

The Design of S-Boxes by Simulated Annealing

John A. Clark, Jeremy L. Jacob, Susan Stepney
Department of Computer Science, University of York
Heslington, York, YO10 5DD, UK
{jac, jeremy, susan}@cs.york.ac.uk

Received 28 March 2004

Revised manuscript received 9 July 2004

Abstract Substitution boxes (S-boxes) are important components in many modern-day symmetric key ciphers. Their study has attracted a great deal of attention over many years. The emergence of a variety of cryptosystem attacks has shown that substitutions must be designed with great care. Some general criteria such as high non-linearity and low autocorrelation have been proposed (providing some protection against attacks such as linear cryptanalysis and differential cryptanalysis). The design of appropriate S-boxes is a difficult task; several criteria must be traded off and the design space is huge. There has been little application of evolutionary search to the development of S-boxes. In this paper we show how a cost function that has found excellent single-output Boolean functions can be generalised to provide improved results for small S-boxes.

Keywords: Cryptography, S-boxes, Nonlinearity, Autocorrelation, Simulated Annealing.

§1 Introduction

Substitution plays a significant role in modern cryptography. Some substitutions are carried out by functions that take several Boolean (i.e. binary) inputs and give a single Boolean output as a result. The classic application is to combine linear feedback shift register streams to produce a binary key stream. Information flows from the input streams to the output key stream since the inputs determine the outputs. However, a key design criterion is that someone observing the output stream should not be able to practically exploit this information to compromise the system (for example, by deducing the initial states of each of the input registers which collectively form the secret key). The functions must be designed very carefully to resist attack. The design of Boolean functions




Fig. 1 Round of Data Encryption Standard

for cryptographic use has received significant attention from cryptographers for decades. Meta-heuristic search has emerged as a potentially very powerful tool for the design of such functions.^{2,11~13)} Most recently, metaheuristic search in combination with theory has found functions with properties unattained by any other means.^{3,5)}

Substitution functions may also have multiple outputs. In many cryptosystems they are termed *substitution boxes* (S-boxes for short). Perhaps the most famous (notorious) S-boxes are those of the Data Encryption Standard (DES).¹⁴⁾ Like many modern ciphers DES is an iterated block cipher; the algorithm is implemented by repeating a smaller and simpler cipher a number of times or ‘rounds’. Within each round the most significant contribution to security is made by eight 6-input, 4-output functions, shown as S1–S8 in Fig. 1. These are specified via lookup tables. The DES algorithm has been subject to a great deal of controversy. Much of this has revolved around the particular substitutions implemented by the eight S-boxes. (Another controversial aspect was the reduction of the initially suggested key-length to 56 bits.) The S-box idea has a firm hold in modern day cryptography. The new international symmetric key cryptography standard, the Advanced Encryption Standard (AES), also uses S-boxes to perform substitutions.

Unlike single-output Boolean functions, there has been little application of evolutionary or other metaheuristic search to the design of S-boxes. In this paper we show how an unusual cost function developed for evolving single-output Boolean functions can be generalised for the case of S-boxes to produce significant improvements on previous work. (Results of attempts to find highly nonlinear small S-boxes have been reported Millan et al.^{9,10)})

§2 Preliminaries

2.1 Background

This section provides some definitions of relevance to Boolean functions with cryptographic application. We denote the substitution table of an n -input k -output Boolean function by $f : \mathbf{B}^n \rightarrow \mathbf{B}^k$, mapping each combination of n Boolean (binary) input values to some combination of k Boolean output values.

For single-output functions if the number of combinations mapping to 0 is the same as the number mapping to 1 then the function is said to be *balanced*. For the multiple-output case, if each k -bit output value appears the same number of times, we say that the function is *regular*.

For the single-output case the substitution table is generally referred to as a ‘truth table’. The *polarity truth table* is a particularly useful representation for our purposes. It is defined by

$$\hat{f}(x) = (-1)^{f(x)}. \quad (1)$$

Two functions f and g are said to be uncorrelated when $\sum_{x \in \mathbf{B}^n} \hat{f}(x)\hat{g}(x) = 0$. If so, if you try to approximate f by using g , you will be right half the time and wrong half the time.

