Poster: Regret Minimizing Audits

Jeremiah Blocki
Student
Carnegie Mellon University

Nicolas Christin
Faculty
Carnegie Mellon University

Anupam Datta
Faculty
Carnegie Mellon University

Arunesh Sinha
Student
Carnegie Mellon University

ABSTRACT

Audits complement access control and are essential for enforcing privacy and security policies in many situations. The importance of audit as an *a posteriori* enforcement mechanism has been recognized in the computer security literature. For example, Lampson [1] takes the position that audit logs that record relevant evidence during system execution can be used to detect violations of policy, establish accountability and punish the violators. More recently, Weitzner et al. [2] also recognize the importance of audit and accountability, and the inadequacy of preventive access control mechanisms as the sole basis for privacy protection in today's open information environment. However, unlike access control, which has been the subject of a significant body of foundational work, there is comparatively little work on the foundations of audit.

Our focus is on policies that cannot be mechanically enforced in their entirety. Privacy regulations, such as the HIPAA for electronic medical record, provide one set of relevant policies of this form. For example, HIPAA allows transmission of protected health information about an individual from a hospital to a law enforcement agency if the hospital believes that the death of the individual was suspicious. Such beliefs cannot, in general, be checked mechanically either at the time of transmission or in an *a posteriori* audit; the checking process requires human auditors to inspect evidence recorded on audit logs.

In practice, organizations like hospitals use ad hoc audits in conjunction with access control mechanisms to protect patient privacy. Typically, the access control policies are quite permissive: all employees who might need patient information to perform activities related to treatment, payment and operations may be granted access to patient records. These permissive policies are necessary to ensure that no legitimate access request is ever denied, as denying such requests could have adverse consequences on the quality of patient care. Unfortunately, a permissive access control regime opens up the possibility of records being inappropriately accessed and transmitted. Audit mechanisms help detect such violations of policy. This is achieved by recording accesses made by employees in an audit log that is then examined by human auditors to determine whether accesses and transmissions were appropriate and to hold individuals accountable for violating policy. Recent studies reveal that many policy violations occur in the real world as employees inappropriately access records of celebrities and family members motivated by general curiosity, financial gain and other considerations [3]. Thus, there is a pressing need to develop audit mechanisms with well understood properties that effectively detect policy violations.

This work presents the first principled learning-theoretic foundation for audits of this form. Our first contribution is a **game-theoretic model** that captures the interaction between the defender (e.g., hospital auditors) and the adversary (e.g., hospital employees). The model takes pragmatic considerations into account, in particular, the periodic nature of audits, a budget that constrains the number of actions that the defender can inspect thus reflecting the imperfect nature of audit-based enforcement, and a loss function that captures the economic impact of detected and missed violations on the organization. We assume that the adversary is worst-case as is standard in other areas of computer security. We also formulate a desirable property of the audit mechanism in this model based on the concept of *regret* in learning theory [4]. Our second contribution is a novel **audit mechanism** that provably minimizes regret for the defender. The mechanism learns from experience and provides operational guidance to the human auditor about which and how many of the accesses to inspect. The regret bound is significantly better than prior results in the learning literature.

Overview of Results

Mirroring the periodic nature of audits in practice, we use a repeated game model [5] that proceeds in rounds. A round represents an audit cycle and, depending on the application scenario, could be a day, a week or even a quarter.

Adversary model: In each round, the adversary performs a set of actions (e.g., access patient records) of which a subset violates policy. Actions are classified into types. For example, accessing celebrity records could be a different type of action from accessing non-celebrity records. The adversary capabilities are defined by parameters that impose upper bounds on the number of actions of each type that she can perform in any round. We place no additional restrictions on the adversary's behavior. In particular, we do not assume that the adversary violates policy following a fixed probability distribution; nor do we assume that she is

rational. Furthermore, we assume that the adversary knows the defender's strategy (audit mechanism) and can adapt her strategy accordingly.

