SAND2018-13805C

Sandia National Laboratories

A Multi-Agent System for Cyber-Physical Network Security

Introduction

Industrial Control Systems/Internet of Things/Operational Technology (I/OT) networks and their variants all share similarities in how their devices and networks bridge the cyber-physical domain. Not only is the primitive functionality shared, but also the lack of conventional cybersecurity techniques. These systems are often accessed remotely by a variety of entities including utility workers, multiple third-party vendors, consumers, brokers, and other machines, where vetting and control of access may be cumbersome or impossible based on the equipment used. As a consequence, these Internet-connected devices that control and monitor physical processes are at risk of disruption by cyber-initiated attacks, and may provide additional paths through which attacks may be carried out. Enterprise and cloud networks may enjoy the availability of resources to support cyber-hardening, -visibility, and -response; the constrained resources of I/OT networks do not readily accommodate upgrades, replacements or bolt-on solutions for cyber security. Furthermore, given the safety requirements of some systems, downtime to implement changes may not be acceptable. Thus, as adversaries have begun to recognize the minimal workfactor required to attack these networks, have cyber practitioners now begun to observe the effects from lackluster security. From smarthomes, to medical devices, to national powergrids, the attack space has seen signal emulation (man-in-the-middle), sensor influence/hijacking, eavesdropping, malware/ransomware, denial-of-service, device destruction. Historically, methods to secure I/OT networks have pointed to solutions that cannot be reasonably implemented due to legacy equipment, vendor complicity, or cost. A new approach is needed that can address these issues, but still be flexible to grow with new cybersecurity techniques and the advancement of network infrastructure. A successful security paradigm in enterprise/cloud networks is largely based on agents. The solution described herein borrows from this mindset and acts in the research space of Multi-Agent Systems (MAS).

MAS and AA

We apply the notion of MAS/agents in I/OT networks through security and mission decoupling. For those networks whose devices are fixed, Autonomous Agents (AA) will not require installation on the endpoints, but may integrated into the I/OT network-space (tied to a shared fieldbus, wireless network, or via bump-in-the-wire). For those that are software-based, an AA may be installed in user-space. AAs may provide: (1) Passive listening/active probing (where applicable); (2) Data/metadata collection; (3) Behavioral analysis and majority voting schemes; (4) AA self-policing; (5) Active defense techniques; (6) Security policy enforcement. The goal being that the AA/MAS shall not affect the realtime communication requirements of the system (ICS), nor the functionality of the devices (IoT). Leveraging multiple sources from academia and industry, we devised a generalized architecture of typical cyber-physical enterprise networks that can address the constructs of legacy networks, is malleable to fit IoT network, and extendable to fit other network types (enterprise, cloud, mobile, tactical), with the option to adapt to future operations and management models as needed. The mapping of the different network models to our generalized architecture is show in the table below.

Network Model	B-Tier	C-Tier	D-Tier	P-Tier
ICS	Business, Logistics (Level 4)	Control software, HMI, Operations (Level 2/3)	Remote devices, collection (Level 1/2)	Physical domain, devices (Level 0/1)
IoT	Business, User Access (Actions)	Storage, Processing, Reporting, Cloud (Insights)	Gateways, Hub	Things, sensors
IIoT	Business Integration	Information, operations, applications	Control	Proximity and physical systems
Enterprise/Cloud	Data center, edge, cloud	Core routing, boundary	Concentrator, distribution	Access, mobile, endpoints

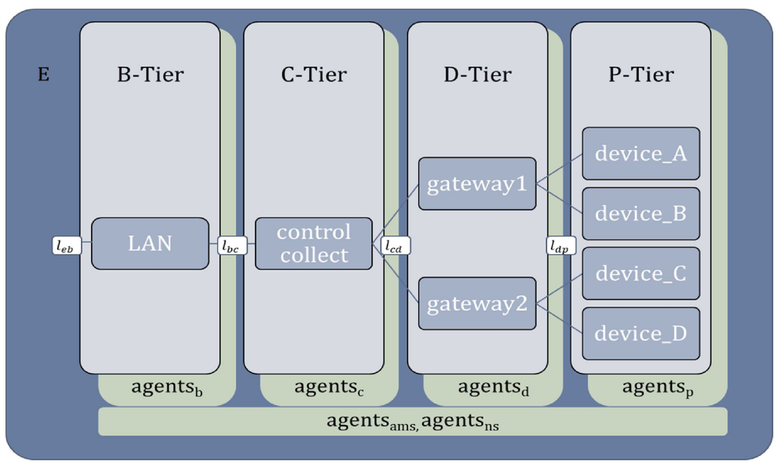

$S_n = \langle E|B|C|D|P \rangle$

$L = \langle l_{eb}|l_{bc}|l_{cd}|l_{dp} \rangle$

$S_a = \langle a_{\{b,c,d,p\}}|ns_j|ams_k \rangle$

$a_i = \langle \kappa|\sigma|\alpha|\rho|kb \rangle$
$\kappa = developer\ or\ AMS\ knowledge$
$\sigma = sensors$
$\alpha = actuators$
$\rho = percept$
$kb = knowledge\ base$