An area of particular importance for cryptanalysts is the ability to approximate a function f by a simple linear function. One of the cryptosystem designer’s tasks is to make such approximation as difficult as possible (by making the function f suitably *nonlinear*). Linearity is a form of structure cryptodesigners clearly strive to avoid. One form of attack that exploits linearity is known as linear cryptanalysis, introduced by Matsui.⁸⁾ It has attracted a great deal of attention. Another form of structure that is to be avoided is *differential* structure. Essentially, particular differences in input words (difference defined by simple bitwise XOR) may be associated with particular differences of output words (again defined by bitwise XOR) with some strong bias (i.e. the output difference is not uniform for a particular input difference). This can often be exploited by a form of attack known as differential cryptanalysis, introduced by Biham and Shamir.¹⁾ It has also attracted a great deal of attention. An excellent introduction to linear and differential cryptanalysis can be found in Heys’ tutorial.⁶⁾

Substitution boxes are essentially n -input k -output functions. These can be viewed as a combination of k individual single-output Boolean functions. Several important security criteria are actually defined in terms of single-output function criteria and so it is essential to understand first the basic Boolean function definitions and concepts. We then extend these to cater for the multiple-output case.

2.2 Cryptographic Criteria for Single-output Functions

Two important formal criteria have emerged for the single-output case to capture some aspects of resilience to the sorts of attacks indicated above. These are high nonlinearity and low autocorrelation and are defined below together with other terminology used in this paper.

A **Linear Boolean Function**, selected by $\omega \in \mathbf{B}^n$, is denoted by

$$L_\omega(x) = \omega_1 x_1 \oplus \omega_2 x_2 \oplus \cdots \oplus \omega_n x_n \quad (2)$$

where $w_i x_i$ denotes the bitwise AND of the i th bits of ω and x , and \oplus denotes bitwise XOR. The set of **affine functions** is the set of linear functions and their complements

$$A_{\omega,c}(x) = L_{\omega}(x) \oplus c \tag{3}$$

where $c \in \mathbf{B}$. The set of 2^n linear functions in polar form are an orthogonal basis for the space R^{2^n} .

For a Boolean function f the **Walsh Hadamard Transform** \hat{F}_f is defined by

$$\hat{F}_f(\omega) = \sum_{x \in \mathbf{B}^n} \hat{f}(x) \hat{L}_{\omega}(x) \tag{4}$$

where $\omega \in \mathbf{B}^n$. $\hat{F}_f(\omega)$ is the vector dot product of the polar forms of f and L_{ω} . The Walsh Hadamard values are essentially (scaled and signed) magnitudes of projections onto the various linear functions. Each value $\hat{F}_f(\omega)$ is a measure of how correlated the function f is with the relevant linear function L_{ω} , i.e. it indicates how well the linear function L_{ω} approximates f . We denote the maximum absolute value taken by any transform value by

$$WH_{max}(f) = \max_{\omega \in \mathbf{B}^n} \left| \hat{F}_f(\omega) \right|. \tag{5}$$

It is related to the nonlinearity of f .

The **nonlinearity** N_f of a Boolean function f is its minimum distance to any affine function. It is given by

$$N_f = \frac{1}{2}(2^n - WH_{max}(f)). \tag{6}$$

We can see that nonlinearity is maximised when the greatest magnitude of any Walsh Hadamard value is minimised. The less a function f can be approximated by *any* linear function, the higher its nonlinearity. There are limits to how resilient a Boolean function can be against linear approximation as demonstrated by Parseval.

Parseval's Theorem states:

$$\sum_{\omega \in \mathbf{B}^n} (\hat{F}_f(\omega))^2 = 2^{2n}. \tag{7}$$

A consequence of this result is that $WH_{max}(f) \geq 2^{n/2}$. A Boolean function \hat{f} (in polar form) is a vector of 2^n elements, each of which is +1 or -1. Parseval's Theorem is a form of Pythagoras' Theorem; the squares of the magnitudes of the projections of \hat{f} onto the orthonormal basis defined by the set of (normalised) linear functions must be the square of the magnitude of the function f . Thus, we can see how Parseval's theorem limits resilience to linear approximation. Some functions, known as Bent functions, achieve the maximum non-linearity. This occurs when $\hat{F}_f(\omega) = 2^{\frac{n}{2}}$ uniformly across all ω . In such cases f is never balanced.

