$HB^{\#}$: increasing the security and efficiency of HB^{+}

Henri Gilbert, Matt Robshaw, and Yannick Seurin

Eurocrypt 2008 – April 16, 2008

the context

- pervasive computing (RFID tags . . .)
- the issue: protection against duplication and counterfeiting $\Longrightarrow$ authentication
- pervasive = very low cost $\Longrightarrow$ very few gates for security
- current proposed solutions use *e.g.*
    - light-weight block ciphers (AES, PRESENT . . .)
    - dedicated asymmetric cryptography (crypto-GPS, SQUASH)
    - protocols based on abstract hash functions and PRFs
- recent proposal HB$^+$ at Crypto '05 by Juels and Weis: very simple, security proof

outline

- HB$^+$: strengths and weaknesses
- introducing RANDOM-HB$^\#$
- introducing HB$^\#$
- Ouafi et al. ’s MIM attack
- conclusions

the ancestor HB [Hopper and Blum 2001]

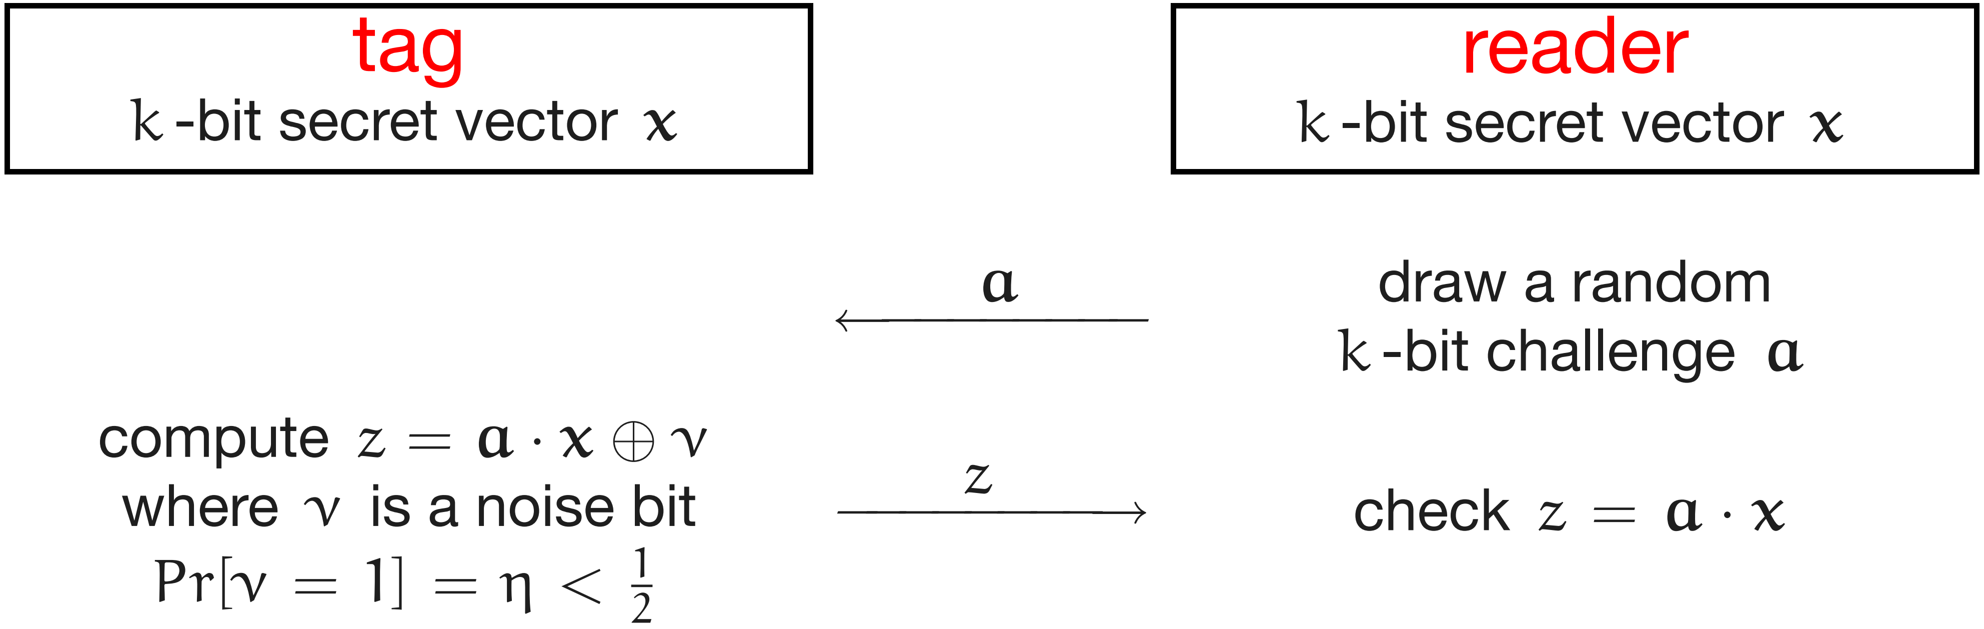

tag

k-bit secret vector $\boldsymbol{x}$

reader

k-bit secret vector $\boldsymbol{x}$

draw a random k-bit challenge $\boldsymbol{a}$

$\xleftarrow{\quad \boldsymbol{a} \quad}$

compute $z = \boldsymbol{a} \cdot \boldsymbol{x} \oplus \nu$
where ν is a noise bit
$\Pr[\nu = 1] = \eta < \frac{1}{2}$

$\xrightarrow{\quad z \quad}$

check $z = \boldsymbol{a} \cdot \boldsymbol{x}$

- this is repeated for r rounds
- the authentication is successful iff at most t rounds have been rejected ($t > \eta r$)

the protocol HB$^+$ [Juels and Weis 2005]

