unrestricted

Good Variants of HB^+ are Hard to Find

(The Cryptanalysis of HB^{++}, HB^* and HB-MP)

Henri Gilbert, Matt Robshaw, and Yannick Seurin

Financial Crypto 2008 – January 29, 2008

ftgroup

orange™

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 (GPS)
  - 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
- cryptanalysis of HB-MP
- cryptanalysis of HB^*
- cryptanalysis of HB^{++}
- conclusions . . . and a trailer

the ancestor HB [Hopper and Blum 2001]

tag		**reader**
k-bit secret vector $\mathbf{x}$		k-bit secret vector $\mathbf{x}$
	$\xleftarrow{\mathbf{a}}$	draw a random k-bit challenge $\mathbf{a}$
compute $z = \mathbf{a} \cdot \mathbf{x} \oplus \nu$ where ν is a noise bit $\Pr[\nu = 1] = \eta < \frac{1}{2}$	$\xrightarrow{z}$	check $z = \mathbf{a} \cdot \mathbf{x}$

- 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		**reader**
k-bit secret vectors $\mathbf{x}$ and $\mathbf{y}$		k-bit secret vectors $\mathbf{x}$ and $\mathbf{y}$
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

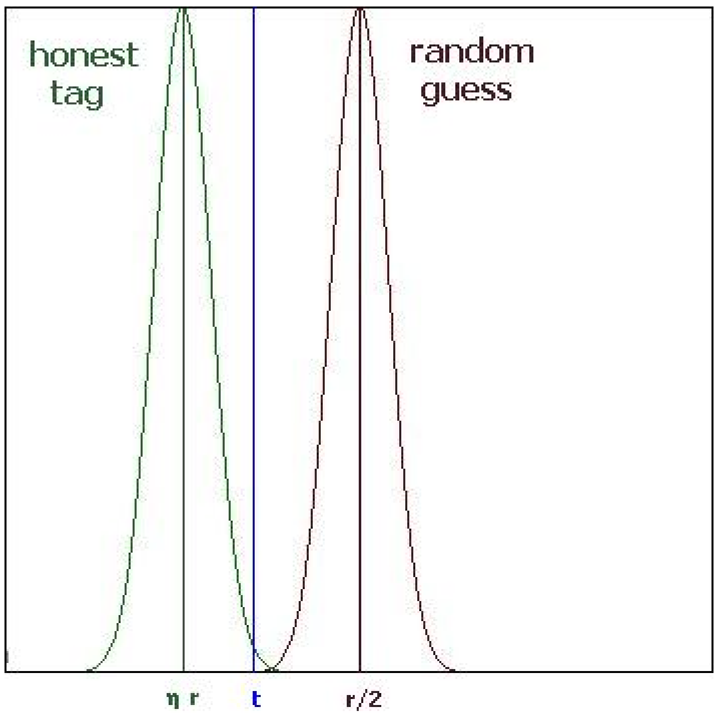

distribution of the number of errors

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 interact 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 man-in-the-middle attack against HB$^+$ [GRS 2005]

tag		reader
k-bit secret vectors $\mathbf{x}$ and $\mathbf{y}$		k-bit secret vectors $\mathbf{x}$ and $\mathbf{y}$
draw a random k-bit blinding vector $\mathbf{b}$	$\xrightarrow{\mathbf{b}}$	
	$\xleftarrow{\mathbf{a}'=\mathbf{a}\oplus\delta}$ Adv! $\xleftarrow{\mathbf{a}}$	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 $z' = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$
		accept? $\rightarrow \delta \cdot \mathbf{x} = 0$
		reject? $\rightarrow \delta \cdot \mathbf{x} = 1$

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

a man-in-the-middle attack against HB^+ [GRS 2005]

