Understanding Adaptivity: Random Systems Revisited

Dimitar Jetchev[1], Onur Özen[1], and Martijn Stam[2]

[1] EPFL IC IIF LACAL, Station 14, CH-1015 Lausanne, Switzerland
dimitar.jetchev@epfl.ch, oezen.onur@gmail.com
[2] Department of Computer Science, University of Bristol, Merchant Venturers Building, Woodland Road, Bristol BS8 1UB, UK
stam@compsci.bristol.ac.uk

Abstract. We develop a conceptual approach for probabilistic analysis of adaptive adversaries via Maurer's methodology of random systems (Eurocrypt'02). We first consider a well-known comparison theorem of Maurer according to which, under certain hypotheses, adaptivity does not help for achieving a certain event. This theorem has subsequently been misinterpreted, leading to a misrepresentation with one of Maurer's hypotheses being omitted in various applications. In particular, the only proof of (a misrepresentation of) the theorem available in the literature contained a flaw. We clarify the theorem by pointing out a simple example illustrating why the hypothesis of Maurer is necessary for the comparison statement to hold and provide a correct proof. Furthermore, we prove several technical statements applicable in more general settings where adaptivity might be helpful, which can be seen as the random system analogue of the game-playing arguments recently proved by Jetchev, Özen and Stam (TCC'12).

1 Introduction

One of the key concepts in cryptographic security definitions and proofs is the notion of indistinguishability [3]. In the information-theoretic setting, the simplest example is how easy it is for a computationally unbounded adversary to distinguish two random variables X and Y based on a single sample from either of the two variables. It is not hard to see that the success probability of the optimal distinguishing algorithm (the distinguisher's advantage) is simply the statistical distance of the two probability distributions for X and Y. Yet, the analysis of current cryptographic systems typically requires much more than distinguishing two random variables. For instance, the related cryptographic primitive of a pseudo-random function allows an adversary to make multiple queries and hence, obtain multiple related samples in order to distinguish between either a truly random function or a pseudo-random one. Moreover, the distinguisher can interact with the system by choosing the queries *adaptively*, i.e., based on the previous queries and corresponding responses. Adversarial adaptivity is notoriously difficult to deal with, not only in the context of pseudorandomness, but across the cryptologic landscape.

With the increasing number of sophisticated cryptographic schemes appearing in the literature (e.g., authenticated encryption, compression functions, message authentication codes), the level of complexity of proving even relatively straightforward security notions such as pseudorandomness or collision resistance becomes ever more

X. Wang and K. Sako (Eds.): ASIACRYPT 2012, LNCS 7658, pp. 313–330, 2012.

involved and complicated. Even though the building blocks of the proofs rarely extend beyond basic notions such as conditional probabilities, Bayes' rule or basic concepts from stochastic processes, combining these building blocks into a rigorous proof poses a challenge in many cases. Consequently, developing a more conceptual approach towards rigorous security analyses of adaptive adversaries is an important challenge in theoretical cryptology.

Games and Random Systems. One of the general methods for security proofs is based on "game-playing" [2,8,16]. A common technique involves the introduction to the game of a flag bad (initially set to `false`). The *fundamental lemma of game playing* [2, §3.4] states that for games that are identical until bad, distinguishing between these games is at most as hard as setting bad to `true`. Several common and a few new techniques employed to prove preimage and collision security of compression functions based on ideal primitives were recently abstracted using game playing by Jetchev, Özen and Stam [7].

A different approach to indistinguishability and probabilistic analysis of adaptive adversaries is through the concept of *random systems*, as introduced by Maurer [11]. This abstraction unifies many existing security proofs and it allows for proving new indistinguishability results. Intuitively, a random system takes a generally unbounded sequence of inputs (queries) and produces an output (response) for each input using a specific source of randomness. Random systems are rigorously modeled in such a way that they exploit the input-output behavior via specifying (abstractly) a set of conditional probability distributions (see Definition 1 for more details).

A *distinguisher* (see Definition 4) can be thought of as another random system that is allowed to query either one of the two random systems and that outputs a binary decision bit at the end. Estimating the advantage in the case of non-adaptive adversaries is often much simpler than estimating the advantage for adaptive ones. Maurer gave a two step approach to deal with adaptive distinguishers effectively.

First, in analogy with the fundamental lemma of game playing, it is always possible to rephrase the problem of upper bounding the advantage of any adversary in distinguishing two arbitrary random systems into one where an adversary has to provoke an event instead [11, 14, Thm.1]. Most of the indistinguishability proofs indeed follow along these lines.

Next, Maurer [11, Thm.2] presented a result stating that, under certain hypotheses, adaptivity does not help to cause an event. Throughout the paper, we often refer to this statement as the adaptive–non-adaptive (ANA) switching lemma (see Section 4.1). It can also be used in the context of events that are meaningful in their own right, such as finding collisions for a hash function.

Our Contribution. In this paper, we revisit and refine the currently existing techniques based on random systems for bounding the advantage of an adaptive adversary for provoking a certain event. Our contribution is twofold. On the one hand, we show that Maurer's phrasing of the ANA switching lemma has been been misinterpreted, in the sense that an essential hypothesis has been omitted in subsequent applications.

This applies to the only proof given in the literature (by Pietrzak [15, §3.2]) which consequently contains an incorrect step. We restate and prove a corrected version that luckily works for most uses of the lemma in the literature. We explain why the original hypothesis is indeed necessary by providing a simple example where adaptivity does help, yet, where the remaining hypotheses have been satisfied. On the other hand, we examine existing techniques to bound the advantage of adaptive adversaries directly in the context of random systems. This can be seen as a generalization of the earlier work by Jetchev, Özen and Stam [7].

The example is rather simple and intuitive: finding a fixed point in a uniformly random permutation. Here, one can easily see that adaptivity is helpful after the first query/response pair is obtained since (assuming that the first query has not produced a fixed point) an adaptive adversary can choose its second query based on the response to the first query and the condition that there is no fixed point yet (see Section 4.1). Indeed, an adaptive adversary can already eliminate one choice for the second query (two for the third and so on), as opposed to a non-adaptive adversary who commits all of its queries in advance. Thus the best adaptive adversary will have a significantly better advantage than any non-adaptive one. Nevertheless, as we demonstrate, the hypotheses of Pietrzak's (mis)interpretation of the ANA switching lemma *are* satisfied, thus completing our counterexample.

We proceed to examine Pietrzak's proof of the lemma to determine what underlies the mistake and whether the proof can be fixed. To some extent, the problem originates from the elliptical notation that the theory of random systems occasionally suffers from. We propose a restatement of the lemma (Theorem 12) together with a correct proof. We then perform the important (if somewhat tedious) task of investigating known examples in the literature where an incorrect version of the ANA switching lemma has been exploited (see the full version). Fortunately, to the best of our knowledge, the flaw uncovered by us does not lead to a violation of any security claim based on the incorrect ANA switching lemma (as the modified hypotheses are still satisfied).

Our second contribution is a string of technical statements, all phrased in the language of random systems, that are applicable in the more general setting where adaptivity might be helpful in triggering an event. The first result (Proposition 9) is the random system interpretation of a well-known technique, where a union bound is computed over the subevent that an adversary provokes the event at the jth step, where the required "stepwise" probabilities (for the subevents) are maximized in a greedy-type manner. This is a standard and often-used argument from security proofs that has not been previously linked to random systems. It makes derivation of the overall bound relatively easy. Yet, in many cases the overall upper bound is not tight enough due to the maximal probabilities occuring for rather unlikely query/response histories or due to overcounting.

