AACOE Updates on Protecting Information and Computation Privacy

Kevin Butler, University of Florida

AACOE Program Review

26 April 2023

Human-Machine Interaction

- Current state of deployed UASs involve significant human interaction (<=L3 autonomy)
- Autonomous systems will potentially learn from simulation data informed by human interaction
- Augmented reality (AR) systems can assist near-term operations while virtual reality (VR) simulators are standard for training
- What risks to privacy are incurred in these systems?

Re-Identification Risk

Gaze Datasets:
- **ET-DK2 (N = 18)**
- **360_em (N = 13)**

360◦ VR Viewing

k*-anonymity: 1 / *k

Attack Success Rate
- **Raw gaze data**
- **Anonymized gaze data**

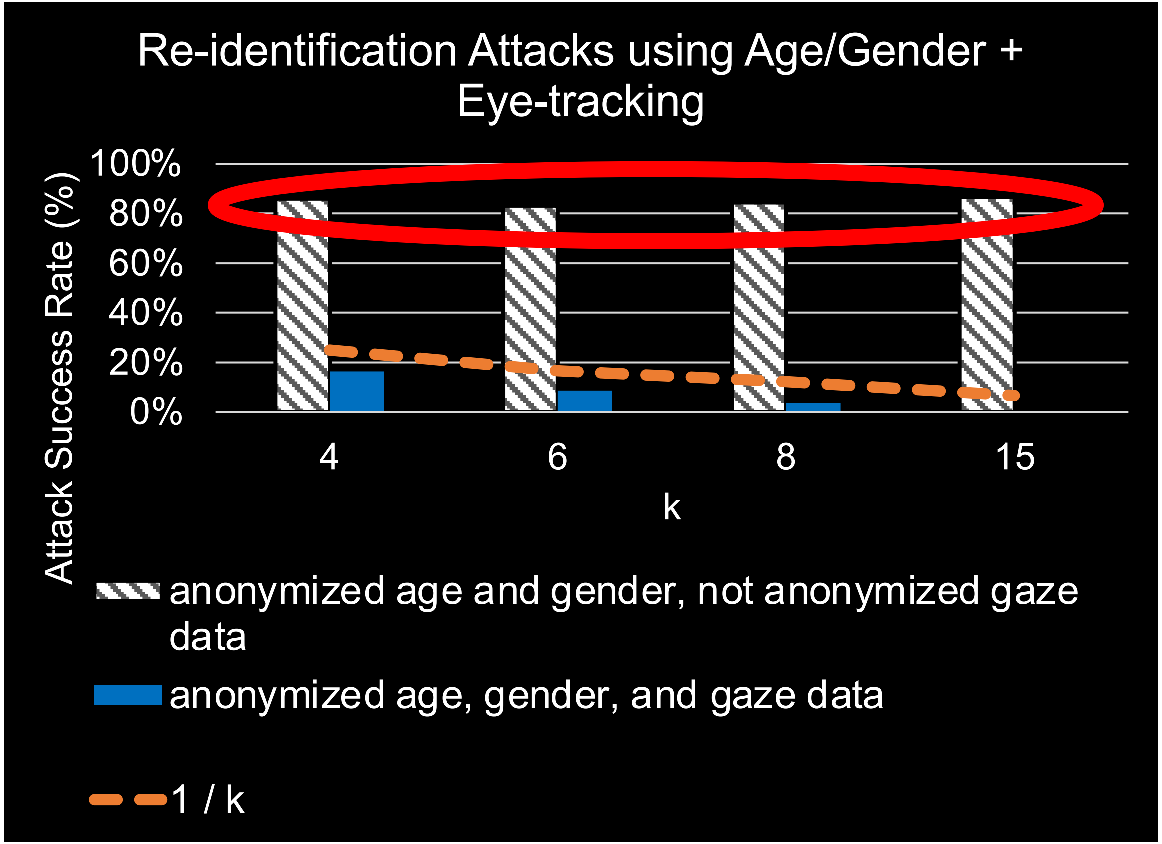

kalεido-DP [Li et al., USENIX 2021]
Spatial Differential Privacy (DP)
Can be applied in real-time or to dataset

Past Work: Noisy samples
Real-time, No Privacy Guarantee

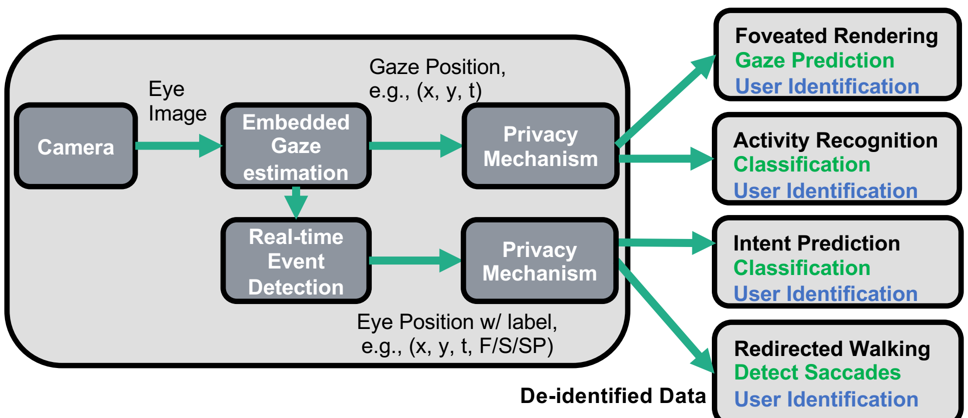

[David-John et al., TVCG 2021]

This Work: Synthetic samples
Dataset, Privacy Guarantee

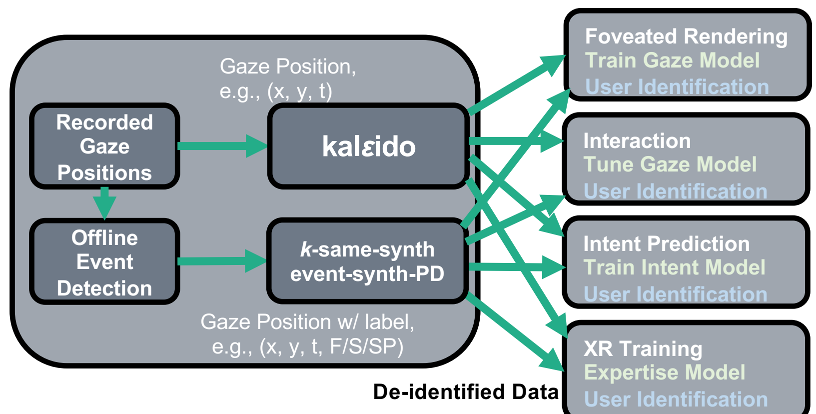

[David-John et al., TVCG 2023]

