# Distinguishing attacks on ISAAC

Jean-Philippe Aumasson

FHNW, 5210 Windisch, Switzerland

**Abstract.** This paper presents two strong distinguishers for the deterministic random bit generator ISAAC, requiring $2^{48}$ and $2^{64}$ samples of respectively 64 and 32 bits, based on the observation that more than $2^{8\,167}$ initial states among the $2^{8\,192}$ ones induce a strongly non-uniform distribution of the bits produced at the first round of the algorithm. A previous attack on ISAAC presented at Asiacrypt'06 by Paul and Preneel is demonstrated to be non relevant, since relies on an erroneous algorithm. The results of this paper stress the unsecurity of ISAAC, both as a pseudo-random generator and as a stream cipher. A modification of the algorithm is proposed to fix the weaknesses discovered.

ISAAC [3] is a deterministic random bits generator presented at FSE'96 by Jenkins, who claims that it has "*no bad initial states, not even the state of all zeros*". We contradict this affirmation, presenting more than $2^{8\,167}$ weak states, in Section 2, after a short description of ISAAC and the observation of some minor weaknesses, in Section 1. Recall that, as a source of non-uniform randomness, weak states might distort simulations, and harm cryptographic applications, and so generators with many such states should not be used. In Section 3, we exploit some weak states to construct two strong distinguishers, requiring respectively $2^{48}$ 64-bit samples and $2^{64}$ 32-bit samples. Sections 4 and 5 respectively propose a modification of ISAAC's algorithm to avoid the design flaws presented, and point out an error in a previous analysis of ISAAC. Section 6 is our conclusion.

## 1 Preliminaries

### 1.1 Presentation of ISAAC

ISAAC is an array-based pseudo-random generator, derived from the generators IA and IBAA, presented in the same paper [3]. Although it is *"designed to be cryptographically secure"* [3], no security proof is given, and only statistical tests argue for its security. Nevertheless, only two publications tackled it until now: one [8] of 2001 by Pudovkina, presenting a state recovery attack running in time $2^{4\,121}$, and a recent one [7] by Paul and Preneel which presents a distinguisher running in time $2^{17}$. However, as we show in Section 5, the authors of the latter attack considered an algorithm slightly distinct from the real one, that makes their attack unrelevant.

We follow the description of the algorithm provided in Figure 4 of [3]; the internal state is an array of 256 32-bit words, and at each round, the algorithm computes another array of 256 32-bit words. In the following, $\alpha$ denotes the initial state, and $\alpha_i$ its $i$th element, while $\omega$ denotes the first output, and $\omega_i$ its $i$th element, for $i \in \{0, \ldots, 255\}$. The generation algorithm takes as parameters the initial values of the three variables $a$, $b$ and $c$; $a$ (32-bit) is used as an entropy accumulator, $b$ (32-bit) contains the previous pseudo-random word, and $c$ (8-bit) is a simple counter, incremented at each round of the algorithm. Their initial values are public, and are not part of the secret initial state.

We give the generation algorithm in a readable form in Algorithm 1.1, for an arbitrary round, where the variable internal state is $s$, the output array is $k$, and the inputs $a$, $b$, and $c$ are those computed in the previous round. The symbol $\oplus$ denotes the bitwise XOR, $+$ stands for the integer addition (modulo $2^k$ when needs to fit a $k$ bit value), and $\ll$ and $\gg$ are the usual shift operators. The value $f(a,i)$ in Algorithm 1.1 is a 32-bit word, defined for all $a$ and $i \in \{0,\ldots,255\}$ as:

$$f(a,i) = \begin{cases} a \ll 13 \text{ if } i \equiv 0 \bmod 4 \\ a \gg 6 \quad \text{if } i \equiv 1 \bmod 4 \\ a \ll 2 \quad \text{if } i \equiv 2 \bmod 4 \\ a \gg 16 \text{ if } i \equiv 3 \bmod 4 \end{cases}.$$

**Input:** $a$, $b$, $c$, and the internal state $s$, an array of 256 32-bit words
**Output:** an array $r$ of 256 32-bit words
1: $c \leftarrow c+1$
2: $b \leftarrow b+c$
3: **for** $i = 0,\ldots,255$ **do**
4: $\quad x \leftarrow s_i$
5: $\quad a \leftarrow f(a,i) + s_{i+128 \bmod 256}$
6: $\quad s_i \leftarrow a + b + s_{x \gg 2 \bmod 256}$
7: $\quad r_i \leftarrow x + s_{s_i \gg 10 \bmod 256}$
8: $\quad b \leftarrow r_i$
9: **end for**
10: **return** $r$