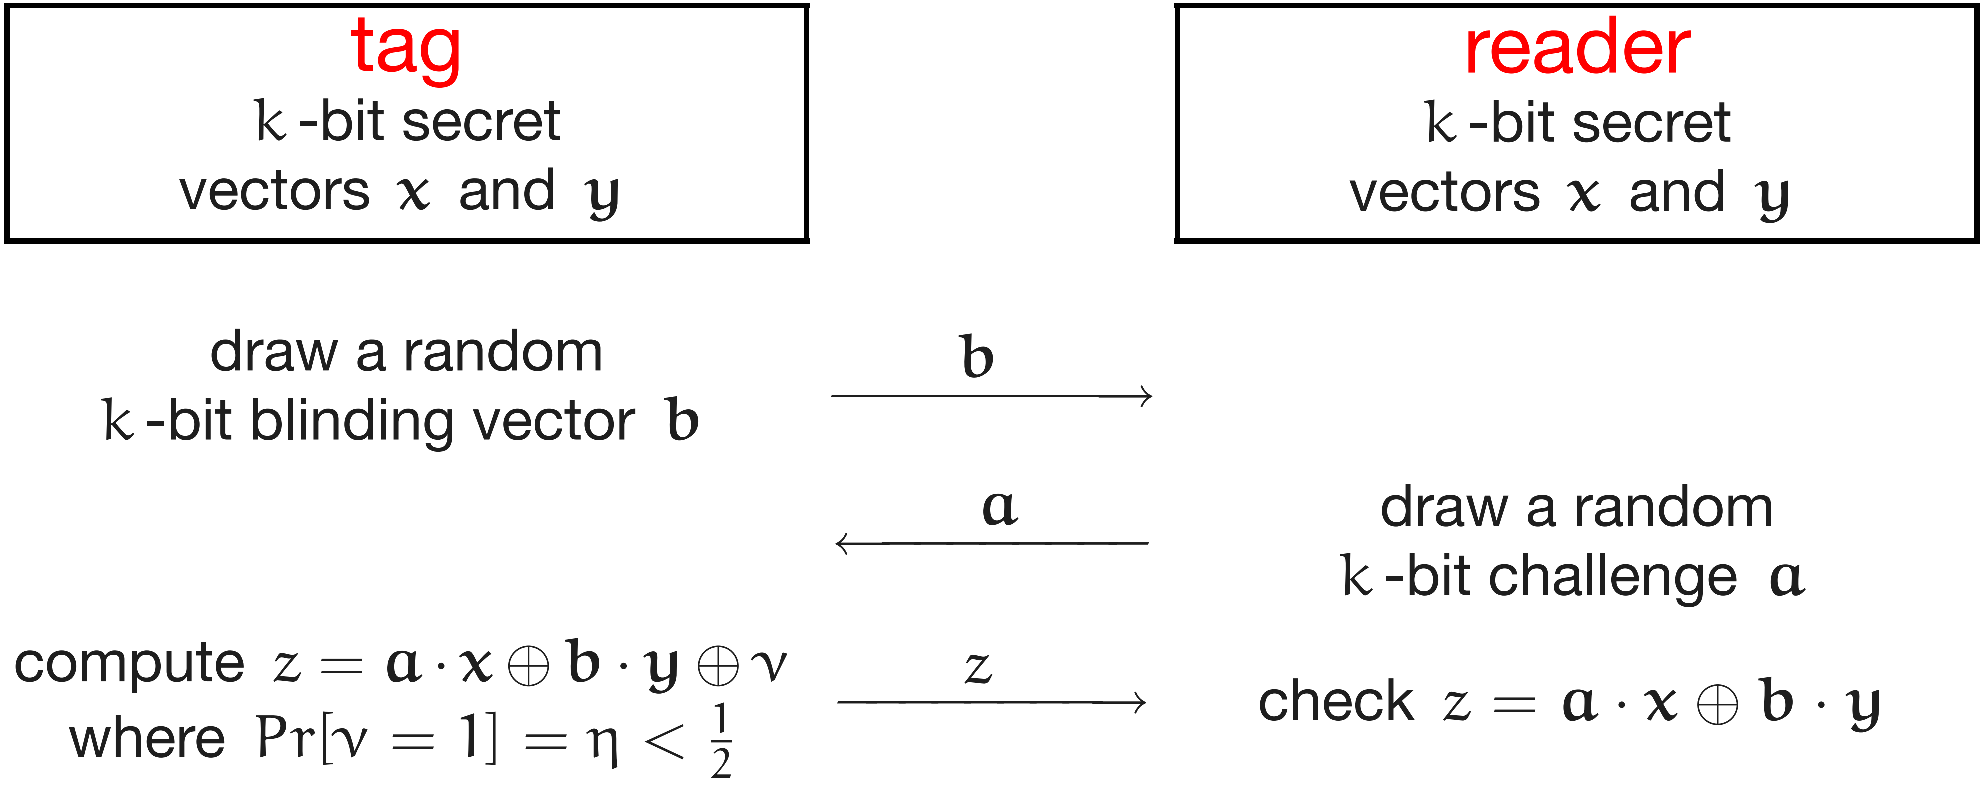

tag

k-bit secret vectors x and y

reader

k-bit secret vectors x and y

draw a random k-bit blinding vector $\mathbf{b}$ $\xrightarrow{\mathbf{b}}$

$\xleftarrow{\mathbf{a}}$ draw a random k-bit challenge $\mathbf{a}$

compute $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y} \oplus \nu$ where $\Pr[\nu = 1] = \eta < \frac{1}{2}$ $\xrightarrow{z}$ check $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$

- this is repeated for r rounds
- the authentication is successful iff at most t rounds have been rejected $(t > \eta r)$