Evaluation of XR Applications

DGaze dataset (**N** = 43)
Gaze Prediction (100 ms), 3D Scene

Mechanism	Identification Rate (↓)	Runtime (↓)
Raw Data	2%	N/A
k-same synth (ours)	1.1%	52 sec
Event-synth-PD (ours)	1.3%	4 min
kalɛido-DP	2.1%	2 min

EHTask dataset (**N** = 30)
Activity Classification, 360° Video

Mechanism	Identification Rate (↓)	Runtime (↓)
Raw Data	28%	N/A
k-same synth (ours)	7.5%	2 min
Event-synth-PD (ours)	9.2%	15 min
kalɛido-DP	6.0%	5 min

UF UNIVERSITY of FLORIDA · Duke UNIVERSITY · TEXAS The University of Texas at Austin · UC SANTA CRUZ

Evaluation of XR Applications

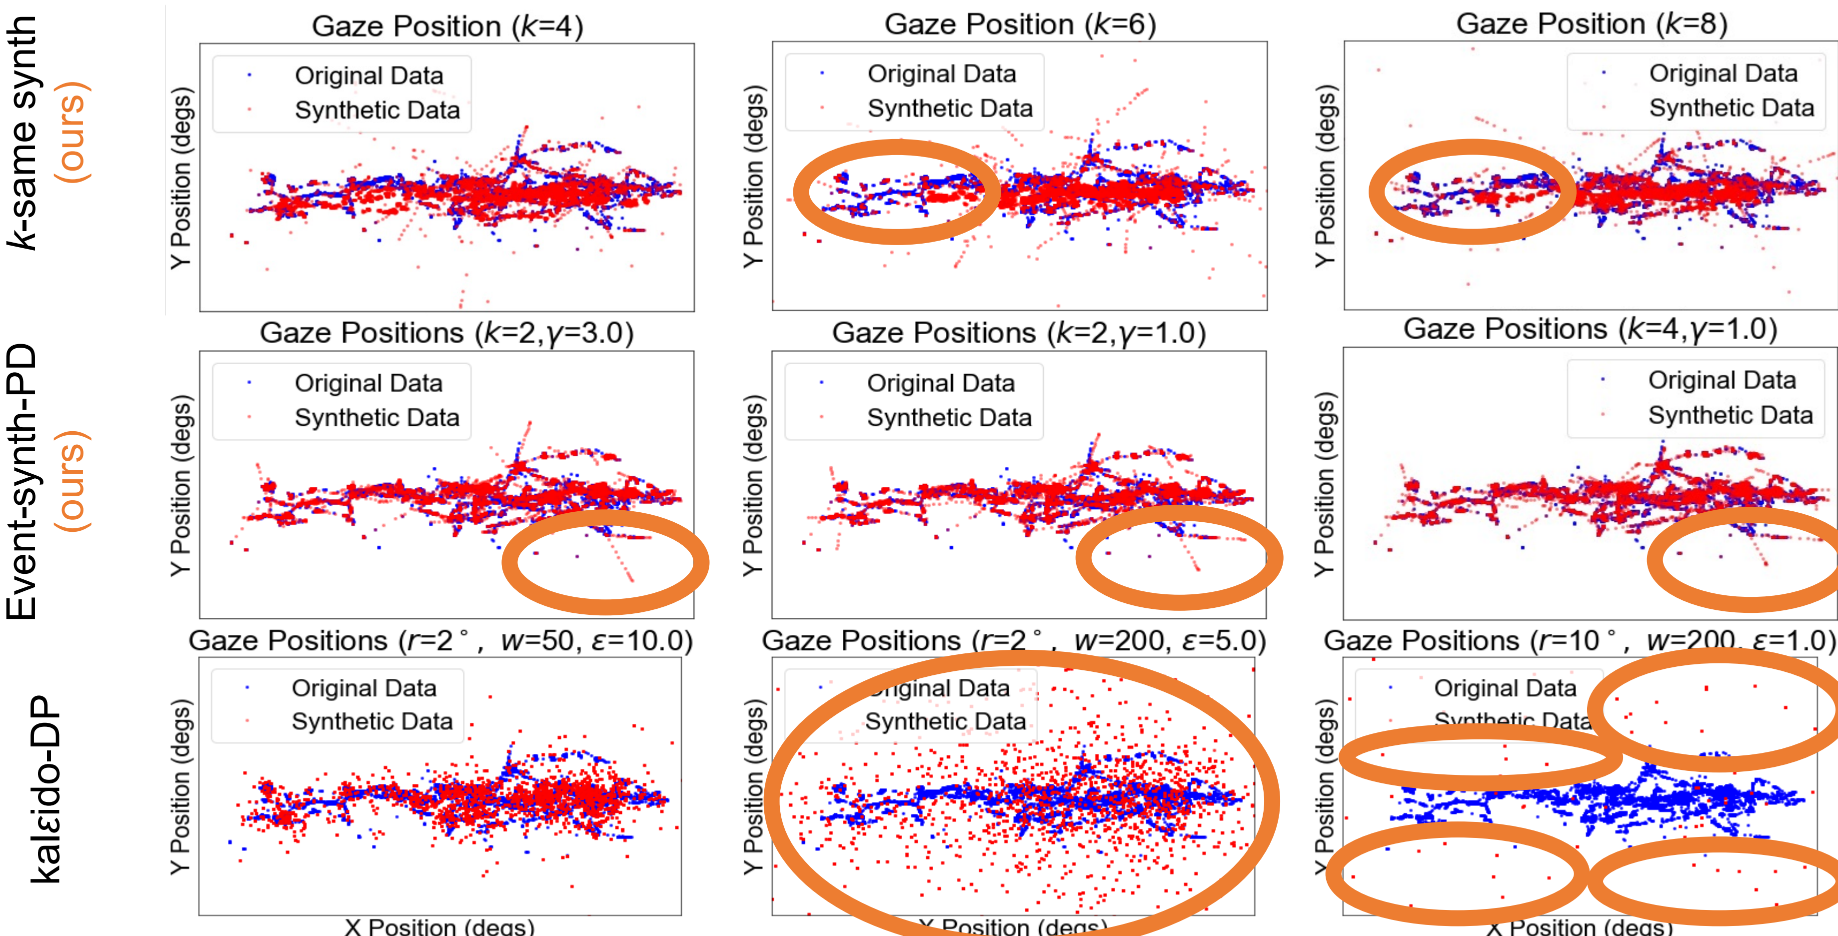

Hard-Label Attacks

Hard-label adversarial machine learning attacks are a "grand-prize":

- Adversary only needs ***query access*** to generate "label-flipped" samples (e.g., through compromised user)
- Hard-label attacks are gaining popularity, but not well characterized apart from convergence guarantees.

