A Characterization of Authenticated-Encryption as a Form of Chosen-Ciphertext Security

T. Shrimpton *

October 18, 2004

Abstract

In this note we introduce a variation of the standard definition of chosen-ciphertext security, which we call IND-CCA3, and prove that IND-CCA3 is equivalent to authenticated-encryption.

1 Introduction

An *authenticated-encryption scheme*, formalized by Bellare and Namprempre [1], Bellare and Rogaway [2], and Katz and Yung [4], achieves two security goals: privacy and authenticity. We refer to this combination of goals as *authenticated-encryption*. When privacy is defined as indistinguishability under a chosen-plaintext attack, and authenticity as existential unforgeablility of ciphertexts, then authenticated-encryption is the strongest known notion of security for a symmetric encryption scheme [1]. Specifically, it is stronger than the Rackoff-Simon notion of chosen-ciphertext security [5], dubbed IND-CCA2 in [3].

In this note we introduce a new and succinct definition of authenticated-encryption. We call it IND-CCA3, as it is a variation, and a strengthening, of IND-CCA2. Informally, the IND-CCA3 notion says that it should be hard for an adversary to distinguish between two worlds. In the first world, the adversary is given an a pair of oracles $\mathcal{E}_K(\cdot), \mathcal{D}_K(\cdot)$ that perform real encryption and decryption of strings of its choosing under some secret key K. In the second, the adversary is given a pair of oracles $\mathcal{E}_K(\$^{|\cdot|}), \perp(\cdot)$. The first of these oracles returns the encryption of random strings of the appropriate length, and the second tells the adversary that every ciphertext it queries is invalid. The presence of this bogus decryption oracle $\perp(\cdot)$ in the second world is, syntactically, all that differentiates the IND-CCA3 and IND-CCA2 definitions. But this small difference seems crucial to achieving authenticated-encryption in a chosen-ciphertext attack model: our main result proves that the IND-CCA3 and authenticated-encryption are equivalent notions.

Our new notion effectively binds together the adversarial goals of *distinguishing* and *forging*, which had previously been considered separately. The intuition is this: if an adversary cannot easily distinguish between these two pairs of oracles, then ciphertexts produced by $\mathcal{E}$ do not help the adversary to forge new, valid ciphertexts; similarly, the decryption algorithm $\mathcal{D}$ does not help the adversary to distinguish real ciphertexts from bogus ones.

2 Syntax and Notation

Fix a *key space* Key, a *message space* $\mathsf{Message} \subseteq \{0,1\}^*$, and a *ciphertext space* $\mathsf{Ciphertext} \subseteq \{0,1\}^*$. The set Key is finite or is otherwise endowed with a distribution (the understood distribution on

* Department of Computer Science, Portland State University Portland, Oregon, 97201, USA. E-mail: teshrim@cs.pdx.edu WWW: www.cs.pdx.edu/~teshrim/

a finite set being the uniform one). We insist that Message has the structure that if $M \in \mathsf{Message}$ then $\{0,1\}^{|M|} \subseteq \mathsf{Message}$. An *encryption scheme* $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is a triple of algorithms. The probabilistic *key generation* algorithm $\mathcal{K}$ returns a key $K \in \mathsf{Key}$; we write $K \xleftarrow{\$} \mathcal{K}$. The *encryption* algorithm $\mathcal{E}$ could be probabilistic or stateful. It takes a key $K \in \mathsf{Key}$ and a message $M \in \mathsf{Message}$ to return a *ciphertext* $C \in \{0,1\}^* \cup \{\text{INVALID}\}$; we write $C \xleftarrow{\$} \mathcal{E}_K(M)$. (If randomized, it flips new coins on each invocation. If stateful, it uses and then updates state that is maintained across invocations.) The distinguished symbol INVALID is returned if $M \notin \mathsf{Message}$. The *decryption* algorithm $\mathcal{D}$ is deterministic. It takes a key $K \in \mathsf{Key}$ and a string $C \in \{0,1\}^*$ to return some $M \in \mathsf{Message} \cup \{\text{INVALID}\}$; we write $M \leftarrow \mathcal{D}_K(C)$. When $\mathcal{D}_K(C)$ returns $M = \text{INVALID}$ it denotes that C is not authentic. We require that $\mathcal{D}_K(\mathcal{E}_K(M)) = M$ for any $K \in \mathsf{Key}$ and $M \in \mathsf{Message}$.

