



Consensus with Byzantine failures and asynchrony: the Ben-Or's algorithm, revisited

Distributed Algorithms

Recap

Positive Result: *There exists a deterministic synchronous protocol that solves consensus, while tolerating crash failures*

FLP Theorem: No deterministic protocol can solve consensus, while tolerating **1 crash** and **asynchrony**



Today

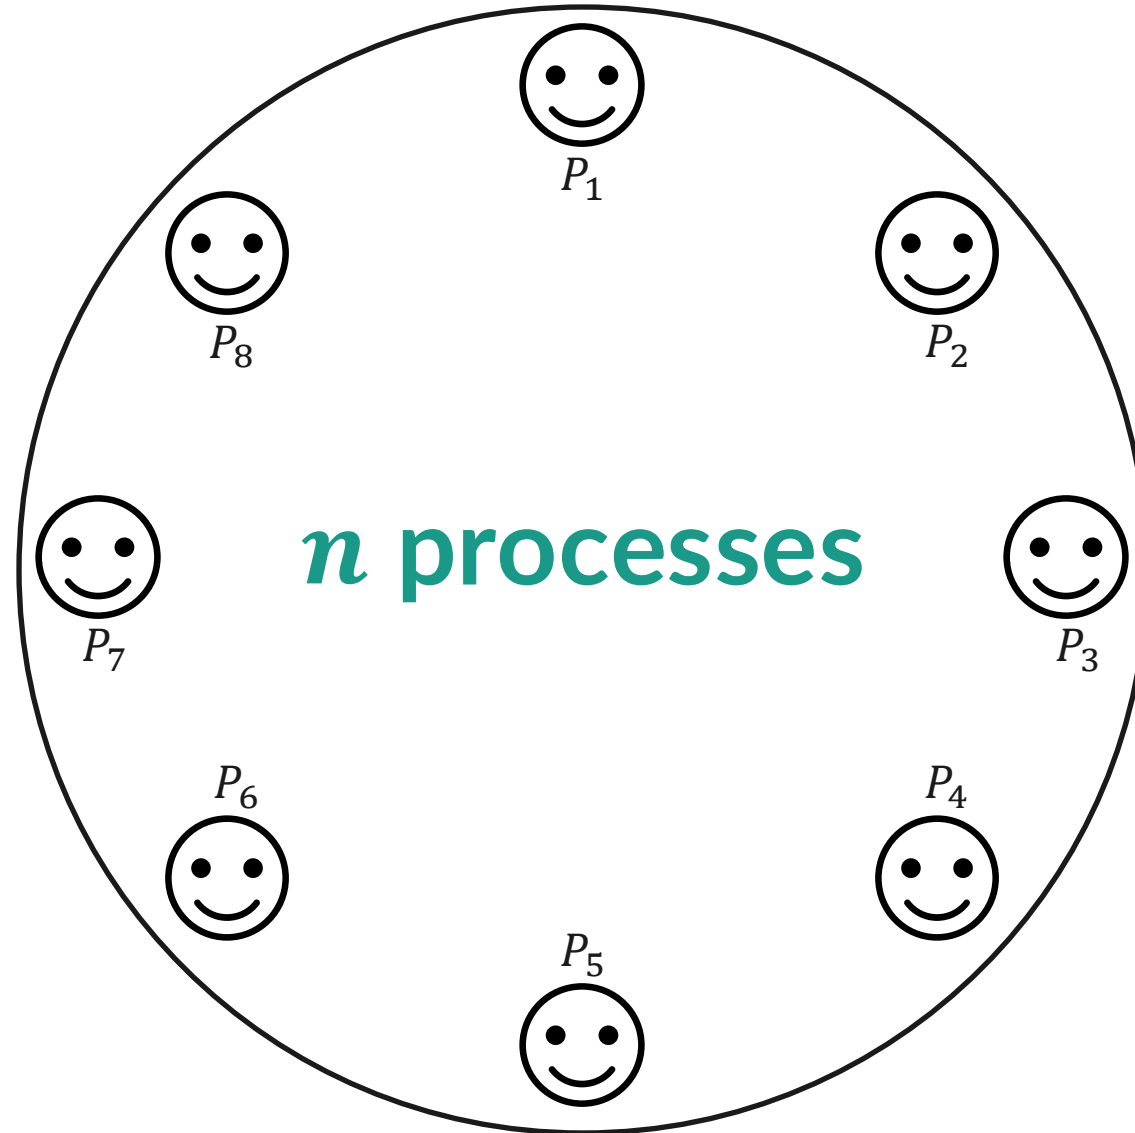
***Positive Result 1:** There exists a
randomized asynchronous protocol
that solves consensus, while
tolerating arbitrary (**Byzantine**) failures*

***Positive Result 2:** There exists a
deterministic synchronous protocol
that solves consensus, while
tolerating arbitrary (**Byzantine**) failures*

Consensus

Problem definition

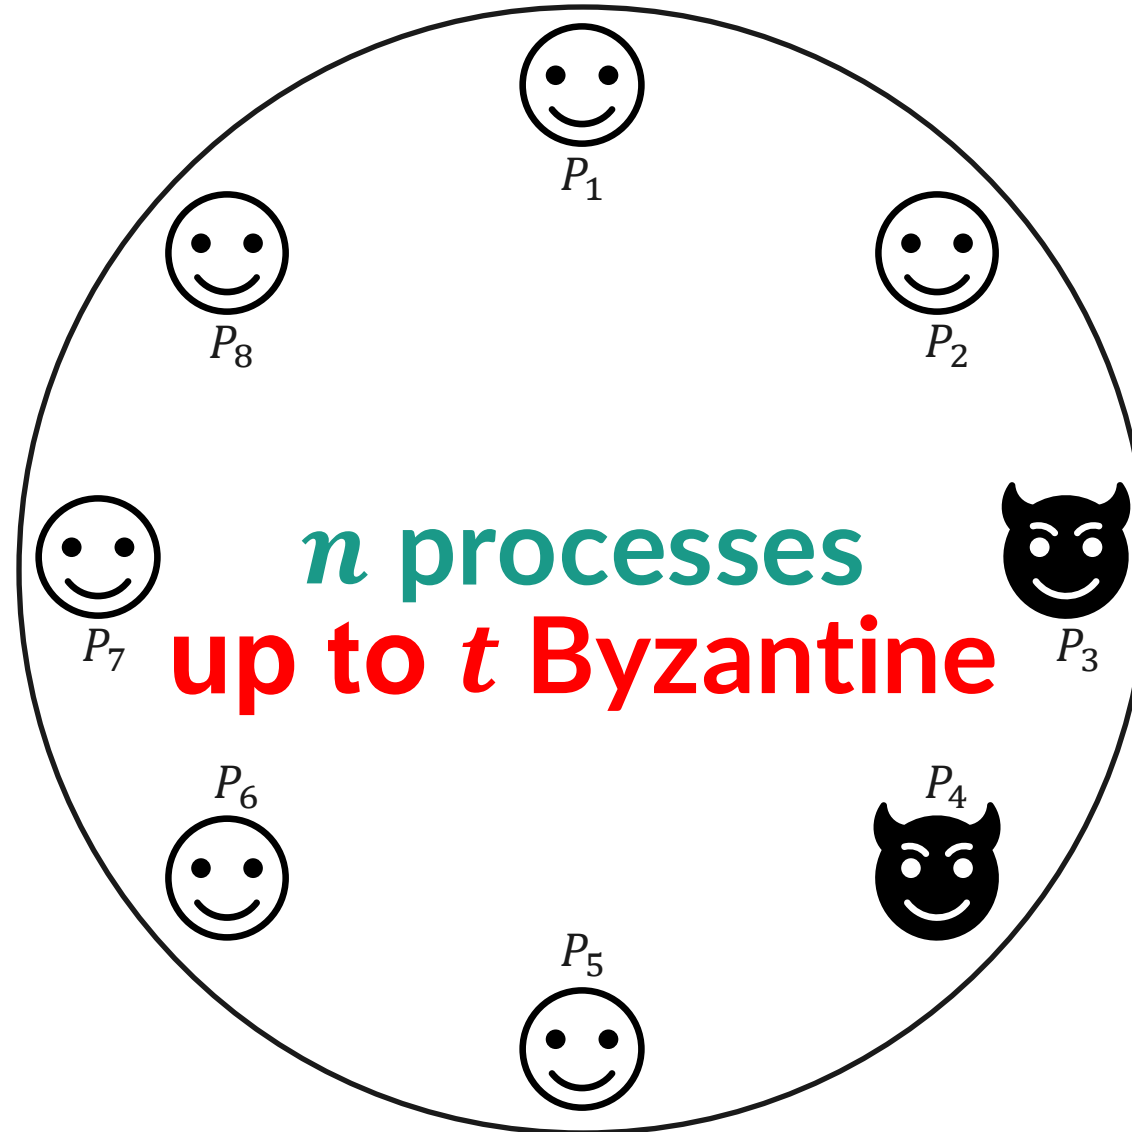
System



Static set of n publicly known identities

Message passing through reliable point-to-point channels

System




Static set of n publicly known identities

Message passing through reliable authenticated point-to-point channels

Interface of Consensus

Propose
Operation



Decide
Callback



Interface of Consensus

Propose
Operation



Decide
Callback



Interface of Consensus

Propose
Operation



Decide
Callback





Properties of Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Agreement:** No two correct processes decide different values.
- **Validity:** A decided value is proposed by a correct process.

