Flip the Cloud: Cyber-Physical Signaling Games in the Presence of Advanced Persistent Threats

Jeffrey Pawlick(✉), Sadegh Farhang, and Quanyan Zhu

Department of Electrical and Computer Engineering,
Polytechnic School of Engineering, New York University, New York, USA
{jpawlick,farhang,quanyan.zhu}@nyu.edu

Abstract. Access to the cloud has the potential to provide scalable and cost effective enhancements of physical devices through the use of advanced computational processes run on apparently limitless cyber infrastructure. On the other hand, cyber-physical systems and cloud-controlled devices are subject to numerous design challenges; among them is that of security. In particular, recent advances in adversary technology pose Advanced Persistent Threats (APTs) which may stealthily and completely compromise a cyber system. In this paper, we design a framework for the security of cloud-based systems that specifies when a device should trust commands from the cloud which may be compromised. This interaction can be considered as a game between three players: a cloud defender/administrator, an attacker, and a device. We use traditional signaling games to model the interaction between the cloud and the device, and we use the recently proposed `FlipIt` game to model the struggle between the defender and attacker for control of the cloud. Because attacks upon the cloud can occur without knowledge of the defender, we assume that strategies in both games are picked according to prior commitment. This framework requires a new equilibrium concept, which we call *Gestalt Equilibrium*, a fixed-point that expresses the interdependence of the signaling and `FlipIt` games. We present the solution to this fixed-point problem under certain parameter cases, and illustrate an example application of cloud control of an unmanned vehicle. Our results contribute to the growing understanding of cloud-controlled systems.

1 Introduction

Advances in computation and information analysis have expanded the capabilities of the physical plants and devices in cyber-physical systems (CPS)[4,13]. Fostered by advances in cloud computing, CPS have garnered significant attention from both industry and academia. Access to the cloud gives administrators the opportunity to build virtual machines that provide to computational resources with precision, scalability, and accessibility.

Despite the advantages that cloud computing provides, it also has some drawbacks. They include - but are not limited to - accountability, virtualization, and security and privacy concerns. In this paper, we focus especially on providing

MHR Khouzani et al. (Eds.): GameSec 2015, LNCS 9406, pp. 289–308, 2015.
DOI: 10.1007/978-3-319-25594-1_16

accurate signals to a cloud-connected device and deciding whether to accept those signals in the face of security challenges.

Recently, system designers face security challenges in the form of *Advanced Persistent Threats* (*APTs*) [19]. APTs arise from sophisticated attackers who can infer a user's cryptographic key or leverage zero-day vulnerabilities in order to completely compromise a system without detection by the system administrator [16]. This type of stealthy and complete compromise has demanded new types of models [6,20] for prediction and design.

In this paper, we propose a model in which a device decides whether to trust commands from a cloud which is vulnerable to APTs and may fall under adversarial control. We synthesize a mathematical framework that enables devices controlled by the cloud to intelligently decide whether to obey commands from the possibly-compromised cloud or to rely on their own lower-level control.

We model the cyber layer of the cloud-based system using the recently proposed `FlipIt` game [6,20]. This game is especially suited for studying systems under APTs. We model the interaction between the cloud and the connected device using a signaling game, which provides a framework for modeling dynamic interactions in which one player operates based on a belief about the private information of the other. A significant body of research has utilized this framework for security [7–9,15,21]. The signaling and `FlipIt` games are coupled, because the outcome of the `FlipIt` game determines the likelihood of benign and malicious attackers in the robotic signaling game. Because the attacker is able to compromise the cloud without detection by the defender, we consider the strategies of the attacker and defender to be chosen with *prior commitment*. The circular dependence in our game requires a new equilibrium concept which we call a *Gestalt equilibrium*[1]. We specify the parameter cases under which the Gestalt equilibrium varies, and solve a case study of the game to give an idea of how the Gestalt equilibrium can be found in general. Our proposed framework has versatile applications to different cloud-connected systems such as urban traffic control, drone delivery, design of smart homes, etc. We study one particular application in this paper:ef control of an unmanned vehicle under the threat of a compromised cloud.

Our contributions are summarized as follows:

(i) We model the interaction of the attacker, defender/cloud administrator, and cloud-connected device by introducing a novel game consisting of two coupled games: a traditional signaling game and the recently proposed `FlipIt` game.
(ii) We provide a general framework by which a device connected to a cloud can decide whether to follow its own limited control ability or to trust the signal of a possibly-malicious cloud.
(iii) We propose a new equilibrium definition for this combined game: Gestalt equilibrium, which involves a fixed-point in the mappings between the two component games.
(iv) Finally, we apply our framework to the problem of unmanned vehicle control.

[1] Gestalt is a noun which means something that is composed of multiple arts and yet is different from the combination of the parts [2].

In the sections that follow, we first outline the system model, then describe the equilibrium concept. Next, we use this concept to find the equilibria of the game under selected parameter regimes. Finally, we apply our results to the control of an unmanned vehicle. In each of these sections, we first consider the signaling game, then consider the `FlipIt` game, and last discuss the synthesis of the two games. Finally, we conclude the paper and suggest areas for future research.

2 System Model

We model a cloud-based system in which a cloud is subject to APTs. In this model, an *attacker*, denoted by $\mathcal{A}$, capable of APTs can pay an attack cost to completely compromise the cloud without knowledge of the cloud defender. The *defender*, or cloud administrator, denoted by $\mathcal{D}$, does not observe these attacks, but has the capability to pay a cost to reclaim control of the cloud. The cloud transmits a message to a *robot* or other device, denoted by $\mathcal{R}$. The device may follow this command, but it is also equipped with an on-board control system for autonomous operation. It may elect to use its autonomous operation system rather than obey commands from the cloud.

This scenario involves two games: the `FlipIt` game introduced in [20], and the well-known signaling game. The `FlipIt` game takes place between the attacker and cloud defender, while the signaling game takes place between the possibly-compromized cloud and the device. For brevity, denote the `FlipIt` game by $\mathbf{G_F}$, the signaling game by $\mathbf{G_S}$, and the combined game - call it `CloudControl` - by $\mathbf{G_{CC}}$ as shown in Fig. 1. In the next subsections, we formalize this game model.

2.1 Cloud-Device Signaling Game

Let θ denote the type of the cloud. Denote *compromized* and *safe* types of clouds by $\theta_{\mathcal{A}}$ and $\theta_{\mathcal{D}}$ in the set Θ. Denote the probabilities that $\theta = \theta_{\mathcal{A}}$ and that $\theta = \theta_{\mathcal{D}}$ by p and $1-p$. Signaling games typically give these probabilities *apriori*, but in `CloudControl` they are determined by the equilibrium of the `FlipIt` game $\mathbf{G_F}$.

Let m_H and m_L denote messages of high and low risk, respectively, and let $m \in M = \{m_H, m_L\}$ represent a message in general. After $\mathcal{R}$ receives the message, it chooses an action, $a \in A = \{a_T, a_N\}$, where a_T represents *trusting the cloud* and a_N represents *not trusting the cloud*.

For the device $\mathcal{R}$, let $u_{\mathcal{R}}^S : \Theta \times M \times A \to \mathscr{U}_{\mathcal{R}}$, where $\mathscr{U}_{\mathcal{R}} \subset \mathbb{R}$. $u_{\mathcal{R}}^S$ is a utility function such that $u_{\mathcal{R}}^S(\theta, m, a)$ gives the device's utility when the type is θ, the message is m, and the action is a. Let $u_{\mathcal{A}}^S : M \times A \to \mathscr{U}_{\mathcal{A}} \subset \mathbb{R}$ and $u_{\mathcal{D}}^S : M \times A \to \mathscr{U}_{\mathcal{D}} \subset \mathbb{R}$ be utility functions for the attacker and defender. Note that these players only receive utility in $\mathbf{G_S}$ if their own type controls the cloud in $\mathbf{G_F}$, so that type is not longer a necessary argument for $u_{\mathcal{A}}^S$ and $u_{\mathcal{D}}^S$.

