Artificial Intelligence Methods for Social Good

M2-2 [Game Theory]:

Security Games

Fei Fang
feifang@cmu.edu
Wean Hall 4126

Quiz 1: Recap: Nash Equilibrium

- In Rock-Paper-Scissors, which of the following is a Nash Equilibrium?
  - $s_1 = (1,0,0), s_2 = (1,0,0)$
  - $s_1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), s_2 = (1,0,0)$
  - $s_1 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), s_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
  - $s_1 = (1,0,0), s_2 = (0,1,0)$

Player 2

Player 1

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissor	-1,1	1,-1	0,0

Quiz 2: Recap: Strong Stackelberg Equilibrium

- In Power of Commitment, what is player 1's utility in Strong Stackelberg Equilibrium?
  - 3.75
  - 2
  - $\frac{11}{3}$
  - 3.5

Player 2

Player 1

	c	d
a	2,1	4,0
b	1,0	3,2

Societal Challenges: Security and Sustainability

Ansbach attack

A suicide bomb injured at least 12 in Germany's Ansbach, near Nuremberg, on July 24. This is the fourth violent incident in Germany in a week.

Source: Reuters

J. Wu, 25/07/2016

REUTERS

Societal Challenges: Security and Sustainability

Societal Challenges: Security and Sustainability

Physical Infrastructure

Transportation Networks

Cyber Systems

Environmental Resources

Endangered Wildlife

Fisheries

Societal Challenges: Security and Sustainability

- Improve tactics of patrol, inspection, screening etc

Game Theoretic Reasoning

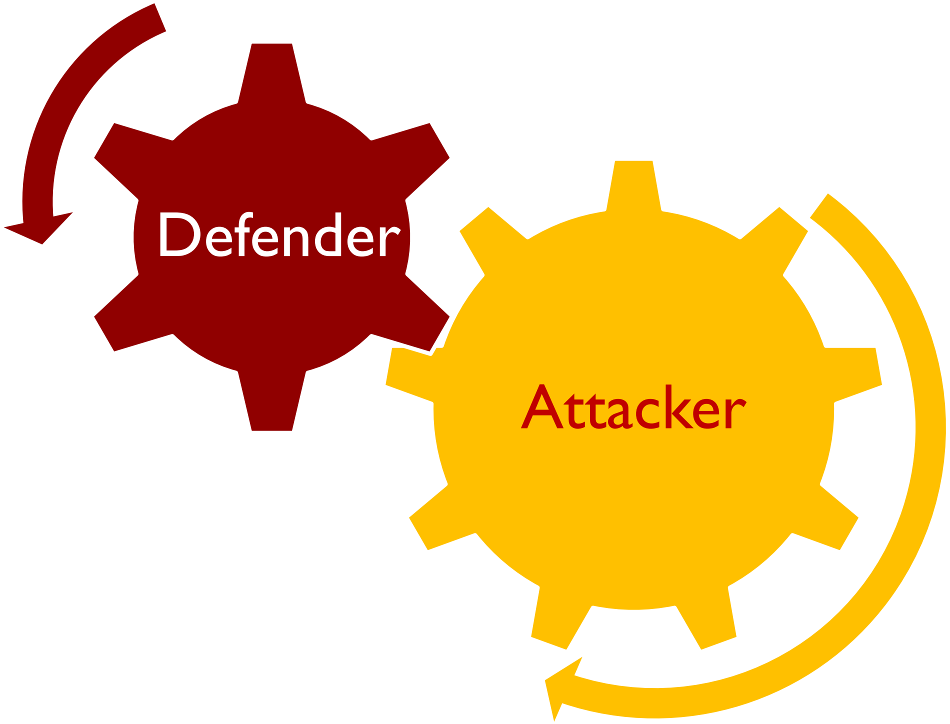

Defender

Attacker

Outline

- Basic model
- Deal with continuous timeline
- Fine-grained planning with practical constraints

Model Security Problem as a Stackelberg Game

- Limited resource allocation
- Adversary surveillance

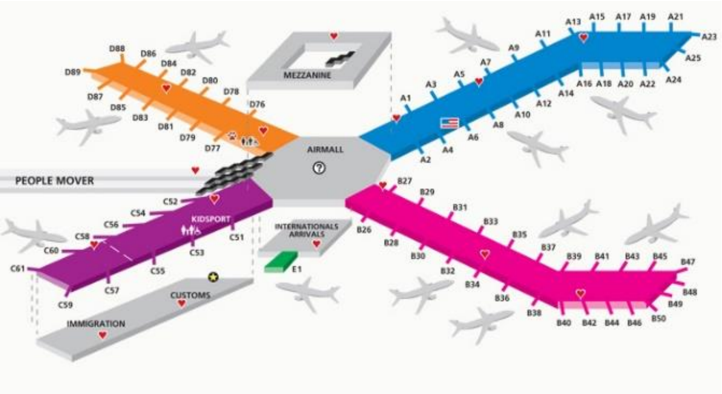

Adversary

Defender

	Target #1	Target #2
Target #1	5, -3	-1, 1
Target #2	-5, 4	2, -1

Model Security Problem as a Stackelberg Game

- Limited resource allocation
- Adversary surveillance

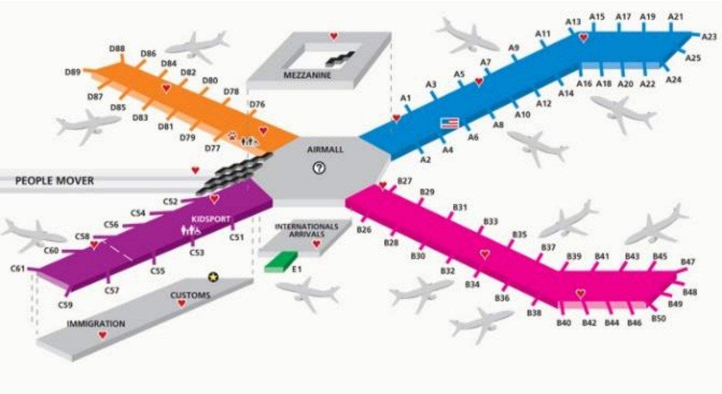

Adversary

Defender

	Target #1	Target #2
Target #1	5, -3	-1, 1
Target #2	-5, 4	2, -1

Model Security Problem as a Stackelberg Game

- Randomization make defender unpredictable
- Stackelberg Security game
  - Defender: Commits to mixed strategy
  - Adversary: Conduct surveillance and best responds

Adversary

Defender

		Target #1	Target #2
55.6%	**Target #1**	5, -3	-1, 1
44.4%	**Target #2**	-5, 4	2, -1

Model Security Problem as a Stackelberg Game

- Strong Stackelberg Equilibrium
  - Attacker break tie in favor of defender
  - AttEU1=0.556*(-3)+0.444*4=0.11
  - AttEU2=0.556*1+0.444*(-1)=0.11
  - DefEU1=0.556*5+0.444*(-5)=0.56
  - DefEU2=0.556*(-1)+0.444*2=0.332
  - Equilibrium: DefStrat=(0.556,0.444), AttStrat=(1,0)

Adversary

Defender

		Target #1	**Target #2**
55.6%	**Target #1**	5, -3	-1, 1
44.4%	**Target #2**	-5, 4	2, -1

Computing SSE

- General-sum
  - Multiple LP or MILP
  - Assume attacks target i^*

$$\min_{p_1,p_2,\dots,p_N} \quad AttEU(i^*)$$
$$\text{s.t. } AttEU(i^*) \geq AttEU(i), \forall i = 1 \dots N$$
$$\sum_i p_i \leq 1$$
$$AttEU(i) = p_i P_i^a + (1 - p_i) R_i^a$$

- Zero-sum
  - Single LP
  - SSE=NE

$$\min_{p_1,p_2,\dots,p_N} \quad v$$
$$\text{s.t. } v \geq AttEU(i), \forall i = 1 \dots N$$
$$\sum_i p_i \leq 1$$

Adversary

		Target #1	Target #2
55.6%	**Target #1**	5, -3	-1, 1
44.4%	**Target #2**	-5, 4	2, -1

Defender

Compute optimal defender strategy

- Polynomial time solvable in games with finite actions and simple structures [Conitzer06]
- NP-Hard in general settings [Korzhyk10]
- SSE=NE for zero-sum games, SSE⊂NE for games with special properties [Yin10]

- Research Challenges
  - Massive scale games with constraints
  - Plan/reason under uncertainty
  - Repeated interaction

Outline

- Basic model
- Deal with continuous timeline
- Fine-grained planning with practical constraints

Game Theoretic Reasoning

Problem

- Optimize the use of patrol resources
  - Moving targets: Fixed schedule
  - Potential attacks: Any time
  - Continuous time

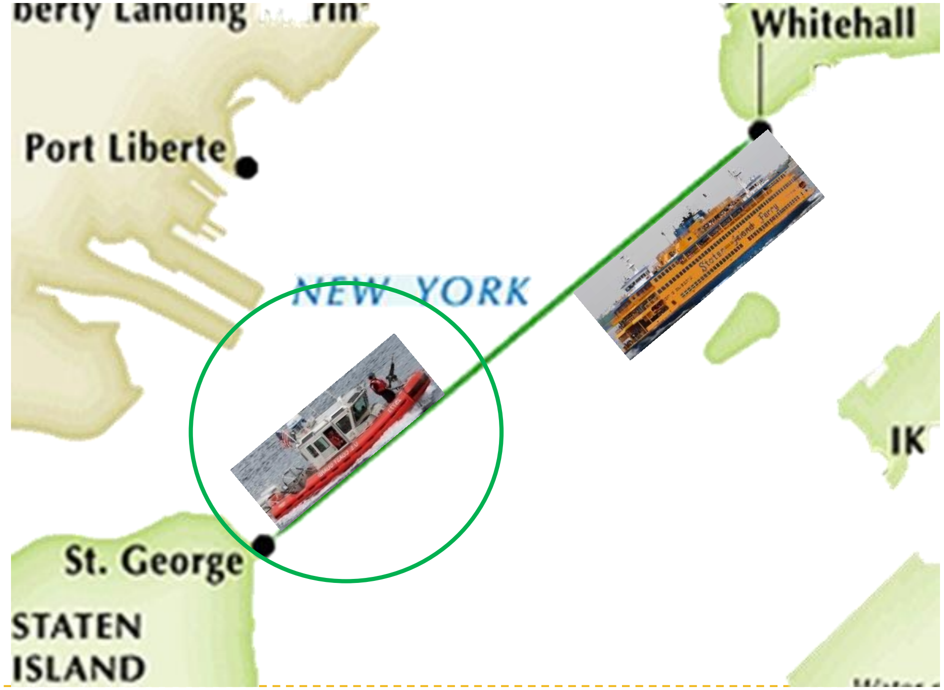

Model

- Attacker: Which target, when to attack
- Defender: Choose a route for patrol boat
- Payoff value for attacker: $u_i(t)$ if not protected, 0 if protected
- Minimax: Minimize attacker's expected utility assume attacker best responds

Attacker's Expected Utility = Target Utility × Probability of Success

Adversary

Defender

		10:00:00 AM **Target 1**	**10:00:01 AM** **Target 1**	…	**10:30:00 AM** **Target 3**	…
30%	Purple Route	-5, 5	-4, 4		0, 0	
40%	Orange Route					
20%	Blue Route					
	……					

HOW TO FIND OPTIMAL DEFENDER STRATEGY

- Step 1: Compact representation for defender

Adversary

Defender

	10:00:00 AM Target 1	**10:00:01 AM Target 1**	…	**10:30:00 AM Target 3**	…
Purple Route	-5, 5	-4, 4		0, 0	
Orange Route					
Blue Route					
……					