Asynchronous Model

No notion of time



Asynchronous Model (Informal)

- No shared global clock: no shared notion of time
- Arbitrary (but finite) message delays
- Model is purely event-driven: reception → sending

Synchronous Model

Known upper bound Δ on message delays

→ rounds of the form compute, send, receive

Graded Consensus

« Stay safe »

Graded Consensus

Specification

Interface of graded consensus

Propose
Operation



Decide
Callback



Interface of graded consensus

Propose
Operation



Decide
Callback



Interface of graded consensus

Propose
Operation



Decide
Callback



Interface of graded consensus

Propose
Operation



Decide
Callback



Grade =
'confidence'

$\mathcal{I}_{min}, \dots, \mathcal{I}_{max}$

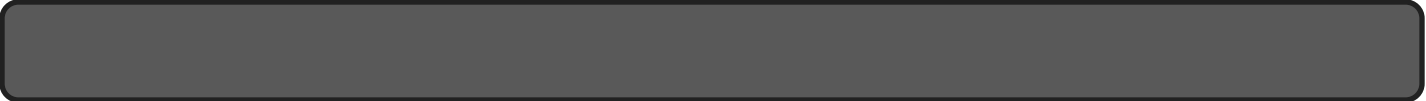
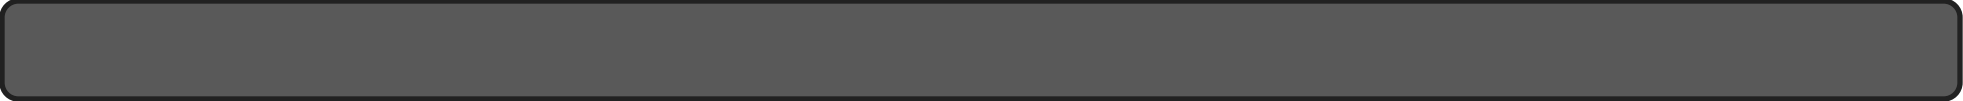


Properties of Graded Consensus

- Termination: All correct processes decide.
- Integrity: No process decides more than once.
- Unanimity: If only v is proposed, then only can be decided.
- Consistency (~~Agreement~~):



Properties of Graded Consensus

- Termination: All correct processes decide.
 - Integrity: No process decides more than once.
 - Unanimity: If only v is proposed, then only (v, g_{max}) can be decided.
 - Consistency (Agreement): 
- 

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have .

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g-g'| \leq$.

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g-g'| \leq 1$.

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g-g'| \leq 1$, and (b) .

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g-g'| \leq 1$, and (b) if $v \neq v'$, then .



Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g-g'| \leq 1$, and (b) if $v \neq v'$, then $g=g'=0$.


Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have either (i) or (ii) , and .

Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have either (i) $g=g'=0$ or (ii) , and .


Properties of Graded Consensus

- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have **either** (i) $g=g'=0$ or (ii) $|g-g'| \leq 1$, and .




Properties of Graded Consensus


- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have **either** (i) $g=g'=0$ or (ii) $|g-g'| \leq 1$, and $v=v'$.



Consistency + Termination + Unanimity




Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.

Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \implies g = g' = 0)$.

Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then either (i) $g = g' = 0$ or (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.

Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \implies g = g' = 0)$.

Consistency + Termination + Unanimity



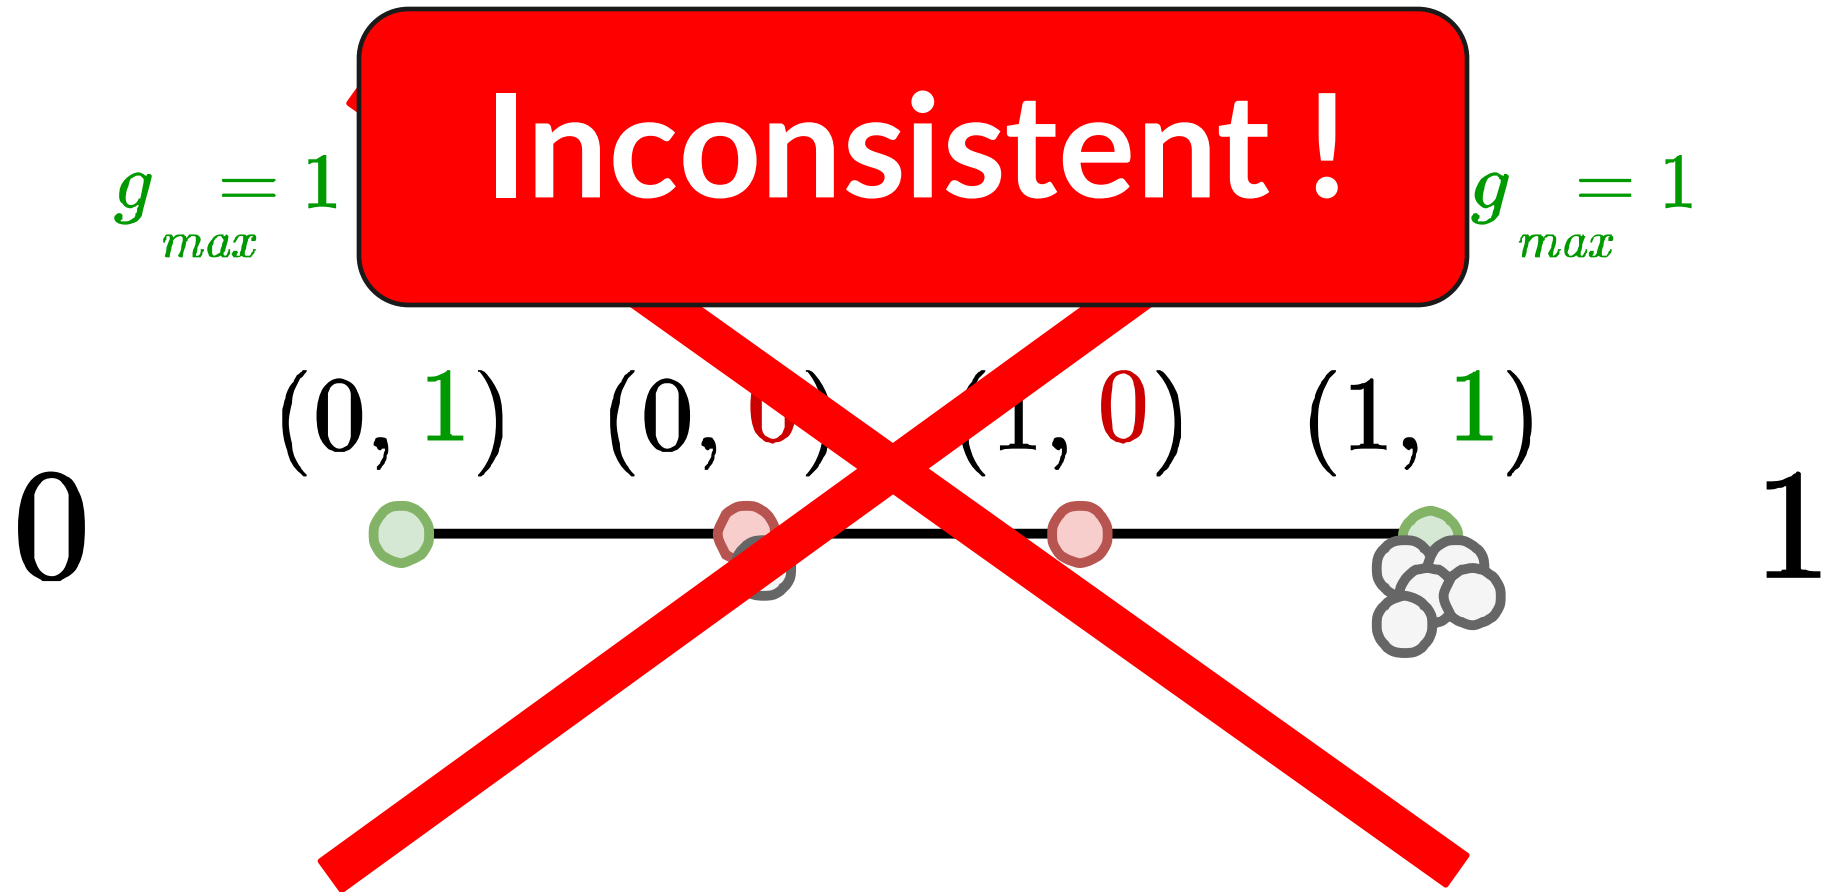
- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.

Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \implies g = g' = 0)$.

Consistency + Termination + Unanimity




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.

Consistency + Termination + Unanimity



All correct processes eventually decide the unanimous proposal with high confidence value.

Consistency + Termination + Unanimity



All correct processes eventually decide the unanimous proposal with high confidence value.

Graded Consensus

Asynchronous implementation with $t < n/3$ Byzantine failures

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination:

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $\square$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\square$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\square$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives \square PROPOSAL messages.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\blacksquare$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\blacksquare$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives $(n - t)$ PROPOSAL messages.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$  where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\blacksquare$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\blacksquare$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives $(n - t)$ PROPOSAL messages. Consequently, every correct process eventually broadcasts an ECHO message

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $\blacksquare$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\blacksquare$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives $(n - t)$ PROPOSAL messages. Consequently, every correct process eventually broadcasts an ECHO message and receives \blacksquare ECHO messages.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives $(n - t)$ PROPOSAL messages. Consequently, every correct process eventually broadcasts an ECHO message and receives $(n - t)$ ECHO messages.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\blacksquare$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Termination: Every correct process eventually triggers a proposal. Thus, every correct process eventually broadcasts a PROPOSAL message and receives $(n - t)$ PROPOSAL messages. Consequently, every correct process eventually broadcasts an ECHO message and receives $(n - t)$ ECHO messages before triggering a decision.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least   ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Unanimity Property:

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ )
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\lceil (n - t)/2 \rceil$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v .

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least   ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v .

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\lceil (n - t)/2 \rceil$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\lceil (n - t)/2 \rceil$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3:   upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:     broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5:   upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:     if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $t$  ECHO messages contain value  $v''$ :
7:       trigger decide( $v'', 1$ )
8:     else:
9:       trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$  where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring that every correct process broadcasts an ECHO message with value v .

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring that every correct process broadcasts an ECHO message with value v . Thus, each correct process receives $(n - t)$ ECHO messages,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $\lceil (n - t)/2 \rceil$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring that every correct process broadcasts an ECHO message with value v . Thus, each correct process receives $(n - t)$ ECHO messages, including at least $\lceil (n - t)/2 \rceil$ with value v ,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring that every correct process broadcasts an ECHO message with value v . Thus, each correct process receives $(n - t)$ ECHO messages, including at least $(n - 2t)$ with value v ,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v''$ , 1)
8:   else:
9:     trigger decide( $v^*$ , 0), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Unanimity Property: Suppose all correct processes propose the same value v . Then, each correct process broadcasts a PROPOSAL message with value v . Consequently, each process eventually receives $(n - t)$ PROPOSAL messages, including at least $(n - 2t)$ with value v and at most t with value $1 - v$. Given that $n > 3t$, we have $(n - 2t) > t$, ensuring that every correct process broadcasts an ECHO message with value v . Thus, each correct process receives $(n - t)$ ECHO messages, including at least $(n - 2t)$ with value v , and consequently decides on $(v, 1)$.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Consistency:

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate).

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j ,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) .

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*


PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| =$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/X$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages.
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages.
```



LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and n/X -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j|$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$.

Therefore, process p_j receives at least ████████ ECHO messages with value w , which ensures that ████████ if ████████

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$.

Therefore, process p_j receives at least $n - 4t$ ECHO messages with value w , which ensures that if

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$. Therefore, process p_j receives at least $n - 4t$ ECHO messages with value w , which ensures that $w' = w$ if

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$. Therefore, process p_j receives at least $n - 4t$ ECHO messages with value w , which ensures that $w' = w$ if $n - 4t > \blacksquare$

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/3$ Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/3$ -resiliency.*

PROOF.

Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$. Therefore, process p_j receives at least $n - 4t$ ECHO messages with value w , which ensures that $w' = w$ if $n - 4t > (n - t)/2$,

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/7$: Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/7$ -resiliency.*

PROOF.


Consistency: Assume a correct process p_i decides $(w, 1)$ (otherwise, consistency is immediate). Process p_i must have received ECHO messages with value w from a set Q_i of $|Q_i| = n - 2t$ distinct processes. Let another correct process, p_j , decide on some value (w', \cdot) . We aim to show that $w' = w$.

Process p_j 's decision was based on receiving ECHO messages from a set Q_j with $|Q_j| = n - t$ processes. The overlap between Q_i and Q_j is $|Q_i \cap Q_j| = |Q_i| + |Q_j| - |Q_i \cup Q_j| \geq (n - 2t) + (n - t) - n = n - 3t$, so $|Q_i \cap Q_j \cap \text{Corrects}| \geq n - 4t$. Therefore, process p_j receives at least $n - 4t$ ECHO messages with value w , which ensures that $w' = w$ if $n - 4t > (n - t)/2$, i.e., if $n > 7t$.

Algorithm 1 Asynchronous Binary Graded Consensus with refinement $R = 2$, and $t < n/7$: Pseudocode (for process p_i)