$e = event\ or\ observation$

$kb = knowledge\ base$

$kb = \bigcup_i^n \rho_i \cup \kappa$

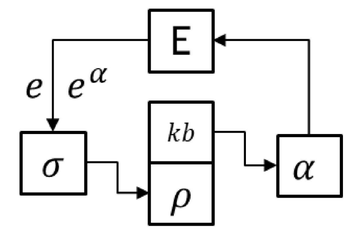

Simple Reflex	:	R, r
Model-based Reflex	:	$M(E)$
Goal-based Agent	:	G, g
Utility-based Agent	:	U

$G = immutable\ goals, based\ on\ \kappa$

$g = mutable\ goals\ to\ obtain\ G$

$R = immutable\ rules\ to\ obtain\ G, g$

$r = mutable\ rules\ to\ obtain\ G, g$

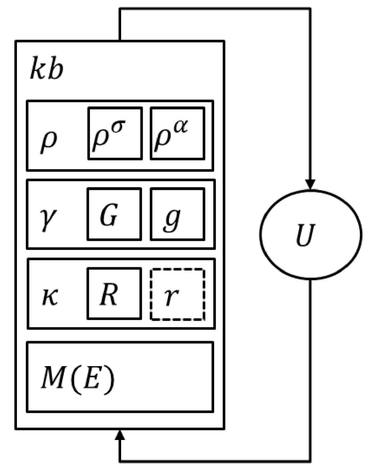

Specifying the MAS

The mission architecture S_n describes the network where data or commands are pushed-to or pulled-from a Control tier *C*, to a Distribution tier *D*, and finally to a Physical tier *P*; business-oriented operations *B* connect to *C*. External entities E exist outside of the purview of S_n, and are adjacent to *B*. The underlying communication planes, or links *L*, are described broadly by the connections between the tiers. In the *C* tier, operational devices collect data from or control devices in the *P* tier. The *D* tier provides the aggregation, normalization, filtering or summarization of data from devices in the *P* tier, any may also send command and control signals to the *P* tier endpoints. Devices in the *P* tier interact directly with their environment, either sensing or performing physical actions. AAs interface directly with the *P* tier (on the broadcast medium, last physical hop or as an embedded agent), and the *C, D, B* tiers, to form a cogent MAS Agent Platform (AP) S_a, as a self-policing out-of-band network. Agents in the *B* tier may collect data or serve as a Directory Facilitator (DF) nameserver ns_j for the S_a. Overall MAS control, developer/maintainer interface, and data push-pull mechanisms are handled by the Agent Management System (AMS), *ams* (where more than one, *k* may exist in the environment). The distributed agents in *C*, *D* and *P* (and *B* as required) collect data, correlate as needed, learn, and send summary or raw data to a AMS for additional reasoning, or may interface with an OT safety system to provide alerting for anomalous observations between the *P*, *D* and/or *C* tiers.

On-going Efforts

Agents communicate via a Message Transport Service (MTS), that may be supported by S_n through L, or an alternate transport network. We are leveraging FIPA to produce an Agent Communication Language (ACL) for message context description; a Communicative Act (CA) for communicating functions or action; and Semantic Language (SL) to define semantics for a CA as a logic of attitudes and actions. Using OSBrain, rapid-prototype development is underway (Python3) using a simplified communication architecture (ACL mapping): PUSH-PULL, REQUEST-REPLY, PUBLISH-SUBSCRIBE. The Agent base class been built, AMS base class is underway, as we continue researching model and learning techniques to build into an ICS Emulation Platform.

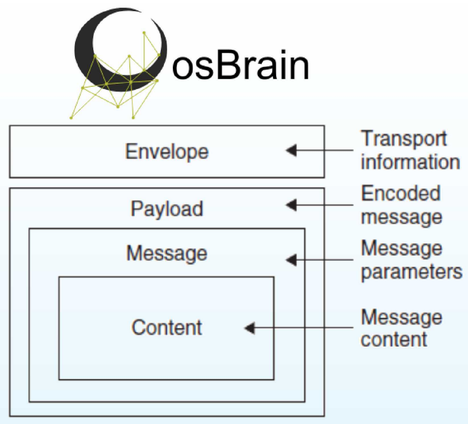

U.S. DEPARTMENT OF ENERGY
NNSA National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. _____

William M.S. Stout

Sandia National Laboratories