STEP 1: COMPACT REPRESENTATION FOR DEFENDER

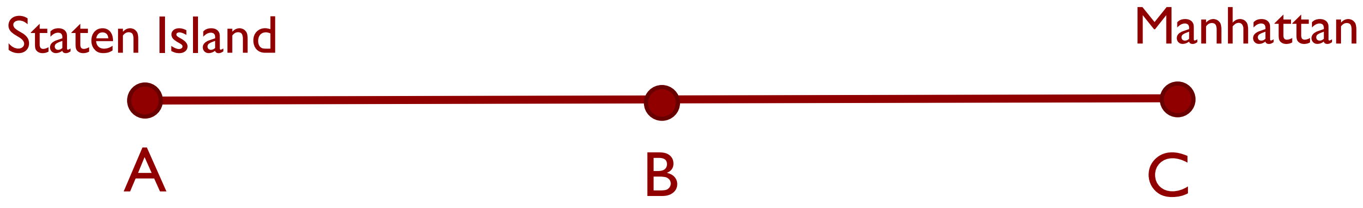

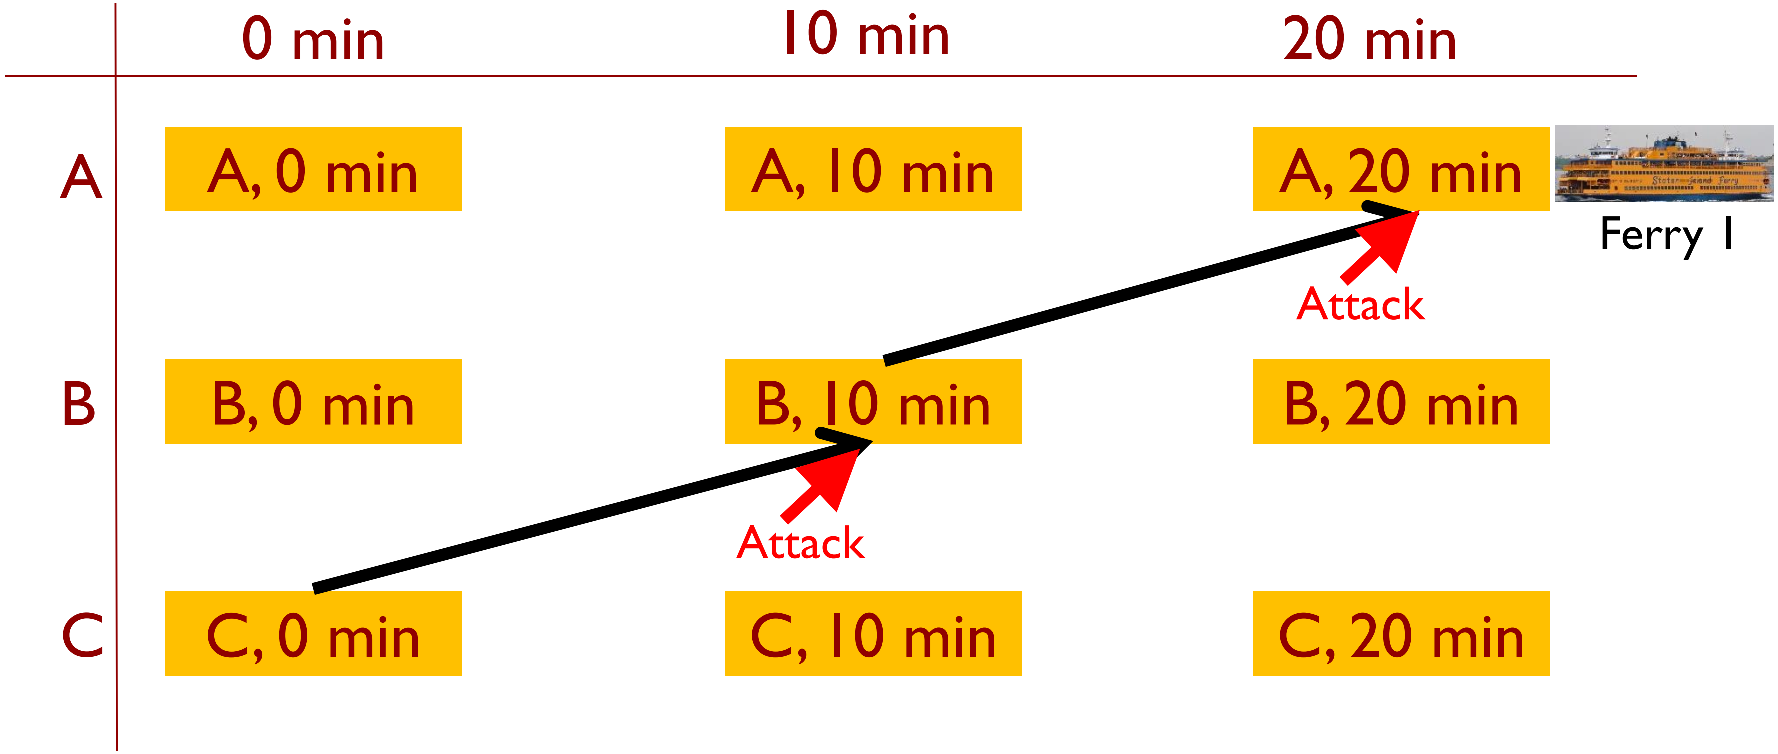

	0 min	10 min	20 min
A	A, 0 min	A, 10 min	A, 20 min
B	B, 0 min	B, 10 min	B, 20 min
C	C, 0 min	C, 10 min	C, 20 min

STEP 1: COMPACT REPRESENTATION FOR DEFENDER

- Full representation: Focus on routes (N^T)
  - Prob(Orange Route) = 0.37 Prob(Green Route) = 0.33
  - Prob(Blue Route) = 0.17 Prob(Purple Route) = 0.13

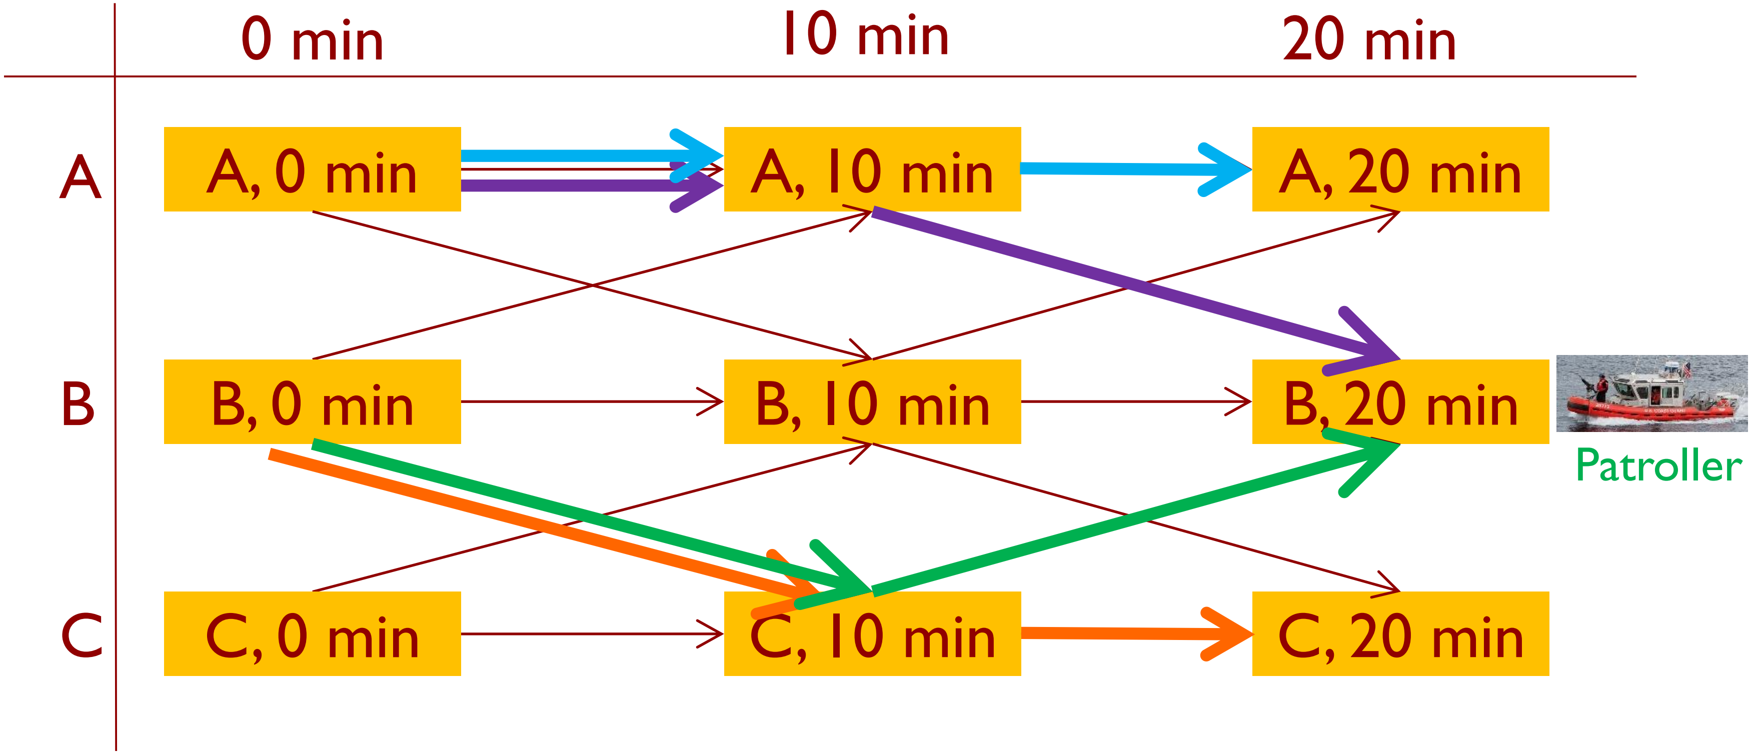

STEP I: COMPACT REPRESENTATION FOR DEFENDER

- Full representation: Focus on routes (N^T)
  - Prob(Orange Route) = 0.37 Prob(Green Route) = 0.33
  - Prob(Blue Route) = 0.17 Prob(Purple Route) = 0.13
- Linear program

$$\min_{\boxed{p_1,p_2,\ldots,p_R}} v$$

s.t. $\boxed{v \geq AttEU(i,\hat{t}),}$

For all target i, time point $\hat{t}$

Probability of route (N^T variables)

Best response

STEP I: COMPACT REPRESENTATION FOR DEFENDER

- **Compact representation**: Focus on edges (N^2T)
  - Probability flow over each edge

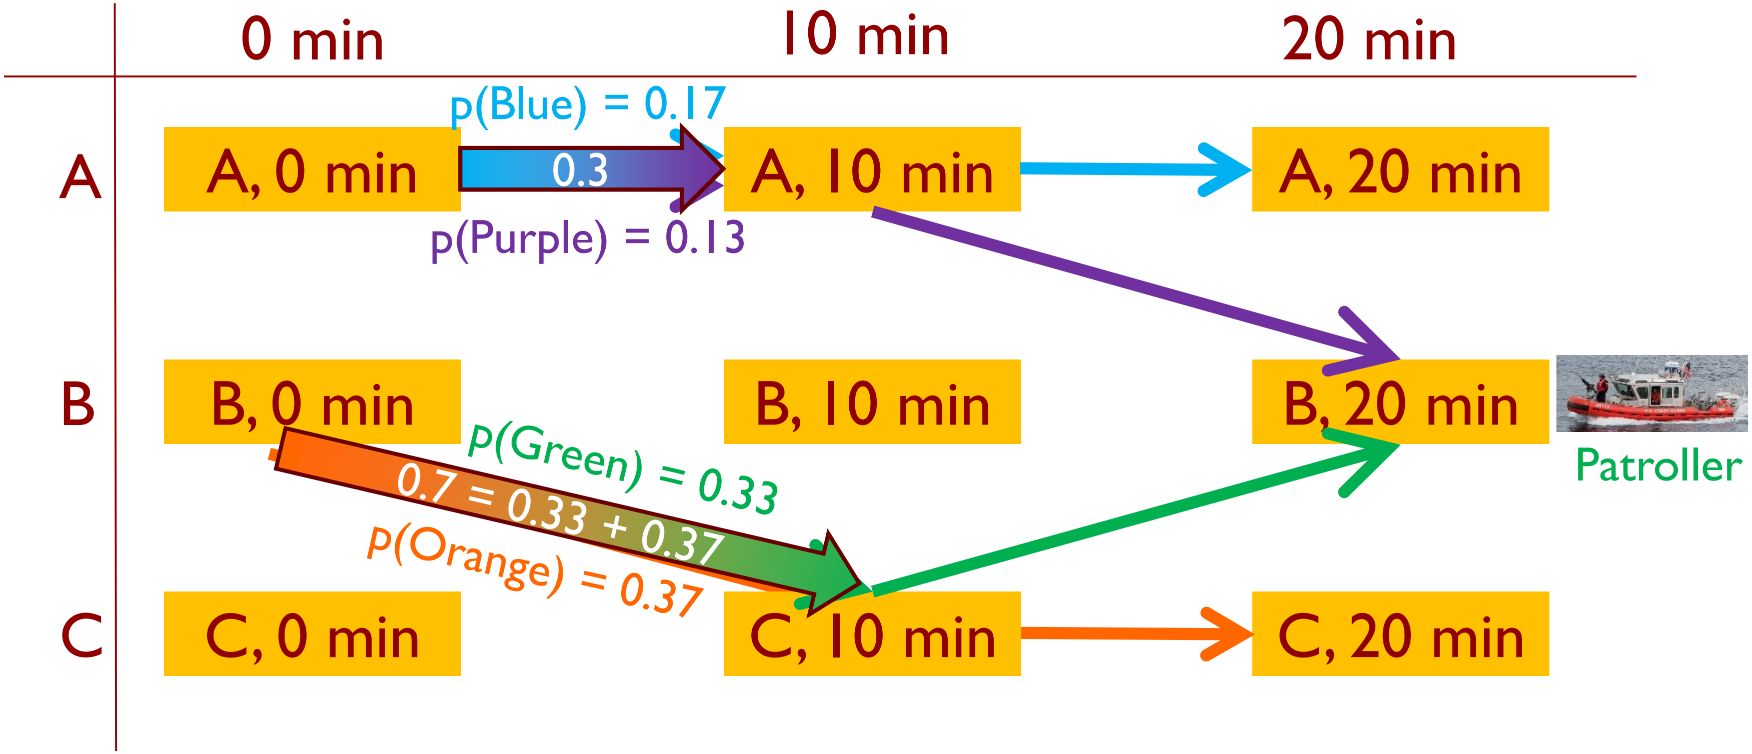

STEP 1: COMPACT REPRESENTATION FOR DEFENDER

- **Theorem 1**: Let p, p' be two defender strategies in full representation, and the compact representation for both strategies is f, then $AttEU_p^i(t) = AttEU_{p'}^i(t)$, and $DefEU_p^i(t) = DefEU_{p'}^i(t), \forall t$

