

Managing and Securing Critical Infrastructure - A Semantic Policy and Trust driven approach*

Wenjia Li: Department of Computer Sciences, College of Information
Technology, Georgia Southern University, Statesboro, GA 30460
Palanivel Kodeswaran: IBM Research - India, Bangalore, India
Pramod Jagtap: Amazon.com Inc, Seattle, WA
Anupam Joshi and Tim Finin: Department of Computer Science and
Electrical Engineering, University of Maryland, Baltimore County,
Baltimore, MD 21250

August 2011

Abstract

Cyber physical systems (CPS) and cyber infrastructure are a key elements of the national infrastructure, and securing them is of vital importance to national security. There is ample evidence that these systems are vulnerable to disruption and damage due to natural disasters social crises, and terrorism. CPS applications are becoming more widespread, ranging from healthcare patient monitoring systems to autonomous vehicles to integrated electrical power grids. often the new application domains cross administrative boundaries and are not under the supervisory control of a single domain. This introduces critical issues of policy and trust that have not been traditionally addressed in their design and management. Most work in securing CPS and cyber infrastructure has focused on security of the communication links between the sensing and actuating elements. We describe a more holistic approach that is based on the concepts of *situation awareness* for monitoring the state of a CPS system and *high-level policies* to manage their functioning and security. Such a framework can manage the trust relationship among entities as well as external contextual information when detecting, evaluating and responding to threats. We illustrate the framework by showing how it can protect the traditional Internet backbone by automatically configuring BGP router systems, defending against attacks and recovering from accidental or malicious damage. We also illustrate how the same framework can be used to secure devices and information in mobile networks.

Key Words: Cyber-Physical System, Critical Infrastructure, security, trust, policy, semantic web

1 Introduction

As the world has become more developed, industrialized and globalized, its reliance on critical physical and cyber infrastructure has increased. This infrastructure includes many systems such as electrical power generation and distribution, roads, bridges and tunnels that make up our ground transportation system, airports and air traffic control supporting airline transportation, communication networks, both wired and wireless, systems for storing and distributing water and food supplies, medical and healthcare delivery systems, and financial, banking and commercial transaction assets.

These systems are vulnerable to disruption and damage due to natural disasters such as earthquakes or hurricanes, social crises like wars and riots, and terrorism that deliberately targets infrastructure to injure, disrupt and frighten citizens. Our economy, public safety, and national security rely on our ability

*This work was supported in part by an STTR grant from the Defense Advanced Research Agency (W31P4Q-06-C-0395), the Air Force Office of Scientific Research under MURI award FA9550-08-1-0265 and the NSF under Grant Number 0910838. This work was done when the first three authors were with the Department of Computer Science and Electrical Engineering, University of Maryland, Baltimore County

to monitor and protect these systems and to quickly remediate and repair any damage that might be done to them.

Critical infrastructure, as its name suggests, generally refers to a network of assets that are important for the proper functioning of a society and economy [1]. Cyber-Physical System (CPS) include those systems that support critical infrastructure and are characterized by a tight coupling between the computational and physical components of the system [2]. CPS have a variety of emerging application domains, ranging from smart grid to patient monitoring to intelligent transportation system.

Securing CPS and critical infrastructure is of great importance to national security. Ample evidence has been found in recent years of serious vulnerabilities of these systems, and many of them have even been discussed in the popular media [3, 4]. Most of current methods to address these vulnerabilities either aim to secure the communication links between the sensing and actuating elements using encryption, or focus on end device access control. However, CPS and critical infrastructure systems have a number of characteristics that require more than these traditional, communication-oriented approaches. In particular,

- Data are heterogeneous.
- Data come from a variety of autonomous sensors.
- Physical effects using actuators are involved.
- Control is split amongst autonomous systems.

Due to these unique features of CPS and critical infrastructure, the current point security solutions are not able to completely meet their security needs. We argue that a more holistic approach that is -driven and aware of various contexts is essential to secure the emerging cyber-physical systems. We next present some scenarios that illustrate the problems above, and show why we need a holistic approach to secure CPS and critical infrastructure is demonstrated by the following examples.

1.1 Security for Smart Grid

One major security challenge for CPS is the heterogeneity of data. Since data in CPS generally come from multiple types of reporting devices (such as water meter, gas meter, electric meter, and weather sensor, etc.), the format and scale of these data are different. Moreover, these various types of reporting devices are normally under different administrative domains. Figure 1 illustrate an example scenario in Advanced Metering Infrastructure (AMI).

We find from Figure 1 that there are four types of reporting devices in this system: water meter, gas meter, , electric meter, and weather sensor. These reporting devices are likely managed by different utility companies, and they generally sense and report data that are in distinct formats. Therefore, it is not feasible to simply apply the traditional security mechanisms, such as the outlier detection technique [5, 6], to process these data in one dimension. For example, we observe from Figure 1 that the gas meter for house C reports an abnormally low gas usage, which is only 30% when compared to the average daily gas usage in winter. On the other hand, the report from the weather sensor for house C shows that the weather condition is extremely bad there (heavy snow, very chilly and very windy). In this case, we want to find out why the gas meter has reported such a low reading for house C .

Therefore, we can define some policy rules to help better understand the context in which data are obtained. In this example, assuming we know that the house has gas heating, we conclude that either the gas meter for house C reported an incorrect reading, or the gas supply for house C encountered technical malfunction. Either could represent a deliberate attack (on the meter or the supply). However, data from other appliance level sensors, or the electric meter sensor, might also make us conclude that the owner is away, and so lower than normal consumption is normal behavior in this context. Similar cross utility scenarios can be envisaged for NASPINET [7] the interconnected network of synchrophasors which can be used for power distribution optimization.

1.2 Policy-based Framework for BGP Routing

In the large-scale and complex distributed infrastructure systems, it is essential to ensure that all the heterogeneous entities behave appropriately. However, users, services and access rights may change frequently in these distributed systems. Therefore, policy based security is likely be the most effective

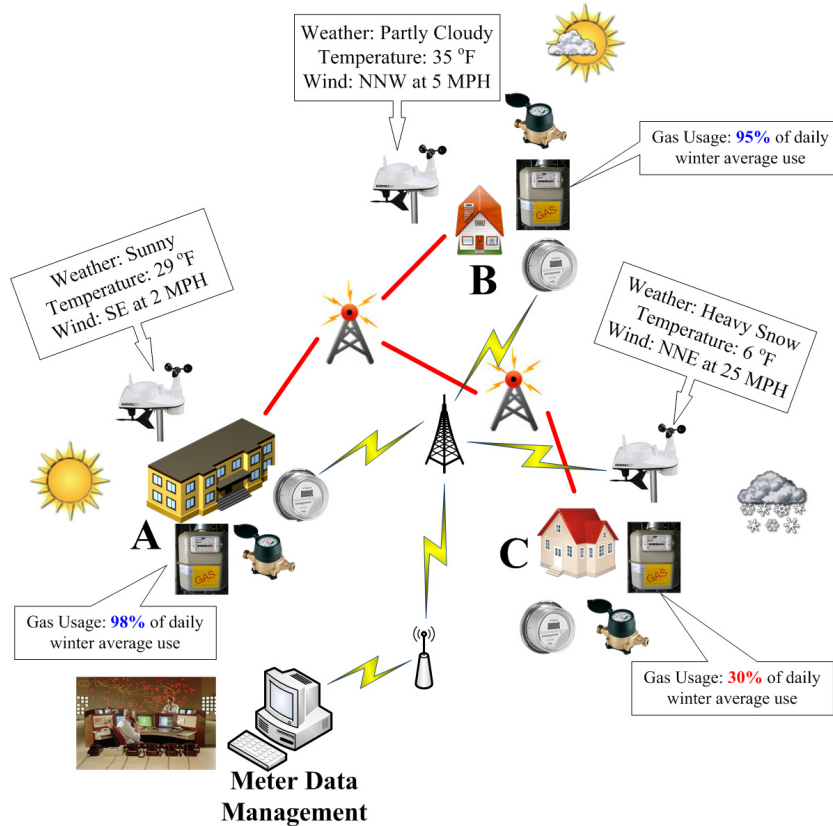


Figure 1: Different types of reporting devices in advanced metering infrastructure

mechanism for distributed systems, because it is possible to specify how different entities act without modifying their internal mechanisms [8].