**Algorithm 1.1.** ISAAC algorithm for an arbitrary round.

For a better understanding of the following developments, we give the redundant Algorithm 1.2, which shows more clearly how the initial state $\alpha$ is used to produce the first array $\omega$,

**Input:** $a$, $b$, $c$, and the initial state $\alpha$, an array of 256 32-bit words
**Output:** an array $\omega$ of 256 32-bit words
1: $b \leftarrow b+c+1$
2: **for** $i = 0,\ldots,255$ **do**
3: $\quad s_i \leftarrow \alpha_i$
4: **end for**
5: **for** $i = 0,\ldots,255$ **do**
6: $\quad a \leftarrow f(a,i) + s_{i+128 \bmod 256}$
7: $\quad s_i \leftarrow a + b + s_{\alpha_i \gg 2 \bmod 256}$
8: $\quad \omega_i \leftarrow \alpha_i + s_{s_i \gg 10 \bmod 256}$
9: $\quad b \leftarrow \omega_i$
10: **end for**
11: **return** $\omega$

**Algorithm 1.2.** ISAAC algorithm computing the first ouput $\omega$ from the initial state $\alpha$.

### 1.2 Observations

We report here some undesirable properties of ISAAC at the origin of the weak states presented in the next section, verified experimentally with the source code provided by ISAAC's author [2]. From now, $\equiv$ symbolizes the equivalence modulo $2^{32}$.

**Fact 1.** *For a random initial state $\alpha$, and fixed $a$, $b$, and $c$, the following statements are verified.*

$$\Pr\left[\exists i \in \{1,\ldots,255\}, \omega_0 \equiv \alpha_0 + \alpha_i\right] \geq \frac{255}{256}. \tag{1}$$

$$\Pr\left[\exists i \in \{1,\ldots,255\}, \omega_0 - \omega_1 \equiv \alpha_0 - \alpha_i\right] \geq \frac{254}{256^2}. \tag{2}$$

*Proof.* (1): let $\mu = f(a,0) + \alpha_{128} + b + c + 1 + \alpha_{(\alpha_0 \gg 2) \bmod 256}$, the value obtained at line 7 of Algorithm 1.2 at the first iteration ($i = 0$). At line 8 , when $i = 0$, we get $\omega_0 = \alpha_0 + \lambda$, where $\lambda = \mu$ if $(\mu \gg 10) \bmod 256 \neq 0$, and $\lambda = \alpha_{(\mu \gg 10) \bmod 256}$ otherwise. Since $\alpha_0$ is random, $(\alpha_0 \gg 2) \bmod 256$ is a random value in $\{0,\ldots,255\}$. Since $\alpha_{128}$ is random, then $\mu$ is a random value in $\{0,\ldots,2^{32}-1\}$. Hence $\mu \gg 10 \bmod 256 \neq 0$ with probability $255/256$, which proves the result.
(2): the result is straightforward, one simply needs to apply the previous reasoning to the two following situations.

- $\omega_0 \equiv \alpha_0 + \alpha_j$ and $\omega_1 \equiv \alpha_1 + \alpha_j$, for some $j \in \{2,\ldots,255\}$.
- $\omega_0 \equiv \alpha_0 + \alpha_1$ and $\omega_1 \equiv \alpha_1 + \alpha_j$, for some $j \in \{2,255\}$.

□

**Fact 2.** *When there exists $i \in \{2,\ldots,255\}$ such that $\omega_0 = \alpha_0 + \alpha_i$, $\alpha_0$ and $i$ are correctly guessed with probability respectively $2^{-32}$ and $1/255$. Thus for a random $\alpha$, one recovers $\alpha_0$ and $\alpha_i$ for a certain $i$, with probability $2^{-32} \cdot 1/255 \cdot 255/256 = 2^{-40}$, whereas ideally this probability should be $2^{-64}$.*

**Fact 3.** *Let $N \in \{0,\ldots,127\}$, and set $\alpha_i = X$ for all $i > N$, and $\alpha_i = Y$ for all $i \leq N$, with fixed positive integers $X < 2^9$ and $Y < 2^{10}$. If $a = b = c = 0$, then*

$$\omega_0 = \begin{cases} X + 2Y + 1 \text{ if } Y \in \{0,\ldots,M\} \\ 2X + Y + 1 \text{ if } Y \in \{M+1,\ldots,2^{10}-1\} \end{cases}, \text{with } M = \max_{0<m<2^9} \{m, (m \gg 2) < N\}.$$

The above result directly follows from Algorithm 1.2; the limitation of $X$ to a 9-bit value comes from the fact that above this bound, $\alpha_i \gg 10 \neq 0$ (*cf.* line 8 of Algorithm 1.2). We also need $Y < 2^{10}$ so that, at line 7, we do not pick an index less than $N$, that is, for which $\alpha_i = Y$. For the general case, the bound $M$ comes from the fact that, at the line 7, we shall pick the value $Y$ as soon as $Y \gg 2$ is less than $N-1$, and $X$ otherwise. Finally, we need $N < 128$ in order to get $i + 128 > N \bmod 256$ for all $i \in \{0,\ldots,N-1\}$ (line 6), and so $a = X$. We obtain exactly $2^9 \cdot 2^{10} \cdot 2^7 = 2^{26}$ such states.