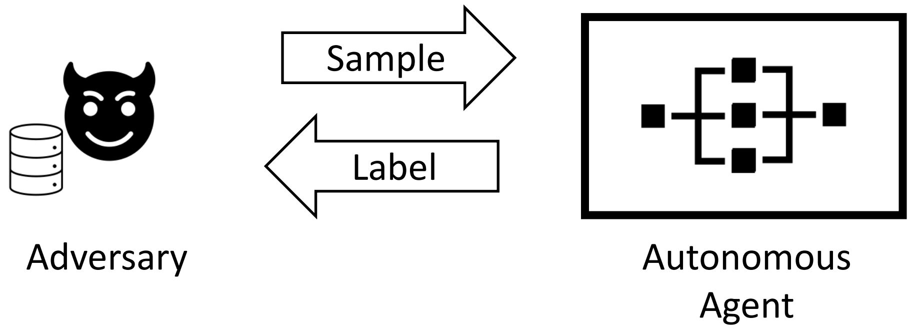

Subsampling and query efficiency

- Questions we sought to answer:
  - What advantages does search subsampling give the adversary?
  - How can we generalize the idea of search subsampling?
- We addressed this as an information-theoretic problem, leveraging the data processing inequality to derive a close—form solution of manifold-gradient mutual information

$$I(\mathcal{G},\mathcal{M})_\epsilon = 2\int_{\mathcal{M}^+} p(1,x^+)\log(\frac{p(1,x^+)}{p_{\mathcal{G}}(1)p_{\mathcal{M}}(x^+)})\,dx^+ + 2\int_{\mathcal{M}^+} p(-1,x^+)\log(\frac{p(-1,x^+)}{p_{\mathcal{G}}(-1)p_{\mathcal{M}}(x^+)})\,dx^+.$$

$$I(\mathcal{G},\mathcal{M})_\epsilon = \frac{2}{\sqrt{2\pi\sigma^2}}\sum_{i=1}^{||\mathcal{M}^+||}\exp(-\frac{(x_i^*-\theta)^2}{2\sigma^2})\cdot\beta_i^+\Delta_i + \frac{2}{\sqrt{2\pi\sigma^2}}\sum_{i=1}^{||\mathcal{M}^+||}\exp(-\frac{(x_i^*+\theta)^2}{2\sigma^2})\cdot\beta_i^-\Delta_i.$$