An *adversary* is a probabilistic algorithm that may have access to oracles. When it is necessary to make an algorithm's oracles explicit, we write them as superscripts. When we write $A \Rightarrow b$ we are referring to the event that adversary A outputs the bit b. We insist that once an adversary outputs it also halts.

3 Preliminaries

PRIVACY. We begin by giving two standard definitions of privacy, following [3]. Let A be an adversary and let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. We define

$$\mathbf{Adv}^{\text{ind-cpa}}_{\Pi}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot)} \Rightarrow 1\right] - \Pr\left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|})} \Rightarrow 1\right]$$

$$\mathbf{Adv}^{\text{ind-cca2}}_{\Pi}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1\right] - \Pr\left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\mathcal{D}_K(\cdot)} \Rightarrow 1\right]$$

The oracle $\mathcal{E}_K(\cdot)$, on input M, returns the encryption $\mathcal{E}_K(M)$. The oracle $\mathcal{E}_K(\$^{|\cdot|})$, on input M, returns the encryption of $|M|$ random bits. The oracle $\mathcal{D}_K(\cdot)$, on input C, returns the decryption $\mathcal{D}_K(C)$. Informally, IND-CPA captures the intuition that a good encryption scheme should produce ciphertexts that look like the encryption of random plaintexts. The IND-CCA2 notion says that encryptions should still look random even in the presence of a decryption oracle. Alternatively, IND-CCA2 says that it should be hard for an adversary to distinguish between a "real" world in which it interacts with a pair of *real* oracles for encryption and decryption, and a "bogus" world in which it interacts with a pair consisting of a *bogus* encryption oracle and a real decryption oracle.

For these definitions we make the following assumptions: the adversary A makes only well-formed queries, never repeats a query, and when given two oracles it never queries C of its right oracle if C was the result of previous left-oracle query. The first two assumptions are without loss of generality, and the third is made to prevent a trivial distinguishing adversary for IND-CCA2.

AUTHENTICITY. Here we give a notion of *authenticity of ciphertexts* for encryption schemes, following [1, 2, 4]. Fix an encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. Let A be a adversary having an encryption oracle $\mathcal{E}_K$. We say that A *forges* if it outputs a ciphertext C such that C was not the response to any $\mathcal{E}_K(M)$ query and $\mathcal{D}_K(C) \neq \text{INVALID}$. We write

$$\mathbf{Adv}^{\text{auth}}_{\Pi}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot)} \text{ forges}\right]$$

We assume that A makes only well-formed queries, and outputs a well-formed ciphertext C. These assumptions are without loss of generality.

Note that the exact formalization of authenticity given in [1] is called INT-CTXT and it is different than AUTH. We chose to use the AUTH formalization of [2] because it is simpler to state,

and since INT-CTXT is qualitatively equivalent to AUTH it makes no qualitative difference to our results.

Since we think of an authenticated-encryption scheme as one achieving both IND-CPA and AUTH, we will often write AE for IND-CPA $\wedge$ AUTH.

Resource-parameterized definitions. If Π is a scheme and A is an adversary and $\mathbf{Adv}^{\mathrm{xxx}}_{\Pi}(A)$ is a measure of adversarial advantage already defined, then we write $\mathbf{Adv}^{\mathrm{xxx}}_{\Pi}(\mathcal{R})$ to mean the maximal value of $\mathbf{Adv}^{\mathrm{xxx}}_{\Pi}(A)$ over all adversaries A that use resources bounded by $\mathcal{R}$. Here $\mathcal{R}$ is a list of variables specifying the resources of interest for the adversary in question. Adversarial resources to which we pay attention are: t—the worst case running time of the adversary; q—the total number of queries asked by the adversary of its oracles; and μ—the aggregate length of these queries, plus the length of the adversary's output, measured in bits. By convention, the running time of an algorithm includes the description size of that algorithm, relative to some standard encoding.