```
1: upon propose( $v_i \in \text{Binary\_Value}$ ):
2:   broadcast  $\langle \text{PROPOSAL}, v_i \rangle$ 
3: upon  $\langle \text{PROPOSAL}, \cdot \rangle$  is received from  $n - t$  processes:
4:   broadcast  $\langle \text{ECHO}, v' \rangle$ , where  $v'$  denotes the value with the highest frequency among the PROPOSAL messages. ▷  $\#(v') > (n - t)/2$ 
5: upon  $\langle \text{ECHO}, \cdot \rangle$  is received from  $n - t$  processes:
6:   if  $\exists v'' \in \text{Binary\_Value}$ , s.t. at least  $n - 2t$  ECHO messages contain value  $v''$ :
7:     trigger decide( $v'', 1$ )
8:   else:
9:     trigger decide( $v^*, 0$ ), where  $v^*$  denotes the value with the highest frequency among the ECHO messages. ▷  $\#(v^*) > (n - t)/2$ 
```

LEMMA 1.1. *Algorithm 1 implements graded consensus with the refinement parameter $R = 2$ and $n/7$ -resiliency.*



One more refinement




Properties of Graded Consensus


- **Termination:** All correct processes decide.
- **Integrity:** No process decides more than once.
- **Consistency (~~Agreement~~):** Assume two correct processes decide (v, g) and (v', g') . We have (a) $|g - g'| \leq 1$, and (b) if $v \neq v'$, then $g = g' = 0$.
- **Unanimity:** If only v is proposed, then only (v, g_{max}) can be decided.

$$g_{max} = 2$$


$$g_{min} = 0$$

$$g_{max} = 2$$







- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, or (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.




- (1) If two correct processes decide (v, g) and (v', g') , then either (i) $g = g' = 0$ or (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.



- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.




- (1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.
- (2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$, and (b) $(v \neq v' \Rightarrow g = g' = 0)$.



(1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g = g' = 0$, **or** (ii) $|g - g'| \leq 1$, and $v = v'$.



(2) If two correct processes decide (v, g) and (v', g') , then (a) $|g - g'| \leq 1$ and (b) $(v \neq v' \Rightarrow g = g' = 0)$.

Inconsistent !




(1) If two correct processes decide (v, g) and (v', g') , then **either** (i) $g=g'=0$, **or** (ii) $|g-g'| \leq 1$, and $v=v'$.

(2) If two correct processes decide (v, g) and (v', g') , then (a) $|g-g'| \leq 1$ and (b) $(v \neq v' \Rightarrow g=g'=0)$.



All correct processes eventually decide the unanimous proposal with high confidence value.



All correct processes eventually decide the unanimous proposal with high confidence value.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: [REDACTED]

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose([REDACTED])

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: [REDACTED]

▸ Update the grade (confidence)

12: **trigger** decide([REDACTED])

▸ Decide the final value and grade

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: ██████████

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: ██████████

▸ Update the grade (confidence)

12: **trigger** decide($v_i^2, \span style="background-color: black; color: black;">████)$

▸ Decide the final value and grade

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: [REDACTED]

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: [REDACTED]

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**
2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$ ▶ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**
4: Integer $g_i \leftarrow 0$ ▶ Grade

5: **upon** propose($v_i \in \text{Value}$):
6: **invoke** \mathcal{GC}_1 .propose(v_i) ▶ Propose input value to the 1st graded consensus instance
7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1): ▶ Received decision from the 1st graded consensus instance
8: $g_i \leftarrow g_i + g_i^1$ ▶ Update the grade (confidence)
9: **invoke** \mathcal{GC}_2 .propose(v_i^1) ▶ Propose input value to the 2nd graded consensus instance
10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2): ▶ Received decision from the 2nd graded consensus instance
11: $g_i \leftarrow g_i + g_i^2$ ▶ Update the grade (confidence)
12: **trigger** decide(v_i^2, g_i) ▶ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

PROOF.

- Termination follows directly from the termination of \mathcal{GC}_1 and \mathcal{GC}_2 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**
2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$ ▶ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**
4: Integer $g_i \leftarrow 0$ ▶ Grade

5: **upon** propose($v_i \in \text{Value}$):
6: **invoke** \mathcal{GC}_1 .propose(v_i) ▶ Propose input value to the 1st graded consensus instance
7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1): ▶ Received decision from the 1st graded consensus instance
8: $g_i \leftarrow g_i + g_i^1$ ▶ Update the grade (confidence)
9: **invoke** \mathcal{GC}_2 .propose(v_i^1) ▶ Propose input value to the 2nd graded consensus instance
10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2): ▶ Received decision from the 2nd graded consensus instance
11: $g_i \leftarrow g_i + g_i^2$ ▶ Update the grade (confidence)
12: **trigger** decide(v_i^2, g_i) ▶ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

PROOF.

- Termination follows directly from the termination of \mathcal{GC}_1 and \mathcal{GC}_2 .
- Unanimity property follows directly from the unanimity property of \mathcal{GC}_1 and \mathcal{GC}_2 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**
2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$ ▶ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**
4: Integer $g_i \leftarrow 0$ ▶ Grade

5: **upon** propose($v_i \in \text{Value}$):
6: **invoke** \mathcal{GC}_1 .propose(v_i) ▶ Propose input value to the 1st graded consensus instance
7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1): ▶ Received decision from the 1st graded consensus instance
8: $g_i \leftarrow g_i + g_i^1$ ▶ Update the grade (confidence)
9: **invoke** \mathcal{GC}_2 .propose(v_i^1) ▶ Propose input value to the 2nd graded consensus instance
10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2): ▶ Received decision from the 2nd graded consensus instance
11: $g_i \leftarrow g_i + g_i^2$ ▶ Update the grade (confidence)
12: **trigger** decide(v_i^2, g_i) ▶ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

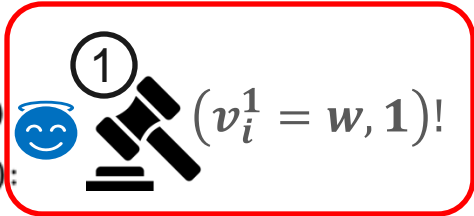
▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)  ($v_i^1 = w, 1$)!

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*


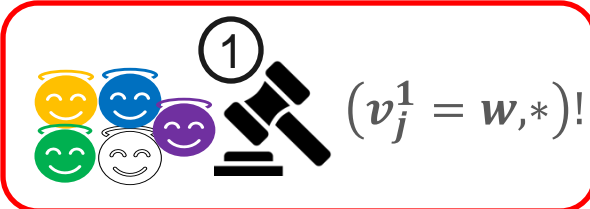
- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 1: $j = 1$.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**
 2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$ ▶ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**
 4: Integer $g_i \leftarrow 0$ ▶ Grade

5: **upon** propose($v_i \in \text{Value}$):
 6: **invoke** \mathcal{GC}_1 .propose(v_i)  $(v_i^1 = w, 1)!$  $(v_j^1 = w, *)!$ ▶ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1): ▶ Received decision from the 1st graded consensus instance
 8: $g_i \leftarrow g_i + g_i^1$ ▶ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1) ▶ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2): ▶ Received decision from the 2nd graded consensus instance
 11: $g_i \leftarrow g_i + g_i^2$ ▶ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i) ▶ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 1: $j = 1$. By consistency, each correct process outputs (w, \cdot) from \mathcal{GC}_1 for some value w , and thus proposes w to \mathcal{GC}_2 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

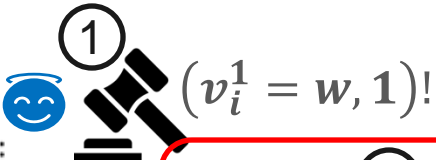
▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)  $(v_i^1 = w, 1)!$

 $(v_j^1 = w, *)!$

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

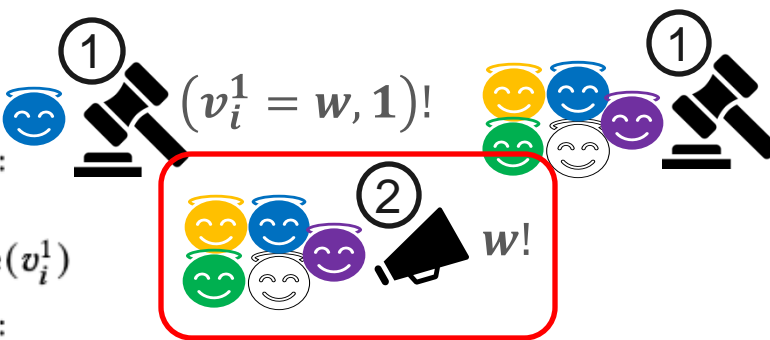
▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 1: $j = 1$. By consistency, each correct process outputs (w, \cdot) from \mathcal{GC}_1 for some value w , and thus proposes w to \mathcal{GC}_2 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

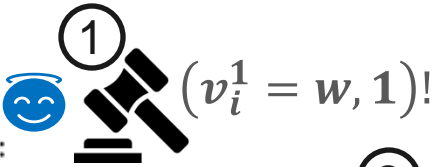
▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)  $(v_i^1 = w, 1)!$

 $(v_j^1 = w, *)!$ ▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)



▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):


▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



LEMMA 1.2. Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.



- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 1: $j = 1$. By consistency, each correct process outputs (w, \cdot) from \mathcal{GC}_1 for some value w , and thus proposes w to \mathcal{GC}_2 . Therefore, due to the unanimity property of graded consensus \mathcal{GC}_2 , every correct process returns $(w, 1)$ from \mathcal{GC}_2 .


Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**
 2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$ ▶ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**
 4: Integer $g_i \leftarrow 0$ ▶ Grade

5: **upon** propose($v_i \in \text{Value}$):
 6: **invoke** \mathcal{GC}_1 .propose(v_i)  $(v_i^1 = w, 1)!$  $(v_j^1 = w, *)!$ ▶ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1): ▶ Received decision from the 1st graded consensus instance
