Public-Key Encryption in a Multi-User Setting: Privacy, Anonymity and Efficiency

Alexandra Boldyreva

Georgia Institute of Technology

Plan

- Encryption, a tool for data privacy
- Provable security
- Standard definitions of data privacy
- The need to consider the multi-user setting
  - Security
  - Efficiency
  - Anonymity

Encryption is

- a tool for achieving data-privacy,
- is very important nowadays,
- used by many people, often without realizing it, when:
  - doing on-line banking and shopping
  - talking on cell phones
  - watching satellite TV and pay-per-view movies

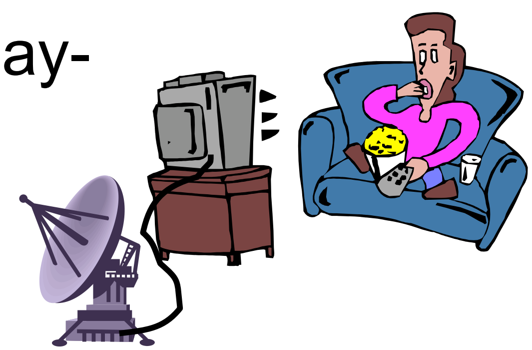

On-line shopping, banking rely on encryption

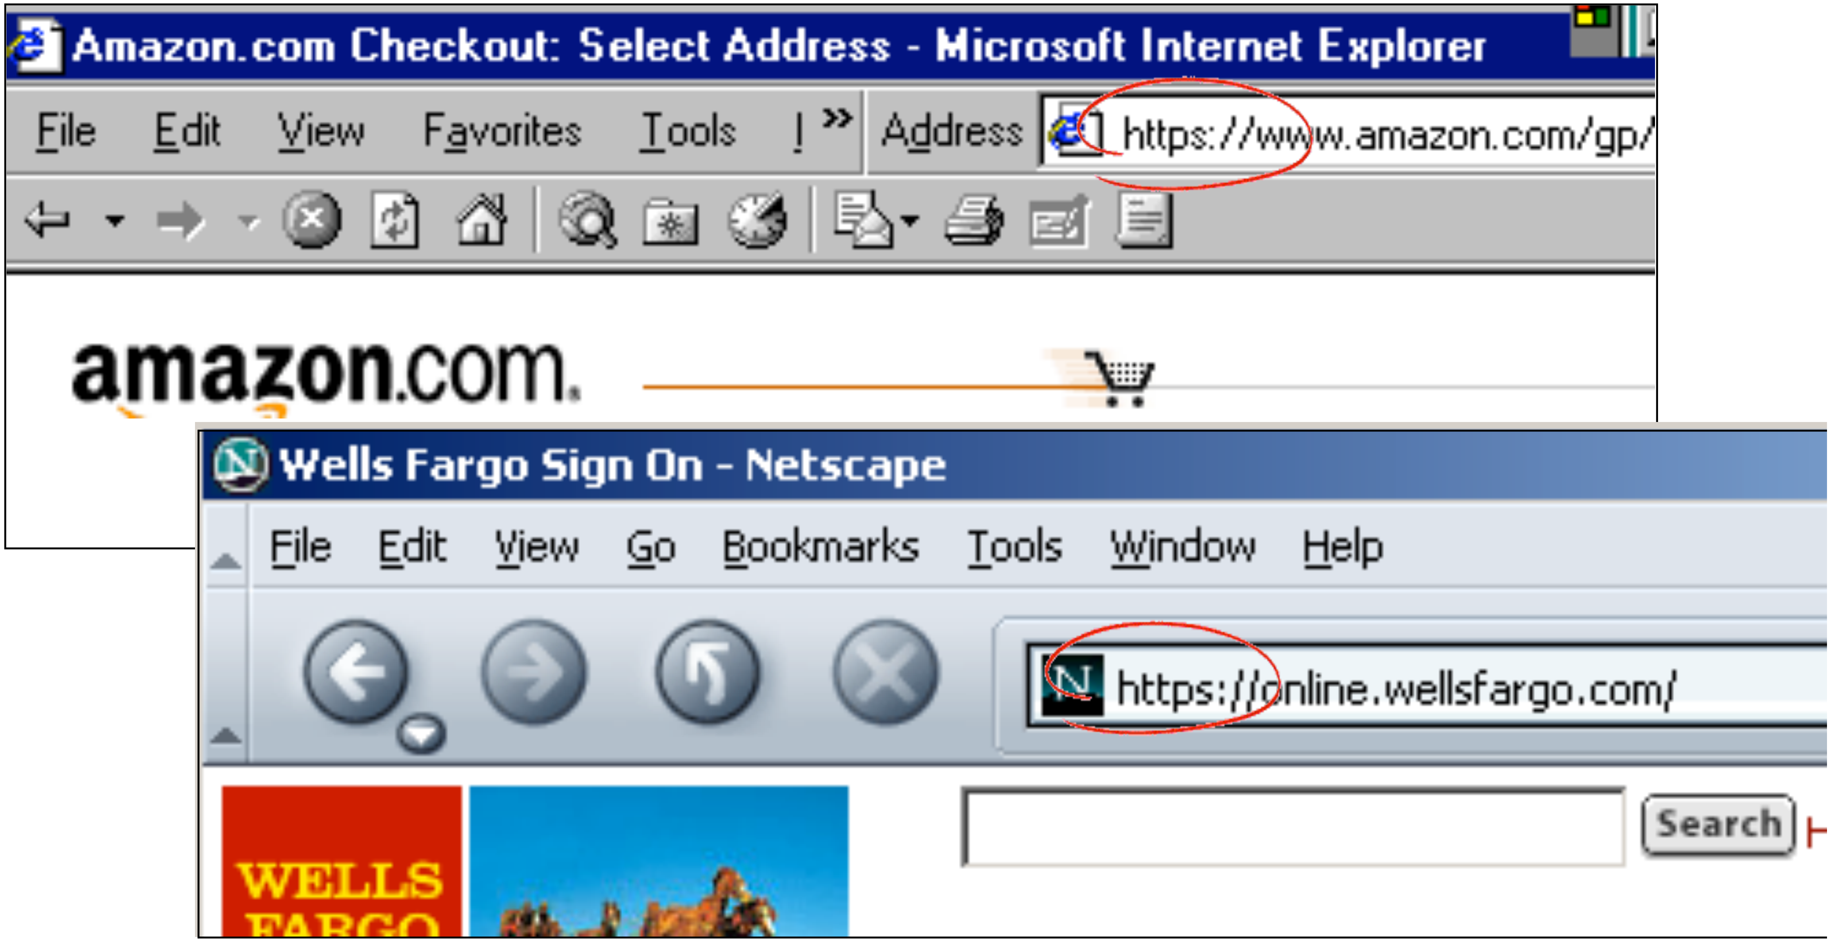

SSL protocol ensures privacy of communicated data (uses RSA-OAEP encryption scheme [BR])

Two settings

1. Symmetric-key setting

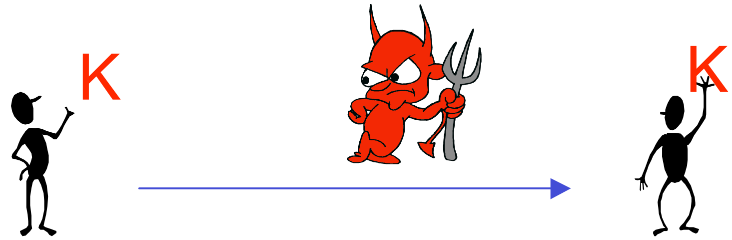

2. Asymmetric (public-key) setting

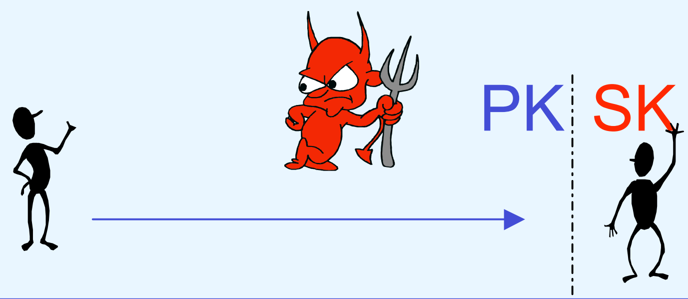

Public-key encryption

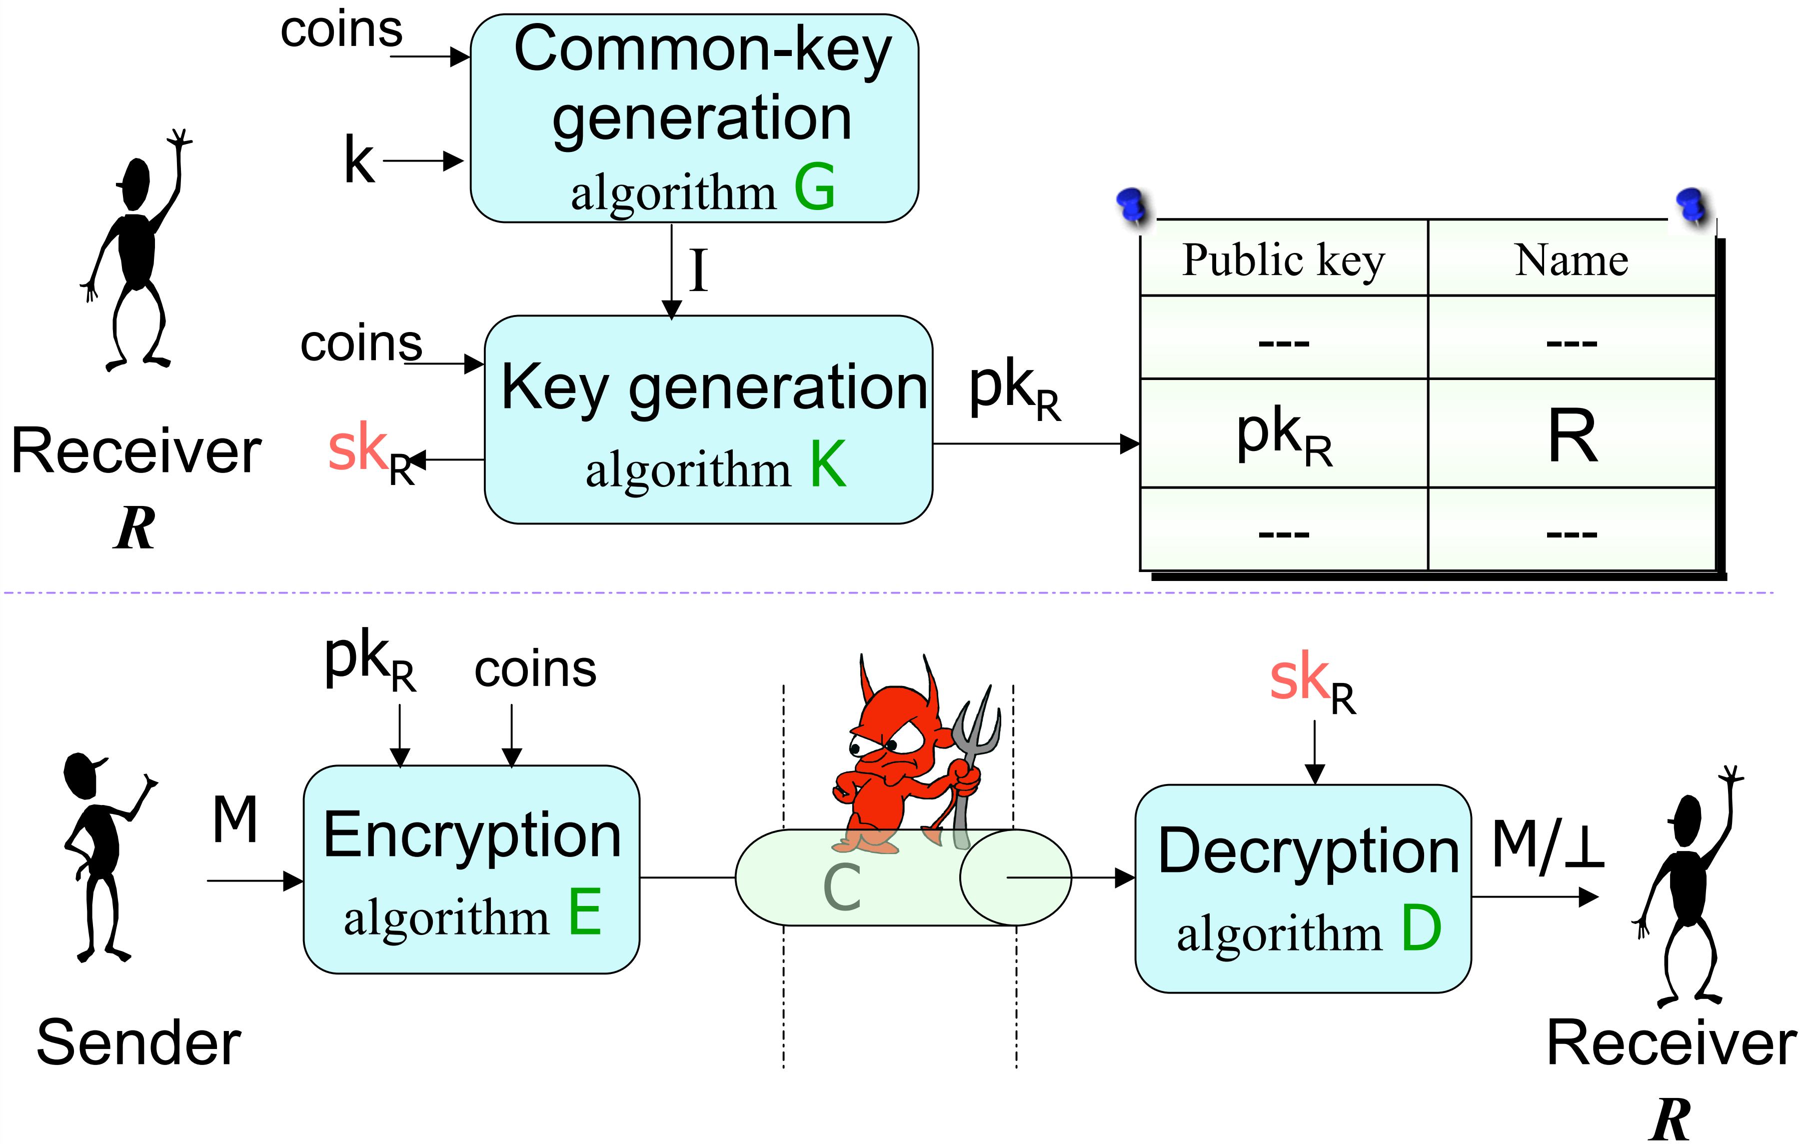

How can we be confident that a given encryption scheme is secure?

