POSTER: Can You Still See Me?: Identifying Robot Operations Over End-to-End Encrypted Channels

Ryan Shah, Chuadhry Mujeeb Ahmed, Shishir Nagaraja
University of Strathclyde, Glasgow
UK

ABSTRACT

Connected robots play a key role in automating industrial workflows. Robots can expose sensitive operational information to remote adversaries. Despite the use of end-to-end encryption, a passive adversary could fingerprint and reconstruct the entire workflows being carried out and developing a detailed understanding of how facilities operate. In this paper, we investigate whether a remote passive attacker can accurately fingerprint robot movements and reconstruct operational workflows. Using a neural network-based traffic analysis approach, we found that attackers can predict TLS-encrypted robot movements with around ~60% accuracy, increasing to near perfect accuracy in realistic settings. Ultimately, simply adopting best cybersecurity practices is not enough to stop even weak (passive) adversaries.

CCS CONCEPTS

• **Security and privacy → Systems security**; **Side-channel analysis and countermeasures**.

KEYWORDS

robotics, security, side channel, traffic analysis, SDN, neural network

ACM Reference Format:
Ryan Shah, Chuadhry Mujeeb Ahmed, Shishir Nagaraja. 2022. POSTER: Can You Still See Me?: Identifying Robot Operations Over End-to-End Encrypted Channels. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '22), May 16–19, 2022, San Antonio, TX, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3507657.3529659

1 INTRODUCTION

Prior to the realisation of Industry 4.0, many organisations relied on various rigid models for end-to-end supply chains to meet consumer demand. The nature of these workflows progressively required higher levels of responsiveness, flexibility and efficiency [10]. A shift to meet these new demands in order to remain relevant led to digitising supply chains, logistics and asset management, and ensuring end-to-end visibility. It is under this push into the realm of Industry 4.0 where connected robotics systems play a pivotal role. For example, in manufacturing facilities and product warehouses, robots are used to transform logistics management and meet increasing demands by working alongside existing human operators. This includes the likes of automated cross-docking and product stocking [17].

While previous robotic implementations made use of configured, pre-planned operations [2, 18], many industrial robotics systems use a teleoperated architecture [5, 13, 16, 19] (Figure 1). In such systems, a teach pendant (controller) is operated by a human which translates commands or movements into instructions the system can understand. A set of input devices (i.e. sensors, buttons on controller, etc.) and output devices (i.e. actuators) are linked together via a control system and a network in which the robot operates [14].

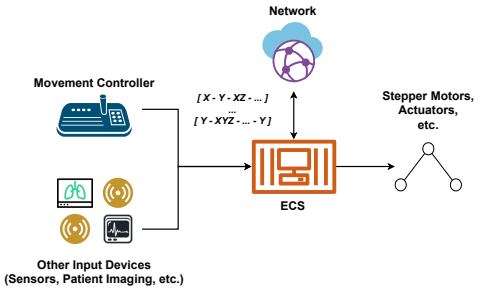

Figure 1: Typical Robot System Architecture

Prior art has demonstrated a number of different attacks on robotics systems, ranging from active targeted attacks such as modifying/dropping commands in-flight and modifying feedback to operators [1, 3, 6, 13]. However, there is little focus on reconnaissance aspects. Although attacks that fall under this umbrella, such as eavesdropping and device fingerprinting, are considered passive, the resulting compromise can still lead to severe consequences. In healthcare domains, for example, legislation mandates the use of TLS as a protective measure for confidential information [11, 12]. However, in industrial settings clear govern is lacking.

In this work, we investigate whether passive attackers can fingerprint industrial robot movements even when channel security measures are in place. Further, we also investigate whether this attack can also be used to accurately reconstruct operations typically carried out by industrial robots. By doing this, it could be possible to identify potentially confidential supply chain workflows which, for example, could be leaked to competitors. Employing a neural network to classify movements from collected traffic flows, we found that individual movements can be fingerprinted with at least 60% accuracy, increasing to near perfect accuracy under simulated network conditions. In terms of operation reconstruction, we observe average success rates also around 60%.

2 SYSTEM DESIGN AND TRAFFIC ANALYSIS

We made use of uFactory's uARM Swift Pro [20], which is operated by an Arduino Mega 2560 running MicroPython. The robot is connected to a teleoperated teach pendant (controller) running on a Windows 10 laptop using the uARM Python SDK. We route TLS-encrypted traffic through a software-defined network (SDN) using Mininet 2.3.0 for simulating realistic network conditions. Wireshark was used to capture traffic flows of movement operations along permutations of X, Y and Z movements for varying distances (1-50mm)

WiSec '22, May 16–19, 2022, San Antonio, TX, USA

ACM ISBN 978-1-4503-9216-7/22/05.
https://doi.org/10.1145/3507657.3529659

Movement	Precision	Recall
X	70%	85%
Y	69%	54%
Z	80%	63%
XY	21%	60%
XZ	68%	92%
YZ	81%	31%
XYZ	72%	97%

Table 1: Baseline Classification Results

and speeds (12.5-100mm/s). Further, we also captured traffic flows under varying network conditions (packet loss and link delay) to observe the impact of the attack in realistic settings. Within our dataset of around 150k traffic samples, we collected the prominent features for analysis, including: packet times, frame lengths, header lengths (IP/TCP/TLS), bytes-in-flight, and round trip times. Traffic samples were normalised using a min-max scaler and stratified to ensure even distribution of samples across each of the movements.

