Assuring Autonomy in Contested Environments

Attack-Resilient Design

Miroslav Pajic

Cyber-Physical Systems Lab (CPSL)

Pratt School of Engineering

Duke University

Attack Surfaces for Autonomous Systems

- Cyber attack surfaces
  - e.g., communication, networks, computers, databases, ...
- Physical attack surfaces
  - e.g., locks, casings, cables, ...
- Environmental attack surfaces
  - e.g., GPS signal, electro-magnetic interference, battery draining/cycling/heating, ...
- Human attack surfaces
  - e.g., phishing, bribing, blackmail

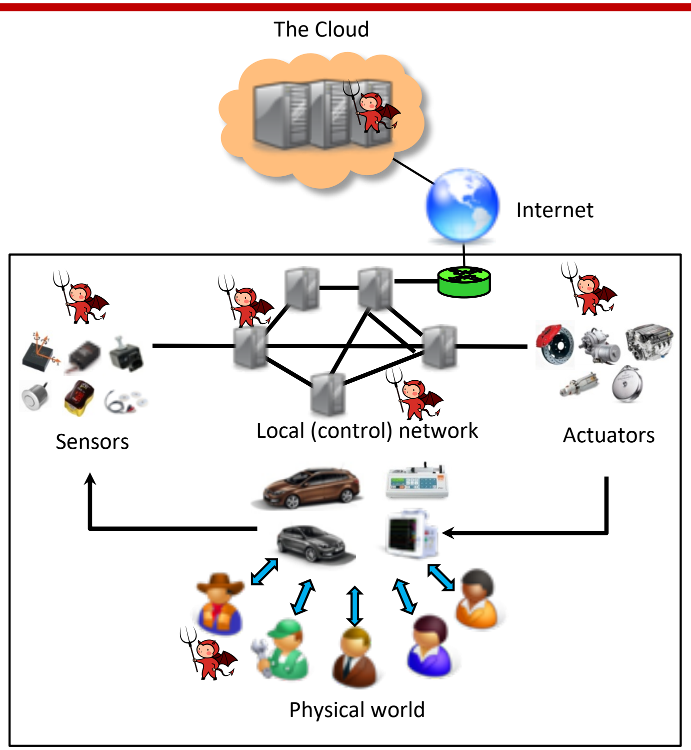

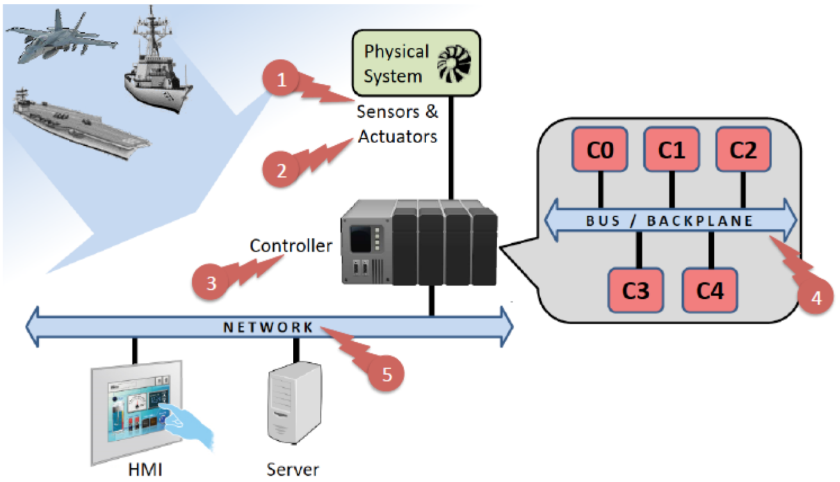

Attacks on Autonomy

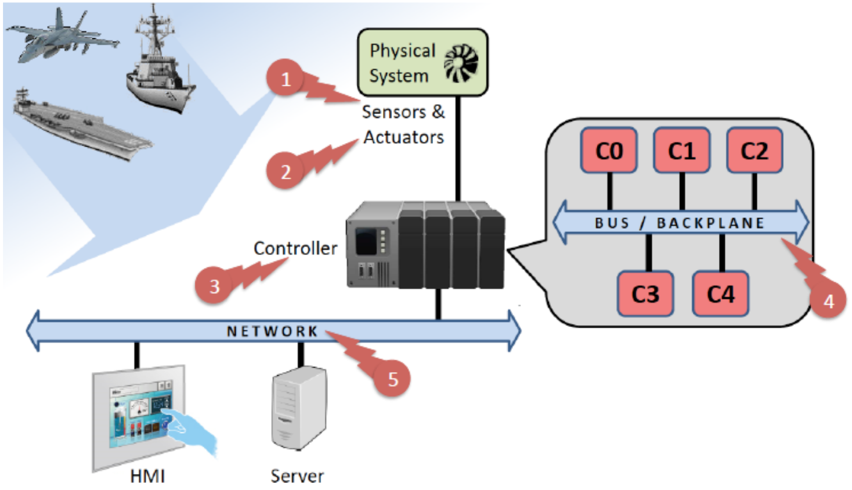

1. Sensor attacks
   - The attacker can arbitrarily change sensor measurements
2. Actuator attacks
   - The attacker can arbitrarily change actuator values
3. Controller attacks
   - The attacker can change the controllers' parameters, resources (e.g., execution model) or even the controllers' code

4-5. Communication attacks

- The attacker can change messages: sensors -> controllers, controllers -> actuators/controllers

Most of these attacks manifest themselves as malicious interference signals, and the defenses against them have to be introduced in the control/autonomy design.

Security-Aware Control for Autonomous Systems

Duke
PRATT SCHOOL of ENGINEERING

Control Stack	Control view	Modeling view	Adding Resiliency
Mission Planner	Long-horizon views	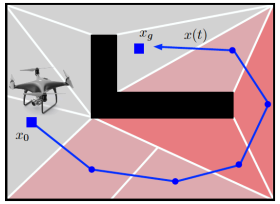	[ICRA19, ICRA20a, ICRA20b, CAV'19a, THMS19]
Tactical Planner	Short-horizon views	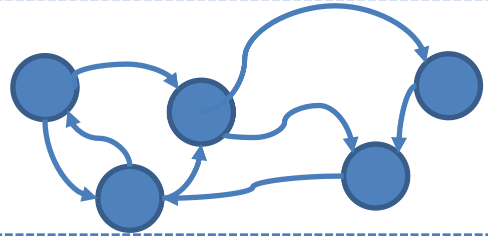	[CDC19a,CDC19b, TAC20*, TII19, TASE20*]
Low-level Control	Continuous/discrete control with constraints	$f_r(x(t)) = \int_0^T \varrho(x(t), h(t))dt + \int_0^T \lVert x(t)\rVert^2 dt,$ $\min \ f_r(x_r(t)) + f_h(x_h(t))$ $\text{s.t.} \ x_r(t) = x_h(t), \ u_r(t) = u_h(t),$	[TAC19a ,TAC19b, TCPS20, ACC20, AUT20a*, AUT20, AUT19, AUT18, TECS17, RTSS17, TCNS17, CSM17, CDC17, CDC18,...]
Vehicle	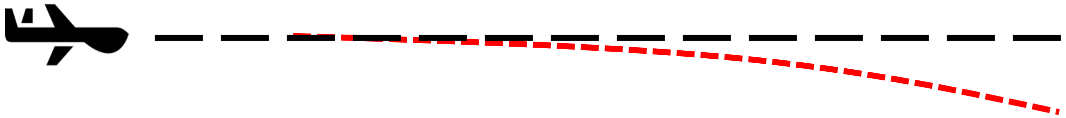		

Our Goal: Add resiliency to controls across different/all levels of control stack

Platform-aware Execution/Integration of Cyber-Physical Security Components

Control view

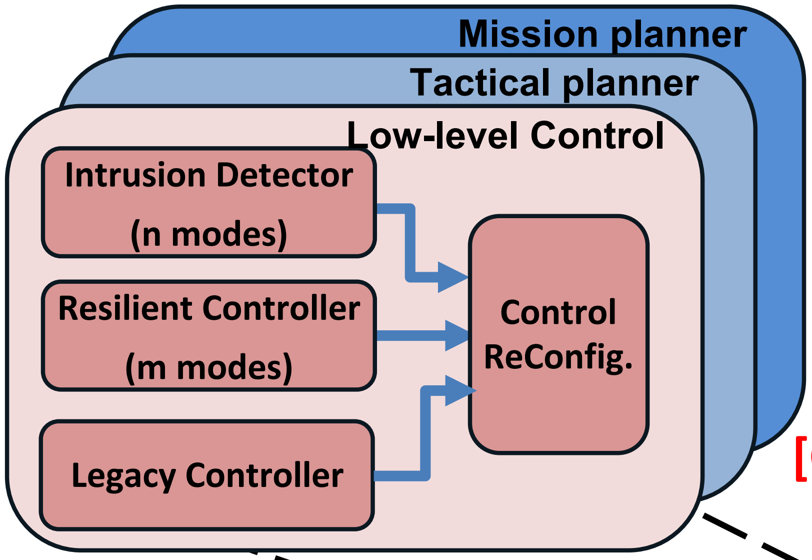

Constrained computation and communication resources limit the full use of developed cyber-physical techniques

Our Goal: Provide quantitative tradeoff procedure to map security-aware modules onto available architecture

[CMS17, TCNS17, TECS/EMSOFT17, RTSS17, TCPS20, TECS/EMSOFT'19, TAC'19]

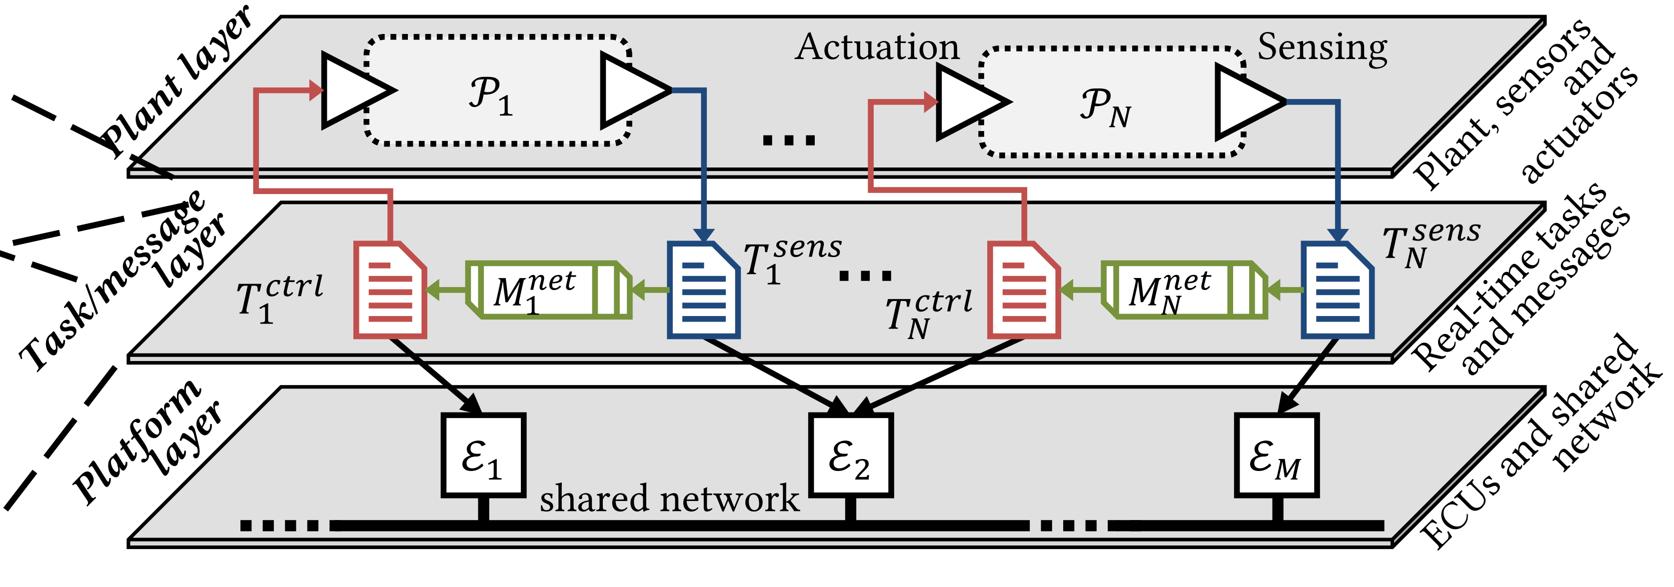

Runtime/platform support

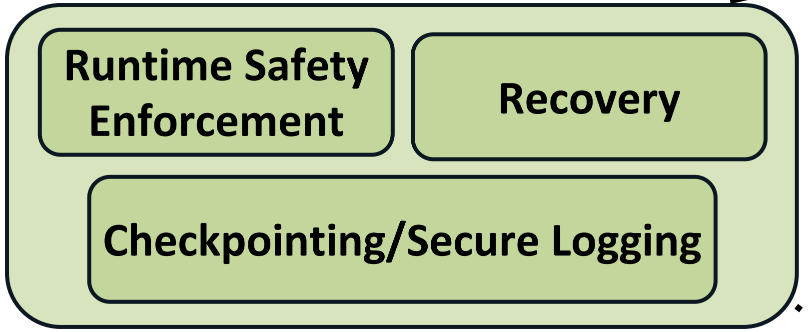

Attack Resilient Design – Some Major Contributions

Duke PRATT SCHOOL of ENGINEERING

- Security for network systems (strong connection with RT 3&4) via a novel moving target defense strategy that randomly changes the availability of sensor data
- Integrating security on resource-constrained platforms/environments (strong connection to RT3)
- Attack resilience supervisory control of discrete event systems (strong connection to RT1)
- Security-aware planning via delay-actions games and reinf. learning (strong connection to RT2)
- Design of security-aware human-autonomy interaction
- Resilient distributed hypothesis testing
- Modeling, design and analysis for security- and privacy-aware systems using (probabilistic) hyperproperties (strong connection to RT6)
- Open-source tool/testbed development
- Working with NATO Science and Technology IST-164 RTG Securing Unmanned and Autonomous Vehicles For Mission Assurance

