

The Step Complexity of Multidimensional Approximate Agreement

Hagit Attiya, Technion

Faith Ellen, University of Toronto


Approximate Agreement

$\varepsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in \mathbb{R}$.

Each nonfaulty process p_i outputs $y_i \in \mathbb{R}$ such that:

- all outputs are within distance ε of each other and
- $y_i \in [\min\{x_1, \dots, x_n\}, \max\{x_1, \dots, x_n\}]$.




Multidimensional Approximate Agreement

$\varepsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in \mathbb{R}^k$, where $k \in \mathbb{Z}^+$.

Each nonfaulty process p_i outputs $y_i \in \mathbb{R}^k$ such that:

- all outputs are within distance ε of each other and
- $y_i \in \text{convex hull of } \{x_1, \dots, x_n\}$.



There are wait-free algorithms among n processes for approximate agreement using only registers:

- [Moran, 1995] $O(n \log(S/\epsilon))$ step complexity using single-writer registers

S = spread of the inputs


$$= \max \{ |x_i - x_j| : 1 \leq i < j \leq n \}.$$

- [Schenk, 1995] $O(\log(M/\epsilon))$ step complexity using multi-writer registers

M = magnitude of the inputs

$$= \max \{ |x_i| : 1 \leq i \leq n \}.$$

$S \leq 2M$ but M can be significantly larger than S



Step complexity lower bounds for approximate agreement using single-writer registers:

- $\Omega(n)$
- [Herlihy, 1991] $\frac{1}{2} \log_3(S/\epsilon)$

OUR RESULTS

Step complexity lower bounds for approximate agreement using multi-writer registers:

- $\Omega\left(\min\left\{\frac{\log M}{\log \log M}, \frac{\sqrt{\log n}}{\log \log n}\right\}\right)$
- $\frac{1}{2} \log_{1+\sqrt{2}}(S/\varepsilon)$

A wait-free algorithm for multi-dimensional approximate agreement using multi-writer registers

- $O(\log n (\log n + \log(S/\epsilon)))$ step complexity

S = spread of the inputs

$$= \max \{ \|x_i - x_j\| : 1 \leq i < j \leq n \}.$$

A wait-free algorithm for approximate agreement using multi-writer registers

- $O(\min\{\log n (\log n + \log(S/\epsilon)), \log(M/\epsilon)\})$ step complexity

M = magnitude of the inputs

$$= \max \{ |x_i| : 1 \leq i \leq n \}.$$

Outline

- half
 - r-agree
- } Schenk
- ApproxAgreeDU
 - TwoGroupApproxAgree
 - Multidimensional Approximate Agreement

half

Each process p_i has an input $x_i \in [a,b]$.

Each nonfaulty process p_i outputs $y_i \in [a,b]$ such that:

- all outputs are within distance $\varepsilon = (b-a)/2$ of each other and
- $y_i \in [\min\{x_1, \dots, x_n\} , \max\{x_1, \dots, x_n\}]$.

step complexity: $O(1)$

half(a,b,x)

uses two single-bit registers r_0 and r_1 , both initially 0

if $x < (a+b)/2$

then $k \leftarrow 0$

else $k \leftarrow 1$

$r_k \leftarrow$ write 1

if $r_{1-k} = 0$

then return x

else return $(a+b)/2$

half(a,b,x)

if $x < (a+b)/2$

then $k \leftarrow 0$

else $k \leftarrow 1$

$r_k \leftarrow$ write 1

if $r_{1-k} = 0$

then return x

else return $(a+b)/2$

Consider the first register r_i that is written to.

All processes that terminate, but don't write to r_i

output $(a+b)/2$.

approximate agreement on $[0,1]$

$\varepsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in [0,1]$.


Each nonfaulty process p_i outputs $y_i \in [0,1]$ such that:

- all outputs are within distance ε of each other and
- $y_i \in [\min\{x_1, \dots, x_n\} , \max\{x_1, \dots, x_n\}]$.

step complexity: $O(\log(1/\varepsilon))$


r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels




r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels




r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels






r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels




r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels



r-agree(x, ϵ)

uses a tree of instances of half with $\lceil \log_2(1/\epsilon) \rceil$ levels



ϵ -unknown approximate agreement on $[0,1]$

Each process p_i has an input $x_i \in [0,1]$ and an accuracy parameter $\epsilon_i > 0$.