- Compact representation does not lead to any loss

Quiz 3: Deal with Continuous Timeline

- How many variables are needed to compute the optimal defender strategy in compact representation?
  - A: $O(N^2T)$
  - B: $O(N^T)$
  - C: $O(NT^2)$
  - D: $O(NT)$

HOW TO FIND OPTIMAL DEFENDER STRATEGY

- Step I: Compact representation for defender
- Step II: Compact representation for attacker

Adversary

Defender

	10:00:00 AM Target 1	**10:00:01 AM Target 1**	…	**10:30:00 AM Target 3**	…
Purple Route	-5, 5	-4, 4		0, 0	
Orange Route					
Blue Route					
……					

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- **Partition attacker action set**
- Only need to reason about a few attacker actions

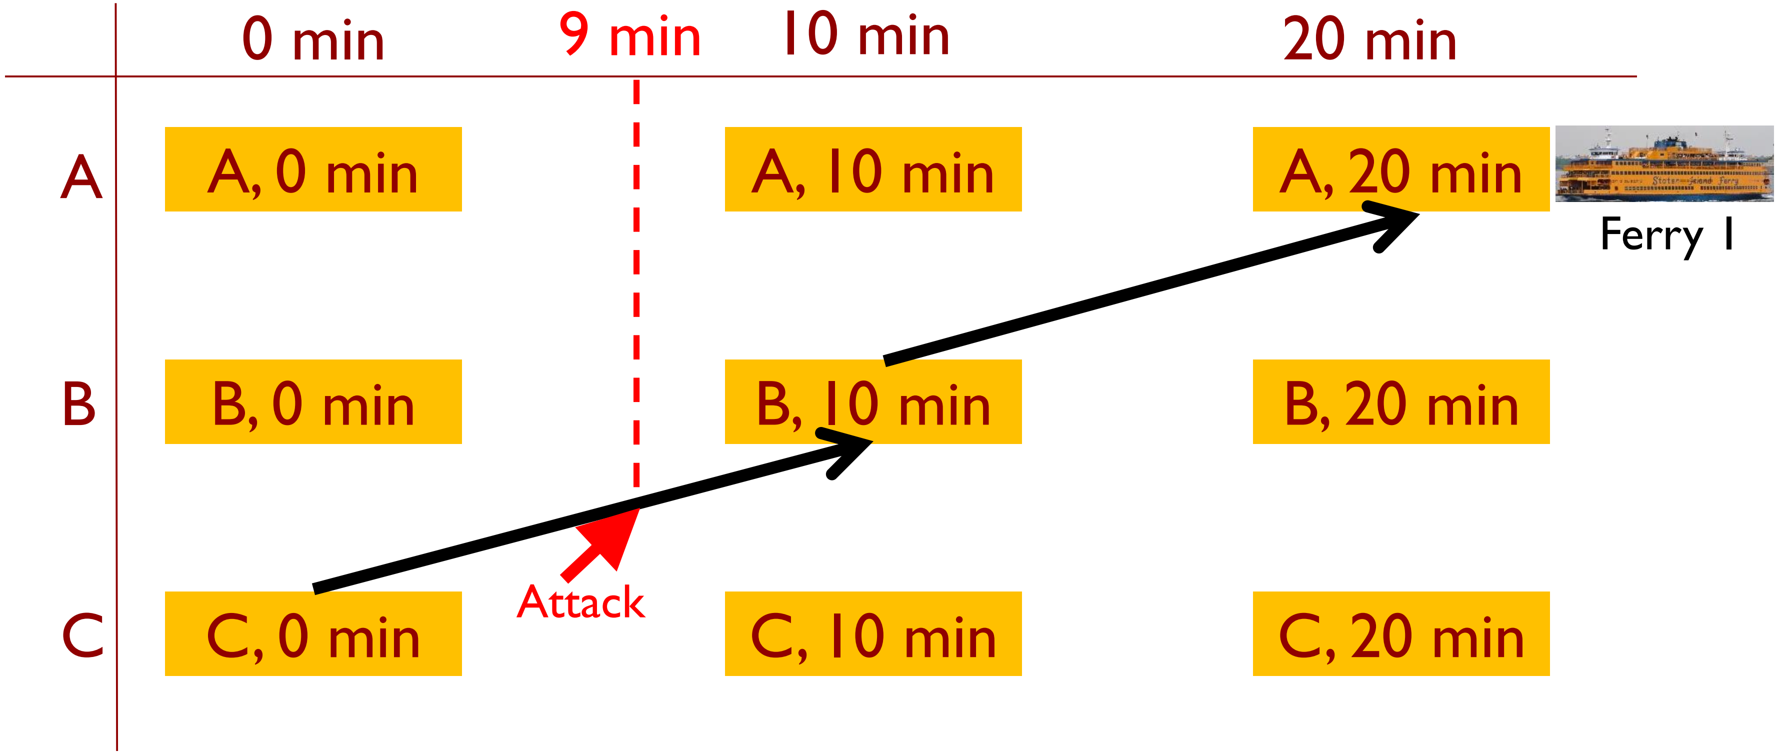

	0 min	9 min	10 min	20 min
A	A, 0 min		A, 10 min	A, 20 min
B	B, 0 min		B, 10 min	B, 20 min
C	C, 0 min		C, 10 min	C, 20 min

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- Partition points θ^k: When protection status changes

Unprotected

Enter θ^1

Protected

Leave θ^2

Unprotected

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- Partition points θ^k: When protection status changes

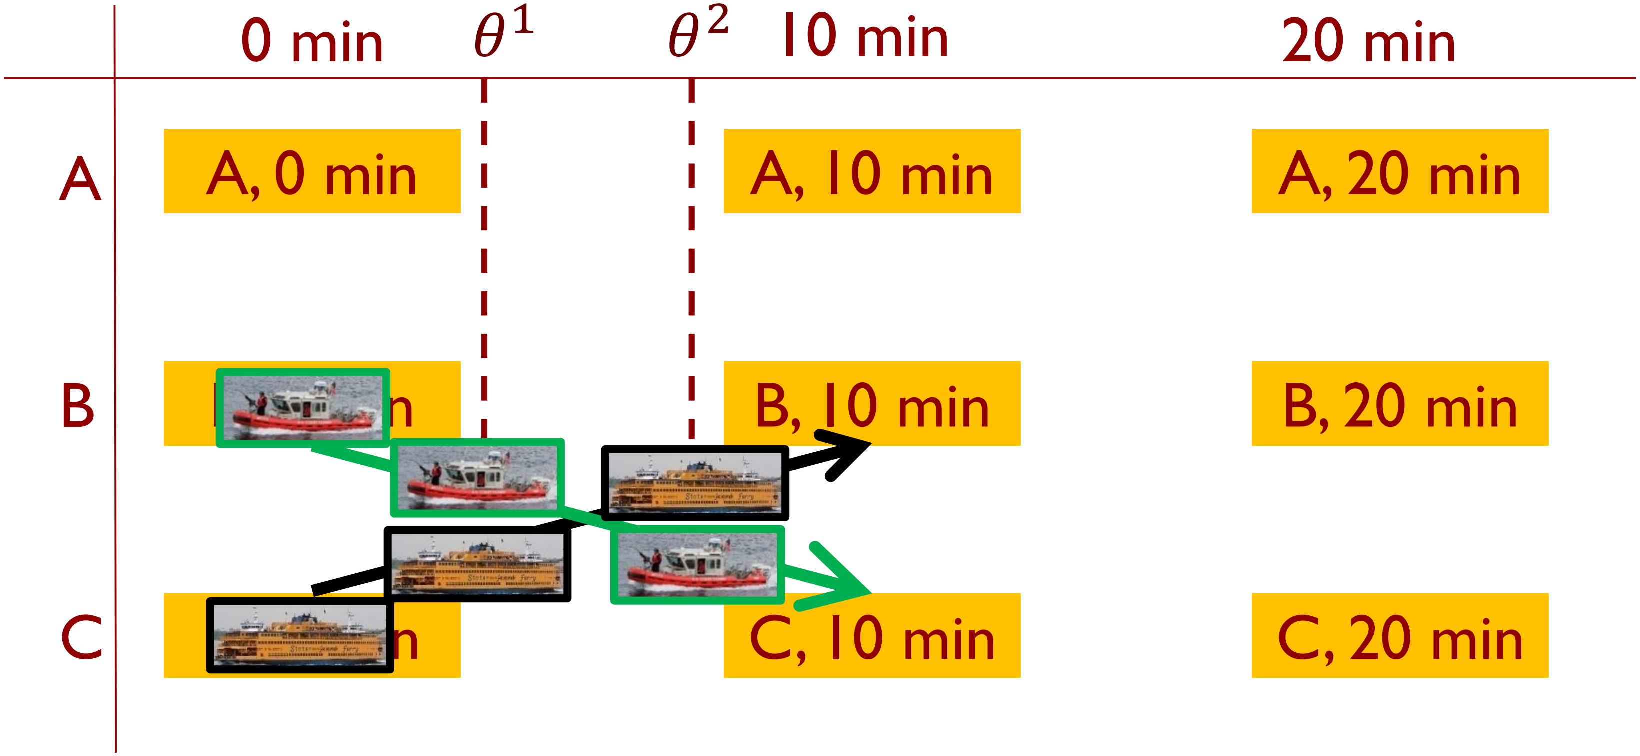

0 min θ^1 θ^2 10 min 20 min

A

A, 0 min

A, 10 min

A, 20 min

B

B, 10 min

B, 20 min

C

C, 10 min

C, 20 min

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- $AttEU = \boxed{\text{Target Utility(t)}} \times \boxed{\text{Probability of Success}}$
- **One best time point in each zone**

