

SETH

A VLSI CHIP FOR THE REAL-TIME INFORMATION DISPERSAL AND RETRIEVAL FOR SECURITY AND FAULT-TOLERANCE

AZER BESTAVROS*

Department of Computer Science
Harvard University
Cambridge, MA 02138.

Abstract: In this paper, SETH,^(a) a hardwired implementation of the recently proposed “Information Dispersal Algorithm” (IDA) [7], is presented. SETH allows the real-time dispersal of information into pieces as well as the retrieval of the original information from the minimal intact subset of these pieces. We begin the paper by introducing IDA and overviewing SETH operation. Next, the different algorithms used are described and schematics of varying levels of details are presented. Next, we present an implementation of SETH [1]-[2] using Scalable CMOS technology that has been fabricated using a MOSIS 3-micon process. We conclude the paper with potential applications and extensions.

Introduction

The storage and transmission of data in computer systems raises significant security and reliability problems. In particular, data might be lost due to hardware failures, it might be accidentally (or even maliciously) garbled or destroyed, and it might be read and interpreted by unauthorized users.

The common solution to the aforementioned security problems is to have users store and communicate their data using some form of encryption, where only authorized users are enabled to decrypt the information through the use of appropriate *Secret Keys* [8]. The proven difficulty of decrypting the information without knowing the secret key guarantees a high level of *security*. On the other hand, to protect against possible failures, redundancy is often used as the only alternative to achieve fault-tolerance. This is usually done by having the users replicate their (possibly encrypted) data into n different machines [5]. The independency of the failure modes of these machines guarantees a high level of *availability*. The main disadvantages of this *encryption and replication* strategy is that it results in an n -fold blowup of the total storage required in the system. Moreover, the existence of *all* the information (possibly encrypted) in one site^(b) for long periods might

make it possible for adversaries to break the secret key.

Recently, Michael Rabin [7] proposed the “Information Dispersal Algorithm (IDA)” as a potential technique that would achieve the required security and reliability using a much smaller level of redundancy and by keeping only *partial* information at a specific site. The main idea is to use a secret key to disperse the information of a file F into n pieces which are transmitted and stored on n different machines (or disks) such that the contents of the original file, F , can be reconstructed from any m of its pieces, where $m \leq n$. On the one hand, the proposed technique guarantees the confidentiality of the dispersed information. As a matter of fact, it is hard to get any clue about the original information unless at least m pieces from the dispersed file are collected. This makes the task of the adversaries more difficult, since they have to control m of the sites and not only one. Even if this happens, it is provably very difficult to reconstruct the original file unless the secret keys are known. On the other hand, the proposed technique guarantees a higher availability, since it tolerates up to $(n - m)$ failures. The salient feature of the “Information Dispersal Algorithm” is that each of the dispersed pieces is of size $\frac{|F|}{m}$, where $|F|$ is the size of the original file. Hence, the added redundancy is $(\frac{n}{m} - 1) * 100\%$.

In this paper, we present a hardwired implementation of the “Information Dispersal Algorithm” that would allow the execution of the algorithm in real-time. A chip, which we called SETH^(c), has been designed in the VLSI Lab of Harvard University using Scalable CMOS technology and has been fabricated by the MOSIS 3-micron process. SETH accepts a stream of data and a set of vector-keys (see the description below) along with the necessary controls so as to produce the streams of encrypted data to be stored on (or communicated with) the different sinks.

^(c)According to an old Egyptian legend, SETH (an Egyptian prince) killed his brother OSIRIS and cut his body into small pieces and “dispersed” it all over the Eastern Mediterranean. Later, his loving wife, ISIS, collected all the pieces together and “reconstructed” OSIRIS, bringing him to life again as the Evil God of the Underworld.

*This work was supported by DARPA N00039-88-C-0163

^(a)©1988 Harvard University. All rights reserved.

^(b)whether stored in or communicated through that site

The chip might as well accept the streams of data from the different machines along with the necessary controls and vector-keys so as to reconstruct the original information.

We begin this paper by introducing the Information dispersal algorithm and presenting an overview of SETH design. Next, we describe how the different operations are done using “irreducible polynomial arithmetic.” In particular, we outline general methods for efficient addition and multiplication in these systems. Next, functional units are described and system block diagrams of varying levels of details are presented. We conclude the paper by examining potential applications and extensions for SETH. In particular, we are considering the use of IDA in the design of I/O subsystems, Redundant Arrays of Inexpensive Disks (RAID) systems, reliable communication, and routing in distributed/parallel systems. SETH demonstrates that using IDA in these applications is feasible.