Results of Dimensionality Reduction

HSJA | BiLN+HSJA | AE+Sign-OPT

Adversarial

Diff

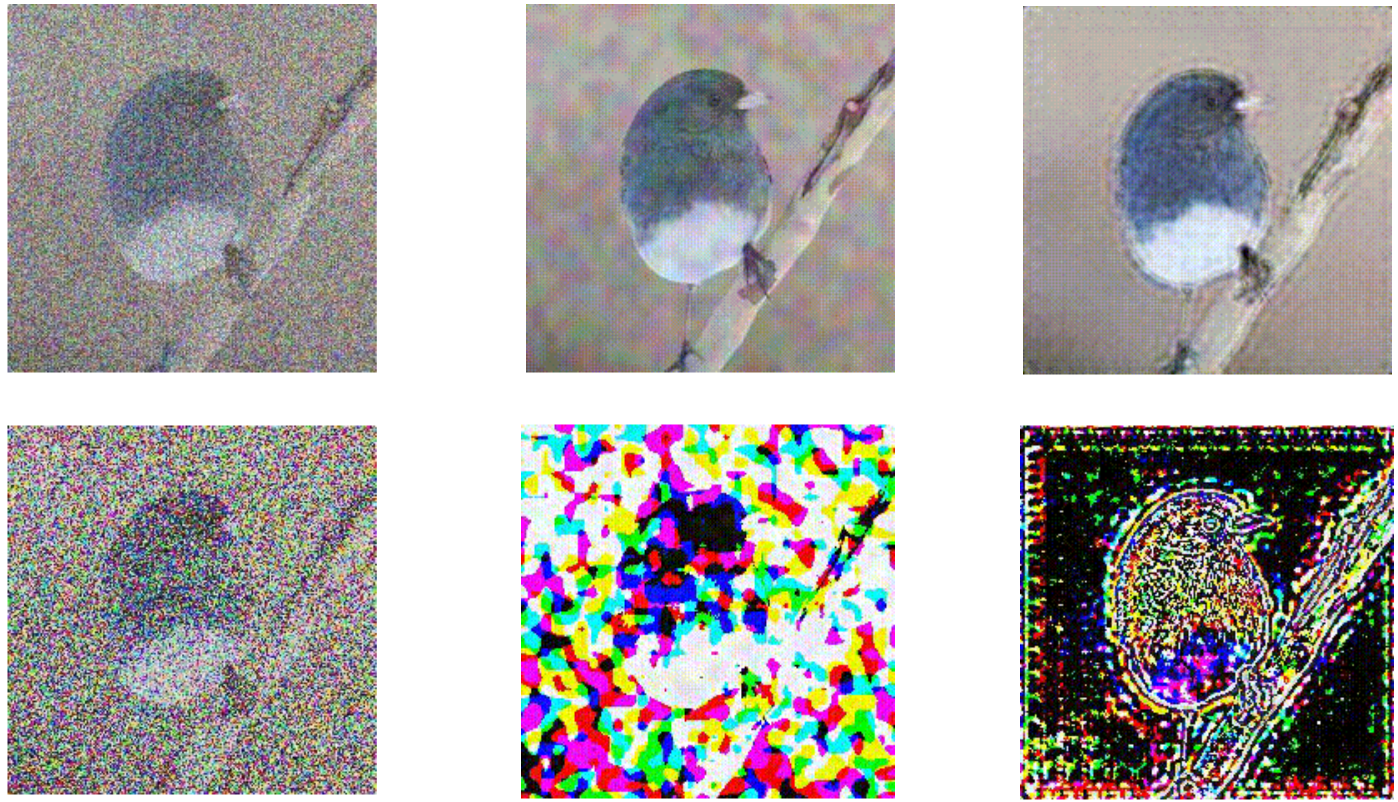

Geometric Interpretation

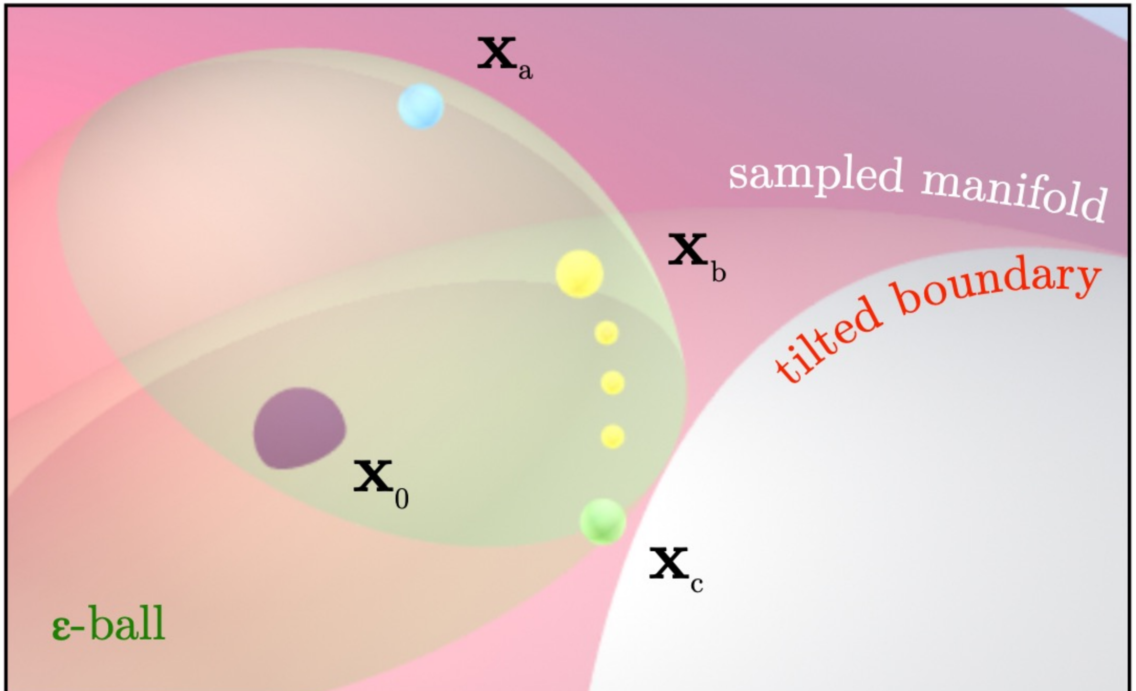

UF UNIVERSITY of FLORIDA

Duke UNIVERSITY

TEXAS The University of Texas at Austin

UC SANTA CRUZ

Protecting Satellite Proximity Operations via Secure Multi-Party Computation

Caroline Fedele

PhD student, University of Florida

26 April 2023

Project Overview

General goal: provide evaluation of secure satellite proximity operations using privacy-preserving computation

Demonstrate use of **secure multiparty computation (SMC)**, a method of operating on encrypted data, allowing **satellite operations** to be conducted between mutually-distrustful agents without leaking information about satellites' capabilities

- Investigating existing tools and SMC approaches with which to implement SMC
- Determining relevant problems in space/satellite research where privacy is a concern
- Prototyping SMC setup for satellites on embedded boards for autonomous operations
- Evaluating algorithms with and without SMC: matrix multiplication, RPO algorithms
- Benchmarking overhead added by SMC
- Broader characterization problems

Motivation: RPO

Rendezvous and Proximity Operations (RPO):

- On-board trajectory operation and replanning
  - E.g. docking, on-orbit servicing/refueling, formation flying
- RPO occurs on-board, autonomously
  - housed in guidance navigation and control (GNC) unit
- Needed at scales of < 500km between satellites

RPO example: docking

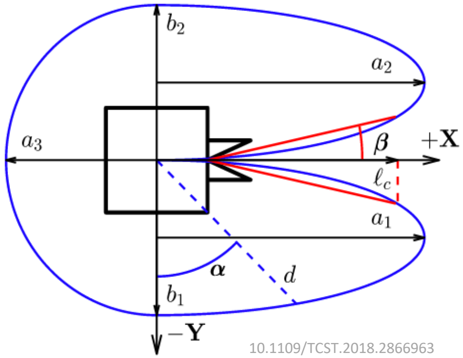

10.1109/TCST.2018.2866963

Ground station vs On-board Control

	Ground station	On-Board
Distance between satellites	1-10 Mm	< 500 km
Time needed	Days-weeks	< 1 day
Speed	km /sec	m /sec
Approach	conjunction analysis	RPO

Problem: Capability Inference

Example: Collision Avoidance in RPO

- Minimum data to share with other satellites
  - position, velocity, ***covariance***

Stochastic systems

- Probabilistic, not deterministic
- Covariance matrices = quantify uncertainty
  - defined by ellipsoid
- Measure of TRUST, decisions based on accuracy

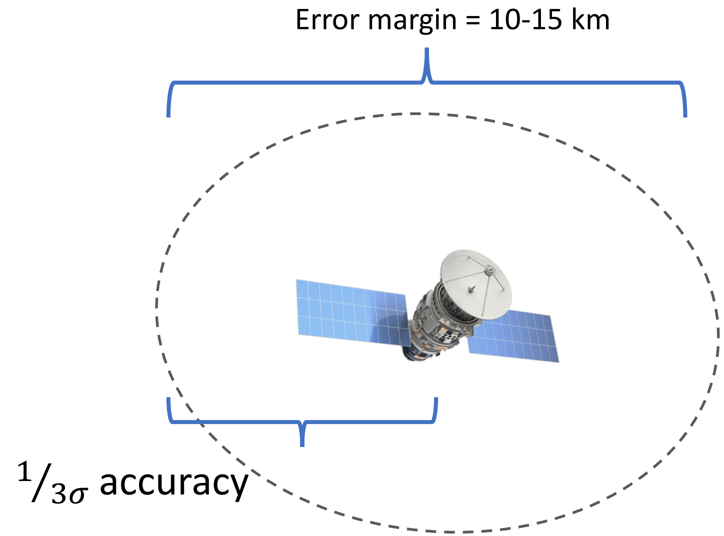

Problem: knowledge of error margins (covariance matrices) can lead to inferences on satellite capabilities, purpose, etc.