## 2 The weak states

Basically, the weak states considered have a fraction of random elements, and the remaining elements are fixed to the same value. We divide them into four non-disjoint sets: $W_1, W_2, W_3$ and $W_4$. This section defines each set, then presents the bias induced by its elements, and provides a few comments. We keep the notation $\alpha$ for the initial state, and $\omega$ for the first array that the algorithm outputs.

### 2.1 Set $W_1$

**Definition.** $\alpha \in W_1 \iff \alpha_0 = \alpha_1$.

**Bias.** For a random $\alpha \in W_1$,

$$\Pr[\omega_0 = \omega_1] \geq 254/256^2.$$

Indeed, for states of $W_1$, $\omega_0 = \omega_1$ holds as soon as a same element of index greater than 2 is picked at the first and second rounds (first has index $\geq 2$ with probability 254/256, then is picked again with conditional probability 1/256).

**Comments.** We will meet these states when building the distinguisher $D_1$.

There are $2^{32 \cdot 254} \cdot 2^{32} = 2^{8\,160}$ states in $W_1$.

### 2.2 Set $W_2$

**Definition.** $\alpha \in W_2 \iff \exists N \in \{2, \ldots, 256\}, \exists X \in \{0, \ldots, 2^{32} - 1\}, \alpha_0 = X, \#\{0 < i < 256, \alpha_i = X\} = N - 1$.

**Bias.** For a random $\alpha \in W_2$,

$$\Pr[\omega_0 = 2X] \geq \frac{N-1}{256}.$$

Indeed, at the first round of the algorithm, a random value $v$ of the state is picked, which is $X$ with probability $(N-1)/256$, then $\omega_0 = \alpha_0 + v$ is returned.

**Comments.** A high statistical bias appears in the distribution of the first 32 bits. For example, if $N$ is set to 6, $\Pr[\omega_0 \equiv 2X] \approx 0.02$, and there are $2^{8\,033}$ states of $W_2$ with $N = 5$.

There are more than $255 \cdot 2^{32 \cdot 254} \cdot 2^{32} \geq 2^{8\,167.99}$ states in $W_2$.

### 2.3 Set $W_3$

**Definition.** $\alpha \in W_3 \iff \exists N \in \{2, \ldots, 256\}, \exists X \in \{0, \ldots, 2^{32}-1\}, \forall i \in \{0, \ldots, N-1\}, \alpha_i = X$.

**Bias.** For a random $\alpha \in W_3$,

$$\Pr[\omega_i \equiv 2X] \geq \frac{N-1-i}{256}, i = 0, \ldots, N-1.$$

Indeed, at line 8 of Algorithm 1.2, $x = X$ holds, and so $\alpha_{\alpha_i \gg 10 \bmod 256}$ is equal to $X$ if $\alpha_i \gg 10 \bmod 256$ is greater than $i$ and strictly less than $N$, which occurs with probability greater than $(N-1-i)/256$, *cf.* Fact 1.

**Comments.** Clearly, $W_3 \subset W_2$. Again, the value $2X$ shall appear with high probability, compared to a random bitstream, but not only in $\omega_0$. For example, if $N = 64$ and $X = 0$: the last 192 elements of $\alpha$ are random, and the 64 first ones set to 0, then $\Pr[\omega_0 = \omega_1 = 0] \approx 0.06 \approx 2^{-4}$. If $N$ is as small as 2, $\Pr[\omega_0 \equiv 2X] \approx 2^{-8}$, much higher than the $2^{-32}$ of an ideal generator. If $N$ is greater than, say, 216, then $2X$ appears in average more than 90 times, thus $X$ is recovered with high probability, and the random elements remaining can be computed by exhaustive search in $2^{48}$.

There are more than $2^{32 \cdot 254} \cdot 2^{32} = 2^{8\,160}$ states in $W_2$.