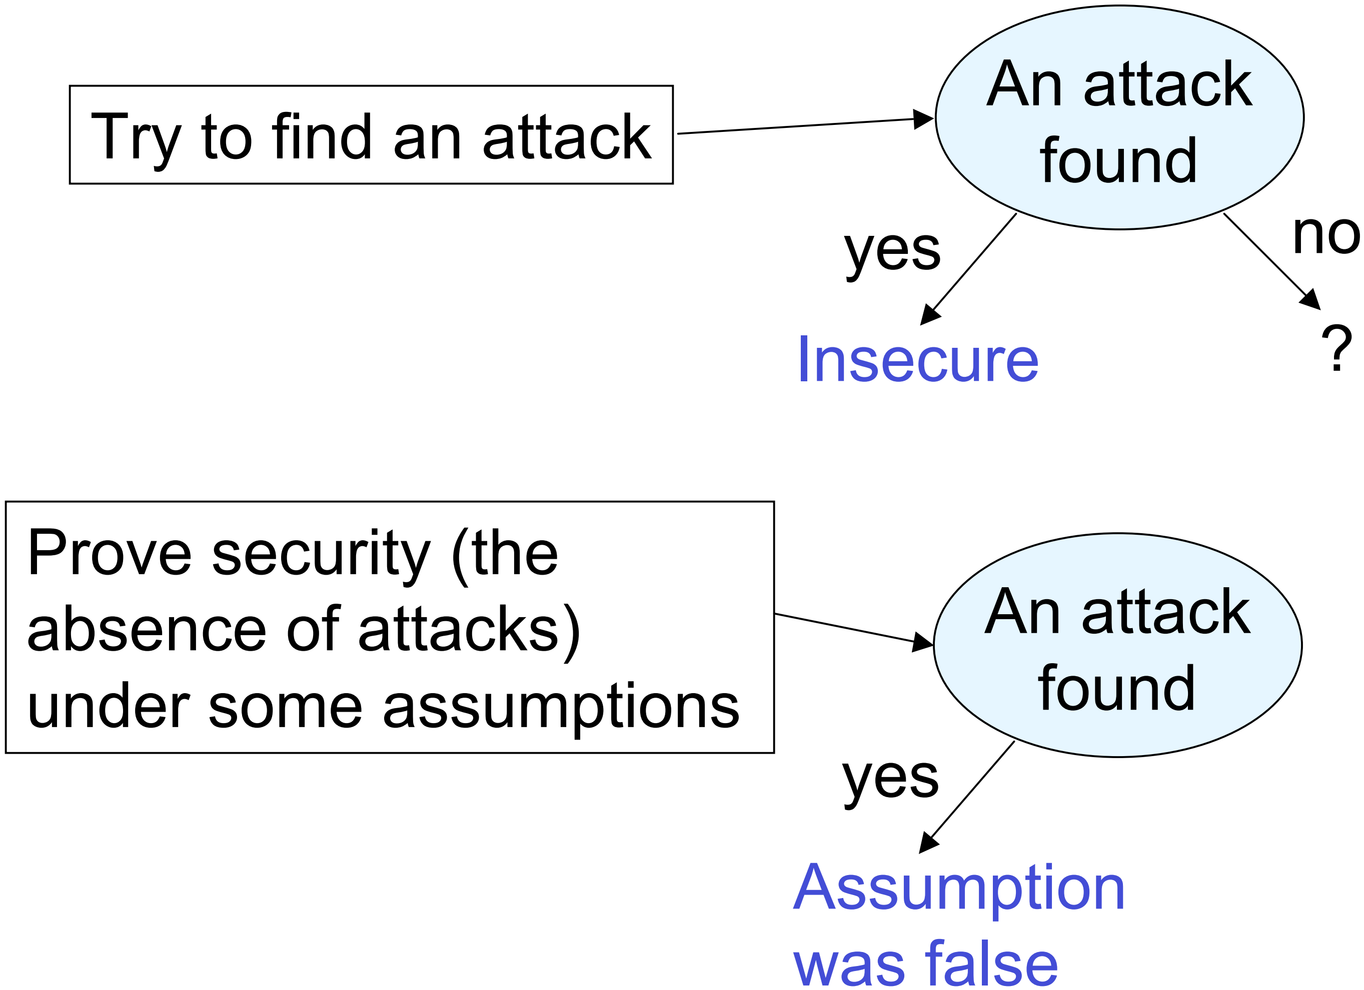

What does “security” mean?

Security means that given a public key and a ciphertext it is infeasible to recover:	But…
the secret key	Can be true if the plaintext is sent in the clear
the plaintext	Can be true if all but the last bit are leaked
etc.	etc.
Any partial information about the plaintext	

Encryption security definition, IND-CPA [GM]

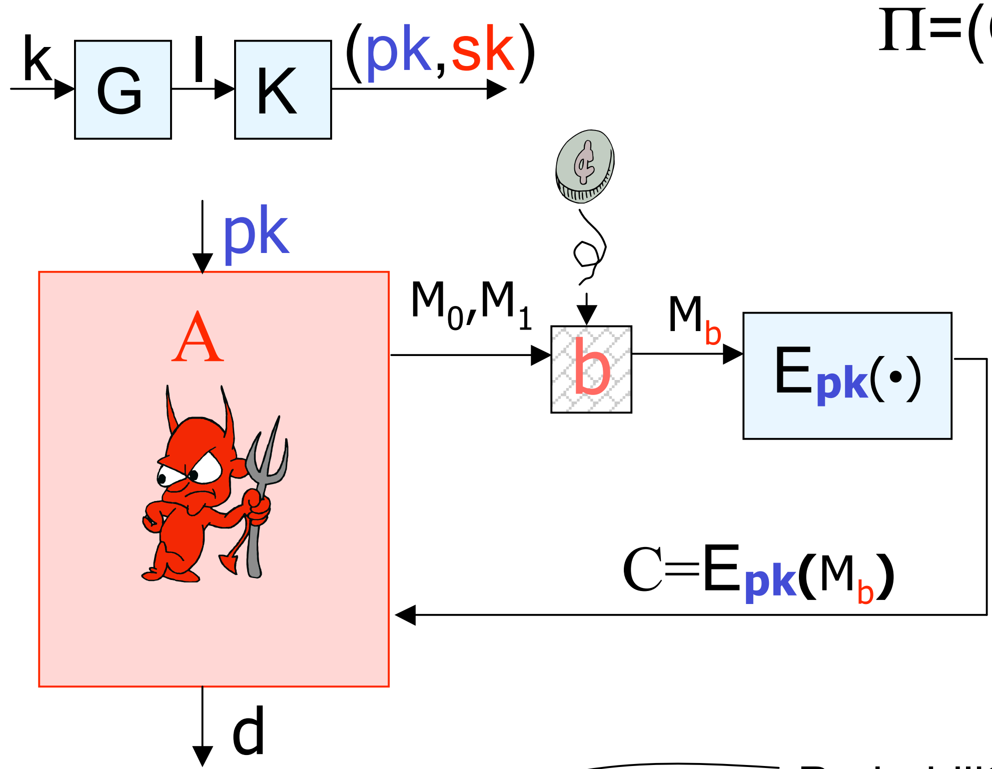

$\Pi=(G,K,E,D)$

Adversary wins if b=d

$\mathrm{Adv}^{1\text{-ind-cpa}}_{A,\Pi}(k)=2\Pr[\mathrm{win}]-1$

Probability is over coins of G,K,E,A and choice of b.

An encryption scheme Π is IND-CPA in the single user setting if for any PPT adversary A, $\mathrm{Adv}^{1\text{-ind-cpa}}_{A,\Pi}(k)$ is negligible in k.

Why IND-CPA?

The definition guarantees that the secret key, plaintexts, or any partial information about the plaintexts are not leaked.

IND-CPA is not always enough

Bleichenbacher's attack on a previous version of SSL:

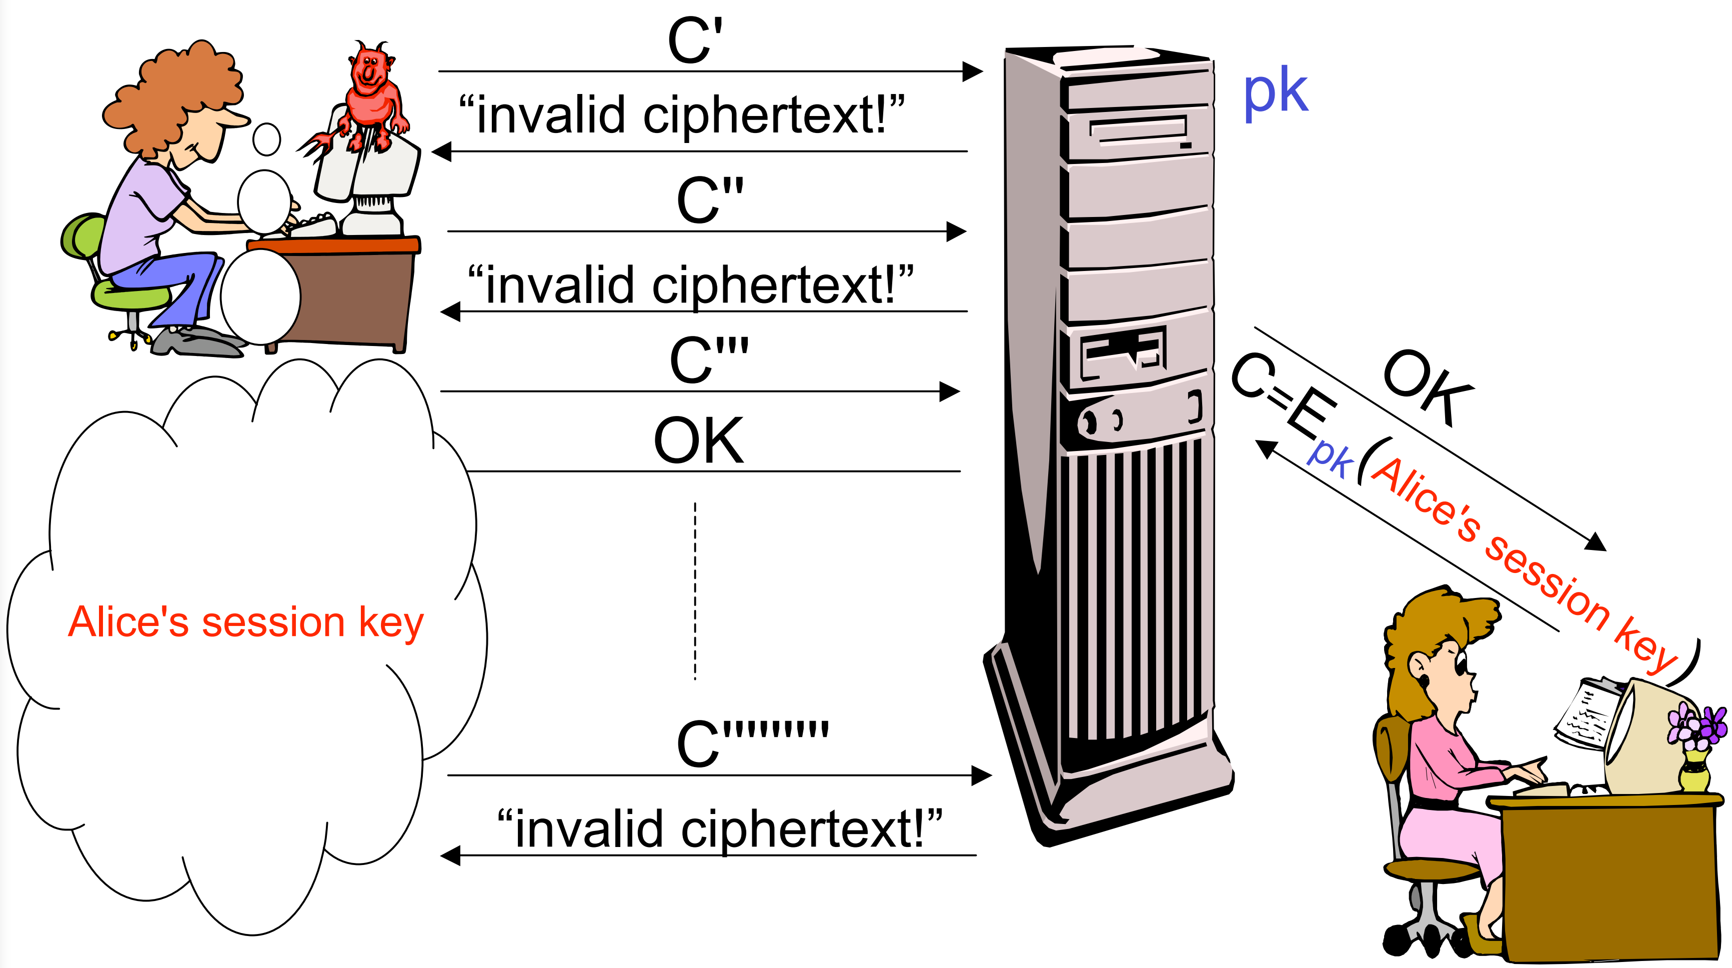

Encryption security definition, IND-CCA

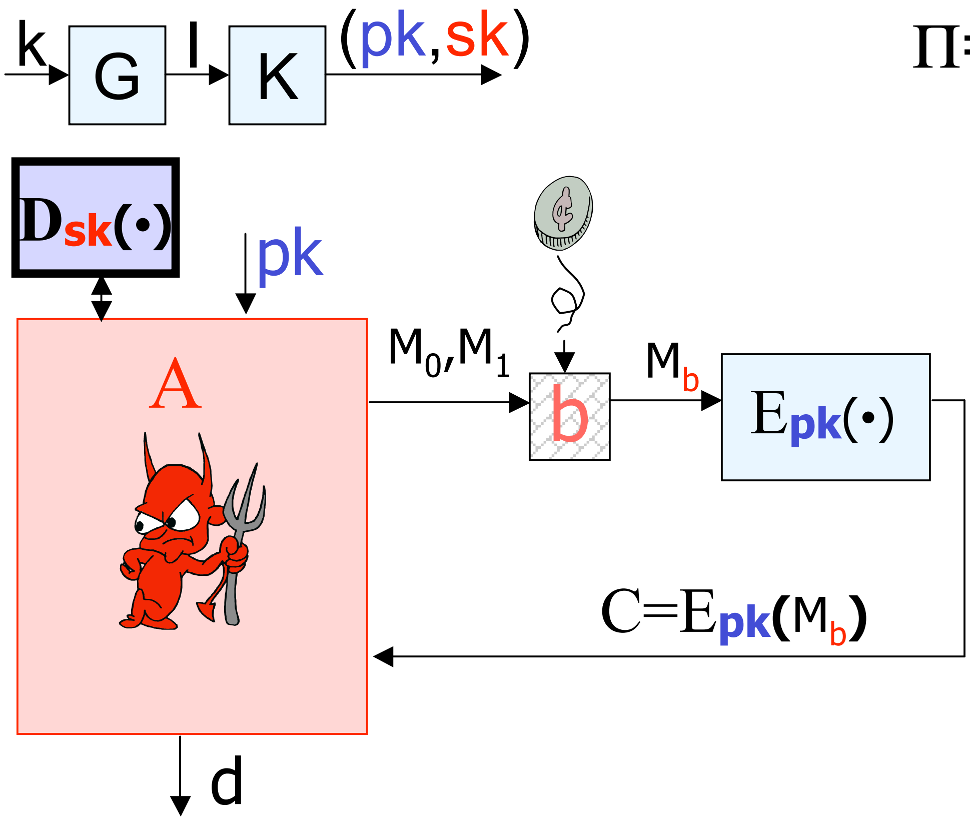

$\Pi=(G,K,E,D)$

A is not allowed to query C to $D_{sk}(\cdot)$

A wins if b=d

$Adv^{1\text{-}ind\text{-}cca}_{A,\Pi}(k)=2Pr[win]-1$

An encryption scheme Π is IND-CCA in the single user setting if for any PPT adversary A, $Adv^{1\text{-}ind\text{-}cca}_{A,\Pi}(k)$ is negligible in k.