Solution: protect error margins using **privacy-preserving computation**

Privacy-Preserving Computation

Privacy-Preserving Computation (PPC)

- Allows for data to remain encrypted during computation
- Protects **physical integrity** of satellite during RPO and **data privacy** keeping data encrypted

Secure Multiparty Computation (SMC):

- Promising, well-developed method of PPC
- Cryptographic protocol that allows set of mutually-distrusting parties to jointly compute a function on their inputs, without revealing information about inputs (millionaire's problem)
- uses a) *garbled circuits* (2 parties) or b) *linear secret sharing* (>2 parties)

Linear Secret Sharing (LSS) scheme:

- keyless distributed encryption process.
- divides the "secret" (inputs) into randomly-generated shares
  and distributes to computing parties.

Background: What is SMC?

Secret Sharing

Distribute secret (input) among *n* parties, i.e. covariance matrices of 2 satellites. Predefined authorized subsets of *n* can reconstruct secret and return to user

Threshold Secret Sharing

- k-out-of-n scheme
- secret *S* divided into *n* shares: $S = (s_1,\ldots,s_n)$
  - *S* = element of finite field
  - shares = mapping to *S* + several random elements
- compromise of k-1 shares gives no info about S

Secret Sharing on Satellites

- Donors/data users = satellites participating in collision avoidance (at least 2)
- Miners = 3 computation servers
- Challenges: latency, bandwidth, small overhead on limited-resource system

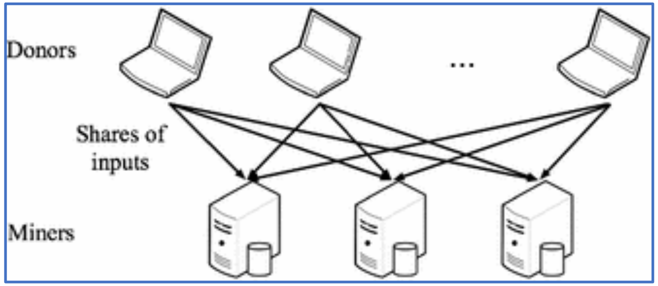

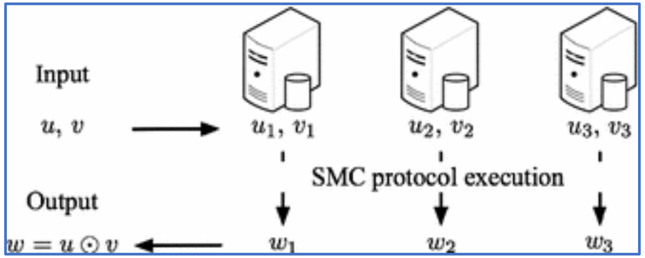

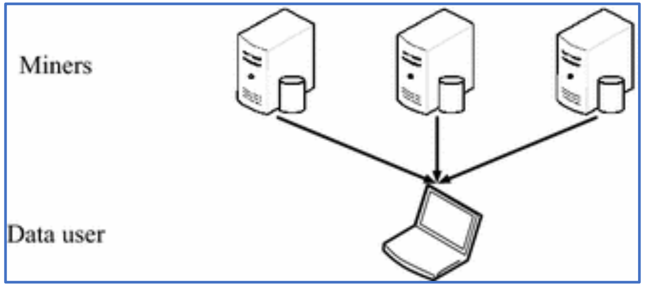

DOI 10.1007/s10207-014-0271-8

UF UNIVERSITY of FLORIDA · Duke UNIVERSITY · TEXAS The University of Texas at Austin · UC SANTA CRUZ

Methodology: Software

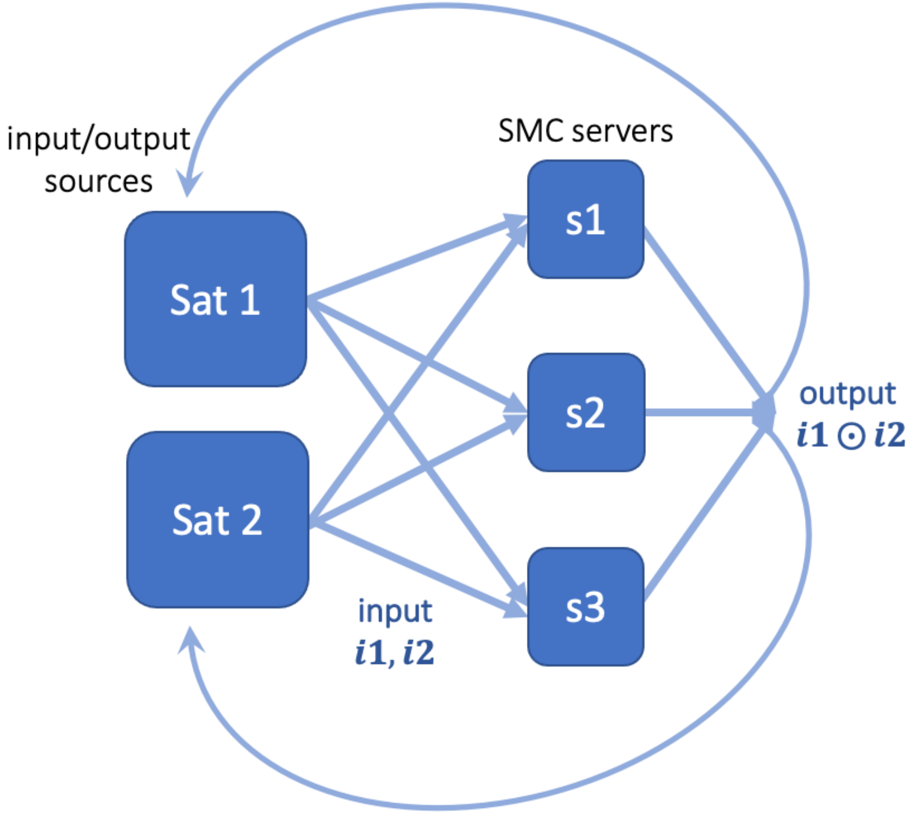

Integrating SMC into satellite operations

- Testing different algorithms
  - Matrix multiplication
  - Artificial potential function
  - Attitude Optimization

Software toolkit

- Sharemind MPC platform
  - 3-party linear secret sharing
  - Provides host for SMC operations
  - System of libraries compatible with C/C++ and proprietary SecreC code

Methodology: hardware

Finding hardware for deployment in space

- Considerations:
  - Commercial off-the-shelf (COTS)
  - Sufficient radiation tolerance
  - Sufficient power & efficiency with limited resources
- Current findings:
  - NVIDIA TX2/nano boards (ARM processors)
  - AMD Ryzen embedded boards (x86 processors)

Emulate satellite cluster

- Prototype with 3 Intel NUC boards
- Networked to communicate with each other
- 3 satellites minimum needed for SMC