Several proofs in the literature tackle the problem of "bad" query/response history by the introduction of an auxiliary event explicitly bounding such a bad history occuring (e.g., [10,17]). Subsequent bounding of the probabilities of on the one hand the auxiliary event and on the other of the actual event conditioned on the auxiliary bad event not occurring, leads to a tighter bound. Proposition 13 generalizes this method in the context of random systems.

Lee et al. [9] recently introduced "wish lists" to the analysis of adaptive adversaries to limit the effect of overcounting. The idea is to cut up the analysis in two parts. First, one upper bounds the maximum size W of the wish list, i.e., the total number of query-response pairs that could ever lead to an adversarial win (to get useful bounds, one typically needs to introduce an auxiliary flag as in the discussion above). Next, one upper bounds the probability p of any particular wish to be granted, i.e., the probability that a query on the wish list gets to the wished for response when actually being asked by the adversary. Finally, one observes that in order to win, at some point an adversary needs to have some wish granted. Intuitively, a union bound over all wishes in the list means the advantage of an adaptive adversary is then at most pW. We formalize this approach in Proposition 14, which assumes as a hypothesis an upper bound on the sum of the stepwise probabilities of success for each query/response history and thus avoids the greedy-type argument. We refine this in Proposition 15 by adding an auxiliary flag event.

Yet, the most subtle and useful (in terms of applications) bounds are provided in Proposition 16. Here, an adaptive adversary is trying to achieve a certain event more than once. A simple example is an adversary trying to obtain more than κ fixed points in a random permutation, but it could also relate to a scenario where an adversary needs to see multiple wishes being granted. The techniques we develop here are very similar to those used for the analysis of a recent incidence-based compression function construction [7]. We illustrate the usefulness of our result by revisiting the analysis of an auxiliary collinearity event needed for the security proof of that construction (see the full version). The strong emphasis on conditional probabilities in the random systems methodology makes it very natural to express the various bounds on an adaptive adversary's advantage, providing a different and arguably clearer perspective on the original proof.

Related Work. Modification of the adversary is an important technique, orthogonal to our work, that is often used to bound the advantage of an adaptive adversary. In particular free queries have been used to great effect in the analysis of double length hash functions [1,6,9]. A typical proof will first modify the adversary—adding the free queries with the somewhat paradoxic effect of taking away some of the adaptivity of the adversary by making it more it more powerful—followed by an analysis of the advantage of this modified adversary. For bounding the advantage of the modified adversary our work comes into play.

Very recently, during their analysis of key-alternating ciphers, Bogdanov et al. [4] uncovered an interesting scenario where a distinguisher surprisingly benefits from adaptivity. While it would be straightforward to describe their problem (and the supporting counterexample) in the random systems framework and subsequently applying the first step of Maurer's two step approach to move it from distinguishing to causing an event, the resulting event cannot be expressed as a predicate, ruling out direct application of many of our theorems. It is an interesting open problem to see if our approach can be extended to improve upon the bounds already obtained by Steinberger [18] and Bogdanov et al.

2 Preliminaries

Notation. Following the terminology and notation of [11, 15], we denote random variables by capital letters (e.g., X), their values by lower-case letters (e.g., x) and their finite[1] sampling spaces by calligraphic letters (e.g., $\mathcal{X}$). For a fixed sample space $\mathcal{X}$, let $\mathcal{X}^k$ be k-fold Cartesian product of $\mathcal{X}$. The corresponding random variables and their values are denoted analogously (i.e., X^k and x^k, respectively). For brevity, we use $\mathsf{P}_A[a]$ to denote the probability $\Pr[A = a]$ and similarly, $\mathsf{P}_{A|BC}[a; b, c]$ for $\Pr[A = a|B = b \wedge C = c]$. If it is clear from the context, we sometimes omit the specific values and simply use, e.g., $\mathsf{P}_{A|BC}$ to denote $\Pr[A = a|B = b \wedge C = c]$.

Random Systems. Various cryptographic systems can be seen as random systems [11] that are modeled as the mathematical abstraction of interactive systems: an $(\mathcal{X}, \mathcal{Y})$-random system takes the inputs $X_1, X_2, \ldots \in \mathcal{X}$ and for each input X_i it generates an output $Y_i \in \mathcal{Y}$ depending probabilistically on $X^i = (X_1, \ldots, X_i)$ and $Y^{i-1} = (Y_1, \ldots, Y_{i-1})$. Random systems have been used in the literature (see e.g., [11–14]) to unify, simplify, generalize, and in some cases strengthen security proofs.

Definition 1 (Random System). *An $(\mathcal{X}, \mathcal{Y})$-random system $\mathbf{F}$ is a (possibly infinite) sequence of conditional probability distributions $\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$ for $i \geq 1$; specifically, the distribution of the outputs Y_i conditioned on $X^i = x^i$ (i.e., the ith query x_i and all previous queries $x^{i-1} = (x_1, \ldots, x_{i-1})$) and $Y^{i-1} = y^{i-1}$ (i.e., all previous outputs $y^{i-1} = (y_1, \ldots, y_{i-1})$). Define*

$$\mathsf{P}^{\mathbf{F}}_{Y^i|X^i} := \prod_{j=1}^{i} \mathsf{P}^{\mathbf{F}}_{Y_j|X^jY^{j-1}},$$

where, for completeness, $\mathsf{P}^{\mathbf{F}}_{Y_1|X^1Y^0} := \mathsf{P}^{\mathbf{F}}_{Y_1|X^1} = \mathsf{P}^{\mathbf{F}}_{Y_1|X_1}$. Two $(\mathcal{X}, \mathcal{Y})$-random systems $\mathbf{F}$ and $\mathbf{G}$ are said to be equivalent (denoted by $\mathbf{F} \equiv \mathbf{G}$) if $\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}} = \mathsf{P}^{\mathbf{G}}_{Y_i|X^iY^{i-1}}$ for all $i \geq 1$ and all arguments $(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i$.

Example 2 (Random system). Random functions and random permutations are special cases of random systems. If $(\mathcal{X}, \mathcal{Y})$ is any pair of sets, a *random function* $\mathcal{X} \to \mathcal{Y}$ is a random variable whose values are functions $\mathcal{X} \to \mathcal{Y}$. For any finite set $\mathcal{X}$, a random permutation is a random variable taking values in the set of permutations of $\mathcal{X}$. A *uniformly random function* $\mathbf{R}$ is a random function with uniform distribution over all functions $\mathcal{X} \to \mathcal{Y}$. Using random systems, we have the following:

$$\mathsf{P}^{\mathbf{R}}_{Y_i|X^iY^{i-1}}[y_i; x^i, y^{i-1}] = \begin{cases} 1 & \text{if } x_i = x_j \text{ for some } j < i \text{ and } y_i = y_j\ , \\ 0 & \text{if } x_i = x_j \text{ for some } j < i \text{ and } y_i \neq y_j\ , \\ 1/|\mathcal{Y}| & \text{else.} \end{cases} \quad (1)$$

A uniformly random permutation is defined analogously.

[1] Most of the results and arguments in this paper generalize to infinite sampling spaces; for simplicity, we restrict to finite spaces as the latter are the ones relevant for cryptographic applications.

Distinguishing Random Systems. In order to distinguish two $(\mathcal{X}, \mathcal{Y})$-random systems $\mathbf{F}$ and $\mathbf{G}$, we use the notion of a distinguisher that can be regarded as a random system itself. A distinguisher interacts with random systems by making queries to either $\mathbf{F}$ or $\mathbf{G}$ and outputs a binary decision bit after a certain number of queries. In the sequel, we consider information-theoretic distinguishers only; they are computationally unbounded and the only measure of complexity is the number of queries made by them.