The Internet is just such a large-scale and complex distributed system. The Border Gateway Protocol (BGP) is the de facto routing protocol used in the Internet today for advertising network reachability. Unlike other routing protocols, BGP is a policy based routing protocol that allows operators to control which routes are chosen in the routing protocol. This flexibility provides operators with the ability to tweak BGP to enforce the high level goals of their organizations. However, configuring BGP routers correctly to enforce organizational goals is a formidable task. The lack of a high level language for modeling and enforcing network wide routing policies forces operators to manually configure BGP routers at the lowest level detail. The resulting configurations have no usable semantics associated with them, and consequently cannot be verified for correctness. Furthermore, the configurations do not always reflect the organization's high level goals. Most configuration files run into hundreds of lines, further making debugging harder.

This problem gets much more severe in the case of military networks where there is a lack of skilled network operators on the field. Furthermore, given the dynamic nature of the environment, organization's routing goals change quickly over time, requiring a corresponding rapid change in the BGP configurations. Expecting operators on the field to manually configure the routers correctly in a short span of time seems unrealistic for medium to large sized networks.

There is a growing need for a high-level language to model and configure BGP routing policies. The goal of such a high level language is to allow operators focus on the policy decisions rather than on the low-level implementation details. For example, operators would like to automatically configure routers to implement existing trust relationships between Autonomous Systems (AS) by merely stating the relationship type such as customer-provider between them without having to manually construct the associated export and import filters. These high-level languages are particularly important in dynamic environments where operators neither have the time nor the skills to manually configure routers.




Figure 2: True Alarm VS False Alarm in Traffic Monitoring System

1.3 Intelligent Transportation System

Transportation is a critical infrastructure, and Intelligent Transport Systems are a key remedy being proposed for sprawl and crawl related problems in large urban environments. Consider a traffic monitoring system, which is depicted in Figure 2. Present generation monitoring systems are based on ground sensors and cameras. However, with increasing computing and communication capabilities embedded in vehicles, their onboard sensors themselves can be used to monitor traffic. From Figure 2(a), we find that a vehicle observes an accident ahead, and it reports this accident to the system. Therefore, the traffic alarm shown in Figure 2(a) is true. In contrast, Figure 2(b) shows two conflicting traffic alarms. Given that there is no accident in this scenario, the vehicle that reports accident to the system is either *faulty* or *malicious*. In this case, the system needs to rely on more data, such as reports from other vehicles or other types of sensors to decide which alarm is trustworthy.

In addition to the reports from vehicles, we may obtain data from other types of sensors, such as surveillance cameras and floating cars [9]. As we discuss in Section 1.1, policies need to be applied in this scenario in order to better represent the context in which the reports are generated.

1.4 Assured Information Sharing

A key element of our critical infrastructure are the systems used for information dissemination and sharing, whether in financial systems such as banks or national security systems such as those that connect the various intelligence and armed forces entities. It is essential to take both security and privacy into account when various kinds of information are shared in CPS and critical infrastructure.

In order to fight the war against non state extremist actors, the DoD, federal agencies, coalition partners and first responders, among others have to proactively share information and make effective decisions. An example scenario pertains to the Distributed Common Ground System (DCGS) being developed by the US DoD. It will ensure the horizontal integration of joint intelligence, surveillance and reconnaissance (ISR) sensor platforms for improving time critical targeting. While the Air Force is developing DCGS (with Raytheon Corporation as the prime contractor), the Navy is developing its version called DCGS-N and the Army is developing its version called DCGS-A. The three organizations must share information for combat operations via DCGS as well as with foreign intelligence services [10].

Yet in doing so, one must protect the confidentiality of sensitive information and appropriately respect the privacy of individuals. Traditional security policies are often based on the concept of need to know and are typified by predefined and often rigid specifications of which principals and roles are pre-authorized to access what information.

Policies determine what and how much to reveal in the discovery stage for both information seekers and providers, and can drive the process of negotiation in the acquisition and release stage. Fine-grained policy integration algorithms are needed to support dynamic coalitions and virtual organizations that need to rapidly share and integrate information. The policy framework will need to be flexible enough to include several classes of policies such as those for confidentiality and privacy, accountability, trust, identity management, multilevel security and compliance among others.

1.5 Situation Awareness

Situation Awareness (SA), as is implied by its name, is “the perception of environmental elements within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future” [11].

One important example scenario for situation awareness is intrusion detection [12]. Traditionally, intrusion detection is a point solution, based on signatures. The “fusion” of multiple sources of attack information into an operational picture is still dependant on the analyst. Also, the characterization and classification of computer attacks and other intrusive behaviors have been limited to taxonomies. Taxonomies, however, lack the necessary and essential constructs needed to enable an intrusion detection system (IDS) to reason over an instance that is representative of the domain of a computer attack. Ontologies provide powerful constructs that include machine interpretable definitions of the concepts within a domain and the relations between them. Ontologies, therefore, provide software systems with the ability to share a common understanding of the information at issue, in turn empowering the software system with a greater ability to reason over and analyze this information [13]. So for instance, such ontologies can reason over “instances” provided signature based point IDS tools, and fuse them together with background knowledge about vulnerabilities in a system and likely attack vectors to create a higher level picture for an analyst rather than a series of lower level alarms. Similar scenarios that deal with sensed data being fused to create an operational picture can be sketched for other domains, such as medical informatics[14].

These examples illustrate the key challenges for securing critical infrastructure that we’ve articulated earlier – data are heterogeneous, and are sensed by autonomous sensors. Any “reaction” or response also has to be coordinated across heterogeneous domains. So situation aware, policy driven, distributed security solutions are more suited to such systems than point solutions that focus on encryption or access control alone. In the sections that follow, we show such systems in two different contexts – that of a cyberphysical systems and internet routing.

2 Related Work

2.1 Security and Trust Management for Cyber-Physical System

The research on Cyber-Physical System attracts increasing attention in recent years. Some research efforts have been made to cope with security threats to CPS. We briefly summarize these research efforts as follows.

One major category of security solutions for CPS is the authentication and encryption techniques. In [15], Rogers et al. proposed an authenticated control framework for distributed voltage support on the smart grid. In this framework, various authentication techniques, such as digital signature [16] and HMAC [17], are used to secure the control of end-user reactive-power-capable devices to mitigate low voltage problems at the transmission system level. Metke et al. [18] discusses key security technologies for a smart grid system, including public key infrastructures and trusted computing.

Another security solution for CPS is the access control method. In [19], a role-based access control (RBAC) system is proposed for the distributed resources in a cyber-physical system. The RBAC system uses Shibboleth [20], which is an attribute authorization service currently being used in Grids.

In a latest research work, Tang et al. [21] presents a method called *Tru-Alarm*, which finds out trustworthy alarms and consequently increases the feasibility of CPS. *Tru-Alarm* first estimates the locations of objects that cause alarms, and then it constructs an object-alarm graph. Finally, trustworthiness inferences are performed based on linked information in the object-alarm graph. To the best of our knowledge, *Tru-Alarm* is the most similar work to our CARE-CPS scheme. However, this method does not take the heterogeneity of data sources into consideration, and all the sensor data (such as sensor data from the *battle-network* system) are processed using the same data processing algorithm. On the contrary, our propose CARE-CPS scheme uses policies to specify trust evaluation in different contexts. In this way, different types of sensor data can be better understood and utilized in our scheme.

In our previous work [5, 6, 22, 23, 24], we have made some efforts to identify abnormal node behaviors and manage the trustworthiness of nodes in Mobile Ad-hoc Networks (MANETs). Because CPS

also partially relies on wireless links to exchange data, and also sensor data are noisy in CPS, we believe that the mechanisms for securing MANETs can also be used to secure CPS with some adjustments.