Denote the strategy of $\mathcal{R}$ by $\sigma_{\mathcal{R}}^S : A \to [0,1]$, such that $\sigma_{\mathcal{R}}^S(a \,|\, m)$ gives the mixed-strategy probability that $\mathcal{R}$ plays action a when the message is m. The role of the sender may be played by $\mathcal{A}$ or $\mathcal{D}$ depending on the state of the cloud,

determined by $\mathbf{G_F}$. Let $\sigma_{\mathcal{A}}^S : M \to [0,1]$ denote the strategy that $\mathcal{A}$ plays when she controls the cloud, so that $\sigma_{\mathcal{A}}^S(m)$ gives the probability that $\mathcal{A}$ sends message m. (The superscript S specifies that this strategy concerns the signaling game.) Similarly, let $\sigma_{\mathcal{D}}^S : M \to [0,1]$ denote the strategy played by $\mathcal{D}$ when he controls the cloud. Then $\sigma_{\mathcal{D}}^S(m)$ gives the probability that $\mathcal{D}$ sends message m. Let $\Gamma_{\mathcal{R}}^S$, $\Gamma_{\mathcal{A}}^S$, and $\Gamma_{\mathcal{D}}^S$ denote the sets of mixed strategies for each player.

For $\mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}$, define functions $\bar{u}_{\mathcal{X}}^S : \Gamma_{\mathcal{R}}^S \times \Gamma_{\mathcal{X}}^S \to \mathscr{U}_{\mathcal{X}}$, such that $\bar{u}_{\mathcal{X}}^S\left(\sigma_{\mathcal{R}}^S, \sigma_{\mathcal{X}}^S\right)$ gives the expected utility to sender $\mathcal{X}$ when he or she plays mixed-strategy $\sigma_{\mathcal{X}}^S$ and the receiver plays mixed-strategy $\sigma_{\mathcal{R}}^S$. Equation (1) gives $\bar{u}_{\mathcal{X}}^S$.

$$\bar{u}_{\mathcal{X}}^S\left(\sigma_{\mathcal{R}}^S, \sigma_{\mathcal{X}}^S\right) = \sum_{a \in A} \sum_{m \in M} u_{\mathcal{X}}^S(m,a)\, \sigma_{\mathcal{R}}^S(a \mid m)\, \sigma_{\mathcal{X}}^S(m), \; \mathcal{X} \in \{\mathcal{A}, \mathcal{D}\} \tag{1}$$

Next, let $\mu : \Theta \to [0,1]$ represent the belief of $\mathcal{R}$, such that $\mu(\theta \mid m)$ gives the likelihood with which $\mathcal{R}$ believes that a sender who issues message m is of type θ. Then define $\bar{u}_{\mathcal{R}}^S : \Gamma_{\mathcal{R}}^S \to \mathscr{U}_{\mathcal{R}}$ such that $\bar{u}_{\mathcal{R}}^S\left(\sigma_{\mathcal{R}}^S \mid m, \mu(\bullet \mid m)\right)$ gives the expected utility for $\mathcal{R}$ when it has belief μ, the message is m, and it plays strategy $\sigma_{\mathcal{R}}^S$. $\bar{u}_{\mathcal{R}}^S$ is given by

$$\bar{u}_{\mathcal{R}}^S\left(\sigma_{\mathcal{R}}^S \mid m, \mu\right) = \sum_{\theta \in \Theta} \sum_{a \in A} u_{\mathcal{R}}^S(\theta, m, a)\, \mu_R(\theta \mid m)\, \sigma_{\mathcal{R}}^S(a \mid m). \tag{2}$$

The expected utilities to the sender and receiver will determine their incentives to control the cloud in the game $\mathbf{G_F}$ described in the next subsection.

2.2 `FlipIt` Game for Cloud Control

The basic version of `FlipIt` [20][2] is played in continuous time. Assume that the defender controls the resource - here, the cloud - at $t = 0$. Moves for both players obtain control of the cloud if it is under the other player's control. In this paper, we limit our analysis to *periodic* strategies, in which the moves of the attacker and the moves of the defender are both spaced equally apart, and their phases are chosen randomly from a uniform distribution. Let $f_{\mathcal{A}} \in \mathbb{R}_+$ and $f_{\mathcal{D}} \in \mathbb{R}_+$ (where $\mathbb{R}_+$ represents non-negative real numbers) denote the attack and renewal frequencies, respectively.

Players benefit from controlling the cloud, and incur costs from moving. Let $w_{\mathcal{X}}(t)$ denote the average proportion of the time that player $\mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}$ has controlled the cloud up to time t. Denote the number of moves up to t per unit time of player $\mathcal{X}$ by $z_{\mathcal{X}}(t)$. Let $\alpha_{\mathcal{D}}$ and $\alpha_{\mathcal{A}}$ represent the costs of each defender and attacker move. In the original formulation of `FlipIt`, the authors consider a fixed benefit for controlling the cloud. In our formulation, the benefit depends on the equilibrium outcomes of the signaling game $\mathbf{G_S}$. Denote these

[2] See [20] for a more comprehensive definition of the players, time, game state, and moves in `FlipIt`. Here, we move on to describing aspects of our game important for analyzing $\mathbf{G_{CC}}$.

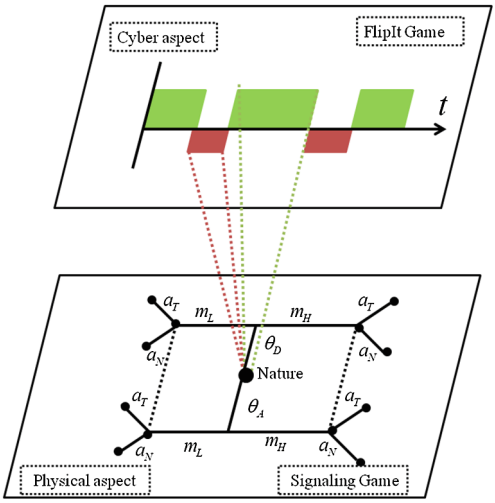

Fig. 1. The `CloudControl` game. The `FlipIt` game models the interaction between an attacker and a cloud administrator for control of the cloud. The outcome of this game determines the type of the cloud in a signaling game in which the cloud conveys commands to the robot or device. The device then decides whether to accept these commands or rely on its own lower-level control. The `FlipIt` and signaling games are played concurrently.

equilibrium utilities of $\mathcal{D}$ and $\mathcal{A}$ by $\bar{u}_{\mathcal{D}}^{S*}$ and $\bar{u}_{\mathcal{A}}^{S*}$. These give the expected benefit of controlling the cloud. Finally, let $u_{\mathcal{D}}^{F}(t)$ and $u_{\mathcal{A}}^{F}(t)$ denote the time-averaged benefit of $\mathcal{D}$ and $\mathcal{A}$ up to time t in $\mathbf{G_F}$. Then

$$u_{\mathcal{X}}^{F}(t) = \bar{u}_{\mathcal{X}}^{S*} w_{\mathcal{X}}(t) - \alpha_{\mathcal{X}} z_{\mathcal{X}}(t), \ \mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}, \tag{3}$$

and, as time continues to evolve, the average benefits over all time become

$$\liminf_{t\to\infty} \bar{u}_{\mathcal{X}}^{S*} w_{\mathcal{X}}(t) - \alpha_{\mathcal{X}} z_{\mathcal{X}}(t), \ \mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}. \tag{4}$$

We next express these expected utilities over all time as a function of periodic strategies that $\mathcal{D}$ and $\mathcal{A}$ employ. Let $\bar{u}_{\mathcal{X}}^{F} : \mathbb{R}_+ \times \mathbb{R}_+ \to \mathbb{R}$, $\mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}$ be expected utility functions such that $\bar{u}_{\mathcal{D}}^{F}(f_{\mathcal{D}}, f_{\mathcal{A}})$ and $\bar{u}_{\mathcal{A}}^{F}(f_{\mathcal{D}}, f_{\mathcal{A}})$ give the average utility to $\mathcal{D}$ and $\mathcal{A}$, respectively, when they play with frequencies $f_{\mathcal{D}}$ and $f_{\mathcal{A}}$. If $f_{\mathcal{D}} \geq f_{\mathcal{A}} > 0$, it can be shown that

$$\bar{u}_{\mathcal{D}}^{F}(f_{\mathcal{D}}, f_{\mathcal{A}}) = \bar{u}_{\mathcal{D}}^{S*}\left(1 - \frac{f_{\mathcal{A}}}{2f_{\mathcal{D}}}\right) - \alpha_{\mathcal{D}} f_{\mathcal{D}}, \tag{5}$$

$$\bar{u}_{\mathcal{A}}^{F}(f_{\mathcal{D}}, f_{\mathcal{A}}) = \bar{u}_{\mathcal{A}}^{S*}\frac{f_{\mathcal{A}}}{2f_{\mathcal{D}}} - \alpha_{\mathcal{A}} f_{\mathcal{A}}, \tag{6}$$

while if $0 \leq f_{\mathcal{D}} < f_{\mathcal{A}}$, then

$$\bar{u}_{\mathcal{D}}^{F}\left(f_{\mathcal{D}}, f_{\mathcal{A}}\right) = \bar{u}_{\mathcal{D}}^{S*} \frac{f_{\mathcal{D}}}{2f_{\mathcal{A}}} - \alpha_{\mathcal{D}} f_{\mathcal{D}}, \tag{7}$$

$$\bar{u}_{\mathcal{A}}^{F}\left(f_{\mathcal{D}}, f_{\mathcal{A}}\right) = \bar{u}_{\mathcal{A}}^{S*} \left(1 - \frac{f_{\mathcal{D}}}{2f_{\mathcal{A}}}\right) - \alpha_{\mathcal{A}} f_{\mathcal{A}}, \tag{8}$$

and if $f_{\mathcal{A}} = 0$, we have

$$\bar{u}_{\mathcal{A}}^{F}\left(f_{\mathcal{D}}, f_{\mathcal{A}}\right) = 0, \quad \bar{u}_{\mathcal{D}}^{F}\left(f_{\mathcal{D}}, f_{\mathcal{A}}\right) = \bar{u}_{\mathcal{D}}^{S*} - \alpha_{\mathcal{D}} f_{\mathcal{D}}. \tag{9}$$

