Many-player Inspection Games in Networked Environments

Gabriele Gianini, Ernesto Damiani
Dipartimento di Informatica
Universita degli Studi di Milano, Italy

Tobias R. Mayer, David Coquil, Harald Kosch
Chair of Distributed Information Systems
University of Passau, Germany

Lionel Brunie
Laboratoire LIRIS
INSA de Lyon, France

***Abstract*—In communication architectures, nodes are expected to spend their own resources so as to relay other nodes' messages or perform other services for the common good. However any selfish node, if given the opportunity, would typically prefer – to spare its own resources – to avoid serving the other nodes. This creates a potential problem to any collaborative protocol. A possible approach towards this issue consists in performing audits on the actions of the individual nodes, and applying some form of sanction to those whose misbehaviour has been detected during an inspection. However typically, auditing is costly and due to limited resources it can be carried on only on a sampling basis. It is clear that the rate of inspection has to be adapted to the rate of misbehavior, so as to strike a balance, from the point of view of the inspector, between the audit costs and the avoided damage to the system. Since the misbehaviour rate of rational agents is not predefined or fixed, but in turn depends from inspection rate, the overall behavior of the system made by inspectors and inspectees fits into a typical interdependent interaction landscape and can be modeled using Game Theory. The above described audit situation corresponds to a class of games known as Inspection Games. In this paper, we model several versions of Inspection Games (IGs), up to the most general case involving m inspectors and n inspectees. We resolve each game by computing the strategy that rational players would follow. Moreover, we also extend the IG model by taking into account the possibility of undetected violations, i.e. false negatives in the inspections.**

I. Introduction

Until recently, distributed systems were typically realized by means of the enactment of a collaborative protocol over networked nodes, being enriched by fault-tolerance mechanisms. In that approach, the protocol designer needs to find a balance in the trade-off between the degree of attainment of the network goals and the level of resource consumption as a whole.

That approach is not adequate to face the challenges of today's distributed systems nor, more specifically, to Digital Ecosystems, as a consequence of the fact that nodes are very likely to be operated by selfish parties. Selfish nodes could make use of their specific knowledge to perform undetected violations with regard to the protocol in order to follow personal interests. Such a selfishness in the collaborative interaction may affect the correct system operation (and is equivalent to an attack [1], [2]). Thus, Digital Ecosystems should not only consider fault-tolerance issues but also selfishness issues.

Securing the correct system operation against selfish behaviour could be realized by directly hardening the collaborative protocol. However, in order to take into account real world constraints, a modification of the protocol may not be possible or not wanted. Therefore, we assume the system designer's has the possibility of deploying timely ex post monitoring with the collaboration of some rational network participants.

Game Theory (GT) enables the modelling of such an interdependent decision landscape, where a system consists of selfish peers or players with non-aligned interests. By means of GT, we are able to predict the players' behaviour under specified circumstances and to calculate the solution of the game, i.e. the *Nash equilibrium*. That consists in a *strategy profile* (a collection of strategies, one for each player), from which no player has incentives to deviate unilaterally, since this would not increase its personal payoff.

The above drafted misbehaviour/auditing scenario, can be mapped to a specific class of GT models: *Inspection Games* (IGs). In this type of games an *inspector* controls the correct behaviour of an *inspectee*, and administers a sanction if a misbehaviour is detected during the inspection. GT modeling allows to find the rate of violation and the rate inspection at equilibrium as a function of the parameters of the problem, i.e. cost of individual inspections, quantitative damage inflicted by a violation, benefit to the violator and value of the sanction.

A designer or system administrator, by adjusting the parameters which are under his control – such as positive and negative incentives or further details of the inspection procedure – can influence the equilibrium. Notice that the task of choosing the target of the audit and performing the audit can be either assigned by the designer to trusted parties or to other selfish nodes in exchange for a reward for the cases of misbehavior detection.

By GT modeling one can find the equilibria of a rich spectrum of configurations from simple ones in which every potentially misbehaving agent has a dedicated inspector, to others where a single inspector is associated to several inspectees (inspectors coordinated so as to partition the pool of inspectees), to further ones where many uncoordinated inspectors inspect many inspectees (an inspectee can undergo inspection by several inspectors, who can in turn audit also other inspectees).

Thanks to the results of the GT modeling the system administrator can assume the role of a game designer: he will be able to tune the game parameters so as to shift the equilibrium to a desired strategy profile; the parameters include the different inspectors-inspectees cardinalities.