Proven secure schemes

Scheme	Security	Proven assuming	Usage
ElGamal	IND-CPA	Decision Diffie-Hellman (DDH)	
Cramer-Shoup	IND-CCA	DDH	
RSA-OAEP [BR]	IND-CCA	One-wayness of RSA, RO	PKCS #1 2.1
DHIES [ABR]	IND-CCA	ODH	IEEE P1363a

Data-privacy in the multi-user setting

Motivation

All provably-secure encryption schemes are proven secure in the single-<u>user</u> setting

Person with a public key, able to **receive** ciphertexts

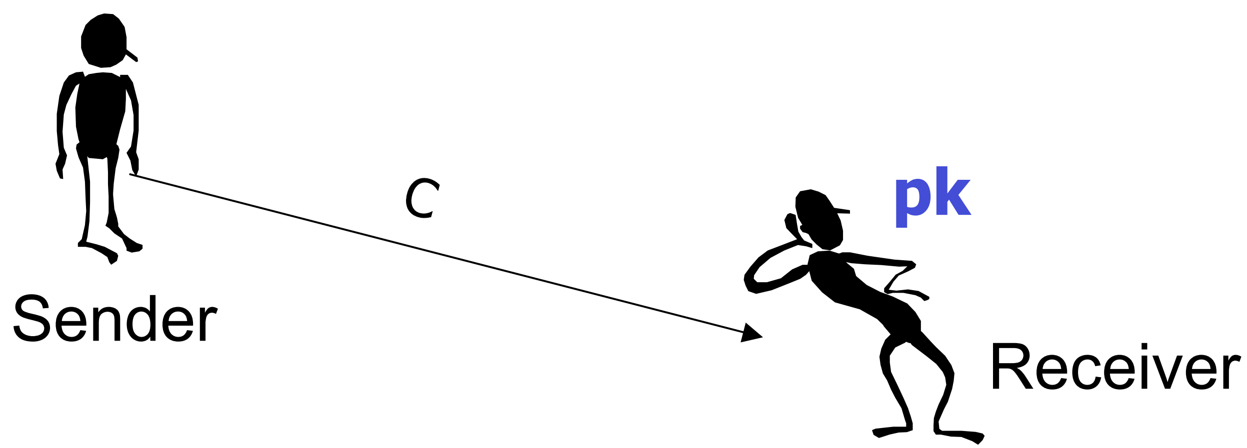

All ciphertexts seen by the adversary are under a single public key

But the single-user setting is very different from practice, where there are many users sending each other encrypted messages:

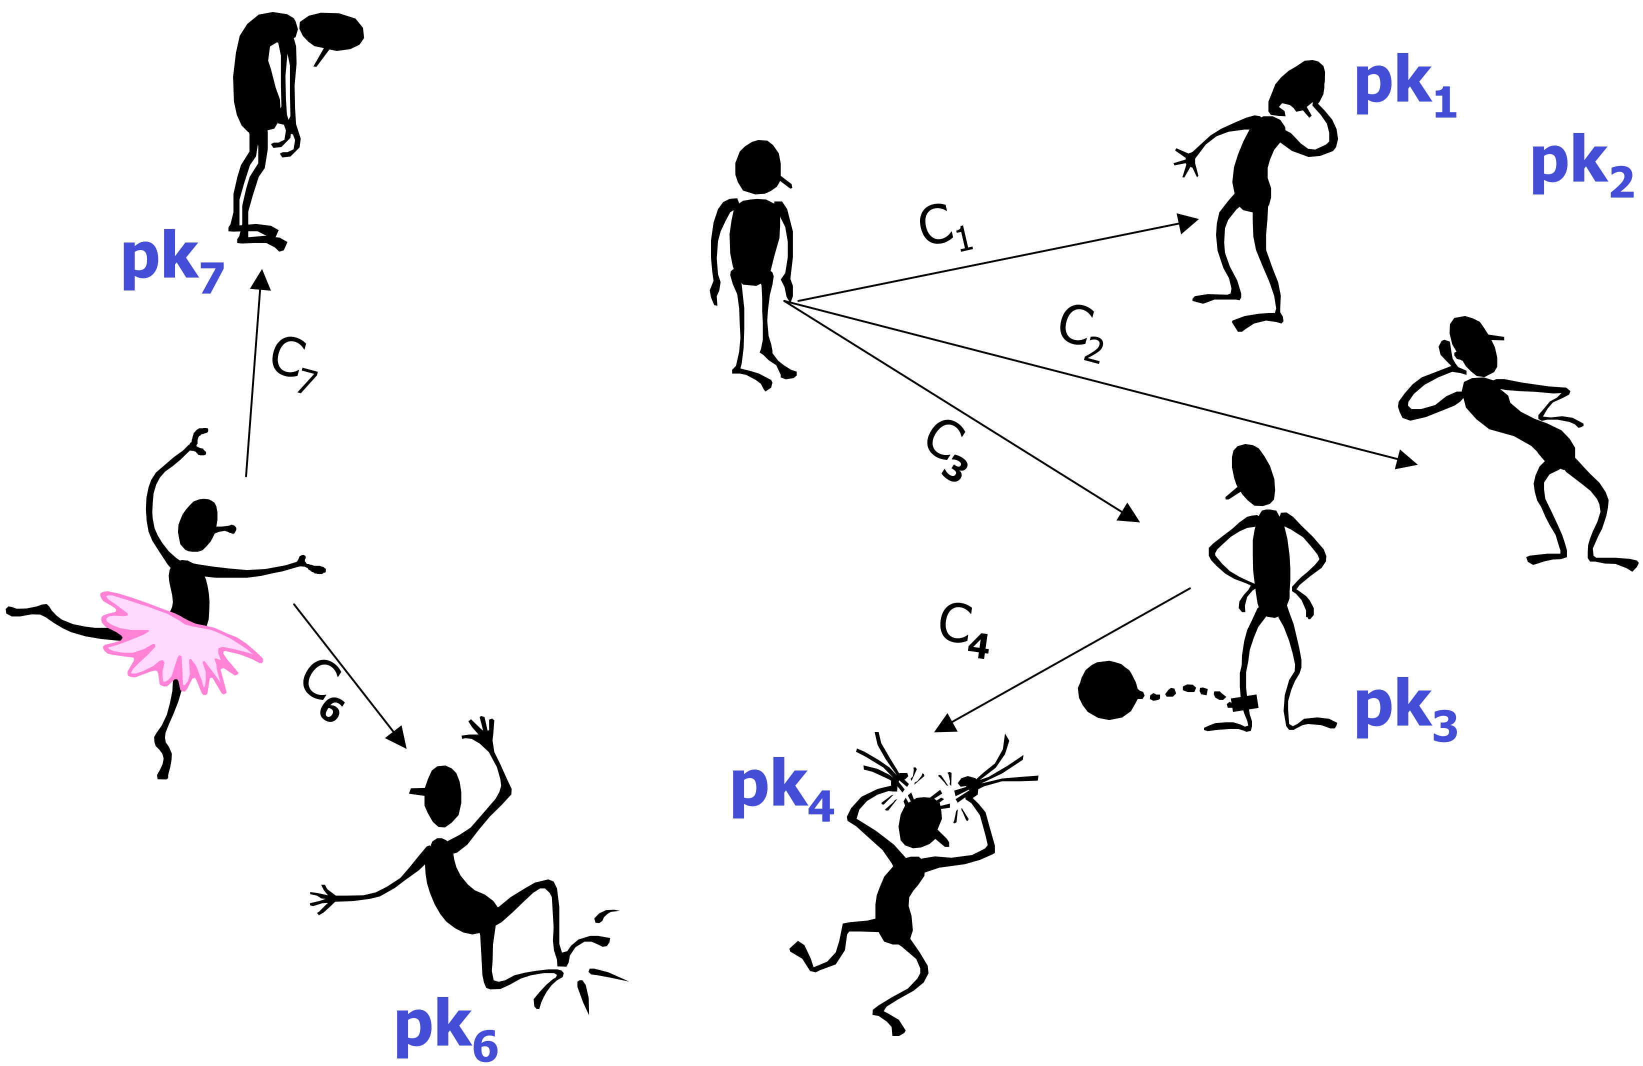

Plain RSA encryption [RSA]

$G(k)$ $k \in \mathbb{N}$	$K(k)$	$E_{pk}(M)$ $M \in Z_N^*$	$D_{sk}(C)$
Return k	$p,q \stackrel{\$}{\leftarrow}$ k-bit primes $N \leftarrow p \cdot q$ $e \stackrel{\$}{\leftarrow} Z^*_{(p-1)(q-1)}$ $d \leftarrow Z^*_{(p-1)(q-1)}$ s.t. $e \cdot d \equiv 1 \bmod (p-1)(q-1)$ $pk \leftarrow (N, e)$ $sk \leftarrow (N, d)$ Return (pk, sk)	$C \leftarrow M^e \bmod N$ Return C	$M \leftarrow C^d \bmod N$ Return M

Believed to be one-way:

easy

$C^d = M$ → $C = M^e$

easy given d

hard not given d

Håstad-type attack on Plain RSA

Plain RSA: $pk=(N, e)$; $E_{pk}(M)=M^e \bmod N$

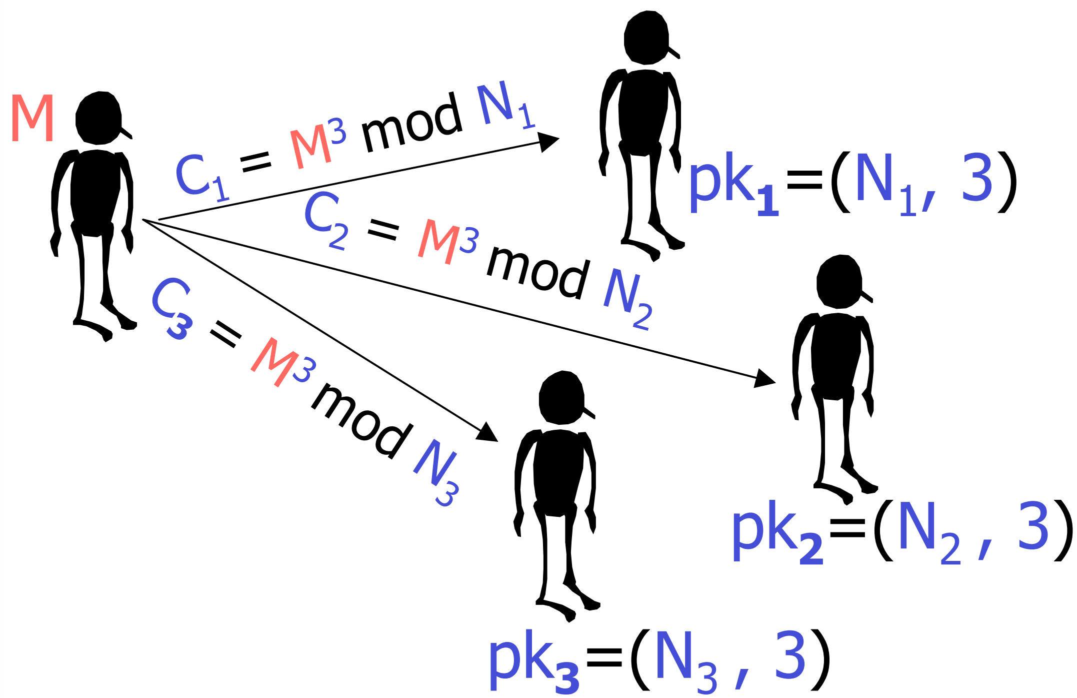

$C_1, C_2, C_3, pk_1, pk_2, pk_3$

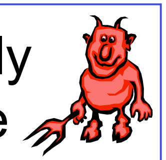

If N_1,N_2,N_3 are relatively prime then by Chinese Remainder Theorem can combine

$$\begin{cases} C_1 = M^3 \bmod N_1 \\ C_2 = M^3 \bmod N_2 \\ C_3 = M^3 \bmod N_3 \end{cases}$$

to find $C = M^3 \bmod N_1 N_2 N_3$

Since $M^3 < N_1 N_2 N_3$ then

$M \leftarrow \sqrt[3]{C}$

- Plain RSA:
  - Is one-way in the single-user setting.
  - Is not one-way in the multi-user setting.
  - However, it is not IND-CPA in the single-user setting.

Plain RSA is not IND-CPA secure

(as well as any deterministic scheme Π)

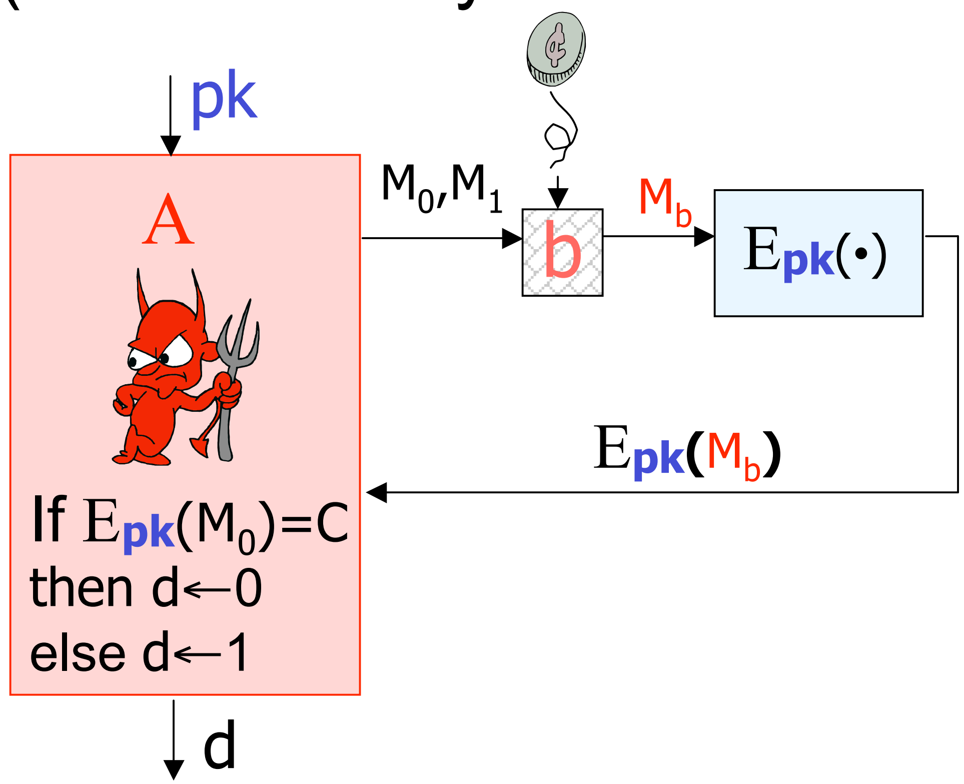

A always wins,

$$\mathrm{Adv}^{1}(A) = 1 \quad \mathrm{Adv}^{1}(t) = 1$$

A crucial question

Are the “provably-secure” schemes (e.g. ElGamal, RSA-OAEP) really secure in the practical (multi-user) setting?