 8: $g_i \leftarrow g_i + g_i^1$ ▶ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)  $w!$ ▶ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2): ▶ Received decision from the 2nd graded consensus instance
 11: $g_i \leftarrow g_i + g_i^2$ ▶ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)  $(v_*^2 = w, 1)!$ ▶ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 1: $j = 1$. By consistency, each correct process outputs (w, \cdot) from \mathcal{GC}_1 for some value w , and thus proposes w to \mathcal{GC}_2 . Therefore, due to the unanimity property of graded consensus \mathcal{GC}_2 , every correct process returns $(w, 1)$ from \mathcal{GC}_2 . Hence, consistency follows directly from the consistency of \mathcal{GC}_1 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the **first instance** of graded consensus where **some process outputs $(\cdot, 1)$ from \mathcal{GC}_j** . If no such j exists, the result is immediate.

Case 2: $j = 2$.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

12: **trigger** decide(v_i^2, g_i)



▸ Propose input value to the 1st graded consensus instance

▸ Received decision from the 1st graded consensus instance

▸ Update the grade (confidence)

▸ Propose input value to the 2nd graded consensus instance

▸ Received decision from the 2nd graded consensus instance

▸ Update the grade (confidence)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

12: **trigger** decide(v_i^2, g_i)



▸ Propose input value to the 1st graded consensus instance

▸ Received decision from the 1st graded consensus instance

▸ Update the grade (confidence)

▸ Propose input value to the 2nd graded consensus instance

▸ Received decision from the 2nd graded consensus instance

▸ Update the grade (confidence)

▸ Decide the final value and grade



LEMMA 1.2. Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 .

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)



▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 ,

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):



▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then $|g_i - g_j| \leq 1$.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then $|g_i - g_j| \leq 1$. Moreover,

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):



▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade

LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then $|g_i - g_j| \leq 1$. Moreover, if $g_i \neq 0$,

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$


1: **Uses:**
 2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**
 4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):
 6: **invoke** \mathcal{GC}_1 .propose(v_i)
 7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):
 8: $g_i \leftarrow g_i + g_i^1$
 9: **invoke** \mathcal{GC}_2 .propose(v_i^1)
 10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):
 11: $g_i \leftarrow g_i + g_i^2$
 12: **trigger** decide(v_i^2, g_i)



- Propose input value to the 1st graded consensus instance
- Received decision from the 1st graded consensus instance
 - Update the grade (confidence)
- Propose input value to the 2nd graded consensus instance
- Received decision from the 2nd graded consensus instance
 - Update the grade (confidence)
 - Decide the final value and grade



LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then $|g_i - g_j| \leq 1$. Moreover, if $g_i \neq 0$, $v_i = v_j$.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, instances $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$


3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)



▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



LEMMA 1.2. *Algorithm 2 implements graded consensus with the refinement parameter $R = 3$.*

- To prove consistency, let j be the first instance of graded consensus where some process outputs $(\cdot, 1)$ from \mathcal{GC}_j . If no such j exists, the result is immediate.

Case 2: $j = 2$. By construction, each correct process outputs $(\cdot, 0)$ from \mathcal{GC}_1 . Therefore, due to the consistency property of graded consensus \mathcal{GC}_2 , if two correct processes p_i and p_j decide on (v_i, g_i) and (v_j, g_j) , respectively, then $|g_i - g_j| \leq 1$. Moreover, if $g_i \neq 0, v_i = v_j$. This implies consistency.

Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):


▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade



Algorithm 2 Binary Graded Consensus (BGC) with refinement $R = 3$ on top of BGC with refinement $R = 2$

1: **Uses:**

2: Binary Graded Consensus, **instances** $\mathcal{GC}_1, \mathcal{GC}_2$

▸ 2 instances of the Binary Graded Consensus protocol with Refinement $R' = 2$

3: **Local Variables:**

4: Integer $g_i \leftarrow 0$

▸ Grade

5: **upon** propose($v_i \in \text{Value}$):

6: **invoke** \mathcal{GC}_1 .propose(v_i)

▸ Propose input value to the 1st graded consensus instance

7: **upon** \mathcal{GC}_1 .decide(v_i^1, g_i^1):

▸ Received decision from the 1st graded consensus instance

8: $g_i \leftarrow g_i + g_i^1$

▸ Update the grade (confidence)

9: **invoke** \mathcal{GC}_2 .propose(v_i^1)

▸ Propose input value to the 2nd graded consensus instance

10: **upon** \mathcal{GC}_2 .decide(v_i^2, g_i^2):


▸ Received decision from the 2nd graded consensus instance

11: $g_i \leftarrow g_i + g_i^2$

▸ Update the grade (confidence)

12: **trigger** decide(v_i^2, g_i)

▸ Decide the final value and grade




Common Coin

Interface of the common coin


Flip
Operation




Yield
Callback



Interface of the common coin



Interface of the common coin





Properties of the Common Coin

- **Termination:** All correct processes eventually yield a binary value.
- **Agreement:** All correct processes agree on 0 (or 1) with probability >0 .
- **Unpredictability:** As soon no correct process has triggered 'flip', the adversary cannot predict the output with probability greater than $1/2$.



A naive implementation of the Common Coin

Algorithm 3 Common Coin

1: **upon** flip():
2: $b \stackrel{\$}{\leftarrow} \{0, 1\}$
3: **trigger** yield(b)


► Choose either 0 or 1 with probability 1/2

A naive implementation of the Common Coin


Algorithm 3 Common Coin

1: **upon** flip():
2: $b \xleftarrow{\$} \{0, 1\}$
3: **trigger** yield(b)

► Choose either 0 or 1 with probability 1/2



***Consensus = Stay safe + Try
(and try again)***



Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
1: Uses:
2:   Binary "Extended" Graded Consensus, instances  $\mathcal{EGC}_1, \mathcal{EGC}_2, \dots$  ▶  $\infty$  instances of the Binary Graded Consensus with Refinement  $R = 3$ 
3: Constants:
4:   Integer  $g_{min} \leftarrow 0$ 
5:   Integer  $g_{max} \leftarrow 2$ 
6: Local Variables:
7:   Binary_Value  $est_i \leftarrow 0$  ▶ Estimate Value
8:   Integer  $g_i \leftarrow g_{min}$  ▶ Grade (Confidence) in  $\{0, 1, 2\}$ 
9:   Integer  $attempt \leftarrow 0$ 
10:  Integer  $halt \leftarrow \infty$ 
11:  Boolean  $decided \leftarrow false$ 
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ 
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if ████████████████████
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if ████████
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
1: Uses:
2:   Binary "Extended" Graded Consensus, instances  $\mathcal{EGC}_1, \mathcal{EGC}_2, \dots$  ▶  $\infty$  instances of the Binary Graded Consensus with Refinement  $R = 3$ 
3: Constants:
4:   Integer  $g_{min} \leftarrow 0$ 
5:   Integer  $g_{max} \leftarrow 2$ 
6: Local Variables:
7:   Binary_Value  $est_i \leftarrow 0$  ▶ Estimate Value
8:   Integer  $g_i \leftarrow g_{min}$  ▶ Grade (Confidence) in  $\{0, 1, 2\}$ 
9:   Integer  $attempt \leftarrow 0$ 
10:  Integer  $halt \leftarrow \infty$ 
11:  Boolean  $decided \leftarrow false$ 
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ 
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if ██████████
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
1: Uses:  
2:   Binary "Extended" Graded Consensus, instances  $\mathcal{EGC}_1, \mathcal{EGC}_2, \dots$  ▶  $\infty$  instances of the Binary Graded Consensus with Refinement  $R = 3$   
3: Constants:  
4:   Integer  $g_{min} \leftarrow 0$   
5:   Integer  $g_{max} \leftarrow 2$   
6: Local Variables:  
7:   Binary_Value  $est_i \leftarrow 0$  ▶ Estimate Value  
8:   Integer  $g_i \leftarrow g_{min}$  ▶ Grade (Confidence) in  $\{0, 1, 2\}$   
9:   Integer  $attempt \leftarrow 0$   
10:  Integer  $halt \leftarrow \infty$   
11:  Boolean  $decided \leftarrow false$   
12: upon propose( $v_i \in \text{Value}$ ):  
13:    $est_i \leftarrow v_i$ :  
14:   while  $halt \geq attempt$ :  
15:     //safety guard  
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus  
17:     if  $g_i == g_{max} \wedge decided = false$ :  
18:       trigger decide( $est_i$ ) ▶ Decide  
19:        $decided \leftarrow true$   
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide  
21:     //try to converge  
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin  
23:     if  $g_i == g_{min}$ :  
24:        $est_i \leftarrow b_i$   
25:      $attempt \leftarrow attempt + 1$ 
```