In the literature, distinguishers are classified based on how they interact with the random systems. For instance, adaptive distinguishers choose their ith query X_i depending on the history (i.e., all previous query-response pairs), whereas non-adaptive distinguishers commit all their queries in advance. Throughout, we let Ad and NAd be the classes of all adaptive and non-adaptive distinguishers, respectively. Definition 4 formally introduces the concept of a distinguisher as well as its interaction with random systems via probability theory.

Definition 3 (Distinguisher). *An $(\mathcal{X}, \mathcal{Y})$-distinguisher $\mathbf{D}$ is a $(\mathcal{Y}, \mathcal{X})$-random system defined by a sequence of conditional probability distributions $\mathsf{P}^{\mathbf{D}}_{X_i|Y^{i-1}X^{i-1}}$. That is, it is a $(\mathcal{Y}, \mathcal{X})$-random system that is one query ahead. A $(\mathcal{X}, \mathcal{Y})$-distinguisher $\mathbf{D}$ and an $(\mathcal{X}', \mathcal{Y}')$-random system $\mathbf{F}$ are said to be compatible if $\mathcal{X}' = \mathcal{X}$ and $\mathcal{Y}' = \mathcal{Y}$.*

One models the interaction of a distinguisher with a random system via a random experiment that is a sequence of conditional probability distributions. This is denoted by $\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X_iY_i|X^{i-1}Y^{i-1}}$ and defined simply as

$$\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X_iY_i|X^{i-1}Y^{i-1}} = \mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}\mathsf{P}^{\mathbf{D}}_{X_i|X^{i-1}Y^{i-1}} \ .$$

Intuitively, this models the probabilities of the distinguisher choosing a given query x_i at the ith step and the random system returning a given response y_i conditioned on the history. Moreover, we define

$$\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X^iY^i} = \prod_{j=1}^{i} \mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X_jY_j|X^{j-1}Y^{j-1}} \ .$$

We are now interested in distinguishing two random systems $\mathbf{F}$ and $\mathbf{G}$ where we assume that both systems are compatible with the distinguisher $\mathbf{D}$. The performance of $\mathbf{D}$ (known as the *advantage* of $\mathbf{D}$ in distinguishing $\mathbf{F}$ from $\mathbf{G}$) is generally measured as follows:

Definition 4. *Let $\mathbf{F}$ and $\mathbf{G}$ be two $(\mathcal{X}, \mathcal{Y})$-random systems that are compatible with a distinguisher $\mathbf{D}$. Given an integer $i > 0$, the advantage of $\mathbf{D}$ in distinguishing $\mathbf{F}$ from $\mathbf{G}$ in i queries is defined to be*

$$\boldsymbol{\Delta}^{\mathbf{D}}_i(\mathbf{F}, \mathbf{G}) := \frac{1}{2} \sum_{(x^i,y^i)\in\mathcal{X}^i\times\mathcal{Y}^i} |\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X^iY^i} - \mathsf{P}^{\mathbf{D}\Diamond\mathbf{G}}_{X^iY^i}| \ .$$

Let $\mathcal{C}$ be a class of distinguishers trying to distinguish $\mathbf{F}$ from $\mathbf{G}$. We define the advantage of the best $\mathcal{C}$-distinguisher making i queries to $\mathbf{F}$ and $\mathbf{G}$ as

$$\boldsymbol{\Delta}^{\mathcal{C}}_i(\mathbf{F}, \mathbf{G}) := \max_{\mathbf{D}\in\mathcal{C}} \left\{ \boldsymbol{\Delta}^{\mathbf{D}}_i(\mathbf{F}, \mathbf{G}) \right\} \ .$$

Random Systems with Monotone Conditions. One of the similarities between random systems and game-playing is a notion known as *monotone condition* or *monotone event*. Intuitively, it represents an event that once set, it cannot be "reset" by additional queries. The notion of *monotone event/condition* is more general and should not be confused with *monotone predicate* (or *monotone binary output* as discussed in [5, §2.3]). To explain the difference, let $\mathcal{A} = \{\mathbf{a}_i\}$ be the sequence of events $\mathbf{a}_1, \mathbf{a}_2, \ldots$.

Monotone predicates (or binary outputs) are simpler and less general since the query/response pairs (x^i, y^i) at step i uniquely determine whether the corresponding event $\mathbf{a}_i$ holds or not, whereas the former could be more complex (e.g., $\mathbf{a}_i$ could be the (monotone) event that a certain flag is set in at most 10 steps). In other words, in the case of monotone predicates, the conditional probability of $\mathbf{a}_i$ occurring conditioned on $X^i = x^i \wedge Y^i = y^i$ is binary, whereas monotone events could be more general. For simplicity, we assume that our monotone events are monotone predicates and consider a sequence of boolean predicates a_i indicating whether $\mathbf{a}_i$ holds (i.e., $a_i \Leftrightarrow \mathbf{a}_i$ holds; equivalently, $\neg a_i \Leftrightarrow \mathbf{a}_i$ does not hold) with the property that $\neg a_i \Rightarrow \neg a_{i+1}$ (the latter guarantees monotonicity).

As an example, consider the monotone event $\mathbf{a}_i$ that after the ith query to a uniformly random function, all distinct inputs result in distinct outputs (i.e., there exists no output collisions). It is not difficult to see that $\mathcal{A} = \{\mathbf{a}_i\}$ is a monotone binary output as $\neg a_i \Rightarrow \neg a_{i+1}$ and a_i is completely determined from (x^i, y^i). Equivalently, if there is an output collision for the ith step, there is also an output collision for all the subsequent steps. The monotonicity condition gives rise to a sequence of binary probabilities $\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i} \in \{0, 1\}$ with the property that

$$\forall i \geq 1, \qquad \mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i} = 1 \Rightarrow \mathsf{P}^{\mathbf{F}}_{a_{i-1}|X^{i-1}Y^{i-1}} = 1\,. \tag{2}$$

Associated to a random system with a monotone binary output, we have the following data:

- **D.0 (data defining F):** these are simply the probability distributions $\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$,
- **D.1 (binary probabilities for $\mathcal{A}$):** these are the binary probabilities $\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i}$ (describing the predicates a_i and $\neg a_i$) satisfying (2).

Remark 5. In the case of monotone conditions, the defining probabilities $\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i}$ can be arbitrary real numbers in the interval $[0, 1]$.

We can derive various other probabilities using conditional probabilities/Bayes' rule as well as $\mathbf{D}.\mathbf{0}$ and $\mathbf{D}.\mathbf{1}$:

Event Probabilities for $\mathcal{A}$: These are the probabilities denoted by $\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}}$. Intuitively, $\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}}$ models the probability of the predicate a_i conditioned on the query/response history, as well as on the predicate a_{i-1}. We derive it from $\mathbf{D}.\mathbf{0}$ and $\mathbf{D}.\mathbf{1}$ as follows:

$$\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}} = \sum_{y_i} \mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i} \mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}.$$

Here, one can also derive the probability distributions $\mathsf{P}^{\mathbf{F}}_{\neg a_i|a_{i-1}X^iY^{i-1}}$ simply as $1 - \mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}}$. It is important to note that if the condition $a_{i-1}X^iY^{i-1}$ evaluates to `false` for all y_i for a given (x^i, y^{i-1}), this probability is set to zero (for reasons that will become clear later). We remark that in a similar manner, one can adjoin yet another monotone condition $\mathcal{B}$ to a random system with a monotone condition $\mathcal{A}$.

