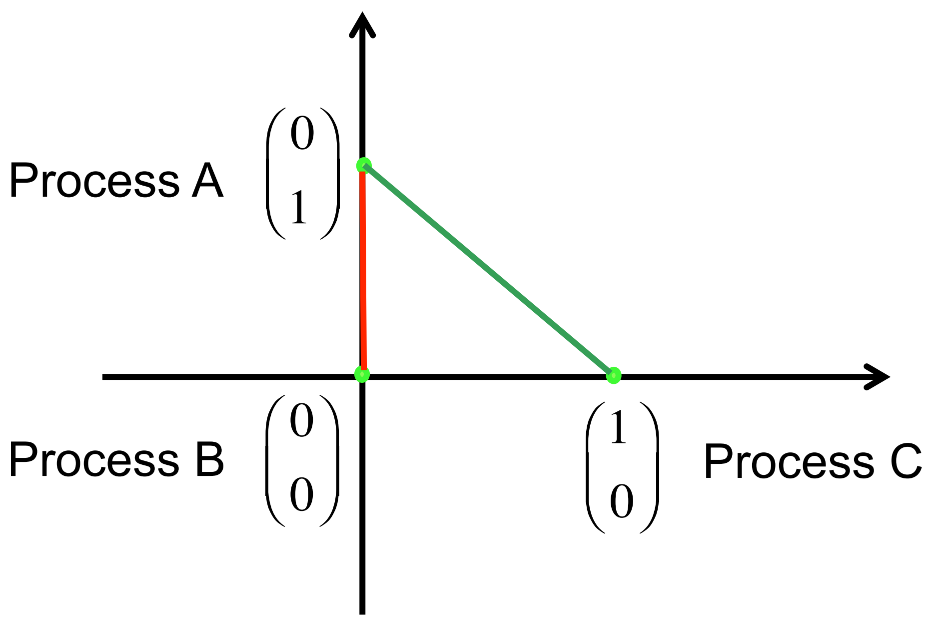

Process A's Viewpoint

- If B faulty : output on green segment (for validity)
- If C faulty : output on red segment

➔ Output must be on both segments = initial state

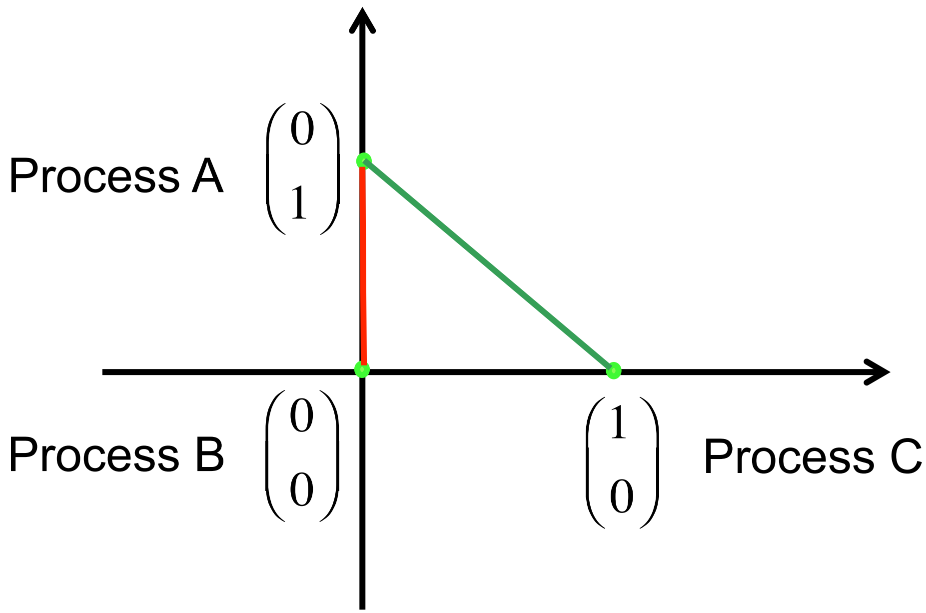

d = 2

- Validity forces each process to choose output = own input

➔ No agreement

➔ n = (d+1) insufficient when f = 1

➔ By simulation, (d+1)f insufficient

Proof generalizes to all d

Talk Outline

	Necessity	Sufficiency
Synchronous	max(3,d+1) f +1	max(3,d+1) f +1
Asynchronous	(d+2) f +1	(d+2) f +1