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

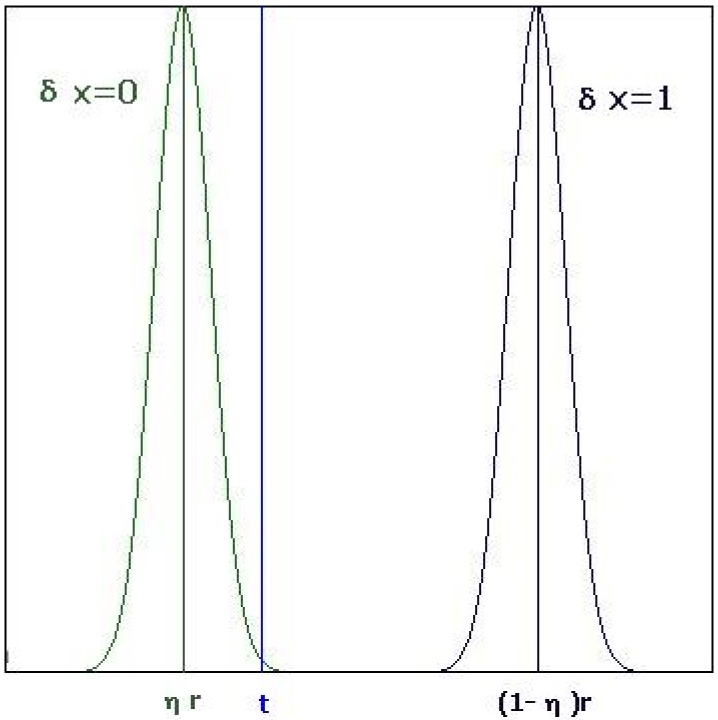

distribution of the number of errors

we need a variant of HB$^+$ resisting MIM attacks

- three recent proposals:
  - HB-MP
  - HB*
  - HB^{++}
- we show how to cryptanalyse them

cryptanalysis of HB-MP

- HB-MP was introduced by Munilla and Peinado
- aim: obtain a more simple (2-pass) protocol but at least as secure as HB^+
- however, there is a *passive* attack against HB-MP
- please see the paper for the details

HB* [Duc and Kim 2007]

tag	reader
k-bit secret vectors **x**, **y** and **s**	k-bit secret vectors **x**, **y** and **s**

draw a random $\mathbf{b} \in_R \{0,1\}^k$
draw $\gamma \in_R \{0,1\} \mid \Pr[\gamma = 1] = \eta'$
compute $w = \mathbf{b} \cdot \mathbf{s} \oplus \gamma$

$\xrightarrow{(\mathbf{b}, w)}$

$\xleftarrow{\mathbf{a}}$ draw a random $\mathbf{a} \in_R \{0,1\}^k$

if $\gamma = 0$ compute
$z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y} \oplus \nu$
else compute $z = \mathbf{a} \cdot \mathbf{y} \oplus \mathbf{b} \cdot \mathbf{x} \oplus \nu$

$\xrightarrow{z}$

if $\mathbf{b} \cdot \mathbf{s} = \mathbf{w}$ check $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$
else check $z = \mathbf{a} \cdot \mathbf{y} \oplus \mathbf{b} \cdot \mathbf{x}$

- this is repeated for r rounds
- the authentication is successful iff at most t rounds have been rejected

a MIM attack on HB*

- try the GRS attack: add a constant δ to the challenges $\mathbf{a}$; then:
- if η' is to low, most of rounds will use equation $\mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$: this is equivalent to HB$^+$ (true when $\eta' \leqslant \frac{t-\eta r}{r(1-2\eta)}$)
- conversely, if η' is close to $1/2$, the following will happen:
  - if $\delta \cdot \mathbf{x} = 0$ and $\delta \cdot \mathbf{y} = 0$ then the reader will accept
  - in all other cases the reader will reject ($\delta \cdot \mathbf{x} = 1$ or $\delta \cdot \mathbf{y} = 1$)
  - hence the adversary is able to learn the vector space $< \mathbf{x}, \mathbf{y} >$

a MIM attack on HB*

- the attack proceeds as follows:
  - find lin. ind. values $\delta_1, \ldots, \delta_{k-2}$ such that the authentication succeeds
  - with overwhelming probability this gives the unordered set $\{\mathbf{c_1}, \mathbf{c_2}, \mathbf{c_3}\} = \{\mathbf{x}, \mathbf{y}, \mathbf{x} \oplus \mathbf{y}\}$
  - identify $\mathbf{x} \oplus \mathbf{y}$ in $\{\mathbf{c_1}, \mathbf{c_2}, \mathbf{c_3}\}$ by querying the honest tag with $\mathbf{a} = \mathbf{b}$ at each round $\Rightarrow z = \mathbf{a} \cdot (\mathbf{x} \oplus \mathbf{y}) \oplus \nu$
  - first impersonation succeeds with proba $1/2$
  - following impersonations succeed with proba 1
- linear complexity: $O(4k)$ authentications are required

