Pruning and Extending the HB$^+$ Family Tree

Henri Gilbert, Matt Robshaw, and Yannick Seurin

Orange Labs

Outline

- HB^+ [Juels and Weis 05]: strengths and weaknesses
- Cryptanalysis of HB^+ variants
  - HB-MP [Munilla and Peinado 07]
  - HB^* [Duc and Kim 07]
  - HB^{++} [Bringer, Chabanne, and Dottax 06]
- A novel variant: $HB^\#$
  - RANDOM-$HB^\#$
  - $HB^\#$

Pervasive devices

- Issue: protection of single memory chips…
  - RFID Tags (Radio Frequency Identification)
  - very low cost cards without microprocessor

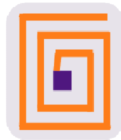

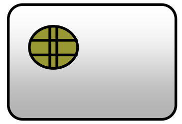

- … against chip cloning and replay attacks…
  - protection against duplication (tickets, banknotes)
  - protection against conterfeiting

- … by means of symmetric authentication
  - limited computing resource (~1000 gates/chip)

    => non-standard symmetric authentication

HB

[Hopper and Blum 01]: secure against passive attacks only

Tag

secret vector **x**

(k bits)

Reader

secrets vector **x**

(k bits)

r rounds

a (k-bit challenge vector) $\leftarrow$

$z = \mathbf{a} \bullet \mathbf{x} \oplus \nu$ (1-bit response) $\rightarrow$

where $\Pr[\nu=1] = \eta < \frac{1}{2}$ (e.g. ¼)

accept round if

$z = \mathbf{a} \bullet \mathbf{x}$

Authentication is successful iff at most t rounds (e.g. ηr) are rejected

The HB$^+$ protocol

[Juels and Weis, Crypto'05]: sequential version

[Katz and Shin, Eurocrypt'06]: parallel version

Tag

secrets **x** and **y**

[k bits each]

Reader

secrets **x** and **y**

[k bits each]

r rounds

b (k-bit blinding vector) $\longrightarrow$

$\longleftarrow$ **a** (k-bit challenge vector)

$z = \mathbf{a \bullet x} \oplus \mathbf{b \bullet y} \oplus \nu$ (1bit) $\longrightarrow$

Pr [noise bit ν =1] = η (e.g. ¼)

accept round if

$z = \mathbf{a \bullet x} \oplus \mathbf{b \bullet y}$

Authentication is successful iff at most t rounds (e.g. ηr) are rejected

k: vectors length [224]; r: #rounds [60]; η:noise rate [1/4]; t: acceptance threshold [25]

Security of HB$^+$

☺ HB$^+$ is provably secure against active attacks [JW05, KS06]

Reduction to the conjectured intractability of the LPN problem:

Given:	a known random $q \times k$ matrix $\mathbf{A}$
	a noise parameter η used to draw bits of a q-bit noise vector $\boldsymbol{\nu}$
	a q-bit vector $\mathbf{z} = \mathbf{A} \cdot \mathbf{x} \oplus \nu$ (k-bit vector $\mathbf{x}$ and ν are unknown)
Find:	the k-bit vector $\mathbf{x}$

- best solving algorithms: [BKW03], later on improved in [LF06]

  $\Rightarrow$ The initially suggested value $k \approx 250$ [JW05, KS06] is too small.

☹ The security model underlying the proofs is restricted

- and there is an efficient attack outside from this model (see hereafter)

Practical limitations of HB⁺

	r	η	k	false reject rate P_{FR}	false accept rate P_{FA}	trans.cost (bits) initial k	k=512
[JW]	60	0.25	224	43% (!)	6×10^{-6}	26984	82000
[KS]	40	0.125	200	38% (!)	7×10^{-9}	16040	41000

☹ Error rates: false rejection rates P_{FR} are unacceptably high

- this is partly due to the unfortunate choice $t=\eta r$
- $t > \eta r$ improves the P_{FR}-P_{FA} balance but the order of magnitude of $\max(P_{FA}, P_{FR})$ remains too high (1%)

☹ Transmision costs are unacceptably high

- 2 k-bit vectors have to be exchanged to get a 1-bit response
- transmission payload: $r(2k+1)$ bits