Low-Level Control in the Presence of Attacks

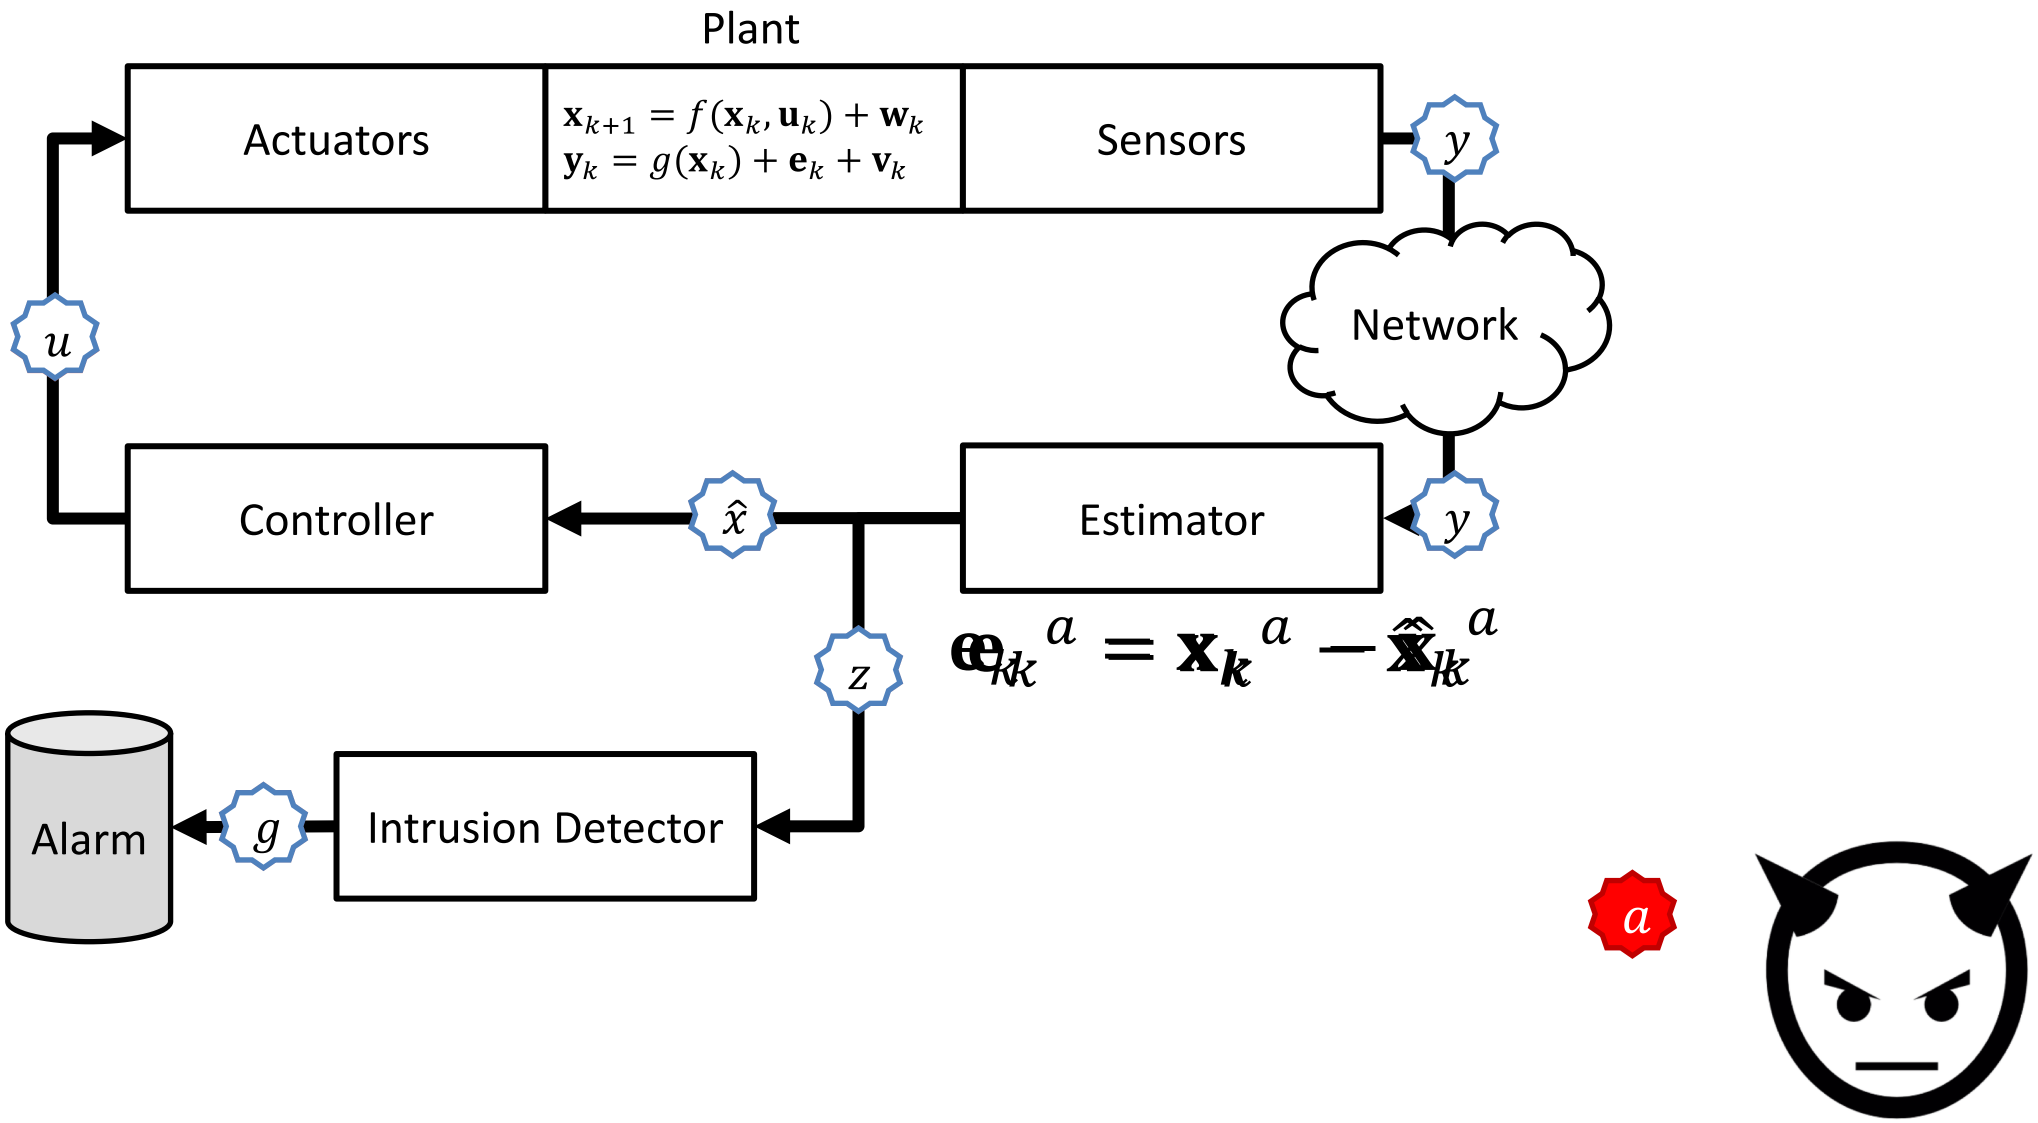

Can Attacker Reach Any State?

Duke PRATT SCHOOL of ENGINEERING

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k$$
$$\mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \mathbf{a}_k + \mathbf{v}_k$$

$$supp(\mathbf{a}_k) = \mathcal{K}$$
$$\mathbf{a}_{k,i} = 0, \forall i \in \mathcal{K}^C$$

Theorem 1 [1,2,3,**4,5***]:
A system presented above is perfectly attackable if and only if it is unstable, and at least one eigenvector **v** corresponding to an unstable mode satisfies $supp(\mathbf{Cv}) \subseteq \mathcal{K}$ and **v** is a reachable state of the dynamic system.

Physical detectors cannot always protect us from an intelligent attacker..

[1] Y. Mo and B. Sinopoli, "*False data injection attacks in control systems*," in First Workshop on Secure Control Systems, 2010
[2] C. Kwon, W. Liu, and I. Hwang, "*Analysis and design of stealthy cyber attacks on unmanned aerial systems*", Journal of Aerospace Information Systems, 1(8), 2014
[3] I. Jovanov and M. Pajic, "*Relaxing Integrity Requirements for Attack-Resilient Cyber-Physical Systems*", IEEE Trans. on Automatic Control, 2019
[4] A. Khazraei and M. Pajic, "*Perfect Attackability of Linear Dynamical Systems with Bounded Noise*," ACC 2020.
[5] A. Khazraei and M. Pajic, "*Attack-Resilient State Estimation with Intermittent Data Authentication*," Automatica, submitted

Can Attacker Reach Any State?

Duke PRATT SCHOOL of ENGINEERING

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k$$
$$\mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \mathbf{a}_k + \mathbf{v}_k$$

$$supp(\mathbf{a}_k) = \mathcal{K}$$
$$\mathbf{a}_{k,i} = 0, \forall i \in \mathcal{K}^C$$

Theorem 1 [1,2,3,**4,5**]:
A system presented above is perfectly attackable if and only if it is unstable, and at least one eigenvector **v** corresponding to an unstable mode satisfies $supp(\mathbf{Cv}) \subseteq \mathcal{K}$ and **v** is a reachable state of the dynamic system.

Theorem [3,4,5]: A system Σ with a global data integrity police $\mu(L)$ is not ***perfectly attackable***.

[1] Y. Mo and B. Sinopoli, "*False data injection attacks in control systems*," in First Workshop on Secure Control Systems, 2010
[2] C. Kwon, W. Liu, and I. Hwang, "*Analysis and design of stealthy cyber attacks on unmanned aerial systems*", Journal of Aerospace Information Systems, 1(8), 2014
[3] I. Jovanov and M. Pajic, "*Relaxing Integrity Requirements for Attack-Resilient Cyber-Physical Systems*", IEEE Trans. on Automatic Control, 2019
[4] A. Khazraei and M. Pajic, "*Perfect Attackability of Linear Dynamical Systems with Bounded Noise,*" ACC 2020.
[5] A. Khazraei and M. Pajic, "*Attack-Resilient State Estimation with Intermittent Data Authentication,*" Automatica, submitted

Why Resources might be a problem
Data Authentication Example

Two transmitters sharing a network:

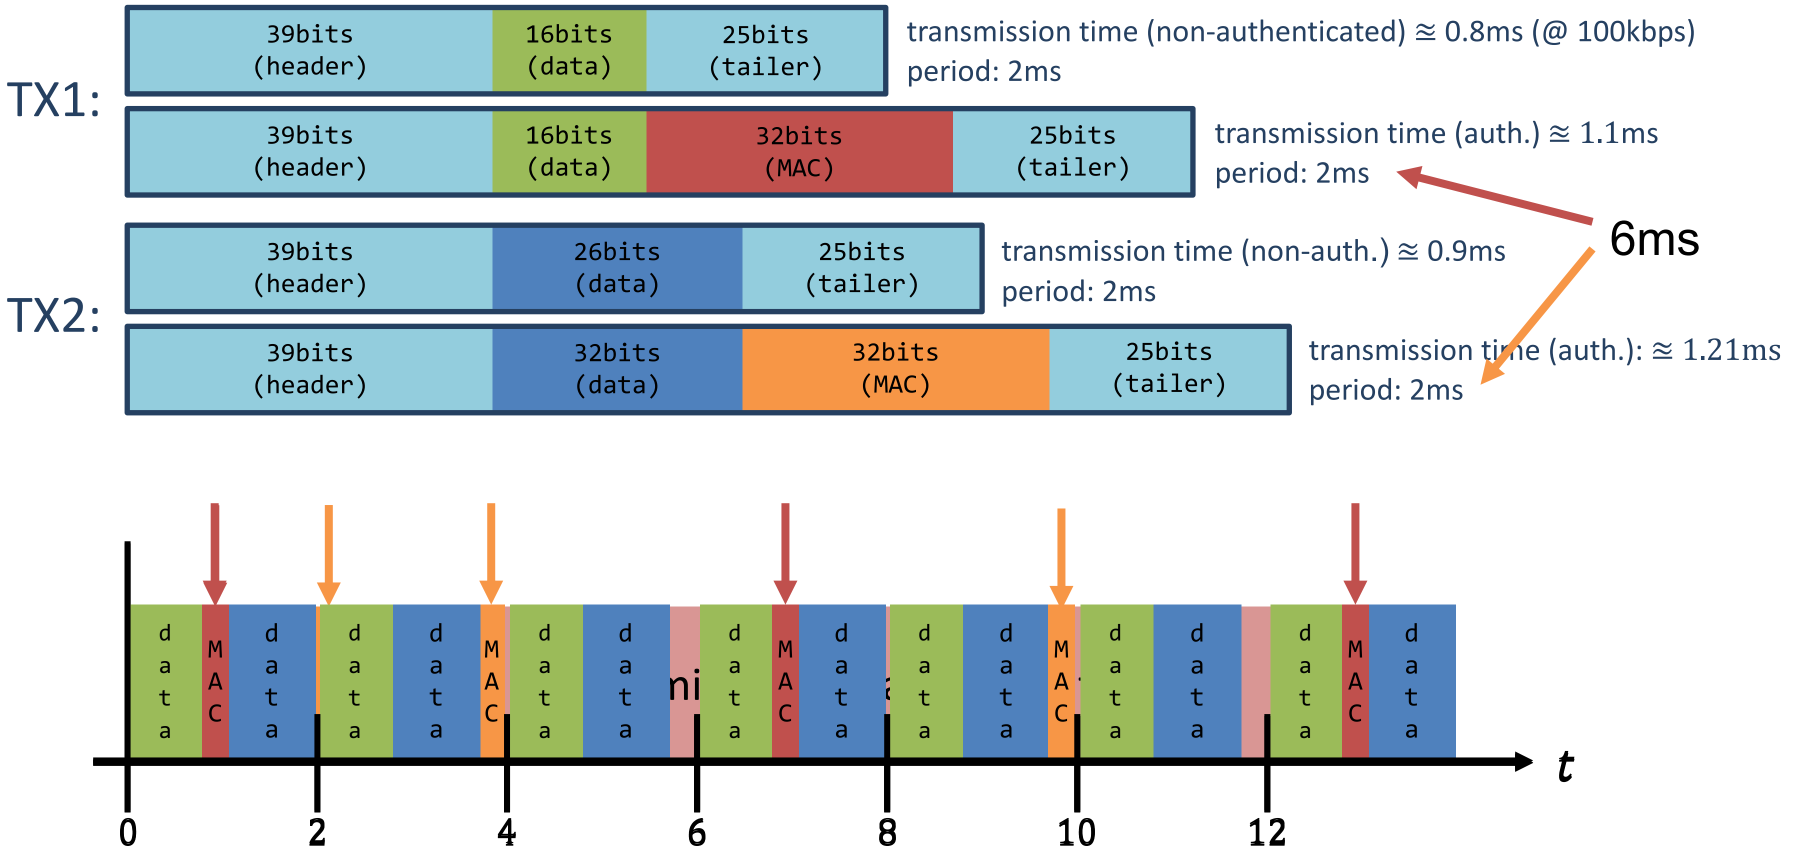

Security-per-$: Communication and computation resources are shared.
So how to add security mechanisms without affecting `normal' operation?

State Estimation Error
In the Presence of Stealthy Attacks

Reachable region of the state estimation error under attack [3]

$$\mathcal{R}[k] = \left\{ \boldsymbol{e} \in \mathbb{R}^n \middle| \begin{array}{c} \boldsymbol{e}\boldsymbol{e}^{\mathrm{T}} \preccurlyeq E[\boldsymbol{e}^a[k]]E[\boldsymbol{e}^a[k]]^{\mathrm{T}} + \gamma Cov(\boldsymbol{e}_k^a) \\ \boldsymbol{e}^a[k] = \boldsymbol{e}_k^a(\mathbf{a}_{1\ldots k}), \mathbf{a}_{1\ldots k} \in \mathcal{A}_k \end{array} \right\}$$

$\mathbf{a}_{1\ldots k} = [\mathbf{a}[1]^{\mathrm{T}} \ldots \mathbf{a}[k]^{\mathrm{T}}]^{\mathrm{T}}$

$\mathcal{A}_k$ is the set of all stealthy attacks

$\boldsymbol{e}_k^a(\mathbf{a}_{1\ldots k})$ is the estimation error evolution due to attack $\mathbf{a}_{1\ldots k}$