Synchronous System

n ≥ max(3,d+1) f +1

1. Reliably broadcast input vector to all processes
[Lamport,Shostak,Pease]

2. Receive multiset Y containing n vectors

3. Output = a deterministically chosen point in

$$\Gamma(Y) = \cap_{T \subseteq Y,\, |T|=|Y|-f} \ \mathrm{Hull}(T)$$

$d = 2, \quad f = 1, \quad n = 4$

- Y contains 4 points, one from faulty process

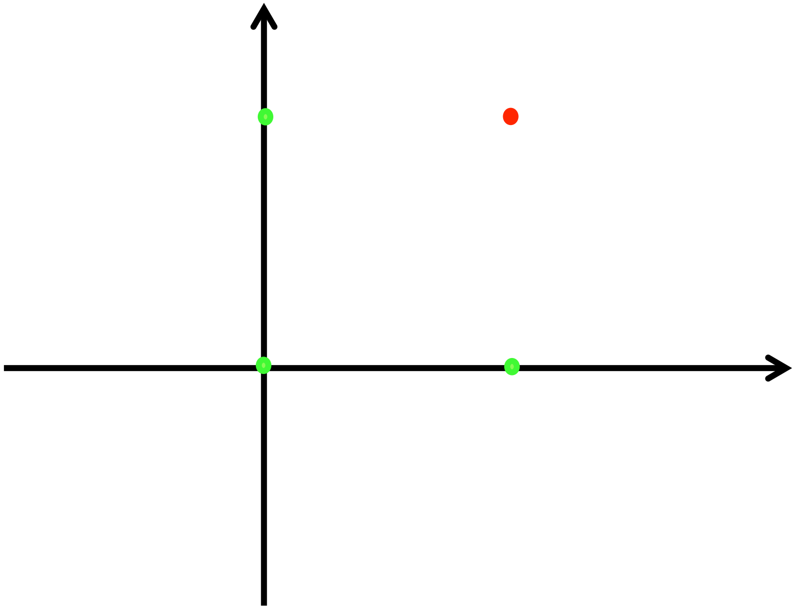

n-f = 3

- Y contains 4 points, one from faulty process
- Output in intersection of hulls of (n-f)-sets in Y

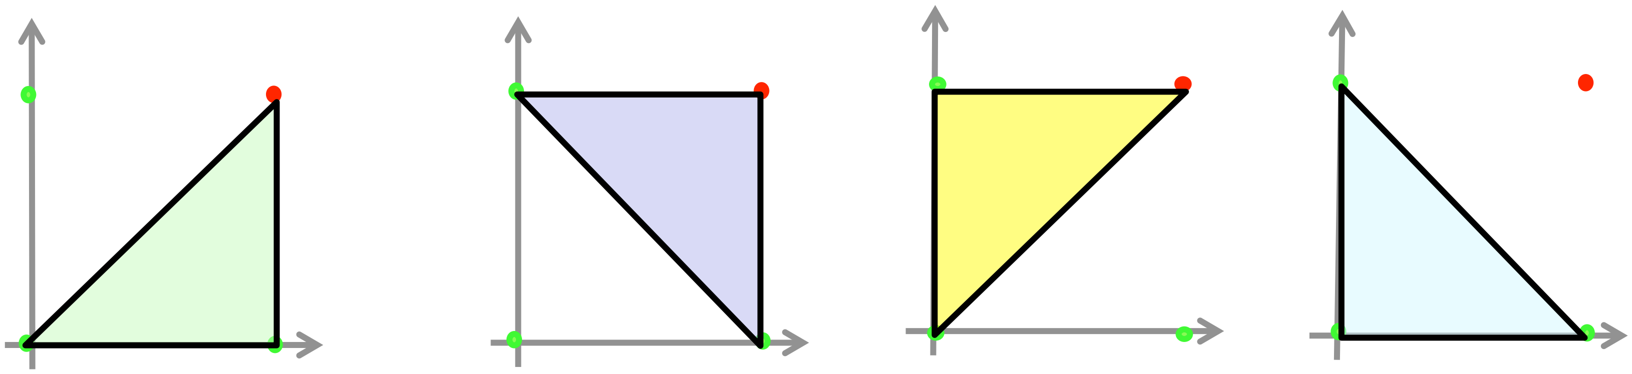

Proof of Validity

Output in $\Gamma(Y) = \cap_{T \subseteq Y,\ |T|=|Y|-f} \ \mathrm{Hull}(T)$

- Claim 1 : Intersection is non-empty
- Claim 2 : All points in intersection are in convex hull of fault-free inputs