MIM Attack on HB$^+$ [GRS05]

If an adversary can (1) modify challenge vectors and (2) know wheter auth. succeeds then any linear comb. $\delta \bullet x$ of the x bits can be derived:

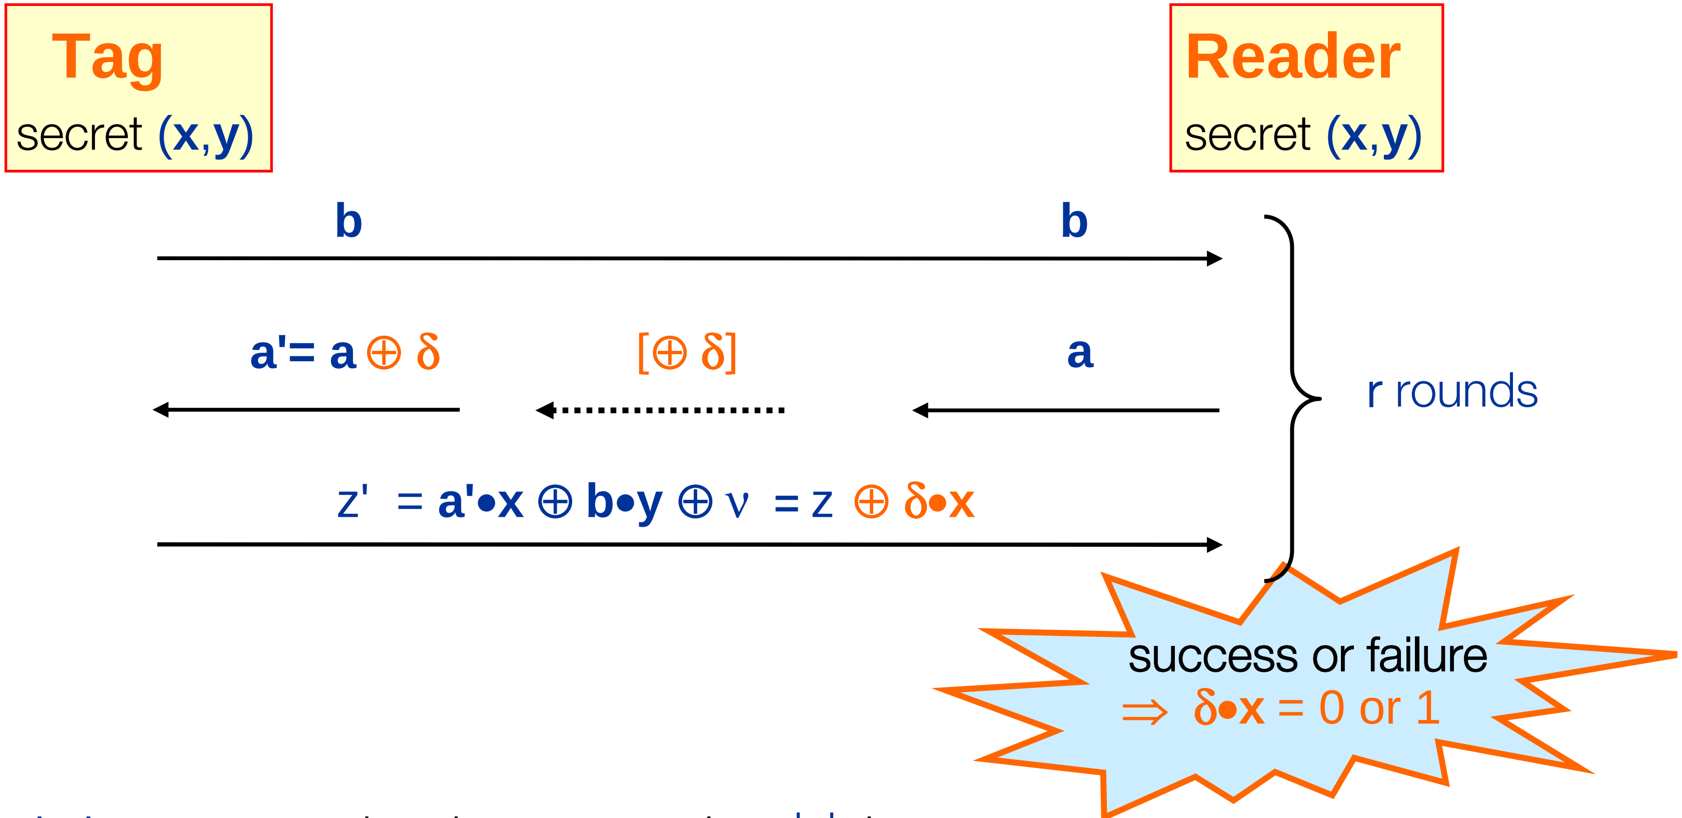

→ To derive x: repeat the above procedure |x| times

→ To derive y: now trivial using a false tag (use constant b)

→ To impersonate the tag: use x and y (or even x only).

Security models

- DET model (detection-based) used in security proofs of [JW],[KS]
  - Phase 1: adversary first interacts q times with a legitimate tag.
  - Phase 2: she interacts once with a reader to impersonate the tag.
- GRS-MIM model (GRS-like man in the middle)
  - Phase 1: adversary interacts with a legitimate tag and a legitimate reader during q authentication exchanges and can observe all messages:
    - she can modify any message sent by the reader to the tag.
    - she has access to the authentication success/failure information.
  - Phase 2: she interacts once with a reader to impersonate the tag.
- MIM model
  - Same as GRS-MIM except adversary can modify all tag-reader messages

(MIM security $\Rightarrow$ GRS-MIM security $\Rightarrow$ DET security)

HB-MP [Munilla and Peinado 07]

HB-MP': simplified version of HB-MP

Tag		**Reader**
secret **x** (k bits)		secret **x** (k bits)
compute $z = \mathbf{a} \bullet \mathbf{x} \oplus \nu$ where **Pr** $[\nu = 1] = \eta < 1/2$	$\xleftarrow{\mathbf{a}}$	choose $\mathbf{a} \in_R \{0,1\}^k$