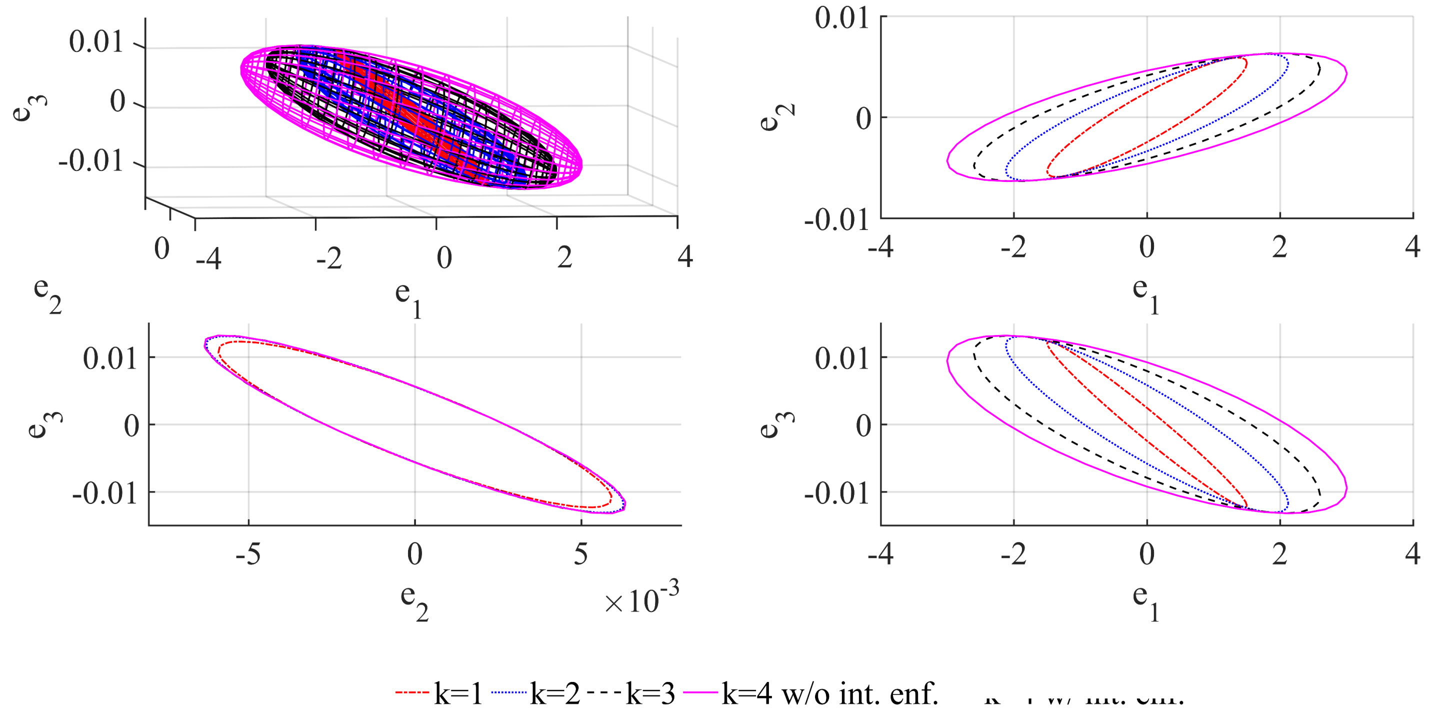

[3] I. Jovanov and M. Pajic, "*Relaxing Integrity Requirements for Attack-Resilient Cyber-Physical Systems*", IEEE Trans. on Automatic Control, 2019

Integrity Enforcement Policy

Integrity enforcement policy ensures attacker's influence is zeroed at enforcement points

Data integrity enforcement policy (μ, l) where $\mu = \{t_k\}_{k=0}^{\infty}$, with $t_{k-1} < t_k, \forall k > 0$
and $l = \sup_{k>0} t_k - t_{k-1}$ ensures that $\mathbf{a}_{1\ldots k} = 0, \forall k \geq 0$

This means that at points of authentication $\boldsymbol{y}_i^{net,a}[k] = \boldsymbol{y}_i^a[k]$

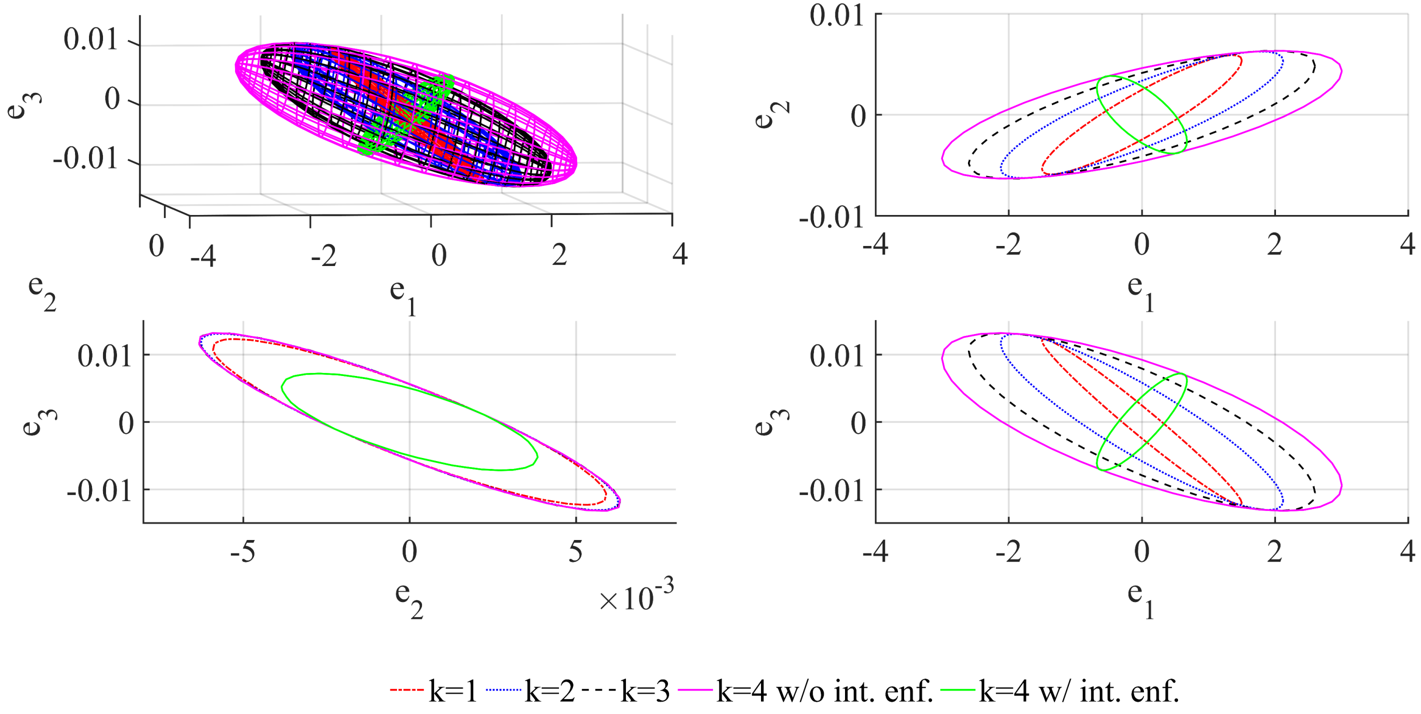

QoC Under Attack as Function of Resources?

Evolution of the state-estimation error due to attack is a sound performance metric

$$\mathcal{J}(l) = sup\{\|\boldsymbol{e}^a\|_2 \big| \boldsymbol{e}^a \in \mathcal{R}^l\} \qquad \mathcal{R}^l = \bigcup_{k=0}^{\infty} \mathcal{R}^l[k]$$

where $\mathcal{R}^l[k]$ denotes $\mathcal{R}[k]$ computed over all integrity enforcement policies with parameter l

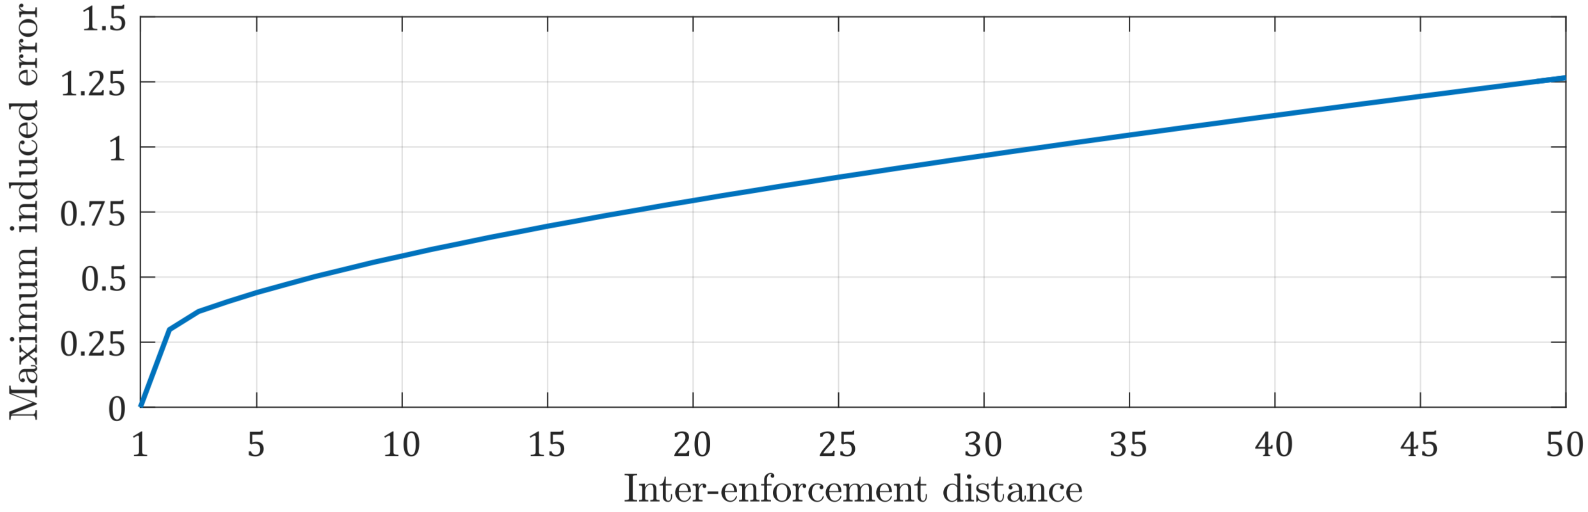

V. Lesi, I. Jovanov, and M. Pajic, "Integrating Security in Resource-Constrained Cyber-Physical Systems", *ACM Transactions on Cyber-Physical Systems*, 2020, **accepted**

QoC Under Attack as Function of Resources?

Evolution of the state-estimation error due to attack is a sound performance metric

$$\mathcal{J}(l) = sup\{\|\boldsymbol{e}^a\|_2 \big| \boldsymbol{e}^a \in \mathcal{R}^l\} \qquad \mathcal{R}^l = \bigcup_{k=0}^{\infty} \mathcal{R}^l[k]$$

where $\mathcal{R}^l[k]$ denotes $\mathcal{R}[k]$ computed over all integrity enforcement policies with parameter l

Piecewise-linear approximation of the QoC-degradation curves

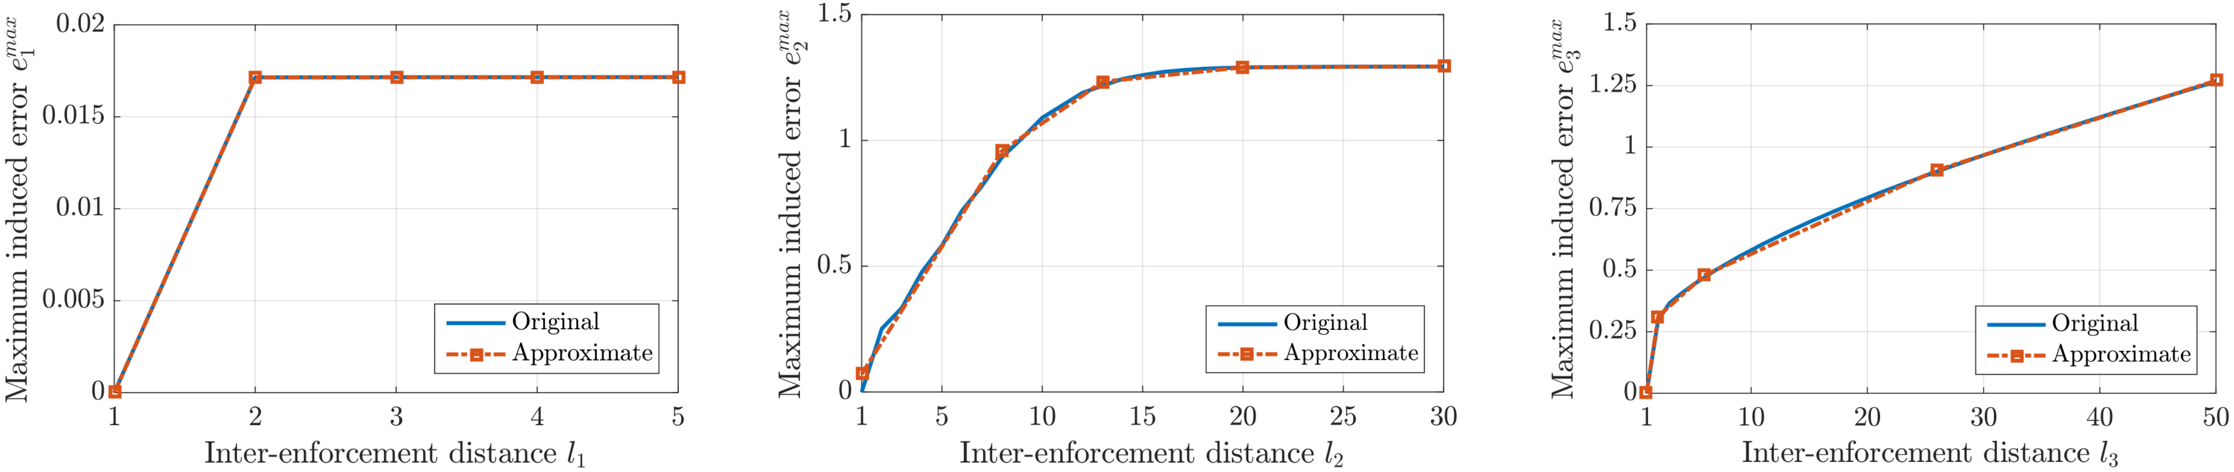

V. Lesi, I. Jovanov, and M. Pajic, "Integrating Security in Resource-Constrained Cyber-Physical Systems", *ACM Transactions on Cyber-Physical Systems*, 2020, **accepted**

Security-Aware Design Framework

Security-aware task modeling: Two-frame, implicit deadline tasks with peak frame offsets

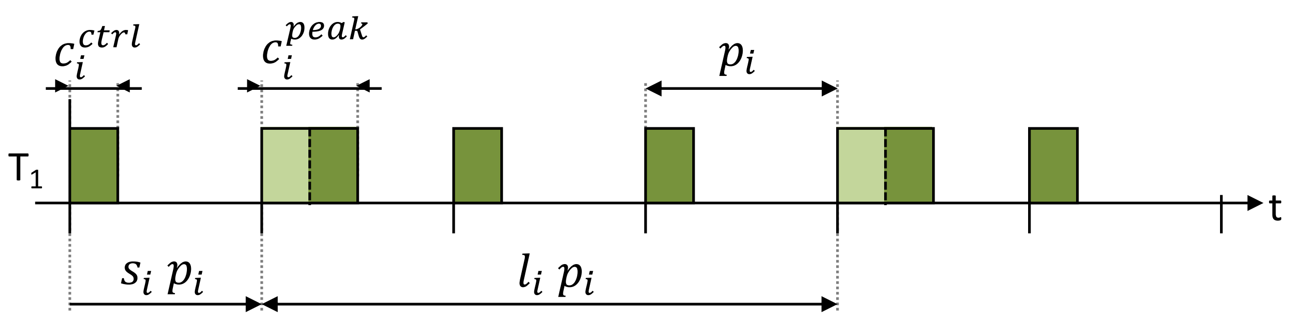

$T_i(C_i, p_i, D_i, l_i, s_i)$ where $C_i=[c_i^{ctrl}, c_i^{peak}]$, $D_i=[d_i^{ctrl}, d_i^{peak}]$, $1 \le i \le N$, and:

c_i^{ctrl}	WCET time of a control frame
c_i^{peak}	WCET time of both control − and security − related parts
p_i	task period
$d_i^{ctrl} = p_i$	relative deadline of a normal frame
$d_i^{peak} = p_i$	relative deadline of a peak frame
l_i	inter − peak frame distance
s_i	offset of the initial peak frame in the range $[0, l_i - 1]$