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$                          ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.3. *If all correct processes begin attempt k with the same estimate value v , they will all decide on v by attempt k and halt by attempt $k + 1$.*

PROOF.

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.3. *If all correct processes begin attempt k with the same estimate value v , they will all decide on v by attempt k and halt by attempt $k + 1$.*

PROOF. By the unanimity property of \mathcal{EGC}_k ,

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.3. *If all correct processes begin attempt k with the same estimate value v , they will all decide on v by attempt k and halt by attempt $k + 1$.*

PROOF. By the unanimity property of \mathcal{EGC}_k , all correct processes return (v, g_{max}) from \mathcal{EGC}_k .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.3. *If all correct processes begin attempt k with the same estimate value v , they will all decide on v by attempt k and halt by attempt $k + 1$.*

PROOF. By the unanimity property of \mathcal{EGC}_k , all correct processes return (v, g_{max}) from \mathcal{EGC}_k . The rest follows directly from the protocol. □

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger  $decide(est_i)$  ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k , updating its estimate est_j to v .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k , updating its estimate est_j to v . Thus, $g_j > g_{min}$,

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k , updating its estimate est_j to v . Thus, $g_j > g_{min}$, so all correct processes ignore the output of CC_k and retain $est_j = v$.

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k , updating its estimate est_j to v . Thus, $g_j > g_{min}$, so all correct processes ignore the output of CC_k and retain $est_j = v$. Consequently, all correct processes begin attempt $k + 1$ with estimate value v . □

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

LEMMA 1.4. *If a correct process decides on v in attempt k , then all correct processes will decide on v by attempt $k + 1$.*

PROOF. Let p_i be the first correct process to decide, and assume it decides on v at attempt k . This implies that p_i returned (v, g_{max}) from \mathcal{EGC}_k . By the consistency property of \mathcal{EGC}_k , every correct process p_j returns $(v, g_j \in \{1, 2\})$ from \mathcal{EGC}_k , updating its estimate est_j to v . Thus, $g_j > g_{min}$, so all correct processes ignore the output of CC_k and retain $est_j = v$. Consequently, all correct processes begin attempt $k + 1$ with estimate value v . Lemma 1.3 then completes the proof. \square

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$                              ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF.

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$                          ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from :

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$                          ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3,

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$                          ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4.


Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$            ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )                                     ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$                                ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt}.flip()$                          ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4. We now prove **Termination**.




Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4. We now prove **Termination**. Let p_i be the first correct process calling CC_k for some attempt k .




Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4. We now prove **Termination**. Let p_i be the first correct process calling CC_k for some attempt k . Let (b, g_i) the pair returned by p_i from \mathcal{EGC}_k .




Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4. We now prove **Termination**. Let p_i be the first correct process calling CC_k for some attempt k . Let (b, g_i) the pair returned by p_i from \mathcal{EGC}_k . With non-zero probability ρ , all correct processes return b from CC_k .




Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

PROOF. **Validity** follows from Lemma 1.3, and **Agreement** follows from Lemma 1.4. We now prove **Termination**. Let p_i be the first correct process calling CC_k for some attempt k . Let (b, g_i) the pair returned by p_i from \mathcal{EGC}_k . With non-zero probability ρ , all correct processes return b from CC_k . We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ : ▶ Decide
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*






We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$ 
17:     if  $g_i == g_{max} \wedge decided = \text{false}$ :
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow \text{true}$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt} \cdot \text{flip}()$ 
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 

```



 $(b, g_i)!$


 $(b, g_j > g_{min})!$


▶ Execute instance of extended graded consensus
 ▶ Decide
 ▶ Halt after the next attempt after having helped the remaining processes to decide
 ▶ Execute instance of common coin

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.




- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$:

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}$  propose( $est_i$ )
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$ 
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:        $attempt \leftarrow attempt + 1$ 

```

$(b, g_i)!$
 $(b, g_j > g_{min})!$

▶ Execute instance of extended graded consensus
 ▶ Decide
 ▶ Halt after the next attempt after having helped the remaining processes to decide
 ▶ Execute instance of common coin








THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency,

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )    $(b, g_j > g_{min})!$  ▶ Decide
19:       decided  $\leftarrow true$ 
20:       halt  $\leftarrow attempt + 1$     $(b, *)!$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt} \cdot \text{flip}()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:     attempt  $\leftarrow attempt + 1$ 

```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency, (a) all processes return (b, \cdot) from \mathcal{EGC}_k ,

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$ 
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt} \cdot \text{flip}()$ 
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 

```

- ▶ Execute instance of extended graded consensus
- ▶ Decide
- ▶ Halt after the next attempt after having helped the remaining processes to decide
- ▶ Execute instance of common coin

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency, (a) all processes return (b, \cdot) from \mathcal{EGC}_k , so they will all have the same estimate value at attempt $k + 1$, either by adopting CC_k 's output or by retaining the value from \mathcal{EGC}_k .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$ 
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt} \cdot \text{flip}()$ 
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 

```

- ▶ Execute instance of extended graded consensus
- ▶ Decide
- ▶ Halt after the next attempt after having helped the remaining processes to decide
- ▶ Execute instance of common coin





THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*

We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency, (a) all processes return (b, \cdot) from \mathcal{EGC}_k , so they will all have the same estimate value at attempt $k + 1$, either by adopting CC_k 's output or by retaining the value from \mathcal{EGC}_k .

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$    $(*, 0)!$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt} \cdot \text{flip}()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 

```

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*




We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency, (a) all processes return (b, \cdot) from \mathcal{EGC}_k , so they will all have the same estimate value at attempt $k + 1$, either by adopting CC_k 's output or by retaining the value from \mathcal{EGC}_k .

- **Case** $\forall p_j \in \text{Correct}, g_j = g_{min}$:

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt}.propose(est_i)$     $(b, g_i)!$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow CC_{attempt}.flip()$   ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 

```

Note: In the original image, lines 16-18 are enclosed in a purple box with a gavel icon and the text $(, g_{min})!$, and lines 22-24 are enclosed in a red box.*

THEOREM 1.5. *Algorithm 4 implements binary Byzantine consensus with probability 1 and has the same resiliency as the underlying graded consensus object.*


We consider two cases depending on the grades obtained by correct processes from \mathcal{EGC}_k . In both cases, all correct processes start the next attempt with the same estimate value, allowing us to apply Lemma 1.3.

- **Case** $\exists p_j \in \text{Correct}, g_j > g_{min}$: By \mathcal{EGC}_k 's consistency, (a) all processes return (b, \cdot) from \mathcal{EGC}_k , so they will all have the same estimate value at attempt $k + 1$, either by adopting CC_k 's output or by retaining the value from \mathcal{EGC}_k .

- **Case** $\forall p_j \in \text{Correct}, g_j = g_{min}$: All processes adopt the value provided by the common coin, which is identical across processes. □

Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```
12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$  ▶ Execute instance of extended graded consensus
17:     if  $g_i == g_{max} \wedge decided = false$ :
18:       trigger decide( $est_i$ ) ▶ Decide
19:        $decided \leftarrow true$ 
20:        $halt \leftarrow attempt + 1$  ▶ Halt after the next attempt after having helped the remaining processes to decide
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt} \cdot \text{flip}()$  ▶ Execute instance of common coin
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 
```




Algorithm 4 Byzantine Agreement Protocol with Extended Graded Consensus

```

12: upon propose( $v_i \in \text{Value}$ ):
13:    $est_i \leftarrow v_i$ :
14:   while  $halt \geq attempt$ :
15:     //safety guard
16:      $(est_i, g_i) \leftarrow \mathcal{EGC}_{attempt} \cdot \text{propose}(est_i)$ 
17:     if  $g_i == g_{max} \wedge decided = \text{false}$ :
18:       trigger decide( $est_i$ )
19:        $decided \leftarrow \text{true}$ 
20:        $halt \leftarrow attempt + 1$ 
21:     //try to converge
22:      $b_i \leftarrow \mathcal{CC}_{attempt} \cdot \text{flip}()$ 
23:     if  $g_i == g_{min}$ :
24:        $est_i \leftarrow b_i$ 
25:      $attempt \leftarrow attempt + 1$ 


```


▶ Execute instance of extended graded consensus
 ▶ Decide
 ▶ Halt after the next attempt after having helped the remaining processes to decide
 ▶ Execute instance of common coin

