

Double-Blind Reputation vs. Intelligent Fake VIP Attacks in Cloud-Assisted Interactions

Jerzy Konorski
Faculty of Electronics, Telecommunications and Informatics
Gdansk University of Technology
Gdansk, Poland
jekon@eti.pg.gda.pl

Post-print of: J. Konorski, "Double-Blind Reputation vs. Intelligent Fake VIP Attacks in Cloud-Assisted Interactions," 2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE), New York, NY, 2018, pp. 1637-1641, doi: 10.1109/TrustCom/BigDataSE.2018.00241.

© 2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Abstract—We consider a generic model of Client-Server interactions in the presence of Sender and Relay, conceptual agents acting on behalf of Client and Server, respectively, and modeling cloud service providers in the envisaged "QoS as a Service paradigm". Client generates objects which Sender tags with demanded QoS level, whereas Relay assigns the QoS level to be provided at Server. To verify an object's right to a QoS level, Relay detects its signature that neither Client nor Sender can modify. Since signature detection is costly Relay tends to occasionally skip it and trust an object; this prompts Sender to occasionally launch a Fake VIP attack, i.e., demand undue QoS level. In a Stackelberg game setting, Relay employs a trust strategy in the form of a double-blind reputation scheme so as to minimize the signature detection cost and undue QoS provision, anticipating a best-response Fake VIP attack strategy on the part of Sender. The paper addresses the question whether the double-blind reputation scheme, previously proved resilient to a probabilistic Fake VIP attack strategy, is equally resilient to more intelligent Sender behavior. Two intelligent attack strategies are proposed and analyzed using two-dimensional Markov chains.

Keywords—QoS, cloud computing, Fake VIP attack, reputation, trust, Stackelberg game

I. INTRODUCTION

Many computer communication scenarios can be modeled as Client-Server interactions: Client generates a sequence of *objects* (e.g., data files, queries, or transactions), each of some *intrinsic class* entitling the object to a specific quality of service (QoS) level at Server. For simplicity assume that only L (low) and H (high) classes are distinguished. Each object is passed to Server with a *demand class* tag. Server decides the *assigned class* (i.e., determines the QoS level) based on the object's demanded class and detected *signature*. The latter refers to a set of intrinsic class-dependent features that Client cannot modify. Signature detection is costly (e.g., may involve computationally intensive pattern matching or context analysis), therefore Server gladly outsources this functionality to her local cloud, a third agent called Relay. Likewise, clever class demanding may bring Client undue benefits (while depleting Server's resources) and requires a separate functionality, which Client gladly outsources to her local cloud, a fourth agent called Sender. Thus we envisage a *QoS as a Service* paradigm analogous to Security as a Service [1]. In this paradigm, Client and Server outsource their QoS-related decision-making functionalities to local clouds, respectively Sender and Relay, and any strategic considerations are restricted to the latter two agents (Fig. 1).




Fig. 1. Cloud-assisted Client-Server interactions.

As signature detection is costly, Relay may occasionally skip it for an arrived object. Aware of this, Sender is tempted to occasionally launch a *Fake VIP attack* upon (demand class H for) an object bearing class L signature according to some *Fake VIP attack strategy*. Relay assigns class L or H according to some *trust strategy*. Following [2], the trust strategy employs a *double-blind reputation scheme* that neither can observe an object's intrinsic class nor reveals to Sender the assigned class or current reputation state. If an object finds on arrival a high enough reputation state, it is *trusted*, i.e., Relay skips signature detection. Thus a noncooperative game arises between Sender and Relay (acting on behalf of Client and Server, respectively). In a Stackelberg game setting, Sender seeks a best response to Relay's trust strategy and selects her Fake VIP attack strategy so as to maximize her expected utility related to a long-term perception of received QoS; anticipating this, Relay selects her trust strategy so as to minimize a utility reflecting the expected combined cost of signature detection and high QoS provision by Server. At the resulting *Stackelberg equilibrium* (SE) both players' expected utilities can be compared to *reputation-free* (RF) play where Relay always trusts an object if the signature detection cost is high enough and never trusts otherwise.

The following postulates [2] reflect some practical aspects of the above scenario and make it nontrivial to analyze: (i) a Fake VIP attack is costless (e.g., Sender is not charged merely for demanding class H), (ii) Relay never gets to know an object's intrinsic class, (iii) signature detection cannot be conditioned upon the demanded class (which may be part of the signature), (iv) high QoS provision is costly for Server, (v) Sender cannot observe individual objects' assigned classes, and (vi) Relay may not cheat, i.e., assign class L upon detection of class H signature. These postulates constitute an *information-separation framework* in which information accessible by both parties is kept to a minimum; we believe such philosophy should underlie cautious design of any security relationship.