The **Autocorrelation Transform** of a Boolean function f is given by

$$\hat{r}_f(s) = \sum_x \hat{f}(x) \hat{f}(x \oplus s) \tag{8}$$

where $s \in \mathbf{B}^n \setminus \{0\}$. We denote the maximum absolute value (over $s \neq 0$) in the autocorrelation spectrum of a function f by AC_f , i.e.,

$$AC_f = \max_s \left| \sum_x \hat{f}(x) \hat{f}(x \oplus s) \right|. \tag{9}$$

Here x ranges over \mathbf{B}^n and s ranges over $\mathbf{B}^n \setminus \{0\}$.

2.3 Extensions to S-boxes

For each k -output S-box, we can extract a single-output Boolean function by simply XOR-ing some subset of the output bits together. If $f(x) : \mathbf{B}^n \rightarrow \mathbf{B}^k$ is an n -input k -output S-box then each $\beta \in \mathbf{B}^k$ defines a function that is a linear combination $f_\beta(x)$ of the k outputs of f . This is given by

$$f_\beta(x) = \beta_1 f_1(x) \oplus \dots \oplus \beta_k f_k(x). \tag{10}$$

For each such function f_β the Walsh-Hadamard values $\hat{F}_{f_\beta}(\omega)$ and autocorrelation values $\hat{r}_{f_\beta}(s)$ are defined in the usual way. (Each such function f_β is now a single-output function defined over the n inputs.)

There are $2^k - 1$ non-trivial functions obtainable in this way. The notions of non-linearity and autocorrelation are readily extended to the multiple-output case. For the k -output case the non-linearity is the worst (lowest) non-linearity of all the $2^k - 1$ non-trivial single output functions f_β obtained as indicated above. Similarly, the autocorrelation is the worst (highest) over all such functions.

§3 Cost Functions

3.1 Traditional Cost Functions

Virtually all optimisation-based work aimed at producing highly nonlinear functions has used nonlinearity itself as the fitness function, i.e. the fitness of a function f on n input variables is given by

$$fitness(f) = N_f = \frac{1}{2} \left(2^n - \max_\omega |\hat{F}(\omega)| \right) \tag{11}$$

or, when viewed as a minimisation problem, the cost function is given by

$$cost(f) = WH_{max}(f) = \max_\omega |\hat{F}(\omega)| \tag{12}$$

Similarly, with low autocorrelation as the target, the autocorrelation itself has been used as the cost function, i.e. the cost function is given by

$$cost(f) = AC(f) = \max_{s \neq 0} \left| \sum_x \hat{f}(x) \hat{f}(x \oplus s) \right| = \max_{s \neq 0} |\hat{r}_f(s)|. \tag{13}$$

Previous optimisation approaches to evolving Boolean functions with desirable cryptographic properties have been generalised to the multiple-output case. Millan has compared random generation and hill-climbing as means of

evolving highly nonlinear bijective S-boxes.⁹⁾ Millan *et al.* have investigated the use of genetic algorithms and hill-climbing to evolve regular S-boxes.¹⁰⁾ Both high nonlinearity and low autocorrelation were targets. The fitness and cost measures for an S-box were the nonlinearity and autocorrelation values of that S-box. For the S-box case, the researchers above have used extensions of the basic definitions as cost functions. For non-linearity the cost function was:

$$cost(f) = \max_{\beta \in \mathbf{B}^k \setminus \{0\}, \omega \in \mathbf{B}^n} \left| \hat{F}_\beta(\omega) \right| \tag{14}$$

where \hat{F}_β is the Walsh Hadamard transform of the function f_β . For autocorrelation the cost function was:

$$cost(f) = \max_{\beta \in \mathbf{B}^k \setminus \{0^k\}, s \in \mathbf{B}^n \setminus \{0^n\}} \left| \hat{r}_\beta(s) \right| \tag{15}$$

where \hat{r}_β is the autocorrelation transform of the function f_β .