4 A New Notion: IND-CCA3

Here we formalize our new notion, IND-CCA3. The definition is identical to that for IND-CCA2 except for one thing: we replace the decryption oracle in the bogus world with an oracle that always returns Invalid.

Let adversary A be an algorithm with access to an oracle and let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. We define the IND-CCA3 advantage measure as

$$\mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(A) \;=\; \Pr\left[K \xleftarrow{\$} \mathcal{K} :\; A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1\right] - \Pr\left[A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1\right]$$

As before, the oracle $\mathcal{E}_K(\cdot)$, on input M, returns the encryption $\mathcal{E}_K(M)$, and the oracle $\mathcal{E}_K(\$^{|\cdot|})$, on input M, returns the encryption of $|M|$ random bits. The oracle $\mathcal{D}_K(\cdot)$, on input C, returns the decryption $\mathcal{D}_K(C)$. The oracle $\perp(\cdot)$ returns Invalid on any input.

As we did for IND-CCA2, we assume that an IND-CCA3 adversary A makes only well-formed queries, never repeats a query, and it never queries C of its right oracle if C was the result of previous left-oracle query. The first two assumptions are without loss of generality, and the third is made to prevent a trivial distinguishing adversary.

5 IND-CCA3 is Authenticated-Encryption

We now prove that any scheme that achieves security relative to our IND-CCA3 definition *is* an authenticated-encryption scheme. We first show that every encryption scheme meeting the IND-CCA3 notion also meets both the IND-CPA notion and the AUTH notion; symbolically, we write IND-CCA3 $\Rightarrow$ AE. Next, we prove the converse is also true, so that any scheme meeting the IND-CPA and AUTH notions also meets the IND-CCA3 notion. For this statement we write AE $\Rightarrow$ IND-CCA3. While the first implication is tight, the second loses a factor of q — the number of queries made — in the security bound. This is due to the fact that the IND-CCA3 definition effectively allows for q forgery attempts, while the AUTH definition allows for just one. In this way, IND-CCA3 is closer to capturing our intuition about authenticated-encryption than the combination IND-CPA $\wedge$ AUTH.

Both results rely on Lemma 3, whose statement and proof we postpone until the end of this section.

Theorem 1 (IND-CCA3 $\Rightarrow$ AE) Fix $t, q, \mu > 0$. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. Then

$$\begin{aligned} \mathbf{Adv}^{\text{ind-cpa}}_{\Pi}(t,q,\mu) &\leq \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(t,q,\mu) \\ \mathbf{Adv}^{\text{auth}}_{\Pi}(t,q,\mu) &\leq 2 \cdot \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(t',q+1,\mu+1) \end{aligned}$$

where $t' = t + O(\mu)$.

Proof: The first result is trivial, so we focus on the second. Let A gain $\delta = \mathbf{Adv}^{\text{auth}*}_{\Pi}(A)$, run in time t and ask q queries of its oracle, these totaling μ bits. Now,

$$\begin{aligned} \delta &= \Pr[K \leftarrow \mathsf{Key} : A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1] \\ &= \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1]\Big) + \\ &\quad \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1]\Big) \\ &= \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(A) + \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1]\Big) \end{aligned}$$

To bound the second addend above, let B be an IND-CCA3 adversary for Π that runs A, answers A's left-oracle queries by querying its own oracle and returning the result to A, answers A's right-oracle queries with INVALID. When A outputs bit b, adversary B outputs $1 - b$. Clearly

$$\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1] \leq \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(B)$$

and so $\delta \leq \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(A) + \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(B)$. Considering the resources used by A and B, we can conclude $\mathbf{Adv}^{\text{auth}*}_{\Pi}(t,q,\mu) \leq 2 \cdot \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(t,q,\mu)$. Appealing to Lemma 3 finishes the proof. ∎

Theorem 2 (AE $\Rightarrow$ IND-CCA3) Fix $t, q, \mu > 0$. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. Then

$$\mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(t,q,\mu) \leq \mathbf{Adv}^{\text{ind-cpa}}_{\Pi}(t,q,\mu) + q \cdot \mathbf{Adv}^{\text{auth}}_{\Pi}(t',q,\mu)$$

where $t' = t + O(q)$.