Usurpation attacks similar to Fake VIP, called *traffic re-mapping attacks*, were studied in [3] and shown to threaten wireless networks employing the Enhanced Distributed Channel Access mode of the IEEE 802.11 MAC protocol. Various protocol-specific defense measures have been proposed in security frameworks quite different from ours, typically violating postulates (ii), (iii), (v), or (vi). Complex context analysis-based measures such as [4] are environment specific too, and typically violate postulate (ii). Against the proposed double-blind reputation scheme, Sender might attempt to learn the present reputation state via online prediction [5], repeated games [6] or some form of reinforcement learning. However, such approaches are inapplicable or highly ineffective in an information-separation framework, especially due to postulate (v). Game-theoretic Intrusion Detection System models such as [7] capitalize on Sender's fear of attack detection; however, due to postulate (vi), a Fake VIP attack cannot be directly punished.

In a previous paper [2] we studied a probabilistic Fake VIP attack strategy called pFVIP, whereby an object bearing class L signature is attacked with a fixed intrinsic class-dependent probability. We found that surprisingly, SE play against pFVIP brings Relay utility gains compared to RF play regardless of the signature detection cost, even if the object generation process is memoryless. This may indicate that the proposed double-blind reputation scheme is a powerful defense—or that pFVIP is not sophisticated enough and the question is whether Relay's defense would be resilient to more intelligent Fake VIP attack strategies. In this paper we study two such strategies to find that while they may improve Sender's utility, Relay's utility is not visibly worsened.

II. MODEL

Object generation at Client follows a binary stationary memoryless random process $(c^{(k)})_{k=1,2,\dots}$ where $c^{(k)} \in \{L, H\}$ denotes the intrinsic class of the k th generated object. Let $\rho = \Pr[c^{(k)} = H]$. Each generated object bears a signature $s^{(k)} \in \{L, H\}$, i.e., a set of intrinsic class-dependent features that Relay recognizes as relevant to the class to be assigned. As explained in the introduction, an object of a given intrinsic class may accidentally bear a "wrong" signature. Let $\varepsilon_L = \Pr[s^{(k)} = H \mid c^{(k)} = L]$ and $\varepsilon_H = \Pr[s^{(k)} = L \mid c^{(k)} = H]$, be the stationary signature error rates. Note that the signature detection scheme Relay employs to decide assigned class for an object may not be known to Sender (moreover, signatures may incur non-deterministic corruption in the communication channel between Sender and Relay). Thus Sender is unaware of objects' signatures and employs a *signature prediction scheme*. Denote the stationary prediction error rates by $\zeta_L = \Pr[p^{(k)} = H \mid s^{(k)} = L]$ and $\zeta_H = \Pr[p^{(k)} = L \mid s^{(k)} = H]$. Let us introduce two auxiliary quantities: $\omega = \Pr[s^{(k)} = L] = (1 - \rho)(1 - \varepsilon_L) + \rho\varepsilon_H$ and $\Omega = \Pr[p^{(k)} = L] = \omega(1 - \zeta_L) + (1 - \omega)\zeta_H$. We will also use a common superscript notation for vectors of attributes of the same object, e.g., $(s,d)^{(k)} = (s^{(k)}, d^{(k)})$.

A. Fake VIP Attack and Trust Strategies

Sender sets the k th object's demanded class $d^{(k)}$ based on its *predicted class* $p^{(k)}$, i.e., predicted outcome of the signature detection at Relay:

$$d^{(k)} = \begin{cases} H, & p^{(k)} = H \vee \sigma^{(k)}, \\ L, & \text{otherwise,} \end{cases} \quad (1)$$

where $\sigma^{(k)}$ is defined by Sender's Fake VIP attack strategy $\sigma(\cdot)$.

Relay decides the assigned class $a^{(k)}$ for the k th object based on signature detection and current *reputation state* $r^{(k)} \in \{1, \dots, R\}$, where $R \geq 2$. (Recall that $d^{(k)}$ may not be known at decision time, cf. postulate (iii).) The reputation state reflects Relay's perception of Sender's recent behavior. In the *trust state* R , Relay trusts the object and passes it to Server as is, implying $a^{(k)} = d^{(k)}$ being set by Server; in such a case we

conservatively assume that Relay does not observe $d^{(k)}$ or $a^{(k)}$, i.e., gets no feedback from Server. In the non-trust states $1, \dots, R - 1$, Relay performs signature detection and decides $a^{(k)} = s^{(k)}$. According to postulate (v), Sender does not observe $a^{(k)}$ and only has a (possibly imperfect) perception of the long-term frequency of $a^{(k)} = H$, i.e., of the probability $\Pr[a^{(k)} = H]$.