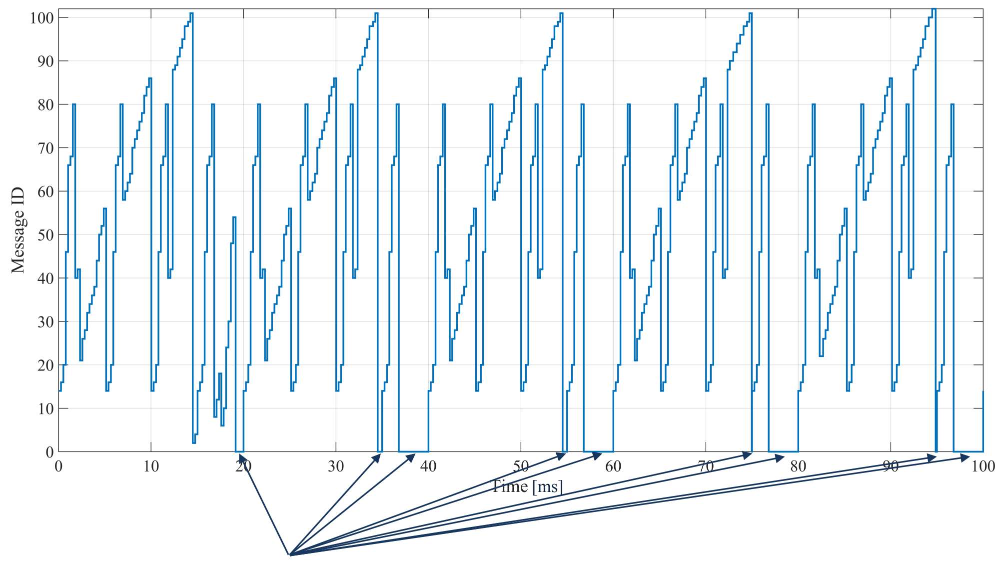

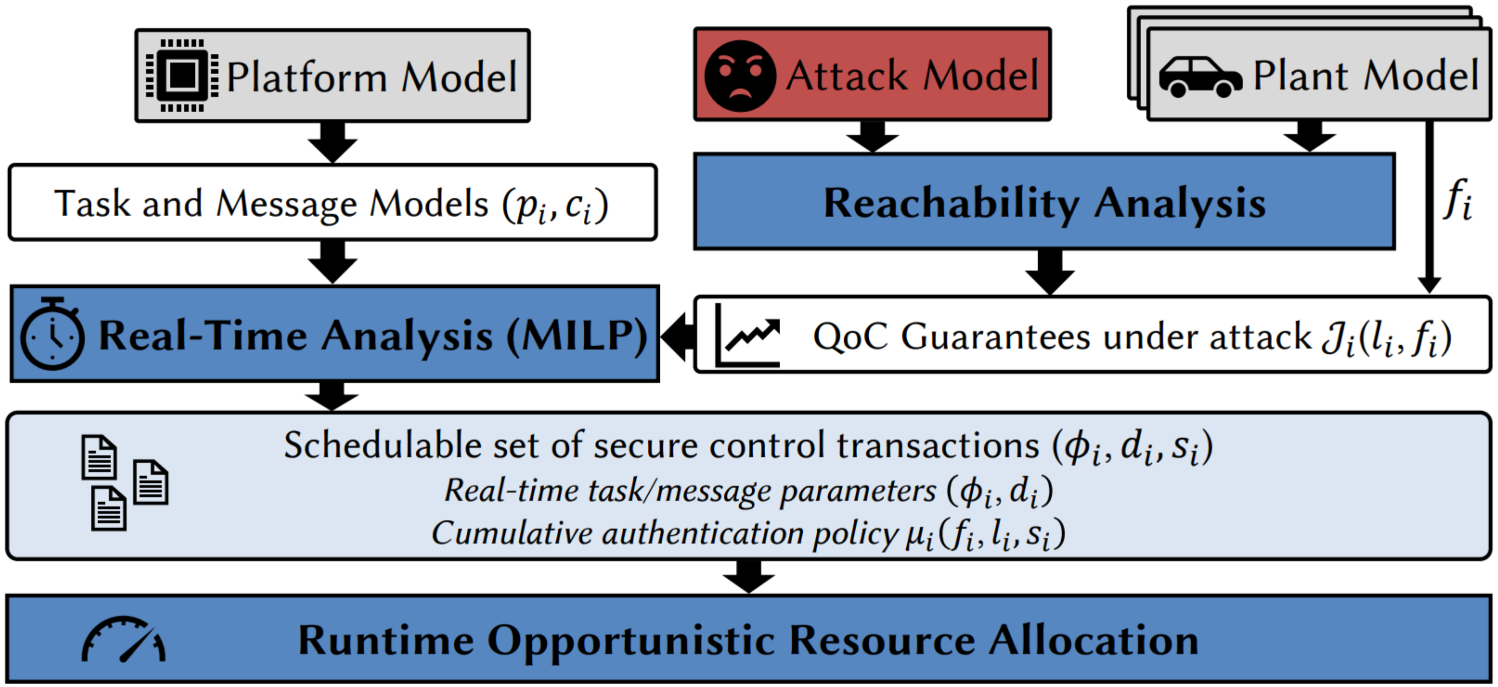

Security-Aware Control for Autonomous Systems

Duke PRATT SCHOOL of ENGINEERING

Control Stack	Control view	Modeling view
Mission Planner	Long-horizon views	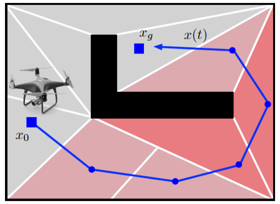
Tactical Planner	Short-horizon views	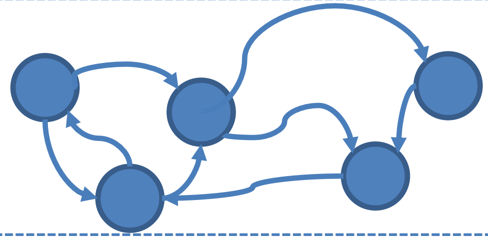
Low-level Control	Continuous/discrete control with constraints	$f_r(x(t)) = \int_0^T \varrho(x(t), h(t))dt + \int_0^T \|x(t)\|^2 dt,$ $\min \; f_r(x_r(t)) + f_h(x_h(t))$ $\text{s.t.} \; x_r(t) = x_h(t), \; u_r(t) = u_h(t),$
Vehicle	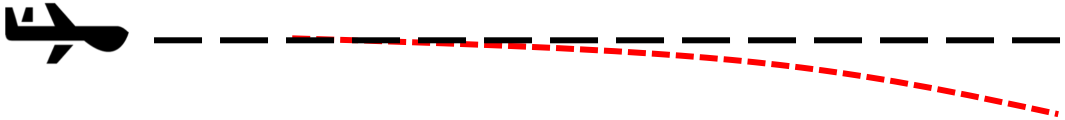	

ARSC: A tool for design of *attack-resilient supervisory controllers*, 2020.
https://gitlab.oit.duke.edu/cpsl/arsc

Y. Wang, A. Bozkurt, and M. Pajic, "Attack-Resilient Supervisory Control of Discrete Event Systems", *IEEE Transactions on Automatic Control*, submitted.
Z. Jakovljevic, V. Lesi, and M. Pajic, "Attacks on Distributed Sequential Control in Manufacturing Automation", *IEEE Transactions on Industrial Informatics*, accepted.
V. Lesi, Z. Jakovljevic and M. Pajic, "Security-Analysis for Distributed IoT-Based Industrial Automation", IEEE Trans. on Automation Science and Engineering, submitted.
Y. Wang and M. Pajic, "Supervisory Control of Discrete Event Systems in the Presence of Sensor and Actuator Attacks", IEEE CDC, 2019.
Y. Wang and M. Pajic, " Attack-Resilient Supervisory Control with Intermittent Authentication", IEEE CDC, 2019.
V. Lesi, Z. Jakovljevic and M. Pajic, "Reliable Industrial IoT-Based Distributed Automation", 4th ACM/IEEE IoTDI, 2019.

Security-Aware Planning for Autonomous Systems

Duke
PRATT SCHOOL *of* ENGINEERING

Control Stack	Control view	Modeling view
Mission Planner	Long-horizon views	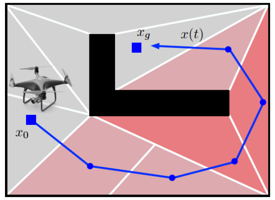
Tactical Planner	Short-horizon views	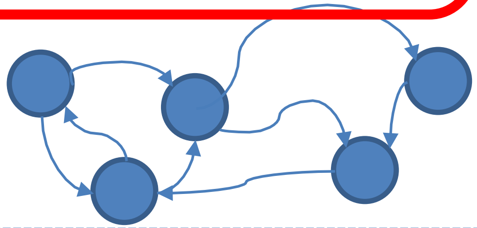
Low-level Control	Continuous/discrete control with constraints	$f_r(x(t)) = \int_0^T \varrho(x(t), h(t))dt + \int_0^T \|x(t)\|^2 dt,$ $\min \; f_r(x_r(t)) + f_h(x_h(t))$ $\text{s.t.} \; x_r(t) = x_h(t), \; u_r(t) = u_h(t),$
Vehicle		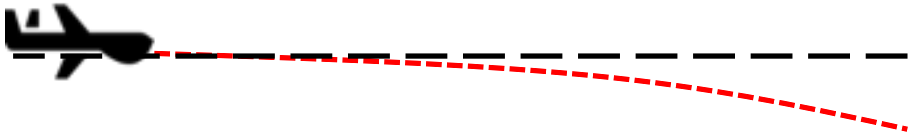

