T-79.512 Cryptology: Special Topics
PRP to PRF constructions II: The Sum Construction

Lauri Tarkkala

April 8, 2002

1 Introduction

A permutation is a function by definition, but it is sometimes undesirable to use a pseudo-random permutation directly as a pseudo-random function. The reason for this is that by applying the birthday paradox the expected time required to distinguish a pseudo-random permutation from a a pseudo-random function is proportional to $2^{n/2}$ instead of 2^n.

This relation ofcourse does not affect polynomial relations between primitives, but this may be a cause for concern when designing cryptographic primitives or protocols if one is interested in the exact security of a scheme.

If a pseudo-random permutation (PRP) is used directly as a pseudo-random function (PRF) then to avoid the birthday paradox, the length of the blocks input to the PRP should be twice as long in bits compared to what would be required if a PRF would be used. This obviously incurs a performance penalty that is desirable to avoid.

The paper by Bellare and Imagliazzo attempt to provide a solution for using PRP's as PRF's using a relatively efficient construction for transforming PRP's into PRF's. This survey presents the main results of that paper. [1]

2 Ratio Based Theorem

Let $\mathrm{Dist}(B,T)$ denote the statistical difference between two distributions.

Definition 2.1 If D is a probabilistic distribution on $D : R^q \rightarrow [0,1]$ then define $D(b_1,\cdots,b_i) = Pr_{(a_1,\cdots,a_q)\leftarrow D}[(a_1,\cdots,a_i) = (b_1,\cdots,b_i)]$.

This is the probability that of the variables $a_i \in R$ taking on values s.t. $a_i = b_i$ when the a_i take on values according to the distribution D.

Definition 2.2 Let $D(b_1,\cdots,b_{i-1};b_i) = Pr_{(a_1,\cdots,a_q)\leftarrow D}(a_i = b_i|(a_1,\cdots,a_{i-1}) = (b_1,\cdots,b_{i-1}))]$.

This is the conditional probability according to the distribution D that $a_i = b_i$ when $a_1,\cdots,a_{i-1}$ has taken on values s.t. $a_j = b_j$ for $1 \leq j < i$.

Let B and T be probabilistic distributions over the set R^q $(B,T : R^q \rightarrow [0,1])$.

Definition 2.3 $\mathrm{Diff}_{B,T}(a_1,\cdots,a_{i-1};\delta) = max_{a_i\in R}\{|T(a_1,\cdots,a_{i-1};a_i) - B(a_1,\cdots,a_{i-1};a_i)| - \delta B(a_1,\cdots,a_{i-1};a_i)\}$.

$\mathrm{Diff}_{B,T}(a_1,\cdots,a_{i-1};\delta)$ is the maximum difference in probability for any output value of the distributions B and T after they have output $a_1,\cdots,a_{i-1}$. This is then scaled using δ relative to the probability that B outputs this character.

Definition 2.4 $\mathrm{Dev}_{B,T;\delta} = Pr_{a_1,\cdots,a_q)\leftarrow B}(\exists i \in \{1,q\} : \mathrm{Diff}_{B,T}(a_1,\cdots,a_{i-1};\delta) > 0)$

If q elements are output using the probability distribution B then $Dev_{B,T,\delta}$ is the probability that there is a sequence of values in B and T s.t. both distributions have first output $a_1,\cdots,a_{i-1}$ in sequence and then the difference in probability that the distributions output a_i is significantly larger than the probability that B outputs this character. Here significantly is defined by the constant δ.

Lemma 2.5 *Let B,T be two continuous probability distributions over R^q.*

- *Assume* $\lambda > 0$.
- *Let* $0 \leq \delta < 1$
- *Let* $\delta_1 = \delta^2/2$. *This is the coefficient of the linear term.*
- *Let* $\delta_2 = 3\delta/2$. *This is the coefficient of the* $q^{1/2}$ *term.*
- *Let* $\epsilon = \delta_1 q + \delta_2 + q^{1/2}\lambda$.
- *Let* $\gamma = \ Dev_{B,T;\delta} + 2e^{-\lambda^2/2}$.

Then the following bound holds.
$Pr_{(a_1,\cdots,a_q)\to B}[|\ln(T(a_1,\cdots,a_q)/B(a_1,\cdots,a_q))| \leq \epsilon] \geq 1-\gamma$.

The Lemma above provides a tool for bounding the ratio of two distributions as a function of q. This Lemma is crucial in the proof of the following theorem which allows one to bound the statistical difference between to distributions as a function of the information available to the adversary (q outputs of the distribution). This lemma shows an alternative approach to analysing constructions such as the sum construction. This method does not perform the analysis by condition on any certain "bad events" such as those considered earlier in this seminar. The theorem below shows an application of this Lemma.