978-1-4799-0786-1/13/$31.00 ©2013 IEEE

In GT the IG model has been discussed under various forms; the contribution of this work consists in the generalization to m inspectors and to the introduction of the possibility of undetected violations, i.e. false negatives in the inspections, so as to suit more realistic scenarios in networking environments.

The remainder of this paper is structured as follows. In Section II we outline at first a basic two-player Inspection Game. Then, in Section III, this game is generalized step by step to a game with m inspectors and n inspectees, and the corresponding Nash equilibria are provided. In Section IV the games are extended with the possibility of false negatives. A short discussion of the related work concludes the paper.

II. Inspection Games definition

Game Theory (GT) is a branch of applied mathematics that models multi-person decision-making situations in order to account for interactions among strategies of rational decision makers. It is principally aimed at determining the preferred combination of strategies that will be adopted by rational agents trying to maximize their payoffs.

A game is defined by a set of players, and, for each player, a set of possible strategies and a player's utility function – mapping any possible state of affairs in the game into a payoff for the player. A *strategy* for a player is a complete plan of actions in all possible situations throughout the game, the goal of every player consists in adopting the strategies maximizing his own payoff, by taking into account that they depend, through the state of affairs, also upon the other players' chosen strategy.

A *Nash equilibrium* is a solution that describes a steady state condition of the game; it corresponds to a combination of strategies (a *strategy profile*) such that no individual player would be better off by changing his own strategy unilaterally.

In this context, an *Inspection Game* represents a specific class of games. In one of its simplest forms – the two-player simultaneous single-round Inspection Game – the set of players consists in $\{Inspector\,,\,Inspectee\}$; since the game is single-round, each player chooses only once a strategy for the game, which is done without the knowledge of the other player's chosen strategy (simultaneous): the inspector can choose between inspecting or not inspecting, i.e. the set of inspector's choices is $\{Inspect\,,\,Do\ not\ inspect\}$, similarly, the inspectee can choose between violating or not violating, i.e. an operation of the set $\{Violate\,,\,Do\ not\ violate\}$. This determines four possible states of affairs.

Hereafter, we will indicate the number of inspectees by n, the number of inspectors by m: an Inspection Game with n inspectees and m inspectors will be indicated by $G(m, n)$. In m inspectors n inspectees games, each time the game is played an inspectee has the choice between violating or not (the collaborative protocol of a networked architecture). An undetected violation will bring him a benefit, at the same time an inspector has the choice between performing or not an inspection: if he does and finds evidence of the violation then the inspectee receives some form of sanction. However, the inspection, whatever the inspection findings, has a cost.

TABLE I
The payoff matrix for a two-player Inspection Game shows possible game states and payoffs.

		Inspector	
		Inspect	*Do not insp.*
Inspectee	*Violate*	$(b-a\,,\,-c)$	$(b\,,\,-d)$
	Do not viol.	$(0\,,\,-c)$	$(0\,,\,0)$

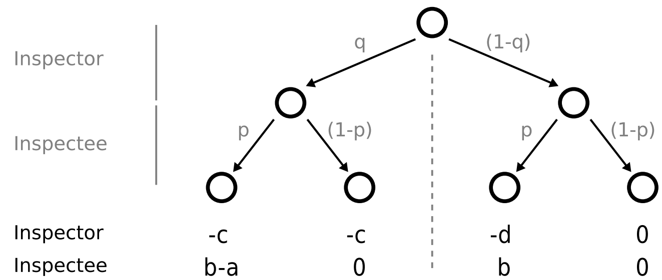

Fig. 1. Inspection Game $G(1, 1)$ for one inspector and one inspectee in extensive form.

For all the versions of the game the following rules for the payoffs will hold:

(i) The state of affairs without violation and without inspection does not bring any damage nor benefit to any player.
(ii) Violation will bring the inspectee a positive benefit b if not detected, but, if detected, it will bring him also a loss $-a$ with $|a| > |b|$.
(iii) The inspection has a fixed cost $-c$ for the inspector, but not detecting a violation would cost him a damage $-d$ with $|d| > |c|$.

Notice that the player's preferences are determined by differences in payoffs; hence, the addition of a constant to the utility function does not influence the solution of the game.

With this payoff schema, the preferences have a circular structure: each inspectee would prefer to violate when not inspected, and each of inspectors would prefer to be inspecting when there is a violation. If the strategy choices have to be taken by the parties simultaneously (or equivalently if they do not have any hint about the other party's move before their own move), the parties cannot determine in advance which one is their own best pure strategy, and they will have to resort to a suitable randomization between the two choices. In other words there is no Nash equilibrium in pure strategies: all the pure strategy profiles have at least one player which would benefit from switching strategy unilaterally.