2.2 Policies for Security in Distributed Systems

According to Sloman, policies define a relationship between subjects and targets [25]. Policy-based security is often used in systems where flexibility is required as users, services and access rights change frequently, such as wireless networks and other large-scale distributed systems. In these distributed systems, it is essential to ensure that all the heterogeneous entities behave appropriately. Therefore, policy based security should be the most effective mechanism for distributed systems, because it is possible to specify how different entities act without modifying their internal mechanisms [8].

Multiple policy languages have been studied in the past decade, such as Extensible Access Control Markup Language (XACML) [26] and the *Rei* policy language [8]. XACML [26] is a language in XML for expressing access policies. XACML allows control over actions and supports resolution of conflicts. On the other hand, *Rei* is a policy language designed for pervasive computing applications that is based on deontic concepts and grounded in a semantic language.

We have also made some efforts to utilize policies in malicious peer detection for Mobile Ad-hoc Networks [22]. Given that wireless networks are widely used in CPS, this method may also help secure CPS.

3 A Policy and Trust Framework to Secure CPS

In this section, we depict the policy and trust driven framework to secure CPS and Critical Infrastructure in details. The goal of the framework is to properly assess the trustworthiness of each reporting device (meter, sensor, measurement unit, etc) in different contexts using policies.

3.1 Scheme Overview

In the policy and trust driven framework, there are three major functional units, namely Data Collection, Policy Management, and Trust Management. Figure 3 illustrates the framework.

The Data Collection unit is responsible for collecting and sending data to either the Policy Management unit or the Trust Management Unit. Sensor data that are relevant to various contextual information are sent to Policy Management unit, such as temperature, weather condition, altitude, geolocation, wireless signal strength, speed, etc. On the other hand, meter readings, such as readings from water meter, gas meter or electric meter are sent to the Trust Management unit so that the trustworthiness of these reporting devices can be determined based on these readings.

3.2 Policy Management

In the Policy Management unit, all the contextual information will be used in policies. For example, if a smart meter is found to report abnormal readings, then the contextual information is used in this case to determine whether these abnormal readings are possibly caused by environmental factors or not. Table 1 describes various contextual information collected by sensors and sent to Data Collection unit, such as the current weather conditions, geolocation, temperature, and signal strength. The contextual information is then reported to the Policy Management unit. Then Policy Management unit analyses the reported contextual information and uses policies to determine whether the meter is intentionally reporting fake readings or the current environmental conditions cause those faulty meter readings.

The system can have multiple policies to consider the effects of various environmental factors. For instance, policies can be declared as (i) *If surrounding temperature is beyond range 0F-120F then there is a possibility of faulty reading*, (ii) *If the signal strength is weak then there is a possibility of faulty reading*, (iii) *If the current weather conditions are either of heavy raining, snowing or foggy then there is a possibility of faulty reading* and (iv) *If the altitude is higher than 2000 feet, weather conditions are snowing and temperature is below 32F then there is a possibility of faulty reading*. These policies are represented in Jena's rules syntax specification in Table 2, Table 3, Table 4 and Table 5.

Table 1: Dynamic environmental data received from sensors. It consists of temperature, weather conditions, location details - latitude, longitude, altitude and signal strength.

```

CPS:Sensor_Device a CPS:Sensor ;
  CPS:has_sensor_id "1" ;
  CPS:has_sensor_type "X" ;
  CPS:has_sensed_information CPS:Sensed_Data.
CPS:Sensed_Data a CPS:Sensor_Information ;
  CPS:has_temperature "20F" ;
  CPS:has_signal_strength "medium" ;
  CPS:has_location_information CPS:Location_Data ;
  CPS:has_weather_information CPS:Weather_Data .
CPS:Location_Data a CPS:Location_Information ;
  CPS:has_latitude "39.253525";
  CPS:has_longitude "-76.710706";
  CPS:has_altitude "456".
CPS:Weather_Data a CPS:Weather_Information ;
  CPS:has_weather_condition "4" ;

```

Table 2: Policy to report the possibility of faulty readings if surrounding temperature is beyond range 0F-120F.

```

[TemperatureRule:
  (?sensorDevice a CPS:Sensor)
  (?sensorDevice CPS:has_sensed_information ?sensedData)
  (?sensedData CPS:has_temperature ?temperature)
  lessThan(?temperature, 0) greaterThan(?temperature, 120)
  ->
  (?sensorDevice CPS:faulty_device "true") [
]

```

Table 3: Policy to report the possibility of faulty readings if signal strength is weak.

```

[SignalStrengthRule:
  (?sensorDevice a CPS:Sensor)
  (?sensorDevice CPS:has_sensed_information ?sensedData)
  (?sensedData CPS:has_signal_strength ?signal_strength)
  equal(?signal_strength, CPS:Signal_Strength_Weak)
  ->
  (?sensorDevice CPS:faulty_device "true")
]

```

Table 4: Policy to report the possibility of faulty readings if current weather conditions are either of Heavy raining, Snow or Foggy.

```

#For convinience, conditions are mapped to numerical values as
#Clear = 1, Sunny = 2, Heavy raining = 3, Heavy snow = 4, Foggy = 5
[WeatherConditionsRule:
  (?sensorDevice a CPS:Sensor)
  (?sensorDevice CPS:has_sensed_information ?sensedData)
  (?sensedData CPS:has_weather_information ?weatherData)
  (?weatherData a CPS:Weather_Information)
  (?weatherData CPS:has_weather_condition ?weatherCondition)
  greaterThan(?weatherData, 2)
->
  (?sensorDevice CPS:faulty_device "true")
]

```

Table 5: Policy to report the possibility of faulty readings incase of higher altitude.

```

#Sensor device can report faulty readings if altitude is greater than
#2000 ft,weather conditions are snowing and temperature is below 32F.
[AltitudeRule:
  (?sensorDevice a CPS:Sensor)
  (?sensorDevice CPS:has_sensed_information ?sensedData)
  (?sensedData CPS:has_weather_information ?weatherData)
  (?sensedData CPS:has_location_information ?locationData)
  (?weatherData CPS:has_weather_condition ?weatherCondition)
  (?sensedData CPS:has_altitude ?altitude)
  (?sensedData CPS:has_temperature ?temperature)
  equal(?weatherData, 4) lessThan(?temperature, 32)
  greaterThan(?altitude, 2000)
->
  (?sensorDevice CPS:faulty_device "true")
]