If $r^{(k)} < R$ and $(s,d)^{(k)} = (L,H)$, i.e., a Fake VIP attack is suspected (though not certain because of a possible signature error), then Sender's reputation state is lowered, whereas a perceived honest demand, i.e., $(s,d)^{(k)} = (L,L)$, raises the reputation state. In the trust state R , where neither $s^{(k)}$ nor $d^{(k)}$ is observed, a cautionary action is to lower the reputation state with stationary probability $\Pr[d^{(k)} = H]$ estimated from objects arrived in non-trust states. In other situations the reputation state remains unchanged. The reputation scheme defines a parameter $\delta \in [0, 1]$ that measures the tendency to lower the current reputation state. The pair (R, δ) is Relay's private information and fully characterizes her trust strategy $\pi(\cdot)$. Formally, let $\Phi = (1 < r^{(k)} < R \wedge (s,d)^{(k)} = (L,H))$, $\Theta = r^{(k)} < R \wedge (s,d)^{(k)} = (L,L)$, $\Psi = (r^{(k)} = R \wedge \text{rand}(\Pr[d^{(k)} = H]))$, $\text{rand}(\cdot)$ be a random event occurring with the specified probability, and $\mathbf{1}_{(\cdot)}$ be the indicator function; then

$$r^{(k+1)} = r^{(k)} - \mathbf{1}_{(\Phi \wedge \text{rand}(\delta)) \vee \Psi} + \mathbf{1}_{\Theta \wedge \neg \text{rand}(\delta)}. \quad (2)$$

(Note the information separation: only $d^{(k)}$ is observed by both Sender and Relay, and only if $r^{(k)} < R$.)

B. Utilities

Sender's expected utility, calculated across a long sequence of generated objects, equals $\Pr[(c,a)^{(k)} = (L,H)] - \Pr[(c,a)^{(k)} = (H,L)]$ (i.e., the k th object produces a unit benefit in the case of a successful Fake VIP attack, and a unit loss if Relay has wrongly assigned class L). For convenience, it is rescaled using two baseline scenarios: *never-trust* ($a^{(k)} \equiv s^{(k)}$) and *always-trust*, ($a^{(k)} \equiv d^{(k)} \equiv H$), which yield Sender expected utilities $1 - \rho - \omega$ and $1 - \rho$, respectively:

$$Eu_{\text{Sender}} = \frac{Eu_{\text{Sender,rep}} - Eu_{\text{Sender,never-trust}}}{Eu_{\text{Sender,always-trust}} - Eu_{\text{Sender,never-trust}}}, \quad (4)$$

A little algebra yields $\Pr[a^{(k)} = H] = 1 - \omega(1 - Eu_{\text{Sender}})$ [2], hence the perception of the long-term frequency of $a^{(k)} = H$ permits Sender to observe (4). Relay's expected utility is

$$Eu_{\text{Relay}} = \beta(1 - \pi_R) + \Pr[a^{(k)} = H], \quad (5)$$

where $\pi_R = \Pr[r^{(k)} = R]$ and $\beta > 0$ is the relative signature detection cost. This cost-type utility reflects the fact that the k th object costs Relay a single signature detection when $r^{(k)} < R$, whereas $a^{(k)} = H$ implies high QoS provision effort at Server.

It is useful to compare (4) and (5) to reputation-free (RF) play, where Relay chooses the never-trust or always-trust scenario, whichever produces a smaller cost (5), i.e., the former when $\beta < \omega$ (to which Sender responds with any Fake VIP attack strategy), and the latter when $\beta \geq \omega$ (to which Sender responds with $\sigma^{(k)} \equiv \text{TRUE}$). The *modified* utilities are defined as the differences between (4) and (5), and the corresponding RF utilities, thus represent gains and losses compared to RF play.

C. Stackelberg Equilibrium

In a non-cooperative game between Sender and Relay, let S_{Sender} and S_{Relay} be Sender's and Relay's strategy spaces whose generic elements are respectively $\sigma(\cdot)$ and $\tau(\cdot) \equiv (R, \delta)$, and $\text{Eu}_{\text{Sender}}(\sigma(\cdot), \tau(\cdot))$ and $\text{Eu}_{\text{Relay}}(\sigma(\cdot), \tau(\cdot))$ given by (4) and (5) be the two players' utility functions. Since R has little impact upon the expected utilities [2], we assume $\tau(\cdot) \equiv \delta$. In a Stackelberg game framework [8], Relay (playing Leader) sets her optimum trust strategy anticipating a best-response Fake VIP strategy of a selfish Sender (playing Follower); in this way the players reach a Stackelberg equilibrium (SE):

$$\begin{aligned} \sigma^*(\cdot)|_{\delta} &= \arg \max_{\sigma(\cdot) \in S_{\text{Sender}}} \text{Eu}_{\text{Sender}}(\sigma(\cdot), \delta) \quad \forall \delta \in [0, 1], \\ \delta^* &\in \arg \min_{\delta \in [0, 1]} \text{Eu}_{\text{Relay}}(\sigma^*(\cdot)|_{\delta}, \delta). \end{aligned} \quad (6)$$