3.2 Spectrum Based Cost Functions

In 2000 a new cost function family was proposed that offered significant improvements for the single-output case.²⁾ This cost function, after significant later experimentation, was shown capable of producing functions with exceptional profiles of security criteria.⁵⁾ Rather than base the cost on extreme values (as per the definition of non-linearity and aurocorrelation), it defined a cost over the whole Walsh-Hadamard spectrum. The motivation is provided by Parseval’s equation; if the magnitudes of the Walsh Hadamard projections are very similar, then the greatest magnitude will be kept low. The details of the experiments for the single output case and the detailed motivation for the cost function adopted are given in Reference.⁴⁾ For present purposes we simply provide the cost function family below:

$$cost(f) = \sum_{\omega} \left| \left| \hat{F}_f(\omega) \right| - X \right|^R \tag{16}$$

where X and R are real-valued parameters. It is difficult to predict what the best such parameter values should be and considerable experimentation is needed. However, as indicated above, they have produced some exceptional results (effectively equalling the best results of theoreticians for functions of 8-inputs or less). A similar cost function obtained by substituting $\hat{r}_f(s)$ for $\hat{F}_f(\omega)$ was later used to similar effect.^{4,5)}

Since spectrum-based approaches generated interesting results for the single-output case an obvious question to pose is ‘Can the spectrum-based approaches be generalised to allow S-boxes to be evolved with desirable properties?’ Two cost functions can now be defined for use in S-box evolution. A cost function based on Walsh-Hadamard spectra is given by

$$cost(f) = \sum_{\beta \in \mathbf{B}^k \setminus \{0\}} \sum_{\omega \in \mathbf{B}^n} \left| \left| \hat{F}_\beta(\omega) \right| - X \right|^R \tag{17}$$

and a similar cost function based on autocorrelation spectra is given by

$$cost(f) = \sum_{\beta \in \mathbf{B}^k \setminus \{0\}} \sum_{s \in \mathbf{B}^n} \left| |\hat{r}_\beta(s)| - X \right|^R. \quad (18)$$

The single output cost functions have been applied to each function defined as a linear combination of the outputs and the results summed over all such combinations.

§4 The General Approach

All the S-boxes we are concerned with are regular: all outputs appear an equal number of times. In the case of bijective S-boxes, each output appears only once. (They are effectively permutations on the input space). In the case of n -input k -output regular S-boxes, each output appears 2^{n-k} times. A variety of move strategies are possible, but we adopt the simplest: simply swap the output values associated with two input values. The result is guaranteed to maintain regularity and bijectivity. In addition, we swap only dissimilar output values. A search starts with a regular (but otherwise random) function and moves by a sequence of such swaps around the search space.

The approach is as follows:

1. Use an annealing-based search to minimise the value of the new cost function (suitably parametrised) given in Equation 17. Let the best solution produced during the search be f_{sa} .
2. Hill-climb from f_{sa} with respect to nonlinearity (or autocorrelation), i.e. use the cost function of Equation 14 or Equation 15 to produce the final solution f_{sahc} .
3. Measure the nonlinearity, autocorrelation and algebraic degree of f_{sahc} .

Although nonlinearity, autocorrelation and algebraic degree are all of interest, our approach is somewhat unusual in that Stage 1 targets none of the criteria directly, Stage 2 considers only one of the first two, and algebraic degree is never considered at all (it is simply measured at the end). The motivation for the application of this function to the single output case is explained in detail in Reference.⁴⁾ In this paper, however, we wish to investigate its extension to the derivation of S-boxes.

We omit here a description of the annealing algorithm used for Stage 1. It is very much a ‘vanilla’ annealing. A full description is given in the Appendix.

§5 Experiments Performed

Two approaches have been used in experiments. In the first, the second-stage hill-climbing is with respect to nonlinearity. We refer to this approach as the NLT (Non-Linearity Targeted) approach. In the second, the second-stage hill-climbing is with respect to autocorrelation. We refer to this as the ACT (Auto-Correlation Targeted) approach.

