Oil and water
Hybrid encryption techniques

Foundations of Cryptography
Computer Science Department
Wellesley College

Fall 2016

Table of contents

Encrypting arbitrary-length messages: Efficiency issues

- Last time we showed how any CPA-secure ℓ-bit public-key encryption scheme can be used to obtain a CPA-secure encryption scheme for messages of arbitrary length. This works, but is hopelessly inefficient.
- Encrypting an ℓ-bit message requires $\gamma \stackrel{\text{def}}{=} \lceil \ell/\ell' \rceil$ invocations of the original scheme.
- We can do better for messages that are sufficiently long, by using private-key encryption in tandem with public-key encryption.

Hybrid encryption

1. The sender chooses a random secret key k, and encrypts k using the public key of the receiver. The resulting ciphertext, c_1 is sent to the receiver establishing a shared secret between the two.
2. The sender then encrypts the message m using a *private-key encryption scheme* and the secret key k. The ciphertext, c_2 is sent to the receiver and recovered using the k.

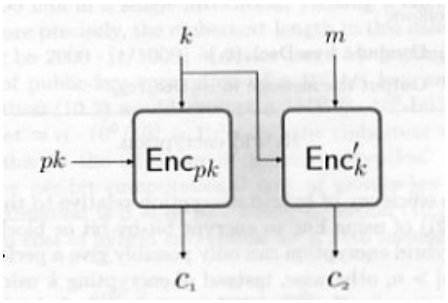

Key-encapsulation mechanisms (KEM)

- A more direct approach is to use a public-key primitive called a *key-encapsulation mechanism* (KEM) to accomplish both of these "in one shot."
- This is conceptually cleaner and more efficient in the bargain.

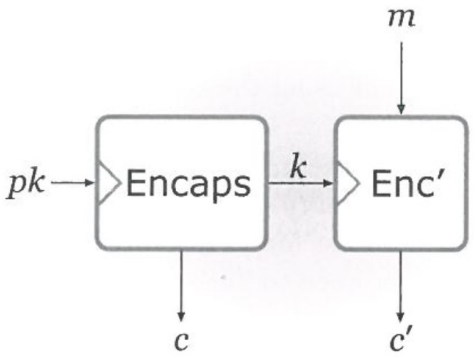

KEM: A formal introduction

Definition 11.9.
A *key-encapsulation mechanism* (KEM) is a tuple of probabilistic polynomial-time algorithms (Gen, Encaps, Decaps) such that

1. The *key-generation algorithm Gen* takes as input the security parameter 1^n and outputs a public-/private-key pair (pk, sk) whose lengths are at least n.
2. The deterministic *Encapsulation algorithm Encaps* takes as input a public key pk and the security parameter 1^n. It outputs a ciphertext c and a key $k \in \{0,1\}^{\ell(n)}$ where ℓ is the *key length*. We write $(c, k) \to \text{Encaps}_{pk}(1^n)$.
3. The *decapsulation algorithm Decaps* takes as input a private key sk and a ciphertext c, and outputs a key k or a special symbol $\perp$ denoting failure. We write $k := \text{Decaps}_{sk}(c)$.

It is required that with all by negligible probability over (sk, pk) output by $\text{Gen}(1^n)$, if $\text{Encap}_{pk}(1^n)$ outputs (c, k), then $\text{Decaps}_{sk}(c)$ outputs k.

Data-encapsulation mechanisms (DEM)

- We implement hybrid encryption using KEM The sender runs $\text{Encap}_{pk}(1^n)$ to obtain c with a key k; it uses a private-key encryption scheme to encrypt its message m.
- The private-key encryption scheme used here is calla *data-encapsulation mechanism* (DEM).
- The ciphertext sent to the receiver includes both c and c'.