the protocol HB^+

- typical parameter values are:
  - $k \simeq 250$ (length of the secret vectors)
  - $\eta \simeq 0.125$ to 0.25 (noise level)
  - $r \simeq 80$ (number of rounds)
  - $t \simeq 30$ (acceptance threshold)
- necessary trade-off between false acceptance rate, false rejection rate and efficiency
- rounds can be parallelized [Katz, Shin, 2006]

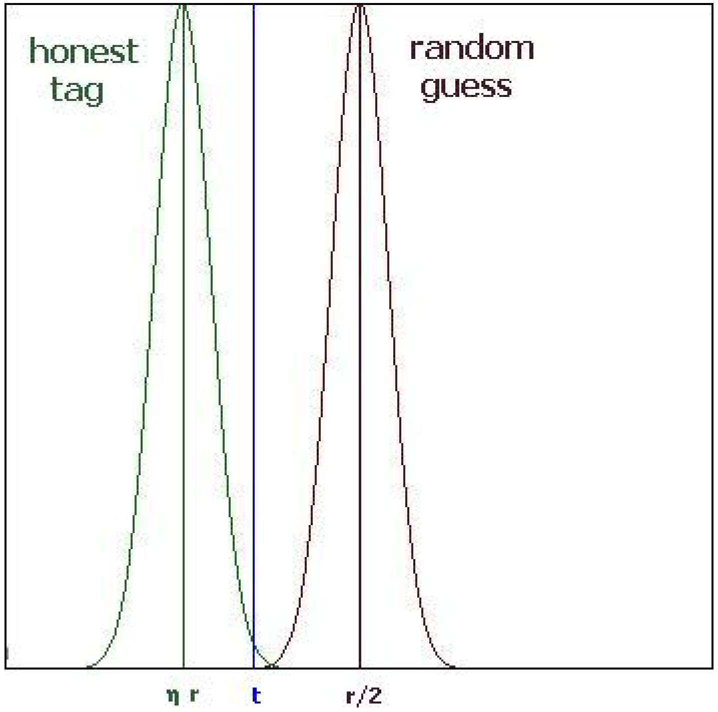

distribution of the number of errors

- practical limitation: transmission costs ($2kr + r$ bits, = tens of thousands)

