Slide 1

COMP3152/9152
Lecture 9
Simultaneous Byzantine Agreement
Ron van der Meyden

Slide 2

Simultaneous Byzantine Agreement

Suppose there are n generals, t of them are traitors, the rest are loyal. But initially, nobody knows who the traitors are. There are no broadcast actions, only message passing. Every general has a preference about whether to attack.

Can we design a protocol so that

1. At some point, all the loyal generals either attack, or they all retreat.
2. If all the generals prefer to atack, then the agreement is to attack.

Even though the traitors may misbehave (e.g., tell one general they want to attack, and another that they want to retreat.)

Motivation: fault-tolerant protocols

Slide 3

(γ, π) is a ba-compatible interpreted context if

1. Each agent i has an action $\texttt{decide}_i(y)$ for $y \in 0, 1$
2. The environments actions include actions $(a_{e1}, \ldots, a_{en})$, where the a_{ei} are tuples that describe
   (a) which messages sent by process j are delivered to process i in that round
   (b) whether or not i fails in that round ($\texttt{fail}_i$), and the nature of the failure
   (c) γ is a recording context
3. Process i's initial state is a tuple of the form $(x_i, \ldots)$, where x_i is i's preference for the decision.
4. The environment's initial state also contains x_i

Slide 4

5. There is a proposition $\mathit{decided}_i(y)$ for $y \in \{0, 1\}$ that is true if i tried to perform $\texttt{decide}_i(y)$ at some previous round
6. There is a proposition $\exists y$ for $y \in \{0, 1\}$ that is true if some process i has $x_i = y$.

Slide 5

Notation:

$deciding_i(y)$ for $\neg decided_i(y) \wedge \bigcirc\, decided_i(y)$

At a point (r, m), let $\mathcal{N}(r, m)$ be the set of nonfaulty agents (for which the enviornment has not yet performed $\mathtt{fail}_i$).

$(\mathcal{I}, r, m) \models deciding_{\mathcal{N}}(y)$ if $(\mathcal{I}, r, m) \models deciding_i(y)$ for all $i \in \mathcal{N}(r, m)$

Slide 6

Specification for SBA

A system $\mathcal{I}$ satisfies the SBA specification if for every run r:

1. *Decision:* Every process that is nonfaulty in r performs exactly one $\mathtt{decide}_i(y)$ action in r
2. *Agreement:* If i is nonfaulty at (r, m) and is about to decide y at (r, m) and j is nonfaulty at (r, m') and is about to decide y' at (r, m) then $y = y'$.
3. *Validity:* If all the processes have the same initial preference x then all the nonfaulty processes decide x
4. *Simultaneity:* the nonfaulty processes decide simultaneously, i.e., if i and j are nonfaulty at (r, m) and i is about to decide at (r, m), then so is j.

Slide 7

Failure Modes

The following are possible failure modes:

1. *Crash Failures:* A faulty process follows its protocol up to the time when it fails (sending a subset of messages) after which it sends no messages.
2. *Omission Failures:* A faulty process follows its protocol, but in any round the set of messages it sends or receives is a subset of what it should be.
3. *Byzantine Failures:* faulty processes may deviate from the protocol in any way: send a subset of messages, send false messages, collude with other faulty processes to deceive the non-faulty processes, etc

Slide 8

Knowledge of preferences is not enough (n=3, t=1)

	1	2	3
x_1	0	0	0
x_2	0	0	0
x_3	0	0	0

$\sim_1$

	1	2	3
x_1	0	X	0
x_2	0	X	$*$
x_3	0	X	0

$\sim_3$

	1	2	3
x_1	0	X	0
x_2	1	X	$*$
x_3	0	X	0

$\sim_1$

	1	2	3
x_1	0	0	0
x_2	1	1	1
x_3	0	0	0

$\sim_{\{1,2,3\}} \ldots\ldots \sim_{\{1,2,3\}}$

	1	2	3
x_1	1	1	1
x_2	1	1	1
x_3	1	1	1

Slide 9

Non-rigid sets of agents

We have defined common knowledge $C_G\phi$, and distributed knowledge $D_G\phi$ with respect to a *fixed* set G of agents.

In SBA, we need to consider sets of agents that depend on the point.

In particular $\mathcal{N}(r, m)$, the set of agents that have not failed at point (r, m).

Slide 10

An attempt to define non-rigid group knowledge

Let $\mathcal{S} : \mathit{Points}(\mathcal{I}) \to \mathit{Agents}$ be a non-rigid set of agents.