III. INTELLIGENT FAKE VIP ATTACKS

In a probabilistic Fake VIP attack strategy, denoted pFVIP [2], fixed probabilities σ_L and σ_H are defined with which intrinsic class L or H objects are attacked, i.e., $\sigma^{(k)} \equiv \text{rand}(\sigma_{\sigma^{(k)}})$ in (1) and so $\sigma(\cdot) \equiv (\sigma_L, \sigma_H)$. The proposed double-blind reputation scheme is then invariably beneficial for Relay (regardless of β) and produces a win-win at $\beta < \omega$, at $\beta > \omega$ Sender incurs losses compared to always-trust. Thus Relay's defense is effective against pFVIP, however, it needs to be tested against more intelligent Fake VIP attack strategies whereby Sender has some rudimentary idea of the workings of Relay's reputation scheme. We will refer to such a strategy *intelligent*, and denote it iFVIP.

iFVIP is assumed to be unaware of R , δ , current reputation state, or assigned class; however, it is aware of the rules (2). It launches an attack when currently suspecting the trust state, while refraining from attack when suspecting a non-trust state. We will say that Sender perceives herself as *trustworthy* in the former case, and as *untrustworthy* in the latter case. Sender's trustworthy and untrustworthy self-perception states will be denoted *TSP* and *USP*, respectively. Let $s\text{-}p^{(k)}$ be the self-perception state at the time of passing

the k th object to Relay, then the Fake VIP attack strategy defines $\sigma^{(k)} \equiv (s\text{-}p^{(k)} = \text{TSP})$ in (1). Uncertain about Relay's signature detection and reputation scheme details, Sender can use the following heuristic to develop a self-perception. For a given object, Sender looks at the predicted class $p^{(k)}$ and demanded class $d^{(k)}$ as dictated by (1). If $p^{(k)} = H$, the self-perception remains unchanged, since Sender then hopes $s^{(k)} = H$ is true, in which case, regardless of $d^{(k)}$, Relay does not lower the reputation state (except perhaps at the trust state). Therefore, demanding $d^{(k)} = H$ is *subjectively neutral* to Sender. If $p^{(k)} = L$ and the current self-perception is *TSP*, Sender demands $d^{(k)} = H$, which is *subjectively dishonest* (since $s^{(k)} = L$ can be surmised and a lowering of the reputation state can be anticipated) and prescribes that Sender downgrade her self-perception to *USP*. Conversely, a prediction $p^{(k)} = L$ under *USP* implies $d^{(k)} = L$, a *subjectively honest* behavior (since a raising of the reputation state can be anticipated) that prescribes that Sender upgrade her self-perception to *TSP*. Note that Sender cannot know if the lowering or raising of the reputation state indeed has taken place (being unaware of $s^{(k)}$ and because of the randomization controlled by the parameter δ of the reputation scheme). To reflect this lack of knowledge, whenever an upgrade or downgrade of self-perception is prescribed, the former is effected with probability φ and the latter with probability $1 - \varphi$, where the parameter φ is part of Sender's attack strategy (a large φ marks a strong drift toward *TSP* and fosters more frequent attacks).

A. Two-Dimensional Markov Chain under iFVIP

In this section we use a two-dimensional Markov chain to calculate Sender's and Relay's expected utilities for iFVIP and to compare the resulting characteristics with pFVIP described earlier. Fig. 3 depicts transitions between *TSP* and *USP*, and between Relay's reputation states $i = 1, \dots, R$ (self-loops are omitted); specified are also the events triggering these transitions. While the transitions at Relay are triggered as before, those at Sender are governed by $\text{rand}(\varphi)$ and $p^{(k)} = L$ according to the above description. To analyze the players' expected utilities, we combine both graphs in Fig. 3 into a two-dimensional Markov chain over the state space $\{\text{TSP}, \text{USP}\} \times \{1, \dots, R\}$, whose generic states have the form (j, i) , with $j \in \{\text{TSP}, \text{USP}\}$. That is, if $ev_{jj'}$ and $\overline{ev}_{i'}$ are the events triggering transitions from state j to state j' in the lower graph and from state i to state i' in the upper graph, respectively, then $ev_{jj'} \wedge \overline{ev}_{i'}$ triggers the transition from state (j, i) to state (j', i') in the combined graph (note that in general these two events are not independent). Of the two components of each state, the first is only known to Sender and the second to Relay.