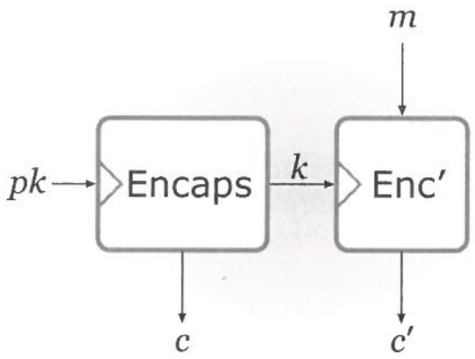

Construction of a hybrid encryption scheme using KEM/DEM

Construction 11.10

Let $\Pi = (\text{Gen}, \text{Encaps}, \text{Decaps})$ be a KEM with key length n, and let $\Pi' = (\text{Gen}', \text{Enc}', \text{Dec}')$ be a private-key encryption scheme. Construct a public-key encryption scheme $\Pi^{\text{hy}} = (\text{Gen}^{\text{hy}}, \text{Enc}^{\text{hy}}, \text{Dec}^{\text{hy}})$ as follows:

- Gen^{hy}: On input 1^n run $\text{Gen}(1^n)$ and use the public and private keys (pk, sk) that are output.
- Enc^{hy}: On input a public key pk and a message $m \in \{0,1\}^*$, proceed as follows:
  1. Compute $(c, k) \leftarrow \text{Encaps}_{pk}(1^n)$.
  2. Compute $c' \leftarrow \text{Enc}'_k(m)$.
  3. Output the ciphertext $\langle c, c' \rangle$.
- Dec^{hy}: On input a private key sk and a ciphertext $\langle c, c' \rangle$ do:
  1. Compute $k := \text{Decaps}_{sk}(c)$.
  2. Output the message $m := \text{Dec}'_k(c')$.

Efficiency of the hybrid encryption scheme

Remark. If $|m| < |n|$ we might as well use pk to encrypt m directly. However, when $|m| \gg |n|$ the hybrid scheme gives a substantial improvement assuming Enc' is more efficient than Enc.

Analysis. For fixed n, let α denote the cost of encrypting an n-bit key using Encaps, and let β denote the per bit cost of encryption using Enc'. Then per bit cost of plaintext using Π^{hy} is

$$\frac{\alpha + \beta \cdot |m|}{|m|} = \frac{\alpha}{|m|} + \beta.$$

which approaches β for sufficiently long m.

*A similar analysis can be used to calculate hybrid encryption ciphtertext length.

Security of hybrid encryption

We will show that

- If Π is CPA-secure KEM and the private-key scheme Π' has indistinguishable encryptions in the presence of an eavesdropper, then it turns out the hybrid scheme is Π^{hy} is CPA-secure.
- The fact the Π' need only have indistinguishable encryptions in the presence of an eavesdropper means weaker, but more efficient, schemes (such as stream ciphers) can be used on the private-key side.

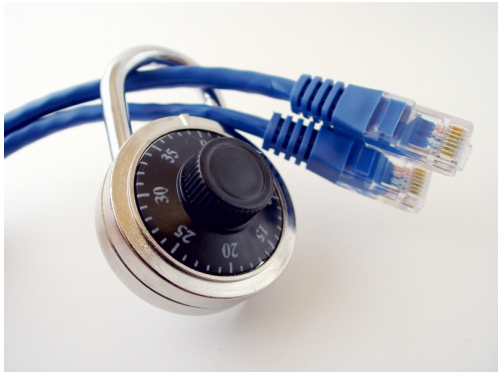

CPA-secure key-encapsulation mechanisms

The CPA indistinguishability experiment $\mathsf{KEM}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n)$*:*

1. $\mathsf{Gen}(1^n)$ is run to obtain keys (pk, sk). Then $\mathsf{Encaps}_{pk}(1^n)$ is run to generate (c, k) with $k \in \{0,1\}^n$.
2. A uniform bit $b \leftarrow \{0,1\}$ is chosen. if $b = 0$ set $\hat{k} := k$. If $b = 1$ then choose a uniform $\hat{k} \in \{0,1\}^n$.
3. Given $(pk, c, \hat{k})$ to $\mathcal{A}$ who outputs a bit b'. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise.