5.1 Experiments and Results

Table 1 records the best nonlinearity values achieved in Millan’s experiments comparing the ability of random search and hill-climbing to evolve 5×5 , 6×6 , 7×7 and 8×8 bijective S-boxes. (For the 8×8 case Table 2 provides more detailed results of Millan’s improvement technique applied to 10000 randomly generated bijective S-boxes. The distributions of the initial and final nonlinearities obtained are shown with linear and logarithmic scales in Fig. 2.) The cost functions defined by equations 17 and 18 have been used to evolve S-boxes of similar dimensions. At the end of each run hill-climbing was carried out with respect to nonlinearity and autocorrelation respectively. 50 runs were carried out for each value of X in the set $-4, -3, -2, -1, 0, 1, 2, 3, 4$. $R = 3.0$ was used throughout. Table 1 records the best *joint* values of nonlinearity and autocorrelation achieved by either technique (i.e. functions were generated which possessed both the indicated nonlinearity value and the indicated autocorrelation value).

Table 1 Summary Results for Bijective n (autocorrelation results also shown for annealing)

n	Millan ⁹⁾		Annealing	
	rnd	HC	SA	AC-SA
5	8	10	10	16
6	20	20	22	32
7	44	46	48	48
8	98	100	102	80

Table 2 Millan’s improvement of random bijective 8×8 S-boxes, from a sample of 10000

Original Non-linearity	Final Non-linearity					
	90	92	94	96	98	100
80				2	2	
82			1	4		
84			2	29	7	
86			5	80	34	
88	1	3	20	258	127	
90		1	65	886	447	
92		2	91	1919	1112	
94			26	1946	1826	
96				246	827	3
98					17	

The results for the bijective S-boxes are not optimal. 6×6 boxes with nonlinearity of 24 have been provided by construction but they seem quite rare (Millan⁹⁾ attempted one million random generation and hill-climbing attempts and found only a nonlinearity of 20). Deriving bijective S-boxes is not an easy task for annealing. As k increases by 1 the number of derived linear combinations to check doubles. An 8×8 bijective S-box with the parameter values shown takes about 20 minutes on 1.4 GHz Pentium PC. However, again this is not easy. Only one (102, 80) function was generated from 200 runs. Similarly, for $n = 7$ only one (48, 48) function was generated. Does this matter? We address this issue below.




Fig. 2 Original and final non-linearity distributions (shown on linear and logarithmic scales) of Millan’s random bijective 8×8 S-boxes, from a sample of 10000




Fig. 3 Nonlinearity and Autocorrelation Values Achieved for $8 \times m$ S-boxes

Burnett *et al.* applied genetic algorithms followed by hill-climbing to evolve $8 \times k$ regular S-boxes (for $k = 2, \dots, 8$). The new cost functions were again used to evolve regular S-boxes of similar dimensions (with $R = 3.0$ and the same range of X as before). The best values achieved for nonlinearity and autocorrelation by each of four techniques (random, Burnett’s, SNLT and SACT) are shown in Fig. 3. We present similar data in Table 3 but here the best *joint* values of nonlinearity and autocorrelation achieved by our new approaches are shown (i.e. functions attained the indicated nonlinearity and autocorrelation values simultaneously).

Burnett *et al.* presented their results as their ‘current conjectures for the

Table 3 Nonlinearity and Autocorrelation Values Achieved for $8 \times m$ S-boxes

n	m	Burnett et al. ¹⁰⁾				Spectrum Based	
		NL		AC		Joint (NL,AC)	
		rnd	GAs	rnd	GAs	SNLT	SACT
8	2	108	110	56	48	(114,32)	(114,32)
8	3	106	108	64	56	(112,40)	(112,40)
8	4	104	106	72	64	(110,56)	(110,48)
8	5	102	104	72	72	(108,64)	(108,56)
8	6	100	104	80	80	(106,64)	(106,64)
8	7	98	102	80	80	(104,80)	(104,72)

achievable bounds'. The results of applying the annealing-based approaches with the new cost functions is fairly dramatic (the hill-climbing second stage with respect to nonlinearity or autocorrelation rarely improves matters). As m increases the same general patterns of declining nonlinearity and increasing autocorrelation are witnessed as by Burnett et al. However, the new cost functions and annealing-based searches have found functions that simultaneously improve nonlinearity and autocorrelation. Most typically, for the best functions, nonlinearity is 4 higher and autocorrelation is 16 lower.