SETH was designed using MAGIC. The design was tested using ESIM and SPICE. The MAGIC layout of the chip as well as the simulation results are available in the VLSI Lab, Harvard University [1]. The chip was fabricated by MOSIS on a 3-micron SCMOS-technology process. Twelve packages were returned and the functionality of the chip was verified.

Overview of SETH

SETH is an interface that allows the reliable and secure storage and communication of information between the different units of a computing system. Reliability is guaranteed by using redundant communication and/or storage, while security is achieved using encryption. A major objective in SETH design was *versatility* and *flexibility*. We view SETH as a *basic building block*. By using many of these blocks in different configurations, one can achieve different design objectives (namely different levels of fault-tolerance, recoverability, redundancy and security).

SETH can be operated in two modes; *Disperse-mode* or *Reconstruct-mode*. In disperse-mode, the data given to SETH is encrypted, dispersed and sent to the different sinks. In reconstruct-mode, the data from a number of intact sinks is collected and recombined to yield the original information. The encrypt/decrypt and disperse/recombine functions are achieved using a set of *secret keys* that must be supplied to SETH. Figure 1, shows the SETH disperse/reconstruct modes.

Keeping the aforementioned versatility and flexibility in mind, we decided to make the SETH chip interface an 8-bit bus to four 4-bit buses – using four 8-bit keys. Hence, a stream of 8 bits can be dispersed into four streams of 4 bits each, so that any two of these four streams are sufficient to reconstruct the original 8-bit stream. Different series/parallel configurations using SETH are possible. For instance, two SETH chips in parallel can be used to interface a 16-bit bus to eight 4-bit buses or to four 8-bit buses or to two 16-bit buses. In any case the level of redundancy is 100% and the achievable fault-tolerance is 50% (that is we tolerate the loss of 50% of the storage sinks or chan-

Figure 1: SETH Disperse and Reconstruct modes

nels). These, however, are still controllable. For instance by using only three of the four output buses, the redundancy drops to 50%, whereas the level of fault-tolerance becomes 33%. In any of the above cases security is guaranteed since the dispersed data is actually encrypted. Moreover, by using exactly two of the output buses, SETH can be used just for encryption and dispersal (with no added redundancy and no support for failures). Another degree of freedom (both in the redundancy/fault-tolerance and security levels) can be achieved by using series configuration of SETH.

Information Dispersal and Retrieval

In this section we explain how the IDA works. We single out the different operations to be performed and which must be carried out in real-time by SETH. For a thorough presentation of the algorithm, we refer the reader to the original paper on IDA, [7], and to the data scattering and gathering examples in the Appendix of [2].

Let F be the original file (information) we want to disperse. The main idea behind IDA is to split this file into n different pieces so that recombining any m of these, $m \leq n$, is sufficient to reconstruct F , whereas any number of pieces less than m would not be sufficient to reveal any information about the contents of any portion of F .

The original file F can be viewed as a sequence (or stream) of data on the form $F = b_1b_2b_3b_4...etc.$, where each b_i in this stream can be viewed as an integer. In order to disperse this stream, we choose a set of n vectors, *secret keys* (V_1, V_2, \dots, V_n) each of length m . These keys have to meet certain (easily satisfiable) linear independence conditions (see [7] and [2] for details.) Let A_{nm} be the array

whose rows are the selected vectors. A represents a mapping from an m -dimensional space to an n -dimensional space, or in other words, from a sequence of m elements to another sequence of n elements. To disperse the file F , we simply map each sequence of m elements from F into a new sequence of n elements using the transformation A . Each element from the resulting sequence is sent to a different site and kept there. So, for each m elements of F we send *one* element to each of the n sites. Therefore, to disperse the whole file F we will need to send $\frac{|F|}{m}$ elements to each of the n sites.

Now, suppose that we want to reconstruct the original file F from the pieces dispersed as described above. This is done by reading any m of these pieces.^(d) Let the pieces be from sites $s_1, s_2, s_3, \dots, s_m$. Let B_{mm} be the array whose rows are $(V_{s_1}, V_{s_2}, V_{s_3}, \dots, V_{s_m})$. Thus, B maps sequences of m elements from F into sequences of m elements, which by virtue of the dispersion step above, are kept at sites $s_1, s_2, s_3, \dots, s_m$. To reconstruct the first m elements of F we need simply to collect the first element from each of m different sites and use the appropriate inverse transformation ($T = B^{-1}$). Note that if the keys were appropriately chosen, such an inverse is guaranteed to exist.