Fixed

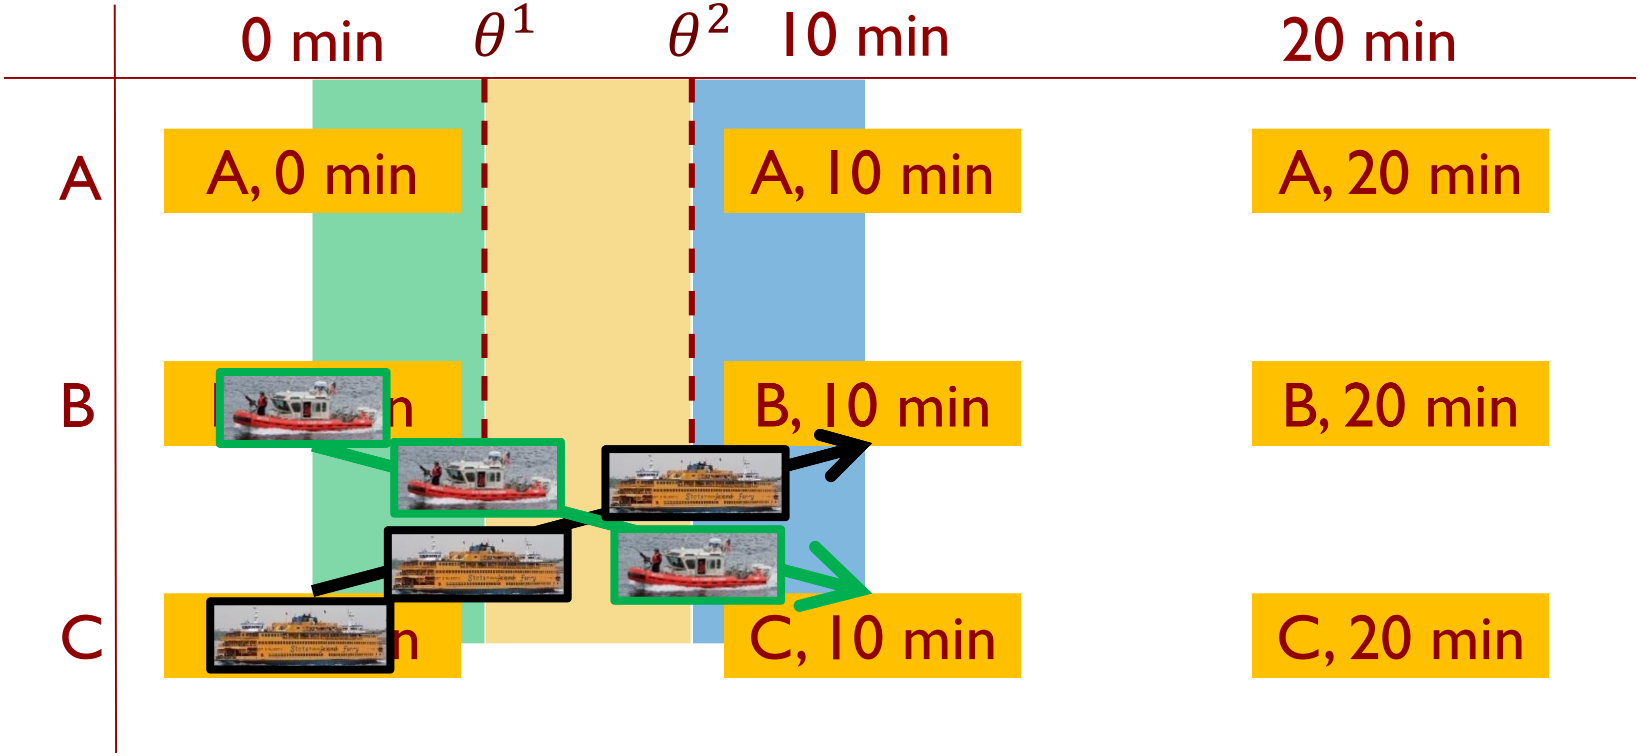

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- $AttEU = \boxed{\text{Target Utility(t)}} \times \boxed{\text{Probability of Success}}$
- **One best time point in each zone**

Fixed

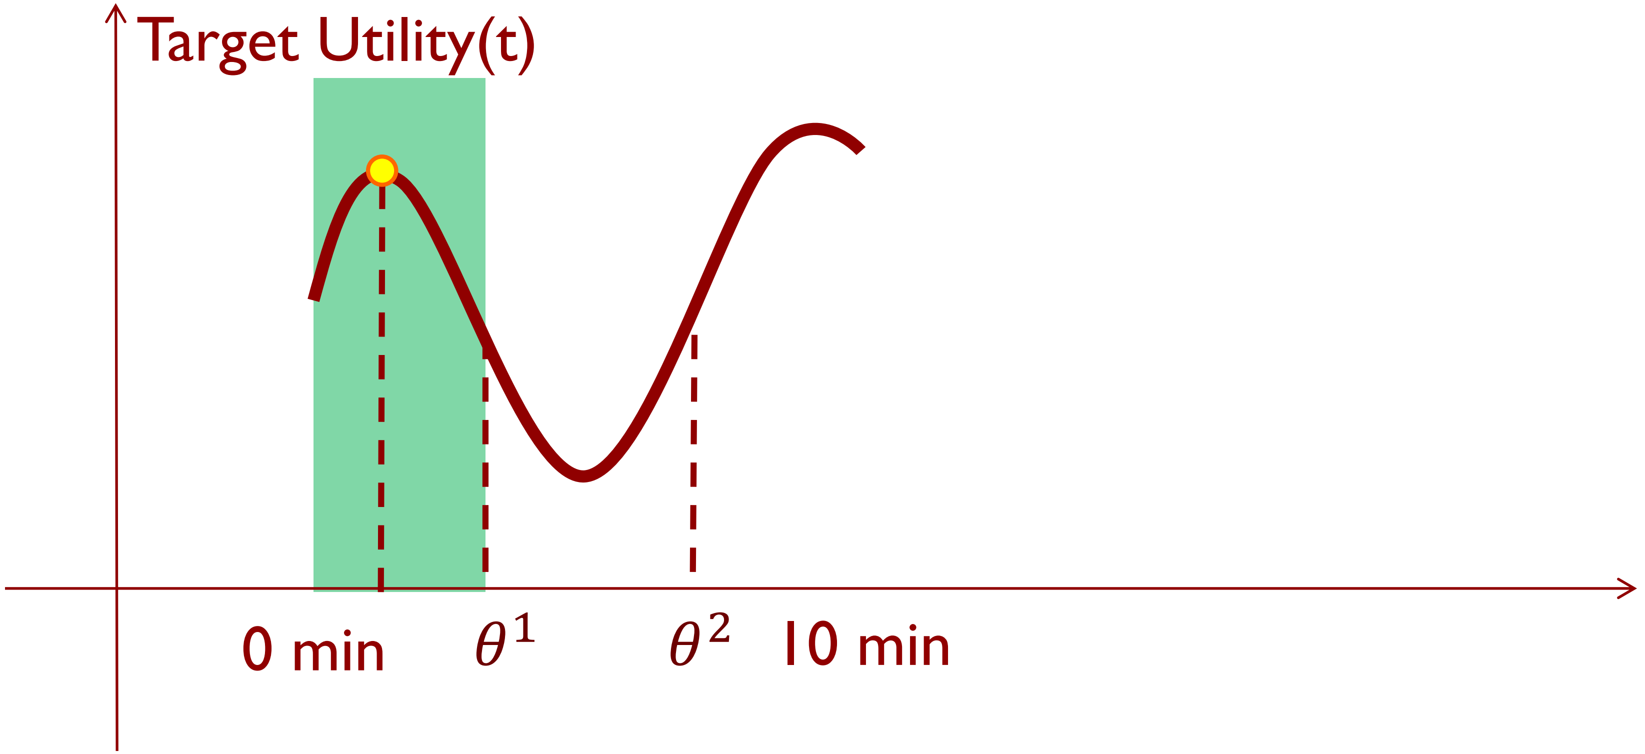

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- $AttEU = \boxed{\text{Target Utility(t)}} \times \boxed{\text{Probability of Success}}$
- **One best time point in each zone**

Fixed

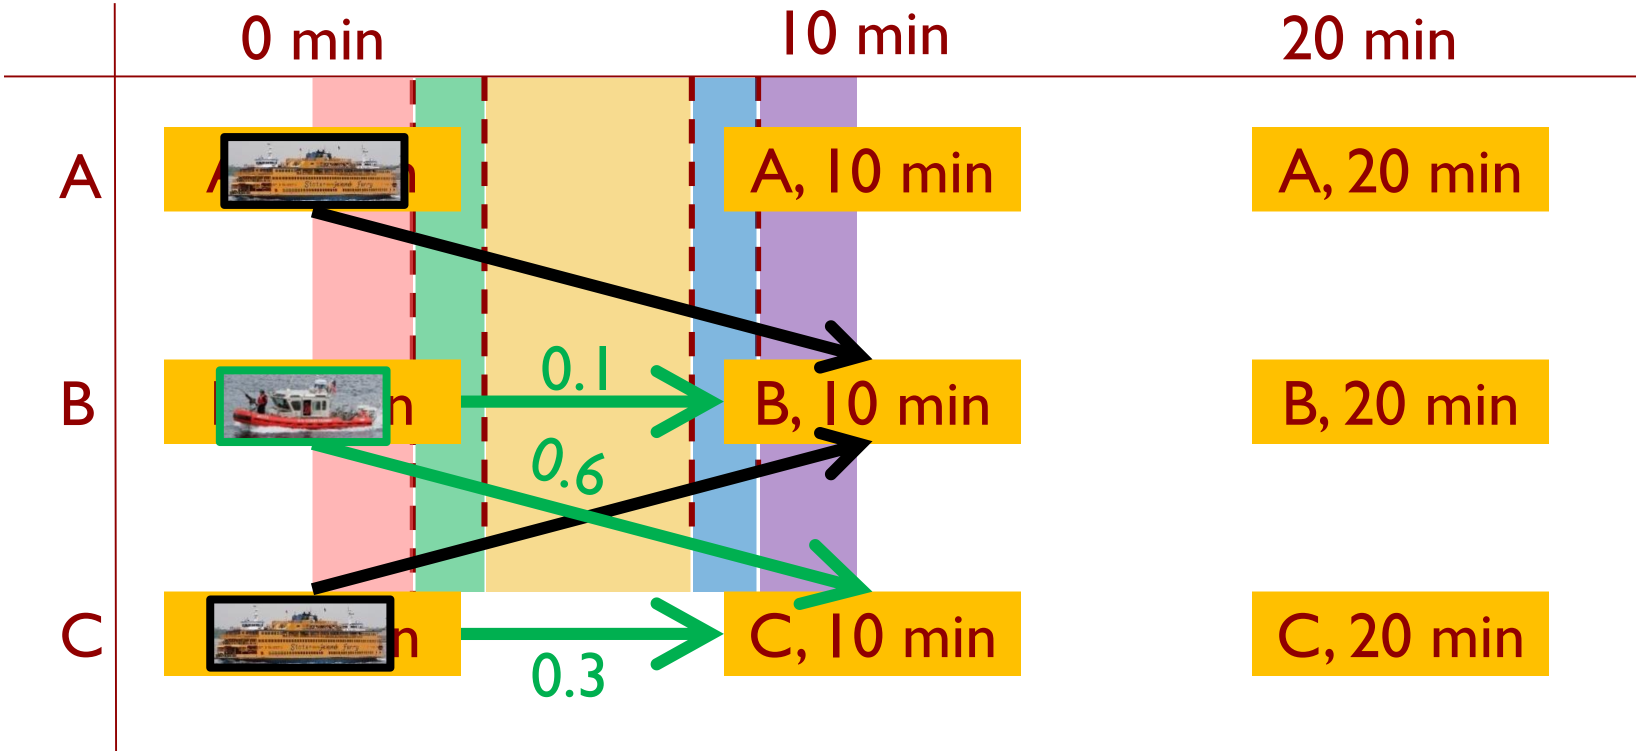

STEP II: COMPACT REPRESENTATION FOR ATTACKER

- **Theorem 2**: Given target utility function $u_i(t)$, assume the defender's pure strategy is restricted to be a mapping from $\{\hat{\boldsymbol{t}}\}$ to $\{\widehat{\boldsymbol{d}}\}$, then in the attacker's best response, attacking time $t^* \in \{\boldsymbol{t}^*\} = \{t | \exists i, j \text{ such that } t = arg \max_{t' \in [\theta_j, \theta_{j+1}]} u_i(t')\}$

- Only considering the best time points does not lead to any loss when attacker best responds

- $\infty \rightarrow O(N^2 T)$

HOW TO FIND OPTIMAL DEFENDER STRATEGY

- Step I: Compact representation for defender
- Step II: Compact representation for attacker
- Step III: Consider infinite defender action set
- Step IV: Equilibrium refinement

EVALUATION: SIMULATION RESULTS

- Randomly chosen utility function
- Attacker's expected utility (lower is better)

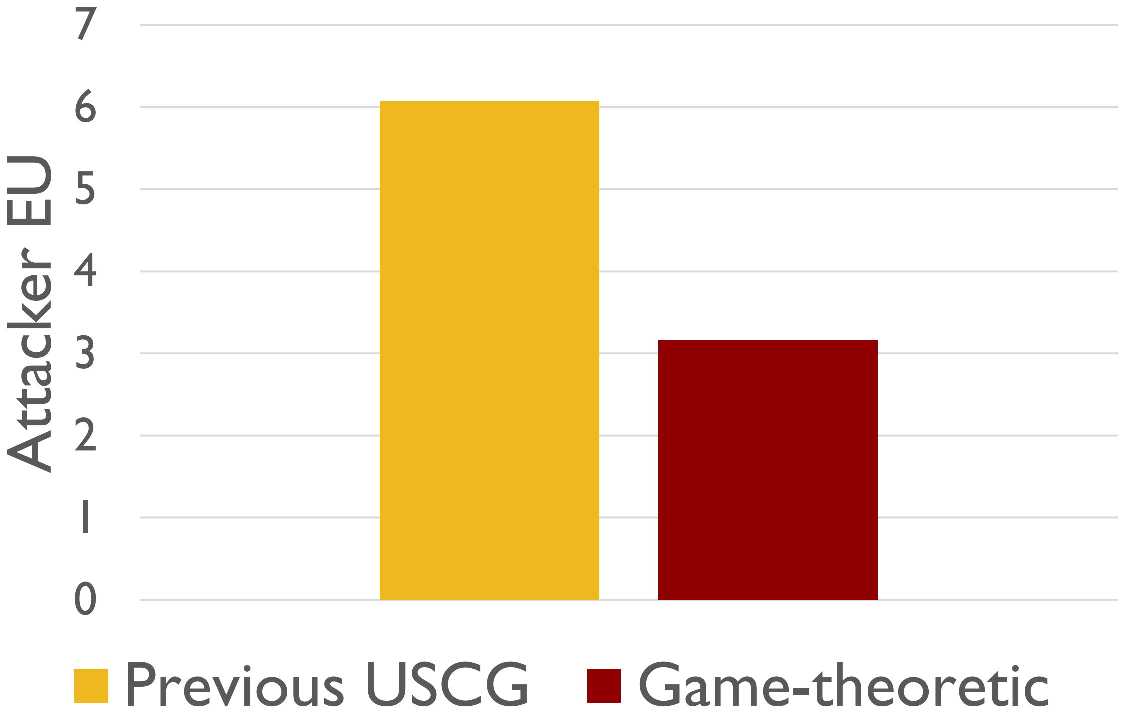

EVALUATION: FEEDBACK FROM REAL-WORLD

- US Coast Guard evaluation
  - Point defense to zone defense
  - Increased randomness
  - Mock attacker