Theorem 2.6 *If* $0 \leq \delta < 1$ *and* $\lambda > 0$ *then*
$\text{Dist}(B,T) \leq \text{Dev}(B,T;\delta) + (e-1)/2(\delta^2 q + 3\delta q^{1/2}\lambda) + 2e^{-\lambda^2/2}$

The significant part of the theorem is the latter parts of the sum, which represents the amount of leaked information regarding the statistical difference on the distributions.

The following corollary follows directly from the above.

Corollary 2.7 *Let* $\lambda(\delta,q) = max[1,(2\ln 1/(\delta q^{1/2}))^{1/2}]$.
$\text{Dist}(B,T) \leq Dev(B,T;\delta) + 4.3\lambda(\delta,q)\delta q^{1/2}$.

The proof presented in [1] for the theorem above total approximately 11 pages and is therefore not presented here. If the reader is interested he or she should consult the paper the proof was originally published in.

The idea is to build a construction s.t. λ turns out to be a logarithmic function of q. This allows one to bound the statistical difference as a linear function of $\delta q(1/2)$. This allows one to approach the "optimal" security of a PRF with a PRP. (The term $\lambda(\delta,q)$ is considered negligible).

3 Sum Construction

The sum construction builds a PRF from a PRP. Let S denote the constructed PRF and let P denote the PRP.

Let $P : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a PRP on n-bit bit-strings using a k-bit key.

The sum-construction can be performed using only one key, by letting $S_k(x) = P_k(x||0) \oplus P_k(x||1)$. Note that the input to S is one bit shorter than the input and output of P.

The parameters in analysing the security of the construction are q, t and n. Let q be the number of queries an adversary can make to to a black-box implementing the constructed PRF. Let n be the size of the output of the constructed PRF. Let t be the amount of deterministic time available to an adversary.

Theorem 3.1 *The adversary is bounded to making at most q queries to an oracle representing the function being examined and it must run in deterministic time t.*

Let $InSec_S(q,t)$ *denote the maximum advantage an adversary can gained in distinguishing between* S *and the set of all functions of the form* $f_i : \{0,1\}^{n-1} \to \{0,1\}^n$.

Let $InSec_P(q,t)$ *denotes the maximum advantage an adversary can gained in distinguishing between* P *and the set of all permutations of the form* $p_i : \{0,1\}^n \to \{0,1\}^n$.

Let $D(n,q) = q/2^n + \mathcal{O}(n)(q/2^n)^{3/2}$.

The advantage against the sum construction is then bounded above as [1]:

$InSec_S(q,t) \leq InSec_P(2q, t + \mathcal{O}(q(n+k))) + D(n,q)$.

The theorem informally states that given a PRF created using the sum construction from a PRP P the maximum advantage an adversary could gain against the PRF can be bounded above in terms of queries and running time.

The last term decreases roughly relative to $(q/2^n)$, which means in practice this term can be made negligible.

Proof.

Let FB denote the family of random functions and let FT denote the family of of functions produced by the sum construction (assuming that a PRP is used in the construction).

The assumption is that the adversary gains access to both $P(x||0)$ and $P(x||1)$ upon performing a query. This obviously does not decrease the adversaries powers, and hence the bound is valid for $P(x||0) \oplus P(x||1)$.

The constructed function S can never output 0 (due to the output being a XOR of the outputs of the same permutation with different inputs), and hence the value 0 is not considered. The probability of guessing a 0 of any function in f_i is $q/2^n$ and this provides the term in $D(n,q)$. The 0 is henceforth removed from the range of all functions.

Let $S_i = \{f(x_j||b)|b \in \{0,1\}, j = 1, \cdots, i\}$ be a set of elements queried by the adversary.

If f is a permutation, then $|S_i| = 2i$.

Let $v \in \{0,1\}^n \setminus \{0^n\}$ be an output of S.

Let $I = |\{u, u \oplus v | u \in S_i\}|$. This is the number of pairs $u, u \oplus v \in S_i$.

Let $E = 2(q^2/2^n)$. This approximates the expected size of I assuming it was created using q queries.

Let $D = I - E$.

The number of ways the output v of S could be constructed now are $2^n/2 - 2i + I$. Substituting and computing one sees that the ratio of the probabilities of an output of v under T (family of sum constructions given a PRP) and under B (family of PRF's) is $1 + \mathcal{O}(D/N)$.

If we regard E as a frequency, a straightforward application of the Chernoff bounds provides us with the result that with probability $1-1/N$ for all i and v: $|D| = \mathcal{O}(n^2 + q(n/2^n)^{1/2})$.

If $q > (2^n)^{1/2}n^2$ let $\delta = \mathcal{O}(nq/2^{3n/2})$ otherwise let $\delta = n^2/2^{3n/2}$.