Equations (5)–(9) with Eq. (1) for $\bar{u}_{\mathcal{X}}^{S}$, $\mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}$ and Eq. (2) for $\bar{u}_{\mathcal{R}}^{S}$ will be main ingredients in our equilibrium concept in the next section.

3 Solution Concept

In this section, we develop a new equilibrium concept for our `CloudControl` game $\mathbf{G_{CC}}$. We study the equilibria of the `FlipIt` and signaling games individually, and then show how they can be related through a fixed-point equation in order to obtain an overall equilibrium for $\mathbf{G_{CC}}$.

3.1 Signaling Game Equilibrium

Signaling games are a class of dynamic Bayesian games. Applying the concept of *perfect Bayesian equilibrium* (as it *e.g.*, [10]) to $\mathbf{G_S}$, we have Definition 1.

Definition 1. *Let the functions* $\bar{u}_{\mathcal{X}}^{S}\left(\sigma_{\mathcal{R}}^{S}, \sigma_{\mathcal{X}}^{S}\right)$, $\mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}$ *and* $\bar{u}_{\mathcal{R}}^{S}\left(\sigma_{\mathcal{R}}^{S}\right)$ *be formulated according to Eqs. (1) and (2), respectively. Then a* perfect Bayesian equilibrium *of the signaling game* $\mathbf{G_S}$ *is a strategy profile* $\left(\sigma_{\mathcal{D}}^{S*}, \sigma_{\mathcal{A}}^{S*}, \sigma_{\mathcal{R}}^{S*}\right)$ *and posterior beliefs* $\mu\left(\bullet \,|\, m\right)$ *such that*

$$\forall \mathcal{X} \in \{\mathcal{D}, \mathcal{A}\}, \, \sigma_{\mathcal{X}}^{S*}\left(\bullet\right) \in \arg\max_{\sigma_{\mathcal{X}}^{S}} \bar{u}_{\mathcal{X}}^{S}\left(\sigma_{\mathcal{R}}^{S*}, \sigma_{\mathcal{X}}^{S}\right), \tag{10}$$

$$\forall m \in M, \, \sigma_{\mathcal{R}}^{S*}\left(\bullet \,|\, m\right) \in \arg\max_{\sigma_{\mathcal{R}}^{S}} \bar{u}_{\mathcal{R}}^{S}\left(\sigma_{\mathcal{R}}^{S} \,|\, m, \mu\left(\bullet \,|\, m\right)\right), \tag{11}$$

$$\mu\left(\theta \,|\, m\right) = \frac{1\left\{\theta = \theta_{\mathcal{A}}\right\} \sigma_{\mathcal{A}}^{S*}\left(m\right) p + 1\left\{\theta = \theta_{\mathcal{D}}\right\} \sigma_{\mathcal{D}}^{S*}\left(m\right)\left(1-p\right)}{\sigma_{\mathcal{A}}^{S*}\left(m\right) p + \sigma_{\mathcal{D}}^{S*}\left(m\right)\left(1-p\right)}, \tag{12}$$

if $\sigma_{\mathcal{A}}^{S*}\left(m\right) p + \sigma_{\mathcal{D}}^{S*}\left(m\right)\left(1-p\right) \neq 0$, *and*

$$\mu\left(\theta \,|\, m\right) = \textit{any distribution on } \Theta, \tag{13}$$

if $\sigma_{\mathcal{A}}^{S*}\left(m\right) p + \sigma_{\mathcal{D}}^{S*}\left(m\right)\left(1-p\right) = 0$.

Next, let $\bar{u}_{\mathcal{D}}^{S*}$, $\bar{u}_{\mathcal{A}}^{S*}$, and $\bar{u}_{\mathcal{R}}^{S*}$ be the utilities for the defender, attacker, and device, respectively, when they play according to a strategy profile $\left(\sigma_{\mathcal{D}}^{S*}, \sigma_{\mathcal{A}}^{S*}, \sigma_{\mathcal{R}}^{S*}\right)$ and belief $\mu\left(\bullet \mid m\right)$ that satisfy the conditions for a perfect Bayesian equilibrium. Define a set-valued mapping $T^S : [0,1] \to 2^{\mathcal{U}_{\mathcal{D}} \times \mathcal{U}_{\mathcal{A}}}$ such that $T^S\left(p; G_S\right)$ gives the set of equilibrium utilities of the defender and attacker when the prior probabilities are p and $1-p$ and the signaling game utilities are parameterized by G_S[3]. We have

$$\left\{\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}\right)\right\} = T^S\left(p; G_S\right). \tag{14}$$

We will employ T^S as part of the definition of an overall equilibrium for $\mathbf{G_{CC}}$ after examining the equilibrium of the `FlipIt` game.

3.2 `FlipIt` Game Equilibrium

The appropriate equilibrium concept for the `FlipIt` game, when $\mathcal{A}$ and $\mathcal{D}$ are restricted to periodic strategies, is *Nash equilibrium* [14]. Definition 2 applies the concept of Nash Equilibrim to $\mathbf{G_F}$.

Definition 2. *A* Nash equilibrium *of the game* $\mathbf{G_F}$ *is a strategy profile* $(f_{\mathcal{D}}^*, f_{\mathcal{A}}^*)$ *such that*

$$f_{\mathcal{D}}^* \in \arg\max_{f_{\mathcal{D}}} \bar{u}_{\mathcal{D}}^F\left(f_{\mathcal{D}}, f_{\mathcal{A}}^*\right), \tag{15}$$

$$f_{\mathcal{A}}^* \in \arg\max_{f_{\mathcal{A}}} \bar{u}_{\mathcal{D}}^F\left(f_{\mathcal{D}}^*, f_{\mathcal{A}}\right), \tag{16}$$

where $\bar{u}_{\mathcal{D}}^F$ *and* $\bar{u}_{\mathcal{A}}^F$ *are computed by Eqs. (5) and (6) if* $f_{\mathcal{D}} \geq f_{\mathcal{A}}$ *and Eqs. (7) and (8) if* $f_{\mathcal{D}} \leq f_{\mathcal{A}}$.

To find an overall equilibrium of $\mathbf{G_{CC}}$, we are interested in the proportion of time that $\mathcal{A}$ and $\mathcal{D}$ control the cloud. As before, denote these proportions by p and $1-p$, respectively. These proportions (as in [6]) can be found from the equilibrium frequencies by

$$p = \begin{cases} 0, \text{ if} & f_{\mathcal{A}} = 0 \\ \frac{f_{\mathcal{A}}}{2f_{\mathcal{D}}}, \text{ if} & f_{\mathcal{D}} \geq f_{\mathcal{A}} > 0 \\ 1 - \frac{f_{\mathcal{D}}}{2f_{\mathcal{A}}}, \text{ if} & f_{\mathcal{A}} > f_{\mathcal{D}} \geq 0 \end{cases} \tag{17}$$

Let G_F parameterize the `FlipIt` game. Now, we can define a mapping $T^F : \mathcal{U}_{\mathcal{D}} \times \mathcal{U}_{\mathcal{A}} \to [0,1]$ such that the expression $T^F\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}; G_F\right)$ gives the proportion of time that the attacker controls the cloud in equilibrium from the values of controlling the cloud for the defender and the attacker. This mapping gives

$$p = T^F\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}; G_F\right). \tag{18}$$

[3] Since $\mathcal{R}$ does not take part in $\mathbf{G_S}$, it is not necessary to include $\bar{u}_{\mathcal{R}}^{S*}$ as an output of the mapping.