In SETH we decided to pick $n = 4$ and $m = 2$, so that F is dispersed into 4 different pieces and using no less than 2 of these pieces would be sufficient to reconstruct F . Moreover, we decided to represent F as a stream of hexadecimal digits, *nibbles*, (integers in the range [0..15]). Hence, each *byte* (8 bits) of F can be viewed as 2 nibbles and we use the IDA described above to produce 4 nibbles that are dispersed to the 4 different sites. This is illustrated in Figure 1. It is to be noted, however, that the design techniques that we present in this paper are independent of the choices above and can be easily applied to any other choices.

The Disperse operation

The *Disperse* is simply a 4×2 transformation A ,

$$A = \begin{pmatrix} V_0^T \\ V_1^T \\ V_2^T \\ V_3^T \end{pmatrix} = \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \\ a_{20} & a_{21} \\ a_{30} & a_{31} \end{pmatrix}$$

where, V_i is the i^{th} key-vector and each a_{ij} is a hexadecimal (4-bit) number. Let $(b_0b_1)^T$ represent the two hexadecimal digits (a total of 8 bits) from the input stream. To disperse this piece of information, we use the transformation A as follows:

$$\begin{pmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \\ a_{20} & a_{21} \\ a_{30} & a_{31} \end{pmatrix} \times \begin{pmatrix} b_0 \\ b_1 \end{pmatrix}$$

^(d)If less than m pieces are available then the file cannot be reconstructed, and if more than m pieces are available then the use of any subset of m pieces will suffice.

where, c_i is the hexadecimal digit sent to the i^{th} site.

The Reconstruct operation

Given that the dispersed data at sites s_1 and s_2 is intact, the *Reconstruct* transformation is simply a 2×2 matrix T ,

$$T = \begin{pmatrix} V_{s_1}^T \\ V_{s_2}^T \end{pmatrix}^{-1} = \begin{pmatrix} t_{00} & t_{01} \\ t_{10} & t_{11} \end{pmatrix}$$

Let $(c_0c_1)^T$ represent the two hexadecimal digits (a total of 8 bits) from the two intact sites. To reconstruct the original byte of information, we use the transformation T as follows:

$$\begin{pmatrix} b_0 \\ b_1 \end{pmatrix} = \begin{pmatrix} t_{00} & t_{01} \\ t_{10} & t_{11} \end{pmatrix} \times \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}$$

Irreducible Polynomial Arithmetic

All operations (namely additions and multiplications) needed for the IDA have to be carried-out using the *irreducible polynomial arithmetic* (IPA) where integers are viewed as polynomials over some finite field Z_p , p is a prime number (see the Appendix of [2] for a more detailed discussion and for a complete example). Taking $p = 2$ (an obvious choice for digital applications), integers are represented as polynomials with binary coefficients. For instance, the hexadecimal digits from 0 to 15 are represented in IPA as follows:-

$$\begin{pmatrix} 0000 \\ 0001 \\ 0010 \\ 0011 \\ 0100 \\ 0101 \\ 0110 \\ 0111 \\ 1000 \\ 1001 \\ 1010 \\ 1011 \\ 1100 \\ 1101 \\ 1110 \\ 1111 \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ x \\ x+1 \\ x^2 \\ x^2+1 \\ x^2+x \\ x^2+x+1 \\ x^3 \\ x^3+1 \\ x^3+x \\ x^3+x+1 \\ x^3+x^2 \\ x^3+x^2+1 \\ x^3+x^2+x \\ x^3+x^2+x+1 \end{pmatrix}$$

In an IPA with 16 elements, all operations are done modulo an irreducible 4th degree polynomial (one that cannot be divided by any polynomial of 3rd or lesser degree.) For instance, it can be easily shown that (x^4+x+1) is an irreducible 4th degree polynomial. Indeed, this is the one used in SETH.

Addition: Addition is straightforward. To add two integers, we add their corresponding polynomials. This is done by adding (modulo-2) the coefficients of the corresponding powers. The following is an example of addition done in IPA:

$$5 + 6 = (x^2 + 1) + (x^2 + x) = (x + 1) = 3$$

It is obvious that addition is just the bitwise “exclusive-or” of the binary representation of the integers. This makes the hardware implementation straightforward.