A Random System Conditioned on $\mathcal{A}$ not Failing (denoted by $\mathbf{F}|\mathcal{A}$): These are probability distributions $\mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}}$ and can be derived from Bayes' rule as follows:

$$\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i} = \mathsf{P}^{\mathbf{F}}_{a_iY_i|X^iY^{i-1}} = \mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}}\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^{i-1}}\ , \tag{3}$$

where the middle term (which has not been defined yet) is a formal symbol for the corresponding probability. Assuming that $\mathsf{P}^{\mathbf{F}}_{a_{i-1}|X^{i-1}Y^{i-1}} = 1$ together with the monotonicity of $\mathcal{A}$, we see that $\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^{i-1}} = \mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}} \neq 0$. One can thus derive the conditional probabilities

$$\mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}} = \frac{\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}\mathsf{P}^{\mathbf{F}}_{a_i|X^iY^i}}{\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}}}.$$

Intuitively, this looks like a random system except that we have conditioned on the predicate a_i. Note that this need not be a probability distribution: for instance, consider the example of a random function $\mathbf{R}\colon \{0,1\}^n \to \{0,1\}^n$ and define the $\mathbf{a}_i$ as the event of having a collision between an input and an output. It might occur that $x_2 = y_1$ in which case $a_2X^2Y^1$ will always evaluate to `false` and thus, the probability $\mathsf{P}^{\mathbf{R}}_{Y_2|a_2X^2Y^1} = 0$ for all y_2, so it will not represent a well-defined distribution on the variable Y_2. In cases when this degeneracy does not occur, we can consider $\mathbf{F}|\mathcal{A}$ as a true random system $\mathbf{G}$ (see Hypothesis 8), denoted $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$. Note that this particular notion of equivalence of a random system and a random system with a monotone condition can be extended slightly in the case of degeneracies too. As described in [11, Defn.6], we say that $\mathbf{F}|\mathcal{A}$ is *equivalent* to $\mathbf{G}$ if $\mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}} = \mathsf{P}^{\mathbf{G}}_{Y_i|X^iY^{i-1}}$ for any i and any values of the parameters for which $\mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}}$ is not identically zero (i.e.., is a distribution). Finally, we note that $\mathbf{F}|\mathcal{A}$ appeared in, e.g., [15, Defn.7].

A Random System with a Condition $\mathcal{A}$ (denoted by $\mathbf{F}^{\mathcal{A}}$): This is the random system corresponding to [15, Defn.6]) and can be derived by

$$\mathsf{P}^{\mathbf{F}}_{a_iY_i|a_{i-1}X^iY^{i-1}} := \mathsf{P}^{\mathbf{F}}_{Y_i|a_iX^iY^{i-1}}\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}X^iY^{i-1}}\ .$$

We also define

$$\mathsf{P}^{\mathbf{F}}_{a_iY^i|X^i} := \prod_{j=1}^{i} \mathsf{P}^{\mathbf{F}}_{a_jY_j|a_{j-1}X^jY^{j-1}}\ .$$

Moreover, we consider distinguishers trying to provoke the negated event $\neg\mathbf{a}_i$ again via a sequence of probability distributions. To indicate the link with $\mathbf{a}_i$, we denote these distributions by $\mathsf{P}^{\mathbf{D}}_{X_i|a_{i-1}X^{i-1}Y^{i-1}}$. As in the case of true random systems, this models the probability distribution of an adversary choosing the ith query based on the

previous responses and the predicate a_{i-1} (meaning that the desired event $\neg\mathbf{a}_{i-1}$ has not occurred after the $(i-1)$st query/response pair).

Using this data, we can derive various probabilities and distributions by imposing Bayes' rule. We define the probabilities for the random experiment $\mathbf{D}\Diamond\mathbf{F}$ by

$$\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{a_i X_i Y_i | a_{i-1} X^{i-1} Y^{i-1}} := \mathsf{P}^{\mathbf{F}}_{a_i Y_i | a_{i-1} X^i Y^{i-1}} \mathsf{P}^{\mathbf{D}}_{X_i | a_{i-1} X^{i-1} Y^{i-1}} \ .$$

Intuitively, this models the probability of choosing a particular query, obtaining a particular response and the predicate a_i (resp., $\neg a_i$) conditioned on the history and the predicate a_{i-1}. Finally, let

$$\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{a_i X^i Y^i} := \prod_{j=1}^{i} \mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{a_j X_j Y_j | a_{j-1} X^{j-1} Y^{j-1}} .$$

Similarly, we define an expression for $\neg a_i$. We are now ready to define the advantage of the distinguisher (adversary) $\mathbf{D}$ in provoking the desired event $\neg\mathbf{a}_i$:

Definition 6. *Let $\mathcal{C}$ be a class of distinguishers $\mathbf{D}$ that are trying to provoke $\neg\mathbf{a}_i$. Given $i > 0$, define $\nu^{\mathbf{D}}(\mathbf{F}, \neg a_i)$ to be the advantage of the distinguisher $\mathbf{D}$ in provoking the event $\neg\mathbf{a}_i$ in the random experiment $\mathbf{D}\Diamond\mathbf{F}$. That is*

$$\nu^{\mathbf{D}}(\mathbf{F}, \neg a_i) = \sum_{(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i} \mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{\neg a_i X^i Y^i} \ .$$

Furthermore, for all $i \geq 1$, define $\nu^{\mathcal{C}}(\mathbf{F}, \neg a_i) := \max_{\mathbf{D} \in \mathcal{C}} \nu^{\mathbf{D}}(\mathbf{F}, \neg a_i)$ to be the maximum advantage over all distinguishers in the class $\mathcal{C}$ trying to provoke $\neg\mathbf{a}_i$.

Finally, we explain the analogue (in the context of random systems) of the fundamental lemma of game-playing and comment on why the random-system statement is more general. Suppose that $\mathbf{F}$ is a random system with a monotone condition $\mathcal{A}$ and let $\mathbf{G}$ be another random system. The analogue of the hypothesis of the fundamental lemma of game-playing (that two games are equivalent up to statements that are evaluated only if $\mathbf{a}_i$ is set to `true`) is simply $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$. Under that hypothesis, we expect that one can bound the distinguishing advantage $\boldsymbol{\Delta}^{\mathbf{D}}_i(\mathbf{F}, \mathbf{G})$ via the advantage $\nu^{\mathbf{D}}(\mathbf{F}, \neg a_i)$ of an adversary to provoke $\neg\mathbf{a}_i$. Interestingly enough, one can deduce the latter from a weaker hypothesis, namely the hypothesis that

$$\mathsf{P}^{\mathbf{F}}_{a_j Y^j | X^j} \leq \mathsf{P}^{\mathbf{G}}_{Y^j | X^j}, \ \forall j \leq i.$$

The following lemma is proven in [15, Lem.6] (see also [11, Thm.1]):

Lemma 7. *Assume that $\mathsf{P}^{\mathbf{F}}_{a_j Y^j | X^j} \leq \mathsf{P}^{\mathbf{G}}_{Y^j | X^j}$ holds for all $j \leq i$. Then for any distinguisher $\mathbf{D}$,*

$$\boldsymbol{\Delta}^{\mathbf{D}}_i(\mathbf{F}, \mathbf{G}) \leq \nu^{\mathbf{D}}(\mathbf{F}, \neg a_i) \ .$$

In the following sections, we develop techniques to upper bound $\nu^{\mathbf{D}}(\mathbf{F}, \neg a_i)$.

3 A Standard Method for Probabilistic Analysis of Adaptive Adversaries

Let $\mathcal{A}$ be a monotone condition and let $\mathbf{F}$ be an $(\mathcal{X}, \mathcal{Y})$-random system. Our goal is to compute an upper bound for $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i)$. The standard way to deal with the overall probability of setting $\neg a_i$ is to bound it by a sum (over $j \leq i$) of the maximum (over all adversaries) probability of winning at the jth step, where these "stepwise" probabilities are only taken over the probability distributions describing $\mathbf{F}$. In other words, for each j, we maximize individually the probability of winning at the jth step assuming that we have not won at step $j-1$. This greedy-type approach for producing an upper bound can be formalized in Proposition 9 (see Appendix of the full version for its proof). We first state an hypothesis that is commonly used throughout the paper.