Definition 11.11. A key-encapsulation mechanism Π is *CPA-secure* if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there exists a negligible function negl such that

$$\Pr[\mathsf{KEM}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n).$$

Theorem and proof goals

Theorem 11.12. If Π is a CPA-secure KEM and Π' is a private-key encryption scheme that has indistinguishable encryptions in the presence of an eavesdropper, then Π^{hy} as in Construction 11.10 is a CPA-secure public-key encryption scheme.

Proof goals.

Let $\mathsf{Encaps}^{(1)}_{pk}(1^n)$ (resp., $\mathsf{Encaps}^{(2)}_{pk}(1^n)$) denote the ciphertext (resp. key) output by Encaps. The fact that Π is CPA-secure means that

$$\left(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Encaps}^{(2)}_{pk}(1^n)\right) \overset{\mathsf{c}}{\equiv} \left(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), k'\right)$$

Similarly Π' has indistinguishable encryption means that for any m_0, m_1 output by $\mathcal{A}$, $\mathsf{Enc'}_k(m_0) \overset{\mathsf{c}}{\equiv} \mathsf{Enc'}_k(m_1)$. To prove CPA-security of Π^{hy} we show

$$\left(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_k(m_0)\right) \overset{\mathsf{c}}{\equiv} \left(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_k(m_1)\right)$$

The proof proceeds in three steps

Theorem 11.12. If Π is a CPA-secure KEM and Π' is a private-key encryption scheme that has indistinguishable encryptions in the presence of an eavesdropper, then Π^{hy} as in Construction 11.10 is a CPA-secure public-key encryption scheme.

Proof idea.

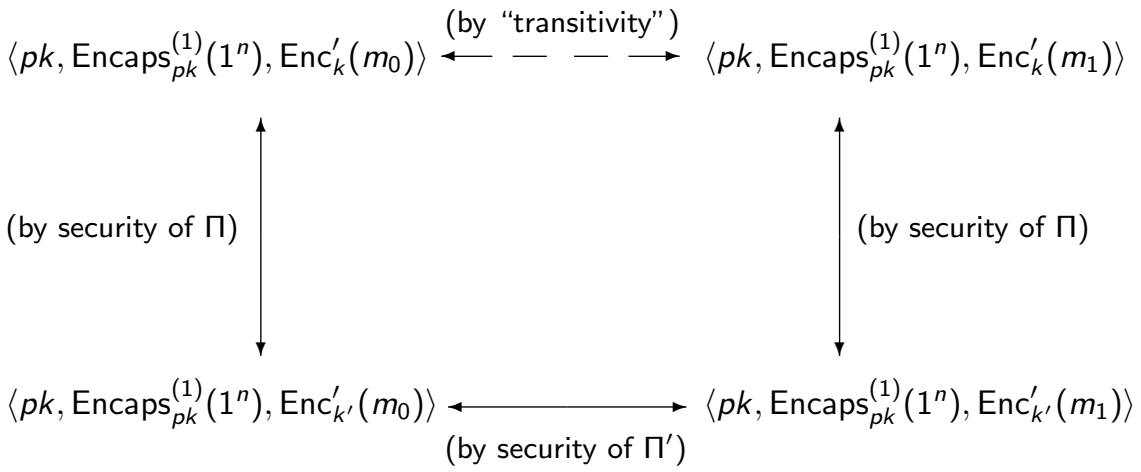

The theorem and proof

Theorem 11.12. If Π is a CPA-secure KEM and Π' is a private-key encryption scheme that has indistinguishable encryptions in the presence of an eavesdropper, then Π^{hy} as in Construction 11.10 is a CPA-secure public-key encryption scheme.