Multiplication: Multiplication is a little bit more complicated. To multiply two integers, we multiply their corresponding polynomials. If the resulting polynomial is of degree less than the order of the irreducible polynomial (4 in our case) then we got the polynomial representation of the result. Otherwise, we have to get the residue of the result when divided by the irreducible polynomial. Again, all additions and multiplications are done (modulo-2). The following is an example of a multiplication done using $(x^4 + x + 1)$ as the irreducible polynomial:

$$\begin{aligned} 3 * 10 &= (x + 1)(x^3 + x) \\ &= (x^4 + x^3 + x^2 + x) \text{mod}(x^4 + x + 1) \\ &= (x^3 + x^2 + 1) \\ &= 13 \end{aligned}$$

Implementing multiplication in IPA is no big deal!!.. The most straightforward implementation would be using table lookup. A more efficient, elegant and scalable implementation, however, can be done using “shifts” and selective “exclusive-or’s”. We discuss this in the following section.

A Disperse/Reconstruct example

Let $F = 2, 4, 1, 14, 6, 8, \dots \text{etc.}$ be the stream to be dispersed, and assume that we selected the secret key-vectors to be:

$$\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \begin{pmatrix} 1 \\ 3 \end{pmatrix}$$

Hence, the transformation A is:

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix}$$

To disperse F we divide it into sequences of 2 elements as shown below:

$$F = \underbrace{2, 4}, \underbrace{1, 14}, \underbrace{6, 8}, \dots \text{etc.}$$

Next, each of these sequences is transformed using A and we obtain the new sequences shown below:

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} \times \begin{pmatrix} 2 \\ 4 \end{pmatrix} = \begin{pmatrix} 2 \\ 6 \\ 10 \\ 14 \end{pmatrix}$$

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} \times \begin{pmatrix} 1 \\ 14 \end{pmatrix} = \begin{pmatrix} 1 \\ 15 \\ 14 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix} \times \begin{pmatrix} 6 \\ 8 \end{pmatrix} = \begin{pmatrix} 6 \\ 14 \\ 5 \\ 13 \end{pmatrix}$$

... etc. ...

From the above, the resulting sequence will be:

$$F' = \underbrace{2, 6, 10, 14}, \underbrace{1, 15, 14, 0}, \underbrace{6, 14, 5, 13}, \dots \text{etc.}$$

The first sink will be given the 1st element from each resulting sequence, i.e. (2, 1, 6, ...etc.) Similarly, the second sink will be given the 2nd element from each resulting sequence, i.e. (6, 15, 14, ...etc.), ...etc.

$$\begin{aligned} F'_1 &= 2, 1, 6, \dots \\ F'_2 &= 6, 15, 14, \dots \\ F'_3 &= 10, 14, 5, \dots \\ F'_4 &= 14, 9, 13, \dots \end{aligned}$$

Now, suppose that the second and fourth site fail and we want to reconstruct the original file F . This is done by first computing the transformation B which consists of the key-vectors for the available sites (namely the first and third) as shown below:

$$B = \begin{pmatrix} 1 & 0 \\ 1 & 2 \end{pmatrix}$$

Second, we compute the inverse of this transformation $T = B^{-1}$ as follows:^(e)

$$T = B^{-1} = \begin{pmatrix} 1 & 0 \\ 9 & 9 \end{pmatrix}$$

Now, to reconstruct the first sequence of m elements of original file, we transform the sequence consisting of the first element from the first and third sites, namely (2, 6), using T as follows:

$$\begin{pmatrix} 1 & 0 \\ 9 & 9 \end{pmatrix} \times \begin{pmatrix} 2 \\ 6 \end{pmatrix} = \begin{pmatrix} 2 \\ 4 \end{pmatrix}$$

Thus obtaining the first two elements (2, 4) of the original file.

Similarly, to get the next two elements of F , we transform the sequence consisting of the second element from the first and third sites, namely (6, 14), using T as follows:

$$\begin{pmatrix} 1 & 0 \\ 9 & 9 \end{pmatrix} \times \begin{pmatrix} 6 \\ 14 \end{pmatrix} = \begin{pmatrix} 1 \\ 14 \end{pmatrix}$$

Thus obtaining the third and fourth elements of F , namely (1, 14). The process is basically the same for the rest of the file.

Notice that the reconstruction needs not be done from the same sinks. The set of intact sinks can be changing dynamically, provided that the appropriate keys are supplied.