Hypothesis 8. Let $\mathbf{F}$ be an $(\mathcal{X}, \mathcal{Y})$-random system and let $\mathcal{A}$ be a monotone condition on $\mathbf{F}$. There exists an $(\mathcal{X}, \mathcal{Y})$-random system $\mathbf{G}$ such that $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$, i.e., for all $i \geq 1$ and all $(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i$,

$$\mathsf{P}^{\mathbf{F}}_{Y_i|a_i X^i Y^{i-1}} = \mathsf{P}^{\mathbf{G}}_{Y_i|X^i Y^{i-1}}.$$

Proposition 9. *Let $\mathbf{F}$ be an $(\mathcal{X}, \mathcal{Y})$-random system and let $\mathcal{A}$ be a monotone condition on $\mathbf{F}$. Assuming that $\sum_{j=1}^{i} \max_{(x^j, y^{j-1})} \left\{ \mathsf{P}^{\mathbf{F}}_{\neg a_j|a_{j-1}X^jY^{j-1}} \right\} < 1$, we have*

$$\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i) \leq \sum_{j=1}^{i} \max_{(x^j, y^{j-1})} \left\{ \mathsf{P}^{\mathbf{F}}_{\neg a_j|a_{j-1}X^jY^{j-1}} \right\}.$$

4 When Adaptivity Does Not Help

4.1 Revisiting the Result of Maurer and Pietrzak

Maurer [11] and Pietrzak [15] provide a general method for proving that under certain hypotheses, adaptive strategies are no better than non-adaptive ones in forcing a condition to fail. In other words, if these hypotheses are satisfied, the advantage of the best Ad- and NAd-distinguisher are equal. Here, we show that the hypothesis (Hypothesis 8) used by Pietrzak is not sufficient for the comparison result of [11, 15] (ANA switching lemma) to hold by providing a particular counterexample in Proposition 10 where the hypothesis is clearly satisfied and where adaptivity does help. We then explain the problem in the ANA switching lemma in detail and suggest different ways to remedy it in Section 4.2. The following statement appears in [15, Lem.6]:

Adaptive–Non-Adaptive (ANA) Switching Lemma. Let $\mathcal{A}$ be a monotone condition and let $\mathbf{F}$ be an $(\mathcal{X}, \mathcal{Y})$-random system. If Hypothesis 8 holds for $\mathbf{F}$, $\mathcal{A}$ and an $(\mathcal{X}, \mathcal{Y})$-random system $\mathbf{G}$, then adaptivity does not help in provoking $\neg a_i$. More precisely,

$$\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i) = \nu^{\mathsf{NAd}}(\mathbf{F}, \neg a_i).$$

Now, we present an example of a random system where adaptive adversaries have *better* advantage than the adaptive ones in provoking a welldefined monotone event.

Proposition 10. *Let $\mathcal{X} = \{0,1\}^n$ and let $\mathbf{P}\colon \mathcal{X} \to \mathcal{X}$ be a uniformly random permutation. Let $\mathbf{a}_i$ be the event that $y_j \neq x_j$ for all $j \leq i$ where $y_j = \mathbf{P}(x_j)$. Then $\mathcal{A}$ is monotone and*

$$\nu^{\mathsf{Ad}}(\mathbf{P}, \neg a_i) > \nu^{\mathsf{NAd}}(\mathbf{P}, \neg a_i)\,.$$

Proof. The sequence of predicates $\{a_i\}$ is monotone by definition. We calculate the probability of obtaining a fixed point for $\mathbf{P}$ after at most two queries; the case for general i follows by inspection. After querying $\mathbf{P}$ with any $X_1 = x_1 \in \{0,1\}^n$, the response $y_1 \in \{0,1\}^n$ is uniformly random. Thus, with probability $1/2^n$ a fixed point is found after the first query. Hence,

$$\begin{aligned}\mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \vee Y_1 = x_1] &= \mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \wedge Y_1 \neq x_1] + \mathsf{P}^{\mathbf{P}}[Y_1 = x_1] = \\ &= \mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \wedge Y_1 \neq x_1] + 1/2^n\,.\end{aligned}$$

The distinction between an adaptive and a non-adaptive strategy shows up after the second query: the latter commits the second query in advance whereas the former chooses it adaptively based on the first query and its response.

Case 1: Non-adaptive adversary. If the adversary were non-adaptive, she would have fixed $x_2 \neq x_1$ prior to obtaining the response y_1 and since $\mathbf{P}(x_2) \neq \mathbf{P}(x_1)$ and $\mathbf{P}$ is a uniformly random permutation, $\mathbf{P}(x_2) \in \{0,1\}^n - \{y_1\}$. Note however that if $x_2 = y_1$, no y_2 could lead to a fixed point. Hence (by Bayes' rule),

$$\mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \wedge Y_1 \neq x_1] = \mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \mid Y_1 \neq x_1, x_2]\mathsf{P}^{\mathbf{P}}[Y_1 \neq x_1, x_2]\,.$$

Clearly, $\mathsf{P}^{\mathbf{P}}[Y_1 \neq x_1, x_2] = (2^n - 2)/2^n$. Moreover, y_2 is uniformly random among $\{0,1\}^n - \{y_1\}$, so

$$\begin{aligned}\mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \wedge Y_1 \neq x_1] = \frac{1}{2^n - 1} \cdot \frac{2^n - 2}{2^n} \Rightarrow \mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \vee Y_1 = x_1] = \\ = \frac{1}{2^n - 1} \cdot \frac{2^n - 2}{2^n} + \frac{1}{2^n} < \frac{1}{2^{n-1}}\,.\end{aligned}$$

Since the above analysis holds for any non-adaptive adversary, we conclude that $\nu^{\mathsf{NAd}}(\mathbf{P}, \neg a_2) < 1/2^{n-1}$.

Case 2: Adaptive adversary. Knowing x_1, y_1 and $y_1 \neq x_1$ from the first query, an adaptive adversary can eliminate one choice for the second query x_2 different from x_1, namely $x_2 = y_1$. Thus, a clever adversary will choose $x_2 \in \{0,1\}^n - \{x_1, y_1\}$ so that the chance of finding a fixed point after the second step is $1/(2^n - 1)$. Thus,

$$\begin{aligned}\mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \wedge Y_1 \neq x_1] = \mathsf{P}^{\mathbf{P}}[Y_2 = x_2 \mid Y_1 \neq x_1 \wedge Y_1 \neq x_2]\mathsf{P}^{\mathbf{P}}[Y_1 \neq x_1 \wedge Y_1 \neq x_2] = \\ = \frac{1}{2^n - 1} \cdot \frac{2^n - 1}{2^n},\end{aligned}$$

and we conclude that

$$\nu^{\mathsf{Ad}}(\mathbf{P}, \neg a_2) \geq \frac{1}{2^n} + \frac{1}{2^n - 1} \cdot \frac{2^n - 1}{2^n} = \frac{1}{2^{n-1}} > \nu^{\mathsf{NAd}}(\mathbf{P}, \neg a_2) \; .$$

□

We now explain why Hypothesis 8 holds for the monotone event sequence $\mathcal{A}$ and the random system $\mathbf{P}$.