```

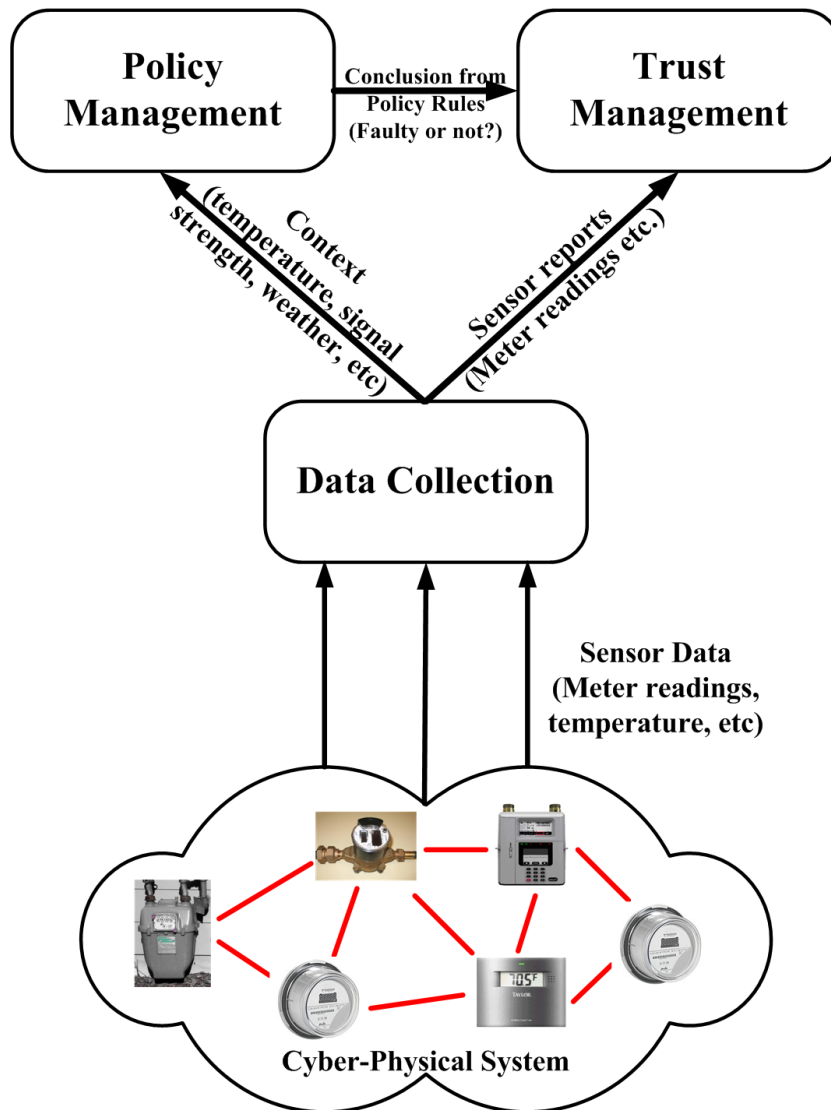



Figure 3: Schematic Diagram of the Policy and Trust Driven Framework

Based on the example policy rules shown in these tables, the Policy Management unit can determine whether the abnormal readings are caused by environmental factors, or they are deliberately sent by compromised devices. Then, this conclusion is used by the Trust Management unit to evaluate the trustworthiness of the reporting devices (meters, sensors, measurement units, etc.).

3.3 Trust Management

In CPS, not only do we care about how trustworthy each reporting device is, but we also need to find out whether the reported data are true or not. For example, in intelligent transportation system we need to know whether or not there is any accident on that road. If so, it is necessary that more actions should be taken, such as sending out an ambulance, issuing traffic alerts, and redirecting incoming traffic, etc. In addition, if a smart meter reports abnormal readings, then the energy company needs to know if the abnormal readings are true or not, so that it can properly respond to the incident. Therefore, in addition to the well-studied *Device Trust*, we introduce two new types of trust, namely *Report Trust* and *Event Trust*, to indicate the trustworthiness of the reported events themselves. Figure 4 shows how *Device trust* and *Report trust* are evaluated.

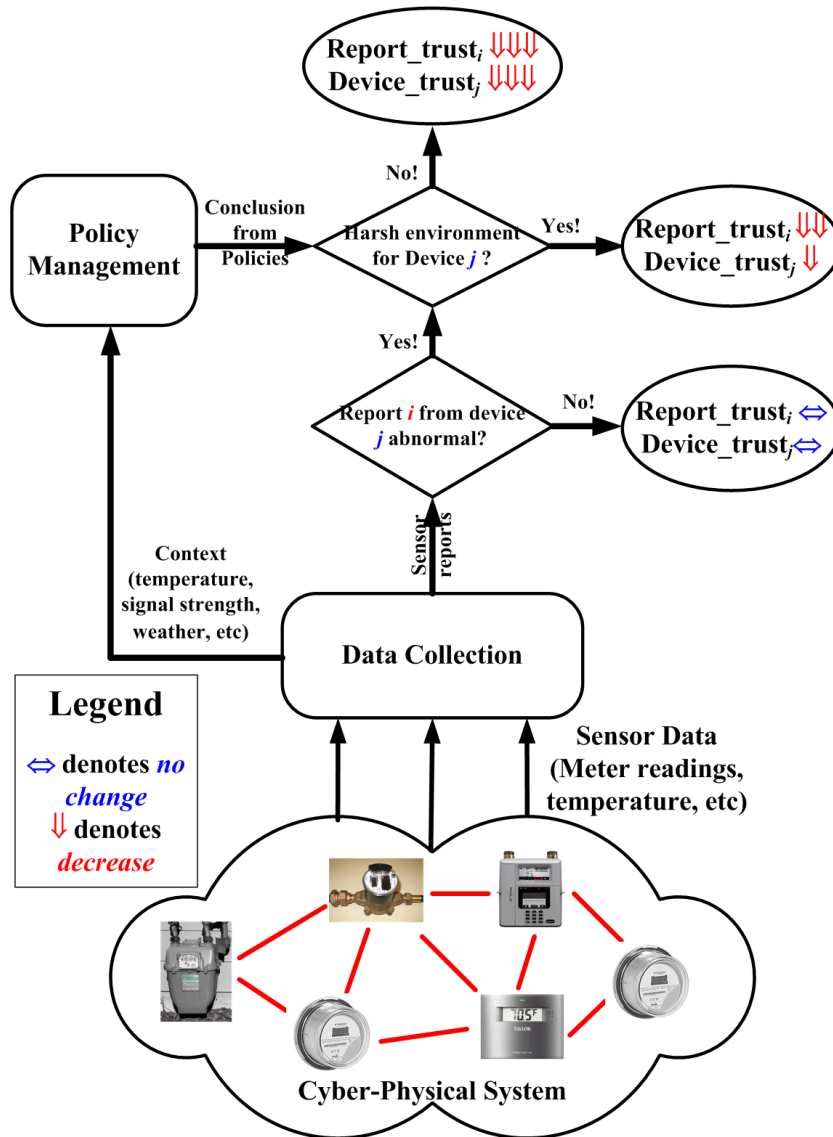


Figure 4: Device Trust VS. Report Trust

Figure 4 shows that Device Trust and Report Trust are evaluated based on both the correctness of the reports and the environmental factors. For the evaluation of *Report Trust*, we first check if there is any abnormal report, which is a report that significantly deviates from other reports regarding the same event. If not, then the report is trustworthy because the reports are consistent from all relevant devices. However, if there is any abnormal report, then the next step is to identify if there is any environmental factor that causes the abnormal report. For example, if the signal strength for the wireless link is weak, then it is possible that the report has been changed or damaged during transmission. In this case, the report trust is decreased because the report may be changed even if it is caused by environmental factors. On the other hand, if the abnormal reports are NOT caused by environmental factors, then it is very likely that the corresponding device has been compromised and controlled by an adversary. In this case, the report trust is significantly decreased because it has been intentionally tampered by the adversary.

Similarly, we can evaluate *Device Trust* based on the results of these two judgements. As for *Event Trust*, because multiple sensors can report different values for the same event ("Accident Ahead!" VS. "No Accident!" for the same portion of a road, etc.), we can decide the value of Event Trust based on the evaluation results for all Report Trust values that are associated with this event. If the majority of reports

Table 6: Simulation Parameters

<i>Parameter</i>	<i>Value</i>
Simulation area	600m × 600m
Num. of nodes	50, 100, 200
Transmission range	120m
Node placement	<i>Random</i>
Num. of bad nodes	5, 10, 20
Simulation time	900s

support this event, then it is trustworthy; otherwise, it is not so trustworthy.

4 Prototype Implementations

In this section, present two example scenarios to demonstrate that our proposed framework will work well for critical infrastructure system protection.

4.1 Security and Trust Management for Wireless Networks

The first scenario is use of the framework to secure devices and information in wireless networks that are components of Cyber-Physical Systems.

First, we obtain some simulation results to evaluate the performance of the policy-driven and context-aware framework. In addition, we declare and execute some example policies on mobile platforms such as Android phones. In this way, we know that the framework can be deployed on real devices.

4.1.1 Simulation Results and Analysis

We use GloMoSim 2.03 [27] as the simulation platform, and table 6 lists the parameters used in the simulation scenarios. Note that the simple trust evaluation method without policy management (such as trust management scheme discussed in [5, 6]) acts as the Baseline method when we evaluate the performance of CARE-CPS.