Each nonfaulty process p_i outputs $y_i \in [0,1]$ such that:

- all outputs are within distance $\max\{\epsilon_1, \dots, \epsilon_n\}$ of each other and
- $y_i \in [\min\{x_1, \dots, x_n\} , \max\{x_1, \dots, x_n\}]$.

step complexity: $O(\log \max\{1/\epsilon_1, \dots, 1/\epsilon_n\})$

Approximate Agreement with Domain Uncertainty

$\epsilon > 0$ is an accuracy parameter known to all processes.


Each process p_i has an input $x_i = u_i + t_i (v_i - u_i) \in \mathbb{R}^k$,

between two points $u_i, v_i \in \mathbb{R}^k$, where $t_i \in [0, 1]$.

Precondition: $\|u_i - u_j\| \leq \epsilon/5$ and $\|v_i - v_j\| \leq \epsilon/5$ for all p_i and p_j .

Each nonfaulty process p_i outputs $y_i \in \mathbb{R}^k$ such that:

- all outputs are within distance ϵ of each other and
- $y_i = u_i + t'_i (v_i - u_i)$ is between u_i and v_i .



Approximate Agreement with Domain Uncertainty

$\epsilon > 0$ is an accuracy parameter known to all processes.


Each process p_i has an input $x_i = u_i + t_i (v_i - u_i) \in \mathbb{R}^k$,

between two points $u_i, v_i \in \mathbb{R}^k$, where $t_i \in [0, 1]$.

Precondition: $\|u_i - u_j\| \leq \epsilon/5$ and $\|v_i - v_j\| \leq \epsilon/5$ for all p_i and p_j .

Each nonfaulty process p_i outputs $y_i \in \mathbb{R}^k$ such that:

- all outputs are within distance ϵ of each other and
- $y_i = u_i + t'_i (v_i - u_i)$ is between u_i and v_i .



Approximate Agreement with Domain Uncertainty

ApproxAgreeDU(u, v, t, ϵ)

$s \leftarrow \|u - v\|$

if $s \leq 2\epsilon/5$ then return $u + t(v-u)$

$\epsilon' \leftarrow \epsilon/5s$

$t' \leftarrow \text{r-agree}(t, \epsilon')$

return $u + t'(v-u)$



step complexity: $O(\log \max\{5s_1/\epsilon, \dots, 5s_n/\epsilon\})$

Two Group Multidimensional Approximate Agreement

$\varepsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in \mathbb{R}^k$.

Preconditions: Processes are divided into two groups **0** and **1** and $\|x_i - x_j\| \leq \varepsilon/5$ for all p_i and p_j in the same group.



Two Group Multidimensional Approximate Agreement


$\epsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in \mathbb{R}^k$.

Preconditions: Processes are divided into two groups **0** and **1** and $\|x_i - x_j\| \leq \epsilon/5$ for all p_i and p_j in the same group.

Each nonfaulty process p_i outputs $y_i \in \mathbb{R}^k$ such that:

- all outputs are within distance ϵ of each other and
- $y_i \in$ convex hull of $\{x_1, \dots, x_n\}$.



TwoGroupApproxAgree(g, x, ϵ)

$a[g] \leftarrow x$

$A[g] \leftarrow \text{write}(a[g])$

$a[1-g] \leftarrow \text{read}(A[1-g])$

if $a[1-g] = \perp$

then $B[g] \leftarrow \text{write}(1)$

$a[1-g] \leftarrow \text{read}(A[1-g])$

 if $a[1-g] = \perp$

 then return x

 else return $\text{ApproxAgreeDU}(a[0], a[1], g, \epsilon, \epsilon/5)$

else $b \leftarrow \text{read}(B[1-g])$

 if $b = 0$

 then return $\text{ApproxAgreeDU}(a[0], a[1], g, \epsilon, \epsilon/5)$

 else return $\text{ApproxAgreeDU}(a[0], a[1], 1-g, \epsilon, \epsilon/5)$

step complexity: $O(\log(S/\epsilon))$

Multidimensional Approximate Agreement

$\varepsilon > 0$ is an accuracy parameter known to all processes.

Each process p_i has an input $x_i \in \mathbb{R}^k$.


Each nonfaulty process p_i outputs $y_i \in \mathbb{R}^k$ such that:


- all outputs are within distance ε of each other and
- $y_i \in$ convex hull of $\{x_1, \dots, x_n\}$.