§6 Conclusion

Comparison with theoretical approaches is difficult. On specific criteria it is clear that the derived S-boxes are not optimal. Nyberg, for example, has demonstrated 8×8 S-boxes with nonlinearity 112. For present purposes we note that spectrum-based cost functions have promise and have provided significant improvements on previous optimisation-based work.

Single-output Boolean functions have been a highly important application area for evolutionary and other metaheuristic search. Our techniques have already equalled and exceeded the combined achievements of theoreticians for the case of functions on small numbers of inputs (for $n \leq 8$). For some higher number of inputs the techniques have also produced functions unattained by other means. For larger numbers of inputs theoreticians clearly have the upper hand.

S-boxes, the multiple-output variant, are relatively unexplored and present significant challenges for the evolutionary computing community. The search spaces involved are clearly vast. For example, although an 8×8 S-box would generally be considered 'small', there are actually 256^{256} possible S-boxes of this size (each of 256 inputs can take any of 256 possible outputs). The message is very clear: S-box design gets hard very quickly. There has been considerable interest in S-boxes in the cryptography community for many years; the problem is a real-world design one. The design problems can be clearly stated and the success criteria can be easily measured. They seem ideal tests problems for evaluation of evolutionary and other nature inspired techniques.

We encourage the evolutionary computing community to attack this problem. The results reported here are targets to attack.

References

- 1) Biham, E. and Shamir, A., "Differential Cryptanalysis of DES-like Cryptosystems (Extended Abstract)," in *Advances in Cryptology – Crypto '90* (Menezes, A.J. and Vanstone, S.A. eds.), LNCS, 537, pp. 2-21, Springer-Verlag, Berlin, 1990.
- 2) Clark, J.A. and Jacob, J.L., "Two Stage Optimisation in the Design of Boolean Functions," in *5th Australasian Conference on Information Security and Privacy* (Dawson, E., Clark, A. and Boyd, C. eds.), ACISP 2000, LNCS, 1841, pp. 242-254. Springer Verlag, July 2000.
- 3) Clark, J.A., Jacob, J.L., Maitra, S. and Stănică, P., "Almost Boolean func-

- tions: the Design of Boolean Functions by Spectral Inversion,” *Computational Intelligence*, 20, 3, pp. 450-462, August 2004.
- 4) Clark, J.A., Jacob, J.L. and Stepney, S., “Searching for Cost Functions,” in *Conference on Evolutionary Computation – CEC-04*, July 2004.
 - 5) Clark, J.A., Jacob, J.L., Stepney, S., Maitra, S. and Millan, W., “Evolving Boolean Functions Satisfying Multiple Criteria,” in *Progress in Cryptology – INDOCRYPT 2002, LNCS, 2551*, pp. 246-259, Springer Verlag, 2002.
 - 6) Heys, H.M., “A Tutorial on Linear and Differential Cryptanalysis,” *Technical report, Electrical and Computer Engineering*, University of Newfoundland, St. John’s, Newfoundland, Canada, 2004.
 - 7) Kirkpatrick, S., Gelatt, Jr. C.D. and Vecchi, M.P., “Optimization by Simulated Annealing,” *Science*, 220, 4598, pp. 671-680, May 1983.
 - 8) Matsui, M., “Linear Cryptanalysis Method for DES Cipher,” in *Advances in Cryptology – EuroCrypt ’93* (Hellesest, T. ed), *LNCS, 765*, pp. 386-397, Springer-Verlag, Berlin, 1993.
 - 9) Millan, W., “How to Improve the Non-linearity of Bijective S-boxes,” in *3rd Australian Conference on Information Security and Privacy* (Boyd, C. and Dawson, E. eds.), *LNCS, 1438*, pp. 181-192, Springer-Verlag, April 1998.
 - 10) Millan, W., Burnett, L., Carter, G., Clark, A. and Dawson, E., “Evolutionary Heuristics for Finding Cryptographically Strong S-Boxes,” in *ICICS 99, LNCS, 1726*, pp. 263-274, Springer-Verlag, 1999.
 - 11) Millan, W., Clark, A. and Dawson, E., “Smart Hill Climbing Finds Better Boolean Functions,” in *Workshop on Selected Areas on Cryptography, SAC 97*, pp. 50-63, 1997.
 - 12) Millan, W., Clark, A. and Dawson, E., “Boolean Function Design Using Hill Climbing Methods,” in *4th Australian Conference on Information Security and Privacy* (Schneier, B. ed), *LNCS, 1587*, pp. 1-11, Springer-Verlag, April 1999.
 - 13) Millan, W., Clark, A. and Dawson, E., “Heuristic Design of Cryptographically Strong Balanced Boolean Functions,” in *Advances in Cryptology EUROCRYPT’98, LNCS, 1403*, pp. 489-499, Springer-Verlag, 1998.
 - 14) National Bureau of Standards, Data Encryption Standard, *NBS FIPS PUB 46*, 1976.