the security of HB$^+$

- HB is provably secure against *passive* (eavesdropping) attacks
- HB$^+$ is provably secure against *active* (in some sense) attacks
- the security relies on the hardness of the *Learning from Parity with Noise* (LPN) problem:

  Given q noisy samples $(\mathbf{a}_i, \mathbf{a}_i \cdot \mathbf{x} \oplus \nu_i)$, where $\mathbf{x}$ is a secret k-bit vector and $\Pr[\nu_i = 1] = \eta$, find $\mathbf{x}$.

- similar to the problem of decoding a random linear code (NP-complete)
- best solving algorithms require $T, q = 2^{\Theta(k/\log(k))}$: BKW [2003], LF [2006]
- numerical examples:
  - for $k = 512$ and $\eta = 0.25$, LF requires $q \simeq 2^{89}$
  - for $k = 768$ and $\eta = 0.01$, LF requires $q \simeq 2^{74}$

security models

- *passive attacks*: the adversary can only eavesdrop the conversations between an honest tag and an honest reader, and then tries to impersonate the tag
- *active attacks on the tag only* (a.k.a. active attacks in the *detection* model): the adversary first interacts with an honest tag (actively, but without access to the reader), and then tries to impersonate the tag
- *man-in-the-middle attacks* (a.k.a. active attacks in the *prevention* model): the adversary can manipulate the tag-reader conversation and observe whether the authentication is successful or not

	passive	active (TAG)	active (MIM)
HB	OK	KO	KO
HB^+	OK	OK	KO

a MIM attack against HB^+ [GRS 2005]

tag
k-bit secret
vectors x and y

reader
k-bit secret
vectors x and y

draw a random
k-bit blinding vector b

$$\xrightarrow{\quad b \quad}$$

draw a random
k-bit challenge a

$$\xleftarrow{a' = a \oplus \delta} \text{ Adv! } \xleftarrow{\;a\;}$$

compute
$z' = a' \cdot x \oplus b \cdot y \oplus \nu$
where $\Pr[\nu = 1] = \eta < \frac{1}{2}$

$$\xrightarrow{z' = z \oplus \delta \cdot x}$$

check $z' = a \cdot x \oplus b \cdot y$

accept? $\rightarrow \delta \cdot x = 0$
reject? $\rightarrow \delta \cdot x = 1$

- at each round, the noise bit ν_i is replaced by $\nu_i \oplus \delta \cdot x$

a MIM attack against HB$^+$ [GRS 2005]

- one authentication enables to retrieve one bit of x
- repeating the procedure with $|x|$ linearly independent δ's enables to derive x
- impersonating the tag is then easy (use $b = 0$)
- note that the authentication fails $\simeq$ half of the time: this may raise an alarm (hence the name detection-based model)