Fig. 4 depicts the underlying state transition graph and transition probabilities; for clarity, self-loops are not drawn, and states (TSP, i) and (USP, i) are respectively labeled i and i' . Consider first transitions out of state i , which lead to states $i - 1$, i' and $(i - 1)'$; transition to $i + 1$ is impossible, since $d^{(k)} = L$ is impossible at *TSP*. The event $(s, p)^{(k)} = (L, H)$ causes $d^{(k)} = H$, a subjectively neutral behavior to Sender, hence not

prescribing a downgrade of self-perception. Yet at $i < R$, Relay detects a Fake VIP attack and is inclined to lower the reputation state (which happens if $\text{rand}(\delta)$ occurs). The event $(s,p)^{(k)} = (L,L)$ moreover prescribes a downgrade of self-perception (which happens if $\neg\text{rand}(\varphi)$ occurs), since $d^{(k)} = H$ is then subjectively dishonest. Transitions out of state i' lead to states $(i-1)'$, i and $i+1$. The first transition is analogous to the one out of state i . The other two are triggered by $p^{(k)} = L$, since the demand $d^{(k)} = L$ that follows is subjectively honest and prescribes that Sender upgrade her self-perception if $\text{rand}(\varphi)$ occurs, while (at $i < R$) it inclines Relay to raise the reputation state if $\neg\text{rand}(\delta)$ occurs. Transition to $(i+1)'$ is impossible, since $d^{(k)} = L$ and $p^{(k)} = L$ coincide at *USP*, hence raising the reputation state must be accompanied by raising the self-perception.

Analysis of the combined events in Fig. 3a and 3b yields the transition probabilities labeled (a) through (f) in Fig. 4. For example, transition (a) is triggered by $\neg(\neg\text{rand}(\varphi) \wedge p^{(k)} = L) \wedge \text{rand}(\delta) \wedge (s,d)^{(k)} = (L,H)$; since $d^{(k)} = H$ is certain, this is equivalent of $\text{rand}(\delta) \wedge ((s^{(k)} = L \wedge \text{rand}(\varphi)) \vee (s,p)^{(k)} = (L,H))$, hence the transition probability is $\delta\varphi\text{Pr}[s^{(k)} = L] + (1-\varphi)\text{Pr}[(s,p)^{(k)} = (L,H)]$. Transition (e) is triggered by $\text{rand}(\varphi) \wedge p^{(k)} = L \wedge \neg((\text{rand}(\delta) \wedge (s,d)^{(k)} = (L,H)) \vee (\neg\text{rand}(\delta) \wedge d^{(k)} = L))$; since $d^{(k)} = H$ and $p^{(k)} = H$ coincide, this simplifies to $\text{rand}(\delta) \wedge \text{rand}(\varphi) \wedge p^{(k)} = L$ and gives the transition probability $\delta\varphi\text{Pr}[p^{(k)} = L]$. The other transition probabilities are obtained similarly. Based on these probabilities, arranged into a matrix $\mathbf{T} = [t_{ij}]_{i,j \in \{TSP, USP\} \times \{1, \dots, R\}}$, a stationary probability distribution π over the state space can be derived in the usual way.

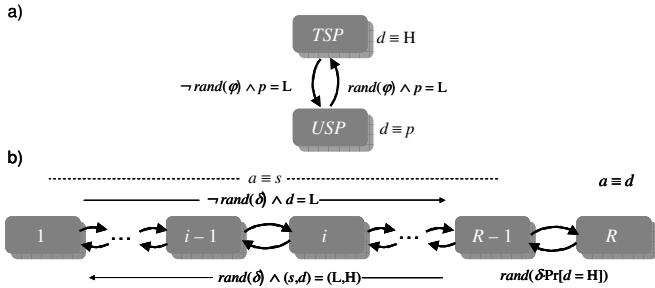


Fig. 3. Transitions of (a) Sender's self-perception, (b) reputation states.

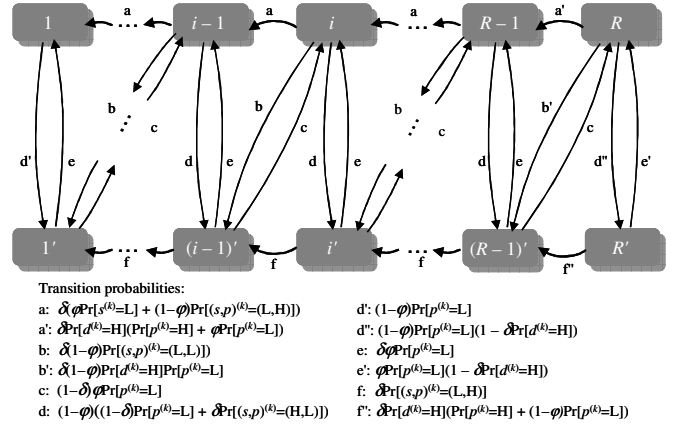


Fig. 4. Two-dimensional Markov chain for iFVIP.