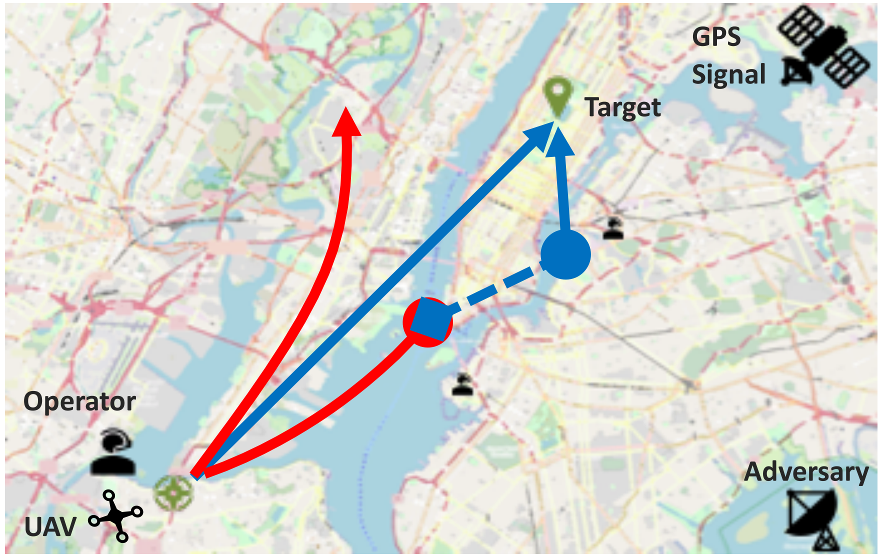

DAG | Hidden-Information Semantics

Duke PRATT SCHOOL of ENGINEERING

UAV Model

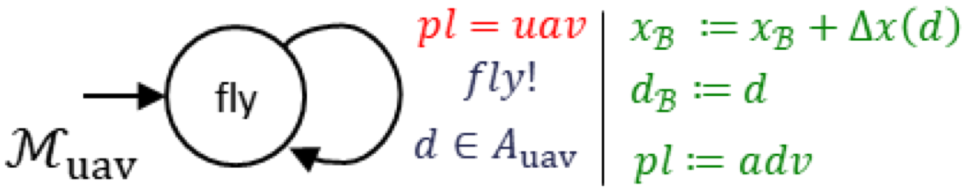

Adversary Model

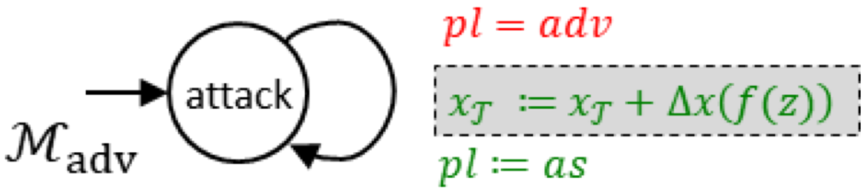

Advisory System Model

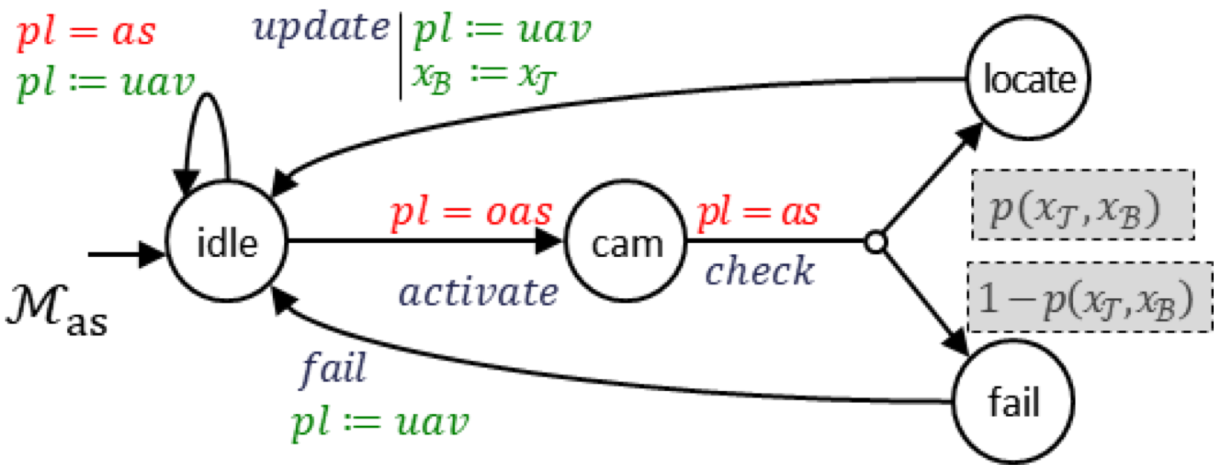

Information inside this box is oftentimes unknown, i.e., **hidden**

Off-the-shelf model checkers do NOT support hidden variables

Strategies CANNOT be synthesized based on hidden information

Approach: Delaying Actions

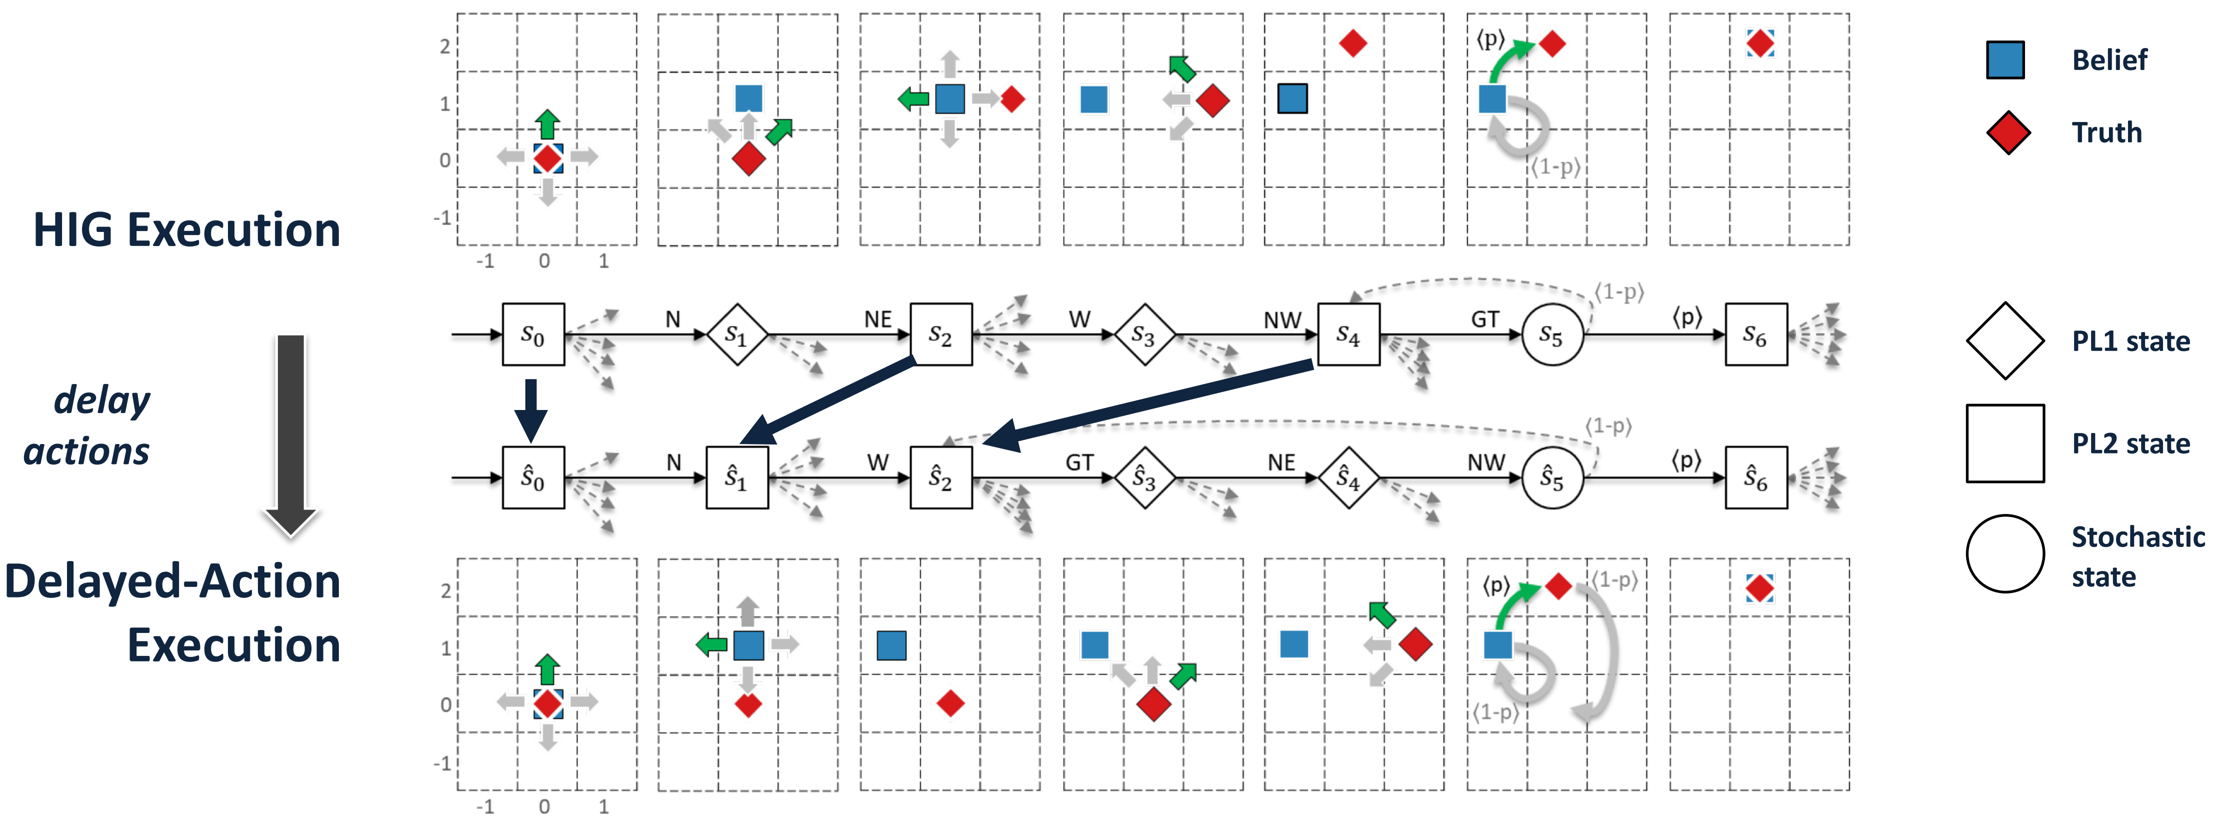

Information is hidden from one player (H-UAV) by delaying the actions of the other player (ADV)

Approach: Delaying Actions

Hidden Information Game

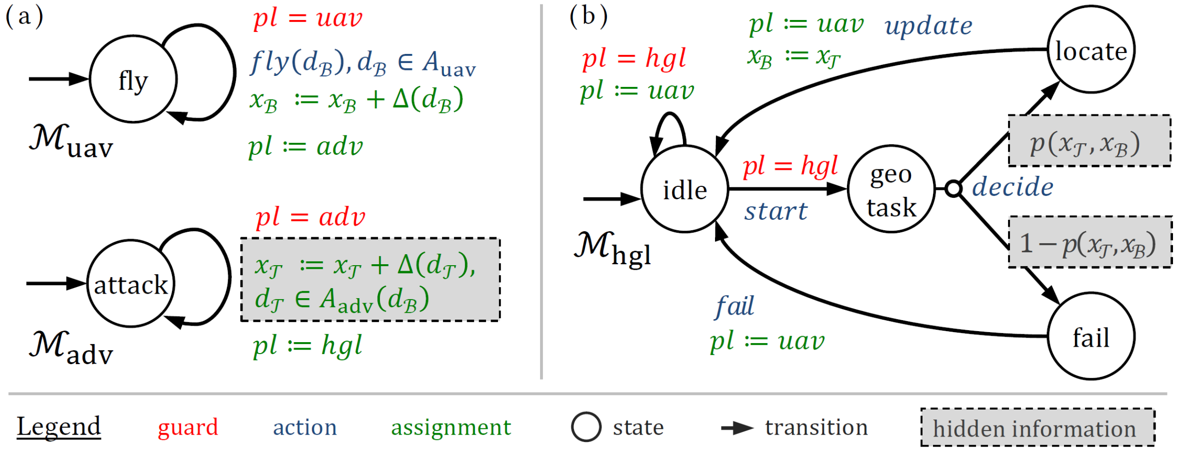

Delayed Action Game

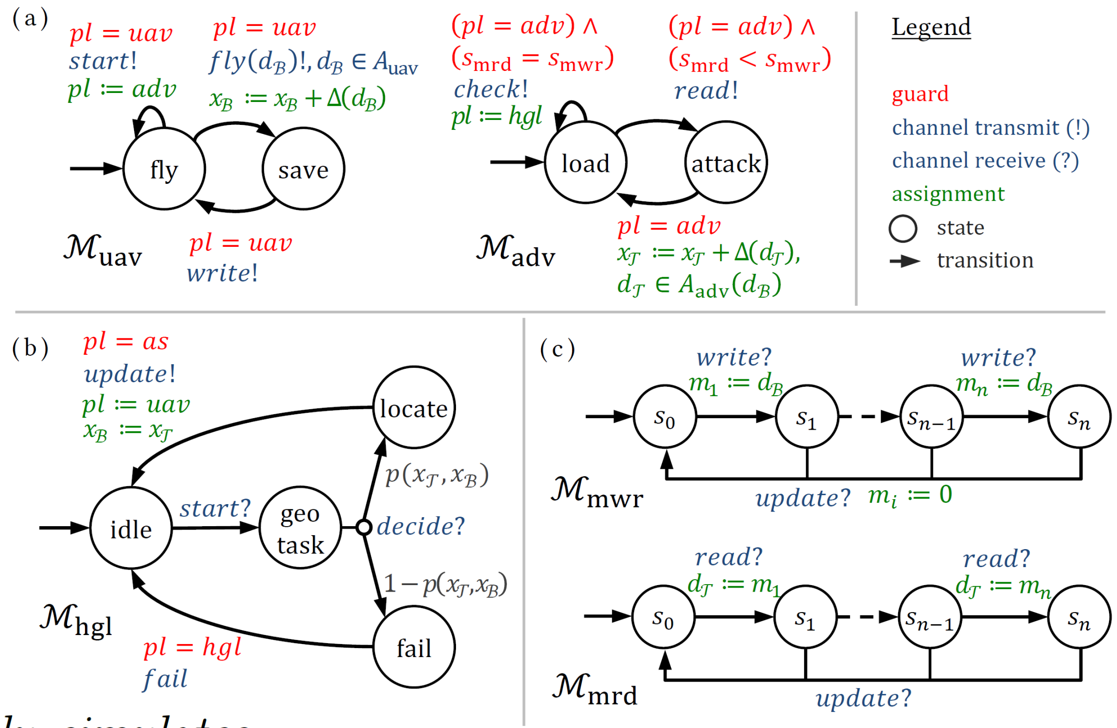

Definition 9 (Game Proper Simulation). *A game $\mathcal{G}_{\mathrm{D}}$ properly simulates $\mathcal{G}_{\mathrm{H}}$, denoted by $\mathcal{G}_{\mathrm{D}} \rightsquigarrow \mathcal{G}_{\mathrm{H}}$, iff* $\forall \varrho \in \mathrm{Prop}(\mathcal{G}_{\mathrm{H}})$, $\exists \hat{\varrho} \in \mathrm{Prop}(\mathcal{G}_{\mathrm{D}})$ *such that* $\varrho \sim \hat{\varrho}$.

Theorem 1 (Probabilistic Simulation). *For any* $s_0 \simeq \hat{s}_0$ *and* $\varrho \in \mathrm{Prop}(\mathcal{G}_{\mathrm{H}})$ *where* $first(\varrho) = s_0$, *it holds that*

$$\Pr\left[last(\varrho) = s'\right] = \Pr\left[\left(\overline{move(\varrho)}\right)(\hat{s}_0) = \hat{s}'\right] \quad \forall s', \hat{s}' \;\; s.t. \;\; s' \simeq \hat{s}'.$$

Theorem 2 (DAG-HIG Simulation). *For any HIG $\mathcal{G}_{\mathrm{H}}$ there exists a DAG* $\mathcal{G}_{\mathrm{D}} = \mathfrak{D}[\mathcal{G}_{\mathrm{H}}]$ *such that* $\mathcal{G}_{\mathrm{D}} \rightsquigarrow \mathcal{G}_{\mathrm{H}}$ *(as defined in Def. 9).*

M. Elfar, Y. Wang, and M. Pajic, "Security-Aware Synthesis using Delayed Action Games", *31st Int. Conference on Computer-Aided Verification* (**CAV**), 2019.

DAG-Based Synthesis

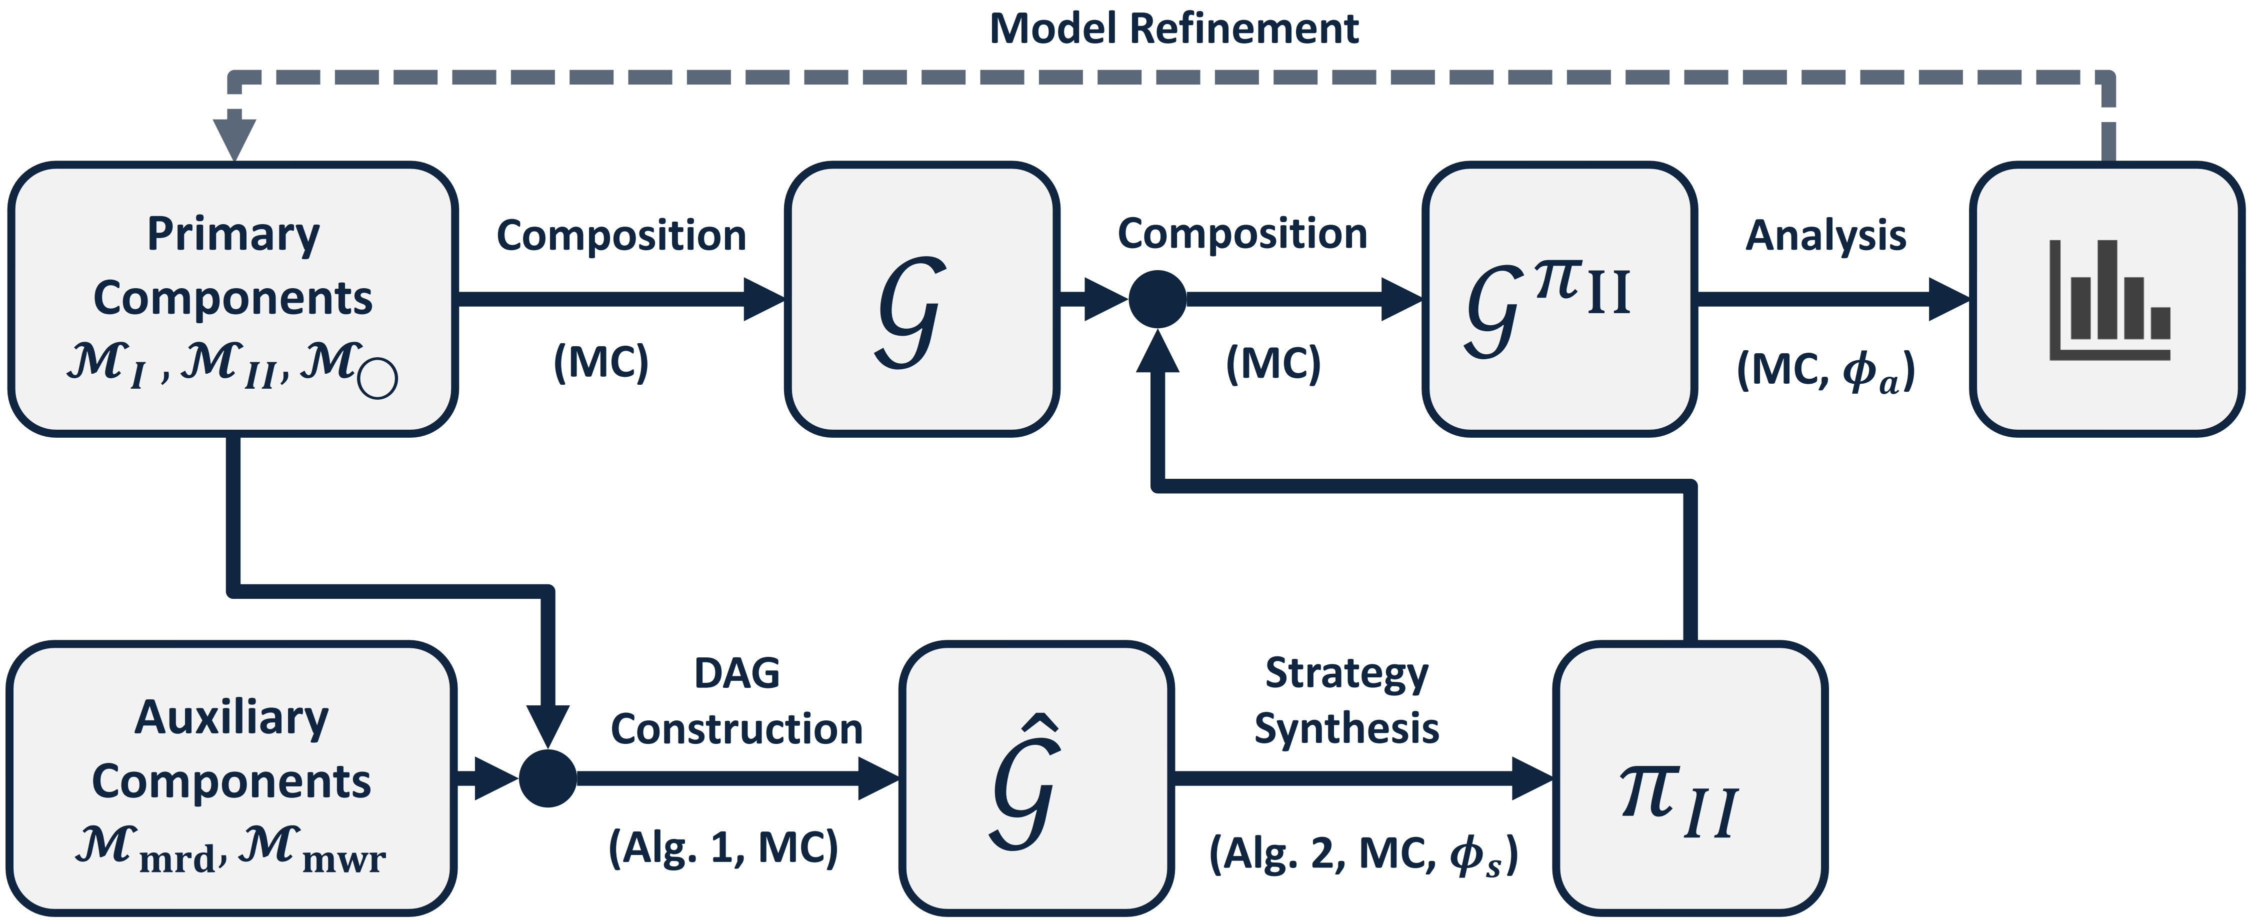

MC: Model Checker
ϕ_s : Synthesis query
ϕ_a : Analysis query

Model-free Control Synthesis from LTLs

Problem Statement

Given an MDP $\mathrm{M} = (S, A, P, s_0, AP, L)$ where P is ***fully*** unknown and an LTL specification φ, design a model-free RL algorithm that finds a finite-memory objective policy π^{φ} that satisfies

$$Pr_{\pi^{\varphi}}(s \vDash \varphi) = Pr_{max}(s \vDash \varphi),$$

where $Pr_{max}(s \vDash \varphi) = max_{\pi} Pr_{\pi}(s \vDash \varphi)$ for all $s \in S$.