Proof. We show Π^{hy} has indistinguishable encryption in the presence of an eavesdropper and use Proposition 11.3.
Fix an arbitrary PPT $\mathcal{A}^{\text{hy}}$, and consider experiment $\text{PubK}^{\text{eav}}_{\mathcal{A}^{\text{hy}},\Pi^{\text{hy}}}(n)$. We show

$$\Pr[\text{PubK}^{\text{eav}}_{\mathcal{A}^{\text{hy}},\Pi^{\text{hy}}}(n) = 1] \leq \frac{1}{2} + \text{negl}(n).$$

By definition of the experiment,

$$\begin{aligned}\Pr[\text{PubK}^{\text{eav}}_{\mathcal{A}^{\text{hy}},\Pi^{\text{hy}}}(n) = 1] &= \frac{1}{2}\cdot \Pr[\mathcal{A}^{\text{hy}}(\text{Encaps}_{pk}^{(1)}(1^n), \text{Enc}'_k(m_0)) = 0] \\ &\quad + \frac{1}{2}\cdot \Pr[\mathcal{A}^{\text{hy}}(\text{Encaps}_{pk}^{(1)}(1^n), \text{Enc}'_k(m_1)) = 1]\end{aligned}$$

where $k = \text{Encaps}_{pk}^{(2)}(1^n)$.

First we establish the left-hand indistinguishability

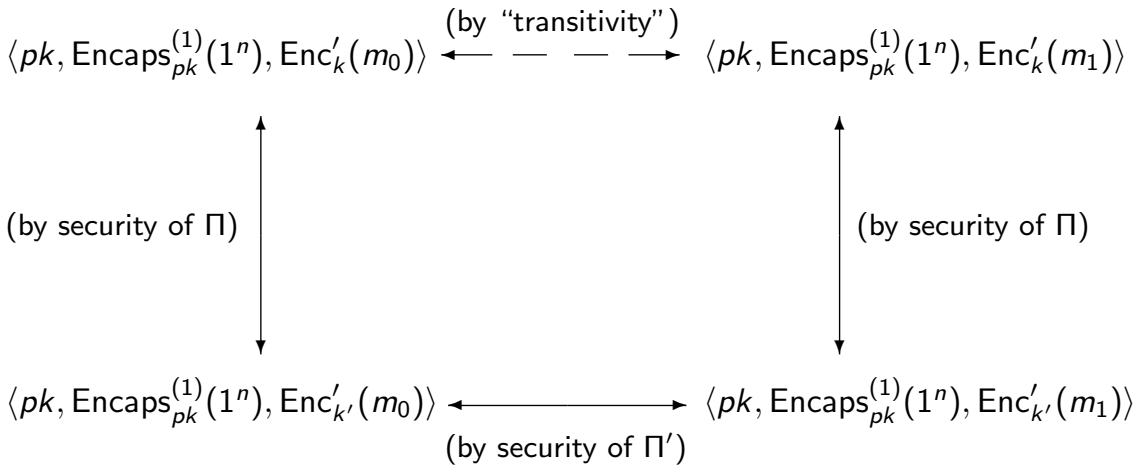

A public-key eavesdropping adversary

Adversary $\mathcal{A}_1$:

1. $\mathcal{A}_1$, given $(pk, c, \hat{k})$.
2. $\mathcal{A}_1$ runs $\mathcal{A}^{\text{hy}}(pk)$ to obtain two messages m_0, m_1. Then $\mathcal{A}_1$ computes $c' \leftarrow \text{Enc}_{\hat{k}}(m_0)$, then runs $\mathcal{A}^{\text{hy}}(c, c')$ and outputs the bit b' that is output by $\mathcal{A}^{\text{hy}}$.