$E_\mathcal{S}\phi = \bigwedge_{i\in\mathcal{S}} K_i\phi$

$C_\mathcal{S}\phi = E_\mathcal{S}\phi \wedge E_\mathcal{S}E_\mathcal{S}\phi \wedge \ldots$

Problem: In general, we can have $(\mathcal{I}, r, m) \models i \in \mathcal{S} \wedge \neg K_i(i \in \mathcal{S})$

Slide 11

A definition that works better

Define

$B_i^\mathcal{S}\phi$ as $K_i(i \in \mathcal{S} \Rightarrow \phi)$

$E_\mathcal{S}\phi$ as $\bigwedge_{i\in\mathcal{S}} B_i^\mathcal{S}\phi$

$C_\mathcal{S}\phi = E_\mathcal{S}\phi \wedge E_\mathcal{S}E_\mathcal{S}\phi \wedge \ldots$

$(\mathcal{I}, r, m) \models D_\mathcal{S}\phi$ if $(\mathcal{I}, r, m) \models D_G(G \subseteq \mathcal{S} \Rightarrow \phi)$ for $G = \mathcal{S}(r, m)$

Remark: these definitions make the following valid:

$$i \in \mathcal{S} \Rightarrow B_i^\mathcal{S}(C_\mathcal{S}(i \in \mathcal{S}))$$

Slide 12

Relating $C_\mathcal{S}$ to reachability

Define (r', m') to be $\mathcal{S}$-reachable from (r, m) if there exists a sequence of points $(r_0, m_0) \ldots (r_k, m_k)$ such that

1. $(r, m) = (r_0, m_0)$ and $(r', m') = (r_k, m_k)$ and
2. for all $l = 0 \ldots k-1$, there exists $i \in \mathcal{S}(r_l, m_l) \cap \mathcal{S}(r_{l+1}, m_{l+1})$ such that $(r_l, m_l) \sim_i (r_{l+1}, m_{l+1})$

Lemma: $(\mathcal{I}, r, m) \models C_\mathcal{S}\phi$ iff $(\mathcal{I}, r', m') \models \phi$ for all (r', m') that are $\mathcal{S}$-reachable from (r, m).

Slide 13

A Property of Byzantine Agreement Solutions

Theorem: Let (γ, π) be a ba-compatible interpreted context and let P be a deterministic protocol. If $\mathcal{I} = \mathcal{I}^{rep}(P, \gamma, \pi)$ satisfies the SBA specification, then

$$\mathcal{I} \models \mathit{deciding}_{\mathcal{N}}(y) \Rightarrow B_i^{\mathcal{N}}(C_{\mathcal{N}}(\mathit{deciding}_{\mathcal{N}}(y)))$$

Corollary: Let (γ, π) be a ba-compatible interpreted context and let P be a deterministic protocol. If $\mathcal{I} = \mathcal{I}^{rep}(P, \gamma, \pi)$ satisfies the SBA specification, then

$$\mathcal{I} \models \mathit{deciding}_{\mathcal{N}}(y) \Rightarrow B_i^{\mathcal{N}}(C_{\mathcal{N}}(\exists y))$$

Slide 14

A knowledge-based program for SBA

Case of

  if $\neg \mathit{decided}_i \wedge B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 0)$ **do** $\texttt{decide}_i(0)$

  if $\neg \mathit{decided}_i \wedge \neg B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 0) \wedge B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 1)$ **do** $\texttt{decide}_i(1)$

  $\vdots$(communication actions)

end case

Slide 15

Assumptions on the faulty processes

Let

1. n be the number of processes
2. t be the maximum number of faulty processes in any run ($t \leq n$)
3. γ^{cr} be the ba-context for crash failures (only)
4. γ^{som} be the ba-context for sending omission failures
5. γ^{gom} be the ba-context for general (send and receive) omission failures
6. $\Gamma^{SBA} = \{\gamma^{cr}, \gamma^{som}, \gamma^{gom}\}$

Slide 16

Solvability with up to t failures

Theorem: There are deterministic protocols that attain SBA in $t+1$ rounds in each of the contexts in Γ^{SBA}.

Theorem: If P is a deterministic protocol that satisfies the SBA specification in a context $\gamma \in \Gamma^{SBA}$, r is a failure free run in $\mathcal{R}^{rep}(P, \gamma)$ and P attains SBA in t' rounds in run r, then $t' \geq t+1$.

Slide 17

Comparing the rate at which different protocols reach agreeement

Let P and P' be two protocols for a context $\gamma \in \Gamma^{SBA}$.

Say that runs $r \in \mathcal{R}^{rep}(P, \gamma)$ and $r' \in \mathcal{R}^{rep}(P', \gamma)$ are *corresponding* if

1. $r(0) = r'(0)$
2. for all rounds m, the environment performs the same actions (e.g. fail a process, block delivery/transmission of a message) in round m of r as in round m of r'.

Slide 18

P *dominates* P' if for every run $r \in \mathcal{R}^{rep}(P, \gamma)$ and corresponding run $r' \in \mathcal{R}^{rep}(P', \gamma)$, if the nonfaulty processes decide in round m of r then the nonfaulty processes decide in round m or later in r'.