^(e)This inverse is with respect to the irreducible polynomial 10011

SETH functional units

Given the 2 hexadecimal digits from the input stream b_0, b_1 , and given the 8 hexadecimal digits which define the transformation $A = [a_{ij}], 0 \leq i, j \leq 3$ described above, SETH should be able to compute 4 hexadecimal digits c_0, c_1, c_2, c_3 using the relation:

$$\begin{pmatrix} c_0 \\ c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \\ a_{20} & a_{21} \\ a_{30} & a_{31} \end{pmatrix} \times \begin{pmatrix} b_0 \\ b_1 \end{pmatrix}$$

Moreover, using the appropriate inverted transformation $T = [t_{ij}], 0 \leq i, j \leq 1$, the data c_0, c_1 returned from any two different sinks can be recombined using the relation:

$$\begin{pmatrix} b_0 \\ b_1 \end{pmatrix} = \begin{pmatrix} t_{00} & t_{01} \\ t_{10} & t_{11} \end{pmatrix} \times \begin{pmatrix} c_0 \\ c_1 \end{pmatrix}$$

In the above matrix operations, all the additions and multiplications are to be done using IPA. It is obvious that the main functional units in SETH are the adder and multiplier (which themselves might be realized using other basic building blocks.) By using an array of eight multipliers and by adding the appropriate results together using 4 adders, a data path that implements matrix multiplication can be realized. To be able to use the same data path for both the disperse and reconstruct modes of SETH, some control logic needs to be added.

The Adder

Let $X = x_3x_2x_1x_0$ and $Y = y_3y_2y_1y_0$ be the binary representation of the two hexadecimal digits X, Y to be added. As we have stated earlier, addition using IPA in the binary field Z_2 , reduces to the bitwise “exclusive-or.” Thus, the logical equations for the result $Z = z_3z_2z_1z_0$ is given by:

$$\begin{aligned} z_3 &= x_3 \oplus y_3 \\ z_2 &= x_2 \oplus y_2 \\ z_1 &= x_1 \oplus y_1 \\ z_0 &= x_0 \oplus y_0 \end{aligned}$$

The Multiplier

Let $X = x_3x_2x_1x_0$ and $Y = y_3y_2y_1y_0$ be the binary representation of the two hexadecimal digits to be multiplied, where the result is $Z = z_3z_2z_1z_0$.

First, we notice that multiplication of binary numbers actually reduces to successive shifts and adds; in our case exactly four shift/add stages are required (one stage for each bit of Y). In each of these stages, the accumulated value $W = w_3w_2w_1w_0$, initially being 0000, is shifted left one bit and if the corresponding bit of Y is 1, X is added to the accumulated value and the result is propagated to the next “lower” stage. Of course all additions have to be done using the technique described above. Second, we notice that the resulting number (using the described four-stage shift/add) cannot be used directly since the polynomial representation of this result might now be

of fourth (or higher) degree and the residue of this result should be computed using $x^4 + x + 1 (\equiv 10011)$. To compute this residue means that we need to perform successive subtractions. Fortunately, subtraction (modulo-2) is just the same as addition (modulo-2). Moreover, we note that these subtractions can be actually done within each stage of the above shift/add steps (and hence need not be accumulated till the end of the multiplication. Finally, we notice that at most one subtraction is needed within each stage of the multiplier since in each such stage, the degree of the accumulated result cannot be increased by more than one, and consequently, by subtracting the irreducible polynomial 10011 just once (whenever an overflow is detected) guarantees that the accumulated result will remain of the third degree (or less).

An algorithm to implement the above technique is shown in Figure 2. In this algorithm, each iteration corresponds to one stage of the multiplier. Several optimizations can be applied to the algorithm. First, we notice that the test for whether y_i is equal to 0 or not and doing the shift-only or shift-and-add operation accordingly can be replaced by a shift and add to the bitwise product (“and”) of y_i and X . Second, we notice that the test for whether a subtraction is needed or not can be replaced by always doing an “exclusive-or” of the result with w_4 (since if $w_4 = 0$ the “exclusive-or” won’t change anything). Finally, we notice that $a \oplus 0 = a$. The optimized version of the algorithm is given in Figure 3. In Figure 4 we illustrate the use of this optimized and systematic method in computing the product of 1011 by 1100.