To answer this one needs to define security in the multi-user setting

Towards a security definition for encryption in the multi-user setting

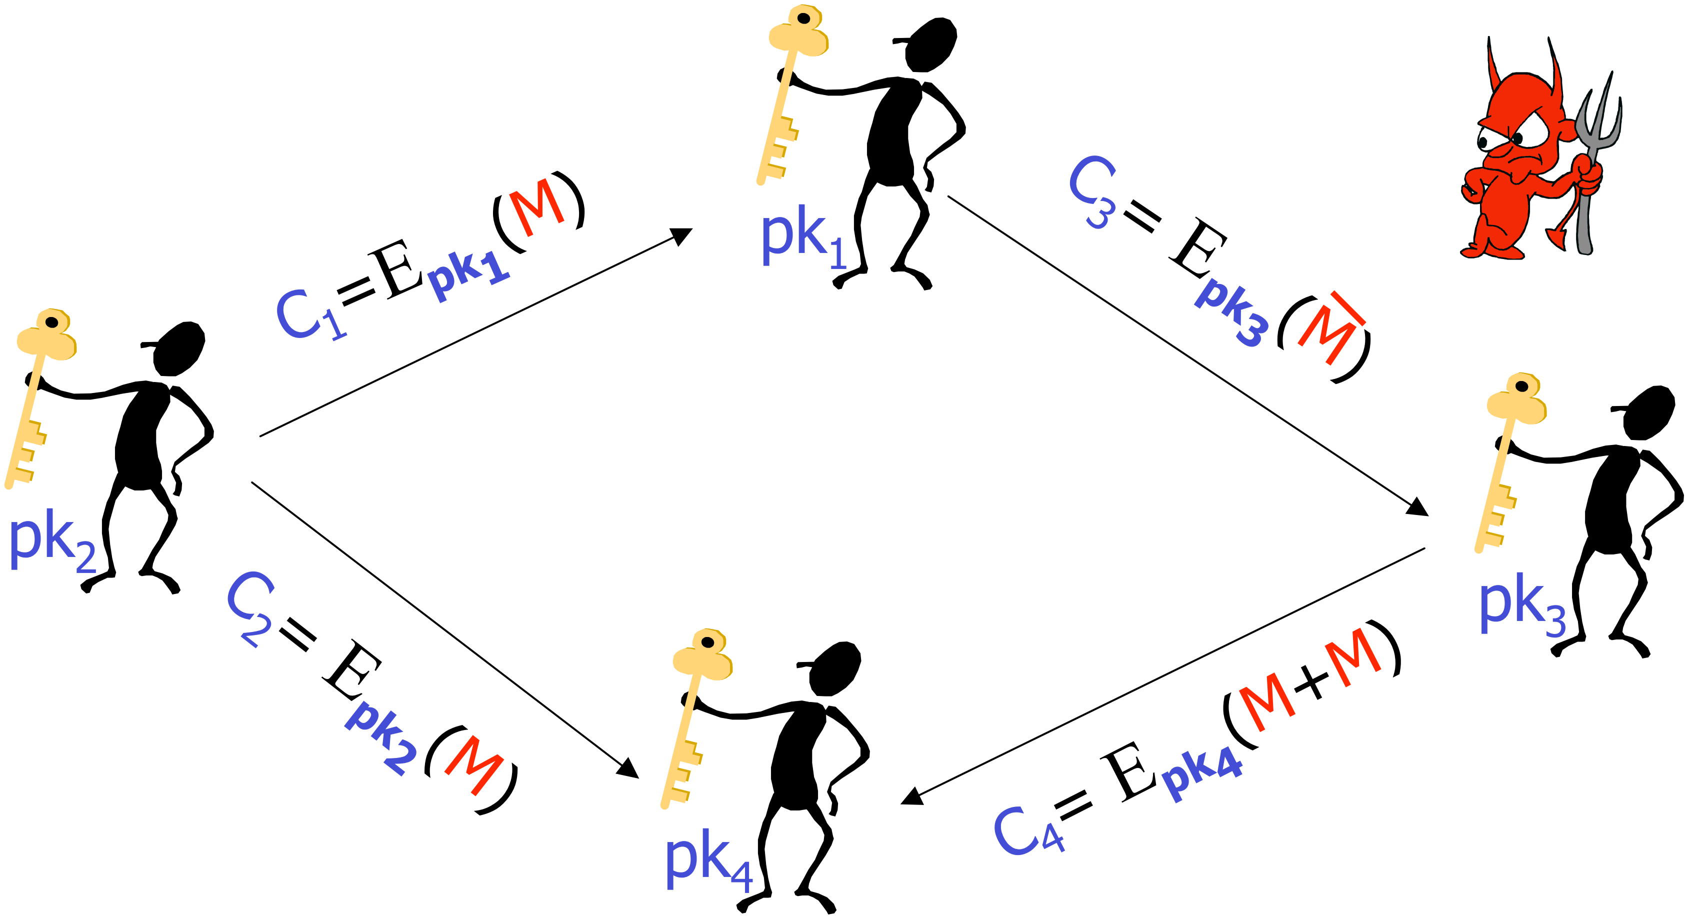

Danger: the adversary can see encryptions of related messages under different public keys.

Security definition (many users, CPA) [BBM]

$\Pi=(G,K,E,D)$

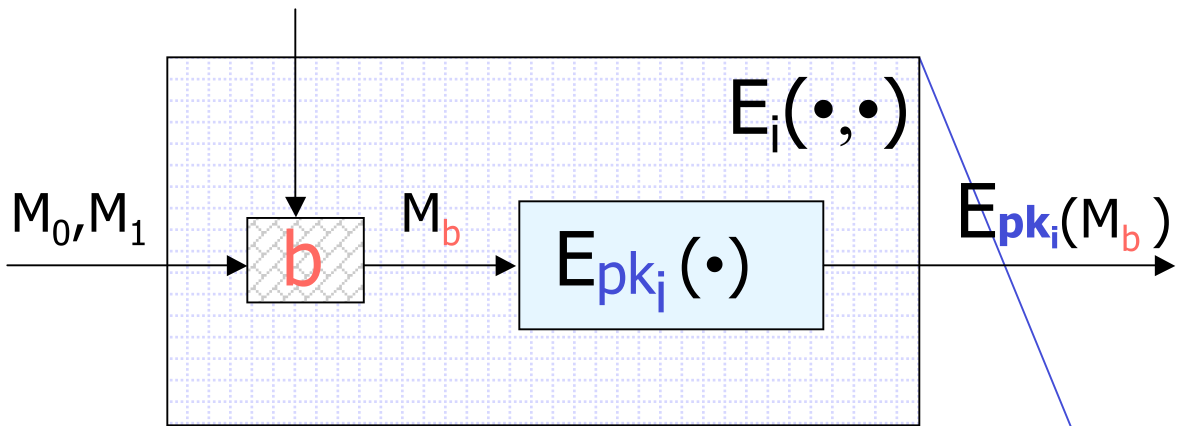

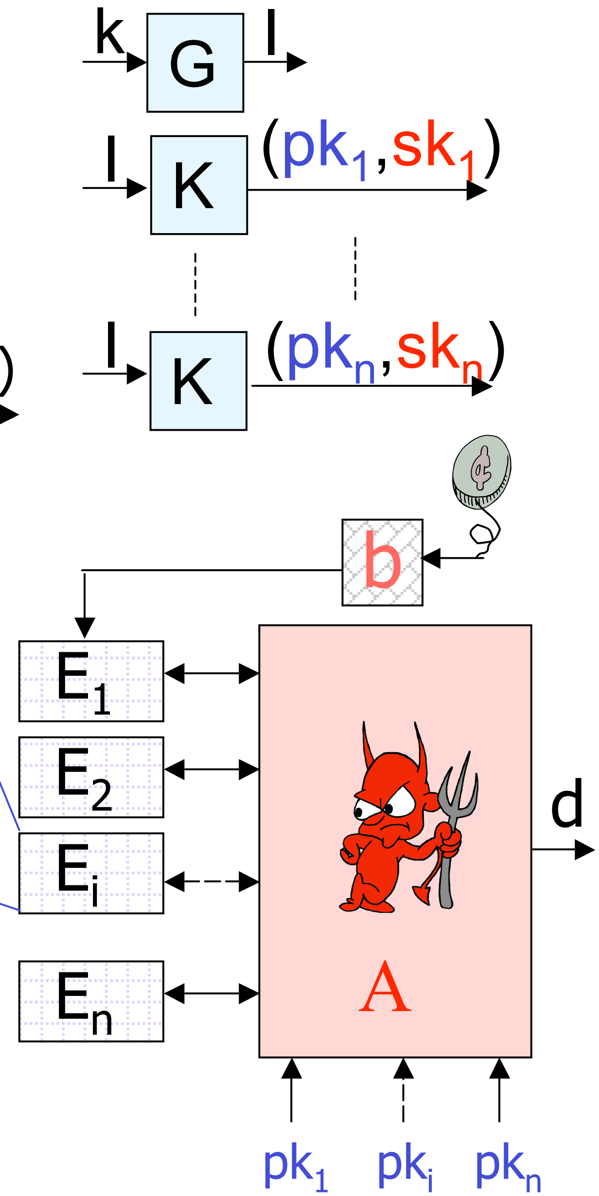

Adversary wins if b = d

$\mathrm{Adv}^{\text{n-ind-cpa}}_{A,\Pi}(k) = 2\Pr[\text{win}] - 1$

An encryption scheme Π is IND-CPA in the multi-user setting if for any polynomial n and PPT adversary A, $\mathrm{Adv}^{\text{n-ind-cpa}}_{A,\Pi}(k)$ is negligible in k.

Reminder

$\text{Adv}^{\text{attack type}}(\text{resourses of attacker})$

= max probability of any attacker breaking the scheme

Small = good

Big = bad

General reduction [BBM]

Theorem. Let Π = (K, E, D) be a public-key encryption scheme. Then

$$\mathrm{Adv}^{\mathrm{n}}(t, q_e) \leq q_e n \cdot \mathrm{Adv}^{1}(t')$$

where $t' \approx t$

Corollary. Encryption schemes polynomially-secure in the single-user setting are polynomially-secure in the multi-user setting.

General reduction

- implies schemes like El Gamal, RSA-OAEP are polynomially-secure in the multi-user setting.
- shows benefits of targeting **strong, well-defined** security definitions in the single-user setting: security in extended settings follows automatically.

The need for concrete security improvements

Consider a public-key encryption scheme Π with $\mathrm{Adv}^{1}(t') \leq 2^{-60}$

Assume in a real setting the number of users n = 200 000 000.

Allow $q_e = 2^{30}$ messages be encrypted under each public key.

Then $\mathrm{Adv}^{n}(t, q) \approx 0.2$

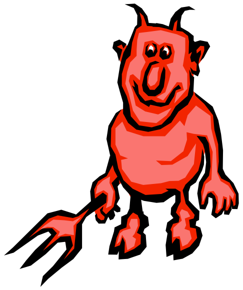

Tightness of the general reduction

Question. Is there a better reduction? No!

Proposition. [BBM] There exists a public-key encryption scheme Π with

$$\mathrm{Adv}^{\mathrm{n}}(t, q_e) = \Theta(q_e n) \cdot \mathrm{Adv}^{1}(t')$$

So, loss in security cannot be prevented in general. But we can hope to do better for specific schemes.

ElGamal encryption scheme

$G(k)$ $k \in \mathbb{N}$	$K()$	$E_{pk}(M)$ $M \in G$	$D_{sk}(C)$
$p \leftarrow k$-bit prime $g \leftarrow$ generator of a group G of order p Return (g, p)	$x \stackrel{\$}{\leftarrow} Z_p$ $X \leftarrow g^x$ $pk \leftarrow (g, p, X)$ $sk \leftarrow (g, p, x)$ Return (pk, sk)	$r \stackrel{\$}{\leftarrow} Z_p$ Return $(g^r, X^r \cdot M)$	$K \leftarrow Y^x$ $M \leftarrow T \cdot K^{-1}$ Return M

ElGamal in the multi-user setting

Our general reduction implies

$$\mathrm{Adv}^{\mathrm{n}}(t, q_e) \leq 2 q_e n \cdot \mathrm{Adv}^{1}(t')$$

Theorem [BBM]: improved reduction

$$\mathrm{Adv}^{\mathrm{n}}(t, q_e) \leq 2\mathrm{Adv}^{1}(t') + \frac{1}{p}$$

ElGamal scheme in the multi-user setting as secure as it is in the single user setting

Towards better efficiency of encryption in the multi-user setting

Consider a scenario where a sender needs to encrypt messages for several recipients:

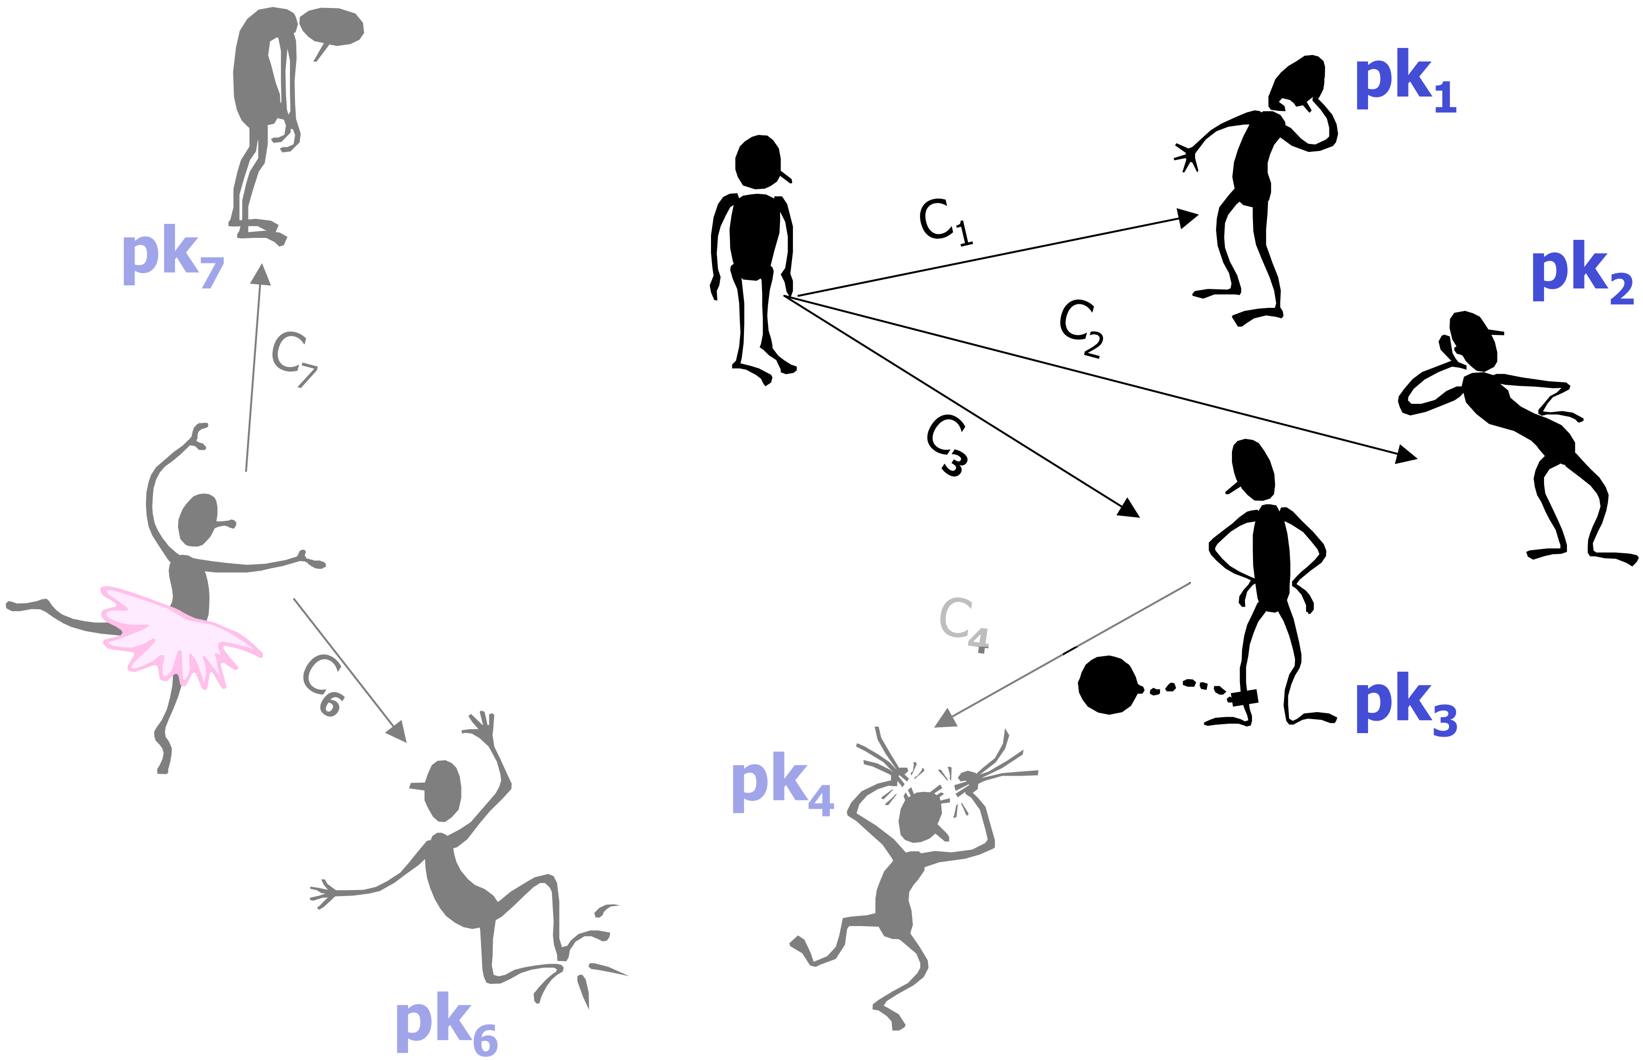

Simple solution

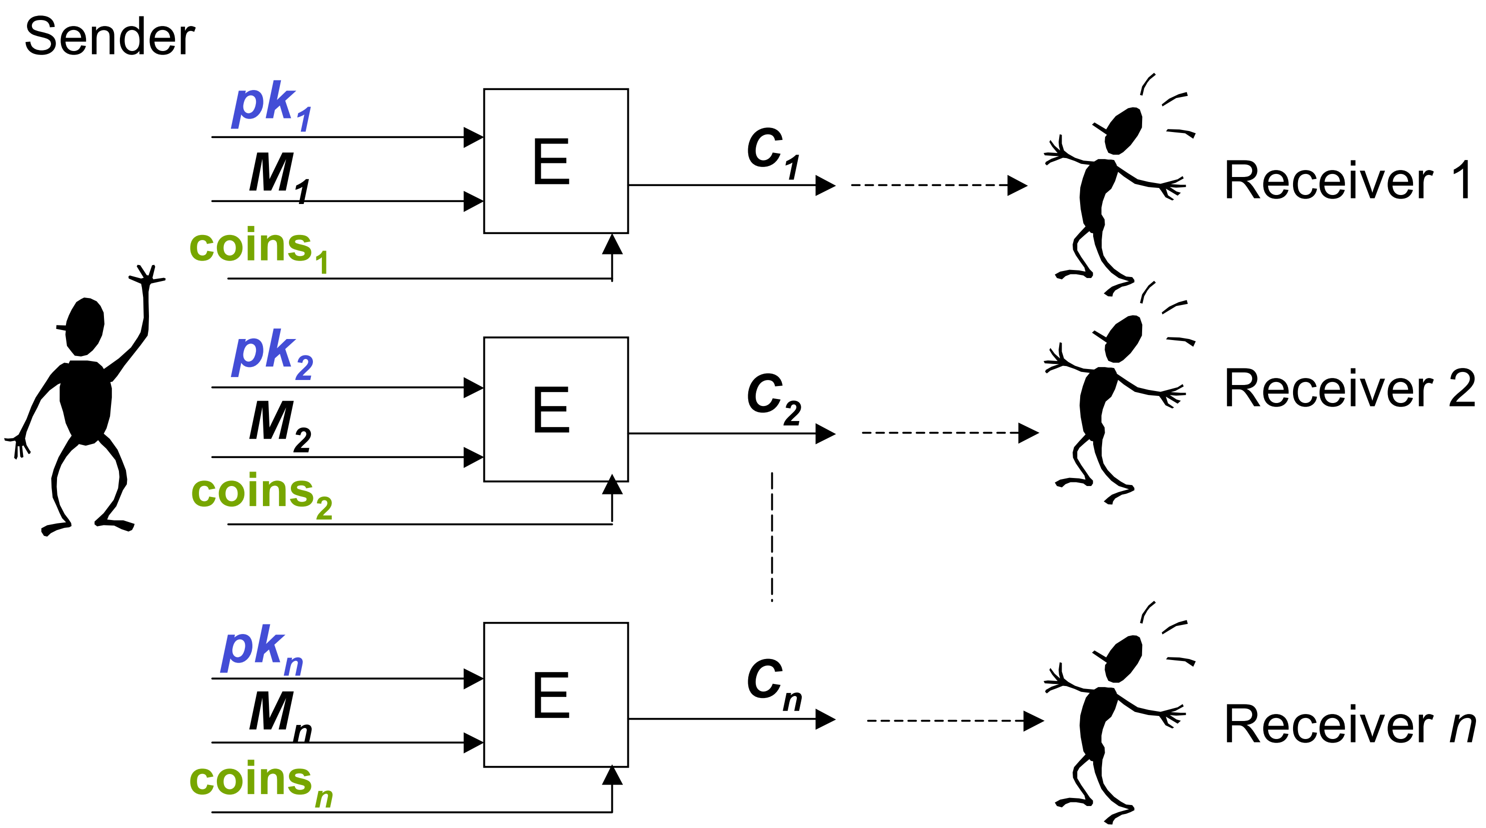

- Computational cost is n times that of the standard scheme
- Total length of all ciphertexts $C_1,..,C_n$ has size n times the size of a ciphertext in the standard scheme
- Can we do it more efficiently?

An application. Pay-TV.

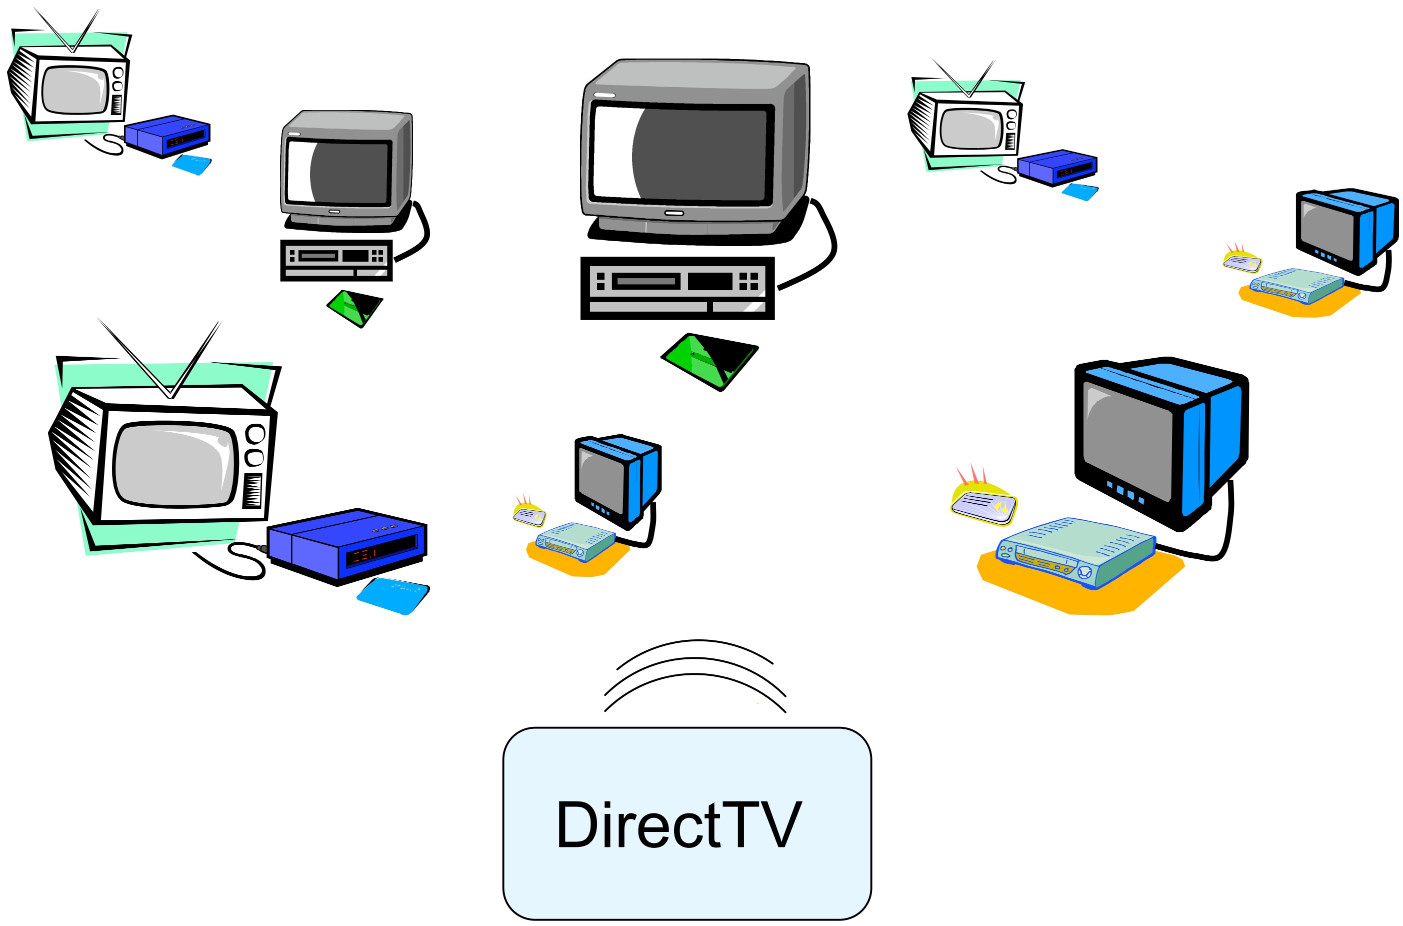

Encrypted messages are being broadcast such that only legitimate recipients can decrypt them.

It is desirable to shorten the broadcast communication

MRES

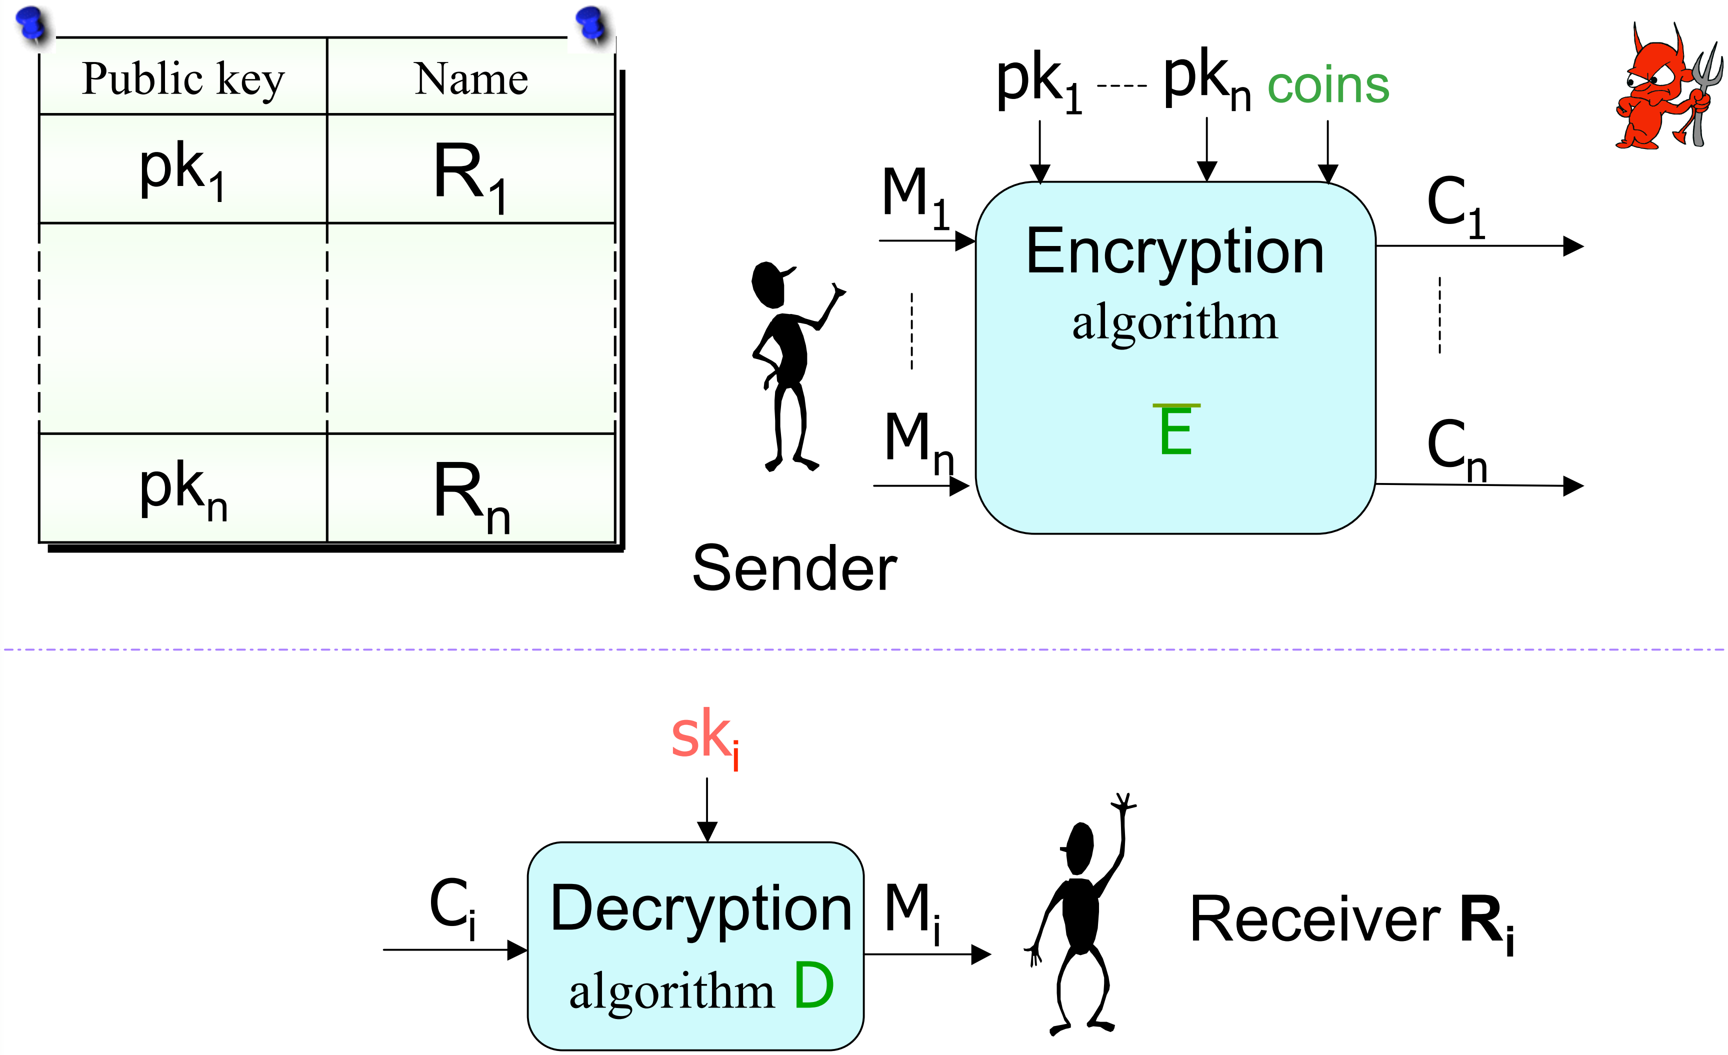

Public key
Name
pk1
R1
pkn
Rn
pk1 ---- pkn coins
M1
Mn
Encryption algorithm E
C1
Cn
Sender
ski
Ci
Decryption algorithm D
Mi
Receiver Ri