### 2.4 Set $W_4$

**Definition.** $\alpha \in W_4 \iff \exists X \in \{0,\ldots,2^{32}-1\}, \forall i \in \{0,\ldots,255\}, \alpha_i = X$.

**Bias.** For a random $\alpha \in W_4$,

$$\Pr[\omega_i \equiv 2X] \geq= 1 - \frac{i+1}{256}.$$

This result comes as a particular case of $W_1$ states. Moreover, the expected number of $i$ such that $\omega_i \equiv 2X$ is greater than

$$\sum_{i=0}^{255}(1 - \frac{i+1}{256}) = 127.5,$$

that is, more than half of the elements produced at the first round are $\equiv 2X$ in average, when $\alpha_i = X$ for $i = 0,\ldots,255$.

**Comments.** It is straightforward to distinguish between a real random bitstream and a one produced by ISAAC initialised with a state with constant value, since the latter shall have about half of the $\omega_i$ equal to $2X$. The full state can even be trivially recovered in a few seconds with a paper and a pen.

There are exactly $2^{32}$ states in $W_4$.

## 3 The strong distinguishers

Briefly, a *strong distinguisher* (see Chapter 3 of [1]) is a probabilistic polynomially bounded algorithm, querying two black boxes, each one returning a bit sample of fixed length; for one box this sample is truly random, while the other's is produced by a pseudo-random generator with a random (unknown) initial state. The algorithm returns either 0 or 1 to designate the box which it "believes" to be the pseudo-random generator. An estimation of the number of samples required for a distinguisher to get a significant probability of success is given in the following theorem by Mantin and Shamir.

**Theorem 1 ([4]).** *Let $\mathcal{D}$ and $\mathcal{D}'$ be distributions, and suppose that the event $E$ happens in $\mathcal{D}$ with probability $p$ and in $\mathcal{D}'$ with probability $p(1+q)$. Then for small $p$ and $q$, $\mathcal{O}(\frac{1}{pq^2})$ samples suffice to distinguish $\mathcal{D}$ from $\mathcal{D}'$ with constant probability of success.*

We present below the distinguishers $D_1$ and $D_2$, and evaluate the number of samples required with regard to this theorem.

### 3.1 Distinguisher $D_1$

Recall that for a random state in $W_1$, $\Pr[\omega_0 = \omega_1] \approx 2^{-8}$. Thus for a random state of ISAAC,

$$\Pr[\omega_0 = \omega_1] \geq 2^{-32}(2^{-8} + 2^{-32}) + (1 - 2^{-32})2^{-32} = 2^{-32} + 2^{-40},$$

whereas this probability is $2^{-32}$ for a truly random bitstream.

Here the boxes shall output 64-bit samples at each query, and the algorithm shall select as the "ISAAC box" the one where the first 32 bits are the most frequently equal to the last 32's (that is, when $\omega_0 = \omega_1$ in ISAAC), and a random box if there is equality of occurences. Applying Theorem 1, we get $p = 2^{-32}$ and $q = 2^{-8}$, so the distinguisher requires about $2^{48}$ samples to get a significant advantage. At most 112 bits of memory are necessary (64 to read the black boxes' output, and at most 48 to count the occurences).

### 3.2 Distinguisher $D_2$

For a random state in $W_2$ with $N = 2$ ($\alpha_0 = \alpha_i$ for some $i > 0$), $\Pr[\omega_0 = 2X] \geq 2^{-8}$. Since $2X$ is even, the least significant bit of $\omega_0$ is 0 with probability $\frac{1}{2} + 2^{-8}$. Let $\zeta$ be this bit, for a random ISAAC state, we get

$$\Pr[\zeta = 0] \geq (1 - 2^{-25})\frac{1}{2} + 2^{-25}(\frac{1}{2} + 2^{-8}) = \frac{1}{2} + 2^{-33},$$

since a random state is in $W_2$ with probability $2^{-25}$. The distinguisher algorithm shall query for 32-bit samples, and choose as ISAAC box the one where the 32-th bit is the most often 0. By Theorem 1 we get $p = \frac{1}{2}$ and $q = 2^{-32}$, thus $2^{64}$ samples of 32-bits are required by the algorithm, so it runs in time about $2^{64}$, and requires at most 96 bits of memory.

## 4 ISAAC+

To fix the weaknesses presented, we modify ISAAC's algorithm, and get Algorithm 1.3. We call the corresponding pseudo-random generator ISAAC+. The modifications: we add $\oplus a$ (line 7 of Algorithm 1.3) to avoid the biases observed, perform rotations (symbols $\lll$, $\ggg$) instead of shifts, so as to get more diffusion from the state bits, and replace an addition by a XOR (line 6) to reduce the linearity over $\mathbb{Z}_{2^{32}}$.