The above algorithm is the basis for our multiplier design. As a matter of fact the whole multiplier is constructed using four identical stages stacked together. Each stage accepts the accumulated value computed from the previous stage (stage #3 being fed with 0000). Also each stage accepts the value of X and one of the bits of Y . The result of the multiplication is the accumulated result from stage #0. The design of each of the multiplier stages requires exactly four “and” gates and five “exclusive-or” gates. The multiplier unit and the multiplication stage are shown in Figure 5.

The matrix multiplication unit

The matrix multiplication unit can be simply built using 8 multipliers to multiply in parallel the 2-element input vector with the corresponding elements of the transformation matrix, yielding 8 different products. Each couple of these products is added in parallel using a separate adder to produce the required 4-element output. Figure 6 shows the matrix multiplication unit.

Data flow Control

SETH is a bidirectional interface between two data streams. These streams are fed to SETH using two bidirectional ports B & C. Both the disperse and reconstruct modes of SETH involve matrix multiplication. This suggests that the same matrix multiplication unit could be

```

Begin
 $w_4 = 0$  ;           Initialize result
 $w_3 = 0$  ;
 $w_2 = 0$  ;
 $w_1 = 0$  ;
 $w_0 = 0$  ;
for ( $i = 3, i \geq 0, i--$ ) {   For each stage
  if ( $y_i = 0$ ) {             Is  $y_i = 0$  ?
     $w_4 = w_3$  ;             If Yes then just shift left
     $w_3 = w_2$  ;
     $w_2 = w_1$  ;
     $w_1 = w_0$  ;
     $w_0 = 0$  ;
  }
  else {
     $w_4 = w_3 \oplus 0$  ;     If No then shift left and add
     $w_3 = w_2 \oplus x_3$  ;
     $w_3 = w_1 \oplus x_2$  ;
     $w_3 = w_0 \oplus x_1$  ;
     $w_3 = 0 \oplus x_0$  ;
  }
  if ( $w_4 = 1$ ) {           Is there a need to subtract ?
     $w_4 = w_3 \oplus 1$  ;     If Yes then do it !!
     $w_3 = w_3 \oplus 0$  ;
     $w_2 = w_2 \oplus 0$  ;
     $w_1 = w_1 \oplus 1$  ;
     $w_0 = w_0 \oplus 1$  ;
  }
}
End

```

Figure 2: IPA multiplication

used for both modes with the appropriate control to forward the data in and out from the matrix multiplication unit.

In disperse-mode (dispersal operation), SETH accepts 8-bit inputs (2 hexadecimal digits) from its B-port. These inputs, as well as the 32-bit (8 hexadecimal digits) transformation matrix are forwarded to the matrix multiplication unit to produce the 16-bit (4 hexadecimal digits) output which is forwarded to SETH's C-port. In reconstruct-mode, SETH accepts 16-bit inputs (4 hexadecimal digits) from its C-port. Depending on the control lines applied, only 2 hexadecimal digits from these 4 are forwarded to the matrix multiplication unit along with the appropriate inverted secret keys provided to the chip. Finally 8 of the 16-bit output produced are routed to SETH's B-port (since only half of the matrix multiplication unit is used in this case).

The SETH chip

Figure 7 shows the simulated delays (worst and best and average), the power consumption, the number of devices used (FETs) and the area of the layout for the different units of SETH. This data does not reflect the delays and power consumed by i/o, Vdd, and Gnd pads. Twelve SETH chips were tested in the VLSI Lab. at Harvard University, of which, three proved to be defective (a yield

```

Begin
 $w_4 = 0$  ;           Initialize result
 $w_3 = 0$  ;
 $w_2 = 0$  ;
 $w_1 = 0$  ;
 $w_0 = 0$  ;
for ( $i = 3, i \geq 0, i--$ ) {   For each stage
   $w_4 = w_3$  ;             Add  $X$  if  $y_i = 1$ , otherwise add 0
   $w_3 = w_2 \oplus (x_3.y_i)$  ;
   $w_3 = w_1 \oplus (x_2.y_i)$  ;
   $w_3 = w_0 \oplus (x_1.y_i)$  ;
   $w_3 = (x_0.y_i)$  ;
   $w_1 = w_1 \oplus w_4$  ;     Adjust result (if necessary)
   $w_0 = w_0 \oplus w_4$  ;
}
End

```

Figure 3: Optimized IPA multiplication

of 75%.) The maximum speed for the correct operation was found to be approximately 4 Mhz. With a more elaborate design (for example, using pipelining) and a more advanced fabrication technology, we believe that this figure can be improved by at least one order of magnitude.