Using our previous notation, we express the relevant probabilities in Fig. 4 as follows: $\text{Pr}[s^{(k)} = L] = \omega$, $\text{Pr}[p^{(k)} = L] = 1 - \text{Pr}[p^{(k)} = H] = \Omega$, $\text{Pr}[(s,p)^{(k)} = (L,H)] = \omega\zeta_L$, $\text{Pr}[(s,p)^{(k)} = (L,L)] = \omega(1 - \zeta_L)$, $\text{Pr}[(s,p)^{(k)} = (H,L)] = (1 - \omega)\zeta_H$. The probability $\text{Pr}[d^{(k)} = H]$, used by Relay to trigger transitions out of the trust states R and R' , is estimated based on demanded classes of objects arrived so far in non-trust states; this probability can be analytically calculated as $1 - \text{Pr}[p^{(k)} = L] \cdot \Pi_{USP}$,

where $\Pi_{USP} = \frac{\sum_{i=1}^{R-1} \pi_i}{\sum_{i=1}^{R-1} (\pi_i + \pi_{i'})}$. Thus \mathbf{T} depends on π , and π can be obtained by numerically solving $\pi = \pi\mathbf{T}(\pi)$ subject to $\sum_{i=1}^R (\pi_i + \pi_{i'}) = 1$. From the viewpoint of both play-

ers' expected utilities, of interest are the stationary trust state probabilities, π_R and $\pi_{R'}$, which quantify the vulnerability to Fake VIP attacks. Note that some values of δ and φ (notably $\delta, \varphi = 0$ or 1) yield nonergodic Markov chains, for which π_R and $\pi_{R'}$ should be stated separately depending on the initial state of the Markov chain; we omit the full analysis for lack of space.

B. Sender's and Relay's Expected Utilities

Under the reputation scheme, Sender's utility is

$$\begin{aligned}
 \bar{E}u_{\text{Sender, reput}} &= \text{Pr}[(c,a)^{(k)} = (L,H)] - \text{Pr}[(c,a)^{(k)} = (L,H)] \\
 &= \sum_{X,X',X''=H \text{ or } L} \text{Pr}[(c,s,p,d,a)^{(k)} = (L,X,X',X'',H)] \\
 &\quad - \sum_{X,X',X''=H \text{ or } L} \text{Pr}[(c,s,p,d,a)^{(k)} = (H,X,X',X'',L)].
 \end{aligned}
 \tag{7}$$

The calculation of the probabilities in (7) is straightforward, e.g., $\text{Pr}[(c,s,p,d,a)^{(k)} = (L,L,H,H,H)] = (1 - \rho)(1 - \varepsilon_L)\zeta_L(\pi_R + \pi_{R'})$, $\text{Pr}[(c,s,p,d,a)^{(k)} = (H,L,L,H,L)] = \rho\varepsilon_H(1 - \zeta_L)\sum_{i=1}^{R-1} \pi_i$, and so on. After some algebra, the utility (4) becomes

$$Eu_{\text{Sender}} = \pi_R + \left(1 - \frac{\Omega}{\omega}\right) \pi_{R'}.$$

(8)

Sender can deduce her utility (8) by observing $\Pr[a^{(k)} = H] = 1 - \omega(1 - Eu_{\text{Sender}})$. This also yields

$$Eu_{\text{Relay}} = \beta(1 - \pi_R - \pi_{R'}) + 1 - \omega(1 - Eu_{\text{Sender}}).$$

(9)

Regarding postulate (vi), note that $(s,a)^{(k)} = (H,L)$ requires a coincidence of *USP* at Sender, the trust state at Relay, and $(s,p)^{(k)} = (H,L)$; hence, $\Pr[(s,a)^{(k)} = (H,L)] = (1 - \omega)\zeta_H\pi_R$. E.g., assuming SE play, $\omega = 0.8$, $R = 5$, $\zeta_L = \zeta_H = 0.3$, and $\beta = 0.3$, this probability is 0 (since (6) then yields $\delta^* = 1$ and so $\pi_{R'} = 0$).

Fig. 5 compares the modified utilities (gains and losses with respect to RF play) under pFVIP and iFVIP for $R = 5$. iFVIP improves Sender's gains for a range of $\beta < \omega$ and reduces her losses for a range of $\beta > \omega$ to an extent dependent on ζ_L and ζ_H , while uniformly reducing Relay's gains for $\beta < \omega$ and a range of $\beta > \omega$. Hence, the little intelligence at Sender helps iFVIP to outperform pFVIP at times.