In addition to interpreting p as the proportion of time that the attacker controls the cloud, we can view it as the likelihood that, at any random time, the cloud will be controlled by the attacker. Of course, this is precisely the value p of interest in $\mathbf{G_S}$. Clearly, $\mathbf{G_F}$ and $\mathbf{G_S}$ are coupled by Eqs. (14) and (18). These two equations specify the overall equilibrium for the `CloudControl` game $\mathbf{G_{CC}}$ through a fixed-point equation, which we describe next.

3.3 Gestalt Equilibrium of $\mathbf{G_{CC}}$

When the `CloudControl` game $\mathbf{G_{CC}}$ is in equilibrium the mapping from the parameters of $\mathbf{G_S}$ to that game's equilibrium and the mapping from the parameters of $\mathbf{G_F}$ to that game's equilibrium are simultaneously satisfied as shown in Fig. 2. Definition 3 formalizes this equilibrium, which we call *Gestalt equilibrium.*

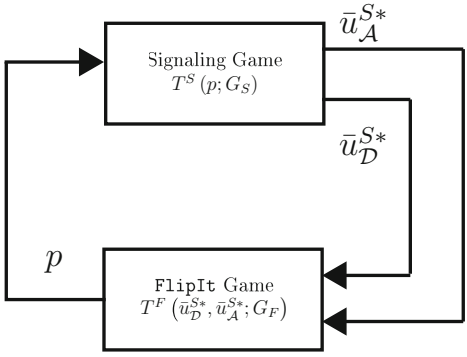

Fig. 2. $\mathbf{G_S}$ and $\mathbf{G_F}$ interact because the utilities in the `FlipIt` game are derived from the output of the signaling game, and the output of the `FlipIt` game is used to define prior probabilities in the signaling game. We call the fixed-point of the composition of these two relationships a Gestalt equilibrium.

Definition 3 ***(Gestalt Equilibrium).*** *The cloud control ratio* $p^\dagger \in [0,1]$ *and equilibrium signaling game utilities* $\bar{u}_{\mathcal{D}}^{S\dagger}$ *and* $\bar{u}_{\mathcal{A}}^{S\dagger}$ *constitute a Gestalt equilibrium of the game* $\mathbf{G_{CC}}$ *composed of coupled games* $\mathbf{G_S}$ *and* $\mathbf{G_F}$ *if the two components of Eq. (19) are simultaneously satisfied.*

$$\left(\bar{u}_{\mathcal{D}}^{S\dagger}, \bar{u}_{\mathcal{A}}^{S\dagger}\right) \in T^S\left(p^\dagger; G_S\right), \quad p^\dagger = T^F\left(\bar{u}_{\mathcal{D}}^{S\dagger}, \bar{u}_{\mathcal{A}}^{S\dagger}; G_F\right) \tag{19}$$

In short, the signaling game utilities $\left(\bar{u}_{\mathcal{D}}^{S\dagger}, \bar{u}_{\mathcal{A}}^{S\dagger}\right)$ *must satisfy the fixed-point equation*

$$\left(\bar{u}_{\mathcal{D}}^{S\dagger}, \bar{u}_{\mathcal{A}}^{S\dagger}\right) \in T^S\left(T^F\left(\bar{u}_{\mathcal{D}}^{S\dagger}, \bar{u}_{\mathcal{A}}^{S\dagger}; G_F\right); G_S\right) \tag{20}$$

In this equilibrium, $\mathcal{A}$ *receives* $\bar{u}_{\mathcal{A}}^{F}$ *according to Eq. (6), Eq. (8), or Eq. (9),* $\mathcal{D}$ *receives* $\bar{u}_{\mathcal{D}}^{F}$ *according to Eq. (5), Eq. (7), or Eq. (9), and* $\mathcal{R}$ *receives* $\bar{u}_{\mathcal{R}}^{S}$ *according to Eq. (2).*

Solving for the equilibrium of $\mathbf{G_{CC}}$ requires a fixed-point equation essentially because the games $\mathbf{G_F}$ and $\mathbf{G_S}$ are played according to *prior committment*. Prior commitment specifies that players in $\mathbf{G_S}$ do not know the outcome of $\mathbf{G_F}$. This structure prohibits us from using a sequential concept such as sub-game perfection and suggests instead a fixed-point equation.

4 Analysis

In this section, we analyze the game proposed in Sect. 2 based on our solution concept in Sect. 3. First, we analyze the signaling game and calculate the corresponding equilibria. Then, we solve the `FlipIt` game for different values of expected payoffs resulting from signaling game. Finally, we describe the solution of the combined game.

4.1 Signaling Game Analysis

The premise of $\mathbf{G_{CC}}$ allows us to make some basic assumptions about the utility parameters that simplifies the search for equilibria. We expect these assumptions to be true across many different contexts.

(A1) $u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_T) > u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_N)$: It is beneficial for the receiver to trust a low risk message from the defender.

(A2) $u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_T) < u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_N)$: It is harmful for the receiver to trust a high risk message from the attacker.

(A3) $\forall m, m' \in M,\ u_{\mathcal{A}}(m, a_T) > u_{\mathcal{A}}(m', a_N)$ and $\forall m, m' \in M\ , u_{\mathcal{D}}(m, a_T) > u_{\mathcal{D}}(m', a_N)$: Both types of sender prefer that either of their messages is trusted rather than that either of their messages is rejected.

(A4) $u_{\mathcal{A}}(m_H, a_T) > u_{\mathcal{A}}(m_L, a_T)$: The attacker prefers an outcome in which the receiver trusts his high risk message to an outcome in which the receiver trusts his low risk message.

Pooling equilibria of the signaling game differ depending on the prior probabilities p and $1-p$. Specifically, the messages on which $\mathcal{A}$ and $\mathcal{D}$ pool and the equilibrium action of $\mathcal{R}$ depend on quantities in Eqs. (21) and (22) which we call *trust benefits.*

$$TB_H(p) = \begin{array}{c} p\left[u_{\mathcal{R}}\left(\theta_{\mathcal{A}}, m_H, a_T\right) - u_{\mathcal{R}}\left(\theta_{\mathcal{A}}, m_H, a_N\right)\right] \\ +(1-p)\left[u_{\mathcal{R}}\left(\theta_{\mathcal{D}}, m_H, a_T\right) - u_{\mathcal{R}}\left(\theta_{\mathcal{D}}, m_H, a_N\right)\right] \end{array} \tag{21}$$

$$TB_L(p) = \begin{array}{c} p\left[u_{\mathcal{R}}\left(\theta_{\mathcal{A}}, m_L, a_T\right) - u_{\mathcal{R}}\left(\theta_{\mathcal{A}}, m_L, a_N\right)\right] \\ +(1-p)\left[u_{\mathcal{R}}\left(\theta_{\mathcal{D}}, m_L, a_T\right) - u_{\mathcal{R}}\left(\theta_{\mathcal{D}}, m_L, a_N\right)\right] \end{array} \tag{22}$$

$TB_H(p)$ and $TB_L(p)$ give the benefit of trusting (compared to not trusting) high and low messages, respectively, when the prior probability is p. These quantities specify whether $\mathcal{R}$ will trust a message that it receives in a pooling equilibrium. If $TB_H(p)$ (respectively, $TB_L(p)$) is positive, then, in equilibrium, $\mathcal{R}$ will trust all messages when the senders pool on m_H (respectively, m_L).

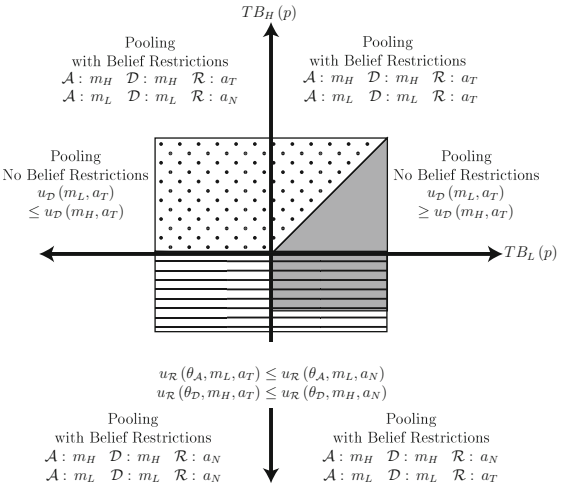

Fig. 3. The four quadrants represent parameter regions of $\mathbf{G_S}$. The regions vary based on the types of pooling equilibria that they support. For instance, quadrant IV supports a pooling equilibrium in which $\mathcal{A}$ and $\mathcal{D}$ both send m_H and $\mathcal{R}$ plays a_N, as well as a pooling equilibrium in which $\mathcal{A}$ and $\mathcal{D}$ both send m_L and $\mathcal{R}$ plays a_T. The shaded regions denote special equilibria that occur under further parameter restrictions.

