

Behind Closed Doors

Process-Level Rootkit Attacks in Cyber-Physical Microgrid Systems

Rath, Suman ; Zografopoulos, Ioannis ; Vergara , Pedro P.; Nikolaidis, Vassilis C. ; Konstantinou, Charalambos

DOI

[10.1109/PESGM48719.2022.9916907](https://doi.org/10.1109/PESGM48719.2022.9916907)

Publication date

2022

Document Version

Final published version

Published in

Proceedings of the 2022 IEEE Power & Energy Society General Meeting (PESGM)

Citation (APA)

Rath, S., Zografopoulos, I., Vergara , P. P., Nikolaidis, V. C., & Konstantinou, C. (2022). Behind Closed Doors: Process-Level Rootkit Attacks in Cyber-Physical Microgrid Systems. In *Proceedings of the 2022 IEEE Power & Energy Society General Meeting (PESGM)* (pp. 1-10). IEEE.
<https://doi.org/10.1109/PESGM48719.2022.9916907>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Behind Closed Doors: Process-Level Rootkit Attacks in Cyber-Physical Microgrid Systems

Suman Rath*, Ioannis Zografopoulos[†], Pedro P. Vergara[‡],
Vassilis C. Nikolaidis[§], Charalambos Konstantinou[†]

*School of Electrical & Computer Engineering, Oklahoma State University

[†]CEMSE Division, King Abdullah University of Science and Technology (KAUST)

[‡]Intelligent Electrical Power Grids, Delft University of Technology (TU Delft)

[§]Dept. of Electrical and Computer Engineering, Democritus University of Thrace

E-mail: suman.rath@okstate.edu, p.p.vergarabarrios@tudelft.nl, vnikolai@ee.duth.gr

{ioannis.zografopoulos, charalambos.konstantinou}@kaust.edu.sa

Abstract—Embedded controllers, sensors, actuators, advanced metering infrastructure, etc. are cornerstone components of cyber-physical energy systems such as microgrids (MGs). Harnessing their monitoring and control functionalities, sophisticated schemes enhancing MG stability can be deployed. However, the deployment of ‘smart’ assets increases the threat surface. Power systems possess mechanisms capable of detecting abnormal operations. Furthermore, the lack of sophistication in attack strategies can render them detectable since they blindly violate power system semantics. On the other hand, the recent increase of process-aware rootkits that can attain persistence and compromise operations in undetectable ways requires special attention. In this work, we investigate the steps followed by stealthy rootkits at the process level of control systems pre- and post-compromise. We investigate the rootkits’ precompromise stage involving the deployment to multiple system locations and aggregation of system-specific information to build a neural network-based virtual data-driven model (VDDM) of the system. Then, during the weaponization phase, we demonstrate how the VDDM measurement predictions are paramount, first to orchestrate crippling attacks from multiple system standpoints, maximizing the impact, and second, impede detection blinding system operator situational awareness.

Index Terms—Rootkit, cyber-physical microgrid, intelligent malware, data-driven prediction, virtual twin.

I. INTRODUCTION

Microgrids (MGs), among others, foster the inclusion of renewable energy sources. Closely coupled cyber and physical layers guarantee the increased flexibility and robust operation of MGs. The information and communication cyber layer receives inputs from sensors and issues controls in the physical layer. As a result, adversaries can exploit cyber vulnerabilities (e.g., insecure protocols, software bugs, etc.) to port their attacks impacting MG’s control and stability [1]. Attacks able to manipulate sensor data can have significant impacts (e.g., blackouts, human safety, equipment damage, etc.) [2].

Existing literature focuses on various cyberattacks, their exploitation techniques, and schemes to detect and mitigate them [3]. The severity and risks associated with malware infections within industrial systems, especially rootkits, has captured the attention of researchers [4]. Rootkits represent a class of malware which can intelligently hide their presence inside their targets [5]. They can eavesdrop system data and allow attackers to collect real-time system information via

remotely accessible connections [6]. Adversaries can exploit these connections to issue malicious commands, manipulate the infected device(s), and stealthily control their operation.