Naïve MRES

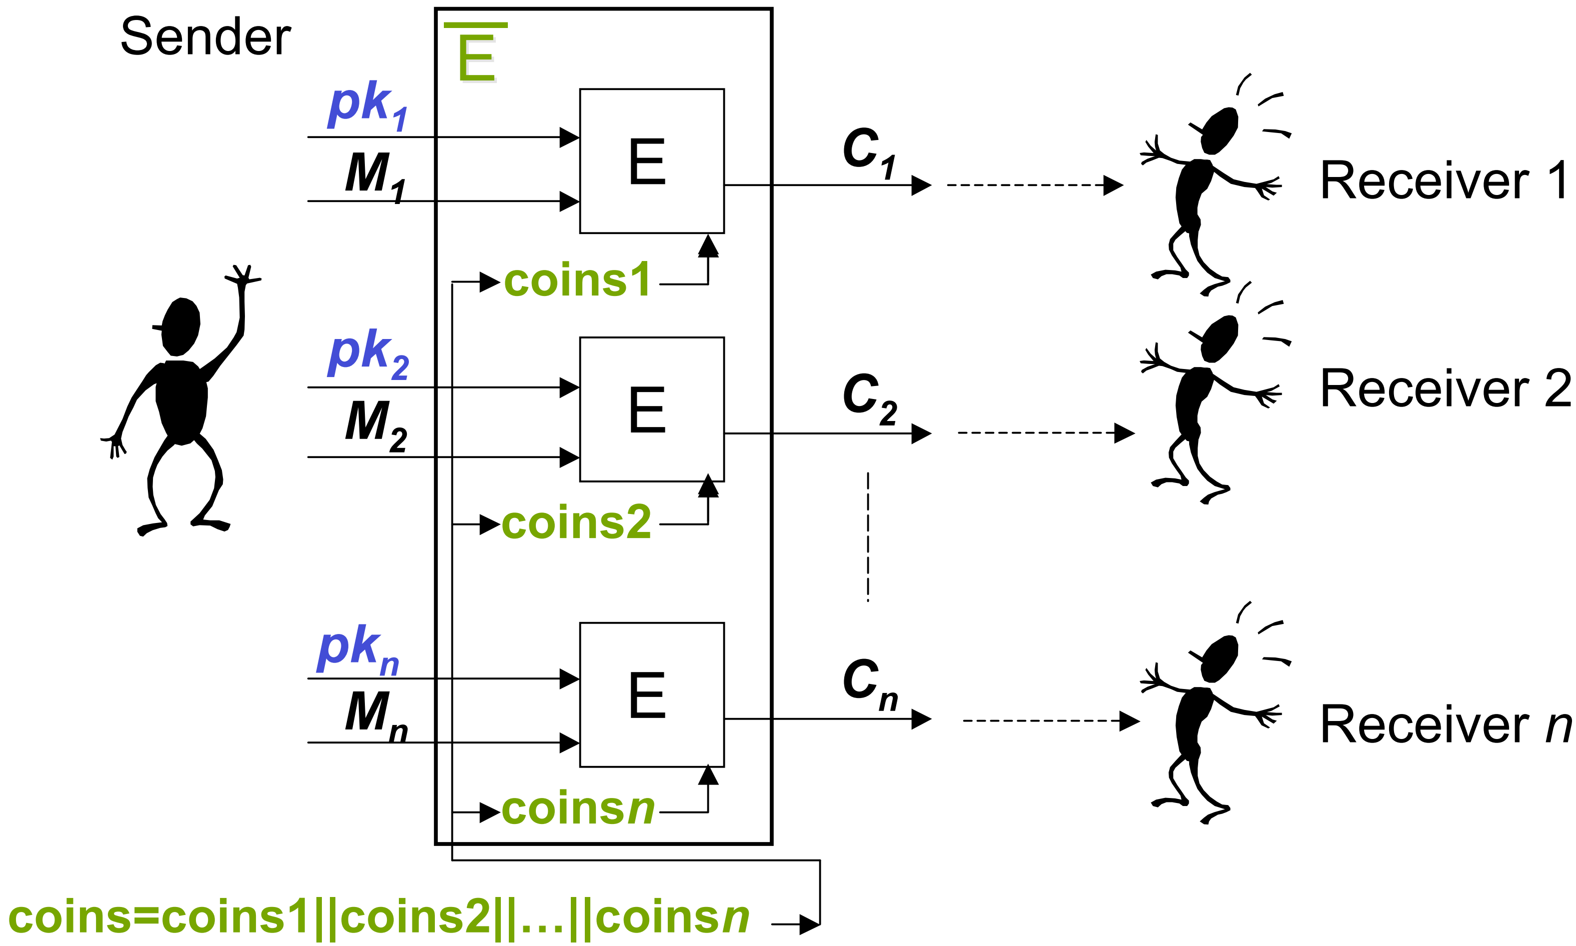

Sender
$\overline{E}$
pk_1
M_1
E
C_1
Receiver 1
coins1
pk_2
M_2
E
C_2
Receiver 2
coins2
pk_n
M_n
E
C_n
Receiver n
coins$n$
coins=coins1||coins2||…||coins$n$

Our suggestion

We suggest a possibility to "reuse" random coins used in the naïve MRES encryption:

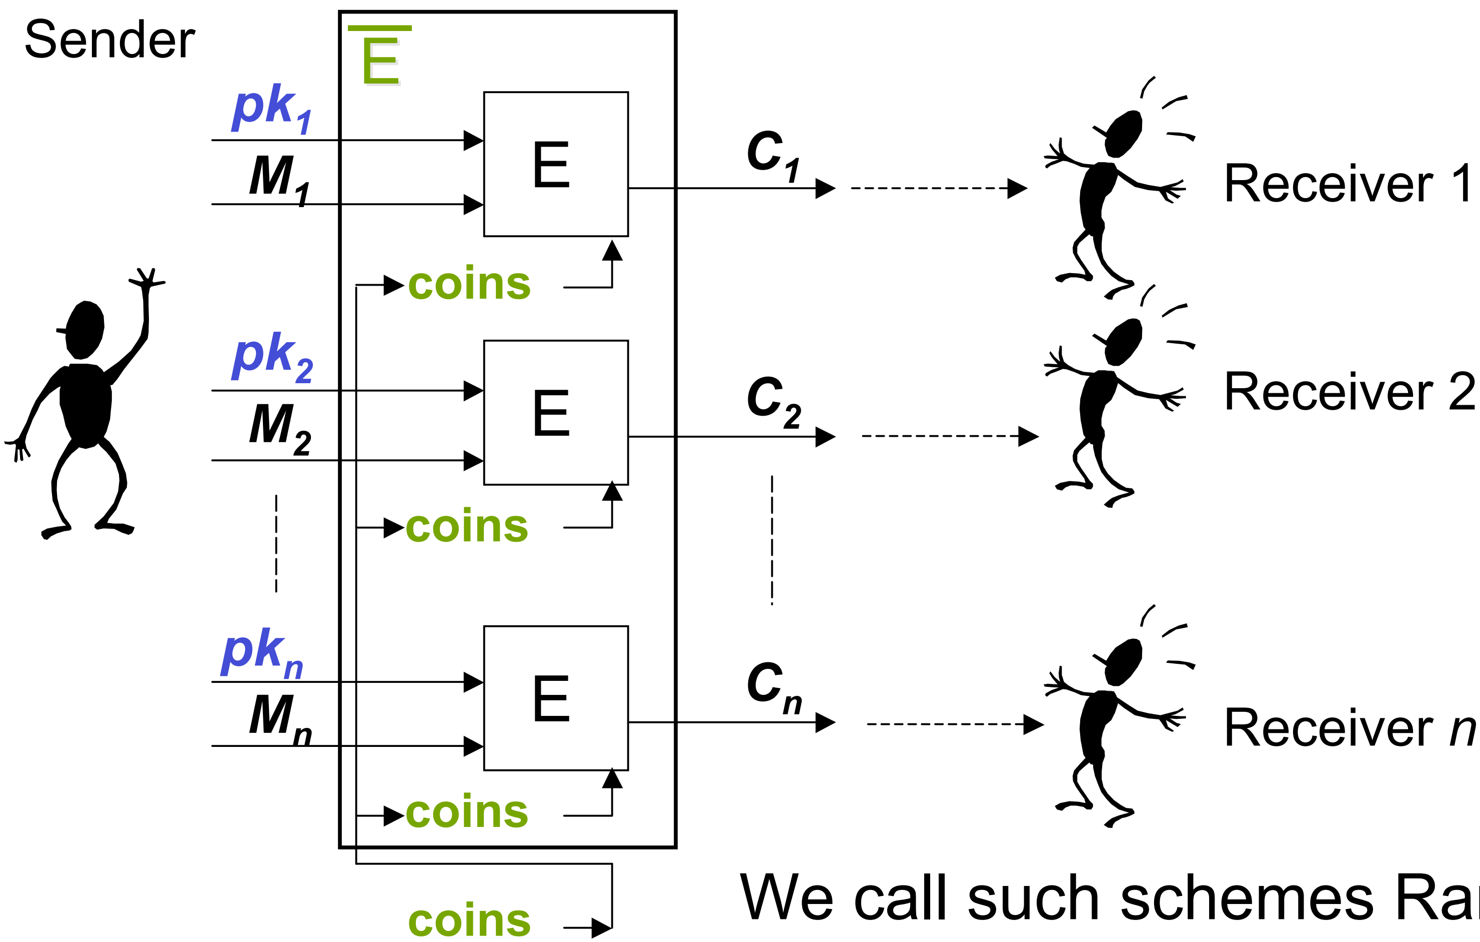

We call such schemes Randomness Reusing MRES (RR-MRES)

Why are RR-MRESs interesting?

Consider ElGamal-based RR-MRES:

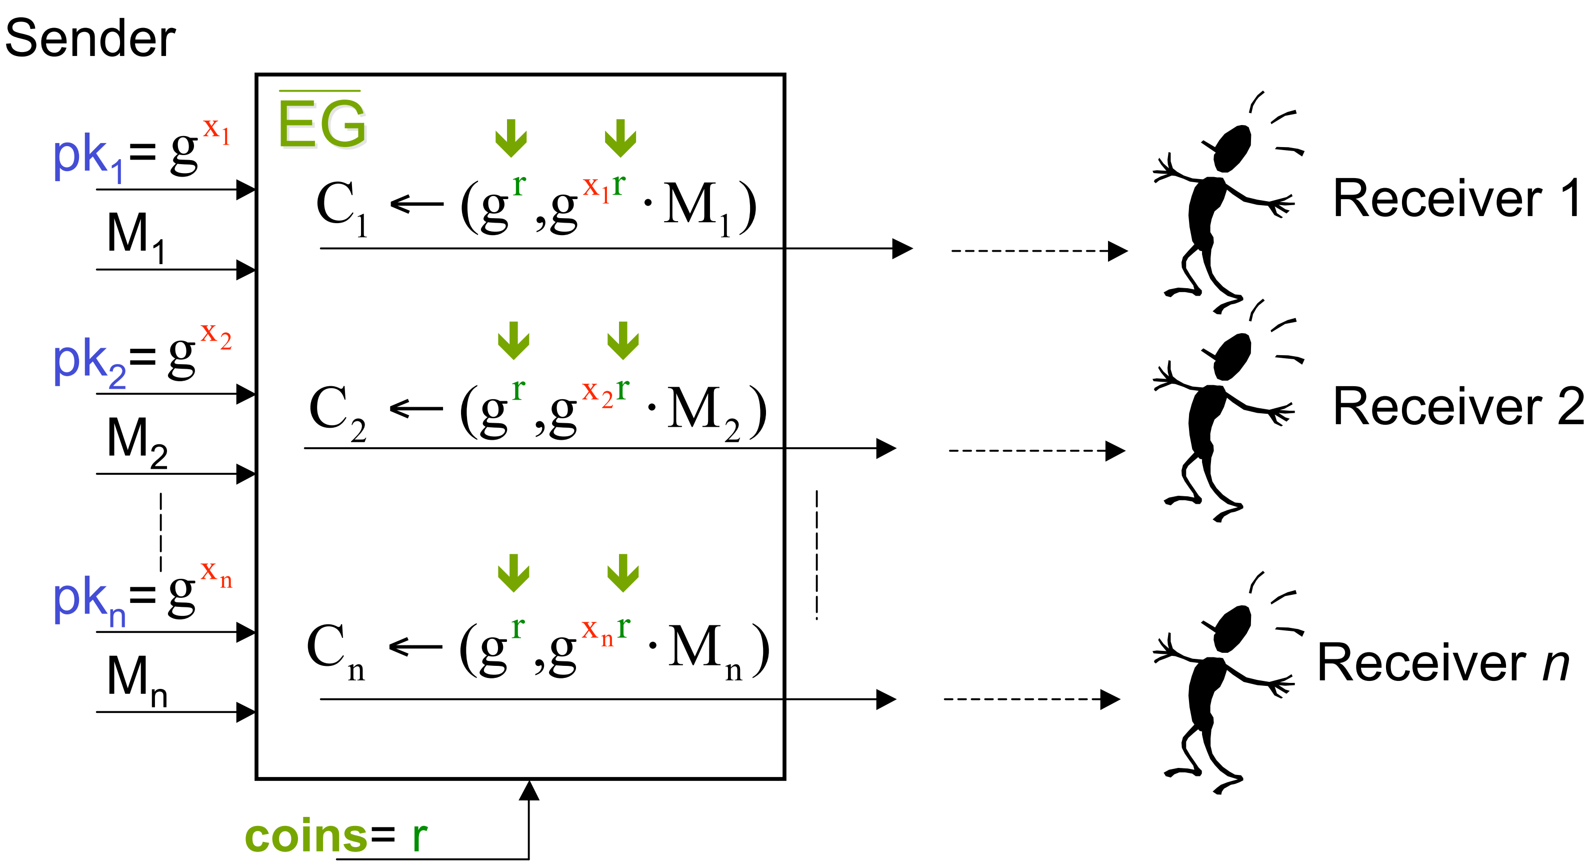

Why are RR-MRESs interesting?

- Half the number of exponentiations used by the naive ElGamal-based MRES.
- If ciphertexts are broadcast, need only send $(g^r, g^{x_1 r} \cdot M_1, .., g^{x_n r} \cdot M_n)$ which is half the length of the broadcast vector for the naive MRES.
- Saving 50% in computation is important: exponentiations are still relatively slow operations, people struggle to get any improvements.
- Our results serve as a proof of concept, people could try to achieve even better savings.

But are these schemes secure?

- To analyze security of MRESs one needs an appropriate security definition.
  - The security definitions for the multi-user setting are incompatible with syntax of MRESs.
  - New types of attacks arise in the multi-recipient setting.