We illustrate the different possible combinations of $TB_H(p)$ and $TB_L(p)$ in the quadrants of Fig. 3. The labeled messages and actions for the sender and receiver, respectively, in each quadrant denote these pooling equilibria. These pooling equilibria apply throughout each entire quadrant. Note that we have not listed the requirements on belief μ here. These are addressed in the Appendix A.2, and become especially important for various equilibrium refinement procedures.

The shaded regions of Fig. 3 denote additional special equilibria which only occur under the additional parameter constraints listed within the regions. (The geometrical shapes of the shaded regions are not meaningful, but their overlap and location relative to the four quadrants are accurate.) The dotted and uniformly shaded zones contain equilibria similar to those already denoted in the equilibria for each quadrant, except that they do not require restrictions on μ. The zone with horizontal bars denotes the game's only separating equilibrium. It is a rather unproductive one for $\mathcal{D}$ and $\mathcal{A}$, since their messages are not trusted. (See the derivation in Appendix A.1.) The equilibria depicted in Fig. 3 will become the basis of analyzing the mapping $T^S(p; G_S)$, which will be crucial for forming our fixed-point equation that defines the Gestalt equilibrium. Before studying this mapping, however, we first analyze the equilibria of the `FlipIt` game on its own.

4.2 `FlipIt` Analysis

In this subsection, we calculate the Nash equilibrium in the `FlipIt` game. Equations (5)–(9) represent both players' utilities in `FlipIt` game. The solution of

this game is similar to what has presented in [6,20], except that the reward of controlling the resource may vary. To calculate Nash equilibrium, we normalize both players' benefit with respect to the reward of controlling the resource. For different cases, the frequencies of move at Nash equilibrium are:

- $\frac{\alpha_{\mathcal{D}}}{\bar{u}_{\mathcal{D}}^{S*}} < \frac{\alpha_{\mathcal{A}}}{\bar{u}_{\mathcal{A}}^{S*}}$ and $\bar{u}_{\mathcal{A}}^{S*}, \bar{u}_{\mathcal{D}}^{S*} > 0$:

$$f_{\mathcal{D}}^{*} = \frac{\bar{u}_{\mathcal{A}}^{S*}}{2\alpha_{\mathcal{A}}},\ f_{\mathcal{A}}^{*} = \frac{\alpha_{\mathcal{D}}}{2\alpha_{\mathcal{A}}^{2}} \times \frac{(\bar{u}_{\mathcal{A}}^{S*})^{2}}{\bar{u}_{\mathcal{D}}^{S*}}, \tag{23}$$

- $\frac{\alpha_{\mathcal{D}}}{\bar{u}_{\mathcal{D}}^{S*}} > \frac{\alpha_{\mathcal{A}}}{\bar{u}_{\mathcal{A}}^{S*}}$ and $\bar{u}_{\mathcal{A}}^{S*}, \bar{u}_{\mathcal{D}}^{S*} > 0$:

$$f_{\mathcal{D}}^{*} = \frac{\alpha_{\mathcal{A}}}{2\alpha_{\mathcal{D}}^{2}} \times \frac{(\bar{u}_{\mathcal{D}}^{S*})^{2}}{\bar{u}_{\mathcal{A}}^{S*}},\ f_{\mathcal{A}}^{*} = \frac{\bar{u}_{\mathcal{D}}^{S*}}{2\alpha_{\mathcal{D}}}, \tag{24}$$

- $\frac{\alpha_{\mathcal{D}}}{\bar{u}_{\mathcal{D}}^{S*}} = \frac{\alpha_{\mathcal{A}}}{\bar{u}_{\mathcal{A}}^{S*}}$ and $\bar{u}_{\mathcal{A}}^{S*}, \bar{u}_{\mathcal{D}}^{S*} > 0$:

$$f_{\mathcal{D}}^{*} = \frac{\bar{u}_{\mathcal{A}}^{S*}}{2\alpha_{\mathcal{A}}},\ f_{\mathcal{A}}^{*} = \frac{\bar{u}_{\mathcal{D}}^{S*}}{2\alpha_{\mathcal{D}}}, \tag{25}$$

- $\bar{u}_{\mathcal{A}}^{S*} \leq 0$:

$$f_{\mathcal{D}}^{*} = f_{\mathcal{A}}^{*} = 0, \tag{26}$$

- $\bar{u}_{\mathcal{A}}^{S*} > 0$ and $\bar{u}_{\mathcal{D}}^{S*} \leq 0$:

$$f_{\mathcal{D}}^{*} = 0 \quad f_{\mathcal{A}}^{*} = 0^{+}. \tag{27}$$

In the case that $\bar{u}_{\mathcal{A}}^{S*} \leq 0$, the attacker has no incentive to attack the cloud. In this case, the defender need not move since we assume that she controls the cloud initially. In the case that $\bar{u}_{\mathcal{A}}^{S*} > 0$ and $\bar{u}_{\mathcal{D}}^{S*} \leq 0$, only the attacker has an incentive to control the cloud. We use $f_{\mathcal{A}}^{*} = 0^{+}$ to signify that the attacker moves only once. Since the defender never moves, the attacker's single move is enough to retain control of the cloud at all times.

Next, we put together the analysis of $\mathbf{G_S}$ and $\mathbf{G_F}$ in order to study the Gestalt equilibria of the entire game.

4.3 $\mathbf{G_{CC}}$ Analysis

To identify the Gestalt Equilibrium of $\mathbf{G_{CC}}$, it is necessary to examine the mapping $T^{S}(p; G_S)$ for all $p \in [0,1]$. As noted in Sect. 4.1, this mapping depends on $TB_H(p)$ and $TB_L(p)$. From assumptions A1-A4, it is possible to verify that $(TB_L(0), TB_H(0))$ must fall in Quadrant I or Quadrant IV and that $(TB_L(1), TB_H(1))$ must lie in Quadrant III or Quadrant IV. There are numerous ways in which the set $(TB_L(p), TB_H(p))$, $p \in [0,1]$ can transverse different parameter regions. Rather than enumerating all of them, we consider one here.

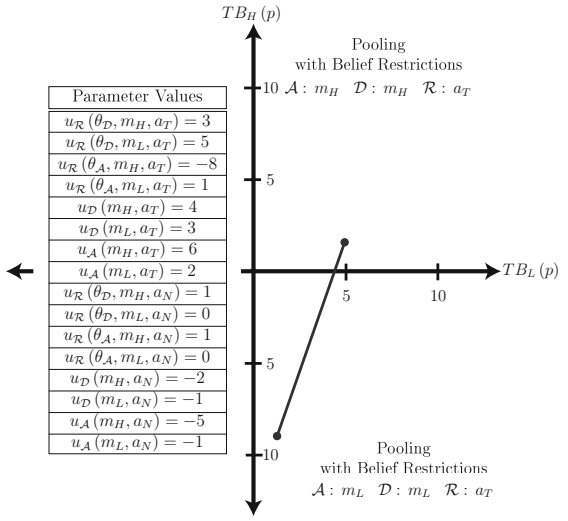

Parameter Values
$u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_T) = 3$
$u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_T) = 5$
$u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_T) = -8$
$u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_T) = 1$
$u_{\mathcal{D}}(m_H, a_T) = 4$
$u_{\mathcal{D}}(m_L, a_T) = 3$
$u_{\mathcal{A}}(m_H, a_T) = 6$
$u_{\mathcal{A}}(m_L, a_T) = 2$
$u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_N) = 1$
$u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_N) = 0$
$u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_N) = 1$
$u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_N) = 0$
$u_{\mathcal{D}}(m_H, a_N) = -2$
$u_{\mathcal{D}}(m_L, a_N) = -1$
$u_{\mathcal{A}}(m_H, a_N) = -5$
$u_{\mathcal{A}}(m_L, a_N) = -1$

Fig. 4. For the parameter values overlayed on the figure, as p ranges from 0 to 1, $TB_H(p)$ and $TB_L(p)$ move from Quadrant I to Quadrant IV. The equilibria supported in each of these quadrants, as well as the equilibria supported on the interface between them, are presented in Table 1.