HB^{++} [Bringer, Chabanne, and Dottax 2005]

tag		reader
k-bit session secret vectors $\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}'$		k-bit session secret vectors $\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}'$
draw a random $\mathbf{b} \in_R \{0,1\}^k$	$\xrightarrow{\mathbf{b}}$	
	$\xleftarrow{\mathbf{a}}$	draw a random $\mathbf{a} \in_R \{0,1\}^k$
compute $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y} \oplus \nu$ and $z' = (f(\mathbf{a})^{\lll i}) \cdot \mathbf{x}' \oplus (f(\mathbf{b})^{\lll i}) \cdot \mathbf{y}' \oplus \nu'$	$\xrightarrow{(z,z')}$	check $z = \mathbf{a} \cdot \mathbf{x} \oplus \mathbf{b} \cdot \mathbf{y}$ and $z' = (f(\mathbf{a})^{\lll i}) \cdot \mathbf{x}' \oplus (f(\mathbf{b})^{\lll i}) \cdot \mathbf{y}'$

- this is repeated for r rounds
- let N (resp. N') be the number of errors on z (resp. z'), the authentication is successful iff $N \leqslant t$ and $N' \leqslant t$

HB^{++} [Bringer, Chabanne, and Dottax 2005]

- uses a k-bit to k-bit permutation f made of a layer of 5-bit S-box S to compute the second response bit $z' = (f(\mathbf{a})^{\ll i}) \cdot \mathbf{x}' \oplus (f(\mathbf{b})^{\ll i}) \cdot \mathbf{y}'$
- the secrets $\mathbf{x}$, $\mathbf{y}$, $\mathbf{x}'$, $\mathbf{y}'$ are renewed before each authentication with a master secret $\mathbf{Z}$ and a universal hash function h

tag		reader
K-bit master secret $\mathbf{Z}$		K-bit master secret $\mathbf{Z}$
draw a random $\mathbf{B} \in_R \{0,1\}^{K'}$	$\xrightarrow{\mathbf{B}}$	
	$\xleftarrow{\mathbf{A}}$	draw a random $\mathbf{A} \in_R \{0,1\}^{K'}$
compute $(\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}') = h(\mathbf{Z}, \mathbf{A}, \mathbf{B})$		compute $(\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}') = h(\mathbf{Z}, \mathbf{A}, \mathbf{B})$

a MIM attack on HB^{++}: phase 1

- aims at gathering approximate equations on (a subset of the bits of) $\mathbf{x}$
- a simple GRS attack fails: the error vector on z'_i is

$$\nu'_i \oplus (f(\mathbf{a_i} \oplus \delta) \oplus f(\mathbf{a_i}))^{\lll i} \cdot \mathbf{x}$$

$\Rightarrow$ randomized, hence $N' \simeq r/2$ and the reader always rejects

- however, what happens if one disturbs $s < r$ rounds?

a MIM attack on HB^{++}: phase 1

- if s is to low, the distributions of N when $\delta \cdot \mathbf{x} = 0$ and when $\delta \cdot \mathbf{x} = 1$ are not well distributed around t
- if s is to high, the expected value of N' is to high and the reader always rejects
- but for s such that $E(N') \simeq t$, it's OK!
- when the reader accepts $(p = 1/4)$, $\delta \cdot \mathbf{x} = 0$ with high probability