- Patrollers feedback
  - More dynamic (speed change, U-turn)

- Professional mariners' observation
  - Apparent increase in Coast Guard patrols

- Used by USCG (without being forced)

PUBLIC FEEDBACK

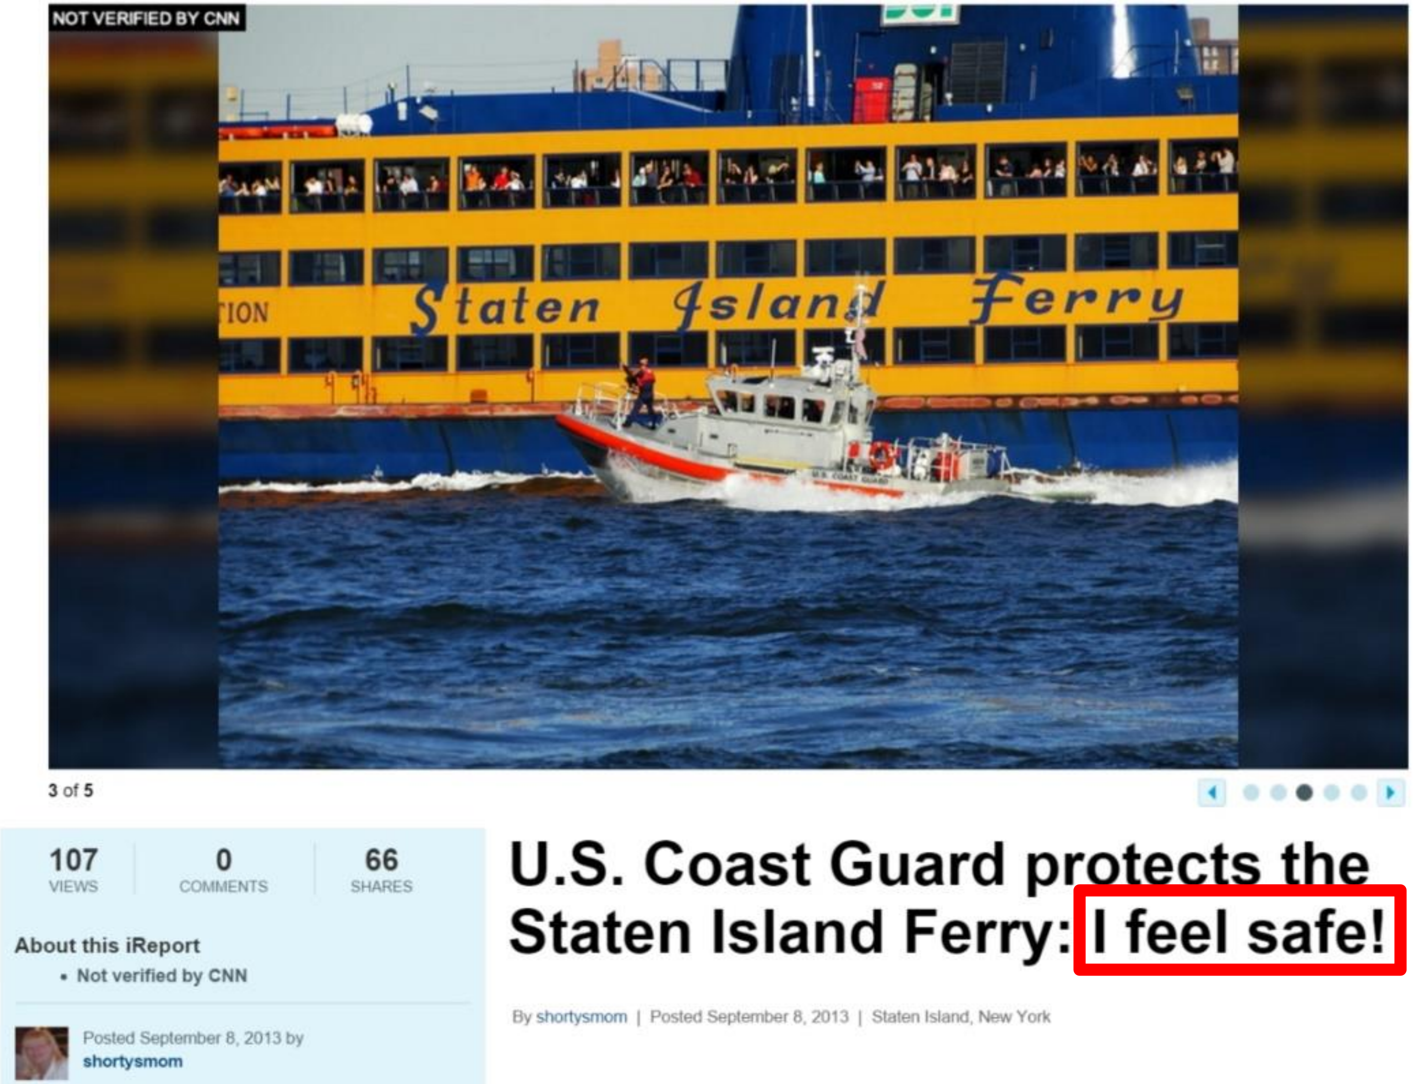

3 of 5

107	0	66
VIEWS	COMMENTS	SHARES

About this iReport

- Not verified by CNN

Posted September 8, 2013 by **shortysmom**

U.S. Coast Guard protects the Staten Island Ferry: I feel safe!

By shortysmom | Posted September 8, 2013 | Staten Island, New York

EXTEND TO 2-D NETWORK

- Complex ferry system: Seattle, San Francisco
- Calculate partition points in 3D space

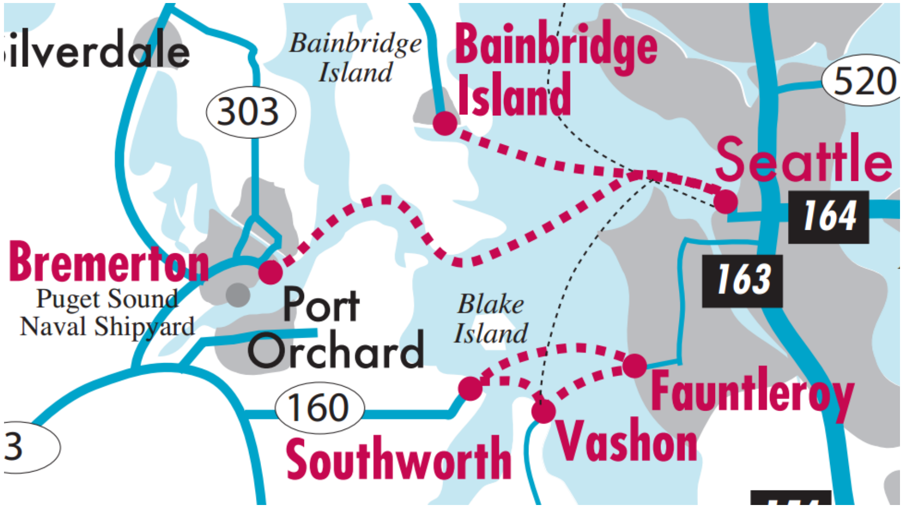

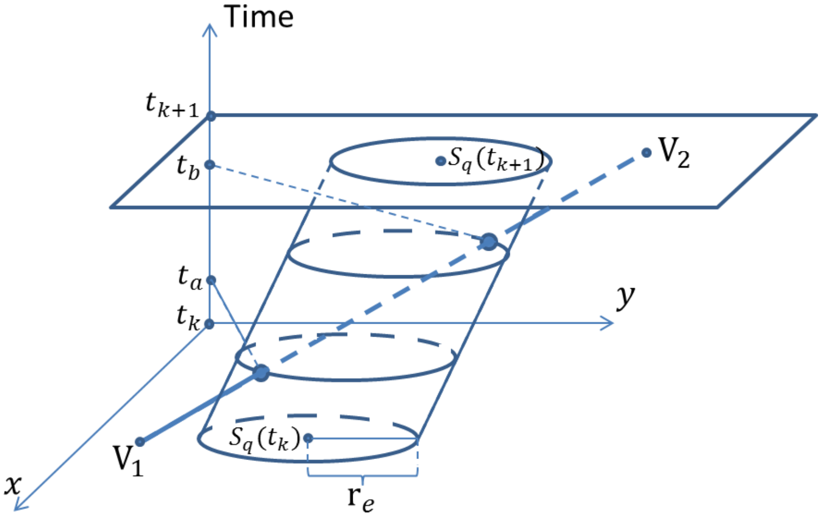

Outline

- Basic model
- Deal with continuous timeline
- Fine-grained planning with practical constraints