Compared to $\zeta_L = \zeta_H = 0$, the presence of $\zeta_L > 0$ is slightly beneficial for Sender and damaging for Relay, since Sender is more inclined to demand class H, whereas the presence of $\zeta_H > 0$ is not (cf. plots B and C vs. A and D). Raising R to 10 improves Relay's gains, but very slightly (within a few percent).

IV. KNOWLEDGEABLE FAKE VIP ATTACKS

iFVIP leverages some rudimentary knowledge about the events triggering reputation state transitions, though the parameters (R, δ) remain unknown to her. Would the knowledge of them help? Suppose that Sender is able to correctly estimate ω , ζ_L and ζ_H . Suppose also that she has used iFVIP for a while estimating her expected utility based on observed $\Pr[a^{(k)} = H]$. Having discovered the functional dependence of π_R and $\pi_{R'}$ on (R, δ) through the Markovian analysis, as well as the fact that these two probabilities are very weakly sensitive to R , she can infer δ and then fit R so that the calculated and observed $\Pr[a^{(k)} = H]$ match. This is a serious, but not unthinkable departure from the information-separation framework (the relevant parameters can also be obtained by hacking into Relay's private data). From Relay's perspective, it presents a worst case.




Fig. 5. Expected utility gains and losses under piFVIP and iFVIP; A: $\zeta_L = \zeta_H = 0$; B: $\zeta_L = 0.3$, $\zeta_H = 0$; C: $\zeta_L = \zeta_H = 0.3$; D: $\zeta_L = 0$, $\zeta_H = 0.3$.

In light of the above, the assumption underlying our worst-case analysis is that Sender can recreate Relay's reputation state transition graph in Fig. 3a. Hence, she can run the corresponding Markov chain in parallel with Relay, building a multilevel self-perception (rather than binary *USP* vs. *TSP* in iFVIP), and only attempting Fake VIP attacks when a trust state is perceived. We will call such a Fake VIP attack strategy *knowledgeable* and denote kFVIP. In what follows we use again a two-dimensional Markov chain to analyze kFVIP and compare the resulting utilities with iFVIP.

A. Two-Dimensional Markov Chain under kFVIP

We combine two one-dimensional Markov chains each with reputation state transitions depicted in Fig. 3a into a two-dimensional Markov chain over the state space $\{1, \dots, R\} \times \{1, \dots, R\}$. A generic state has the form (j, i) , where j represents Sender's self-perception and i represents the current reputation state as viewed by Relay. Again, if ev_{ji} and \overline{ev}_{ji} are the (in general not independent) events triggering transitions from state j to state j' in one graph and from state i to state i' in the other graph, respectively, then $ev_{ji'} \wedge \overline{ev}_{ji}$ triggers the transition from state (j, i) to state (j', i') in the combined graph. Since $rand(\delta)$ represents an independent random draw with probability δ , we take $rand(\delta) \wedge rand(\delta) = rand(\delta^2)$. As before, the first component of each state is only known to Sender and the second to Relay. Applying kFVIP to the k th generated object whose predicted class is $p^{(k)}$, Sender sets $d^{(k)} = p^{(k)}$ if her current self-perception $j^{(k)}$ is non-trust and $d^{(k)} = H$ otherwise.

That is, $\sigma^{(k)} \equiv (j^{(k)} = R)$ in (1). At the trust state, Relay estimates the probability $\Pr[d^{(k)} = H]$ based on demanded classes of objects arrived so far in non-trust states: $\Pr[d^{(k)} = H] = 1 - \Pr[p^{(k)} = L] \cdot \sum_{j=1}^{R-1} \sum_{i=1}^{R-1} \pi_{j,i} / \sum_{j=1}^R \sum_{i=1}^{R-1} \pi_{j,i}$.

Fig. 6 depicts the state transition graph and transition probabilities; for clarity, self-loops are not drawn. The probabilities $\Pr[s^{(k)} = L]$, $\Pr[p^{(k)} = L]$ and $\Pr[(s,p)^{(k)} = (L,H)]$ can be calculated as before. For $0 < \delta < 1$, the two-dimensional Markov chain is ergodic and the derivation of the stationary state probabilities $\pi_{j,i}$ follows in the same way as for iFVIP.

B. Sender's and Relay's Expected Utilities

Under the reputation scheme, Sender's and Relay's utilities can be found analogously to (8) and (9), yielding

$$Eu_{\text{Sender}} = \pi_{R,R} + \left(1 - \frac{\Omega}{\omega}\right) \sum_{j=1}^{R-1} \pi_{j,R}, \quad (11)$$

(11)

$$Eu_{\text{Relay}} = \beta(1 - \pi_{R,R}) + 1 - \omega(1 - Eu_{\text{Sender}}). \quad (12)$$

(12)

(note that in the no-uncertainty model, $Eu_{\text{Sender}} = \pi_{R,R}$). Further calculation Sender's and Relay's modified utilities (gains and losses compared to RF play). They are depicted in Fig. 7. It is assumed that Relay uses the equilibrium δ^* defined by (6) that she has configured against iFVIP.