Proposition 11. *Let $\mathbf{P}$ and $\mathcal{A}$ be as in Proposition 10. Then, Hypothesis 8 holds using the monotone condition $\mathcal{A}$, along with taking $\mathbf{P}$ as $\mathbf{F}$.*

Proof. Let $i = 2$. We simply need to define the distributions (i) $\mathsf{P}^{\mathbf{G}}_{Y_1|X^1}$ for all $y_1 \in \mathcal{Y}$ and $x^1 \in \mathcal{X}^1$, and (ii) $\mathsf{P}^{\mathbf{G}}_{Y_2|X^2Y^1}$ for all $y_2 \in \mathcal{Y}$, $x^2 \in \mathcal{X}^2$ and $y^1 \in \mathcal{Y}^1$. For (i), define

$$\mathsf{P}^{\mathbf{G}}_{Y_1|X^1} = \begin{cases} \frac{1}{2^n - 1} & \text{if } y_1 \neq x_1, \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, $\mathsf{P}^{\mathbf{F}}_{Y_1|a_1X^1} = \mathsf{P}^{\mathbf{G}}_{Y_1|X^1}$. For (ii), assuming $y_1 \neq x_1$ and $x_1 \neq x_2$, define the distribution in the following two cases:

Case 1: $x_2 = y_1$. There are $2^n - 1$ possible values for $y_2 = \mathbf{P}(x_2)$ occurring with equal probabilities and none of these values can lead to a fixed point, so we have

$$\mathsf{P}^{\mathbf{G}}_{Y_2|X^2Y^1} = \begin{cases} 0 & \text{if } y_2 = y_1 = x_2, \\ \frac{1}{2^n - 1} & \text{otherwise .} \end{cases}$$

Case 2: $x_2 \neq y_1$. Here, the case of $y_2 = x_2 \neq y_1$ causes $\neg a_2$, so one can define:

$$\mathsf{P}^{\mathbf{G}}_{Y_2|X^2Y^1} = \begin{cases} 0 & \text{if } y_2 = y_1 \text{ or } y_2 = x_2 \neq y_1 \\ \frac{1}{2^n - 2} & \text{otherwise .} \end{cases}$$

We easily verify that in all cases, $\mathsf{P}^{\mathbf{F}}_{Y_2|a_2X^2Y^1} = \mathsf{P}^{\mathbf{G}}_{Y_2|X^2Y^1}$. □

4.2 Another Look at the Comparison of Adaptive vs. Non-adaptive Adversaries

Propositions 10 and 11 show that the ANA switching lemma cannot hold as stated in [15, Lem.6]. We now analyze in detail the proof of the ANA switching lemma given in [15], identify the step that causes the discrepancy and propose a fix.

The Mistake in the Original Proof [15]. The ANA switching lemma first appears in [11, Thm.2] with the correct hypothesis (see (1) of loc. cit.), but without a proof. A slightly different version referring to the original claim is given in [12, Prop.2] (again

without a proof). The only proof, to the best of our knowledge, appears in [15, Lem.6] and is based on a chain of equalities and inequalities starting with

$$1-\nu^{\mathsf{Ad}}(\mathbf{F},\neg a_i)=\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{(x^i,y^i)}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{a_jY_j|X^jY^{j-1}}\mathsf{P}^{\mathbf{D}}_{X^j|Y^{j-1}}\right)\right\}\ .$$

Similarly to Proposition 9, the proof is based on applying Bayes' rule to $\mathsf{P}^{\mathbf{F}}_{a_jY_j|X^jY^{j-1}}$. The application of the Bayes' rule in [15, Lem.6] is, however, incorrect[2]. The correct application yields (assuming that the conditional distributions are well-defined)

$$\mathsf{P}^{\mathbf{F}}_{Y_ja_j|X^jY^{j-1}}=\mathsf{P}^{\mathbf{F}}_{Y_j|a_jX^jY^{j-1}}\mathsf{P}^{\mathbf{F}}_{a_j|X^jY^{j-1}}\ .$$

The problem is that the term $\mathsf{P}^{\mathbf{F}}_{a_i|X^i}=\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{a_j|X^jY^{j-1}}$ is assumed to be independent of Y^{i-1} (see the top line of [15, p.30] - step (2.26)). There is no reason why (for a fixed x^i) this term should be independent of Y^{i-1}; yet, this is used implicitly in the argument. We have seen in Proposition 10 that the probability $\mathsf{P}^{\mathbf{F}}_{a_2|X^2Y^1}$ depends on Y^1, so the ANA switching lemma does not apply.

Strengthening the Hypotheses. We now propose a simple fix to the ANA switching lemma by adding an extra hypothesis, essentially stating that the probability of achieving a success on the jth query is independent of the *answers* to all the previous queries. This statement (albeit in a different formulation) already appears as (1) in Maurer's original [11, Thm.2], as well as a rephrased reproduction [12, Prop.2]. Neither of these statements comes with a proof and both omit mention of Hypothesis 8, although in [11, Thm.2] an alternative condition (2) is given such that (2) is claimed to imply both (1) and Hypothesis 8.

Our proof of Theorem 12 follows largely along the lines of the (incorrect) proof of Pietrzak, but obviously with fixes applied where necessary. Here, Hypothesis 8 is needed to guarantee that all conditional probabilities $\mathsf{P}^{\mathbf{F}}_{Y_j|a_jX^jY^{j-1}}$ are well-defined and are also distributions when considered as functions on $y_j\in\mathcal{Y}$. The second hypothesis simply says that if there is no dependency of the conditionals $\mathsf{P}^{\mathbf{F}}[a_j|a_{j-1}\wedge X^j=x^j\wedge Y^{j-1}=y^{j-1}]$ on the previous outputs then adaptivity should not help at all.

Theorem 12. *Let $\mathbf{F}$ be an $(\mathcal{X},\mathcal{Y})$-random system and let $\mathcal{A}$ be a monotone condition on $\mathbf{F}$. Let $i>0$ be an integer. Suppose that Hypothesis 8 holds for $\mathbf{F}$ and $\mathcal{A}$. If, in addition, for every $j\le i$ and $x^j\in\mathcal{X}^j$, $\mathsf{P}^{\mathbf{F}}[a_j|a_{j-1}\wedge X^j=x^j\wedge Y^{j-1}=y^{j-1}]$ is independent of $y^{j-1}\in\mathcal{Y}^{j-1}$, then adaptivity does not help in provoking $\neg a_i$, i.e.,*

$$\nu^{\mathsf{Ad}}(\mathbf{F},\neg a_i)=\nu^{\mathsf{NAd}}(\mathbf{F},\neg a_i).$$

Proof of Theorem 12. We first note that $\nu^{\mathsf{Ad}}(\mathbf{F},\neg a_i)\ge\nu^{\mathsf{NAd}}(\mathbf{F},\neg a_i)$ holds. The rest of the proof follows by showing the other direction of the inequality; we have that $1-\nu^{\mathsf{Ad}}(\mathbf{F},\neg a_i)$ equals

[2] Furthermore, the argument in [15, Lem.6] does not state whether the conditional probabilities $\mathsf{P}^{\mathbf{F}}_{Y_j|a_jX^jY^{j-1}}$ are well-defined, for all $j\le i$.