Use a binary tree of height $\lceil \log_2 n \rceil$, with one leaf for each process, and a separate instance of **TwoGroupApproxAgree** at every other node.

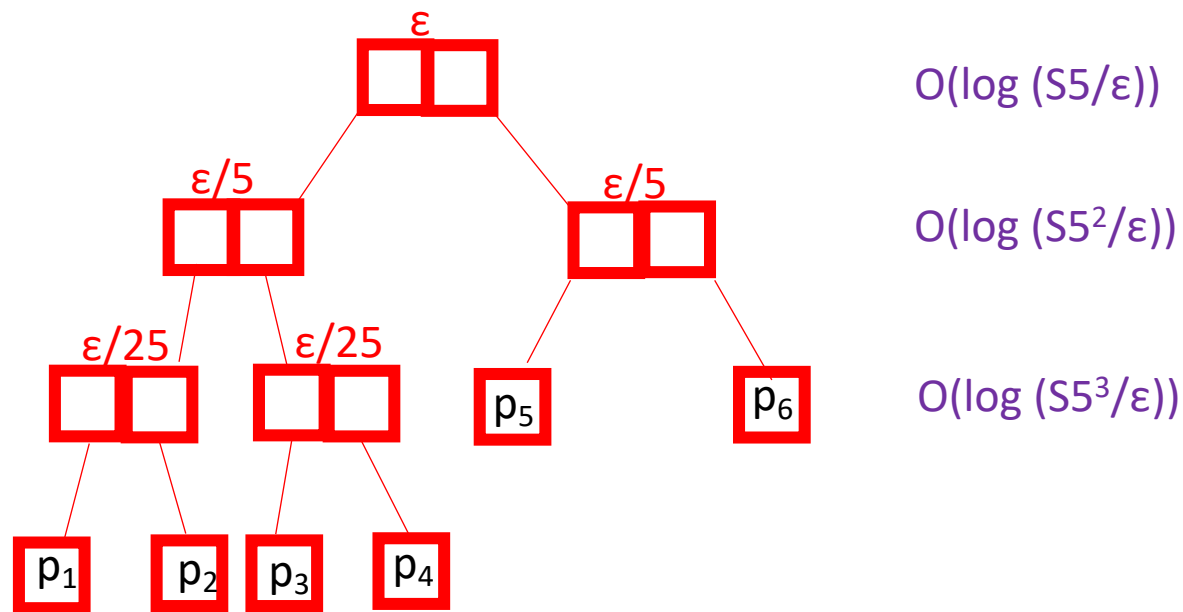
Use accuracy parameter $\varepsilon/5^d$ at internal nodes of depth d .

Each process p_i traverses a path from its leaf to the root, using x_i as its input to the first instance and, for each subsequent instance, using the output from the previous instance as its input.





Lemma For $0 \leq d < \lceil \log_2 n \rceil$, at each node of depth d , the inputs of the processes in each group are within distance $\epsilon/5^{d+1}$ of one another.



Lemma For $0 \leq d < \lceil \log_2 n \rceil$, at each node of depth d , the inputs of the processes in each group are within distance $\epsilon/5^{d+1}$ of one another.

Theorem Multidimensional approximate agreement among n processes whose inputs are within distance S of one another can be solved with step complexity $O(\log n (\log n + \log(S/\epsilon)))$.

A wait-free algorithm for multi-dimensional approximate agreement using multi-writer registers

- $O(\log n (\log n + \log(S/\epsilon)))$ step complexity

S = spread of the inputs

$$= \max \{ \|x_i - x_j\| : 1 \leq i < j \leq n \}.$$

A wait-free algorithm for approximate agreement using multi-writer registers

- $O(\min\{\log n (\log n + \log(S/\epsilon)), \log(M/\epsilon)\})$ step complexity

M = magnitude of the inputs

$$= \max \{ |x_i| : 1 \leq i \leq n \}.$$

Open Problems

- Close the gap for approximate agreement between the upper bound $O(\min\{\log n (\log n + \log(S/\varepsilon)), \log(M/\varepsilon)\})$ and the lower bounds

$$\Omega\left(\min\left\{\frac{\log M}{\log \log M}, \frac{\sqrt{\log n}}{\log \log n}\right\}\right) \text{ and}$$

$$\frac{1}{2} \log_{1+\sqrt{2}}(S/\varepsilon).$$

- Extend the algorithms for multidimensional approximate agreement to asynchronous message-passing systems.