Tverberg's Theorem

$\geq (d+1)f+1$ points can be partitioned into $(f+1)$ sets such that their convex hulls intersect

$d = 2$

$f = 2$

$n = 8$

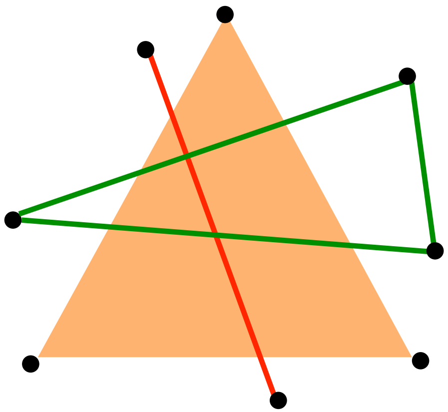

Tverberg's Theorem

≥ (d+1)f+1 points can be partitioned into (f+1) sets such that their convex hulls intersect

d = 2

f = 2

n = 8

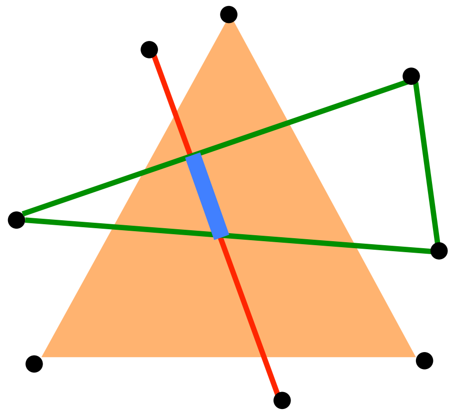

Tverberg points

Claim 1: Intersection is Non-Empty

$$\Gamma(Y) = \cap_{T \subseteq Y,\, |T|=|Y|-f} \ \mathrm{Hull}(T)$$

- Each T contains one set in Tverberg partition of Y

Claim 1: Intersection is Non-Empty

$$\Gamma(Y) = \cap_{T \subseteq Y,\, |T|=|Y|-f} \text{ Hull}(T)$$

- Each T contains one set in Tverberg partition of Y

➔ Intersection contains all Tverberg points of Y

Claim 1: Intersection is Non-Empty

$$\Gamma(Y) = \cap_{T \subseteq Y,\, |T|=|Y|-f} \ \mathrm{Hull}(T)$$

- Each T contains one set in Tverberg partition of Y

➔ Intersection contains all Tverberg points of Y

➔ Non-empty by Tverberg theorem when ≥ (d+1)f+1

Claim 2:
Intersection in Convex Hull of Fault-Free Inputs

$$\Gamma(Y) = \cap_{T \subseteq Y, \, |T| = |Y| - f} \ \mathrm{Hull}(T)$$

- At least one T contains inputs of only fault-free processes

➔ Claim 2

Talk Outline

	Necessity	Sufficiency
Synchronous	max(3,d+1) f +1	max(3,d+1) f +1
Asynchronous	(d+2) f +1	(d+2) f +1

Asynchronous System
$n \geq (d+2) f + 1$ is Necessary

- Suppose f=1, n=d+2

- One process very slow
  … remaining d+1 must terminate on their own

- d+1 processes choose output = own input
  (as in synchronous case)

Talk Outline

	Necessity	Sufficiency
Synchronous	max(3,d+1) f +1	max(3,d+1) f +1
Asynchronous	(d+2) f +1	(d+2) f +1

Asynchronous System
$n \geq (d+2) f + 1$

- Algorithm executes in asynchronous rounds
- Process i computes $v_i[t]$ in its round t
- Initialization: $v_i[0]$ = input vector

Asynchronous System
$n \geq (d+2) f + 1$

- Algorithm executes in asynchronous rounds
- Process i computes $v_i[t]$ in its round t
- Initialization: $v_i[0]$ = input vector