Consider parameters such that $TB_L(0), TB_H(0) > 0$ and $TB_L(1) > 0$ but $TB_H(1) < 0$[4]. This leads to an $\mathscr{L}$ that will traverse from Quadrant I to Quadrant IV. Let us also assume that $u_{\mathcal{D}}(m_L, a_T) < u_{\mathcal{D}}(m_H, a_T)$, so that Equilibrium 5 is not feasible. In Fig. 4, we give specific values of parameters that satisfy these conditions, and we plot $(TB_L(p), TB_H(p))$ for $p \in [0,1]$. Then, in Table 1, we give the equilibria in each region that the line segment traverses. The equilibrium numbers refer to the derivations in the Appendix A.2.

If p is such that the signaling game is played in Quadrant I, then both senders prefer pooling on m_H. By the *first mover advantage*, they will select Equilibrium 8. On the border between Quadrants I and IV, $\mathcal{A}$ and $\mathcal{D}$ both prefer an equilibrium in which $\mathcal{R}$ plays a_T. If they pool on m_L, this is guaranteed. If they pool on m_H, however, $\mathcal{R}$ receives equal utility for playing a_T and a_N; thus, the senders cannot guarantee that the receiver will play a_T. Here, we assume that the senders maximize their worst-case utility, and thus pool on m_L. This is Equilibrium 3. Finally, in Quadrant IV, both senders prefer to be trusted, and so select Equilibrium 3. From the table, we can see that the utilities will have a jump at the border between Quadrants I and IV. The solid line in Fig. 5 plots the ratio $\bar{u}_{\mathcal{A}}^{S*}/\bar{u}_{\mathcal{D}}^{S*}$ of the utilities as a function of p.

[4] These parameters must satisfy $u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_T) > u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_N)$ and $u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_T) > u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_N)$. Here, we give them specific values in order to plot the data.

Table 1. Signaling game equilibria by region for a game that traverses between Quadrant I and Quadrant IV. Some of the equilibria are feasible only for constrained beliefs μ, specified in Appendix A.2. We argue that the equilibria in each region marked by (*) will be selected.

Region	Equilibria
Quadrant I	Equilibrium 3: Pool on m_L; μ constrained; $\mathcal{R}$ plays a_T
	*Equilibrium 8: Pool on m_H; μ unconstrained; $\mathcal{R}$ plays a_T
$TB_H\,(p) = 0$ Axis	*Equilibrium 3: Pool on m_L; μ constrained; $\mathcal{R}$ plays a_T
	Equilibrium 8: Pool on m_H; μ unconstrained; $\mathcal{R}$ plays a_T
	Equilibrium 6: Pool on m_H; μ constrained; $\mathcal{R}$ plays a_N
Quadrant IV	*Equilibrium 3: Pool on m_L; μ constrained; $\mathcal{R}$ plays a_T
	Equilibrium 6: Pool on m_H; μ constrained; $\mathcal{R}$ plays a_N

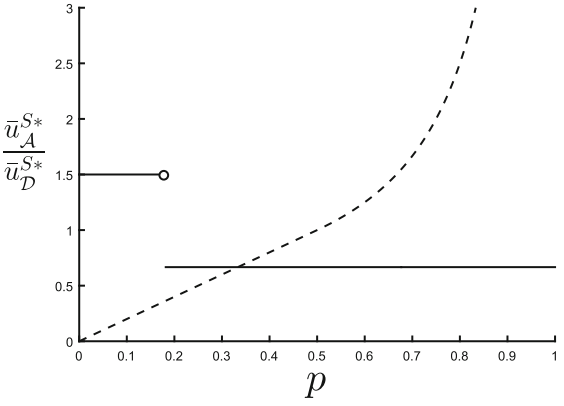

Fig. 5. T^F and T^S are combined on a single set of axis. In T^S (the solid line), the independent variable is on the horizontal axis. In T^F (the dashed line), the independent variable is on the vertical axis. The intersection of the two curves represents the Gestalt equilibrium.

Next, consider the mapping $p = T^F\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}\right)$. As we have noted, p depends only on the ratio $\bar{u}_{\mathcal{A}}^{S*}/\bar{u}_{\mathcal{D}}^{S*}$[5]. Indeed, it is continuous in that ratio when the outcome at the endpoints is appropriately defined. This mapping is represented by the dashed line in Fig. 5, with the independent variable on the vertical axis.

We seek a fixed-point, in which $p = T^F\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}\right)$ and $\left(\bar{u}_{\mathcal{D}}^{S*}, \bar{u}_{\mathcal{A}}^{S*}\right) = T^S\,(p)$. This shown by the intersection of the solid and dashed curves plotted in Fig. 5.

[5] When $\bar{u}_{\mathcal{A}}^{S*} = \bar{u}_{\mathcal{D}}^{S*} = 0$, we define that ratio to be equal to zero, since this will yield $f_{\mathcal{A}} = 0$ and $p = 0$, as in Eqs. (9) and (17). When $\bar{u}_{\mathcal{D}}^{S*} = 0$ and $\bar{u}_{\mathcal{A}}^{S*} > 0$, it is convenient to consider the ratio to be positively infinite since this is consistent with $p \to 1$.

At these points, the mappings between the signaling and the `FlipIt` games are mutually satisfied, and we have a Gestalt equilibrium.[6]

5 Cloud Control Application

In this section, we describe one possible application of our model: a cyber-physical system composed of autonomous vehicles with some on-board control but also with the ability to trust commands from the cloud. Access to the cloud can offer automated vehicles several benefits [12]. First, it allows access to massive computational resources - *i.e., infrastructure as a service* (*IaaS*). (See [5].) Second, it allows access to large datasets. These datasets can offer super-additive benefits to the sensing capabilities of the vehicle itself, as in the case of the detailed road and terrain maps that automated cars such as those created by Google and Delphi combine with data collected by lidar, radar and vision-based cameras [1,11]. Third, interfacing with the cloud allows access to data collected or processed by humans through crowd-sourcing applications; consider, for instance, location-based services [17,18] that feature recommendations from other users. Finally, the cloud can allow vehicles to collectively learn through experience [12].

Attackers may attempt to influence cloud control of the vehicle through several means. In one type of attack, adversaries may be able to *steal or infer cryptographic keys* that allow them authorization into the network. These attacks are of the *complete compromise* and *stealth* types that are studied in the `FlipIt` framework [6,20] and thus are appropriate for a `CloudControl` game. `FlipIt` also provides the ability to model *zero-day exploit*s, vulnerabilities for which a patch is not currently available. Each of these types of attacks on the cloud pose threats to unmanned vehicle security and involve the complete compromise and steathiness that motivate the `FlipIt` framework.

5.1 Dynamic Model for Cloud Controlled Unmanned Vehicles

In this subsection, we use a dynamic model of an autonomous car to illustrate one specific context in which a cloud-connected device could be making a decision of whether to trust the commands that it would receive or to follow its own on-board control.

We consider a car moving in two-dimensional space with a fixed speed v_0 but with steering that can be controlled. (See Fig. 6, which illustrates the "bicycle model" of steering control from [3].) For simplicity, assume that we are interested in the car's deviation from a straight line. (This line might, *e.g.*, run along the

[6] Note that this example featured a discontinuity in signaling game utilities on the border between equilibrium regions. Interestingly, even when the pooling equilibria differ between regions, it is possible that the equilibrium on the border admits a mixed strategy that provides continuity between the different equilibria in the two regions, and thus makes T^S continuous. This could allow $\mathbf{G_{CC}}$ to have multiple Gestalt equilibria.