MG-based process level rootkit attacks should be considered when designing and implementing detection and mitigation strategies. Rootkit attacks should represent a crucial part of reliability assessment procedures for certification laboratories, tasked with identifying and eliminating vulnerabilities in embedded control devices and designing experimental testbeds for evaluation of cyber-physical power systems [4], [7]. This paper extends our previous work-in-progress on stealthy rootkit attacks [8], and highlights a potential rootkit attack path after the installation of the malware at multiple locations within the MG. After its deployment, the rootkit aggregates system measurements to build an accurate system replica allowing the estimation of the MG states trajectories; in our case, using a neural network-based approach. We demonstrate different approaches that rootkits could utilize to conceal their presence (during the system information aggregation phase) and disguise the attack impact overcoming existing security fortifications exploiting system state estimations. The paper presents simulation results to demonstrate the attack methodology and achievable system impact.

II. ATTACK MODEL

Rootkits can leverage the vulnerabilities (cyber layer) of operating system architectures encountered in power system workstations, HMIs, etc. during their deployment. Persistence is then achieved by masking their presence through system information and log modifications [9]. In most cases, the compromised systems exhibit vulnerabilities similar to consumer computers, and rootkits can exploit them to deploy on different levels of the hardware or software stack, e.g., firmware, bootloader, memory, kernel -based rootkits, etc. In Fig. 1, we present a MG model infected with the rootkit. Our threat model assumes that the rootkit can access and stealthily modify process level sensor measurements and controller strategies at different locations within the MG, i.e., a strong adversary model [3]. The rootkit achieves persistence and disguises its operation, by collecting sufficient system data to anticipate the MG state trajectories. By reporting the expected

the secondary controller introduces $\delta\omega$ and δv terms in the power controller equations:

$$\omega_i^* = \omega_n - D_{P_i} P_i + \delta\omega \quad (5)$$

$$v_i^* = v_n - D_{Q_i} Q_i + \delta v \quad (6)$$

In Eqs. (5) and (6), $\delta\omega$ and δv negate the droop caused as a result of the primary control. They are defined as follows [1]:

$$\begin{aligned} \delta\omega = K_1 & \left(\sum_{j \in N(i)} a_{ij}(\omega_j - \omega_i) + g_i(\omega_n - \omega_i) \right. \\ & \left. + \sum_{j \in N(i)} a_{ij}(D_{P_j} P_j - D_{P_i} P_i) \right) \end{aligned} \quad (7)$$

$$\delta v = K_2 \left(\sum_{j \in N(i)} a_{ij}(D_{Q_j} Q_j - D_{Q_i} Q_i) \right) \quad (8)$$

where a_{ij} represents the element in the adjacency matrix of the bidirectional connected communication graph, g represents the pinning gain, and K_1, K_2 are constants [1]. The control targets for the secondary controller are shown below:

$$\lim_{t \rightarrow \infty} \|\omega_i(t) - \omega_n\| = 0 \quad \forall i \quad (9)$$

$$\lim_{t \rightarrow \infty} \|D_{P_i} P_i - D_{P_j} P_j\| = 0 \quad \forall i, j \quad (10)$$

$$\lim_{t \rightarrow \infty} \|D_{Q_i} Q_i - D_{Q_j} Q_j\| = 0 \quad \forall i, j \quad (11)$$

The rootkit's long-term target is to disrupt the aforementioned control objectives by introducing small manipulations whose consequences are unnoticeable by bad data detectors (BDDs). The rootkit methodology and the steps required for the attack are discussed in the following sections.

A. VDDM: Target Identification and State Prediction

In recent days, virtual power plants (VPPs) have been widely used for the modeling of demand-response management schemes, participation in day-ahead markets, etc. [14], [15]. Open source data enable adversaries to access and build such VPPs for malicious gains [16]. In our use case, the adversary uses the VDDM of the MG for two purposes: (i) performance evaluation of an attack vector before its introduction in the physical system, and (ii) prediction of system behaviour under normal conditions to hide the rootkit's manipulations from the system, e.g., BDDs.

Before initiating an attack, the adversary attempts to infect the maximum possible agents at the sensor, communication, and controller levels to maximize possible data collection. However, in a practical scenario, the attacker's ability to install malware at different stages may be limited by physical, hardware, or software constraints [17], [18]. In that case, the attacker can construct a virtual replica of the system with limited state information. Thus, the virtual system may not be a precise replica of the actual system and its accuracy will be determined by the quality of the data captured and the adversary's knowledge about the system [3]. The replica, in our work, uses a combination of Kalman filtering (KF) and neural networks to predict the future states of the MG. Specifically, KF is used to train an artificial neural network

(ANN) that acts as the state estimator. It improves the quality and size of the training dataset by estimating missing values of data [19]. KF can also generate a sufficiently large dataset from a relatively smaller set of accurate data points, eliminating any potential uncertainty associated with the rootkit's data collection abilities (which may affect training of the VDDM). The hybrid estimator – combining KF and ANN – improves prediction accuracy and convergence speed when compared with traditional KF/NN approaches [20], [21]. The steps involved in the generation of the training and test datasets for the ANN are described below.