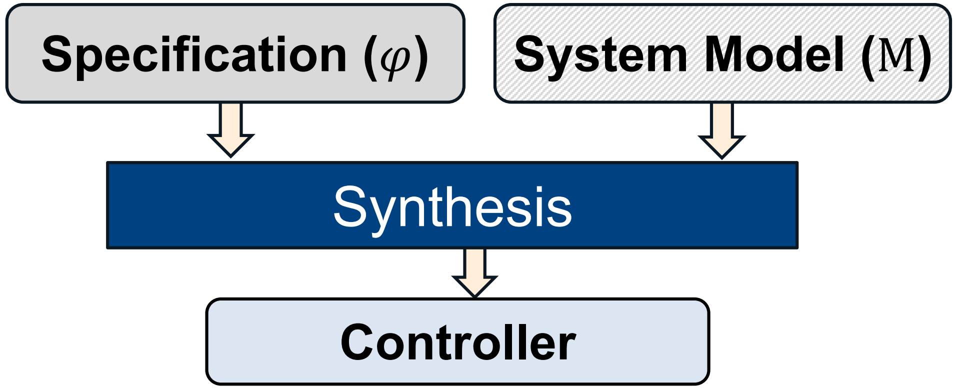

[1] A. Bozkurt, Y. Wang, M. Zavlanos, and M. Pajic, "Control Synthesis from Linear Temporal Logic Specifications using Model-Free Reinforcement Learning", *IEEE International Conference on Robotics and Automation* (**ICRA**), 2020, **accepted**.

[2] Q. Gao, M. Pajic, and M. Zavlanos, "Deep Imitative Reinforcement Learning for Temporal Logic Robot Motion Planning with Noisy Semantic Observations", *IEEE International Conference on Robotics and Automation* (**ICRA**), 2020, **accepted**.

Model-free Control Synthesis from LTLs

Duke
PRATT SCHOOL *of* ENGINEERING

Problem Statement

Given an MDP $\mathrm{M} = (S, A, P, s_0, AP, L)$ where P is ***fully*** unknown and an LTL specification φ, design a model-free RL algorithm that finds a finite-memory objective policy π^{φ} that satisfies

$$Pr_{\pi^{\varphi}}(s \vDash \varphi) = \ Pr_{max}(s \vDash \varphi),$$

where $Pr_{max}(s \vDash \varphi) = max_{\pi} Pr_{\pi}(s \vDash \varphi)$ for all $s \in S$.

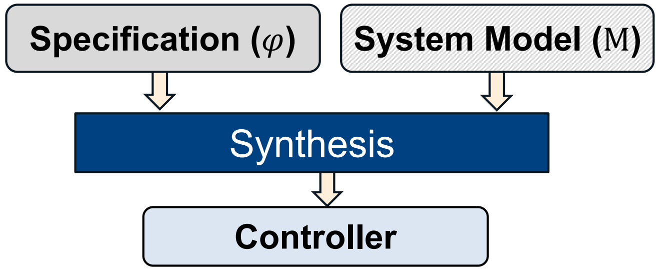

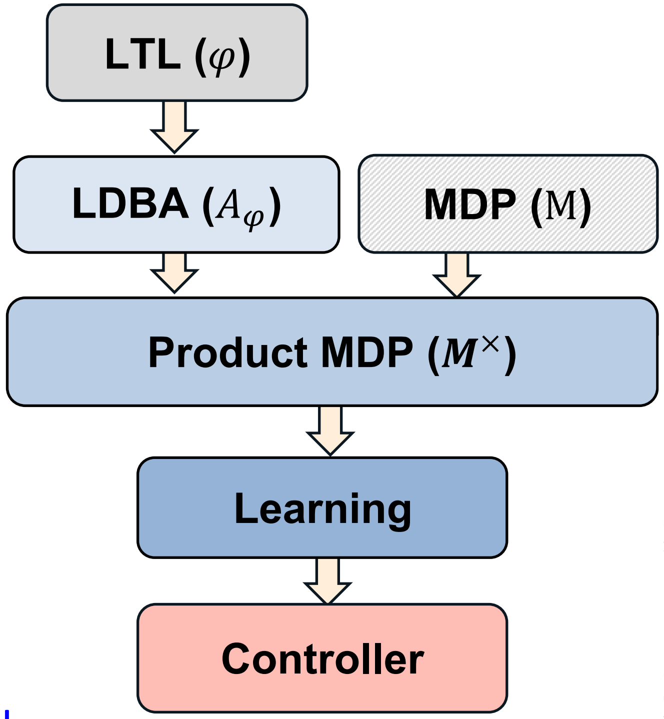

Limit-Deterministic Büchi Automata (LDBA) – consist of two deterministic components the ***initial*** and ***accepting***. The only nondeterministic transitions are the ϵ-moves from the initial component to the accepting components.

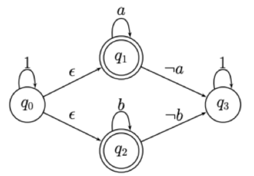

(a) A derived LDBA $\mathcal{A}$ for the LTL formula $\varphi = \Diamond\Box a \vee \Diamond\Box b$

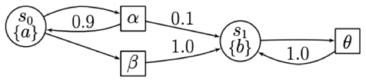

(b) An example MDP $\mathcal{M}$; the circles denote MDP states, rectangles denote actions, and numbers transition probabilities

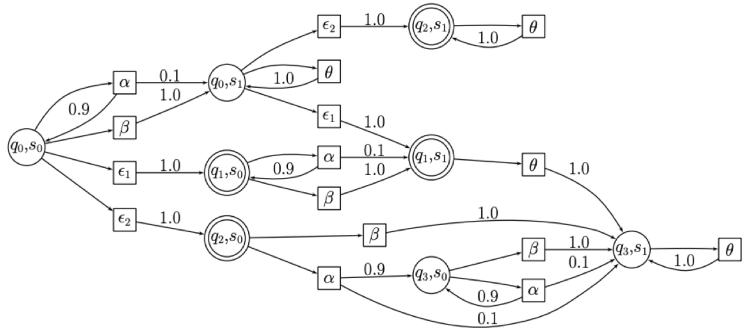

(c) The obtained product MDP

CSRL – Control synthesis for LTL objectives via model-free reinforcement learning,2020.
https://gitlab.oit.duke.edu/cpsl/csrl

Model-free Learning for Stochastic Buchi Games

Duke PRATT SCHOOL *of* ENGINEERING

Problem Statement

Given an MDP $M = (S, A, P, s_0, AP, L)$ where P is ***fully unknown*** and an LTL specification φ, design a model-free RL algorithm that finds a ***finite-memory*** objective policy π^{φ} that satisfies

$$Pr_{\pi^{\varphi}}(s \vDash \varphi) = Pr_{max}(s \vDash \varphi),$$

where

$$Pr_{max}(s \vDash \varphi) = max_{\pi} Pr_{\pi}(s \vDash \varphi)$$

for all $s \in S$.

Theorem 1: For a given two-player stochastic Buchi game $\mathcal{G}$ with $B \subseteq S$, the value of the game $v^{\gamma}_{c,e}$ for the strategy pair (c, e) and the discount factor γ satisfies

$$\lim_{\gamma \to 1^-} v^{\gamma}_{c,e}(s) = Pr_{c,e}(s \models \Box\Diamond B) \tag{8}$$

for all states $s \in S$, if the return of a path is defined as

$$G_t(\sigma) := \sum_{i=0}^{\infty} R_B(\sigma[t+i]) \cdot \prod_{j=0}^{i-1} \Gamma_B(\sigma[t+j]) \tag{9}$$

where $\prod_{j=0}^{-1} := 1$, $R_B : S \to [0,1)$ and $\Gamma_B : S \to (0,1)$ are the reward and the discount functions defined as:

$$R_B(s) := \begin{cases} 1-\gamma_B & s \in B \\ 0 & s \notin B \end{cases}, \quad \Gamma_B(s) := \begin{cases} \gamma_B & s \in B \\ \gamma & s \notin B \end{cases} \tag{10}$$

Here, we set $\gamma_B = \gamma_B(\gamma)$ as a function of γ such that

$$\lim_{\gamma \to 1^-} \frac{1-\gamma}{1-\gamma_B(\gamma)} = 0. \tag{11}$$

Humans, Systems, and Security

- Human-on-the-Loop Autonomy
  - Complex systems that involve both autonomous and human agents with overlapping roles
- Research Question
  - How to build security-aware human-autonomy interaction with performance guarantees?
- Motive
  - Collaboration rather than complete autonomy
  - Ignoring human factors during design phase may impact system performance
  - How does human presence impact various system performance measures?
  - Human context awareness (in real-time) as part of security analysis/design?

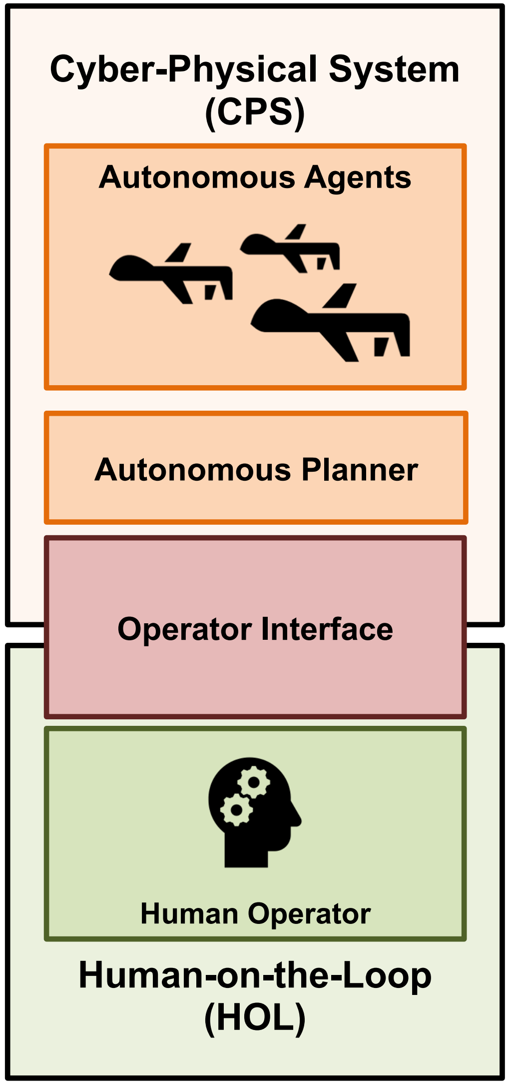

RESCHU-SA Testbed

- Simulation environment for human-UAVs command and control systems
- Extendable, open source

Security-Aware Features

- ***Live Camera Feed***
  Camera always streams the ground truth

- ***Attack Engine***
  Attack specifications: attack goals, when & where to attack a UV
  Attack model: aggressive vs stealthy

- ***Randomized Map***
  Randomly-generated map to ensure unbiased experiments and diverse features

