

CSC 144V: Autonomous Vehicles

Lecture 4.3

Latency and Availability Attacks Exercises Solution

Department of Computer Science,
Hofstra University

Exercise: Delay becomes distance

An AV drives at **25 m/s** (90 km/h) and brakes with $a = 6 \text{ m/s}^2$. Normal perception latency is **50 ms**; a latency attack adds **250 ms** (300 ms in total). An obstacle becomes detectable at **58 m**.

- (a) How far does the AV travel blind because of the extra 250 ms ($d = v \cdot \Delta t$)?
- (b) Stopping distance (latency distance + $v^2/2a$) without and with the attack?
- (c) Does the AV hit the obstacle in each case? If so, at what speed?

Exercise Solution: Delay becomes distance

- (a) $d = 25 \times 0.25 = 6.25 \text{ m}$.
- (b) Braking: $25^2 / (2 \times 6) = 52.1 \text{ m}$. Normal: $25 \times 0.05 + 52.1 = 53.3 \text{ m}$. Attack: $25 \times 0.3 + 52.1 = 59.6 \text{ m}$.
- (c) Normal: $53.3 \text{ m} < 58 \text{ m}$, **stops**. Attack: $59.6 \text{ m} > 58 \text{ m}$, **collision**. Remaining braking distance $58 - 7.5 = 50.5 \text{ m}$, so $v^2 = 25^2 - 2 \times 6 \times 50.5 = 19$, $v \approx 4.4 \text{ m/s}$ ($\approx 16 \text{ km/h}$).
- A late output is as dangerous as a wrong output when the margin is small.

Exercise: Amplification and deadlines

Pipeline latency $T = 30 \text{ ms}$, deadline $D = 45 \text{ ms}$. A stage with benign share $f = 5 \%$ of T is amplified by $A = 50$. For sequential stages: $A_{\text{e2e}} = 1 + f(A - 1)$; the deadline is missed only if $A > 1 + (D/T - 1)/f$.

- (a) Compute the threshold on A .
- (b) Compute A_{e2e} and the new latency. Is the deadline missed?
- (c) An admission cap limits A to 5. New A_{e2e} and latency?
- (d) Repeat (a)–(b) if the amplified stage has $f = 1 \%$.

Exercise Solution: Amplification and deadlines

- (a) Threshold = $1 + (45/30 - 1)/0.05 = 1 + 0.5/0.05 = 11$.
- (b) $A_{e2e} = 1 + 0.05 \times 49 = 3.45$; latency = $30 \times 3.45 = 103.5 \text{ ms}$ > 45 ms: **missed** ($50 > 11$).
- (c) $A_{e2e} = 1 + 0.05 \times 4 = 1.2$; latency = **36 ms** $\leq$ 45 ms: **not missed**.
- (d) $f = 1 \%$: threshold = $1 + 0.5/0.01 = 51$ > 50; $A_{e2e} = 1.49$, latency **44.7 ms**: just **not missed**.
- Neither f nor D/T is a property of the model: they belong to the system.

Exercise: Stage vs. pipeline factor

Reported numbers: **SlowTrack**: tracking stage 453.8×, whole camera pipeline 28.4×. **Groundswell**: ≈ 188× more NMS candidates, only 2.7× end-to-end latency. Use $A_{e2e} = 1 + f(A - 1)$.

- (a) Derive the benign share f of the amplified stage for each attack.
- (b) For a pipeline with $T = 30$ ms and $D = 45$ ms, what stage factor A is needed to miss the deadline for each f ? Is the reported A enough?

Attack	A (stage)	A_{e2e}	f	Threshold A
SlowTrack	453.8	28.4		
Groundswell	188	2.7		

Exercise Solution: Stage vs. pipeline factor

- (a) $f = (A_{e2e} - 1)/(A - 1)$: SlowTrack $27.4/452.8 \approx 6.1 \%$; Groundswell $1.7/187 \approx 0.91 \%$.
- (b) Threshold $= 1 + 0.5/f$: SlowTrack ≈ 9.3 ; Groundswell ≈ 56 . Both reported factors (453.8, 188) are far above their thresholds, so both would miss the 45 ms deadline.
- Unamplified stages dilute the effect: stage-level factors overstate pipeline-level risk, but a large enough factor still breaks the deadline.

Attack	A (stage)	A _{e2e}	f	Threshold A
SlowTrack	453.8	28.4	6.1 %	9.3
Groundswell	188	2.7	0.91 %	56

Exercise: Tracker pool accumulation

A tracker has memory: an unmatched false track survives for $R = 30$ frames after birth. The scene has **20 real objects**. The attacker creates **10 fake detections per frame**. Association cost $\propto M \times N$ (M tracks, N detections).

- (a) Steady-state track pool M under attack, and N ?
- (b) Association cost vs. the benign case ($M = N = 20$)?
- (c) An admission cap allows at most **2 track births per frame**. New M and cost?

Case	M (tracks)	N (detections)	$M \times N$	vs. benign
Benign	20	20	400	1×
Attack, no cap				
Attack, 2 births/frame				

Exercise Solution: Tracker pool accumulation

- (a) $M = 20 + R \times 10 = 20 + 300 = 320$ tracks; $N = 20 + 10 = 30$ detections.
- (b) $320 \times 30 = 9,600$ vs. $20 \times 20 = 400$: **24×** more work.
- (c) $M = 20 + 30 \times 2 = 80$; cost = $80 \times 30 = 2,400$ (6× benign): **4× lower** than uncapped.
- A per-frame cap on detections bounds one dimension only; the persistent pool is bounded by capping track **admission** (TrackShield).

Case	M (tracks)	N (detections)	M × N	vs. benign
Benign	20	20	400	1×
Attack, no cap	320	30	9,600	24×
Attack, 2 births/frame	80	30	2,400	6×