Since the attacker has access to sensors and controllers, real-time measurements from a set of infected physical sensors ζ can be captured. The state of the system is estimated from the captured measurements using KF. Let the vector containing the measurements (obtained from ζ) be m . The system state is:

$$x(t) = F(t)x(t-1) + B(t)c(t) + n(t) \quad (12)$$

where $x(t)$, $F(t)$, $B(t)$, $c(t)$, $n(t)$ represent the state vector, the state transition matrix, the control input matrix, the control input vector, and the unknown process noise vector corresponding to $x(t)$, respectively. $n(t)$ is given by a zero mean normal distribution provided by the covariance matrix $C_0(t)$, $E[n(t)n(t)'] = C_0(t)$. The equation depicting the measurement can be written as follows.

$$m(t) = H(t)x(t) + v(t) \quad (13)$$

where $H(t)$ and $v(t)$ represent the transformation matrix mapping states to measurements and the known process noise vector, respectively. $v(t)$ is obtained from the covariance matrix $C_1(t)$, $E[v(t)v(t)'] = C_1(t)$. The state in the next time step can be estimated as follows:

$$x(t+1|t) = F(t)x(t|t) + B(t)c(t) \quad (14)$$

$$m(t+1|t) = H(t)x(t+1|t) \quad (15)$$

$$n(t+1) = m(t+1) - m(t+1|t) \quad (16)$$

Further,

$$x(t+1|t+1) = x(t+1|t) + N(t+1)n(t+1) \quad (17)$$

where $N(t+1)$ represents the Kalman gain utilized for state covariance. The state covariance can be estimated as:

$$P(t+1|t) = F(t)P(t|t)F(t)' + C_0(t) \quad (18)$$

$$S(t+1) = H(t+1)P(t+1|t)H(t+1)' + C_1(t+1) \quad (19)$$

$$N(t+1) = P(t+1)H(t+1)'S(t+1)^{-1} \quad (20)$$

$$P(t+1|t+1) = P(t+1|t) - N(t+1)S(t+1)N(t+1)' \quad (21)$$

where $P(t+1|t)$, $S(t+1)$, $N(t+1)$ and $P(t+1|t+1)$ are the state prediction covariance, measurement covariance, filter gain, and updated state covariance, respectively.

The measurements collected by the rootkit, are used to analyze the current state and estimate the future states of the system. Furthermore, the obtained data is used to perform




Fig. 2. Mean squared error (MSE) during training, validation and testing of the artificial neural network (ANN).

supervised training of an ANN with a l -layer feed-forward architecture that can predict the future state X_P of the system once it receives the measurement vector m and the requisite estimation time T as inputs. Let the initial set of weights (defined during the training phase) be W_0 . The initial output X_{P0} can be defined as:

$$X_{P0} = f_A(W_0, l, m, T) \quad (22)$$

As the training progresses, the weights are updated through back-propagation (post-completion of each iteration) to minimize the error in the outputs per the predefined loss function.

$$W_{i+1} = f_B(\eta, e, W_i, l) \quad (23)$$

where η is the learning rate, and e represents the error between output of the network (after the i_{th} iteration) and the estimated output from the training data.

The predicted state (output) after the i_{th} iteration can be defined as a function mapping the obtained measurements (input) to the future state, given a set of weights W_i .

$$X_{Pi} = f_A(W_i, l, m, T) \quad (24)$$

After convergence, the network can be deployed for state estimation in the malicious system. Before the attack vector α introduction, the attacker tries to estimate the future system state $X_{P\alpha, T_\alpha}$ using the malicious VDDM.

$$X_{P\alpha, T_\alpha} = f_A(W_f, l, m + \alpha, T_\alpha) \quad (25)$$

where T_α represents the end of the attack period. The magnitude of α must satisfy the following constraint: $|\alpha| \leq |\alpha_{max}|$, where $|\alpha_{max}|$ represents the maximum allowable magnitude of the attack vector to evade BDDs. During the weaponization phase, the attacker estimates and reports the expected normal system state to conceal the rootkit actions.

$$X_{N\alpha, T_\alpha} = f_A(W_f, l, m, T_\alpha) \quad (26)$$

where $X_{N\alpha, T_\alpha}$ is the predicted normal trajectory of the system in the absence of the attack vector α .