select $\mathbf{b} \neq \mathbf{a}$ with $\mathbf{b} \bullet \mathbf{x} = z$	$\xrightarrow{\mathbf{b}}$	accept round if $\mathbf{b} \bullet \mathbf{x} = \mathbf{a} \bullet \mathbf{x}$

Authentication is successful iff at most t rounds are rejected

HB-MP

Aim: immunity against passive and active attacks, including GRS-like attacks

Tag		Reader
k-bit secrets $\mathbf{x}, \mathbf{y}$		k-bit secrets $\mathbf{x}, \mathbf{y}$
At round $i \in [0;r-1]$:	$\xleftarrow{\mathbf{a}\ (m < k \text{ bits})}$	choose $\mathbf{a} \in_R \{0,1\}^m$
$\mathbf{x} = \text{rotate}(\mathbf{x}, \mathbf{y}_i)$		
compute $\mathbf{z} = \mathbf{a} \bullet \lfloor \mathbf{x} \rfloor_m \oplus \nu$		
select $\mathbf{b} \neq \mathbf{a}$ with $\mathbf{b} \bullet \lfloor \mathbf{x} \rfloor_m = \mathbf{z}$	$\xrightarrow{\mathbf{b}\ (m < k \text{ bits})}$	accept round if $\mathbf{b} \bullet \lfloor \mathbf{x} \rfloor_m = \mathbf{a} \bullet \lfloor \mathbf{x} \rfloor_m$

Notation: $\lfloor \mathbf{x} \rfloor_m$ = m least significant bits of $\mathbf{x}$;

rotate($\mathbf{x}$,ρ) = bitwise rotation of $\mathbf{x}$ by ρ bits to the left.

$\mathbf{y}_i$ = bit i of $\mathbf{y}$;

Authentication is successful iff at most t rounds are rejected

A passive attack against HB-MP

The verification equations can be written: $(a^i \oplus b^i) \bullet \lfloor x^i \rfloor_m = 0$

(where a^i, b^i, and x^i denote the values of a, b and x at rount i.)