We use the following two parameters to evaluate the accuracy of our CARE-CPS scheme: Precision (P) and Recall (R). These two parameters are defined as follows.

$$P = \frac{\text{Num of Truly Malicious Devices Caught}}{\text{Total Num of Untrustworthy Devices Caught}}$$

$$R = \frac{\text{Num of Truly Malicious Devices Caught}}{\text{Total Num of Truly Malicious Devices}}$$

Each simulation scenario has 20 runs with distinct random seeds, which ensures a unique initial node placement for each run. Each experimental result is the average over the 20 runs for this simulation scenario. The simulation results are shown in Figure 5 and Figure 6.

We find from Figure 5(a) that CARE-CPS always achieves a higher precision score than the Baseline method when node density varies. Moreover, when the device density is higher, both methods yield a better precision. This is the case because it is more likely to receive true reports from sensors when there are a higher number of well-behaved sensors.

Similarly, Figure 5(b) shows that CARE-CPS outperforms the Baseline method in terms of recall. Also, the recall value is higher when the device density is higher.

Figure 6(a) and Figure 6(b) depict the precision and recall values for CARE-CPS and the Baseline method. We find that both the precision and recall values decrease when there are a higher percentage of malicious devices, which is pretty obvious. In addition, CARE-CPS is able to produce a better performance than the Baseline method in terms of both precision and recall values.

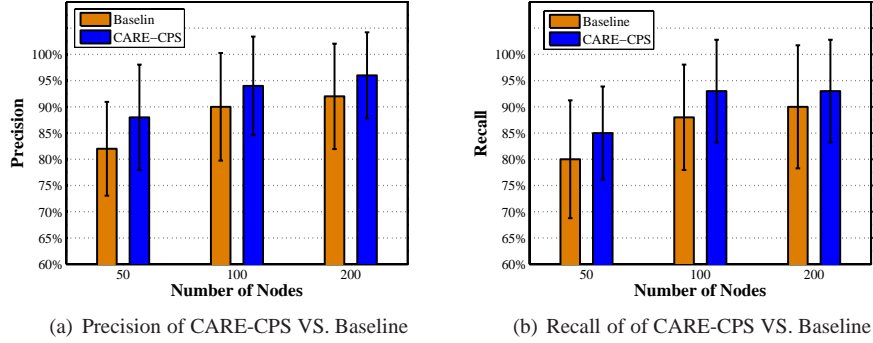


Figure 5: Effect of Device Density on CARE-CPS and Baseline

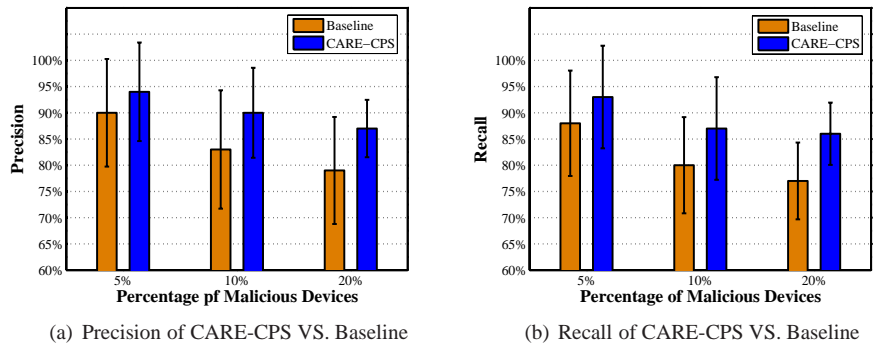


Figure 6: Effect of Adversary Percentage on CARE-CPS and Baseline

4.1.2 Experimental Results on Android Phones

In addition to the simulation, we build an android application which treat smartphones as components of a Cyber Physical System. We use the device capabilities to collect sensor data and to perform reasoning over sensed data and contextual information using Jena. The experimental results are displayed in the following Figure 7 and Figure 8.

Figure 7(a) shows the environmental factors for an abnormal sensor report. According to the policy rule, this abnormal report is caused by the out-of-bound temperature. Therefore, we conclude that the abnormal report is caused by the environmental conditions, which is shown in Figure 7(b).

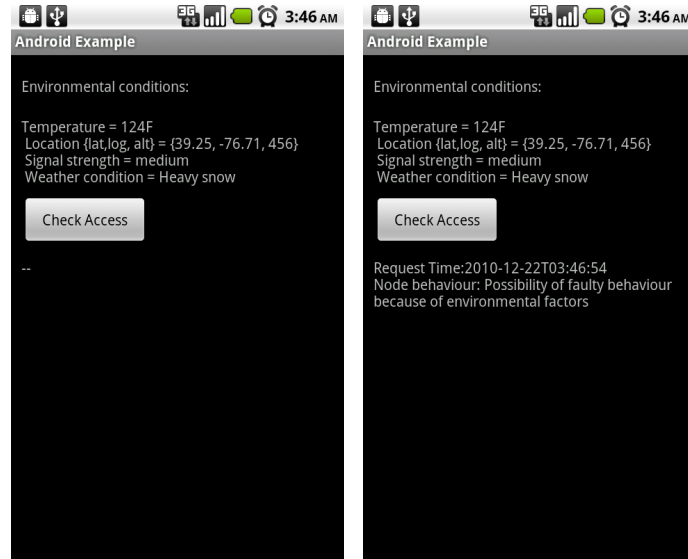
On the other hand, Figure 8(a) illustrates the environment condition for another abnormal sensor report. According to this environment condition, the policy rule concludes that the abnormal report is NOT caused by the environment condition, which is displayed in Figure 8(b).

4.2 Using Semantic Policies for Managing BGP Route Dissemination

The second scenario is use of the framework to protect the traditional internet backbone by automatically configuring BGP systems.

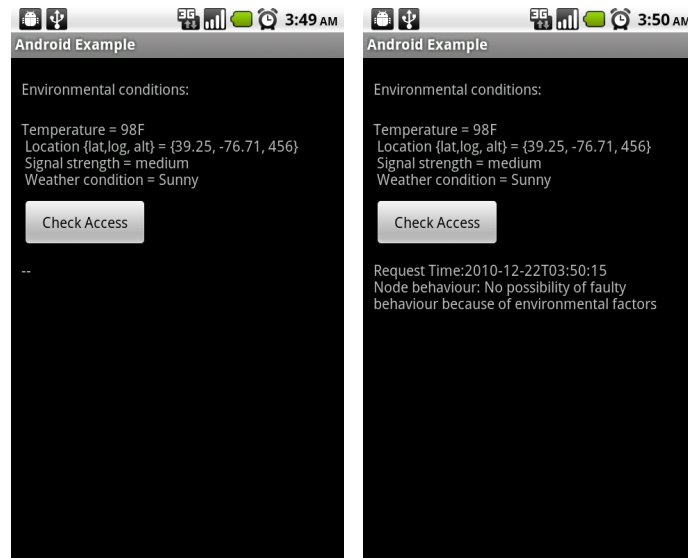
4.2.1 Introduction and Motivation

Border Gateway Protocol (BGP) was originally designed as a simple path vector protocol to share routing information between autonomous systems (AS) which has today, become the de-facto inter-domain routing protocol enabling the Internet. Autonomous systems (ISPs, enterprises etc) use policies, which are driven by various factors such as commercial peering agreements, security considerations, load balancing requirements etc., to define how the routes are to be shared and among which peers. These policies are then implemented in the network routers as configuration parameters that control the protocol behavior.



(a) Environmental Factors in Faulty Case (b) Result of Policy Execution in Faulty Case

Figure 7: Policy Execution in Faulty Case



(a) Environmental Factors in Normal Case (b) Result of Policy Execution in Normal Case

Figure 8: Policy Execution in Normal Case

One of the main challenges is in ensuring that network configuration settings are applied consistently throughout the network so that the correct actions are taken by the network devices both within an autonomous system and across boundaries. Current approaches to configuring BGP routers are operator dependent, device centric, and do not consider overall network objectives. Even under fairly static organizational policies, BGP misconfiguration has been the major cause for internet outage in recent years. Furthermore since routes are expressed as mere IP prefixes such as 127.0/16 with no additional metadata, there is an inherent inflexibility in specifying high level policies such as “Share route with tier I partners”. Furthermore, implementing these configuration changes requires time and highly skilled personnel and is not suitable for scenarios such as emergency response and army battlefield operations where minimizing

deployment time and complexity is vital.