$$
\begin{aligned}
&\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{(x^i,y^i)}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{Y_j|a_jX^jY^{j-1}}\mathsf{P}^{\mathbf{F}}_{a_j|a_{j-1}X^jY^{j-1}}\mathsf{P}^{\mathbf{D}}_{X_j|a_{j-1}X^{j-1}Y^{j-1}}\right)\right\}\\
\overset{(*)}{=}&\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{(x^i,y^i)}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{G}}_{Y_j|X^jY^{j-1}}\mathsf{P}^{\mathbf{F}}_{a_j|a_{j-1}X^j}\mathsf{P}^{\mathbf{D}}_{X_j|a_{j-1}X^{j-1}Y^{j-1}}\right)\right\}\\
=&\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{x^i}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{a_j|a_{j-1}X^j}\right)\sum_{y^i}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{G}}_{Y_j|X^jY^{j-1}}\mathsf{P}^{\mathbf{D}}_{X_j|a_{j-1}X^{j-1}Y^{j-1}}\right)\right\}\\
=&\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{x^i}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{a_j|a_{j-1}X^j}\right)\right\}\\
\geq&\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{x^i}\left(\prod_{j=1}^{i}\mathsf{P}^{\mathbf{F}}_{a_j|a_{j-1}X^j}\mathsf{P}^{\mathbf{D}}_{X_j|a_{j-1}}\right)\right\}=\min_{\mathbf{D}\in\mathsf{Ad}}\left\{\sum_{x^i}\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{a_iX^i}\right\}\\
\geq&=(1-\nu^{\mathsf{NAd}}(\mathbf{F},\neg a_i))\,.
\end{aligned}
$$

Here, (*) uses Hypothesis 8, as well as the extra hypothesis that $\mathsf{P}^{\mathbf{F}}_{a_j|X^jY^{j-1}}$ is independent of y^{j-1}. Hence, $\nu^{\mathsf{NAd}}(\mathbf{F},\neg a_i)) \geq \nu^{\mathsf{Ad}}(\mathbf{F},\neg a_i))$ and the claim follows. □

5 Towards Obtaining Better Bounds

5.1 Using an Auxiliary Flag

The standard approach given in Section 3 has the disadvantage that for more complex constructions, the maximal probabilities can get too large. This is often due to the fact that the maximum is achieved for rather degenerate values of (x^i, y^i) that occur with very low probability. Assuming that one can bound the probability of the degeneracy, one way to refine the analysis of the adaptive adversary is to introduce an auxiliary event (flag) that is set only for non-degenerate pairs (x^i, y^i). More precisely, if $\mathbf{a}_i$ is the monotone event to be studied, we introduce a flag event $\mathbf{b}_i$ (together with a corresponding predicate b_i indicating whether $\mathbf{b}_i$ has occurred or not) and we use the fact that

$$\neg\mathbf{a}_i \Leftrightarrow (\neg\mathbf{a}_i \wedge \mathbf{b}_i) \vee (\neg\mathbf{a}_i \wedge \neg\mathbf{b}_i) \Rightarrow (\neg\mathbf{a}_i \wedge \mathbf{b}_i) \vee \neg\mathbf{b}_i.$$

Now, bounding the advantage of achieving $\neg\mathbf{a}_i$ amounts to bounding the advantage of achieving $\neg\mathbf{a}_i \wedge \mathbf{b}_i$ together with bounding the probability of degeneracy (or, of $\neg\mathbf{b}_i$). The latter can be done via Proposition 9; yet for the former we need to introduce new definitions.

All this can be rigorously modeled using random systems as follows: suppose that $\mathbf{F}$ is a random system with a monotone condition $\mathcal{B}$ (here, $\mathcal{B}$ represents the flag event). Suppose further that $\mathbf{F}|\mathcal{B}$ is equivalent to another random system $\mathbf{G}$ (i.e., $\mathbf{F}|\mathcal{B} \equiv \mathbf{G}$). Now, we simply impose a monotone condition $\mathcal{A}$ on $\mathbf{G}$. Equivalently, we need to specify the corresponding probabilities and distributions from Section 2. Suppose that we are given the following data:

- Event probabilities $\mathsf{P}^{\mathbf{G}}_{a_i|a_{i-1}X^iY^{i-1}}$ also denoted by $\mathsf{P}^{\mathbf{F}}_{a_i|a_{i-1}b_iX^iY^{i-1}}$ (to indicate better what they are supposed to model),
- The random system $\mathbf{G}|\mathcal{A}$, namely, probabilities $\mathsf{P}^{\mathbf{G}}_{Y_i|a_iX^iY^{i-1}}$ that we also denote by $\mathsf{P}^{\mathbf{F}}_{Y_i|a_ib_iX^iY^{i-1}}$,
- Distinguisher relative to $\mathcal{A}$, namely, probability distributions denoted by $\mathsf{P}^{\mathbf{D}}_{X_i|a_{i-1}b_{i-1}X^{i-1}Y^{i-1}}$.

This data allows us to upper bound the advantage $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i \wedge b_i)$ (by defining $\mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{\neg a_i|b_i} = \mathsf{P}^{\mathbf{D}\Diamond\mathbf{G}}_{\neg a_i}$) following exactly the same steps as in Section 2 (for the random system $\mathbf{G}$ and the monotone event $\mathcal{A}$). Moreover, we assume all the corresponding notation. The following proposition provides an upper bound on the adaptive advantage (see Appendix of the full version for its proof):

Proposition 13. *Let* $\mathbf{F}$ *be a random system with a monotone condition* $\mathcal{B}$ *with the property that there exists a random system* $\mathbf{G}$ *such that* $\mathbf{F}|\mathcal{B} \equiv \mathbf{G}$. *Let* $\mathcal{A}$ *be a monotone condition on* $\mathbf{G}$. *Assuming that*

$$\sum_{j=1}^{i} \max_{(x^j,y^{j-1})} \left\{ \mathsf{P}^{\mathbf{F}}_{\neg a_j|a_{j-1}b_jX^jY^{j-1}} \right\} < 1 \,,$$

we have

$$\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i \wedge b_i) \leq \sum_{j=1}^{i} \max_{(x^j,y^{j-1})} \left\{ \mathsf{P}^{\mathbf{F}}_{\neg a_j|a_{j-1}b_jX^jY^{j-1}} \right\} \,.$$

5.2 Improving the Bounds Obtained from Step-Specific Maximization

The greedy approach based on step-specific maximization often has limitations in the sense that the produced bounds are not tight enough. One can obtain better bounds via the simple observation that the advantage of an adversary in provoking $\neg\mathbf{a}_i$ for a monotone event $\mathcal{A}$ can be bounded by the sum of the event probabilities for the negated events $\neg\mathbf{a}_j$ for $j \leq i$ that are part of the data defining the monotone condition $\mathcal{A}$. Consequently, if one is able to provide upper bounds on these sums, one would automatically obtain an upper bound on the adaptive advantage.

In order to carry out this idea rigorously, we consider two methods that are formally stated in Propositions 14 and 15 (see Appendix of the full version for the proof of the former; the proof of the latter follows from the proof of Propositions 13 and 14). We first give ourselves an upper bound B_Σ on the sum of the event probabilities and then show that the same B_Σ bounds the adaptive advantage as well. The second method is a variation of the first where one uses an auxiliary event. These two techniques are important whenever the bounds given in Propositions 9 and 13 are not sufficiently tight). A good example of that is the analysis an adaptive adversary trying to achieve a collision in the compression function of [7] (see Appendix of the full version for the details).

Proposition 14. *Let* $\mathbf{F}$ *be a random system with a monotone event* $\mathcal{A}$*. If there exists a value* $B_\Sigma \in (0,1)$ *such that for all* $(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i$

$$\sum_{j=1}^{i} \mathsf{P}^{\mathbf{F}}_{\neg a_j | a_{j-1} X^j Y^{j-1}} \leq B_\Sigma \ ,$$

then $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i) \leq B_\Sigma$.

The following proposition shows the natural generalization of the above proposition to the case of auxiliary events (its proof follows from the proof of Propositions 13 and 14):

Proposition 15. *Let* $\mathbf{F}$ *be a random system with a monotone condition* $\mathcal{B}$ *with the property that there exists a random system* $\mathbf{G}$ *such that* $\mathbf{F}|\mathcal{B} \equiv \mathbf{G}$*. Let* $\mathcal{A}$ *be a monotone condition on* $\mathbf{G}$*. Suppose that there exists a value* $B_\Sigma \in (0,1)$ *such that for all* $(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i$

$$\sum_{j=1}^{i} \mathsf{P}^{\mathbf{F}}_{\neg a_j | a_{j-1} b_j X^j Y^{j-1}} = \sum_{j=1}^{i} \mathsf{P}^{\mathbf{G}}_{\neg a_j | a_{j-1} X^j Y^{j-1}} \leq B_\Sigma \ .$$

Then $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_i \wedge b_i) \leq B_\Sigma$.