Proof: Let A be an adversary that gains $\delta = \mathbf{Adv}^{\text{ind-cca3}}_{\Pi}(A)$, runs in time t, asks q queries, these totaling μ bits. Then,

$$\begin{aligned} \delta &= \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1] \\ &= \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1]\Big) + \\ &\quad \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1]\Big) \\ &= \mathbf{Adv}^{\text{auth}*}_{\Pi}(A) + \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1]\Big) \end{aligned}$$

To bound the second addend above, let B be an ind-cpa-adversary for Π that runs A, answers A's left-oracle queries by querying its own oracle and returning the result to A, answers A's right-oracle queries with INVALID, and outputting whatever bit A does. Clearly

$$\Pr[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1] - \Pr[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\$^{|\cdot|}),\perp(\cdot)} \Rightarrow 1] \leq \mathbf{Adv}^{\text{ind-cpa}}_{\Pi}(B)$$

and so $\delta \leq \mathbf{Adv}^{\text{auth}*}_{\Pi}(A) + \mathbf{Adv}^{\text{ind-cpa}}_{\Pi}(B)$. Appealing to Lemma 3, we have the claimed result. ∎

The following lemma says that the experiment used to define AUTH (an adversary, given an $E_K(\cdot)$ oracle, attempts to forge a new ciphertext given) can be recast as an experiment in which the adversary, given an $E_K(\cdot)$ oracle, attempts to distinguish between a real decryption oracle and a bogus decryption oracle that returns INVALID on every input ciphertext.

Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. We define the following advantage measure,

$$\mathbf{Adv}^{\text{auth}*}_{\Pi}(A) = \Pr\left[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1\right] - \left[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1\right]$$

where $\perp(\cdot)$, as before, returns INVALID on every query. The adversary is forbidden, as usual, from asking C of its right oracle if C was the result of a previous left-oracle query. We lift this definition to a resource-parameterized one in the usual way.

Lemma 3 Fix $t, q, \mu > 0$. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an encryption scheme. Then

$$\begin{aligned} \mathbf{Adv}^{\text{auth}}_{\Pi}(t, q, \mu) &\leq \mathbf{Adv}^{\text{auth}*}_{\Pi}(t', q+1, \mu+1) \\ \mathbf{Adv}^{\text{auth}*}_{\Pi}(t, q, \mu) &\leq q \cdot \mathbf{Adv}^{\text{auth}}_{\Pi}(t'', q, \mu) \end{aligned}$$

where $t' = t + O(\mu)$ and $t'' = t + O(q)$.

Proof: We begin with the first result. Let A gain $\delta = \mathbf{Adv}^{\text{auth}}_{\Pi}(A)$, run in time t and ask q queries, these totaling μ bits. Let B be an auth$*$ adversary for Π that runs A and answers all of its queries by querying its own left oracle (which is an $\mathcal{E}_K$ oracle). When A halts with its forgery attempt C, let B ask its right oracle C: if this is answered by INVALID, B outputs 0; otherwise it outputs 1. Clearly $\delta \leq \mathbf{Adv}^{\text{auth}*}_{\Pi}(B)$ and the resource of B are those claimed in the first part of lemma statement.

To prove the second result, let A gain $\delta = \mathbf{Adv}^{\text{auth}*}_{\Pi}(A)$, run in time t, ask $q_{\mathcal{E}}$ queries to its left oracle and $q_{\mathcal{D}}$ queries of its right oracle, these totaling μ bits. Let $q = q_{\mathcal{E}} + q_{\mathcal{D}}$. We construct an adversary B for attacking Π in the auth sense; this adversary is implicitly parameterized by an integer $j > 0$. Let B run A, answering left-oracle queries with its own $\mathcal{E}_K$ oracle. Adversary B will maintain a counter (initialized to 0) of A's right-oracle queries. When the counter is less than j, A's right-oracle queries are answered with INVALID. When A makes its jth right-oracle query C, B outputs C as its forgery. To finish the specification of B, it remains to fix a value for j, which we do now.