Our previous findings remain valid, i.e., prediction errors can be beneficial for Sender and damaging for Relay if they induce more frequent Fake VIP attacks ($\zeta_L > 0$ and $\zeta_H = 0$), and increasing R does not affect the plots visibly. Comparison of iFVIP and kFVIP reveals that the knowledge of the reputation scheme parameters does not help Sender to improve her utility or reduce Relay's; in fact, the opposite is true. Owing to the randomization induced by the parameter δ , as well as her uncertainty as to the occurrence of $s^{(k)} = L$ and $(s,p)^{(k)} = (L,H)$, Sender's self-perception has only statistical chances to keep up with the current reputation state at Relay. Therefore kFVIP cannot properly coordinate Fake VIP attacks with the occurrences of the trust state. Here, iFVIP is more successful.




Fig. 6. Two-dimensional Markov chain for kFVIP.




Fig. 7. Expected utility gains and losses under iFVIP and kFVIP; A: $\zeta_L = \zeta_H = 0$; B: $\zeta_L = 0.3, \zeta_H = 0$; C: $\zeta_L = \zeta_H = 0.3$; D: $\zeta_L = 0, \zeta_H = 0.3$.

V. CONCLUSION

Regarding the central question of this paper, whether a double-blind reputation scheme makes an adequate defense

against more intelligent Fake VIP attack strategies than pFVIP, we offer the following findings:

- For a range of β , Sender may increase her SE utility beyond what she can achieve using pFVIP by adopting a more intelligent Fake VIP attack strategy. The conceived strategy we call iFVIP exploits some rudimentary idea of the workings of Relay's reputation scheme. Sender can then develop a self-perception of her trustworthiness to better choose an object on which to launch a Fake VIP attack. However, Sender's increased intelligence only very slightly diminishes Relay's utility.
- Sender's knowledge of the reputation scheme parameters at Relay, whether acquired by learning, by chance or illegally, seemingly gives her even more edge on Relay. However, a resulting strategy we call kFVIP turns out not to visibly improve Sender's utility or worsen Relay's. This is because the randomization of reputation state transitions prevents Sender from accurately coordinating her self-perception with the trust or non-trust states at Relay.

These findings naturally do not yield an ultimate answer to the above question. Rather, they stimulate further research into even more intelligent Fake VIP attack strategies and the resiliency of the double-blind reputation scheme to them. The reputation scheme should be systematically designed so that at SE it minimizes Relay's expected cost against an arbitrarily intelligent Sender's attack strategy (perhaps developed using advanced AI learning algorithms based on the observed sequence of $(c,p,d)^{(k)}$ and the $\Pr[a^{(k)} = H]$ feedback from Relay). How far such a strategy is from the presented heuristic ones remains an open question. On the modeling side, the model needs to be extended to correlated object generation processes, such as TCP, streamlined video, or online gaming traffic. Convergence to SE in dynamic game scenarios also needs to be studied. These issues are left for future work.

REFERENCES

- [1] V. Varadarajan and U. Tupakula, "Security as a service model for cloud environment," *IEEE Trans Network and Service Management*, vol. 11, pp. 60–75, 2014.
- [2] J. Konorski, "Fake VIP Attacks and their mitigation via double-blind reputation," *Proc. ITNAC 2017*, Melbourne, Australia.
- [3] S. Szott and J. Konorski, "Traffic remapping attacks in ad hoc networks," *IEEE Comm. Magazine*, Jan. 2018.
- [4] W. Li, A. Joshi, and T. Finin, "CAST: Context-aware security and trust framework for mobile ad-hoc networks using policies," *Distributed and Parallel Databases*, vol. 31, Issue 2, pp 353–376, June 2013.
- [5] P. L. Bartlett, Online Prediction. Lecture notes 2016 (online). Available at: stat.berkeley.edu/~bartlett/papers/b-ol-16.pdf.
- [6] Y. Freund, M. Kearns, Y. Mansour, D. Ron, R. Rubinfeld, and R. E. Schapire, "Efficient algorithms for learning to play repeated games against computationally bounded adversaries," *Proc. 36th Annual Symp. Foundations of Computer Science*, Milwaukee, WI, Nov. 1995.
- [7] A. Patcha and Jung-Min Park, "A game theoretic formulation for intrusion detection in mobile ad hoc networks," *Int. J. of Network Security*, vol. 2, 2, pp.131–137, March 2006.
- [8] E. Rasmussen, *Games and Information: An Introduction to Game Theory*, 3rd ed., Blackwell Publ. 200