Apply the corollary from above with these values of δ and the proof is complete.

■

4 Truncation Construction

Another, albeit simpler, construction for converting from a PRF to a PRP is the so-called truncation construction, which merely discards certain bits of the output of a PRP. The result is clearly not a PRP. Let n denote the amount of input bits and m the amount of output bits ($m < n$).

Using the notation from the previous section, the advantage an adversary can gain against such a construction is bounded as[1]:

$InSec_S(q,t) \leq InSec_P(q, t + \mathcal{O}(q(n+k))) + D(n,q)$.

$D(n,q) = \mathcal{O}(n)2^m/2^n q^{1/2} if 0 < q \leq 2^m$.

$D(n,q) = \mathcal{O}(n)q/2^n(2^m)^{1/2} if 2^m < q < 2^n/(2^m)^{1/2}$.

The truncation construction provides better security if a significant amount of bits of the input of a PRP (e.g. half) are truncated, but if only a few bits are truncated then the sum construction is superior.

This is shown in the figure belowq which compares the bound of a sum construction (the function $s(x)$) for a 16-bit PRF with the truncation construction bounds for 30-bit ($t3(x)$), 26-bit ($t1(x)$) and 24-bit ($t2(x)$) when producing a 16-bit PRF.

The efficiency of the sum construction in terms of preserving security should be quite clear from the figure. Note that the truncation construction does requite the additional input bits that the sum construction does not.

References

[1] M. Bellare. A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP→ PRF conversion. FOCS 99 submission. Februrary 1999.