Fine-Grained Planning

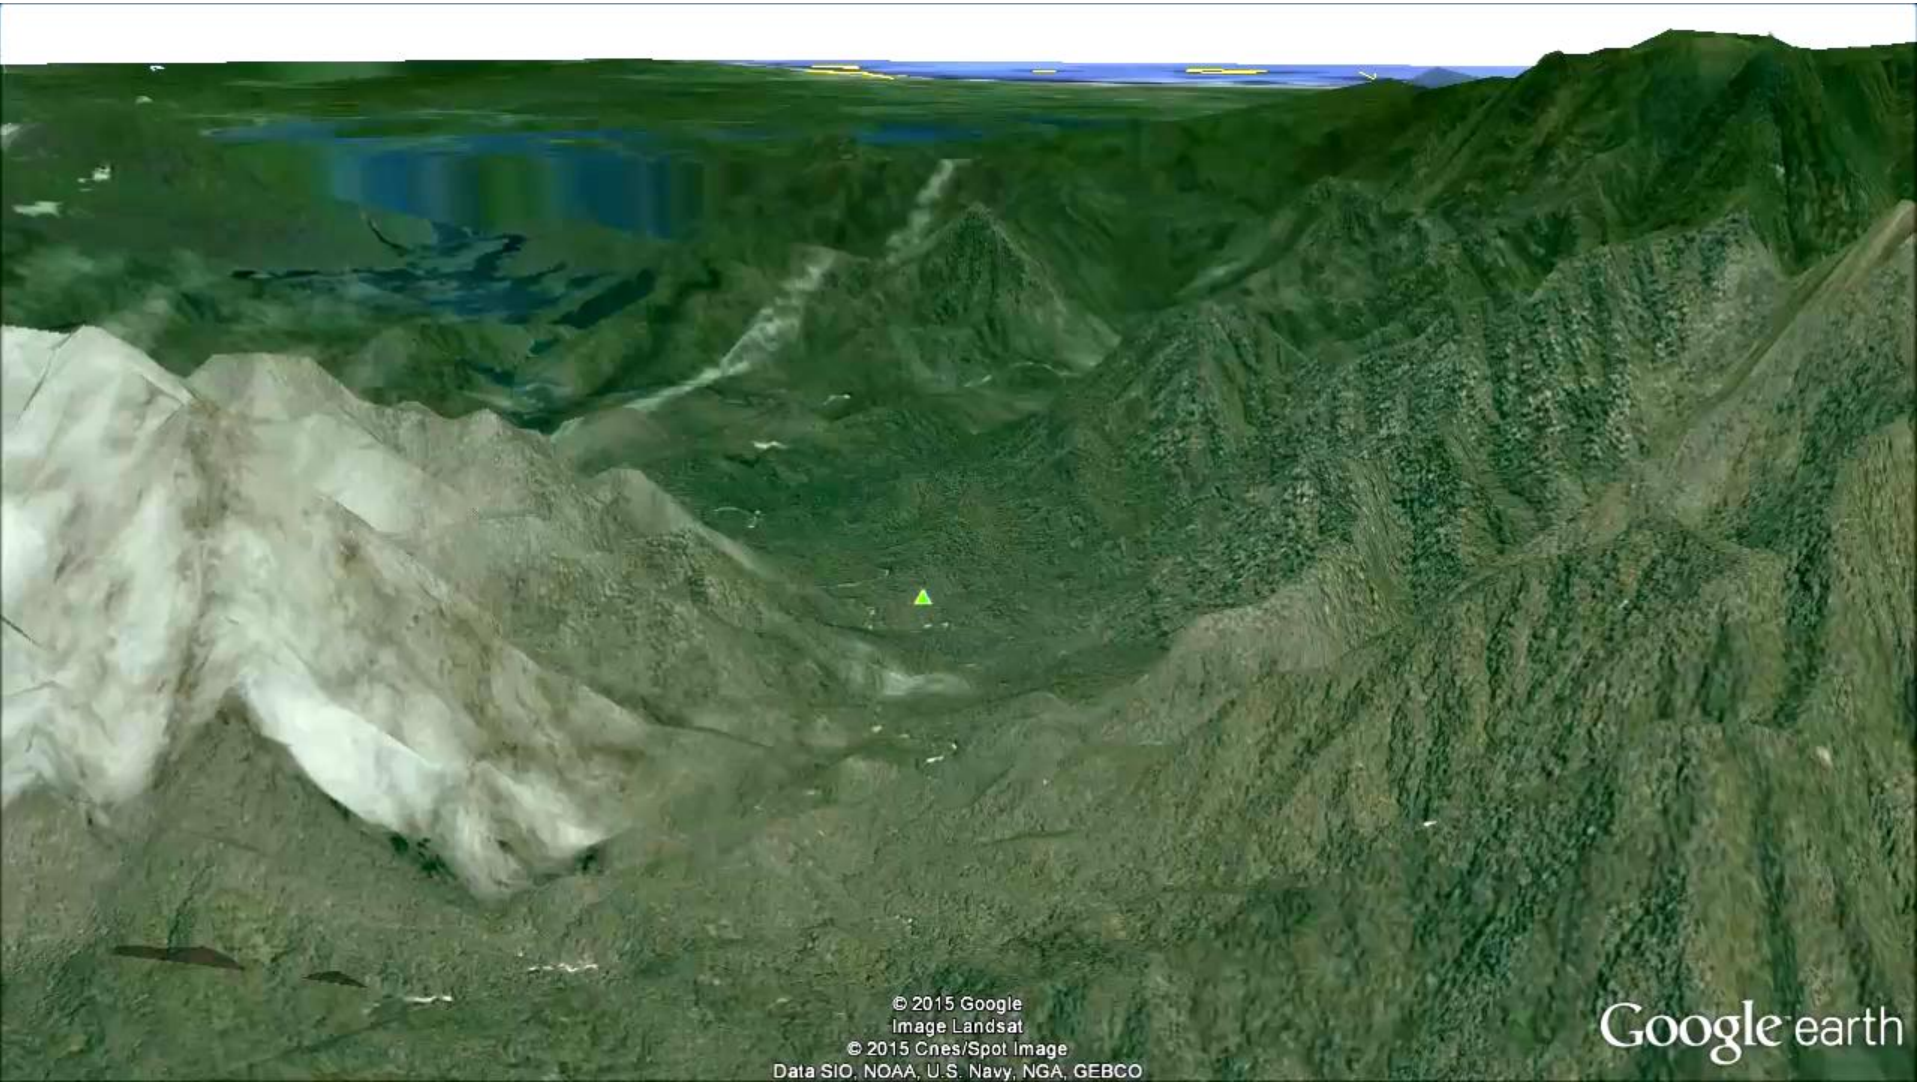

Fine-Grained Planning

(Not) Fine-Grained Planning

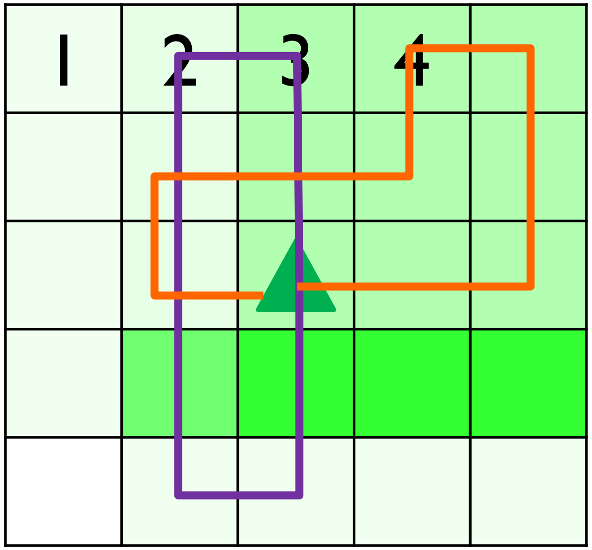

- Animal density (utility) represented by color
- Max patrol length=10
- Attack two cells

Adversary

Defender

		Cell1&2	Cell 2&3	…	Cell 3&4	…
30%	Purple Route	-2, 2	0, 0		-5, 5	
40%	Orange Route					
20%	Blue Route					
	……					

(Not) Fine-Grained Planning

- Option 1: Go back to time-location graph
  - Only apply to integer-valued distance
  - Generalizable to general-sum games

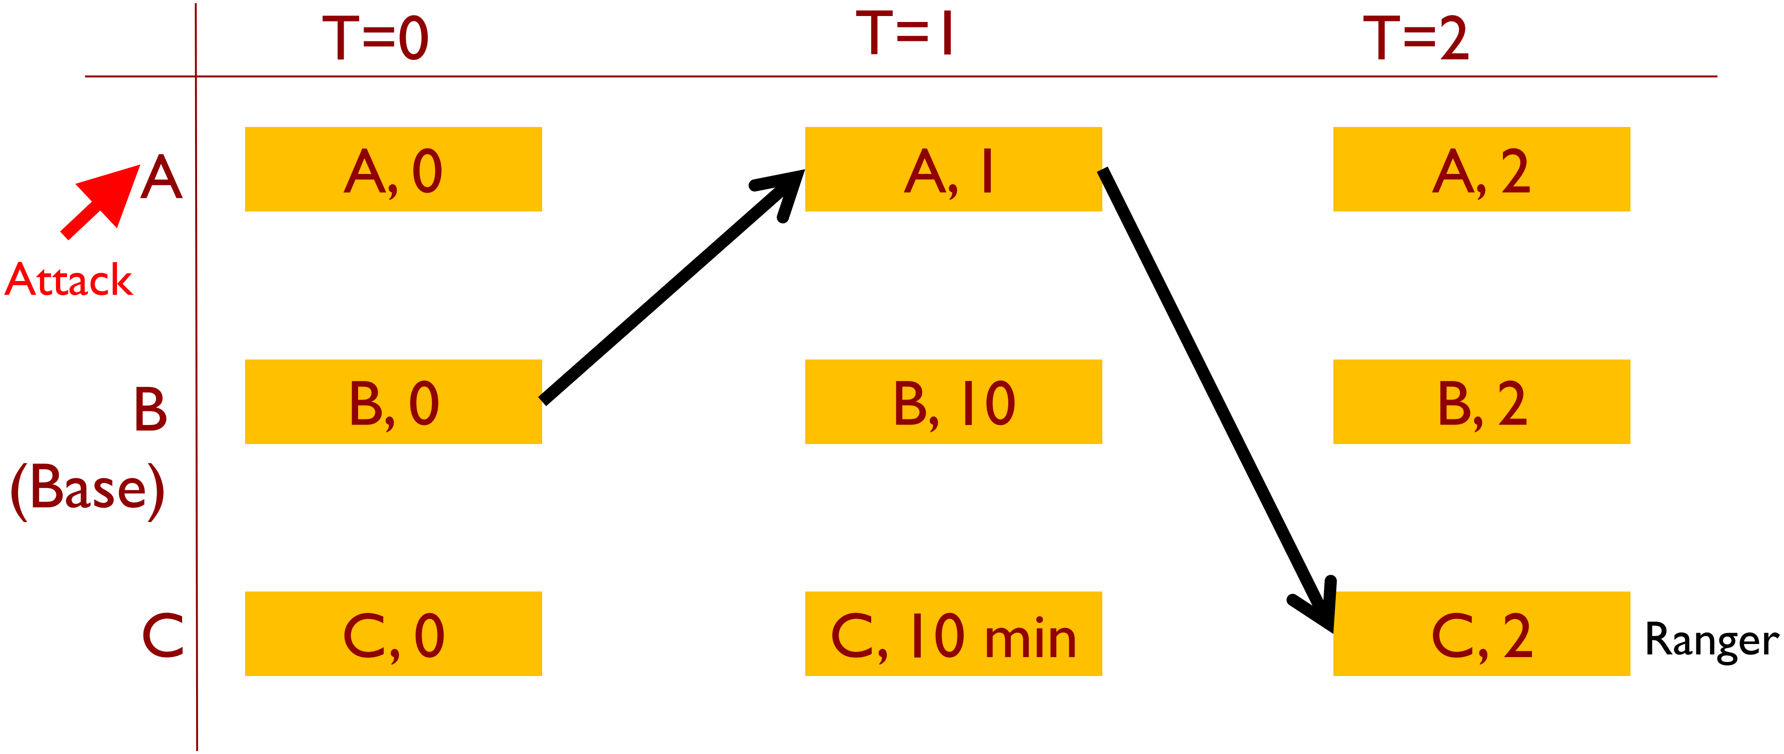

	T=0	T=1	T=2
A	A, 0	A, 1	A, 2
B (Base)	B, 0	B, 10	B, 2
C	C, 0	C, 10 min	C, 2

(Not) Fine-Grained Planning

- Option 1: Go back to time-location graph
  - Only apply to integer-valued distance
  - Generalizable to general-sum games
- Option 2: Incremental strategy generation
  - Generalizable to fine-grained planning
  - Only apply to zero-sum games

Incremental Strategy Generation

- Start with a subset of actions for each player
- Compute NE strategy for both players
  - In zero-sum games, SSE=NE for defender
- Fix attacker strategy, compute best route for defender among all possible routes (coin collection problem), add to the matrix
- Fix defender strategy, compute best cells for attacker among all possible choices (greedy), add to the matrix
- Re-compute NE
- Repeat until best responses already in the matrix

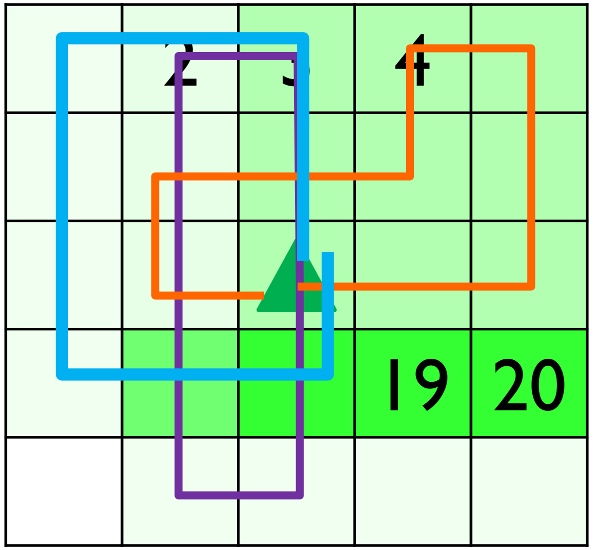

Adversary

Defender

		60%	40%	
		Cell1&2	**Cell 2&3**	**Cell 18&19**
30%	Purple Route	-2, 2	0, 0	
70%	Orange Route			
	Blue Route			

(Not) Fine-Grained Planning

- Option 1: Go back to time-location graph
  - Only apply to integer-valued distance
  - Generalizable to general-sum games
- Option 2: Incremental strategy generation
  - Generalizable to fine-grained planning
  - Only apply to zero-sum games
- Option 3: Cutting plane
  - Generalizable to fine-grained planning
  - Generalizable to general-sum games

Cutting Plane

- Focus on the coverage probability
  - $c_1 = 0, c_2 = 0.3, c_7 = 0.3 + 0.7 = 1, \ldots$