When $b = 0$ in $\text{KEM}^{\text{cpa}}_{\mathcal{A}_1,\Pi}(n)$, $\mathcal{A}_1$ is given $(pk, c, \hat{k})$ where $c, \hat{k}$ were output by $\text{Encaps}^{(1)}_{pk}(1^n)$. So, $\mathcal{A}^{hy}$ is given $\langle c, \text{Enc}'_k(1^n)\rangle$ where k is encapped by c. In this case,

$$\Pr[\mathcal{A}_1 \text{ outputs } 0 \mid b = 0] = \Pr[\mathcal{A}^{\text{hy}}(\text{Encaps}^{(1)}_{pk}(1^n), \text{Enc}'_k(m_0)) = 0].$$

When $b = 1$ in $\text{KEM}^{\text{cpa}}_{\mathcal{A}_1,\Pi}(n)$ is given $(pk, c, \hat{k})$ where $\hat{k}$ is uniform and independent of c. So, $\mathcal{A}^{hy}$ is given $\langle c, \text{Enc}'_{k'}(1^n)\rangle$ where k' is encapped by c.

$$\Pr[\mathcal{A}_1 \text{ outputs } 1 \mid b = 1] = \Pr[\mathcal{A}^{\text{hy}}(\text{Encaps}^{(1)}_{pk}(0^n), \text{Enc}'_{k'}(m_0)) = 1].$$

Completing the left-hand argument

Since Π is a CPA-secure KEM, there exists a negligible function negl_1 such that:

$$\begin{aligned}\frac{1}{2}+\mathsf{negl}_1 &\geq \Pr[\mathsf{KEM}^{\mathsf{cpa}}_{\mathcal{A}_1,\Pi}(n)=1]\\ &= \frac{1}{2}\cdot\Pr[\mathcal{A}_1 \text{ outputs } 0 \mid b=0]+\frac{1}{2}\cdot\Pr[\mathcal{A}_1 \text{ outputs } 1 \mid b=1]\\ &= \frac{1}{2}\cdot\Pr[\mathcal{A}^{\mathsf{hy}}(\mathsf{Encaps}_{pk}(1^n),\mathsf{Enc}'_k(m_0))=0]\\ &\quad+\frac{1}{2}\cdot\Pr[\mathcal{A}^{\mathsf{hy}}(\mathsf{Encaps}_{pk}(1^n),\mathsf{Enc}'_{k'}(m_0))=1].\end{aligned}$$

where $K=\mathsf{Encaps}^{(2)}_{pk}(1^n)$ and k' is uniform and independent.

Right-hand indistinguishability is proven similarly

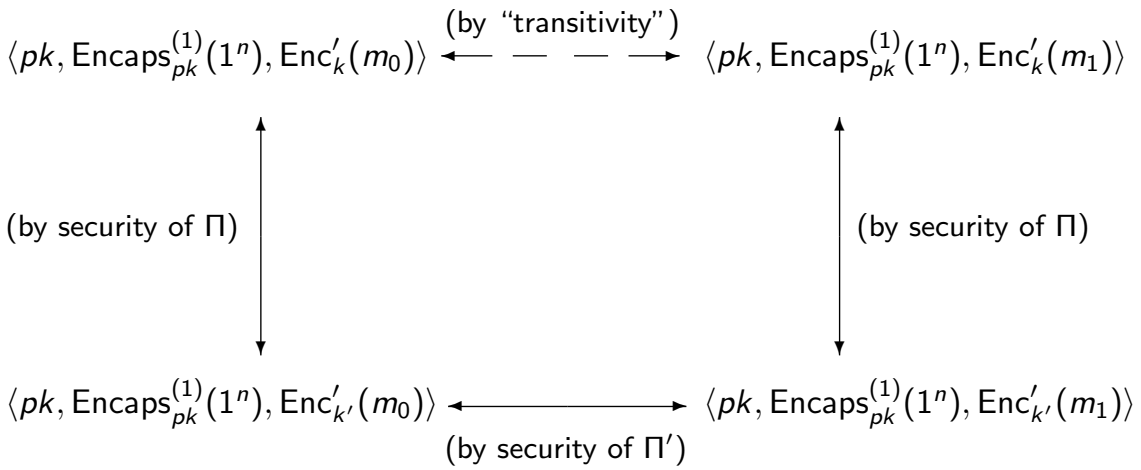

Results of the right-hand argument

For any PPT adversary $\mathcal{A}_2$ that eavesdrops on a message encrypted using public-key scheme Π. there exists a negligible function negl_2 such that:

$$\begin{aligned}\frac{1}{2}+\mathsf{negl}_1 &\geq \Pr[\mathsf{KEM}^{\mathsf{cpa}}_{\mathcal{A}_1,\Pi}(n)=1]\\ &= \frac{1}{2}\cdot\Pr[\mathcal{A}_2 \text{ outputs } 0 \mid b=0]+\frac{1}{2}\cdot\Pr[\mathcal{A}_2 \text{ outputs } 1 \mid b=\\ &= \frac{1}{2}\cdot\Pr[\mathcal{A}^{\mathsf{hy}}(pk,\mathsf{Encaps}_{pk}(1^n),\mathsf{Enc}'_k(m_1))=1]\\ &\quad+\frac{1}{2}\cdot\Pr[\mathcal{A}^{\mathsf{hy}}(pk,\mathsf{Encaps}_{pk}(1^n),\mathsf{Enc}'_{k'}(m_1))=0].\end{aligned}$$

Next we establish the bottom indistinguishability

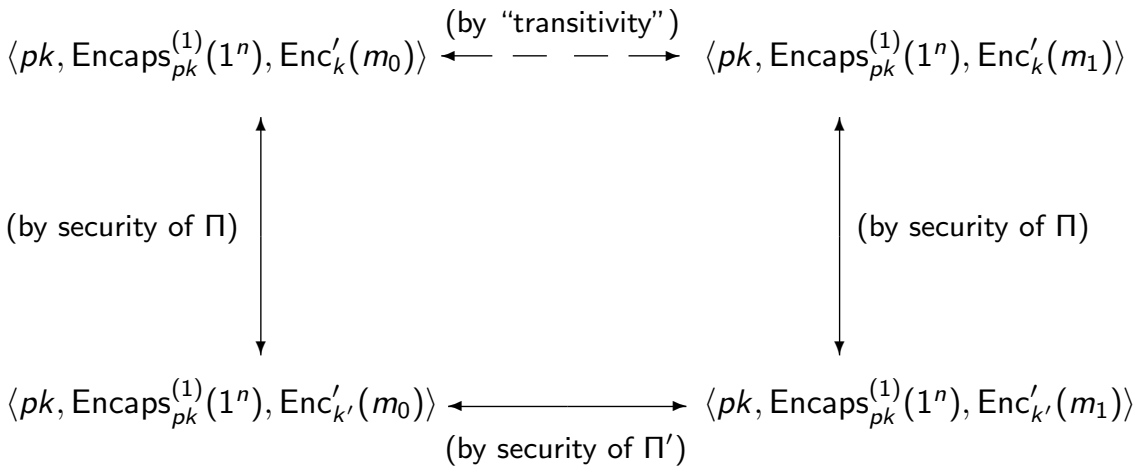

A private-key eavesdropping adversary

Adversary $\mathcal{A}'$:

1. $\mathcal{A}'$ runs $\mathsf{Gen}(1^n)$ on its own to generate keys (pk, sk). Also compute $c \leftarrow \mathsf{Encaps}^{(1)}_{pk}(1^n)$.
2. $\mathcal{A}'$ runs $\mathcal{A}^{\mathsf{hy}}(pk)$ to obtain two messages m_0, m_1 which it then outputs and receives in return a ciphertext c'.
3. $\mathcal{A}'$ runs $\mathcal{A}^{\mathsf{hy}}(c, c')$ and outputs the bit b' that is output by $\mathcal{A}^{\mathsf{hy}}$.