- Step 1: observation of one authentication exchange of a legitimate tag:

  $\longrightarrow$ record (a^i, b^i) pairs or $a^i \oplus b^i$ values

- Step 2: impersonation of the tag:

  $\longrightarrow$ on challenge a'^i answer b'^i such that $a'^i \oplus b'^i = a^i \oplus b^i$

  i .e. $b'i = a'^i \oplus a^i \oplus b^i$

The attack works exactly in the same way against HB-MP'.

HB* [Duc and Kim 07]

Aim: resistance to all active attacks, including GRS-like attacks

Tag		**Reader**
k-bit secrets **x**, **y**, and **s**		k-bit secrets **x**, **y** and **s**
draws $\nu \in \{0,1\} \mid \Pr[\nu = 1] = \eta$		
draws $\gamma \in \{0,1\} \mid \Pr[\gamma = 1] = \eta'$		
choose $\mathbf{b} \in_R \{0,1\}^k$ $w = \mathbf{b} \bullet \mathbf{s} \oplus \gamma$	$\xrightarrow{\mathbf{b},\ w}$	
	$\xleftarrow{\mathbf{a}}$	choose $\mathbf{a} \in_R \{0,1\}^k$
if $\gamma = 0$: $z = \mathbf{a} \bullet \mathbf{x} \oplus \mathbf{b} \bullet \mathbf{y} \oplus \nu$ else: $z = \mathbf{a} \bullet \mathbf{y} \oplus \mathbf{b} \bullet \mathbf{x} \oplus \nu$	$\xrightarrow{z}$	if $\mathbf{b} \bullet \mathbf{s} = w$: check $z = \mathbf{a} \bullet \mathbf{x} \oplus \mathbf{b} \bullet \mathbf{y}$ else: check $z = \mathbf{a} \bullet \mathbf{y} \oplus \mathbf{b} \bullet \mathbf{x}$

Authentication is successful iff at most t rounds are rejected

A MIM attack against HB* (1/2)

The attack is a close variant of the GRS attack against HB+.

At each round, the challenge vector **a** is replaced by **a** $\oplus$ **δ**, and consequently:

- If γ=0: z is replaced by z $\oplus$ δ • x
- If γ=1: z is replaced by z $\oplus$ δ • y

The ratio between both events is governed by η'.

1 If η' is sufficiently small ($\eta' < \frac{t - \eta r}{r(1-2\eta)}$) the original HB+ attack still works.

The disturbed authentication is:

- likely to succeed if **δ•x** = 0
- unlikely to succeed if **δ•x** = 0.

2 Otherwise

The disturbed authentication is:

- likely to succeed if δ•x = 0 and δ•y = 0 (z is then never affected)
- unlikely to succeed if δ•x =1 or δ•y =1

A MIM attack against HB* (2/2)

Step 1: find lin. ind. values $\delta_1, \delta_2, \delta_{k-2}$ such that the authentication succeeds.

$\rightarrow$ with high proba. $(\delta_1, \delta_2, \delta_{k-2})$ is a basis of $\langle x,y\rangle^{\perp}$, i.e. $\langle\delta_1, \delta_2, \delta_{k-2}\rangle^{\perp} = \langle x, y\rangle$.

$\rightarrow$ we get the unordered set $\{c_1,c_2,c_3\} = \{x, y, x \oplus y\}$

Step 2: identify $x \oplus y$ in $\{c_1, c_2, c_3\}$

query honest tag with $a = b$ at each round

$\Rightarrow z = a \bullet (x \oplus y) \oplus \nu$ at each round

$\rightarrow \#\{\text{rounds} \mid z = a \bullet c_i\}$ is maximal for $c_i = x \oplus y$

Step 3: first impersonation attempt with success proba. ½

Step 4: later impersonation attempts have success proba. ≈ 1

Low complexity: approximately 4k authentications required

HB^{++} [Bringer, Chabanne, and Dottax 05]

Aim: keep HB^+ security in restricted model and prevent MIM attacks

Outline:

- uses a k-bit to k-bit function f based on a [5-bit] s-box S

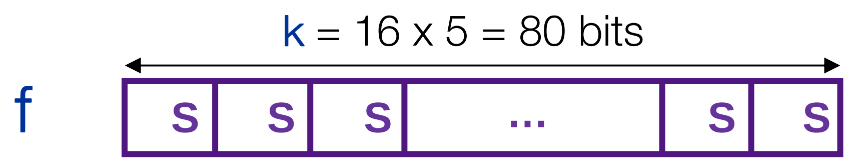

- 4 secret key vectors x, x', y, y' instead of x and y
- 2 response bits instead of 1 at each round, namely (at round i)
  - $z = a \bullet x \oplus b \bullet y \oplus \nu$ as before
  - $z' = f(a)^{<<i} \bullet x' \oplus f(b)^{<<i} \bullet y' \oplus \nu'$ s-box and rotation by i bits
- x, x', y, y' are renewed at each authentication

HB^{++} [Bringer, Chabanne, and Dottax 05]

Tag
secret **Z**

Reader
secret **Z**

Stage 1: renewal of authentication keys x, x', y, y'

choose **B** $\in_R \{0,1\}^k$ $\xrightarrow{\mathbf{B}}$

$\xleftarrow{\mathbf{A}}$ choose **A** $\in_R \{0,1\}^k$

Tag: (x, x', y, y')= h (Z, A, B)

Reader: (x, x', y, y')= h (Z, A, B)

Stage 2: actual authentication

choose **b** $\in_R \{0,1\}^k$ $\xrightarrow{\mathbf{b}}$

$\xleftarrow{\mathbf{a}}$ choose **a** $\in_R \{0,1\}^k$

compute:

$z = \mathbf{a} \bullet \mathbf{x} \oplus b \bullet y \oplus \nu$

$z' = (f(a)^{<<i}) \bullet x' \oplus (f(b)^{<<i}) \bullet y' \oplus \nu'$

$\xrightarrow{(z, z')}$

accept round if:

$z = a \bullet x \oplus b \bullet y$ and

$z' = (f(a)^{<<i}) \bullet x' \oplus (f(b) << i) \bullet y'$

Authentication is successful iff at most t rounds are rejected

Attack on HB^{++} without key renewal

- Attack scenario **almost identical to the GRS attack** on HB^{+}
- But the adversary only disturbs the challenge vectors of the $s < r$ **first rounds** using a fixed **disturbance vector** δ. Other rounds are not disturbed.