Applications of SETH

There are basically two areas where SETH may be used; Data storage and retrieval, and Data communication. In a storage system, SETH would be located between the storage device and the data bus. In the case of a communication system, SETH would be placed at each end of a SETH bus with the system bus at each of the remaining ends. A beneficial side effect of using IDA is improved load balancing in both storage and communication.

Data Storage and Retrieval using SETH

The integrity of stored information could be improved greatly by using SETH. Mechanical storage devices are the weakest components in a computer system due to their intolerance to shocks, vibration, dust, and their inherently unreliable moving parts.

Two SETH chips configured in parallel, and three disks will make a fault tolerant system with two times the storage of a single disk. In addition the system is secure from information thefts. During normal operation, data could be read from any two of the three drives on the system, while data (dispersed using SETH) will be written to all three disks at one time. In the case of a drive failure (or even a bad track or sector on a single drive), the SETH chips can still read data from the remaining two good drives until the bad drive can be replaced (or reformatted). When a new drive is installed, reading and writing all data back to the three disks will result in an initialization of the newly installed disk. As a matter of fact, adding a new drive can be done on-line,

```

0 0 0 0           ; Stage #3
1 0 1 1
-----
1 0 1 1
  0 0
-----
1 0 1 1           ; Stage #2
1 0 1 1
-----
1 1 0 1
  1 1
-----
1 1 1 0           ; Stage #1
0 0 0 0
-----
1 1 0 0
  1 1
-----
1 1 1 1           ; Stage #0
0 0 0 0
-----
1 1 1 0
  1 1
-----
1 1 0 1           ; Result

```

Figure 4: Multiplication of 1011 by 1100 using IPA

with no service interruption. It is important to realize that if the failure modes of the three disks are independent, which is normally the case, then the probability that the SETH-based design will fail is extremely small. For instance assume that the probability of losing a specific track (or sector) is P (P is usually in the order of 10^{-6}). The SETH-based design will fail to read a specific track if and only if the *same* track (or sector) in two or more of the disks will be lost. This probability can be shown to be $3P^2(1 - P)$ (in the order of 10^{-12}). A detailed analysis of the potential gains in I/O subsystems when using IDA are discussed in [4].

Data storage and retrieval using SETH is secure as well. When the system needs to be secured, the three disks can be locked in three different places (maybe under different machines – or at different sites.) An adversary will need to access at least two of the disks (assuming that he knows what SETH is really doing !!) before retrieving any “meaningful” information. In addition, to do that, the conversion matrix used for dispersal needs to be known. This makes it more difficult especially if such information is generated randomly by the operating system or even the underlying hardware.

The use of IDA in the design of RAID systems (Redundant Arrays of Inexpensive Disks) has been investigated in [3]. It provides unparalleled gains in performance, availability, and required redundancy. We have demonstrated that such an approach is superior to previously suggested techniques, namely shadowing and parity [6].

Data communication using SETH

Placing SETH chips on both sides of an information bus will increase the reliability of the bus and make it harder for information thieves to tap the information thereon. For example, one SETH chip may interface with an 8-bit data bus and send 16 bits to the other end, where another

Figure 5: The multiplier Unit and Stage

SETH chip would be used to recombine the data. The security of the bus results from the difficulty to decipher the information. The coding matrices used in SETH can be changed at frequent intervals to make deciphering still more difficult. The bus is highly fault-tolerant since the original information may be reconstructed from any two of the four groups of nibbles available on the bus.

In [7], IDA has been used in routing packets on cube-based architectures. The suggested technique is fully described in that paper. An obvious extension of this work would be to consider network topologies other than the n -cube. Also, fine tuning the routing algorithms to the system parameters^(f) is another interesting problem.

Conclusion

In this paper, we have presented “SETH” – a hardwired implementation of the Information Dispersal Algorithm. SETH allows the real-time dispersal of information into different pieces as well as the retrieval of the original information from the available pieces. SETH accepts a stream of data and a set of “secret keys” so as to produce the required streams of dispersed data to be stored on (or communicated with) the different machines. SETH might as well accept the streams of data from the different machines along with the necessary controls and keys so as to reconstruct the original information. The design of SETH involved finding efficient techniques for computing using IPA. In particular, we have outlined general methods for addition and multiplication in these systems.