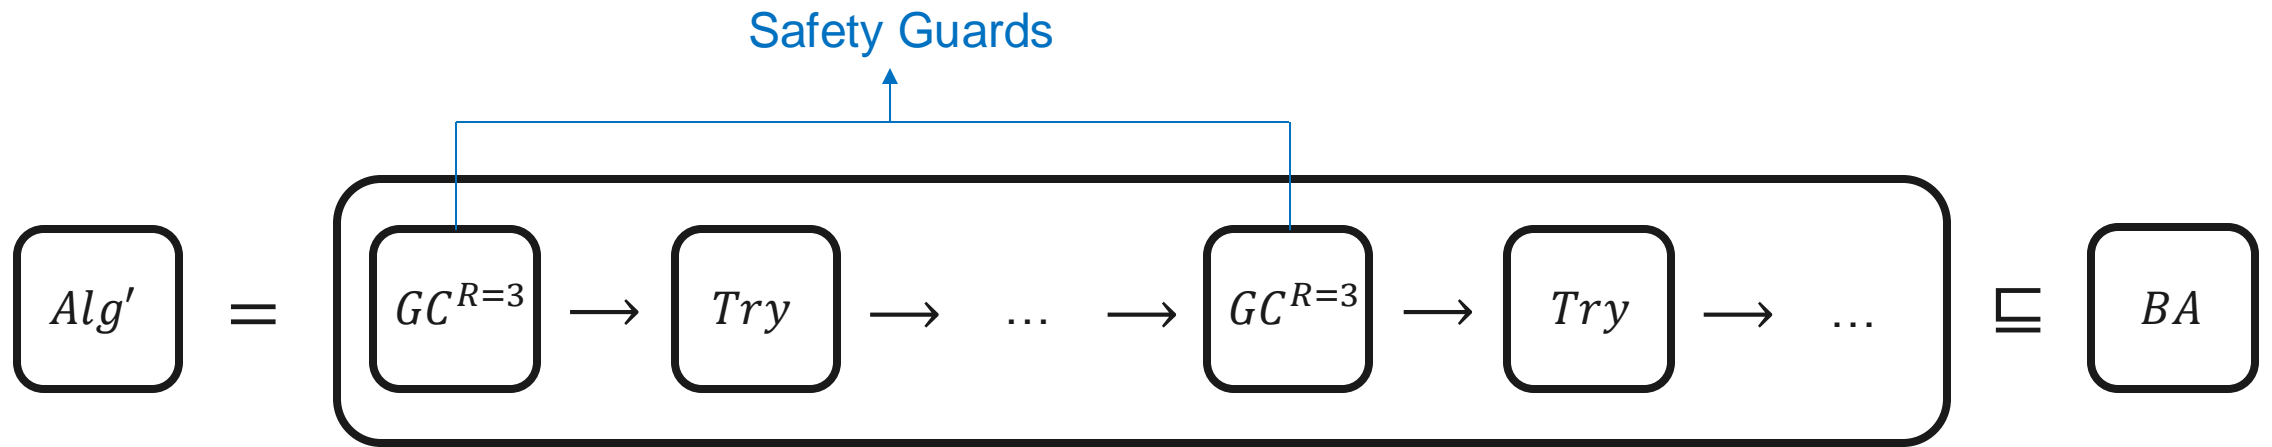


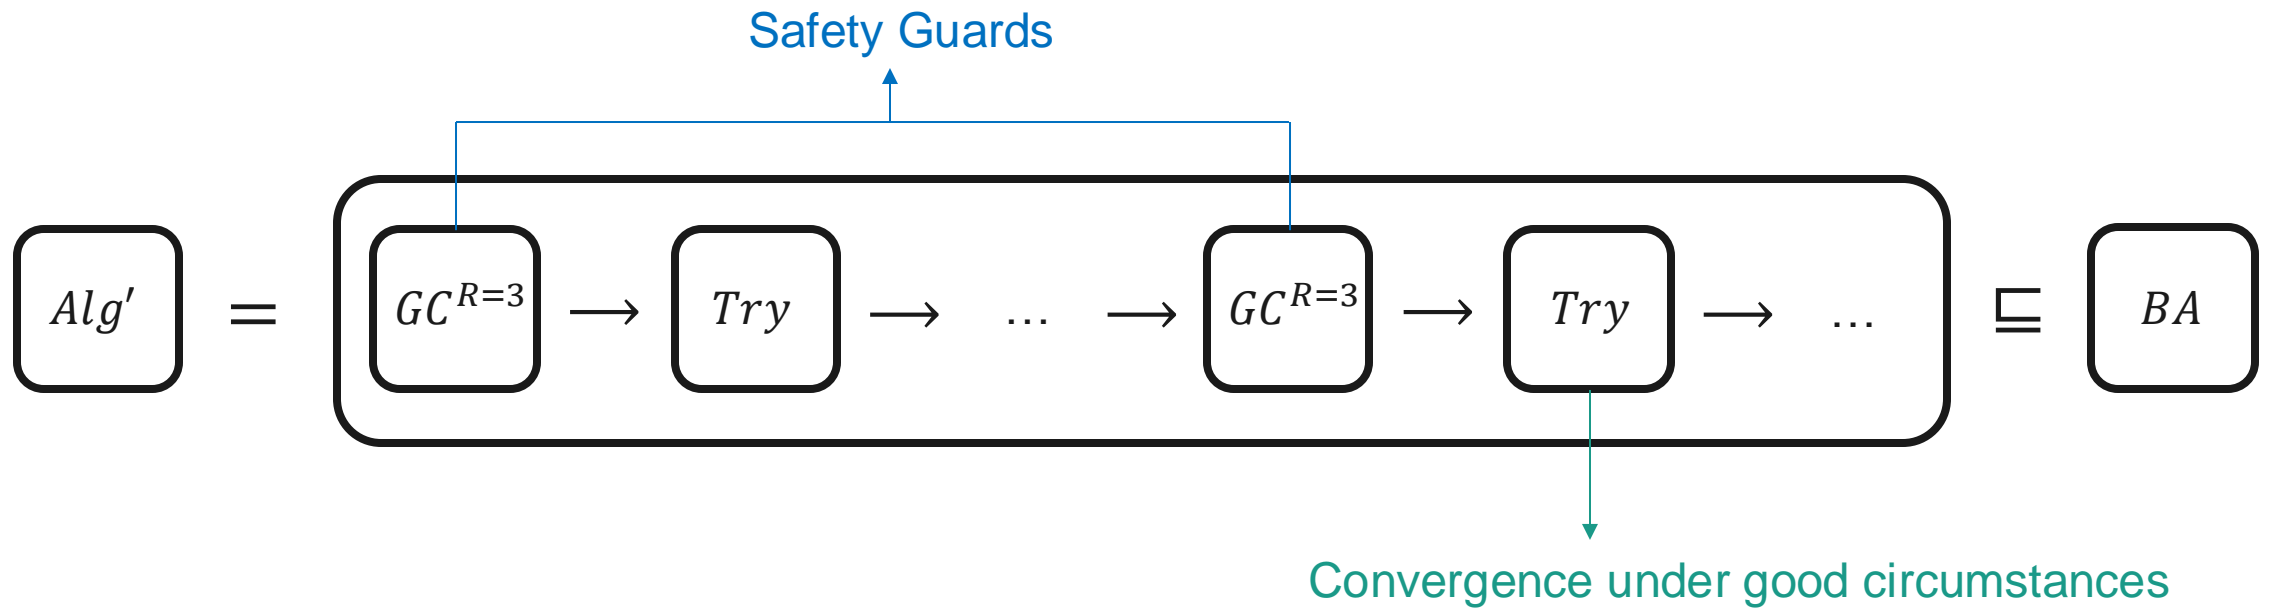
***Positive Result 1:** There exists a
randomized asynchronous protocol
that solves consensus, while
tolerating arbitrary (**Byzantine**) failures*