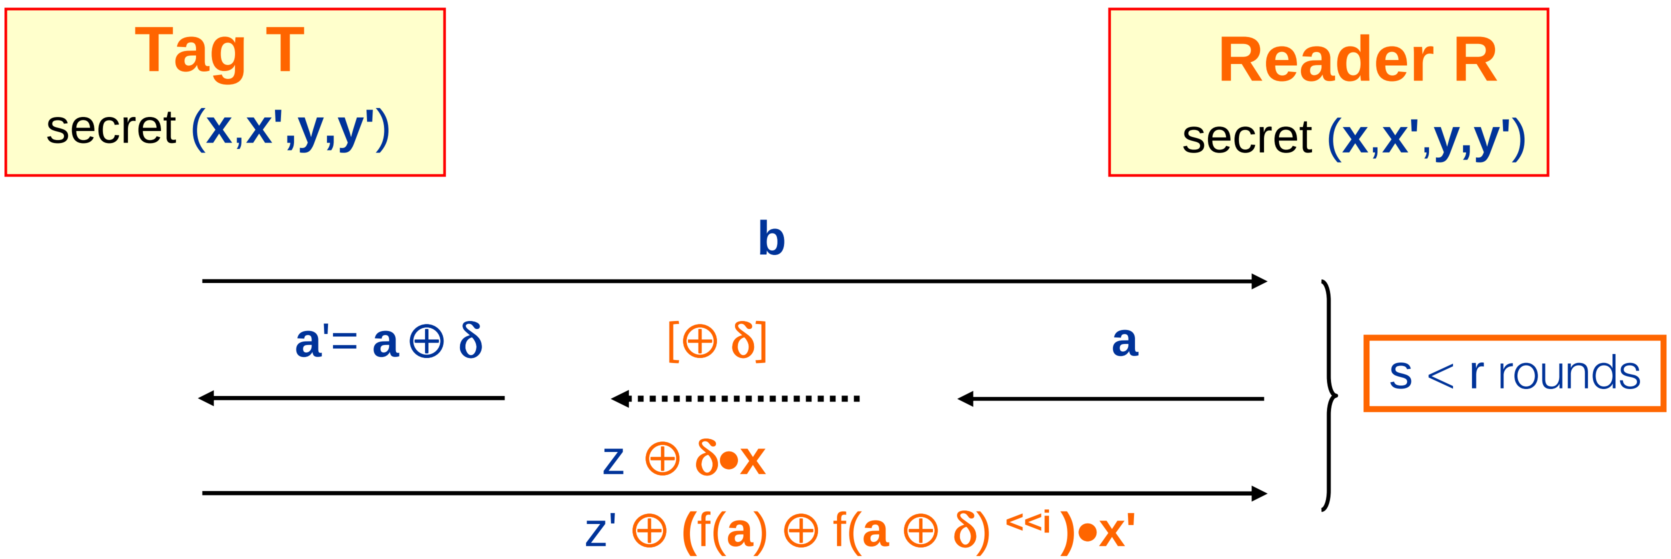

- If s is well chosen
  - p_0 = **Pr[R accepts| $\delta \bullet x$ = 0]** is non-negligible
  - p_1 = **Pr[R accepts| $\delta \bullet x$ = 1]** is negligible

  $\rightarrow$ therefore **Pr[$\delta \bullet x \neq 0$ |R accepts] = $p_1/(p_0+p_1)$** is very small.
- Example **If k=80, r=80, t=30, η=1/4, for s=40:**

  Pr[R accepts] $\approx$30% and **Pr[$\delta \bullet x \neq 0$ |R accepts] $\approx$ 0.007**

Detail of the function h [WH: KYS05]

- Inputs:
  - $\mathbf{Z = (Z_1, \ldots, Z_{48})}$ 48 16-bit words = 768 bits
  - $\mathbf{M = (A,B)}$ 2 80-bit words = 160 bits
    $\mathbf{= (M_1, \ldots, M_{10})}$ 10 16-bit words
- Output:
  - $\mathbf{h(Z, A, B) = \underbrace{(g_{Z1\ldots Z10}(M),\ g_{Z3\ldots Z13}(M), \ldots,\ g_{Z39\ldots Z48}(M))}_{\text{20 16-bit words}} = (x,x',y,y')}$ 4 80-bit words
- g is defined as follows:

  $\mathbf{g}$: $\{0,1\}^{16x10} \times \{0,1\}^{16x10} \rightarrow \{0,1\}^{16}$

  GF(2^{16}) constant

  $$\mathbf{g_{K1, \ldots, K10}(M_1, \ldots, M_{10}) = \sum_{i=1}^{5} (M_{2i-1} + K_{2i-1})(M_{2i} + K_{2i}) \cdot c_i}$$

$\rightarrow$ over **GF(2^{16})**: if (A,B) is known, each 16-bit component of **h(Z,A,B)** is a known affine function of 15 unknown 16-bit values $\mathbf{Z_{2j-1}.Z_{2j},\ Z_{2j-1}, Z_{2j}}$

$\rightarrow$ over **GF(2)**: if (A,B) is known, each bit of **h(Z,A,B)** is a known affine function of 16x15 = 240 expanded key bits.

Attack on the complete HB^{++} scheme

- Step 1: we disturb the authentication protocol with δ values that fit one single 16-bit word of x (e.g. $\delta = (\delta_0,\ldots,\delta_{15}, 0, \ldots 0)$
  - each successful authentication provides one equation $\boldsymbol{\delta \bullet x} \approx \boldsymbol{0}$ in one word of x
    i.e. one approximate equation in 240 expanded key bits
  - → 5 low complexity LPN problems: 240 unknowns, $\varepsilon < 1\%$
  - → we derive the expanded key part allowing to derive x

- Step 2: we derive the expanded key part allowing to derive x'
  - → we get and solve 5 additional LPN problems.

- Step 3: we record the quartets (a, b, z, z') of a successful authentication, we can reuse the vectors b and correct z and z' according to Δa to impersonate the tag.