In this example scenario, we address the problem of secure route exchange among peers in a battlefield scenario where deployment time is critical and there are no guarantees of skilled network operators. We propose an alternate model to achieve policy based routing that can provide fine grained policy specification to automate network configuration and ease network management. Specifically, in this chapter, we focus on import and export policies concerning route exchange among peers belonging to different Autonomous Systems (ASes). The model relies on two key components; namely a tagging mechanism that allows routes to convey higher level semantic information that can be used in conjunction with information about the participating BGP peers and a framework for specifying rules in an easy to use, formal model that can be checked for consistency. In our model, ASes encode routes that they originate with descriptions conveying semantics such as what this route represents, who this route can be shared with, traffic type limitations etc using RDF/OWL as a special option and transitive path attribute in BGP. Our motivation for using OWL [28] (specifically, OWL-DL), besides being a W3C standard, is mainly its capabilities for expressing formal semantics, defining class hierarchies and their relationships, associated properties, cardinality restrictions while still retaining decidability and computational completeness. Using OWL for ontology specification makes the framework generic, flexible and more scalable than using proprietary labeling schemes that raise interoperability issues.

Using the framework, BGP speakers can run a reasoning engine that can reason over the RDF descriptions of the various routes and invoke rules depending on the correct set of actions that need to be enforced. Our framework utilizes SWRL [29] as the rule language which provides an easy to use mechanism for specifying event-condition-action rules which constitute the majority of rules envisioned for a typical network. Using this framework, we can control route exchanges at a finer granularity that also enables us to control the traffic flowing in the network.

We show how our architecture can be used to provide fine grained levels of control that is simple to implement and easy to verify for correctness. We have developed a network ontology to be used to describe BGP protocol packets with attributes to describe the route meta-data and example policies to fine tune the BGP decision process. We have also developed a simulation toolkit in NS2 to implement aspects of our proposed architecture allowing us to simulate various scenarios and how policies can be expressed to offer desired behavior.

4.2.2 BGP Routing and Configuration Management

BGP is the de facto routing protocol used in the internet today. BGP started out as a simple path-vector protocol, and with the growing commercialization of the internet, router owners wanted better control over the routing process. Consequently, a number of mechanisms were added to BGP to allow the routing process to be configured in a way that best suits the needs of the organizations in terms of route selection and propagation. For example, import and export policies allow specifying which routes can be accepted and exported respectively. These policies can be used to control the amount of traffic flowing in and out of an Autonomous System (AS). BGP follows a decision process that compares the attributes of BGP routes in selecting a route from among multiple routes to a destination. The list of attributes in descending order of importance are local preference, AS path length, Origin Type, Multi Exit Descriptor (MED), exterior vs interior BGP learned, IGP cost to edge router and router id. By setting appropriate values to the various attributes, operators can tune the decision process to satisfy their economic, political and other policies. For example, ASes can secure themselves by employing appropriate import and export policies. Import policies decide which routes can be accepted by a router. By using a policy that requires routes to be validated before acceptance, ASes can avoid learning invalid routes. Similar policies can be used to reject routes originating from previously know malicious address spaces. On the other hand, export policies determine which routes are exchanged with BGP routers from neighboring ASes. Export policies can be specified that prevent private addresses from being exported outside the AS. Similarly, key infrastructure services could be hidden away from the external world.

BGP policies are implemented as appropriate router configurations in the network. For example, an import policy to reject routes from the prefix 168.10 would be expressed as follows.

```
router bgp 10
neighbor 20.200.1.1 remote-as 20
```

```
neighbor 20.200.1.1 prefix-list PEER-IN in
!
ip prefix-list PEER-IN deny 168.10.0.0/16
```

However, mapping BGP policies to network configurations is not a trivial process. Current approaches are operator dependent, device centric, do not consider the overall network objectives and need to be co-ordinated among multiple routers and ASes. Furthermore, these approaches are not scalable considering the large number of prefixes (131000) and ASes (around 16500), the variety of emerging applications and the dynamically changing network conditions in terms of network traffic and link usage. Furthermore, BGP misconfigurations have often been cited as a major cause for the internet routing architecture going down [30]. Clearly, there is a need for a mechanism that can automatically map high level policies to appropriate network level services without much human intervention. In this work, we propose a semantics driven policy based network that can aid in building such a mechanism

4.2.3 Semantics Driven Policy Based Networks

Policy based networks employ mechanisms that allow network operators to specify at a high level, rules defining how packet flows are handled within a network, how network resources are allocated, access control restrictions and levels of service. All these policies are then enforced by configuring the network devices with the requisite primitives so that the desired actions are performed on the data streams. For example, BGP allows specifying policies that decide whether a router can accept a route from a neighboring router or not.

In previous work [31, 32], we have proposed an architecture for policy based networks that involves semantically tagging packets (in OWL/RDF) to convey higher level meta data about the content being carried in the packets. This semantic information can then be reasoned over at the network elements to provide specialized services in the network. Our policy based network is a multi-tier system with hierarchical policy enforcement with the highest level of the hierarchy being the central NOC for an ISP and the lowest level being an adaptation layer that is responsible for translating the high level policies into low level protocol specific configuration routines that can be applied to the various network elements being managed.

In this work, we have adapted the above framework to handle BGP interactions and use it to specify routing policies. We limit our discussion to how the various components of our general architecture work to drive the BGP decision process. More details on the architecture itself is available in [31, 32].

The *Network Ontology* (NetOnto) is the OWL ontology that we define to mark up the routes being exchanged. By using OWL rather than simple XML, the language is semantically richer and highly extensible which is very important especially when we have interdomain interactions (such as peering arrangements, SLAs etc). Policies are written using the concepts defined in NetOnto using SWRL as the rule language. OWL has axiomatic and model-theoretic semantics, which allows for verification of knowledge expressed in OWL constructs. OWL + SWRL can be used to define ontologies, using which one can declaratively define facts, policies, and rules in terms of what needs to be true or false for a policy to hold. The route descriptions are carried in the BGP updates as optional transitive attributes either as directly embedded in a bit efficient format, or contain a URL to the description or use UUIDs that imply a certain well known description. A Policy Enforcement Point (PEP) extracts this description and adds to it, any extra contextual information including aspects such as peer identity, network state (congestion, link failures etc), network technology (wired, hybrid, MANET, cellular) etc. This information is then sent to the Policy Decision Point (PDP) for reasoning. The response back from the PDP will cause specific configurations to be installed by the PEP on the device (in this work, as we are dealing with import/export policies, the PDP filters appropriately the routes that are exchanged).

4.2.4 Securing BGP through Route Filtering - A Use Case

We describe how our framework can be used to secure BGP route exchange through appropriate import and export policies. To apply the above framework to provide BGP route dissemination that takes into account the security credentials and external relationships, we needed to make two modifications to the protocol. The first modification is aimed at establishing the identity of the BGP peers in a secure and verifiable manner. For this purpose, we assume the BGP session establishment process is extended to

include the sharing of signed credentials to validate the identity of the BGP peers and their affiliations. Prior work such as S-BGP [33] have shown that this is feasible using a public key infrastructure and signed certificates.

This modification is necessary as it is important for a BGP router to establish the identity of its peer so that the routes learned from and advertised to this peer can be handled correctly. The second modification is to include with the route advertisement in the BGP update messages, an additional optional and transitive attribute that conveys semantic meta-data about that NLRI. The intent here is for the originating AS to provide this meta-data so that other nodes can handle the route appropriately. The interim routers are also allowed to add to this description as necessary (keeping the original intact) in a manner that is secure and cannot be repudiated. In this work, we are concerned about the import/export policies in use in the BGP decision process. The modifications allow nodes in our framework to, for each route that is being advertised to or learned from, contact a PDP that will reason over the semantic information provided for that route and the policies that need to be enforced, and communicate to the node whether or not, the route can be shared or accepted.