Appendix

§A1 Description of Simulated Annealing

In 1983 Kirkpatrick et al.⁷⁾ proposed *simulated annealing*, a new search technique inspired by the cooling processes of molten metals. It merges hill-climbing with the probabilistic acceptance of non-improving moves. The basic algorithm is shown in Fig. A1. The search starts at some initial state $S := S_0$. There is a control parameter T known as the temperature. This starts ‘high’ at T_0 and is gradually lowered. At each temperature, a number *MIL* (Moves in Inner Loop) of moves to new states are attempted. A candidate state Y is randomly selected from the neighborhood $N(S)$ of the current state. The change

in value, δ , of f is calculated. If it improves the value of $f(S)$ (i.e., if $\delta < 0$ for a minimisation problem) then a move to that state is taken ($S = Y$); if not, then it is taken with some probability. The worse a move is, the less likely it is to be accepted. The lower the temperature T , the less likely is a worsening move to be accepted. Probabilistic acceptance is determined by generating a random value U in the range $(0 \dots 1)$ and performing the indicated comparison. Initially the temperature is high and virtually any move is accepted.

As the temperature is lowered it becomes ever more difficult to accept worsening moves. Eventually, only improving moves are allowed and the process becomes 'frozen'. The algorithm terminates when the stopping criterion is met. Common stopping criteria, and the ones used for the work in this paper, are to stop the search after a fixed number *MaxIL* of inner loops have been executed, or else when some maximum number *MUL* of consecutive unproductive inner loops have been executed (i.e., without a single move having been accepted). Generally the best state achieved so far is also recorded (since the search may actually move out of it and subsequently be unable to find a state of similar quality). At the end of each inner loop the temperature is lowered. The simplest way of lowering the temperature is to multiply by a constant cooling factor α in the range $(0 \dots 1)$; this is known as *geometric cooling*. The basic simulated annealing algorithm has proven remarkably effective over a range of problems.

```

S := S0
T := T0
repeat
{
  for (int i = 0; i < MIL; i++)
  {
    select Y ∈ N(S)
    δ := f(Y) - f(S)
    if (δ < 0) then
      S := Y
    else
      generate U := rnd(0, 1)
      if (U < exp(-δ/T)) then S := Y
  }
  T = T × α
}
until stopping criterion is met

```

Fig. A1 Basic Simulated Annealing for Minimization Problems



John A. Clark: He is Professor of Critical Systems at the University of York, where he leads the software testing, security and cryptography work. Much of this has been concerned with the application of meta-heuristic search.



Jeremy L. Jacob: He has a BSc. in Mathematics from the University of Hull, England, M.Sc. and D.Phil. in Computation from the University of Oxford, England and now works for the University of York. His research interests include modelling secure systems and software engineering practices for secure systems.



Susan Stepney: She is Professor of Computer Science at the University of York, and leads the Non-Standard Computation research group there. She is a member of the ACM, Fellow of the British Computer Society, and moderator of the UKCRC Grand Challenge in Non-Classical Computation. Her main research interests include novel applications of nature-inspired computation, modelling self-organising complex systems and designing and reasoning about emergent properties.