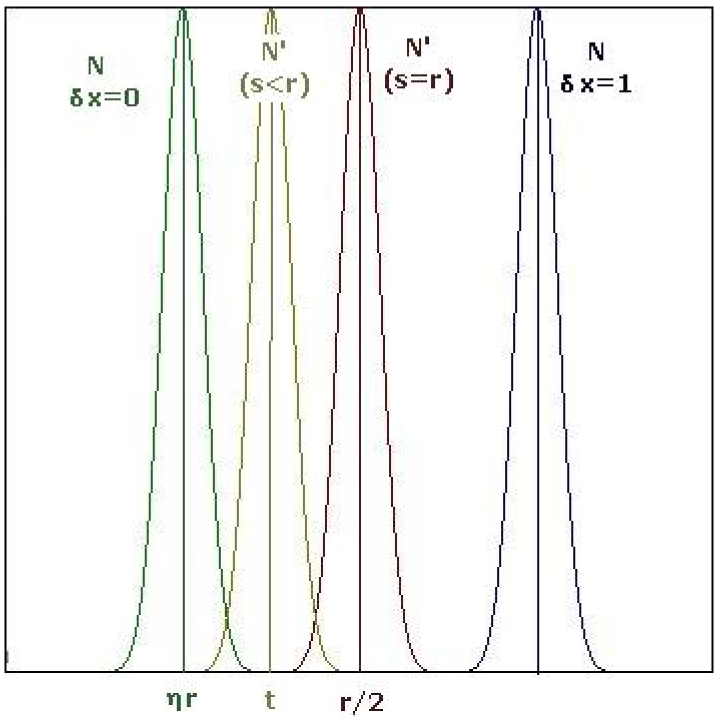

- example: for $k = 80, r = 80, \eta = 0.25$, $t = 30$, by disturbing $s = 40$ rounds, $\Pr[\text{false guess}] \simeq 0.01$

a MIM attack on HB^{++}: phase 2

- getting into the details of $h(\mathbf{Z}, \mathbf{A}, \mathbf{B})$:
  - $\mathbf{Z} = (\mathbf{Z}_1, \ldots, \mathbf{Z}_{48})$: 48 16-bit words = 768 bits in total
  - $\mathbf{M} = (\mathbf{A}, \mathbf{B}) = (\mathbf{M}_1, \ldots, \mathbf{M}_{10})$: 10 16-bit words = 160 bits in total
  - $h(\mathbf{Z}, \mathbf{A}, \mathbf{B}) = (\mathbf{x}, \mathbf{y}, \mathbf{x}', \mathbf{y}')$
    $= (g_{Z_1 \ldots Z_{10}}(\mathbf{M}), g_{Z_3 \ldots Z_{13}}(\mathbf{M}), \ldots, g_{Z_{39} \ldots Z_{48}}(\mathbf{M}))$: 20 16-bit words
- if $(\mathbf{A}, \mathbf{B})$ is known, each of these 20 16-bit words is an affine function of 160 $\mathbf{Z}$ bits and 80 quadratic functions of $\mathbf{Z}$ bits = 240 expanded key bits
- thanks to the approximate equations of phase 1, solve an LPN problem with key length 240 and low noise parameter

a MIM attack on HB^{++}: summary

- step 1: disturb the authentication protocol with δ's affecting one single 16-bit word of $\mathbf{x}$ and get approximate equations on the secret bits allowing to derive $\mathbf{x}$ $\Rightarrow$ 5 LPN problems to solve
- step 2: derive the expanded key bits allowing to derive $\mathbf{x}'$ (5 additional LPN problems)
- step 3: impersonate the tag by reusing previous blinding vectors **b**
- complexity estimate: for for $k = 80, r = 80, \eta = 0.25, t = 30$, by disturbing $s = 40$ rounds, $4 \times 10 \times 2^{30} \simeq 2^{35}$ authentications needed

conclusions...

	passive	active (TAG)	active (MIM)
HB	OK	KO	KO
HB^+	OK	OK	KO
HB-MP	KO	KO	KO
HB^*	OK	OK	KO
HB^{++}	OK	OK	KO
?	OK	OK	OK

- HB^+ remains the most attractive member of the family...
- but still has some practical problems: MIM attack, high communication complexity (50 to 100 Kbit / auth.)
- a (simple) variant resistant to MIM attacks would be highly interesting

…and a trailer

- introducing: HB$^{\#}$ [Gilbert, Robshaw, and Seurin, Eurocrypt 2008]
- main idea: generalize the form of the secrets from vectors to matrices
- main advantages: reduced communication complexity, *provable security* against a large class of MIM attacks
- drawback: more storage required, but remains practical
- see you in Istanbul for more details ;-) (in the meanwhile, the paper is available on e-print)

thanks for your attention!

questions?