P *strictly dominates* P' if P dominates P' and there exists a run of P where the nonfaulty processes decide strictly earlier than in the corresponding run of P'.

P is *optimal* for SBA in context γ if it is not strictly dominated by any other protocol for SBA in this context.

P is *optimum* for SBA in context γ if it dominates all protocols for SBA in this context.

Slide 19

The full-information protocol

The *full-information protocol FIP* is the joint protocol $(FIP_1, \ldots, FIP_n)$ defined by the following rule:

> In round m, FIP_i sends a message containing *all* of agent i's local state to all other agents.

Slide 20

Basic formulas

Say that a formula ϕ is *determined by the initial state* in a system $\mathcal{I}$ if for every point (r, m) of $\mathcal{I}$, we have $(\mathcal{I}, r, m) \models \phi$ iff $(\mathcal{I}, r, 0) \models \phi$.

A formula is *basic* if it is of the form $K_i\psi$, $D_{\mathcal{N}}\psi$, $C_{\mathcal{N}}\psi$ or $B_i^{\mathcal{N}}\psi$ where ψ is determined by the initial state.

lide 21

FIP is an optimum protocol

Theorem: Assume that $\gamma \in \Gamma^{SBA}$ and that ϕ is a basic formula. Also assume that P is a deterministic protocol, $\mathcal{I} = \mathcal{I}^{rep}(P, \gamma, \pi^{sba})$ and $\mathcal{I}^{FIP} = \mathcal{I}^{rep}(FIP, \gamma, \pi^{sba})$. Let $r \in \mathcal{R}^{rep}(P, \gamma)$ and corresponding run $r^{FIP} \in \mathcal{R}^{rep}(FIP, \gamma)$ be corresponding runs. Then for all $m \geq 0$, if $(\mathcal{I}, r, m) \models \phi$ then $(\mathcal{I}^{FIP}, r^{FIP}, m) \models \phi$.

lide 22

Corollary: If $\gamma \in \Gamma^{SBA}$ and $\mathcal{I}^{FIP} = \mathcal{I}^{rep}(FIP, \gamma, \pi^{sba})$ then for all runs r of $\mathcal{I}^{FIP}$, there is a time $m \leq t+1$ such that $(\mathcal{I}^{FIP}, r, m) \models C_{\mathcal{N}}(\exists 0) \vee C_{\mathcal{N}}(\exists 1)$.

Slide 23

Clean Rounds

Consider γ^{cr} (Crash failures) and the full information protocol *FIP*.

When does initial information become common knowledge?

Define the proposition $\mathit{faulty}(i)$, for i an agent, to hold at (r, m) if i has failed in some round $m' \leq m$ in r

Say that round m is *clean* if for every process i, if $(\mathcal{I}, r, m) \models D_{\mathcal{N}}(\mathit{faulty}(i))$ then $(\mathcal{I}, r, m-1) \models D_{\mathcal{N}}(\mathit{faulty}(i))$.

(i.e., no new faults discovered by nonfaulty processes)

Define proposition *clean* to hold at (r, m) if some round $m' \leq m$ is clean.

Slide 24

let $\mathcal{I}^{cr} = \mathcal{I}(FIP, \gamma^{cr}, \pi^{sba})$

Theorem: If ϕ is determined by the initial state, then $\mathcal{I}(FIP, \gamma^{cr}, \pi^{sba}) \models C_{\mathcal{N}}(\mathit{clean}) \wedge D_{\mathcal{N}}\phi \Rightarrow C_{\mathcal{N}}\phi$.

If there are at most t failures then one of the first $t+1$ rounds of every run must be clean, and at all points $D_{\mathcal{N}}(\exists 0)$ or $D_{\mathcal{N}}(\exists 1)$, so

Corollary: Let r be a run of $\mathcal{I}^{cr}$. Then $(\mathcal{I}^{cr}, r, t+1) \models C_{\mathcal{N}}(\exists 0) \vee C_{\mathcal{N}}(\exists 1)$.

So the condition for making a decision is always attained by time $t+1$.

lide 25

Exact timing of decision point

Let

$\#Failed$ be the number of processes that have failed

$\#KnownFailed(r,m) = max\{k \mid (\mathcal{I}^{cr}, r, m) \models D_{\mathcal{N}}(\#Failed \geq k)\}$

$diff(r,m) = \#KnownFailed(r,m) - m$

$\mathcal{W}(r) = \max_{m \geq 0} diff(r,m)$

lide 26

Then

Theorem: Let r be a run of $\mathcal{I}^{cr}$ and let $T = min(t, n-2)$. If ϕ is determined by the initial state, then
$(\mathcal{I}^{cr}, r, T+1-\mathcal{W}(r)) \models D_{\mathcal{N}}(\phi) \Rightarrow C_{\mathcal{N}}(\phi)$.