IV. RESULTS AND DISCUSSION

An AC MG model is designed using MATLAB following the architecture of the model presented in [1]. The model consists of four DGs. Details of various system parameters can be obtained from [1]. Data collected from this test model is used to generate a VDDM, which can predict the future states of the system using the KF-ANN hybrid state estimator model discussed in Section III. The ANN is trained through




Fig. 3. Rootkit attack impact on the DG frequency.

supervised learning using labelled datasets generated from the designed MG model and the KF algorithm. Fig. 2 shows the accuracy of the VDDM by depicting the mean squared error (MSE) during the training, validation and testing phases.

The MG model is also used as a test system to demonstrate the possible impact that a rootkit malware with partial/complete access can cause to disrupt nominal behavior. Three target objectives are assigned to the malicious rootkit: (i) frequency manipulation, (ii) voltage manipulation, and (iii) disruption of load sharing. For these objectives, the rootkit uses its controller- and sensor-level access at the PCC side to first create a false fault alarm, cause defensive islanding of the system, and then compromise the DG operation.

A. Target: Frequency Manipulation

When the adversary instructs the rootkit to initiate the frequency manipulation attack, the malware evaluates its access level and uses the malicious VDDM to identify the vulnerable agents, which can be exploited to achieve the desired level of deviation. In this case, the adversarial objective entails achieving a higher steady state frequency. To accomplish this, the rootkit manipulates load devices and sensors. It modifies ω_n affecting $\delta\omega$ as per Eq. (7) and further influences ω_i^* as per Eq. (5). Fig. 3 depicts the normal and manipulated system trajectories after forced islanding. The rootkit can successfully distort the system frequency trajectory through the manipulation of load sensors and the local controller associated with the leader, i.e., DG 1. Fig. 3 shows that the frequency attains steady state at 50.05 Hz, however, the deviation could be further increased depending on the introduced attack vector.

B. Target: Voltage Manipulation

In the voltage manipulation case, the rootkit malware can either modify the voltage levels at any particular bus or create voltage instability affecting the whole MG system. After the target assignment, the malware uses the VDDM to determine the manipulation strategy required to achieve the set objective. In our case, the rootkit introduces false reactive power demand by thoroughly crafting bias injections to the system sensors. Fig. 4 and Fig. 5 show the normal voltage trajectory and the voltage trajectory after the rootkit manipulations are initiated. The experimental results demonstrate that the rootkit attack can cause the gradual increase of the three-phase voltage magnitude at DG 1.




Fig. 4. Three-phase voltage magnitude at DG 1 after islanding occurs.




Fig. 5. Abnormal voltage increase at DG 1 after the rootkit attack.

C. Target: Disturbance in Load Sharing

In the attack-free scenario, all DGs in the simulated system share the total load demand equally (load sharing pattern shown in Fig. 6). However, during the attack phase, the rootkit disrupts the existing load sharing pattern. After the vulnerable agent identification, the malware establishes a time frame ($t = 5s$) within which manipulations will be introduced. For instance, an attack could be initiated during peak demand when the DGs operate close to their generation limits. Fig. 7 shows the disruption in load sharing caused by the rootkit through controller-level manipulation. Continued manipulations destabilize the system (after $t = 5s$) creating instabilities.

V. CONCLUSIONS AND FUTURE WORK

The impact of process-aware rootkit attacks represents a severe threat for power systems. The presented rootkit can learn the MG's operational patterns. We demonstrate how such information can help mount crippling attacks against the local DGs and the PCC level. Rootkits can achieve persistence by masking their presence. Our future work will focus on developing detection and mitigation mechanisms, e.g., moving target defences, stable kernel representations [22], etc. thwarting the adversarial learning behavior of rootkits.