As a consequence each party will have to adopt a *mixed strategy* (defined by a probability distribution over the pure strategies) fulfilling the so called *indifference condition*: the mixed strategy of one player must be a mixing of his own pure strategies such that the other player's expected payoff will not change whatever mix of his own pure strategy is adopted. This joint mixed strategy (the mixed strategy profile) will represent the Nash equilibrium of the game.

For instance for the two-player Inspection Game, the individual mixed strategy is represented by probability values: an inspector chooses an *inspection probability* q and the inspectee

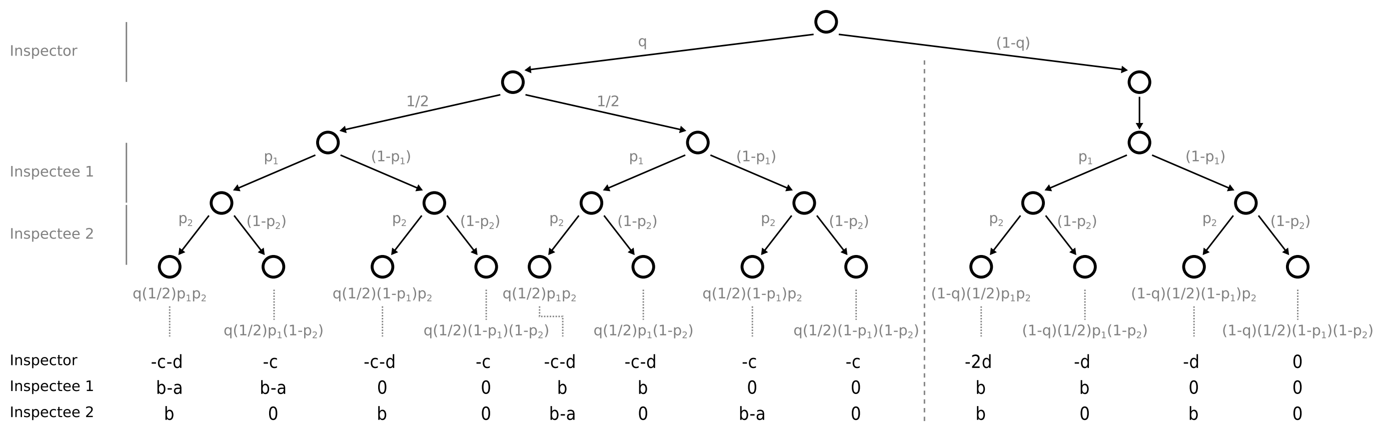

Fig. 2. Inspection Game $G(1,2)$ in extensive form.

a *violation probability* p. Suppose that the inspectee adopts the violation choice with probability p, and that the inspector adopts the inspection choice with probability q. Then, the solution of the game can be found by computing the pair (p,q) such that neither the inspectee can improve his expected payoff by deviating from p, nor the inspector can improve his expected payoff by deviating from q. In the remainder of this work, the indifference strategies will denoted by q^* and p^* respectively, the solution of a game by (p^*, q^*).

III. From a Two-Player Game to an Inspection Games with Several Players

Now, we will detail the basic two-player simultaneous single-round Inspection Game $G(1,1)$ outlined in the section before more formally, generalize it up to a $G(m,n)$ game and provide solutions for all game types.

A. Game $G(1,1)$ - One Inspector, One Inspectee

1) Game Setup: Table I represents the two-player Inspection Game in the so called Normal Form: the rows correspond to the possible moves of the inspectee (the inspectee's pure strategies), the columns correspond to the possible moves of the inspector (the inspector's pure strategies). The cells represent the four possible states of affairs and the corresponding payoffs for the players: in each cell, the pair (x,y) means that from that state of affairs the first player (the inspectee) obtains a total payoff x, while the second (the inspector) obtains a total payoff y. The structure of the game can also be represented in extensive form as shown in Fig. 1. Since $b > 0$ while $0 > (b-a)$ the inspectee will prefer to violate when the other does not inspect and will prefer not to violate when the other inspects. Conversely, since $-c > -d$, the inspector will prefer to inspect when the other violates and not to inspect when the other does not violate.