… 2 steps per round

Step 1 in Round t

- Reliably broadcast state $v_i[t-1]$
- Primitive from [Abraham, Amit, Dolev] ensures that

  each pair of fault-free processes receives
  (n-f) identical messages

Step 2 in Round t

- Process i receives multiset B_i of vectors in step 1

$$|B_i| \geq n-f$$

Step 2 in Round t

- Process i receives multiset B_i of vectors in step 1

$$|B_i| \geq n\text{-}f$$

- For each (n-f)-subset Y of B_i … choose a point in $\Gamma(Y)$

Step 2 in Round t

- Process i receives multiset B_i of vectors in step 1

$$|B_i| \geq n-f$$

- For each (n-f)-subset Y of B_i … choose a point in $\Gamma(Y)$

- New state $v_i[t]$ = average over these points

Validity

- $|B_i| \geq n-f$

  $n \geq (d+2) f + 1 \rightarrow n-f \geq (d+1) f + 1 \rightarrow$ Tverberg applies

- Validity proof similar to synchronous

ε-Agreement

Recall from Step 2

- For each (n-f)-subset Y of B_i … choose a point in Γ(Y)
- New state $v_i[t]$ = average over these points

ε-Agreement

Recall from Step 2

- For each (n-f)-subset Y of B_i … choose a point in Γ(Y)
- New state $v_i[t]$ = average over these points

Because i and j receive identical n-f messages in step 1, they choose at least one identical point above

ε-Agreement

Recall from Step 2

- For each (n-f)-subset Y of B_i … choose a point in Γ(Y)
- New state $v_i[t]$ = average over these points

Because i and j receive identical n-f messages in step 1, they choose at least one identical point above

$$\mathbf{v}_i[t] = \sum \alpha_k \, \mathbf{v}_k[t-1]$$

$$\mathbf{v}_j[t] = \sum \beta_k \, \mathbf{v}_k[t-1]$$

$v_i[t]$ and $v_i[t]$ as convex combination of fault-free states, with non-zero weight for an identical process

ε-Agreement

$$\mathbf{v}_i[t] \quad = \quad \sum \alpha_k \, \mathbf{v}_k[t-1]$$

$$\mathbf{v}_j[t] \quad = \quad \sum \beta_k \, \mathbf{v}_k[t-1]$$

$v_i[t]$ and $v_i[t]$ as convex combination of fault-free states, with non-zero weight for an identical process

Rest of the argument standard in convergence proofs

ε-Agreement

$$\mathbf{v}_i[t] \quad = \quad \sum \alpha_k \, \mathbf{v}_k[t-1]$$

$$\mathbf{v}_j[t] \quad = \quad \sum \beta_k \, \mathbf{v}_k[t-1]$$

$v_i[t]$ and $v_i[t]$ as convex combination of fault-free states, with non-zero weight for an identical process

Rest of the argument standard in convergence proofs

➔ Range of each vector element shrinks by a factor < 1 in each round

➔ ε-Agreement after sufficient number of rounds

Summary

- Necessary and sufficient n for vector consensus
- Synchronous & asynchronous systems

Matrix Form

$$\mathbf{v}_i[t] \quad = \quad \sum \alpha_k \, \mathbf{v}_k[t-1]$$

$$\mathbf{v}_j[t] \quad = \quad \sum \beta_k \, \mathbf{v}_k[t-1]$$

$v_i[t]$ and $v_i[t]$ as convex combination of fault-free states, with non-zero weight for an identical process

v[t] = M[t] v[t-1]

where M[t] is row stochastic with a coefficient of ergodicity < 1

Matrix Form

$$\mathbf{v}_i[t] = \sum \alpha_k \mathbf{v}_k[t-1]$$

$$\mathbf{v}_j[t] = \sum \beta_k \mathbf{v}_k[t-1]$$

$v_i[t]$ and $v_i[t]$ as convex combination of fault-free states, with non-zero weight for an identical process

v[t] = M[t] v[t-1]   where M[t] is row stochastic with a coefficient of ergodicity < 1

➔ Consensus because ΠM[t] has a limit with identical rows

Hajnal 1957
Wolfowitz 1963

Matrix Form

- Popular tool in decentralized control literature on *fault-free* iterative consensus [Tsitsiklis,Jadbabaei]
- Allows derivation of stronger results
  - Incomplete graphs
  - Time-varying graphs

Thanks!