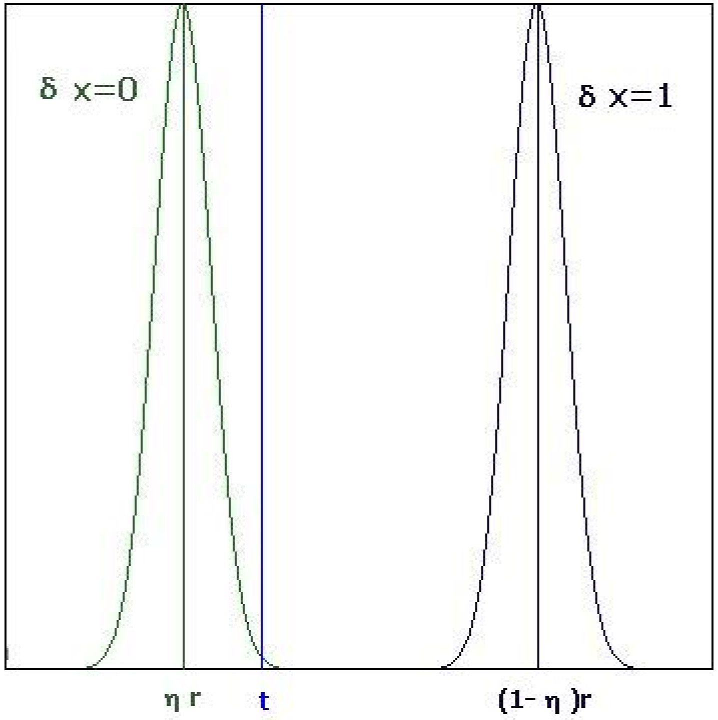

distribution of the number of errors

previous variants of HB$^+$

- three recent proposals aiming at thwarting MIM attacks:
  - HB-MP [Munilla and Peinado, 2007]
  - HB* [Duc and Kim, 2007]
  - HB^{++} [Bringer, Chabanne and Dottax, 2006]
- these three variants have been cryptanalysed recently [Gilbert, Robshaw and Seurin (FC ’08)]
- latest proposals . . .
  - Trusted-HB [Bringer, Chabanne, 2008]
  - PUF-HB [Hammouri, Sunar, ACNS 2008]

introducing RANDOM-HB$^{\#}$

tag

$k_X \times m$ and $k_Y \times m$-bit
secret **matrices** X and Y

reader

$k_X \times m$ and $k_Y \times m$-bit
secret **matrices** X and Y

Tag: draw a random
k_Y-bit blinding vector $\mathbf{b}$

$$\xrightarrow{\mathbf{b}}$$

Reader: draw a random
k_X-bit challenge $\mathbf{a}$

$$\xleftarrow{\mathbf{a}}$$

Tag: compute $\mathbf{z} = \mathbf{a} \cdot X \oplus \mathbf{b} \cdot Y \oplus \mathbf{\nu}$
where $\Pr[\mathbf{\nu}[i] = 1] = \eta < \frac{1}{2}$

$$\xrightarrow{\mathbf{z}}$$

Reader: check
$\mathrm{Hwt}(\mathbf{z} \oplus \mathbf{a} \cdot X \oplus \mathbf{b} \cdot Y) \leqslant t$

- one single pass
- accept iff the number of errors is less than some threshold $t > \eta m$

introducing RANDOM-HB$^{\#}$

- HB^{+} = many blinding vector/challenge pairs (a_i, b_i), one secret pair (x, y)
- RANDOM-HB$^{\#}$ = one blinding vector/challenge pair (a, b), many secret pairs (x_i, y_i)
- $\Rightarrow$ effectively reduces the communication complexity

security models: refinement

- recall the three models:
  - passive attacks (eavesdropping)
  - TAG attacks (the adversary can actively query an honest tag)
  - MIM attacks (man-in-the-middle attacks, the adversary can manipulate the tag-reader conversation and observe whether the authentication is successful or not)
- we refine the MIM model and define the **GRS-MIM** attacks: the adversary can only manipulate the messages *from the reader to the tag*
- HB^+ is susceptible to linear-time GRS-MIM attacks (hence the name)

security proof for RANDOM-HB$^{\#}$

- relies on the MHB-puzzle:

> Given q noisy samples $(\mathbf{a}_i, \mathbf{a}_i \cdot X \oplus \boldsymbol{\nu}_i)$, where X is a secret $k \times m$ matrix and $\Pr[\boldsymbol{\nu}_i[j] = 1] = \eta$, and a random challenge $\mathbf{a}$, find $\mathbf{a} \cdot X$.

- LPN is hard implies that no efficient adversary can guess $\mathbf{a} \cdot X$ with probability noticeably greater than $\frac{1}{2^m}$
- this is proved using results on *weakly verifiable puzzles* [CHS05] ; see the full version of the paper

security proof for RANDOM-HB$^{\#}$

- we reduce the security of RANDOM-HB$^{\#}$ in the GRS-MIM model to the LPN problem:

security against GRS-MIM attacks $\xrightarrow{3}$ security against TAG attacks $\xrightarrow{2}$ MHB puzzle $\xrightarrow{1}$ LPN problem

  - 1: weakly verifiable puzzles
  - 2: technical . . . (see the paper)
  - 3: if the adversary adds δ to the challenge α, the additional error vector $\delta \cdot X$ will have very high Hamming weight (because of the high minimal distance of X) and the reader will always reject
- general MIM adversaries are not handled by our security proof . . .