Adversary

Defender

		60%	40%
		Cell1&2	**Cell 2&3**
30%	Purple Route	-2, 2	0, 0
70%	Orange Route		

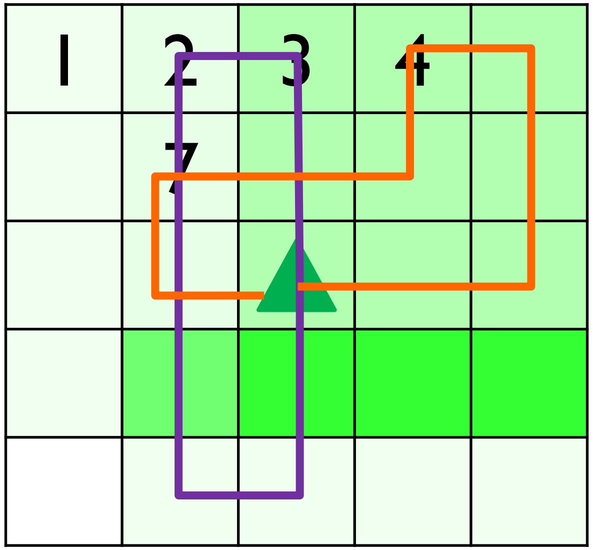

Cutting Plane

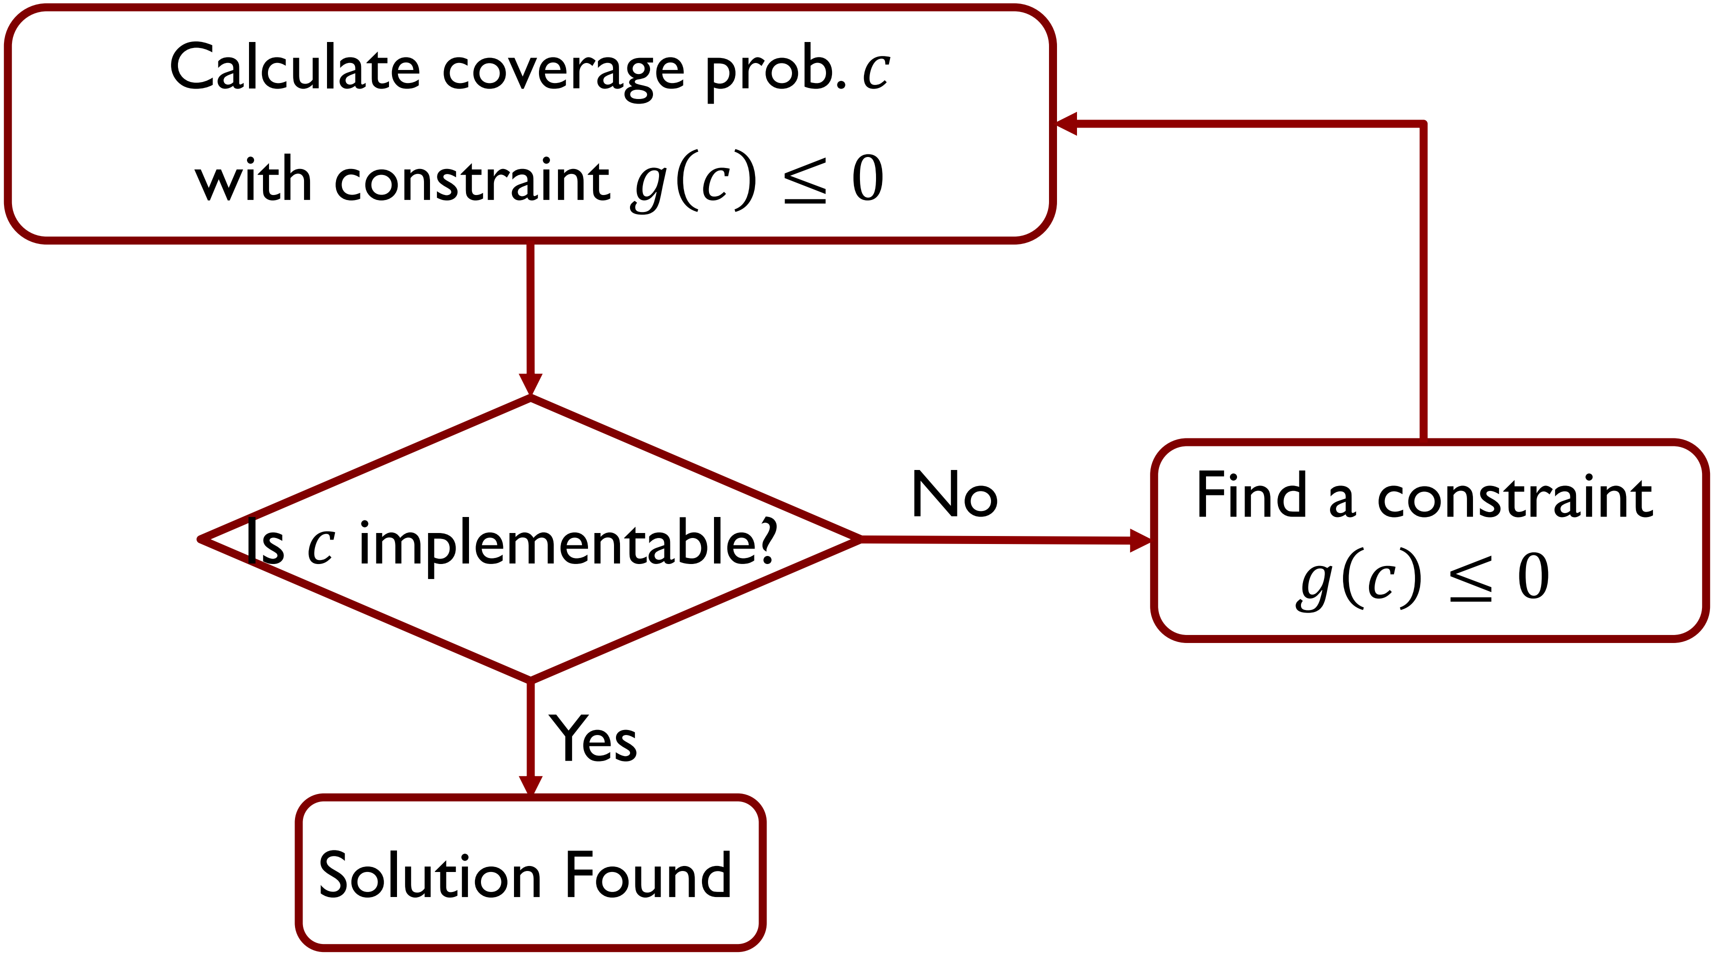

Calculate coverage prob. c with constraint $g(c) \leq 0$

Is c implementable?

No

Find a constraint $g(c) \leq 0$

Yes

Solution Found

Rong Yang, Albert Xin Jiang, Milind Tambe, Fernando Ordonez. *Scaling-up Security Games with Boundedly Rational Adversaries: A Cutting-plane Approach.* IJCAI'13

Cutting Plane

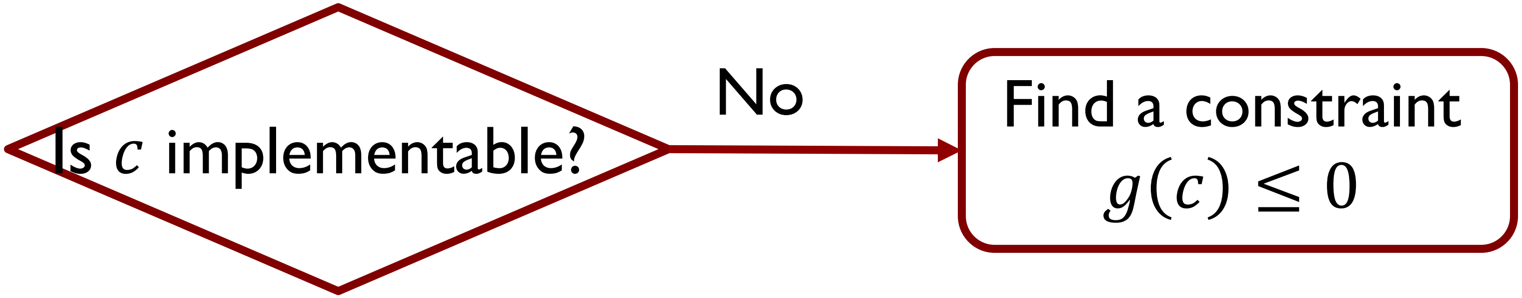

Is c implementable?

No

Find a constraint $g(c) \leq 0$

$\exists p$, such that $c_i = \sum_j p_j A_{ji}$ $\Leftrightarrow$ $z = \min_{\boxed{p}} \|c - A^T p\|_1$

Prob. of taking each route

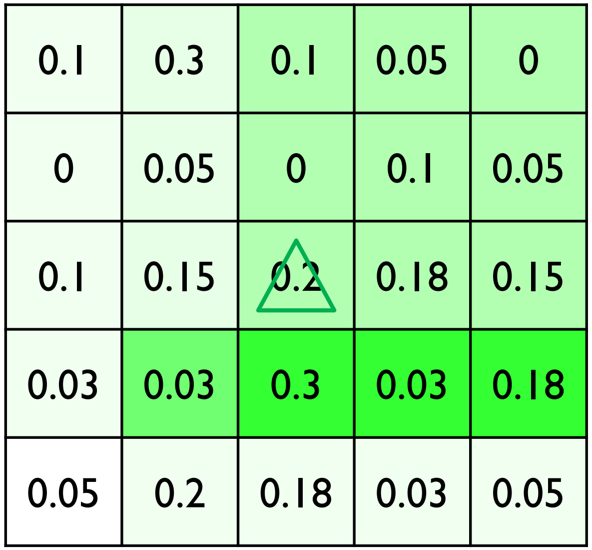

0.1	0.3	0.1	0.05	0
0	0.05	0	0.1	0.05
0.1	0.15	△ 0.2	0.18	0.15
0.03	0.03	0.3	0.03	0.18
0.05	0.2	0.18	0.03	0.05

if $z = 0$, implementable
if $z > 0$, found p^* and g

Cutting Plane

$$z = \min_{\boxed{p}} ||c - A^T p||_1$$

Prob. of taking each route

Not enumerate all routes?
Column generation!