Defender model: In each round, the defender inspects a subset of actions of each type performed by the adversary. The defender has to take two competing factors into account. First, inspections incur cost. The defender has an audit budget that imposes upper bounds on how many actions of each type she can inspect. We assume that the cost of inspection increases linearly with the number of inspections. So, if the defender inspects fewer actions, she incurs lower cost. Note that, because the defender cannot know with certainty whether the actions not inspected were malicious or benign, this is a game of imperfect information [6]. Second, the defender suffers a loss in reputation for detected violations. The loss is higher for violations that are detected externally (e.g., by an Health and Human Services audit, or because the breach is publicized by the media) than those that are caught by the defender's audit mechanism, thus incentivizing the defender to inspect more actions.

In addition, the loss incurred from a detected violation depends on the type of violation. For example, inappropriate access of celebrities' patient records might cause higher loss to a hospital than inappropriate access of other patients' records. Also, to account for the evolution of public memory, we assume that violations detected in recent rounds cause greater loss than those detected in rounds farther in the past. The defender's audit mechanism has to take all these considerations into account in prescribing the number of actions of each type that should be inspected in a given round, keeping in mind that the defender is playing against the powerful strategic adversary described earlier.

Note that for adequate privacy protection, the economic and legal structure must ensure that it is in the best interests of the organization to audit significantly. Our abstraction of the reputation loss from policy violations that incentivizes organizations to audit can, in practice, be achieved through penalties imposed by government audits as well as through market forces, such as brand name erosion and lawsuits.

Regret property: We formulate a desirable property for the audit mechanism by adopting the concept of regret from online learning theory. The idea is to compare the loss incurred when the real defender plays according to the strategy prescribed by the audit mechanism to the loss incurred by a hypothetical defender with perfect knowledge of the number of violations of each type in each round. The hypothetical defender is allowed to pick a fixed strategy to play in each round that prescribes how many actions of each type to inspect. The *regret* of the real defender in hindsight is the difference between the loss of the hypothetical defender and the actual loss of the real defender averaged over all rounds of game play. We require that the regret of the audit mechanism quickly converges to a small value and, in particular, to zero as the number of rounds tends to infinity.

Intuitively, this definition captures the idea that although the defender does not know in advance how to allocate her audit budget to inspect different types of accesses (e.g., celebrity record accesses vs. non-celebrity record accesses), the recommendations from the audit mechanism should have the desirable property that over time the budget allocation comes close to the optimal fixed allocation. For example, if the best strategy is to allocate 40% of the budget to inspect celebrity accesses and 60% to non-celebrity accesses, then the algorithm should quickly converge towards these values.

Audit mechanism: We develop a new audit mechanism that provably minimizes regret for the defender. The algorithm is efficient and can be used in practice. In each round, the algorithm prescribes how many actions of each type the defender should inspect. It does so by maintaining weights for each possible defender action and picking an action with probability proportional to the weight of that action. The weights are updated based on a loss estimation function, which is computed from the observed loss in each round. Intuitively, the algorithm learns the optimal distribution over actions by increasing the weights of actions that yielded better payoff than the expected payoff of the current distribution and decreasing the weight otherwise.

The use of a loss estimation function and the characterization of its properties is a novel contribution of this paper that allows us to achieve significantly better bounds than prior work in the regret minimization literature. Our main technical result is that the exact bound on regret for this algorithm is approximately $2\sqrt{2\frac{\ln N}{T}}$ where N is the number of possible defender actions and T is the number of rounds (audit cycles). This bound improves the best known bounds of $O\left(\frac{N^{1/3}\log N}{\sqrt[3]{T}}\right)$ for regret minimization over games of imperfect information. The better bounds are important from a practical standpoint because they imply that the algorithm converges to the optimal fixed strategy much faster.

References

[1] B. W. Lampson, "Computer security in the real world," *IEEE Computer*, vol. 37, no. 6, pp. 37–46, 2004.

[2] D. J. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. A. Hendler, and G. J. Sussman, "Information accountability," *Commun. ACM*, vol. 51, no. 6, pp. 82–87, 2008.

[3] G. Hulme, "Steady Bleed: State of HealthCare Data Breaches," September 2010, InformationWeek.

[4] A. Blum and Y. Mansour, "Learning, regret minimization, and equilibria," *Algorithmic Game Theory*, pp. 79–102, 2007.

[5] D. Fudenberg and J. Tirole, *Game theory.* MIT Press, 1991.

[6] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, "The nonstochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2003.