2) Game Solution: If the inspectee wants to induce the indifference in the inspector, he will have to set his own parameter p so as to equalize the expected inspector's payoff for an inspection to the inspector's payoff for lack of inspection. Similarly, if the inspector wants to induce the indifference in the inspectee, he will have to set his own parameter q so as to equalize the expected inspectee's payoff for a violation to the inspectee's payoff for lack of violation. Altogether $(-c) = p(-d)$ and $q(b-a) + (1-q)b = 0$, from which we get the simple solutions $p^* = \frac{c}{d}$ and $q^* = \frac{b}{a}$. Notice that, by construction, q^* is determined by the quantities defining the payoffs of the inspectee and that in the expression, as expected, the benefit b for an undetected violation, at the numerator, compete with the loss a for the detected one. Similarly, p^* is determined by the quantities defining the payoffs of the inspector and the cost for an inspection plays the opposite role to the avoided damage d.

B. Game $G(1,n)$ – One Inspector, n Inspectees

1) Game Setup: Let us consider an $(n+1)$-player simultaneous single-round Inspection Game, with one inspector and n inspectees, with the same payoff assumptions *(i),(ii),(iii)* as above. Additionally we assume that each individual violation causes a distinct damage (to the inspector) so that there is a maximum damage of nd. If the inspector decides for the inspection, then the inspection will be performed on a single randomly chosen inspectee: given the inspection, each inspectee will have probability $\frac{1}{n}$ to be inspected. The inspectees, from now on inspectee $1, \cdots$, inspectee n, have respectively probability $p_1 \cdots p_n$ of violating the rule. The solution of this game is represented by the values of q^*, $p_1^*, \cdots p_n^*$ of the above $(n+1)$ parameters at the Nash equilibrium. The tree diagram in Fig. 2 shows the different game result possibilities for $n = 2$.

2) Game Solution: An important point to observe here is that there is no coupling between inspectees: the payoff of one inspectee does not depend on the other inspectee's choices. *The inspectees are indifferent to the strategies of one another*.

Inspector's Indifference: For symmetry between the inspectees, their individual parameters will correspond to the same value, that we call p^* (i.e. $p^* = p_1^* = \cdots = p_n^*$): the value of p^* will have to satisfy a simple equation, equalizing the expected value (value times probability) on the inspector for undetected violations due to lack of inspection, with the balance between the expected value of an unfruitful inspection and the one of a fully or partially successful inspection. The expected value for no-inspection is given by the expected number of the violations of n inspectees times the damage d created by each one: i.e. by $np \times d$. The expected value for inspection is given by the constant cost c plus the expected value of the violations which have gone undetected. Since in this case the inspector is securing with certainty only one inspectee, the expected value of the undetected violations is given by the expected number of violations p of the remaining inspectees. Hence the following expected value for the inspection $c+(n-1)pd$. The resulting indifference equation is $npd = c+d(n-1)p$ hence $p^* = \frac{c}{d}$ Notice that the optimal p is the same as the one for a single inspectee: *the presence of further inspectees does not change the best strategy of one inspectee*. This is is a natural consequence of the lack of coupling between inspectees.

Inspectees' Indifference: Looking at the structure of the game one can observe that, since inspectees are not coupled to one another by the game's payoffs, they consider the inspection to another inspectee as equivalent to no inspection at all. Hence, in order to make each inspectee indifferent, the inspector has to behave as if each of them were playing against him an effective two-player one-inspector-one-inspectee $G(1,1)$ game with rescaled parameters. We can describe this effective game by introducing an effective probability of inspection $q_{eff} = \frac{q}{n}$. The extensive form two player effective game is the same as the one shown in Fig. 2 except that the probability q is substituted by $q_{eff} = \frac{q}{n}$. The inspectee's indifference is obtained equalizing the expected value for non violation, which is null, to the expected value of violation, given by the balance between the detected one and the undetected one. In case of violation there will always be a benefit for the inspectee, so the expected value is given by b added to the expected value of the loss (loss times probability of inspection $q_{eff} = \frac{q}{n}$). The indifference equation is $b - a\frac{q}{n} = 0$ hence $q^* = \frac{b}{\frac{a}{n}}$ The factor $\frac{1}{n}$ results from the fact that one inspector is shared by n inspectees.

C. Game G(m,1) – m Inspectors, One Inspectee

1) Game Setup: Beside the assumptions *(i),(ii),(iii)* already adopted, we are forced also to postulate that inspectors *share the damage* of any occurring violation which goes undetected (i.e. detected by none): this introduces some coupling among inspectors. The game is represented in Fig. 3 in extensive form for the case of $m = 2$ inspectors, $n = 1$ inspectee.