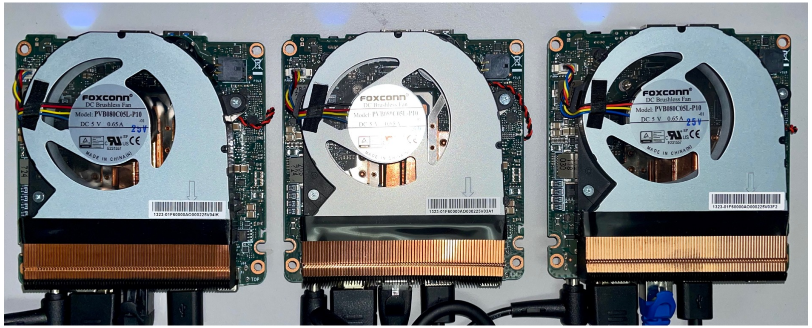

Hardware setup

Cluster of satellites (Hawkeye 360)

Press Release, 2020. https://www.he360.com/hawkeye-360-completes-milestone-in-preparation-to-launch-second-cluster/

Docking Algorithm

Another example: Artificial Potential Function (APF)

- Scenario: docking & collision avoidance at close range
    - On-board trajectory control
- Linear (relative) equations of motion

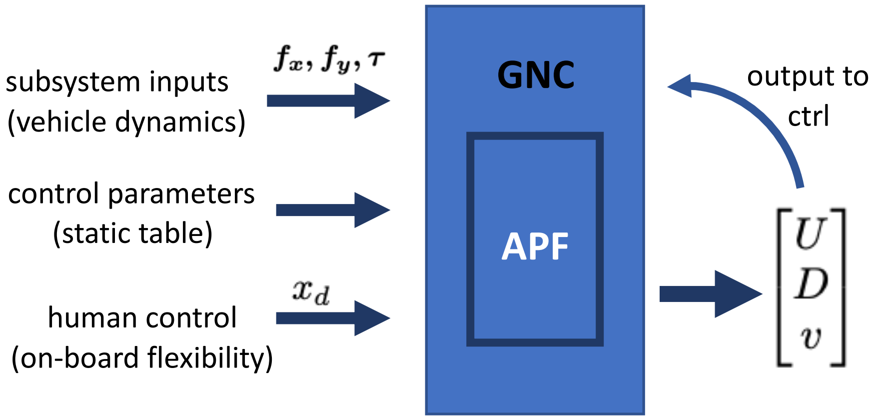

$$\begin{bmatrix} U \\ D \\ v \end{bmatrix} = apf(f_x, f_y, \tau, x_d, ..., \text{static params})$$

encrypted values

Keep-out zone potential

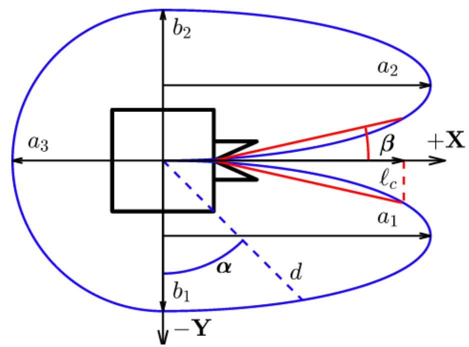

$$\phi_r = \sum_{i=1}^{N} \psi_i \exp\left[-\frac{\mathbf{r}_{co_i}^T \mathbf{N}_i \mathbf{r}_{co_i}}{\sigma_i}\right]$$

10.1109/TCST.2018.2866963

Optimization Algorithm

RPO example: Attitude Optimization

- Example scenario: inspecting downed satellite for sake of servicing, cooperation between satellites of different agencies/countries
- command torque to guide attitude of system to zero
- Blended cost approach:
    - Optimize **fuel** & **ending state**

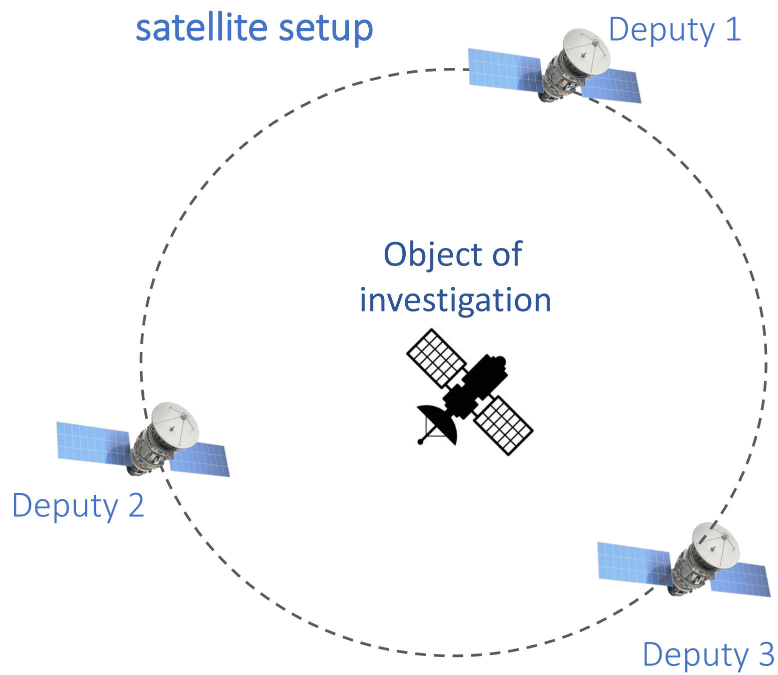

Privatized shared parameters

- Initial states: $\omega_1, \omega_2, \omega_3,\ v_1, v_2, v_3$
- Principle inertia: J, J_2, J_3