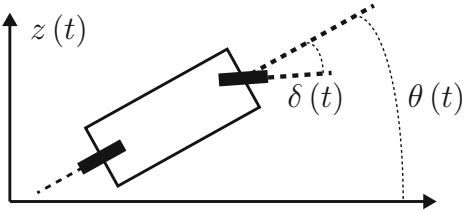

Fig. 6. A bicycle model is a type of representation of vehicle steering control. Here, $\delta(t)$ is used to denote the angle between the orientation of the front wheel and the heading $\theta(t)$. The deviation of the vehicle from a straight line is given by $z(t)$

center of the proper driving lane.) Let $z(t)$ denote the car's vertical distance from the horizontal line, and let $\theta(t)$ denote the heading of the car at time t. The state of the car can be represented by a two-dimensional vector $w(t) \triangleq \left[z(t)\ \theta(t) \right]^T$. Let $\delta(t)$ denote the angle between the orientation of the front wheel - which implements steering - and the orientation of the length of the car. We can consider $\delta(t)$ to be the input to the system. Finally, let $y(t)$ represent a vector of outputs available to the car's control system. The self-driving cars of both Google and Delphi employ radar, lidar, and vision-based cameras for localization. Assume that these allow accurate measurement of both states, such that $y_1(t) = z(t)$ and $y_2(t) = \theta(t)$. If the car stays near $w(t) = \left[0\ 0 \right]^T$, then we can approximate the system with a linear model. Let a and b denote the distances from the rear wheel to the center of gravity and the rear wheel to the front wheel of the car, respectively. Then the linearized system is given in [3] by the equations:

$$\frac{d}{dt}\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} = \begin{bmatrix} 0 & v_0 \\ 0 & 0 \end{bmatrix}\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} + \begin{bmatrix} \frac{av_0}{b} \\ \frac{v_0}{b} \end{bmatrix}\delta(t), \tag{28}$$

$$\begin{bmatrix} y_1(t) \\ y_2(t) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} \tag{29}$$

5.2 Control of Unmanned Vehicle

Assume that the unmanned car has some capacity for automatic control without the help of the cloud, but that the cloud typically provides more advanced navigation.

Specifically, consider a control system onboard the unmanned vehicle designed to return it to the equilibrium $w(t) = \left[0\ 0 \right]^T$. Because the car has access to both of the states, it can implement a state-feedback control. Consider a linear, time-invariant control of the form

$$\delta_{car}(t) = -\left[k_1\ k_2 \right]\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix}. \tag{30}$$

This proportional control results in the closed-loop system

$$\frac{d}{dt}\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} = \left(\begin{bmatrix} 0 & v_0 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} \frac{av_0}{b} \\ \frac{v_0}{b} \end{bmatrix} \begin{bmatrix} k_1 & k_2 \end{bmatrix} \right) \begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} \tag{31}$$

The unmanned car $\mathcal{R}$ may also elect to obtain data or computational resources from the cloud. Typically, this additional access would improve the control of the car. The cloud administrator (defender $\mathcal{D}$), however, may issue faulty commands or there may be a breakdown in communication of the desired signals. In addition, the cloud may be compromised by $\mathcal{A}$ in a way that is stealthy. Because of these factors, $\mathcal{R}$ sometimes benefits from rejecting the cloud's command and relying on its own navigational abilities. Denote the command issued by the cloud at time t by $\delta_{cloud}(t) \in \delta_{\mathcal{A}}(t), \delta_{\mathcal{D}}(t)$, depending on who controls the cloud. With this command, the system is given by

$$\frac{d}{dt}\begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} = \begin{bmatrix} 0 & v_0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z(t) \\ \theta(t) \end{bmatrix} + \begin{bmatrix} \frac{av_0}{b} \\ \frac{v_0}{b} \end{bmatrix} \delta_{cloud}(t). \tag{32}$$

5.3 Filter for High Risk Cloud Commands

In cloud control of an unmanned vehicle, the self-navigation state feedback input given by $\delta_{car}(t)$ in Eq. (30) represents the control that is expected by the vehicle given its state. If the signal from the cloud differs significantly from the signal given by the self-navigation system, then the vehicle may classify the message as "high-risk." Specifically, define a *difference threshold* τ, and let

$$m = \begin{cases} m_H, \text{ if } & |\delta_{cloud}(t) - \delta_{car}(t)| > \tau \\ m_L, \text{ if } & |\delta_{cloud}(t) - \delta_{car}(t)| \leq \tau \end{cases} \tag{33}$$

Equation (33) translates the actual command from the cloud (controlled by $\mathcal{D}$ or $\mathcal{A}$) into a message in the cloud signaling game.

Equations (31) and (32) give the dynamics of the unmanned car electing to trust and not trust the cloud. Based on these equations, Fig. 7 illustrates the combined self-navigating and cloud controlled system for vehicle control.

6 Conclusion and Future Work

In this paper, we have proposed a general framework for the interaction between an attacker, cloud administrator/defender, and cloud-connected device. We have described the struggle for control of the cloud using the `FlipIt` game and the interaction between the cloud and the connected device using a traditional signaling game. Because these two games are played by prior commitment, they are coupled. We have defined a new equilibrium concept - *i.e., Gestalt equilibrium*, which defines a solution to the combined game using a fixed-point equation. After illustrating various parameter regions under which the game may be

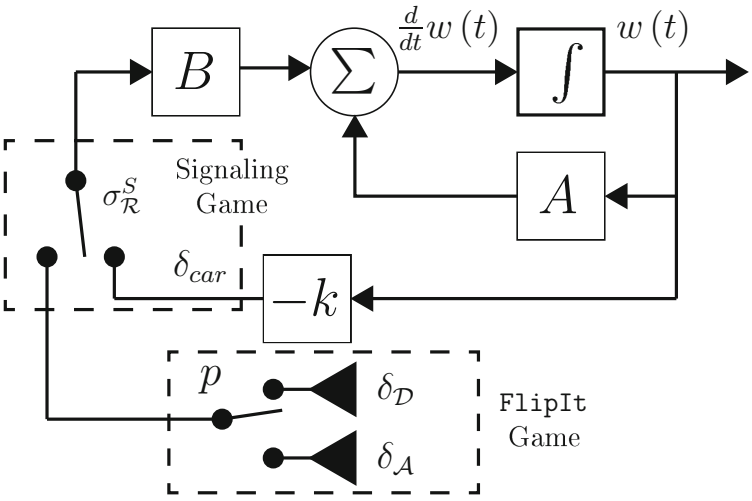

Fig. 7. Block-diagram model for unmanned vehicle navigation control. At any time, the vehicle uses strategy $\sigma_{\mathcal{R}}^{S}$ to decide whether to follow its own control or the control signal from the cloud, which may be $\delta_{\mathcal{A}}$ or $\delta_{\mathcal{D}}$, depending on the probabilities p, $1-p$ with which $\mathcal{A}$ and $\mathcal{D}$ control the cloud. Its own control signal, δ_{car}, is obtained via feedback control.

played, we solved the game in a sample parameter region. Finally, we showed how the framework may be applied to unmanned vehicle control.

Several directions remain open for future work. First, the physical component of the cyber-physical system can be further examined. Tools from optimal control such as the linear-quadratic regulator could offer a rigorous framework for defining the costs associated with the physical dynamic system, which in turn would define the payoffs of the signaling game. Second, future work could search for conditions under which a Gestalt equilibrium of the `CloudControl` game is guaranteed to exist. Finally, devices that use this framework should be equipped to learn online. Towards that end, a learning algorithm could be developed that is guaranteed to converge to the Gestalt equilibrium. Together with the framework developed in the present paper, these directions would help to advance our ability to secure cloud-connected and cyber-physical systems.

A Derivation of Signaling Game Equilibria

In this appendix, we solve for the equilibria of $\mathbf{G_S}$.

A.1 Separating Equilibria

First, we search for separating equilibria of $\mathbf{G_S}$. In separating equilibria, $\mathcal{R}$ knows with certainty the type of the cloud.

$\mathcal{D}$ plays m_L and $\mathcal{A}$ plays m_H. If $\mathcal{D}$ plays m_L (as a pure strategy) and $\mathcal{A}$ plays m_H, then the receiver rejects any m_H according to assumption A2. The best action for $\mathcal{A}$ is to deviate to m_L. Thus, this is not an equilibrium.

$\mathcal{D}$ plays m_H and $\mathcal{A}$ plays m_L. If $\mathcal{D}$ plays m_H and $\mathcal{A}$ plays m_L, the $\mathcal{R}$'s best response depends on the utility parameters. If $u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_T) \leq u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_N)$ and $u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_T) \leq u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_N)$, then $\mathcal{R}$ plays a_N in response to both messages. There is no incentive to deviate. Denote this separating equilibrium as *Equilibrium #2*.

If $u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_T) \leq u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_N)$ and $u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_T) > u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_N)$, then a_N is within the set of best responses to m_L, whereas a_T is the unique best response to m_H. Assuming that he prefers to certainty receive a higher utility, $\mathcal{A}$ deviates to m_H.

If $u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_T) > u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_N)$ and $u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_T) \leq u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_N)$, then a_N is within the set of best responses to m_H, whereas a_T is the unique best response to m_L. Thus, $\mathcal{D}$ deviates to m_L.