Exact Consensus

- Agreement: Fault-free processes agree *exactly*
- Validity: Agreed value in convex hull of inputs at fault-free processes
- Termination: In finite time

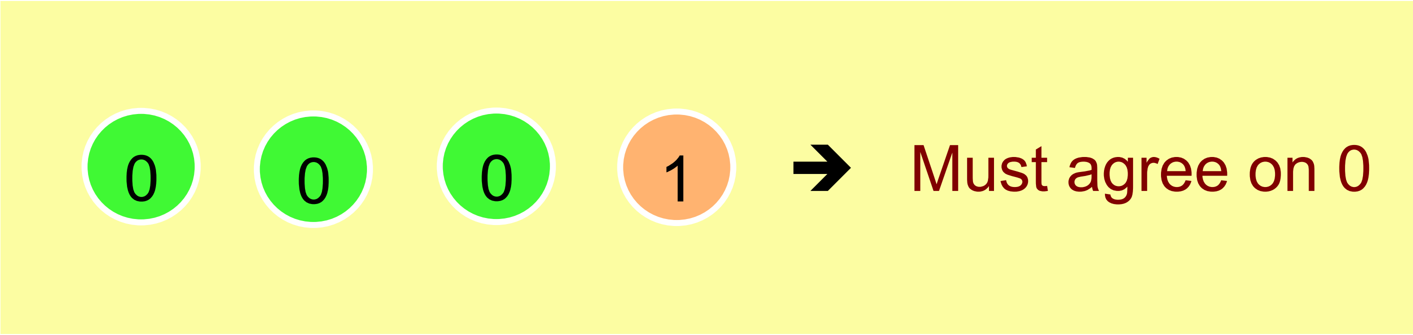

Exact Consensus

- Agreement: Fault-free processes agree *exactly*
- Validity: Agreed value in convex hull of inputs at fault-free processes
- Termination: In finite time

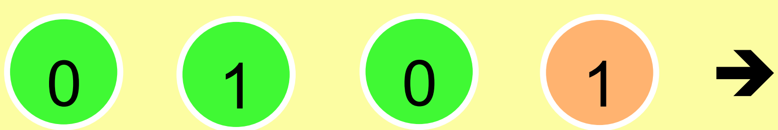

May agree on .4

Exact Consensus

Impossible with asynchrony [FLP]

Approximate Consensus

- Agreement: Fault-free processes agree *approximately*
- Validity: …
- Termination: …

Approximate Consensus

- Agreement: Fault-free processes agree *approximately*
- Validity: …
- Termination: …

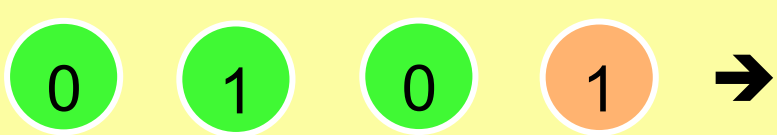

May agree on ≈ .4

Necessary & Sufficient Condition (Complete Graphs)

- $n \geq 3f+1$

Necessary & Sufficient Condition (Complete Graphs)

- $n \geq 3f+1$

for

- Exact consensus with synchrony
- Approximate consensus with asynchrony

Necessary & Sufficient Condition (Complete Graphs)