What should an adversary be allowed to do?

- An adversary can see encryptions of related messages under different public keys
- An adversary can be one of the recipients:
  - Learn the corresponding secret key, decrypt the ciphertexts
  - Register its own public key, which possibly depends on public keys of honest users

Rogue-key attacks: An example

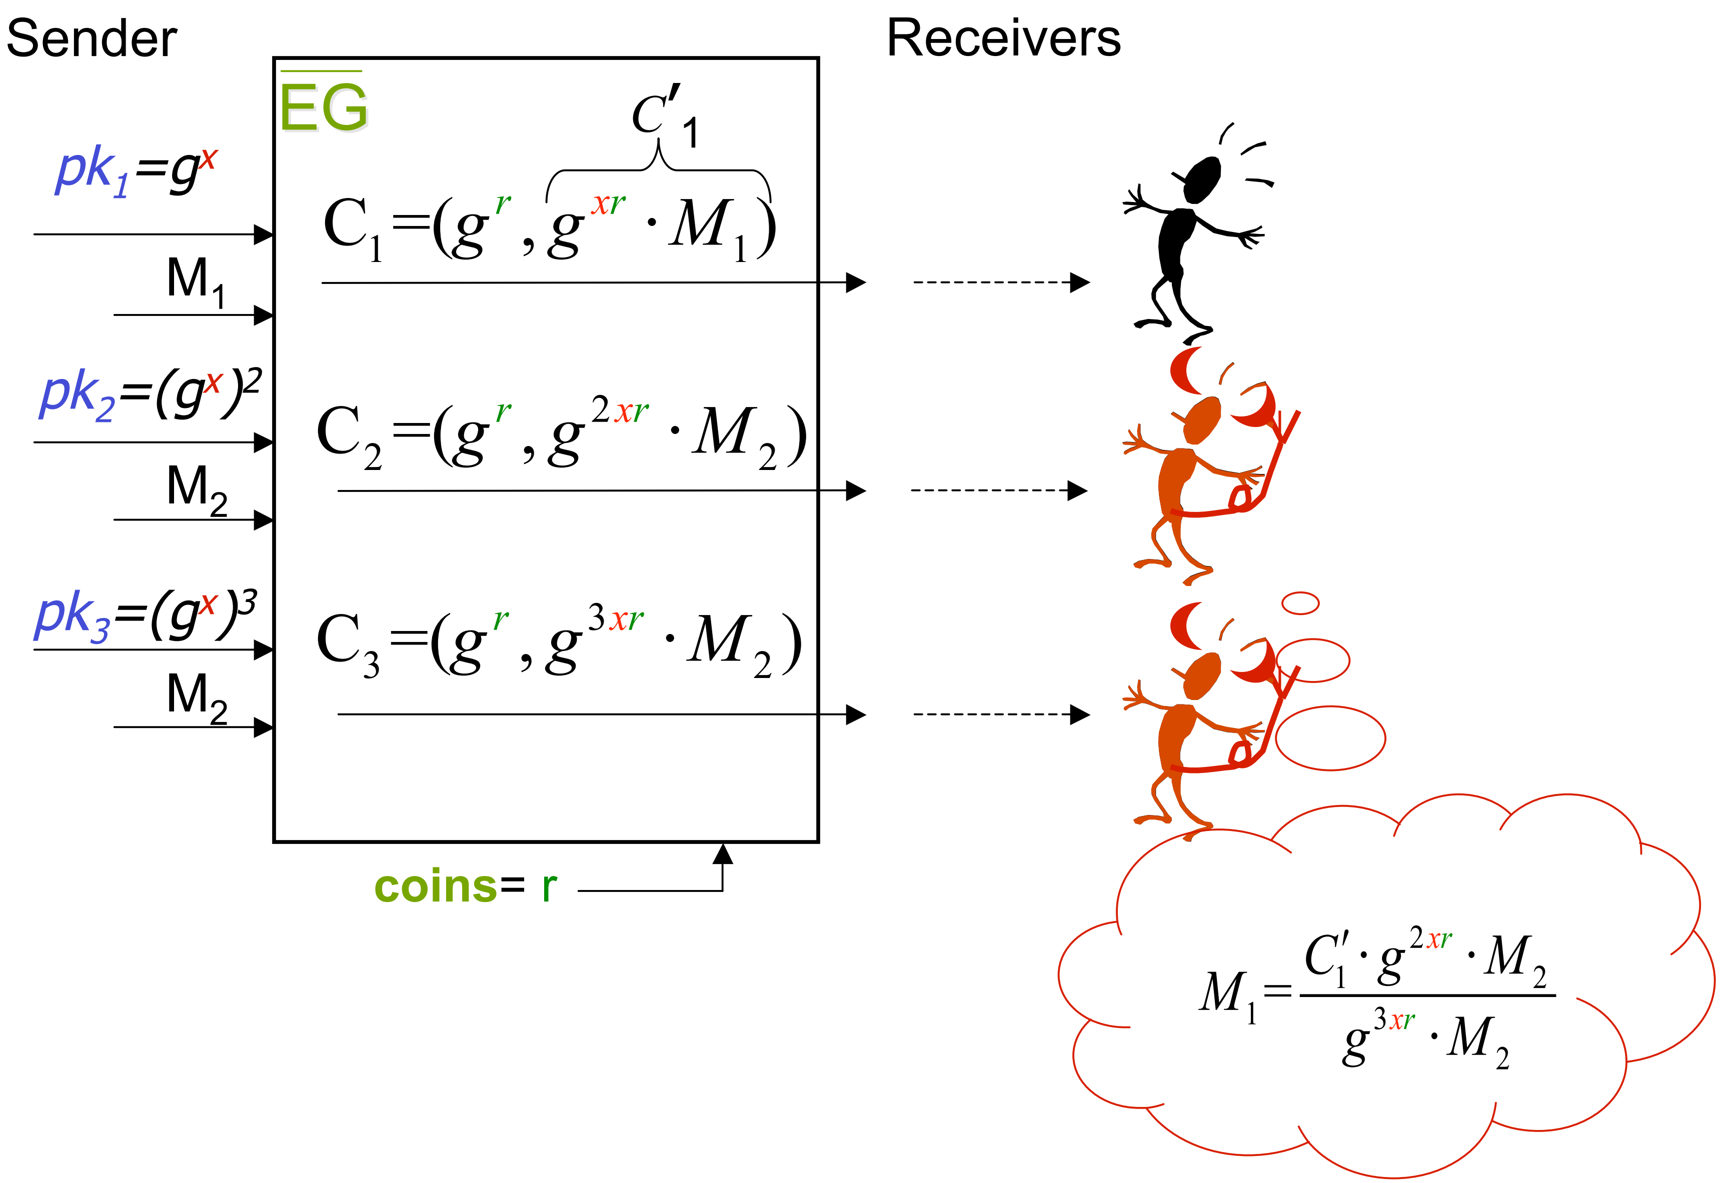

Towards a security definition for MRESs

- Our model allows an adversary to corrupt some recipients and obtain their secret keys.
- The model also allows an adversary to choose public keys of corrupted recipients as a function of the public keys of honest recipients.
- But: Only if it also outputs valid corresponding secret keys. This abstraction avoids consideration of explicit proofs of knowledge of secret keys, which are done (should be done) when users register their public keys with the CAs.
- Security requires it still be unable to obtain even partial information about messages sent to uncorrupted recipients.

MRESs security definition (against CPA) [BBS]

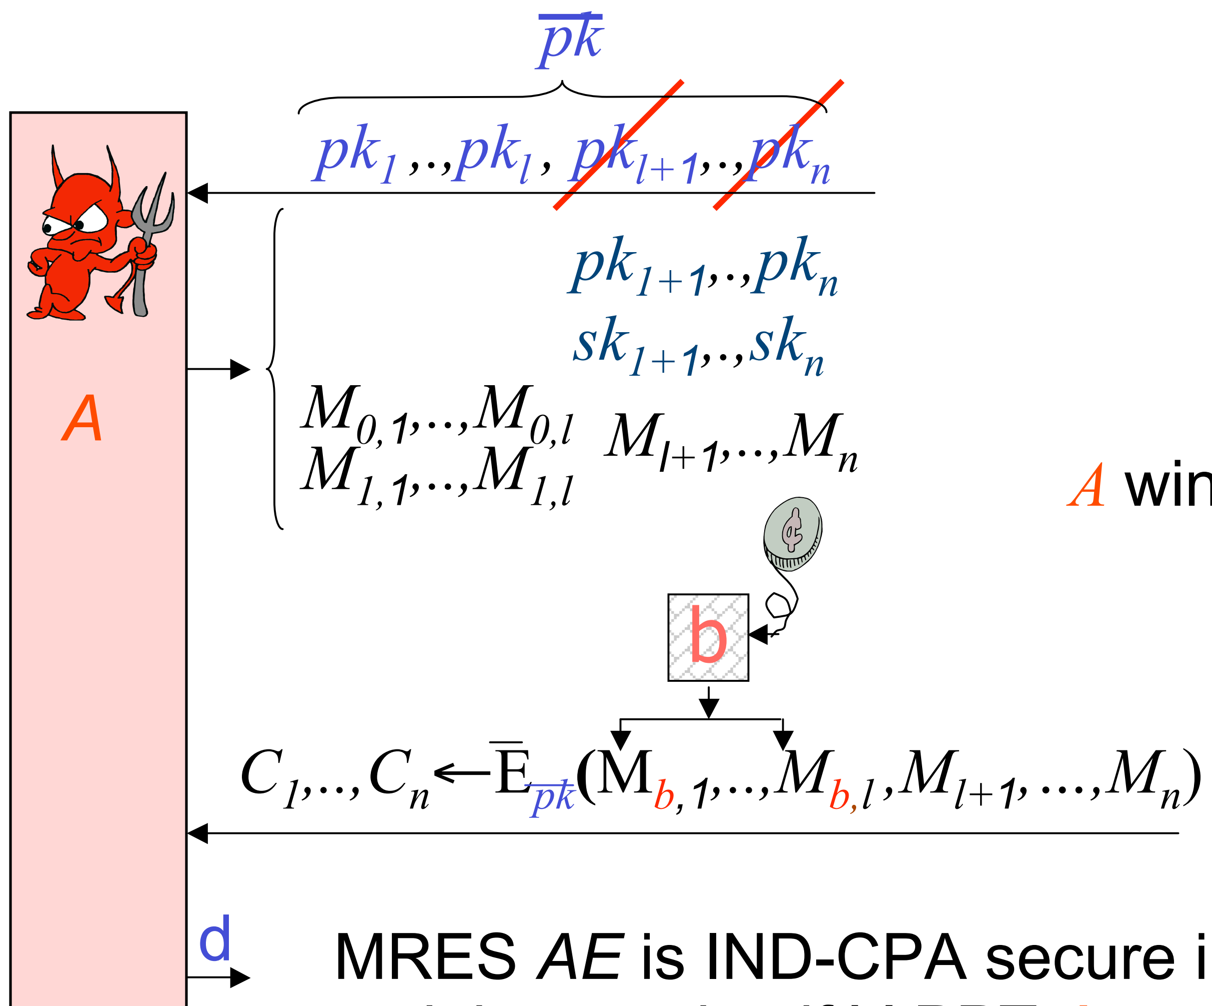

A wins if b=d

MRES *AE* is IND-CPA secure in the multi-recipient setting if ∀ PPT *A*

$$Adv^{n-mr-cpa}_{A,AE}(k) = 2\Pr[A \text{ wins}] - 1$$

is negligible in k

The possibility of insider and rogue-key attacks exists for both standard and multi-recipient encryption schemes.

The definition of security for MRESs takes them into account while the one for encryption in the multi-user setting does not. Why?

It is not necessary:

Claim. For ATK∈{CPA,CCA}, an encryption scheme $\mathcal{AS}$ is IND-ATK in the multi-user setting iff $\mathcal{AS}$-based naïve MRES is IND-ATK in the multi-recipient setting.

Not all RR-MRESs are secure. Example. RSA-OAEP

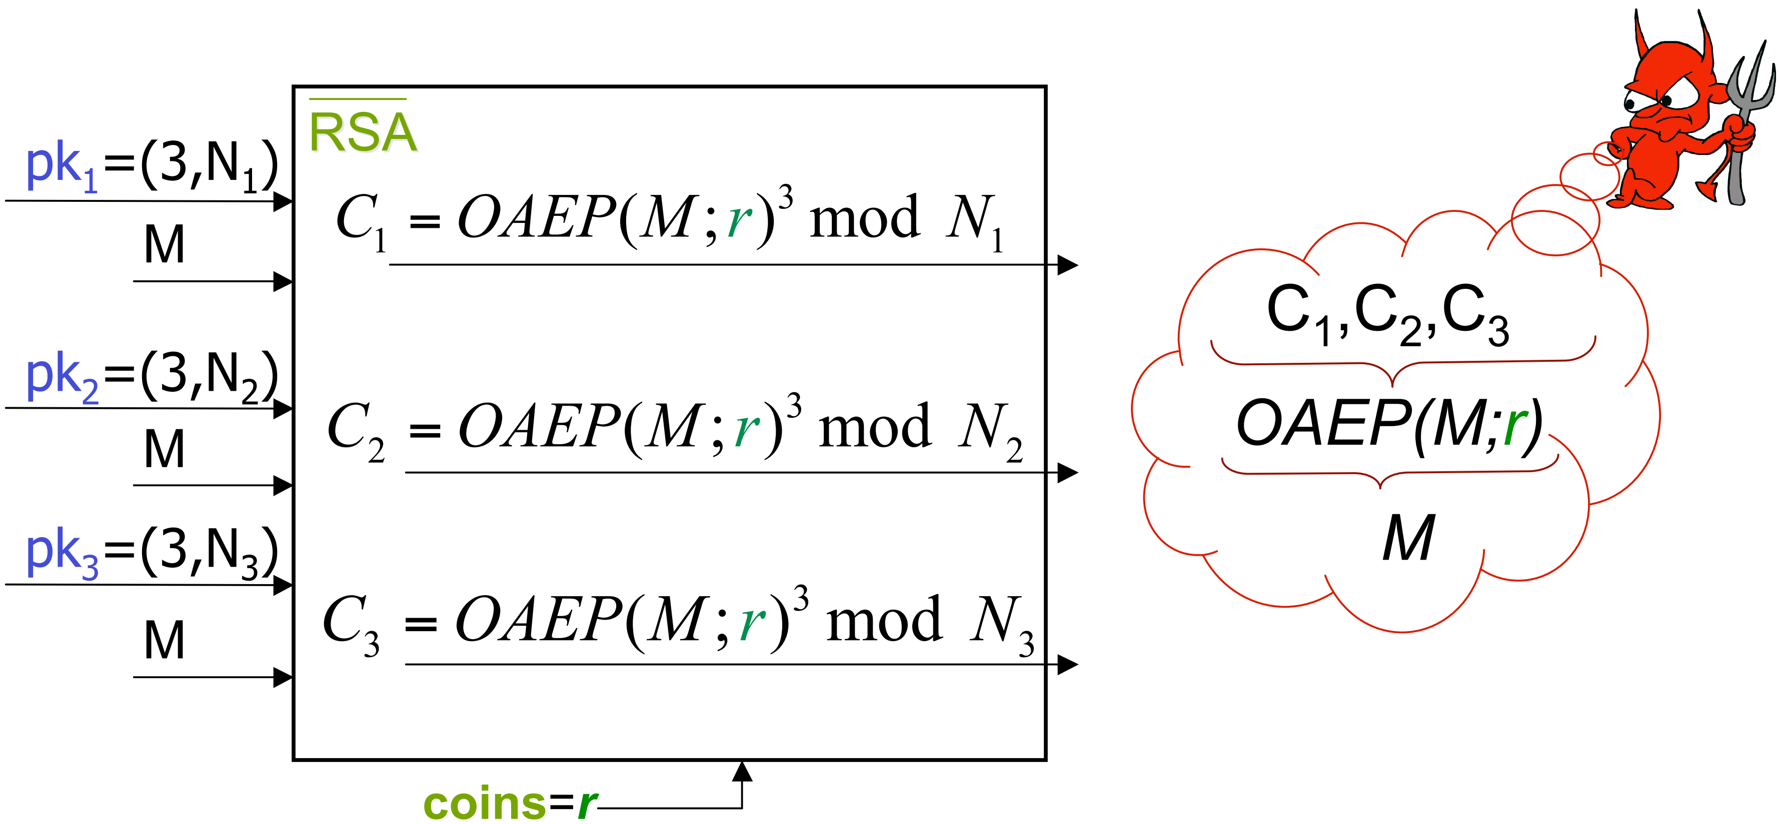

To facilitate finding secure RR-MRESs

we suggest a special test:

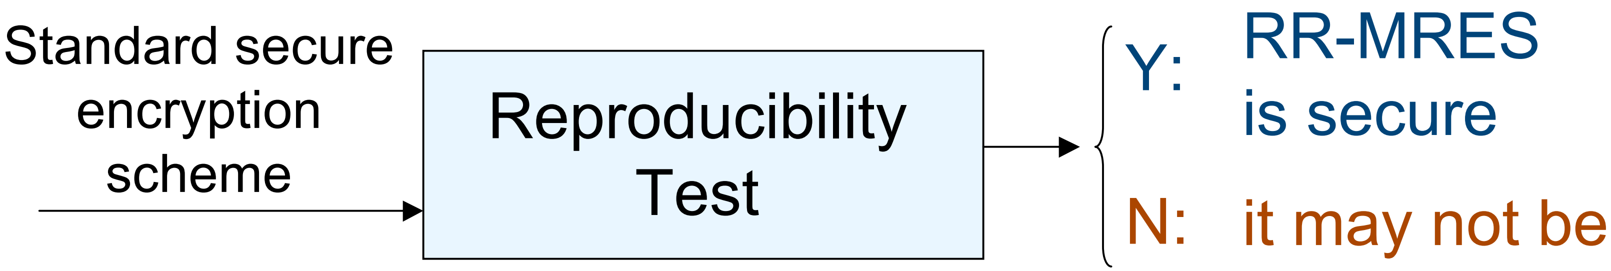

Avoids case-by-case security analysis of MRESs.

Reproducibility test and theorem

A standard encryption scheme AE is reproducible if ∃ a PPT algorithm *R*:

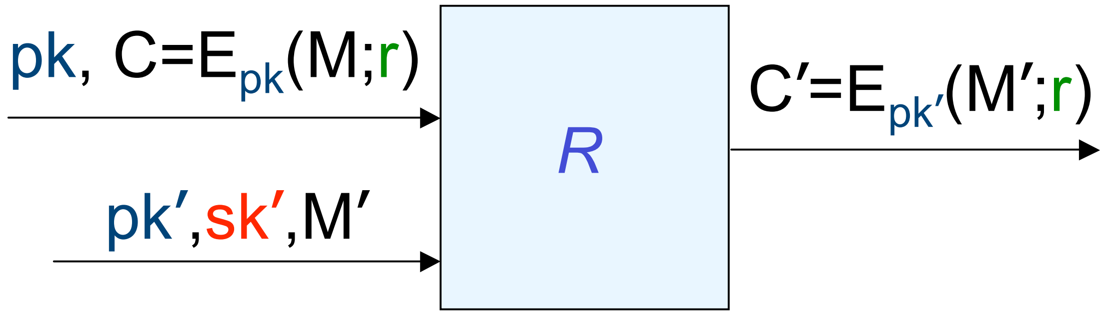

Theorem. For ATK∈{CPA,CCA}, if a standard encryption scheme AE is

1. reproducible
2. IND-ATK secure

then the corresponding RR-MRES is IND-ATK secure.

Security of ElGamal-based RR-MRES

Lemma. El Gamal encryption scheme *EG* is reproducible.

Proof.

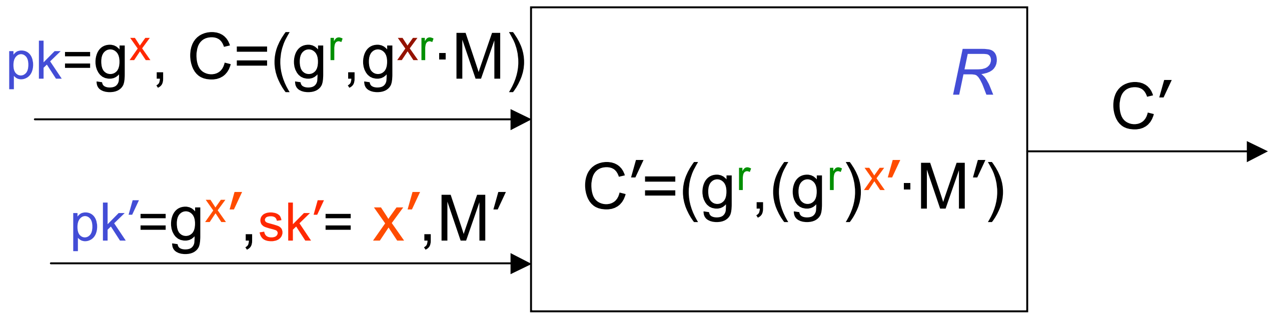

Fact. *DDH* is hard => *EG* is IND-CPA secure

Corollary.

Lemma
\+
Fact
\+
Reproducibility theorem

=> El Gamal-based RR-MRES is IND-CPA secure

Anonymity (key-privacy) in the multi-user setting

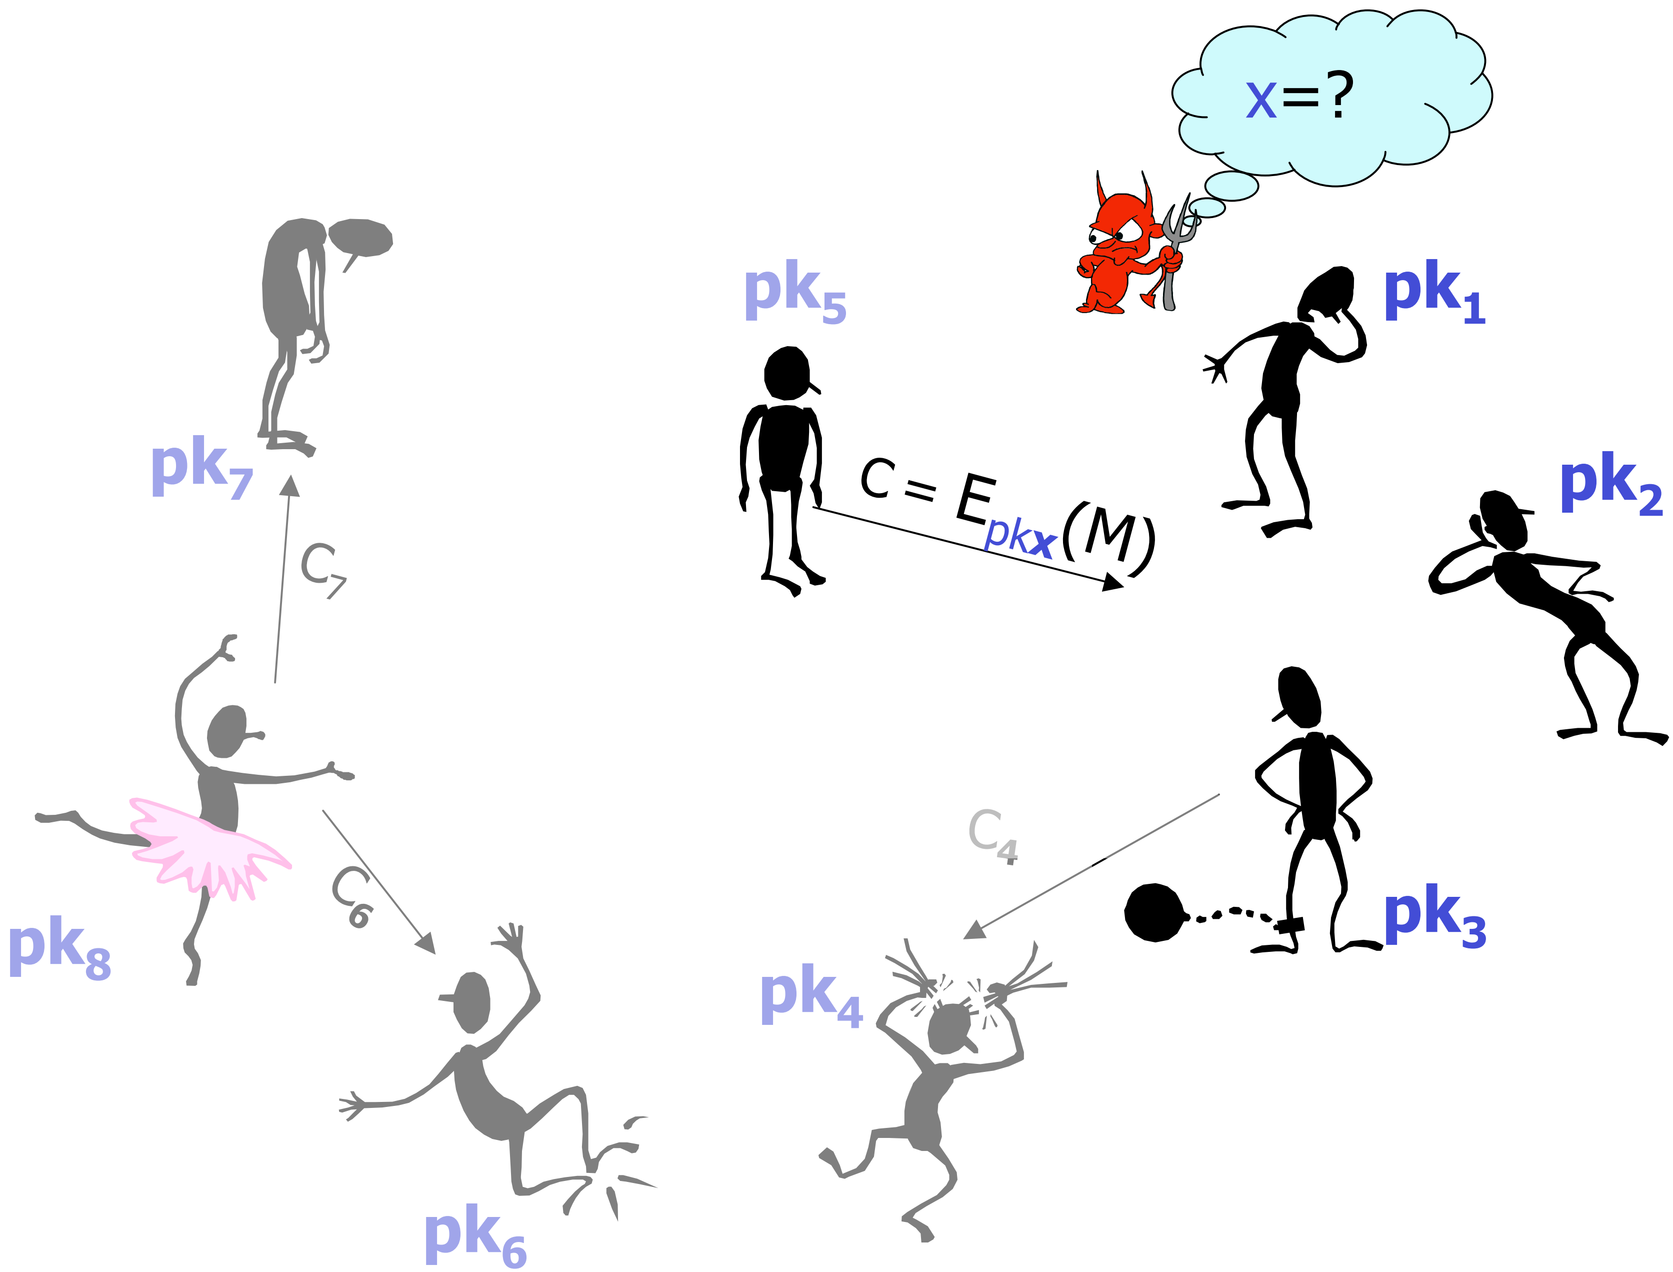
x=?
pk1
pk2
pk3
pk4
pk5
pk6
pk7
pk8
C = E_pkx(M)
C7
C6
C4

Anonymity (key privacy)

- Data privacy was considered the sole goal of encryption
- Key privacy is another, previously overlooked goal

$C = E_{pk_2}(M)$

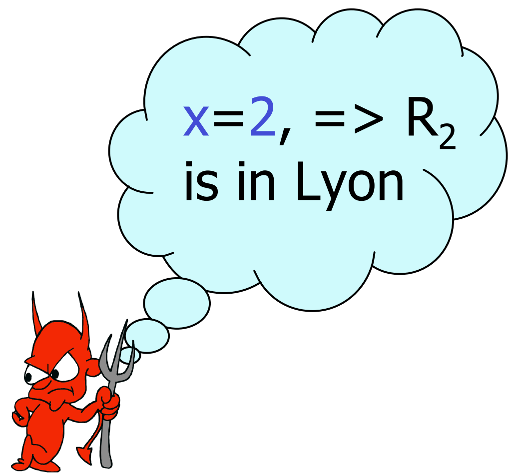

Public key	Name
pk_1	R_1
pk_n	R_n

This property of encryption is required be various protocols, such as anonymous credentials, mix-nets, private keyword search, etc.

RSA is not anonymous

Public key	Name
$pk1=(e_1,N_1)$	R_1
$pk2=(e_2,N_2)$	R_2

$E_{pk_i}(M)$: $M^{e_i} \bmod N_i$

If $C > N_1$, then it's a ciphertext addressed to R_2

The same attack applies to all popular variations of RSA scheme, including RSA-OAEP.

Anonymity. [BBDP] Summary of contributions

- Defined an appropriate security definition
- Proved that ElGamal and Cramer-Shoup provide anonymity under the same assumptions they provide data-privacy
- Show how to modify RSA to provide anonymity

References.

- [BBM] With M. Bellare and S. Micali, "Public-Key Encryption in a Multi-User Setting: Security Proofs and Improvements." In Eurocrypt 2000 Proceedings, (LNCS) Vol. 1807, pp. 259-274, 2000.
- [BBDP] With M. Bellare, A. Desai and D. Pointcheval, "Key-Privacy in Public-Key Encryption." In Asiacrypt 2001 Proceedings, LNCS Vol. 2248, pp. 566-582, 2001.
- [BBS] With M. Bellare and J. Staddon, "Randomness Re-use in Multi-Recipient Encryption Schemes." In Public Key Cryptography (PKC) 2003 Proceedings, LNCS Vol. 2567, pp. 85-99, 2003.
- [BBKS] With M. Bellare, K. Kurosawa and J. Staddon, "Multi-Recipient Encryption Schemes: Security and Optimization." Work in progress.

Available at
http://www.cc.gatech.edu/~aboldyre/publications.html

Thank you!