Complexity estimate: **if k=80, r=80, t=30, η=1/4, for s= 40,** $\varepsilon \approx 1\%$

- Authentications needed: $4 \times 10 \times 2^{30} \approx 2^{35}$
- Complexity: $4 \times 2^{41} = 2^{44}$

RANDOM-HB#

Aim: render HB+ resistant to MIM attacks

- replace **x** by a random k_x x m binary matrix X
- replace **y** by a random k_y x m binary matrix Y
- authentication has now one single round

tag secret **X, Y**		reader secret **X, Y**
	b (k_y-bit blinding vector) →	
	← **a** (k_x-bit challenge vector)	
Draw m-bit noise $\nu = (\nu_i)$ s.t. $\forall i\ \Pr[\nu_i = 1] = \eta < 1/2$	$\mathbf{z} = \mathbf{a}\cdot\mathbf{X} \oplus \mathbf{b}\cdot\mathbf{Y} \oplus \nu$ → (m-bit response vector)	Check $Hwt(\mathbf{z} \oplus \mathbf{a}\cdot\mathbf{X} \oplus \mathbf{b}\cdot\mathbf{Y}) \leq t$ where $t \in [\eta m, m/2[$

☺ Transmission costs and error rates become realistic

due to the better balance between challenge and response lenghts

☺ Provable security against a larger class of attacks

☹ Storage requirements for matrices X and Y

→ solved by HB#

Security of RANDOM-HB#

- HB+ — HB puzzle

Given:	q k-bit random vectors a_i q noisy bits $a_i \cdot x^t + \nu_i$ where $\Pr[\nu_i=1] = \eta < ½$ a random challenge vector a	Guess $a \cdot x^t$

LPN hardness $\xrightarrow{\text{[HB,JW]}}$ HB puzzle (1-1/2) hardness $\xrightarrow{\text{[JW]}}$ HB+ DET security

- Here: — MHB puzzle

Given:	q k-bit random vectors a_i q noisy m-bit vectors $a_i \cdot X + \nu_i$ where $\Pr[\nu_{i\,j}=1] = \eta < ½$ a random challenge vector a	Guess $a \cdot X$

LPN hardness $\xrightarrow{\approx\text{[CHS05]}}$ MHB puzzle $(1-1/2^m)$ hardness $\longrightarrow$ RANDOM-HB# DET security $\longrightarrow$ RANDOM-HB# GRS-MIM security

+ informal security argument in the general MIM model

$HB^{\#}$ (1/2)

- Definition: a $k \times m$ matrix M is a Toeplitz matrix iff it has constant coefficients on all upper left to bottom right diagonals.

  $\rightarrow$ M is determined by the $k+m-1$ coefficients of column 1 and row 1

Toeplitz matrix

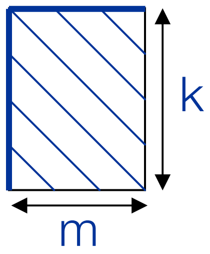

- $HB^{\#}$ is identical to RANDOM-$HB^{\#}$ (the tag's answer is still: $z = a \cdot X \oplus b \cdot Y \oplus \nu$)

  ... except X and Y are now random binary Toeplitz matrices.

  $\rightarrow$ low storage requirements: k_x+k_y+m-2 bits instead of $(k_x+k_y)m$

  $\rightarrow$ efficient on tag computations

HB$^{\#}$ (2/2)

Security

- Conj: the Toeplitz-MHB puzzle is hard and HB$^{\#}$ is secure in the **DET model**
- Th: if HB$^{\#}$ is secure in the DET model, then (under easy to meet conditions on the parameters set) it is also secure in the **GRS-MIM model**
- Strong arguments for HB$^{\#}$ security in the **general MIM model** using the fact that the set of k x m Toeplitz matrices is a $1/2^m$-balanced family of hash functions.

Parameter values for HB$^{\#}$

k_X	k_Y	m	η	t	P_{FR}	P_{FA}	com. (bits)	stor. (bits)
80	512	1164	0.25	405	2^{-45}	2^{-83}	1756	2918
80	512	441	0.125	113	2^{-45}	2^{-83}	1033	1472

Conclusions

- $HB^{\#}$ attains a truly practical performance profile

- some further optimisations of $HB^{\#}$ might be of practical value:
  - test weight of noise vector ν before using it and reduce m
  - use larger noise level η and reduce k_Y?

- the use of LPN and matricial variants (MHB /Toeplitz MHB) in symmetric cryptography deserves further exploration.