- $n \geq 3f+1$

for

- Exact consensus with synchrony
- Approximate consensus with asynchrony

with scalar inputs

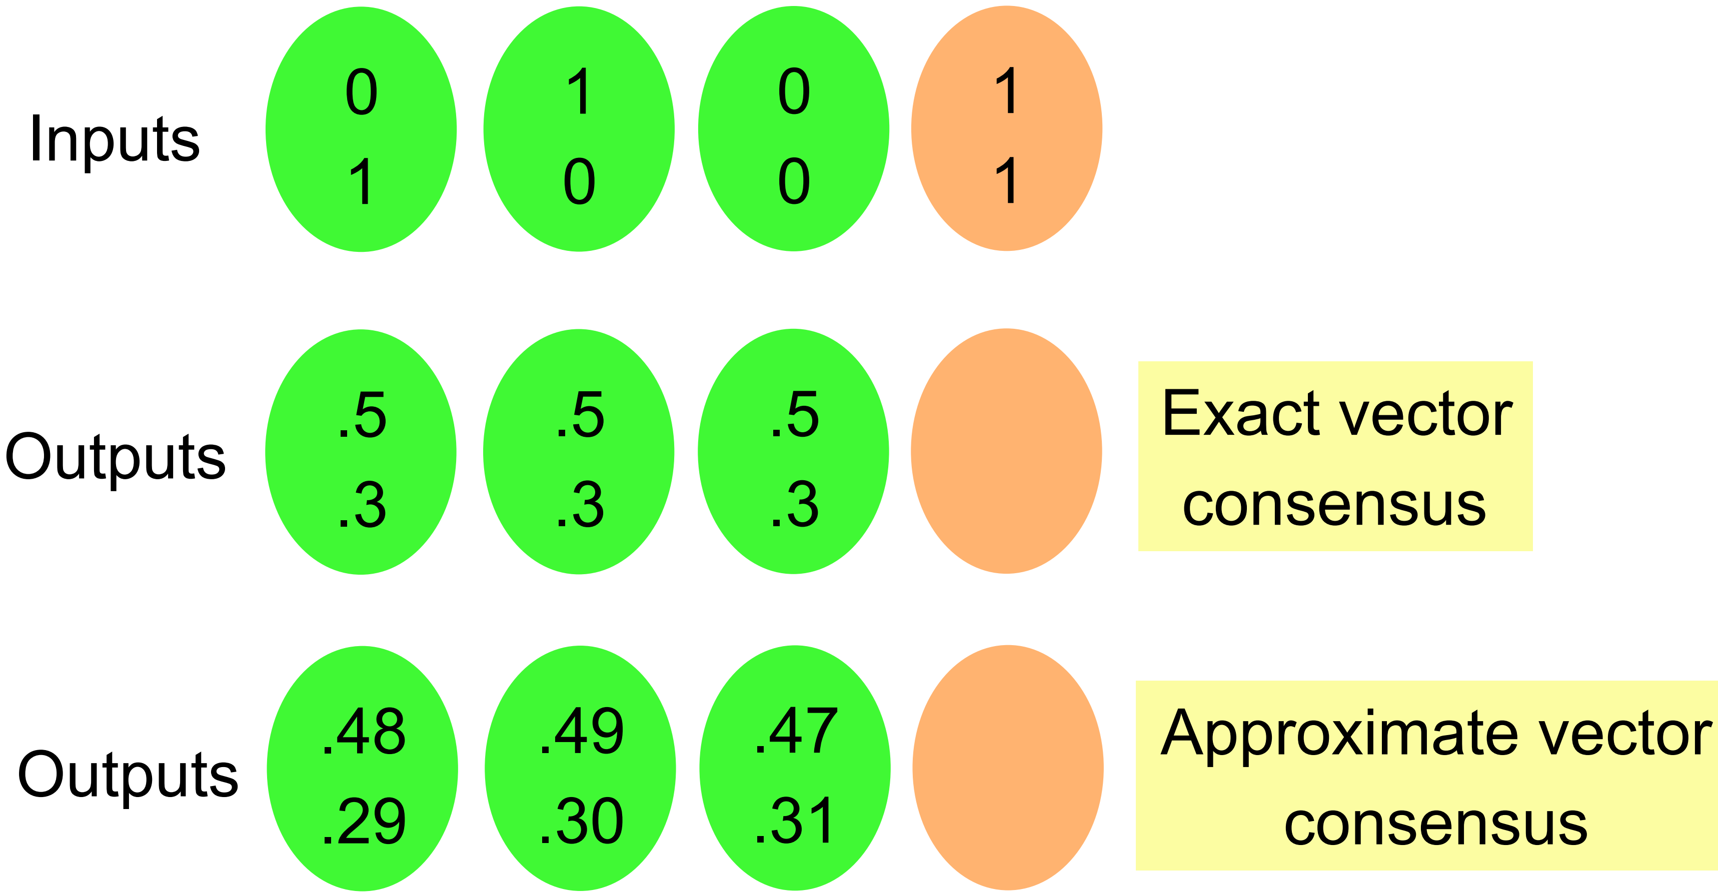

Inputs
0 1
1 0
0 0
1 1
Outputs
.5 .3
.5 .3
.5 .3
Exact vector consensus
Outputs
.48 .29
.49 .30
.47 .31
Approximate vector consensus