The use case we consider in this example scenario is that of a secure version of BGP where there are constraints on route exchanges between BGP peers. As with the real Internet, BGP nodes are owned by different agencies that have different affiliations. During the initial session establishment, nodes exchange their identity information to indicate the agencies to which they belong. These agencies or organizations have external socio-economic, political or financial relationships that will influence the BGP nodes in their exchanges. Routes advertised by each AS are tagged with additional semantic information that describe aspects such as its confidentiality, sharing restrictions etc. For such a use case, the following policies would be appropriate:

- Routes marked as “ShareWithFriendly” can only be exchanged between routers that belong to organizations that have a collaborative relationship
- Routes marked as “Restricted” can only be shared between nodes that belong to the same parent organization (even if they are different divisions of that organization)
- Routes marked to be used only for data backup traffic are installed only during non-peak hours
- Allow a route to be used only for data traffic that has a specified or higher clearance level.

We used the ns-BGP [34] extension to NS2 to implement our framework. The network topology considered is a linear network as shown in Figure 9 with nodes grouped into various ASes. Each node is initialized with credentials that specify what organization the node belongs to. We modified the BGP session establishment process to allow the exchange of these credentials so that the BGP nodes can establish the identity and affiliation of the peers with which they are interacting with. We added an additional optional transitive attribute to the BGP update messages to convey additional semantic information about the route. For the network ontology, we used Protege as the editor for specifying our ontology. Jess was used as the reasoning engine. The choice of Jess was mainly motivated by its easy integration with Protege. Other reasoning engines can be used as a replacement if desired.

To begin, we defined an ontology [35] to use for our BGP example. We modeled the various BGP protocol messages and constructs. Since we are dealing with import/export policies, we modeled special instances of classes representing the various actions that a BGP router (PEP) should take such as whether a route should be advertised or not, whether a route should be accepted or not etc. These special instances contain the low level primitive commands that need to be invoked to realize the necessary behavior. In our case, we implemented handlers in the NS2 implementation to handle the response coming back from the reasoner to determine whether a route should be included in an advertisement or whether a route that was received, should be accepted (these commands are expressed as snippets of Tcl code that are evaluated by NS2).

Using this framework, we implemented our typical use case scenario focusing on the import/export policies for BGP. For our example, we consider a network of four autonomous domains with five BGP routers. The Autonomous Domain AS0 belongs to UK forces. The Autonomous Domains AS1 and AS2 belong to two organizations within the US military. Finally, the last Autonomous Domain AS3 belongs to Russian military. During the initial BGP session establishment, the identity of each of the peers is established. This indicates the organization that the router belongs to (US_{Milcom}, UK_{Milcom}, Russian_{Milcom}




Figure 9: A simple network topology was used for testing and evaluation, consisting of a linear network with nodes grouped into various autonomous systems.

```

balaji@pegasus: ~/NS2-28COCONET/ns-2.28/BGPsScripts
File Edit View Terminal Tabs Help
iBGP Validation Test:

Three ASes connected in a line, the middle one containing two
BGP routers, the others just one each.
AS0      AS1      AS2      AS3
n0)-----{ n1 ... n2 }-----{n3}-----{n4
  eBGP    iBGP    eBGP    eBGP
  UK      US      US      Russia

Simulation starts...

time: 0.3
n0 (ip_addr 10.0.0.1) defines a network 10.0.0.0/24 (ShareWithFriendly).
Reasoner OKs announcement of route 10.0.0.0/24 by AS0:10.0.0.1/32 to peer AS1:10.0.1.1/32
Reasoner OKs announcement of route 10.0.0.0/24 by AS1:10.0.1.1/32 to peer AS1:10.0.2.1/32
Reasoner OKs announcement of route 10.0.0.0/24 by AS1:10.0.2.1/32 to peer AS2:10.0.3.1/32
Reasoner denies announcement of route 10.0.0.0/24 by AS2:10.0.3.1/32 to peer AS3:10.0.4.1/32

time: 1.3
n3 (ip_addr 10.0.3.1) defines a network 10.0.3.0/24 (Restricted).
Reasoner OKs announcement of route 10.0.3.0/24 by AS2:10.0.3.1/32 to peer AS1:10.0.2.1/32
Reasoner denies announcement of route 10.0.3.0/24 by AS2:10.0.3.1/32 to peer AS3:10.0.4.1/32
Reasoner OKs announcement of route 10.0.3.0/24 by AS1:10.0.2.1/32 to peer AS1:10.0.1.1/32
Reasoner denies announcement of route 10.0.3.0/24 by AS1:10.0.1.1/32 to peer AS0:10.0.0.1/32

time: 2.3
n4 (ip_addr 10.0.4.1) defines a network 10.0.4.0/24 (None).

```

Figure 10: Simulation Output

etc) which is tracked in the “owner” property of the network devices. Some of these organizations have external relationships (such as NATO to which US_{Milcom} and UK_{Milcom} belong). Such external relationships are modeled through OWL restrictions on properties. For example, a device that is part of NATO is modeled as one where there is a necessary and sufficient constraint that the owner is either an instance of US_{Milcom} , UK_{Milcom} or $France_{Milcom}$. Each router that originates a route includes a description that at the least, indicates the sharing restrictions for that route. In the current version, we have values such as None (which is similar to the “internet” community attribute in BGP), Restricted and ShareWithFriendly as examples. The intention here is that a route marked as “ShareWithFriendly” can only be shared with a peer who can be considered friendly. For example, if we considered forces within NATO to be friendly, a SWRL policy to permit the routes marked as “ShareWithFriendly” to be exchanged could be written as:

```

BGP_Update(?adv) ^
interimRouter(?adv, ?routeradvertising) ^
dest(?adv, ?peer) ^
NATO_Forces(?routeradvertising) ^
NATO_Forces(?peer) ^
routeRestriction(?adv, ?restriction) ^
ShareWithFriendly(?restriction) ^
AllowRouteAdvertisement(?allow)
→ inferredAction(?adv, ?allow)

```

Once the simulation starts, each router advertises its routes with its peers in order to compute its routing table. The simulation proceeds until all routes are computed and the routers converge on their tables. Note that when two routers belonging to UK_{Milcom} and US_{Milcom} (AS0 and AS1) are in a BGP session and while none of the routers have explicitly been identified as belonging to NATO, the

reasoner can deduce this relationship and allow route exchanges between them. Similarly the reasoner can deduce that the route exchange cannot be allowed between AS2 and AS3 as they do not have an explicit relationship that permits this. Figure 10 is a snapshot of the system with the nodes contacting the reasoner to determine if routes can be exchanged and the responses received.

In this manner, we can now setup arbitrary relationships between routers and can specify policies through higher level rule based mechanisms to implement fine grained control over the protocol. This example can be easily extended to scenarios where the relationships are short lived and arbitrary such as in emergency response scenarios (where organizations may temporarily want to share information for providing quick response), application need driven (such as for supporting live event feeds) etc. by extending on the ontology and defining the desired policies.

5 Conclusion and Future Work

In this chapter, we identify why securing critical infrastructure systems such as CPS present challenges beyond what traditional security mechanisms can handle. Such systems have point solutions that either encrypt communication or provide end-device access control. We show how to build distributed, context aware, policy driven systems better suited to protect critical infrastructure using two specific domain examples.

The framework that we have discussed in Section 3 provides a good starting point to cope with security threats in the cyber-physical critical infrastructure. However, the two prototyped systems that we have described in Section 4 merely address a few instances of the broad security/vulnerability problem formulated in Section 1. Thus, there is enough space that remains to be further explored.