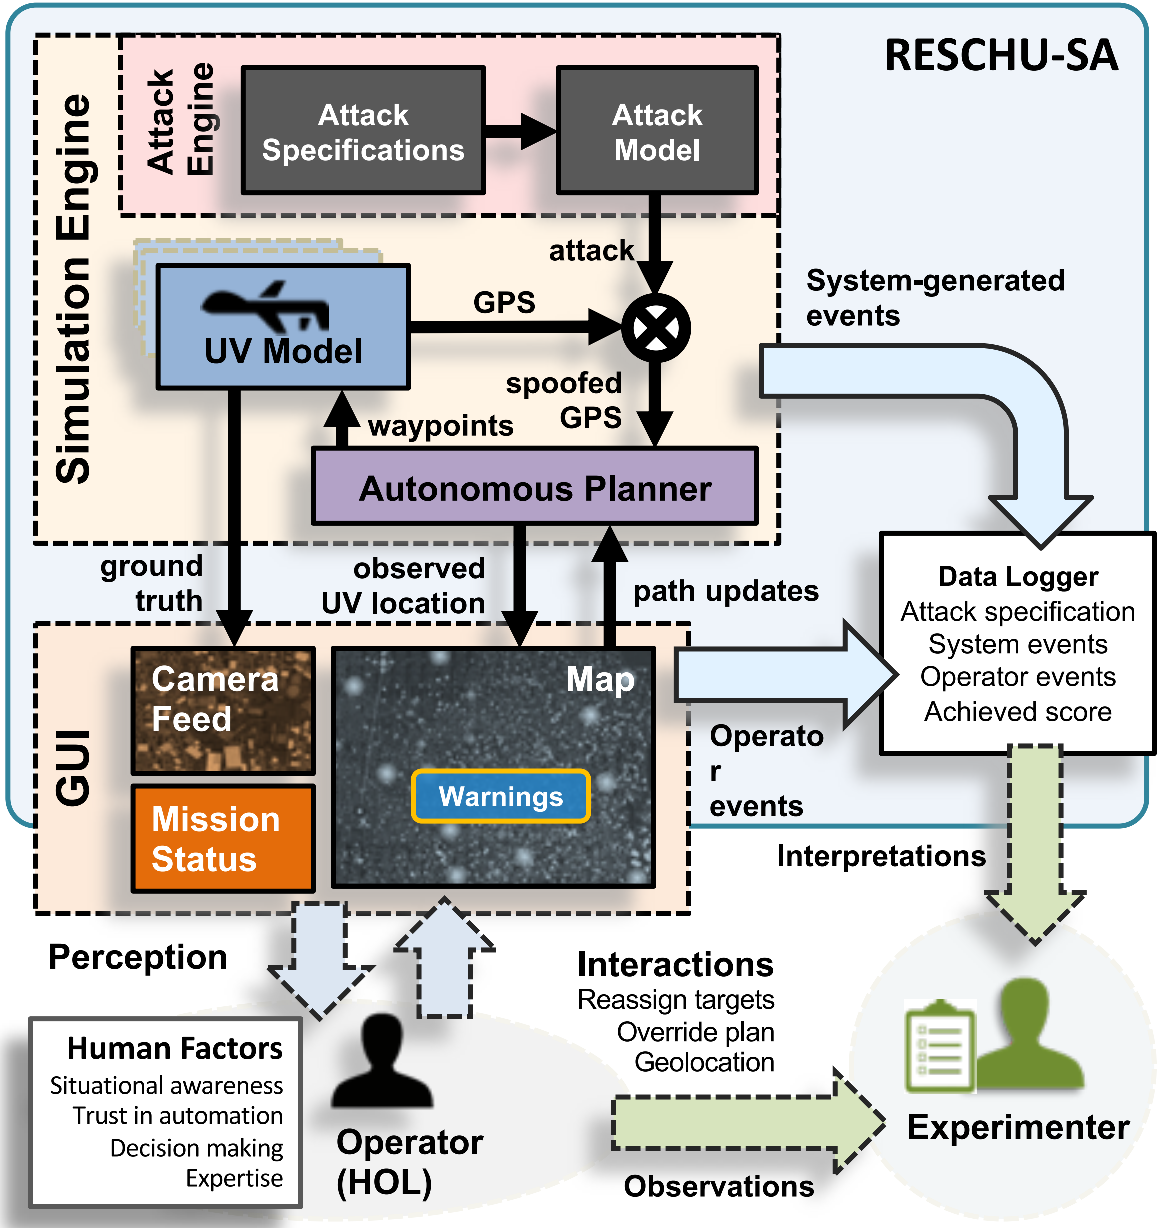

Security-aware Human-on-the-Loop Planning

Duke PRATT SCHOOL of ENGINEERING

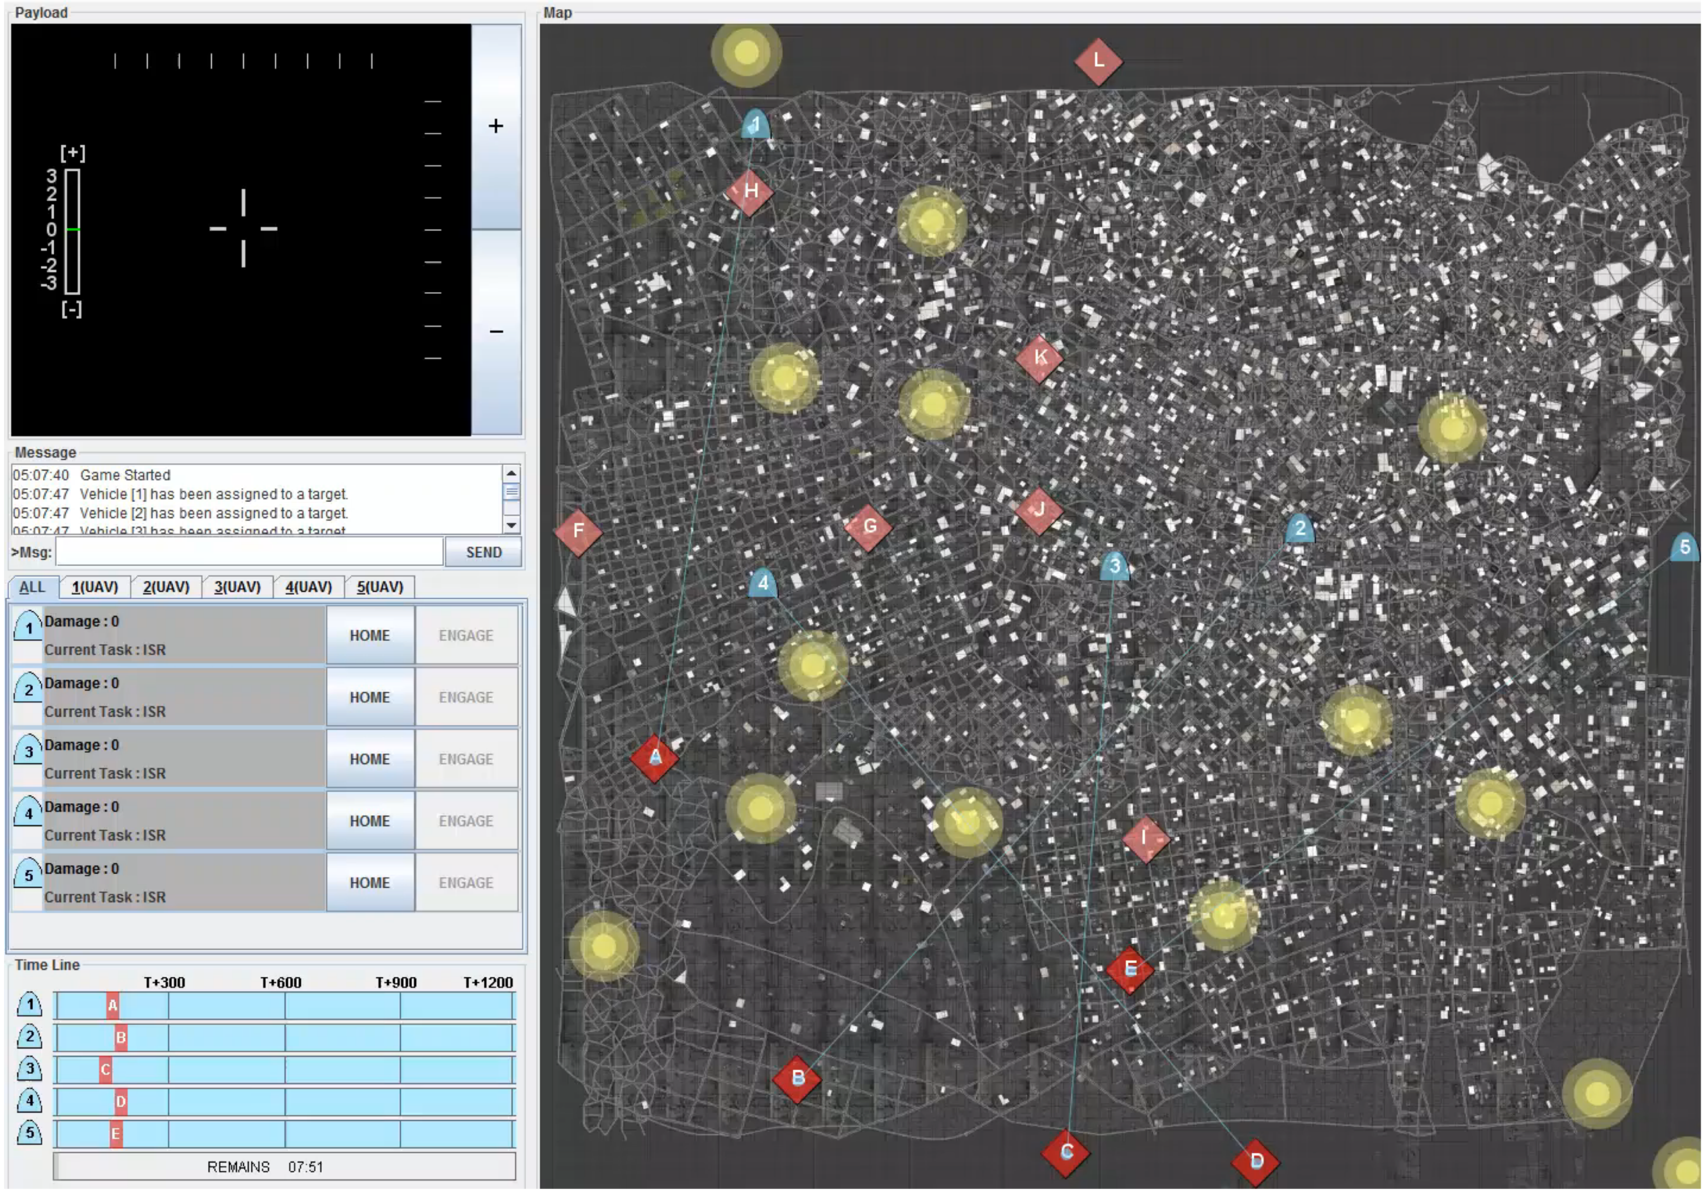

[1] M. Elfar, H. Zhu, M. L. Cummings, and M. Pajic, “Security-Aware Synthesis of Human-UAV Protocols”, *IEEE Int. Conf. on Robotics and Automation* (**ICRA**), 2019.

[2] H. Zhu, M. Cummings, M. Elfar, Z. Wang, and M. Pajic, “Operator Strategy Model Development in UAV Hacking Detection”, *IEEE Trans. on Human-Machine Systems*, Dec. 2019.

Experimental Setup – Understanding Human Geolocation Strategies and Context-Awareness

Duke
PRATT SCHOOL of ENGINEERING

Two missions

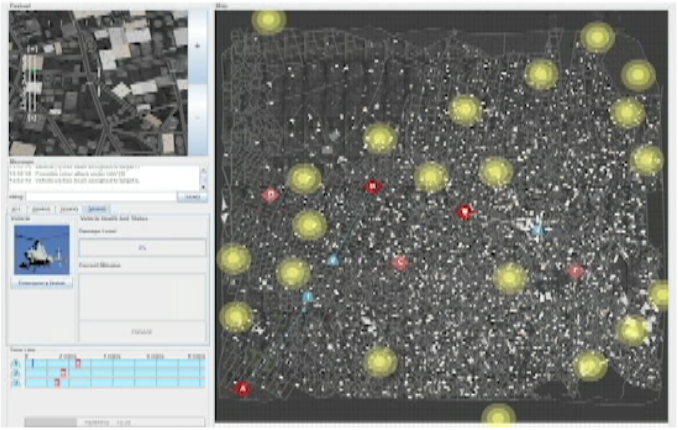

Low Workload

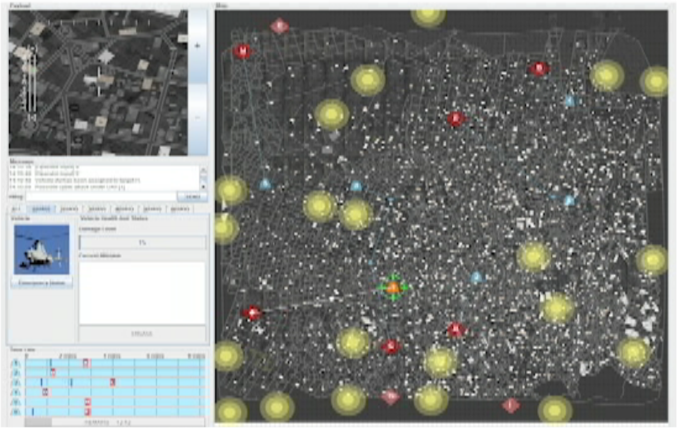

High Workload

[Footage from actual experiments at speed 5x]

36 Subjects

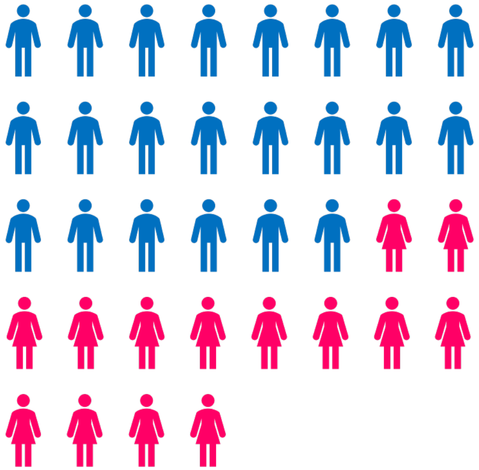

Scoring System

- Visit Target Locations
- Avoid Hazard Zones
- Detect Cyber Attacks
- Finish within the time limit

- Sustain damage from hazard zones
- Lose assets due to cyber attacks
- Miss targets by the time limit

[1] M. Elfar, H. Zhu, M. L. Cummings, and M. Pajic, ``Security-Aware Synthesis of Human-UAV Protocols", *2019 International Conference on Robotics and Automation* (**ICRA**), 2019.

[2] H. Zhu, M. Cummings, M. Elfar, Z. Wang, and M. Pajic, “Operator Strategy Model Development in UAV Hacking Detection”, IEEE Transactions on Human-Machine Systems, 2019.

Development of Operator Behavior Models in Human Supervisory Control Scenarios [IEEE THMS19]

Duke PRATT SCHOOL of ENGINEERING

Index	Observations
1	Add waypoint
2	Move waypoint
3	Delete waypoint
4	Move endpoint
5	Switch target
6	Engage task
7	Select UAV
8	Acknowledge notification
9	Ignore notification
10	Consider UAV hacked
11	Consider UAV not hacked
12	Adjust zoom level