2) Game Solution: We will exploit the symmetry between the inspectors, since we know that at the equilibrium $q_1^* = \cdots = q_m^* = q^*$.

Inspectee's Indifference: The equation for the inspectee's indifference should equalize the expected value for no violation, which is null, to the expected value for violation. This in turn is given by the balance between the expected value of detection and that of non detection: since the benefit for violation is always present, be the violation detected or not, the balance is obtained by subtracting from b the expected value of

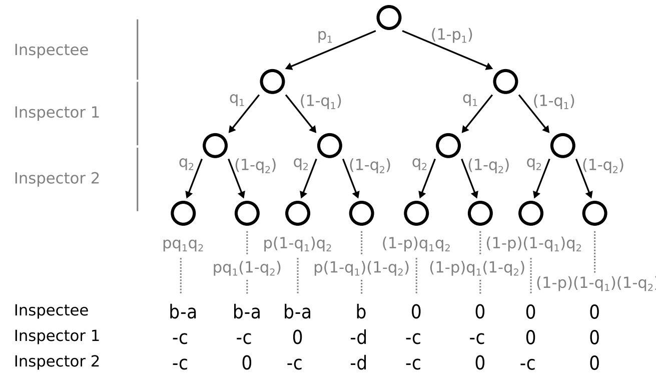

Fig. 3. Inspection Game $G(2,1)$ in extensive form. None of the eight possible strategy profiles can represent a pure strategy Nash equilibrium: in each of the above column at least one of the players could improved his payoff by unilaterally deviating from the pure strategy.

loss only (probability times value of loss). The probability of detection by at least one of the inspectors is $1-(1-q)^m$, hence the overall indifference equation is $b - a(1-(1-q)^m) = 0$ which has solution for q^* such that $(1-q^*)^m = 1 - \frac{b}{a}$ or $q^* = 1 - (1-\frac{b}{a})^{\frac{1}{m}}$

Inspectors' Indifference: The inspectors' indifference equation should equalize the expected value of inspection, which is given by a constant cost, to the expected value of no inspection. The latter corresponds to the expected value of the number of violations by the only inspectee when no other inspector is inspecting. Hence, the indifference equation is $c = dp(1-q)^{m-1}$ which has solution for $p^* = \frac{\frac{c}{d}}{(1-q^*)^{m-1}}$ or explicitly – taking into account that at the equilibrium value $(1-q^*) = (1-\frac{b}{a})^{\frac{1}{m}}$ – for $p^* = \frac{\frac{c}{d}}{(1-\frac{b}{a})^{\frac{m-1}{m}}}$

D. Game G(m,n) - m Inspectors, n Inspectees

In the game with m uncoordinated inspectors and n (non interacting) inspectees, the presence of n inspectees reduces the probability of any inspector visiting the i-th inspectee from q_i to $q_i^{eff} = \frac{q_i}{n}$. Hereafter, exploiting the symmetry among inspectors, we will use q in place of q_i and exploiting the symmetry among inspectees, we will use p in place of p_i.

Inspectees' Indifference: The indifference equation for each inspectee, which is used here to determine q, is $b-a(1-(1-\frac{q}{n})^m) = 0$ which has solution for $(1-\frac{q^*}{n})^m = 1 - \frac{b}{a}$ or $q^* = n\big(1-(1-\frac{b}{a})^{\frac{1}{m}}\big)$ It is as in the previous case, that of the game $G(m,1)$ except that q is replaced by the effective $q_{eff} = \frac{q}{n}$.

Inspectors' Indifference: The inspectors' indifference equation which is used here to determine p, should equalize the expected value of no inspection – this corresponds to d times the expected value of the number n of inspectees' violations going undetected by the other (m-1) inspectors – with the expected value of inspection (on a single inspector) – this is given by a constant cost plus the individual damage times the expected value of the number (n-1) of inspectees' violations going undetected by the other (m-1) inspectors. In both cases, the answer depends on the expected number of undetected violations when each inspectees violates the rule with

TABLE II
SUMMARY OF THE RESULTS FOR THE INSPECTION GAMES. $G(m,n)$ INDICATES AN INSPECTION GAME WITH m INSPECTORS AND n INSPECTEES; p^* INDICATES THE EQUILIBRIUM VIOLATION PROBABILITY FOR THE INSPECTEE, q^* INDICATES THE EQUILIBRIUM INSPECTION PROBABILITY OF THE INSPECTOR.