From a preliminary analysis, we found *simple traffic analysis* approaches (i.e. basic frequency analysis) do not provide enough interpretable information for accurate fingerprinting. Because of this, we employed a shallow neural network. This consisted of an input layer with one neuron for each of our collected features, one dense hidden layer of 108 neurons and an output layer with 7 neurons corresponding to our movement classes. The dense layer used the ReLU activation function and output layer used softmax with categorical cross-entropy as the loss function. For our dataset, we randomly split the traffic samples in a 60/20/20 split for training, testing and validation respectively.

We first experimented with a baseline set of samples (distance of 1mm and speed 12.5mm/s). As seen in Table 1, we observe an average of ~60% accuracy and found Y-based movements show lower success, possibly due to lower variance in traffic features. Next, we investigated the impact of our distance and speed parameters on classification. For distance, we found that an increase in distance leads to slightly lowered accuracy among most movements. This is potentially due to changes in the payload lowering variance among movement classes. For speed, we found that increasing movement speeds improves the classification accuracy, specifically at 50mm/s.

Given that our traffic flows are TLS encrypted, simulating realistic network conditions is crucial to our analysis for teleoperated environments using WAN networks [9, 14]. By introducing even a low link delay we see significant improvement in classification accuracy (near pefect). This may be due to increased variations in round-trip and packet interarrival times leading to higher accuracy, unlike distance and speed which primarily impact the traffic payload characteristics. As well as delay, we experimented with various packet losses. In realistic settings, failures or inefficiencies of network components can lead to packet loss. In comparison with delay, we also observe a significant increase in classification accuracy among all movements, potentially due to drops in packet arrivals resulting in increased interarrival times.

2.1 Operation Reconstruction

A further extension to our attack evaluation aimed at investigating whether higher-level operations (such as cross-docking and stocking) can be reconstructed from encrypted traffic flows. Specifically, we evaluated pick-and-place, push, pull and packing operations. For these operations, we took inspiration for movement trajectories from existing industrial robot datasets, such as the *Forward Dynamics Dataset Using KUKA LWR and Baxter* [**?**] for pick and place and the *Inverse Dynamics Dataset Using KUKA* [15] for push/pull. At the heart of these workflows is the actual dynamic movements themselves which may be aided by additional input (i.e. from sensors). Ultimately, given that movement patterns are the primary factor which establishes specific workflows, it is reasonable to conduct our experiment on reconstructing movements from traffic patterns solely using movement information. In this experiment, we observed that these operations can be reconstructed with at least an accuracy of around 60%. This is important as continuous monitoring of movement patterns can reveal potentially confidential workflows and could be given to competing facilities, for exampl. Further, this information can even combined with other side channels such as acoustic or EMF which may provide another level of information leakage (i.e. identifying the weight of products via EMF being cross-docked could leak information about contents).

3 COUNTERMEASURES

In this work, we examine Tor as a primary countermeasure to mitigate against our traffic analysis attack. Given its success as a defence in other areas (i.e. website fingerprinting), it is sensible to examine its efficacy in the face of our attack. We setup a Tor hidden service which receives control commands sent by our controller over HTTPS and monitored the incoming traffic on the hidden service host. The traffic was routed through multiple ASes over ~20 different circuits. In comparison with our baseline, we found that Tor leads to around a 20% decrease in accuracy across movements. In the context of operation reconstruction, we find significant decreases in recovery rate with drops of at least 30% present among the operations.

Examining the Tor traffic features shows that latency does not (overall) present a big problem for many cases. However, multiple robotics systems operating in a single space may leave such wait times undesirable at scale. Thus, other countermeasures can be considered as a point of future work. This includes padding robot traffic and mixing in background traffic. In the case of padding, techniques such as constant-rate and VIT-based approaches [4, 7] have shown success in various applications. These can be evaluated as countermeasures for future work to mitigate downsides (i.e. delays) associated with Tor. With regard to mixing, regularizing traffic (padding) may impose larger overheads, and thus a more light-weight approach which does not require additional infrastructure (i.e. adding dummy traces) [8] may prove more successful as a defence.

4 CONCLUSION

In conclusion, we present a case for evaluating whether a passive adversary can still identify robot movements, even when the traffic between a robot and controller (in a teleoperated architecture) is encrypted under TLS. We propose a shallow learning approach, which shows that it is possible for an adversary to successfully classify our robot's movements when protected by TLS with around 60% accuracy. Furthermore, we demonstrate that taking into account more fine-grained movement details such as distance of movement the accuracy increases, and when factors that impact network traffic, such as packet loss and link delay, are taken into account, we can achieve perfect accuracy (100%) for classifying our robot's movements.