Counting Successes. In Proposition 14, we are mainly interested in estimating the maximal probability of the event (success) occurring once. Nevertheless, in some cases the major monotone event $\mathcal{A}$ might depend on an auxiliary condition that intrinsically requires an event (success) to occur more than once. As a simple example, consider a generalization of the case studied in Proposition 10: let $\mathbf{P}$ be a uniformly random permutation $\mathbf{P}\colon \mathcal{X} \to \mathcal{X}$ for $\mathcal{X} = \{0,1\}^n$ and let $\neg \mathbf{a}_i$ be the event that $y_j = x_j$ for more than κ values of $j \leq i$ where $y_j = \mathbf{P}(x_j)$ and κ is a positive integer. More precisely, a_i is the predicate that there exist at most κ fixed points after the ith query.

Such a general problem can be modeled and studied using random systems as follows: suppose that $\mathbf{F}$ is an $(\mathcal{X}, \mathcal{Y})$-random system. We then attach an event called hit_i to the random system $\mathbf{F}$ - this is the success event at step i. Note that hit_i is not monotone. Moreover, we introduce a random variable ctr_i to indicate the number of successes up to step i. In other words, $\mathsf{ctr}_0 = 0$ and for every $j \geq 1$, $\mathsf{ctr}_j = \mathsf{ctr}_{j-1} + 1$ if hit_j occurs and $\mathsf{ctr}_j = \mathsf{ctr}_{j-1}$ otherwise. Finally, we can associate monotone events $\mathcal{A}_\kappa = \{\mathbf{a}_{\kappa,i}\}$ for every integer $\kappa \geq 0$, so that $\mathbf{a}_{\kappa,i}$ is event that there are at most κ successes after the ith query. In other words, $\mathbf{a}_{\kappa,i}$ is the event that $\mathsf{ctr}_i \leq \kappa$.

In order to attach the success event to the random system, we provide the following additional data to $\mathbf{D}.\mathbf{0}$:

$\mathbf{H}.\mathbf{1}$: Binary probabilities $\mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_i | X^i Y^i}$ for every $x^i \in \mathcal{X}^i$ and $y^i \in \mathcal{Y}^i$.

We can derive the following probabilities from $\mathbf{D}.\mathbf{0}$ and $\mathbf{H}.\mathbf{1}$ via Bayes' rule:

- Probabilities $\mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_i | X^i Y^{i-1}}$ for every $x^i \in \mathcal{X}^i$ and $y^{i-1} \in \mathcal{Y}^{i-1}$ defined by

$$\mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_i | X^i Y^{i-1}} = \sum_{y_i} \mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_i | X^i Y^i} \mathsf{P}^{\mathbf{F}}_{Y_i | X^i Y^{i-1}} .$$

- The data for each of the monotone events $\mathcal{A}_\kappa$.

Proposition 16 sets an upper bound on $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_{i,\kappa})$ (see Appendix of the full version for its proof).

Proposition 16. *Let κ be a non-negative integer and suppose that there exists a value $B_\Sigma \in (0,1)$ such that for all $(x^i, y^i) \in \mathcal{X}^i \times \mathcal{Y}^i$,*

$$\sum_{j=0}^{i} \mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_j | X^j Y^{j-1}} \leq B_\Sigma \qquad \textit{and} \qquad \mathsf{P}^{\mathbf{F}}_{\mathsf{hit}_i | X^i Y^{i-1}} > 0.$$

Then $\nu^{\mathsf{Ad}}(\mathbf{F}, \neg a_{i,\kappa}) \leq B_\Sigma^{\kappa+1}$.

Remark 17. We should indicate the analogy between Proposition 14 and Proposition 16 with [7, Prop.7] and [7, Prop.9], respectively. We believe that having such statements and techniques developed in the general context of random systems could serve as a guiding tool for more conceptual security proofs for other constructions in the future.

Acknowledgement. The authors gratefully acknowledge Krzysztof Pietrzak for insightful discussions and the Crypto'12 and Asiacrypt'12 program committees for their useful feedback. This work has been supported in part by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II. Jetchev and Özen were supported by a grant of the Swiss National Science Foundation, 200021-122162. The work was initiated while Stam was at EPFL.

References

1. Armknecht, F., Fleischmann, E., Krause, M., Lee, J., Stam, M., Steinberger, J.: The Preimage Security of Double-Block-Length Compression Functions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 233–251. Springer, Heidelberg (2011)
2. Bellare, M., Rogaway, P.: The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)
3. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo random bits. In: FOCS, pp. 112–117. IEEE Computer Society (1982)
4. Bogdanov, A., Knudsen, L.R., Leander, G., Standaert, F.-X., Steinberger, J., Tischhauser, E.: Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations (Extended Abstract). In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 45–62. Springer, Heidelberg (2012)
5. Gaži, P., Maurer, U.: Free-Start Distinguishing: Combining Two Types of Indistinguishability Amplification. In: Kurosawa, K. (ed.) ICITS 2009. LNCS, vol. 5973, pp. 28–44. Springer, Heidelberg (2010)
6. Hirose, S.: Some Plausible Constructions of Double-Block-Length Hash Functions. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, pp. 210–225. Springer, Heidelberg (2006)
7. Jetchev, D., Özen, O., Stam, M.: Collisions Are Not Incidental: A Compression Function Exploiting Discrete Geometry. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 303–320. Springer, Heidelberg (2012)
8. Kilian, J., Rogaway, P.: How to Protect DES Against Exhaustive Key Search (an Analysis of DESX). J. Cryptology 14(1), 17–35 (2001)

9. Lee, J., Stam, M., Steinberger, J.: The Collision Security of Tandem-DM in the Ideal Cipher Model. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 561–577. Springer, Heidelberg (2011)
10. Lucks, S.: A collision-resistant rate-1 double-block-length hash function. In: Biham, E., Handschuh, H., Lucks, S., Rijmen, V. (eds.) Symmetric Cryptography. No. 07021 in Dagstuhl Seminar Proceedings, Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, Dagstuhl, Germany (2007), http://drops.dagstuhl.de/opus/volltexte/2007/1017
11. Maurer, U.M.: Indistinguishability of Random Systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002)
12. Maurer, U., Pietrzak, K.: The Security of Many-Round Luby–Rackoff Pseudo-Random Permutations. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 544–561. Springer, Heidelberg (2003)
13. Maurer, U.M., Pietrzak, K.: Composition of Random Systems: When Two Weak Make One Strong. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 410–427. Springer, Heidelberg (2004)
14. Maurer, U.M., Pietrzak, K., Renner, R.S.: Indistinguishability Amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007)
15. Pietrzak, K.: Indistinguishability and Composition of Random Systems. ETH Zurich, Ph.D. thesis (2005), http://homepages.cwi.nl/%7Epietrzak/publications/thesis05.ps
16. Shoup, V.: Sequences of Games: A Tool for Taming Complexity in Security Proofs. Cryptology ePrint Archive, Report 2004/332 (2004), http://eprint.iacr.org/
17. Steinberger, J.P.: The Collision Intractability of MDC-2 in the Ideal-Cipher Model. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 34–51. Springer, Heidelberg (2007)
18. Steinberger, J.P.: Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Report 2012/481 (2012), http://eprint.iacr.org/