	p^*	q^*
$G(1,1)$	$P \equiv \frac{c}{d}$	$Q \equiv \frac{b}{a}$
$G(1,n)$	P	nQ
$G(m,1)$	$\frac{P}{(1-Q)^{\frac{m-1}{m}}}$	$1-(1-Q)^{\frac{1}{m}}$
$G(m,n)$	$\frac{P}{(1-Q)^{\frac{m-1}{m}}}$	$n\big(1-(1-Q)^{\frac{1}{m}}\big)$

probability p and each inspector performs an inspection with probability q – a quantity which we can call $u(n,p,m,q)$. The indifference equation will equate the following two expected values $d\,u(m-1,q,n,p) = c + d\,u((m-1),q,(n-1),p)$ which can be rearranged so that $\frac{c}{d} = u(n,p,(m-1),q) - u((n-1),p,(m-1),q)$ The difference at the second member represents the expected number of *extra* undetected violations, which occur when an inspector does not inspect. The missing inspection does not produce any extra undetected violations if the peer, which would be inspected, does not violate the rule, or if that peer is already inspected by at least one of the other inspectors. In other words, the missing inspection leaves one extra inspectee violating the rule undetected only when that inspectee does perform the violation and the other $(m-1)$ inspectors do not detect it: the former event happens with probability p and the latter with probability $(1-\frac{q}{n})^{m-1}$ (since each inspector has probability $\frac{q}{n}$ of falling over that inspectee). Hence the indifference equation reads $\frac{c}{d} = p(1-\frac{q}{n})^{m-1}$ and has solution for $p^* = \frac{\frac{c}{d}}{(1-\frac{q}{n})^{m-1}}$. Overall, substituting q^*, we have $p^* = \frac{\frac{c}{d}}{(1-\frac{b}{a})^{\frac{m-1}{m}}}$

The results are summarized in Table II. Notice that the p^* of the various $G(.,n)$ is equal to that of the corresponding $G(.,1)$: adding or removing inspectees does not change the p^* because there is no coupling between inspectees. On the contrary, the q^* of the various $G(.,n)$ is n times larger than that of the corresponding $G(.,1)$: multiplying the inspectees' number by n does change q^* because it requires a proportional increase in the inspectors' effort. Notice as well that both p^* and q^* of the $G(m,.)$ are reduced with respect to the corresponding $G(1,.)$: this is coherent with an increased and joint inspectors' pressure.

IV. INSPECTION GAMES WITH FALSE NEGATIVES

Notice that due to the limited resources available in general to inspection mechanisms there exist the possibility that a violation, which has occurred, goes undetected (the possibility of false negatives). False positives are still not allowed: when an inspection detects a violation, there is no doubt that the violation has actually occurred. Let us indicate by γ the probability that an inspection does detect a violation which has actually occurred. The enriched version of the game, including the possibility of false negatives, in the case of one inspectee and one inspector $G(1,1)$ is shown in Fig. 4. Similar extensions can be devised for the other $G(1,n)$, $G(m,1)$ and $G(m,n)$: we will indicate the corresponding games with false negatives by $\Gamma(\cdot,\cdot)$. Their Nash equilibria can be found by straightforward considerations. A first key observation for the development of the more general cases $G(\cdot,\cdot)$ concerns the inspectee's indifference equation used by the inspector to determine q^*: whenever an inspector sets the probability of inspection to the value q, the inspectee, due to false negatives (which corresponds to an inefficiency), perceives an effective probability γq; due to this fact, wherever there was a q in the equations for the $G(\cdot,\cdot)$ there is a γq in the equations for the $\Gamma(\cdot,\cdot)$. Therefore, the equilibrium values q^* for the inspectors in *all* the games $G(\cdot,\cdot)$ will be rescaled by a factor $1/\gamma$. For this reason, in order to find the equilibria for the $\Gamma(\cdot,\cdot)$'s we need to discuss in detail only the inspector's indifference condition, used to determine p^*.

A. Game $\Gamma(1,1)$ – One Inspector, One Inspectee

The equilibrium equation for the inspector in $\Gamma(1,1)$ changes slightly with respect to $G(1,1)$: the payoff for the inspection is not simply $(-c)$, but is decremented by the term $(-d)(1-\gamma)p$, due to possible inspection failure. The overall indifference equation is thus $-c + (-d)(1-\gamma)p = (-d)p$, or $c = pd\gamma$, which has the following solution $p^* = \frac{c}{\gamma d}$., valid for $\gamma d \geq c$. As anticipated above, the solution value for q is instead $q^* = \frac{b}{\gamma a}$. Notice that the two solution values p^* and q^* are equal to the solution values for $G(1,1)$ rescaled by a factor $1/\gamma$, which represents an increased violation rate and a correspondingly increased inspection rate.