Theorem: Let r be a run of $\mathcal{I}^{cr}$ and let $T = min(t, n-2)$. If ϕ is determined by the initial state, and $m < T+1-\mathcal{W}(r)$, then
$(\mathcal{I}^{cr}, r, m) \models D_{\mathcal{N}}(\phi) \Rightarrow C_{\mathcal{N}}(\phi)$ iff $(\mathcal{I}^{cr}, r, 0) \models D_{\mathcal{N}}(\phi) \Rightarrow C_{\mathcal{N}}(\phi)$.

Since $\exists 0/\exists 1$ are not initially common knowledge, whichever is true in r becomes common knowledge to nonfaulty processes in r after k rounds if the waste is $t+1-k$.

Slide 27

Optimizing the Full-Information Protocol

The full-information protocol is wasteful in the amount of information it sends.

Suppose initial states consist of one bit and all messages are delivered. Then the size $size(k)$ of local states of the agents after round k is defined by

$s(0) = 1$

$s(k+1) = s(k) + (n-1) \cdot s(k-1)$

So $s(k) \leq (n-1)^{k-1}$.

Local states grow exponentially in size.

Slide 28

We can represent a run of the full-information protocol up to a point (r,m) by a labelled graph $G(r,m)$:

1. vertices are pairs (i,k) with i an agent and $0 \leq k \leq m$
2. $(i,0)$ is labelled with agent i's initial state
3. there is an edge from (i,k) to $(j,k+1)$ labelled $+$ if the message from i to j in round $k+1$ was delivered
4. there is an edge from (i,k) to $(j,k+1)$ labelled $-$ if the message from i to j in round $k+1$ was delivered

lide 29

We can represent agent i's state of information at (r, m) as a subgraph $G(r_i(m))$ of $G(r, m)$ in which

1. some edges are missing (if i does not know whether a message was delivered)
2. some initial states are

$G(r, m)$ has size $O(mn^2)$, so for the first n rounds, size $O(n^3)$.

Let FIP' be the protocol in which, instead of sending its complete local state, agent i sends $G(r_i(m))$.

Then the first n rounds, the local states of i have size $O(n \cdot n \cdot n^3) = O(n^5)$.

lide 30

Theorem: Assume that $\gamma \in \Gamma^{sba}$, and let ϕ be a basic formula. Let r and r' be corresponding runs of $\mathcal{I} = \mathcal{I}^{rep}(FIP, \gamma, \pi^{sba})$ and $\mathcal{I}' = \mathcal{I}^{rep}(FIP', \gamma, \pi^{sba})$, respectively. Then $(\mathcal{I}, r, m) \models \phi$ iff $(\mathcal{I}', r', m) \models \phi$.

Slide 31

Theorem: There is an algorithm that, given input $r_i(m)$, with $m < n$, decides in time polynomial in n whether $(\mathcal{I}^{rep}(FIP', \gamma, \pi^{sba}), r, m) \models B_i C_{\mathcal{N}}(\exists y)$, when γ is the ba-context for either crash-failures or sending omission failures.

For General Omission failures, the problem is NP-hard (hence not in polynomial time).

Slide 32

An efficient knowledge-based program for SBA

SBA=

Case of

  if $\neg decided_i \wedge B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 0)$ **do** $\texttt{decide}_i(0)$;

  if $\neg decided_i \wedge \neg B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 0) \wedge B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 1)$ **do** $\texttt{decide}_i(1)$;

  if $\neg decided_i \wedge \neg B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 0) \wedge \neg B_i^{\mathcal{N}} C_{\mathcal{N}}(\exists 1)$

   do $\texttt{send}_i(G(local\ state))$

end case

lide 33

This ensures satisfaction of SBA...

Theorem: If (γ, π) is a ba-compatible interpreted context, $\mathcal{I}$ is consistent with the knowledge-based program SBA in (γ, π), and $C_{\mathcal{N}}(\exists 0) \vee C_{\mathcal{N}}(\exists 1)$ is attained in every run of $\mathcal{I}$, then $\mathcal{I}$ satisfies the SBA specification. Moreover, the processes decide in a run r if $\mathcal{I}$ at the round following the first time that $C_{\mathcal{N}}(\exists 0) \vee C_{\mathcal{N}}(\exists 1)$ is attained.

lide 34

Theorem: for $\gamma \in \Gamma^{sba}$, there is a unique interpreted system representing SBA in (γ, π^{sba}). Moreover, the corresponding protocol is an optimum protocol for SBA in γ.

Corollary: There are polynomial time optimum protocols for SBA in the crash failure and sending omission failure contexts.

(If P≠NP then no polynomial-time protocol can be optimum for SBA in the general omission failure model.)