Exercises

1. Name three practical examples of cyber-physical critical infrastructure that you observe in your daily life. Identify the specific security need for each example and list out correspond solutions to address the need.
2. Consider the scenario shown in Figure 1, in which different types of sensors are reporting conflicting data. Suppose you are the technician of the utility company, how might you write policies to catch the situation and resolve the problem?
3. Suppose you are assigned to design a traffic alert collection and processing system for the Interstate highway system. What design goals you need to set to make it robust and efficient? How can context and policies be properly used here to help achieve these design goals?
4. Consider the sensor data shown in Table 1. In addition to the example rules that have been discussed, can you come up with any rule that can make use of these sensor data? Also, can you name any additional type of sensor data that may be meaningful to collect in this application?
5. Smart (Power) Grid is a common application of cyber-physical critical infrastructure. From the description in the chapter, and from online sources, identify the types of sensors that may be used in power grid. Articulate the main security challenges that a smart grid would face.
6. Stuxnet is a computer worm discovered in July 2010 that attacks industrial control systems, including those commonly used in critical infrastructures. Please describe the basic features of Stuxnet. What policiess can potentially protect critical infrastructure against such attacks.
7. The route failure for YouTube service that was caused by Pakistan Telecom in 2008 is a good example of BGP routing misconfiguration. Search on Internet and try to come up with policies that would protect from such errors, and write them down as rules in English.
8. For situational awareness application in battlefield, which types of sensor data can help us better make the decision? How will you define policies to properly catch the contextual infomation that sensors collect? Please write a couple of sample policies.

References

- [1] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly, "Identifying, understanding, and analyzing critical infrastructure interdependencies," *IEEE control systems magazine.*, vol. 21, pp. 11–25, 2001.
- [2] Wikipedia, "Cyber-physical system." http://en.wikipedia.org/wiki/Cyber-physical_system
- [3] M. Shiels, "Spies 'infiltrate us power grid'." <http://news.bbc.co.uk/2/hi/technology/7990997.stm>, 2009.
- [4] S. Gorman, "Electricity grid in u.s. penetrated by spies." <http://online.wsj.com/article/SB123914805204099085.html>, 2009.
- [5] W. Li, J. Parker, and A. Joshi, "Security through collaboration in manets," in *Proceedings of 4th International Conference on Collaborative Computing: Networking, Applications and Worksharing, CollaborateCom 2008*, vol. 10 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering (LNICST)*, pp. 696–714, Springer, 2008.
- [6] W. Li and A. Joshi, "Outlier detection in ad hoc networks using dempster-shafer theory," in *Proceedings of the Tenth International Conference on Mobile Data Management: Systems, Services and Middleware, 2009. MDM '09.*, pp. 112–121, IEEE Computer Society, May 2009.
- [7] P. N. N. Laboratory, "Naspinet." <http://www.naspi.org/naspinet.stm>
- [8] L. Kagal, T. Finin, and A. Joshi, "A policy language for a pervasive computing environment," in *Proceedings of IEEE 4th International Workshop on Policies for Distributed Systems and Networks. POLICY 2003.*, 2003.
- [9] Wikipedia, "Floating car data." http://en.wikipedia.org/wiki/Floating_car_data.
- [10] N. R. C. Committee on C4ISR for Future Naval Strike Groups, "C4isr for future naval strike groups," 2006.
- [11] Wikipedia, "Situation awareness." http://en.wikipedia.org/wiki/Situation_awareness.
- [12] Wikipedia, "Intrusion detection." http://en.wikipedia.org/wiki/Intrusion_detection.
- [13] J. Undercoffer, A. Joshi, T. Finin, and J. Pinkston, "A Target-Centric Ontology for Intrusion Detection," in *The 18th International Joint Conference on Artificial Intelligence*, July 2003.
- [14] S. Agarwal, A. Joshi, T. Finin, Y. Yesha, and T. Ganous, "A pervasive computing system for the operating room of the future," *Mobile Networks and Applications*, vol. 12, pp. 215–228, 2007. 10.1007/s11036-007-0010-8.
- [15] K. Rogers, R. Klump, H. Khurana, A. Aquino-Lugo, and T. Overbye, "An authenticated control framework for distributed voltage support on the smart grid," *IEEE Transactions on Smart Grid*, vol. 1, pp. 40–47, june 2010.
- [16] S. Goldwasser, S. Micali, and R. L. Rivest, "A digital signature scheme secure against adaptive chosen-message attacks," *SIAM J. Comput.*, vol. 17, pp. 281–308, April 1988.
- [17] H. Krawczyk, M. Bellare, and R. Canetti, "Hmac: keyed-hashing for message authentication," *RFC*, vol. 2104, pp. 1–12, 1997.
- [18] A. Metke and R. Ekl, "Security technology for smart grid networks," *IEEE Transactions on Smart Grid*, vol. 1, pp. 99–107, june 2010.
- [19] V. Muppavarapu and S. M. Chung, "Role-based access control for cyber-physical systems using shibboleth," in *Proceedings of DHS Workshop on Future Directions in Cyber-Physical Systems Security*, pp. 57–60, 2009.
- [20] M. Erdos and S. Cantor, "Shibboleth-architecture draft v05," 2002.

- [21] L.-A. Tan, X. Yu, S. Kim, J. Han, C.-C. Hung, and W.-C. Peng, "Tru-alarm: Trustworthiness analysis of sensor networks in cyber-physical systems," in *Proceeding of the 10th IEEE International Conference on Data Mining. ICDM '10.*, 2010.
- [22] W. Li, A. Joshi, and T. Finin, "Policy-based malicious peer detection in ad hoc networks," in *Proceedings of the International Conference on Computational Science and Engineering, 2009. CSE '09.*, vol. 3, pp. 76–82, IEEE Computer Society, Aug. 2009.
- [23] W. Li, A. Joshi, and T. Finin, "Coping with node misbehaviors in ad hoc networks: A multi-dimensional trust management approach," in *Proceedings of the 11th International Conference on Mobile Data Management. MDM '10.*, pp. 112–121, IEEE Computer Society, May 2010.
- [24] W. Li, J. Parker, and A. Joshi, "Security through collaboration and trust in manets," *ACM/Springer Mobile Networks and Applications (MONET)*, pp. 1–11, 2010 (Online First).
- [25] M. Sloman, "Policy driven management for distributed systems," *Journal of Network and Systems Management*, vol. 2, pp. 333–360, 1994.
- [26] S. Godik and T. Moses, "Oasis extensible access control markup language (xacml)," 2002.
- [27] X. Zeng, R. Bagrodia, and M. Gerla, "Glomosim: a library for parallel simulation of large-scale wireless networks," *ACM SIGSIM Simulation Digest*, vol. 28, no. 1, pp. 154–161, 1998.
- [28] D. L. McGuinness and F. van Harmelen, "Owl web ontology language overview," tech. rep., W3C Recommendation 10 February 2004, 2004.
- [29] I. Horrocks, P. F. Patel-Schneider, H. Boley, S. Tabet, B. Grosz, and M. Dean, "Swrl: A semantic web rule language combining owl and ruleml," tech. rep., W3C Member submission 21 may 2004, 2004.
- [30] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding bgp misconfiguration," *SIGCOMM Comput. Commun. Rev.*, vol. 32, no. 4, pp. 3–16, 2002.
- [31] S. B. Kodeswaran, O. Ratsimor, A. Joshi, and F. Perich, "Utilizing Semantic Tags for Policy Based Networking," in *Globecom 2007 (accepted for publication)*, November 2007.
- [32] S. B. Kodeswaran and A. Joshi, "Content and context aware networking using semantic tagging," in *ICDEW '06: Proceedings of the 22nd International Conference on Data Engineering Workshops (ICDEW'06)*, (Washington, DC, USA), p. 77, IEEE Computer Society, 2006.
- [33] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol(s-bgp)," *IEEE Journal on Selected Areas in Communication*, vol. 18, pp. 582–592, 2000.
- [34] T. Feng, R. Ballantyne, and L. Trajkovic, "Implementation of bgp in a network simulator," in *Proc. Applied Telecommunications Symposium, ATS'04*, pp. 149–154, April 2004.
- [35] "<http://www.cs.umbc.edu/kodeswar/ontologies/NetworkOnto.owl>."