Evaluation: Matrix Multiplication

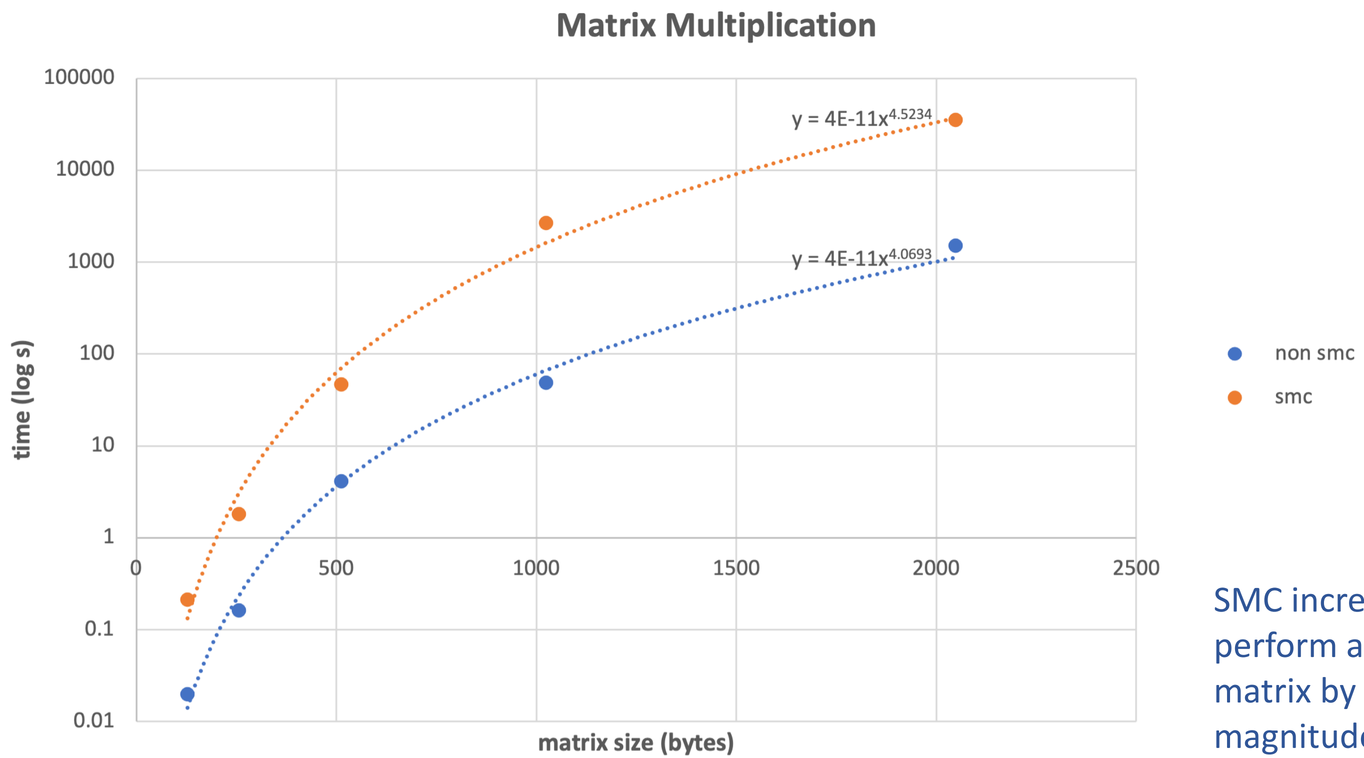

SMC increases time to perform algorithm on each matrix by 1-1.5 orders of magnitude

Evaluation: APF

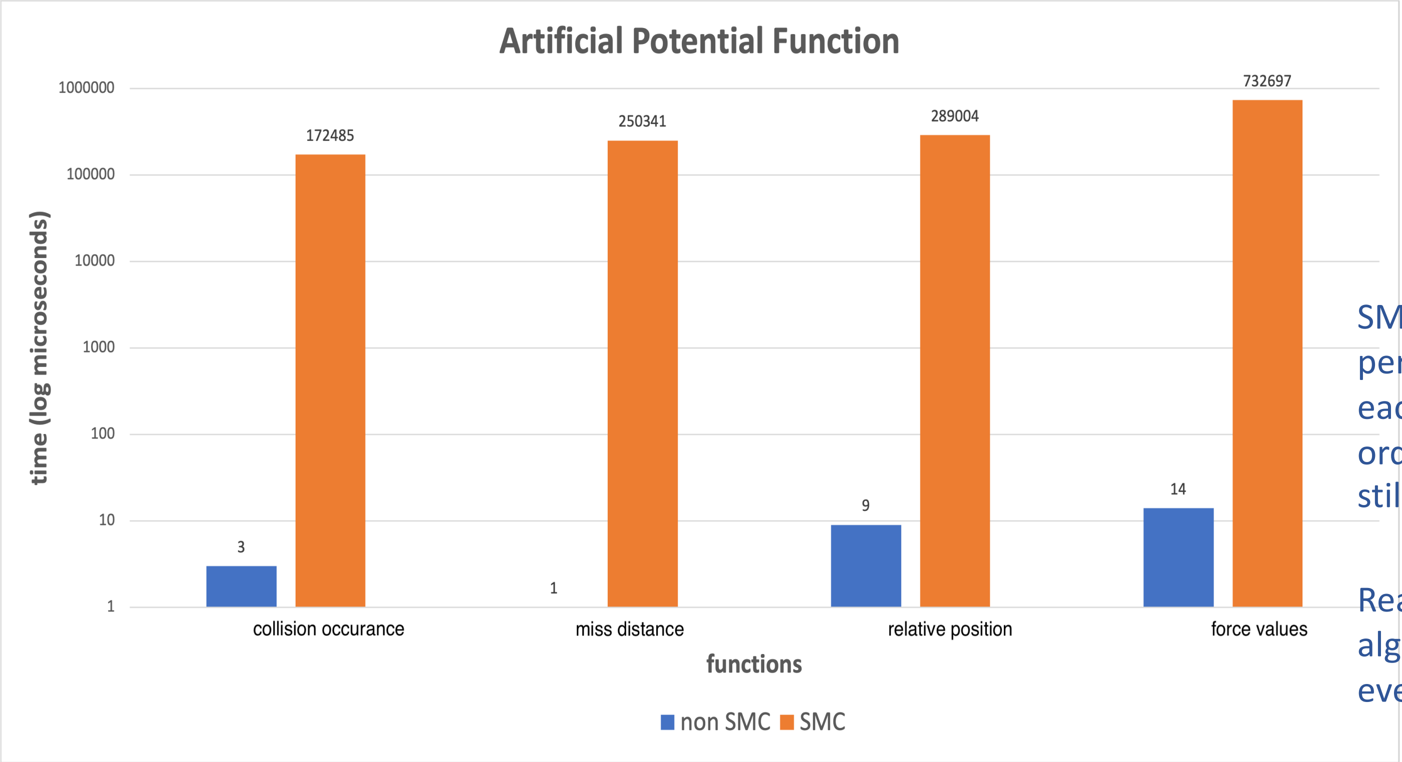

SMC increases time to perform algorithm on each function by 4-5 orders of magnitude, still < 1s to execute