B. Game $\Gamma(1,n)$ – One Inspector, n Inspectees

In the indifference equation for the inspector, used to determine p^*, in the game $\Gamma(1,n)$ we have the non-inspection side, npd, of the equality – representing the expected value of the damage from a set of n independent inspectee choosing to violate with probability p – and the inspection side, consisting on the terms also present in $G(1,n)$, i.e. $c + d(n-1)p$, plus the failed inspection term $(1-\gamma)pd$, hence $pd = c + (1-\gamma)pd$, which is equivalent to $\gamma pd = c$ and gives the solution $p^* = \frac{c}{\gamma d}$ As anticipated above the solution value for q is instead $q^* = n\frac{b}{\gamma a}$. The two values p^* and q^* are equal to the solution values for $G(1,n)$ rescaled by a factor $1/\gamma$.

C. Game $\Gamma(m,1)$ – m Inspectors, One Inspectee

As anticipated above the solution value for q is such that $(1-\gamma q^*)^m = 1 - \frac{b}{a}$ and its explicit form can be found in Table III. As for p, the inspector's indifference equation should equalize the expected value of inspection, to the expected value of no inspection. The latter term corresponds to the expected value of the violation (probability times expected value) by the only inspectee, when no other inspector perform as successful inspection, i.e. is $(-d)p(1-\gamma q)^{m-1}$, the successful inspection has probability γq.

The former term is given by the constant cost $(-c)$ plus the inspection failure expected expected value: in case of

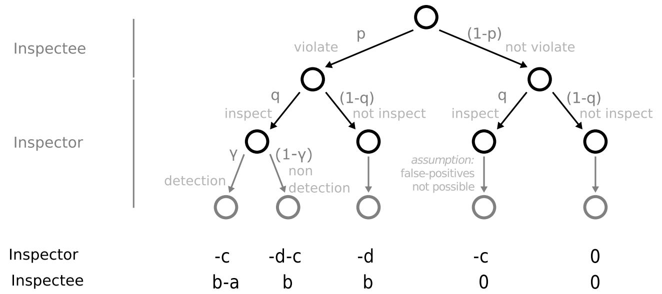

Fig. 4. Inspection Game, extended by false negatives, in extensive form with payoffs for inspector and inspectee.

inspection there is a damage $(-d)$ only if the inspectee has performed the violation (which happens with probability p) and this inspection fails (which happens with probability $1-\gamma$) while at the same time no other inspector has performed a successful inspection (probability $(1-\gamma q)^{m-1}$). Hence, the indifference equation is $c + dp(1-\gamma)(1-\gamma q)^{m-1} = dp(1-\gamma q)^{(m-1)}$ with solution $p^* = \frac{c}{\gamma d(1-\gamma q)^{(m-1)}}$. In terms of the solution q^* $p^* = \frac{c}{\gamma d(1-\frac{b}{a})^{(m-1)/m}}$. The two values p^* and q^* are equal to the solution values for $G(m, 1)$ rescaled by a factor $1/\gamma$.

TABLE III
SUMMARY OF THE RESULTS FOR THE INSPECTION GAMES $\Gamma(\cdot, \cdot)$

	$\gamma\, p^*$	$\gamma\, q^*$
$\Gamma(1, 1)$	$P \equiv \frac{c}{d}$	$Q \equiv \frac{b}{a}$
$\Gamma(1, n)$	P	nQ
$\Gamma(m, 1)$	$\frac{P}{(1-Q)^{\frac{m-1}{m}}}$	$1-(1-Q)^{\frac{1}{m}}$
$\Gamma(m, n)$	$\frac{P}{(1-Q)^{\frac{m-1}{m}}}$	$n\big(1-(1-Q)^{\frac{1}{m}}\big)$