Consider the advantage δ gained by A. Let E be the event that A asks at least one valid right-oracle query C (i.e., $\mathcal{D}_K(C) \neq$ INVALID) during its execution. We can write

$$\delta \ = \ \Big(\Pr[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \wedge \mathsf{E}] + \Pr[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \wedge \overline{\mathsf{E}}]\Big) -$$

$$\begin{aligned}
&\Big(\Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \wedge \mathsf{E}] + \Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \wedge \overline{\mathsf{E}}]\Big) \\
= \quad &\Big(\Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \wedge \mathsf{E}] - \Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \wedge \mathsf{E}]\Big) + \\
&\Big(\Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \wedge \overline{\mathsf{E}}] - \Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \wedge \overline{\mathsf{E}}]\Big)
\end{aligned}$$

Notice that if event E does not occur, then all right oracle queries are answered with INVALID whether A had been provided a $\mathcal{D}_K$ oracle or a $\perp$ oracle. Thus, the second addend above is zero. Continuing,

$$\begin{aligned}
\delta \quad &= \quad \Big(\Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \wedge \mathsf{E}] - \Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \wedge \mathsf{E}]\Big) \\
&= \quad \Big(\Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\mathcal{D}_K(\cdot)} \Rightarrow 1 \mid \mathsf{E}] - \Pr[K \xleftarrow{\$} \mathcal{K} :\ A^{\mathcal{E}_K(\cdot),\perp(\cdot)} \Rightarrow 1 \mid \mathsf{E}]\Big) \Pr[\mathsf{E}] \\
&\leq \quad \Pr[\mathsf{E}]
\end{aligned}$$

Let E_j, $j \in [1..q_{\mathcal{D}}]$, be the event that E occurs on the jth right-oracle query. Then $\Pr[\mathsf{E}] = \Pr[\mathsf{E}_1 \vee \cdots \vee \mathsf{E}_{q_{\mathcal{D}}}] \leq \sum_{j=1}^{q_{\mathcal{D}}} \Pr[\mathsf{E}_j]$. It must be the case that $\Pr[\mathsf{E}_j] \geq \delta/q_{\mathcal{D}}$ for some j: fix this value of j in the description of B. Thus, for this value of j

$$\begin{aligned}
\delta \quad &\leq \quad q_{\mathcal{D}} \cdot \Pr[\mathsf{E}_j] \\
&\leq \quad q_{\mathcal{D}} \cdot \Pr[B \text{ forges}] \\
&= \quad q_{\mathcal{D}} \cdot \mathbf{Adv}^{\text{auth}}_{\Pi}(B) \\
&\leq \quad q \cdot \mathbf{Adv}^{\text{auth}}_{\Pi}(B)
\end{aligned}$$

It is easy to verify that B uses at most the resources claimed in the second result. ∎

References

[1] M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Advances in Cryptology— Asiacrypt '00*, Lecture Notes in Computer Science, vol. 1976, T. Okamoto, ed., Springer-Verlag, 2000.

[2] M. Bellare and P. Rogaway. Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography. *Advances in Cryptology— Asiacrypt '00*. Lecture Notes in Computer Science, vol. 1976, T. Okamoto, ed., Springer-Verlag, 2000.

[3] M. Bellare, A. Desai, E. Jokipii and P. Rogaway. A concrete security treatment of symmetric encryption: Analysis of the DES modes of operation. *Proceedings of 38th Annual Symposium on Foundations of Computer Science* (FOCS 97), IEEE, 1997.

[4] J. Katz and M. Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. *Fast Software Encryption* (FSE 2000), Lecture Notes in Computer Science, vol 1978, B. Schneier, ed., Springer, pp. 284–299, 2001.

[5] C. Rackoff and D. Simon Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. *Advances in Cryptology— CRYPTO '91*, Lecture Notes in Computer Science, vol. 576, J. Feigenbaum, ed., Springer-Verlag, 1991