REFERENCES

- [1] S. Rath *et al.*, "A cyber-secure distributed control architecture for autonomous ac microgrid," *IEEE Systems Journal*, vol. 15, no. 3, 2021.
- [2] C. Konstantinou *et al.*, "Gps spoofing effect on phase angle monitoring and control in a real-time digital simulator-based hardware-in-the-loop environment," *IET Cyber-Physical Systems: Theory & Applications*, vol. 2, no. 4, pp. 180–187, 2017.
- [3] I. Zografopoulos *et al.*, "Cyber-physical energy systems security: Threat modeling, risk assessment, resources, metrics, and case studies," *IEEE Access*, vol. 9, pp. 29 775–29 818, 2021.
- [4] J. Powers *et al.*, "Whitelist malware defense for embedded control system devices," in *Saudi Arabia Smart Grid (SASG)*. IEEE, 2015.
- [5] P. Krishnamurthy *et al.*, "Stealthy rootkits in smart grid controllers," in *International Conference on Computer Design (ICCD)*, 2019, pp. 20–28.




Fig. 6. Load sharing among DGs during nominal operation.




Fig. 7. Impact of rootkit manipulations on the load sharing scheme.

- [6] C. Xenofontos *et al.*, "Consumer, commercial and industrial iot(in) security: attack taxonomy and case studies," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 199–221, 2022.
- [7] J. Reeves *et al.*, "Intrusion detection for resource-constrained embedded control systems in the power grid," *International Journal of Critical Infrastructure Protection*, vol. 5, no. 2, pp. 74–83, 2012.
- [8] S. Rath, I. Zografopoulos, and C. Konstantinou, "Stealthy rootkit attacks on cyber-physical microgrids: Poster," in *Proceedings of the 12th ACM Int'l Conference on Future Energy Systems*. ACM, 2021, p. 294–295.
- [9] MITRE. Rootkit. [Online]. Available: <https://tinyurl.com/2p83p89n>
- [10] A. P. Kuruvela *et al.*, "Hardware-assisted detection of firmware attacks in inverter-based cyberphysical microgrids," *International Journal of Electrical Power & Energy Systems*, vol. 132, p. 107150, 2021.
- [11] M. Leng, S. Sahoo, and F. Blaabjerg, "Stability investigation of dc microgrids under stealth cyber attacks," in *2021 IEEE Energy Conversion Congress and Exposition (ECCE)*, 2021, pp. 1427–1432.
- [12] J. Ospina *et al.*, "On the feasibility of load-changing attacks in power systems during the covid-19 pandemic," *IEEE Access*, vol. 9, 2021.
- [13] A. Bidram *et al.*, *Cooperative synchronization in distributed microgrid control*. Springer, 2017.
- [14] A. Mnatsakanyan and S. W. Kennedy, "A novel demand response model with an application for a virtual power plant," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 230–237, 2014.
- [15] A. Baringo, L. Baringo, and J. M. Arroyo, "Day-ahead self-scheduling of a virtual power plant in energy and reserve electricity markets under uncertainty," *IEEE Transactions on Power Systems*, vol. 34, no. 3, 2018.
- [16] A. Keliris *et al.*, "Open source intelligence for energy sector cyberattacks," in *Critical infrastructure security and resilience*. Springer, 2019.
- [17] X. Wang *et al.*, "Malicious firmware detection with hardware performance counters," *IEEE Transactions on Multi-Scale Computing Systems*, vol. 2, no. 3, pp. 160–173, 2016.
- [18] O. M. Anubi and C. Konstantinou, "Enhanced resilient state estimation using data-driven auxiliary models," *IEEE Transactions on Industrial Informatics*, vol. 16, no. 1, pp. 639–647, 2019.
- [19] R. Habtom and L. Litz, "Estimation of unmeasured inputs using recurrent neural networks and the extended kalman filter," in *Proceedings of International Conference on Neural Networks*, vol. 4. IEEE, 1997.
- [20] P. M. Sieberg *et al.*, "Hybrid state estimation-a contribution towards reliability enhancement of artificial neural network estimators," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [21] R. Zhan and J. Wan, "Neural network-aided adaptive unscented kalman filter for nonlinear state estimation," *IEEE Signal Processing Letters*, vol. 13, no. 7, pp. 445–448, 2006.
- [22] I. Zografopoulos and C. Konstantinou, "Detection of malicious attacks in autonomous cyber-physical inverter-based microgrids," *IEEE Transactions on Industrial Informatics*, 2021.