D. Game $\Gamma(m, n)$ – m Inspectors, n Inspectees

As before, the solution for q yielded by the inspectee indifference equation is such that $\left(1-\gamma\frac{q^*}{n}\right)^m = 1-\frac{b}{a}$ i.e. as in the case $\Gamma(m, 1)$ but with q replaced by q/n, the explicit form is shown in Table III. The inspectors' indifference equation to determine p, should equalize the expected expected value (on a single inspector) of no inspection (this corresponds to d, times the expected value of the number n of inspectees' violations going undetected by the other (m-1) inspectors) with the expected expected value (on a single inspector) of inspection. Following (but not retracing) the derivation used in the game $G(m, n)$ we observe that after some manipulations, one will have to equate the value $(-c)$ of the certain cost for an inspection to the expected expected value of the extra detected violation (damage $(-d)$ times the violation probability p times successful detection probability γ of a violation undetected by the other inspectors). Hence the equation $c = d\,p\,\gamma\,(1-\gamma\frac{q}{n})^{m-1}$ or $\frac{c}{\gamma d} = p\,(1-\gamma\frac{q}{n})^{m-1}$ and has solution for $p^* = \frac{\frac{c}{\gamma d}}{(1-\gamma\frac{q}{n})^{m-1}}$ Overall, substituting q^*, we have $p^* = \frac{\frac{c}{\gamma d}}{(1-\frac{b}{a})^{\frac{m-1}{m}}}$ Again, the solution values p^* and q^* are equal to those of $G(m, 1)$ rescaled by a factor $1/\gamma$.

The results for the games $\Gamma(\cdot, \cdot)$ are summarized in Table III, they are equal to the solution values for the corresponding $G(\cdot, \cdot)$ rescaled by a factor $1/\gamma$, which represents an increased violation rate and a correspondingly increased inspection rate.

V. DISCUSSION AND CONCLUSIONS

In this paper, we introduced the Inspection Games and their solutions (Nash equilibria) in the context of an applications to networked architectures. To this end, we started with a two-player Inspection Game $G(1, 1)$ and developed the generalized versions $G(1, n)$, $G(m, 1)$ and $G(m, n)$. Then we adapted them to the modeling needs of communication architectures by means of versions $\Gamma(\cdot, \cdot)$, taking false negatives (not detected violation during an inspection) into account. Inspection Games have been defined by Dresher [3], [4], and have undergone several developments (see Kolokoltsov [5] for a very recent survey). There were already some works devoted to the case with one inspector and several inspectees. Already Avenhaus and Kilgour [6] have studied a three-person non-zero-sum game with one inspector and two inspectees in a setting richer than those considered here (there, the probability of detecting the inspectee's illegal action is a given function of inspection effort). Hohzaki [7] provides a generalization to the case of n inspectees, to the complex case where they are characterized by different attributes and as such may belong to different categories (e.g. countries) and studies how to optimally partition the effort. In our simpler case the effort cannot be partitioned and the probability of detection is not function of detection. The main contribution of this paper lies in the modelization of systems with m selfish inspectors and in the addition of the possibility of false negatives. Among the possible future developments are the following: multi-round inspection games with memory and changing payoffs; various spatial settings (static lattices, evolving random graphs, mobility…), bounded rationality behaviors and evolutionary versions of the game.

Acknowledgments. This work was partly supported by the EU-funded projects ASSERT4SOA (contract n. FP7-257351) and CUMULUS (c.n. FP7-318580), by the italian MIUR project SecurityHorizons (c.n. 2010XSEMLC) and by the UFI (Université Franco-Italienne) prog. Vinci2011 project C4-9. One of the authors (GG) was supported by the CNRS c.n.30022/2011.

REFERENCES

[1] S. Nielson, S. Crosby, and D. Wallach, "A taxonomy of rational attacks," in *Proc. 4th Int. Workshop on P2P Systems (IPTPS'05)*, 2005, pp. 36–46.

[2] T. Mayer, L. Brunie, D. Coquil, and H. Kosch, "Evaluating the robustness of publish/subscribe systems," in *Proc. VI Int. Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC'11)*, 2011.

[3] M. Dresher, "A sampling inspection problem in arms control agreements," Memorandum RM-2972-ARPA, The RAND Corporation, 1962.

[4] G. Owen, *Game theory*. Academic Press, 1995. [Online]. Available: http://books.google.it/books?id=psLDQgAACAAJ

[5] V. Kolokoltsov, H. Passi, and W. Yang, "Inspection and crime prevention: an evolutionary perspective," *ArXiv/e-prints/math.OC/1306.4219*, 2013.

[6] R. Avenhaus and D. Kilgour, "Efficient distributions of arms-control inspection effort," *Nav.Res.Logis.(NRL)*, vol. 51, no. 1, pp. 1–27, 2004.

[7] R. Hohzaki, "An inspection game with multiple inspectees," *European Journal of Operational Research*, vol. 178, no. 3, pp. 894 – 906, 2007.