If $u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_T) > u_{\mathcal{R}}^S(\theta_{\mathcal{A}}, m_L, a_N)$ and $u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_T) > u_{\mathcal{R}}^S(\theta_{\mathcal{D}}, m_H, a_N)$, then $\mathcal{R}$ plays a_T in response to both messages. We have assumed, however, that $\mathcal{A}$ prefers to be trusted on m_H compared to being trusted on m_L (A4), so $\mathcal{A}$ deviates and this is not an equilibrium.

A.2 Pooling Equilibria

Next, we search for pooling equilibria of $\mathbf{G_S}$. In pooling equilibria, $\mathcal{R}$ relies only on the prior probabilities p and $1-p$ in order to form his belief about the type of the cloud. The existence of pooling equilibria depend essentially on the trust benefits $TB_H(p)$ and $TB_L(p)$.

Pooling on m_L. If $TB_L(p) < 0$, then $\mathcal{R}$'s best response is a_N. This will only be an equilibrium if his best response to m_H would also be a_N. This is the case only when the belief satisfies

$$\begin{aligned} &\mu(\theta_{\mathcal{A}} \mid m_H)\, u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_T) + (1-\mu(\theta_{\mathcal{A}} \mid m_H))\, u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_T) \\ &\leq \mu(\theta_{\mathcal{A}} \mid m_H)\, u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_H, a_N) + (1-\mu(\theta_{\mathcal{A}} \mid m_H))\, u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_H, a_N) \end{aligned} \tag{34}$$

Moreover, this can only be an equilibrium when neither $\mathcal{A}$ nor $\mathcal{D}$ have an incentive to deviate: *i.e.*, when

$$u_{\mathcal{A}}^S(m_H, a_N) \leq u_{\mathcal{A}}^S(m_L, a_N) \text{ and } u_{\mathcal{D}}^S(m_H, a_N) \leq u_{\mathcal{D}}^S(m_L, a_N) \tag{35}$$

If these conditions are satisfied, then denote this equilibrium by *Equilibrium #1*.

If $TB_L(p) \geq 0$, then $\mathcal{R}$'s best response us a_T. Whether this represents an equilibrium depends on if $\mathcal{A}$ or $\mathcal{D}$ have incentives to deviate from m_L. If $u_{\mathcal{D}}^S(m_H, a_T) \leq u_{\mathcal{D}}^S(m_L, a_T)$ and $u_{\mathcal{A}}^S(m_H, a_T) \leq u_{\mathcal{A}}^S(m_L, a_T)$, then neither has an incentive to deviate. This is *Equilibrium #5*. If one of these inequalities does

not hold, then the player who prefers m_H to m_L will deviate if $\mathcal{R}$ would play a_T in response to the deviation. The equilibrium condition is narrowed to when the belief makes $\mathcal{R}$ not trust m_H; when Eq. (34) is satisfied. Call this *Equilibrium #3*.

Pooling on m_H. The pattern of equilibria for pooling on m_H follows a similar structure to the pattern of equilibria for pooling on m_L.

If $TB_H(p) < 0$, then $\mathcal{R}$'s best response is a_N. This will only be an equilibrium if his best response to m_L would also be a_N. This is the case only when the belief satisfies

$$\begin{aligned} &\mu(\theta_{\mathcal{A}} \mid m_L)\, u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_T) + (1 - \mu(\theta_{\mathcal{A}} \mid m_L))\, u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_T) \\ &\leq \mu(\theta_{\mathcal{A}} \mid m_L)\, u_{\mathcal{R}}(\theta_{\mathcal{A}}, m_L, a_N) + (1 - \mu(\theta_{\mathcal{A}} \mid m_L))\, u_{\mathcal{R}}(\theta_{\mathcal{D}}, m_L, a_N) \end{aligned} \tag{36}$$

To guarantee that $\mathcal{A}$ and $\mathcal{D}$ do not deviate, we require

$$u^S_{\mathcal{A}}(m_H, a_N) \geq u^S_{\mathcal{A}}(m_L, a_N) \text{ and} u^S_{\mathcal{D}}(m_H, a_N) \geq u^S_{\mathcal{D}}(m_L, a_N) \tag{37}$$

If these conditions are satisfied, then we have *Equilibrium #6*.

If $TB_H \geq 0$, then $\mathcal{R}$'s best response is a_T. If $u^S_{\mathcal{D}}(m_H, a_T) \geq u^S_{\mathcal{D}}(m_L, a_T)$ and $u^S_{\mathcal{A}}(m_H, a_T) \geq u^S_{\mathcal{A}}(m_L, a_T)$, then neither $\mathcal{A}$ nor $\mathcal{D}$ have an incentive to deviate. Call this *Equilibrium #8*. If one of these inequalities does not hold, then the belief must satisfy Eq. (36) for an equilibrium to be sustained. Denote this equilibrium by *Equilibrium #7*.

References

1. Delphi drive, Delphi Automotive. http://www.delphi.com/delphi-drive
2. Gestalt, Mirium-Webster. http://www.merriam-webster.com/dictionary/gestalt
3. Aström, K.J., Murray, R.M.: Feedback Systems: An Introduction for Scientists and Engineers. Princeton University Press, Princeton (2010)
4. Baheti, R., Gill, H.: Cyber-physical systems. In: The Impact of Control Technology, vol. 12, pp. 161–166 (2011)
5. Bhardwaj, S., Jain, L., Jain, S.: Cloud computing: A study of infrastructure as a service (IAAS). Int. J. Eng. Inf. Technol. **2**(1), 60–63 (2010)
6. Bowers, K.D., van Dijk, M., Griffin, R., Juels, A., Oprea, A., Rivest, R.L., Triandopoulos, N.: Defending against the unknown enemy: applying FlipIt to system security. In: Grossklags, J., Walrand, J. (eds.) GameSec 2012. LNCS, vol. 7638, pp. 248–263. Springer, Heidelberg (2012)
7. Carroll, T.E., Grosu, D.: A game theoretic investigation of deception in network security. In: Security and Communication, Networks vol. 4(10), pp. 1162–1172 (2011)
8. Casey, W., Morales, J.A., Nguyen, T., Spring, J., Weaver, R., Wright, E., Metcalf, L., Mishra, B.: Cyber security via signaling games: toward a science of cyber security. In: Natarajan, R. (ed.) ICDCIT 2014. LNCS, vol. 8337, pp. 34–42. Springer, Heidelberg (2014)

9. Farhang, S., Manshaei, M.H., Esfahani, M.N., Zhu, Q.: A dynamic bayesian security game framework for strategic defense mechanism design. In: Poovendran, R., Saad, W. (eds.) GameSec 2014. LNCS, vol. 8840, pp. 319–328. Springer, Heidelberg (2014)
10. Fudenberg, D., Tirole, J.: Game Theory, vol. 393. MIT press, Cambridge (1991)
11. Guizzo, E.: How googles self-driving car works. IEEE Spectrum Online, 18 October
12. Kehoe, B., Patil, S., Abbeel, P., Goldberg, K.: A survey of research on cloud robotics and automation. IEEE Trans. Autom. Sci. Eng. **12**(2), 398–409 (2015)
13. Lee, E.A.: Cyber physical systems: design challenges. In: 2008 11th IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC), pp. 363–369. IEEE (2008)
14. Nash, J.F., et al.: Equilibrium points in n-person games. Proc. Nat. Acad. Sci. USA **36**(1), 48–49 (1950)
15. Pawlick, J., Zhu, Q.: Deception by design: Evidence-based signaling games for network defense. arXiv preprint arXiv:1503.05458 (2015)
16. Portokalidis, G., Slowinska, A., Bos, H.: Argos: an emulator for fingerprinting zero-day attacks for advertised honeypots with automatic signature generation. ACM SIGOPS Operating Syst. Rev. **40**(4), 15–27 (2006)
17. Sampigethaya, K., Huang, L., Li, M., Poovendran, R., Matsuura, K., Sezaki, K.: Caravan: Providing location privacy for vanet. Technical report, DTIC Document (2005)
18. Sampigethaya, K., Li, M., Huang, L., Poovendran, R.: Amoeba: Robust location privacy scheme for vanet. IEEE J. Sel. Areas Commun. **25**(8), 1569–1589 (2007)
19. Tankard, C.: Advanced persistent threats and how to monitor and deter them. Netw. Secur. **2011**(8), 16–19 (2011)
20. van Dijk, M., Juels, A., Oprea, A., Rivest, R.L.: Flipit: The game of "stealthy takeover". J. Cryptol. **26**(4), 655–713 (2013)
21. Zhuang, J., Bier, V.M., Alagoz, O.: Modeling secrecy and deception in a multiple-period attacker-defender signaling game. Eur. J. Oper. Res. **203**(2), 409–418 (2010)