Reasonable since algorithm refreshes every 30 sec– 15 min

Evaluation: Optimization

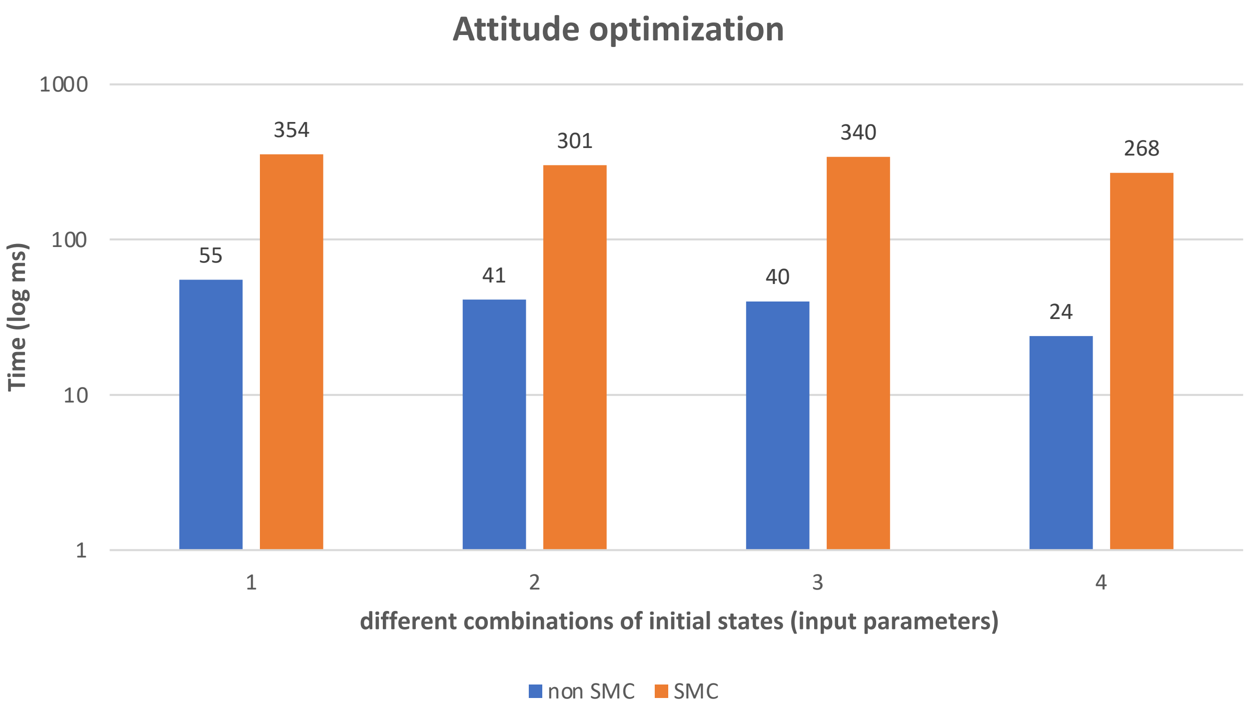

SMC increases time to perform algorithm on each function by ~1 order of magnitude

Algorithm refreshes every 10 seconds, still reasonable to use SMC

Results and Next Steps

Results:

- SMC adds 1-1.5 orders of magnitude of overhead in matrix multiplication
- SMC adds 4-5 orders of magnitude in APF algorithm functions
    - Each operation is still <1 second in this environment, **promising for SMC in practice**
- SMC adds ~1 order of magnitude in attitude optimization code

Next Steps:

- Test prototype on space-related hardware, i.e. NVIDIA and/or AMD boards
- Look into **efficiency** improvements
    - parallelization to increase efficiency
    - SIMD vectorization to improve scalability
- Investigate **characterization** problem beyond covariance matrices...

Characterization Problem

From Hardware

From Filters

From Learning

Concrete Observables
- Trajectory/Pose
- Reflectance/IR
- ....

Inferred Properties
- Propulsion System
- Material Properties
- ...

Inferred Application
- Science
- Position, Nav, & Timing
- ...

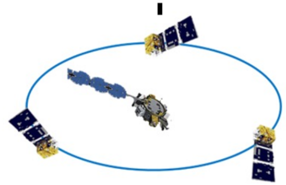

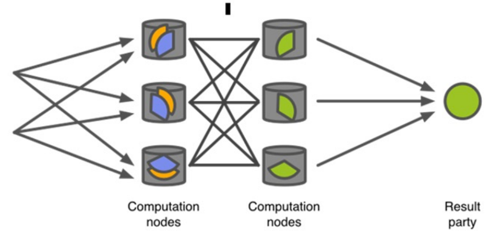

Secure Information To Protect Capability

Impact

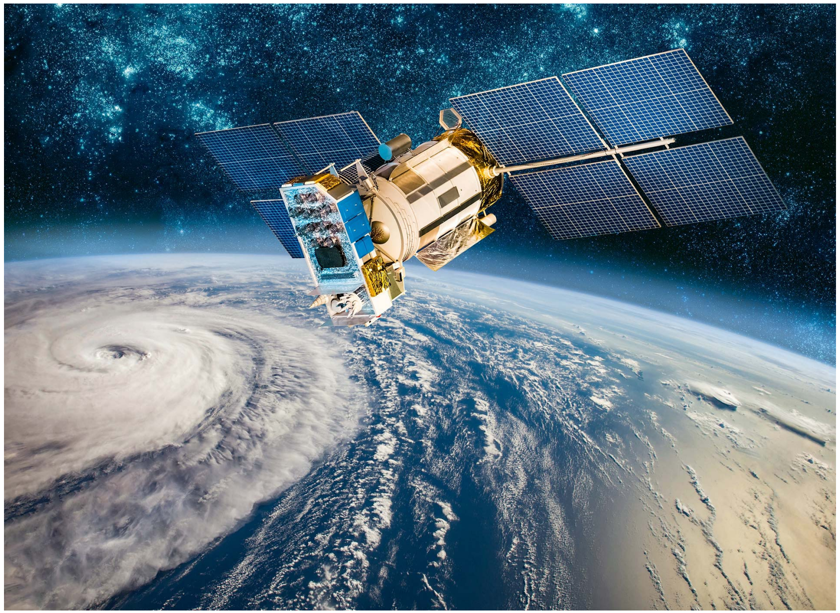

Source: verdict.co.uk

Action:

- Testing prototype of different satellite operations, specifically in RPO settings
- Integrating SMC into relevant space applications

Impact:

- Enhancing security in space, specifically problems where privacy is a concern
- Expanding applied cryptography/SMC to a new domain with these space applications

UF | UNIVERSITY of FLORIDA

Duke UNIVERSITY

TEXAS The University of Texas at Austin

UC SANTA CRUZ

Acknowledgements

Kevin Butler

Tyler Lovelly

Chris Petersen

Harsh Avathale & Shivam Gupta

FICS lab (UF)

SPACER lab (AFRL)