Attack detection functional group

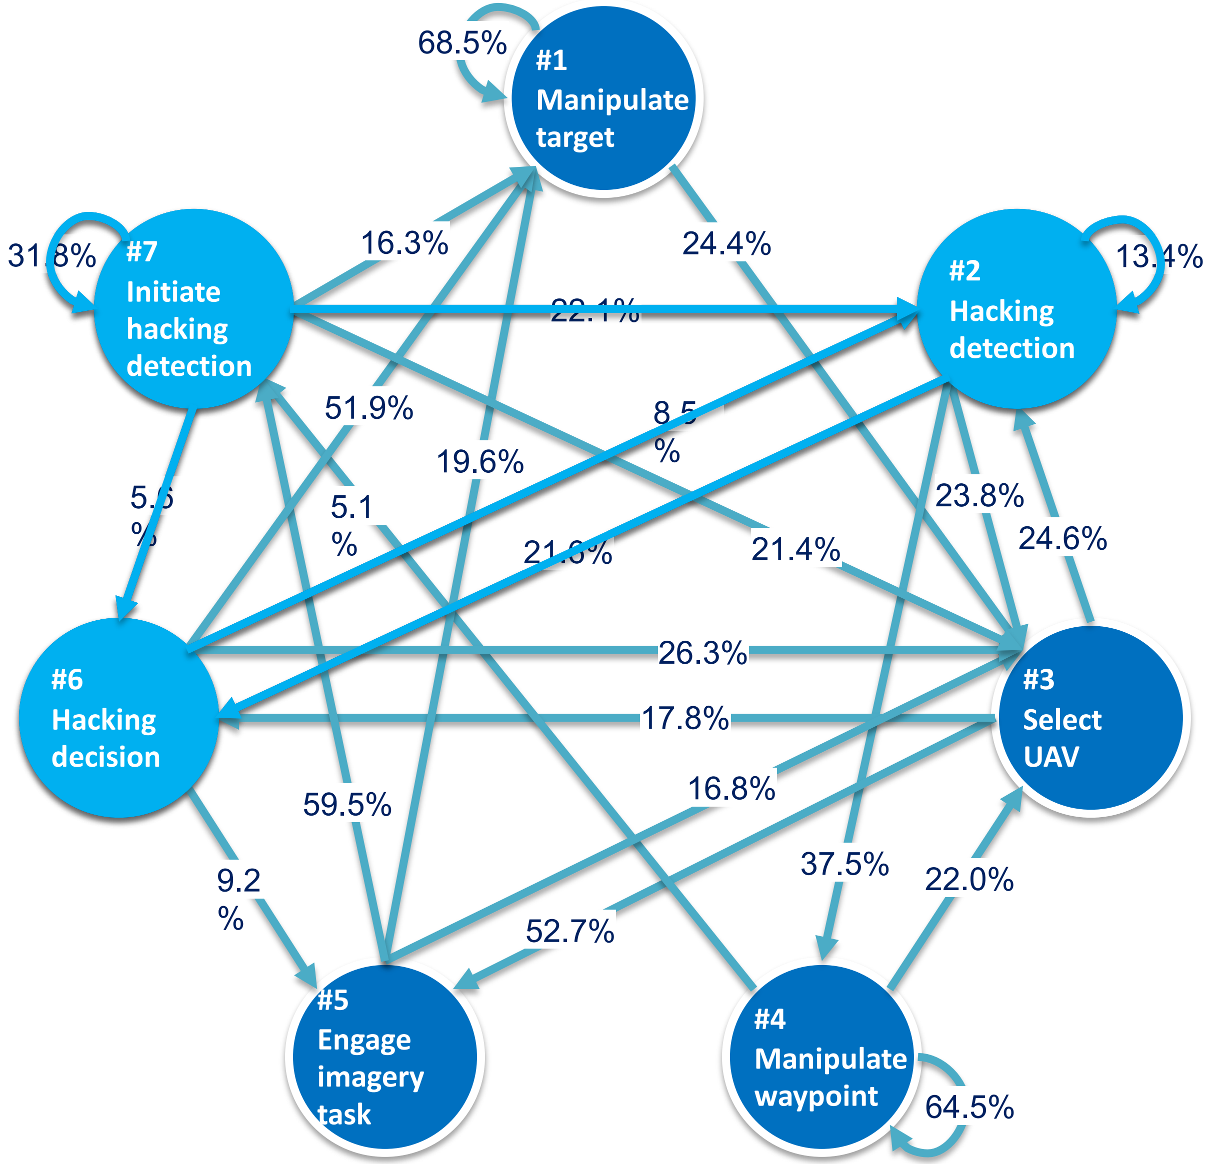

Navigation and imagery task functional group

Security-aware Human-on-the-Loop Planning

Environment Setup

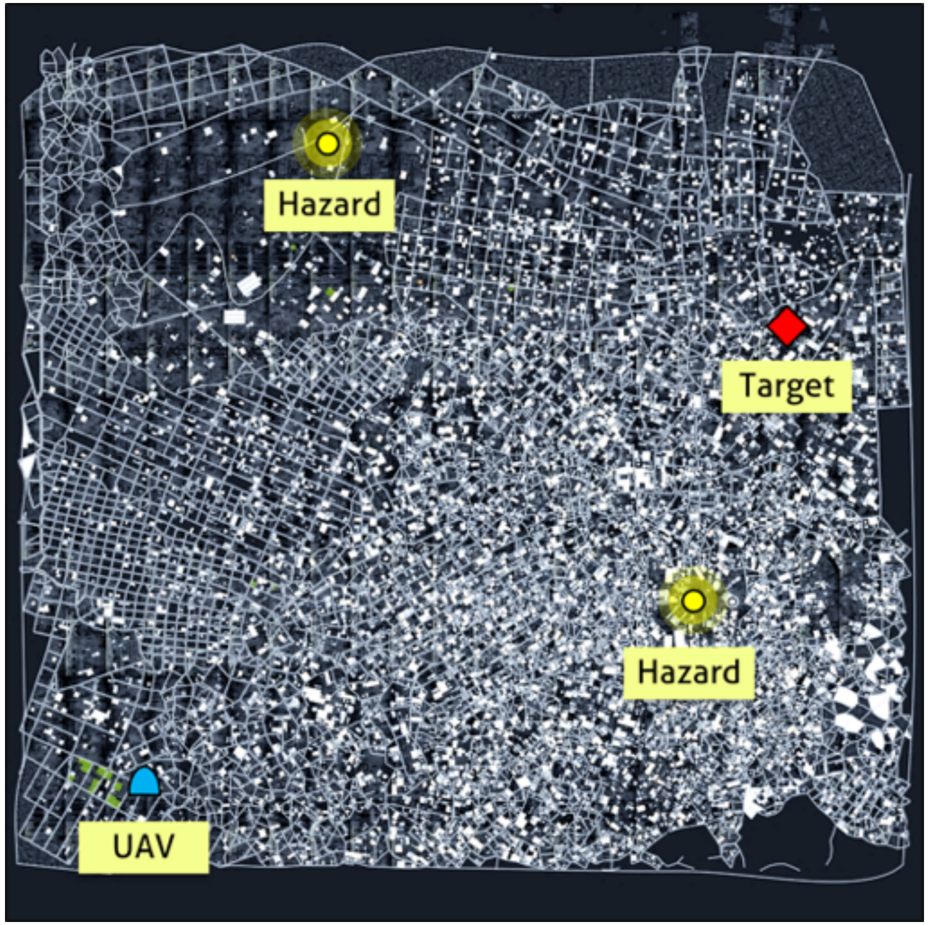

Synthesis Objectives

$$\phi_{\text{syn}}(k) := \langle\langle \text{uav} \rangle\rangle \Pr_{\max=?} \left[\neg \mathit{hazard} \, \mathsf{U}^{\leqslant k} \, (\mathit{locate} \wedge \mathit{reach}) \right]$$

$$\hat{\mathcal{G}}^{\{\pi_i\}_{i=0}^{q}} : \ \phi_{\text{ana}}(n) := \langle\langle \text{adv} \rangle\rangle \Pr_{\min=?} \left[\mathsf{F}^{\leqslant n} \, \mathit{target} \right]$$

Synthesis Procedure

Algorithm 2: Protocol synthesis procedure

Input: Initial location x_0, synthesis query ϕ_{syn}, max horizon $h_{\max}$
Output: H-UAV protocols $\Pi = \{(\pi_{\text{uav}}, \pi_h)\}$

$X \leftarrow \{x_0\}$ initialize set of initial locations (subgames)
foreach *unexplored initial location* $x_i \in X$ **do**
  $\hat{s}_0 \leftarrow (\text{UAV}, x_i, \epsilon)$ set subgame initial state
  $\mathit{stop} \leftarrow \bot$, $h \leftarrow 1$ reset stopping flag and horizon
  while $h \leqslant h_{\max} \ \wedge \ \neg \mathit{stop}$ **do**
    $(\pi_{\text{uav}}, \varphi) \leftarrow \mathit{synth}\left(\hat{\mathcal{G}}^{\pi_h}_{\hat{s}_0}, \phi_{\text{syn}}\right)$ find a winning strategy
    if π_{uav} *exists* **then**
      $\Pi \leftarrow \Pi \cup (\pi_{\text{uav}}, \pi_h, \varphi)$ add to the protocol
      $X \leftarrow X \cup \mathit{reach}(\pi_{\text{uav}})$ update reachability set
      $h \leftarrow h + 1$ explore next horizon
    else $\mathit{stop} \leftarrow \top$
  $\mathit{prune}(\Pi)$

M. Elfar, H. Zhu, M. L. Cummings, and M. Pajic, "Security-Aware Synthesis of Human-UAV Protocols", *IEEE Int. Conf. on Robotics and Automation* (**ICRA**), 2019.

Security-aware Human-on-the-Loop Planning

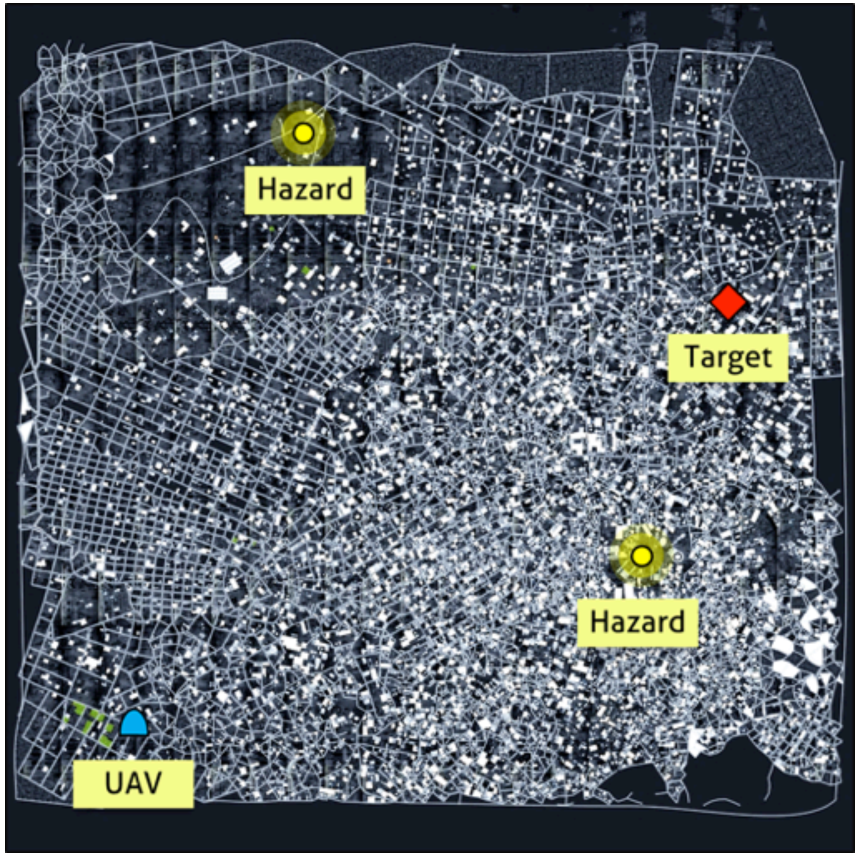

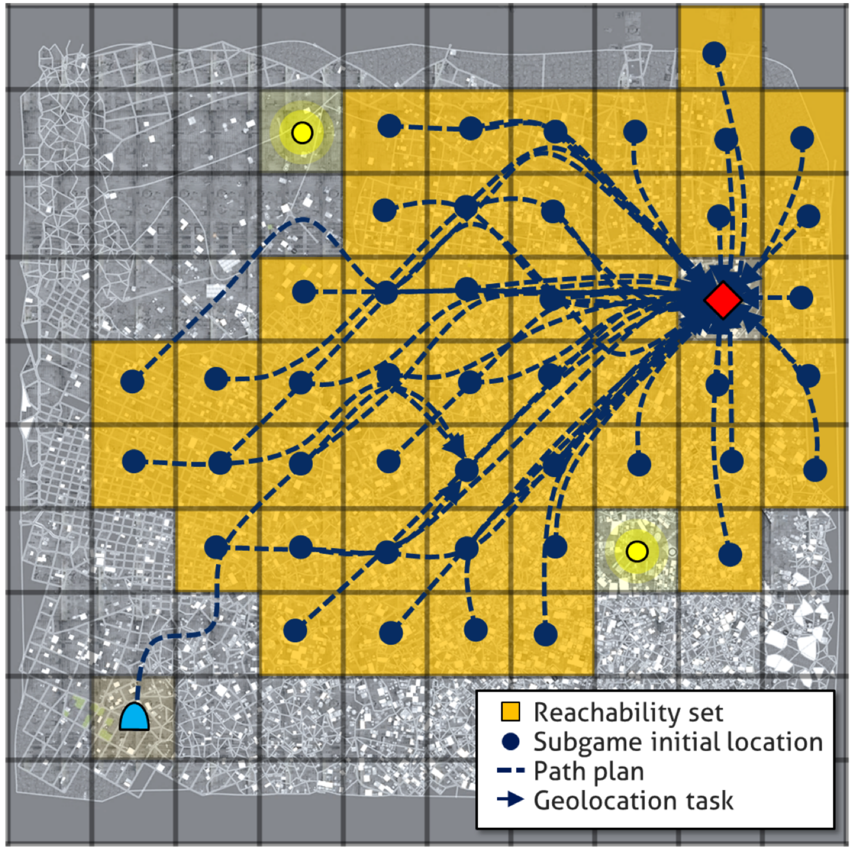

M. Elfar, H. Zhu, M. L. Cummings, and M. Pajic, "Security-Aware Synthesis of Human-UAV Protocols", *IEEE Int. Conf. on Robotics and Automation* (**ICRA**), 2019.

Security-aware Human-on-the-Loop Planning

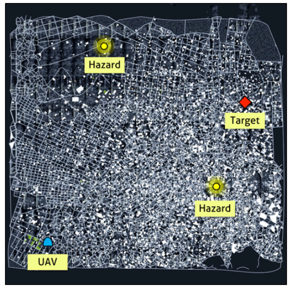

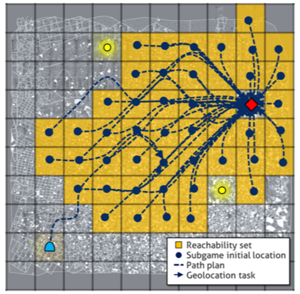

Analysis Results (PRISM-games)

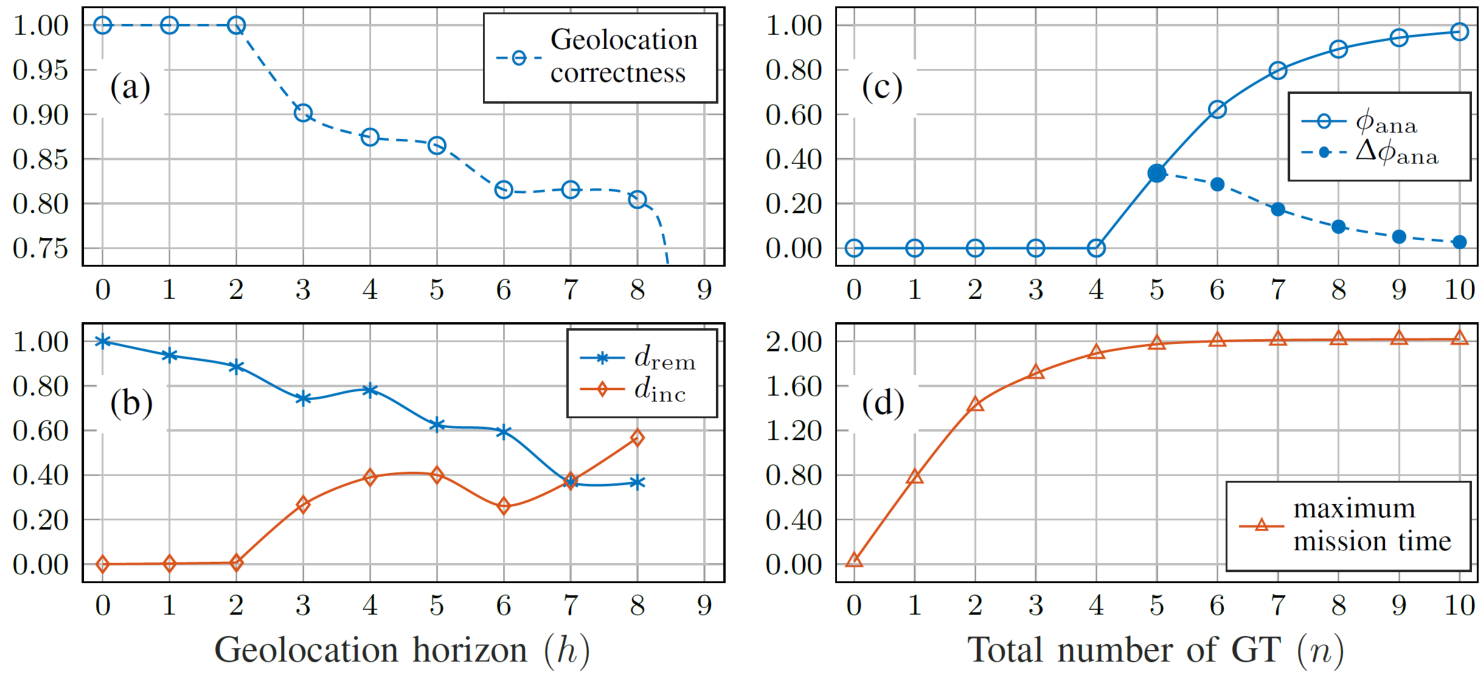

User Evaluation

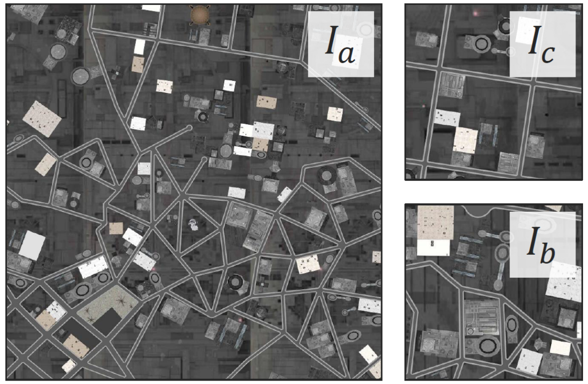

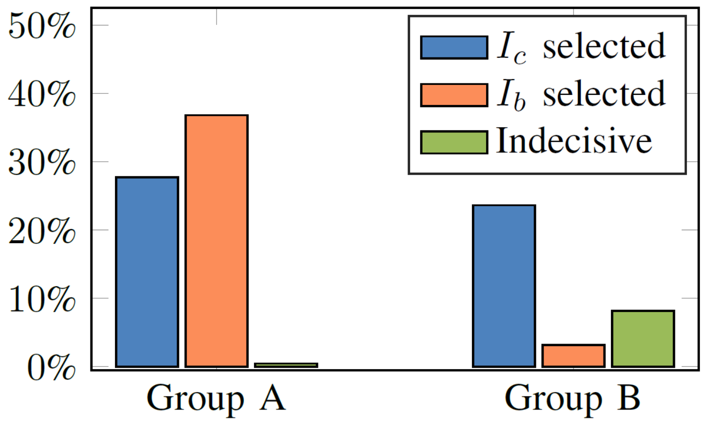

M. Elfar, H. Zhu, M. L. Cummings, and M. Pajic, “Security-Aware Synthesis of Human-UAV Protocols”, *IEEE Int. Conf. on Robotics and Automation* (**ICRA**), 2019.

Thank you

UF | UNIVERSITY of FLORIDA

Duke UNIVERSITY

TEXAS The University of Texas at Austin

UC SANTA CRUZ