introducing HB$^{\#}$

- main drawback of RANDOM-HB$^{\#}$ is storage: $(k_X + k_Y) \cdot m$ bits, *i.e.* tens of Kbits
- HB$^{\#}$ is identical to RANDOM-HB$^{\#}$ except for the form of the matrices: it uses **Toeplitz** matrices
- reduces the storage requirements to $(k_X + k_Y + 2m - 2)$ bits: practical ($\simeq 1.5$ Kbits)

$$\begin{pmatrix} & & t_3 & t_2 & t_1 \\ & & & t_3 & t_2 \\ & \iddots & & & t_3 \\ t_{k+m-1} & & & & \end{pmatrix}$$

- Toeplitz matrices have good randomization properties: $(x \to x \cdot T)_T$ is a $1/2^m$-balanced function family (for any non-zero vector a, $a \cdot T$ is uniformly distributed)

security of HB$^{\#}$

- no formal reduction for HB$^{\#}$, only heuristic arguments using the previously mentioned property of Toeplitz matrices
- however we proved that

  HB$^{\#}$ secure against TAG attacks $\Rightarrow$ HB$^{\#}$ secure against GRS-MIM attack

general MIM attacks (!one-night slides!)

- at the rump session, Ouafi et al. outlined a (non GRS-) MIM attack against (RANDOM-)HB$^{\#}$
- idea: use an eavesdropped communication $(\alpha, \beta, \gamma = \alpha \cdot X \oplus \beta \cdot Y \oplus \nu)$ between the tag and the reader, add it to subsequent communications with a few more perturbations and use the reader decision to "remove" the noise ν
- breaks the proposed parameters with less authentications that we expected

general MIM attacks (!one-night slides!)

- asymptotic complexity?
- polynomial only for ill-chosen parameters, namely when the XOR of two random noise vectors is still below the threshold:

$$\eta_2 m < t, \quad \text{where} \quad \eta_2 = 2\eta(1-\eta)$$

- when the parameters are such that $\eta_2 m > t$, the attack becomes exponential
- this may be the missing condition to complete the security proof . . .

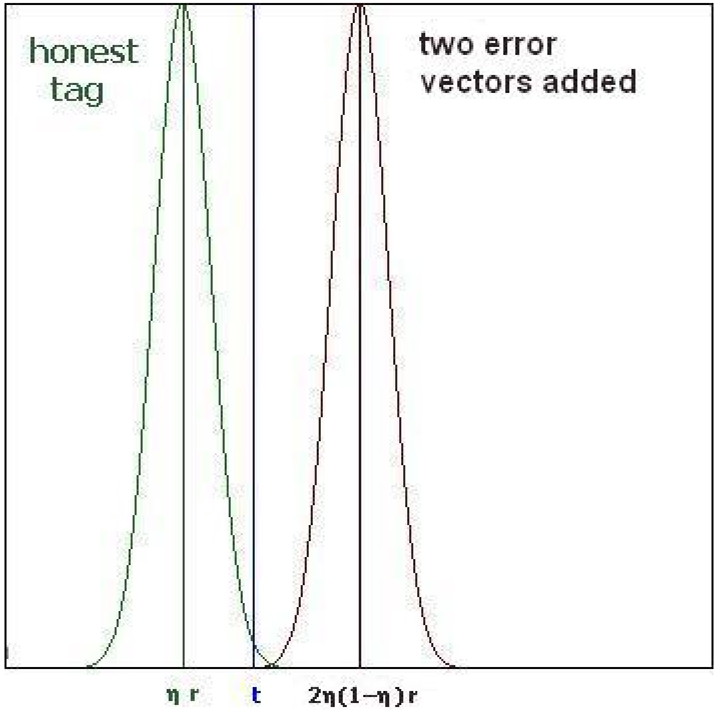

distribution of the number of errors

conclusions. . .

	HB^+	RANDOM-$\mathrm{HB}^\#$	$\mathrm{HB}^\#$
Storage (bits)	500	150 000	1 500
Transmission (bits/auth.)	50 000	1 000	1 000
Entropy gen. by the tag (bits/auth.)	25 000	500	500
TAG attack	OK	OK	? (prob. OK) $(*)$
GRS-MIM attack	KO	OK	? (prob. OK) (implied by $(*)$)
MIM attack	KO	??	??

- full paper available from http://eprint.iacr.org/2008/028

. . . and a trailer

- what other cryptographic primitive can you build from LPN?
- we propose a *symmetric encryption scheme* whose security can be reduced to the LPN problem
- this is LPN-C, to be presented at ICALP 2008 . . .

thanks for your attention!

questions?