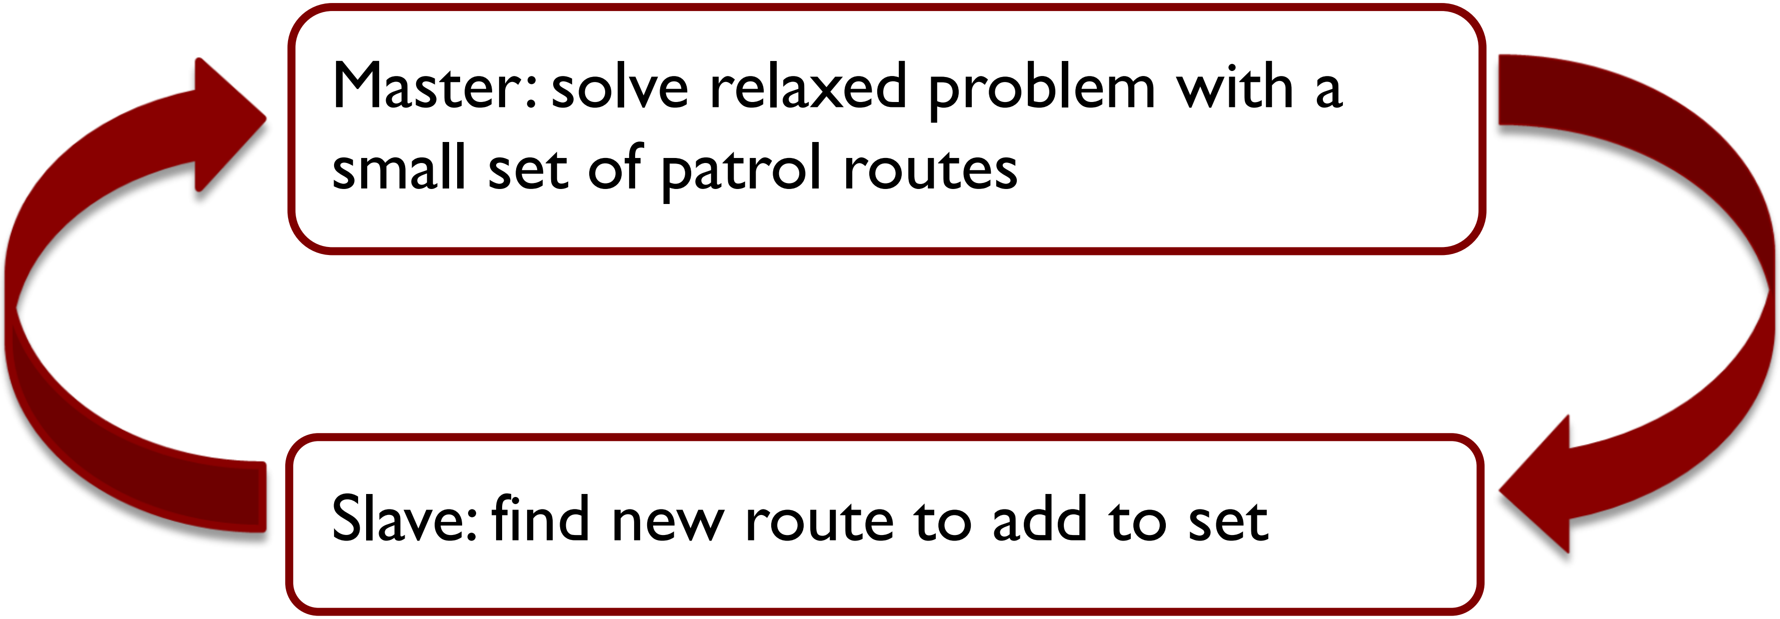

Master: solve relaxed problem with a small set of patrol routes

Slave: find new route to add to set

Cutting Plane

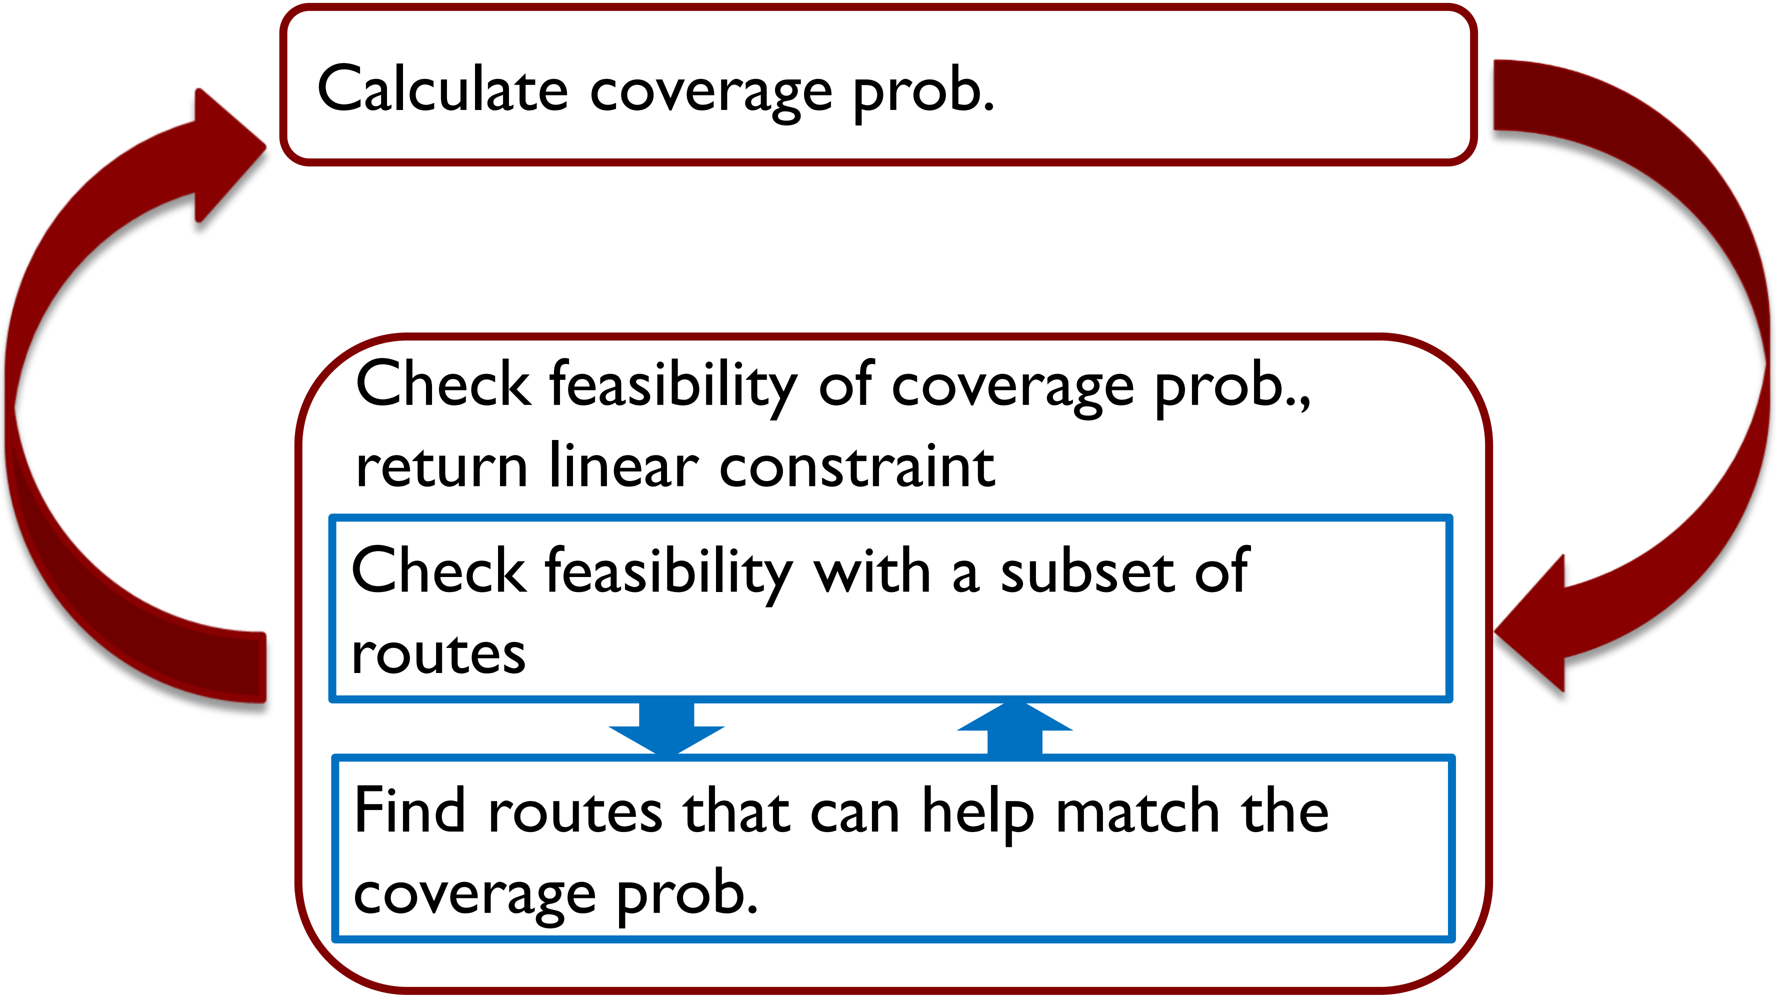

Calculate coverage prob.

Check feasibility of coverage prob., return linear constraint

Check feasibility with a subset of routes

Find routes that can help match the coverage prob.

Behind the Scene

- Hierarchical Modeling

- Find implementable game-theoretic solutions
  - Incremental strategy generation
  - Cutting plane

PAWS (Protection Assistant for Wildlife Security)

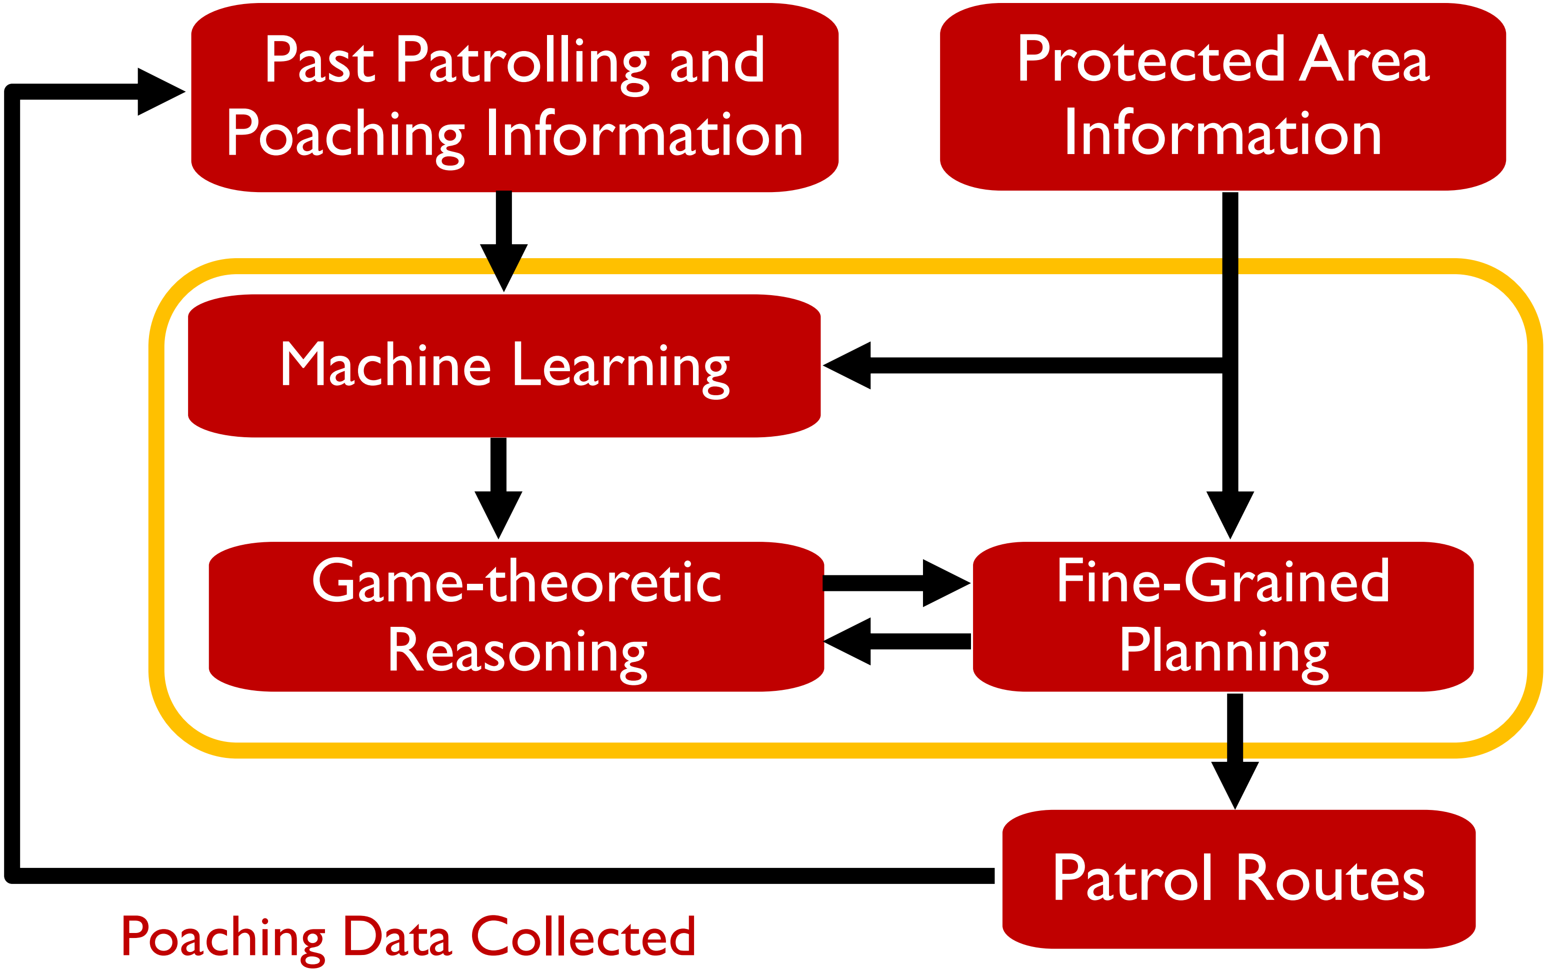

Past Patrolling and Poaching Information

Protected Area Information

Machine Learning

Game-theoretic Reasoning

Fine-Grained Planning

Patrol Routes

Poaching Data Collected

Real-World Deployment

- In collaboration with Panthera, Rimba
- Regular deployment since July 2015 (Malaysia)


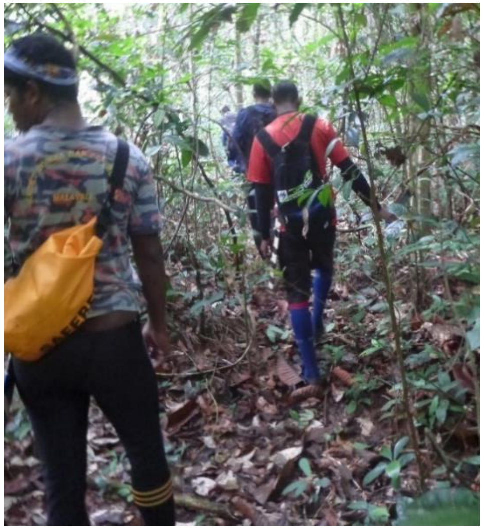


Real-World Deployment

Animal Footprint

Tiger Sign

Tree Mark

Camping Sign

Lighter

Summary

- Basic model
- Deal with continuous timeline
- Fine-grained planning with practical constraints

- Key take-aways
  - Game theory can be used to model security/sustainability challenges
  - Practical challenges void simple models
  - Evaluation through real-world deployment is challenging