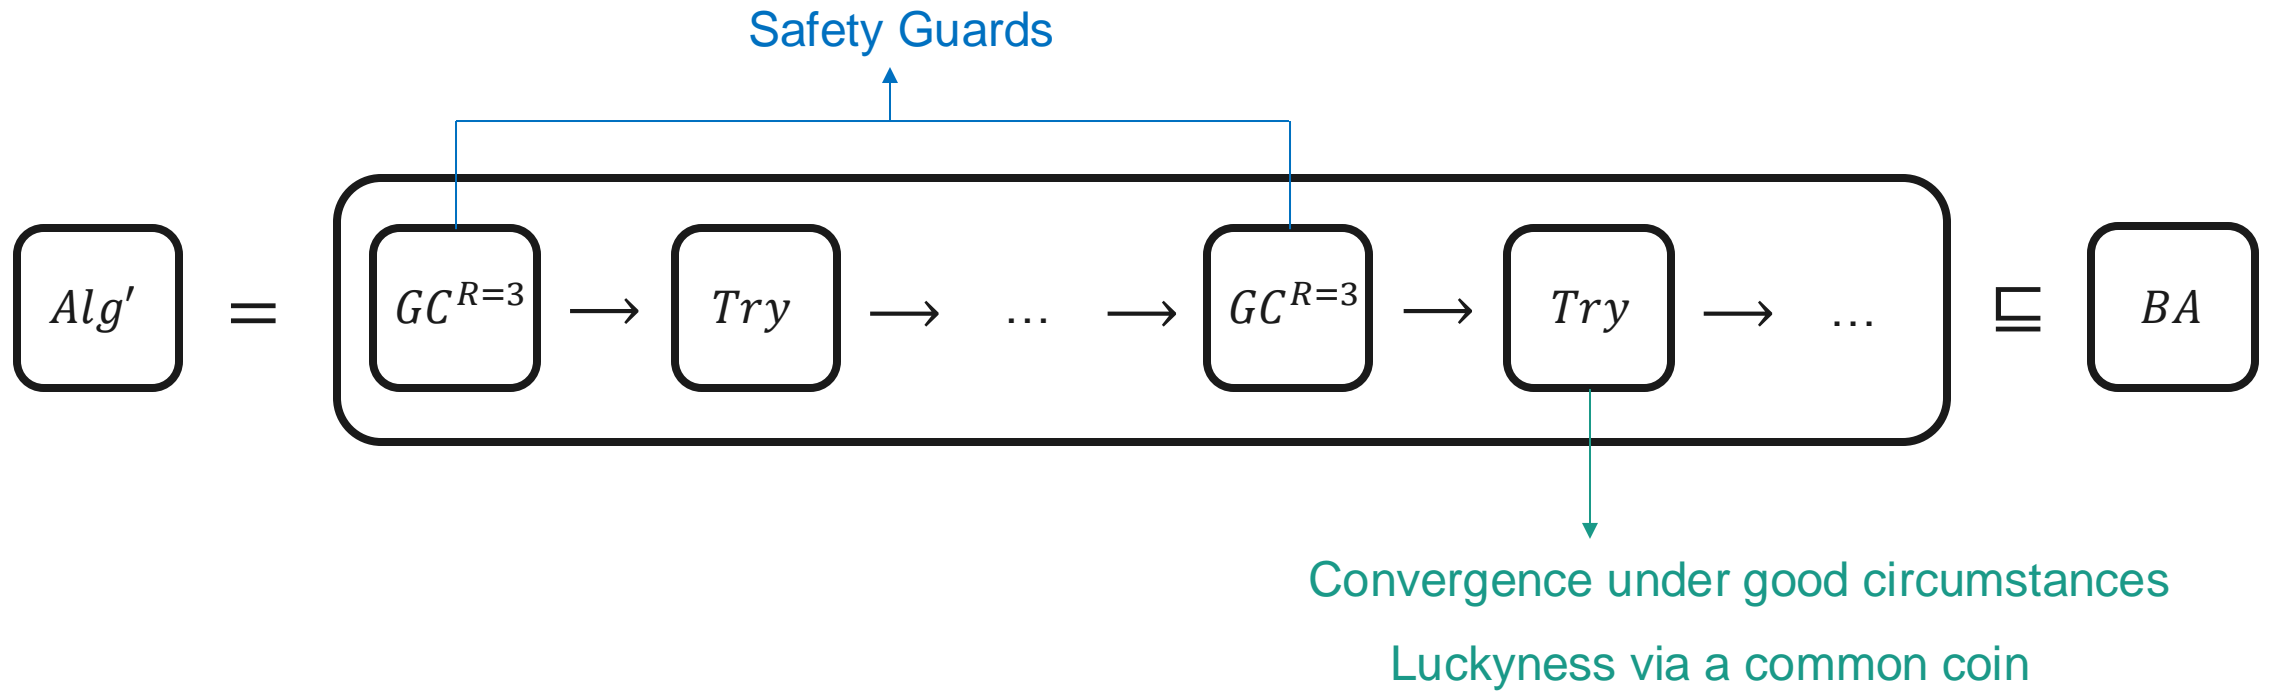
A general perspective











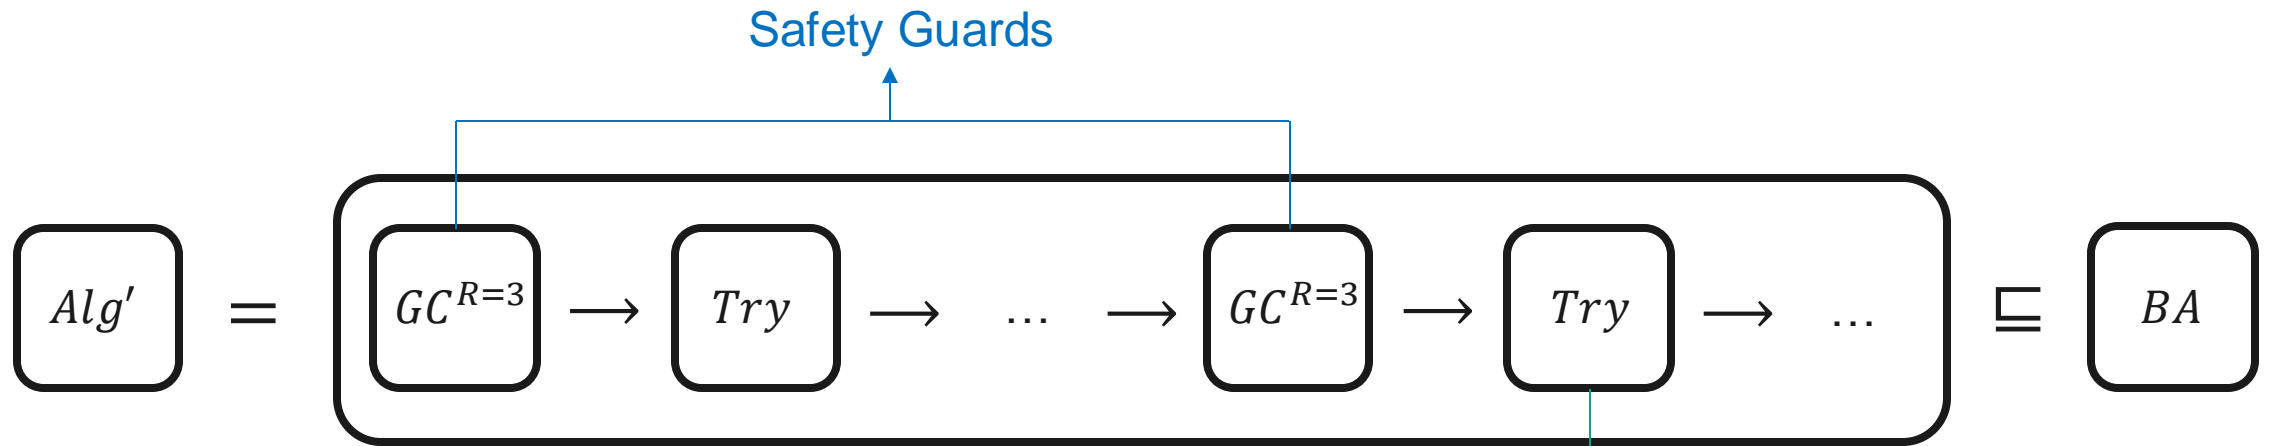




Convergence under good circumstances
 Luckyness via a common coin
 Eventual Synchrony + Synchronization



- Convergence under good circumstances
- Luckyness via a common coin
 - Eventual Synchrony + Synchronization
 - Unreliable Failure Detectors




Convergence under good circumstances

Luckyness via a common coin

Eventual Synchrony + Synchronization

Unreliable Failure Detectors

Fair scheduling / Noisy Environment



Convergence under good circumstances


Luckyness via a common coin

Eventual Synchrony + Synchronization

Unreliable Failure Detectors

Fair scheduling / Noisy Environment

Synchrony + round-robin rotating leader



Correct Leader



Synchrony

***Positive Result 2:** There exists a
deterministic synchronous protocol
that solves consensus, while
tolerating arbitrary (**Byzantine**) failures*

—