REFERENCES

[1] Homa Alemzadeh, Daniel Chen, Xiao Li, Thenkurussi Kesavadas, Zbigniew T Kalbarczyk, and Ravishankar K Iyer. 2016. Targeted attacks on teleoperated

surgical robots: Dynamic model-based detection and mitigation. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 395–406.

[2] William L Bargar, André Bauer, and Martin Börner. 1998. Primary and revision total hip replacement using the Robodoc® system. *Clinical Orthopaedics and Related Research®* 354 (1998), 82–91.

[3] Tamara Bonaci, Jeffrey Herron, Tariq Yusuf, Junjie Yan, Tadayoshi Kohno, and Howard Jay Chizeck. 2015. To make a robot secure: An experimental analysis of cyber security threats against teleoperated surgical robots. *arXiv preprint arXiv:1504.04339* (2015).

[4] Jonas Bushart and Christian Rossow. 2020. Padding Ain't Enough: Assessing the Privacy Guarantees of Encrypted {DNS}. In *10th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 20)*.

[5] Mohsen Moradi Dalvand and Saeid Nahavandi. 2014. Improvements in teleoperation of industrial robots without low-level access. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2170–2175.

[6] Nicholas DeMarinis, Stefanie Tellex, Vasileios Kemerlis, George Konidaris, and Rodrigo Fonseca. 2018. Scanning the internet for ros: A view of security in robotics research. *arXiv preprint arXiv:1808.03322* (2018).

[7] Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. 2003. On countermeasures to traffic analysis attacks. In *IEEE Systems, Man and Cybernetics SocietyInformation Assurance Workshop, 2003.* IEEE, 188–195.

[8] Jiajun Gong and Tao Wang. 2020. Zero-delay lightweight defenses against website fingerprinting. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*. 717–734.

[9] Satoru Goto, Takuya Naka, Yoshitaka Matsuda, and Naruto Egashira. 2010. Teleoperation System of Robot arms combined with remote control and visual servo control. In *Proceedings of SICE Annual Conference 2010*. IEEE, 1975–1981.

[10] Heiner Lasi, Peter Fettke, Hans-Georg Kemper, Thomas Feld, and Michael Hoffmann. 2014. Industry 4.0. *Business & information systems engineering* 6, 4 (2014), 239–242.

[11] U.S. Department of Health and Human Services (HHS). 2013. Guidance to Render Unsecured Protected Health Information Unusable, Unreadable, or Indecipherable to Unauthorized Individuals. https://www.hhs.gov/hipaa/for-professionals/breach-notification/guidance/index.html.

[12] THE EUROPEAN PARLIAMENT and THE COUNCIL OF THE EUROPEAN UNION. 2017. "Medical Device Regulations – Official Journal of the European Union. https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R0745.

[13] Davide Quarta, Marcello Pogliani, Mario Polino, Federico Maggi, Andrea Maria Zanchettin, and Stefano Zanero. 2017. An experimental security analysis of an industrial robot controller. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 268–286.

[14] Reza Rahimi, C Shao, Malathi Veeraraghavan, Andrea Fumagalli, Jorge Nicho, J Meyer, S Edwards, C Flannigan, and P Evans. 2017. An industrial robotics application with cloud computing and high-speed networking. In *2017 First IEEE International Conference on Robotic Computing (IRC)*. IEEE, 44–51.

[15] Elmar Rueckert, Moritz Nakatenus, Samuele Tosatto, and Jan Peters. 2017. Learning inverse dynamics models in o (n) time with lstm networks. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, 811–816.

[16] Timo Salmi, Jari M Ahola, Tapio Heikkilä, Pekka Kilpeläinen, and Timo Malm. 2018. Human-Robot Collaboration and Sensor-Based Robots in Industrial Applications and Construction. In *Robotic Building*. Springer, 25–52.

[17] SAP. 2021. Learn How Industry 4.0 And Robots Strengthen Warehouse Logistics. https://www.forbes.com/sites/sap/2021/11/22/learn-how-industry-40-and-robots-strengthen-warehouse-logistics.

[18] Arndt P Schulz, Klaus Seide, Christian Queitsch, Andrea Von Haugwitz, Jan Meiners, Benjamin Kienast, Mohamad Tarabolsi, Michael Kammal, and Christian Jürgens. 2007. Results of total hip replacement using the Robodoc surgical assistant system: clinical outcome and evaluation of complications for 97 procedures. *The International Journal of Medical Robotics and Computer Assisted Surgery* 3, 4 (2007), 301–306.

[19] Ashutosh Tewari, James Peabody, Richard Sarle, Guruswami Balakrishnan, Ashok Hemal, Alok Shrivastava, and Mani Menon. 2002. Technique of da Vinci robot-assisted anatomic radical prostatectomy. *Urology* 60, 4 (2002), 569–572.

[20] uFactory. [n. d.]. UFACTORY – uARM. https://www.ufactory.cc/pages/uarm. Last Accessed: 03/12/2020.