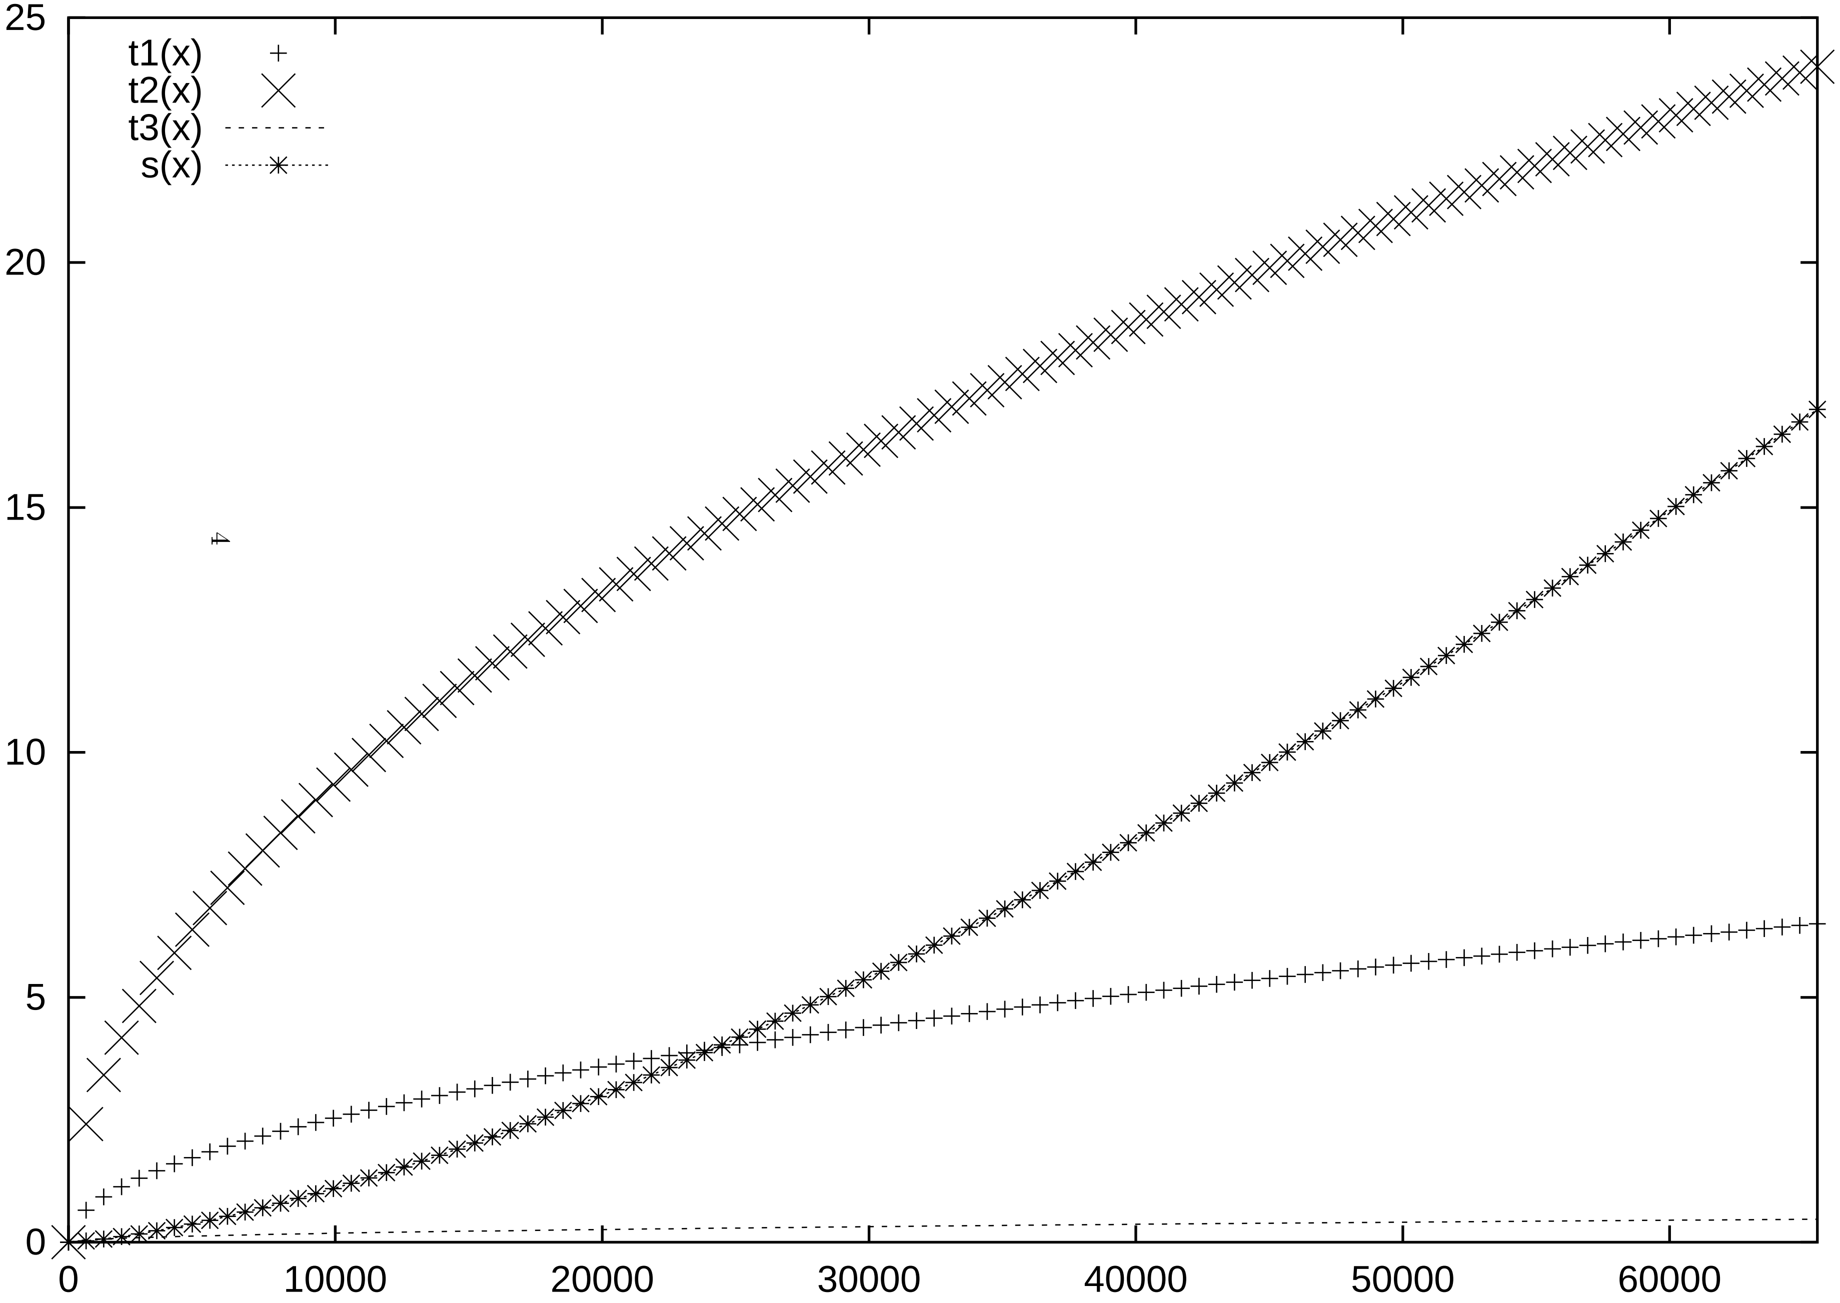
t1(x)
t2(x)
t3(x)
s(x)
0
5
10
15
20
25
0
10000
20000
30000
40000
50000
60000