When $b = 0$ in $\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A}',\Pi'}(n)$, $\mathcal{A}'$ is given $\mathsf{Enc}'_k(m_0)$ where k' was chosen at random and independent of everything. In this case, $\mathcal{A}^{hy}$ is given a ciphertext of the form $\langle c, \mathsf{Enc}'_{k'}(m_0)\rangle$. So

$$\Pr[\mathcal{A}' \text{ outputs } 0 \mid b = 0] = \Pr[\mathcal{A}^{\mathsf{hy}}(\mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_k(m_0)) = 0].$$

When $b = 1$ in $\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A}',\Pi'}(n)$, $\mathcal{A}'$ is given $\mathsf{Enc}'_{k'}(m_1)$. In this case,

$$\Pr[\mathcal{A}_1 \text{ outputs } 1 \mid b = 1] = \Pr[\mathcal{A}^{\mathsf{hy}}(\mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_{k'}(m_1)) = 1].$$

Completing the bottom argument

Since Π' has indistinguishable encryption in the presence of an eavesdropper, there exists a negligible function negl' such that:

$$\begin{aligned}
\frac{1}{2} + \mathsf{negl}' &\geq \Pr[\mathsf{PubK}^{\mathsf{eav}}_{\mathcal{A}',\Pi'}(n) = 1] \\
&= \frac{1}{2} \cdot \Pr[\mathcal{A}' \text{ outputs } 0 \mid b = 0] + \frac{1}{2} \cdot \Pr[\mathcal{A}' \text{ outputs } 1 \mid b = 1] \\
&= \frac{1}{2} \cdot \Pr[\mathcal{A}^{hy}(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_{k'}(m_0)) = 0] \\
&\quad + \frac{1}{2} \cdot \Pr[\mathcal{A}^{hy}(pk, \mathsf{Encaps}^{(1)}_{pk}(1^n), \mathsf{Enc}'_{k'}(m_1)) = 1].
\end{aligned}$$

Completing the proof of Theorem 11.12

Summing these three inequalities and using the fact that the sum of three negligible functions is negligible, we have

$$\begin{aligned}\frac{3}{2} + \mathsf{negl} \geq \\ \frac{1}{2} \cdot (&\Pr[\mathcal{A}^{\mathrm{hy}}(pk, c, \mathsf{Enc}'_k(m_0)) = 0] + \Pr[\mathcal{A}^{\mathrm{hy}}(pk, c, \mathsf{Enc}'_{k'}(m_0)) = 1] \\ &+ \Pr[\mathcal{A}^{hy}(pk, c, \mathsf{Enc}'_{k'}(m_0)) = 0] + \Pr[\mathcal{A}^{hy}(pk, c, \mathsf{Enc}'_{k'}(m_1)) = 1] \\ &+ \Pr[\mathcal{A}^{hy}(pk, c, \mathsf{Enc}'_{k'}(m_1)) = 1] + \Pr[\mathcal{A}^{hy}(pk, c, \mathsf{Enc}'_{k'}(m_1)) = 0])\end{aligned}$$

Using the fact that several of these probabilities sum to 1, we have

$$\begin{aligned}&\frac{1}{2} + \mathsf{negl} \\ &\quad \geq \frac{1}{2} \cdot (\Pr[\mathcal{A}^{\mathrm{hy}}(\mathsf{Enc}_{pk}(k), \mathsf{Enc}'_k(m_0)) = 0] + \Pr[\mathcal{A}^{\mathrm{hy}}(\mathsf{Enc}_{pk}(k), \mathsf{Enc}'_k(m_1)) = 1]) \\ &\quad = \Pr[\mathsf{PubK}^{\mathit{sfeav}}_{\mathcal{A}^{hy}, \Pi^{hy}}(n) = 1].\end{aligned}$$

□