The design of SETH can be extended in several ways. In our current implementation, an outside mechanism is

^(f)number of processors, size of packets, failure rates, . . . , etc.

icantly enhance the security of the dispersal algorithm at the expense of a blowup in the size of the chip as well as an increase in the propagation delay. The increase in the propagation delay is not a critical factor, since by using pipelining the overall delay in communicating messages (especially long ones) can be downplayed. The number of pins required for the chip is another crucial factor. In particular, if the chip is to be used for parallel communication, then the number of pins required for data i/o is $(n + m) \times k$. For large values of n , m , and k , serial communication is likely to be used. Still, another alternative would be to partition the computations so as more than one chip could be used in parallel.

The potential applications of IDA are numerous. In particular, the use of IDA in I/O systems, RAID designs, distributed communication and routing is promising. SETH demonstrates that using IDA in these applications is feasible.

Acknowledgements:

This project has been a collaboration involving many people. I would like to thank them all. In particular, I would like to thank Steve Morss and Adam Strassberg who helped me with the layout and simulation of SETH. I am grateful to Prof. Thomas Cheatham, Prof. Michael Rabin and Prof. James Clark for their advice and support during the course of this work.

References

- [1] A. Bestavros, S. Morss, and A. Strassberg, *SETH: a chip for reliable and secure communication using the Information Dispersal Algorithm*, Internal report, VLSI Lab., Harvard University, (May, 1988.)
- [2] A. Bestavros, *SETH: A VLSI chip for the real-time Information Dispersal and retrieval for security and fault-tolerance*, Technical Report, TR-06-89, Harvard University, (January, 1989.)
- [3] A. Bestavros, *IDA-based disk array systems*, Technical Memorandum 45312-890707-01TM, AT&T Bell Laboratories, (July, 1989.)
- [4] A. Bestavros, D. Chen, and W. Wong, *Reliability and performance of parallel disk systems*, Technical Memorandum 45312-891206-01TM, AT&T Bell Laboratories, (December, 1989.)
- [5] G. Gibson, L. Hellerstein, R. Karp, R. Katz and D. Patterson, *Coding Techniques for Handling Failures in Large Disk Arrays*, Technical Report UCB/CSD 88/477, Computer Science Division, University of California, (July, 1988.)
- [6] D. Patterson, P. Chen, G. Gibson, and R. Katz, "Introduction to Redundant Arrays of Inexpensive Disks (RAID)," *Proceedings of COMPCON-89, the Thirty-fourth IEEE Computer Society International Conference*, (March, 1989.)
- [7] M. Rabin, *Efficient Dispersal of Information for Security, Load Balancing and Fault Tolerance*, Technical Report, TR-02-87, Department of Computer Science, DAS, Harvard University, (April, 1987.) Also appeared in the *Journal of the Association for Computing Machinery*, Vol. 36, No. 2, (April, 1989), pp. 335-348.
- [8] A. Shamir, "How to share a secret?," *Communication of the ACM* 22, 11, (November, 1979), Pp. 612-613.

Figure 6: The matrix multiplication unit

Name of Cell	Average Delay ns.	Worst-case Power UWatt/MHz.	Total No of FET's	Area in Lambda ²
Inverter	1.44	2.33	2	038x011
And gate	2.07	0.76	6	036x036
Xor gate	1.99	0.36	8	049x052
Multi-stage	6.02	4.84	64	155x144
Multiplier	17.19	19.28	256	172x621
Adder	1.99	5.76	128	050x800
Matrix unit	19.98	160.00	2176	1565x891
Control	16.01	84.49	46	275x158
SETH	34.85	509.00	2428	Pad:64P46x68

Figure 7: Delay, power, and area of SETH units

responsible for providing the secret keys (inverse keys) to be used in dispersing (recombining) the data. This outside mechanism, however, can be relieved from the burden of computing the inverse keys if the design is made so that the inverse keys are computed on the fly (given the original secret keys and the set of intact sinks). This is quite feasible. Moreover, the possibility of "automatically" generating the secret keys for each file (or set of packets) is very attractive. In this case, information about these keys has to be included in the dispersed data so as to be used later in the computation of the inverse keys.

Our choices for n (the number of sinks), m (the minimum number of sinks required to reconstruct F) and k (the character size) can also be changed. It can be shown that the size of the chip scales linearly with $n \times m \times k^2$ and that the propagation delay scales linearly with k . Increasing n and m would result in a more flexible design in terms of the achievable levels of redundancy and fault-tolerance. On the other hand, increasing the value of k would signif-