Resilience of Multi-Robot Systems to Physical Masquerade Attacks

Kacper Wardega, Roberto Tron & Wenchao Li
{ktw,tron,wenchao}@bu.edu

#1 Background

- Cloud robotics, e.g. FetchCore and Kiva.
- Security of industrial robot controllers (Quarta et. al. S&P'17).
- Spoof-resilience of multi-robot systems (Gil et. al Aut. Rob.'17).

#2 Questions

- Are Physical Masquerade Attacks a legitimate concern?
- How can system designers defend against such attacks?

#3 Threat Model

- Full plan-time control of a single robot.
- Inherits the control actions of non-compromised robots.
- Full information of planned routes and sensor capabilities.

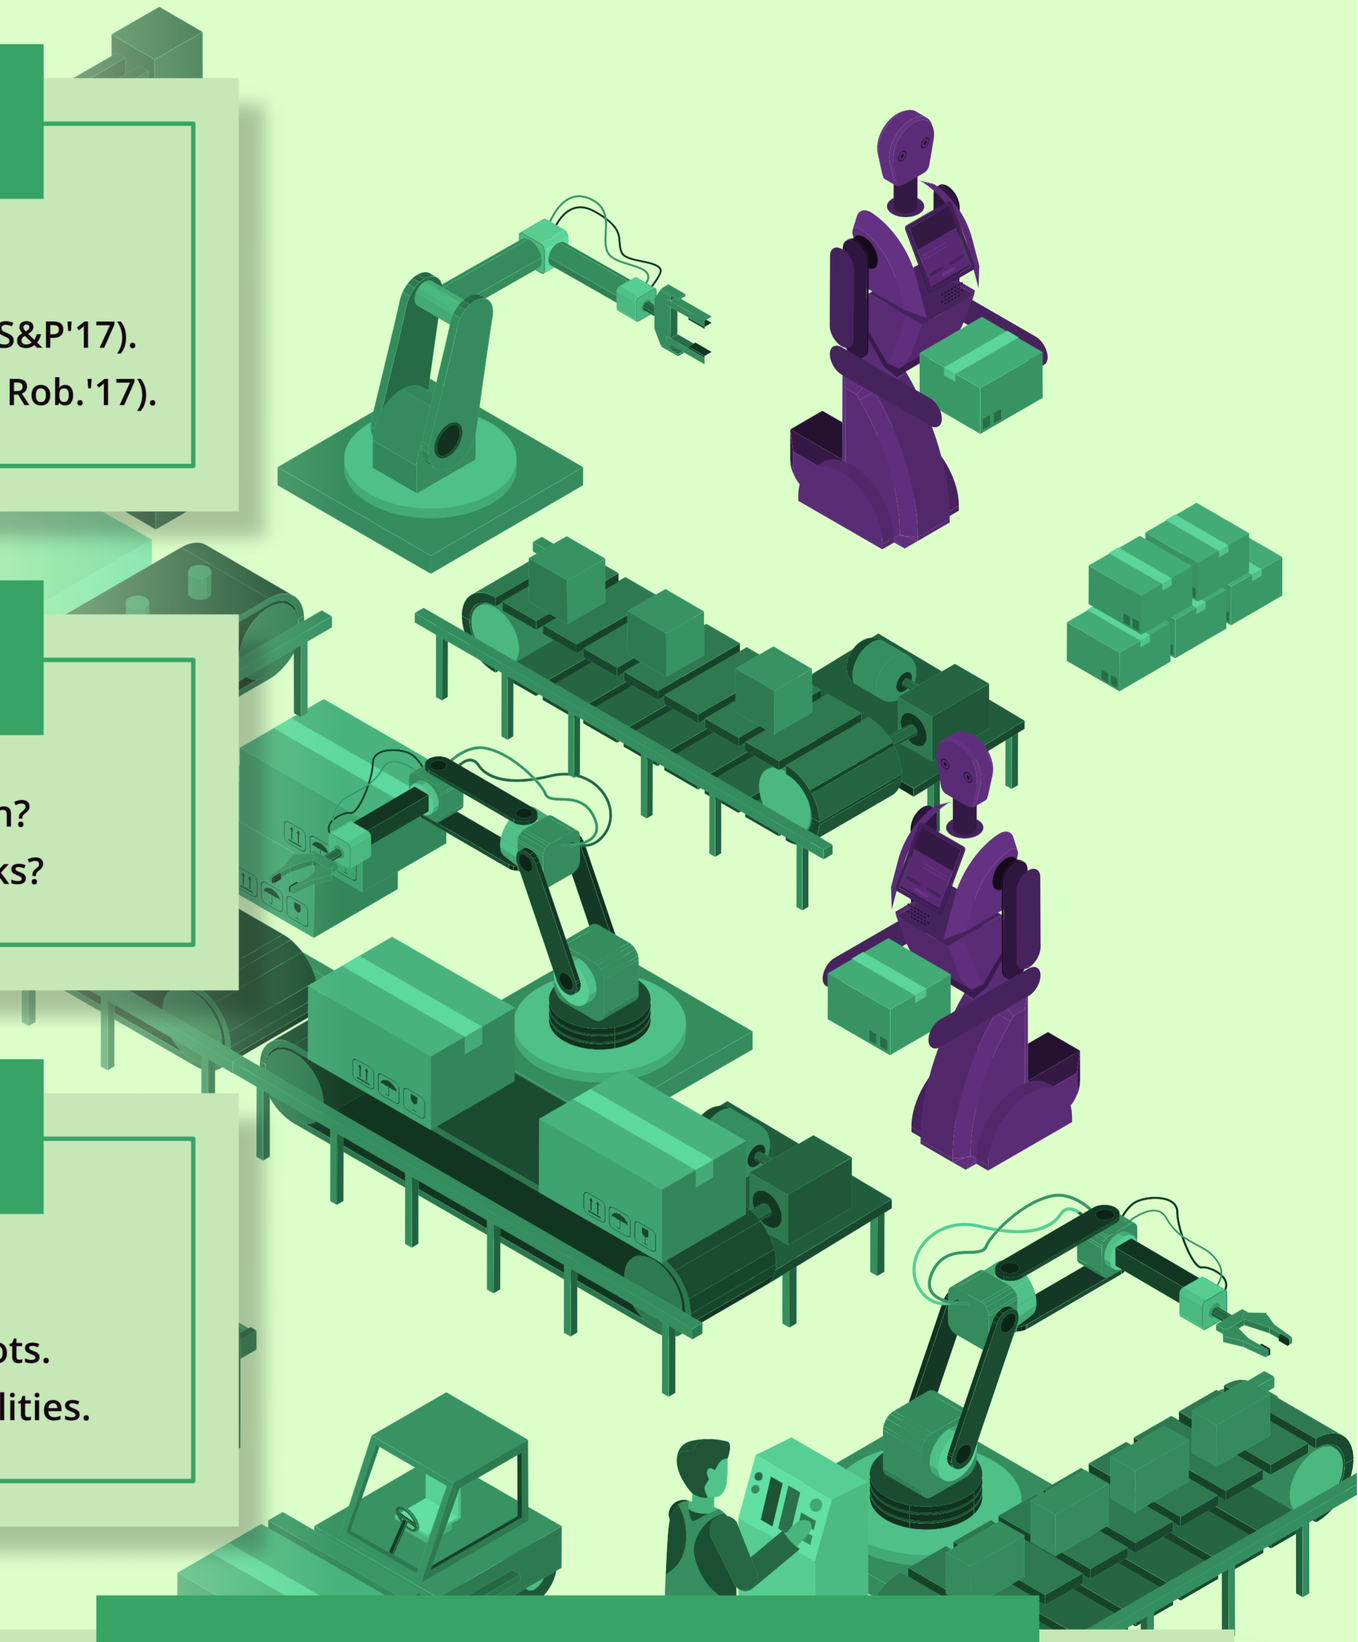

#4 Results

>90%

∃∀SMT

- Over 90% of conventionally-obtained MAPF solutions are vulnerable to physical masquerade attacks.
- We have obtained a complete and optimal defense algorithm that leverages inter-agent observations through a constraint-based encoding.

#5 Conclusions

- Malicious agents can reach secure locations undetected.
- Planning inter-agent observations limits the efficacy of PMA.

#6 Future Work

- Exploration of sub-optimal defense-oriented planning.
- Extending the current formulation to decentralized settings.

[1] D. Quarta, M. Pogliani, M. Polino, F. Maggi, A. M. Zanchettin and S. Zanero, "An Experimental Security Analysis of an Industrial Robot Controller," *IEEE Symposium on Security and Privacy (S&P)*, 268–286, 2017.
[2] S. Gil, S. Kumar, M. Mazumder, D. Katabi and D. Rus, "Guaranteeing spoof-resilient multi-robot networks," *Autonomous Robots*, 41(6), 1383–1400, Aug. 2017.
Illustration designed by macrovector / Freepik

BOSTON UNIVERSITY