**Input:** $a$, $b$, $c$, and the internal state $s$, an array of 256 32-bit words
**Output:** an array $r$ of 256 32-bit words
1: $c \leftarrow c + 1$
2: $b \leftarrow b + c$
3: **for** $i = 0, \ldots, 255$ **do**
4: $\quad x \leftarrow s_i$
5: $\quad a \leftarrow f'(a, i) + s_{i+128 \bmod 256}$
6: $\quad s_i \leftarrow a \oplus b + s_{x \ggg 2 \bmod 256}$
7: $\quad r_i \leftarrow x + a \oplus s_{s_i \ggg 10 \bmod 256}$
8: $\quad b \leftarrow r_i$
9: **end for**
10: **return** $r$

**Algorithm 1.3.** ISAAC+'s algorithm for an arbitrary round.

ISAAC+has the following properties.

- The properties stated in Section 1.2 do not hold: we get $\omega_0 = \alpha_0 + \alpha_i \oplus (a \ll 13 + \alpha_{128})$, for a random state, $\alpha_{128}$ is random in $\{0, \ldots, 2^{32} - 1\}$, thus so is $a \ll 13 + \alpha_{128}$. This contradicts the first proposition, and thereby the followings.
- The states presented in Section 2 lose their undesirable biases, for analog reasons. Consequently, the distinguishers $D_1$ and $D_2$ do not apply since the bias is deleted.
- ISAAC+ runs with roughly the same algorithmic complexity.
- Like ISAAC, ISAAC+ successfully passes all the Diehard [5] and NIST [6] statistical tests (this guarantees a minimal statistical quality of the pseudo-random bitstream).

## 5 Comment on a Previous Attack

At Asiacrypt'06, Paul and Preneel presented [7] distinguishers for several stream ciphers and pseudo-random generators with RC4-like construction, including ISAAC. However their analysis is based on a incorrect version of the algorithm, probably due to the hardly understandable code given in [3]: in their paper, at line 4 of Algorithm 3, the internal state updated is not the current one, but the next; they wrote "$4 : m[i+1] = \ldots$" instead of "$4 : m[i] = \ldots$". In ISAAC's code, the statement `*(m++) = <some expression>` indeed affects the current value pointed by $m$ at the expression given, *then* increments the pointer.

Based on this incorrect algorithm, the authors observe that the output at iteration $i$ comes equal to $2s_i$ with probability $\frac{1}{2}(1 + 2^{-8})$. From the bias over the parity they construct a distinguisher running in time $\approx 2^{17}$. However this does not apply to the real algorithm of ISAAC, where the value $s_i$ (denoted $m[i]$ in [7]) is updated *before* picking the output (*cf.* line 7 of Algorithm 1.2), and so the previous value of $s_i$ is not picked with the probability they considered.

## 6 Conclusion

We have shown that the bits generated by random states of ISAAC could be distinguished from true random bits in reasonable time, due to the existence of huge sets of weak states. Those results also apply to the generators IA and IBAA [3], whose algorithms are very close to ISAAC's. Although we managed to repair the problems pointed out, the new generator ISAAC+ does not offer much more security guarantees than its brother, and so should not be considered as a proposal for a new pseudo-random generator. We hope that our results will help to fill the lack of study of ISAAC.

## References

1. Oded Goldreich. *Foundations of Cryptography*, volume 1. Cambridge University Press, 2001.
2. Robert J. Jenkins. `http://www.burtleburtle.net/bob/rand/isaacafa.html`.
3. Robert J. Jenkins. ISAAC. In D. Gollmann, editor, *FSE'96*, volume 1039 of *Lecture Notes in Computer Science*, pages 41–49. Springer, 1996.
4. Itsik Mantin and Adi Shamir. A practical attack on broadcast RC4. In M. Matsui, editor, *FSE'01*, volume 2355 of *Lecture Notes in Computer Science*, pages 152–164. Springer, 2001.
5. Georges Marsaglia. *The Diehard Battery of Tests of Randomness*, 1995. Available at `http://stat.fsu.edu/pub/diehard/`.
6. National Institue of Standards and Technology. *Statistical Test Suite 1.8*, 2005. Available at `http://http://csrc.nist.gov/rng/`.
7. Souradyuti Paul and Bart Preneel. On the (in)security of stream ciphers based on arrays and modular addition. In Xuejia Lai, editor, *ASIACRYPT'06*, Lecture Notes in Computer Science, page ? Springer, 2006.
8. Marina Pudovkina. A known plaintext attack on the ISAAC keystream generator. IACR ePrint Archive, Report 2001/049, 2001. Available at `http://eprint.iacr.org/2001/049`.